IBM

Implementing the IBM Storwize V5000

- Easily manage and deploy systems with embedded GUI
- Experience rapid and flexible provisioning
- Protect data with remote mirroring

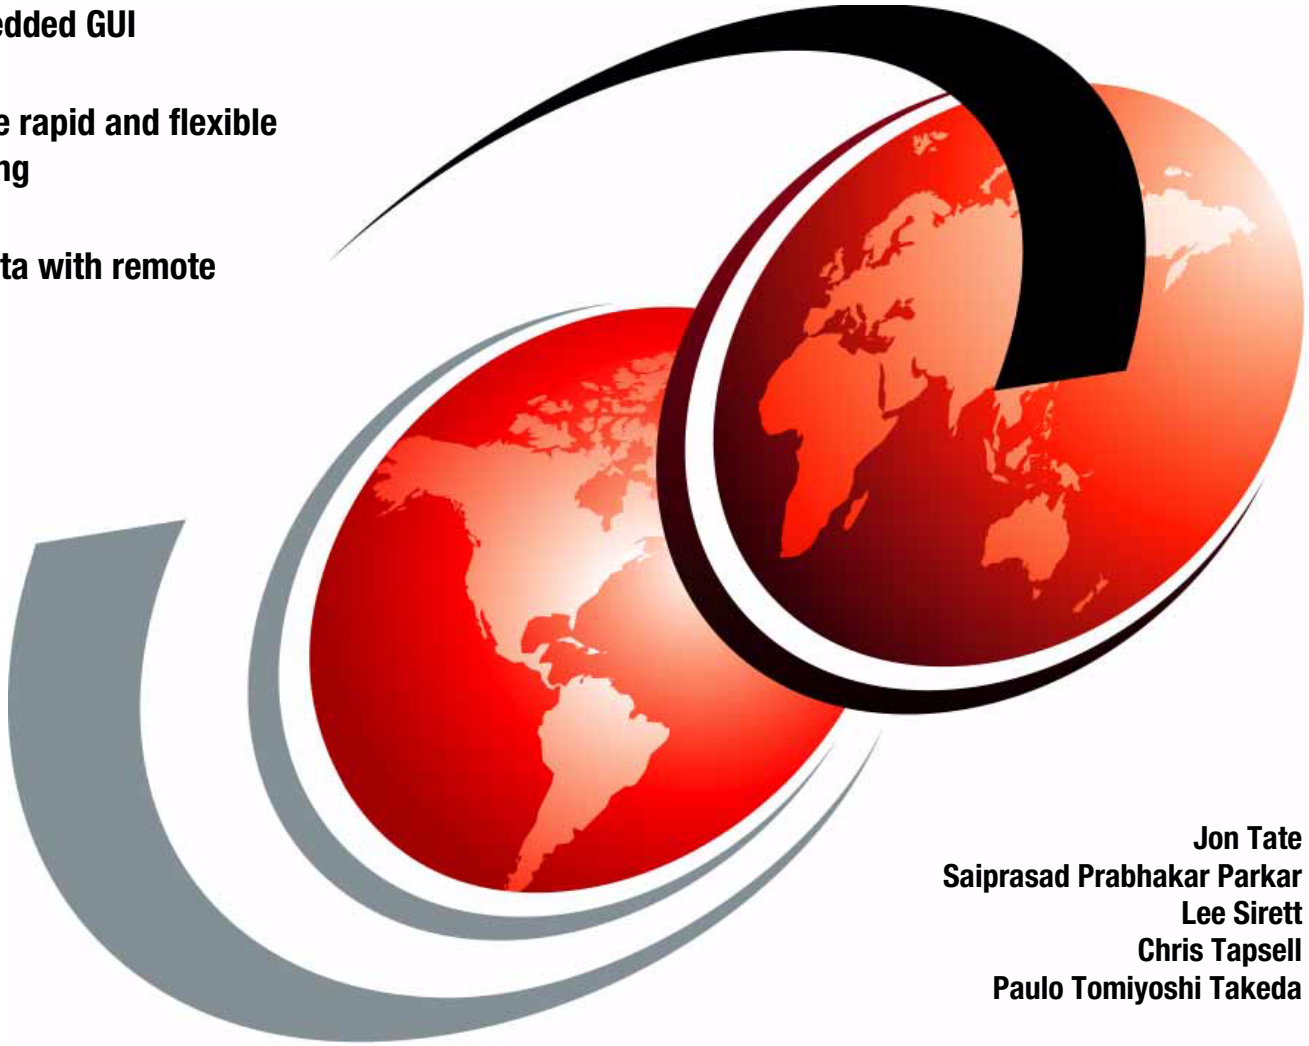

Jon Tate
Saiprasad Prabhakar Parkar
Lee Sirett
Chris Tapsell
Paulo Tomiyoshi Takeda

Redbooks

ibm.com/redbooks

International Technical Support Organization

Implementing the IBM Storwize V5000

October 2013

SG24-8162-00

Note: Before using this information and the product it supports, read the information in “Notices” on page ix.

First Edition (October 2013)

This edition applies to Version 7 Release 1 of IBM Storwize V5000 machine code.

Contents

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®
DS8000®
Easy Tier®
FlashCopy®
IBM®
Netfinity®
Power Systems™
Redbooks®
Redbooks (logo) ®
Storwize®
System i®
System Storage®
Tivoli®
VIA®
XIV®
xSeries®

The following terms are trademarks of other companies:

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Preface

Organizations of all sizes are faced with the challenge of managing massive volumes of increasingly valuable data. But storing this data can be costly, and extracting value from the data is becoming more difficult. IT organizations have limited resources but must stay responsive to dynamic environments and act quickly to consolidate, simplify, and optimize their IT infrastructures. The IBM® Storwize® V5000 system provides a smarter solution that is affordable, easy to use, and self-optimizing, which enables organizations to overcome these storage challenges.

Storwize V5000 delivers efficient, entry-level configurations that are specifically designed to meet the needs of small and midsize businesses. Designed to provide organizations with the ability to consolidate and share data at an affordable price, Storwize V5000 offers advanced software capabilities that are usually found in more expensive systems.

This IBM Redbooks® publication is intended for pre-sales and post-sales technical support professionals and storage administrators.

The concepts in this book also relate to the IBM Storwize V3700.

This book was written at a software level of Version 7 Release 1.

Authors

This book was produced by a team of specialists from around the world working at the IBM Manchester Lab, UK.

Jon Tate is a Project Manager for IBM System Storage® SAN Solutions at the International Technical Support Organization, San Jose Center. Before joining the ITSO in 1999, he worked in the IBM Technical Support Center, providing Level 2/3 support for IBM storage products. Jon has over 27 years of experience in storage software and management, services, and support, and is an IBM Certified Consulting IT Specialist and an IBM SAN Certified Specialist. He is also the UK Chairman of the Storage Networking Industry Association.

Saiprasad Prabhakar Parkar is a senior IT Specialist for IBM at the ISTL Pune, India. He has worked for IBM for five years and provides Level 3 support for UNIX, IBM Power Systems, and storage products. Sai has a total of 10 years of experience in UNIX, and Power System and Storage. He also is an IBM Certified Solution Specialist.

Lee Sirett is a Storage Technical Advisor for the European Storage Competency Centre (ESCC) in Mainz, Germany. Before he joined the ESCC, he worked in IBM Technical Support Services for 10 years and providing support for various IBM Products, including Power Systems™. Lee has 24 years experience in the IT Industry. He is IBM Storage Certified and an IBM Certified XIV® Administrator and Certified XIV Specialist.

Chris Tapsell is a Presales Storage Technical Specialist for IBM Systems & Technology Group. Before his current role, in his 25+ years at IBM, he worked as a CE covering products such as, typewriters up to AS400 (System i®), as a Support Specialist for all of the IBM Intel server products (PC Server, Netfinity®, xSeries®, and System x), PCs and notebooks, and as a Presales Technical Specialist for System x.

Chris holds a number of IBM Certifications covering System x and Storage products.

Paulo Tomiyoshi Takeda is a SAN and Storage Disk Specialist at IBM Brazil. He has over eight years of experience in the IT arena. He holds a Bachelors degree in Information Systems from Universidade da Fundacao Educacional de Barretos and is IBM Certified for IBM DS8000® and IBM Storwize V7000. His areas of expertise include planning, configuring, and troubleshooting DS8000, SAN Volume Controller, and IBM Storwize V7000. He was involved in storage-related projects such as, capacity growth planning, SAN consolidation, storage microcode upgrades, and copy services in the Open Systems environment.

Thanks to the following people for their contributions to this project:

- Martyn Spink
- Djihed Afifi
- Karl Martin
- Imran Imtiaz
- Doug Neil
- David Turnbull
- Stephen Bailey

IBM Manchester Lab

- John Fairhurst
- Paul Marris
- Paul Merrison

IBM Hursley

- Mary Connell

IBM STG

Thanks to the following authors of the previous edition of this book:

- Uwe Dubberke
- Justin Heather
- Andrew Hickey
- Imran Imtiaz
- Nancy Kinney
- Dieter Utesch

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience by using leading-edge technologies. Your efforts helps to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at this website:

`http://www.ibm.com/redbooks/residencies.html`

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form that is found at this website:

  `http://www.ibm.com/redbooks`

- Send your comments in an email to:

  `redbooks@us.ibm.com`

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

1

Overview of the IBM Storwize V5000 system

This chapter provides an overview of the IBM Storwize V5000 architecture and includes a brief explanation of storage virtualization.

This chapter includes the following topics:

- IBM Storwize V5000 overview
- IBM Storwize V5000 terminology
- IBM Storwize V5000 models
- IBM Storwize V5000 hardware
- IBM Storwize V5000 terms
- IBM Storwize V5000 features
- Problem management and support

1.1 IBM Storwize V5000 overview

The IBM Storwize V5000 solution provides a modular storage system that includes the capability to virtualize its own internal storage and external SAN-attached storage. The IBM Storwize V5000 system is a virtualizing Redundant Array of Independent Disk (RAID) entry and midrange storage system.

IBM Storwize V5000 features the following benefits:

- Brings enterprise technology to entry and midrange storage
- Speciality administrators are not required
- Easy client setup and service
- Ability to grow the system incrementally as storage capacity and performance needs change
- Simple integration into the server environment

The IBM Storwize V5000 addresses the block storage requirements of small and midsize organizations and consists of one 2U control enclosure and, optionally, up to six 2U expansion enclosures, which are connected via serial-attached Small Computer Systems Interface (SCSI SAS) cables that make up one system that is called an I/O Group.

Two I/O Groups can be connected to form a cluster.

The control and expansion enclosures are available in the following form factors and can be intermixed within an I/O group:

- 12 x 3.5-inch drives in a 2U unit
- 24 x 2.5-inch drives in a 2U unit

Within each enclosure, there are two canisters. Control enclosures contain two node canisters, and expansion enclosures contain two expansion canisters.

The IBM Storwize V5000 supports up to 168 x 3.5-inch or 336 x 2.5-inch or a combination of both drive form factors for the internal storage in a two I/O group cluster.

SAS, NL-SAS and solid-state drives (SSDs) types are supported.

The IBM Storwize V5000 is designed to accommodate the most common storage network technologies to enable easy implementation and management. It can be attached to hosts via a SAN fabric, an iSCSI infrastructure, or via SAS. Hosts can be SAN or Direct attached.

Important: IBM Storwize V5000 can be direct-attached to a host. For more information about restrictions, see the IBM System Storage Interoperation Center (SSIC), which is available at this website:

`http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss`

Information also is available at this website:

`http://www-01.ibm.com/support/docview.wss?uid=ssg1S1004233`

The IBM Storwize V5000 is a virtualized storage solution that groups its internal drives into RAID arrays (called Managed Disks or MDisks). MDisks can also be created by importing LUNs from external FC SAN-attached storage. These MDisks are then grouped into storage pools. Volumes are created from these storage pools and provisioned out to hosts. Storage pools are normally created with MDisks of the same type and capacity of drive. Volumes can be moved non-disruptively between storage pools with differing performance characteristics. For example, a volume can be moved between a storage pool that is made up of NL-SAS drives to a storage pool made up of SAS drives.

The IBM Storwize V5000 system also provides several configuration options that are aimed at simplifying the implementation process. It also provides configuration presets and automated wizards called *Directed Maintenance Procedures* (DMP) to help resolve any events that might occur.

Included with an IBM Storwize V5000 system is a simple and easy to use graphical user interface (GUI) that is designed to allow storage to be deployed quickly and efficiently. The GUI runs on any supported browser. The management GUI contains a series of preestablished configuration options that are called *presets* that use commonly used settings to quickly configure objects on the system. Presets are available for creating volumes and IBM FlashCopy® mappings and for setting up a RAID configuration.

You can also use the command-line interface (CLI) to set up or control the system.

1.2 IBM Storwize V5000 terminology

The IBM Storwize V5000 system introduced some terminology, which is consistent with the entire IBM Storwize family and SAN Volume Controller. The terms are defined in Table 1-1.

Table 1-1 IBM Storwize V5000 terminology

IBM Storwize V5000 term	Definition
Battery	Each control enclosure node canister in a IBM Storwize V5000 contains a battery.
Canister	Canisters are hardware units that are subcomponents of a IBM Storwize V5000 enclosures. Each enclosure contains two canisters.
Chain	A set of enclosures that is attached to provide redundant access to the drives that are inside the enclosures. Each control enclosure has two chains.
Clone	A copy of a volume on a server at a particular point. The contents of the copy can be customized while the contents of the original volume are preserved.
Control enclosure	A hardware unit that includes the chassis, node canisters, drives, and power sources.
Data migration	By using IBM Storwize V5000, you can migrate data from existing external storage to its internal volumes.
Drive	IBM Storwize V5000 supports a range of hard disk drives (HDDs) and SSDs.

IBM Storwize V5000 term	Definition
Enclosure	An enclosure is the basic housing unit for the IBM Storwize V5000. It is the rack-mounted hardware that contains all the main components of the system: canisters, drives, and power supplies.
Event	An occurrence that is significant to a task or system. Events can include completion or failure of an operation, a user action, or the change in the state of a process.
Expansion canister	A hardware unit that includes the SAS interface hardware that enables the node hardware to use the drives of the expansion enclosure.
Expansion enclosure	A hardware unit that includes expansion canisters, drives, and power supply units.
External storage	MDisks that are SCSI logical units (LUs) presented by storage systems that are attached to and managed by the clustered system.
Fibre Channel port	Fibre Channel ports are connections for the hosts to get access to the IBM Storwize V5000.
Host mapping	The process of controlling which hosts can access specific volumes within a IBM Storwize V5000.
Internal storage	Array MDisks and drives that are held in enclosures and nodes that are part of the IBM Storwize V5000.
iSCSI (Internet Small Computer System Interface)	Internet Protocol (IP)-based storage networking standard for linking data storage facilities.
Managed disk (MDisk)	A component of a storage pool that is managed by a clustered system. An MDisk is part of a RAID array of internal storage or a SCSI LU for external storage. An MDisk is not visible to a host system on the storage area network.
Node canister	A hardware unit that includes the node hardware, fabric, and service interfaces, SAS, expansion ports, and battery.
PHY	A single SAS lane. There are four PHYs in each SAS cable.
Power Supply Unit	Each enclosure has two power supply units (PSU).
Quorum disk	A disk that contains a reserved area that is used exclusively for cluster management. The quorum disk is accessed when it is necessary to determine which half of the cluster continues to read and write data.
Serial-Attached SCSI (SAS) ports	SAS ports are connections for the host to get direct attached access to the IBM Storwize V5000 and expansion enclosure.
Snapshot	An image backup type that consists of a point-in-time view of a volume.
Storage pool	A collection of storage capacity that provides the capacity requirements for a volume.
Strand	The SAS connectivity of a set of drives within multiple enclosures. The enclosures can be control enclosures or expansion enclosures.

IBM Storwize V5000 term	Definition
Thin provisioning or thin provisioned	The ability to define a storage unit (full system, storage pool, or volume) with a logical capacity size that is larger than the physical capacity that is assigned to that storage unit.
Volume	A discrete unit of storage on disk, tape, or other data recording medium that supports some form of identifier and parameter list, such as, a volume label or input/output control.
Worldwide port names	Each Fibre Channel port is identified by their physical port number and worldwide port name (WWPN).

1.3 IBM Storwize V5000 models

The IBM Storwize V5000 platform consists of a number of different models.

More information: For more information about the features, benefits, and specifications of IBM Storwize V5000 models, see this website:

http://www.ibm.com/systems/storage/disk/storwize_v5000/index.html

The information in this book is accurate at the time of writing. However, as the IBM Storwize V5000 matures, expect to see new features and enhanced specifications.

The IBM Storwize V5000 models are described in Table 1-2. All models have two node canisters. C models are control enclosures and E models are expansion enclosures.

Table 1-2 IBM Storwize V5000 models

Model	Cache	Drive slots
One-Year Warranty		
2077-12C	16 GB	12 x 3.5-inch
2077-24C	16 GB	24 x 2.5-inch
2077-12E	N/A	12 x 3.5-inch
2077-24E	N/A	24 x 2.5-inch
Three-Year Warranty		
2078-12C	16 GB	12 x 3.5-inch
2078-24C	16 GB	24 x 2.5-inch
2078-12E	N/A	12 x 3.5-inch
2078-24E	N/A	24 x 2.5-inch

Figure 1-1 shows the front view of the 2077/2078-12C and 12E enclosures.

Figure 1-1 IBM Storwize V5000 front view for 2077/2078-12C and 12E enclosures

The drives are positioned in four columns of three horizontal-mounted drive assemblies. The drive slots are numbered 1 - 12, starting at upper left and going left to right, top to bottom.

Figure 1-2 shows the front view of the 2077/2078-24C and 24E enclosures.

Figure 1-2 IBM Storwize V5000 front view for 2077/2078-24C and 24E enclosure

The drives are positioned in one row of 24 vertically mounted drive assemblies. The drive slots are numbered 1 - 24, starting from the left. There is a vertical center drive bay molding between slots 12 and 13.

1.4 IBM Storwize V5000 hardware

The IBM Storwize V5000 solution is a modular storage system that is built on a common enclosure (control enclosure and expansion enclosure).

Figure 1-3 shows an overview of the hardware components of the IBM Storwize V5000 solution.

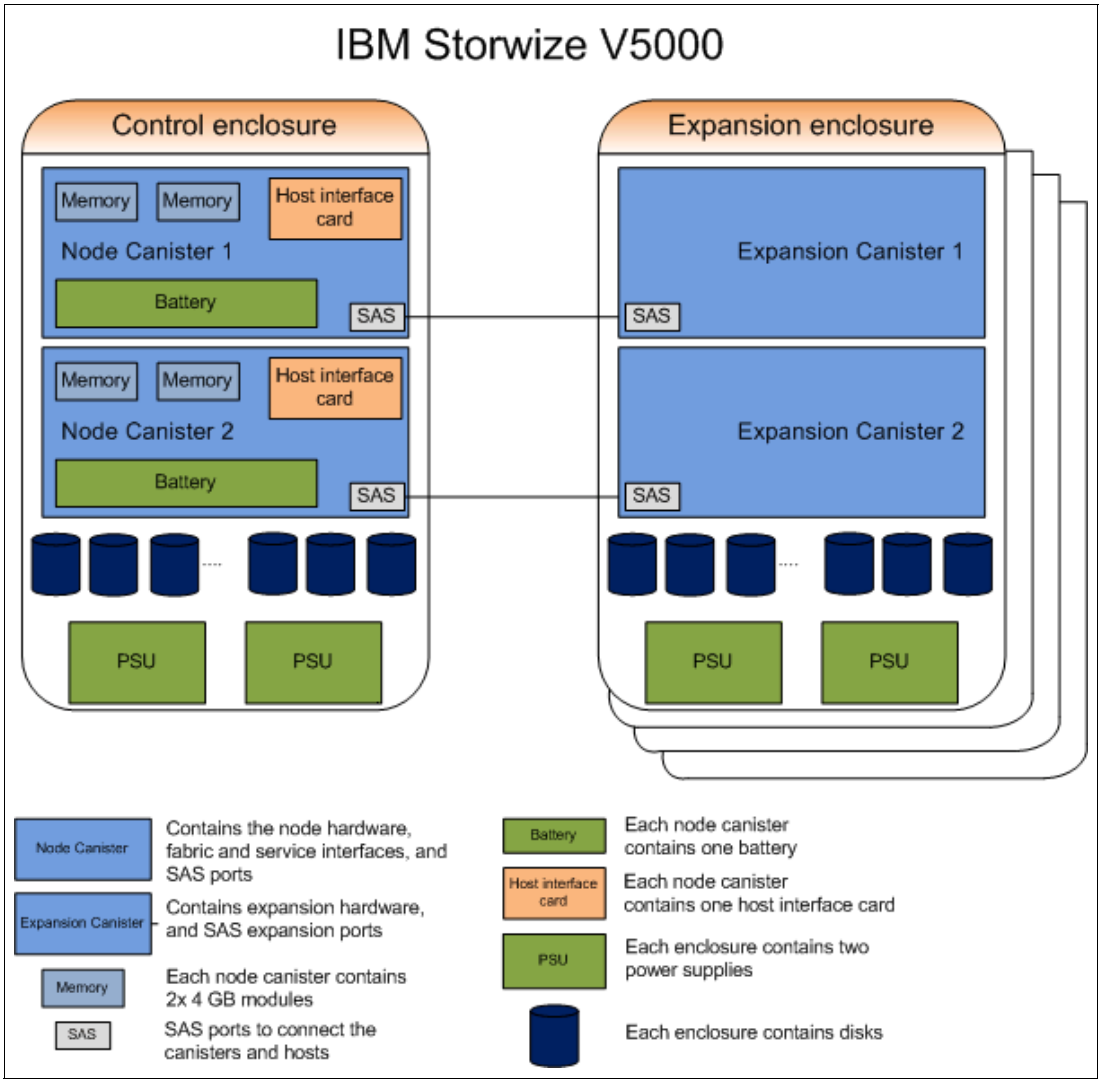

Figure 1-3 IBM Storwize V5000 hardware components

Figure 1-4 shows the controller rear view of IBM Storwize V5000 models 12C and 24C.

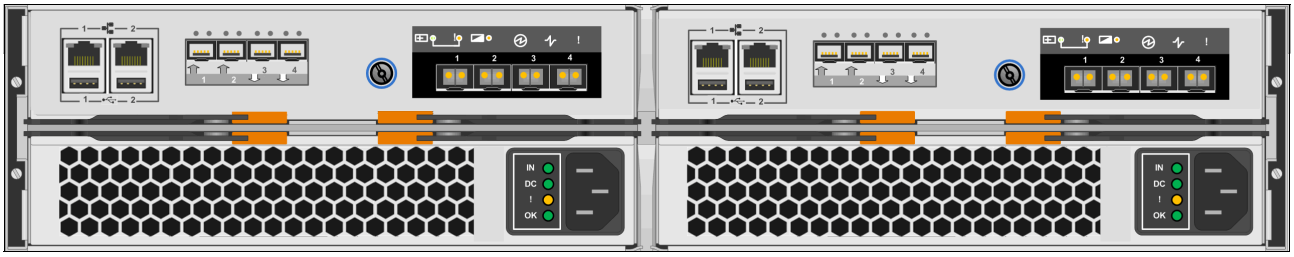

Figure 1-4 IBM Storwize V5000 controller rear view of models 12C and 24C

In Figure 1-4 on page 7, you can see that there are two power supply slots at the bottom of the enclosure. The power supplies are identical and exchangeable. There are two canister slots at the top of the chassis.

In Figure 1-5, you can see the rear view of an IBM Storwize V5000 expansion enclosure.

Figure 1-5 IBM Storwize V5000 expansion enclosure rear view - models 12E and 24E

You can see that the only difference between the node enclosure and the expansion enclosure are the canisters. The canisters of the expansion have only the two SAS ports.

For more information about the expansion enclosure, see 1.4.2, “Expansion enclosure” on page 9.

1.4.1 Control enclosure

Each IBM Storwize V5000 system has one control enclosure that contains two node canisters, disk drives, and two power supplies.

Figure 1-6 shows a single node canister.

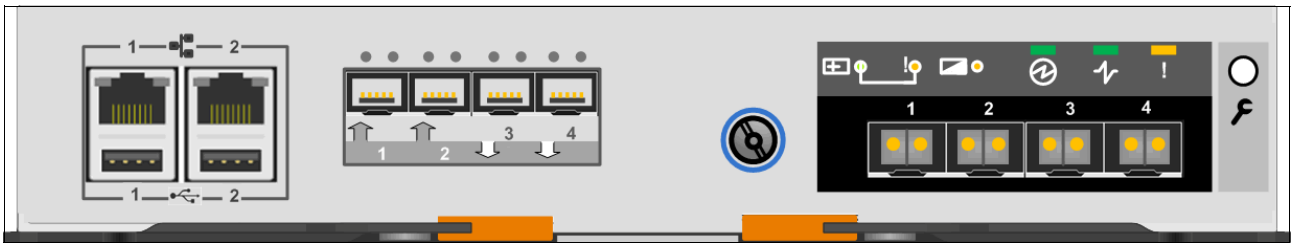

Figure 1-6 IBM Storwize V5000 node canister

Each node canister contains the following hardware:

- Battery
- Memory: 8 GB memory
- 8 Gb Fibre Channel Host interface card
- Four 6 Gbps SAS ports
- Two 10/100/1000 Mbps Ethernet ports
- Two USB 2.0 ports (one port is used during installation)
- System flash

The battery is used in case of power loss. The IBM Storwize V5000 system uses this battery to power the canister while the cache data is written to the internal system flash. This memory dump is called a *fire hose memory dump*. After the system is up again, this data is loaded back to the cache for destage to the disks.

Figure 1-6 on page 8 also shows the following that are provided by the IBM Storwize V5000 node canister:

- Two 10/100/1000 Mbps Ethernet ports, which are used for management. Port 1 (left port) must be configured. The second port is optional and is used for management. Both ports can be used for iSCSI traffic. For more information, see Chapter 4, “Host configuration” on page 153.
- Two USB ports. One port is used during the initial configuration or when there is a problem. They are numbered 1 on the left and 2 on the right. For more information about usage, see Chapter 2, “Initial configuration” on page 27.
- Four serial attached SCSI (SAS) ports. They are numbered 1 on the left to 4 on the right. The IBM Storwize V5000 uses ports 1 and 2 for host connectivity and ports 3 and 4 to connect to the optional expansion enclosures. The IBM Storwize V5000 incorporates two SAS chains and three expansion enclosures can be connected to each chain.
- Four Fibre Channel ports, which operate at 2 Gbps, 4 Gbps, or 8 Gbps. The ports are numbered from left to right starting with 1.

Service port: Do not use the port marked with a wrench. This port is a service port only.

The two nodes act as a single processing unit and form an I/O group that is attached to the SAN fabric, an iSCSI infrastructure or directly attached to hosts via FC or SAS. The pair of nodes is responsible for serving I/O to a volume. The two nodes provide a highly available fault-tolerant controller so that if one node fails, the surviving node automatically takes over. Nodes are deployed in pairs that are called *I/O groups*.

One node is designated as the configuration node, but each node in the control enclosure holds a copy of the control enclosure state information.

The IBM Storwize V5000 supports two I/O groups in a clustered system.

The terms *node canister* and *node* are used interchangeably throughout this book.

1.4.2 Expansion enclosure

The optional IBM Storwize V5000 expansion enclosure contains two expansion canisters, disk drives, and two power supplies.

Figure 1-7 shows an overview of the expansion enclosure.

Figure 1-7 Expansion enclosure of the IBM Storwize V5000

The expansion enclosure power supplies are the same as the control enclosure. There is a single power lead connector on each power supply unit.

Figure 1-8 shows the expansion canister ports.

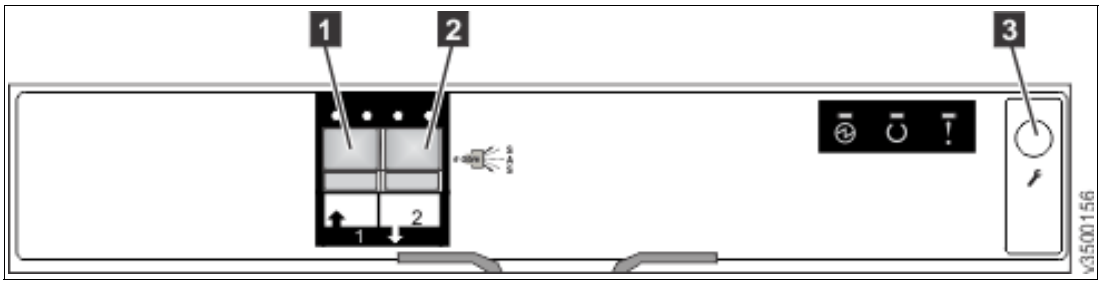

Figure 1-8 Expansion canister ports

As shown in Figure 1-8, each expansion canister provides two SAS interfaces that are used to connect to the control enclosure and any optional expansion enclosures. The ports are numbered 1 on the left and 2 on the right. SAS port 1 is the IN port and SAS port 2 is the OUT port.

Use of the SAS connector 1 is mandatory because the expansion enclosure must be attached to a control enclosure or another expansion enclosure. SAS connector 2 is optional because it is used to attach to more expansion enclosures.

Each port includes two LEDs to show the status. The first LED indicates the link status and the second LED indicates the fault status.

For more information about LED and ports, see this website:

http://pic.dhe.ibm.com/infocenter/storwize/v5000_ic/index.jsp

1.4.3 Host connectivity

With 1 Gb iSCSI, 8 Gb FC, and 6 Gb SAS host interfaces supported as standard, the IBM Storwize V5000 is designed to accommodate the most common storage networks. This broad networking support enables deployment of IBM Storwize V5000 in existing storage network infrastructures.

The 1 Gb iSCSI and 6 Gb SAS interfaces are built into the node canister hardware and the 8 Gb FC interface is supplied by a host interface card (HIC). As of this writing, the 8 Gb FC HIC is the only HIC that is available and is supplied as standard.

1.4.4 Disk drive types

IBM Storwize V5000 enclosures support SSD, SAS, and Nearline SAS drive types. Each drive has two ports (two PHYs) and I/O can be issued down both paths simultaneously.

Table 1-3 shows the IBM Storwize V5000 Disk Drive types that are available at the time of writing.

Table 1-3 IBM Storwize V5000 Disk Drive types

Drive type		Speed	Size
2.5-inch form factor	Solid-state disk	N/A	200 and 400 GB
2.5-inch form factor	SAS	10,000 rpm	600 GB, 900 GB, and 1.2 TB
2.5-inch form factor	SAS	15,000 rpm	146 and 300 GB
2.5-inch form factor	Nearline SAS	7,200 rpm	1 TB
3.5-inch form factor	SAS	10,000 rpm	900 GB and 1.2 TB[a]
3.5-inch form factor	SAS	15,000 rpm	300 GB[b]
3.5-inch form factor	Nearline SAS	7,200 rpm	2 TB, 3 TB, and 4 TB

a. 2.5-inch drive in a 3.5-inch drive carrier
b. 2.5-inch drive in a 3.5-inch drive carrier

1.5 IBM Storwize V5000 terms

In this section, we introduce the terms that are used for the IBM Storwize V5000 throughout this book.

1.5.1 Hosts

A host system is a server that is connected to IBM Storwize V5000 through a Fibre Channel connection, an iSCSI connection, or through a SAS connection.

Hosts are defined on IBM Storwize V5000 by identifying their WWPNs for Fibre Channel and SAS hosts. iSCSI hosts are identified by using their iSCSI names. The iSCSI names can be iSCSI qualified names (IQNs) or extended unique identifiers (EUIs). For more information, see Chapter 4, “Host configuration” on page 153.

Hosts can be Fibre Channel attached via an existing Fibre Channel network infrastructure or direct attached, iSCSI attached via an existing IP network, or directly attached via SAS. A significant benefit of having direct attachment is that you can attach the host directly to the IBM Storwize V5000 without the need for an FC or IP network.

1.5.2 Node canister

A node canister provides host interfaces, management interfaces, and SAS interfaces to the control enclosure. A node canister has the cache memory, the internal storage to store software and logs, and the processing power to run the IBM Storwize V5000 virtualizing and management software. A clustered system consists of a one or two node pairs or I/O groups.

One of the nodes within the system is known as the *configuration node* that manages configuration activity for the clustered system. If this node fails, the system nominates the other node to become the configuration node.

1.5.3 I/O groups

Within IBM Storwize V5000, there are one or two pairs of node canisters, which are known as *I/O groups*. The IBM Storwize V5000 supports four node canisters in the clustered system, which provides two I/O groups.

When a host server performs I/O to one of its volumes, all the I/Os for a specific volume are directed to the I/O group. Also, under normal conditions, the I/Os for that specific volume are always processed by the same node within the I/O group.

Both nodes of the I/O group act as preferred nodes for their own specific subset of the total number of volumes that the I/O group presents to the host servers (a maximum of 2048 volumes per hosts). However, both nodes also act as a failover node for its partner node within the I/O group. Therefore, a node takes over the I/O workload from its partner node (if required) without affecting the server's application.

In a IBM Storwize V5000 environment (which uses active-active architecture), the I/O handling for a volume can be managed by both nodes of the I/O group. Therefore, servers that are connected through Fibre Channel connectors must use multipath device drivers to handle this capability.

The IBM Storwize V5000 I/O groups are connected to the SAN so that all application servers that access volumes from the I/O group have access to them. Up to 1024 host server objects can be defined to one I/O group or 2048 in a two I/O group system.

Important: The active/active architecture provides availability to process I/Os for both controller nodes and allows the application to continue running smoothly, even if the server has only one access route or path to the storage controller. This type of architecture eliminates the path/LUN thrashing that is typical of an active/passive architecture.

1.5.4 Clustered system

A clustered system consists of one or two pairs of node canisters or I/O groups. All configuration, monitoring, and service tasks are performed at the system level. The configuration settings are replicated across all node canisters in the clustered system. To facilitate these tasks, one or two management IP addresses are set for the clustered system. By using this configuration, you can manage the clustered system as a single entity.

There is a process to back up the system configuration data on to disk so that the clustered system can be restored in the event of a disaster. This method does not back up application data; only IBM Storwize V5000 system configuration information is backed up.

System configuration backup: After the system configuration is backed up, save the backup data on your hard disk (or at the least outside of the SAN). If you are unable to access the IBM Storwize V5000, you do not have access to the backup data if it is on the SAN. Perform this configuration backup after each configuration change to be safe.

The system can be configured by using the IBM Storwize V5000 management software (GUI), CLI, or the USB key.

1.5.5 RAID

The IBM Storwize V5000 contains a number of internal drives, but these drives cannot be directly added to storage pools. The drives must be included in a RAID array to provide protection against the failure of individual drives.

These drives are referred to as members of the array. Each array has a RAID level. RAID levels provide different degrees of redundancy and performance and have different restrictions regarding the number of members in the array.

IBM Storwize V5000 supports hot spare drives. When an array member drive fails, the system automatically replaces the failed member with a hot spare drive and rebuilds the array to restore its redundancy. Candidate and spare drives can be manually exchanged with array members.

Each array has a set of goals that describe the required location and performance of each array. A sequence of drive failures and hot spare takeovers can leave an array unbalanced, that is, with members that do not match these goals. The system automatically rebalances such arrays when the appropriate drives are available.

The following RAID levels are available:

- RAID 0 (striping, no redundancy)

  RAID 0 arrays stripe data across the drives. The system supports RAID 0 arrays with one member, which is similar to traditional JBOD attach. RAID 0 arrays have no redundancy, so they do not support hot spare takeover or immediate exchange. A RAID 0 array can be formed by one to eight drives.

- RAID 1 (mirroring between two drives, which is implemented as RAID 10 with two drives)

  RAID 1 arrays stripe data over mirrored pairs of drives. A RAID 1 array mirrored pair is rebuilt independently. A RAID 1 array can be formed by two drives only.

- RAID 5 (striping, can survive one drive fault, with parity)

  RAID 5 arrays stripe data over the member drives with one parity strip on every stripe. RAID 5 arrays have single redundancy. The parity algorithm means that an array can tolerate no more than one member drive failure. A RAID 5 array can be formed by 3 - 16 drives.

- RAID 6 (striping, can survive two drive faults, with parity)

  RAID 6 arrays stripe data over the member drives with two parity stripes (which is known as the P-parity and the Q-parity) on every stripe. The two parity strips are calculated by using different algorithms, which give the array double redundancy. A RAID 6 array can be formed by 5 - 16 drives.

- RAID 10 (RAID 0 on top of RAID 1)

  RAID 10 arrays have single redundancy. Although they can tolerate one failure from every mirrored pair, they cannot tolerate two-disk failures. One member out of every pair can be rebuilding or missing at the same time. A RAID 10 array can be formed by 2 - 16 drives.

1.5.6 Managed disks

An MDisk refers to the unit of storage that IBM Storwize V5000 virtualizes. This unit can be a logical volume on an external storage array that is presented to the IBM Storwize V5000 or a RAID array that consists of internal drives. The IBM Storwize V5000 then can allocate these MDisks into storage pools.

An MDisk is invisible to a host system on the storage area network because it is internal to the IBM Storwize V5000 system.

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 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 being used by the IBM Storwize V5000 storage system.

- Managed

  Managed MDisks are LUNs 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 might be on these LUNs when they are added 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. For more information, see Chapter 6, “Storage migration wizard” on page 237.

1.5.7 Quorum disks

A quorum disk is an MDisk that contains a reserved area for use exclusively by the system. In the IBM Storwize V5000, internal drives can be considered as quorum candidates. The clustered system uses quorum disks to break a tie when exactly half the nodes in the system remain after a SAN failure.

The clustered system automatically forms the quorum disk by taking a small amount of space from an MDisk. It allocates space from up to three different MDisks for redundancy, although only one quorum disk is active.

To avoid the possibility of losing all of the quorum disks because of a failure of a single storage system if the environment has multiple storage systems, you should allocate the quorum disk on different storage systems. It is possible to manage the quorum disks by using the CLI.

1.5.8 Storage pools

A storage pool is a collection of MDisks (up to 128) that are grouped to provide capacity for volumes. All MDisks in the pool are split into extents of the same size. Volumes are then allocated out of the storage pool and are mapped to a host system.

MDisks can be added to a storage pool at any time to increase the capacity of the pool. MDisks can belong in only one storage pool. For more information, see Chapter 7, “Storage pools” on page 295.

Each MDisk in the storage pool is divided into a number of extents. The size of the extent is selected by the administrator when the storage pool is created and cannot be changed later. The size of the extent ranges from 16 MB - 8 GB.

Default extent size: The GUI of IBM Storwize V5000 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. The GUI cannot change the extent size. Therefore, if you want to create storage pools with a different extent size, this must be done via the CLI by using the **mkmdiskgrp** and **mkarray** commands.

The extent size directly affects the maximum volume size and storage capacity of the clustered system.

A system can manage 2^22 (4,194,304) extents. For example, with a 16 MB extent size, the system can manage up to 16 MB x 4,194,304 = 64 TB of storage.

The effect of extent size on the maximum volume and cluster size is shown in Table 1-4.

Table 1-4 Maximum volume and cluster capacity by extent size

Extent size	Maximum volume capacity for normal volumes (GB)	Maximum storage capacity of cluster
16	2048 (2 TB)	64 TB
32	4096 (4 TB)	128 TB
64	8192 (8 TB)	256 TB
128	16384 (16 TB)	512 TB
256	32768 (32 TB)	1 PB
512	65536 (64 TB)	2 PB
1024	131072 (128 TB)	4 PB
2048	262144 (256 TB)	8 PB
4096	262144 (256 TB)	16 PB
8192	262144 (256 TB)	32 PB

Use the same extent size for all storage pools in a clustered system, which is a prerequisite if you want to migrate a volume between two storage pools. If the storage pool extent sizes are not the same, you must use volume mirroring to copy volumes between storage pools, as described in Chapter 7, “Storage pools” on page 295.

A storage pool can have a threshold warning set that automatically issues a warning alert when the used capacity of the storage pool exceeds the set limit.

Single-tiered storage pool

MDisks that are used in a single-tiered storage pool should have the following characteristics to prevent performance and other problems:

- They should have the same hardware characteristics, for example, the same RAID type, RAID array size, disk type, and disk revolutions per minute (RPMs).
- The disk subsystems providing the MDisks must have similar characteristics, for example, maximum input/output operations per second (IOPS), response time, cache, and throughput.
- Use MDisks of the same size, and ensure that the MDisks provide the same number of extents. If this configuration is not feasible, you must check the distribution of the volumes' extents in that storage pool.

Multitiered storage pool

A multitiered storage pool has a mix of MDisks with more than one type of disk tier attribute, for example, a storage pool that contains a mix of generic_hdd AND generic_ssd MDisks.

A multitiered storage pool contains MDisks with different characteristics unlike the single-tiered storage pool. However, each tier should have MDisks of the same size and MDisks that provide the same number of extents.

A multitiered storage pool is used to enable automatic migration of extents between disk tiers by using the IBM Storwize V5000 Easy Tier® function, as described in Chapter 9, "Easy Tier" on page 411.

1.5.9 Volumes

A volume is a logical disk that is presented to a host system by the clustered system. In our virtualized environment, the host system has a volume that is mapped to it by IBM Storwize V5000. IBM Storwize V5000 translates this volume into a number of extents, which are allocated across MDisks. The advantage with storage virtualization is that the host is decoupled from the underlying storage, so the virtualization appliance can move around the extents without impacting the host system.

The host system cannot directly access the underlying MDisks in the same manner as it can access RAID arrays in a traditional storage environment.

The following types of volumes are available:

- Striped

  A striped volume is allocated one extent in turn from each MDisk in the storage pool. This process continues until the space that is required for the volume is satisfied.

  It also is possible to supply a list of MDisks to use.

Figure 1-9 shows how a striped volume is allocated, assuming 10 extents are required.

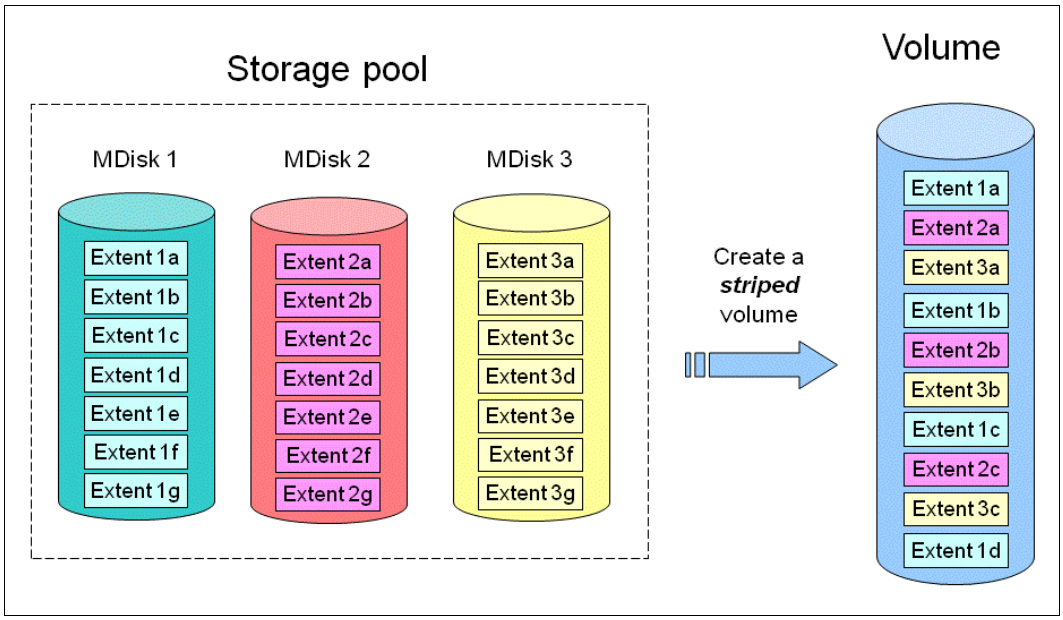

Figure 1-9 Striped volume

- Sequential

A sequential volume is a volume in which the extents are allocated one after the other from one MDisk to the next MDisk, as shown in Figure 1-10.

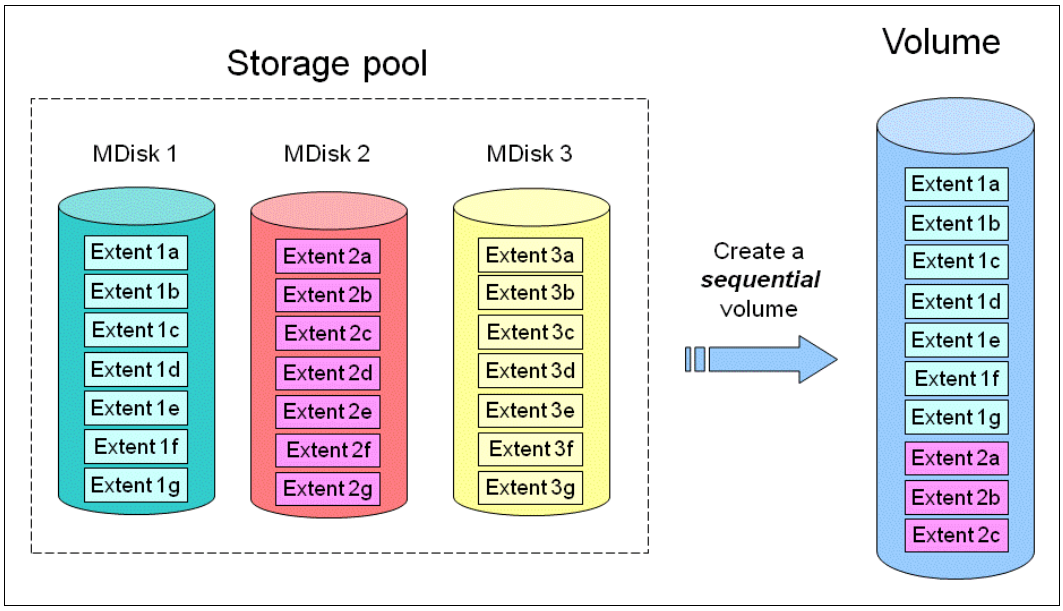

Figure 1-10 Sequential volume

- Image mode

  Image mode volumes are special volumes that have a direct relationship with one MDisk. They are used to migrate existing data into and out of the clustered system to or from external FC SAN-attached storage.

  When the image mode volume is created, a direct mapping is made between extents that are on the MDisk and the extents that are on the volume. The logical block address (LBA) x on the MDisk is the same as the LBA x on the volume, which ensures that the data on the MDisk is preserved as it is brought into the clustered system, as shown in Figure 1-11.

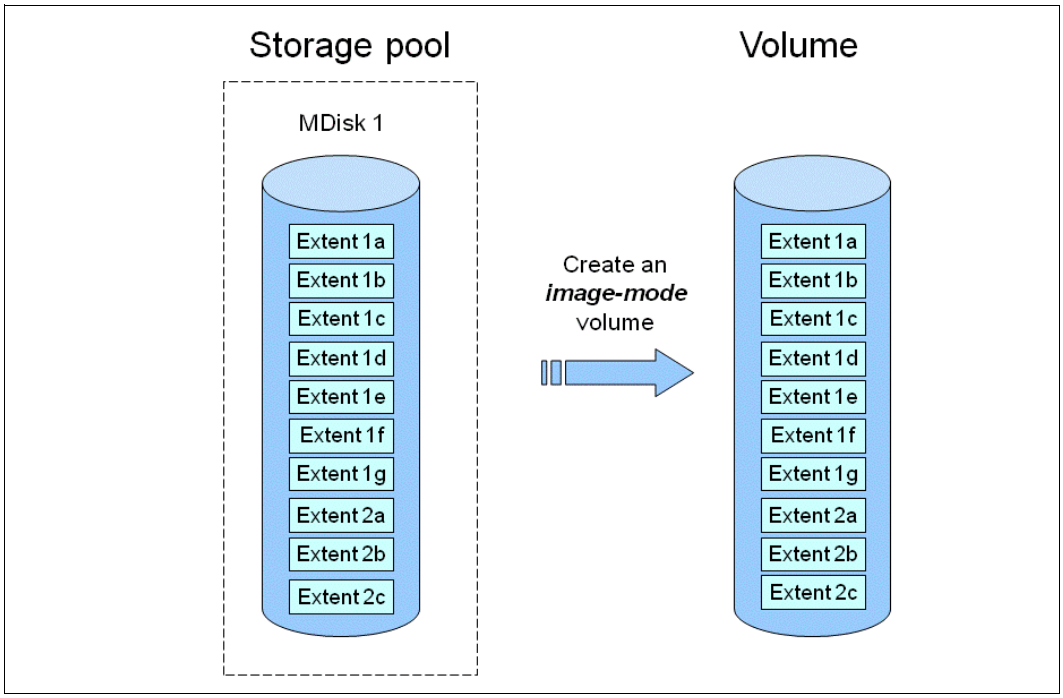

Figure 1-11 Image mode volume

Some virtualization functions are not available for image mode volumes, so it is often useful to migrate the volume into a new storage pool. After it is migrated, the MDisk becomes a managed MDisk.

If you want to migrate data from an existing storage subsystem, use the Storage Migration wizard, which guides you through the process.

For more information, see Chapter 6, “Storage migration wizard” on page 237.

If you add an MDisk that contains data to a storage pool, any data on the MDisk is lost. If you are presenting externally virtualized LUNs that contain data to a IBM Storwize V5000, import them as image mode volumes to ensure data integrity or use the migration wizard.

1.5.10 iSCSI

iSCSI is an alternative method of attaching hosts to the IBM Storwize V5000. The iSCSI function is a software function that is provided by the IBM Storwize V5000 code, not hardware.

In the simplest terms, iSCSI allows the transport of SCSI commands and data over an Internet Protocol network that is based on IP routers and Ethernet switches. iSCSI is a block-level protocol that encapsulates SCSI commands into TCP/IP packets and uses an existing IP network instead of requiring FC HBAs and a SAN fabric infrastructure.

Concepts of names and addresses are carefully separated in iSCSI.

An iSCSI name is a location-independent, permanent identifier for an iSCSI node. An iSCSI node has one iSCSI name, which stays constant for the life of the node. The terms initiator name and target name also refer to an iSCSI name.

An iSCSI address specifies the iSCSI name of an iSCSI node and a location of that node. The address consists of a host name or IP address, a TCP port number (for the target), and the iSCSI name of the node. An iSCSI node can have any number of addresses, which can change at any time, particularly if they are assigned by way of Dynamic Host Configuration Protocol (DHCP). An IBM Storwize V5000 node represents an iSCSI node and provides statically allocated IP addresses.

Each iSCSI node, that is, an initiator or target, has a unique IQN, which can have a size of up to 255 bytes. The IQN is formed according to the rules that were adopted for Internet nodes. The IQNs can be abbreviated by using a descriptive name, which is known as an alias. An alias can be assigned to an initiator or a target.

For more information about configuring iSCSI, see Chapter 4, “Host configuration” on page 153.

1.5.11 SAS

The SAS standard is an alternative method of attaching hosts to the IBM Storwize V5000. The IBM Storwize V5000 supports direct SAS host attachment providing easy-to-use, affordable storage needs. Each SAS port device has a worldwide unique 64-bit SAS address.

1.6 IBM Storwize V5000 features

In this section, we describe the features of the IBM Storwize V5000.

1.6.1 Mirrored volumes

IBM Storwize V5000 provides a function that is called *storage volume mirroring*, which enables a volume to have two physical copies. Each volume copy can belong to a different storage pool, be generic or a thin-provisioned, and be on different physical storage systems, which provides a high-availability solution.

When a host system issues a write to a mirrored volume, IBM Storwize V5000 writes the data to both copies. When a host system issues a read to a mirrored volume, IBM Storwize V5000 requests it from the primary copy. If one of the mirrored volume copies is temporarily unavailable, the IBM Storwize V5000 automatically uses the alternative copy without any outage for the host system. When the mirrored volume copy is repaired, IBM Storwize V5000 resynchronizes the data.

A mirrored volume can be converted into a non-mirrored volume by deleting one copy or by splitting away one copy to create a non-mirrored volume.

The mirrored volume copy can be any type: image, striped, sequential, and thin-provisioned or not. The two copies can be different volume types.

The use of mirrored volumes also can assist with migrating volumes between storage pools that have different extent sizes. Mirrored volumes also can provide a mechanism to migrate fully allocated volumes to thin-provisioned volumes without any host outages.

The Volume Mirroring feature is included as part of the base software and no license is required.

1.6.2 Thin provisioning

Volumes can be configured to be *thin-provisioned* or *fully allocated*. Concerning application reads and writes, a thin-provisioned volume behaves as though they were fully allocated. When a volume is created, the user specifies two capacities: the real capacity of the volume and its virtual capacity.

The real capacity determines the quantity of MDisk extents that are allocated for the volume. The virtual capacity is the capacity of the volume that is reported to IBM Storwize V5000 and to the host servers.

The real capacity is used to store the user data and the metadata for the thin-provisioned volume. The real capacity can be specified as an absolute value or a percentage of the virtual capacity.

The thin provisioning feature can be used on its own to create over-allocated volumes, or it can be used with FlashCopy. Thin-provisioned volumes can be used with the mirrored volume feature as well.

A thin-provisioned volume can be configured to *autoexpand*, which causes the IBM Storwize V5000 to automatically expand the real capacity of a thin-provisioned volume as its real capacity is used. This parameter prevents the volume from going offline. Autoexpand attempts to maintain a fixed amount of unused real capacity on the volume. This amount is known as the *contingency capacity*. The default setting is 2%.

The contingency capacity initially is set to the real capacity that is assigned when the volume is created. If the user modifies the real capacity, the contingency capacity is reset to be the difference between the used capacity and real capacity.

A volume that is created with a zero contingency capacity goes offline when it must expand. A volume with a non-zero contingency capacity stays online until it is used up.

Autoexpand does not cause the real capacity to grow much beyond the virtual capacity. The real capacity can be manually expanded to more than the maximum that is required by the current virtual capacity and the contingency capacity is recalculated.

To support the autoexpansion of thin-provisioned volumes, the storage pools from which they are allocated have a configurable warning capacity. When the used free capacity of the group exceeds the warning capacity, a warning is logged. For example, if a warning of 80% is specified, the warning is logged when 20% of the free capacity remains.

A thin-provisioned volume can be converted to a fully allocated volume by using volume mirroring (and vice versa).

The Thin Provisioning feature is included as part of the base software and no license is required.

1.6.3 Easy Tier

IBM Easy Tier provides a mechanism to seamlessly migrate hot spots to the most appropriate tier within the IBM Storwize V5000 solution. This migration can be to different tiers of internal drive within IBM Storwize V5000 or to external storage systems that are virtualized by IBM Storwize V5000.

The Easy Tier function can be turned on or off at the storage pool and volume level.

It is possible to demonstrate the potential benefit of Easy Tier in your environment before installing SSDs by using the IBM Storage Advisor Tool.

For more information about Easy Tier, see Chapter 9, “Easy Tier” on page 411.

The IBM Easy Tier feature is licensed per enclosure.

1.6.4 Storage Migration

By using the IBM Storwize V5000 Storage Migration feature, you can easily move data from other Fibre Channel attached external storage to the internal capacity of the IBM Storwize V5000. Migrating data from other storage to the IBM Storwize V5000 storage system provides benefit from more functionality, such as, the easy-to-use GUI, internal virtualization, thin provisioning, and Copy Services.

The Storage Migration feature is included as part of the base software and no license is required.

1.6.5 FlashCopy

FlashCopy copies a source volume on to a target volume. The original contents of the target volume is lost. After the copy operation starts, the target volume has the contents of the source volume as it existed at a single point in time. Although the copy operation takes time, the resulting data at the target appears as though the copy was made instantaneously.

FlashCopy is sometimes described as an instance of a time-zero (T0) copy or a point-in-time (PiT) copy technology.

FlashCopy can be performed on multiple source and target volumes. FlashCopy permits the management operations to be coordinated so that a common single point-in-time is chosen for copying target volumes from their respective source volumes.

IBM Storwize V5000 also permits multiple target volumes to be FlashCopied from the same source volume. This capability can be used to create images from separate points in time for the source volume, and to create multiple images from a source volume at a common point in time. Source and target volumes can be thin-provisioned volumes.

Reverse FlashCopy enables target volumes to become restore points for the source volume without breaking the FlashCopy relationship and without waiting for the original copy operation to complete. IBM Storwize V5000 supports multiple targets and thus multiple rollback points.

The FlashCopy feature is licensed per enclosure.

For more information about FlashCopy copy services, see Chapter 10, “Copy services” on page 449.

1.6.6 Remote Copy

The Remote Copy can be maintained in one of two modes: synchronous or asynchronous.

With the IBM Storwize V5000, Metro Mirror and Global Mirror are the IBM branded terms for the functions that are synchronous Remote Copy and asynchronous Remote Copy.

By using the Metro Mirror and Global Mirror Copy Services features, you can set up a relationship between two volumes so that updates that are made by an application to one volume are mirrored on the other volume. The volumes can be in the same system or on two different systems.

For both Metro Mirror and Global Mirror copy types, one volume is designated as the primary and the other volume is designated as the secondary. Host applications write data to the primary volume and updates to the primary volume are copied to the secondary volume. Normally, host applications do not perform I/O operations to the secondary volume.

The Metro Mirror feature provides a synchronous copy process. When a host writes to the primary volume, it does not receive confirmation of I/O completion until the write operation completes for the copy on the primary and secondary volumes. This ensures that the secondary volume is always up-to-date with the primary volume if a failover operation must be performed.

The Global Mirror feature provides an asynchronous copy process. When a host writes to the primary volume, confirmation of I/O completion is received before the write operation completes for the copy on the secondary volume. If a failover operation is performed, the application must recover and apply any updates that were not committed to the secondary volume. If I/O operations on the primary volume are paused for a brief time, the secondary volume can become an exact match of the primary volume.

Global Mirror can operate with or without cycling. When it is operating without cycling, write operations are applied to the secondary volume as soon as possible after they are applied to the primary volume. The secondary volume is less than one second behind the primary volume, which minimizes the amount of data that must be recovered in the event of a failover. However, this requires that a high-bandwidth link is provisioned between the two sites.

When Global Mirror operates with cycling mode, changes are tracked and where needed copied to intermediate change volumes. Changes are transmitted to the secondary site periodically. The secondary volumes are much further behind the primary volume, and more data must be recovered in the event of a failover. Because the data transfer can be smoothed over a longer time period, however, lower bandwidth is required to provide an effective solution.

For more information about the IBM Storwize V5000 Copy Services, see Chapter 10, “Copy services” on page 449).

The IBM Remote Copy feature is licensed per enclosure.

Copy Services configuration limits

For the most up-to-date list of these limits, see the following website:

http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003702&myns=s028&mynp=familyind5402112&mync=E

1.6.7 External virtualization

By using this feature, you can consolidate FC SAN-attached disk controllers from various vendors into pools of storage. In this way, the storage administrator can manage and provision storage to applications from a single user interface and use a common set of advanced functions across all the storage systems under the control of the IBM Storwize V5000.

The External Virtualization feature is licensed per disk enclosure.

1.7 Problem management and support

In this section, we introduce problem management and support topics.

1.7.1 IBM Assist On-site and remote service

The IBM Assist On-site tool is a remote desktop-sharing solution that is offered through the IBM website. With it, the IBM service representative can remotely view your system to troubleshoot a problem.

You can maintain a chat session with the IBM service representative so that you can monitor this activity and understand how to fix the problem yourself or allow the representative to fix it for you.

To use the IBM Assist On-site tool, the management PC that is used to manage the IBM Storwize V5000 must have access the Internet. For more information about this tool, see this website:

http://www.ibm.com/support/assistonsite/

When you access the website, you sign in and enter a code that the IBM service representative provides to you. This code is unique to each IBM Assist On-site session. A plug-in is downloaded on to your PC to connect you and your IBM service representative to the remote service session. The IBM Assist On-site contains several layers of security to protect your applications and your computers.

You also can use security features to restrict access by the IBM service representative.

Your IBM service representative can provide you with more information about the use of the tool, if required.

1.7.2 Event notifications

IBM Storwize V5000 can use Simple Network Management Protocol (SNMP) traps, syslog messages, and a Call Home email to notify you and the IBM Support Center when significant events are detected. Any combination of these notification methods can be used simultaneously.

Each event that IBM Storwize V5000 detects is sent to a different recipient. You can configure IBM Storwize V5000 to send each type of notification to specific recipients or only the alerts that are important to the system.

1.7.3 SNMP traps

SNMP is a standard protocol for managing networks and exchanging messages. IBM Storwize V5000 can send SNMP messages that notify personnel about an event. You can use an SNMP manager to view the SNMP messages that IBM Storwize V5000 sends. You can use the management GUI or the IBM Storwize V5000 CLI to configure and modify your SNMP settings.

You can use the Management Information Base (MIB) file for SNMP to configure a network management program to receive SNMP messages that are sent by the IBM Storwize V5000. This file can be used with SNMP messages from all versions of IBM Storwize V5000 software.

1.7.4 Syslog messages

The syslog protocol is a standard protocol for forwarding log messages from a sender to a receiver on an IP network. The IP network can be IPv4 or IPv6. IBM Storwize V5000 can send syslog messages that notify personnel about an event. IBM Storwize V5000 can transmit syslog messages in expanded or concise format. You can use a syslog manager to view the syslog messages that IBM Storwize V5000 sends. IBM Storwize V5000 uses the User Datagram Protocol (UDP) to transmit the syslog message. You can use the management GUI or the CLI to configure and modify your syslog settings.

1.7.5 Call Home email

The Call Home feature transmits operational and error-related data to you and IBM through a Simple Mail Transfer Protocol (SMTP) server connection in the form of an event notification email. When configured, this function alerts IBM service personnel about hardware failures and potentially serious configuration or environmental issues. You can use the call home function if you have a maintenance contract with IBM or if the IBM Storwize V5000 is within the warranty period.

To send email, you must configure at least one SMTP server. You can specify as many as five other SMTP servers for backup purposes. The SMTP server must accept the relaying of email from the IBM Storwize V5000 clustered system IP address. You can then use the management GUI or the CLI to configure the email settings, including contact information and email recipients. Set the reply address to a valid email address. Send a test email to check that all connections and infrastructure are set up correctly. You can disable the Call Home function at any time by using the management GUI or the CLI.

1.8 More information resources

This section describes resources that are available for more information.

1.8.1 Useful IBM Storwize V5000 websites

For more information about Storwize V5000, see the following websites:

- The IBM Storwize V5000 home page:

  http://www.ibm.com/storage/support/storwize/v5000
- IBM Storwize V5000 Online Information Center:

  http://pic.dhe.ibm.com/infocenter/storwize/v5000_ic/index.jsp

1.8.2 IBM Storwize learning videos on YouTube

Videos are available on YouTube that describe the IBM Storwize GUI and are available at the URLs that are shown in Table 1-5.

Table 1-5 Videos available on YouTube

Video description	URL
IBM Storwize V7000 Storage Virtualization Terminology Overview	http://www.youtube.com/watch?v=I2rzt3m2gPO
IBM Storwize V7000 Interface tour	http://www.youtube.com/watch?v=FPbNRs9HacQ
IBM Storwize V7000 Volume management	http://www.youtube.com/watch?v=YXeKqH8Sd9o
IBM Storwize V7000 Migration	http://www.youtube.com/watch?v=dXxnUN6dk74
IBM Storwize V7000 Introduction to FlashCopy	http://www.youtube.com/watch?v=MXWgGWjBzG4
VMware data protection with Storwize V7000	http://www.youtube.com/watch?v=vecOap-qwbA
IBM SAN Volume Controller and Storwize V7000 Performance Panel Sped-up! (HD)	http://www.youtube.com/watch?v=7noC71tLkWs
IBM Storwize V3700 Hardware Installation	http://www.youtube.com/watch?v=VuEfmfXihrs
IBM Storwize V3700 - Effortless Management	http://www.youtube.com/watch?v=BfGbKWcCsR4
Introducing IBM Storwize V3700	http://www.youtube.com/watch?v=AePPKiXE4xM
IBM Storwize V3700 Initial Setup	http://www.youtube.com/watch?v=oj9uhTYe6gg
Storwize V7000 Installation	http://www.youtube.com/watch?v=kCCFxM5ZMV4

These videos are applicable to IBM Storwize V5000 because the GUI interface on the IBM Storwize V3700 and IBM Storwize V7000 is similar. The IBM Storwize V3700 hardware also is similar and the videos provide a good overview of the functions and features.

2

Initial configuration

This chapter provides a description of the initial configuration steps for the IBM Storwize V5000.

This chapter includes the following topics:

- Planning for IBM Storwize V5000 installation
- First time setup
- Initial configuration steps
- Call Home, email event alert, and inventory settings

2.1 Hardware installation planning

Proper planning before the actual physical installation of the hardware is required. The following checklist of requirements can be used to plan your installation:

- ☐ Install the hardware as described in *IBM Storwize V5000 Quick Installation Guide Version 6.4.1*, GC27-4219
- ☐ For more information about planning the IBM Storwize V5000 environment, see this website:

  `http://pic.dhe.ibm.com/infocenter/storwize/V5000_ic/index.jsp?topic=%2Fcom.ibm.storwize.V5000.641.doc%2Fsvc_webplanning_21pb8b.html`
- ☐ An appropriate 19-inch rack with 2 - 14 U of space should be available, depending on the number of enclosures to install. Each enclosure measures 2 U and a single control enclosure with up to six expansion enclosures constitutes an IBM Storwize V5000 system.
- ☐ There should be redundant power outlets in the rack for each of the two power cords that are included per enclosure. In all, 2 - 14 outlets are required, depending on the number of enclosures to install. The power cords conform to the IEC320 C13/C14 standards.
- ☐ A minimum of four Fibre Channel ports that are attached to the fabric are required. However, it is a best practice to use eight 2-Gbps, 4-Gbps, or 8-Gbps Fibre Channel ports.

  Fibre Channel ports: Fibre Channel (FC) ports are required only if you are using FC hosts. You can use the Storwize V5000 with Ethernet-only cabling for iSCSI hosts or use serial-attached SCSI (SAS) cabling for direct attach hosts.

- ☐ You should have eight 2-Gbps, 4-Gbps, or 8-Gbps compatible Fibre Channel cable drops.
- ☐ Up to four hosts can be directly connected by using SAS ports 1 and 2 on each node canister, with SFF-8644 mini SAS HD cabling.
- ☐ You should have a minimum of two Ethernet ports on the LAN, with four preferred for more configuration access redundancy or iSCSI host access.
- ☐ You should have a minimum of two Ethernet cable drops, with four preferred for more configuration access redundancy or iSCSI host access. Ethernet port one on each node canister must be connected to the LAN, with port two as optional.

  Ports: Port 1 on each node canister must be connected to the same physical LAN or be configured in the same VLAN and be on the same subnet or set of subnets.

- ☐ Verify that the default IP addresses that are configured on Ethernet port 1 on each of the node canisters (192.168.70.121 on node one and 192.168.70.122 on node 2) do not conflict with existing IP addresses on the LAN. The default mask that is used with these IP addresses is 255.255.255.0 and the default gateway address that is used is 192.168.70.1.
- ☐ You should have a minimum of three IPv4 or IPv6 IP addresses for system configuration. One is for the clustered system and is what the administrator uses for management, and one for each node canister for service access as needed.

  IP addresses: A fourth IP address should be used for backup configuration access. This other IP address allows a second system IP address to be configured on port 2 of either node canister, which the storage administrator can also use for management of the IBM Storwize V5000 system.

- A minimum of one and up to four IPv4 or IPv6 addresses are needed if iSCSI attached hosts access volumes from the IBM Storwize V5000.
- A single 1-meter, 3-meter, or 6-meter SAS cable per expansion enclosure is required. The length of the cables depends on the physical rack location of the expansion relative to the control enclosure or other expansion enclosures. Locate the control enclosure so that up to six enclosures can be located, as shown in Figure 2-1 on page 30. The IBM Storwize V5000 supports two external SAS chains by using SAS ports 3 and 4 on the control enclosure node canisters.

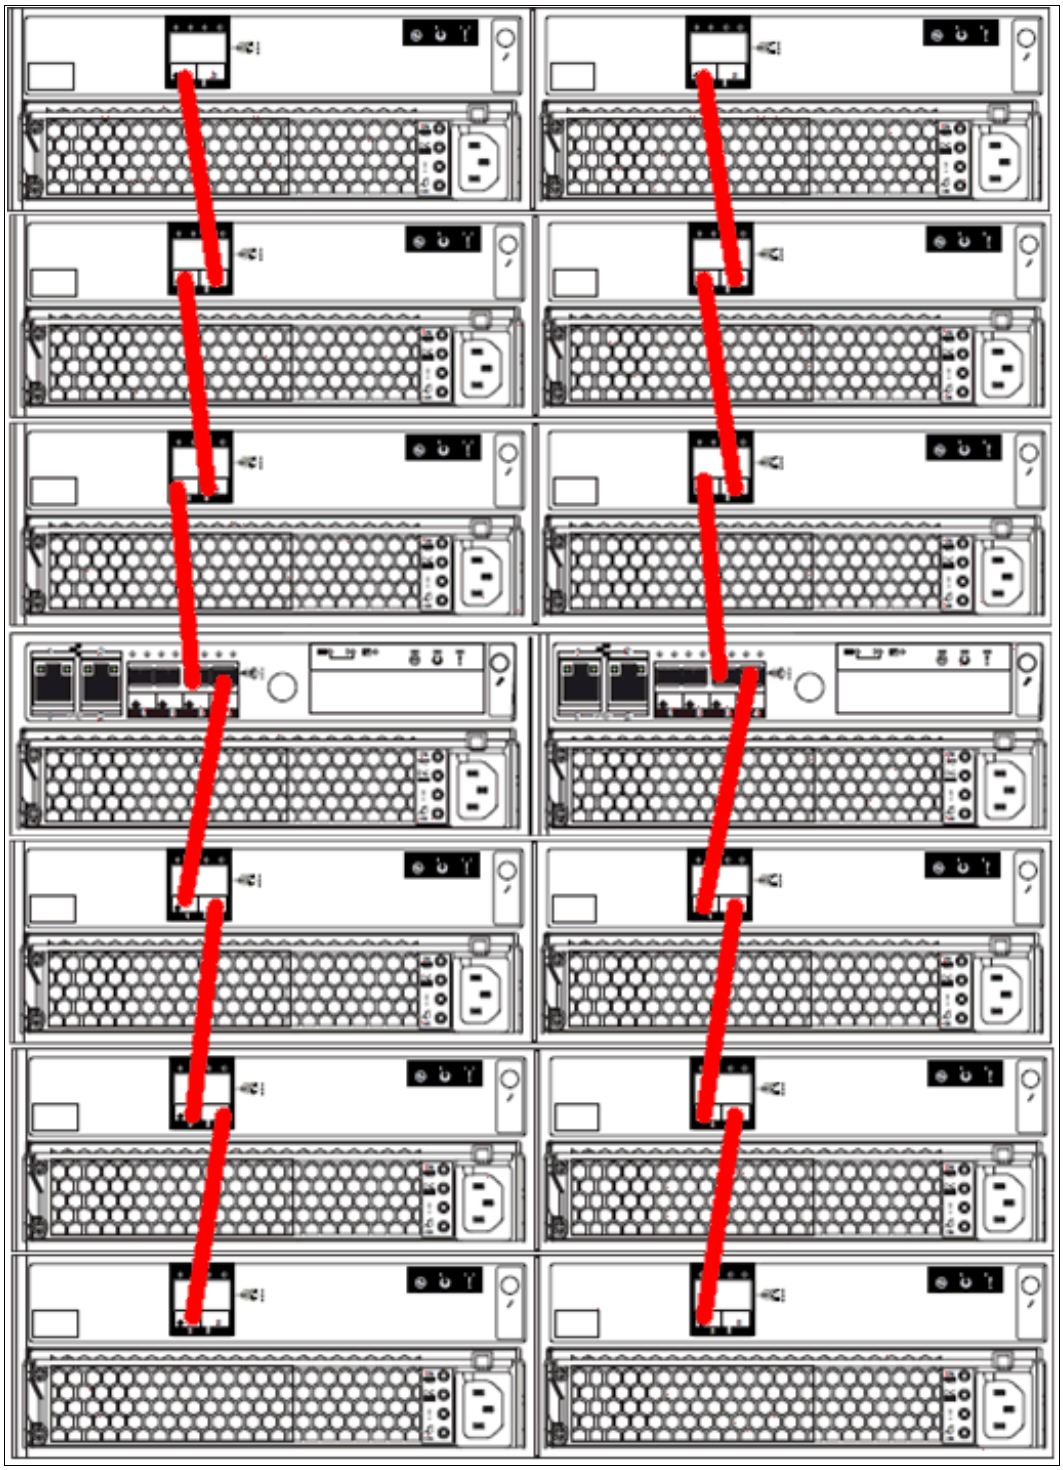

Figure 2-1 Connecting the SAS expansion cables example

The following connections must be made:

- Connect SAS port 4 of the left node canister in the control enclosure to SAS port 1 of the left expansion canisters in the first expansion enclosure.
- Connect SAS port 4 of the right node canister in the control enclosure to SAS port 1 of the right expansion canisters in the first expansion enclosure.

– Connect SAS port 3 of the left node canister in the control enclosure to SAS port 1 of the left expansion canister in the second enclosure (above the control enclosure, as shown in Figure 2-1 on page 30).
– Connect SAS port 3 of the right node canister in the control enclosure to SAS port 1 of the right expansion canister in the second enclosure (above the control enclosure, as shown in Figure 2-1 on page 30).

Continue to add expansion controllers alternately on the two different SAS chains that originate at ports 3 and 4 on the control enclosure node canisters. No expansion enclosure should be connected to both SAS chains.

Disk drives: The disk drives that are included with the control enclosure (model 2077-12C or 2077-24C) are part of the single SAS chain. The expansion enclosures should be connected to the SAS chain as shown in Figure 2-1 on page 30 so that they can use the full bandwidth of the system.

2.2 SAN configuration planning

The recommended SAN configuration is composed of a minimum of two fabrics that encompass all host ports and any ports on external storage systems that are to be virtualized by IBM Storwize V5000. The IBM Storwize V5000 ports are evenly split between the two fabrics to provide redundancy if one of the fabrics goes offline (planned or unplanned).

Virtualized Storage: External storage systems that are to be virtualized are used for migration purposes only.

Zoning must be implemented after the IBM Storwize V5000, hosts, and optional external storage systems are connected to the SAN fabrics.

To enable the node canisters to communicate with each other in band, create a zone with only the IBM Storwize V5000 WWPNs (two from each node canister) on each of the two fabrics. If an external storage system is to be virtualized, create a zone in each fabric with the IBM Storwize V5000 WWPNs (two from each node canister) with up to a maximum of eight WWPNs from the external storage system. Assuming every host has a Fibre Channel connection to each fabric, create a zone with the host WWPN and one WWPN from each node canister in the IBM Storwize V5000 system in each fabric. The critical point is that there should only ever be one initiator (host HBA) in any zone. For load balancing between the node ports on the IBM Storwize V5000, alternate the host Fibre Channel ports between the ports of the Storwize V5000.

There should be a maximum of eight paths through the SAN from each host to the IBM Storwize V5000. Hosts where this number is exceeded are not supported. The restriction is there to limit the number of paths that the multi-pathing driver must resolve. A host with only two HBAs should not exceed this limit with proper zoning in a dual fabric SAN.

Maximum ports or WWPNs: IBM Storwize V5000 supports a maximum of 16 ports or WWPNs from a virtualized external storage system.

Figure 2-2 shows how to cable devices to the SAN. Refer to this example as the zoning is described.

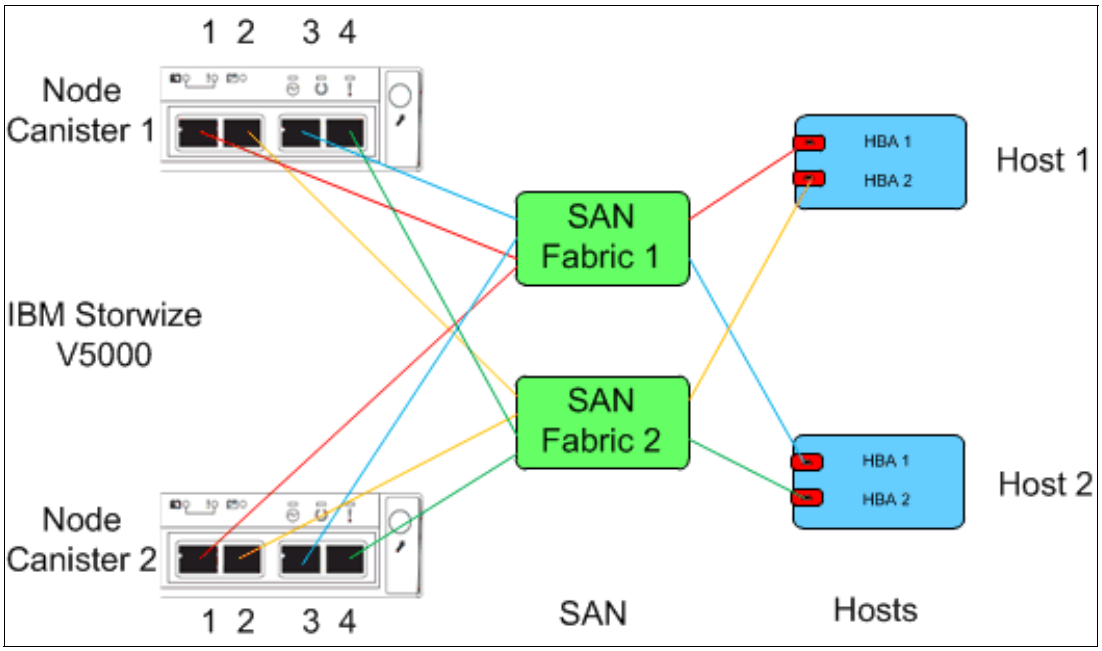

Figure 2-2 SAN cabling and zoning diagram

Create a host/IBM Storwize V5000 zone for each server that volumes are mapped to and from the clustered system, as shown in the following examples in Figure 2-2:

- Zone Host 1 port 1 (HBA 1) with both node canister ports 1
- Zone Host 1 port 2 (HBA 2) with both node canister ports 2
- Zone Host 2 port 1 (HBA 1) with both node canister ports 3
- Zone Host 2 port 2 (HBA 2) with both node canister ports 4

Similar zones should be created for all other hosts with volumes on the Storwize V5000.

Verify interoperability with which the IBM Storwize V5000 connects to SAN switches or directors by following the requirements that are provided at this website:

`http://www-01.ibm.com/support/docview.wss?uid=ssg1S1004111`

Switches or directors are at the firmware levels that are supported by the IBM Storwize V5000.

Important: IBM Storwize V5000 port login maximum that is listed in the restriction document must not be exceeded. The document is available at this website:

`http://www-01.ibm.com/support/docview.wss?uid=ssg1S1004233`

Connectivity issues: If you have any connectivity issues between IBM Storwize V5000 ports and Brocade SAN Switches or Directors at 8 Gbps, see this website for the correct setting of the fillword port config parameter in the Brocade operating system:

`http://www-01.ibm.com/support/docview.wss?rs=591&uid=ssg1S1003699`

2.3 FC Direct-attach planning

IBM Storwize V5000 can be used with a direct-attach Fibre Channel host configuration. The recommended configuration for direct attachment is to have at least one Fibre Channel cable from the host that is connected to each node of the IBM Storwize V5000 to provide redundancy if one of the nodes goes offline, as shown in Figure 2-3.

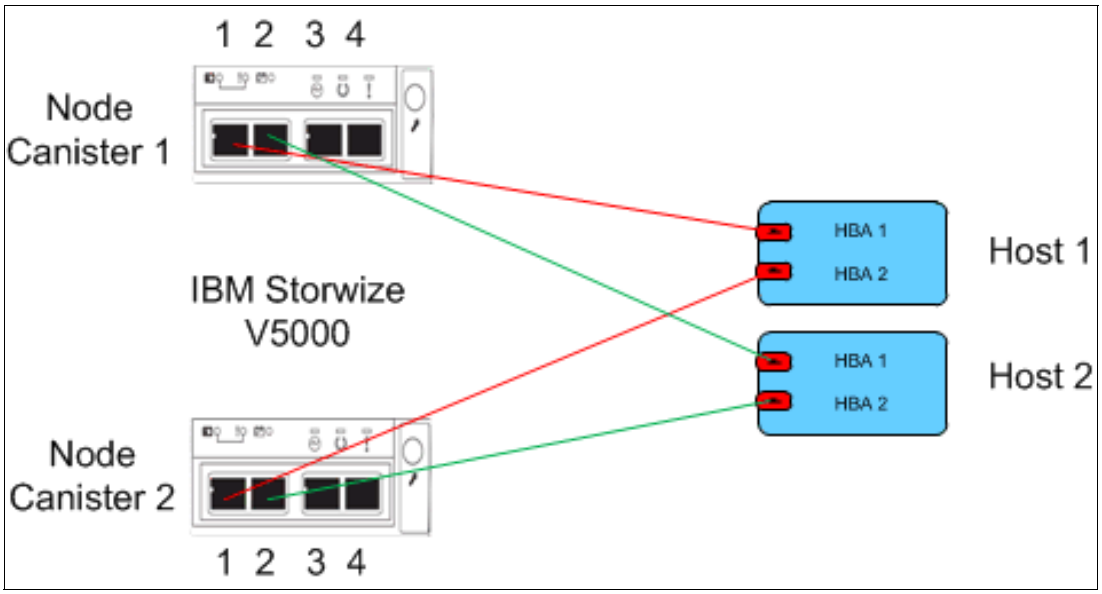

Figure 2-3 FC Direct-attach host configuration

Verify direct attach interoperability with the IBM Storwize V5000 and the supported server operating systems by following the requirements that are provided at this website:

`http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss`

2.4 SAS Direct-attach planning

There are two SAS ports per node canister that are available for direct host attach on an IBM Storwize V5000. These are ports 1 and 2. Do not use ports 3 and 4 because they are reserved for expansion enclosure connectivity only. Refer to Figure 2-4 on page 34 to correctly identify ports 1 and 2. Also, note the keyway in the top of the SAS connector.

Inserting cables: It is possible to insert the cables upside down despite the keyway. Ensure that the blue tag on the SAS connector is underneath when you are inserting the cables.

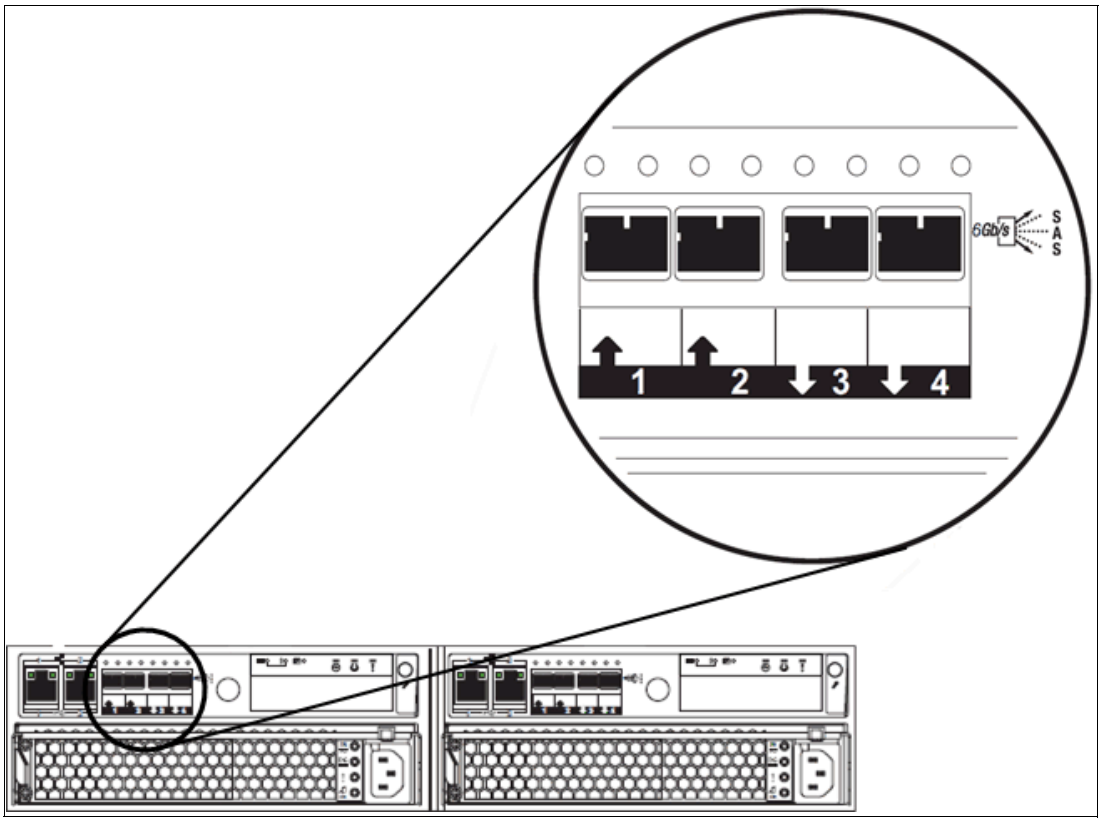

Figure 2-4 SAS port identification

Although it is possible to attach four hosts (one to each of the two available SAS ports on the two node canisters), the recommended configuration for direct attachment is to have at least one SAS cable from the host that is connected to each node of the IBM Storwize V5000. This configuration provides redundancy if one of the nodes goes offline, as shown in Figure 2-5.

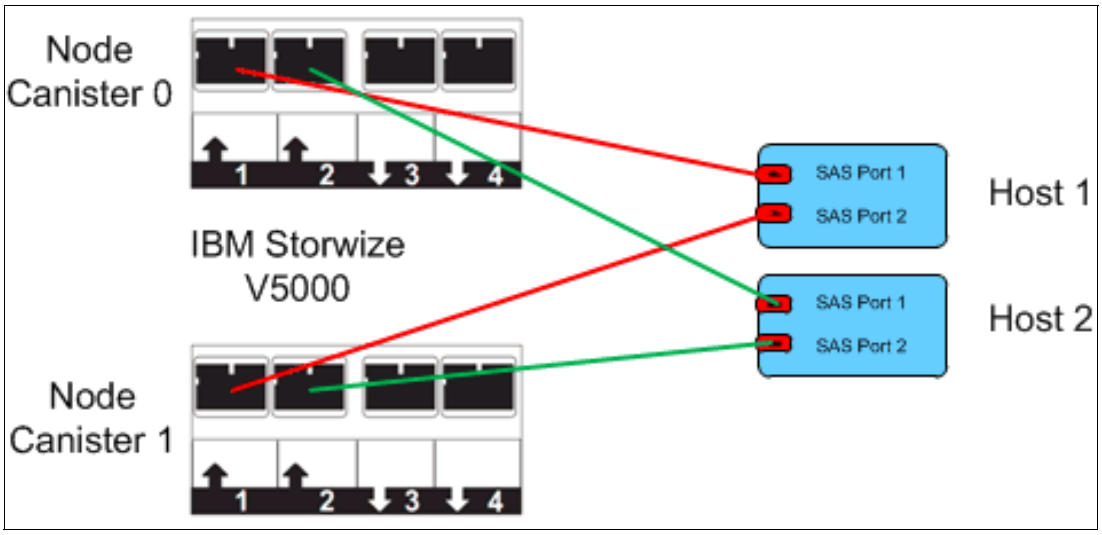

Figure 2-5 SAS host direct-attach

2.5 LAN configuration planning

There are two Ethernet ports per node canister that are available for connection to the LAN on an IBM Storwize V5000 system.

Ethernet port 1 is for accessing the management GUI, the service assistant GUI for the node canister, and iSCSI host attachment. Port 2 can be used for the management GUI and iSCSI host attachment.

Each node canister in a control enclosure connects over an Ethernet cable from Ethernet port 1 of the canister to an enabled port on your Ethernet switch or router. Optionally, you can attach an Ethernet cable from Ethernet port 2 on the canister to your Ethernet network.

Configuring IP addresses: There is no issue with configuring multiple IPv4 or IPv6 addresses on an Ethernet port or the use of the same Ethernet port for management and iSCSI access. However, you cannot use the same IP address for management and iSCSI host use.

Table 2-1 shows possible IP configuration of the Ethernet ports on the IBM Storwize V5000 system.

Table 2-1 Storwize V5000 IP address configuration options per node canister

Storwize V5000 Management Node Canister 1		**Storwize V5000 Partner Node Canister 2**	
IPv4/6 management address	ETH PORT 1	IPv4/6 service address	ETH PORT 1
IPv4/6 service address		IPv4/6 iSCSI address	
IPv4/6 iSCSI address			
IPv4/6 management address	ETH PORT 2	IPv4/6iSCSI address	ETH PORT 2
IPv4/6 iSCSI address			

IP management addresses: The IP management address that is shown on Node Canister 1 in Table 2-1 is an address on the configuration node. If a failover occurs, this address transfers to Node Canister 2 and this node canister becomes the new configuration node. The management addresses are managed by the configuration node canister only (1 or 2; in this case, by Node Canister 1).

2.5.1 Management IP address considerations

Because Ethernet port 1 from each node canister must be connected to the LAN, a single management IP address for the clustered system is configured as part of the initial setup of the IBM Storwize V5000 system.

The management IP address is associated with one of the node canisters in the clustered system and that node then becomes the configuration node. Should this node go offline (planned or unplanned), the management IP address fails over to the other node's Ethernet port 1.

For more clustered system management redundancy, you should connect Ethernet port 2 on each of the node canisters to the LAN, which allows for a backup management IP address to be configured for access, if necessary.

Figure 2-6 shows a logical view of the Ethernet ports that are available for configuration of the one or two management IP addresses. These IP addresses are for the clustered system and therefore are associated with only one node, which is then considered the configuration node.

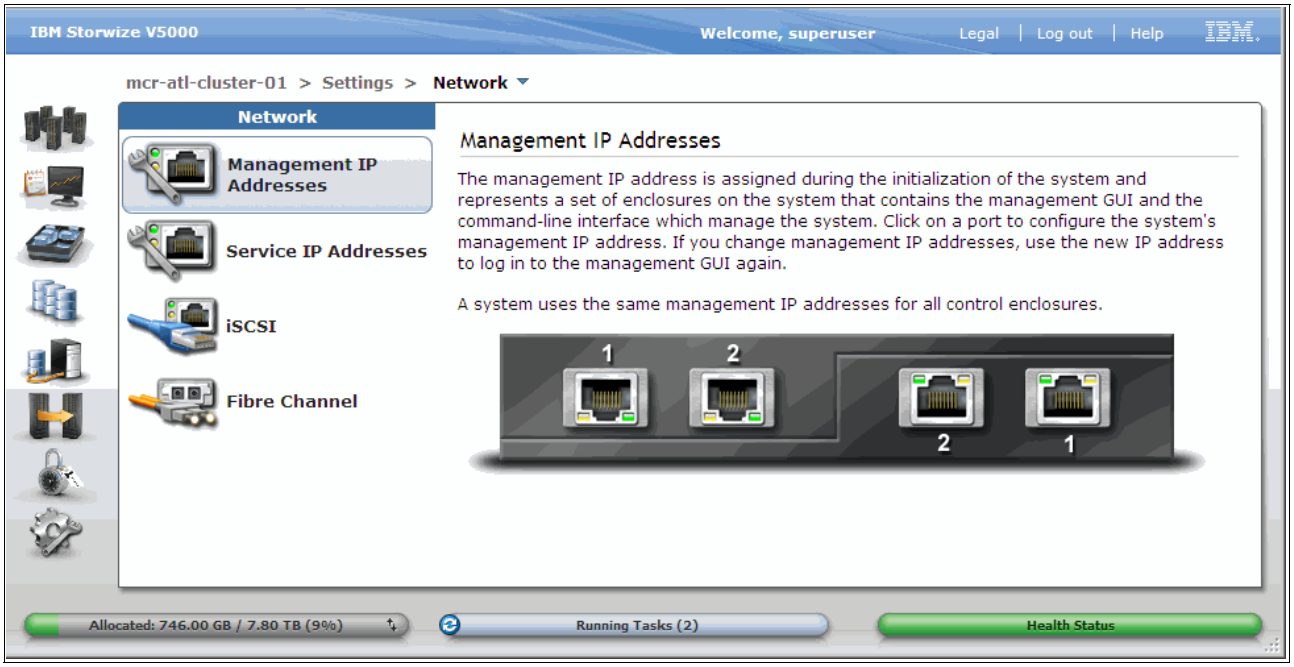

Figure 2-6 Ethernet ports available for configuration

2.5.2 Service IP address considerations

Ethernet port 1 on each node canister is used for system management and for service access, when required. In normal operation, the service IP addresses are not needed. However, if there is a node canister problem, it might be necessary for service personnel to log on to the node to perform service actions.

Figure 2-7 on page 37 shows a logical view of the Ethernet ports that are available for configuration of the service IP addresses. Only port one on each node can be configured with a service IP address.

Figure 2-7 Service IP addresses available for configuration

2.6 Host configuration planning

Hosts should have two Fibre Channel connections for redundancy, but the IBM Storwize V5000 also supports hosts with a single HBA port connection. However, if that HBA, its link to the SAN fabric or the fabric fails, the host loses access to its volumes. Even with a single connection to the SAN, the host has multiple paths to the IBM Storwize V5000 volumes because that single connection must be zoned with at least one Fibre Channel port per node. Therefore, a multipath driver is required. This is also true for direct-attach SAS hosts. They can be connected by using a single host port that allows up to four hosts to be configured, but for redundancy two SAS connections per host are recommended. If two connections per host are used, a multipath driver also is required on the host. If iSCSI host is to be employed, they also require an MPIO driver. Both node canisters should be configured and connected to the network so any iSCSI hosts see at least two paths to volumes and an MPIO driver are required to resolve these.

SAN Boot is supported by IBM Storwize V5000. For more information, see the IBM Storwize V5000 Information Center at this website:

`http://pic.dhe.ibm.com/infocenter/storwize/v5000_ic/index.jsp`

Verify that the hosts that access volumes from the IBM Storwize V5000 meet the requirements that are found at this website:

`http://www-947.ibm.com/support/entry/portal/overview/hardware/system_storage/disk_systems/entry-level_disk_systems/ibm_storwize_v3700`

Multiple operating systems are supported by IBM Storwize V5000. For more information about HBA/Driver/multipath combinations, see this website:

`http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss`

As per the IBM System Storage Interoperation Center (SSIC), keep the following items under consideration:

- Host operating systems are at the levels that are supported by the IBM Storwize V5000.
- HBA BIOS, device drivers, firmware, and multipathing drivers are at the levels that are supported by IBM Storwize V5000.
- If boot from SAN is required, ensure that it is supported for the operating systems that are deployed.
- If host clustering is required, ensure that it is supported for the operating systems that are deployed.
- All direct connect hosts should have the HBA set to point-to-point.

For more information, see Chapter 4, “Host configuration” on page 153.

2.7 Miscellaneous configuration planning

During the initial setup of the IBM Storwize V5000 system, the installation wizard asks for various information that you should have available during the installation process. Several of these fields are mandatory to complete the initial configuration.

The information in the following checklist is helpful to have before the initial setup is performed. The date and time can be manually entered, but to keep the clock synchronized, use a network time protocol (NTP) service:

- ☐ Document the LAN NTP server IP address that is used for synchronization of devices.
- ☐ For alerts to be sent to storage administrators and to set up Call Home to IBM for service and support, you need the following information:
  - ☐ Name of primary storage administrator for IBM to contact, if necessary.
  - ☐ Email address of the storage administrator for IBM to contact, if necessary.
  - ☐ Phone number of the storage administrator for IBM to contact, if necessary.
  - ☐ Physical location of the IBM Storwize V5000 system for IBM service (for example, Building 22, first floor).
  - ☐ SMTP or email server address to direct alerts to and from the IBM Storwize V5000.
  - ☐ For the Call Home service to work, the IBM Storwize V5000 system must have access to an SMTP server on the LAN that can forward emails to the default IBM service address: callhome1@de.ibm.com for Americas-based systems and callhome0@de.ibm.com for the rest of the World.
  - ☐ Email address of local administrators that must be notified of alerts.
  - ☐ IP address of SNMP server to direct alerts to, if required (for example, operations or Help desk).

After the IBM Storwize V5000 initial configuration, you might want to add more users who can manage the system. You can create as many users as you need, but the following roles generally are configured for users:

- Security Admin
- Administrator
- CopyOperator
- Service
- Monitor

The user in the Security Admin role can perform any function on the IBM Storwize V5000.

The user in the Administrator role can perform any function on the IBM Storwize V5000 system, except create users.

User creation: The create users function is allowed by the Security Admin role only and should be limited to as few users as possible.

The user in the Copyoperator role can view anything in the system, but the user can configure and manage only the copy functions of the FlashCopy capabilities.

The user in the Monitor role can view object and system configuration information but cannot configure, manage, or modify any system resource.

The only other role that is available is the service role, which is used if you create a user ID for the IBM service representative. This user role allows IBM service personnel to view anything on the system (as with the monitor role) and perform service-related commands, such as, adding a node back to the system after it is serviced or including disks that were excluded.

2.8 System management

The graphical user interface (GUI) is used to configure, manage, and troubleshoot the IBM Storwize V5000 system. It is used primarily to configure RAID arrays and logical drives, assign logical drives to hosts, replace and rebuild failed disk drives, and expand the logical drives.

It allows for troubleshooting and management tasks, such as, checking the status of the storage server components, updating the firmware, and managing the storage server.

The GUI also offers advanced functions, such as, FlashCopy, Volume Mirroring, Remote Mirroring, and EasyTier. A command-line interface (CLI) for the IBM Storwize V5000 system also is available.

This section describes system management by using the GUI and CLI.

2.8.1 GUI

A web browser is used for GUI access. You must use a supported web browser to access the management GUI. For more information about supported web browsers, see *Checking your web browser settings for the management GUI*, which is available at this website:

http://pic.dhe.ibm.com/infocenter/storwize/V5000_ic/index.jsp?topic=%2Fcom.ibm.storwize.V5000.641.doc%2Fsvc_configuringbrowser_1obg15.html

Complete the following steps to open the Management GUI from any web browser:

1. Browse to one of the following locations:
   a. http(s)://*host name of your cluster*/
   b. http(s)://*cluster IP address of your cluster*/ Example: https://192.168.70.120

2. Use the following default login information:
   – User ID: superuser
   – Password: passw0rd

For more information about how to use this interface, see this website:

http://pic.dhe.ibm.com/infocenter/storwize/V5000_ic/index.jsp?topic=%2Fcom.ibm.storwize.V5000.641.doc%2Ftbrd_usbgui_1936tw.html

More information also is available in Chapter 3, “Graphical user interface overview” on page 75.

After the initial configuration that is described in 2.10, “Initial configuration” on page 49 is completed, the IBM Storwize V5000 Welcome window opens, as shown in Figure 2-8.

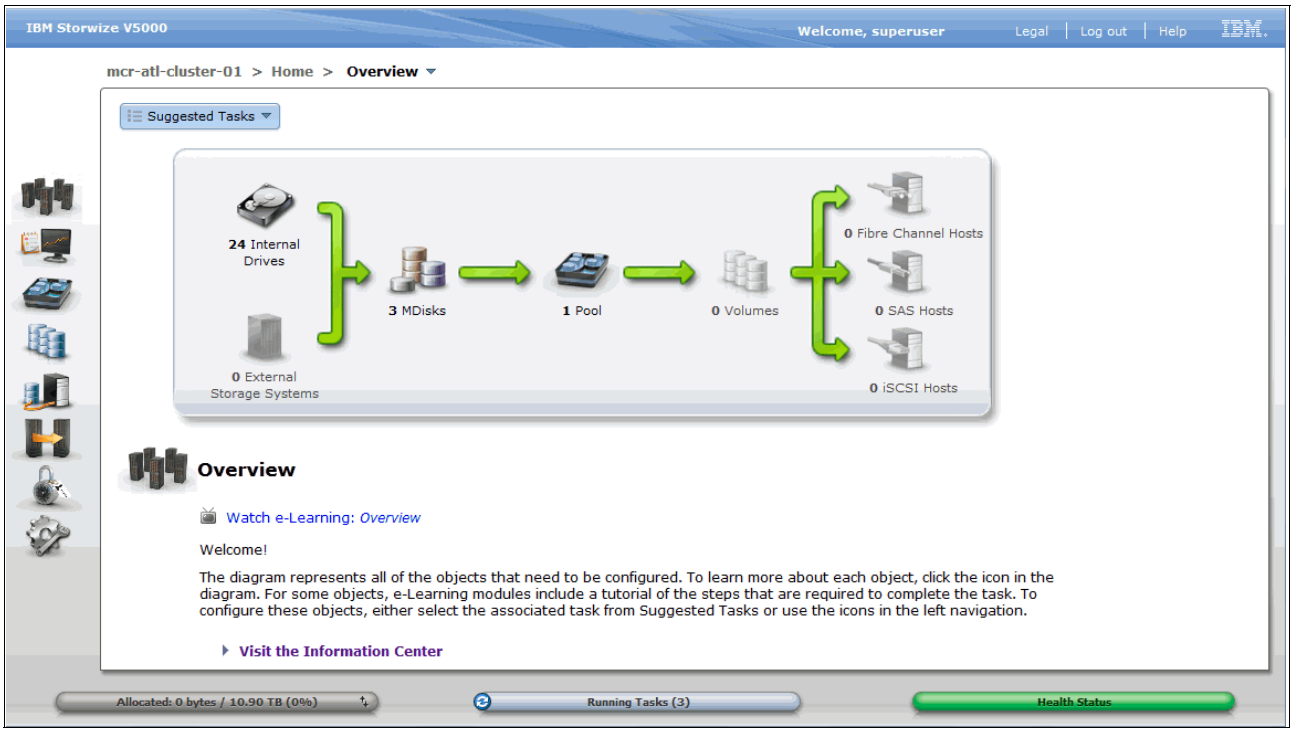

Figure 2-8 Setup wizard: Welcome window

2.8.2 CLI

The CLI is a flexible tool for system management that uses the SSH protocol. A public/private SSH key pair is optional for SSH access. For more information about setting up SSH Access for Windows, Linux, or UNIX systems, see Appendix A, “Command-line interface setup and SAN Boot” on page 609. The storage system can be managed by using the CLI, as shown in Example 2-1.

Example 2-1 System management by using the CLI

```
IBM_Storwize:mcr-atl-cluster-01:superuser>lsenclosureslot
enclosure_id slot_id port_1_status port_2_status drive_present drive_id
1            1       online        online        yes           20
1            2       online        online        yes           22
1            3       online        online        yes           21
1            4       online        online        yes           23
1            5       online        online        yes           17
1            6       online        online        yes           12
1            7       online        online        yes           10
1            8       online        online        yes           18
1            9       online        online        yes           9
```

```
1            10      online       online       yes          11
1            11      online       online       yes          8
1            12      online       online       yes          14
1            13      online       online       yes          15
1            14      online       online       yes          13
1            15      online       online       yes          16
1            16      online       online       yes          19
1            17      online       online       yes          1
1            18      online       online       yes          3
1            19      online       online       yes          6
1            20      online       online       yes          0
1            21      online       online       yes          4
1            22      online       online       yes          7
1            23      online       online       yes          2
1            24      online       online       yes          5
IBM_Storwize:mcr-atl-cluster-01:superuser>
```

The initial IBM Storwize V5000 system setup should be done by using the process and tools that are described in 2.9, “First-time setup” on page 41.

2.9 First-time setup

This section describes how to perform a first-time IBM Storwize V5000 system setup.

IBM Storwize V5000 uses an initial setup process that is contained within a USB key. The USB key is delivered with each storage system and contains the initialization application file that is called `InitTool.exe`. The tool is configured with your IBM Storwize V5000 system management IP address, the subnet mask, and the network gateway address by first plugging the USB stick into a Windows or Linux system.

The IBM Storwize V5000 starts the initial setup when you plug in the USB key with the newly created file in to the storage system.

USB key: If you cannot find the official USB key that is supplied with the IBM Storwize V5000, you can use any USB key that you have and download and copy the `initTool.exe` application from IBM Storwize V5000 Support at this website:

`http://www.ibm.com/storage/support/Storwize/V5000`

The USB stick contains a readme file that provides details about how to use the tool with various operating systems. The following operating systems are supported:

- Microsoft Windows (R) 7 (64-bit)
- Microsoft Windows XP (32-bit only)
- Apple Mac OS(R) X 10.7
- Red Hat Enterprise Server 5
- Ubuntu (R) desktop 11.04

We use Windows in the following examples.

Complete the following steps to perform the initial setup by using the USB key:

1. Plug the USB key into a Windows system and start the initialization tool. If the system is configured to autorun USB keys, the initialization tool starts automatically; otherwise, open My Computer and double-click the **`InitTool.bat`** file. The opening window of the tool is shown in Figure 2-9. After the tool is started, select **Next** and then select **Create a new system**.

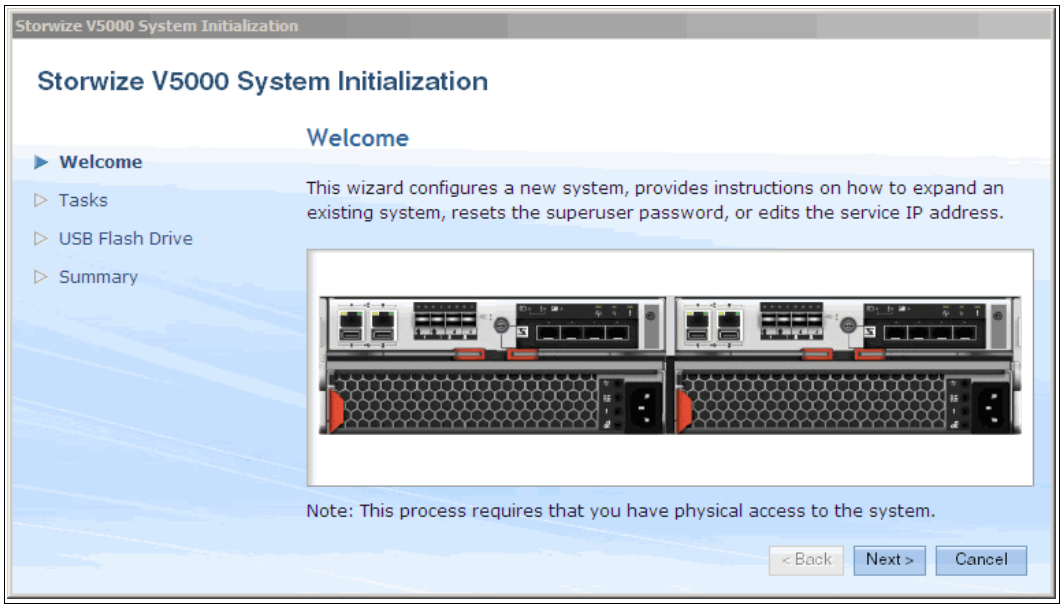

Figure 2-9 System Initialization: Welcome window

Mac OS or Linux: For Mac OS or Linux, complete the following steps:

a. Open a terminal window.

b. Locate the root directory of the USB flash drive:

- For Mac systems, the root directory is often in the `/Volumes/` directory.
- For Linux systems, the root directory is often in the `/media/` directory.
- If an automatic mount system is used, the root directory can be located by entering the mount command.

c. Change the directory to the root directory of the flash drive.

d. Enter: `sh InitTool.sh`

The options for creating a system are shown in Figure 2-10.

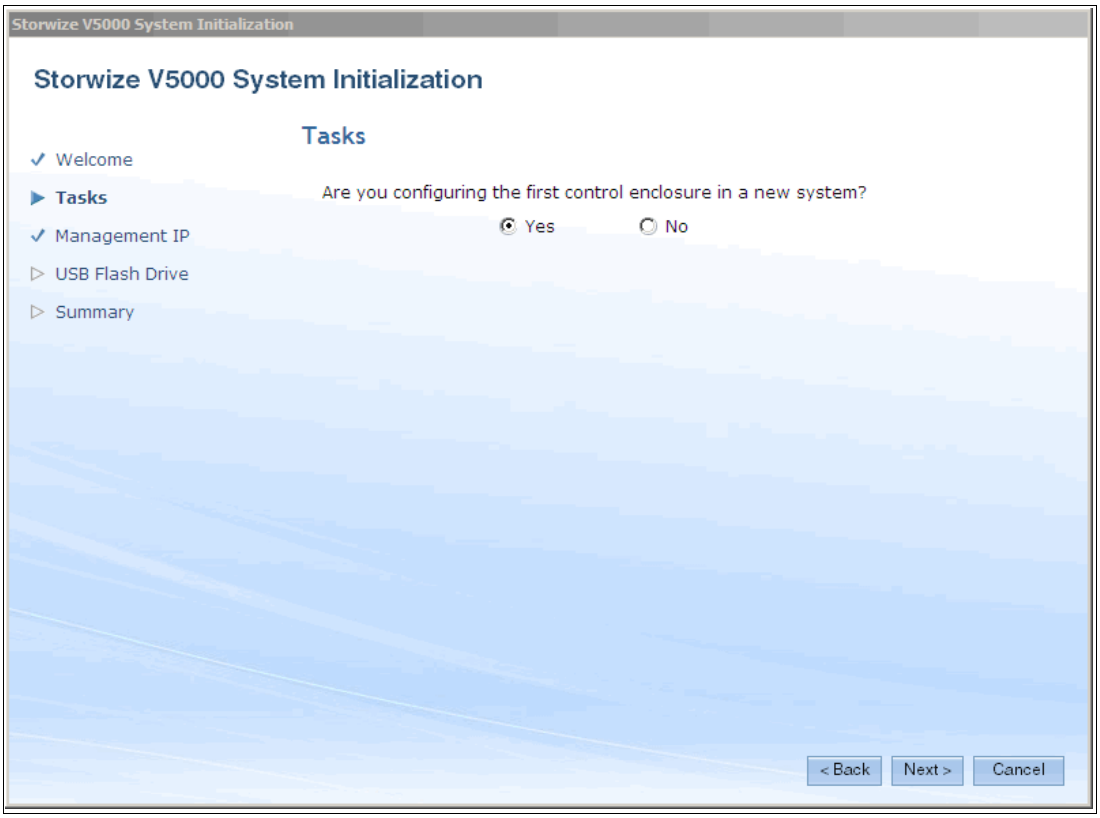

Figure 2-10 System Initialization: Create a system

There are other options available through the Tasks section. However, these options generally are only required after initial configuration. The options are shown in Figure 2-11 on page 44 and are accessed by selecting **No** to the initial question to configure a new system. A second question asks if you want to view instructions on how to expand a system with a new control enclosure appears. Selecting **No** to this question gives the option to reset the superuser password or set the service IP of a node canister. Selecting **Yes** (as shown in Figure 2-10) progresses through the initial configuration of the IBM Storwize V5000.

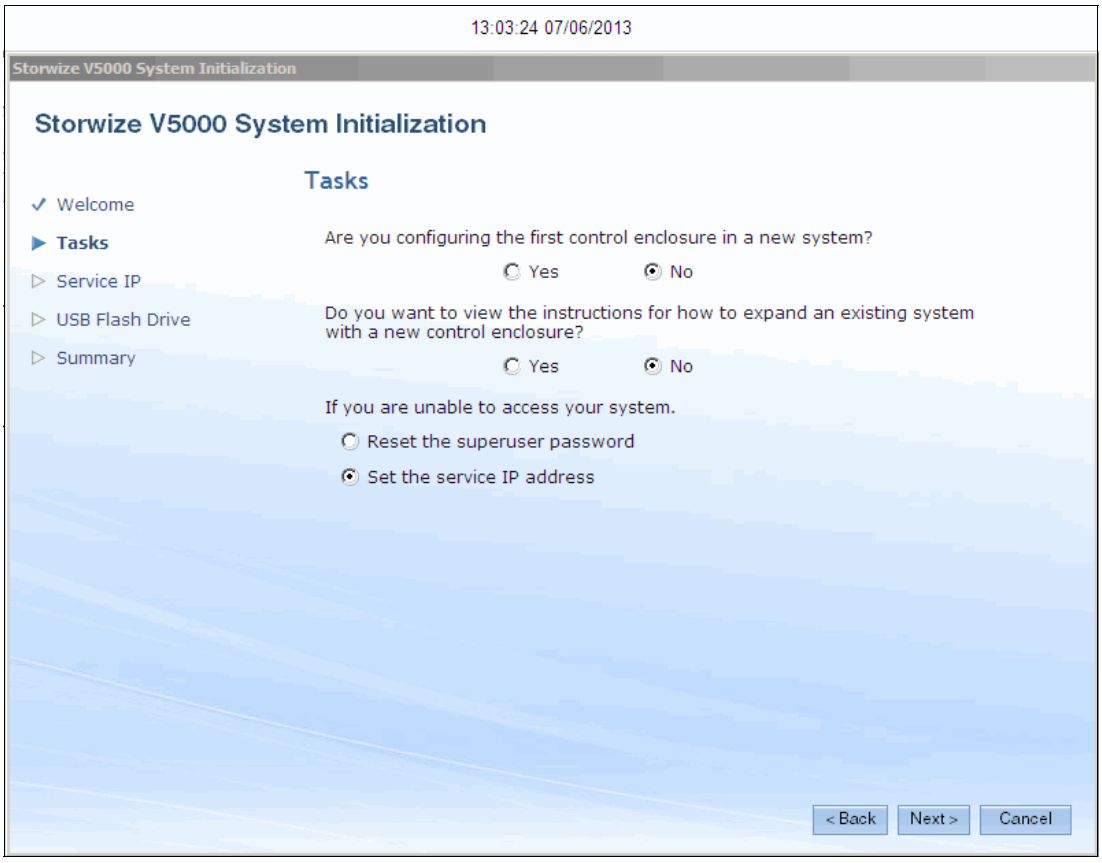

Figure 2-11 Inittool task options

2. Set the Management IP address, as shown in Figure 2-12.

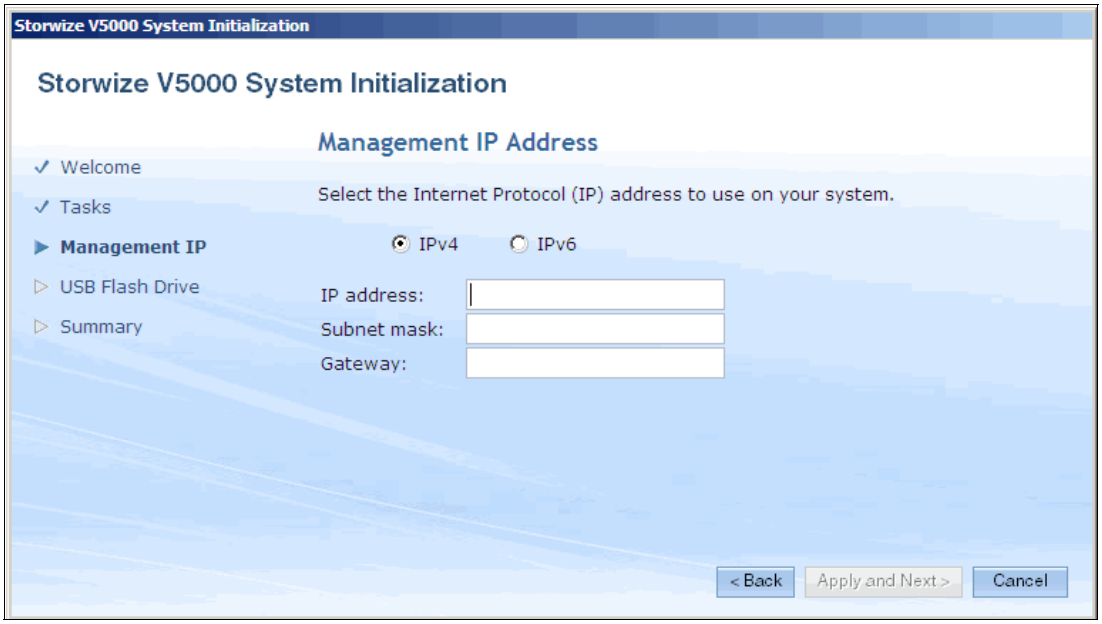

Figure 2-12 System Initialization: Management IP

3. Click **Apply and Next** to display the IBM Storwize V5000 power up instructions, as shown in Figure 2-13.

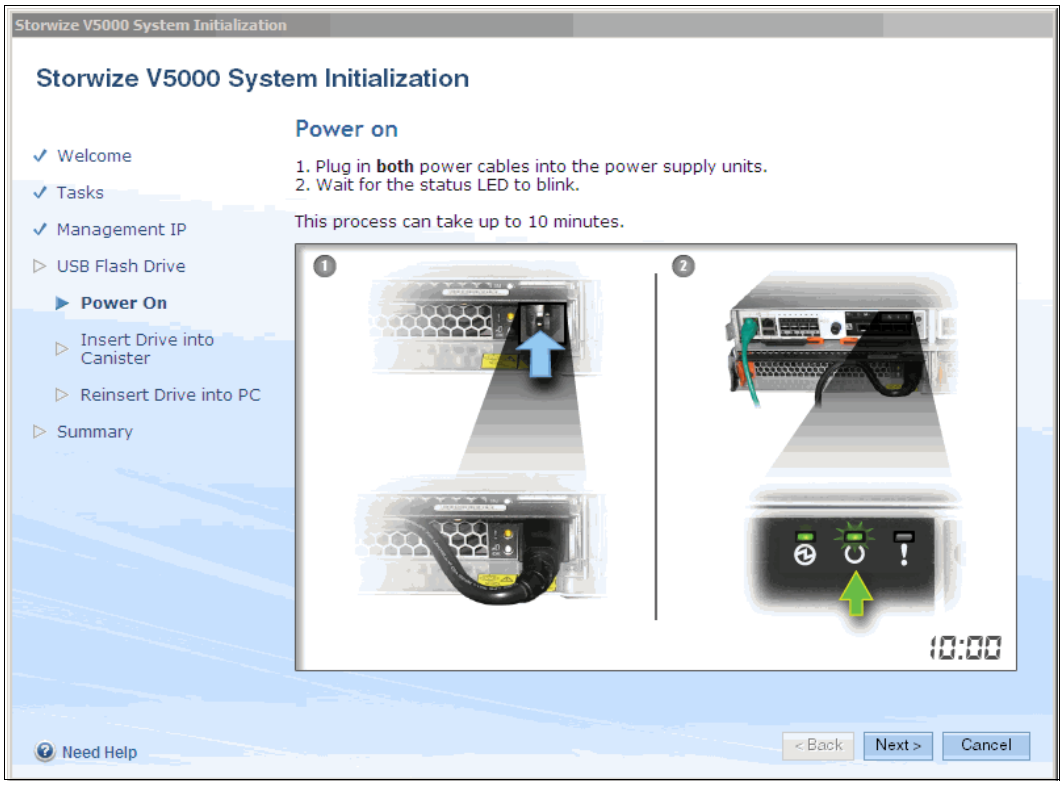

Figure 2-13 Initialization application: V5000 Power up

Any expansion enclosures that are part of the system should be powered up and allowed to come ready before the control enclosure. Follow the instructions to power up the IBM Storwize V5000 and wait for the status LED to flash. Then, insert the USB stick in one of the USB ports on the left side node canister. This node becomes the control node and the other node is the partner node. The fault LED begins to flash. When it stops, return the USB stick to the Windows PC.

Clustered system creation: While the clustered system is created, the amber fault LED on the node canister flashes. When this LED stops flashing, remove the USB key from IBM Storwize V5000 and insert it in your system to check the results.

The wizard then attempts to verify connectivity to the IBM Storwize V5000, as shown in Figure 2-14.

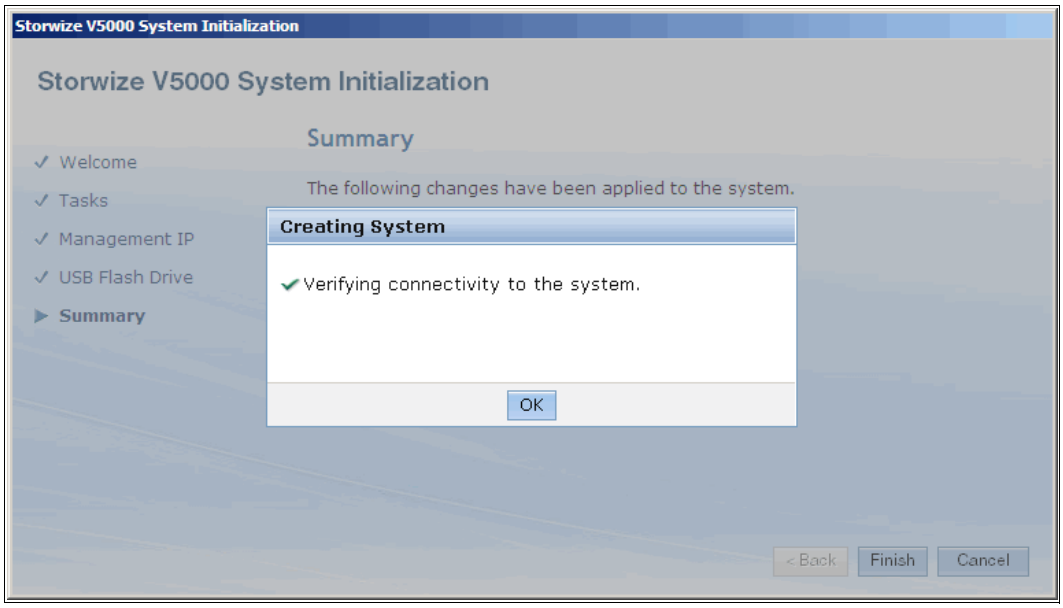

Figure 2-14 Verify system connectivity

If successful, a summary page is displayed that shows the settings that are applied to the IBM Storwize V5000, as shown in Figure 2-15.

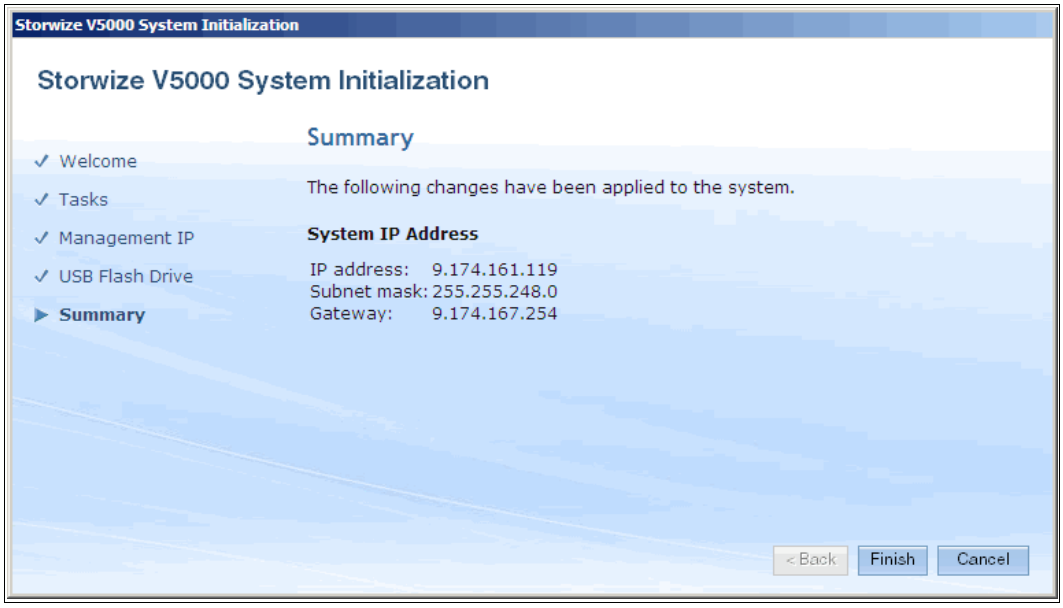

Figure 2-15 Initialization Summary

If the connectivity to the IBM Storwize V5000 cannot be verified, the warning that is shown in Figure 2-16 is displayed.

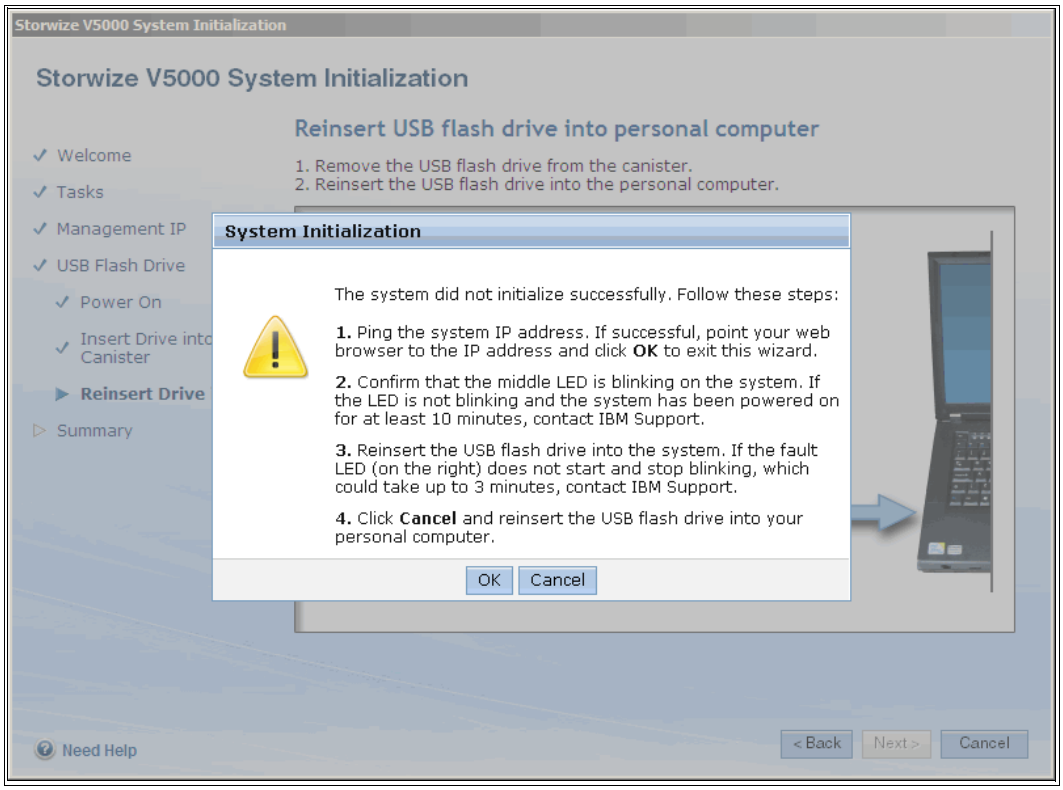

Figure 2-16 Initialization Failure

Follow the on-screen instructions to resolve any issues. The wizard assumes the system that you are using can connect to the IBM Storwize V5000 through the network. If it cannot connect, you must follow step 1 from a machine that does have network access to the IBM Storwize V5000. After the initialization process completes successfully, click **Finish**.

The initial setup is now complete. If you have a network connection to the Storwize system, the wizard redirects you to the system Management GUI, as shown in Figure 2-17.

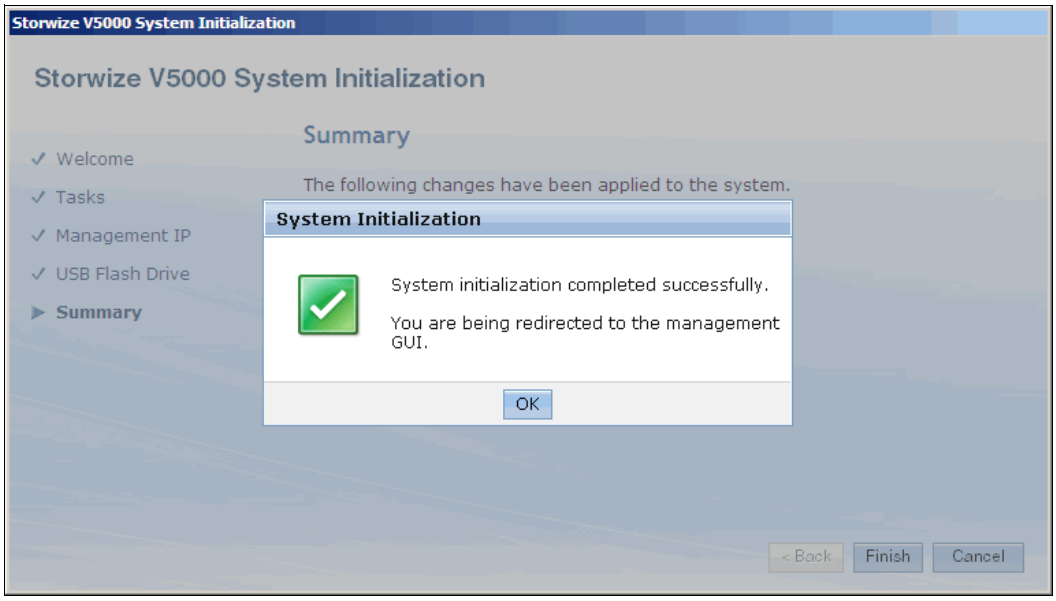

Figure 2-17 System Initialization complete

We describe system initial configuration by using the GUI in 2.10, “Initial configuration” on page 49.

2.10 Initial configuration

This section describes how to complete the initial configuration, including the following tasks:

- Setting name, date, and time
- Initial storage configuration by using the setup wizard

If you just completed the initial setup, that wizard automatically redirects to the IBM Storwize V5000 GUI. Otherwise, complete the following steps to complete the initial configuration process:

1. Start the configuration wizard by using a web browser on a workstation and point it to the system management IP address that was defined in Figure 2-12 on page 44. Enter the default superuser password <passw0rd> (where 0 = zero), as shown in Figure 2-18.

Figure 2-18 Setup wizard: Login

2. After you are logged in, a welcome window opens, as shown in Figure 2-19 on page 50.

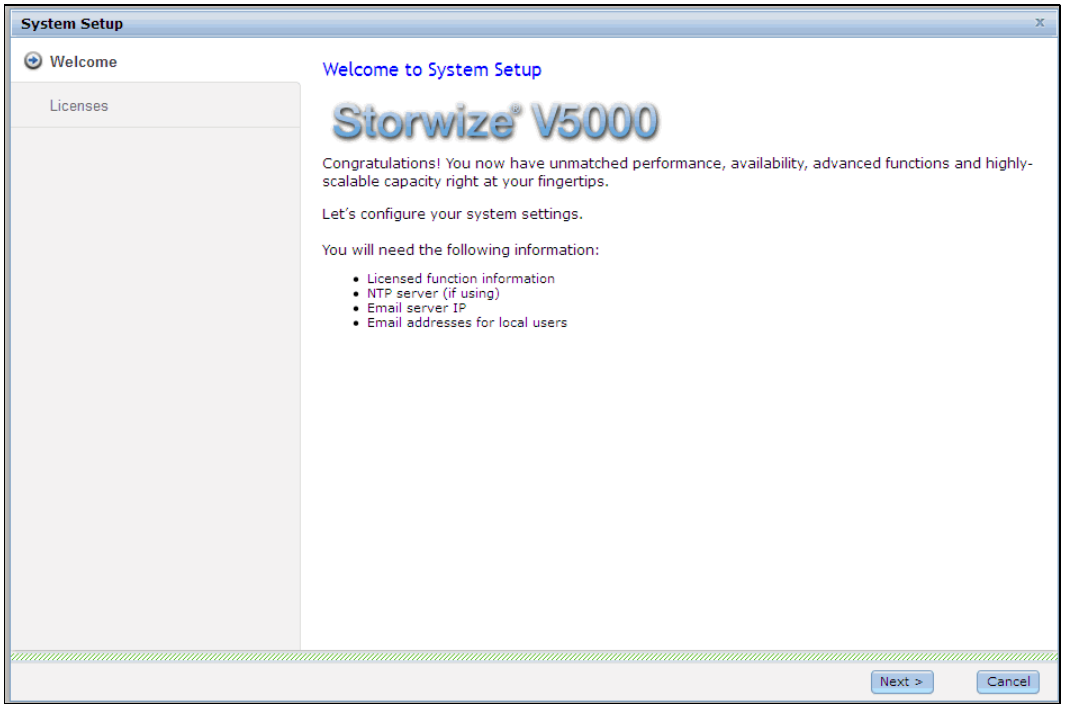

Figure 2-19 Welcome window

Click **Next** to start the configuration wizard.

3. Set up the system name, as shown in Figure 2-20.

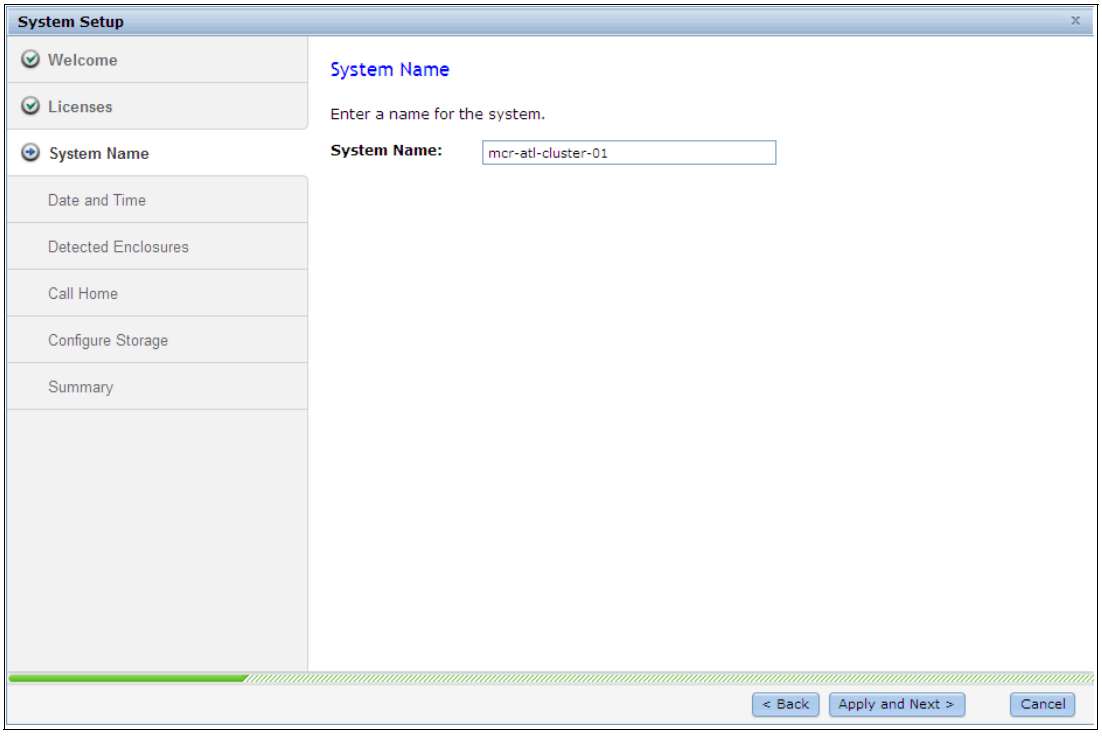

Figure 2-20 Setup wizard: Insert system name

There are two options for configuring the date and time, as shown in Figure 2-21.

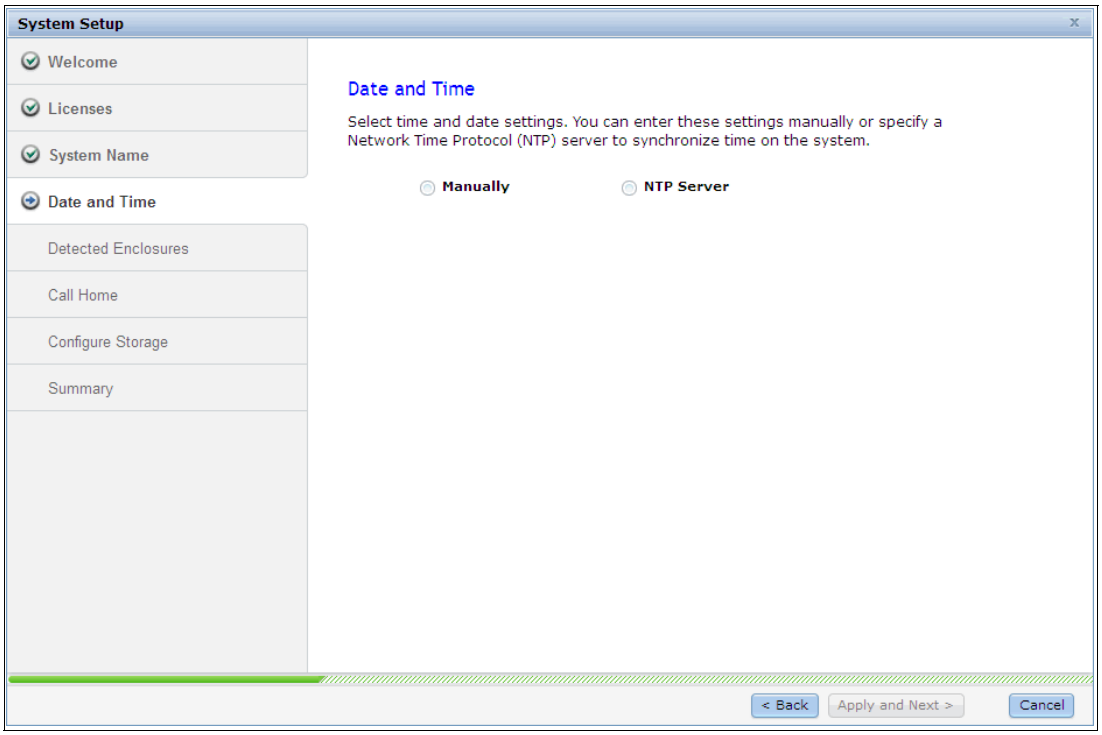

Figure 2-21 Setup wizard: Date and time

Select the required method and enter the date and time manually or specify a network address for an NTP server. After this is done, the **Apply and Next** option becomes active. Click this option to continue.

4. The configuration wizard continues with the hardware configuration. Verify the hardware, as shown in Figure 2-22 on page 52.

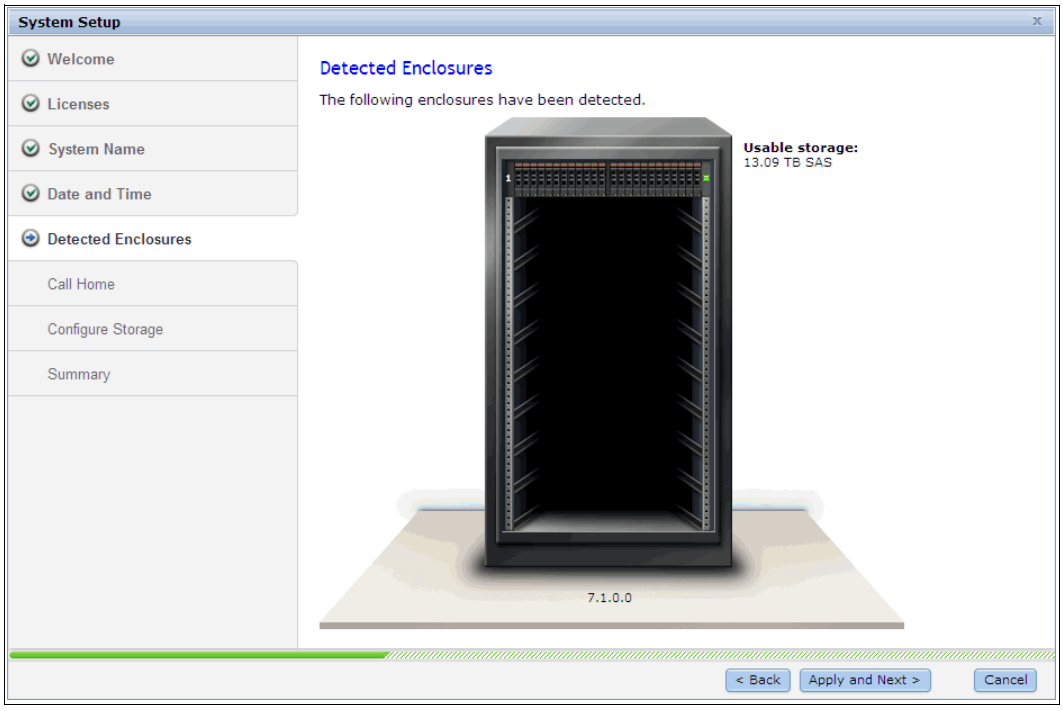

Figure 2-22 Setup wizard: Verify the detected hardware

Click **Apply and Next**.

5. The next window in the configuration process is setting up Call Home, as shown in Figure 2-23.

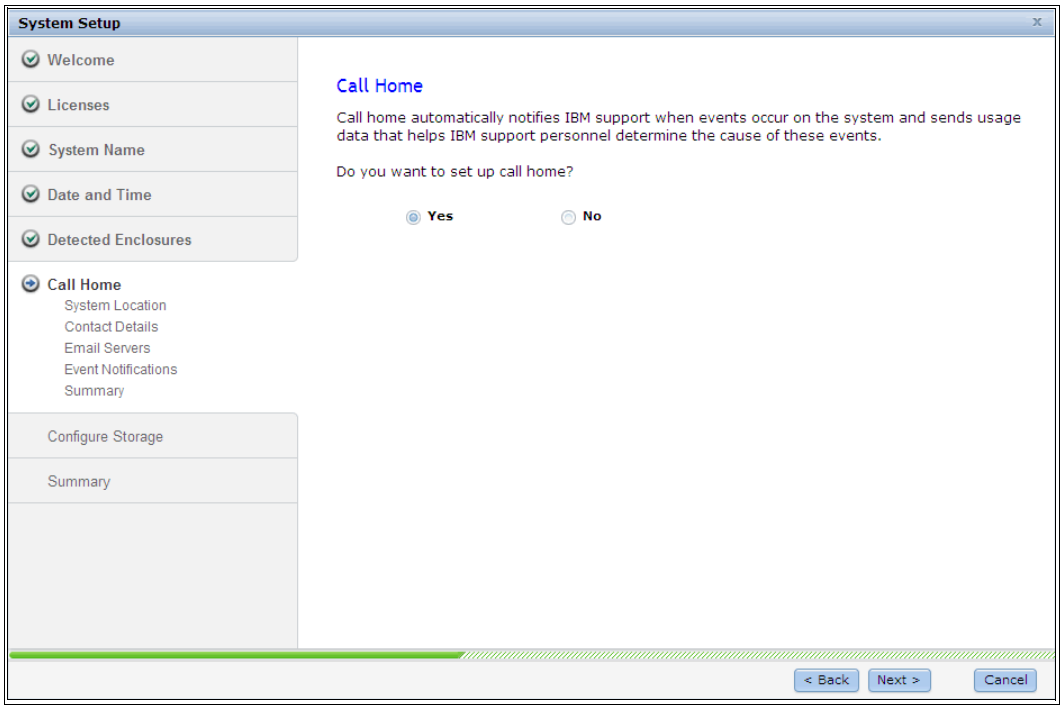

Figure 2-23 Call Home setup

It is possible to configure your system to send email reports to IBM if an issue that requires hardware replacement is detected. This function is called *Call Home*. When this email is received, IBM automatically opens a problem report and contacts you to verify whether replacements parts are required.

Call Home: When Call Home is configured, the IBM Storwize V5000 automatically creates a Support Contact with one of the following email addresses, depending on country or region of installation:

- US, Canada, Latin America, and Caribbean Islands: `callhome1@de.ibm.com`
- All other countries or regions: `callhome0@de.ibm.com`

IBM Storwize V5000 can use Simple Network Management Protocol (SNMP) traps, syslog messages, and Call Home email to notify you and the IBM Support Center when significant events are detected. Any combination of these notification methods can be used simultaneously.

To set up Call Home, you need the location details of the IBM Storwize V5000, Storage Administrators details, and at least one valid SMTP server IP address. If you do not want to configure Call Home now, it can be done later by using the GUI option by clicking **Settings** → **Event Notification** (for more information, see 2.10.2, “Configuring Call Home, email alert, and inventory” on page 69). If your system is under warranty or you have a hardware maintenance agreement to enable pro-active support of the IBM Storwize V5000, it is recommended that Call Home is configured. Selecting **Yes** and clicking **Next** moves to the window that is used to enter the location details, as shown in Figure 2-24.

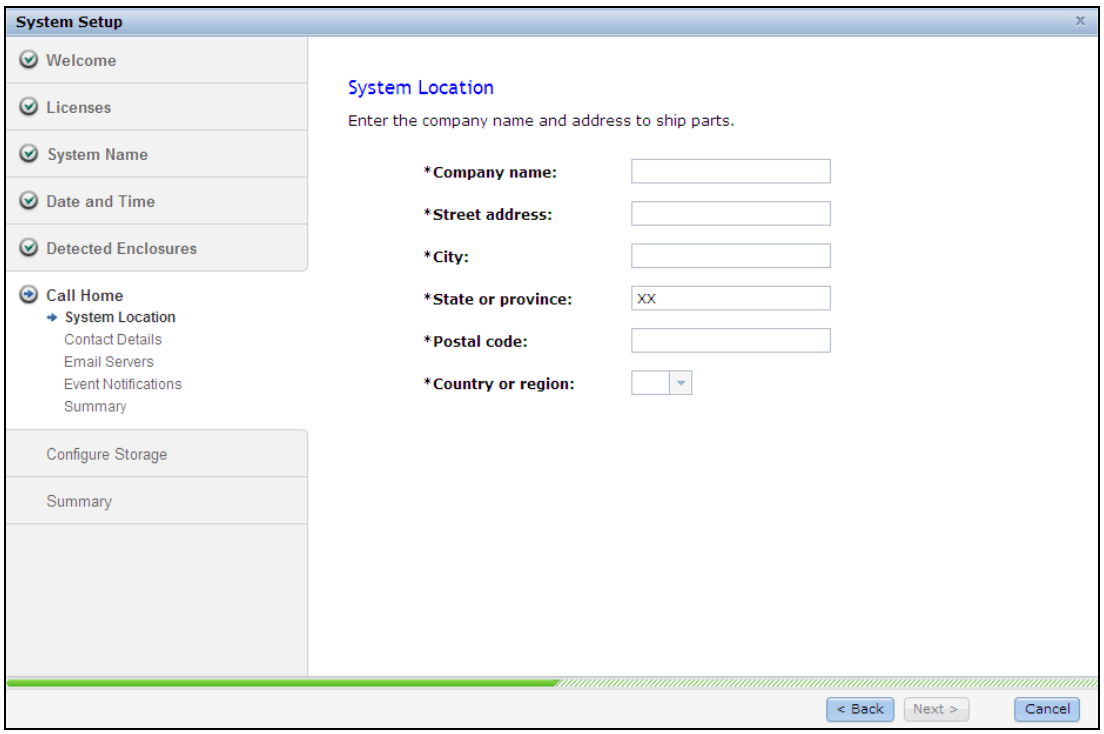

Figure 2-24 Location details

These details appear on the Call Home data to enable IBM Support to correctly identify where the IBM Storwize V5000 is located.

Important: Unless the IBM Storwize V5000 is in the US, the state or province field should be completed by using XX. Follow the help for correct entries for locations inside the US.

You can enter the contact details of the main storage administrator in the nest window, as shown in Figure 2-25. You can choose to enter the details for a 24-hour operations desk. These details also are sent with any Call Home. This information allows IBM Support to contact the correct people to quickly progress any issues.

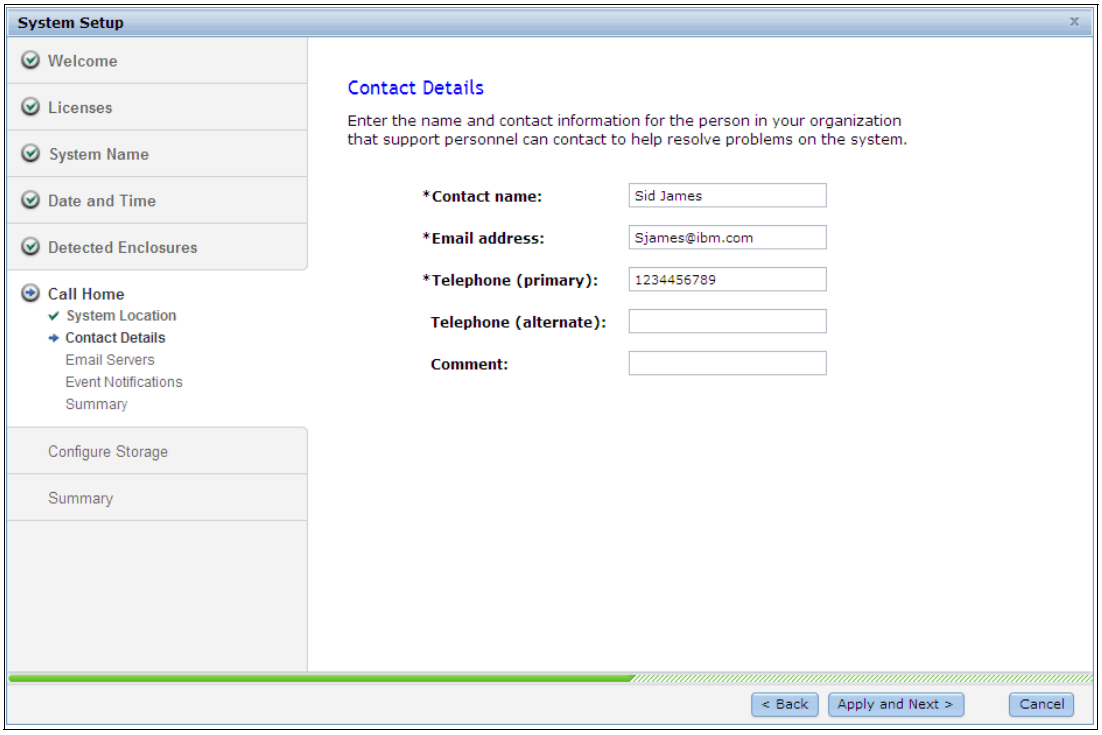

Figure 2-25 Contact details

The next window is for email server details. To enter more than one email server, click the green + icon, as shown in Figure 2-26 on page 55.

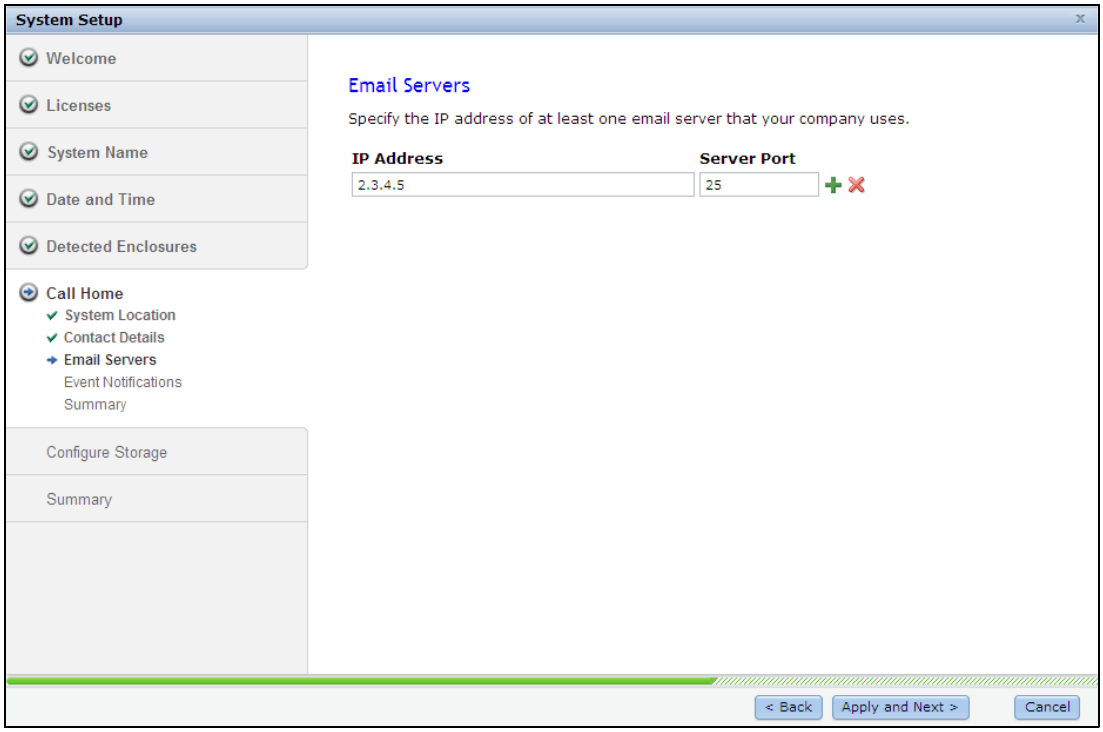

Figure 2-26 Email server details

The IBM Storwize V5000 also can configure local email alerts. These can be sent to a storage administrator or an email alias for a team of administrators or operators. To add more than one recipient, click the green **+** icon, as shown in Figure 2-27.

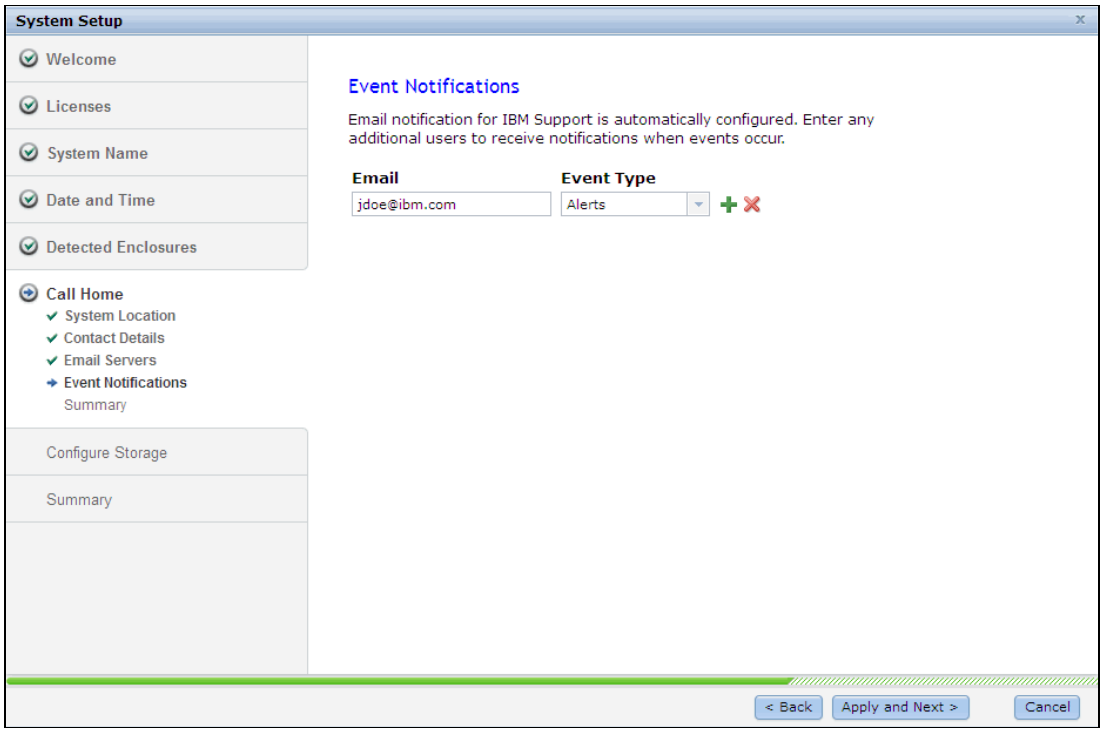

Figure 2-27 Event notification

Clicking **Apply** and **Next** displays the summary window for the call home options, as shown in Figure 2-28.

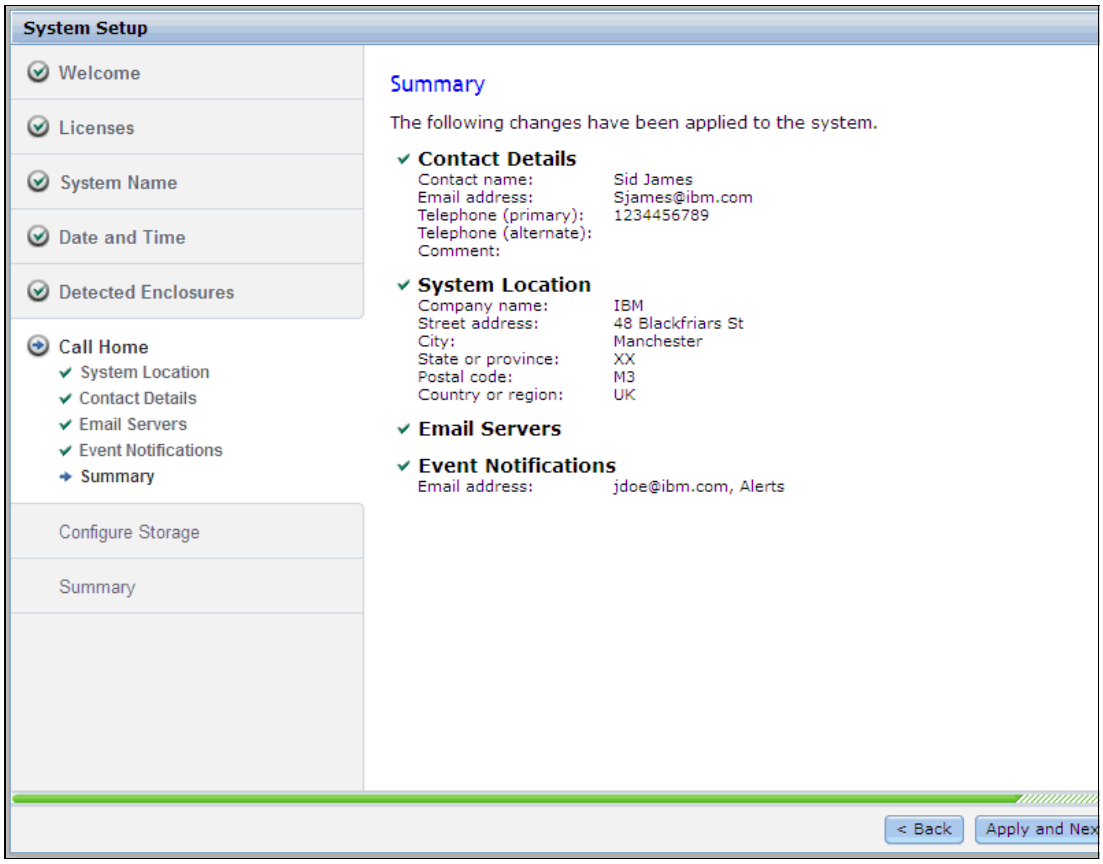

Figure 2-28 Call Home summary

Click **Apply and Next**.

6. The initial configuration wizard moves on to the Configure Storage option next. This option takes all the disks in the IBM Storwize V5000 and automatically configures them into optimal RAID arrays for use as MDisks. If you do not want to automatically configure disks now, select **No** and you exit the wizard to the IBM Storwize V5000 GUI, as shown in Figure 2-29 on page 57.

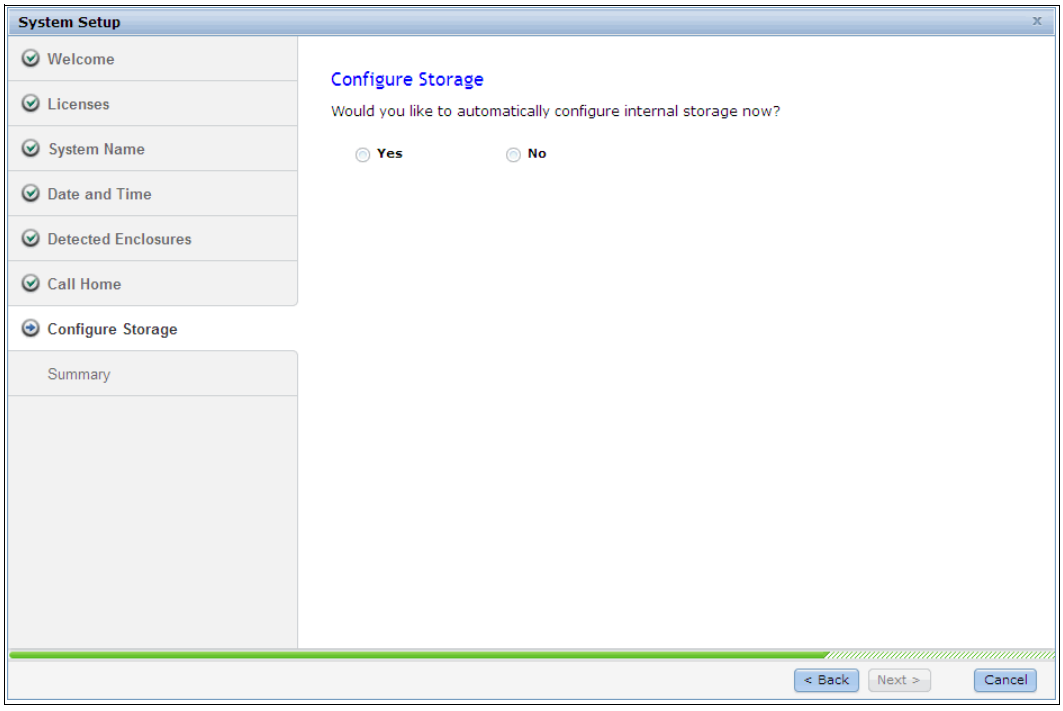

Figure 2-29 Configure Storage option

Selecting **Yes** and clicking **Next** moves to the summary window that shows the RAID configuration that the IBM Storwize V5000 implement, as shown in Figure 2-30.

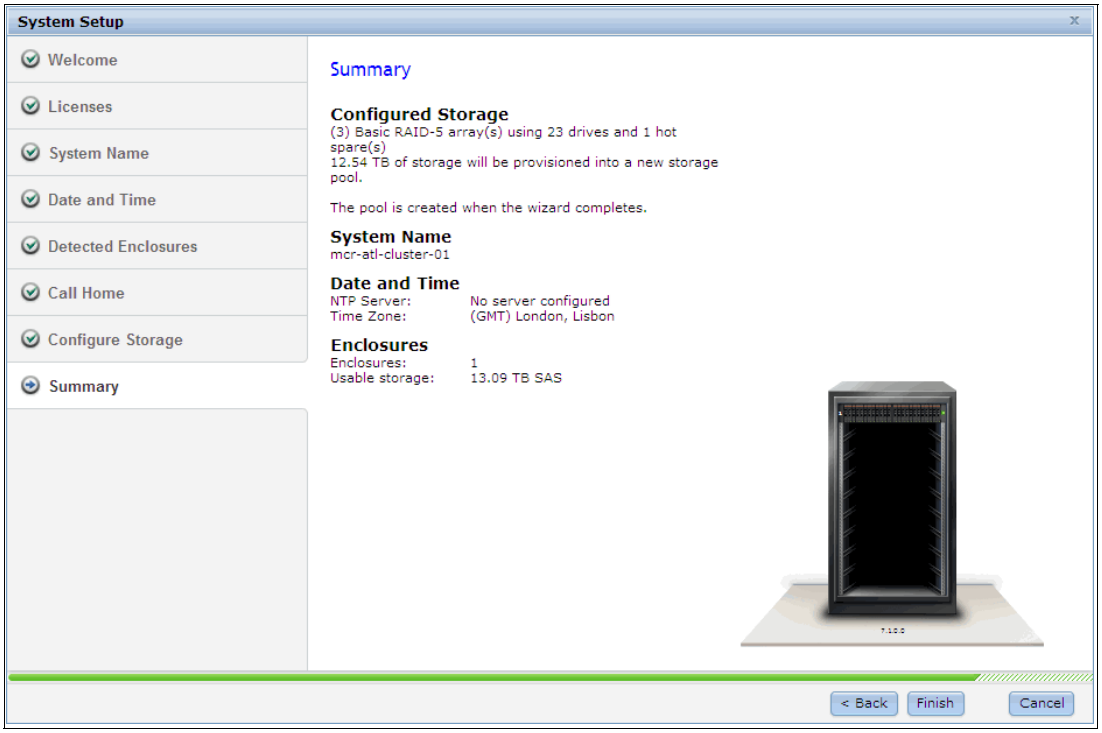

Figure 2-30 Storage Configuration Summary

The storage pool is created when you click, **Finish** as shown in Figure 2-31.

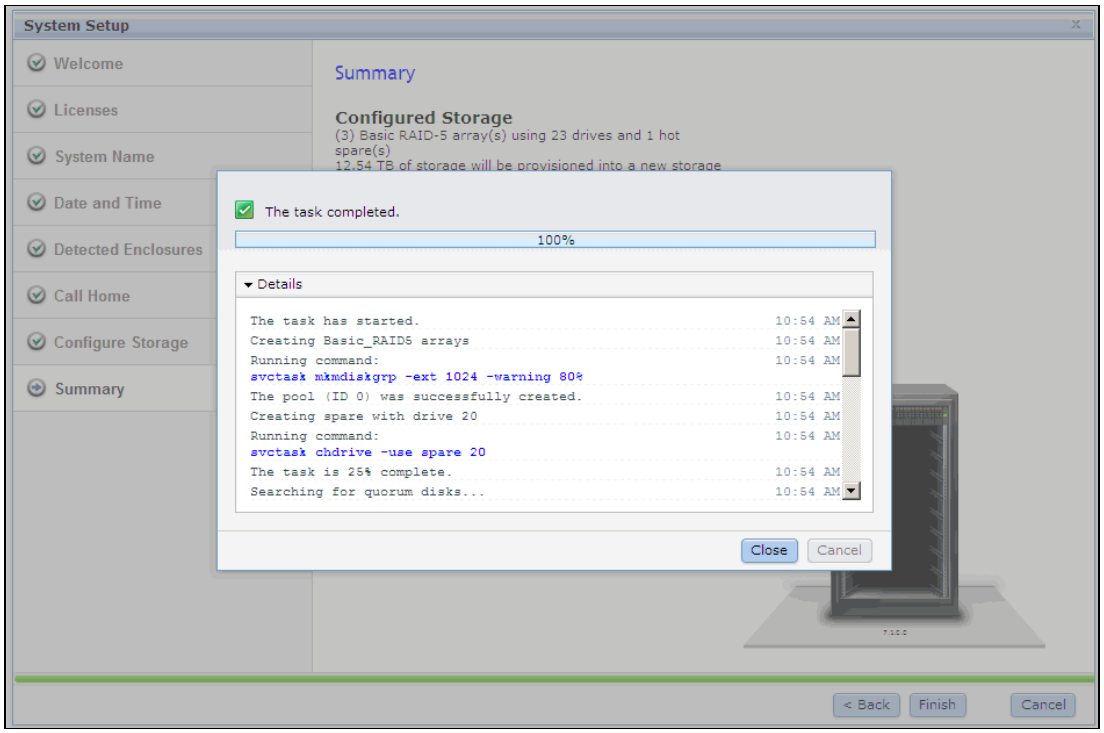

Figure 2-31 Storage array creation

Closing the task box completes the Initial configuration wizard and automatically directs you to the Create Hosts task option on the GUI, as shown in Figure 2-32.

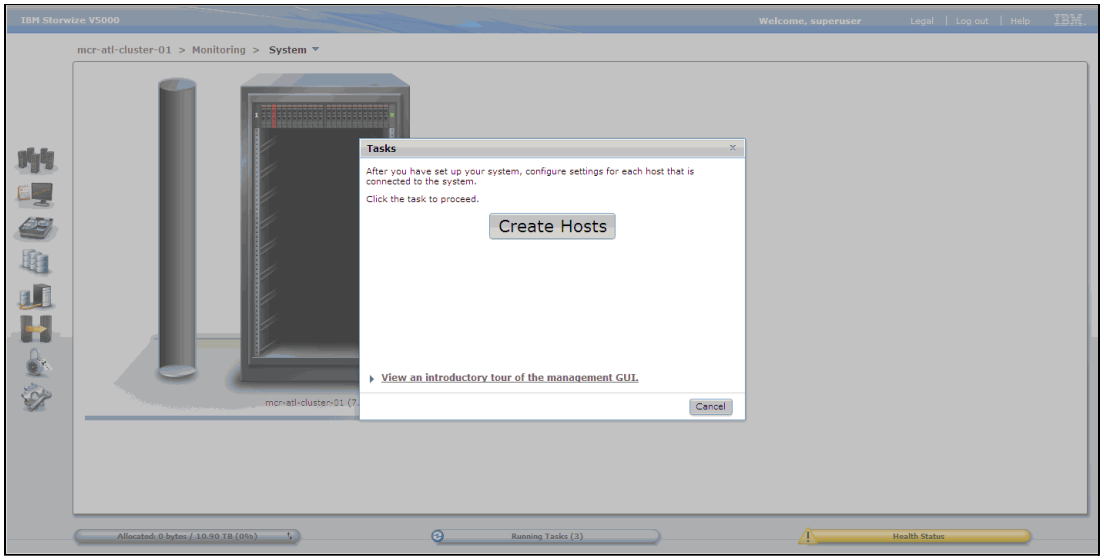

Figure 2-32 Create hosts

If you choose to create hosts at this stage, see Chapter 4, “Host configuration” on page 153 for details.

Selecting **Cancel** exits to the IBM Storwize V5000 GUI. There is also a hot link to the e-Learning tours that are available through the GUI.

2.10.1 Adding Enclosures after initial configuration

When the initial install of the IBM Storwize V5000 is complete, all expansion enclosures and control enclosures that were purchased at that time should be installed as part of the initial configuration. This process enables the system to make the best use of the enclosures and drives that are available.

Adding a control enclosure

If you are expanding the IBM Storwize V5000 after the initial installation by adding a second I/O Group (a second control enclosure) or adding expansion enclosures, follow the physical installation procedures as described in *IBM Storwize V5000 Quick Installation Guide Version 6.4.1*, GC27-4219. For more information about zoning the node canisters, see 2.2, “SAN configuration planning” on page 31.

After the hardware is installed, cabled, and powered on, a second control enclosure is added.

Complete the following steps to use the management GUI to configure the new enclosure:

1. In the Monitoring tab, select **Actions** → **Add Enclosures** → **Control and Expansions**, as shown in Figure 2-33.

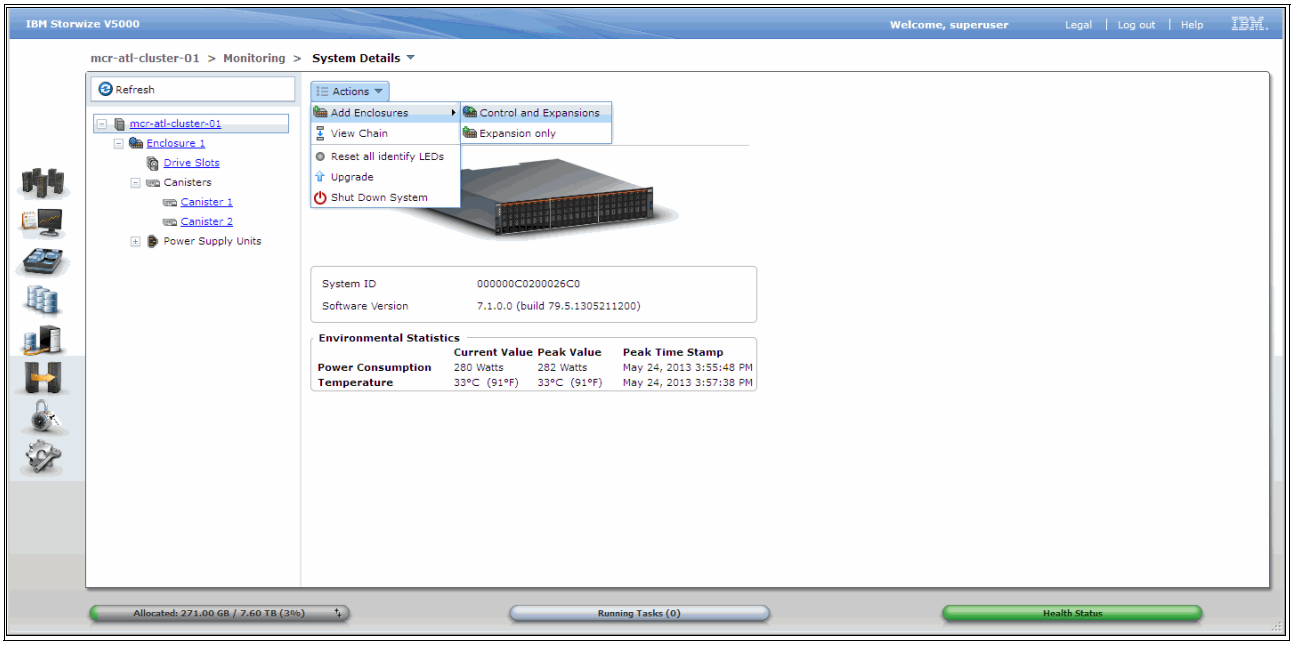

Figure 2-33 Option to add a control enclosure

2. If the control enclosure is properly configured, the new control enclosure is identified in the next window, as shown in Figure 2-34.

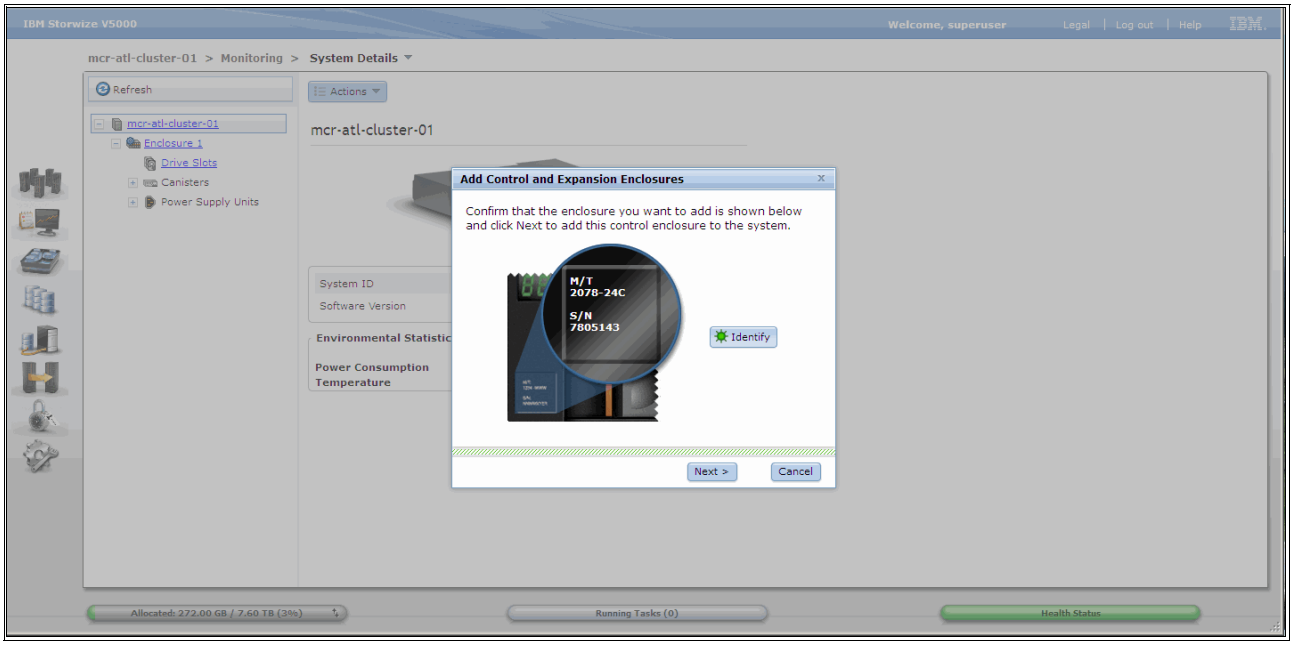

Figure 2-34 New control enclosure identification

Click the **Identify** option to turn on the identify LEDs of the new canister, if required. Otherwise, click **Next.**

3. You might receive a message that indicates the software level of the new control enclosure needs upgrading, as shown in Figure 2-35. This is normal if the new enclosure is at a lower level of code than your existing IBM Storwize V5000. Click **OK**.

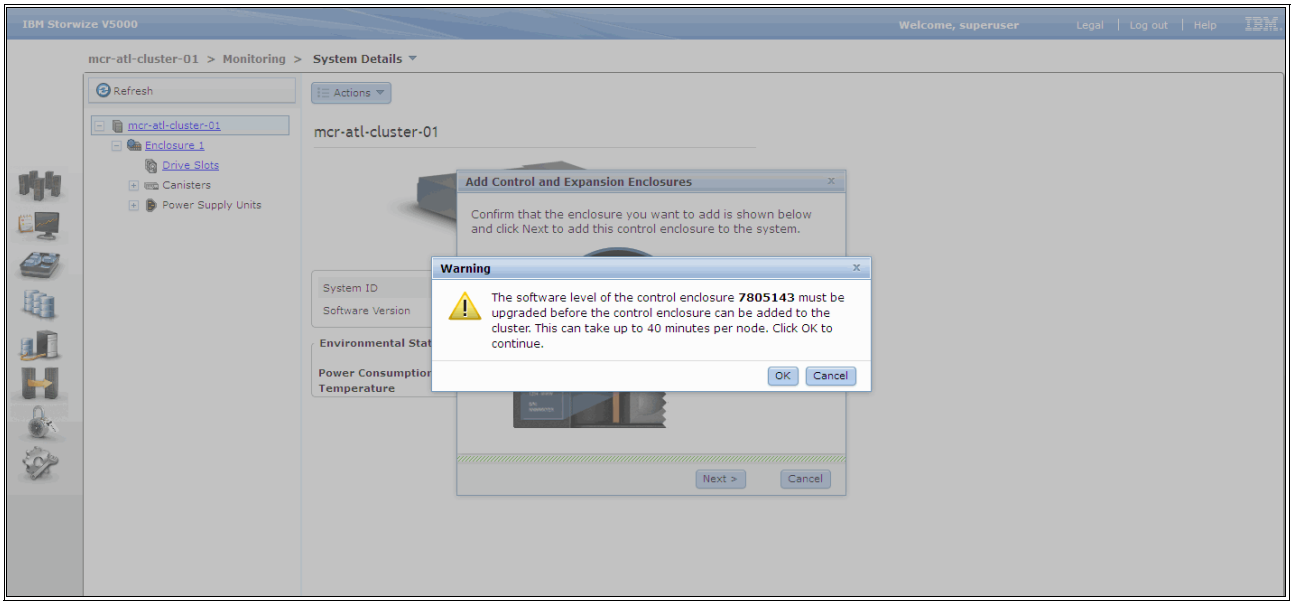

Figure 2-35 New control enclosure software upgrade warning

It can take several minutes for the software upgrade to complete.

Important: It is recommended that you have your system at the latest level of code before any enclosure expansions are done.

After the code upgrade completes or if the new enclosure is already at the same level, the IBM Storwize V5000 adds the new enclosure to the configuration, as shown in Figure 2-36.

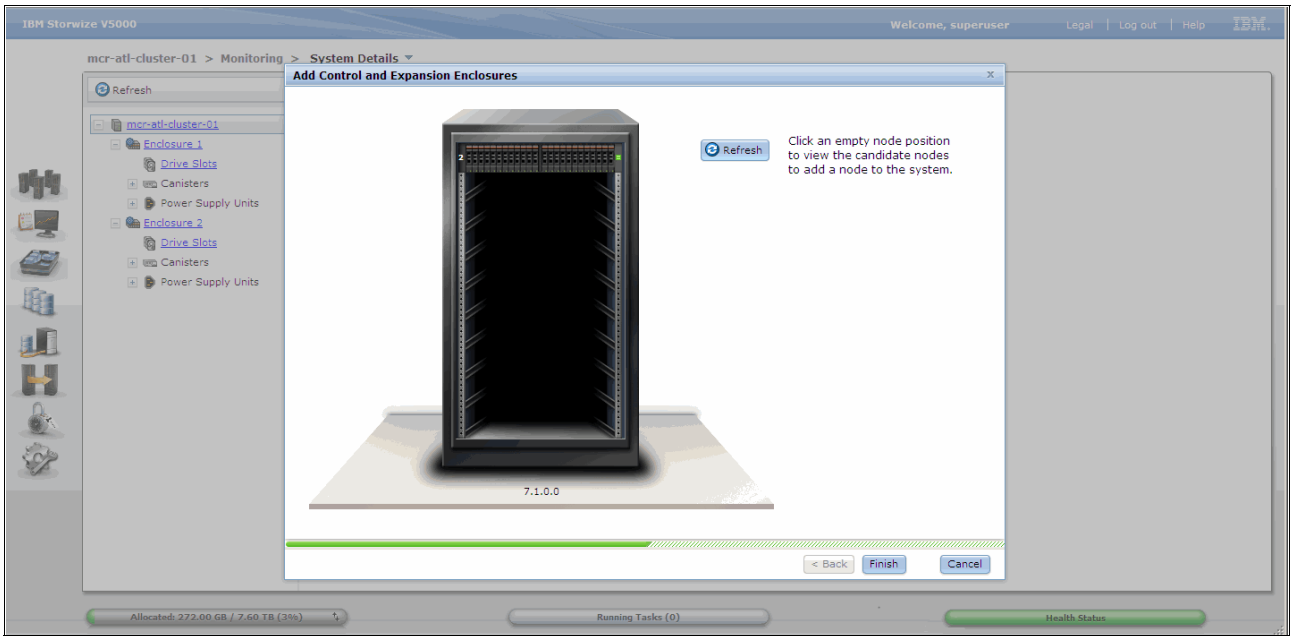

Figure 2-36 Add enclosure complete

Because the new control enclosure forms an I/O Group of its own, it appears as a single enclosure in the rack. The original I/O Group is not shown even though they are part of the same clustered system. The wording in the window is also misleading. By clicking the enclosure that is shown, you see the candidate nodes that are to be added to the system. The empty spaces do not actually do anything. If no new hardware is shown, check your cabling and zoning and use the Refresh option. Be aware the Refresh option also is disabled in subsequent windows if you use it. Therefore, if you still cannot see the new hardware after a refresh, you might have to stop the process by clicking **Cancel**, correcting any physical connectivity issues or hardware issues, and then beginning the process of adding an enclosure again.

4. To add the enclosure, select the new enclosure and click **Finish**. The task to add the enclosure completes, as shown in Figure 2-37 on page 62.

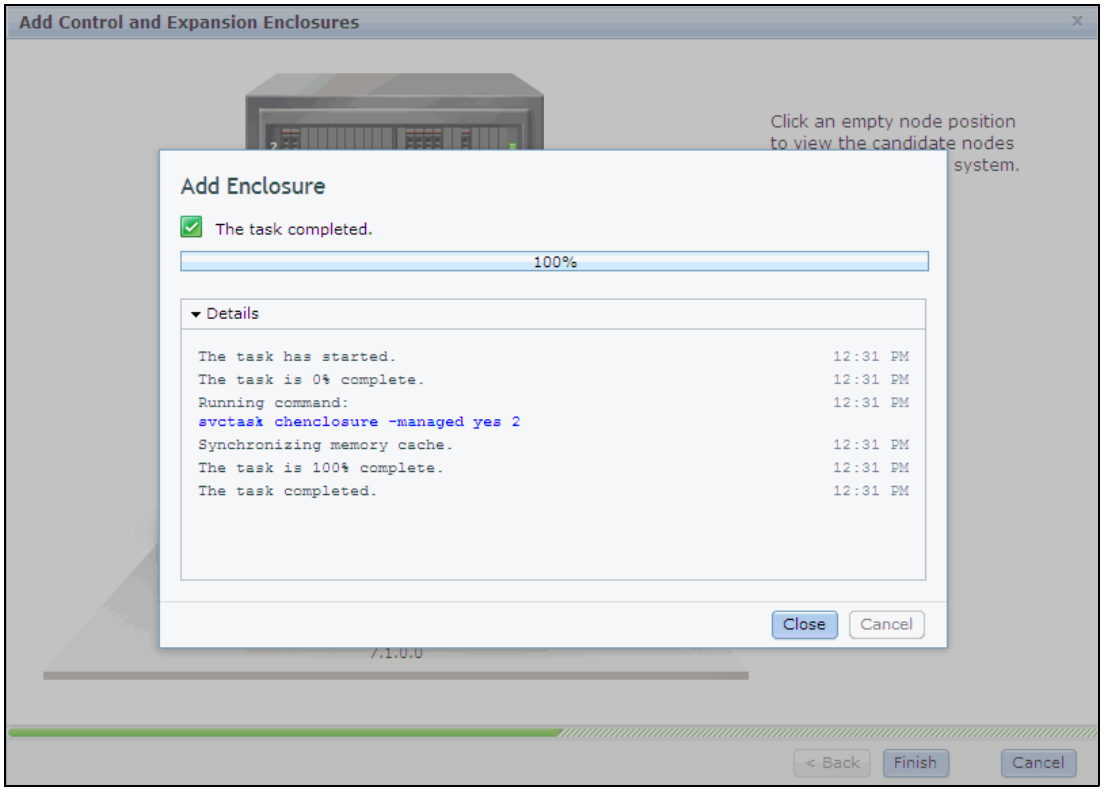

Figure 2-37 Add control enclosure task completion

5. Click **Close** to finish the wizard. You are prompted to configure the new storage, as shown in Figure 2-38.

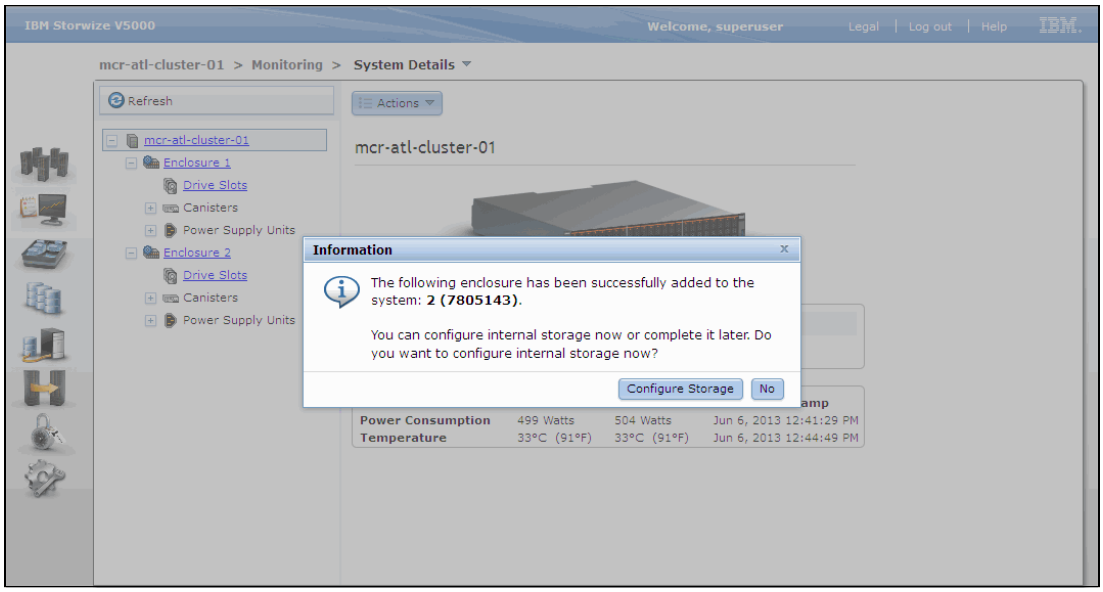

Figure 2-38 New storage configuration prompt

At this point, you can choose **Configure Storage** or **No** to quit the wizard and return to the IBM Storwize V5000 GUI.

If you choose to configure storage, a wizard starts, as shown in Figure 2-39.

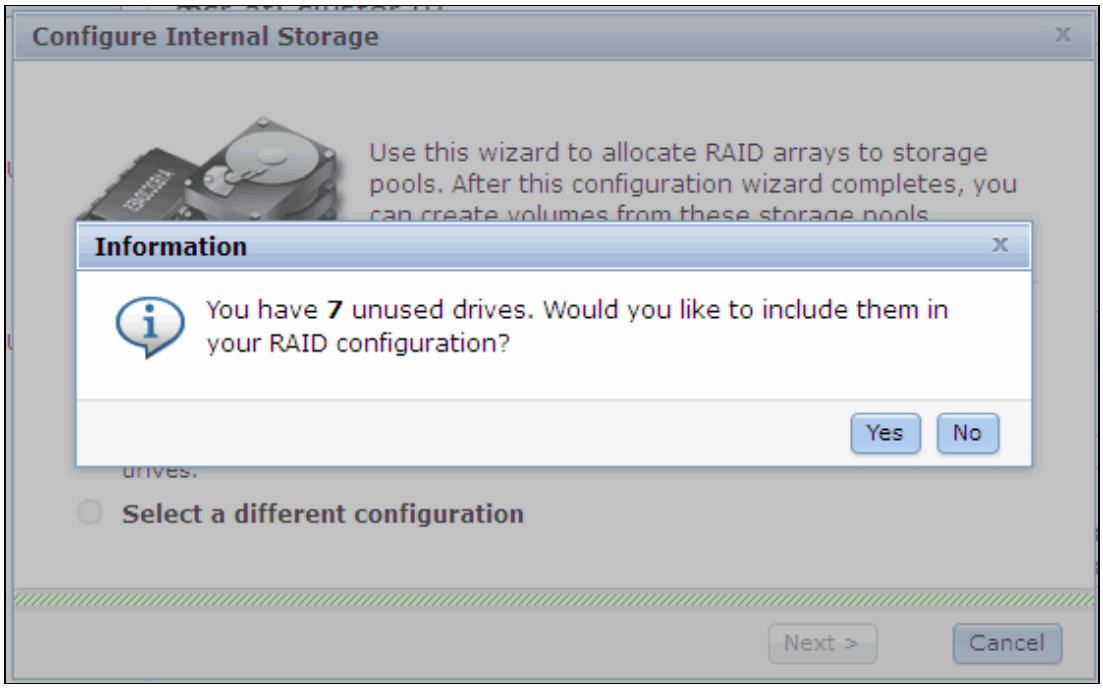

Figure 2-39 Configure new enclosure storage

6. Select **Yes** to have the IBM Storwize V5000 automatically configure the new drives as candidates. Select **No** to exit the wizard.

   The wizard prompts you to configure the new internal storage, as shown in Figure 2-40.

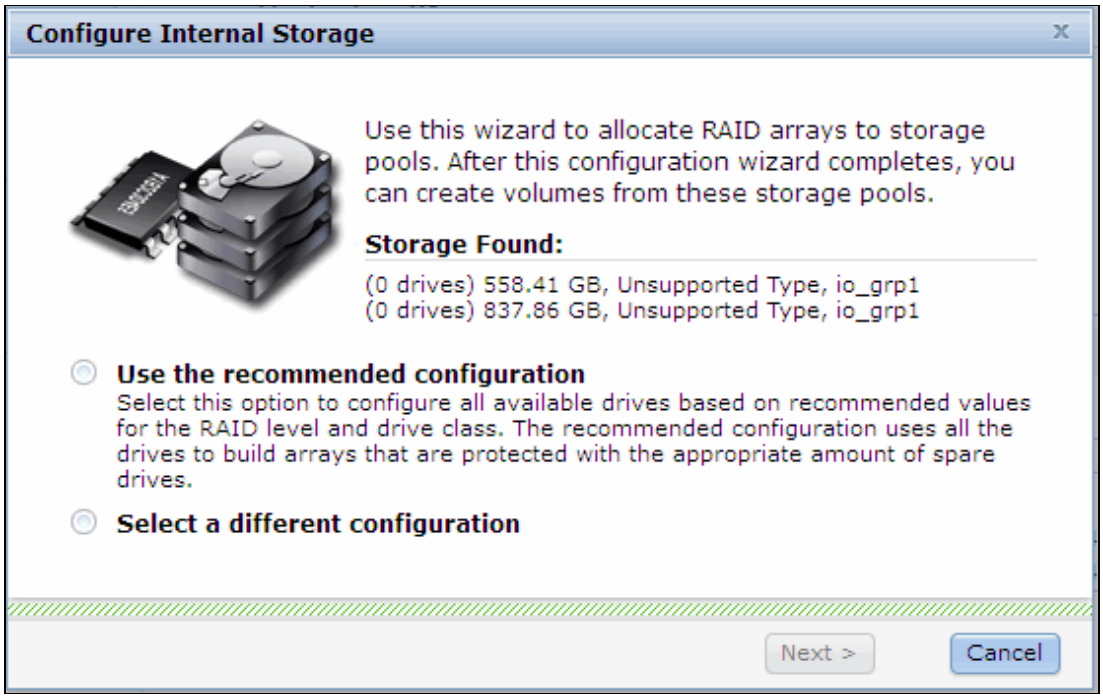

Figure 2-40 Configure new Internal Storage

7. The new enclosure is now be part of the cluster, as shown in Figure 2-41.

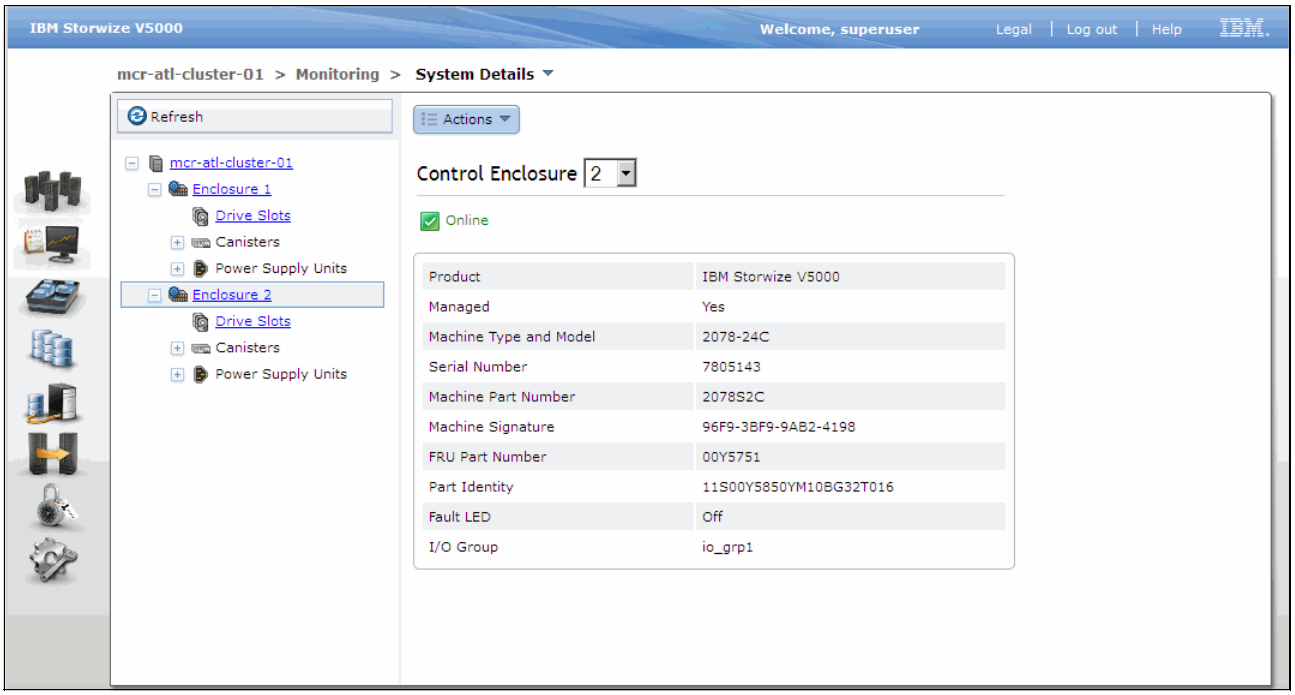

Figure 2-41 New enclosure that is shown as part of existing cluster

Adding a new expansion enclosure

Complete the following steps to add a new expansion controller:

1. In the Monitoring tab, select **Actions** → **Add Enclosures,** as shown in Figure 2-42. If you have a four-node cluster (two control enclosures), the only option that is available is to add an expansion enclosure. If you have a two-node cluster (a single control enclosure), you have the options that are shown in Figure 2-33 on page 59. In this case, select **Expansion only**.

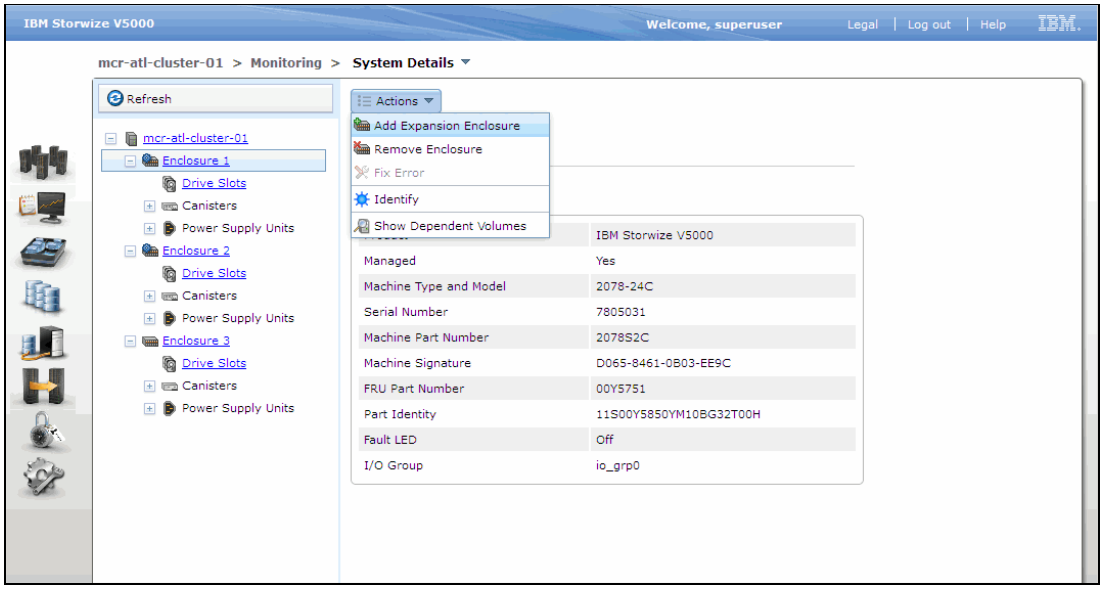

Figure 2-42 Adding an expansion enclosure

2. You are prompted to check and confirm cabling and power to the new expansion enclosure. Click **Next** to continue, as shown in Figure 2-43.

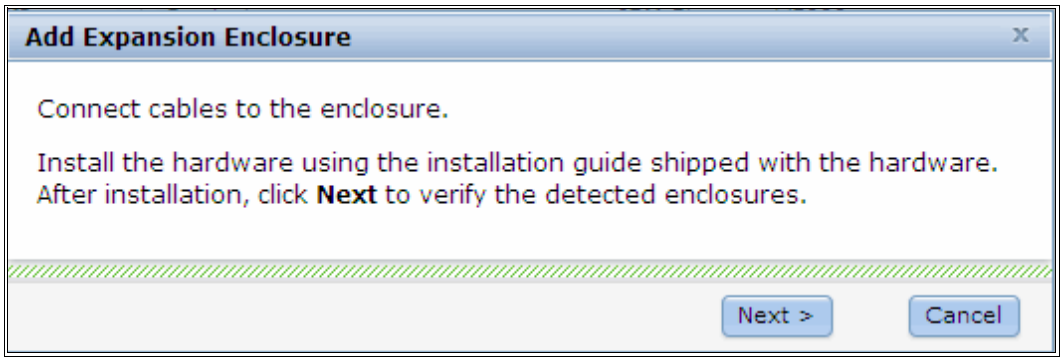

Figure 2-43 Expansion enclosure cable check

3. A task runs and completes to discover the new hardware, as shown in Figure 2-44. Click **Close** to continue.

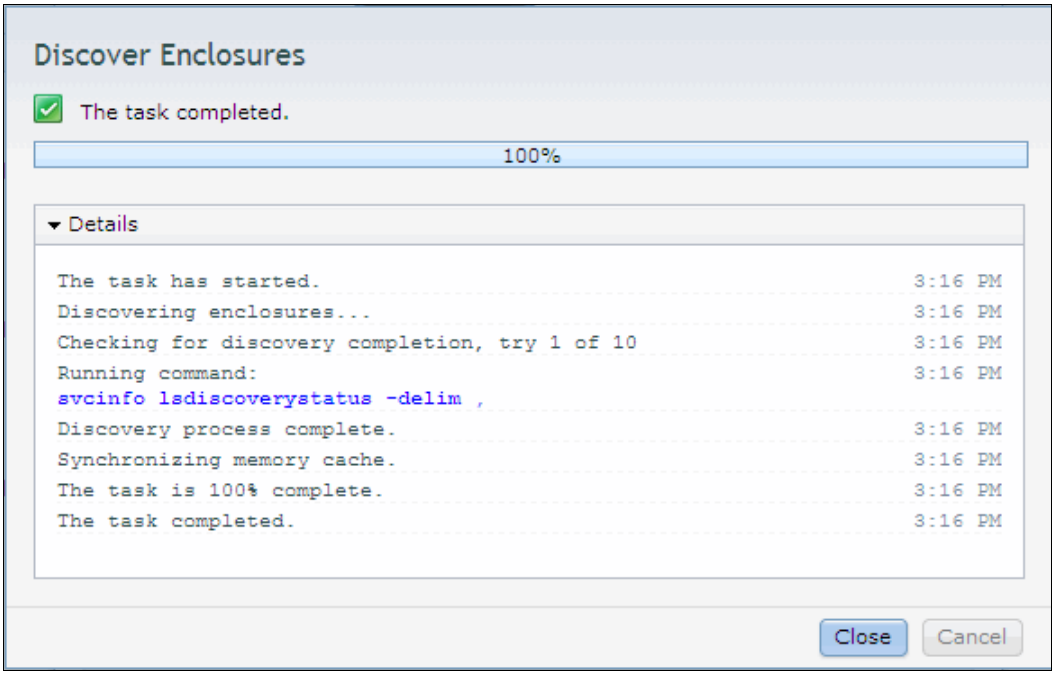

Figure 2-44 New enclosure discovery task

4. A window opens that shows the details of the new hardware to be added, as shown in Figure 2-45. There is an option to identify the new enclosure by flashing the identify light and another option to view the SAS chain that relates to the enclosure.

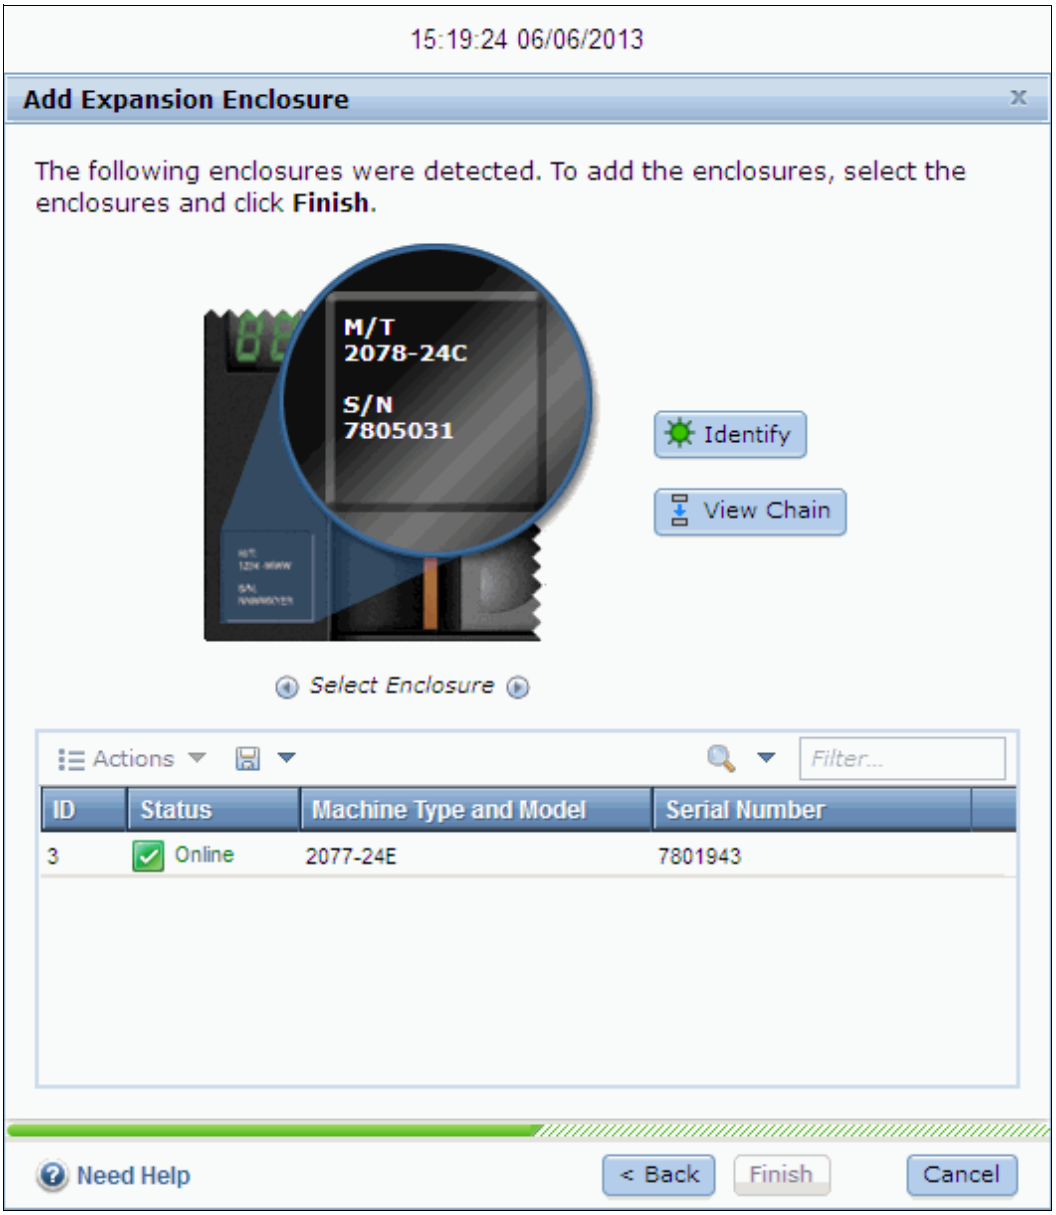

Figure 2-45 New hardware to be added

5. To add the enclosure, highlight it and click **Finish**, as shown in Figure 2-46.

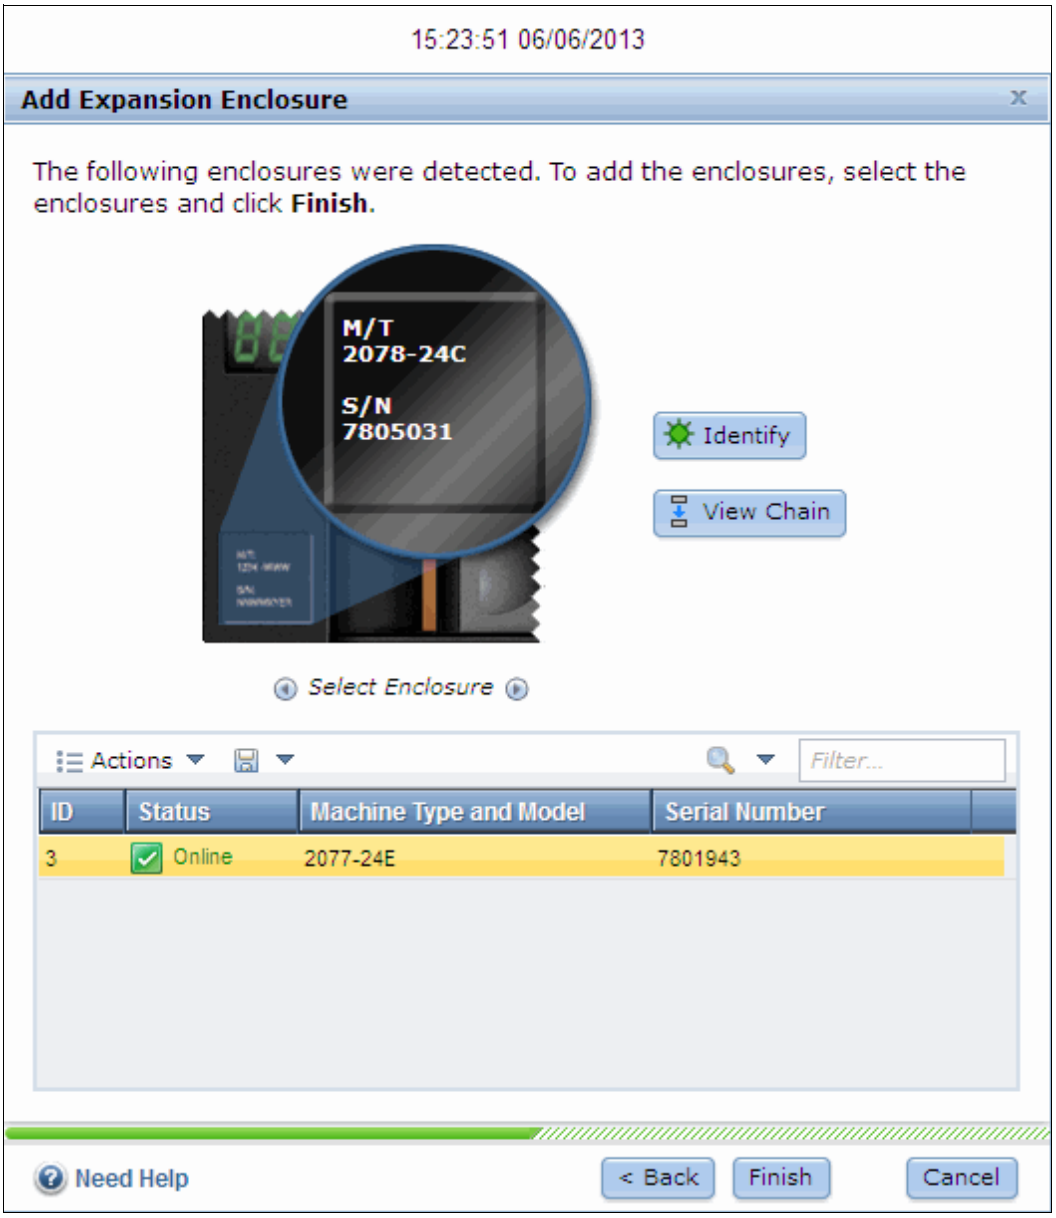

Figure 2-46 Selecting new hardware to be added

6. The task to add the new enclosure runs and completes, as shown in Figure 2-47. Click **Close**.

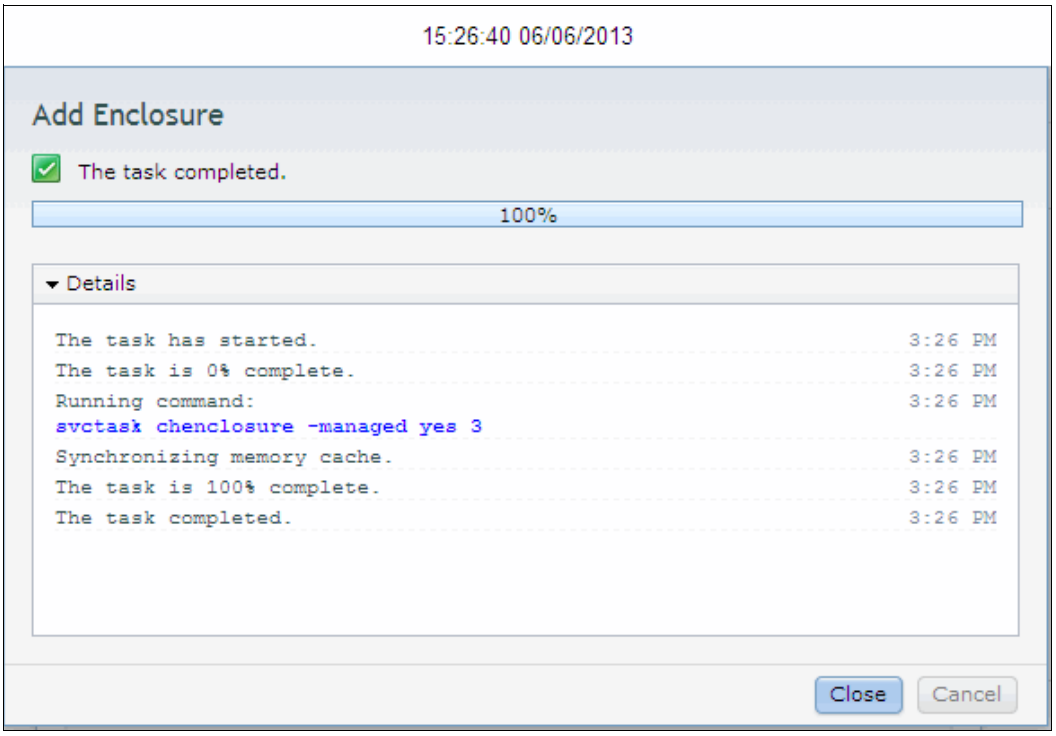

Figure 2-47 Add new enclosure task completion

7. The new expansion enclosure now is shown as part of the cluster that is attached to its control enclosure, as shown in Figure 2-48.

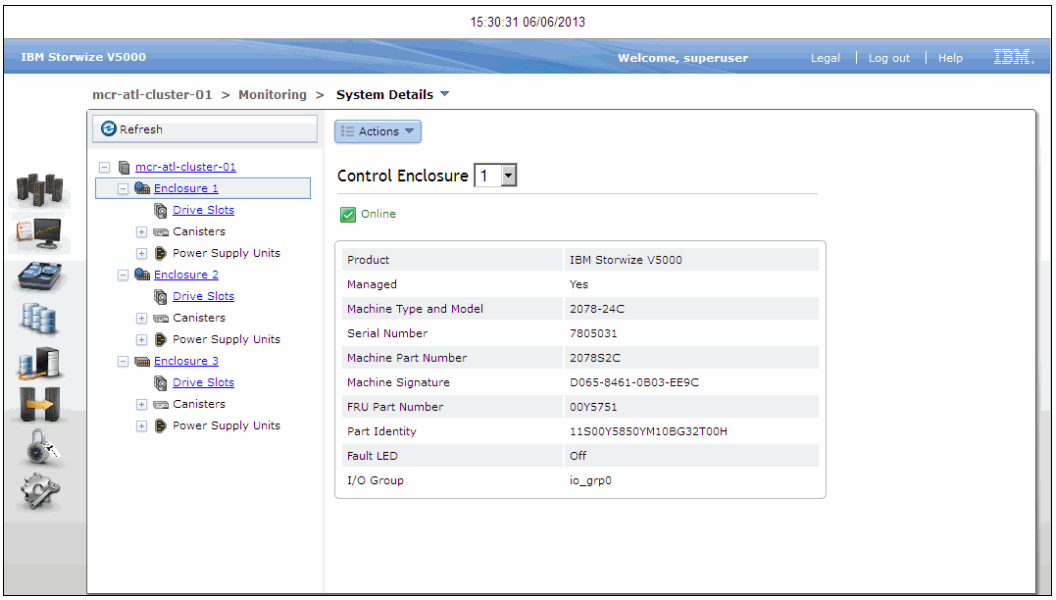

Figure 2-48 New expansion enclosure as part of the cluster

For more information about how to provision the new storage in the expansion enclosure, see Chapter 7, “Storage pools” on page 295.

2.10.2 Configuring Call Home, email alert, and inventory

If your system is under warranty or you have a hardware maintenance agreement, it is recommended to configure your system to send email reports to IBM if an issue that requires hardware investigation is detected. This feature is known as Call Home and is typically configured during the Initial Configuration of the system, as described in item 5 on page 52.

To configure the Call Home and email alert event notification in IBM Storwize V5000 after the Initial Configuration, complete the following steps:

1. Click **Settings** → **Event Notifications**, as shown in Figure 2-49.

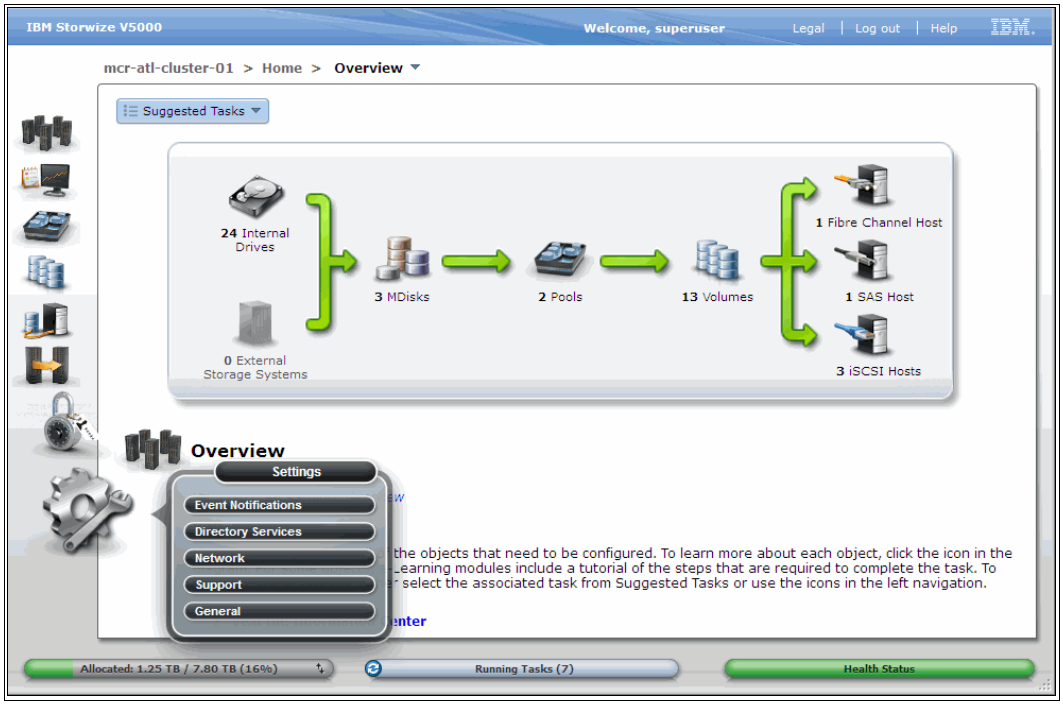

Figure 2-49 Enabling Call Home

2. Click **Email** → **Enable Email Event Notification**, as shown in Figure 2-50.

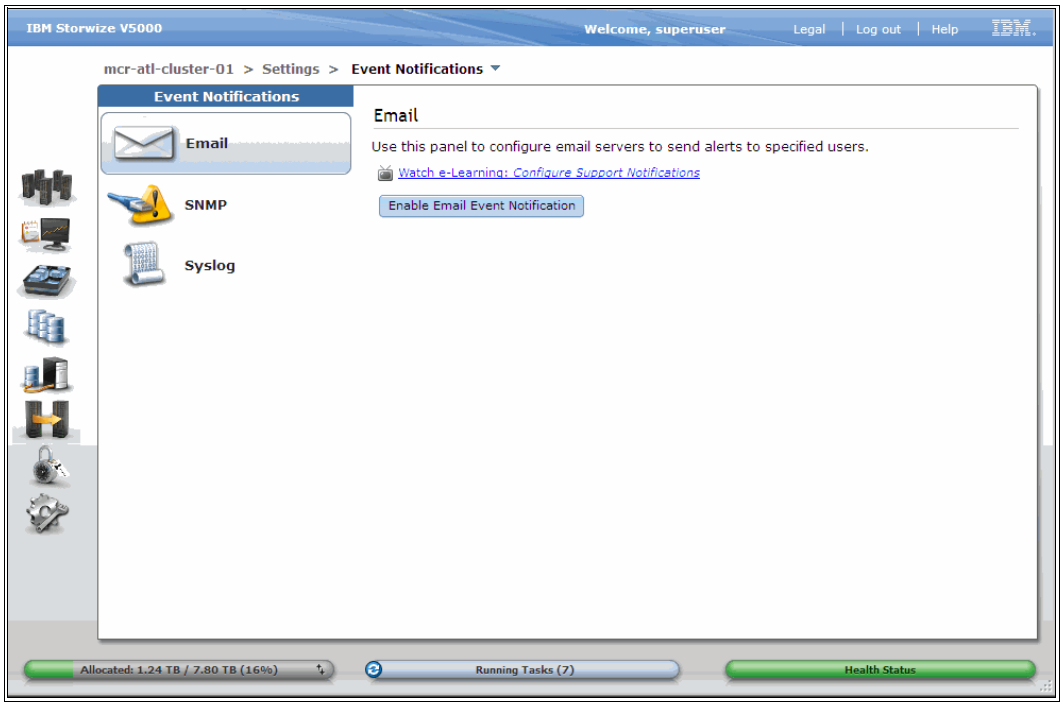

Figure 2-50 Selecting Event Notification

The wizard to configure Call Home starts, as shown Figure 2-51.

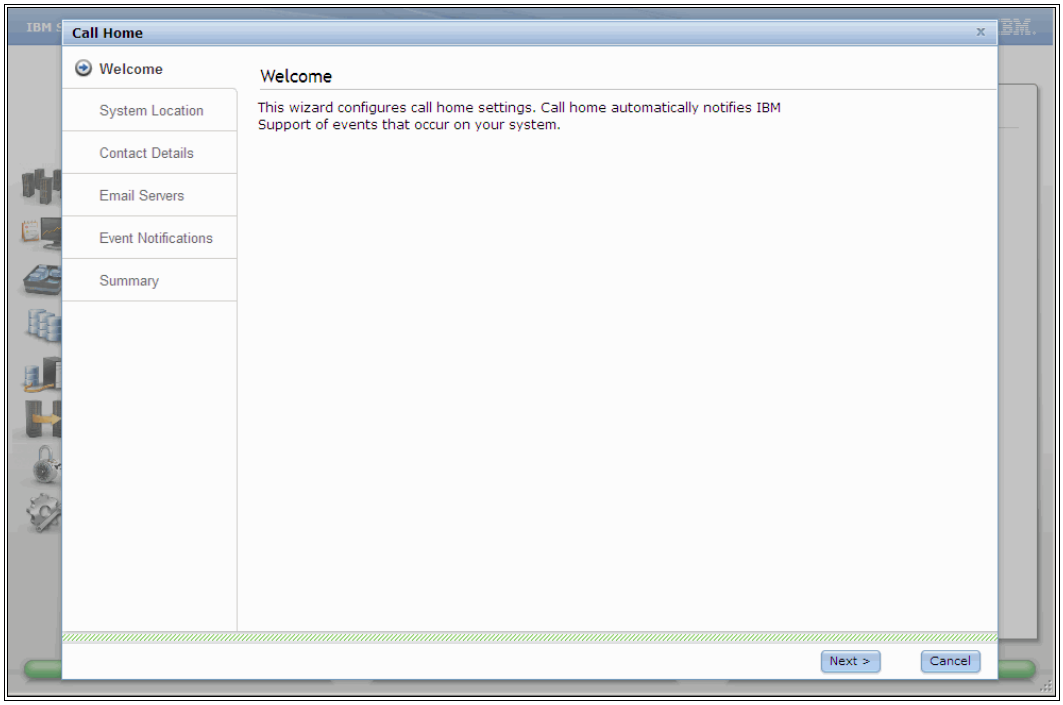

Figure 2-51 Call Home wizard

You are prompted to enter the details of the system, contact, event notification, and email server.

2.10.3 Service Assistant tool

The IBM Storwize V5000 is initially configured with three IP addresses: one service IP address for each node canister and a management IP address, which is set when the cluster is started.

The following methods are available to configure the Storwize V5000 system:

- The Inittool Program, as described in 2.9, “First-time setup” on page 41.
- The Service Assistant tool, which is described next.

Additionally, the Management IP and Service IP addresses can be changed within the GUI as shown in 3.4.8, “Settings menu” on page 137.

The Service Assistant (SA) tool is a web-based GUI that is used to service individual node canisters, primarily when a node has a fault and is in a service state. A node cannot be active as part of a clustered system while it is in a service state. The SA is available even when the management GUI is not accessible. The following information and tasks are included:

- Status information about the connections and the node canister.
- Basic configuration information, such as, configuring IP addresses.
- Service tasks, such as, restarting the common information model object manager (CIMOM) and updating the worldwide node name (WWNN).
- Details about node error codes and hints about what to do to fix the node error.

Important: The Service Assistant tool can be accessed only by using the superuser account.

The Service Assistance GUI is available by using a service assistant IP address on each node. The SA GUI is accessed through the cluster IP addresses by appending service to the cluster management URL. If the system is down, the only other method of communicating with the node canisters is through the SA IP address directly. Each node can have a single SA IP address on Ethernet port 1.It is recommended that these IP addresses are configured on all Storwize V5000 node canisters.

The default IP address of canister 1 is 192.168.70.121 with a subnet mask of 255.255.255.0.

The default IP address of canister 2 is 192.168.70.122, with a subnet mask of 255.255.255.0.

To open the SA GUI, enter one of the following URLs into any web browser:

- http(s)://*cluster IP address of your cluster*/*service*
- http(s)://*service IP address of a node*/service

Example:

- Management address: http://1.2.3.4/service
- SA access address: http://1.2.3.5/service

When you are accessing SA by using the `<cluster address>/service`, the configuration node canister SA GUI login window opens, as shown in Figure 2-52.

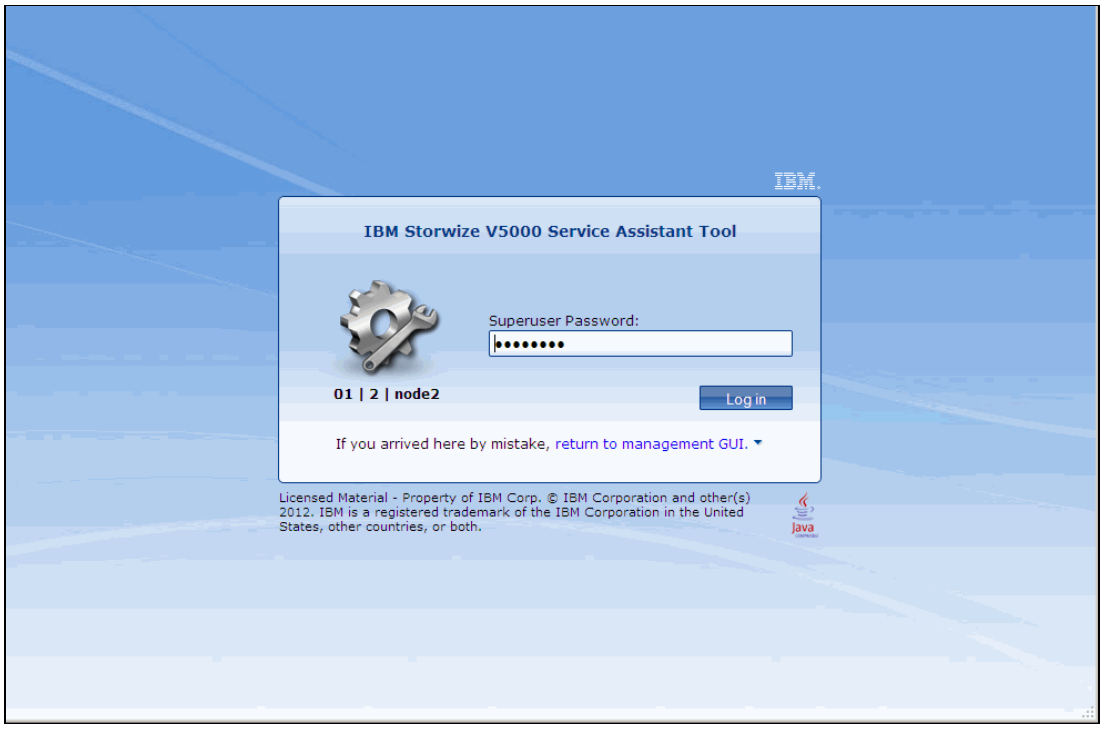

Figure 2-52 Service Assistant Login

The SA interfaces can view status and run service actions on other nodes and the node where user is connected.

After you are logged in, you see the Service Assistant Home window, as shown in Figure 2-53.

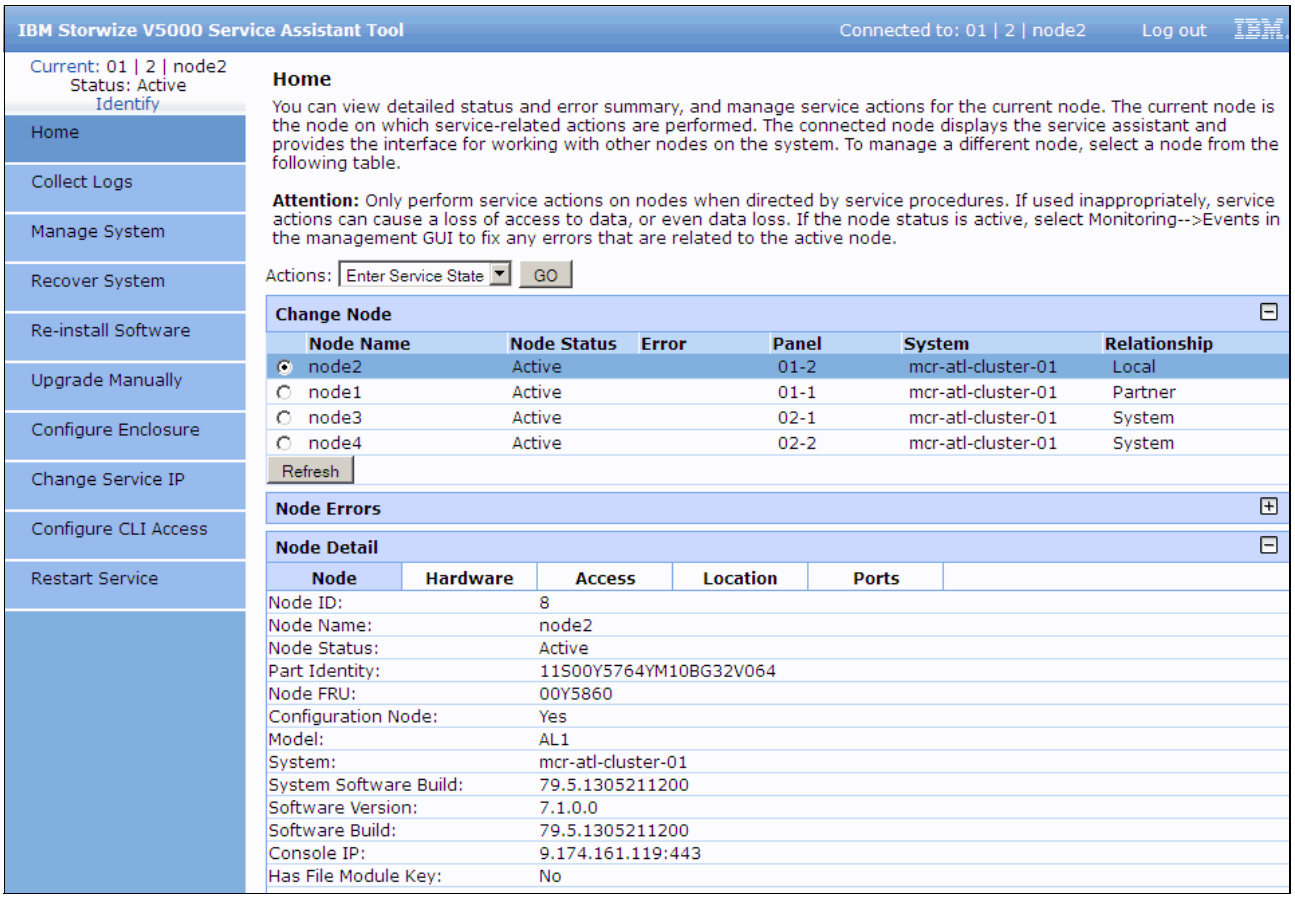

Figure 2-53 Service Assistant Home Window

The current canister node is displayed in the upper left corner of the GUI. In Figure 2-53, this is I/O Group 1 node 2. To change the canister, select the relevant node in the Change Node section of the window. You see the details in the upper left change to reflect the new canister.

The SA GUI provides access to service procedures and displays the status of the node canisters. These procedures should be carried out only if you directed to do so by IBM Support.

For more information about how to use the SA tool, see this website:

http://pic.dhe.ibm.com/infocenter/storwize/V5000_ic/index.jsp?topic=%2Fcom.ibm.storwize.V5000.641.doc%2Ftbrd_sagui_1938wd.html

3

Graphical user interface overview

This chapter provides an overview of the IBM Storwize V5000 graphical user interface (GUI) and shows how to navigate the configuration panels.

This chapter includes the following topics:

- Getting started
- Navigation
- Status Indicators menus
- Function icon menus
- Management GUI help

3.1 Getting started

This section provides information about accessing the IBM Storwize V5000 management GUI. It covers topics such as, supported browsers, log in modes, and the layout of the Overview panel.

3.1.1 Supported browsers

The IBM Storwize V5000 management software is a browser-based GUI. It is designed to simplify storage management by providing a single point of control for monitoring, configuration, and management.

For more information about supported browsers, see the IBM Storwize V5000 Information Center at this website:

`http://pic.dhe.ibm.com/infocenter/storwize/v5000_ic/index.jsp`

3.1.2 Access the management GUI

To access the management GUI, open a supported web browser and enter the management IP address or Hostname of the IBM Storwize V5000. The login panel is displayed, as shown in Figure 3-1.

Figure 3-1 IBM Storwize V5000 login panel

Default user name and password: Use the following information to log in to the IBM Storwize V5000 storage management:

- User Name: `superuser`
- Password: `passw0rd` (a zero replaces the letter O)

A successful login shows the Overview panel by default, as shown in Figure 3-2. Alternatively, the last opened window from the previous session is displayed.

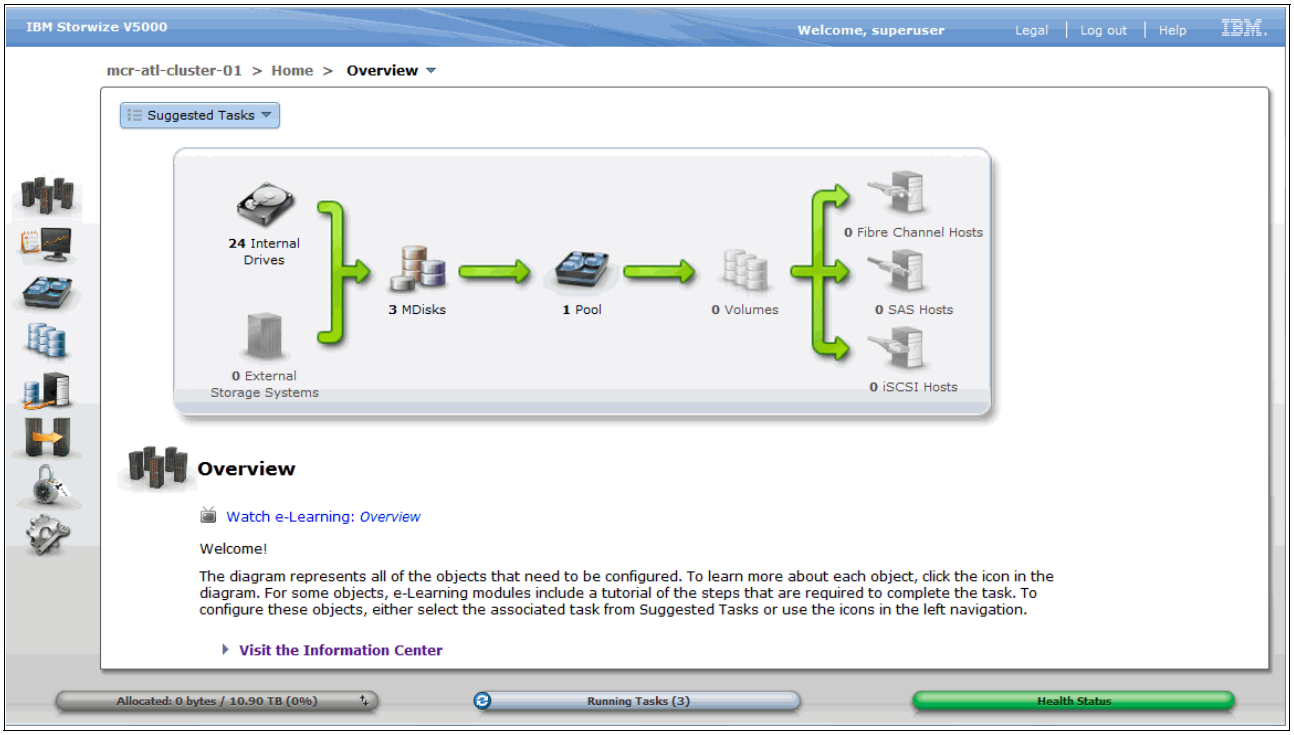

Figure 3-2 IBM Storwize V5000 overview panel

Figure 3-1 on page 76 shows the IBM Storwize V5000 login panel and the option to enable low graphics mode. This feature can be useful for remote access over narrow bandwidth links. The Function Icons no longer enlarge and list the available functions. However, you can navigate by clicking a Function Icon and by using the breadcrumb navigation aid.

For more information about the Function Icons, see 3.1.3, “Overview panel layout” on page 79.

For more information about the breadcrumb navigation aid, see 3.2.3, “Breadcrumb navigation aid” on page 84.

Figure 3-3 shows the management GUI in low graphics mode.

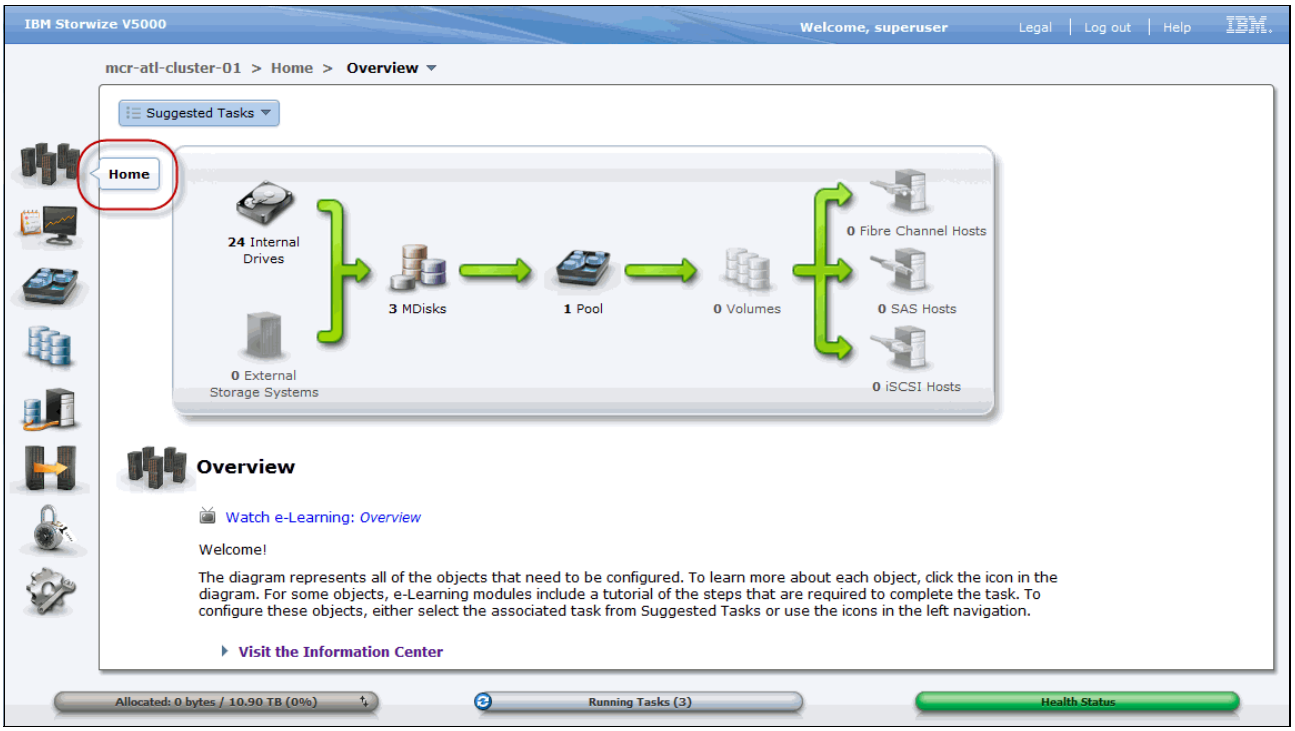

Figure 3-3 Management GUI low graphics mode

3.1.3 Overview panel layout

As shown in Figure 3-4, the Overview panel has three main sections: Function Icons, Extended Help, and Status Indicators.

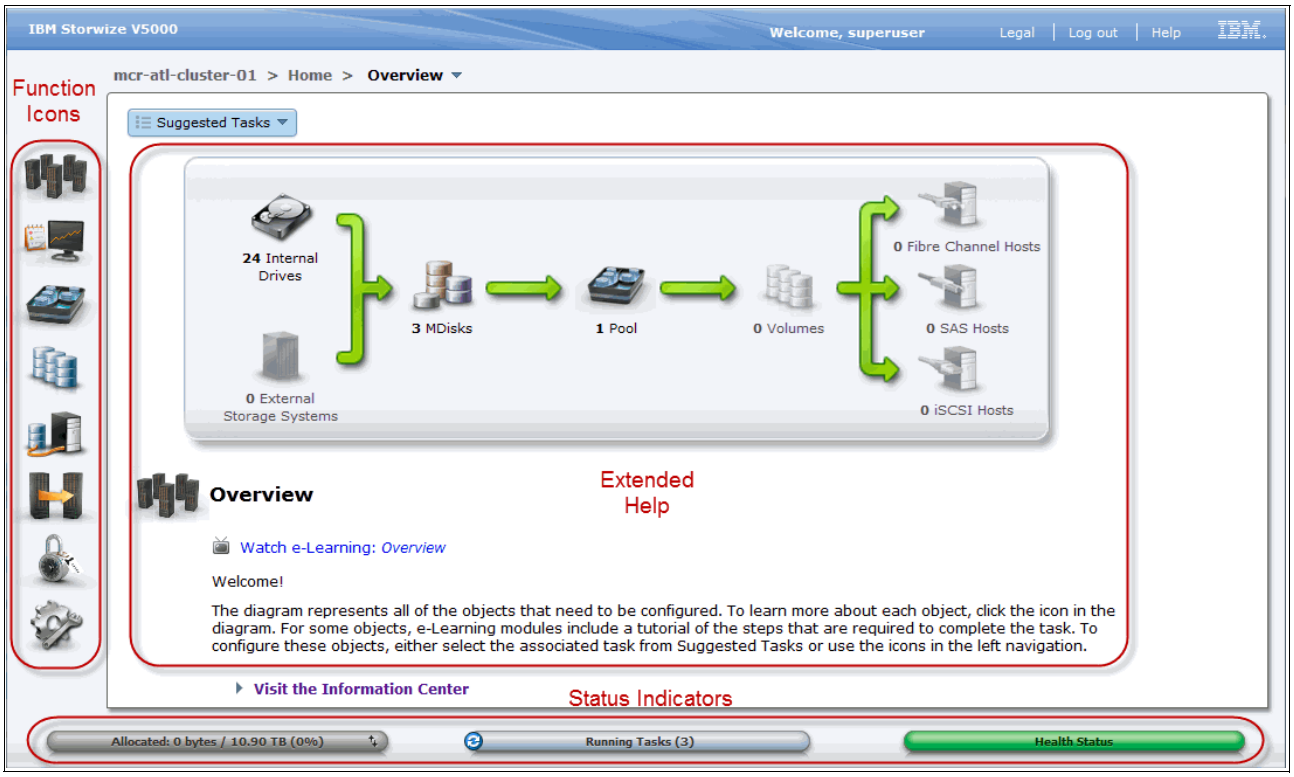

Figure 3-4 Three main sections of the IBM Storwize V5000 overview panel

The Function Icons section shows a column of images. Each image represents a group of interface functions. The icons enlarge with mouse hover and the following menus are shown:

- Home
- Monitoring
- Pools
- Volumes
- Hosts
- Copy Services
- Access
- Settings

The Extended Help section has a flow diagram that shows the available system resources. The flow diagram consists of system resource images and green arrows. The images represent the physical and logical elements of the system. The green arrows show the order to perform storage allocation tasks and highlight the various logical layers between the physical internal disks and the logical volumes.

Clicking the objects in this area shows more information. This information provides Extended Help references, such as, the online version of the Information Center and e-Learning modules. This information also provides direct links to the various configuration panels that relate to the highlighted image.

The Status Indicators section shows the following horizontal status bars:

- Allocated: Status that is related to the storage capacity of the system.
- Running Tasks: Status of tasks that are running and the recently completed tasks.
- Health Status: Status relating to system health, which is indicated by using the following color codes:
  - Green: Healthy
  - Yellow: Degraded
  - Red: Unhealthy

Hovering the mouse pointer and clicking the horizontal bars provides more information and menus, which is described in 3.3, “Status Indicators menus” on page 93.

3.2 Navigation

Navigating the management tool is simple and, like most systems, there are many ways to navigate. The two main methods are to use the Function Icons section or the Extended Help section of the Overview panel. For more information about these sections, see 3.1.3, “Overview panel layout” on page 79.

This section describes the two main navigation methods and introduces the well-known breadcrumb navigation aid and the Suggested Tasks aid. Information regarding the navigation of panels with tables also is provided.

3.2.1 Function icons navigation

Hovering the mouse pointer over one of the eight function icons on the left side of the panel enlarges the icon and provides a menu with which to access various functions. Move the pointer to the required function and click the function. Figure 3-5 shows the results of hovering the mouse pointer over a function icon.

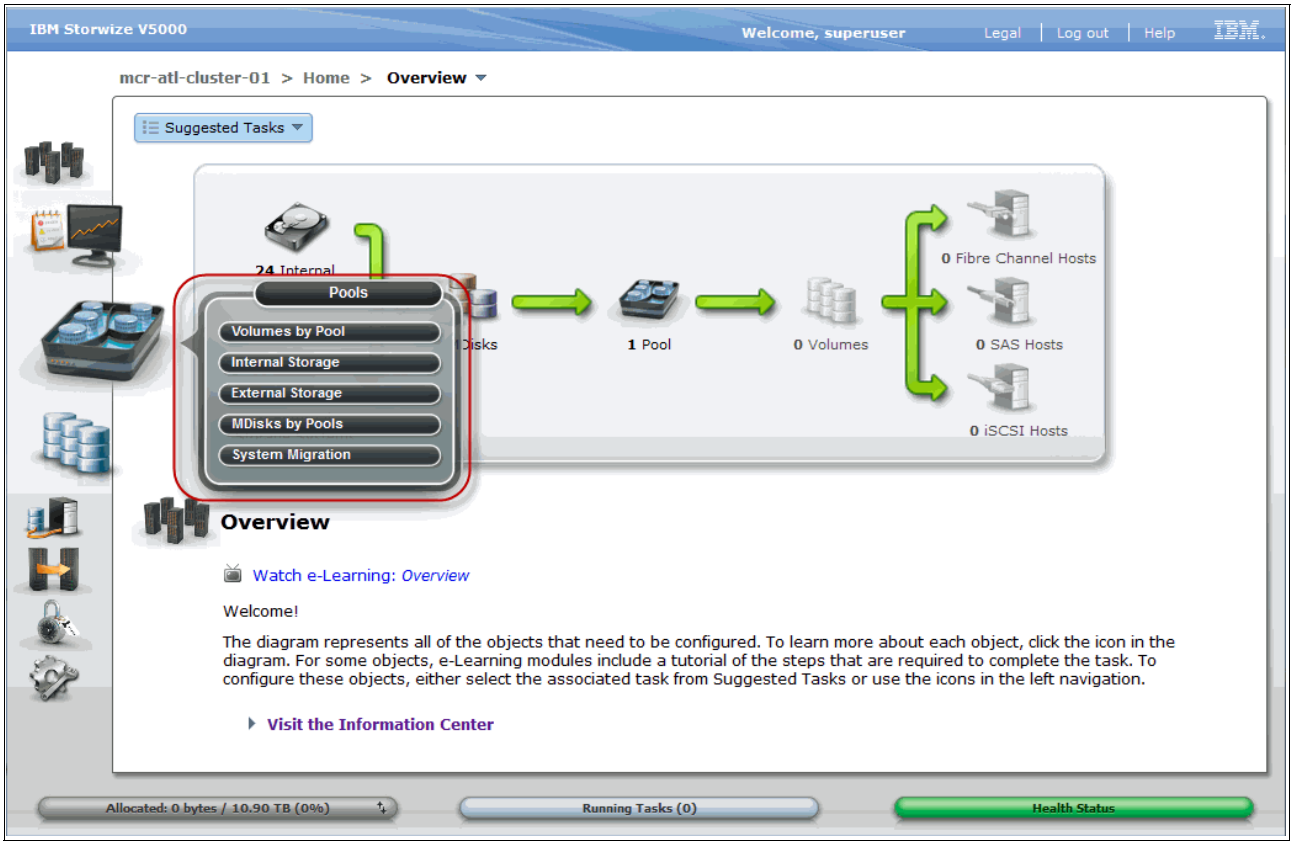

Figure 3-5 Hovering over a function icon

Figure 3-6 shows all of the menus with options under the Function Icons section.

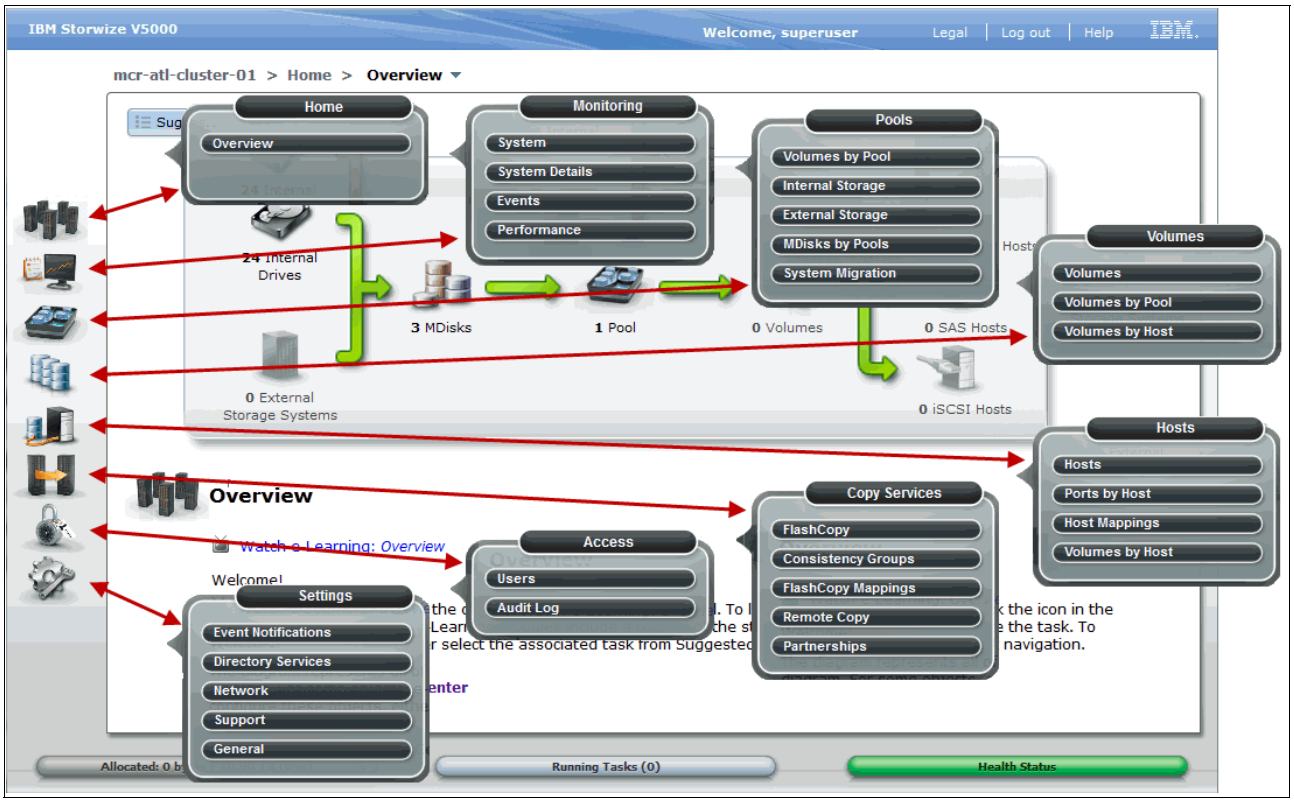

Figure 3-6 Options that are listed under Function Icons section

3.2.2 Extended help navigation

Selecting an image in the flow diagram of the Extended Help section in the Overview panel shows information beneath the flow diagram. This information contains links to e-Learning modules and configuration panels that are related to the selected image. This feature is convenient when the system is implemented because it is possible to work from left to right, following the flow, and select each object in order. Figure 3-7 on page 83 shows the selection of Internal Drives in the flow diagram. The information that is below the flow diagram relates to the internal storage.

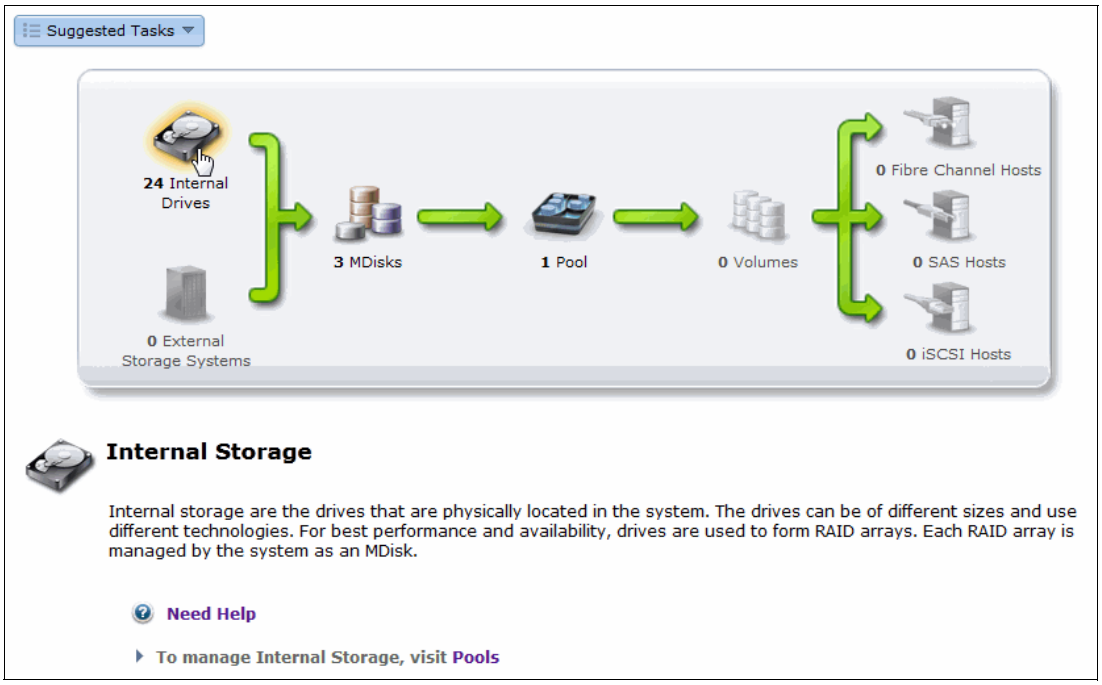

Figure 3-7 Navigating GUI with the extended help section

To access the e-Learning modules, click **Need Help**. To configure the internal storage, click **Pools**. Figure 3-8 shows the selection of Pools in the Extended Help section, which opens the Internal Storage panel.

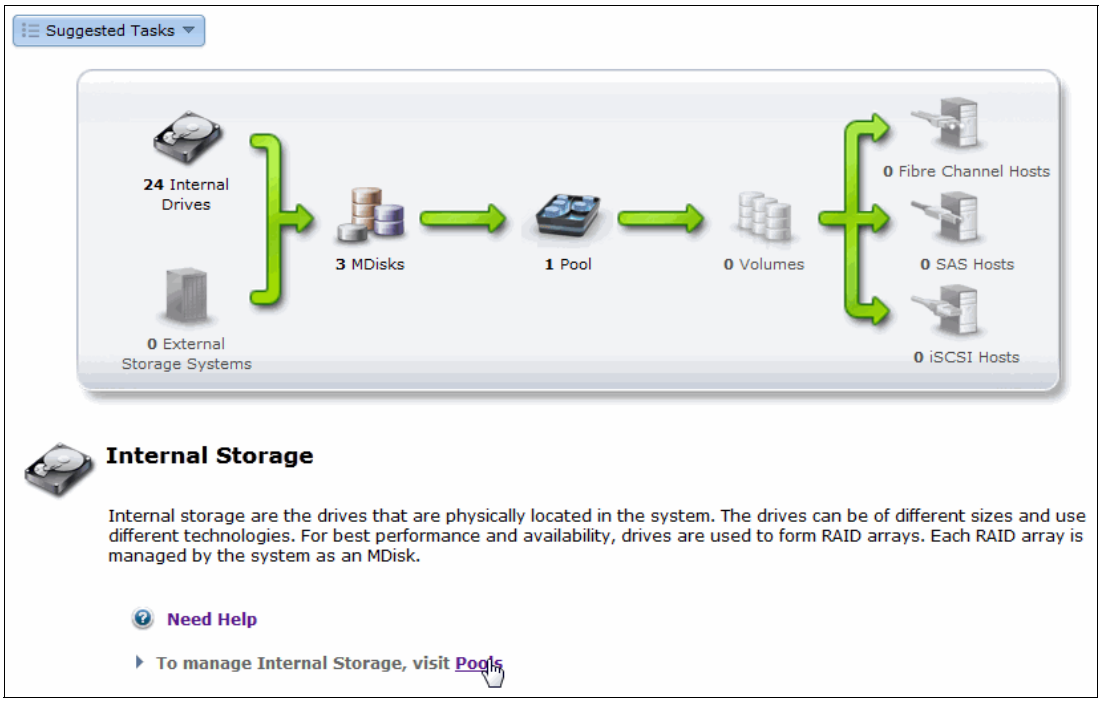

Figure 3-8 Using the extended help section

Figure 3-9 shows the Internal Storage panel, which is shown because Pools was selected in the information area of the Extended Help section.

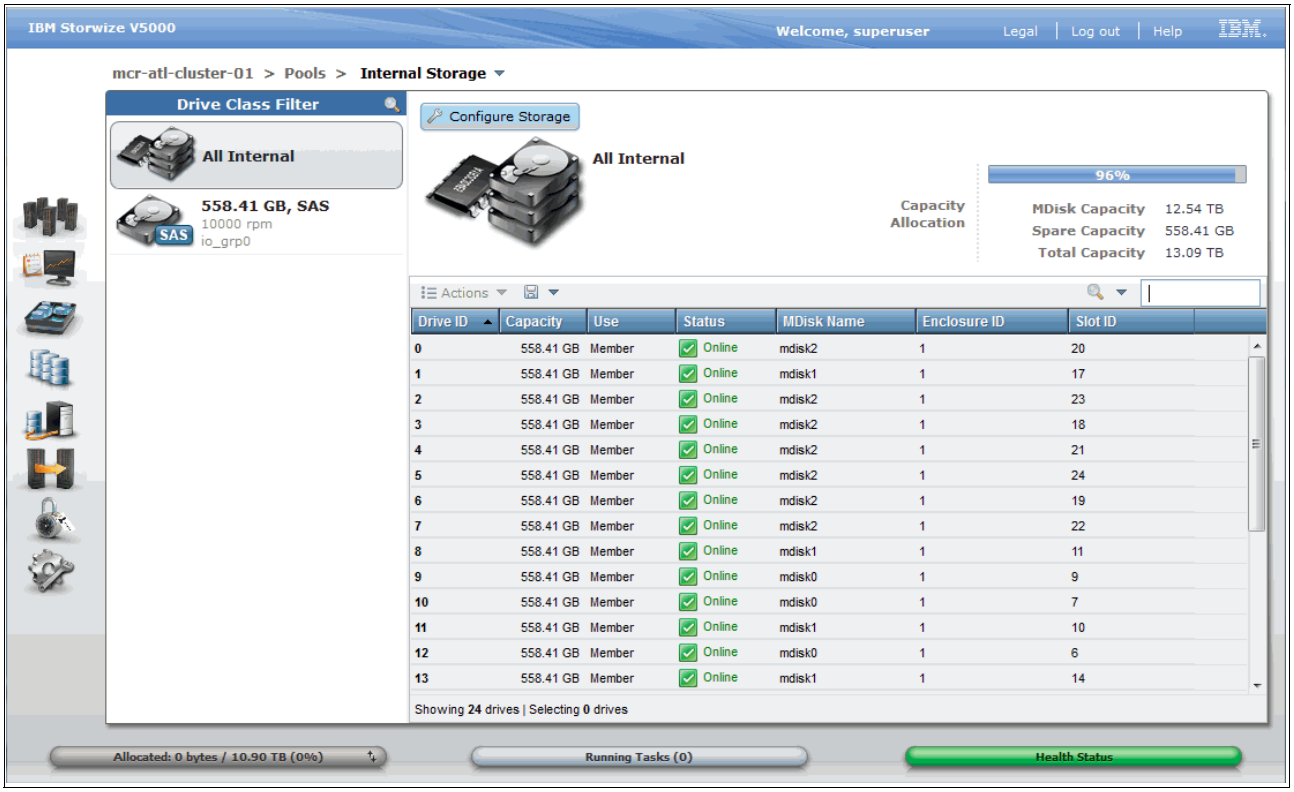

Figure 3-9 The internal storage configuration panel

3.2.3 Breadcrumb navigation aid

The IBM Storwize V5000 panels use the breadcrumb navigation aid to show the trail that was browsed. This breadcrumb navigation aid is in the top area of the panel and includes a System menu on the last breadcrumb. Figure 3-10 on page 85 shows the breadcrumb navigation aid for the System panel.

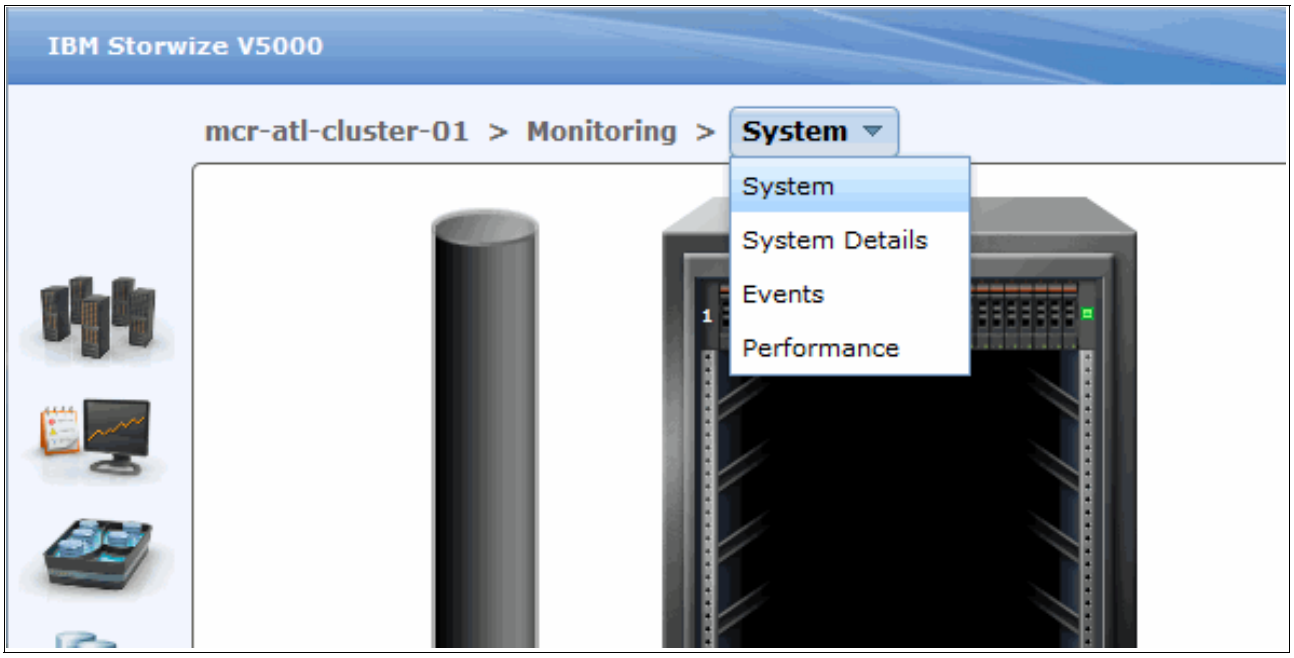

Figure 3-10 Breadcrumb navigation aid

3.2.4 Suggested Tasks feature

The Suggested Tasks feature is a navigation and configuration aid that is in the top area of the Overview panel. The list of suggested tasks changes, depending on the configuration of the system. This aid can be useful to follow during the system installation process. Figure 3-11 shows the Suggested Tasks navigation and configuration aid.

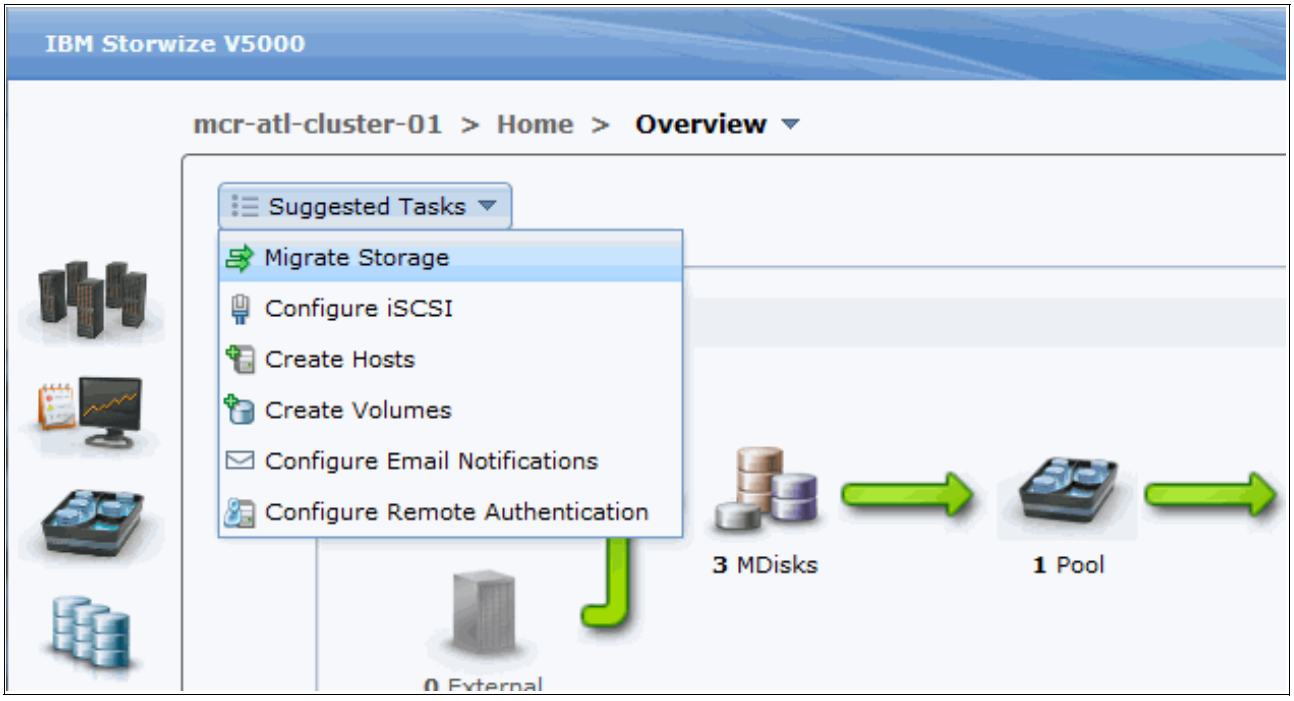

Figure 3-11 Suggested Tasks navigation and configuration aid

3.2.5 Presets

The management GUI contains a series of preestablished configuration options that are called *presets* that use commonly used settings to quickly configure objects on the system. Presets are available for creating volumes and IBM FlashCopy mappings and for setting up a RAID configuration. Figure 3-12 shows the available internal storage presets.

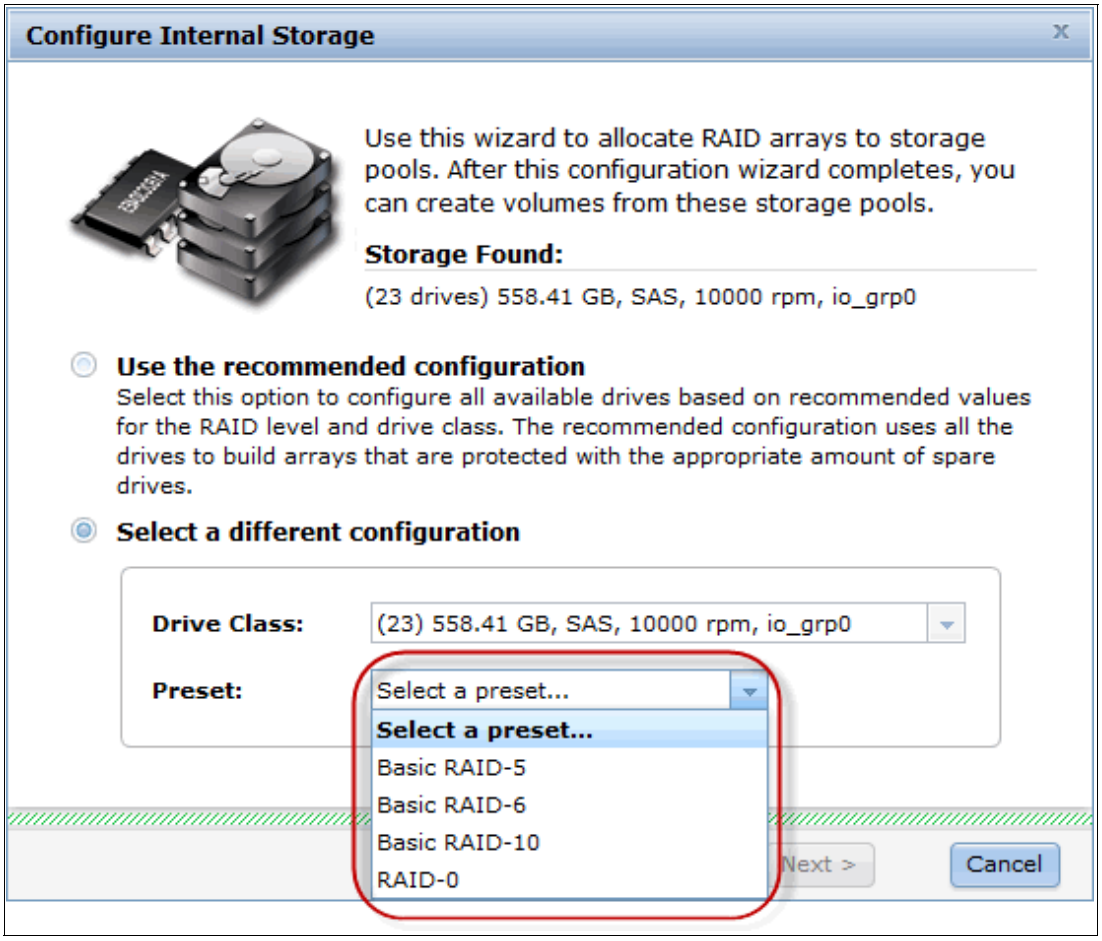

Figure 3-12 Internal storage preset selection

3.2.6 Access actions

The IBM Storwize V5000 functional panels provide access to various actions that can be performed, such as, modify attributes and rename, add, or delete objects. The available actions menus can be accessed by using one of two main methods: highlight the resource and use the Actions drop-down menu (as shown in Figure 3-13), or right-click the resources, as shown in Figure 3-14.

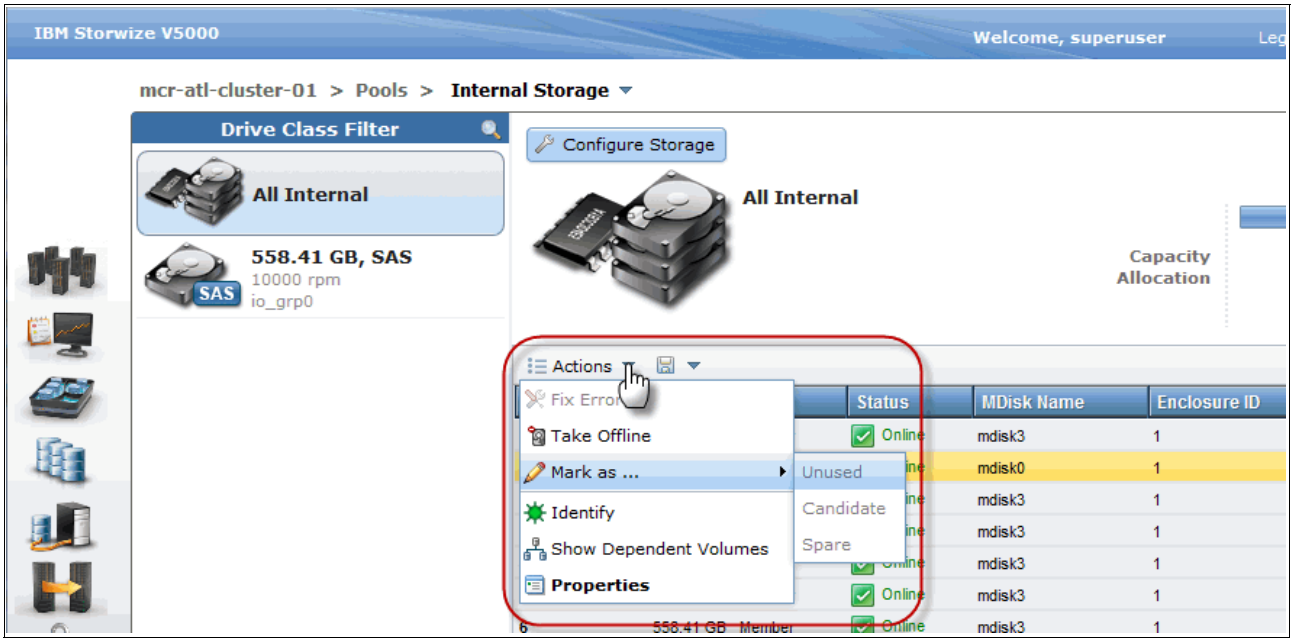

Figure 3-13 Actions menu

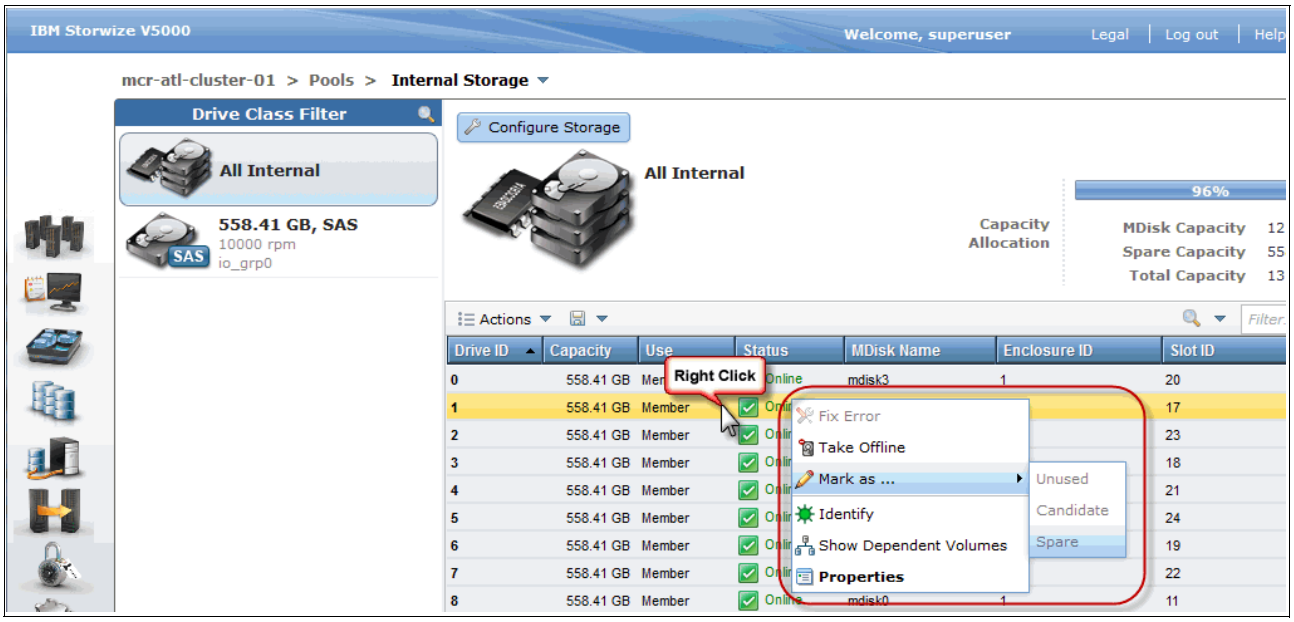

Figure 3-14 Right-clicking the Actions menu

3.2.7 Task progress

An action starts a running task and shows a task progress panel, as shown in Figure 3-15. Click **Details** to show the underlying command-line interface (CLI) commands. The commands are highlighted in blue and can be pasted into a configured IBM Storwize V5000 SSH terminal session, if required. This is useful when you are developing CLI scripts.

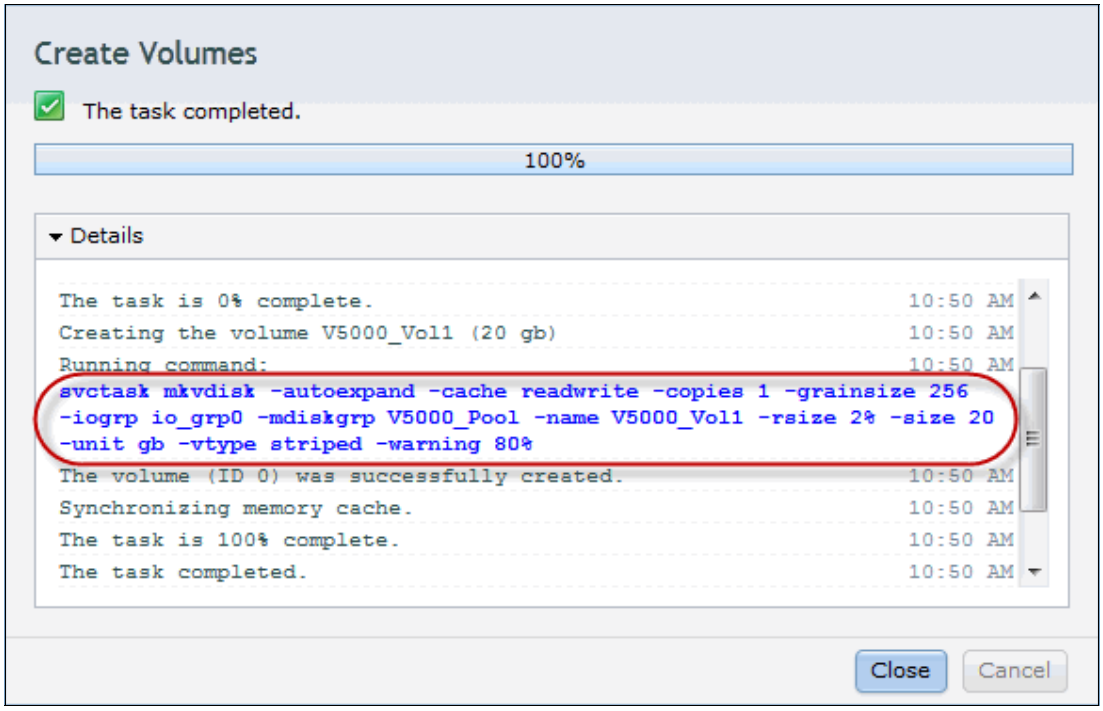

Figure 3-15 Task progress panel

3.2.8 Navigating panels with tables

Many of the configuration and status panels show information in a table format. This section describes the following useful methods to navigate panels with rows and columns:

- Sorting columns
- Reordering columns
- Adding or removing columns
- Multiple selections
- Filtering objects

Sorting columns

Columns can be sorted by clicking the column heading. Figure 3-16 on page 89 shows the result of clicking the heading of the Capacity column. The table is now sorted and lists volumes with the least amount of capacity at the top of the table.

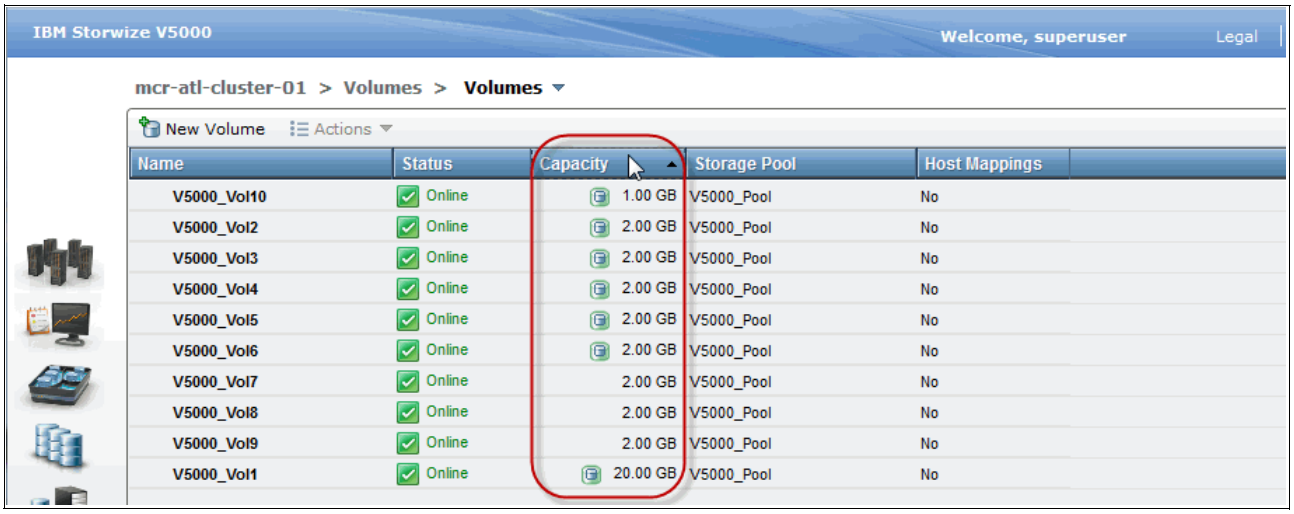

Figure 3-16 Sorting columns by clicking the column heading

Reordering columns

Columns can be reordered by dragging the column to the required location. Figure 3-17 shows the location of the column with the heading Host Mappings positioned in the last column. Dragging this heading reorders the columns in the table.

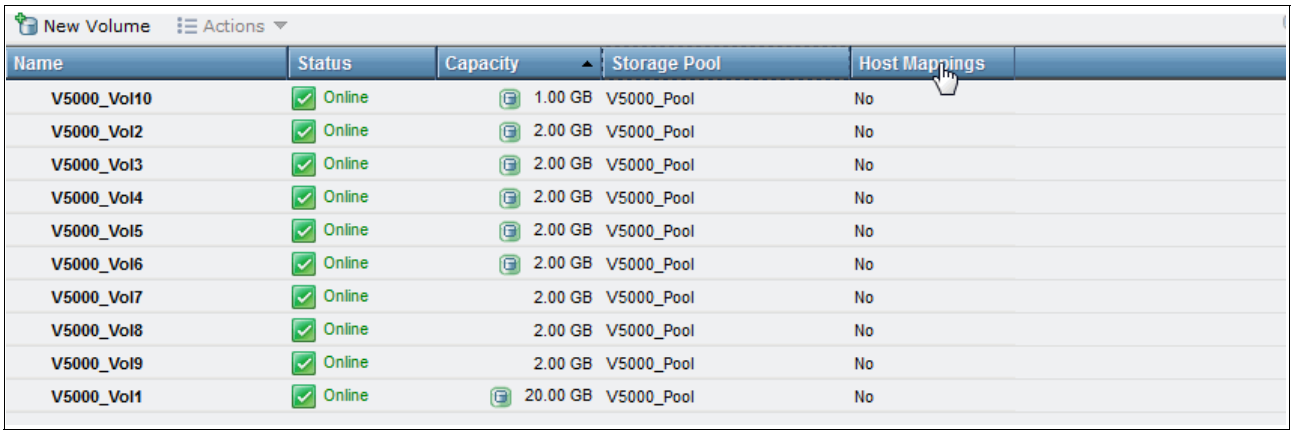

Figure 3-17 Reordering columns

Figure 3-18 shows the column heading Host Mappings as it is dragged to the required location.

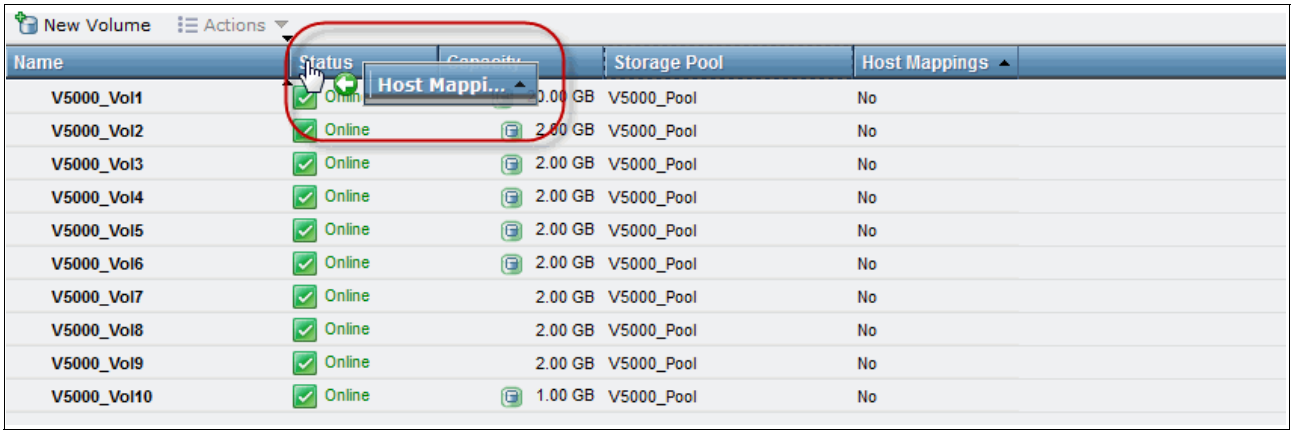

Figure 3-18 Dragging a column heading to the required location

Figure 3-19 shows the result of dragging the column heading Host Mappings to the new location.

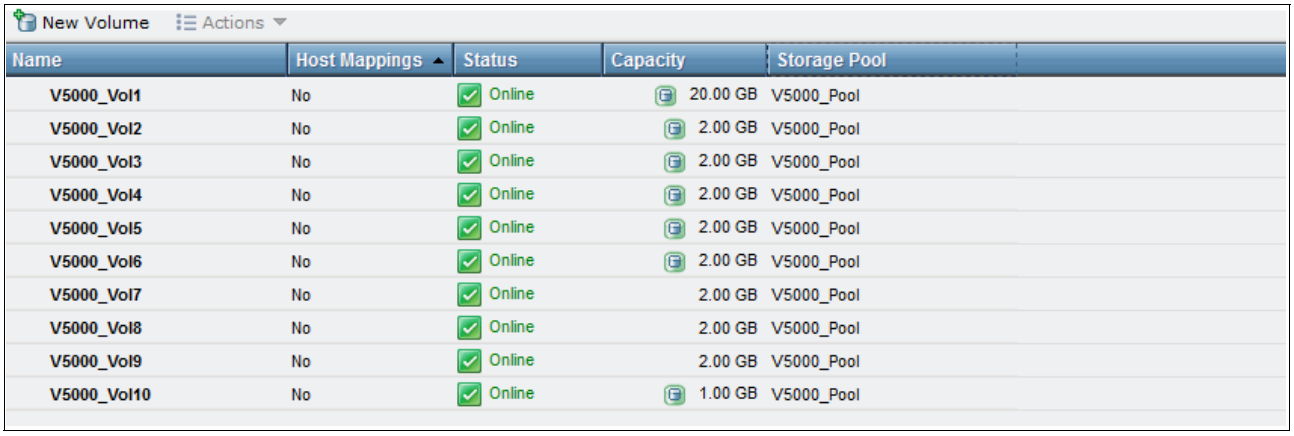

Figure 3-19 Reordering column headings

Adding or removing columns

To add or remove a column, right-click the heading bar and select the required column headings by selecting the box that is next to the heading name. Figure 3-20 shows the addition of the column heading Real Capacity.

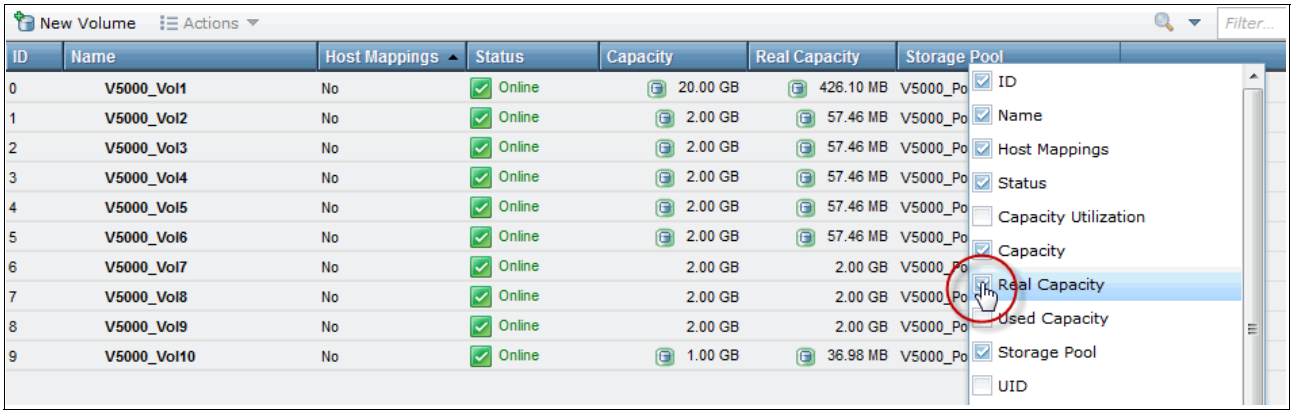

Figure 3-20 Adding column heading Real Capacity

Important: Some users might run into a problem in which a context menu from the Firefox browser is shown by right-clicking to change the column heading. This issue can be fixed by clicking in Firefox: **Tools** → **Options** → **Content** → **Advanced (for Java setting)** → **Select: Display or replace context menus**.

The web browser requirements and recommended configuration settings to access the IBM Storwize V5000 management GUI can be found in the IBM Storwize V5000 Information Center at this website:

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

Multiple selections

By using the management tool, you also can select multiple items in a list by using a combination of the Shift or Ctrl keys.

Using the Shift key

To select multiple items in a sequential order, click the first item that is listed, press and hold the Shift key, and then click the last item in the list. All of the items between the first and last items are selected, as shown in Figure 3-21.

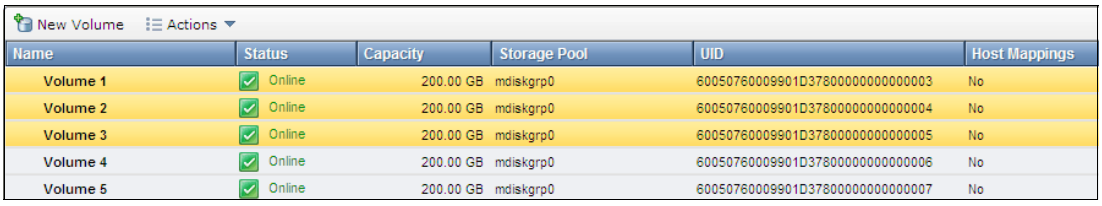

Figure 3-21 Selection of three sequential items

Using the Ctrl key

To select multiple items that are not in sequential order, click the first item, press and hold the Ctrl key, and then click the other items that you need. Figure 3-22 on page 92 shows the selection of two non-sequential items.

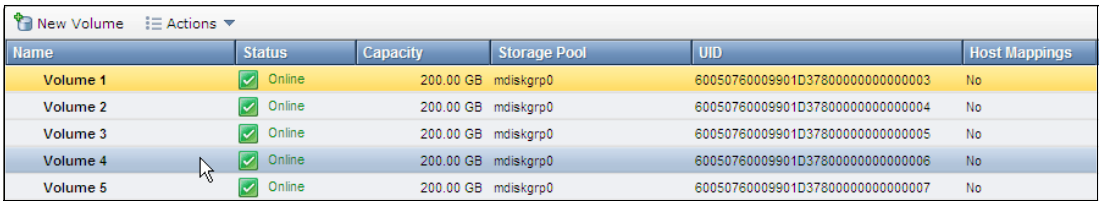

Figure 3-22 Selecting two non-sequential items

Figure 3-23 shows the result of the use of the Ctrl key to select multiple non-sequential items.

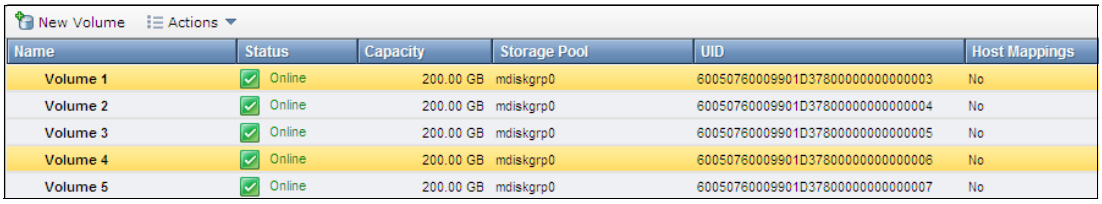

Figure 3-23 Result of selecting two non-sequential items

Filtering objects

To focus on a subset of the listed items that are shown in a panel with columns, use the filter field that is found at the upper right side of the table. This tool shows items that match the value that is entered. Figure 3-24 shows the text **`Vol1`** was entered into the filter field. Now, only volumes with the text **`Vol1`** in any column are listed and the filter word also is highlighted.

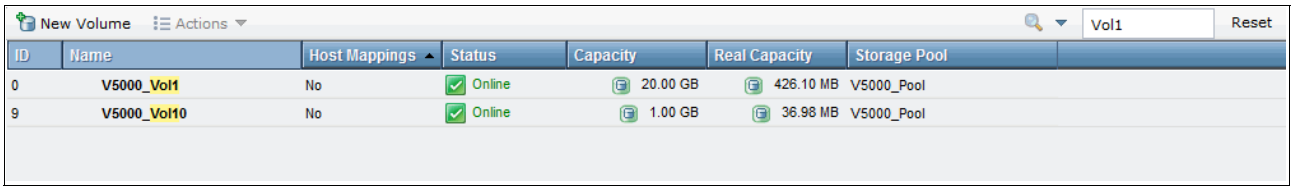

Figure 3-24 Filtering objects to display a subset of the volumes

Filter by column

Click the magnifying glass that is next to the filter field to activate the filter by column feature. Figure 3-25 shows the Filter by Column drop-down menu. This allows the filter field value to be matched to the values of a specific column.

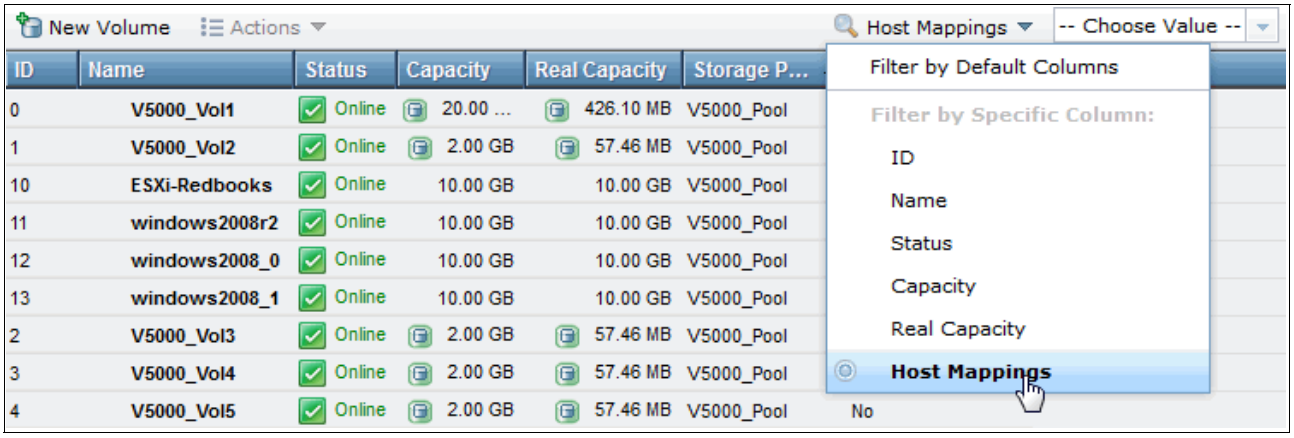

Figure 3-25 Filter by column

Figure 3-26 shows the column filter is set to Host Mappings, the filter value set to Yes, and the resulting Volumes with Hosts mapped.

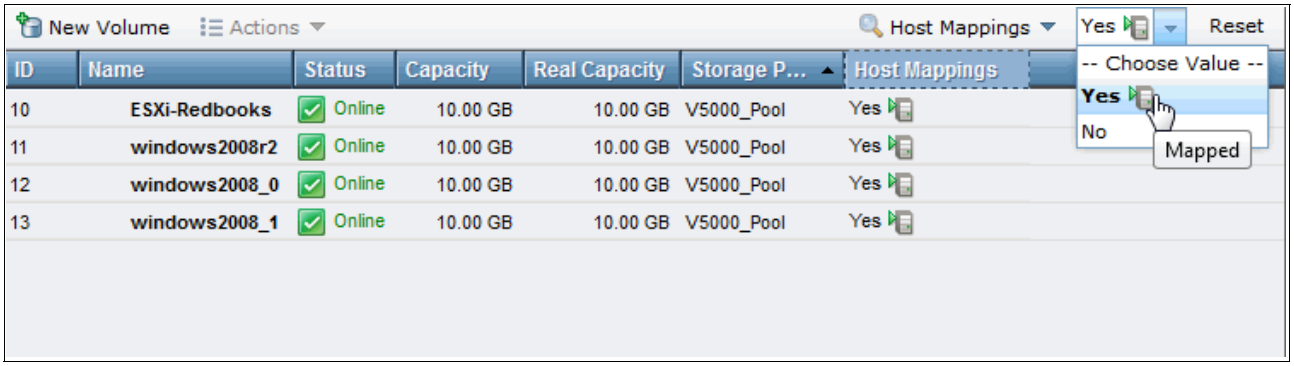

Figure 3-26 Choosing filter value

3.3 Status Indicators menus

This section provides more information about the horizontal bars that are shown at the bottom of the management GUI panels. The bars are *status indicators*, and include associated bar menus. This section describes the Allocated, Running Tasks, and Health Status bar menus.

3.3.1 Horizontal bars

As described in 3.1.3, “Overview panel layout” on page 79, the status indicators include the allocated, running tasks, and health status horizontal bars and are shown at the bottom of the panel. The status indicators are color-coded and draw attention to alerts, events, and errors. Hovering over and clicking the bars shows more menus.

3.3.2 Allocated status bar menu

The allocated status bar shows the capacity status. Hovering over the image of two arrows on the right side of the Allocated status bar shows a description of the allocated menu comparison that is in use. Figure 3-27 on page 94 shows the comparison of the used capacity to the real capacity.

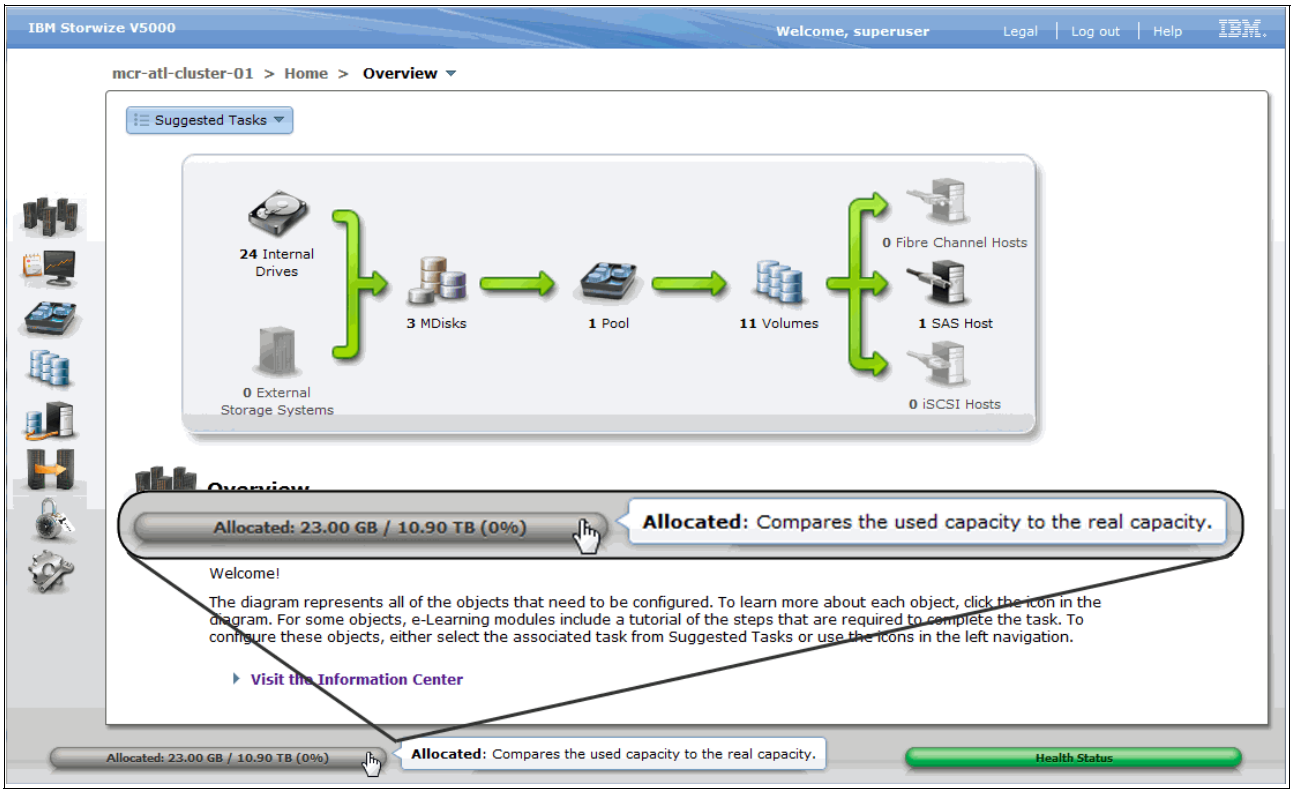

Figure 3-27 Allocated bar that compares used capacity to real capacity

To change the allocated bar comparison, click the image of the two arrows on the right side of the Allocated status bar. Figure 3-28 on page 95 shows the new comparison of virtual capacity to real capacity.

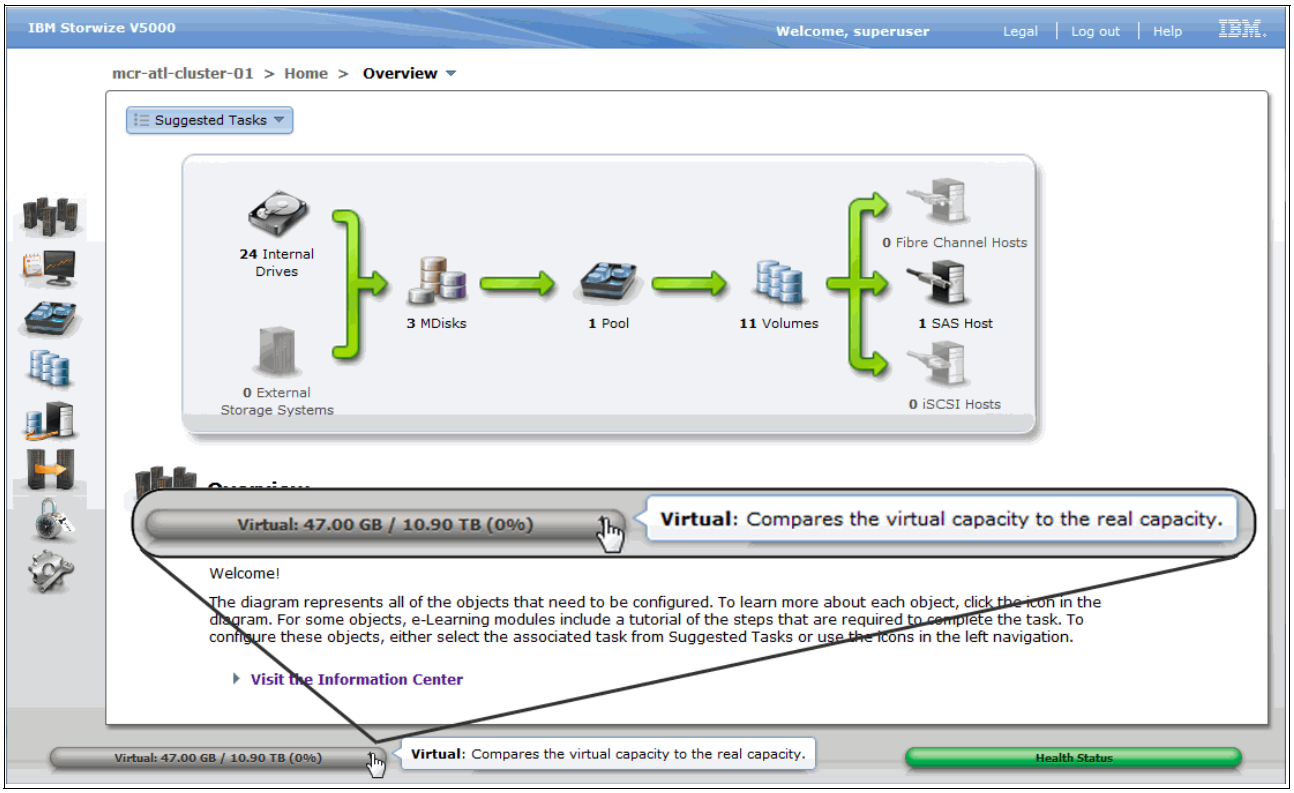

Figure 3-28 Changing the allocated menu comparison, virtual capacity to real capacity

3.3.3 Running tasks bar menu

To show the Running Tasks bar menu, click the circular image to the left of the running tasks status bar. This menu lists running and recently completed tasks and groups similar tasks. Figure 3-29 shows the Running Tasks menu.

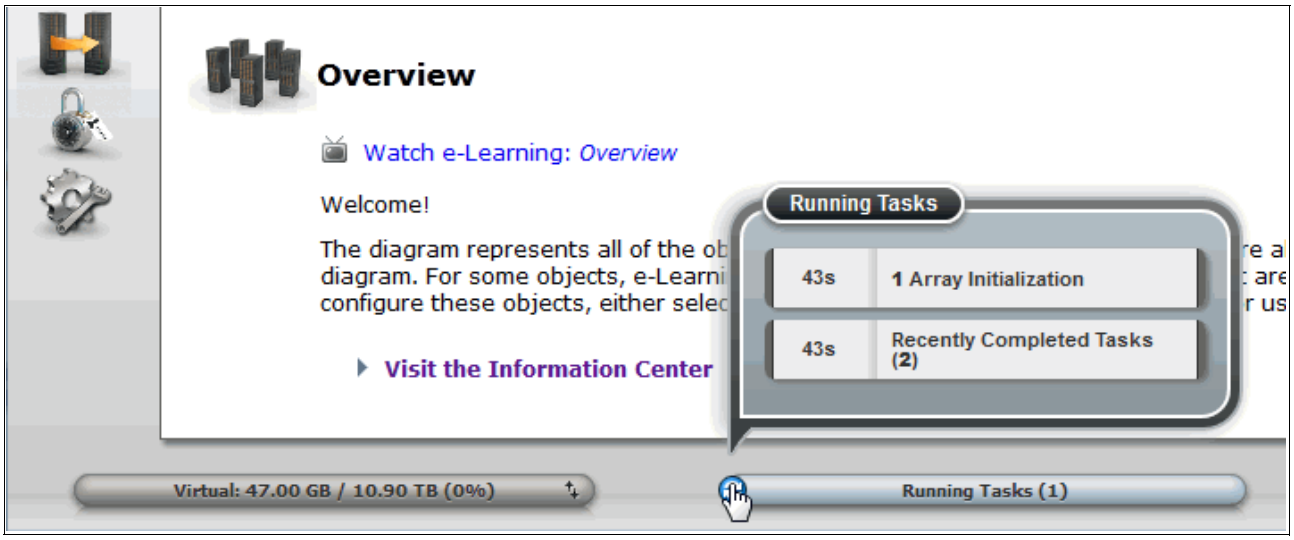

Figure 3-29 Running Tasks menu

For an indication of task progress, browse to the Running Tasks bar menu and click the task. Figure 3-30 shows the selection of a task from the Running Tasks menu.

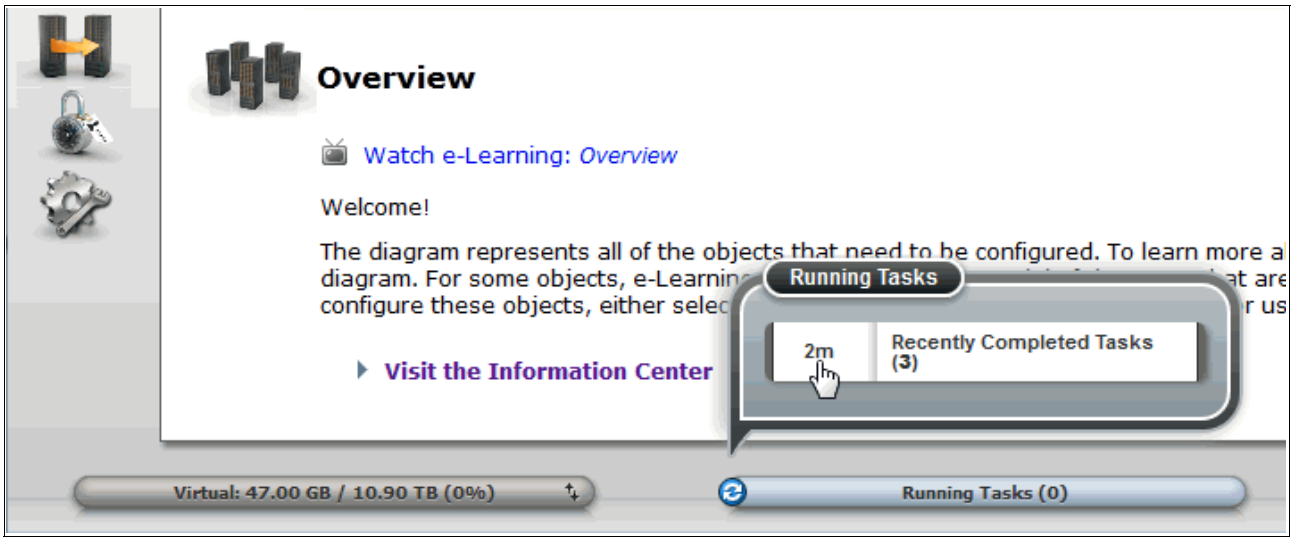

Figure 3-30 Selecting a task from the Running Task menu

Figure 3-31 shows the Recently Completed tasks panel.

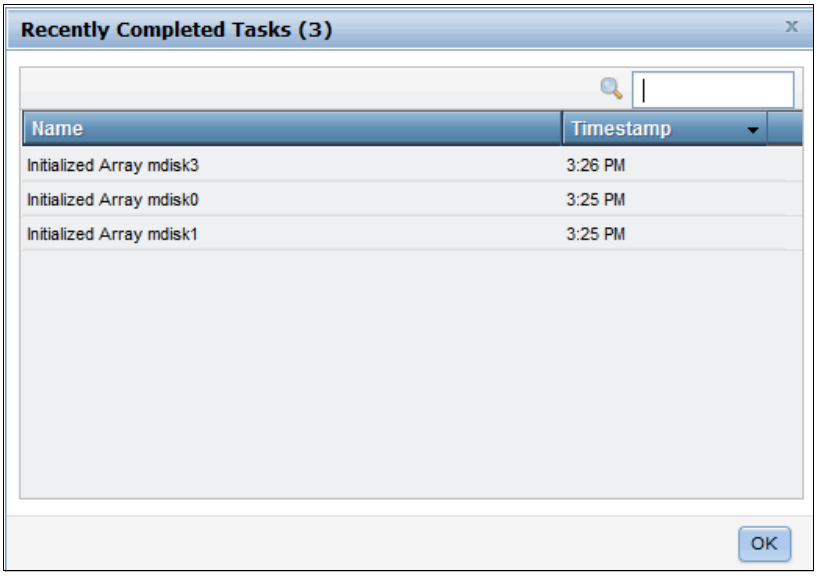

Figure 3-31 Recently Completed tasks panel

3.3.4 Health status bar menu

The health status bar provides an indication of the overall health of the system. The following color of the status bar indicates the state of IBM Storwize V5000:

- Green: Healthy
- Yellow: Degraded
- Red: Unhealthy

If a status alert occurs, the health status bar can turn from green to yellow or to red. To show the health status menu, click the attention icon on the left side of the health status bar, as shown in Figure 3-32.

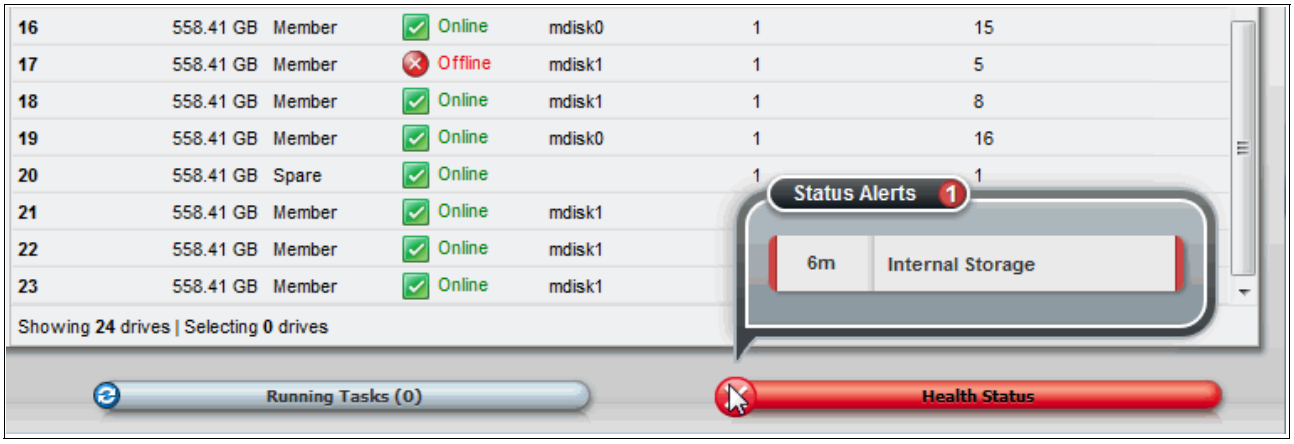

Figure 3-32 Health status menu

The health status bar menu shows the system as Unhealthy and provides a description of Internal Storage for the type of event that occurred. To investigate the event, open the health status bar menu and click the description of the event, as shown in Figure 3-33.

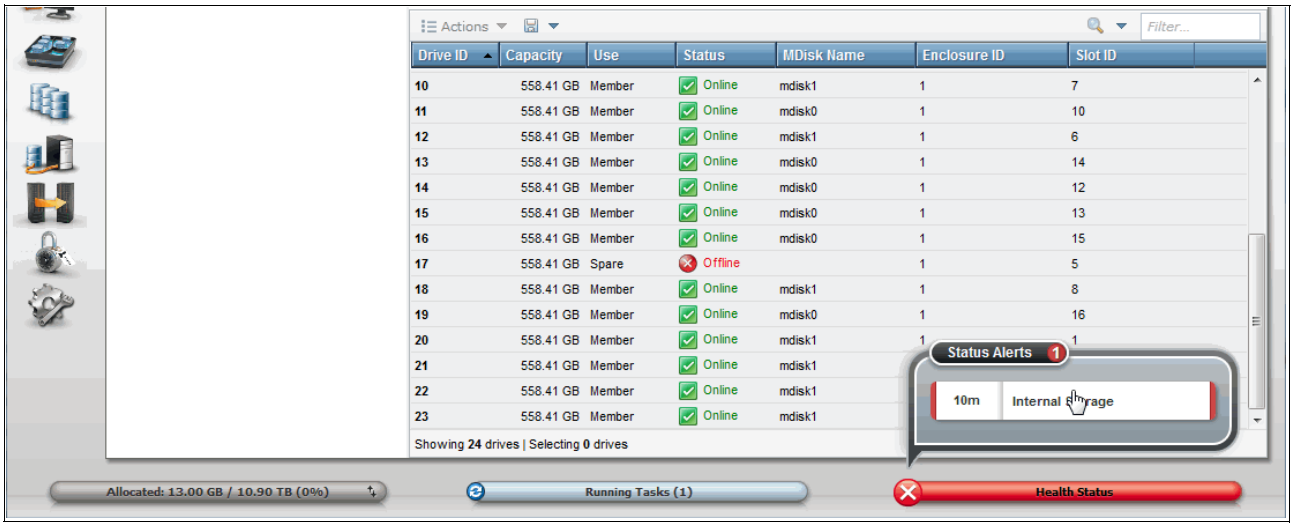

Figure 3-33 Status and description of an alert via the health status menu

Click the description of the event in the health status menu to show the Events panel (**Monitoring** → **Events**), as shown in Figure 3-34. This panel lists all events and provides directed maintenance procedures (DMPs) to help resolve errors. For more information, see “Events panel” on page 105.

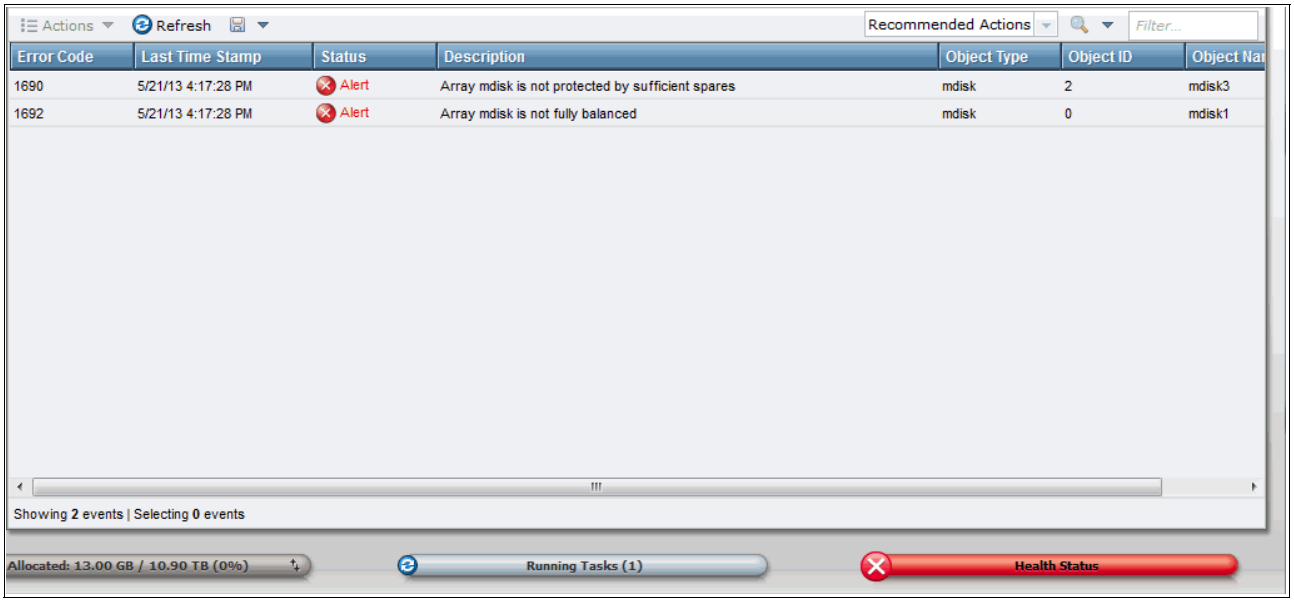

Figure 3-34 Events panel via health status menu

3.4 Function icon menus

The IBM Storwize V5000 management GUI provides function icons that are an efficient and quick mechanism that is used for navigation. As described in section 3.1.3, “Overview panel layout” on page 79, each graphic on the left side of the panel is a function icon that presents a group of interface functions. Hovering over one of the eight function icons shows a menu that lists the functions. Figure 3-35 shows all of the Function Icon menus.

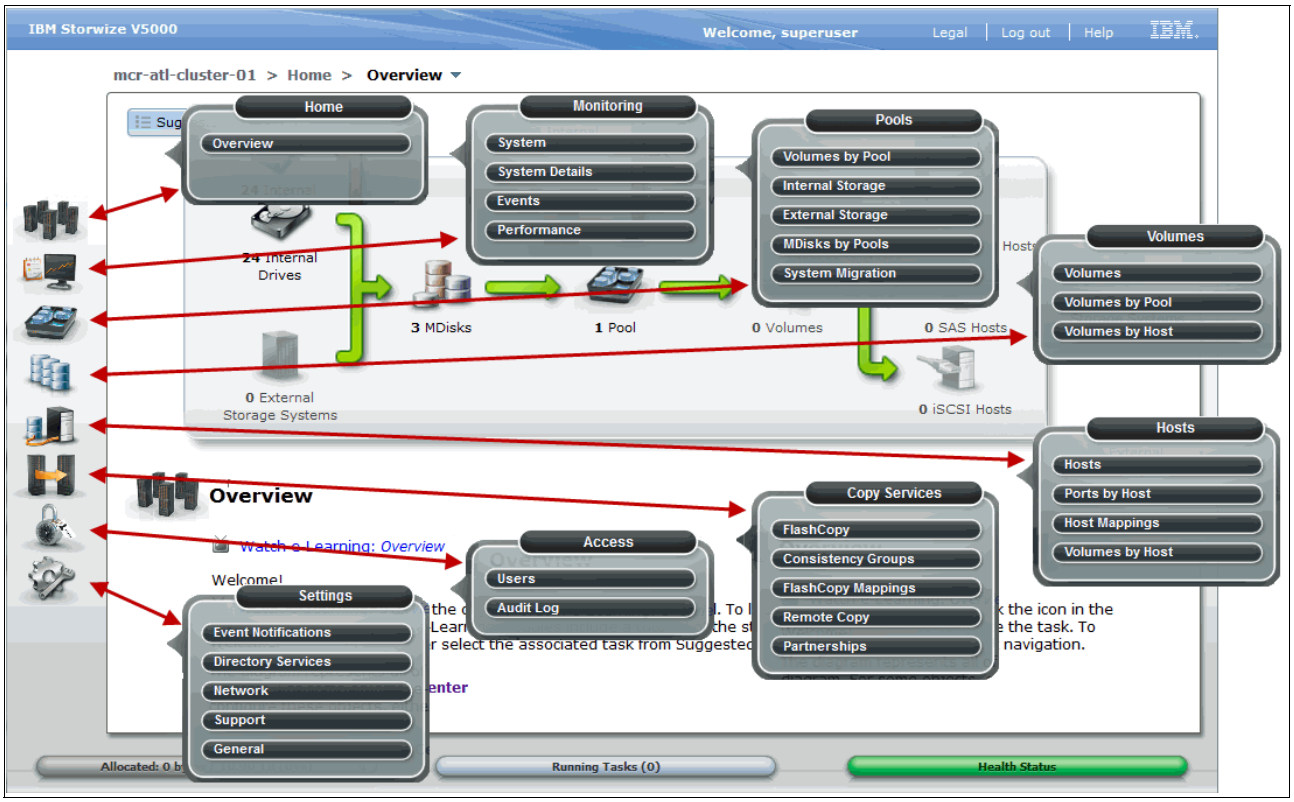

Figure 3-35 All Function Icon menus

3.4.1 Home menu

As shown in Figure 3-36, the Home menu provides access to the Overview panel.

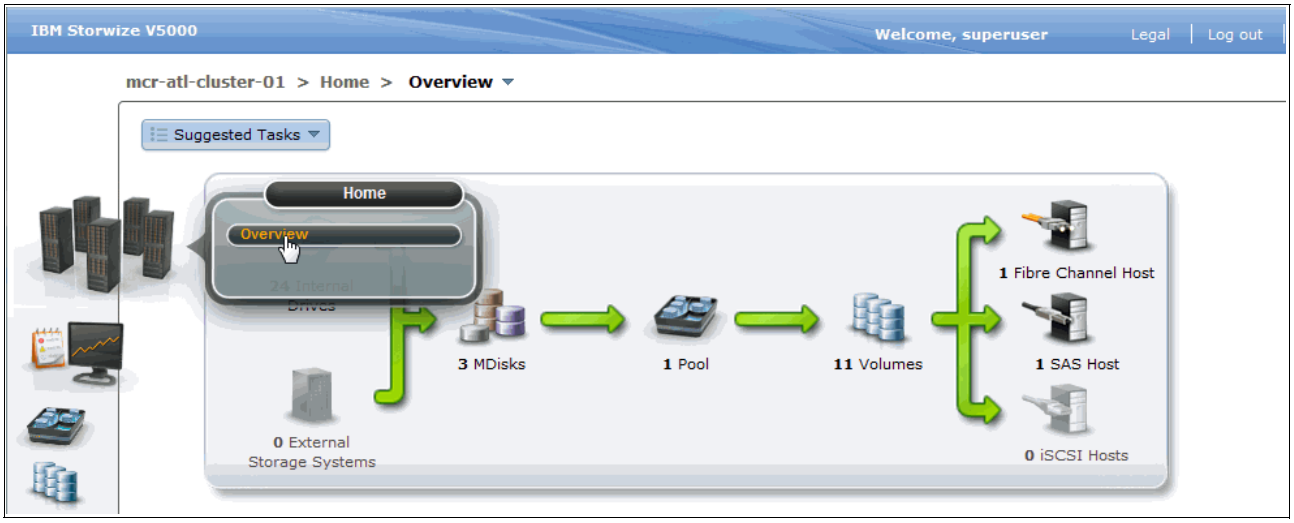

Figure 3-36 Home menu

To see the Overview panel, select **Overview** in the Home menu to open the panel. For more information, see 3.1.3, “Overview panel layout” on page 79.

3.4.2 Monitoring menu

As shown in Figure 3-37, the Monitoring menu provides access to the System, System Details, Events, and Performance panels.

Figure 3-37 Monitoring menu

System panel

Select **System** in the Monitoring menu to open the panel. The System panel (as shown in Figure 3-38 on page 101), shows capacity usage, enclosures, and all drives in the system.

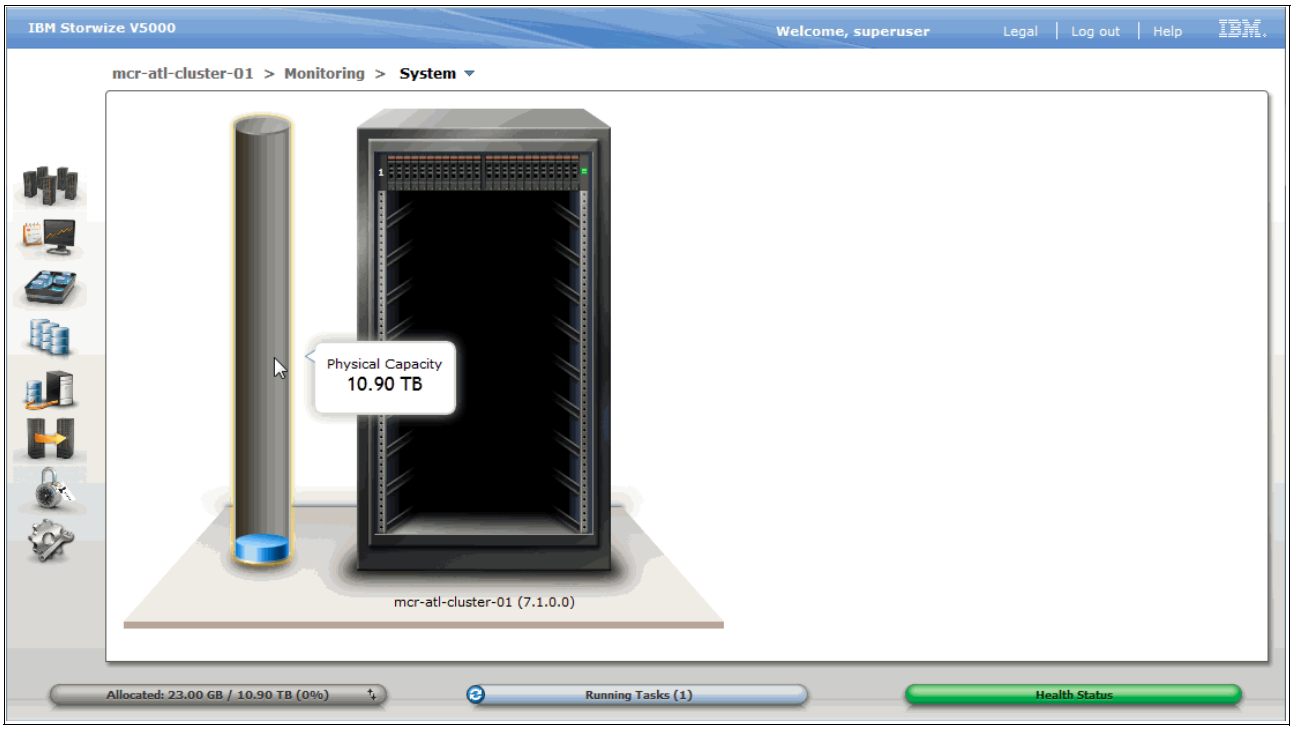

Figure 3-38 The System panel

Selecting the name and version of the system shows more information about storage allocation. The information is presented under two tabs: Info and Manage. Figure 3-39 shows the System panel Info tab.

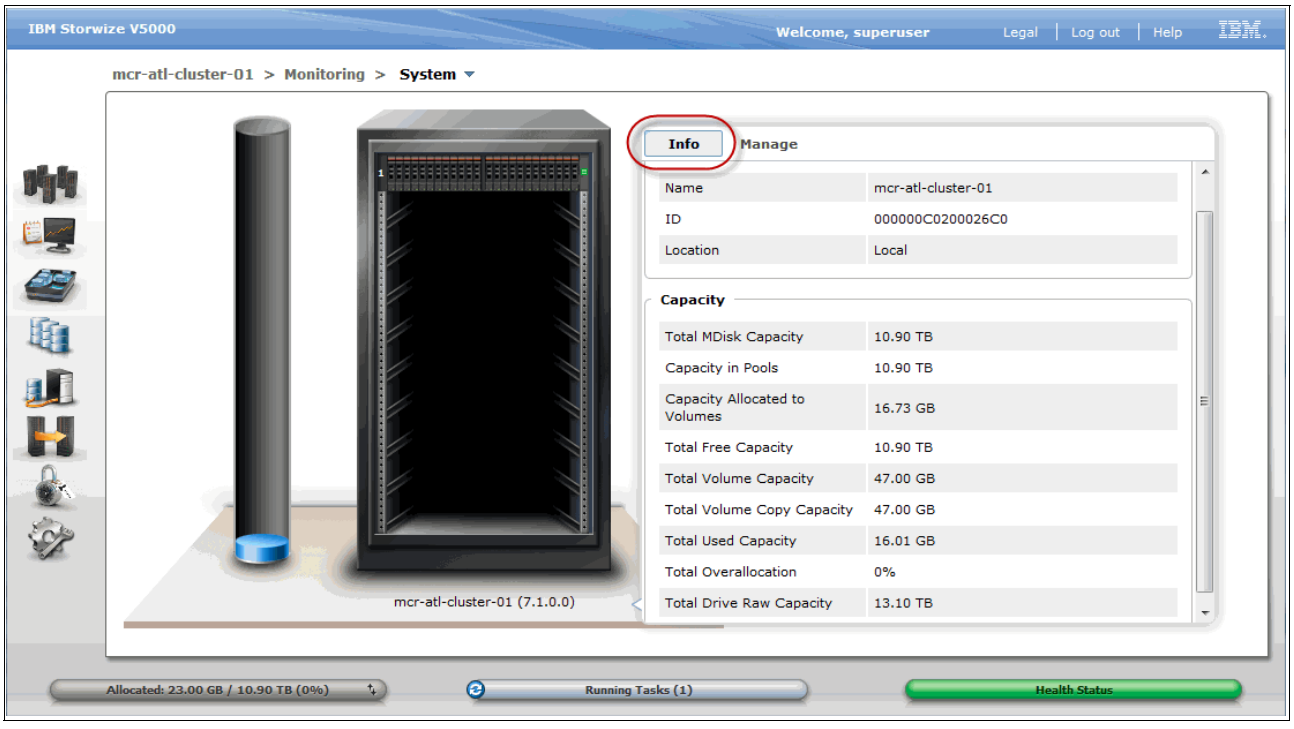

Figure 3-39 System panel Info tab

Select the **Manage** tab to show the name of the system and shutdown and upgrade actions, as shown in Figure 3-40.

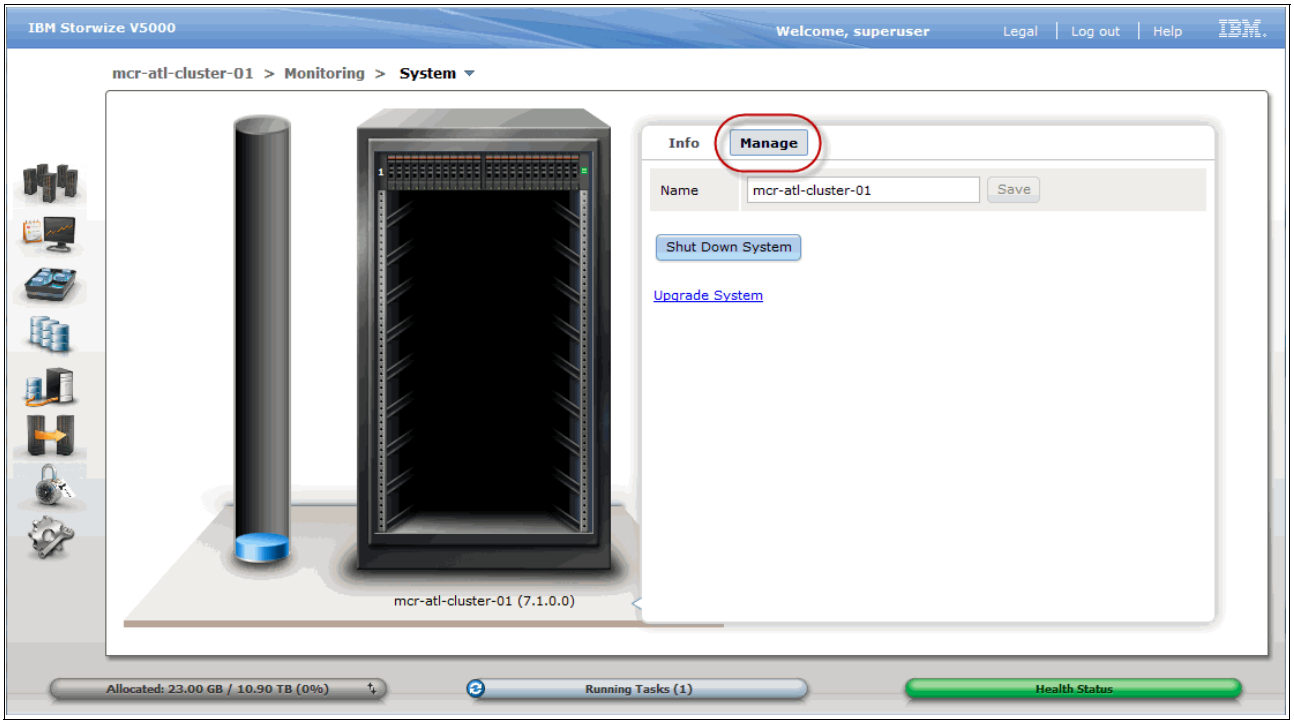

Figure 3-40 System panel Manage tab

Selecting a rack-mounted enclosure shows more information. Hovering over a drive shows the drive status, size, and speed details. Identify starts the blue identification LED on the front of the enclosure. Click **Enclosure 1** to show the System Details panel. For more information, see “System Details panel” on page 103. Figure 3-41 on page 103 shows the System panel enclosure view.

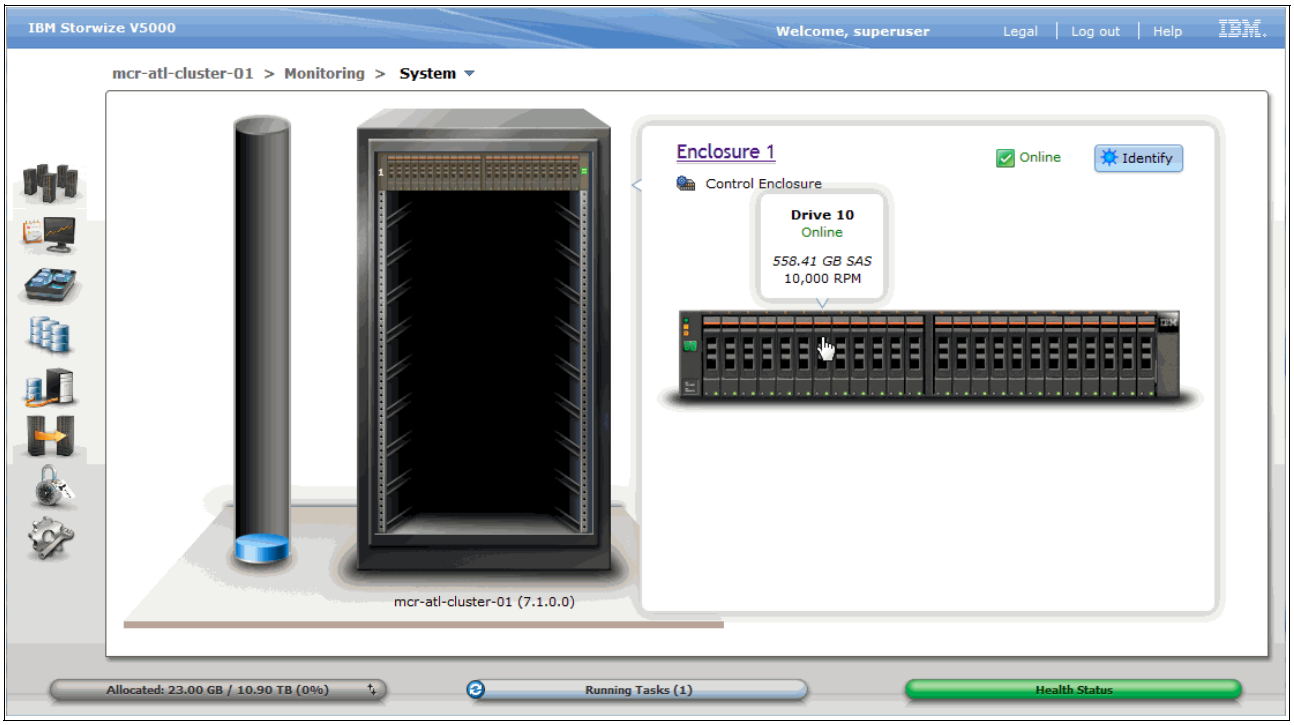

Figure 3-41 System panel enclosure view

System Details panel

Select **System Details** in the Monitoring menu to open the panel. As shown in Figure 3-42, the System Details panel provides the status and details of the components that make up the system.

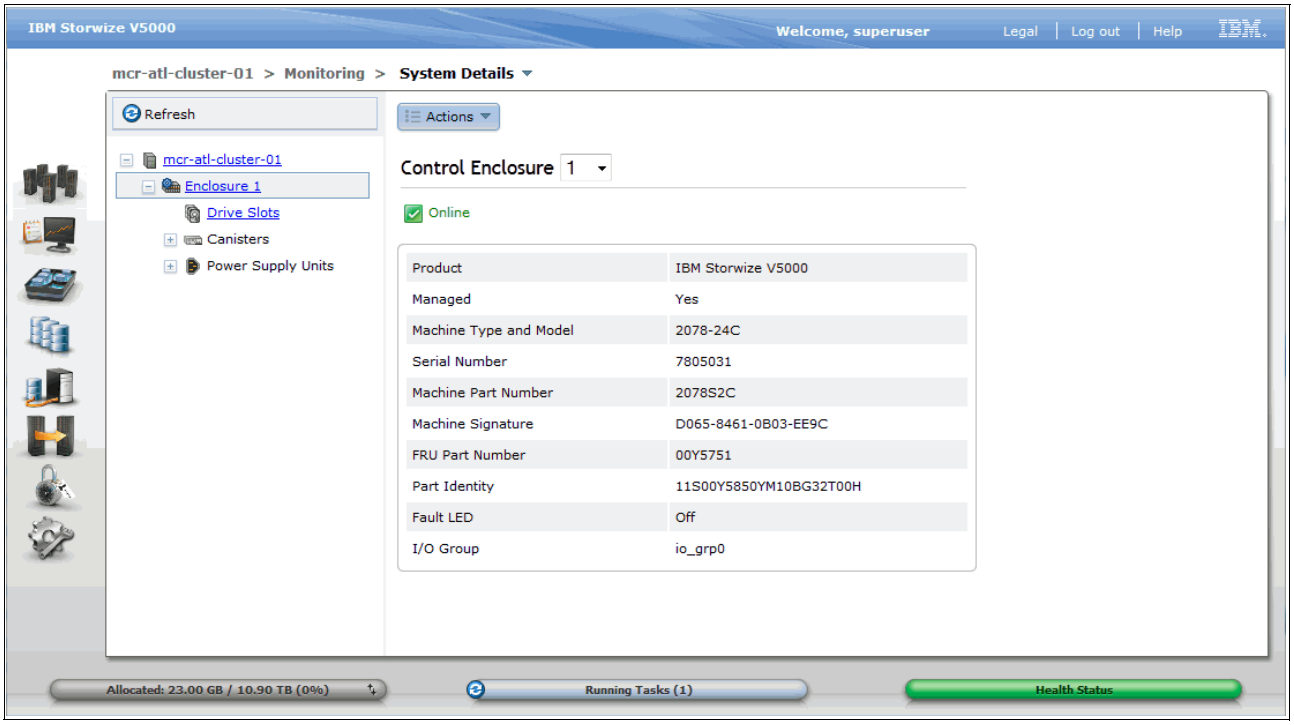

Figure 3-42 System Details panel

Actions and environmental statistics

Actions, such as, adding expansion enclosures, viewing the SAS chain connections, and performing a software upgrade and a system shutdown, can be run from the System Details panel. Information that relates to environmental statistics, such as, power consumption and temperature, is also accessible from this panel. Figure 3-43 shows the available actions for and the environmental statistics of the enclosure.

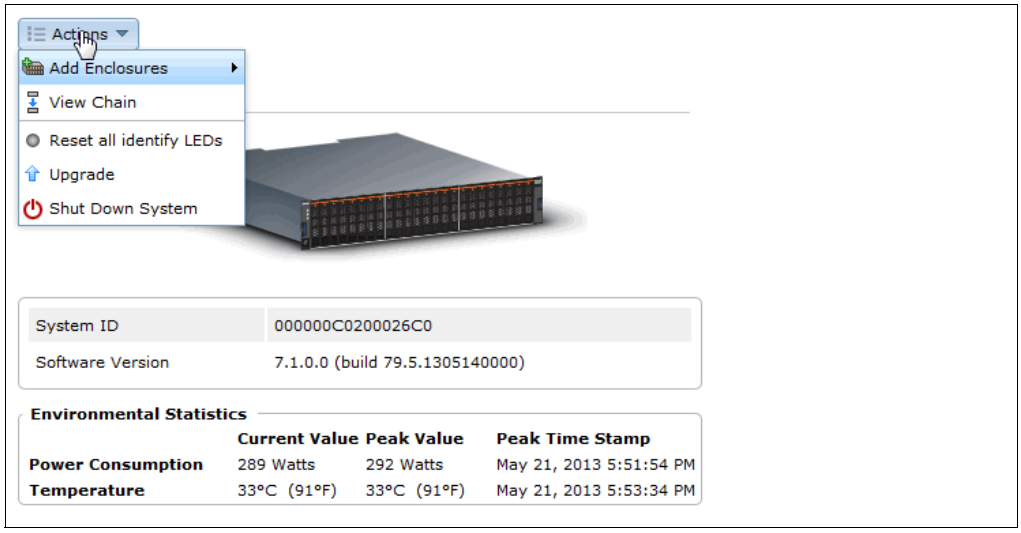

Figure 3-43 System details actions and environmental statistics

Node canister information

Node canister information, such as, FC and SAS WWPNs and iSCSI IQNs, is useful for host attachment purposes. This information is shown by clicking the control enclosure node canister in the System Details panel. Figure 3-44 shows node canister information.

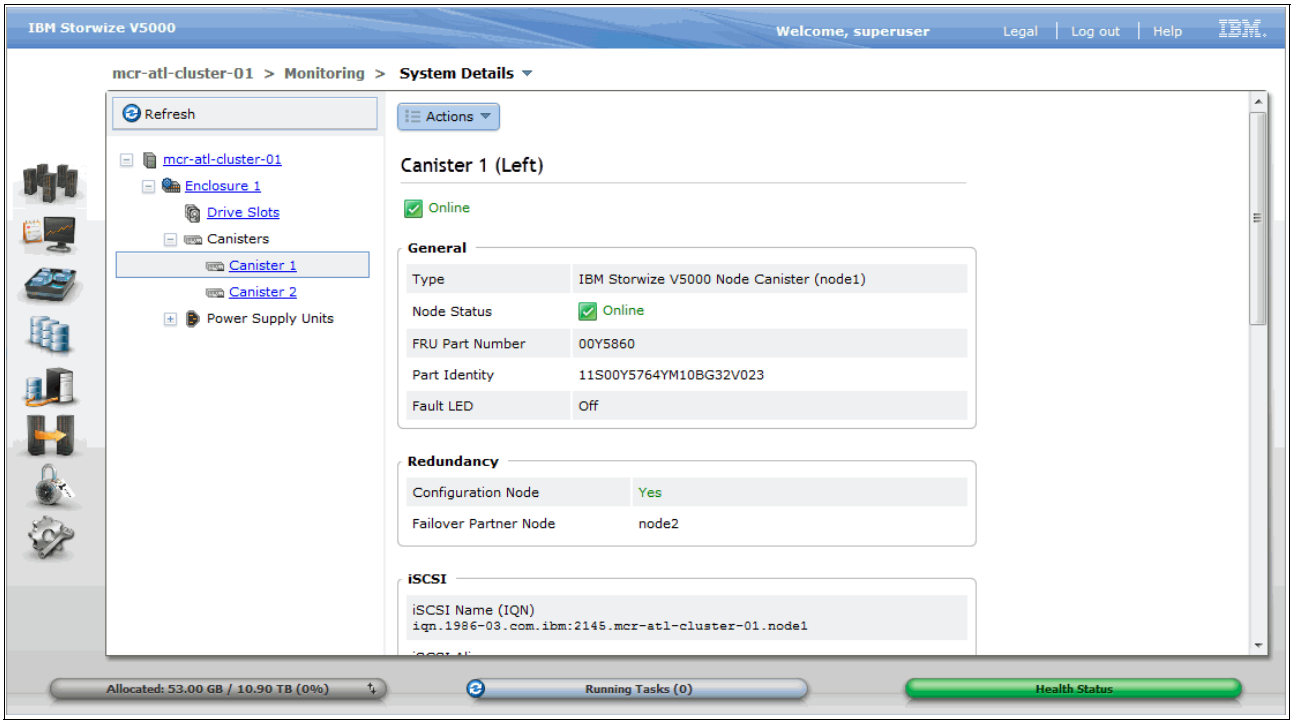

Figure 3-44 Node canister information via system details panel

Events panel

Select **Events** in the Monitoring menu to open the Events panel. The machine is optimal when all errors are addressed and no items are found in this panel, as shown in Figure 3-45.

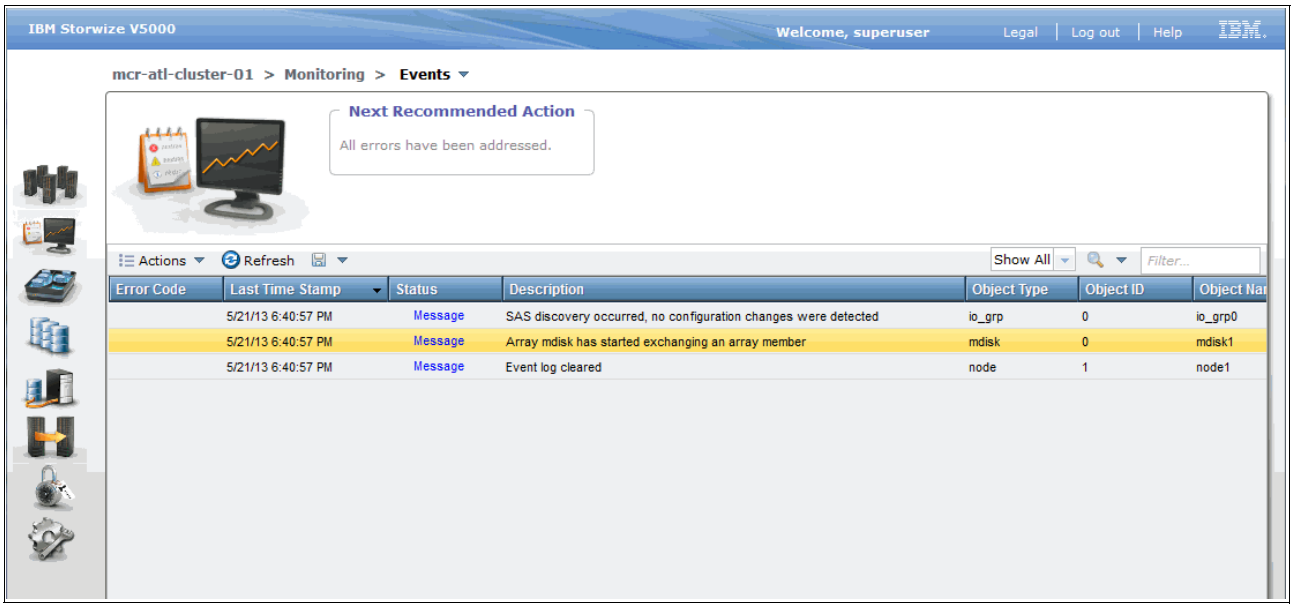

Figure 3-45 Events panel with all errors addressed

Filtering events view

To view Unfixed Messages and Alerts or to Show All, select the appropriate option from the menu that is next to the filter field, as shown in Figure 3-46. For more information, see “Filtering objects” on page 92.

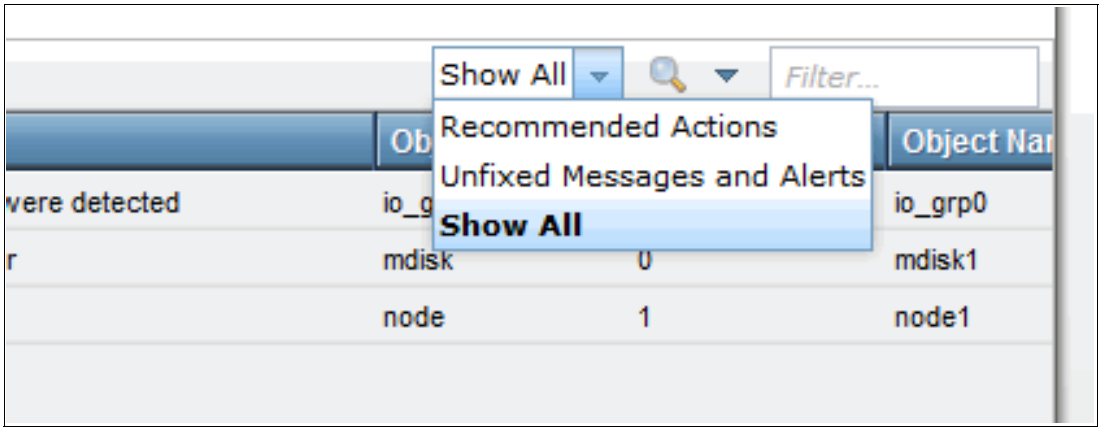

Figure 3-46 Unfixed messages and alerts in the events panel

Event properties

To show actions and properties that are related to an event or to repair an event that is not the Next Recommended Action, right-click the event to show other options. Figure 3-47 on page 106 shows the selection of the **Properties** option.

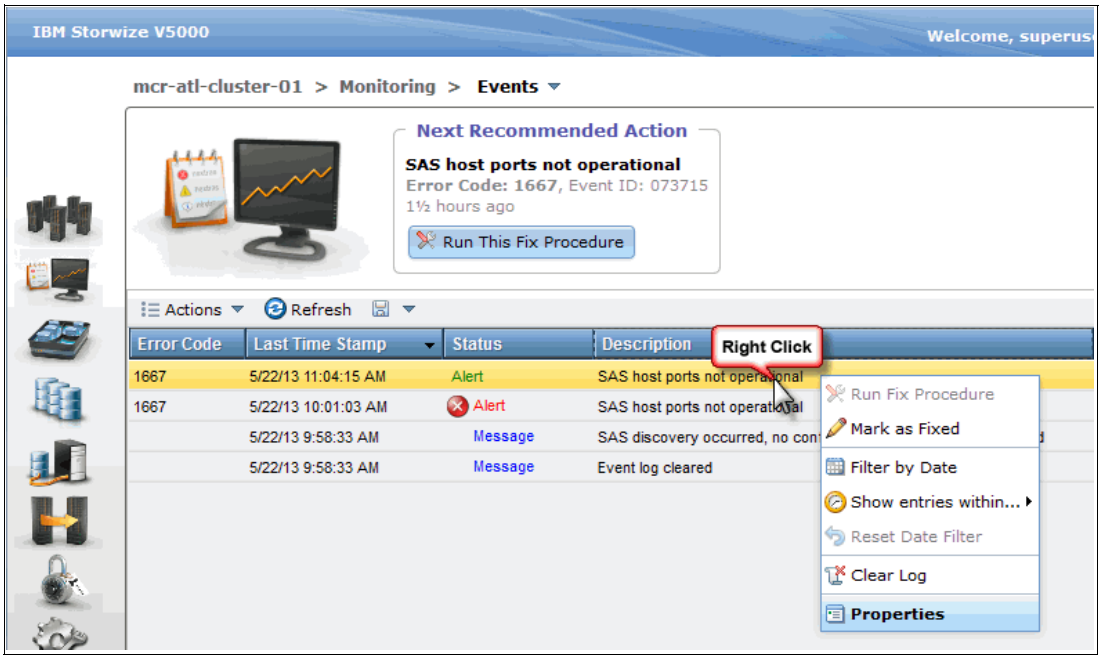

Figure 3-47 Selecting event properties

Figure 3-48 shows the properties and sense data for an event.

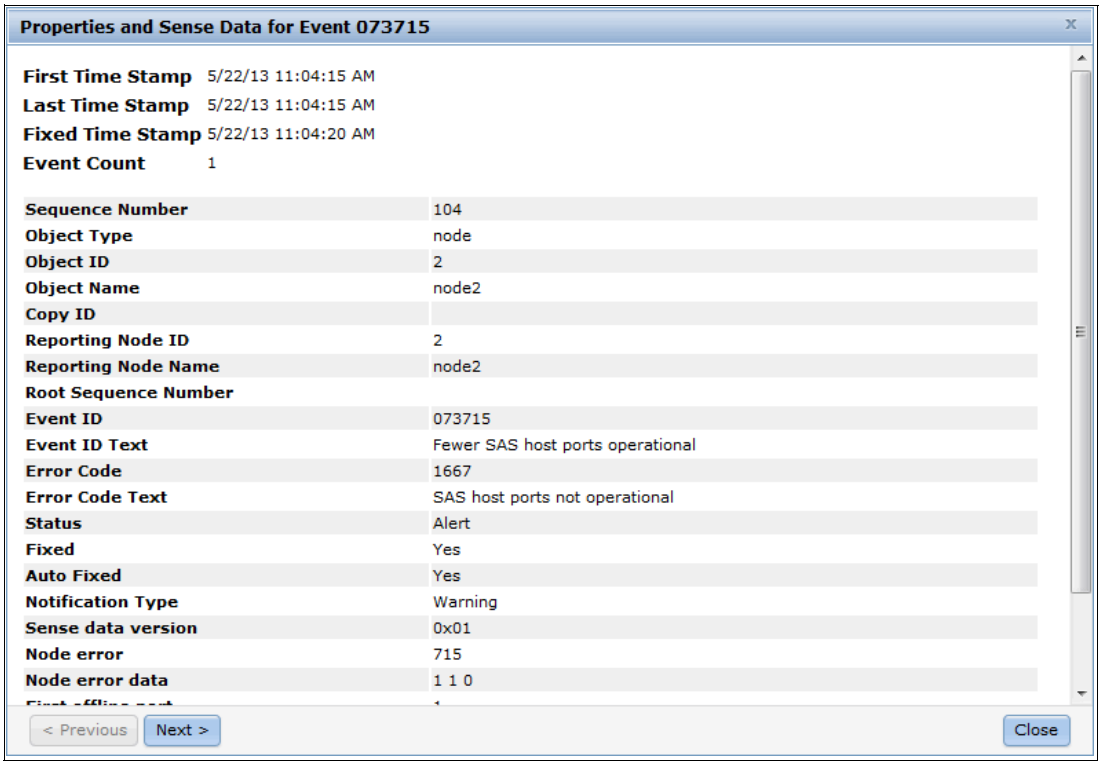

Figure 3-48 Event properties and sense data

Show events entries within

To show events that occurred within a certain time of a particular event, select the required event entry, then select **Show entries within...** from the Actions menu and set the period value. Figure 3-49 shows the selection of the Show entries within... option with a period value of 5 minutes. This shows all events within 5 minutes of the selected event.

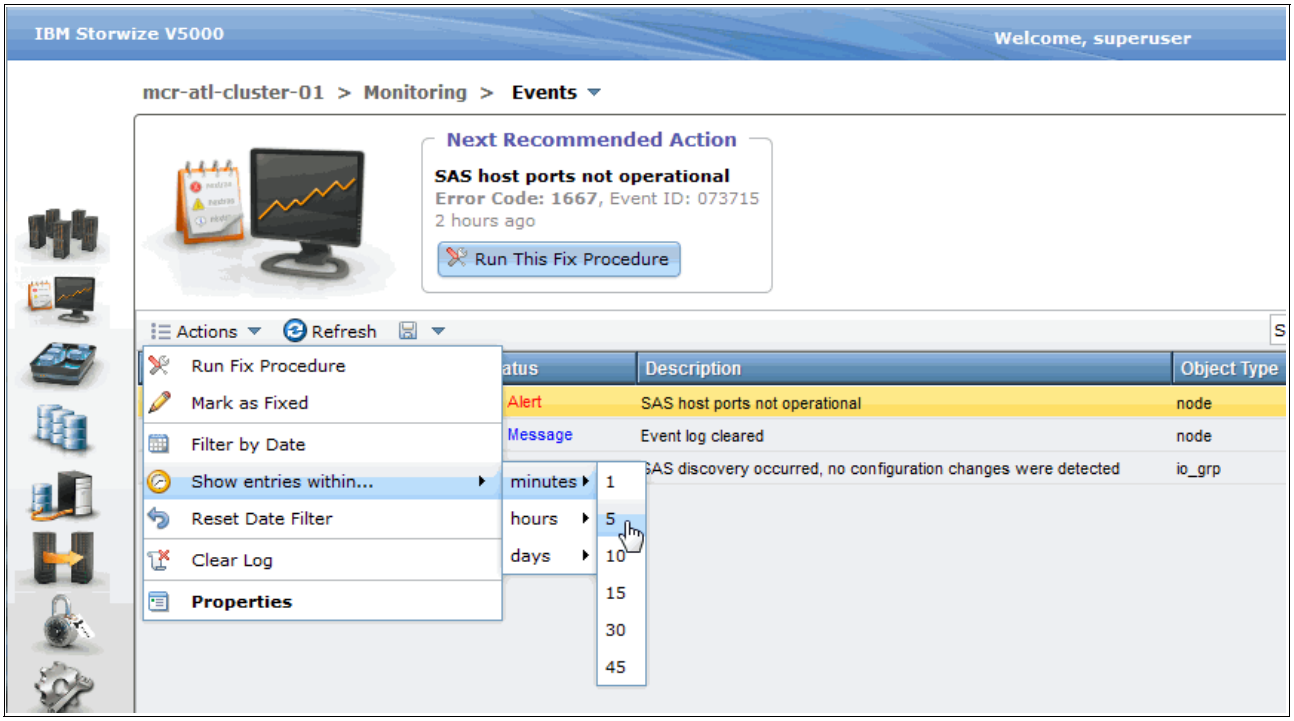

Figure 3-49 Showing events within a set time

Saving events to a file

It is possible to save the events that are listed in the events panel to a file. To do this, click the diskette icon and select the format that you require to save the file. A comma-delimited file is created that can be saved in text format or as a `.csv` file for input to a spreadsheet program, such as, Microsoft Excel.

Figure 3-50 on page 108 shows saving the events as formatted values.

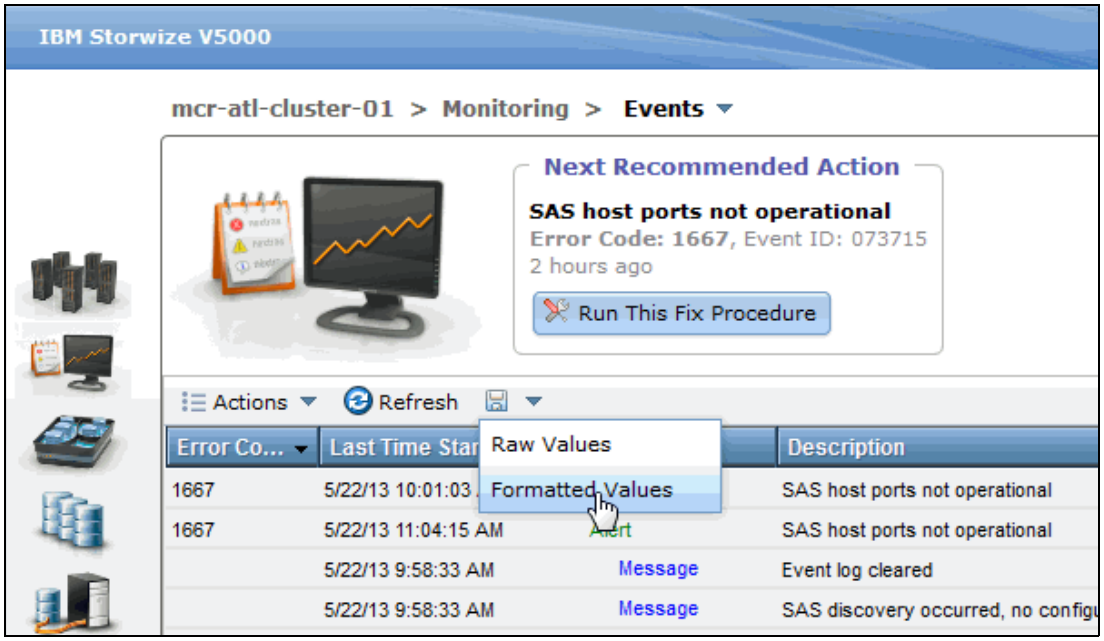

Figure 3-50 Saving events as formatted values

Performance panel

Select **Performance** in the Monitoring menu to open the Performance panel. This panel shows graphs that represent the last 5 minutes of performance statistics. The performance graphs include statistics about CPU Utilization, Volumes, Interfaces, and MDisks. Figure 3-51 shows the Performance panel.

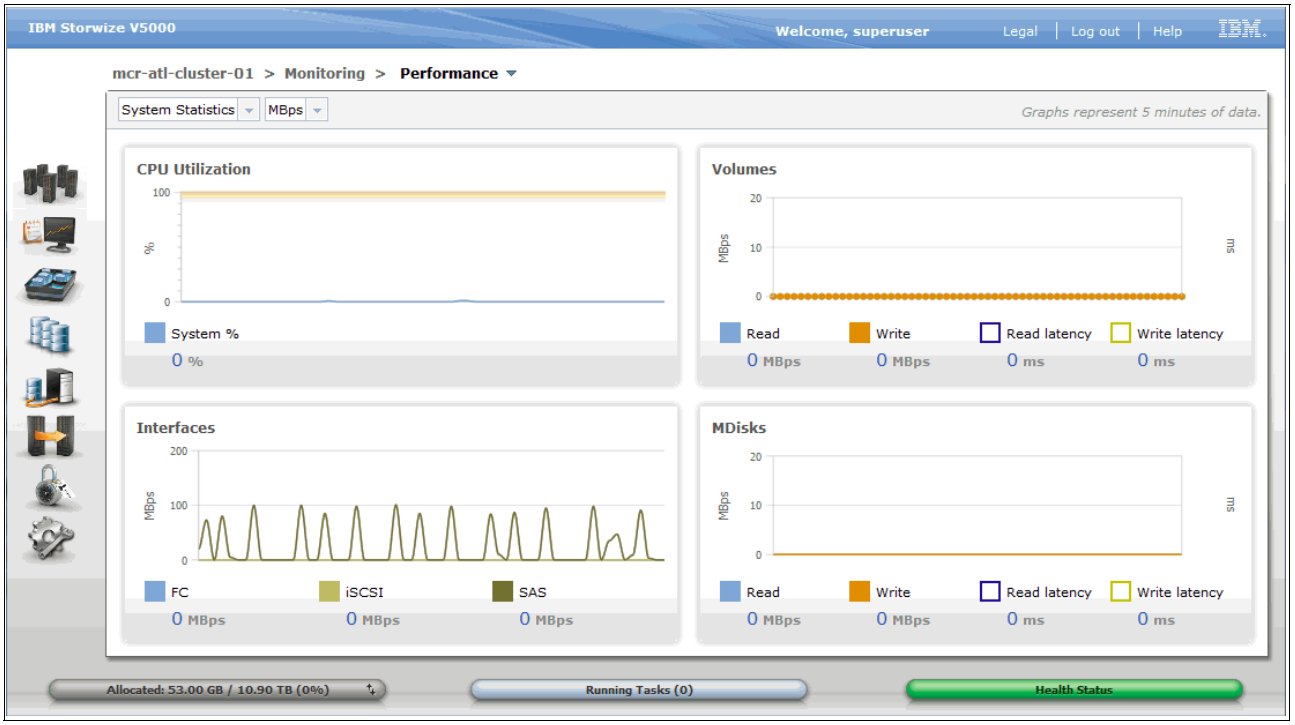

Figure 3-51 Performance panel

Custom-tailoring performance graphs

The Performance panel can be customized to show the workload of a single node, which is useful to help determine whether the system is working in a balanced manner. Figure 3-52 shows the custom-tailoring of the performance graphs by selecting node 1 from the System Statistics menu. The measurement type can also be changed between throughput (MBps) or IOPS by selecting the relevant value.

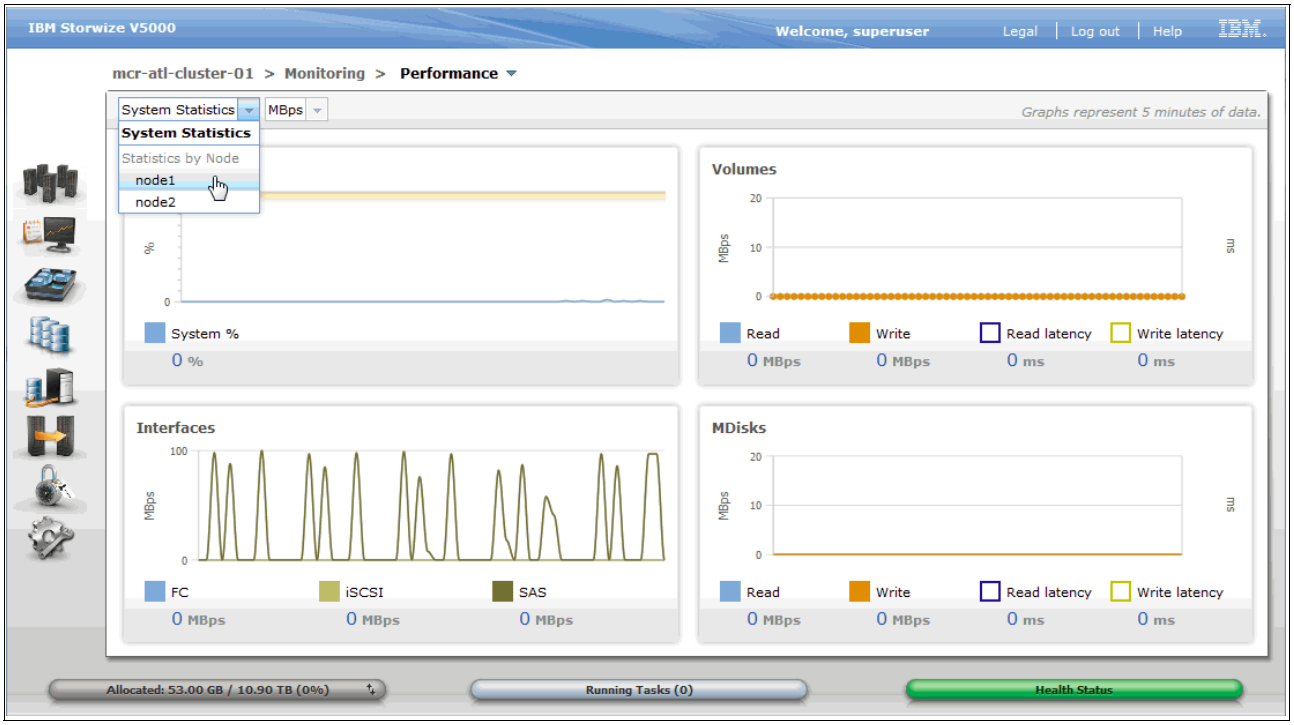

Figure 3-52 Graphs representing performance statistics of a single node

Performance peak value

Peak values over the last 5-minute period can be seen by hovering over the current value, as shown in Figure 3-53 on page 110 for the SAS Interfaces.

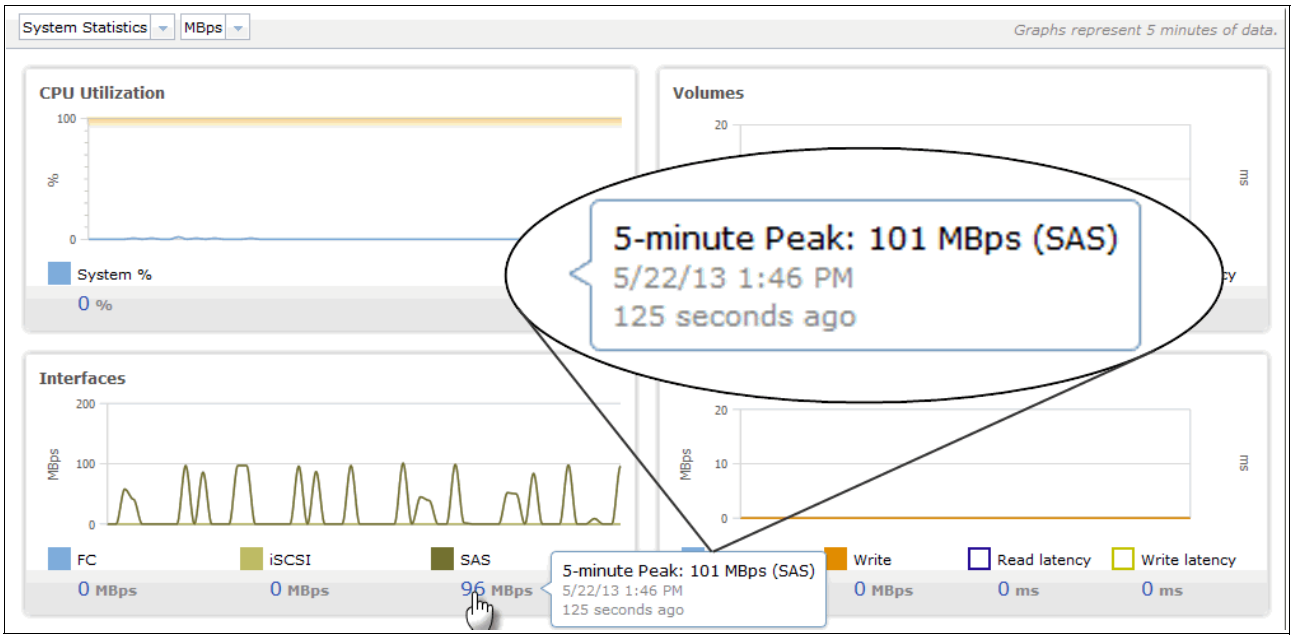

Figure 3-53 Peak SAS Interface usage value over the last 5 minutes

3.4.3 Pools menu

The Pools menu provides access to the Volumes by Pools, Internal Storage, MDisks by Pools, and System Migration functions, as shown in Figure 3-54.

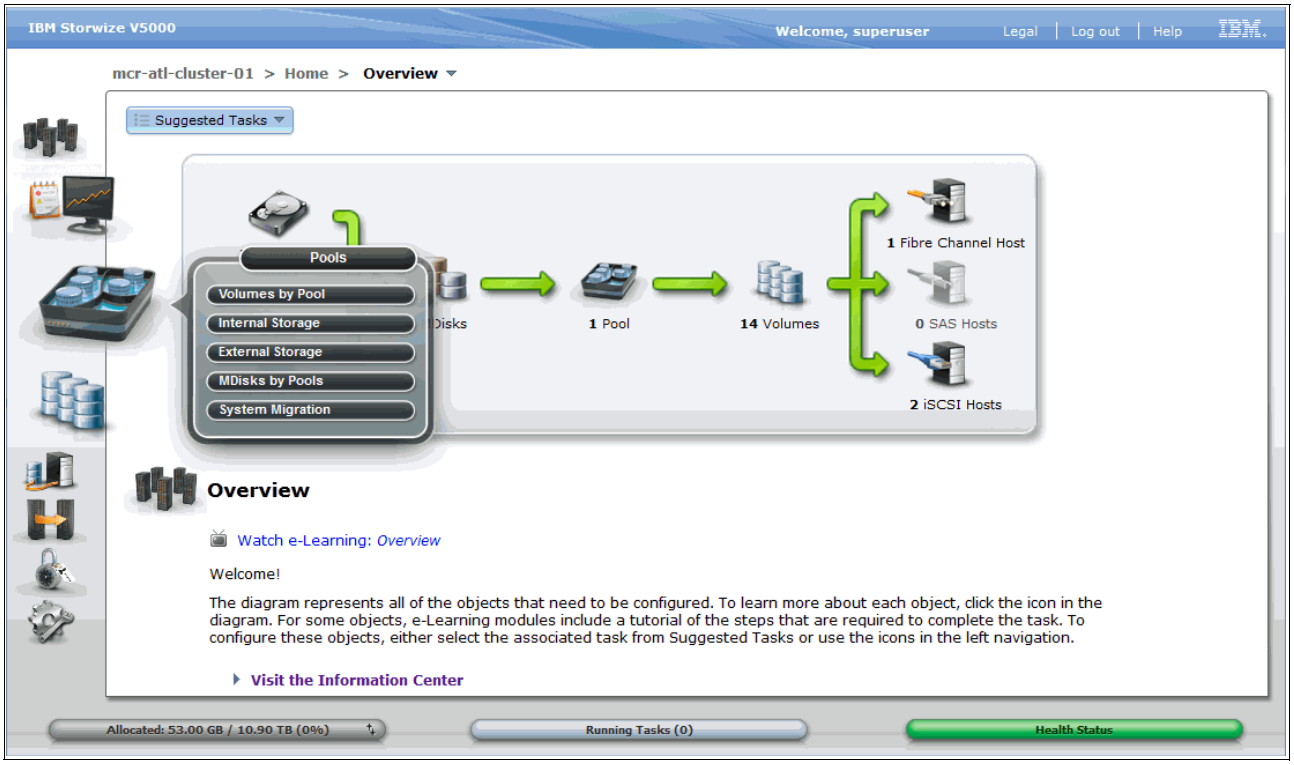

Figure 3-54 Pools menu

Volumes by Pool panel

Select **Volumes by Pool** in the Pools menu to open the panel. By using the Volumes by Pool panel, you can display volumes by using the Pool Filter function. This view makes it easy to manage volumes and determine the amount of real capacity that is available for more allocations. Figure 3-55 shows the Volumes by Pool panel.

Figure 3-55 Volumes by Pools panel

Volume Allocation

The upper right corner of the Volumes by Pool panel shows the Volume Allocation, which, in this example, shows the physical capacity (3.81 TB), the virtual capacity (5.10 TB), and the used capacity (204.00 GB in the green portion). The red bar shows the threshold at which a warning is generated when the used capacity in the pool first exceeds the threshold that is set for the physical capacity of the pool. By default, this threshold is set to 80% but can be altered in the pool properties. Figure 3-56 shows the volume allocation information that is displayed in the Volumes by Pool panel.

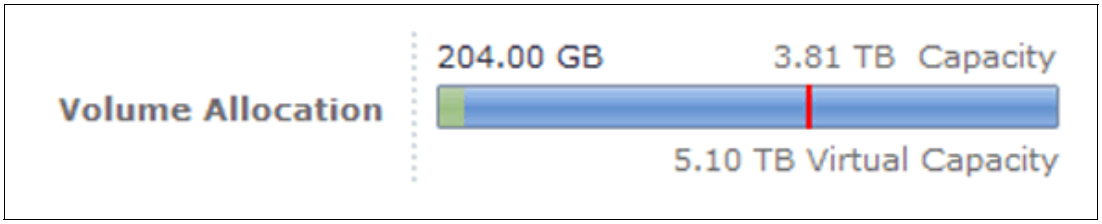

Figure 3-56 Volume Allocation

Renaming pools

To rename a pool, select the pool from the pool filter and click the name of the pool. Figure 3-57 shows that pool `V5000_Pool_1` was selected to be renamed.

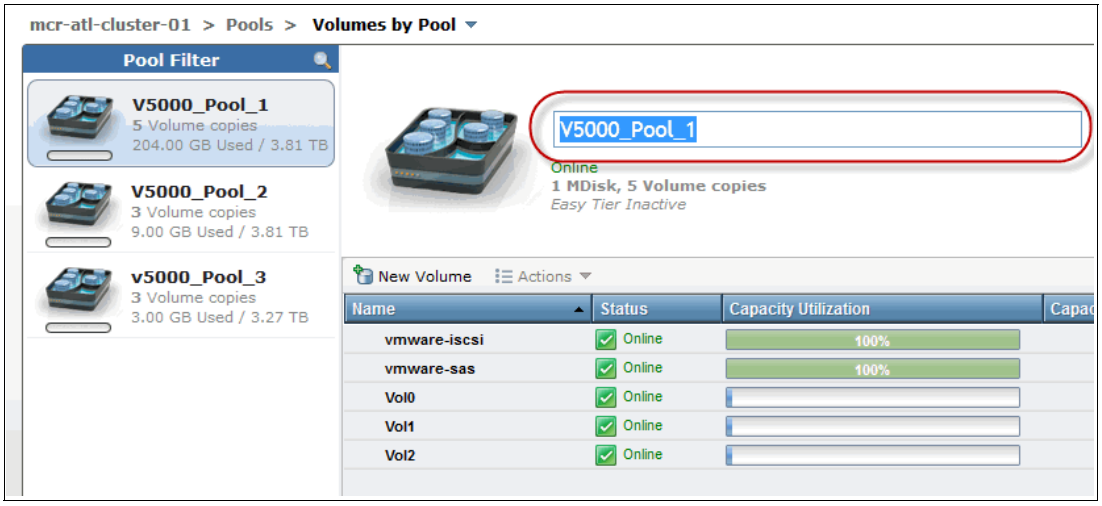

Figure 3-57 Renaming a pool

Changing pool icons

To change the icon that is associated with a pool, select the pool in the pool filter, click the large pool icon that is above New Volume and Actions, then use the Choose Icon buttons to select the wanted image. This change helps to manage and differentiate between the classes of drive or the tier of the storage pool. Figure 3-58 shows the pool change icon panel.

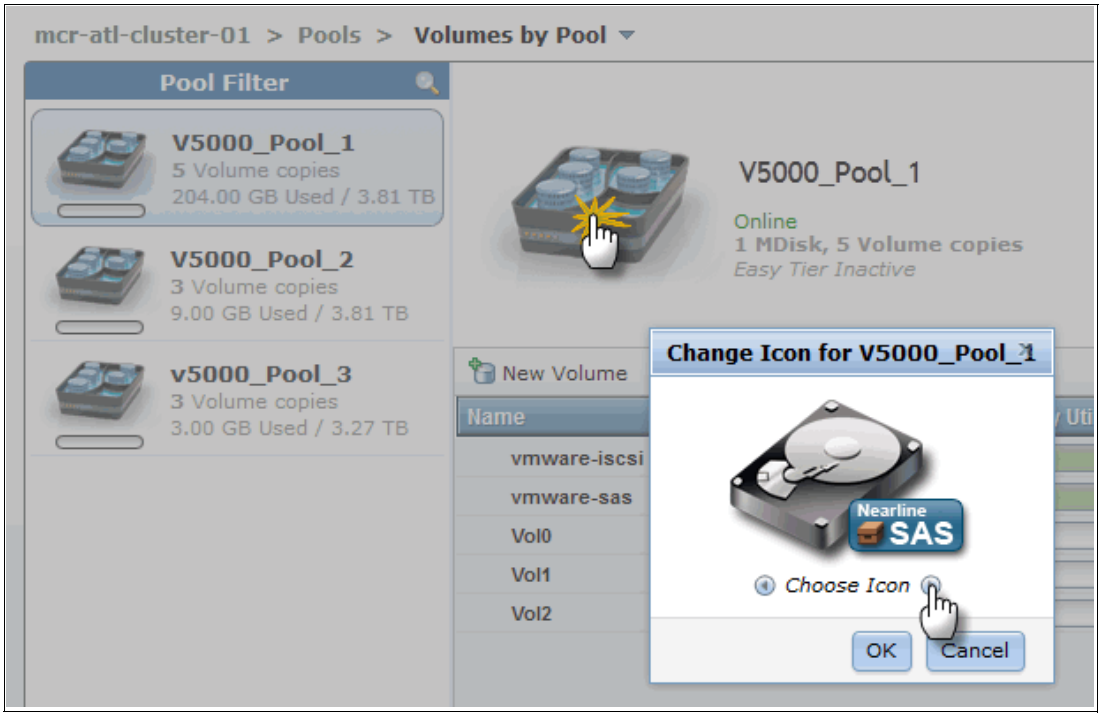

Figure 3-58 Changing a pool icon

Volume functions

The Volumes by Pool panel also provides access to the volume functions via the Actions menu, the New Volume option, and by right-clicking a listed volume. For more information about navigating the Volume panel, see 3.4.4, “Volumes menu” on page 121. Figure 3-59 shows the volume functions that are available via the Volumes by Pool panel.

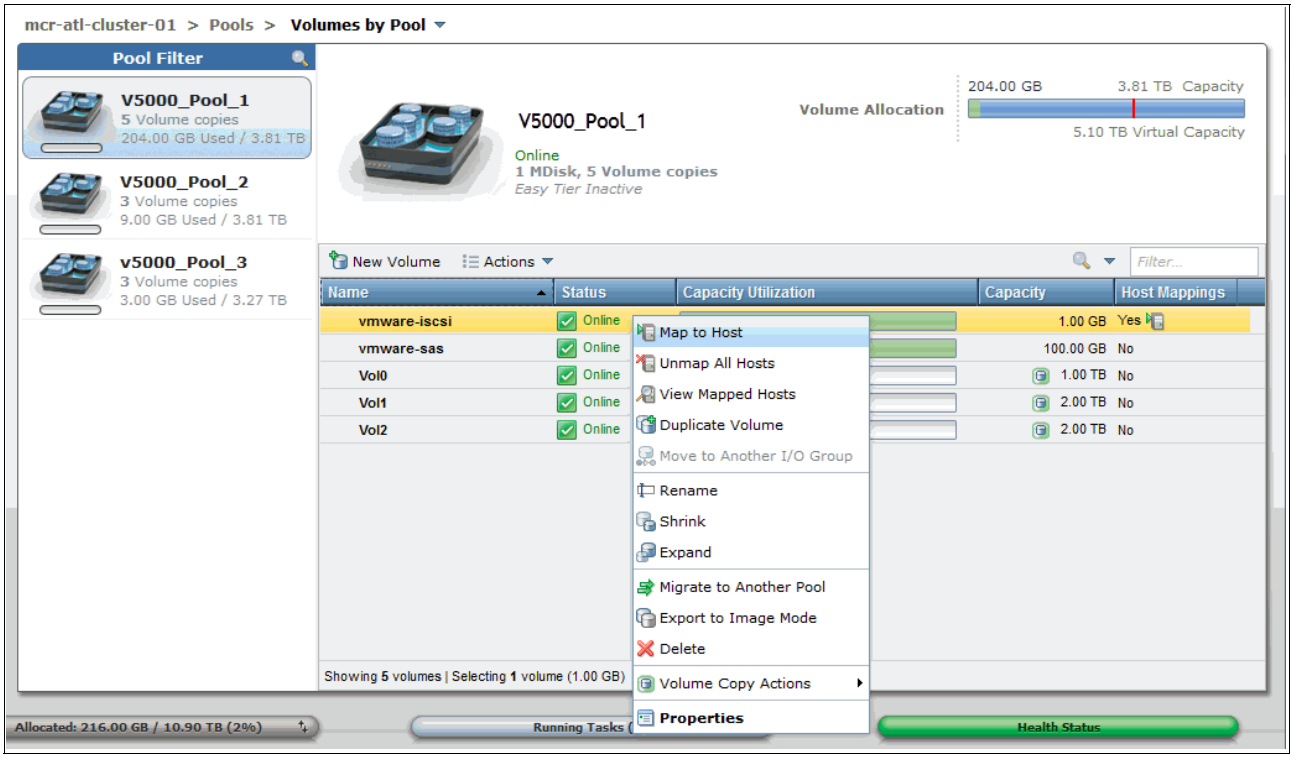

Figure 3-59 Volume functions are available via the Volume by Pools panel

Internal Storage panel

Select **Internal Storage** in the Pools menu to open the Internal Storage panel, as shown in Figure 3-60. The internal storage consists of the drives that are contained in the IBM Storwize V5000 control enclosure and any SAS-attached IBM Storwize V5000 expansion enclosures. By using the Internal Storage panel, you can configure the internal storage into RAID protected storage (MDisks). You can also filter the displayed drive list by drive class.

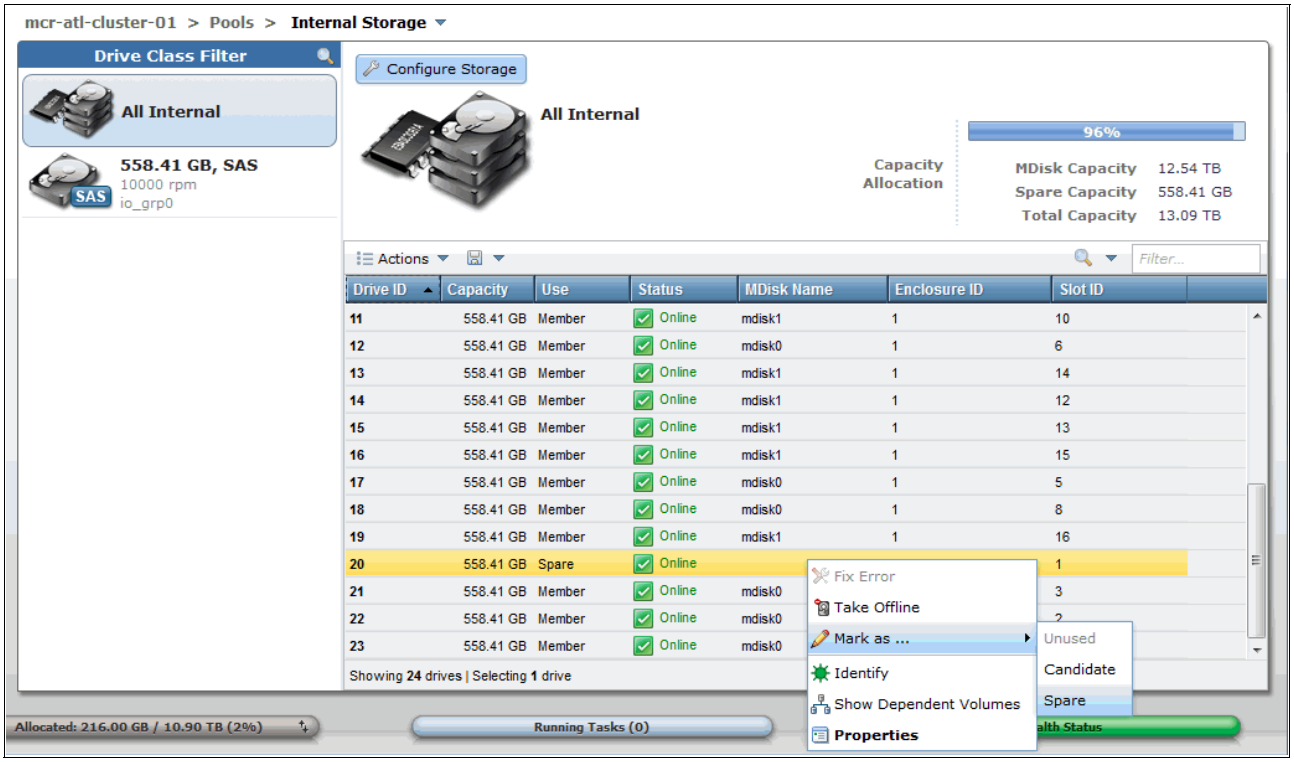

Figure 3-60 Drive actions menu of the internal storage panel

Drive actions

Drive level functions, such as, identifying a drive and marking a drive as offline, unused, candidate, or spare, can be accessed here. Right-click a listed drive to show the Actions menu. Alternatively, the drives can be selected and then the Action menu is used. For more information, see “Multiple selections” on page 91. Figure 3-60 shows the Drive Actions menu.

Drive properties

Drive properties and dependent volumes can be displayed from the Internal Storage panel. Select **Properties** from the Drive Actions menu. The drive Properties panel shows the drive attributes and the drive slot SAS port status. Figure 3-61 on page 115 shows the drive properties with the Show Details option selected.

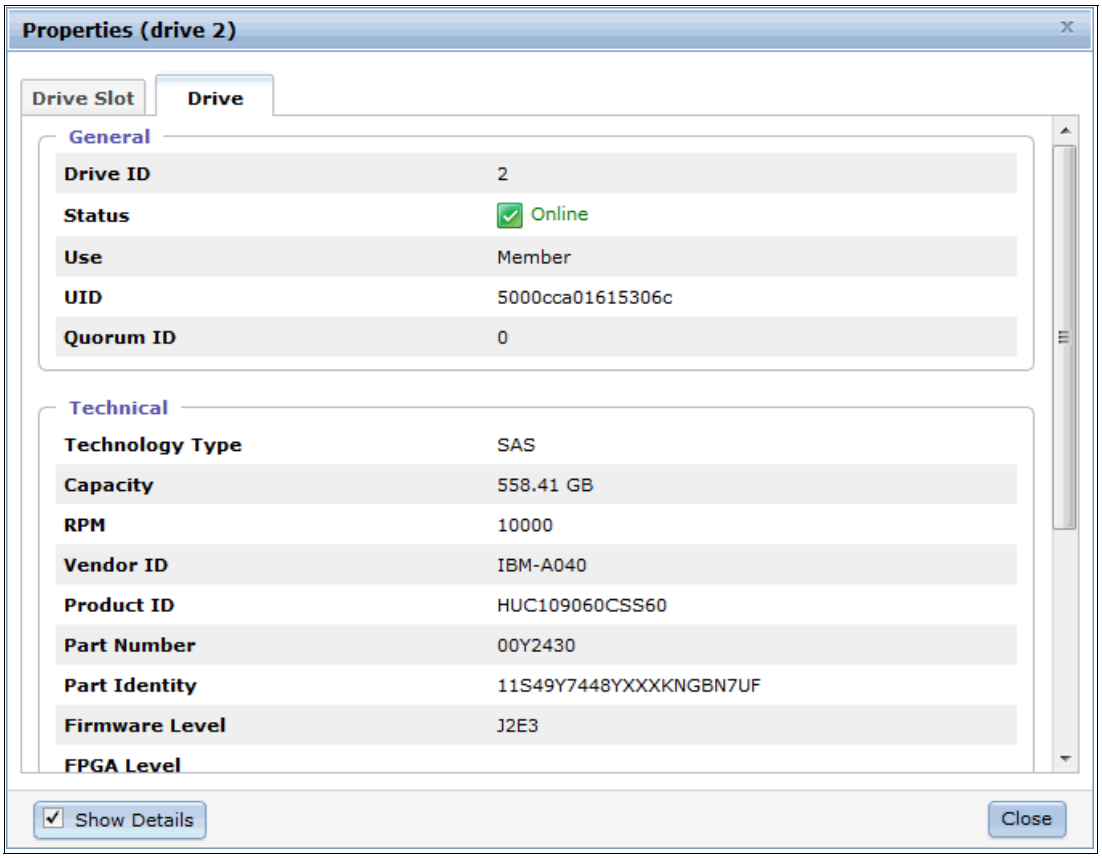

Figure 3-61 Drive properties

Configure internal storage wizard

Click **Configure Storage** to show the Configure Internal Storage wizard, as shown in Figure 3-62.

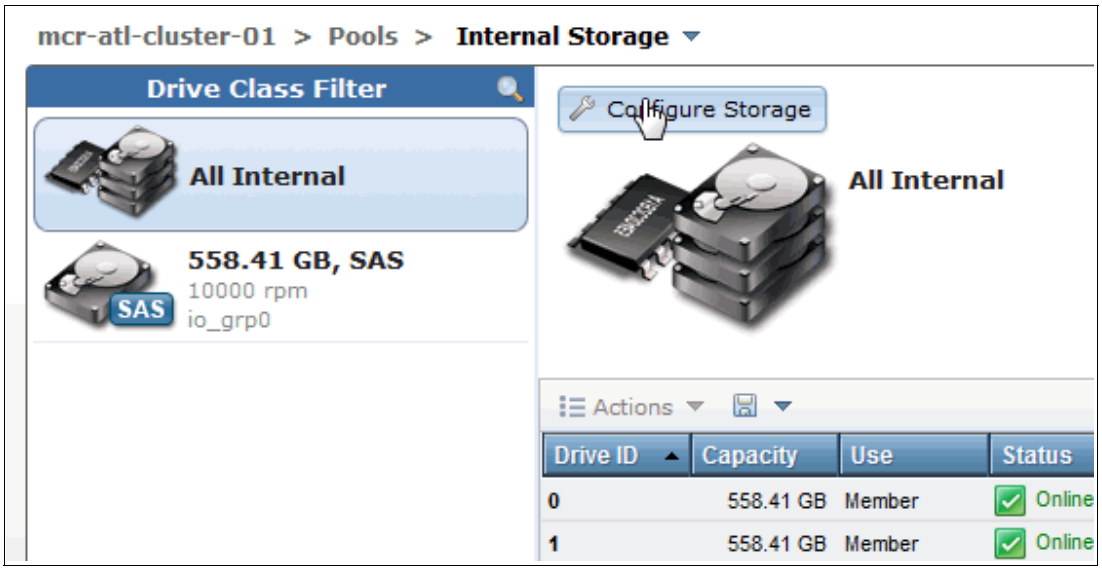

Figure 3-62 Internal Storage panel

By using this wizard, you can configure the RAID properties and pool allocation of the internal storage. Figure 3-63 shows Step 1 of the Configure Internal Storage wizard.

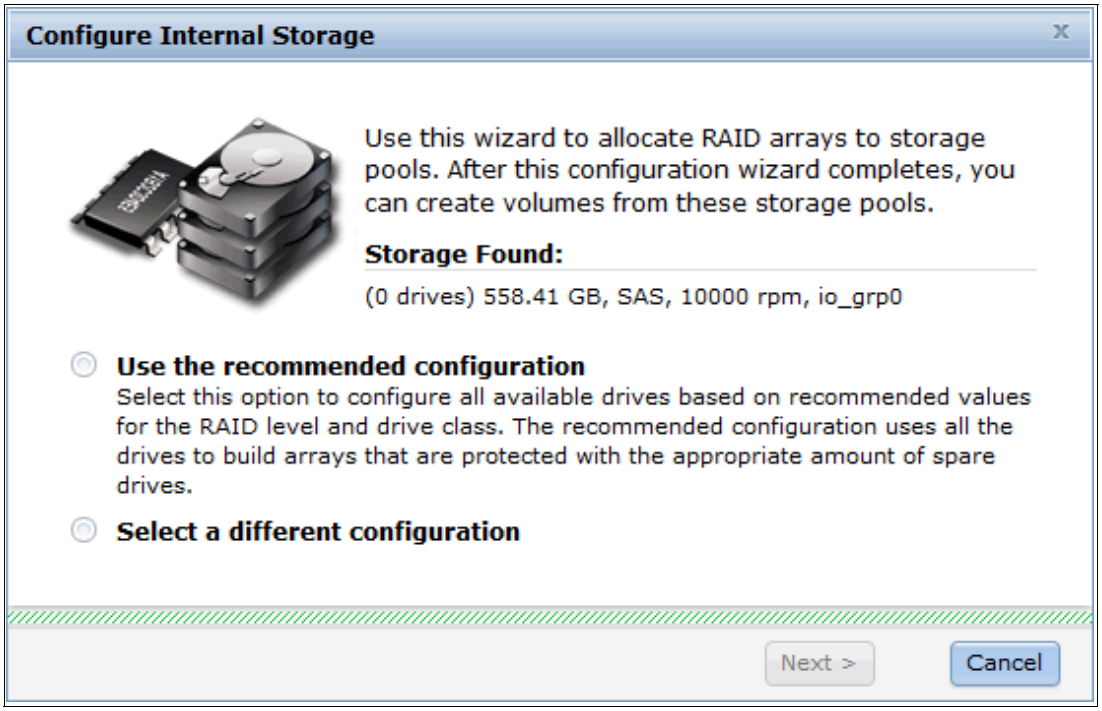

Figure 3-63 Configure Internal Storage wizard: Step 1

Figure 3-64 shows Step 2 of the Configure Internal Storage wizard.

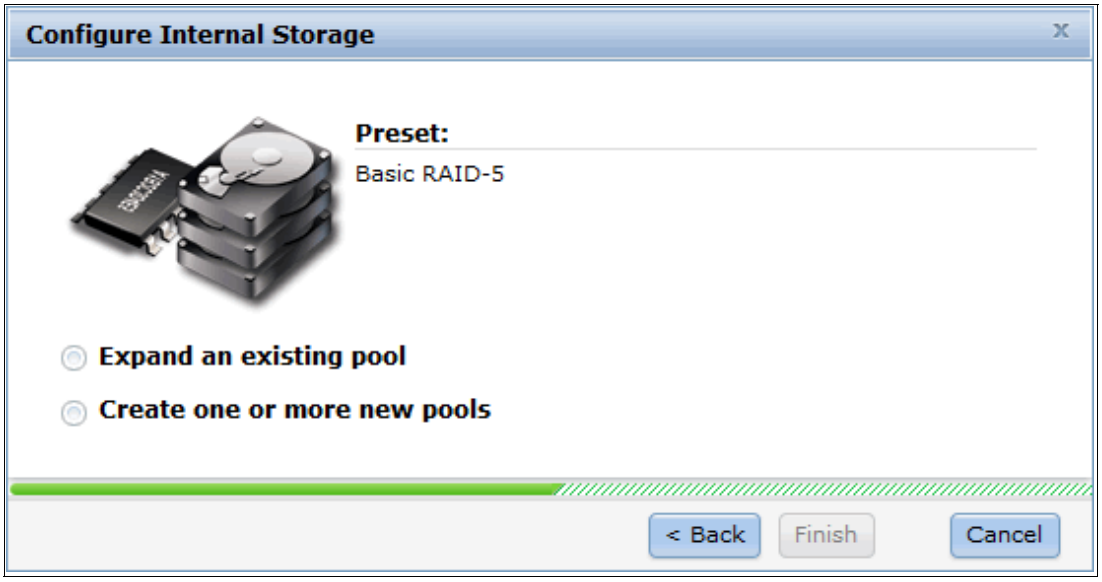

Figure 3-64 Configuring Internal Storage wizard: Step 2

MDisks by Pool panel

Select **MDisks by Pool** in the Pools menu to open the MDisks by Pool panel. By using this panel, you can perform such tasks such as, display MDisks in each pool, create pools, delete pools, and detect externally virtualized storage. Figure 3-65 shows the MDisks by Pool panel.

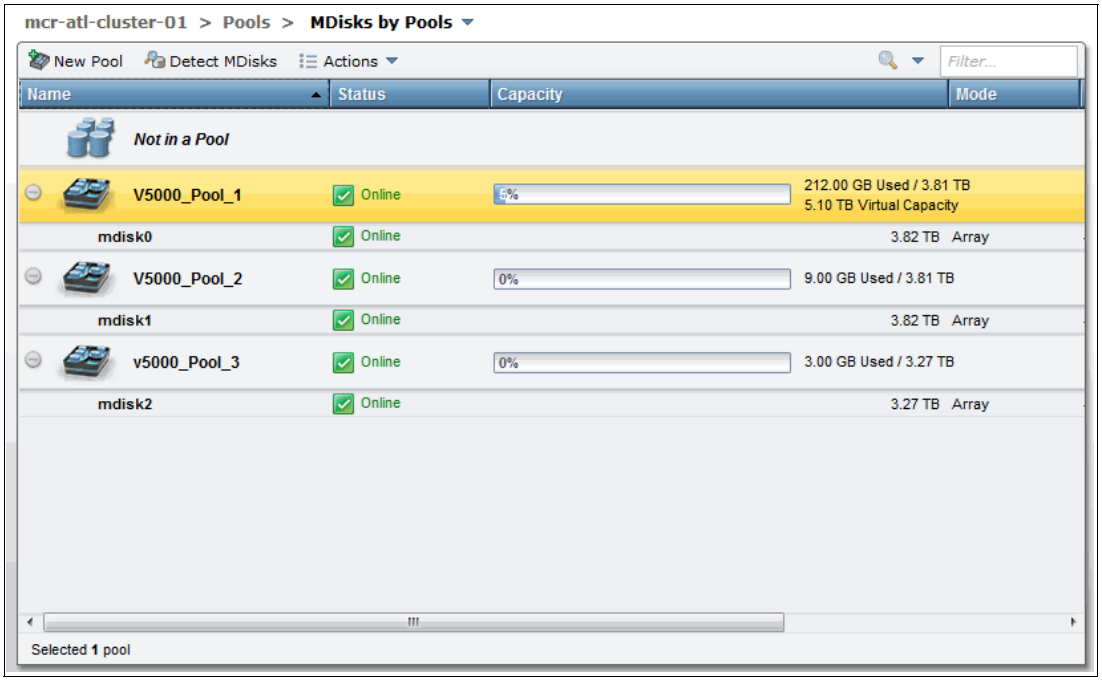

Figure 3-65 MDisks by Pool panel

Pool actions

To delete a pool or change the pool name or icon, right-click the listed pool. Alternatively, the Actions menu can be used. Figure 3-66 shows the pool actions.

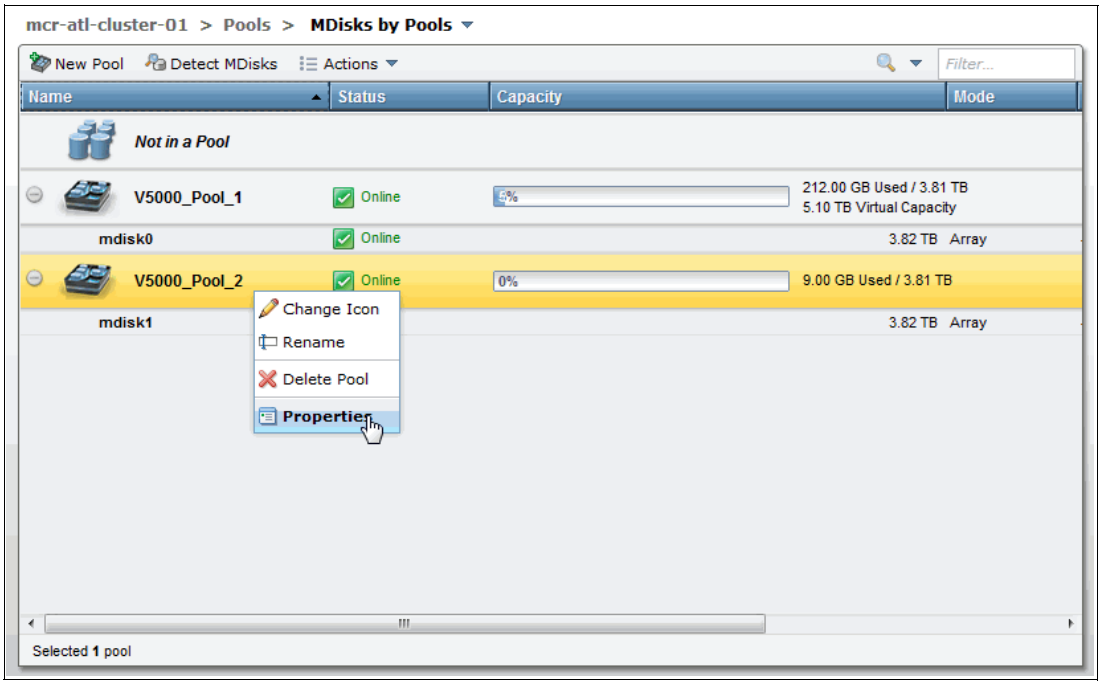

Figure 3-66 Pool actions

RAID actions

By using the MDisks by Pool panel, you can perform MDisk RAID tasks, such as, Set Spare Goal, Swap Drive, and Delete. To access these functions, right-click the MDisk, as shown in Figure 3-67.

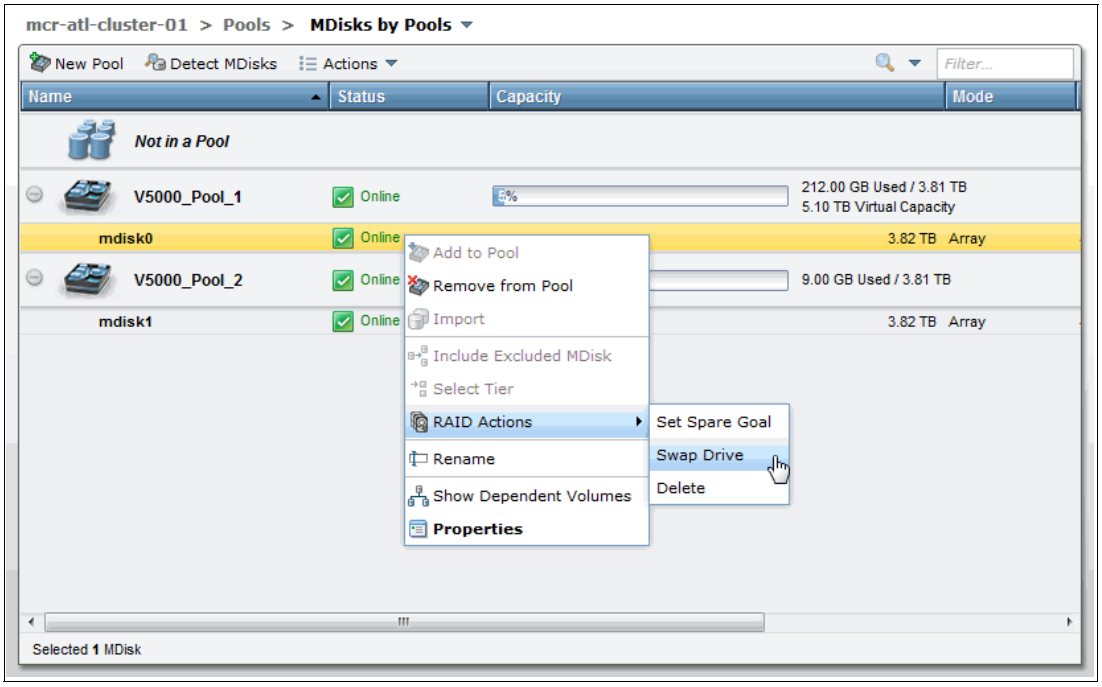

Figure 3-67 RAID actions menu

System Migration panel

Select **System Migration** in the Pools menu to open the System Migration panel, as shown in Figure 3-68. This panel is used to migrate data from externally virtualized storage systems to the internal storage of the IBM Storwize V5000. The panel displays image mode volume information. To begin a migration, click **Start New Migration** and the Start Migration wizard is shown.

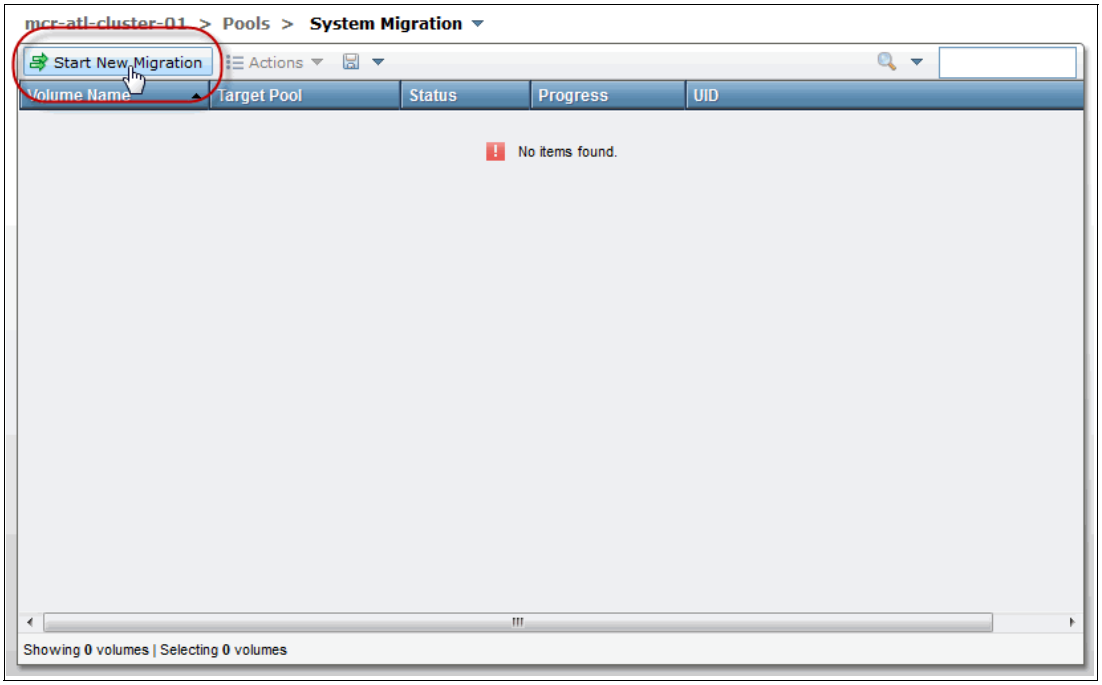

Figure 3-68 System Migration panel

Storage Migration wizard

The Storage Migration wizard is used for data migration from other Fibre Channel-attached storage systems to the IBM Storwize V5000. Figure 3-69 shows Step 1 of the Storage Migration wizard.

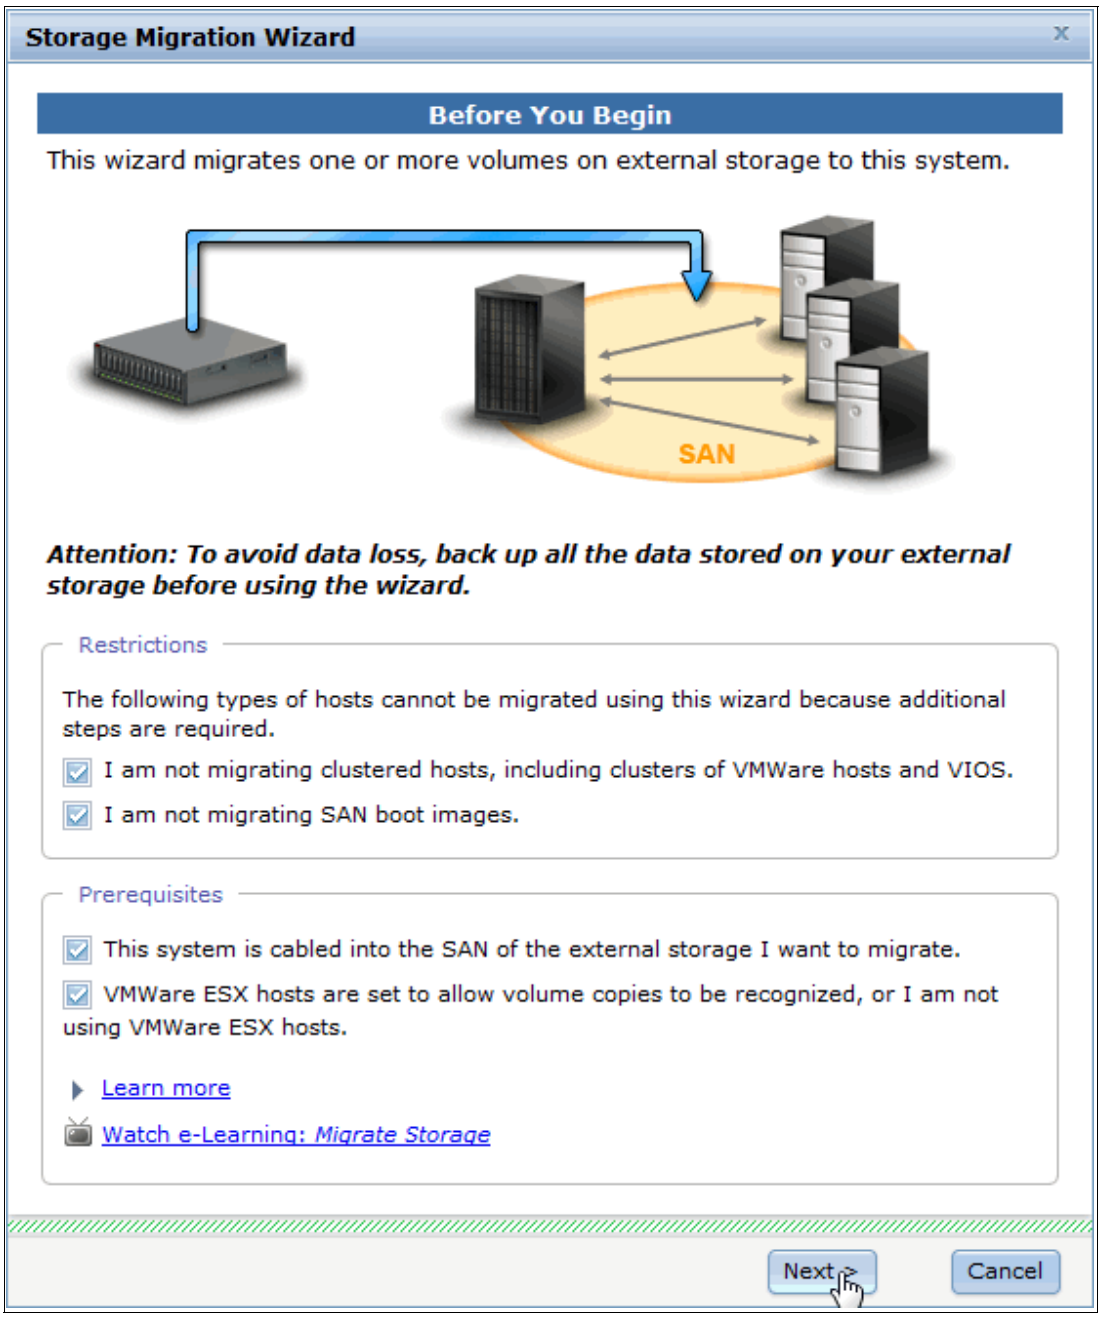

Figure 3-69 Storage Migration wizard

For more information, see Chapter 6, “Storage migration wizard” on page 237.

3.4.4 Volumes menu

As shown in Figure 3-70, the Volumes menu provides access to the Volumes, Volumes by Pool, and Volumes by host functions.

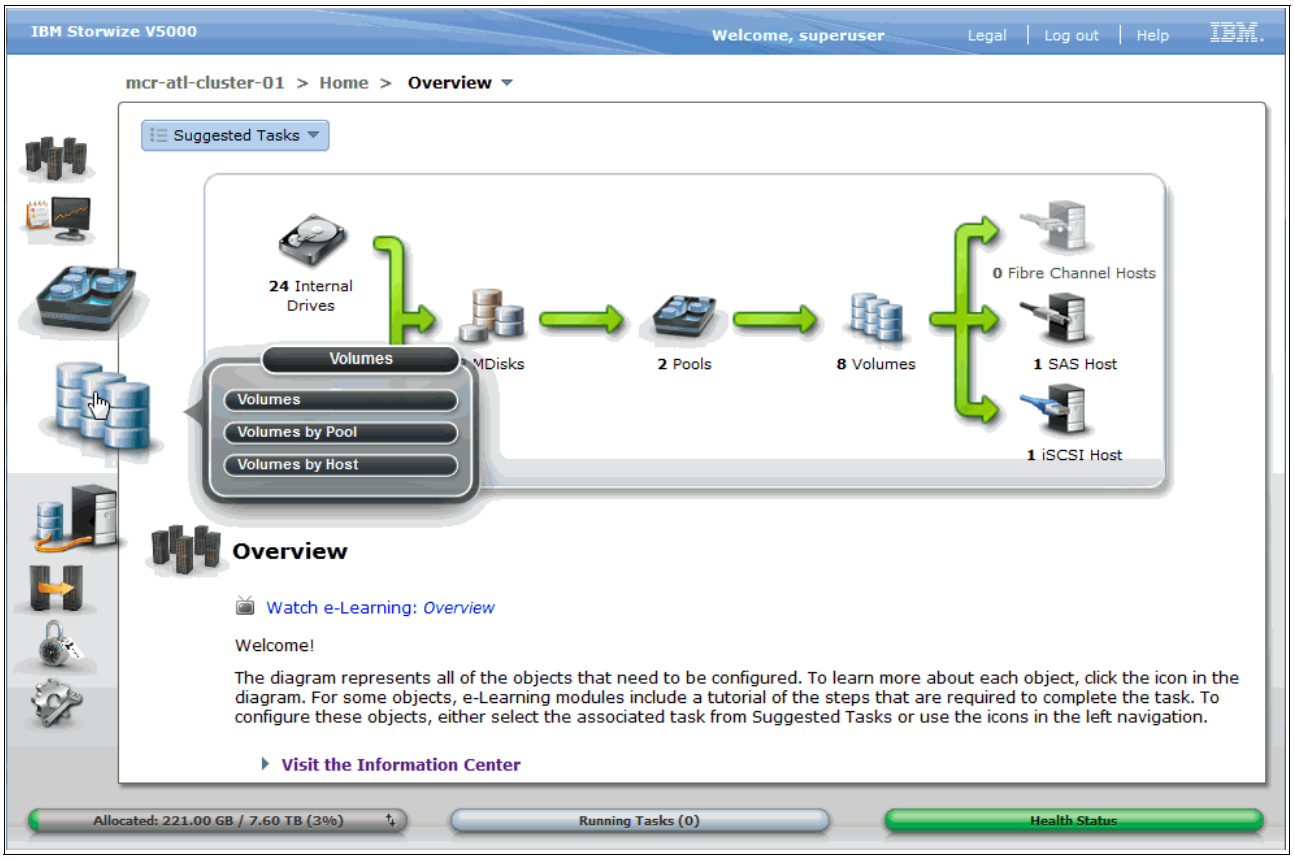

Figure 3-70 Selecting the Volumes menu

Volumes panel

Select **Volumes** in the Volumes menu to open the panel, as shown in Figure 3-71. The Volumes panel shows all of the volumes in the system. The information that is displayed is dependent on the columns that are selected.

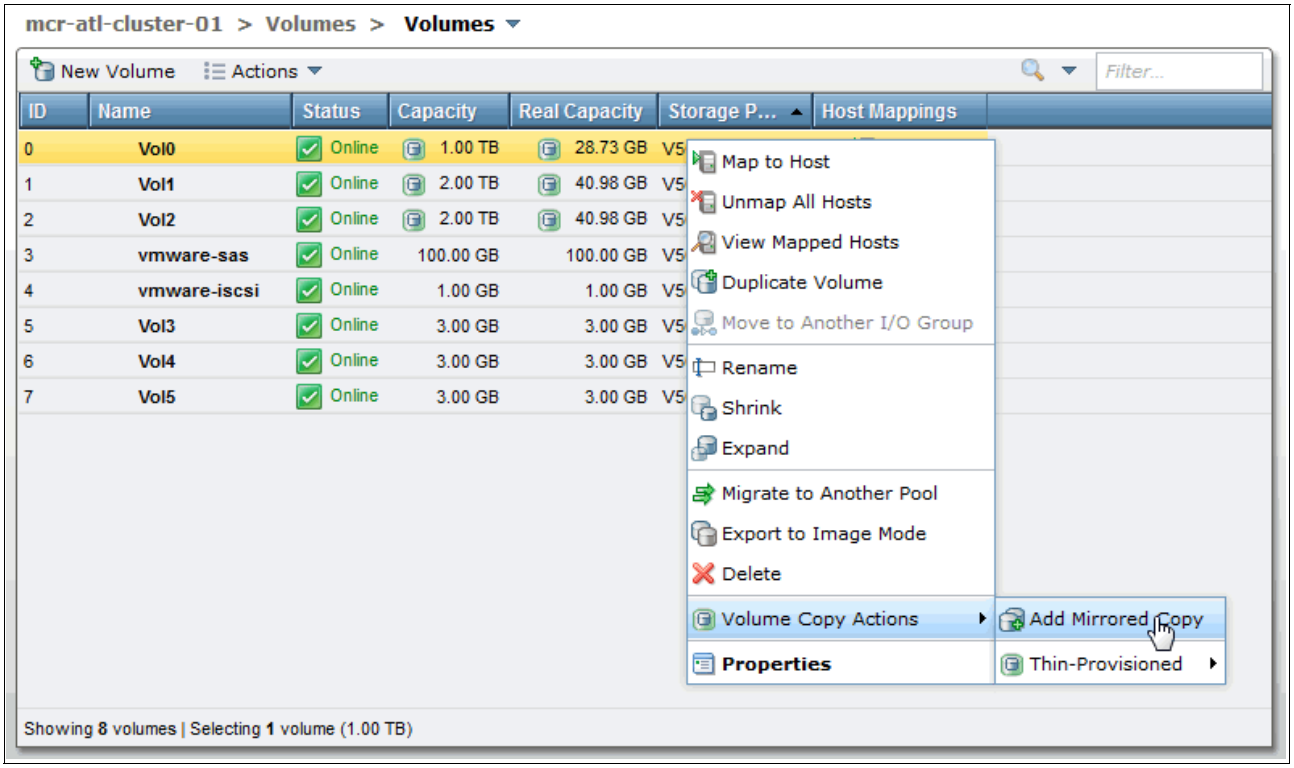

Figure 3-71 Volumes panel

Volume actions

Volume actions such as, Map to Host, Unmap All Hosts, Rename, Shrink, Expand, Migrate to Another Pool, Delete, and Add Mirrored Copy can be performed from this panel.

Create new volumes

Click **New Volume** to open the New Volume panel, as shown in Figure 3-72 on page 123. By using this panel, you can select a preset when a volume is created. The presets are designed to accommodate most user cases. The presets are generic, thin-provisioned, mirror, or thin mirror. After a preset is determined, select the storage pool from which the volumes are allocated. An area to name and size the volumes is shown.

For more information, see Chapter 5, “I/O Group basic volume configuration” on page 161 and Chapter 8, “Advanced host and volume administration” on page 349.

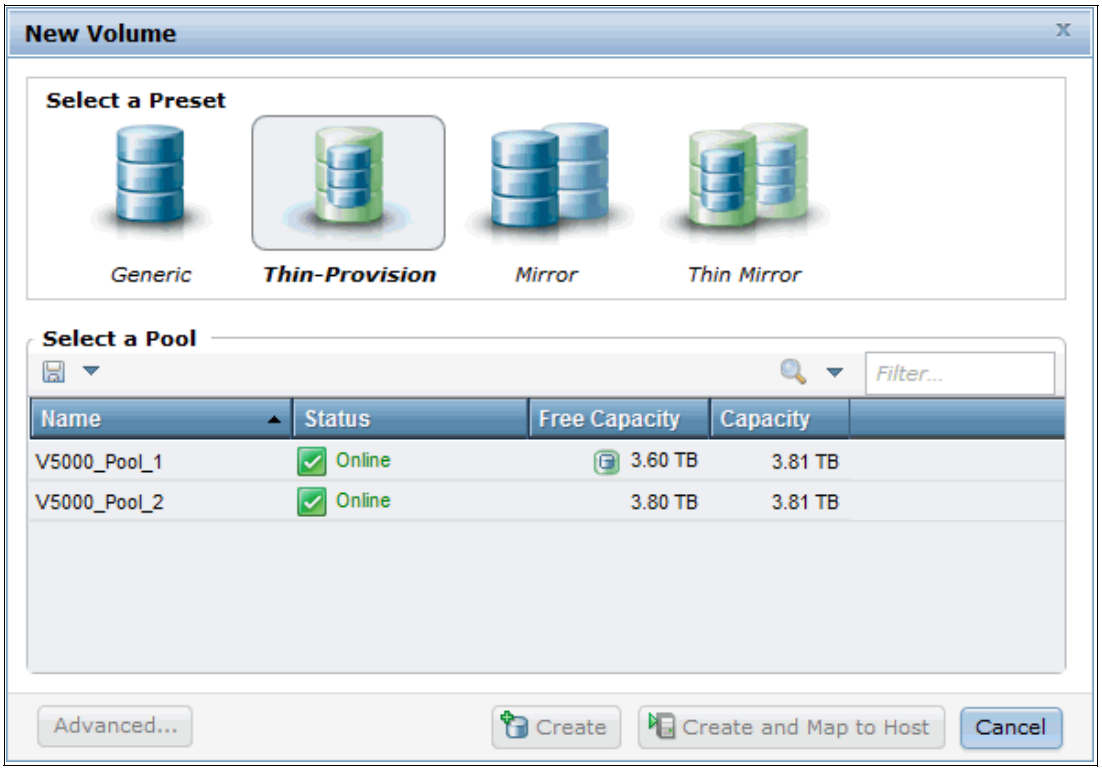

Figure 3-72 New Volume panel

Creating multiple volumes

A useful feature is available for quickly creating multiple volumes of the same type and size. Specify the number of volumes that are required in the Quantity field, then complete the volume capacity and name. A number range can also be specified.

The New Volumes panel displays a summary that shows the real and virtual capacity that is used if the proposed volumes are created. Click **Create** or **Create and Map to Host** to continue.

Figure 3-73 shows the quantity of 3 in the Quantity field.

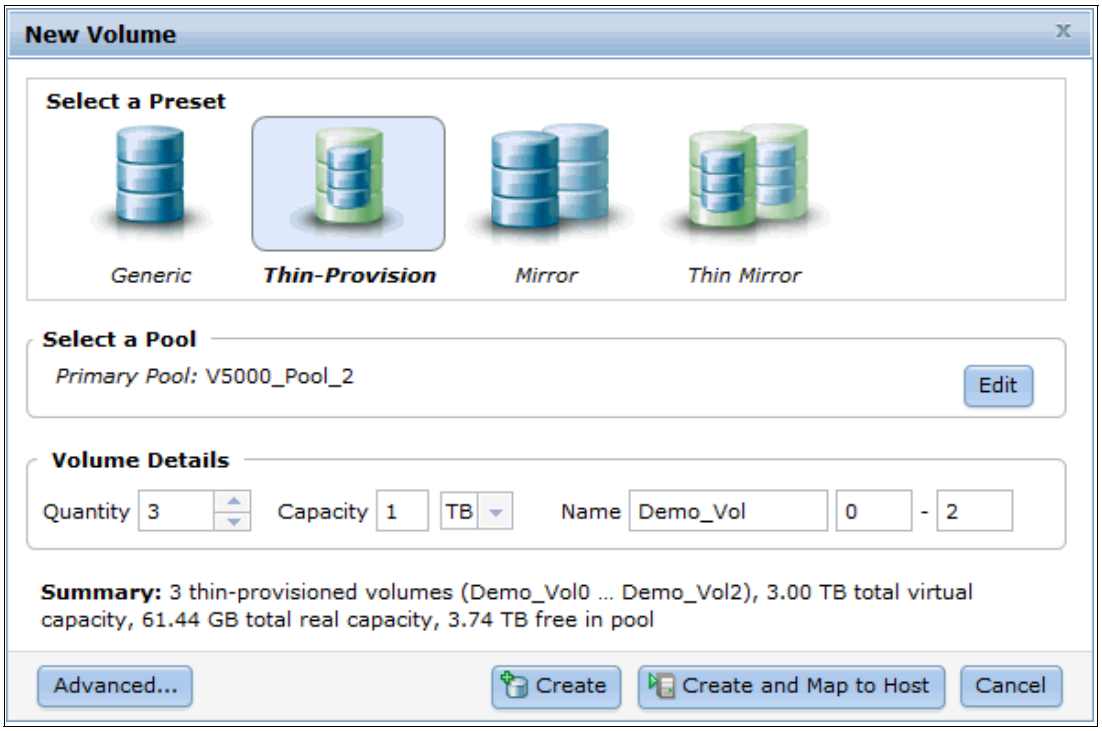

Figure 3-73 Creating multiple volumes quickly

Volume advanced settings

Click **Advanced** to show more volume configuration options. Use this feature when the preset does not meet your requirements. After the advanced settings are configured, click **OK** to return to the New Volumes panel. Figure 3-74 shows the Advanced Settings panel.

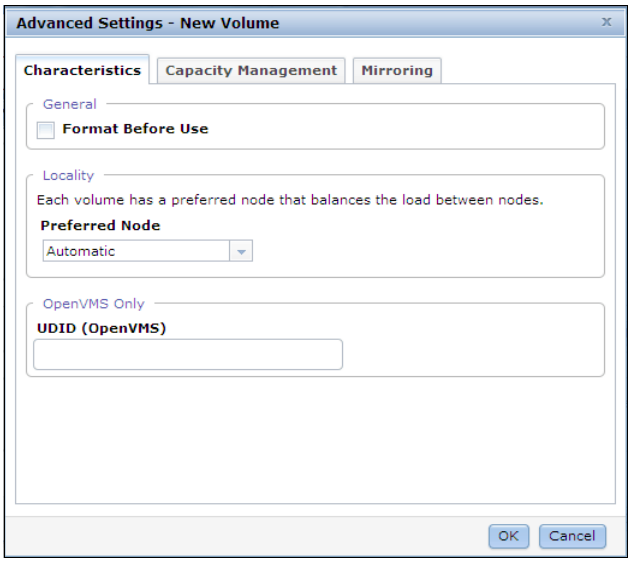

Figure 3-74 Advanced Settings panel

Volumes by Pool panel

For more information, see “Volumes by Pool panel” on page 111.

Volumes by Host panel

Select **Volumes by Host** in the Volumes menu to open the panel. By using the Volume by Hosts panel, you can focus on volumes that are allocated to a particular host by using the host selection filter.

3.4.5 Hosts menu

As shown in Figure 3-75, the Hosts menu provides access to the Hosts, Ports by Host, Host Mappings, and Volumes by Host functions.

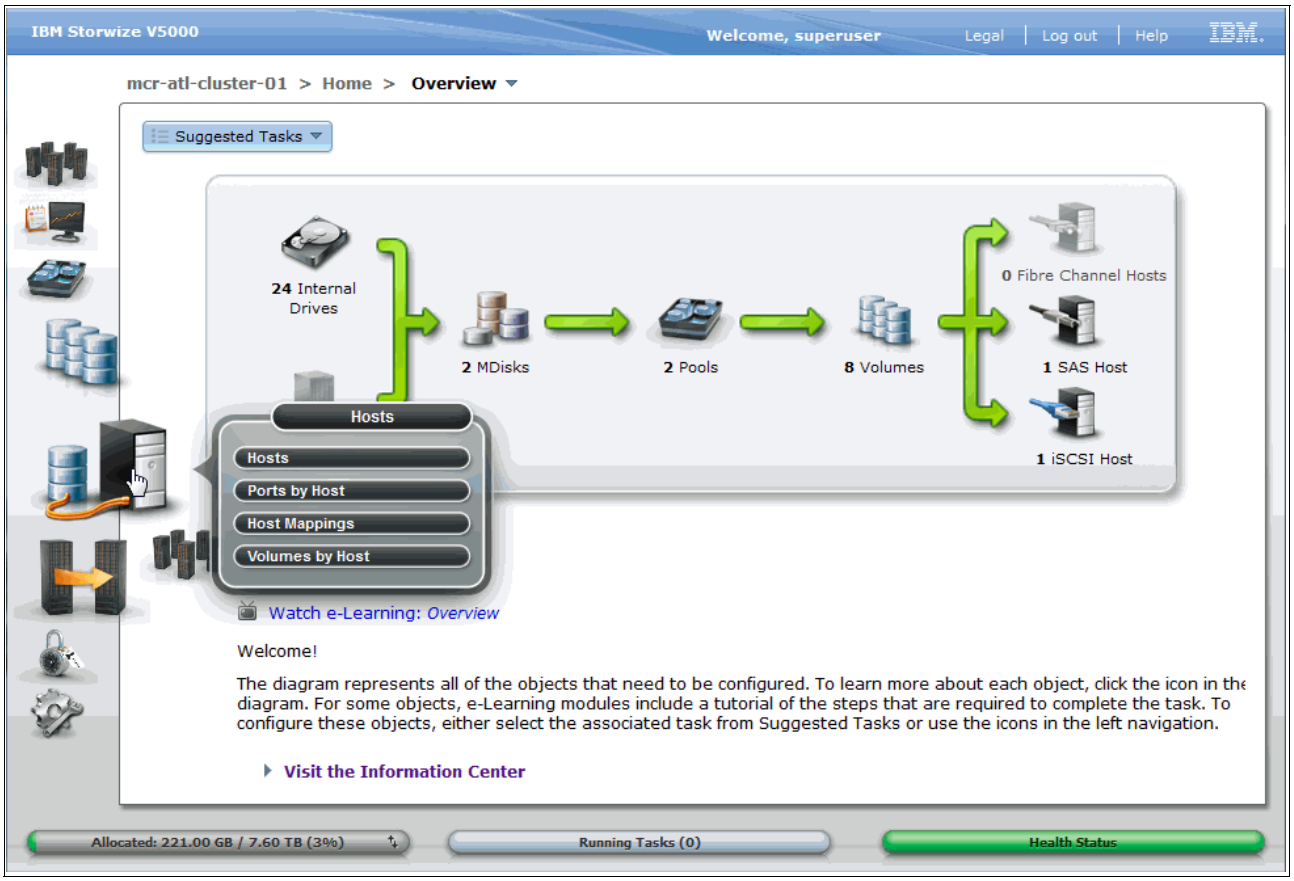

Figure 3-75 Selecting the Hosts menu

Hosts panel

Select **Hosts** in the Hosts menu to open the panel, as shown in Figure 3-76. The Hosts panel shows all of the hosts that are defined in the system.

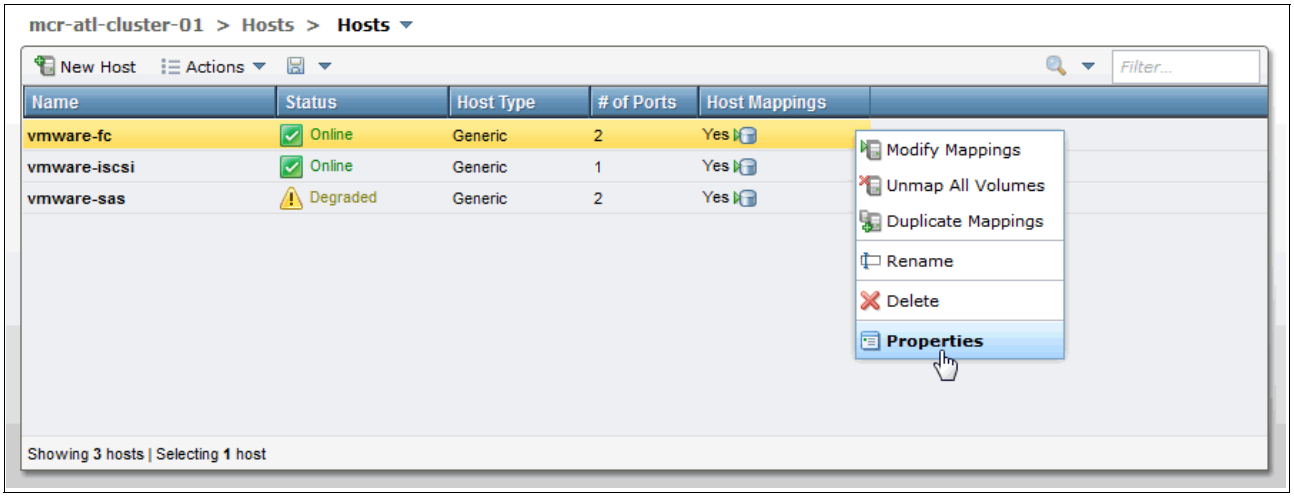

Figure 3-76 Hosts panel

Host Actions

Host Actions, such as, Modify Mappings, Unmap All Volumes, Duplicate Mappings, Rename, Delete, and Properties can be performed from the Hosts panel. Figure 3-76 shows the actions that are available from the Hosts panel.

For more information about the Hosts Actions menu, see 8.1, “Advanced host administration” on page 350.

Creating a host

Click **New Host** and the Create Host panel opens. Choose the host type from Fibre Channel (FC), iSCSI, or SAS host and the applicable host configuration panel is shown. After the host type is determined, the host name and port definitions can be configured. Figure 3-77 on page 127 shows the Choose the Host Type panel of the Create Host window.

For more information about how to create hosts, see Chapter 4, “Host configuration” on page 153.

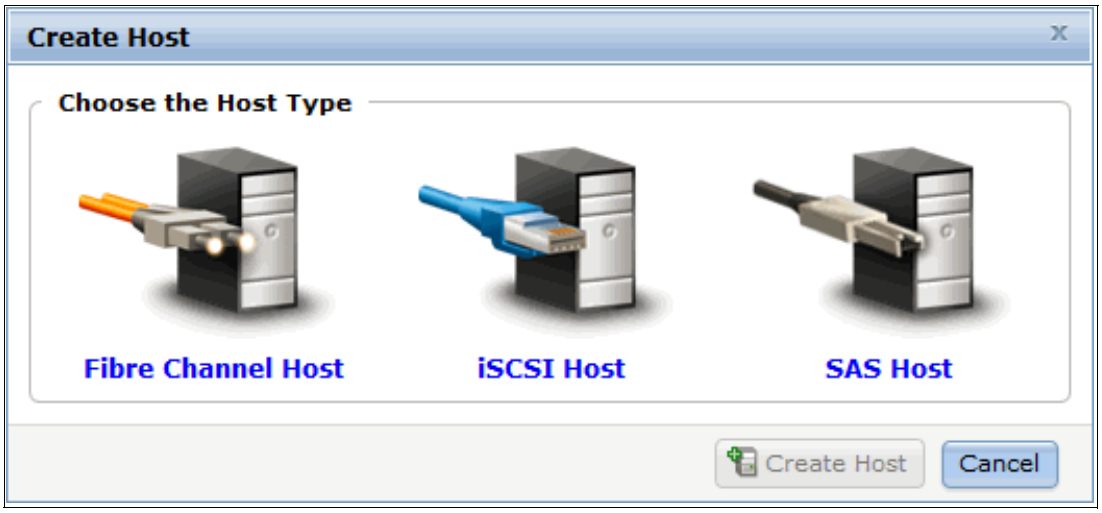

Figure 3-77 Choose the Host Type panel

Ports by Host panel

Select **Ports by Host** in the Hosts menu to open the panel, as shown in Figure 3-78. The panel shows the address, status, and type of ports that are assigned to the host definition. Actions such as, map, unmap, and port deletion can be performed from this panel.

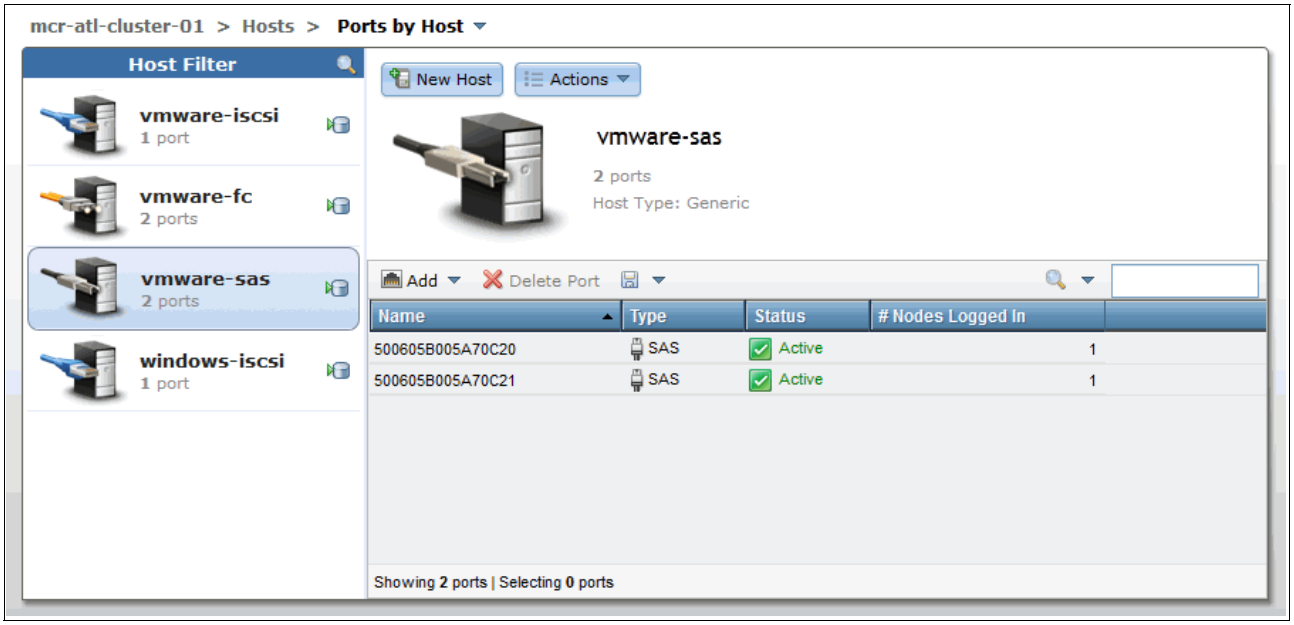

Figure 3-78 Ports by Host panel

Host Mappings panel

Select **Host Mappings** in the Hosts menu to open the panel, as shown in Figure 3-79. This panel shows the volumes that each host can access with the corresponding SCSI ID. The Unmap Volume action can be performed from this panel.

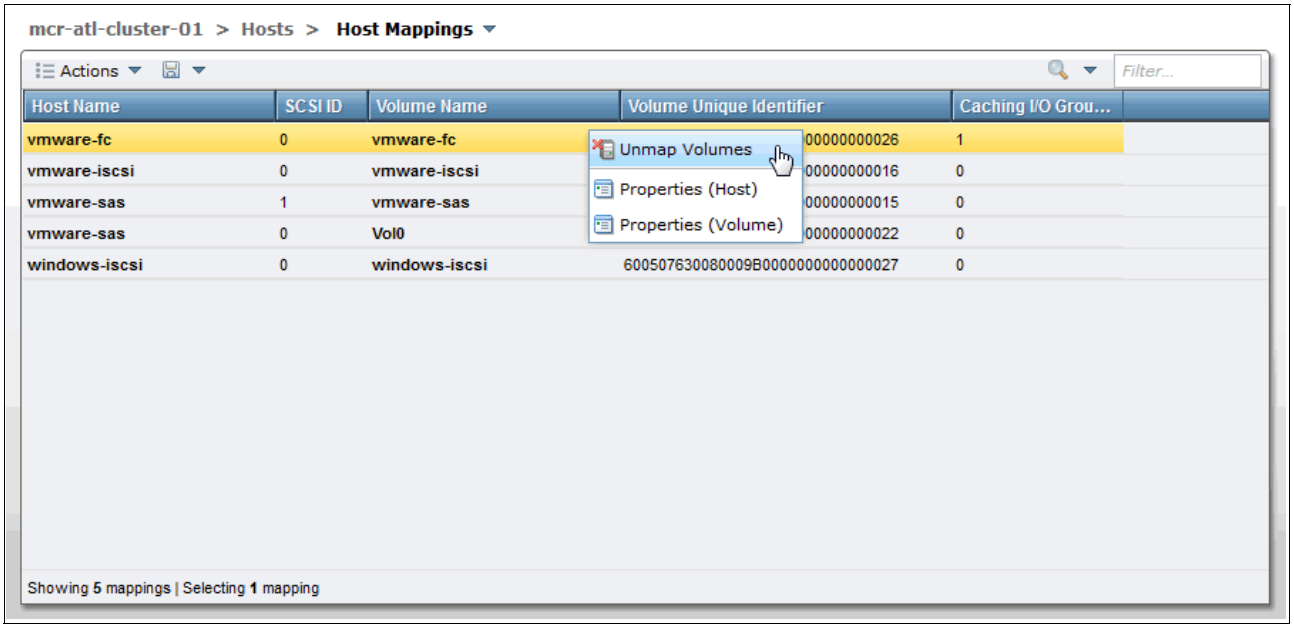

Figure 3-79 Host Mappings panel

Volumes by Host panel

For more information, see “Volumes by Host panel” on page 125.

3.4.6 Copy Services menu

The Copy Services menu provides access to the FlashCopy, Consistency Groups, FlashCopy Mappings, Remote Copy, and Partnership functions. Figure 3-80 on page 129 shows the Copy Services menu.

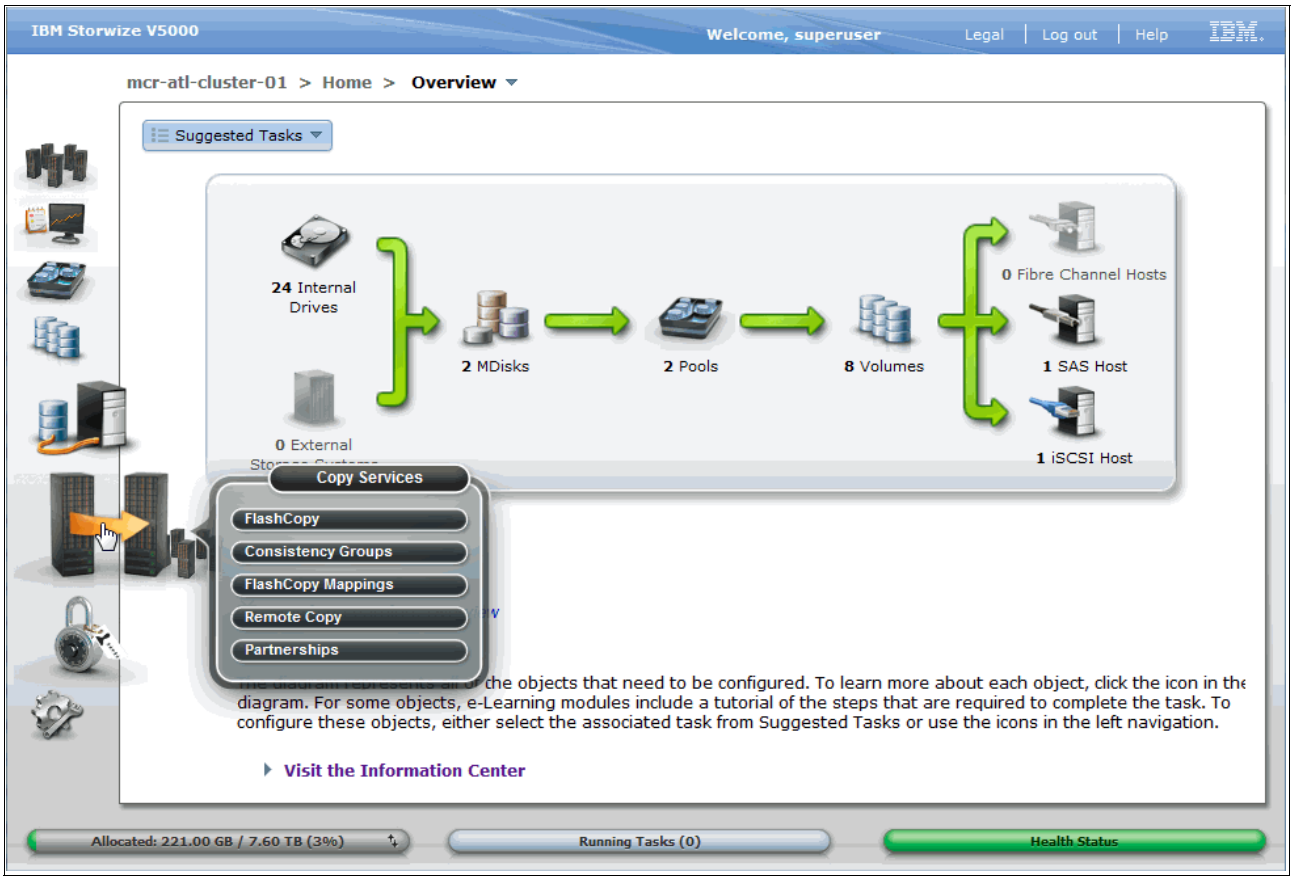

Figure 3-80 Copy Services menu

FlashCopy panel

Select **FlashCopy** in the Copy Services menu to open the panel, as shown in Figure 3-81. The FlashCopy panel displays all of the volumes that are in the system.

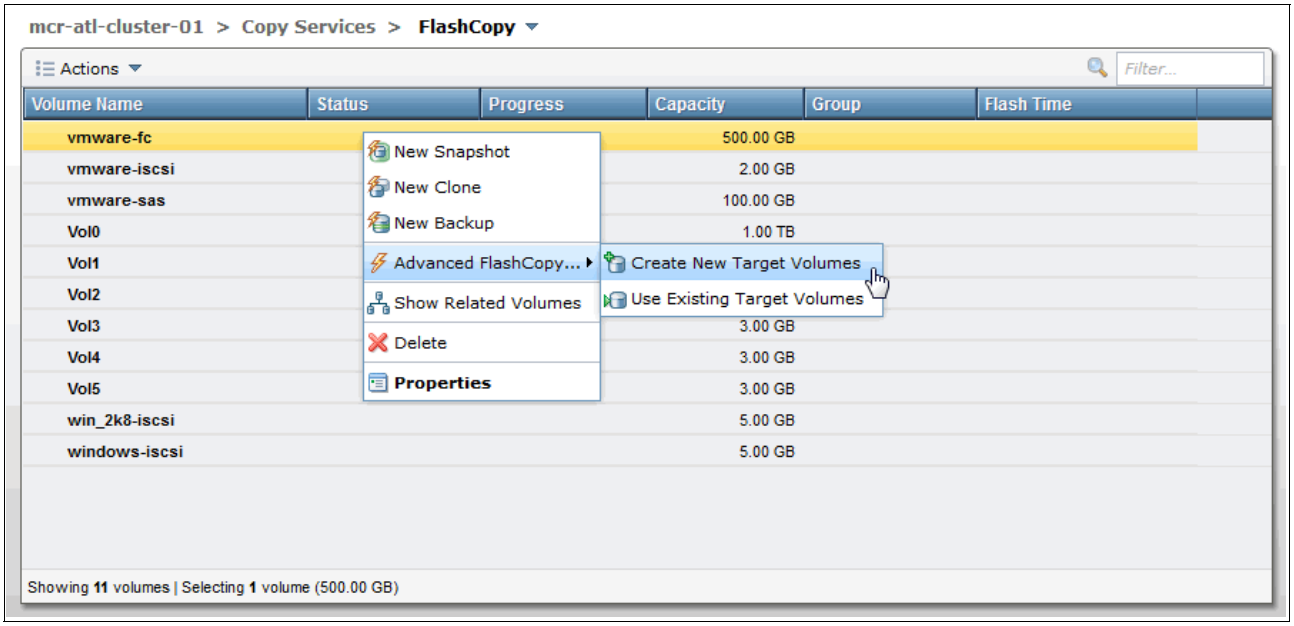

Figure 3-81 FlashCopy panel

FlashCopy actions

FlashCopy actions such as, New Snapshot, New Clone, New Backup, Advanced FlashCopy, and Delete can be performed from this panel. Figure 3-81 on page 129 shows the actions that are available from the FlashCopy panel.

Consistency Groups panel

Select **Consistency Groups** in the Copy Services menu to open the panel. A consistency group is a container for FlashCopy mappings. Grouping allows FlashCopy mapping actions such as, prepare, start, and stop to occur at the same time for the group instead of coordinating actions individually. This feature can help ensure that the group's target volumes are consistent to the same point and remove several FlashCopy mapping administration tasks.

The Consistency Group panel shows the defined groups with the associated FlashCopy mappings. Group Actions such as, FlashCopy Map Start, Stop, and Delete can be performed from this panel. New FlashCopy Mapping also can be selected from this panel. For more information, see "FlashCopy mappings panel". Figure 3-82 shows the Consistency Group panel.

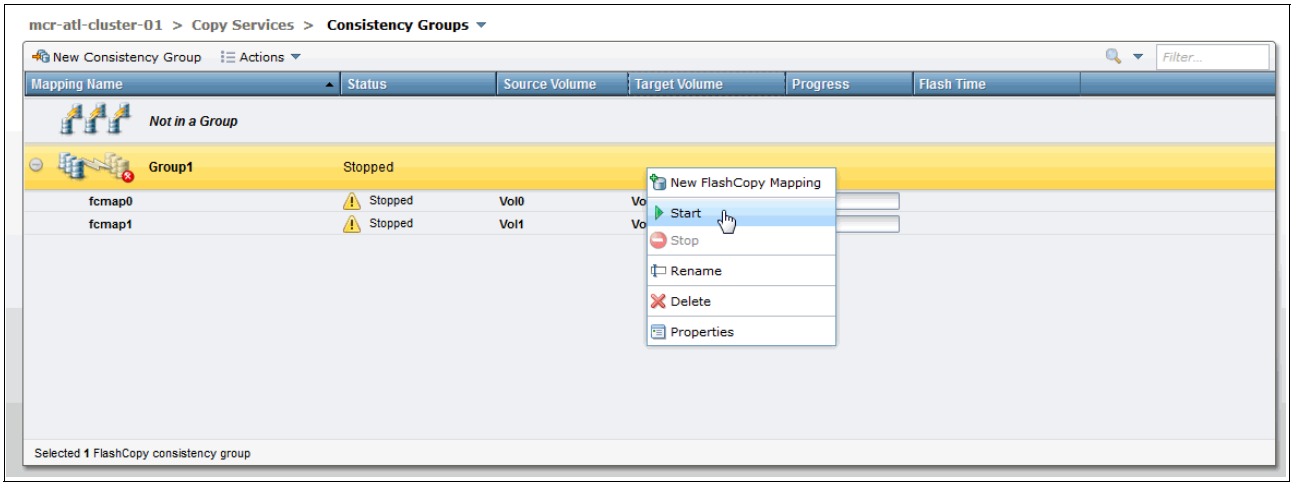

Figure 3-82 Consistency Groups panel

FlashCopy mappings panel

Select **FlashCopy Mappings** in the Copy Services menu to open the panel. FlashCopy mappings define the relationship between source volumes and target volumes. The FlashCopy Mappings panel shows information that relates to each mapping, such as, status, progress, source and target volumes, and flash time. Select **New FlashCopy Mapping** to configure a new mapping or use the Actions menu to administer the mapping. Figure 3-83 on page 131 shows the FlashCopy Mappings panel.

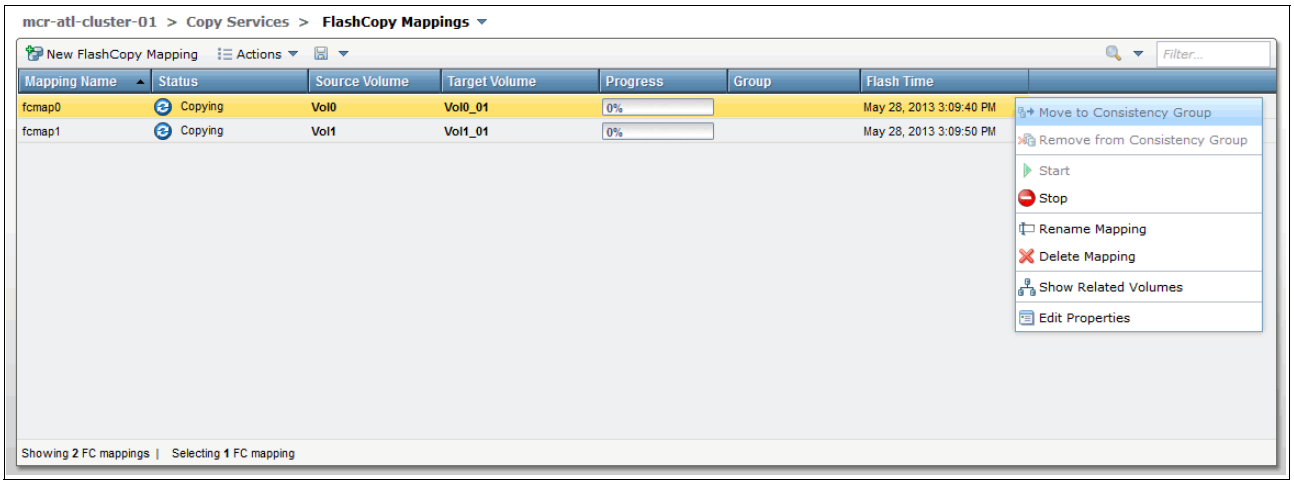

Figure 3-83 FlashCopy Mappings panel

For more information about how to create and administer FlashCopy mappings, see Chapter 8, “Advanced host and volume administration” on page 349.

Remote Copy panel

Clicking **Remote Copy** opens the window that is shown in Figure 3-84. This window shows the existing Remote Copy relationships in which you can set up and modify consistency groups. From this window, you can also start and stop relationships, add relationships to a consistency group, and switch the direction of the mirror.

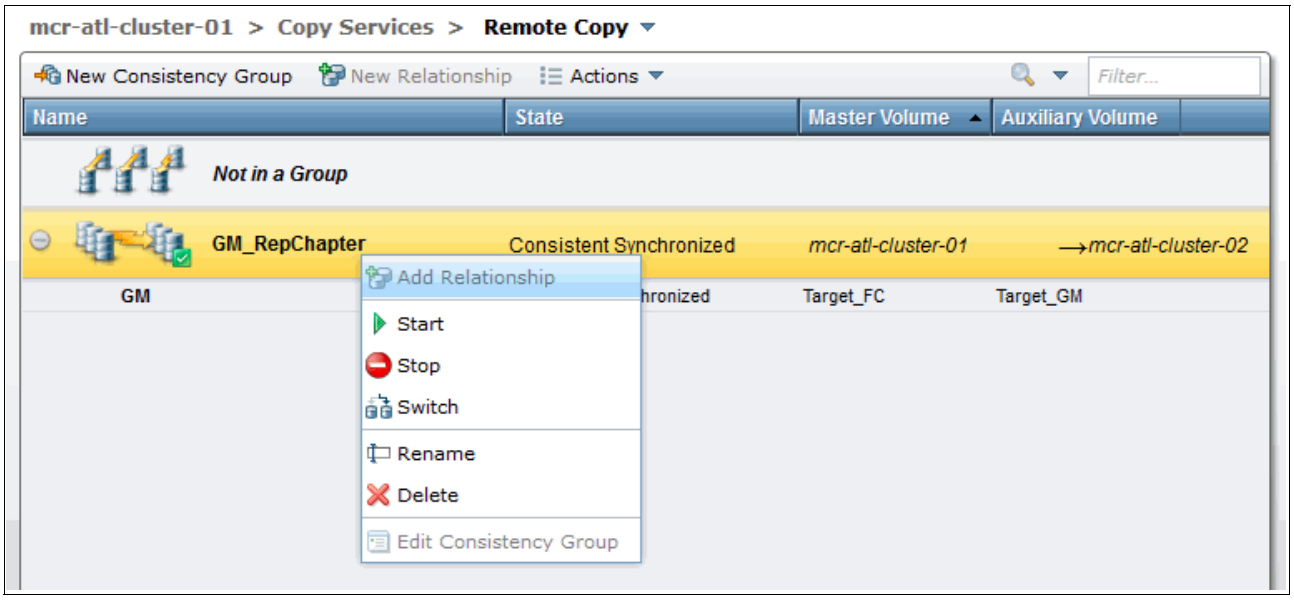

Figure 3-84 Remote Copy window

Partnerships panel

Clicking **Partnerships** opens the window that is shown in Figure 3-85. In this window, you can set up a new partnership or delete an existing partnership with another IBM Storwize or SAN Volume Controller system for the purposes of remote mirroring.

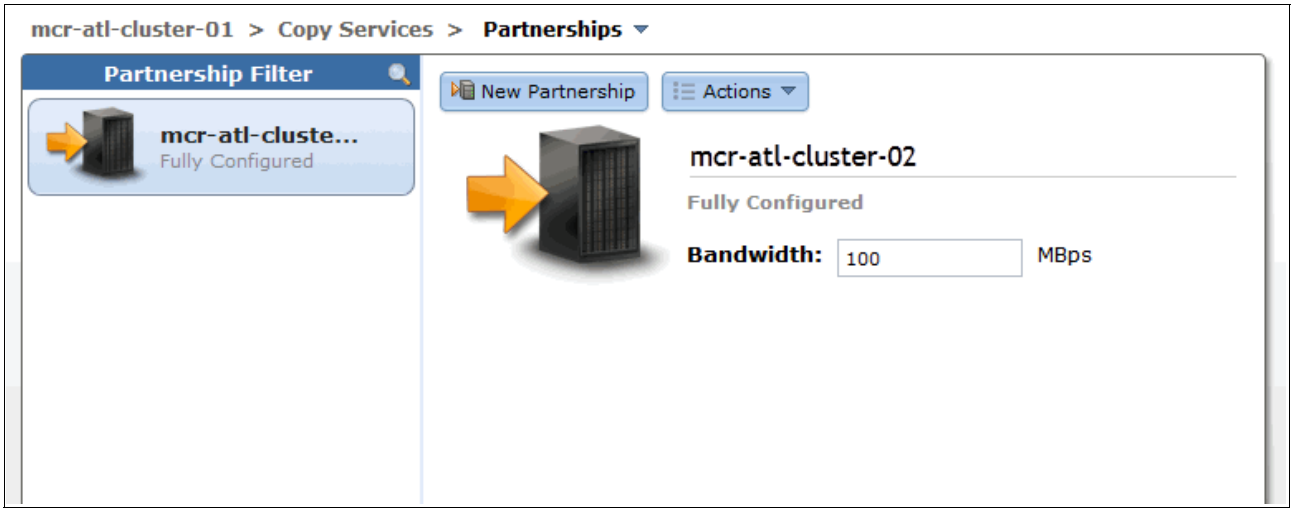

Figure 3-85 Partnerships window

From this window, you can also set the background copy rate. This rate specifies the bandwidth, in megabytes per second (MBps), that is used by the background copy process between the clusters.

3.4.7 Access menu

The Access menu provides access to the Users and Audit Log functions, as shown in Figure 3-86.

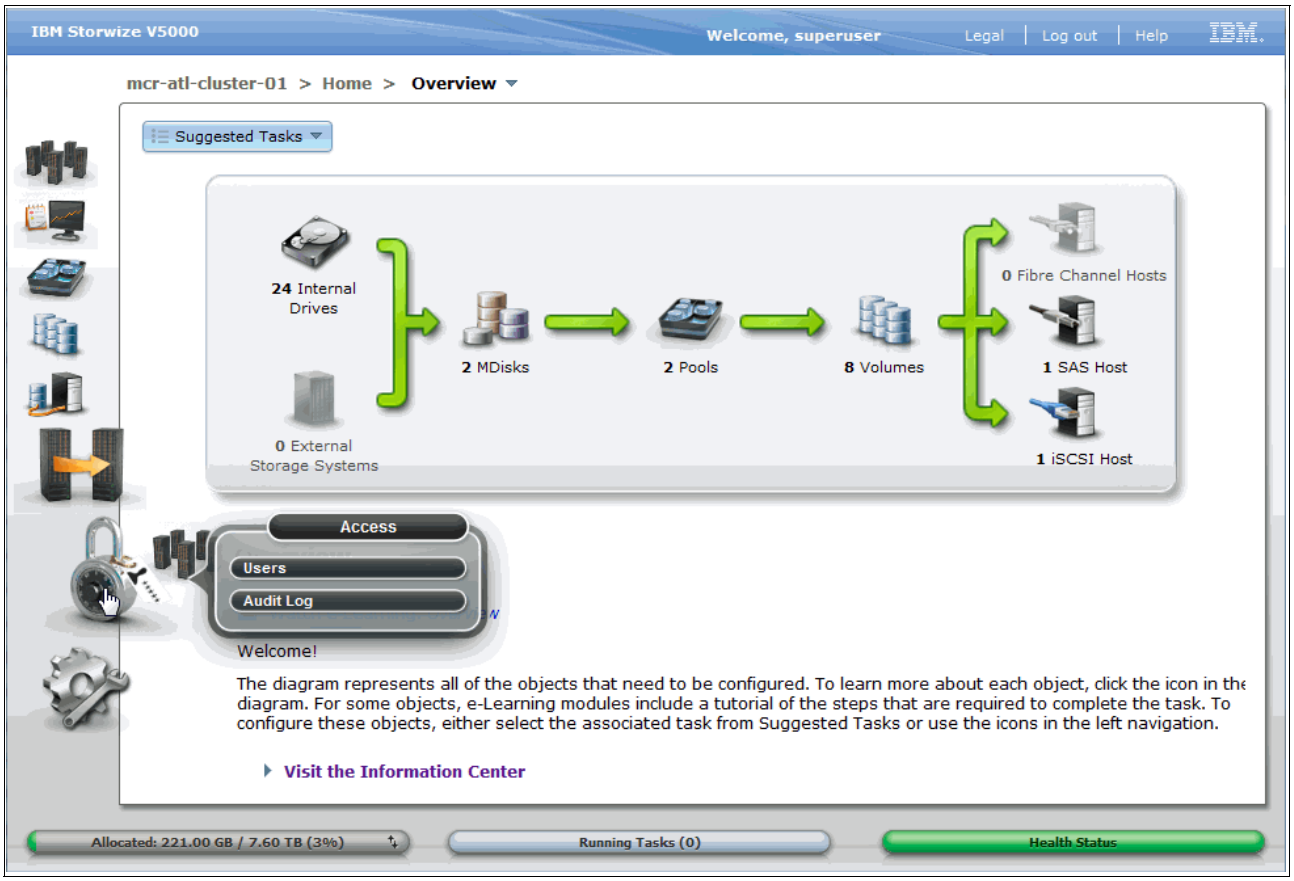

Figure 3-86 Access menu

Users panel

Select **Users** in the Access menu to open the panel. The Users panel shows the defined user groups and users for the system. The users that are listed can be filtered by user group. Click **New User Group** to open the Create a New Group panel. Figure 3-87 shows the Users panel and the Users Actions menu.

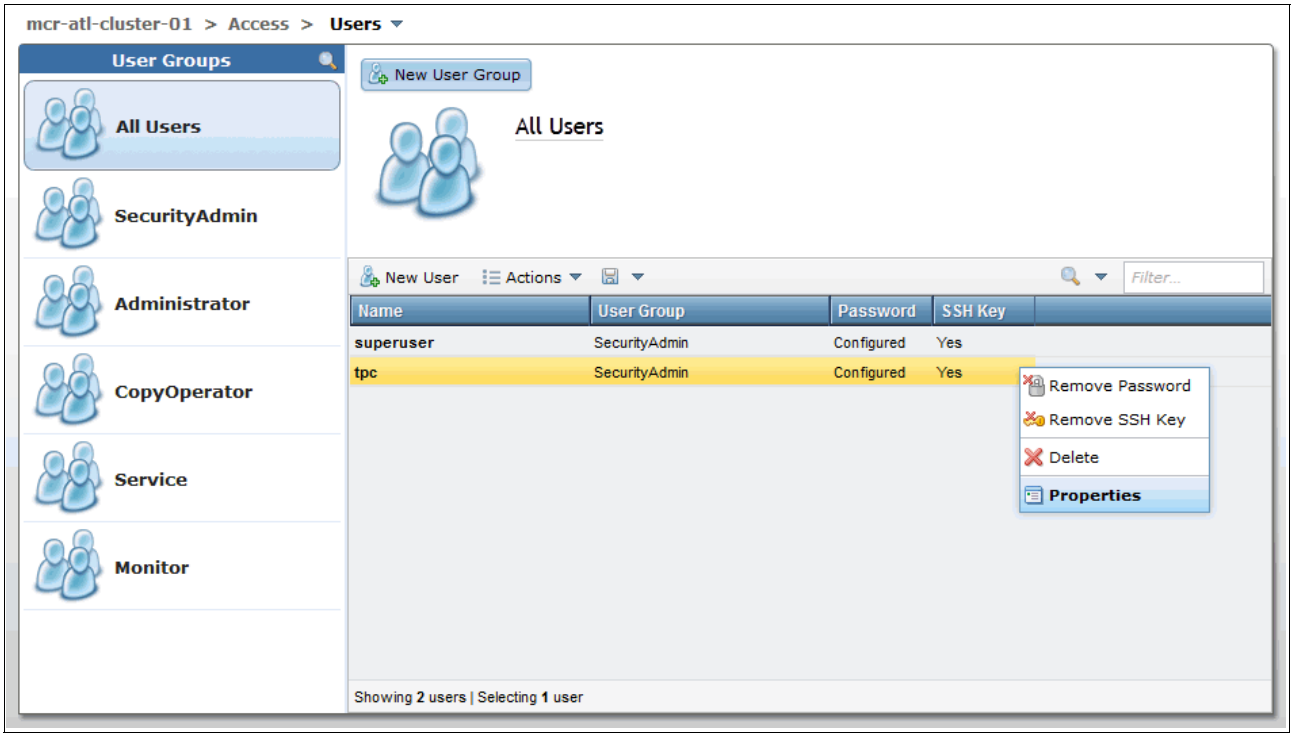

Figure 3-87 Users panel

Creating a user group

By using the New User Group panel, you can configure user groups. Enter the group name, select the role, then click **Create**, as shown in Figure 3-88.

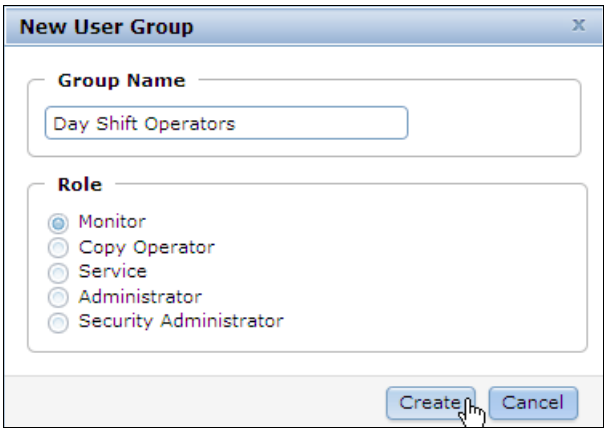

Figure 3-88 New User Group panel

Creating a user

Click **New User** to define a user to the system. Figure 3-89 shows the Users panel and the New User option.

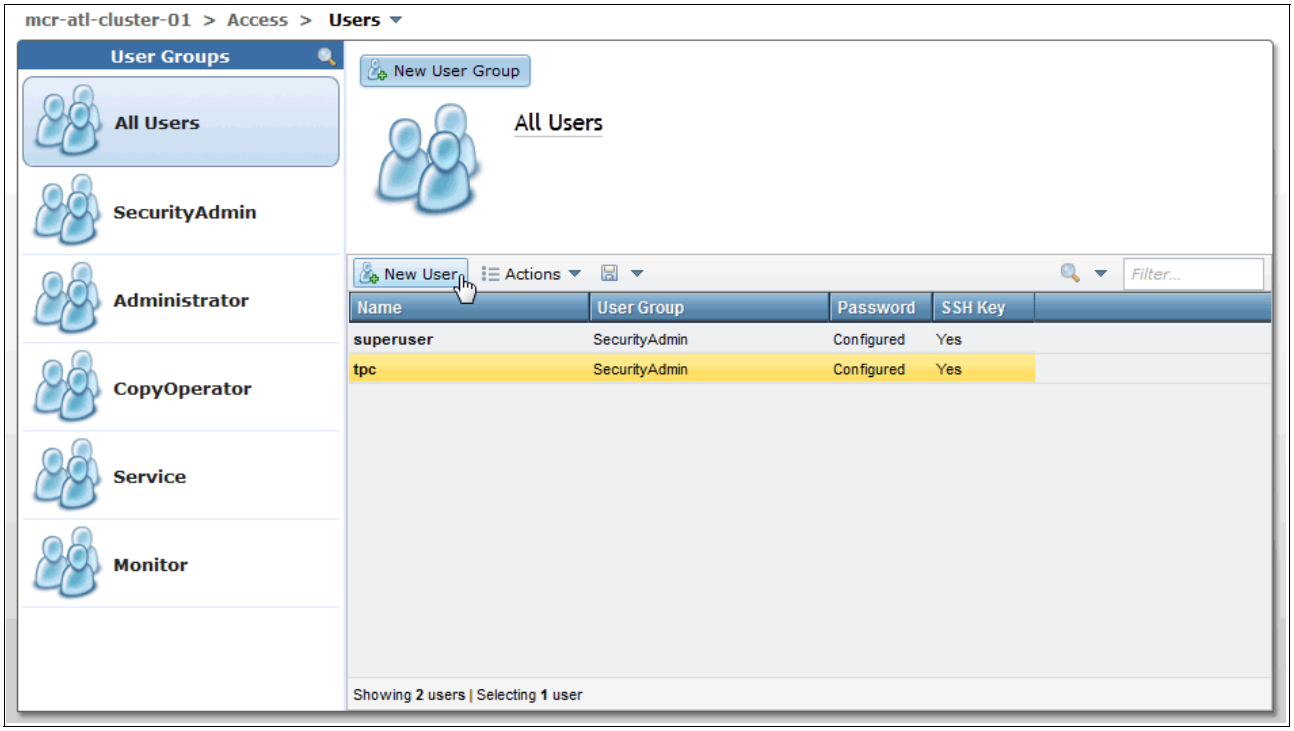

Figure 3-89 Users panel and the New User option

By using the New User panel, you can configure the user name, password, and authentication mode. It is essential to enter the user name, password, group, and authentication mode. The public Secure Shell (SSH) key is optional. After the user is defined, click **Create**.

The authentication mode can be set to local or remote. Select local if the IBM Storwize V5000 performs the authentication locally. Select remote if a remote service such as, an LDAP server authenticates the connection. If remote is selected, the remote authentication server must be configured in the IBM Storwize V5000 by clicking **Settings menu** → **Directory Services panel**.

The SSH configuration can be used to establish a more secure connection to the command-line interface. For more information, see Appendix A, “Command-line interface setup and SAN Boot” on page 609.

Figure 3-90 shows the New User panel.

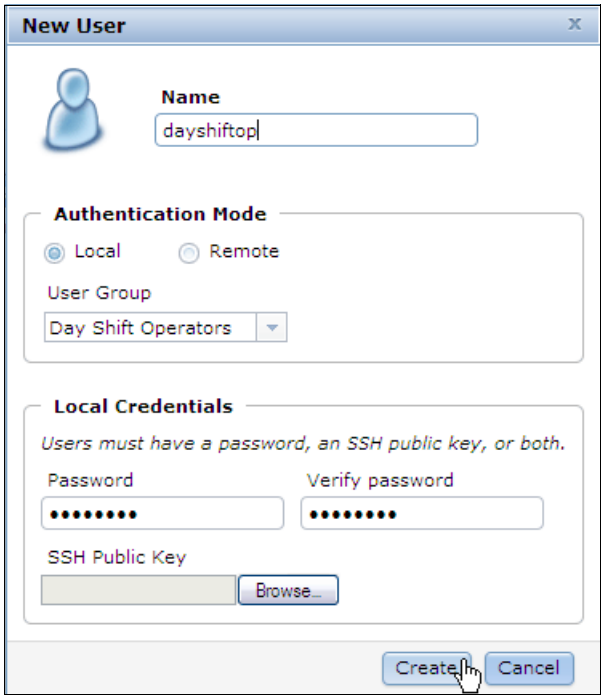

Figure 3-90 New User panel

Audit Log panel

Select **Audit Log** in the Access menu to open the panel. The audit log tracks action commands that are issued through a CLI session or through the management GUI. The Audit Log panel displays information about the command, such as, the user, time stamp, and any associated command parameters. The log can be filtered by date or by the Show entries within... feature to reduce the number of items that are listed. It is not possible to delete or alter the Audit log. Figure 3-91 shows the Audit Log panel.

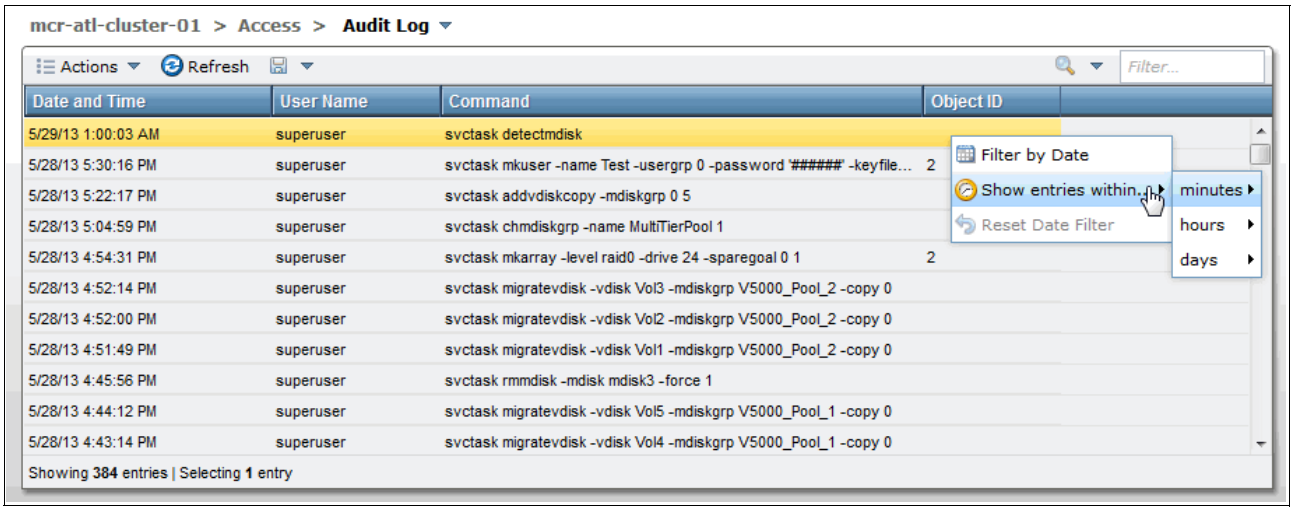

Figure 3-91 Audit Log panel

3.4.8 Settings menu

The Setting menu provides access to the Event Notifications, Directory Services, Network, Support, and General functions. Figure 3-92 shows the Settings menu.

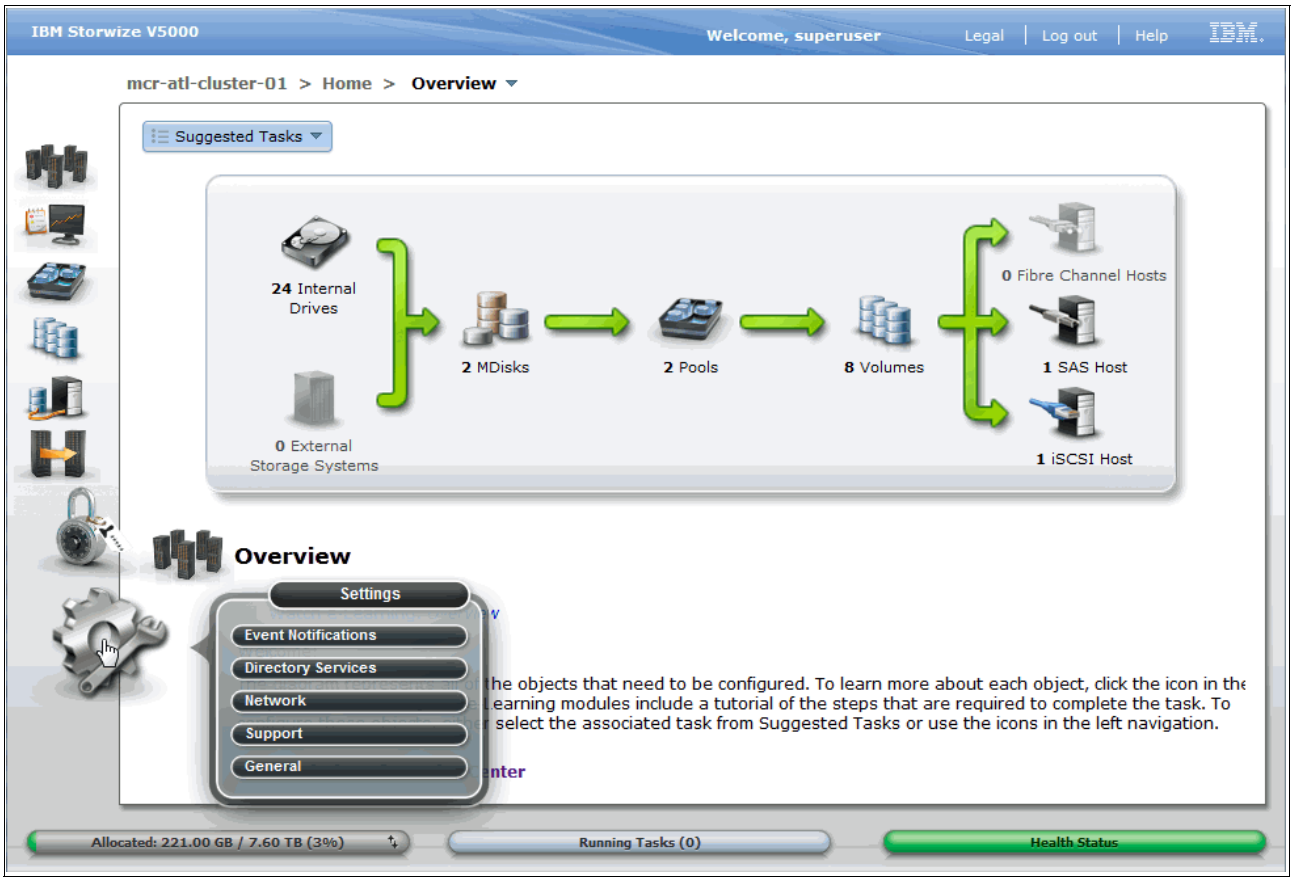

Figure 3-92 Settings menu

Event Notifications panel

Select **Event Notifications** in the Settings menu to open the panel. The IBM Storwize V5000 can use Simple Network Management Protocol (SNMP) traps, syslog messages, emails, and IBM Call Home to notify users when events are detected. Each event notification method can be configured to report all events or alerts. Alerts are the significant events and might require user intervention. The event notification levels are critical, warning, and information.

The Event Notifications panel provides access to the Email, SNMP, and Syslog configuration panels. IBM Call Home is an email notification for IBM Support. It is automatically configured as an email recipient and is enabled when the Email event notification option is enabled by following the Call Home wizard.

Enabling Email Event Notification option

Click **Enable Email Event Notification** to open the Call Home wizard. Figure 3-93 shows the Event Notifications Email configuration panel.

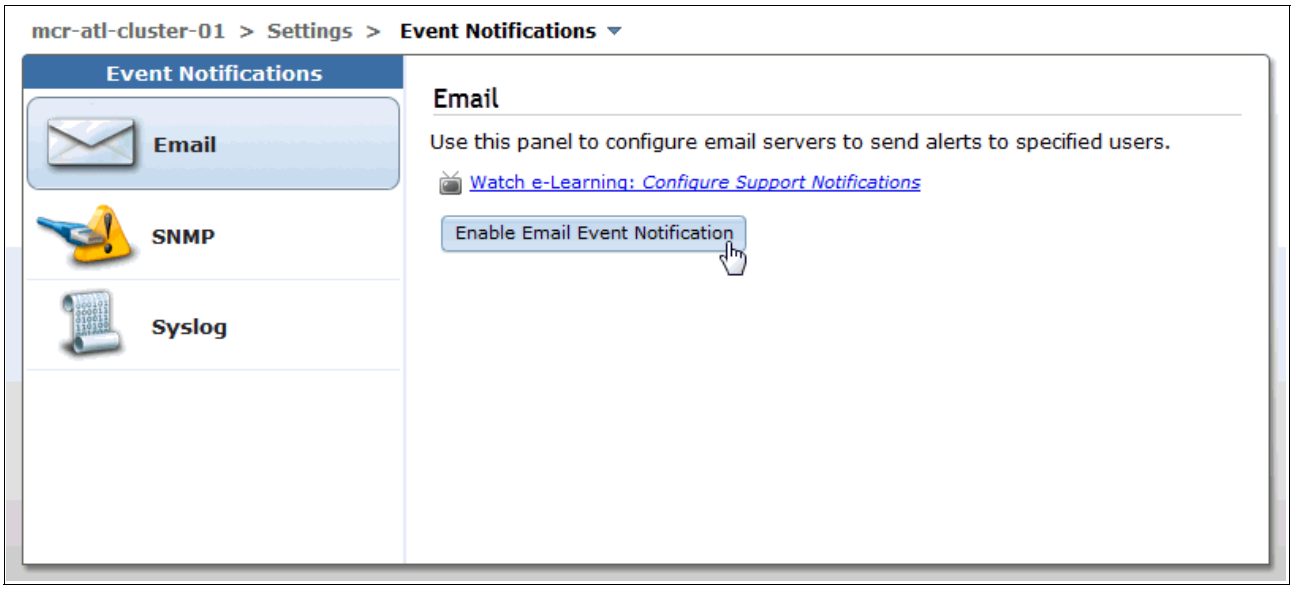

Figure 3-93 Event Notifications panel: Email

Call Home wizard

The Call Home wizard, as shown in Figure 3-94, guides the user through account contact, machine location entry, and email configuration tasks.

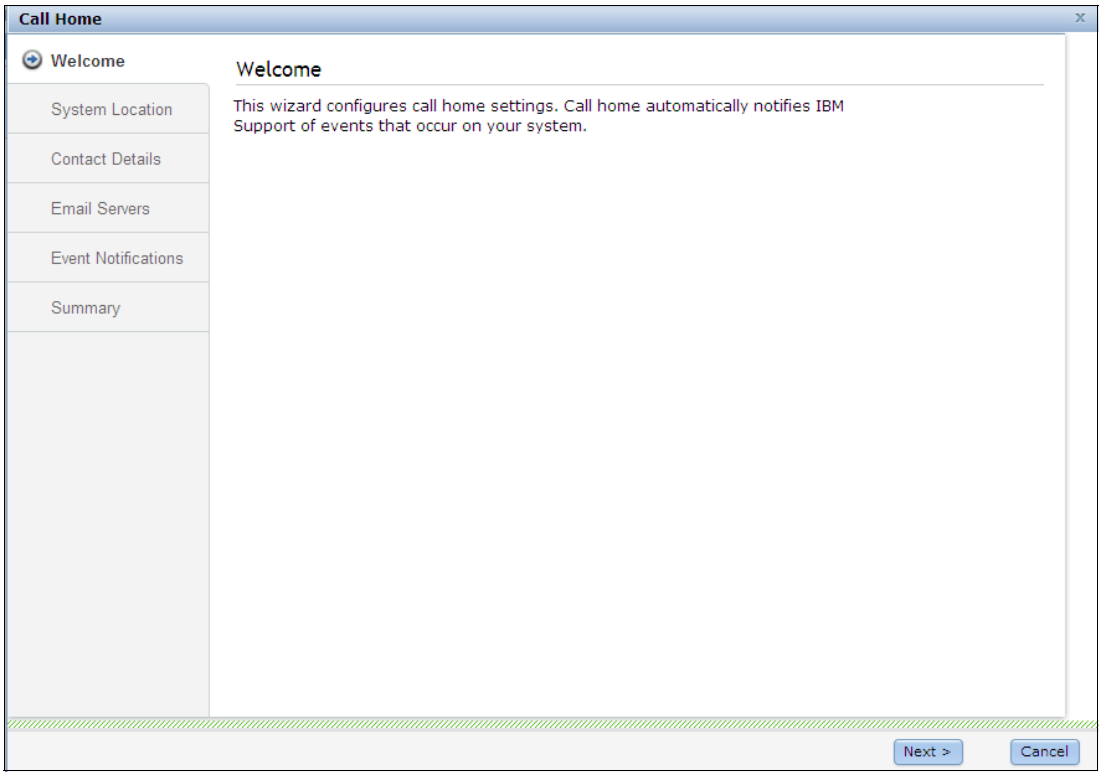

Figure 3-94 Call home wizard

SNMP event notification

As shown in Figure 3-95, the Event Notifications panel provides access to the SNMP configuration panel. Click **SNMP** to open the panel, then enter the server details. Multiple servers can be configured by clicking + to add more servers.

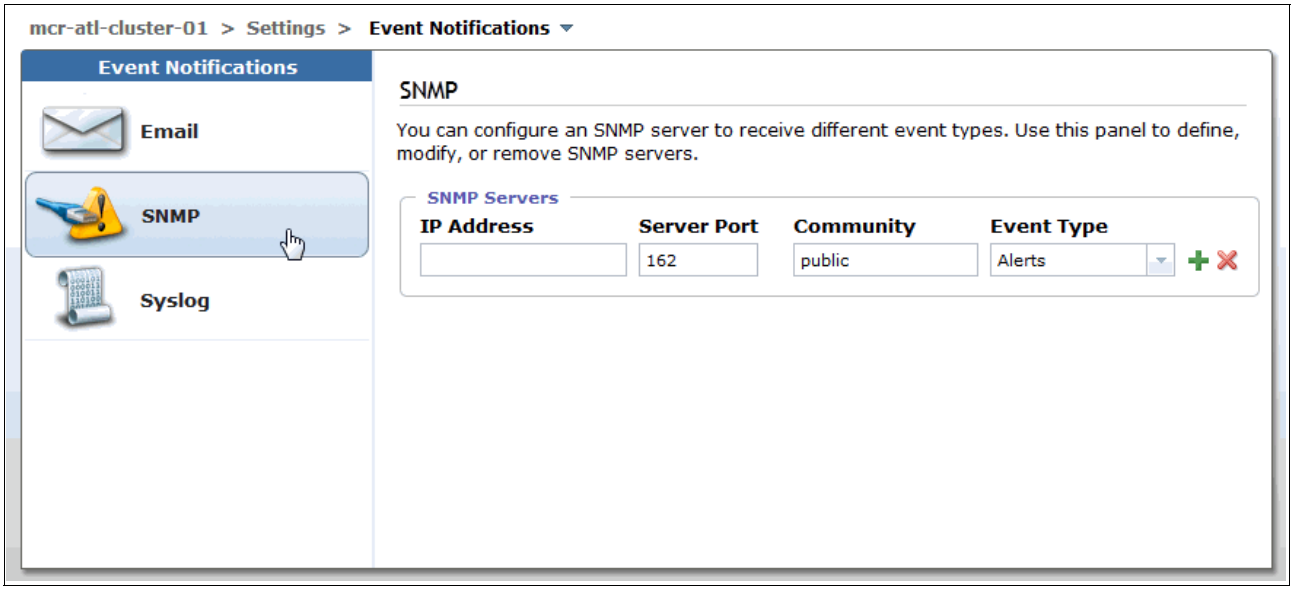

Figure 3-95 SNMP configuration panel

Syslog event notification

The Event Notifications panel provides access to the Syslog configuration panel. Click **Syslog** to open the panel, then enter the server details. Multiple servers can be configured by clicking + to add more servers. Figure 3-96 shows the Syslog configuration panel.

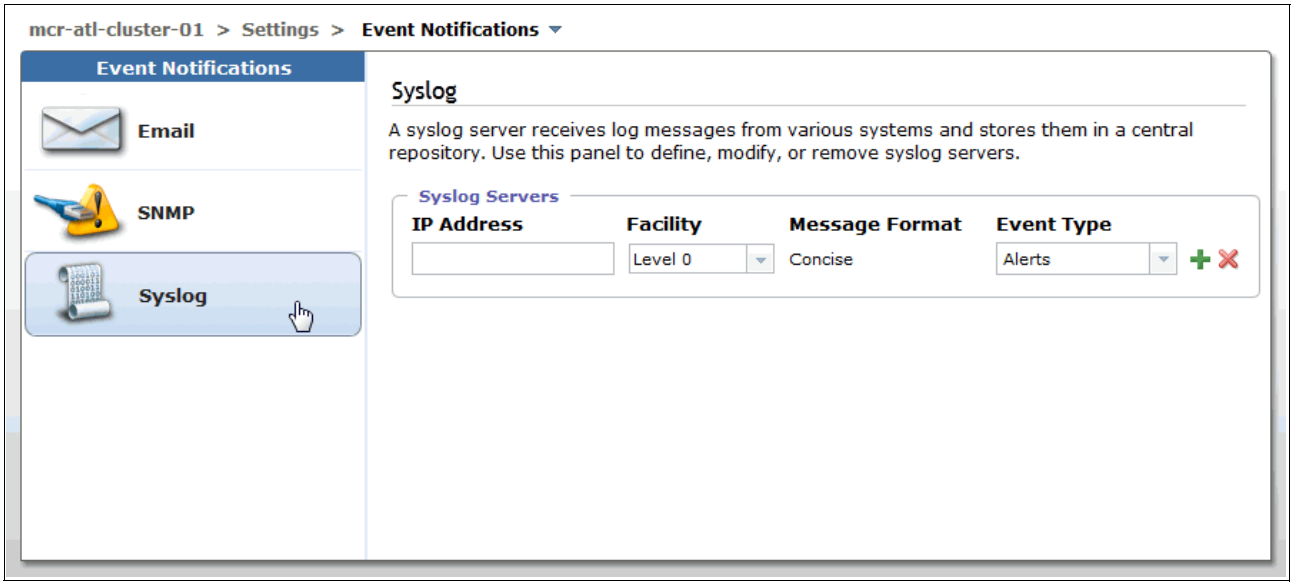

Figure 3-96 Syslog configuration panel

Directory Services panel

Select **Directory Services** in the Settings menu to open the panel. The Directory Services panel provides access to the Remote Authentication wizard. Remote authentication must be configured to create remote users on the IBM Storwize V5000. A remote user is authenticated on a remote service, such as, IBM Tivoli® Integrated Portal or a Lightweight Directory Access Protocol (LDAP) provider.

Enabling Remote Authentication

Click **Configure Remote Authentication** to open the wizard, as shown in Figure 3-97.

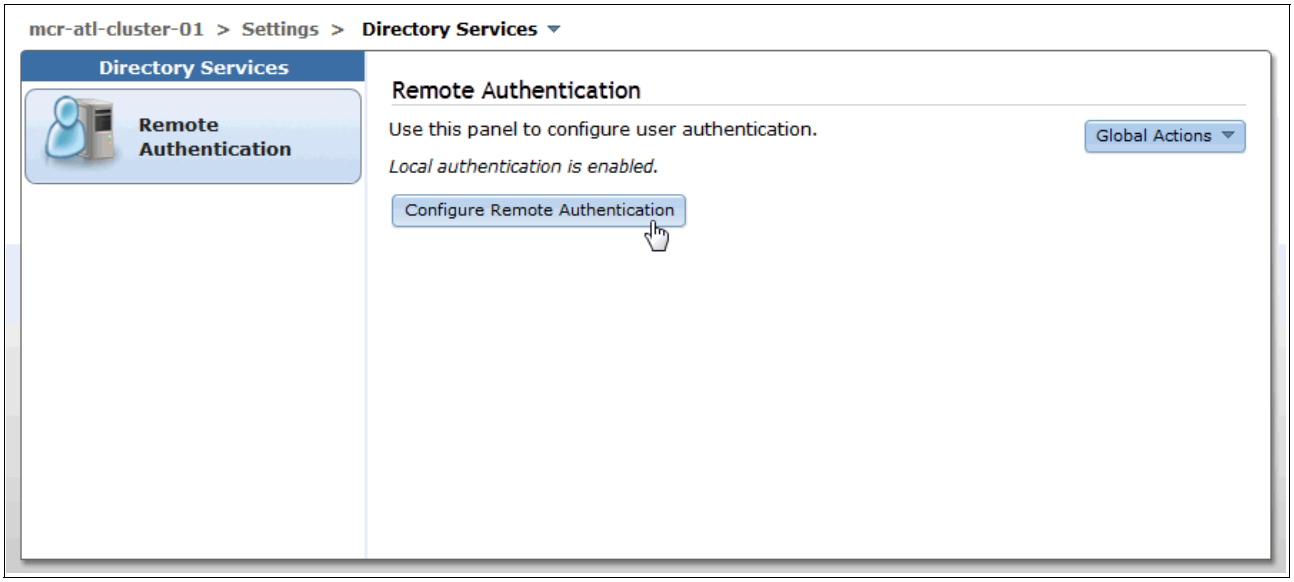

Figure 3-97 Directory Services panel

Network panel

Select **Network** in the General menu to open the panel. As shown in Figure 3-98, the Network panel provides access to the Management IP Addresses, Service IP Addresses, iSCSI, and Fibre Channel configuration panels.

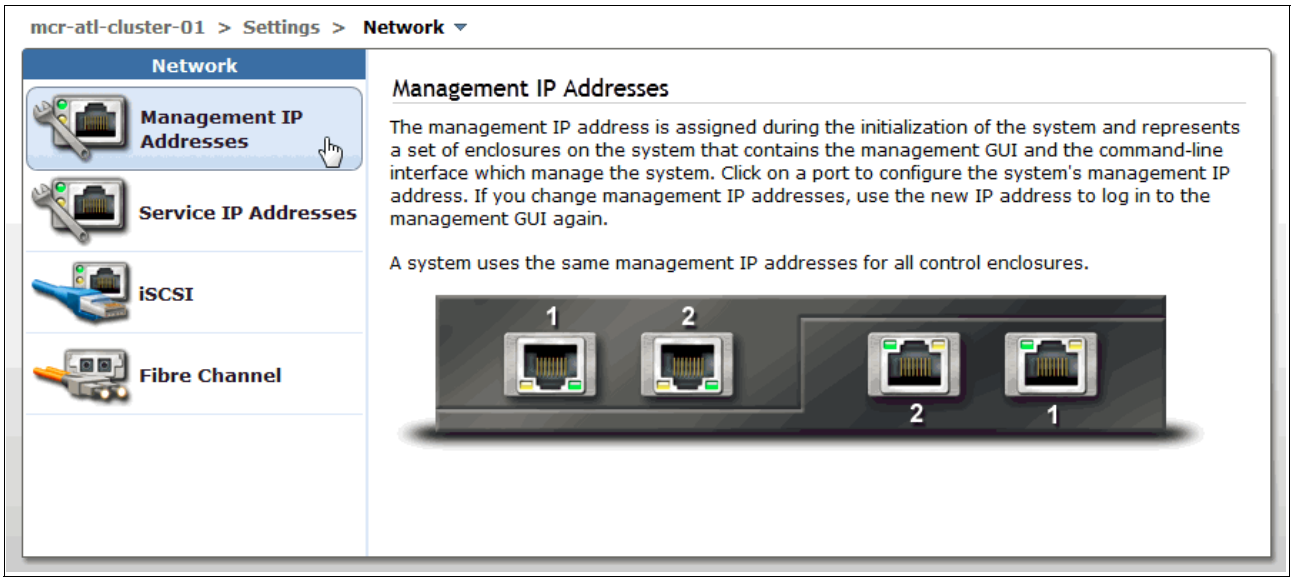

Figure 3-98 Network panel

Management IP addresses

The Management IP address is the IP address of the system and is configured during initial setup. The address can be an IPv4 address, IPv6 address, or both. The Management IP address is logically assigned to Ethernet port 1 of each node canister, which allows for node canister failover.

Another Management IP address can be logically assigned to Ethernet port 2 of each node canister for more fault tolerance. If the Management IP address is changed, use the new IP address to log in to the Management GUI again. Click **Management IP Addresses** and then click the port that you want to configure (the corresponding port on the partner node canister is also highlighted). Figure 3-99 shows Management IP Addresses configuration panel.

Figure 3-99 Management IP Addresses configuration panel

Service IP Addresses

Service IP addresses are used to access the Service Assistant. The address can be an IPv4 address, IPv6 address, or both. The Service IP addresses are configured on Ethernet port 1 of each node canister. Click **Service IP Addresses** and the select the **Control Enclosure** and **Node Canister** to configure. Figure 3-100 on page 142 shows the Service IP addresses configuration panel.

For more information, see 2.10.3, “Service Assistant tool” on page 71.

Figure 3-100 Service IP Addresses configuration panel

iSCSI connectivity

The IBM Storwize V5000 supports iSCSI connections for hosts. Click **iSCSI** and select the node canister to configure the iSCSI IP addresses. Figure 3-101 shows the iSCSI Configuration panel.

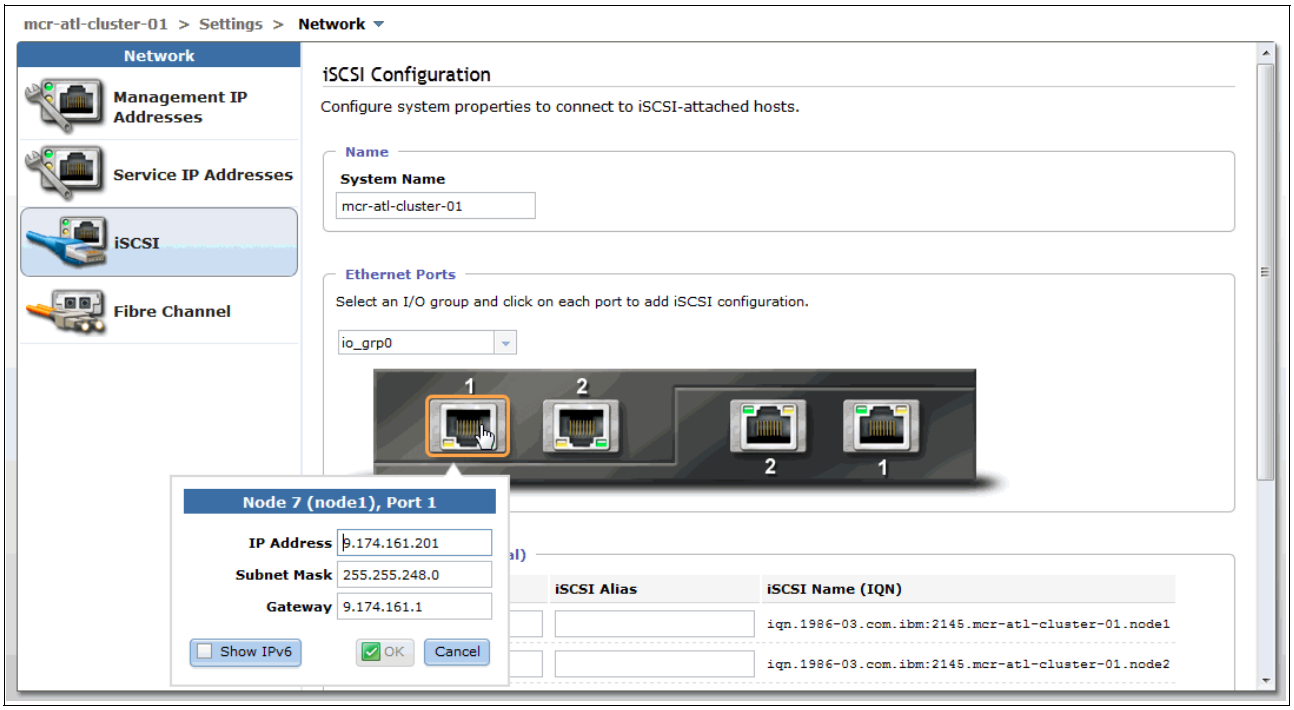

Figure 3-101 iSCSI Configuration panel

Fibre Channel connectivity

The Fibre Channel panel displays Fibre Channel connections that are established between the IBM Storwize V5000 node canisters, other storage systems, and hosts. Click **Fibre Channel** and select the required view from the **View connectivity for:** drop-down menu. Figure 3-102 shows the Fibre Channel panel with the All nodes, storage systems, and hosts option selected.

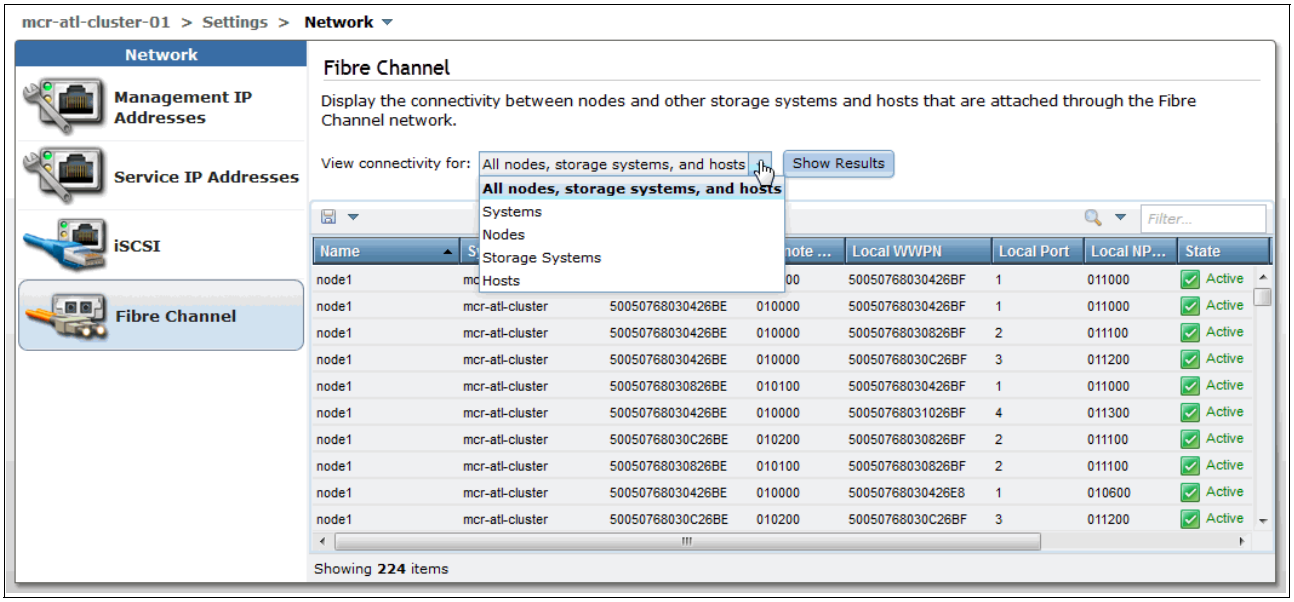

Figure 3-102 Fibre Channel panel

Support panel

Select **Support** in the Settings menu to open the Support panel. As shown in Figure 3-103, this panel provides access to the IBM support package, which is used by IBM to assist with problem determination. Click **Download Support Package** to access the wizard.

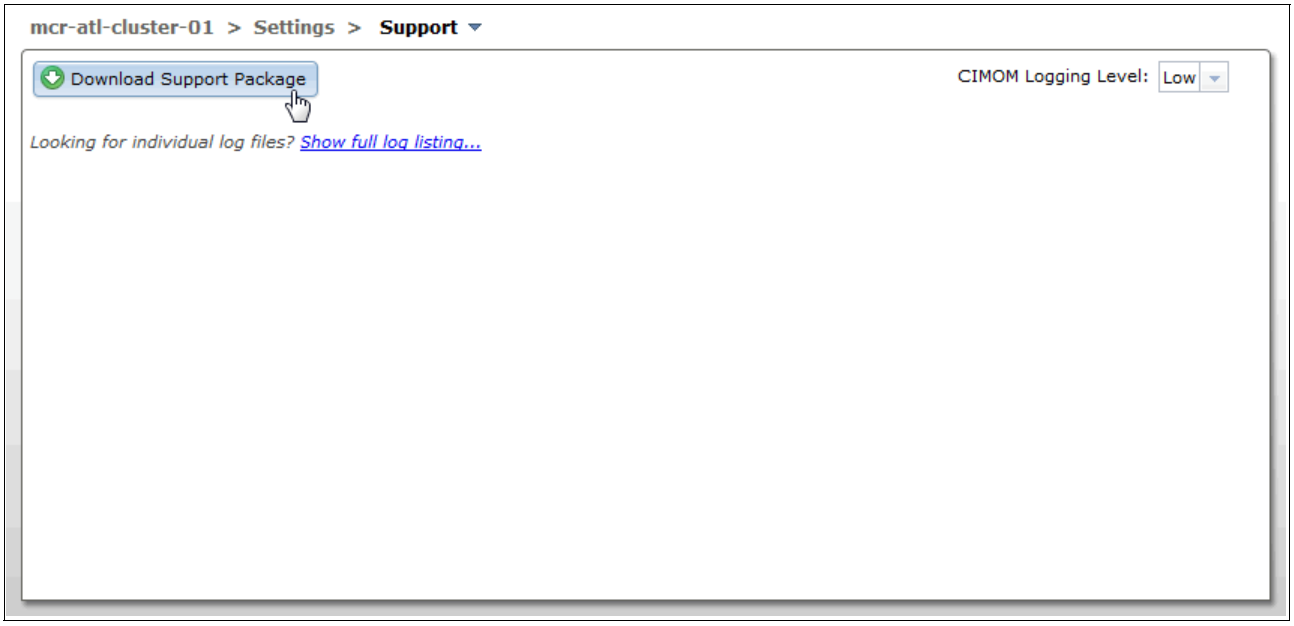

Figure 3-103 Support panel

Download Support Package wizard

The Download Support Package wizard provides a selection of various package types. IBM support provides direction on package type selection as required. To download the package, select the type and click **Download**. The output file can be saved to the user's workstation. Figure 3-104 shows the Download Support Package wizard.

Figure 3-104 Download Support Package wizard

Show full log listing

The Support panel also provides access to the files that are on the node canisters, as shown in Figure 3-105. Click **Show full log listing...** to access the node canister files. To save a file to the user's workstation, select a file, right-click the file, and select **Download**. To change to the file listing to show the files on a partner node canister, select the node canister from the menu that is next to the panel filter.

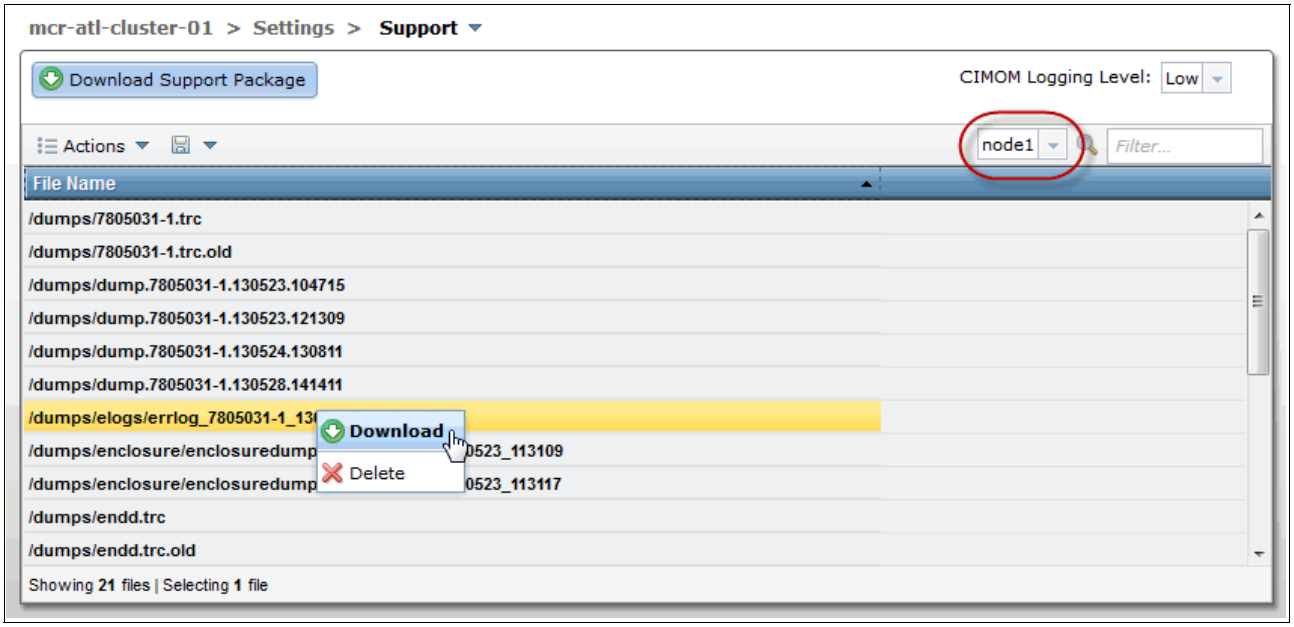

Figure 3-105 Full log listing

General panel

Select **General** in the Settings menu to open the General panel. This panel provides access to the Date and Time, Licensing, Upgrade Software, and GUI Preferences configuration panels.

Date and Time

Click **Data and Time** to configure the date and time manually or via a Network Time Protocol (NTP) server. Figure 3-106 shows the Date and Time function of the General panel.

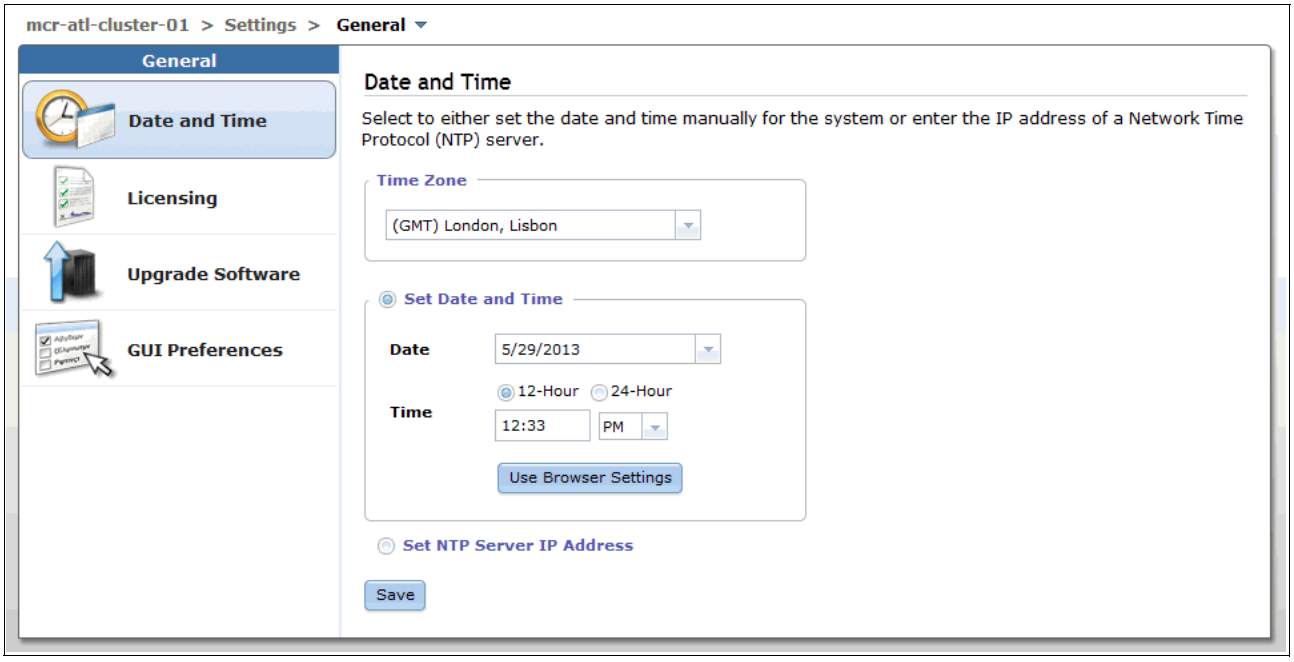

Figure 3-106 General panel

Licensing

The Licensing view shows the current system licensing. The IBM Storwize V5000 uses the same honor-based licensing as the Storwize V7000, which is based on per enclosure licensing.

The following optional licenses are available:

- FlashCopy
- Remote Copy
- Easy Tier
- External Virtualization

Figure 3-107 on page 146 shows the Update License panel within the General panel. In this example, two enclosures are licensed for FlashCopy, Remote Copy, and Easy Tier, while External Virtualization is licensed for 10 external disk trays.

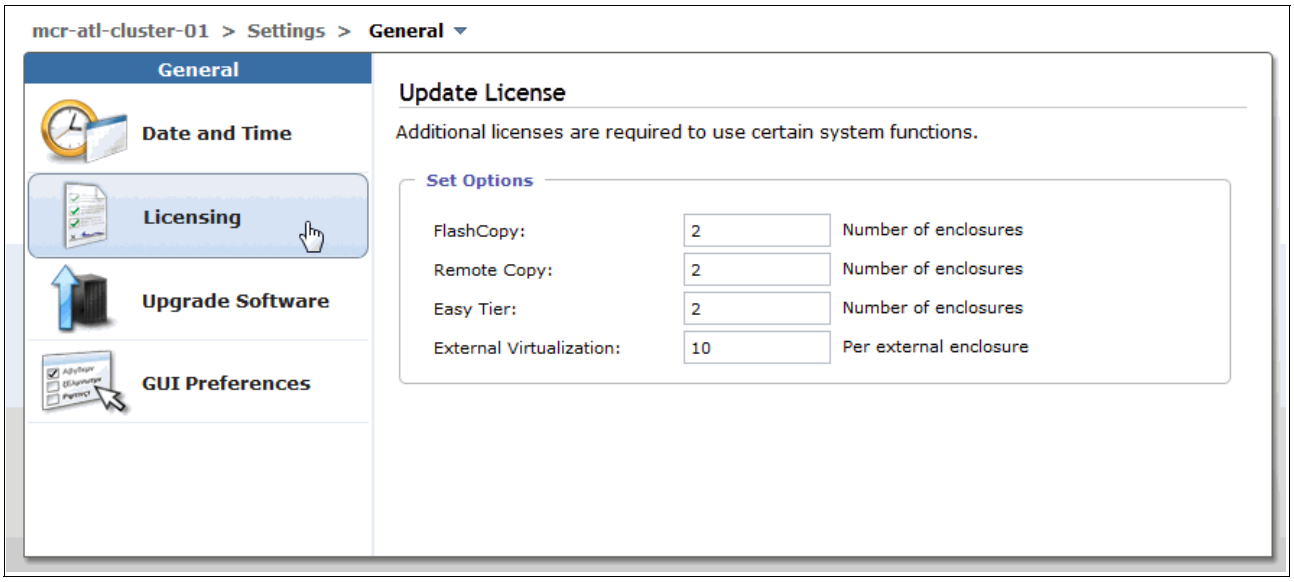

Figure 3-107 Update License panel

Upgrade Software panel

IBM recommends that you use the latest version of software. The Upgrade Software panel shows the current software level. If the system is connected to the Internet, it connects to the IBM upgrade server to check whether the current level is the latest. If an update is available, a direct link to the code is provided to the make code download process easier.

To upgrade the code, the IBM Storwize V5000 Code and the IBM Storwize V5000 Upgrade Test Utility must be downloaded. After the files are downloaded, it is best to check the MD5 checksum to ensure that the files are sound. Read the release notes, verify compatibility, and follow all IBM recommendations and prerequisites.

To upgrade the software of the IBM Storwize V5000, click **Launch Upgrade Wizard**. After the upgrade starts, an Abort option is shown that can be used to stop the upgrade process. Figure 3-108 on page 147 shows the Upgrade Software panel.

For more information, see 12.4, “Upgrading software” on page 580.

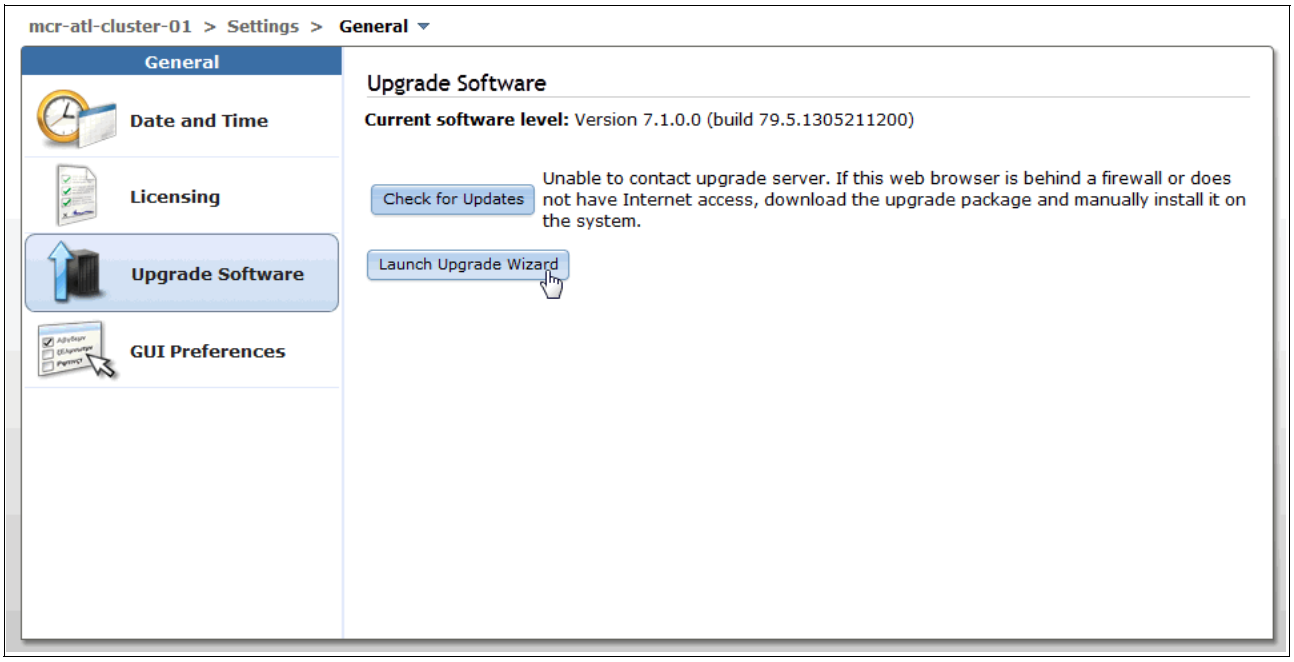

Figure 3-108 Upgrade machine code panel

GUI Preferences panel

By using the GUI Preferences panel (as shown in Figure 3-109), you can refresh GUI objects, restore default browser preferences, set table selection policy, and configure the Information Center web address.

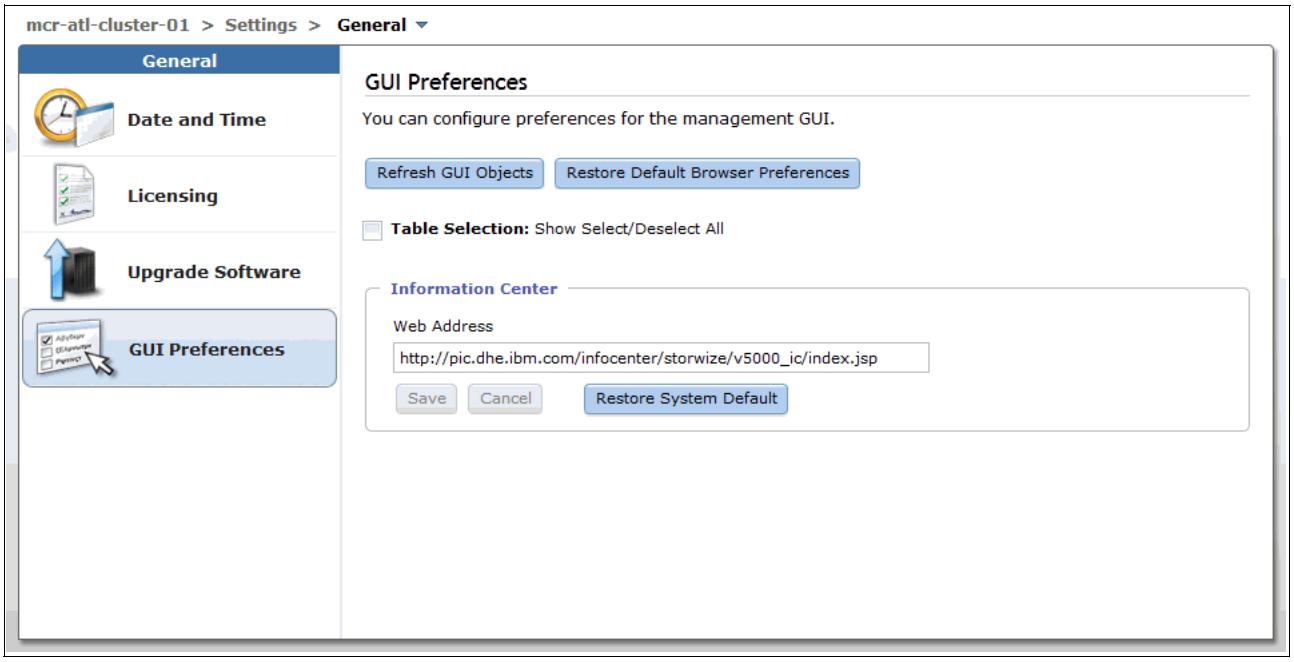

Figure 3-109 GUI Preferences panel

3.5 Management GUI help

This section provides information about the following methods that are available to get help while you use the IBM Storwize V5000 management GUI:

- IBM Storwize V5000 Information Center
- e-Learning modules
- Embedded panel help
- Question mark help
- Hover help
- IBM endorsed YouTube videos

3.5.1 IBM Storwize V5000 Information Center

The best source of information for the IBM Storwize V5000 is the Information Center. Click **Visit the Information Center** for direct access to the online version from the Overview panel, as shown in Figure 3-110.

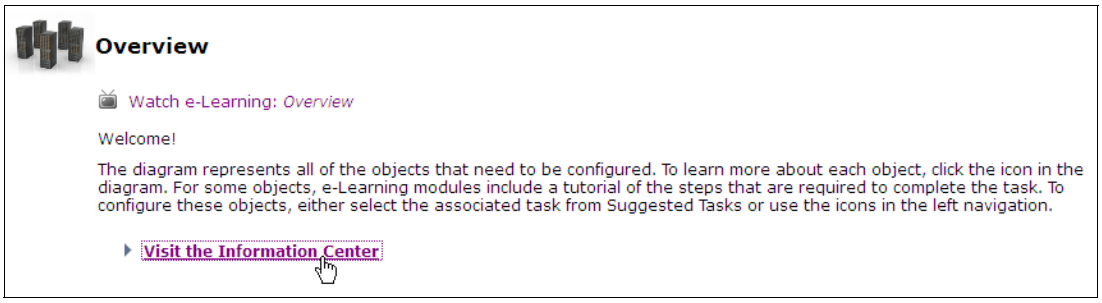

Figure 3-110 Overview panel

3.5.2 Watching an e-Learning video

The IBM Storwize V5000 provides embedded e-Learning videos to watch. The videos provide directions to complete various tasks. Click **Watch e-Learning** to start the video, as shown in Figure 3-111.

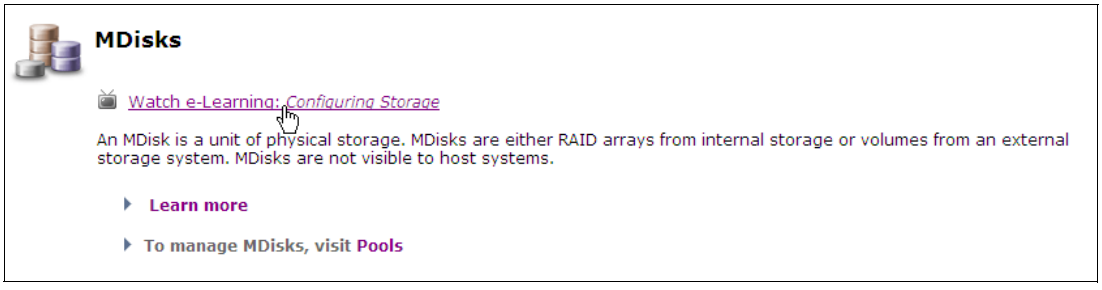

Figure 3-111 Watch e-Learning module

3.5.3 Learning more

The IBM Storwize V5000 provides embedded **Need Help** links to explain important concepts and panels. Click **Need Help** to open the information panel, as shown in Figure 3-112.

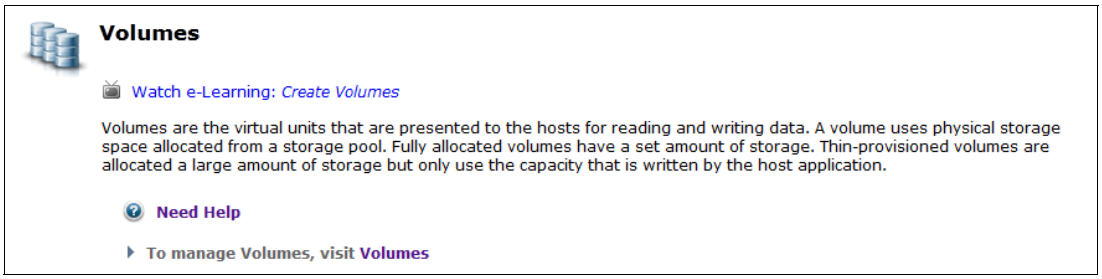

Figure 3-112 Learn more link

Figure 3-113 shows the information panel.

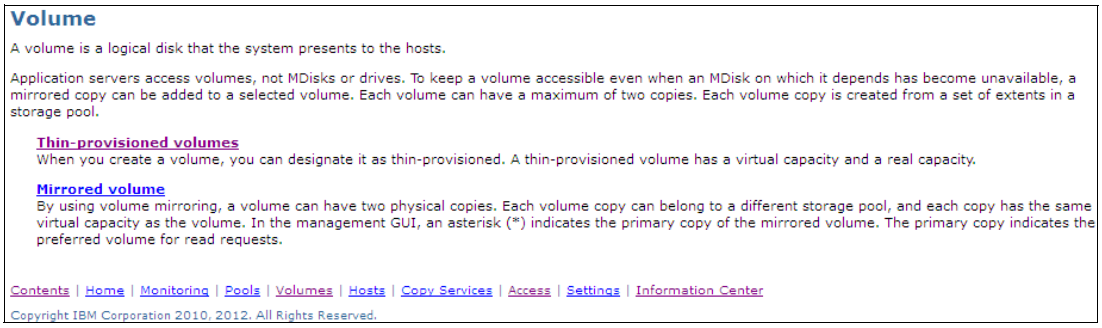

Figure 3-113 Information panel

3.5.4 Embedded panel help

The IBM Storwize V5000 provides embedded help that is available on each panel. Click **Help** to open the information panel, as shown in Figure 3-114.

Figure 3-114 Embedded panel help

Figure 3-115 shows the information panel that is opened from the embedded panel help. The information panel includes links to various other panels, including the Information Center.

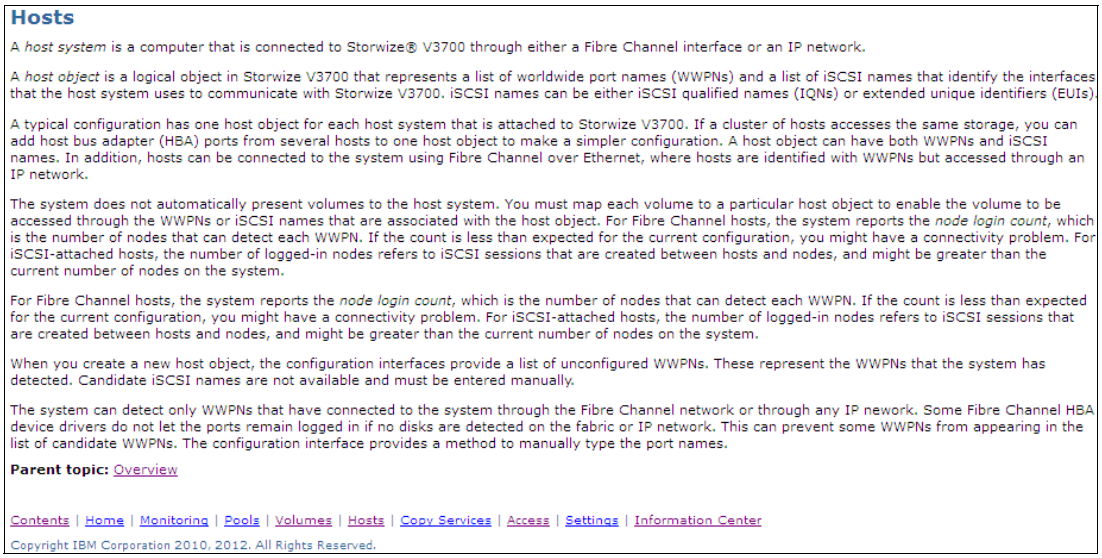

Figure 3-115 Information panel

3.5.5 Hidden question mark help

The IBM Storwize V5000 provides a hidden question mark help feature for some settings or items that are found in various configuration panels. This help feature is accessed by hovering next to an item where the question mark is shown and the help bubble is displayed, as shown in Figure 3-116.

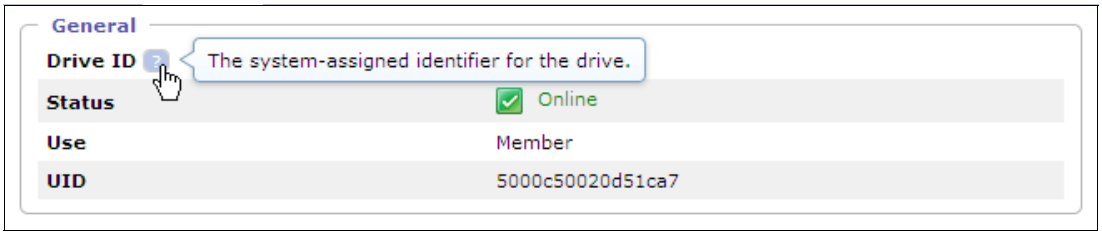

Figure 3-116 Hidden question mark help

3.5.6 Hover help

The IBM Storwize V5000 provides hidden help tags that are shown when you hover over various functions and items, as shown in Figure 3-117.

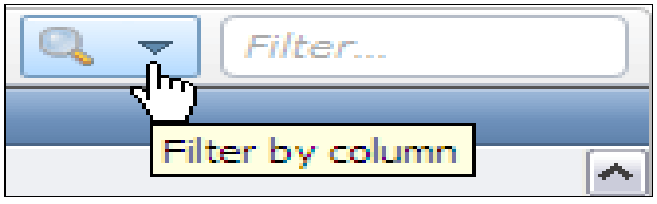

Figure 3-117 Hover help

3.5.7 IBM endorsed YouTube videos

IBM endorses various YouTube videos for the IBM storage portfolio. Client feedback suggests that these videos are a good tool to show management GUI navigation and tasks. Check for new videos from IBM Storage to find useful information at the IBM System Storage Channel at this website:

https://www.youtube.com/user/ibmstoragevideos

4

Host configuration

This chapter provides an overview on how to set up Open System hosts and the different methods that are available in the context of IBM Storwize V5000. It also describes how to use the IBM Storwize V5000 GUI to create hosts connections to access Storage Disk Subsystem volumes. For more information about Volume administration, see Chapter 5, “I/O Group basic volume configuration” on page 161.

This chapter includes the following topics:

- Host attachment overview
- Preparing the host operating system
- Configuring hosts on IBM Storwize V5000

4.1 Host attachment overview

A host system is an open-systems computer that is connected to a switch through a Fibre Channel (FC) or Internet Small Computer System Interface (iSCSI). Because IBM Storwize V5000 is geared towards small to medium scale data center storage solutions, a direct-attached Serial Attached SCSI (SAS) interface is also supported.

IBM Storwize V5000 supports the following host attachment protocols:

- 8 Gb Fibre Channel (FC) Protocol
- 6 Gb SAS Protocol
- 1 Gb iSCSI

In this chapter, we assume that your hosts are ready and attached to your FC and IP network, or directly attached if SAS Host Bus Adapters (HBAs) are used and that you completed the steps that are described in 2.9, “First-time setup”.

Follow basic switch and zoning recommendations and ensure that each host has at least two network adapters, that each adapter is on a separate network (or, at minimum, in a separate zone), and connections to all canisters exist. This setup assures four paths for failover and failback purposes. For SAS connections, ensure that each host has at least two SAS HBA connections to each IBM Storwize V5000 canister for resiliency purposes.

Before new volumes are mapped on the host of your choice, some preparation goes a long way towards ease of use and reliability. There are several steps that are required on a host system to prepare for mapping new IBM Storwize V5000 volumes. Use the System Storage Interoperation Center (SSIC) to check which code levels are supported to attach your host to your storage. SSIC is an IBM web tool that checks the interoperation of host, storage, switches, and multipathing drivers. For more information about IBM Storwize V5000 compatibility, see this website:

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

This chapter focuses on Windows and VMware. If you must attach any other hosts, for example, IBM AIX®, Linux, or even an Apple system, you can find the required information in the IBM Storwize V5000 Information Center at this website:

http://pic.dhe.ibm.com/infocenter/storwize/v5000_ic/index.jsp

4.2 Preparing the host operating system

In this section, we describe how to prepare Microsoft Windows and VMware host side attachment that is required to use an IBM Storwize V5000 with FC, iSCSI, or SAS connectivity.

4.2.1 Windows 2008 R2: Preparing for FC attachment

Complete the following steps to prepare a Windows 2008 (R2) host to connect to an IBM Storwize V5000 by using FC:

1. Make sure that the latest operating system Service Pack, updates, and hotfixes are applied to your Microsoft server.
2. Use the latest firmware and driver levels on your host system.
3. Install the HBAs on the Windows server by using the latest BIOS.
4. Connect the FC Host Adapter ports to the switches by using FC cables.
5. Configure the switches (SAN Zoning).
6. Configure the HBA parameters, if necessary.
7. Set the Windows timeout value.
8. Install the multipath Driver Device Module software.

Downloading and installing the supported drivers and firmware

Install a supported HBA driver for your configuration. Use the Windows Device Manager or vendor tools, such as, SANsurfer for QLogic product, HBAnyware for Emulex, or Brocade HBA Software Installer to install the driver. Also, check and update the BIOS (firmware) level of the HBA by using the tools that were provided by manufacturer. Always check the readme file to see whether there are Windows registry parameters that should be set for the HBA driver.

Configuring Brocade HBAs for Windows

This section applies to Windows hosts that have Brocade HBAs installed. After the device driver and firmware are installed, you must configure the HBAs. To perform this task, use the Brocade host connectivity manager (HCM) software or reboot into the HBA BIOS, load the adapter defaults, and set the following values:

- Host Adapter BIOS: Disabled (unless the host is configured for SAN Boot)
- Queue depth: 4

Configuring QLogic HBAs for Windows

This section applies to Windows hosts that have QLogic HBAs installed.

After the device driver and firmware are installed, you must configure the HBAs. To complete this task, use the QLogic SANsurfer software or reboot into the HBA BIOS, load the adapter defaults, and set the following values:

- Host Adapter BIOS: Disabled (unless the host is configured for SAN Boot)
- Adapter Hard Loop ID: Disabled
- Connection Options: 1 (point-to-point only)
- Logical Unit Numbers (LUNs) Per Target: 0
- Port Down Retry Count: 15

Configuring Emulex HBAs for Windows

This section applies to Windows hosts that have Emulex HBAs installed.

After the device driver and firmware are installed, you must configure the HBAs. To complete this task, use the Emulex HBAnyware software or reboot into the HBA BIOS, load the defaults, and set topology to 1 (10F_Port Fabric).

Setting the Windows timeout value

For Windows hosts, the disk I/O timeout value should be set to 60 seconds as an overall rule, but you must also check the recommended guidelines for your application. To verify this setting, complete the following steps:

1. Click **Start** → **Run**.
2. In the window, enter `regedit` and press Enter.
3. In the registry editor, search for the `HKEY_LOCAL_MACHINE\System\CurrentControlSet\services\Disk\TimeOutValue` key.
4. Confirm that the value for the key is 60 (decimal value), and, if necessary, change the value to 60, as shown in Figure 4-1.

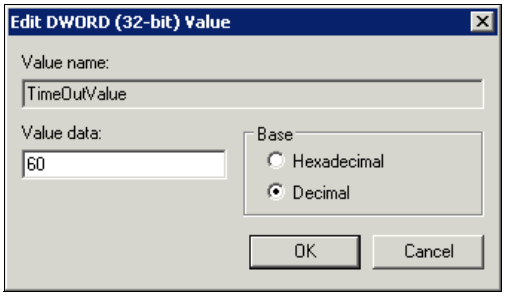

Figure 4-1 Windows timeout value

Installing Microsoft MPIO multipathing software

Microsoft Multipath Input/Output (MPIO) solutions are designed to work with device-specific modules (DSMs) that are written by vendors. The MPIO driver package does not form a complete solution on its own. By using this joint solution, the storage vendors can design device-specific solutions that are tightly integrated with the Microsoft Windows operating system. MPIO in Microsoft Windows 2008 is a DSM that is designed to work with Storage Arrays that support the Asymmetric Logical Unit Access (ALUA) control model (active-active Storage Controllers).

The intent of MPIO is to provide better integration of a multipath storage solution with the operating system. It also allows the use of multipath in the SAN infrastructure during the boot process for SAN Boot hosts.

To install MPIO on a computer that is running Microsoft Windows Server 2008, complete the following steps:

1. Open Server Manager by clicking **Start** → **Administrative Tools** → **Server Manager**.
2. In the Features area, click **Add Features**.
3. Select **MPIO** from the list of available features. Click **Next**.
4. Review and confirm the installation selections and click **Install**.

Before your ESXi host can discover the IBM Storwize V5000 storage, the iSCSI initiator must be configured and authentication might need to be done (depending on customer scenario), as shown Figure 4-2.

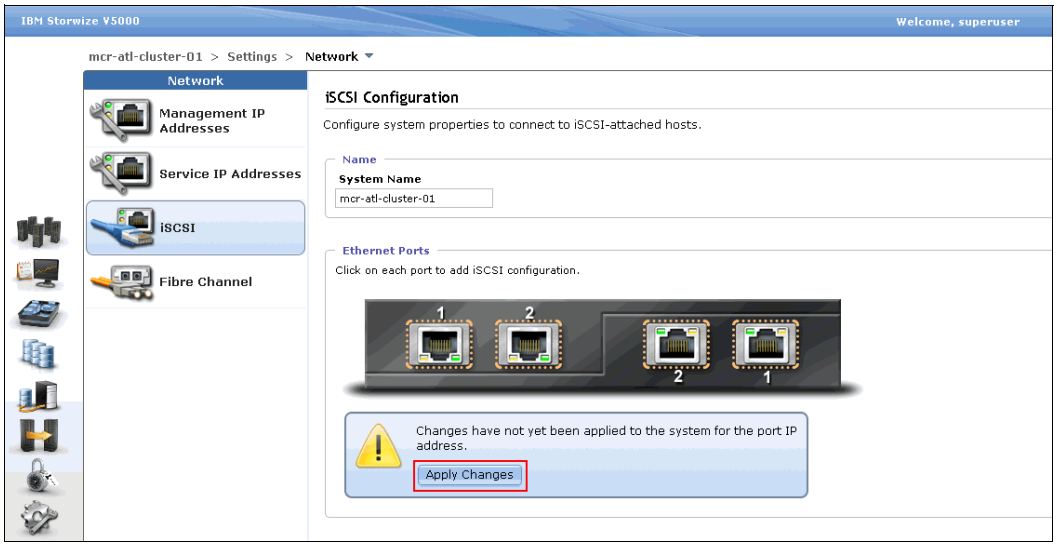

Figure 4-2 iSCSI IP Configuration

You can verify the network configuration by using the `vmkping` utility. If you must authenticate the target, you might need to configure the dynamic or static discovery address and target name of the Storwize V5000 in vSphere.

For more information about creating volumes and mapping them to a host, see Chapter 5, “I/O Group basic volume configuration” on page 161.

4.2.2 Creating SAS hosts

These steps provide guidance on how to setup hosts with SAS HBAs. Complete the following steps by using the IBM Storwize V5000 GUI to create an SAS host:

1. Click **SAS Host**. The Create Host window opens, as shown in Figure 4-3

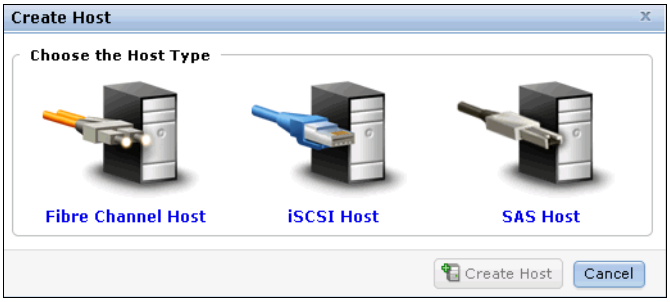

Figure 4-3 Create SAS host

2. Enter the host name and, from the drop-down menu, select the SAS worldwide port name (WWPN) or names that are associated with the host, as shown in Figure 4-4.

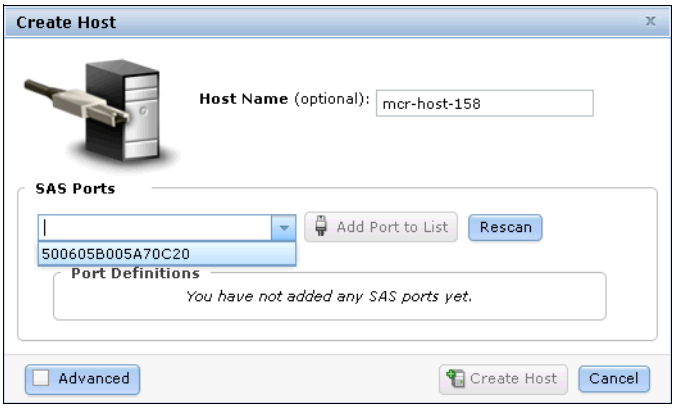

Figure 4-4 Available SAS WWPN or WWPNs

3. Click **Advanced** to expand the Advanced Settings options.
4. As shown Figure 4-5, select HP/UX or TPGS if you are creating one of these types of hosts. In our example, an HP/UX host is created with permissions to access volumes from I/O Group 1.

Important host setting: If this setting is set incorrectly, the host appears in volume mapping options that physically cannot be created. For more information, see Chapter 5, “I/O Group basic volume configuration” on page 161.

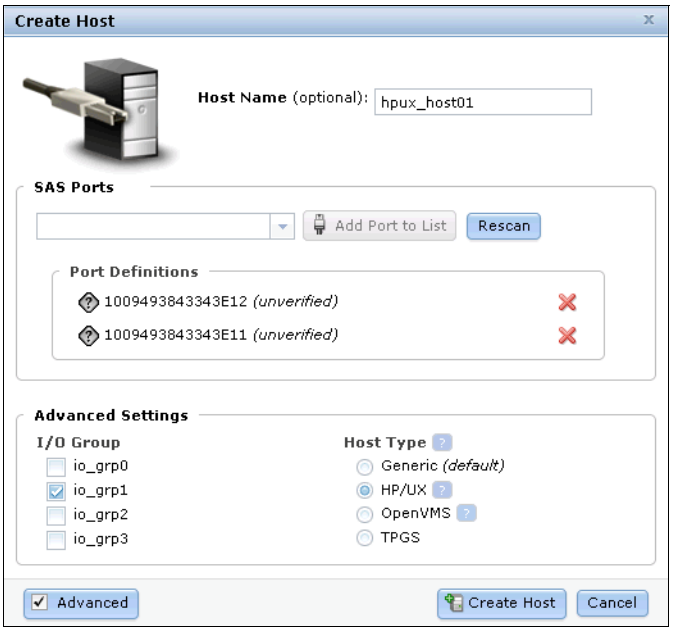

Figure 4-5 Creating HP/UX SAS host

5. Click **Create Host** to create the SAS Host object on IBM Storwize V5000.
6. Click **Close** when the task completes.

The IBM Storwize V5000 shows the host port WWPNs that are available if you prepared the hosts. If they do not appear in the list, scan for new disks in your operating system and click **Rescan** in the configuration wizard. If they still do not appear, check your physical connectivity and pay particular attention to the SAS cable orientation and repeat the scanning. For more information about hosts, see the Information Center at this website:

http://pic.dhe.ibm.com/infocenter/storwize/v5000_ic/index.jsp

The IBM Storwize V5000 is now configured and ready for SAS Host use. For advanced host and volume administration, see Chapter 8, “Advanced host and volume administration” on page 349.

5

I/O Group basic volume configuration

This chapter describes how to use the IBM Storwize V5000 to create a volume and map a volume to a host. A volume is a logical disk on the IBM Storwize V5000 that is provisioned out of a storage pool and is recognized by a host with an identifier UID field and a parameter list.

The first part of the chapter describes how to create volumes of different types and map them to the defined host.

The second part of this chapter describes how to discover those volumes. After you finish this chapter, your basic configuration is complete and you can store data on the IBM Storwize V5000.

For more information about advanced host and volume administration, such as, adding and deleting host ports and creating thin provisioned volumes, see Chapter 8, “Advanced host and volume administration” on page 349.

This chapter includes the following topics:

- Provisioning storage from IBM Storwize V5000 and making it available to the host
- Mapping a volume to the host
- Discovering the volumes from the host and specifying multipath settings

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.

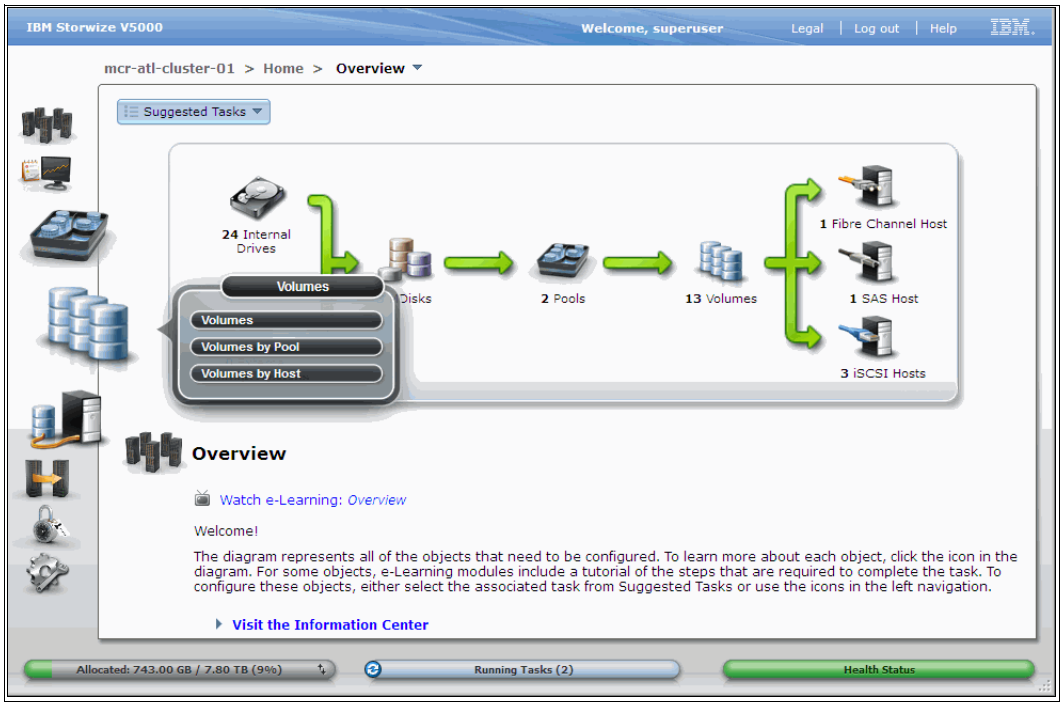

Figure 5-1 GUI Volumes option

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

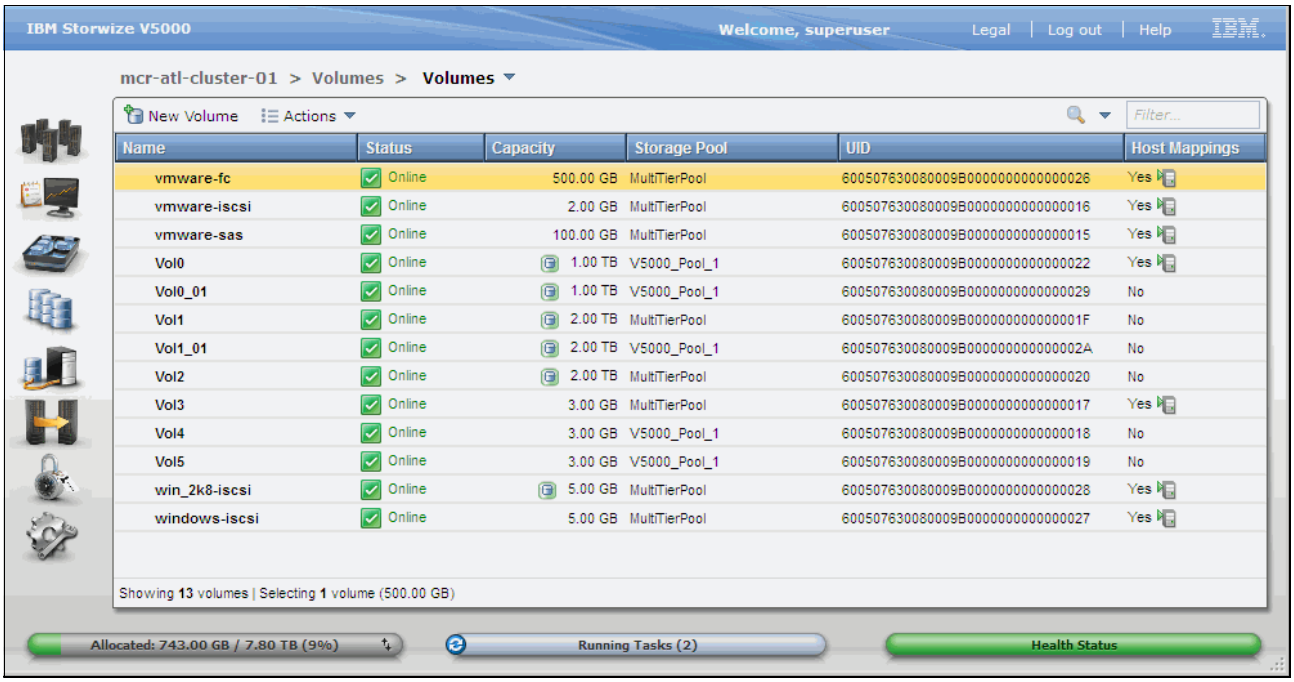

Figure 5-2 Volume listings

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.

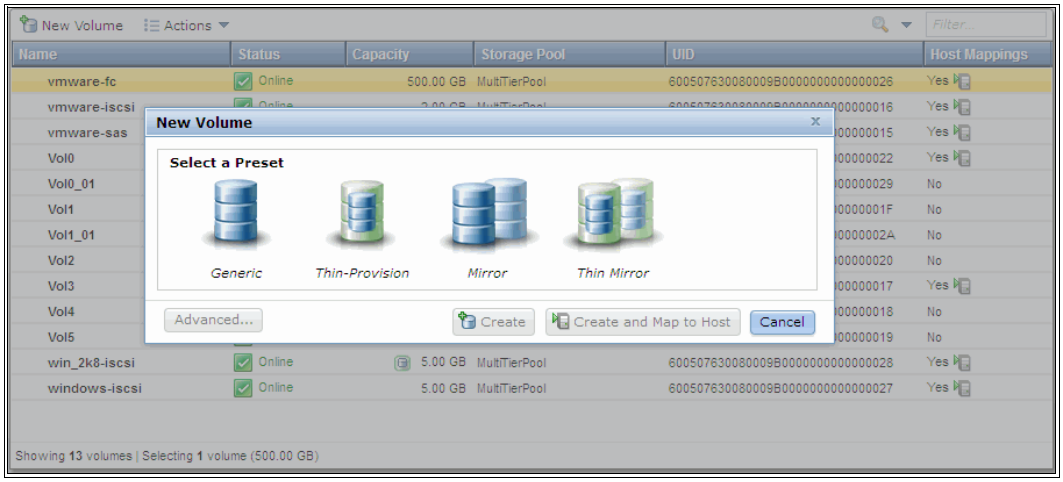

Figure 5-3 New Volume window

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.

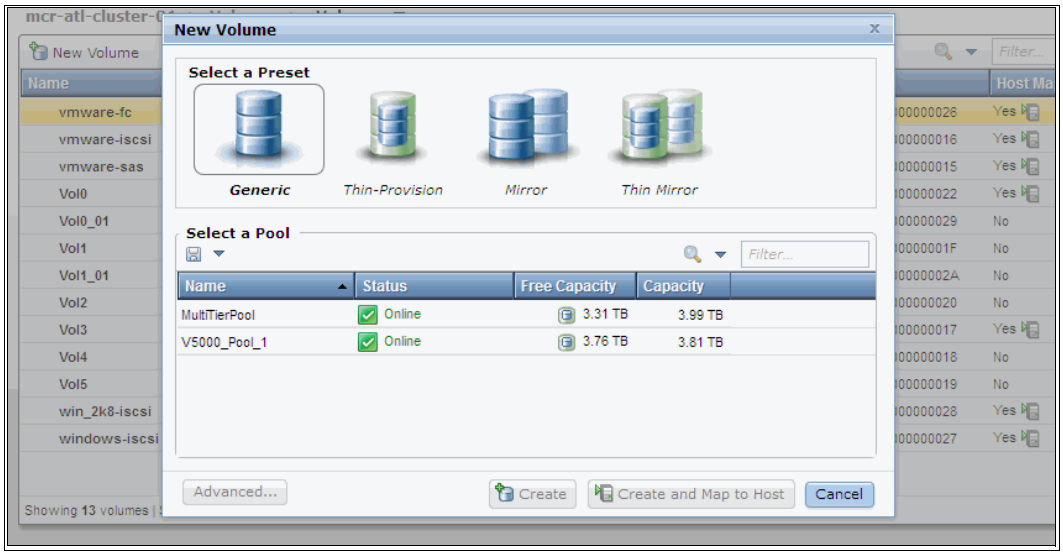

Figure 5-4 Provisioning a Generic volume

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.

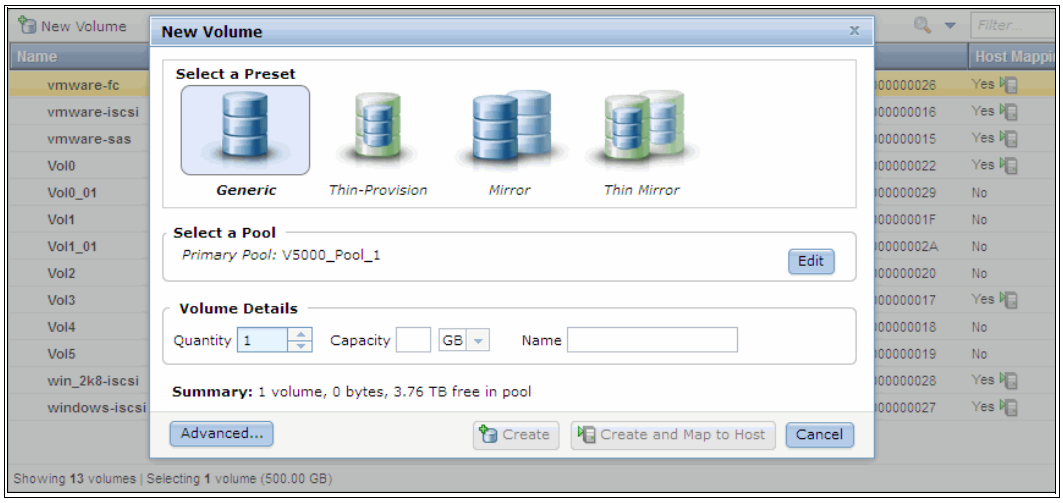

Figure 5-5 Pool selection

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.

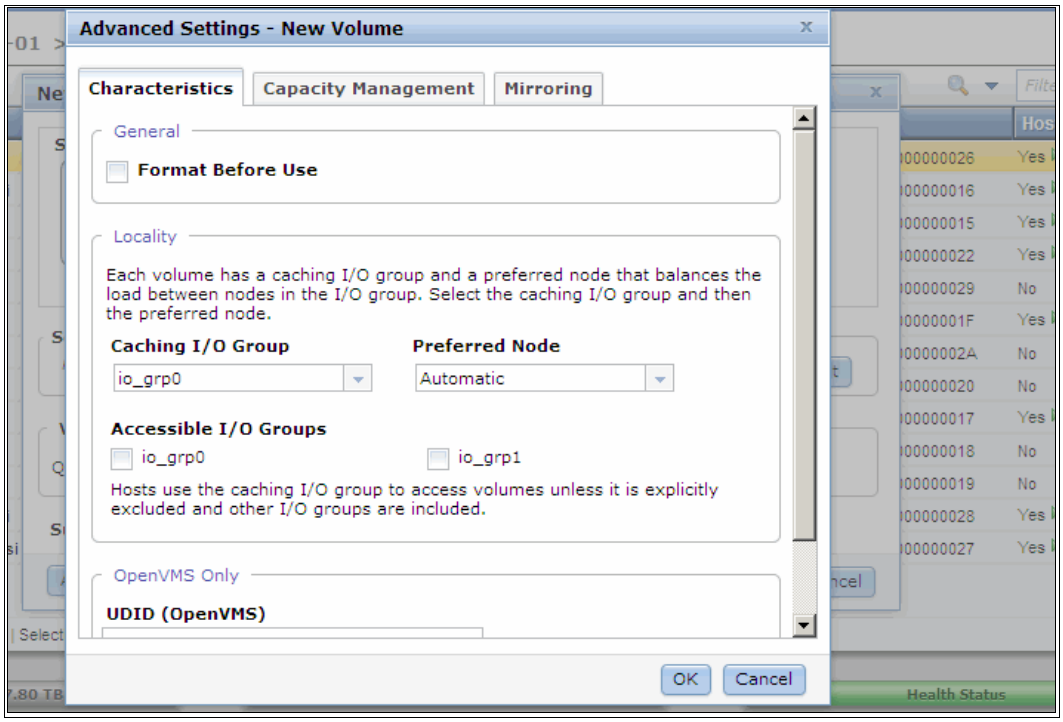

Figure 5-6 Generic advanced options

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.

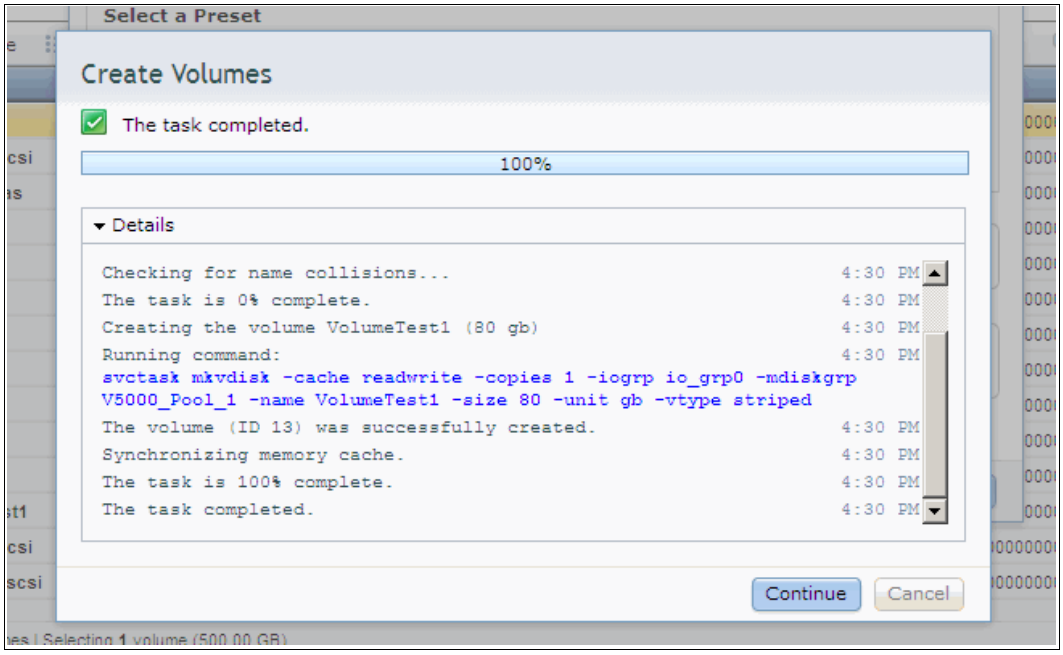

Figure 5-7 Volume creation complete

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.

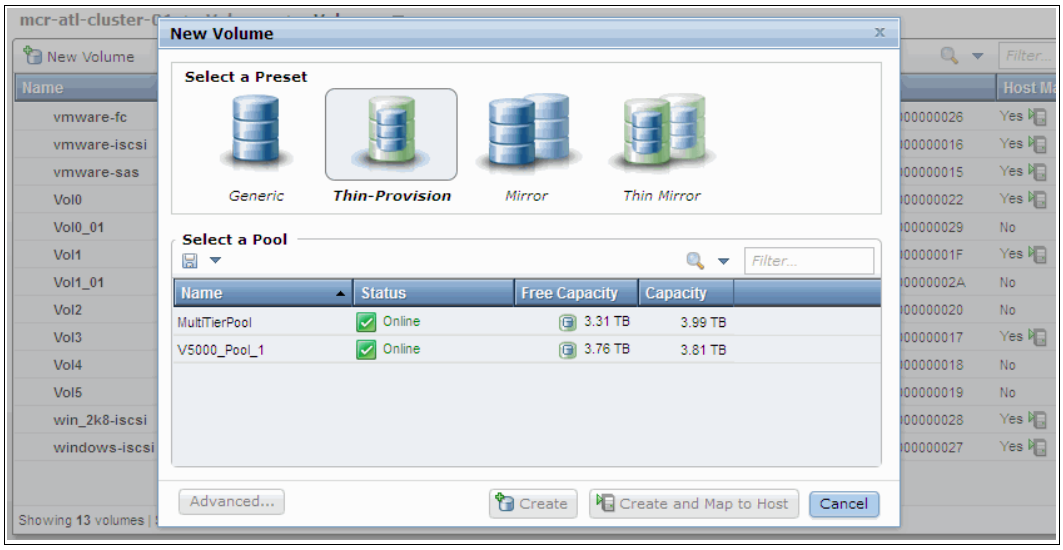

Figure 5-8 Creating a thin-provisioned volume

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.

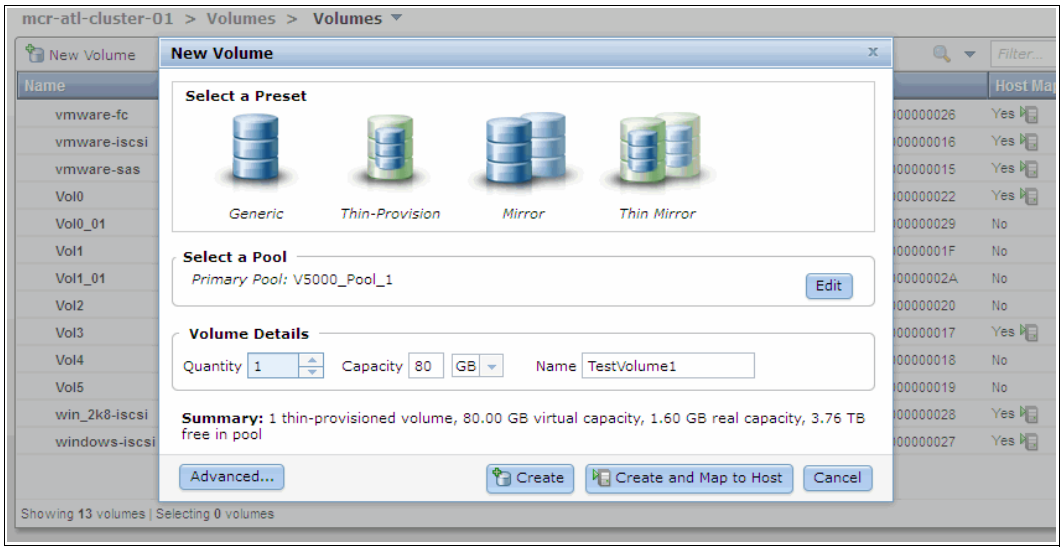

Figure 5-9 Enter the volume name and size

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.

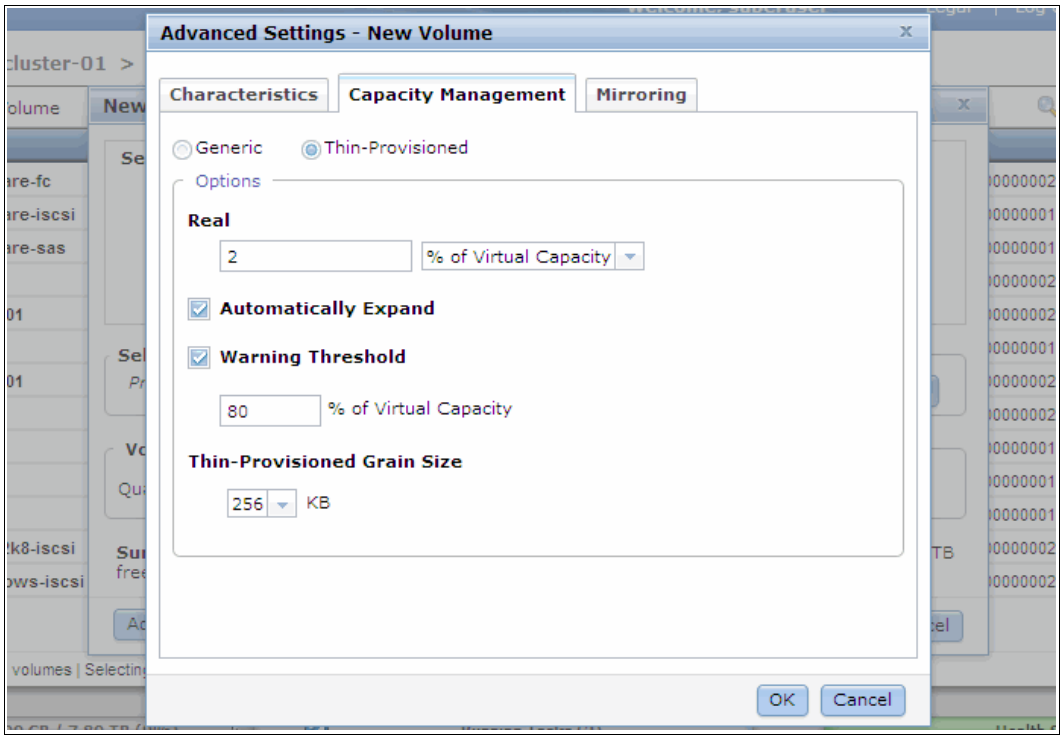

Figure 5-10 Advanced Settings: New Volume

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.

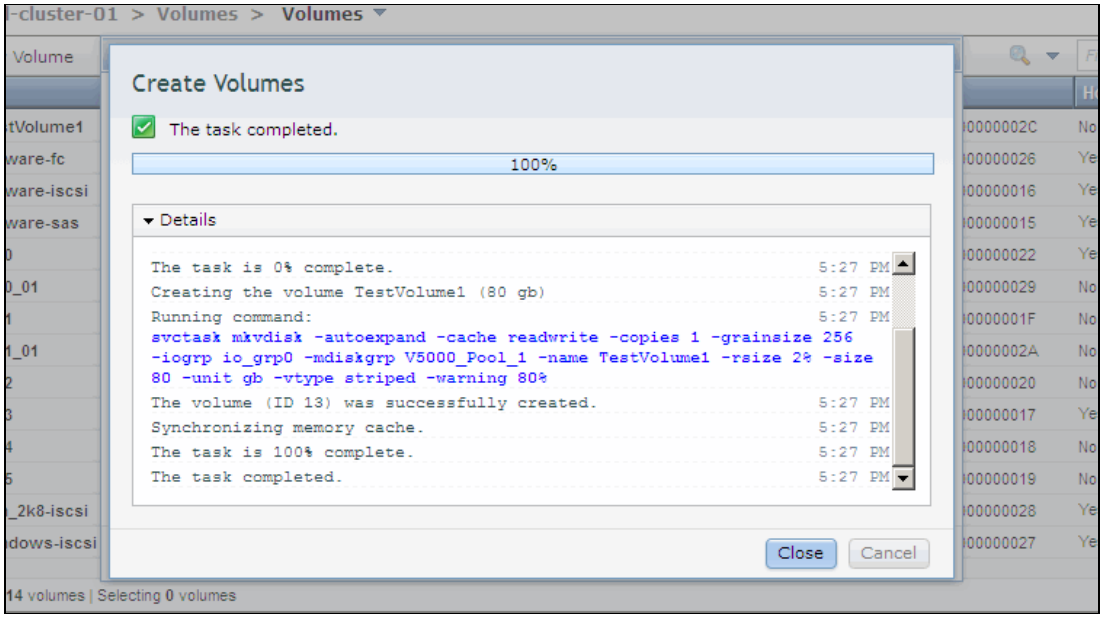

Figure 5-11 Thin volume creation complete

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.

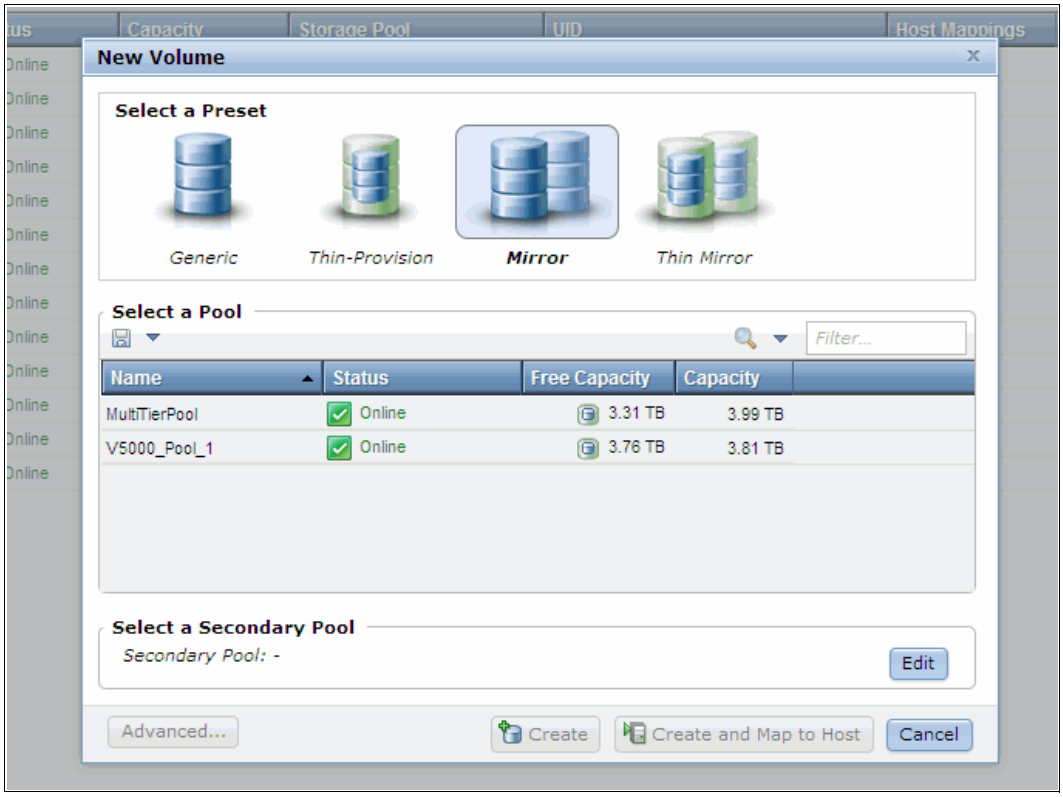

Figure 5-12 Create a mirrored volume

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.

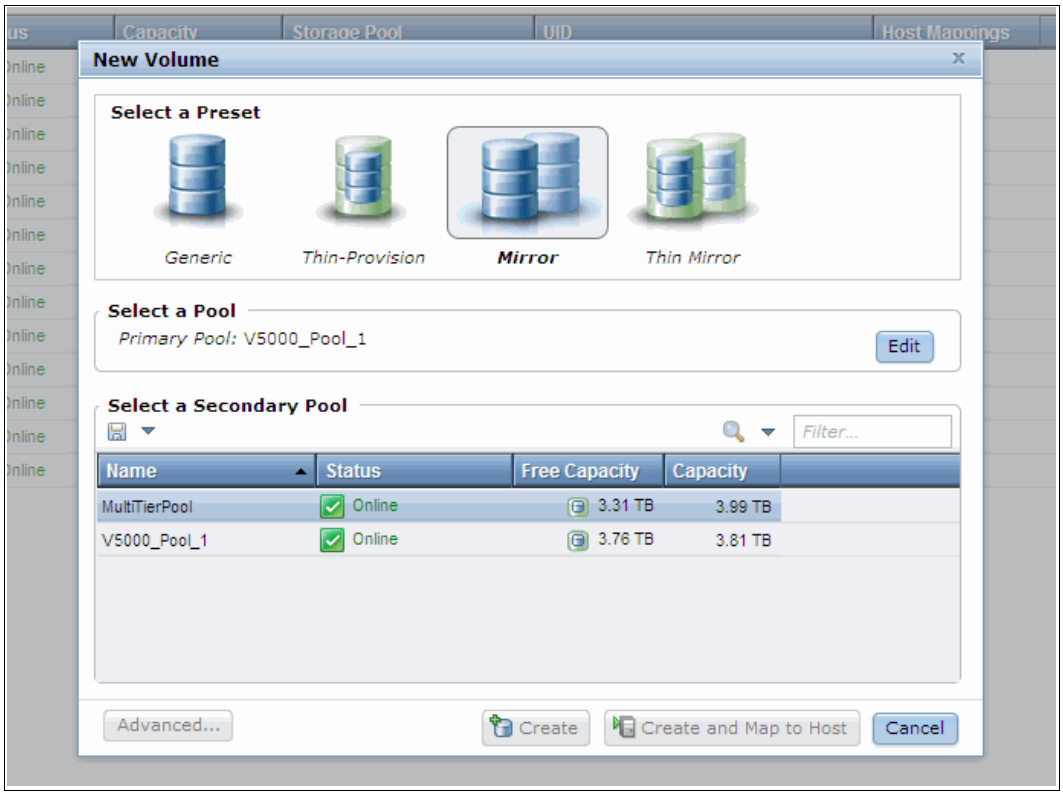

Figure 5-13 Selecting primary storage pool

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

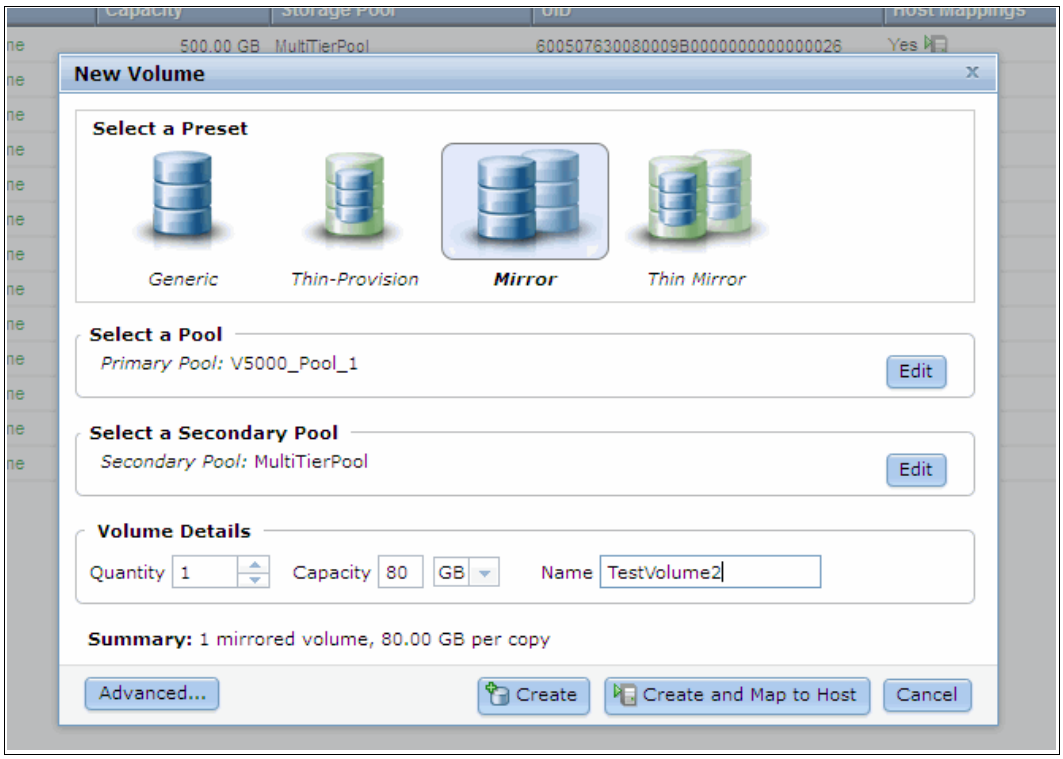

Figure 5-14 Select a secondary pool, volume name, and size

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.

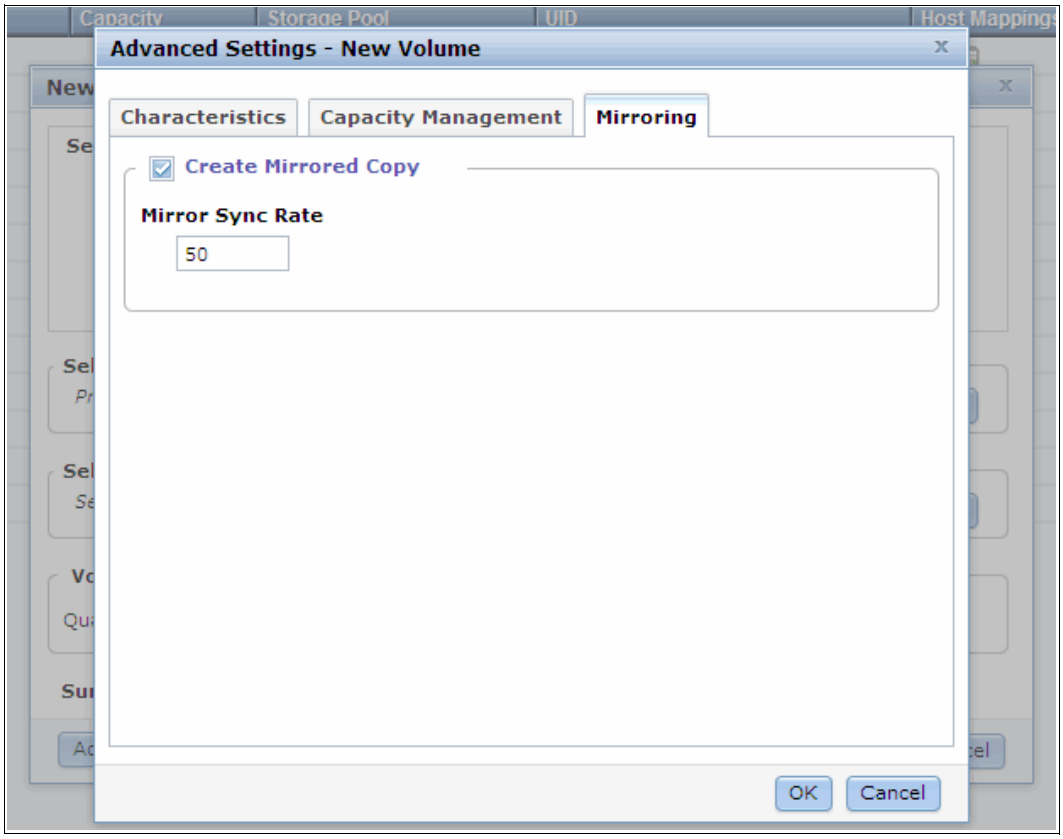

Figure 5-15 Advanced mirroring features

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.

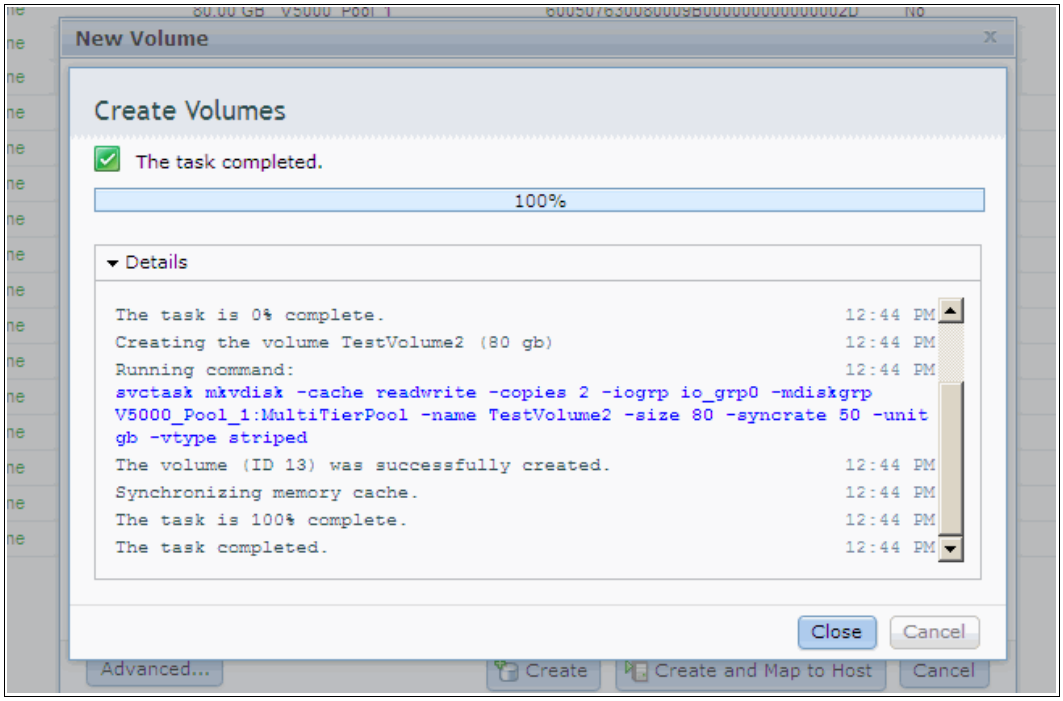

Figure 5-16 Mirrored volume task complete

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.

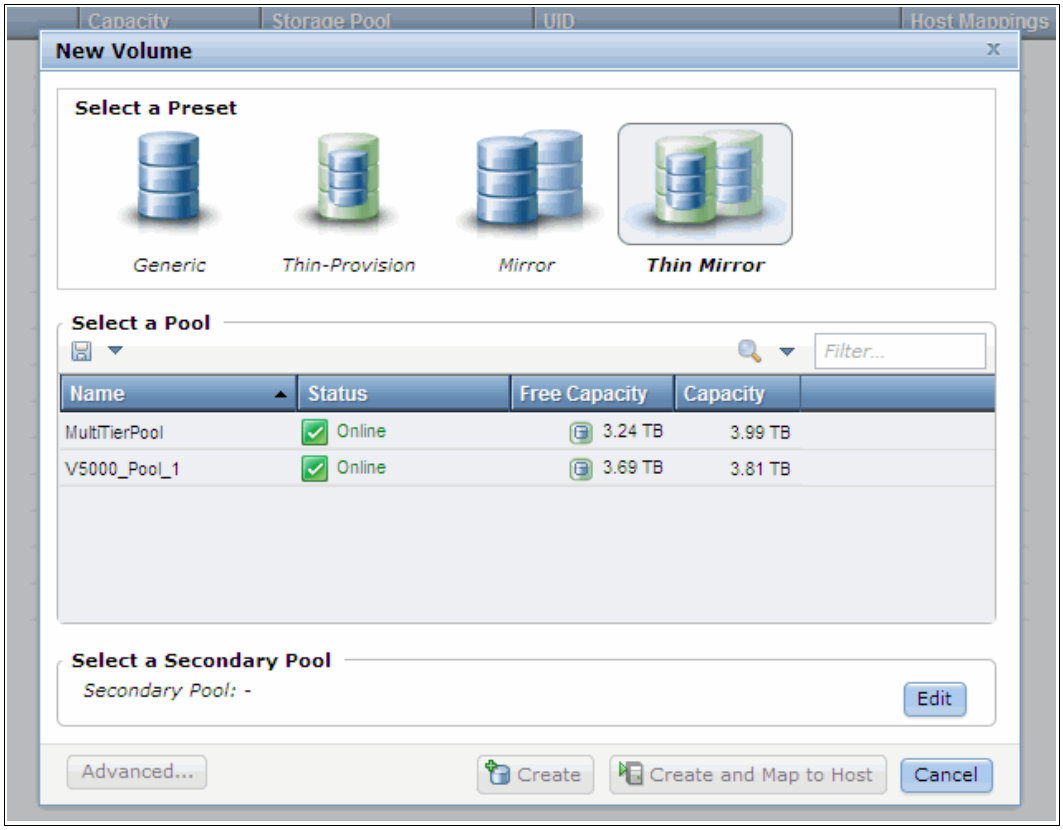

Figure 5-17 Create a Thin Mirror

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

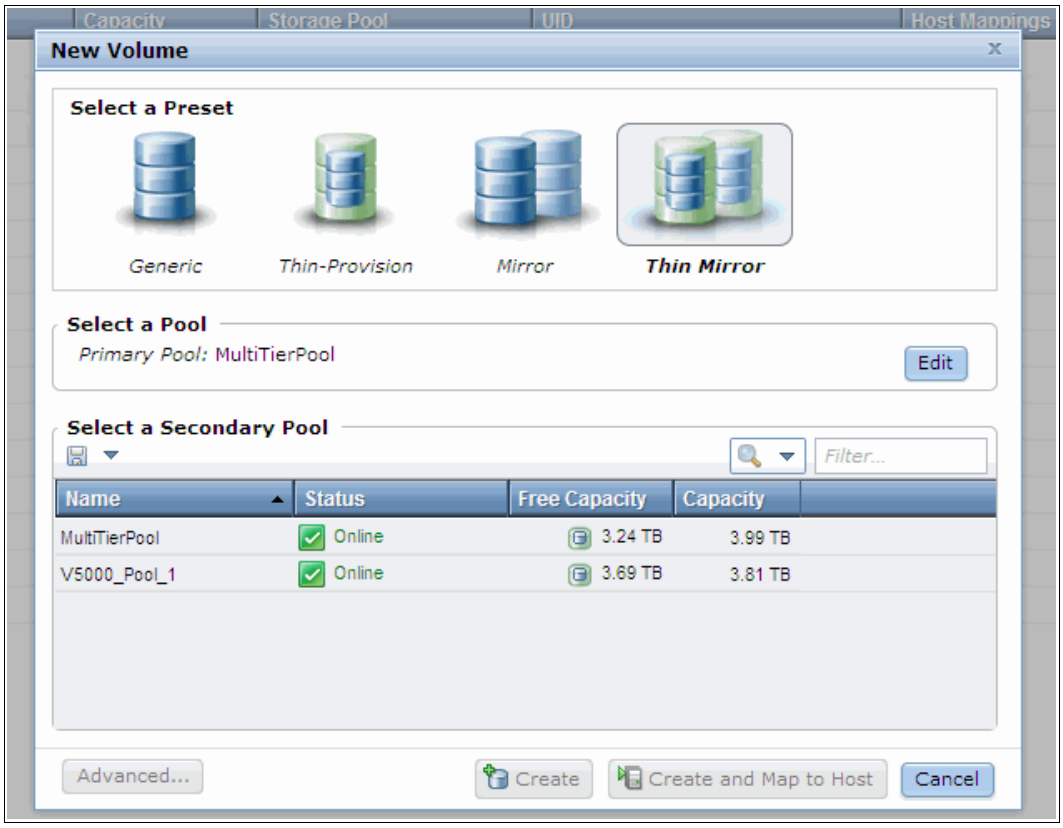

Figure 5-18 Selecting storage pools

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.

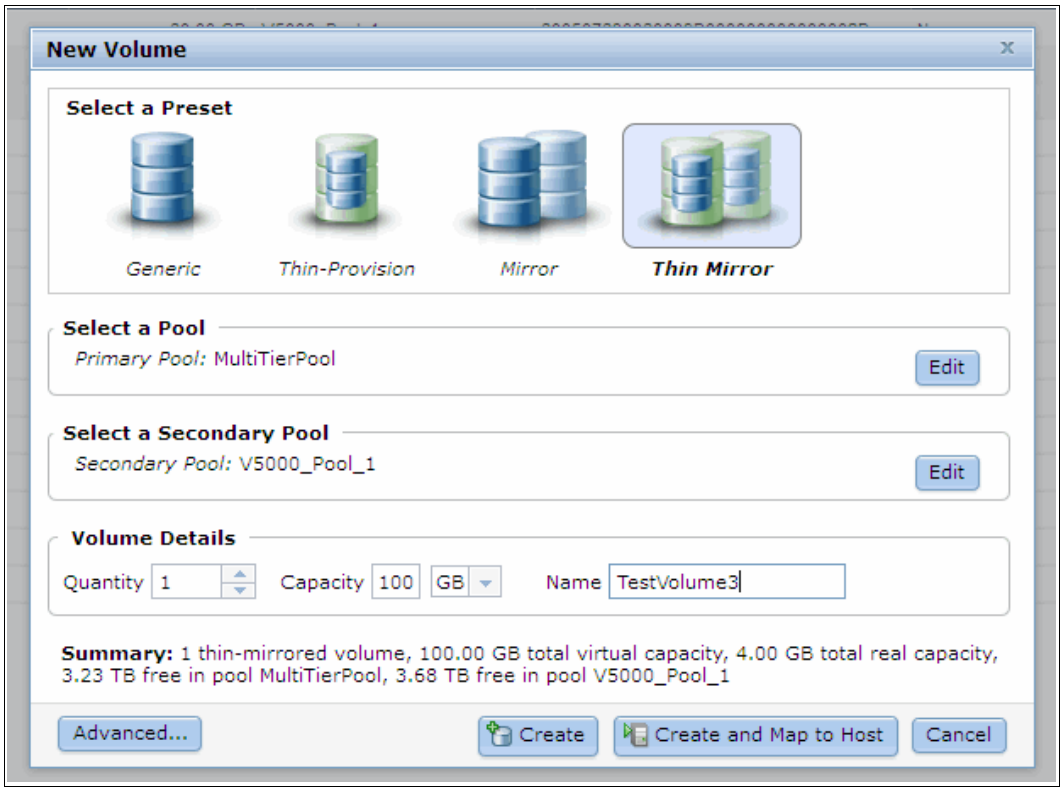

Figure 5-19 Enter a volume name and size

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.

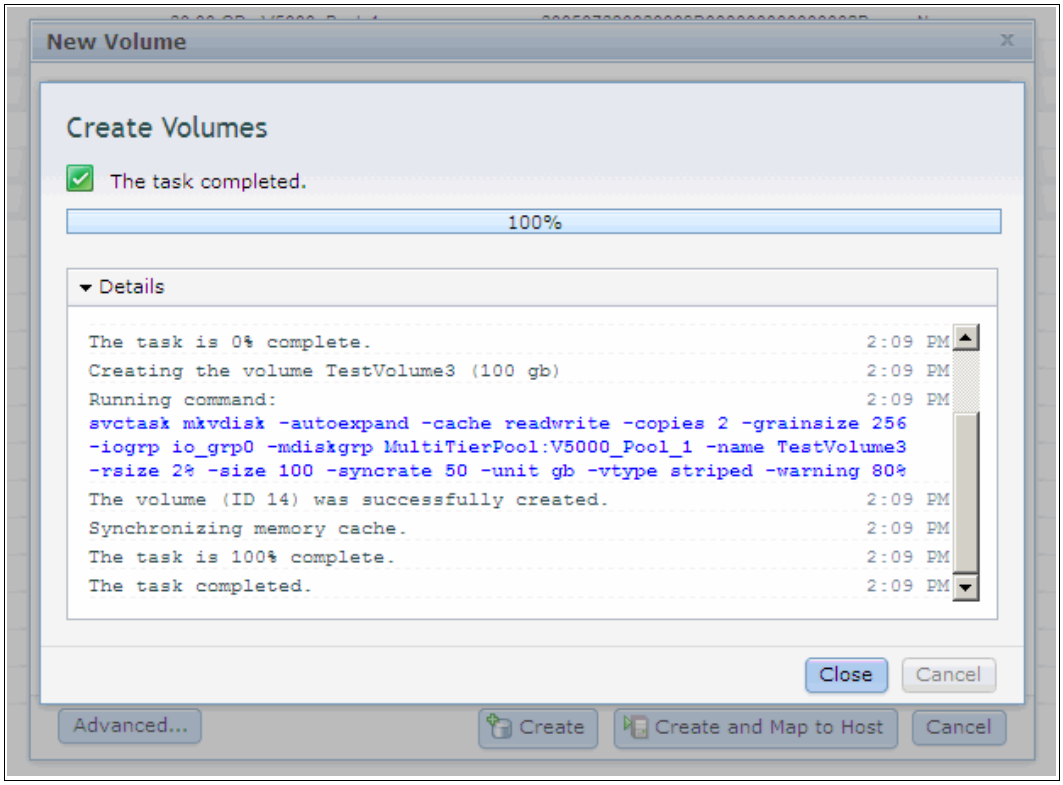

Figure 5-20 Thin Mirror Volume task complete

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.

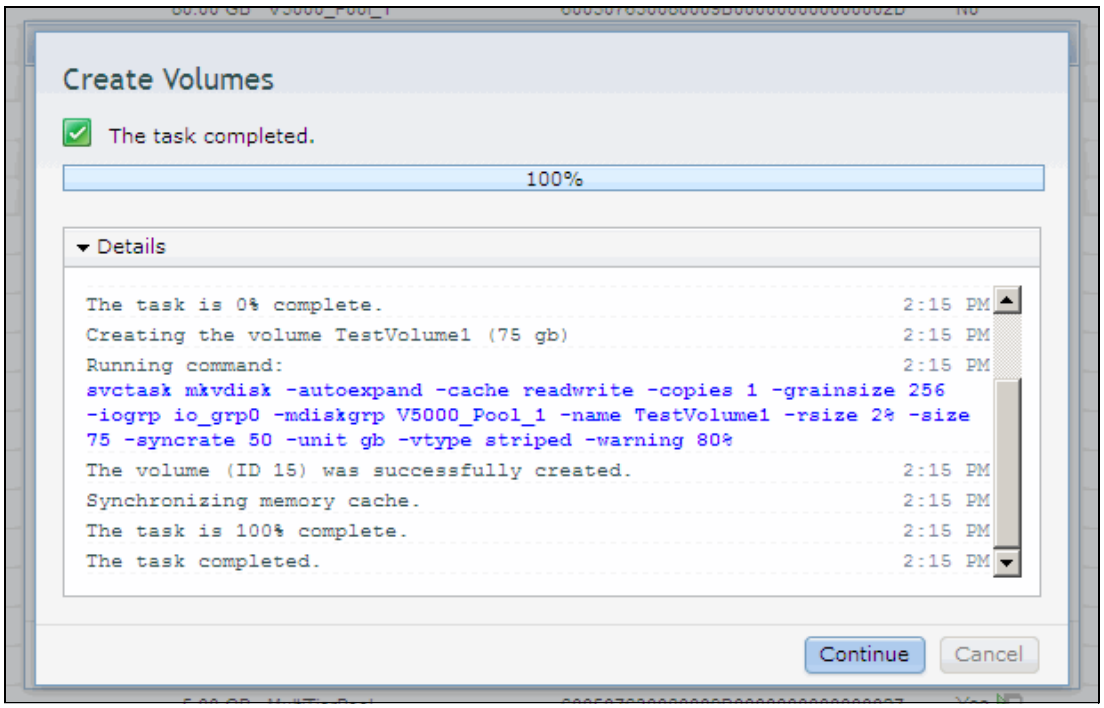

Figure 5-21 Continue to host mapping

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.

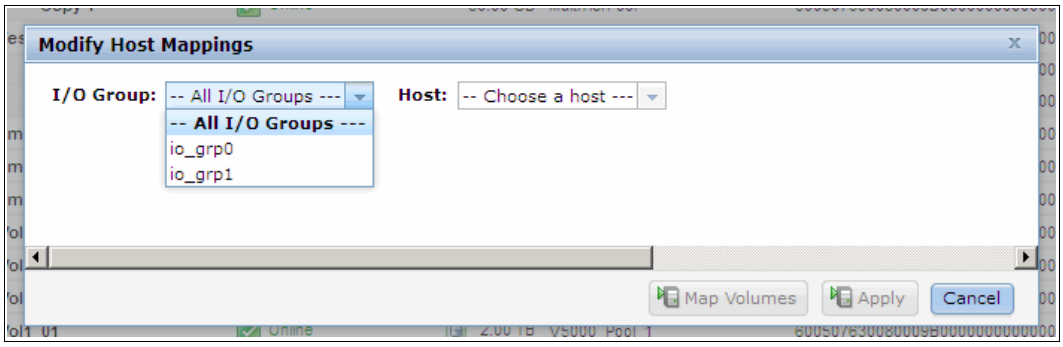

Figure 5-22 Choose an I/O Group

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.

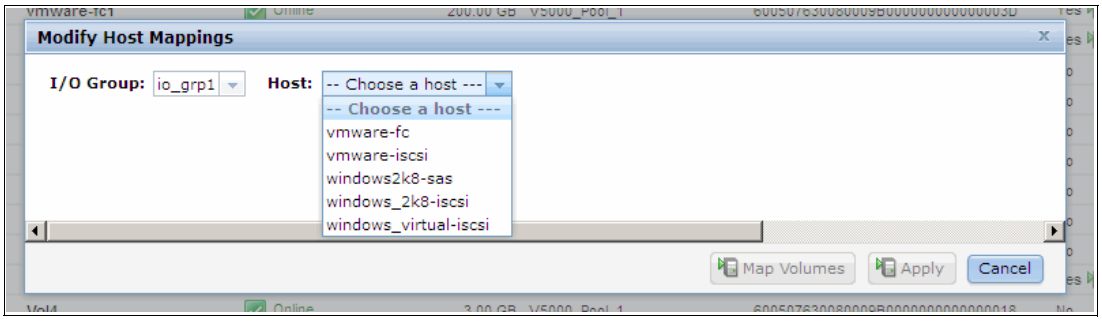

Figure 5-23 Select a host

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.

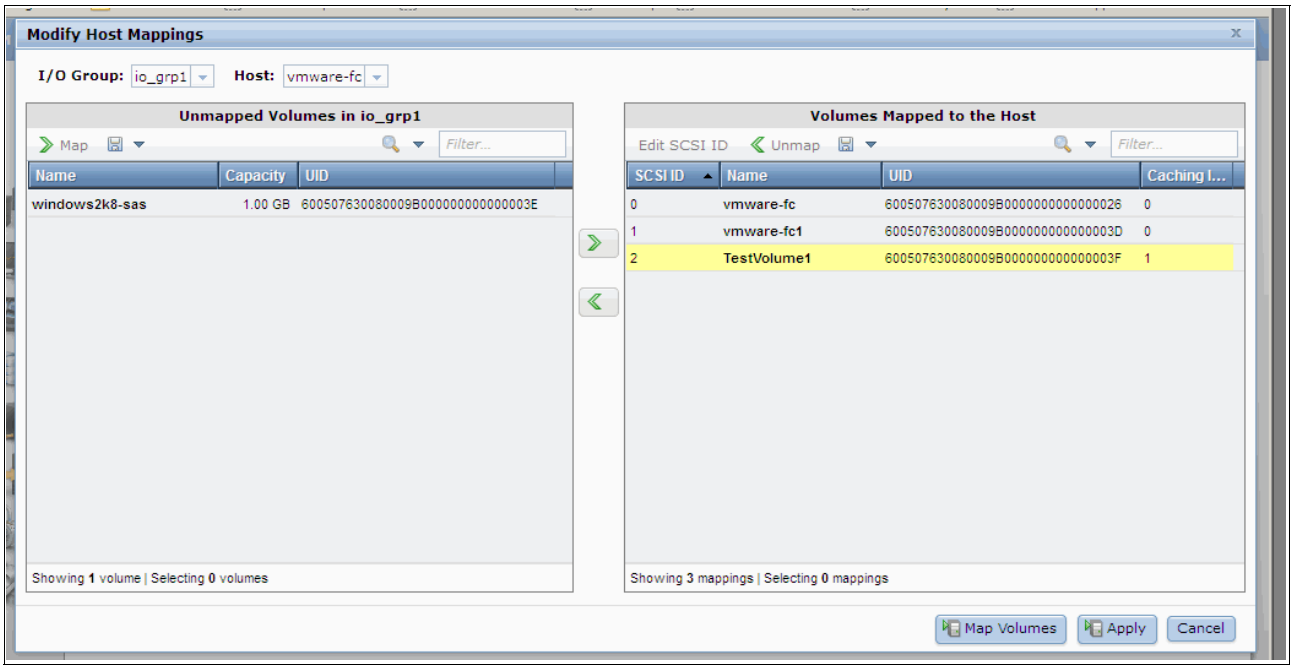

Figure 5-24 Modify Host Mappings window

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.

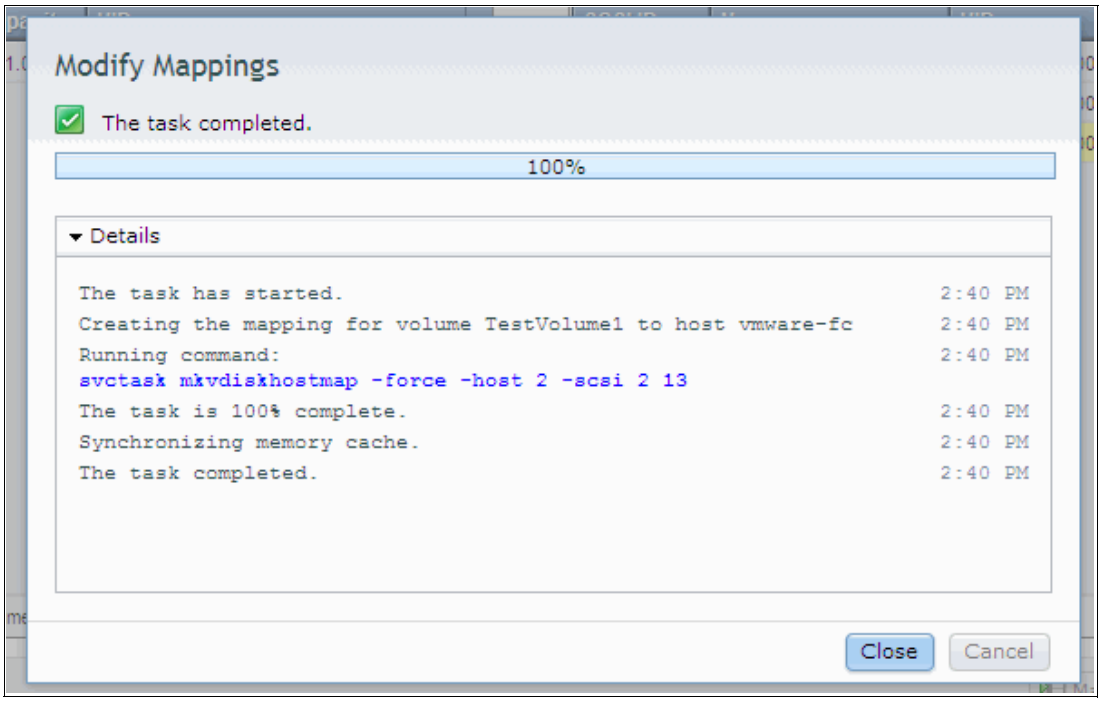

Figure 5-25 Host mapping task complete

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.

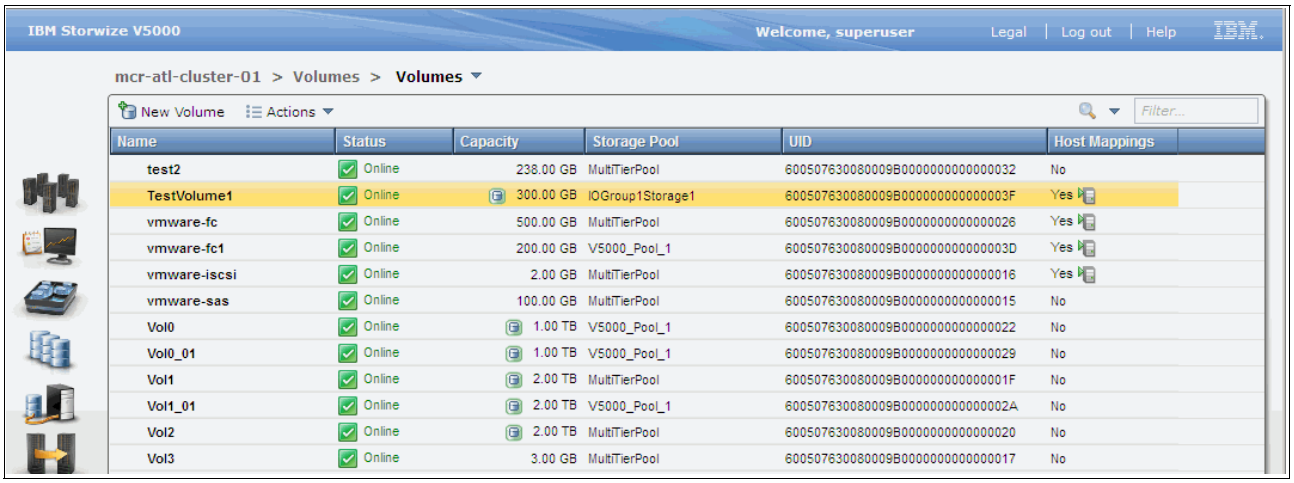

Figure 5-26 New volume that is mapped to host

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.

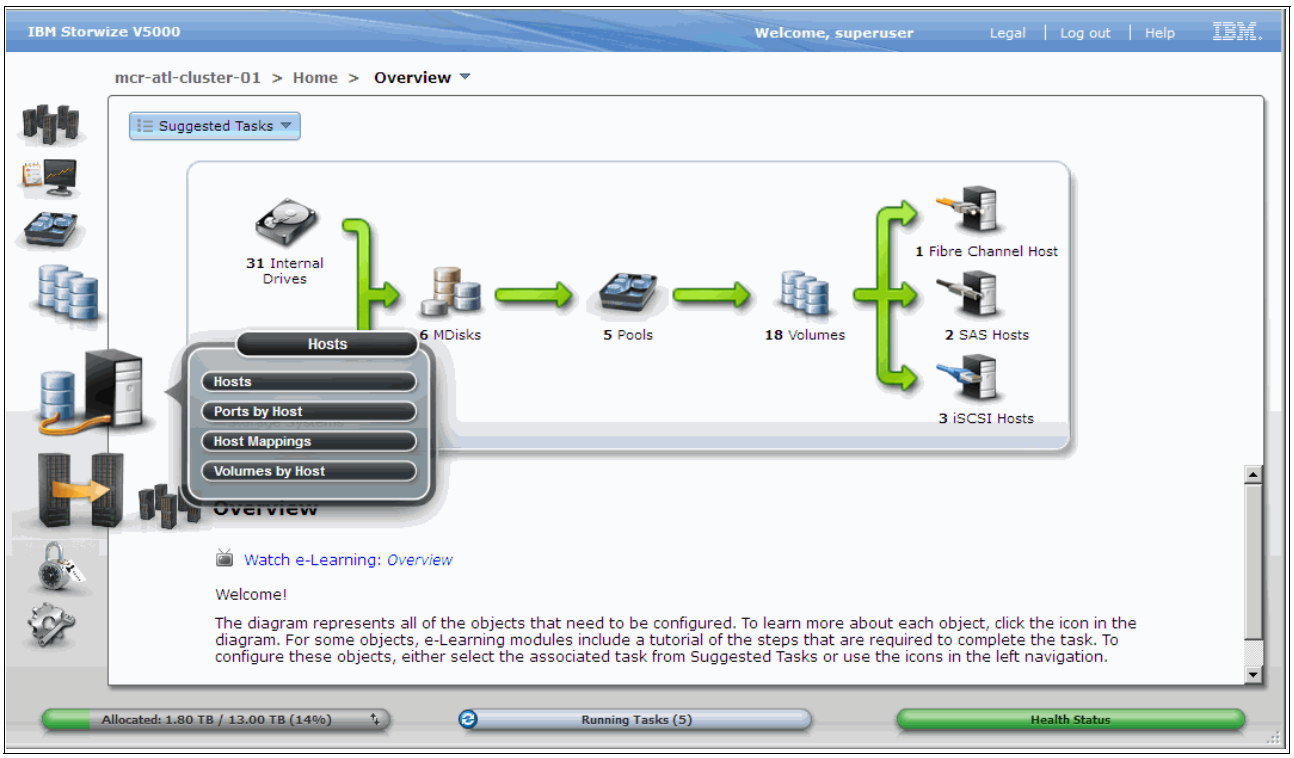

Figure 5-27 Hosts window

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

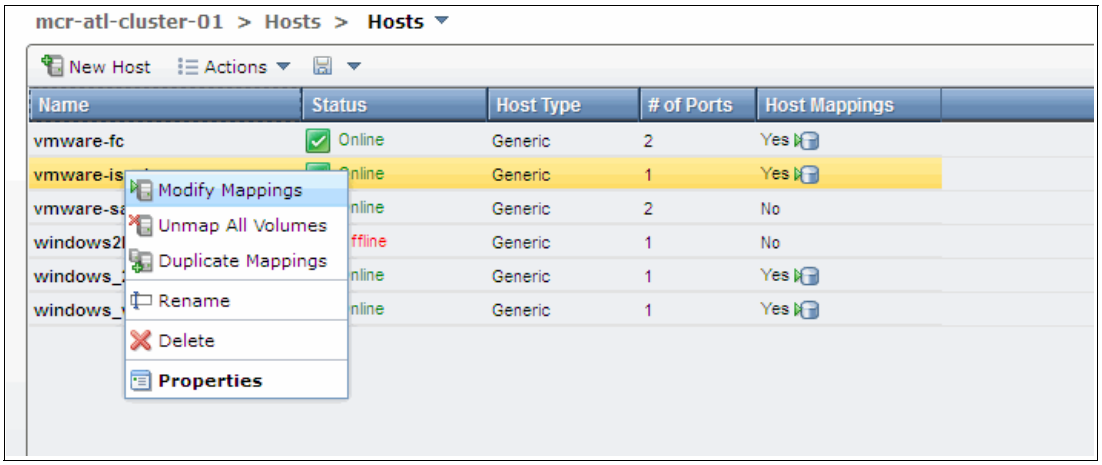

Figure 5-28 Modify mappings selection

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.

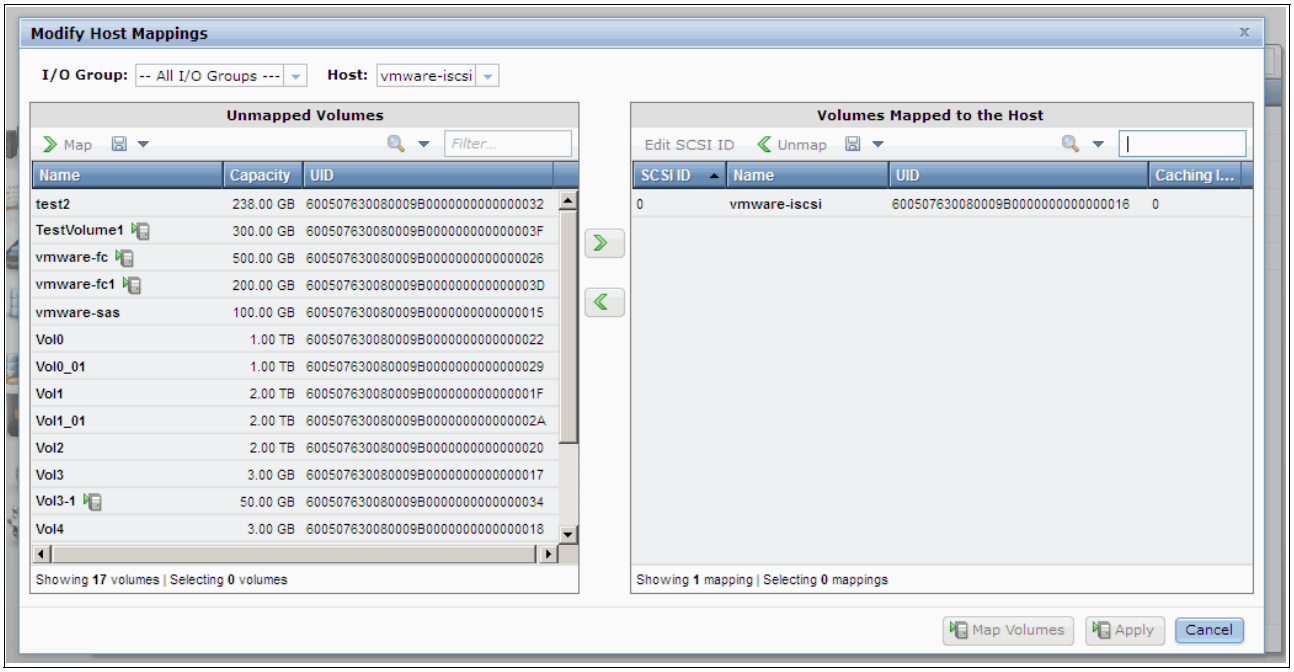

Figure 5-29 Modify host mappings window

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

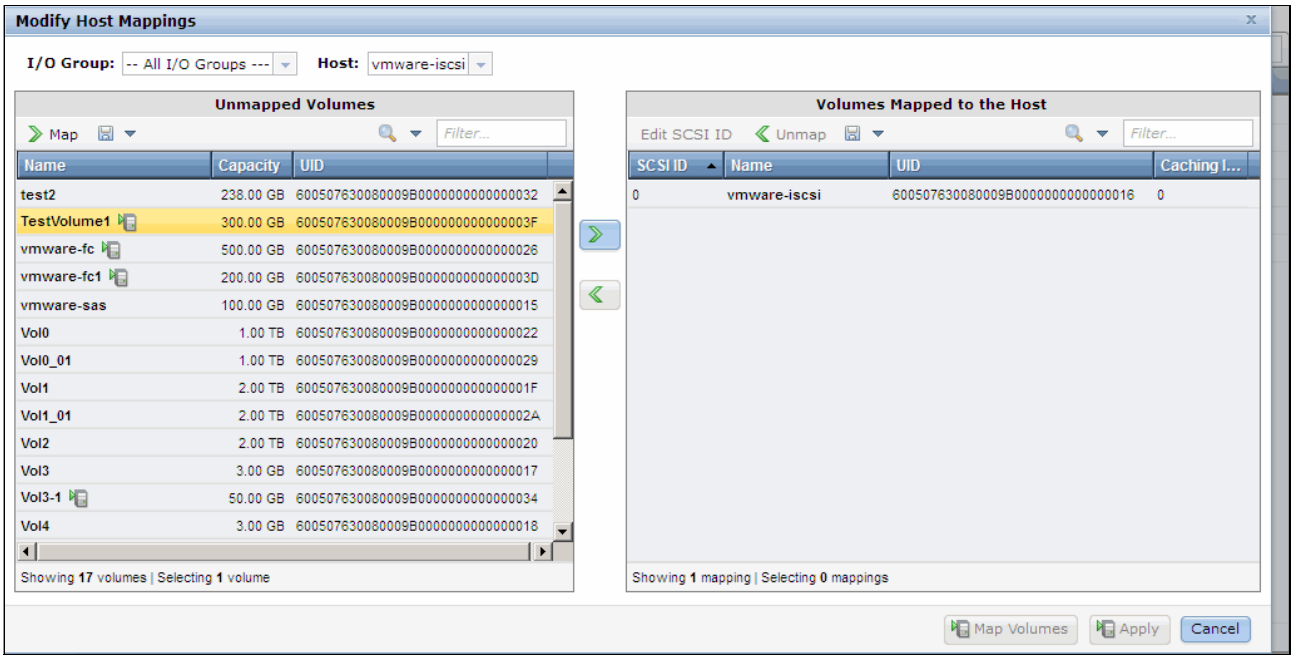

Figure 5-30 Volume mapping selected

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.

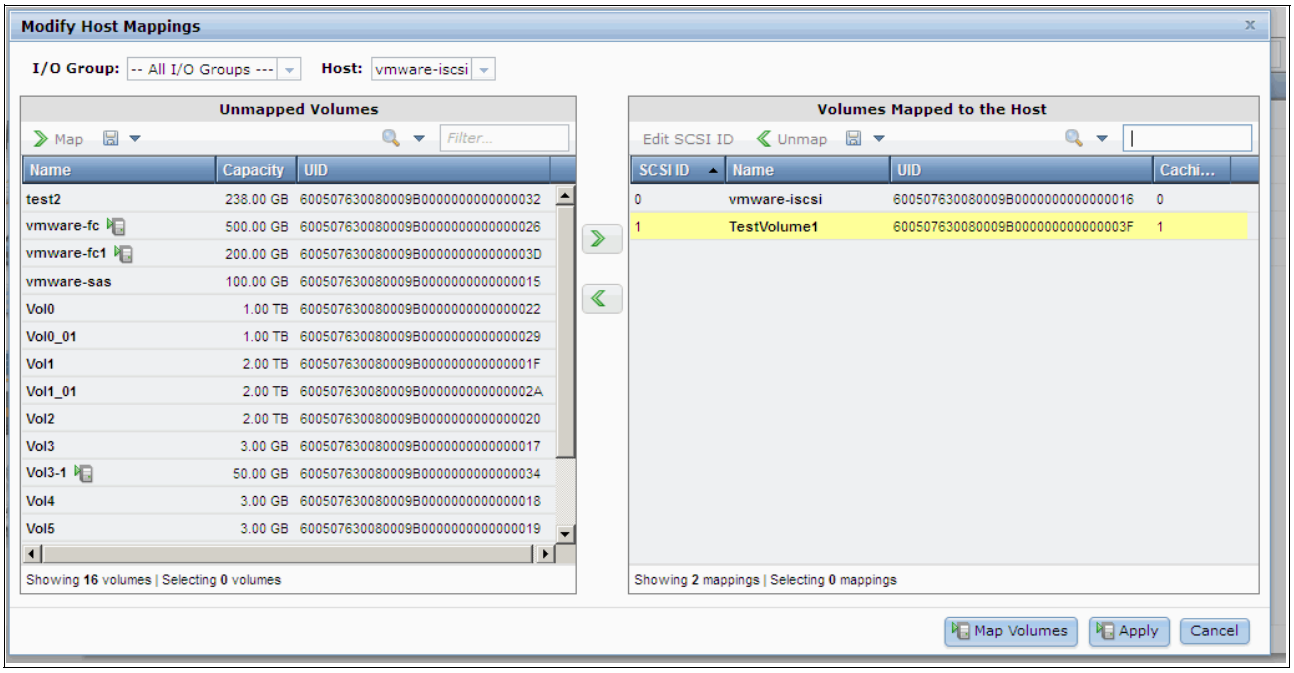

Figure 5-31 Modify host mappings window

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.

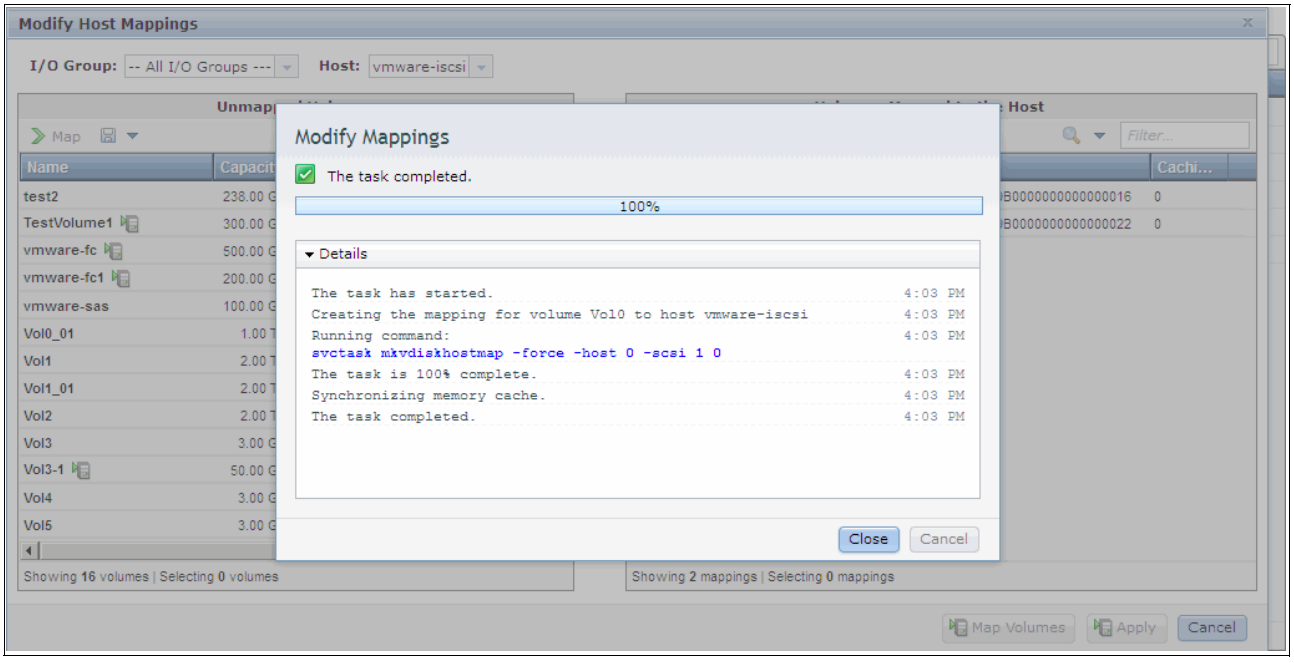

Figure 5-32 Modify mapping complete

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.

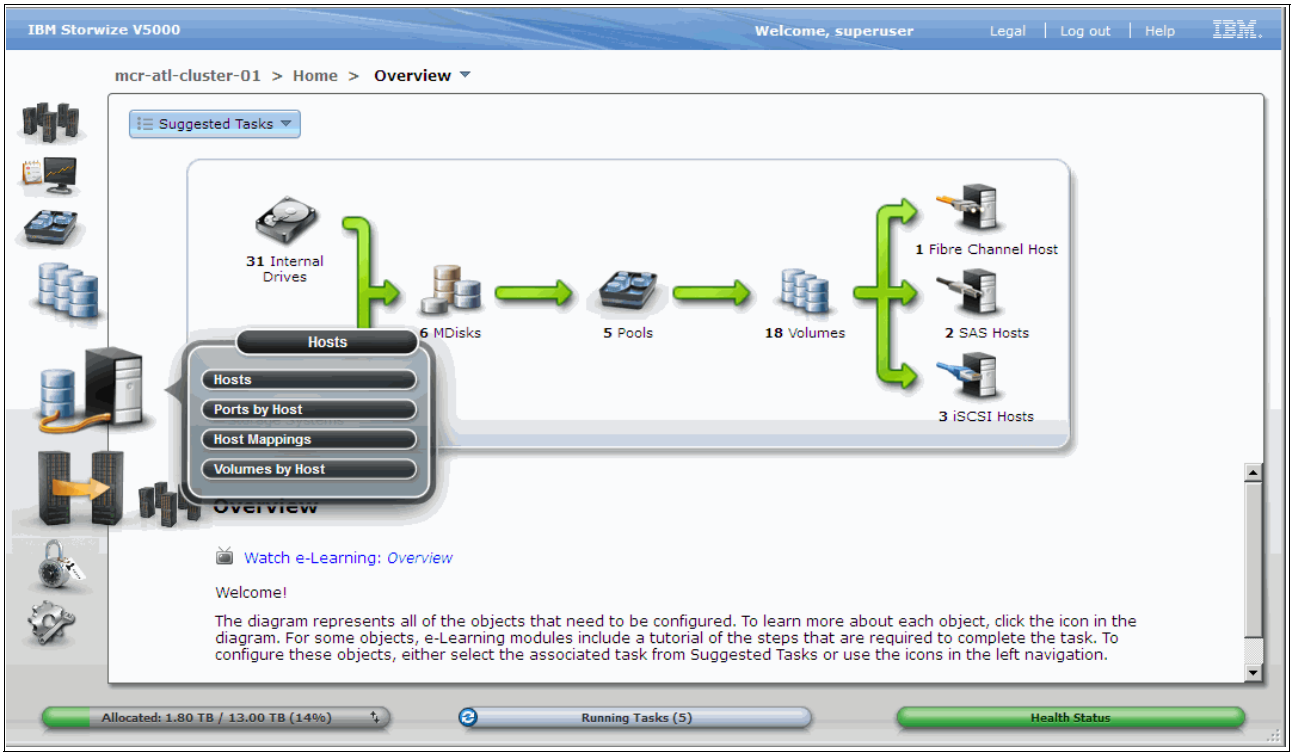

Figure 5-33 Open all hosts

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

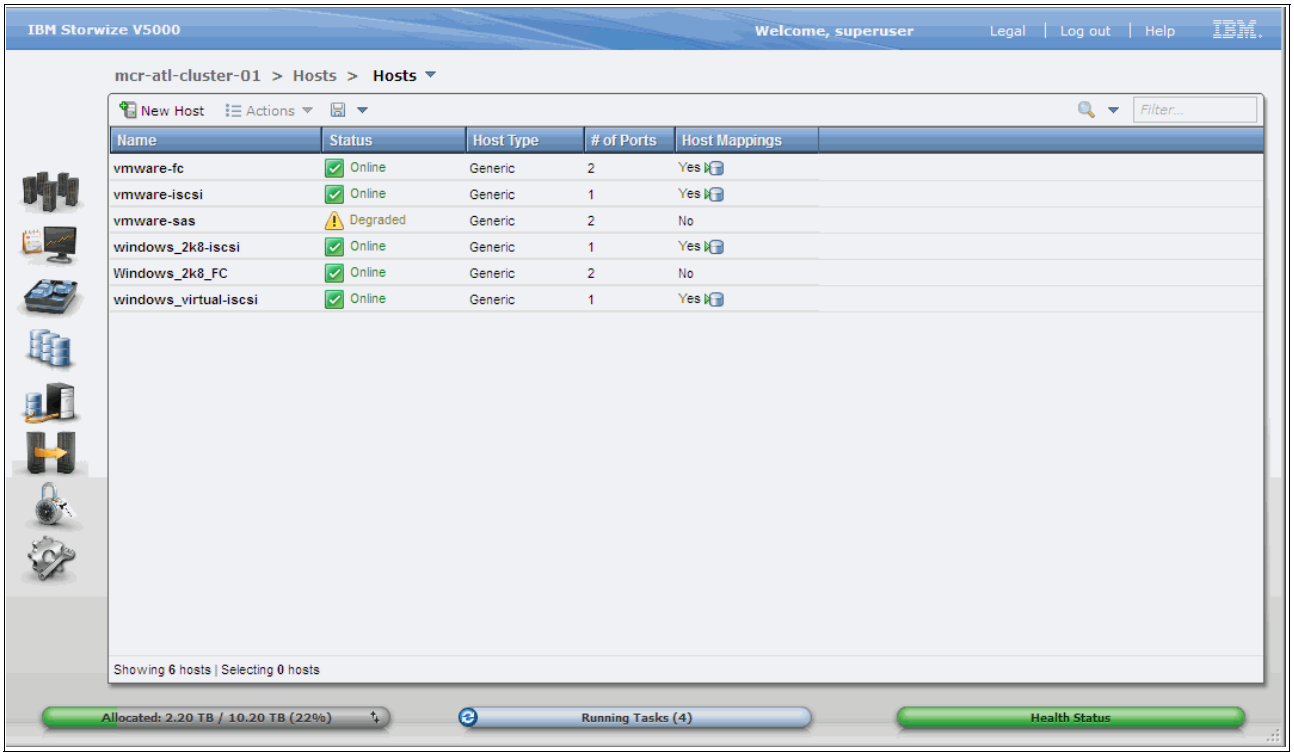

Figure 5-34 All Hosts view

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.

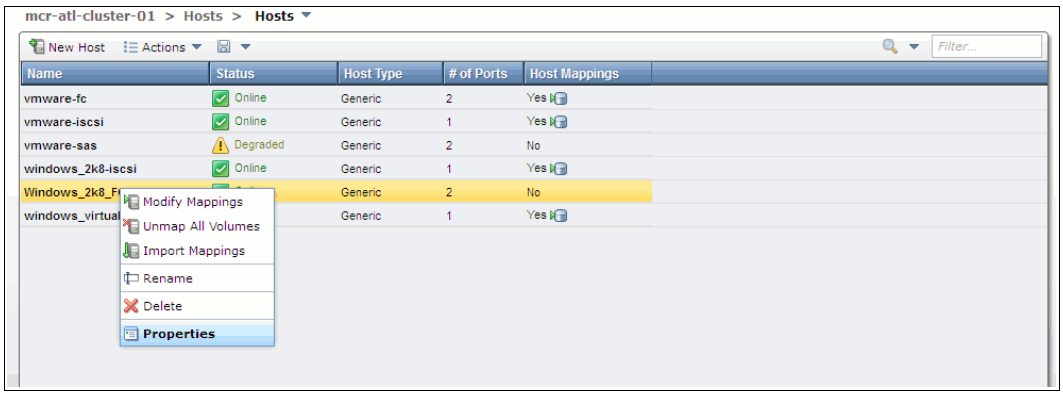

Figure 5-35 Host properties

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

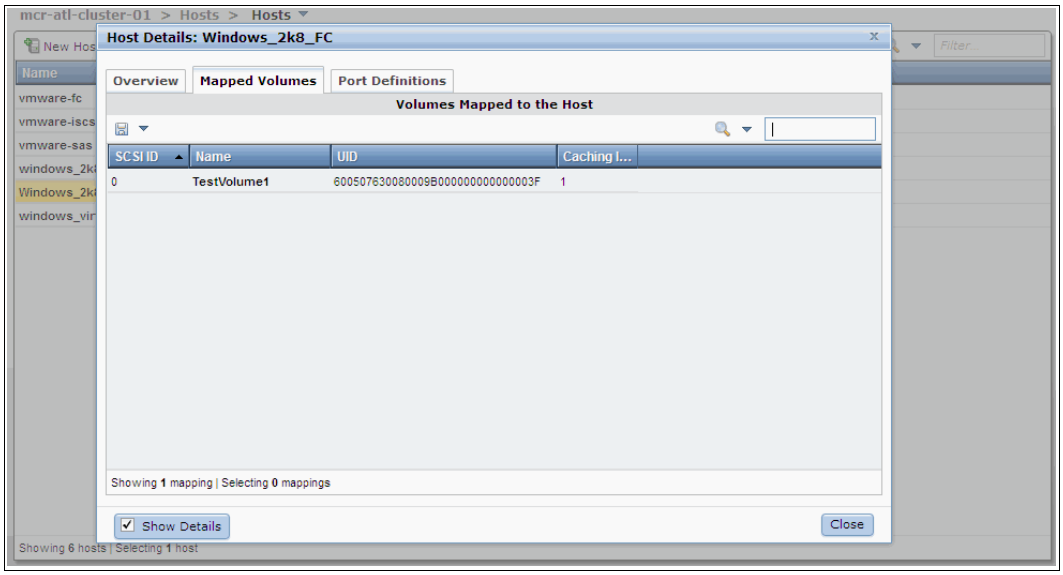

Figure 5-36 Mapped volumes to a host

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**.

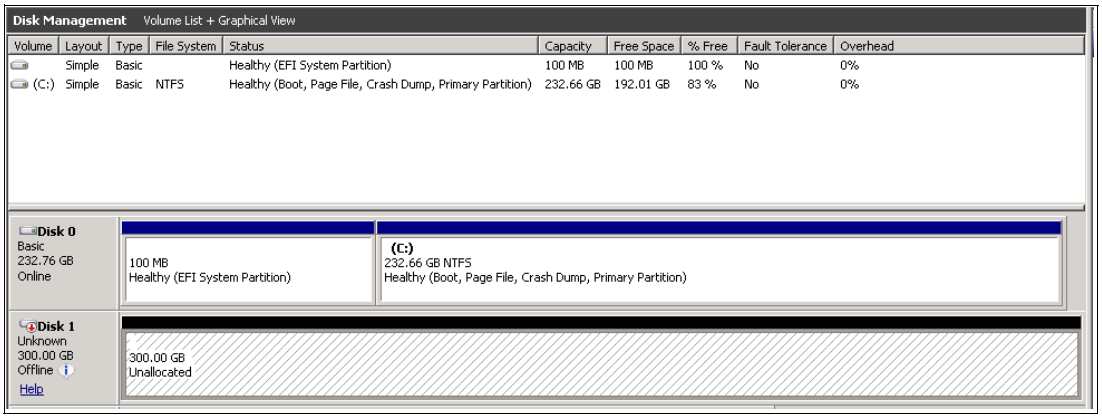

Figure 5-37 Windows Disk Management

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.

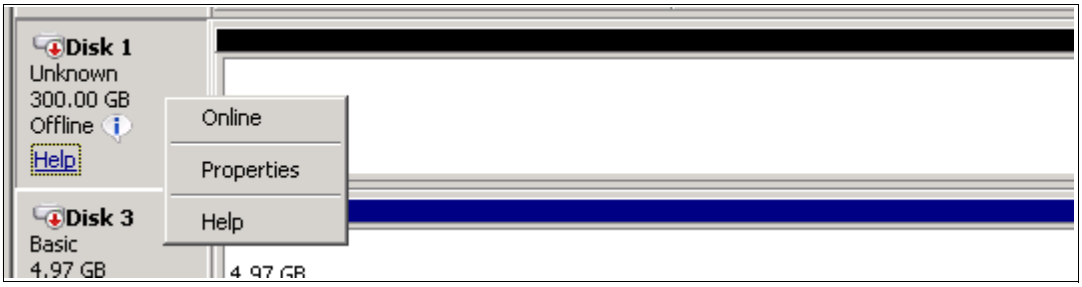

Figure 5-38 Setting a disk online

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

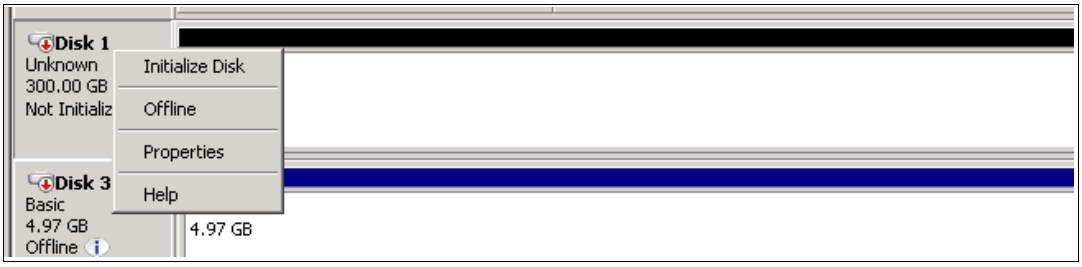

Figure 5-39 Initializing disk

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

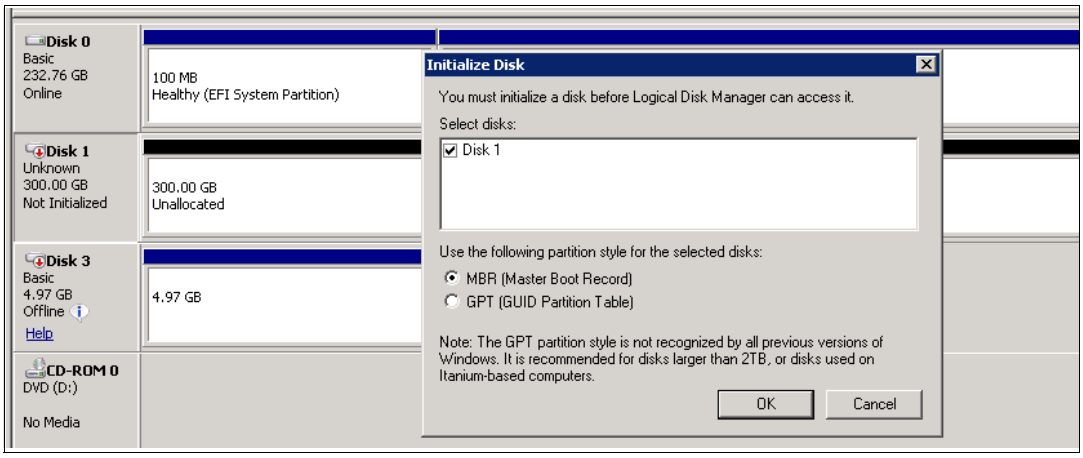

Figure 5-40 Initialize Disk option

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

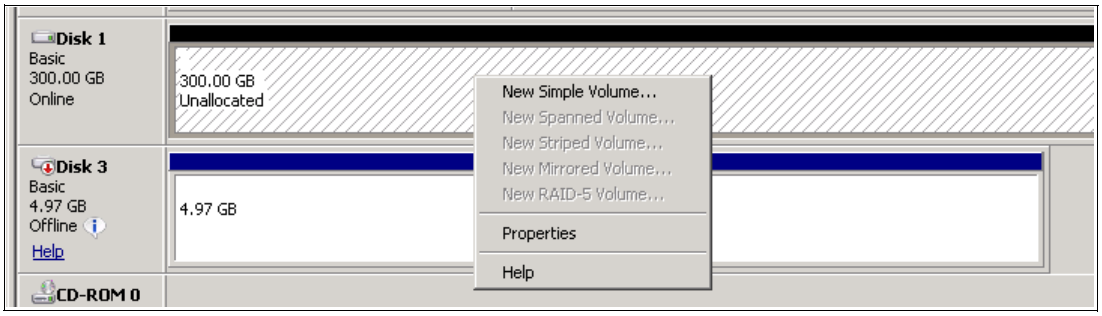

Figure 5-41 New Simple Volume

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.

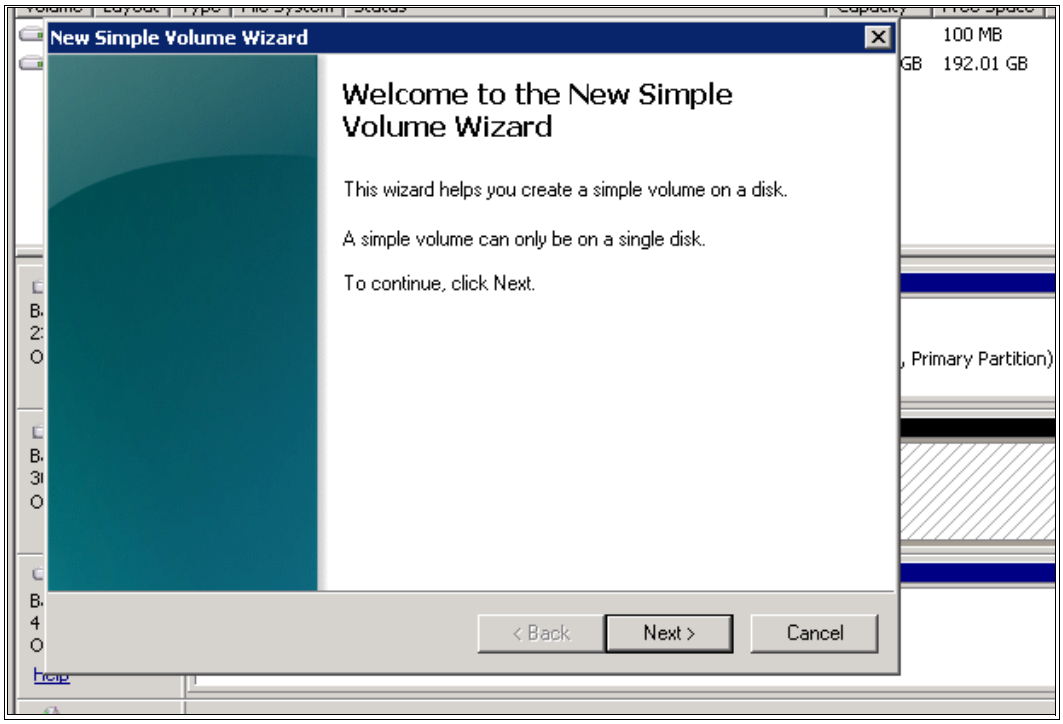

Figure 5-42 New Volume wizard

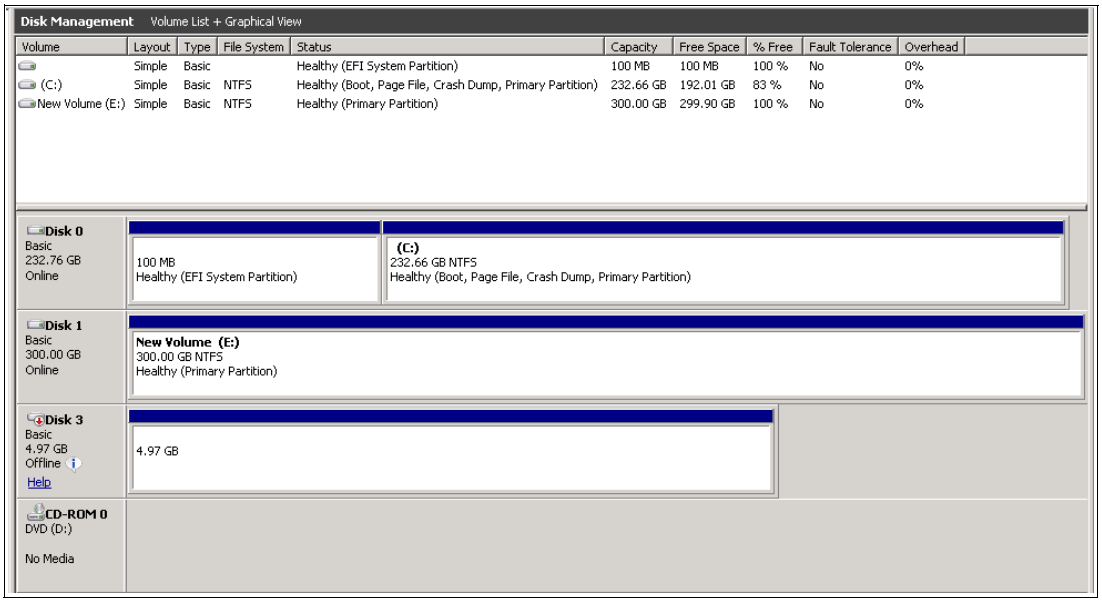

Figure 5-43 Volume is ready to use

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.

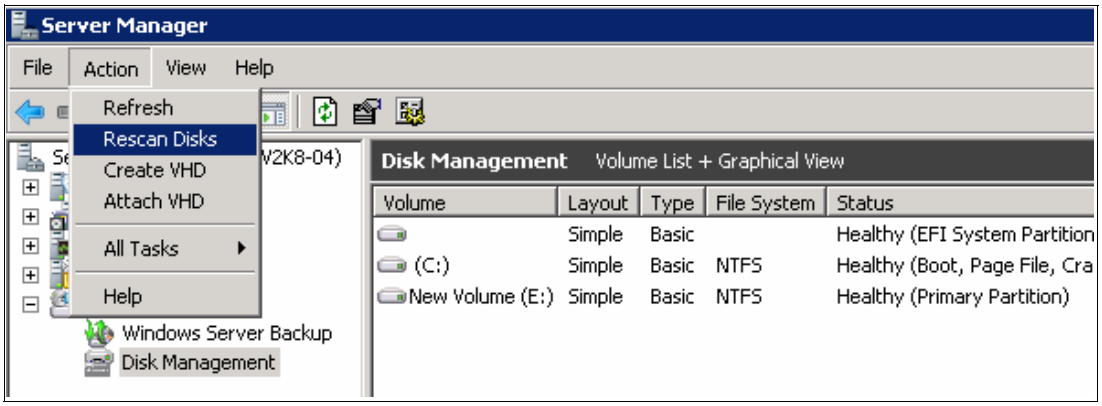

Figure 5-44 Windows disk rescan

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.

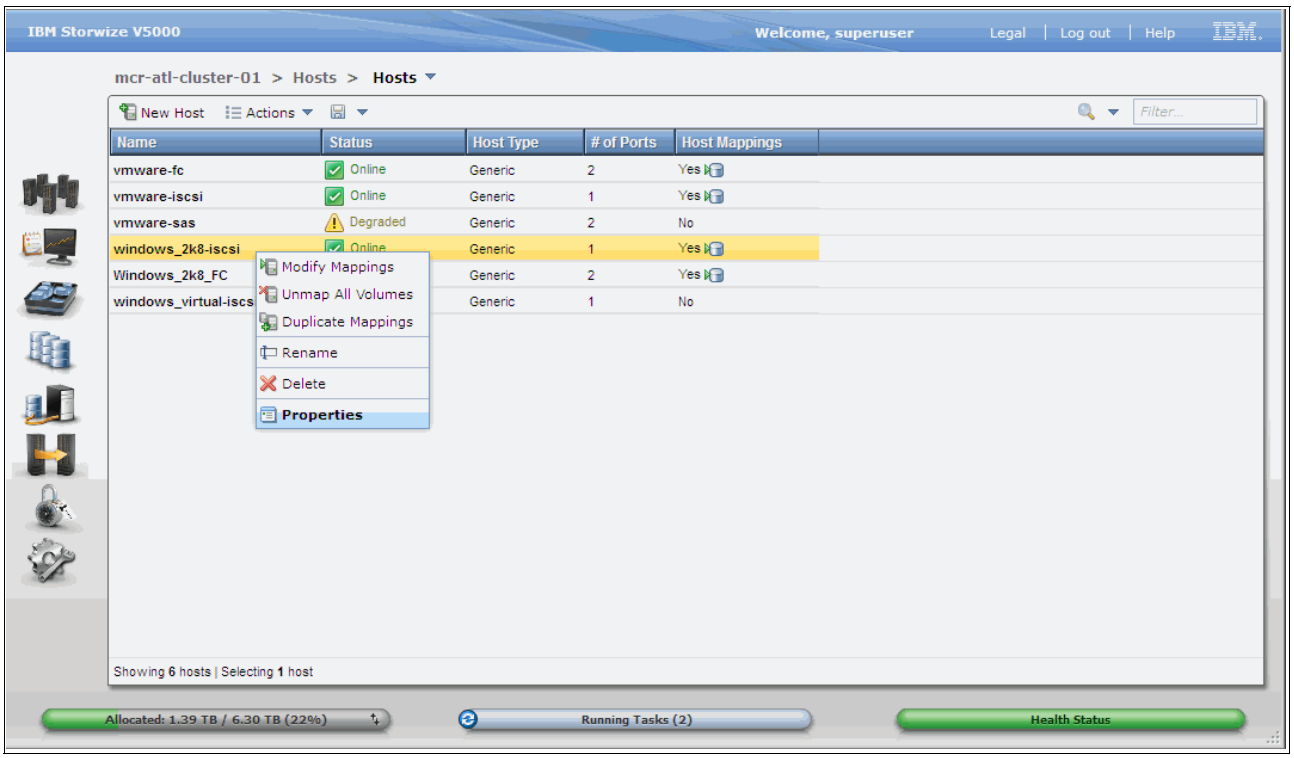

Figure 5-45 All Hosts view

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

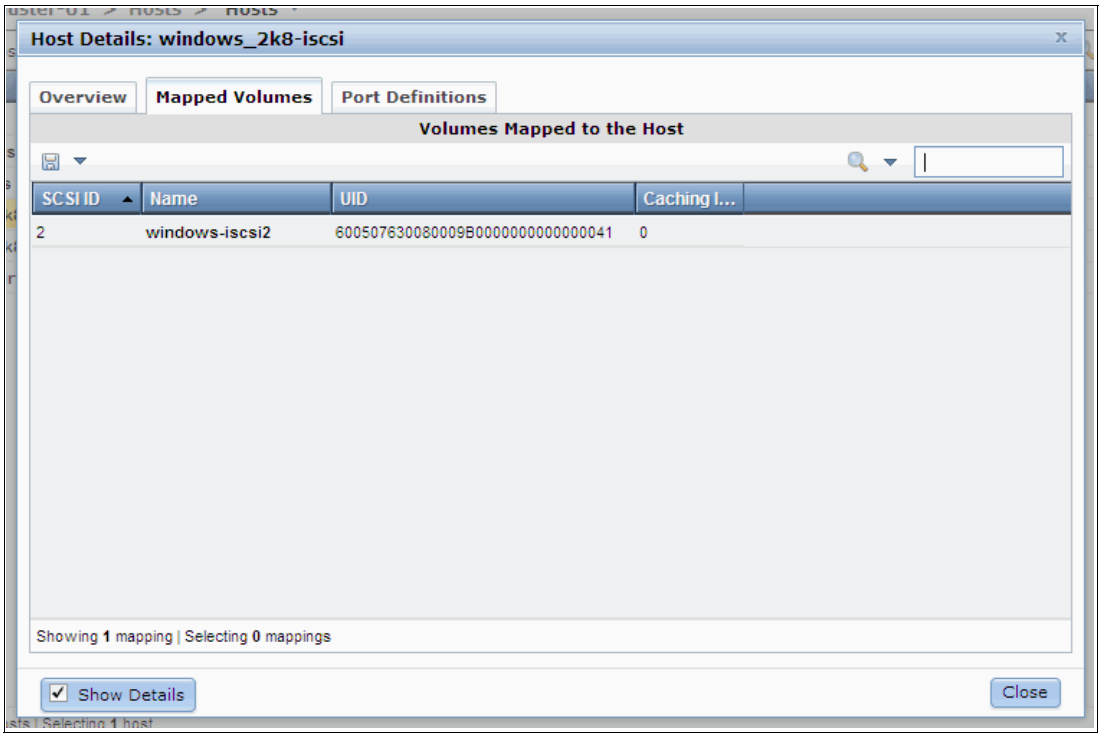

Figure 5-46 Mapped volumes on an iSCSI host

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.

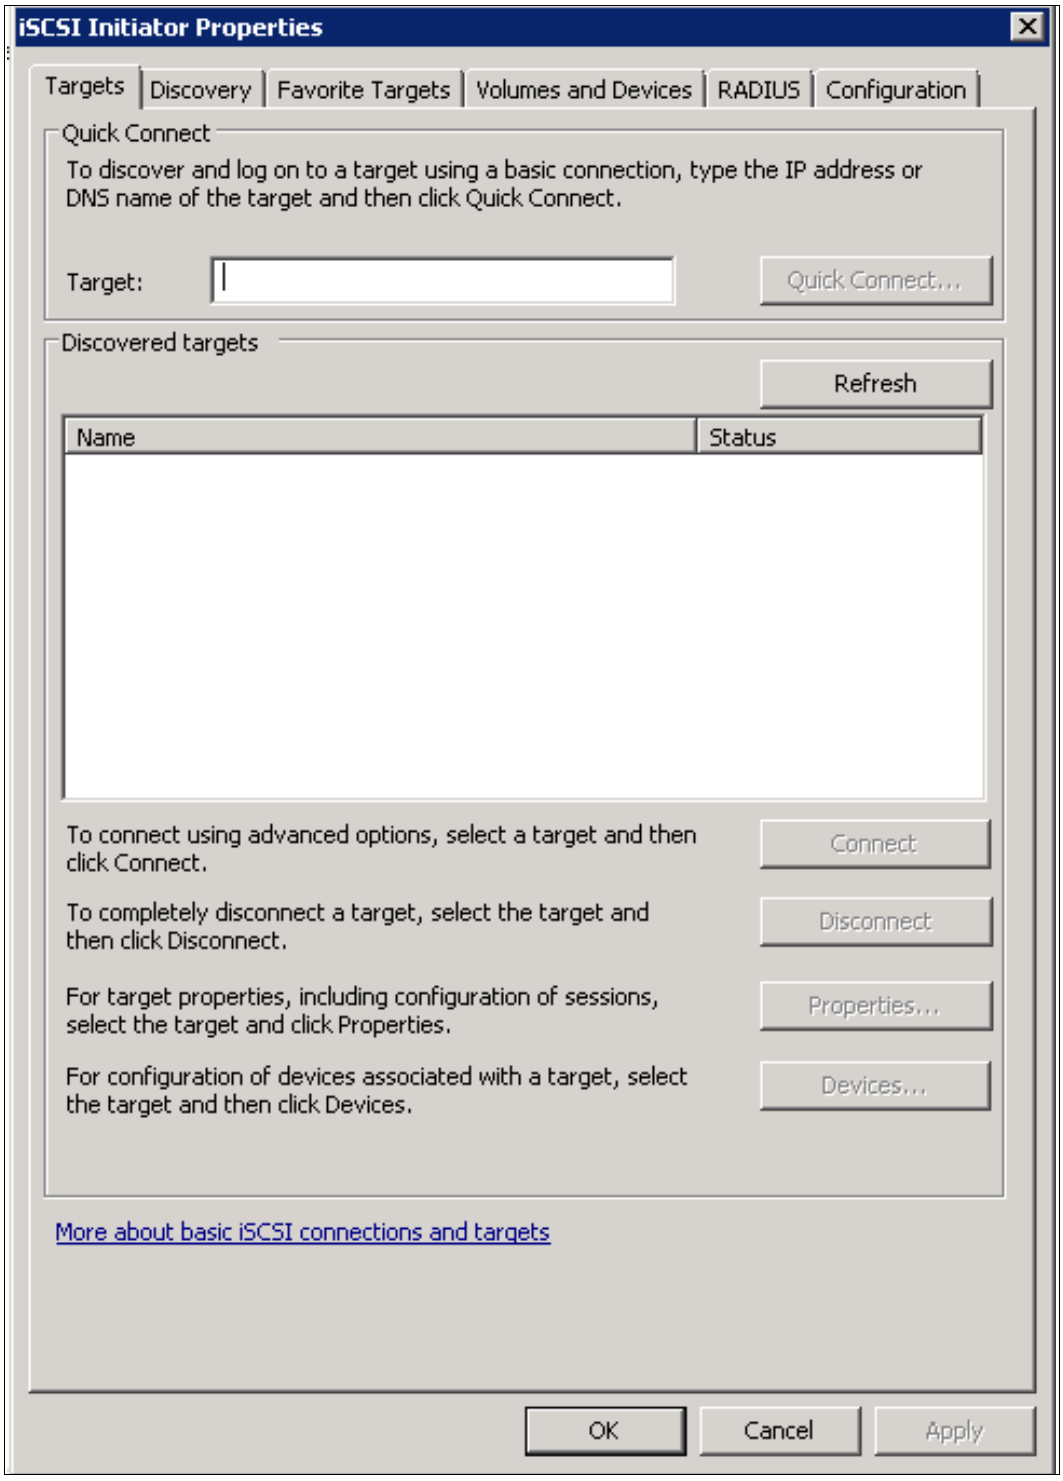

Figure 5-47 Windows iSCSI Configuration tab

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.

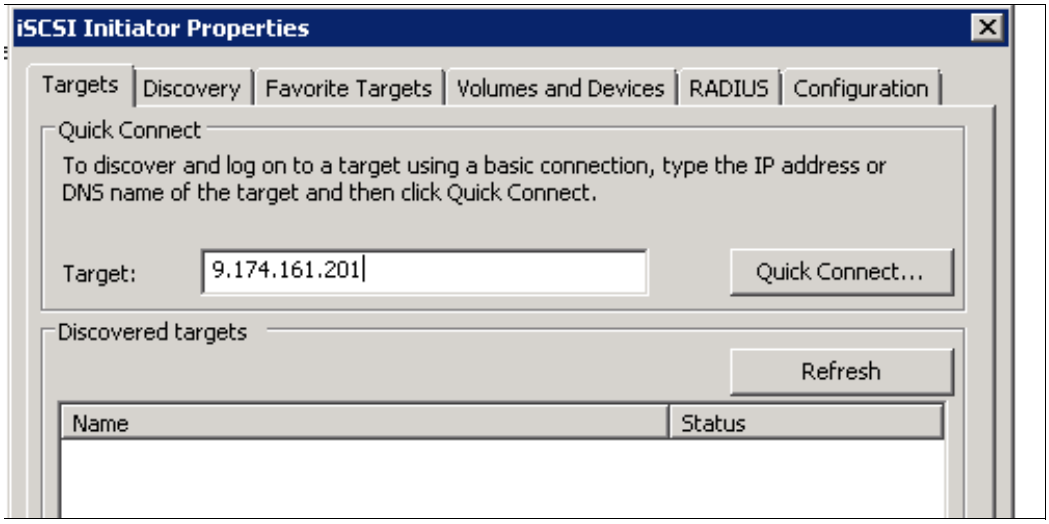

Figure 5-48 iSCSI Quick Connect

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

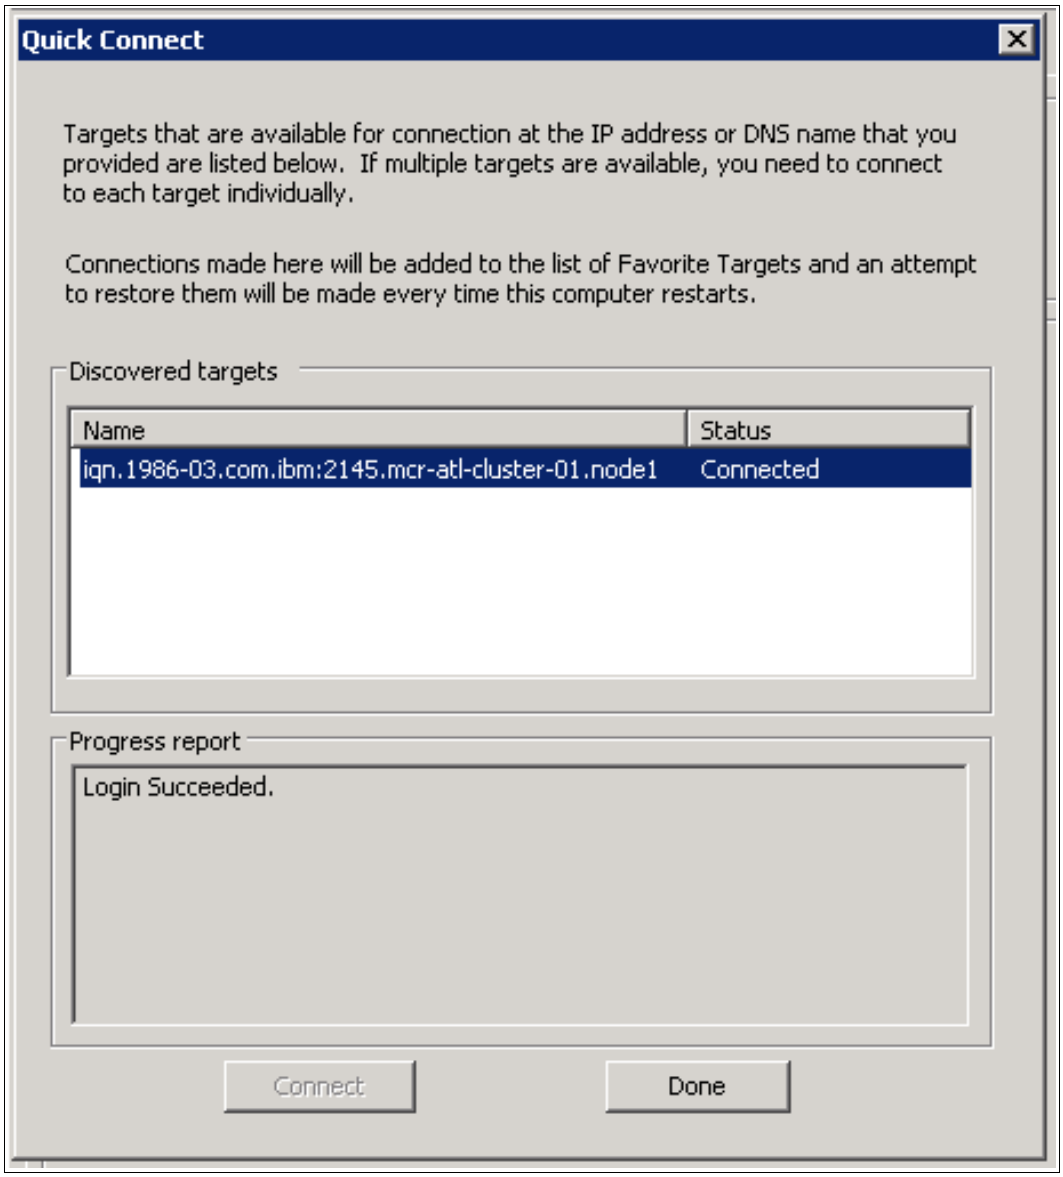

Figure 5-49 iSCSI Initiator target is connected

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.

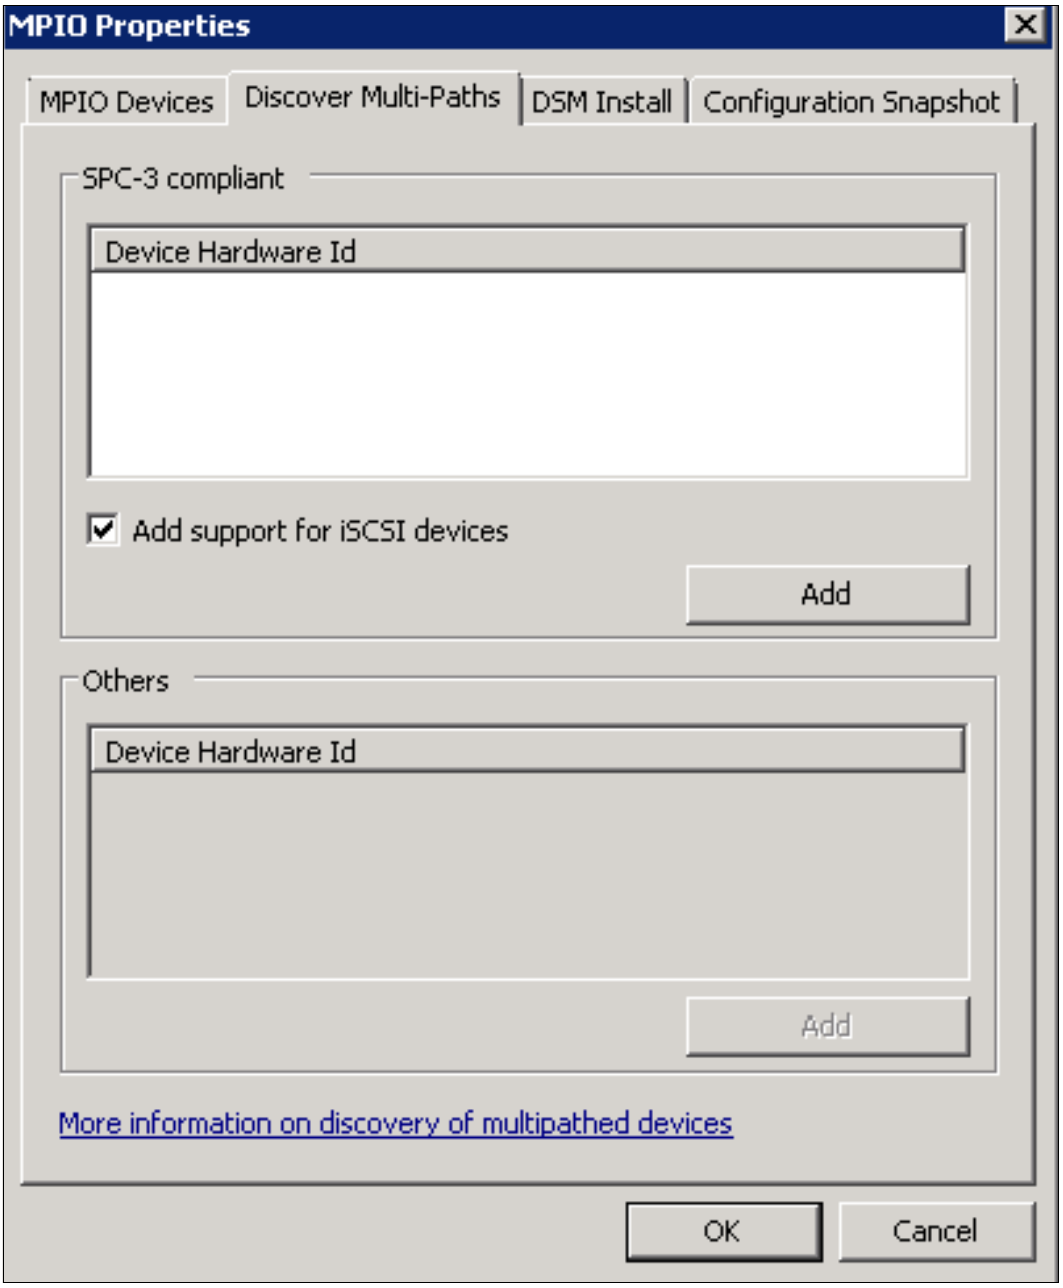

Figure 5-50 Enable iSCSI MPIO

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.

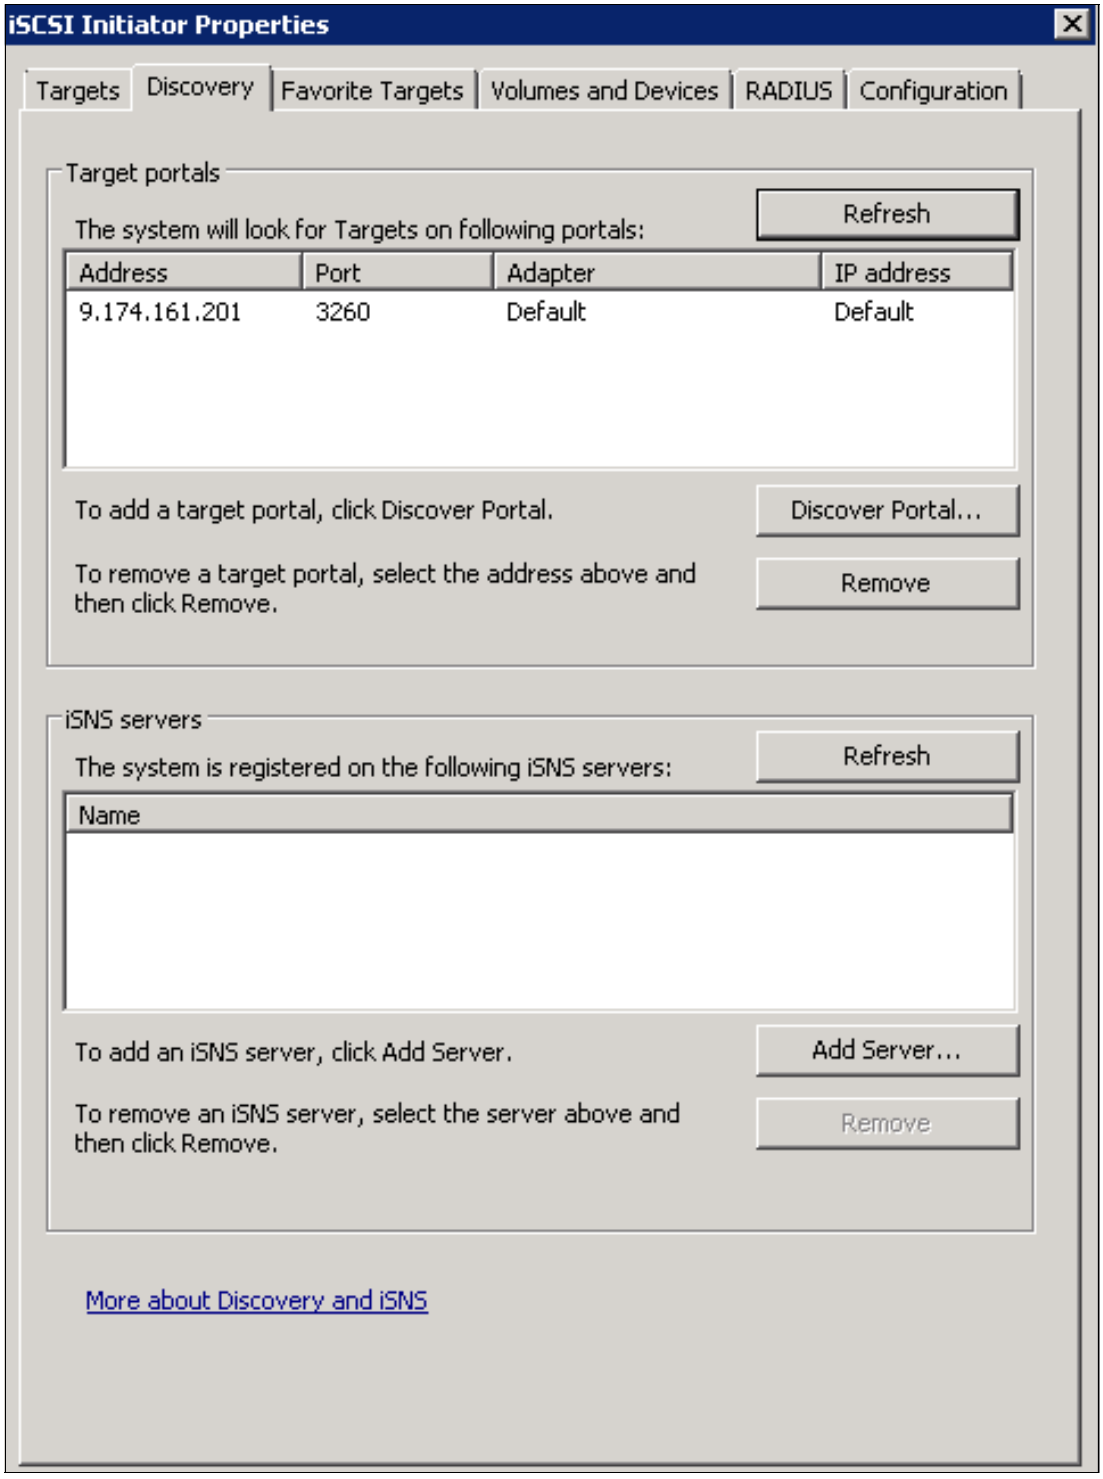

Figure 5-51 iSCSI Properties Discovery tab

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

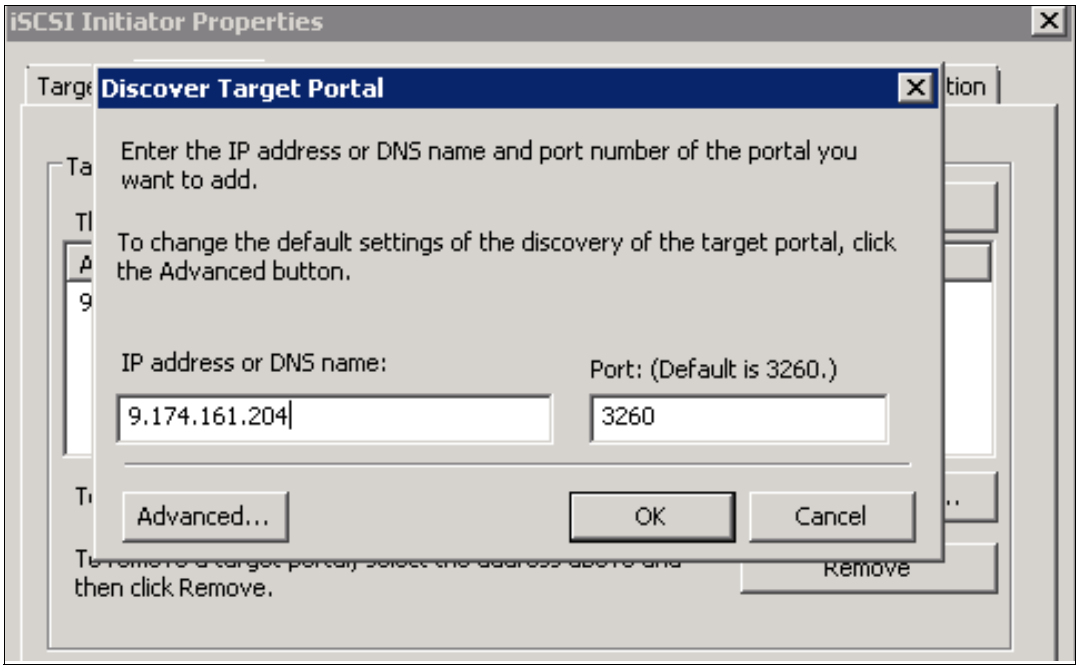

Figure 5-52 Discover Target Portal window

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

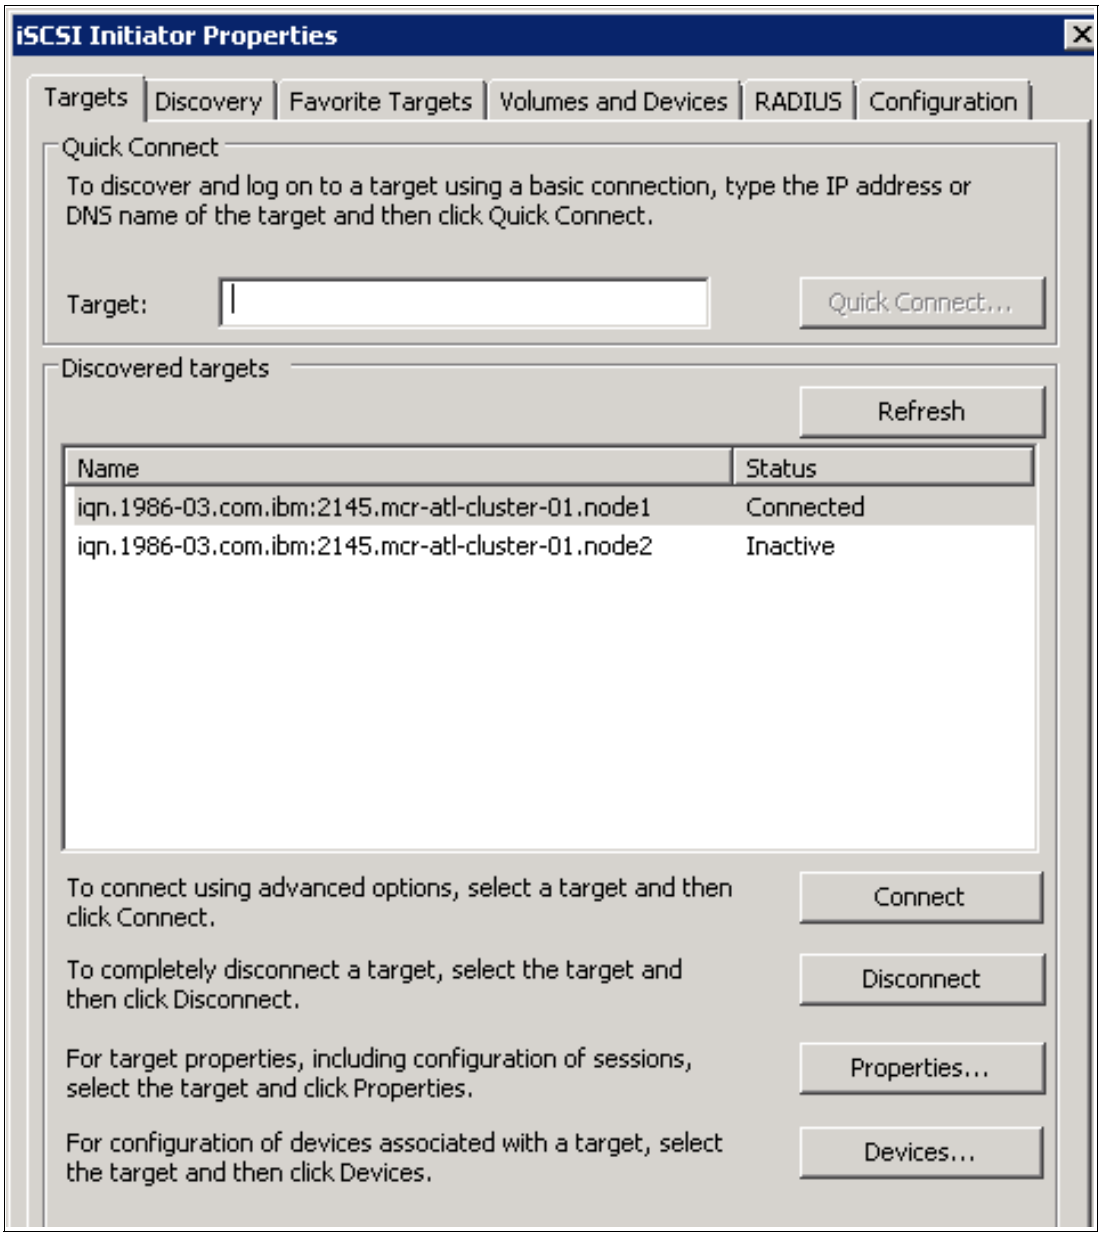

Figure 5-53 Inactive target ports

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

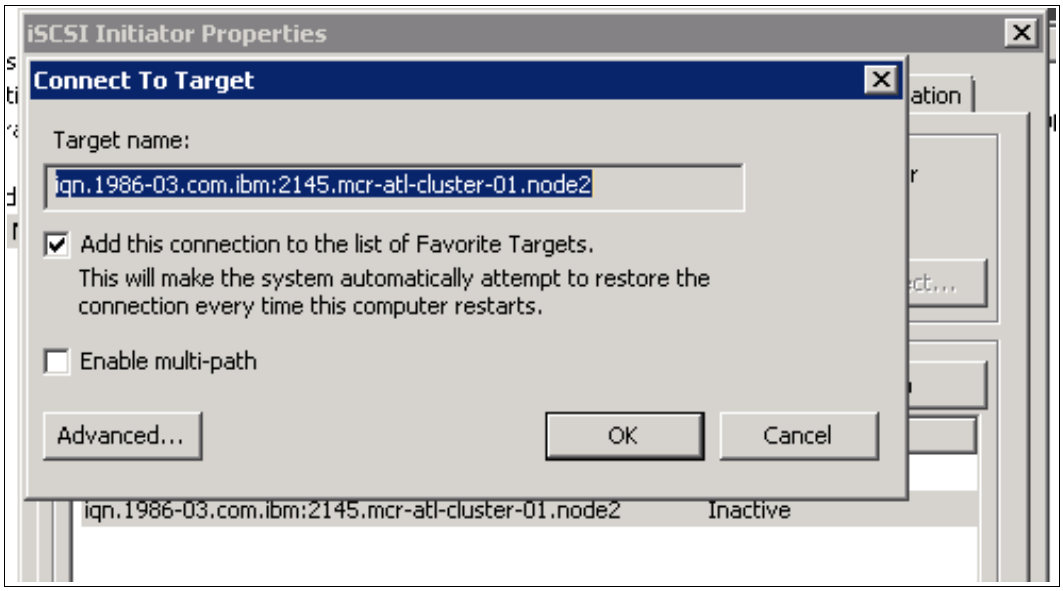

Figure 5-54 Connect to a target

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

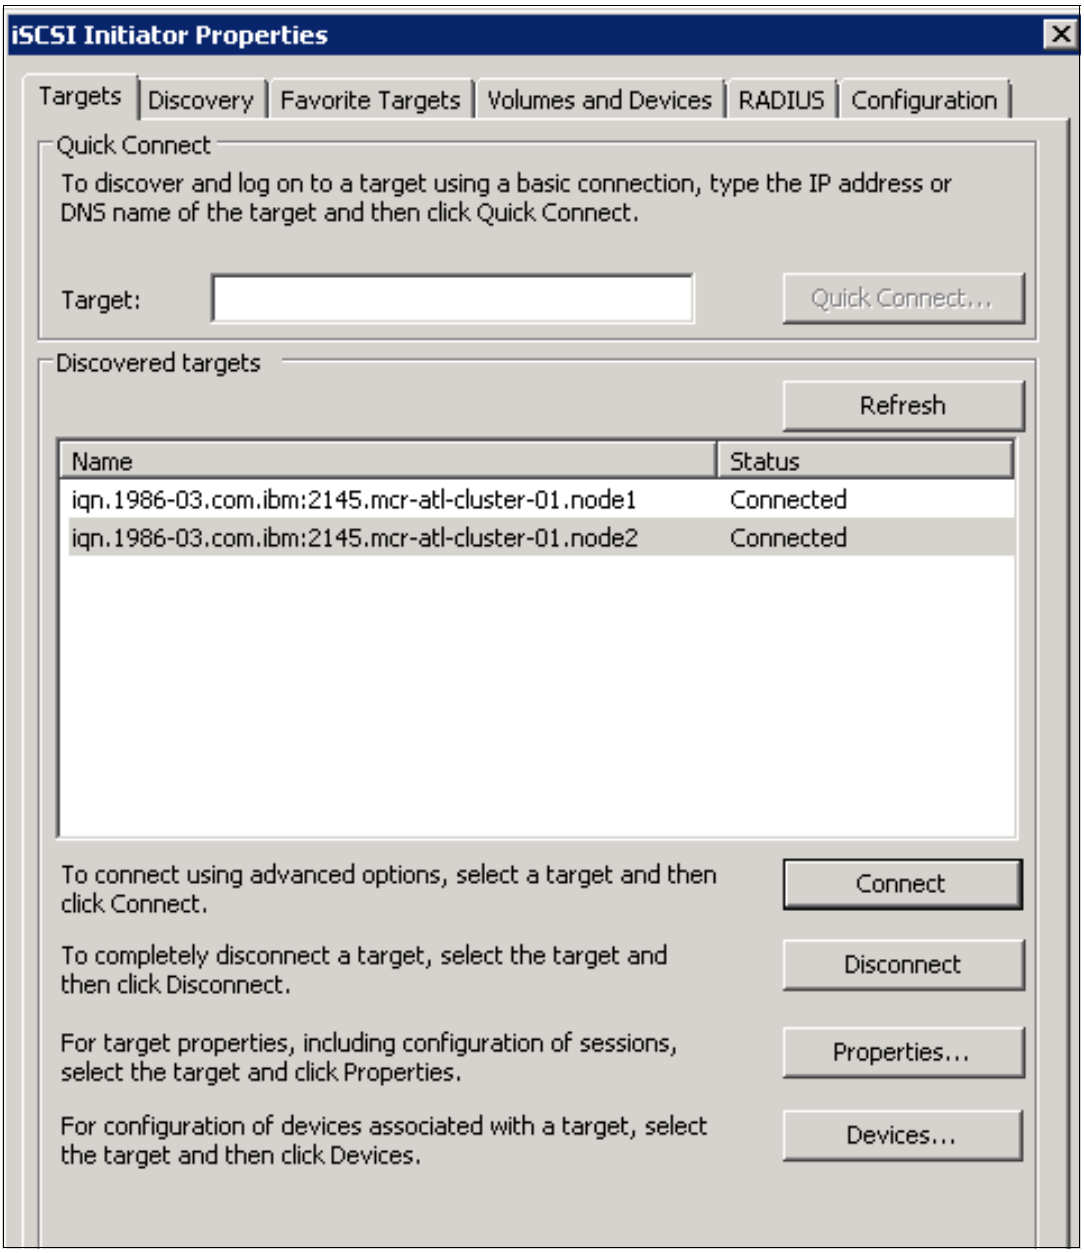

Figure 5-55 Second target port connected

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**.

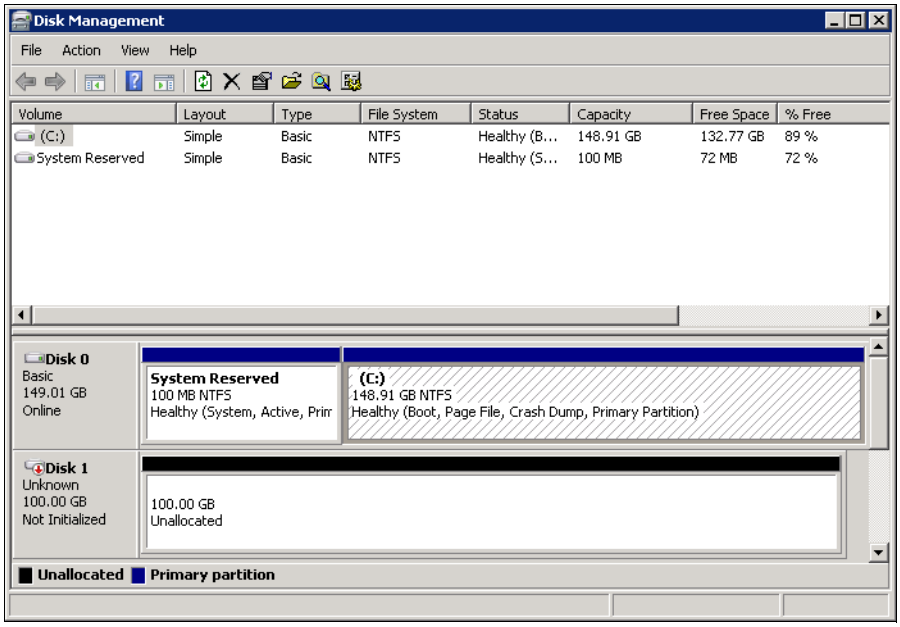

Figure 5-56 Windows Disk Management

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.

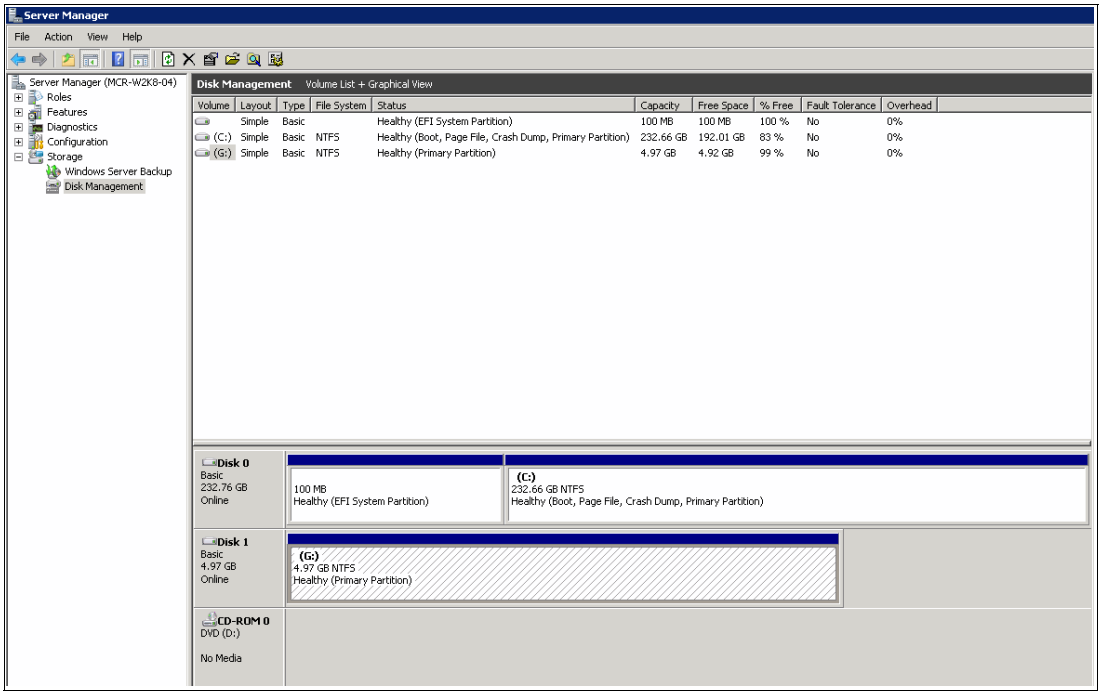

Figure 5-57 Disk is ready to use

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.

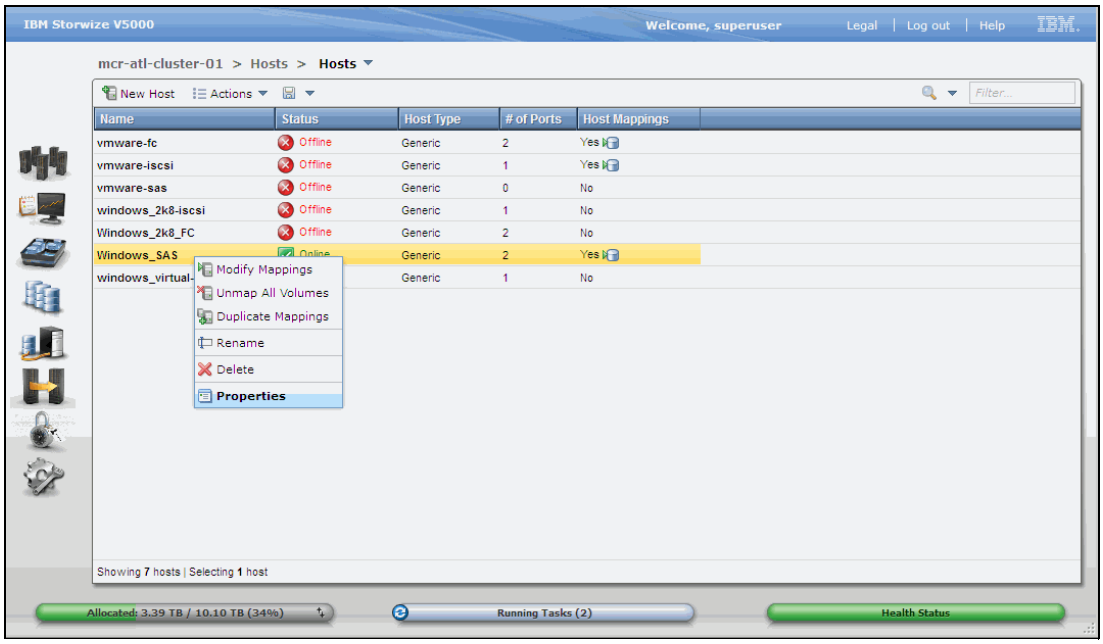

Figure 5-58 Windows SAS host from host view

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

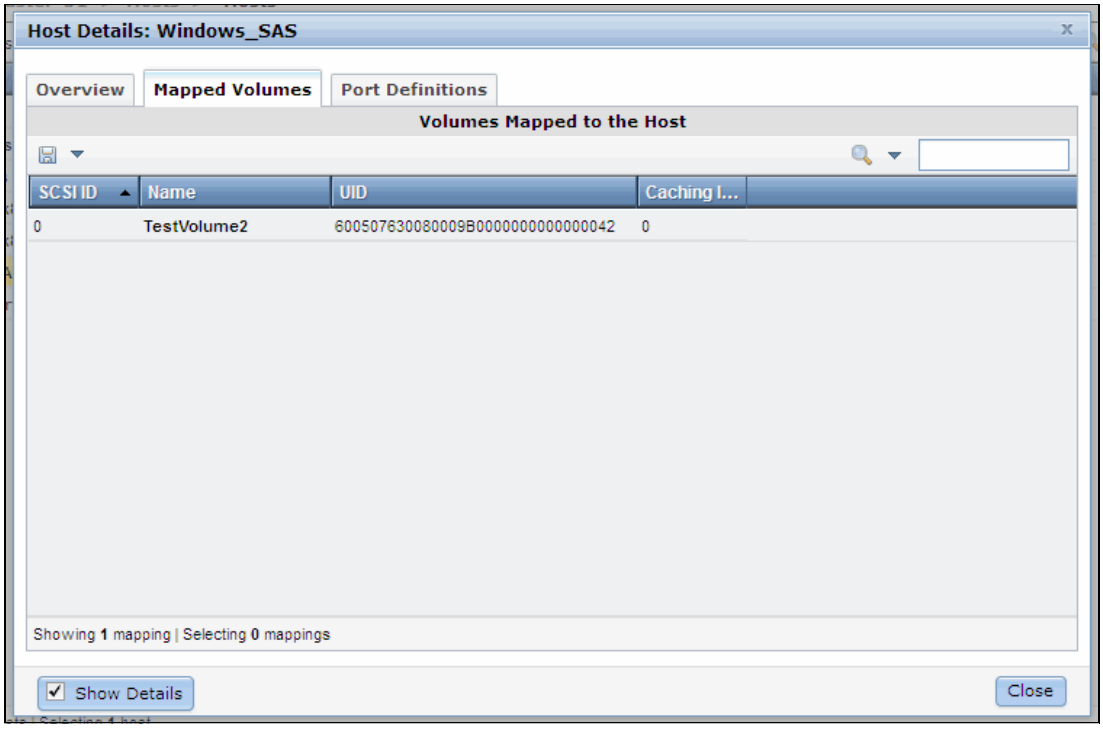

Figure 5-59 SAS host mapped volumes

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**.

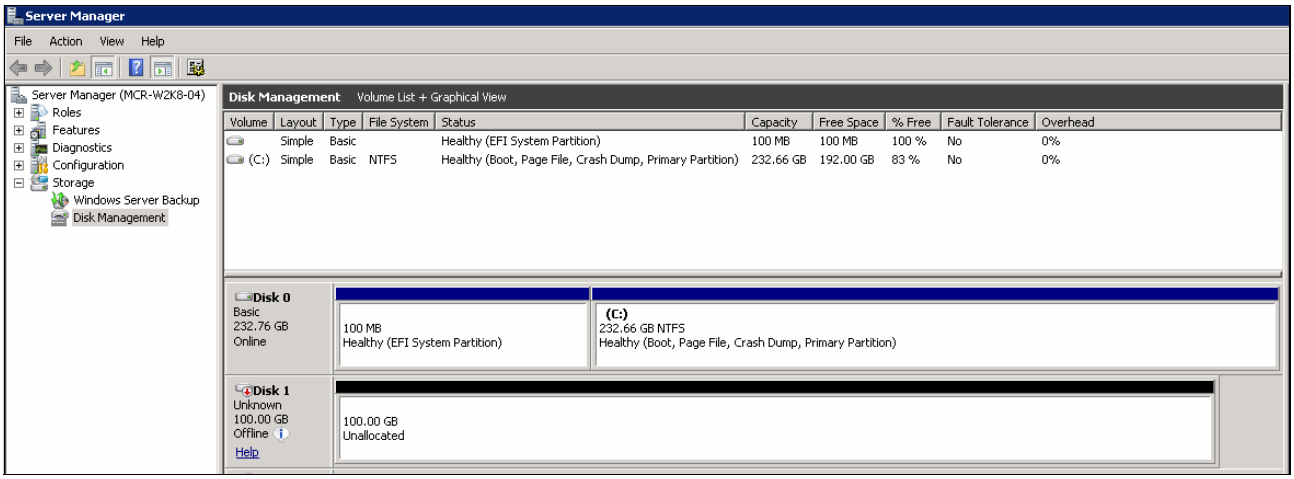

Figure 5-60 Windows Disk Management

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.

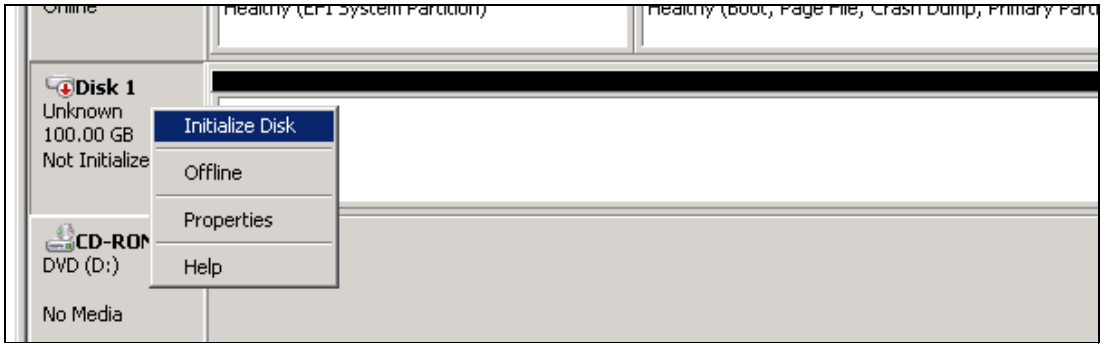

Figure 5-62 Initializing disk

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

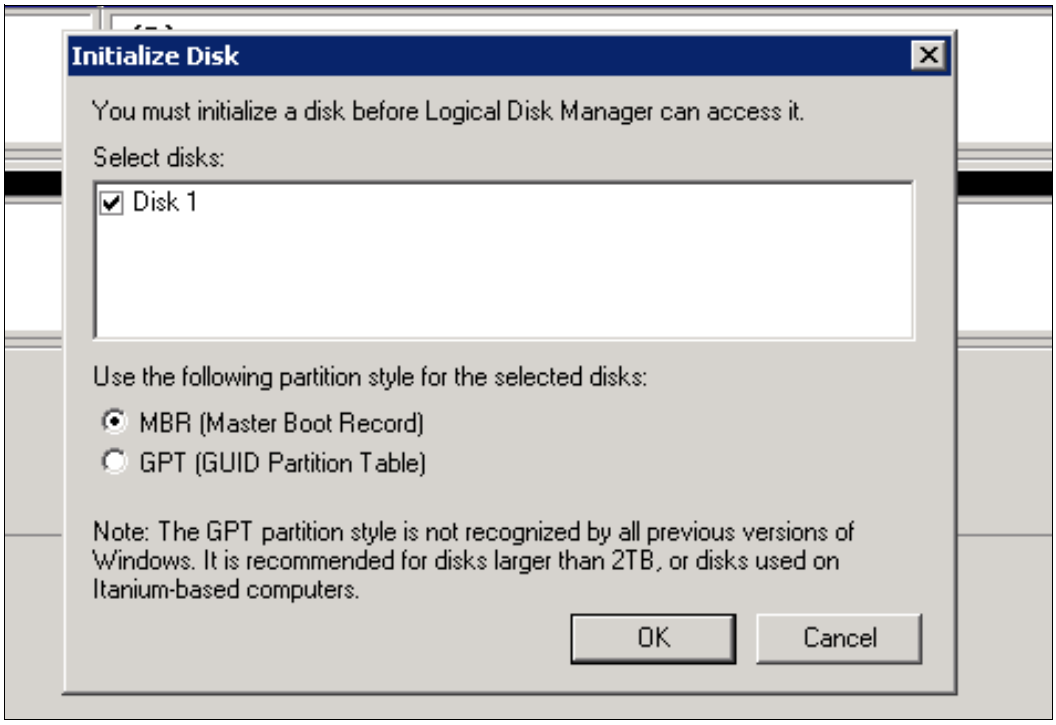

Figure 5-63 Initialize disk option

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

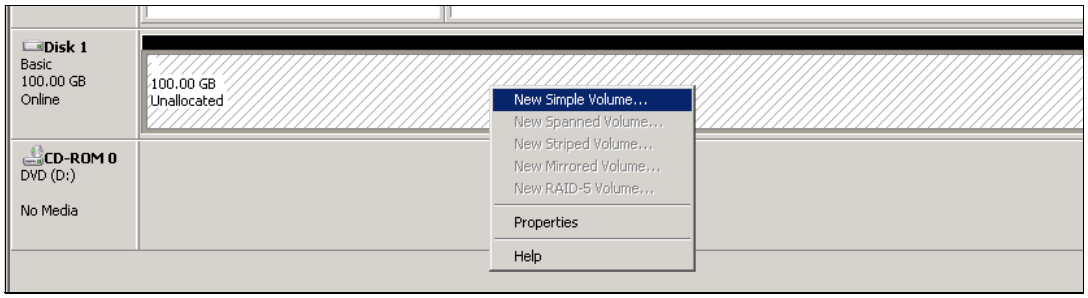

Figure 5-64 New simple volume: SAS attach

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.

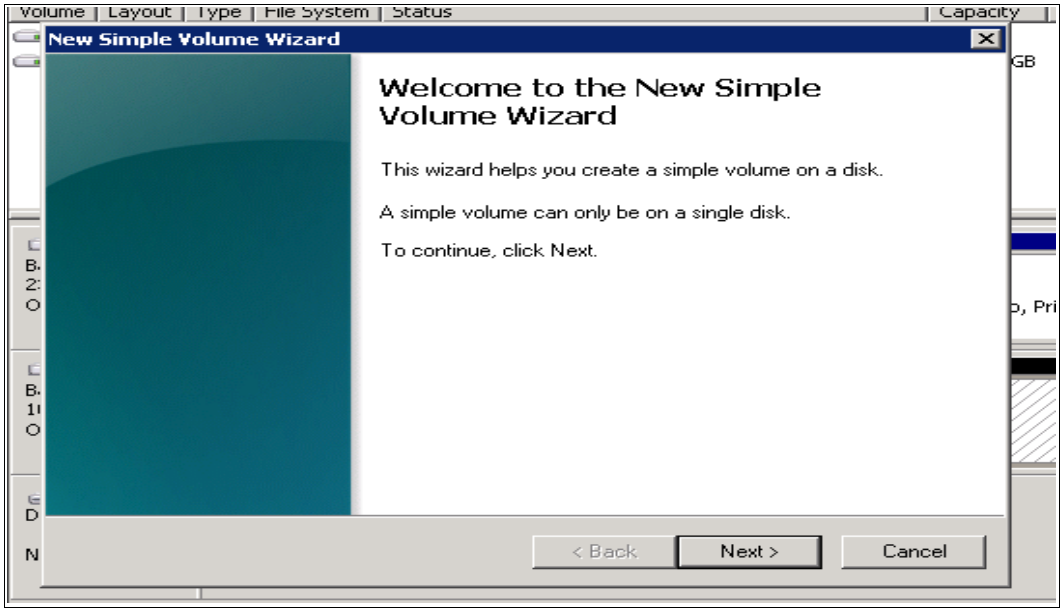

Figure 5-65 Simple volume wizard

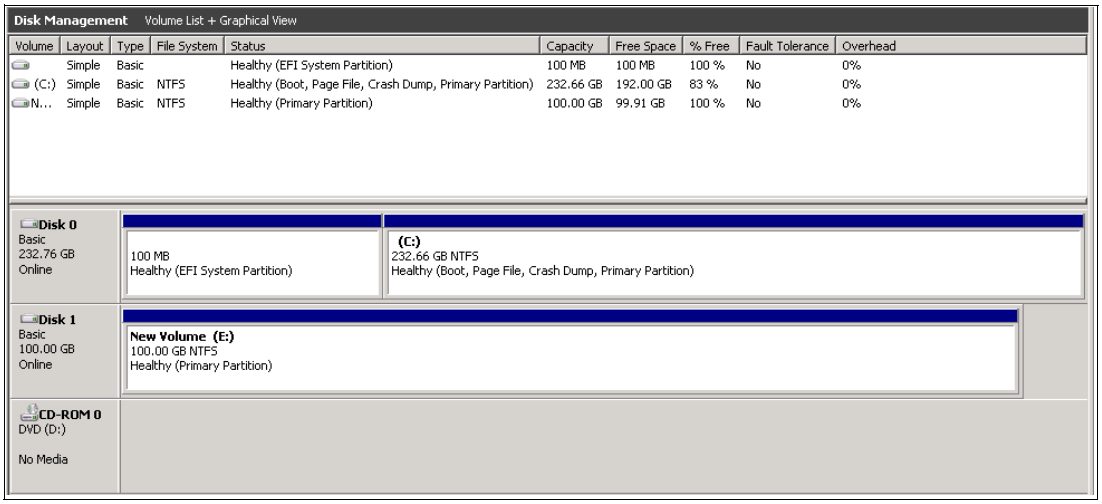

Figure 5-66 SAS attached volume ready to use

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.

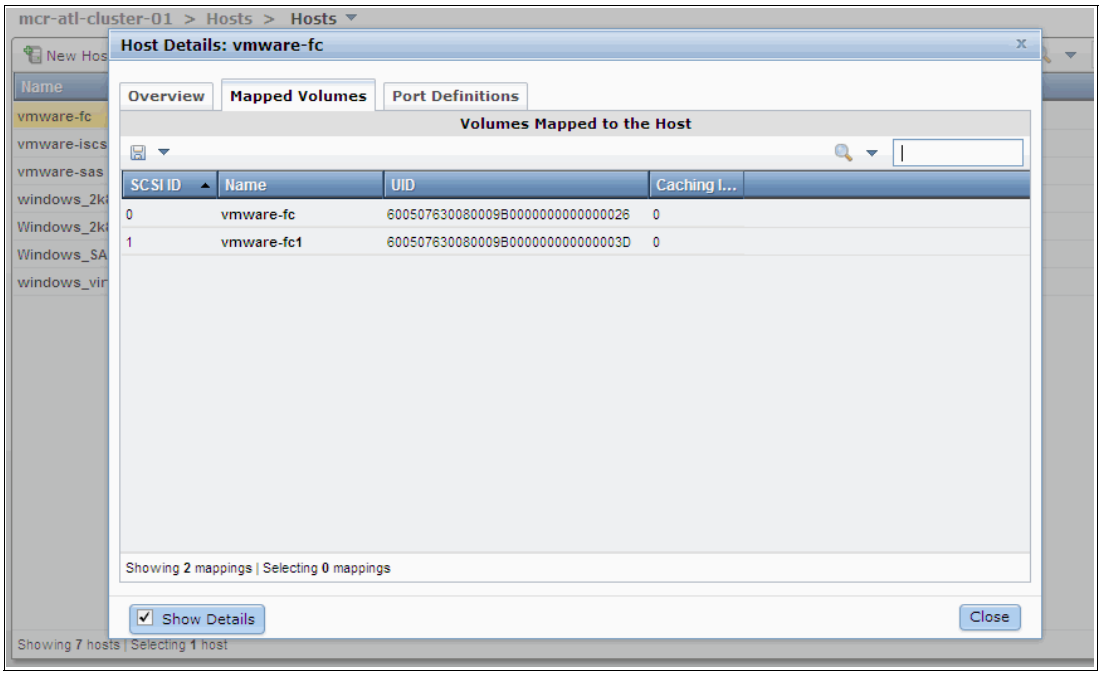

Figure 5-68 Mapped volumes to ESX FC host

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.

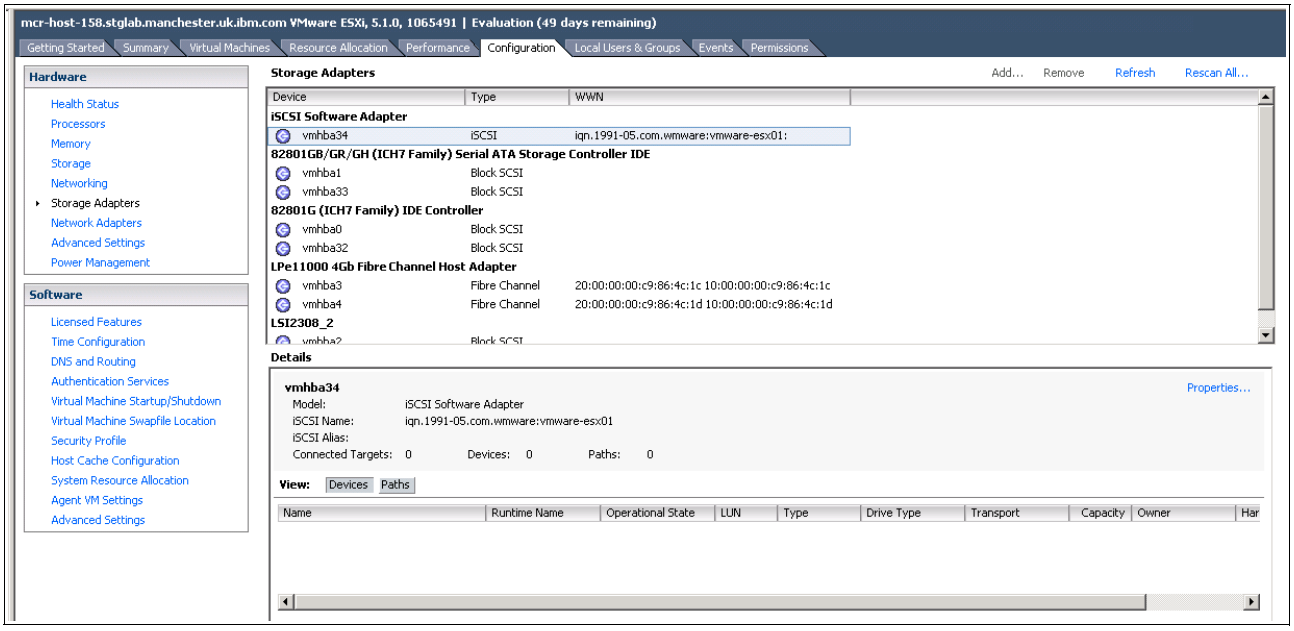

Figure 5-69 vSphere Client: Storage adapters

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.

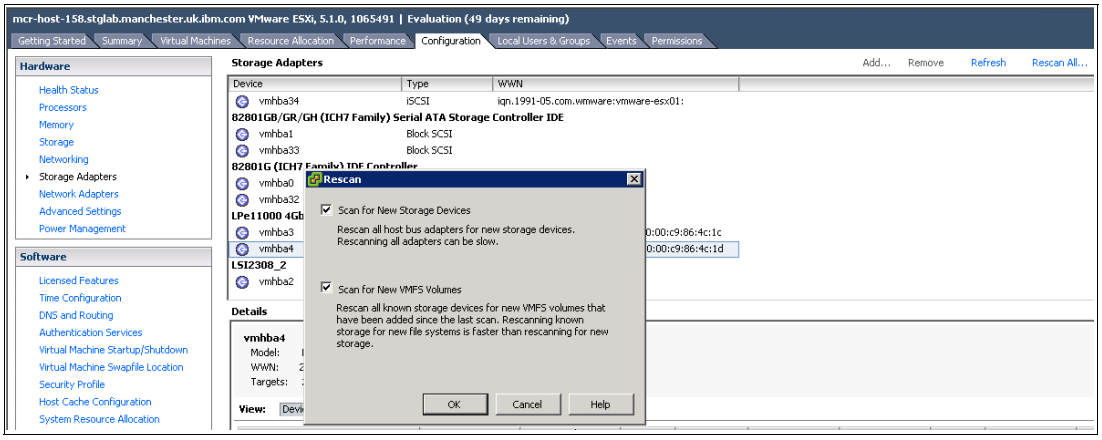

Figure 5-70 Rescan

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.

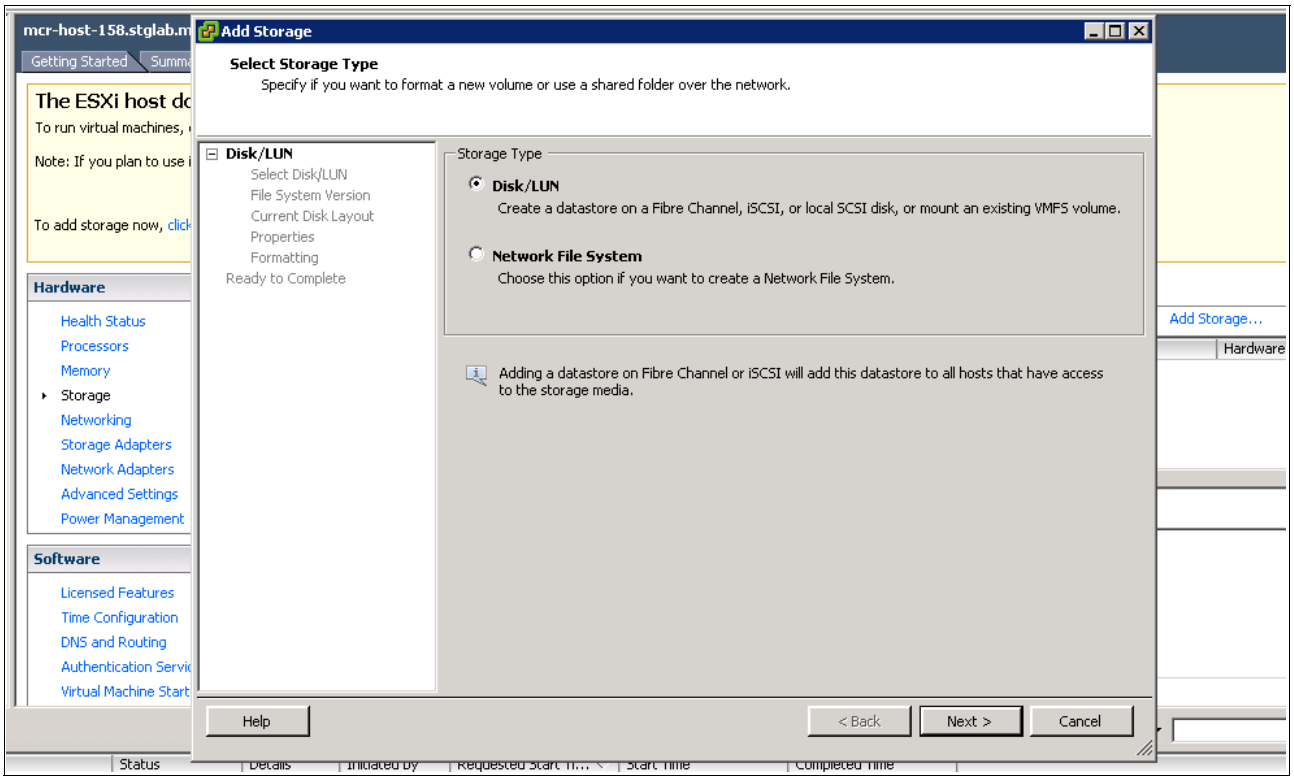

Figure 5-71 vSphere Client: Add Storage

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**.

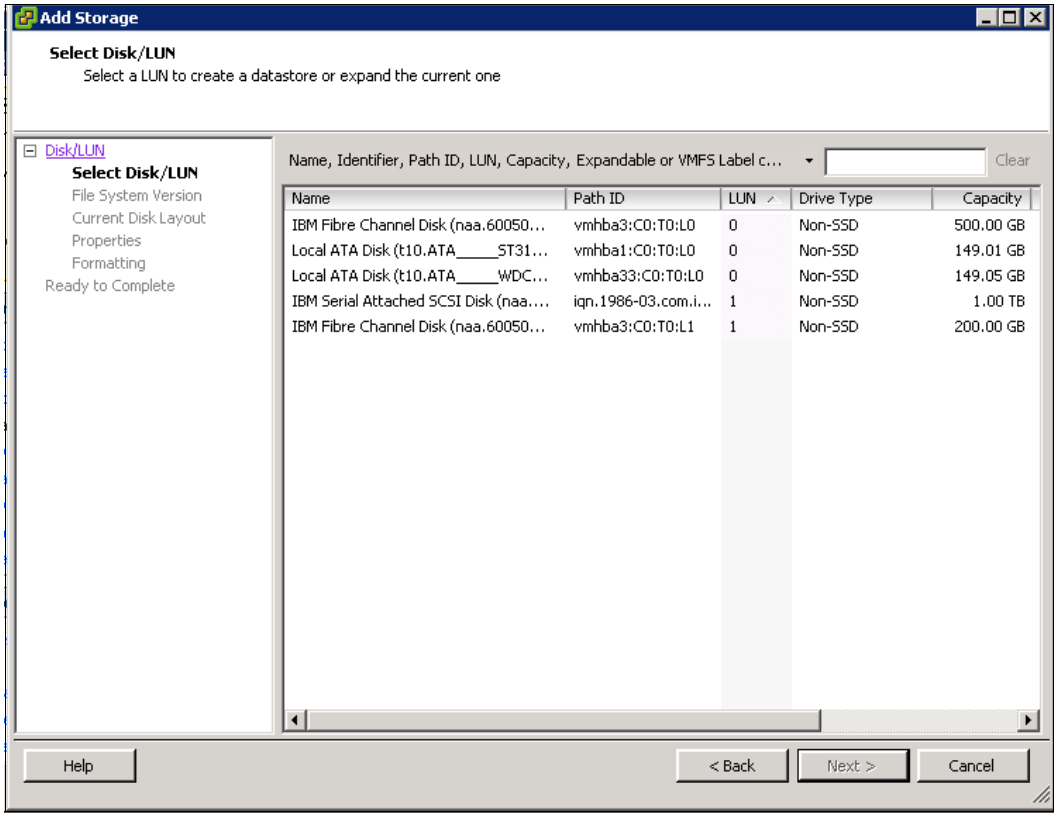

Figure 5-72 Select FC Disk

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

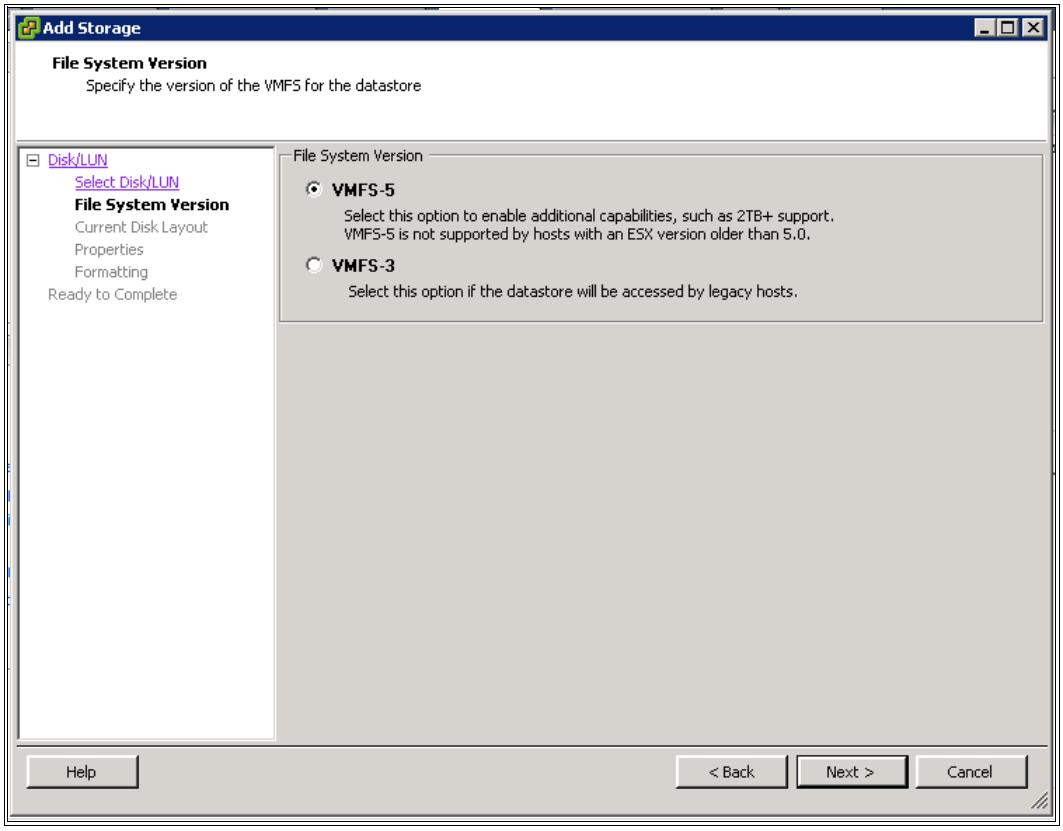

Figure 5-73 Select File System Version

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.

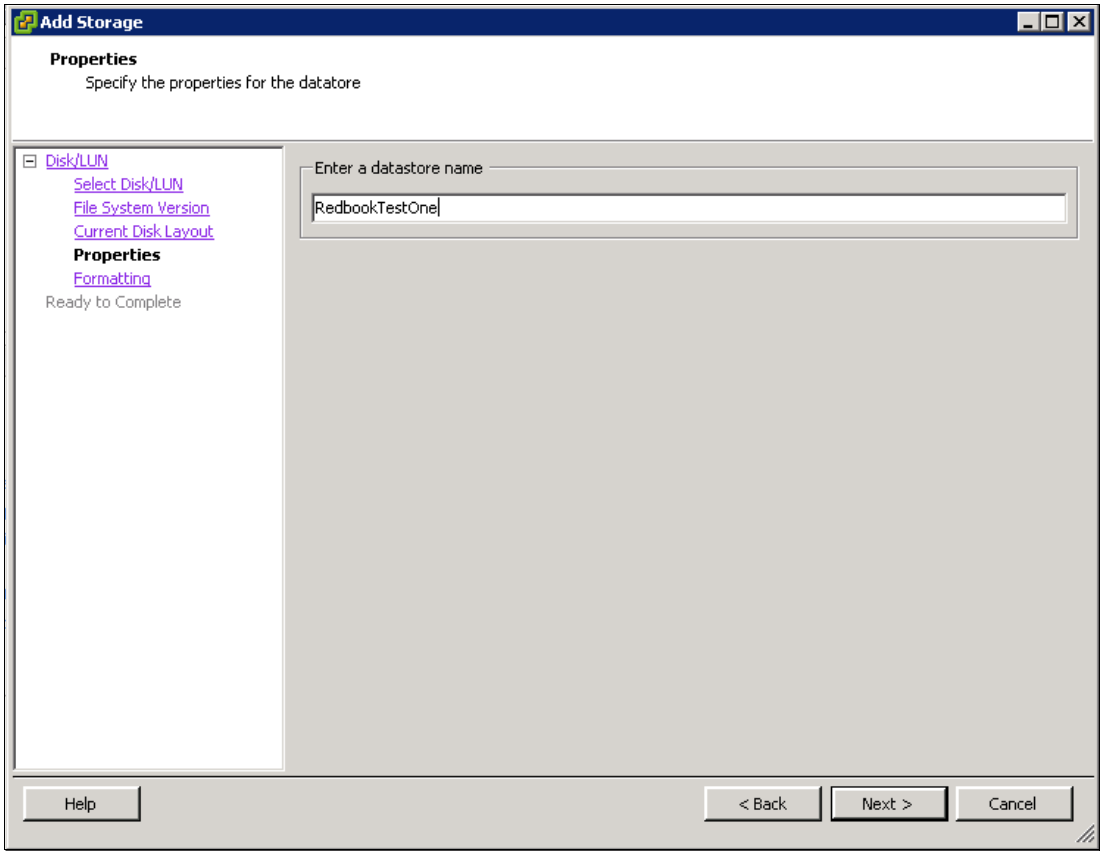

Figure 5-74 Enter a Datastore name

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.

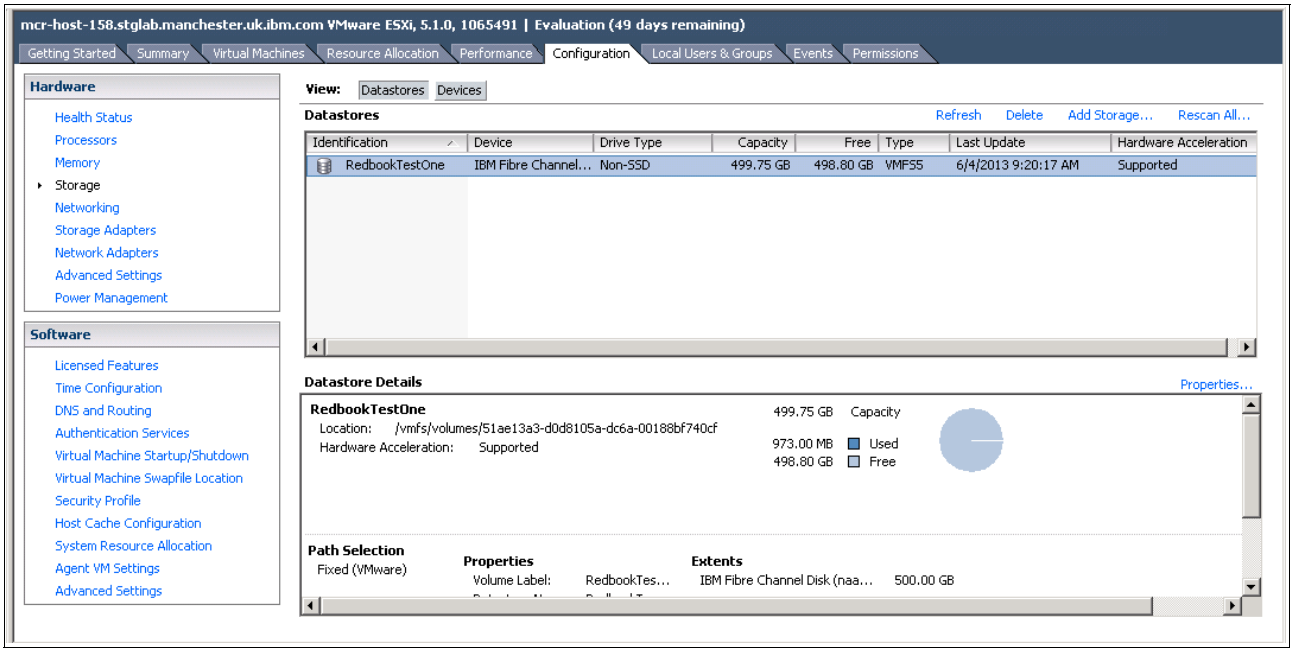

Figure 5-75 Add Storage task complete

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.

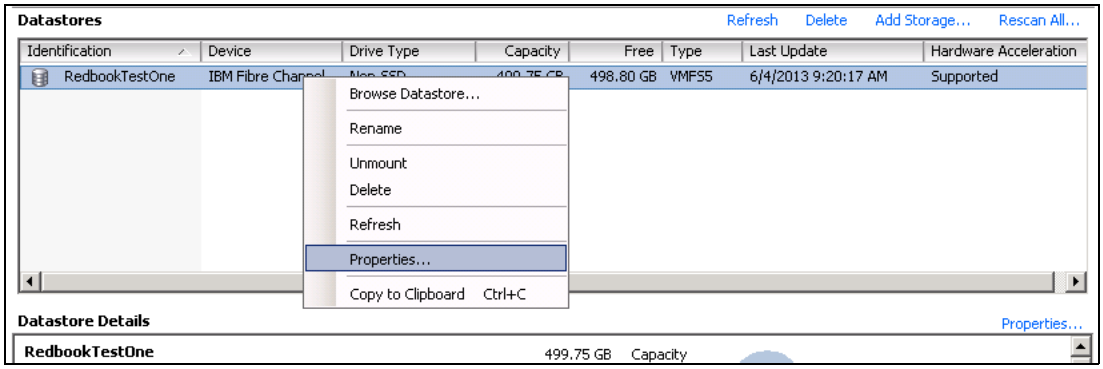

Figure 5-76 Datastore properties

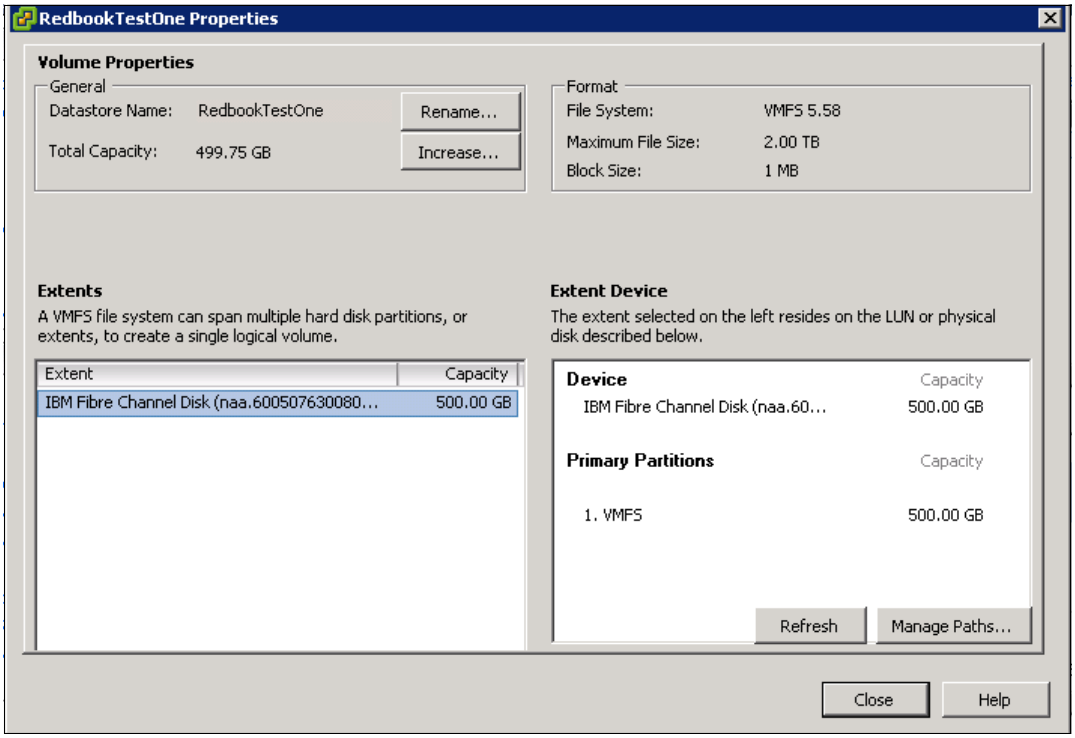

Figure 5-77 Datastore property details

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

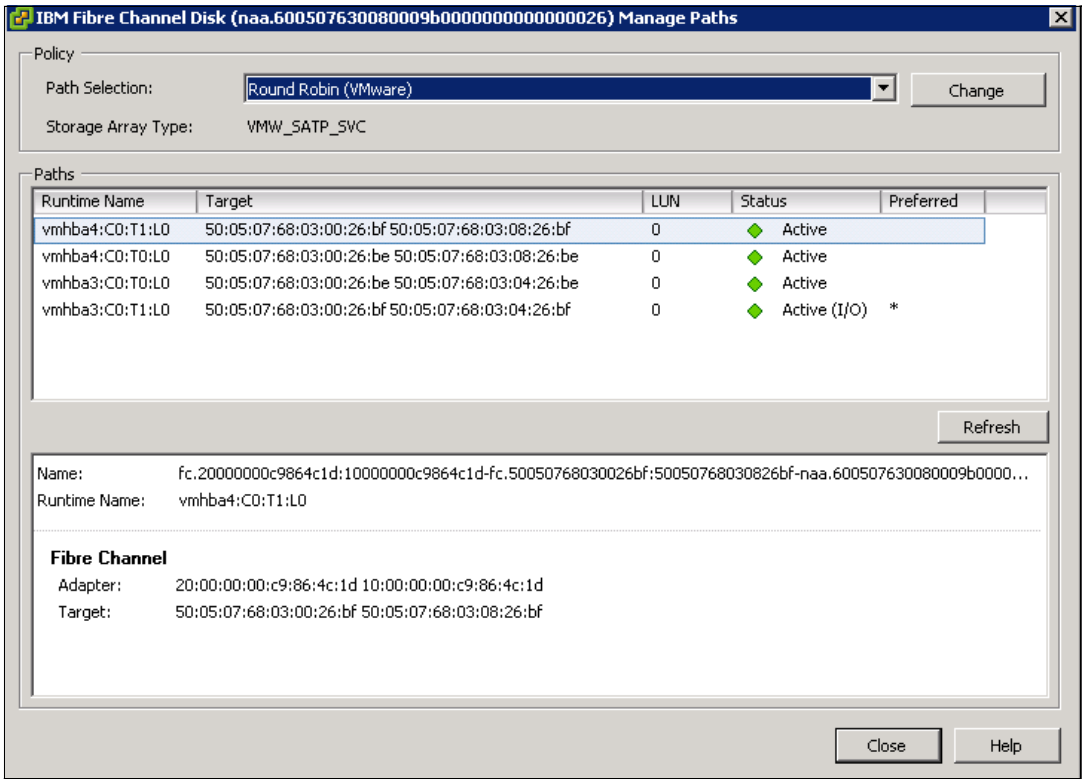

Figure 5-78 Select a Datastore multipath setting

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.

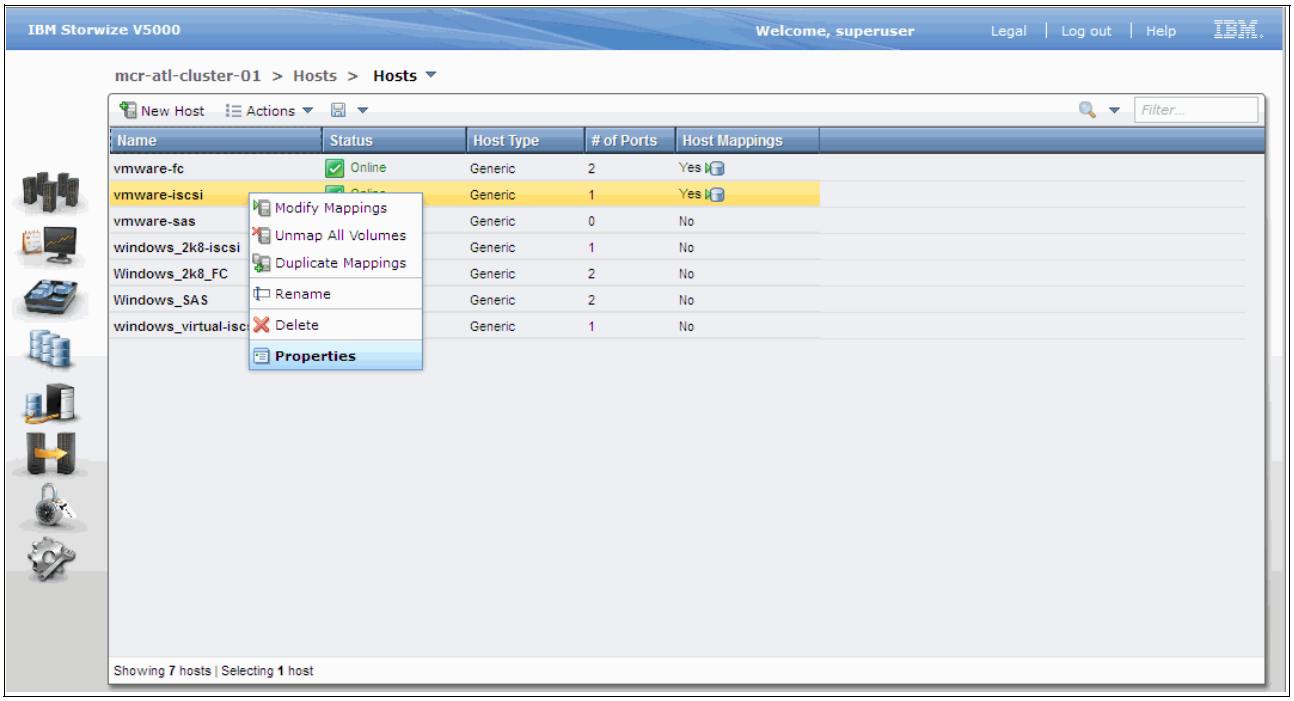

Figure 5-79 Select iSCSI ESX host properties

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

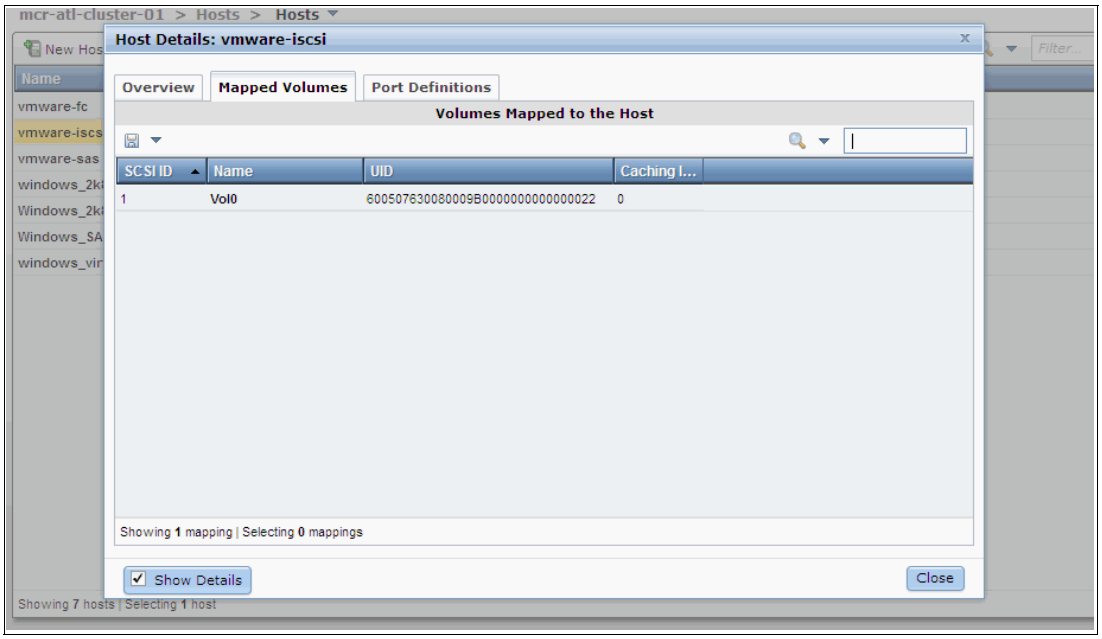

Figure 5-80 iSCSI ESX host properties

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.

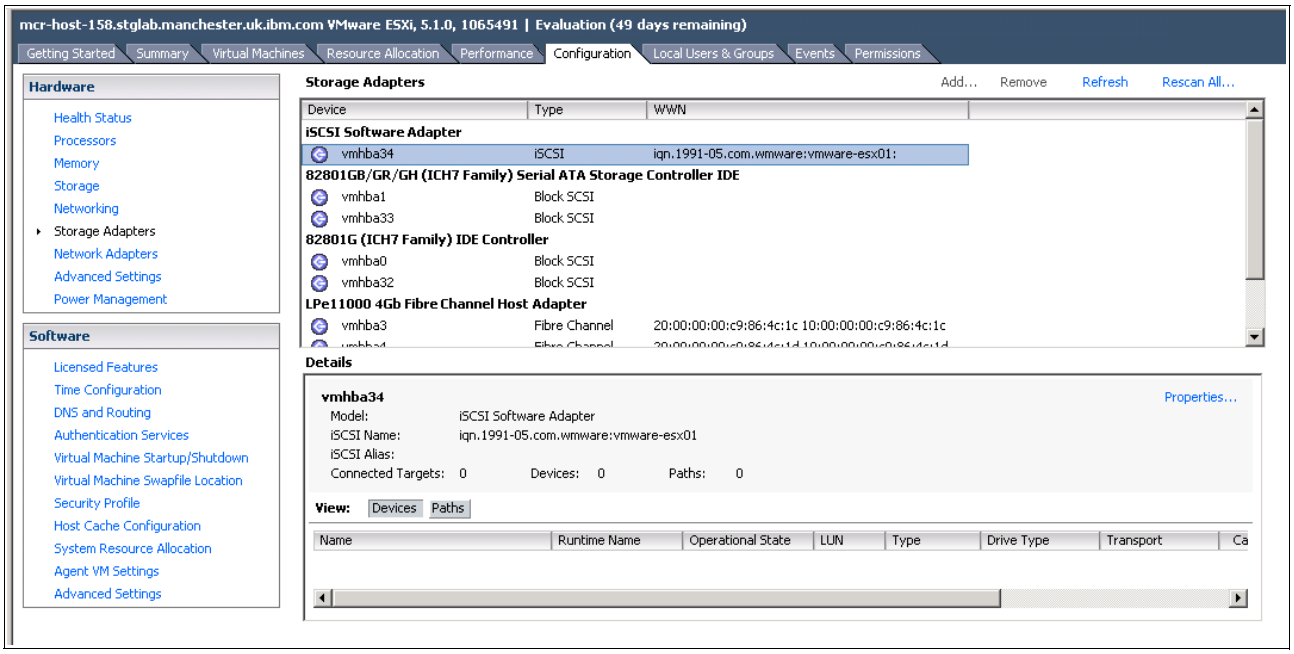

Figure 5-81 vSphere Client: Storage Adapters

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**.

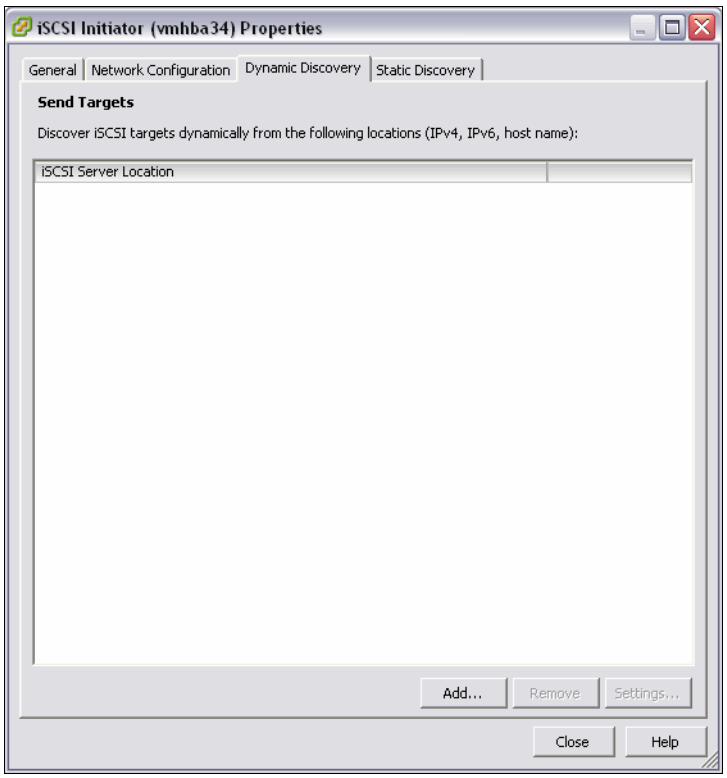

Figure 5-82 iSCSI Initiator properties

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.

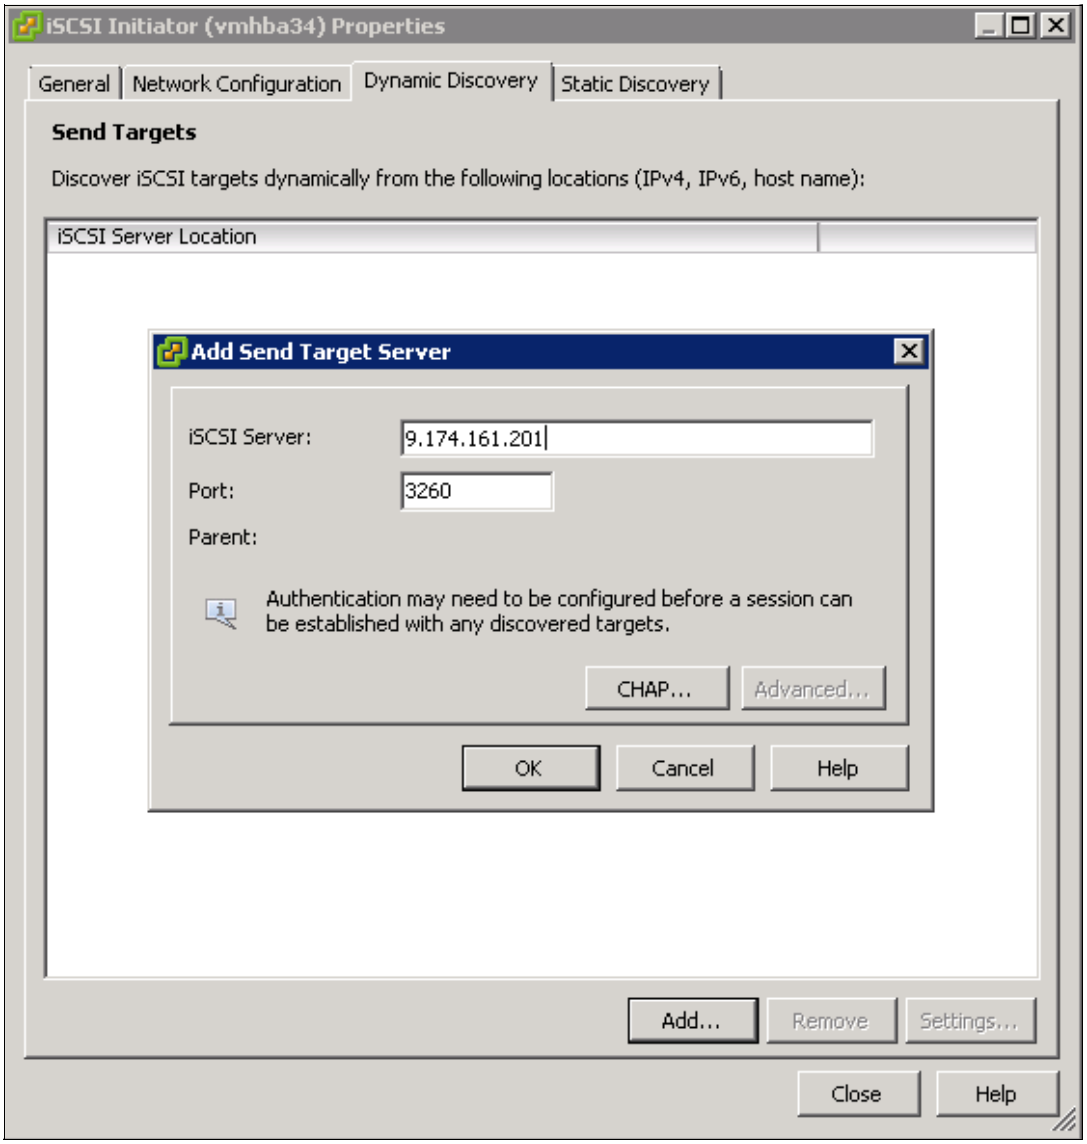

Figure 5-83 Enter a target IP address

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.

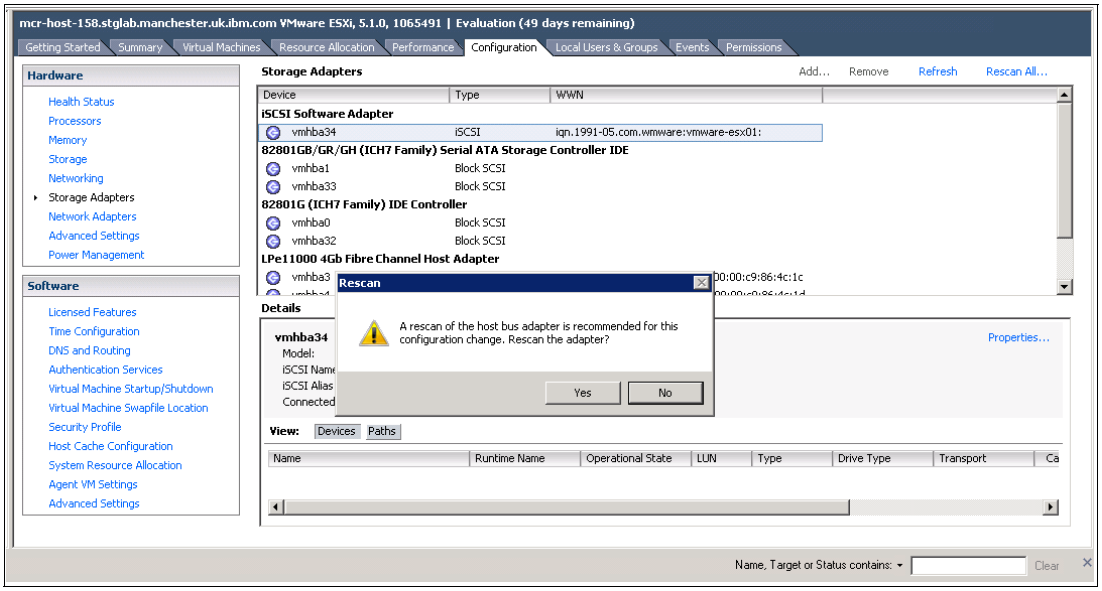

Figure 5-84 Confirm the rescan

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**.

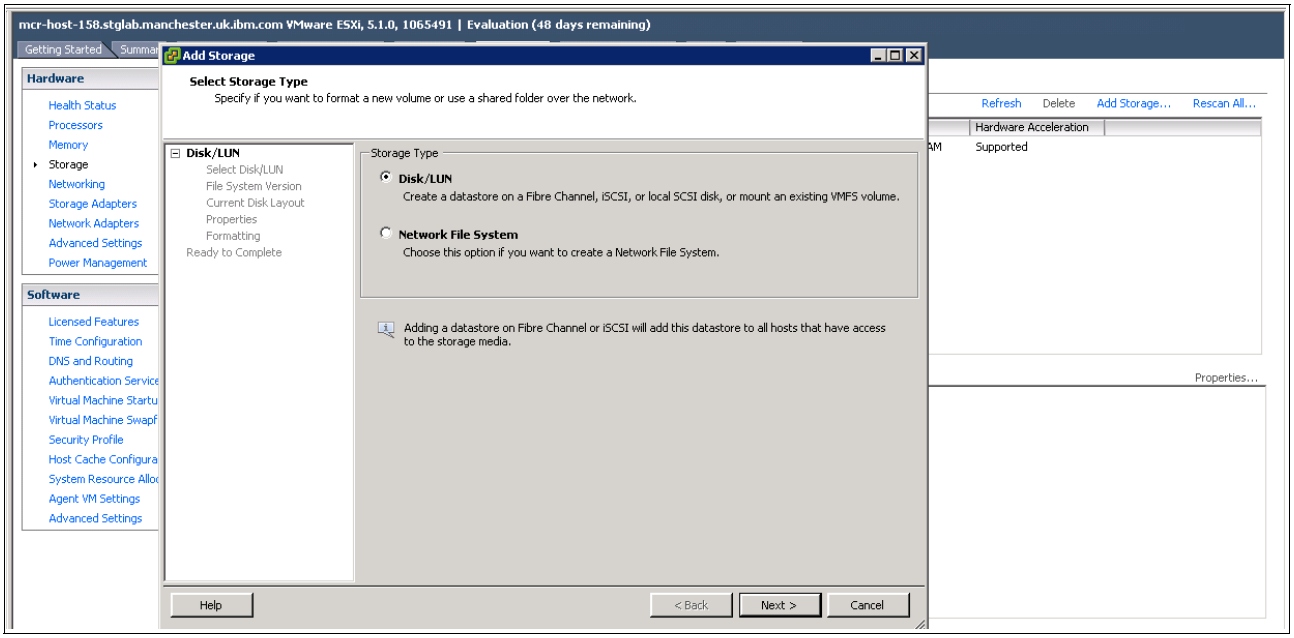

Figure 5-85 vSphere Client: Add Storage

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

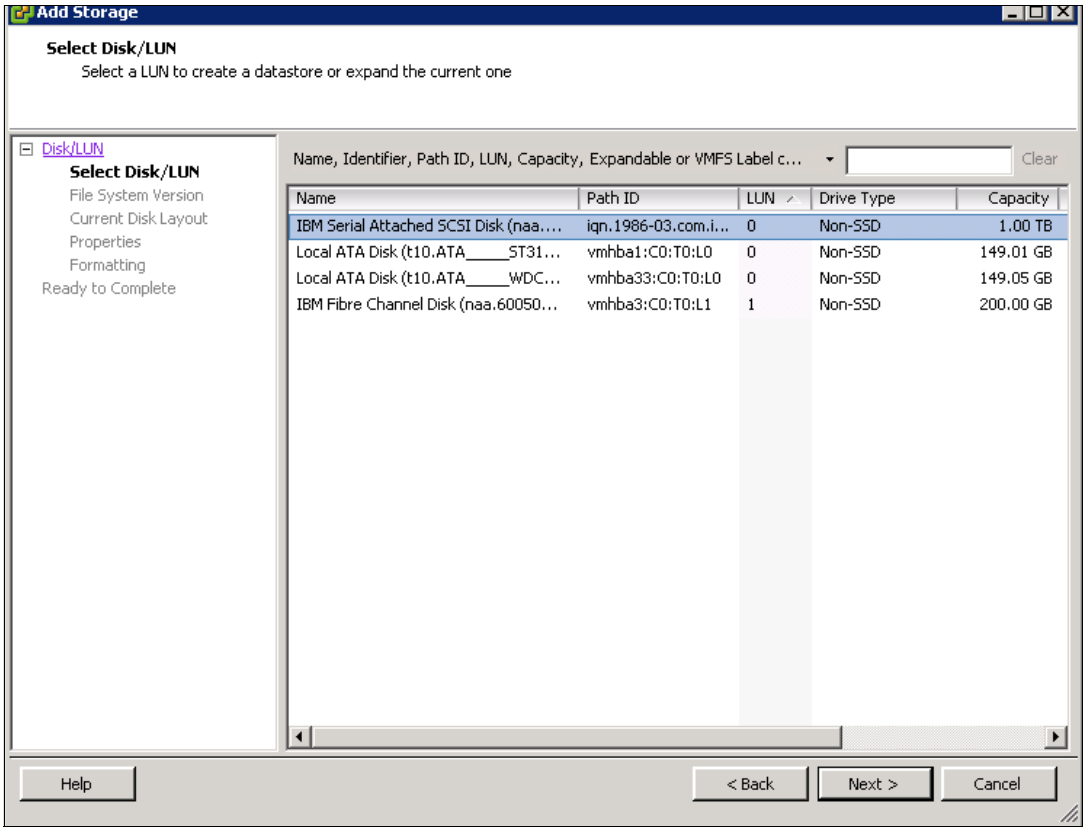

Figure 5-86 Select iSCSI LUN

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

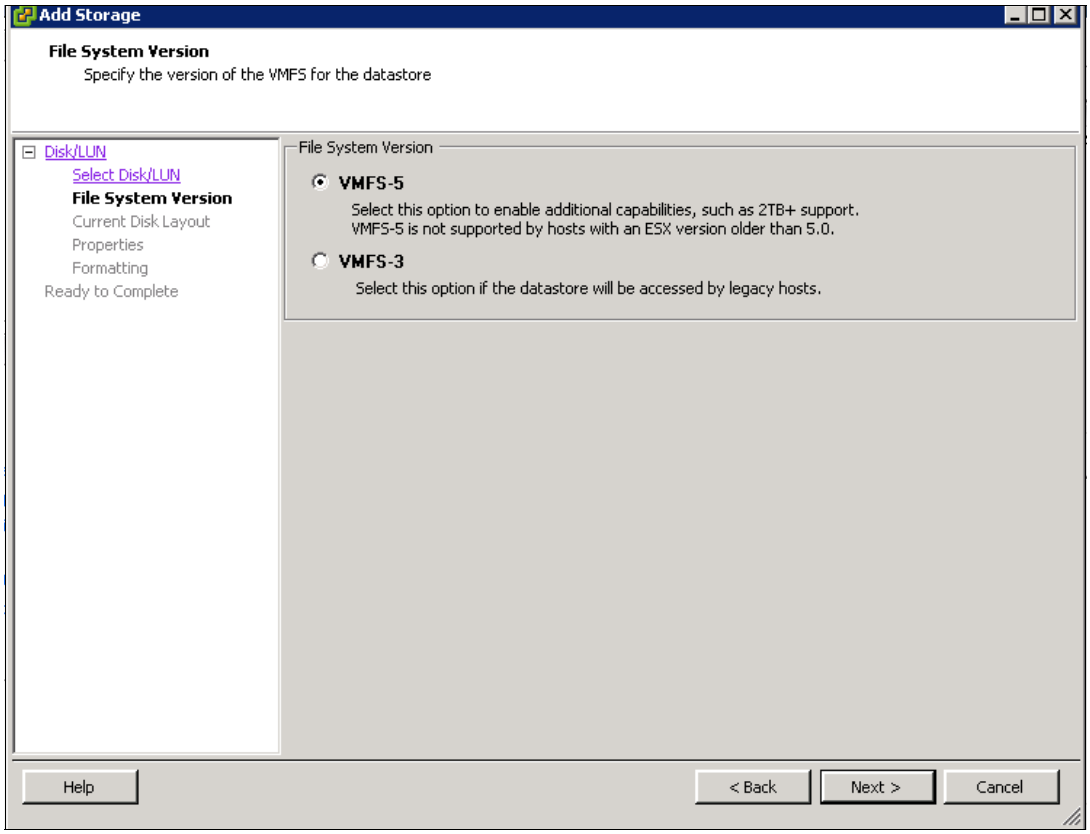

Figure 5-87 Select File System Version

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

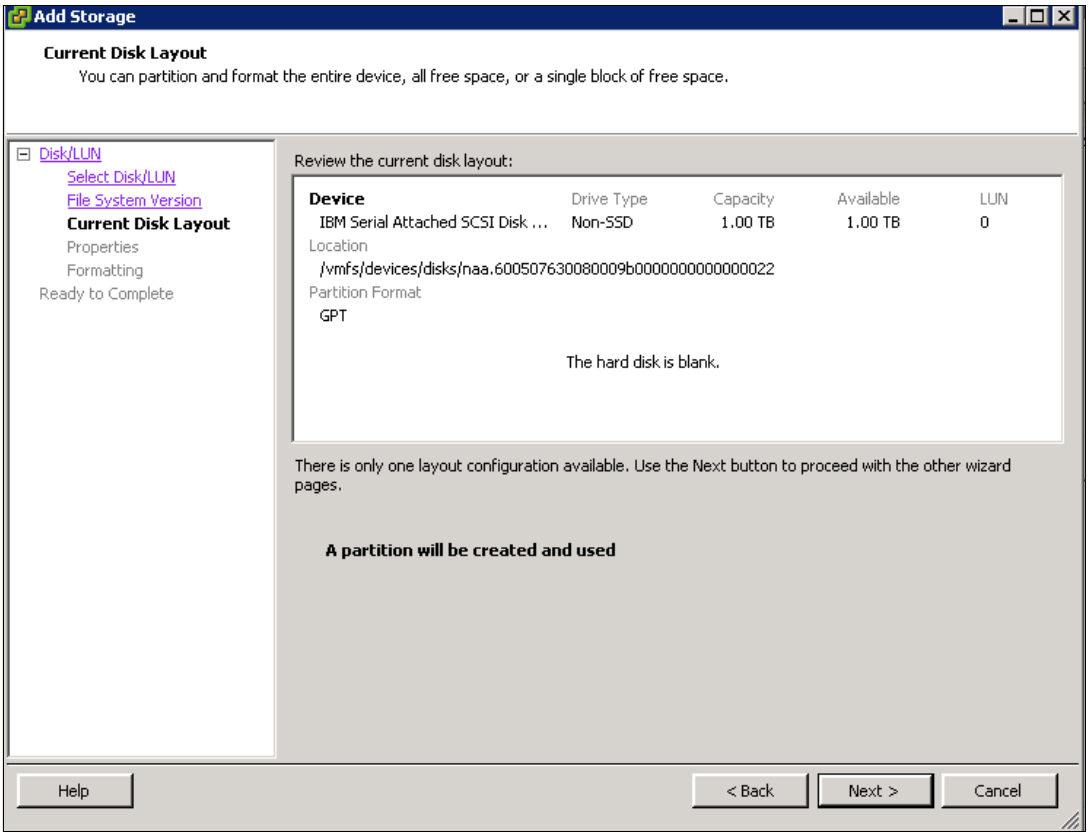

Figure 5-88 Current Disk Layout

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

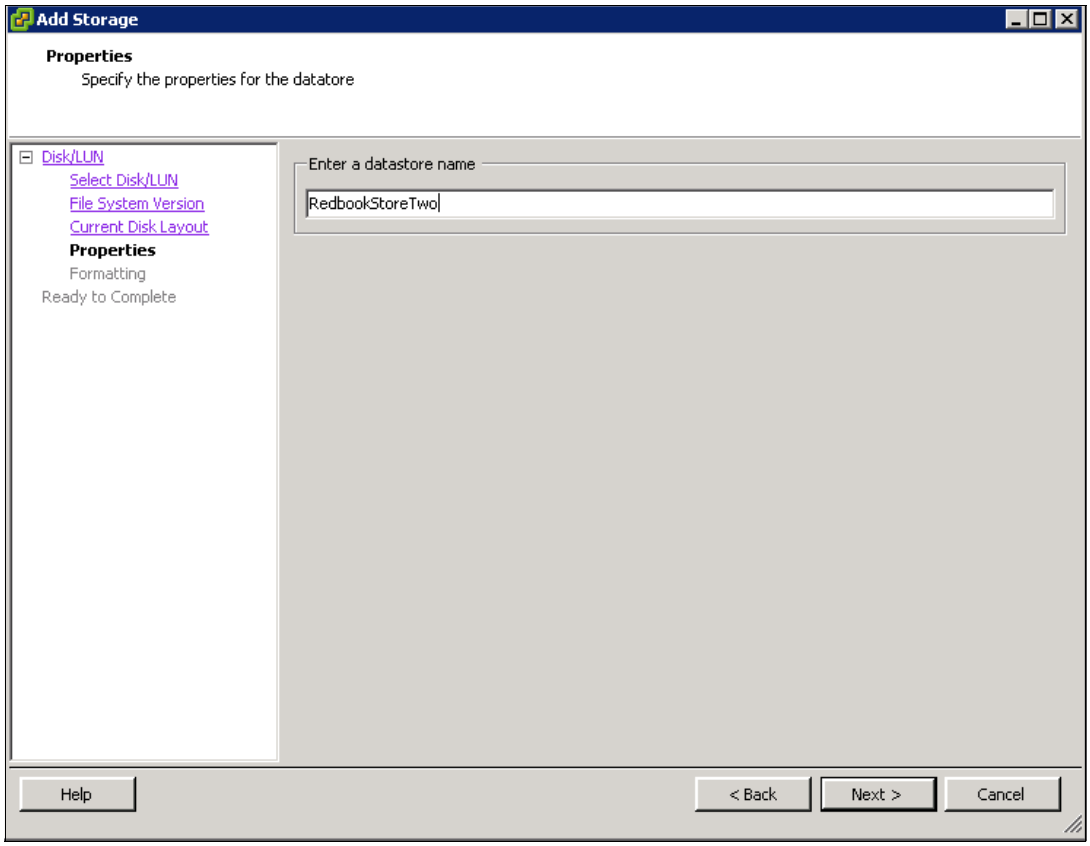

Figure 5-89 Enter a Datastore name

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

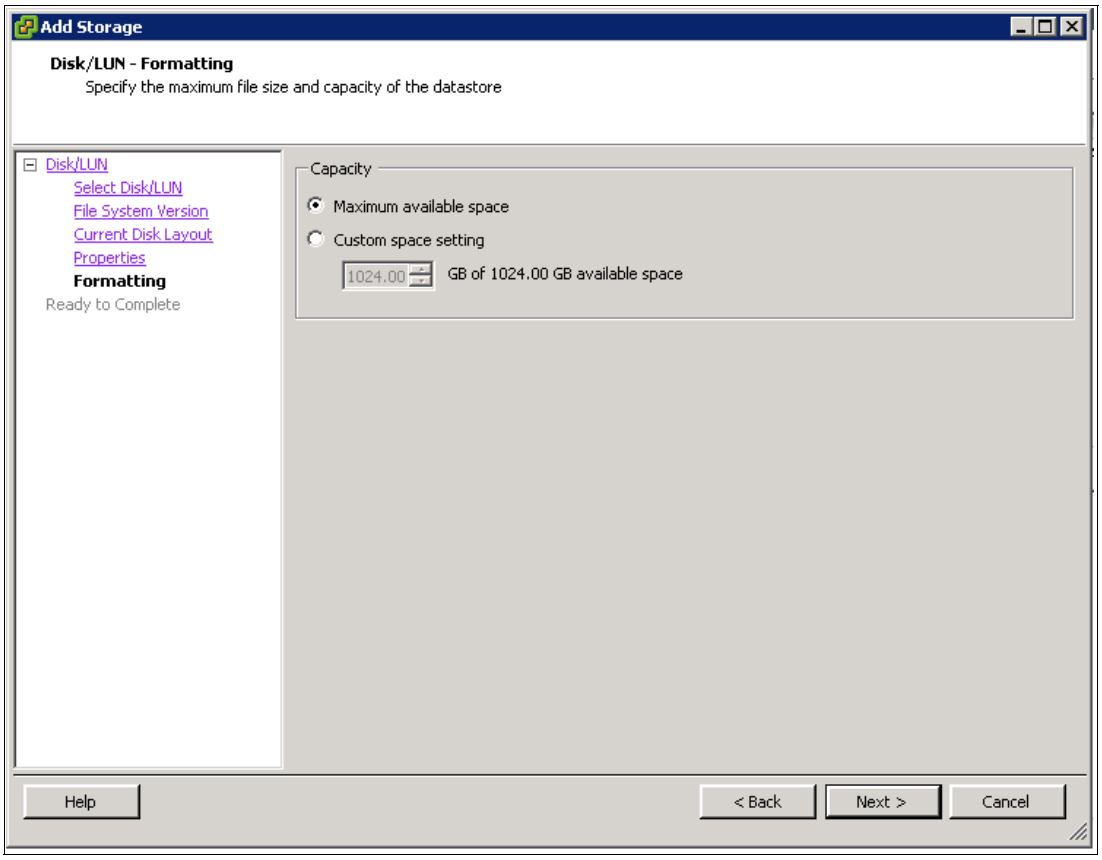

Figure 5-90 Capacity

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

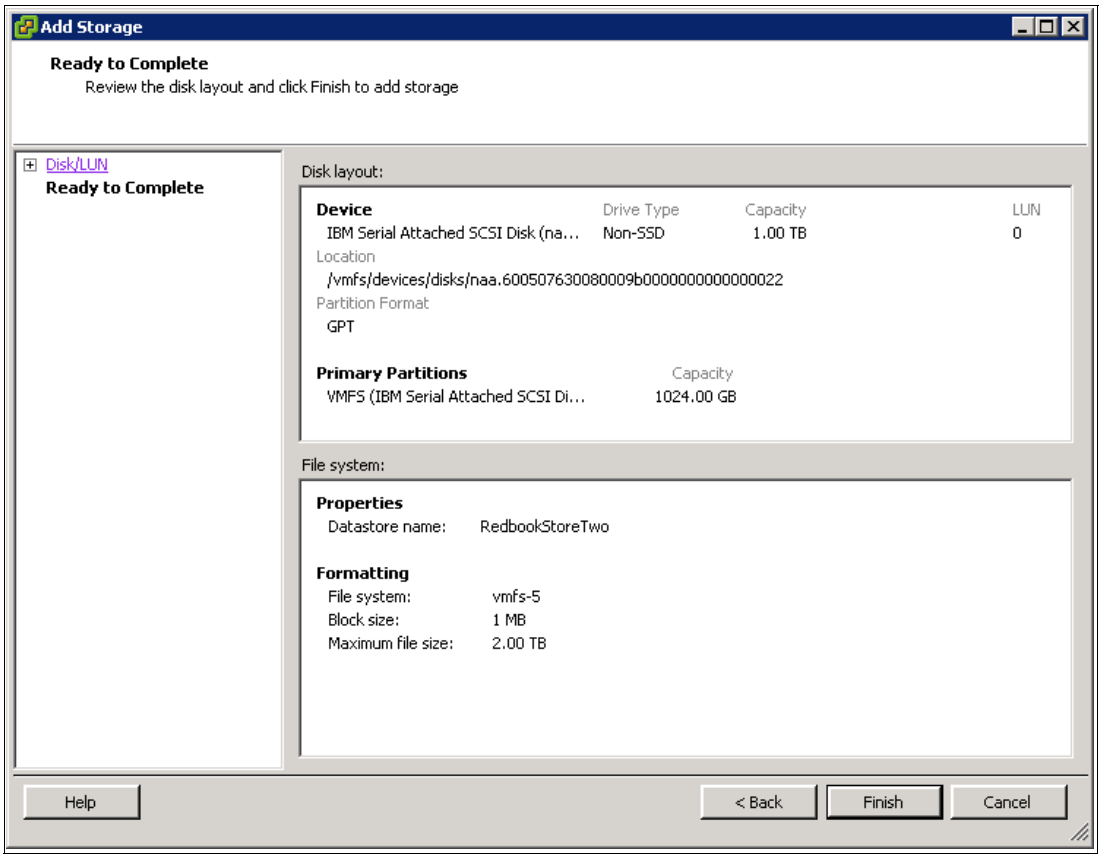

Figure 5-91 Finish the wizard

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.

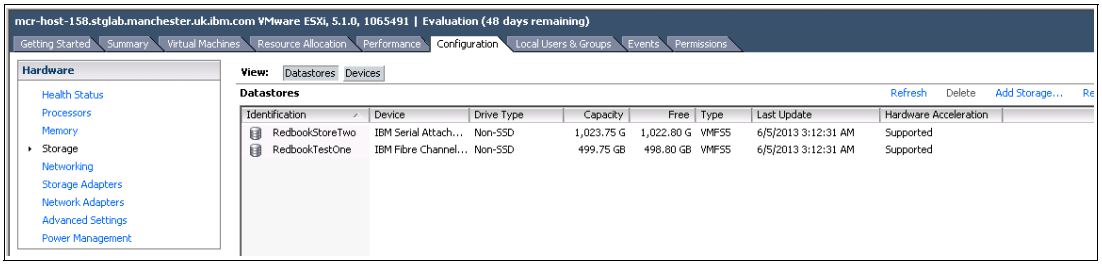

Figure 5-92 New Datastore available

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

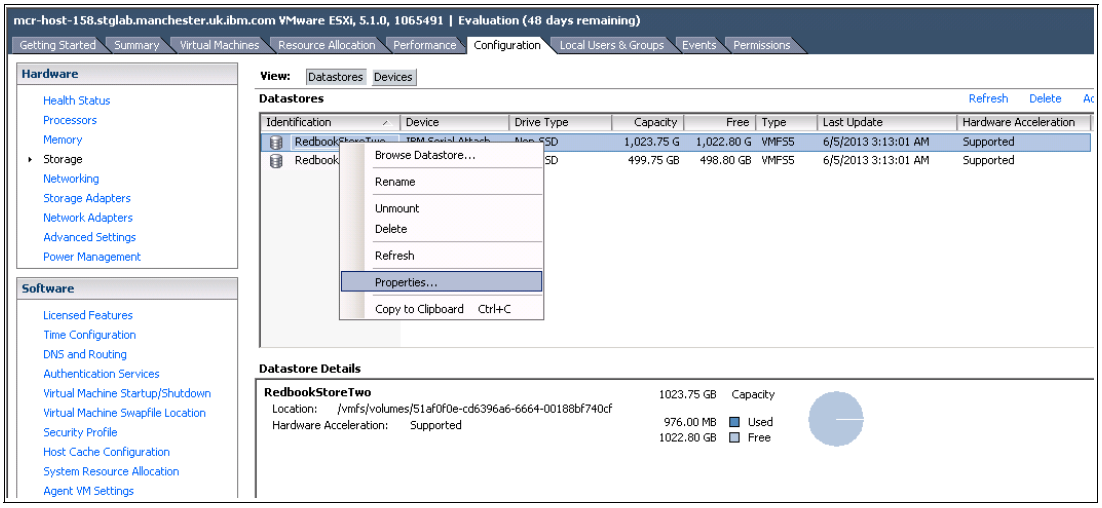

Figure 5-93 Datastore properties

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

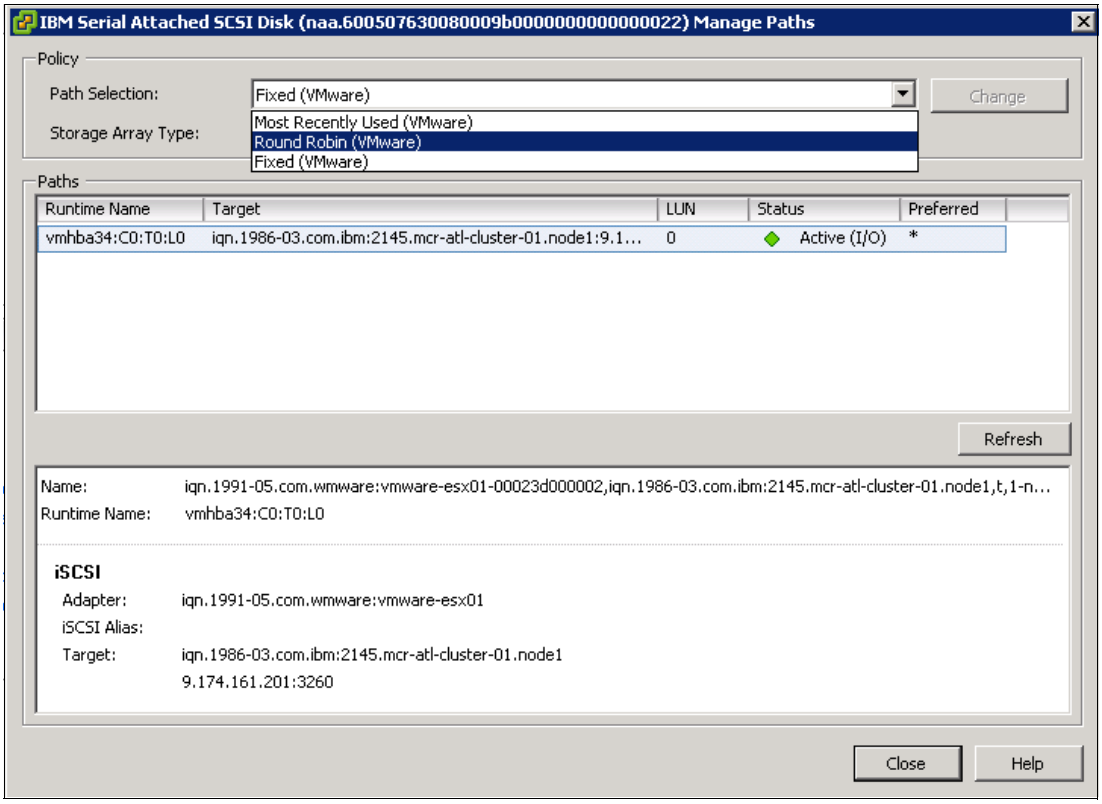

Figure 5-94 Change the multipath policy

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.

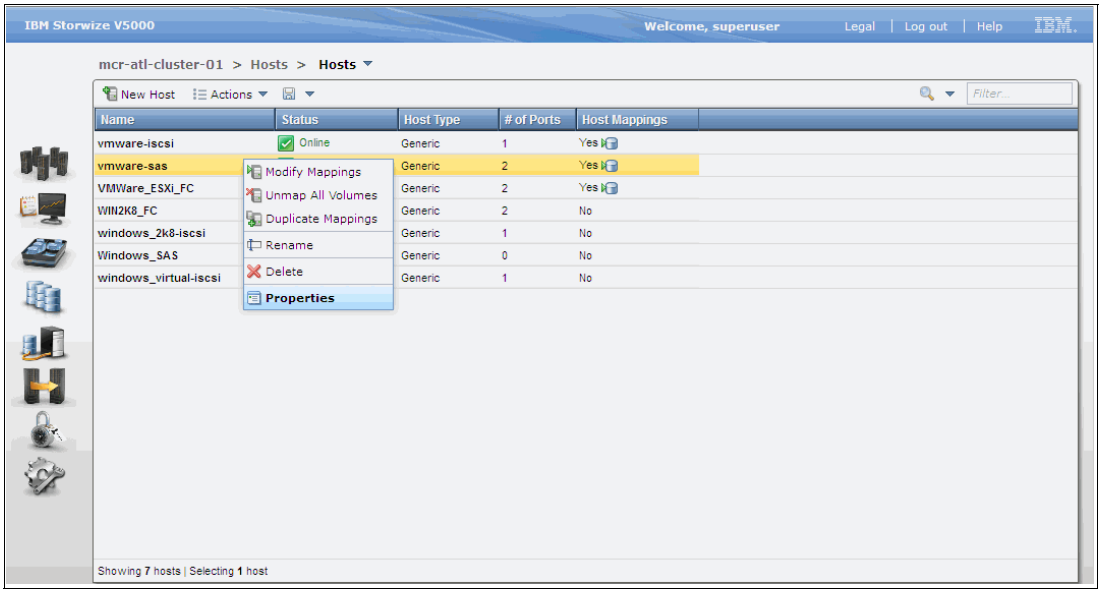

Figure 5-95 Example ESX SAS host

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

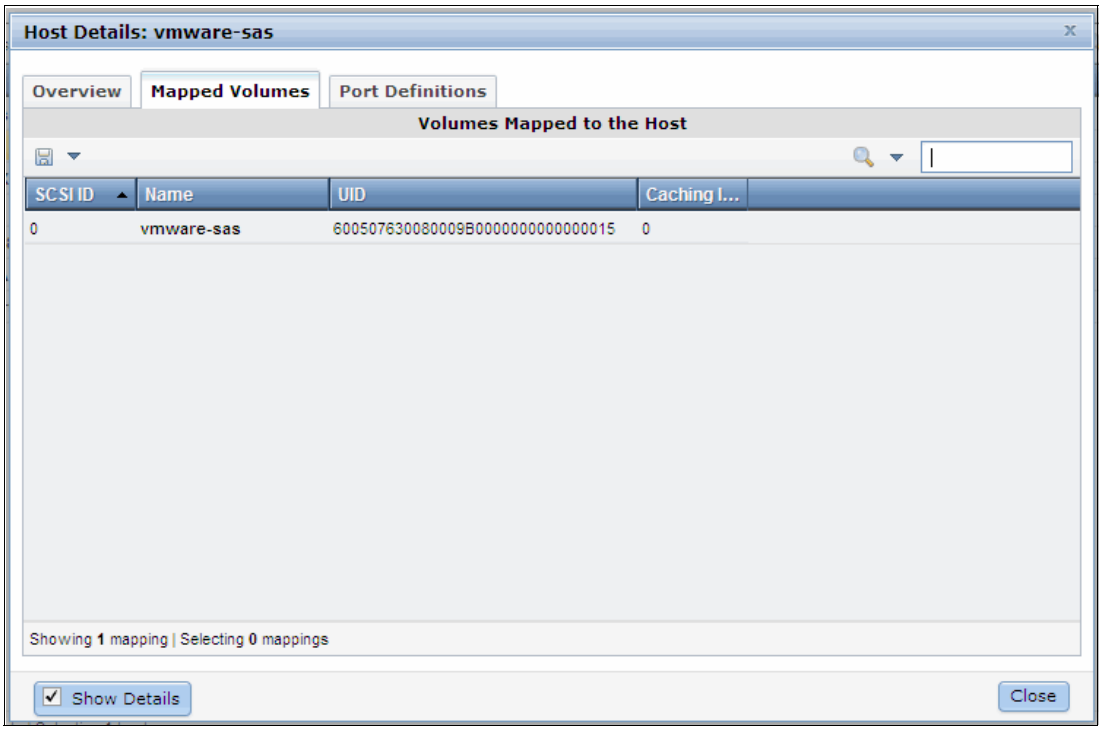

Figure 5-96 Mapped volumes to ESX SAS host

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.

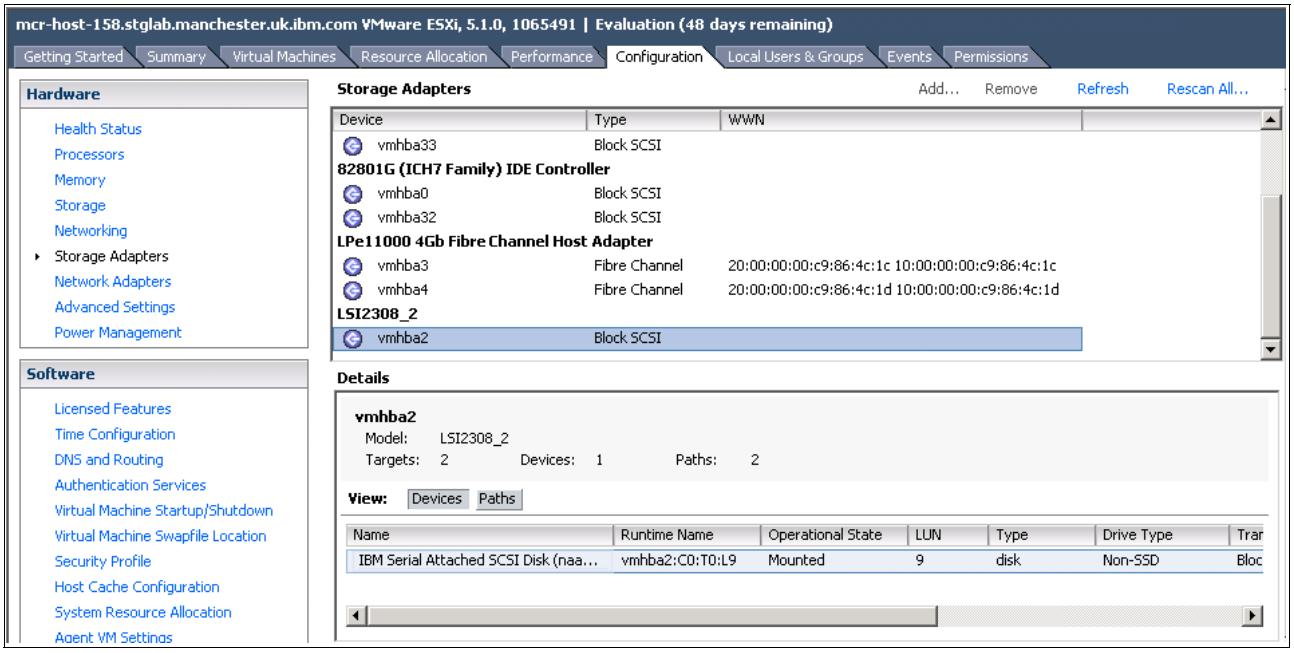

Figure 5-97 vSphere Client: Storage Adapters

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.

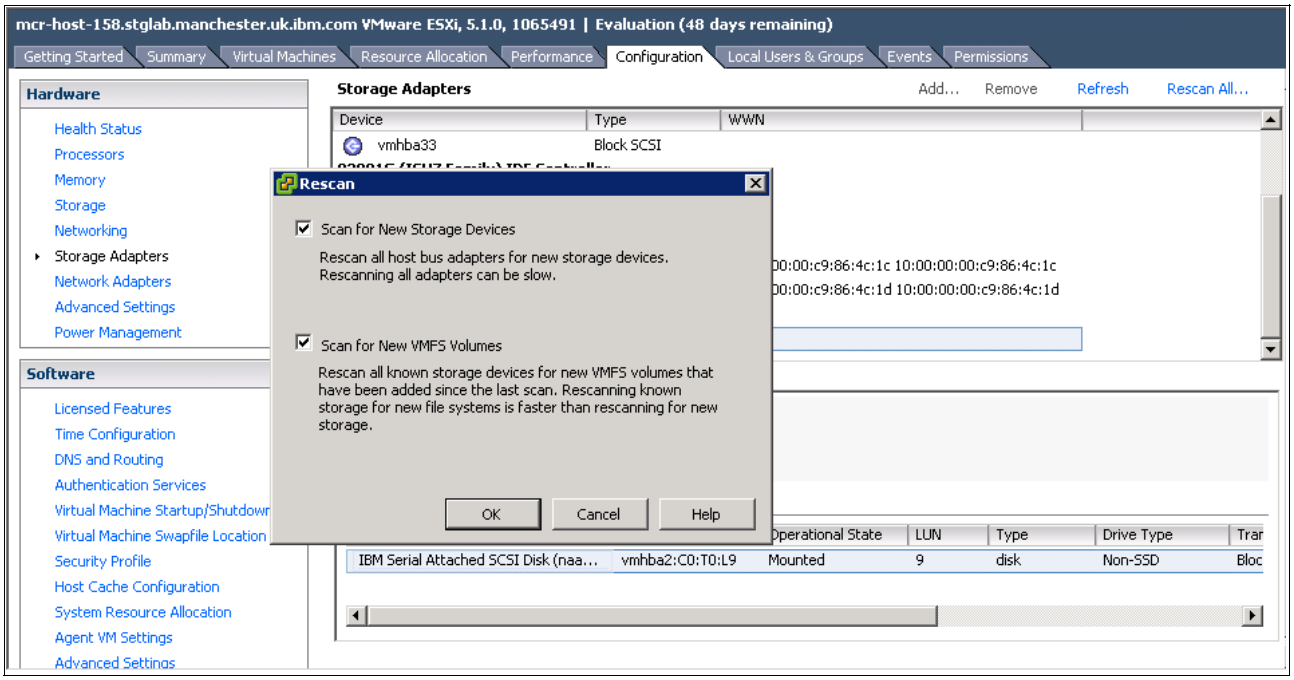

Figure 5-98 Rescan

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.

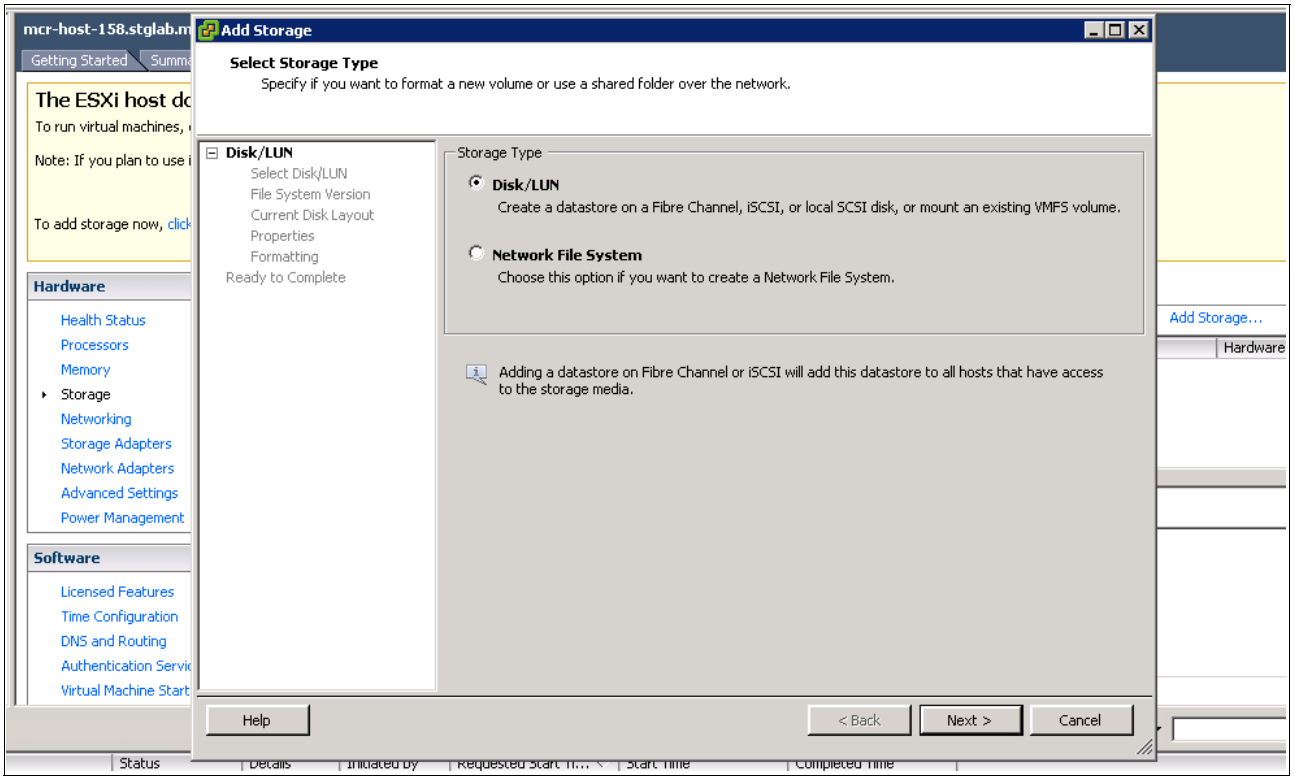

Figure 5-99 vSphere Client: Add Storage

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**.

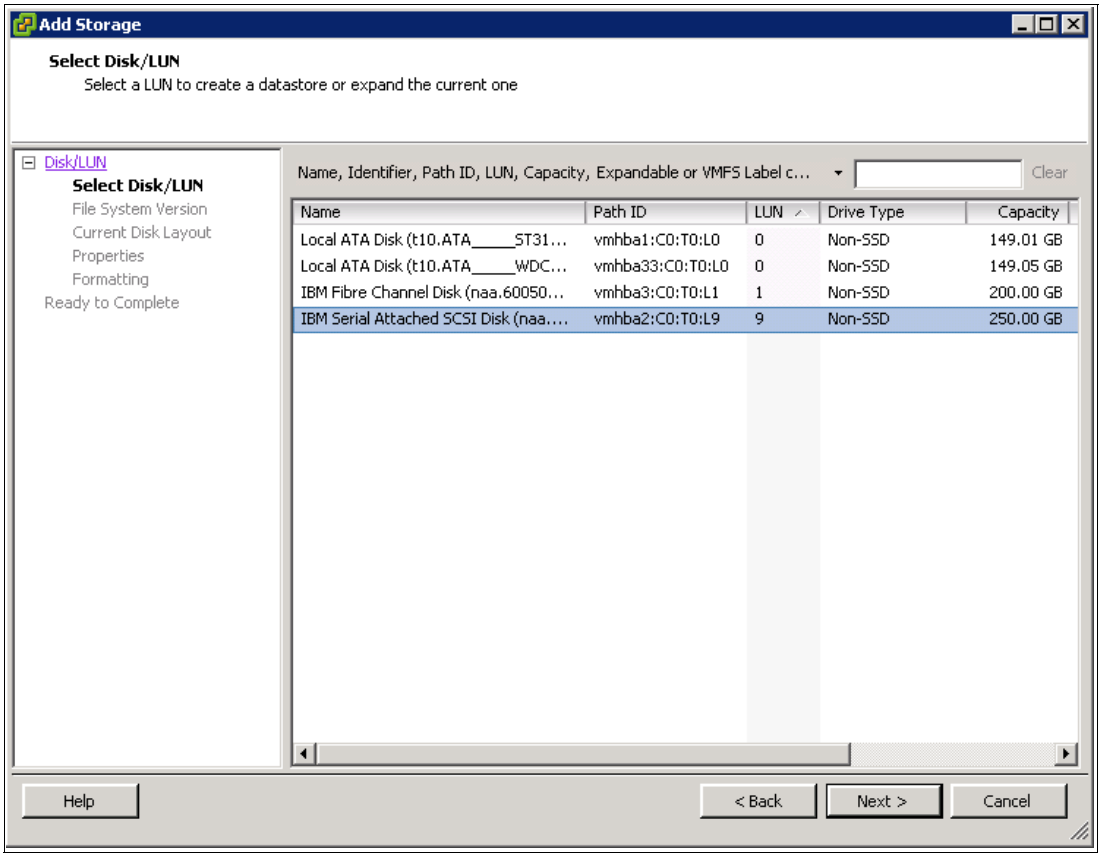

Figure 5-100 Select SAS Disk

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

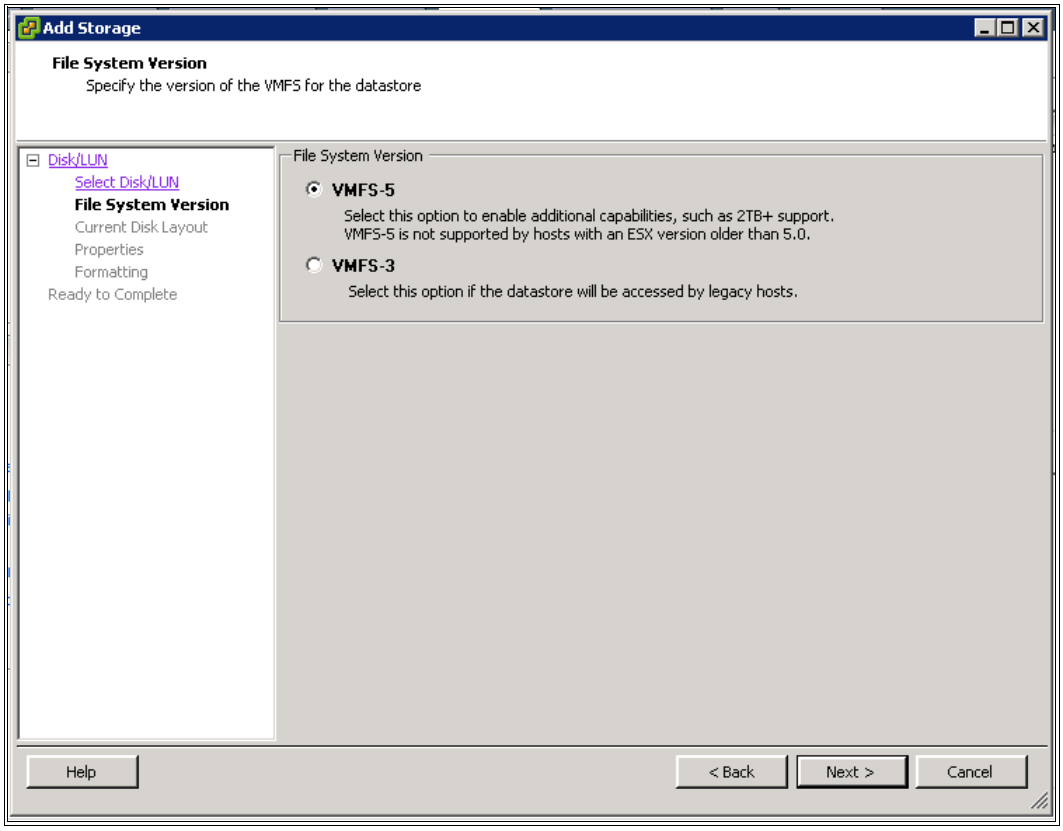

Figure 5-101 Select File System Version

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.

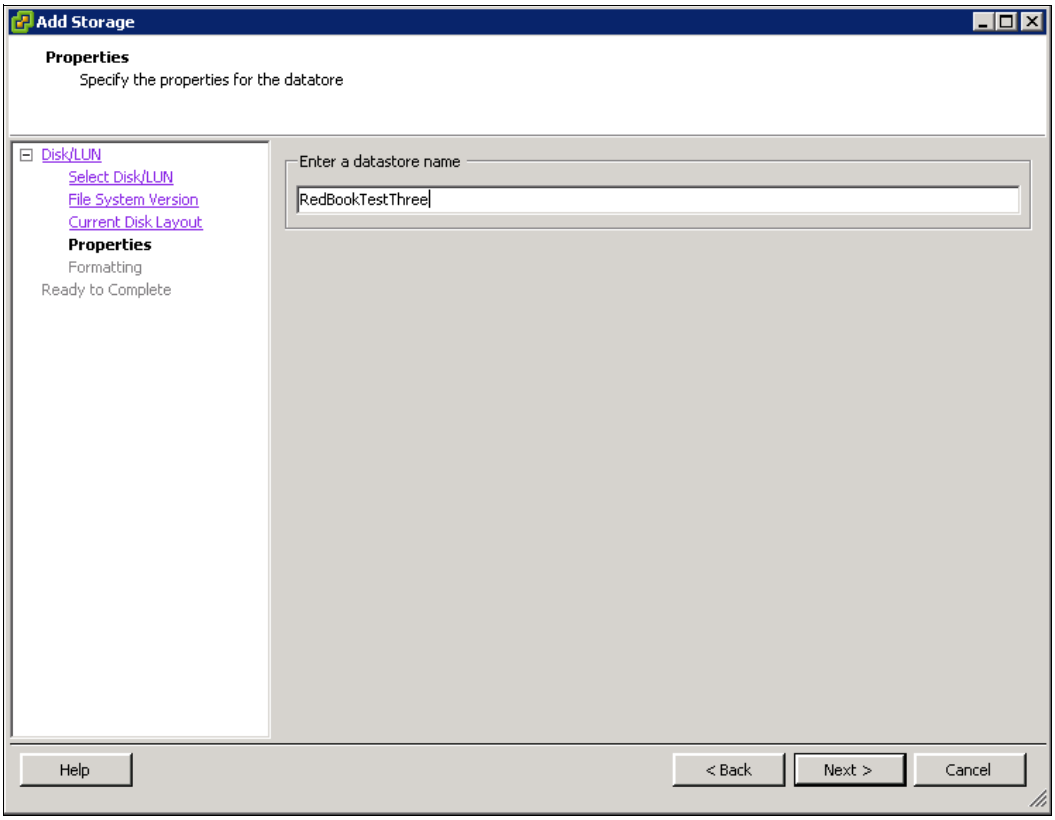

Figure 5-102 Adding Datastore name

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.

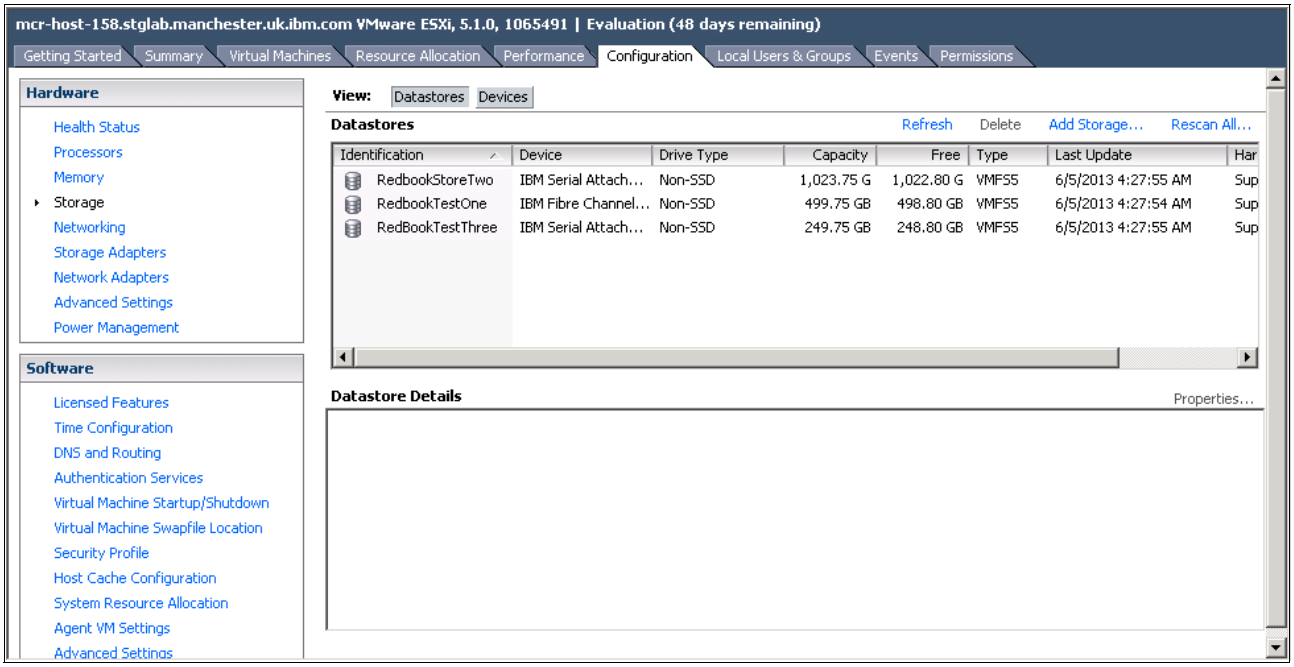

Figure 5-103 Add Storage task complete

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

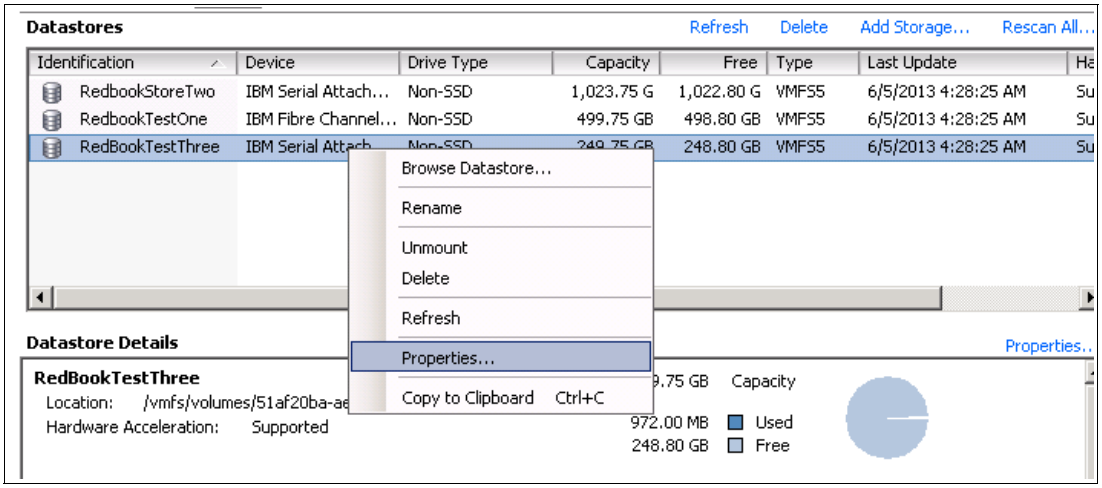

Figure 5-104 Datastore properties

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

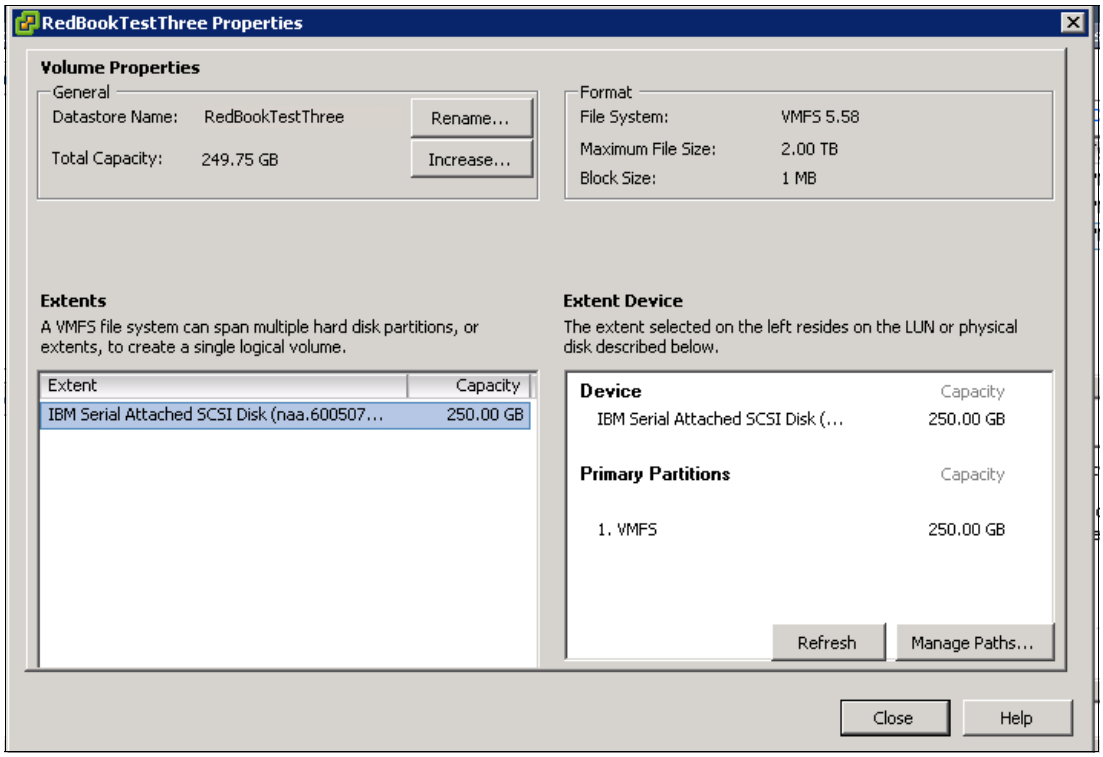

Figure 5-105 Datastore property details

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

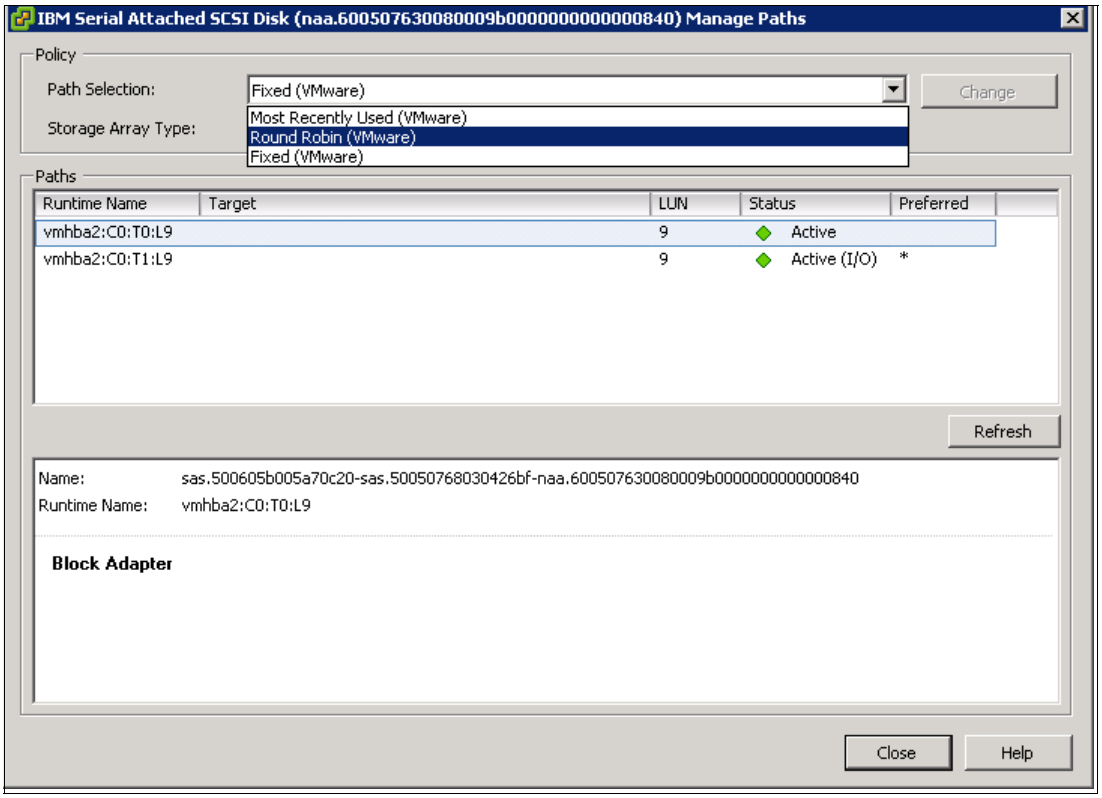

Figure 5-106 Select a Datastore multipath setting

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.

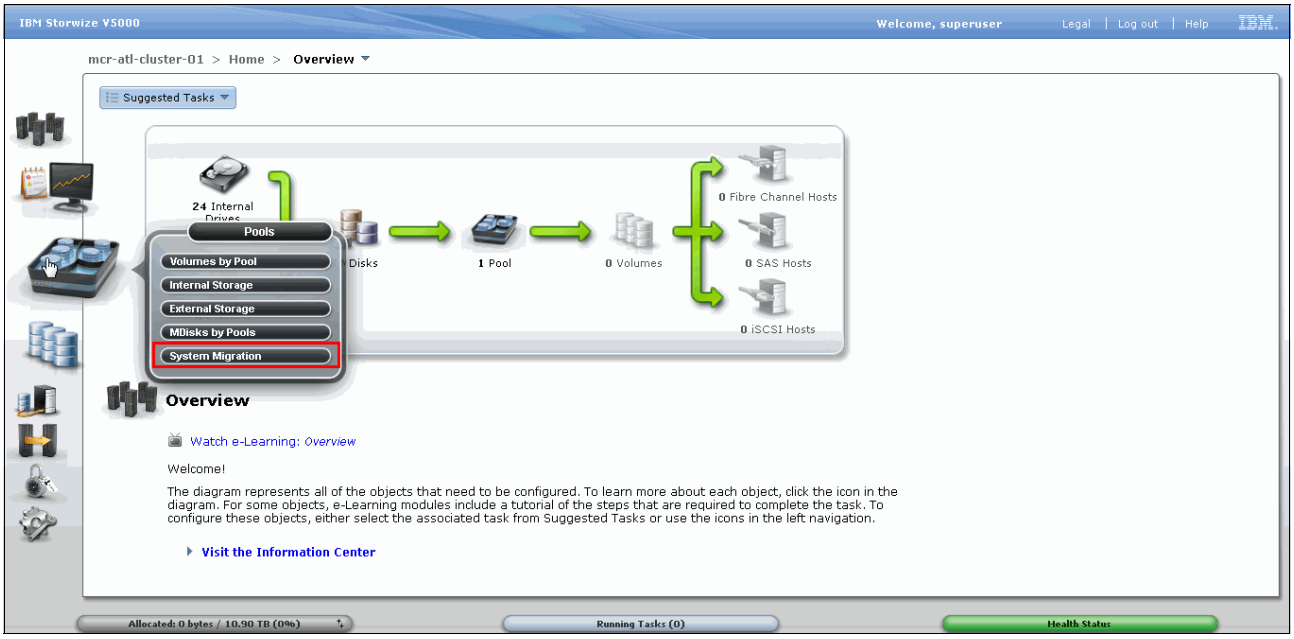

Figure 6-1 Pools menu

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

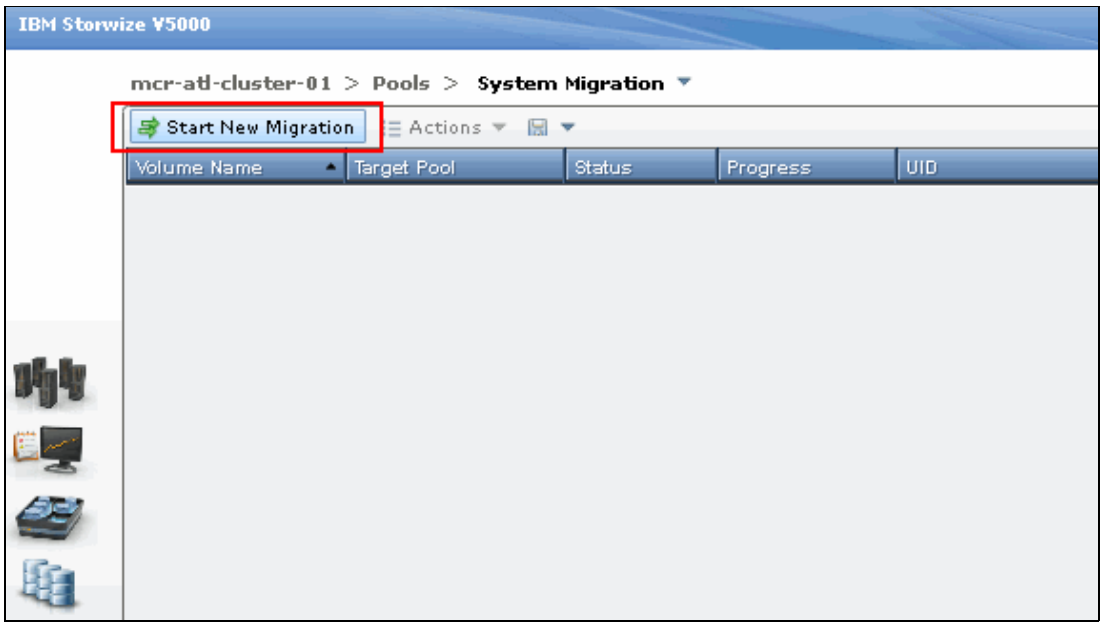

Figure 6-2 System Migration panel

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.

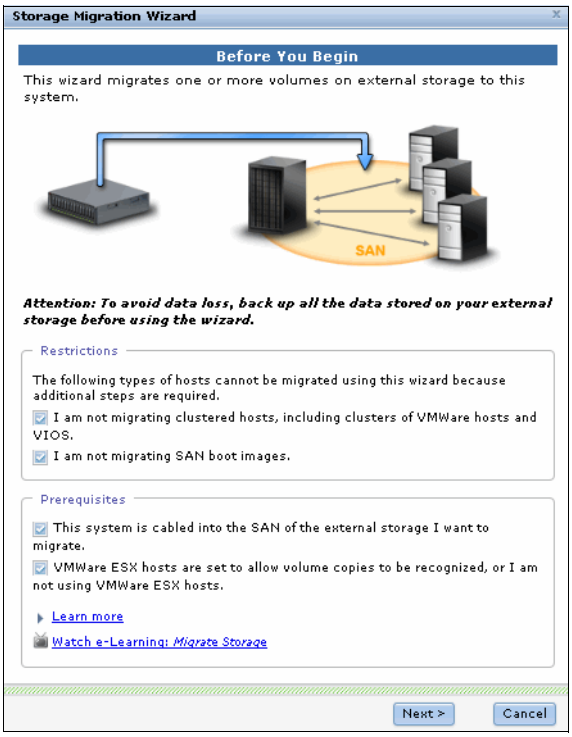

Figure 6-3 Step 1 of the storage migration wizard

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.

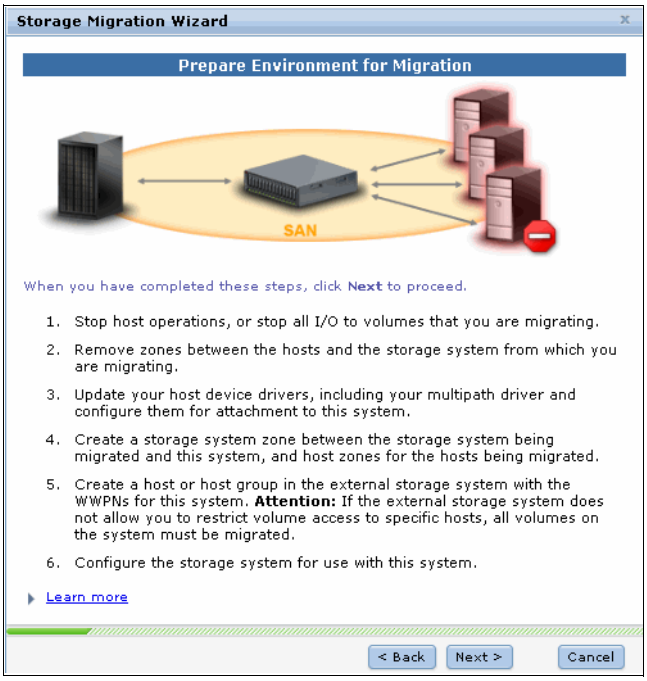

Figure 6-4 Prepare environment

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.

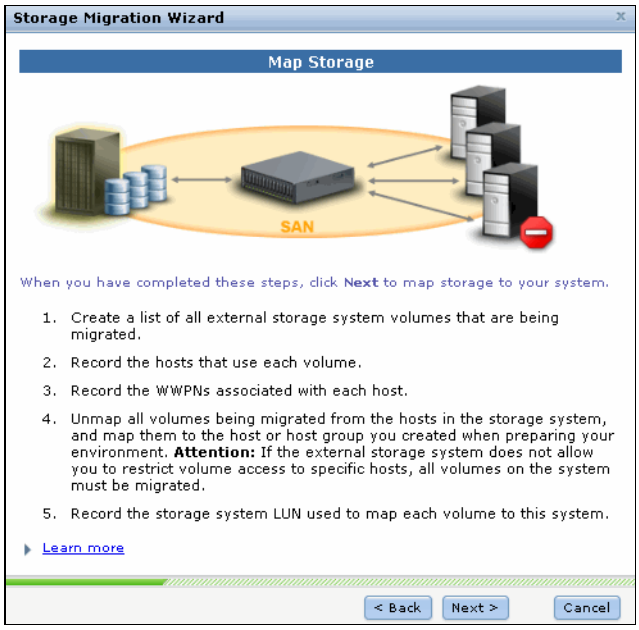

Figure 6-5 Map Storage panel

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.

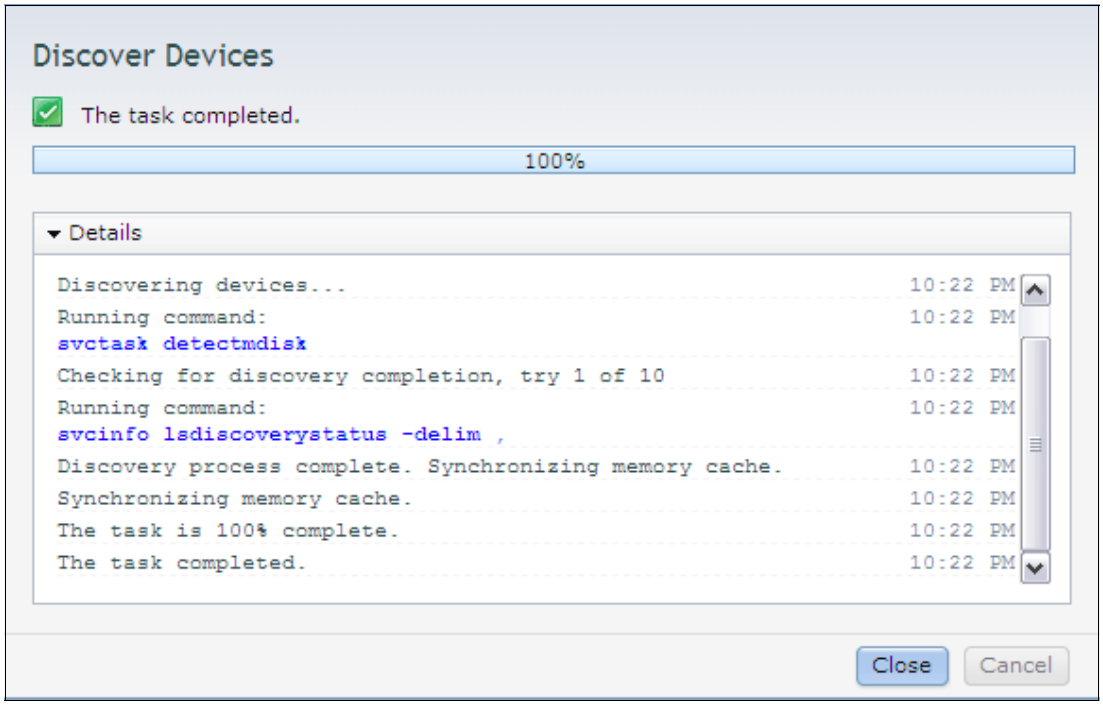

Figure 6-6 Discover Devices task

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.

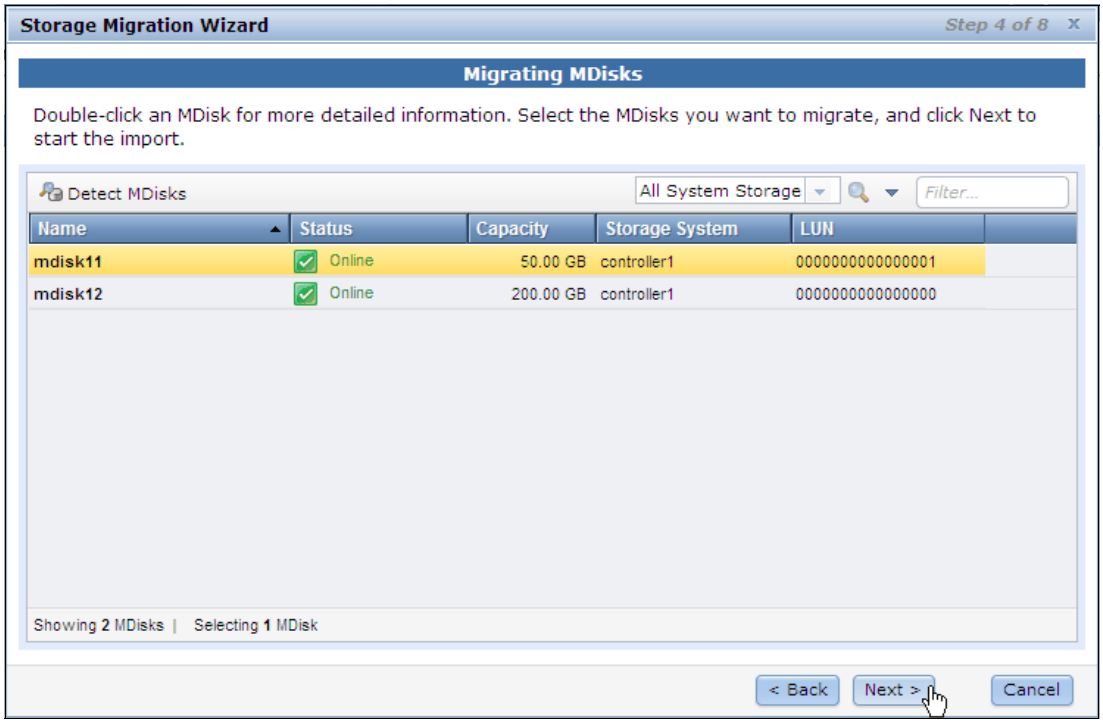

Figure 6-7 Migrating MDisks panel

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.

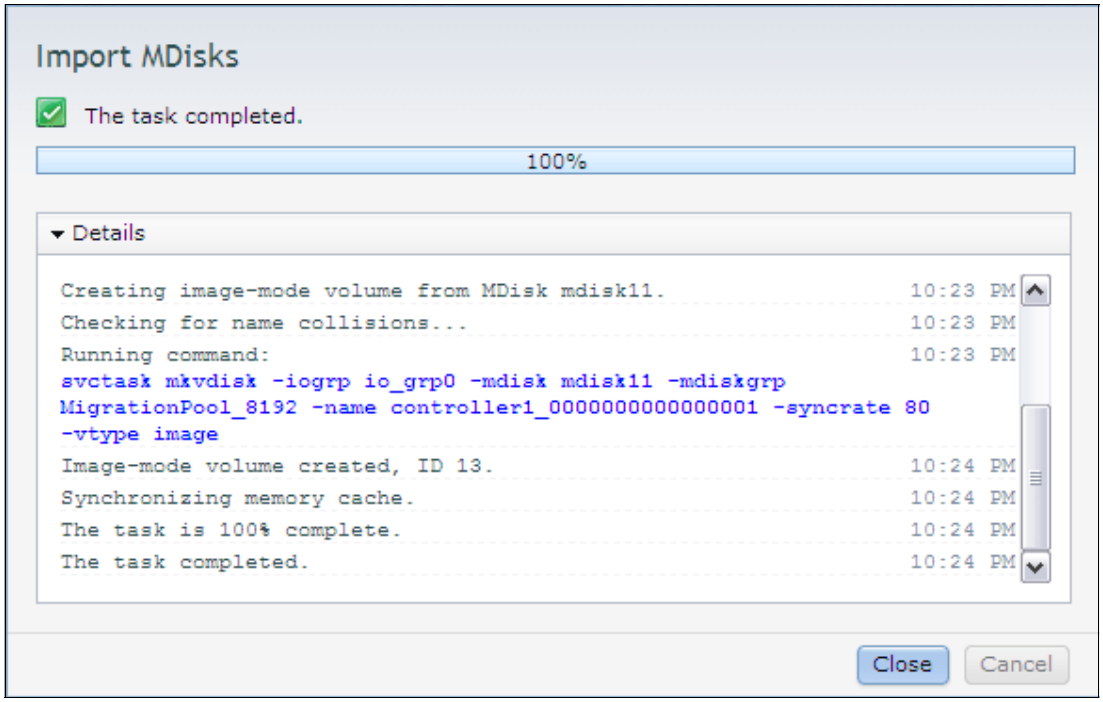

Figure 6-8 Import MDisks task

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.

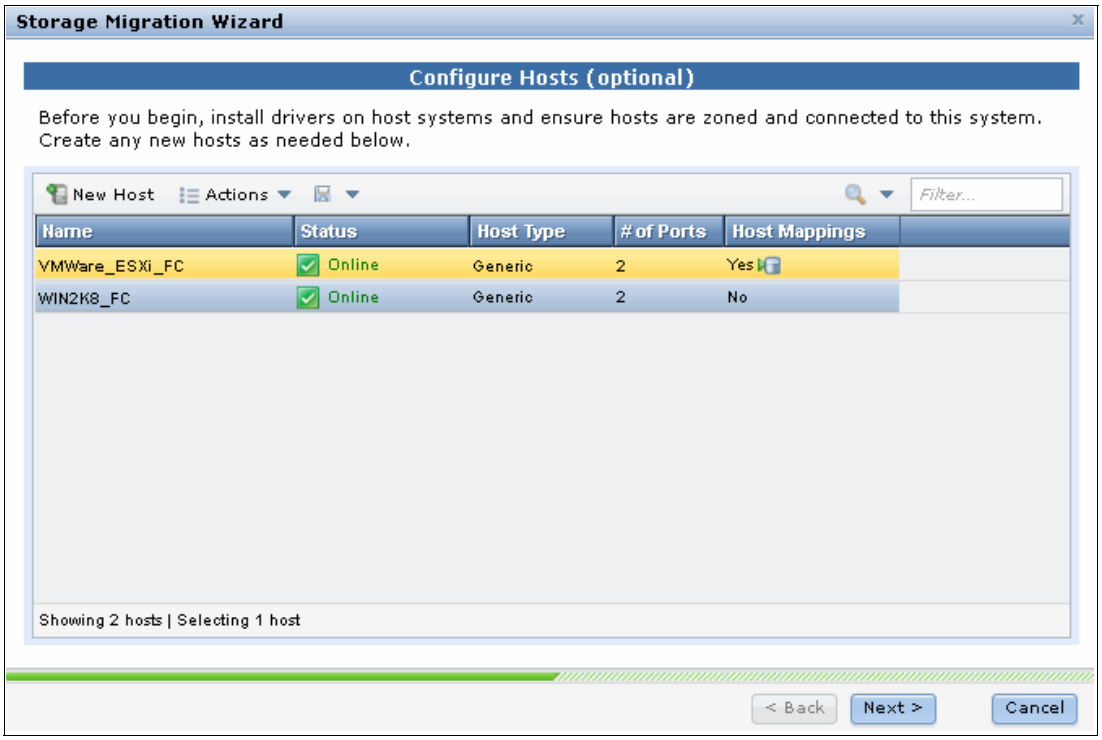

Figure 6-9 Configure Hosts panel

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.

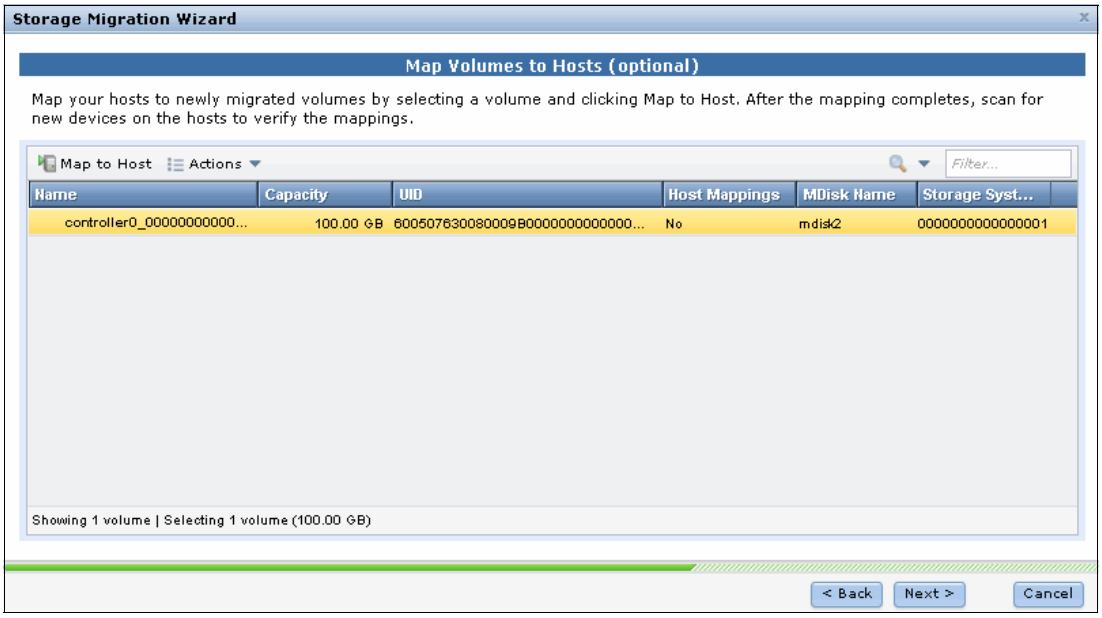

Figure 6-10 Map Volumes to Hosts panel

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.

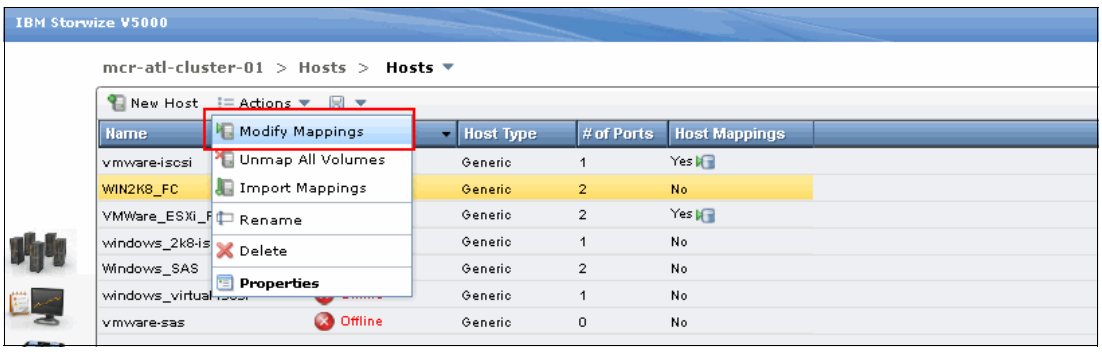

Figure 6-11 Modify Mappings

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.

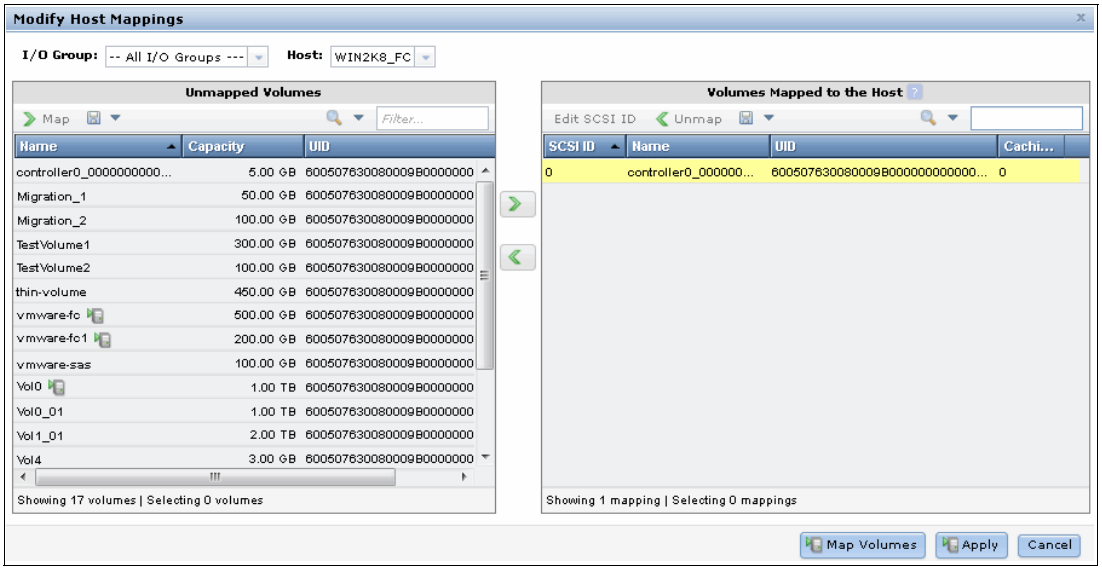

Figure 6-12 Modify Host Mappings panel

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.

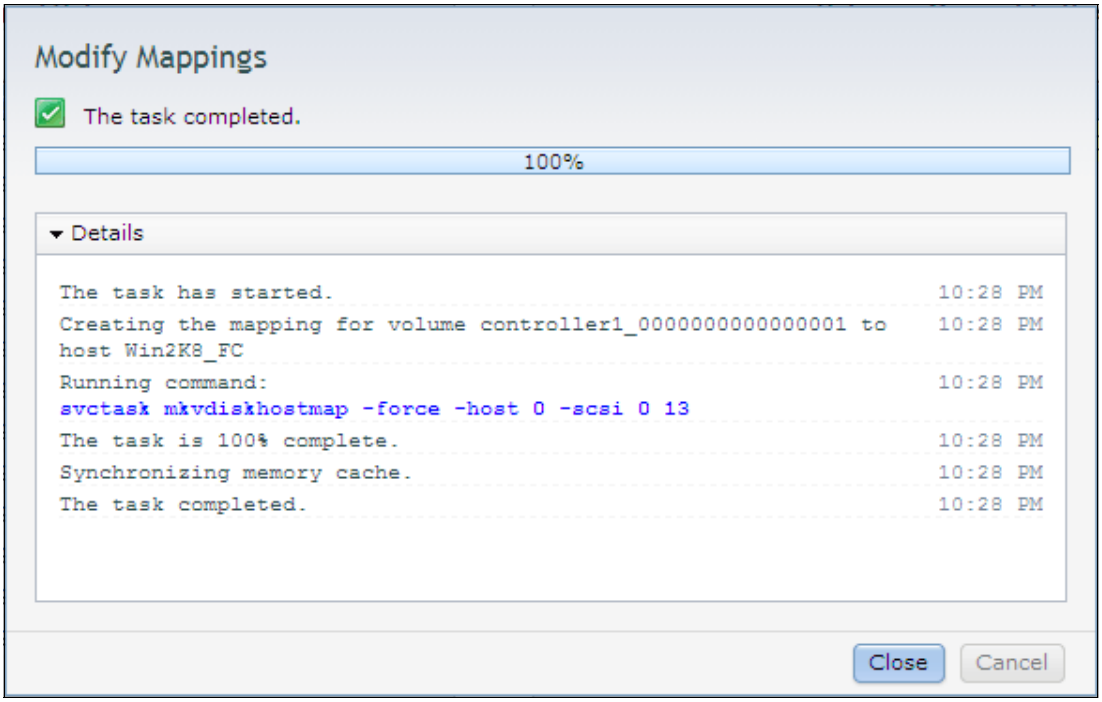

Figure 6-13 Modify Mappings task

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.

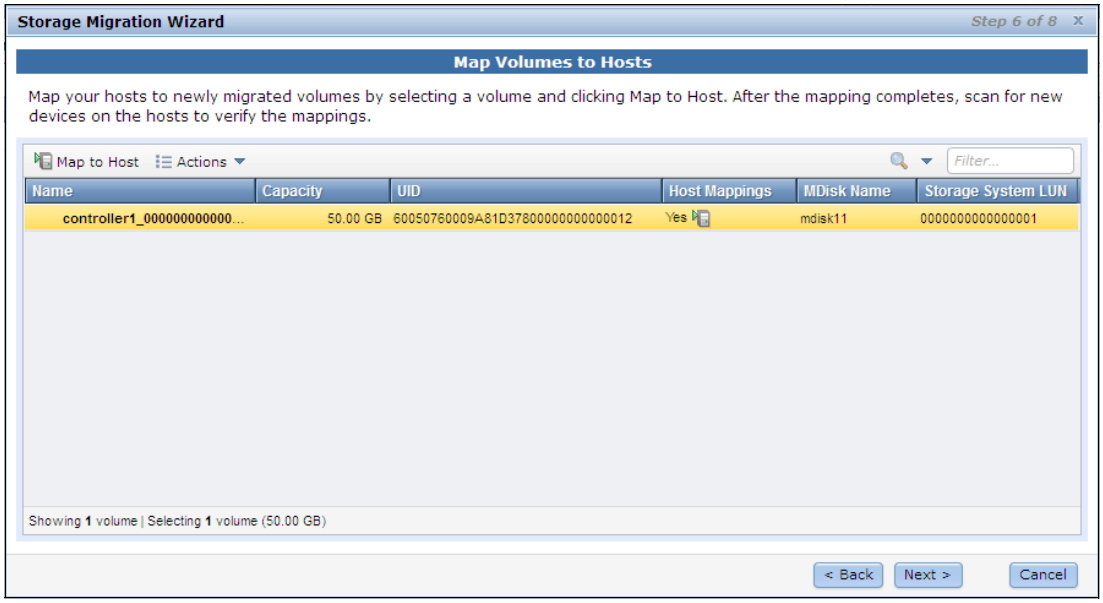

Figure 6-14 Map Volumes to Hosts panel that shows Yes in the Host Mappings column

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.

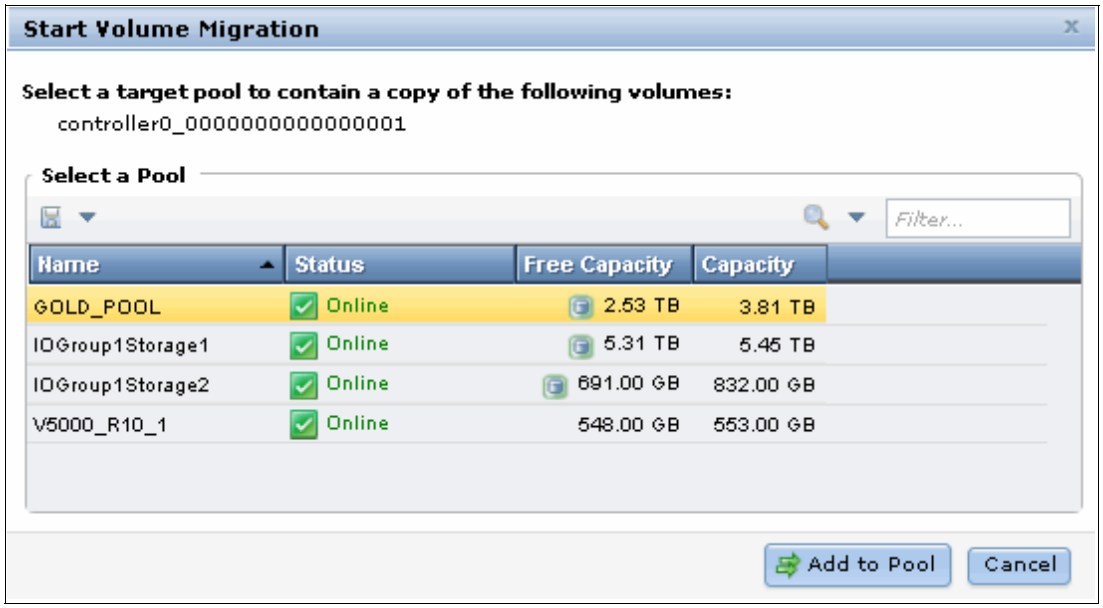

Figure 6-15 Select a Pool panel

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.

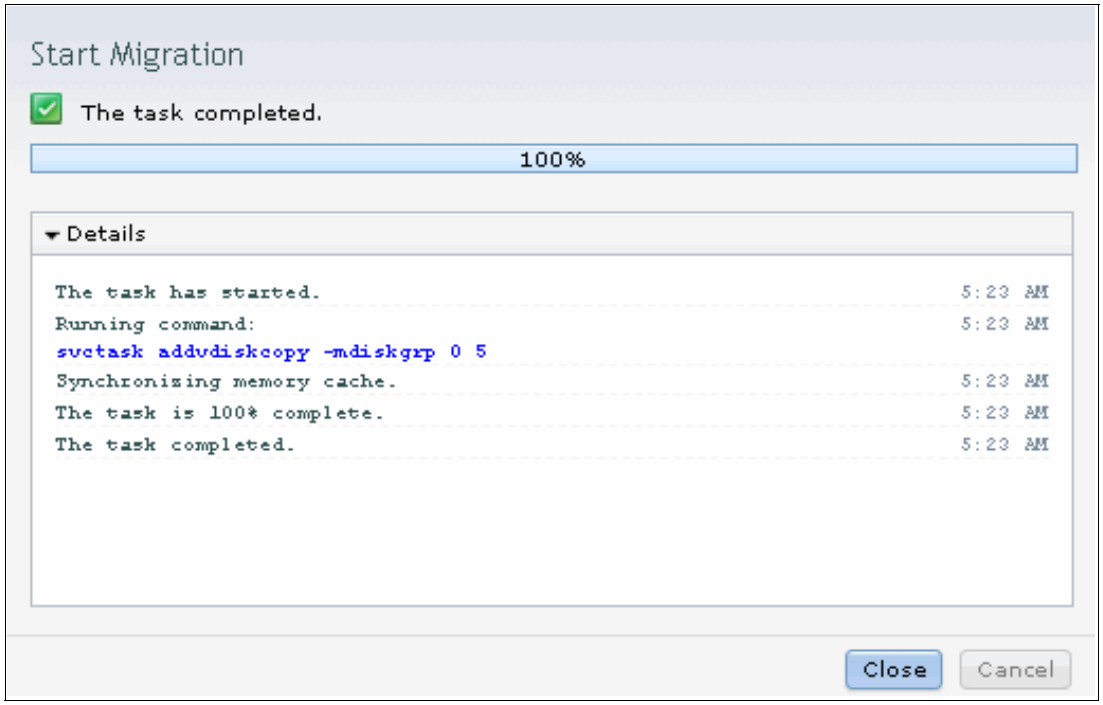

Figure 6-16 Start Migration task

Step 8: Finish the storage migration wizard

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

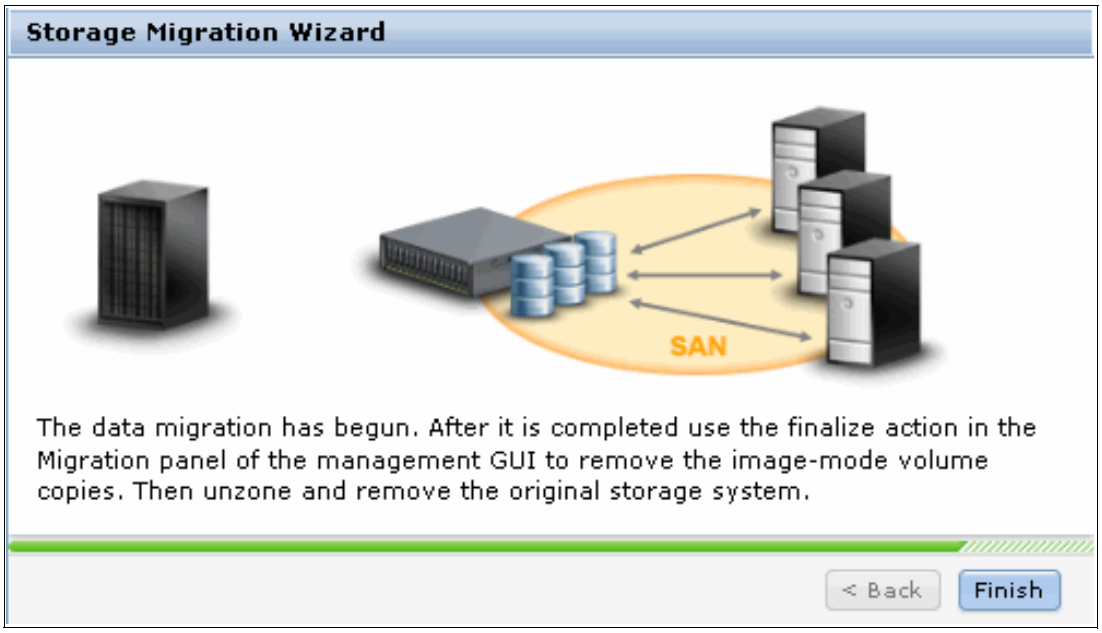

Figure 6-17 Step 8 of the storage migration wizard

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.

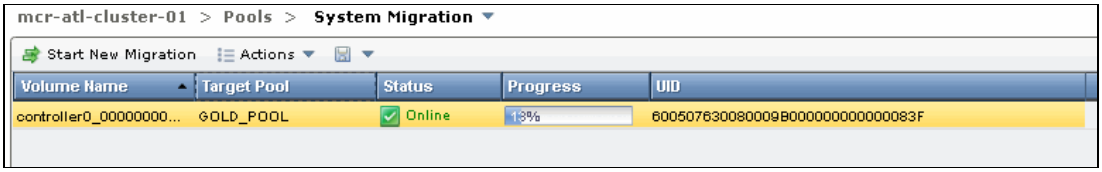

Figure 6-18 Storage Migration panel with a migration in progress

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).

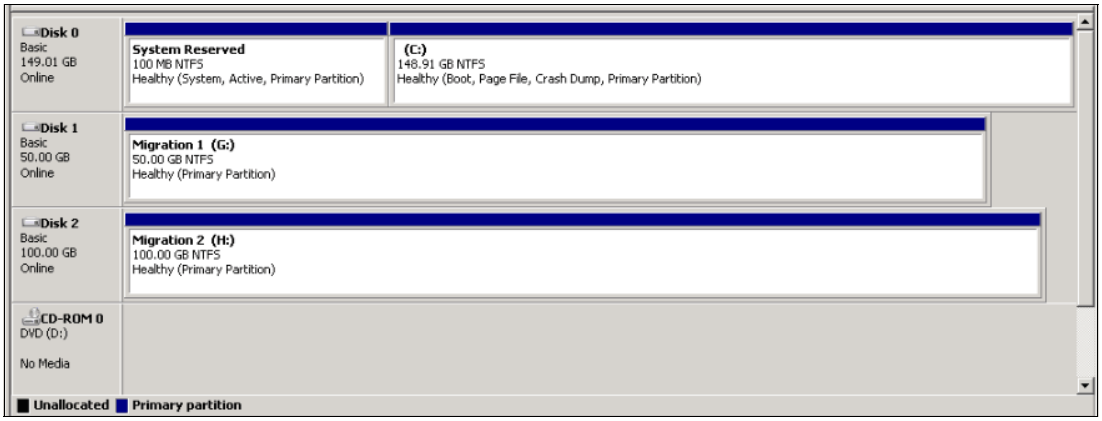

Figure 6-19 Windows 2008 disk management panel

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.

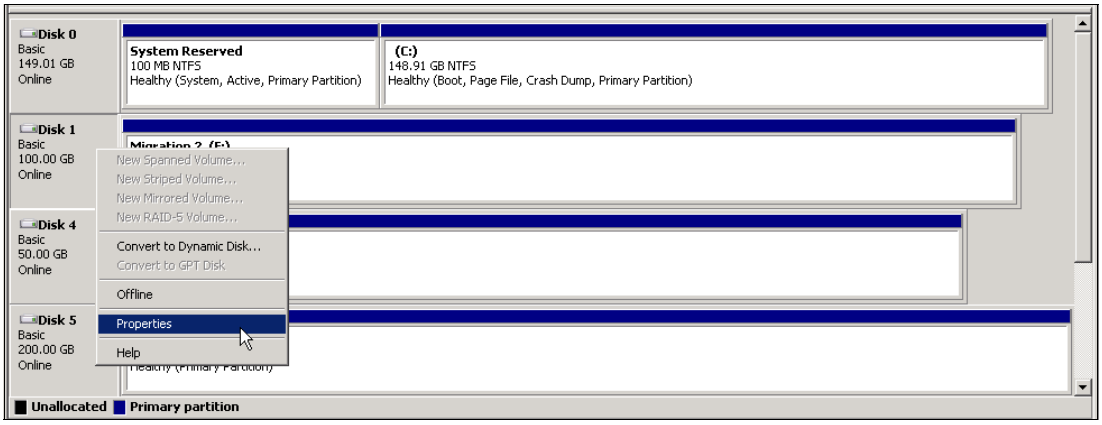

Figure 6-20 Display properties of disk before migration

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.

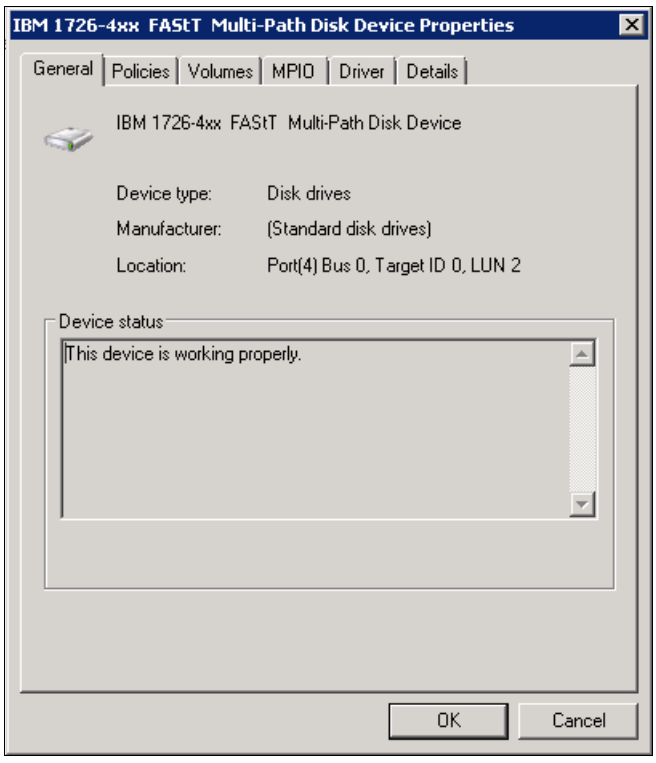

Figure 6-21 Windows 2008 Disk Properties: General tab

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.

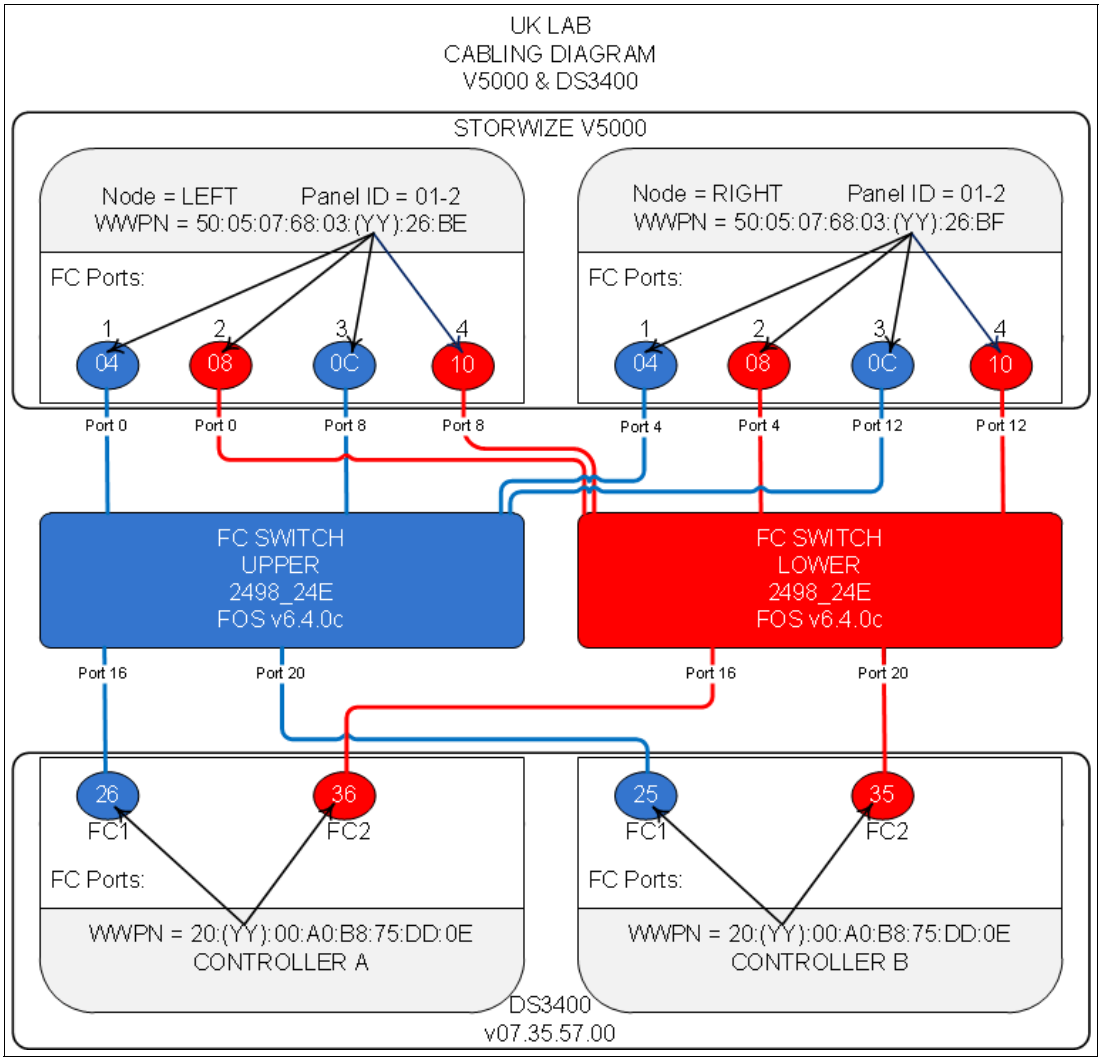

Figure 6-22 Example scenario Fibre Channel cabling layout

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.

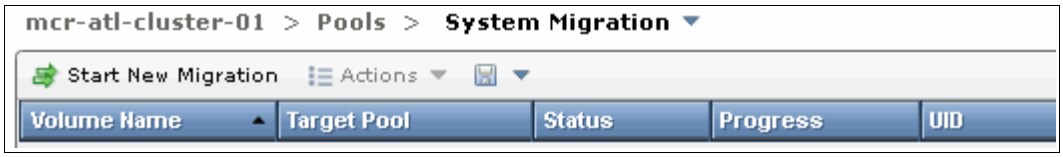

Figure 6-23 Start new migration

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.

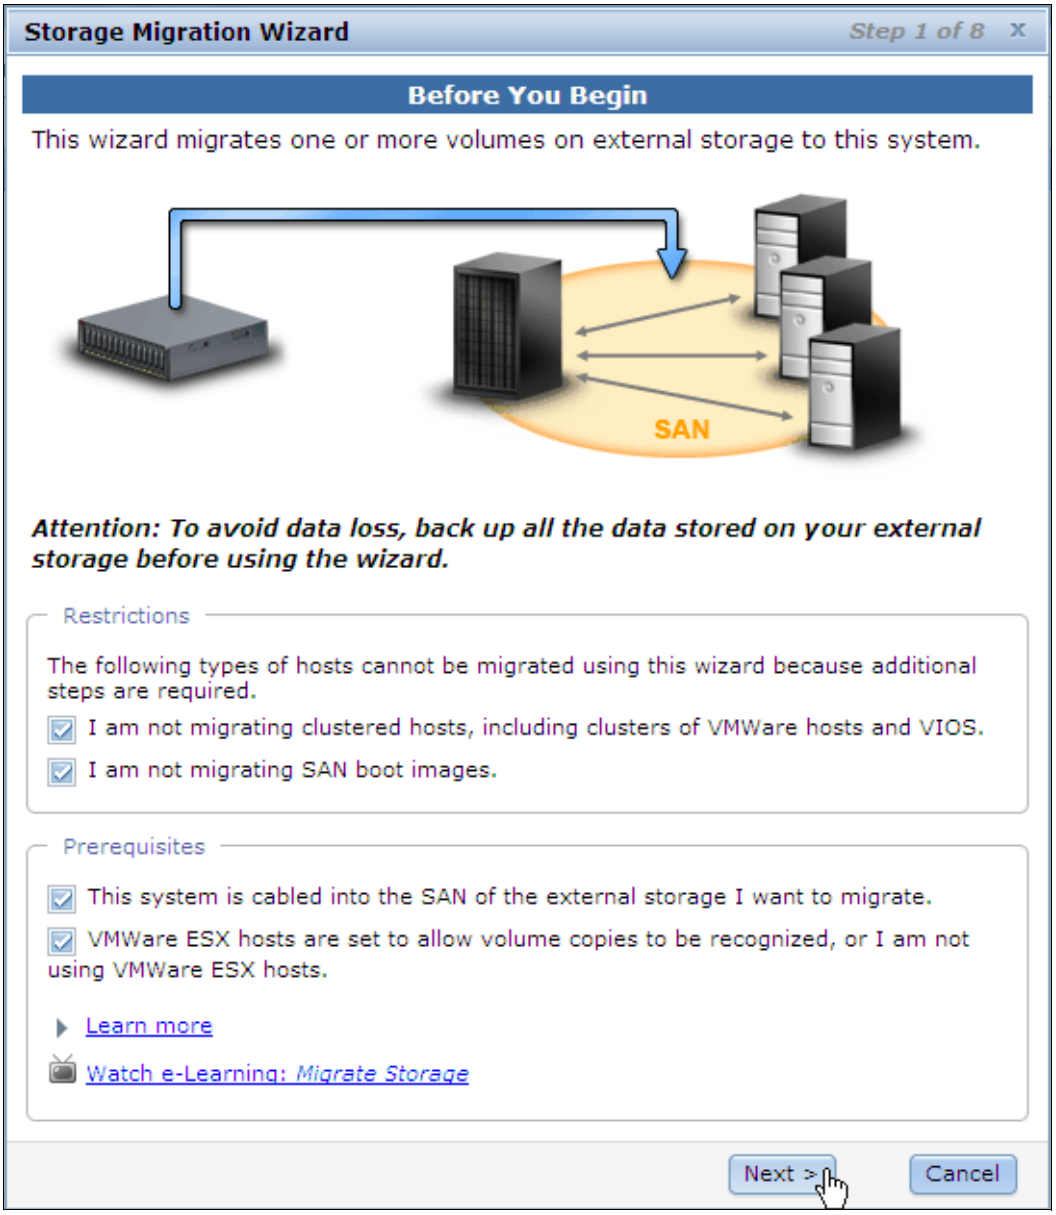

Figure 6-24 Storage Migration wizard: Step 1

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

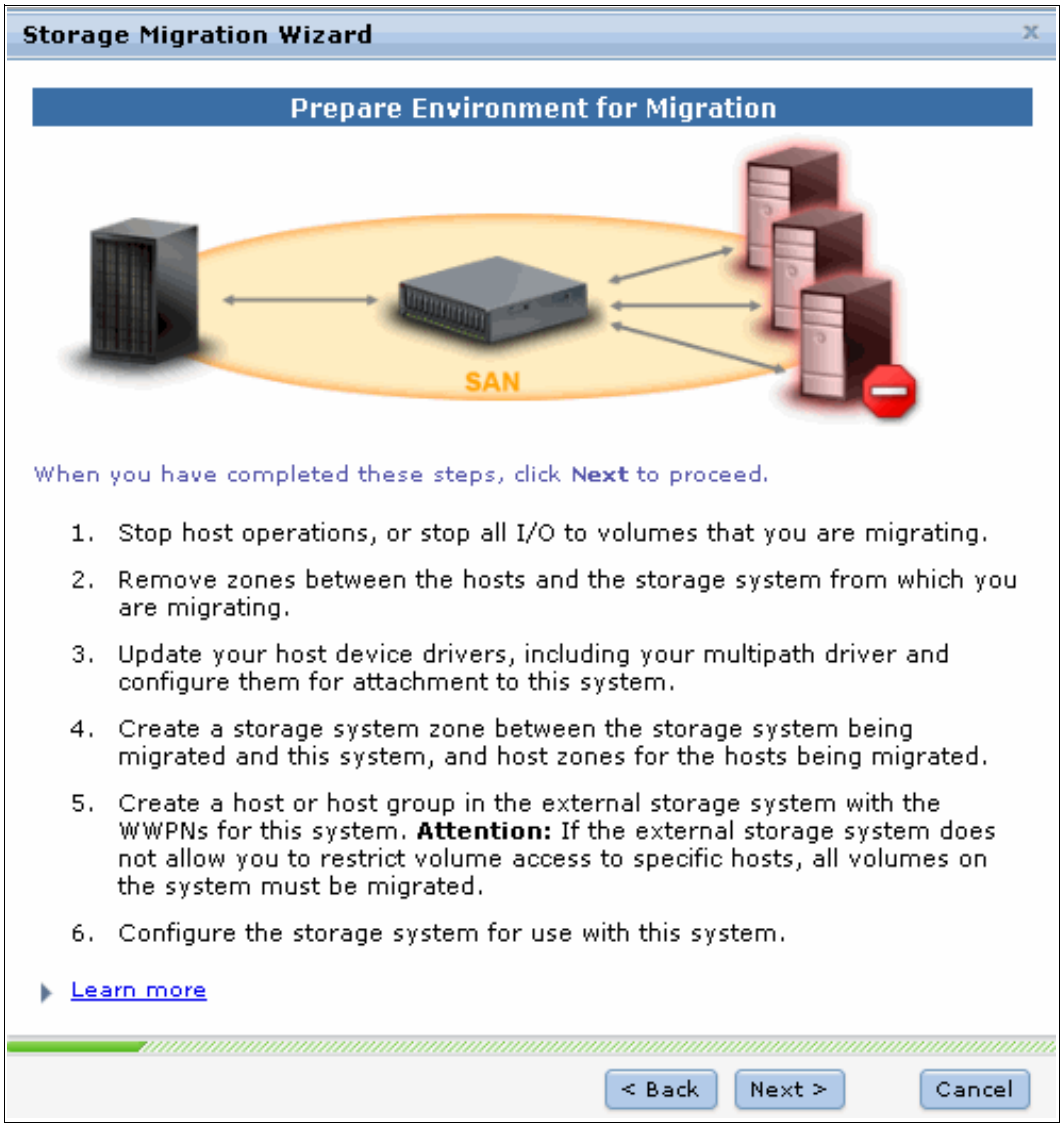

Figure 6-25 Storage Migration wizard: Step 2

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.

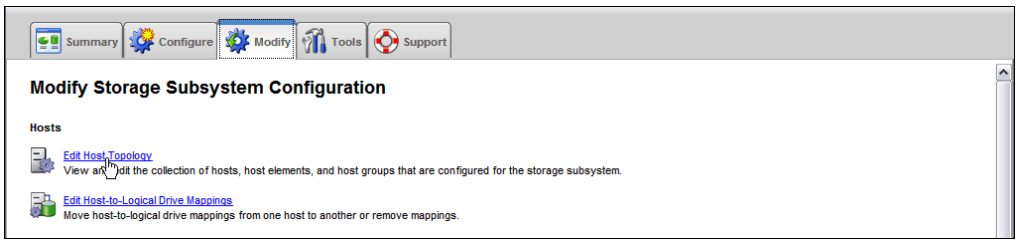

Figure 6-26 IBM DS3400 modify tab: Edit Host Topology

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

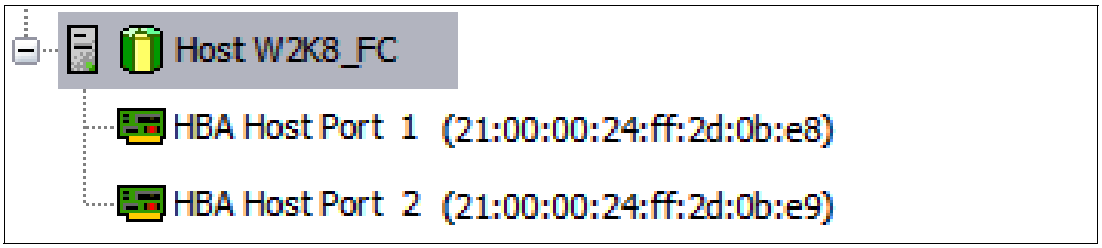

Figure 6-27 IBM DS3400 host definition

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.

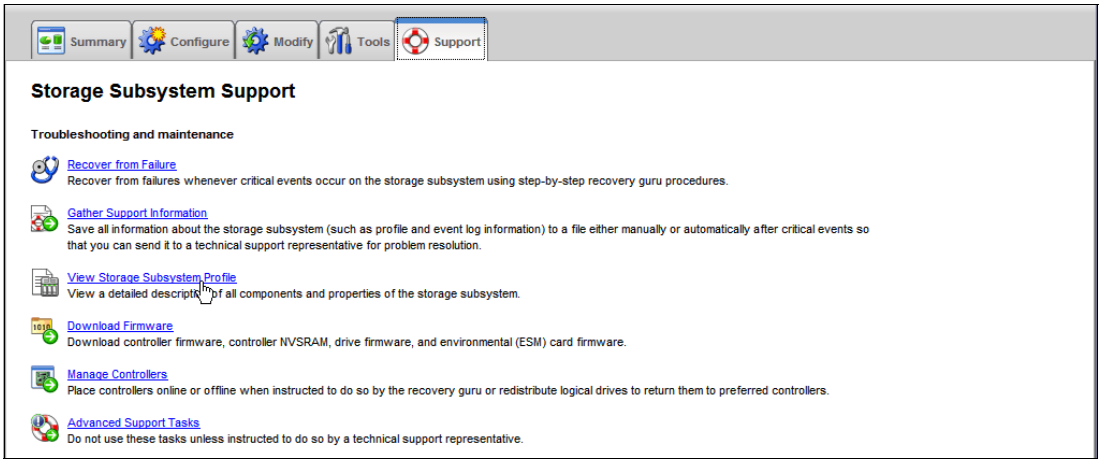

Figure 6-28 Storage Subsystem Support profile

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

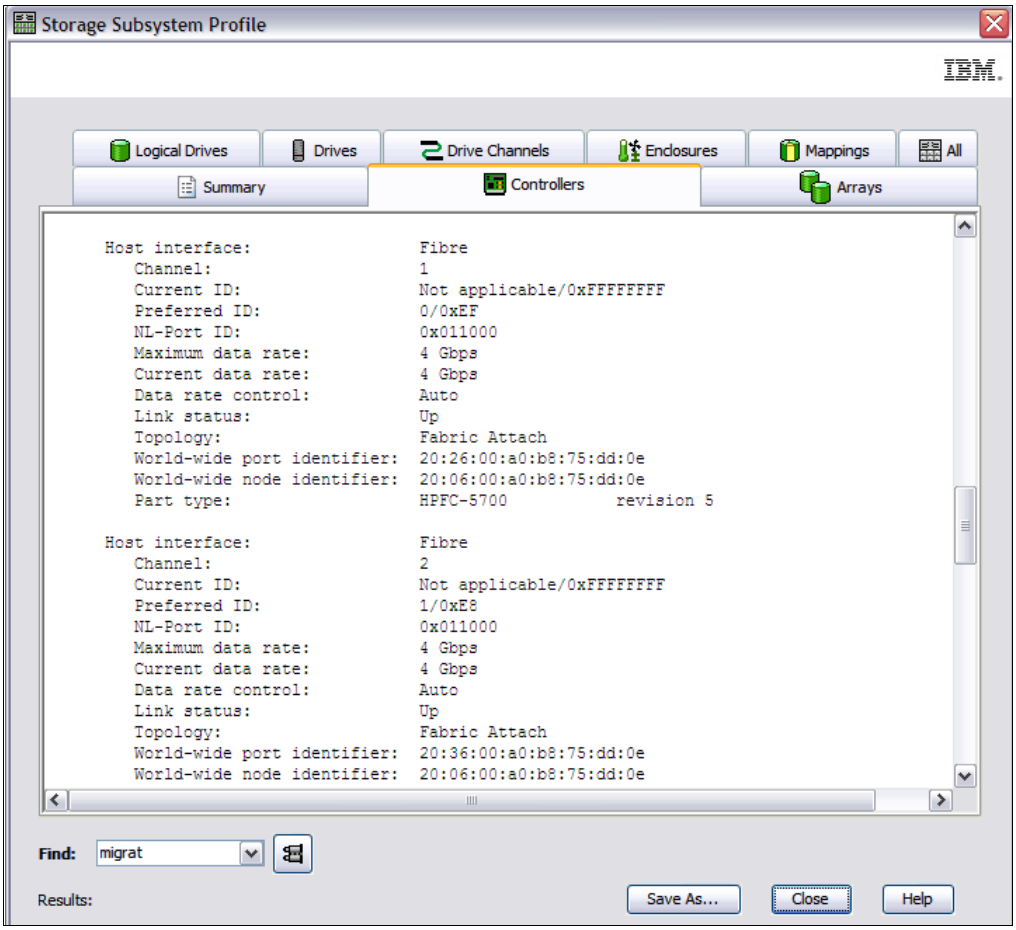

Figure 6-29 Storage Subsystem Profile: Controller WWPNs

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**.

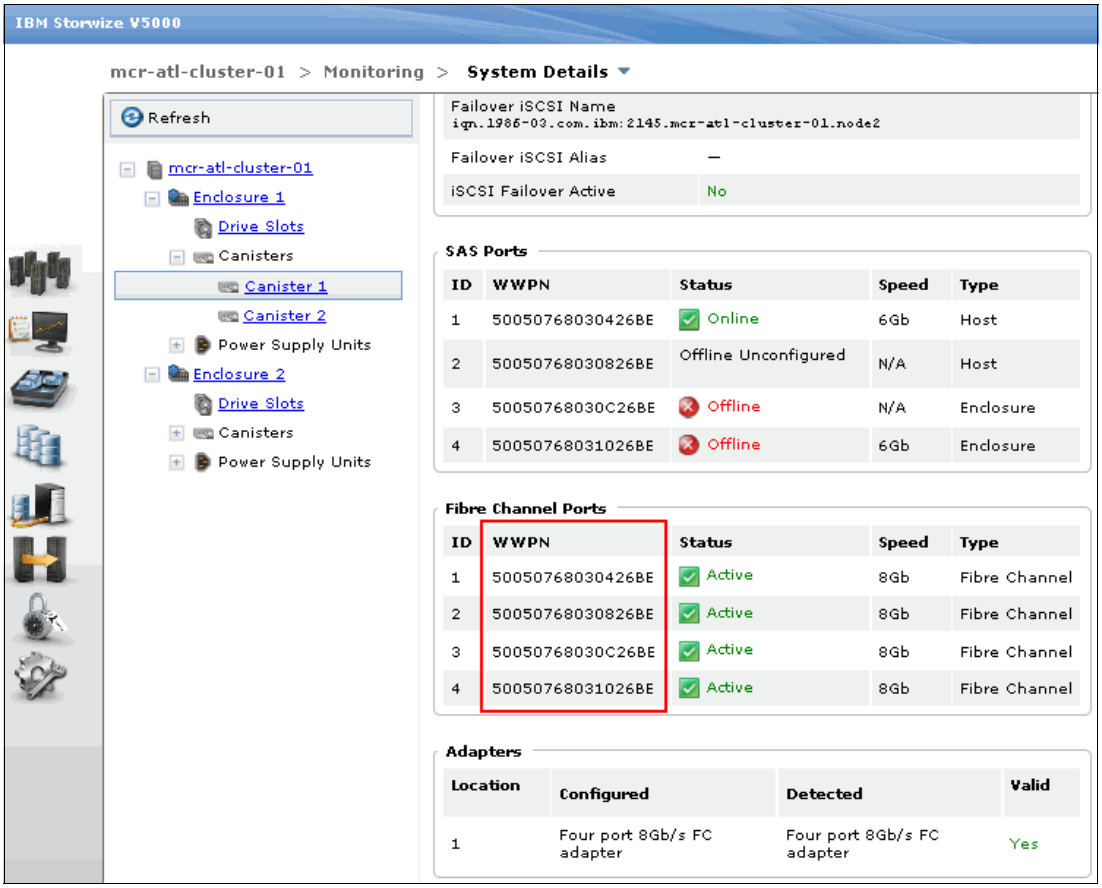

Figure 6-30 Storwize V5000 node canister WWPNs information

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.

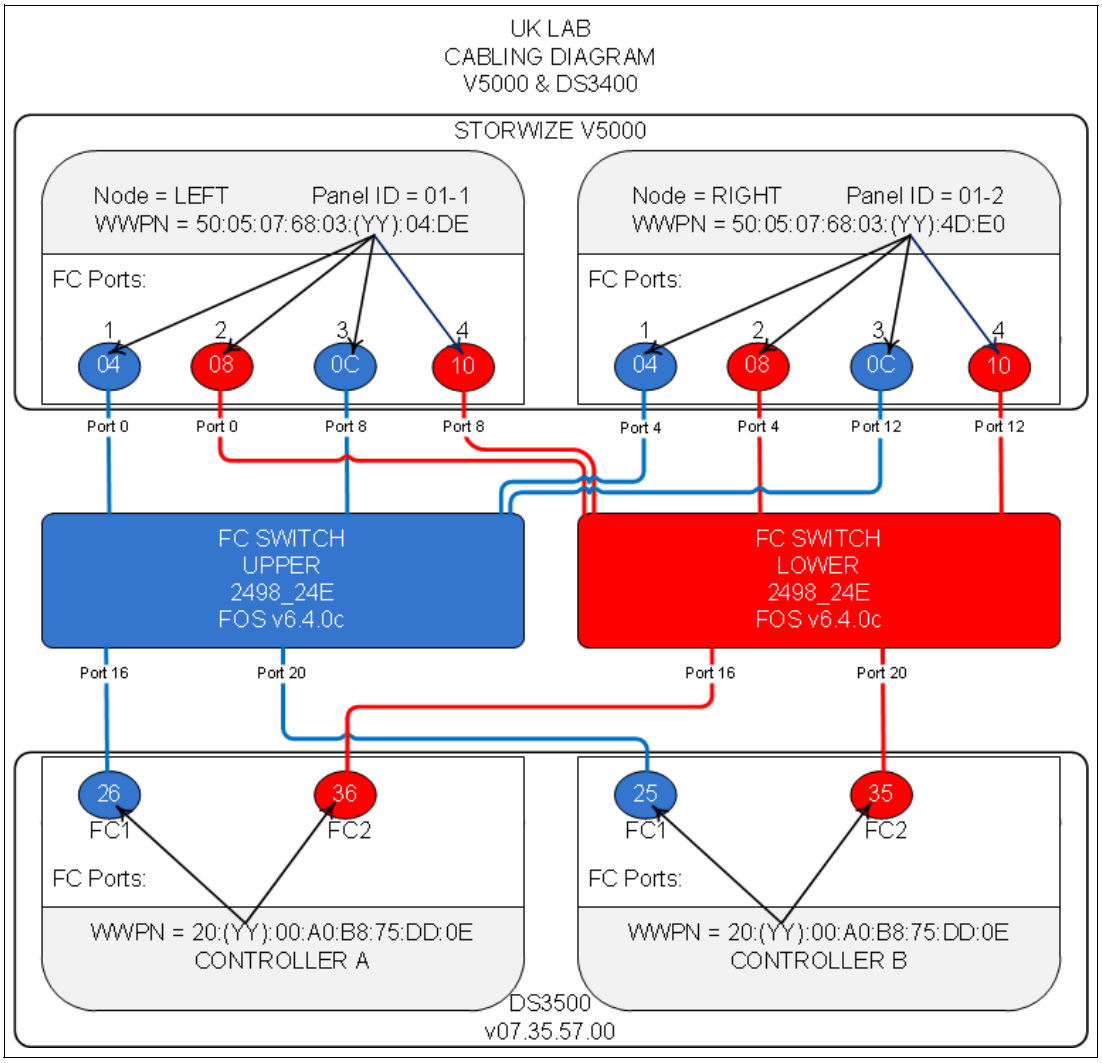

Figure 6-31 Example scenario Storwize V5000 and IBM DS3400 WWPN location diagram

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.

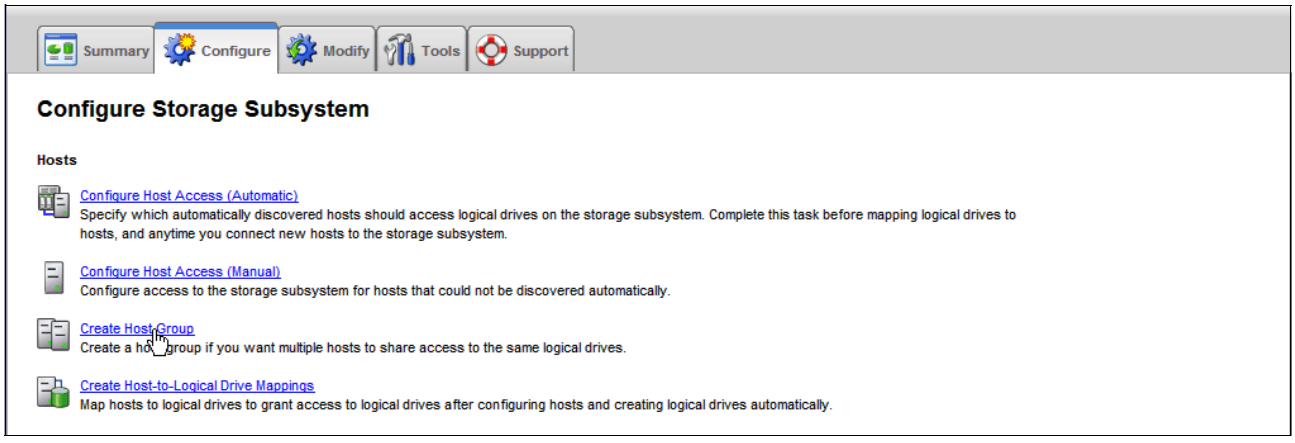

Figure 6-32 Configure Storage Subsystem

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.

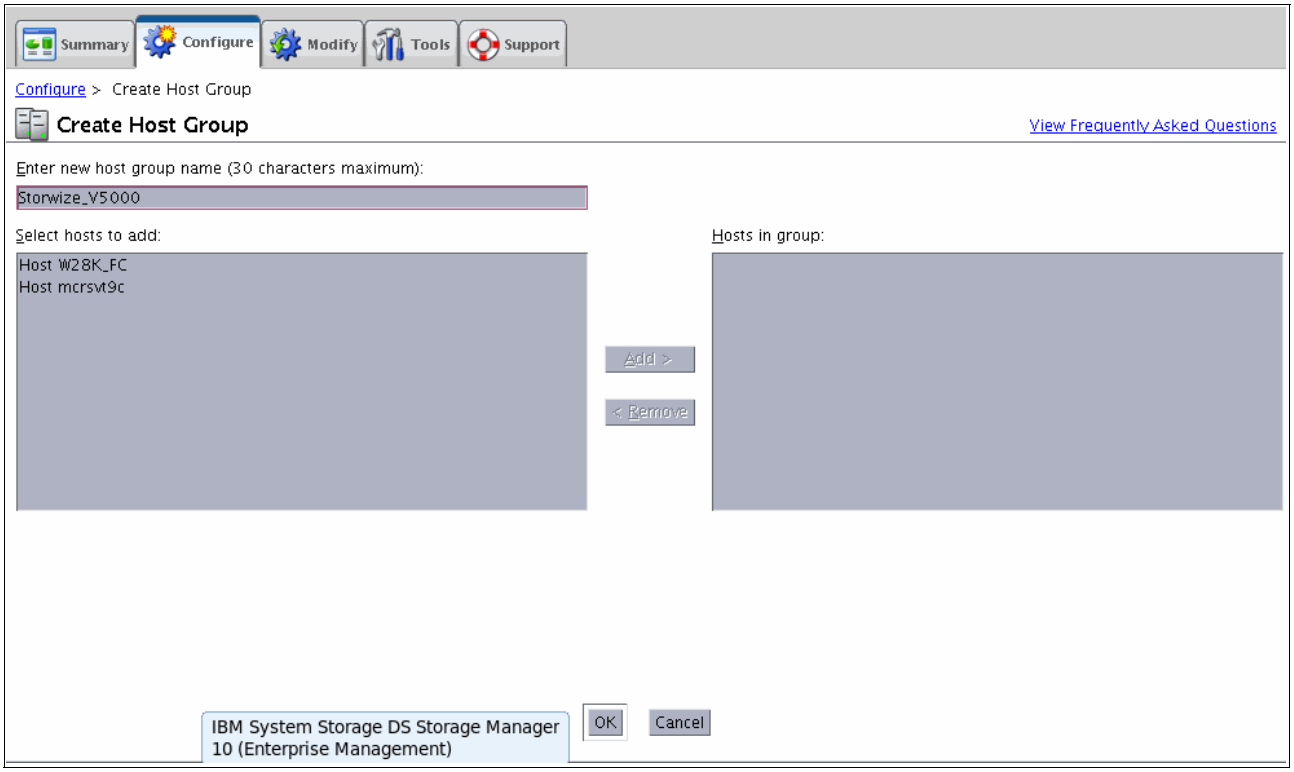

Figure 6-33 IBM DS3400 Create Host Group panel

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.

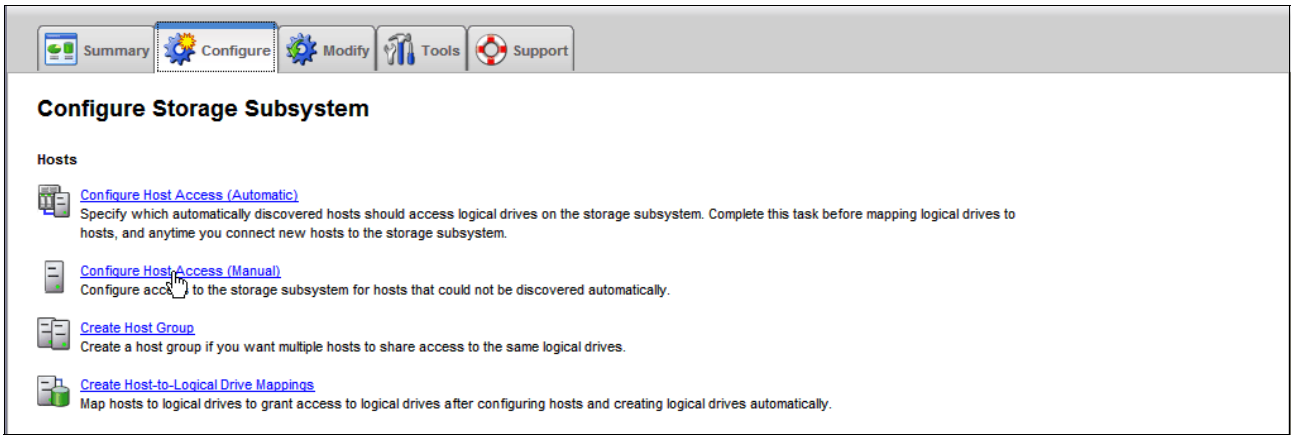

Figure 6-34 Selecting Configure Host-Access (Manual) option

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.

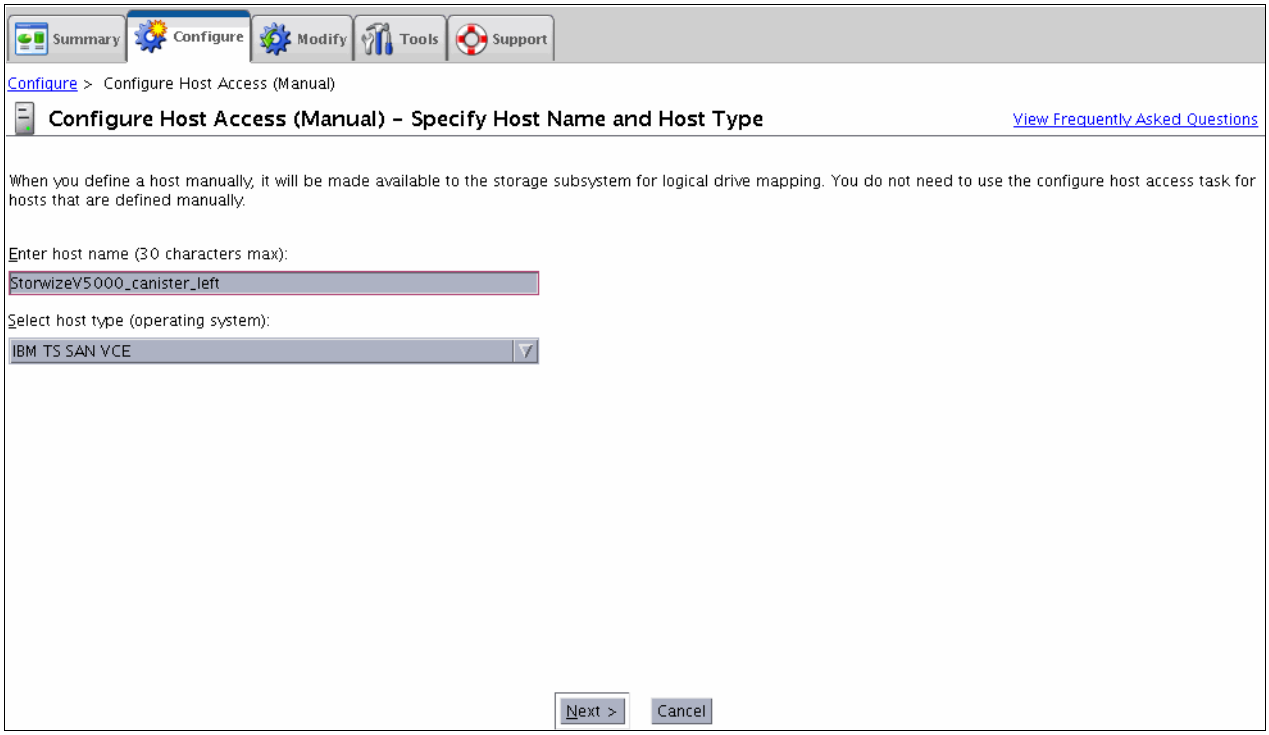

Figure 6-35 IBM DS3400 storage manager Configure tab: Configure host

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.

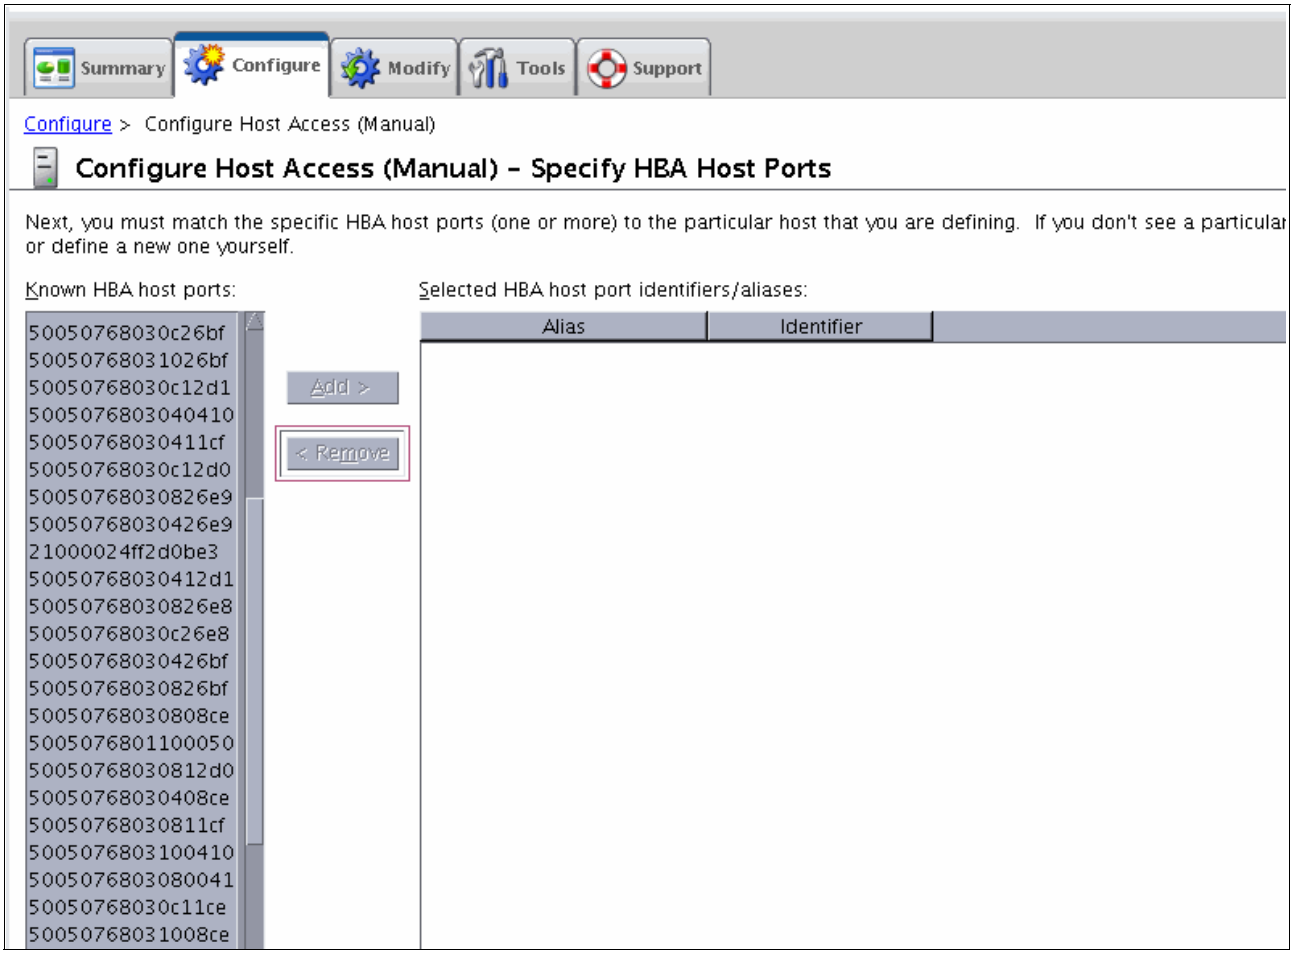

Figure 6-36 IBM DS3400 Specify HBA Host Ports panel

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

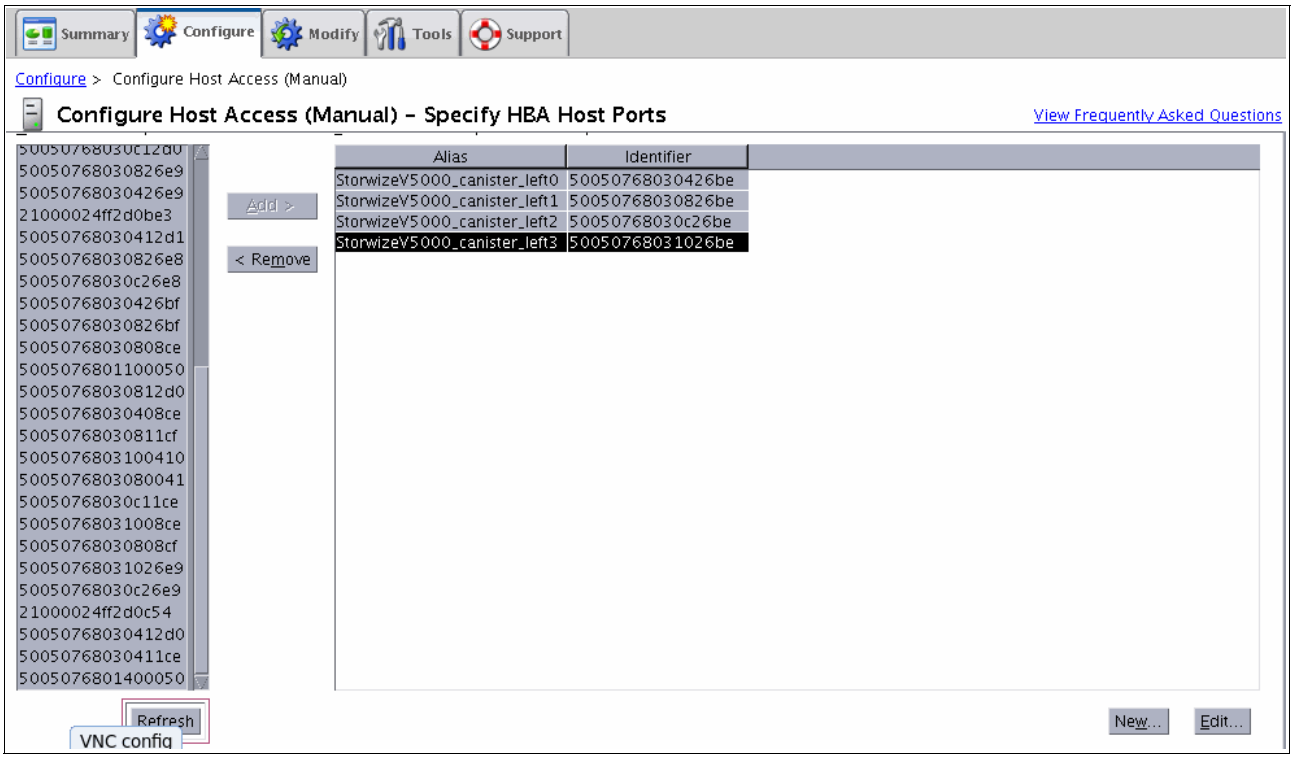

Figure 6-37 IBM DS3400 storage manager specifying HBA host ports: Edit alias

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.

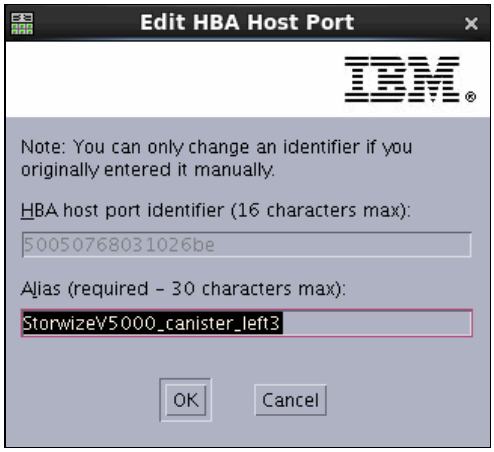

Figure 6-38 IBM DS3400 Edit HBA Host Port panel

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.

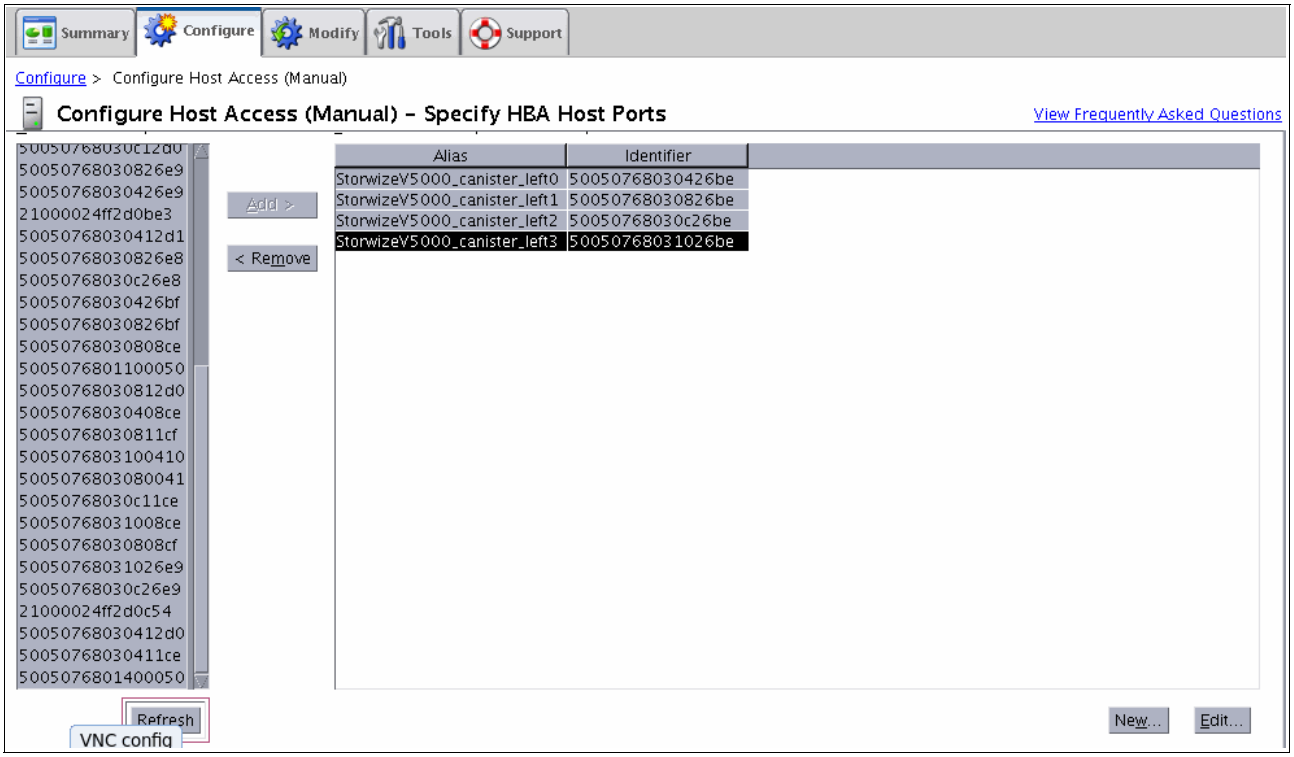

Figure 6-39 IBM DS3400 Specify HBA Host Ports panel

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.

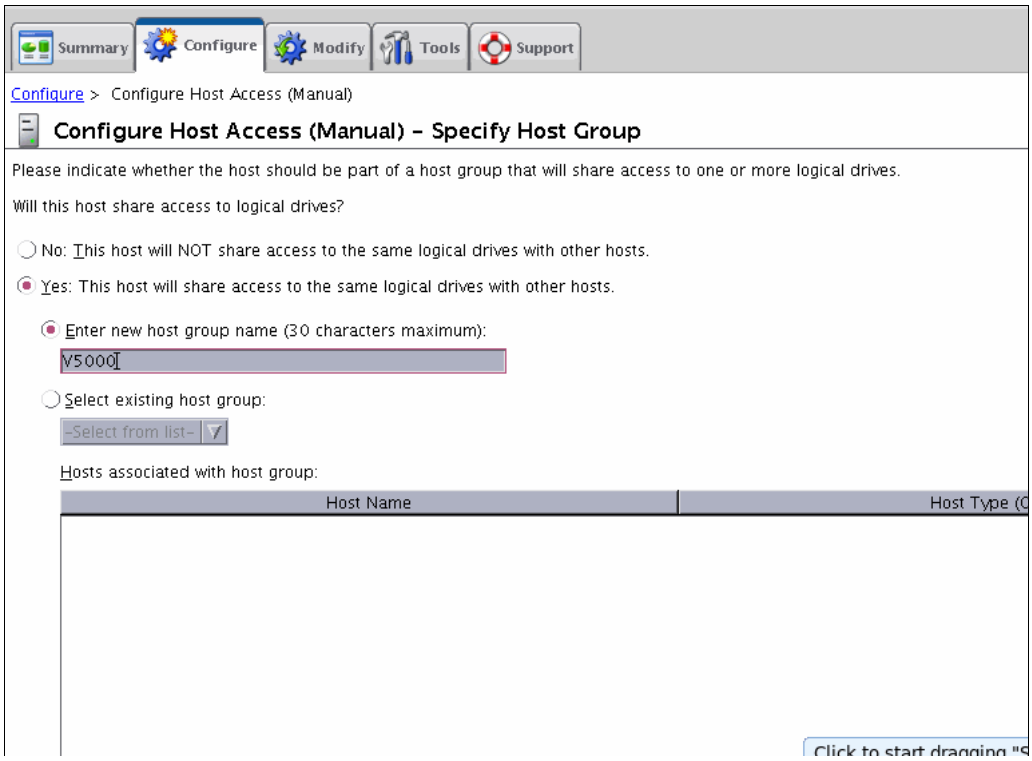

Figure 6-40 IBM DS3400 Specify Host Group panel

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.

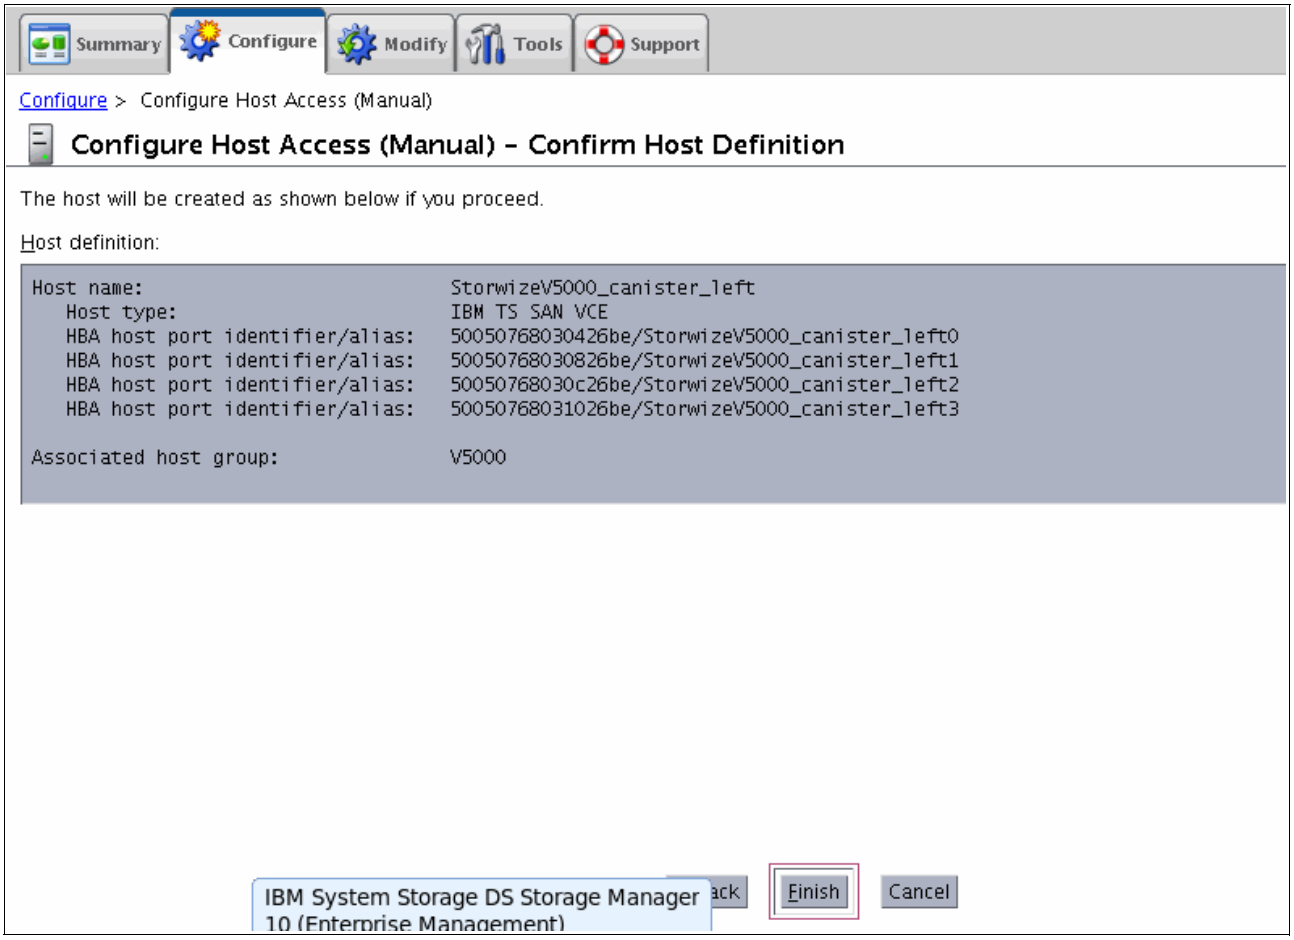

Figure 6-41 IBM DS3400 Confirm Host Definition panel

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.

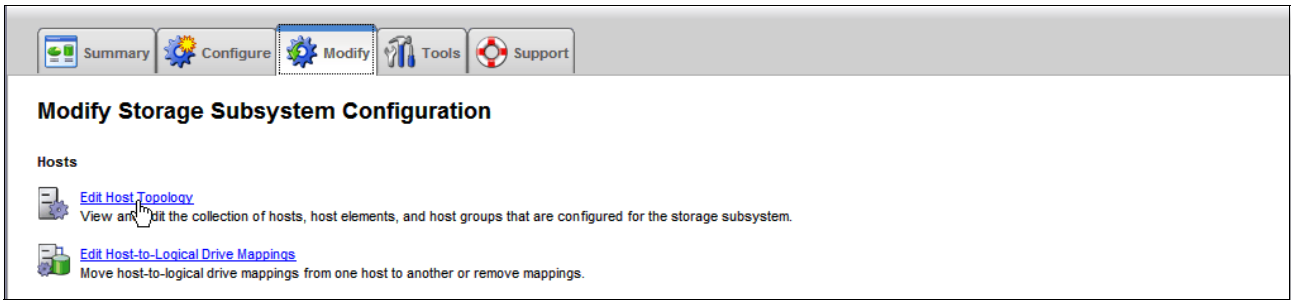

Figure 6-42 Selecting the Edit Host Topology option

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.

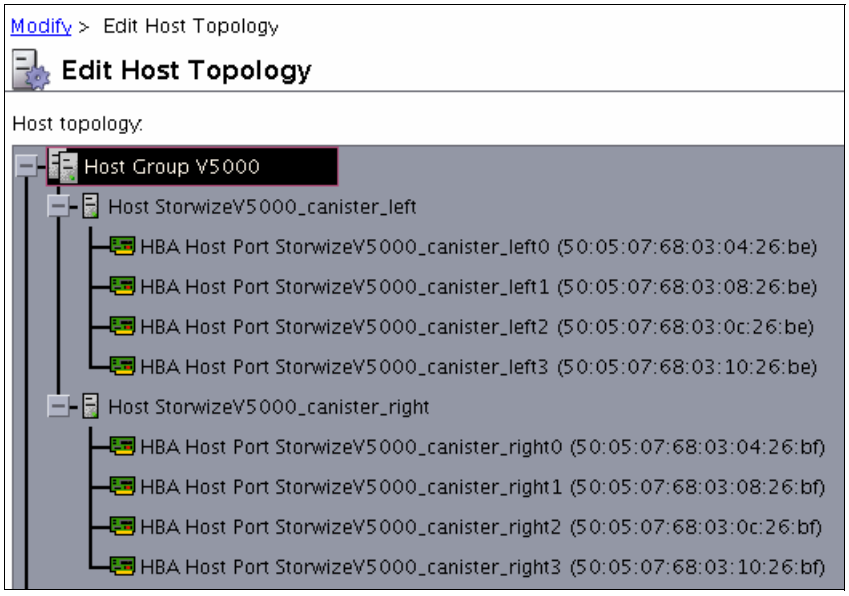

Figure 6-43 IBM DS3400 host group definition for the Storwize V5000

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.

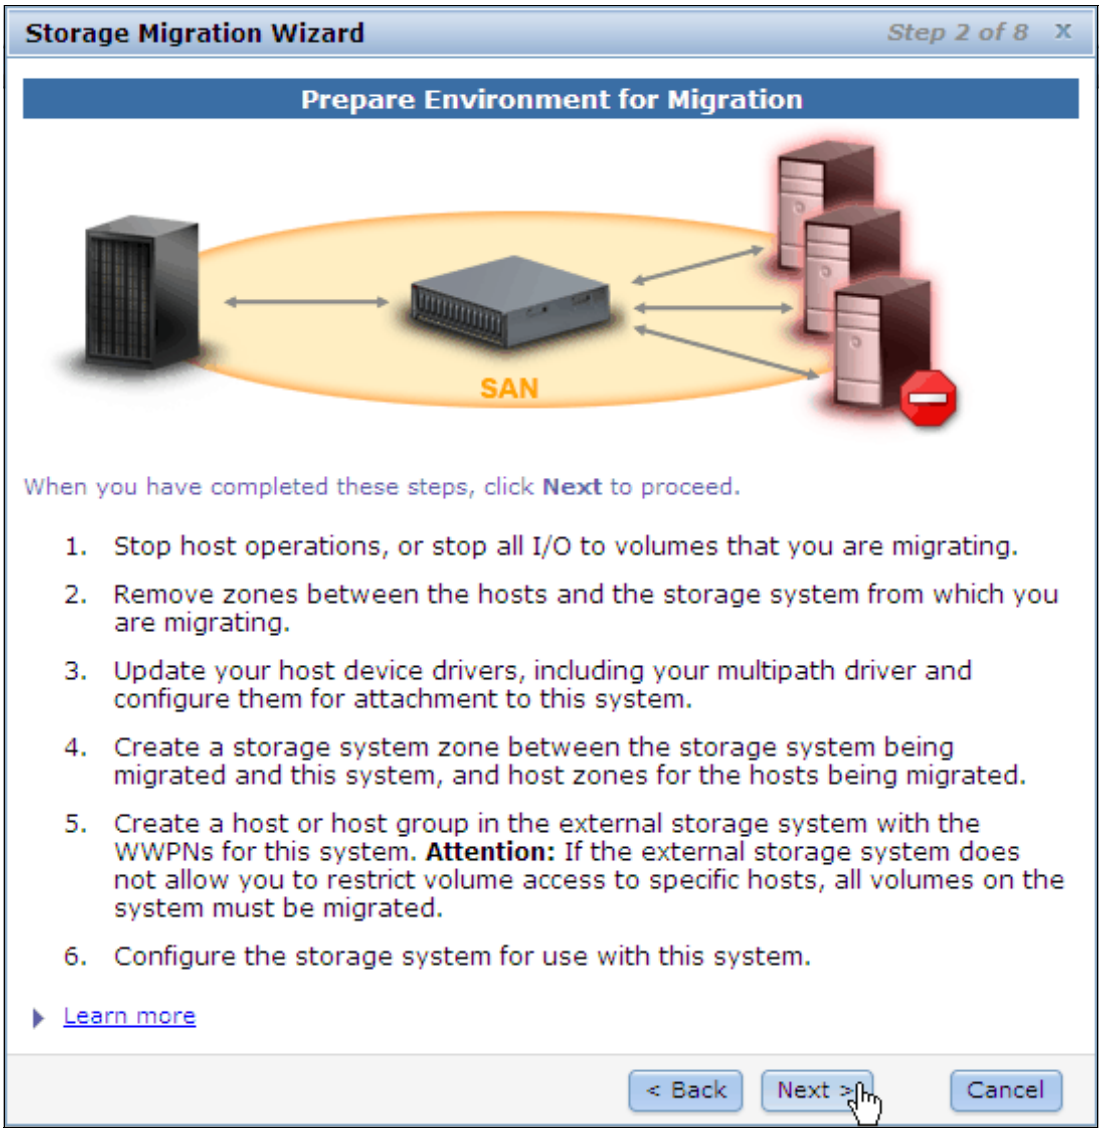

Figure 6-44 Step 2 of the Storage Migration wizard

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

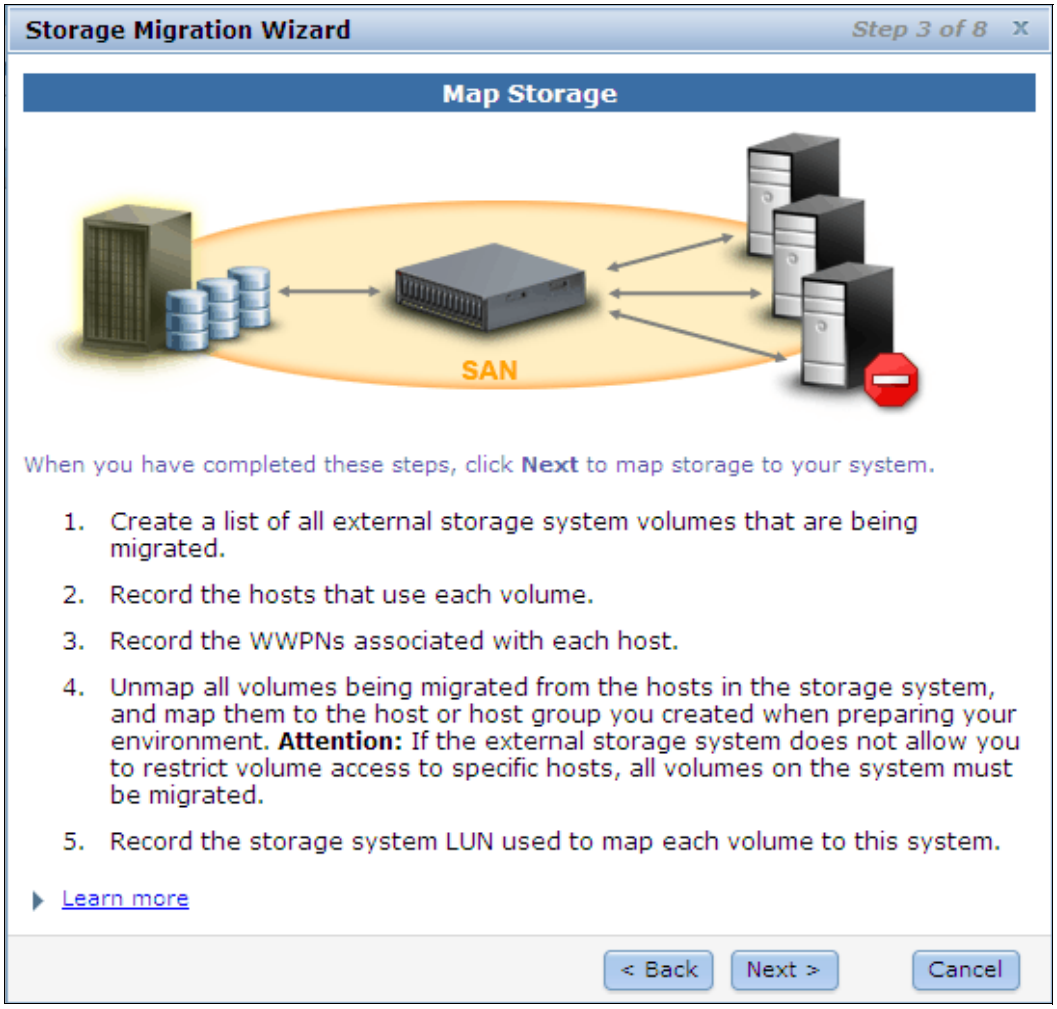

Figure 6-45 Step 3 of the Storage Migration wizard

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.

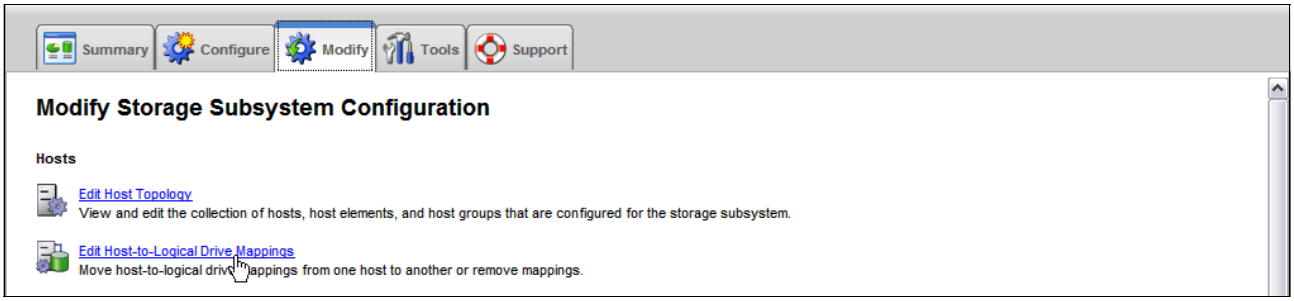

Figure 6-46 IBM DS3400 storage manager Modify tab

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.

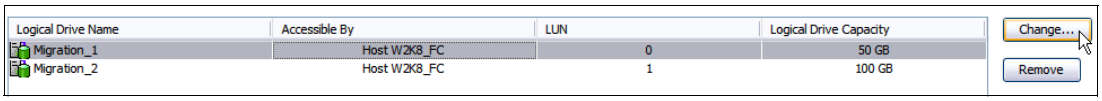

Figure 6-48 IBM DS3400 modify mapping panel: Change mapping

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.

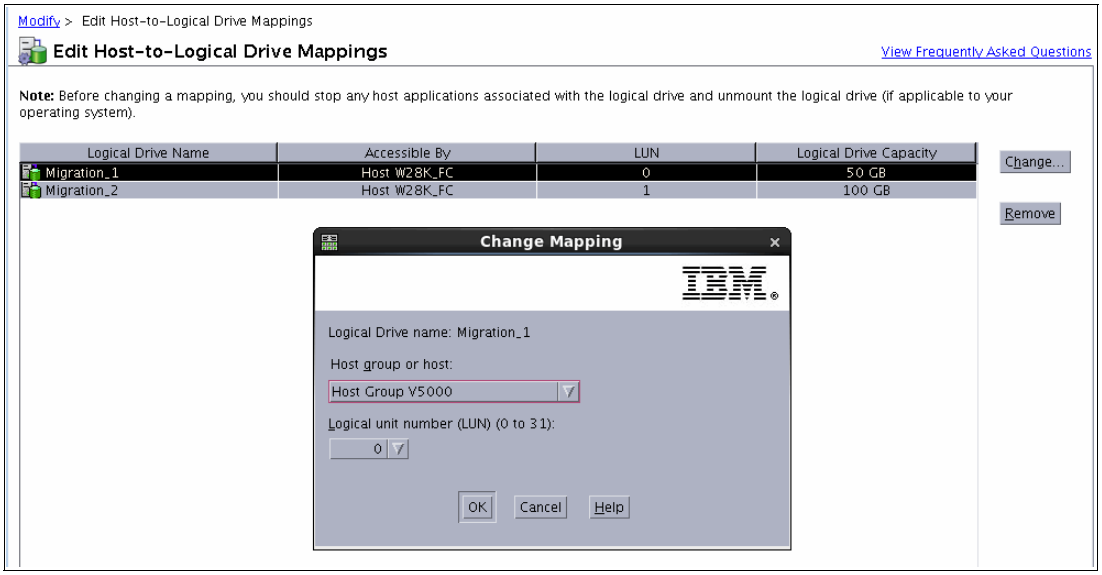

Figure 6-49 IBM DS3400 Change Mapping panel

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

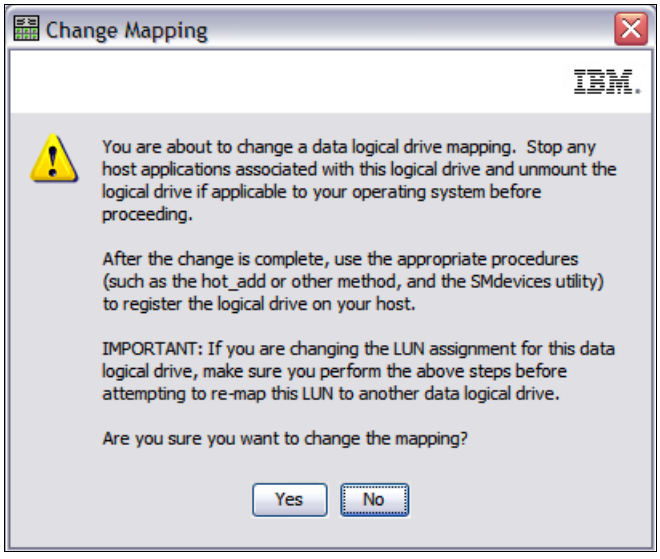

Figure 6-50 Change Mapping confirmation panel

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.

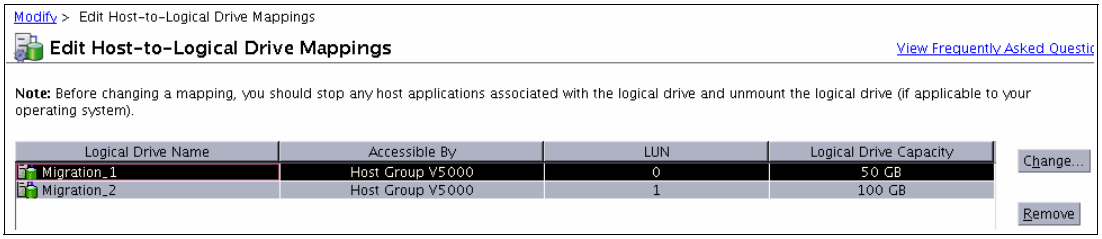

Figure 6-51 Edit Host-to-Logical Drive Mappings panel

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.

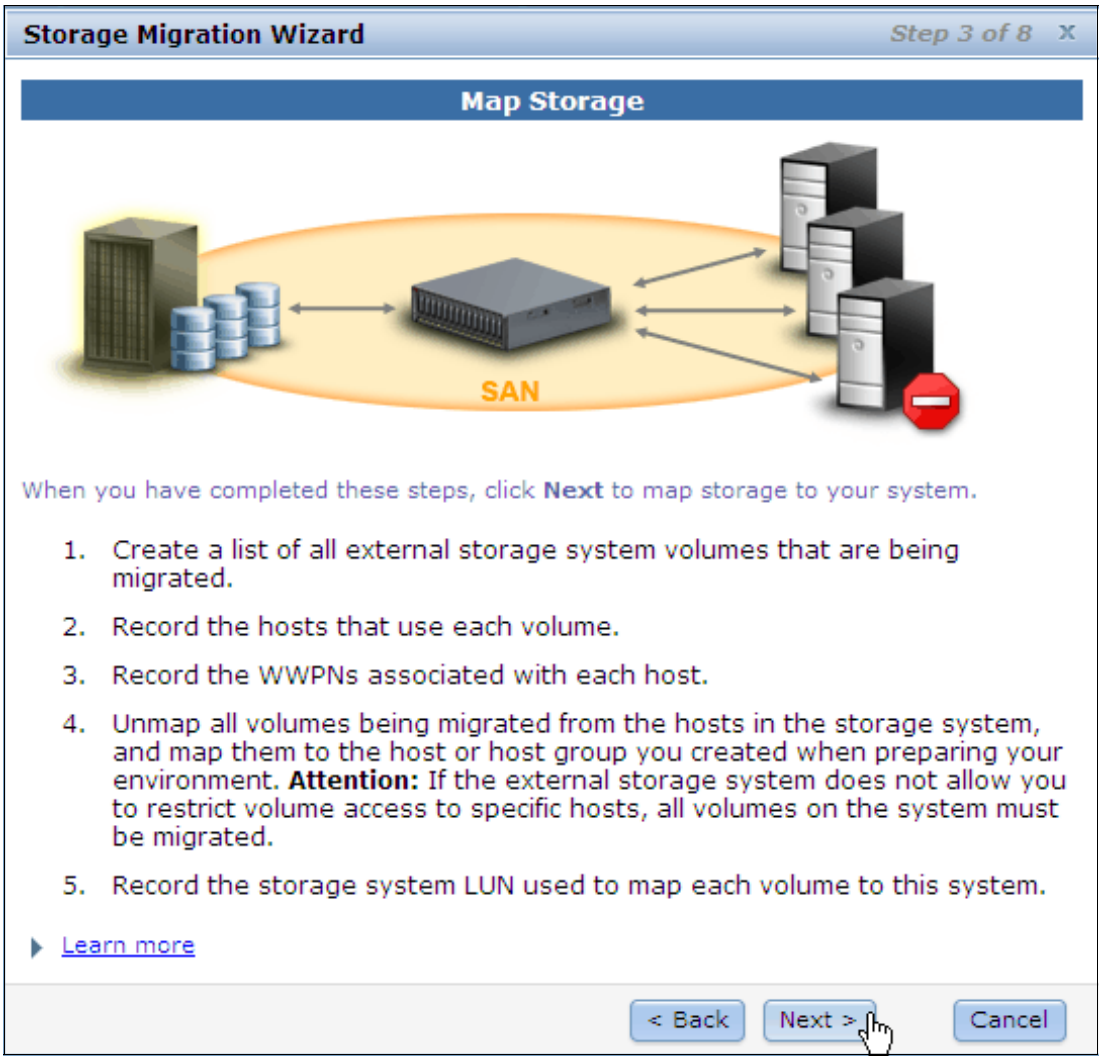

Figure 6-52 Step 3 of the Storage Migration wizard

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

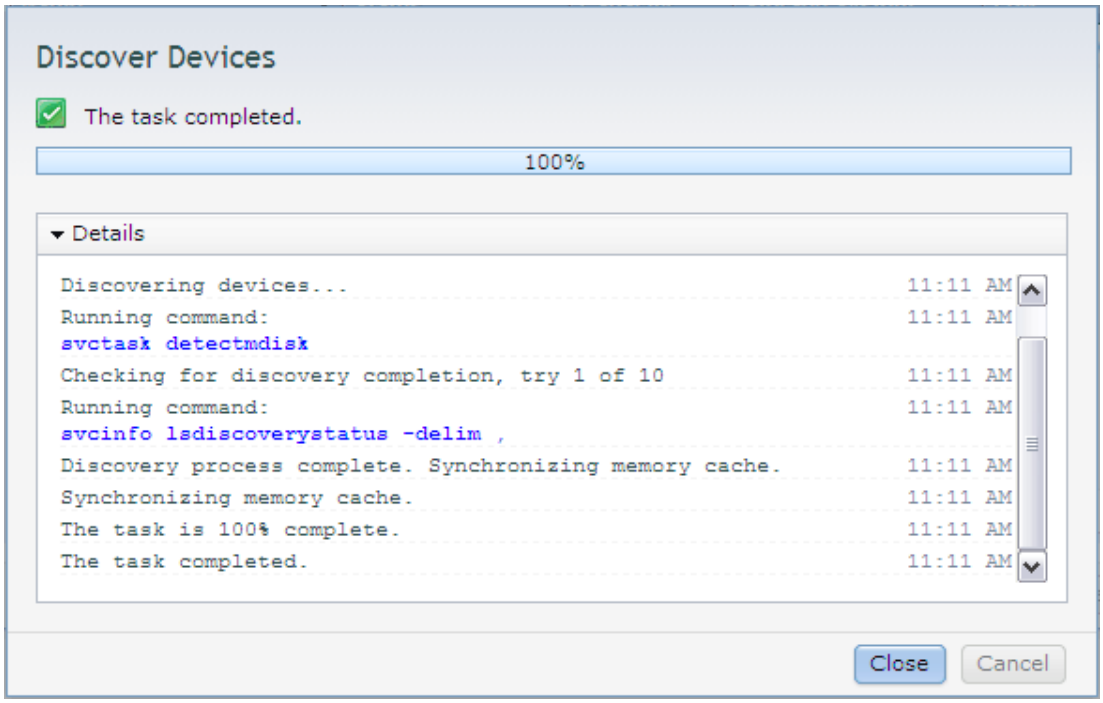

Figure 6-53 Discover Devices panel

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.

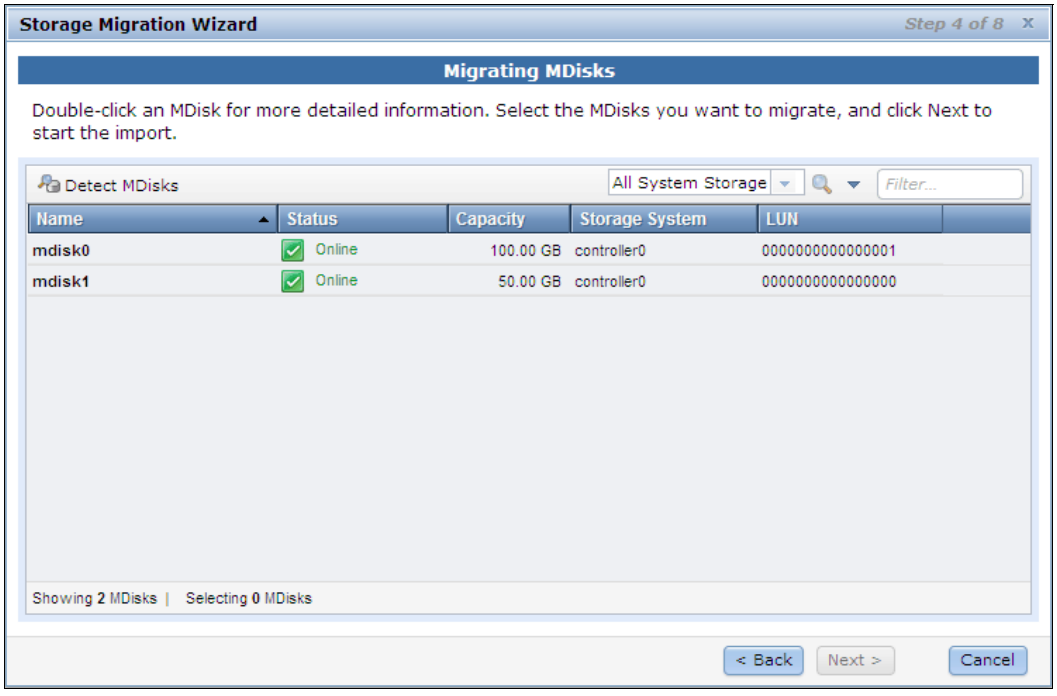

Figure 6-54 Step 4 of the Storage Migration wizard

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

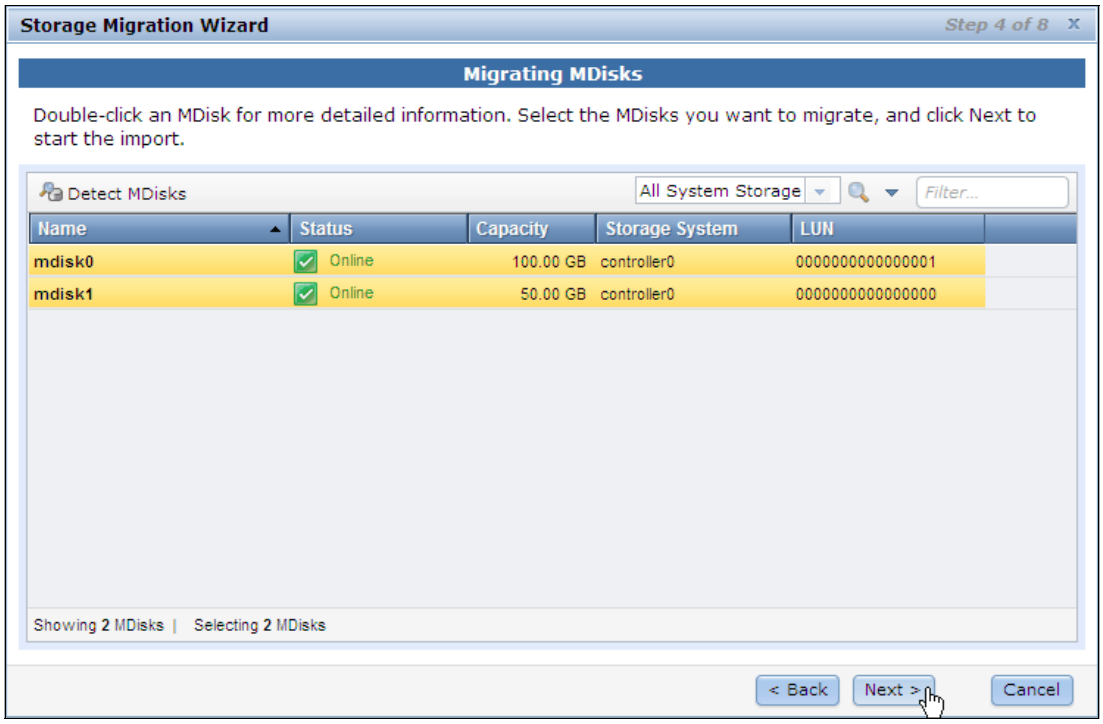

Figure 6-55 Selecting MDisk to migrate

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.

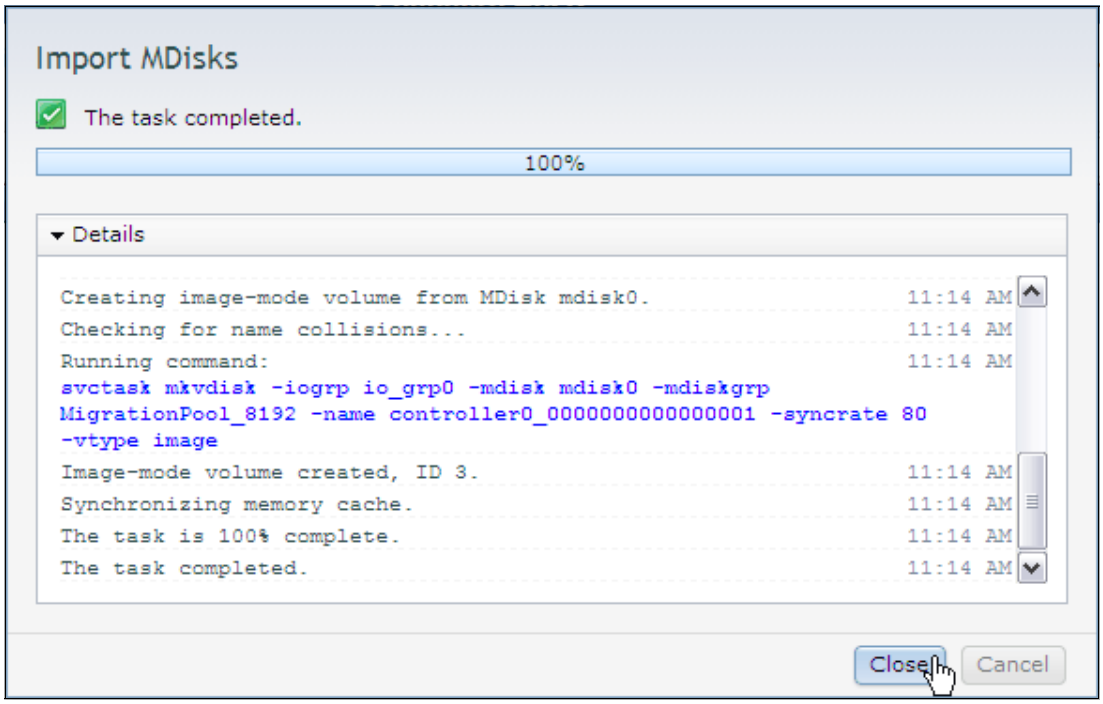

Figure 6-56 Import MDisks panel

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.

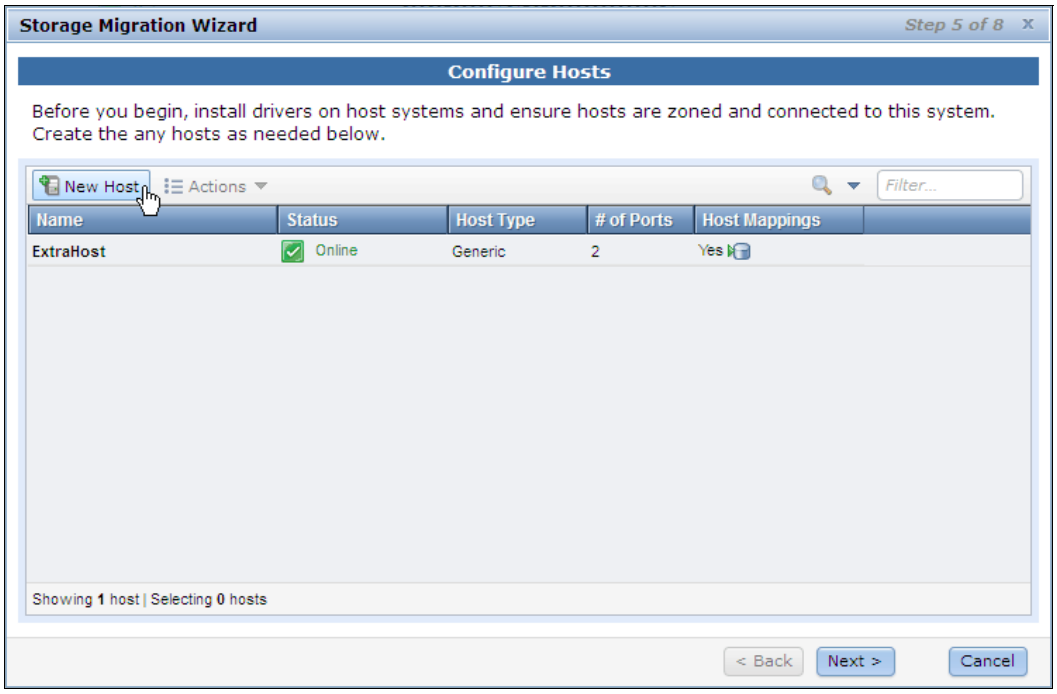

Figure 6-57 Selecting Create Host option

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.

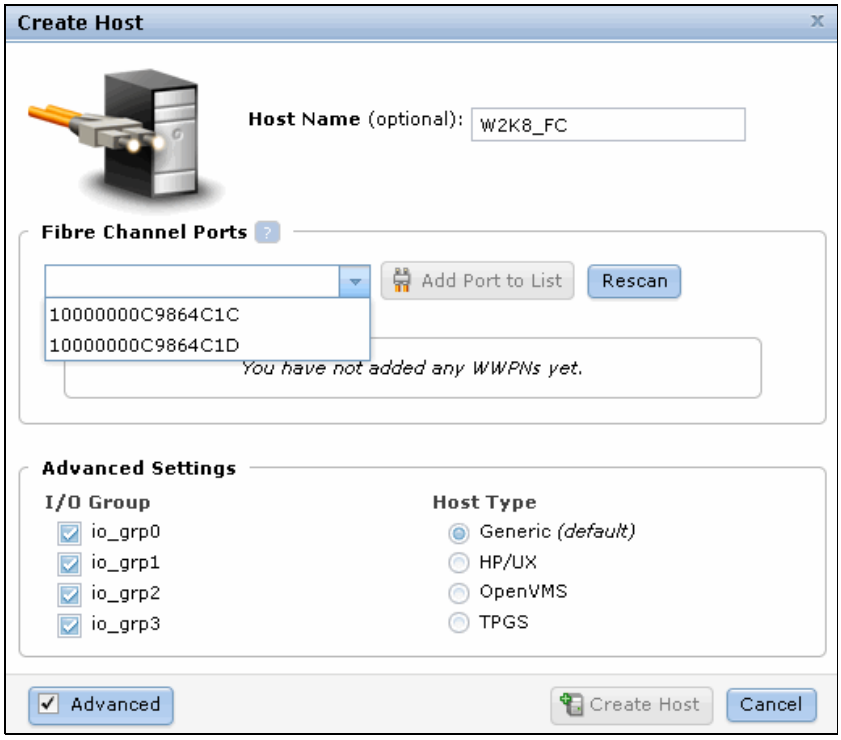

Figure 6-58 Select WWPNs

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.

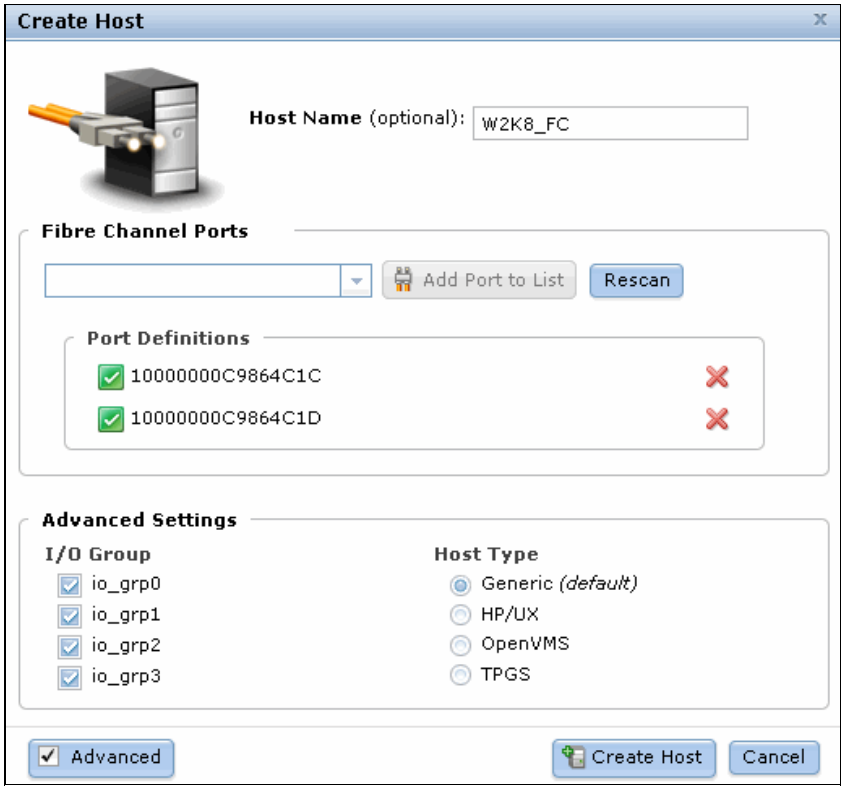

Figure 6-59 Required port definitions listed

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

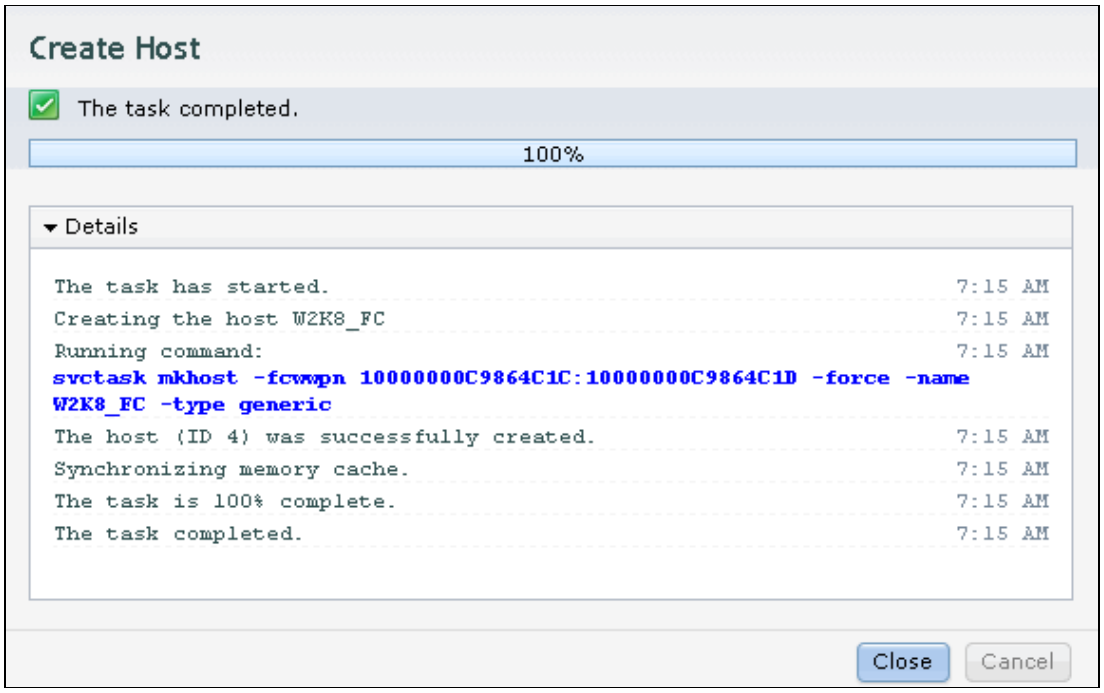

Figure 6-60 Create Host running task panel

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.

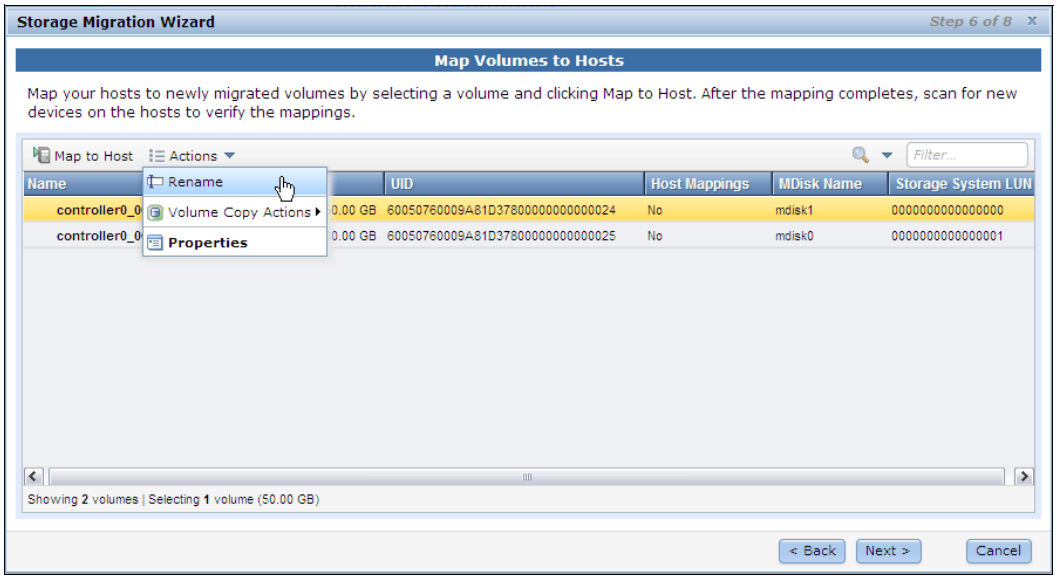

Figure 6-62 Step 6 of the Storage Migration wizard

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.

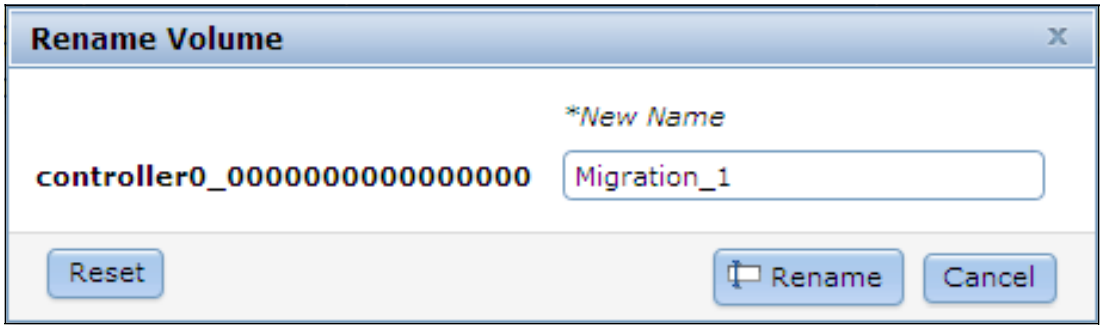

Figure 6-63 Rename volume panel

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

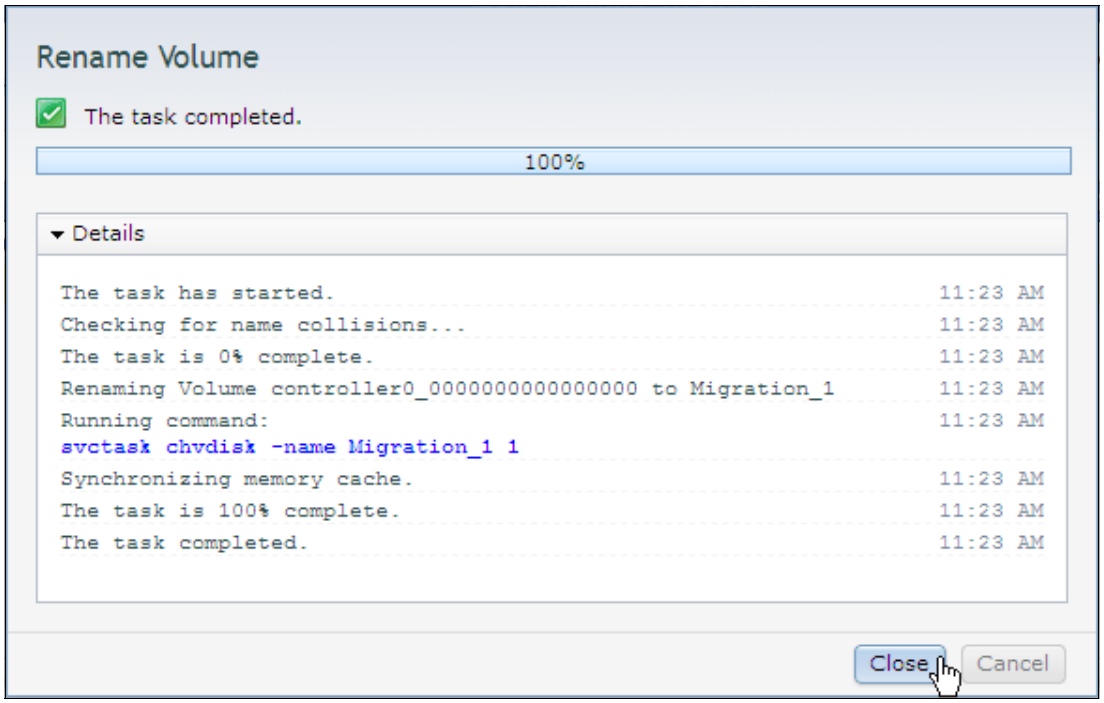

Figure 6-64 Rename Volume running task

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.

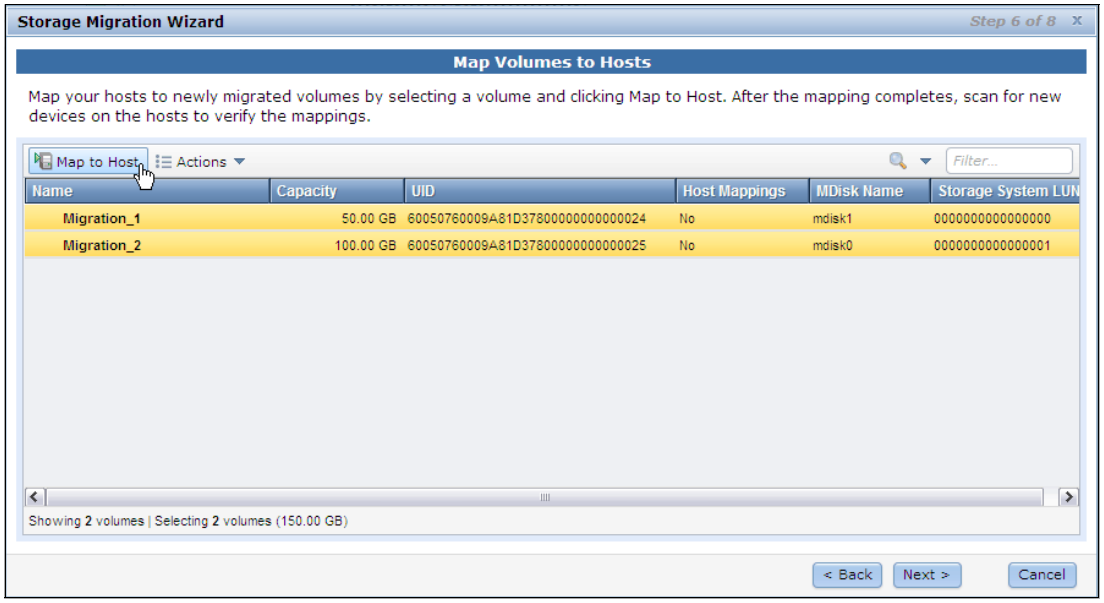

Figure 6-65 Renamed MDisks highlighted for mapping

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.

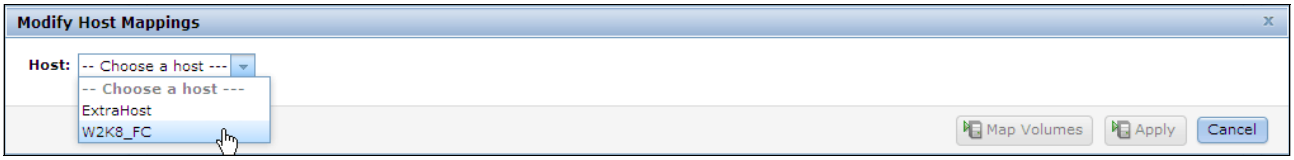

Figure 6-66 Modify Host Mappings panel

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.

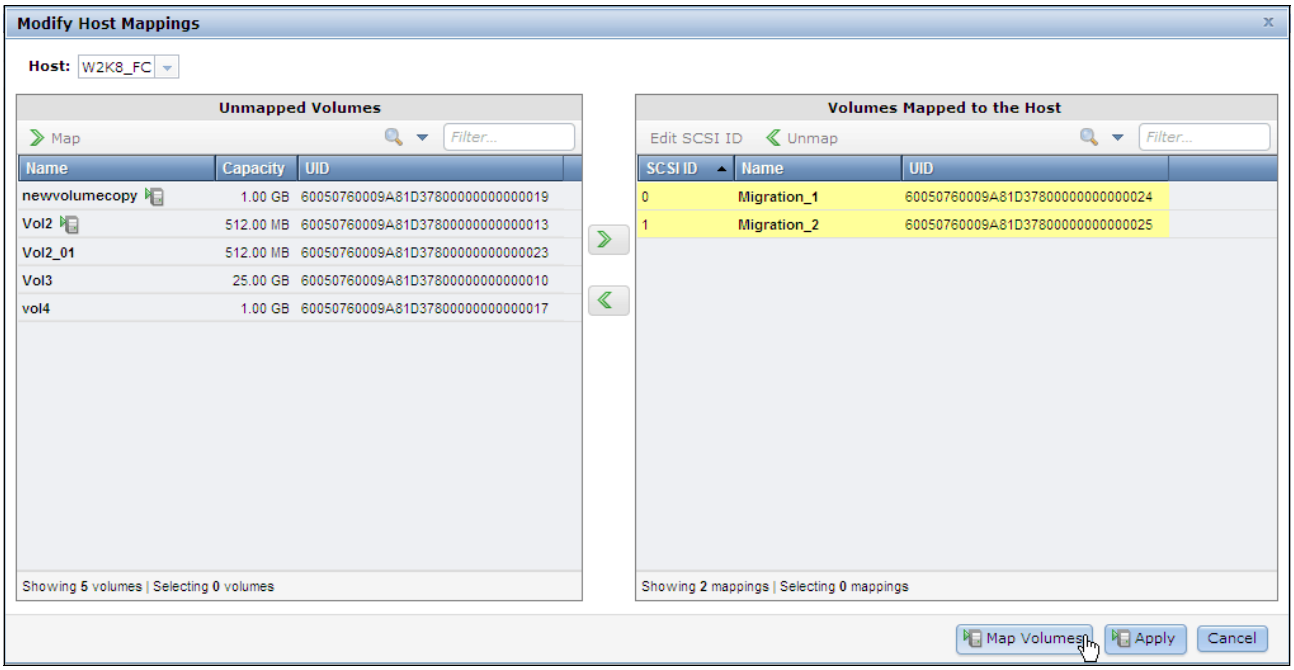

Figure 6-67 Modify Host Mappings panel

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

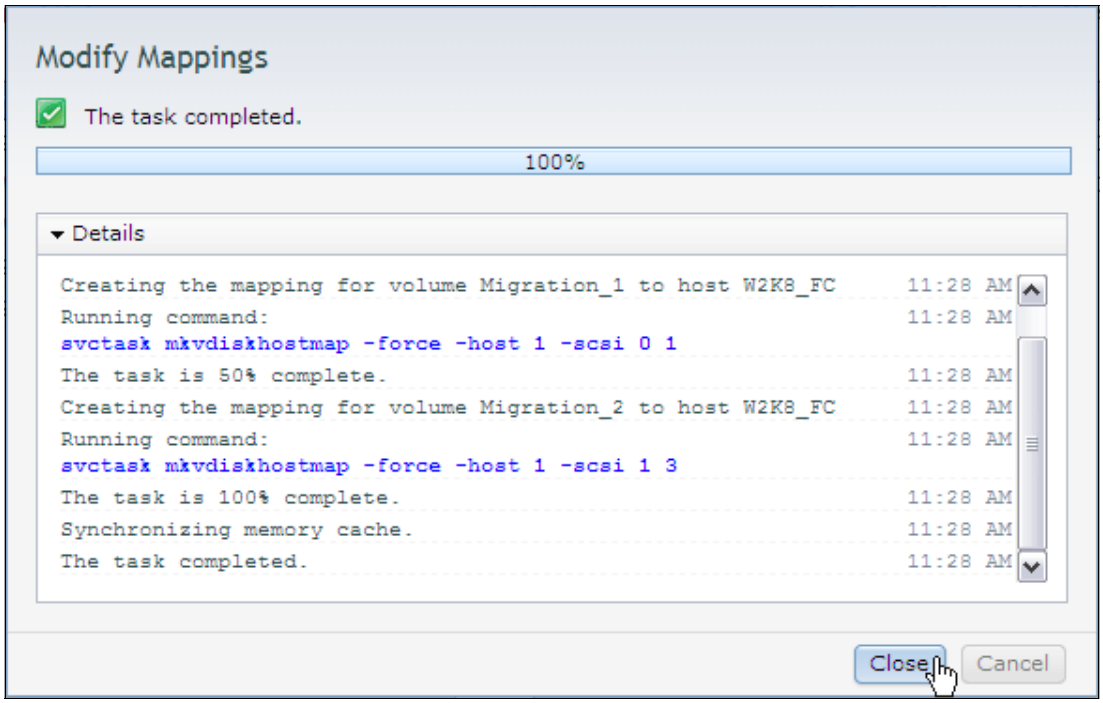

Figure 6-68 Modify Mappings running task

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.

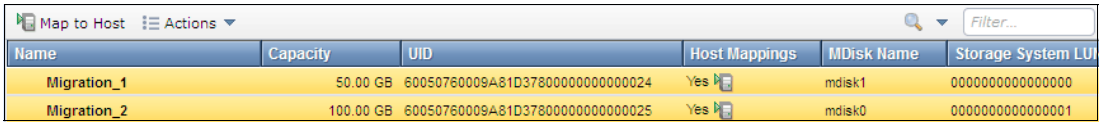

Figure 6-69 MDisks mapped

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.

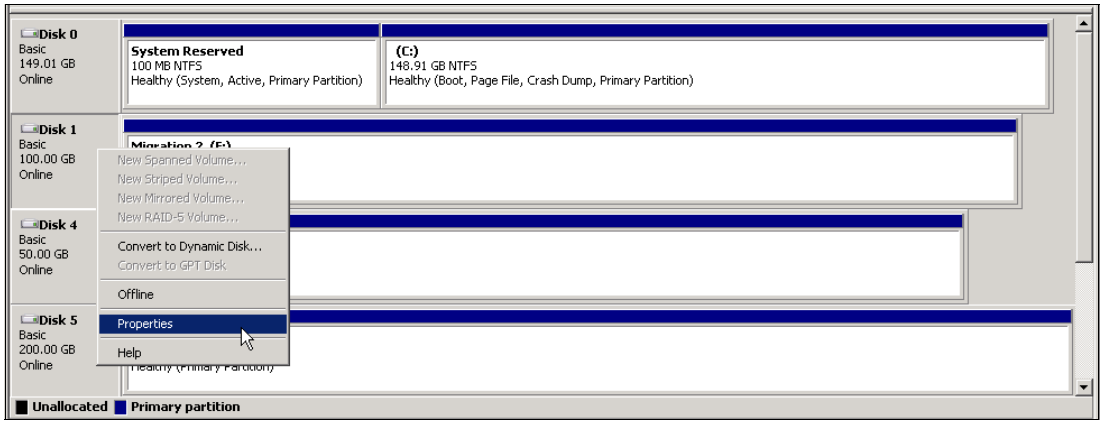

Figure 6-70 Display the disk properties from the Windows 2008 disk migration panel

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

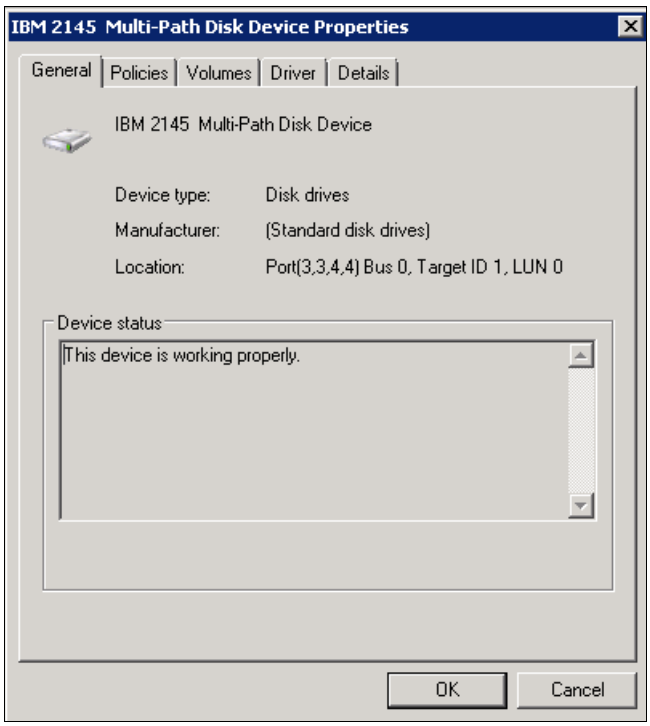

Figure 6-71 Windows 2008 properties General tab

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.

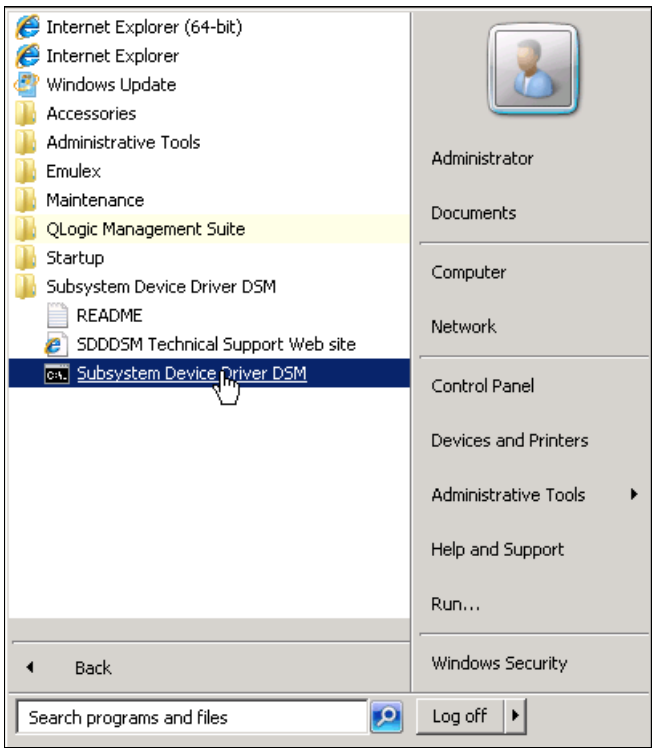

Figure 6-72 Windows 2008 R2 example: Open SSDDSM command line

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.

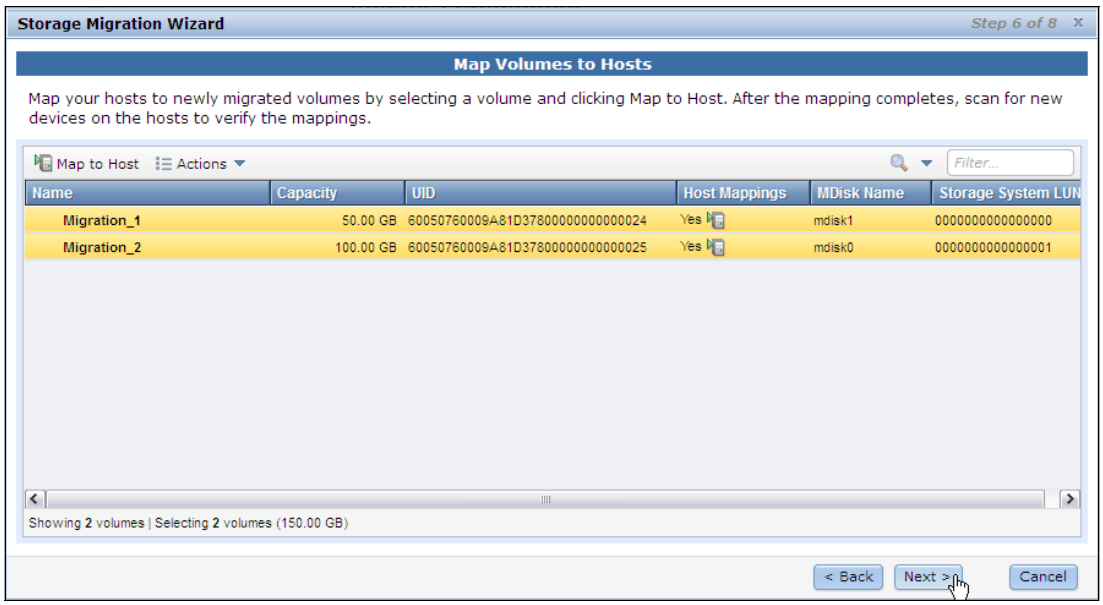

Figure 6-73 Mapped volumes and UIDs

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.

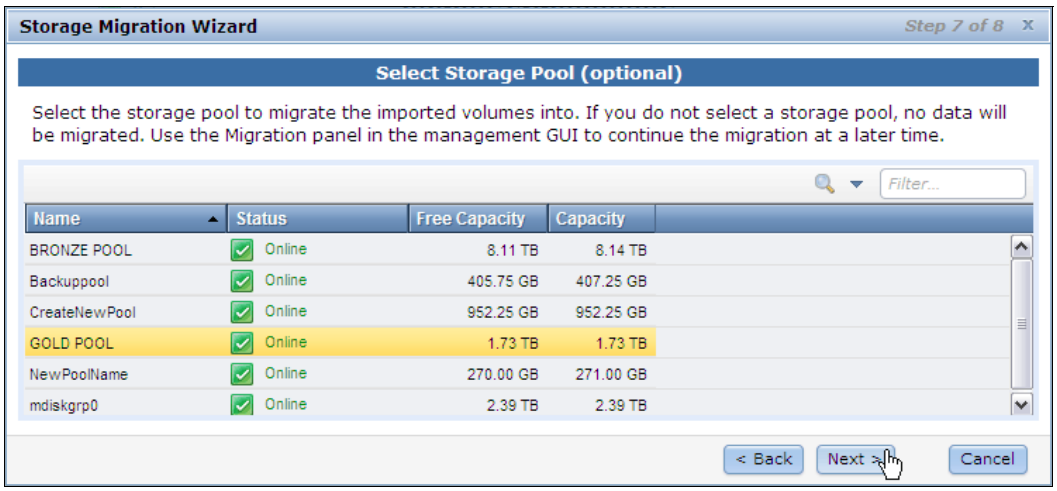

Figure 6-74 Select storage pool

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.

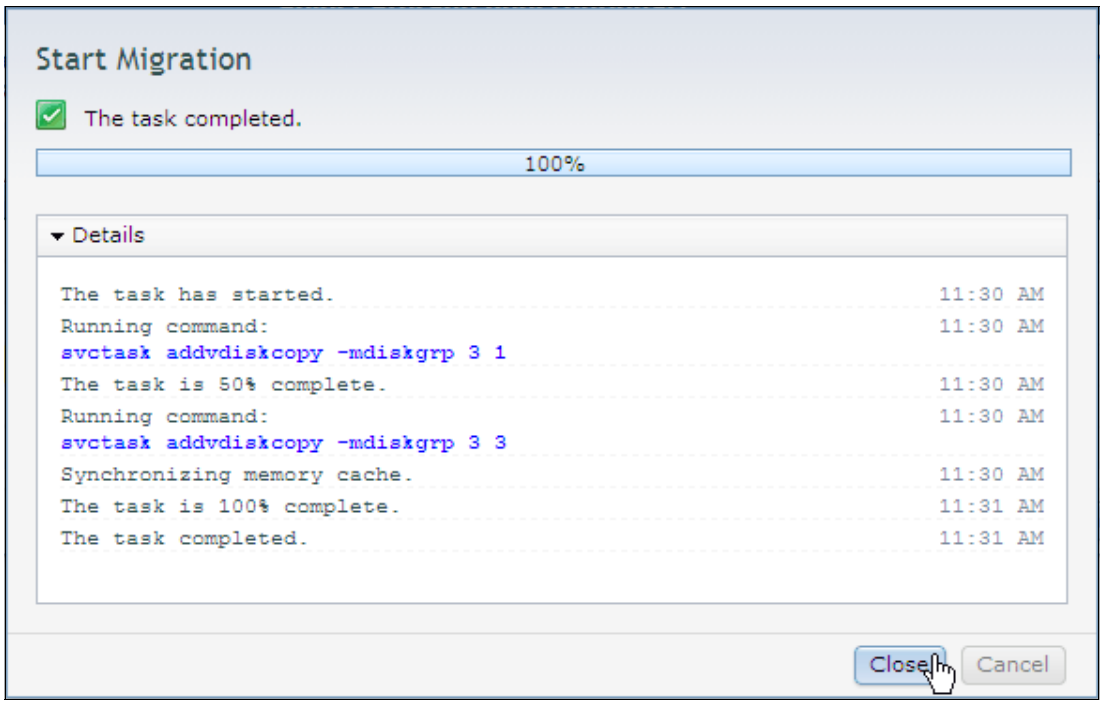

Figure 6-75 Start Migration completed task panel

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

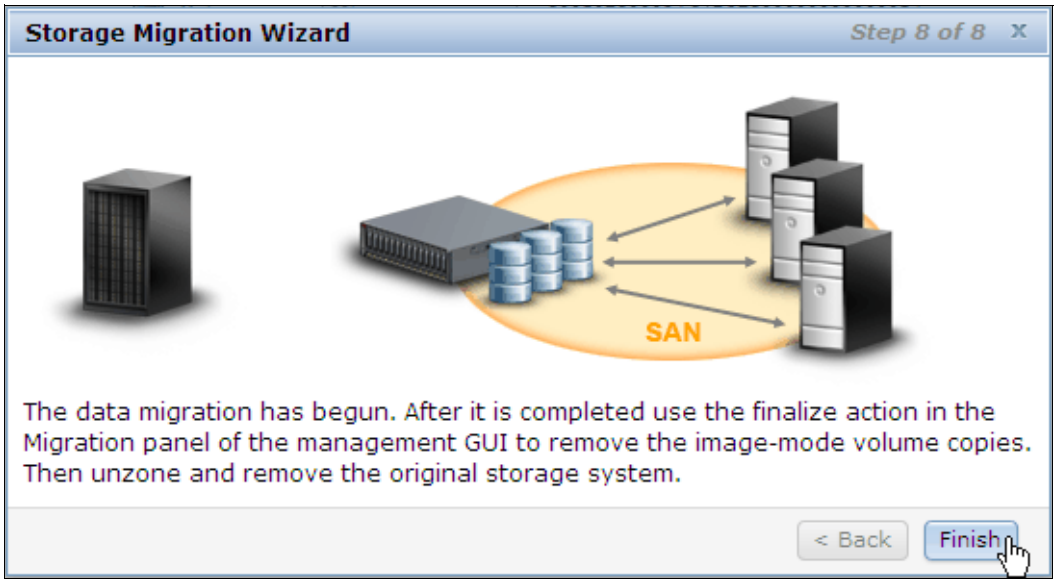

Figure 6-76 Step 8 of the Storage Migration wizard

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.

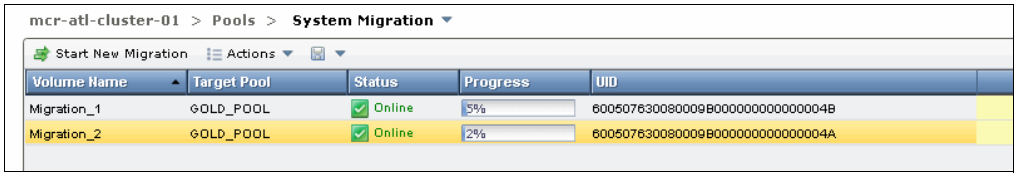

Figure 6-77 Migration progress indicators

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.

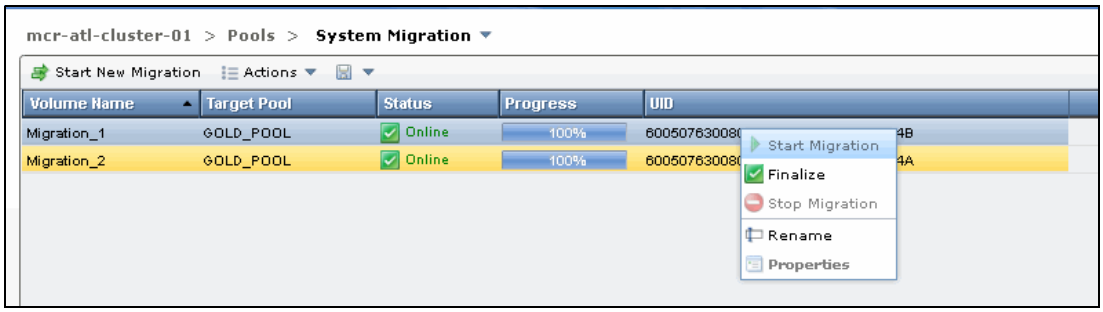

Figure 6-78 Finalize Volume Migrations

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

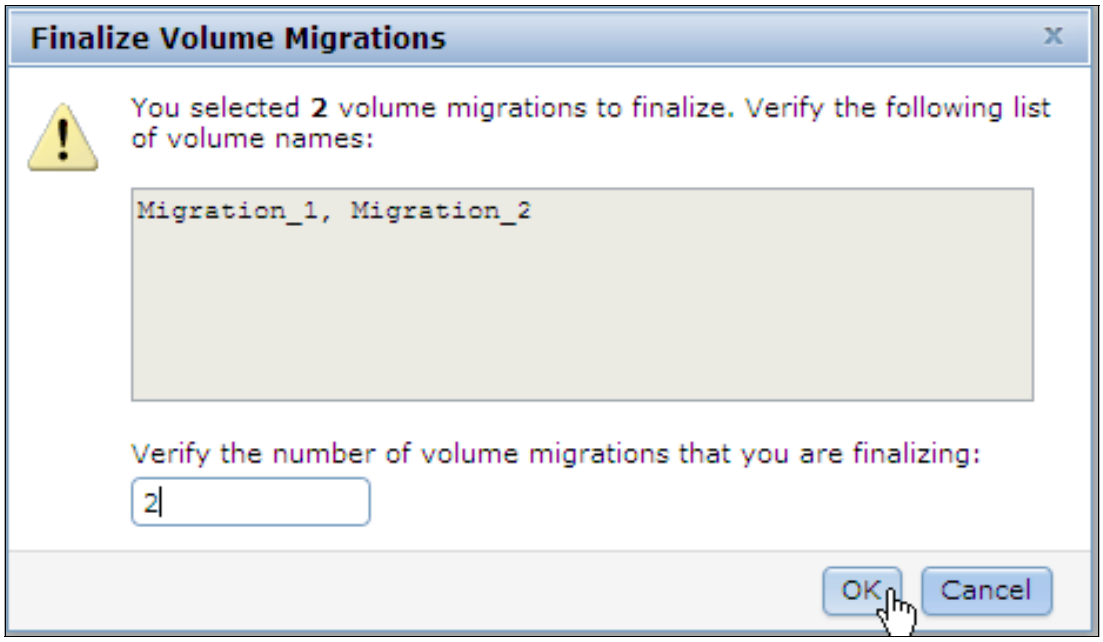

Figure 6-79 Finalize Volume Migrations panel

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.

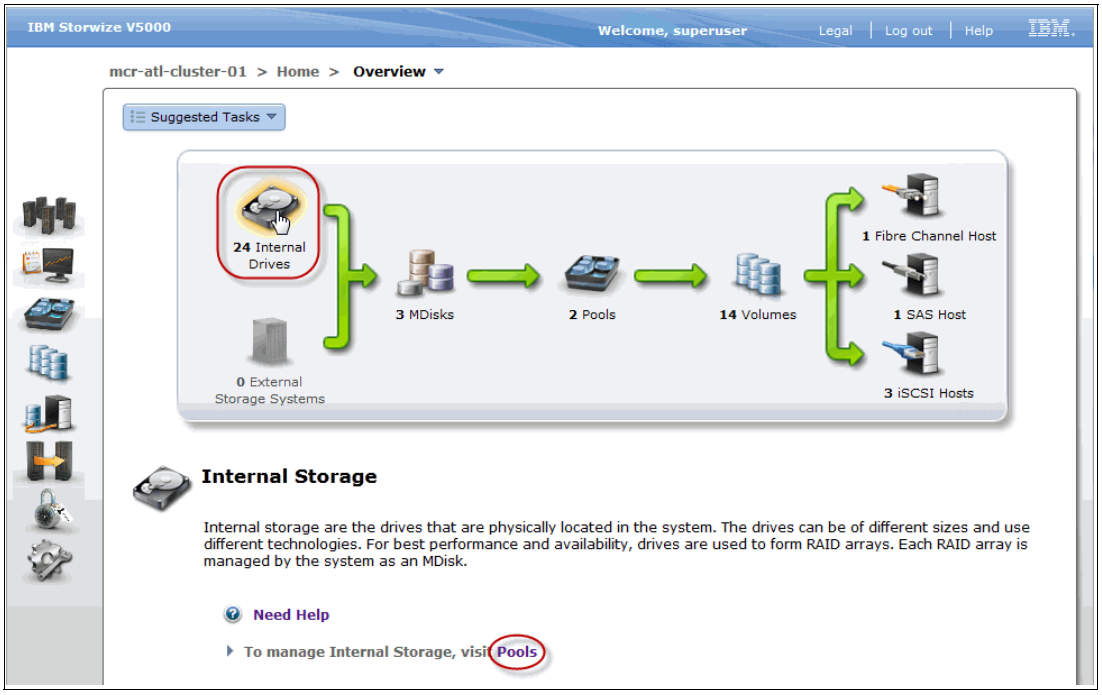

Figure 7-1 Internal Storage via Home Overview

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.

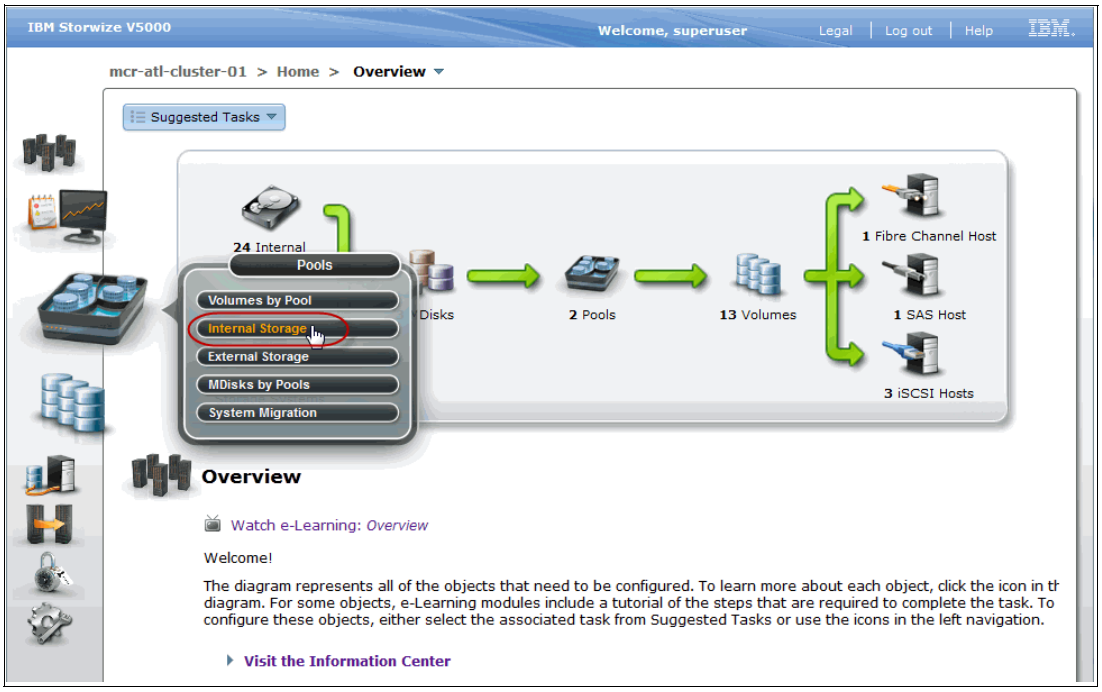

Figure 7-2 Internal Storage Details via Pools icon

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).

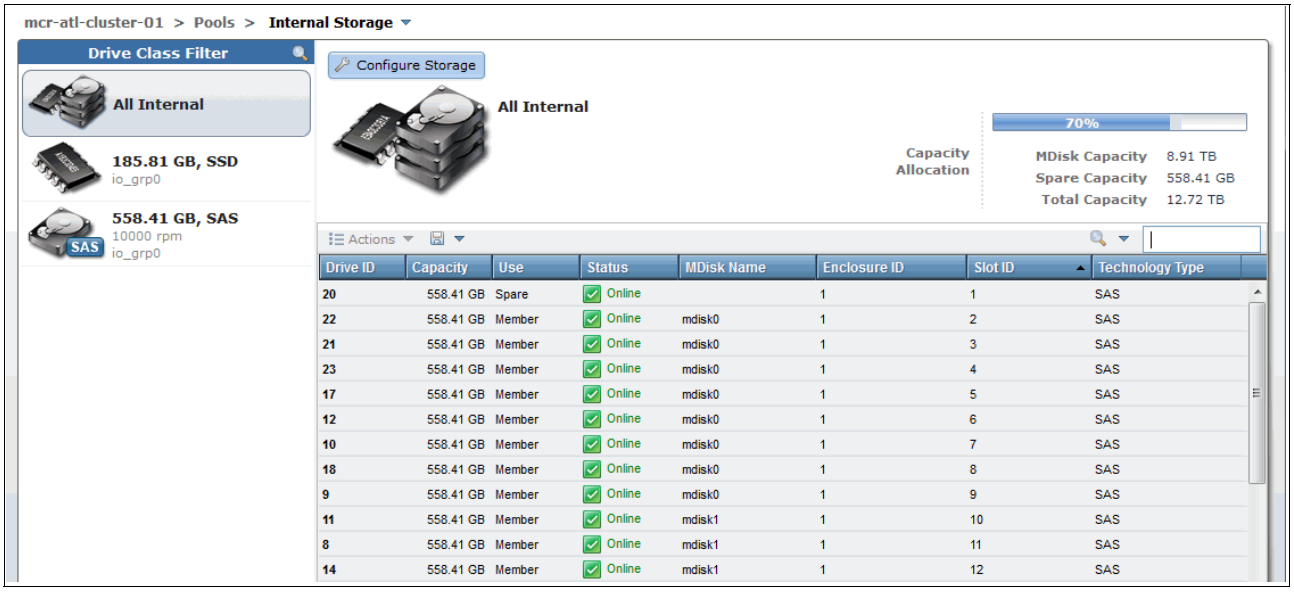

Figure 7-3 Internal storage window

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.

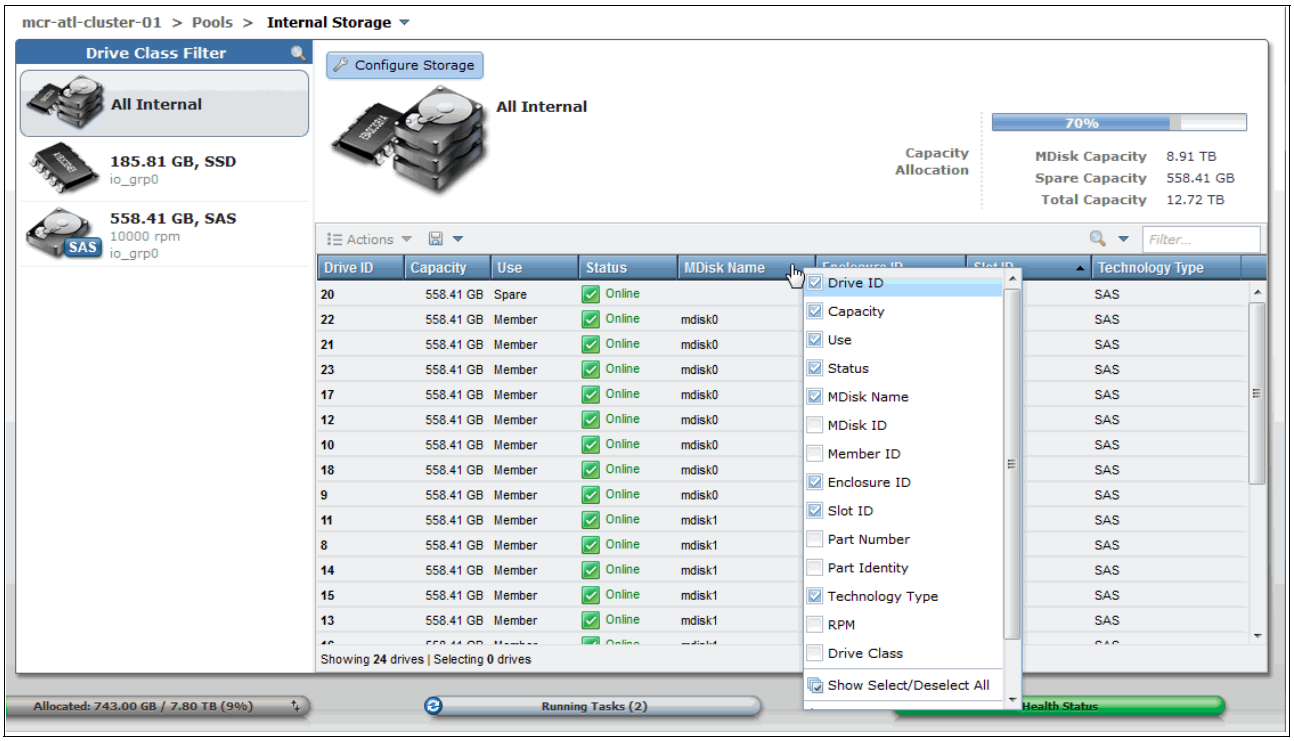

Figure 7-4 Internal storage window details selection

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.

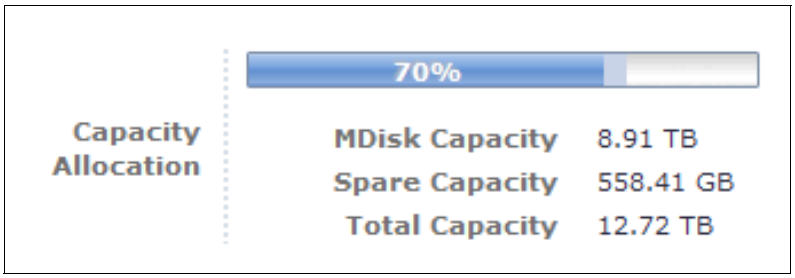

Figure 7-5 Internal storage allocation indicator

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.

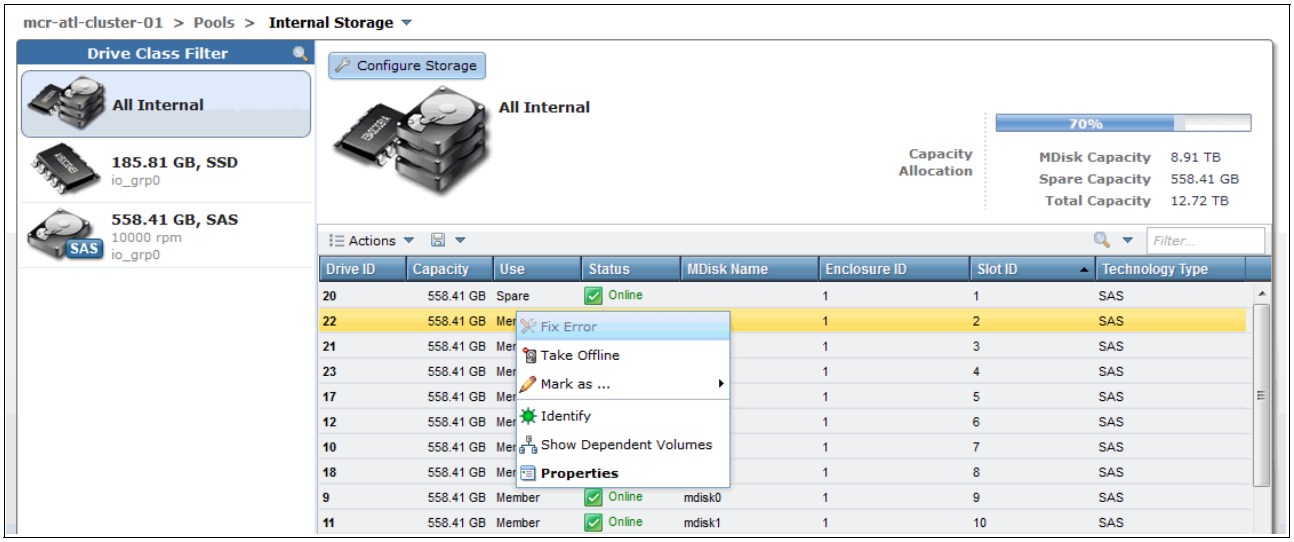

Figure 7-6 Internal drive Actions menu

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.

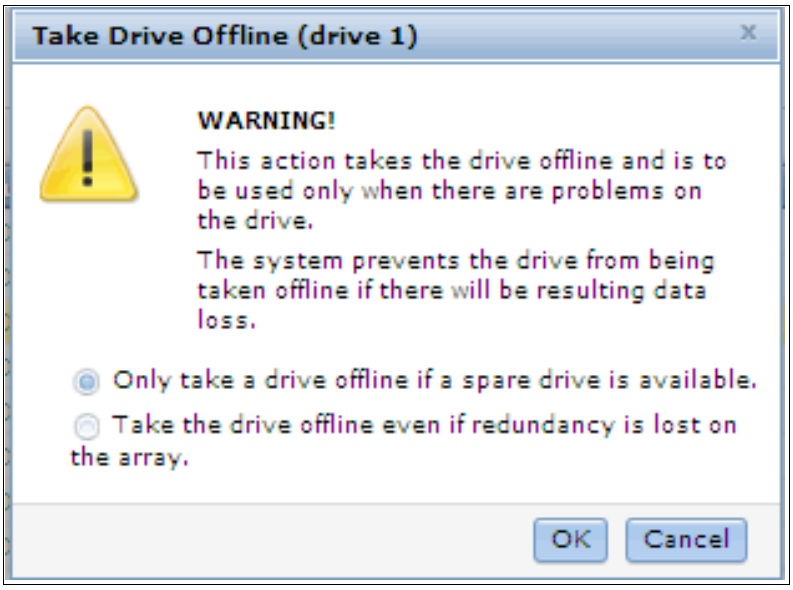

Figure 7-7 Take internal drive offline warning

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.

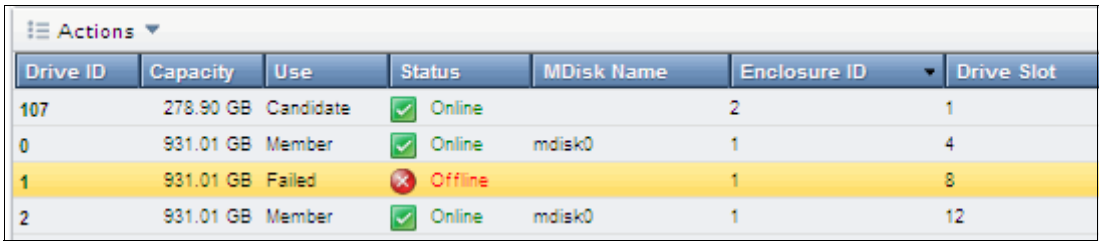

Figure 7-8 Internal drive taken offline

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.

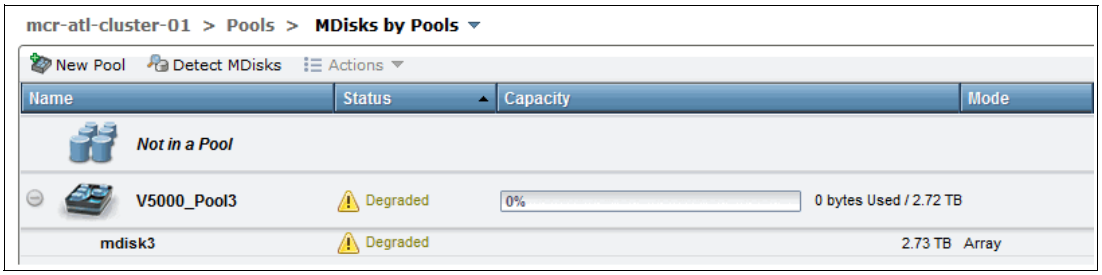

Figure 7-9 Internal drive that is failed with MDisk degraded

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.

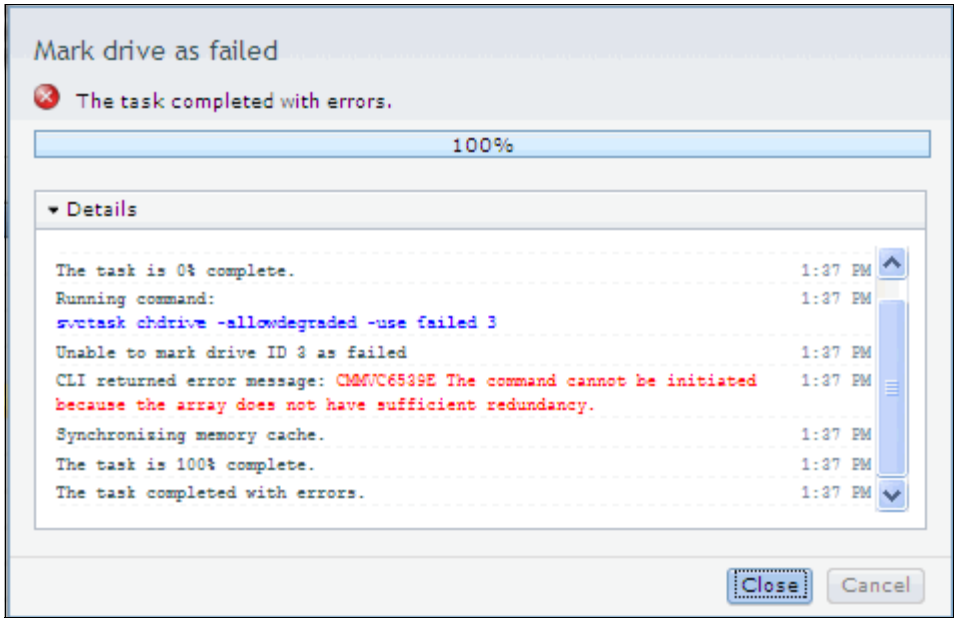

Figure 7-10 Internal drive offline not allowed because of insufficient redundancy

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.

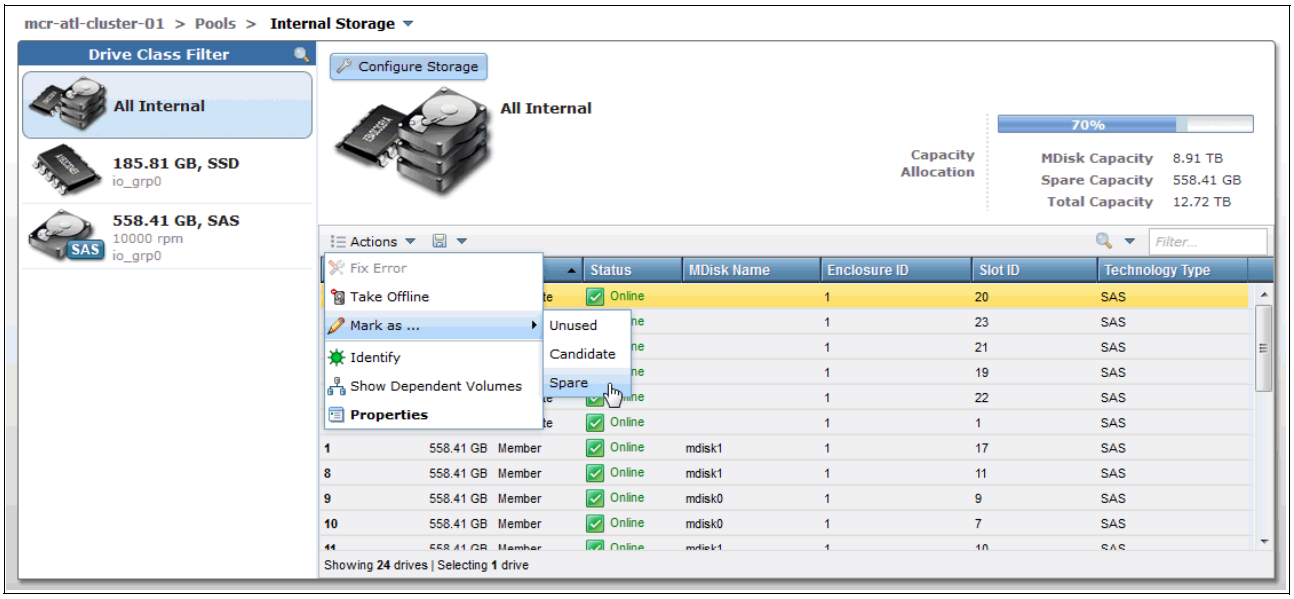

Figure 7-11 Internal drive Mark as... option

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.

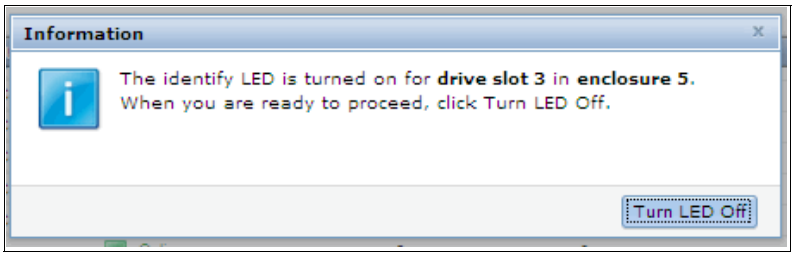

Figure 7-13 Internal drive identification

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.

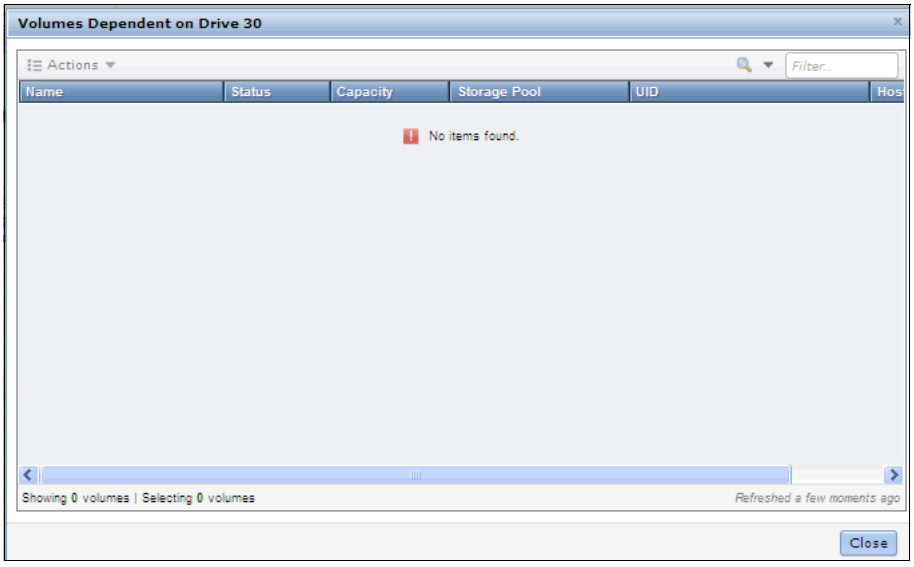

Figure 7-14 Internal drive no dependent volume

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

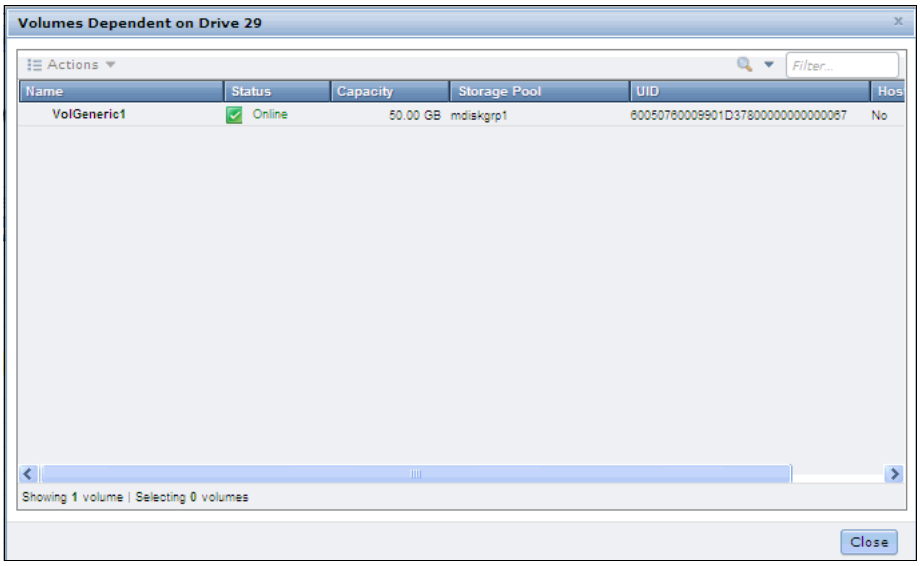

Figure 7-15 Internal drive with dependent volume

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.

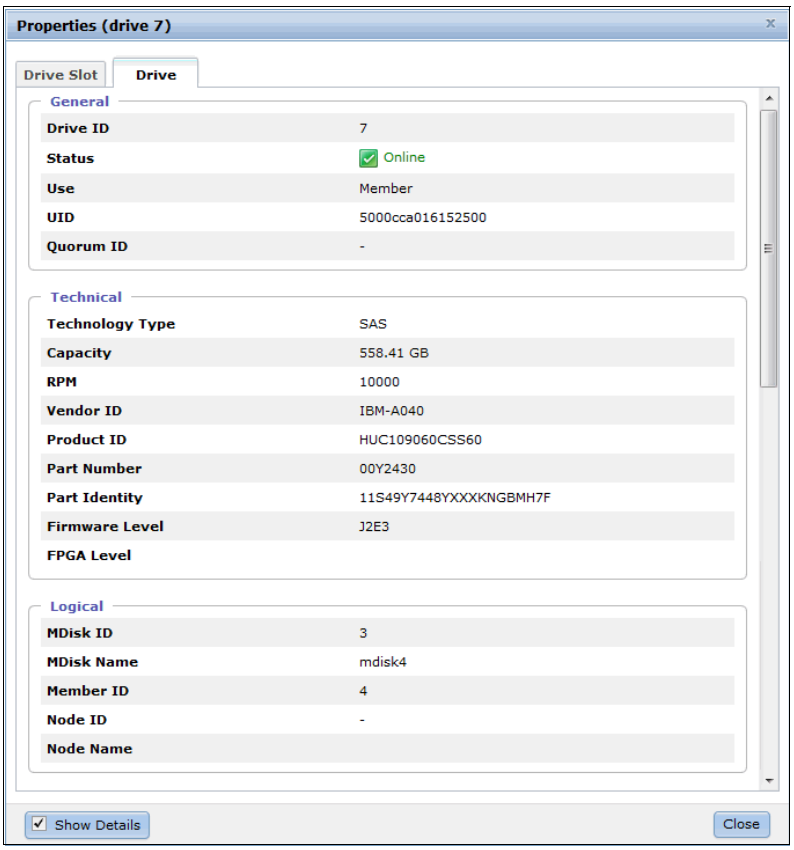

Figure 7-16 Internal drives properties: Part1

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

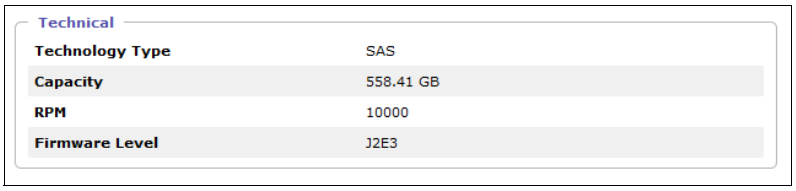

Figure 7-17 Internal drives properties no details

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.

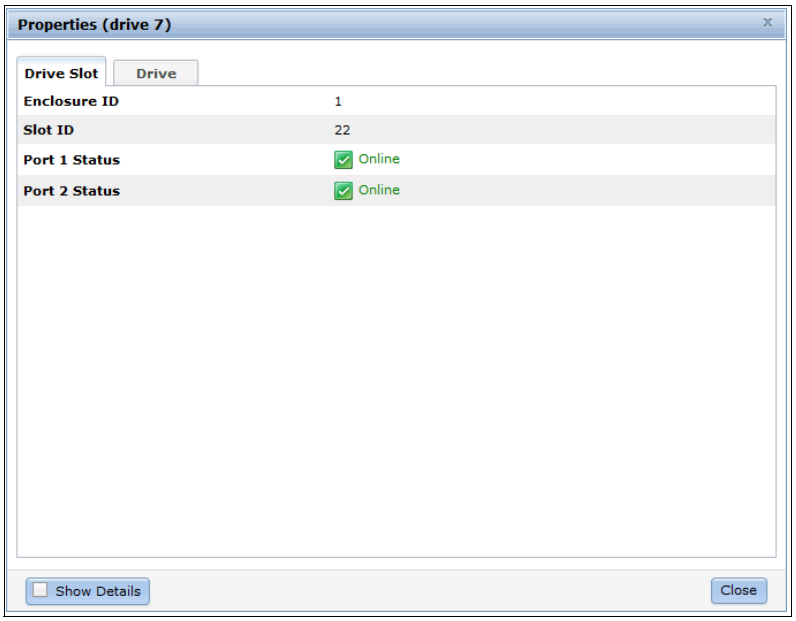

Figure 7-18 Internal drive properties slot

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.

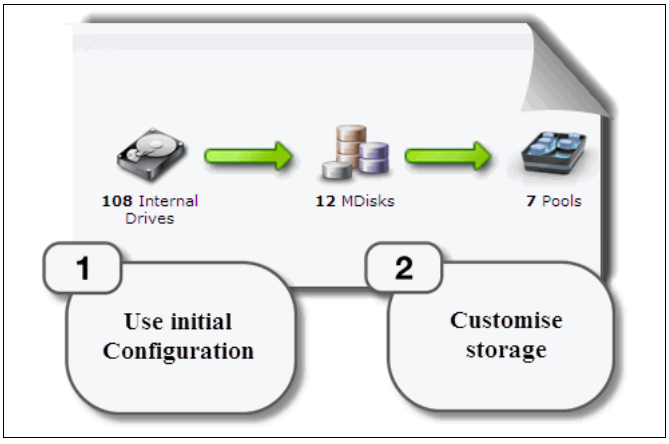

Figure 7-19 Decision to customize storage configuration

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.

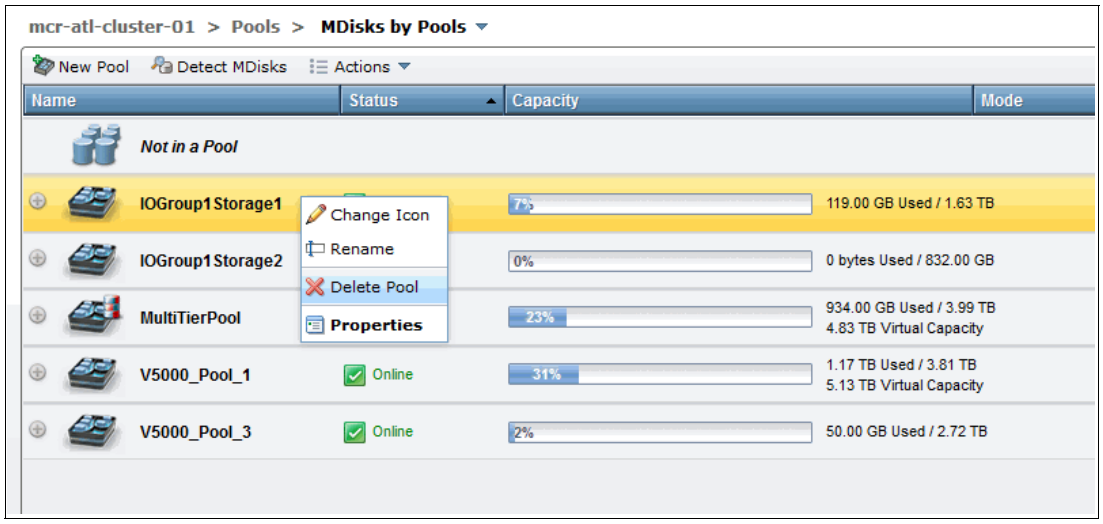

Figure 7-20 Delete selected pool

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.

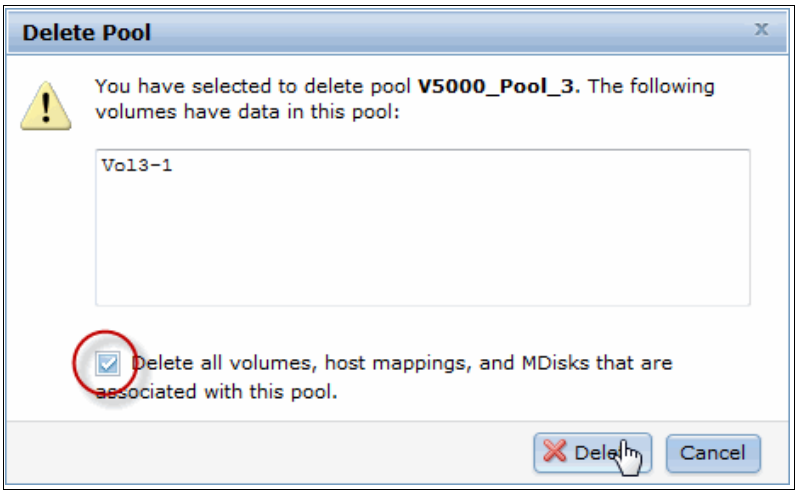

Figure 7-21 Delete pool confirmation

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.

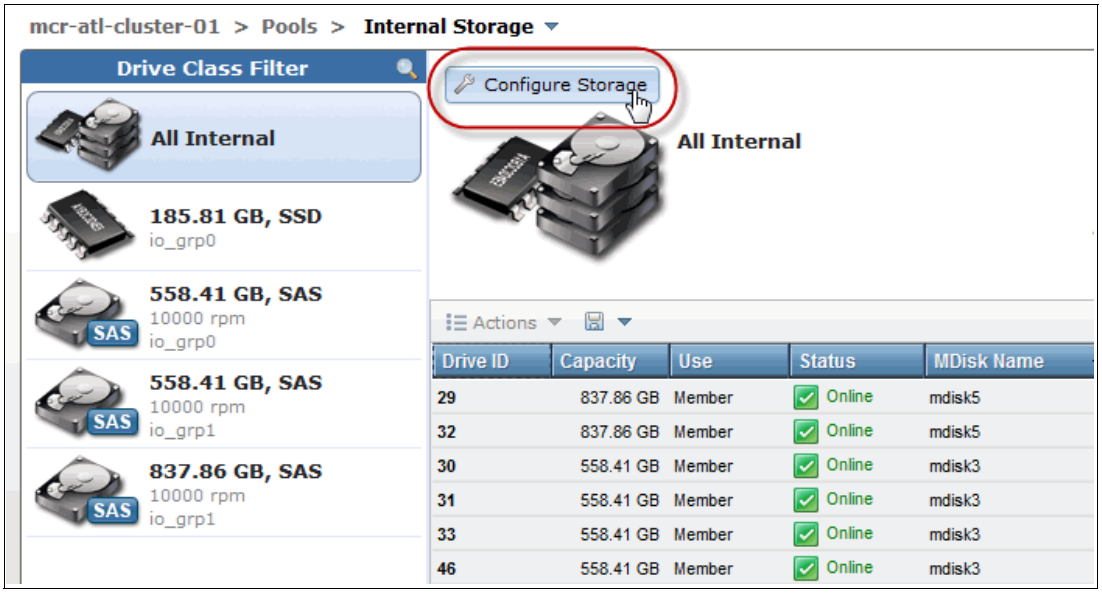

Figure 7-22 Click Configure Storage

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.

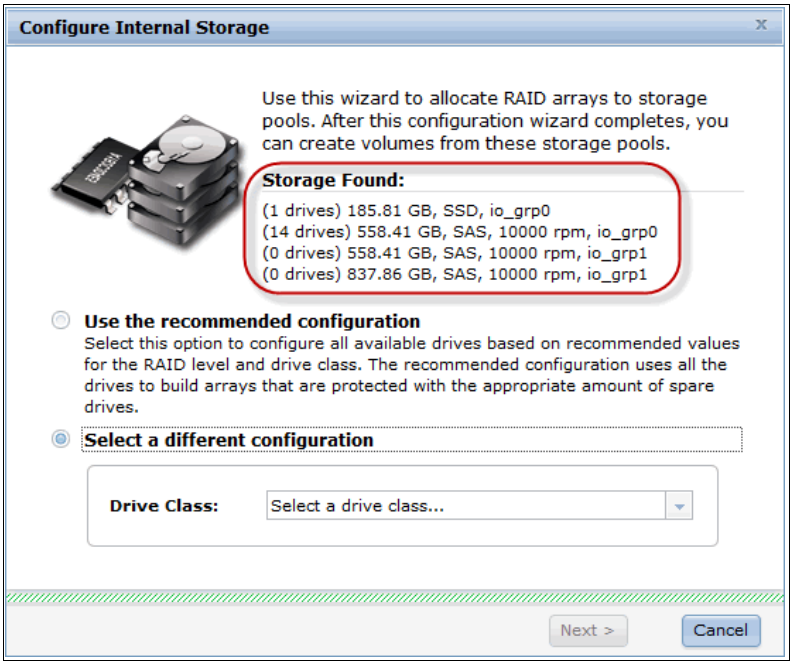

Figure 7-23 Available drives for new MDisk

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.

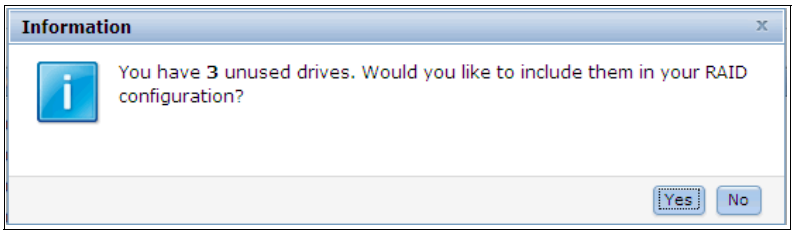

Figure 7-24 Unused drives warning

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.

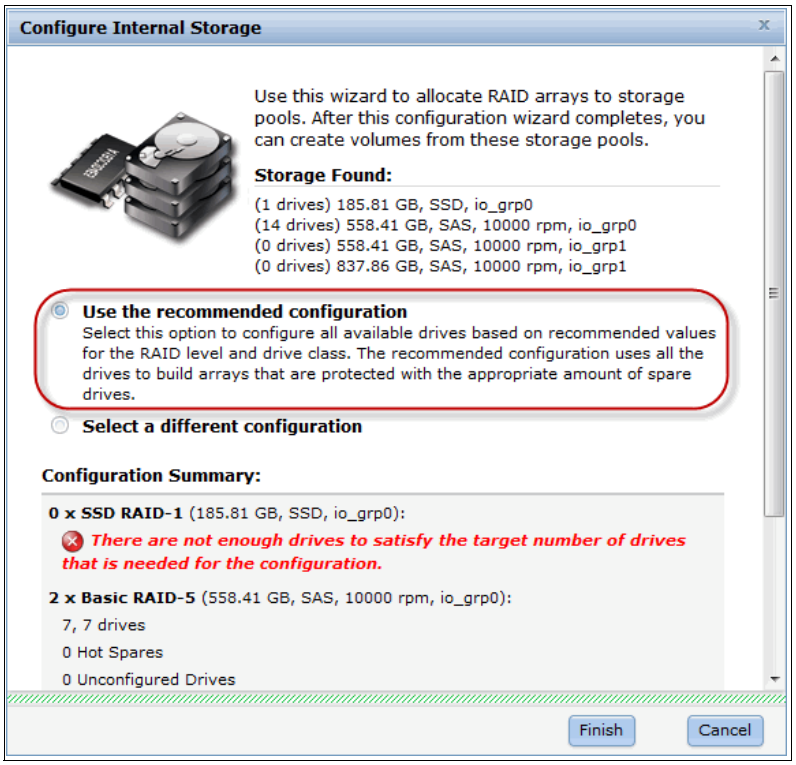

Figure 7-25 The recommended configuration

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.

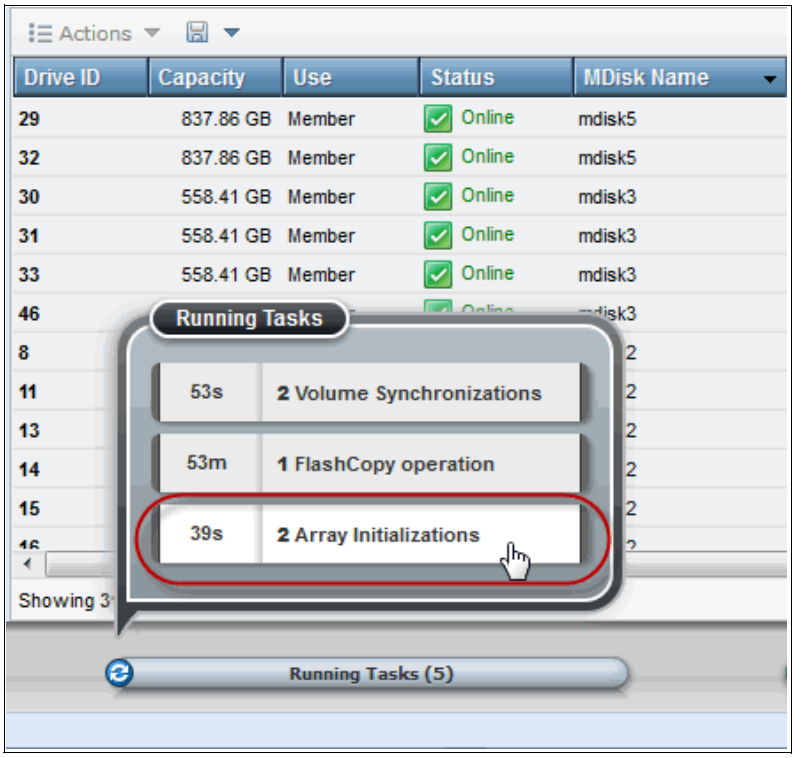

Figure 7-26 Running task panel

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.

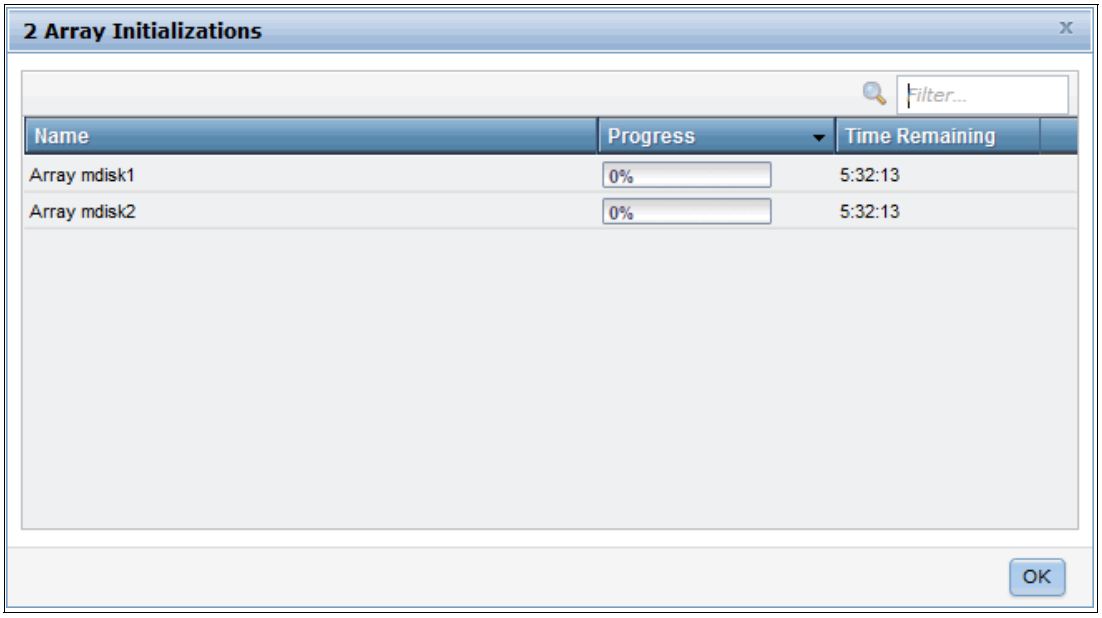

Figure 7-27 Initialization progress view

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.

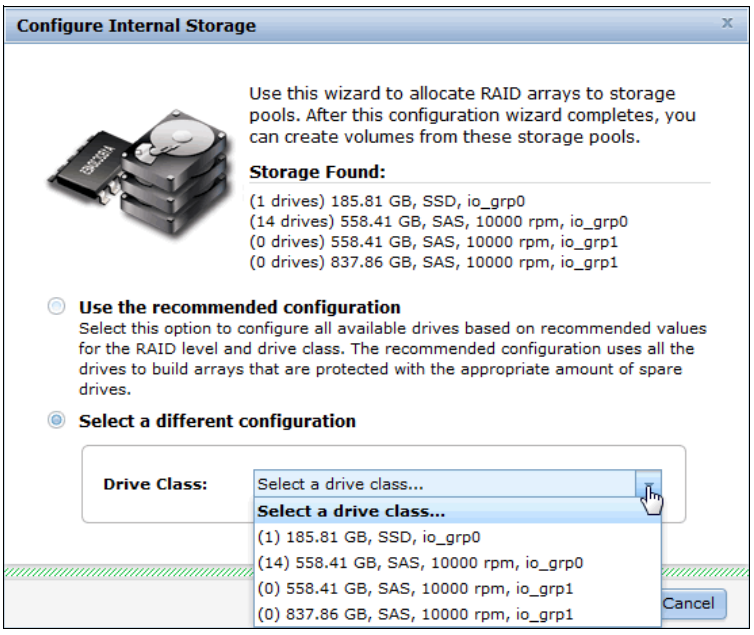

Figure 7-28 Select drive class for new configuration

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

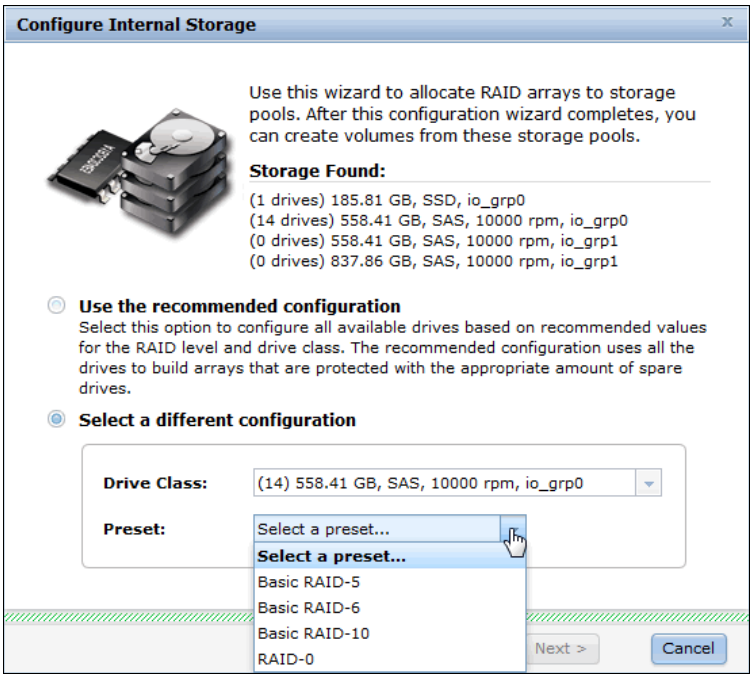

Figure 7-29 Select the RAID preset

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).

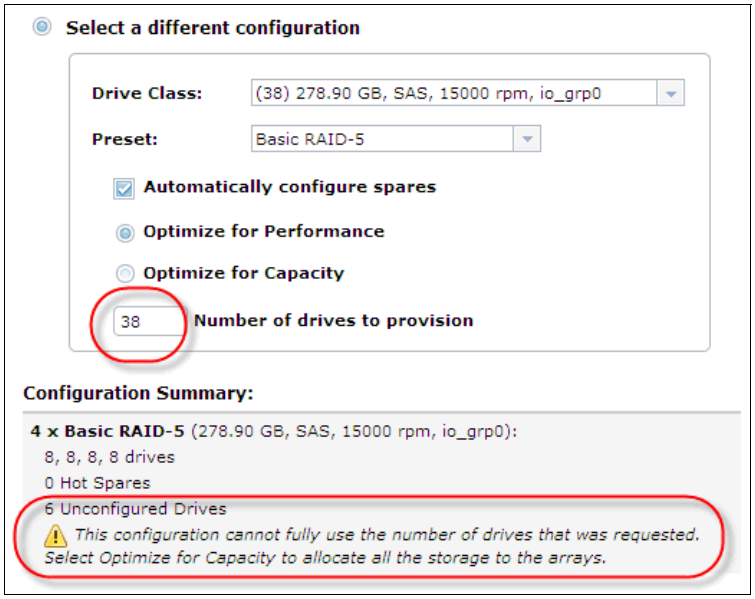

Figure 7-30 Optimization for performance failed

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

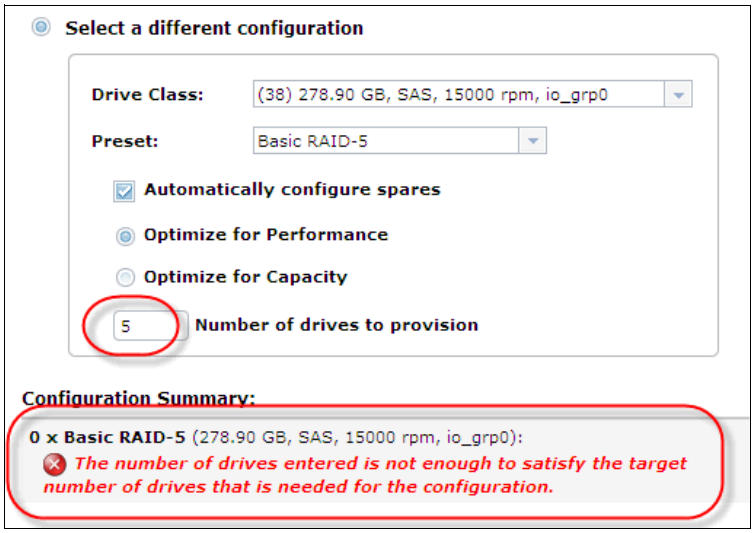

Figure 7-31 Not enough drives for performance optimization

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

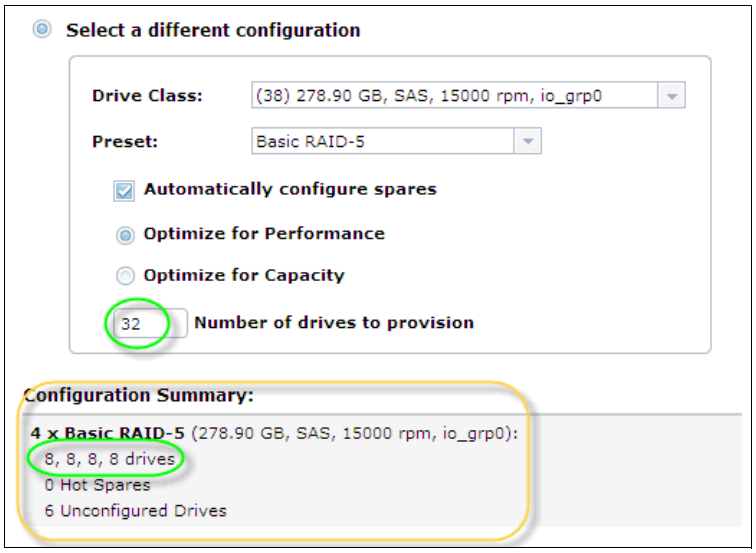

Figure 7-32 Arrays match performance goals

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.

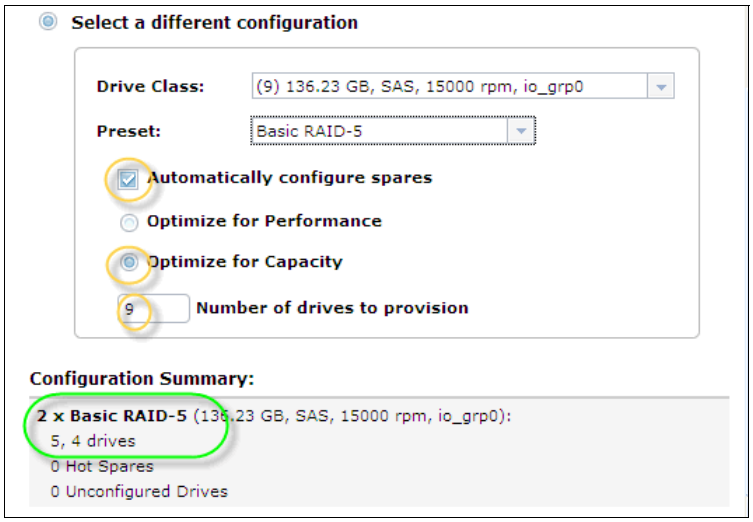

Figure 7-33 Capacity optimized configuration

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.

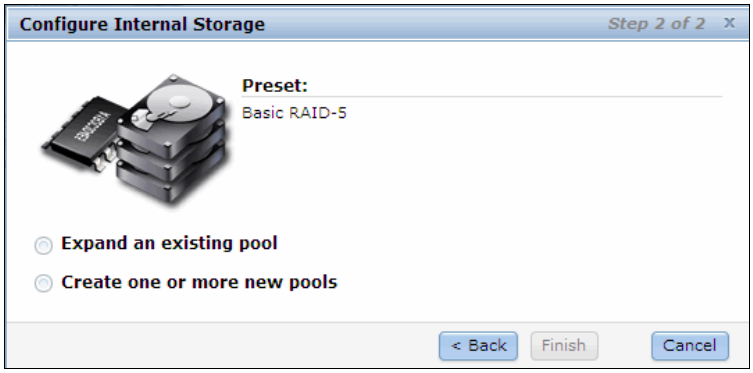

Figure 7-34 Storage pool selection

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.

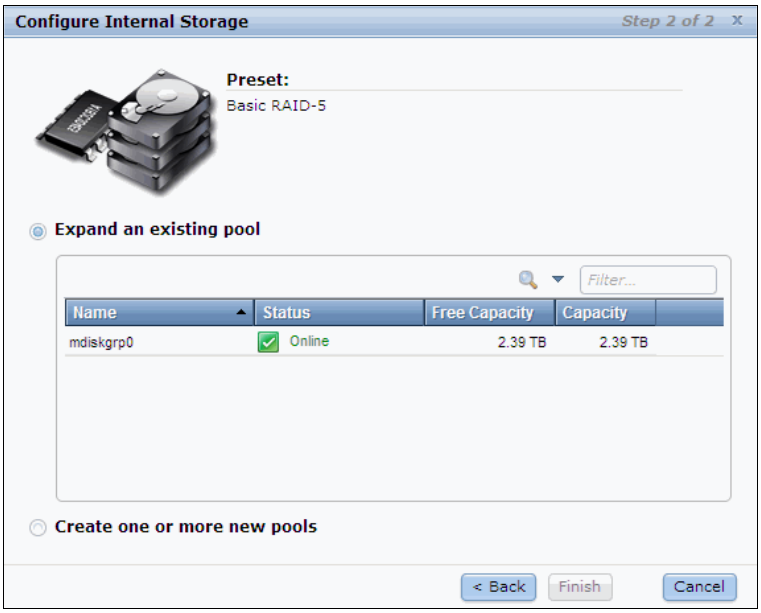

Figure 7-35 List of matching storage pool

b. Create one or more pools.

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

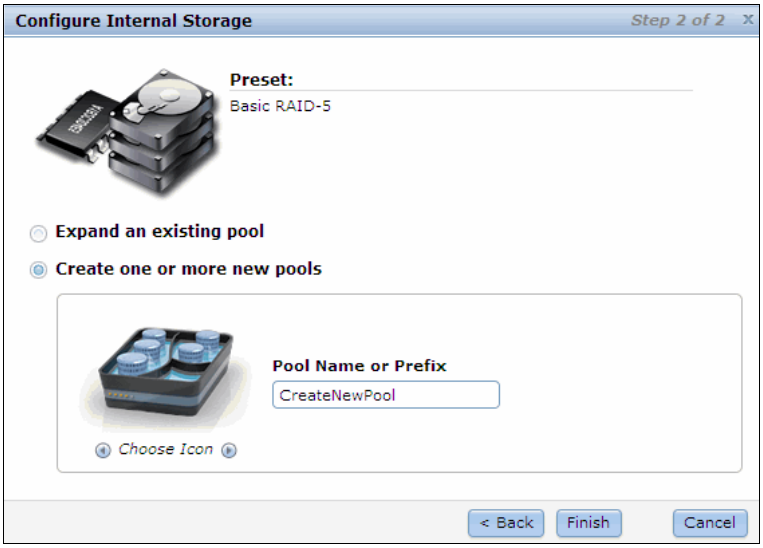

Figure 7-36 Create new pool

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.

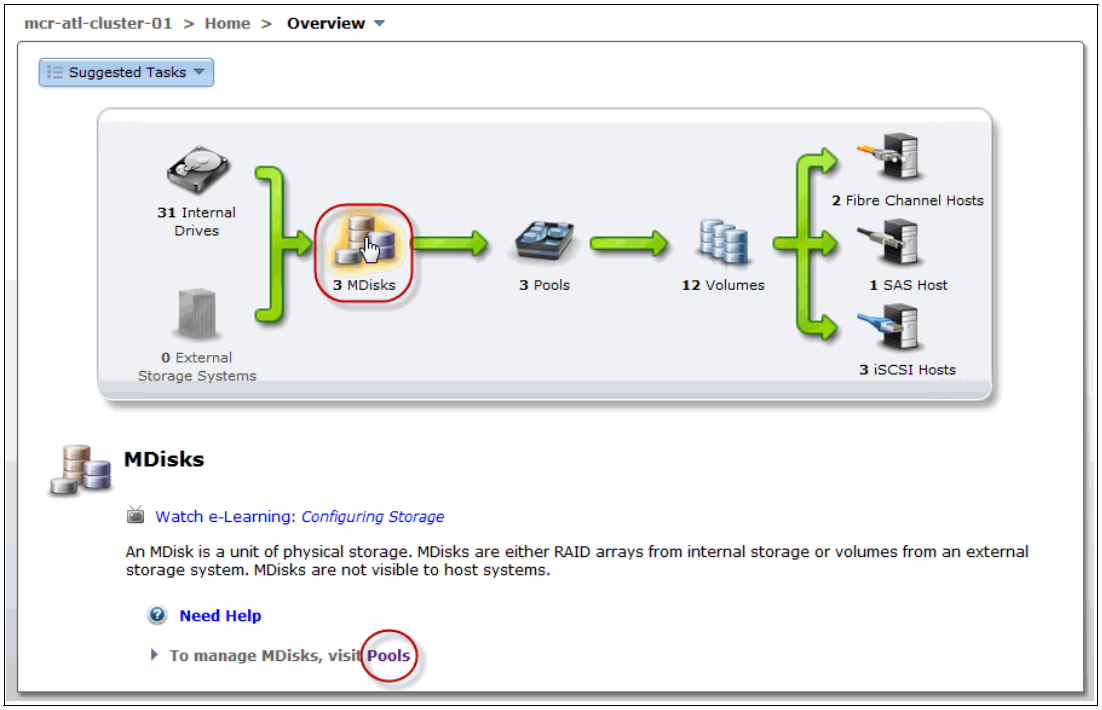

Figure 7-37 MDisk from Overview window

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.

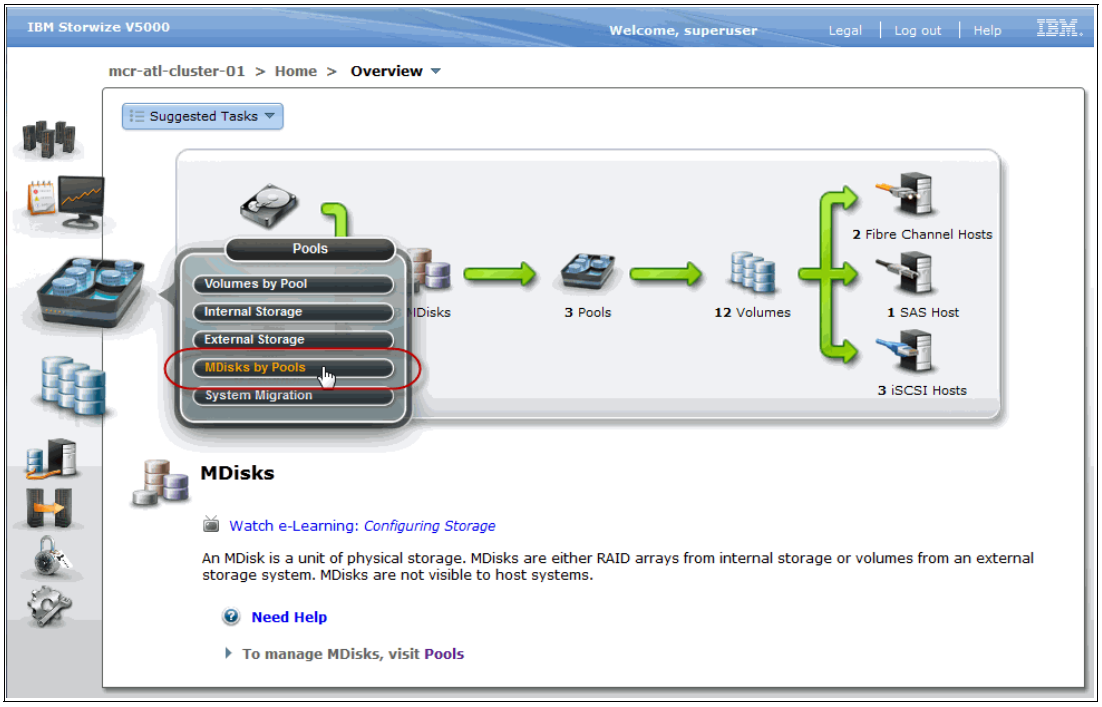

Figure 7-38 MDisk from Pools icon

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.

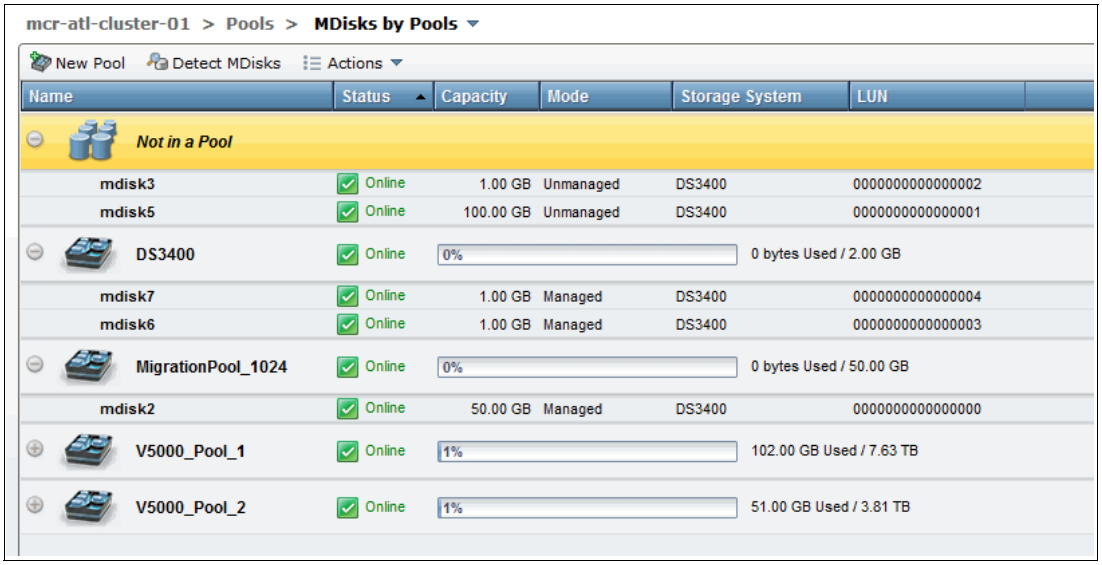

Figure 7-39 MDisks by Pools window

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.

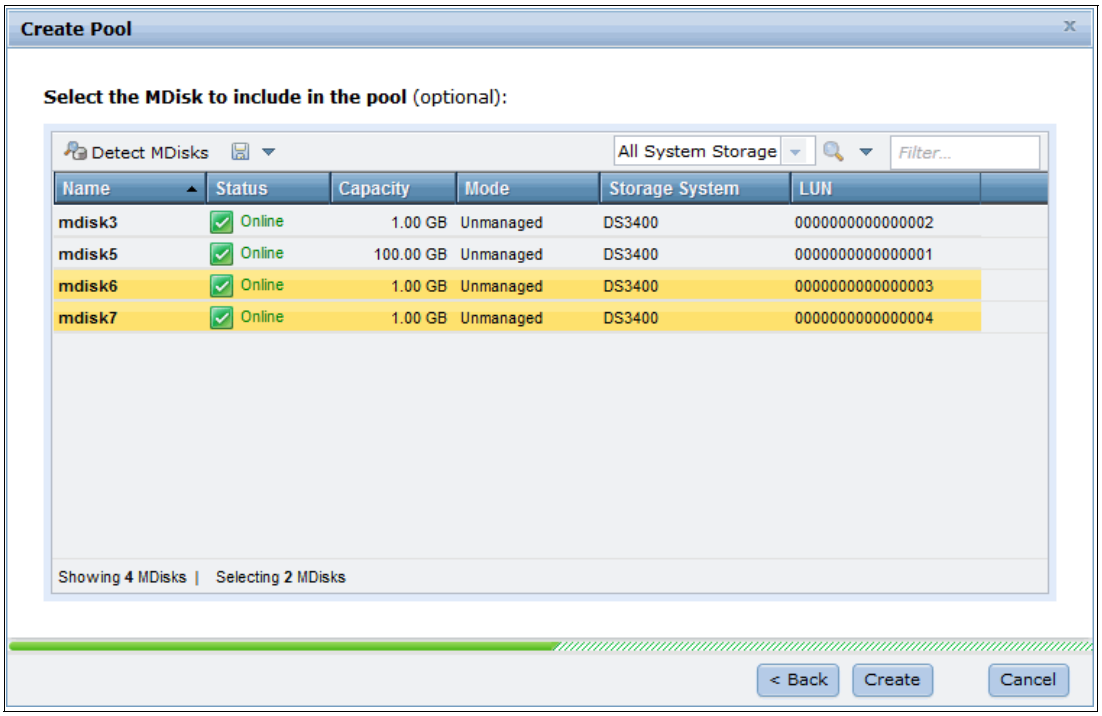

Figure 7-41 Create Pool: Part 2

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.

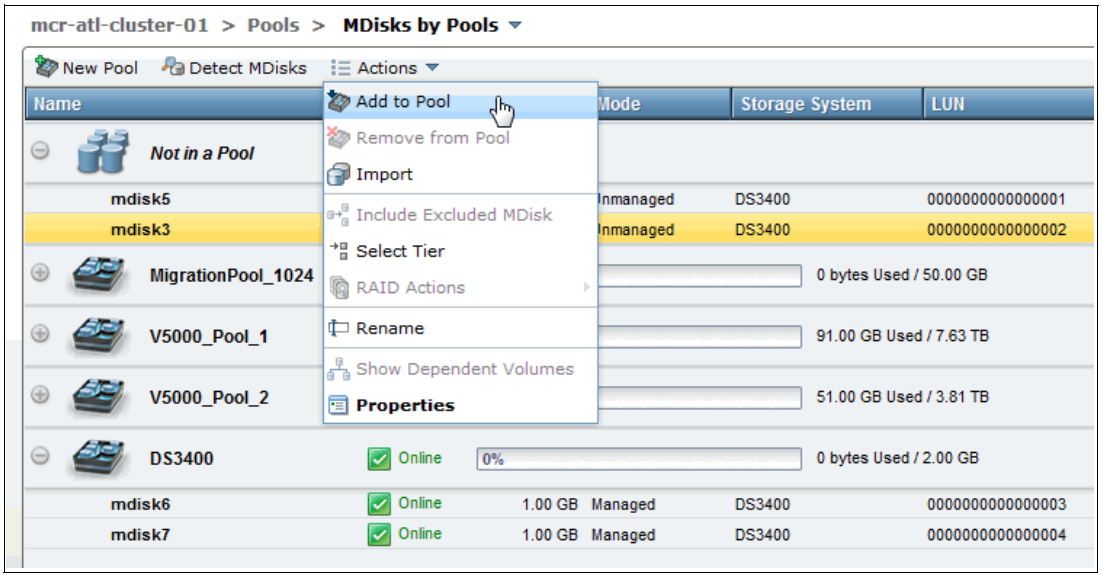

Figure 7-42 Add an unmanaged MDisk to a storage pool

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.

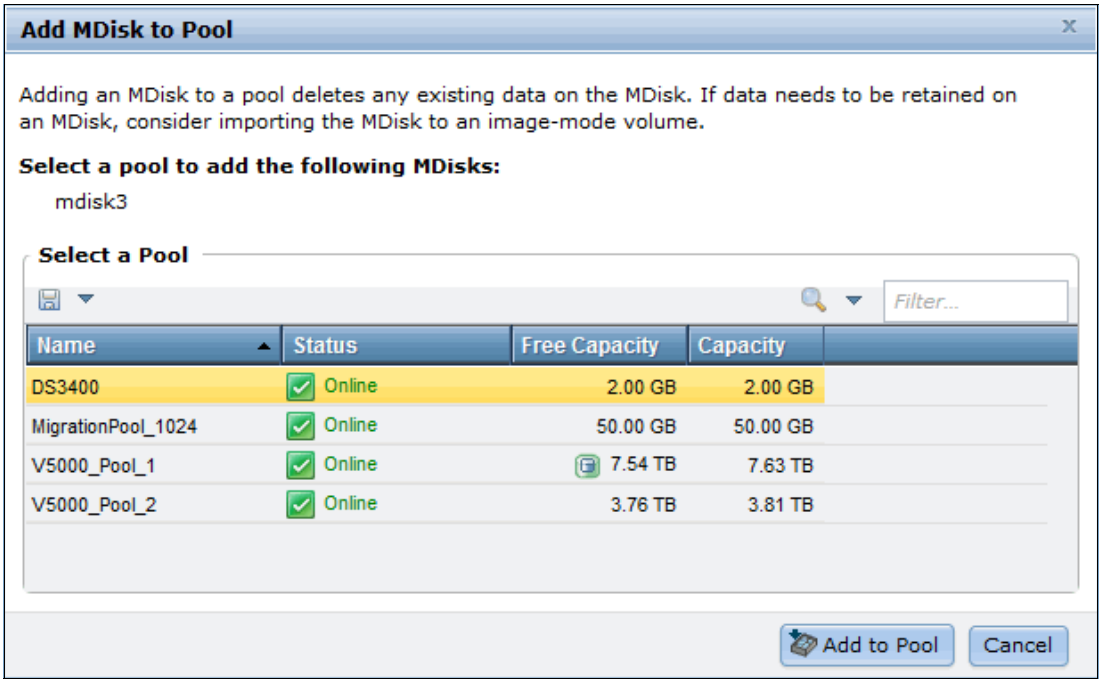

Figure 7-43 Add MDisk to pool

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.

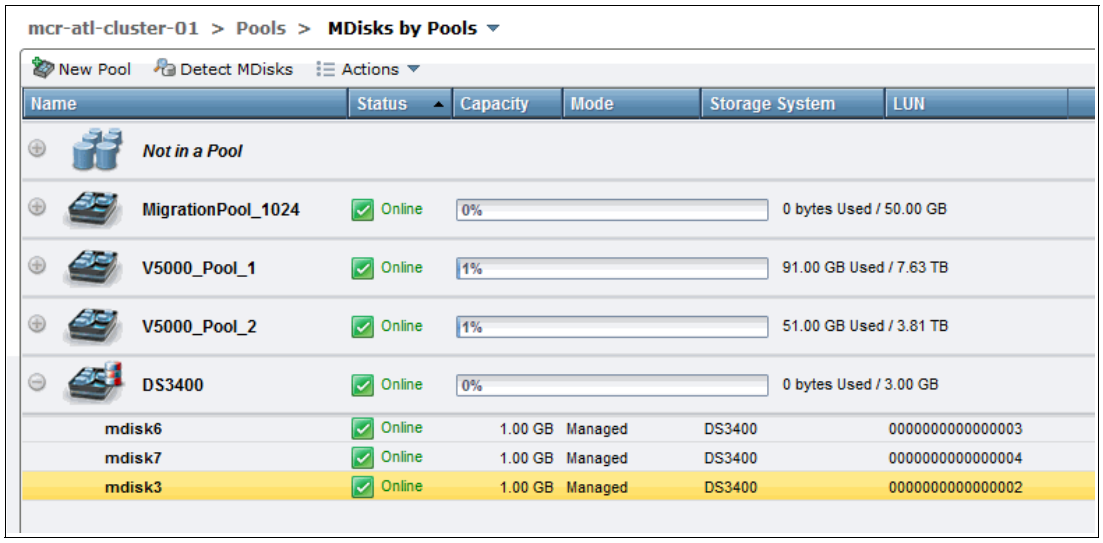

Figure 7-44 MDisk added to pool

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.

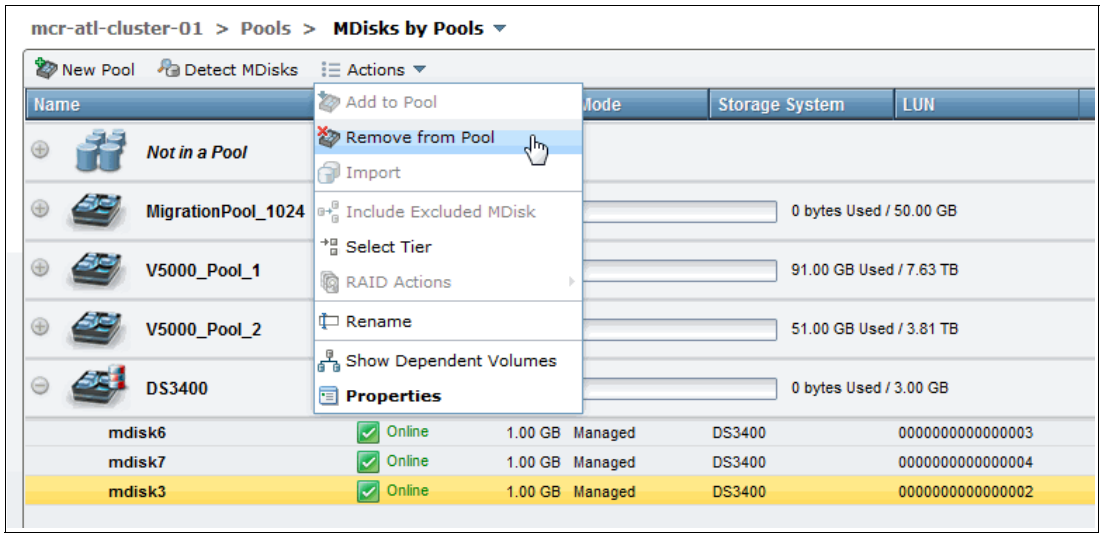

Figure 7-45 Remove an MDisk from the storage pool

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.

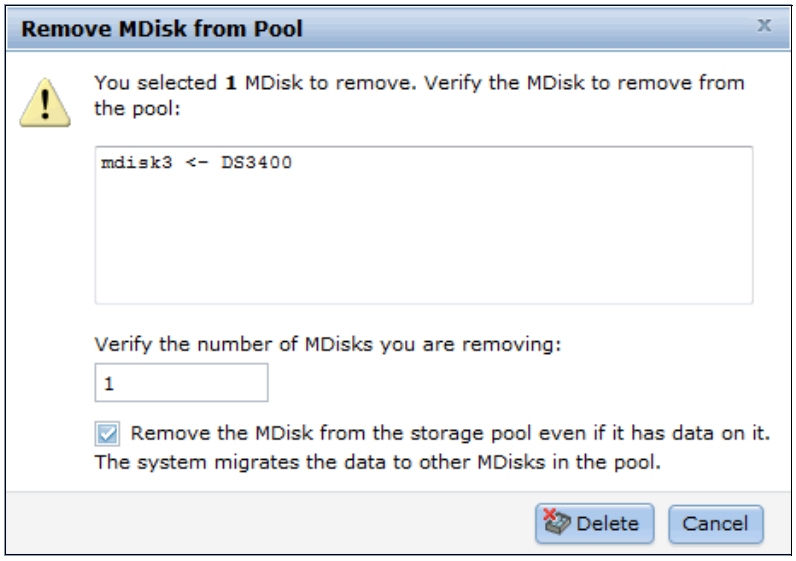

Figure 7-46 Confirm the removal of MDisk from the pool

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.

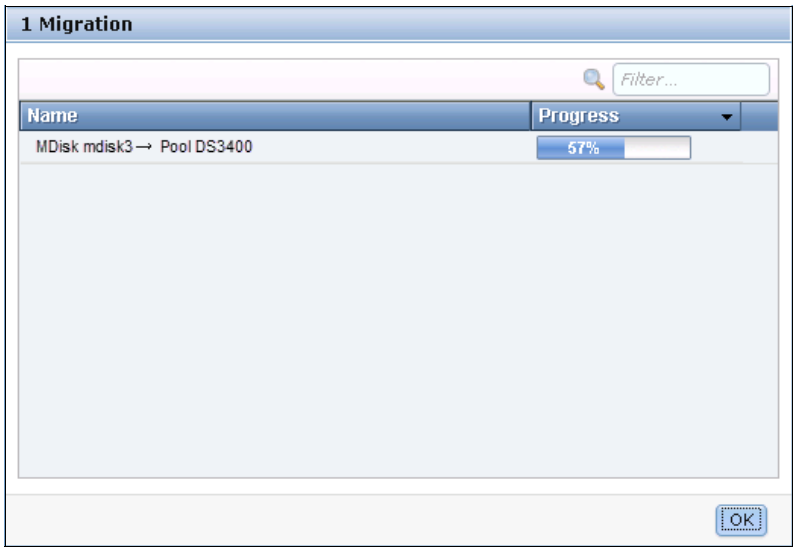

Figure 7-47 Data migration progress when MDisks are removed from the pool

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.

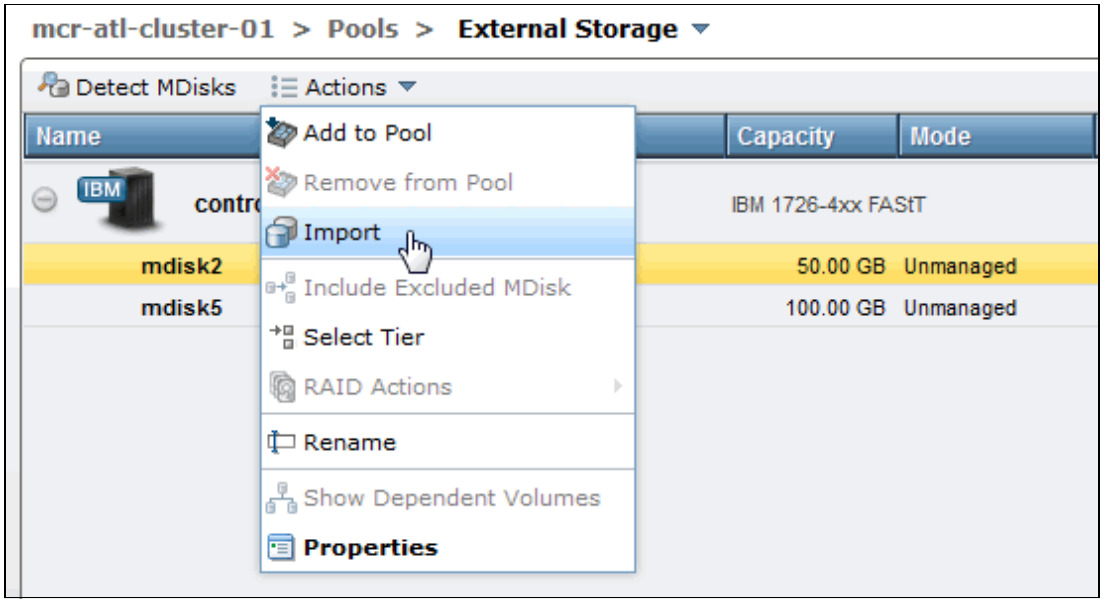

Figure 7-48 Import MDisk

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

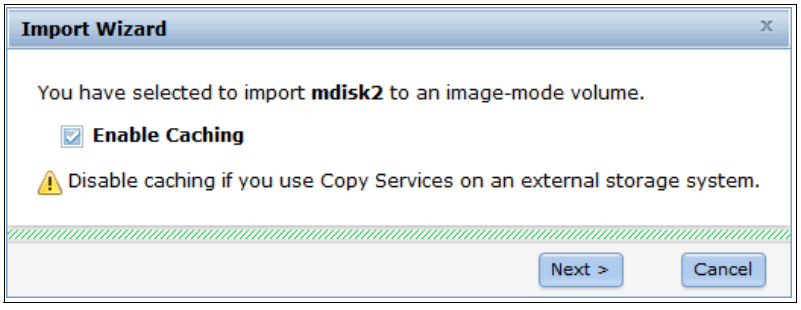

Figure 7-49 Import wizard: Step 1

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.

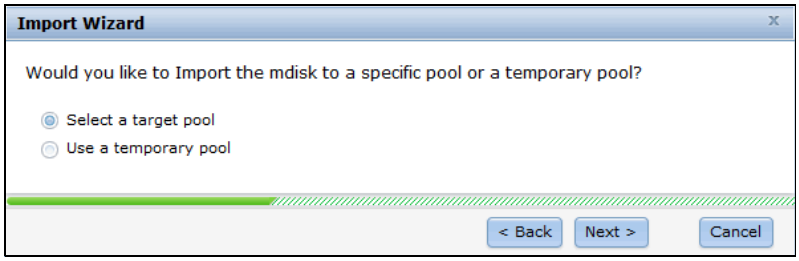

Figure 7-50 Import wizard: Step 2

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.

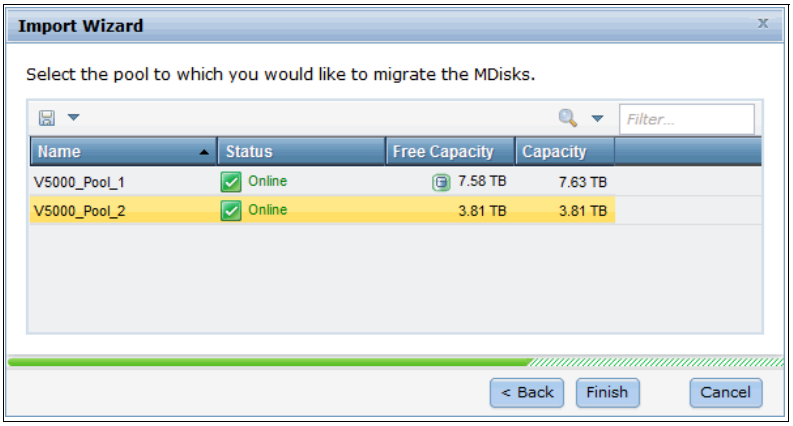

Figure 7-51 Import wizard: Step 3

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

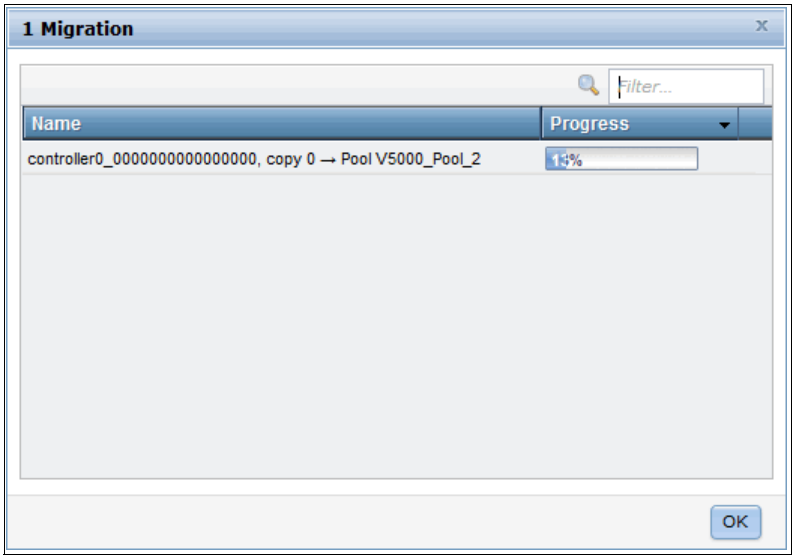

Figure 7-52 Migration progress in the status indicator of Running Tasks

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

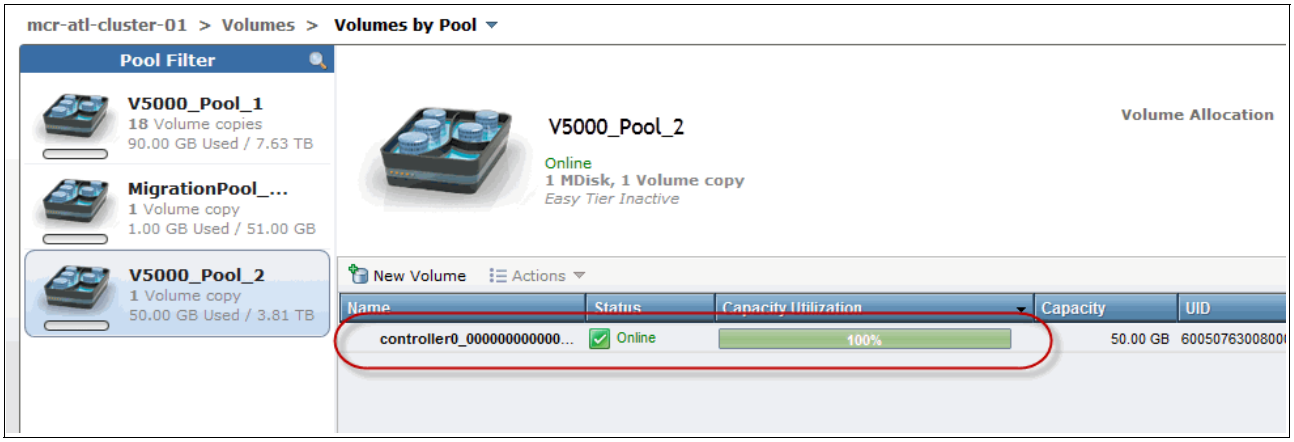

Figure 7-53 Volume migrated to destination pool

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.

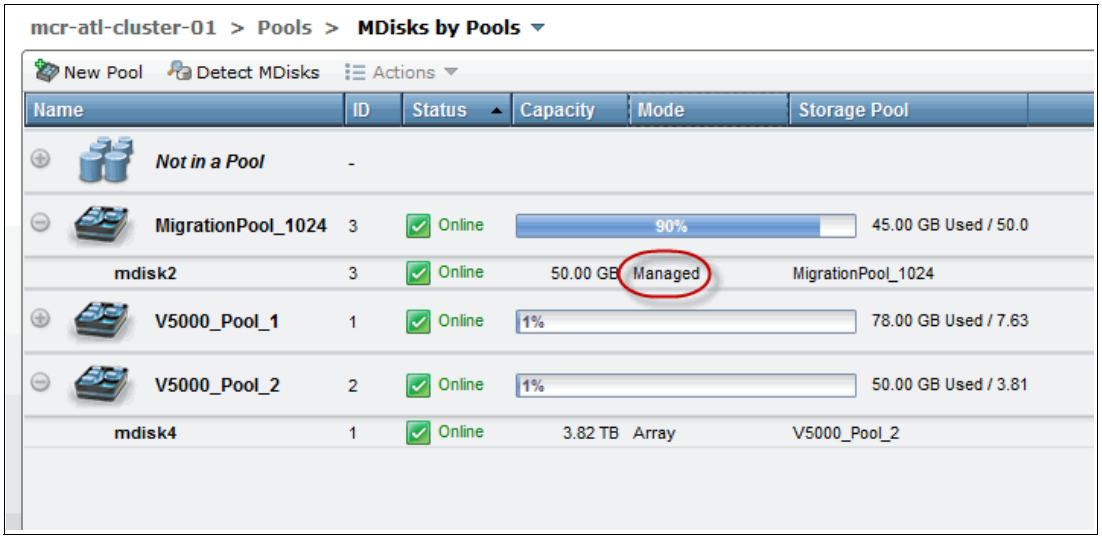

Figure 7-54 MDisk mode that is changed to managed

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.

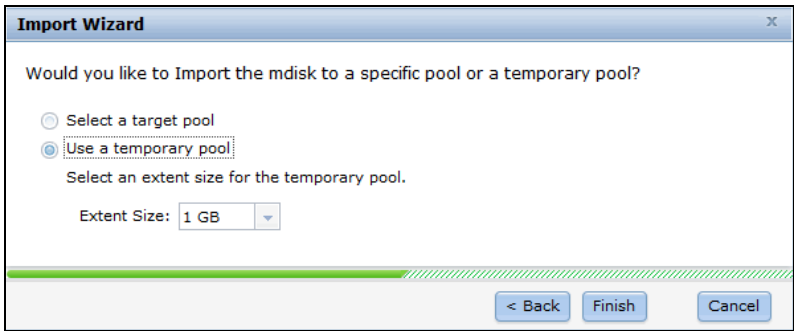

Figure 7-55 Import MDisk to a temporary pool

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

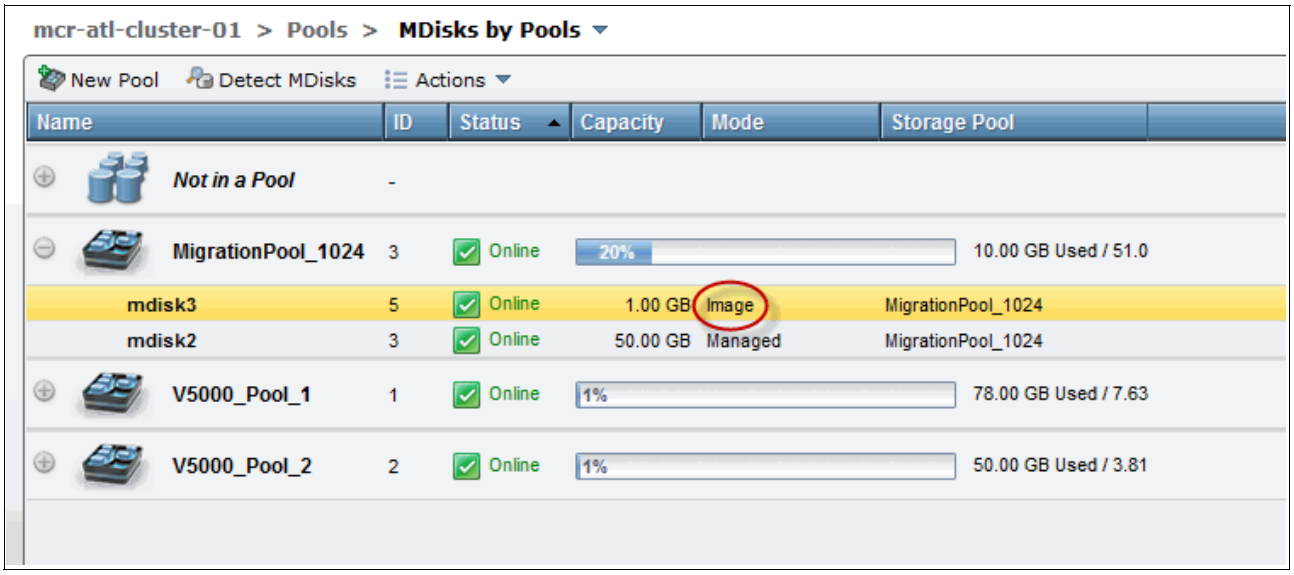

Figure 7-56 MDisk after import

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.

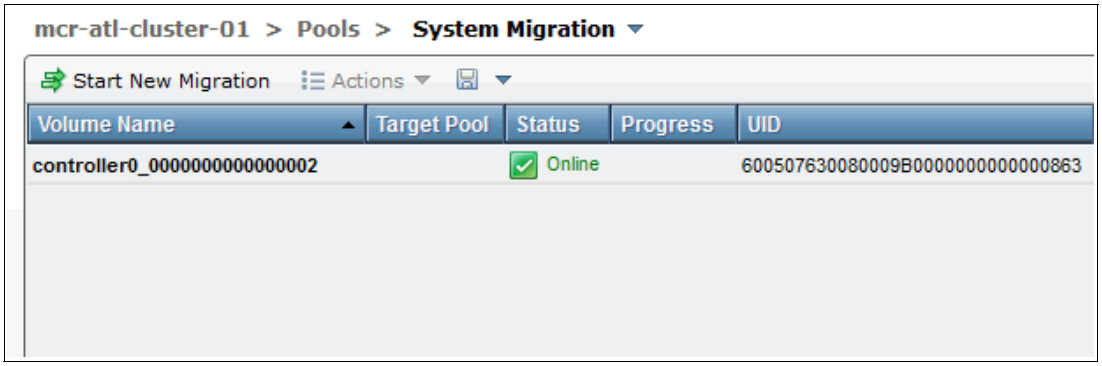

Figure 7-57 Imported MDisk in the System Migration window

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.

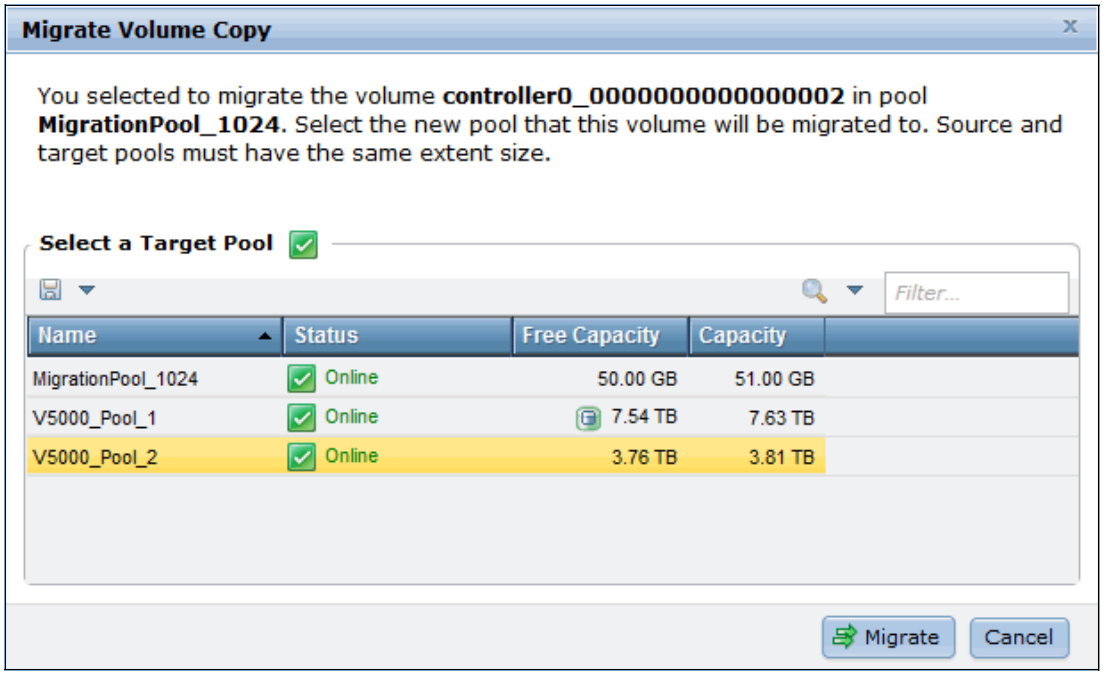

Figure 7-58 Migrate Image Mode Volume into a regular storage pool

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.

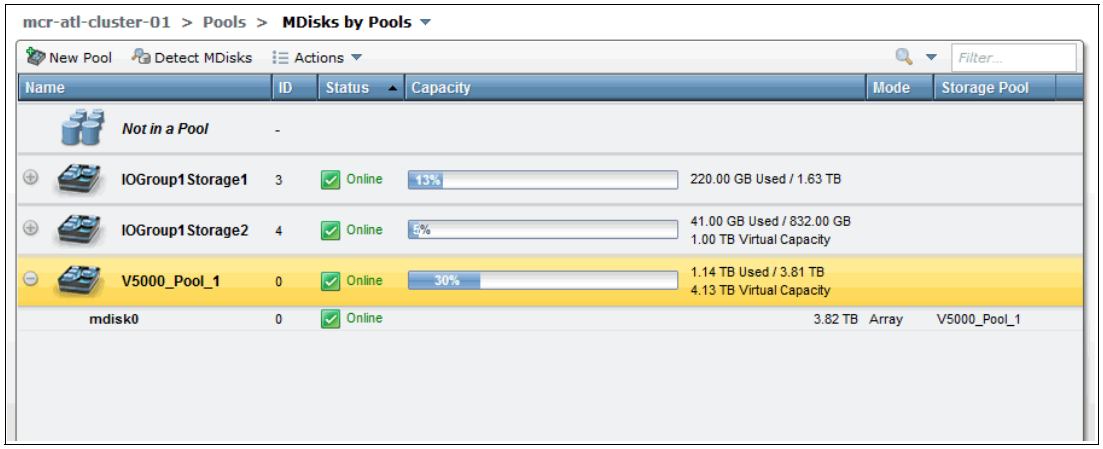

Figure 7-59 MDisk by Pool window

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.

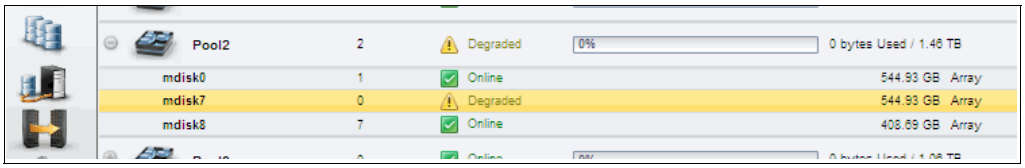

Figure 7-60 One degraded MDisk in pool

  – Offline

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

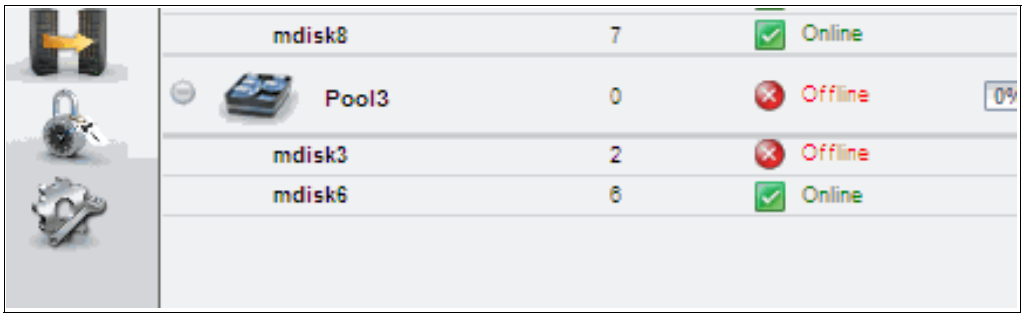

Figure 7-61 Offline MDisk in a pool

- 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.

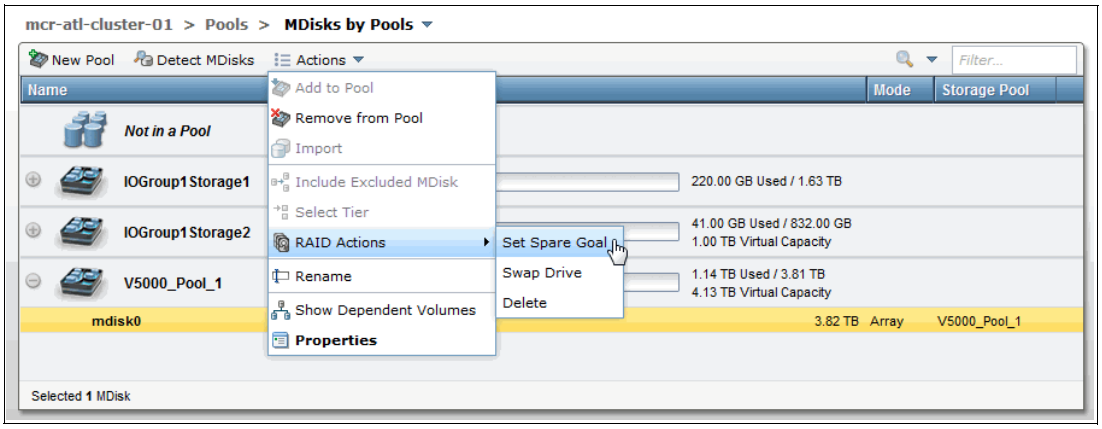

Figure 7-62 MDisk RAID actions

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.

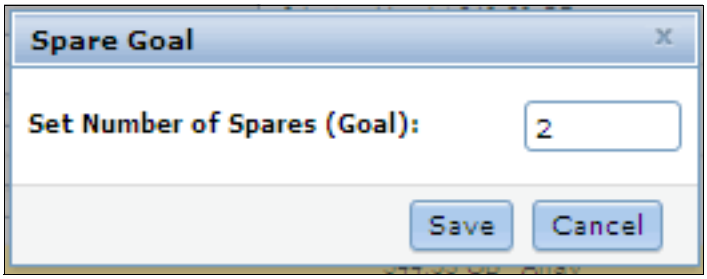

Figure 7-63 MDisk set spare goal

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**.

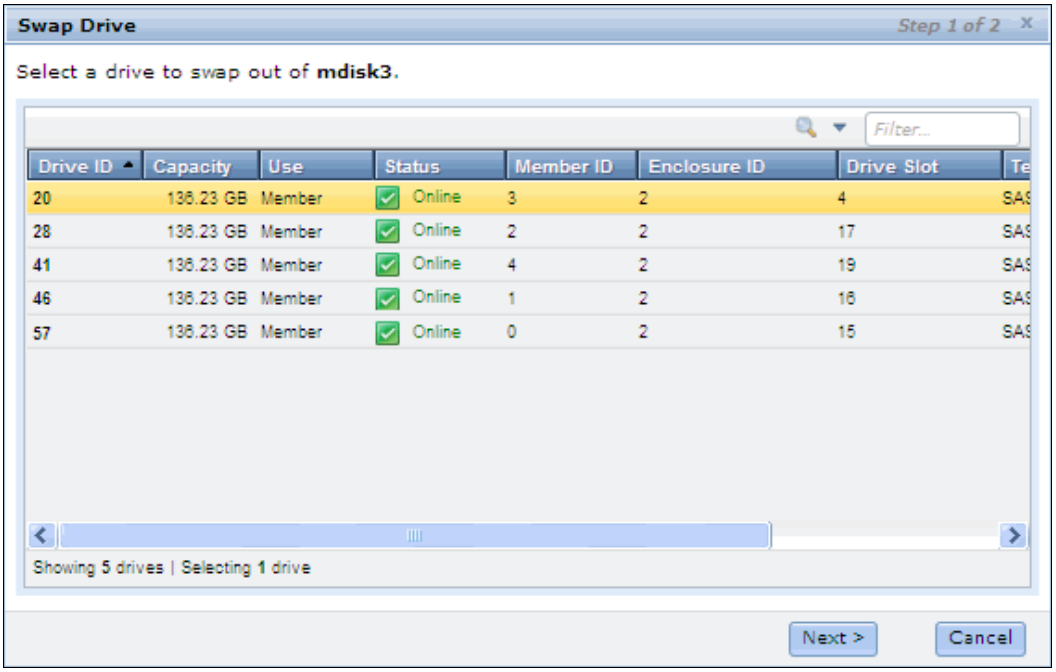

Figure 7-64 MDisk swap drive: Step 1

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**.

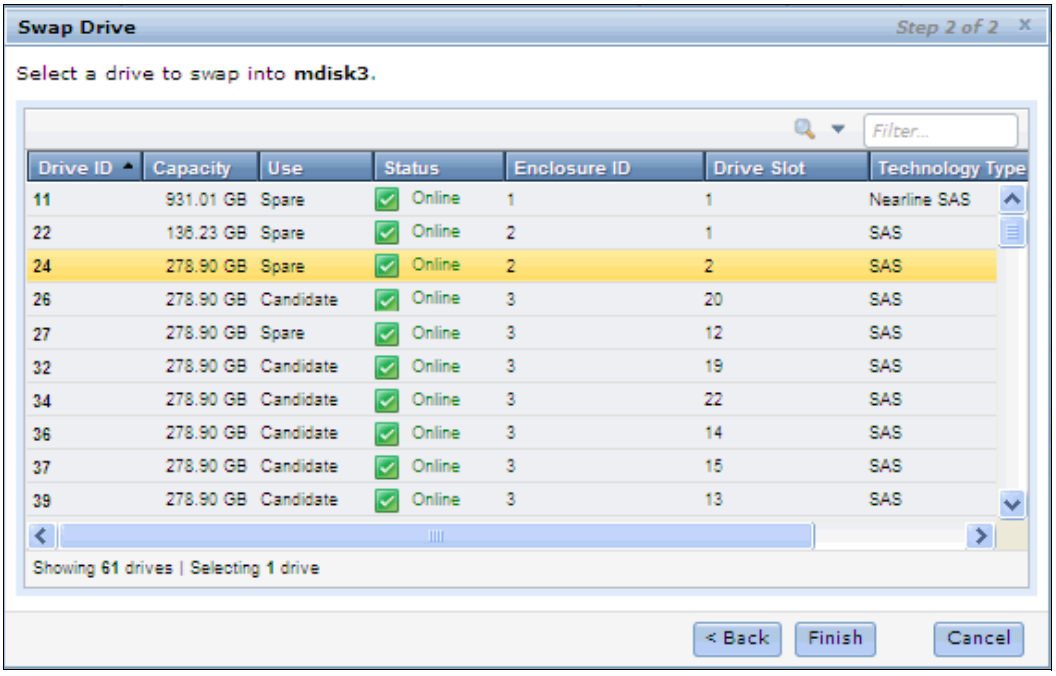

Figure 7-65 MDisk swap drive: Step 2

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.

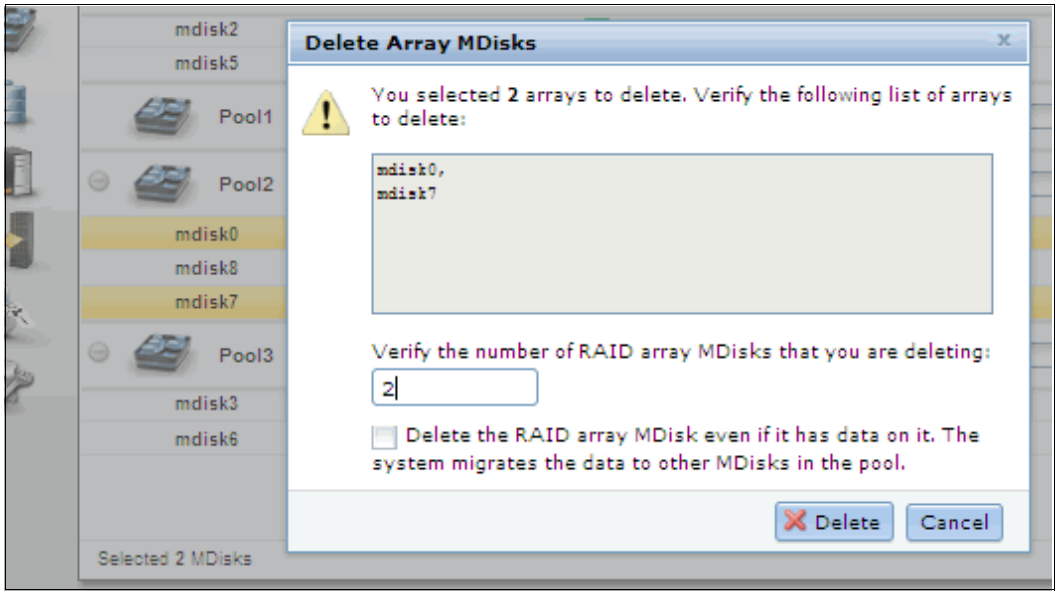

Figure 7-66 MDisk delete confirmation

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.

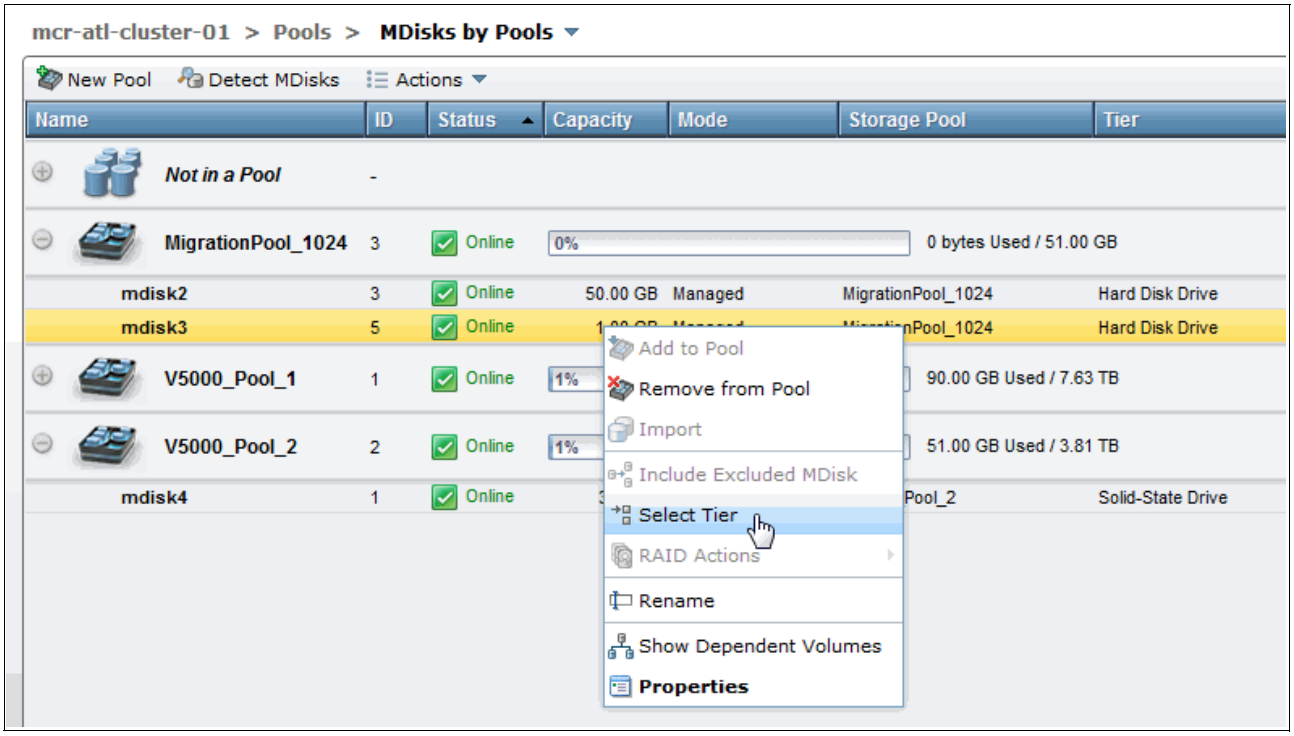

Figure 7-68 Select Tier for an MDisk

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.

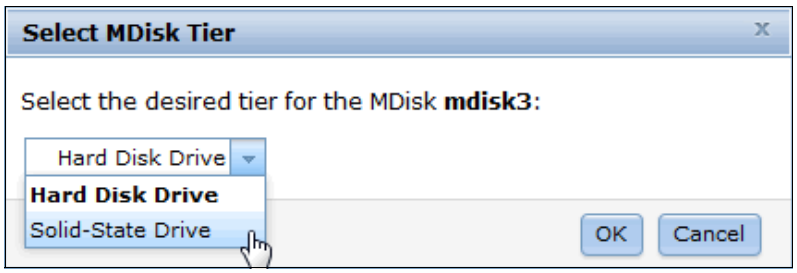

Figure 7-69 Assign wanted tier to an MDisk

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

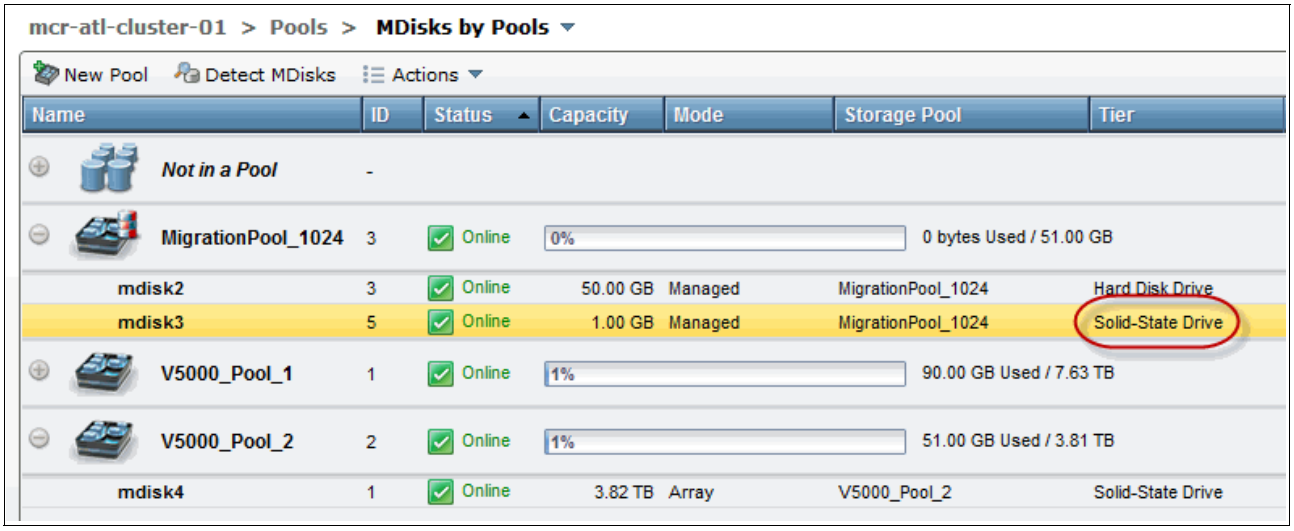

Figure 7-70 Wanted tier that is assigned to the MDisk

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.

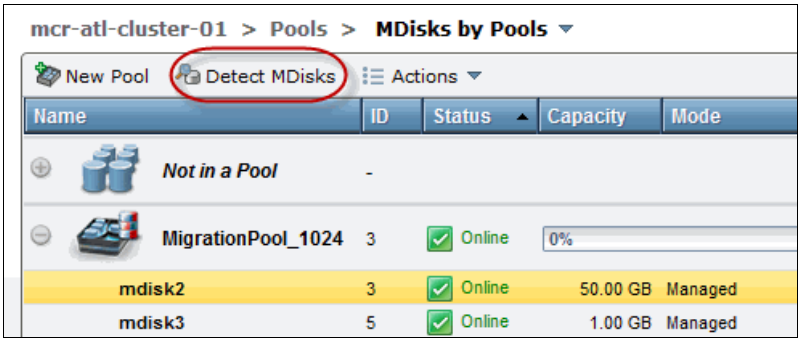

Figure 7-71 Detect MDisks

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.

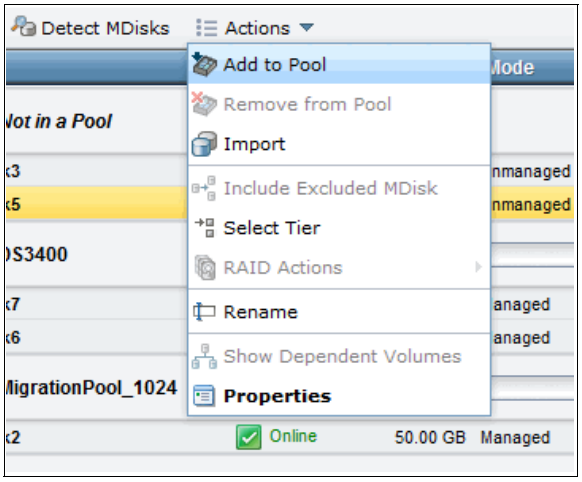

Figure 7-72 MDisk actions on externally virtualized storage

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**.

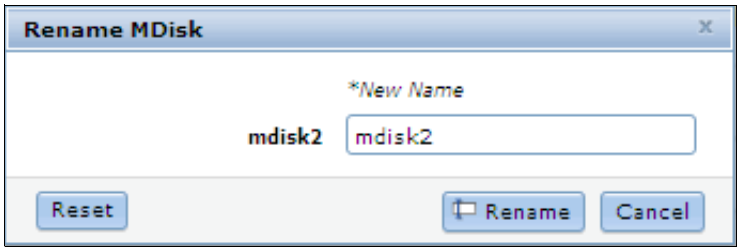

Figure 7-73 Rename MDisk

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.

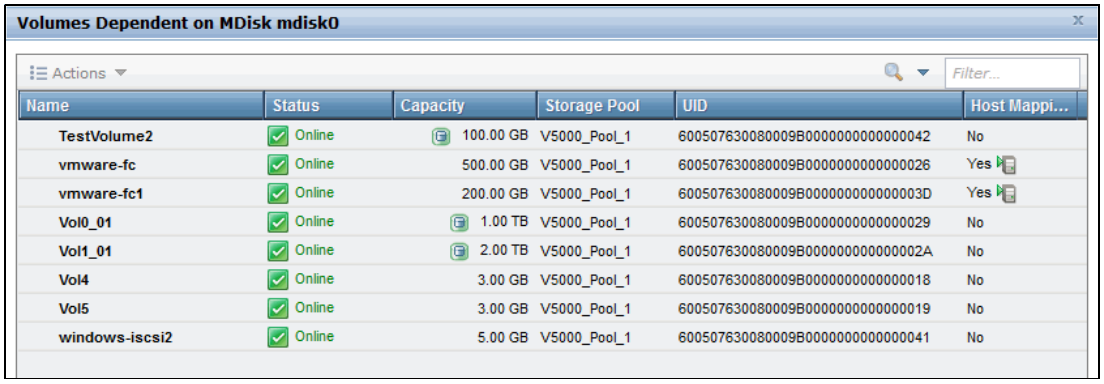

Figure 7-74 Show dependent volumes

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**.

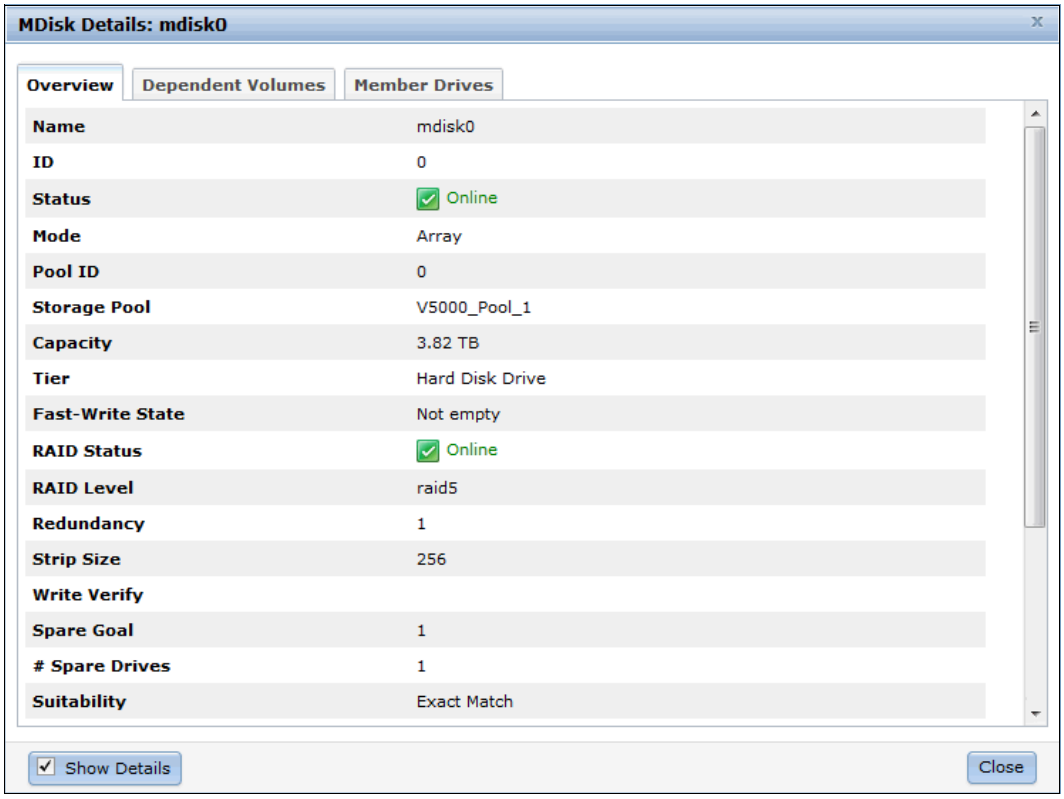

Figure 7-75 MDisk properties overview

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

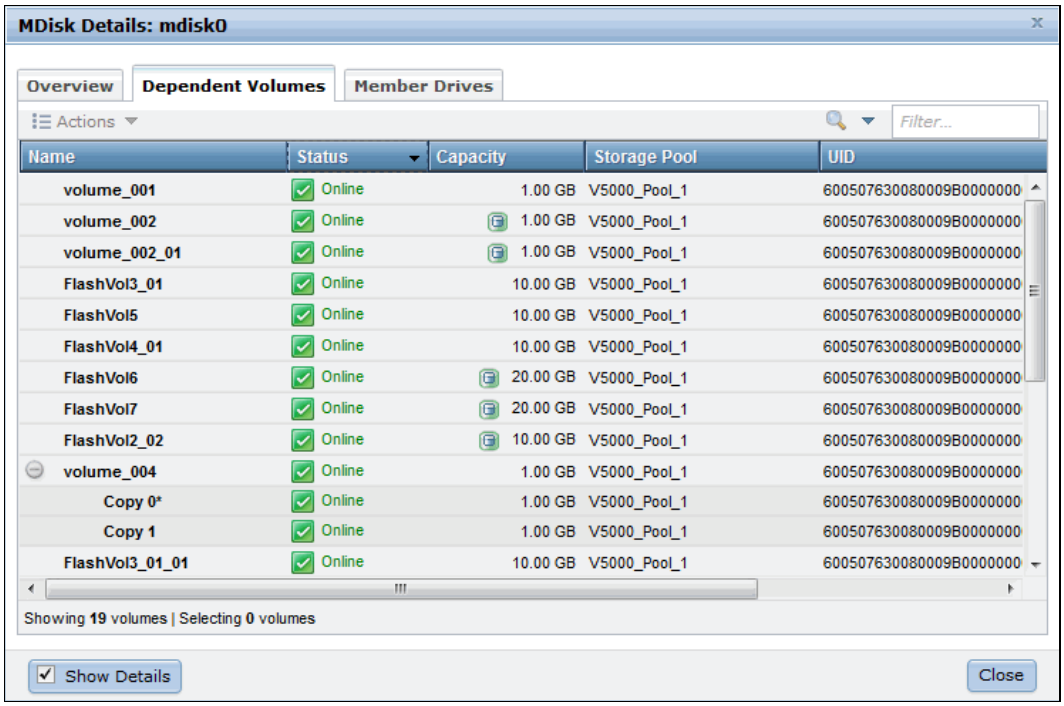

Figure 7-76 MDisk dependent volumes

- 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.

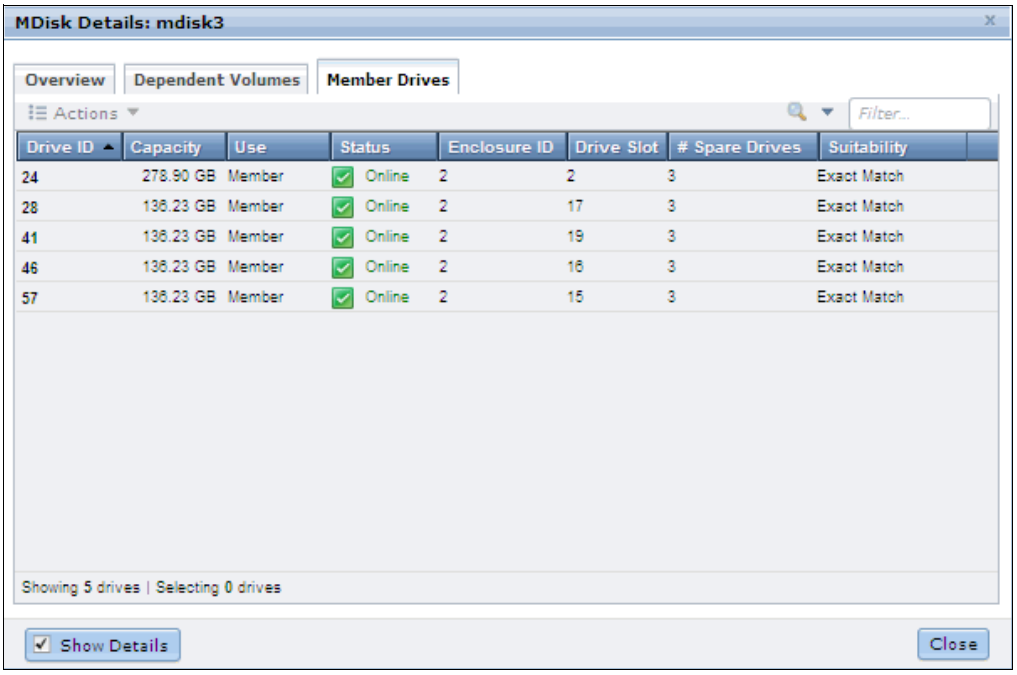

Figure 7-77 MDisk properties member

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.

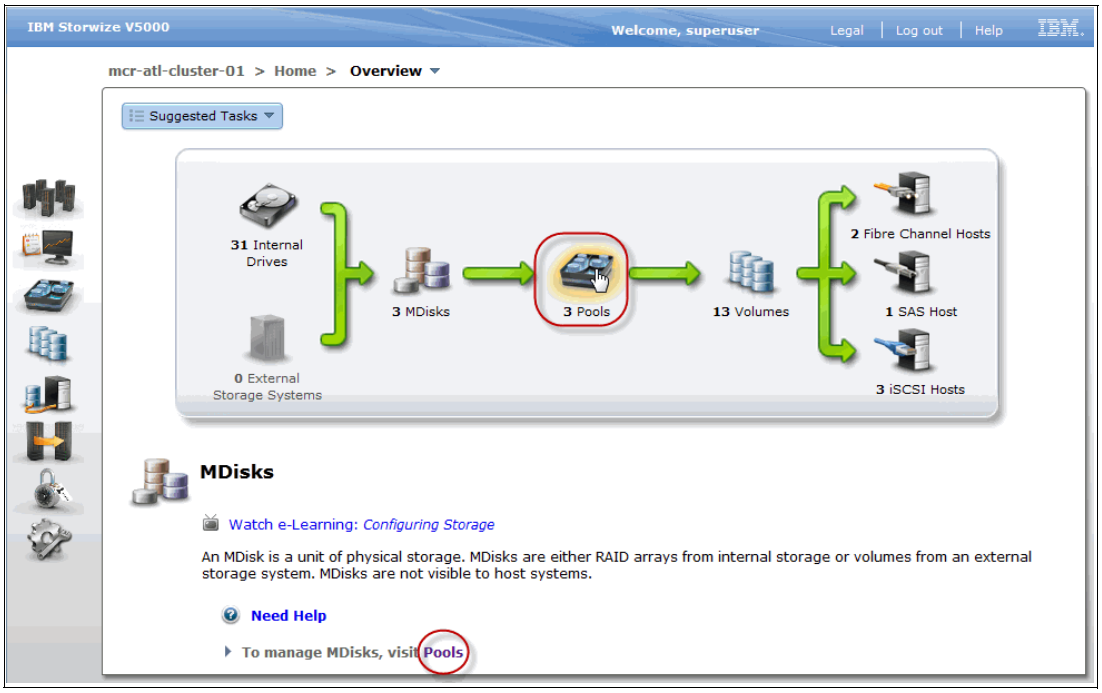

Figure 7-78 Pools from the overview window

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

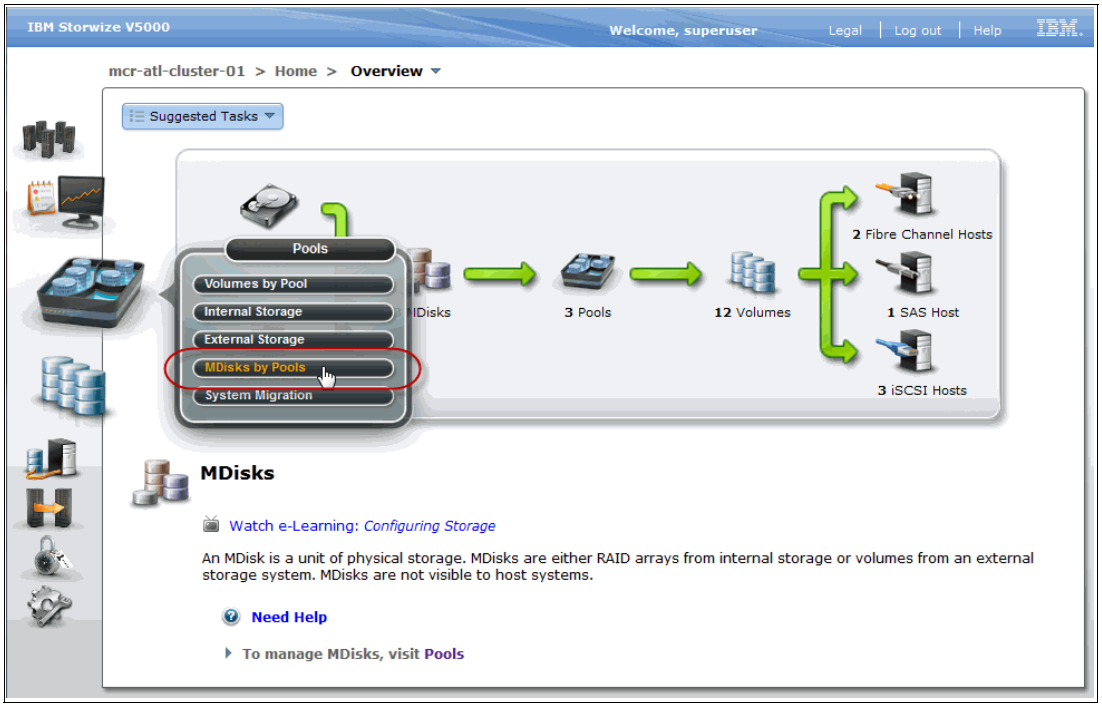

Figure 7-79 Pools from MDisk by Pools window

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.

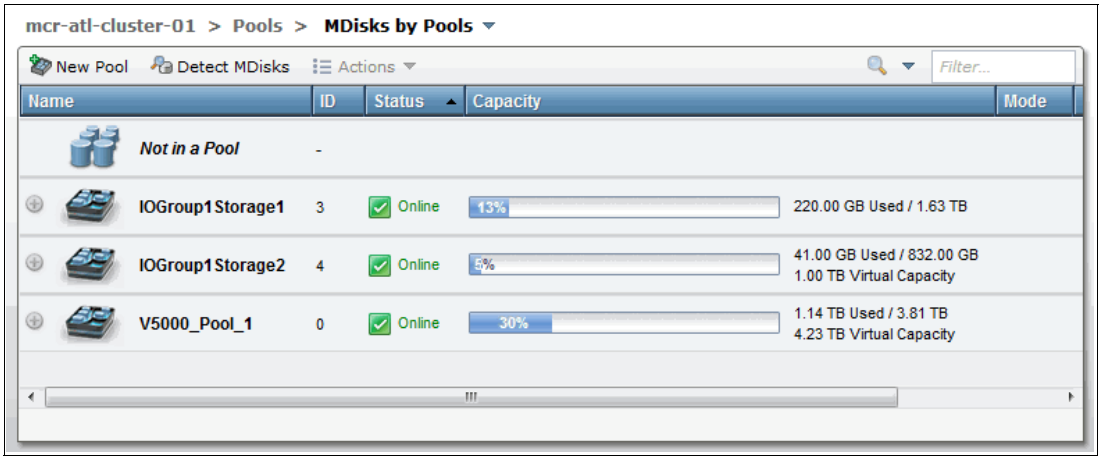

Figure 7-80 Pool window

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.

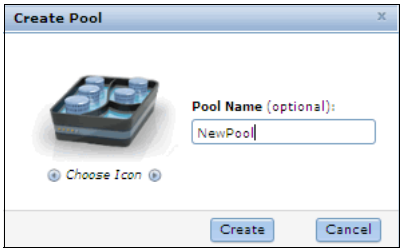

Figure 7-81 Create pool name input

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

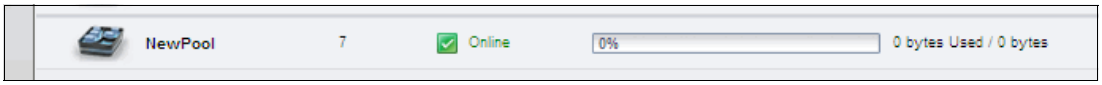

Figure 7-82 Empty pool that is created

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.

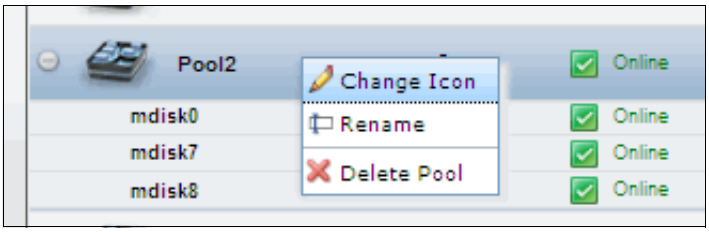

Figure 7-83 Pool action overview

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.

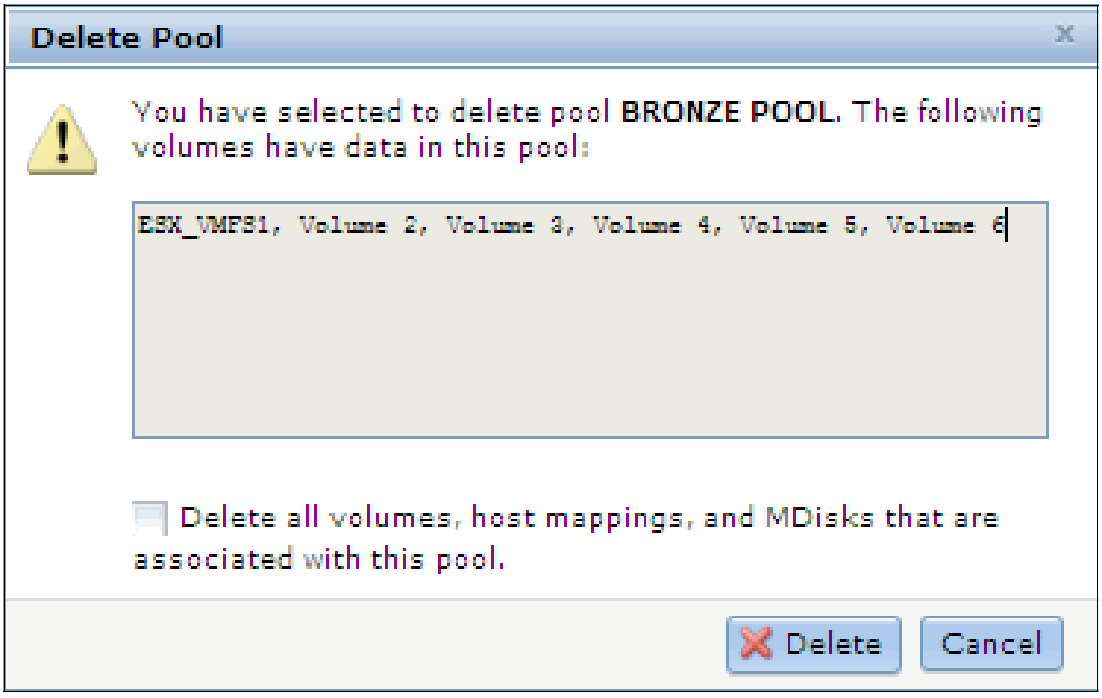

Figure 7-86 Confirmation to delete the storage pool

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.

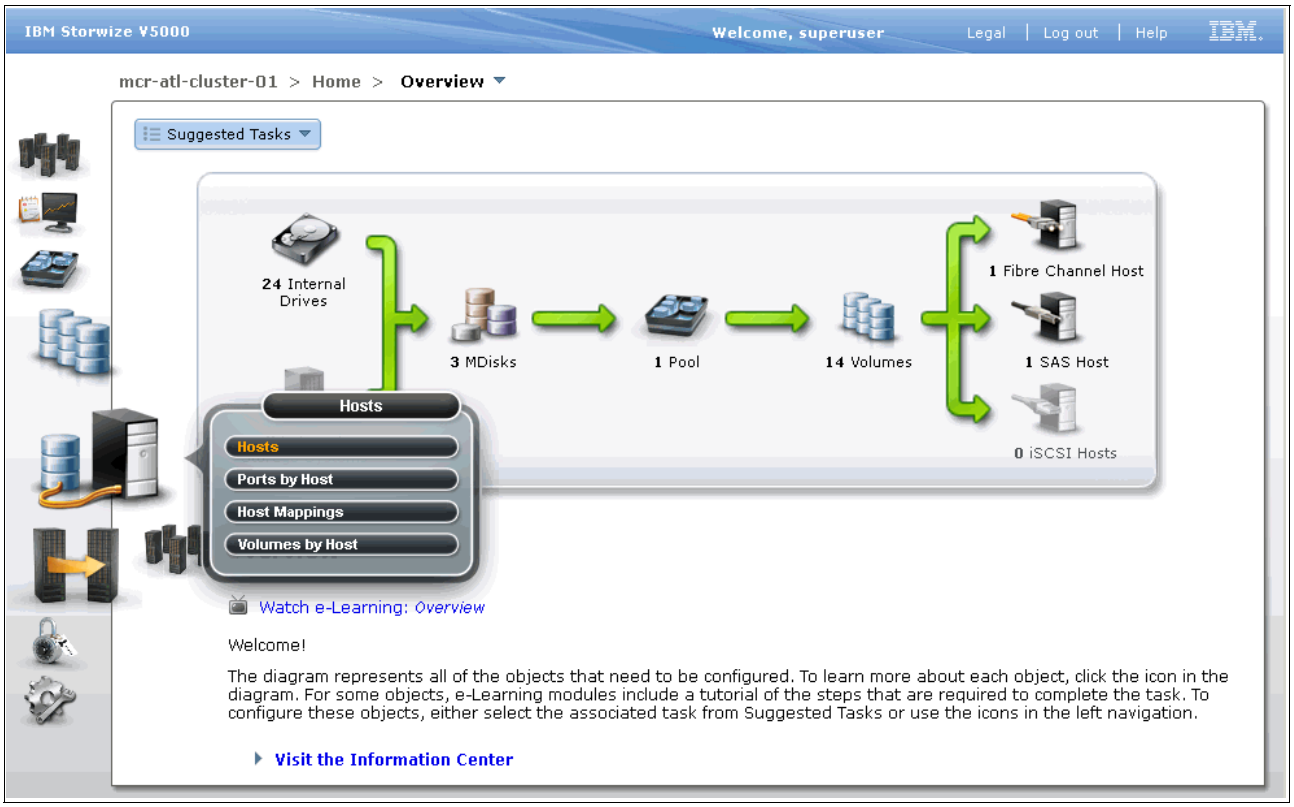

Figure 8-1 Host menu

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

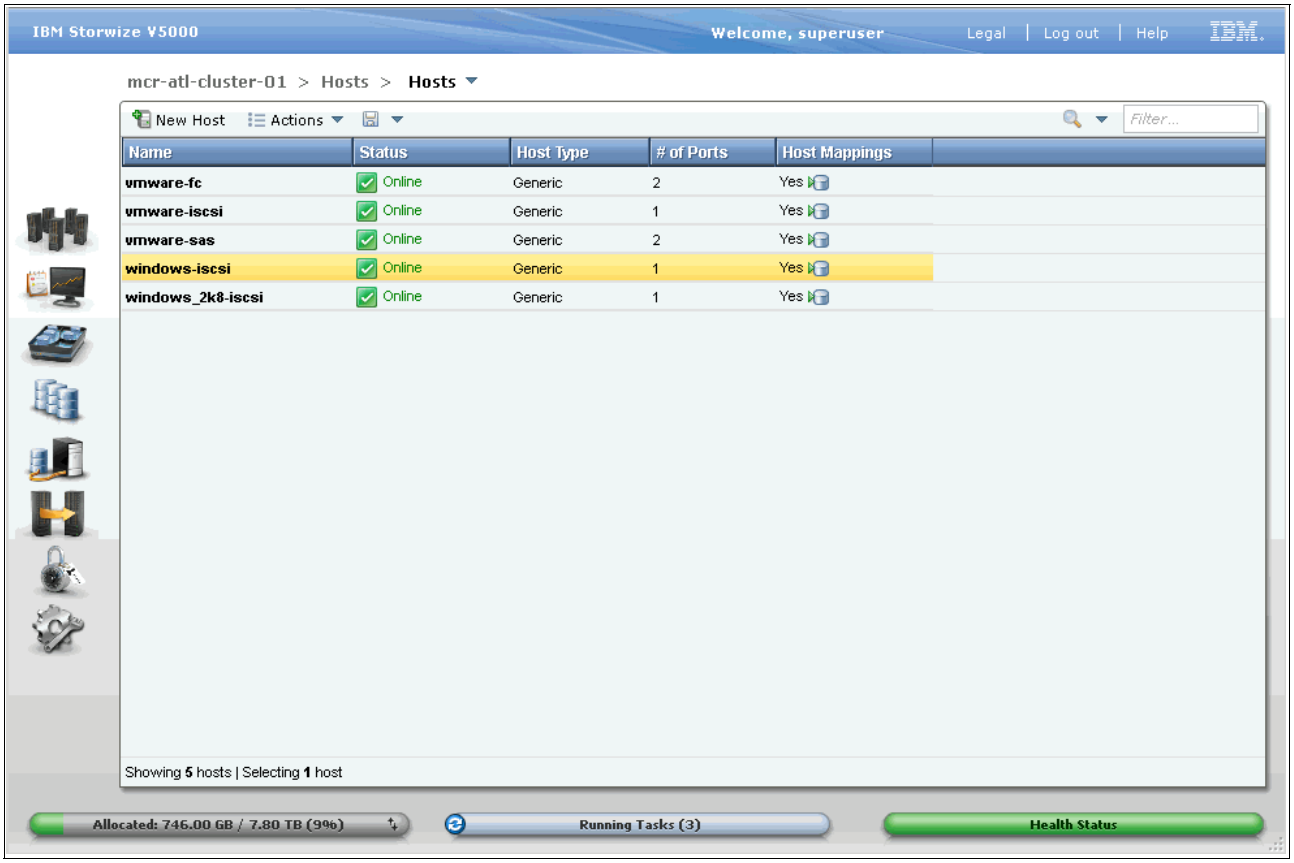

Figure 8-2 Hosts

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.

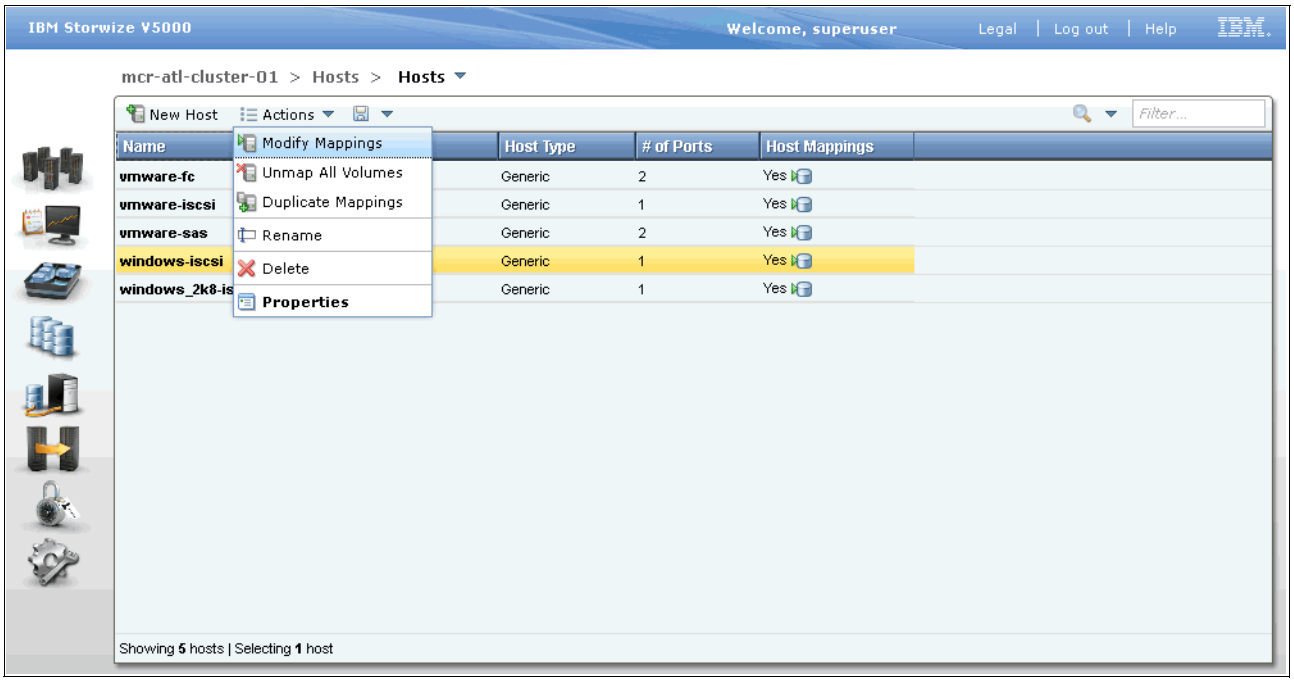

Figure 8-3 Host menu options

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.

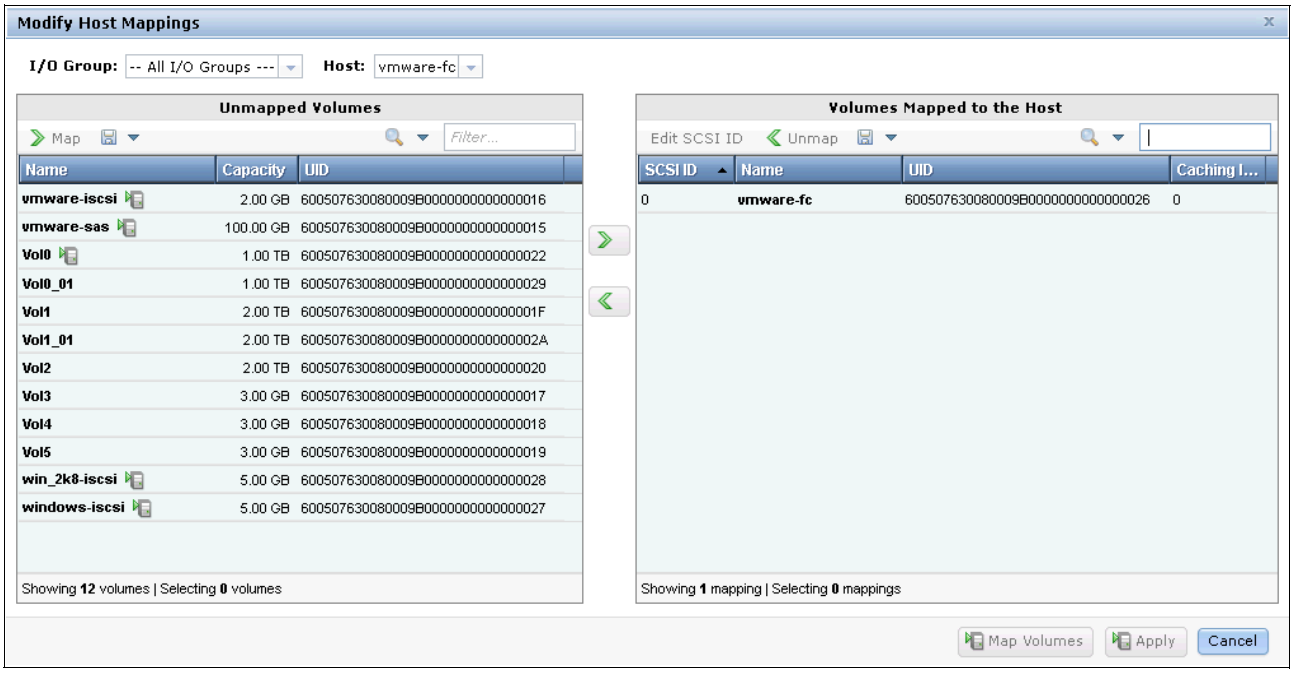

Figure 8-4 Host mappings window

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.

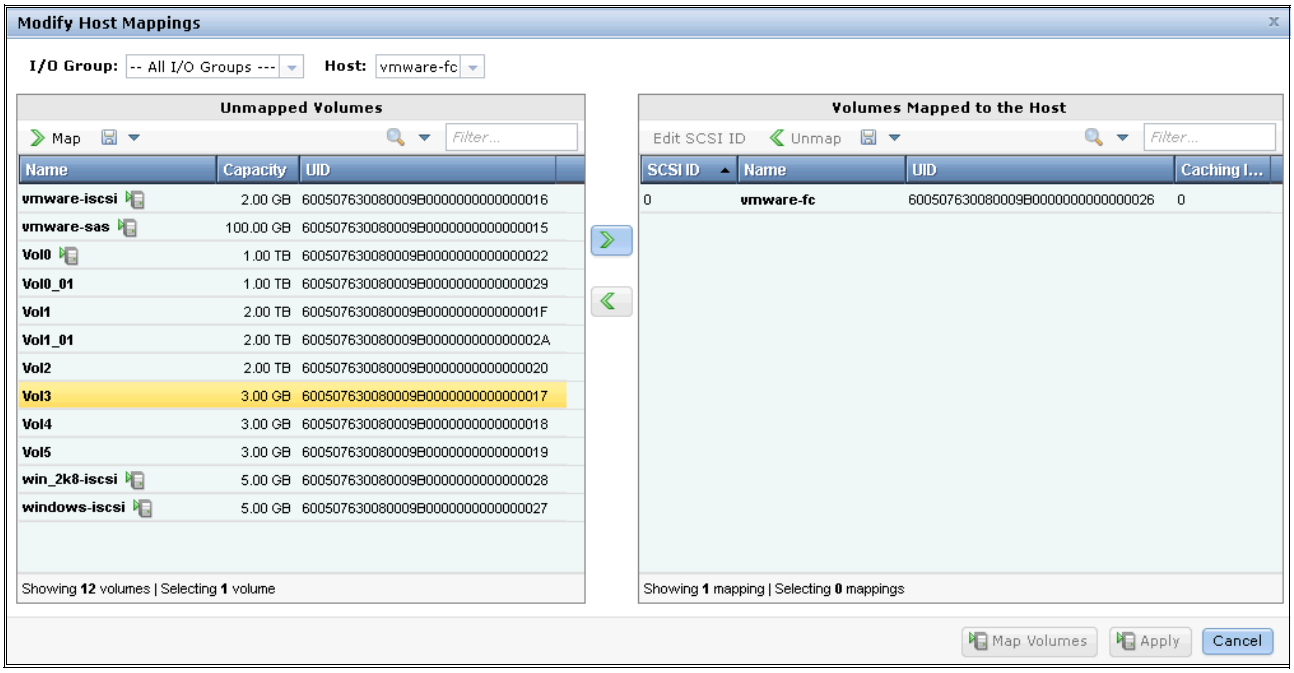

Figure 8-5 Modify Host Mappings

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.

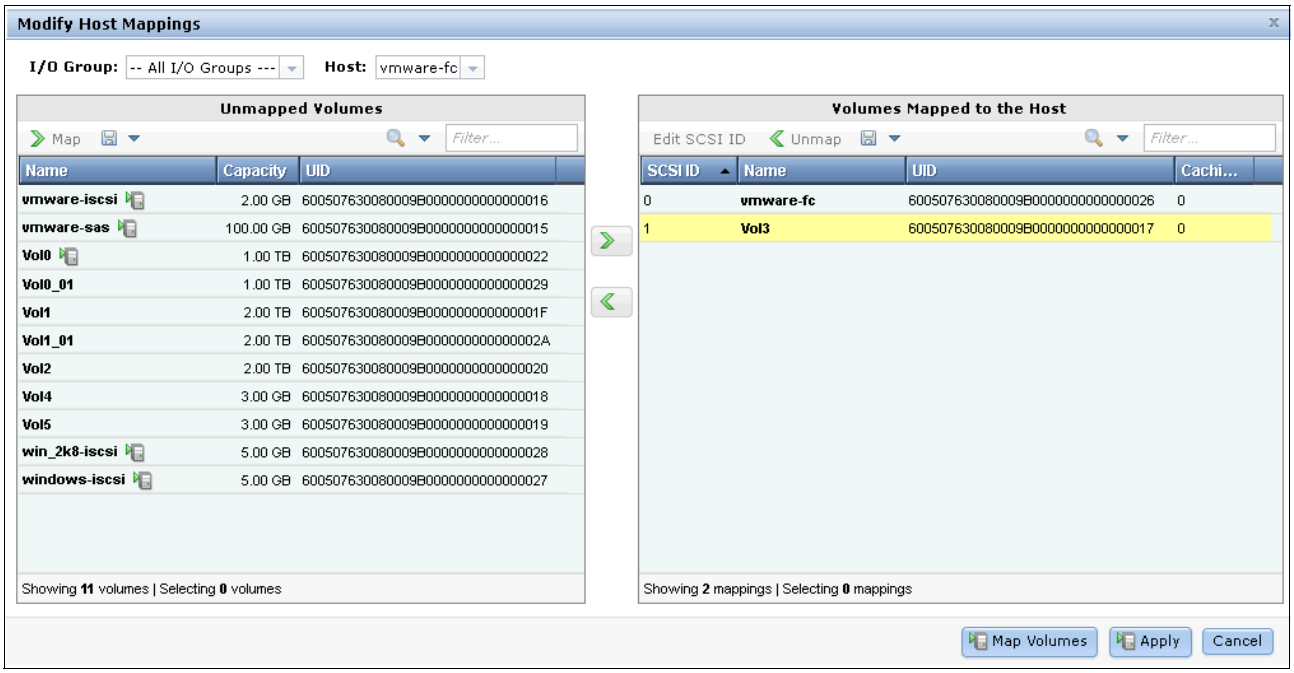

Figure 8-6 Modify Host Mappings

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.

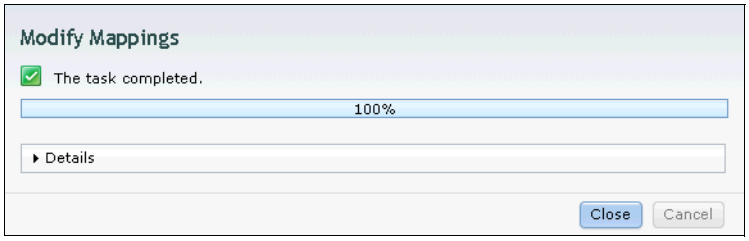

Figure 8-7 Modify Mappings task completed

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.

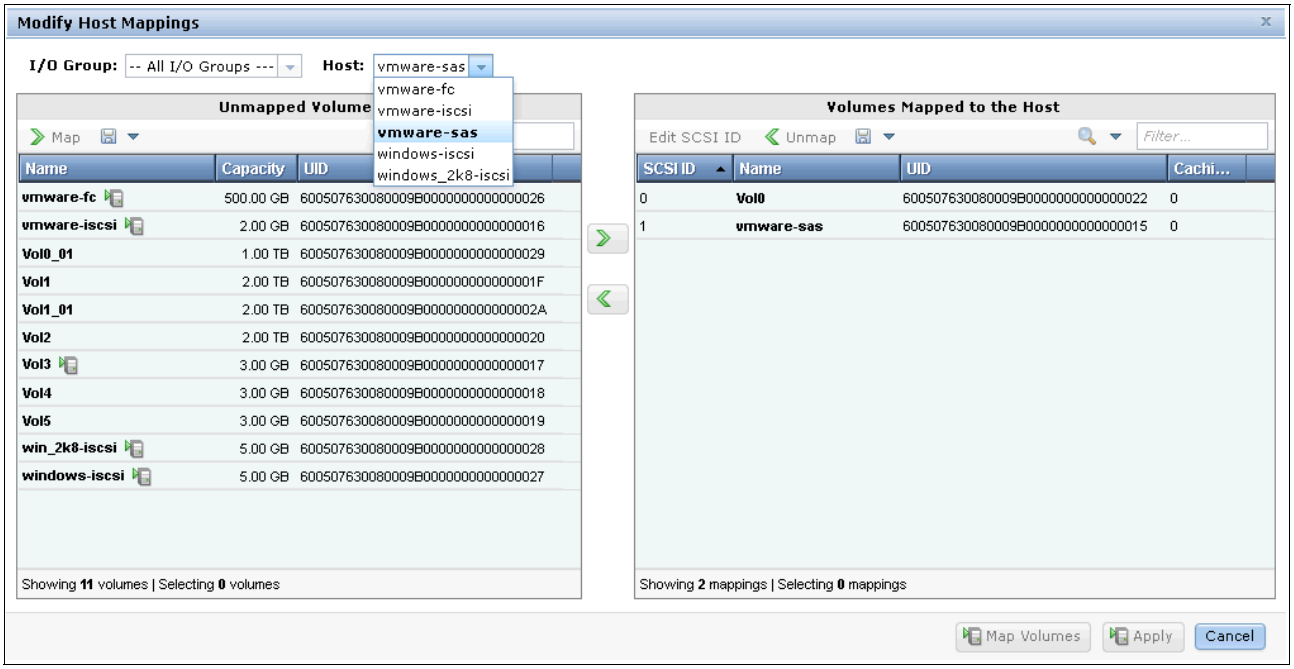

Figure 8-8 Selecting another host to modify

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.

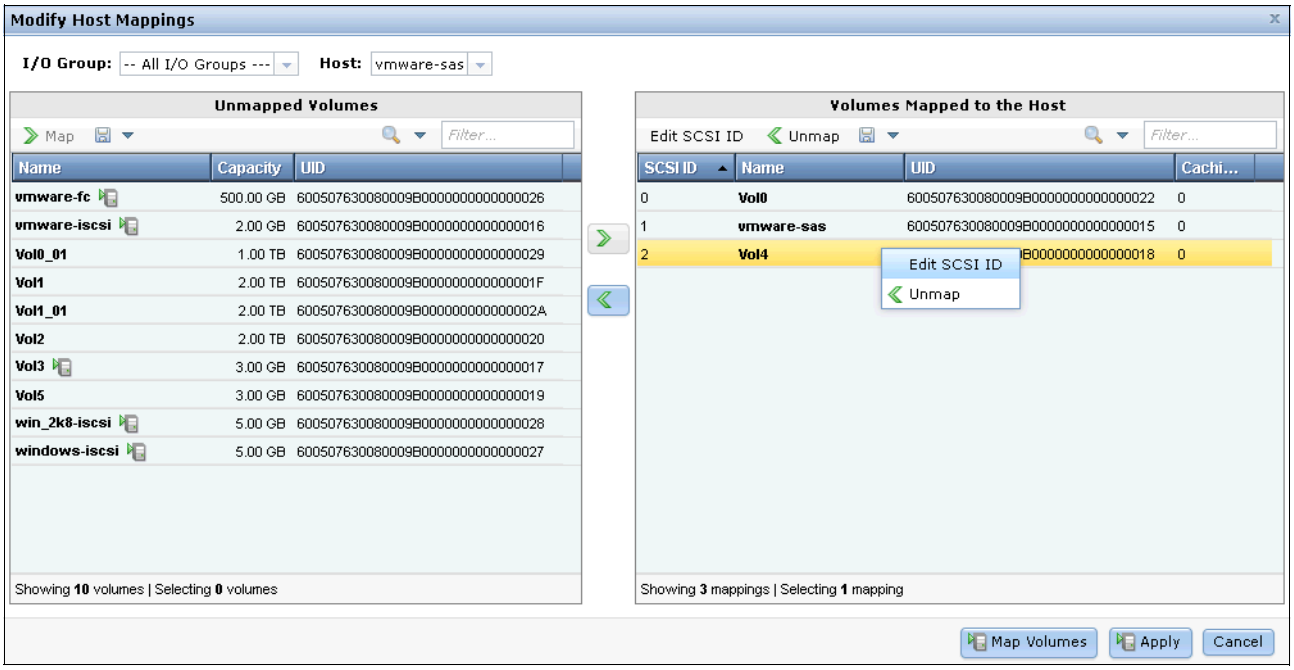

Figure 8-9 Editing iSCSI ID

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.

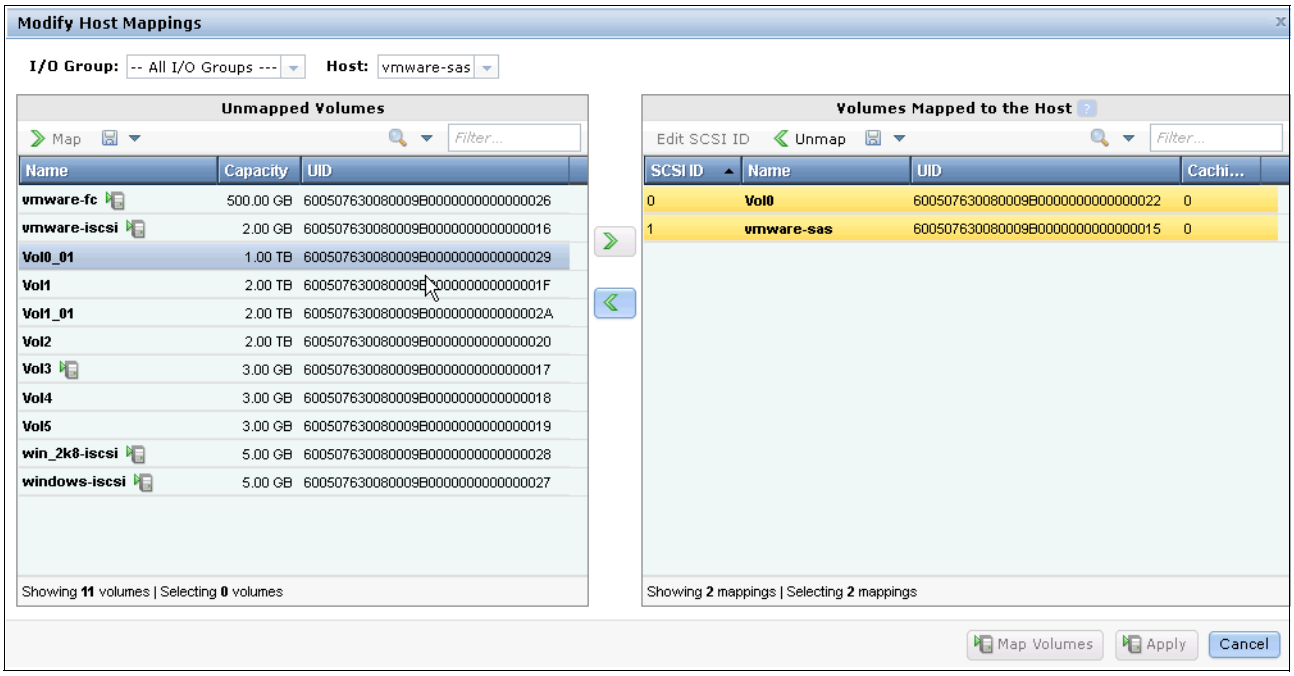

Figure 8-10 Unmapping certain volumes

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.

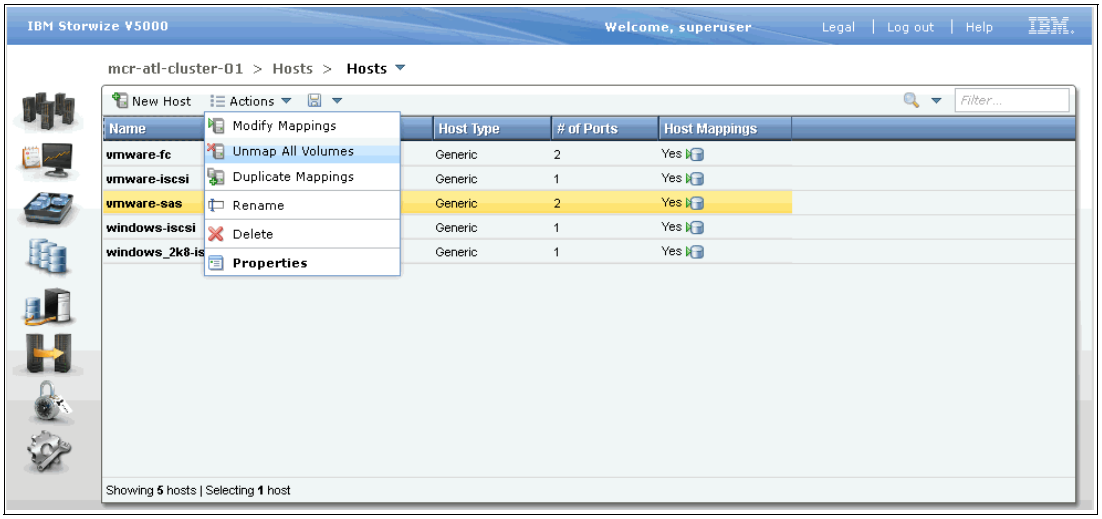

Figure 8-11 Unmap all volumes

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.

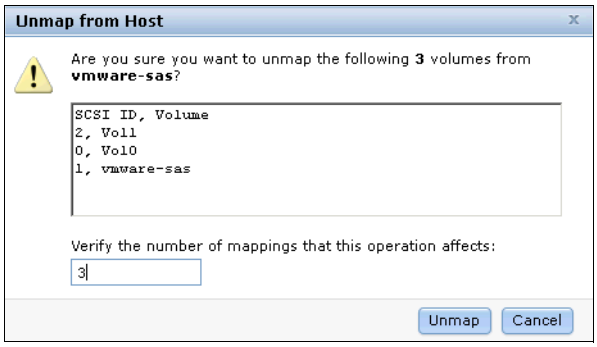

Figure 8-12 Enter the number of mappings to be removed

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.

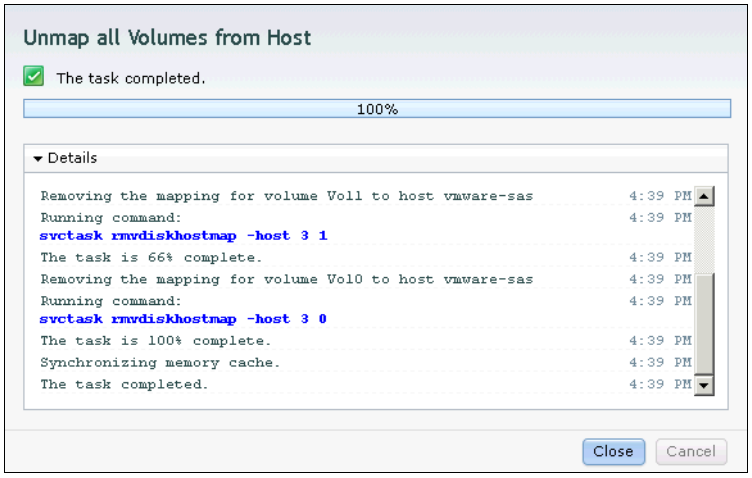

Figure 8-13 Unmapping all volumes from host

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

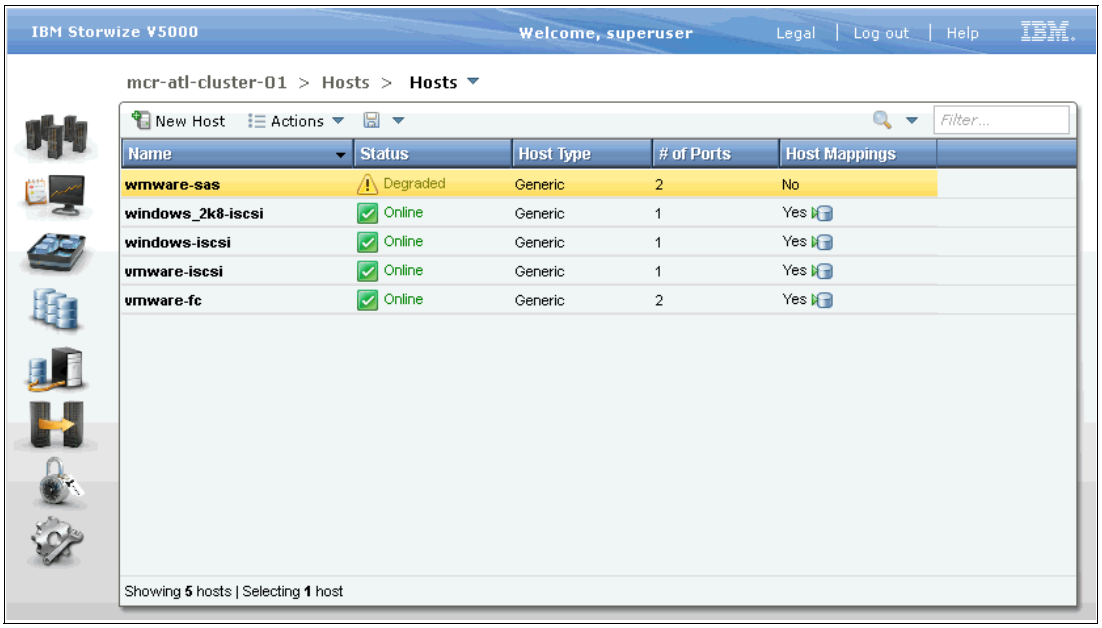

Figure 8-14 Host mapping

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.

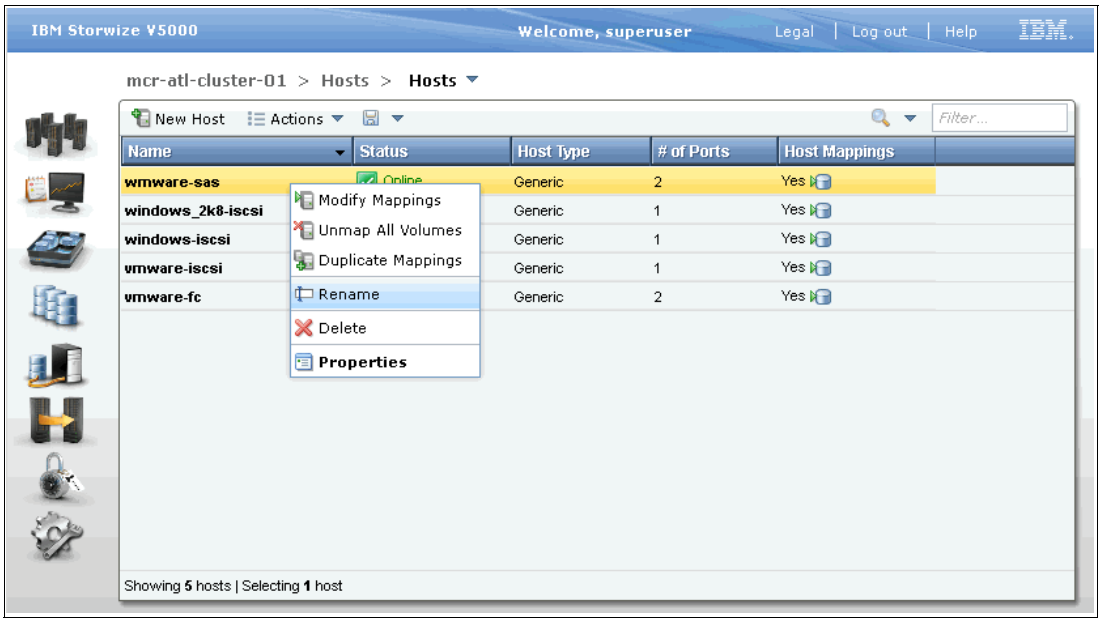

Figure 8-15 Renaming a host

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.

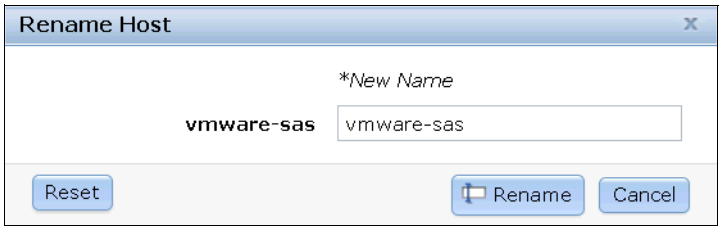

Figure 8-16 Renaming a host window

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

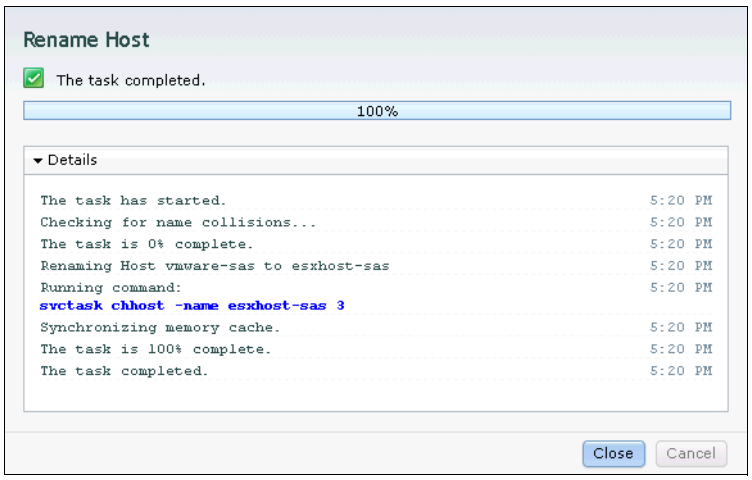

Figure 8-17 Rename a host task completed

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.

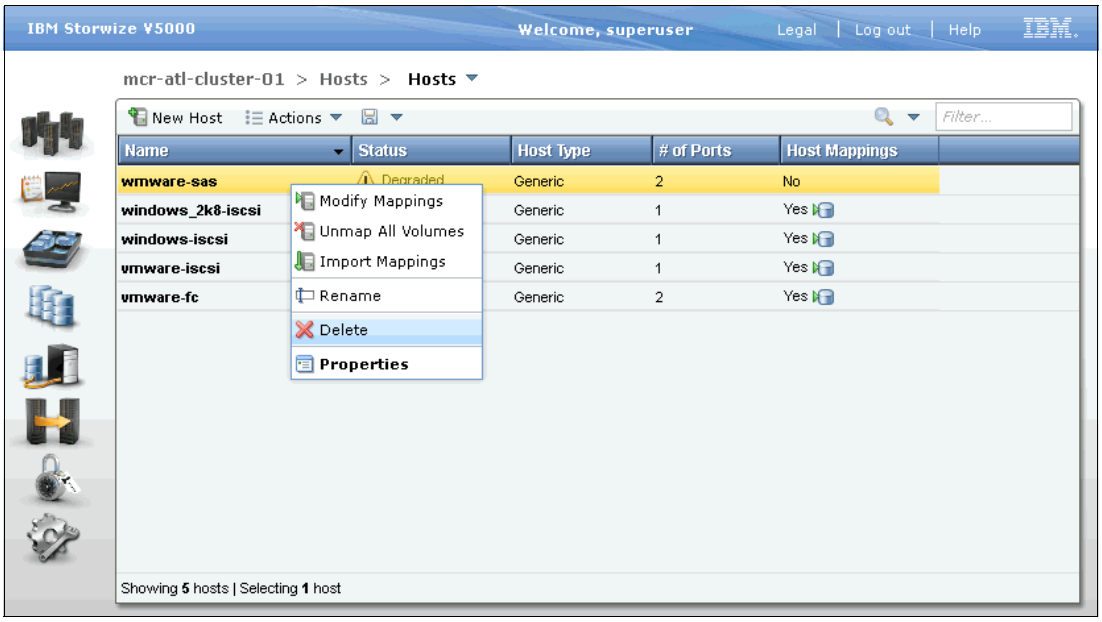

Figure 8-18 Deleting a host

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

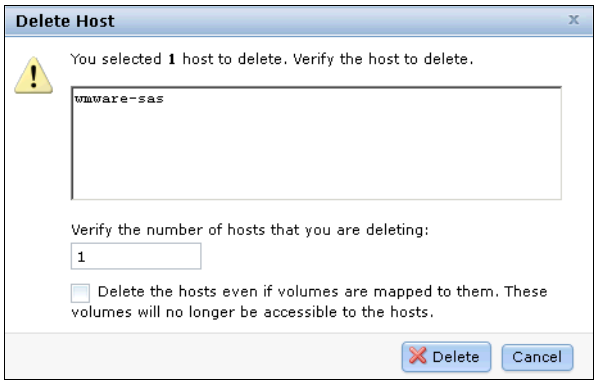

Figure 8-19 Deleting a host

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.

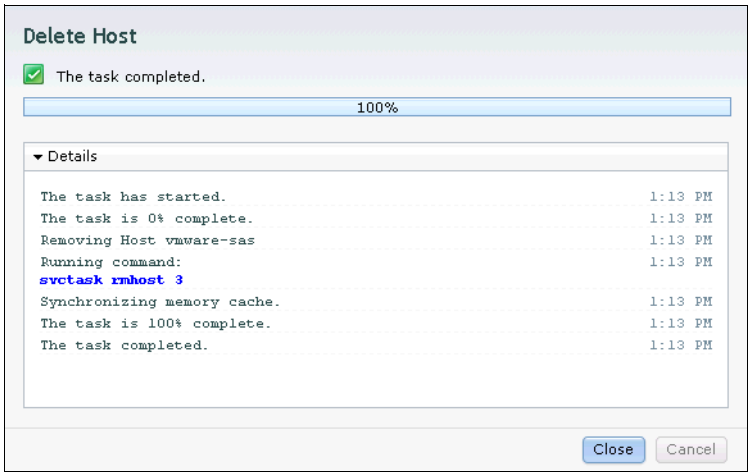

Figure 8-20 Delete host task completed

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.

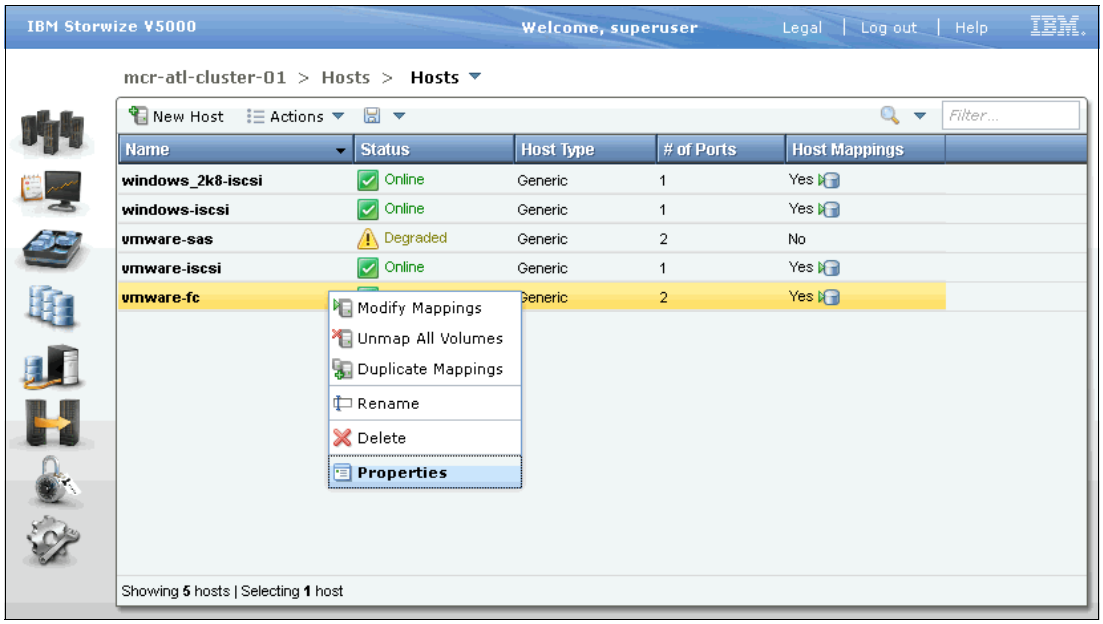

Figure 8-21 Opening host properties

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.

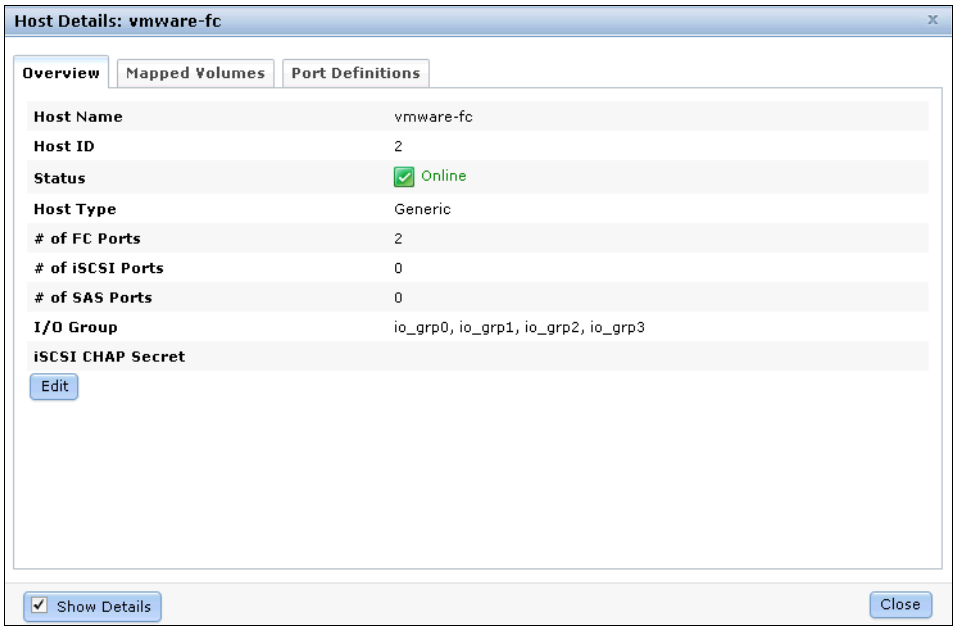

Figure 8-22 Host detail information

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.

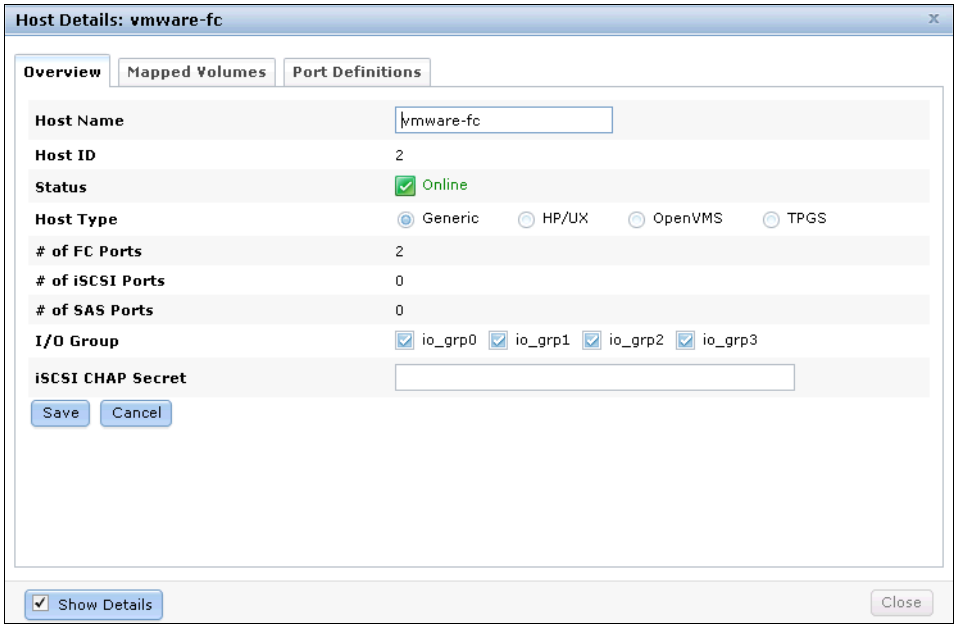

Figure 8-23 Host properties: Editing host information

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.

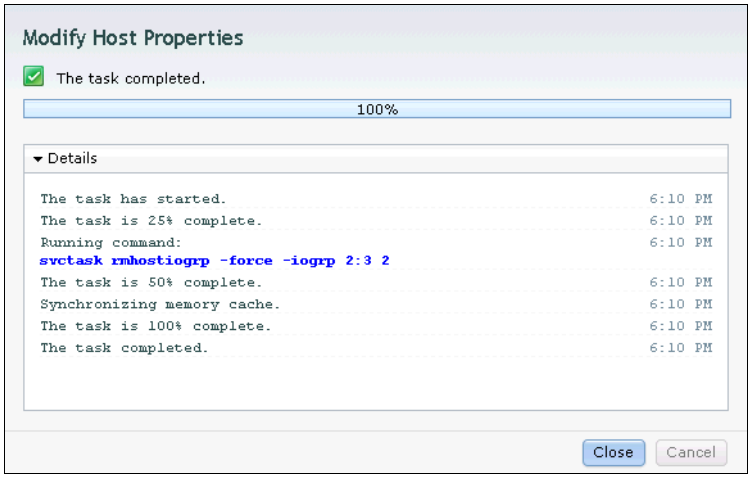

Figure 8-24 Editing host properties task completed

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.

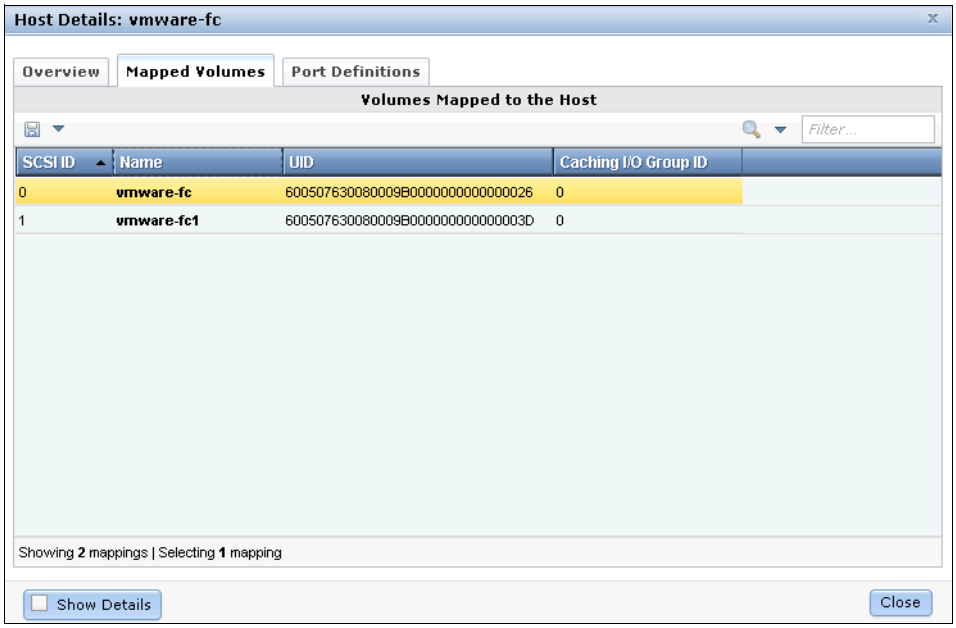

Figure 8-25 Host Details: Mapped volumes information

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.

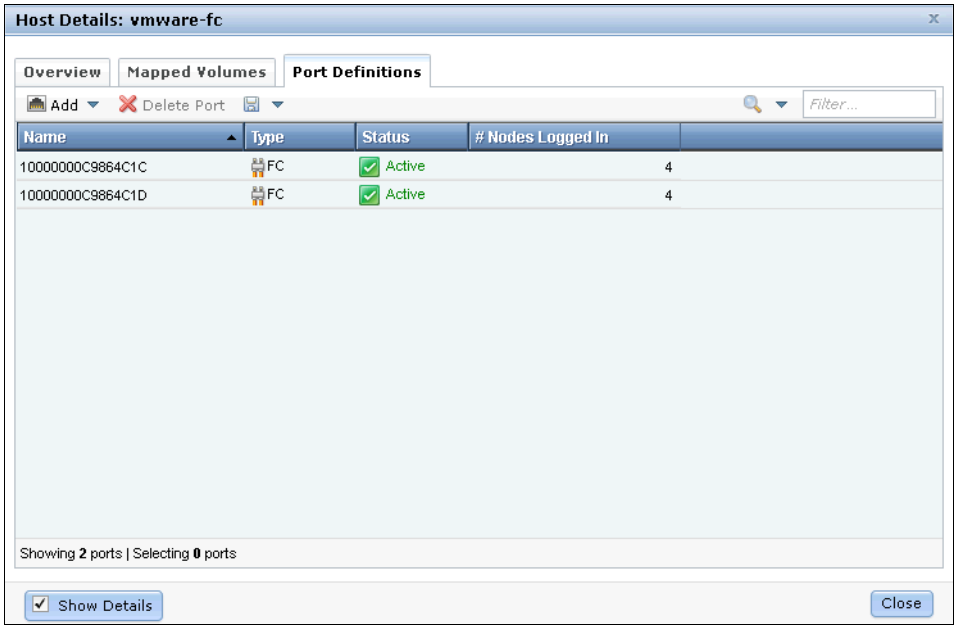

Figure 8-26 Host port details

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.

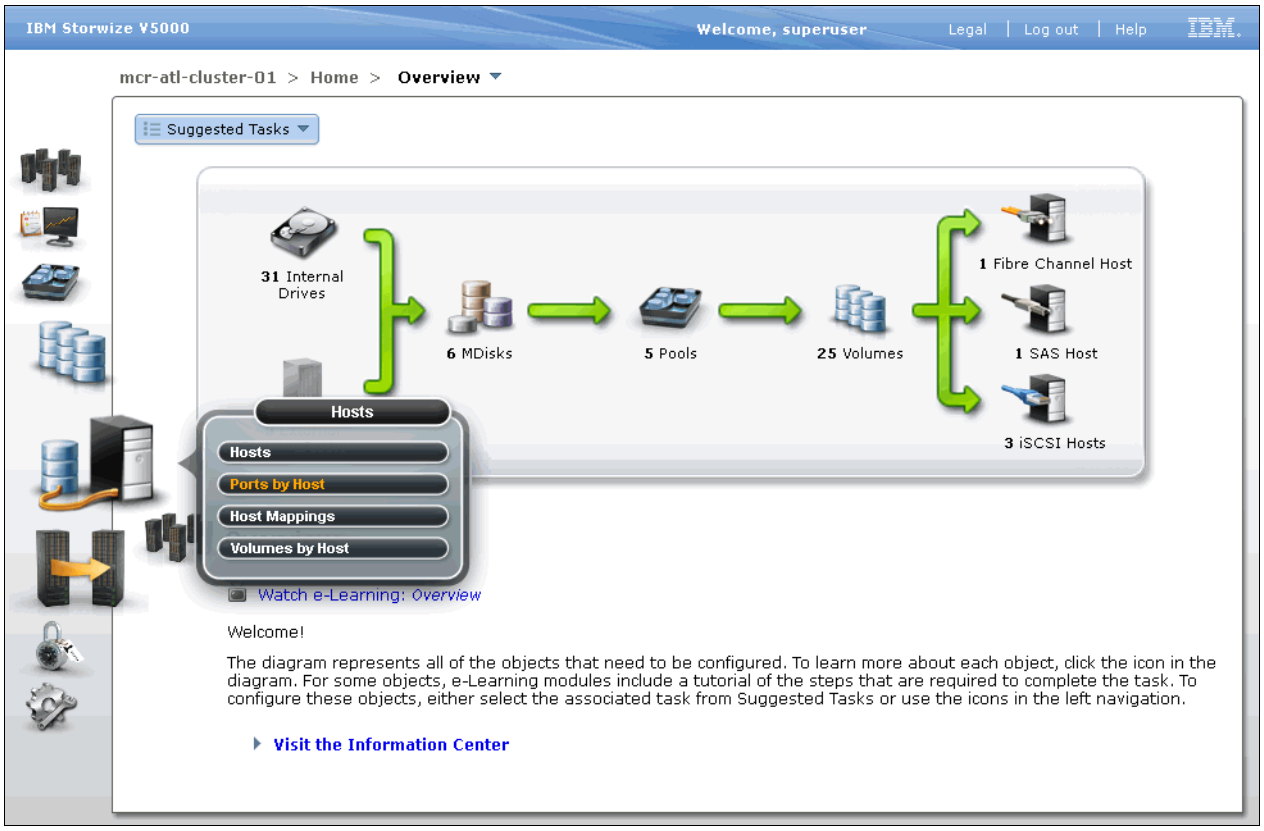

Figure 8-27 Ports by Host window

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.

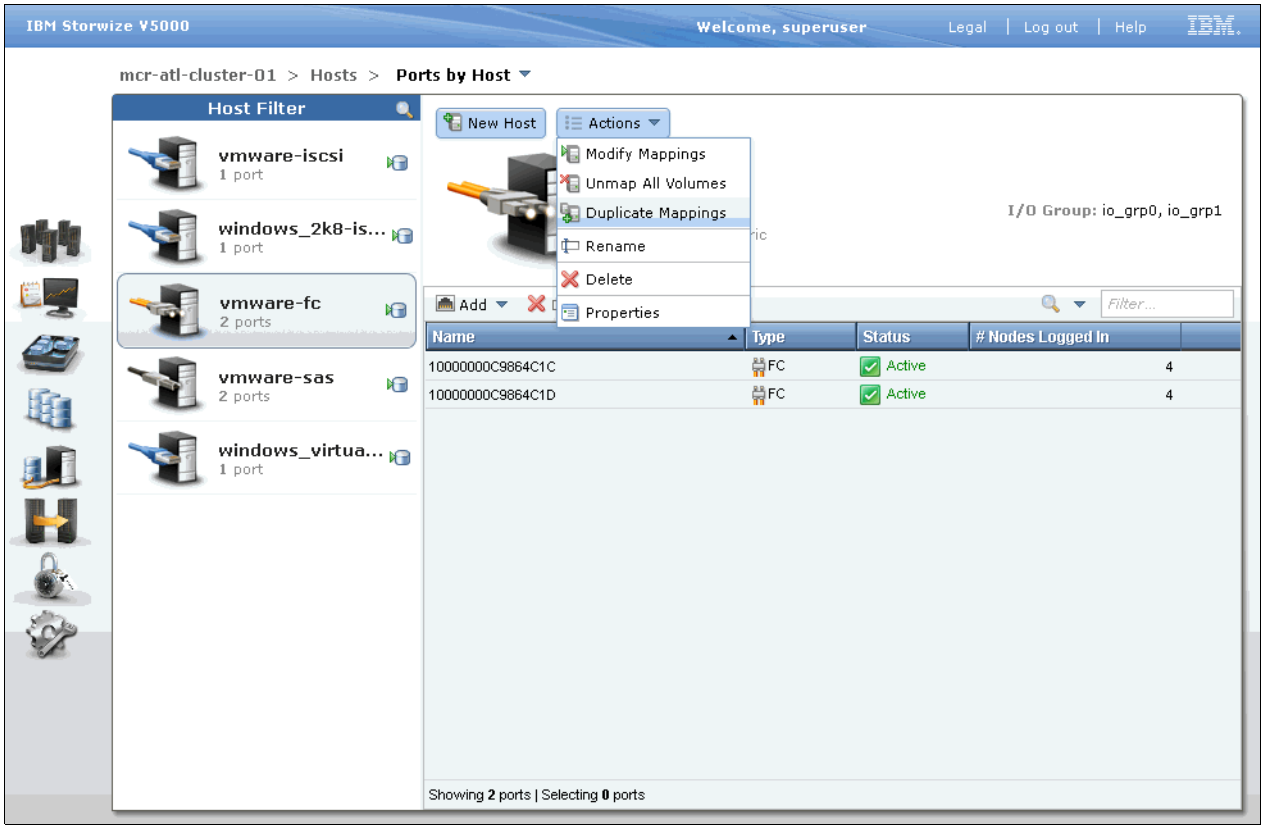

Figure 8-28 Host Action menu

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.

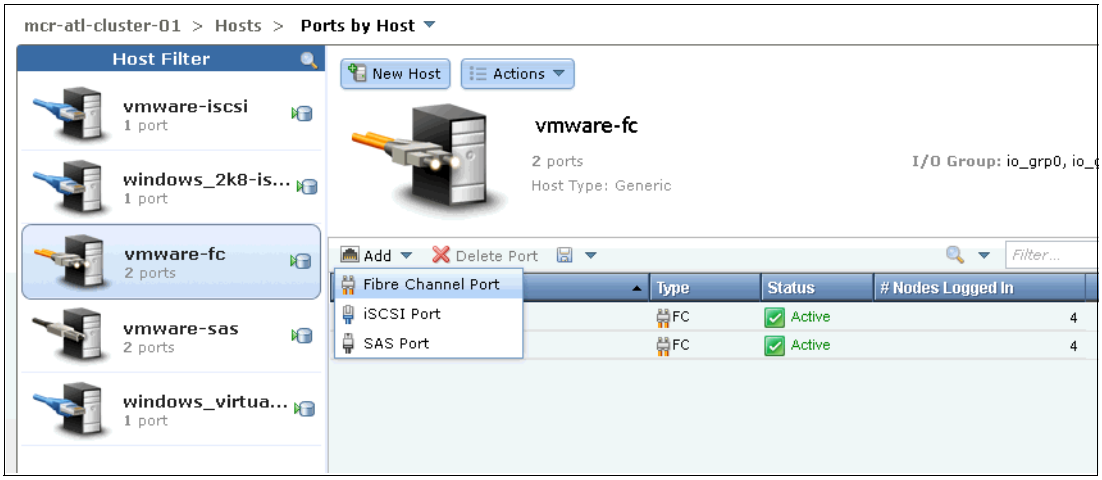

Figure 8-29 Adding a host port

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.

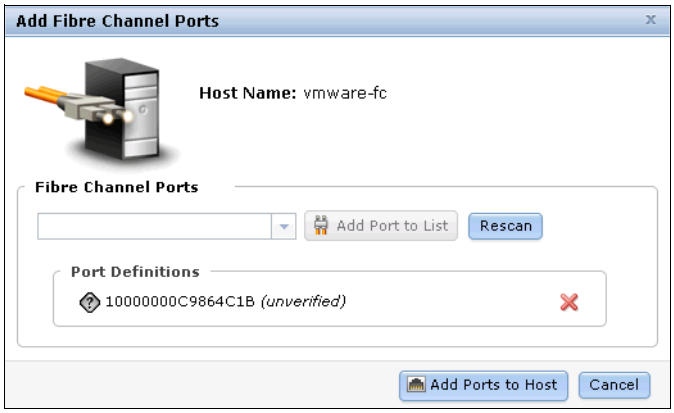

Figure 8-30 Adding offline port

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.

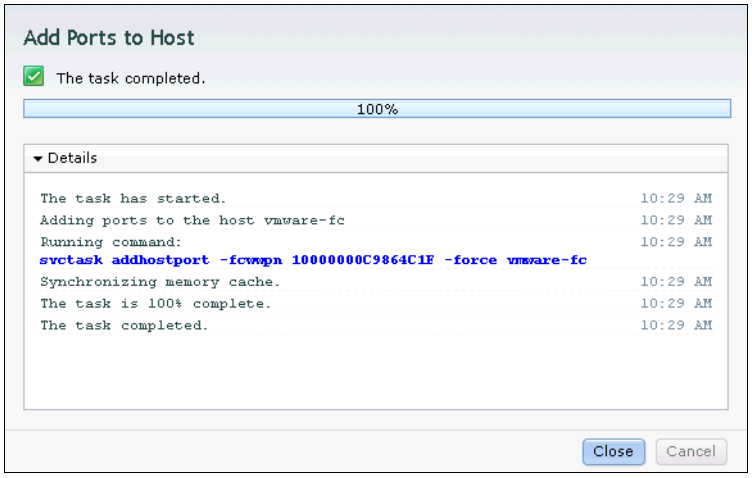

Figure 8-31 Adding a host port

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.

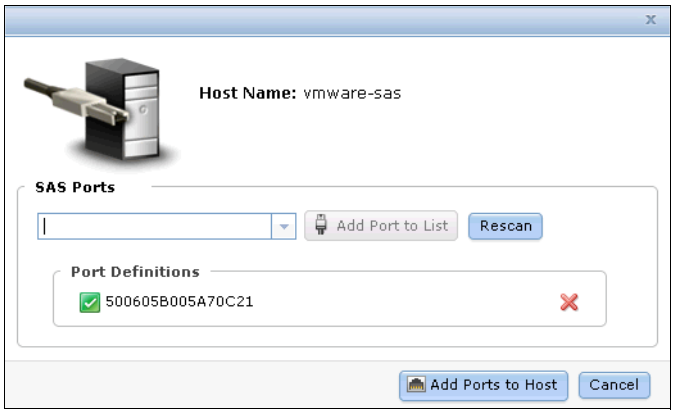

Figure 8-32 Adding an online SAS port

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.

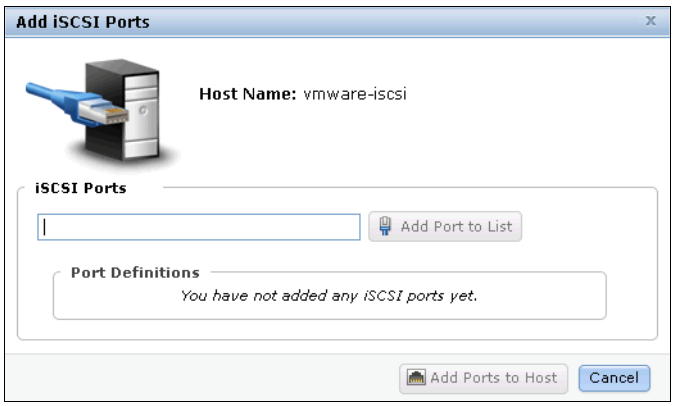

Figure 8-33 Adding iSCSI Host Port

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.

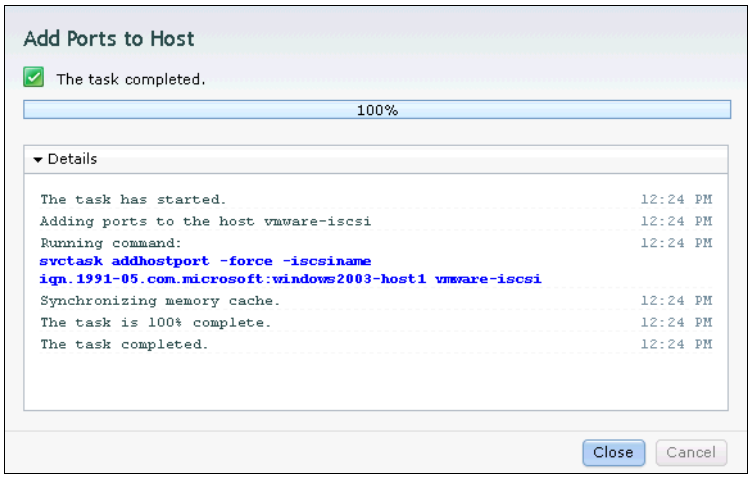

Figure 8-34 Successful iSCSI port addition

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.

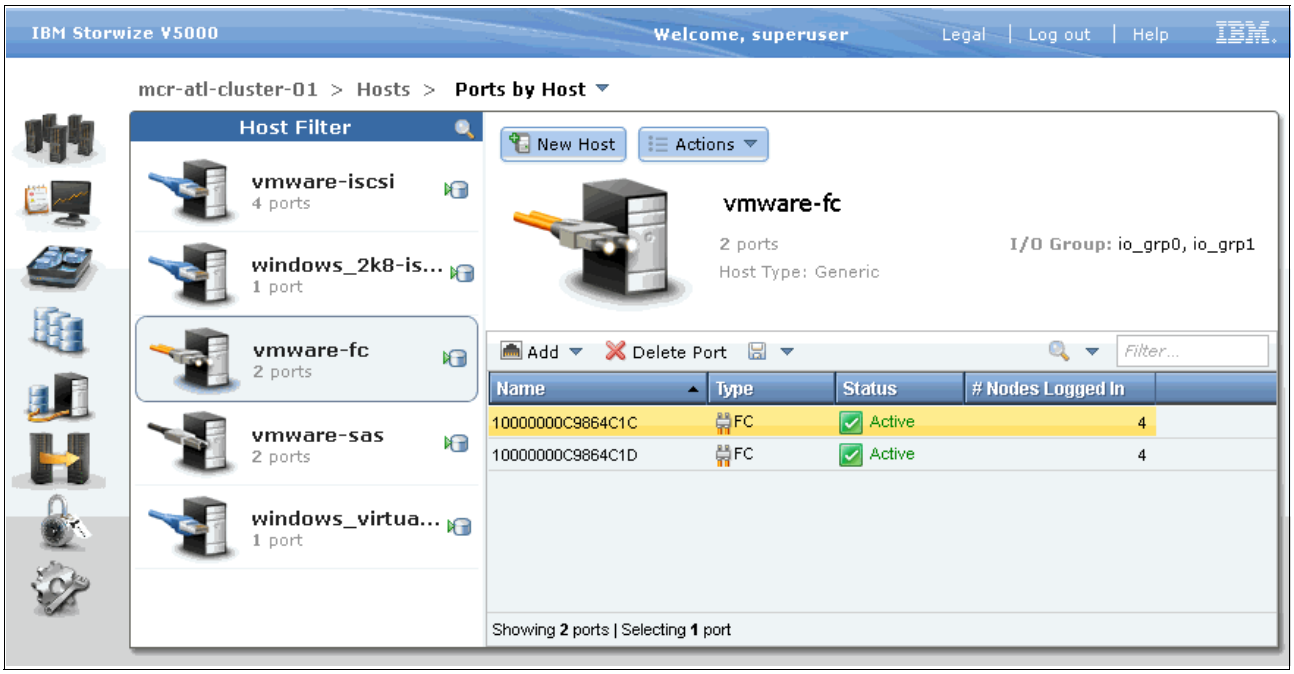

Figure 8-35 Delete host port

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.

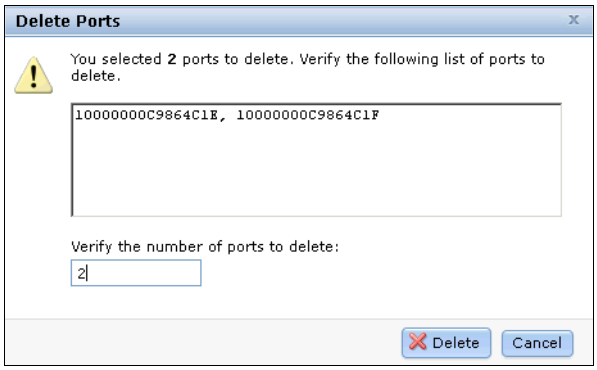

Figure 8-36 Deleting host port

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.

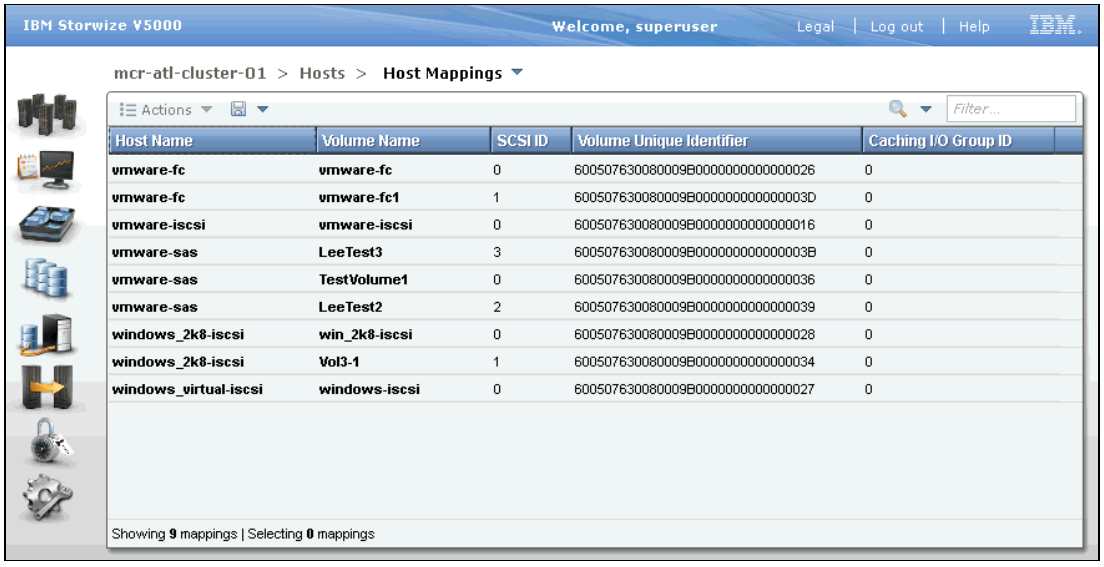

Figure 8-37 Host volume mappings

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)

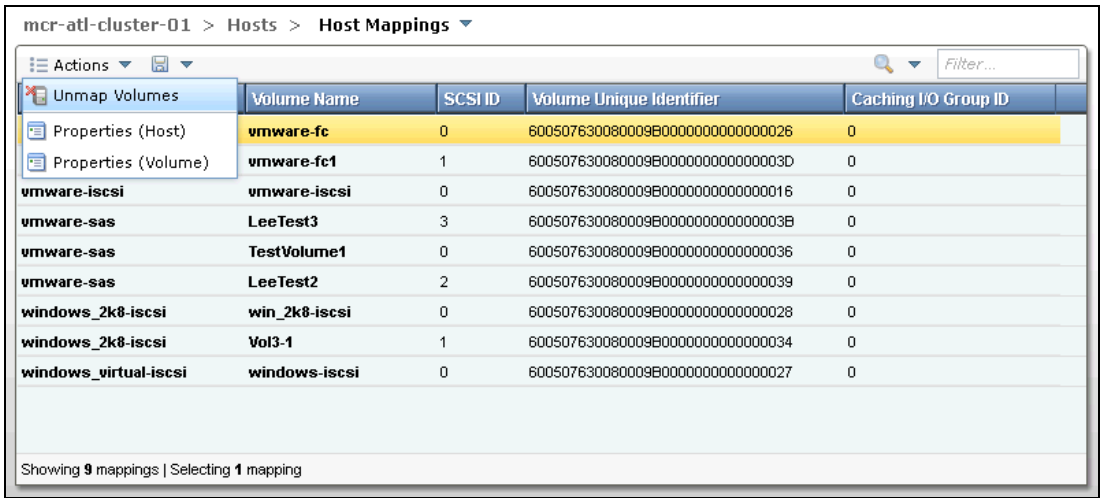

Figure 8-38 Host mapping options

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.

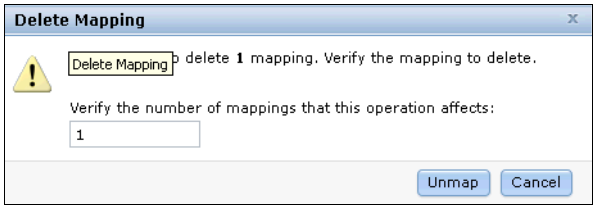

Figure 8-39 Unmapping a volume from host

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`.

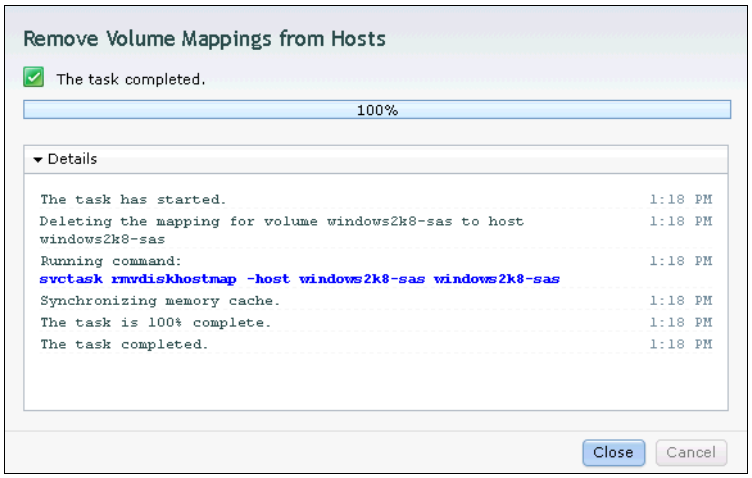

Figure 8-40 Unmapping a volume from host

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

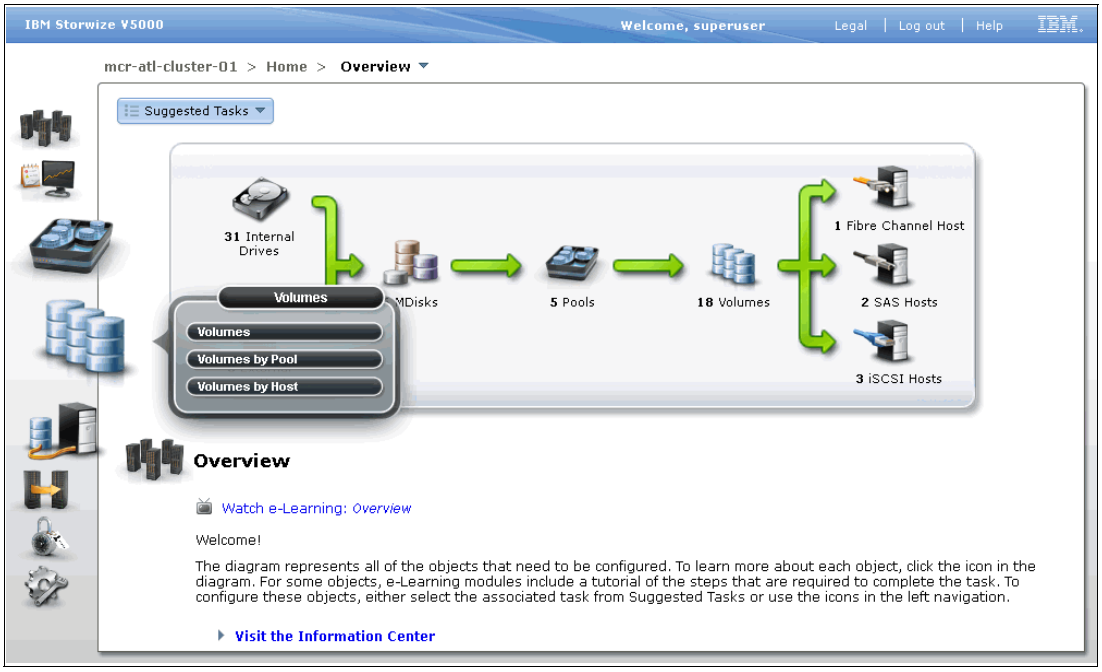

Figure 8-41 Volume options menu

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.

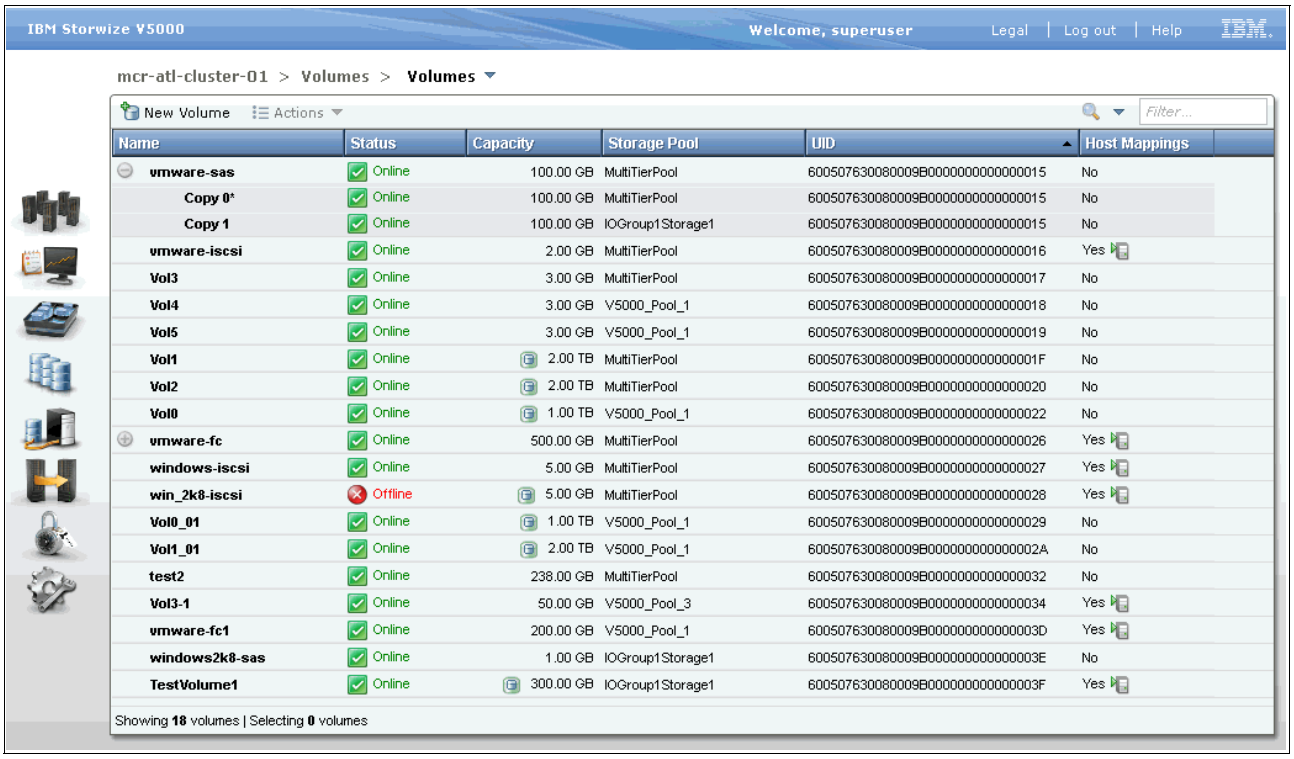

Figure 8-42 Volume window

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.

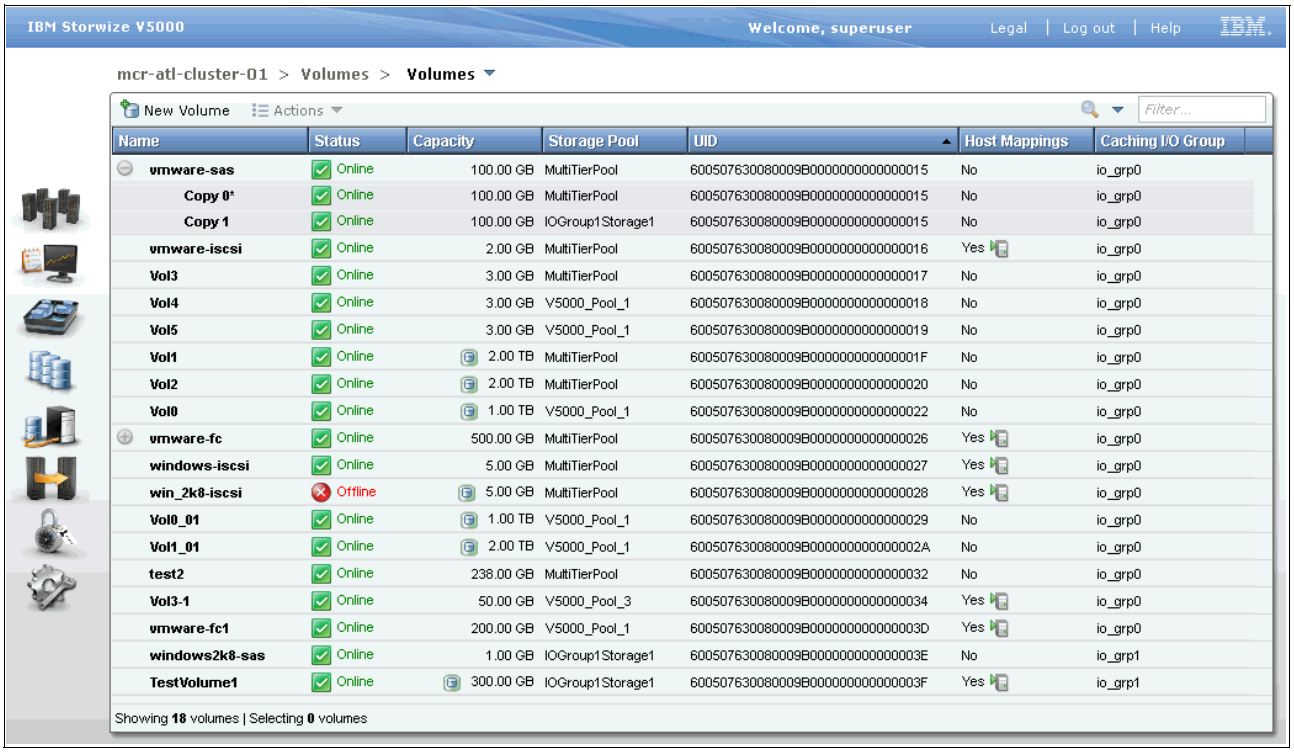

Figure 8-43 Expand volume copies

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.

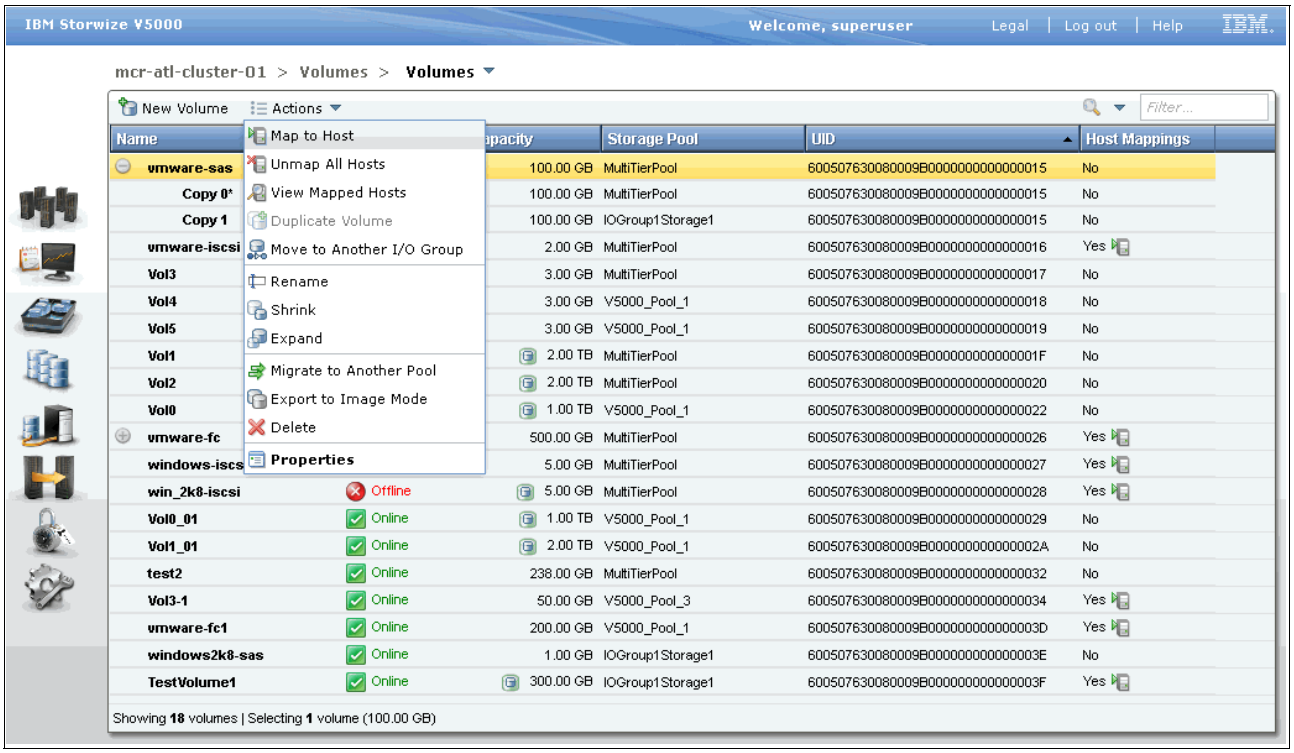

Figure 8-44 Listing the action options for volume.

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.

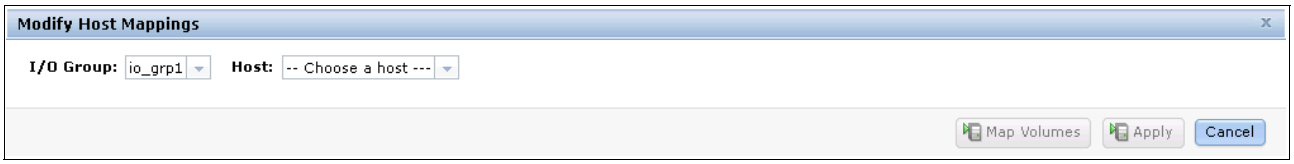

Figure 8-45 Modify Host Mappings menu

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.

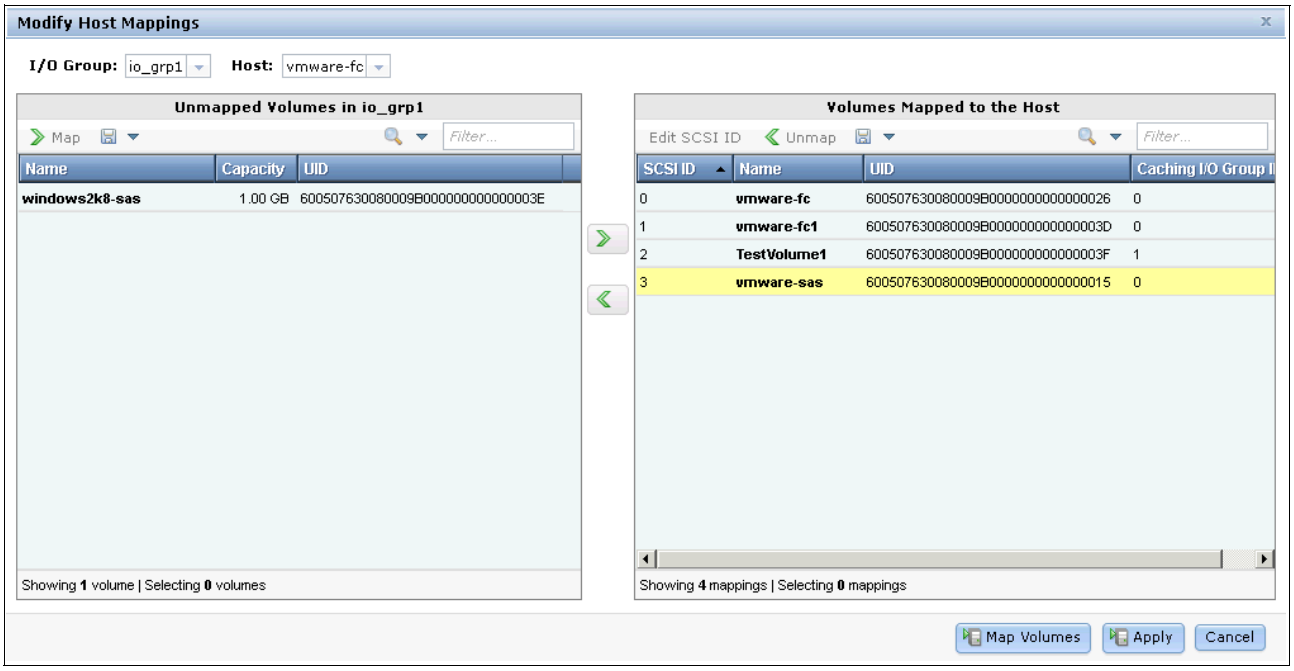

Figure 8-46 Modify Host Mappings

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.

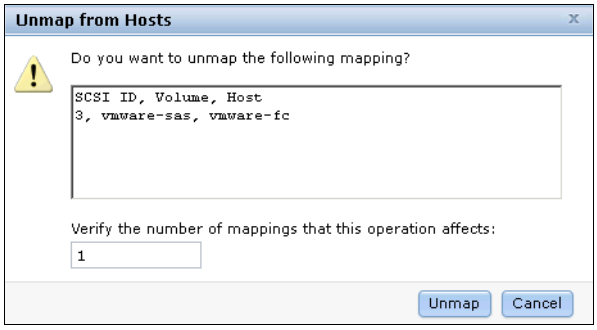

Figure 8-47 Unmapping from host (or hosts)

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.

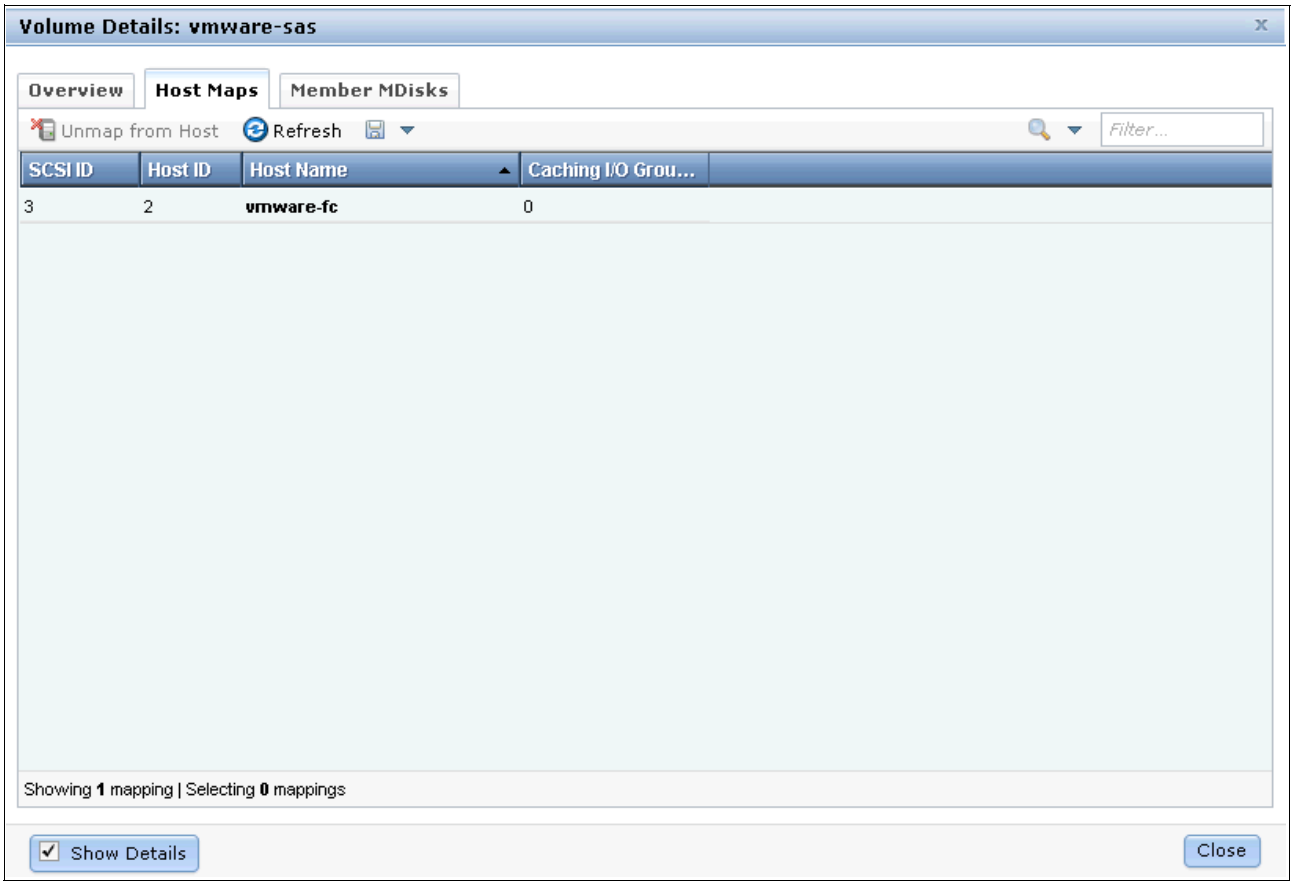

Figure 8-48 Volume to host mapping

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.

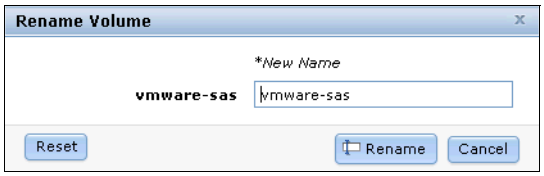

Figure 8-49 Renaming a volume

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.

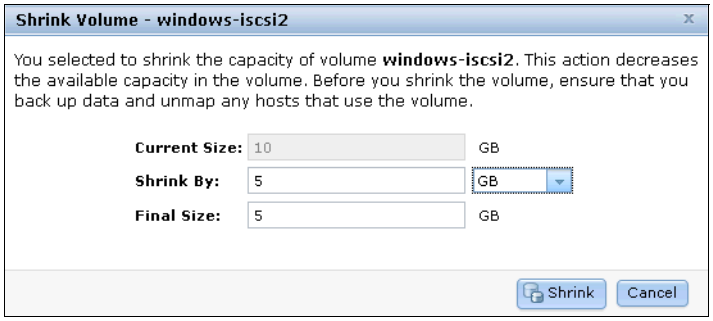

Figure 8-50 Shrink Volume window

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.

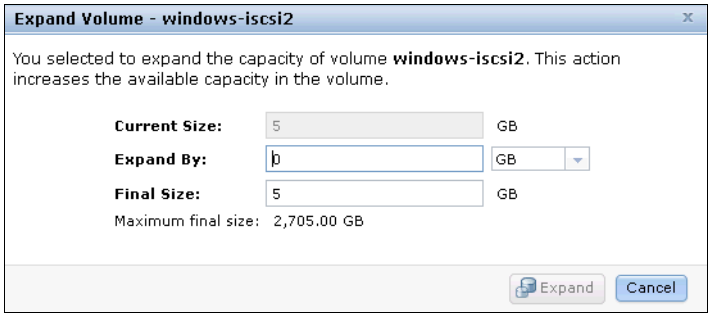

Figure 8-51 Expand Volume window

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.

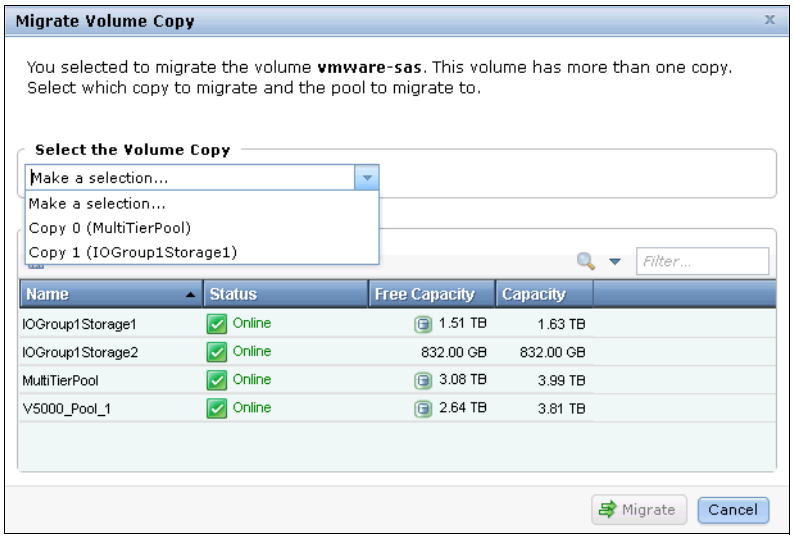

Figure 8-52 Migrate Volume

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.

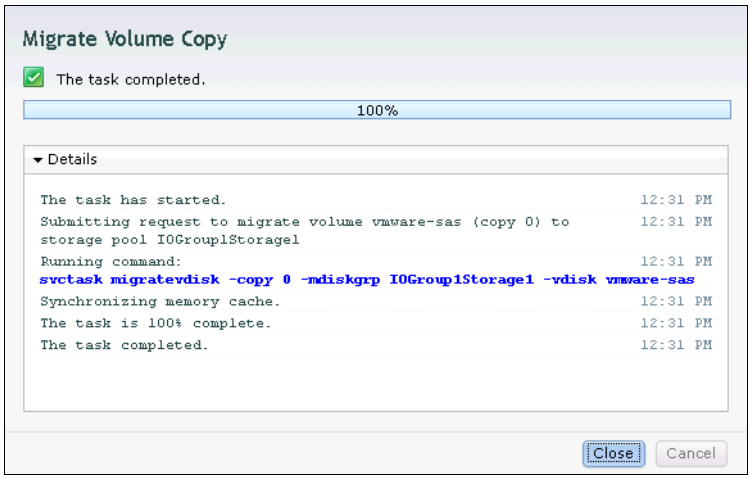

Figure 8-53 Volume Copy Migration starts

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.

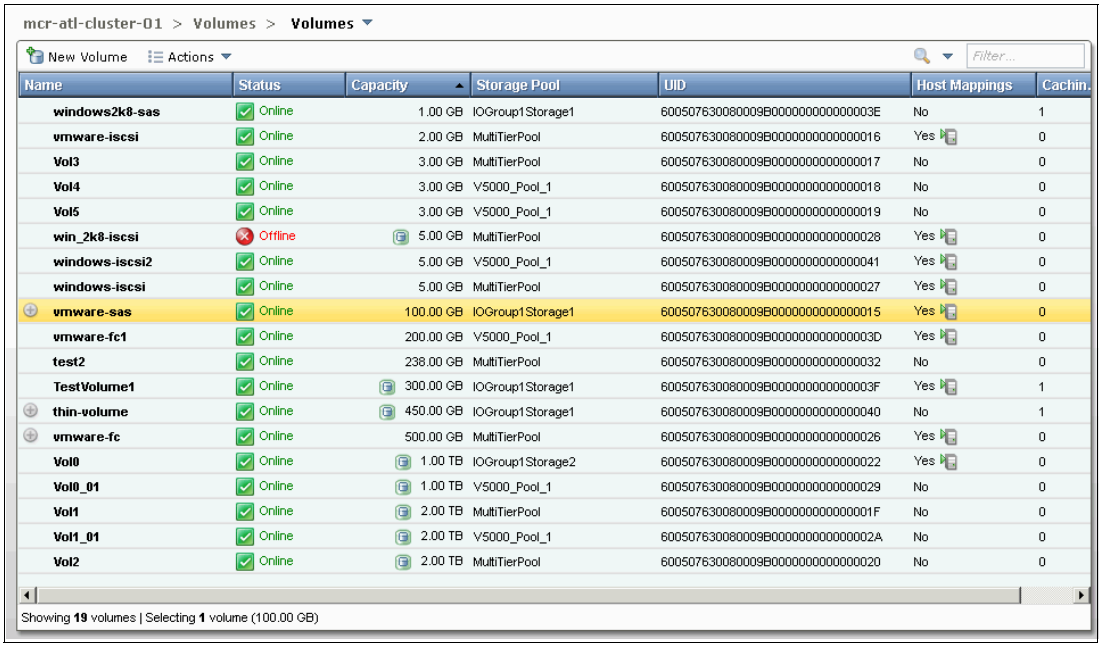

Figure 8-54 Volume showing at new storage pool

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.

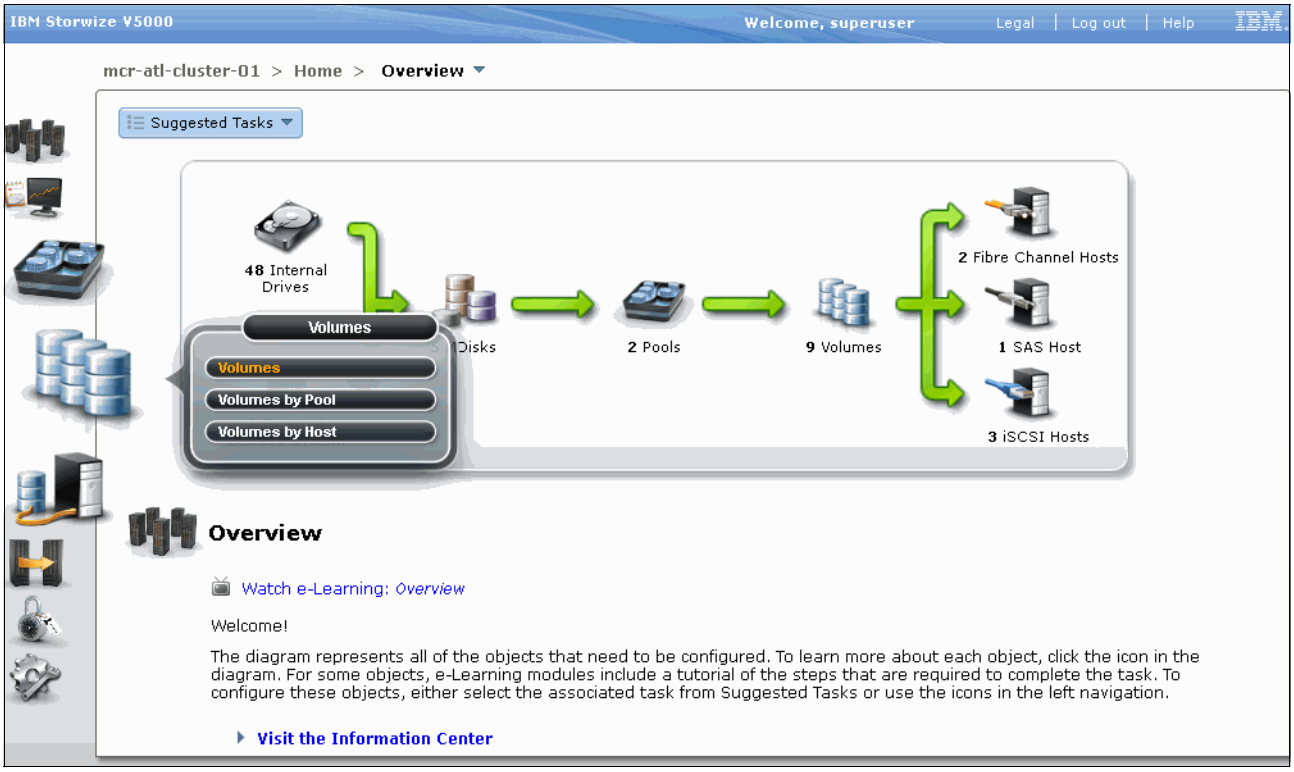

Figure 8-55 Exporting a volume to an image mode

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.

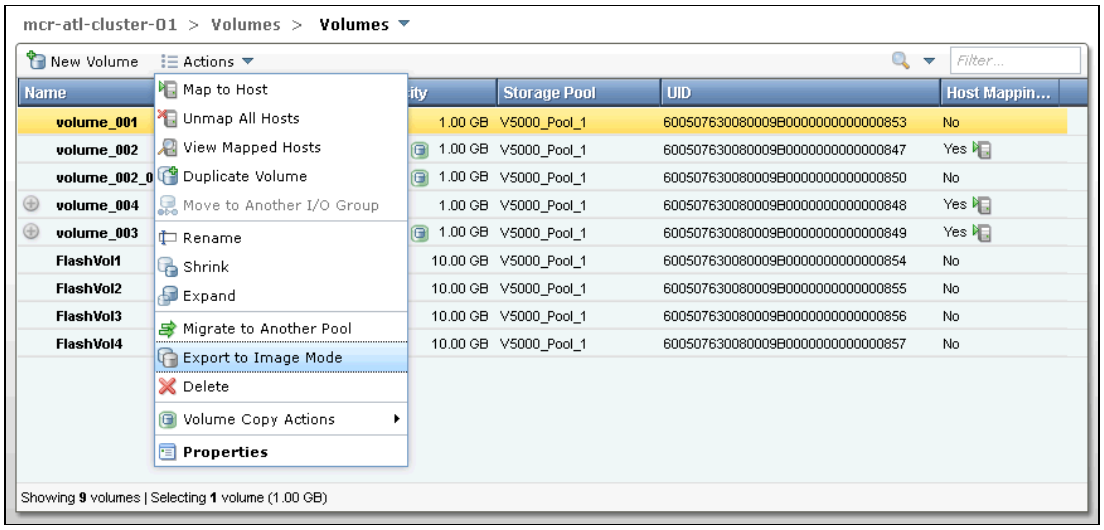

Figure 8-56 Exporting a volume to an image mode

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.

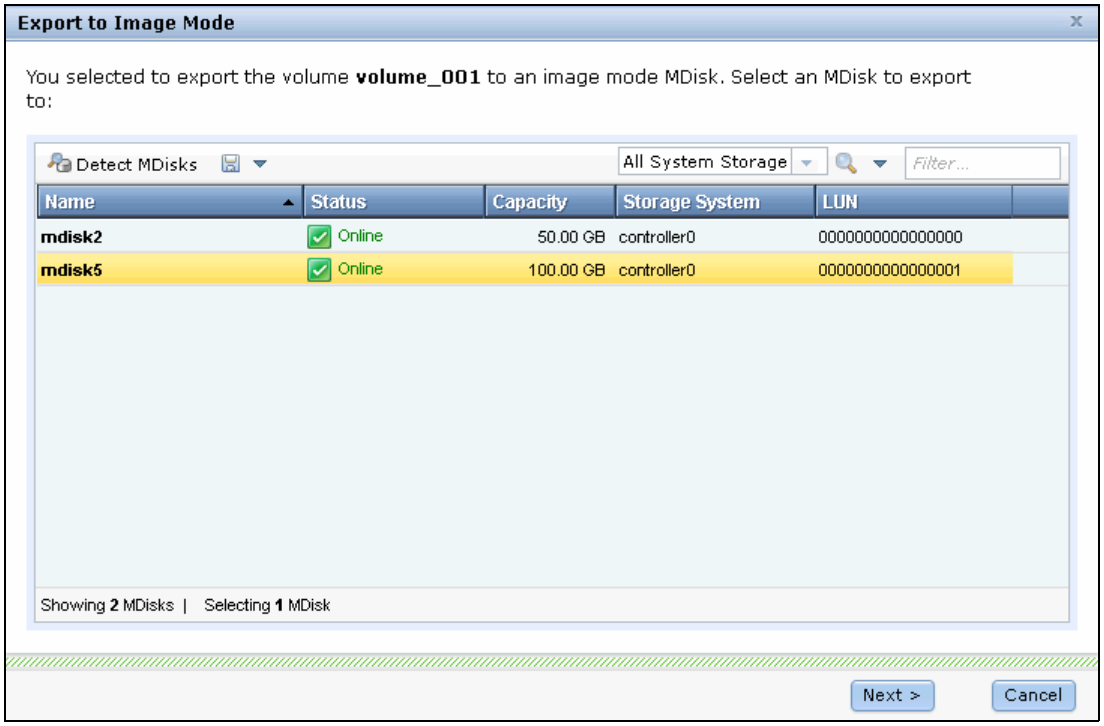

Figure 8-57 Selecting the Manage Disk to export the volume.

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.

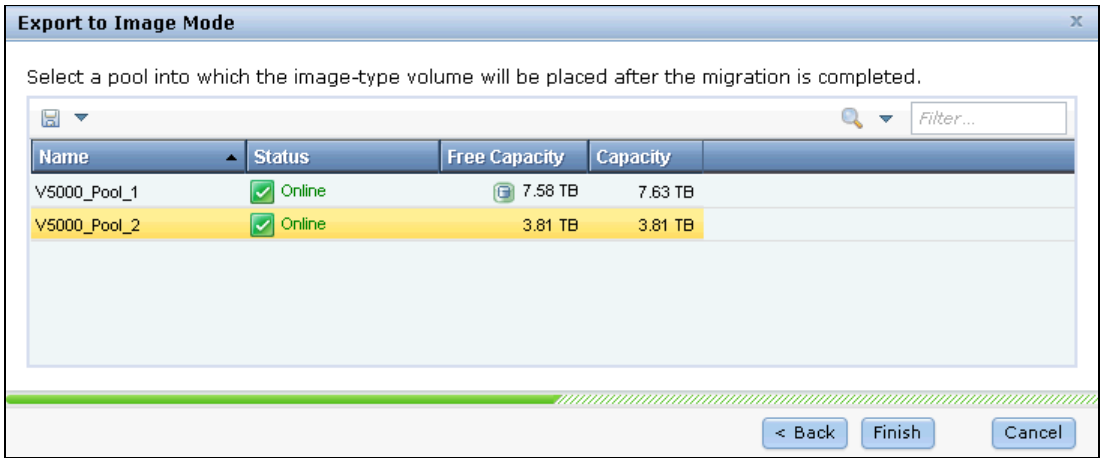

Figure 8-58 Select the Storage Pool

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.

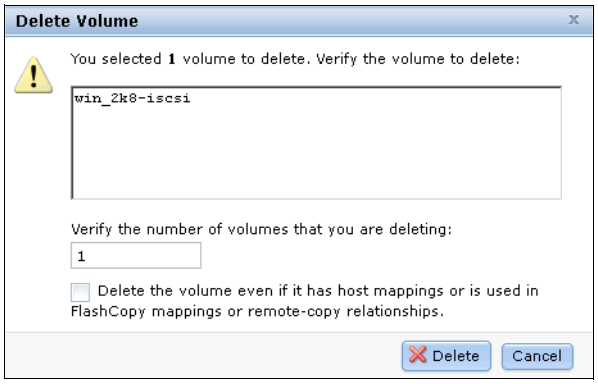

Figure 8-59 Delete Volume window

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

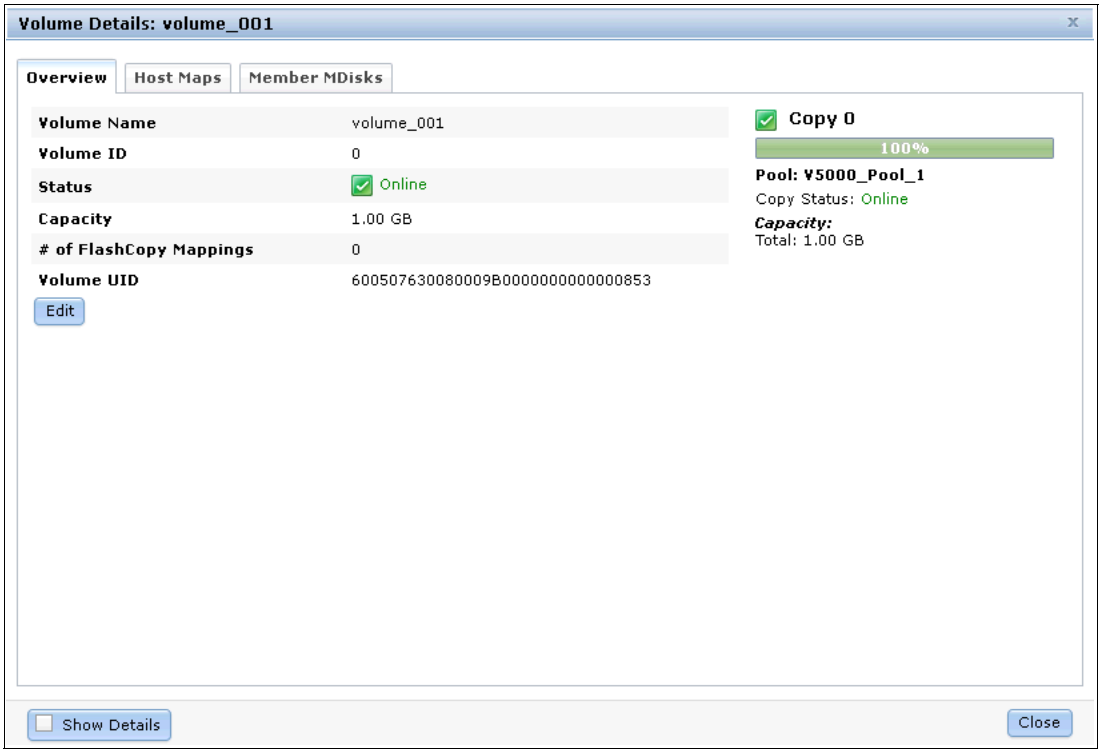

Figure 8-60 Volume Details: Overview tab

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.

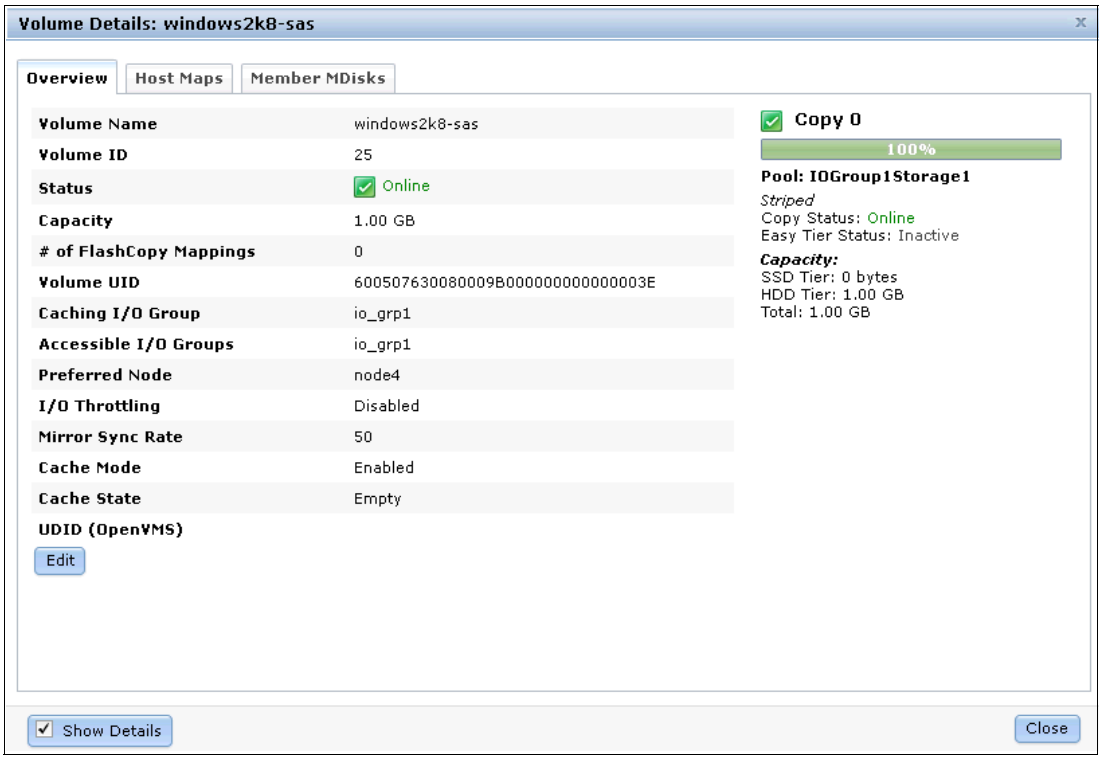

Figure 8-61 Volume properties window

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.

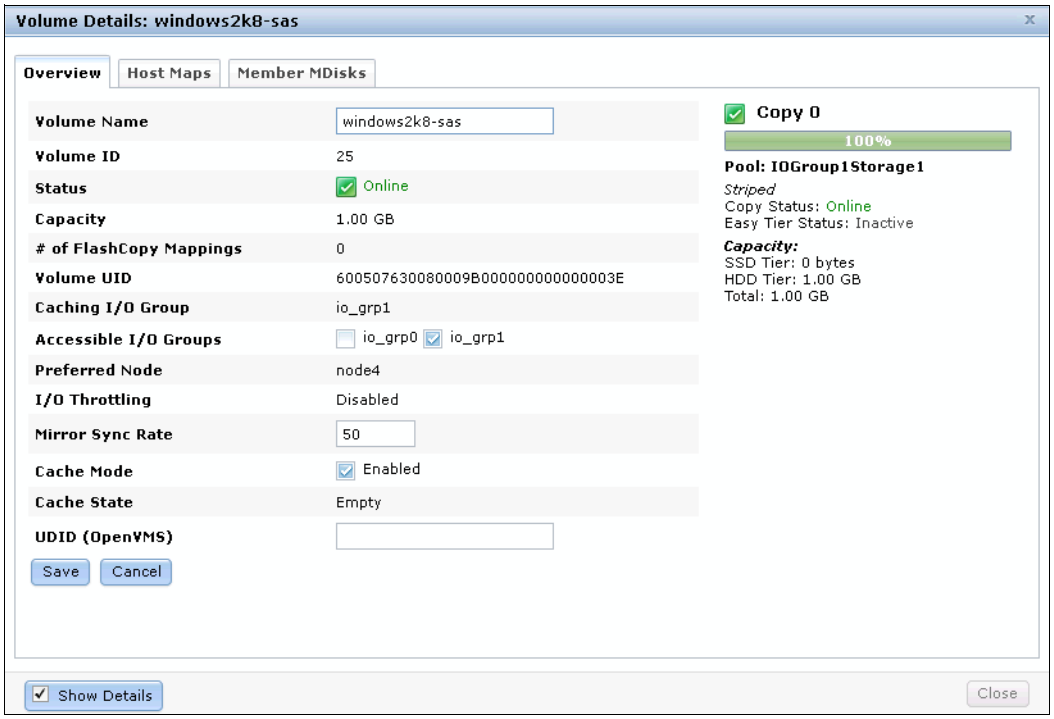

Figure 8-62 Modify Volume Details window

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.

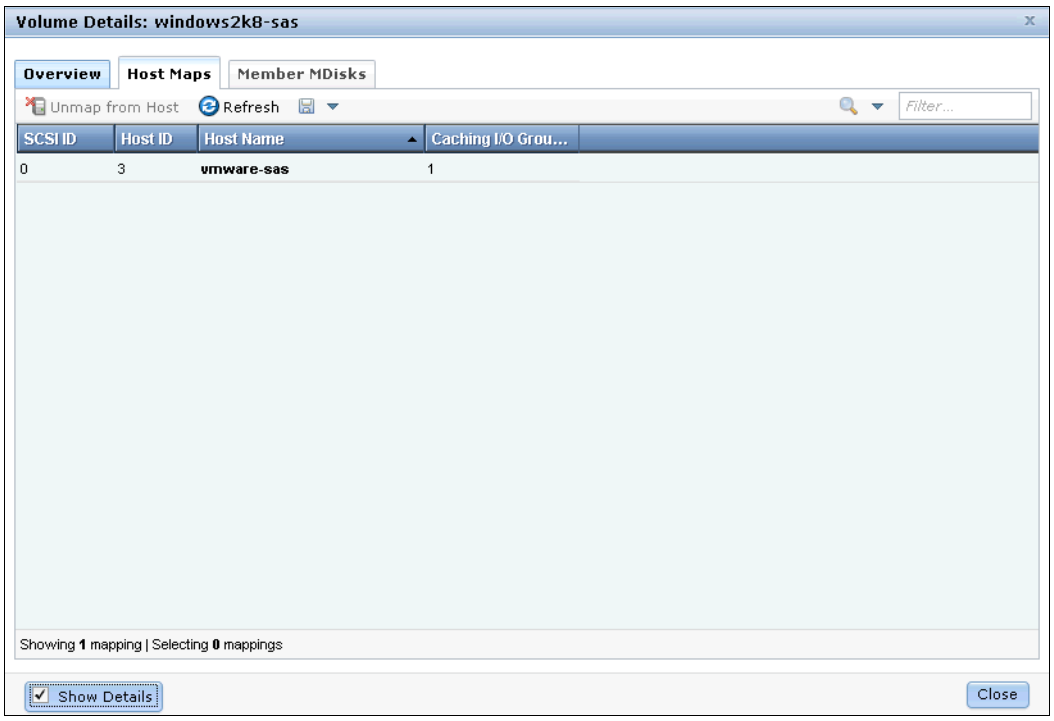

Figure 8-63 Host Maps

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.

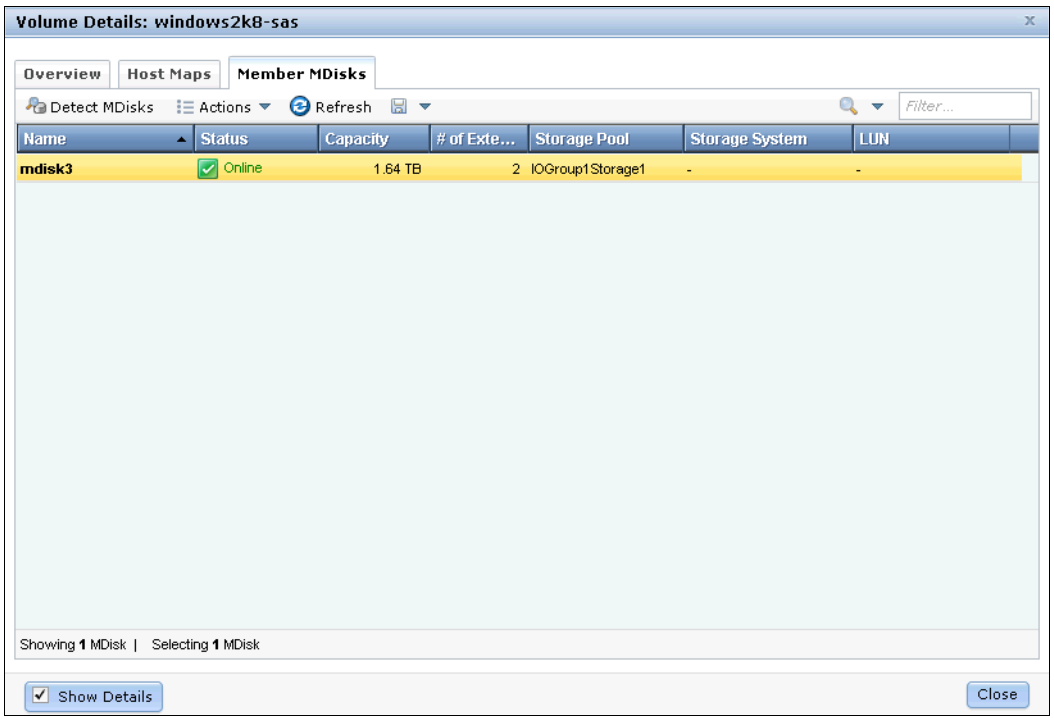

Figure 8-65 Member MDisk tab

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.

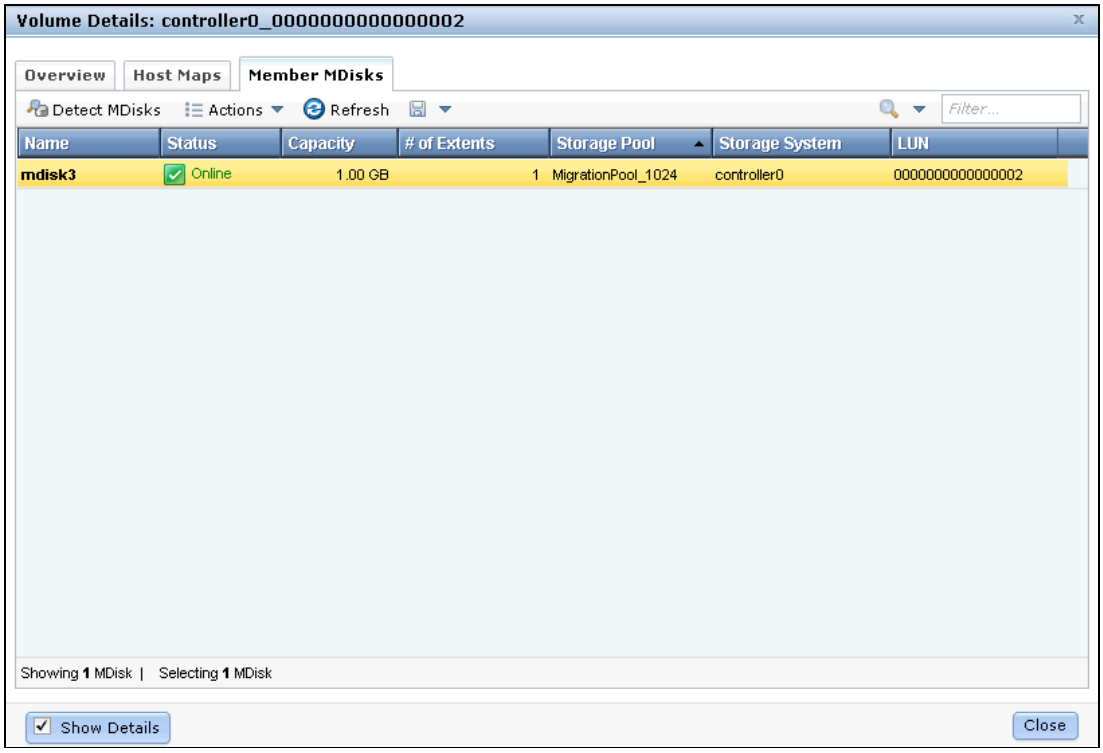

Figure 8-66 Image mode volume details

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.

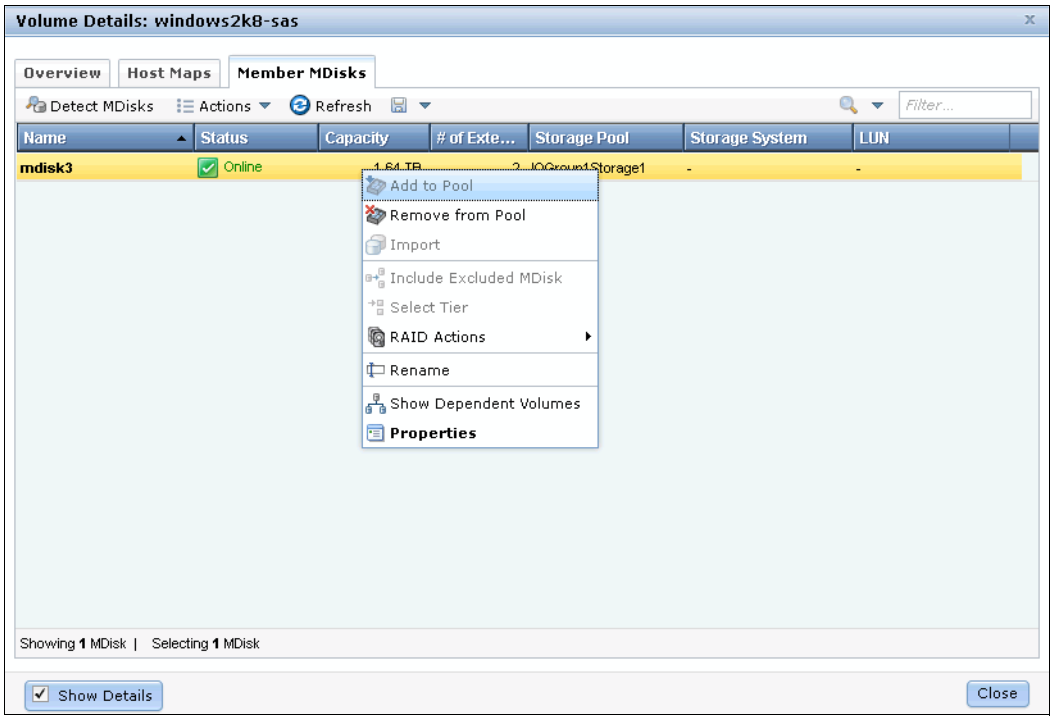

Figure 8-67 MDisk action menu

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.

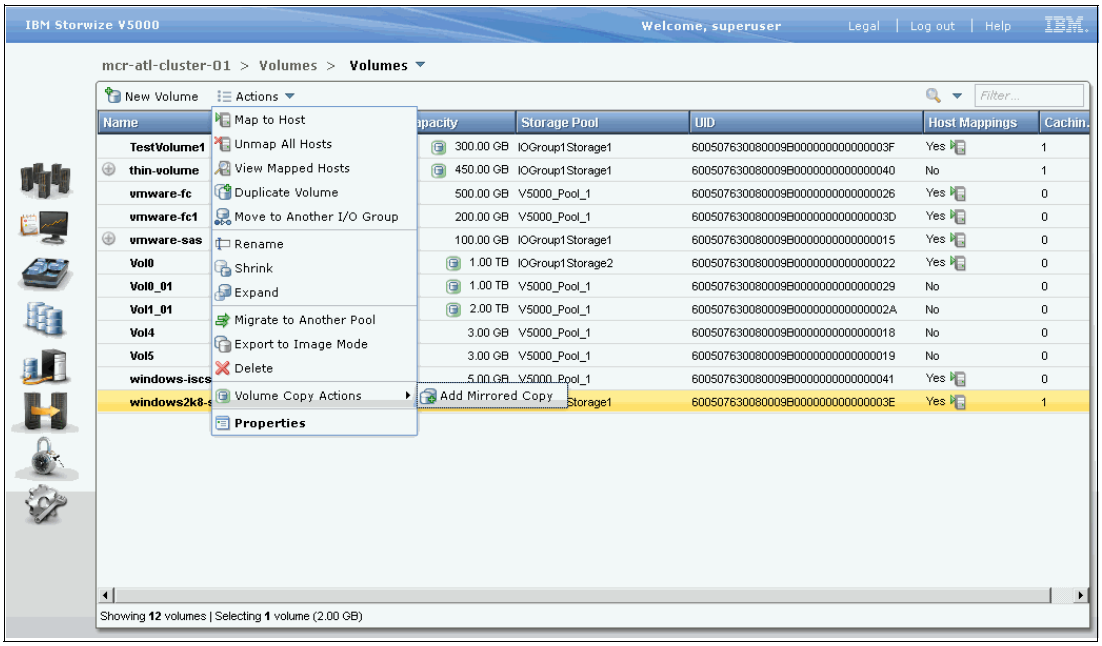

Figure 8-68 Add mirrored copy

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**.

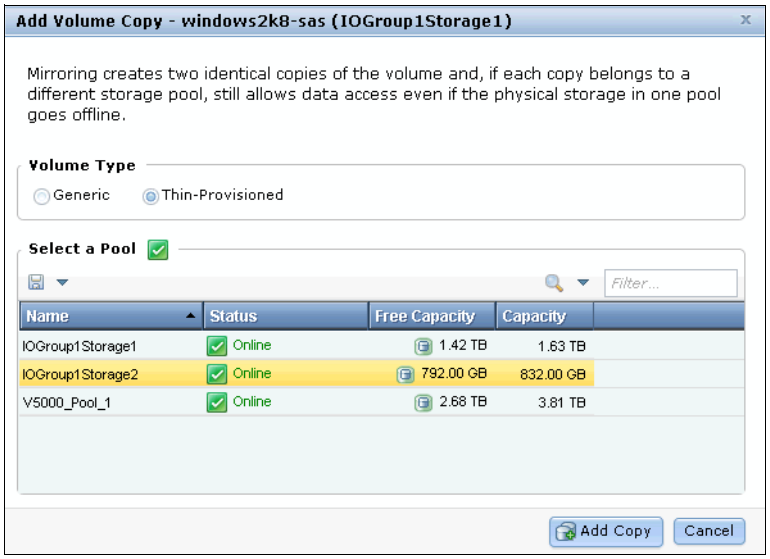

Figure 8-69 Select storage pool

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.

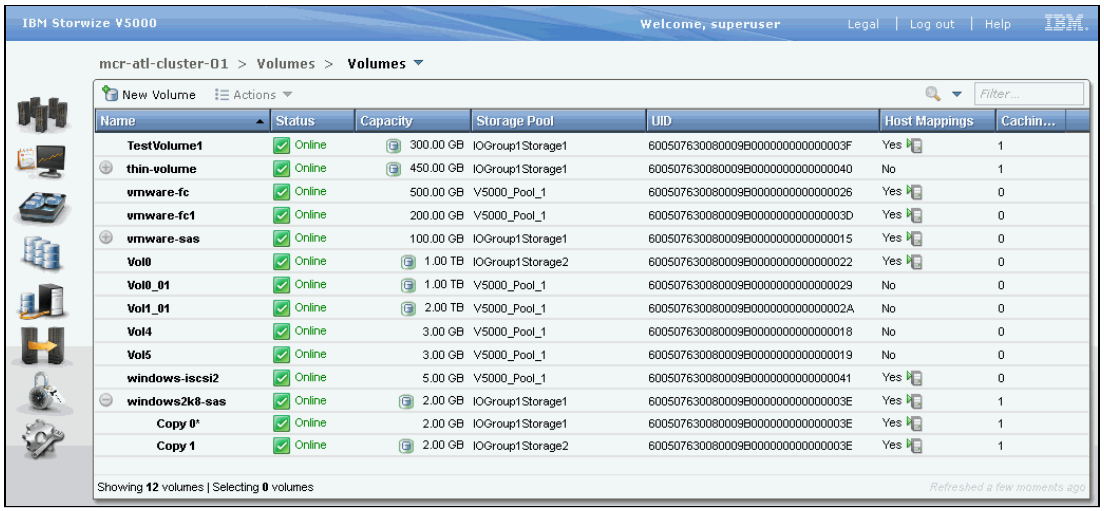

Figure 8-70 Volume containing two copies

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

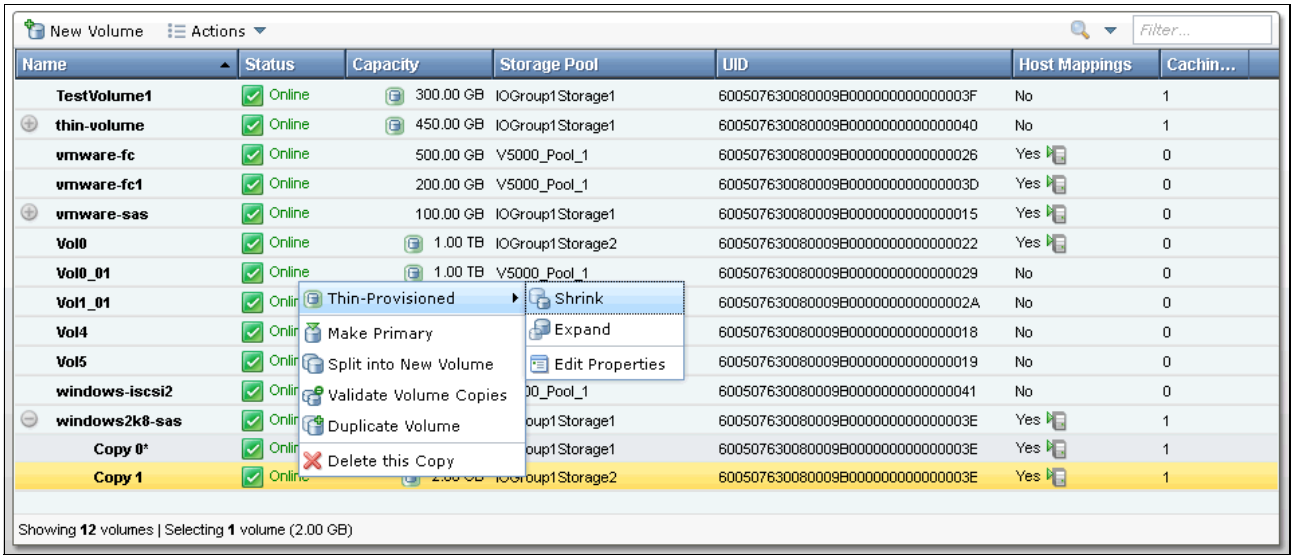

Figure 8-71 Working with thin-provisioned volumes

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.

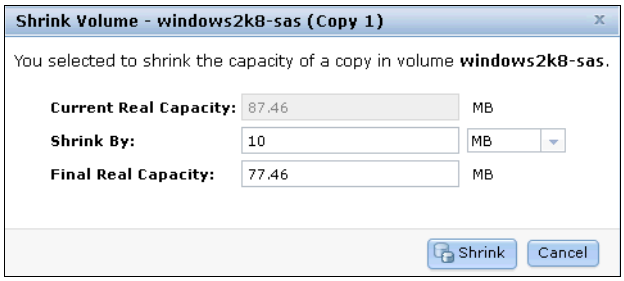

Figure 8-72 Shrink Volume window

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.

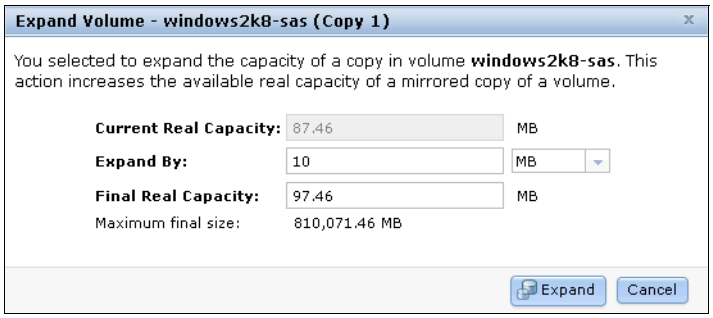

Figure 8-73 Expand Volume window

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.

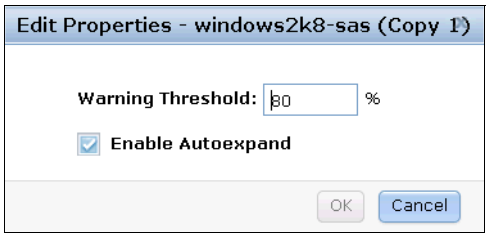

Figure 8-74 Edit Properties window

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

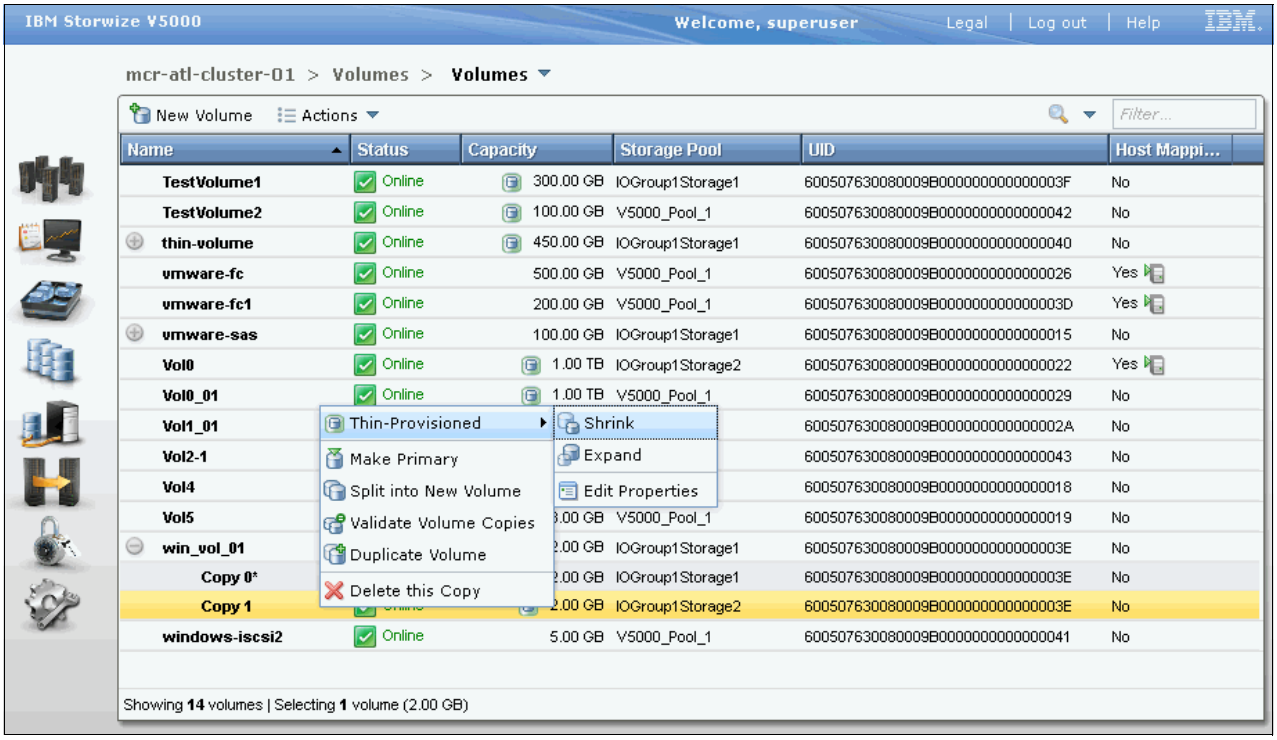

Figure 8-75 Volume copy actions

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.

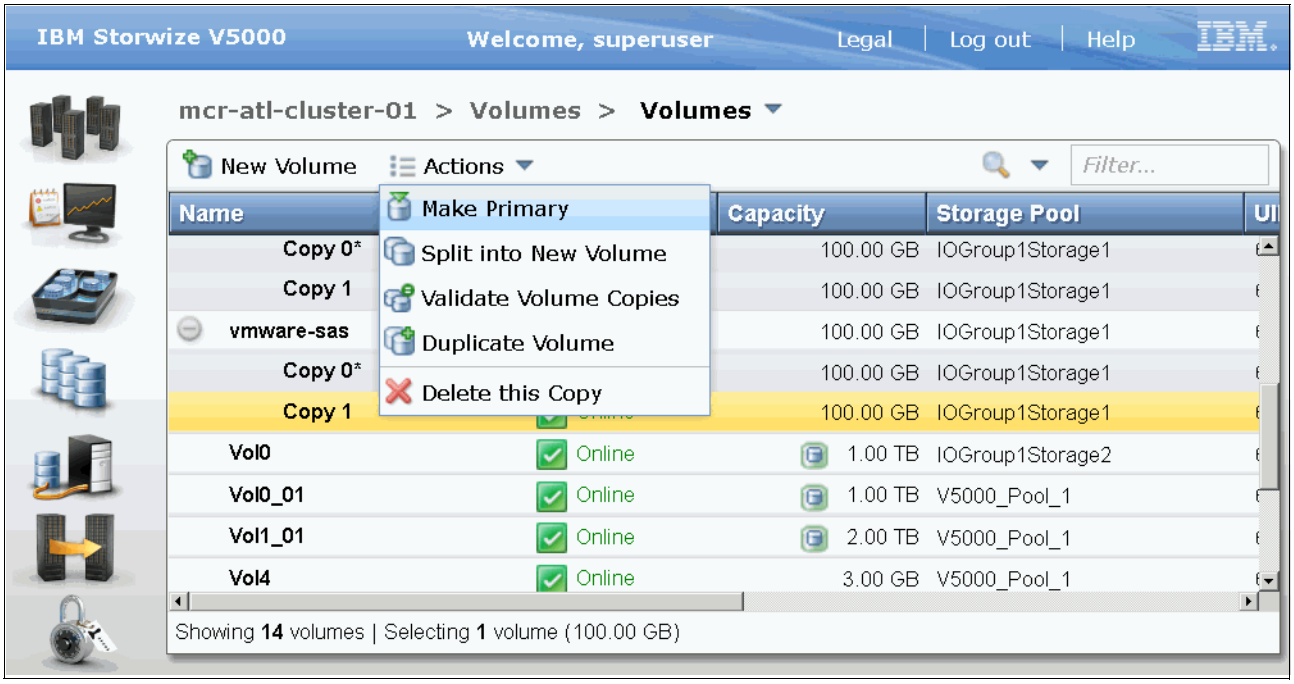

Figure 8-77 Make primary

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.

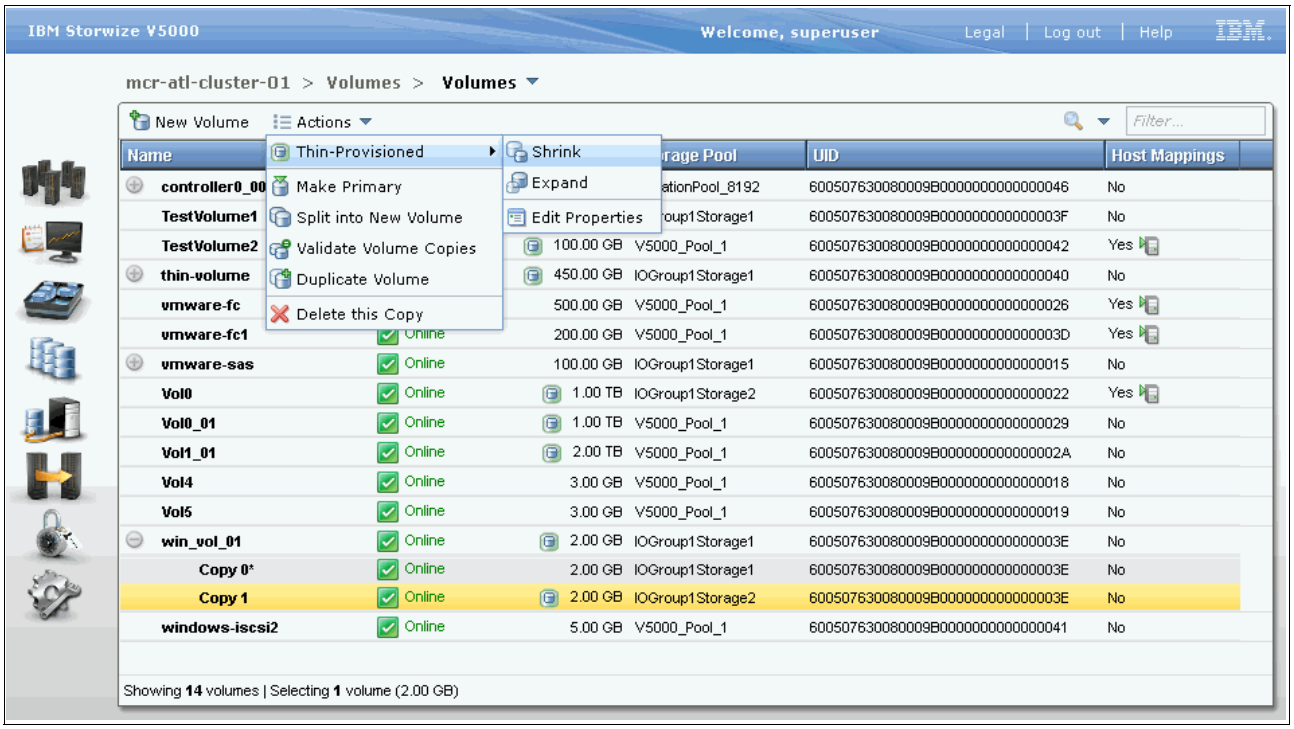

Figure 8-78 Thin-Provisioned menu item

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.

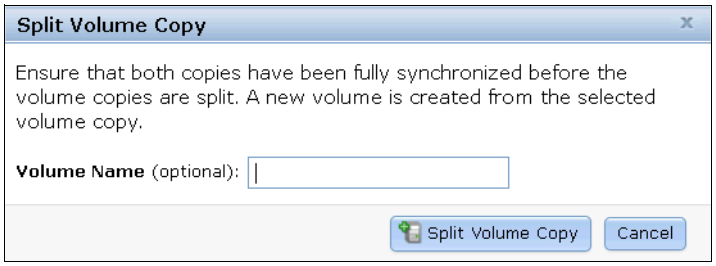

Figure 8-79 Split Volume Copy window

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.

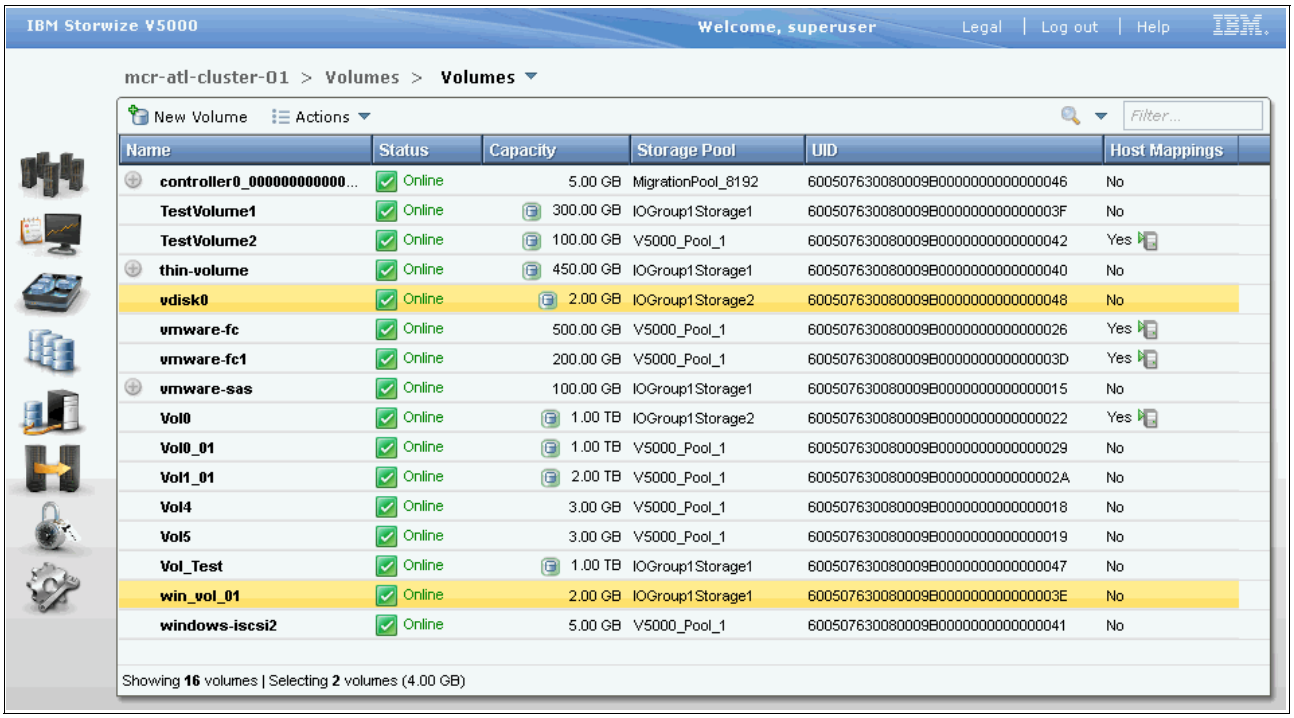

Figure 8-80 All Volumes: New volume from split copy

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.

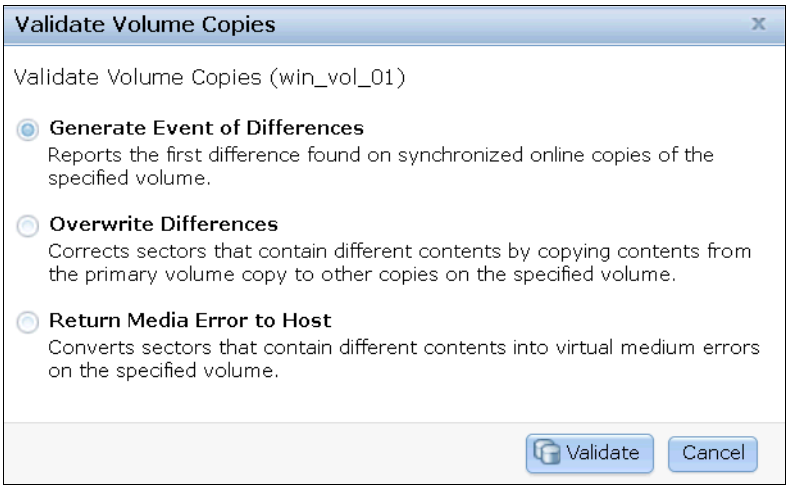

Figure 8-81 Validate Volume Copies window

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.

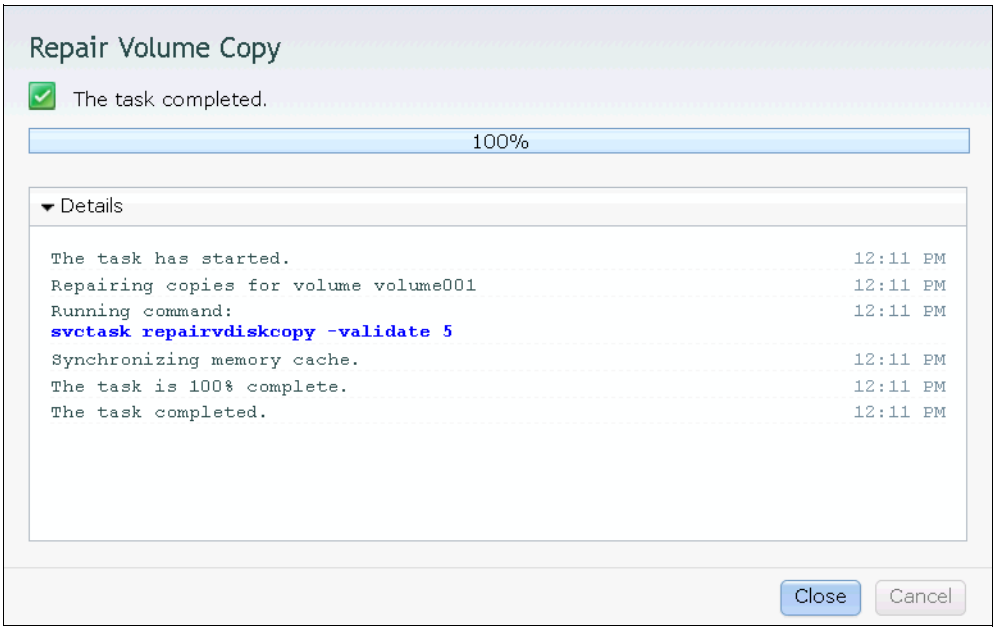

Figure 8-82 Volume copy validation output

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.

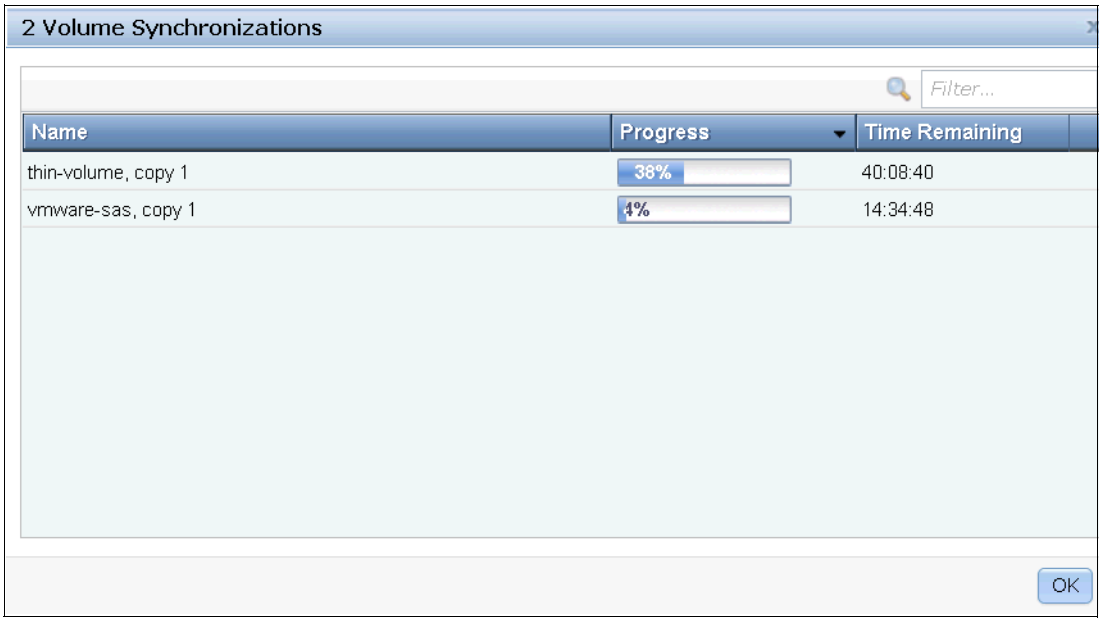

Figure 8-83 Validate Volume Copies: Running Tasks

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.

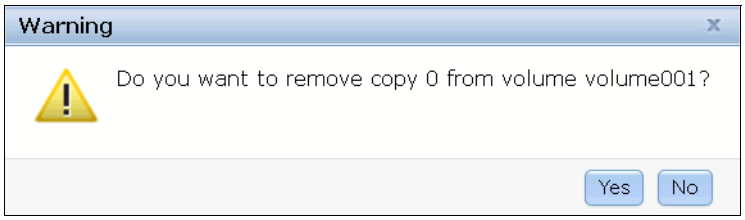

Figure 8-84 Delete a copy

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.

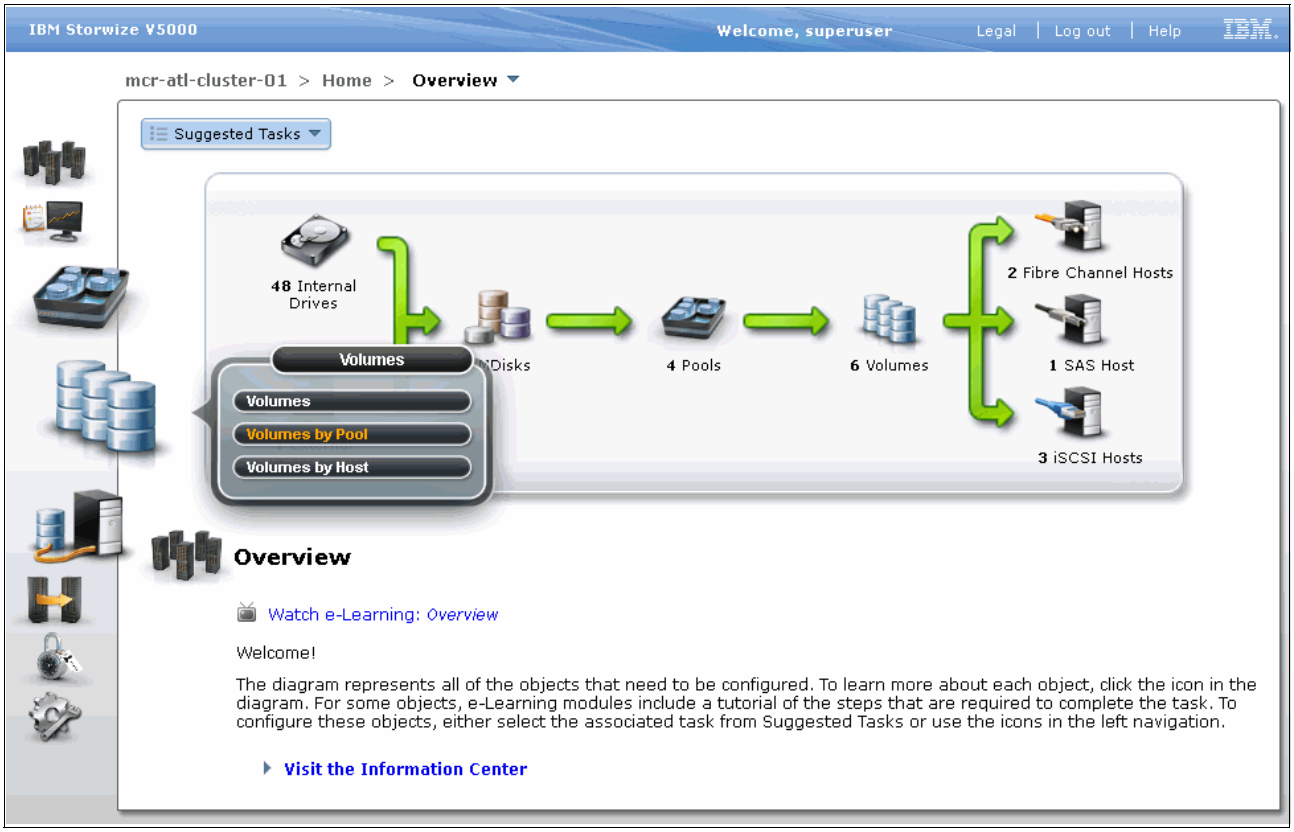

Figure 8-85 Volumes by Pool

The Volumes by Pool window opens, as shown in Figure 8-86.

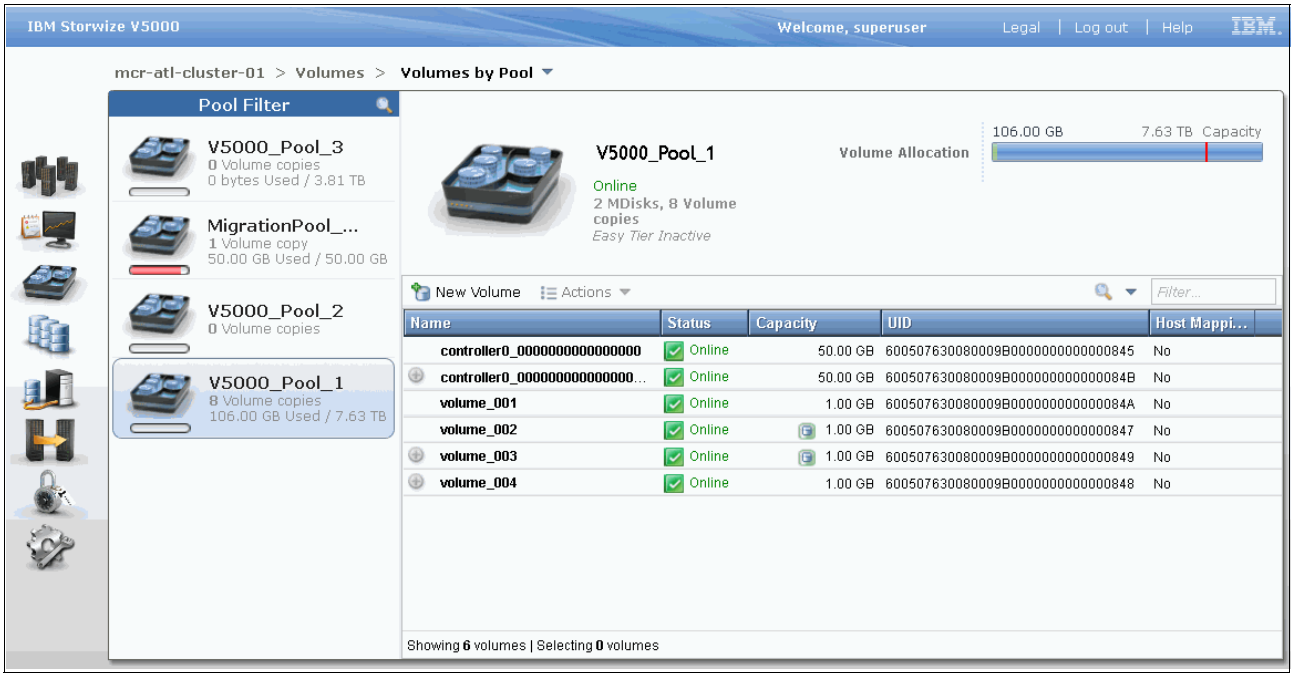

Figure 8-86 Volumes by Pool window

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.

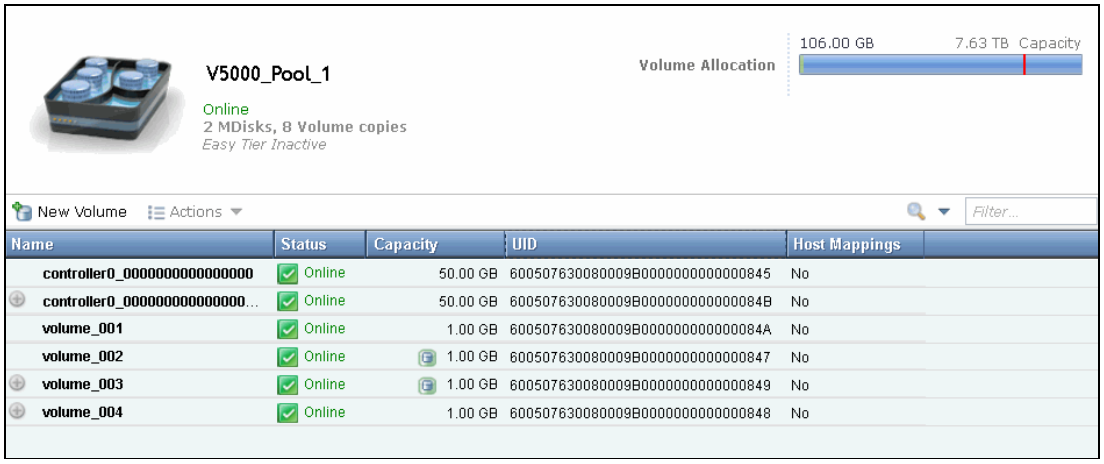

Figure 8-87 Volumes by Storage Pool

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.

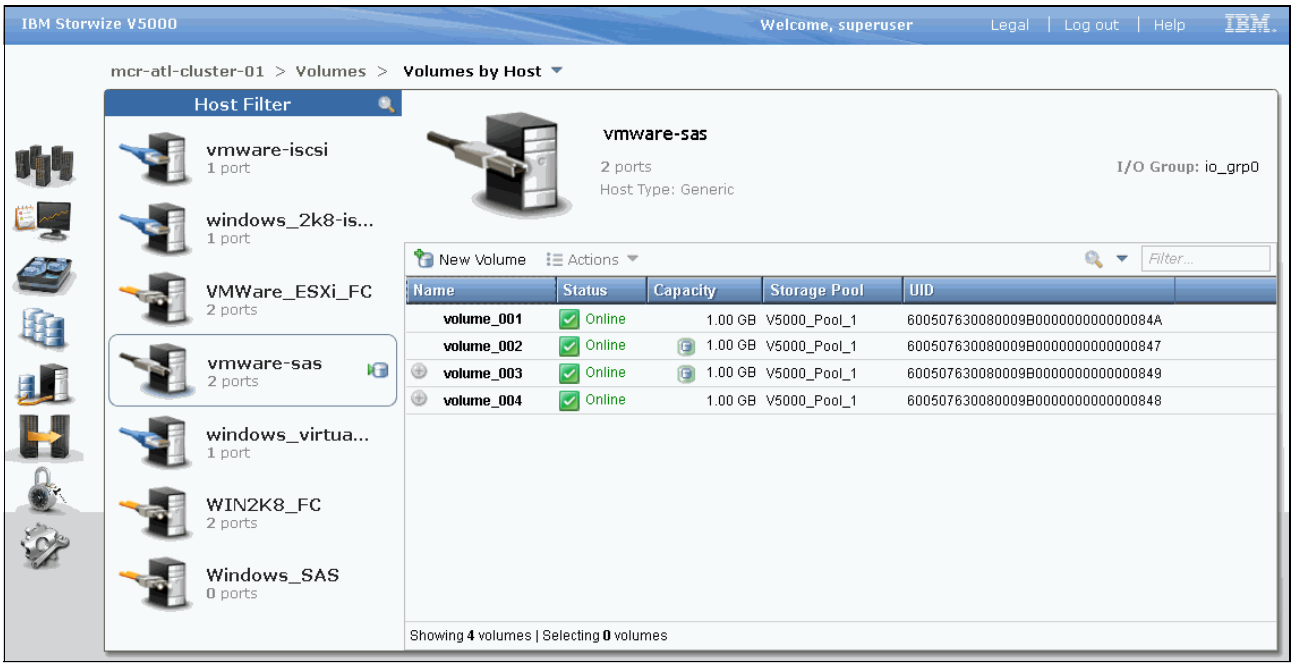

Figure 8-88 Volumes by host

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.

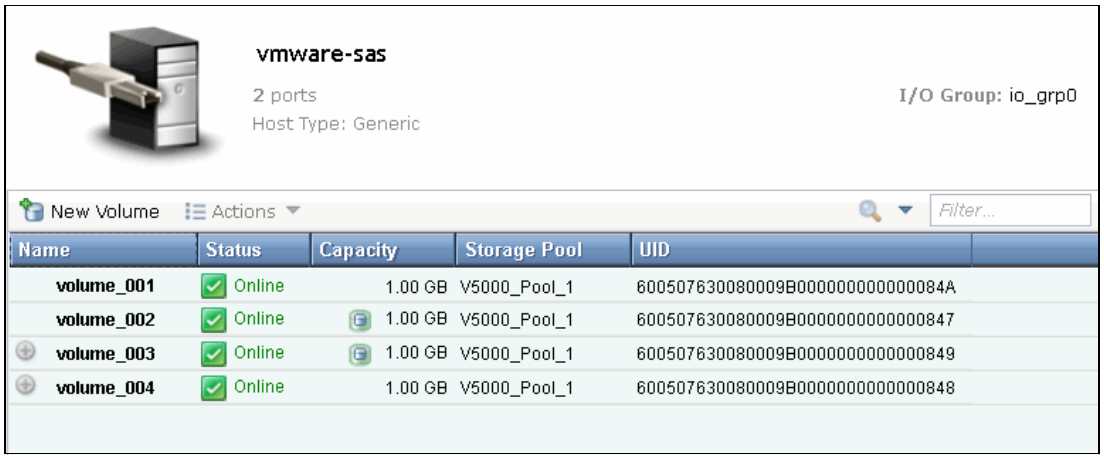

Figure 8-89 Volumes by Host

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

9.1 Easy Tier overview

Easy Tier is an optional licensed feature of IBM Storwize V5000 that brings enterprise storage functions to the midrange segment. It enables automated subvolume data placement throughout different storage tiers to intelligently align the system with current workload requirements and to optimize the usage of SSDs. This functionality includes the ability to automatically and non-disruptively relocate data (at the extent level) from one tier to another tier in either direction to achieve the best available storage performance for your workload in your environment.

Easy Tier reduces the I/O latency for hot spots, but it does not replace storage cache. Easy Tier and storage cache solve a similar access latency workload problem, but these methods weigh differently in the algorithmic construction that is based on “locality of reference”, recency, and frequency. Because Easy Tier monitors I/O performance from the device end (after cache), it can pick up the performance issues that cache cannot solve and complement the overall storage system performance.

In general, the storage environments I/O is monitored on volumes and the entire volume is always placed inside one appropriate storage tier. Determining the amount of I/O is too complex for monitoring I/O statistics on single extents, moving them manually to an appropriate storage tier, and reacting to workload changes.

Easy Tier is a performance optimization function that overcomes this issue because it automatically migrates (or moves) extents that belong to a volume between different storage tiers, as shown in Figure 9-1. Because this migration works at the extent level, it is often referred to as *sub-LUN migration*.

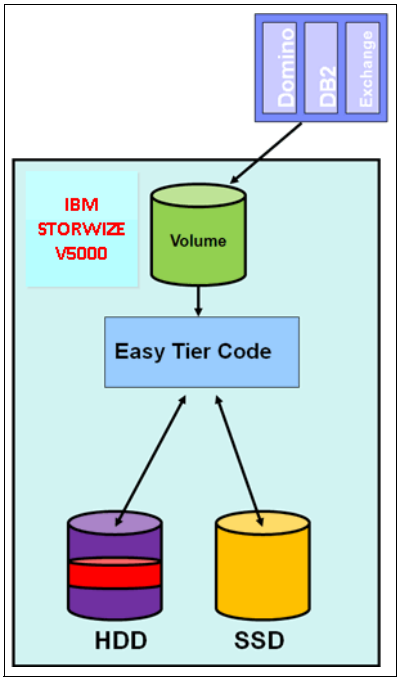

Figure 9-1 Easy Tier

You can enable Easy Tier for storage on a volume basis. It monitors the I/O activity and latency of the extents on all Easy Tier enabled volumes over a 24-hour period. Based on the performance log, it creates an extent migration plan and dynamically moves high activity or hot extents to a higher disk tier within the same storage pool. It also moves extents whose activity dropped off (or cooled) from higher disk tier MDisks back to a lower tier MDisk.

To enable this migration between MDisks with different tier levels, the target storage pool must consist of different characteristic MDisks. These pools are named *multitiered storage pools*. IBM Storwize V5000 Easy Tier is optimized to boost the performance of storage pools that contain HDDs and SSDs.

To identify the potential benefits of Easy Tier in your environment before actually installing higher MDisk tiers (such as, SSDs), it is possible to enable the Easy Tier monitoring volumes in single-tiered storage pools. Although the Easy Tier extent migration is not possible within a single-tiered pool, the Easy Tier statistical measurement function is possible. Enabling Easy Tier on a single-tiered storage pool starts the monitoring process and logs the activity of the volume extents. In this case, Easy Tier creates a migration plan file that can then be used to show a report on the number of extents that is appropriate for migration to higher level MDisk tiers, such as, SSDs.

The STAT is a no-cost tool that helps you to analyze this data. If you do not have an IBM Storwize V5000, use Disk Magic to get a better idea about the required number of SSDs that is appropriate for your workload. If you do not have any workload performance data, a good starting point can be to add approximately 5% of net capacity of SSDs to your configuration. However, this ratio is heuristics-based and changes according to different applications or different disk tier performance in each configuration. For database transactions, a ratio of fast SAS or Fibre Channel (FC) drives to SSD is about 6:1 to achieve the optimal performance, but this ratio depends on the environment on which it is implemented.

Easy Tier is available for IBM Storwize V5000 internal volumes and volumes on external virtualized storage subsystems because the SSDs can be internal or external drives. However, from the fabric point of view, it is a best practice to use SSDs inside the IBM Storwize V5000 (even if the lower tiered disk pool is on external storage) because this configuration reduces the traffic that is traversing the SAN environment.

9.2 Easy Tier for IBM Storwize V5000

This section describes the terms and gives an example implementation of Easy Tier on the IBM Storwize V5000. After reading this section, you should understand the Easy Tier concept as it relates to the IBM Storwize V5000.

9.2.1 Disk tiers

It is likely that IBM Storwize V5000 internal disks and external disks have different performance attributes. As described in Chapter 7, “Storage pools” on page 295, without Easy Tier, it is a best practice to place drives with the same attributes (the number of revolutions per minute, size, and type) in the same storage pool, and not to intermix different drives with different attributes. This configuration is also valid for external MDisks that are grouped into storage pools. All internal HDDs and external MDisks are initially put into the generic_hdd tier by default. An internal SSD is identified as a high-performance tier MDisk by IBM Storwize V5000 and all external SSD MDisks must be changed to the high-performance tier, as described in Chapter 7, “Storage pools” on page 295.

9.2.2 Tiered storage pools

With IBM Storwize V5000, we must differentiate between the following types of storage pools:

- Single-tiered storage pools
- Multitiered storage pools

As shown in Figure 9-2, single-tiered storage pools include one type for disk tier attribute. Each disk should have the same size and performance characteristics. Multitiered storage pools are populated with two different disk tier attributes, which means high-performance tier SSDs and generic HDDs. A volume migration is when the complete volume is migrated from one storage pool to another storage pool. An Easy Tier data migration moves only extents inside the storage pool to different performance attributes.

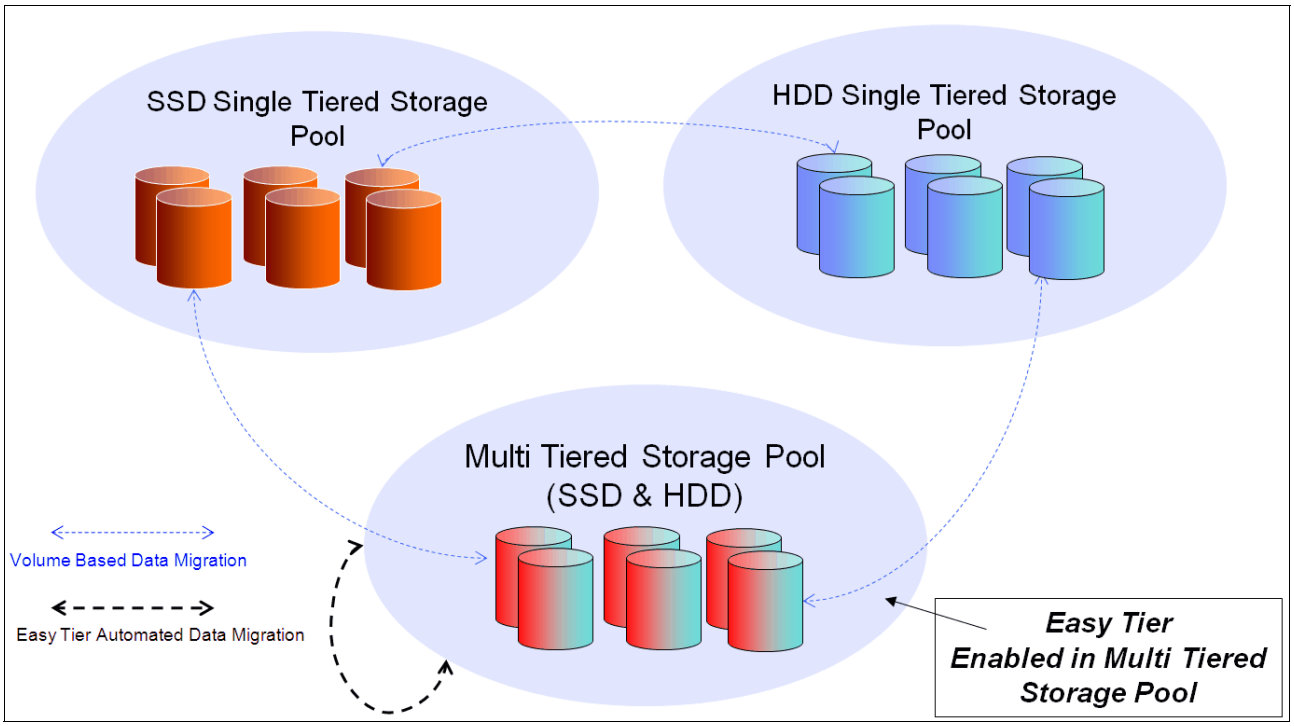

Figure 9-2 Tiered storage pools

9.3 Easy Tier process

The Easy Tier feature consists of four main processes. Figure 9-3 shows the flow between these processes. These processes ensure that the extent allocation in multitiered storage pools is optimized for the best performance that was monitored on your workload in the last 24 hours. Statistics about extent usage are collected at five-minute intervals. A heat map is created every 24 hours that is used by the internal algorithms to generate a migration plan and a summary report. This migration plan contains information about which extents to promote to the upper tier or to demote to the lower tier, and the summary report is used by STAT. For more information, see 9.6, “IBM Storage Tier Advisor Tool” on page 434.

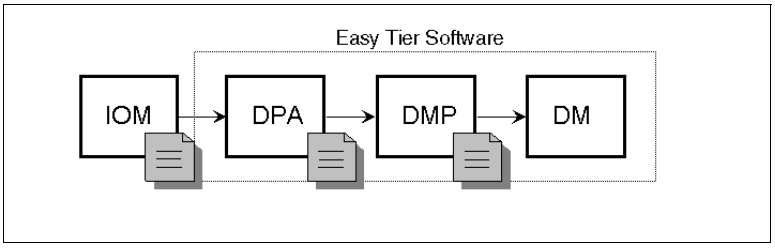

Figure 9-3 Easy Tier process flow

Easy Tier is based on an algorithm with a threshold to evaluate if an extent is cold or hot. If an extent activity is below this threshold, it is not considered by the algorithm to be moved to the SSD tier.

The four main processes and the flow between them are described in the following sections.

9.3.1 I/O Monitoring

The I/O Monitoring (IOM) process operates continuously and monitors host volumes for I/O activity. It collects performance statistics for each extent and derives averages for a rolling 24-hour period of I/O activity.

Easy Tier makes allowances for large block I/Os and thus considers only I/Os of up to 64 KB as migration candidates.

This process is an efficient process and adds negligible processing impact to the IBM Storwize V5000 node canisters.

9.3.2 Data Placement Advisor

The Data Placement Advisor (DPA) uses workload statistics to make a cost benefit decision about which extents should be candidates for migration to a higher performance (SSD) tier.

This process also identifies extents that must be migrated back to a lower (HDD) tier.

9.3.3 Data Migration Planner

By using the previously identified extents, the Data Migration Planner (DMP) process builds the extent migration plan for the storage pool.

9.3.4 Data Migrator

The Data Migrator (DM) process involves scheduling and the actual movement or migration of the volume's extents up to, or down from, the high disk tier. The extent migration rate is capped to a maximum of up to 15 MBps. This rate equates to around 2 TB a day that is migrated between disk tiers, as shown in Figure 9-4.

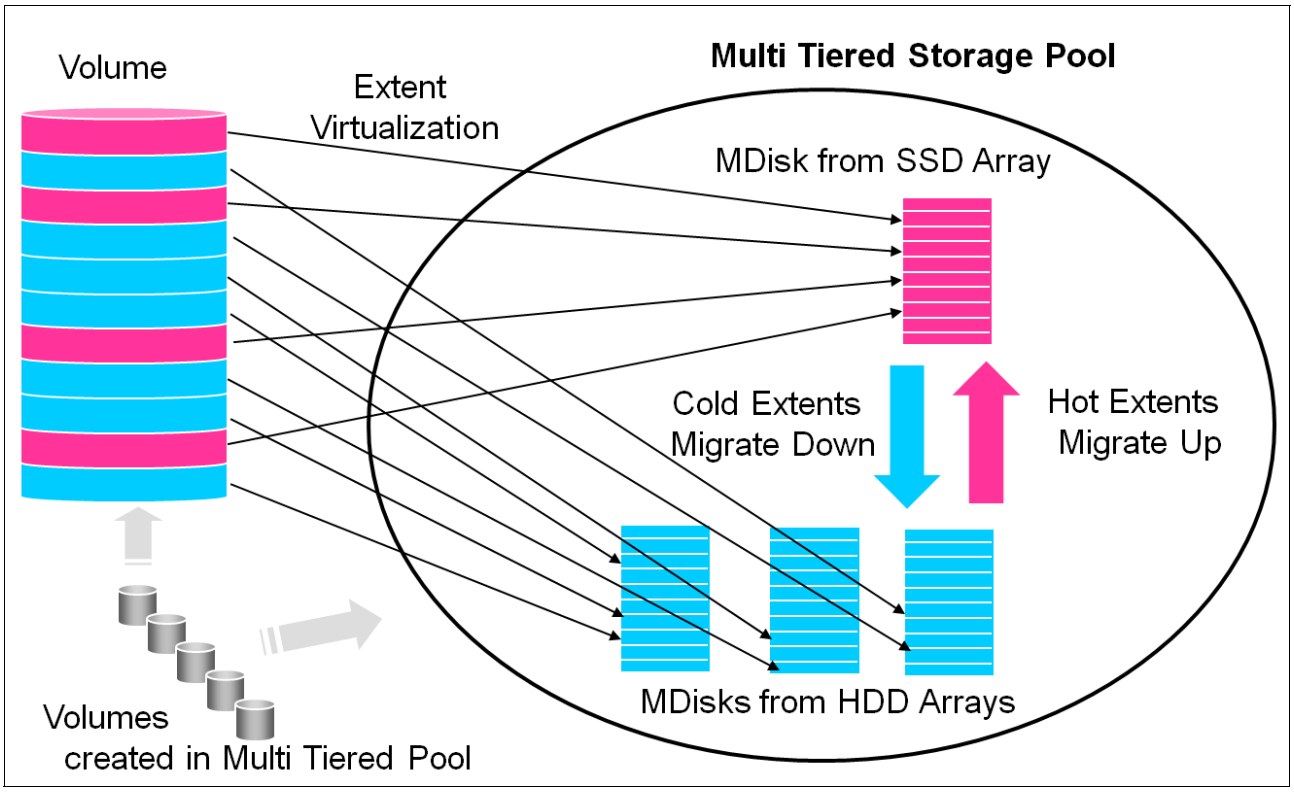

Figure 9-4 Easy Tier Data Migrator

9.3.5 Easy Tier operating modes

IBM Storwize V5000 offers the following operating modes for Easy Tier:

- Easy Tier: OFF

  Easy Tier can be turned off. No statistics are recorded and no extents are moved.

- Evaluation Mode

  If you turn on Easy Tier in a single-tiered storage pool, it runs in Evaluation Mode, which means it measures the I/O activity for all extents. A statistic summary file is created and can be offloaded from the IBM Storwize V5000. This file can be analyzed with the IBM Storage Tier Advisory Tool, as described in 9.6, "IBM Storage Tier Advisor Tool" on page 434. This analysis shows the benefits for your workload if you were to add SSDs to your pool before any hardware is acquired.

- Auto Data Placement Mode

  This operating mode is enabled by default if you create a multitiered storage pool. Easy Tier is also set to On to all volumes inside the multitiered storage pool. The extents are migrated dynamically by the Easy Tier processes to achieve the best performance. The movement is not apparent to the host server and applications, and it provides increased performance only.

If you do want to disable Auto Data Placement Mode for single volumes that are inside a multitiered storage pool, it is possible to turn off the mode at the volume level. This action excludes the volume from Auto Data Placement Mode and measures the I/O statistics only.

The statistic summary file can be offloaded for input to the advisor tool. The tool produces a report about the extents that are moved to SSD and a prediction of performance improvement that can be gained if more SSD was available.

9.3.6 Easy Tier rules

The following operating rules apply when IBM System Storage Easy Tier is used on the IBM Storwize V5000:

- Automatic data placement and extent I/O activity monitors are supported on each copy of a mirrored volume. Easy Tier works with each copy independently of the other copy.

  Volume mirroring: Volume mirroring can have different workload characteristics on each copy of the data because reads are normally directed to the primary copy and writes occur to both. Thus, the number of extents that Easy Tier migrates to SSD tier probably is different for each copy.

- Easy Tier works with all striped volumes, which include the following volumes:
  - Generic volumes
  - Thin-provisioned volumes
  - Mirrored volumes
  - Thin-mirrored volumes
  - Global and Metro Mirror sources and targets
- Easy Tier automatic data placement is not supported for image mode or sequential volumes. I/O monitoring for such volumes is supported, but you cannot migrate extents on such volumes unless you convert image or sequential volume copies to striped volumes.
- If possible, IBM Storwize V5000 creates volumes or volume expansions by using extents from MDisks from the HDD tier. Extents from MDisks from the SSD tier are used if no HDD space is available.
- When a volume is migrated out of a storage pool that is managed with Easy Tier, Automatic Data Placement Mode is no longer active on that volume. Automatic Data Placement is also turned off while a volume is migrated, even if it is between pools that both have Easy Tier Automatic Data Placement enabled. Automatic Data Placement for the volume is re-enabled when the migration is complete.
- SSD performance is dependent on block sizes (small blocks perform much better than larger blocks). Because Easy Tier is optimized to work with SSD, it decides whether an extent is hot by measuring I/O smaller than 64 KB, but it migrates the entire extent to the appropriate disk tier.
- As extents are migrated, the use of smaller extents makes Easy Tier more efficient.
- The first migration of hot data to SSD starts about one hour after Automatic Data Placement Mode is enabled. It takes up to 24 hours to achieve optimal performance.
- In the current IBM Storwize V5000 Easy Tier implementation, it takes about two days before hot spots are considered moved from SSDs, which prevents hot spots from being moved from SSDs if the workload changes over a weekend.
- If you run an unusual workload over a longer period, Automatic Data Placement can be turned off and on online to avoid data movement.

Depending on which storage pool and which Easy Tier configuration is set, a volume copy can have the Easy Tier states that are shown in Table 9-1.

Table 9-1 Easy Tier states

Storage pool	Single-tiered or multitiered storage pool	Volume copy Easy Tier setting	Easy Tier status
Off	Single-tiered	Off	Inactive[a]
Off	Single-tiered	On	Inactive[a]
Off	Multitiered	Off	Inactive[a]
Off	Multitiered	On	Inactive[a]
Auto[b]	Single-tiered	Off	Inactive[a]
Auto[b]	Single-tiered	On	Inactive[a]
Auto[b]	Multitiered	Off	Measured[c]
Auto[b]	Multitiered	On	Active[d] [e]
On	Single-tiered	Off	Measured[c]
On	Single-tiered	On	Measured[c]
On	Multitiered	Off	Measured[c]
On	Multitiered	On	Active[d]

a. When the volume copy status is inactive, no Easy Tier functions are enabled for that volume copy.

b. The default Easy Tier setting for a storage pool is Auto, and the default Easy Tier setting for a volume copy is On. This scenario means that Easy Tier functions are disabled for storage pools with a single tier, and that automatic data placement mode is enabled for all striped volume copies in a storage pool with two tiers.

c. When the volume copy status is measured, the Easy Tier function collects usage statistics for the volume, but automatic data placement is not active.

d. If the volume copy is in image or sequential mode or is being migrated, the volume copy Easy Tier status is measured instead of active.

e. When the volume copy status is active, the Easy Tier function operates in automatic data placement mode for that volume.

9.4 Easy Tier configuration by using the GUI

This section describes how to activate Easy Tier by using the IBM Storwize V5000 GUI.

9.4.1 Creating multitiered pools: Enable Easy Tier

In this section, we describe how to create multitiered storage pools by using the GUI. When a storage pool changes from single-tiered to multitiered, Easy Tier is enabled by default for the pool and on all volume copies inside this pool.

To create multitiered pools, complete the following steps:

1. Click **Volumes** → **Volumes by Pool**. Figure 9-5 shows that the `V5000_Pool_1` storage pool exists and Easy Tier is inactive in our example. An SSD must be added to the `V5000_Pool_2` storage pool to enable Easy Tier.

Figure 9-5 Single-tiered pool

2. Click **Pools** → **Internal Storage**. Figure 9-6 shows that one internal SSD is available and it is in the Unused status. Internal SSDs are assigned the generic_ssd high performance tier automatically by the IBM Storwize V5000.

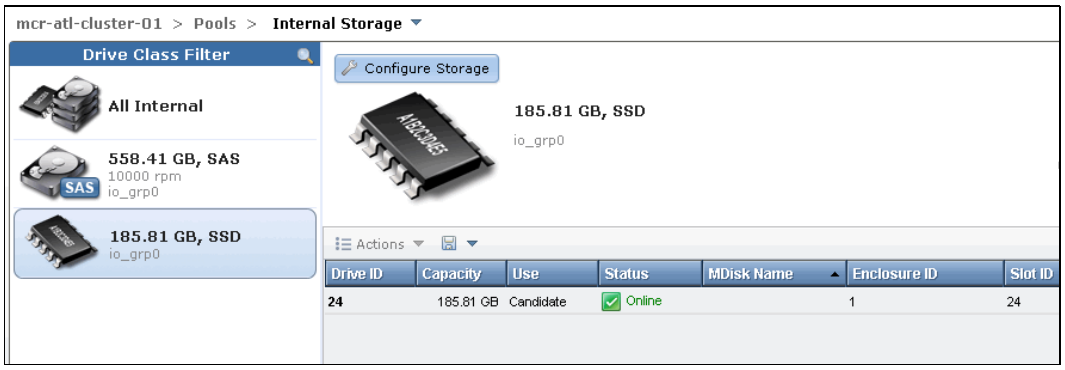

Figure 9-6 Internal SSDs

3. Click **Configure Storage** and the Storage Configuration wizard opens. Because the SSD is in Unused status, an information message is displayed, as shown in Figure 9-7. Click **Yes** to proceed to the next window.

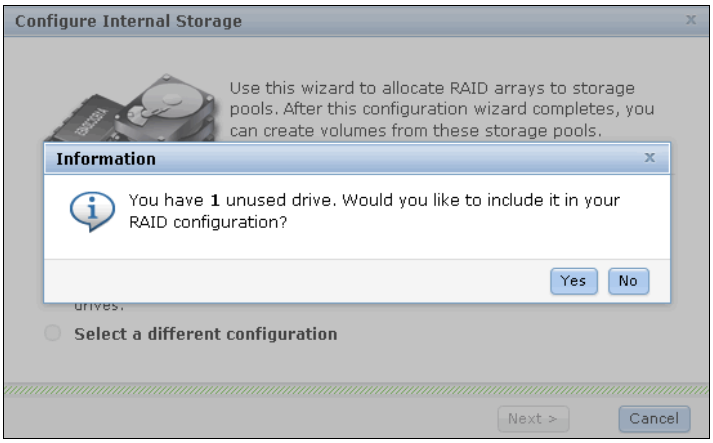

Figure 9-7 Unused drives information message

4. The IBM Storwize V5000 configures the newly detected drive and changes its status to Candidate after the task is completed, as shown in Figure 9-8. Click **Close** to start the SSD configuration.

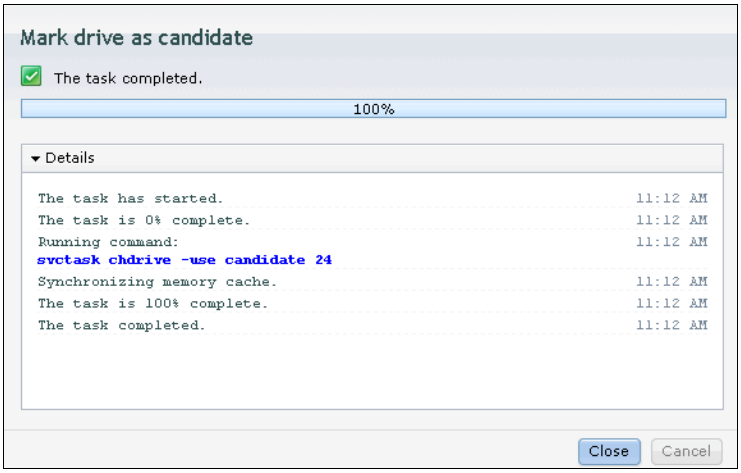

Figure 9-8 Drive configuration output

Figure 9-9 shows the first step of the Configuration wizard.

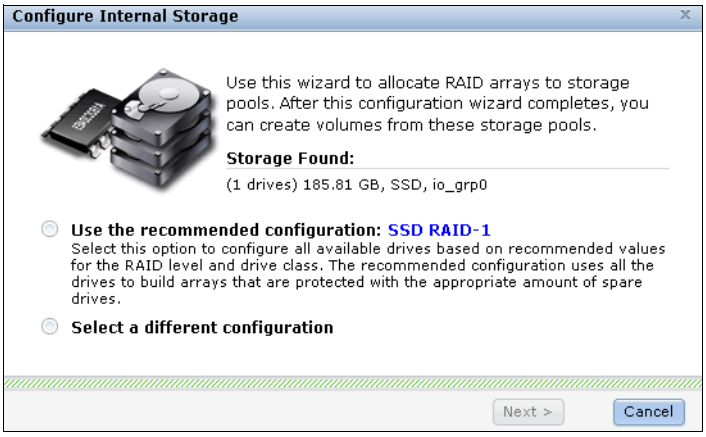

Figure 9-9 Configure Internal Storage window

The wizard recommends the use of the SSD to enable Easy Tier. If you select **Use recommended configuration**, it selects the recommended RAID level and hot spare coverage for your system automatically, as shown in Figure 9-10.

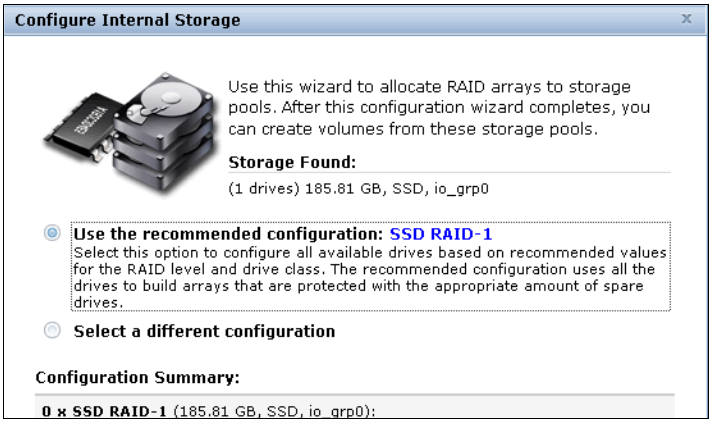

Figure 9-10 Recommended configuration

If you select **Select a different configuration** (as shown in Figure 9-11), you can select the preset.

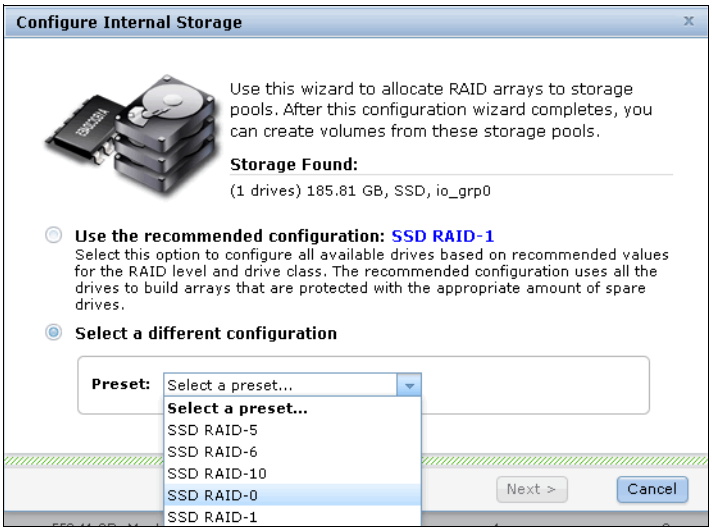

Figure 9-11 Select a preset menu

5. Choose a custom RAID level, or you can also select the SSD Easy Tier preset to review and modify the recommended configuration. Because we do not have enough drives in our configuration, the SSD Easy Tier preset is not available from the preset selection. When it is available, this preset configures a RAID 10 array with a spare goal of one drive. In this example, we create a RAID 0 array, although this is not best practice and is not used in a production environment. Because there are not enough drives, an error message is displayed, as shown in Figure 9-12 on page 423.

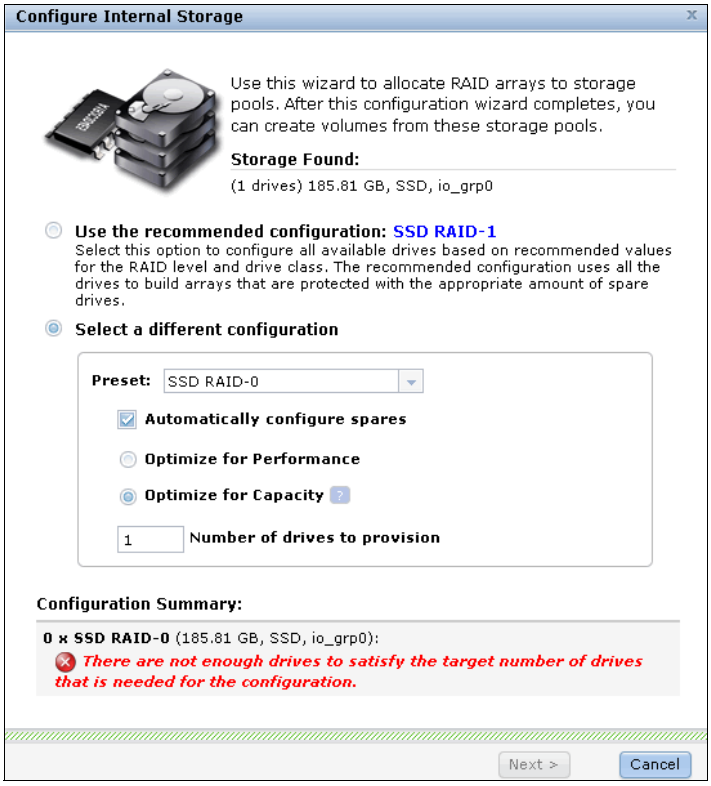

Figure 9-12 Select RAID 0 preset

This error message can be avoided if the Automatically configure spares option is not chosen, as shown in Figure 9-13. A RAID 0 array with one drive and zero spares is created.

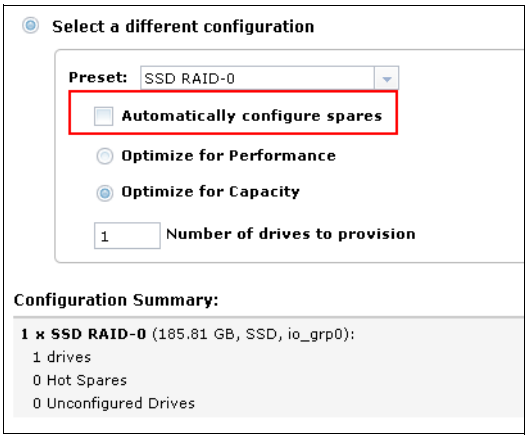

Figure 9-13 Array creation configuration summary

6. To create a multitiered storage pool, the SSD must be added to an existing generic HDD pool. Select **Expand an existing pool** (as shown in Figure 9-14) and select the pool that you want to change to a multitiered storage pool. In our example, `V5000_Pool_2` is selected. Click **Finish**.

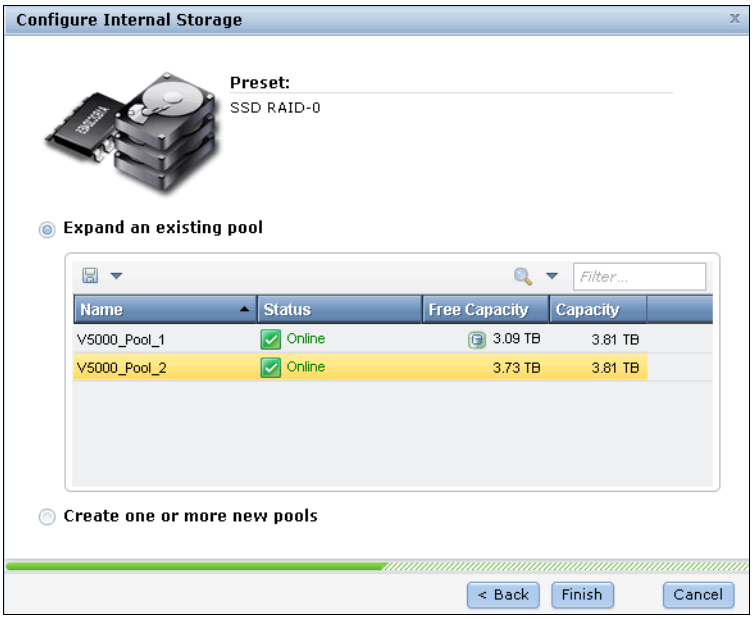

Figure 9-14 Expand an existing pool

7. Now the array is configured on the SSDs and added to the selected storage pool. Click **Close** after the task completes, as shown in Figure 9-15.

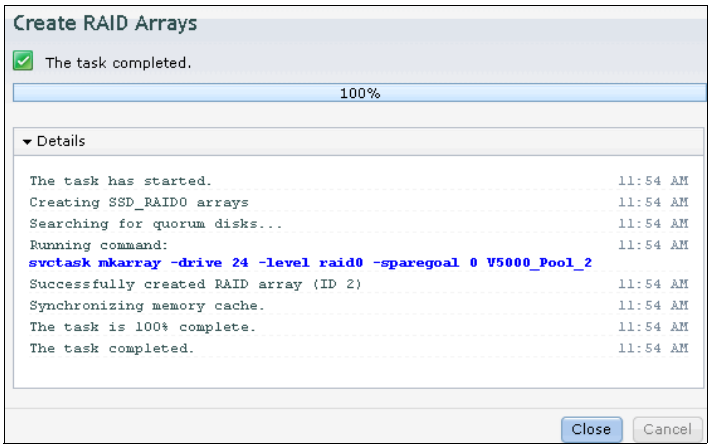

Figure 9-15 Array creation completed task

Figure 9-16 shows that the internal SSDs usage changed to Member and that the wizard created an MDisk that is named `mdisk2`.

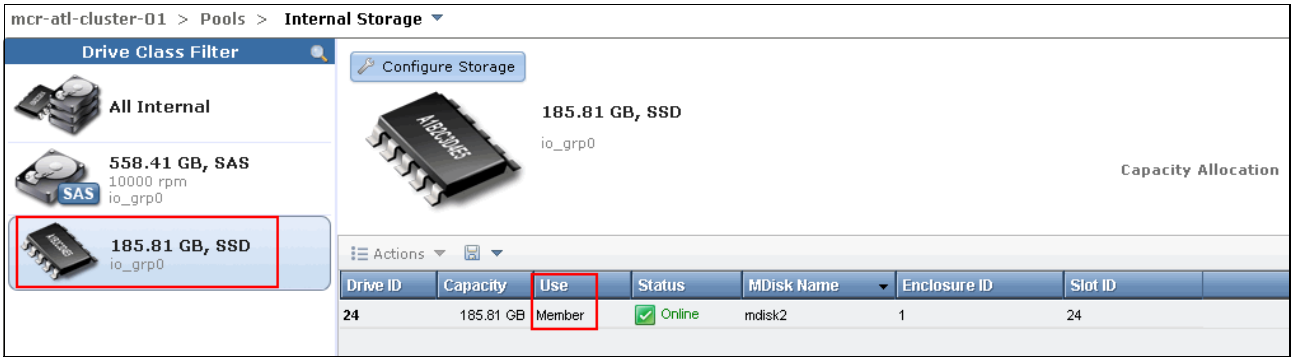

Figure 9-16 SSD usage is changed

In Figure 9-17, you see that the new MDisk is now part of the `V5000_Pool_2` storage pool and that the status of the Easy Tier changed to Active. In this pool, Automatic Data Placement Mode is started and the Easy Tier processes start to work.

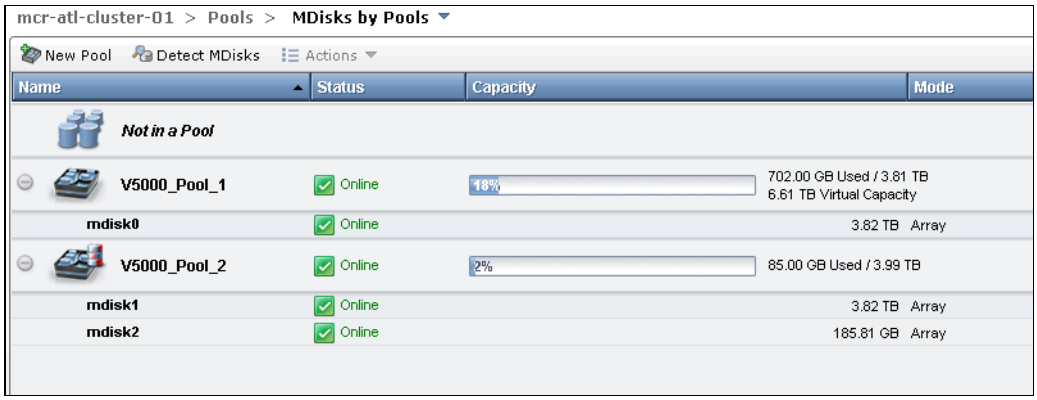

Figure 9-17 Easy Tier active

The storage pool was successfully changed to a multitiered storage pool (as indicated by the icon:) and Easy Tier was activated by default. To reflect this change, we renamed the storage pool and changed the function icon, as described in Chapter 7, “Storage pools” on page 295 and as shown in Figure 9-18.

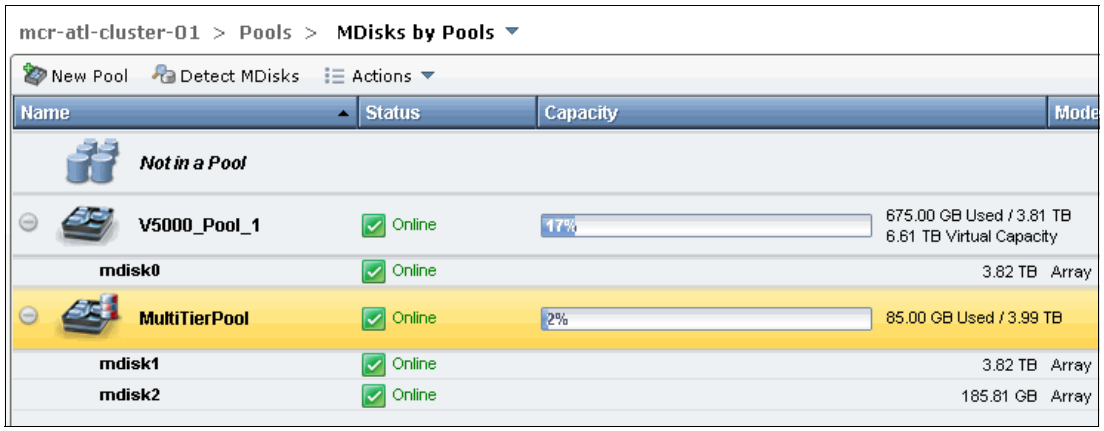

Figure 9-18 Multitiered storage pool

By default, Easy Tier is now active in this storage pool and all its volumes. Figure 9-19 shows three volumes on the multitiered storage pool.

Figure 9-19 Volumes by Pool

If you open the properties of a volume by clicking **Actions** → **Properties**, you can also see that Easy Tier is enabled on the volume by default, as shown in Figure 9-20.

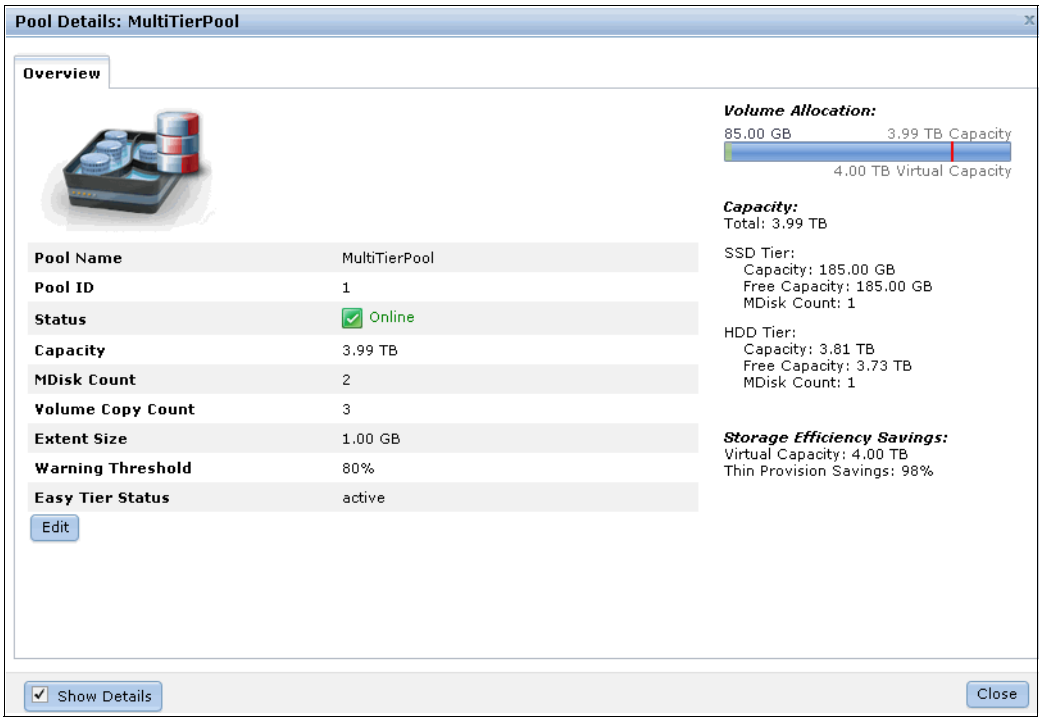

Figure 9-20 Easy Tier enabled volume

If a volume has more than one copy, Easy Tier can be enabled and disabled on each copy separately. This action depends on the storage pool where the volume copy is defined. You can see a volume with two copies that are stored in two different storage pools, as shown in Figure 9-21.

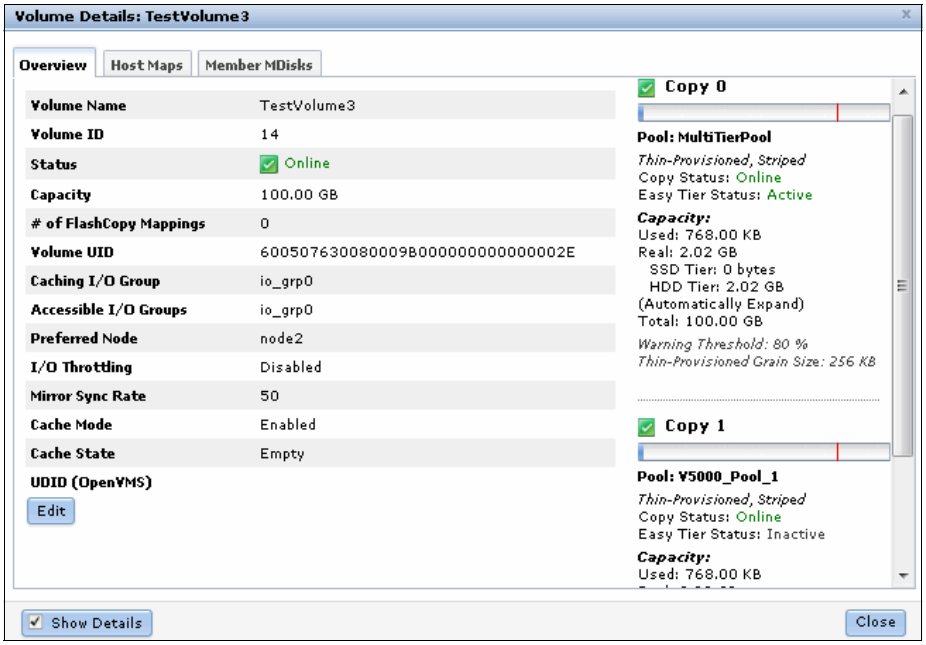

Figure 9-21 Easy Tier by Copy

If you want to enable Easy Tier on the second copy, change the storage pool of the second copy to a multitiered storage pool by repeating these steps.

If an external SSD is used, you must select the tier manually, and then add the external SSD MDisk to a storage pool, as described in as described in Chapter 7, “Storage pools” on page 295. This action also changes the storage pools to multitiered storage pools and enables Easy Tier on the pool and the volumes.

9.4.2 Downloading Easy Tier I/O measurements

Easy Tier is now enabled and Automatic Data Placement Mode is active. Extents are automatically migrated to, or from, high -performance disk tiers, and the statistic summary collection is now active. The statistics log file can be downloaded to analyze how many extents were migrated, and to monitor if it makes sense to add more SSDs to the multitiered storage pool.

To download the statistics file, complete the following steps:

1. Click **Settings** → **Support**, as shown in Figure 9-22.

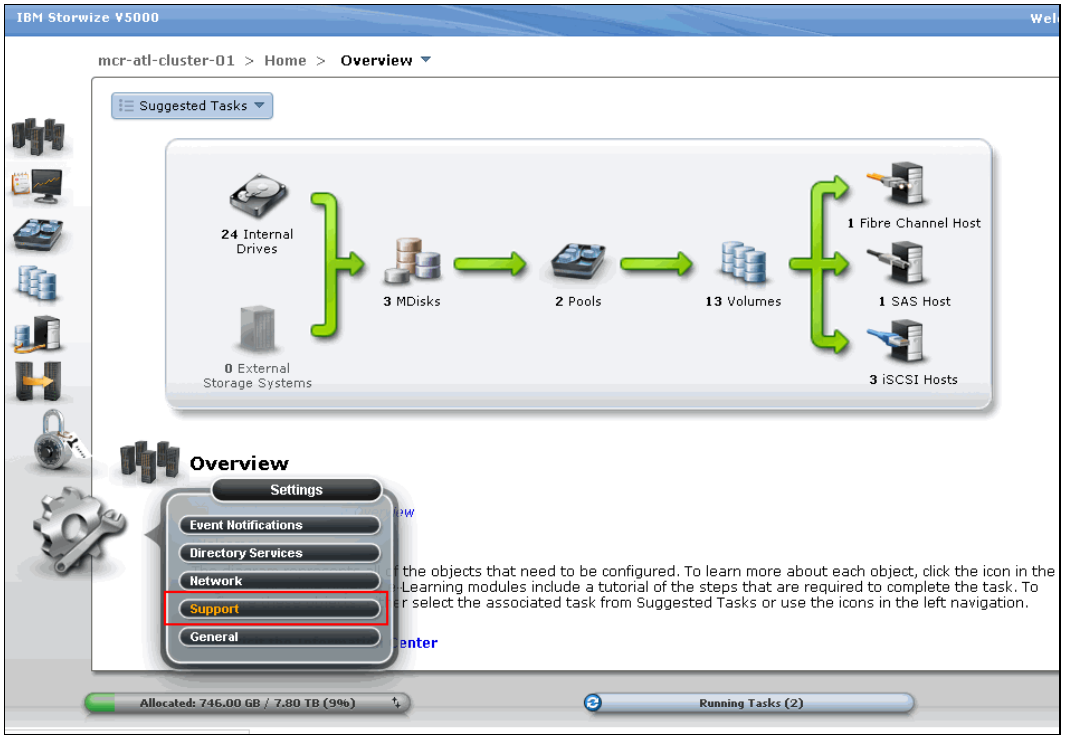

Figure 9-22 Settings menu

2. Click **Show full log listing**, as shown in Figure 9-23.

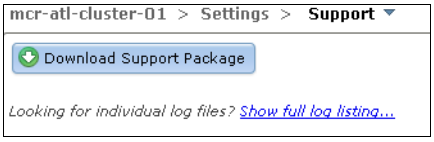

Figure 9-23 Download files menu

This action lists all the log files available to download, as shown in Figure 9-24. The Easy Tier log files are always named `dpa_heat.canister_name_date.time.data`.

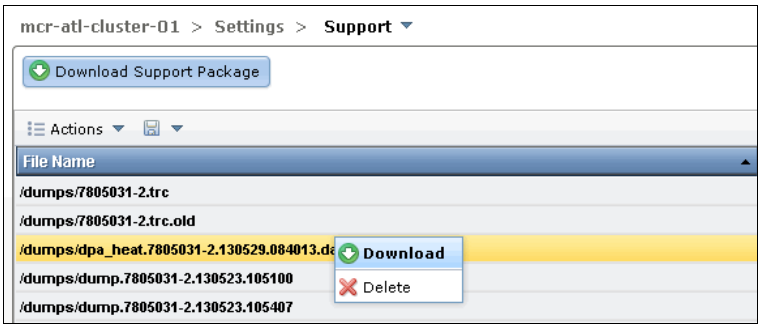

Figure 9-24 Download dpa_heat file

Log file creation: Depending on your workload and configuration, it can take up to 24 hours until a new Easy Tier log file is created.

If you run Easy Tier for a longer period, it generates a heat file at least every 24 hours.

The time and date of the file creation is included in the file name. The heat log file always includes the measured I/O activity of the last 24 hours.

3. Right-click the **dpa_heat.canister_name_date.time.data** file and click **Download**. Select the file for Easy Tier measurement for the most representative period.

   You can also use the search field that is on the right to filter your search, as shown in Figure 9-25.

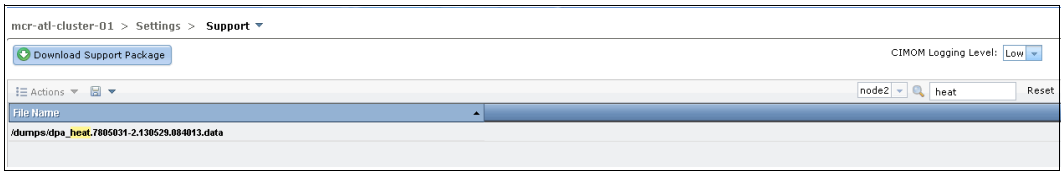

Figure 9-25 Filter your search

Depending on your browser settings, the file is downloaded to your default location, or you are prompted to save it to your computer. This file can be analyzed as described in 9.6, “IBM Storage Tier Advisor Tool” on page 434.

9.5 Easy Tier configuration by using the command-line interface

The process that is used to enable IBM Storwize V5000 Easy Tier by using the GUI is described in 9.4, “Easy Tier configuration by using the GUI” on page 419. Easy Tier can also be configured by using the command-line interface (CLI). For the advanced user, this method offers several more options for Easy Tier configuration.

Before you use the CLI, you must configure CLI access, as described in Appendix A, “Command-line interface setup and SAN Boot” on page 609.

Readability: In most examples that are shown in this section, many lines were deleted in the command output or responses so that we can concentrate only on the information that is related to Easy Tier.

9.5.1 Enabling Easy Tier evaluation mode

If you want to enable Easy Tier in evaluation mode, you must enable Easy Tier on a single-tiered storage pool. Connect to your IBM Storwize V5000 by using the CLI and run the **lsmdiskgrp** command, as shown in Example 9-1 on page 430. This command shows an overview about all configured storage pools and the Easy Tier status of the pool. In our example, there are two storage pools listed: `mdiskgrp0` with Easy Tier inactive, and `Multi_Tier_Pool` with Easy Tier enabled.

Example 9-1 List storage pools

```
IBM_2078:admin>lsmdiskgrp
id name            status mdisk_count ... easy_tier easy_tier_status
0  mdiskgrp0       online 3           ... auto      inactive
1  Multi_Tier_Pool online 3           ... auto      active
```

To get a more detailed view of the single-tiered storage pool, run the **lsmdiskgrp *storage pool name*** command, as shown in Example 9-2.

Example 9-2 Storage Pools details: Easy Tier inactive

```
IBM_2078:admin>lsmdiskgrp mdiskgrp0
id 0
name mdiskgrp0
status online
mdisk_count 3
...
easy_tier auto
easy_tier_status inactive
tier generic_ssd
tier_mdisk_count 0
...
tier generic_hdd
tier_mdisk_count 3
...
```

To enable Easy Tier on a single-tiered storage pool, run the **chmdiskgrp -easytier on *storage pool name*** command, as shown in Example 9-3. Because this storage pool does not have any SSD MDisks, it is not a multitiered storage pool; only measuring is available.

Example 9-3 Enable Easy Tier on a single-tiered storage pool

```
IBM_2078:admin>chmdiskgrp -easytier on mdiskgrp0
IBM_2078:admin>
```

Check the status of the storage pool again by running the **lsmdiskgrp *storage pool name*** command again, as shown in Example 9-4.

Example 9-4 Storage pool details: Easy Tier ON

```
IBM_2078:admin>lsmdiskgrp mdiskgrp0
id 0
name mdiskgrp0
status online
mdisk_count 3
vdisk_count 7
...
easy_tier on
easy_tier_status active
tier generic_ssd
tier_mdisk_count 0
...
tier generic_hdd
tier_mdisk_count 3
...
```

Run the **svcinfo lsmdiskgrp** command again, as shown in Example 9-5. You see that Easy Tier is turned on the storage pool now, but Automatic Data Placement Mode is not active on the multitiered storage pool.

Example 9-5 Storage pool list

```
IBM_2078:admin>lsmdiskgrp
id name            status mdisk_count vdisk_count ... easy_tier easy_tier_status
0  mdiskgrp0       online 3           7           ... on        active
1  Multi_Tier_Pool online 3           0           ... auto      active
```

To get the list of all the volumes defined, run the **lsvdisk** command, as shown in Example 9-6. For this example, we are only interested in the `redhat1` volume.

Example 9-6 All volumes list

```
IBM_2078:admin>lsvdisk
id name    IO_group_id IO_group_name status       mdisk_grp_id
5  redhat1 0           io_grp0       online many  many
...
```

To get a more detailed view of a volume, run the **lsvdisk *volume name*** command, as shown in Example 9-7. This output shows two copies of a volume: Copy 0 is in a multitiered storage pool and Automatic Data Placement is active, Copy 1 is in the single-tiered storage pool, and Easy Tier evaluation mode is active, as indicated by the `easy_tier_status measured` line.

Example 9-7 Volume details

```
IBM_2078:admin>lsvdisk redhat1
id 5
name redhat1
IO_group_id 0
IO_group_name io_grp0
status online
mdisk_grp_id many
mdisk_grp_name many
capacity 10.00GB
...

copy_id 0
status online
sync yes
primary yes
mdisk_grp_id 1
mdisk_grp_name Multi_Tier_Pool
...
easy_tier on
easy_tier_status active
tier generic_ssd
tier_capacity 0.00MB
tier generic_hdd
tier_capacity 10.00GB
...

copy_id 1
status online
sync yes
```

```
primary no
mdisk_grp_id 0
mdisk_grp_name mdiskgrp0
....
easy_tier on
easy_tier_status measured
tier generic_ssd
tier_capacity 0.00MB
tier generic_hdd
tier_capacity 10.00GB
...
```

These changes are also reflected in the GUI, as shown in Figure 9-26. Select the **Show Details** option to view the details of the Easy Tier for each of the volume copies.

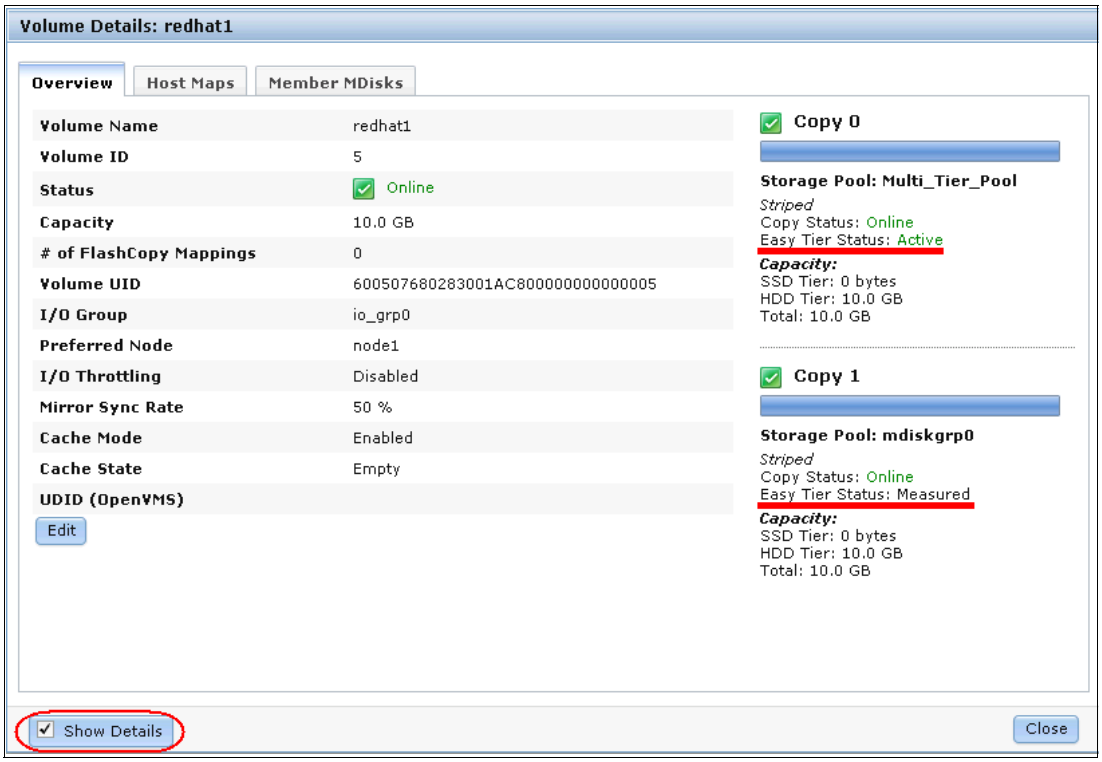

Figure 9-26 Easy Tier status by volume

Easy Tier evaluation mode is now active on the single-tiered storage pool (`mdiskgrp0`), but only for measurement. For more information about downloading and analyzing the I/O statistics, see 9.4.2, “Downloading Easy Tier I/O measurements” on page 427.

9.5.2 Enabling or disabling Easy Tier on single volumes

If you enable Easy Tier on a storage pool, all volume copies inside the Easy Tier pools also have Easy Tier enabled by default. This setting applies to multitiered and single-tiered storage pools. It is also possible to turn Easy Tier on and off for single volume copies.

To disable Easy Tier on single volumes, run the **chvdisk -easytier off *volume name*** command, as shown in Example 9-8.

Example 9-8 Disable Easy Tier on a single volume

```
IBM_2078:admin>chvdisk -easytier off redhat1
IBM_2078:admin>
```

This command disables Easy Tier on all copies of this volume. Example 9-9 shows that the Easy Tier status of the copies did change, even if Easy Tier is still enabled on the storage pool.

Example 9-9 Easy Tier disabled

```
IBM_2078:admin>lsvdisk redhat1
id 5
name redhat1
IO_group_id 0
IO_group_name io_grp0
status online
mdisk_grp_id many
mdisk_grp_name many
capacity 10.00GB
...

copy_id 0
status online
sync yes
primary yes
mdisk_grp_id 1
mdisk_grp_name Multi_Tier_Pool
...
easy_tier off
easy_tier_status measured
tier generic_ssd
tier_capacity 0.00MB
tier generic_hdd
tier_capacity 10.00GB
...

copy_id 1
status online
sync yes
primary no
mdisk_grp_id 0
mdisk_grp_name mdiskgrp0
....
easy_tier off
easy_tier_status measured
tier generic_ssd
tier_capacity 0.00MB
tier generic_hdd
tier_capacity 10.00GB
...
```

To enable Easy Tier on a volume, run the **chvdisk -easytier on *volume name*** command (as show in Example 9-10), and the Easy Tier Status changes back to enabled (as shown in Example 9-7 on page 431).

Example 9-10 Easy Tier enabled

```
IBM_2078:admin>chvdisk -easytier on redhat1
IBM_2078:admin>
```

9.6 IBM Storage Tier Advisor Tool

The STAT is a Windows console tool. If you run Easy Tier in evaluation mode, the tool analyzes the extents and estimates how much benefit you derive if you implement Easy Tier Automatic Data Placement with SSD MDisks. If Automatic Data Placement Mode is already active, the analysis also includes an overview of migrated hot data and recommendations about whether you can derive any benefit by adding more SSD drives. The output provides a graphical representation of the performance data that is collected by Easy Tier over a 24-hour operational cycle.

9.6.1 Creating graphical reports

STAT takes input from the `dpa_heat` log file and produces an HTML file that contains the report. Download the `heat_log` file, as described in 9.4.2, “Downloading Easy Tier I/O measurements” on page 427, and save it to the HDD of a Windows system.

For more information about the tool and to download it, see this website:

`http://www-01.ibm.com/support/docview.wss?uid=ssg1S4000935`

Click **Start** → **Run**, enter `cmd`, and then click **OK** to open a command prompt.

Typically, the tool is installed in the `C:\Program Files\IBM\STAT` directory. Enter the command to generate the report, as shown in Example 9-11.

`C:\Program Files\IBM\STAT>STAT.exe -o c:\`*`directory_where_you_want_the output_to_go`*
`c:\location_of_dpa_heat_data_file`

If you do not specify `-o c:\directory_where_you_want_the output_to_go`, the output goes to the directory where the `STAT.exe` file is located.

Example 9-11 Generate HTML file

```
C:\EasyTier>STAT.exe -o C:\EasyTier C:\StorwizeV5000_Logs\dpa_heat.31G00KV-1.101
209.131801.data

CMUA00019I The STAT.exe command has completed.

C:\EasyTier>
```

Browse the directory where you directed the output file, and there is a file named `index.html`. Open the file by using your browser to view the report.

9.6.2 STAT reports

If you open the `index.html` file of an IBM Storwize V5000 system that is in Easy Tier evaluation mode, a window opens that gives you an estimate of the benefit if you were to add SSDs, as shown in Figure 9-27. The report shows a heading of IBM Storwize V7000. However, you can ignore this heading because this tool was originally available for IBM Storwize V7000 but also works with IBM Storwize V5000.

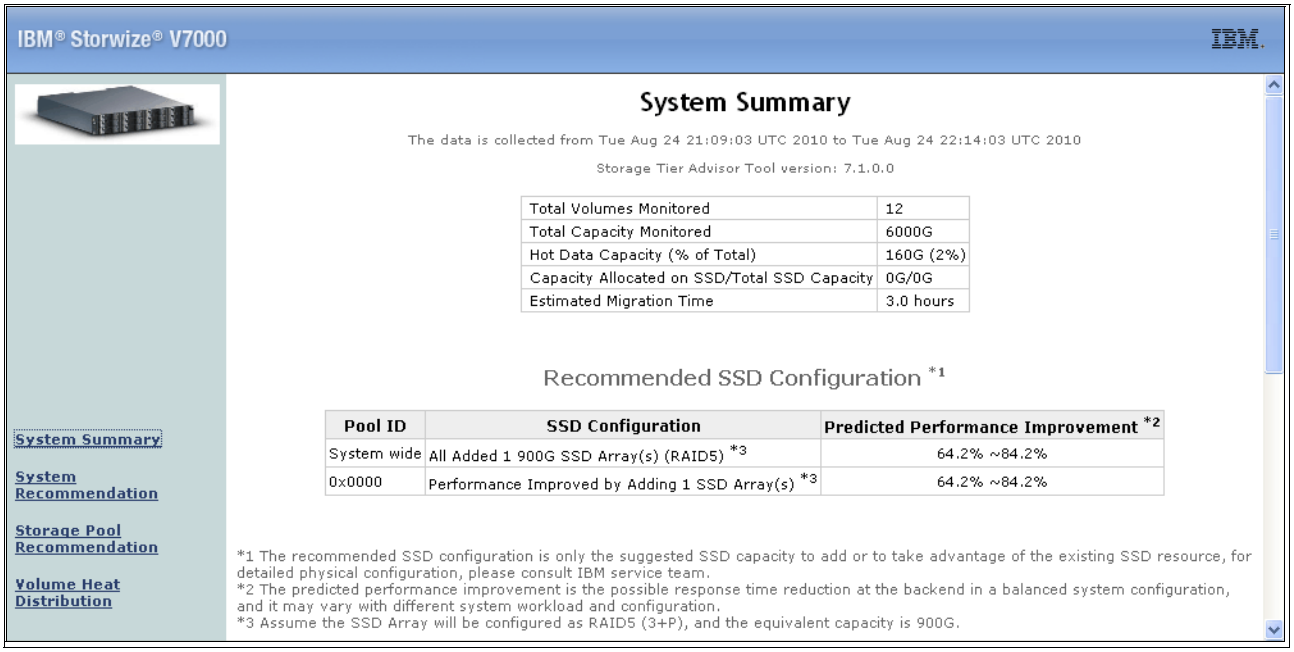

Figure 9-27 STAT report: System Summary

Important: Because this tool was originally available for SAN Volume Controller and IBM Storwize V7000, you can ignore the fact that it is showing IBM Storwize V7000 in the report banner.

The System Summary window provides the most important numbers. In Figure 9-27, we see that 12 volumes were monitored with a total capacity of 6000 GB. The result of the analysis of the hot extents is that about 160 GB (which means 2%) should be migrated to the high-performance disk tier.

It also recommends that one SSD RAID 5 array should be added as a high-performance tier that consists of four SSD drives (3+P). This predicted performance improvement is the possible response time reduction at the back end in a balanced system is 64% - 84%.

Click **Volume Heat Distribution** to change to a more detailed view, as shown in Figure 9-28.

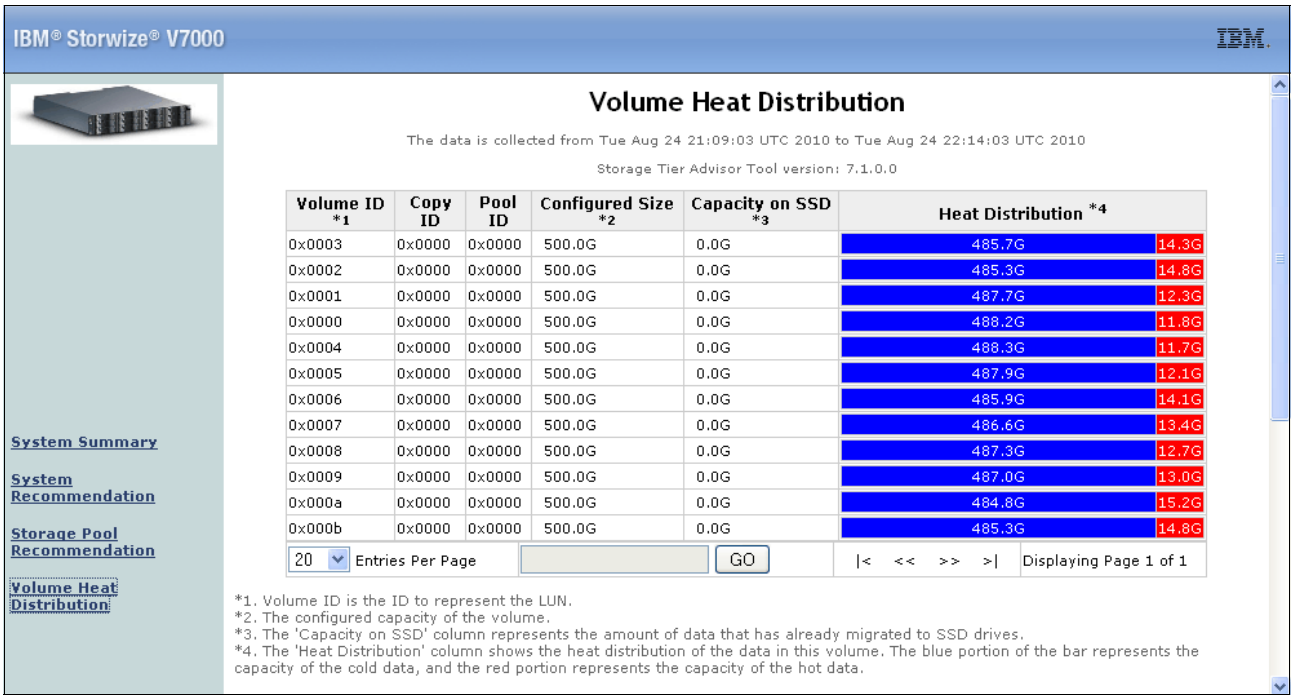

Figure 9-28 Volume Heat Distribution window

The table that is shown in Figure 9-28 gives you a more detailed view as to how the hot extents are distributed across your system. It contains the following information:

- Volume ID: The unique ID of each volume on the IBM Storwize V5000.
- Copy ID: If a volume owns more than one copy, the data is measured for each copy.
- Pool ID: The unique ID of each pool that is configured on the IBM Storwize V5000.
- Configured Size: The configured size of each volume that is represented to the host.
- Capacity on SSD: Capacity of the volumes on high-performance disk tier (even in evaluation mode, volumes can be on high performance disk tiers if they were moved there before).
- Heat Distribution: Shows the heat distribution of the data in this volume. The blue portion of the bar represents the capacity of the cold extents and the red portion represents the capacity of the hot extents. The red hot data is a candidate to be moved to the high-performance disk tier.

9.7 Tivoli Storage Productivity Center

The IBM Tivoli Storage Productivity Center provides a set of policy-driven automated tools for managing storage capacity, availability, events, performance, and assets in your enterprise environment. Tivoli Storage Productivity Center provides storage management from the host and application to the target storage device. It also provides disk and tape subsystem configuration and management, Performance Management, SAN fabric management and configuration, and usage reporting and monitoring. In this section, we describe how to use Tivoli Storage Productivity Center to get usage reporting and to monitor performance data.

Tivoli Storage Productivity Center can help you to identify, evaluate, control, and predict your enterprise storage management assets. Because it is policy-based, it can detect potential problems and automatically make adjustments that are based on the policies and actions that you define. For example, it can notify you when your system is running out of disk space or warn you of an impending storage hardware failure. By alerting you to these and other issues that are related to your stored data, you can prevent unnecessary system and application downtime.

9.7.1 Tivoli Storage Productivity Center benefits

Tivoli Storage Productivity Center includes the following benefits:

- Simplifies the management of storage infrastructures
- Manages, configures, and provisions SAN-attached storage
- Monitors and tracks performance of SAN-attached devices
- Monitors, manages and controls (through zones) SAN fabric components
- Manages the capacity usage and availability of the file systems and databases
- Offers performance monitoring and reporting
- Reports can be viewed by using a web-based GUI

9.7.2 Adding IBM Storwize V5000 in Tivoli Storage Productivity Center

After the Tivoli Storage Productivity Center is installed, it is ready to connect to the IBM Storwize V5000 system.

Complete the following steps to connect Tivoli Storage Productivity Center to the IBM Storwize V5000 system:

1. Open your browser and use the following link to start Tivoli Storage Productivity Center, as show in Figure 9-29:

   `http://TPC_system_Hostname:9550/ITSRM/app/en_US/index.html`

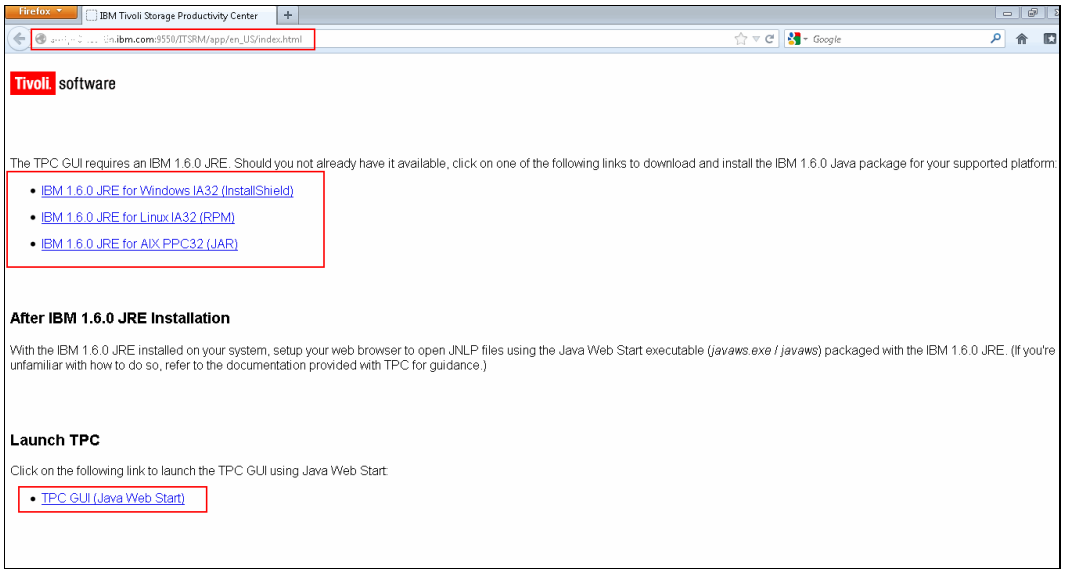

Figure 9-29 Tivoli Storage Productivity Center front page

You also can find a link on the webpage to download IBM Java, if required. To start Tivoli Storage Productivity Center console, click the Tivoli Storage Productivity Center GUI (Java Web Start).

Tivoli Storage Productivity Center starts an application download, as shown in Figure 9-30. If this is the first time you logged in, it takes time to install the required Java packages to the local system.

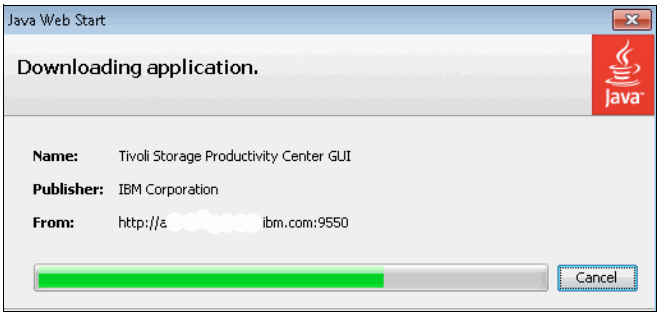

Figure 9-30 Downloading Tivoli Storage Productivity Center application

2. Use your login credentials to access Tivoli Storage Productivity Center, as shown in Figure 9-31.

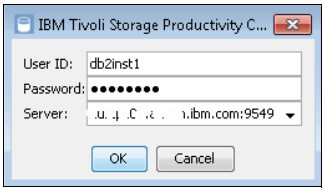

Figure 9-31 Tivoli Storage Productivity Center login access

3. After successfully logging in, you are ready to add storage devices to Tivoli Storage Productivity Center, as shown in Figure 9-32.

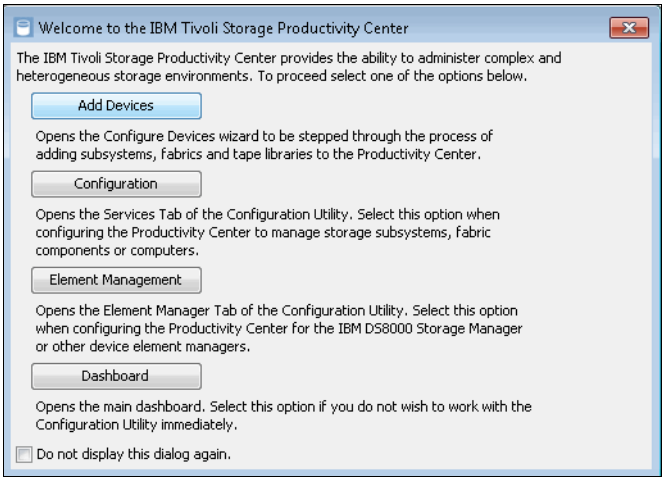

Figure 9-32 Add Devices window

4. Enter the details of your IBM Storwize V5000 in Tivoli Storage Productivity Center, as shown in Figure 9-33.

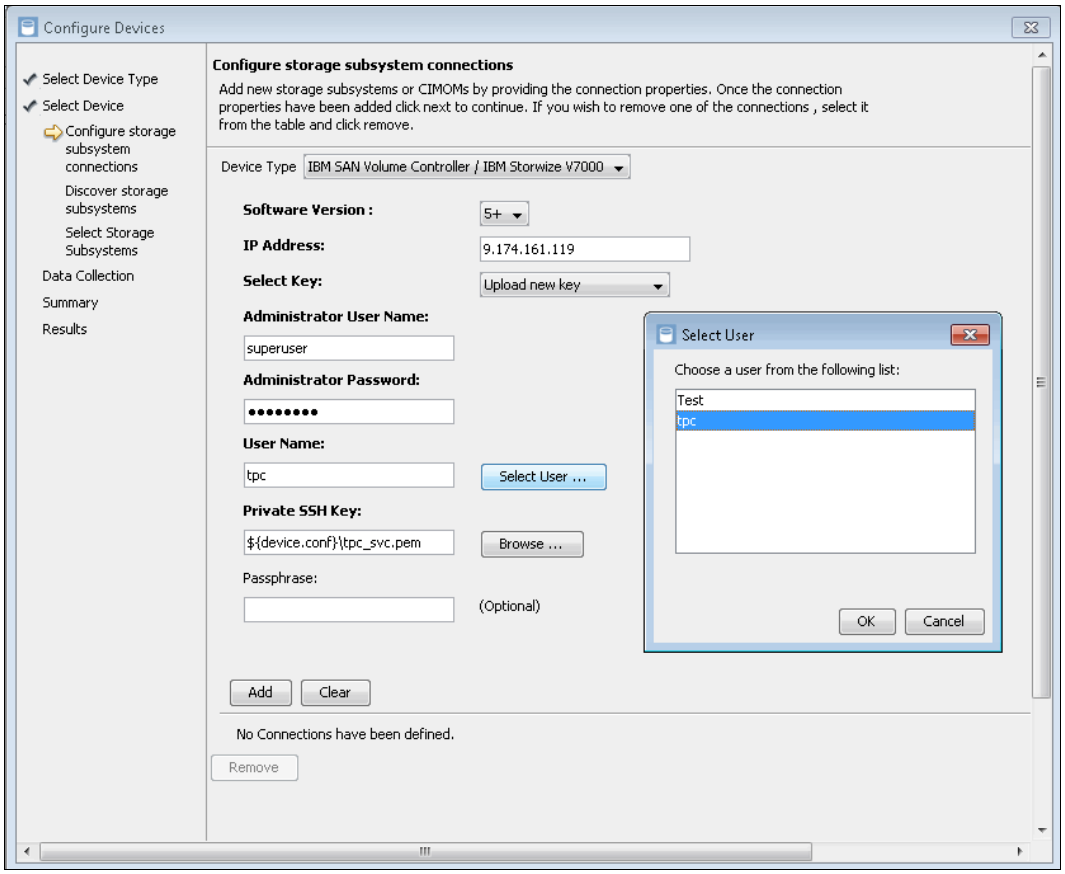

Figure 9-33 Configure device in Tivoli Storage Productivity Center

Continue to following the wizard after you complete all the required fields. After the wizard is completed, Tivoli Storage Productivity Center collects information from IBM Storwize V5000. A summary of details is shown at the end of discovery process.

9.8 Administering and reporting an IBM Storwize V5000 system through Tivoli Storage Productivity Center

This section shows examples of how to use Tivoli Storage Productivity Center to administer, configure, and generate reports for IBM Storwize V5000 system. A detailed description about Tivoli Storage Productivity Center reporting is beyond the intended scope of this book.

9.8.1 Basic configuration and administration

By using Tivoli Storage Productivity Center, you can administer, monitor, and configure your IBM Storwize V5000 system. However, not all of the options that are normally associated with the IBM Storwize V5000 GUI or CLI are available.

After successfully adding your IBM Storwize V5000 system, click **Disk Manager** → **Storage Subsystems** to view your configured devices, as shown in Figure 9-34.

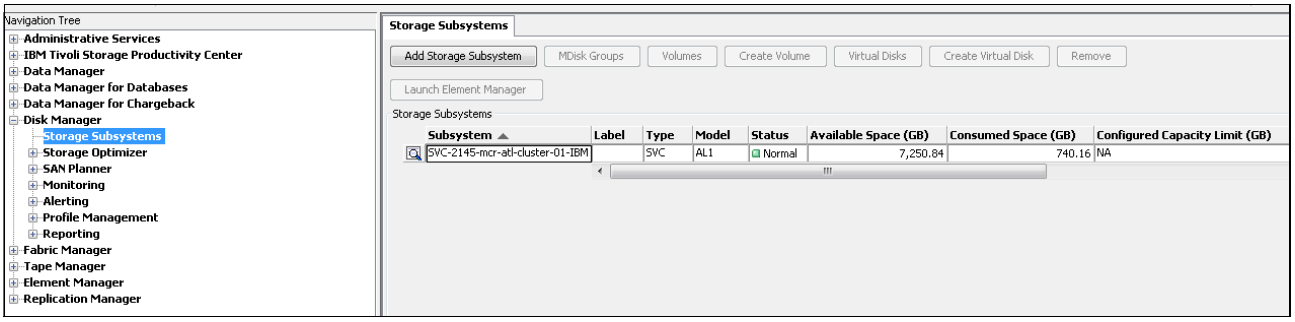

Figure 9-34 Storage Subsystem view

When you highlight the IBM Storwize V5000 system, action buttons become available that you can use to view the device configuration or create virtual disks, as shown in Figure 9-35.

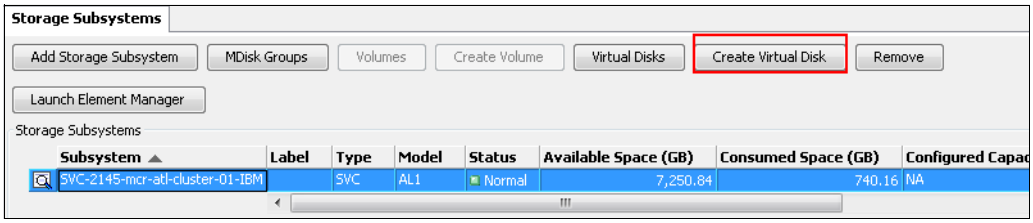

Figure 9-35 Action buttons

The MDisk Groups option provides you with a detailed list of the configured MDisk groups including, pool space, available space, configured space, and Easy Tier Configuration.

The Virtual Disks option lists all the configured disks with the added option to filter them by MDisk Group. The list includes several attributes, such as, capacity, volume type, and type.

Important: Tivoli Storage Productivity Center and SAN Volume Controller use the following terms:

- Virtual Disk: The equivalent of a Volume on a Storwize device
- MDisk Group: The equivalent of a Storage Pool on a Storwize device.

If you click **Create Virtual Disk**, the Create Virtual Disk wizard window opens, as shown in Figure 9-36 on page 441. Use this window to create volumes and specify several options, such as, size, name, thin provisioning, and add MDisks to an MDisk group.

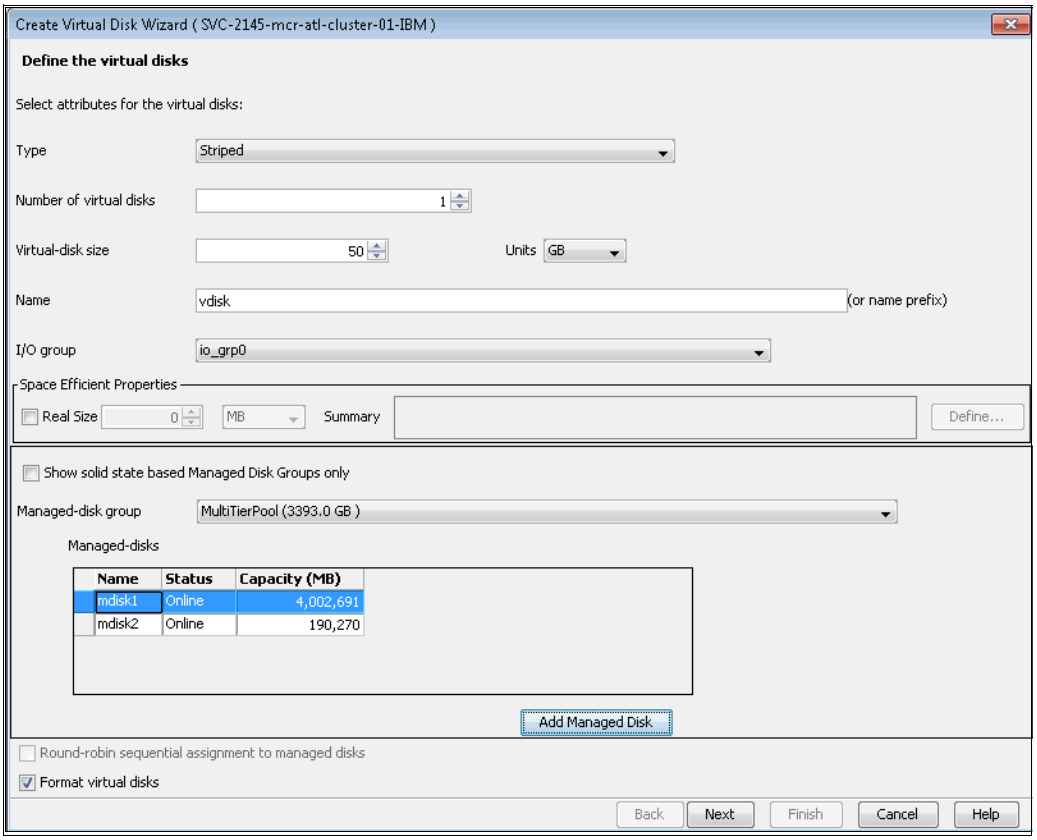

Figure 9-36 Virtual Disk wizard Creation

9.8.2 Generating reports by using Java GUI

In this section, we describe how to generate sample reports by using the GUI. We also create a probe to collect information from IBM Storwize V5000, as shown in Figure 9-37.

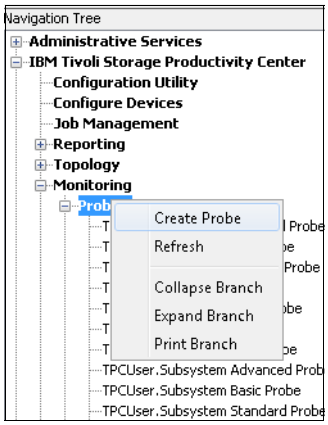

Figure 9-37 Create Probe option

Add the IBM Storwize V5000 probe for collecting information, as shown in Figure 9-38.

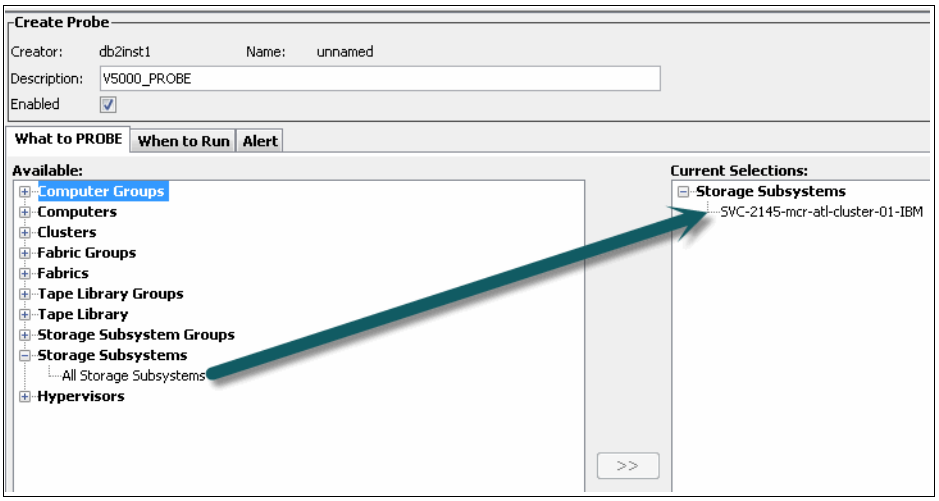

Figure 9-38 Adding IBM Storwize V5000 in probe

After you create a probe, you can click **Create Subsystem Performance Monitor**, as shown in Figure 9-39.

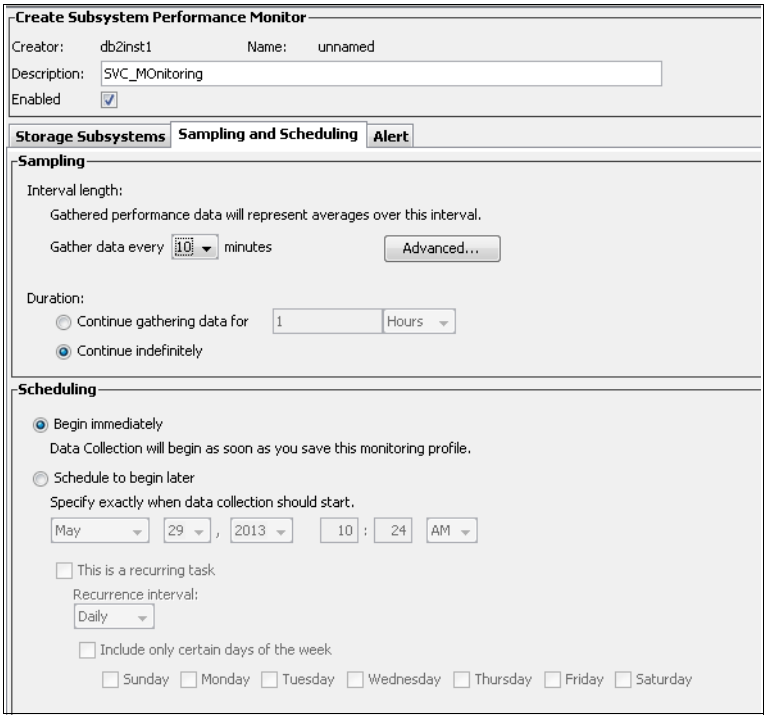

Figure 9-39 Create subsystem performance monitor

To check the MDisk performance, click **Disk Manager** → **Reporting** → **Storage Subsystem Performance** → **By Managed Disk**. You see many options to include in the wizard to check MDisk performance, as shown in Figure 9-40.

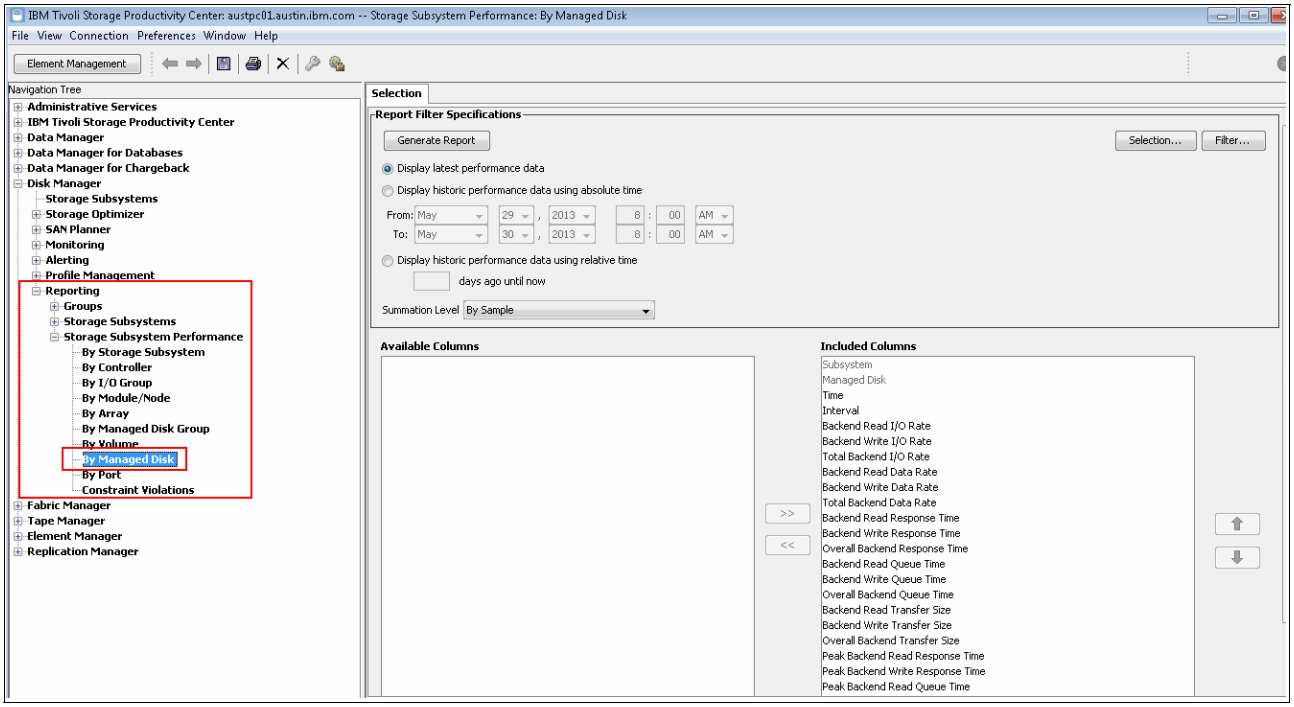

Figure 9-40 Managed disk performance report filter specification

Click **Generate Report** to see a report, as shown in Figure 9-41.

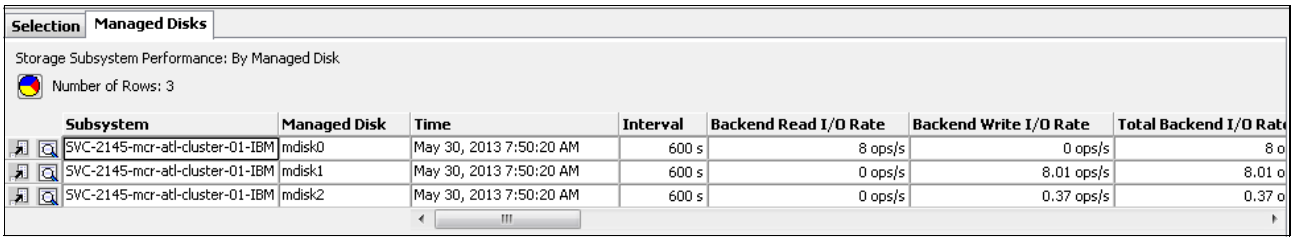

Figure 9-41 MDisk performance report

Click the upper left icon to see a history chart report of the selected MDisk, as shown in Figure 9-42 on page 444.

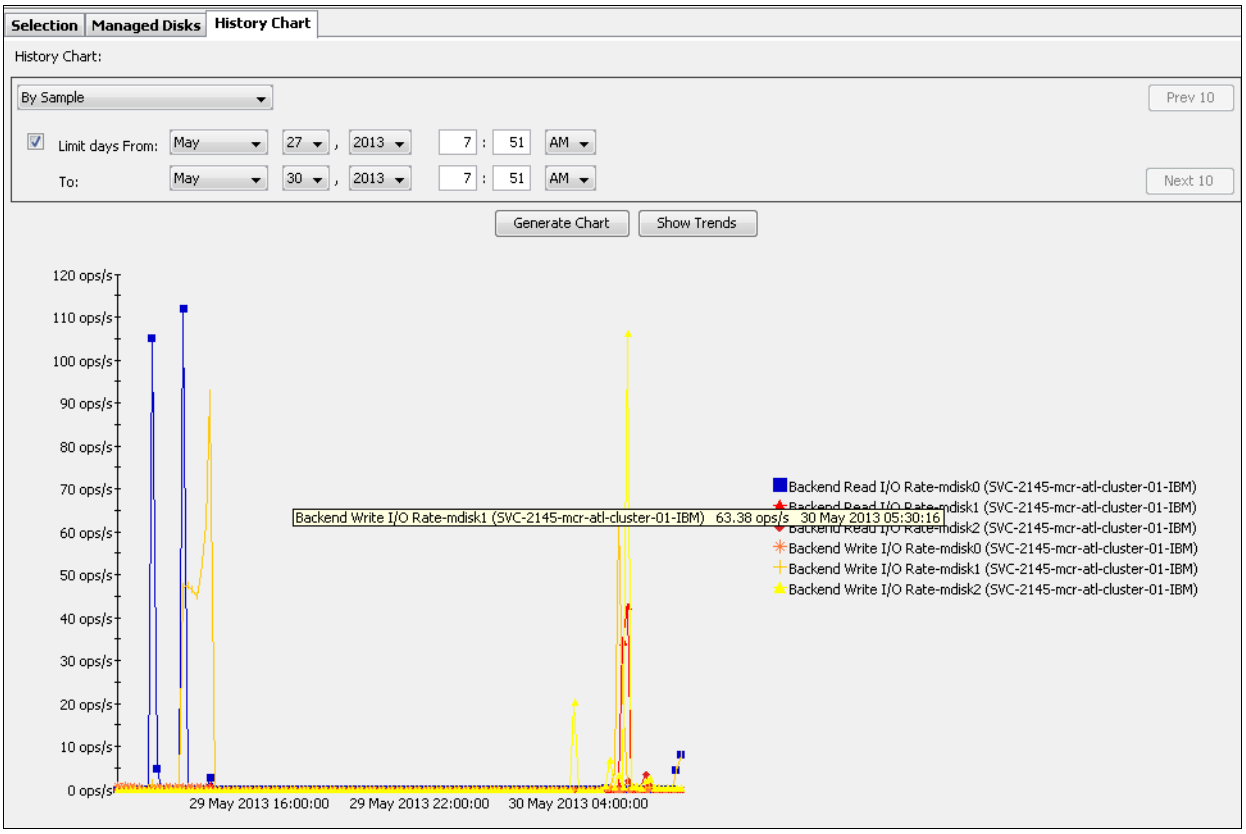

Figure 9-42 MDisk history chart

9.8.3 Generating reports by using Tivoli Storage Productivity Center web console

In this section, we describe how to generate reports by using the Tivoli Storage Productivity Center web console.

To connect to the web page, browse to the following URL:

`https://tpchostname.com:9569/srm/`

You see a login panel (as shown in Figure 9-43) and log in by using your Tivoli Storage Productivity Center credentials.

Figure 9-43 Tivoli Storage Productivity Center login panel

After you log in, you see the Tivoli Storage Productivity Center web dashboard, as shown in Figure 9-44. The Tivoli Storage Productivity Center web-based GUI is used to show information about the storage resources in your environment. It contains predefined and custom reports about performance and storage tiering.

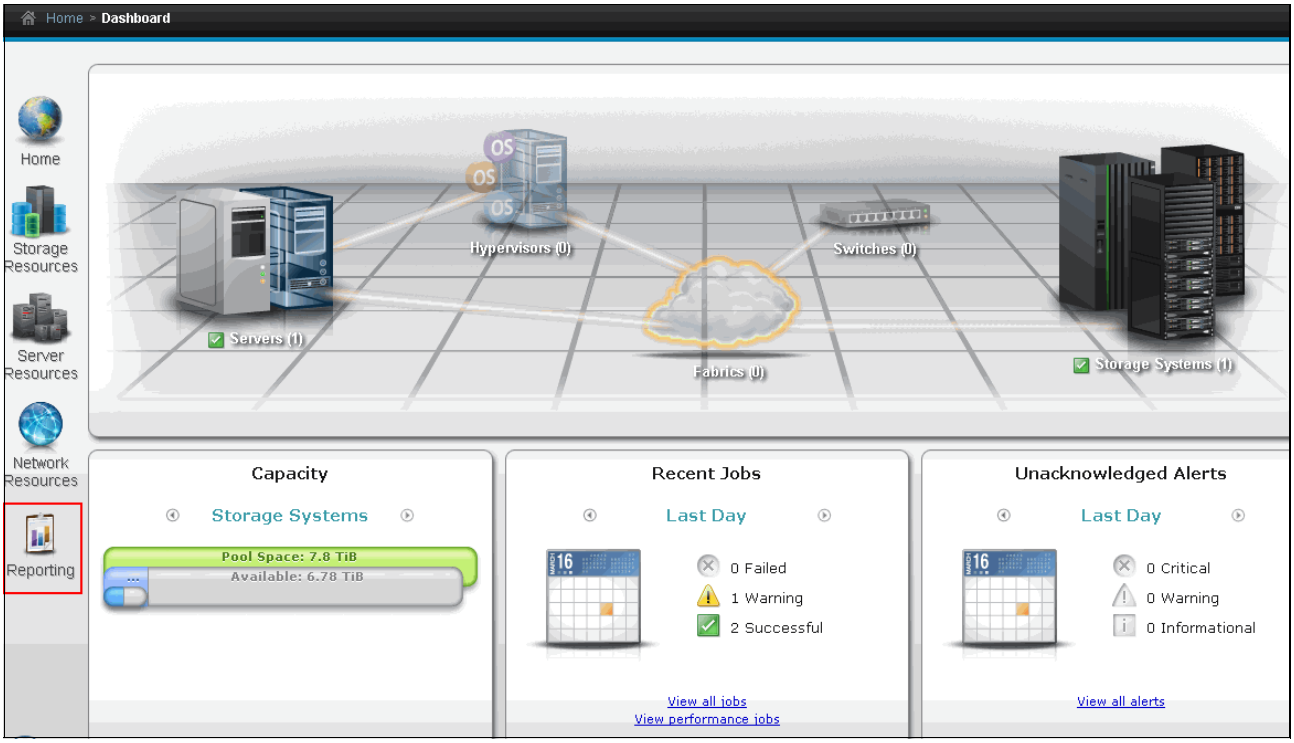

Figure 9-44 Tivoli Storage Productivity Center Dashboard

You can use IBM Tivoli Common Reporting to view predefined reports and create custom reports from the web-based GUI. Predefined reports are also included, as shown in Figure 9-45.

Figure 9-45 Tivoli Storage Productivity Center web-based reporting

Figure 9-46 shows how to select predefined Storage Tiering reporting.

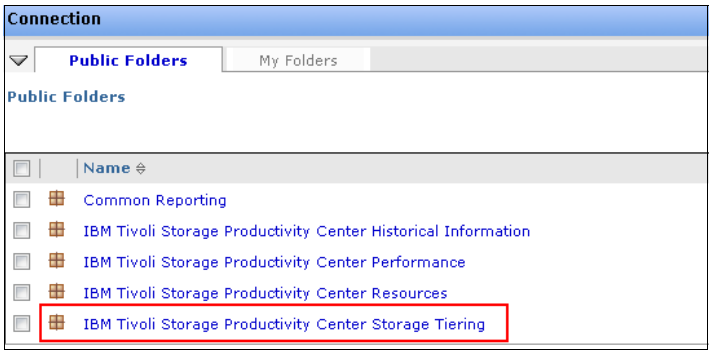

Figure 9-46 Tivoli Storage Productivity Center Storage tiering reporting

Figure 9-47 shows the different report options for Storage Tiering.

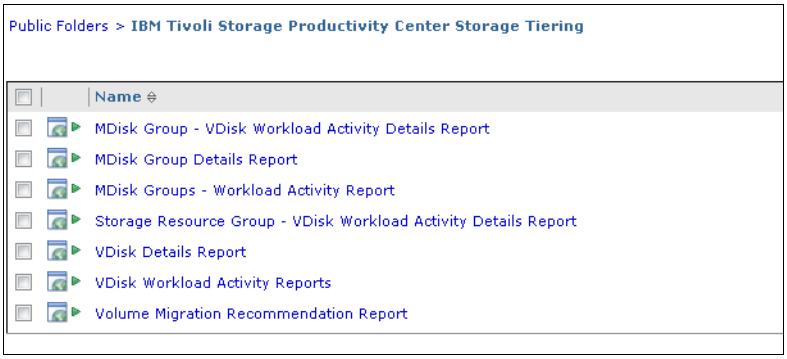

Figure 9-47 Details reports

Figure 9-48 shows the output from VDisk Details report.

IBM® Tivoli®

VDisk Details Report

Storage Virtualizer SVC-2145-mcr-atl-cluster-01-IBM **VDisk** win

Date filter All

MDisk Group	VDisk	VDisk Report Date	VDisk Capacity (GB)	VDisk Thin Provisioning Capacity Status	VDisk Peak Utilization (%)	VDisk Read I/O Rate (ops/s)	VDisk Max Read I/O Rate (ops/s)	VDisk Average Read Response Time (ms/op)	VDisk Max Read Response Time (ms/op)	VDisk Max Read Response Time Threshold of {70 %} Exceeded By (hr.:min.)	VDisk Write I/O Rate (ops/s)	VDisk Average Write Response Time (ms/op)	VDisk Max Write Response Time (ms/op)	VDisk Max Write Response Time Threshold of {90 %} Exceeded By (hr.:min.)	VDisk Max Read Density Rate (ops/s per GB)	VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)	VDi Capa Shar Tot VDi Capa (%
MultiTierPool	win_2k8-iscsi	May	5.00	OK	0.06	0.00	0.06	8.25	8.36	00:00	0.00	1.34	1.34	00:00	0.01	0.08	0

Figure 9-48 VDisk Details report

Figure 9-49 shows the Report Overview in pie-chart format.

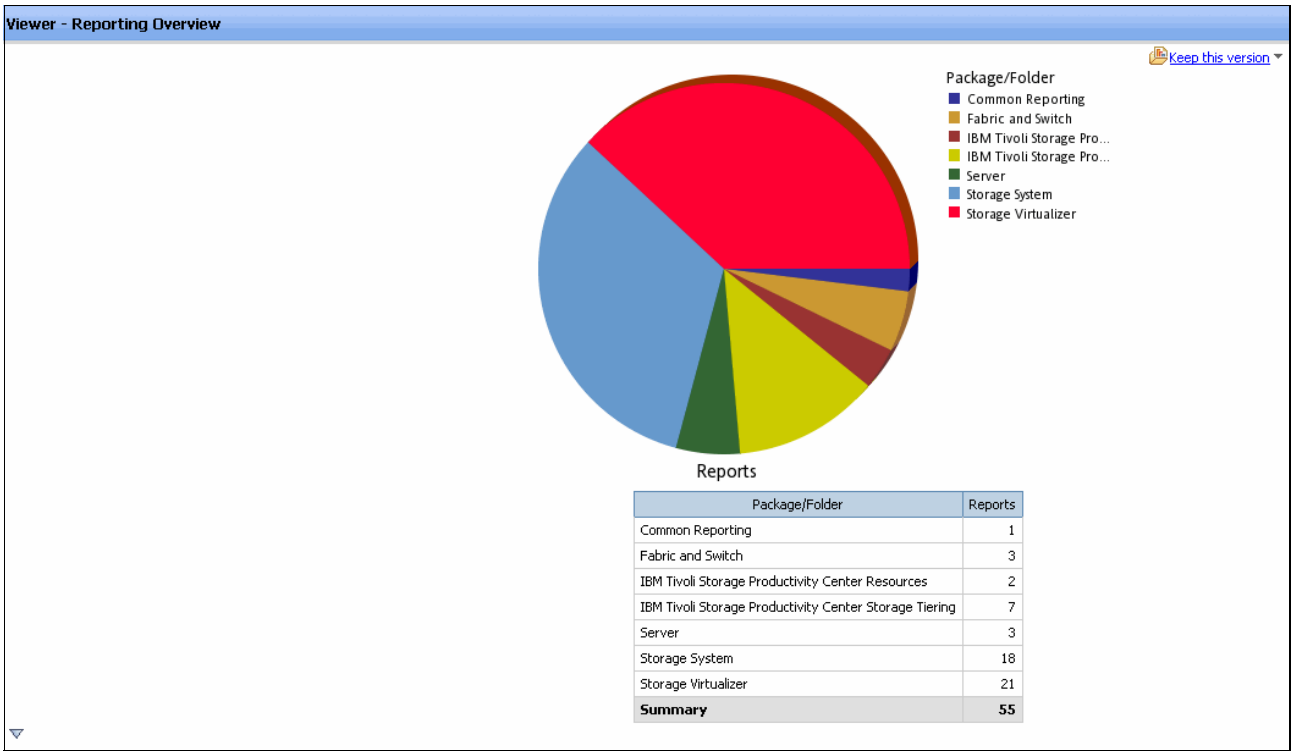

Package/Folder	Reports
Common Reporting	1
Fabric and Switch	3
IBM Tivoli Storage Productivity Center Resources	2
IBM Tivoli Storage Productivity Center Storage Tiering	7
Server	3
Storage System	18
Storage Virtualizer	21
Summary	**55**

Figure 9-49 Reporting Overview

Figure 9-50 shows the Easy Tier usage for volumes. To open this report in Tivoli Storage Productivity Center, click **Storage Resources** → **Volumes**.

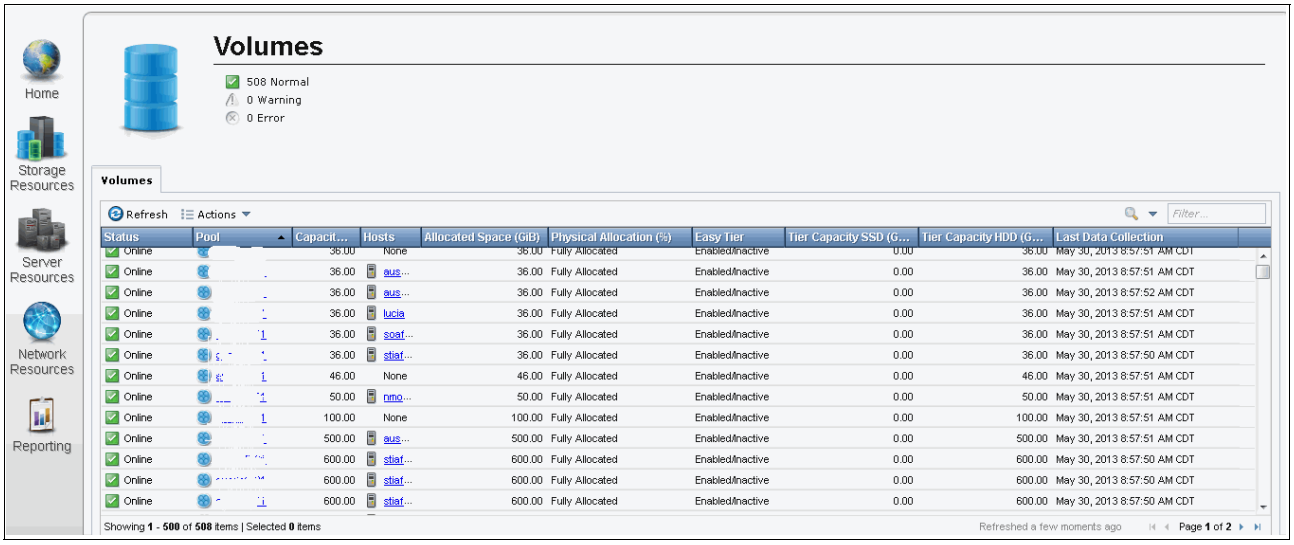

Figure 9-50 Volume Easy Tier usage

Figure 9-51 shows a detailed list of storage pools.

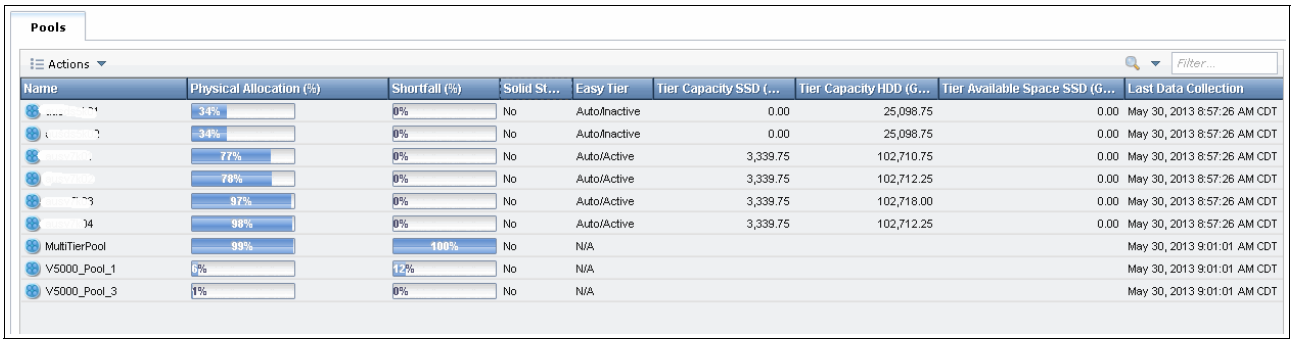

Figure 9-51 Pool Easy Tier information

Figure 9-52 shows Storage Virtualized Pool details in graph format.

Figure 9-52 Pool details

10

Copy services

In this chapter, we describe the copy services functions that are provided by the IBM Storwize V5000 storage system, including FlashCopy and Remote Copy. Copy services functions are useful for making data copies for backup, application test, recovery, and so on. The IBM Storwize V5000 system makes it easy to apply these functions to your environment through its intuitive GUI.

This chapter includes the following topics:

- FlashCopy
- Remote Copy
- Troubleshooting Remote Copy
- Managing Remote Copy by using the GUI

10.1 FlashCopy

By using the FlashCopy function of the IBM Storwize V5000 storage system, you can create a *point-in-time copy* of one or more volumes. In this section, we describe the structure of FlashCopy and provide details about its configuration and use.

You can use FlashCopy to solve critical and challenging business needs that require the duplication of data on your source volume. Volumes can remain online and active while you create consistent copies of the data sets. Because the copy is performed at the block level, it operates below the host operating system and cache and, therefore, is not apparent to the host.

Flushing: Because FlashCopy operates at the block level below the host operating system and cache, those levels do need to be flushed for consistent FlashCopy copies.

While the FlashCopy operation is performed, I/O to the source volume is frozen briefly to initialize the FlashCopy bitmap and then is allowed to resume. Although several FlashCopy options require the data to be copied from the source to the target in the background (which can take time to complete), the resulting data on the target volume copy appears to complete immediately. This task is accomplished by using a bitmap (or bit array) that tracks changes to the data after the FlashCopy is started, and an indirection layer, which allows data to be read from the source volume transparently.

10.1.1 Business requirements for FlashCopy

When you are deciding whether FlashCopy addresses your needs, you must adopt a combined business and technical view of the problems you must solve. Determine your needs from a business perspective, and then determine whether FlashCopy fulfills the technical needs of those business requirements.

With an immediately available copy of the data, FlashCopy can be used in the following business scenarios:

- Rapidly creating consistent backups of dynamically changing data

  FlashCopy can be used to create backups through periodic running; the FlashCopy target volumes can be used to complete a rapid restore of individual files or the entire volume through Reverse FlashCopy (by using the **-restore** option).

  The target volumes that are created by FlashCopy can also be used for backup to tape. By attaching them to another server and performing backups from there, it allows the production server to continue largely unaffected. After the copy to tape completes, the target volumes can be discarded or kept as a rapid restore copy of the data.

- Rapidly creating consistent copies of production data to facilitate data movement or migration between hosts

  FlashCopy can be used to facilitate the movement or migration of data between hosts while it minimizes downtime for applications. FlashCopy allows application data to be copied from source volumes to new target volumes while applications remain online. After the volumes are fully copied and synchronized, the application can be stopped and then immediately started on the new server that is accessing the new FlashCopy target volumes. This mode of migration is faster than other migration methods that are available through the IBM Storwize V5000 because the size and the speed of the migration is not as limited.

- Rapidly creating copies of production data sets for application development and testing

  Under normal circumstances to perform application development and testing, data must be restored from traditional backup media, such as, tape. Depending on the amount of data and the technology in use, this process can easily take a day or more. With FlashCopy, a copy can be created and be online for use in just a few minutes. The time varies based on the application and the data set size.

- Rapidly creating copies of production data sets for auditing purposes and data mining

  Auditing or data mining normally require the usage of the production applications. This situation can cause high loads for databases track inventories or similar data. With FlashCopy, you can create copies for your reporting and data mining activities. This feature reduces the load on your production systems, which increases their performance.

- Rapidly creating copies of production data sets for quality assurance

  Quality assurance is an interesting case for FlashCopy. Because traditional methods involve so much time and labor, the refresh cycle is typically extended. This reduction in time that is required allows much more frequent refreshes of the quality assurance database.

10.1.2 FlashCopy functional overview

FlashCopy occurs between a source volume and a target volume. The source and target volumes must be the same size. Multiple FlashCopy mappings (source-to-target relationships) can be defined, and point-in-time consistency can be maintained across multiple point-in-time mappings by using consistency groups. For more information about FlashCopy consistency groups, see “FlashCopy consistency groups” on page 456.

The minimum granularity that IBM Storwize V5000 storage system supports for FlashCopy is an entire volume; it is not possible to use FlashCopy to copy only part of a volume. Additionally, the source and target volumes must belong to the same IBM Storwize V5000 storage system, but they do not have to be in the same storage pool.

Before you start a FlashCopy (regardless of the type and options that are specified), the IBM Storwize V5000 must put the cache into write-through mode, which flushes the I/O that is bound for the source volume. If you are scripting FlashCopy operations from the CLI, you must run the **prestartfcmap** or **prestartfcconsistgrp** command. However, this step is managed for you and carried out automatically by the GUI. This is not the same as flushing the host cache, which is not required. After FlashCopy is started, an effective copy of a source volume to a target volume is created. The content of the source volume is immediately presented on the target volume and the original content of the target volume is lost. This FlashCopy operation is also referred to as a *time-zero copy* (T_0).

Immediately following the FlashCopy operation, the source and target volumes are available for use. The FlashCopy operation creates a bitmap that is referenced and maintained to direct I/O requests within the source and target relationship. This bitmap is updated to reflect the active block locations as data is copied in the background from the source to target and updates are made to the source.

Figure 10-1 shows the redirection of the host I/O toward the source volume and the target volume.

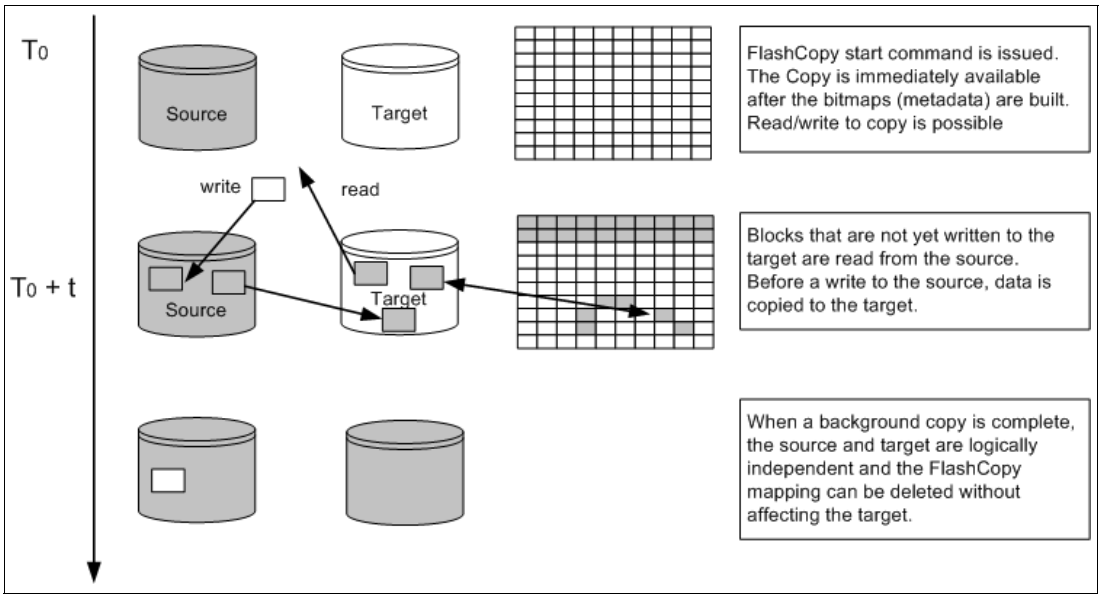

Figure 10-1 Redirection of host I/O

When data is copied between volumes, it is copied in units of address space known as *grains*. Grains are units of data that are grouped to optimize the use of the bitmap that track changes to the data between the source and target volume. You have the option of using 64 KB or 256 KB grain sizes (256 KB is the default). The FlashCopy bitmap contains 1 bit for each grain and is used to track whether the source grain was copied to the target. The 64 KB grain size uses bitmap space at a rate of four times the default 256 KB size.

The FlashCopy bitmap dictates the following read and write behaviors for the source and target volumes:

- Read I/O request to source: Reads are performed from the source volume the same as for non-FlashCopy volumes.
- Write I/O request to source: Writes to the source cause the grains of the source volume to be copied to the target if they were not already and then the write is performed to the source.
- Read I/O request to target: Reads are performed from the target if the grains already were copied; otherwise, the read is performed from the source.
- Write I/O request to target: Writes to the target cause the grain to be copied from the source to the target first, unless the entire grain is being written and then the write completes to the target only.

FlashCopy mappings

A FlashCopy mapping defines the relationship between a source volume and a target volume. FlashCopy mappings can be stand-alone mappings or a member of a consistency group, as described in “FlashCopy consistency groups” on page 456.

Incremental FlashCopy mappings

In an incremental FlashCopy, the initial mapping copies all of the data from the source volume to the target volume. Subsequent FlashCopy mappings copy only data that was modified since the initial FlashCopy mapping. This action reduces the amount of time that it takes to re-create an independent FlashCopy image. You can define a FlashCopy mapping as incremental only when you create the FlashCopy mapping.

Multiple target FlashCopy mappings

You can copy up to 256 target volumes from a single source volume. Each relationship between a source and target volume is managed by a unique mapping such that a single volume can be the source volume for up to 256 mappings.

Each of the mappings from a single source can be started and stopped independently. If multiple mappings from the same source are active (in the copying or stopping states), a dependency exists between these mappings.

If a single source volume has multiple target FlashCopy volumes, the write to the source volume does not cause its data to be copied to all of the targets. Instead, it is copied to the newest target volume only. The older targets refer to new targets first before they refer to the source. A dependency relationship exists between a particular target and all newer targets that share a source until all data is copied to this target and all older targets.

Cascaded FlashCopy mappings

The cascaded FlashCopy function allows a FlashCopy target volume to be the source volume of another FlashCopy mapping. Up to 256 mappings can exist in a cascade. If cascaded mappings and multiple target mappings are used, a tree of up to 256 mappings can be created.

Cascaded mappings differ from multiple target FlashCopy mappings in depth. Cascaded mappings have an association in the manner of A > B > C, while multiple target FlashCopy has an association in the manner A > B1 and A > B2.

Background copy

The background copy rate is a property of a FlashCopy mapping that is defined as a value of 0 - 100. The background copy rate can be defined and dynamically changed for individual FlashCopy mappings. A value of 0 disables background copy. This option is also called the *no-copy option*, which provides pointer-based images for limited lifetime uses.

With FlashCopy background copy, the source volume data is copied to the corresponding target volume in the FlashCopy mapping. If the background copy rate is set to 0 (which means disable the FlashCopy background copy), only data that changed on the source volume is copied to the target volume. The benefit of using a FlashCopy mapping with background copy enabled is that the target volume becomes a real independent clone of the FlashCopy mapping source volume after the copy is complete. When the background copy is disabled, only the target volume is a valid copy of the source data while the FlashCopy mapping remains in place. Copying only the changed data saves your storage capacity (assuming it is thin provisioned and the **`-rsize`** option was correctly setup.)

The relationship of the background copy rate value to the amount of data that is copied per second is shown in Table 10-1 on page 454.

Table 10-1 Background copy rate

Value	Data that is copied per second	Grains per second (256 KB grain)	Grains per second (64 KB grain)
1 - 10	128 KB	0.5	2
11 - 20	256 KB	1	4
21 - 30	512 KB	2	8
31 - 40	1 MB	4	16
41 - 50	2 MB	8	32
51 - 60	4 MB	16	64
61 - 70	8 MB	32	128
71 - 80	16 MB	64	256
81 - 90	32 MB	128	512
91 - 100	64 MB	256	1024

Data copy rate: The data copy rate remains the same regardless of the FlashCopy grain size. The difference is the number of grains that are copied per second. The gain size can be 64 KB or 256 KB. The smaller size uses more bitmap space and thus limits the total amount of FlashCopy space possible. However, it might be more efficient regarding the amount of data that is moved, depending on your environment.

Cleaning rate

The cleaning rate provides a method for FlashCopy copies with dependant mappings (multiple target or cascaded) to complete their background copies before their source goes offline or is deleted after a stop is issued.

When you create or modify a FlashCopy mapping, you can specify a cleaning rate for the FlashCopy mapping that is independent of the background copy rate. The cleaning rate is also defined as a value of 0 - 100, which has the same relationship to data copied per second as the backup copy rate (see Table 10-1).

The cleaning rate controls the rate at which the cleaning process operates. The purpose of the cleaning process is to copy (or flush) data from FlashCopy source volumes upon which there are dependent mappings. For cascaded and multiple target FlashCopy, the source might be a target for another FlashCopy or a source for a chain (cascade) of FlashCopy mappings. The cleaning process must complete before the FlashCopy mapping can go to the stopped state. This feature and the distinction between *stopping* and *stopped* states was added to prevent data access interruption for dependent mappings when their source is issued a stop.

FlashCopy mapping states

A mapping is in one of the following states at any point:

- Idle or Copied

  The source and target volumes act as independent volumes even if a mapping exists between the two. Read and write caching is enabled for the source and the target volumes.

If the mapping is incremental and the background copy is complete, the mapping records only the differences between the source and target volumes. The source and target volumes go offline if the connection to both nodes in the IBM Storwize V5000 storage system that the mapping is assigned to is lost.

- Copying

  The copy is in progress. Read and write caching is enabled on the source and the target volumes.

- Prepared

  The mapping is ready to start. The target volume is online, but is not accessible. The target volume cannot perform read or write caching. Read and write caching is failed by the SCSI front end as a hardware error. If the mapping is incremental and a previous mapping completed, the mapping records only the differences between the source and target volumes. The source and target volumes go offline if the connection to both nodes in the IBM Storwize V5000 storage system that the mapping is assigned to is lost.

- Preparing

  The target volume is online, but not accessible. The target volume cannot perform read or write caching. Read and write caching is failed by the SCSI front end as a hardware error. Any changed write data for the source volume is flushed from the cache. Any read or write data for the target volume is discarded from the cache. If the mapping is incremental and a previous mapping completed, the mapping records only the differences between the source and target volumes. The source and target volumes go offline if the connection to both nodes in the IBM Storwize V5000 storage system that the mapping is assigned to is lost.

- Stopped

  The mapping is stopped because you issued a stop command or an I/O error occurred. The target volume is offline and its data is lost. To access the target volume, you must restart or delete the mapping. The source volume is accessible and the read and write cache is enabled. If the mapping is incremental, the mapping is recording write operations to the source volume. The source and target volumes go offline if the connection to both nodes in the IBM Storwize V5000 storage system that the mapping is assigned to is lost.

- Stopping

  The mapping is copying data to another mapping. If the background copy process is complete, the target volume is online while the stopping copy process completes. If the background copy process did not complete, data is discarded from the target volume cache. The target volume is offline while the stopping copy process runs. The source volume is accessible for I/O operations.

- Suspended

  The mapping did start, but it did not complete. Access to the metadata is lost, which causes the source and target volume to go offline. When access to the metadata is restored, the mapping returns to the copying or stopping state and the source and target volumes return online. The background copy process resumes.

  Any data that was not flushed and was written to the source or target volume before the suspension is in cache until the mapping leaves the suspended state.

FlashCopy consistency groups

Consistency groups address the requirement to preserve point-in-time data consistency across multiple volumes for applications that include related data that spans them. For these volumes, consistency groups maintain the integrity of the FlashCopy by ensuring that dependent writes are run in the application's intended sequence. For more information about dependent writes, see "Dependent writes" on page 456.

When consistency groups are used, the FlashCopy commands are issued to the FlashCopy consistency group, which performs the operation on all FlashCopy mappings that are contained within the consistency group.

Figure 10-2 shows a consistency group that consists of two FlashCopy mappings.

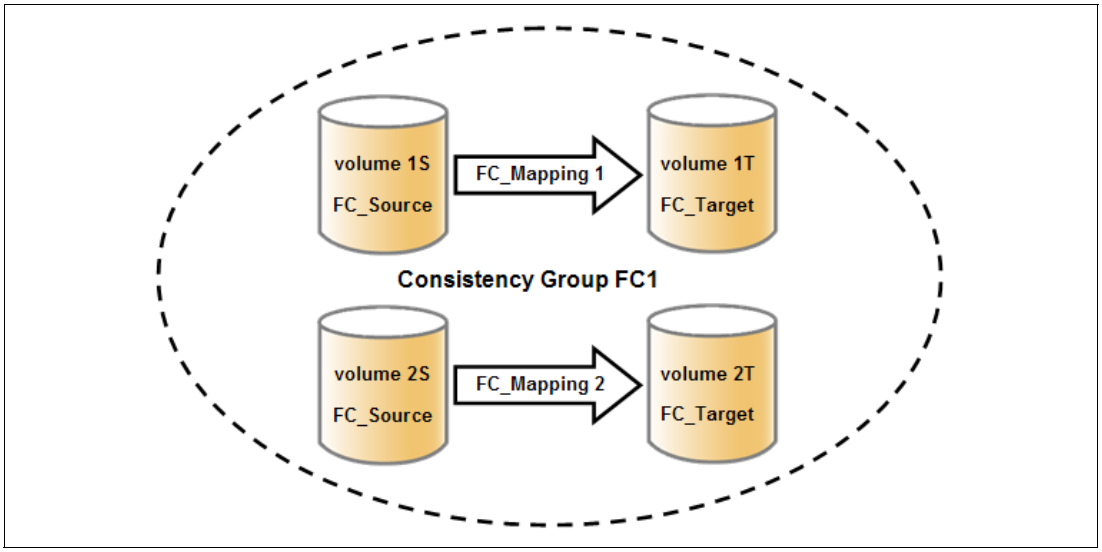

Figure 10-2 FlashCopy consistency group

FlashCopy mapping management: After an individual FlashCopy mapping was added to a consistency group, it can be managed only as part of the group. Operations such as, start and stop are no longer allowed on the individual mapping.

Dependent writes

To show why it is crucial to use consistency groups when a data set spans multiple volumes, consider the following typical sequence of writes for a database update transaction:

1. A write is run to update the database log, which indicates that a database update is about to be performed.
2. A second write is run to complete the actual update to the database.
3. A third write is run to update the database log, which indicates that the database update completed successfully.

The database ensures the correct ordering of these writes by waiting for each step to complete before it starts the next step. However, if the database log (updates 1 and 3) and the database (update 2) are on separate volumes, it is possible for the FlashCopy of the database volume to occur before the FlashCopy of the database log. This situation can result in the target volumes seeing writes (1) and (3) but not (2) because the FlashCopy of the database volume occurred before the write was completed.

In this case, if the database was restarted by using the backup that was made from the FlashCopy target volumes, the database log indicates that the transaction completed successfully when, in fact, it had not. This situation occurs because the FlashCopy of the volume with the database file was started (bitmap was created) before the write completed to the volume. Therefore, the transaction is lost and the integrity of the database is in question.

To overcome the issue of dependent writes across volumes and to create a consistent image of the client data, it is necessary to perform a FlashCopy operation on multiple volumes as an atomic operation by using consistency groups.

A FlashCopy consistency group can contain up to 512 FlashCopy mappings. The more mappings that you have, the more time it takes to prepare the consistency group. FlashCopy commands can then be issued to the FlashCopy consistency group and simultaneously for all of the FlashCopy mappings that are defined in the consistency group. For example, when the FlashCopy for the consistency group is started, all FlashCopy mappings in the consistency group are started at the same time, which results in a point-in-time copy that is consistent across all FlashCopy mappings that are contained in the consistency group.

A consistency group aggregates FlashCopy mappings, not volumes. Thus, where a source volume has multiple FlashCopy mappings, they can be in the same or separate consistency groups. If a particular volume is the source volume for multiple FlashCopy mappings, you might want to create separate consistency groups to separate each mapping of the same source volume. Regardless of whether the source volume with multiple target volumes is in the same consistency group or in separate consistency groups, the resulting FlashCopy produces multiple identical copies of the source data.

The consistency group can be specified when the mapping is created. You can also add the FlashCopy mapping to a consistency group or change the consistency group of a FlashCopy mapping later.

Important: Do not place stand-alone mappings into a consistency group because they become controlled as part of that consistency group.

FlashCopy consistency group states

A FlashCopy consistency group is in one of the following states at any point:

- Idle or Copied

  All FlashCopy Mappings in this consistency group are in the Idle or Copied state.

- Preparing

  At least one FlashCopy mapping in this consistency group is in the Preparing state.

- Prepared

  The consistency group is ready to start. While in this state, the target volumes of all FlashCopy mappings in this consistency group are not accessible.

- Copying

  At least one FlashCopy mapping in the consistency group is in the Copying state and no FlashCopy mappings are in the Suspended state.

- Stopping

  At least one FlashCopy mapping in the consistency group is in the Stopping state and no FlashCopy mappings are in the Copying or Suspended state.

- Stopped

  The consistency group is stopped because you issued a command or an I/O error occurred.
- Suspended

  At least one FlashCopy mapping in the consistency group is in the Suspended state.
- Empty

  The consistency group does not have any FlashCopy mappings.

Reverse FlashCopy

Reverse FlashCopy enables FlashCopy targets to become restore points for the source without breaking the FlashCopy relationship and without waiting for the original copy operation to complete. It supports multiple targets and multiple rollback points.

A key advantage of Reverse FlashCopy is that it does not delete the original target, thus allowing processes that use the target, such as, a tape backup, to continue uninterrupted.

You can also create an optional copy of the source volume that is made before the reverse copy operation is started. This copy restores the original source data, which can be useful for diagnostic purposes.

Figure 10-3 shows an example of the reverse FlashCopy scenario.

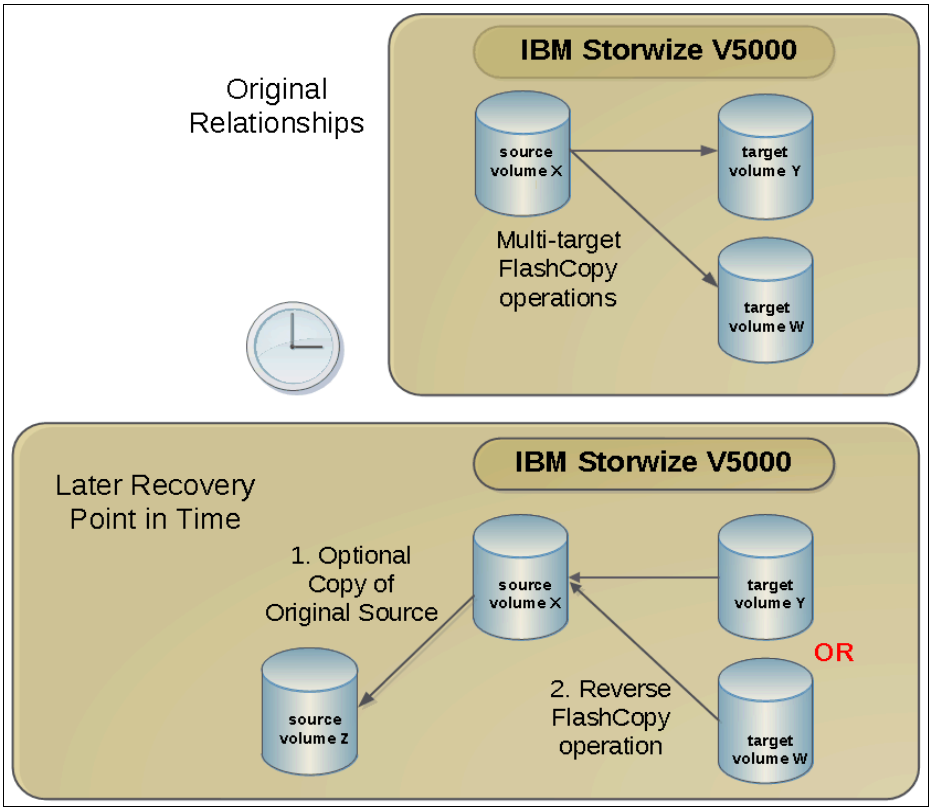

Figure 10-3 Reverse FlashCopy scenario

To restore from a FlashCopy backup by using the GUI, complete the following steps:

1. (Optional) Create a target volume (volume Z) and run FlashCopy on the production volume (volume X) to copy data on to the new target for later problem analysis.
2. Create a FlashCopy map with the backup to be restored (volume Y) or (volume W) as the source volume and volume X as the target volume.
3. Start the FlashCopy map (volume Y to volume X).

> **The -restore option:** In the CLI, you must add the **`-restore`** option to the command **`svctask startfcmap`** manually. For more information about using the CLI, see Appendix A, “Command-line interface setup and SAN Boot” on page 609.

Regardless of whether the initial FlashCopy map (volume X to volume Y) is incremental, the Reverse FlashCopy operation copies only the modified data.

Consistency groups are reversed by creating a set of new “reverse” FlashCopy maps and adding them to a new “reverse” consistency group. Consistency groups cannot contain more than one FlashCopy map with the same target volume.

10.1.3 Planning for FlashCopy

There are several items that must be considered before a FlashCopy is performed, which are described in this section.

Guidelines for FlashCopy implementation

Consider the following guidelines for FlashCopy implementation:

- The source and target volumes must be on the same IBM Storwize V5000 storage system.
- The source and target volumes do not need to be in the same storage pool.
- The FlashCopy source and target volumes can be thin-provisioned.
- The source and target volumes must be the same size. The size of the source and target volumes cannot be altered (increased or decreased) while a FlashCopy mapping is defined.
- FlashCopy operations perform in direct proportion to the performance of the source and target disks. If you have a fast source disk and slow target disk, the performance of the source disk is reduced because it must wait for the write operation to occur at the target before it can write to the source.

Maximum configurations for FlashCopy

Table 10-2 shows some of the FlashCopy maximum configurations. For more information about the latest values, see the IBM Storwize V5000 Information Center at this website:

http://pic.dhe.ibm.com/infocenter/storwize/v5000_ic/index.jsp

Table 10-2 FlashCopy maximum configurations

FlashCopy property	Maximum
FlashCopy targets per source	256
FlashCopy mappings per cluster	4,096
FlashCopy consistency groups per cluster	127
FlashCopy mappings per consistency group	512

FlashCopy presets

The IBM Storwize V5000 storage system provides three FlashCopy presets (Snapshot, Clone, and Backup) to simplify the more common FlashCopy operations, as shown in Table 10-3.

Table 10-3 FlashCopy presets

Preset	Purpose
Snapshot	Creates a point-in-time view of the production data. The snapshot is not intended to be an independent copy. Instead, it is used to maintain a view of the production data at the time the snapshot is created. This preset automatically creates a thin-provisioned target volume with none of the capacity that is allocated at the time of creation. The preset uses a FlashCopy mapping with none of the background copy so that only data written to the source or target is copied to the target volume.
Clone	Creates an exact replica of the volume, which can be changed without affecting the original volume. After the copy operation completes, the mapping that was created by the preset is automatically deleted. This preset automatically creates a volume with the same properties as the source volume and creates a FlashCopy mapping with a background copy rate of 50. The FlashCopy mapping is configured to automatically delete when the FlashCopy mapping reaches 100% completion.
Backup	Creates a point-in-time replica of the production data. After the copy completes, the backup view can be refreshed from the production data, with minimal copying of data from the production volume to the backup volume. This preset automatically creates a volume with the same properties as the source volume. The preset creates an incremental FlashCopy mapping with a background copy rate of 50.

Presets: All of the presets can be adjusted by using the Advanced Settings expandable section in the GUI.

10.1.4 Managing FlashCopy by using the GUI

The IBM Storwize V5000 storage system provides a separate function icon to access copy service management. The following windows are available for managing FlashCopy under the Copy Services function icon:

- FlashCopy
- Consistency Groups
- FlashCopy Mappings

The Copy Services function icon is shown in Figure 10-4.

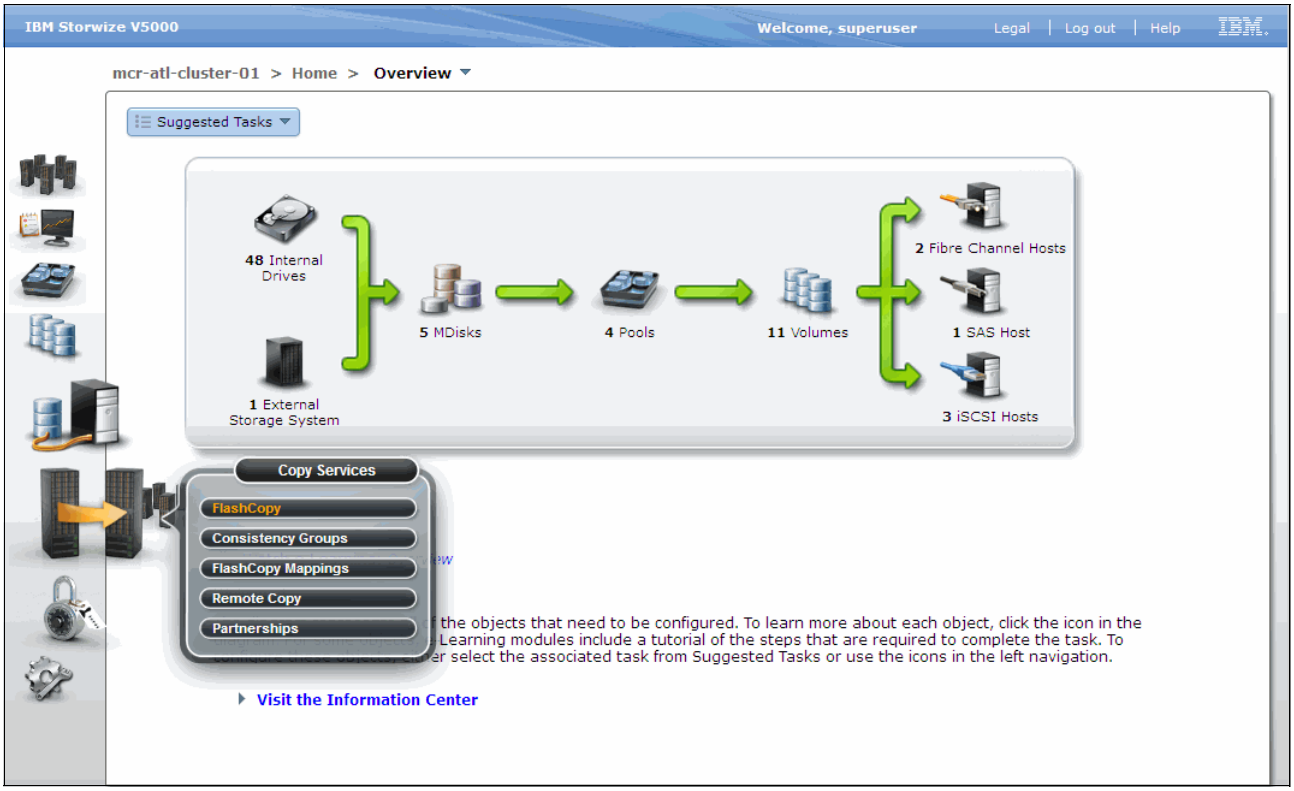

Figure 10-4 Copy Services function icon

Most of the actions to manage the FlashCopy mapping can be done in the FlashCopy window or the FlashCopy Mappings windows, although the quick path to create FlashCopy presets can be found only in the FlashCopy window.

Click **FlashCopy** in the Copy Services function icon menu and the FlashCopy window opens, as shown in Figure 10-5 on page 462. In the FlashCopy window, the FlashCopy mappings are organized by volumes.

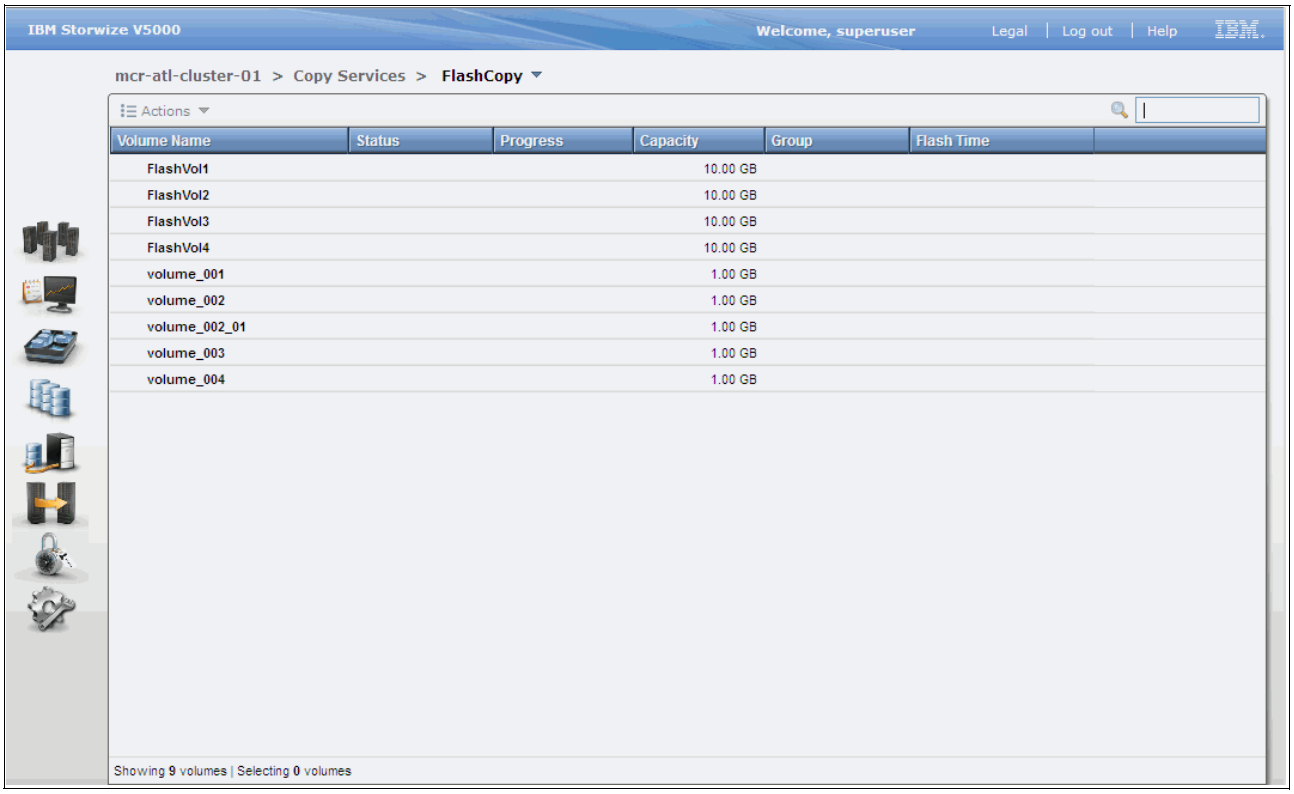

Figure 10-5 FlashCopy window

Click **FlashCopy Mappings** in the Copy Services function icon menu and the FlashCopy Mappings window opens, as shown in Figure 10-6. In the FlashCopy Mappings window, the FlashCopy mappings are listed individually.

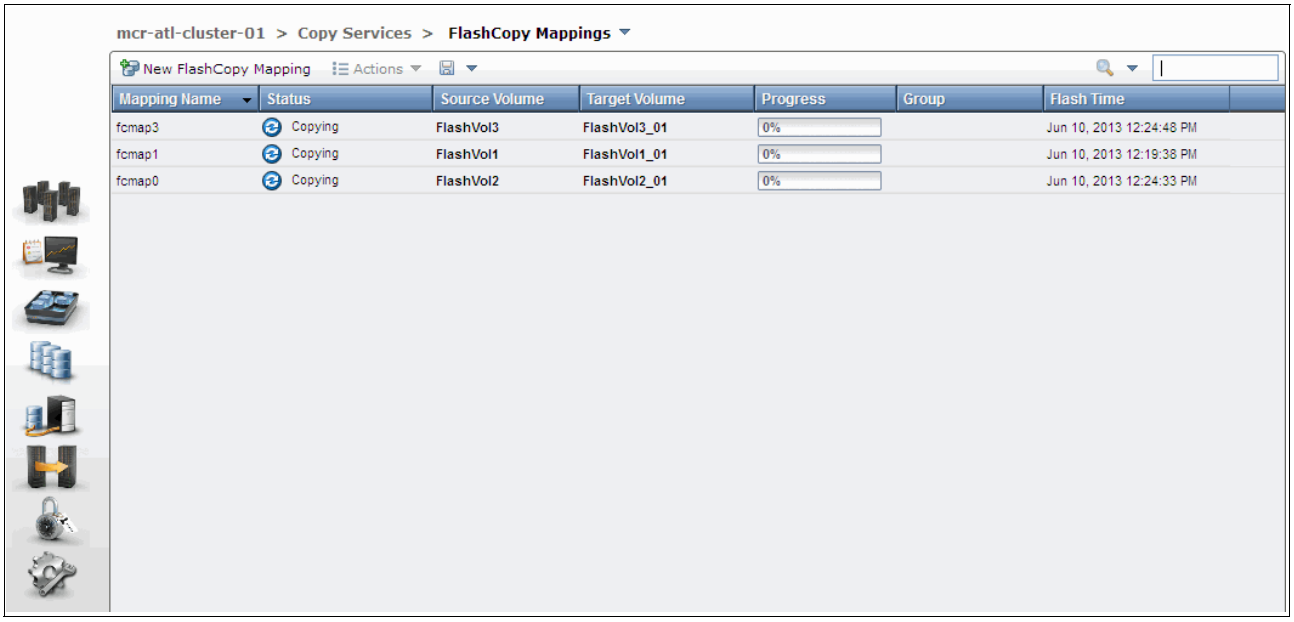

Figure 10-6 FlashCopy Mappings window

The Consistency Groups window is used to manage the consistency groups for FlashCopy mappings. Click **Consistency Groups** in the Copy Services function icon menu and the Consistency Groups window opens, as shown in Figure 10-7.

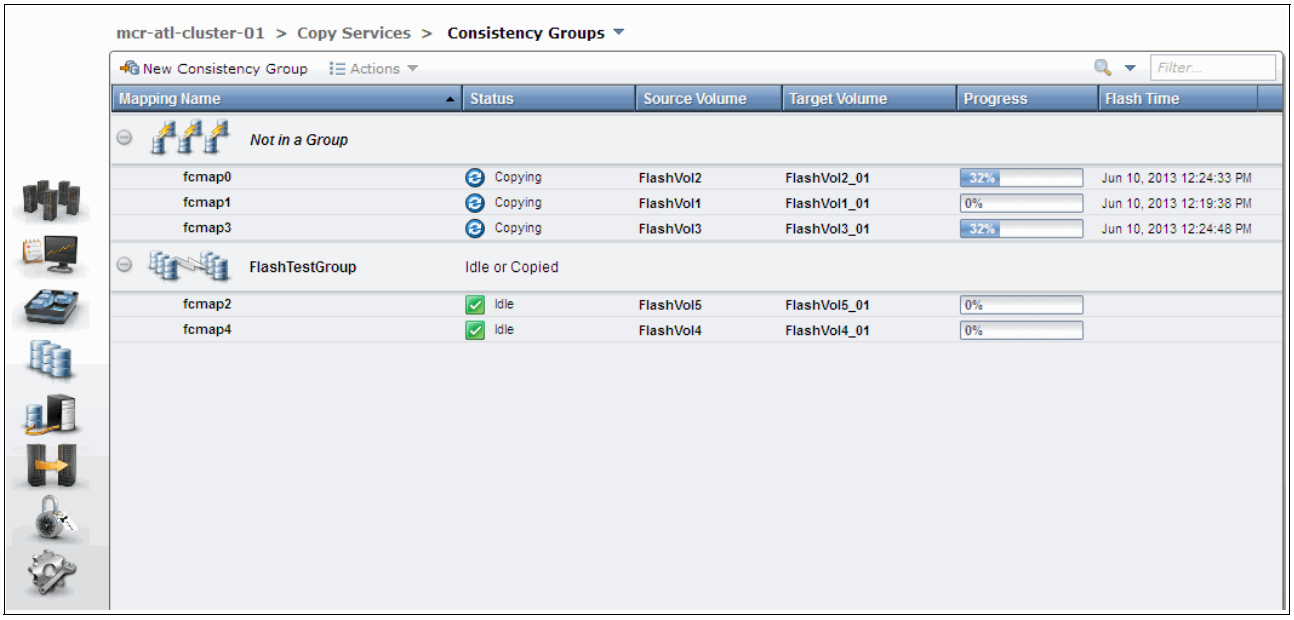

Figure 10-7 Consistency Groups window

Quick path to create FlashCopy presets

It is easy to create a FlashCopy by using the presets in the FlashCopy window.

Creating a snapshot

In the FlashCopy window, choose a volume and click **New Snapshot** from the Actions drop-down menu, as shown in Figure 10-8. Alternatively, you can highlight your chosen volume and right-click to access the Actions menu.

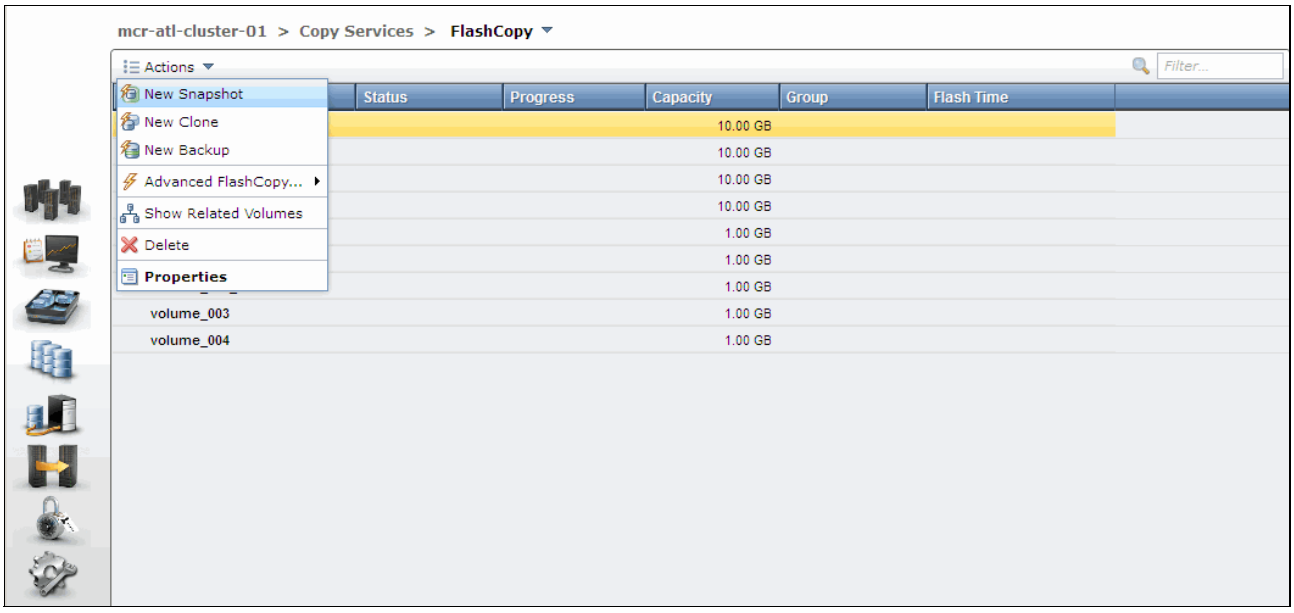

Figure 10-8 Create a snapshot by using the preset

You now have a snapshot volume for the volume you selected.

Creating a clone

In the FlashCopy window, choose a volume and click **New Clone** from the Actions drop-down menu, as shown in Figure 10-9. Alternatively, highlight your chosen volume and right-click to access the Actions menu.

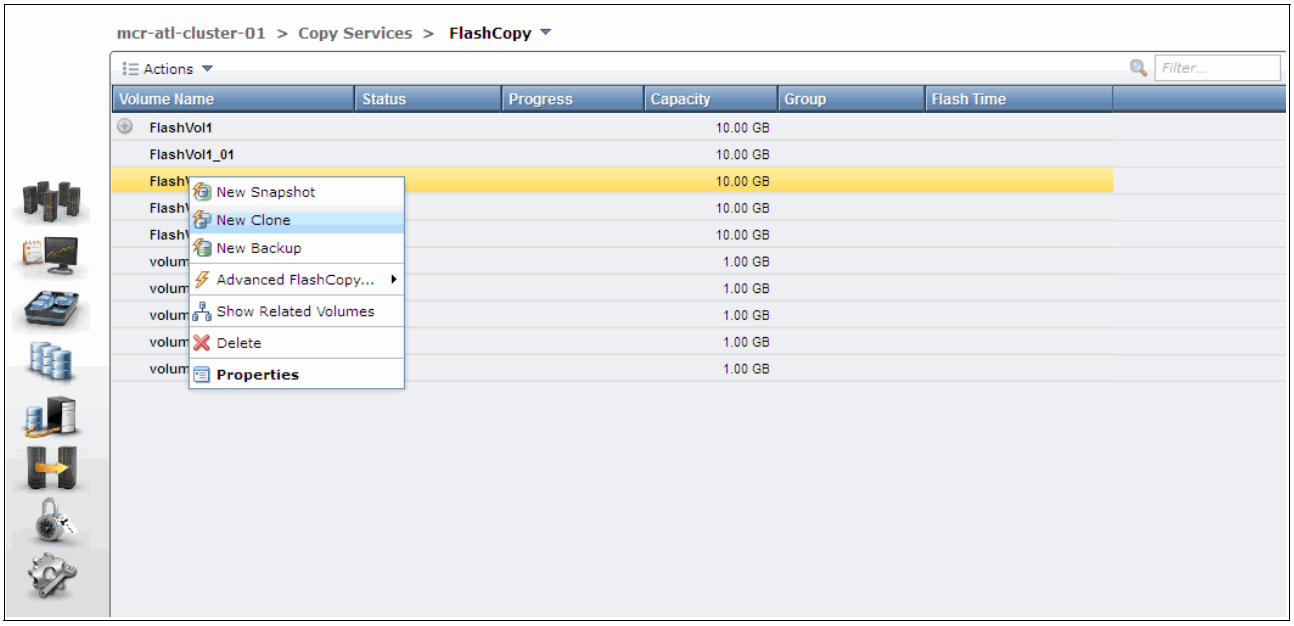

Figure 10-9 Create a clone from the preset

You now have a clone volume for the volume you selected.

Creating a backup

In the FlashCopy window, choose a volume and click **New Backup** from the Actions drop-down menu, as shown in Figure 10-10 on page 465. Alternatively, highlight your chosen volume and right-click to access the Actions menu.

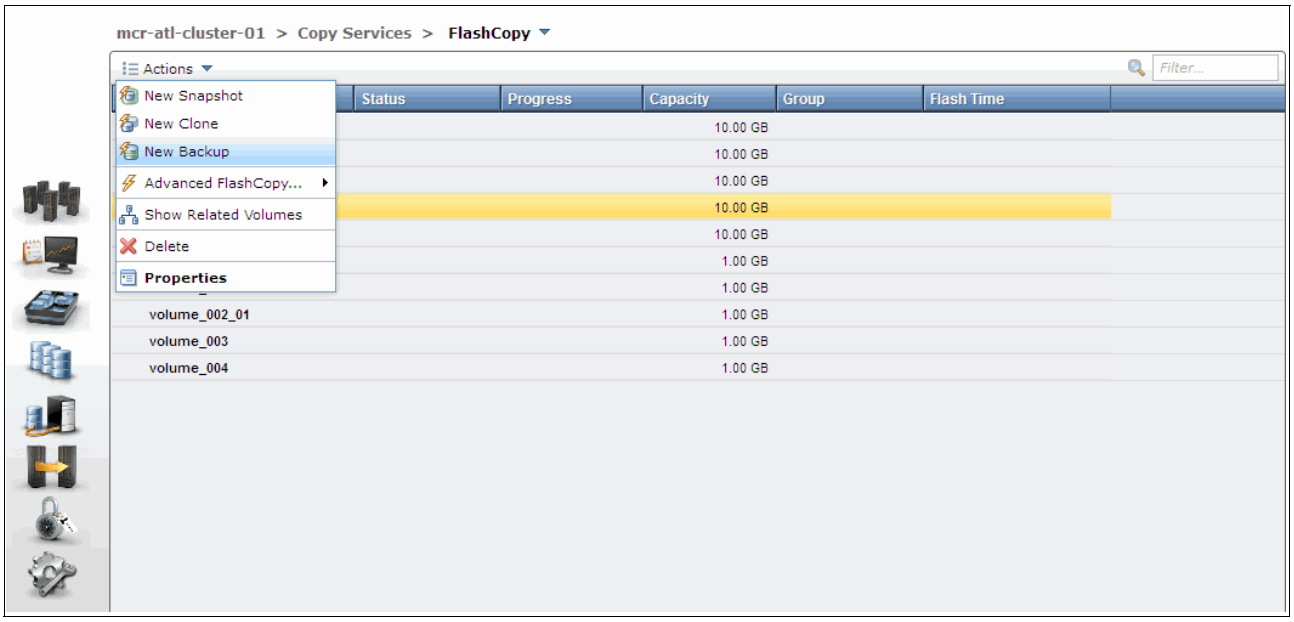

Figure 10-10 Create a backup from the preset

You now have a backup volume for the volume you selected.

In the FlashCopy window and in the FlashCopy Mappings window, you can monitor the progress of the running FlashCopy operations, as shown in Figure 10-11. The progress bars for each target volume indicate the copy progress in percentage. The copy progress remains 0% for snapshots (there is no change until data is written to the target volume). The copy progresses for clone and backup continues to increase until the copy process completes.

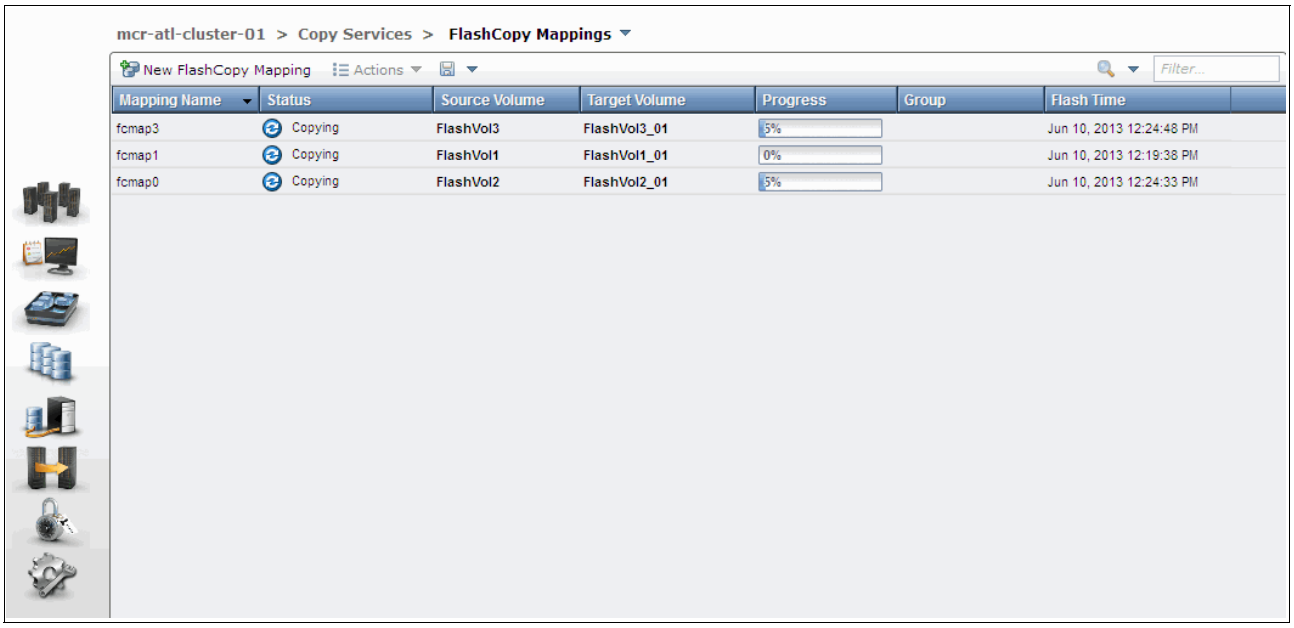

Figure 10-11 FlashCopy in progress that is viewed in the FlashCopy Mappings window

The copy progress can be also found under the Running Tasks status indicator, as shown in Figure 10-12.

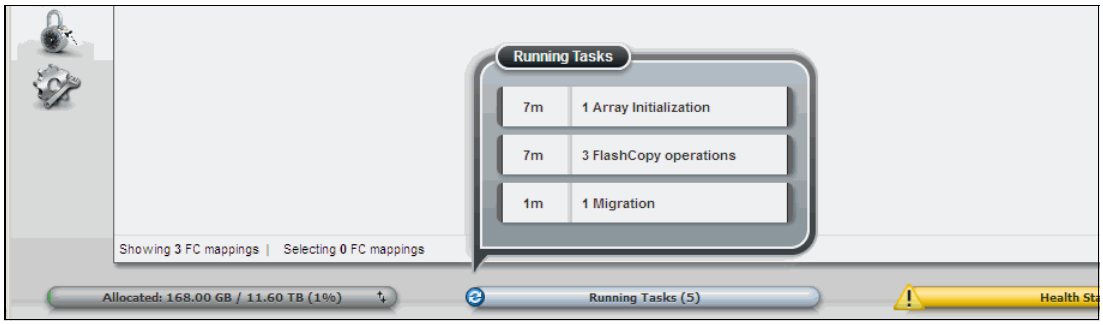

Figure 10-12 Running Tasks bar: FlashCopy operations

This view is slightly different from the FlashCopy and FlashCopy Mappings windows, as shown in Figure 10-13.

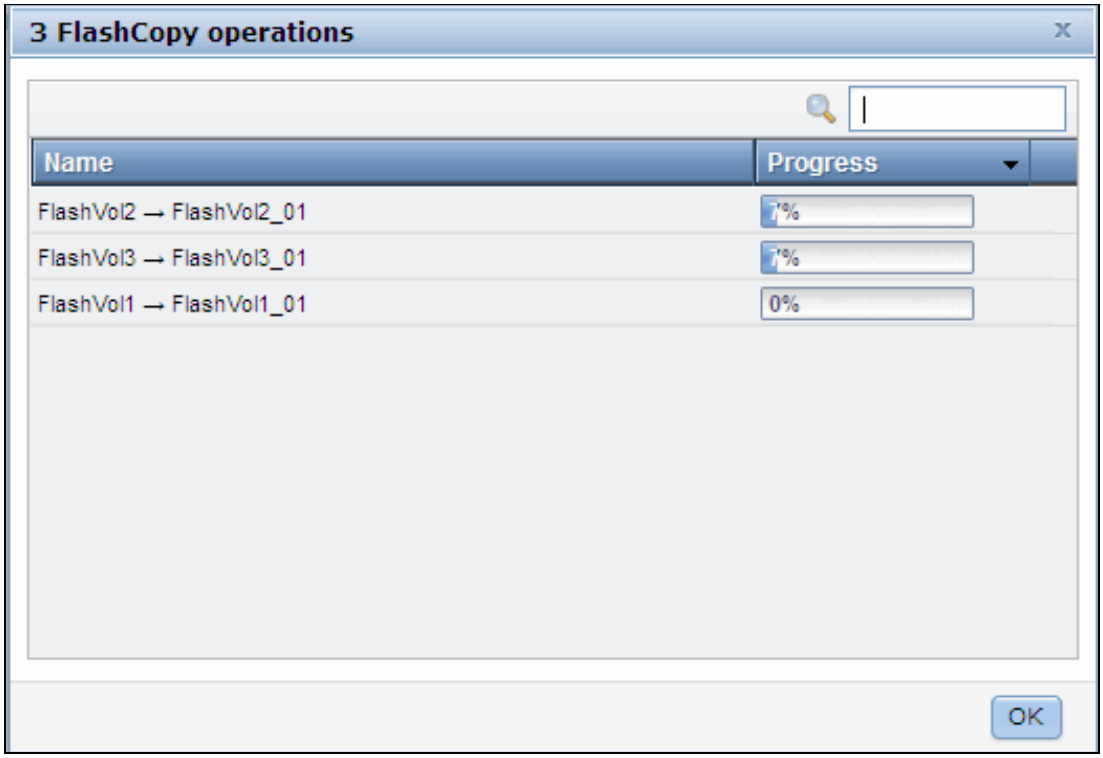

Figure 10-13 FlashCopy operations that are shown through Running Tasks

After the copy processes complete, you find the FlashCopy mapping with the clone preset (`FlashVol2` in our example) was deleted automatically, as shown in Figure 10-14. There are now two identical volumes that are independent of each other.

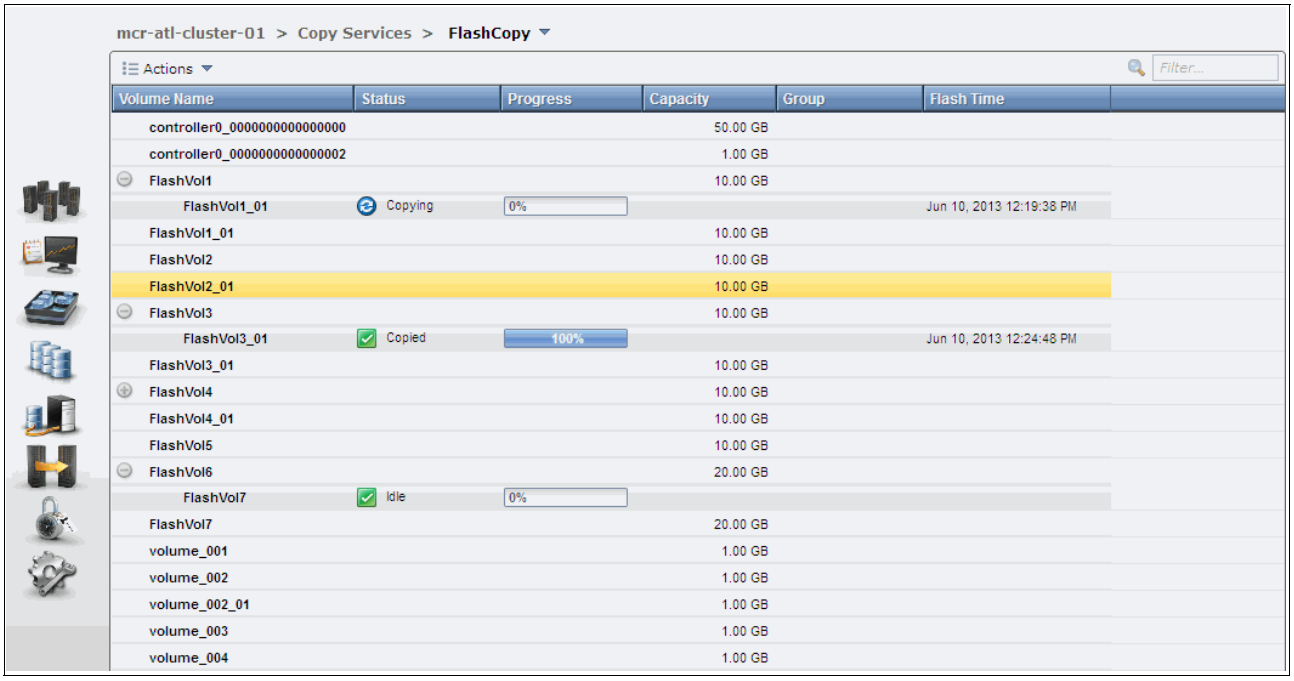

Figure 10-14 FlashCopy progresses complete

10.1.5 Managing FlashCopy mappings

The FlashCopy presets cover the most frequently used FlashCopy configurations for general situations. However, customized FlashCopy mappings are still necessary in some complicated scenarios.

Creating FlashCopy mappings

You can create FlashCopy mappings through the FlashCopy window. Select the volume that you want to be the source volume for the FlashCopy mapping and click **Advanced FlashCopy...** from the Actions drop-down menu, as shown in Figure 10-15 on page 468. Alternatively, select the volume and right-click.

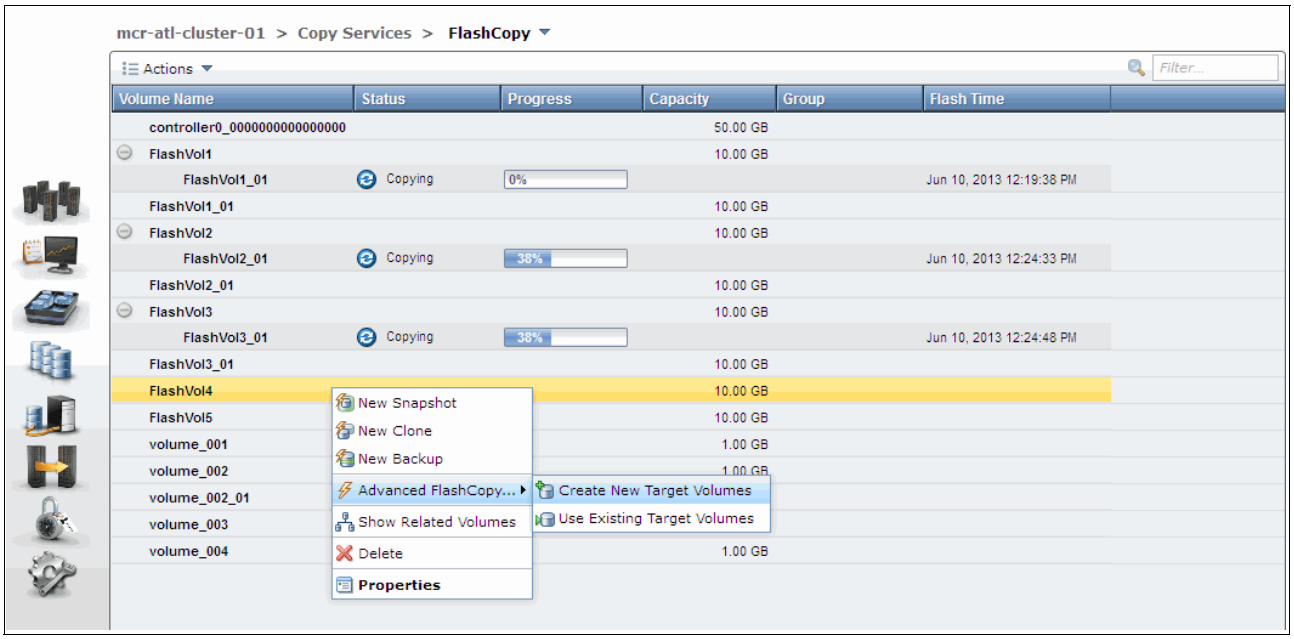

Figure 10-15 Create advanced FlashCopy

You can Create New Target Volumes as part of the mapping process or Use Existing Target Volumes. We describe creating volumes next. To use existing volumes, see “Using existing target volumes” on page 474.

Creating target volumes

Complete the following steps to create target volumes:

1. Click **Create new target volumes** if you have not yet created the target volume.
2. The wizard guides you to choose a preset, as shown in Figure 10-16. Choose one preset that has the most similar configuration to the one that is required and click **Advanced Settings** to make any appropriate adjustments to the advanced settings.

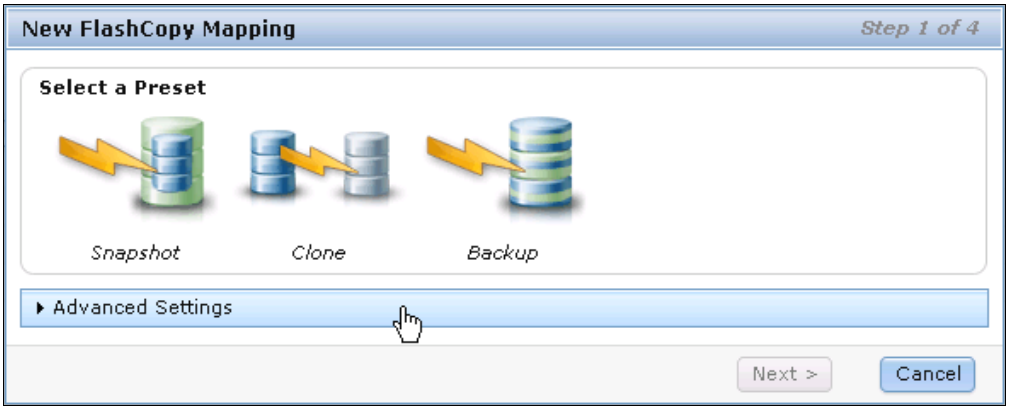

Figure 10-16 Choose a preset most similar to your requirement

The following default advanced settings for the snapshot preset are shown in Figure 10-17:

- Background Copy: 0
- Incremental: No
- Auto Delete after completion: No
- Cleaning Rate: 0

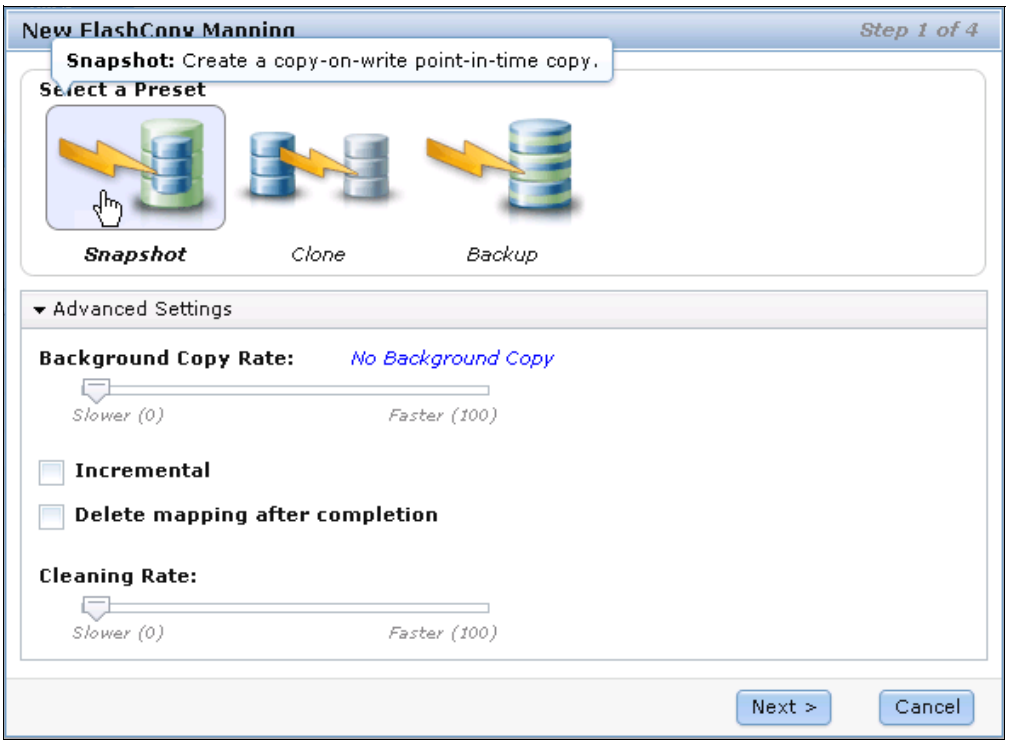

Figure 10-17 Default setting for the snapshot preset

The following Advanced Settings for the Clone Preset are shown in Figure 10-18:

- Background Copy: 50
- Incremental: No
- Auto Delete after completion: Yes
- Cleaning Rate: 50

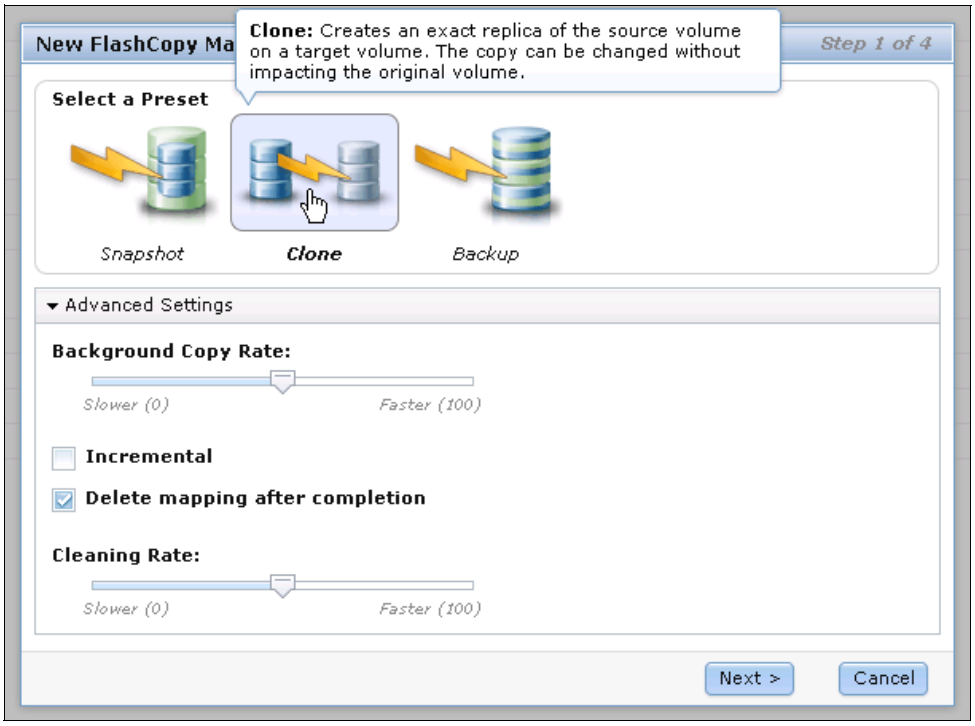

Figure 10-18 Default settings for the clone preset

Figure 10-19 shows the following Advanced Settings for the Backup preset:

- Background Copy: 50
- Incremental: Yes
- Auto Delete after completion: No
- Cleaning Rate: 50

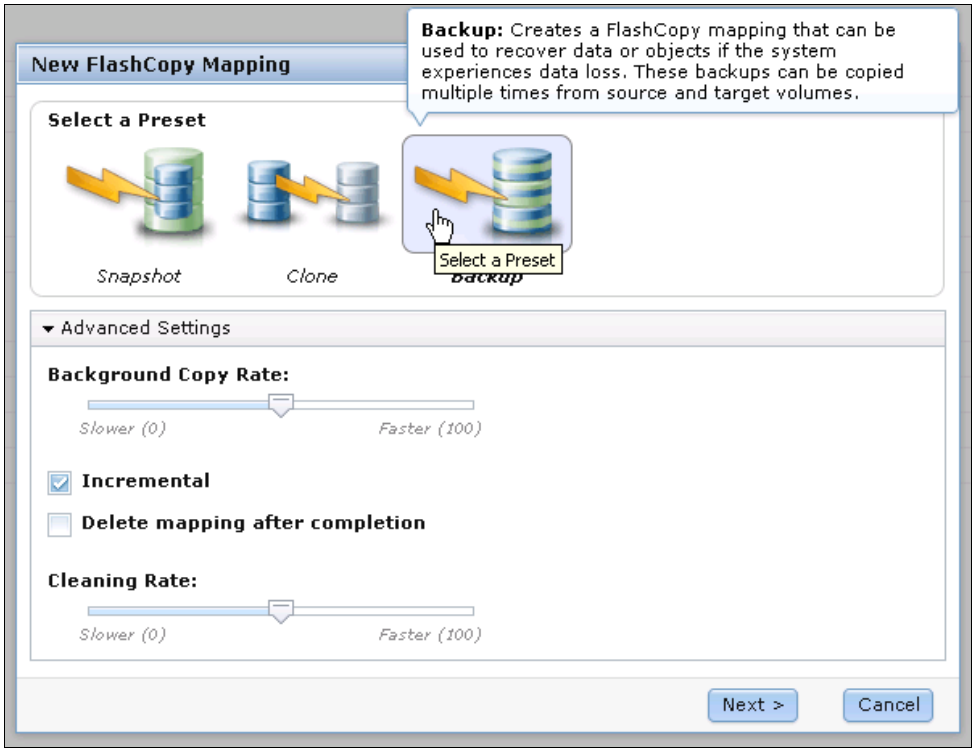

Figure 10-19 Default settings for the backup preset

3. Change the settings of the FlashCopy mapping according to your requirements and click **Next**.
4. In the next step, you can add your FlashCopy mapping to a consistency group, as shown in Figure 10-20. If the consistency group is not ready, the FlashCopy mapping can be added to the consistency group afterward. Click **Next** to continue.

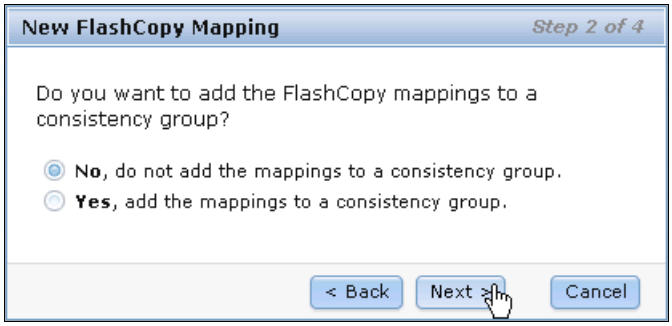

Figure 10-20 Add FlashCopy mapping to a consistency group

5. You can choose from which storage pool you want to create your target volume. As shown in Figure 10-21, you can select the same storage pool that is used by the source volume or a different pool. Click **Next** to continue.

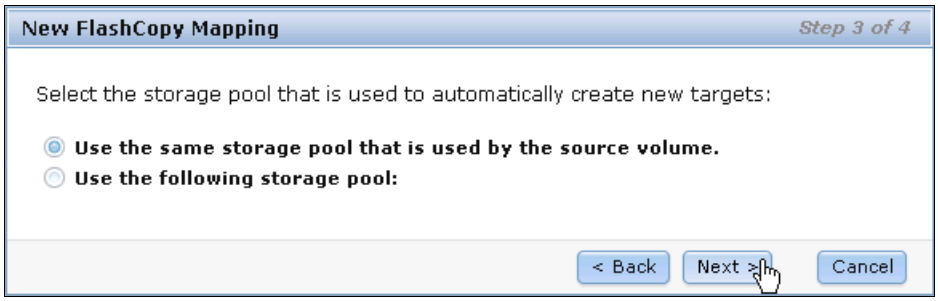

Figure 10-21 Using the same storage pool with the source volume

6. You can define how the new target volumes manage capacity. The Create a generic volume option is your default choice if you selected Clone or Backup as your basic preset. If you select a thin-provisioned volume, more options are available, as shown in Figure 10-22.

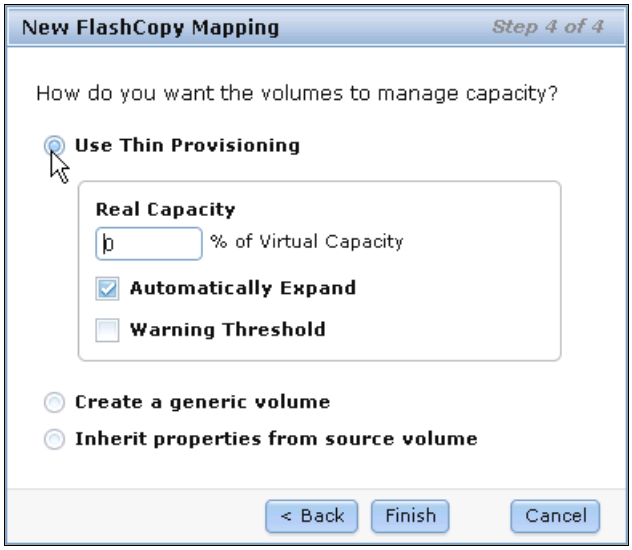

Figure 10-22 Creating a thin provisioned target volume

7. Click **Finish** when you make your decision and the mappings and volume are created, as shown in Figure 10-23.

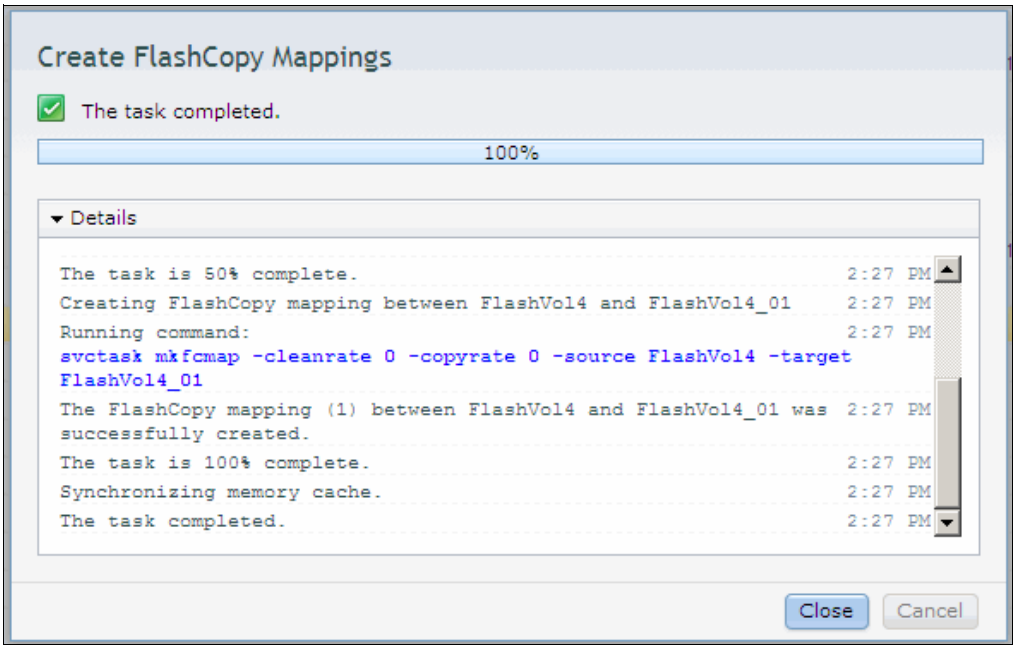

Figure 10-23 Advanced FlashCopy create task complete

8. Close the window to see the FlashCopy mapping that is created on your volume with a new target, as shown in Figure 10-24. The status of the newly created FlashCopy mapping is Idle; it can be started, as described in “Starting a FlashCopy mapping” on page 477.

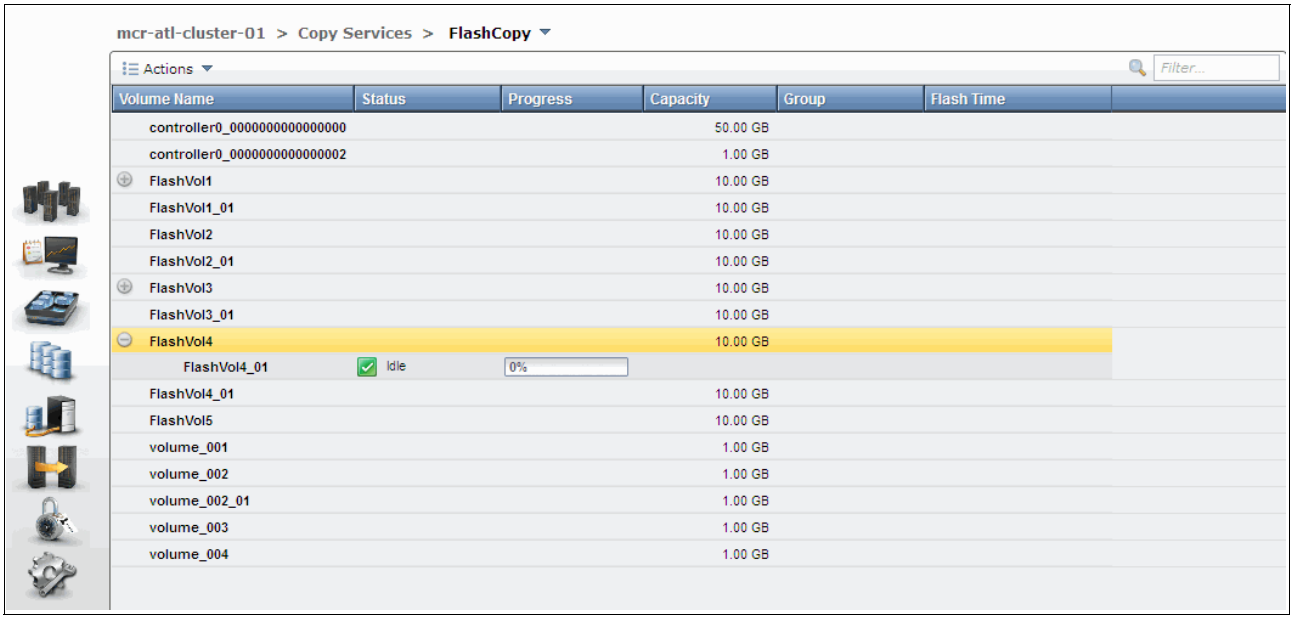

Figure 10-24 New FlashCopy mapping is created with a new target

Using existing target volumes

Complete the following steps to use existing target volumes:

1. If you already have candidate target volumes, select **Use Existing Target Volumes** in the Advanced FlashCopy menu, as shown in Figure 10-25.

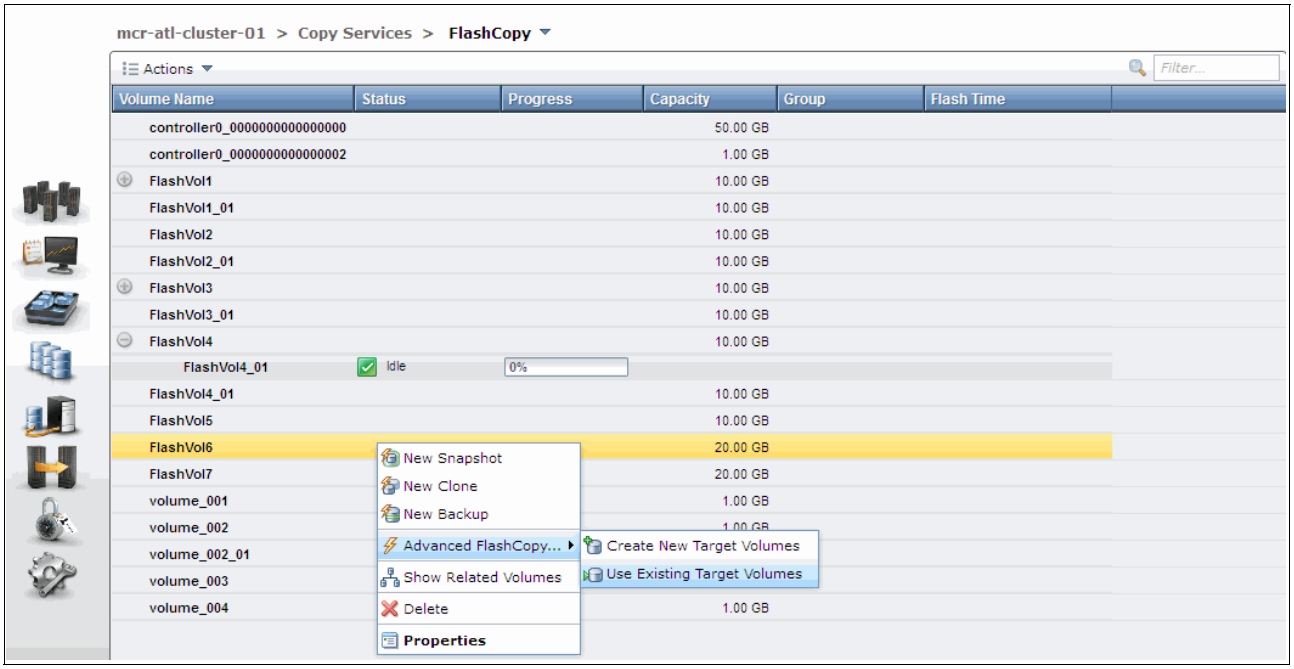

Figure 10-25 Create FlashCopy mapping by using existing target volume

2. You must choose the target volume for the source volume that you selected. Select the target volume from the drop-down menu in the right pane of the window and click **Add**, as shown in Figure 10-26.

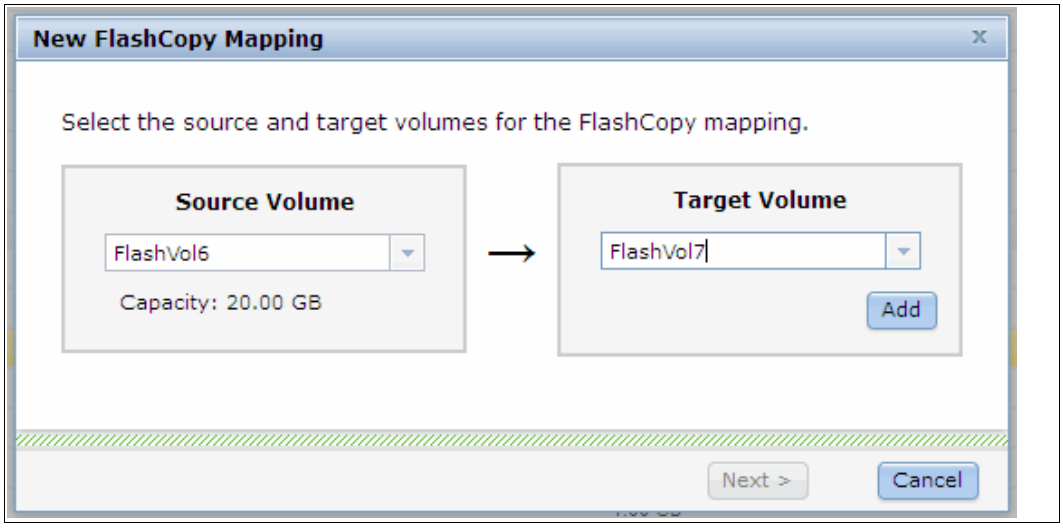

Figure 10-26 Select the target volume

3. After you click **Add**, the FlashCopy mapping is listed, as shown in Figure 10-27 on page 475. Click the red X if the FlashCopy mapping is not the one you want to create. If the FlashCopy mapping is what you want, click **Next** to continue.

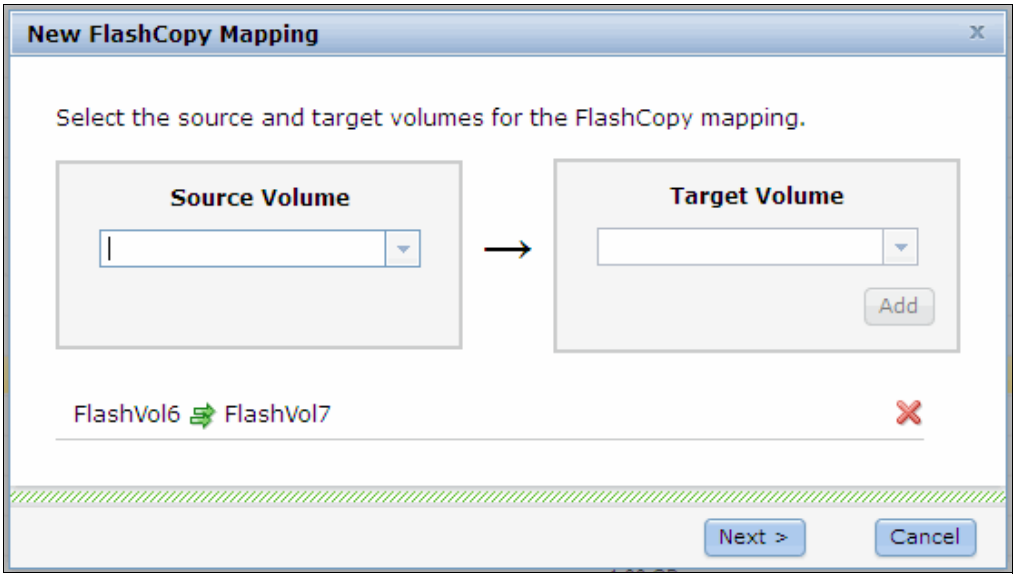

Figure 10-27 Add FlashCopy mapping

4. Select the preset and (if necessary) adjust the settings by using the Advanced Settings section as shown in Figure 10-28. (For more information about the advanced setting, see “Creating target volumes” on page 468.) Confirm that the settings meet your requirements and then click **Next**.

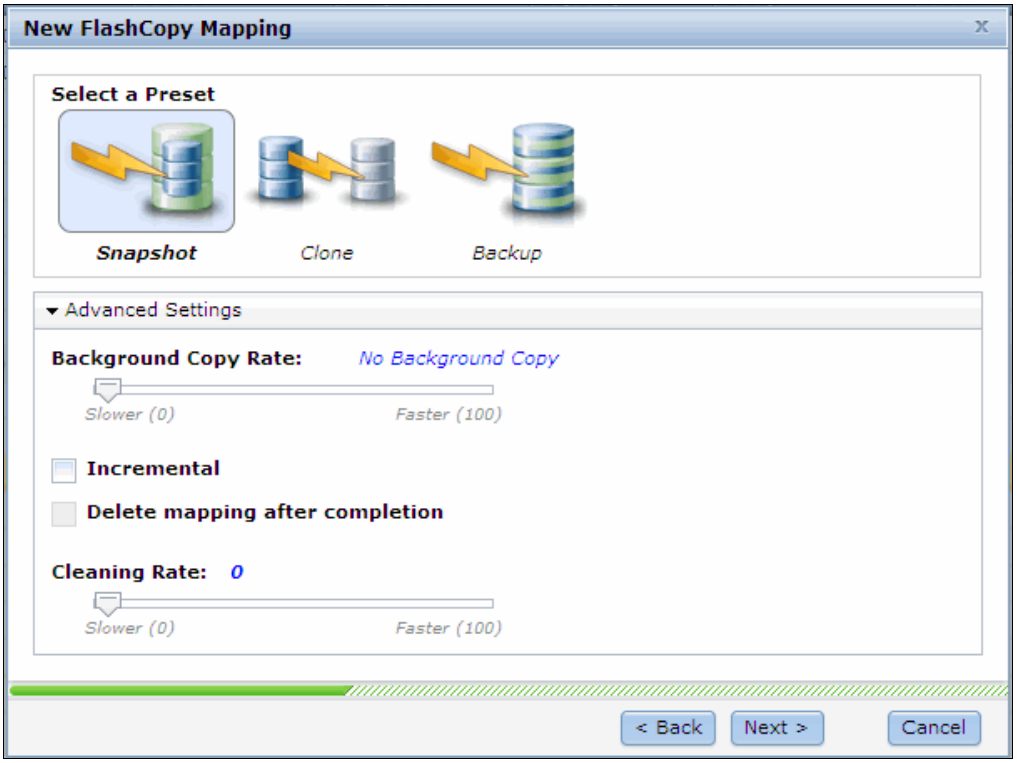

Figure 10-28 Select a preset and make your adjustments

5. You can now add the FlashCopy mapping to a consistency group (if necessary), as shown in Figure 10-29. Selecting **Yes** shows a drop-down menu from which you can select a consistency group. Click **Finish** and the FlashCopy mapping is created with the status of Idle, as shown in Figure 10-24 on page 473.

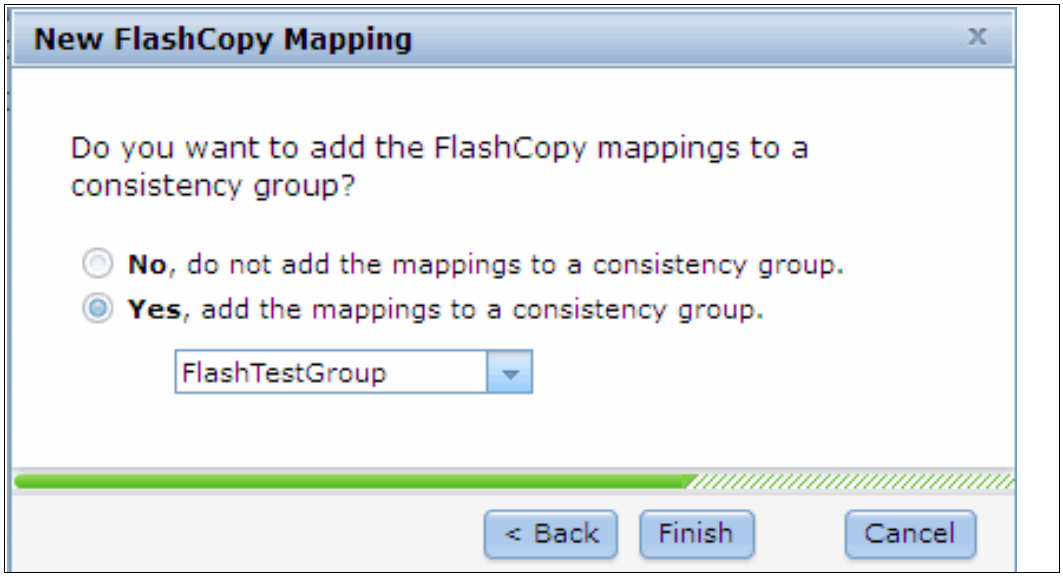

Figure 10-29 Select a consistency group to add the FlashCopy mapping

Creating new FlashCopy mappings

You can also create FlashCopy mappings in the FlashCopy Mappings window by clicking **New FlashCopy Mapping** at the upper left, as shown in Figure 10-30.

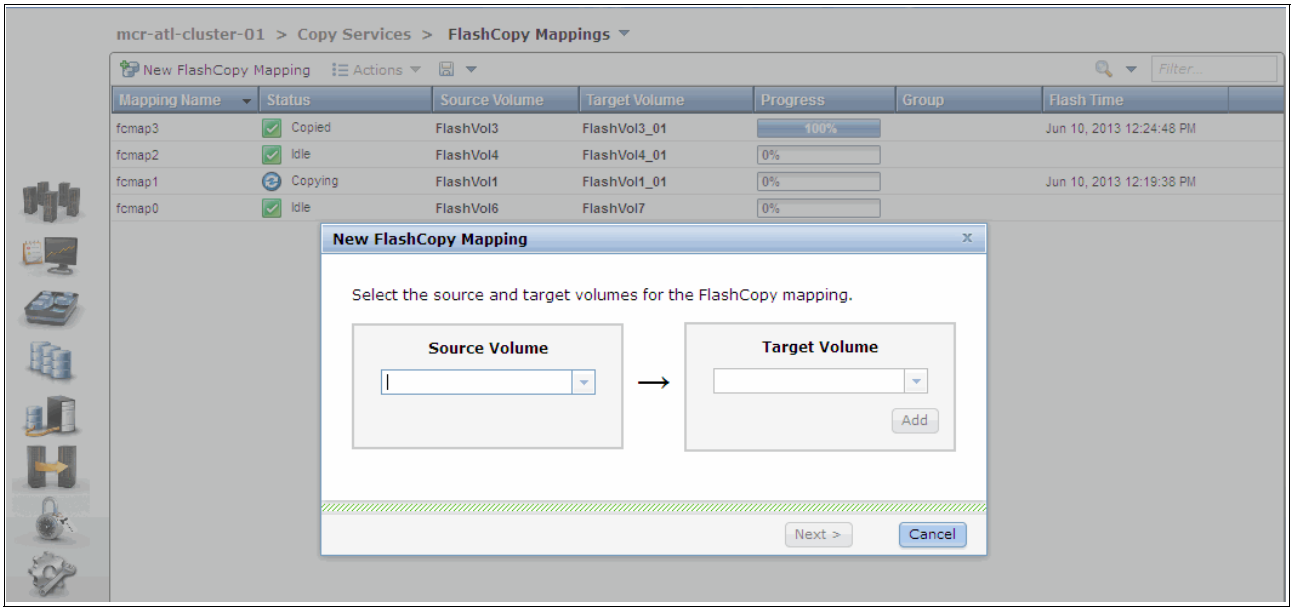

Figure 10-30 Create a FlashCopy mapping in the FlashCopy Mappings window

A wizard guides you through the process to create a FlashCopy mapping. The steps are the same as creating an Advanced FlashCopy mapping by using Existing Target Volumes, as described in “Using existing target volumes” on page 474.

Starting a FlashCopy mapping

Most of the FlashCopy mapping actions can be performed in the FlashCopy window or the FlashCopy Mapping window. For the actions that are available in both windows, we show in the following sections the steps in the FlashCopy window, although the steps are the same if you were to use the FlashCopy Mapping window.

You can start the mapping by selecting the FlashCopy target volume in the FlashCopy window and selecting the **Start** option from the Actions drop-down menu (as shown in Figure 10-31) or by selecting the volume and right-clicking. The status of the FlashCopy mapping changes from Idle to Copying.

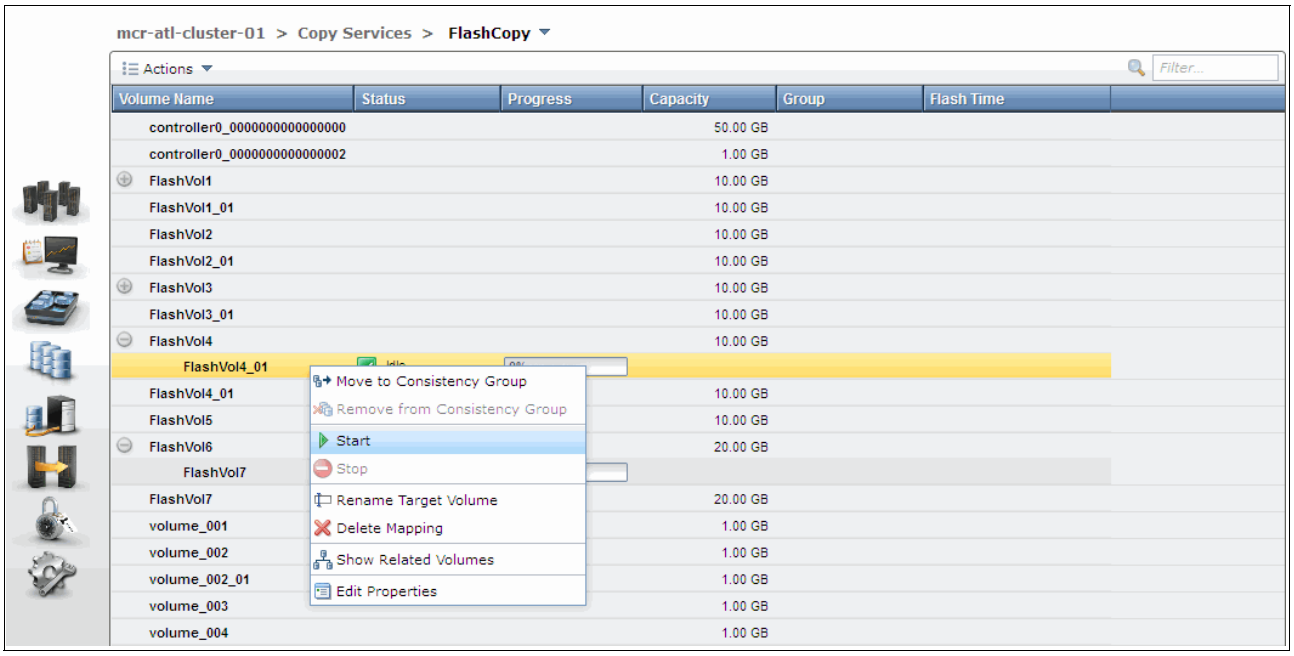

Figure 10-31 Start FlashCopy mapping

Stopping a FlashCopy mapping

The FlashCopy mapping can be stopped by selecting the FlashCopy target volume in the FlashCopy window and clicking the **Stop** option from the Actions drop-down menu, as shown in Figure 10-32 on page 478. After the stopping process completes, the status of the FlashCopy mapping is changed to Stopped.

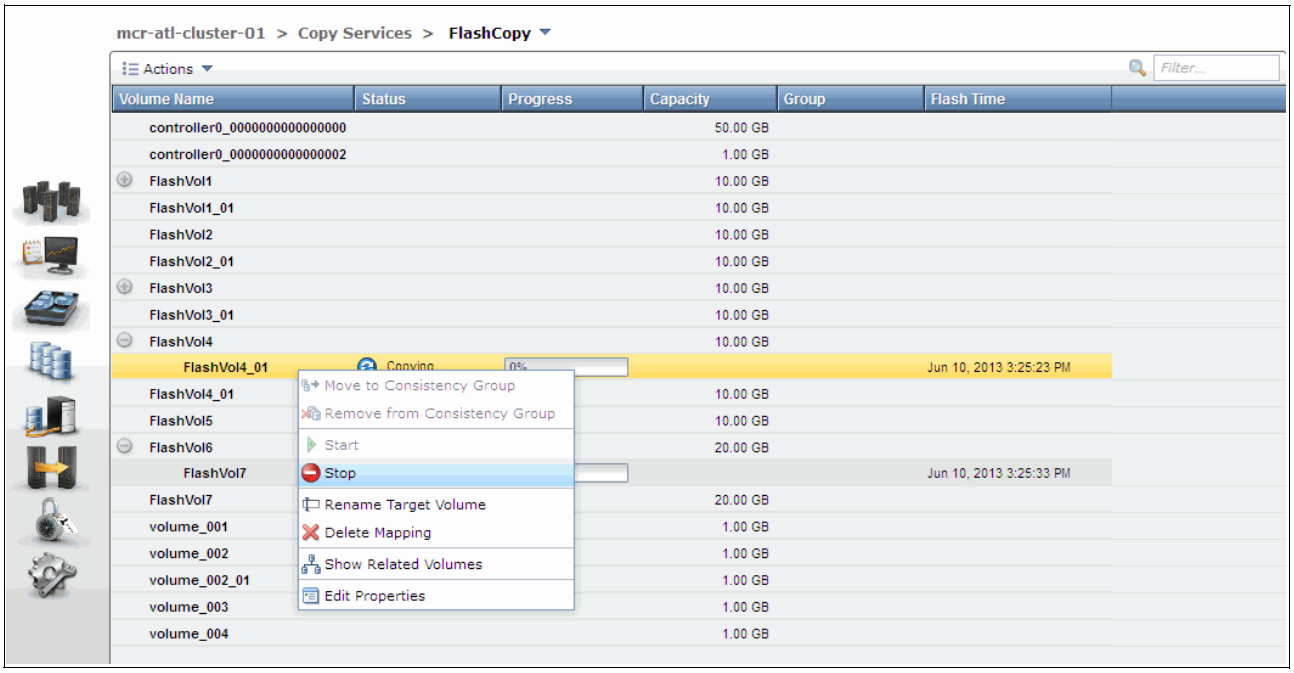

Figure 10-32 Stopping a FlashCopy mapping

Renaming the target volume

If the FlashCopy target volumes were created automatically by IBM Storwize V5000 storage system, the name of the target volume is the source volume name plus a suffix that includes numbers. The name of the target volumes can be changed to be more meaningful in your environment.

To change the name of the target volume, select the FlashCopy target volume in the FlashCopy window and click the **Rename Target Volume** option from the Actions drop-down menu (as shown in Figure 10-33) or right-click the selected volume.

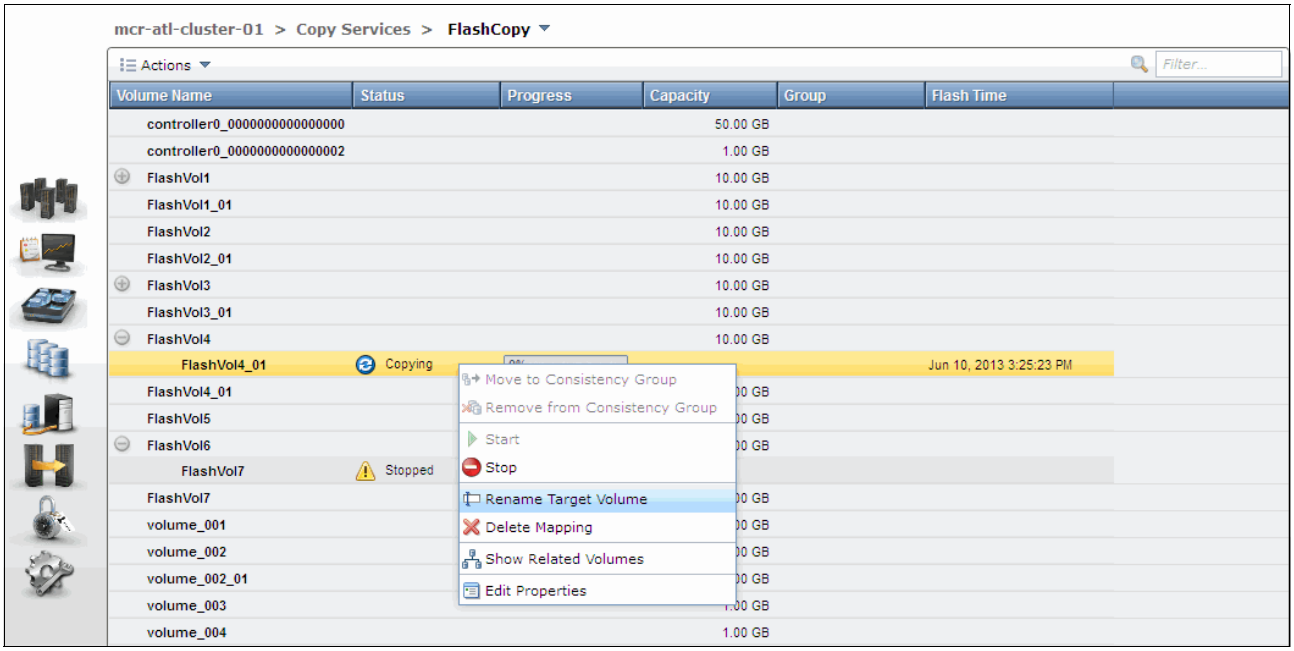

Figure 10-33 Rename a target volume

Enter your new name for the target volume, as shown in Figure 10-34. Click **Rename** to finish.

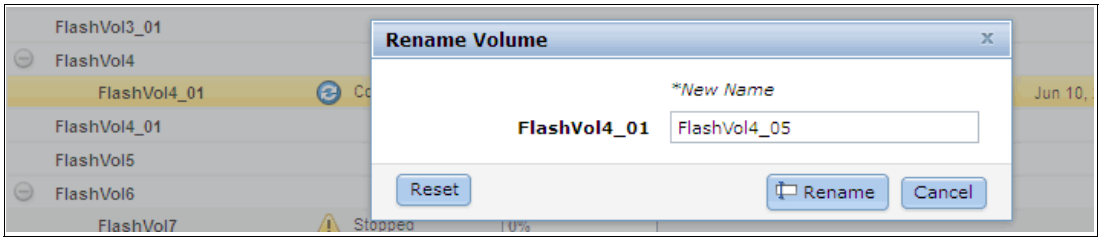

Figure 10-34 Rename a target volume

Renaming a FlashCopy mapping

The FlashCopy mappings are created with names that begin with `fcmap`. The name of FlashCopy mappings can be changed to be more meaningful to you.

To change the name of a FlashCopy mapping, select the FlashCopy mapping in the FlashCopy Mappings window and click **Rename Mapping** in the Actions drop-down menu, as shown in Figure 10-35.

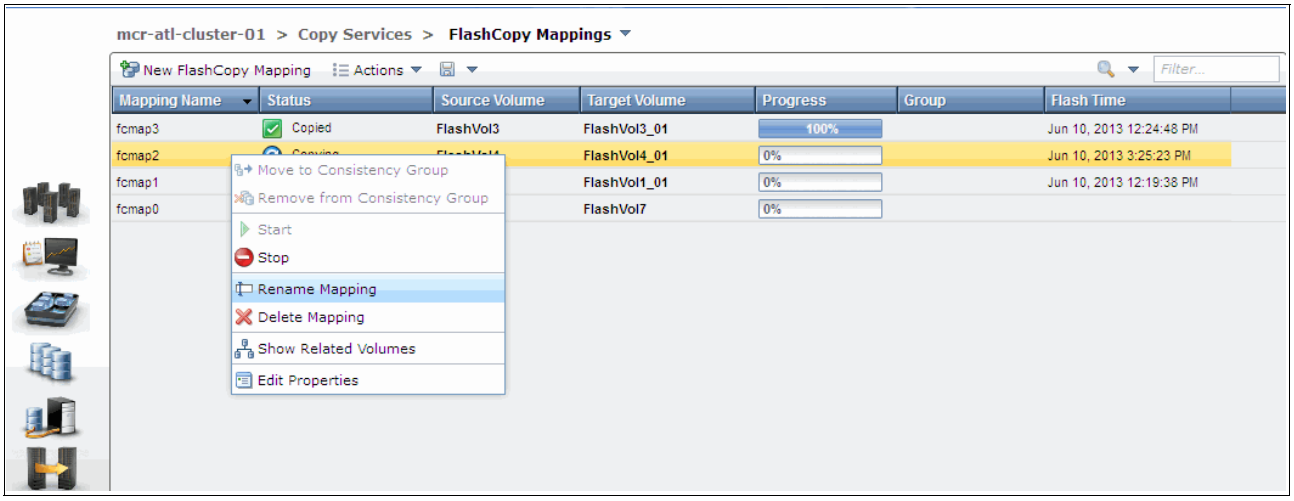

Figure 10-35 Rename a FlashCopy mapping

You must enter the new name for the FlashCopy mapping, as shown in Figure 10-36. Click **Rename** to finish.

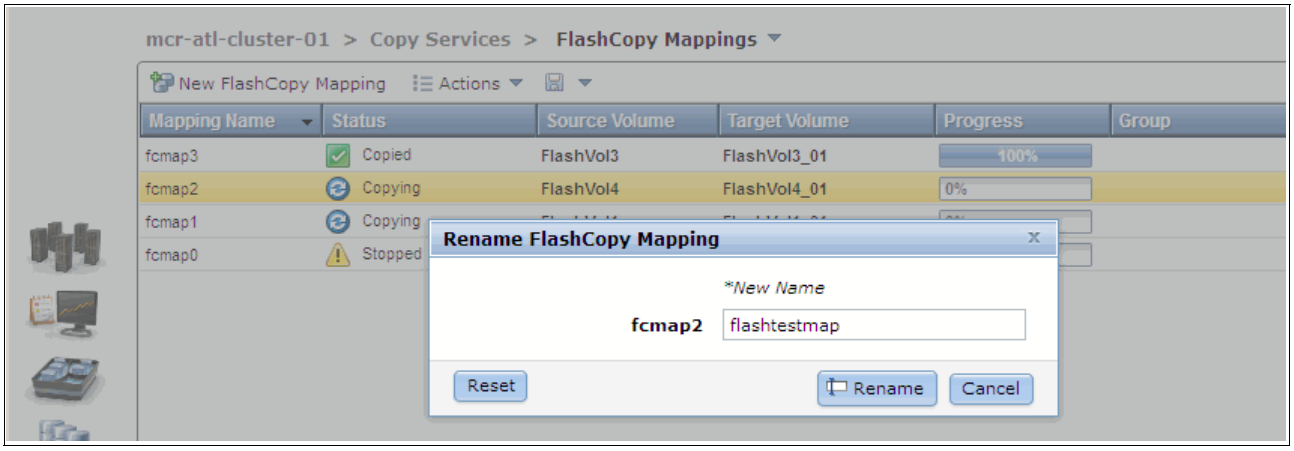

Figure 10-36 Enter a new name for the FlashCopy mapping

Deleting a FlashCopy mapping

The FlashCopy mapping can be deleted by selecting the FlashCopy target volume in the FlashCopy window and clicking **Delete Mapping** in the Actions drop-down menu (as shown in Figure 10-37) or by right-clicking the selected volume.

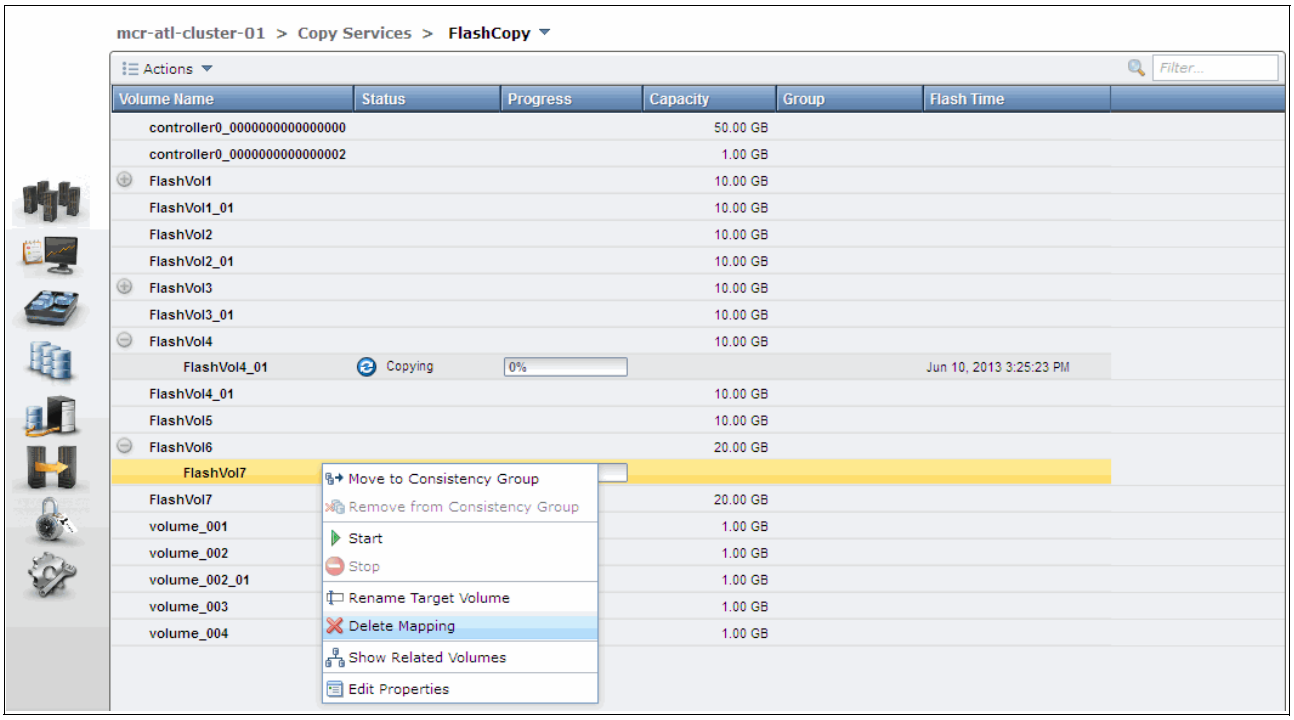

Figure 10-37 Select Delete Mapping

FlashCopy Mapping state: If the FlashCopy mapping is in the Copying state, it must be stopped before it is deleted.

You must confirm your action to delete FlashCopy mappings in the window that opens, as shown in Figure 10-38. Verify the number of FlashCopy mappings that you must delete. If you want to delete the FlashCopy mappings while the data on the target volume is inconsistent with the source volume, select the option to do so. Click **Delete** and your FlashCopy mapping is removed.

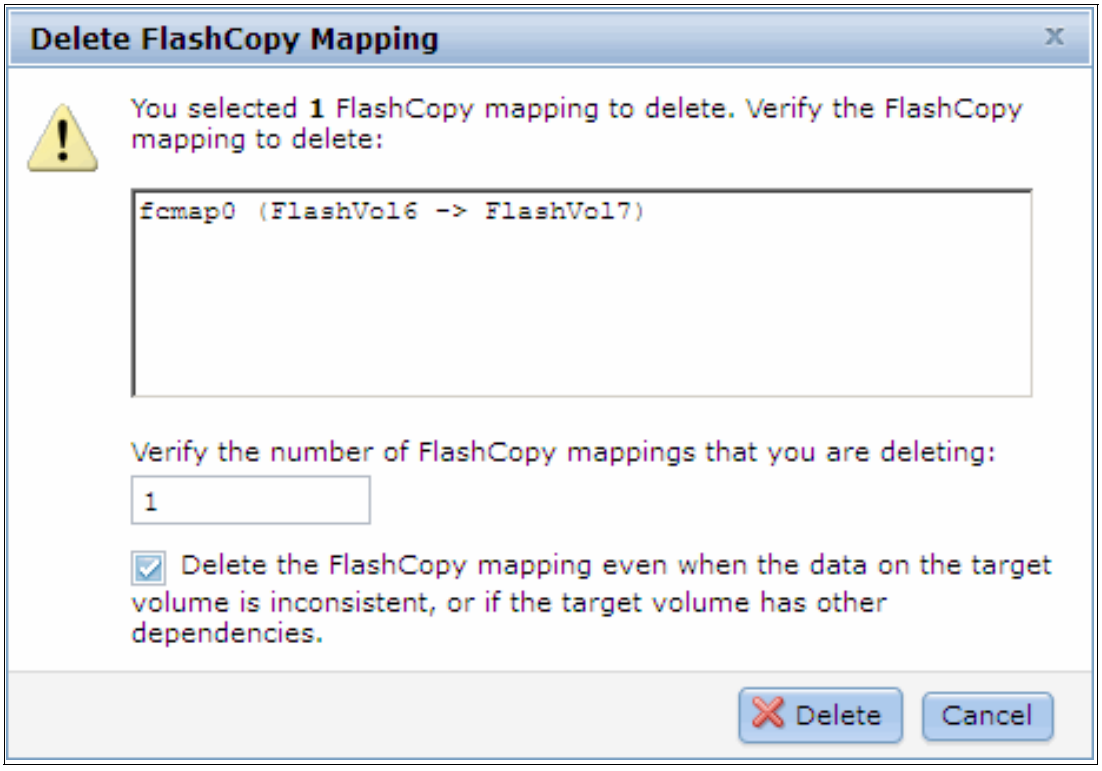

Figure 10-38 Confirm the deletion of FlashCopy mappings

Deleting FlashCopy mapping: Deleting the FlashCopy mapping does not delete the target volume. If you must reclaim the storage space that is occupied by the target volume, you must delete the target volume manually.

Showing related volumes

You can show the FlashCopy mapping dependencies by selecting a target or source volume in the FlashCopy window and clicking **Show Related Volumes** in the Actions drop-down menu (as shown in Figure 10-39) or right-clicking the selected volume.

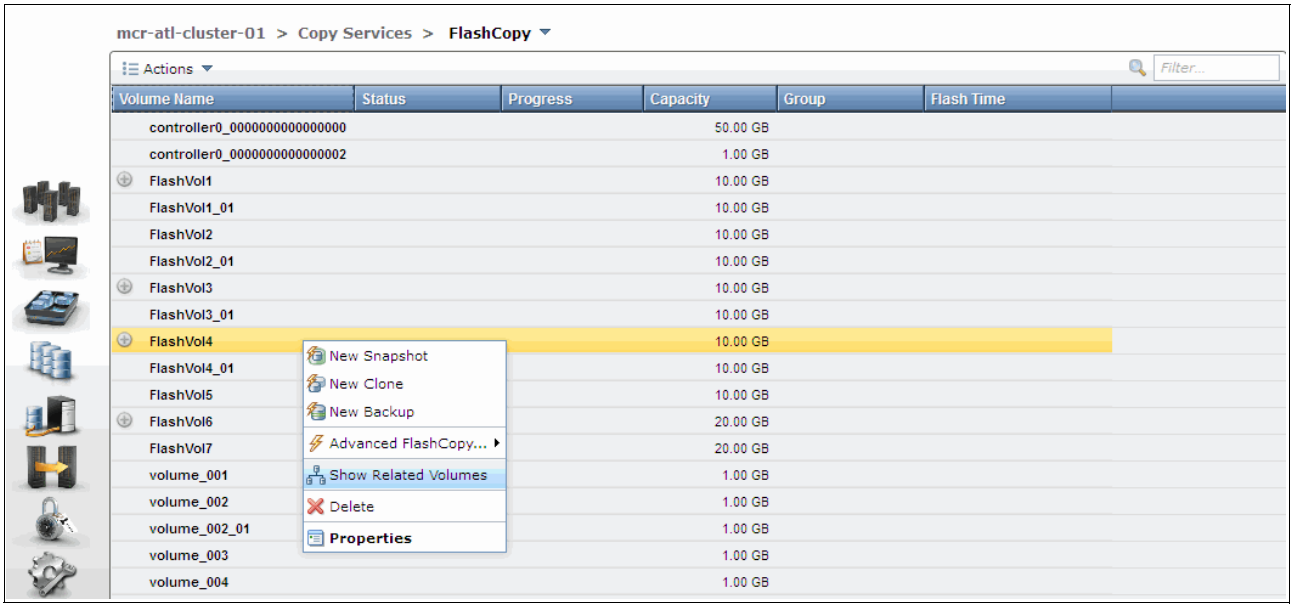

Figure 10-39 Show Related Volumes menu

The FlashCopy mapping dependency tree opens, as shown in Figure 10-40.

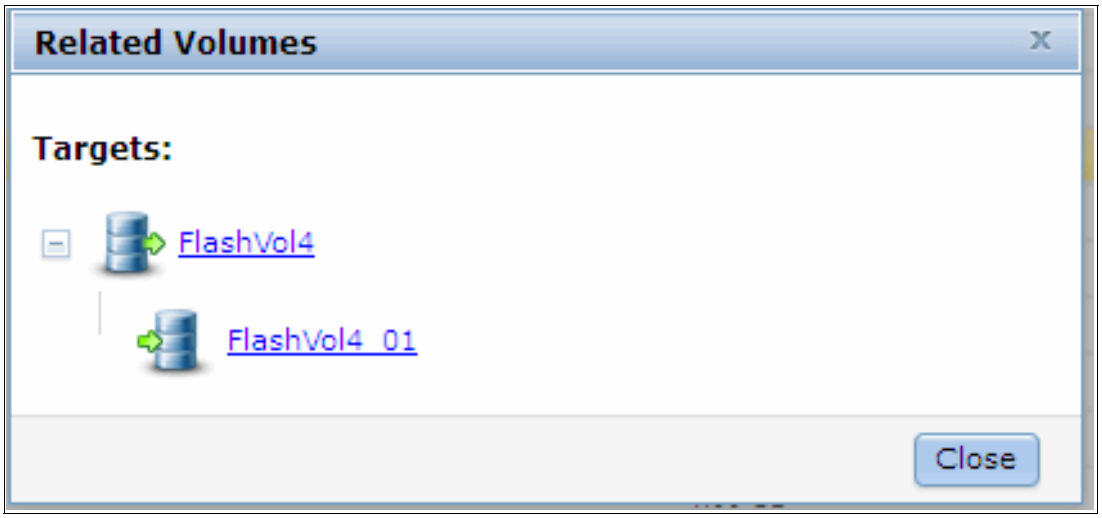

Figure 10-40 FlashCopy mapping dependency

Clicking either volume shows the properties of the volume, as shown in Figure 10-41.

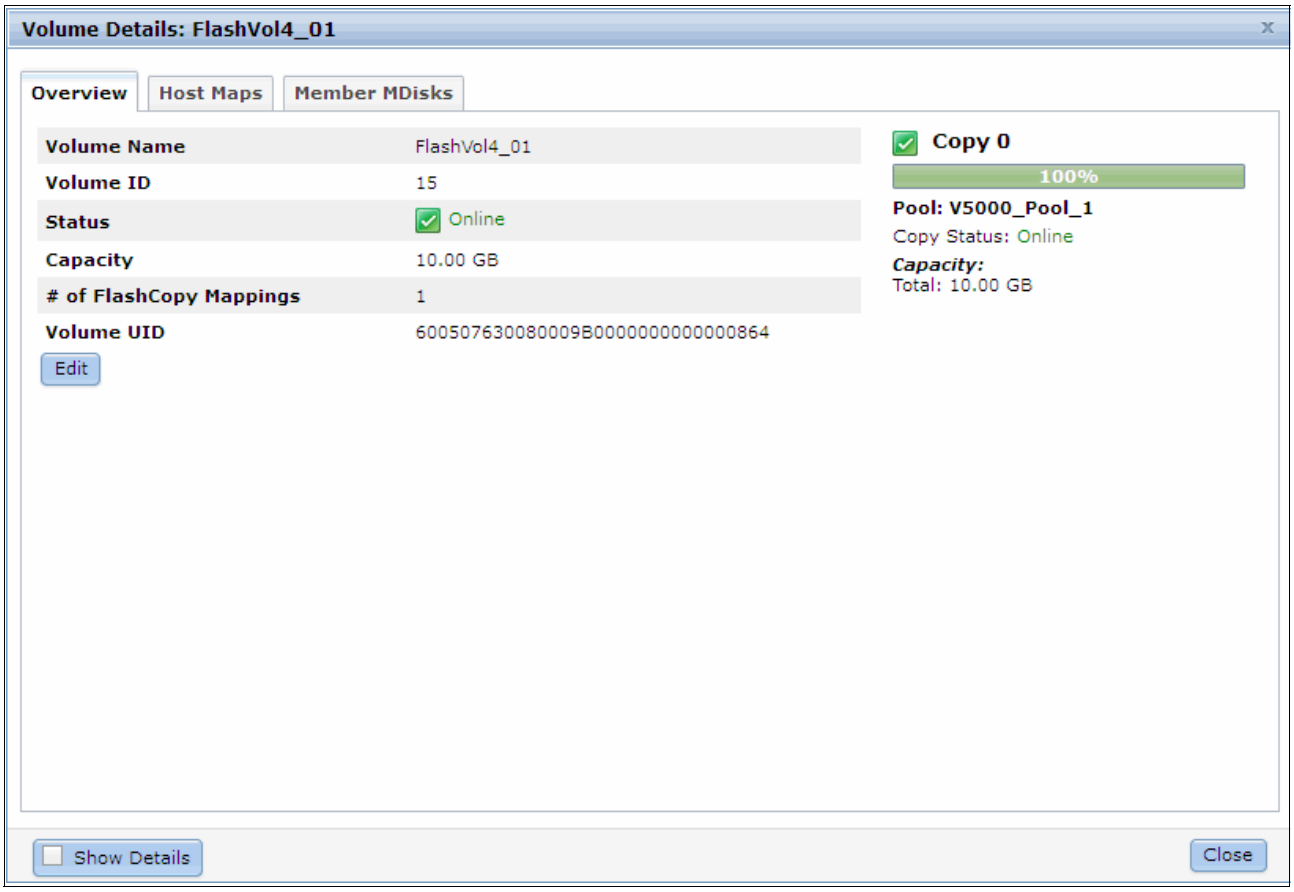

Figure 10-41 Target FlashCopy Volume details

Editing properties

The background copy rate and cleaning rate can be changed after the FlashCopy mapping is created. Select the FlashCopy target mapping in the FlashCopy window and click **Edit Properties** in the Actions drop-down menu (as shown in Figure 10-42) or right-click.

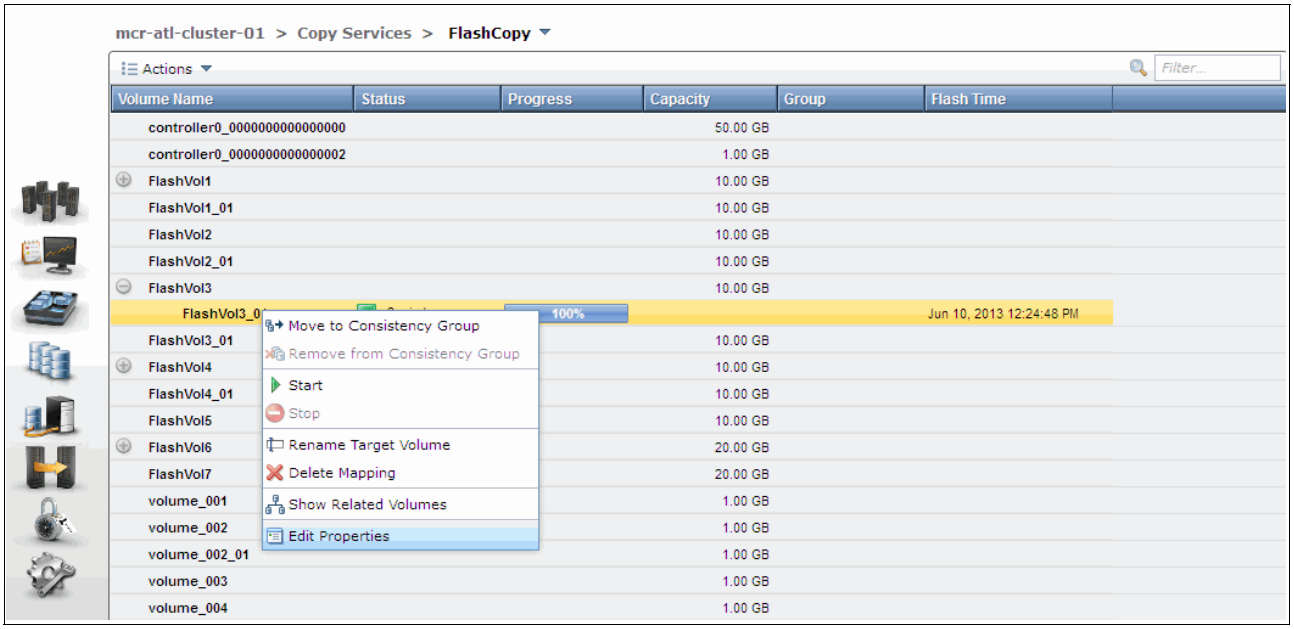

Figure 10-42 Edit Properties menu

You can then modify the value of the background copy rate and cleaning rate by moving the pointers on the bars, as shown in Figure 10-43. Click **Save** to save changes.

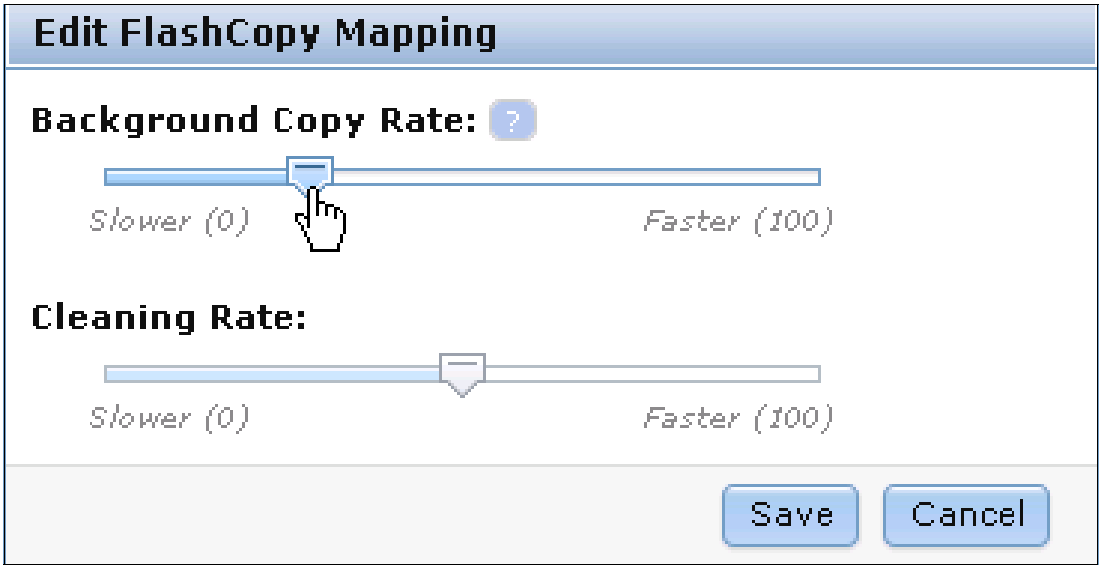

Figure 10-43 Change the copy rate

Restoring from a FlashCopy

Complete the following steps to manipulate FlashCopy target volumes to restore a source volume to a previous known state:

1. Identify the FlashCopy relationship that you want to restore. In our example, we want to restore `FlashVol1`, as shown in Figure 10-44.

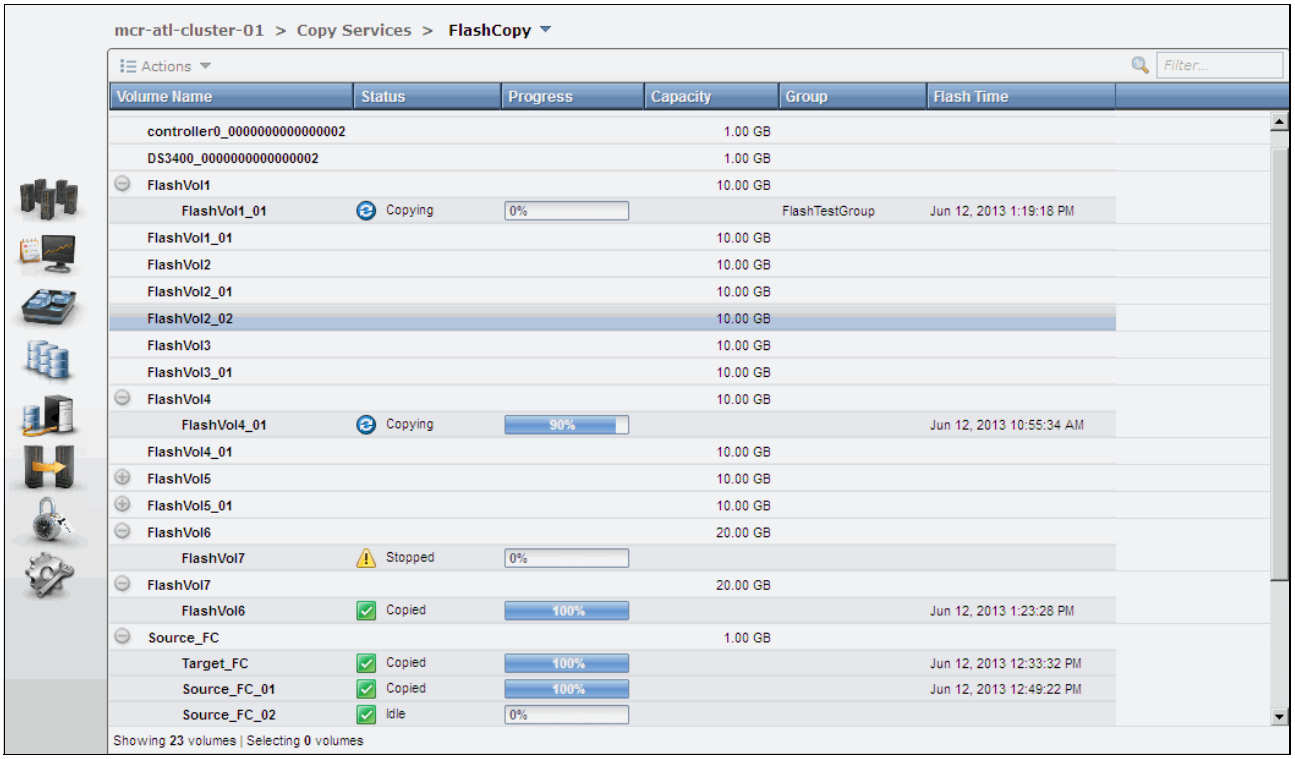

Figure 10-44 Starting FlashCopy restore

2. Create a mapping by using the target volume of the mapping to be restored. In our example, it is `FlashVol1_01`, as shown in Figure 10-45. Select **Advanced FlashCopy** → **Use Existing Target Volumes**.

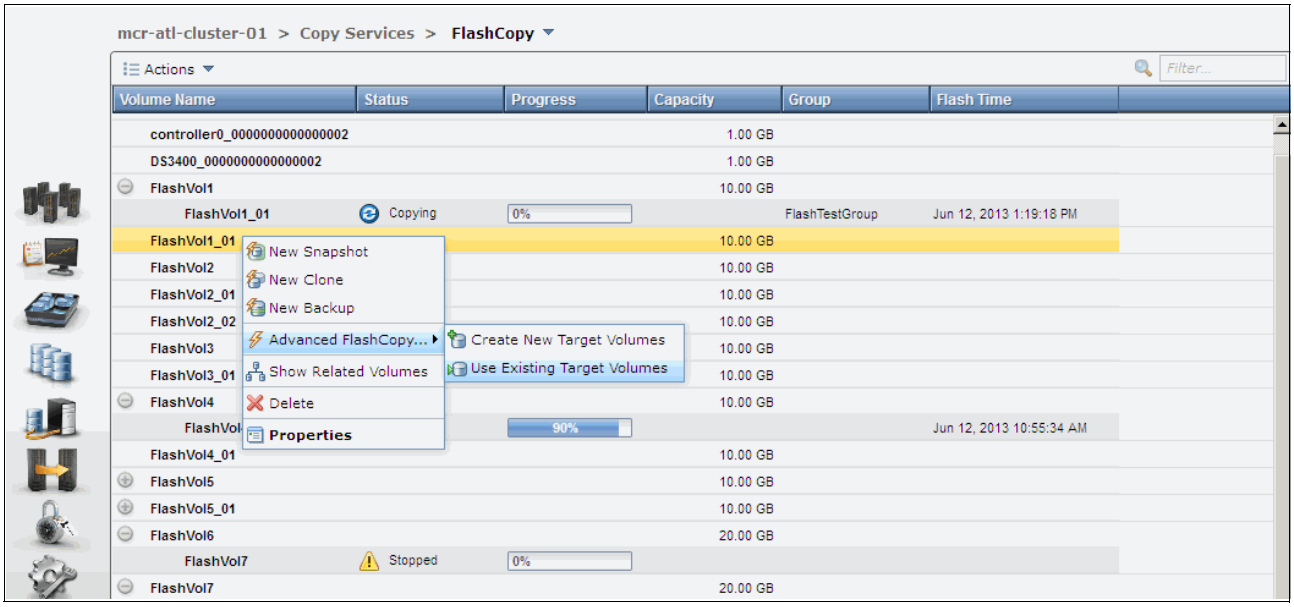

Figure 10-45 Create reverse mapping

3. The Source Volume is preselected with the target volume that was selected in the previous step. Select the Target Volume from the drop-down menu (you select the source volume that you want to restore). In our example, we select `FlashVol1`, as shown in Figure 10-46.

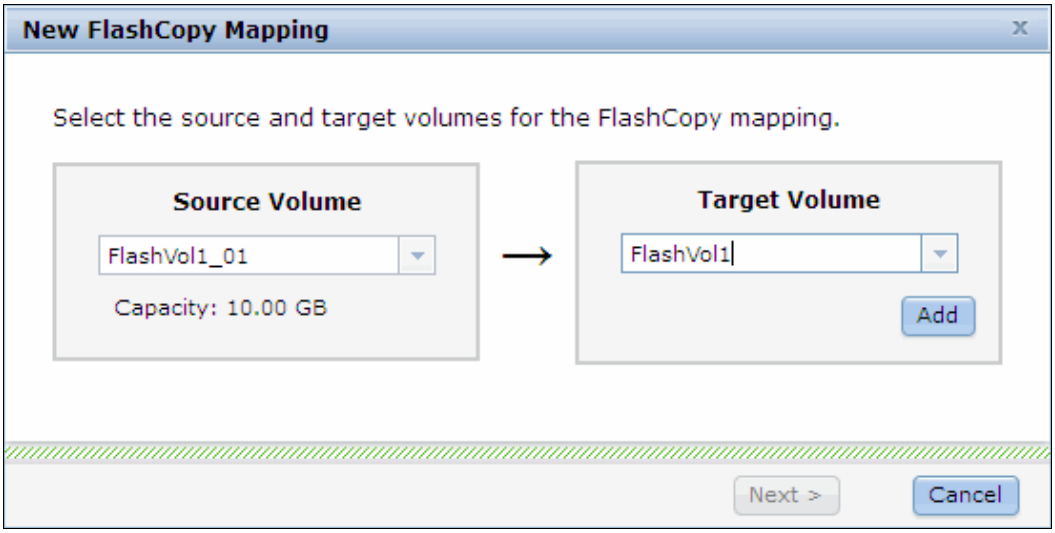

Figure 10-46 Select target volume

4. Click **Add**. A warning message appears, as shown in Figure 10-47. Click **Close.** This message is shown because we are using a source as a target.

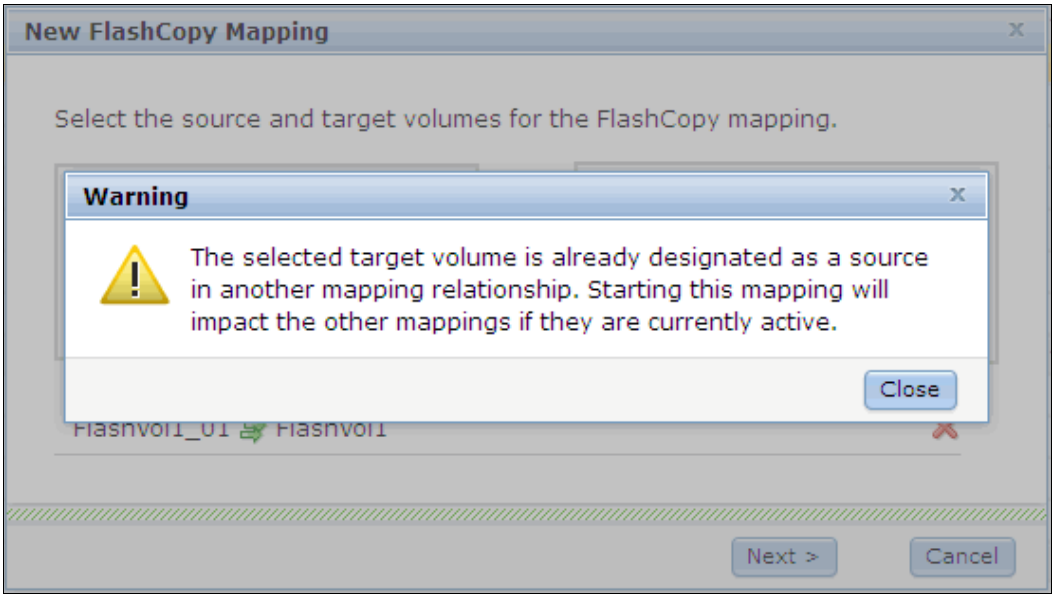

Figure 10-47 Flash restore warning

5. Click **Next** and you see a snapshot preset choice, as shown in Figure 10-48.

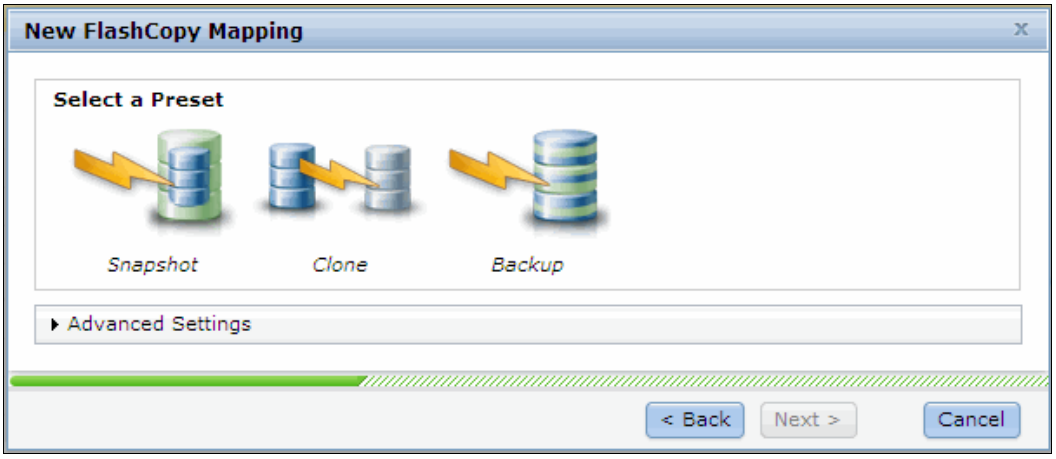

Figure 10-48 Choose snapshot preset

Select **Snapshot** and click **Next**

6. In the next window, you are asked if the new mapping is to be part of a consistency group, as shown in Figure 10-49. In our example, the new mapping is not part of a consistency group, so we click **No** and then **Finish** to create the mapping.

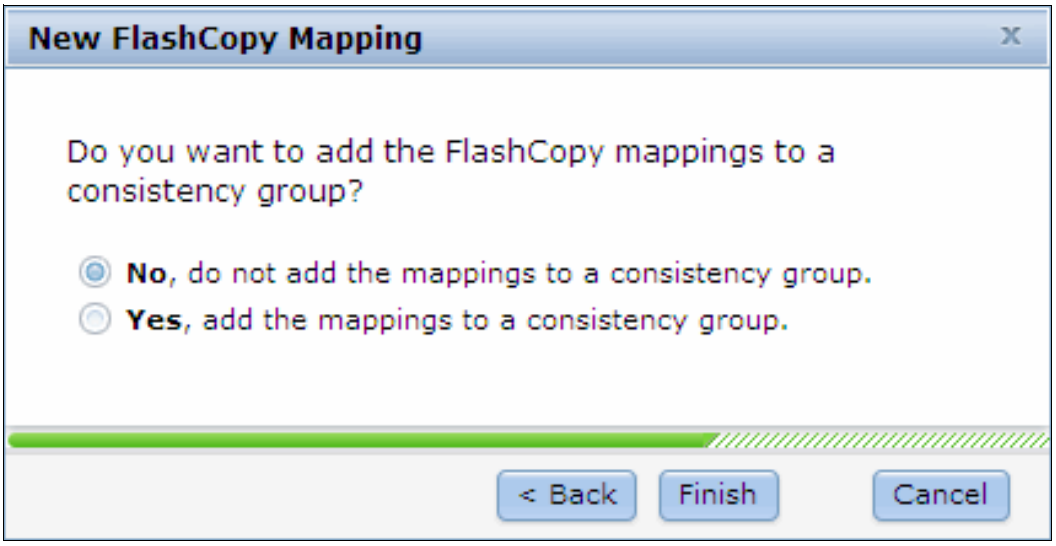

Figure 10-49 Add new mapping to consistency group

7. The new reverse mapping is now created and shown in the Idle state, as shown in Figure 10-50.

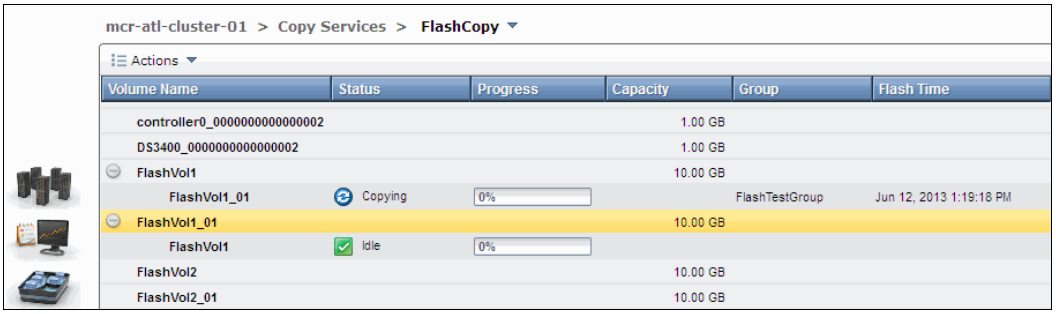

Figure 10-50 New reverse mapping

8. To restore the original source volume `FlashVol1` with the snapshot we took (`FlashVol1_01`), we select the new mapping and right-click to open the Actions menu, as shown in Figure 10-51.

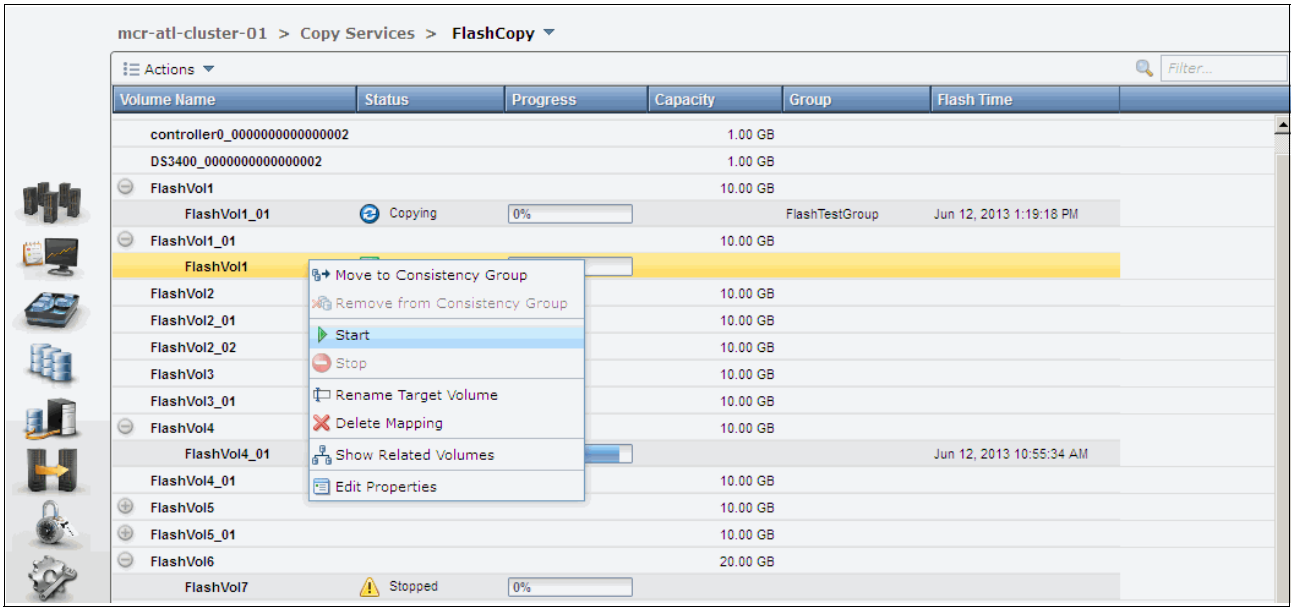

Figure 10-51 Starting the reverse mapping

9. Click **Start** to write over `FlashVol1` with the original bitmap data that was saved in the FlashCopy `FlashVol01_01`. The command then completes, as shown in Figure 10-52.

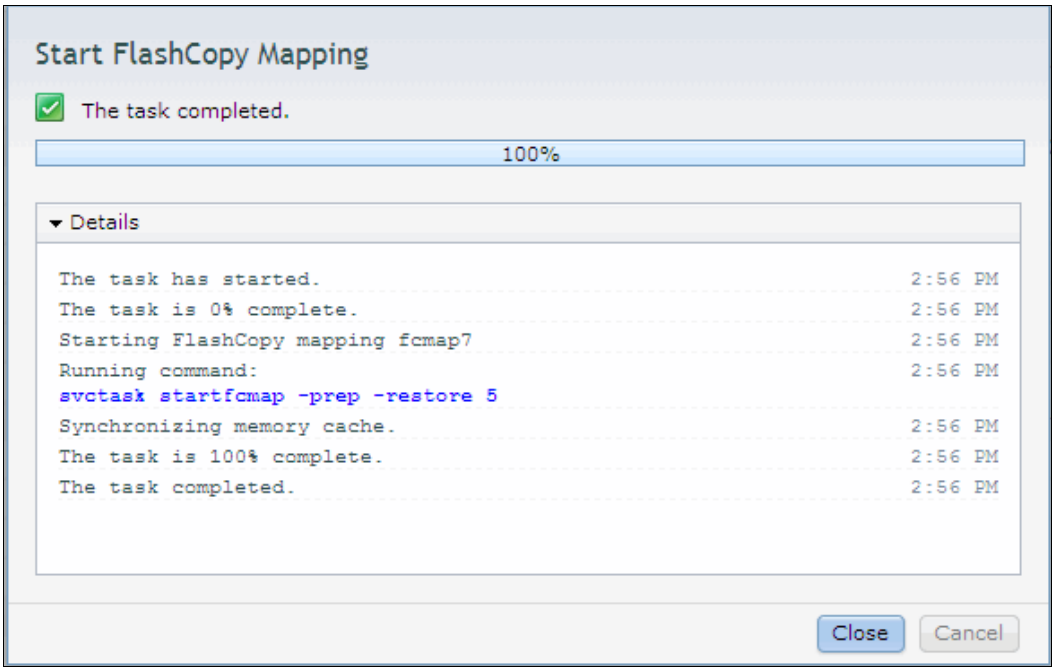

Figure 10-52 Flash Restore command

Important: The underlying command that is run by the IBM Storwize V5000 appends the **`-restore`** option automatically.

10. The reverse mapping now shows as 100% copied, as shown in Figure 10-53.

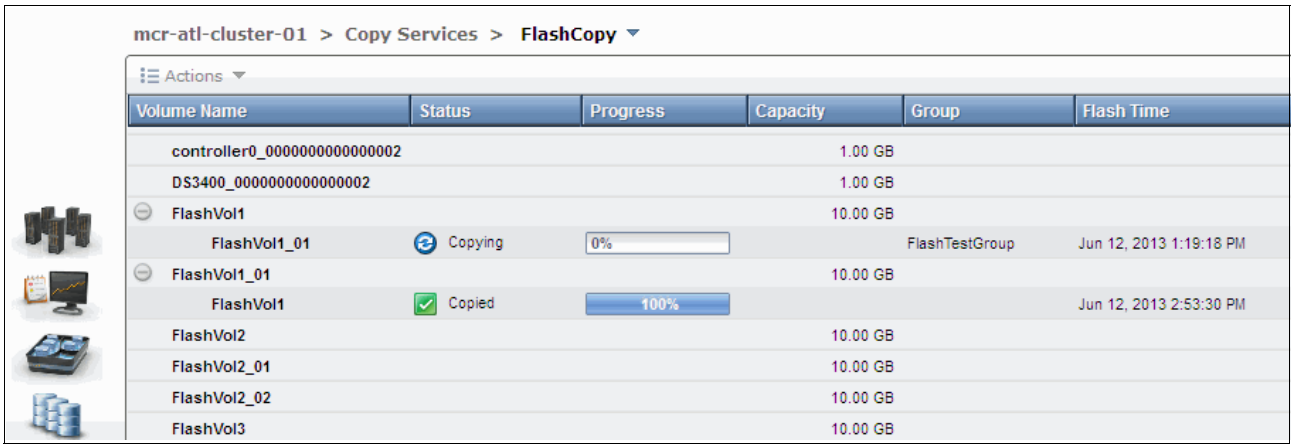

Figure 10-53 Source volume restore complete

10.1.6 Managing a FlashCopy consistency group

FlashCopy consistency groups can be managed by clicking **Consistency Groups** under the Copy Services function icon, as shown in Figure 10-54.

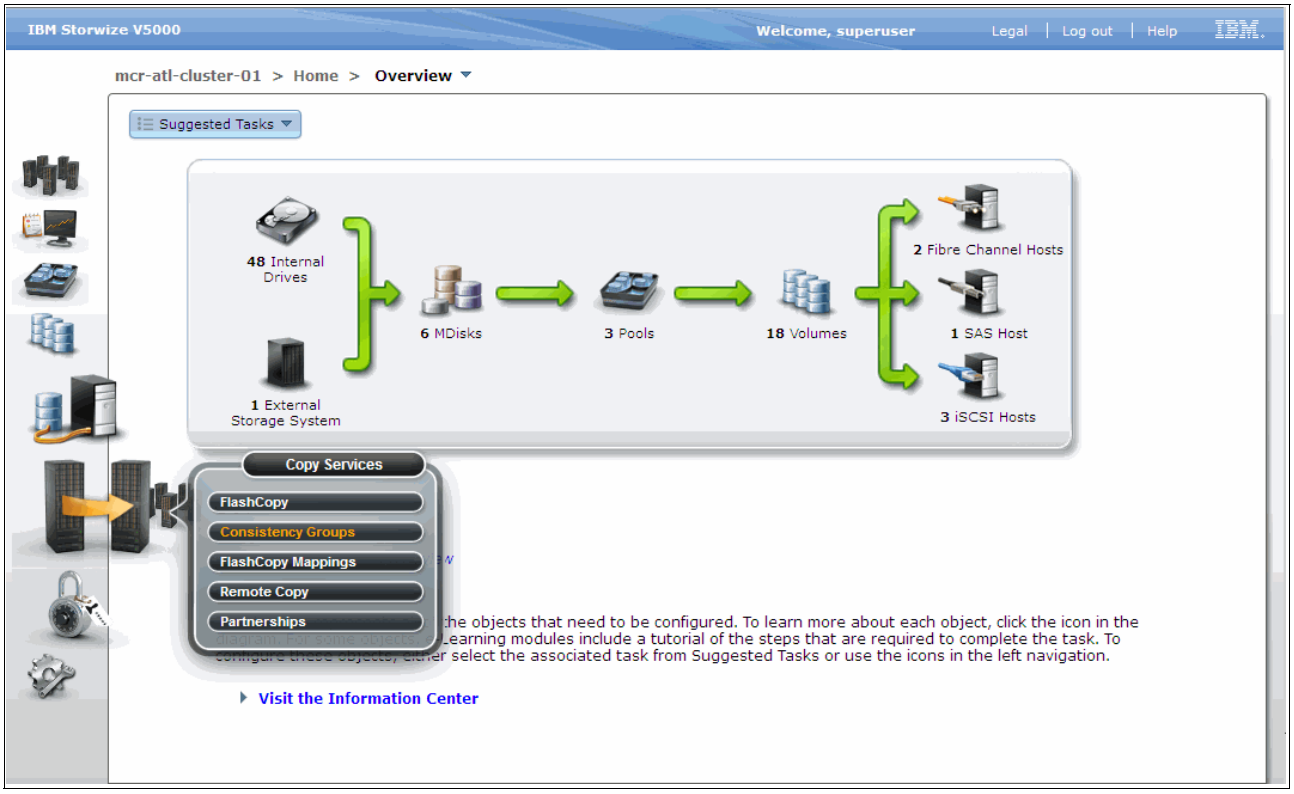

Figure 10-54 Access to the Consistency Groups window

As shown in Figure 10-55, the Consistency Groups window is where you can manage consistency groups and FlashCopy mappings.

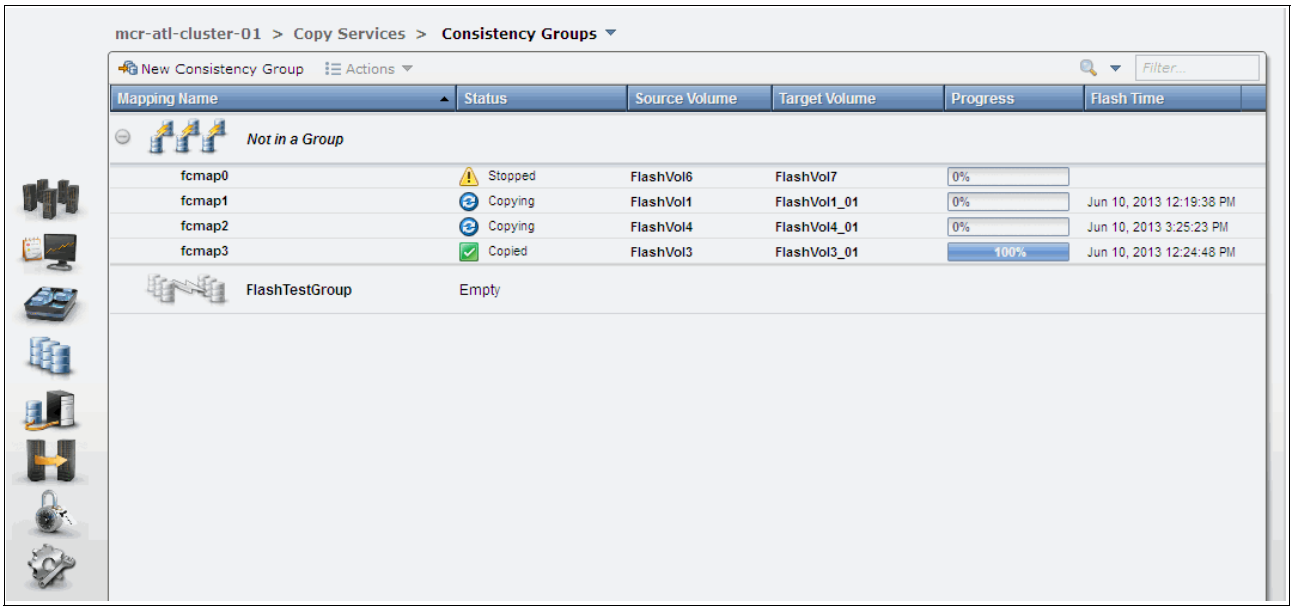

Figure 10-55 Consistency Groups window

In the left pane of the Consistency Groups window, you can list the consistency groups that you need. Click **Not in a Group**, and then expand your selection by clicking the plus (+) icon next to it. All the FlashCopy mappings that are not in any consistency groups are displayed underneath.

In the lower pane of the Consistency Groups window, you can discover the properties of a consistency group and the FlashCopy mappings in it. You can also take action on any consistency groups and FlashCopy mappings within the Consistency Groups window, as allowed by their state. For more information, see 10.1.5, “Managing FlashCopy mappings” on page 467.

Creating a FlashCopy consistency group

To create a FlashCopy consistency group, click **New Consistency Group** at the top of the Consistency Groups window, as shown in Figure 10-56.

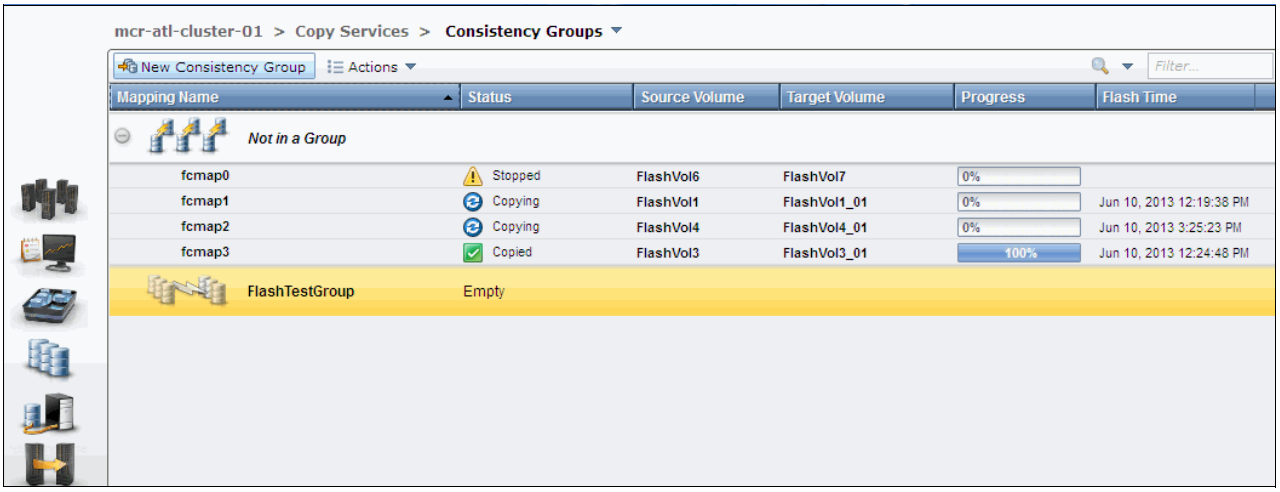

Figure 10-56 New Consistency Group option

You are prompted to enter the name of the new consistency group, as shown in Figure 10-57. Following your naming conventions, enter the name of the new consistency group in the name field and click **Create**.

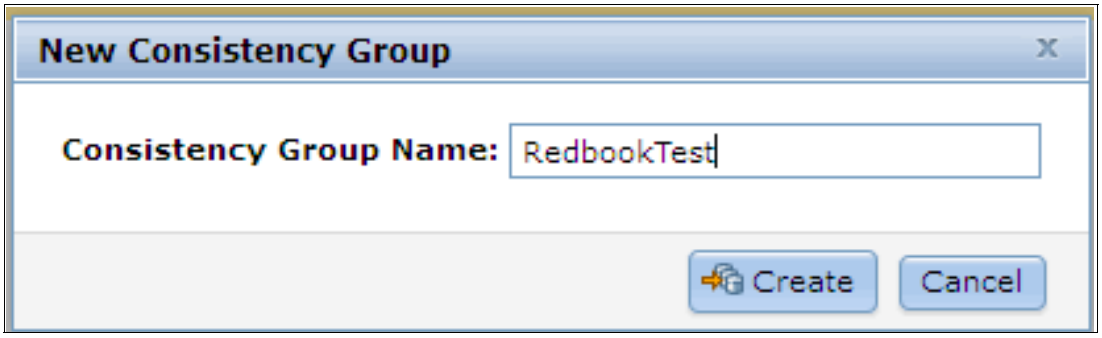

Figure 10-57 Entering the name for the consistency group

After the creation process completes, you find a new consistency group, as shown in Figure 10-58.

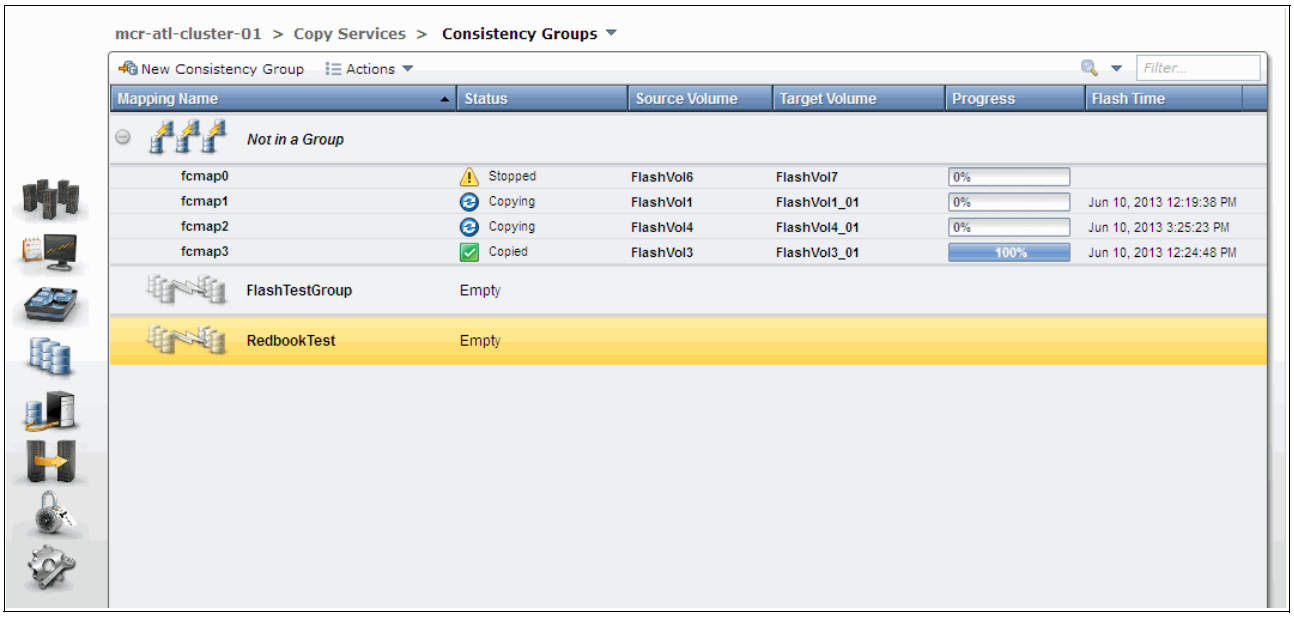

Figure 10-58 New consistency group

You can rename the Consistency Group by selecting it and then right-clicking or by using the Actions drop-down menu. Select **Rename** and enter the new name, as shown in Figure 10-59. Next to the name of the consistency group, the state shows that it is now an empty consistency group with no FlashCopy mapping in it.

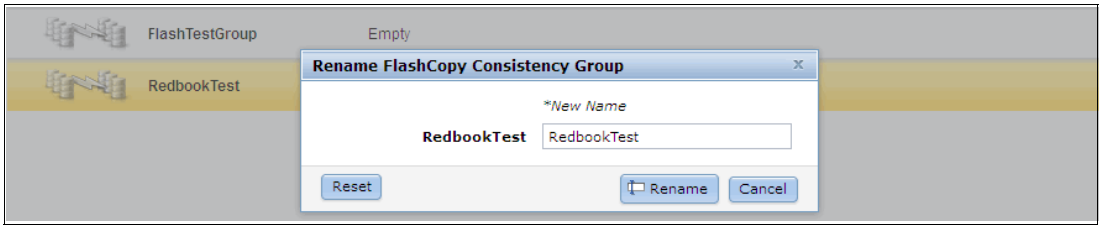

Figure 10-59 Renaming a consistency group

Adding FlashCopy mappings to a consistency group

Click **Not in a Group** to list all the FlashCopy mappings with no Consistency Group. You can add FlashCopy mappings to a Consistency Group by selecting them and clicking the **Move to Consistency Group** option from the Actions drop-down menu, as shown in Figure 10-60 on page 494.

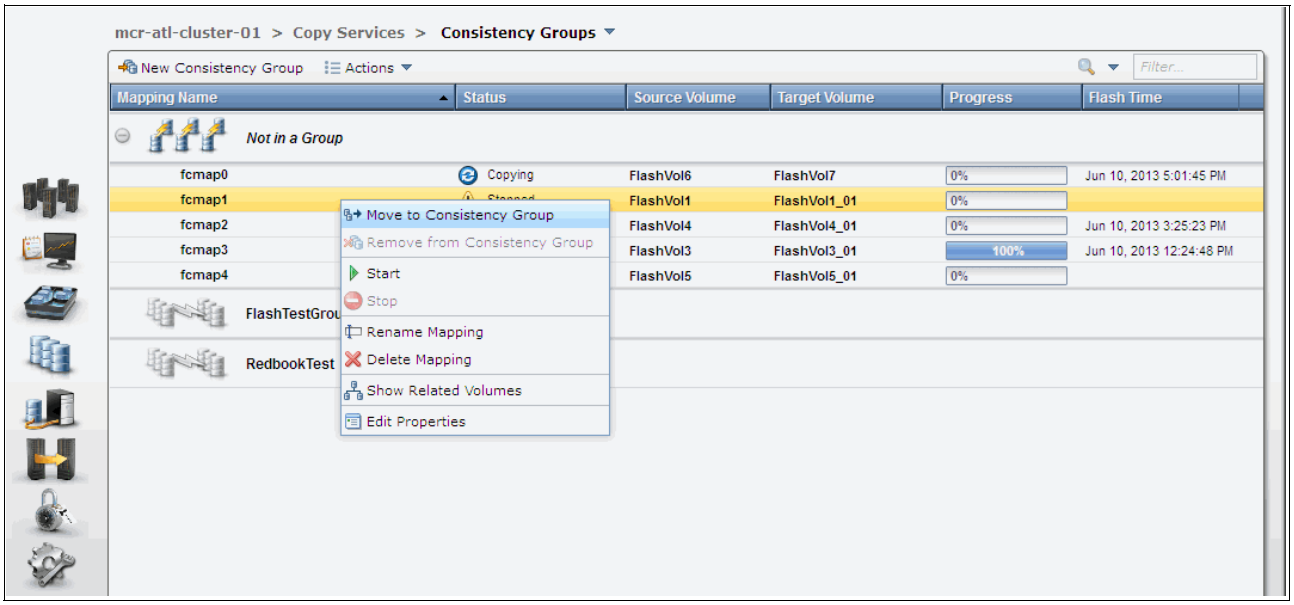

Figure 10-60 Select the FlashCopy mappings to add to a consistency group

Important: You cannot move mappings that are copying. Selecting a snapshot that is already running results in the Move to Consistency Group option being disabled.

Selections of a range are performed by highlighting a mapping, pressing and holding the Shift key, and clicking the last item in the range. Multiple selections can be made by pressing and holding the Ctrl key and clicking each mapping individually. The option is also available by right-clicking individual mappings.

You are prompted to specify which consistency group you want to move the FlashCopy mapping into, as shown in Figure 10-61. Choose from the list in the drop-down menu. Click **Move to Consistency Group** to continue.

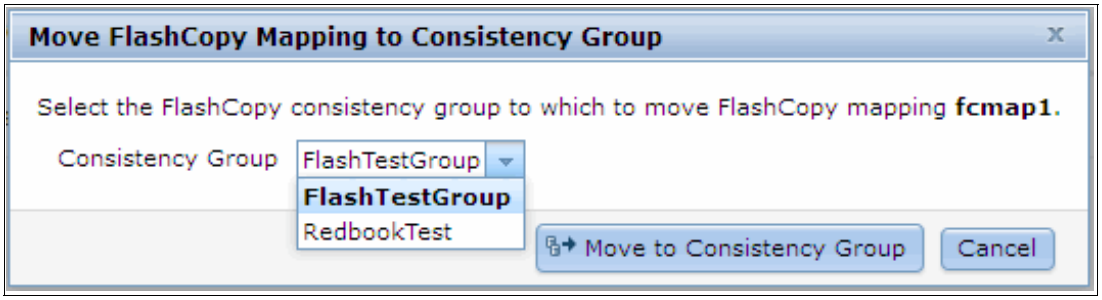

Figure 10-61 Select consistency group

After the action completes, you find that the FlashCopy mappings you selected were removed from the Not In a Group list to the consistency group you chose.

Starting a consistency group

To start a consistency group, highlight the required group and click **Start** from the Actions drop-down menu or right-click the consistency group, as shown in Figure 10-62.

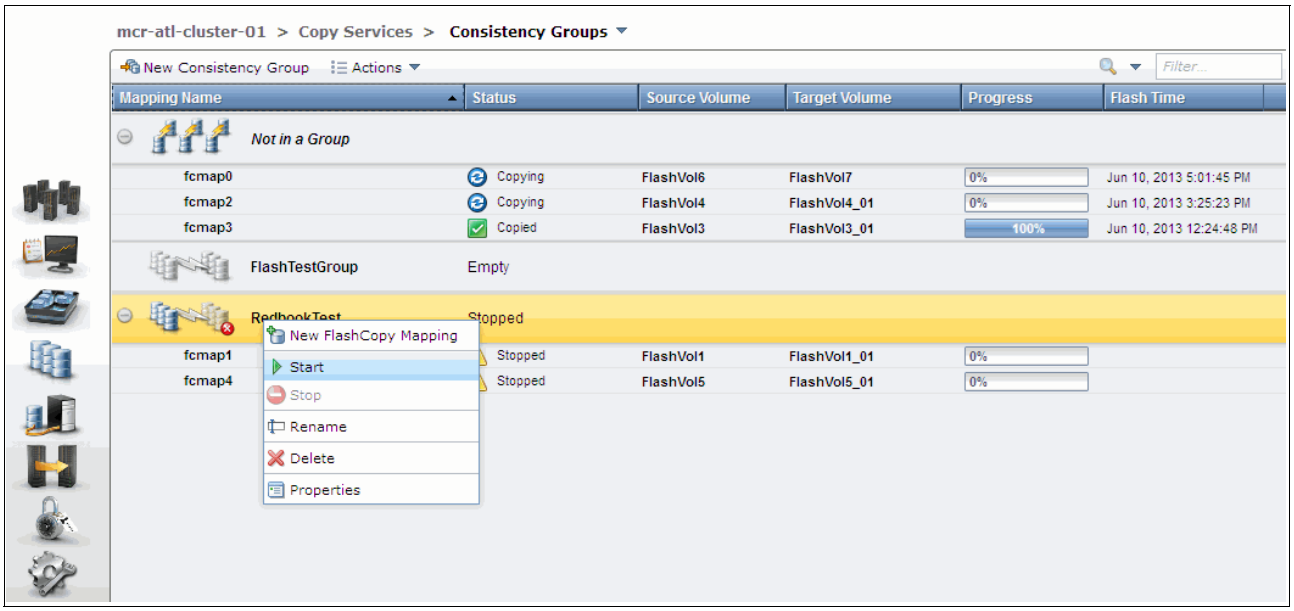

Figure 10-62 Start a consistency group

After you start the consistency group, all the FlashCopy mappings in the consistency group start at the same time. The state of consistency group and all the underlying mappings changes to Copying, as shown in Figure 10-63.

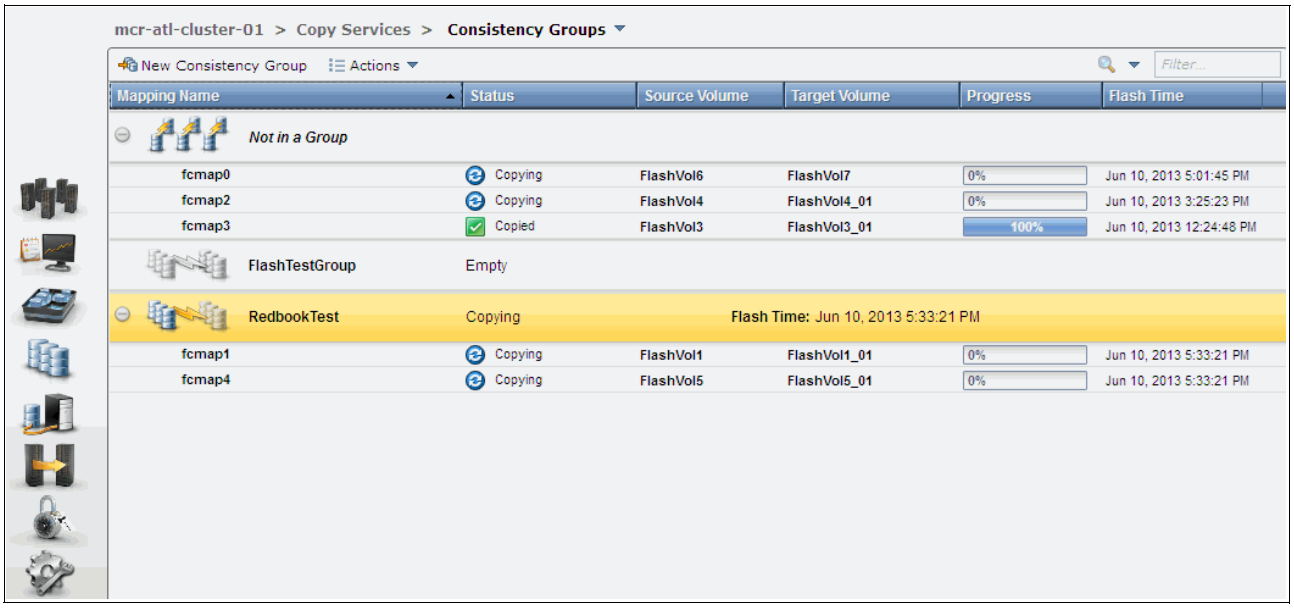

Figure 10-63 Consistency group start completes

Stopping a consistency group

The consistency group can be stopped by selecting **Stop** from the Actions drop-down menu or right-clicking, as shown in Figure 10-64.

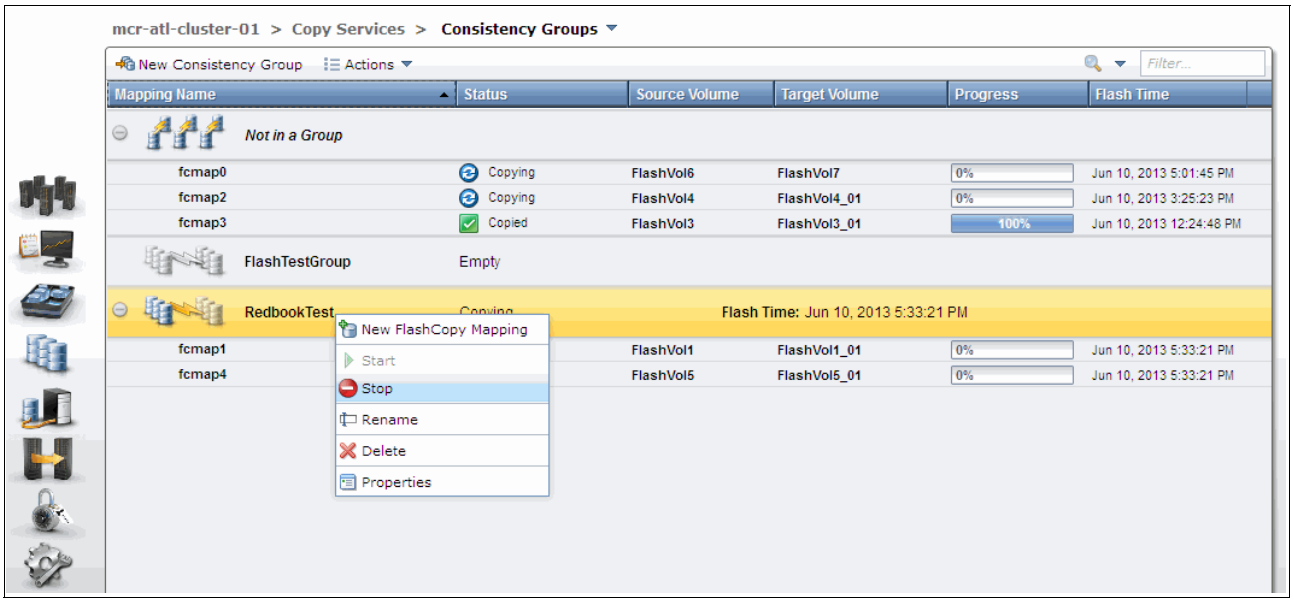

Figure 10-64 Stop a consistency group

After the stop process completes, the FlashCopy mappings in the consistency group are in the Stopped state and a red X icon appears on the function icon of this consistency group to indicate an alert, as shown in Figure 10-65.

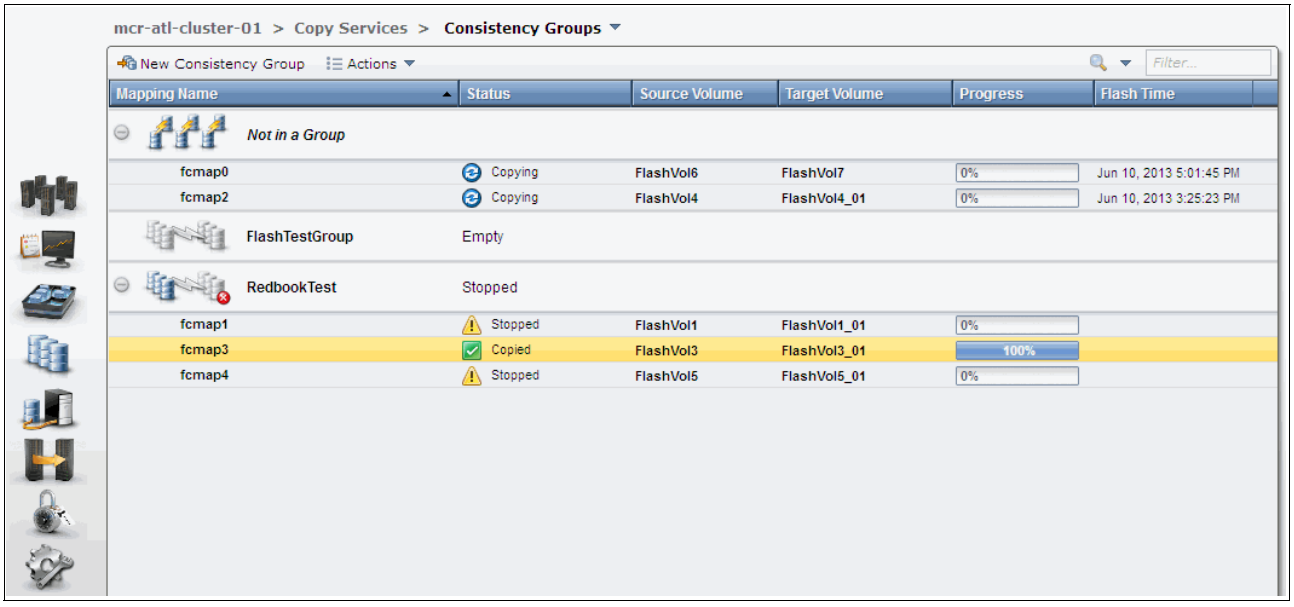

Figure 10-65 Consistency group stop completes

Previously copied relationships that were added to a consistency group that was later stopped before all members of the consistency group completed synchronization remain in the Copied state.

Removing FlashCopy mappings from a consistency group

FlashCopy mappings can be removed from a consistency group by selecting the FlashCopy mappings and clicking **Remove from Consistency Group** from the Actions drop-down menu of the FlashCopy mapping or right-clicking, as shown in Figure 10-66.

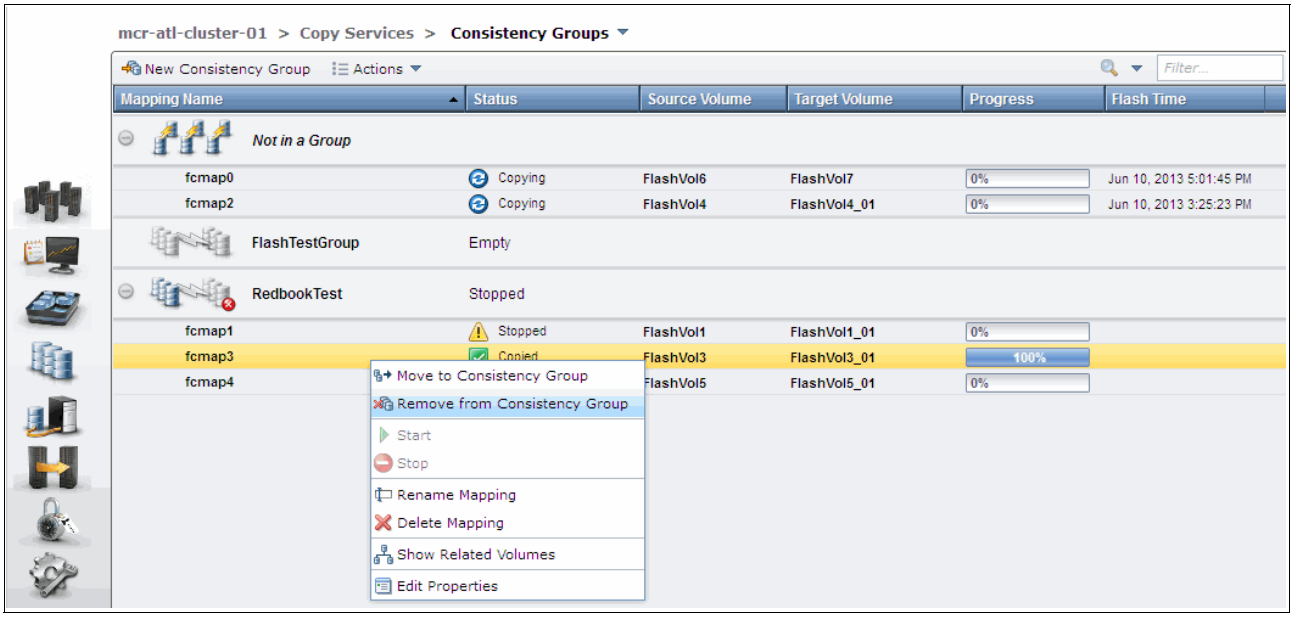

Figure 10-66 Remove from consistency group

The FlashCopy mappings are returned to the Not in a Group list after they are removed from the consistency group.

Deleting a consistency group

A consistency group can be deleted by clicking **Delete** from the Actions drop-down menu or right-clicking the selected group, as shown in Figure 10-67.

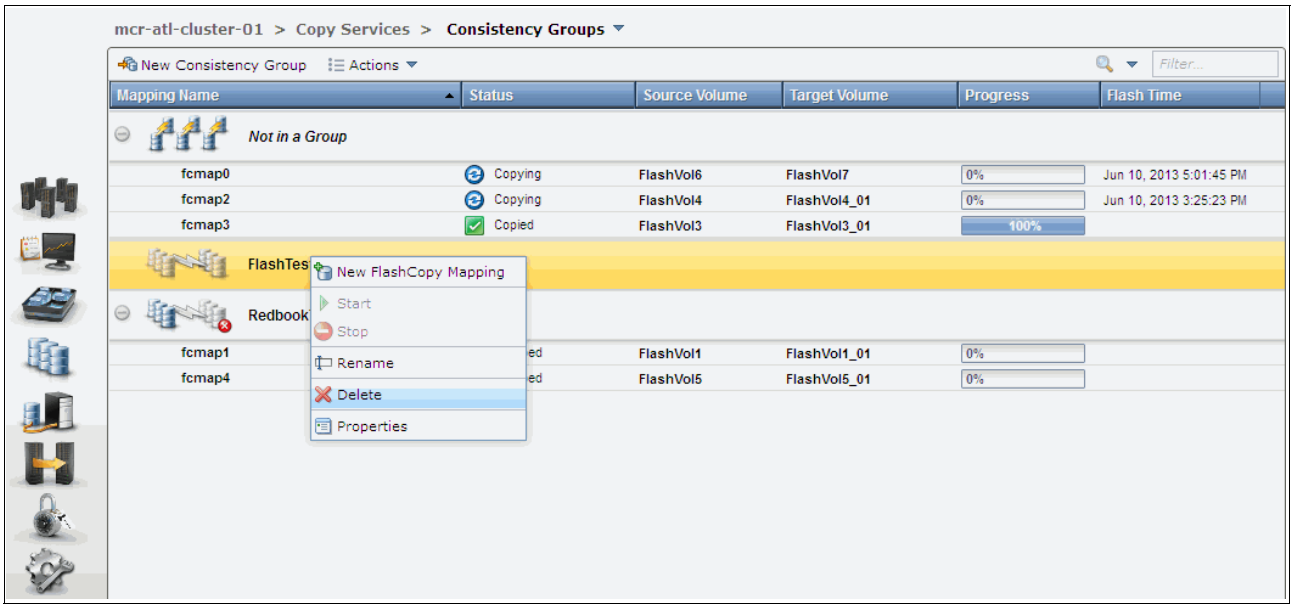

Figure 10-67 Delete a consistency group

Restoring from a FlashCopy Consistency Group

It is possible to manipulate FlashCopy mappings that were captured as part of a consistency group to restore the source volumes of those mappings to the state they were all in at the time the FlashCopy was taken.

To restore a consistency group from a FlashCopy, we must create a reverse mapping of all the individual volumes that are contained within the original consistency group. In our example, we have two FlashCopy mappings (`fcmap1` and `fcmap4`) in a consistency group that is known as FlashTestGroup, as shown in Figure 10-68.

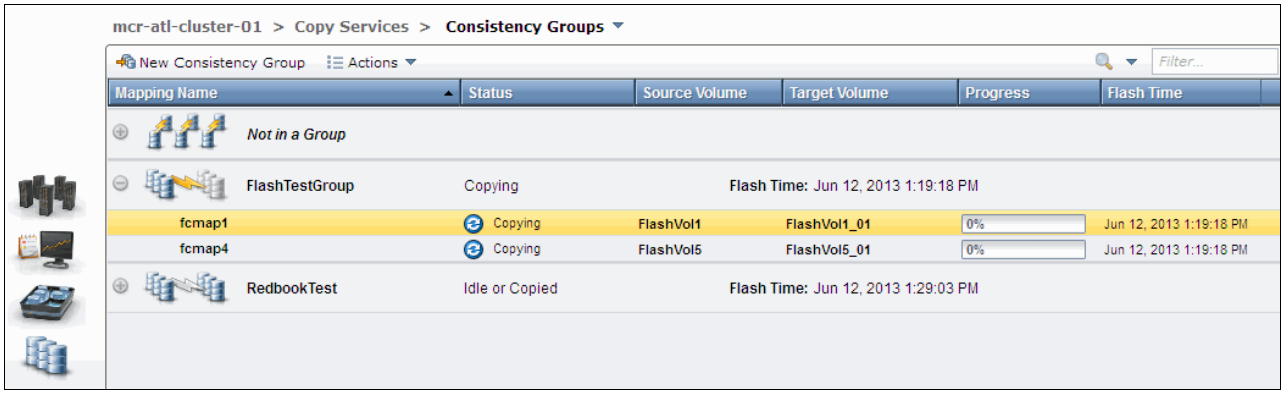

Figure 10-68 Creating FlashCopy reverse mapping

Complete the following steps:

1. Click **New Consistency Group** in the upper left corner (as shown in Figure 10-68) and create a consistency group. In our example, we created a group called `RedBookTest`.
2. Follow the procedure that is described in "Restoring from a FlashCopy" on page 485 to create reverse mappings for each of the mappings that exist in the source consistency group (`FlashTestGroup`). When prompted to add to a consistency group (as shown in Figure 10-49 on page 488), select **Yes** and from the drop-down menu and then select the new "reverse" consistency group that you created in step 2. In our example, this group is `RedBookTest`. The result should be similar to what is shown in Figure 10-69.

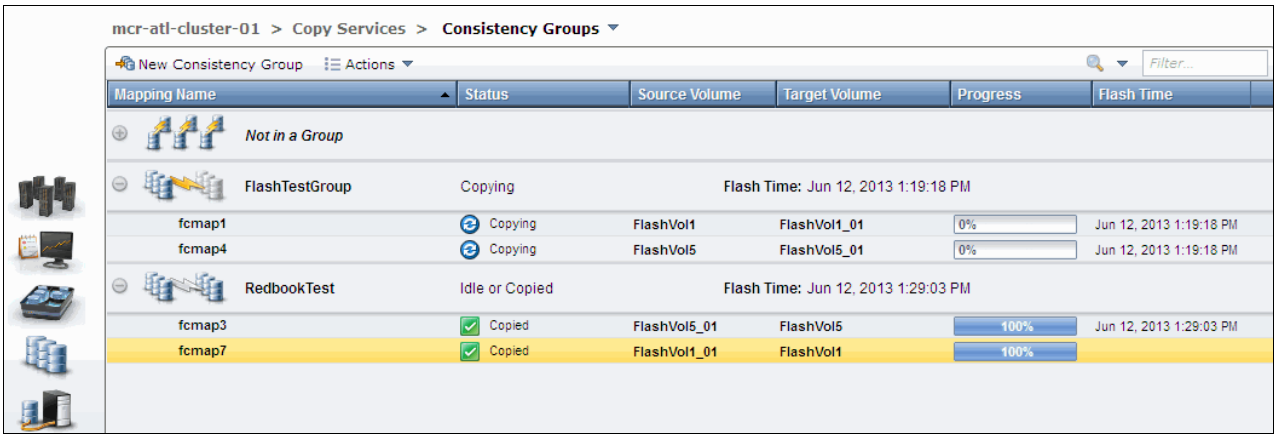

Figure 10-69 Reverse Consistency group populated.

3. To restore the consistency group, highlight the reverse consistency group and click **Start**, as shown in Figure 10-70.

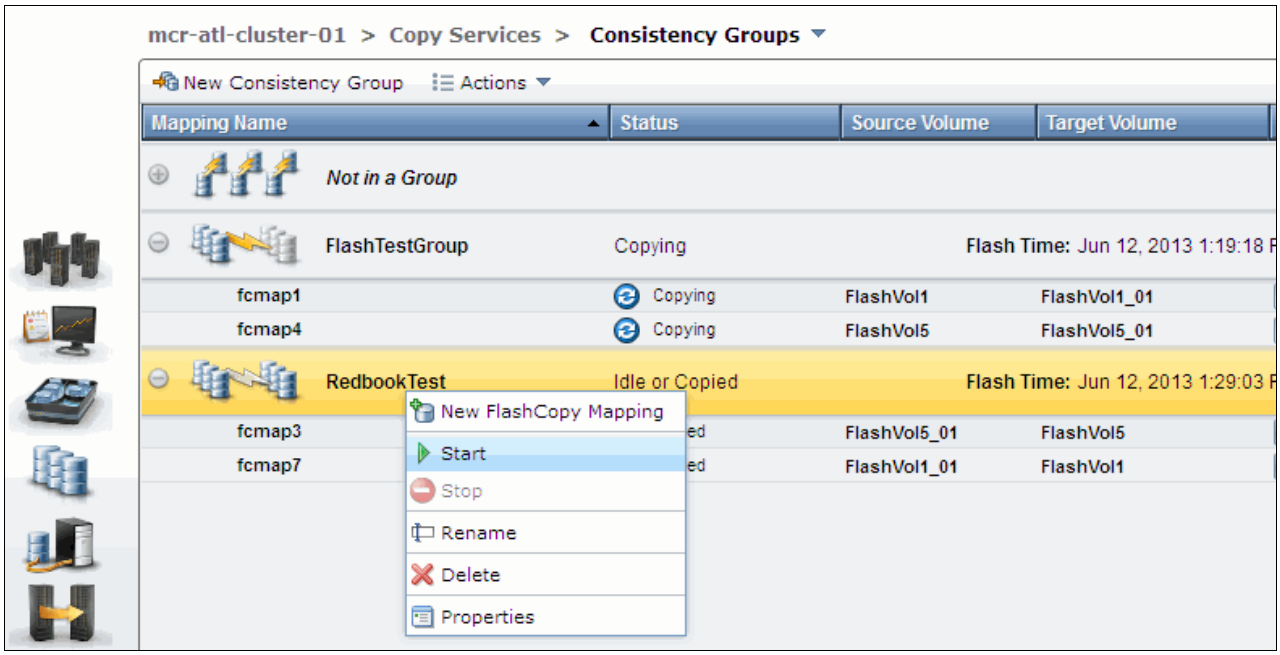

Figure 10-70 Starting Consistency group restore

4. Click **Start** to overwrite `FlashVol1` and `FlashVol5` with the original bitmap data that was saved in the FlashTestGroup FlashCopy consistency group mapping. The command completes, as shown in Figure 10-71.

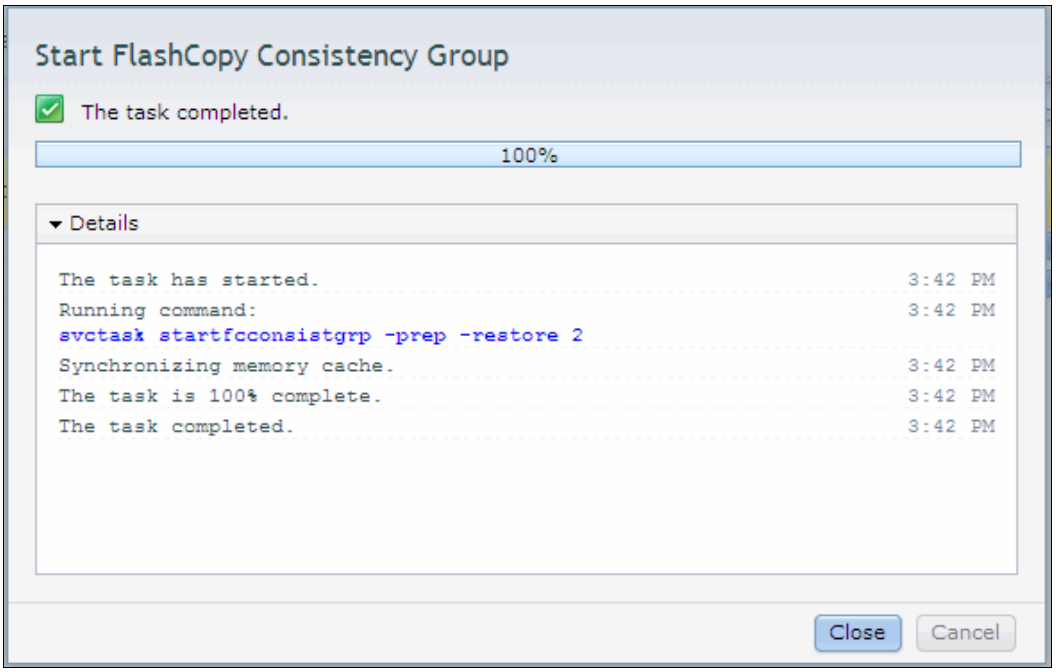

Figure 10-71 Consistency Group restore command

Important: The IBM Storwize V5000 automatically appends the **-restore** option to the command.

5. Click **Close** and the command panel returns to the Consistency Group window. The reverse consistency group now shows as a 100% copied and all volumes in the original FlashTestGroup were restored, as shown in Figure 10-72.

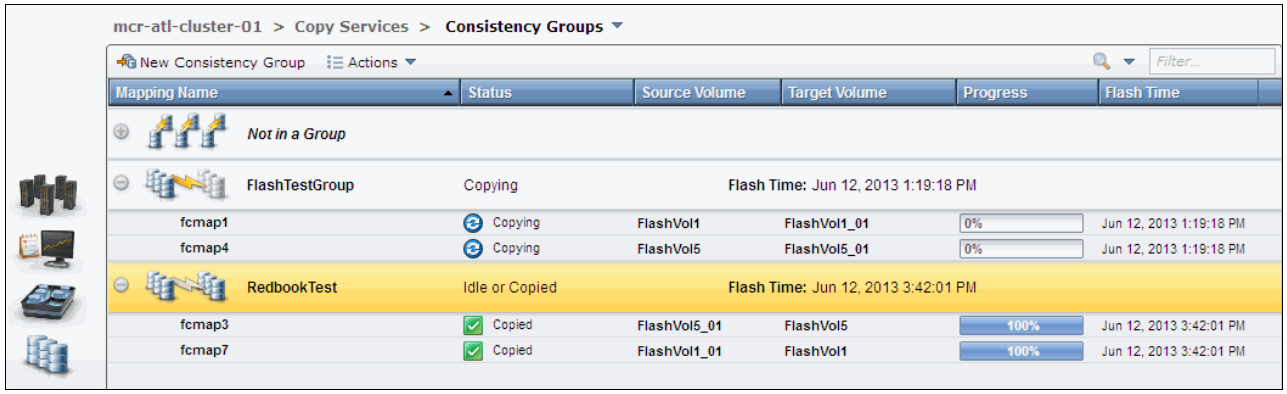

Figure 10-72 Consistency Group restored

10.2 Remote Copy

In this section, we describe how the Remote Copy function works in IBM Storwize V5000. We also provide the implementation steps for Remote Copy configuration and management by using the GUI.

Remote Copy consists of three methods for copying: Metro Mirror, Global Mirror, and Global Mirror with Change Volumes. Metro Mirror is designed for metropolitan distances with a synchronous copy requirement. Global Mirror is designed for longer distances without requiring the hosts to wait for the full round-trip delay of the long-distance link through asynchronous methodology. Global Mirror with Change Volumes is an added piece of functionality for Global Mirror that is designed to attain consistency on lower-quality network links.

Metro Mirror and Global Mirror are IBM branded terms for the functions Synchronous Remote Copy and Asynchronous Remote Copy. Throughout this book, the term “Remote Copy” is used to refer to both functions where the text applies to each term equally.

10.2.1 Remote Copy concepts

Remote Copy concepts are described in this section.

Partnership

When a partnership is created, we connect two separate IBM Storwize V5000 systems or an IBM SAN Volume Controller, Storwize V3700, or Storwize V7000, and an IBM Storwize V5000. After the partnership creation is configured on both systems, further communication between the node canisters in each of the storage systems is established and maintained by the SAN. All inter-cluster communication goes through the Fibre Channel network.

The partnership must be defined on both IBM Storwize V5000 or on the IBM Storwize V5000 and the other IBM SAN Volume Controller, Storwize V3700, or Storwize V7000 storage system to make the partnership fully functional.

Interconnection: Interconnects between IBM Storwize products were introduced in Version 6.3.0. Because IBM Storwize V5000 supports only version 7.10 or higher, there is no problem with support for this functionality. However, any other Storwize product must be at a minimum level of 6.3.0 to connect to the IBM Storwize V5000 and the IBM Storwize V5000 must set the replication layer by using the **`svctask chsystem -layer replication`** limitations that are described next.

Introduction to layers

IBM Storwize V5000 implements the concept of *layers*. Layers determine how the IBM Storwize portfolio interacts with the IBM SAN Volume Controller. Currently, there are two layers: replication and storage.

The replication layer is used when you want to use the IBM Storwize V5000 with one or more IBM SAN Volume Controllers as a Remote Copy partner. The storage layer is the default mode of operation for the IBM Storwize V5000, and is used when you want to use the IBM Storwize V5000 to present storage to an IBM SAN Volume Controller.

The layer for the IBM Storwize V5000 can be switched by running **`svctask chsystem -layer replication`**. Generally, switch the layer while your IBM Storwize V5000 system is not in production. This situation prevents potential disruptions because layer changes are not I/O-tolerant.

Figure 10-73 shows the effect of layers on IBM SAN Volume Controller and IBM Storwize V5000 partnerships.

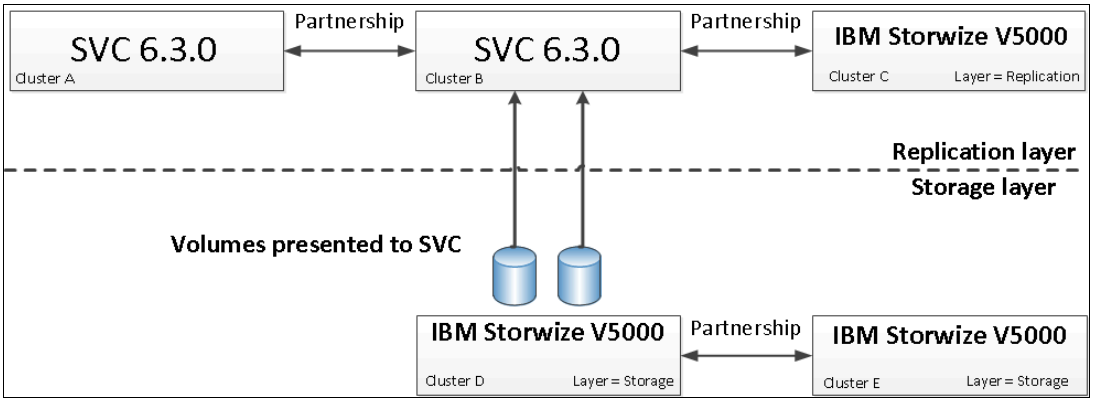

Figure 10-73 IBM Storwize V5000 virtualization layers

The replication layer allows an IBM Storwize V5000 system to be a Remote Copy partner with an IBM SAN Volume Controller. The storage layer allows an IBM Storwize V5000 system to function as back-end storage for an IBM SAN Volume Controller. An IBM Storwize V5000 system cannot be in both layers at the same time.

Limitations on the SAN Volume Controller and Storwize V5000 partnership

IBM SAN Volume Controller and IBM Storwize V5000 systems can be partners in a Remote Copy partnership. However, the following limitations apply:

- The layer for the V5000 must be set to replication. The default is storage.
- If any other SAN Volume Controller or IBM Storwize V5000 ports are visible on the SAN (aside from the ones on the cluster where you are making the changes), you cannot change the layer.
- If any host object is defined to an IBM SAN Volume Controller or IBM Storwize V5000 system, you cannot change the layer.
- If any MDisks from an IBM Storwize V5000 other than the one you are making the layer change on are visible, you cannot change the layer.
- If any cluster partnership is defined, you cannot change the layer.

Partnership topologies

A partnership between up to four IBM Storwize V5000 systems is allowed.

The following typical partnership topologies between multiple IBM Storwize V5000s are available:

- Daisy-chain topology, as shown in Figure 10-74.

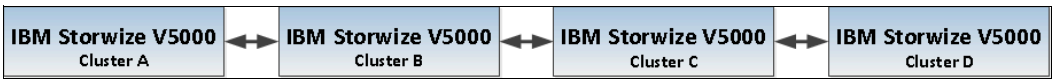

Figure 10-74 Daisy chain partnership topology for IBM Storwize V5000

- Triangle topology, as shown in Figure 10-75.

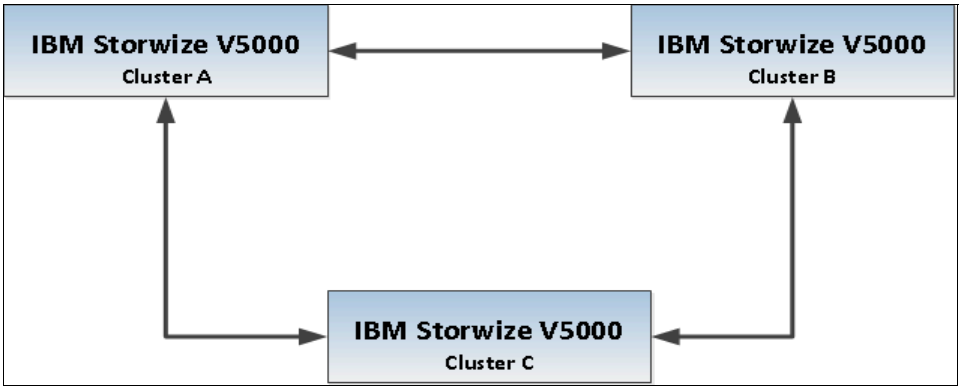

Figure 10-75 Triangle partnership topology for IBM Storwize V5000

- Star topology, as shown in Figure 10-76.

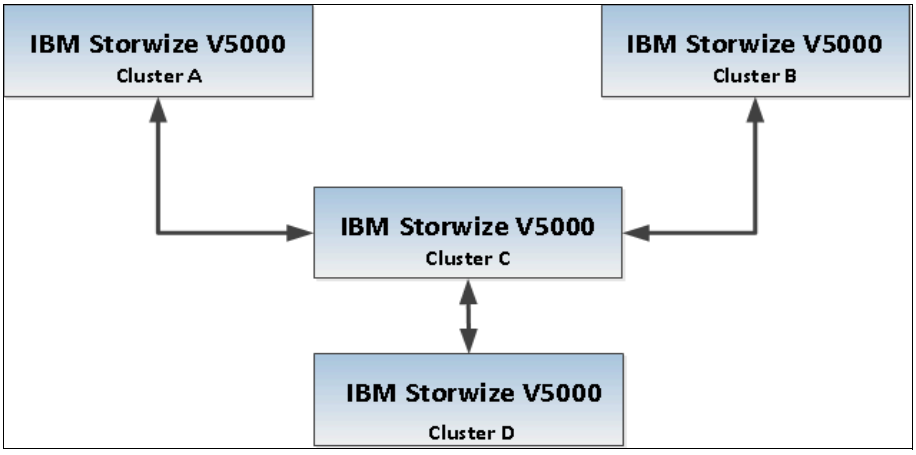

Figure 10-76 Star topology for IBM Storwize V5000

- Full-meshed topology, as shown in Figure 10-77.

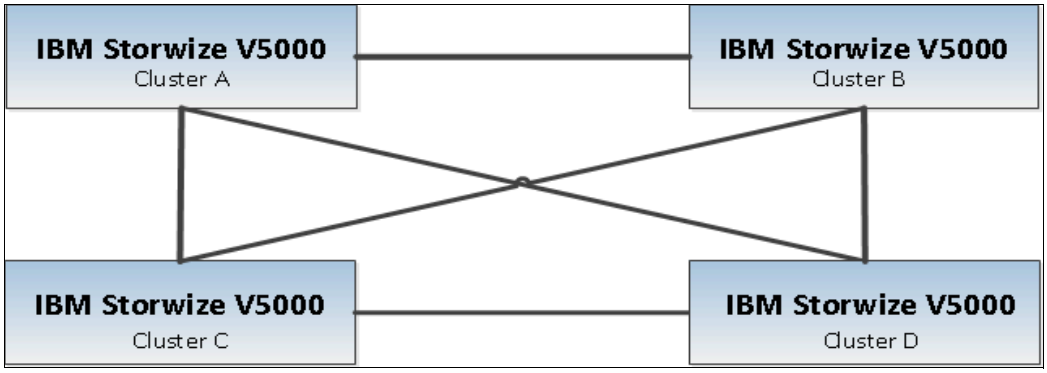

Figure 10-77 Full meshed IBM Storwize V5000

Partnerships: These partnerships are valid for configurations with SAN Volume Controllers and IBM Storwize V5000 systems if the IBM Storwize V5000 systems are using the replication layer. They are also valid for Storwize V3700 and V7000 products.

Partnership states

A partnership has the following states:

- Partially Configured

  Indicates that only one cluster partner is defined from a local or remote cluster to the displayed cluster and is started. For the displayed cluster to be configured fully and to complete the partnership, you must define the cluster partnership from the cluster that is displayed to the corresponding local or remote cluster.

- Fully Configured

  Indicates that the partnership is defined on the local and remote clusters and is started.

- Remote Not Present

  Indicates that the remote cluster is not present for the partnership.

- Partially Configured (Local Stopped)

  Indicates that the local cluster is only defined to the remote cluster and the local cluster is stopped.

- Fully Configured (Local Stopped)

  Indicates that a partnership is defined on the local and remote clusters and the remote cluster is present, but the local cluster is stopped.

- Fully Configured (Remote Stopped)

  Indicates that a partnership is defined on the local and remote clusters and the remote cluster is present, but the remote cluster is stopped.

- Fully Configured (Local Excluded)

  Indicates that a partnership is defined between a local and remote cluster; however, the local cluster was excluded. This state can occur when the fabric link between the two clusters was compromised by too many fabric errors or slow response times of the cluster partnership.

- Fully Configured (Remote Excluded)

  Indicates that a partnership is defined between a local and remote cluster; however, the remote cluster was excluded. This state can occur when the fabric link between the two clusters was compromised by too many fabric errors or slow response times of the cluster partnership.

- Fully Configured (Remote Exceeded)

  Indicates that a partnership is defined between a local and remote cluster and the remote is available; however, the remote cluster exceeds the number of allowed clusters within a cluster network. The maximum of four clusters can be defined in a network. If the number of clusters exceeds that limit, the IBM Storwize V5000 system determines the inactive cluster or clusters by sorting all the clusters by their unique identifier in numerical order. The inactive cluster partner that is not in the top four of the cluster-unique identifiers shows Fully Configured (Remote Exceeded).

Remote Copy relationships

A Remote Copy relationship is a relationship between two individual volumes of the same size. These volumes are called a *master (source) volume* and an *auxiliary (target) volume*.

Typically, the master volume contains the production copy of the data and is the volume that the application normally accesses. The auxiliary volume often contains a backup copy of the data and is used for disaster recovery.

The master and auxiliary volumes are defined when the relationship is created, and these attributes never change. However, either volume can operate in the primary or secondary role as necessary. The primary volume contains a valid copy of the application data and receives updates from the host application, which is analogous to a source volume. The secondary volume receives a copy of any updates to the primary volume because these updates are all transmitted across the mirror link. Therefore, the secondary volume is analogous to a continuously updated target volume. When a relationship is created, the master volume is assigned the role of primary volume and the auxiliary volume is assigned the role of secondary volume. The initial copying direction is from master to auxiliary. When the relationship is in a consistent state, you can reverse the copy direction.

The two volumes in a relationship must be the same size. The Remote Copy relationship can be established on the volumes within one IBM Storwize V5000 storage system, which is called an *intra-cluster relationship*. The relationship can also be established in different IBM Storwize V5000 storage systems or between an IBM Storwize V5000 storage system and an IBM SAN Volume Controller, IBM Storwize V3700, or IBM Storwize V7000, which are called *inter-cluster relationships*.

Important: The use of Remote Copy target volumes as Remote Copy source volumes is not allowed. A FlashCopy target volume can be used as Remote Copy source volume and also as a Remote Copy target volume.

Metro Mirror

Metro Mirror is a type of Remote Copy that creates a synchronous copy of data from a master volume to an auxiliary volume. With synchronous copies, host applications write to the master volume but do not receive confirmation that the write operation completed until the data is written to the auxiliary volume. This action ensures that both volumes have identical data when the copy completes. After the initial copy completes, the Metro Mirror function always maintains a fully synchronized copy of the source data at the target site.

Figure 10-78 shows how a write to the master volume is mirrored to the cache of the auxiliary volume before an acknowledgement of the write is sent back to the host that issued the write. This process ensures that the auxiliary is synchronized in real time if it is needed in a failover situation.

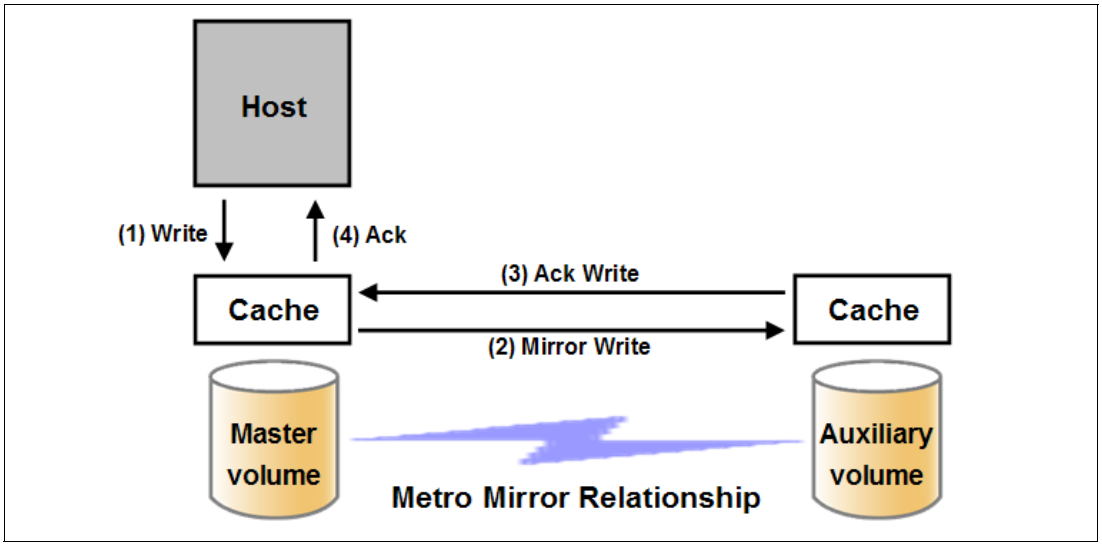

Figure 10-78 Write on volume in a Metro Mirror relationship

The Metro Mirror function supports copy operations between volumes that are separated by distances up to 300 km. For disaster recovery purposes, Metro Mirror provides the simplest way to maintain an identical copy on the primary and secondary volumes. However, as with all synchronous copies over remote distances, there can be a performance impact to host applications. This performance impact is related to the distance between primary and secondary volumes and, depending on application requirements, its use might be limited based on the distance between sites.

Global Mirror

Global Mirror provides an asynchronous copy, which means that the secondary volume is not an exact match of the primary volume at every point. The Global Mirror function provides the same function as Metro Mirror Remote Copy without requiring the hosts to wait for the full round-trip delay of the long-distance link; however, some delay can be seen on the hosts in congested or overloaded environments. Make sure that you closely monitor and understand your workload.

In a synchronous Remote Copy (which Global Mirror provides), write operations are completed on the primary site and the write acknowledgement is sent to the host before it is received at the secondary site. An update of this write operation is sent to the secondary site at a later stage, which provides the capability to perform Remote Copy over distances that exceed the limitations of synchronous Remote Copy.

The distance of Global Mirror replication is limited primarily by the latency of the WAN link that is provided. Global Mirror has a requirement of 80 ms round-trip-time for data that is sent to the remote location. The propagation delay is roughly 8.2 µs per mile or 5 µs per kilometer for Fibre Channel connections. Each device in the path adds more delay of about 25 µs. Devices that use software (such as, some compression devices) adds much more time. The time that is added by software-assisted devices is highly variable and should be measured directly. Be sure to include these times when you are planning your Global Mirror design.

You should also measure application performance that is based on the expected delays before Global Mirror is fully implemented. The IBM Storwize V5000 storage system provides you with an advanced feature of Global Mirror that permits you to test performance implications before Global Mirror is deployed and a long-distance link is obtained. This advanced feature is enabled by modifying the IBM Storwize V5000 storage system parameters **`gmintradelaysimulation`** and **`gminterdelaysimulation`**. These parameters can be used to simulate the write delay to the secondary volume. The delay simulation can be enabled separately for each intra-cluster or inter-cluster Global Mirror. You can use this feature to test an application before the full deployment of the Global Mirror feature. For more information about how to enable the CLI feature, see Appendix A, “Command-line interface setup and SAN Boot” on page 609.

Figure 10-79 on page 507 shows that a write operation to the master volume is acknowledged back to the host that is issuing the write before the write operation is mirrored to the cache for the auxiliary volume.

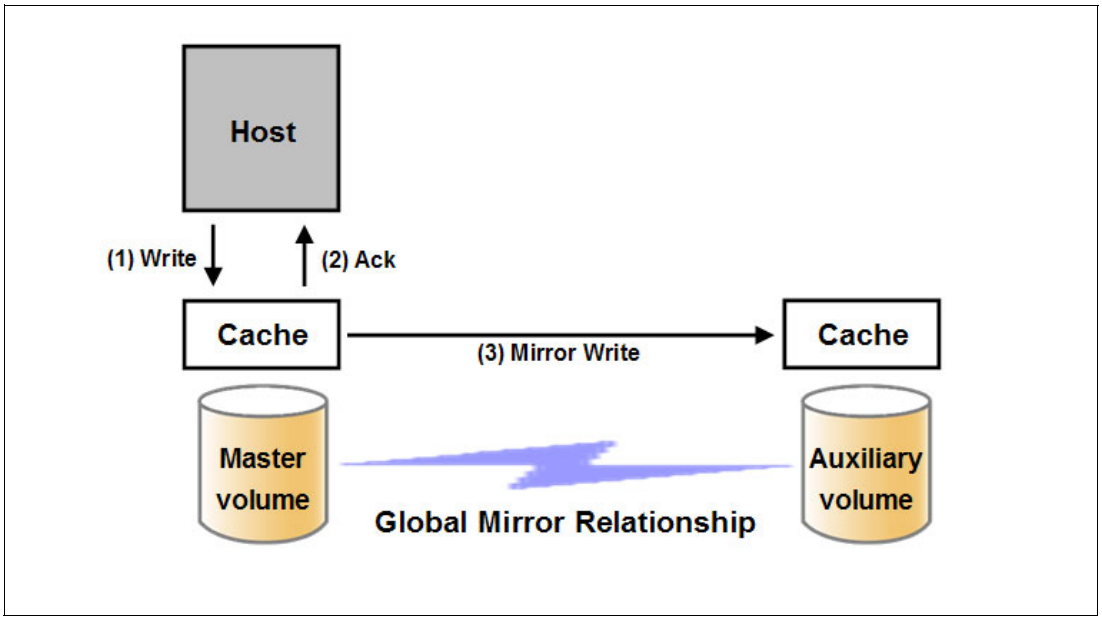

Figure 10-79 Global Mirror write sequence

The Global Mirror algorithms always maintain a consistent image on the auxiliary volume. They achieve this consistent image by identifying sets of I/Os that are active concurrently at the master, assigning an order to those sets, and applying those sets of I/Os in the assigned order at the secondary.

In a failover scenario where the secondary site must become the master source of data (depending on the workload pattern and the bandwidth and distance between local and remote site), certain updates might be missing at the secondary site. Therefore, any applications that use this data must have an external mechanism for recovering the missing updates and reapplying them; for example, a transaction log replay.

10.2.2 Global Mirror with Change Volumes

Global Mirror within the IBM Storwize V5000 is designed to achieve a recovery point objective (RPO) as low as possible so that data is as up-to-date as possible. This capability places some strict requirements on your infrastructure and in certain situations (with low network link quality or congested or overloaded hosts), you might be affected by multiple 1920 (congestion) errors.

Congestion errors happen in the following primary situations:

- Congestion at the source site through the host or network.
- Congestion in the network link or network path.
- Congestion at the target site through the host or network.

Global Mirror includes functionality that is designed to address the following conditions that negatively affect some Global Mirror implementations:

- Estimation of bandwidth requirements tends to be complex.
- It is often difficult to ensure that the latency and bandwidth requirements can be met.
- Congested hosts on the source or target site can cause disruption.
- Congested network links can cause disruption with only intermittent peaks.

To address these issues, *Change Volumes* were added as an option for Global Mirror relationships. Change Volumes use the FlashCopy functionality but cannot be manipulated as FlashCopy volumes because they are special-purpose only. Change Volumes replicate point-in-time images on a cycling period (the default is 300 seconds). This situation means that your change rate must include only the condition of the data at the point-in-time the image was taken instead of all the updates during the period. This situation can provide significant reductions in replication volume.

Figure 10-80 shows a basic Global Mirror relationship without Change Volumes.

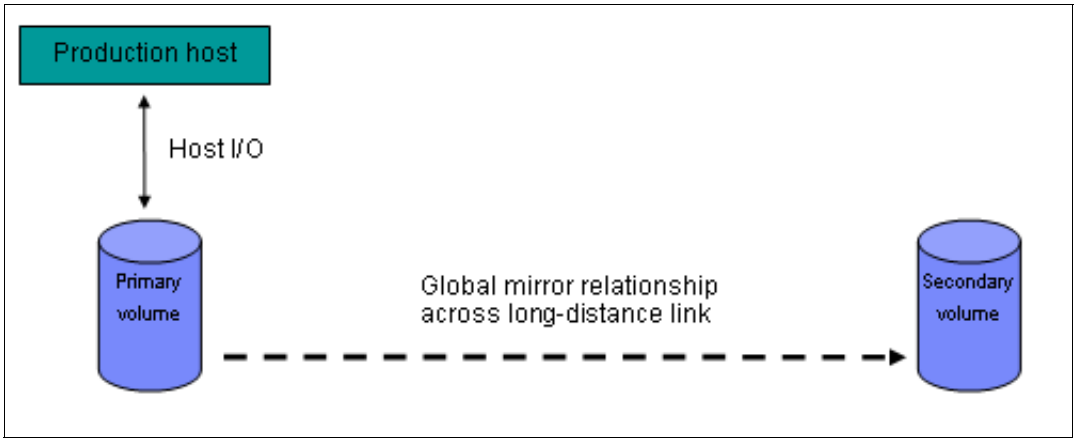

Figure 10-80 Global Mirror without Change Volumes

Figure 10-81 shows a relationship with the Change Volumes.

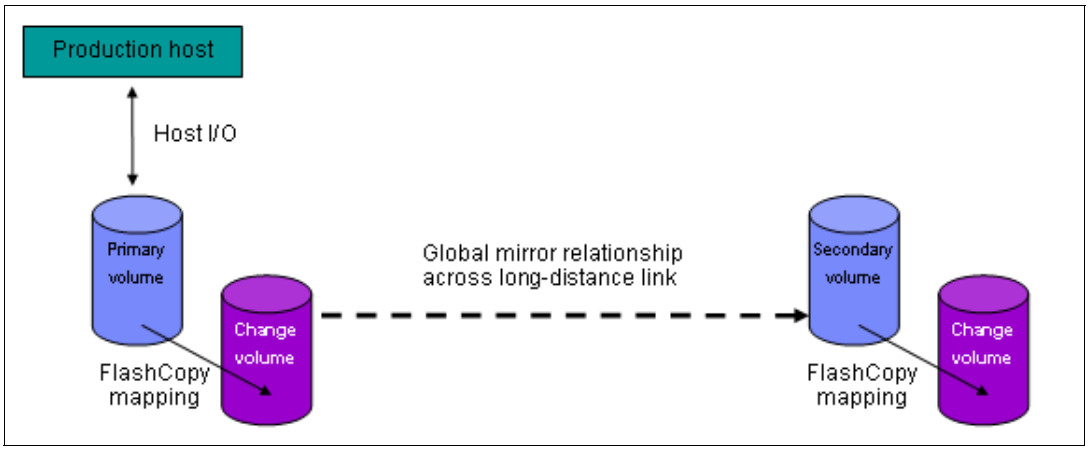

Figure 10-81 Global Mirror with Change Volumes

With Change Volumes, a FlashCopy mapping exists between the primary volume and the primary Change Volume. The mapping is updated during a cycling period (every 60 seconds to one day). The primary Change Volume is then replicated to the secondary Global Mirror volume at the target site, which is then captured in another change volume on the target site. This situation provides a consistent image at the target site and protects your data from being inconsistent during resynchronization.

Figure 10-82 shows a number of I/Os on the source volume, the same number on the target volume, and in the same order. Assuming that this set is the same set of data that is updated over and over, these updates are wasted network traffic and the I/O can be completed much more efficiently, as shown in Figure 10-83.

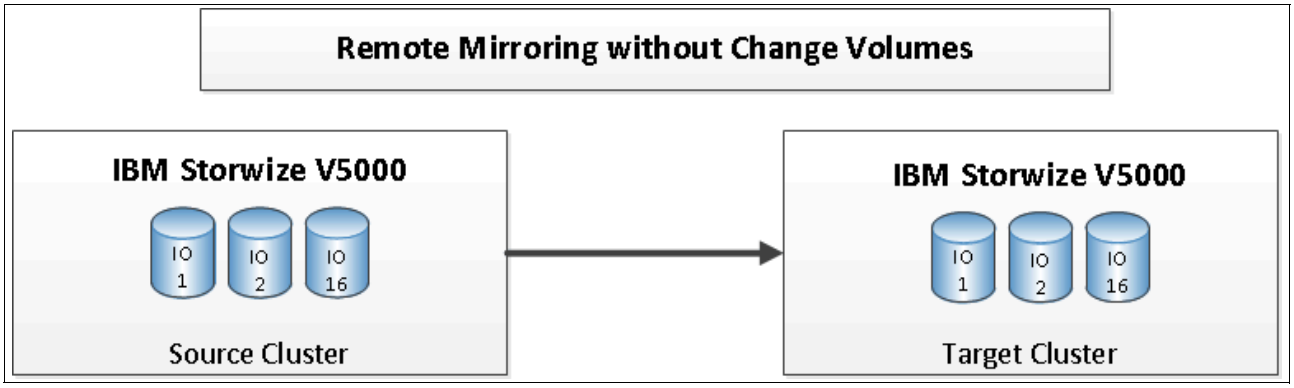

Figure 10-82 Global Mirror I/O replication without Change Volumes

In Figure 10-83, the same data is being updated repeatedly, so Change Volumes demonstrate significant I/O transmission savings because you must send only I/O number 16, which was the last I/O before the cycling period.

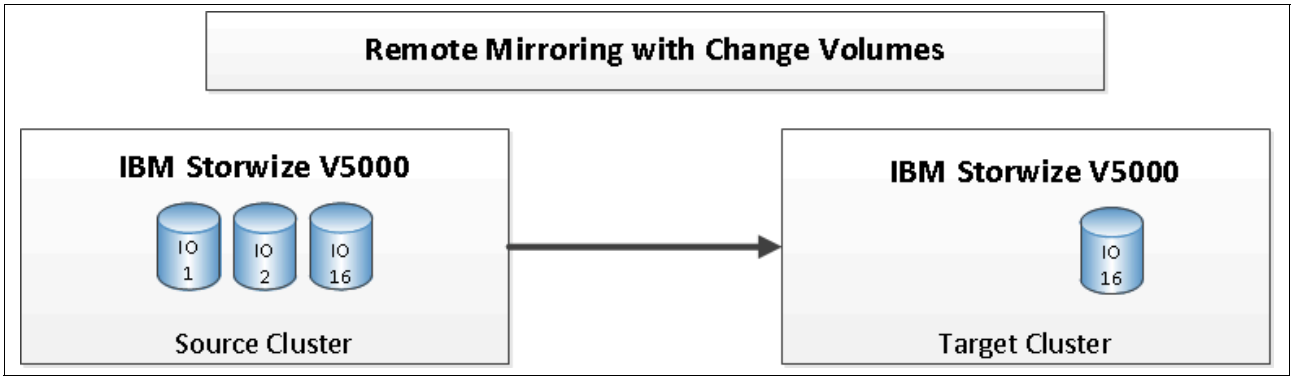

Figure 10-83 Global Mirror I/O replication with Change Volumes

The cycling period can be adjusted by running **`chrcrelationship -cycleperiodseconds <60-86400>`**. If a copy does not complete in the cycle period, the next cycle does not start until the prior cycle completes. It is for this reason that the use of Change Volumes gives you the following possibilities for RPO:

- If your replication completes in the cycling period, your RPO is twice the cycling period.
- If your replication does not complete within the cycling period, your RPO is twice the completion time. The next cycling period starts immediately after the prior period is finished.

Careful consideration should be put into balancing your business requirements with the performance of Global Mirror with Change Volumes. Global Mirror with Change Volumes increases the inter-cluster traffic for more frequent cycling periods, so going as short as possible is not always the answer. In most cases, the default should meet your requirements and perform reasonably well.

Important: When Global Mirror volumes with Change Volumes are used, make sure that you remember to select the Change Volume on the auxiliary (target) site. Failure to do so leaves you exposed during a resynchronization operation.

The GUI automatically creates Change Volumes for you. However, it is a limitation of this initial release that they are fully provisioned volumes. To save space, you should create thin-provisioned volumes in advance and use the existing volume option to select your change volumes.

Remote Copy consistency groups

A consistency group is a logical entity that groups copy relationships. By grouping the relationships, you can ensure that these relationships are managed in unison and the data within the group is in a consistent state. For more information about the necessity of consistency groups, see 10.1.6, “Managing a FlashCopy consistency group” on page 490.

Remote Copy commands can be issued to a Remote Copy consistency group, and, therefore, simultaneously for all Metro Mirror relationships that are defined within that consistency group, or to a single Metro Mirror relationship that is not part of a Metro Mirror consistency group.

Figure 10-84 shows the concept of Remote Copy consistency groups. Because the RC_Relationships 1 and 2 are part of the consistency group, they can be handled as one entity, while the stand-alone RC_Relationship 3 is handled separately.

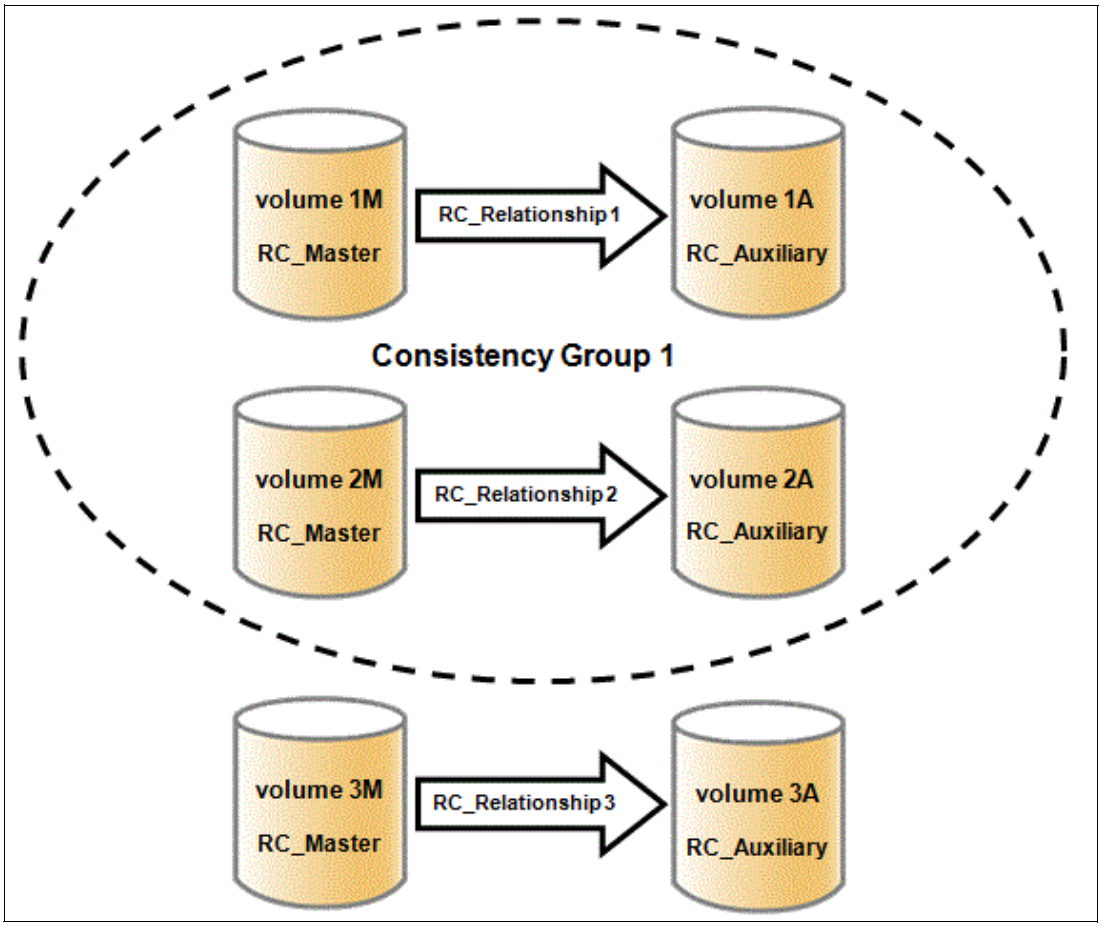

Figure 10-84 Remote Copy consistency group

Remote Copy relationships can belong only to one consistency group, but they do not have to belong to a consistency group. Relationships that are not part of a consistency group are called *stand-alone relationships*. A consistency group can contain zero or more relationships. All relationships in a consistency group must have matching primary and secondary clusters, which are sometimes referred to as *master clusters* and *auxiliary clusters*. All relationships in a consistency group must also have the same copy direction and state.

Metro Mirror and Global Mirror relationships cannot belong to the same consistency group. A copy type is automatically assigned to a consistency group when the first relationship is added to the consistency group. After the consistency group is assigned a copy type, only relationships of that copy type can be added to the consistency group.

Remote Copy and consistency group states

Stand-alone Remote Copy relationships and consistency groups share a common configuration and state model. All of the relationships in a non-empty consistency group have the same state as the consistency group.

The following states apply to the relationships and the consistency groups, except for the Empty state, which is only for consistency groups:

- InconsistentStopped

  The primary volumes are accessible for read and write I/O operations, but the secondary volumes are not accessible for either one. A copy process must be started to make the secondary volumes consistent.

- InconsistentCopying

  The primary volumes are accessible for read and write I/O operations, but the secondary volumes are not accessible for either one. This state indicates that a copy process is ongoing from the primary to the secondary volume.

- ConsistentStopped

  The secondary volumes contain a consistent image, but it might be outdated about the primary volumes. This state can occur when a relationship was in the ConsistentSynchronized state and experiences an error that forces a freeze of the consistency group or the Remote Copy relationship.

- ConsistentSynchronized

  The primary volumes are accessible for read and write I/O operations. The secondary volumes are accessible for read-only I/O operations.

- Idling

  The primary volumes and the secondary volumes are operating in the primary role. Therefore, the volumes are accessible for write I/O operations.

- IdlingDisconnected

  The volumes in this half of the consistency group are all operating in the primary role and can accept read or write I/O operations.

- InconsistentDisconnected

  The volumes in this half of the consistency group are all operating in the secondary role and cannot accept read or write I/O operations.

- ConsistentDisconnected

  The volumes in this half of the consistency group are all operating in the secondary role and can accept read I/O operations but not write I/O operations.

- Empty

  The consistency group does not contain any relationships.

10.2.3 Remote Copy planning

Before you use Remote Copy, you must plan for its usage.

General guidelines for Remote Copy

General guidelines for Remote Copy include the following considerations:

- Partnerships between up to four IBM Storwize V5000 storage systems, IBM SAN Volume Controller systems, IBM Storwize V7000, or IBM Storwize V3700 is allowed. The partnership must be defined on any partnered IBM Storwize storage systems or IBM SAN Volume Controller systems to make it fully functional.
- The two volumes in a relationship must be the same size.
- The Remote Copy relationship can be established on the volumes within one IBM Storwize V5000 storage system or in different IBM Storwize V5000 storage systems. When the two volumes are in the same cluster, they must be in the same I/O group.
- You cannot use Remote Copy target volumes as Remote Copy source volumes. However, a FlashCopy target volume can be used as Remote Copy source volume. Other restrictions are outlined in Table 10-5 on page 514.
- The Metro Mirror function supports copy operations between volumes that are separated by distances up to 300 km.
- One Remote Copy relationship can belong only to one consistency group.
- All relationships in a consistency group must have matching primary and secondary clusters, (master clusters and auxiliary clusters). All relationships in a consistency group must also have the same copy direction and state.
- Metro Mirror and Global Mirror relationships cannot belong to the same consistency group.
- To manage multiple Remote Copy relationships as one entity, relationships can be made part of a Remote Copy consistency group, which ensures data consistency across multiple Remote Copy relationships and provides ease of management.
- An IBM Storwize V5000 storage system implements flexible resynchronization support, which enables it to resynchronize volume pairs that experienced write I/Os to both disks and to resynchronize only those regions that are known to changed.
- Global Mirror with Change Volumes should have Change Volumes that are defined for the master and auxiliary volumes.

Remote Copy configuration limits

Table 10-4 lists the Remote Copy configuration limits.

Table 10-4 Remote Copy configuration limits

Parameter	Value
Number of Remote Copy consistency groups per cluster	256
Number of Remote Copy relationships per consistency group	8,192
Number of Remote Copy relationships per I/O Group	2,048
Total Remote Copy volume capacity per I/O Group	1024 TB (This limit is the total capacity for all master and auxiliary volumes in the I/O group.)

SAN planning for Remote Copy

In this section, we describe some guidelines that can be used for planning for a SAN for Remote Copy.

Zoning recommendation

Node canister ports on each IBM Storwize V5000 must communicate with each other so that the partnership can be created. These ports must be visible to each other on your SAN. Proper switch zoning is critical to facilitating inter-cluster communication.

The following SAN zoning recommendation should be considered:

- For each node canister, exactly two Fibre Channel ports should be zoned to exactly two Fibre Channel ports from each node canister in the partner cluster.
- If dual-redundant inter-switch links (ISLs) are available, the two ports from each node should be split evenly between the two ISLs; that is, exactly one port from each node canister should be zoned across each ISL. For more information, see this website:

  `http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003634&myns=s033&mynp=familyind5329743&mync=E`
- All local zoning rules should be followed. A properly configured SAN fabric is key to not only local SAN performance, but Remote Copy. For more information about these rules, see this website:

  `http://publib.boulder.ibm.com/infocenter/storwize/ic/index.jsp?topic=%2Fcom.ibm.storwize.V5000.doc%2Fsvc_configrulessummary_02171530.html`

  Fabrics: When a local fabric and a remote fabric are connected for Remote Copy purposes, the ISL hop count between a local node and a remote node cannot exceed seven.

Remote Copy link requirements

The following link requirements are valid for Metro Mirror and Global Mirror:

- Round-trip latency

  The total round-trip latency must be less than 80 ms and less than 40 ms in each direction. Latency simulations should be performed with your applications before any network links are put in place to see whether the applications perform at an acceptable level while they meet the round-trip latency requirement.

- Bandwidth

  The bandwidth must satisfy the following requirements:

  – If you are not using Change Volumes, be able to sustain peak write load for all mirrored volumes and background copy traffic.
  – If you are using Change Volumes with Global Mirror, be able to sustain change rate of Source Change Volumes and background copy traffic.
  – Other background copy rate (the best practice is 10% to 20% of maximum peak load) for initial synchronization and resynchronization.
  – Remote Copy internal communication at idle with or without Change Volumes is approximately 2.6 Mbps. This amount is the minimum amount.

Redundancy: If the link between two sites is configured with redundancy so that it can tolerate single failures, the link must be sized so that the bandwidth and latency requirement can be met during single failure conditions.

Interaction between Remote Copy and FlashCopy

Table 10-5 lists which combinations of FlashCopy and Remote Copy are supported.

Table 10-5 FlashCopy and Remote Copy interaction

Component	Remote Copy primary	Remote Copy secondary
FlashCopy source	Supported	Supported. When the FlashCopy relationship is in the Preparing and Prepared states, the cache at the Remote Copy secondary site operates in write-through mode. This process adds more latency to the already latent Remote Copy relationship.
FlashCopy target	This combination is supported and has the following restrictions: ► Running **`stop -force`** might cause the Remote Copy relationship to fully resynchronize. ► The I/O group must be the same.	This combination is supported by the restriction that the FlashCopy mapping cannot be copying, stopping, or suspended. Otherwise, the restrictions are the same as at the Remote Copy primary site.

If you are not using Global Mirror with Change Volumes, for disaster recovery purposes, you can use the FlashCopy feature to create a consistent copy of an image before you restart a Global Mirror relationship.

When a consistent relationship is stopped, the relationship enters the consistent_stopped state. While in this state, I/O operations at the primary site continue to run. However, updates are not copied to the secondary site. When the relationship is restarted, the synchronization process for new data is started. During this process, the relationship is in the inconsistent_copying state.

The secondary volume for the relationship cannot be used until the copy process completes and the relationship returns to the consistent state. When this situation occurs, start a FlashCopy operation for the secondary volume before you restart the relationship. While the relationship is in the Copying state, the FlashCopy feature can provide a consistent copy of the data. If the relationship does not reach the synchronized state, you can use the FlashCopy target volume at the secondary site.

10.3 Troubleshooting Remote Copy

Remote Copy (Global Mirror and Metro Mirror) has the following primary error codes:

- 1920: This error is a congestion error that means that the source, the link between source and target, or the target cannot keep up with the rate of demand.
- 1720. This error is a heartbeat or cluster partnership communication error. This error tends to be more serious because failing communication between your cluster partners involves some extended diagnostic time.

10.3.1 1920 error

A 1920 error (event ID 050010) can have several triggers. The following official probable cause projections are available:

- Primary cluster or SAN fabric problem (10%)
- Primary cluster or SAN fabric configuration (10%)
- Secondary cluster or SAN fabric problem (15%)
- Secondary cluster or SAN fabric configuration (25%)
- Inter-cluster link problem (15%)
- Inter-cluster link configuration (25%)

In practice, the error that is most often overlooked is latency. Global Mirror has a round-trip-time tolerance limit of 80 ms. A message that is sent from your source cluster to your target cluster and the accompanying acknowledgement must have a total time of 80 ms (or 40 ms each way).

The primary component of your round-trip time is the physical distance between sites. For every 1,000 km (621.36 miles), there is a 5 ms delay. This delay does not include the time that is added by equipment in the path. Every device adds a varying amount of time depending on the device, but you can expect about 25 µs for pure hardware devices. For software-based functions (such as, compression that is implemented in software), the added delay tends to be much higher (usually in the millisecond-plus range).

Consider this example. Company A has a production site that is 1,900 km from their recovery site. Their network service provider uses five devices to connect the two sites. In addition to those devices, Company A uses a SAN Fibre Channel Router at each site to provide FCIP to encapsulate the Fibre Channel traffic between sites. There are now seven devices, and 1,900 km of distance delay. All the devices add 200 µs of delay each way. The distance adds 9.5 ms each way, for a total of 19 ms. Combined with the device latency that is 19.4 ms of physical latency at a minimum. This latency is under the 80 ms limit of Global Mirror, but this number is the best case number. Link quality and bandwidth play a significant role here. Your network provider likely ensures a latency maximum on your network link; be sure to stay below the Global Mirror RTT (Round Trip Time) limit. You can easily double or triple the expected physical latency with a lower quality or lower bandwidth network link. As a result, you are suddenly within range of exceeding the limit the moment a large flood of I/O happens that exceeds the bandwidth capacity you have in place.

When you get a 1920 error, always check the latency first. The FCIP routing layer can introduce latency if it is not properly configured. If your network provider reports a much lower latency, this report can be an indication of a problem at your FCIP Routing layer. Most FCIP Routing devices have built-in tools that you can use to check the round-trip delay time (RTT). When you are checking latency, remember that TCP/IP routing devices (including FCIP routers) report RTT by using standard 64-byte ping packets.

Figure 10-85 shows why the effective transit time should be measured only by using packets large enough to hold a Fibre Channel frame. This packet size is 2148 bytes (2112 bytes of payload and 36 bytes of header) and you should allow more capacity to be safe because different switching vendors have optional features that might increase this size.

Packet Size	Link Size	Serialization Delay (Time Required to Send Data)	Unit
64	256 Kbps	2.0E+03	microseconds
64	1.5 Mbps	3.4E+02	microseconds
64	100 Mbps	5.1E+00	microseconds
64	155 Mbps	3.3E+00	microseconds
64	622 Mbps	8.2E-01	microseconds
64	1 Gbps	5.1E-04	microseconds
64	10 Gbps	5.1E-05	microseconds
1500	256 Kbps	4.7E+04	microseconds
1500	1.5 Mbps	8.0E+03	microseconds
1500	100 Mbps	1.2E+02	microseconds
1500	155 Mbps	7.7E+01	microseconds
1500	622 Mbps	1.9E+01	microseconds
1500	1 Gbps	1.2E+01	microseconds
1500	10 Gbps	1.2E+00	microseconds
2148	256 Kbps	6.7E+04	microseconds
2148	1.5 Mbps	1.1E+04	microseconds
2148	100 Mbps	1.7E+02	microseconds
2148	155 Mbps	1.1E+02	microseconds
2148	622 Mbps	2.8E+01	microseconds
2148	1 Gbps	1.7E+01	microseconds
2148	10 Gbps	1.7E-03	microseconds

Figure 10-85 Effect of packet size (in bytes) versus link size

Before you proceed, take a quick look at the second largest component of your round-trip-time; that is, serialization delay. Serialization delay is the amount of time that is required to move a packet of data of a specific size across a network link of a bandwidth. This delay is based on a simple concept that the time that is required to move a specific amount of data decreases as the data transmission rate increases.

In Figure 10-85, there are orders of magnitude of difference between the different link bandwidths. It is easy to see how 1920 errors can arise when your bandwidth is insufficient and why you should never use a TCP/IP **`ping`** to measure RTT for FCIP traffic.

Figure 10-85 compares the amount of time in microseconds that is required to transmit a packet across network links of varying bandwidth capacity. The following packet sizes are used:

- 64 bytes: The size of the common ping packet
- 1500 bytes: The size of the standard TCP/IP packet
- 2148 bytes: The size of a Fibre Channel frame

Your path MTU affects the delay that is incurred in getting a packet from one location to another, when it causes fragmentation, or is too large and causes too many retransmits when a packet is lost. After you verified your latency by using the correct packet size, proceed with normal hardware troubleshooting.

10.3.2 1720 error

The 1720 error (event ID 050020) is the other primary error code of Remote Copy. Because the term “System Partnership” implies that all involved virtualization systems are partners, they must communicate with each other. When a partner on either side stops communicating, you see a 1720 error appear in your error log. According to official documentation, there are no likely field replaceable unit breakages or other causes.

In practice, the source of this error is most often a fabric problem or a problem in the network path between your partners. When you receive this error, if your fabric has more than 64 HBA zoned ports, you should check your fabric configuration for zoning of more than one HBA port for each node per I/O group. One port for each node per I/O group that is associated with the host is the recommended zoning configuration for fabrics. For those fabrics with 64 or more host ports, this recommendation becomes a rule. You must follow this zoning rule or the configuration is technically unsupported.

Improper zoning leads to SAN congestion, which can inhibit remote link communication intermittently. Checking the zero buffer credit timer through IBM Tivoli Storage Productivity Center and comparing its value against your sample interval might reveal potential SAN congestion. When a zero buffer credit timer is above 2% of the total time of the sample interval, it is likely to cause problems.

Next, always ask your network provider to check the status of the link. If the link is okay, watch for repetition of this error. It is possible in a normal and functional network setup to have occasional 1720 errors, but multiple occurrences indicate a larger problem.

If you receive multiple 1720 errors, recheck your network connection and then check the IBM Storwize V5000 partnership information to verify their status and settings. Perform diagnostic tests for every piece of equipment in the path between the two systems. It often helps to have a diagram that shows the path of your replication from logical and physical configuration viewpoints.

If your investigation fails to resolve your Remote Copy problems, you should contact your IBM support representative for a complete analysis.

10.4 Managing Remote Copy by using the GUI

The IBM Storwize V5000 storage system provides a separate function icon for copy service management. The following windows are available for managing Remote Copy, which are accessed through the Copy Services function icon:

- Remote Copy
- Partnerships

As the name implies, these windows are used to manage Remote Copy and the partnership.

10.4.1 Managing cluster partnerships

The Partnership window is used to manage a partnership between clusters. To access the Partnership window, click the **Copy Services** function icon and then click **Partnerships**, as shown in Figure 10-86.

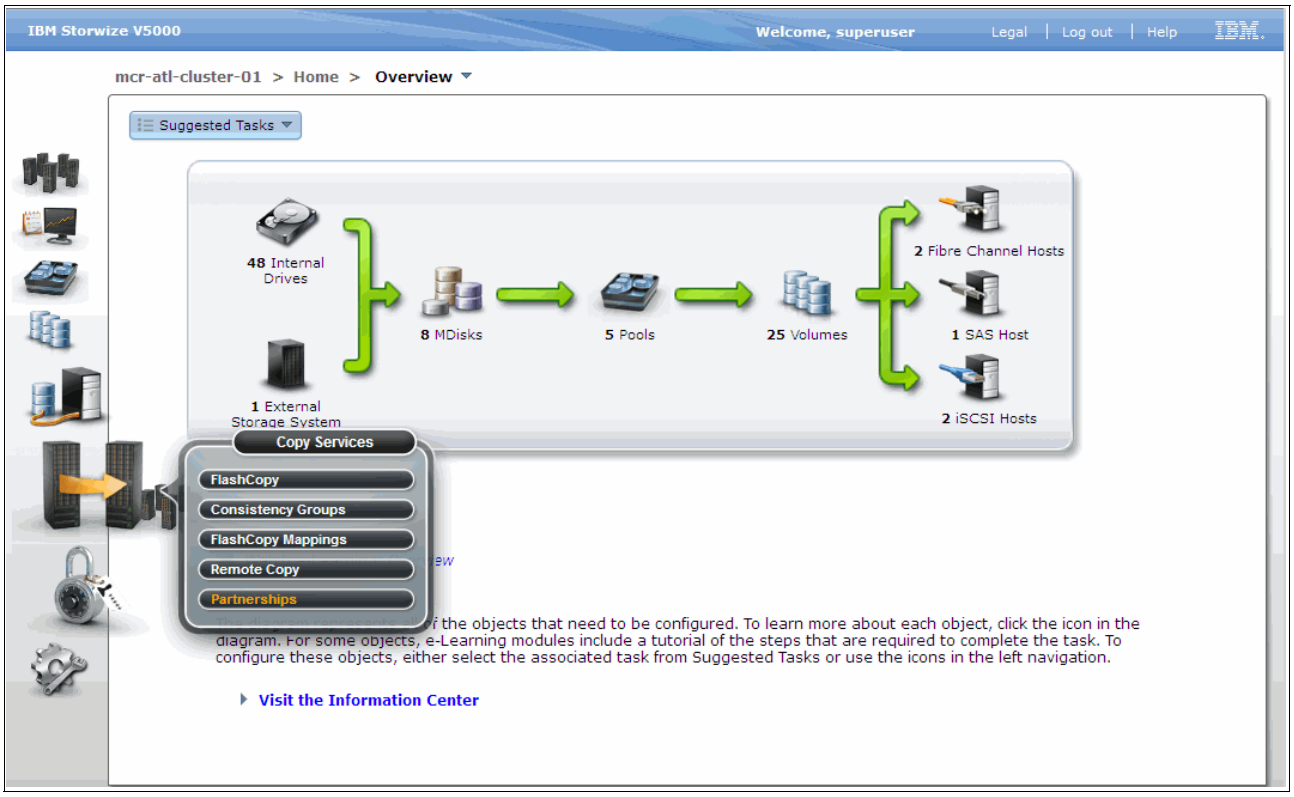

Figure 10-86 Partnership window

Creating a partnership

No partnership is defined in our example (see Figure 10-87), so you must create a partnership between the IBM Storwize V5000 systems. Click **New Partnership** in the Partnership window.

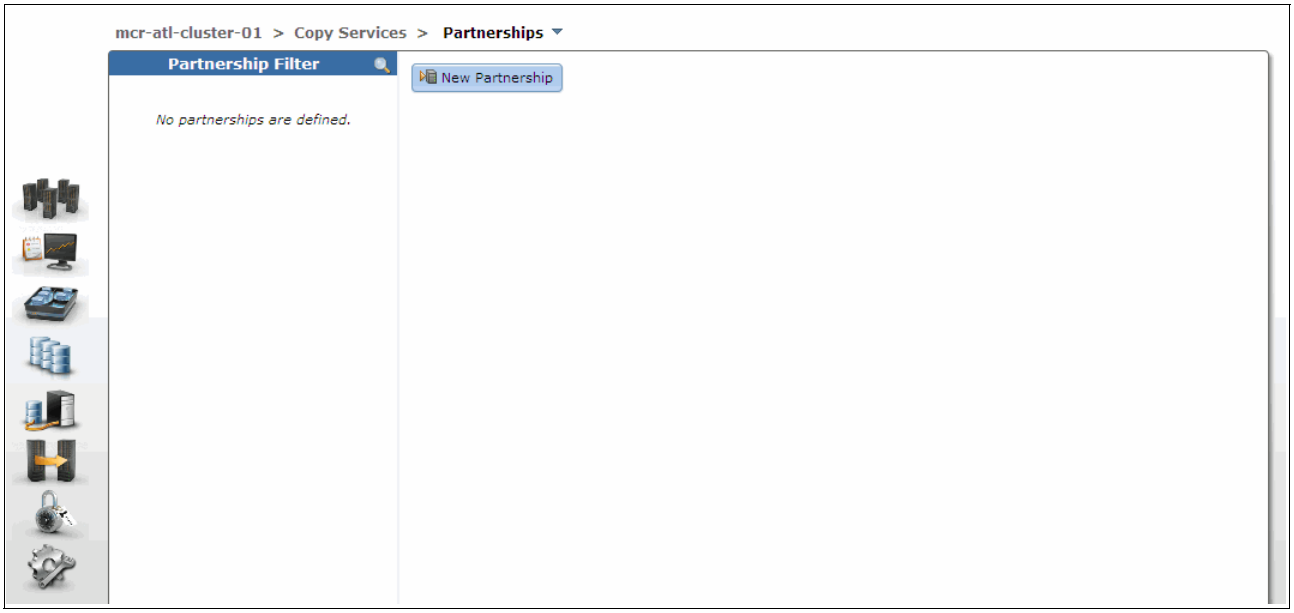

Figure 10-87 Create a cluster partnership

If there is no partnership candidate, an error window opens, as shown in Figure 10-88.

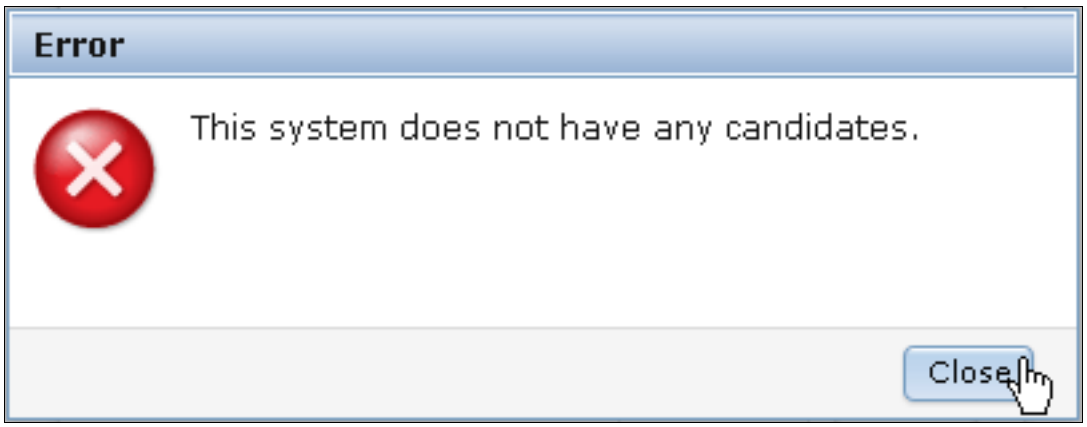

Figure 10-88 No candidates are available to create a partnership

Check the zoning and the system status and make sure that the clusters can see each other. Then, you can create your partnership by selecting the appropriate remote storage system (as shown in Figure 10-89 on page 520), and defining the available bandwidth between both systems.

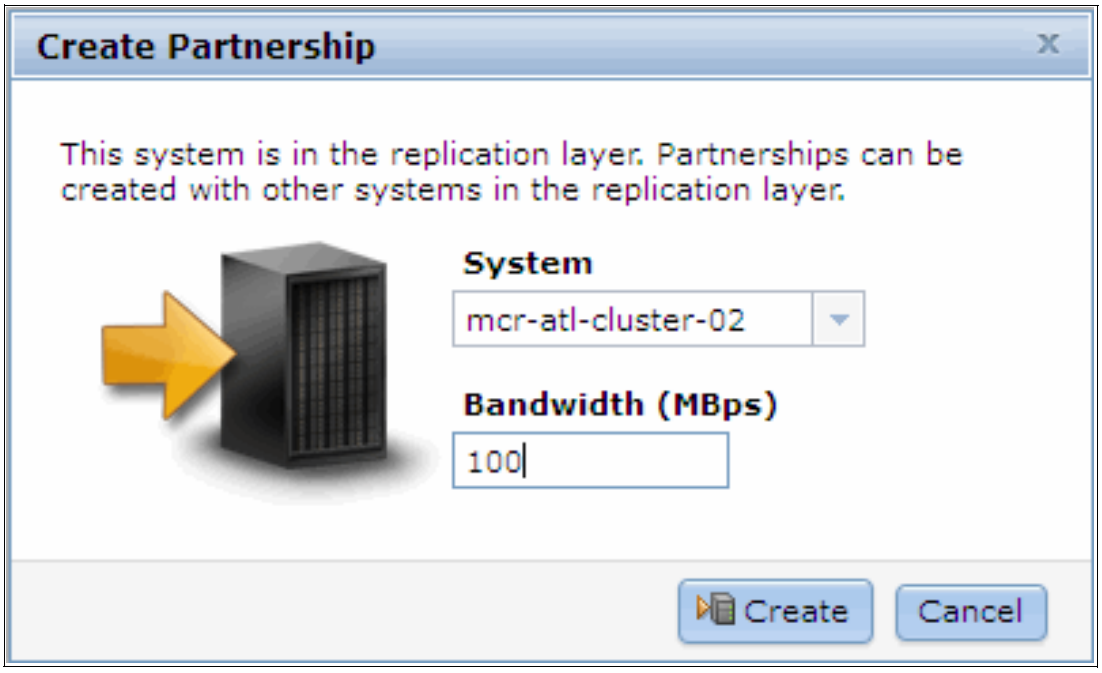

Figure 10-89 Select the remote IBM Storwize storage system for a new partnership

The bandwidth that you must enter here is used by the background copy process between the clusters in the partnership. To set the background copy bandwidth optimally, make sure that you consider all three resources (primary storage, inter-cluster link bandwidth, and auxiliary storage) to avoid overloading them, which affects the foreground I/O latency.

Click **Create** and the partnership definition is complete on the first IBM Storwize V5000 system. You can find the partnership that is listed in the left pane of the Partnership window. If you select the partnership, more information for this partnership is displayed on the right, as shown in Figure 10-90.

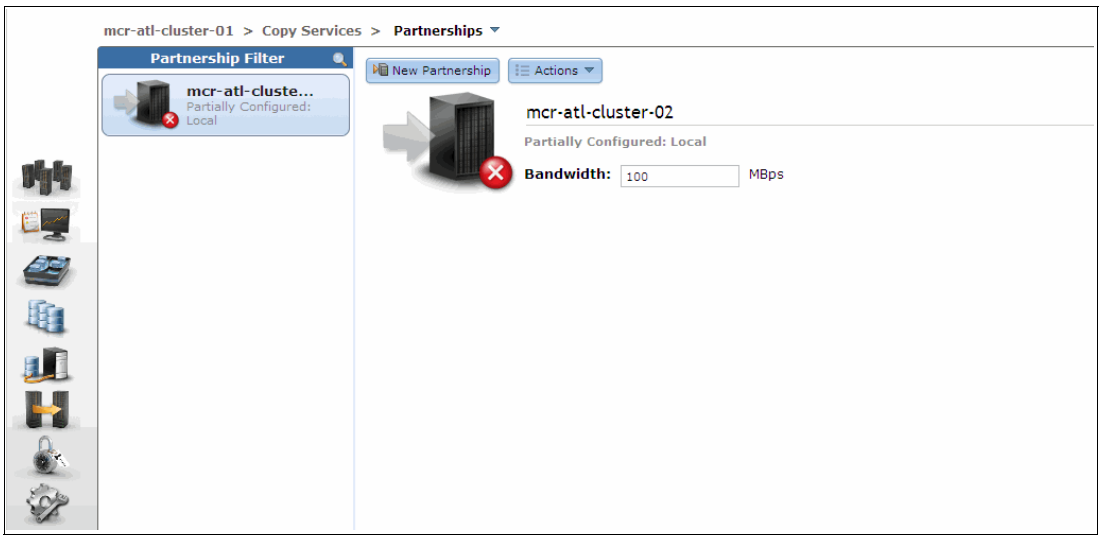

Figure 10-90 Partially configured partnership

Important: The partnership is in the “Partially Configured: Local” state because we did not yet define it on the other IBM Storwize V5000. For more information about partnership states, see “Remote Copy and consistency group states” on page 511.

Complete the same steps on the second storage system for the partnership to become fully configured. The Remote Copy partnership is now implemented between the two IBM Storwize V5000 systems and both systems are ready for further configuration of Remote Copy relationships, as shown in Figure 10-91.

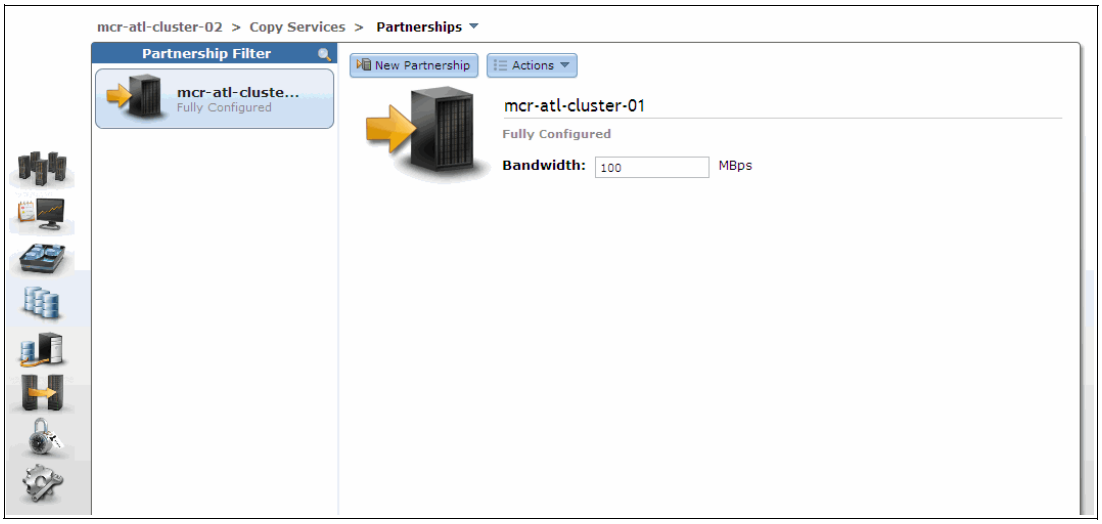

Figure 10-91 Fully configured partnership

You can also change the bandwidth setting for the partnership in the Partnerships window. Click **Apply Changes** to confirm your modification.

Stopping and starting a partnership

You can stop the partnership by clicking **Stop Partnership** from the Actions drop-down menu, as shown in Figure 10-92. If you stop the partnership, the relationship that uses this partnership is disconnected.

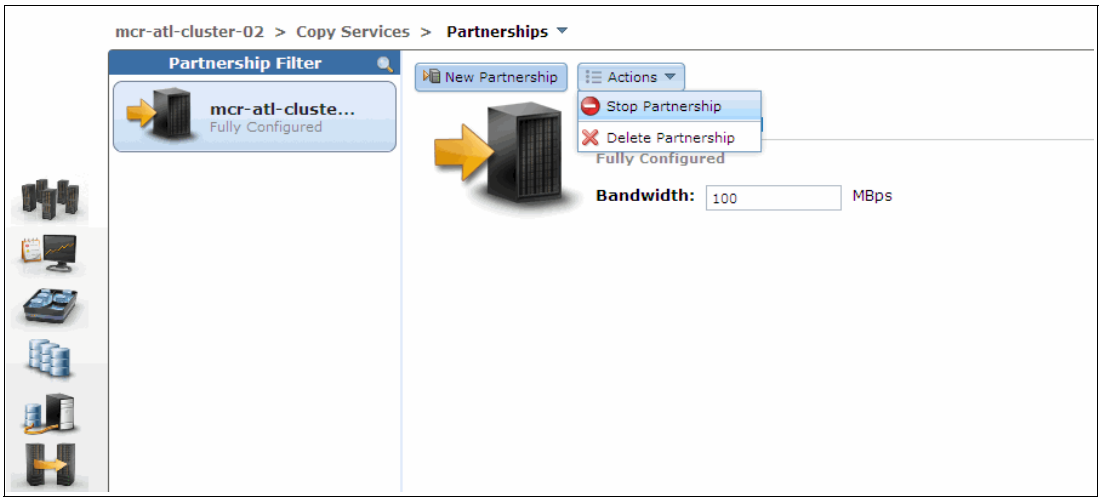

Figure 10-92 Stop the partnership

After you stop the partnership, your partnership is listed as Fully Configured: Stopped, as shown in Figure 10-93.

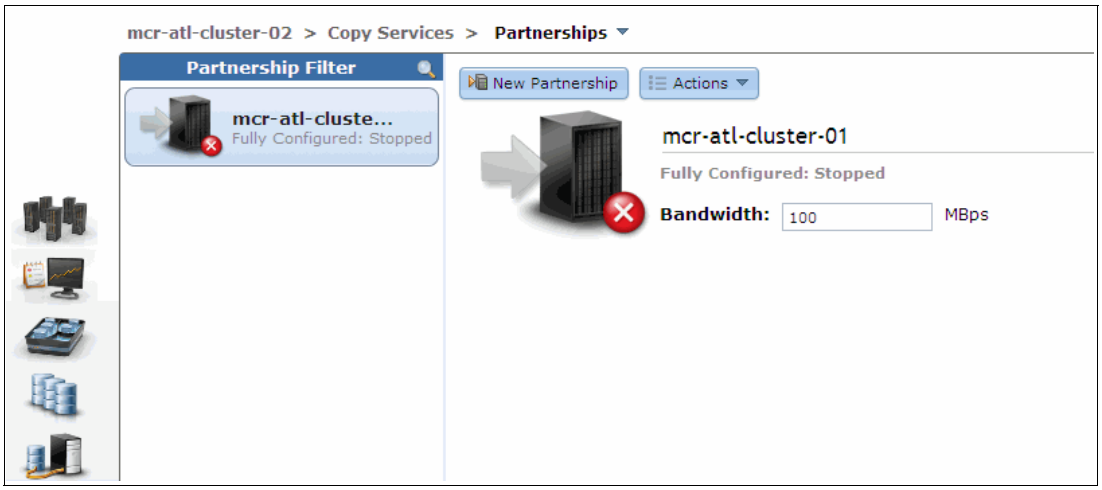

Figure 10-93 Fully configured partnership in Stopped state

You can restart a stopped partnership by clicking **Start Partnership** from the Actions drop-down menu.

The partnership returns to the fully configured status when it is restarted.

Deleting a partnership

You can delete a partnership by selecting **Delete Partnership** from the Actions drop-down menu, as shown in Figure 10-92 on page 521.

10.4.2 Managing stand-alone Remote Copy relationships

A Remote Copy relationship can be defined between two volumes where one is the master (source) and the other one is the auxiliary (target) volume. Use of Remote Copy auxiliary volumes as Remote Copy master volumes is not allowed. Open the Remote Copy window to manage Remote Copy by clicking the **Copy Services** icon and then clicking **Remote Copy**, as shown in Figure 10-94 on page 523.

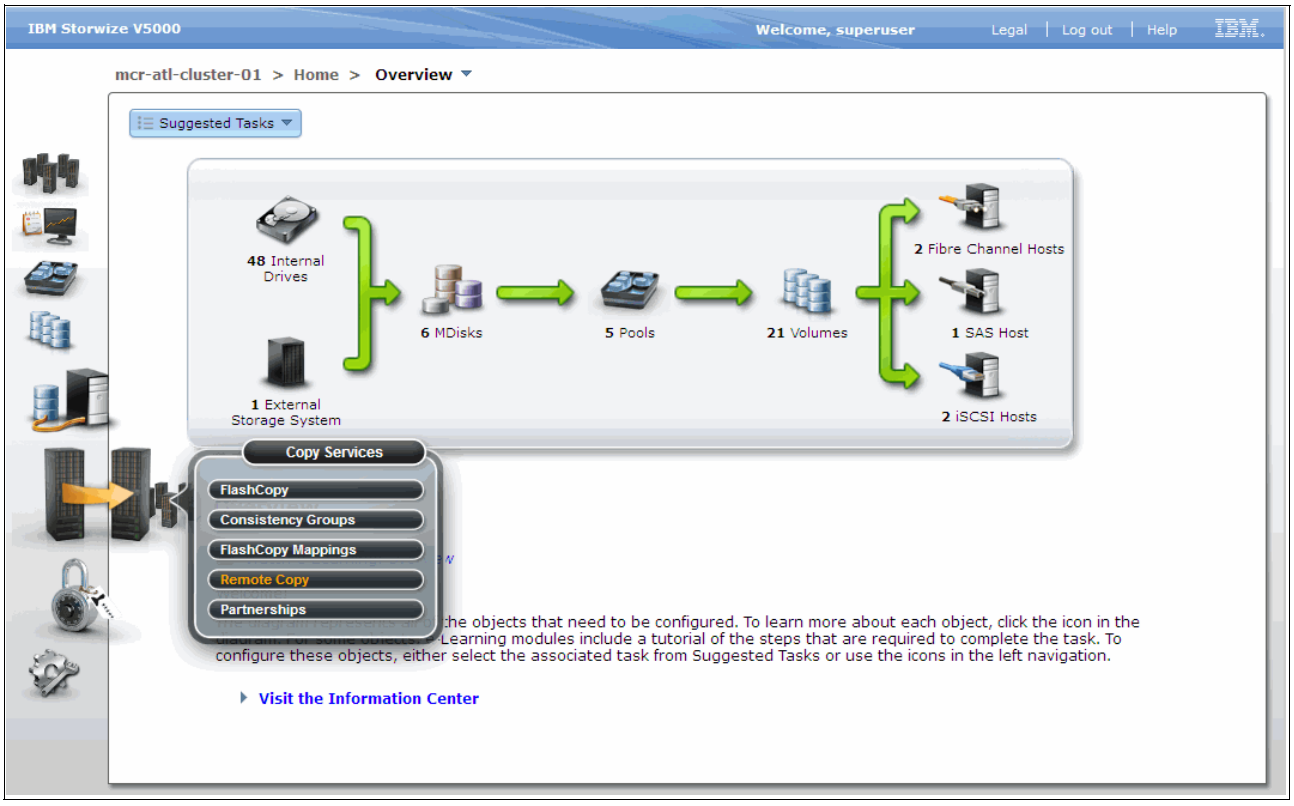

Figure 10-94 Open Remote Copy window

As shown in Figure 10-95, the Remote Copy window is where you can manage Remote Copy relationships and Remote Copy consistency groups.

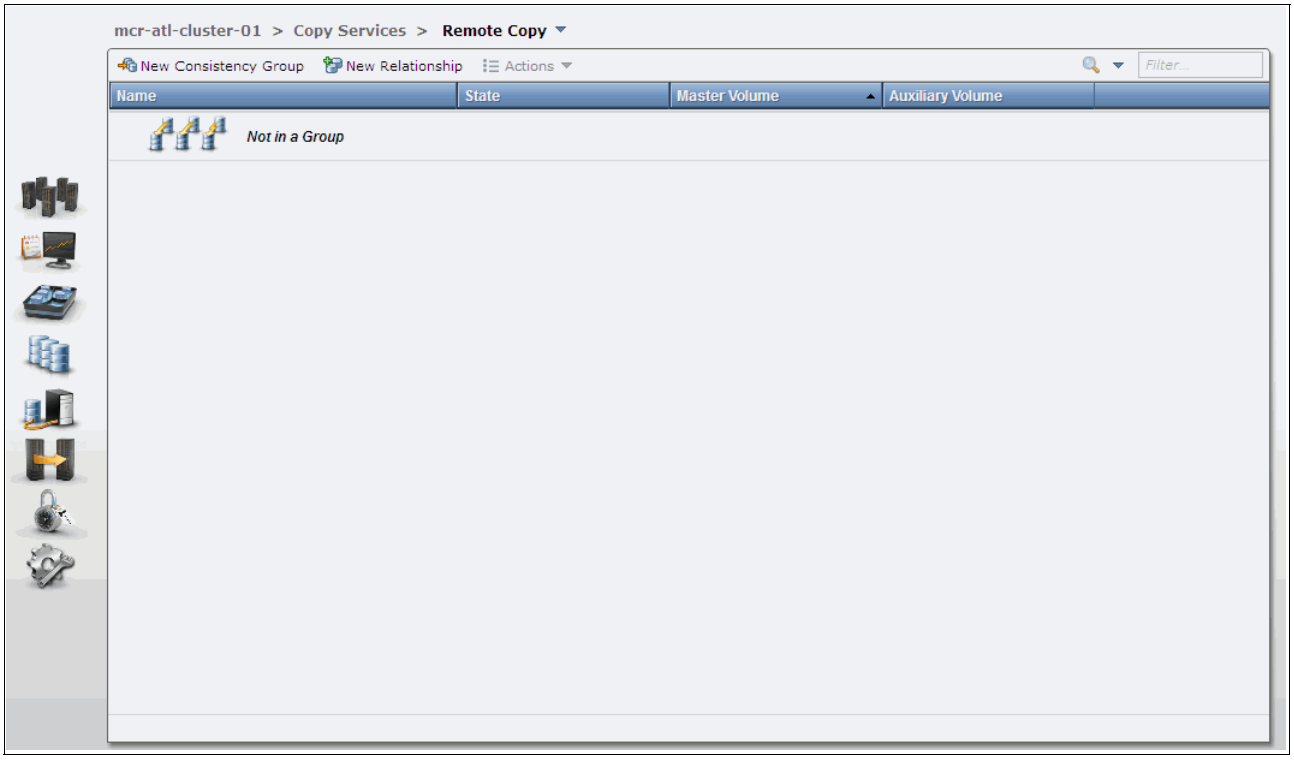

Figure 10-95 Remote Copy window

The Remote Copy window displays a list of Remote Copy consistency groups. You can also take actions on the Remote Copy relationships and Remote Copy consistency groups. Click **Not in a Group** and all the Remote Copy relationships that are not in any Remote Copy consistency groups are displayed. To customize the blue column heading bar and select different attributes of Remote copy relationships, right-click anywhere in the blue bar.

Creating stand-alone Remote Copy relationships

Important: Before a remote copy relationship is created, target volumes that are the same size as the source volumes that you want to mirror must be created. For more information about creating volumes, see Chapter 5, “I/O Group basic volume configuration” on page 161.

To create a Remote Copy relationship, click **New Relationship** at the top of the Remote Copy window, as shown in Figure 10-95 on page 523. A wizard opens and guides you through the Remote Copy relationship creation process.

As shown in Figure 10-96, you must set the Remote Copy relationship type first. Based on your requirements, you can select Metro Mirror (synchronous replication) or Global Mirror (asynchronous replication). Select the appropriate replication type and click **Next**.

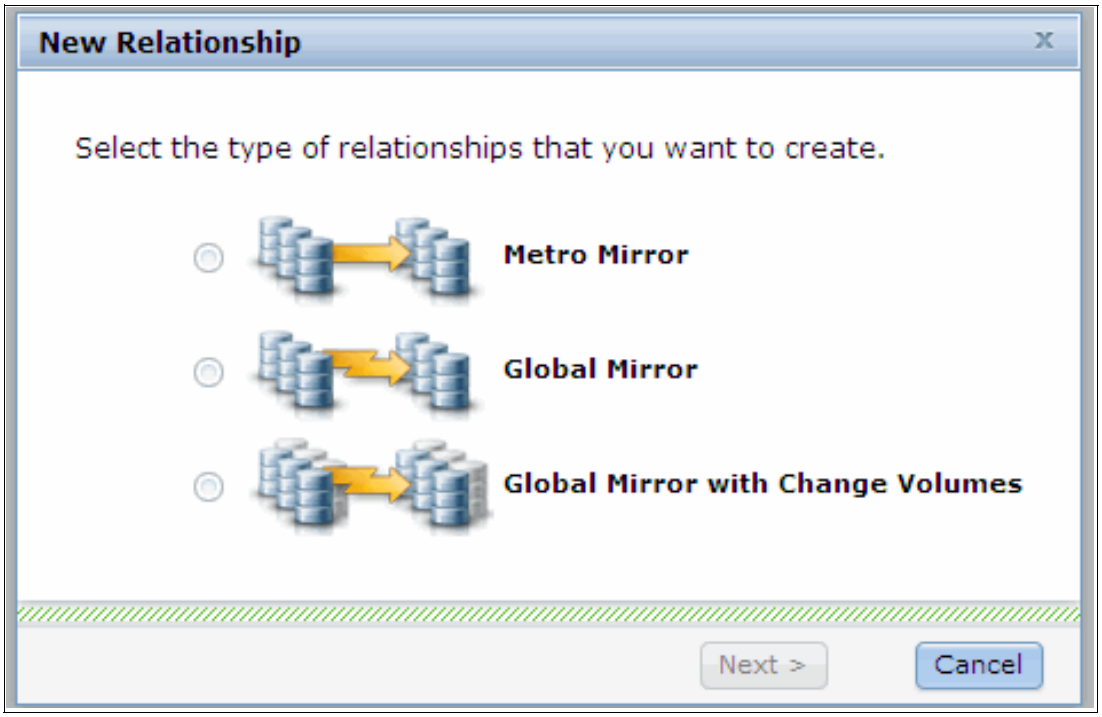

Figure 10-96 Select the appropriate Remote Copy type

You must select where your auxiliary (target) volumes are: the local system or the already defined second storage system. In our example (as shown in Figure 10-97), choose another system to build an inter-cluster relationship. Click **Next** to continue.

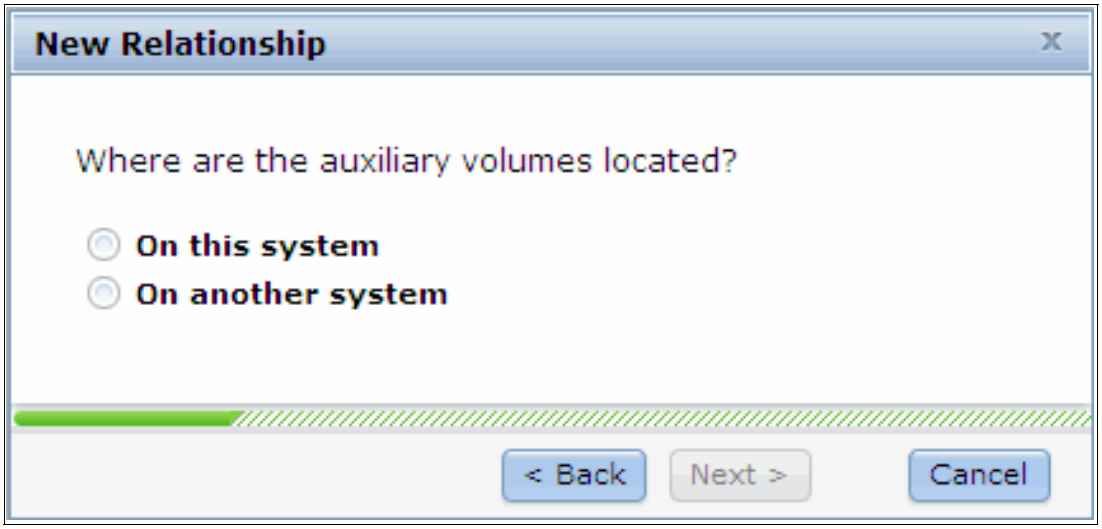

Figure 10-97 Select Remote Copy partner

The Remote Copy master and auxiliary volume must be specified. Both volumes must have the same size. As shown in Figure 10-98, the system offers only appropriate auxiliary candidates with the same volume size as the selected master volume. After you select the volumes that are based on your requirement, click **Add**.

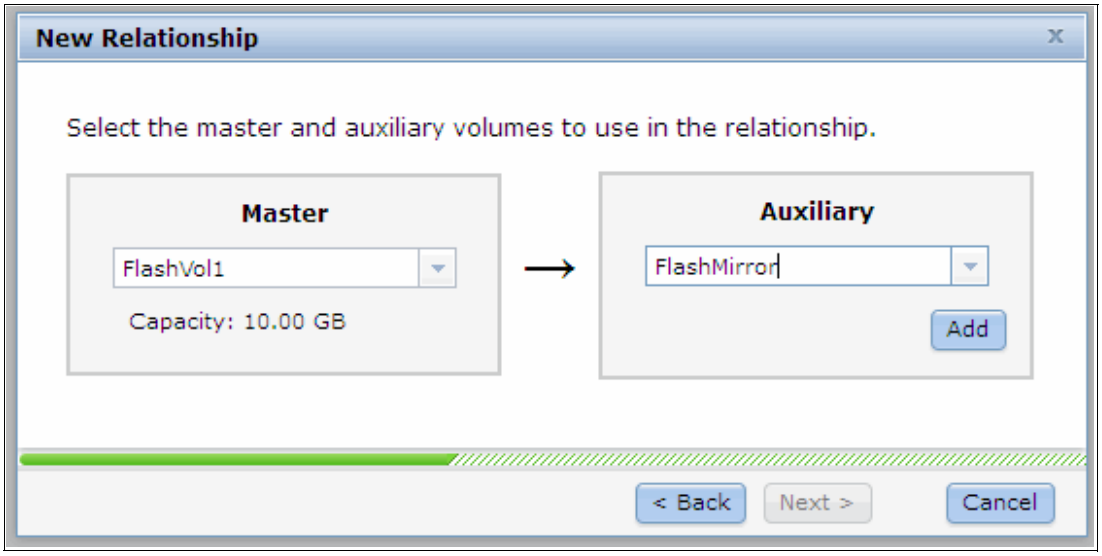

Figure 10-98 Select the master and auxiliary volume

You can define multiple and independent relationships by clicking **Add**. You can remove a relationship by clicking the red cross. In our example, we create two independent Remote Copy relationships, as shown in Figure 10-99.

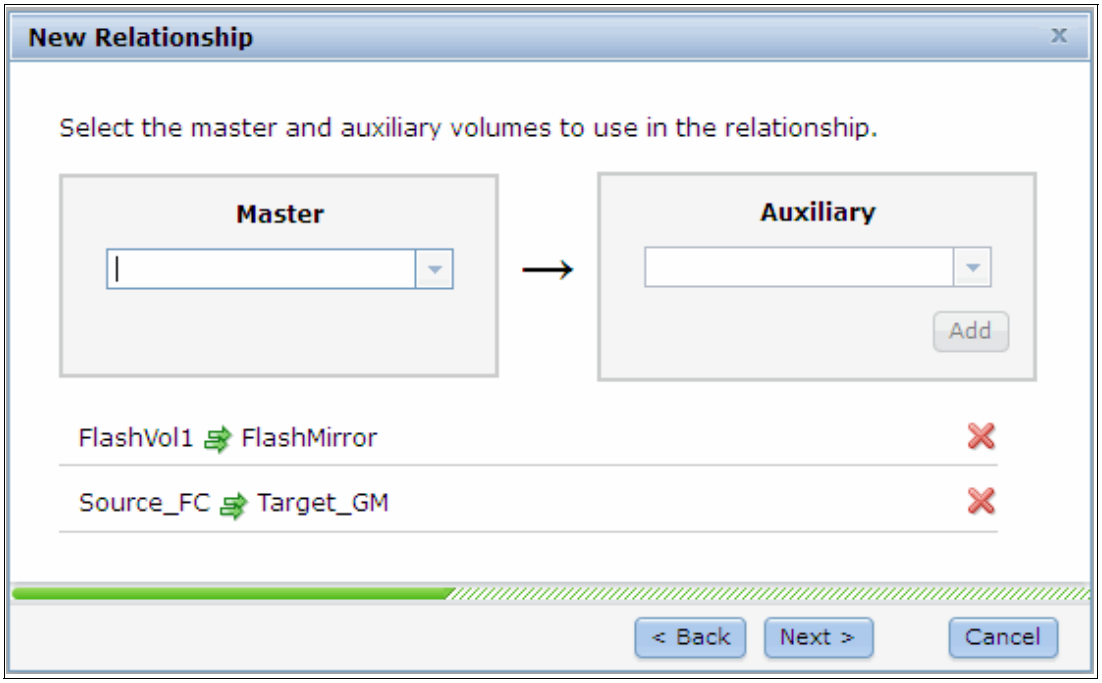

Figure 10-99 Define multiple independent relationships

A window opens and prompts you to select if the volumes in the relationship are already synchronized. In most situations, the data on the master volume and on the auxiliary volume are not identical, so click **No** and then click **Next** to enable an initial copy, as shown in Figure 10-100.

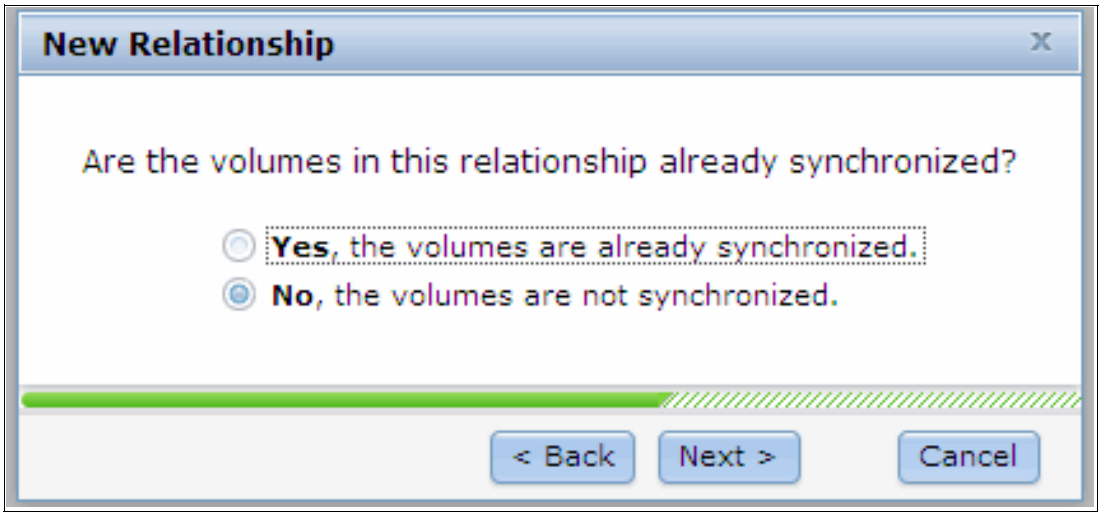

Figure 10-100 Activate initial data copy

If you select **Yes, the volumes are already synchronized** in this step, a warning message opens, as shown in Figure 10-101. Confirm that the volumes are truly identical, and then click **OK** to continue.

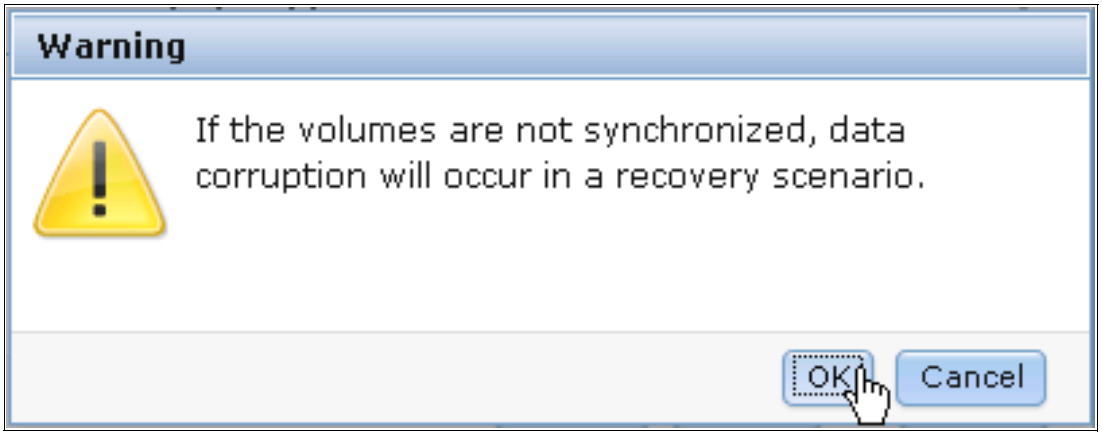

Figure 10-101 Warning message to make sure that the volumes are synchronized

You can choose to start the initial copying progress now or wait to start it later. In our example, select **Yes, start copying now** and click **Finish**, as shown in Figure 10-102.

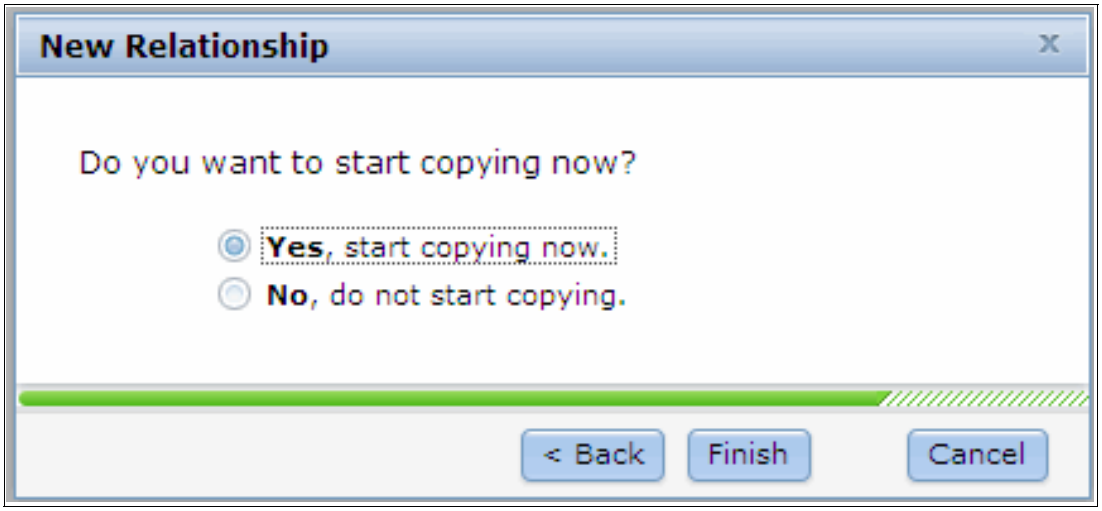

Figure 10-102 Choose if you want to start copying now or later

After the Remote Copy relationships creation completes, two independent Remote Copy relationships are defined and displayed in the Not in a Group list, as shown in Figure 10-103.

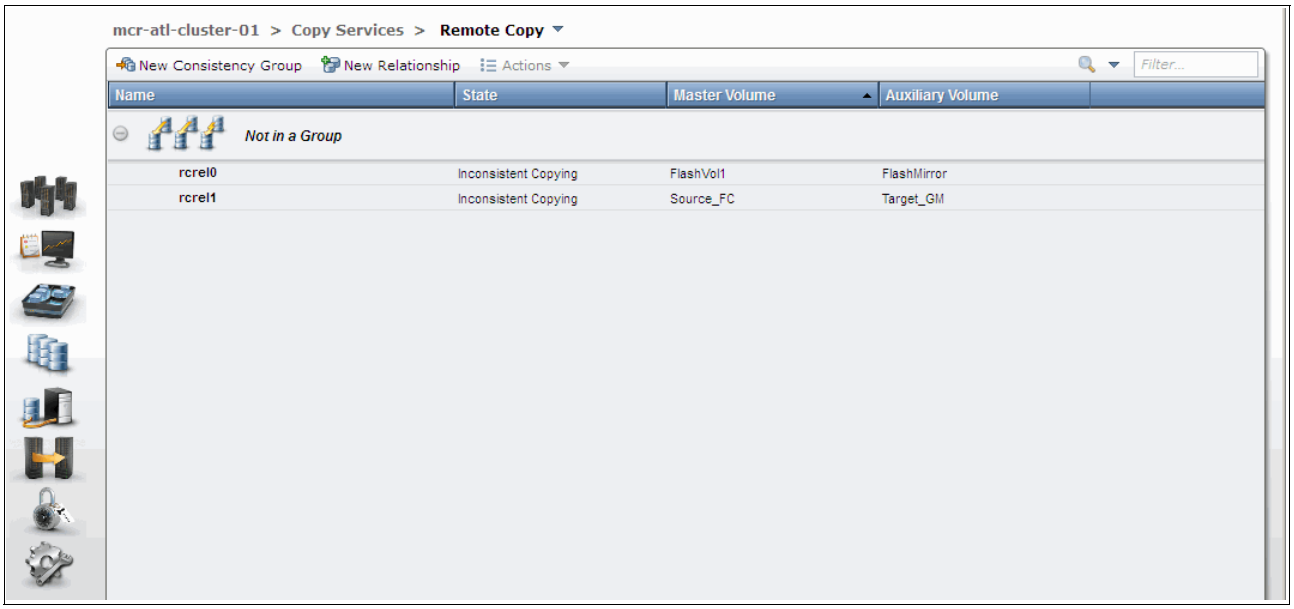

Figure 10-103 Creating a Remote Copy relationship process completes

Optionally, you can monitor the ongoing initial synchronization in the Running Tasks status indicator, as shown in Figure 10-104. Highlight one of the operations and click to see the progress.

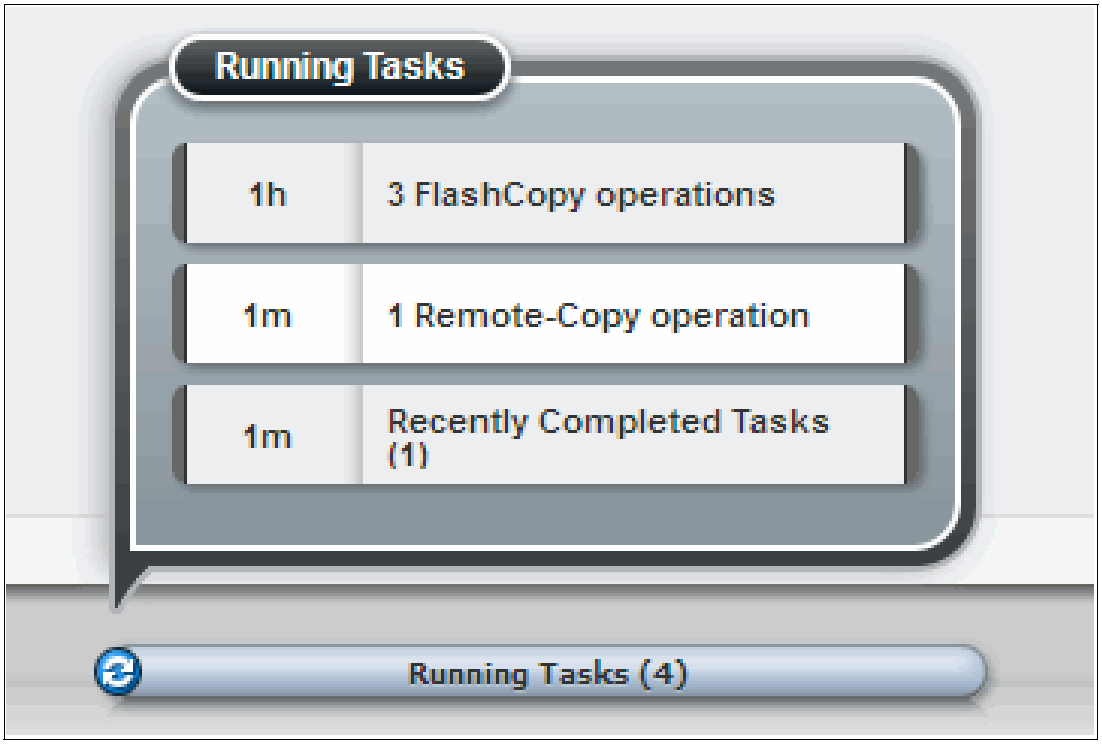

Figure 10-104 Remote copy initialization progress through Running Tasks

Stopping a stand-alone Remote Copy relationship

The Remote Copy relationship can be stopped by selecting the relationship and clicking **Stop** from the Actions drop-down menu, as shown in Figure 10-105.

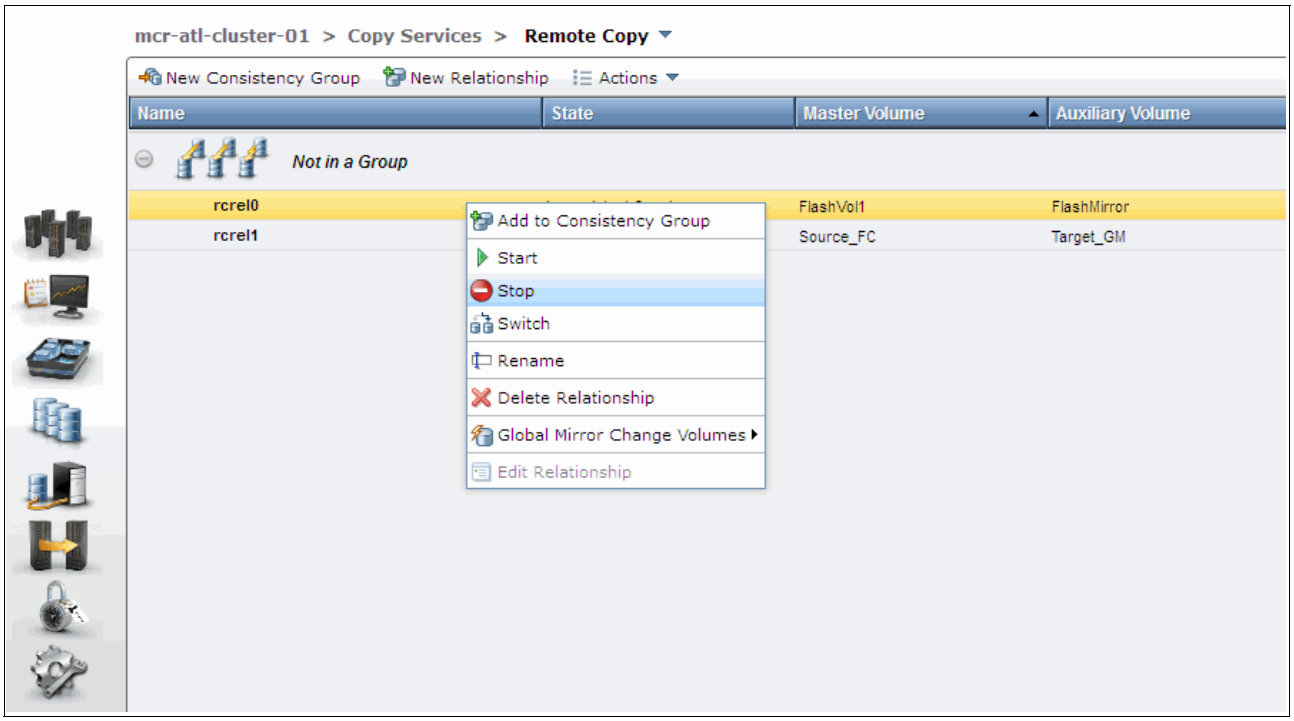

Figure 10-105 Stop Remote Copy relationship

A prompt appears. Click to allow secondary read/write access, if required, and then click **Stop Relationship**, as shown in Figure 10-106.

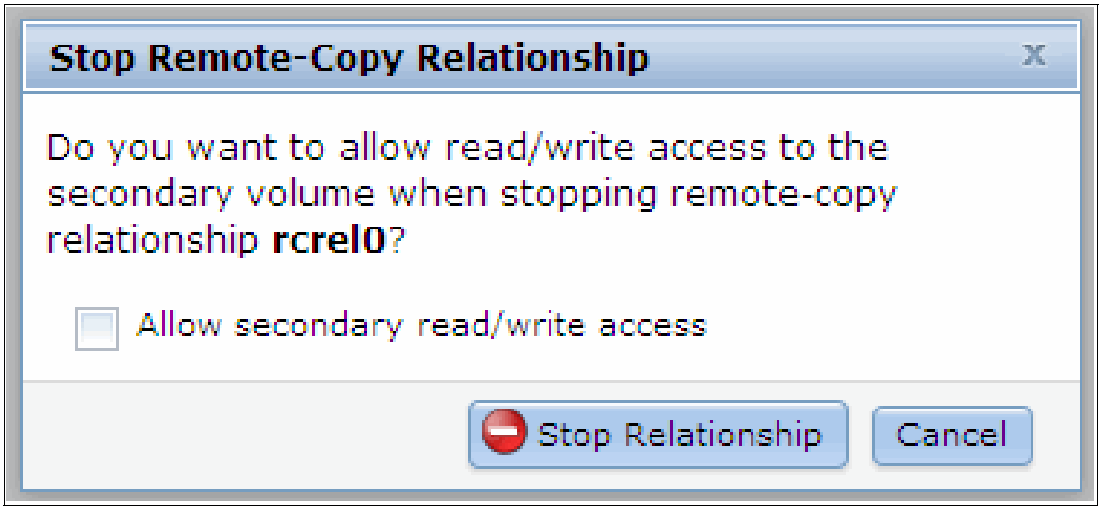

Figure 10-106 Option to allow secondary read/write access

After the stop completes, the state of the Remote Copy relationship is changed from Consistent Synchronized to Idling, as shown in Figure 10-107. Read/write access to both volumes is now allowed unless you selected otherwise.

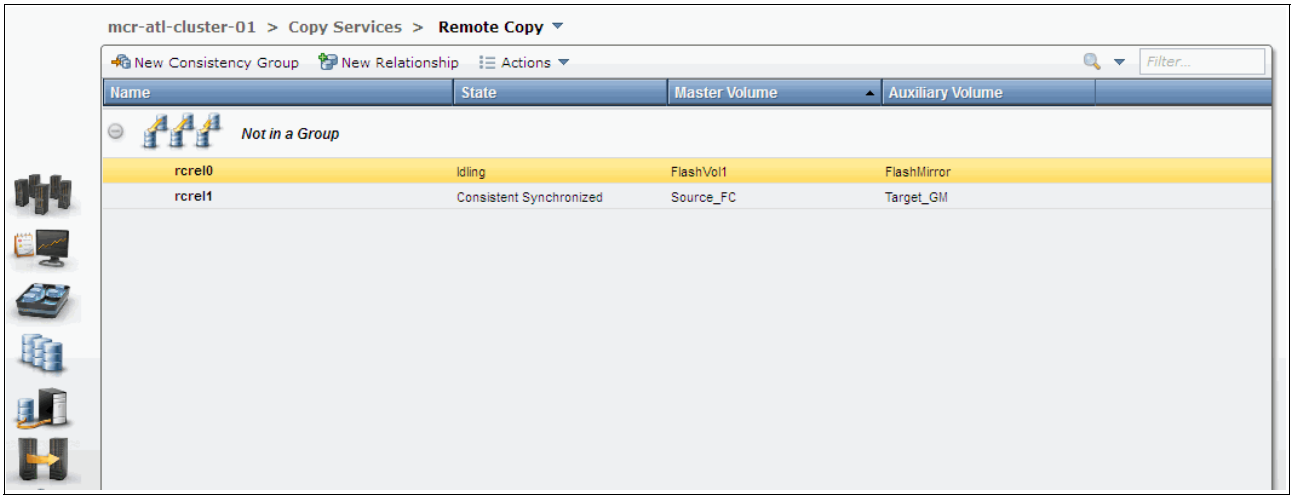

Figure 10-107 Remote Copy relationship stop completes

Starting a stand-alone Remote Copy relationship

The Remote Copy relationship can be started by selecting the relationship and clicking **Start** from the Actions drop-down menu, as shown in Figure 10-108.

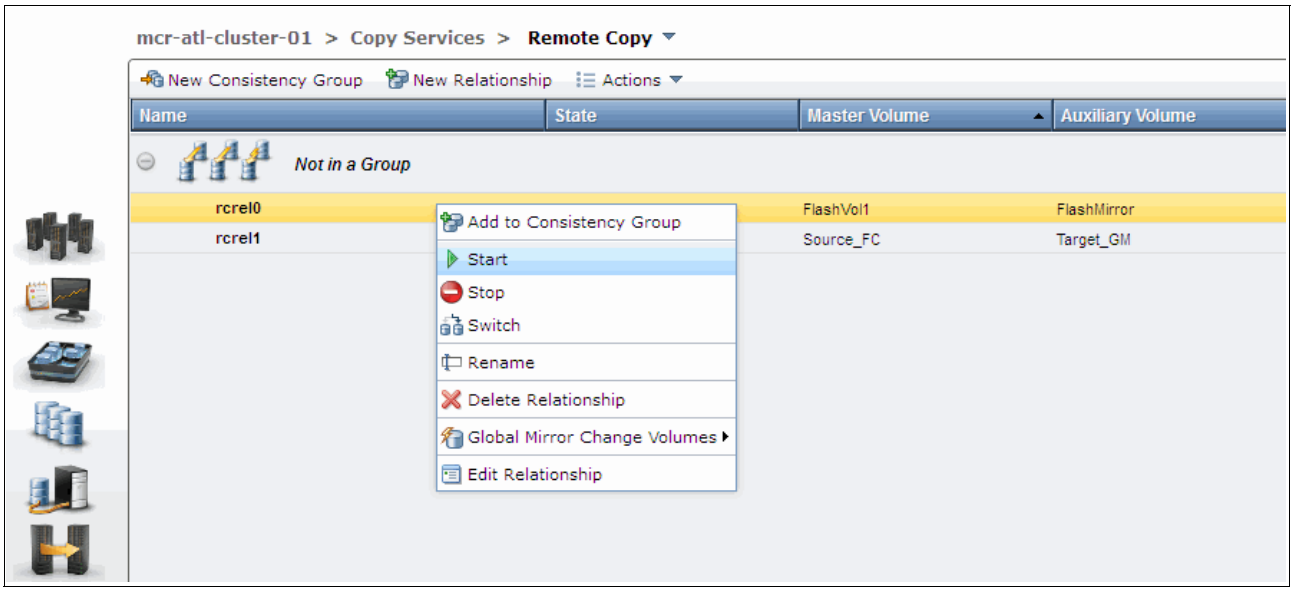

Figure 10-108 Start a Remote Copy relationship

When a Remote Copy relationship is started, the most important item is selecting the copy direction. Both master and auxiliary volumes can be the primary. Make your decision that is based on your requirements and click **Start Relationship**. In our example, choose the master volume to be the primary, as shown in Figure 10-109.

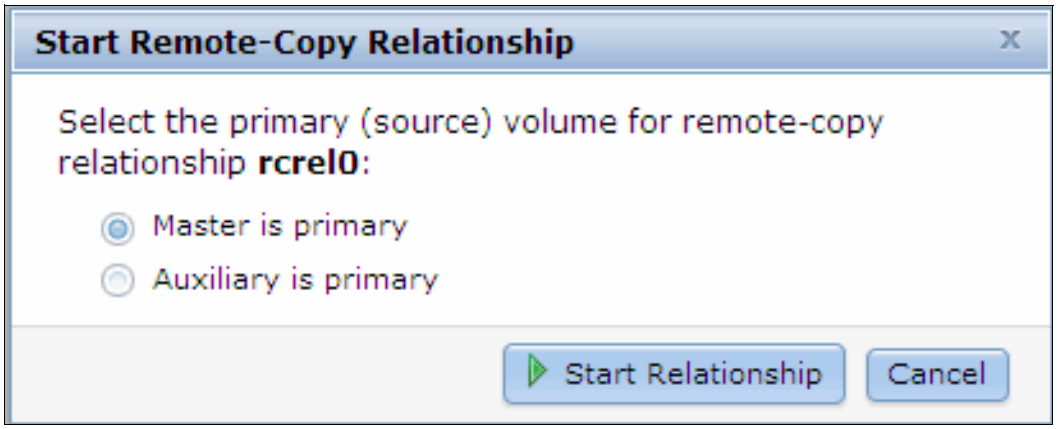

Figure 10-109 Choose the copy direction

Switching the direction of a stand-alone Remote Copy relationship

The copy direction of the Remote Copy relationship can be switched by selecting the relationship and clicking **Switch** from the Actions drop-down menu, as shown in Figure 10-110.

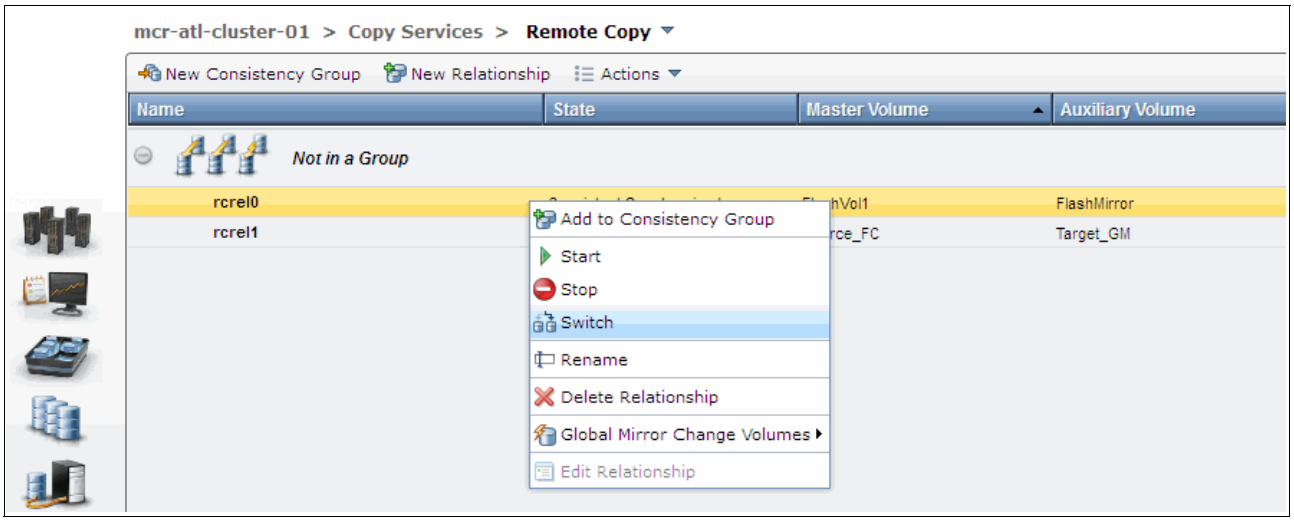

Figure 10-110 Switch Remote Copy relationship

A warning message opens and shows you the consequences of this action, as shown in Figure 10-111 on page 532. If you switch the Remote Copy relationship, the copy direction of the relationship becomes the opposite; that is, the current primary volume becomes the secondary while the current secondary volume becomes the primary. Write access to the current primary volume is lost and write access to the current secondary volume is enabled. If it is not a disaster recovery situation, you must stop your host I/O to the current primary volume in advance. Make sure that you are prepared for the consequences. If so, click **OK** to continue.

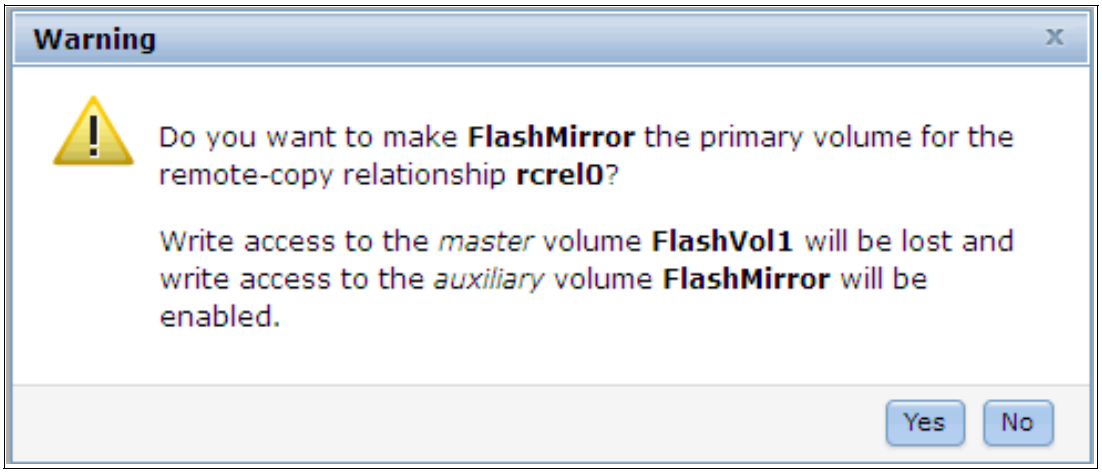

Figure 10-111 Warning message for switching direction of a Remote Copy relationship

After the switch completes, your Remote Copy relationship is tagged (as shown in Figure 10-112), and shows you that the primary volume in this relationship was changed.

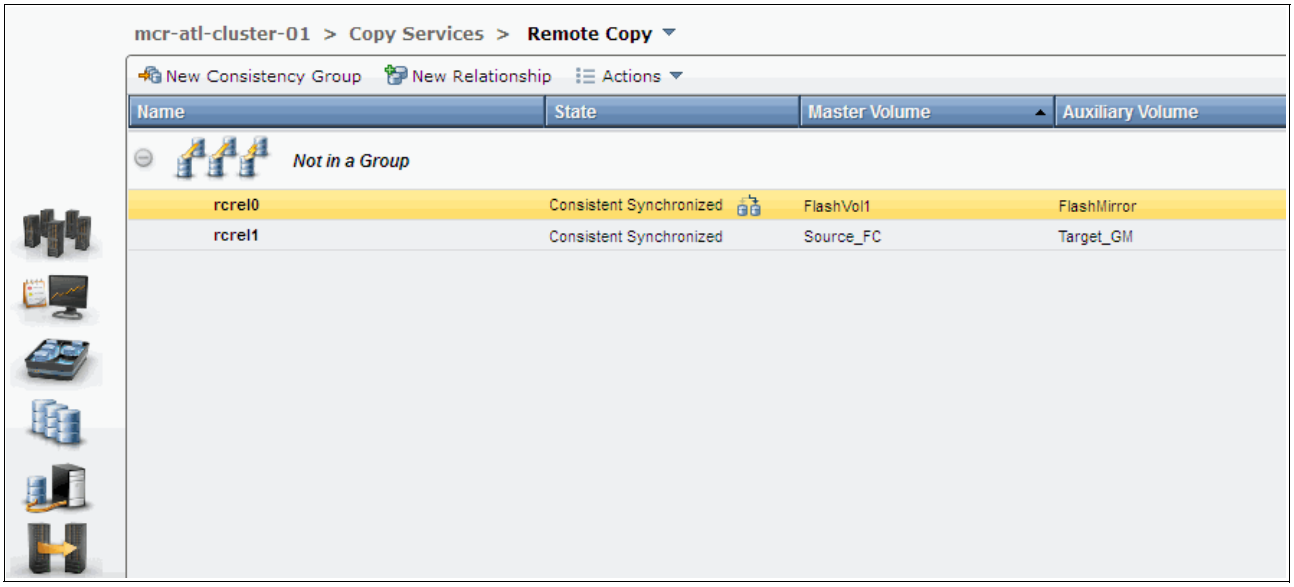

Figure 10-112 Switch icon on the state of the relationship

Renaming a stand-alone Remote Copy relationship

The Remote Copy relationship can be renamed by selecting the relationship and clicking **Rename** from the Actions drop-down menu, as shown in Figure 10-113 on page 533.

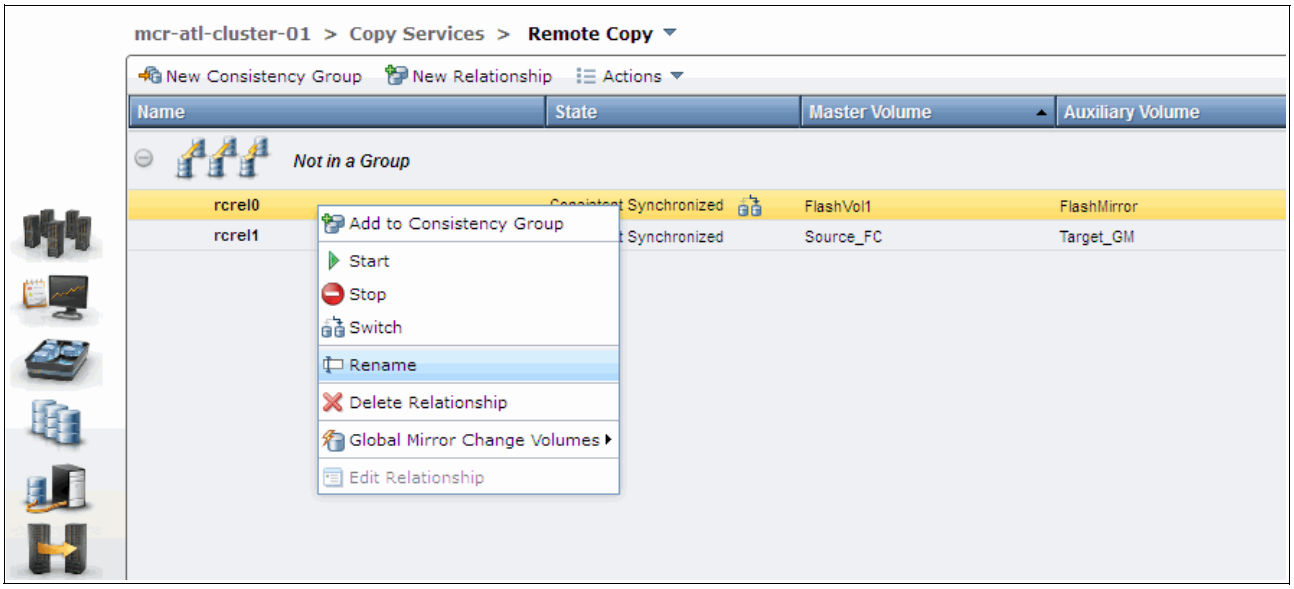

Figure 10-113 Rename the Remote Copy relationship

Enter the new name for the Remote Copy relationship and click **Rename**.

Deleting a stand-alone Remote Copy relationship

The Remote Copy relationship can be deleted by selecting the relationship and clicking **Delete Relationship** from the Actions drop-down menu, as shown in Figure 10-114.

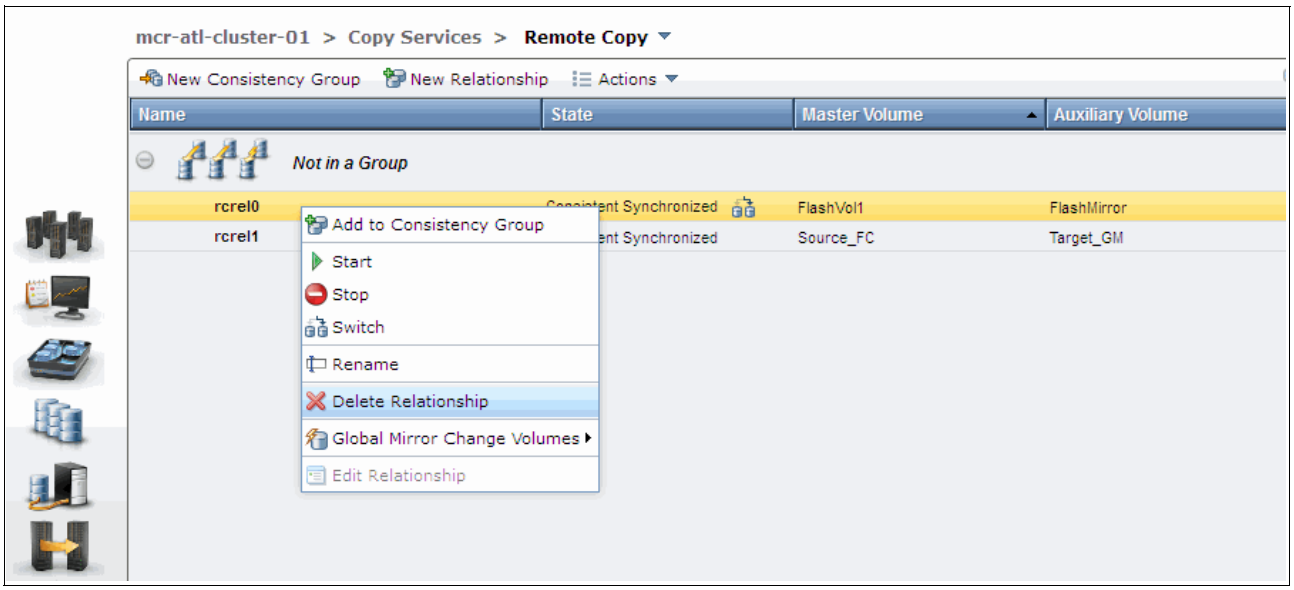

Figure 10-114 Delete a Remote Copy relationship

You must confirm this deletion by verifying the number of relationships to be deleted, as shown in Figure 10-115 on page 534. Click **Delete** to proceed.

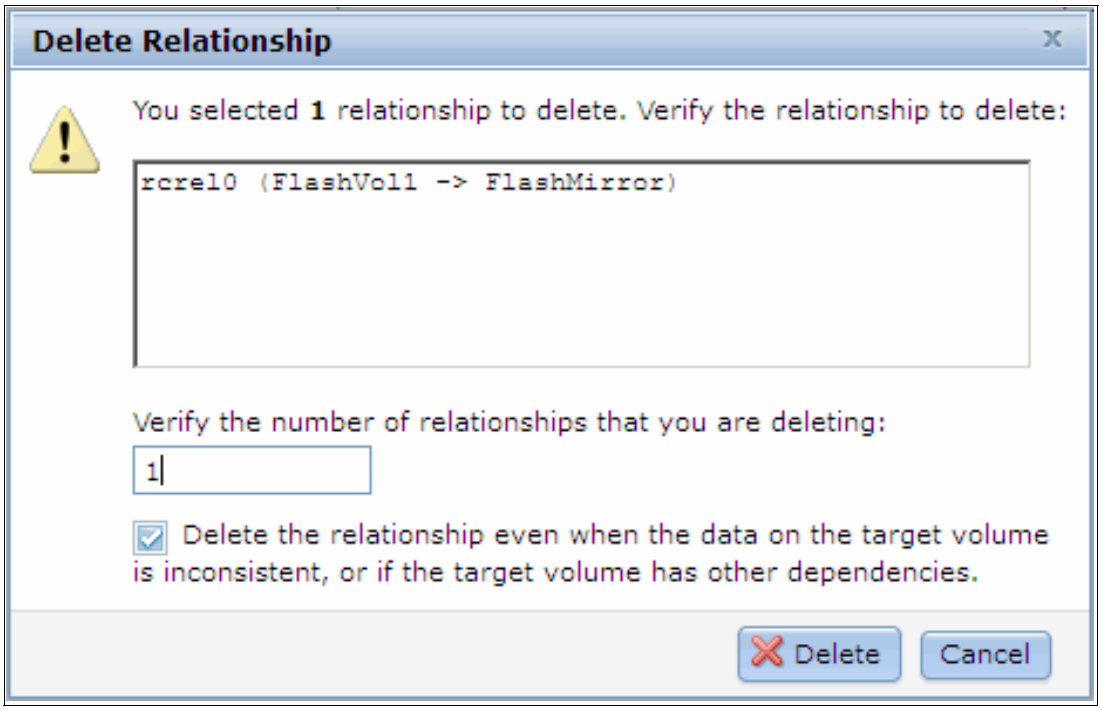

Figure 10-115 Confirm the relationship deletion

10.4.3 Managing a Remote Copy consistency group

A Remote Copy consistency group can be managed from the Remote Copy window as well.

Creating a Remote Copy consistency group

To create a Remote Copy consistency group, click **New Consistency Group**, as shown in Figure 10-116.

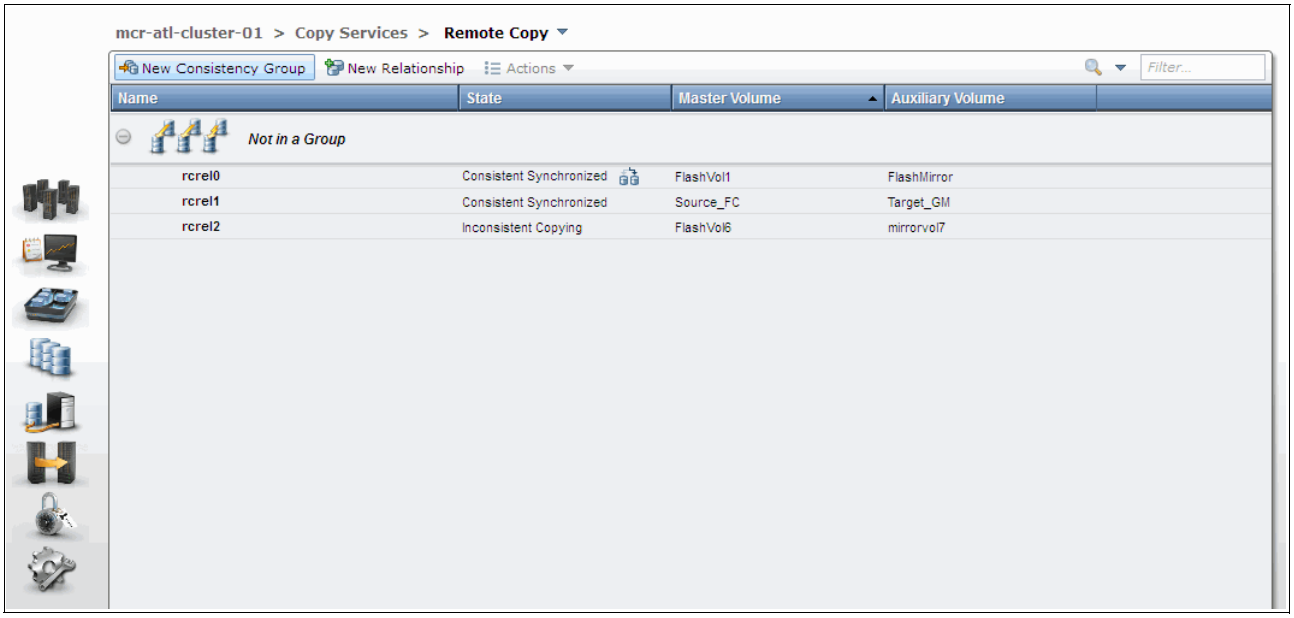

Figure 10-116 Create a Remote Copy consistency group

You must enter a name for your new consistency group, as shown in Figure 10-117.

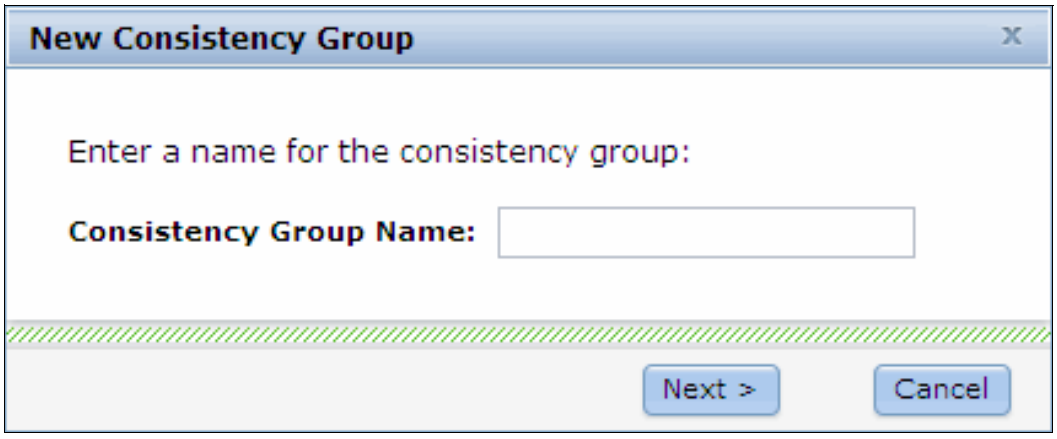

Figure 10-117 Enter a name for the new consistency group

You are prompted for the location of auxiliary volumes, as shown in Figure 10-118. In our example, these volumes are on another system. Select the relevant options and from the drop-down menu, select the correct remote system. In our example, we have only one remote system defined. Click **Next** to continue.

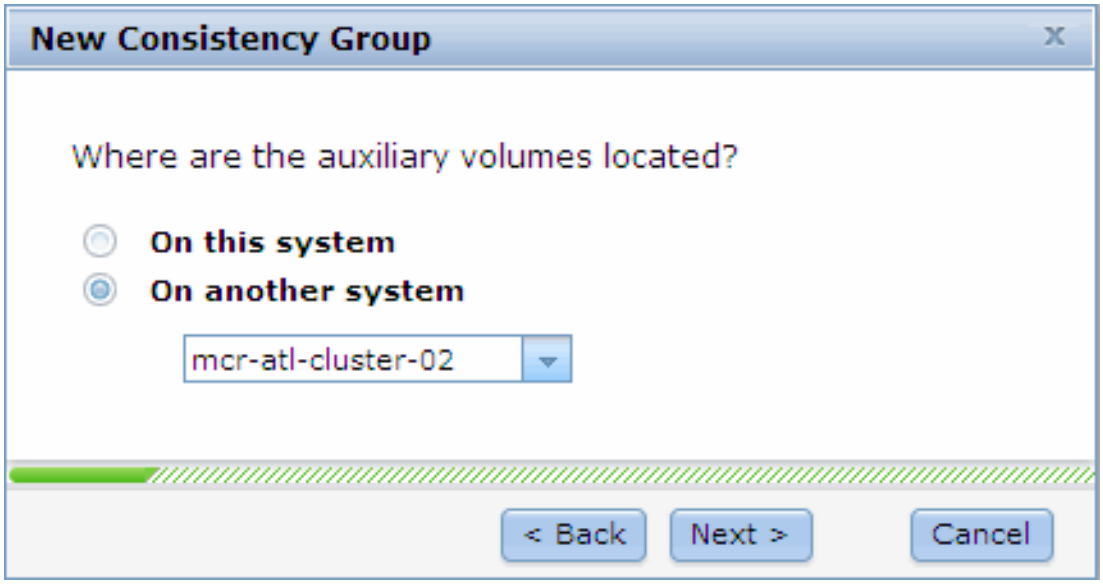

Figure 10-118 Remote Copy consistency group auxiliary volume location window

You are then prompted to create an empty consistency group or add relationships to it, as shown in Figure 10-119.

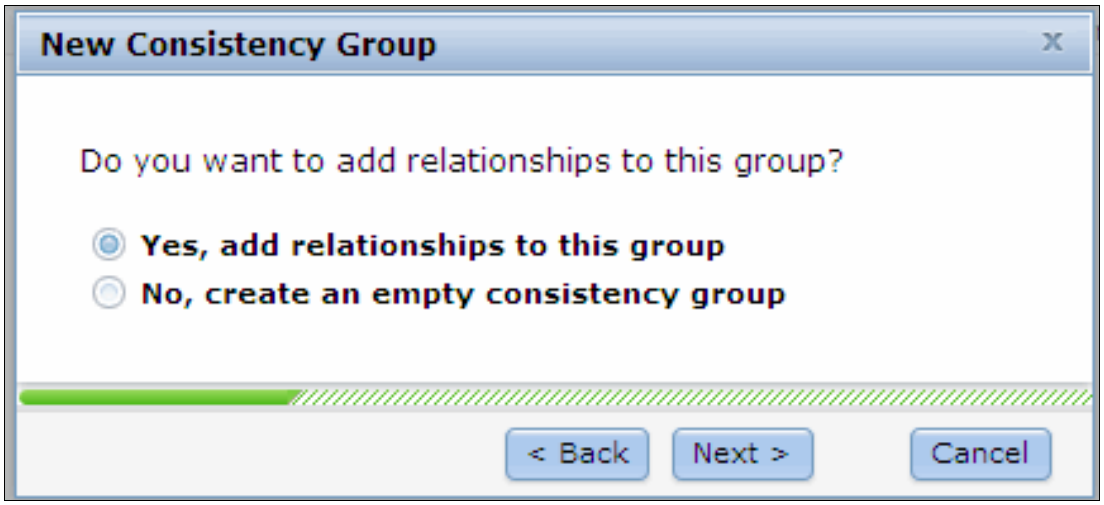

Figure 10-119 Creating an empty consistency group

If you select **No** and click **Finish**, the wizard completes and creates an empty Remote Copy Consistency Group. Selecting **Yes** prompts for the type of copy to create, as shown in Figure 10-120.

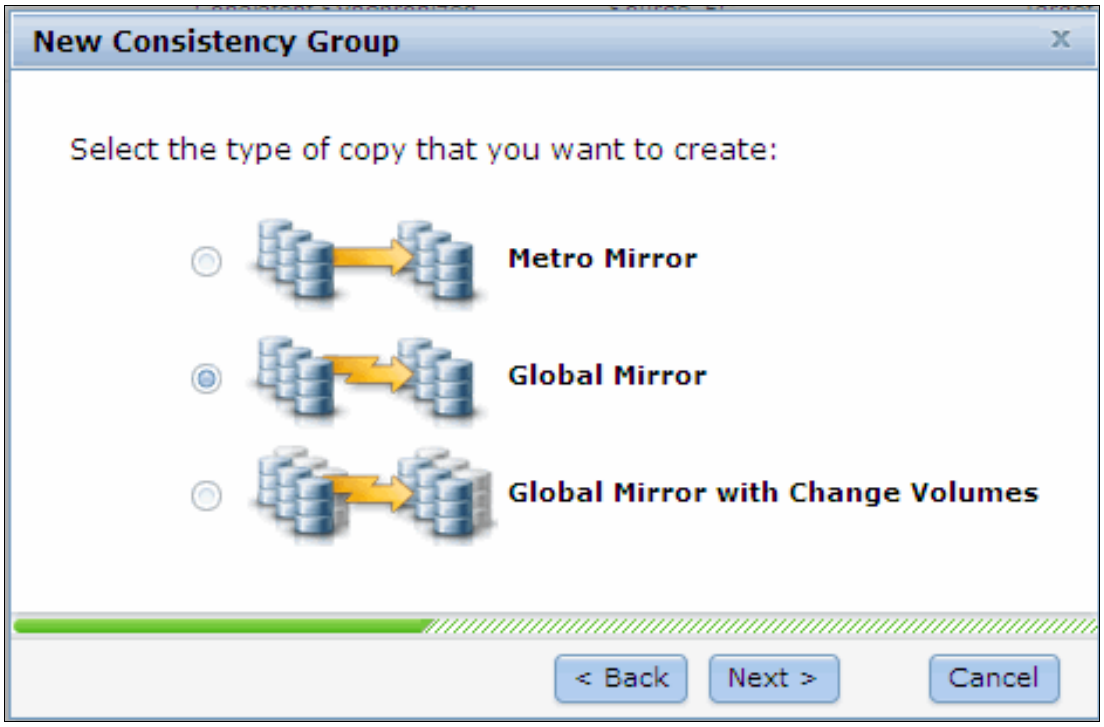

Figure 10-120 Remote Copy consistency group copy type

Choose the relevant copy type and click **Next.** In the following window, you can choose existing relationships to add the new consistency group. This step is optional. Use the Ctrl and Shift keys to select multiple relationships to add. If you decide that you do not want to use any of these relationships but you do want to create other relationships, click **Next**.

However, if you already highlighted a relationship and then decide you do not want any of these relationships, you cannot remove the relationship. You must stop the wizard and start again, as shown in Figure 10-121.

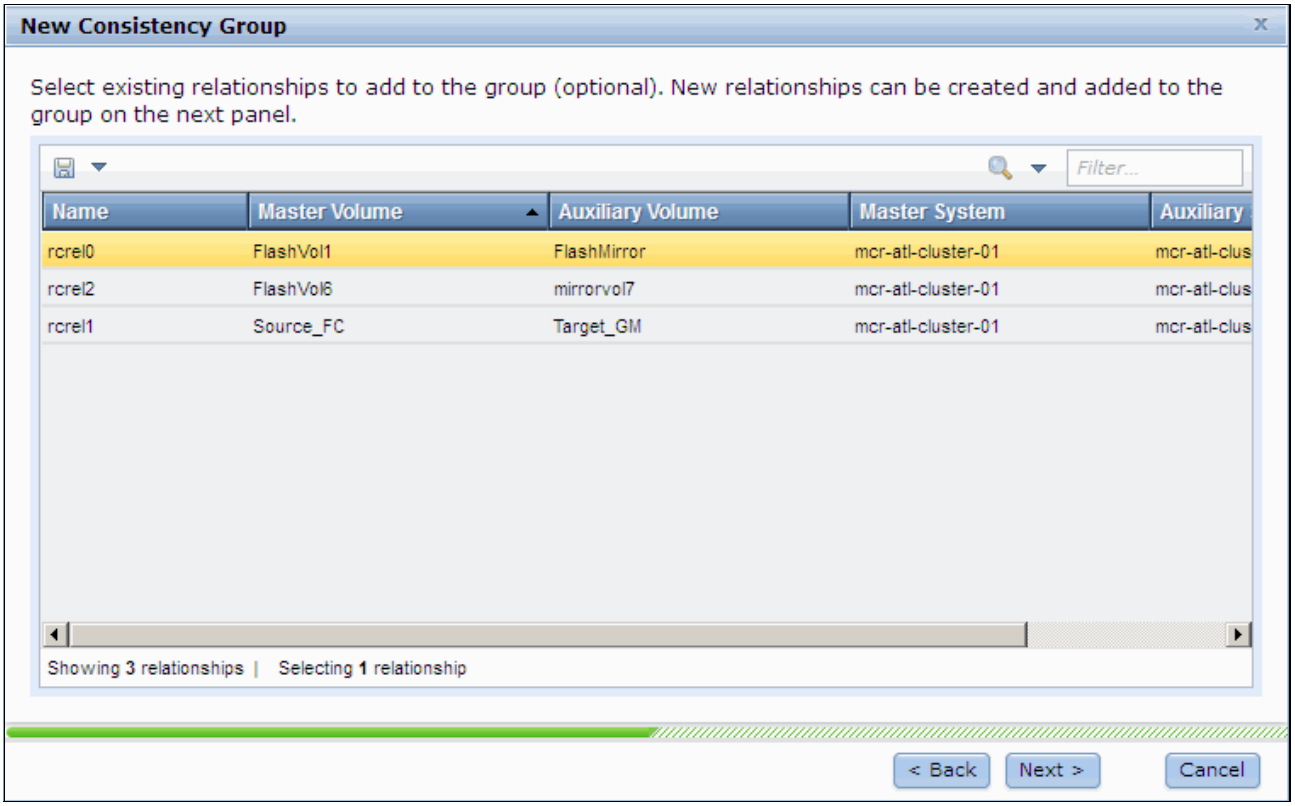

Figure 10-121 Selecting existing relationships

The next window is optional and gives the option to create relationships to add to the consistency group, as shown in Figure 10-122.

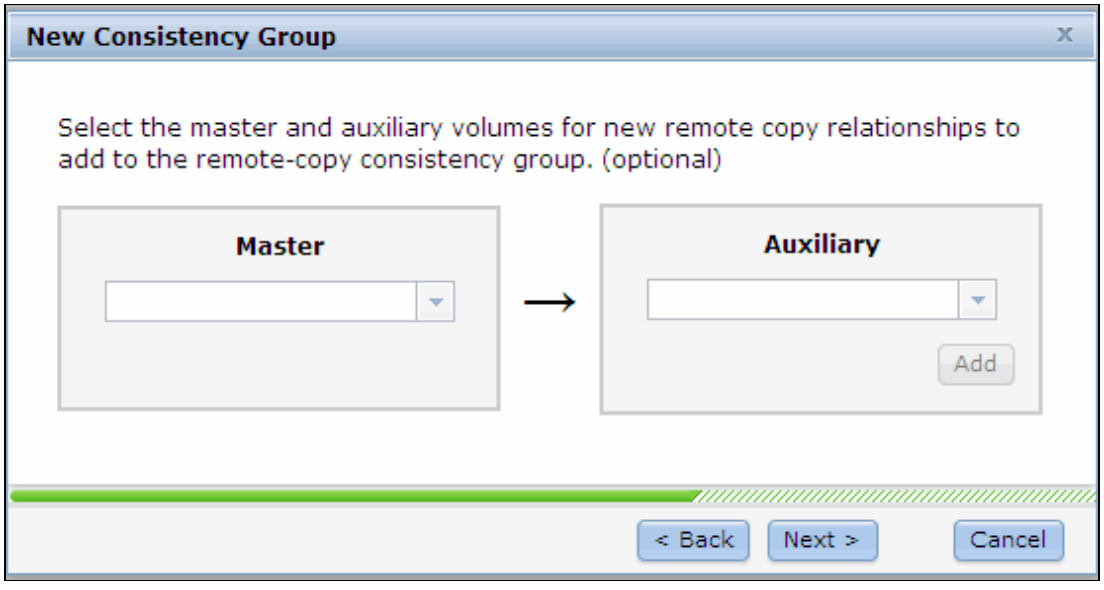

Figure 10-122 Creating relationships for Remote Copy consistency group

Select the relevant Master and Auxiliary volumes for the relationship you want to create and click **Add.** Multiple relationships can be defined by selecting another Master and Auxiliary volume and clicking **Add** again. When you finish, click **Next**. The next window prompts for whether the relationships are synchronized, as shown in Figure 10-123.

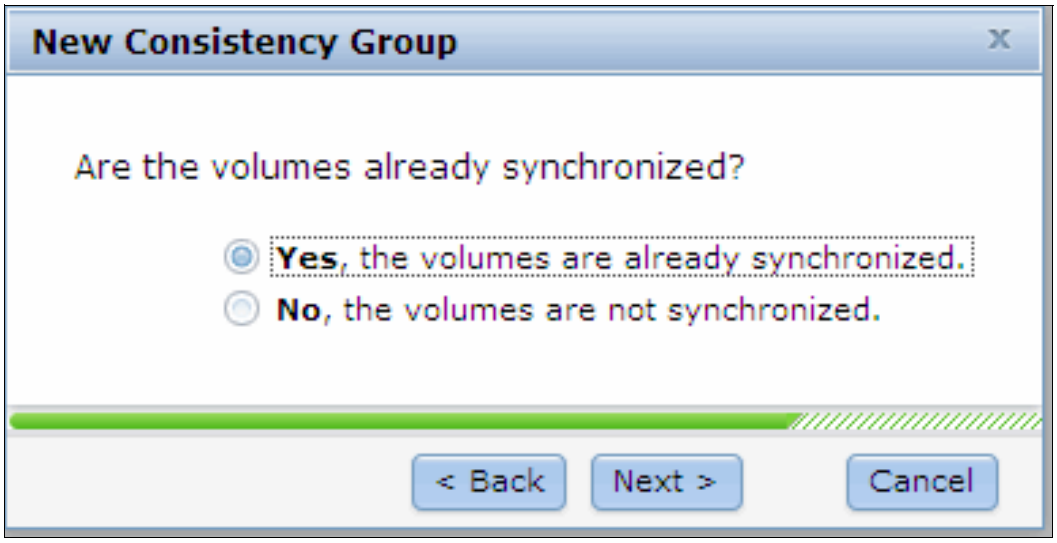

Figure 10-123 Volume synchronization

In the next window, you are asked whether you want to start copying the volumes now, as shown in Figure 10-124.

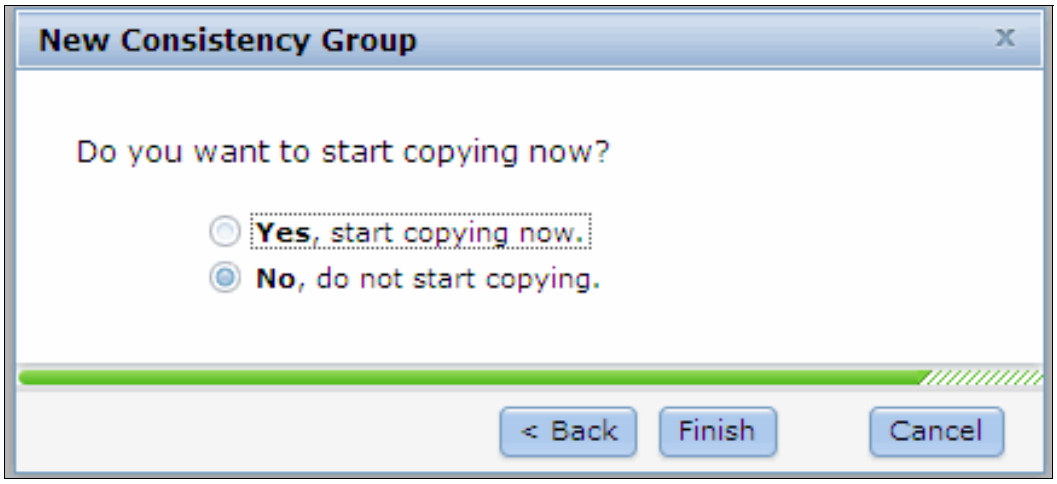

Figure 10-124 Remote Consistency group start copying option

After you select this option, click **Finish** to create the Remote Copy Consistency Group. Click **Close** to close the task window and the new consistency group is now shown in the GUI, as shown in Figure 10-125 on page 539.

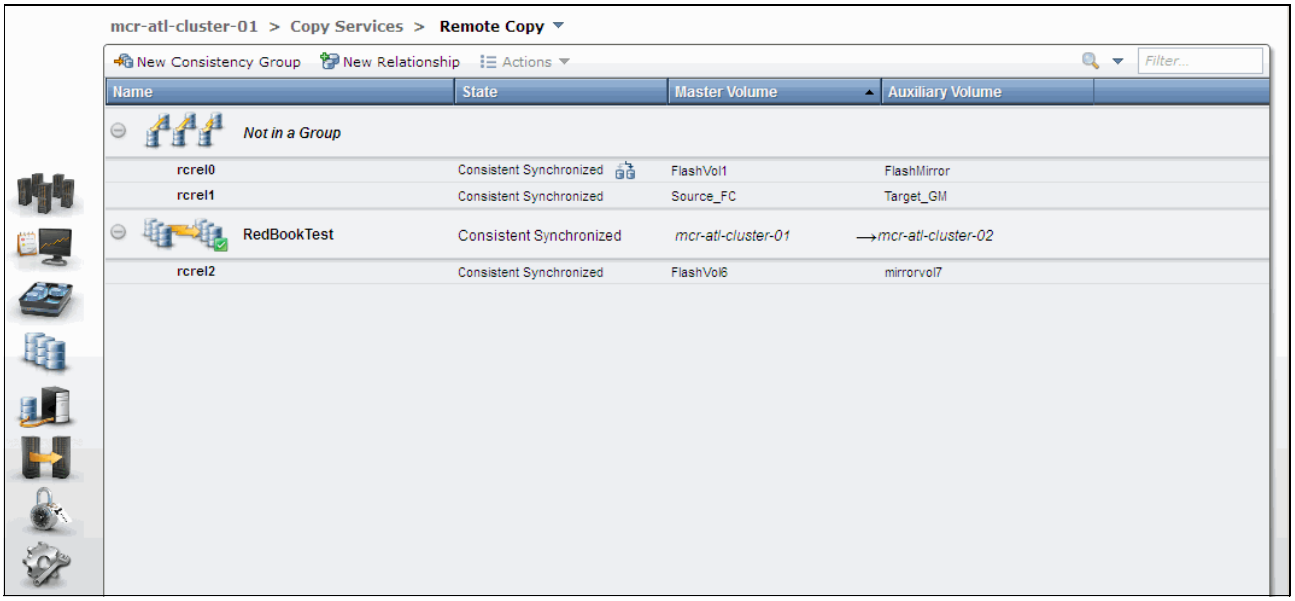

Figure 10-125 New Remote Consistency group created

In our example, we created a consistency group with a single relationship. Other Remote Copy relationships are added to the consistency group later.

You can find the name and the status of the consistency group beside the Relationship function icon. It is easy to change the name of consistency group by right-clicking the name, selecting **Rename** and then entering a new name. Alternatively, highlight the consistency group and select **Rename** from the Actions drop-down menu. Similarly, below the Relationship function icon is the Remote Copy relationships in this consistency group. The actions on the Remote Copy relationships can be applied here by using the Actions drop-down menu or right-clicking the relationships, as shown in Figure 10-126.

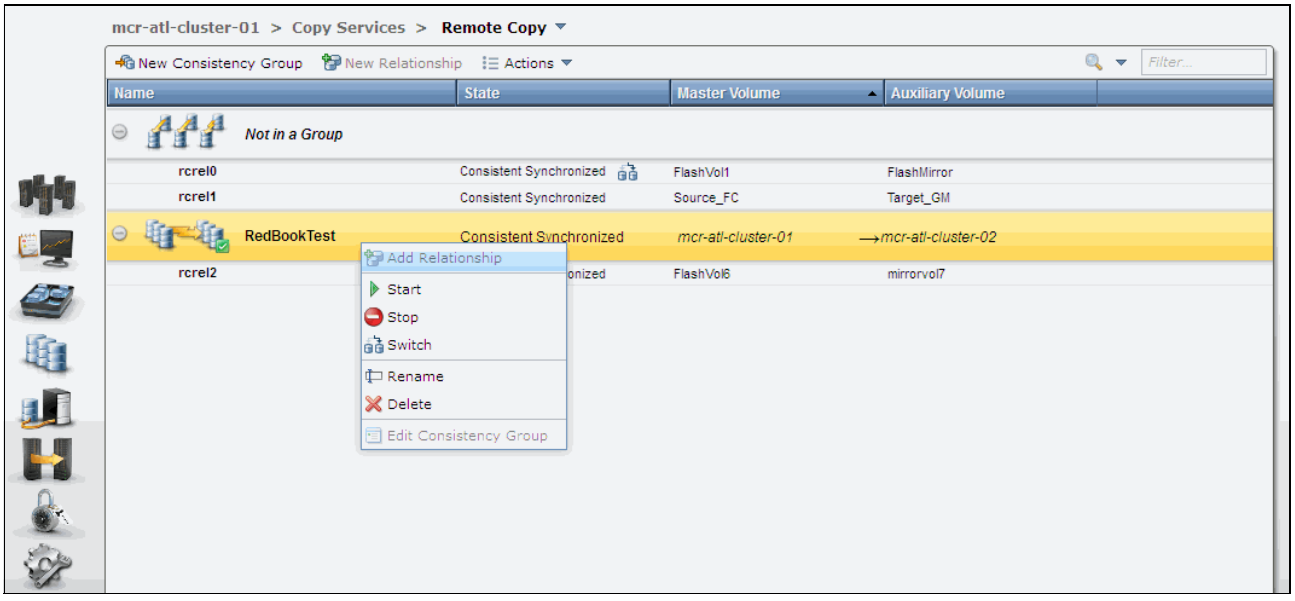

Figure 10-126 Drop-down menu options

Adding Remote Copy to a consistency group

The Remote Copy relationships in the Not in a Group list can be added to a consistency group by selecting the volumes and clicking **Add to Consistency Group** from the Actions drop-down menu, as shown in Figure 10-127.

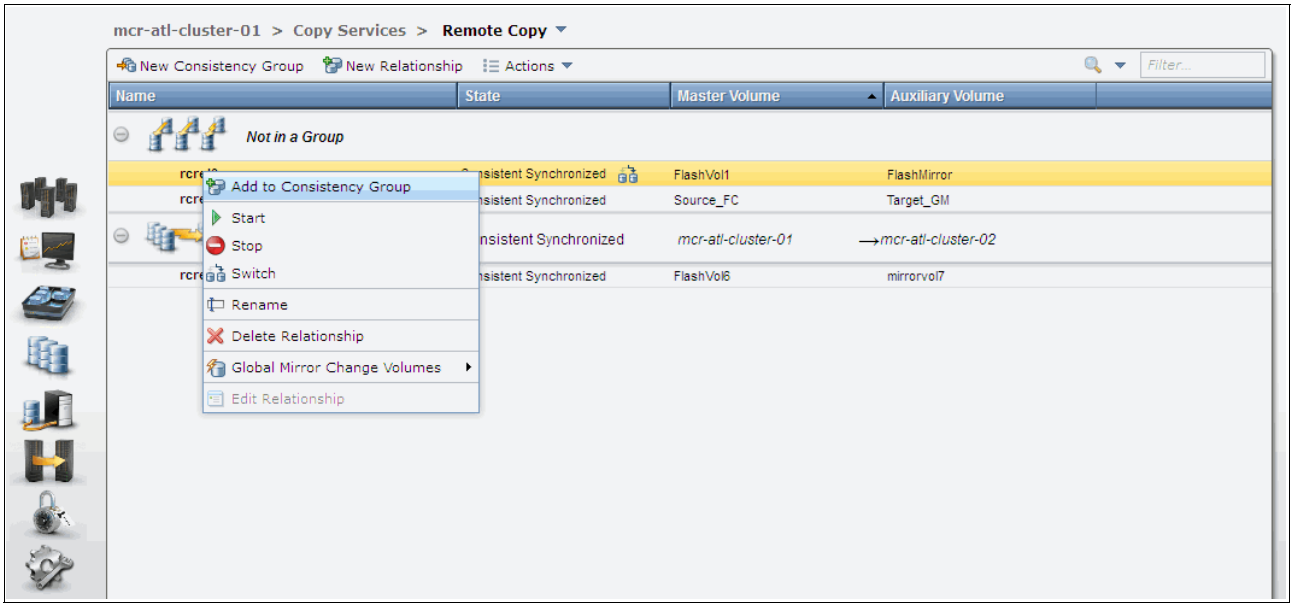

Figure 10-127 Add Remote Copy relationships to a consistency group

You must choose the consistency group to which to add the Remote Copy relationships. Based on your requirements, select the appropriate consistency group and click **Add to Consistency Group**, as shown in Figure 10-128.

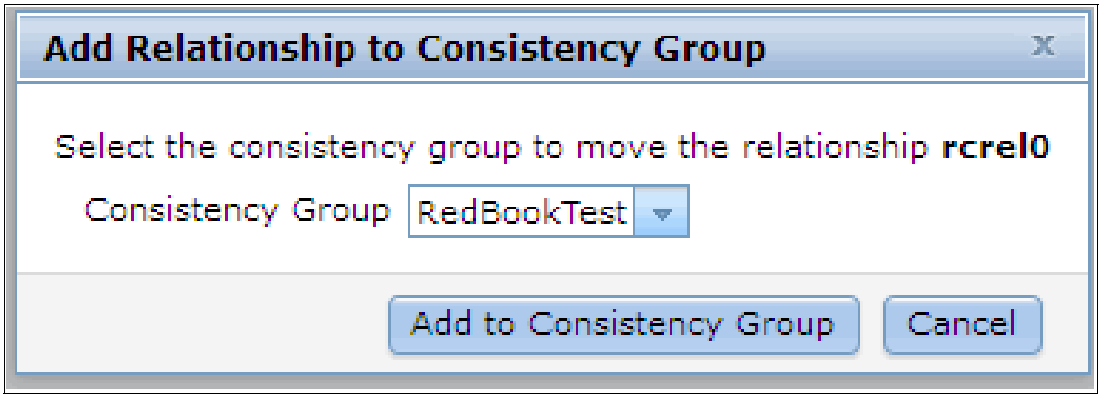

Figure 10-128 Choose the consistency group to add the remote copies

Your Remote Copy relationships are now in the consistency group that you selected.

Starting a consistency group

The Remote Copy relationship can be started by clicking **Start** from the Actions drop-down menu, as shown in Figure 10-129 on page 541.

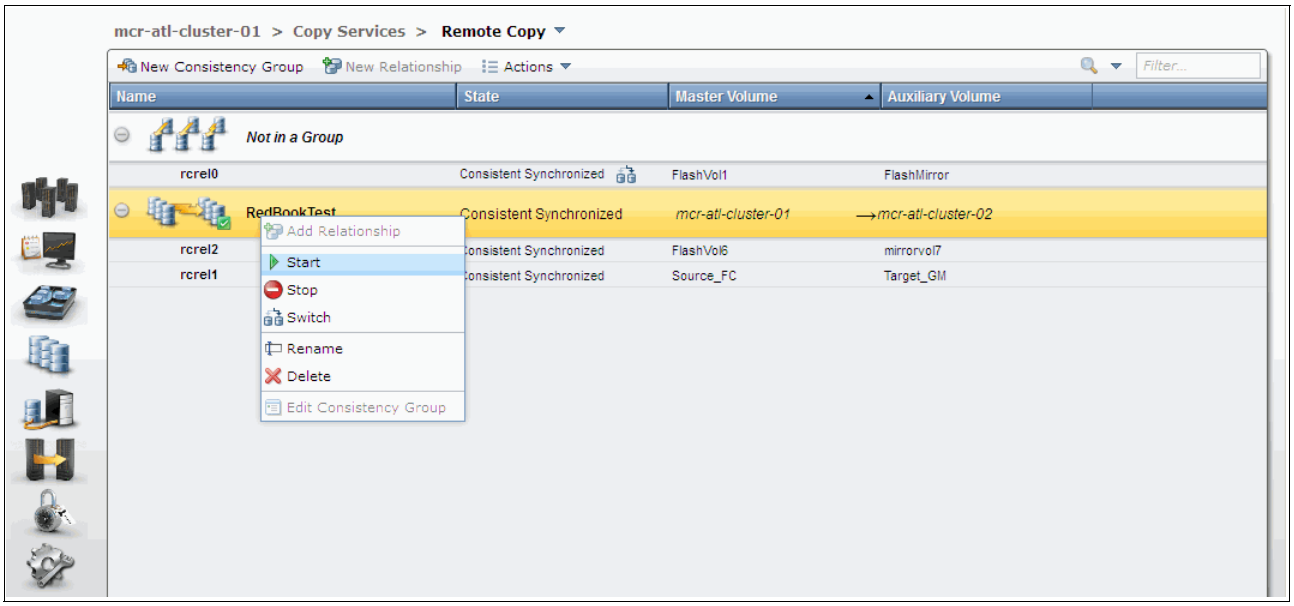

Figure 10-129 Start the consistency group

The consistency group starts copying data from the primary to the secondary.

Stopping a consistency group

The Remote Copy relationship can be stopped by clicking **Stop** in the Actions drop-down menu, as shown in Figure 10-130.

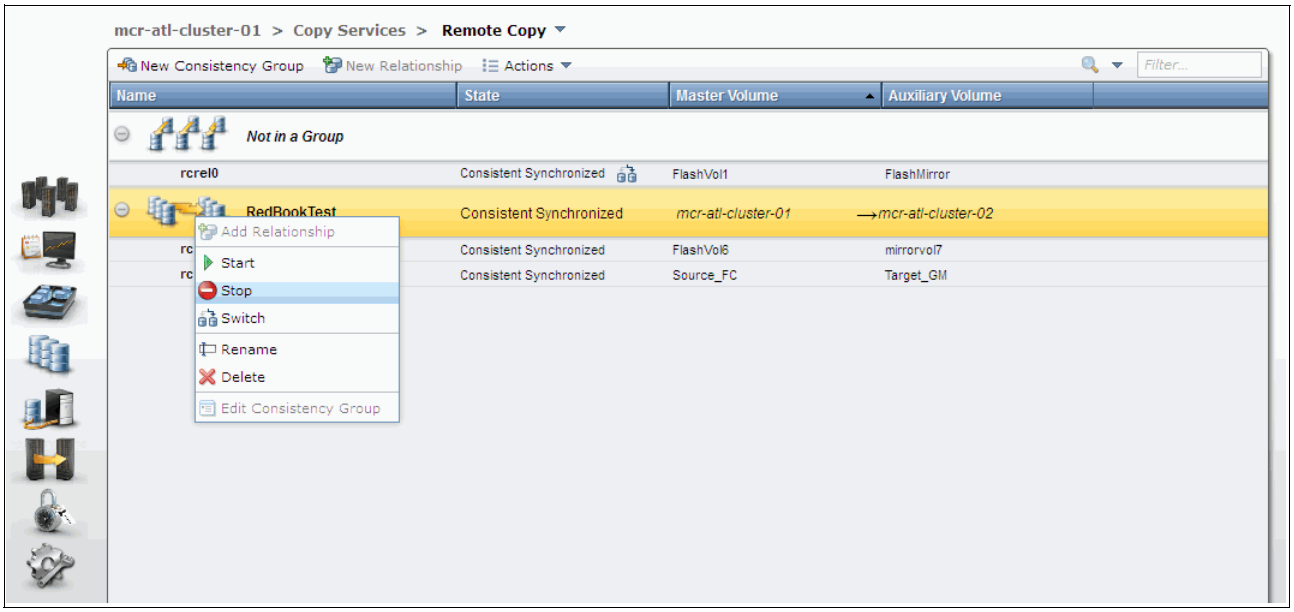

Figure 10-130 Stop the consistency group

You can allow read/write access to secondary volumes by selecting the option (as shown in Figure 10-131) and clicking **Stop Consistency Group**.

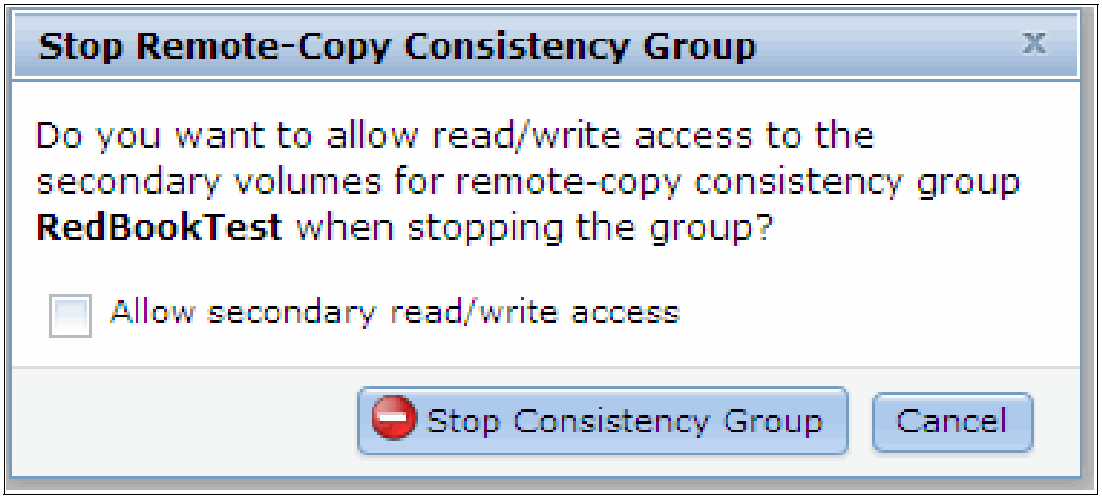

Figure 10-131 Confirm consistency group stop and allow secondary read/write access

Switching a consistency group

As with the switch action on the Remote Copy relationship, you can switch the copy direction of the consistency group. To switch the copy direction of the consistency group, click **Switch** from the Actions drop-down menu, as shown in Figure 10-132.

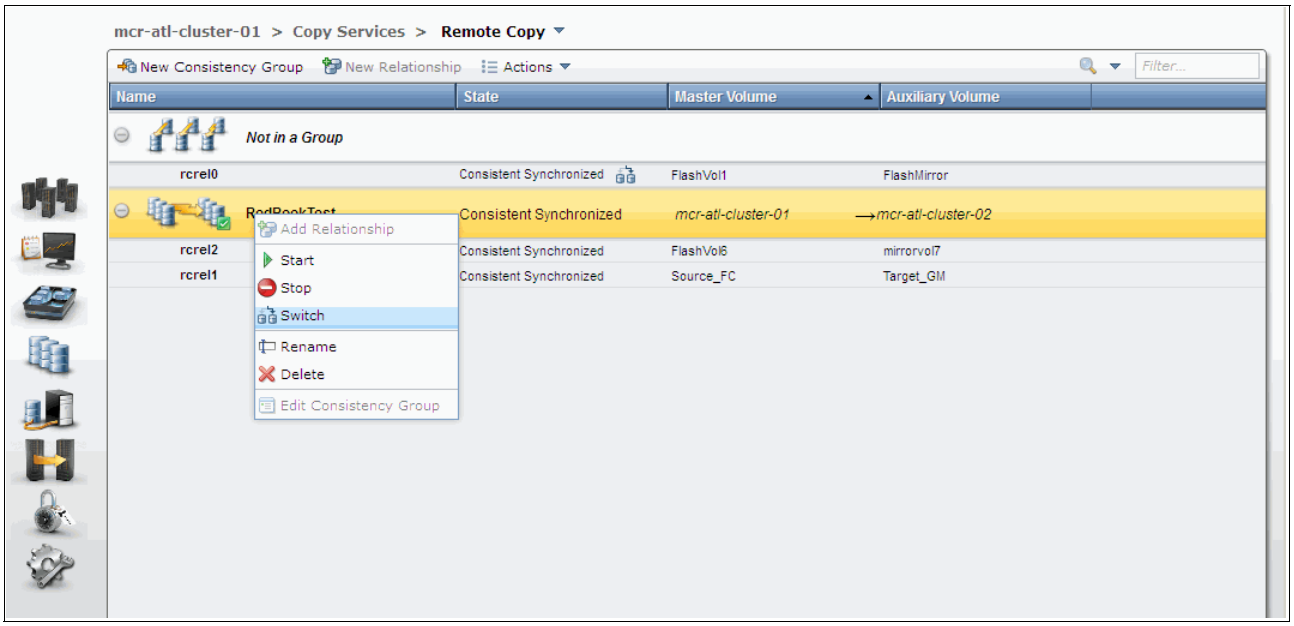

Figure 10-132 Switch the copy direction of a consistency group

A warning message opens, as shown in Figure 10-133. After the switch, the primary cluster in the consistency group changes. Write access to current master volumes is lost, while write access to the current auxiliary volumes is enabled. This change affects host access, so make sure that these settings are what you need, and if so, click **OK** to continue.

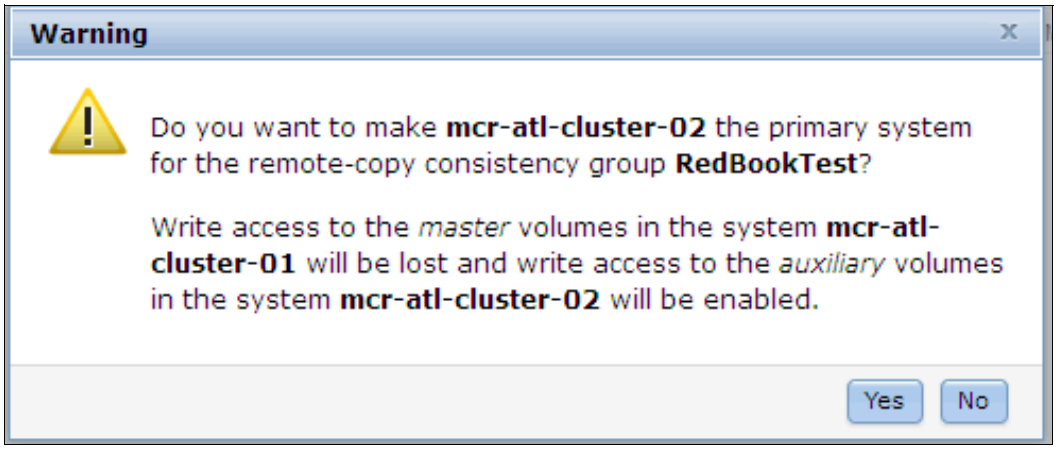

Figure 10-133 Warning message to confirm the switch

Removing Remote Copy relationships from a consistency group

The Remote Copy relationships can be removed from the consistency group by selecting the Remote Copy relationships and clicking **Remove from Consistency Group** from the Actions drop-down menu, as shown in Figure 10-134.

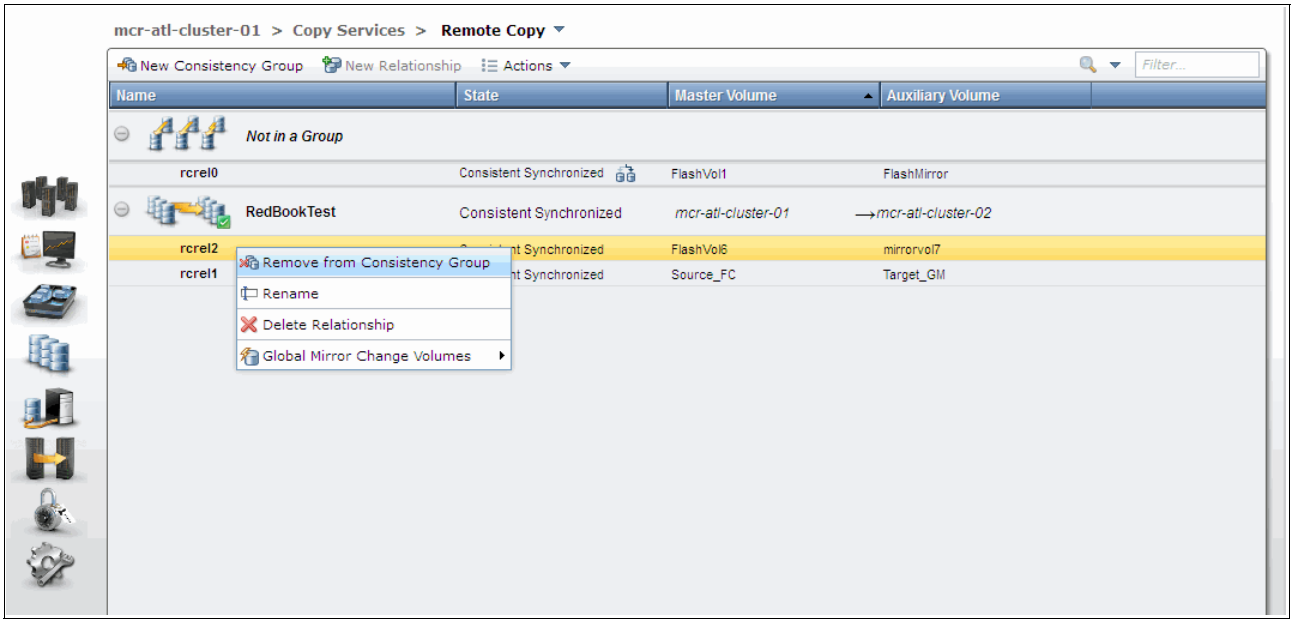

Figure 10-134 Remove Remote Copy relationships from a consistency group

You are prompted to confirm the Remote Copy relationships you want to delete from the consistency group, as shown in Figure 10-135. Make sure the Remote Copy relationships that are shown in the field are the ones that you must remove from the consistency group. Click **Remove** to proceed.

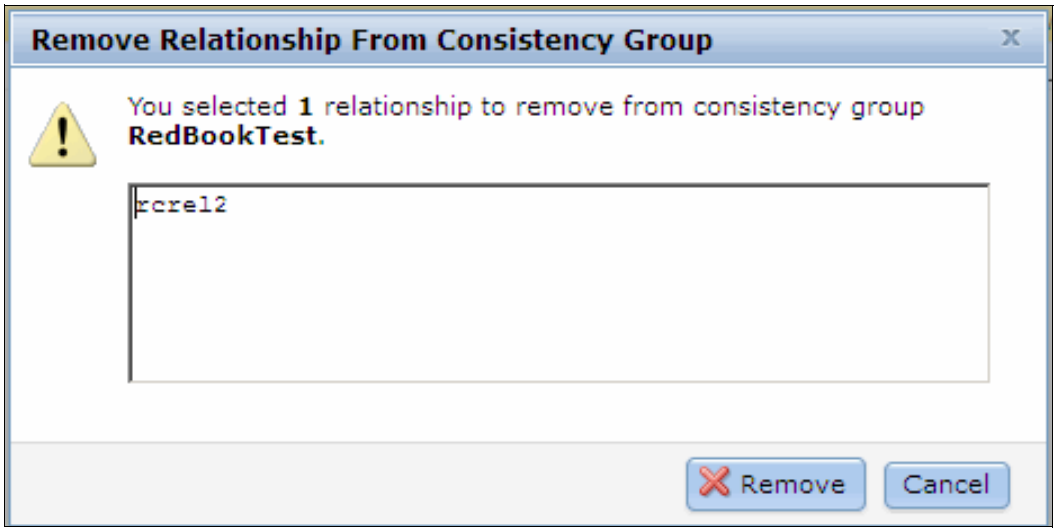

Figure 10-135 Confirm the relationships to remove from the Remote Copy consistency group

After the removal process completes, the Remote Copy relationships are deleted from the consistency group and displayed in the Not in a Group list.

Deleting a consistency group

The consistency group can be deleted by selecting **Delete** from the Actions drop-down menu, as shown in Figure 10-136.

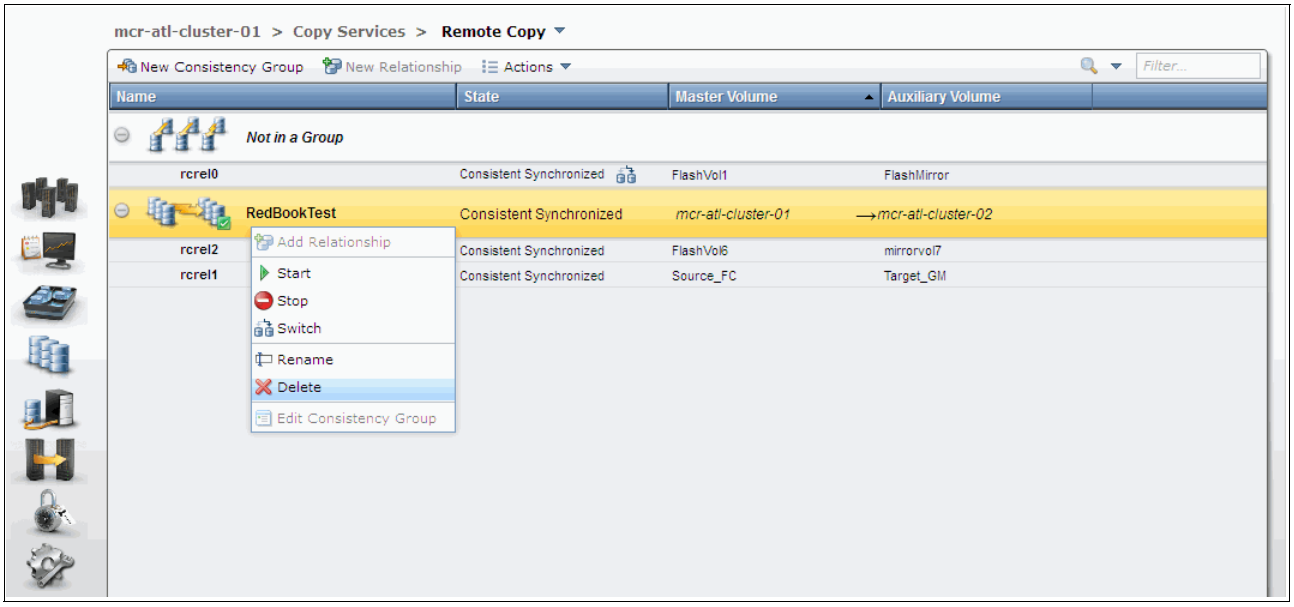

Figure 10-136 Delete a consistency group

You must confirm the deletion of the consistency group, as shown in Figure 10-137. Click **OK** if you are sure that this consistency group should be deleted.

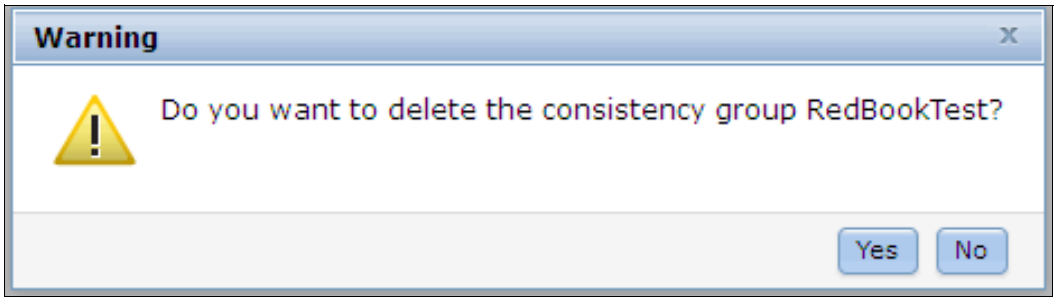

Figure 10-137 Warning to confirm deletion of the consistency group

The consistency group is deleted. Any relationships that were part of the consistency group are returned to the Not in a Group list.

11

External storage virtualization

In this chapter, we describe how to incorporate external storage systems into the virtualized world of the IBM Storwize V5000. A key feature of IBM Storwize V5000 is its ability to consolidate disk controllers from various vendors into pools of storage. In this way, the storage administrator from a single user interface can manage and provision storage to applications, and use a common set of advanced functions across all the storage systems under the control of the IBM Storwize V5000.

This chapter includes the following topics:

- Planning for external storage virtualization
- Working with external storage

11.1 Planning for external storage virtualization

In this section, we describe how to plan for virtualizing external storage with IBM Storwize V5000. Virtualizing the storage infrastructure with IBM Storwize V5000 makes your storage environment more flexible, cost-effective, and easy to manage. The combination of IBM Storwize V5000 and an external storage system allows more storage capacity benefits from the powerful software function within the IBM Storwize V5000.

The external storage systems that are incorporated into the IBM Storwize V5000 environment can be new systems or existing systems. The data on the existing storage systems can be easily migrated to the IBM Storwize V5000 managed environment, as described in Chapter 6, “Storage migration wizard” on page 237, and Chapter 7, “Storage pools” on page 295.

11.1.1 License for external storage virtualization

From a licensing standpoint, when external storage systems are virtualized by IBM Storwize V5000, a per-enclosure External Virtualization license is required. For more information, contact your IBM account team or IBM Business Partner for further assistance.

Migration: If the IBM Storwize V5000 is used as a general migration tool, the appropriate External Virtualization licenses must be ordered. The only exception is if you want to migrate existing data from external storage systems to IBM Storwize V5000 internal storage because you can temporarily configure your External Storage license within 45 days. For a more than 45 day migration requirement from external storage to IBM Storwize V5000 internal storage, an appropriate External Virtualization license must be ordered.

You can configure the IBM Storwize V5000 licenses by clicking the **Settings** icon and then clicking **General** → **Licensing**, as shown in Figure 11-1.

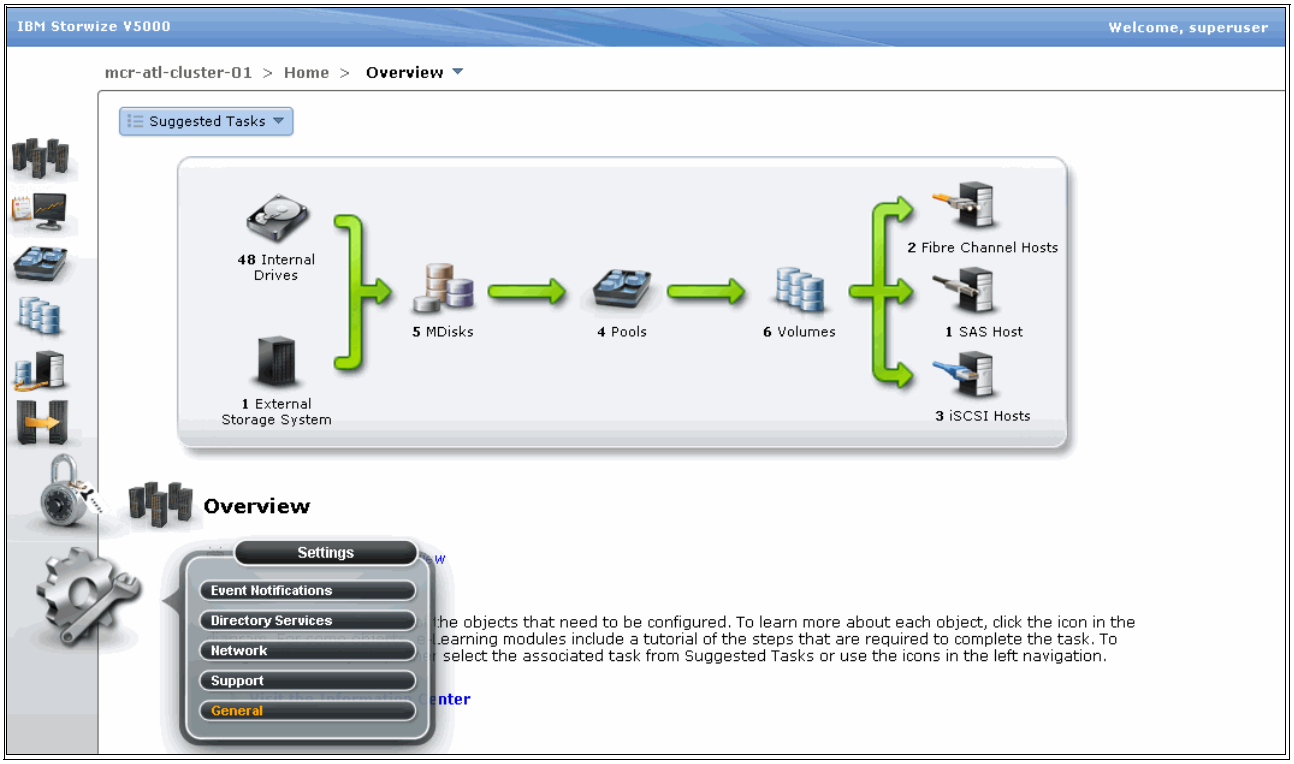

Figure 11-1 General option

In the Advanced window, click **Licensing** and the Update License view opens in the right pane, as shown in Figure 11-2.

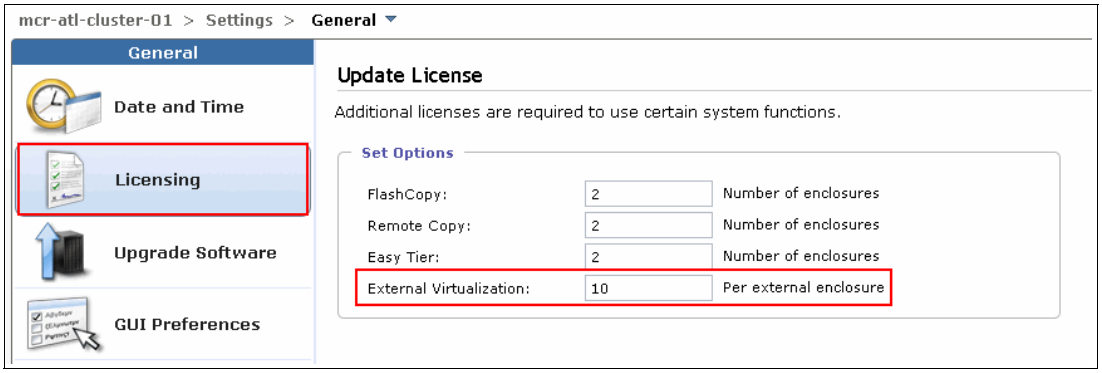

Figure 11-2 Update License window

In the Update License pane, there are two license options you can set: External Virtualization Limit and Remote-Copy Limit. Set these license options to the limit you obtained from IBM.

For assistance with licensing questions or to purchase an External Virtualization or Remote Copy license, contact your IBM account team or IBM Business Partner.

11.1.2 SAN configuration planning

External storage controllers that are virtualized by IBM Storwize V5000 must be connected through SAN switches. A direct connection between the IBM Storwize V5000 and storage controllers or hosts ports is not supported.

Make sure that the switches or directors are at the firmware levels that are supported by the IBM Storwize V5000 and that the IBM Storwize V5000 port login maximums that are listed in the restriction document are not exceeded. The configuration restrictions can be found on the Support home page, which is available at this website:

http://www-947.ibm.com/support/entry/portal/Overview

The recommended SAN configuration is composed of a minimum of two fabrics. The ports on external storage systems are virtualized by the IBM Storwize V5000 and the IBM Storwize V5000 ports and are evenly split between the two fabrics to provide redundancy if one of the fabrics goes offline.

After the IBM Storwize V5000 and external storage systems are connected to the SAN fabrics, zoning must be implemented. In each fabric, create a zone with the four IBM Storwize V5000 worldwide port names (WWPNs), two from each node canister with up to a maximum of eight WWPNs from each external storage system.

Ports: IBM Storwize V5000 supports a maximum of 16 ports or WWPNs from an external storage system that is virtualized.

Figure 11-3 shows an example of how to cable devices to the SAN. Refer to this example as we describe the zoning.

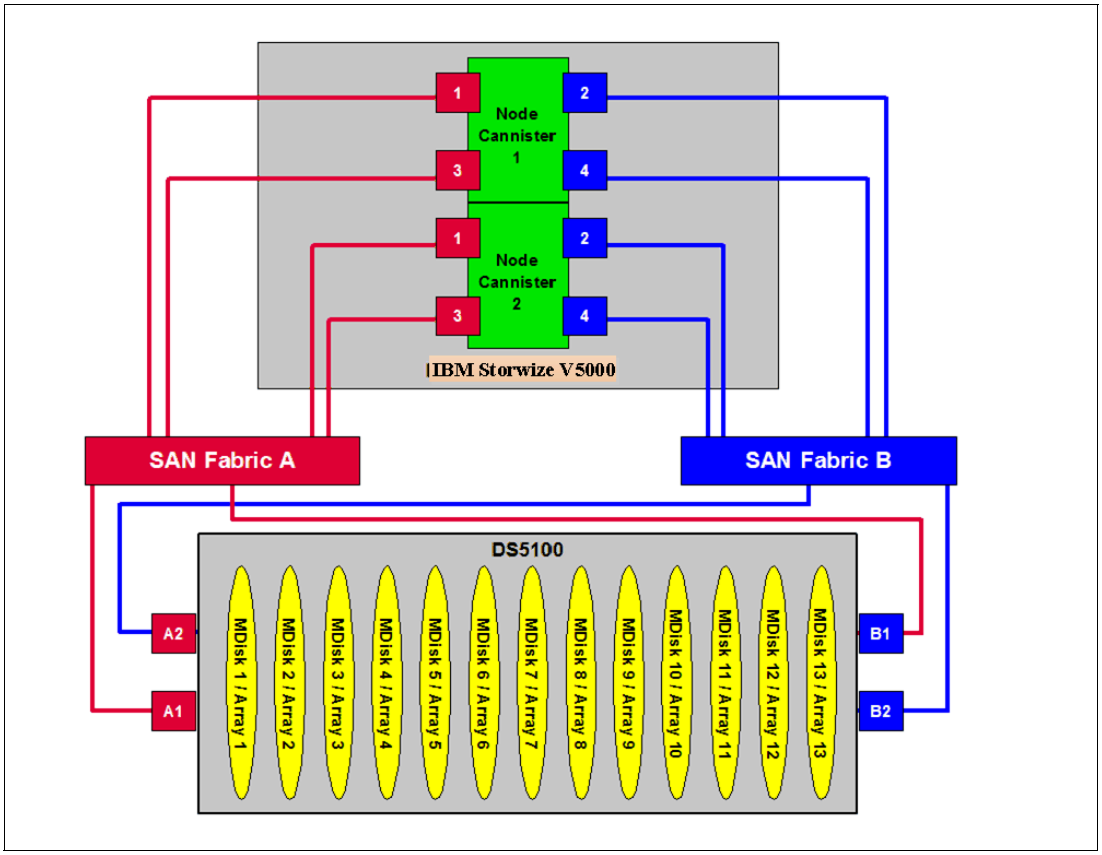

Figure 11-3 SAN cabling and zoning example diagram

Create an IBM Storwize V5000/external storage zone for each storage system to be virtualized, as shown in the following examples:

- Zone DS5100 controller ports A1 and B1 with all node ports 1 and 3 in the RED fabric
- Zone DS5100 controller ports A2 and B2 with all node ports 2 and 4 in the BLUE fabric

11.1.3 External storage configuration planning

External storage systems provide redundancy through various RAID levels, which prevents a single physical disk failure from causing an MDisk, storage pool, or associated host volume from going offline. To minimize the risk of data loss, virtualize storage systems only where logical unit numbers (LUNs) are configured by using a RAID level other than RAID 0 (for example RAID 1, RAID 10, RAID 0+1, RAID 5, or RAID 6).

Verify that the storage controllers to be virtualized by IBM Storwize V5000 meet the requirements. The configuration restrictions can be found on the Support home page, which is available at this website:

`http://www-947.ibm.com/support/entry/portal/Overview`

Make sure that the firmware or microcode levels of the storage controllers to be virtualized are supported by IBM Storwize V5000.

IBM Storwize V5000 must have exclusive access to the LUNs from the external storage system that are mapped to it. LUNs cannot be shared between IBM Storwize V5000s or between an IBM Storwize V5000 and other storage virtualization platforms or between an IBM Storwize V5000 and hosts. However, different LUNs can be mapped from one external storage system to an IBM Storwize V5000 and other hosts in the SAN through different storage ports.

Make sure to configure the storage subsystem LUN masking settings to map all LUNs to all the WWPNs in the IBM Storwize V5000 storage system.

Ensure that you see the IBM Storwize V5000 Information Center and review the “Configuring and servicing external storage system” topic before you prepare the external storage systems for discovery by the IBM Storwize V5000 system. This Information Center can be found at this website:

http://publib.boulder.ibm.com/infocenter/storwize/ic/index.jsp

11.1.4 Guidelines for virtualizing external storage

When external storage is virtualized by using the IBM Storwize V5000, the following guidelines must be followed:

- Avoid splitting arrays into multiple LUNs at the external storage system level. When possible, create a single LUN per array for mapping to the IBM Storwize V5000.
- Except for Easy Tier, do not mix MDisks that vary in performance or reliability in the same storage pool. Always put similarly sized MDisks into one storage pool. For more information about Easy Tier, see Chapter 9, “Easy Tier” on page 411.
- Do not leave volumes in image mode. Use image mode only 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.
- The use of the copy services in Storwize V5000 gives you a unified method to manage data integrity across heterogeneous storage systems.
- The Easy Tier function is included with the IBM Storwize V5000 system. The external storage system can benefit from this powerful storage tiering function to remove hot spots and improve overall performance.

11.2 Working with external storage

In this section, we describe how to manage external storage by using an IBM Storwize V5000.

The basic concepts of managing external storage system are the same as internal storage. IBM Storwize V5000 discovers LUNs from the external storage system as one or more MDisks. These MDisks are added to a storage pool in which volumes are created and mapped to hosts, as needed.

11.2.1 Adding external storage

To add new external storage systems to the IBM Storwize V5000 virtualized environment, complete the following steps:

1. Zone a minimum of two and a maximum of 16 Fibre Channel ports from the external storage system with all eight Fibre Channel ports on the IBM Storwize V5000 system. As a best practice, have two fabrics for redundancy in the SAN. Then, in each fabric, zone two ports from each node canister in the IBM Storwize V5000 system with half the ports from the external system. As the IBM Storwize V5000 is virtualizing your storage, hosts should be zoned with the V5000 controllers WWPNs.
2. By using the storage partitioning or LUN masking feature of the external storage system, create a group that includes all eight IBM Storwize V5000 WWPNs.
3. Create equal size arrays on the external system by using any RAID level except zero.
4. Create a single LUN per RAID array.
5. Map the LUNs to all eight Fibre Channel ports on the IBM Storwize V5000 system by assigning them to the group that was created in step 2.
6. Verify that IBM Storwize V5000 discovered the LUNs as unmanaged MDisks. If they do not show up automatically, click **Detect MDisk** from the MDisk window of the GUI, as described in Chapter 7, “Storage pools” on page 295. You should see the MDisks mapped to the IBM Storwize V5000 under the respective Storage system.
7. Select the storage tier for the MDisks.
8. Create a storage pool.
9. Add the MDisks to the pool.
10. Create volumes and map them to hosts, as needed.

If the external storage systems are not new systems (that is, there is existing data on the LUNs that must be kept after virtualization), complete the steps that are described in Chapter 6, “Storage migration wizard” on page 237 to prepare the environment. You can then migrate the existing data with or without the use of the wizard to IBM Storwize V5000 internal storage or some other external storage system.

Chapter 6, “Storage migration wizard” on page 237 shows how to manually import MDisks and migrate the data to other storage pools. Whether you migrate the data with the wizard, you can select your destination storage pools to be internal storage pools or external storage pools.

11.2.2 Managing external storage

The IBM Storwize V5000 provides an individual external window for managing external storage systems.

You can access the external window by opening the Getting Started window and clicking the External Storage System function icon. Extended help information for external storage appears. Click **Physical Storage** and the external window opens.

Figure 11-4 shows how to access the External Storage window from the Getting Started window.

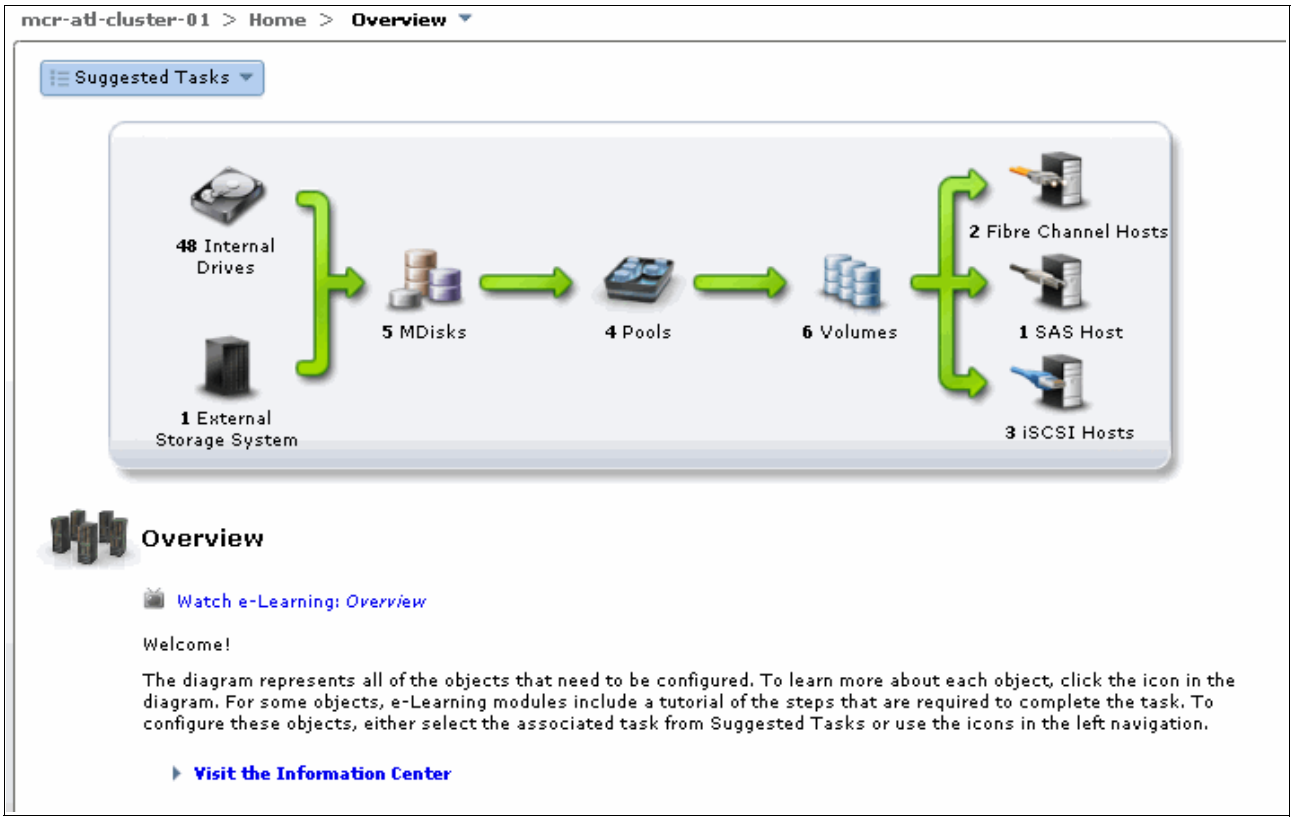

Figure 11-4 Access the External Storage window from the Getting Started window

The other method to access the external window is to use the Physical Storage function icons that are shown in the left pane, as shown in Figure 11-5.

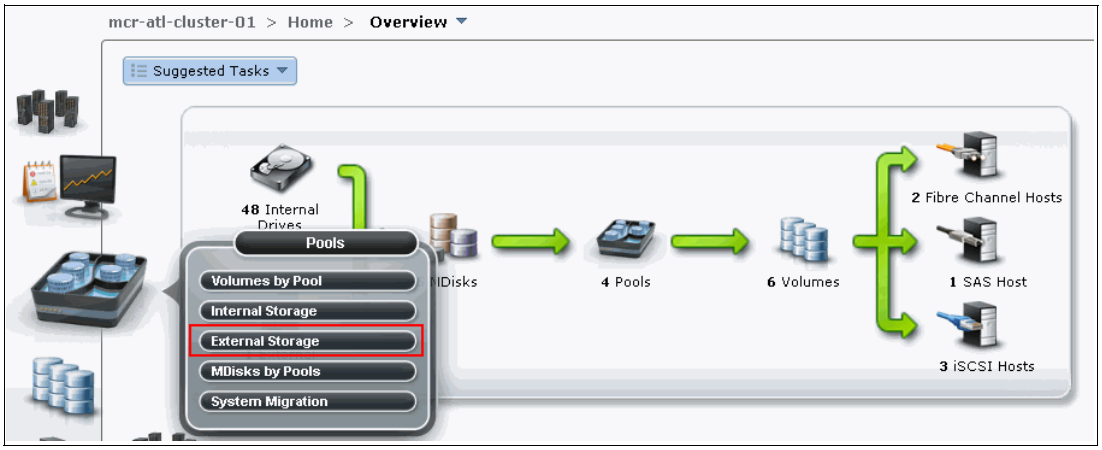

Figure 11-5 Access the External window from the Getting Started window

The External window (see Figure 11-6) gives you an overview of all your external storage systems. There is a list of the external storage systems on the left side of the window. With the help of the filter, you can show only the external storage systems on which you must act. If you click and highlight the external storage system, detailed information is shown in the right pane, including all the MDisks that are provided by it.

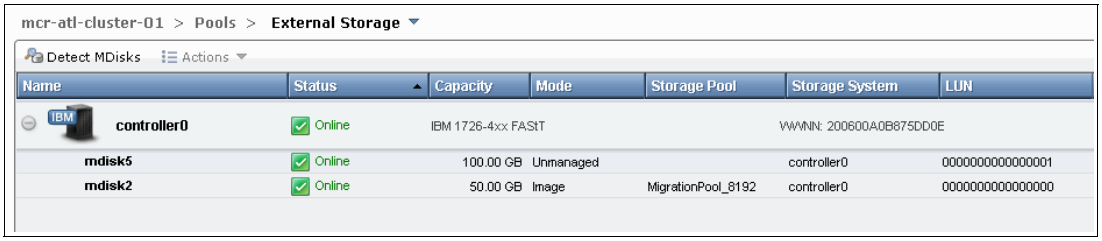

Figure 11-6 External Storage window

On the right side of the window, you can change the name of external storage system by clicking the name beside the picture of the external storage box. The status of the external storage system and its WWPN can also be found under the name.

From the Actions drop-down list (which is found at the top of the name of external storage on the right part of the External window), you can find the Show Dependent Volumes option, as shown in Figure 11-7.

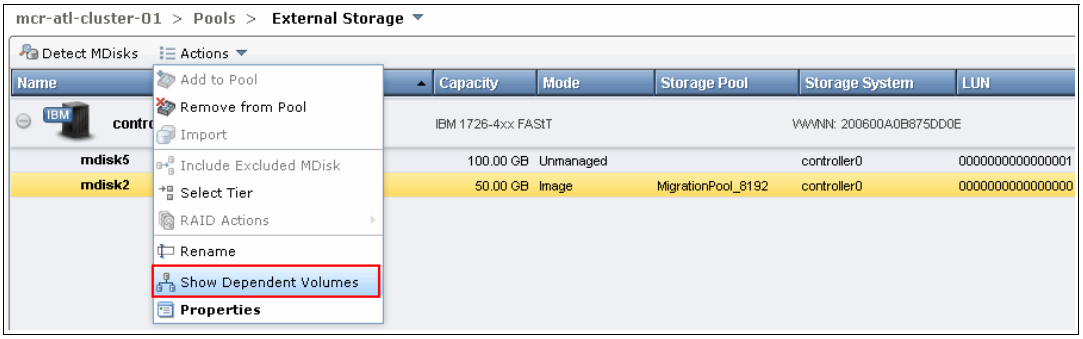

Figure 11-7 Show Dependent Volumes option in the Actions drop-down menu

Clicking the **Show Dependent Volumes** option shows you the volumes in this external storage system, as shown in Figure 11-8.

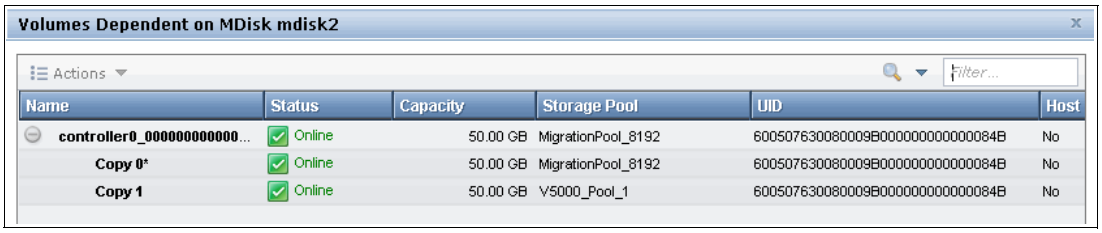

Figure 11-8 Volumes dependent on external storage

In the window that is shown in Figure 11-8 on page 555, you can take volume actions, including Map to Host, Shrink, Expand, Migrate to Another Pool, and Volume Copy Actions, as shown in Figure 11-9.

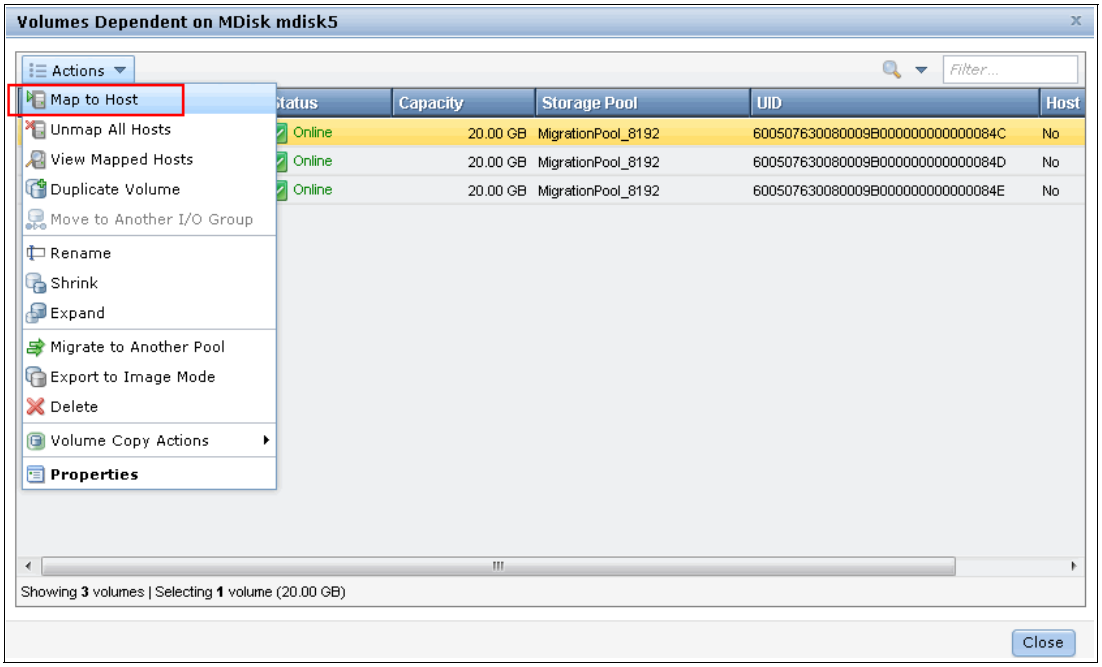

Figure 11-9 Actions that you can take with volumes

One of the features of the IBM Storwize V5000 storage system is that it can be used as a data migration tool. In the IBM Storwize V5000 virtualization environment, you can migrate your application data nondisruptively from one internal or external storage system to another storage system, which makes storage management much simpler with less risk.

Volume copy is another key feature that you can benefit from by using IBM Storwize V5000 virtualization. Two copies can be applied to your data to enhance the availability for a critical application. A volume copy can be also used to generate test data or data migration.

For more information about the volume actions of the IBM Storwize V5000 storage system, see Chapter 8, “Advanced host and volume administration” on page 349.

Returning to the External window, you discover an MDisk menu on the right, including an MDisk list that shows the MDisks that are provided by this external storage system. You can find the name of an MDisk, its capacity, the storage pool, and the storage system it belongs to in the list. The actions on MDisks can also be made through the menu, including Detect MDisks, Add to Pool, and Import. This menu is the same as the one in the MDisks window.

Figure 11-10 shows the MDisk menu for the external storage window.

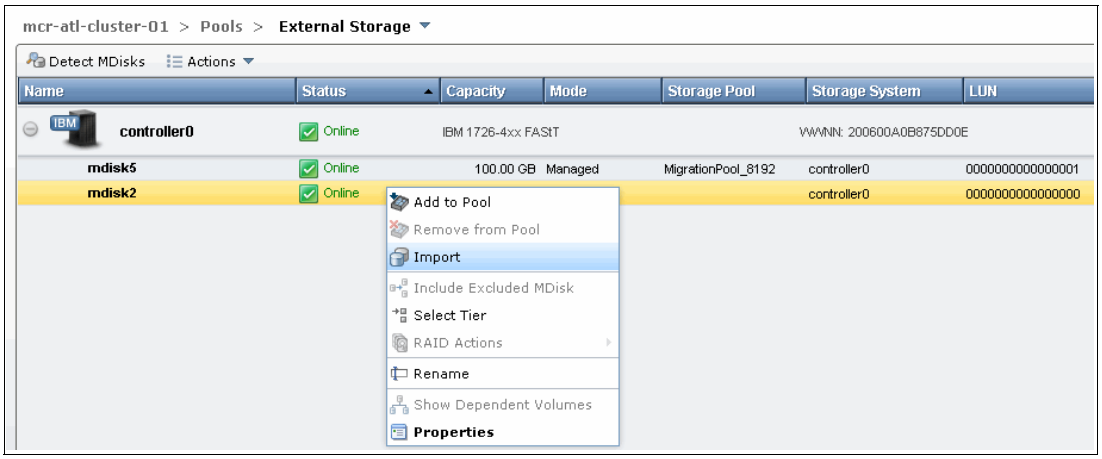

Figure 11-10 MDisk menu in the External window

11.2.3 Removing external storage

If you want to remove the external storage systems from the IBM Storwize V5000 virtualized environment, you have the following options:

- If you want to remove the external storage systems and discard the data on it, complete the following steps:
  a. Stop any host I/O on the volumes.
  b. Remove the volumes from the host file system, logical volume, or volume group, and remove the volumes from the host device inventory.
  c. Remove the host mapping of volumes and the volumes on IBM Storwize V5000.
  d. Remove the storage pools to which the external storage systems belong, or you can keep the storage pool and remove the MDisks of the external storage from the storage pools.
  e. Unzone and disconnect the external storage systems from the IBM Storwize V5000.
  f. Click **Detect MDisks** to make IBM Storwize V5000 discover the removal of the external storage systems.
- If you want to remove the external storage systems and keep the volumes and their data on the IBM Storwize V5000, complete the following steps:
  a. Migrate volumes and their data to the other storage pools that are on IBM Storwize V5000 internal storage or other external storage systems.
  b. Remove the storage pools to which the external storage systems belong, or you can keep the storage pools and remove the MDisks of the external storage from the storage pools.
  c. Unzone and disconnect the external storage systems from the IBM Storwize V5000.
  d. Click **Detect MDisks** to make IBM Storwize V5000 discover the removal of the external storage systems.

- If you want to remove the external storage systems from IBM Storwize V5000 control and keep the volumes and their data on external storage systems, complete the following steps:

  a. Migrate volumes and their data to the other storage pools that are on IBM Storwize V5000 internal storage or other external storage systems, as described in Chapter 6, “Storage migration wizard” on page 237.

  b. Remove the storage pools to which the external storage systems belong, or you can keep the storage pools and remove the MDisks of the external storage from the storage pools.

  c. Export volumes to image mode with the MDisks on the external storage systems. For more information about the restrictions and prerequisites for migration, see Chapter 6, “Storage migration wizard” on page 237.

     You also must record pre-migration information; for example, the original SCSI IDs the volumes used to be mapped to hosts. Some operating systems do not support changing the SCSI ID during the migration. For more information about migration, see the IBM Storwize V5000 Information Center at this website:

     http://publib.boulder.ibm.com/infocenter/storwize/ic/index.jsp

  d. Unzone and disconnect the external storage systems from the IBM Storwize V5000.

  a. Click **Detect MDisks** to make IBM Storwize V5000 discover the removal of the external storage systems.

12

RAS, monitoring, and troubleshooting

There are various ways to monitor and troubleshoot the IBM Storwize V5000. In this chapter, we show the ways in which the IBM Storwize V5000 can be administered from a monitoring and troubleshooting point of view.

This chapter includes the following topics:

- Reliability, availability, and serviceability on the IBM Storwize V5000
- IBM Storwize V5000 components
- Configuration backup procedure
- Upgrading software
- Event log
- Collecting support information
- Powering on and shutting down IBM Storwize V5000

12.1 Reliability, availability, and serviceability on the IBM Storwize V5000

This section describes the Reliability, Availability, and Serviceability (RAS) features of IBM Storwize V5000 monitoring and troubleshooting. RAS features are important concepts in the design of the IBM Storwize V5000. Hardware and software features, design considerations, and operational guidelines all contribute to make the IBM Storwize V5000 reliable.

Fault tolerance and a high level of availability are achieved by the following features:

- The RAID capabilities of the underlying disk subsystems.
- The compass architecture that is used by the IBM Storwize V5000 nodes.
- Auto-restart of nodes that are hung.
- Battery units to provide cache memory protection in the event of a site power failure.
- Host system multipathing and failover support.

High levels of serviceability are achieved by providing the following benefits:

- Cluster error logging
- Asynchronous error notification
- Dump capabilities to capture software detected failures
- Concurrent diagnostic procedures
- Directed maintenance procedures
- Concurrent log analysis and memory dump data recovery tools
- Concurrent maintenance of all IBM Storwize V5000 components
- Concurrent upgrade of IBM Storwize V5000 Software and microcode
- Concurrent addition or deletion of a node canister in a cluster
- Software recovery through the Service Assistant Tool
- Automatic software version correction when a node is replaced
- Detailed status and error conditions that are displayed via the Service Assistant Tool
- Error and event notification through Simple Network Management Protocol (SNMP), syslog, and email
- Node canister support package gathering via USB, in case of network connection problem

At the heart of the IBM Storwize V5000 is a redundant pair of *node canisters*. The two canisters share the data transmitting and receiving load between the attached hosts and the disk arrays.

12.2 IBM Storwize V5000 components

This section describes each of the components that make up the IBM Storwize V5000 system. Components are described in terms of location, function, and serviceability.

12.2.1 Enclosure midplane assembly

The enclosure midplane assembly is the unit that contains the node or expansion canisters and the power supply units. The enclosure midplane assembly initially is generic and configured as a control enclosure midplane or an expansion enclosure midplane. During the basic system configuration, Vital Product Data (VPD) is written to the enclosure midplane assembly, which decides whether the unit is a control enclosure midplane or an expansion enclosure midplane.

Control enclosure midplane

The control enclosure midplane holds node canisters and the power supply units. The control enclosure midplane assembly has specific VPD, such as, WWNN 1, WWNN 2, machine type and model, machine part number, and serial number. The control enclosure midplane must be replaced only by a trained service provider. After a generic enclosure midplane assembly is configured as a control enclosure midplane, it is no longer interchangeable with an expansion enclosure midplane.

Expansion enclosure midplane

The expansion enclosure midplane holds expansion canisters and the power supply units. The expansion enclosure midplane assembly also has specific VPD, such as, machine type and model, machine part number, and serial number. After a generic enclosure midplane assembly is configured as an expansion enclosure midplane, it is no longer interchangeable with a control enclosure midplane. The expansion enclosure midplane must be replaced only by a trained service provider.

Figure 12-1 shows back of the Enclosure Midplane Assembly.

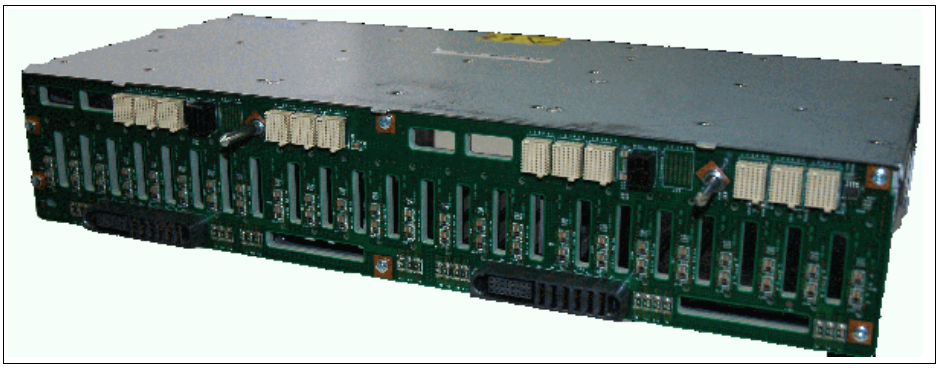

Figure 12-1 Rear view of Enclosure Midplane Assembly

For more information about replacing the control or expansion enclosure midplane, see the IBM Storwize V5000 Information Center at this website:

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

12.2.2 Node canisters: Ports and LED

There are two node canister slots along the top of the unit. The left slot is canister 1 and the right slot is canister 2.

Figure 12-2 shows the back of a fully equipped node enclosure.

Figure 12-2 Node canister

USB ports

There are two USB connectors side-by-side and they are numbered as 1 on the left and as 2 on the right. There are no indicators that are associated with the USB ports. Figure 12-3 shows the USB ports.

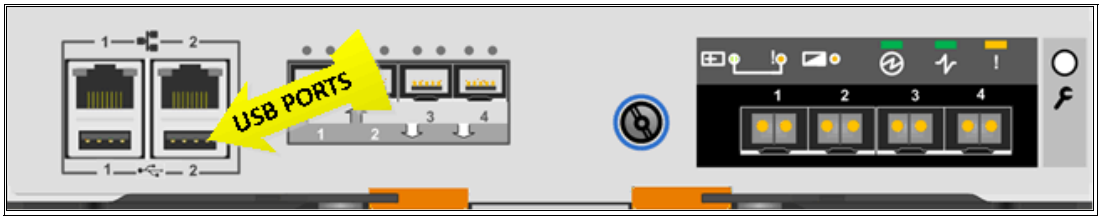

Figure 12-3 Node Canister USB ports

Ethernet ports

There are two 10/100/1000 Mbps Ethernet ports side-by-side on the canister and they are numbered 1 on the left and 2 on the right. Port 1 is required and port 2 optional. The ports are shown in Figure 12-4.

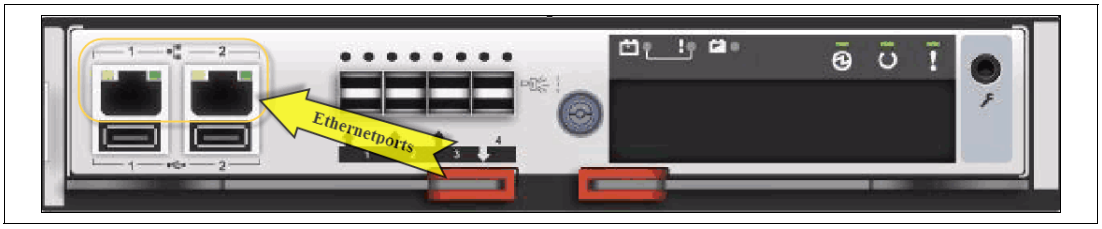

Figure 12-4 Node canister Ethernet ports

Each port has two LEDs and their status is shown in Table 12-1.

Table 12-1 Ethernet LEDs status

LED	Color	Meaning
Link state	Green	On: There is an Ethernet link.
Activity	Yellow	Flashing: There is activity on the link.

SAS ports

There are four 6-Gbps Serial Attached SCSI (SAS) ports side-by-side on the canister. They are numbered 1 on the left to 4 on the right. IBM Storwize V5000 uses port 1 and 2 for host connectivity and ports 3 and 4 to connect optional expansion enclosure. The ports are shown in Figure 12-5.

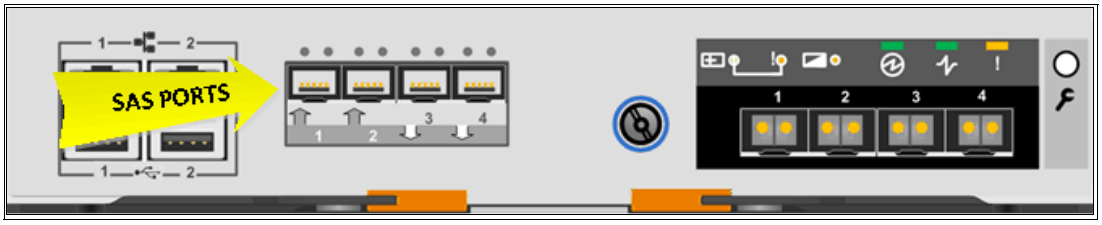

Figure 12-5 Node canister SAS ports

The SAS LED status meanings are described in Table 12-2.

Table 12-2 SAS LED Status

State	Meaning
green	Indicates at least one of the SAS lanes on this connector are operational. If the light is off when it is connected, there is a problem with the connection.
amber	If the light is on, one of the following errors occurred: ► One or more (but not all) of the four lanes are up for this connector (if none of the lanes are up, the activity light is off) ► One or more of the up lanes are running at a different speed to the others ► One or more of the up lanes are attached to a different address to the others

IBM Storwize V5000 uses SFF-8644 mini-SAS HD connector cable to connect enclosures, as shown in Figure 12-6.

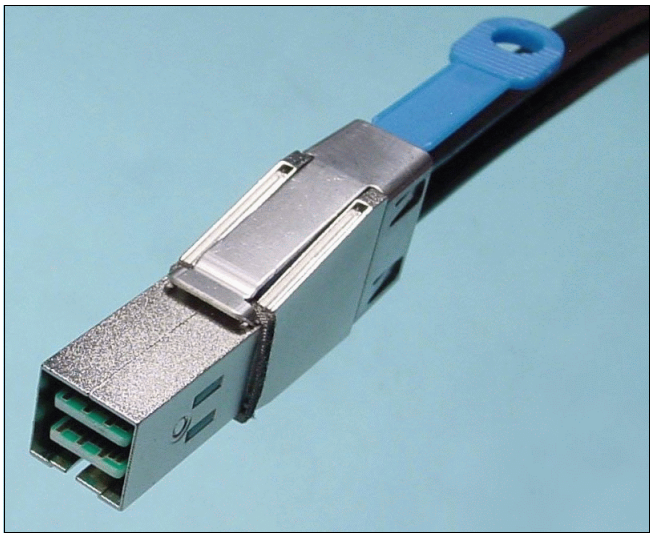

Figure 12-6 Mini-SAS HD SFF 8644 connector

Battery status

Each node canister stores a battery, the status of which is displayed on three LED on the back of the unit, as shown in Figure 12-7.

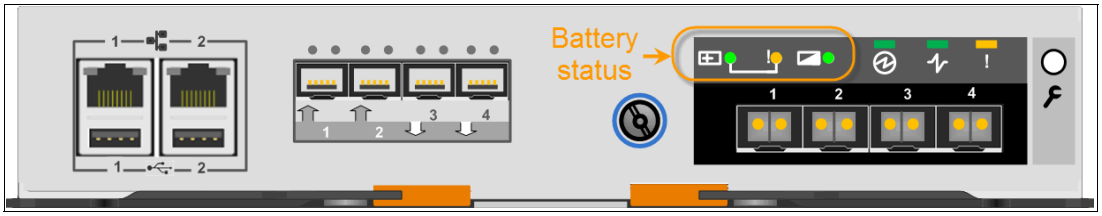

Figure 12-7 Node canister battery status

The battery indicator status meanings are described in Table 12-3.

Table 12-3 Battery indicator on Node canister

Color	Name	Definition
Green (left)	Battery Status	► Fast flash: Indicates that the battery is charging and has insufficient charge to complete a single memory dump. ► Flashing: Indicates that the battery has sufficient charge to complete a single memory dump only. ► Solid: Indicates that the battery is fully charged and has sufficient charge to complete two memory dumps.
Amber	Fault	Indicates a fault with the battery.
Green (right)	Battery in use	Indicates that hardened or critical data is writing to disk.

Canister status

The status of each canister is displayed by three LEDs, as shown in Figure 12-8.

Figure 12-8 System status indicator

The system status LED meanings are described in Table 12-4.

Table 12-4 System status indicator

Color	Name	Definition
Green (left)	System Power	► Flashing: The canister is in standby mode in which case IBM Storwize V5000 is not running. ► Fast flashing: The cannister is running a self test. ► On: The cannister is powered up and the IBM Storwize V5000 code is running.
Green (mid)	System Status	► Off: There is no power to the canister, the canister is in standby mode, Power On SelfTest (POST) is running on the canister, or the operating system is loading. ► Flashing: The node is in candidate or service state; it cannot perform I/O. It is safe to remove the node. ► Fast flash: A code upgrade is running. ► On: The node is part of a cluster.
Amber	Fault	► Off: The node is in candidate or active state. This state does not mean that there is no hardware error on the node. Any error that is detected is not severe enough to stop the node from participating in a cluster (or there is no power). ► Flashing: Identifies the canister. ► On: The node is in service state, or there is an error that is stopping the software from starting.

12.2.3 Node canister replaceable hardware components

The IBM Storwize V5000 node canister contains the following customer-replaceable replaceable components:

- Host Interface Card
- Memory
- Battery

Figure 12-9 shows the location of these parts within the node canister.

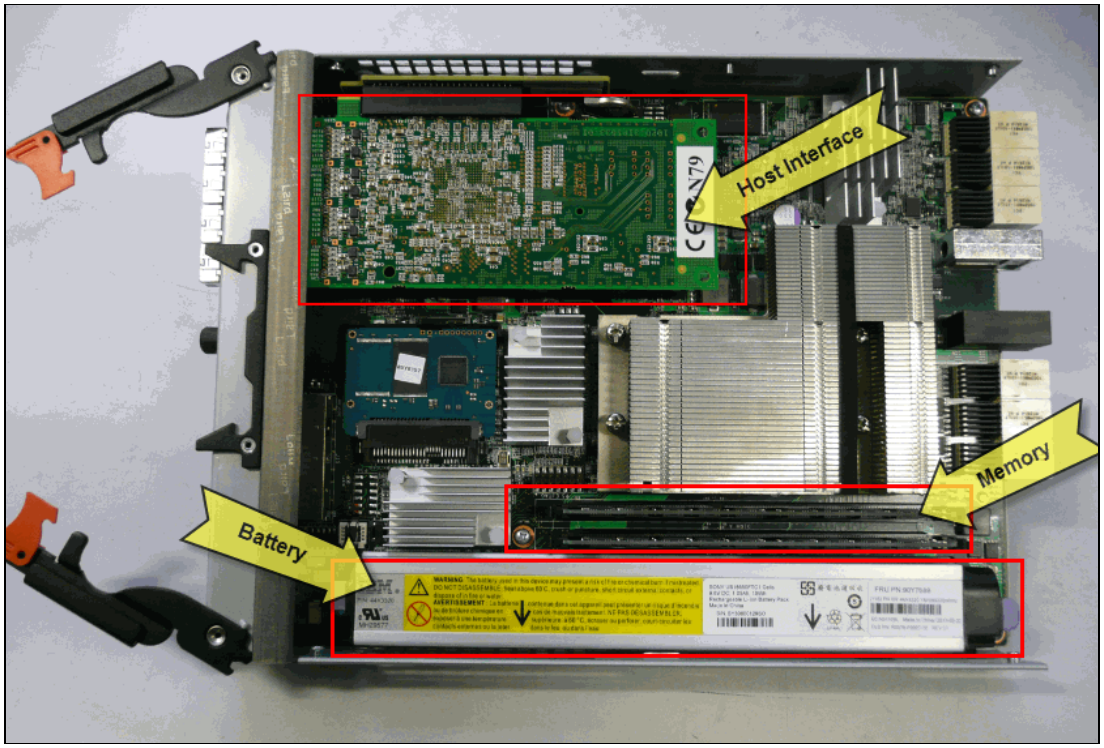

Figure 12-9 Node canister customer replaceable parts

Host interface card replacement

For more information about the replacement process, see the IBM Storwize V5000 Information Center at this website:

http://pic.dhe.ibm.com/infocenter/storwize/V5000_ic/topic/com.ibm.storwize.V5000.641.doc/V5000_rplc_hic.html

At the website, browse to **Troubleshooting Removing and replacing parts** → **Replacing host interface card**.

The host interface card replacement is shown in Figure 12-10.

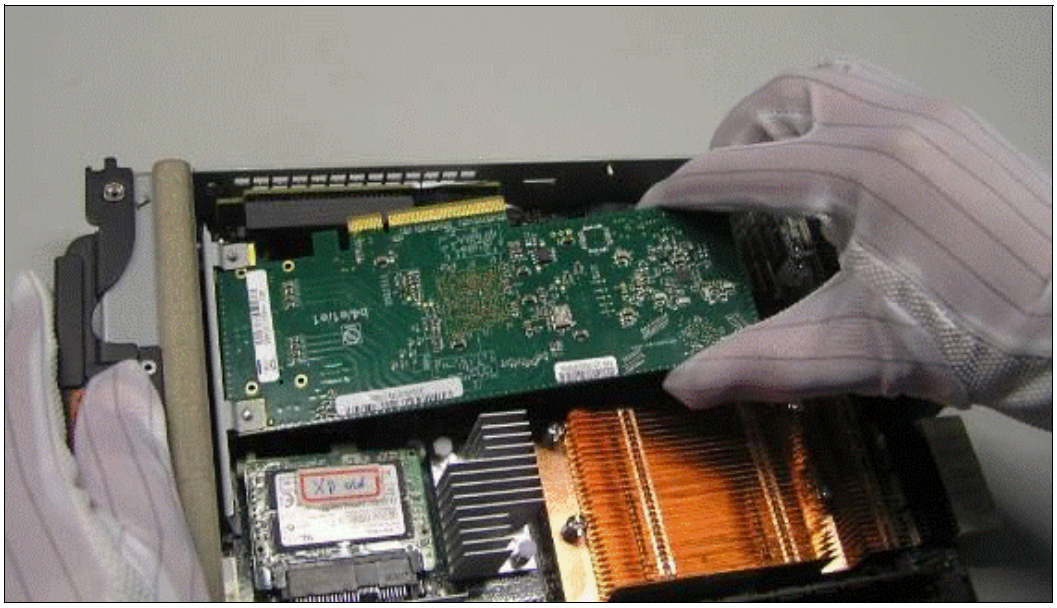

Figure 12-10 Host Interface card replacement

Memory replacement

For more information about the memory replacement process, see the IBM Storwize V5000 Information Center at this website:

http://pic.dhe.ibm.com/infocenter/storwize/V5000_ic/topic/com.ibm.storwize.V5000.641.doc/V5000_rplc_nodecan_dimm.html

At the website, browse to **Troubleshooting Removing and replacing parts** → **Replacing the node canister memory (2x 4 GB DIMM)**.

Figure 12-11 shows the memory location.

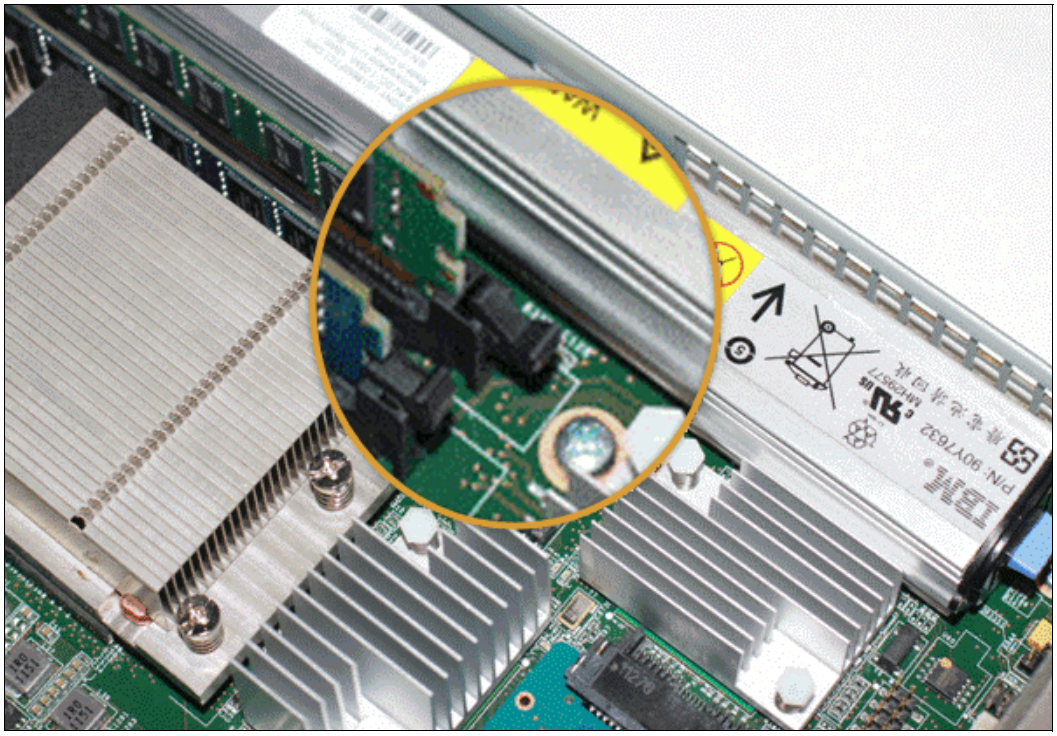

Figure 12-11 Memory replacement

Battery Backup Unit replacement

Caution: The battery is a lithium ion battery. To avoid possible explosion, do not incinerate the battery. Exchange the battery only with the part that is approved by IBM.

Because the Battery Backup Unit (BBU) is seated in the node canister, the BBU replacement leads to a redundancy loss until the replacement is completed. Therefore, it is recommended to replace the BBU only when advised to do so. It is also recommended to follow the Directed Maintenance Procedures (DMP).

For more information about how to replace the BBU, see the Information Center at this website:

http://pic.dhe.ibm.com/infocenter/storwize/V5000_ic/topic/com.ibm.storwize.V5000.641.doc/V5000_rplc_batt_nodecan.html

At the website, browse to **Troubleshooting Removing and replacing parts** → **Replacing battery in a node canister**.

Complete the following steps to replace the BBU:

1. Grasp the blue touch points on each end of the battery, as shown in Figure 12-12.

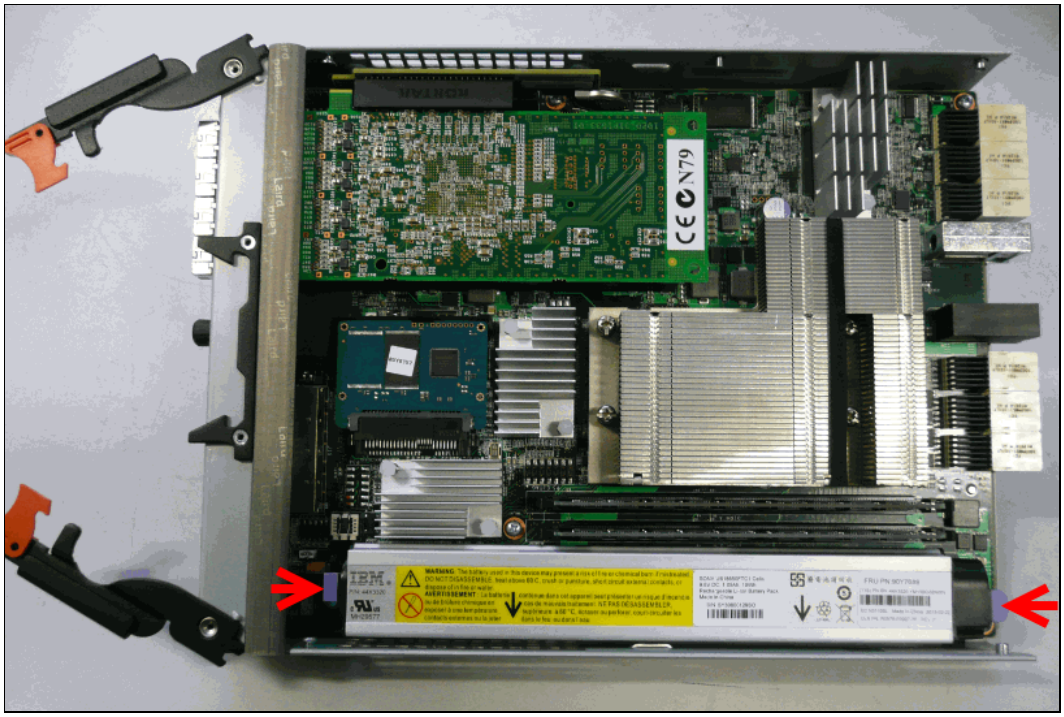

Figure 12-12 BBU replacement: Step 1

2. Lift the battery vertically upwards until the connectors disconnect.

Important: During a BBU change, the battery must be kept parallel to the canister system board while it is removed or replaced, as shown in Figure 12-13. Keep equal force, or pressure, on each end.

Figure 12-13 BBU replacement: Step 2

12.2.4 Expansion canister: Ports and LED

There are two expansion canister slots along with top of the unit.

SAS ports

SAS ports are used to connect the expansion canister to the node canister or to an extra expansion in the chain. Figure 12-14 shows the SAS ports that are on the expansion canister.

Figure 12-14 Expansion canister SAS ports

The meaning of the SAS port LEDs is described in Table 12-5.

Table 12-5 SAS LED status meaning

State	Meaning
Green	Indicates at least one of the SAS lanes on these connectors are operational. If the light is off when connected, there is a problem with the connection.
Amber	If the light is on, one of the following errors occurred: ► One or more (but not all) of the four lanes are up for this connector (if none of the lanes are up, the activity light is off). ► One or more of the up lanes are running at a different speed to the others. ► One or more of the up lanes are attached to a different address to the others.

Canister status

Each expansion canister has its status displayed by three LEDs, as shown in Figure 12-15.

Figure 12-15 Enclosure canister status

The LED status is described in Table 12-6.

Table 12-6 Enclosure canister status

Color	Name	Definition
Green (left)	Power	Indicates that the canister is receiving power.
Green (mid)	Status	If the light is on, the canister is running normally. If the light is flashing, there is an error communicating with the enclosure.
Amber	Fault	If the light is solid, there is an error logged against the canister or the firmware is not running.

12.2.5 Disk subsystem

The IBM Storwize V5000 disk subsystem is made up of control and expansion enclosures. The system can have one or two control enclosures, with each control enclosure attaching to up to six expansion enclosures. Each enclosure contains the drives that are based on the enclosure type.

This section describes the parts of the disk subsystem.

SAS cabling

Expansion enclosures are attached to control enclosures by using SAS cables. There are two supported SAS chains and up to three expansion enclosures can be attached to each chain. The node canister uses SAS ports 3 and 4 for enclosures while ports 1 and 2 for host connectivity.

Important: When an SAS cable is inserted, ensure that the connector is oriented correctly by confirming that the following conditions are met:

- The pull tab must be below the connector.
- Insert the connector gently until it clicks into place. If you feel resistance, the connector is probably oriented the wrong way. Do not force it.
- When inserted correctly, the connector can be removed only by pulling the tab.

The expansion canister has SAS port 1 for channel input and SAS port 2 for output to connect another expansion enclosure.

The SAS cabling is shown in Figure 12-16.

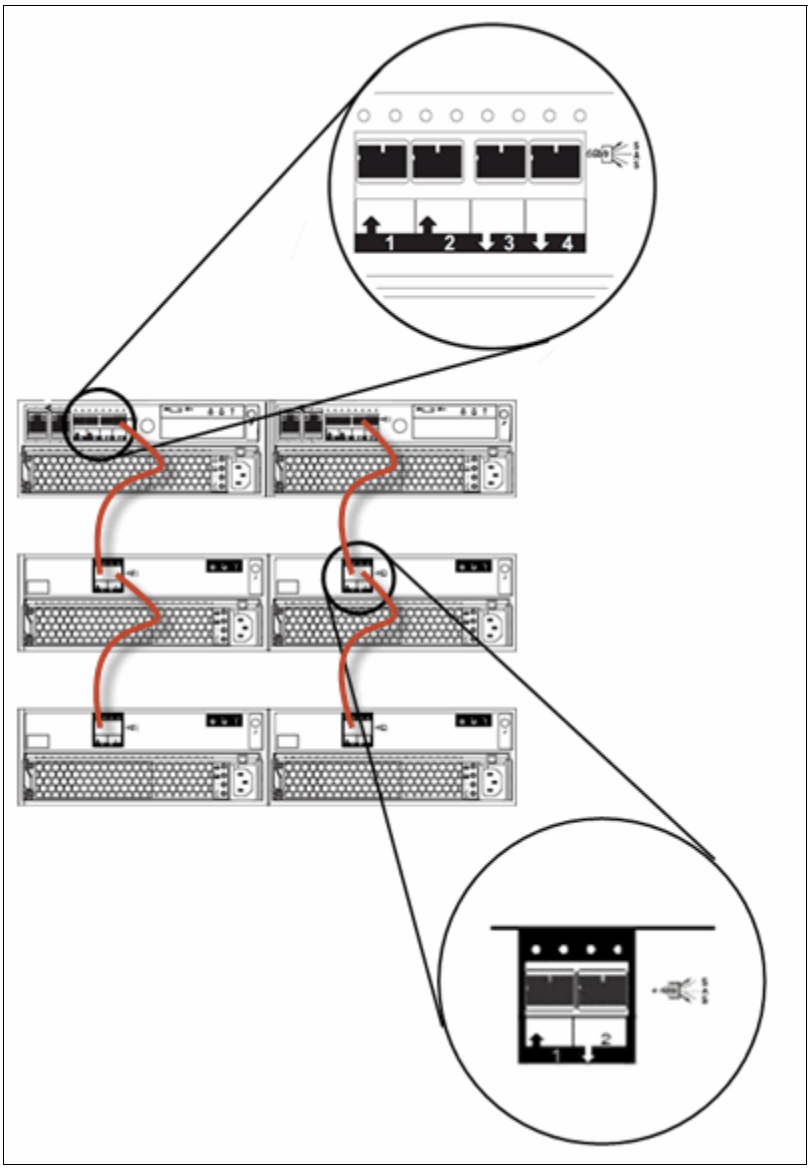

Figure 12-16 SAS cabling for single I/O Group)

A strand starts with an SAS initiator chip inside an IBM Storwize V5000 node canister and progresses through SAS expanders, which connect disk drives. Each canister contains an expander. Each drive has two ports, each of which is connected to a different expander and strand. This configuration means both nodes directly access each drive and there is no single point of failure.

At system initialization when devices are added to or removed from strands (and at other times), the IBM Storwize V5000 Software performs a discovery process to update the state of the drive and enclosure objects.

Slot numbers in enclosures

The IBM Storwize V5000 is made up of enclosures. There are four types of enclosures, as described in Table 12-7.

Table 12-7 Enclosure slot numbering

Enclosure type	Number of slots
Enclosure 12x 3.5-inch drives: ► Control enclosure 2077-12C ► Expansion enclosure 2077-12E	Enclosure with 12 slots. 
Enclosure 24x 2.5-inch drives: ► Control enclosure 2077-24C ► Expansion enclosure 2077-24E	Enclosure with 24 slots. 

Array goal

Each array has a set of goals that describe the wanted location and performance of each array member. A sequence of drive failures and hot spare takeovers can leave an array unbalanced; that is, with members that do not match these goals. The system automatically rebalances such arrays when appropriate drives are available.

RAID level

An IBM Storwize V5000 supports the RAID 0, RAID 1, RAID 5, RAID 6, or RAID 10. Each RAID level is described in Table 12-8.

Table 12-8 RAID levels that are supported by an IBM Storwize V5000

RAID level	Where data is striped	Drive count (Min - Max)
0	Arrays have no redundancy and do not support hot-spare takeover.	1 - 8
1	Provides disk mirroring, which duplicates data between two drives. A RAID 1 array is internally identical to a two-member RAID 10 array.	2
5	Arrays stripe data over the member drives with one parity strip on every stripe. RAID 5 arrays have single redundancy with higher space efficiency than RAID 10 arrays, but with some performance penalty. RAID 5 arrays can tolerate no more than one member drive failure.	3 - 16
6	Arrays stripe data over the member drives with two parity strips on every stripe. A RAID 6 array can tolerate any two concurrent member drive failures.	5 - 16

RAID level	Where data is striped	Drive count (Min - Max)
10	Arrays stripe data over mirrored pairs of drives. RAID 10 arrays have single redundancy. The mirrored pairs rebuild independently. One member out of every pair can be rebuilding or missing at the same time. RAID 10 combines the features of RAID 0 and RAID 1.	2 - 16

Disk scrubbing

The scrub process runs when arrays do not have any other background processes. The process checks that the drive logical block addresses (LBAs) are readable and array parity is synchronized. Arrays are scrubbed independently and each array is entirely scrubbed every seven days.

Solid-state drives

Solid-state drives (SSDs) are treated no differently by an IBM Storwize V5000 than hard disk drives (HDDs) concerning RAID arrays or MDisks. The SSDs in the storage that are managed by the IBM Storwize V5000 are combined into an array, usually in RAID 10 or RAID 5 format. It is unlikely that RAID 6 SSD arrays are used because of the double parity impact, with two SSD logical drives that are used for parity only.

12.2.6 Power supply unit

All enclosures require two power supply units (PSUs) for normal operation. A single PSU can power the entire enclosure for redundancy.

Figure 12-17 shows the power supplies.

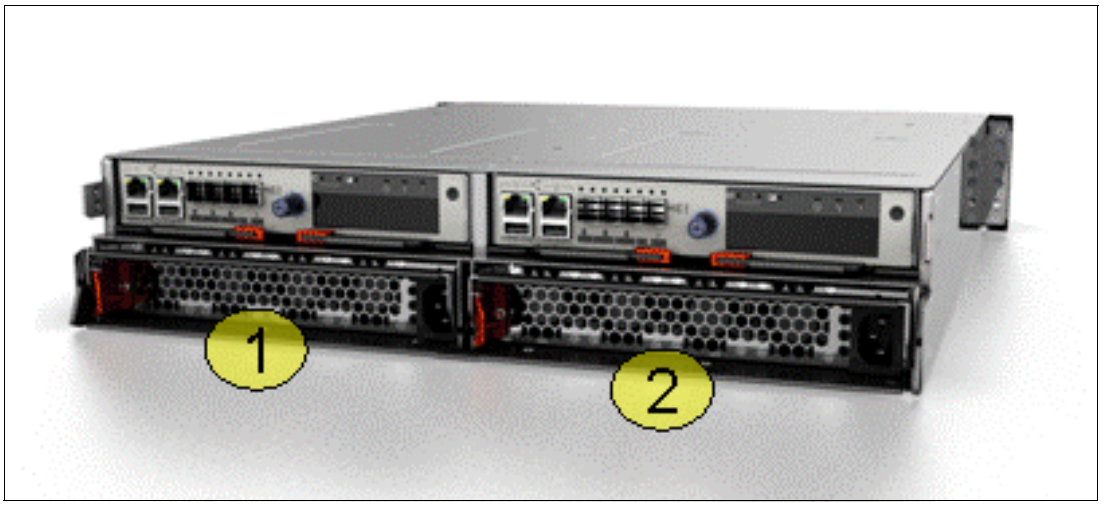

Figure 12-17 Power supply

The left side PSU is numbered 1 and the right side PSU is numbered 2.

PSU LED indicator

The indicators are the same for the control and expansion unit.

Figure 12-18 shows the PSU LED Indicators.

Figure 12-18 PSU LED Indicators

Table 12-9 shows the colors and meaning of the LEDs.

Table 12-9 PSU LED definitions

Position	Color	Marking	Name	Definition
1	Green	In	AC Status	Main power is delivered
2	Green	DC	DC Status	DC power is available
3	Amber	Fault exclamation mark	Fault	Fault on PSU
4	Blue	OK	Service action that is allowed	N/A

12.3 Configuration backup procedure

If there is a serious failure that requires the system configuration must be restored, the configuration backup file must be used. The file contains configuration data such as, arrays, pools, and volumes (but no customer applications data). The backup file is updated by the cluster every day.

Even so, it is important to save the file after you change your system configuration, which requires a command-line interface (CLI) connection to start manual backup.

Regularly saving a configuration backup file on the IBM Storwize V5000 is important and it must be done manually. Download this file regularly to your management workstation to protect the configuration data (a best practice is to automate this download procedure by using a script and saving it daily on a remote system).

12.3.1 Generating a configuration backup by using the CLI

To generate a configuration backup by using the CLI, run the **svcconfig backup** command, as shown in Example 12-1.

Example 12-1 Example for backup CLI command

```
svcconfig backup
```

The progress of the command is detected by dots, as shown in Example 12-2.

Example 12-2 Backup CLI command progress and output

```
.....................................................................................
.....................................................................................
...................
CMMVC6155I SVCCONFIG processing completed successfully
```

The **svcconfig backup** command creates three files that provide information about the backup process and cluster configuration. These files are created in the `/tmp` directory on the configuration node and are listed on the support view.

The three files that are created by the backup process are described Table 12-10.

Table 12-10 File names that are created by the backup process

File name	Description
`svc.config.backup.xml`	This file contains your cluster configuration data.
`svc.config.backup.sh`	This file contains the names of the commands that were issued to create the backup of the cluster.
`svc.config.backup.log`	This file contains details about the backup, including any error information that might be reported.

12.3.2 Downloading a configuration backup by using the GUI

To download a configuration backup file by using the GUI, complete the following steps:

1. Click the **Settings** icon and then click **Support**, as shown in Figure 12-19.

Figure 12-19 Configuration backup open support view

2. Select the configuration node on the support view, as shown in Figure 12-20.

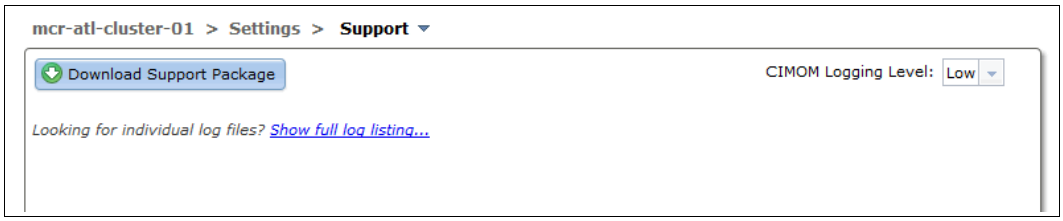

Figure 12-20 Configuration backup select configuration node

3. Select the **Show full log listing...** option (as shown in Figure 12-21) to list all of the available log files that are stored on the configuration node.

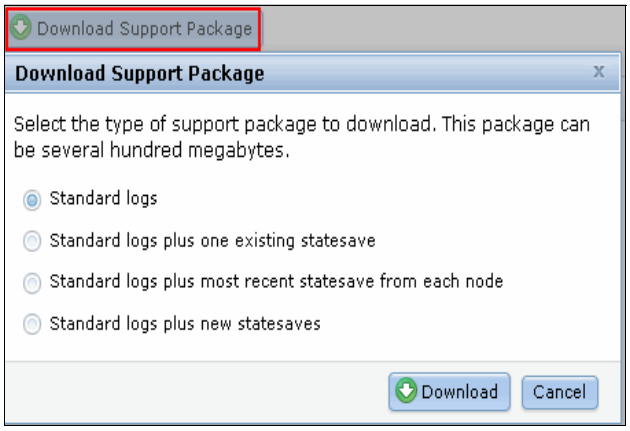

Figure 12-21 Support package selection

4. Search for a file named `/dumps/svc.config.backup.xml_*`, as shown in Figure 12-22. Select the file, right-click it, and then select **Download**.

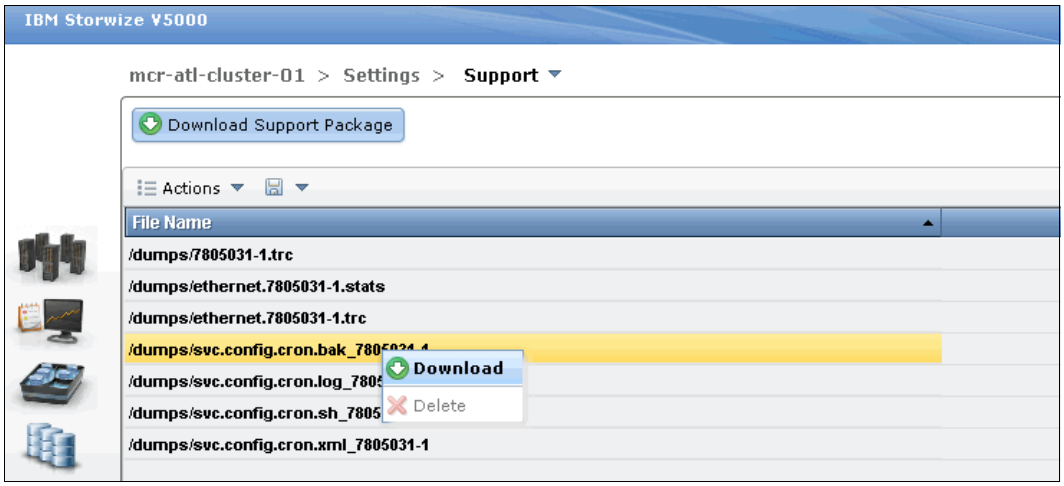

Figure 12-22 Configuration backup start download

5. Save the configuration backup file on your management workstation where it can be found easily, as shown in Figure 12-23.

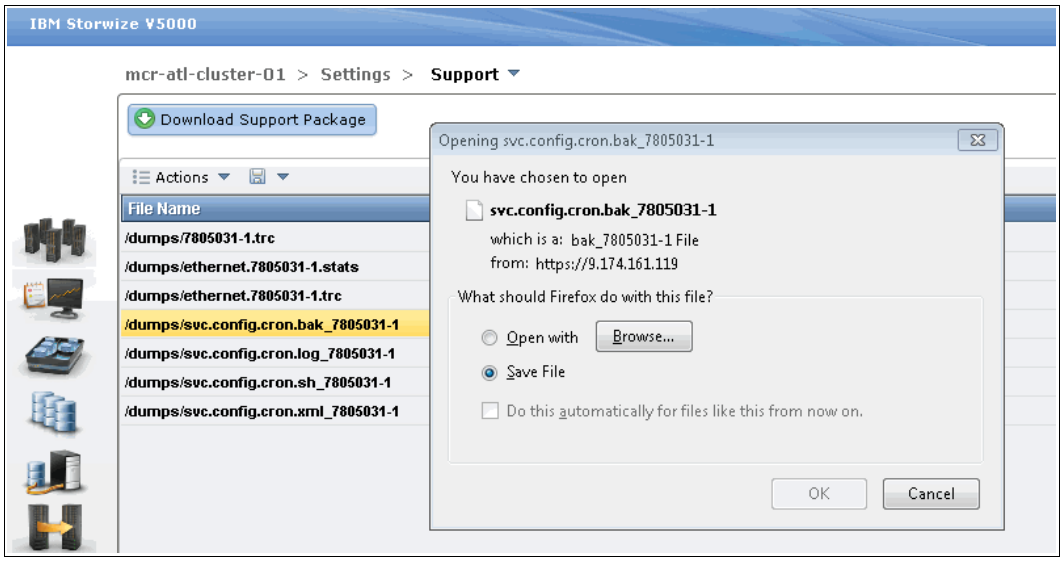

Figure 12-23 Configuration backup save file

Even if the configuration backup file is updated automatically, it might be of interest to verify the time stamp of the actual file. Therefore, the `/dumps/svc.config.backup.xml_xx` file must be opened with an editor, such as, WordPad, as shown in Figure 12-24.

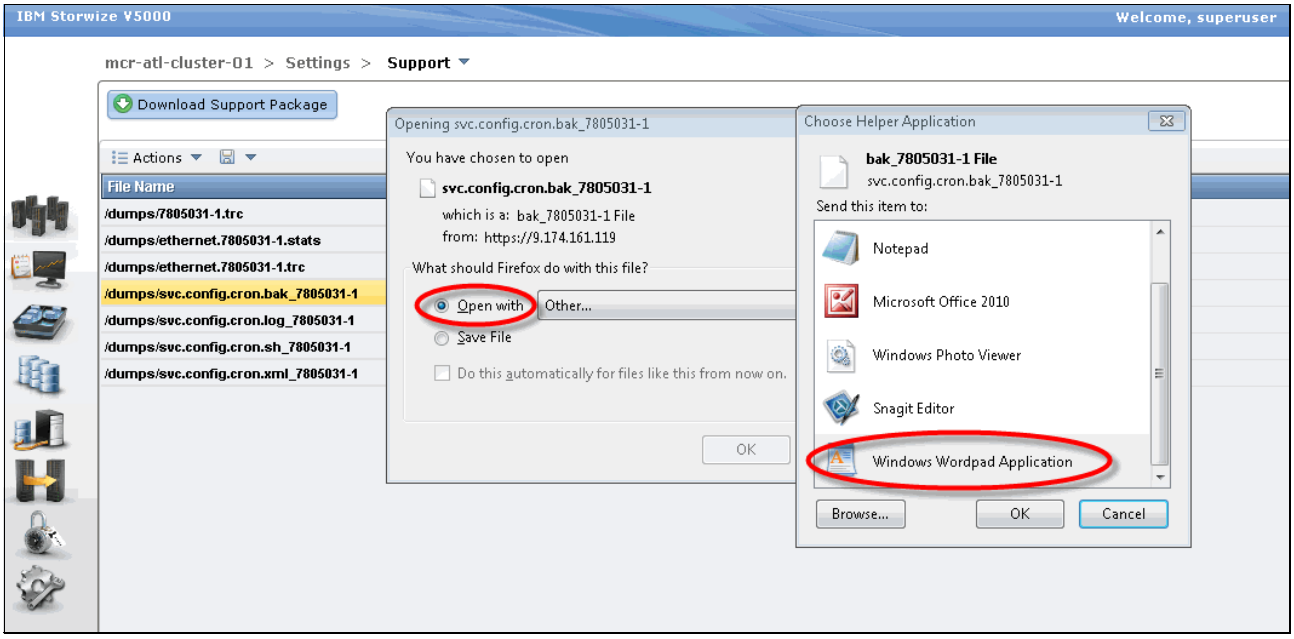

Figure 12-24 Open backup XML file with WordPad

Open the `/dumps/svc.config.backup.xml_xx file` with an editor (we used WordPad) and search for the string `timestamp=`, which is found near of the top of the file. Figure 12-25 shows the file that is opened and the time stamp information in it.

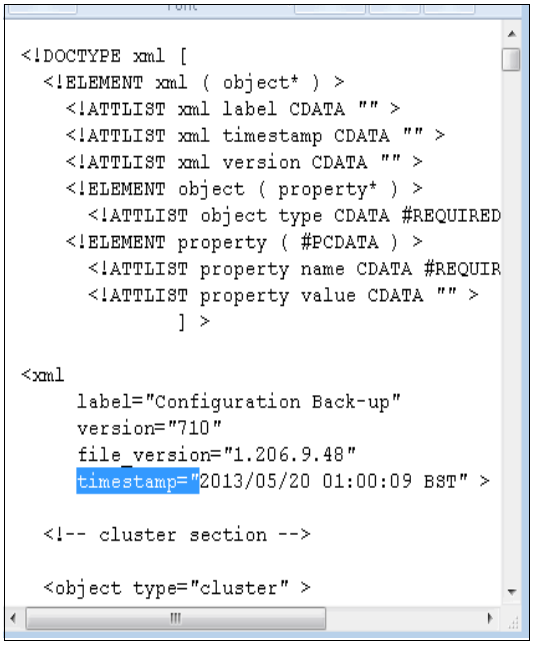

```
<!DOCTYPE xml [
  <!ELEMENT xml ( object* ) >
    <!ATTLIST xml label CDATA "" >
    <!ATTLIST xml timestamp CDATA "" >
    <!ATTLIST xml version CDATA "" >
    <!ELEMENT object ( property* ) >
      <!ATTLIST object type CDATA #REQUIRED
    <!ELEMENT property ( #PCDATA ) >
      <!ATTLIST property name CDATA #REQUIR
      <!ATTLIST property value CDATA "" >
              ] >

<xml
     label="Configuration Back-up"
     version="710"
     file_version="1.206.9.48"
     timestamp="2013/05/20 01:00:09 BST" >

  <!-- cluster section -->

  <object type="cluster" >
```

Figure 12-25 Timestamp in backup XML file

12.4 Upgrading software

The system upgrade process involves the upgrading of your entire IBM Storwize V5000 environment.

Allow sufficient time to plan your tasks, review your preparatory upgrade tasks, and complete the upgrade of the IBM Storwize V5000 environment. The upgrade procedures can be divided into these general processes. Table 12-11 shows the software upgrade tasks.

Table 12-11 Software upgrade tasks

Sequence	Upgrade tasks
1	Decide whether you want to upgrade automatically or manually. During an automatic upgrade procedure, the clustered system upgrades each of the nodes systematically. The automatic method is the preferred procedure for upgrading software on nodes. However, you can upgrade each node manually.
2	Ensure that CIM object manager (CIMOM) clients are working correctly. When necessary, upgrade these clients so that they can support the new version of IBM Storwize V5000 code.
3	Ensure that multipathing drivers in the environment are fully redundant. If you experience failover issues with multipathing driver support, resolve these issues before you start normal operations.
4	Upgrade other devices in the IBM Storwize V5000 environment. Examples might include upgrading hosts and switches to the correct levels.
5	Upgrade your IBM Storwize V5000.

Important: The amount of time it takes to perform an upgrade can vary depending on the amount of preparation work that is required and the size of the environment. Generally, allow more than two hours for an upgrade if you have two I/O Groups.

Some code levels support upgrades only from specific previous levels. If you upgrade to more than one level above your current level, you might be required to install an intermediate level.

Important: Ensure that you have no unfixed errors in the log and that the system date and time are correctly set. Start the fix procedures, and ensure that you have fix any outstanding errors before you attempt to concurrently upgrade the code.

12.4.1 Upgrading software automatically

During the automatic upgrade process, each node in the system upgrades individually and the new code is staged on the nodes. While each node restarts, there might be some degradation in the maximum I/O rate that can be sustained by the system. After all of the nodes in the system are successfully restarted with the new code level, the new level is automatically committed.

The upgraded node is temporarily unavailable and all I/O operations fail to that node. As a result, the I/O error counts increase and the failed I/O operations are directed to the partner node of the working pair. Applications do not see any I/O failures. When new nodes are added to the system, the upgrade package is automatically downloaded to the new nodes from the IBM Storwize V5000 system.

The upgrade can be performed concurrently with normal user I/O operations. However, there is a possibility that performance might be affected.

Multipathing requirement

Before you upgrade, ensure that the multipathing driver is fully redundant with every path that is available and online. You might see errors that are related to the paths, which go away (failover) and the error count that increases during the upgrade. When the paths to the nodes return, the nodes fall back to become a fully redundant system. After an approximate 30-minute delay, the paths to the other node fail.

12.4.2 GUI upgrade process

The automatic upgrade process is started in the GUI by starting the Upgrade wizard, as shown in Figure 12-26. Browse to **Settings** → **General** → **Upgrade Software** → **Launch Upgrade wizard**.

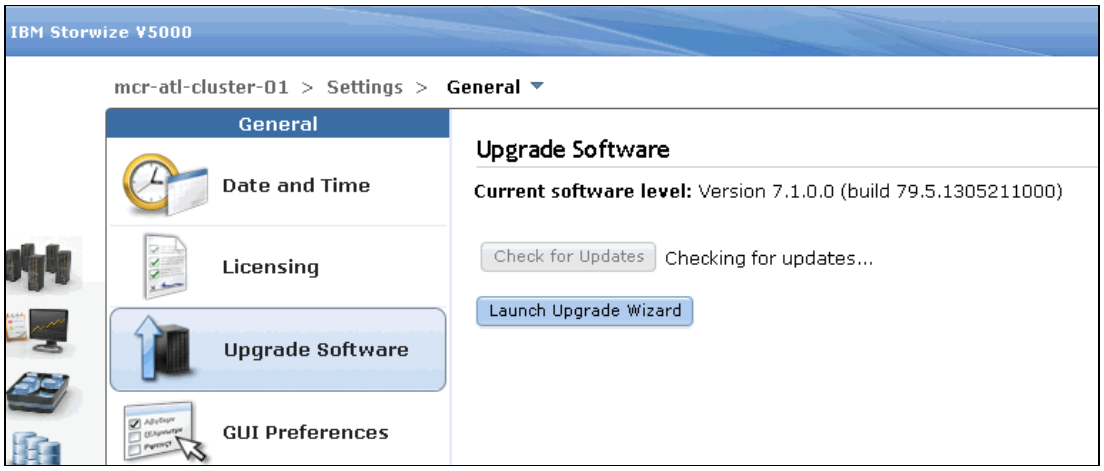

Figure 12-26 Start Upgrade wizard

As a first step, the Upgrade test utility must be downloaded from the Internet (the link is provided within the panel). If the tool was downloaded and stored on the management station, it can be uploaded, as shown in Figure 12-27.

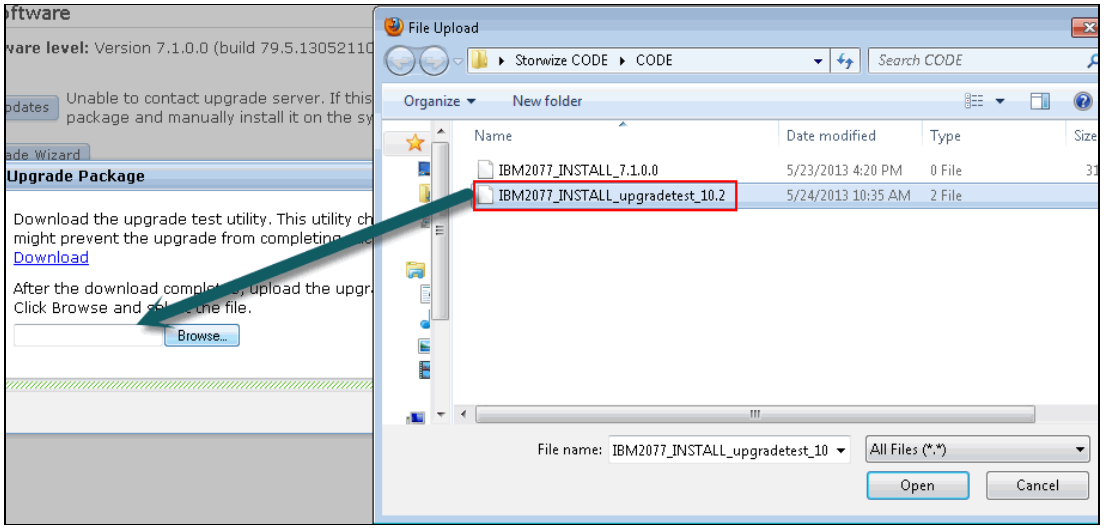

Figure 12-27 Download Upgrade test utility

A confirmation panel opens, as shown in Figure 12-28.

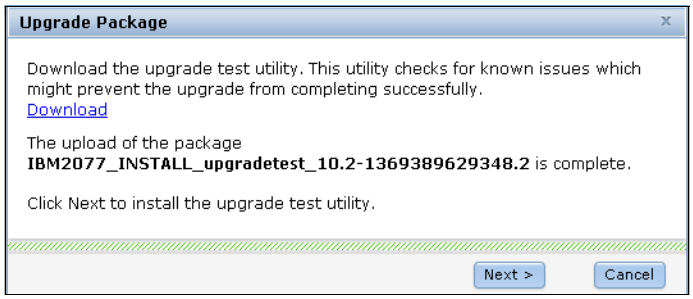

Figure 12-28 Upload test utility completed

The version to which the system should be upgraded must be entered in step 2 of the wizard. By default, the latest code level (at the time of writing) is shown, as shown in Figure 12-29.

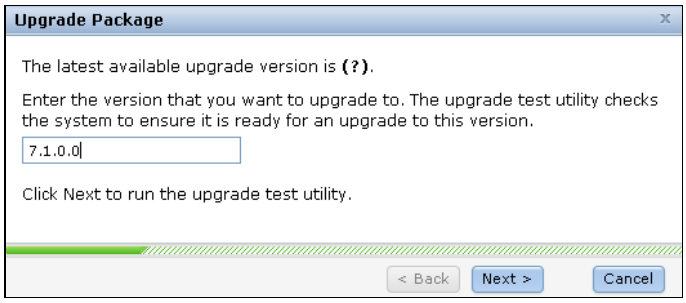

Figure 12-29 Enter version to be checked by tool

Important: You must choose the correct code level because you cannot recheck this information later. The version that is selected is used throughout the rest of the process.

Figure 12-30 shows the panel that indicates the background test task is running.

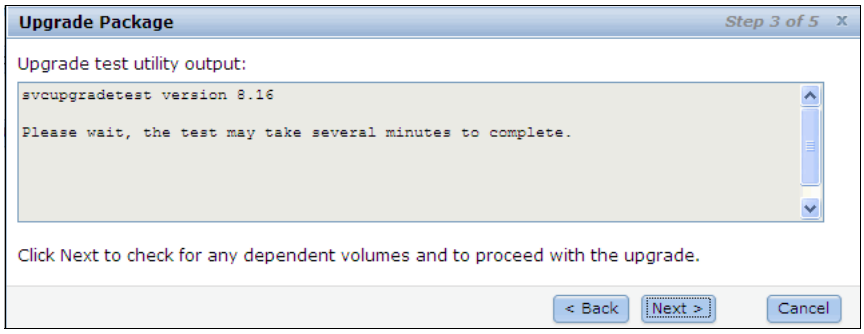

Figure 12-30 Wait utility to complete

The utility can be run as many times as necessary on the same system to perform a readiness check in preparation for a software upgrade.

Next, the code must be downloaded. If the code was downloaded to the management station, it can be directly uploaded, as shown in Figure 12-31. Verify that the correct code file is used.

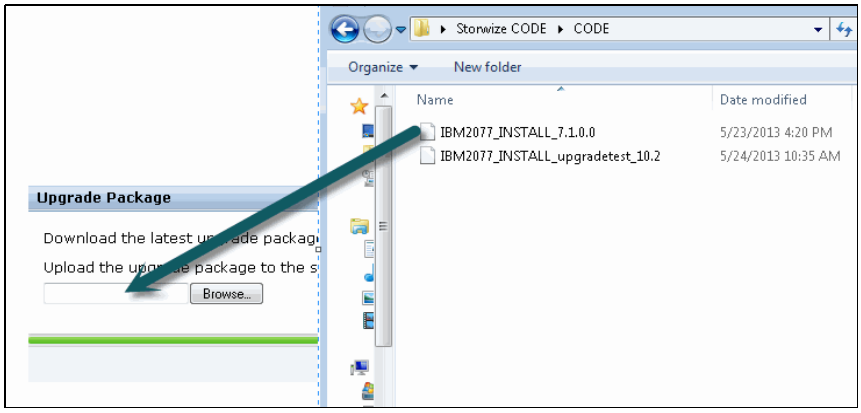

Figure 12-31 Download code

As shown in Figure 12-32, a confirmation window opens.

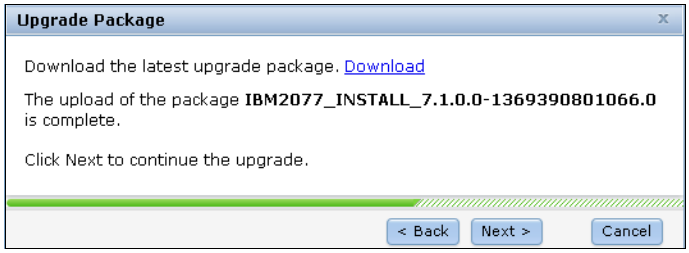

Figure 12-32 Code upload that is completed

The automated code upgrade can be started when the Automatic upgrade option is selected in the panel, as shown in Figure 12-33 (this is the default choice). If the upgrade is done manually for any reason, the selection must be made; however, an automatic upgrade is recommended.

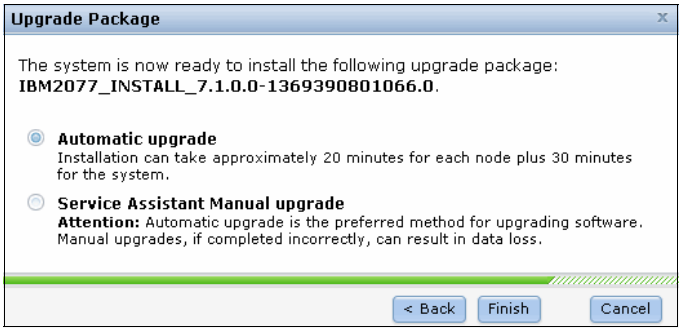

Figure 12-33 Upload mode decision

If you choose to select the **Service Assistant Manual** upgrade option, see 12.4.3, “Upgrading software manually” on page 584.

Select **Finish** to start the upgrade process on the nodes. Messages inform you when the nodes are upgraded. When all nodes are rebooted, the upgrade process is complete. It can take up to two hours to finish this process.

12.4.3 Upgrading software manually

Important: It is highly recommended to upgrade the IBM Storwize V5000 automatically by following the Upgrade wizard. If a manual upgrade is used, make sure that you do not skip any step.

The steps for manual upgrade are shown on the Service Assistant Manual Upgrade panel.

Complete the following steps to manually upgrade the software:

1. In the management GUI, click **Settings** → **General** → **Upgrade Software** and run the Upgrade wizard. In step 5 of the wizard, select **Service Assistant Manual upgrade**, as shown in Figure 12-34.

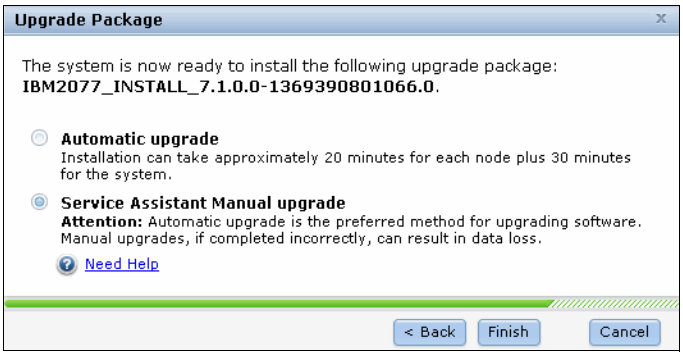

Figure 12-34 Select manual upgrade mode

After you select manual upgrade, a warning appears, as shown in Figure 12-35.

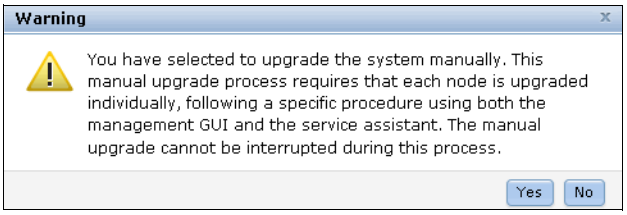

Figure 12-35 Manual upgrade warning

Both nodes are set to **"**Waiting for Upgrade" status in the Upgrade Machine Code panel, as shown in Figure 12-36.

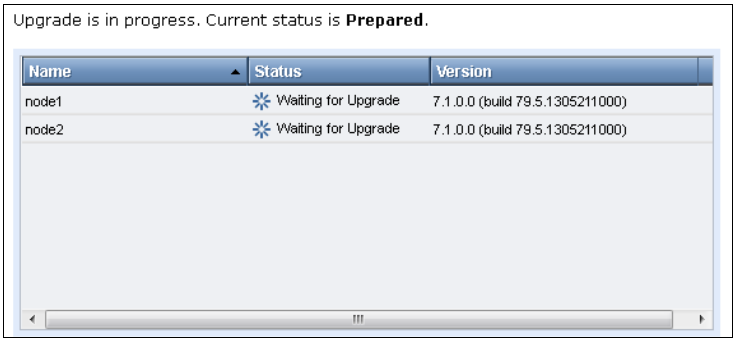

Figure 12-36 Node status to waiting for upgrade

2. In the management GUI, select **System Details** and select the canister that contains the node you want to upgrade next. As shown in Figure 12-37, select **Remove Node** in the Action menu, which shows you a Health Status alert.

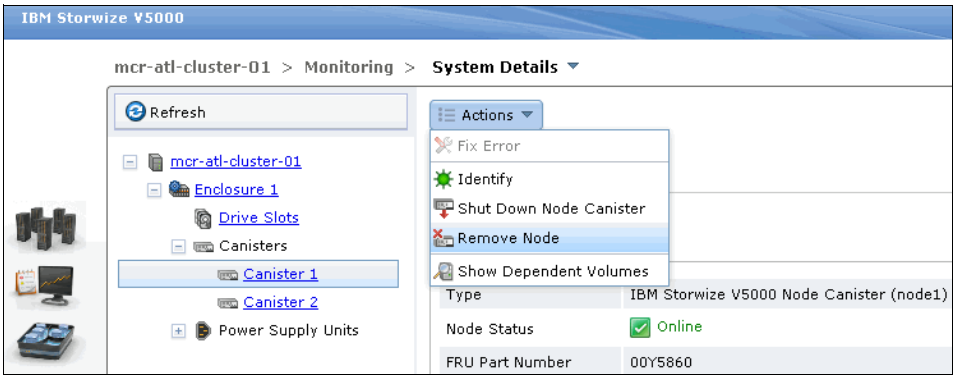

Figure 12-37 Remove the non-config node from cluster

Important: Make sure that you select the non-config node first.

A warning message appears, as shown in Figure 12-38.

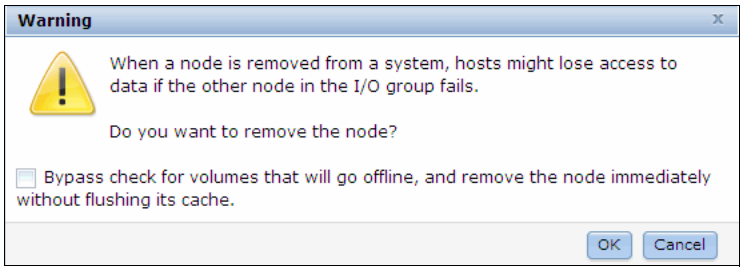

Figure 12-38 Remove node warning message

The non-configuration node is removed from GUI Upgrade Machine Code panel, as shown in Figure 12-39.

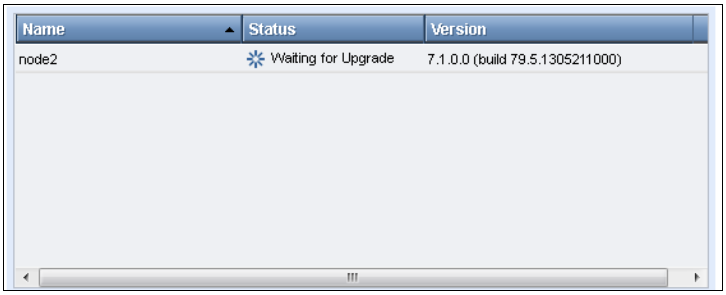

Figure 12-39 Non-configuration node was removed

In the System Details panel, the node is shown as Unconfigured, as shown in Figure 12-40.

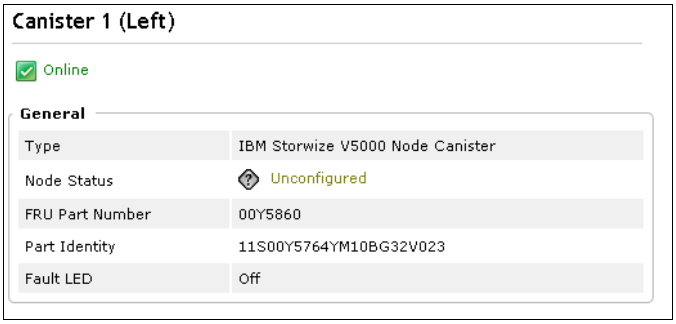

Figure 12-40 Node status shows unconfigured

3. In the Service Assistant panel, the node that is ready for upgrade must be selected. Select the node that shows Node status as service mode and has no available cluster information, as shown in Figure 12-41.

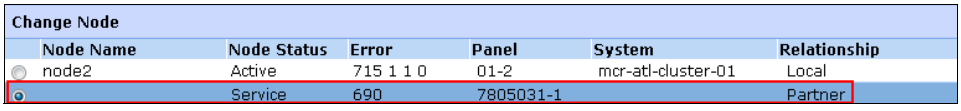

Figure 12-41 Select node in service mode for upgrade

4. In the Service Assistant panel, select **Upgrade Manually** and select the machine code version that you want to upgrade on selected node, as shown in Figure 12-42.

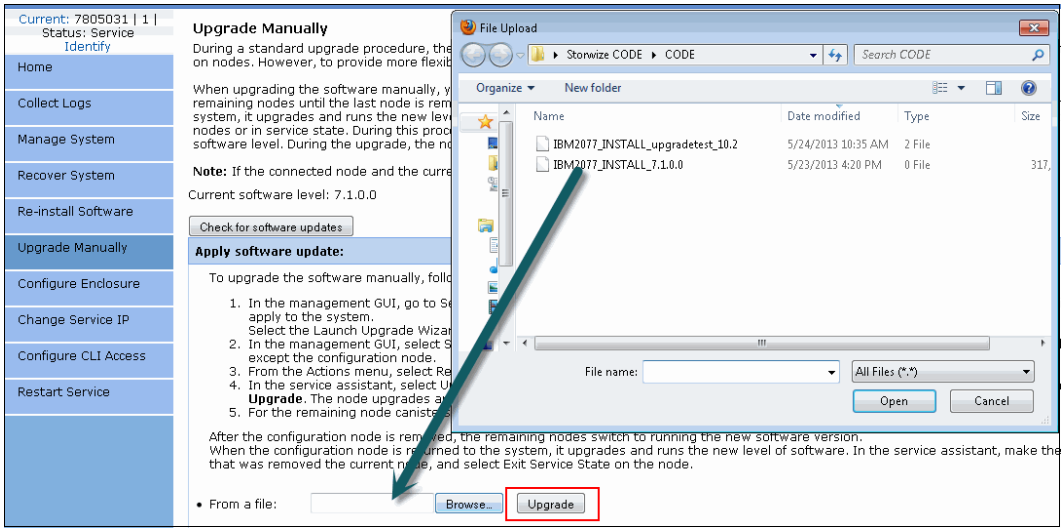

Figure 12-42 Select machine code file for upgrade

5. Click **Upgrade** to start the upgrade process on the first node.

   The node is added automatically into the system after upgrade. Upgrading and adding the node again can take up to 30 minutes, as shown in Figure 12-43.

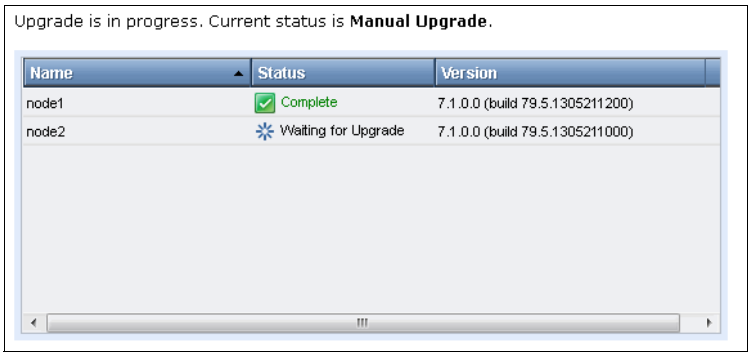

Figure 12-43 Non-config node completed upgrade

6. Repeat steps 2 - 4 for the remaining node (or nodes).

   After you remove the configuration node from the cluster for upgrade, a warning appears, as shown in Figure 12-44.

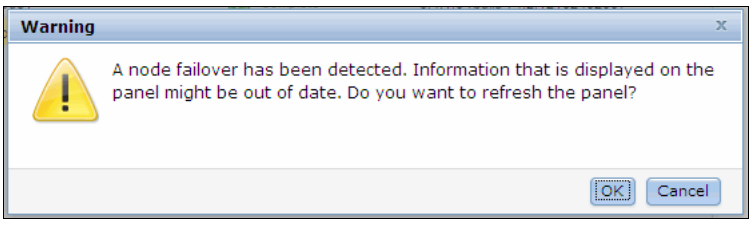

Figure 12-44 Configuration node failover warning

Important: The configuration node remains in Service State when it is added again to the cluster. Therefore, exit Service State manually.

7. To exit from service state, browse to the home panel of the Service Assistant and open the Action menu. Select **Exit Service State**, as shown in Figure 12-45.

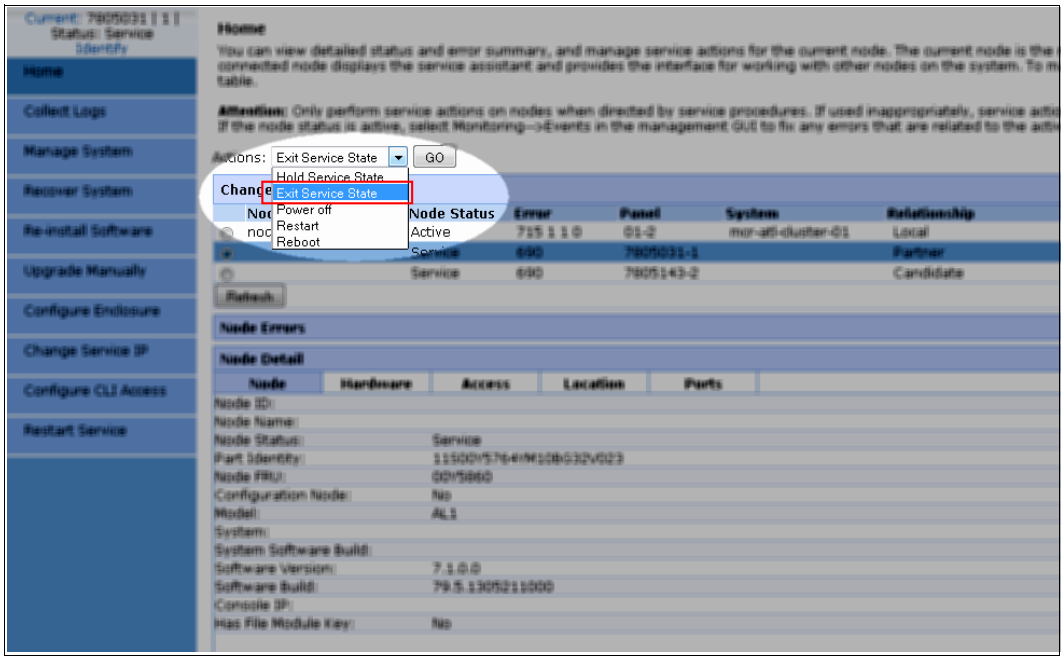

Figure 12-45 Exit service state to add node back in cluster

Both the nodes are now back in the cluster (as shown in Figure 12-46) and the system is running on the new code level.

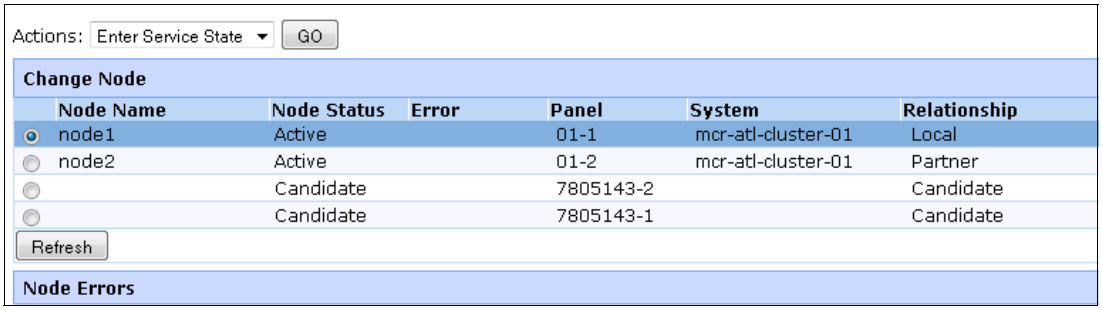

Figure 12-46 Cluster is active again and running new code level

12.5 Event log

Whenever a significant change in the status of IBM Storwize V5000 is detected, an event is submitted to the event log.

All events are classified as *alerts* or *messages*.

An alert is logged when the event requires some action. Some alerts have an associated error code that defines the service action that is required. The service actions are automated through the fix procedures. If the alert does not have an error code, the alert represents an unexpected change in state. This situation must be investigated to see whether it is expected or represents a failure. Investigate an alert and resolve it when it is reported.

A message is logged when a change that is expected is reported; for instance, an IBM FlashCopy operation completes.

The event log panel can be opened via the GUI by clicking **Monitoring** → **Events**, as shown in Figure 12-47.

Figure 12-47 Open eventlog panel

Figure 12-48 shows the event log.

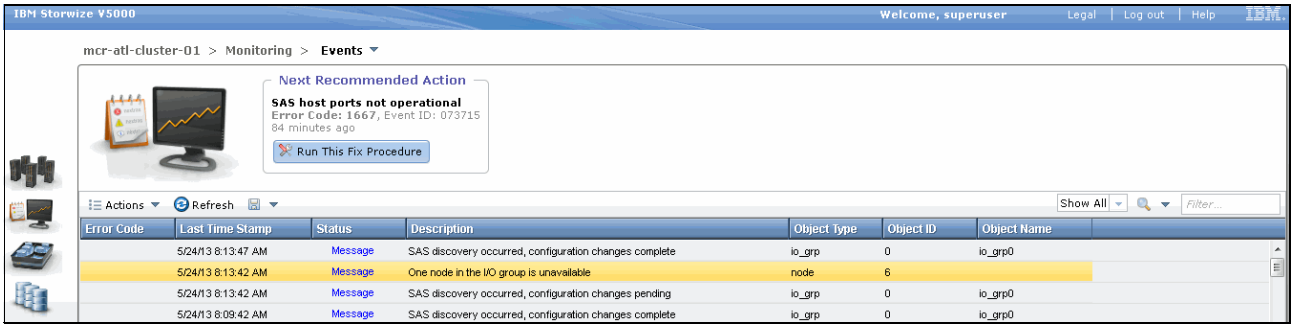

Figure 12-48 The event log view

12.5.1 Managing the event log

The event log features a size limit. After it is full, newer entries replace the older entries, which are not required.

To avoid a repeated event that fills the event log, some records in the event log refer to multiple occurrences of the same event. When event log entries are coalesced in this way, the time stamp of the first occurrence and the last occurrence of the problem is saved in the log entry. A count of the number of times that the error condition occurred also is saved in the log entry. Other data refers to the last occurrence of the event.

Event log panel columns

Right-clicking in any column header opens the option menu in which you can select columns that are shown or hidden.

Figure 12-49 shows all of the possible columns that can be displayed in the error log view.

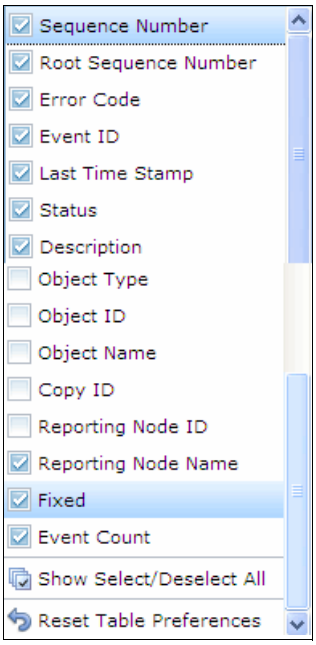

Figure 12-49 Possible event log columns

The following available fields are recommended at a minimum to assist you in diagnosing problems:

- Event ID

  This number precisely identifies the reason why the event was logged.

- Error code

  This number describes the service action that should be followed to resolve an error condition. Not all events have error codes that are associated with them. Many event IDs can have the same error code because the service action is the same for all of the events.

- Sequence number

  A number that identifies the event.

- Event count

  The number of events that are coalesced into this event log record.

- Fixed

  When an alert is shown for an error condition, it indicates whether the reason for the event was resolved. In many cases, the system automatically marks the events that are fixed when appropriate. There are some events that must be manually marked as fixed. If the event is a message, this field indicates that you read and performed the action. The message must be marked as read.

- Last time

  The time when the last instance of this error event was recorded in the log.

- Root sequence number

  If set, this number is the sequence number of an event that represents an error that probably caused this event to be reported. Resolve the root event first.

Event log panel options

Figure 12-50 shows the main Event log panel options, which should be used to handle system events.

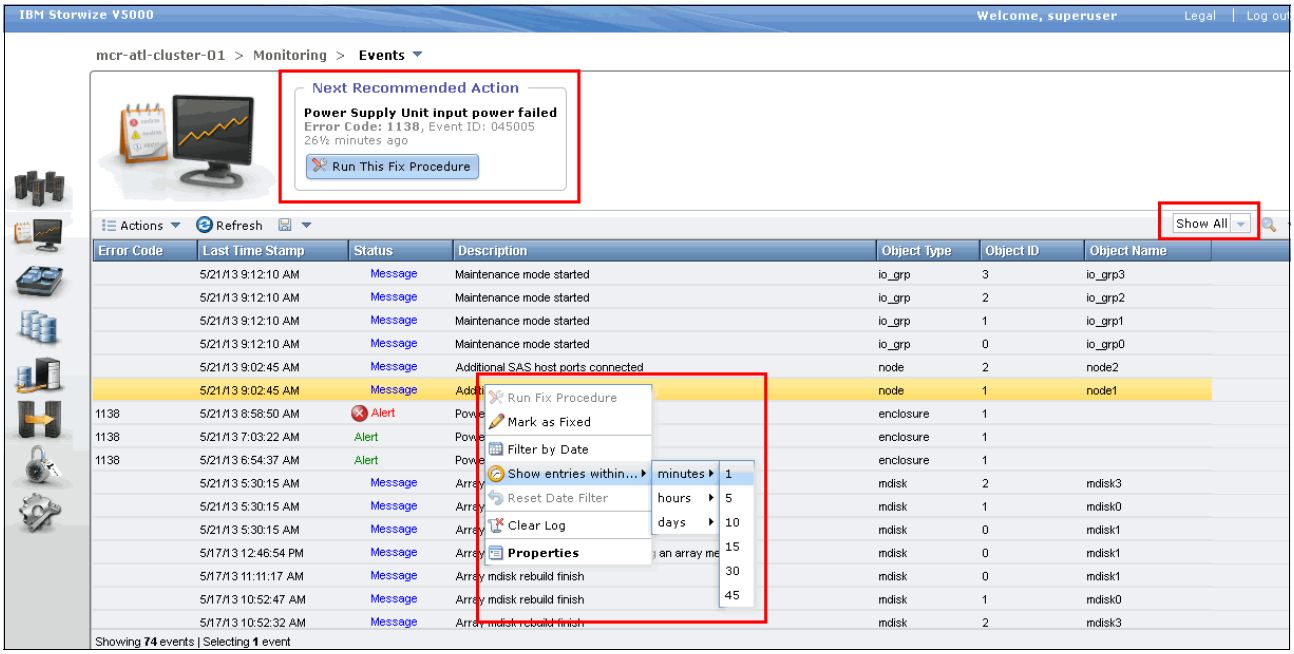

Figure 12-50 Eventlog Panel

Event log filter options

The following log filter options are available:

- Show all

  This option lists all available events.

- Unfixed Messages and Alerts

  This option lists unfixed events. This option is useful to find events that must be handled but no actions are required or recommended.

- Recommended Actions (default)

  Only events with recommended actions (Status Alert) are displayed.

  Important: Check for this filter option if no event is listed. There might be events that are not associated to recommended actions.

Figure 9-51 shows an event log with no items found, which does not necessarily mean that the event log is clear. To check whether the log is clear, use the filter option Show all.

Figure 12-51 No items found in event log

Actions on single event

Right-clicking a single event gives the following options that might be used for that specific event:

- Mark as Fixed

  It is possible to start the Fix Procedure on this specific event, even if it is not the recommended next action.

  Some events, such as, messages, must be set to Mark as Fixed.

- Show entries within... minutes/hours/days

  This option is to limit the error log list to a specific date or a time slot. The following selectable values are available:

  - Minutes: 1, 5, 10, 15, 30, and 45
  - Hours: 1, 2, 5, and 12
  - Days: 1, 4, 7, 15, and 30

- Clear Log

  This option clears the complete error log, even if only one event was selected.

  Important: These actions cannot be undone and might prevent the system from being analyzed when severe problems occur.

- Properties

  This option provides more sense data for the selected event that is shown in the list.

Recommended Actions

A fix procedure is a wizard that helps you to troubleshoot and correct the cause of an error. Some fix procedures reconfigure the system that is based on your responses, ensure that actions are carried out in the correct sequence, and prevent or mitigate loss of data. For this reason, you always must run the fix procedure to fix an error, even if the fix might seem obvious.

To run the fix procedure for the error with the highest priority, go to the Recommended Action panel at the top of the Event page and click **Run This Fix Procedure**. When you fix higher priority events first, the system often can automatically mark lower priority events as fixed.

For more information about how to run a DMP, see 12.5.2, “Alert handling and recommended actions” on page 593.

12.5.2 Alert handling and recommended actions

All events in Alert status require attention. Alerts are listed in priority order and should be fixed sequentially by using the available fix procedures.

Example: SAS cable fault

For this example, we created an error on one SAS cable connection between two expansions by removing the cable from one port.

The following example shows how faults are represented in the error log, how information about the fault can be gathered, and the recommended action (DMP) can be used to fix the error:

- Detect an alert

  The Health Status indicator, which is permanently present on most of the GUI panel (for more information, see Chapter 3, “Graphical user interface overview” on page 75) is showing a yellow alert. Click the indicator to retrieve the specific information, as shown in Figure 12-52.

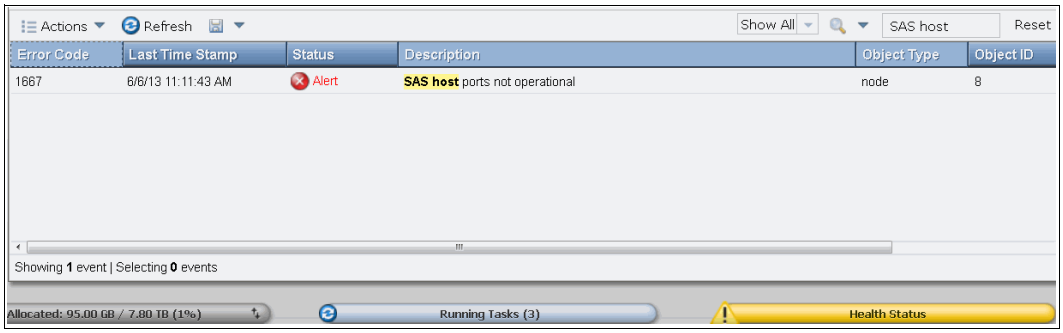

Figure 12-52 Health check shows degraded system status

  Review the event log for more information.

- Find alert in event log

  The default filter in the error log view is Recommended actions. This option lists the alert event only. Figure 12-53 shows the Next Recommended Action list.

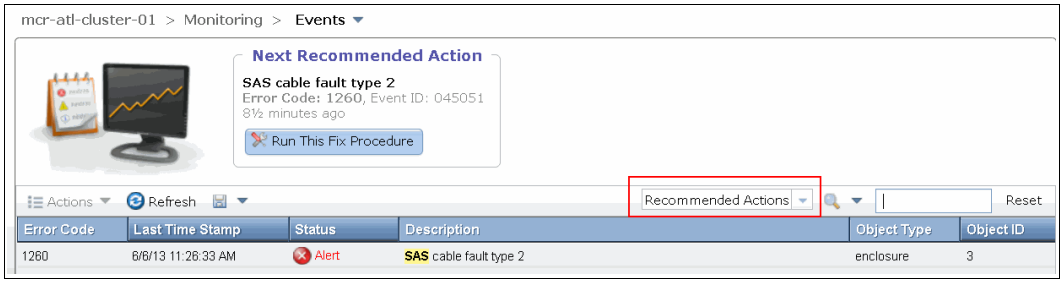

Figure 12-53 Next Recommended Action list

- Gather additional information: Show all

  Find the events that are logged around the alert to understand what happened or find more information for better understanding and to find the original problem. Use the **Show all** filter to see all of the logged events, as shown in Figure 12-54.

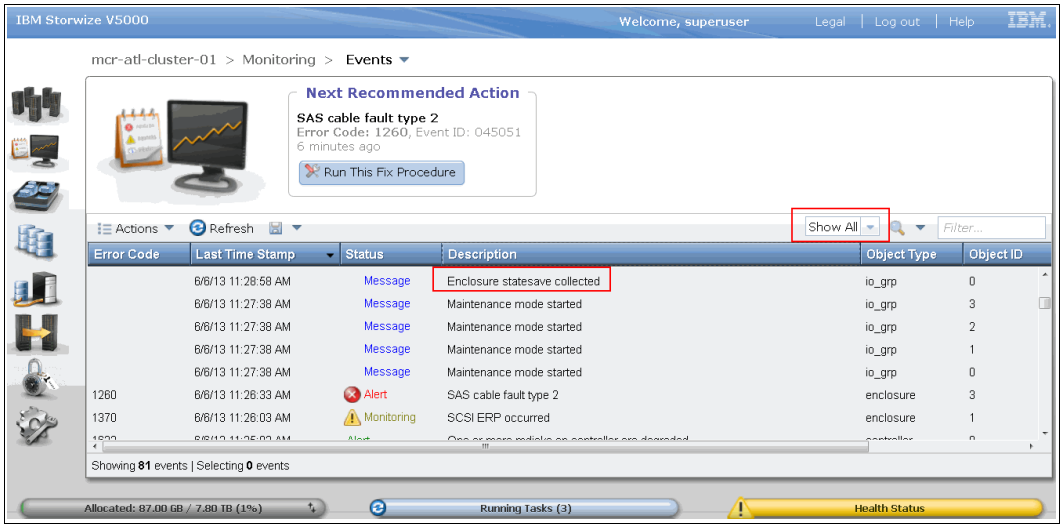

Figure 12-54 Show all events

- Gather additional information: Alert properties

  More details about the event (for example, enclosure ID and canister ID) can be found in the properties option, as shown in Figure 12-55 on page 595. This information might be of interest for problem fixing or for root cause analysis.

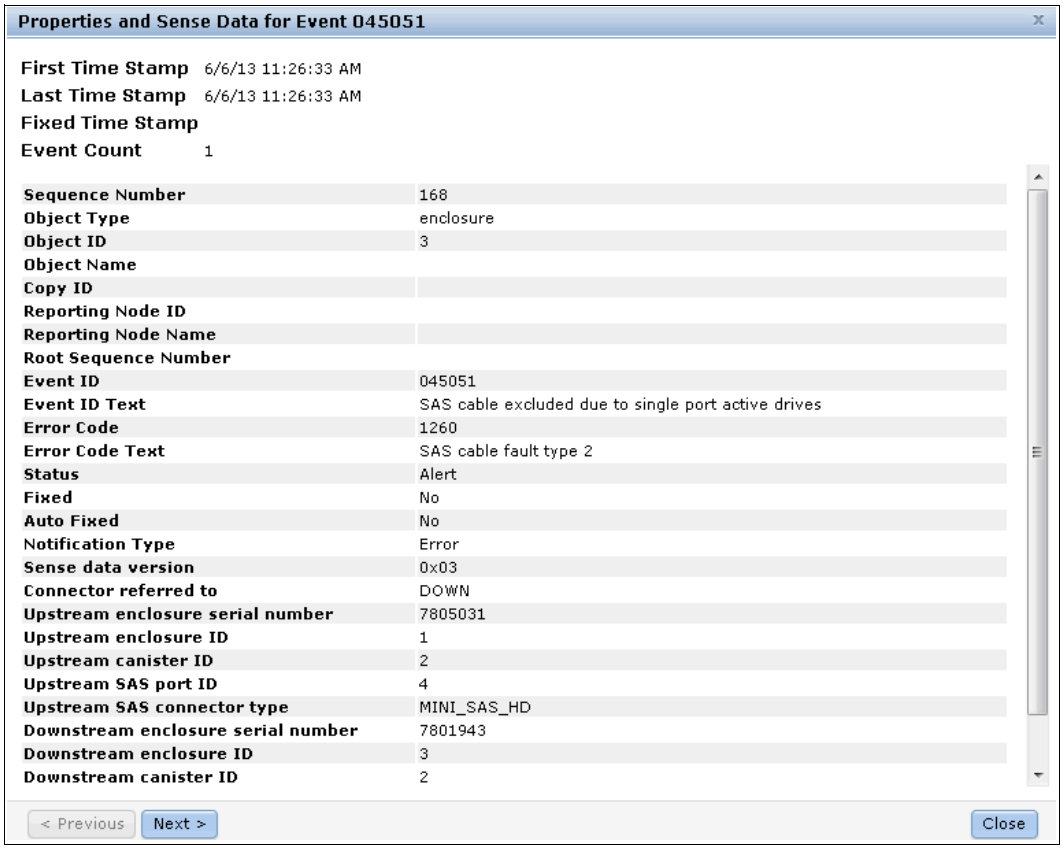

Figure 12-55 Alert properties

- Run recommended action (DMP)

  It is highly recommended to fix alerts under the guidance of the recommended action by using the DMP. There are running tasks in the background that might be missed when the DMP is bypassed. Not all alerts have DMPs available.

  To start the DMP, right-click the alert record or click **Run this fix procedure** at the top of the window.

  The steps and panels of DMP are specific to the error that must be fixed. The following figures represent the recommended action (DMP) for the SAS cable event example.

Figure 12-56 shows step 1 of the DMP SAS cable event.

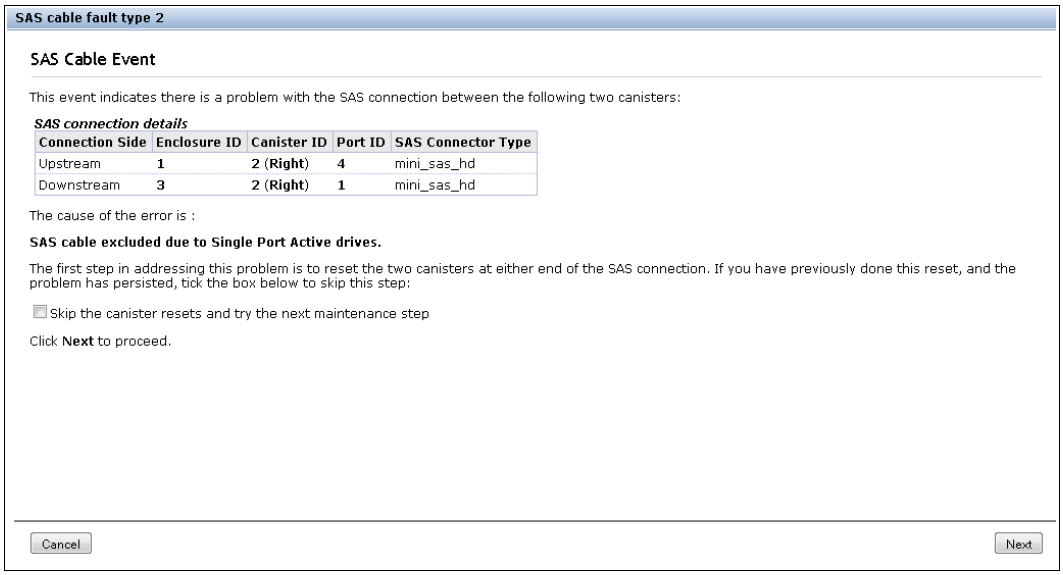

Figure 12-56 SAS cable Recommended action DMP step 1

Figure 12-57 shows step 2 of the DMP SAS cable event.

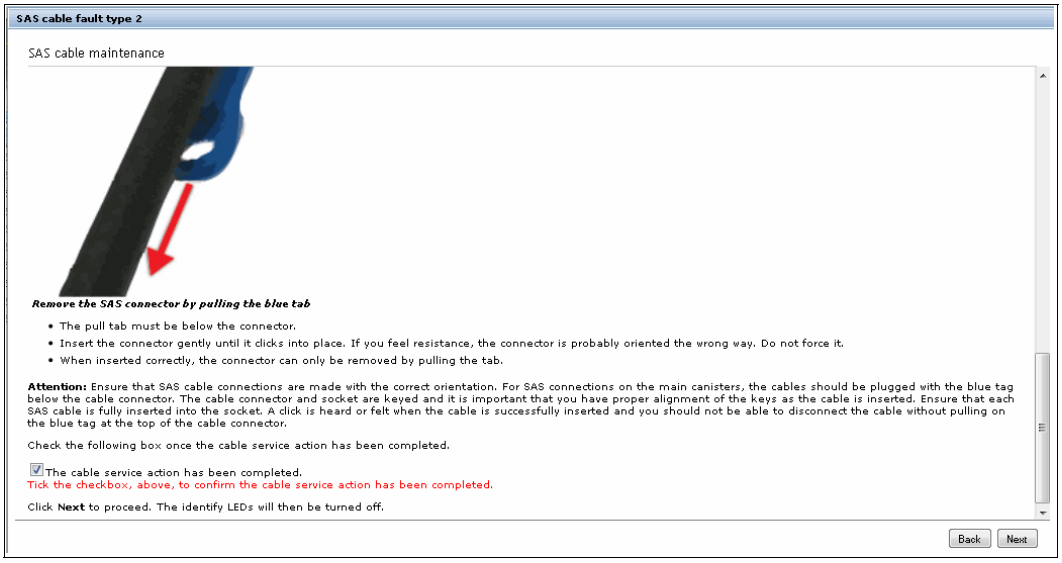

Figure 12-57 SAS cable Recommended action DMP step 2

Figure 12-58 shows step 3 of the DMP SAS cable event.

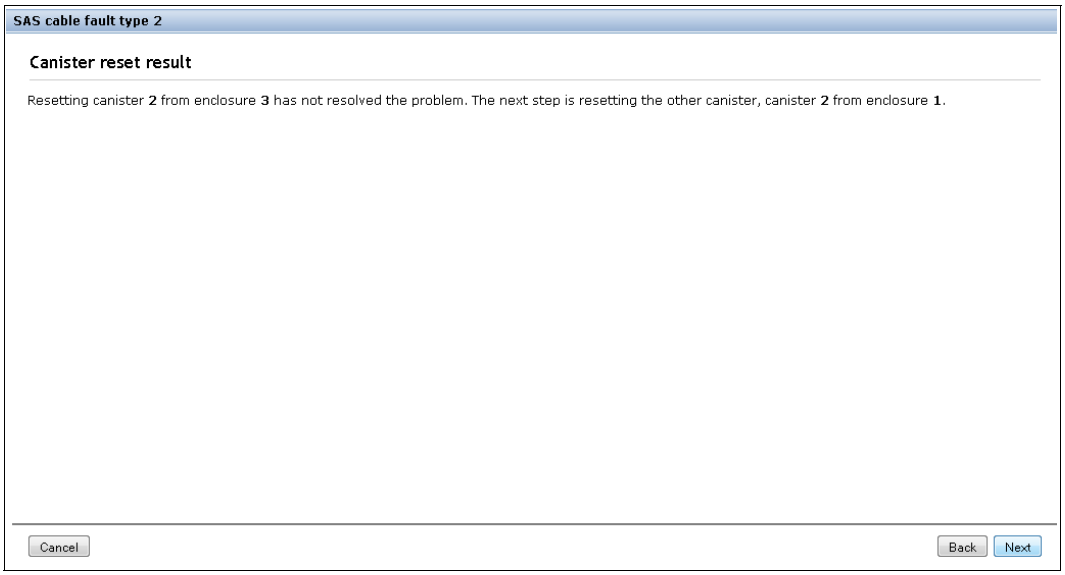

Figure 12-58 SAS cable Recommended action DMP step 3

Figure 12-59 shows step 4 of the DMP SAS cable event.

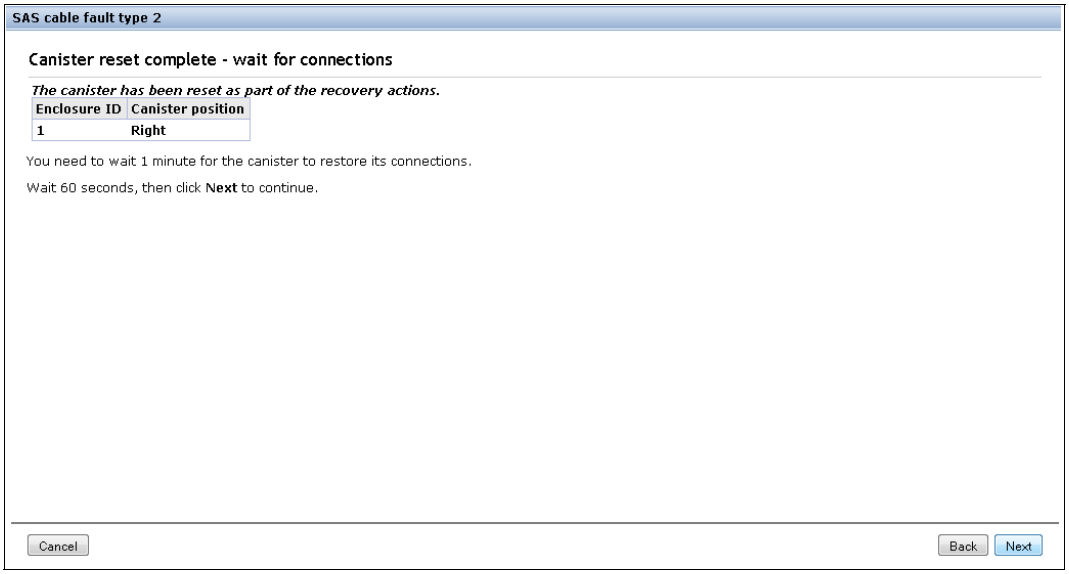

Figure 12-59 SAS cable Recommended action DMP step 4

Figure 12-60 shows step 5 of the DMP SAS cable event.

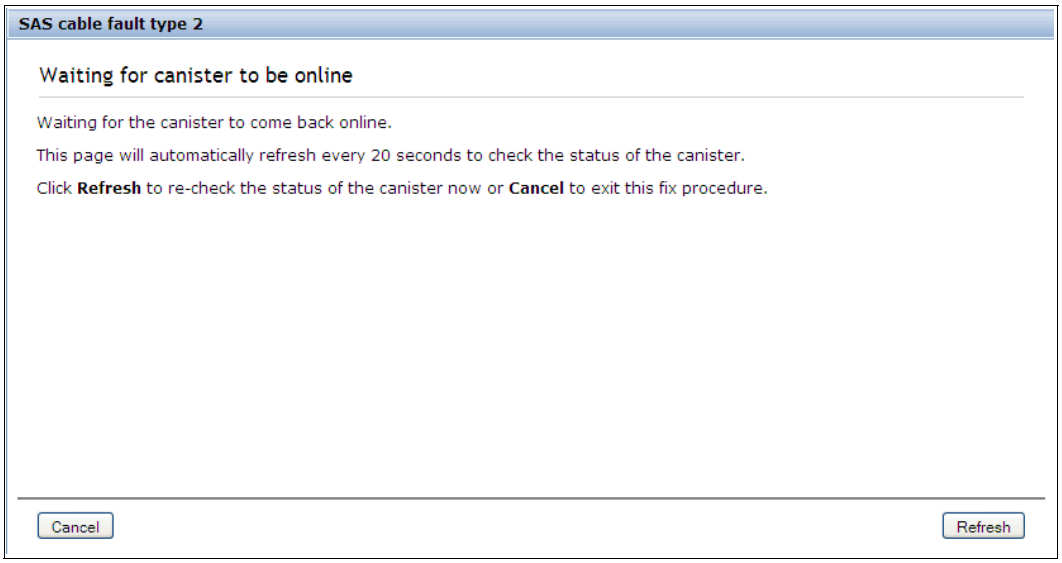

Figure 12-60 SAS cable Recommended action DMP step 5

Figure 12-61 shows step 6 of the DMP SAS cable event.

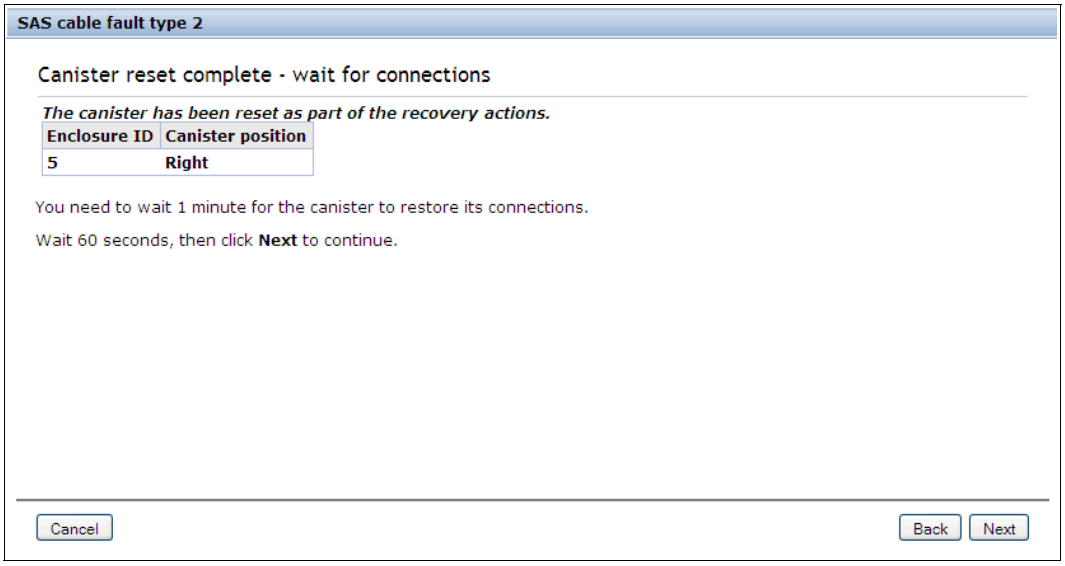

Figure 12-61 SAS cable Recommended action DMP step 6

Figure 12-62 shows step 7 of the DMP SAS cable event.

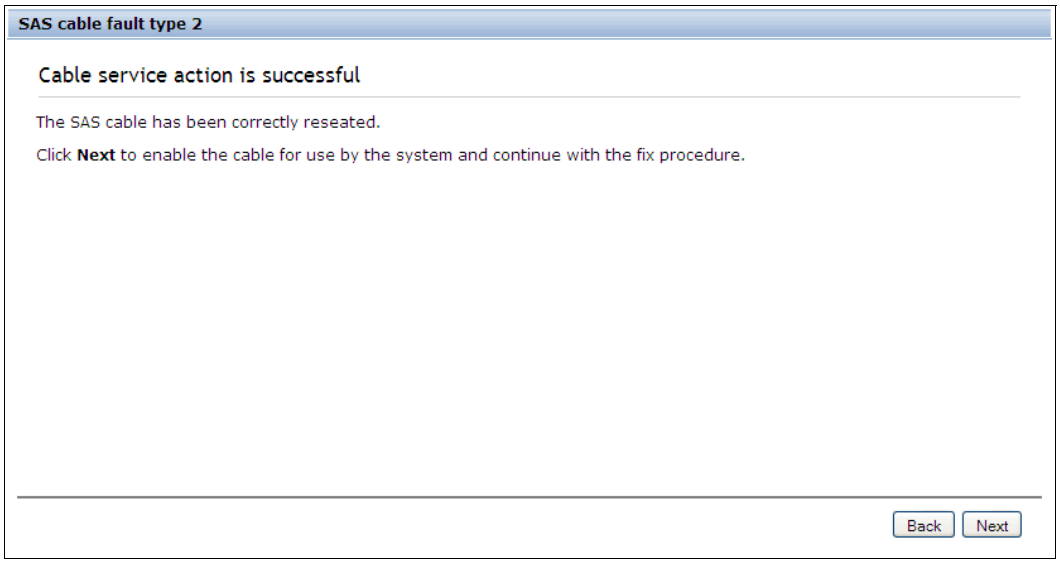

Figure 12-62 SAS cable Recommended action DMP step 7

Figure 12-63 shows step 8 of the DMP SAS cable event.

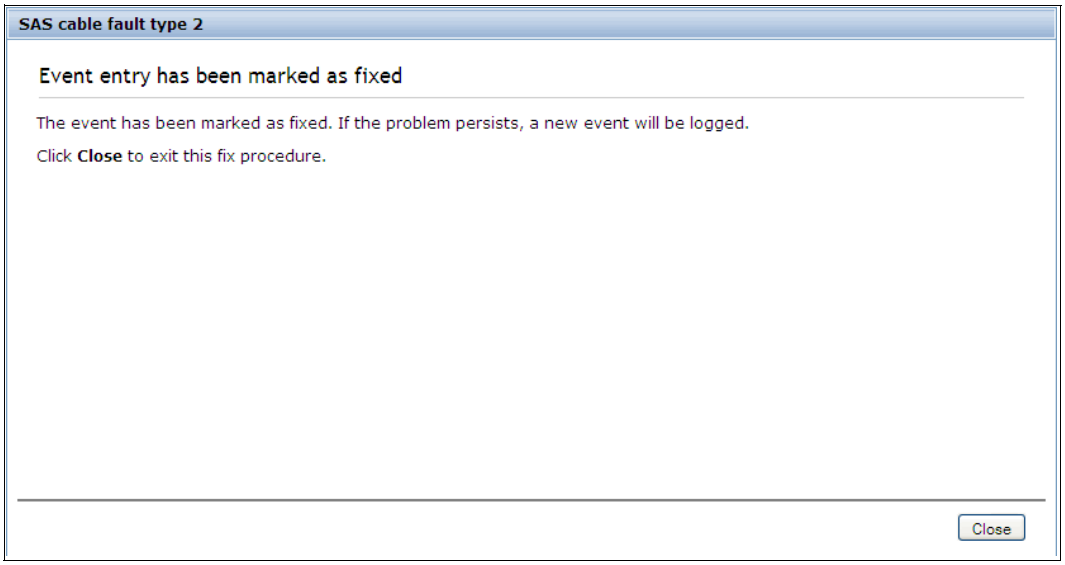

Figure 12-63 SAS cable Recommended action DMP step 8

When all of the steps of the DMP are processed successfully, the recommended action is complete and the problem should be fixed. Figure 12-64 on page 600 shows the red color of the event status changed to green. The system health status is green and there are no other that must be addressed.

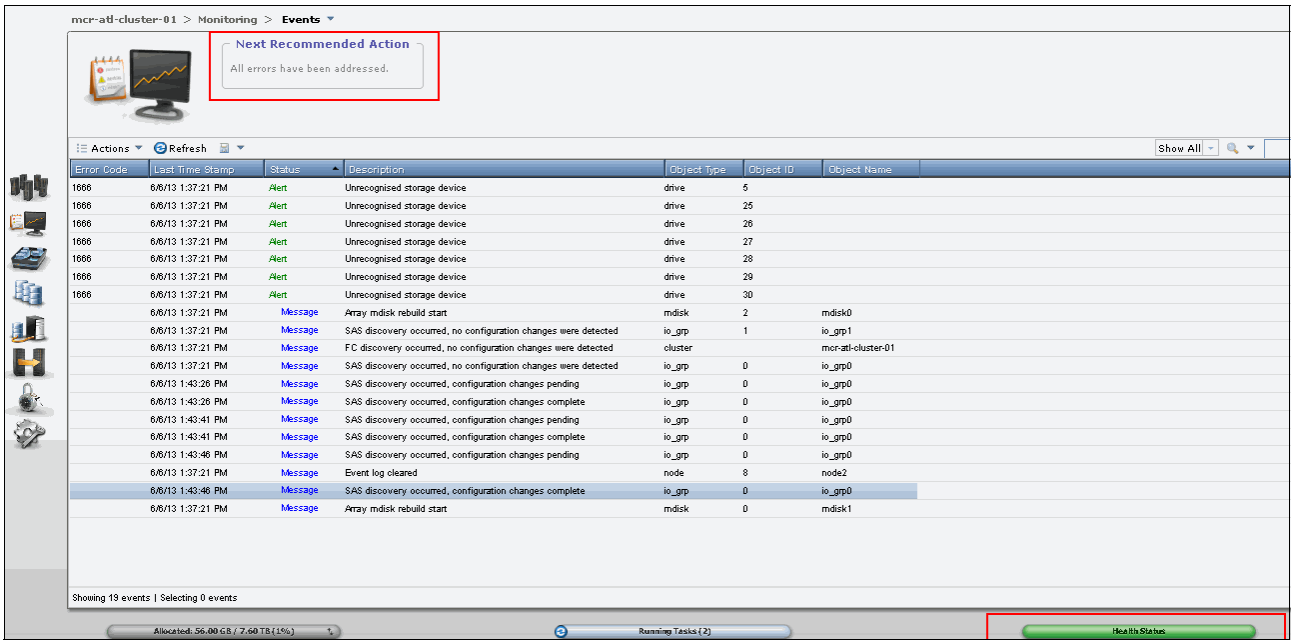

Figure 12-64 Recommended action that is completed

Handling multiple alerts

If there are multiple alerts that are logged, the IBM Storwize V5000 recommends a next action to fix the problem (or problems).

Figure 12-65 shows the event log that displays multiple alert.

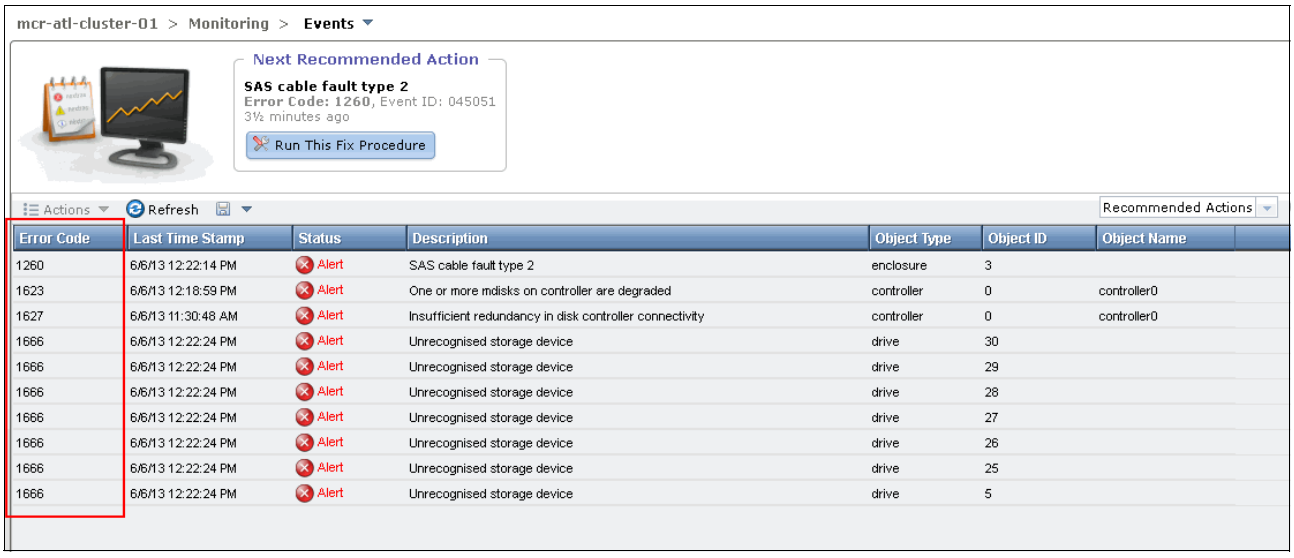

Figure 12-65 Multiple alert events that are displayed in the event log

The Next Recommended Action function orders the alerts by severity and displays the events with the highest severity first. If multiple events have the same severity, they are ordered by date and the oldest event is displayed first.

The following order of severity starts with the most severe condition:

- Unfixed alerts (sorted by error code; the lowest error code has the highest severity)
- Unfixed messages
- Monitoring events (sorted by error code; the lowest error code has the highest severity)
- Expired events
- Fixed alerts and messages

Faults are often fixed with the fixture of the most severe fault.

12.6 Collecting support information

If you have a problem and call the IBM Support Center, you might be asked to provide support data, as described in the next section.

12.6.1 Support information via GUI

Complete the following steps to collect support information by using the GUI:

1. Click **Settings** and then the **Support** tab (as shown in Figure 12-66) to begin the procedure of collecting support data.

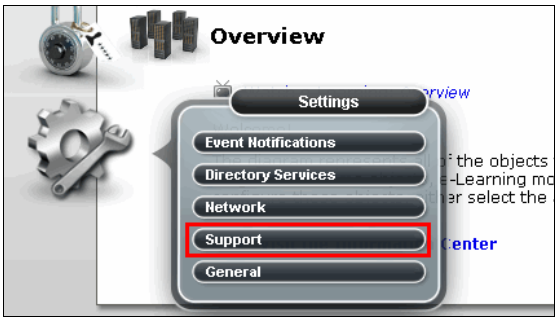

Figure 12-66 Support Files VIA® GUI

2. Click **Download Support Package**, as shown in Figure 12-67.

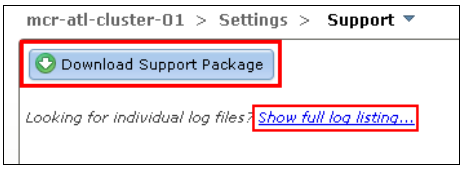

Figure 12-67 Download Support Package

The panel that is shown in Figure 12-68 opens and you can select one of four different versions of the `svc_snap` support package.

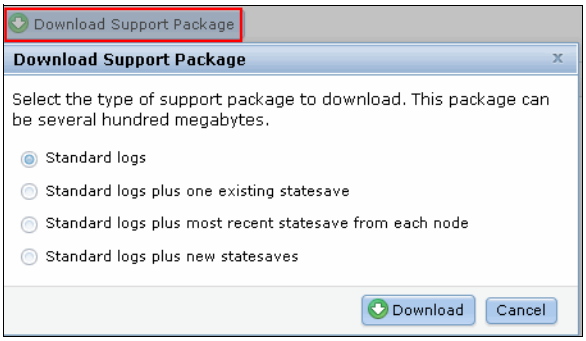

Figure 12-68 Support Package Selection

The version that you download that depends on the event that you are investigating. For example, if you noticed in the event log that a node was restarted, capture the snap with the latest existing statesaves.

The following components are included in the support package:

- Standard logs

  Contains the most recent logs that were collected from the system. These logs are most commonly used by Support to diagnose and solve problems.

- Standard logs plus one existing statesave

  Contains the standard logs from the system and the most recent statesave from any of the nodes in the system. Statesaves are also known as *memory dumps* or *live memory dumps*.

- Standard logs plus most recent statesave from each node

  This option is used most often by the support team for problem analysis. They contain the standard logs from system and the most recent statesave from each node in the system.

- Standard logs plus new statesave

  This option might be requested by the Support team for problem determination. It generates a new statesave (livedump) for all of the nodes and packages them with the most recent logs.

Save the resulting snap file in a directory for later use or to upload to IBM support.

12.6.2 Support information via Service Assistant

The IBM Storwize V5000 management GUI collects information from all the components in the system. The Service Assistant collects information from all node canisters. The snap file is the information that is collected and packaged in a single file.

If the package is collected by using the Service Assistant, ensure that the node from which the logs are collected is the current node, as shown in Figure 12-69.

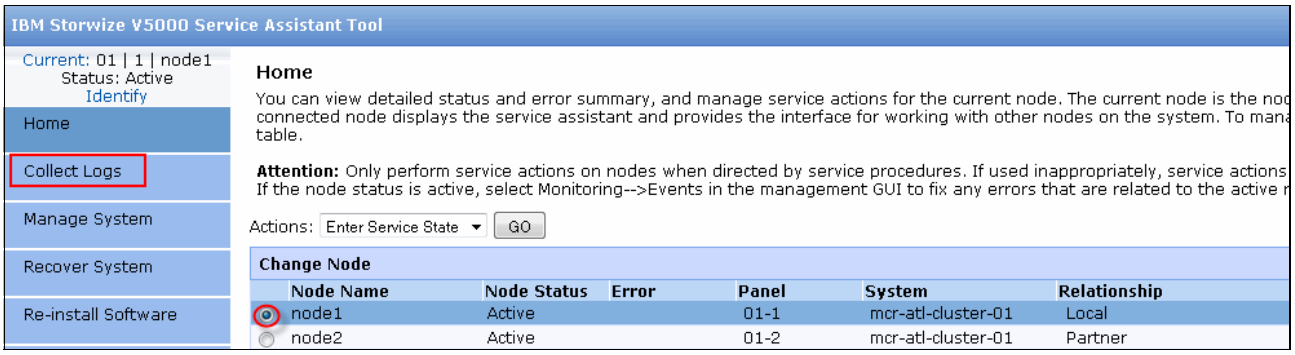

Figure 12-69 Collect logs with Service Assistance

Support information can be downloaded with or without the latest statesave, as shown in Figure 12-70.

Figure 12-70 Download support file via Service Assistant

12.6.3 Support Information onto USB stick

Whenever GUI, Service Assistant, or a remote connection is unavailable, snaps can be collected from each single node by using the USB stick.

Complete the following steps to collect snaps by using the USB stick:

1. Create a text file that includes the following command:

   `satask snap -dump`

2. Save the file as `satask.txt` in the root directory of the USB stick.
3. Insert the USB stick in the USB port of the node from which the support data should be collected.
4. Wait until no write activities are recognized (this process can take 10 minutes or more).

5. Remove the USB stick and check the results, as shown in Figure 12-71.

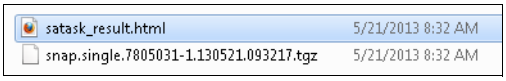

Figure 12-71 Single snap result files on USB stick

satask_result file

The `satask_result.html` file is the general response to the command that is issued via the USB stick. If the command did not run successfully, it is noticed in this file. Otherwise, any general system information is stored here, as shown in Figure 12-72.

satask snap -dump

```
Tue May 21 09:32:12 BST 2013
```

System Status

sainfo lsservicenodes

```
panel_name cluster_id       cluster_name       node_id node_name relation  node_status error_data
01-1       000000C0200026C0 mcr-atl-cluster-01 1       node1     local     Active
01-2       000000C0200026C0 mcr-atl-cluster-01 2       node2     partner   Active
```

Figure 12-72 satask_result.txt on USB stick (header only)

Snap memory dump on USB

A complete statesave of the node where the USB was attached is stored in a `.zip` file. The name of the file includes the node name and the time stamp. The content of the `.zip` file is shown in Figure 12-73.

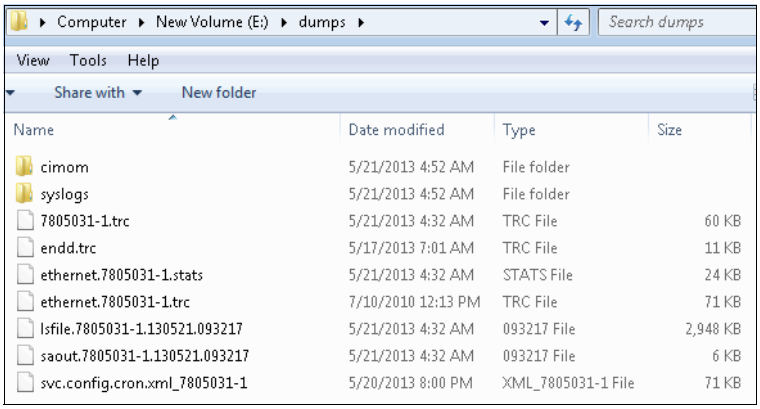

Figure 12-73 Single snap memory dump on USB stick

12.7 Powering on and shutting down IBM Storwize V5000

In the following sections, we describe the process to power on and shut down the IBM Storwize V5000 system by using the GUI and the CLI.

12.7.1 Shutting down the system

In this section, we show how to shut down the IBM Storwize V5000 system by using the GUI and CLI.

Important: You should never shut down your IBM Storwize V5000 by powering off the PSUs, removing both PSUs, or removing both power cables from a running system.

Powering down by using the GUI

You can shut down only one node canister or the entire cluster. When you shut down only one node canister, all of the activities remain active. When you shut down the entire cluster, you must power on locally to restart the system.

To shut down by using the GUI, complete the following steps:

1. Browse to the Monitoring function icon (as shown in Figure 12-74) and click **System Details**.

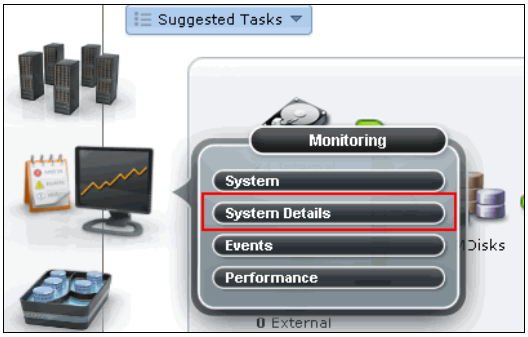

Figure 12-74 Power down via system details

2. Select the root level of the system detail tree, click **Actions**, and then select **Shut Down System**, as shown in Figure 12-75.

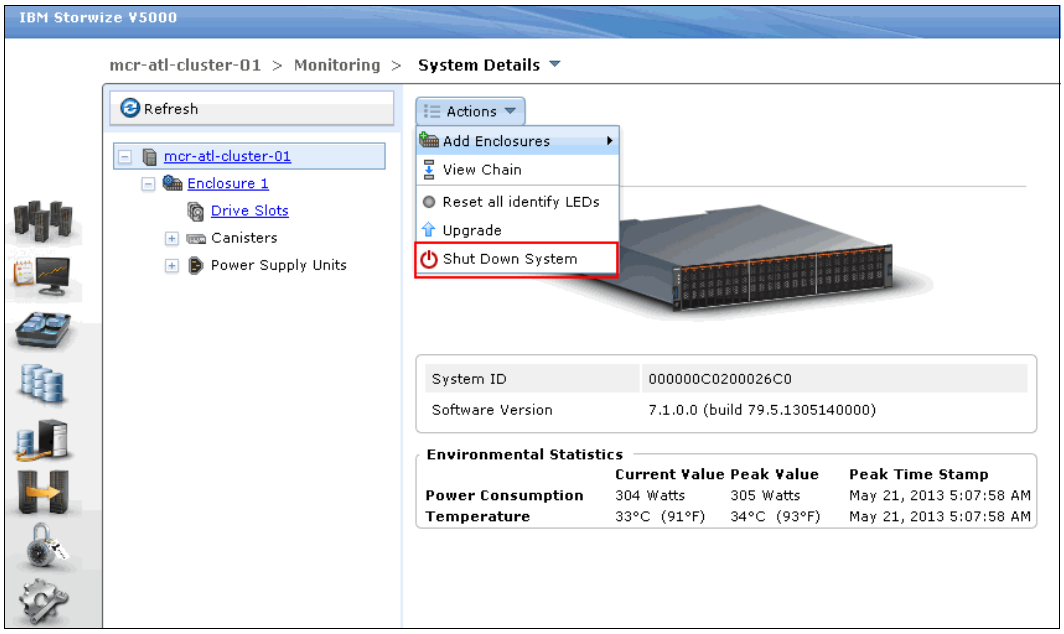

Figure 12-75 Power Down System option

The following process can be used as an alternative to steps 1 and 2, as shown in Figure 12-75:

a. Browse to the Monitoring navigator and open the System view.

b. Click the System that is under the system display.

   An information panel opens.

c. Click the **Manage** tab.

d. Click **Shut Down System** to shut down, as shown in Figure 12-76.

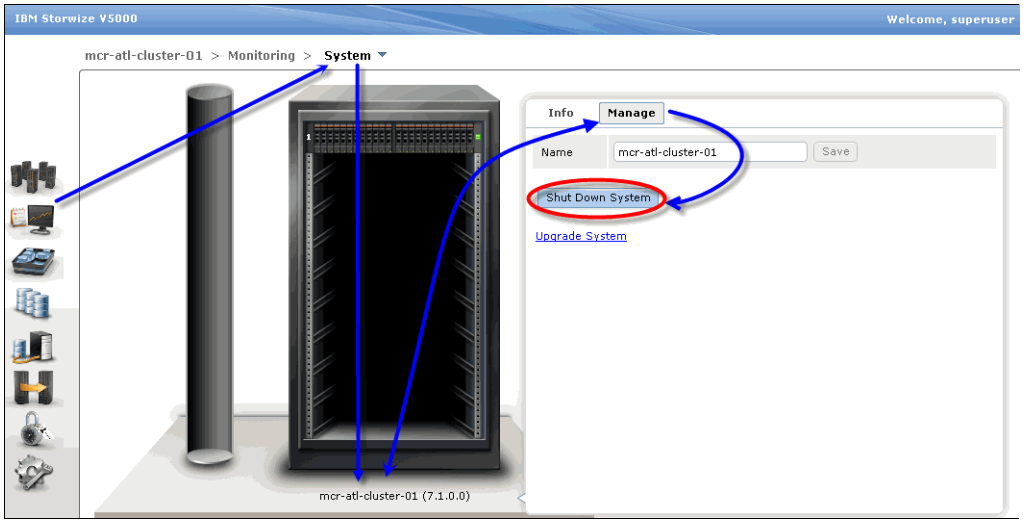

Figure 12-76 Shut down system via Monitoring system GUI

3. The Confirm System Shutdown window opens. A message opens and prompts you to confirm whether you want to shut down the cluster. Ensure that you stopped all FlashCopy mappings, data migration operations, and forced deletions before you continue. Enter `Yes` and click **OK** to begin the shutdown process, as shown in Figure 12-77.

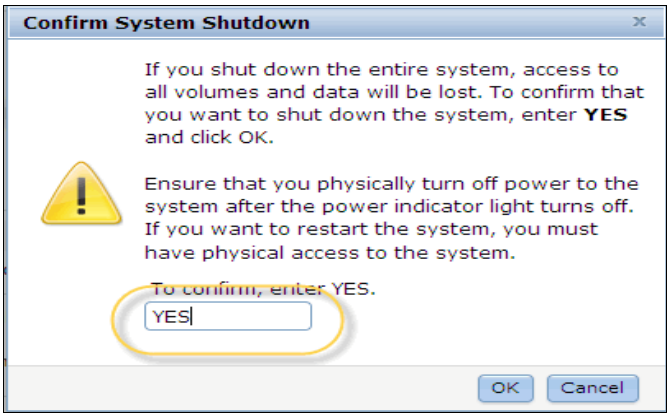

Figure 12-77 Shut Down confirmation

4. Wait for the power LED on both node canisters in the control enclosure to flash at 1 Hz, which indicates that the shutdown operation completed (1 Hz is half as fast as the drive indicator LED).

Tip: When you shut down an IBM Storwize V5000, it does not automatically restart. You must manually restart the system.

Shutting down by using the CLI

The CLI is the other option that can be used to shut down an IBM Storwize V5000. The CLI is accessed via the PuTTY utility.

Warning: If you are shutting down the entire system, you lose access to all volumes that are provided by this system. Shutting down the system also shuts down all IBM Storwize V5000 nodes. This shutdown causes the hardened data to be dumped to the internal HDD.

Run the **`stopsystem`** command to shut down a clustered system, as shown in Example 12-3.

Example 12-3 Shut down

```
stopsystem

Are you sure that you want to continue with the shut down?

# Type y to shut down the entire clustered system.
```

12.7.2 Powering on

Complete the following steps to power on the system:

Important: This process assumes that all power is removed from the enclosure. If the control enclosure is shut down but the power is not removed, the power LED on all node canisters flash at a rate of half of one second on, half of one second off. In this case, remove the power cords from both power supplies and then reinsert them.

1. Ensure that any network switches that are connected to the system are powered on.
2. Power on any expansion enclosures by connecting the power cord to both power supplies in the rear of the enclosure or turning on the power circuit.
3. Power on the control enclosure by connecting the power cords to both power supplies in the rear of the enclosure and turning on the power circuits.

   The system starts. The system starts successfully when all node canisters in the control enclosure have their status LED permanently on, which should take no longer than 10 minutes.
4. Start the host applications.

A

Command-line interface setup and SAN Boot

This appendix describes the setup of the command-line interface (CLI) and provides more information about the SAN Boot function.

This appendix includes the following sections:

- Command-line interface
- SAN Boot

Command-line interface

The IBM Storwize V5000 system has a powerful CLI, which offers even more functions than the GUI. This section is not intended to be a detailed guide to the CLI because that topic is beyond the scope of this book. The basic configuration of the IBM Storwize V5000 CLI and some example commands are covered. However, the CLI commands are the same as in the SAN Volume Controller. In addition, there are more commands that are available to manage internal storage. If a task is completed in the GUI, the CLI command always is displayed in the details, as shown throughout this book.

Detailed CLI information is available in the IBM Storwize V5000 Information Center under the Command Line section, which can be found at this website:

http://pic.dhe.ibm.com/infocenter/storwize/V5000_ic/index.jsp?topic=%2Fcom.ibm.storwize.V5000.641.doc%2Fsvc_clicommandscontainer_229g0r.html

Implementing the IBM Storwize V7000 V6.3, SG24-7938, also has information about the use of the CLI. The commands in that book also apply to the IBM Storwize V5000 system because it is part of the Storwize family.

Basic setup

In the IBM Storwize V5000 GUI, authentication is done by using a user name and a password. The CLI uses a Secure Shell (SSH) to connect from the host to the IBM Storwize V5000 system. A private and public key pair or user name and password is necessary. The following steps are required to enable CLI access with SSH keys:

1. A public key and private key are generated as a pair.
2. A public key is uploaded to the IBM Storwize V5000 system by using the GUI.
3. A client SSH tool is configured to authenticate with the private key.
4. A secure connection is established between the client and IBM Storwize V5000 system.

Secure Shell is the communication vehicle that is used between the management workstation and the IBM Storwize V5000 system. The SSH client provides a secure environment from which to connect to a remote machine. It uses the principles of public and private keys for authentication.

SSH keys are generated by the SSH client software. The SSH keys include a public key, which is uploaded and maintained by the clustered system, and a private key, which is kept private on the workstation that is running the SSH client. These keys authorize specific users to access the administration and service functions on the system. Each key pair is associated with a user-defined ID string that can consist of up to 40 characters. Up to 100 keys can be stored on the system. New IDs and keys can be added, and unwanted IDs and keys can be deleted. To use the CLI, an SSH client must be installed on that system, the SSH key pair must be generated on the client system, and the client's SSH public key must be stored on the IBM Storwize V5000 systems.

The SSH client that is used in this book is PuTTY. There also is a PuTTY key generator that can be used to generate the private and public key pair. The PuTTY client can be downloaded at no cost at the following website:

http://www.chiark.greenend.org.uk

The following tools should be downloaded:

- PuTTY SSH client: `putty.exe`
- PuTTY key generator: `puttygen.exe`

Generating a public and private key pair

To generate a public and private key pair, complete the following steps:

1. Start the PuTTY key generator to generate the public and private key pair, as shown in Figure A-1.

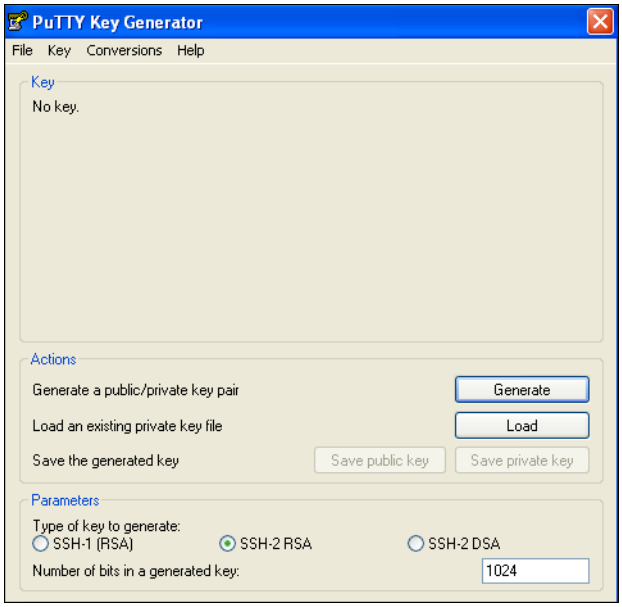

Figure A-1 PuTTY key generator

Make sure that the following options are selected:

- Type of key to generate: SSH2 RSA
- Number of bits in a generated key: 1024

2. Click **Generate** and move the cursor over the blank area to generate the keys, as shown in Figure A-2.

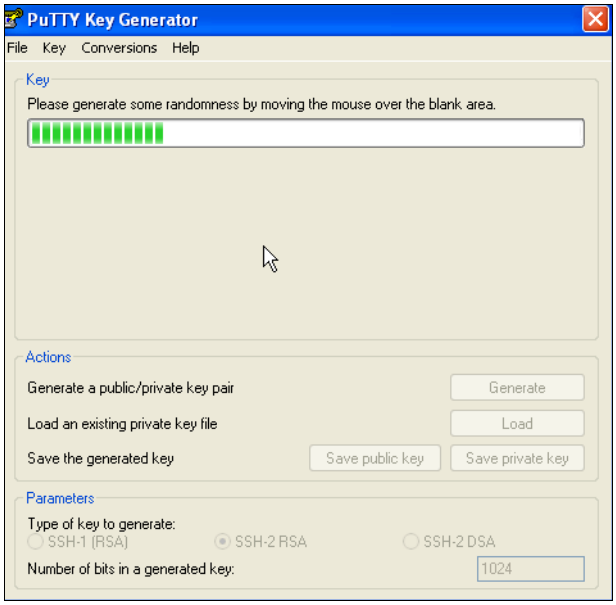

Figure A-2 Generate keys

Generating keys: The blank area that is indicated by the message is the large blank rectangle on the GUI inside the section of the GUI labeled Key. Continue to move the mouse pointer over the blank area until the progress bar reaches the far right side. This action generates random characters to create a unique key pair.

3. After the keys are generated, save them for later use. Click **Save public key**, as shown in Figure A-3.

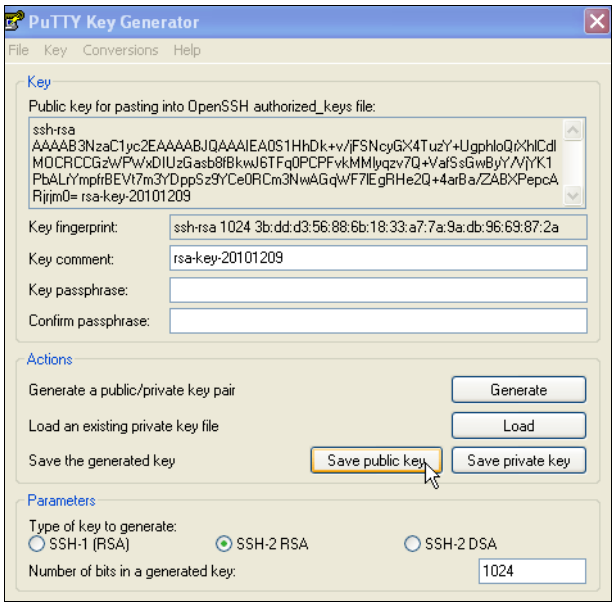

Figure A-3 Save public key

4. You are prompted for a name (for example, `pubkey`) and a location for the public key (for example, `C:\Support Utils\PuTTY`). Click **Save**.

   Be sure to record the name and location of this SSH public key because this information must be specified later.

Public key extension: By default, the PuTTY key generator saves the public key with no extension. Use the string `pub` for naming the public key; for example, `pubkey`, to differentiate the SSH public key from the SSH private key.

5. Click **Save private key**, as shown in Figure A-4.

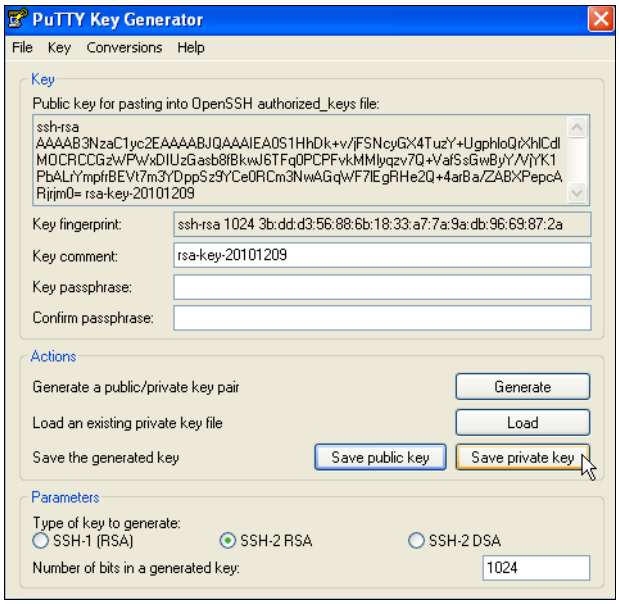

Figure A-4 Save private key

6. You receive a warning message, as shown in Figure A-5. Click **Yes** to save the private key without a passphrase.

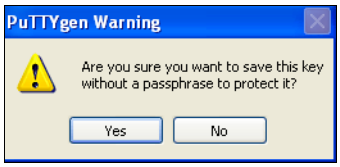

Figure A-5 Confirm the security warning

7. When prompted, enter a name (for example, `icat`), select a secure place as the location, and click **Save**.

> **Key generator:** The PuTTY key generator saves the private key with the PPK extension.

8. Close the PuTTY key generator.

Uploading the SSH public key to the IBM Storwize V5000

After you create your SSH key pair, you must upload your SSH public key onto the SAN Volume Controller system. Complete the following steps to upload the key:

1. Open the user section, as shown in Figure A-6.

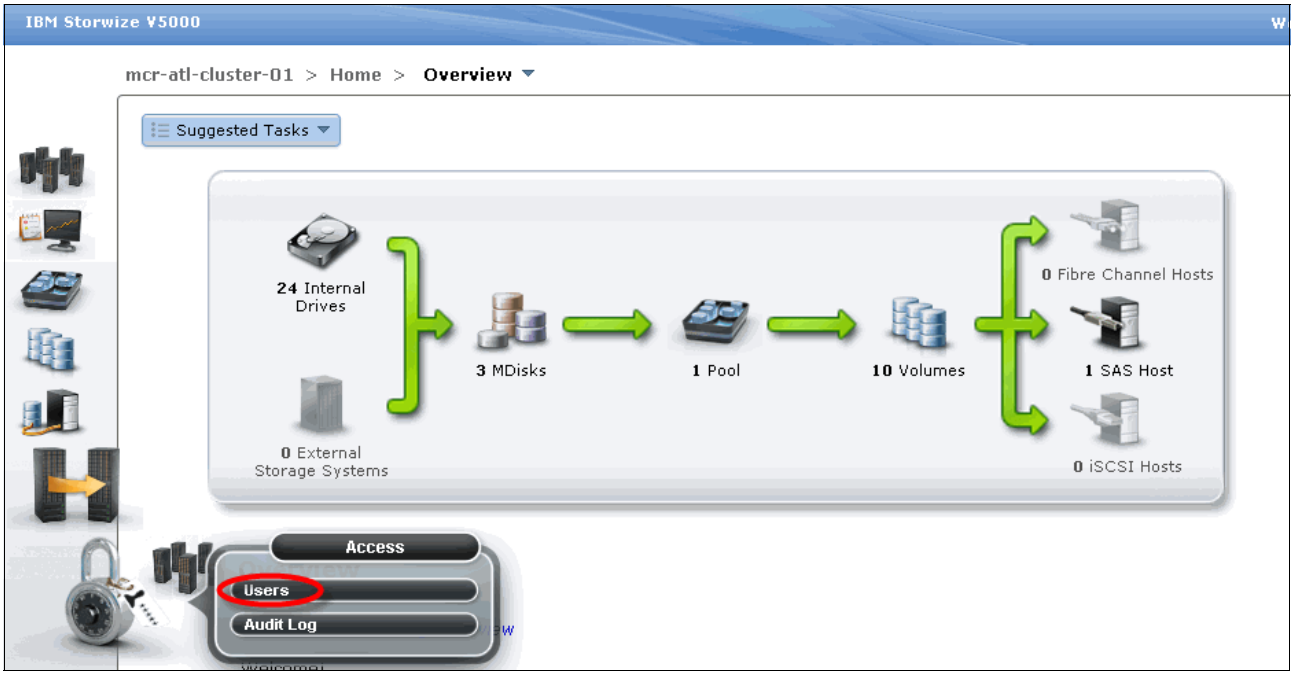

Figure A-6 Open user section

2. Right-click the user for which you want to upload the key and click **Properties**, as shown in Figure A-7.

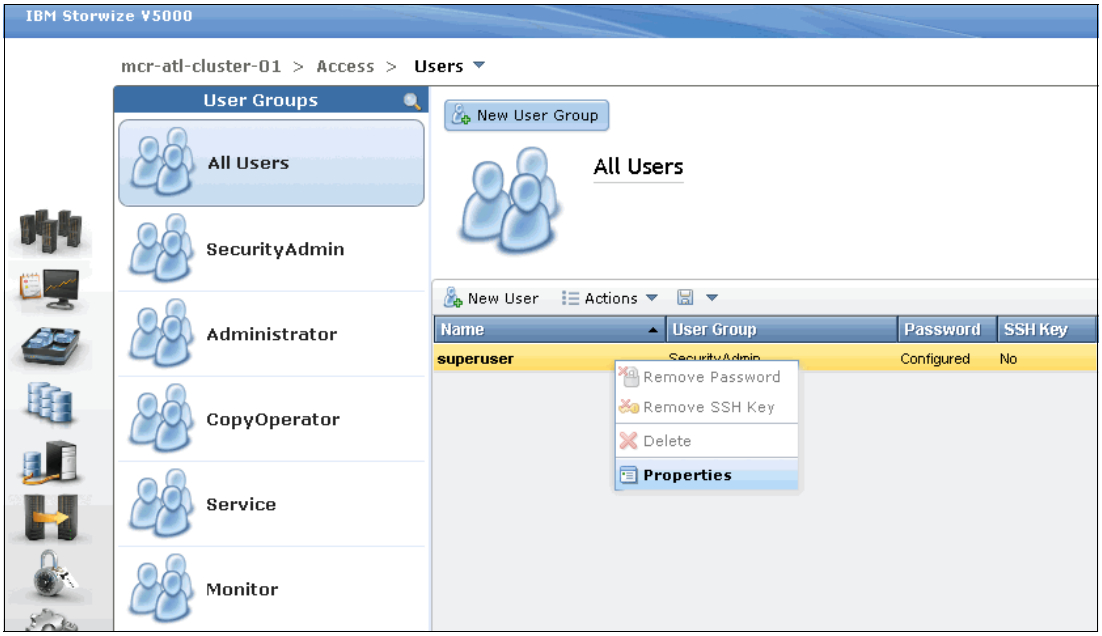

Figure A-7 Superuser properties

3. To upload the public key, click **Browse**, select your public key, and click **OK**, as shown in Figure A-8.

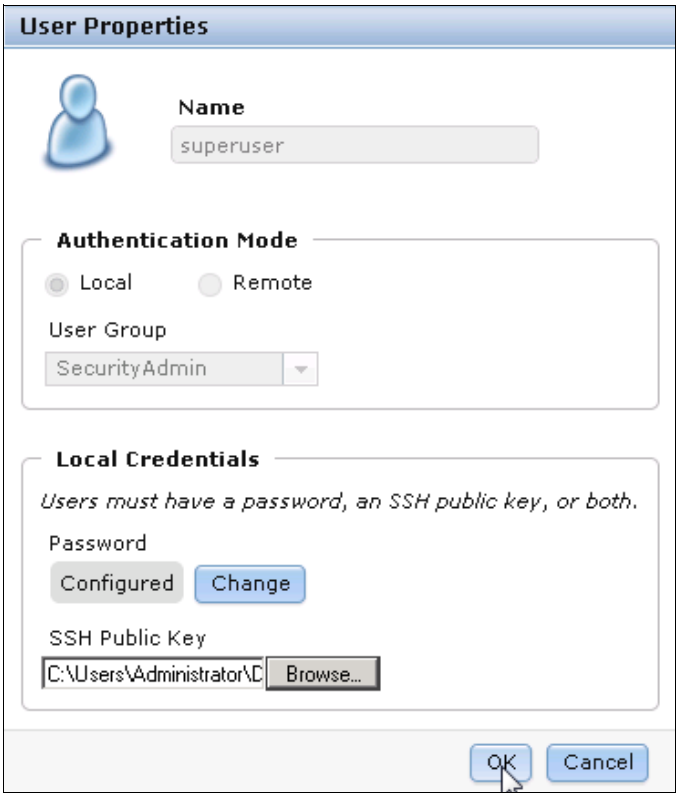

Figure A-8 Select public key

4. Click **OK** and the key is uploaded, as shown in Figure A-9.

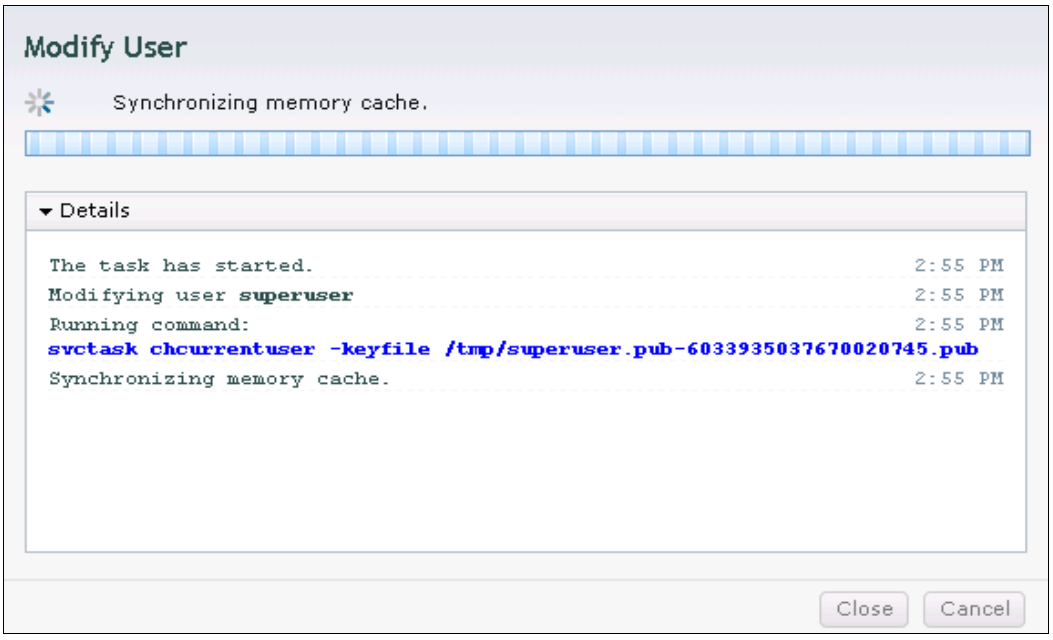

Figure A-9 Public key upload complete

5. Click **Close** to return to the GUI.

Configuring the SSH client

Before the CLI can be used, the SSH client must be configured. Complete the following steps to configure the client:

1. Start PuTTY, as shown in Figure A-10.

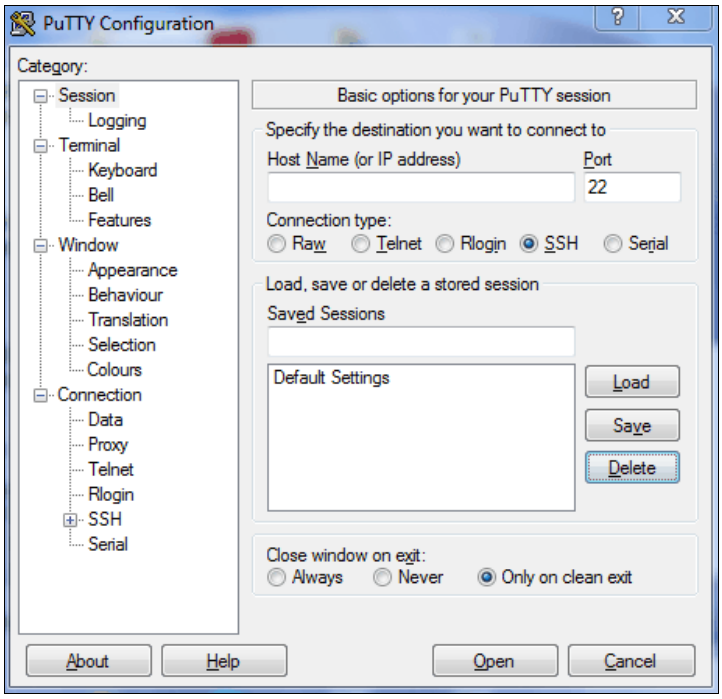

Figure A-10 PuTTY

In the right side pane under the “Specify the destination you want to connect to” section, select **SSH**. Under the “Close window on exit” section, select **Only on clean exit**, which ensures that if there are any connection errors, they are displayed in the user’s window.

2. From the Category pane on the left side of the PuTTY Configuration window, click **Connection** → **SSH** to open the PuTTY SSH Configuration window, as shown in Figure A-11.

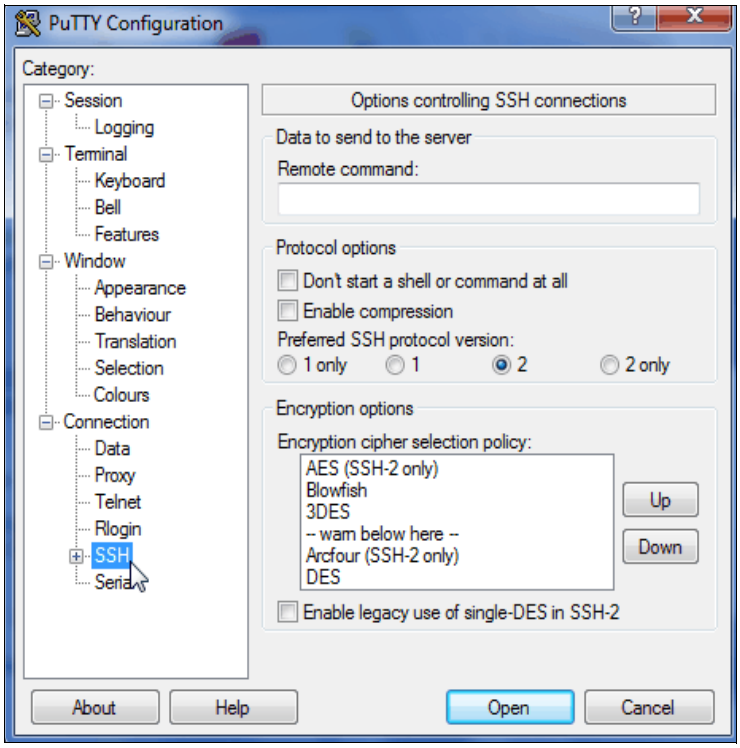

Figure A-11 SSH protocol version 2

3. In the right side pane in the “Preferred SSH protocol version” section, select **2**.

4. From the Category pane on the left side of the PuTTY Configuration window, click **Connection** → **SSH** → **Auth**. As shown in Figure A-12, in the right side pane in the "Private key file for authentication:" field under the Authentication Parameters section, browse to or manually enter the fully qualified directory path and file name of the SSH client private key file that was created earlier (for example, `C:\Support Utils\PuTTY\icat.PPK`).

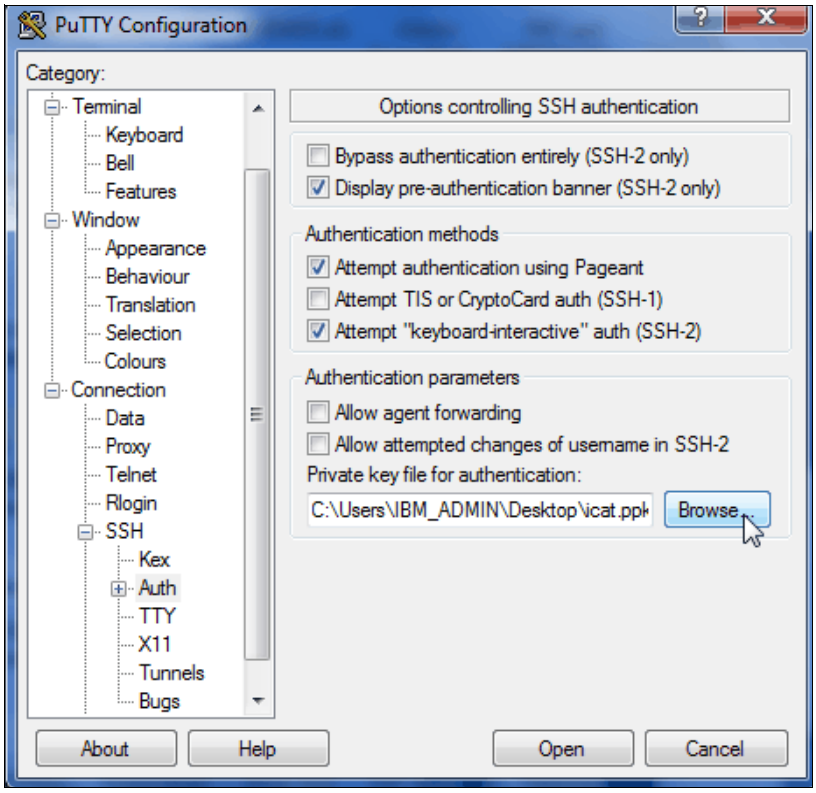

Figure A-12 SSH authentication

5. From the Category pane on the left side of the PuTTY Configuration window, click **Session** to return to the Session view, as shown in Figure A-10 on page 616.

6. In the right side pane, enter the host name or system IP address of the IBM Storwize V5000 clustered system in the Host Name field. Enter a session name in the Saved Sessions field, as shown in Figure A-13.

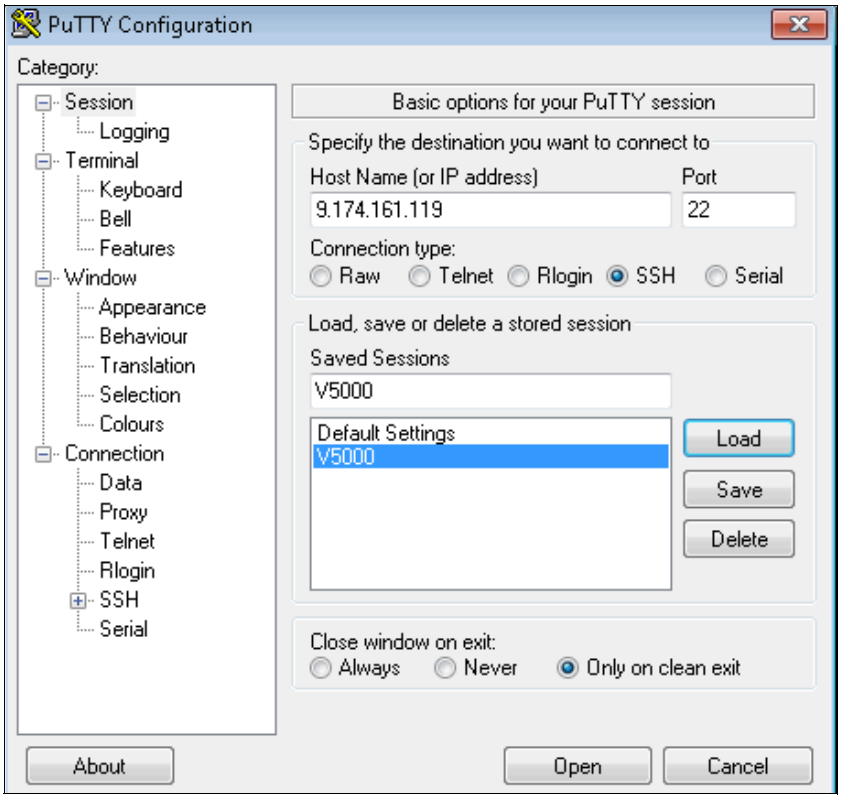

Figure A-13 Enter session information

7. Click **Save** to save the new session, as shown in Figure A-14.

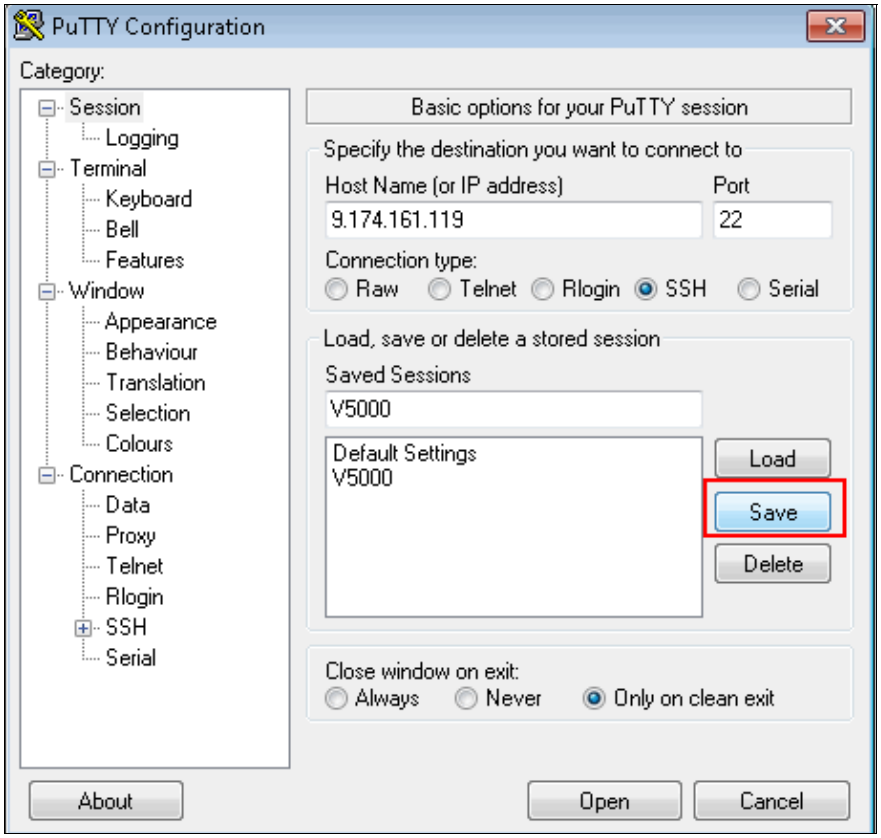

Figure A-14 Save Session

8. Highlight the new session and click **Open** to connect to the IBM Storwize V5000 system.
9. PuTTY now connects to the system and prompts you for a user name. Enter `admin` as the user name and press Enter (see Example A-1).

Example: A-1 Enter user name

```
login as: superuser
Authenticating with public key "rsa-key-20130521"
Last login: Tue May 21 15:21:55 2013 from 9.174.219.143
IBM_Storwize:mcr-atl-cluster-01:superuser>
```

The CLI is now configured for IBM Storwize V5000 administration.

Example commands

A detailed description about all of the available commands is beyond the intended scope of this book. In this section, sample commands that we referenced in this book are presented.

The **`svcinfo`** and **`svctask`** prefixes are no longer needed in IBM Storwize V5000. If you have scripts that use this prefix, they run without problems. If you enter **`svcinfo`** or **`svctask`** and press the Tab key twice, all of the available subcommands are listed. Pressing the Tab key twice also auto-completes commands if the input is valid and unique to the system.

Enter **lsvdisk** (as shown in Example A-2) to list all configured volumes on the system. The example shows that six volumes are configured.

Example: A-2 List all volumes

```
IBM_Storwize:mcr-atl-cluster-01:superuser>lsvdisk
id name        IO_group_id IO_group_name status mdisk_grp_id mdisk_grp_name
capacity type    FC_id FC_name RC_id RC_name vdisk_UID    opy_count
fast_write_state se_copy_count RC_change compressed_copy_count
0  V5000_Vol1  0           io_grp0       online 0            V5000_Pool
20.00GB  striped                          6005076300800         empty
1             no        0
1  V5000_Vol2  0           io_grp0       online 0            V5000_Pool
2.00GB   striped                          6005076300800         empty
1             no        0
2  V5000_Vol3  0           io_grp0       online 0            V5000_Pool
2.00GB   striped                          6005076300800         empty
1             no        0
3  V5000_Vol4  0           io_grp0       online 0            V5000_Pool
2.00GB   striped                          6005076300800         empty
1             no        0
4  V5000_Vol5  0           io_grp0       online 0            V5000_Pool
2.00GB   striped                          6005076300800         empty
1             no        0
5  V5000_Vol6  0           io_grp0       online 0            V5000_Pool
2.00GB   striped                          6005076300800         empty
1             no        0
```

Enter **lshost** to see a list of all configured hosts on the system, as shown in Example A-3.

Example: A-3 List hosts

```
IBM_Storwize:mcr-atl-cluster-01:superuser>lshost
id name          port_count iogrp_count status
0  windows2008r2 2          4           online
```

To map the volume to the hosts, enter **mkvdiskhostmap**, as shown in Example A-4.

Example: A-4 Map volumes to host

```
IBM_Storwize:mcr-atl-cluster-01:superuser>mkvdiskhostmap -host ESXi-1 -scsi 0   -force
ESXi-Redbooks
Virtual Disk to Host map, id [0], successfully created
```

To verify the host mapping, enter **lsvdiskhostmap**, as shown in Example A-5.

Example: A-5 List all hosts that are mapped to a volume

```
IBM_Storwize:mcr-atl-cluster-01:superuser>lshostvdiskmap ESXi-1
id name   SCSI_id vdisk_id vdisk_name    vdisk_UID
4  ESXi-1 0       2        ESXi-Redbooks 600507680185853FF000000000000011
```

In the CLI, there are more options available than in the GUI. All advanced settings can be set; for example, I/O throttling. To enable I/O throttling, change the properties of a volume by using the **changevdisk** command, as shown in Example A-6. To verify the changes, run the **lsvdisk** command.

Example: A-6 Enable advanced properties: I/O throttling

```
IBM_Storwize:mcr-atl-cluster-01:superuser>chvdisk -rate 1200 -unit mb
ESXi-Redbooks
IBM_Storwize:mcr-atl-cluster-01:superuser>
IBM_Storwize:mcr-atl-cluster-01:superuser>lsvdisk ESXi-Redbooks

id 2
name ESXi-Redbooks
.
.
vdisk_UID 600507680185853FF000000000000011
virtual_disk_throttling (MB) 1200
preferred_node_id 2
.
.
IBM_Storwize:mcr-atl-cluster-01:superuser>
```

Command output: The **lsvdisk** command lists all available properties of a volume and its copies. To make it easier to read, lines in Example A-6 were deleted.

If you do not specify the unit parameter, the throttling is based on I/Os instead of throughput, as shown in Example A-7.

Example: A-7 Throttling based on I/O

```
IBM_Storwize:mcr-atl-cluster-01:superuser>chvdisk -rate 4000 ESXi-Redbooks
IBM_Storwize:mcr-atl-cluster-01:superuser>lsvdisk ESXi-Redbooks
id 2
name ESXi-Redbooks
.
.
vdisk_UID 600507680185853FF000000000000011
throttling 4000
preferred_node_id 2
.
.
IBM_Storwize:mcr-atl-cluster-01:superuser>
```

To disable I/O throttling, set the I/O rate to 0, as shown in Example A-8.

Example: A-8 Disable I/O Throttling

```
IBM_Storwize:mcr-atl-cluster-01:superuser>chvdisk -rate 0 ESXi-Redbooks
IBM_Storwize:mcr-atl-cluster-01:superuser>lsvdisk ESXi-Redbooks
id 2
.
.
vdisk_UID 600507680185853FF000000000000011
throttling 0
preferred_node_id 2
IBM_Storwize:mcr-atl-cluster-01:superuser>
```

SAN Boot

IBM Storwize V5000 supports SAN Boot for Windows, VMware, and many other operating systems. SAN Boot support can change, so regularly check the IBM Storwize V5000 interoperability matrix at this website:

http://www-01.ibm.com/support/docview.wss?uid=ssg1S1004111

The IBM Storwize V5000 Information Center has more information about SAN Boot for different operating systems, which is available at this website:

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

For more information about SAN Boot, see the *IBM System Storage Multipath Subsystem Device Driver User's Guide,* GC52- 1309-03, which is available at this website:

ftp://ftp.software.ibm.com/storage/subsystem/UG/1.8--3.0/SDD_1.8--3.0_User_Guide_English_version.pdf

Enabling SAN Boot for Windows

Complete the following steps to install the Windows host by using SAN Boot:

1. Configure the IBM Storwize V5000 system so that only the boot volume is mapped to the host.
2. Configure the Fibre Channel SAN so that the host sees only one IBM Storwize V5000 system node port. Multiple paths during installation are not supported.
3. Configure and enable the host bus adapter (HBA) BIOS.
4. Install the operating system by using the normal procedure and select the volume as the partition on which to install.

   HBAs: You might need to load another HBA device driver during installation, depending on your Windows version and the HBA type.

5. Install SDDDSM after the installation completes.
6. Modify your SAN zoning to allow multiple paths.

7. Check your host to see whether all paths are available.
8. Set redundant boot devices in the HBA BIOS to allow the host to boot when its original path fails.

Enabling SAN Boot for VMware

Complete the following steps to install a VMware ESXhost by using SAN Boot:

1. Configure the IBM Storwize V5000 system so that only the boot volume is mapped to the host.
2. Configure the Fibre Channel SAN so that the host sees only one IBM Storwize V5000 system node port. Multiple paths during installation are not supported.
3. Configure and enable the HBA BIOS.
4. Install the operating system by using the normal procedure and select the volume as the partition on which to install.

   HBAs: You might need to load another I HBA device driver during installation, depending on your ESX level and the HBA type.

5. Modify your SAN zoning to allow multiple paths.
6. Check your host if all paths are available and modify the multipath policy, if required.

Windows SAN Boot migration

If you have a host that runs a Windows 2000 Server, Windows Server 2003, or Windows Server 2008 operating system and have existing SAN Boot images that are controlled by storage controllers, you can migrate these images to image-mode volumes that are controlled by the IBM Storwize V5000 system.

SAN Boot procedures: For more information about SAN Boot procedures for other operating systems, see the IBM Storwize V5000 Information Center at this website:

http://pic.dhe.ibm.com/infocenter/storwize/V5000_ic/index.jsp?topic=%2Fcom.ibm.storwize.V5000.641.doc%2FV5000_ichome_641.html

Complete the following steps to migrate your existing SAN Boot images:

1. If the existing SAN Boot images are controlled by an IBM storage controller that uses the IBM Subsystem Device Driver (SDD) as the multipathing driver, you must use SDD V1.6 or higher. Run the SDD **`datapath set bootdiskmigrate 2076`** command to prepare the host for image migration. See the Multipath SDD matrix to download packages at this website:

   http://www-01.ibm.com/support/docview.wss?rs=540&context=ST52G7&dc=DA400&uid=ssg1S7001350&loc=en_US&cs=utf-8&lang=en#WindowsSDD

2. Shut down the host.
3. Complete the following configuration changes on the storage controller:
   a. Write down the SCSI LUN ID each volume is using (for example, boot LUN SCSI ID 0, Swap LUN SCSI ID 1, and Database LUN SCSID 2).
   b. Remove all of the image-to-host mappings from the storage controller.

c. Map the existing SAN Boot image and any other disks to the IBM Storwize V5000 system.

4. Change the zoning so that the host can see the IBM Storwize V5000 I/O group for the target image mode volume.
5. Complete the following configuration changes on the IBM Storwize V5000 system:
   a. Create an image mode volume for the managed disk (MDisk) that contains the SAN Boot image. Use the MDisk unique identifier to specify the correct MDisk.
   b. Create a host object and assign the host HBA ports.
   c. Map the image mode volume to the host by using the same SCSI ID as before. For example, you might map the boot disk to the host with SCSI LUN ID 0.
   d. Map the swap disk to the host, if required. For example, you might map the swap disk to the host with SCSI LUN ID 1.
6. Change the boot address of the host by completing the following steps:
   a. Restart the host and open the HBA BIOS utility of the host during the booting process.
   b. Set the BIOS settings on the host to find the boot image at the worldwide port name (WWPN) of the node that is zoned to the HBA port.
7. If SDD V1.6 or higher is installed and you ran the **bootdiskmigrate** command in step 1 on page 624, reboot your host, update SDDDSM to the latest level, and go to step 14. If SDD V1.6 is not installed, go to step 8.
8. Modify the SAN Zoning so that the host sees only one path to the IBM Storwize V5000.
9. Boot the host in single-path mode.
10. Uninstall any multipathing driver that is not supported for IBM Storwize V5000 system hosts that run the applicable Windows Server operating system.
11. Install SDDDSM.
12. Restart the host in single-path mode and ensure that SDDDSM was properly installed.
13. Modify the SAN Zoning to enable multipathing.
14. Rescan drives on your host and check that all paths are available.
15. Reboot your host and enter the HBA BIOS.
16. Configure the HBA settings on the host. Ensure that all HBA ports are boot-enabled and can see both nodes in the I/O group that contains the SAN Boot image. Configure the HBA ports for redundant paths.
17. Exit the BIOS utility and finish starting the host.
18. Map any other volumes to the host, as required.

Related publications and information

The publications and information that is listed in this section are considered particularly suitable for a more detailed discussion of the topics that are covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide more information about the topics in this book. Some publications that are referenced in the following list might be available in softcopy only:

- *Implementing the IBM System Storage SAN Volume Controller V6.3*, SG24-7933
- *Implementing the IBM Storwize V7000 V6.3*, SG24-7938
- *SAN Volume Controller: Best Practices and Performance Guidelines*, SG24-7521
- *Implementing an IBM/Brocade SAN with 8 Gbps Directors and Switches*, SG24-6116

You can search for, view, download, or order these documents and other Redbooks, Redpapers, Web Docs, draft, and other materials at this website:

`http://www.ibm.com/redbooks`

IBM Storwize V5000 publications

Storwize V5000 publications are available at this website:

`https://ibm.biz/BdxyDL`

IBM Storwize V5000 support

Storwize V5000 support is available at this website:

`https://ibm.biz/BdxyD9`

Help from IBM

IBM Support and downloads

`https://www.ibm.com/support`

IBM Global Services

`https://www.ibm.com/services`

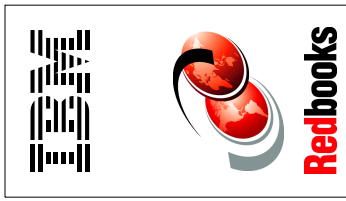

Implementing the IBM Storwize V5000

(1.0” spine)
0.875”<->1.498”
460 <-> 788 pages

Implementing the IBM Storwize V5000

Easily manage and deploy systems with embedded GUI

Experience rapid and flexible provisioning

Protect data with remote mirroring

Organizations of all sizes are faced with the challenge of managing massive volumes of increasingly valuable data. But storing this data can be costly, and extracting value from the data is becoming more difficult. IT organizations have limited resources but must stay responsive to dynamic environments and act quickly to consolidate, simplify, and optimize their IT infrastructures. The IBM Storwize V5000 system provides a smarter solution that is affordable, easy to use, and self-optimizing, which enables organizations to overcome these storage challenges.

Storwize V5000 delivers efficient, entry-level configurations that are specifically designed to meet the needs of small and midsize businesses. Designed to provide organizations with the ability to consolidate and share data at an affordable price, Storwize V5000 offers advanced software capabilities that are usually found in more expensive systems.

This IBM Redbooks publication is intended for pre-sales and post-sales technical support professionals and storage administrators.

The concepts in this book also relate to the IBM Storwize V3700.

This book was written at a software level of Version 7 Release 1.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-8162-00 ISBN 0738438766