IBM

Implementing the Storwize V3500

- Consolidate and share data with low-cost iSCSI
- Deploy storage in minutes and perform storage management
- Experience peace of mind

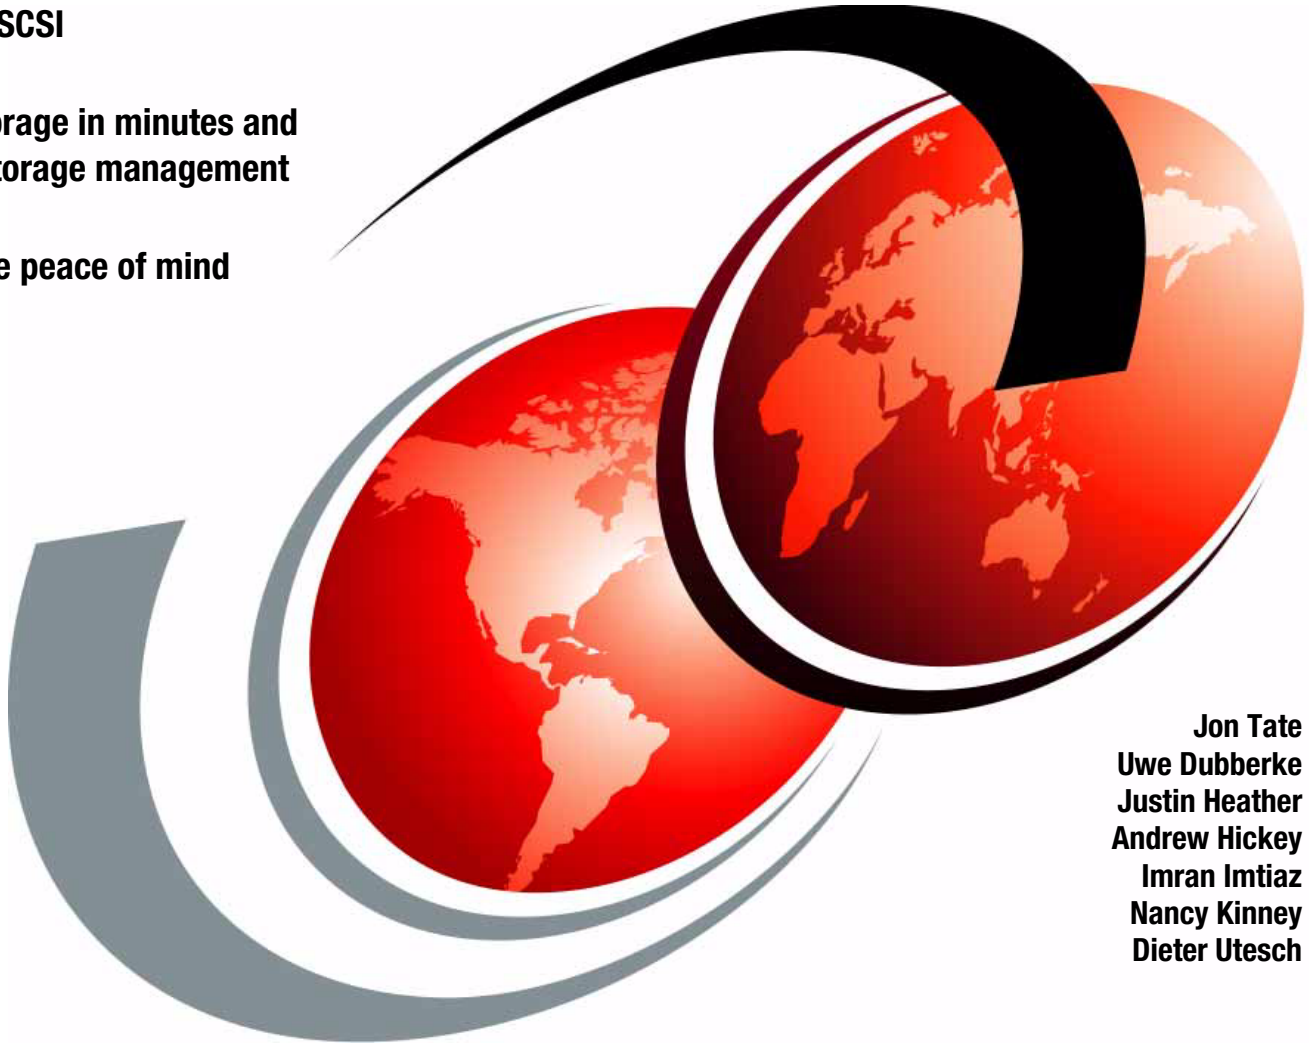

Jon Tate
Uwe Dubberke
Justin Heather
Andrew Hickey
Imran Imtiaz
Nancy Kinney
Dieter Utesch

Redbooks

ibm.com/redbooks

International Technical Support Organization

Implementing the Storwize V3500

May 2013

SG24-8125-00

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

First Edition (May 2013)

This edition applies to Version 6 Release 4 of IBM Storwize V3500 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®
FlashCopy®
Global Technology Services®
IBM Flex System™
IBM®
Redbooks®
Redbooks (logo) ®
Storwize®
System Storage®
Tivoli®

The following terms are trademarks of other companies:

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

Businesses of all sizes are faced with the challenge of managing huge volumes of data that are becoming increasingly valuable. But storing this data can be costly, and extracting value from the data is becoming more and more difficult. IT organizations have limited resources and cannot afford to make investment mistakes. The IBM® Storwize® V3500 system provides a smarter solution that is affordable, simple, and efficient, which enables businesses to overcome their storage challenges.

IBM Storwize V3500 is the most recent addition to the IBM Storwize family of disk systems. It delivers easy-to-use, entry-level configurations that are specifically designed to meet the modest budgets of small and medium-sized businesses.

IBM Storwize V3500 features the following highlights:

- Consolidate and share data with low cost iSCSI storage networking.
- Deploy storage in minutes and perform storage management tasks quickly and easily through a breakthrough graphical user interface.
- Experience peace of mind with proven IBM Storwize family high-availability data protection with snapshot technology and IBM warranty support.
- Optimize efficiency by allocating only the amount of disk space needed at the time it is required with high performance, thin-provisioning capabilities.

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 support for IBM storage products. Jon has 27 years of experience in storage software and management, services, and support. He is an IBM Certified IT Specialist and an IBM SAN Certified Specialist. He also is the UK Chairman of the Storage Networking Industry Association.

Uwe Dubberke is an IBM Certified Specialist for High End Disk Solutions, working as a field specialist for DASD and SAN products in IBM Germany. Starting in 1990, at IBM he was responsible for different high-end customers as an Account CE. He also worked as an SE and since 1999 he has supported the EMEA Central Region Hardware Support Center in Mainz as a virtual team member. Since 2005, he is a virtual member for the SAN and DASD Support Group in Mainz. He holds a degree in Electrical Engineering, with a specialization in communications engineering, from the University of Applied Sciences of Gelsenkirchen (Germany). Uwe has co-authored IBM Redbooks® on DS8000®, SSD and Brocade.

Justin Heather is an IBM Remote Technical Support Engineer for IBM System Storage and SAN Solutions. Based in Brisbane, he is responsible for supporting IBM Storage clients from Australia and New Zealand. Joining IBM in 1995, Justin has spent many years as an IBM SSR, providing on site support for IBM ZIPX Servers, IBM System Storage, and OEM networking products. He also has performed an IBM Advisory SE role by providing pre-sales technical support and on site IBM Lab Services for IBM System Storage and SAN solutions. Justin is studying for a Bachelor of Information Technology degree at CQUniversity, Australia.

Andrew Hickey was a software engineer working in Storage Systems at the IBM Manchester Lab in Manchester, England, at the time of this writing. He worked in the Manchester Lab since it was founded in 2008. He also worked on the IBM Storwize V3500 and V3700 products as a tester, and contributed to their online documentation. He has since left IBM.

Imran Imtiaz is a Software Test Engineer for IBM System Storage SAN Solutions at the IBM Manchester Lab, UK. He joined IBM in 2008 and is responsible for the functional testing of various IBM products. His areas of expertise include storage hardware, storage area networks, virtualization, and system security. He holds a degree in Computer Science with a specialization in Computer Networks, from Manchester Metropolitan University, UK.

Nancy Kinney worked as an IBM Remote Technical Support Engineer for IBM Global Technology Services® in the Austin AIX/IBM System Storage center where she acquired end-to-end experience in troubleshooting FC technologies. She is an Subject Matter Expert in the Columbia Missouri Storage GDC, responsible for supporting SSO Storage clients. She holds an IBM Midrange Specialist Certification and NETAPP/NSeries NCDA Certification. She has a range of experience working with FC Storage technologies from multiple vendors and working with multipathing drivers for operating systems that attach to storage.

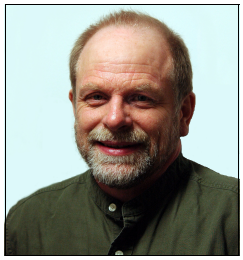

Dieter Utesch is a Product Field Engineer for IBM Midrange Storage Products in IBM Global Technology Services, Germany. Since 2008, he worked as Level 2 support for IBM DS4000/DS5000 products in the European Storage Competence Centre (ESCC). Before that, Dieter joined the Information Infrastructure and Lab Solutions group in Systems Lab Europe, Mainz, Germany. Dieter has 30 years experience of IBM Information Technology.

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

- Martyn Spink
- Djihed Afifi
- David Turnbull

  IBM Manchester Lab

- Paul Marris
- Simon Swindells

  IBM Hursley

- John Bynum

  IBM San Jose

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 using leading-edge technologies. Your efforts will help 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 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

IBM Storwize V3500 overview

This chapter provides an overview of the Storwize V3500.

This chapter includes the following topics:

- Storwize V3500 overview
- Storwize V3500 terminology
- Storwize V3500 hardware
- Storwize V3500 terms
- FlashCopy
- Problem management and support
- Useful Storwize V3500 websites
- Storwize V3500 learning videos on YouTube

1.1 Storwize V3500 overview

The Storwize V3500 solution provides a modular storage system that includes the capability to virtualize its own internal storage. The IBM Storwize V3500 system is a virtualizing RAID entry/midrange storage system.

Storwize V3500 has the following benefits:

- Brings enterprise technology to entry storage
- Specialty administrators are not required
- Easy client setup and service
- Simple integration into the server environment

The Storwize V3500 subsystem consists of one enclosure. Control enclosures contain disk drives and two nodes (an I/O group), which are attached to the SAN fabric.

The simplest use of the Storwize V3500 is as a traditional RAID subsystem. The internal drives are configured into RAID arrays and virtual disks are created from those arrays.

Storwize V3500 supports regular drives only.

A Storwize V3500 system provides a number of configuration options that are aimed at simplifying the implementation process. It also provides automated wizards, called Directed Maintenance Procedures (DMPs), to help resolve any events that might occur. An Storwize V3500 system is a clustered, scalable, and an entry storage system.

Figure 1-1 shows the Storwize V3500 high-level overview that is attached via a SAN.

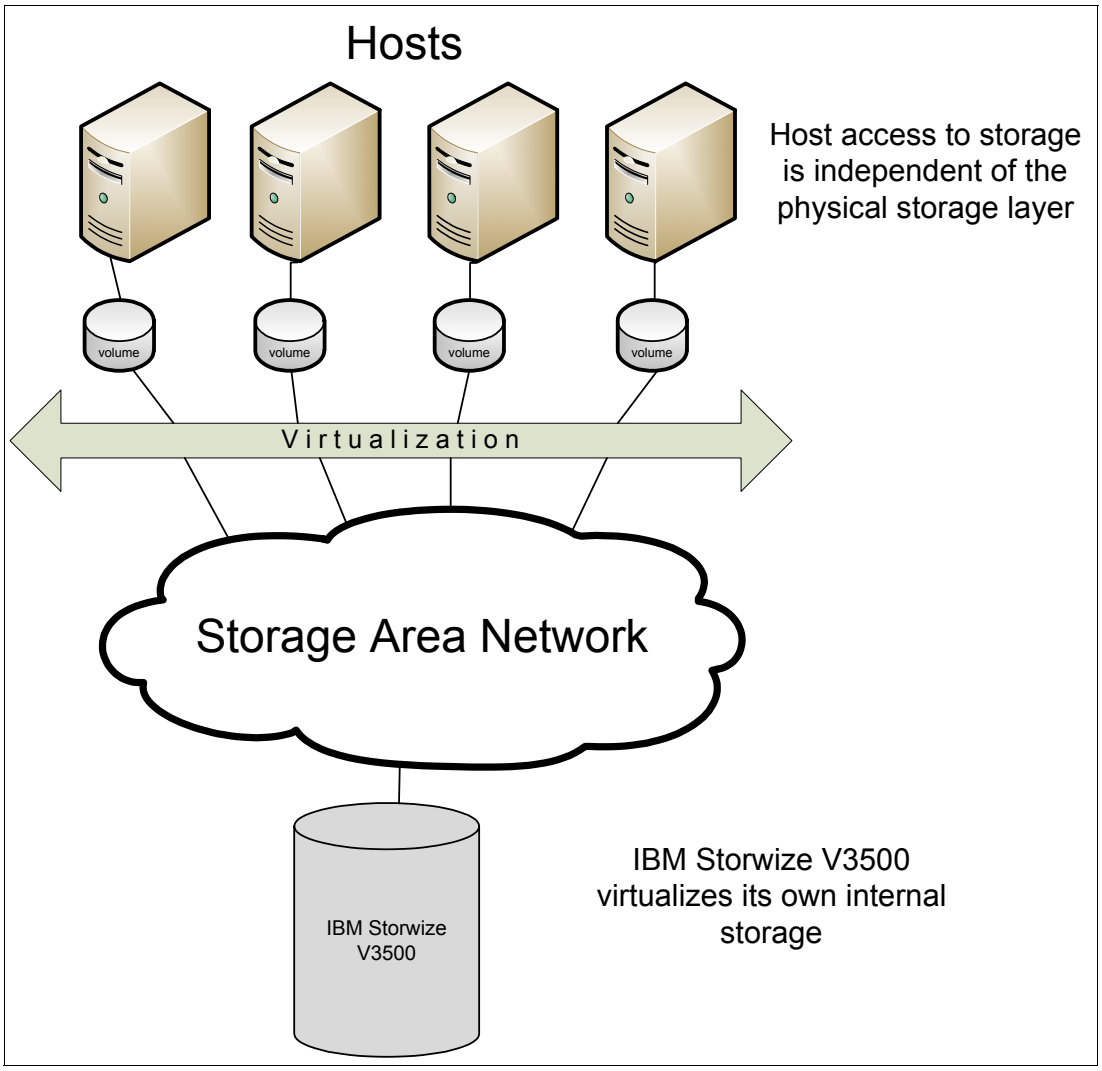

Figure 1-1 Storwize V3500 overview that is attached via SAN

Important: For more information about restrictions, see these websites:

- http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss
- http://www-01.ibm.com/support/docview.wss?uid=ssg1S1004197

A major benefit of the Storwize V3500 is that you can directly attach it to a host, as shown in Figure 1-2.

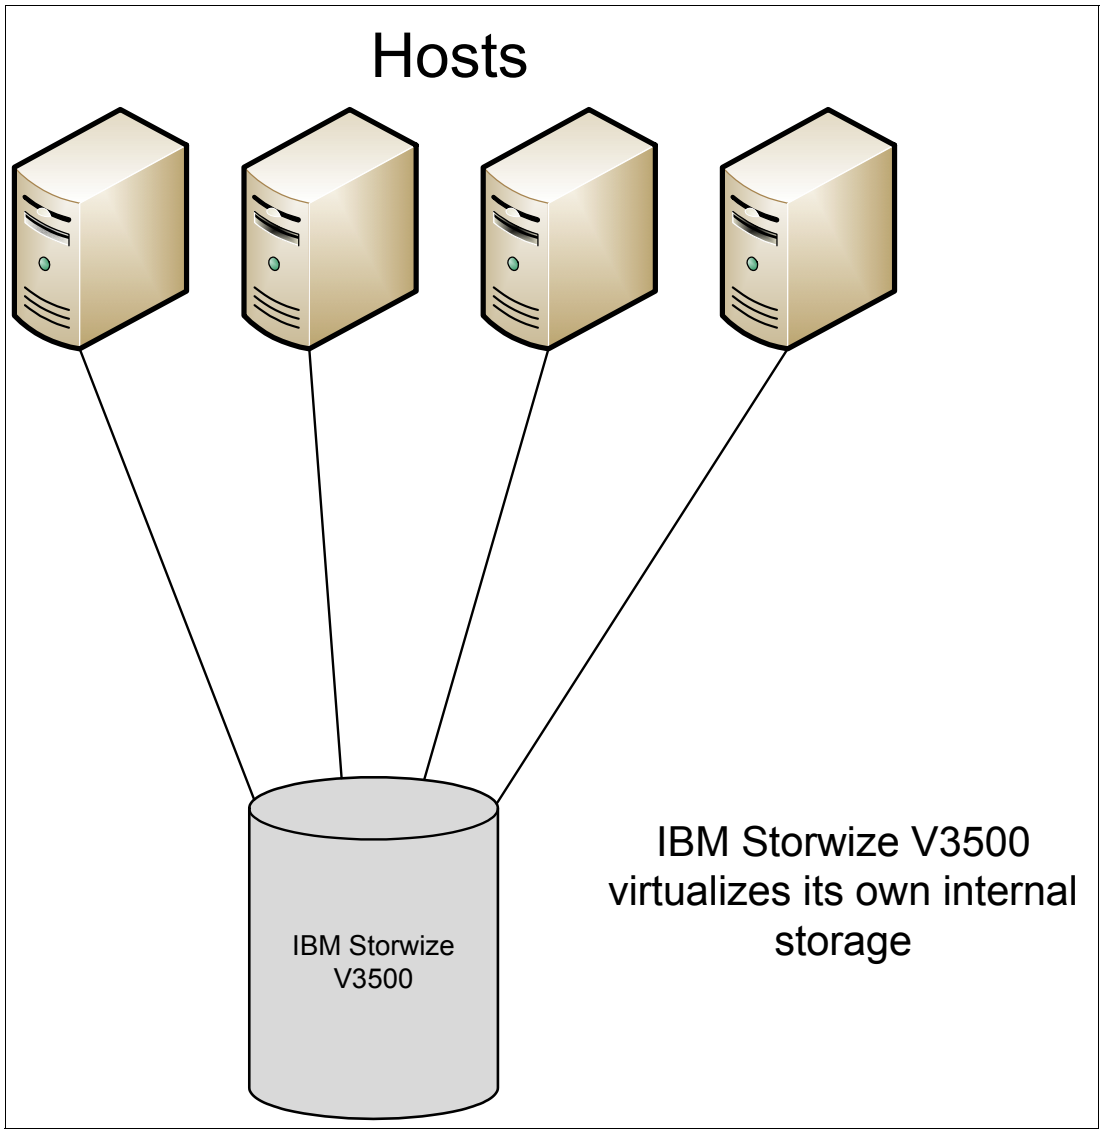

Figure 1-2 Storwize V3500 direct attached

Included with the Storwize V3500 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 pre-established 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.

The Storwize V3500 solution provides a choice of up to 12 x 3.5-inch or 24 x 2.5-inch Serial Attached SCSI (SAS) drives for the internal storage.

The Storwize V3500 solution consists of one control enclosure. The control enclosure contains two node canisters, two power supplies, and the drives.

You can have up to 36 TB of storage capacity one complete enclosure.

1.2 Storwize V3500 terminology

The Storwize V3500 system introduced some terminology that is consistent with the rest of the IBM Storwize family and SAN Volume Controller. The terms are defined in Table 1-1.

Table 1-1 Storwize V3500 terminology

Storwize V3500 term	Definition
Battery	Each control enclosure node canister in a Storwize V3500 contains a battery.
Canister	Canisters are hardware units that are subcomponents of a Storwize V3500 enclosures. Each enclosure contains two canisters.
Clone	A copy of a volume on a server at a particular point in time. 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	You can migrate data from existing external storage to its internal volumes.
Drive	Storwize V3500 supports a range of hard disk drives (HDDs) and solid-state drives (SSDs)
Enclosure	An enclosure is the basic housing unit for the Storwize V3500. 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.
Fibre Channel port	Fibre Channel ports are connections for the hosts to get access to the Storwize V3500.
Host mapping	The process of controlling which hosts have access to specific volumes within a Storwize V3500.
Internal storage	Array managed disks (MDisks) and drives that are held in enclosures and nodes that are part of the Storwize V3500.
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 Small Computer System Interface (SCSI) logical unit (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, serial-attached SCSI (SAS) ports, and battery.
Power Supply Unit	Each enclosure has two power supply units (PSU).

Storwize V3500 term	Definition
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.
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.
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 I/O control.
Worldwide port names	Each Fibre Channel port is identified by their physical port number and by a worldwide port name (WWPN).

1.2.1 Storwize V3500 models

The Storwize V3500 platform consists of two different models. Each model contains two canisters that, although they can be replaced independently, are seen as part of the enclosure.

More information: For the most up-to-date information about the features, benefits, and specifications of Storwize V3500 models, see this website:

http://pic.dhe.ibm.com/infocenter/storwize/v3500_ic/index.jsp

The information in this book is valid at the time of writing. However, as the Storwize V3500 matures, there might be new features and enhanced specifications.

The Storwize V3500 models are described in Table 1-2.

Table 1-2 Storwize V3500 models

Model	Cache	Drive slots	Maximum capacity	Power supply
2071-02A (with two node canisters)	8 GB	12 x 3.5-inch	36 TB	Integrated dual power supplies
2071-10A (with two node canisters)	8 GB	24 x 2.5-inch	24 TB	Integrated dual power supplies

iSCSI and Fibre Channel connectivity with intermix flexibility

With iSCSI host interface support and optional Fibre Channel (FC) support, Storwize V3500 is designed to accommodate the following common storage network technologies:

- By using iSCSI, you can capitalize on your existing IT skills and infrastructure by delivering block-based storage across the IP network. iSCSI provides an alternative to the investment in a storage-only, FC infrastructure.
- A mature and proven technology, FC is the host interface of choice for high-performance environments. By using Storwize V3500 with FC, you can use the benefits that are offered by a SAN environment.

The Storwize V3500 comes with four 1-Gb iSCSI host interface ports, with eight 8-Gb FC ports available as an optional feature.

Figure 1-3 shows the front view of the 2071-02A enclosure.

Figure 1-3 IBM Storwize V3700 front view for 2072-12C, 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 1-4 shows the front view of the 2071-10A enclosure.

Figure 1-4 Storwize V3500 front view for 2072-24C, 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).

1.2.2 Storwize V3500 features

The following features are available with Storwize V3500:

- Thin provisioning (space-efficient disks)

  Traditional fully allocated volumes allocate real physical disk capacity for an entire volume, even if that capacity is never used. Thin-provisioned volumes allocate real physical disk capacity only when data is written to the logical volume.

- Volume mirroring

  Provides a single volume image to the attached host systems while pointers are maintained to two copies of data in separate storage pools. If one copy fails, Storwize V3500 provides continuous data access by redirecting I/O to the remaining copy. When the copy becomes available, automatic resynchronization occurs.

- FlashCopy

  Storwize V3500 provides a volume level point-in-time copy function to its internal virtualized volumes. This function is designed to create copies for backup, parallel processing, testing, and development, and have the copies available almost immediately.

  Storwize V3500 includes the following FlashCopy functions:

  – Full and incremental copy

    This function copies only the changes from the source or target data since the last FlashCopy operation. It is designed to enable completion of point-in-time online backups much more quickly than the use of traditional FlashCopy.

  – Multitarget FlashCopy

    Storwize V3500 supports multiple copies from a single source. Each copy is managed by a unique mapping. In general, each mapping acts independently and is not affected by other mappings that are sharing the source volume.

  – Cascaded FlashCopy

    This function is used to create copies of copies and supports full, incremental, or nocopy operations.

  – Reverse FlashCopy

    The use of this function allows data from an earlier point-in-time copy to be restored with minimal disruption to the host.

  – FlashCopy nocopy with thin provisioning

    This function provides a combination of the use of thin-provisioned volumes and FlashCopy together to help reduce disk space requirements when copies are made. The following variations of this option are available:

    - Space-efficient source and target with background copy, which copies only the allocated space.
    - Space-efficient target with no background copy, which copies only the space that is used for changes between the source and target. This variation generally is referred to as *snapshots*.

    This function can be used with multi-target, cascaded, and incremental FlashCopy.

- Consistency groups

  Consistency groups address the issue where application data is on multiple volumes. By placing the FlashCopy relationships into a consistency group, commands can be issued against all of the volumes in the group. This action enables a consistent point-in-time copy of all of the data, even though it might be on a physically separate volume.

FlashCopy mappings can be members of a consistency group, or they can be operated in a stand-alone manner; that is, not as part of a consistency group. FlashCopy commands can be issued to a FlashCopy consistency group, which affects all FlashCopy mappings in the consistency group, or to a single FlashCopy mapping if it is not part of a defined FlashCopy consistency group.

- Data Migration (no charge for temporary usage)

  Storwize V3500 provides a data migration function that can be used to import external storage systems into the Storwize V3500 system.

  By using this function, you can complete the following tasks:

  - Move volumes nondisruptively onto a newly installed storage system
  - Migrate data from other storage systems to Storwize V3500 managed storage

- The Storwize V3500 can be virtualized from systems such as the IBM SAN Volume Controller, IBM Storwize V7000, or IBM Flex System™ V7000 Storage Node.

1.2.3 Storwize V3500 availability

Availability of this product is limited to the following countries and regions:

- People's Republic of China
- Hong Kong S.A.R. of the PRC
- Macao S.A.R. of the PRC
- Taiwan

1.3 Storwize V3500 hardware

The Storwize V3500 solution is a modular storage system that is built on a common enclosure (control enclosure and expansion enclosure).

Figure 1-5 shows an overview of the hardware components of the Storwize V3500 solution.

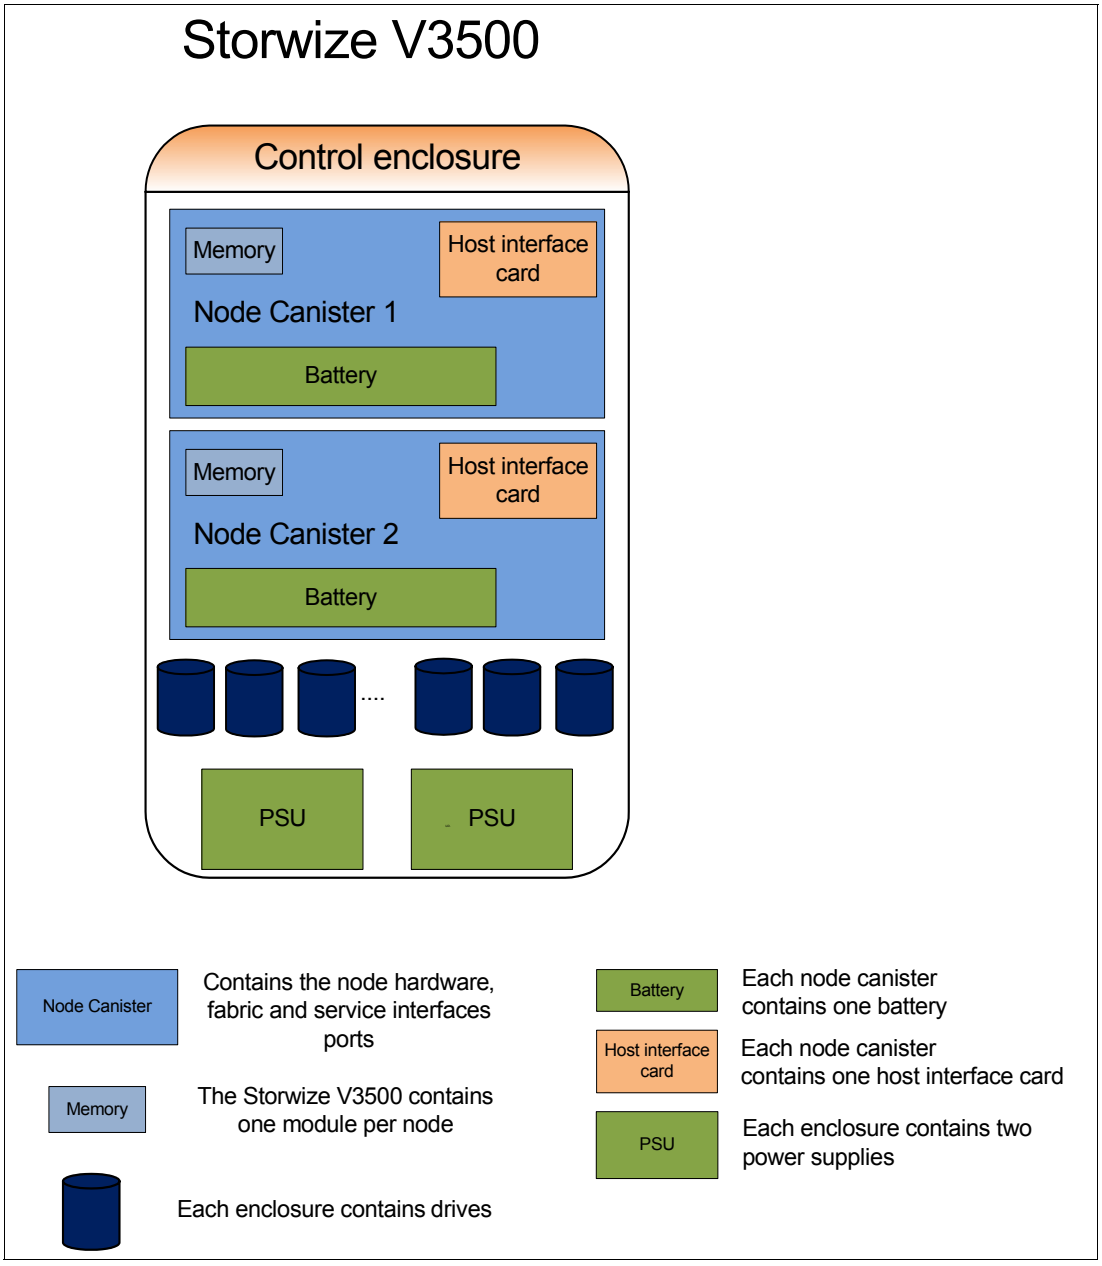

Figure 1-5 Storwize V3500 hardware components

Figure 1-6 shows the controller rear view of Storwize V3500 models 02A and 10A.

Figure 1-6 Storwize V3500 controller rear view - models 02A and 10A

In the rear view, you can see that there are two power supply slots at the bottom of the enclosure. Both slots are the same and they are exchangeable.

There are two canister slots at the top of the chassis.

There are two models of the control enclosure: one features 12 3.5-inch disk slots (model 02A) and the other features 24 2.5-inch disk slots (model 10A).

Each node canister contains the following hardware:

- Battery
- Memory: 4-GB memory
- Host interface card
- Three 6 Gbps SAS ports (for future use)
- 10/100/1000 Mbps Ethernet ports
- Two USB 2.0 ports (one port is used during installation)
- System flash

The battery is used in a power loss. The Storwize V3500 system needs these batteries to write the cache data 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 to destage it to the disks.

In Figure 1-7, you can see the different ports at the rear, which are provided by the Storwize V3500 node canister.

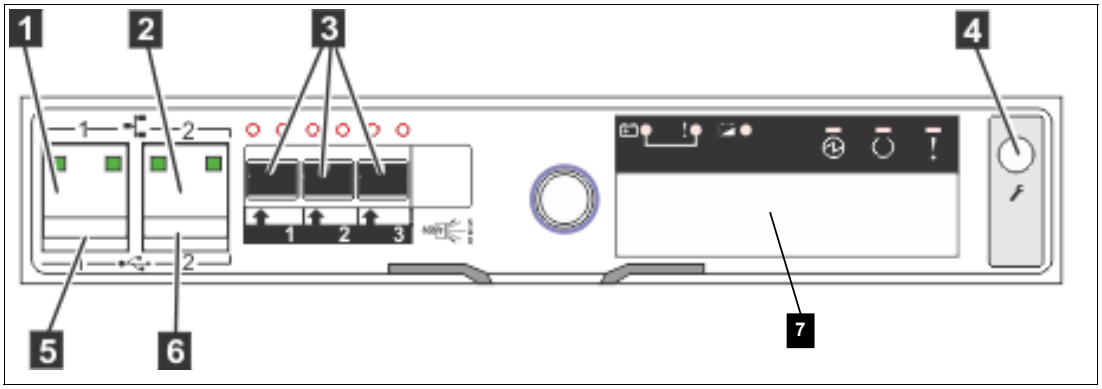

Figure 1-7 Storwize V3500 node ports

In Figure 1-7, the host interface card (HIC) is at position **7**.

The Fibre Channel ports operate at 2, 4, or 8 Gbps. The ports of the host interface cards are always numbered from left to right, starting with 1. For more information, see this website:

http://pic.dhe.ibm.com/infocenter/storwize/v3500_ic/index.jsp

The Storwize V3500 is delivered with three serial attached SCSI (SAS) ports. They are numbered 1 on the left to 3 on the right; however, the Storwize V3500 does not use these ports. They are at position number 3 in Figure 1-7 on page 11.

The two 10/100/1000 Mbps Ethernet ports that are shown in Figure 1-7 on page 11 and numbered 1 and 2 are used for management. Port 1 (left port) must be configured. The second port is optional and can be used for management. You also can use both ports for iSCSI traffic. For more information, see Chapter 4, “Host configuration” on page 141.

There are two USB ports (numbered 5 and 6 in Figure 1-7 on page 11). One is used during the initial configuration or in case of problem. They are numbered 1 on the left and 2 on the right. For more information, see Chapter 2, “Initial configuration” on page 27.

As described previously, the system flash is used to store the dump files in case of power loss and for the operating system.

Important: Do not use the port marked with 4 in Figure 1-7 on page 11. 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. 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 enclosures holds a copy of the control enclosure state information.

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

Storwize V3500 enclosures support SAS and Nearline SAS drive types. Each SAS drive has internal two ports and I/O can be issued down both paths simultaneously.

Table 1-3 shows the Storwize V3500 Disk Drive types that are available at the time of writing.

Table 1-3 Storwize V3500 Disk Drive types

Drive type		Speed	Size
2.5-inch form factor	SAS	10,000 rpm	300,600, and 900 GB
2.5-inch form factor	SAS	15,000 rpm	146 and 300 GB
2.5-inch form factor	Nearline SAS	7,200 rpm	500 GB and 1 TB
3.5-inch form factor	Nearline SAS	7,200 rpm	2 and 3 TB

1.4 Storwize V3500 terms

In this section, we introduce the terminology of the Storwize V3700 that is used throughout this book.

1.4.1 Hosts

A host system is a server that is connected to Storwize V3500 through a Fibre Channel connection or through an iSCSI connection.

Hosts are defined to the Storwize V3500 by identifying their worldwide port names (WWPNs) for Fibre Channel hosts. For iSCSI hosts, they 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 141.

A significant benefit of not having direct attach is that you can attach the host directly to the Storwize V3500. If you do not have any FC switches, you can connect hosts without any FC network directly, as shown in Figure 1-2 on page 4.

1.4.2 Node canister

Storwize V3500 has a hardware component that is called the *node canister*. It is the node that is a logical object within the node canister and provides the virtualization of internal volumes, cache, and a copy service (FlashCopy) function. A clustered system consists of a one node pair.

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

1.4.3 I/O groups

Within Storwize V3500, there are two node canisters that are known as *I/O groups*. The Storwize V3500 supports two node canisters in the clustered system, which provides one 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. However, both nodes also act as a failover node for its partner node within the I/O group. A node takes over the I/O workload from its partner node (if required) without affecting the server’s application.

In a Storwize V3500 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 Storwize V3500 I/O groups are connected to the SAN so that all application servers that are accessing volumes from the I/O group have access to them. Up to 256 host server objects can be defined to one Storwize V3500.

Important: The active-active architecture provides the availability to process I/Os for both controller nodes. It 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.

1.4.4 Clustered system

All configuration, monitoring, and service tasks are performed at the system level. The configuration settings are replicated across both 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 that is provided 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 Storwize V3500 system configuration information.

System configuration backup: After the system configuration is backed up, save the backup data on your hard disk (or at least outside of the SAN). If you are unable to access the Storwize V3500, you do not have access to the backup data if it is on the SAN. Perform this configuration after each configuration change to be on the safe side.

One node is designated as the configuration node canister and it is the only node that activates the system IP address. If the configuration node canister fails, the system chooses the other node to be the new configuration node. The new configuration node takes over the system IP addresses.

The system can be configured by using the Storwize V3500 management software (GUI) or the CLI.

1.4.5 RAID

The Storwize V3500 setup includes several internal drives, but these drives cannot be directly added to storage pools. The drives must be included in a Redundant Array of Independent Disks (RAID) 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. They also have different restrictions regarding the number of members in the array.

The Storwize V3500 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 wanted 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 1 (mirroring between two drives)
- RAID 5 (striping, can survive one drive fault, with parity)
- RAID 6 (striping, can survive two drive faults, with parity)
- RAID 10 (RAID 0 on top of RAID 1)

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 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 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 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 gives the array double redundancy. A RAID 6 array can be formed by 5 - 16 drives.

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.

1.4.6 Managed disks

A managed disk (MDisk) refers to the unit of storage that Storwize V3500 virtualizes. This unit is a RAID array that consists of internal drives. The Storwize V3500 can then allocate these MDisks into various storage pools. The Storwize V3500 also can pick up MDisk from external storage, but only during the migration process.

An MDisk is not visible to a host system on the storage area network because it is internal in the Storwize V3500 system.

An MDisk features the following modes:

- Array

  Array mode MDisks are constructed from drives that have RAID capability. Array MDisks always are associated with storage pools.

- Image

  Image MDisks are 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 217.

1.4.7 Quorum disks

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 a local disk. It allocates space from up to three different disks for redundancy, although only one quorum disk is active. It is possible to manage the quorum disks by using the CLI.

1.4.8 Storage pools

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

MDisks can be added to a storage pool at any time to increase the capacity of the storage pool. MDisks can belong in only one storage pool. For more information, see Chapter 7, “Storage pools” on page 277.

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

The extent size has a direct effect on the maximum volume size and storage capacity of the clustered system. A system can manage 4 million (4 x 1024 x 1024) extents. For example, a system with a 16-MB extent size can manage up to 16 MB x 4 MB = 64 TB of storage.

The effect of extent size on the maximum volume size is shown inTable 1-4, which lists the extent size and the corresponding maximum clustered system size.

Table 1-4 Maximum volume capacity by extent size

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

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 different, you must use volume mirroring to copy volumes between storage pools, as described in Chapter 7, “Storage pools” on page 277.

For most clustered systems, a capacity of 1 PB is sufficient. Use a value of 256 MB.

Default extent size: The GUI of Storwize V35000 has a default extent size value of 256 MB when you define a new storage pool.

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.

1.4.9 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 Storwize V3500. Storwize V3500 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 the extents around without affecting 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 three 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 1-8 shows how a striped volume is allocated, assuming 10 extents are required.

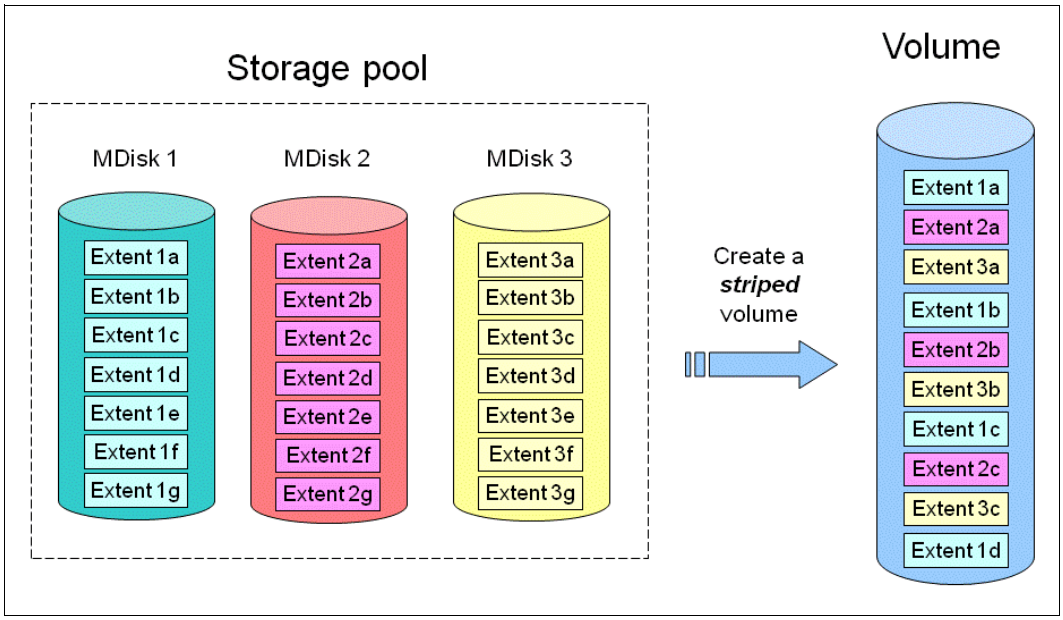

Figure 1-8 Striped volume

- Sequential

  A sequential volume is where the extents are allocated one after the other from one MDisk to the next MDisk, as shown in Figure 1-9.

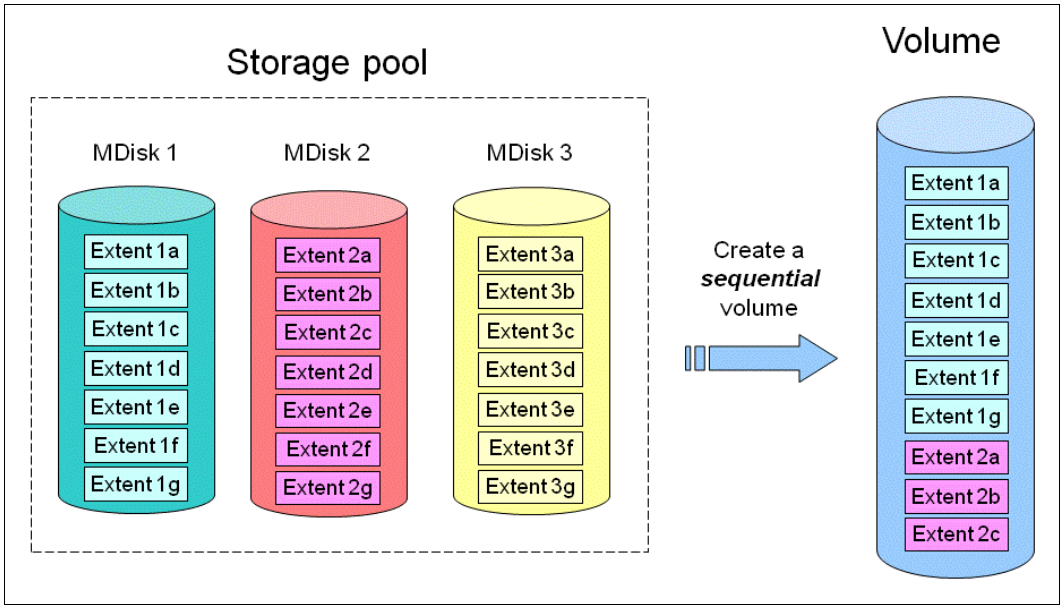

Figure 1-9 Sequential volume

- 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.

  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. This configuration ensures that the data on the MDisk is preserved as it is brought into the clustered system, as shown in Figure 1-10 on page 19.

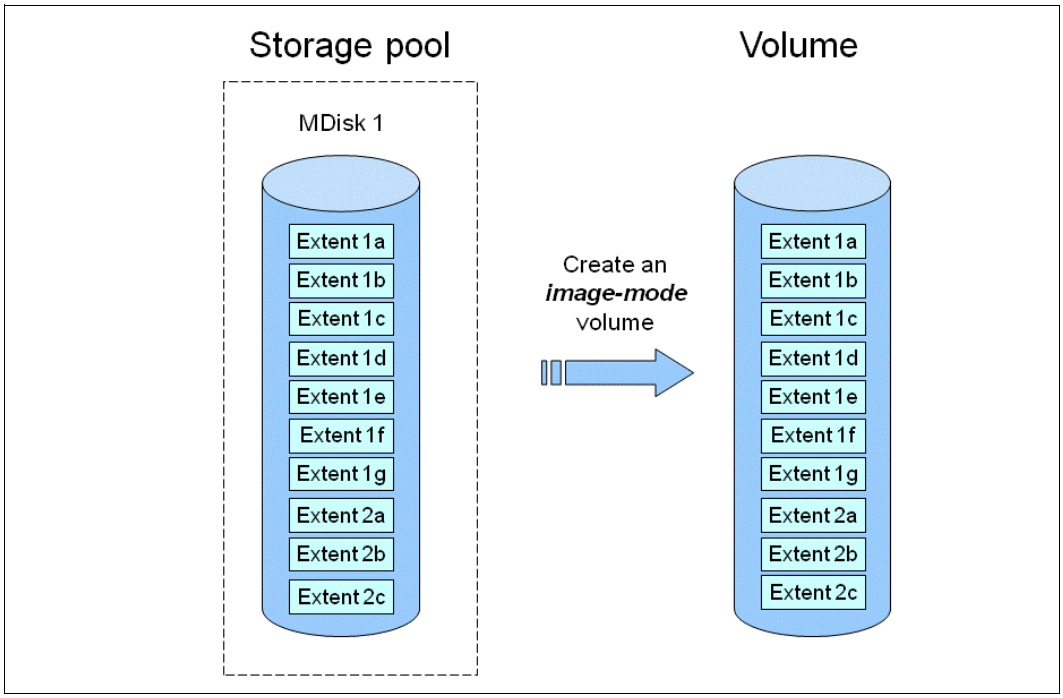

Figure 1-10 Image mode volume

On the Storwize V3500, the image mode volume is used only to migrate the volume into the Storwize V3500. After it is migrated, you can disconnect the migrated subsystem.

If 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 217.

1.4.10 Thin-provisioned volumes

Volumes can be configured to be *thin provisioned* or *fully allocated*. A thin-provisioned volume behaves regarding application reads and writes 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 Storwize V3500 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 overallocated volumes. It also can be used with FlashCopy. Thin-provisioned volumes also can be used with the mirrored volume feature.

A thin-provisioned volume can be configured to *autoexpand*, which causes Storwize V3500 to automatically expand the real capacity of a thin-provisioned volume as its real capacity is used. This parameter prevents the volume to go offline. Autoexpand attempts to maintain a fixed amount of unused real capacity on the volume. This amount is known as the *contingency capacity*.

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 as soon as 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).

1.4.11 Mirrored volumes

Storwize V3500 provides a function that is called *volume mirroring*, which enables a volume to have two physical copies. Each volume copy can belong to a different storage pool.

When a host system issues a write to a mirrored volume, Storwize V3500 writes the data to both copies. When a host system issues a read to a mirrored volume, Storwize V3500 places it into the primary copy. If one of the mirrored volume copies is temporarily unavailable, the Storwize V3500 automatically uses the alternative copy without any outage for the host system. When the mirrored volume copy is repaired, Storwize V3500 resynchronizes the data.

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

The mirrored volume copy can be of 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. It also can provide a mechanism to migrate fully allocated volumes to thin-provisioned volumes without any host outages.

1.4.12 iSCSI

iSCSI is an alternative means of attaching hosts to the Storwize V3500. The iSCSI function is a software function (not hardware) that is provided by the Storwize V3500 code.

In the simplest terms, iSCSI allows the transport of SCSI commands and data over a 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 expensive FC HBAs and a SAN fabric infrastructure.

A pure SCSI architecture is based on the client/server model. A client (for example, server or workstation) initiates read or write requests for data from a target server (for example, a data storage system).

Commands, which are sent by the client and processed by the server, are put into the Command Descriptor Block (CDB). The server runs a command and completion is indicated by a special signal alert.

The major functions of iSCSI include encapsulation and the reliable delivery of CDB transactions between initiators and targets through the Internet Protocol network, especially over a potentially unreliable IP network.

The following 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 V3700 node represents an iSCSI node and provides statically allocated IP addresses.

Each iSCSI node (that is, an initiator or target) has a unique iSCSI Qualified Name (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 iSCSI qualified name format is defined in RFC3720 and contains (in order) the following elements:

- The string `iqn`.
- A date code that specifies the year and month in which the organization registered the domain or subdomain name that is used as the naming authority string.
- The organizational naming authority string, which consists of a valid, reversed domain or a subdomain name.
- Optionally, a colon (:), followed by a string of the assigning organization's choosing, which must make each assigned iSCSI name unique.

For Storwize V3500, the following IQN is used for the iSCSI target:

iqn.1986-03.com.ibm:2145.<*clustername*>.<*nodename*>

On a Windows server, the following IQN (the name for the iSCSI initiator) can be defined:

iqn.1991-05.com.microsoft:<*computer name*>

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. The alias is independent of the name and does not need to be unique. Because it is not unique, the alias must be used in a purely informational way. It cannot be used to specify a target at login or used during authentication. Both targets and initiators can have aliases.

An iSCSI name provides the correct identification of an iSCSI device irrespective of its physical location. The IQN is an identifier, not an address.

Changing names: Before you change system or node names for a Storwize V3500 clustered system that has servers that are connected to it using SCSI, be aware that because the system and node name are part of the IQN for the Storwize V3500, you can lose access to your data by changing these names. The Storwize V3500 GUI shows a specific warning, but the CLI does not.

The iSCSI session, which consists of a *login phase* and a *full feature phase*, is completed with a special command.

The login phase of the iSCSI is identical to the FC port login process (PLOGI). It is used to adjust various parameters between two network entities and to confirm the access rights of an initiator.

If the iSCSI login phase is completed successfully, the target confirms the login for the initiator; otherwise, the login is not confirmed and the TCP connection breaks.

When the login is confirmed, the iSCSI session enters the full feature phase. If more than one TCP connection was established, iSCSI requires that each command/response pair goes through one TCP connection. Thus, each separate read or write command is run without the need to trace each request for passing separate flows. However, separate transactions can be delivered through separate TCP connections within one session.

For more information about configuring iSCSI, see Chapter 4, “Host configuration” on page 141.

1.5 FlashCopy

FlashCopy makes a copy of a source volume on a target volume. The original content of the target volume is lost. After the copy operation starts, the target volume has the contents of the source volume as they existed at a single point. 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.

The Storwize V3500 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. It also can be used 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. It also eliminates the need to wait for the original copy operation to complete. Storwize V3500 supports multiple targets and thus multiple rollback points.

1.6 Problem management and support

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

1.6.1 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 (which is used to manage the Storwize V3500) must have Internet access. 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.

1.6.2 Event notifications

Storwize V3500 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 Storwize V3500 detects is sent to different a recipient. You can configure Storwize V3500 to send each type of notification to specific recipients or only the alerts, which are important to the system.

1.6.3 SNMP traps

SNMP is a standard protocol for managing networks and exchanging messages. Storwize V3500 can send SNMP messages that notify personnel about an event. You can use an SNMP manager to view the SNMP messages that Storwize V3500 sends. You can use the management GUI or the Storwize V3500 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 Storwize V3500. This file can be used with SNMP messages from all versions of Storwize V3500 Software.

1.6.4 Syslog messages

The syslog protocol is a standard protocol that is used for forwarding log messages from a sender to a receiver on an IP network. The IP network can be IPv4 or IPv6. Storwize V3500 can send syslog messages that notify personnel about an event. Storwize V3500 can transmit syslog messages in expanded or concise format. You can use a syslog manager to view the syslog messages that Storwize V3500 sends. Storwize V3500 uses the User Datagram Protocol (UDP) to transmit the syslog message. You can use the management GUI or the Storwize V3500 CLI to configure and modify your syslog settings.

1.6.5 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 Storwize V3500 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 Storwize V3500 clustered system IP address. You can then use the management GUI or the Storwize V3500 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 Storwize V3500 CLI.

1.7 Useful Storwize V3500 websites

The following websites can provide more information about Storwize V3700:

- Storwize V3500 home page:

  http://www-947.ibm.com/support/entry/portal/overview/hardware/system_storage/disk_systems/entry-level_disk_systems/ibm_storwize_v3500

- Storwize V3500 Configuration Limit and Restrictions:

  http://www-01.ibm.com/support/docview.wss?uid=ssg1S1004197

- Storwize V3500 Online Information Center:

  http://pic.dhe.ibm.com/infocenter/storwize/v3500_ic/index.jsp

- IBM Redbooks publications about Storwize V3500:

  http://www.redbooks.ibm.com/cgi-bin/searchsite.cgi?query=V3500&SearchOrder=1&SearchFuzzy=

1.8 Storwize V3500 learning videos on YouTube

Videos are available on YouTube that describe the IBM Storwize GUI and are available at the URLs 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 at not specifically for the Storwize V3500. However, because the interface to the Storwize V7000 and Storwize V3700 is similar, the videos provide an overview of the functions and features.

2

Initial configuration

This chapter provides a description of the initial configuration steps for the IBM Storwize V3500.

This chapter includes the following topics:

- Planning for IBM Storwize V3500 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 V3500 Quick Installation Guide Version 6.4.1*, GC27-4218.
- ☐ For information about planning the IBM Storwize V3500 environment, see this website:

  http://pic.dhe.ibm.com/infocenter/storwize/V3500_ic/index.jsp?topic=%2Fcom.ibm.storwize.V3500.641.doc%2Fsvc_webplanning_21pb8b.html

2.2 SAN configuration planning

The recommended SAN configuration is composed of a minimum of two fabrics with all host ports, and any ports on external storage systems to be virtualized by IBM Storwize V3500. The IBM Storwize V3500 ports are evenly split between the two fabrics to provide redundancy in the event one of the fabrics goes offline (planned or unplanned).

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

After the IBM Storwize V3500, hosts, and optional external storage systems are connected to the SAN fabrics, zoning must be implemented.

In each fabric, create a zone with only the four IBM Storwize V3500 worldwide port names (WWPNs), two from each node canister. If there is an external storage system that is virtualized, create a zone in each fabric with the four IBM Storwize V3500 WWPNs, two from each node canister. Also, include 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 in each fabric with the host WWPN and one WWPN from each node canister in the IBM Storwize V3500 system.

Maximum ports or WWPNs: IBM Storwize V3500 supports a maximum of 16 ports or WWPNs from a virtualized external storage system.

Figure 2-1 shows how to cable devices to the SAN. See this example as we describe the zoning.

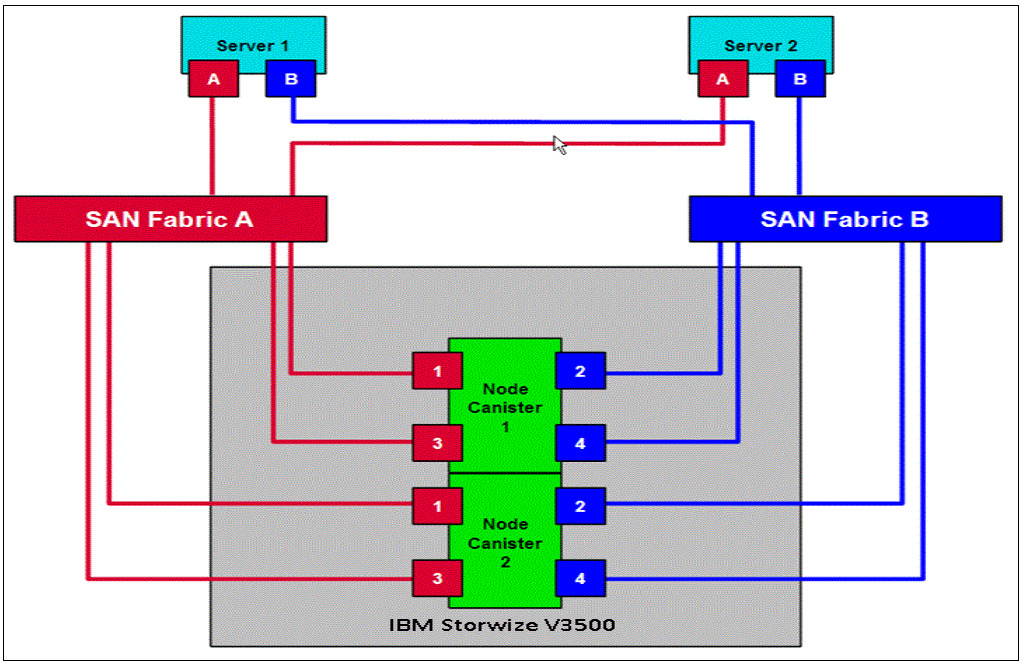

Figure 2-1 SAN cabling and zoning diagram

Create a host/IBM Storwize V3500 zone for each server to which volumes are mapped to and from the clustered system, as shown in the following examples in Figure 2-1:

- Zone Server 1 port A (RED) with all node port 1s
- Zone Server 1 port B (BLUE) with all node port 2s
- Zone Server 2 port A (RED) with all node port 3s
- Zone Server 2 port B (BLUE) with all node port 4s

Verify interoperability with which the IBM Storwize V3500 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 V3500.

IBM Storwize V3500 port login maximum that is listed in restriction document must not be exceeded.

Connectivity issues: If you have any connectivity issues between IBM Storwize V3500 ports and Brocade SAN Switches or Directors at 8 Gbps, see the correct setting of the fillword port config parameter in the Brocade operating system at this website:

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

2.3 Direct attach planning

IBM Storwize V3500 can be used with direct attach Fibre Channel (FC) server configuration. The recommended configuration for direct attachment is to have at least one FC cable from the server that is connected to each node of the IBM Storwize V3500. This configuration provides redundancy if one of the nodes goes offline, as shown in Figure 2-2.

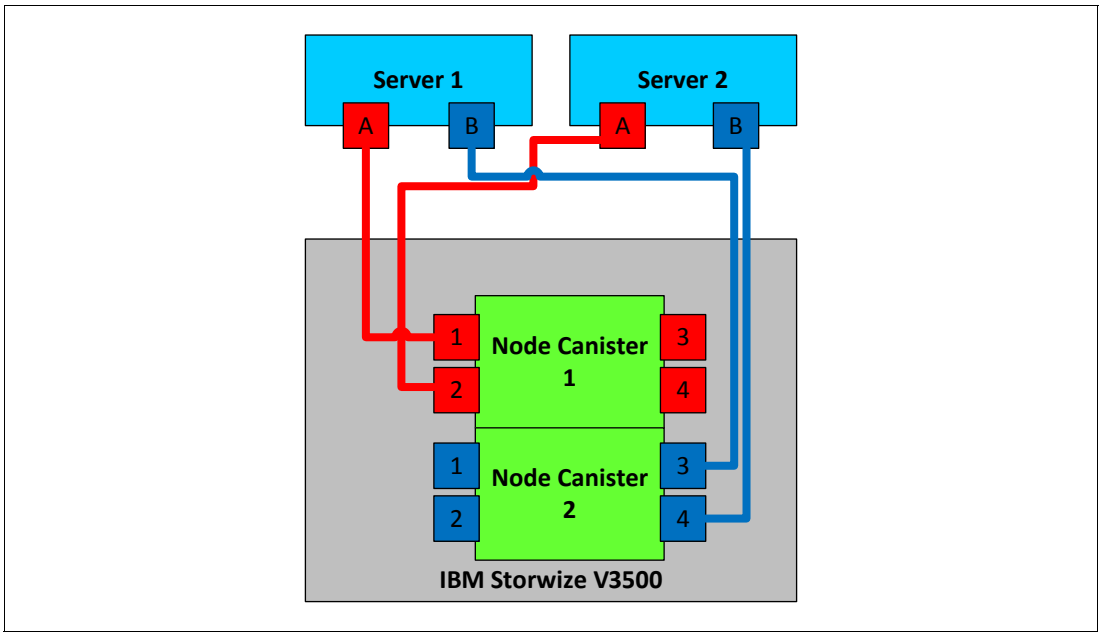

Figure 2-2 Direct attach server configuration

Verify direct attach interoperability with the IBM Storwize V3700 and the supported server operating systems by following the requirements that provided at this website:

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

2.4 LAN configuration planning

There are two Ethernet ports per node canister that are available for connection to the LAN on an IBM Storwize V3500 system.

Ethernet port 1 is used 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 both management and iSCSI host use.

Figure 2-3 shows a possible IP configuration of the Ethernet ports on the IBM Storwize V3500 system.

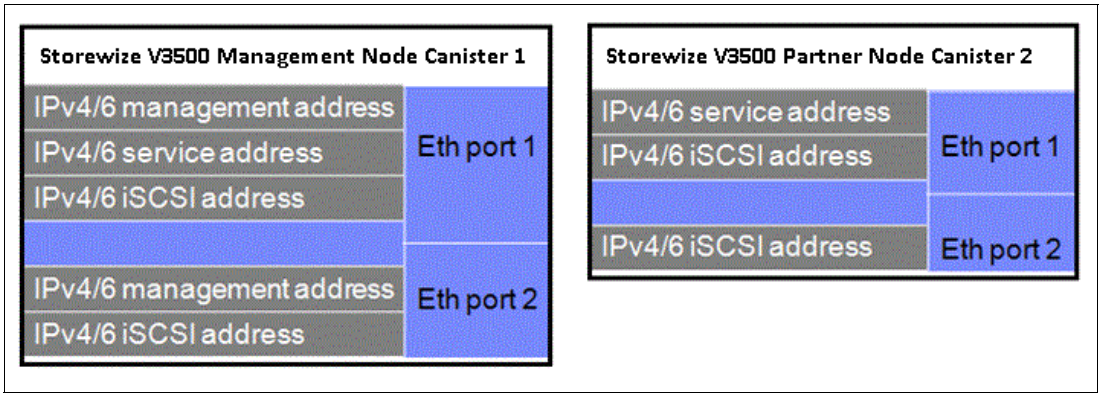

Figure 2-3 Example of IP address configuration options per node canister

IP management addresses: The IP management address that is shown on node canister 1 is an address on the configuration node. If there is a failover, this address transfers to node canister 2. 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).

2.4.1 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 V3500 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. This configuration allows for a backup management IP address to be configured for access if necessary.

Figure 2-4 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. Therefore, they are associated with only one node, which is considered the configuration node.

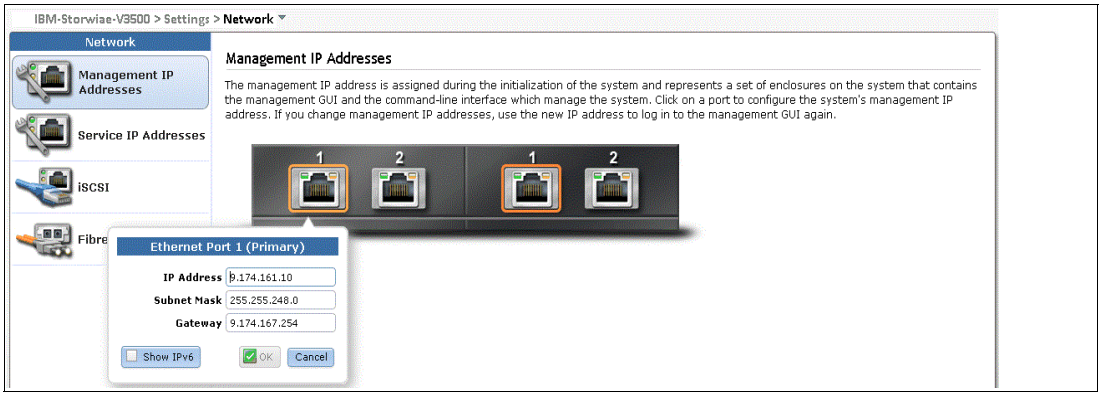

Figure 2-4 Ethernet ports that are available for configuration

2.4.2 Service IP address considerations

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

Figure 2-5 shows a logical view of the Ethernet ports that are available for configuring the service IP addresses. Only port one on each node can be configured with a service IP address.

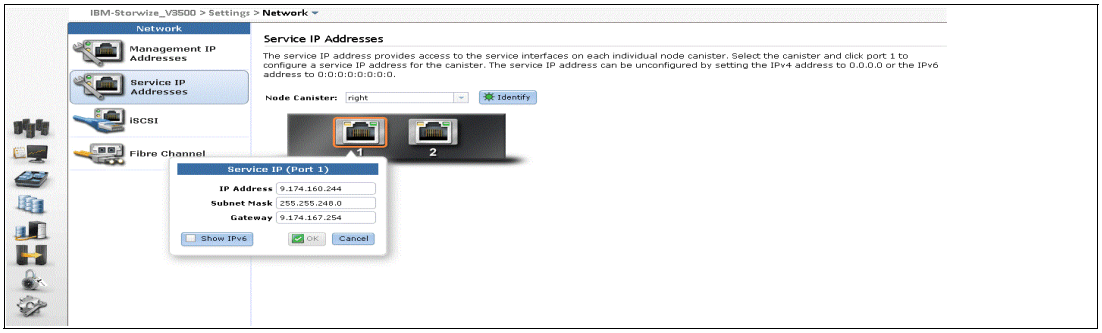

Figure 2-5 Service IP addresses available for configuration

2.5 Host configuration planning

Hosts should have two FC connections for redundancy, although IBM Storwize V3500 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 V3500 volumes. These paths are available because that single connection must be zoned with at least one FC port per node. Therefore, a multipath driver is required.

SAN Boot is supported by IBM Storwize V3500. The requirements are listed on the IBM Storwize V3500 support matrix and configuration instructions are provided in the *IBM Storwize V3500 Host Attachment Guide*, which can be found at this website:

http://pic.dhe.ibm.com/infocenter/storwize/v3500_ic/index.jsp?topic=%2Fcom.ibm.storwize.v3500.641.doc%2Fsvc_hostattachmentmain.html

Verify that the hosts that access volumes from the IBM Storwize V3500 meet the requirements that are found at this website:

http://www-01.ibm.com/support/docview.wss?uid=ssg1S1004197

Multiple OS are supported by IBM Storwize V3500. For more information about HBA/Driver/multipath combinations, see this website:

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

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 V3500.
- HBA BIOS, device drivers, firmware, and multipathing drivers are at the levels that are supported by IBM Storwize V3500.
- If boot from SAN is required, ensure that it is supported for the operating systems to be deployed.
- If host clustering is required, ensure that it is supported for the operating systems to be deployed.
- All direct connect hosts should have the HBA set to point-to-point.

For more information, see Chapter 4, “Host configuration” on page 141.

2.6 Miscellaneous configuration planning

During the initial setup of the IBM Storwize V3500 system, the installation wizard asks for various information that you should have available during the installation process. Several of these fields are mandatory for you to continue the initial configuration.

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

- ☐ Document the current LAN NTP server IP address that is used for synchronizing 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 V3500 system for IBM service (for example, Building 22, first floor).
  - ☐ SMTP or email server address to direct alerts to and from the IBM Storwize V3500.

  For the Call Home service to work, the IBM Storwize V3500 system must have access to an SMTP server on the LAN that can forward emails to the default IBM service address callhome0@de.ibm.com.

- ☐ Email address of local administrators who must be notified of alerts.
- ☐ IP address of SNMP server to direct alerts to if wanted (for example, operations or help desk).

After the IBM Storwize V3500 initial configuration, you might want to add other users who can manage the system. You can create as many users as you need, but only the following roles generally are configured for users:

- Security Admin
- Administrator
- Copyoperator
- Service
- Monitor

The user in the role of Security Admin role can perform any function on the IBM Storwize V3500.

The user in the Administrator role can perform any function on the IBM Storwize V3500 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. However, 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, 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. IBM service personnel in this role can view anything on the system, as with the monitor role, plus perform service-related commands, such as adding a node back to the system after being serviced or including disks that were excluded.

2.7 System management

The graphical user interface (GUI) is used to configure, manage, and troubleshoot the IBM Storwize V3500 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.

IT also offers advanced functions, such as FlashCopy. A command-line interface (CLI) for the IBM Storwize V3500 system also is available.

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

2.7.1 Graphical user interface

A web browser is used for GUI access. You must use a supported web browser to access the management GUI. For a list of supported web browsers, see the *Checking your web browser settings for the management GUI* document at this website:

http://pic.dhe.ibm.com/infocenter/storwize/v3500_ic/index.jsp?topic=%2Fcom.ibm.storwize.v3500.641.doc%2Fsvc_configuringbrowser_1obg15.html

To open the Management GUI from any web browser, complete the following steps:

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*/ (for 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/v3500_ic/index.jsp?topic=%2Fcom.ibm.storwize.v3500.641.doc%2Ftbrd_usbgui_1936tw.html

After the initial configuration that is described in 2.9, “Initial configuration” on page 41 is completed, the IBM Storwize V3500 Welcome window opens, as shown in Figure 2-6.

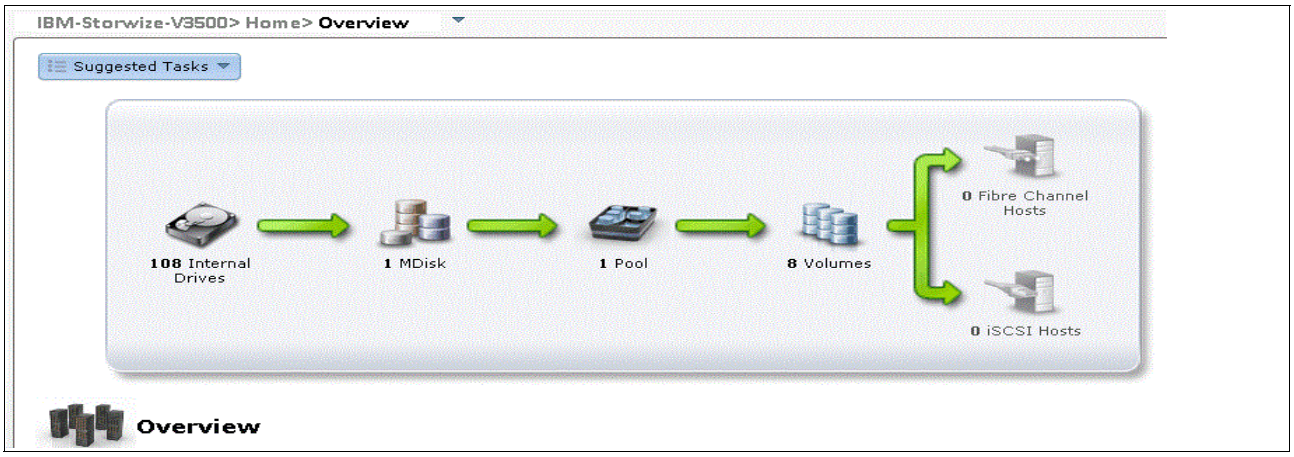

Figure 2-6 Setup Wizard: Welcome window.

2.7.2 Command-line interface

The CLI is a flexible tool for system management and uses the Secure Shell (SSH) protocol. A public/private SSH key pair is optional for SSH access. Appendix A, “Command-line interface setup and SAN Boot” on page 483 describes how to set up SSH access for Windows, Linux, or UNIX systems. 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_2072:IBM-Storwize-V3500:admin>lsuser
id name      password ssh_key remote usergrp_id usergrp_name
0  superuser yes      yes     no     0          SecurityAdmin
1  admin     yes      yes     no     1          Administrator
2  Andrew    yes      no      no     2          CopyOperator
3  Justin    yes      no      no     3          Service
4  Deiter    yes      no      no     4          Monitor
5  Uwe       yes      no      no     0          SecurityAdmin
6  Nancy     yes      no      no     0          SecurityAdmin
IBM_2072:IBM-Storwize-V3500:admin>
```

The initial IBM Storwize V3500 system setup should be done by using the process and tools that are described in 2.8, “First-time setup” on page 36.

2.8 First-time setup

This section describes how to perform a first-time IBM Storwize V3500 system setup.

IBM Storwize V3500 uses an easy-to-use initial setup that is contained within a USB key. The USB key is delivered with each storage system and contains the initialization application called `InitTool.exe`. A system management IP address, the subnet mask, and the network gateway address are required. The initialization application creates a configuration file on the USB key.

The IBM Storwize V3500 starts the initial setup when you plug in the USB key with the new file in to the storage system.

USB key: If you cannot find the official USB key that is supplied with the IBM Storwize V3500, you can use any USB key that you have available. Download and copy the `initTool.exe` application from the IBM Storwize V3500 Support at this website:

`http://www.ibm.com/storage/support/Storwize/V3500`

The USB key contains the `initTool.exe` file, as shown in Figure 2-7.

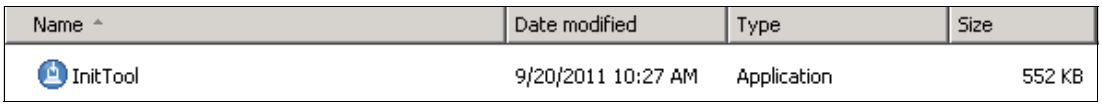

Figure 2-7 InitTool program in the USB key

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

1. Plug the USB key into an 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.exe` file. The opening window of the tool is shown in Figure 2-8. After the tool is started, select **Next** and select **Create a system**, as shown in Figure 2-8.

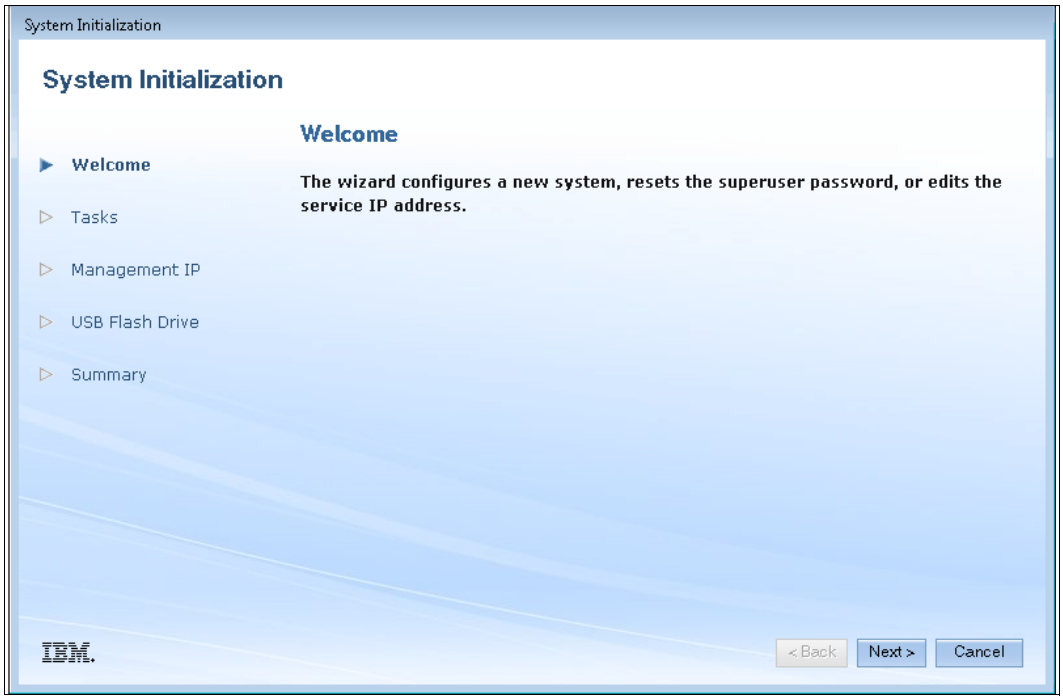

Figure 2-8 System Initialization: Welcome window

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

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

d. Enter: sh InitTool.sh

2. The options for creating a system are shown in Figure 2-9.

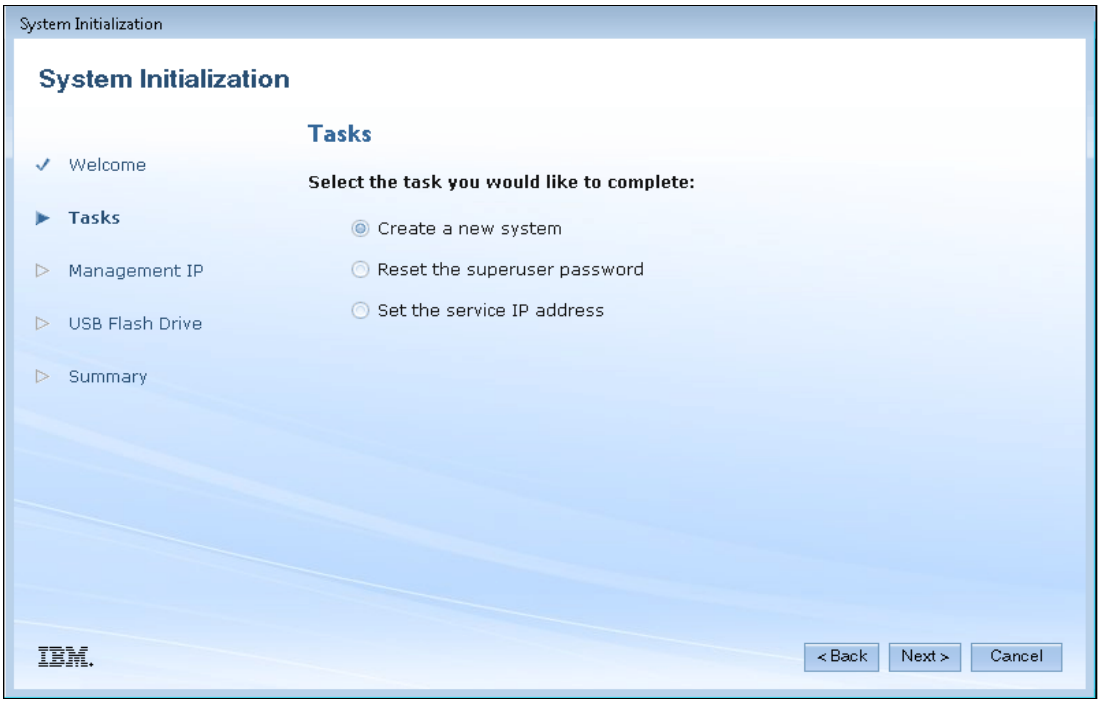

Figure 2-9 System Initialization: Create a system

Important: The other options that are available in the initialization tool are used to recover the service IP address or the superuser password.

3. Set the Management IP address, as shown in Figure 2-10.

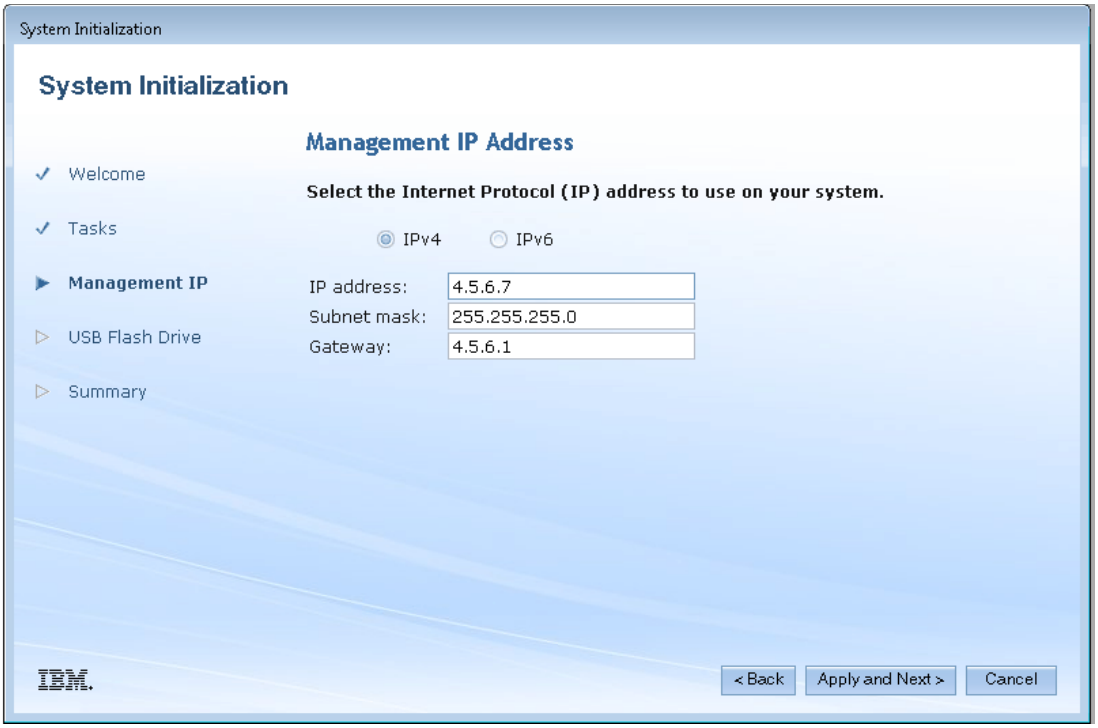

Figure 2-10 System Initialization: Management IP

4. Click **Apply and Next** to show the USB flash drive instructions, as shown in Figure 2-11.

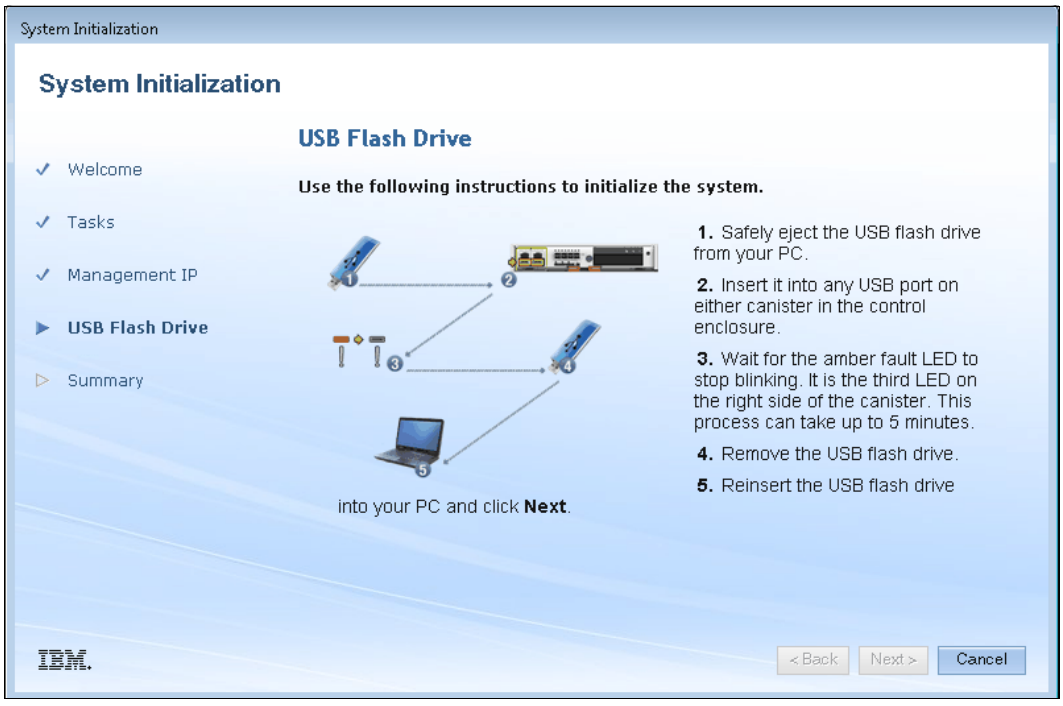

Figure 2-11 Initialization application: Finalizing the initial setup

Follow the instructions through the summary portion. After the initialization completes, click **Finish**. The application creates a file that is called `satask.txt` on the USB key, as shown in Figure 2-12.

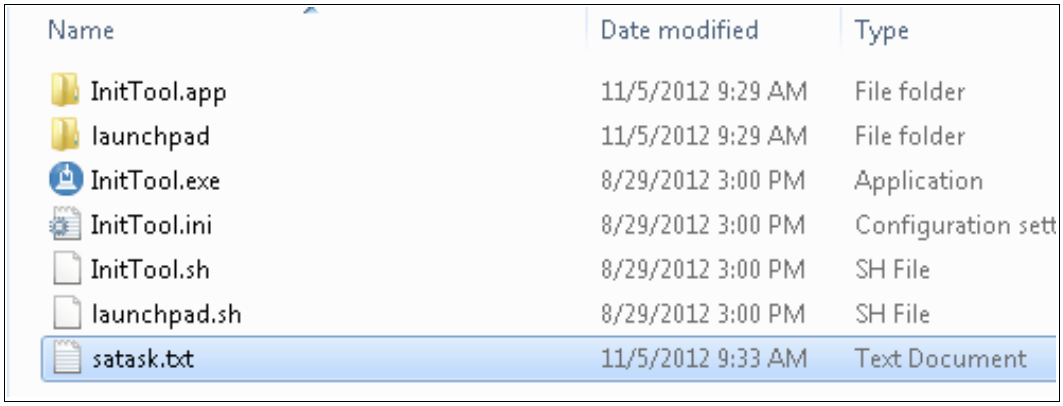

Figure 2-12 The satask.txt file that is created during the InitTool procedure

The contents of this file are similar to the following command:

```
satask mkcluster -clusterip 4.5.6.7 -gw 4.5.6.1 -mask 255.255.255.0
```

5. Unplug the USB key from your Windows system and plug it into the IBM Storwize V3500 USB connector slot. The storage system automatically completes the following tasks:
   a. Detect the key.
   b. Read the `satask.txt` file and run its contents.
   c. Delete the `satask.txt` file.
   d. Write the response to a new file that is called `satask_result.html`

   Clustered system creation: While the clustered system is created, the amber fault LED on the node canister flashes. When the amber fault LED stops flashing, remove the USB key from IBM Storwize V3500 and insert it in your system to check the results.

6. After this task completes successfully, the initial setup is done. If you have a network connection to the Storwize system, the system Management GUI is displayed. The IBM Storwize V3500 is available for further configuration changes by using the newly defined configuration address.

Each node has two Ethernet ports that can be used for system management. Ethernet port 1 is used for system management and must be configured and connected on both nodes. The use of Ethernet port 2 is optional.

Each IBM Storwize V3500 clustered system has one or two system IP addresses. If the configuration node fails, the system IP addresses are transferred to another node in the same clustered system.

Important: The first system management IP address always uses port 1. Always connect port 1 for all node canisters to the management network.

2.9 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

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 is defined in Figure 2-10 on page 39. Enter the default superuser password <passw0rd> (where 0 = zero), as shown in Figure 2-13.

Figure 2-13 Setup Wizard: Login

2. Set up the system name, as shown in Figure 2-14.

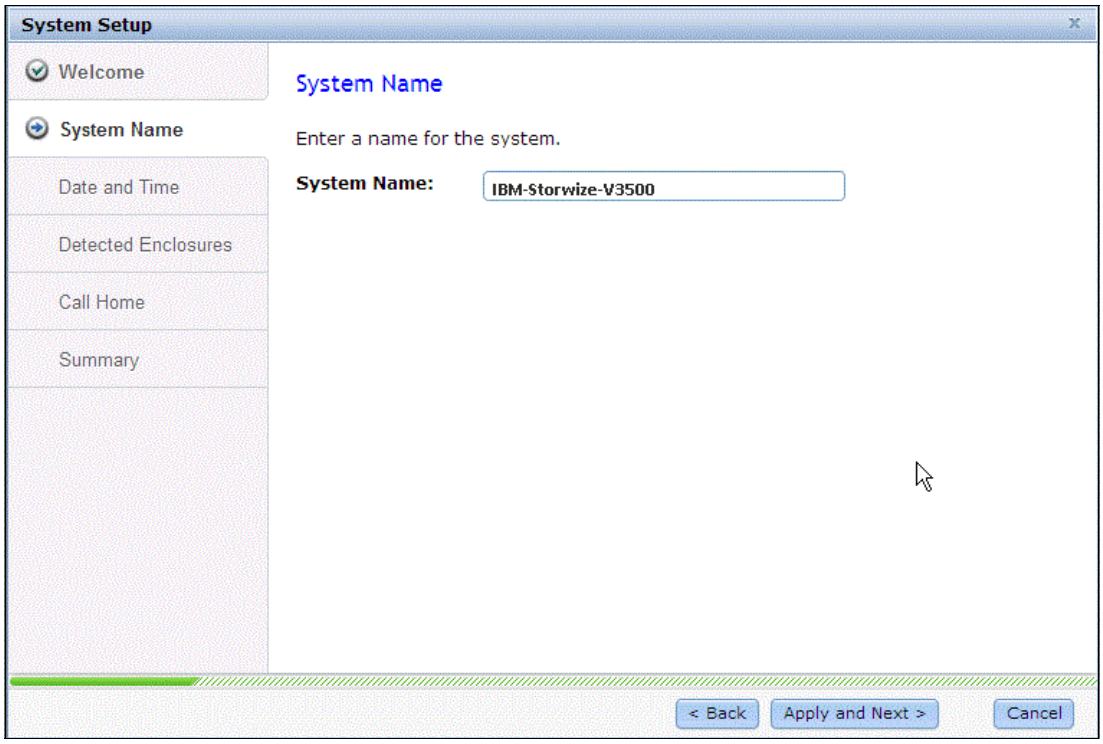

Figure 2-14 Setup Wizard: Insert system name.

USB flash drive: If the GUI is not accessible, you can use the USB key. When a USB key is plugged into a node canister, the node canister software searches for a text file that is named `satask.txt` in the root directory. If the software finds the file, it attempts to run a command that is specified in the file. When the command completes, a file that is called `satask_result.html` is written to the root directory of the USB key. If this file does not exist, it is created.

3. Set the date and time, as shown in Figure 2-15.

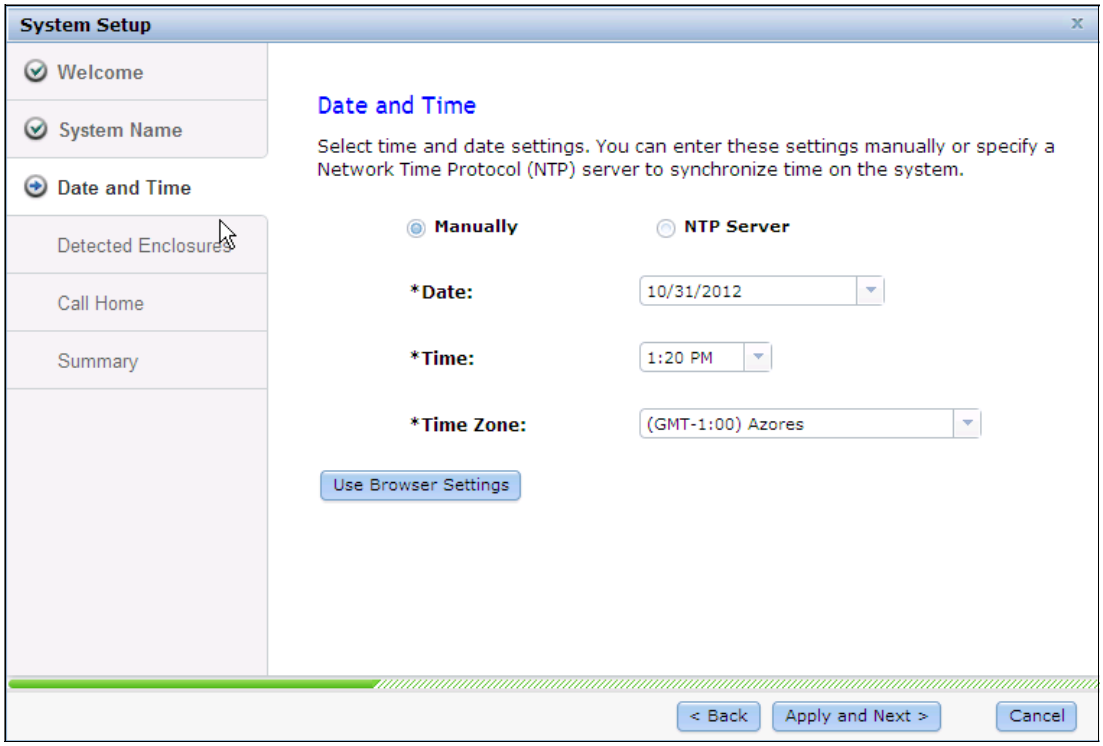

Figure 2-15 Setup Wizard - date and time.

4. The configuration wizard continues with the hardware configuration. Verify the hardware, as shown in Figure 2-16.

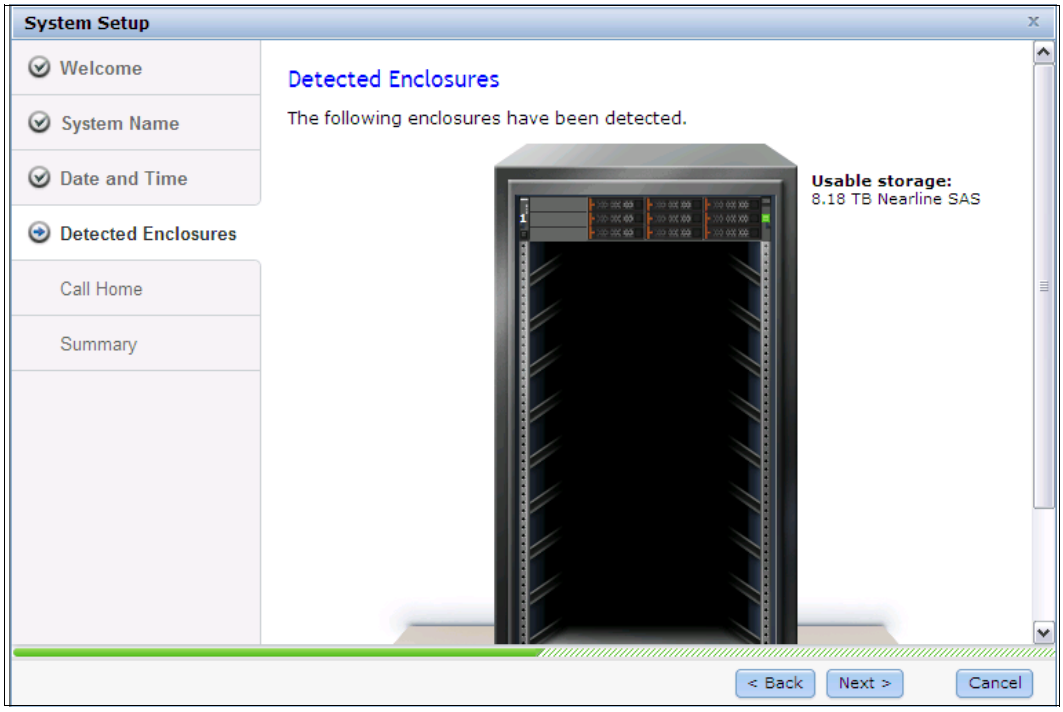

Figure 2-16 Setup wizard: Verify the detected hardware

5. When an array is created, the array members are synchronized with each other by a background initialization process, as shown in Figure 2-17.

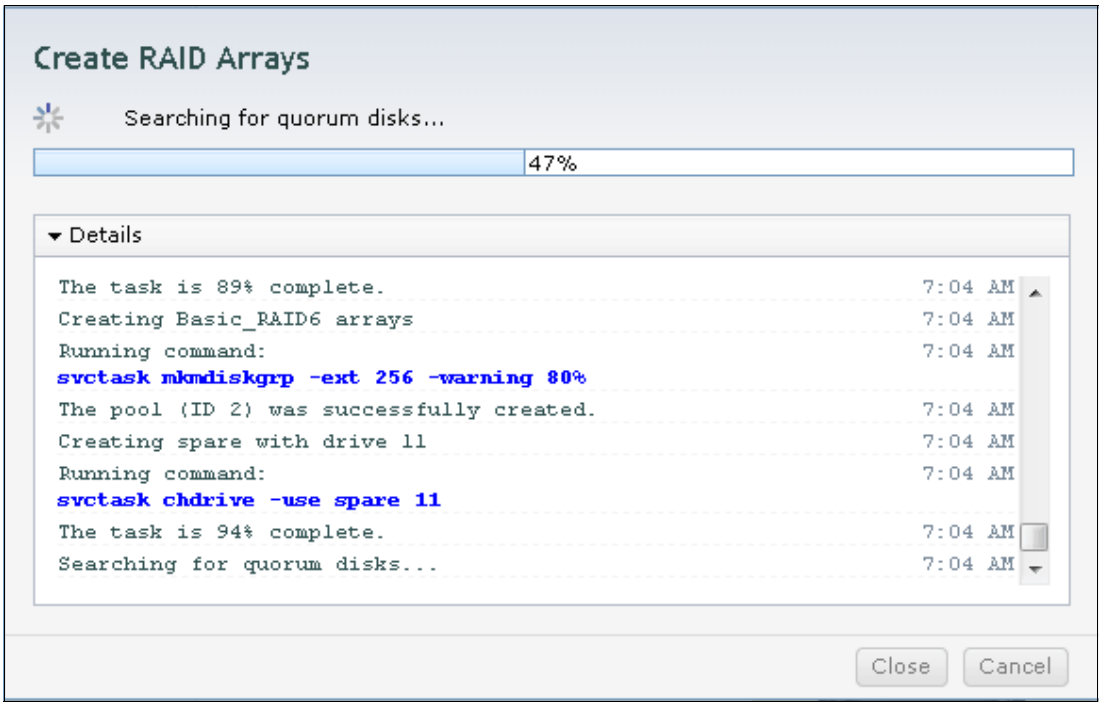

Figure 2-17 Preset RAID configuration

The management GUI includes several pre-established configuration options to help you save time during the configuration process. These preset options incorporate commonly used parameters.

Presets are available for creating volumes and FlashCopy mappings and for setting up RAID configuration.

For a complete description of the management GUI, start the e-Learning tutorial module by selecting *Watch e-Learning: Overview* after the system finishes the configuration process (as shown in Figure 2-18) or see Chapter 7, “Storage pools” on page 277.

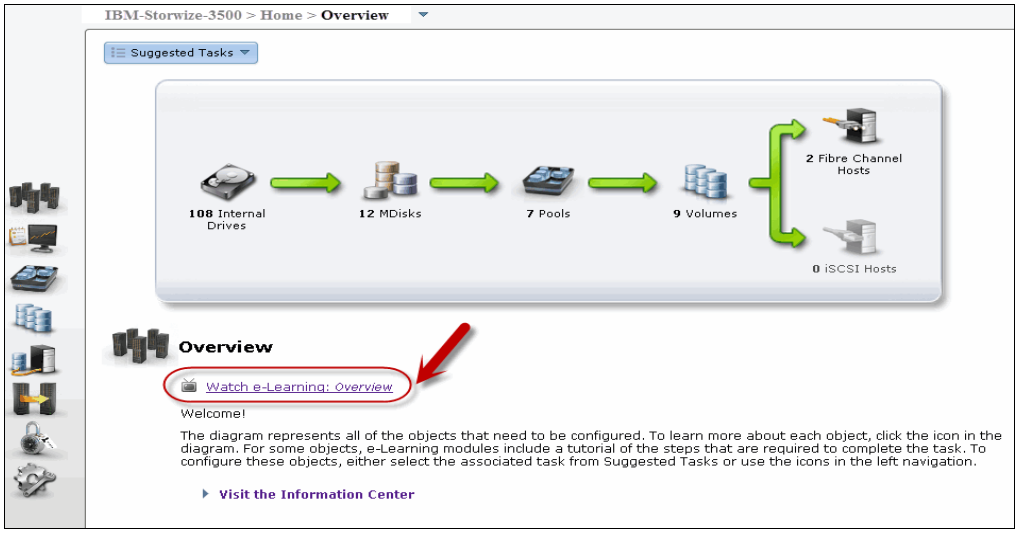

Figure 2-18 Watch e-Learning: Overview

For more information about the presets, see this website:

http://pic.dhe.ibm.com/infocenter/storwize/v3500_ic/index.jsp?topic=%2Fcom.ibm.storwize.v3500.641.doc%2Fsvc_presets_09271645.html

2.9.1 Configuring Call Home, email alert, and inventory

If your system is within warranty, or you have a hardware maintenance agreement, 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 you are configuring Call Home, use 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 V3500 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.

Complete the following steps to configure the Call Home and email alert event notification in IBM Storwize V3500:

1. Click **Call Home**, as shown in Figure 2-19.

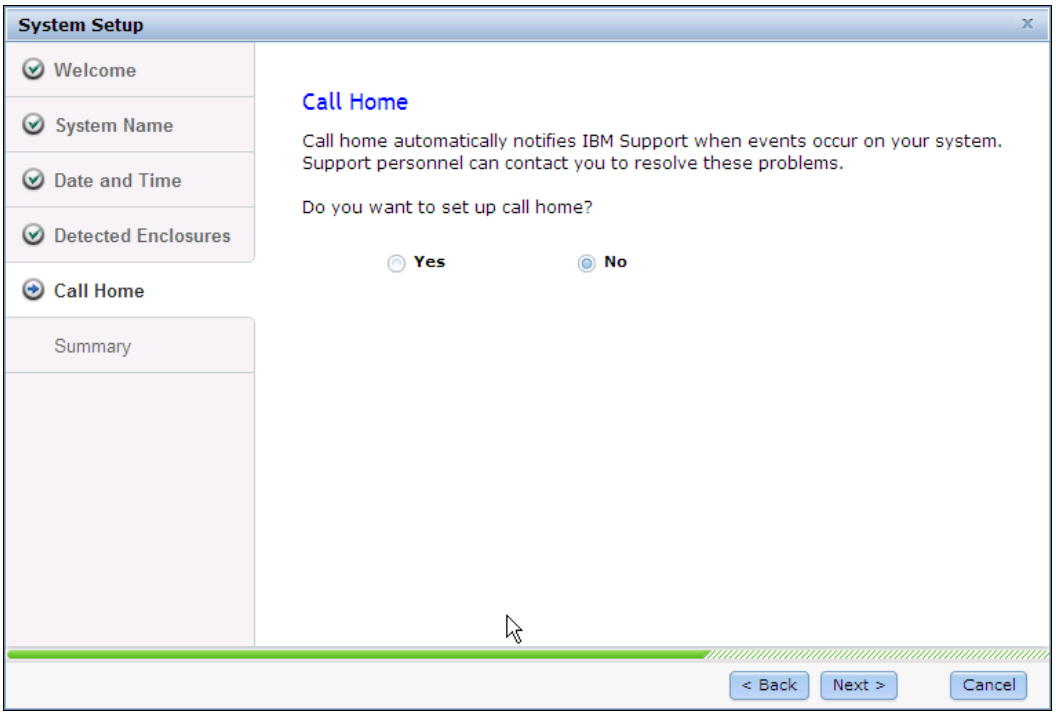

Figure 2-19 System setup: Call Home

2. Follow the prompts for the Call Home system location, as shown in Figure 2-20 (all fields are mandatory).

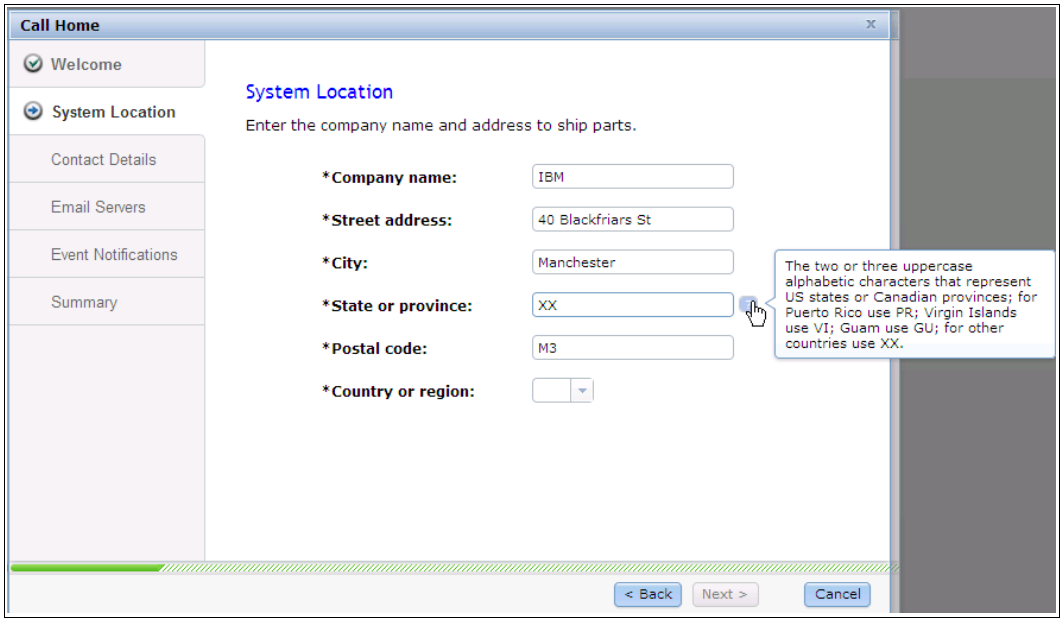

Figure 2-20 Call Home: System location

3. Enter the mandatory contact information, as shown in Figure 2-21.

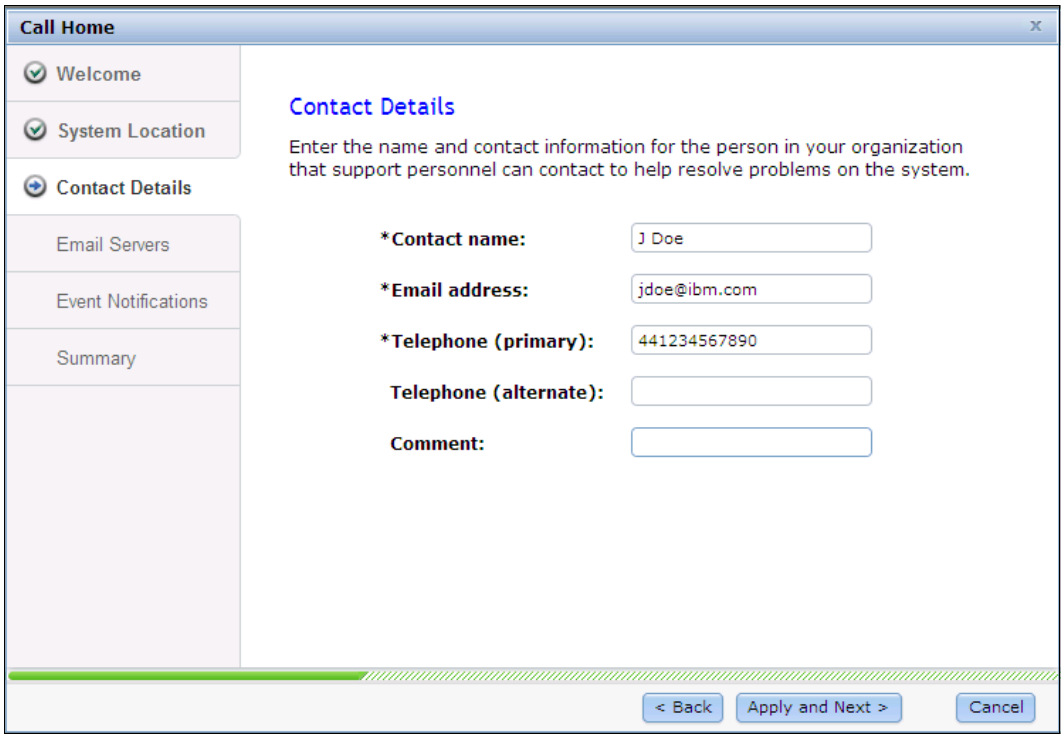

Figure 2-21 Call Home: Contact information

4. Configure your email server IP settings, as shown in Figure 2-22.

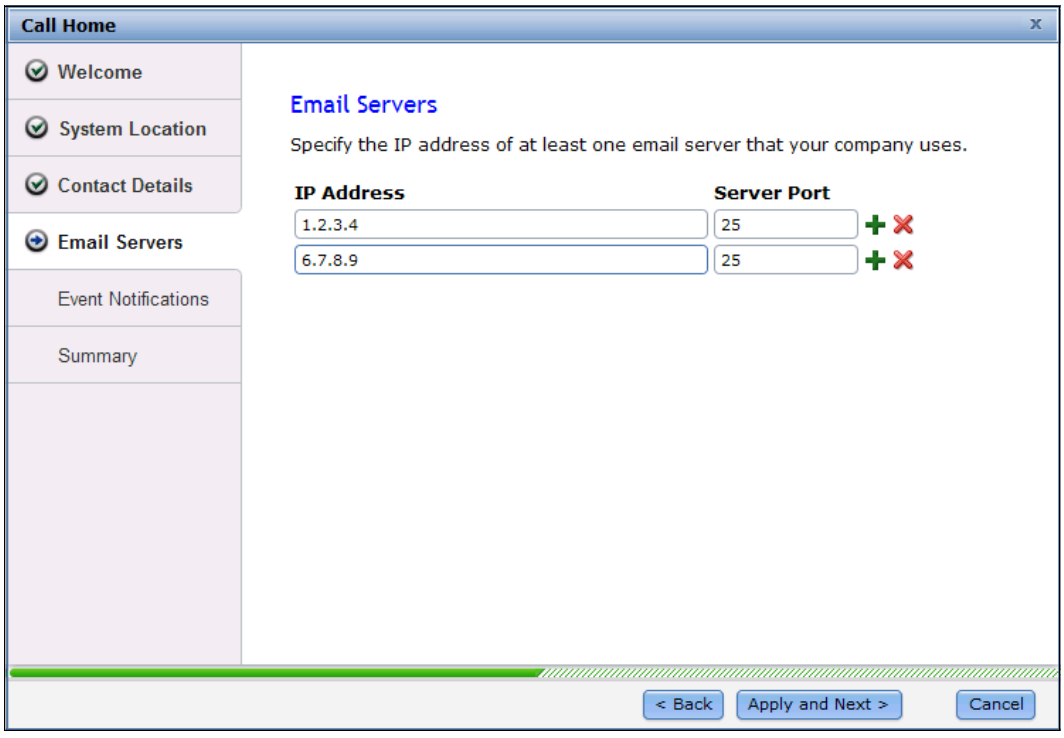

Figure 2-22 Call Home: Email Servers IP address

2.9.2 Configuring event notifications

Each event that IBM Storwize V3500 detects is assigned a notification type of Error, Warning, or Information. Complete the following steps to configure IBM Storwize V7000 to send each type of notification to specific recipients:

1. Set up to whom you want to receive these notifications. Figure 2-23 shows that you can configure alerts or all events to be sent to specific users.

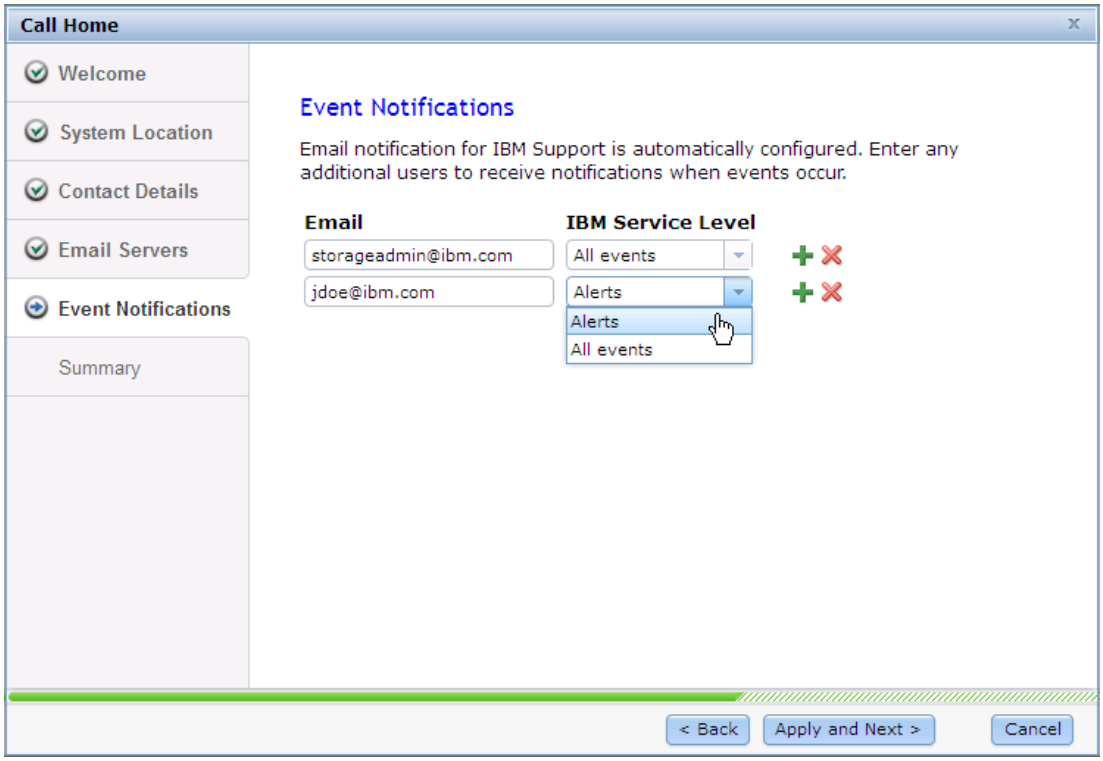

Figure 2-23 Event notifications

2. After you specify your event notification, select **Apply and Next**. The Summary window opens, as shown in Figure 2-24.

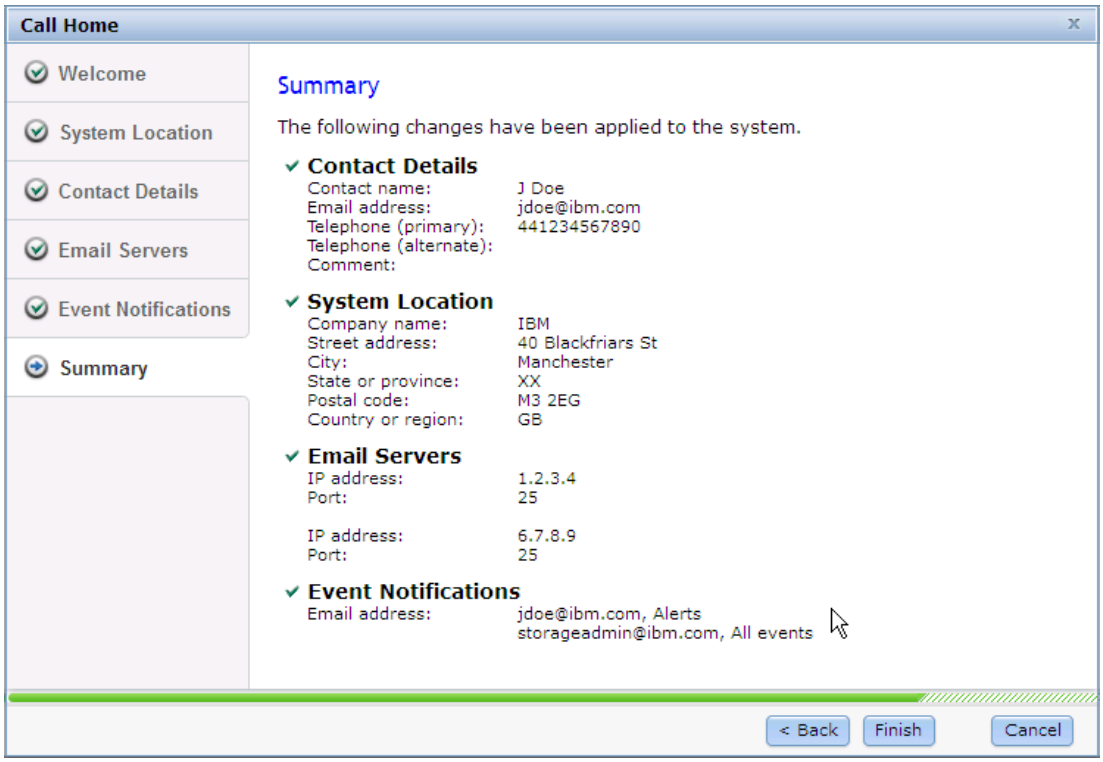

Figure 2-24 Event Notification summary

Choosing finish here sets up the email settings for the event notifications and you see the window that is shown in Figure 2-25.

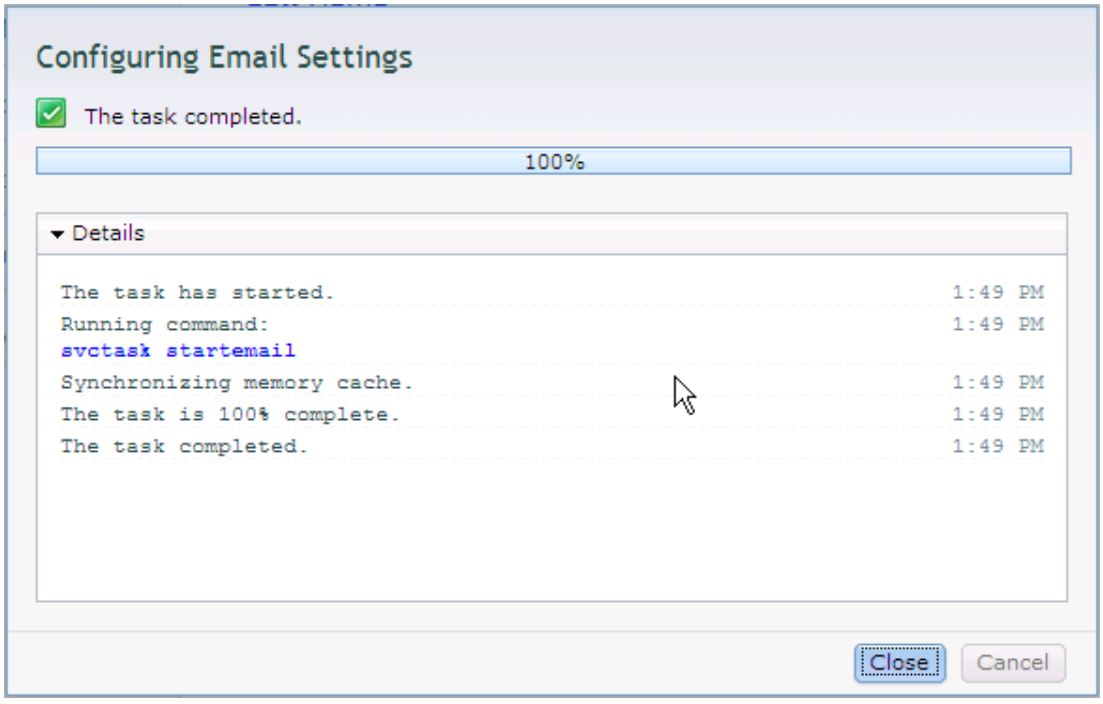

Figure 2-25 Configuring Email Settings

After the email settings are configured, the Summary window opens, as shown in Figure 2-26.

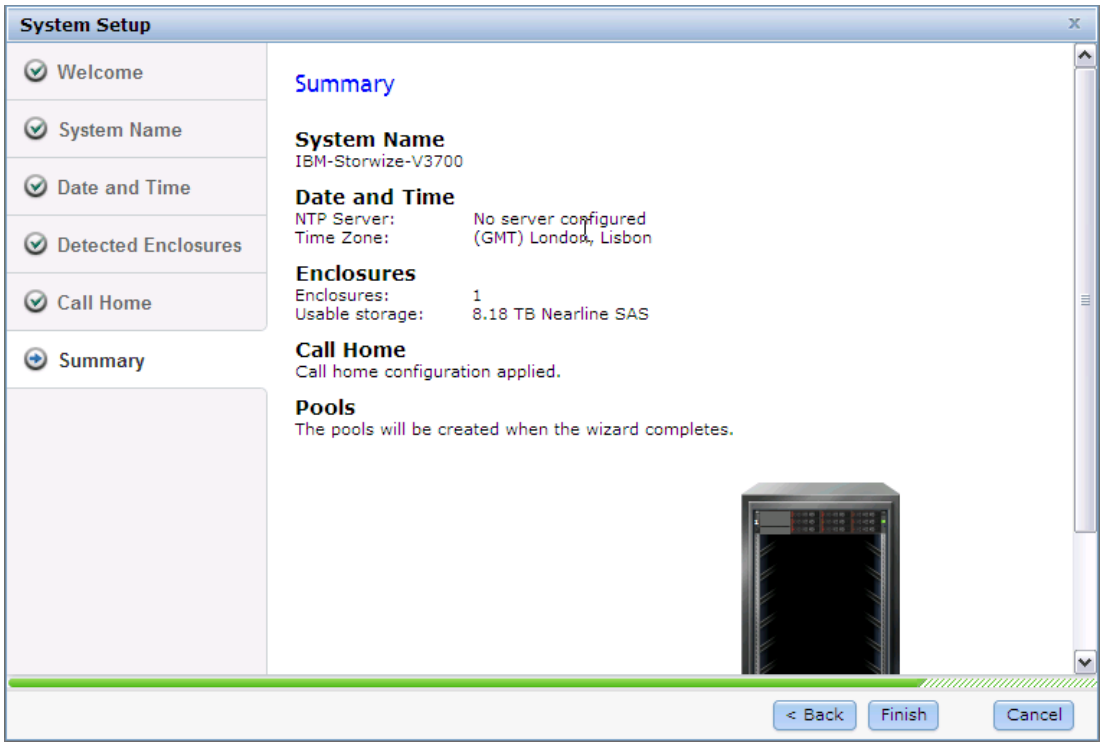

Figure 2-26 Summary

2.9.3 Service Assistant Tool

The V3500 is configured initially 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 initially configure the Storwize V3500 system:

- The Inittool Program (for more information, see 2.8, “First-time setup” on page 36)
- The Service Assistant Tool: The SA interface is a web 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 SA GUI and CLI are used to service a node even if it cannot perform I/O. This interface provides full remote access to the information and tasks that are available in a standard SAN Volume Controller front panel interface. 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 of node error codes and hints about what to do to fix the node error.

Important: The SA can be accessed only by using the superuser account.

The SA GUI is available by using a service assistant IP address on each node. The SA GUI is also accessed through the cluster IP addresses by appending service to the cluster management URL.

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

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

For example:

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

When you are accessing SA by using the `<cluster address>/service`, the configuration node canister SA GUI login window is opened. The SA interfaces can view status and run service tasks on other nodes, in addition to the node where the user is connected.

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

The default value IP addresses of canister 2 is 192.168.70.122 with a subnet mask of 255.255.255.0.

The IBM Storwize V3500 Login Assistant Login Screen is shown Figure 2-27.

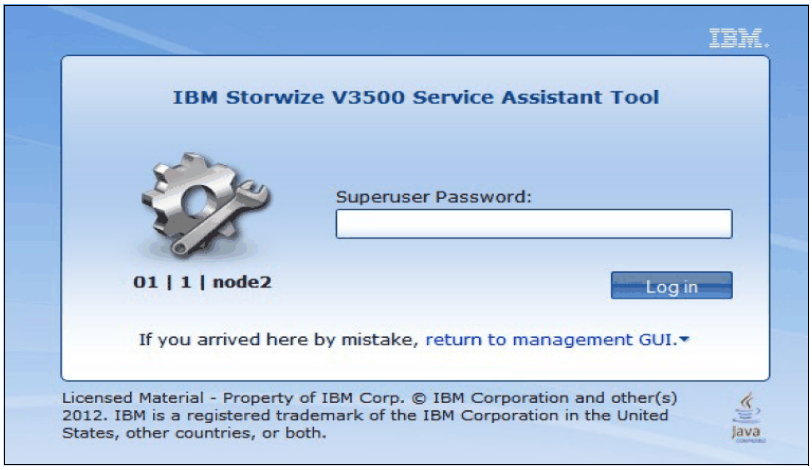

Figure 2-27 Service Assistant Login Screen

After you are logged in, you see the Service Assistant home window, as shown in Figure 2-28. The current canister node is displayed in the upper left corner of the GUI.

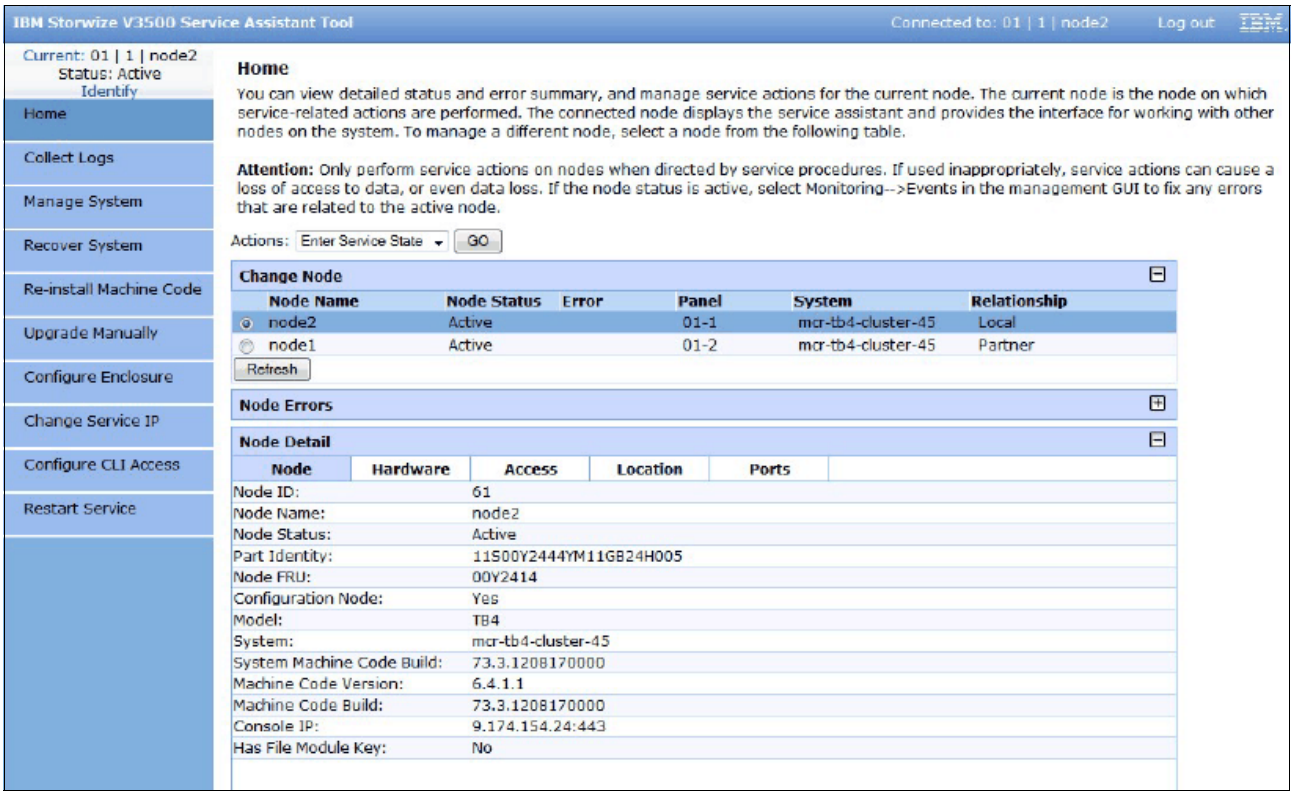

Figure 2-28 Service Assistant Home Screen

Each node can have a single SA IP address on Ethernet port 1. The SA is accessed by the management IP. If the system is down, the only other method of communicating with the node canisters is through the SA IP address.

It is recommended that these IP addresses are configured on all Storwize V3500 node canisters.

The SA GUI provides access to service procedures and displays the status of the node canisters. It is recommended to run these procedures only if directed.

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

http://pic.dhe.ibm.com/infocenter/storwize/v3500_ic/index.jsp?topic=%2Fcom.ibm.storwize.v3500.641.doc%2Ftbrd_sagui_1938wd.html

3

Graphical user interface overview

This chapter provides an overview of the Storwize V3500 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 Storwize V3500 management GUI. It covers topics such as supported browsers, log in modes, and the layout of the Overview panel.

3.1.1 Support browsers

The Storwize V3500 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. Table 3-1 lists the web browsers that are supported by the management GUI. For the latest version of this table see *Quick Installation Guide for the IBM Storwize V3500*, GC27-4218-00.

Table 3-1 Supported web browsers

Browser	Supported versions
Mozilla Firefox	3.5 or later
Microsoft Internet Explorer	8.0 or later

Important: The web browser requirements and recommended configuration settings to access the Storwize V3500 management GUI can be found in the Storwize V3500 Information Center at this website:

http://pic.dhe.ibm.com/infocenter/storwize/v3500_ic/index.jsp

3.1.2 Accessing the management GUI

To log in to the management GUI, open a supported web browser and enter the management IP address or Hostname of the Storwize V3500. The login panel is displayed, as shown in Figure 3-1.

Figure 3-1 Storwize V3500 login panel

Default user name and password: Use the following information to log in to the Storwize V3700 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.

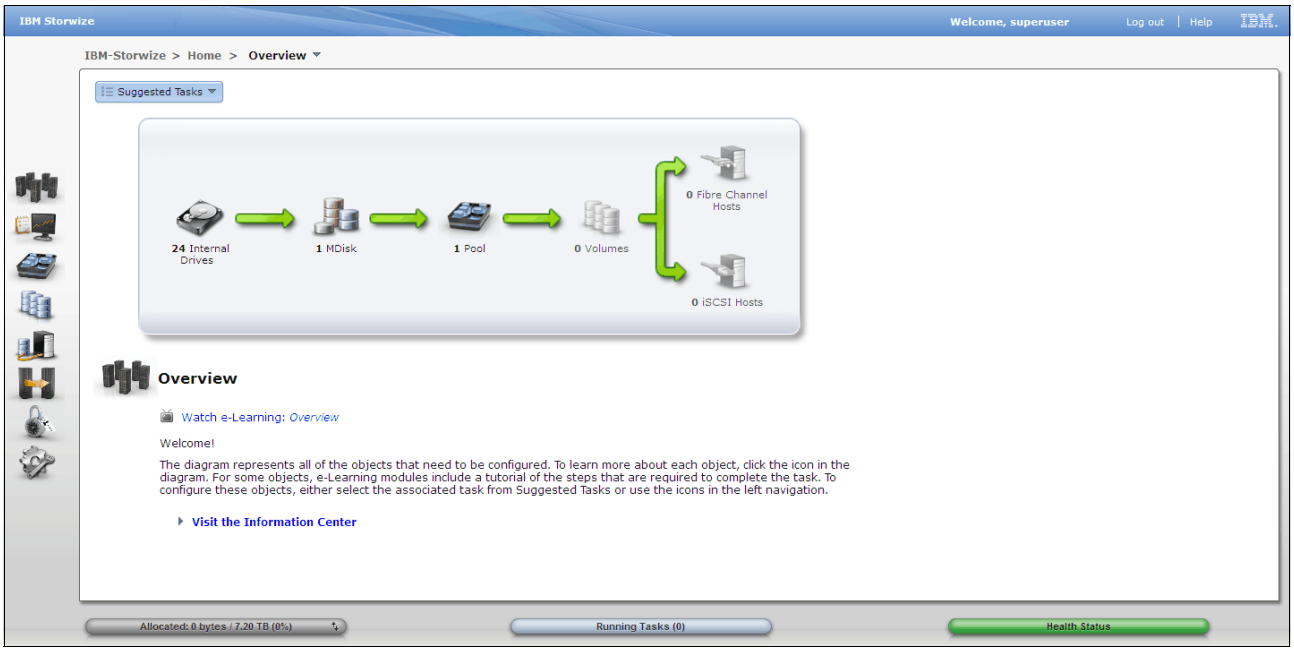

Figure 3-2 Storwize V3500 Overview panel

Low graphics mode

Figure 3-3 on page 57 shows the Storwize V3500 login panel and the check box to enable low graphics mode. This feature can be useful for remote access over narrow band links. The function icons no longer enlarge and list available functions. However, navigation is achieved 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 57.

For more information about the breadcrumb navigation aid, see 3.2.3, “Breadcrumb navigation aid” on page 62.

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

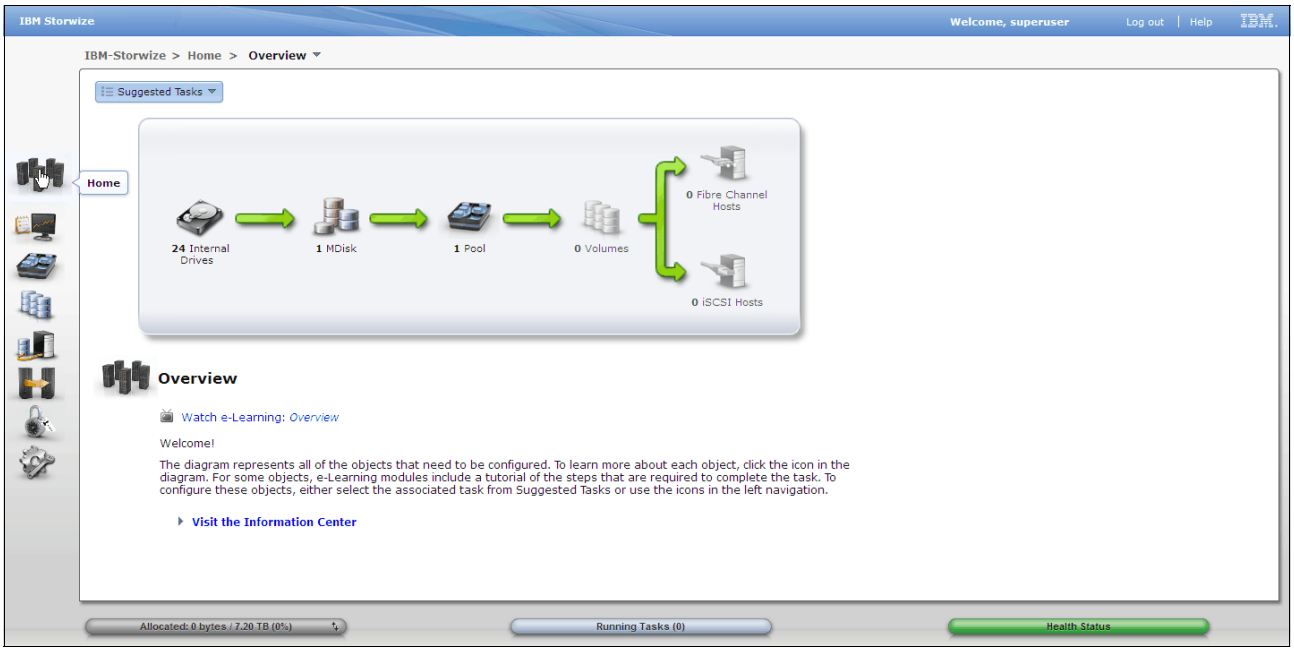

Figure 3-3 Storwize V3500 GUI low graphics mode

3.1.3 Overview panel layout

As shown in Figure 3-4, the Overview panel includes three main sections: Function Icons, Extended Help, and Status Indicators. Each section can be used to browse the management GUI.

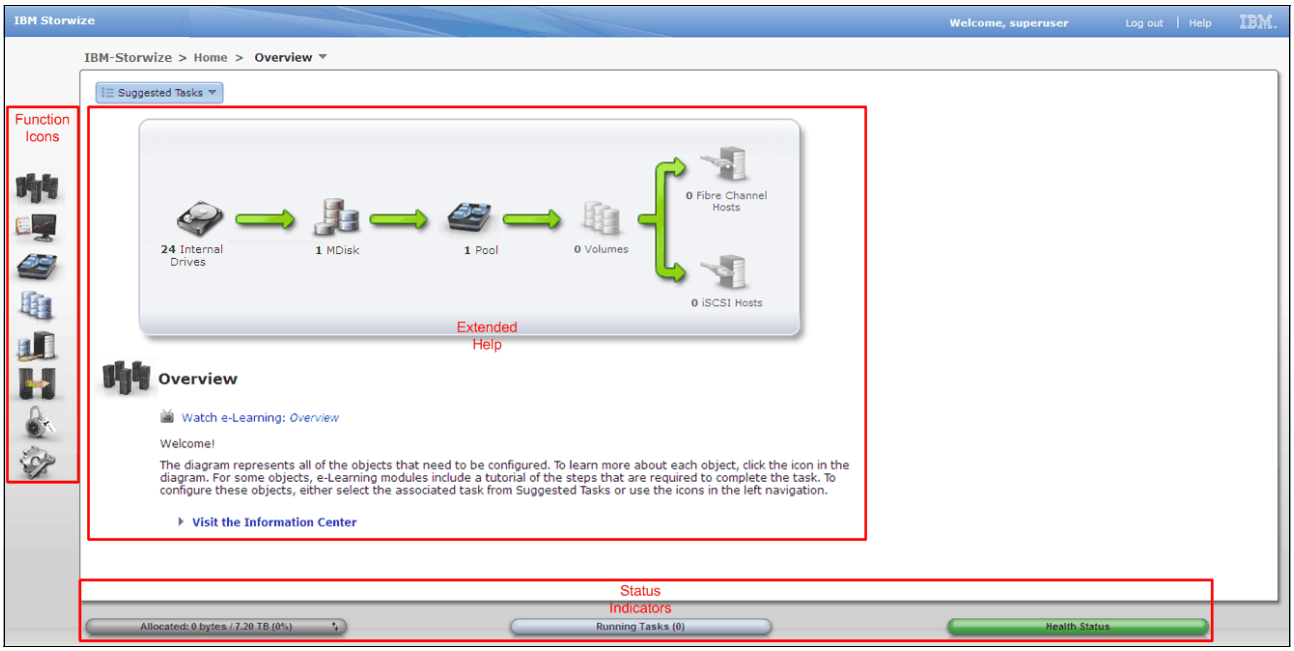

Figure 3-4 Main sections of the Storwize V3500 overview panel

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

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

The Extended Help section includes 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 in the flow diagram show the order in which to perform storage allocation tasks. It also highlights 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 currently running tasks 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 72.

3.2 Navigation

Navigating the management tool is simple. As with 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, which are described in 3.1.3, “Overview panel layout” on page 57.

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 pointer over one of the eight function icons on the left side of the panel enlarges the icon and provide a menu to access various functions. Move the pointer to the wanted function and click the function. Figure 3-5 shows the results of hovering the pointer over a function icon.

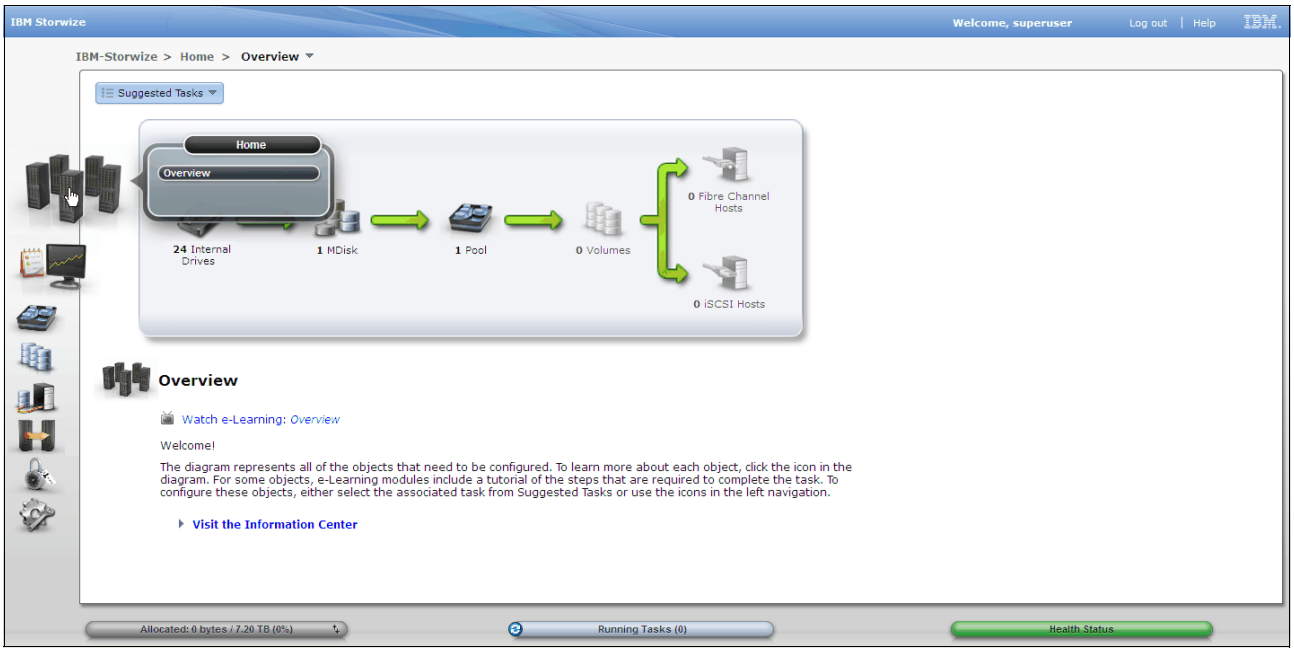

Figure 3-5 Hovering the pointer over a function icon

Figure 3-6 shows all of the menus under the function icons.

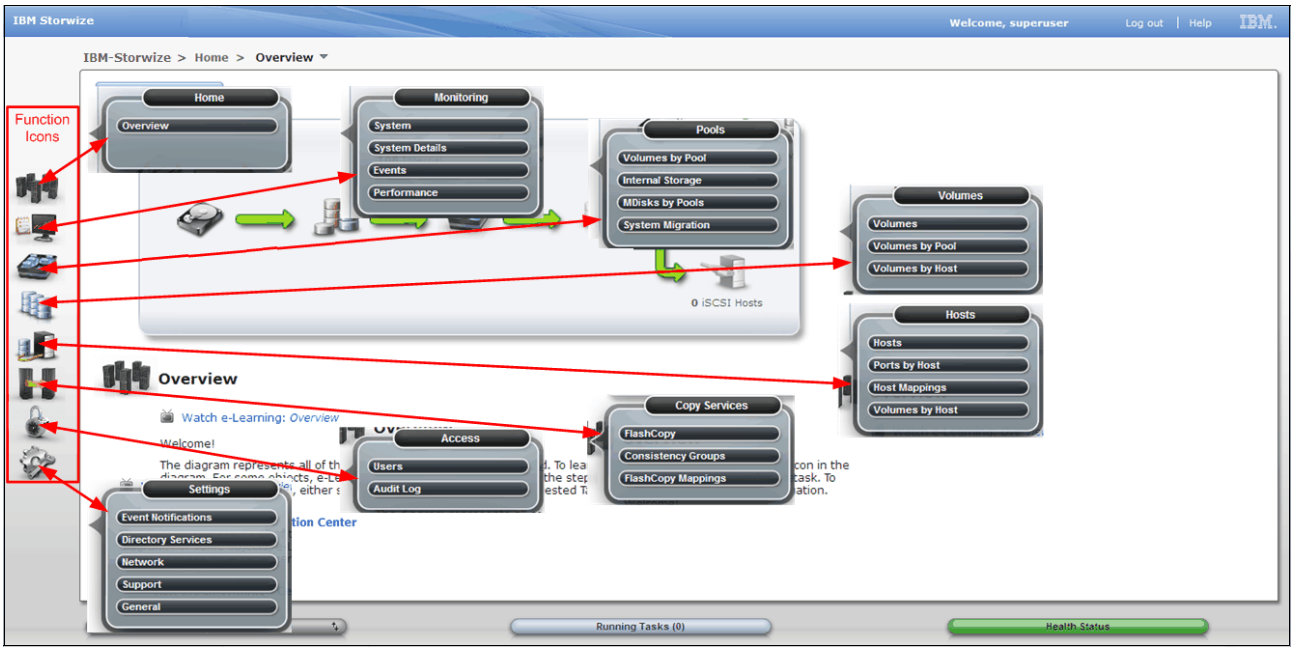

Figure 3-6 Storwize V3500 all options that are listed under function icons

3.2.2 Extended help navigation

Selecting an image in the flow diagram of the Extended Help section of the Overview panel shows more 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 selecting each object in order. Figure 3-7 shows the selection of Internal Drives in the flow diagram. The information that is below the flow diagram relates to the internal storage.

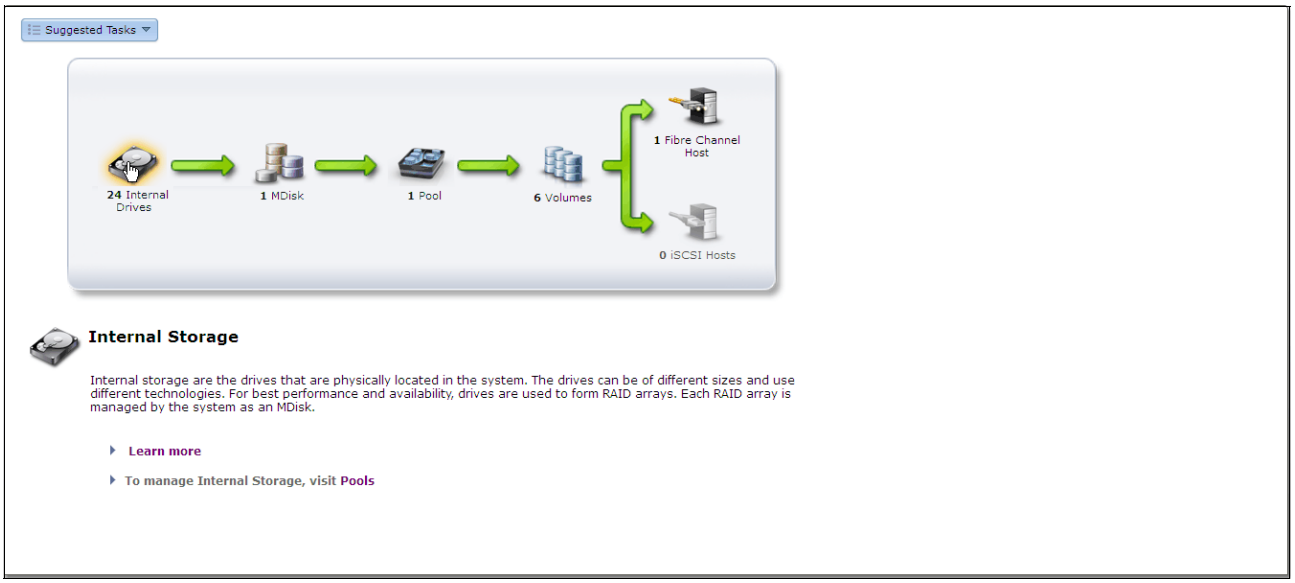

Figure 3-7 Navigating GUI with the extended help section

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

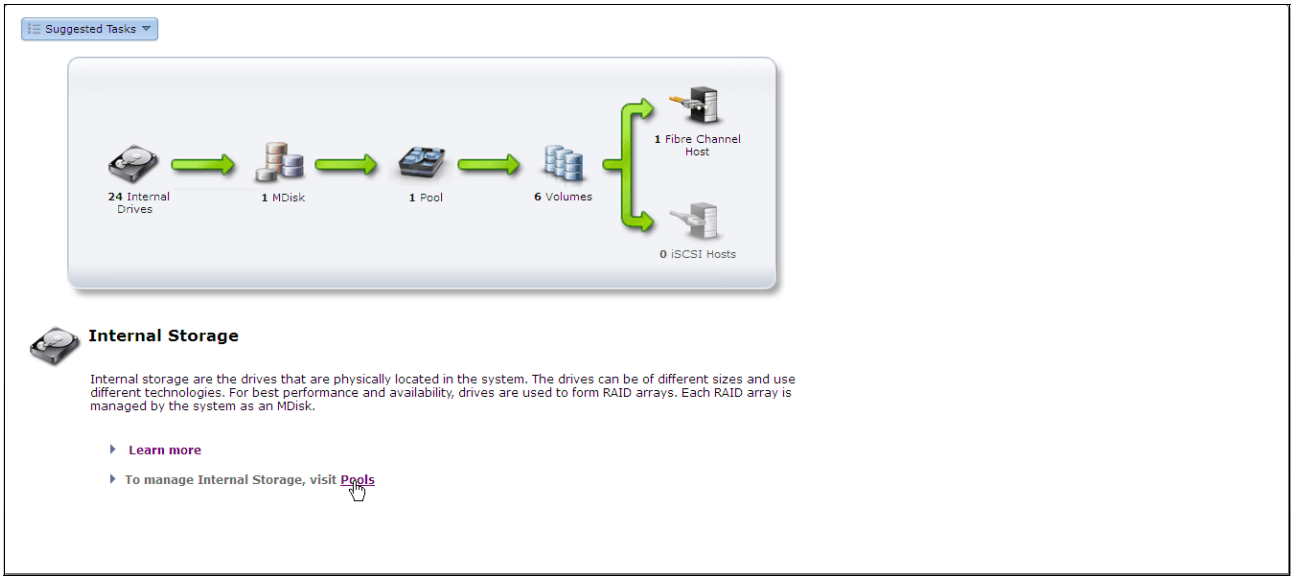

Figure 3-8 Using the extended help section

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

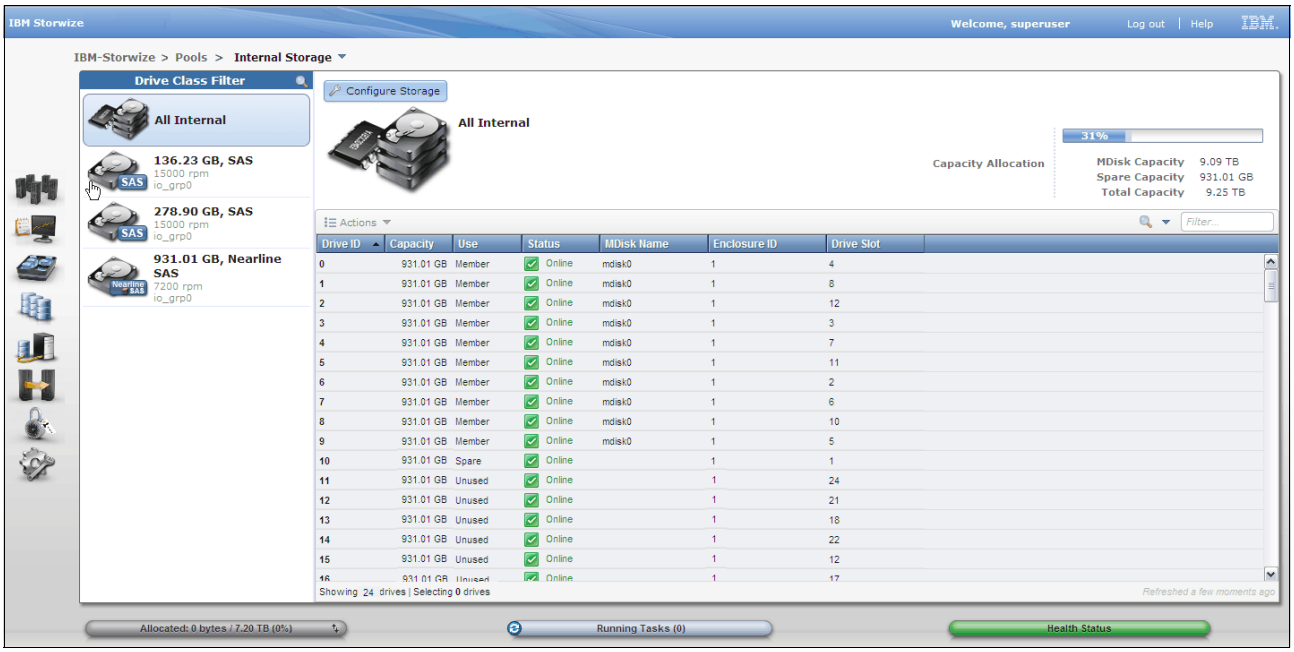

Figure 3-9 The internal storage configuration panel

3.2.3 Breadcrumb navigation aid

The Storwize V3500 panels use the breadcrumb navigation aid to show the trail that was browsed. This breadcrumb navigation aid is at the top area of the panel and includes a drop-down selection menu on the last breadcrumb. Figure 3-10 shows the breadcrumb navigation aid for the System panel.

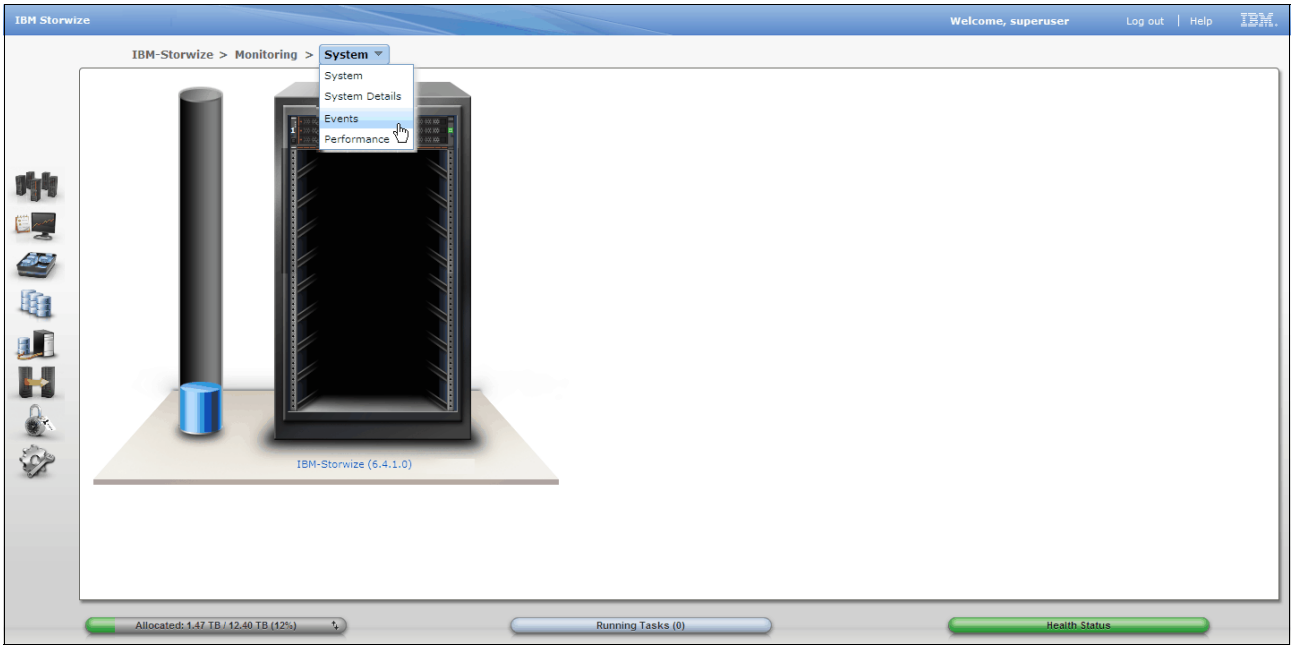

Figure 3-10 Breadcrumb navigation aid

3.2.4 Suggested tasks aid

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.

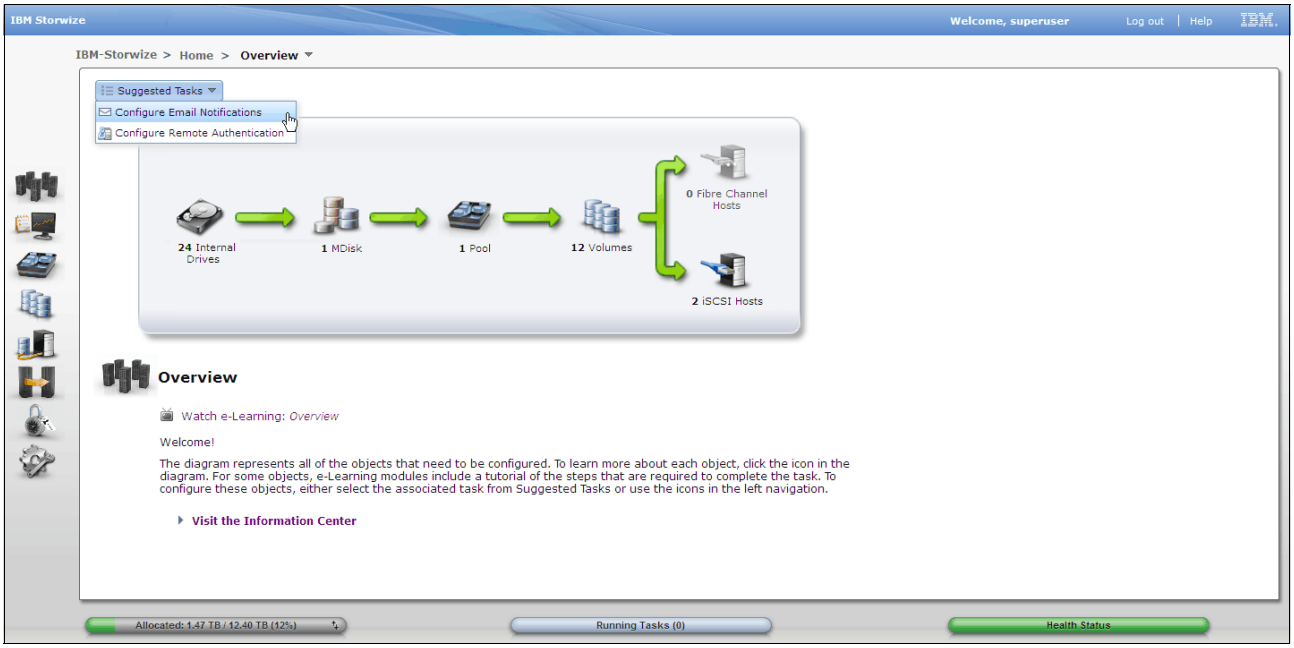

Figure 3-11 Suggested tasks navigation and configuration aid

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. They also are available for setting up a RAID configuration. Figure 3-12 shows the available internal storage presets.

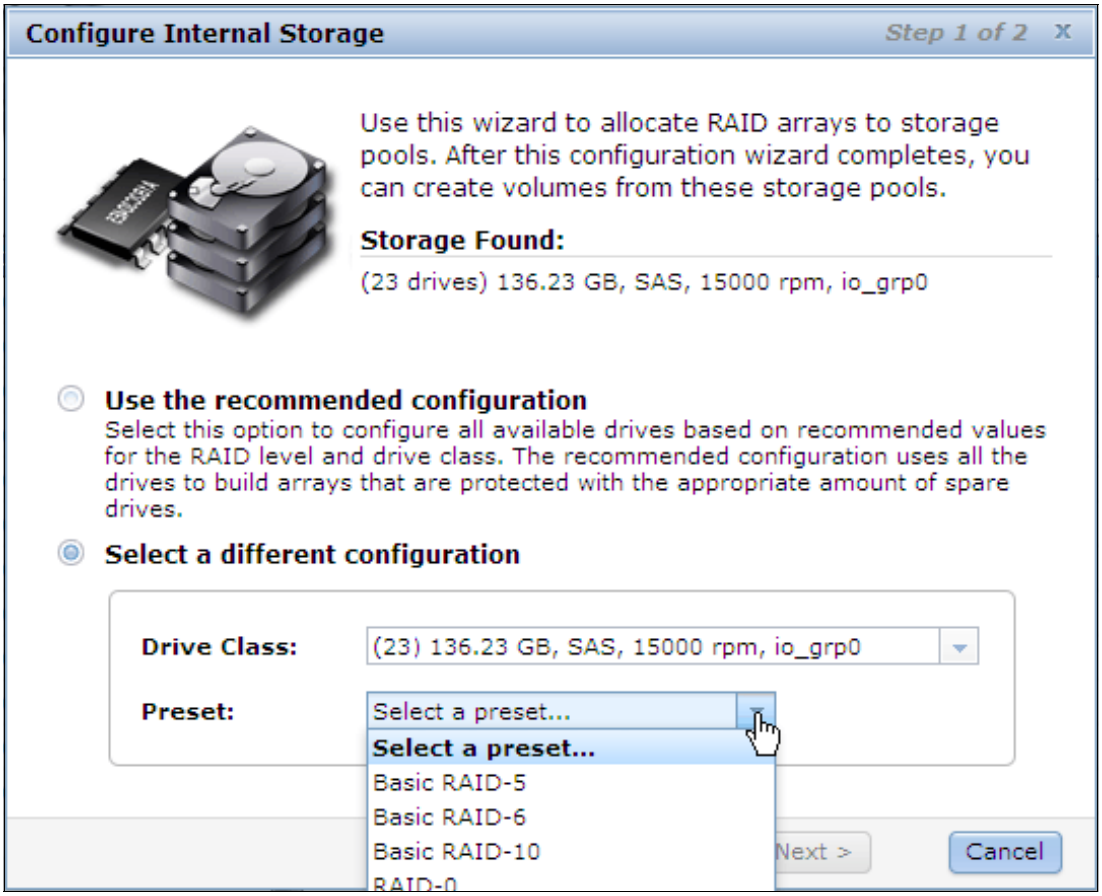

Figure 3-12 Internal storage preset selection

3.2.6 Access actions

The Storwize V3500 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 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.

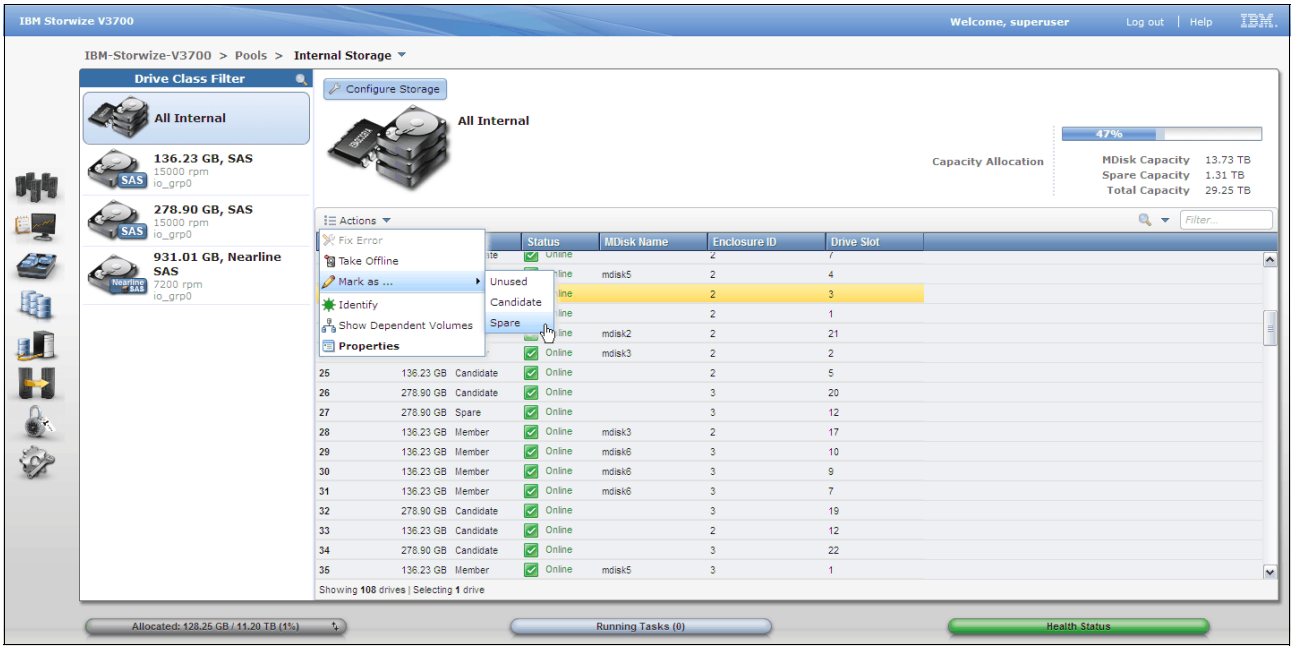

Figure 3-13 Actions drop-down menu

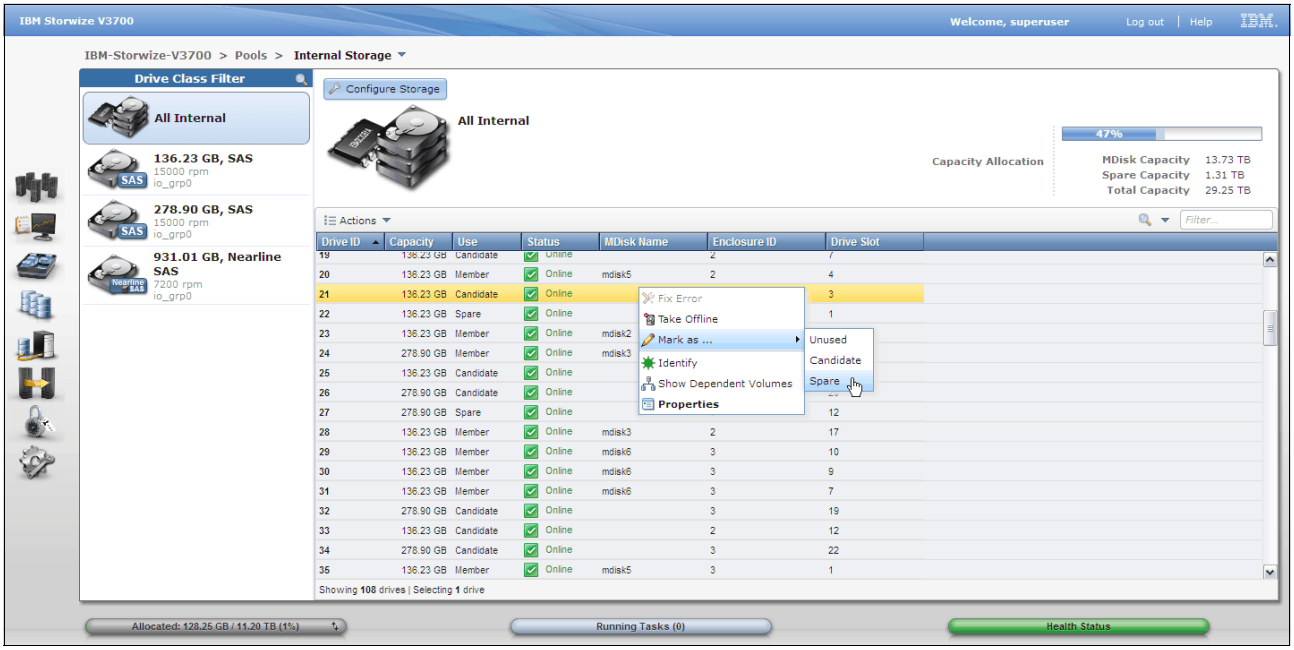

Figure 3-14 Right-click actions menu

3.2.7 Task progress

An action starts running tasks and shows a task progress box, 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 cut and pasted into a configured Storwize V3500 SSH terminal session, if required.

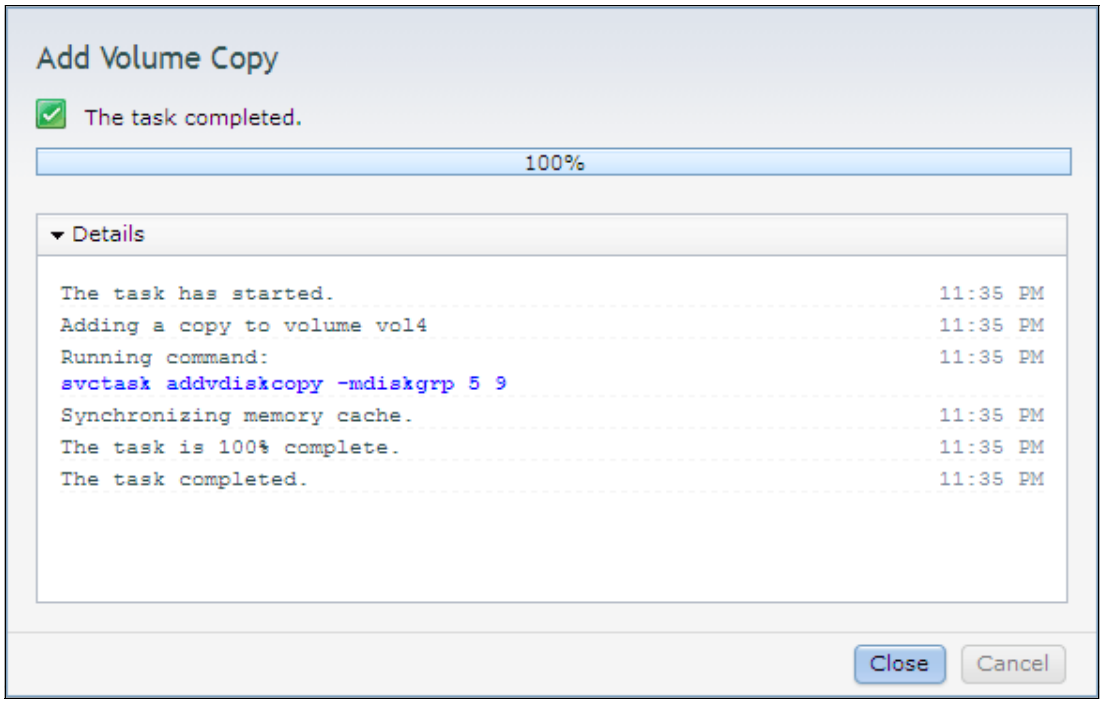

Figure 3-15 Task progress box

3.2.8 Navigating panels with tables

Many of the configuration and status panels show information in a table format with rows and columns. This section describes the following useful methods to browse 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 shows the result of clicking the heading of the Capacity Utilization column. The table is now sorted and lists volumes with the least amount of capacity utilization at the top of the table.

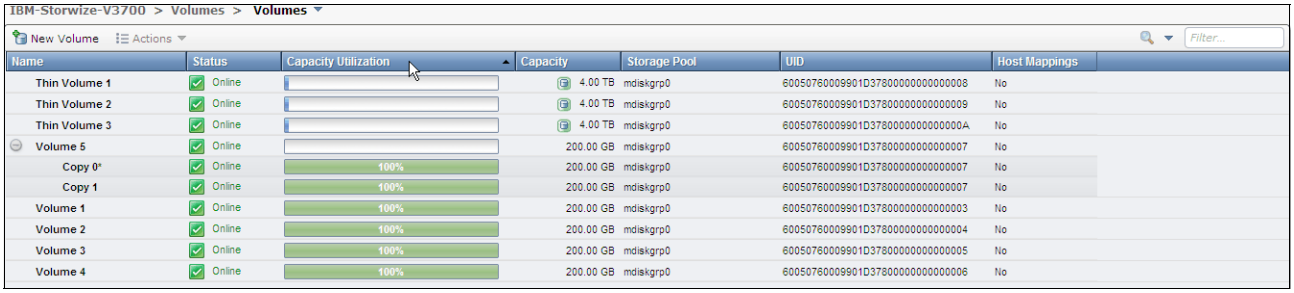

Figure 3-16 Sorting columns by clicking the column heading

Reordering columns

Columns can be reordered by clicking and dragging the column to the wanted location. Figure 3-17 shows the current location of the column with the heading UID positioned between the headings Storage Pool and Host Mappings. Clicking and dragging this heading reorders the columns in the table.

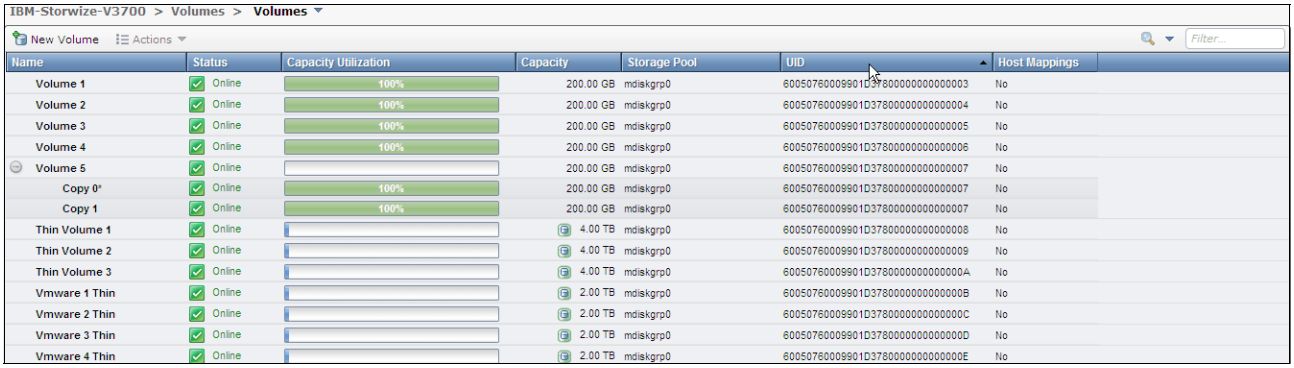

Figure 3-17 Reordering columns part 1

Figure 3-18 shows the column heading UID was dragged to the wanted location.

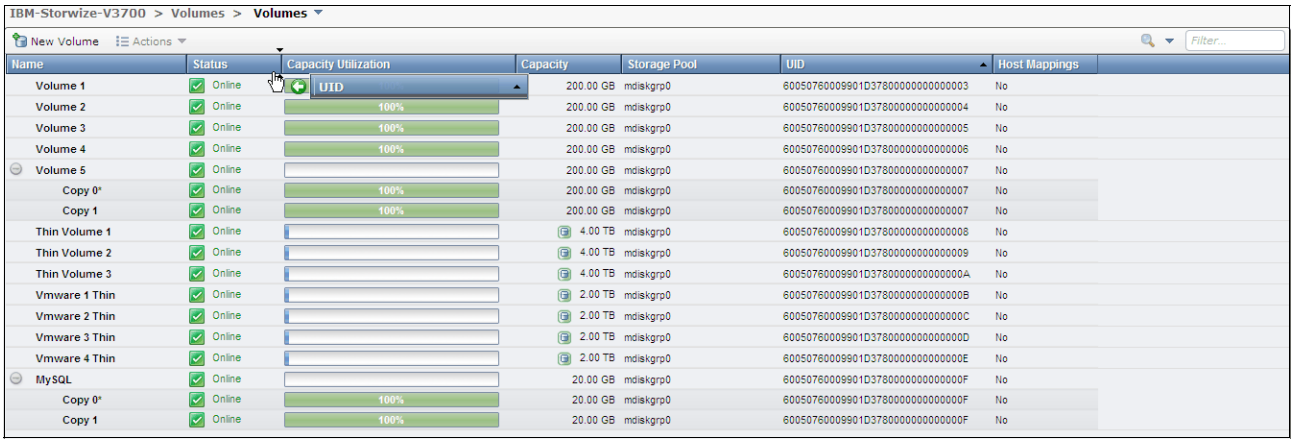

Figure 3-18 Dragging a column heading to the wanted location

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

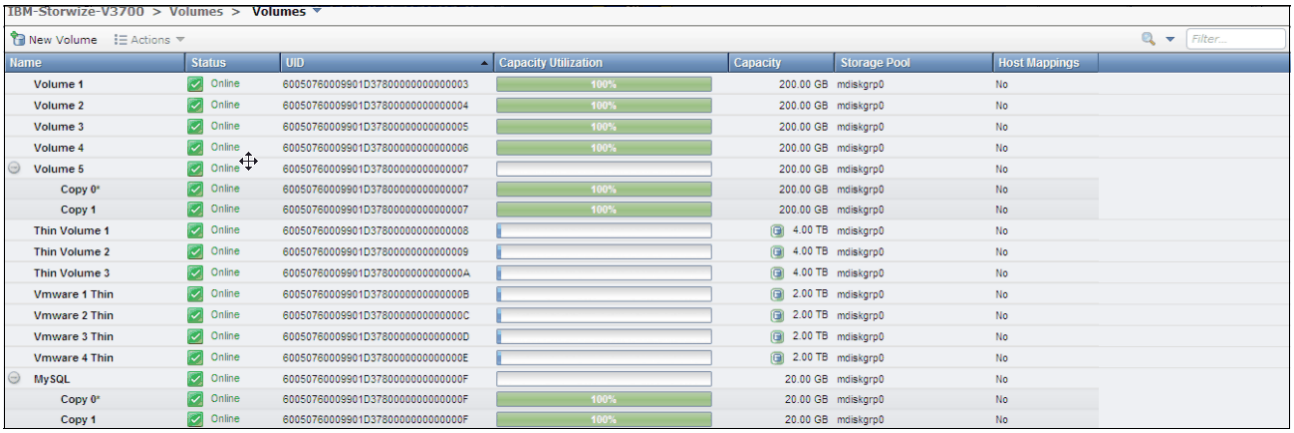

Figure 3-19 Reordering column headings

Adding or removing columns

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

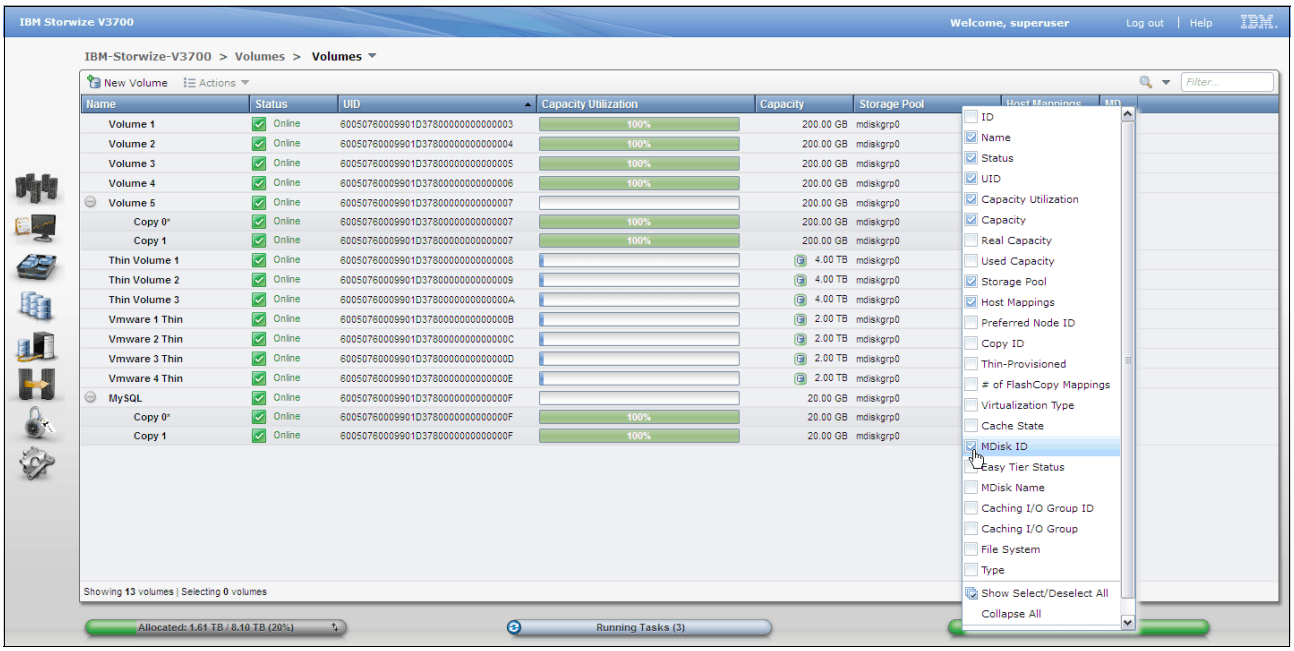

Figure 3-20 Adding column heading MDisk ID

Important: Some users might run into a problem where a context menu from the Firefox browser is shown by right-clicking change to 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 Storwize V3500 management GUI can be found in the Storwize V3500 Information Center at this website:

http://pic.dhe.ibm.com/infocenter/storwize/v3500_ic/index.jsp

Multiple selections

By using the management tool, you 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 the items in between those two items are selected, as shown in Figure 3-21.

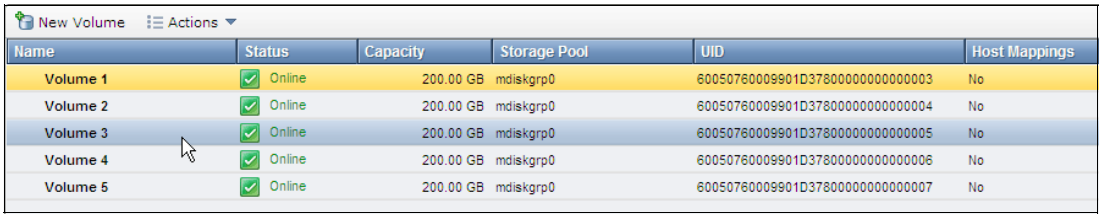

Figure 3-21 Selection of three sequential items: Part 1

Figure 3-22 shows the result of the use of the Shift key to select multiple sequential items.

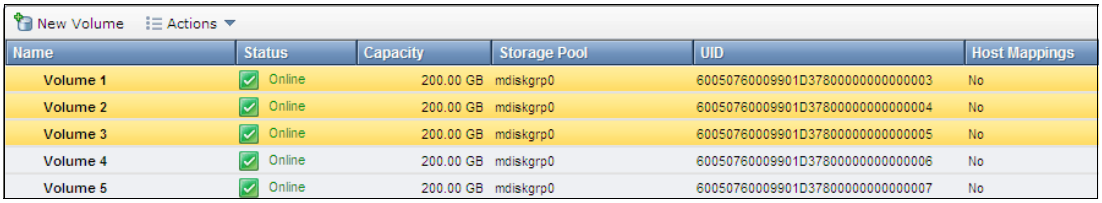

Figure 3-22 Selection of three sequential items: Part 2

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 click the other items that you require. Figure 3-23 shows the selection of two non-sequential items.

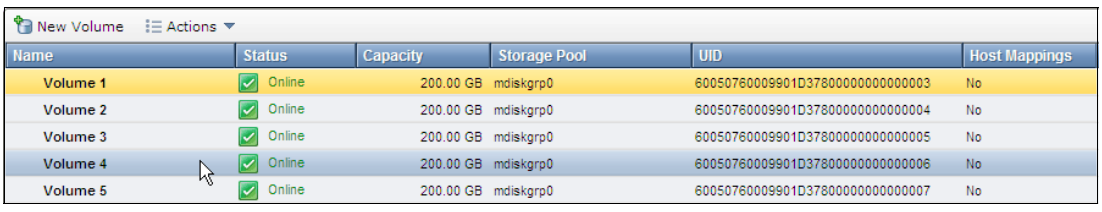

Figure 3-23 Selection of two non-sequential items: Part 1

Figure 3-24 shows the result of the use of the Ctrl key to select multiple non-sequential items.

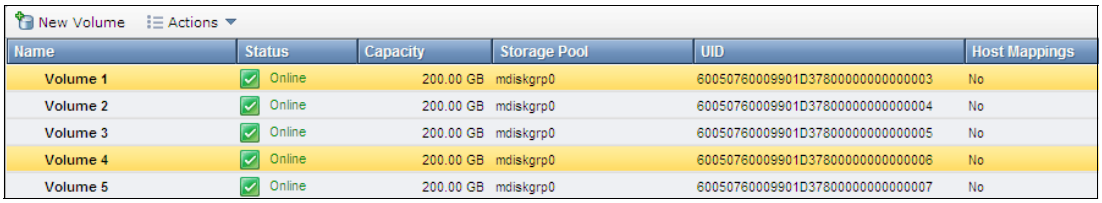

Figure 3-24 Selection of two non-sequential items: Part 2

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 3-25 shows the text `Thin` was entered into the filter field. Now, only volumes with the text `Thin` in any column are listed and the filter word also is highlighted.

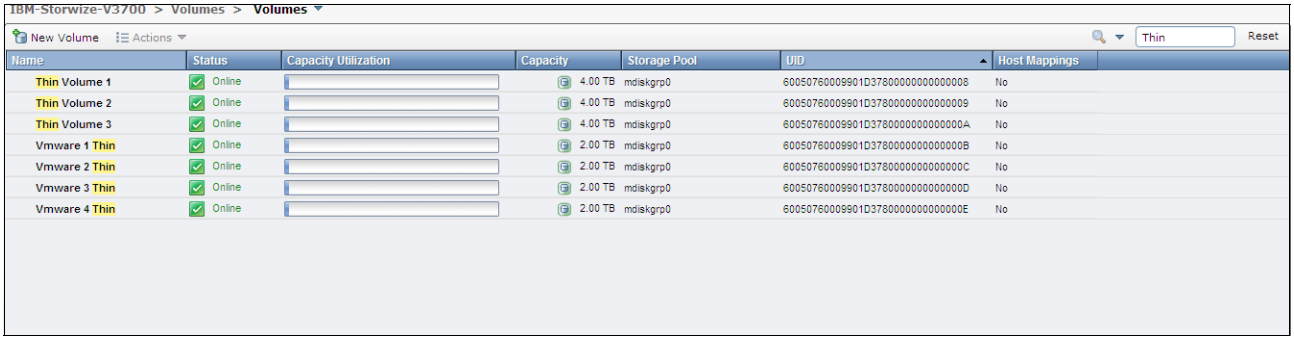

Figure 3-25 Filtering objects to display a subset of the volumes

Filtering by column

Click the magnifying glass that is next to the filter field to activate the filter by column feature. This feature allows the filter field value to be matched to the values of a specific column. Figure 3-26 shows the column filter is set to Host Mappings and the filter value can be selected.

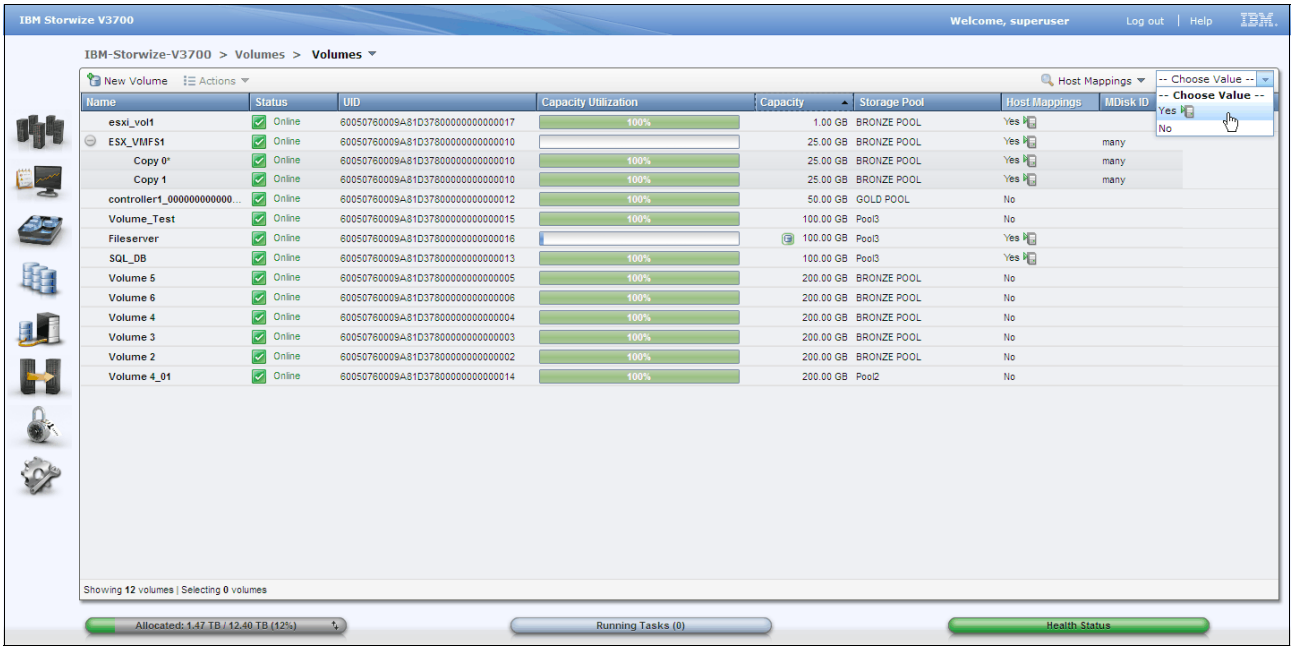

Figure 3-26 Filter by column

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 57, 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 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 shows the comparison of the used capacity to the real capacity.

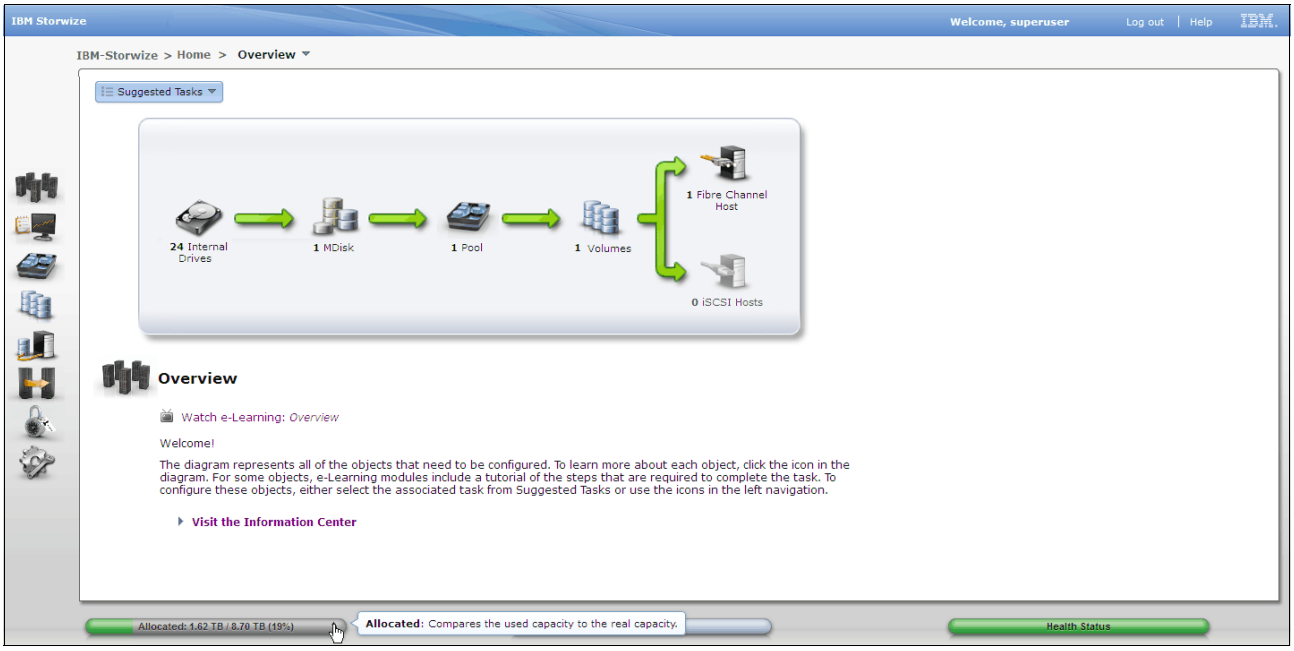

Figure 3-27 Allocated bar, comparing used capacity to real capacity

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

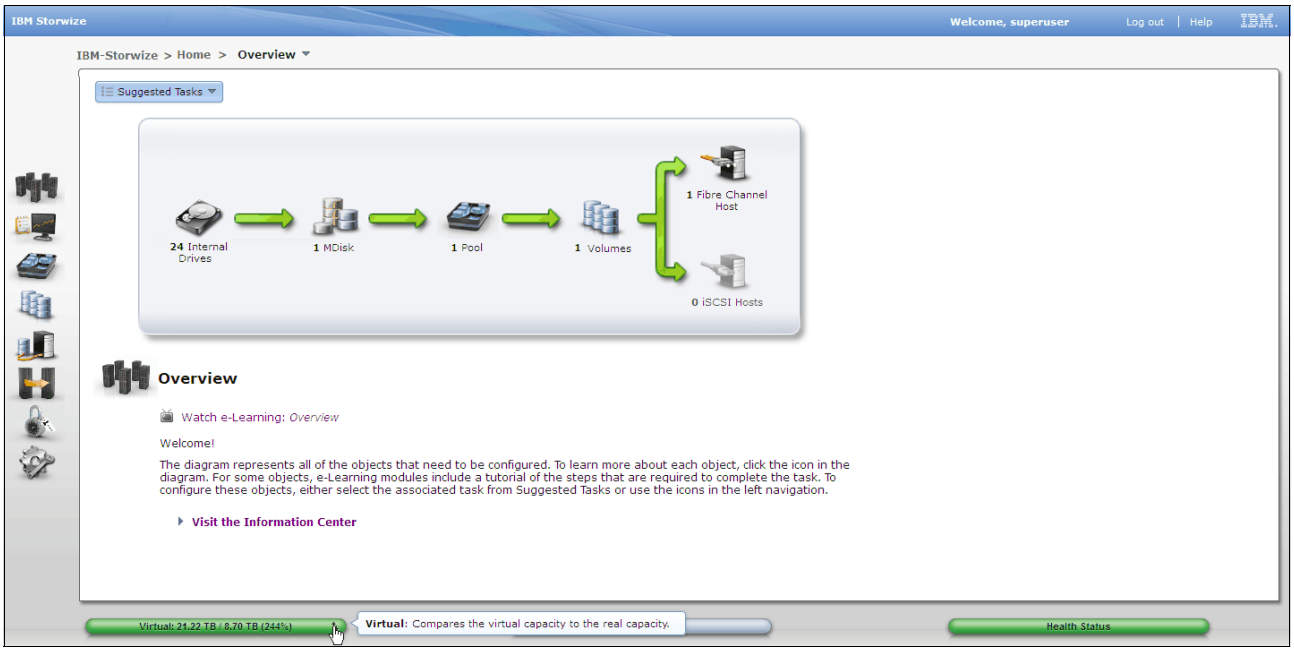

Figure 3-28 Changing the allocated menu comparison, virtual capacity to real capacity

3.3.3 Running tasks bar menu

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

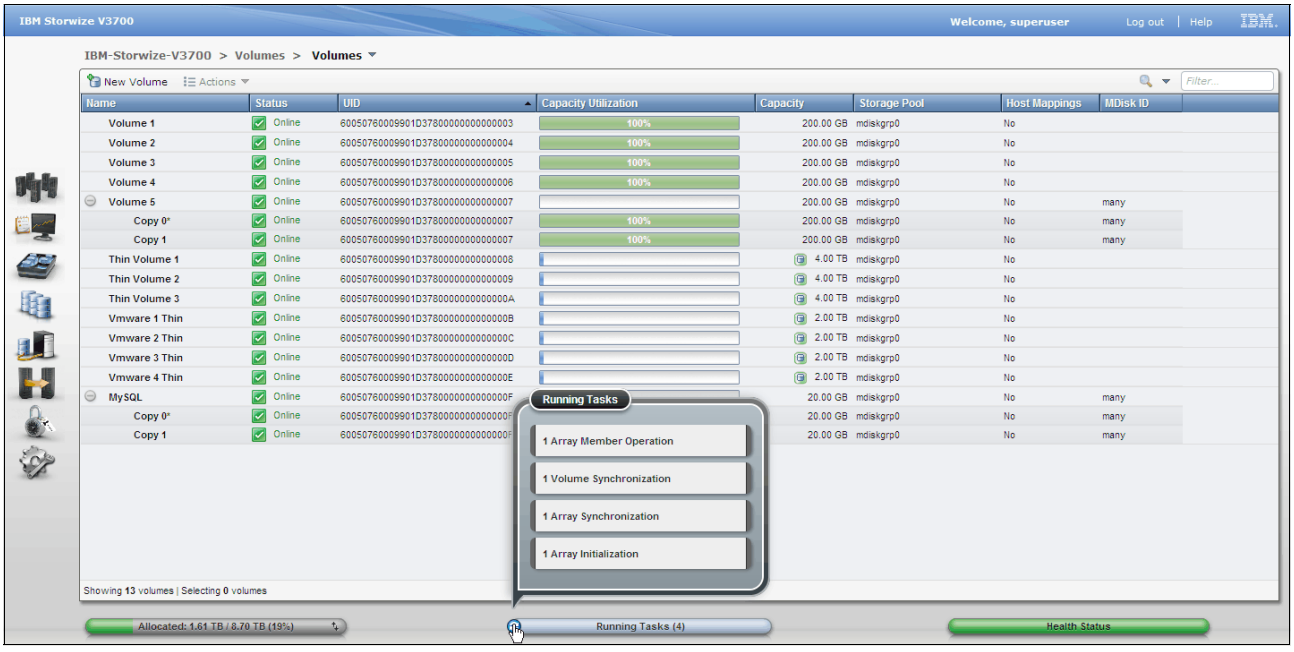

Figure 3-29 Running Tasks menu

For an indication of task progress, browse to the Running Tasks bar menu and click the task, as shown in Figure 3-30.

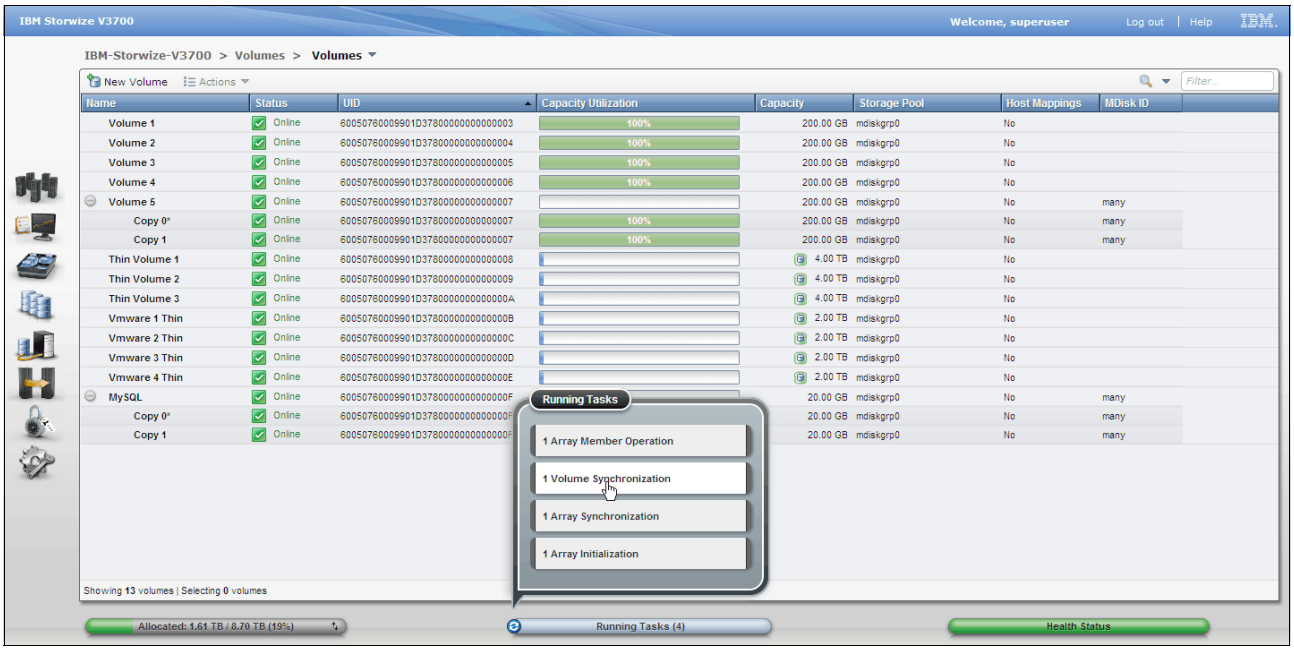

Figure 3-30 Selecting a task from the running task menu for and indication of task progress

Figure 3-31 shows the progress of all running Volume Synchronization tasks.

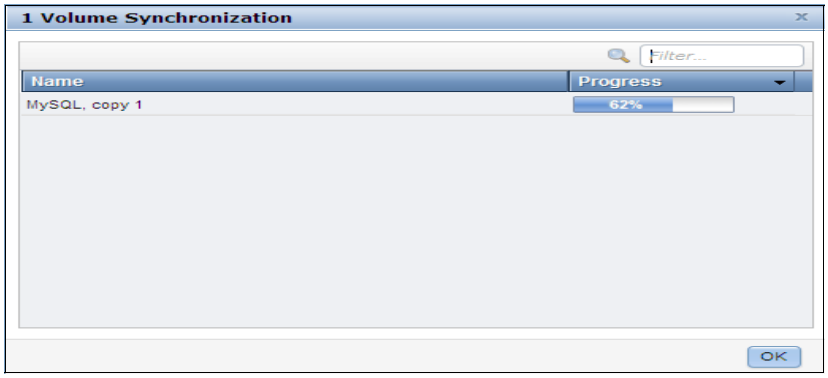

Figure 3-31 Progress of all running Volume Synchronization tasks

3.3.4 Health status bar menu

The health status bar indicates the overall health of the system. The following color of the status bar indicates the state of Storwize V3500:

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

If a status alert occurs, the health status bar can turn yellow or red. To show the health status menu, click the triangular attention image on the left side of the health status bar. Figure 3-32 shows the health status menu.

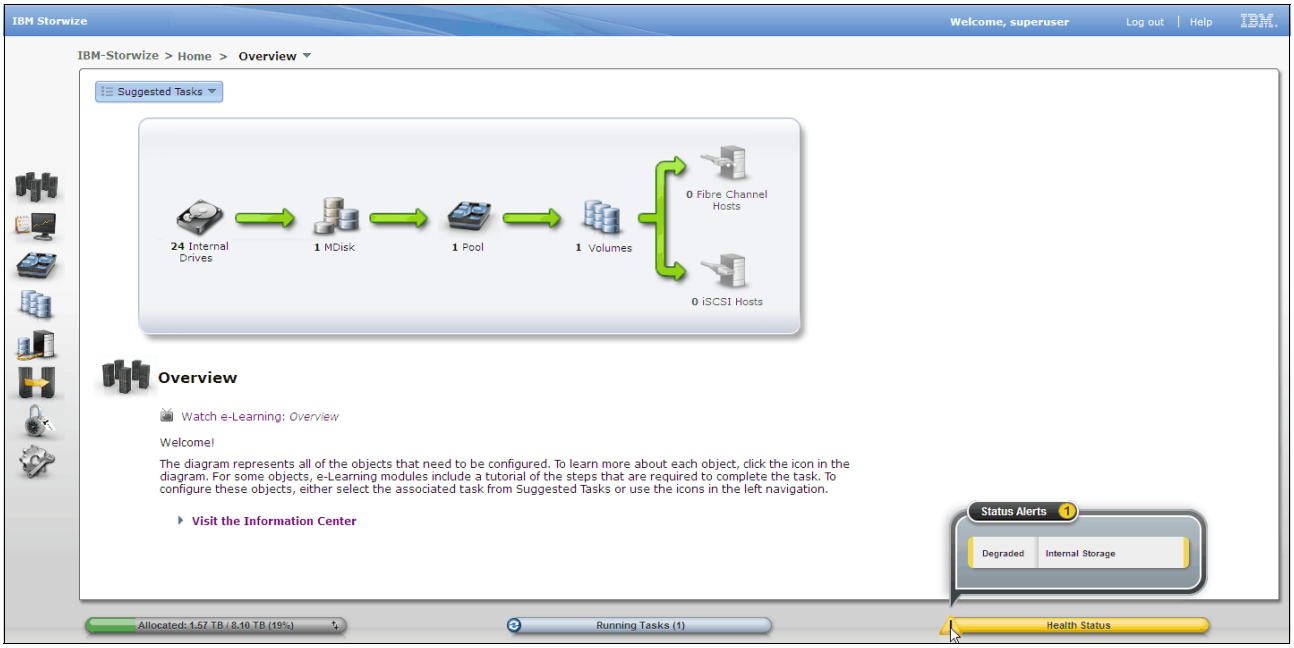

Figure 3-32 Health status menu alert

In Figure 3-33, the health status bar menu reports the system as degraded 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.

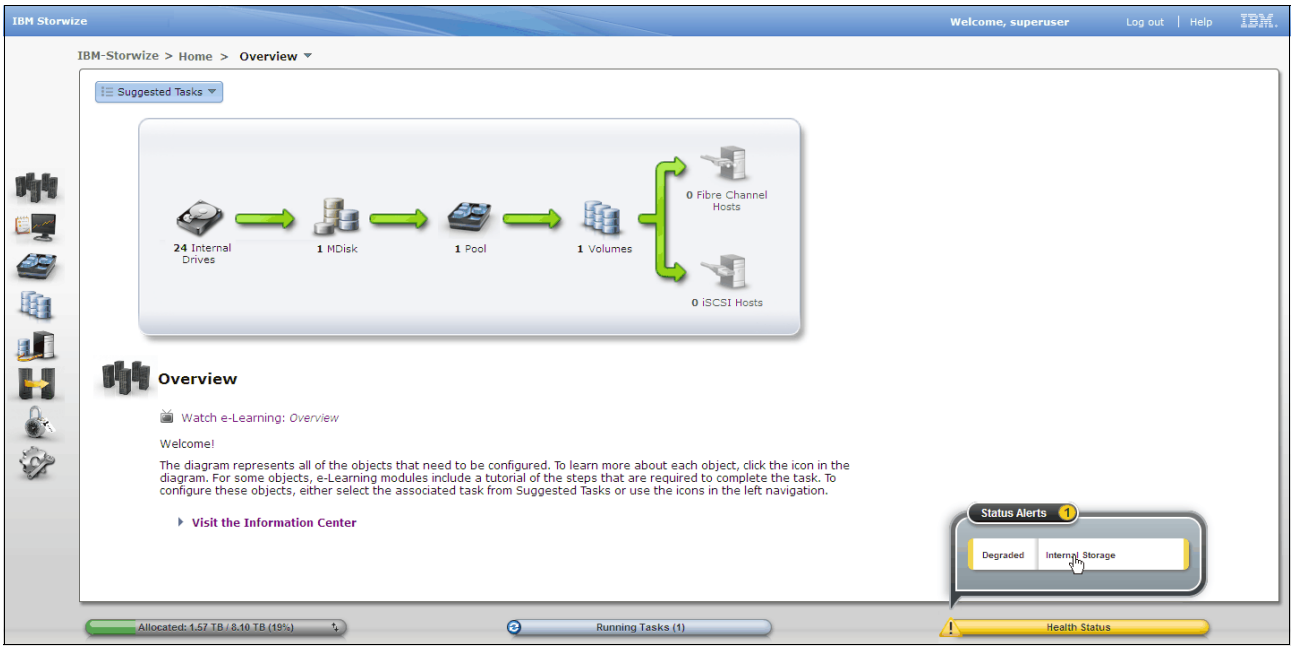

Figure 3-33 The status and description of an alert via the health status menu

Click the description of the event in the health status menu shows the events panel (**Monitoring** → **Events)**, as shown in Figure 3-34. This panel lists all of the events and provides directed maintenance procedures to help resolve errors. For more information, see “Events panel” on page 83.

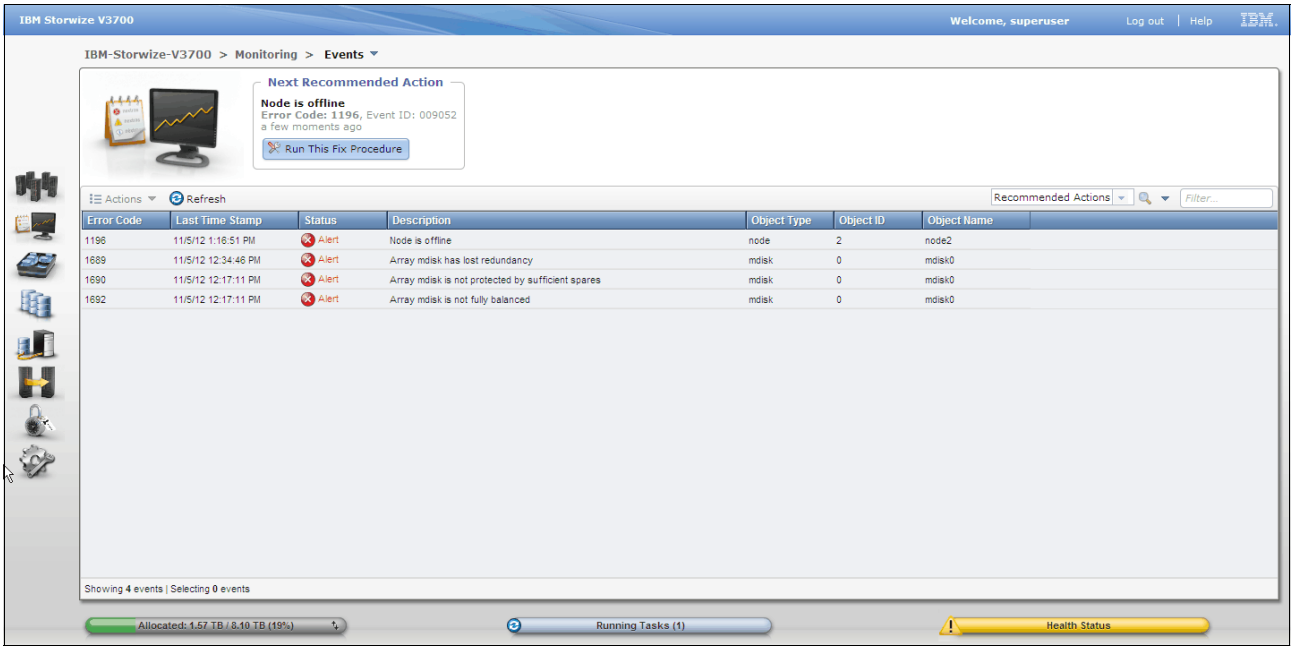

Figure 3-34 Events panel via health status menu

3.4 Function icon menus

The Storwize V3500 management GUI provides function icons that are an efficient and quick mechanism for navigation. As described in section 3.1.3, “Overview panel layout” on page 57, each graphic on the left side of the panel is a function icon that represents 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.

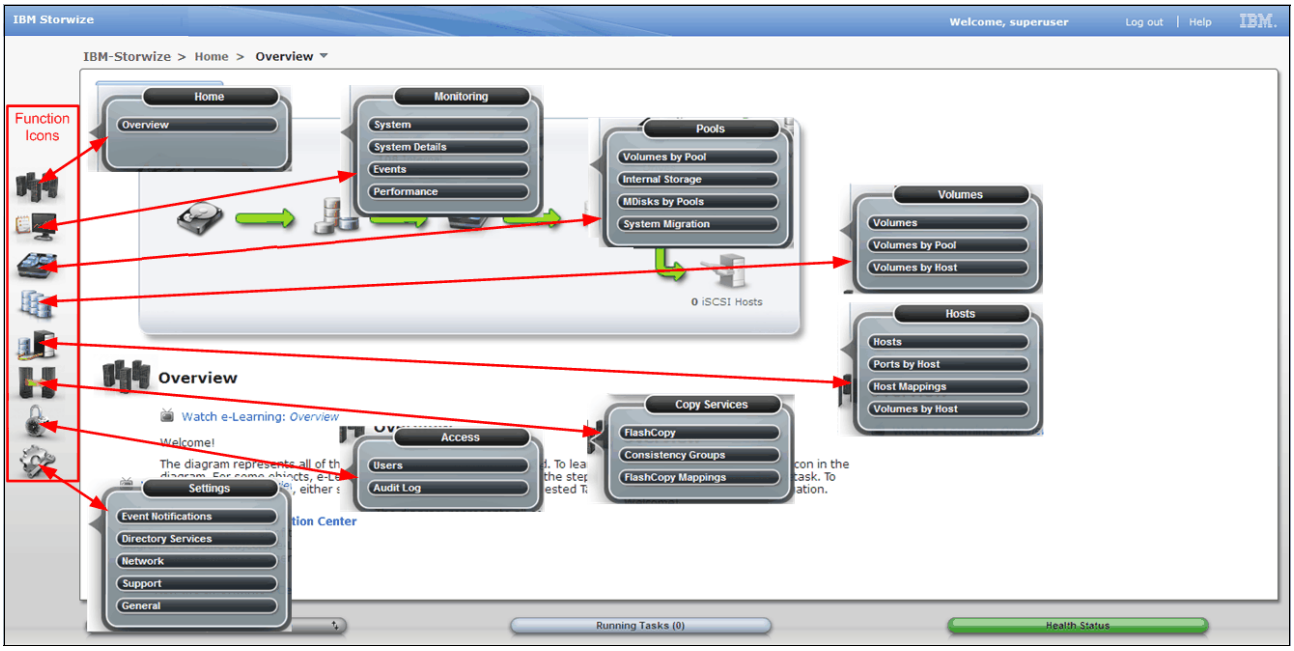

Figure 3-35 All function icon menus

3.4.1 Home menu

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

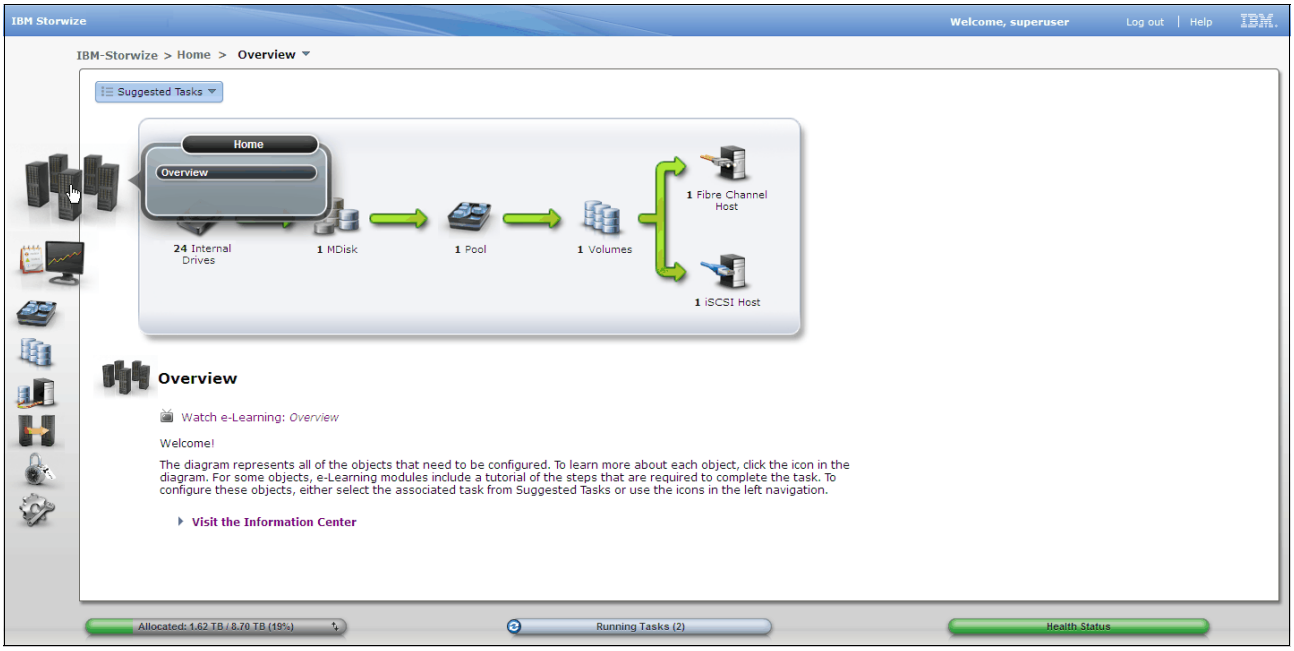

Figure 3-36 Home menu

Overview panel

Select **Overview** in the Home menu to open the panel. For more information about the Overview panel, see 3.1.3, “Overview panel layout” on page 57.

3.4.2 Monitoring menu

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

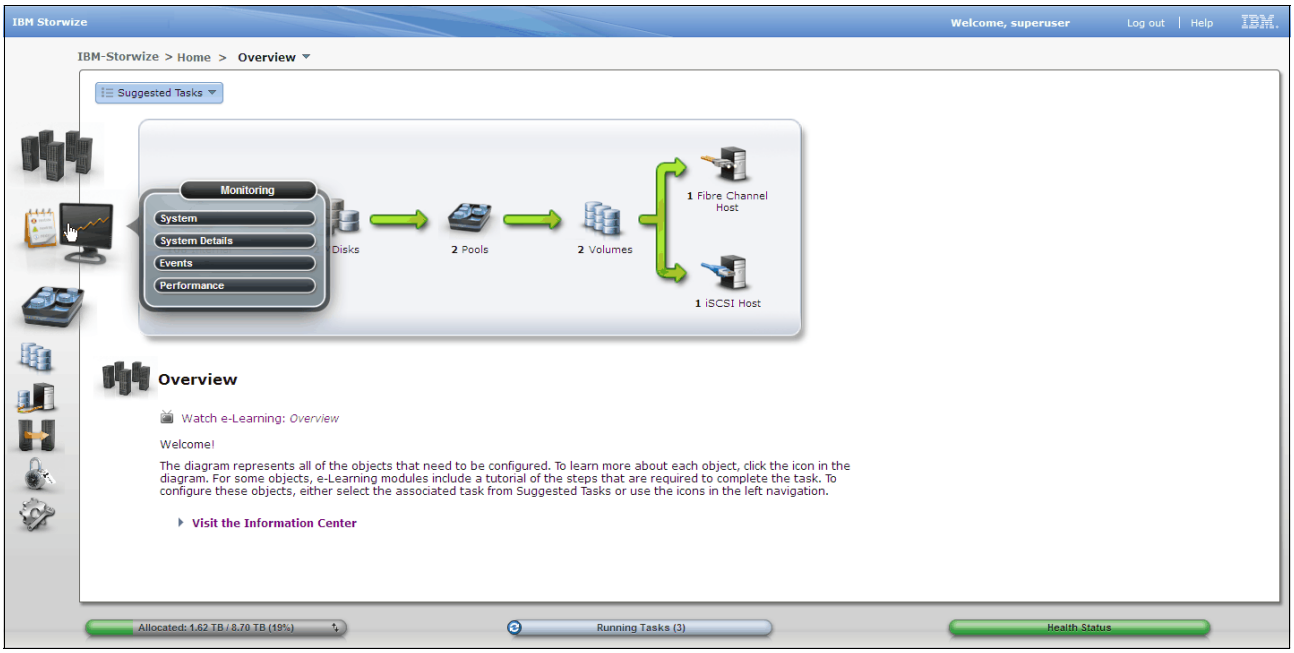

Figure 3-37 Monitoring menu

System panel

Select **System** in the Monitoring menu to open the panel. The System panel, as shown in Figure 3-38, shows capacity usage, the enclosure, and all drives in the system.

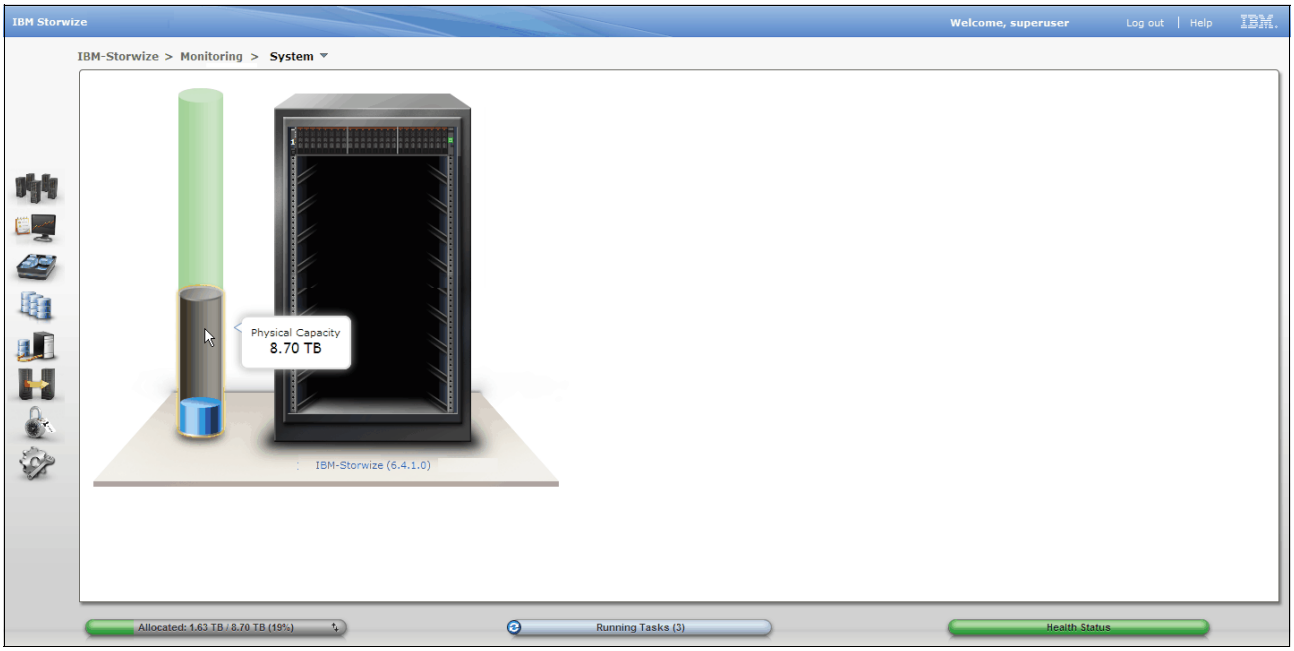

Figure 3-38 The System panel

Selecting the hotlink that includes the name and version of the system shows more information about storage allocation. Figure 3-39 shows the System panel Info tab.

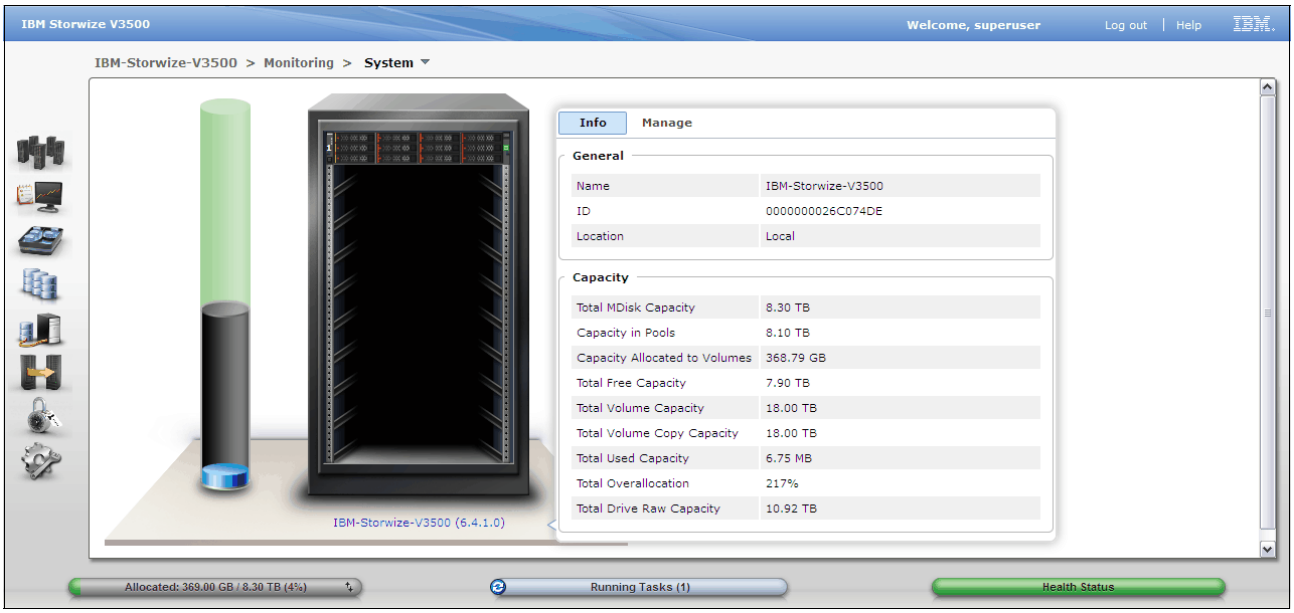

Figure 3-39 System panel Info tab

Select the **Manage** tab to show the name of the system and shutdown and upgrade actions. Figure 3-40 shows the System panel Manage tab.

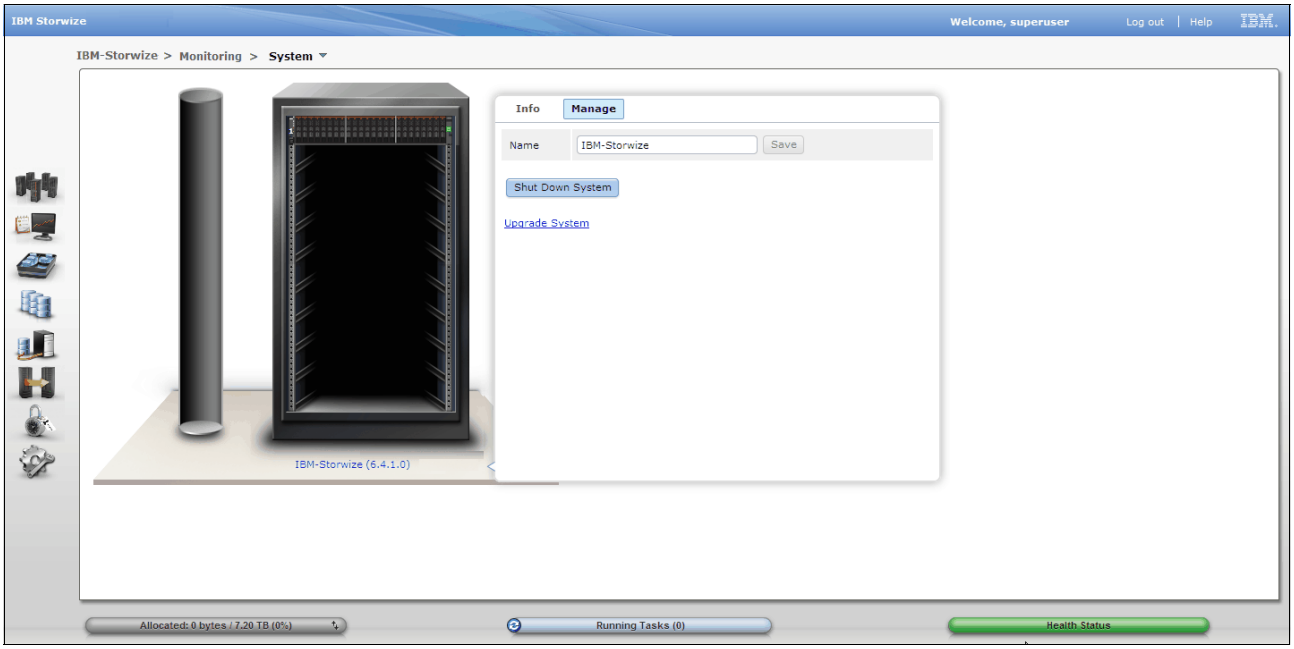

Figure 3-40 System panel Manage tab

Selecting the rack-mounted enclosure shows more information. Hovering over a drive provides the drive status, size, and speed details. Click **Identify** to activate the blue identification LED that is on the front of the enclosure. The Enclosure 1 hotlink shows the System Details panel. For more information, see “System details panel” on page 81. Figure 3-41 shows the System panel enclosure view.

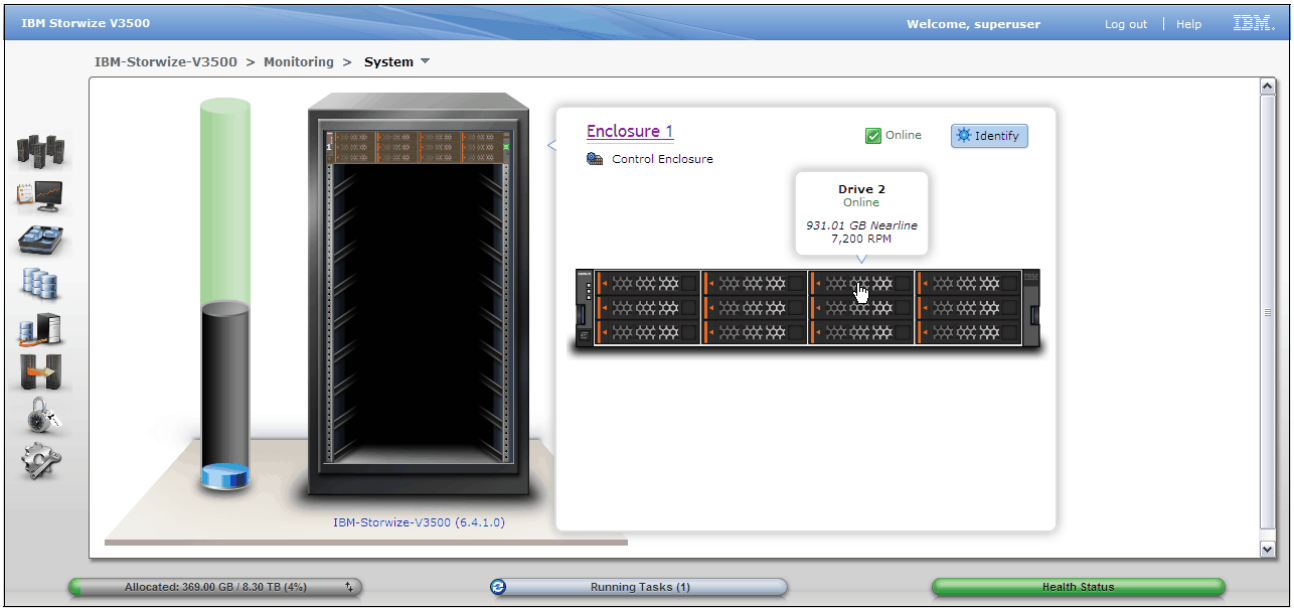

Figure 3-41 System panel enclosure view

System details panel

Select **System Details** in the Monitoring menu to access 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.

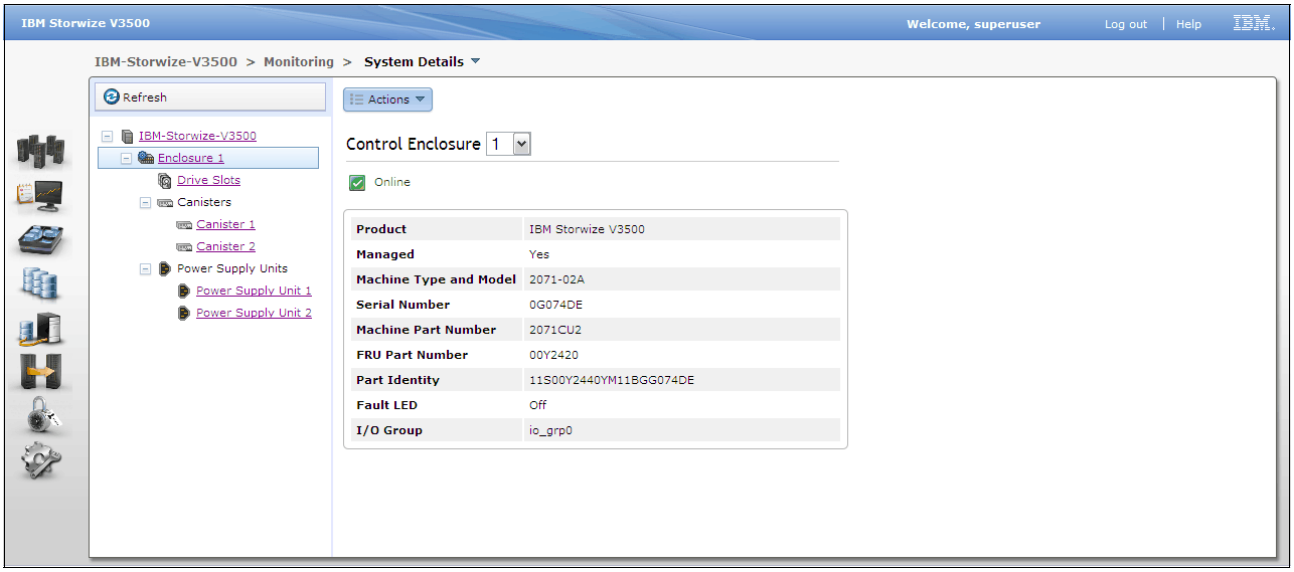

Figure 3-42 System Details panel

Actions and environmental statistics

Tasks such as performing a software upgrade and a system that is shut down can be run from the System Details panel. Information that relates to environmental statistics, such as power consumption and temperature, also is accessible from this panel. Figure 3-43 shows the environmental statistics of a system.

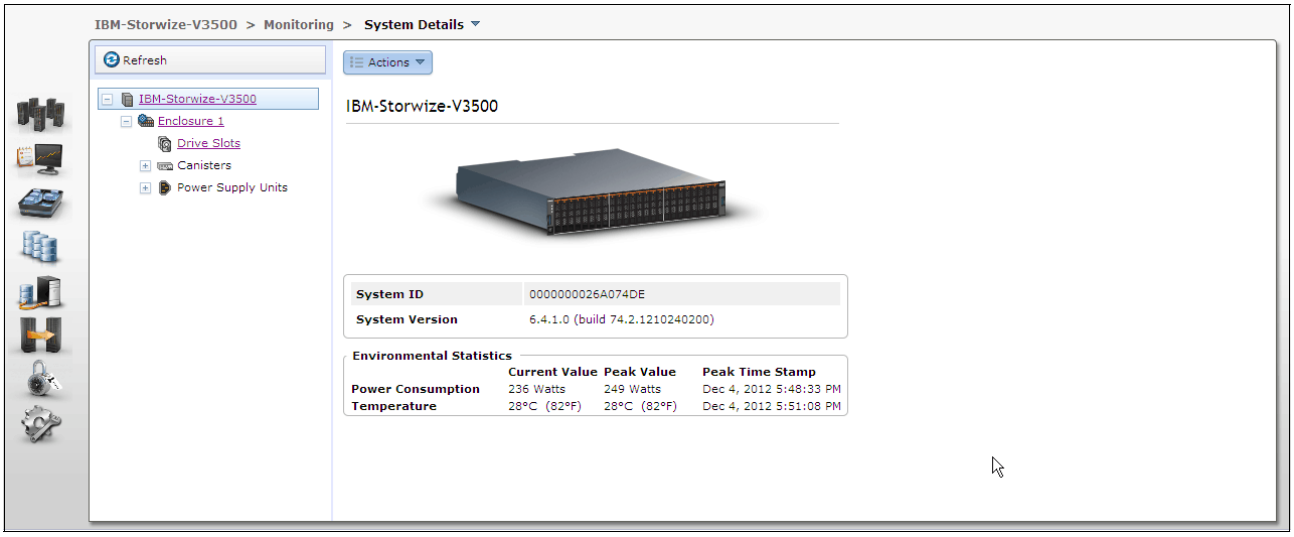

Figure 3-43 Storwize V3500 environmental statistics

Node canister information

Additionally, node canister information, such as FC WWPNs and ISCSI IQNs, is useful for host attachment purposes. Clicking the control enclosure node canister in the system details panel shows this information, as shown in Figure 3-44.

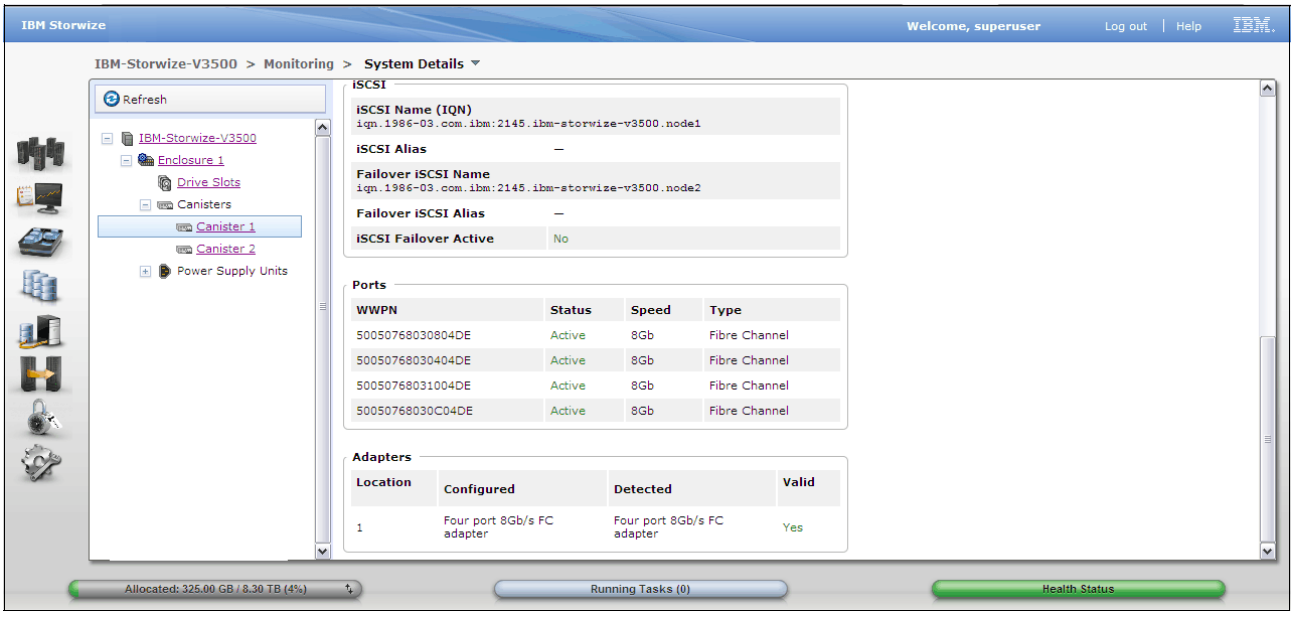

Figure 3-44 Node canister information via System Details panel

Events panel

Select **Events** in the Monitoring menu to open the panel. The machine is optimal when all errors are addressed and there are no items that are found in the Events panel. Figure 3-45 shows the events panel with no recommended actions.

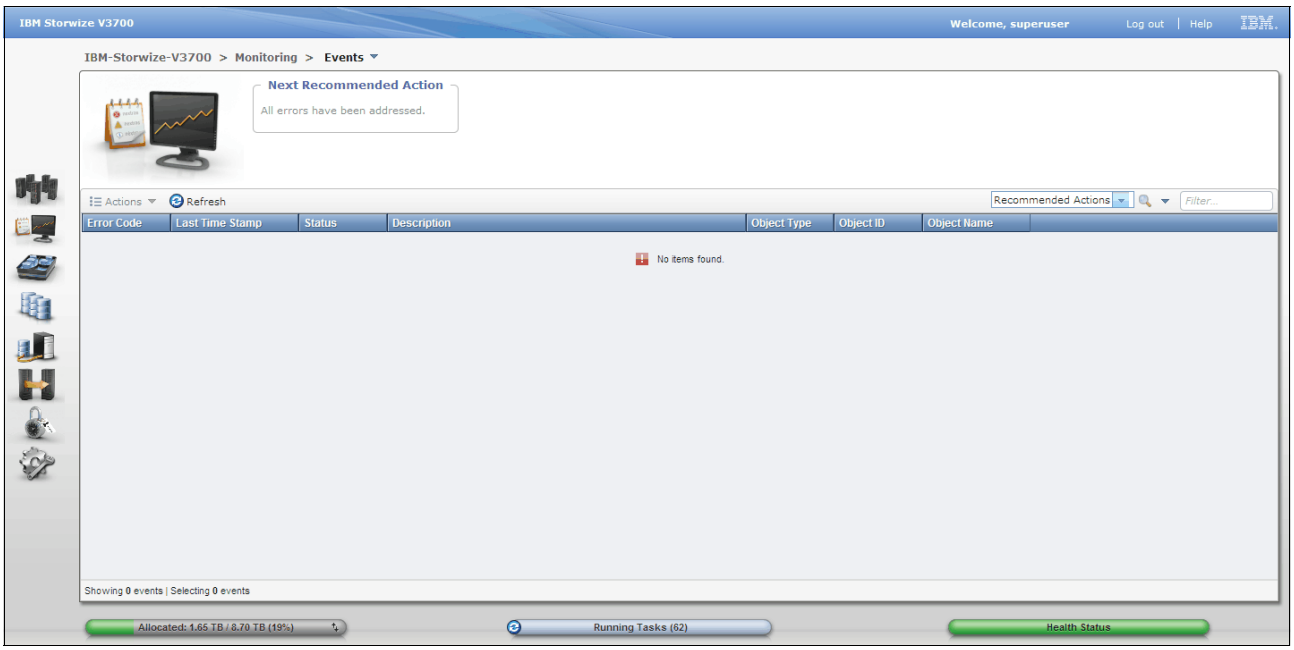

Figure 3-45 Events panel with all errors addressed

Filtering events view

To view unfixed messages and alerts, select the **Unfixed Messages and Alerts** option from the drop-down menu that is next to the filter field. Figure 3-46 shows unfixed messages and alerts. For more information, see “Filtering objects” on page 70.

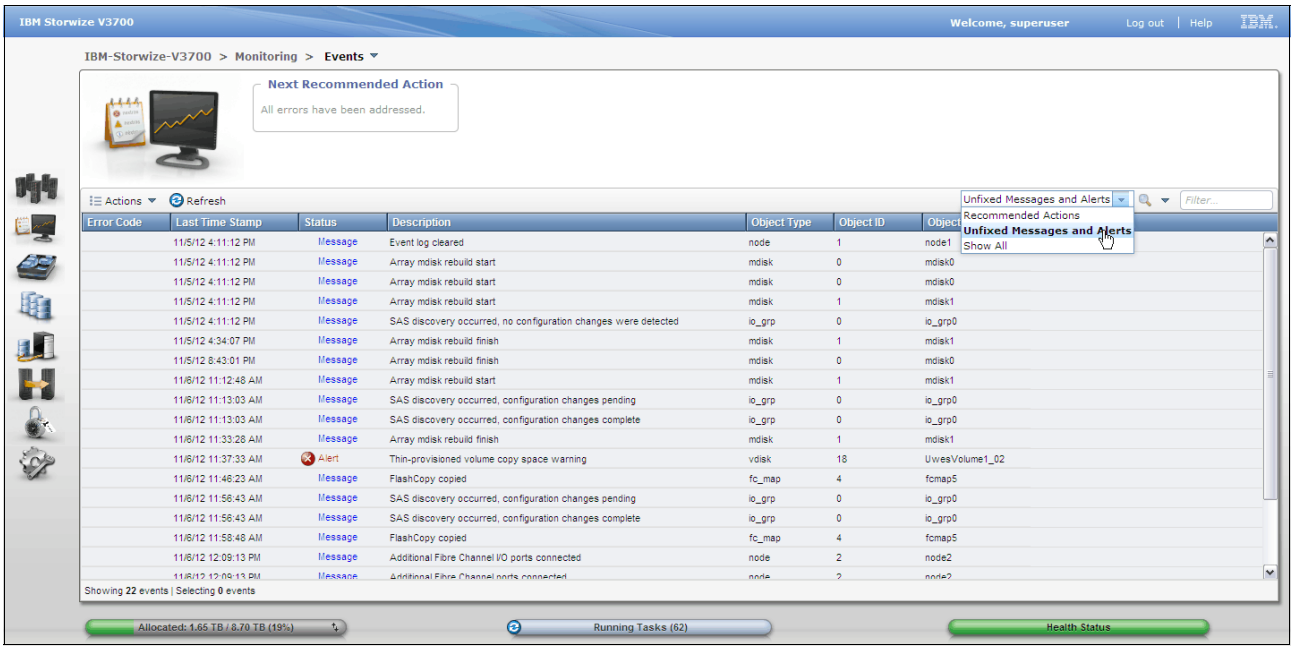

Figure 3-46 Unfixed messages and alerts in the events panel

To show all messages select **Show All** option from the drop-down menu that is next to the filter field. Figure 3-47 shows Events panel.

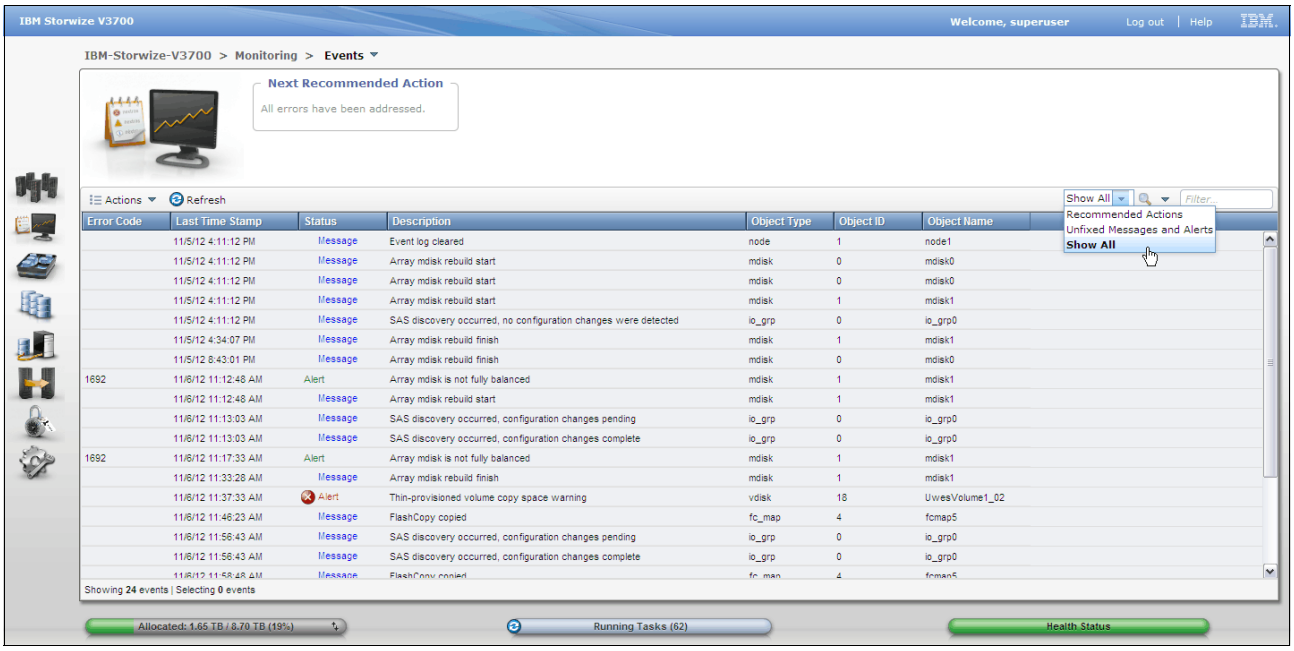

Figure 3-47 Events panel with show all selected

Event properties

To show the 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 3-48 shows the selection of the **Properties** option.

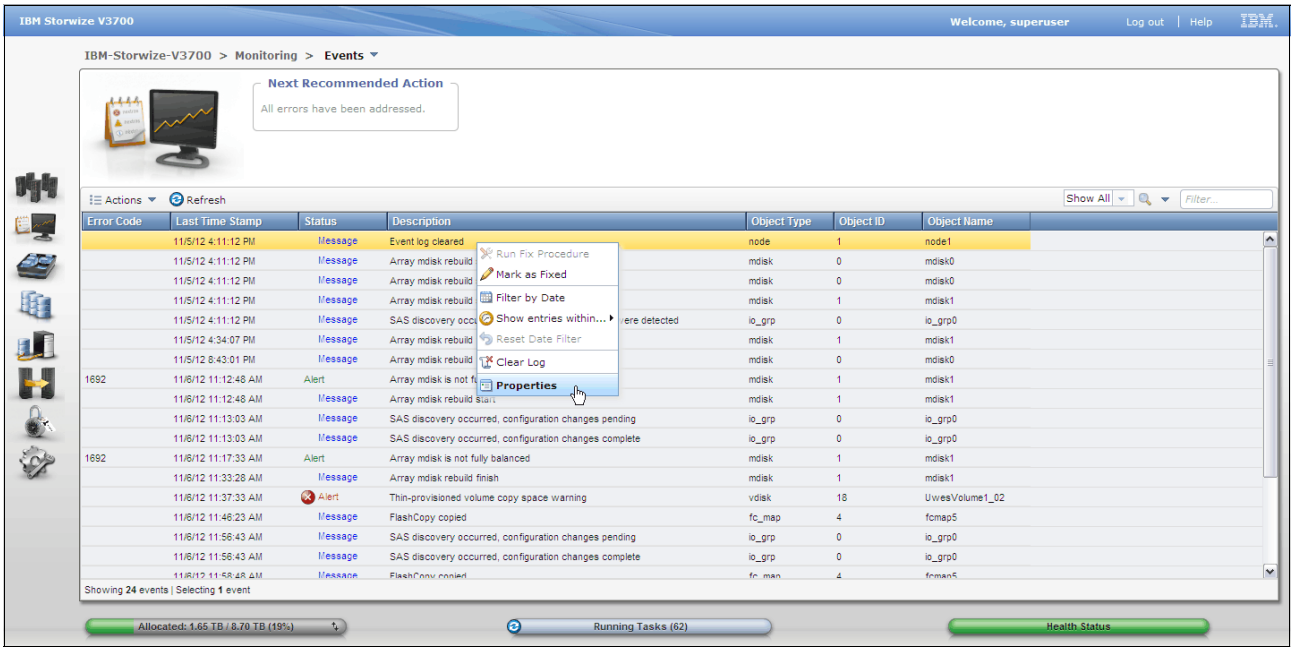

Figure 3-48 Selecting event properties

Figure 3-49 shows the properties of an event.

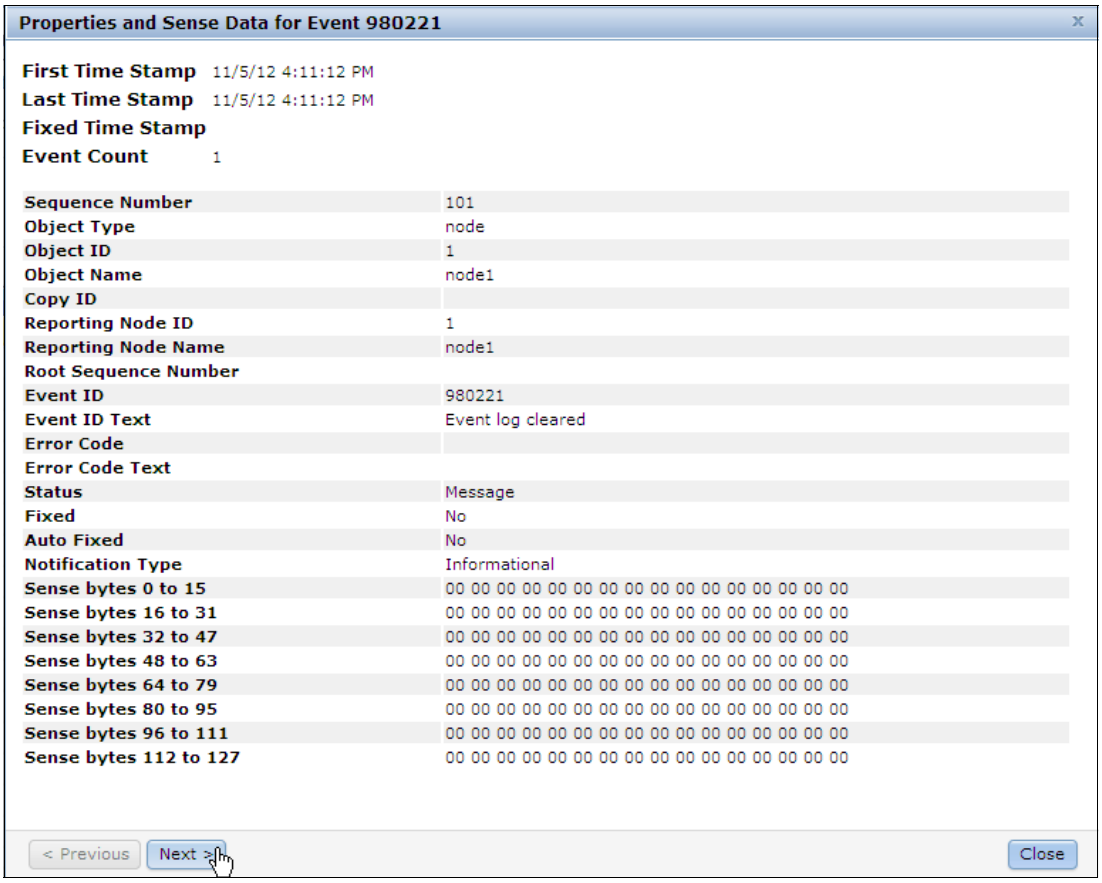

Figure 3-49 Event properties

Show events entries within...

To show events that occurred within a certain time, select **Show entries within...** from the actions drop-down menu and set the period value. Figure 3-50 shows the selection of the Show entries within... option with a period value of 1 minute.

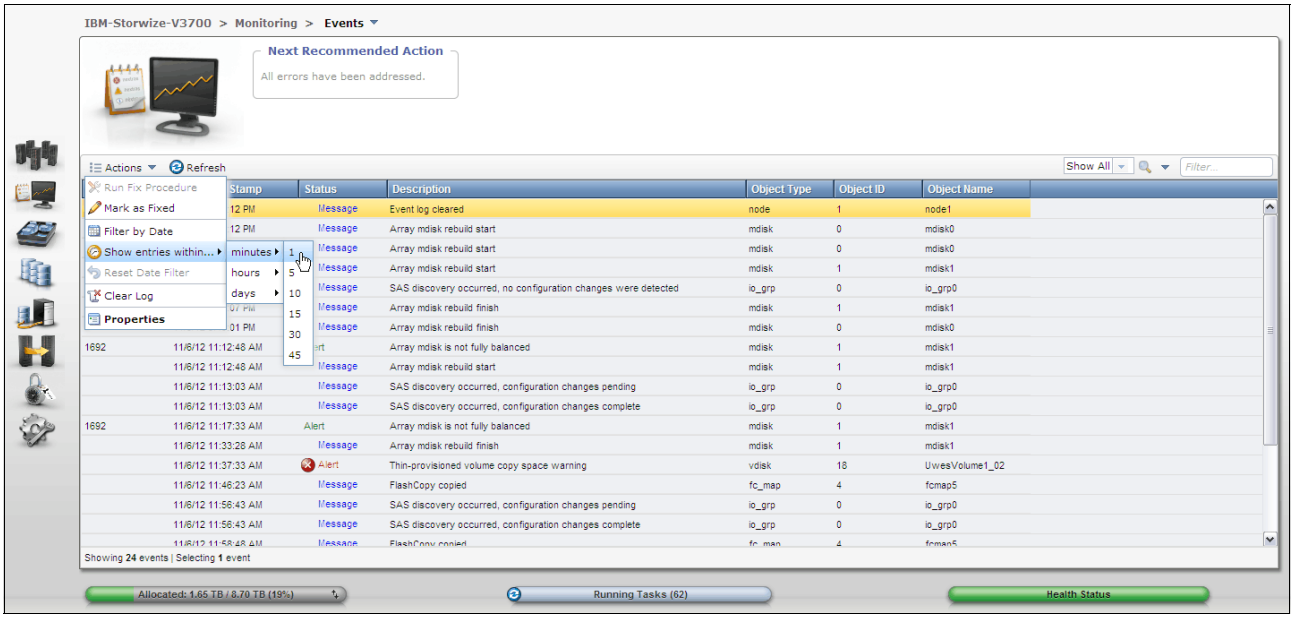

Figure 3-50 Showing events within a set time

Performance panel

Select **Performance** in the Monitoring menu to show the panel, as shown in Figure 3-51. This panel features graphs that represent the last 5 minutes of performance statistics. The performance graphs include statistics that are related to CPU utilization, volumes, MDisks, and interfaces.

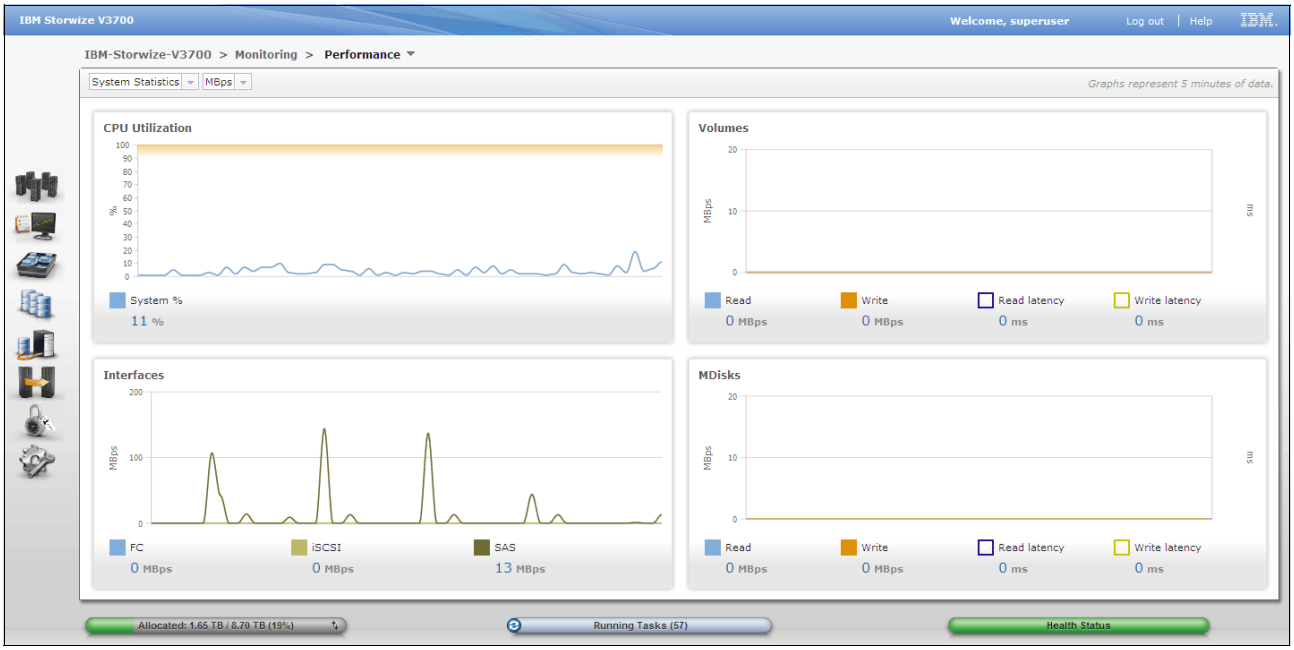

Figure 3-51 Performance panel

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 drop-down menu.

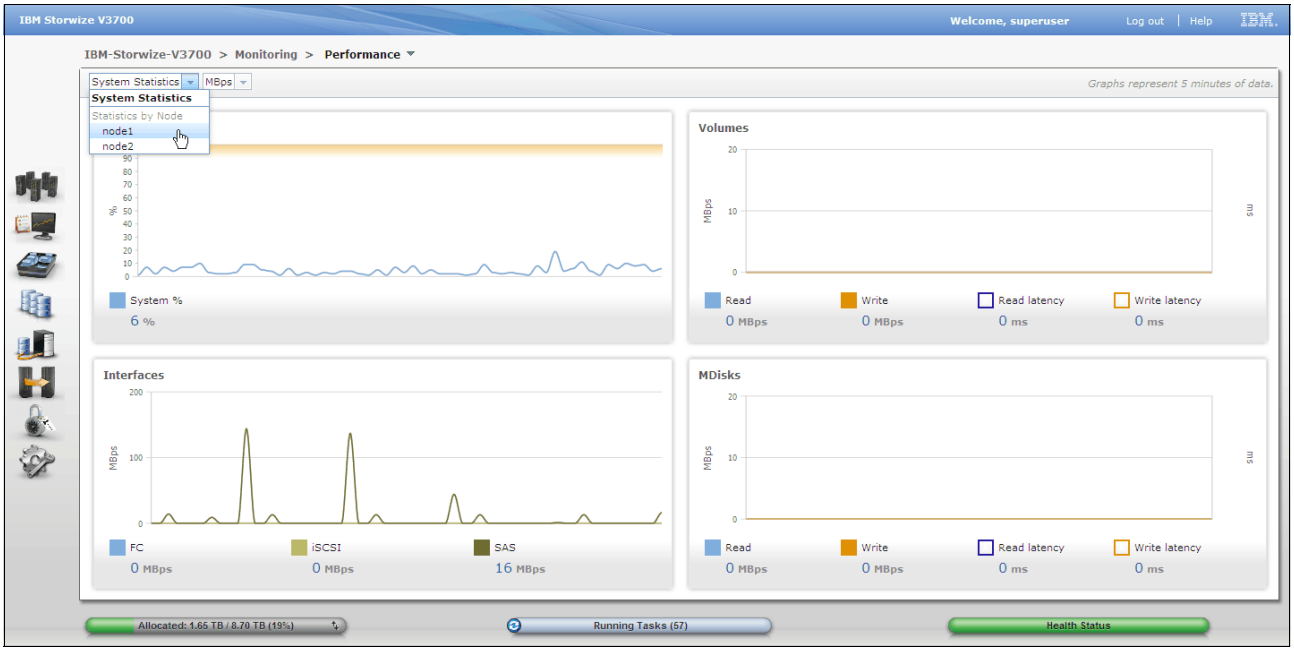

Figure 3-52 Graphs representing performance statistics of a single node

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.

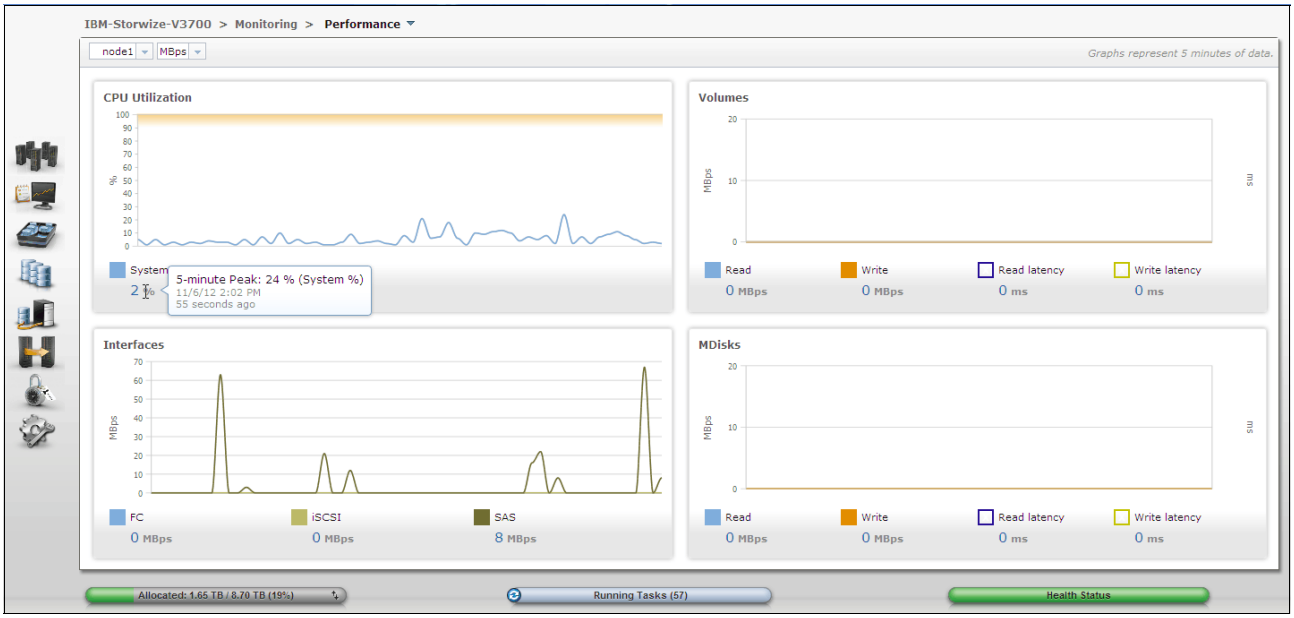

Figure 3-53 Peak CPU utilization value over the last 5 minutes

3.4.3 Pools menu

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

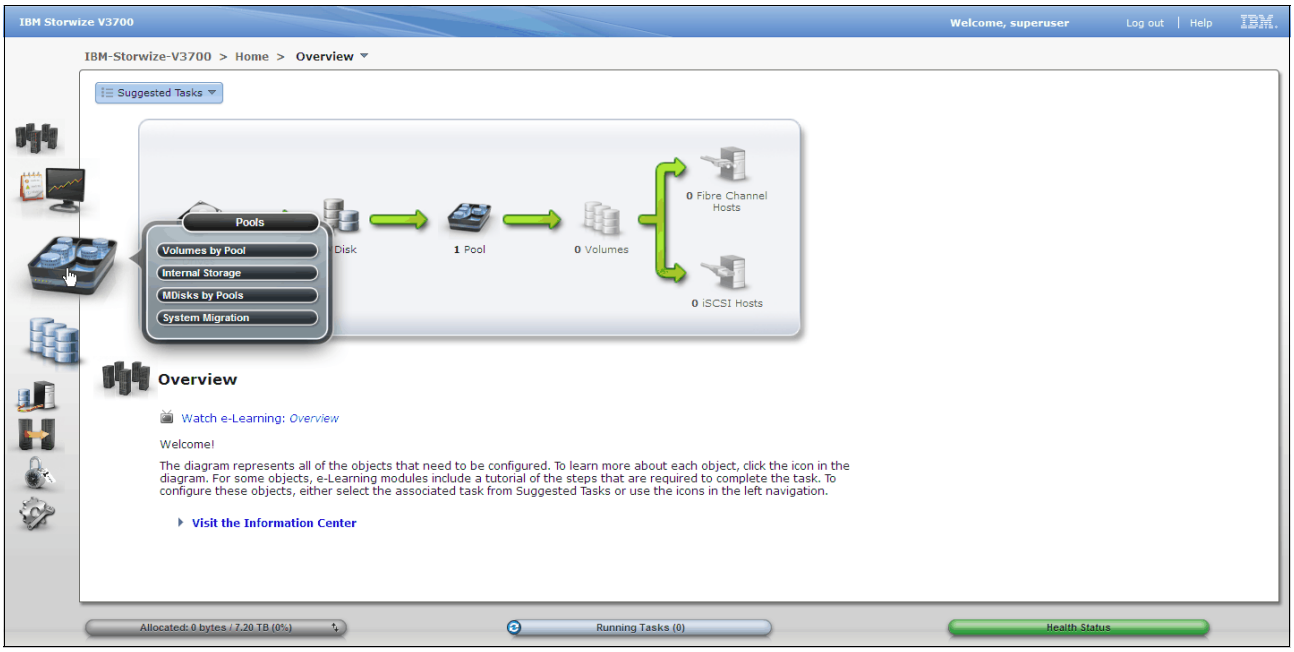

Figure 3-54 Pools menu

Volumes by pools panel

Select **Volumes by Pool** in the Pools menu to show the panel, as shown in Figure 3-55. The Volumes by Pool panel shows 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 further allocations.

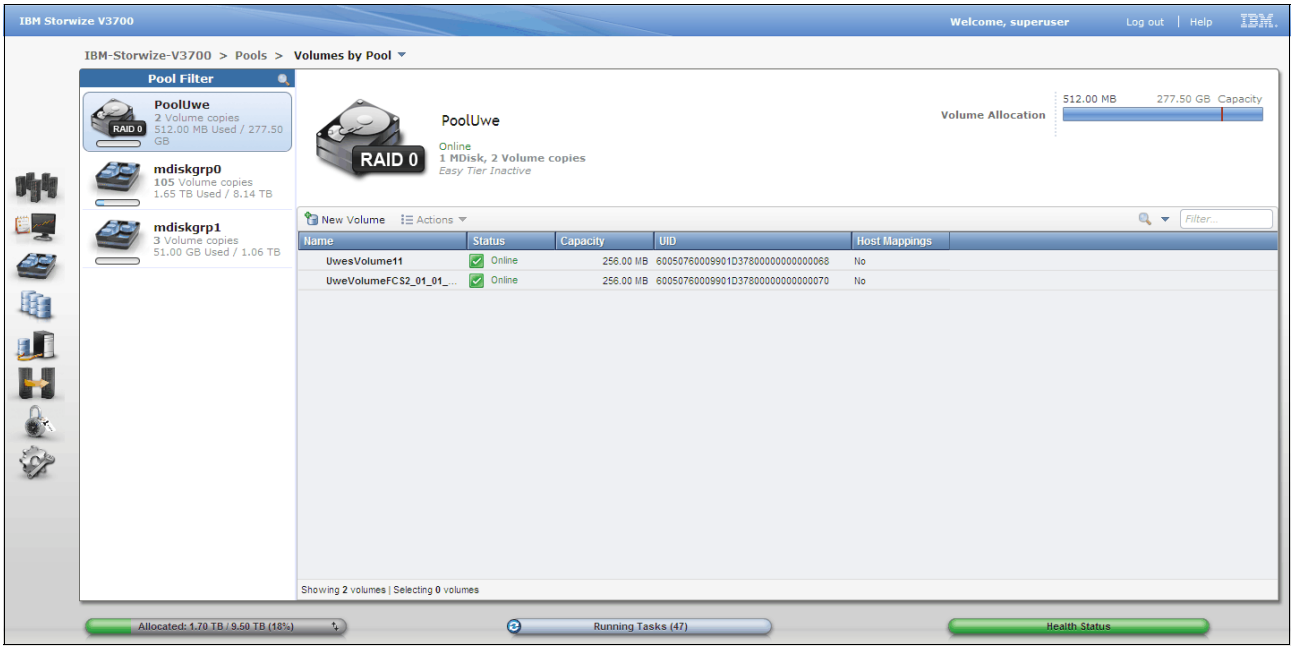

Figure 3-55 Volumes by Pools panel

Volume allocation

The upper right corner of the Volumes by Pool panel shows the actual volume allocation, which is the amount of capacity that is allocated to volumes in this pool. Figure 3-56 highlights the volume allocation information that is shown in the Volumes by Pools panel.

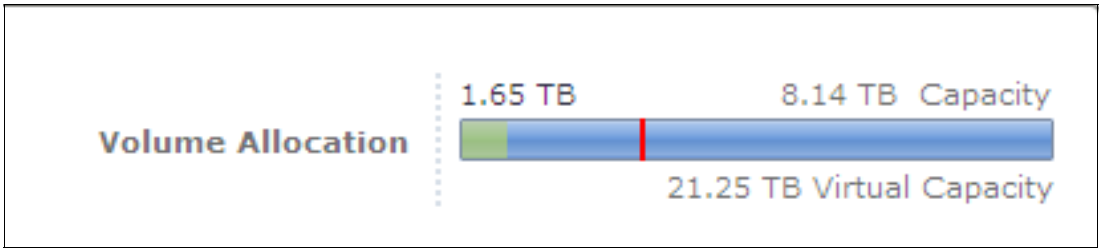

Figure 3-56 Volume Allocation

Renaming pools

To change the name of a pool, select the pool from the pool filter and click the name of the pool. Figure 3-57 shows that pool mdiskgrp0 was renamed to Gold Pool.

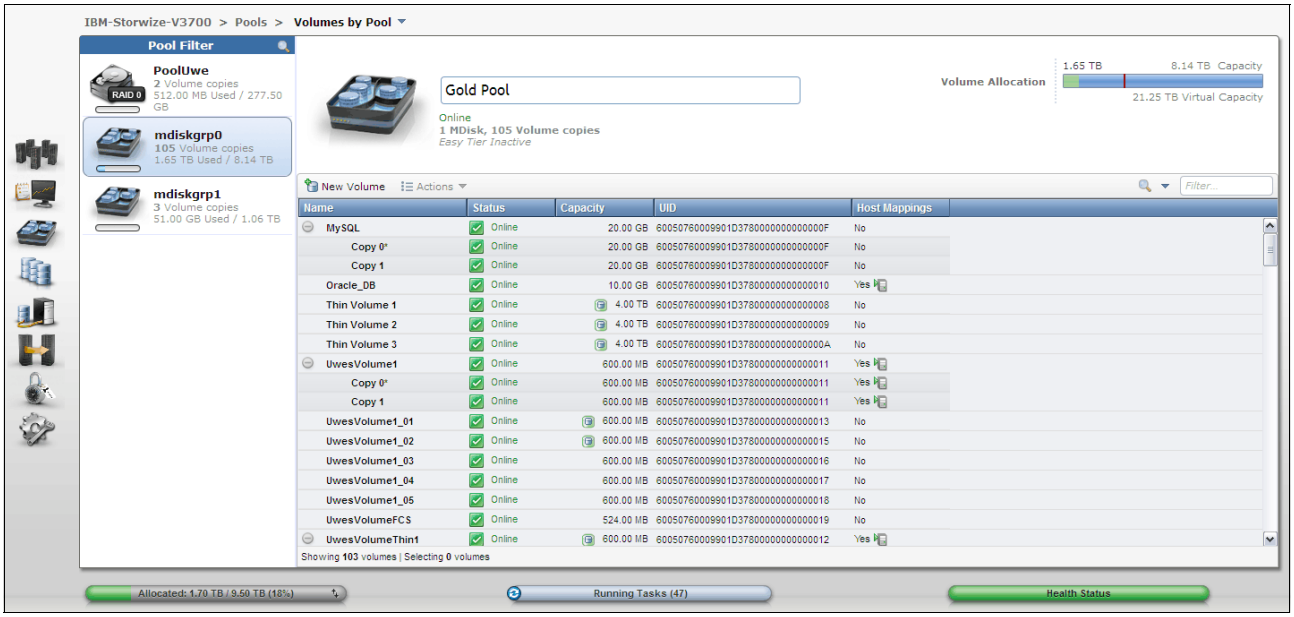

Figure 3-57 Renaming a pool

Changing pool icons

To change the icon that is associated with a pool, select the pool in the pool filter and click the large pool icon that is above New Volume and Actions. This change helps to manage and differentiate between the different 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 Pools panel also provides access to the volume functions via the Actions drop-down menu, the New Volumes option, and via right-clicking a listed volume. For more information about navigating the volume panel, see 3.4.4, “Volumes menu” on page 99. Figure 3-59 shows the volume functions that are available via the Volumes by Pool panel.

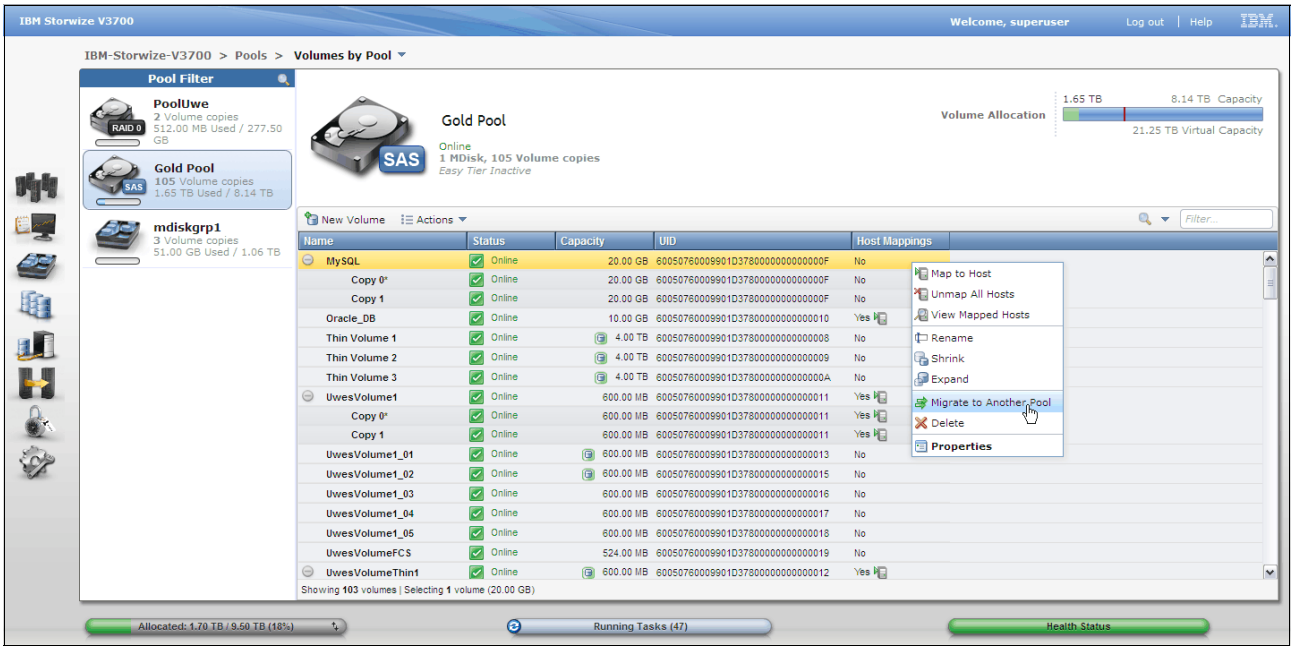

Figure 3-59 Volume functions are available via the Volumes by Pools panel

Internal storage panel

Select **Internal Storage** in the pools menu to show the internal storage panel, as shown in Figure 3-60. The internal storage consists of the drives that are contained in the Storwize V3500 control enclosure. By using the internal storage panel, you can configure the internal storage into RAID-protected storage. This panel also allows the displayed drive list to be filtered by drive class.

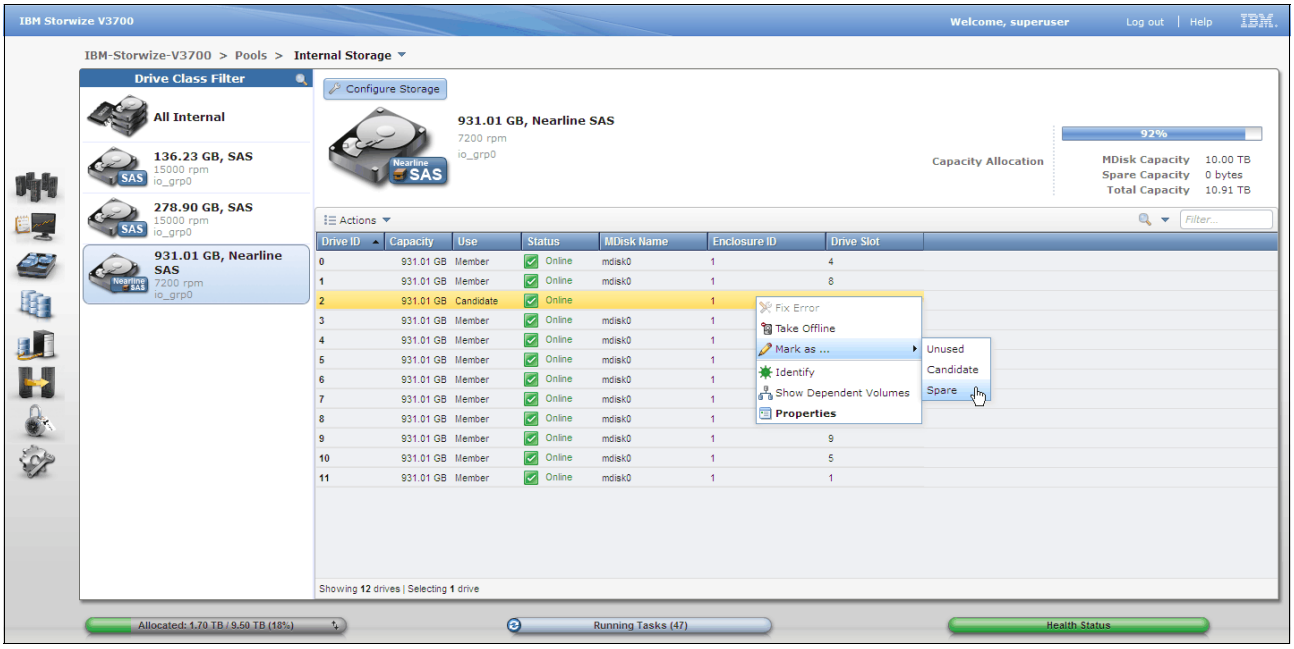

Figure 3-60 Drive actions menu of the internal storage panel

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-clicking a listed drive shows the Action menu. Alternatively, the drives can be selected and the Action menu is used. Multiple drives also can be selected. For more information, see “Multiple selections” on page 69. Figure 3-60 shows the drive actions menu.

Drive properties

Drive properties and dependent volumes can be displayed from the Internal Storage panel. Figure 3-61 shows the selection of drive properties from the Action menu.

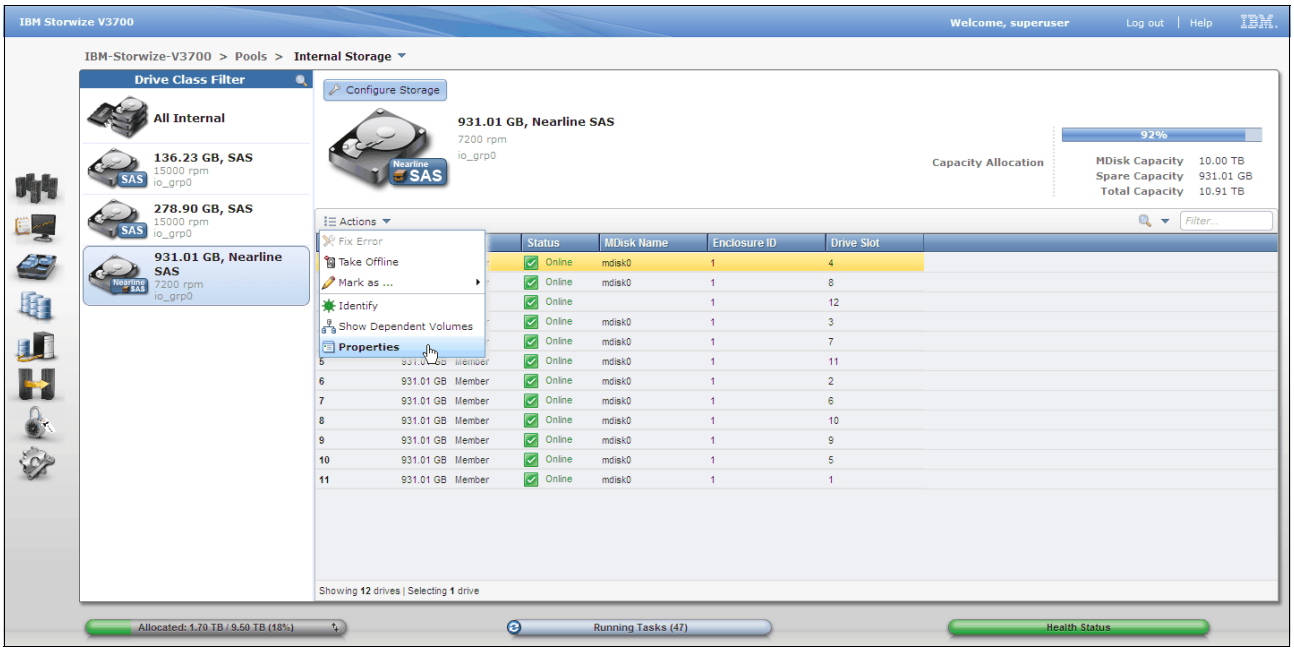

Figure 3-61 Internal storage panels actions

The Drive Properties panel shows drive attributes and the drive slot SAS port status, as shown in Figure 3-62.

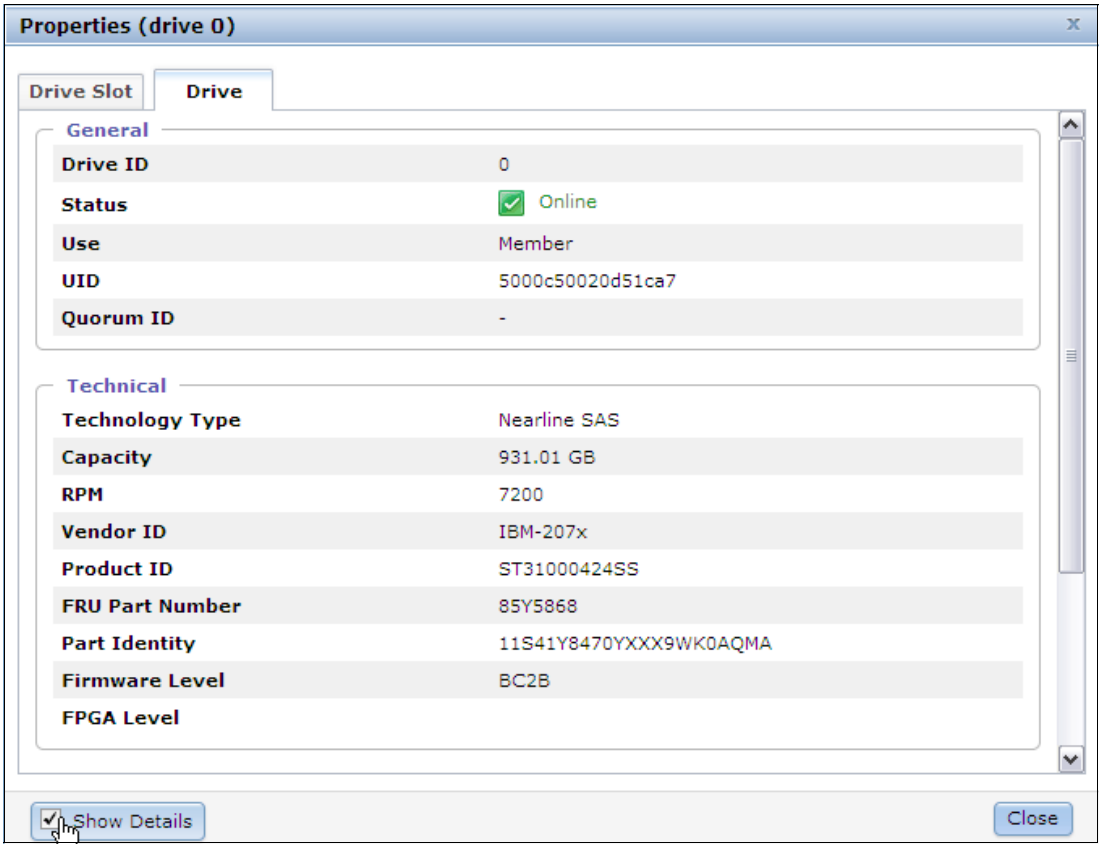

Figure 3-62 Drive properties

Configure Internal Storage wizard

Click **Configure Storage** to show the Configure Internal Storage wizard, as shown in Figure 3-63.

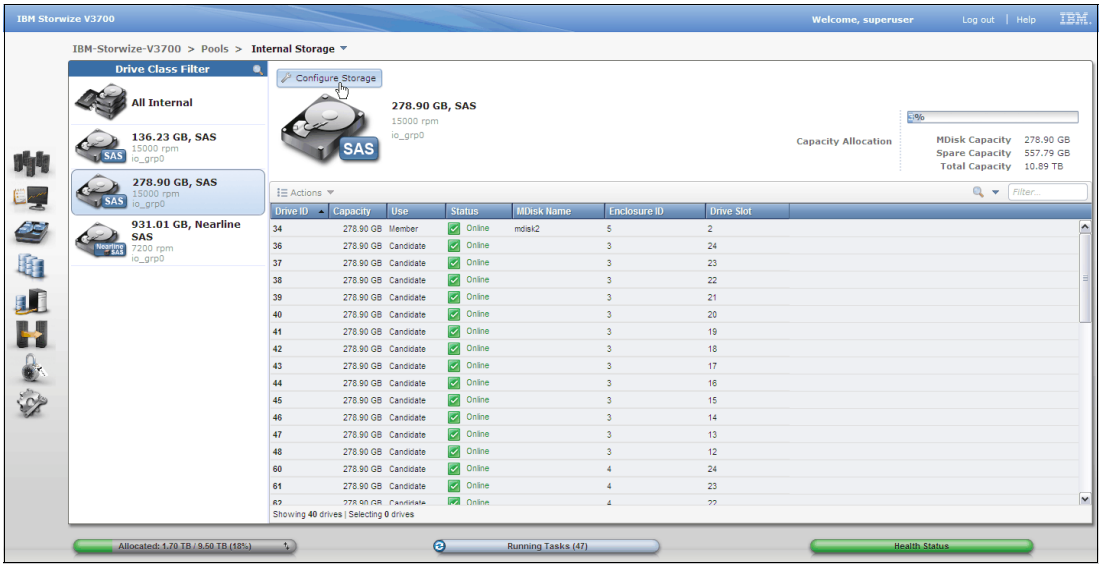

Figure 3-63 Internal Storage panel

By using this wizard, you can configure the RAID-properties and pool allocation of the internal storage. Figure 3-64 shows step 1 of the Configure Internal Storage wizard.

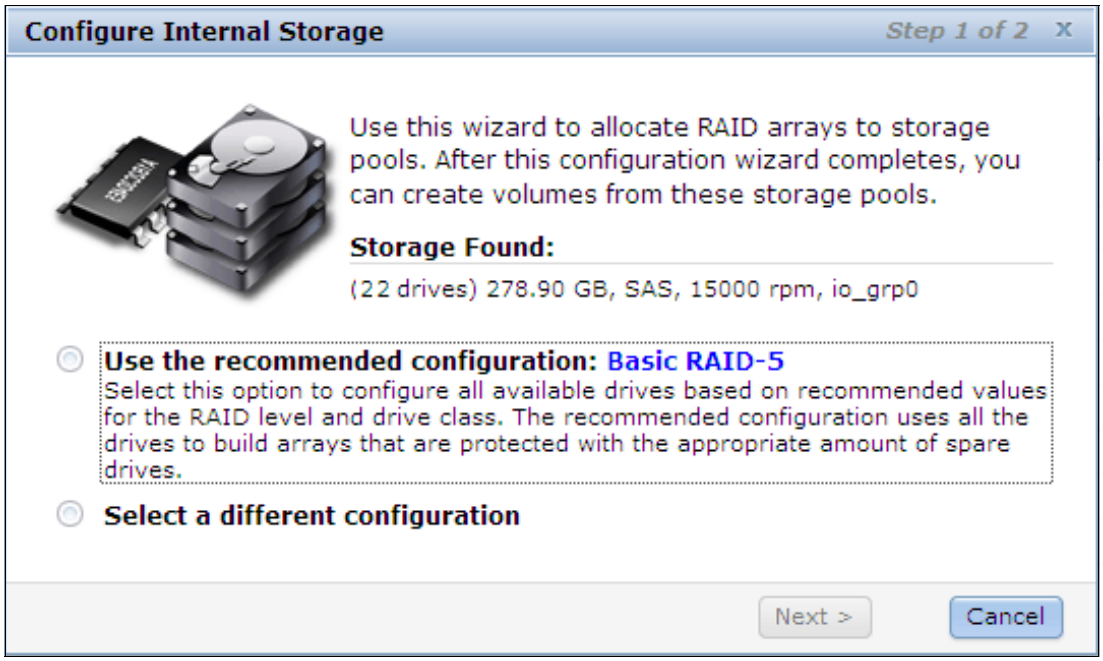

Figure 3-64 Configure internal storage wizard: Step 1

Figure 3-65 shows step 2 of the Configure Internal Storage wizard.

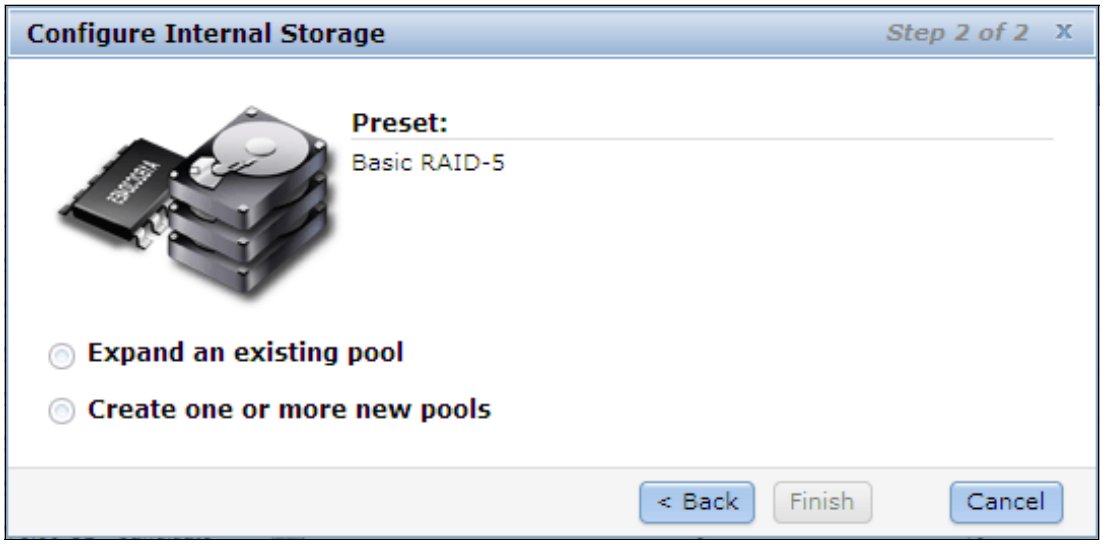

Figure 3-65 Configuring internal storage wizard: Step 2

MDisks by Pools panel

Select **MDisks by Pools** in the Pools menu to show the MDisks by Pools panel, as shown in Figure 3-66. By using this panel, you can perform tasks such as: display MDisks in each pool, create pools, delete pools, and detect externally virtualized storage.

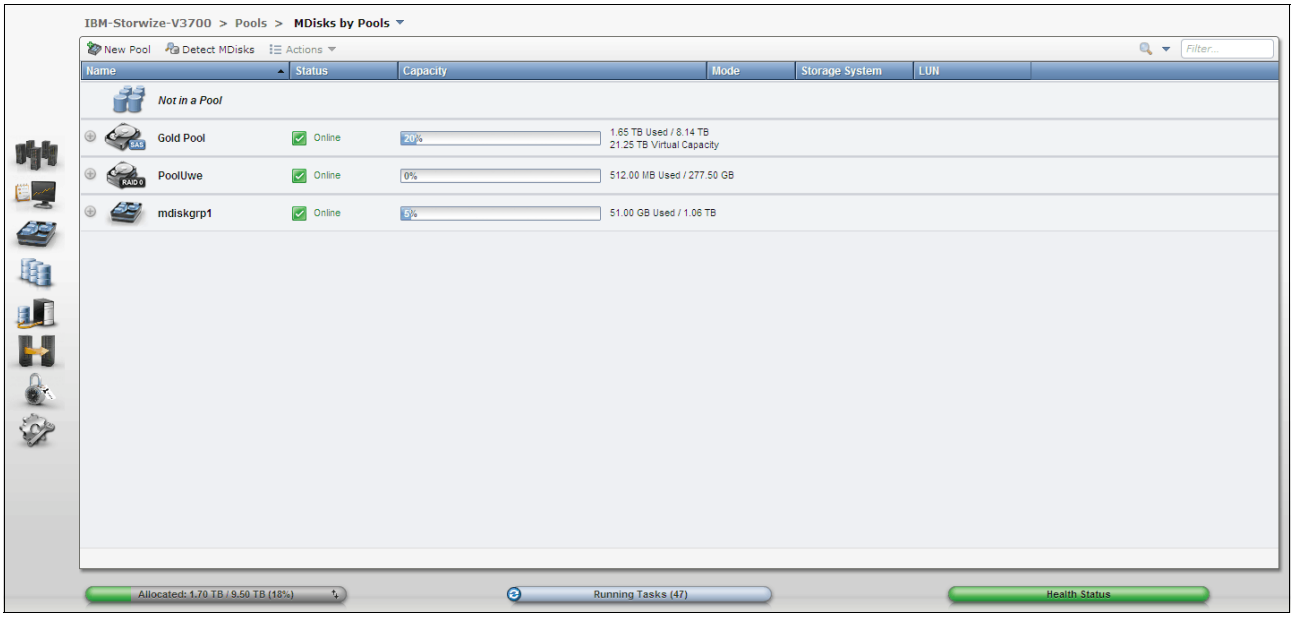

Figure 3-66 MDisks by Pool panel

Pool actions

To delete a pool or change the pool name or icon, right-click the listed pool. Alternatively, the Actions drop-down menu can be used, as shown in Figure 3-67.

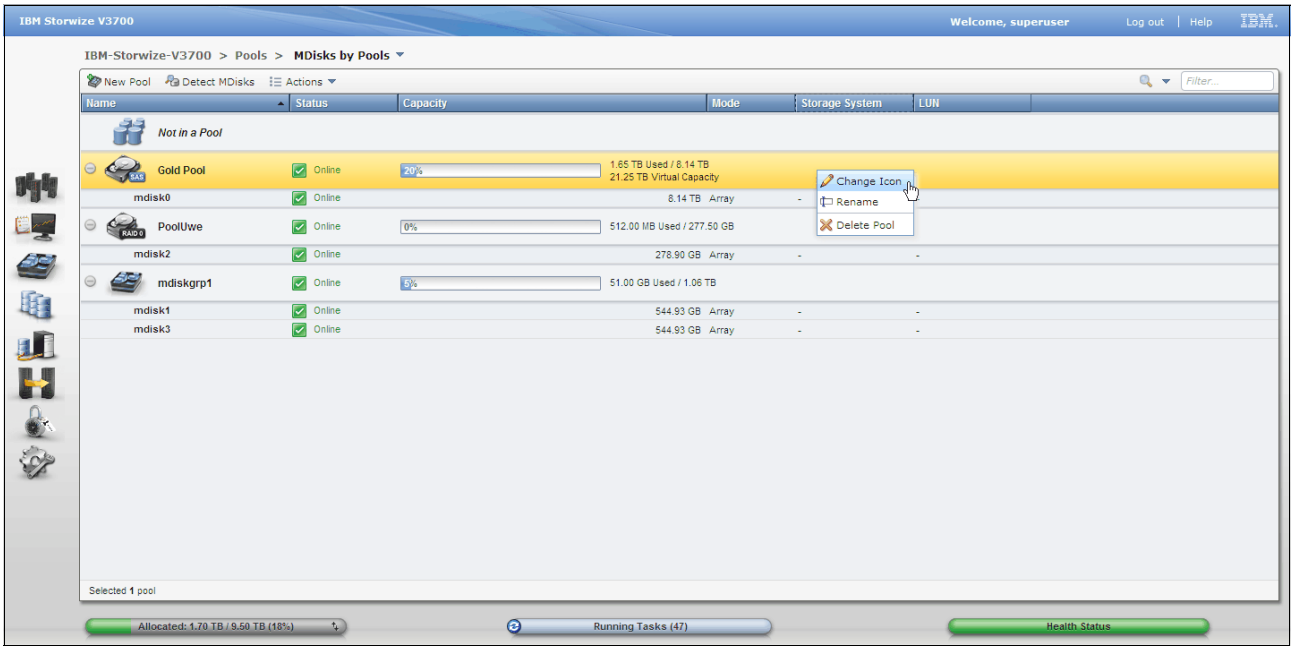

Figure 3-67 Pool actions

RAID actions

By using the MDisks by Pools panel, you can perform MDisk RAID tasks such as: set spare goal, swap drive, and delete. To access these functions, select and right-click the MDisk. Figure 3-68 shows the RAID actions menu.

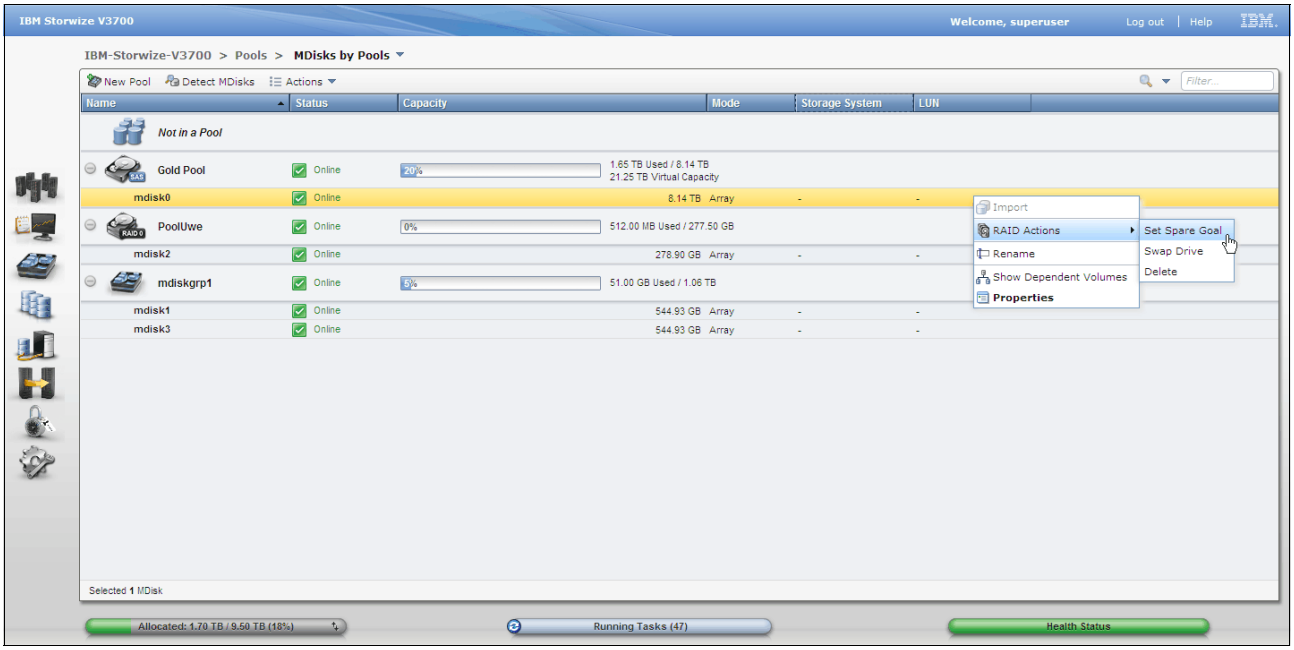

Figure 3-68 RAID actions menu

System Migration panel

Select **System Migration** in the Pools menu to show the System Migration panel, as shown in Figure 3-69. By using this panel, you can migrate data from externally virtualized storage systems to the internal storage of the Storwize V3500. The panel displays image mode volume information. To start a migration, click **Start New Migration** and the Start Migration wizard opens.

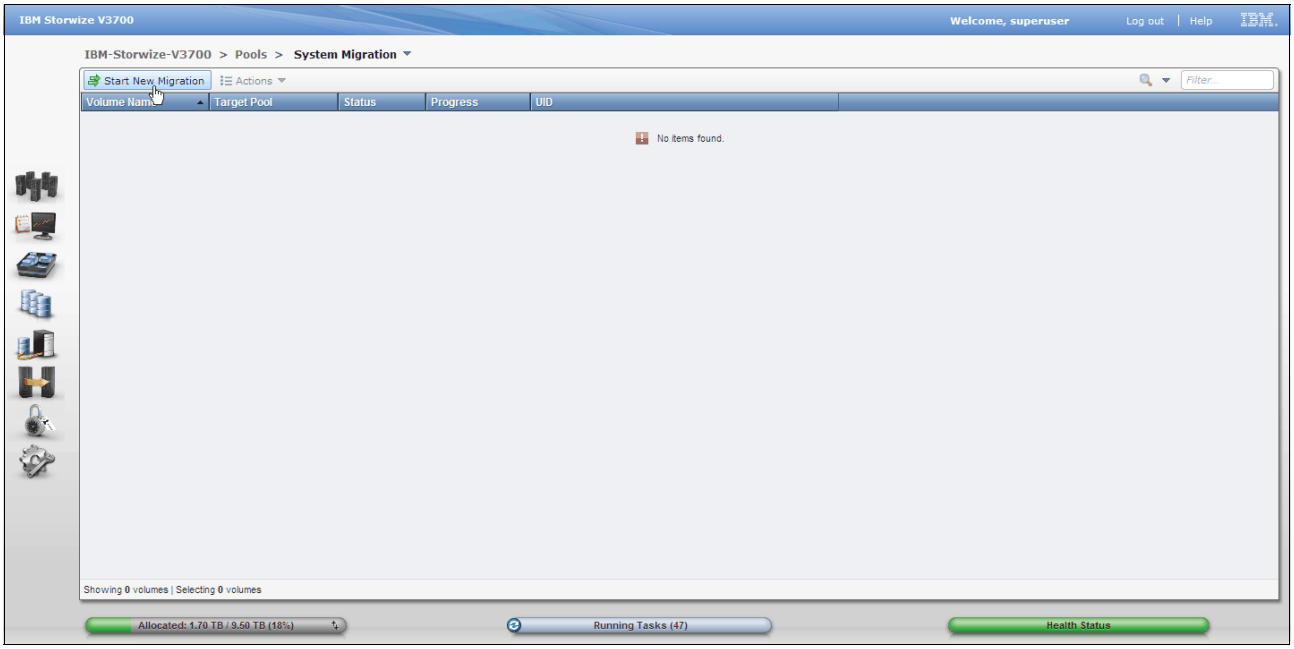

Figure 3-69 System migration panel

Storage Migration wizard

The Storage Migration wizard is used to migrate data from older storage systems to the Storwize V3500. Figure 3-70 shows step 1 of the Storage Migration wizard.

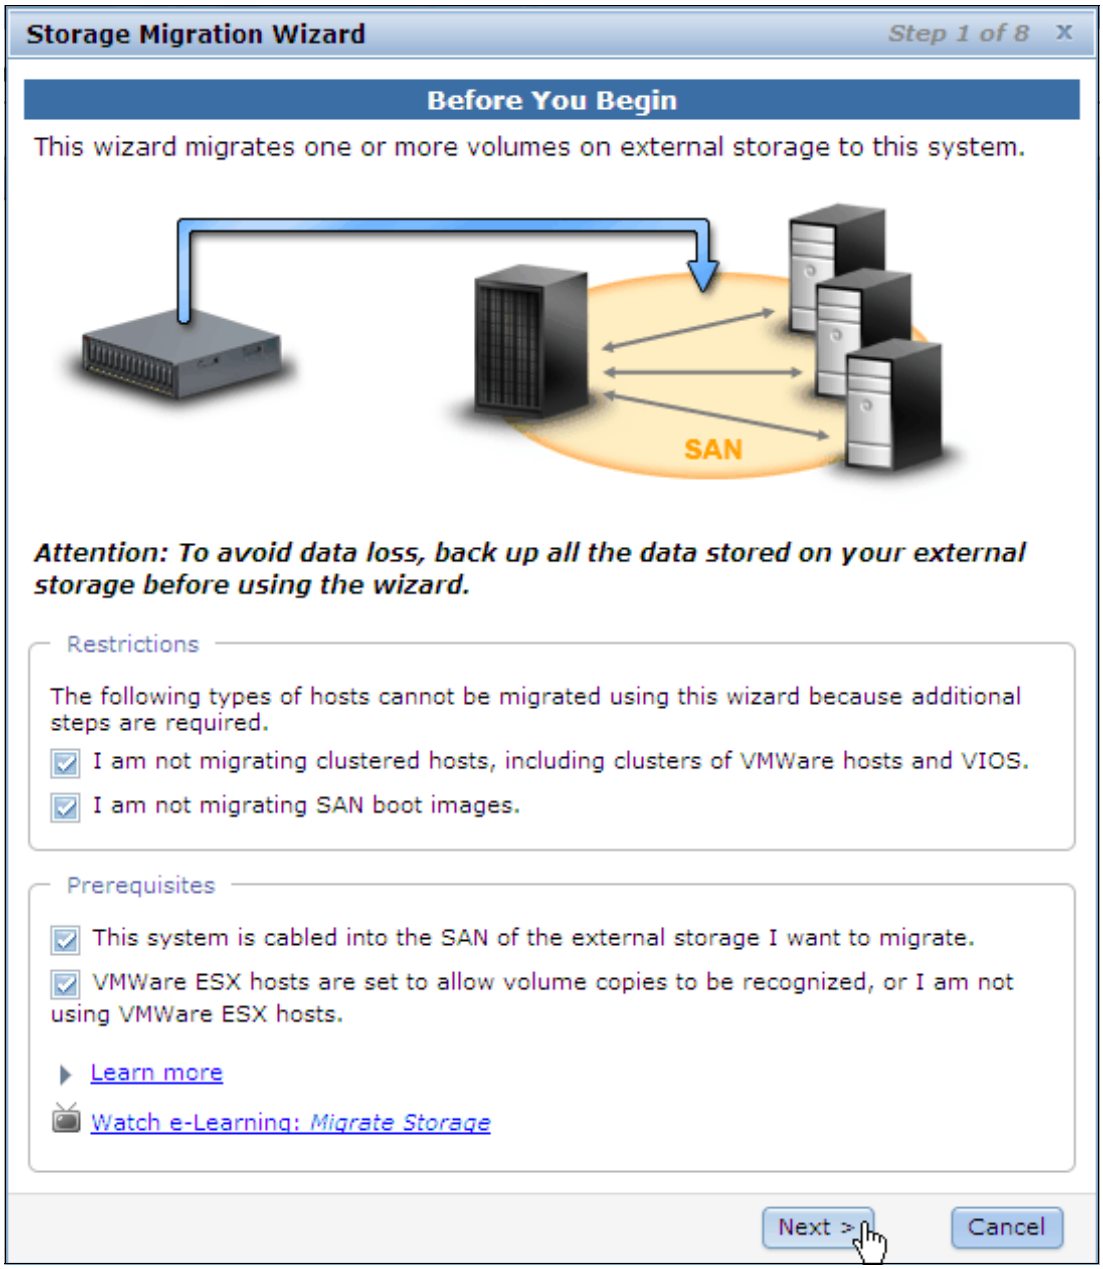

Figure 3-70 Storage migration wizard: Step 1

3.4.4 Volumes menu

As shown in Figure 3-71, the Volumes menu provides access to the Volumes, Volumes by Pool, and Volumes by Host functions.

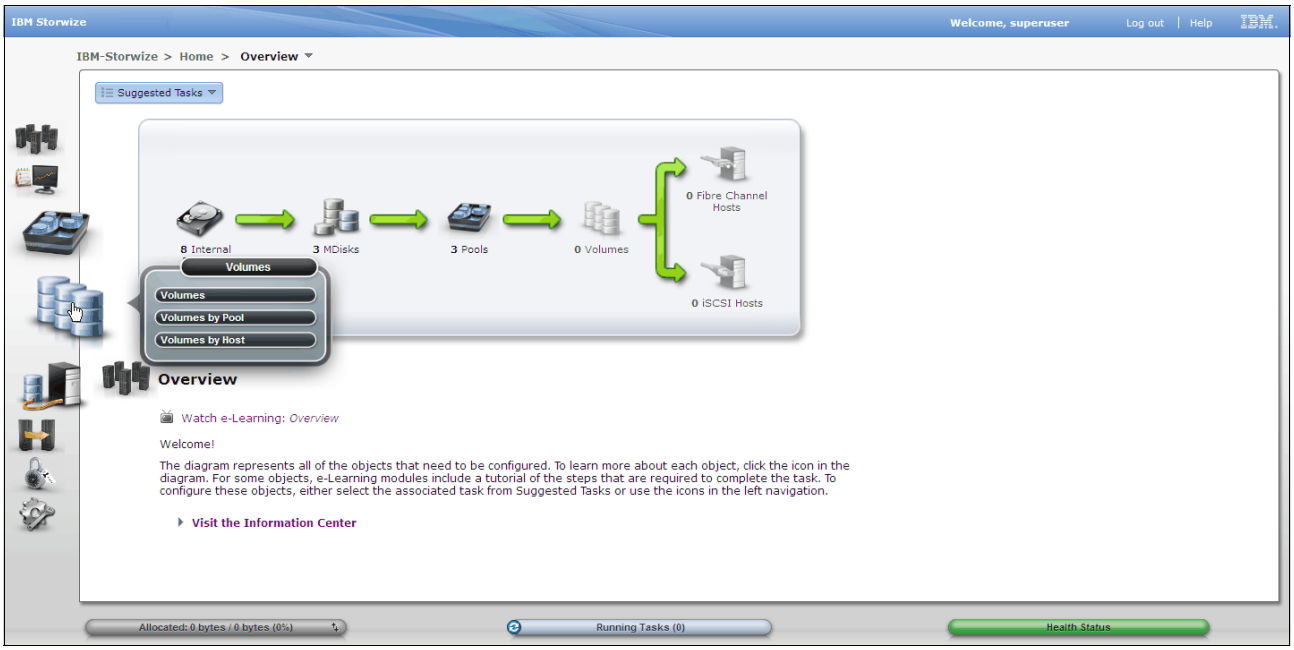

Figure 3-71 Selection of the volumes menu

Volumes panel

Select **Volumes** in the Volumes menu to open the panel, as shown in Figure 3-72. The Volumes panel shows all of the volumes in the system. The information that is displayed is dependent on the columns that are selected.

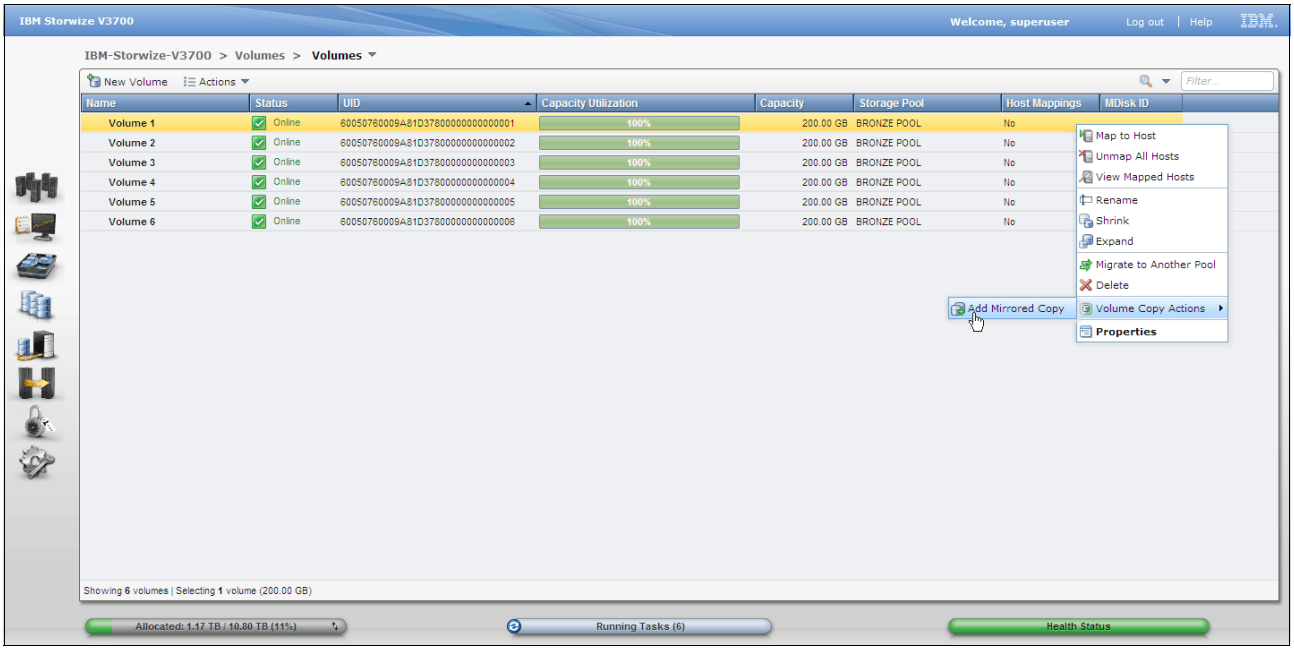

Figure 3-72 The volumes panel

Volume actions

Volume actions such as: map, unmap, rename, shrink, expand, migrate to another pool, delete, and mirror can be performed by using this panel.

Creating volumes

Click **New Volume** to open the New Volume panel, as shown in Figure 3-73. By using this panel, you can select a preset when you are creating a volume. The presets are designed to accommodate most of the cases. The following presets are available: 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 the size the volumes is then shown.

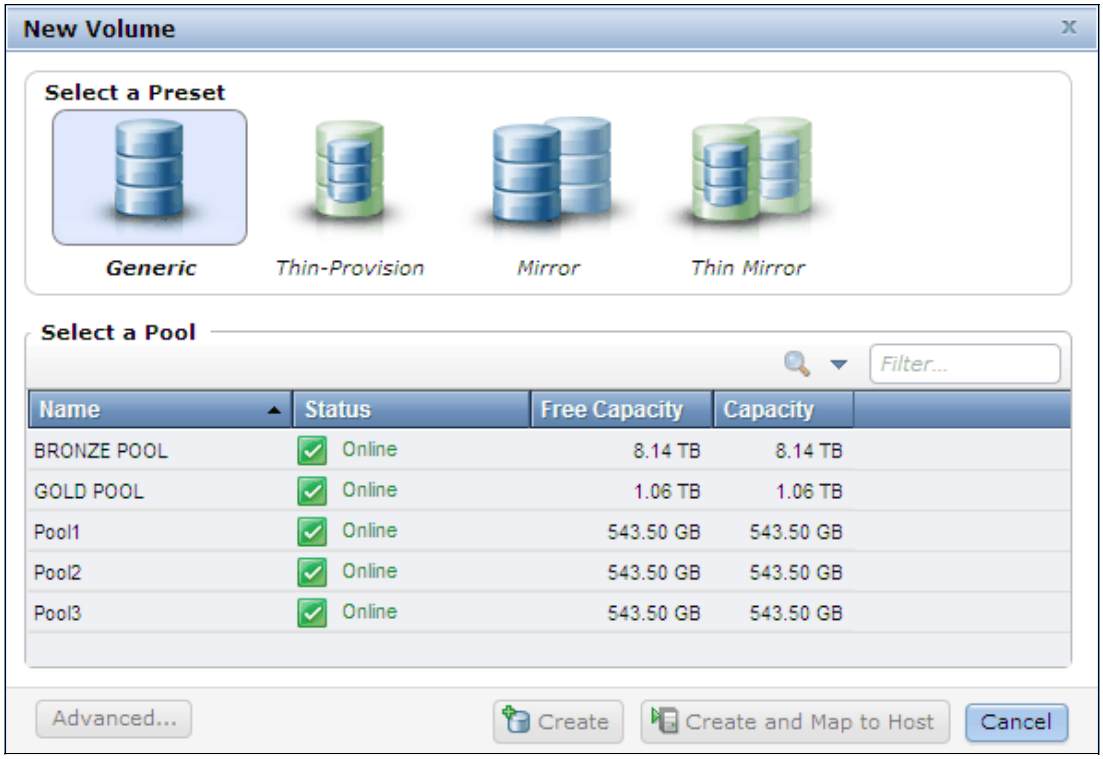

Figure 3-73 New Volume panel

Creating multiple volumes

A useful feature is available for quickly creating multiple volumes of the same type and size. After the name and size of the first volume is entered, adjust the quantity of more volumes of the same type and size that must be created in the Add Volumes field and then select **+** to create the volumes. Figure 3-74 shows the quantity of 5 in the Add Volumes field.

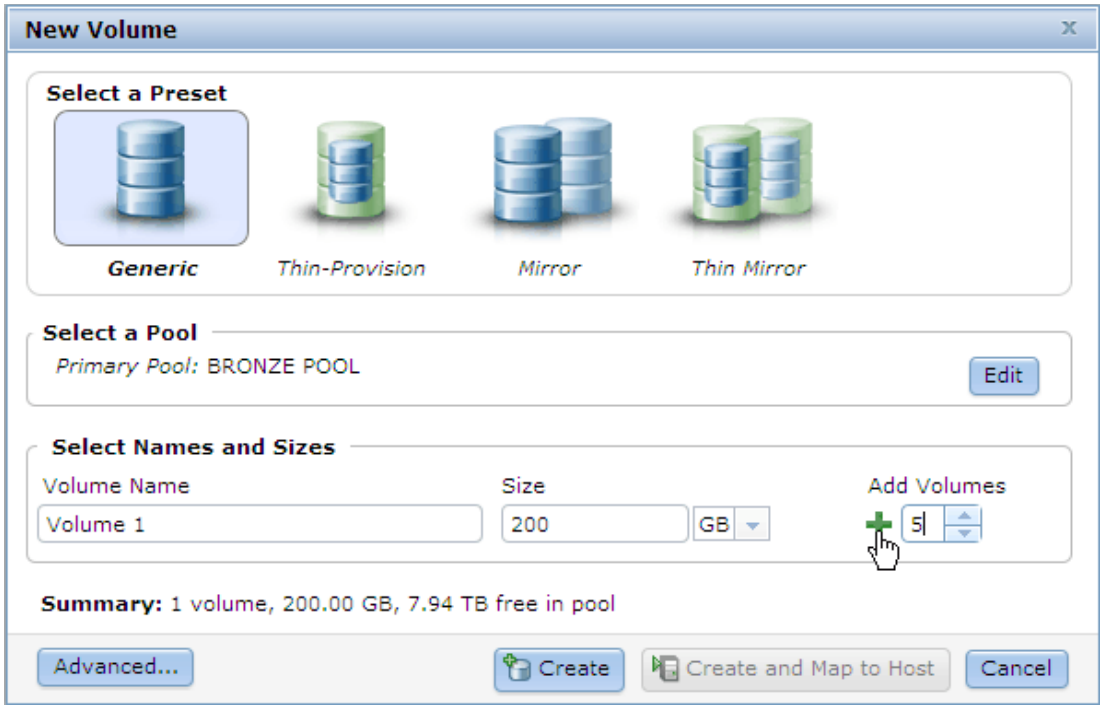

Figure 3-74 Creating multiple volumes quickly

When volumes are added, any number value in the volume name automatically increments to ensure that each volume is unique. To remove a volume from the list, click **X**.

Volume advanced settings

Click **Advanced** to see more volume configuration options. You can 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 3-75 shows the Advanced Settings panel.

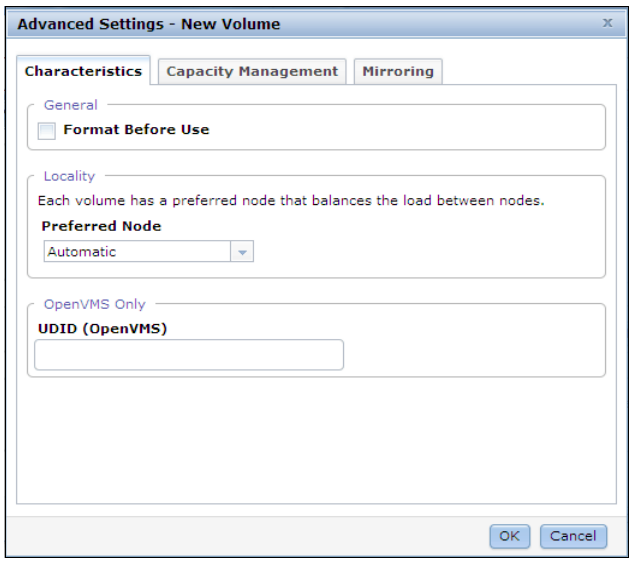

Figure 3-75 New volume advanced settings

Capacity Summary

The New Volumes panel also includes a summary that shows the real capacity that is used if the proposed volumes are created. After all of the volumes in the list meet the volumes requirements, click **Create** or **Create and Map to Host** to continue. Figure 3-76 shows the New Volume panel and the result of adding five volumes by using the add volumes feature.

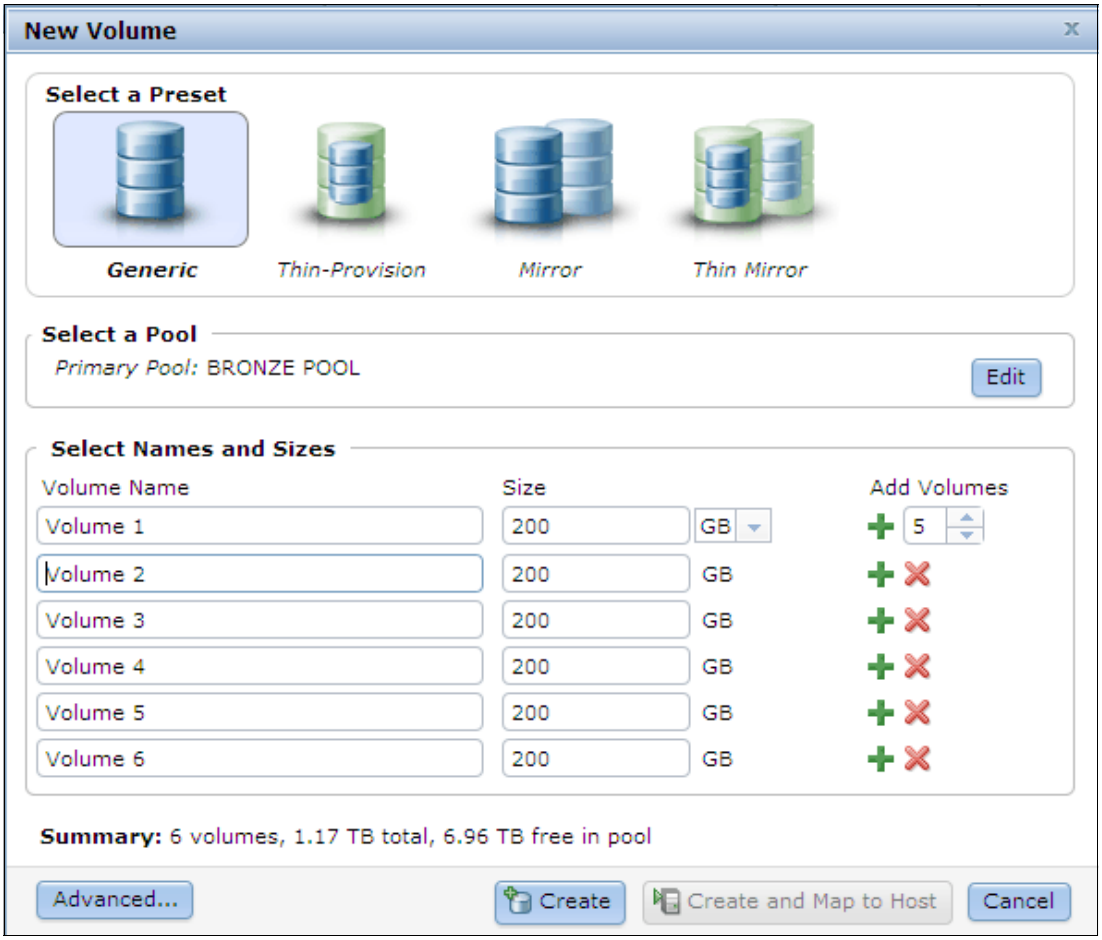

Figure 3-76 Multiple volumes that are created with the add volumes feature

Volumes by Pool panel

This method is an alternative method to access the Volumes by Pool panel. For more information, see “Volumes by pools panel” on page 88.

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 volumes that are allocated to a particular host by using the host selection filter.

Volume actions

After a host is selected, all of the volume actions such as: map, unmap, rename, shrink, expand, migrate to another pool, delete, and mirror can be performed. Figure 3-77 on page 104 shows the Volume by Host panel.

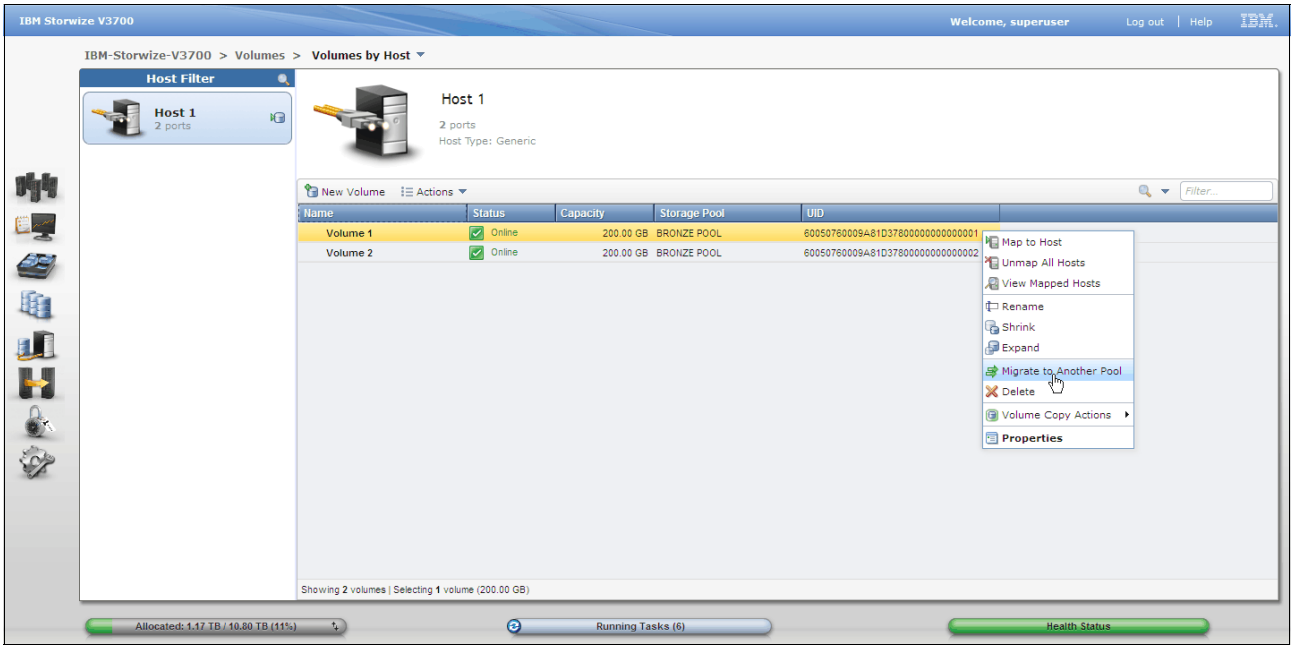

Figure 3-77 Volumes by Host panel

3.4.5 Hosts menu

As shown in Figure 3-78, the Hosts menu provides access to the Hosts, Ports by Host, Host Mappings, and Volumes by Host functions.

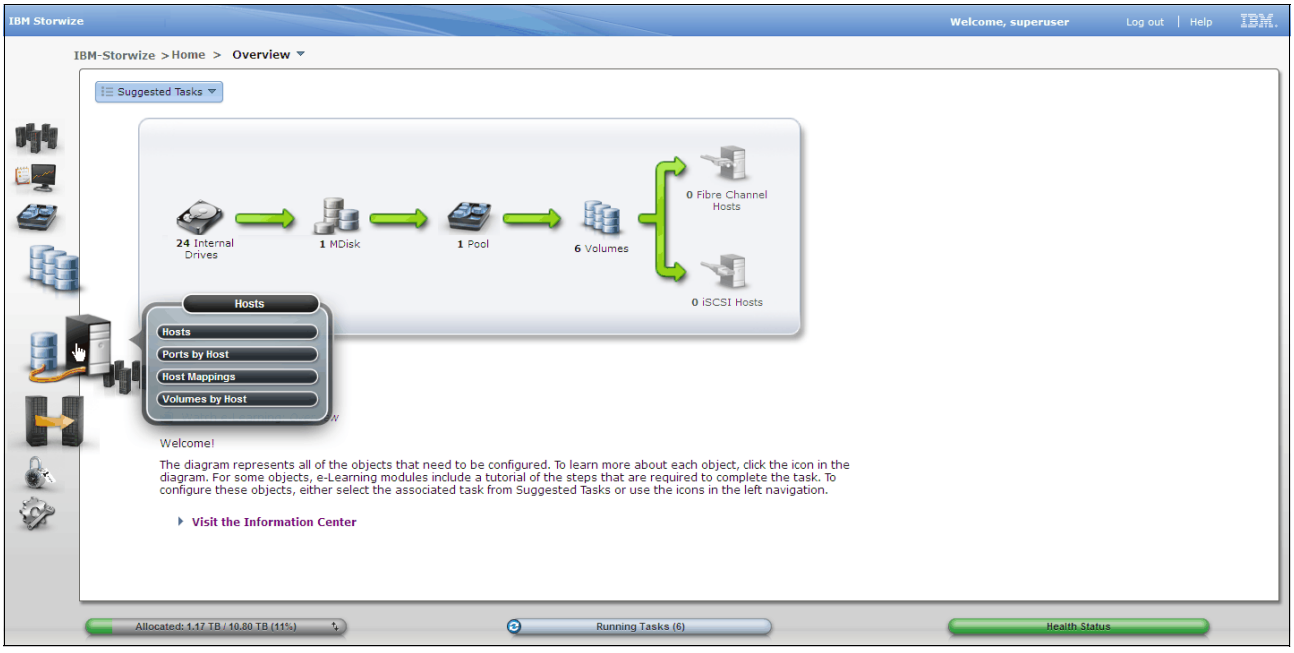

Figure 3-78 Selection of the Hosts menu

Hosts panel

Select the **Hosts** item in the Hosts menu to open the panel, as shown in Figure 3-79. The Hosts panel shows all of the hosts that are defined in the system.

Host Actions

Host Actions such as: map, unmap, rename, and delete can be performed from the Hosts panel, as shown in Figure 3-79.

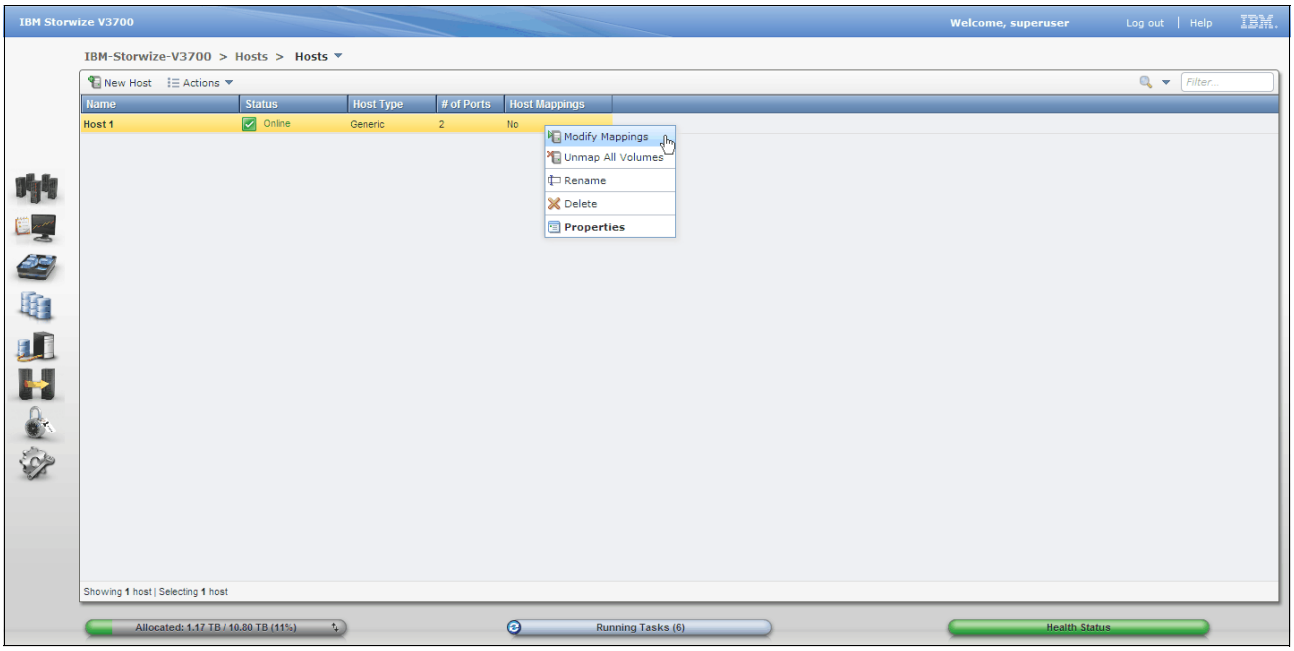

Figure 3-79 The Hosts panel

Creating a host

Click **New Host** and the Create Host window is opened. Choose the host type from Fibre Channel or iSCSI 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 3-80 shows the Choose the Host Type panel of the Create Host window.

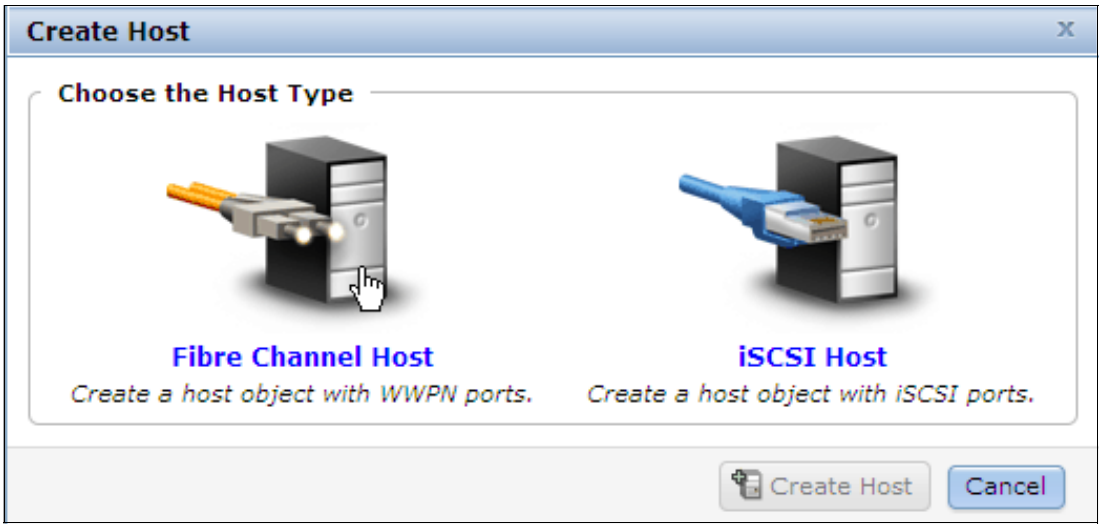

Figure 3-80 Choose the Host Type panel

Creating an FC host

Enter a name for the host in the field that is provided. The FC host port definitions are automatically discovered after the hosts zoning is activated. The discovered port definitions can be selected and added from the Add Port to List drop-down menu.

Select the port definition from the list and click **Add Port to List** to assign the port to the host definition. If the drop-down menu includes the message “No candidate HBA ports were found”, or an expected port definition was not automatically discovered and displayed, check zoning and click **Rescan**.

A port definition with an adjacent green check box means that the port is online and operational. To remove a port definition from the list, click the red **X**.

By using this panel, you can manually enter the port definitions. To perform this task, enter the WWPN in the field that is provided (the WWPN is not case-sensitive) and accept the bytes that are separated by colon. However, do not insert a space between each byte.

Ensure that the port definition you entered contains exactly 16 characters and click **Add Port to List** to assign the port to the host definition. This function is useful when the zoning is not yet activated and the WWPNs are known. In this case, the port includes a question mark symbol and `(unverified)` next to the listed port definition.

After the zoning is complete, the hosts port status can be seen via the Ports by Host panel. For more information, see “Ports by Host panel” on page 111. Advanced Settings options should be selected to allow for the configuration of HP/UX, OpenVMS, and TPGS hosts. At least one port definition must be assigned to activate the Create Host option. Click **Create Host** and the host creation task runs.

Figure 3-81 shows the FC Create Host window.

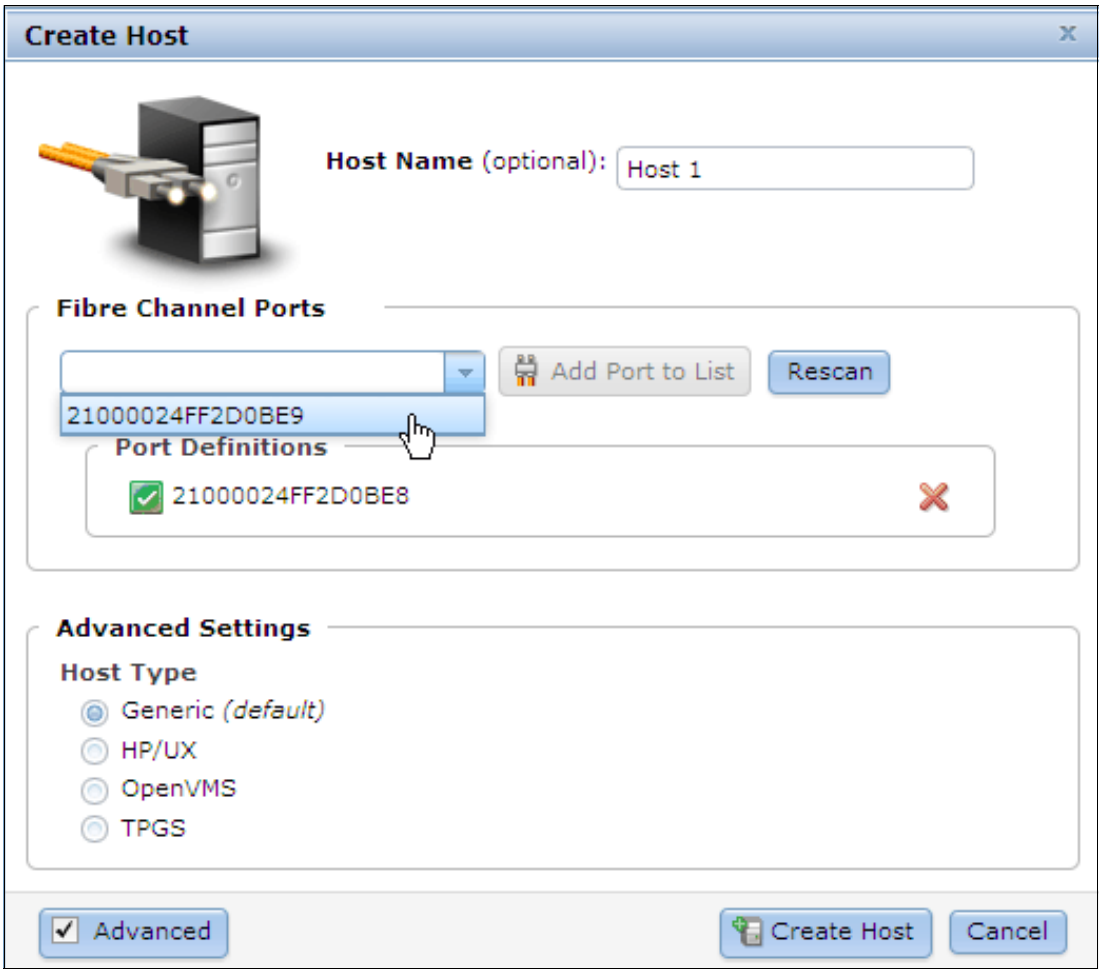

Figure 3-81 FC Create Host window with advanced settings selected

Creating an iSCSI host

Enter a name for the host in the field that is provided. By using this panel, you can enter the host port definitions. Enter the iSCSI IQN in the field that is provided and click **Add Port to List** to assign the port to the host definition.

To remove a port definition from the list, click the red **X**. Advanced Settings should be selected to allow for the configuration of HP/UX, OpenVMS, and TPGS hosts. **Use CHAP authentication** should be selected if it is required. However, to help in problem determination, this step can be delayed until after the first few hosts are configured and their connectivity is tested without authentication that is configured. At least one port definition must be assigned to enable the Create Host option. Click **Create Host** and the host creation task runs.

Figure 3-82 shows the iSCSI Create Host panel.

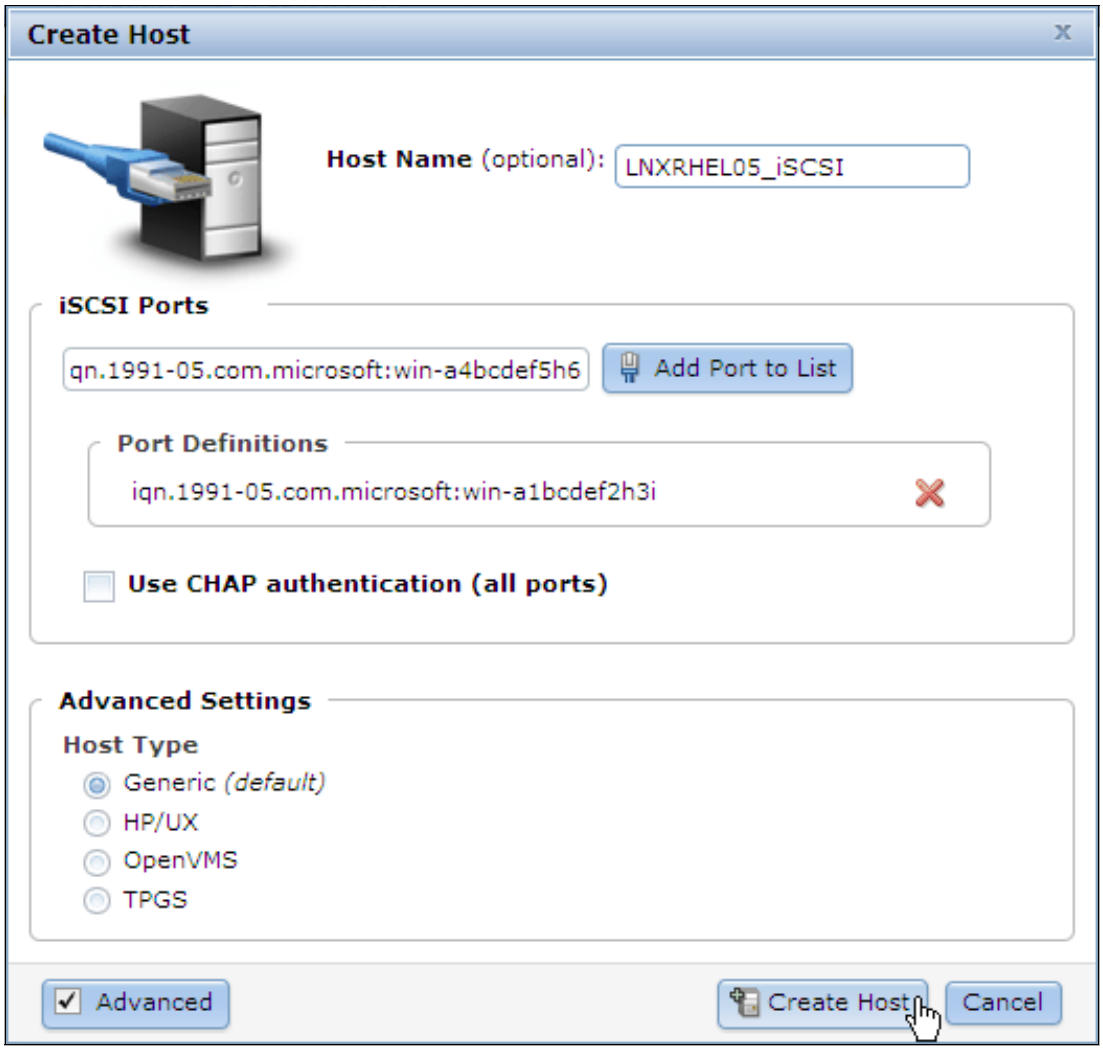

Figure 3-82 The iSCSI Create Host panel with Advanced Settings selected

Modify Host Mappings panel

Host mapping is a method of controlling which hosts can access the volumes. A host can access only volumes that are mapped to it. By using the Modify Host Mappings panel, you can configure the host mappings. Right-click a host that is listed in the Hosts panel and select **Modify Mappings** to open the modify hosts mappings panel.

Figure 3-83 shows the method for opening the Modify Host Mappings panel.

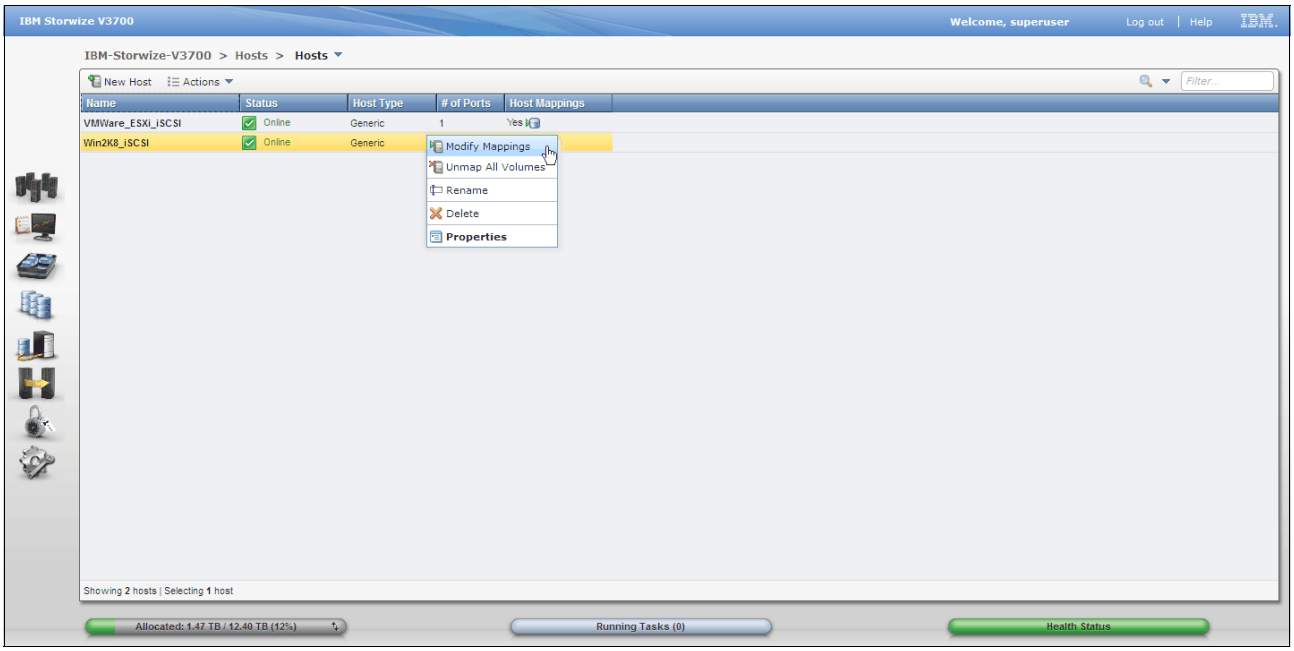

Figure 3-83 Method for opening the Modify Hosts Mappings panel

Mapping volumes

By using the Modify Host Mappings panel, you can map unmapped volumes to the selected host. Verify the correct host is selected in the upper left corner of the panel.

Complete the following steps to map a volume:

1. Select the volume from the Unmapped Volumes section on the left side of the panel.
2. Click the green > to move the volume to Volumes Mapped to the Host section on the right side of the panel.

   The volume is highlighted in yellow, is not mapped to the host and, therefore, is not accessible by the host.
3. Click **Edit SCSI ID**, if required.

   The SCSI ID controls the order in which the volumes are presented to the host. Some host software requires that the SCSI IDs to be in sequence.
4. After the volumes that are highlighted in yellow are as required, click **Apply** to complete the mapping operation. The mapped volumes are not highlighted in yellow and are accessible by the host. Continue to select other hosts and perform other host mappings as required.
5. After the mapping operations are complete, click **Cancel** to close the Modify Host Mappings panel. **Map Volumes** also can be used to close the panel, but it is best to use this option if only a single mapping operation is required because the apply is performed and then the panel is closed. Figure 3-84 shows the Modify Host Mappings panel.

Unmapping volumes

By using the Modify Host Mappings panel, you can unmap volumes from hosts. Stop I/O to the volume and verify that the correct host is selected in the upper left corner of the panel.

Complete the following steps to unmap a volume:

1. Select the volume from the Volumes Mapped to the Host section on the right side of the panel.
2. Click the green < to move the volume back to the Unmapped Volumes section on the left side
3. Click **Apply** to complete the unmapping operation. After all of the unmapping operations are complete, click **Cancel** to close the Modify Host Mappings panel. Figure 3-84 shows the Modify Host Mappings panel.

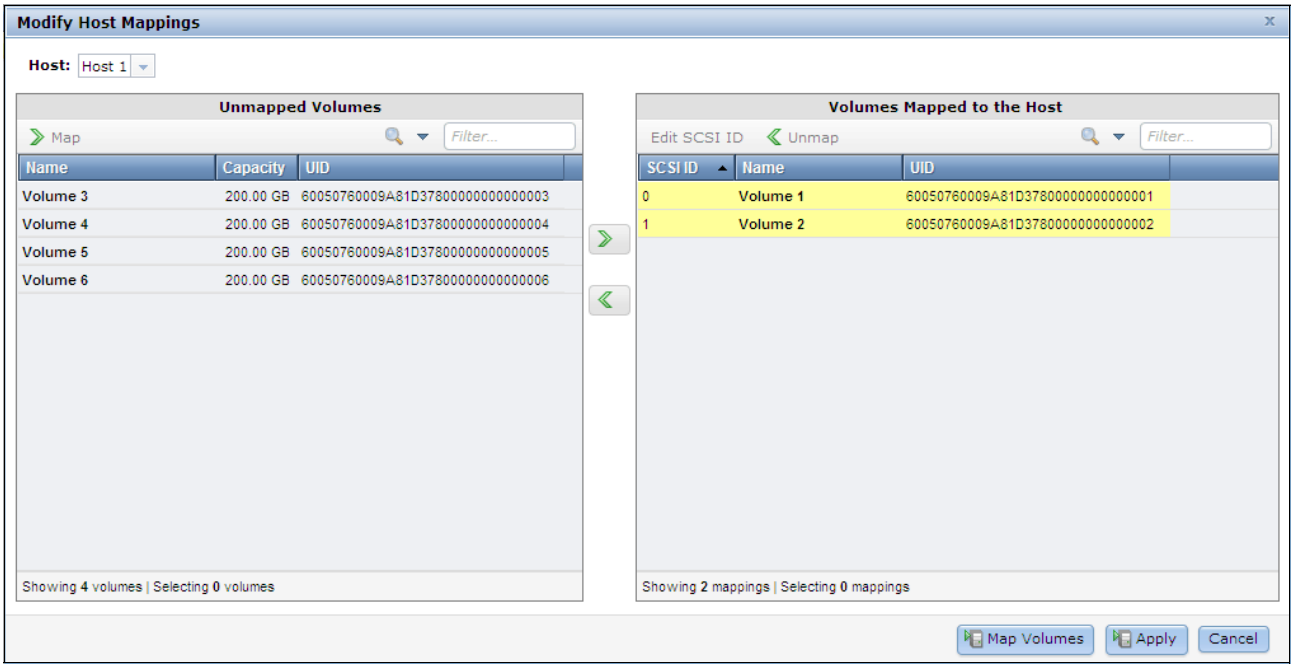

Figure 3-84 Modify host mappings panel

Ports by Host panel

Select **Ports by Host** in the Hosts menu to open the panel, as shown in Figure 3-85. The panel shows the address, status, and type of ports that is assigned to the host definition. Actions such as map, unmap, and port deletion can be performed by using this panel.

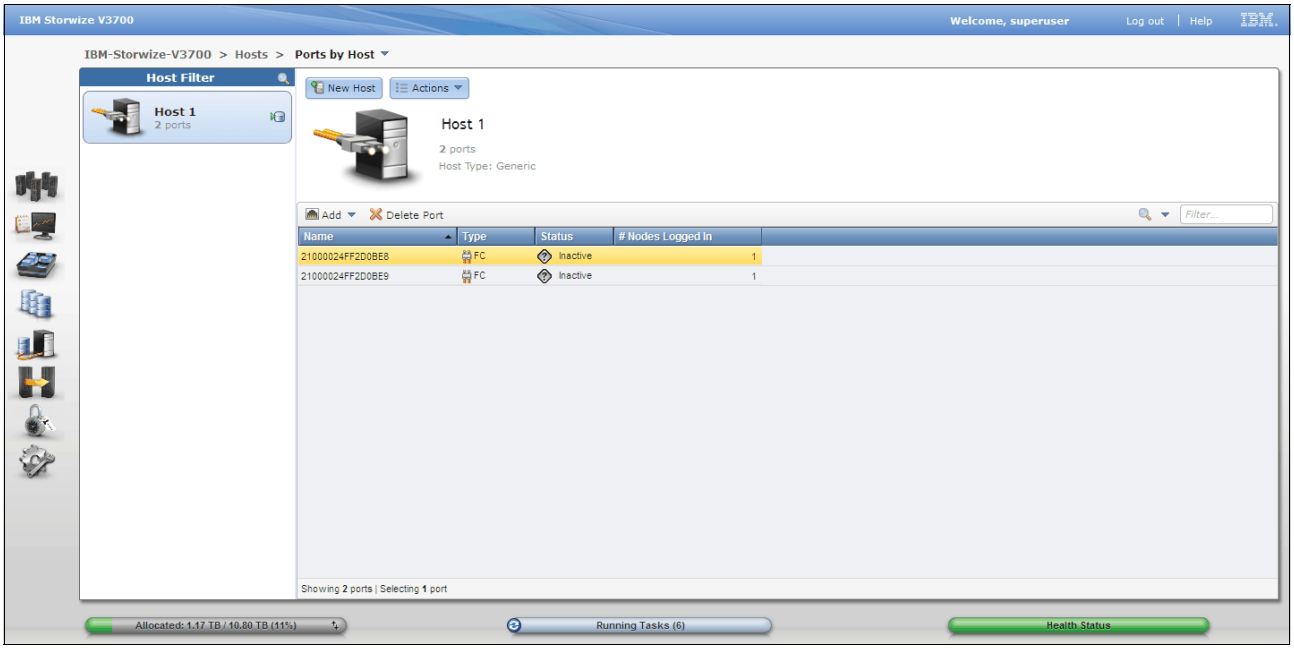

Figure 3-85 Ports by host panel

Host Mappings panel

Select **Host Mappings** in the Hosts menu to open the panel, as shown in Figure 3-86. 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.

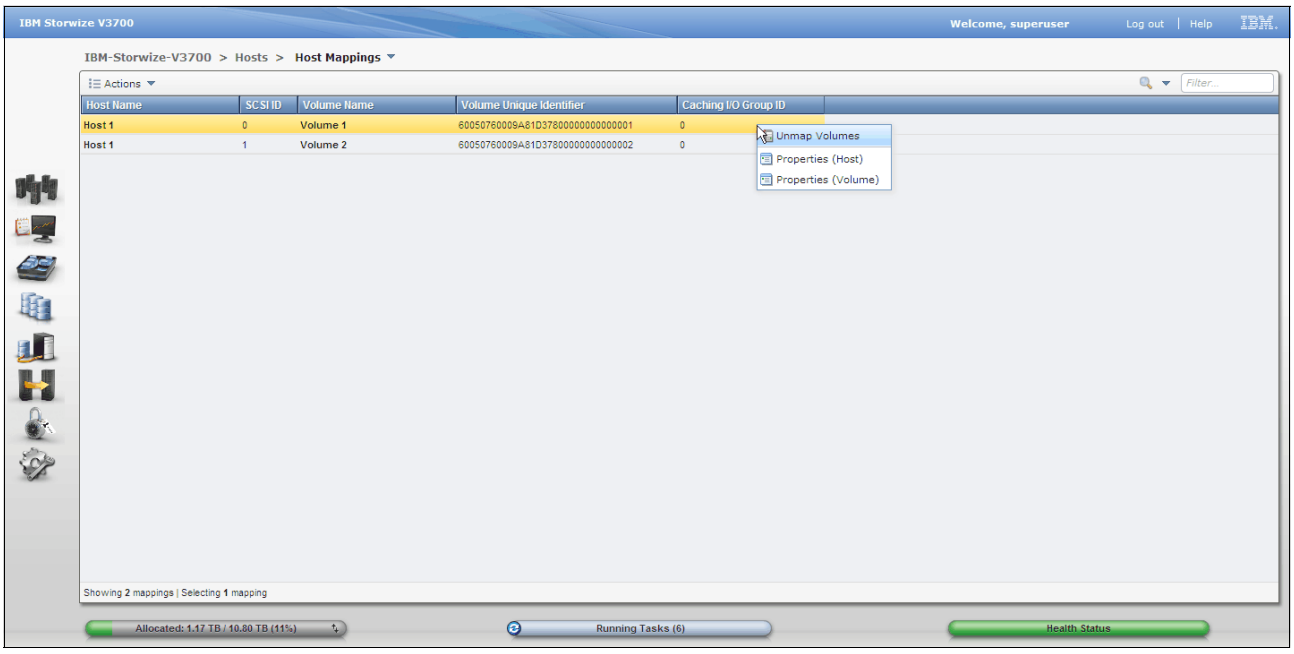

Figure 3-86 Host Mappings panel

Volumes by Host panel

This method is an alternative method to access the Volumes by Host panel. For more information, see “Volumes by Host panel” on page 103.

3.4.6 Copy Services menu

By using the Copy Services menu, you can access the FlashCopy, Consistency Groups, and FlashCopy Mappings functions. Figure 3-87 shows the Copy Services menu.

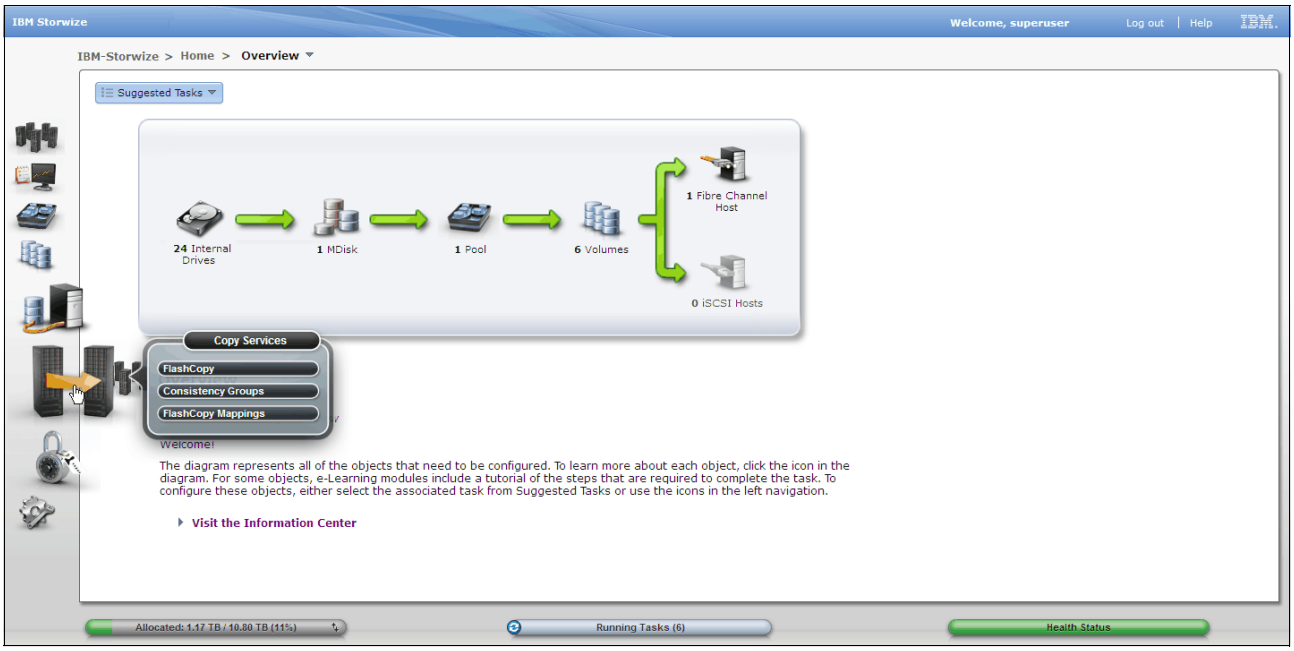

Figure 3-87 Copy Services menu

FlashCopy panel

Select **FlashCopy** in the Copy Services menu to open the panel, as shown in Figure 3-88. The FlashCopy panel displays all volumes in the system. FlashCopy point-in-time is a copy function that copies the contents of a source volume to a target volume.

FlashCopy actions

FlashCopy actions such as, snapshot, clone, backup, target assignment, and deletion can be performed from this panel, as show in Figure 3-88.

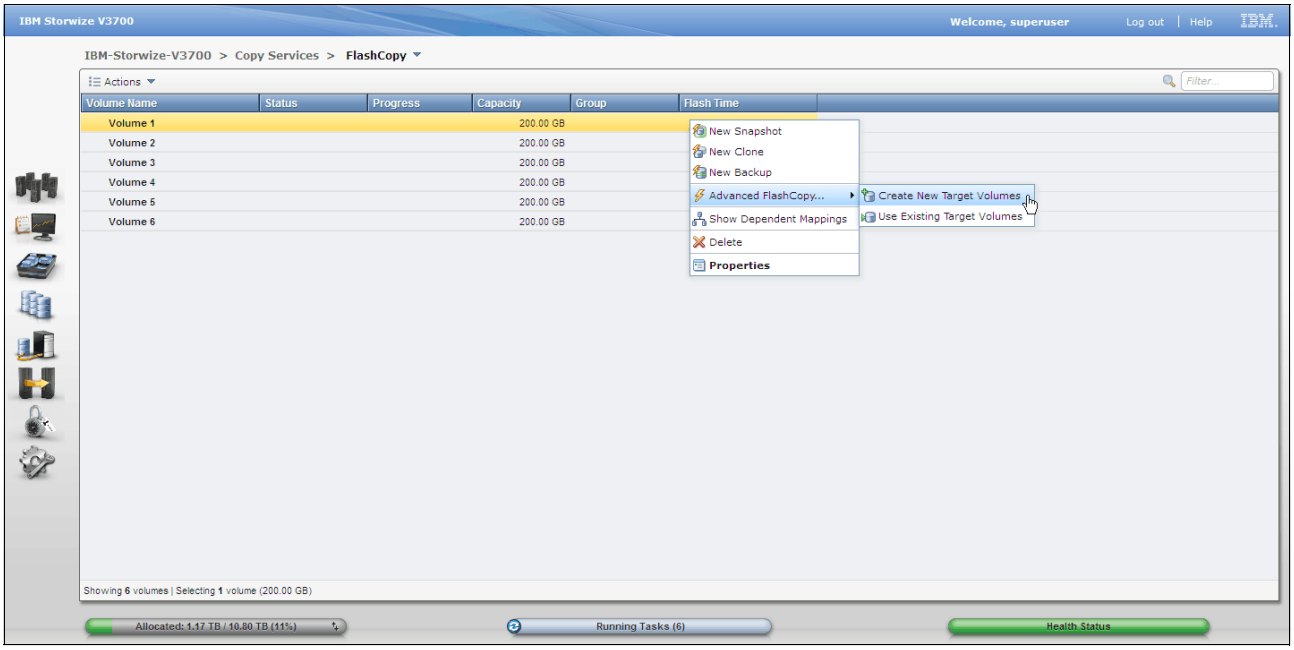

Figure 3-88 FlashCopy panel

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 groups target volumes are consistent to the same point in time and remove several FlashCopy mapping administration tasks.

The Consistency Group panel, as shown in Figure 3-89, 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. The New FlashCopy Mapping option also can be selected from this panel. For more information, see “FlashCopy Mappings panel” on page 115.

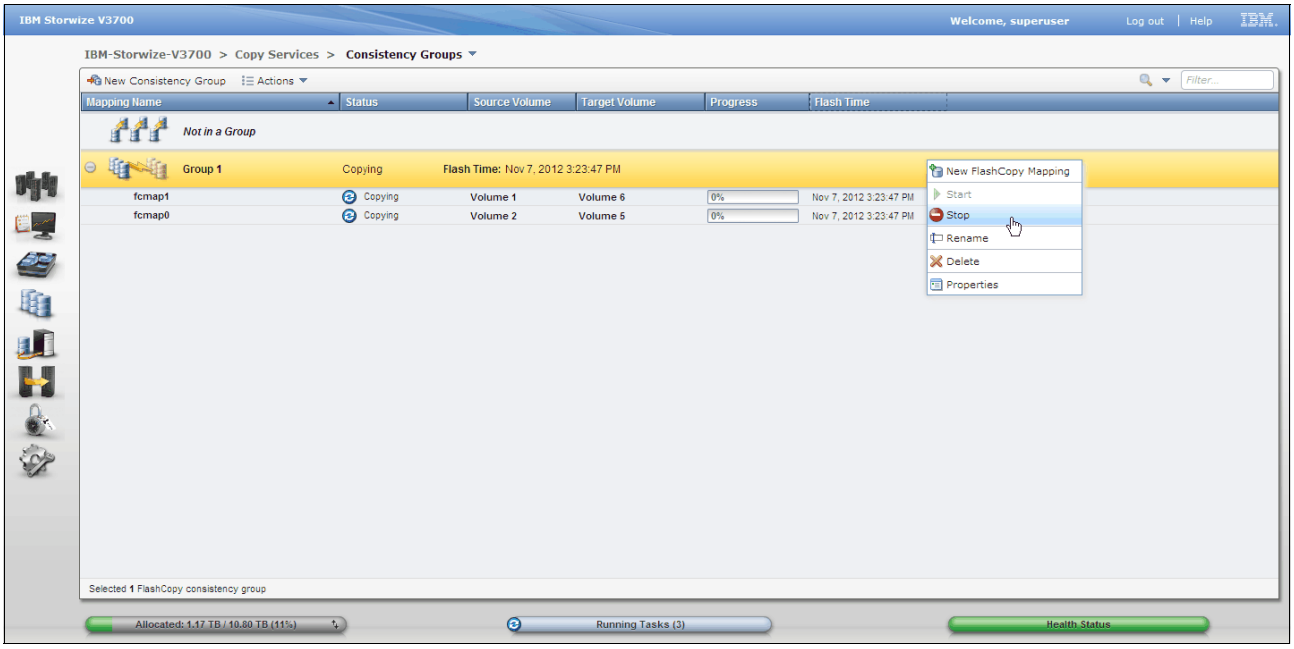

Figure 3-89 Consistency groups panel

FlashCopy Mappings panel

Select **FlashCopy Mappings** in the Copy Services menu to open the panel, as shown in Figure 3-90. FlashCopy mappings define the relationship between source volumes and target volumes. The FlashCopy Mappings panel shows information about each mapping such as: status, progress, source and target volumes, and flash time. Select **New FlashCopy Mapping** to configure a new mapping.

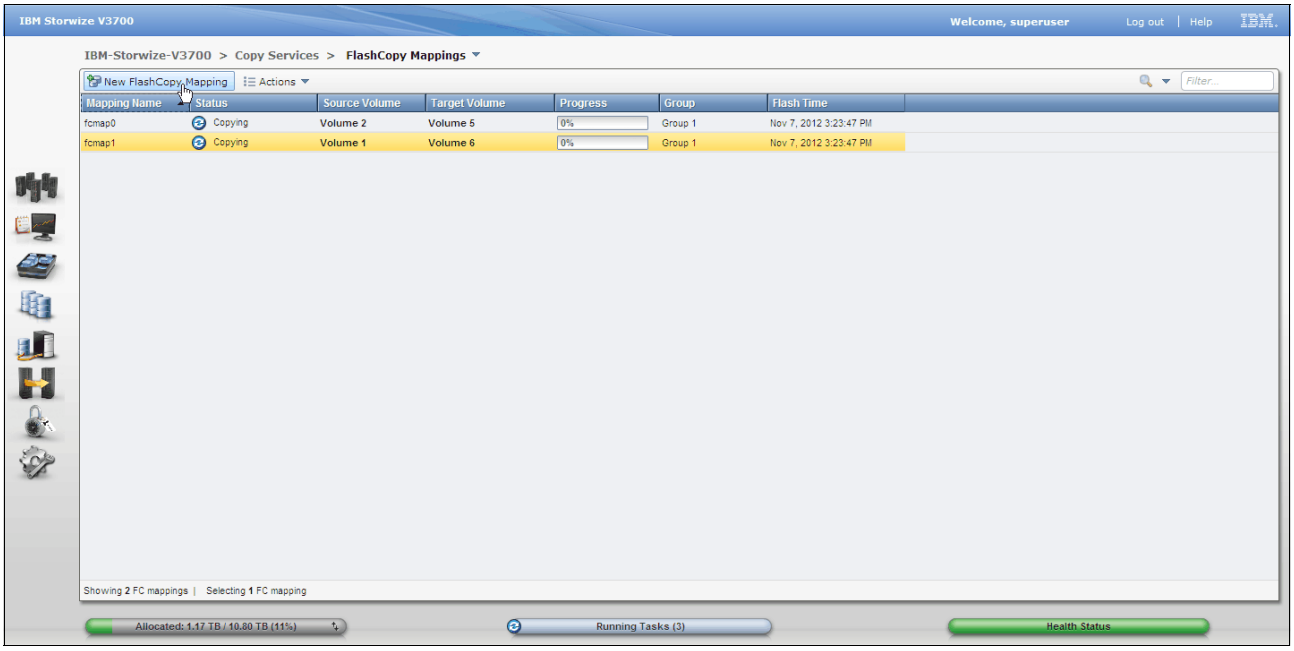

Figure 3-90 FlashCopy mappings panel

Create FlashCopy mapping

Complete the following steps to create a FlashCopy mapping:

1. Select the source and target volumes from the respective drop-down menus and click **Add**. The mapping definition is added to the list.

   To remove a mapping definition from the list, click the red **X**. Click **Next** to continue. Figure 3-91 shows step 1 of the New FlashCopy Mapping wizard.

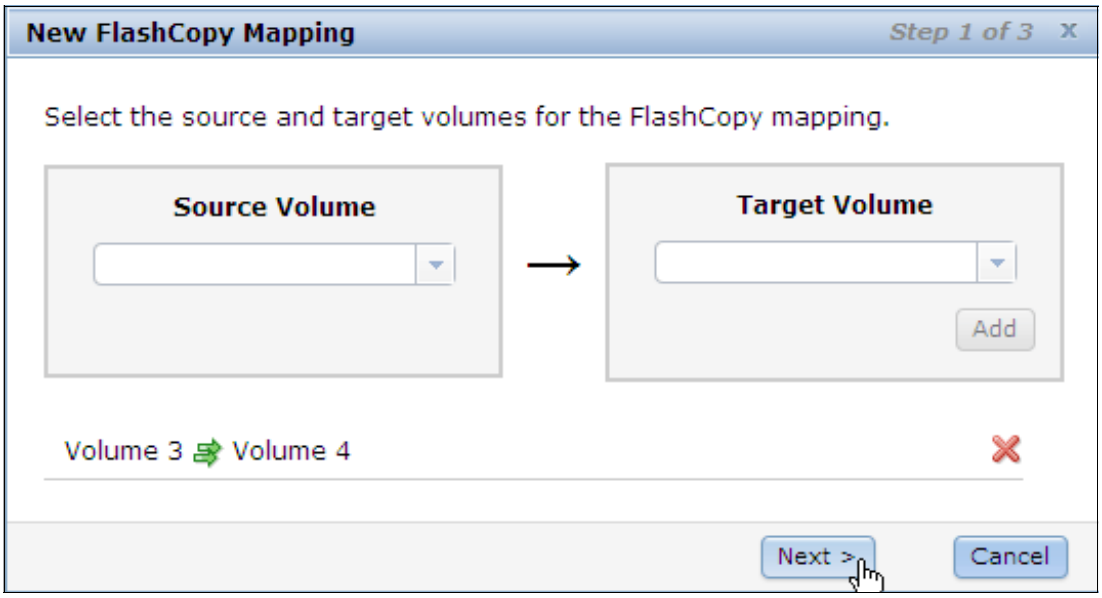

Figure 3-91 Step 1 of the New FlashCopy Mapping wizard

2. Select a FlashCopy preset and modify the advance settings as required. Select **Next** to continue, as shown in Figure 3-92.

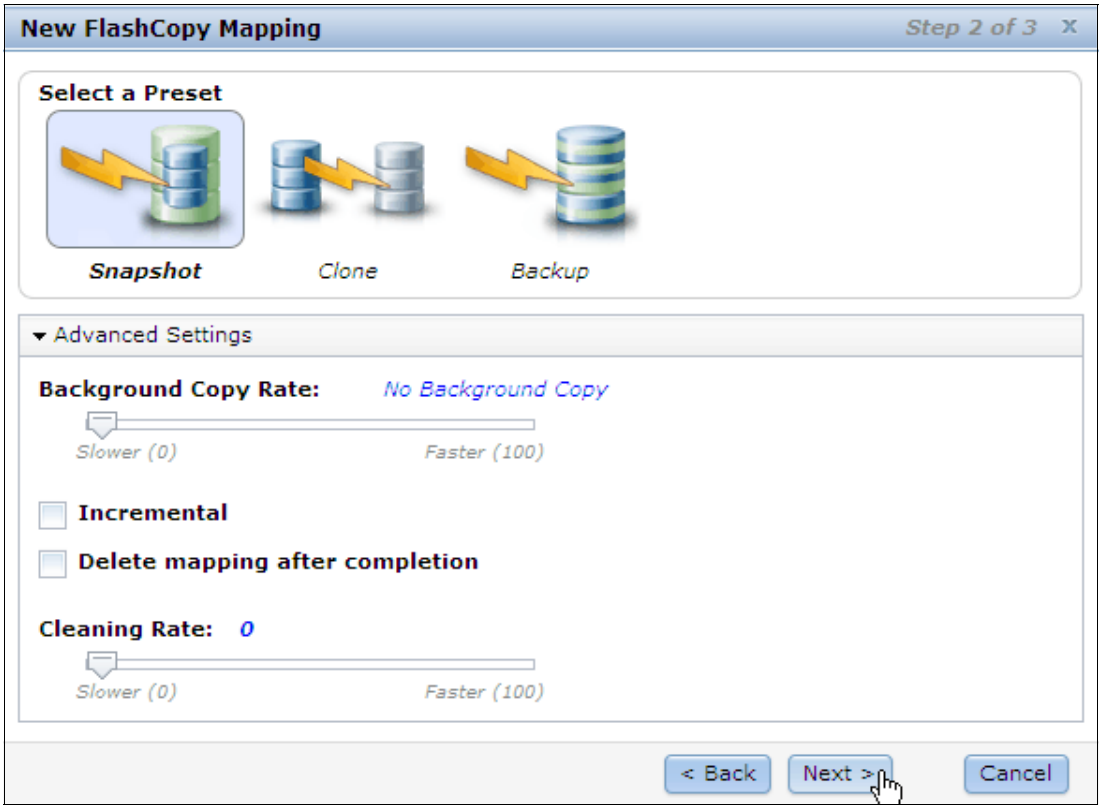

Figure 3-92 Step 2 of the new FlashCopy mapping wizard

More information: For more information about FlashCopy Mappings and Consistency Groups, see the **Product overview** → **Technical overview** → **Copy Services features** section of the IBM Storwize V3500 Information Center at this website:

http://pic.dhe.ibm.com/infocenter/storwize/v3500_ic/index.jsp

3. The new FlashCopy mapping can be added to a consistency group. Choose whether this mapping is grouped and click **Finish**, as shown in Figure 3-93.

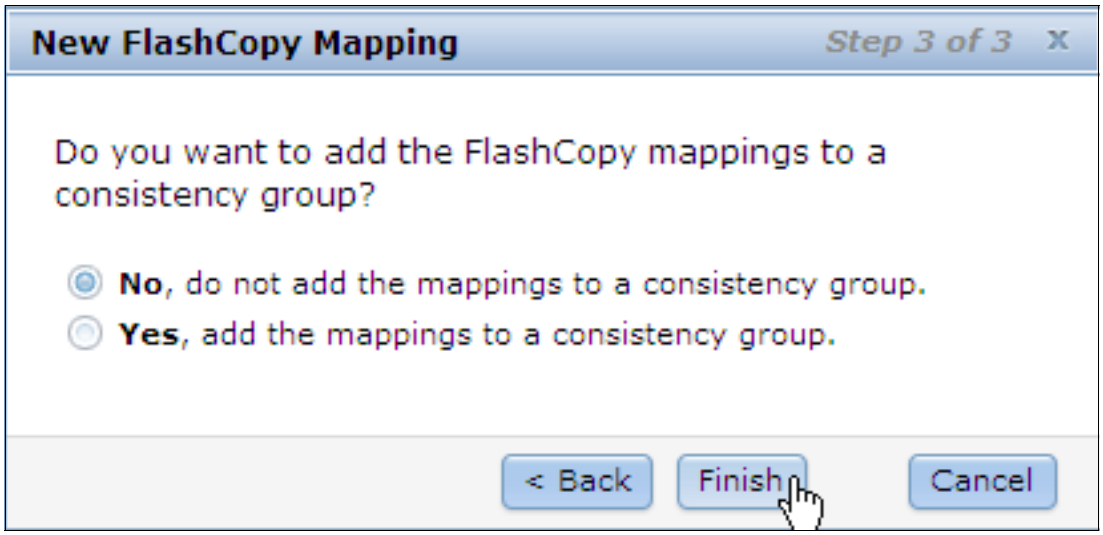

Figure 3-93 Step 3 of the New FlashCopy Mapping wizard

3.4.7 Access menu

As shown in Figure 3-94, the Access menu provides access to the Users and Audit Log functions.

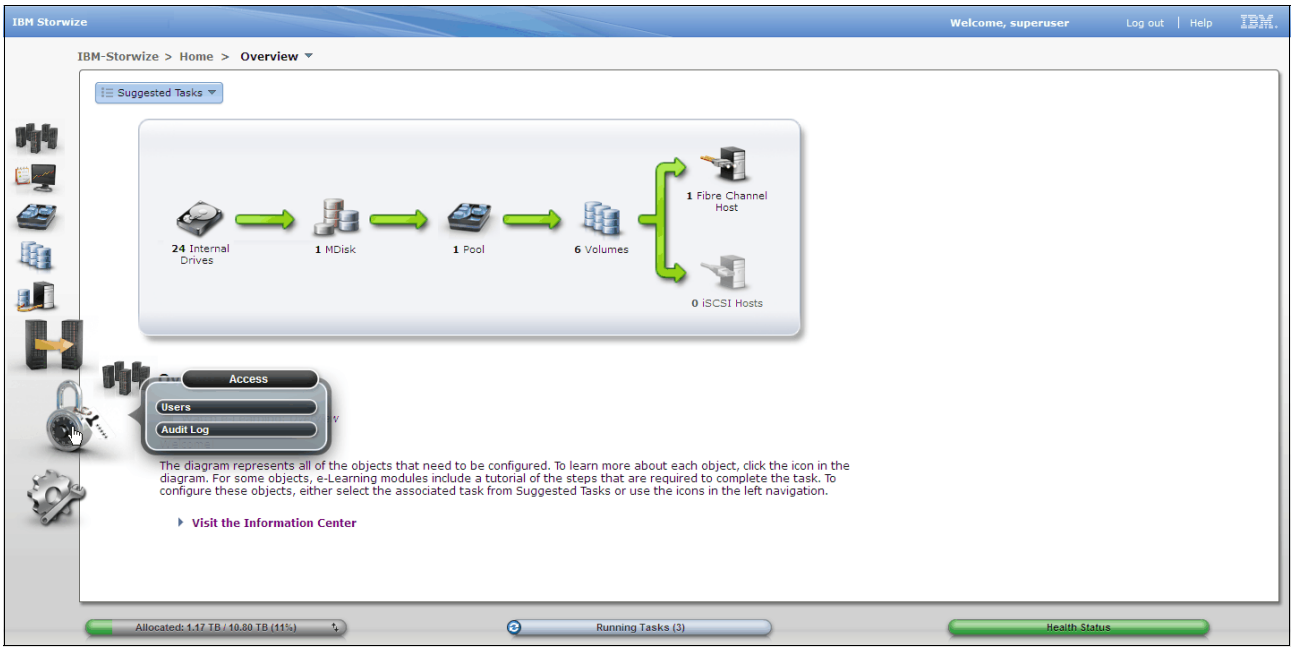

Figure 3-94 Access menu

Users panel

Select **Users** in the Access menu to open the panel. As shown in Figure 3-95, 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 Group panel.

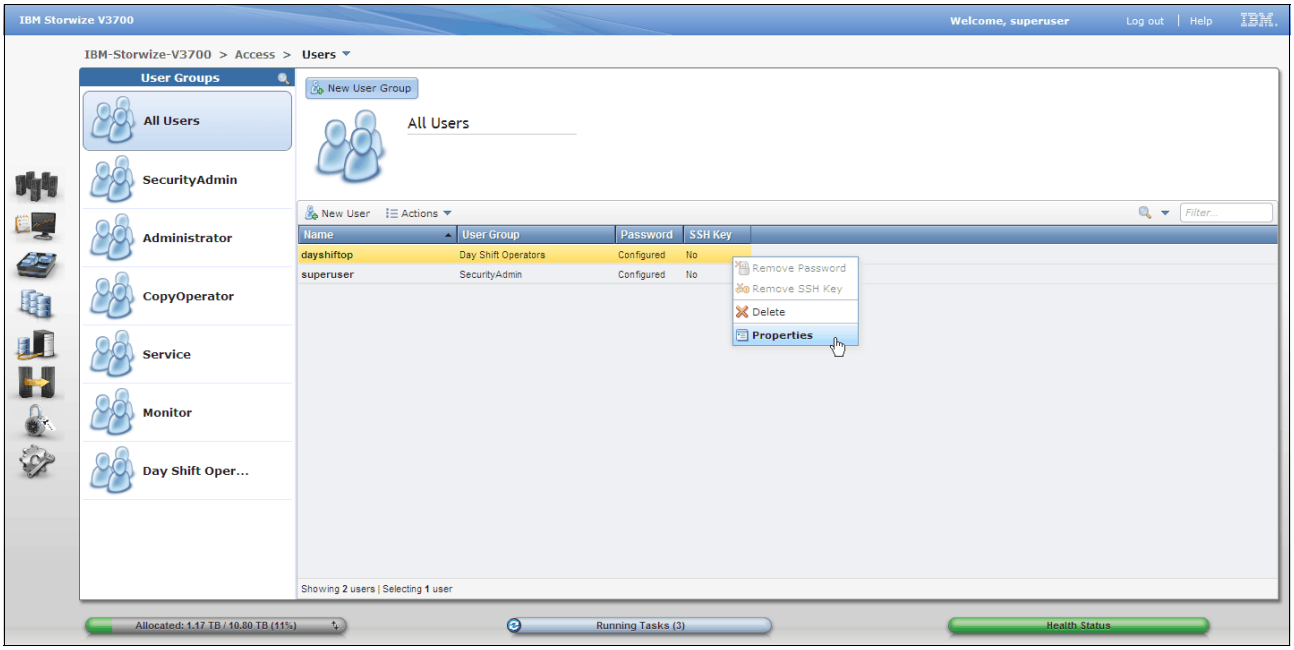

Figure 3-95 Users panel

Creating user group

By using the New User Group panel (as shown in Figure 3-96), you can configure a user group. Enter the group name, select the role, then click **Create** to create a group.

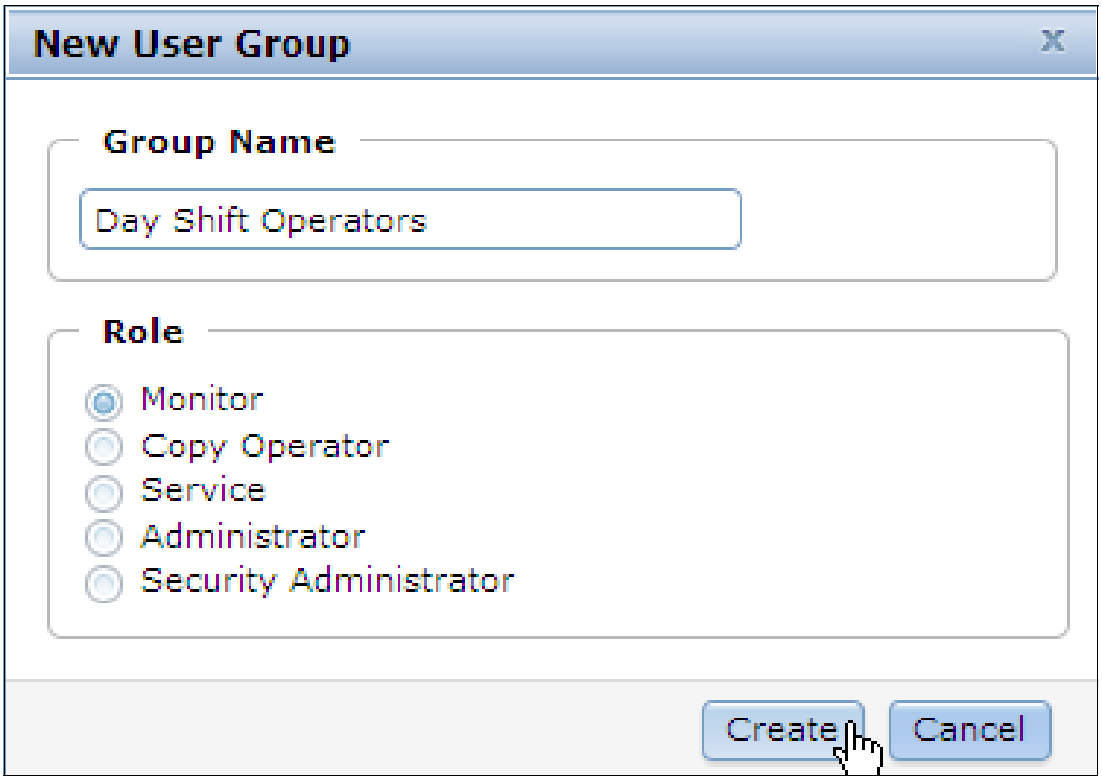

Figure 3-96 New User Group panel

Creating user

Click **New User** to define a user to the system. Figure 3-97 shows the Users panel and the New User option.

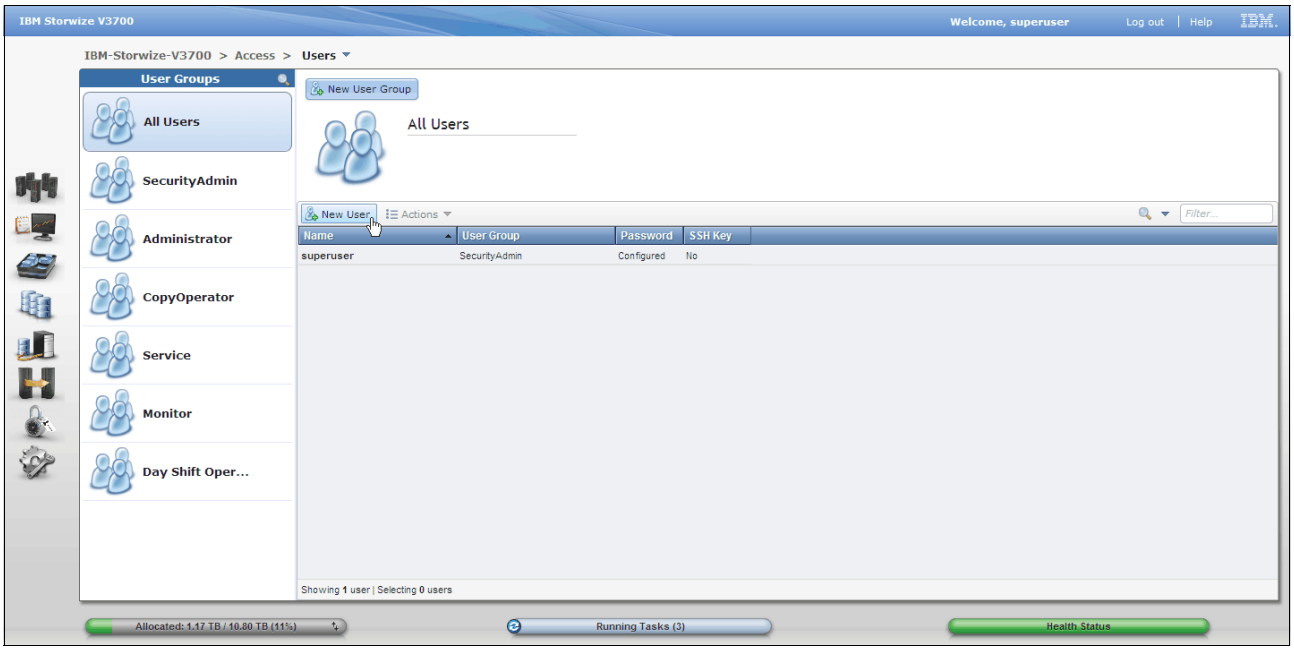

Figure 3-97 Users panel and the New User option

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 Storwize V3500 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 Storwize V3500 by clicking **Settings menu** → **Directory Services panel**.

The SSH configuration can be used to establish a more secure connection to the CLI. It is important to enter the associated SSH private key into terminal session to enable work.

Figure 3-98 shows the New User panel.

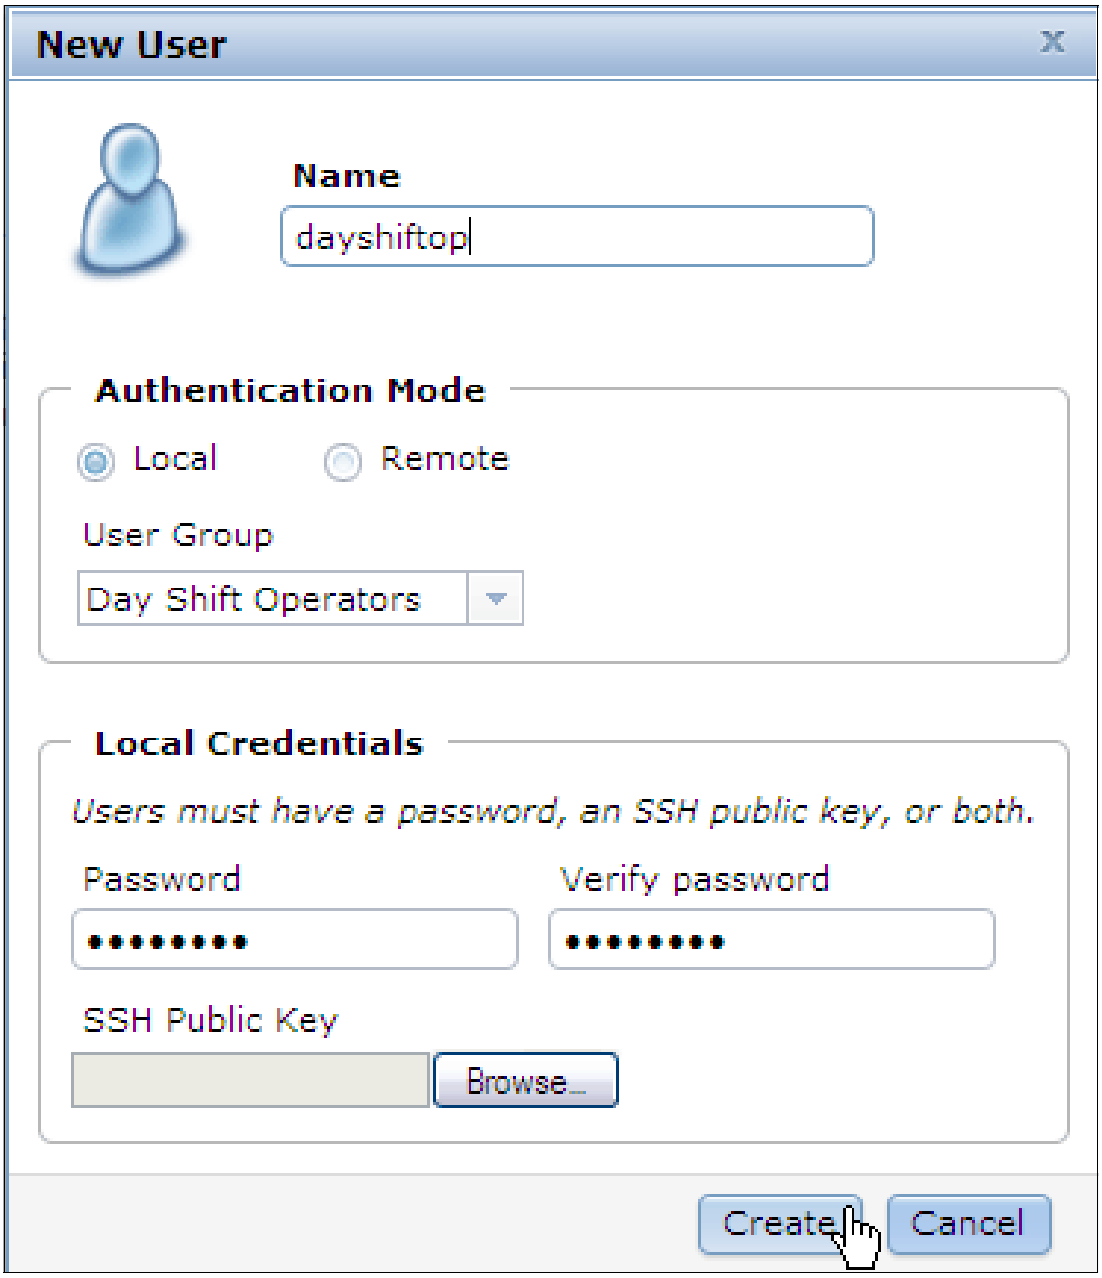

Figure 3-98 New User panel

Audit Log panel

Select **Audit Log** in the Access menu to open the panel, as shown in Figure 3-99. 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 actionable command such as: the user, the time stamp, and any associated command parameters. The log can be filtered to reduce the number of items that are listed.

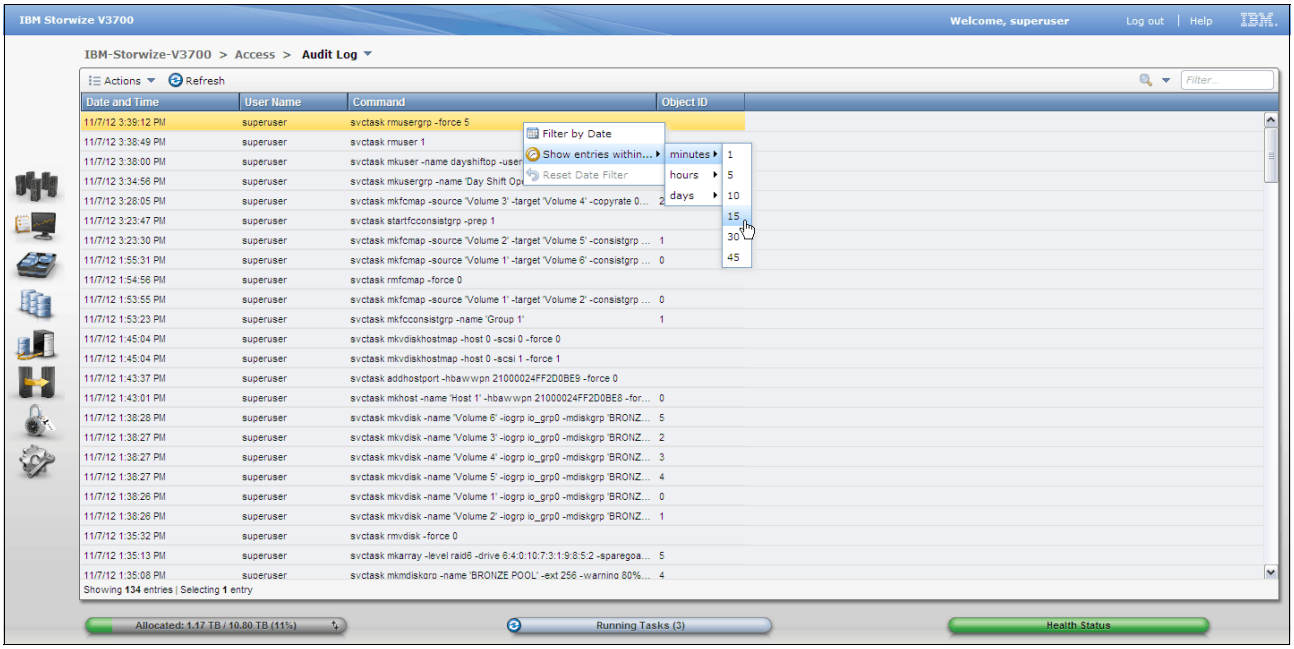

Figure 3-99 Audit Log panel

3.4.8 Settings menu

The Setting menu provides access to the Event Notifications, Directory Services, Network, Support, and General functions, as shown in Figure 3-100.

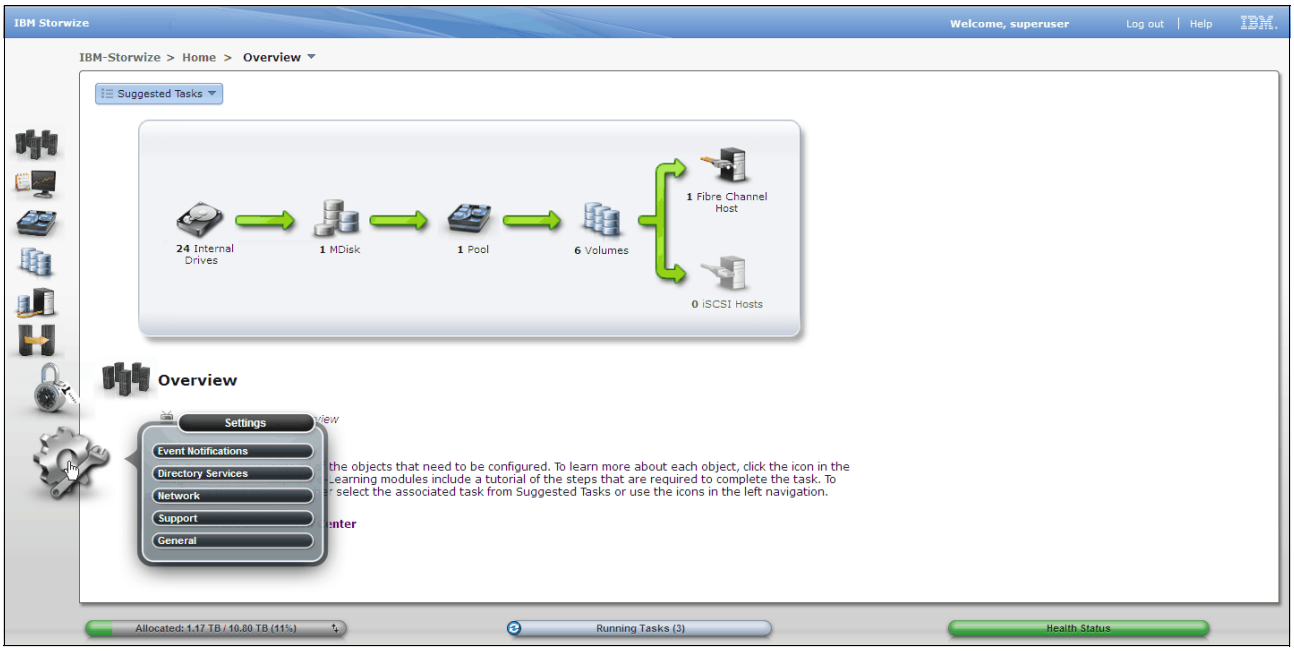

Figure 3-100 Settings menu

Event Notifications panel

Select **Event Notifications** in the Settings menu to open the panel. The Storwize V3500 can use Simple Network Management Protocol (SNMP) traps, syslog messages, emails, and IBM Call Homes to notify parties 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 is enabled by following the Call Home wizard.

Enabling Email Event Notification

Click **Enable Email Event Notification** to open the Call Home wizard. Figure 3-101 shows the Event Notifications Email configuration panel.

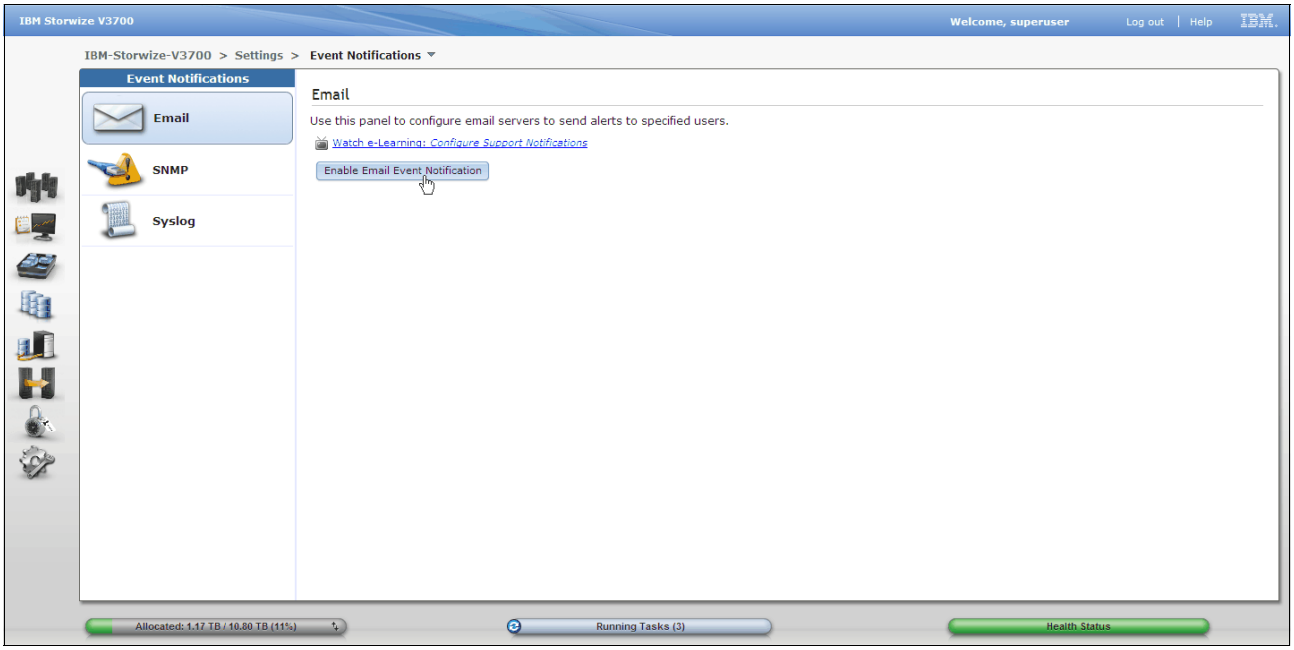

Figure 3-101 Event notification panel

Call Home wizard

The Call Home wizard guides the user through account contact and machine location entry, and email configuration tasks. Figure 3-102 shows the Call Home wizard.

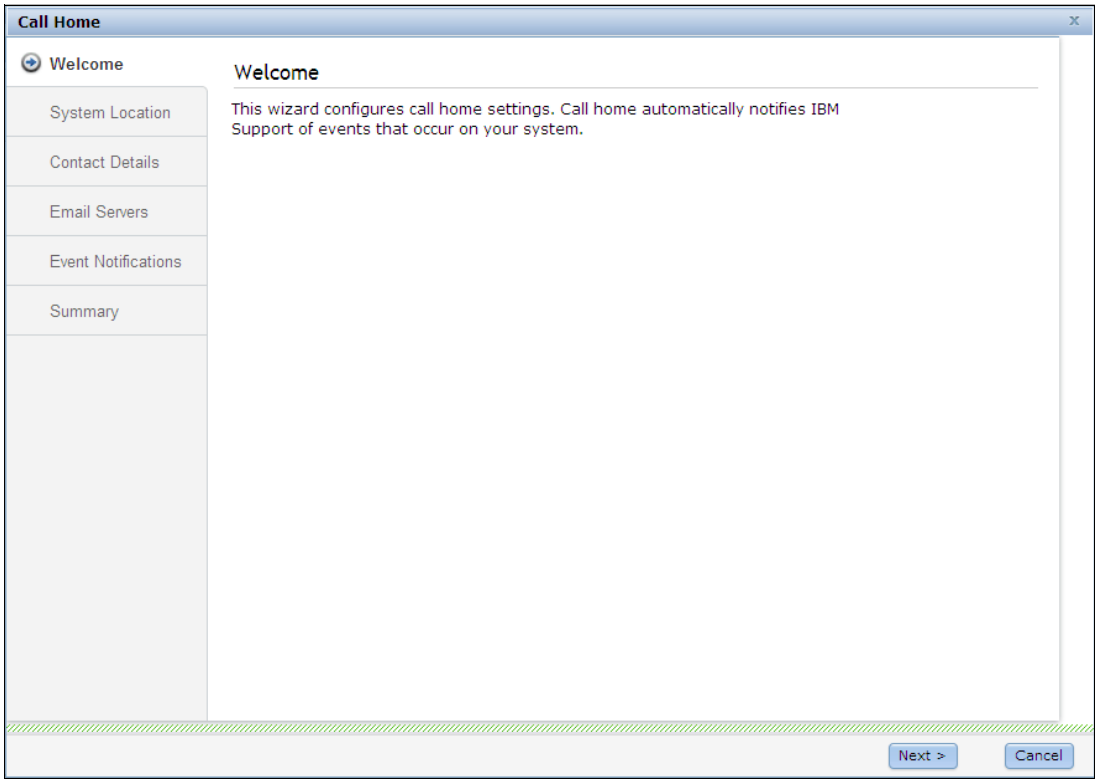

Figure 3-102 Call home wizard

SNMP Event Notification

The Event Notifications panel provides access to SNMP configuration panel. Click **SNMP** to open the panel, then enter the relevant server details. Multiple servers can be configured by clicking +. Figure 3-103 shows the SNMP configuration panel.

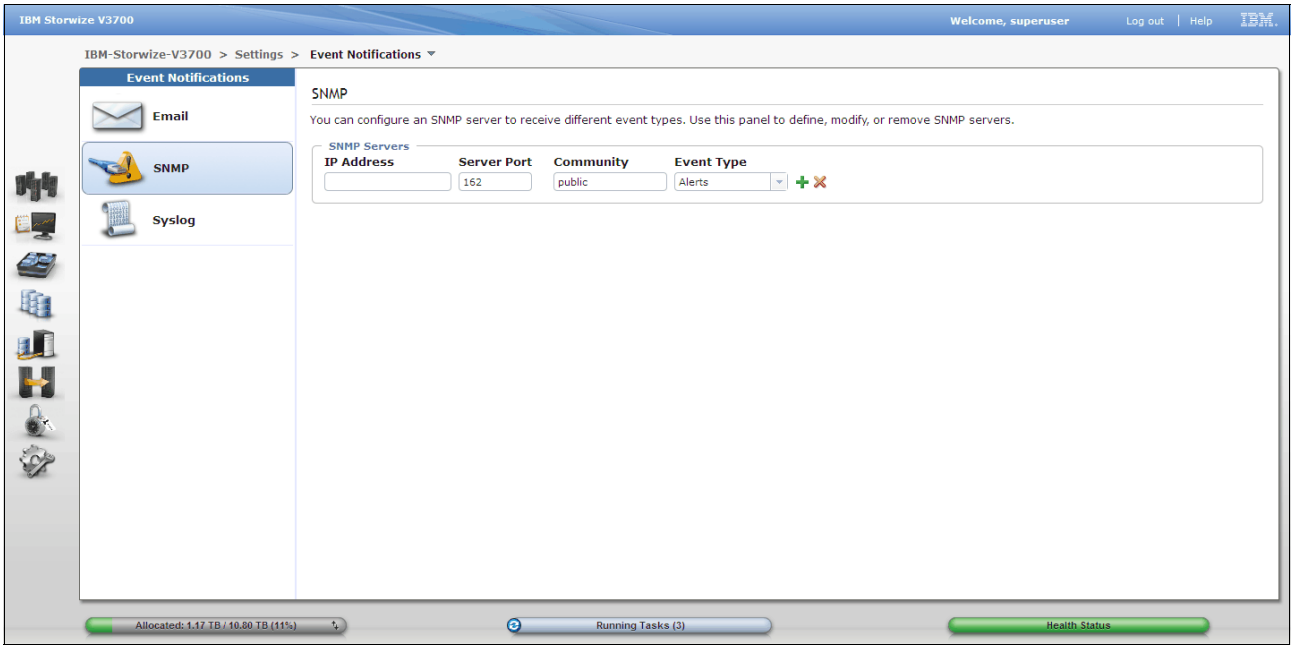

Figure 3-103 SNMP configuration panel

Syslog Event Notification

The Event Notifications panel provides access to Syslog configuration panel. Click **Syslog** to open the panel, then enter the relevant server details. Multiple servers can be configured by clicking +. Figure 3-104 shows the Syslog configuration panel.

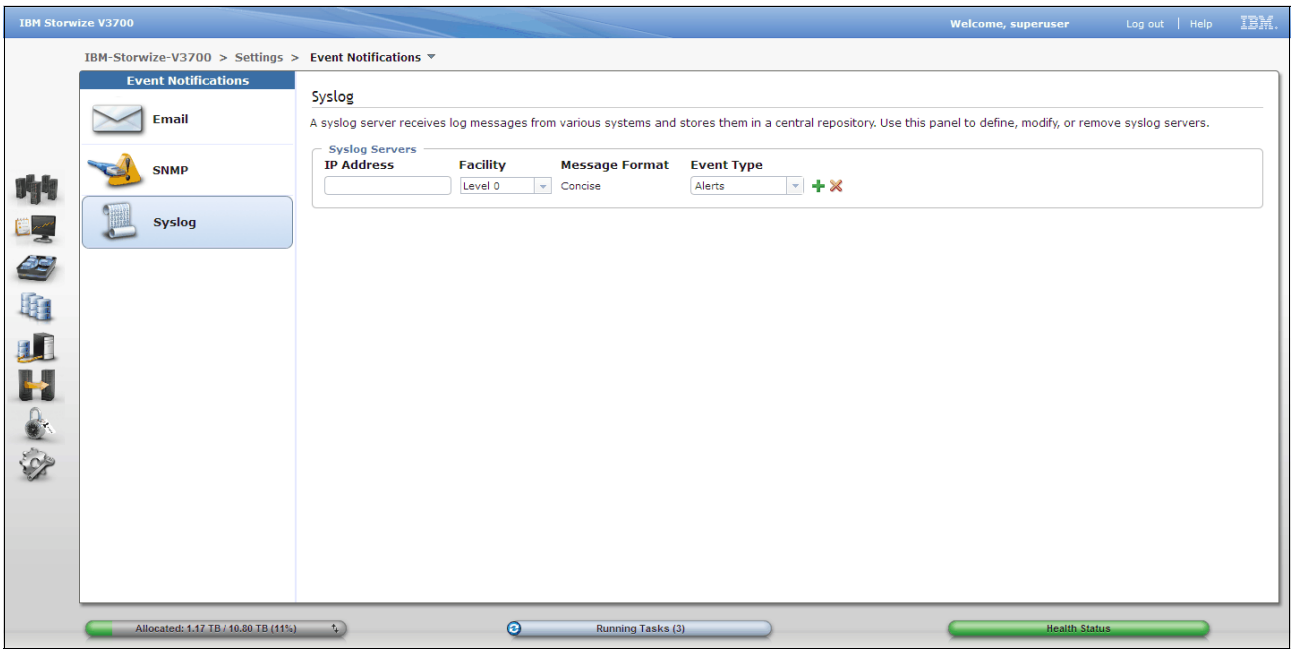

Figure 3-104 Syslog configuration panel

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 Storwize V3500. 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 start the wizard. Figure 3-105 shows the Directory Services panel.

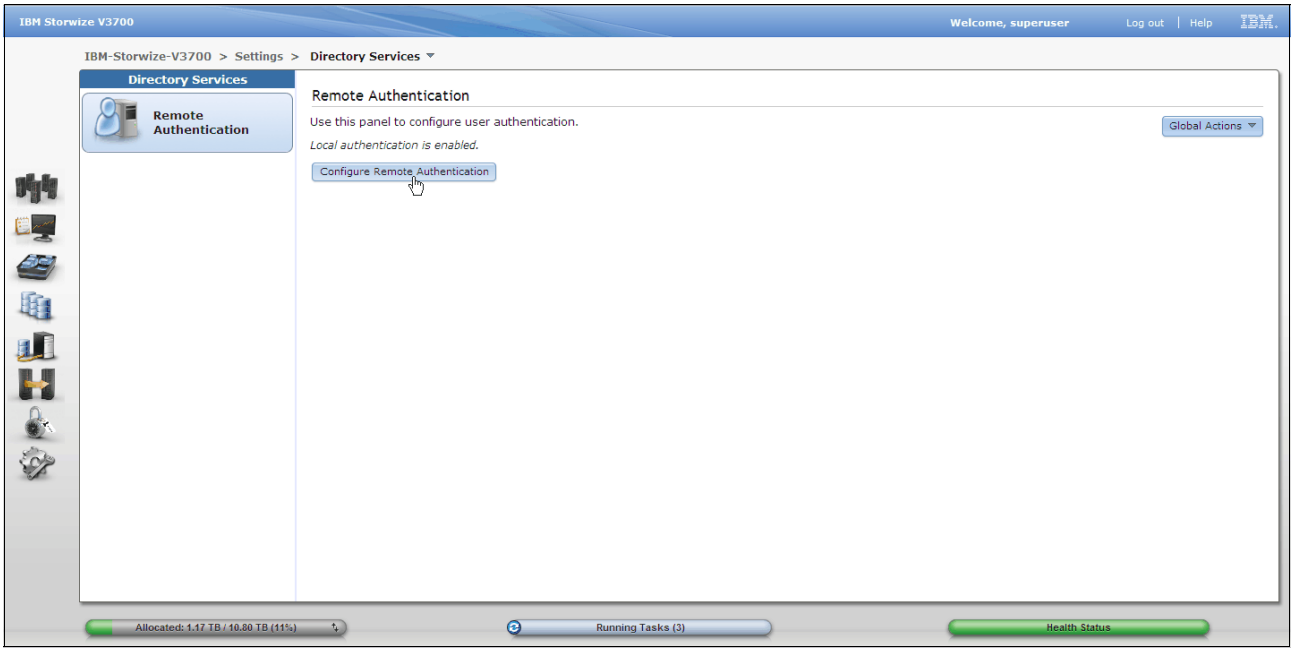

Figure 3-105 Directory services panel

Network panel

Select **Network** in the General menu to open the panel. The Network provides access to the Management IP Addresses, Service IP Addresses, iSCSI, and Fibre Channel configuration panels. Figure 3-106 shows the Network panel.

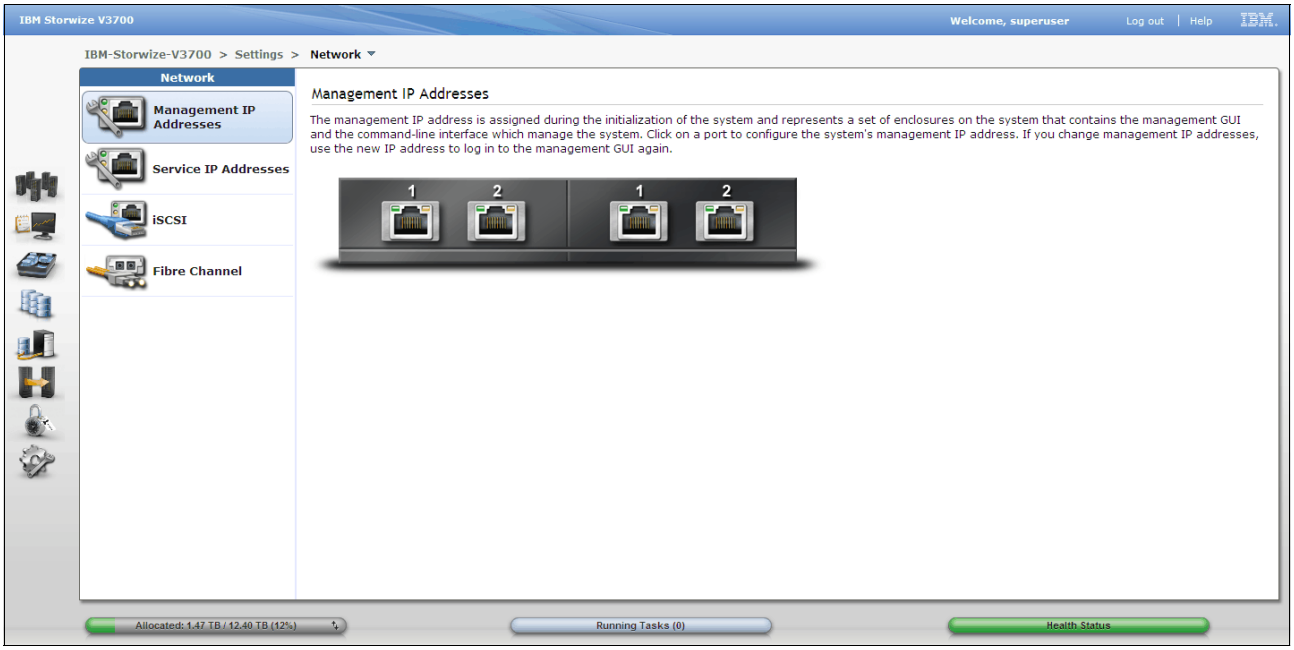

Figure 3-106 Network panel

Management IP Addresses

The Management IP Address is the IP address of the system and was configured during initial setup. It can be an IPv4 address, an 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. Furthermore, 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 wanted Port to configure. The corresponding port on the partner node canister also is highlighted. Figure 3-107 shows Management IP Addresses configuration panel.

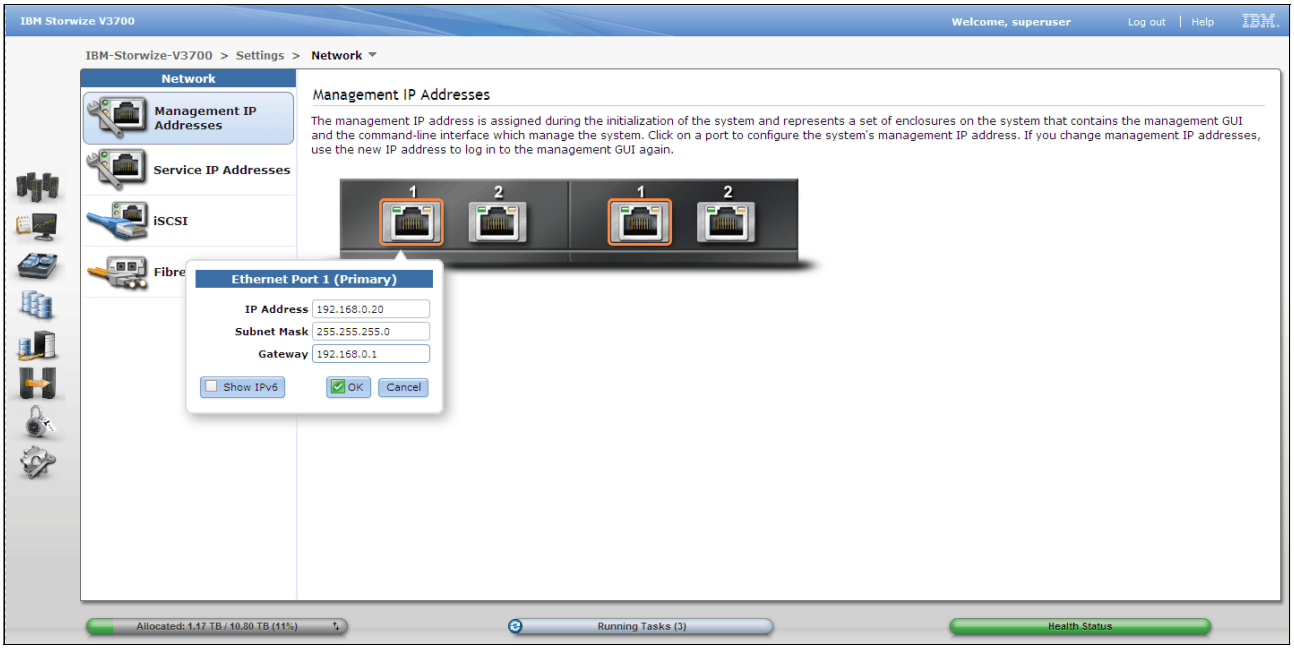

Figure 3-107 Management IP addresses configuration panel

Service IP Addresses

Service IP addresses are used to access the service assistant when a node is not operating correctly as part of the system. It can be an IPv4 address, an 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 node canister to configure. Figure 3-108 shows the Service IP addresses configuration panel.

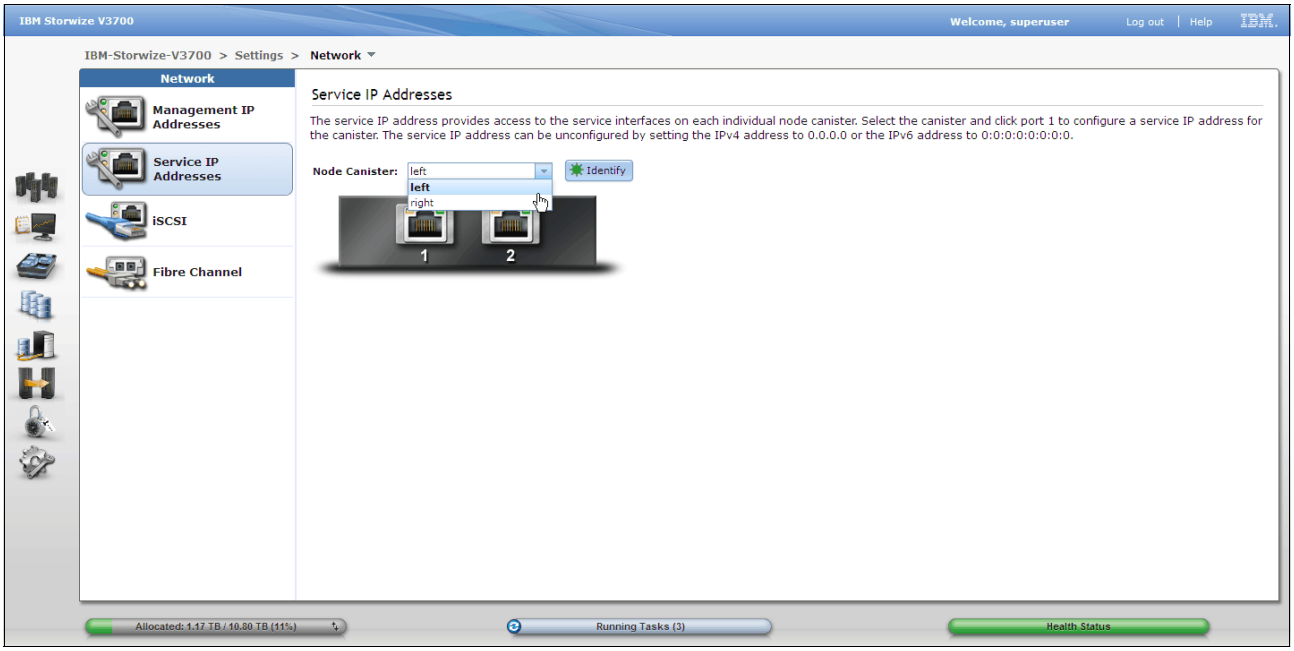

Figure 3-108 Service IP addresses configuration panel

iSCSI connectivity

The Storwize V3500 supports iSCSI connections for hosts. Click **iSCSI** and select the node canister to configure the iSCSI IP addresses. Figure 3-109 shows the iSCSI configuration panel.

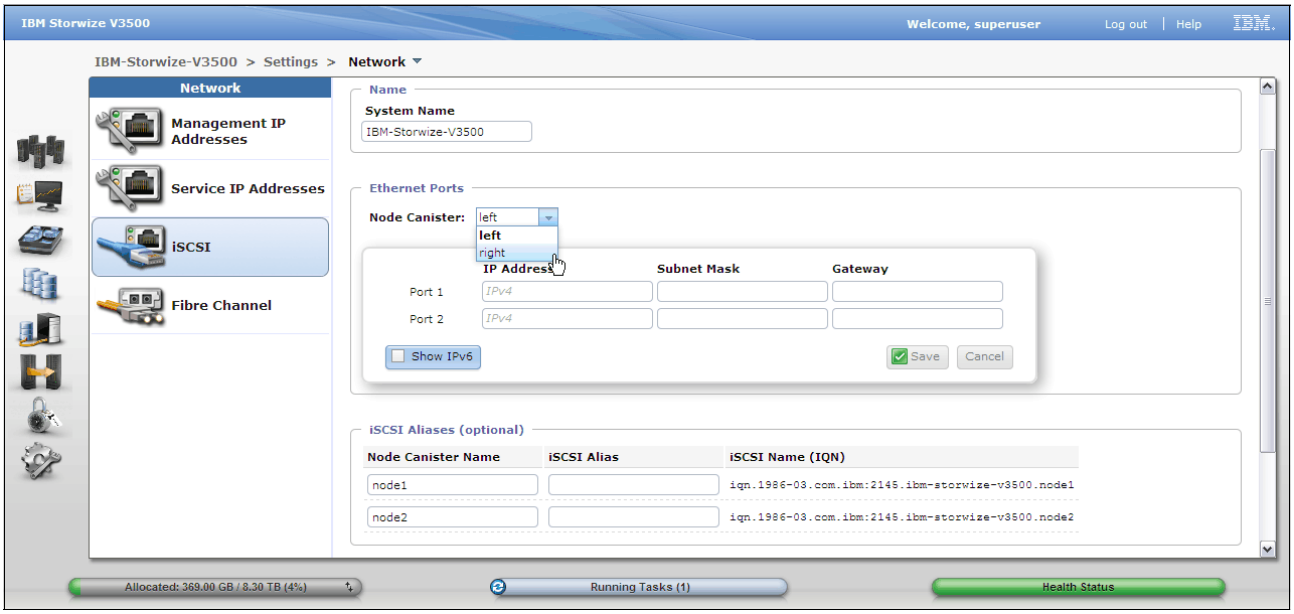

Figure 3-109 The iSCSI configuration panel

Fibre Channel connectivity

The Fibre Channel panel displays established Fibre Channel connections between the Storwize V3500 node canisters and other storage systems and hosts. Click **Fibre Channel** and select the wanted view from the **View connectivity for:** drop-down menu. Figure 3-110 shows the Fibre Channel panel.

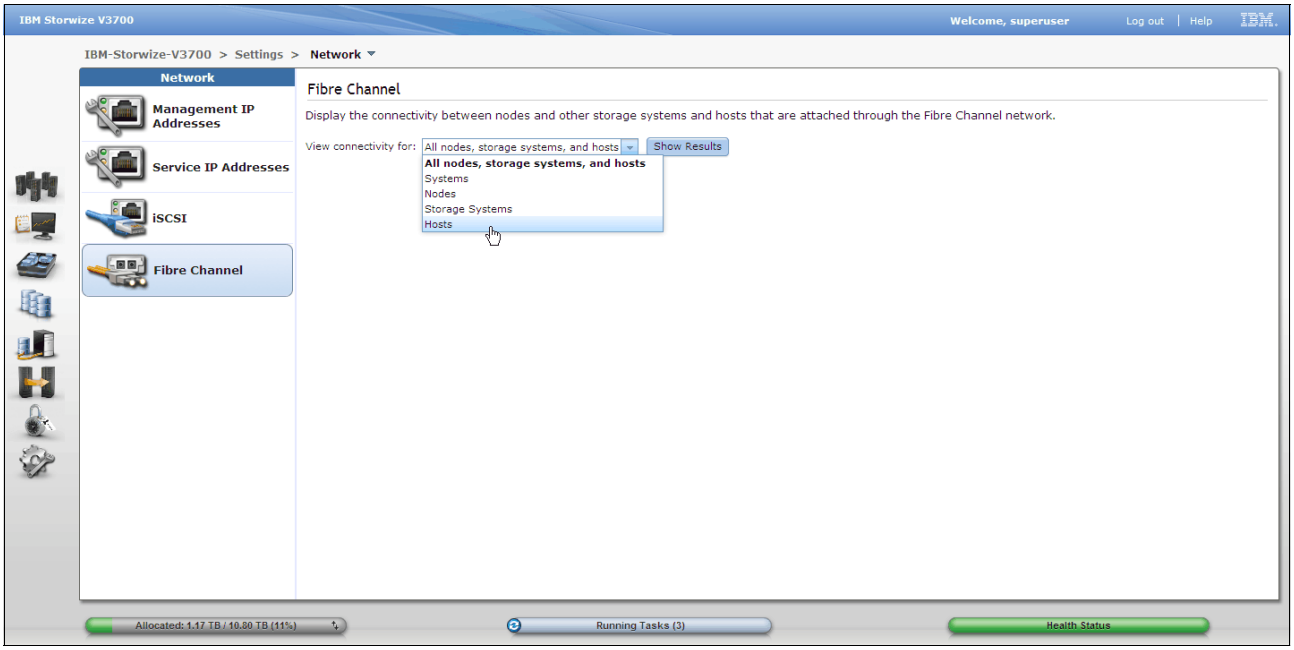

Figure 3-110 Fibre Channel panel

Figure 3-111 shows the results of the selected Fibre Channel connectivity view.

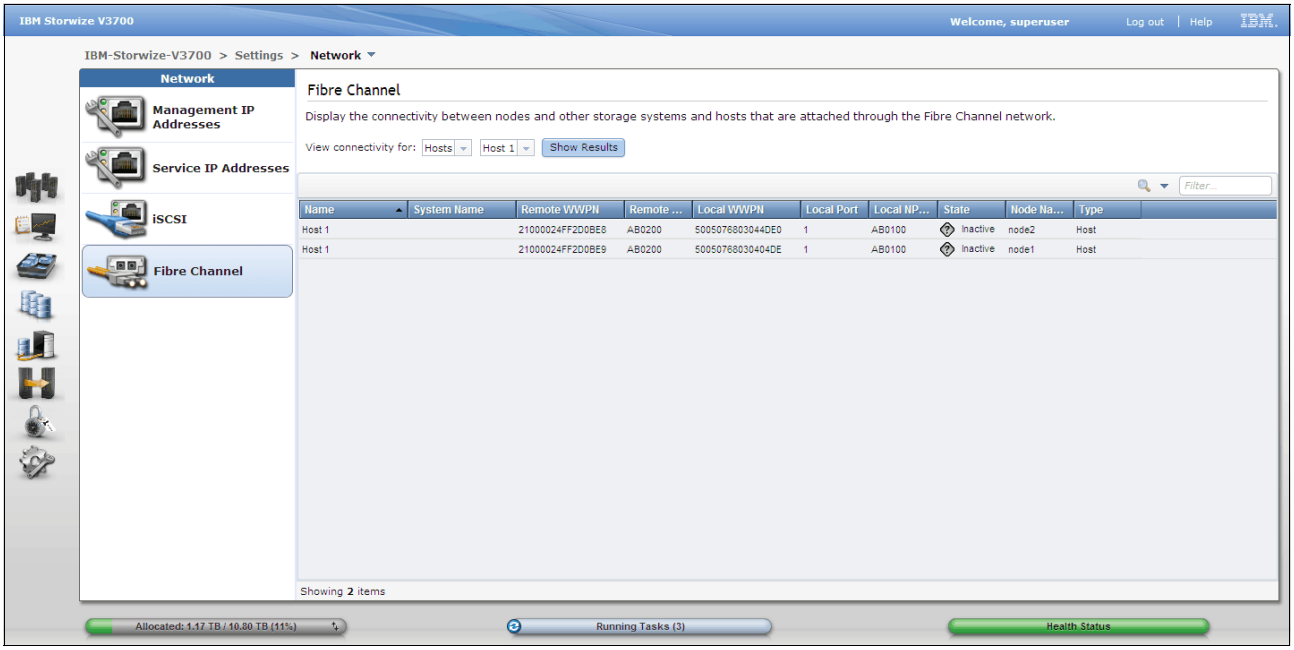

Figure 3-111 Results of the selected Fibre Channel connectivity view

Support panel

Select **Support** in the Settings menu to open the panel. The Support 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 3-112 shows the Support panel.

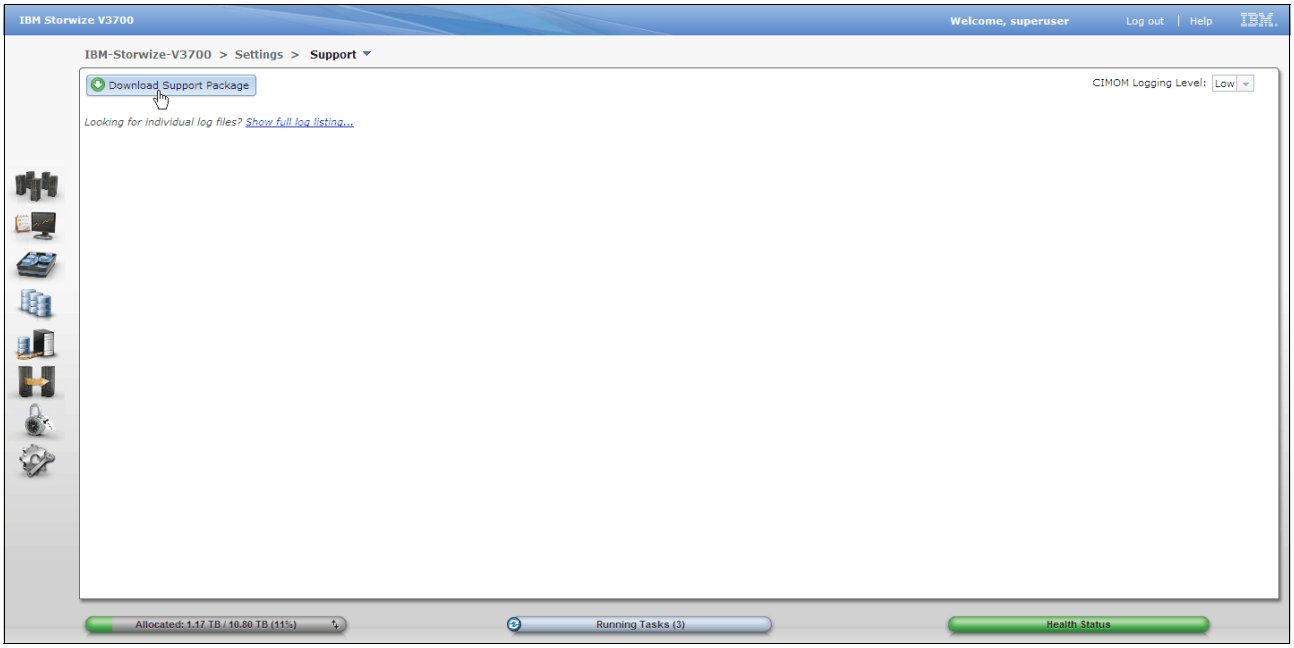

Figure 3-112 Support panel

Downloading the Support Package wizard

The Download Support Package wizard provides various package types. IBM support provides direction on package type selection as required. To download the package, select type and click **Download**. The output file can be saved to the users workstation. Figure 3-113 shows the Download Support Package wizard.

Figure 3-113 Download Support Package wizard

Show full log listing

The Support panel also provides access to the files currently residing on the node canisters, as shown in see Figure 3-112 on page 133. 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 and select **Download**. To change to the file listing to show the files that are residing on a partner node canister, select the node canister from the drop-down menu that is next to the panel filter. Figure 3-114 shows the full log listing.

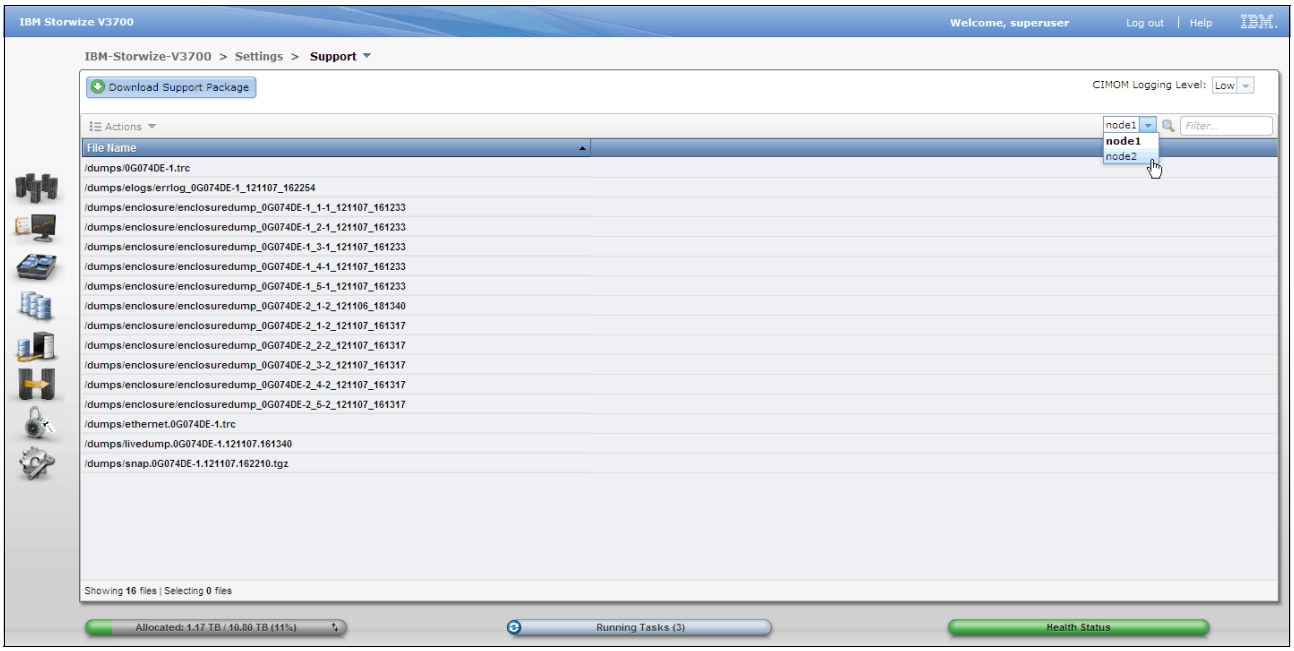

Figure 3-114 Full log listing

General panel

Select **General** in the Settings menu to open the panel. The General panel provides access to the Date and Time, Upgrade Machine Code, and GUI Preferences configuration panels. Figure 3-115 shows the General panel.

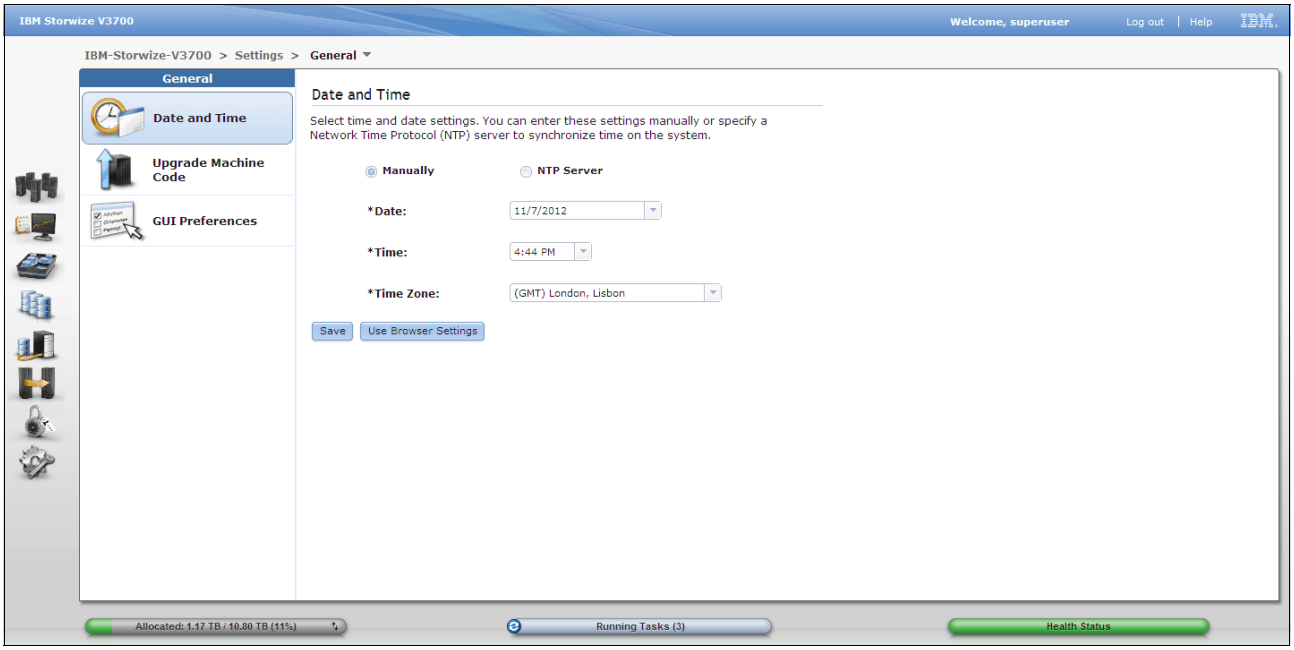

Figure 3-115 General panel

Date and Time

Click **Data and Time** to configure date and time manually or via Network Time Protocol (NTP) server. Figure 3-115 shows the Date and Time function of the General panel.

Upgrade machine code

IBM recommends that you use the latest version of machine code. The Upgrade Machine Code panel displays the current machine code level. If the system has access to the internet, it contacts the IBM upgrade server to check if the current level is the latest level. If an update is available, a direct link to the code is provided to make the code download process easier.

To upgrade the code, the IBM Storwize V3500 Code and the IBM Storwize V3500 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 machine code of the Storwize V3500, click **Launch Upgrade Wizard**. After the upgrade starts, an Abort button is shown, which can be used to abort the upgrade process. Figure 3-116 shows the Upgrade Machine Code panel.

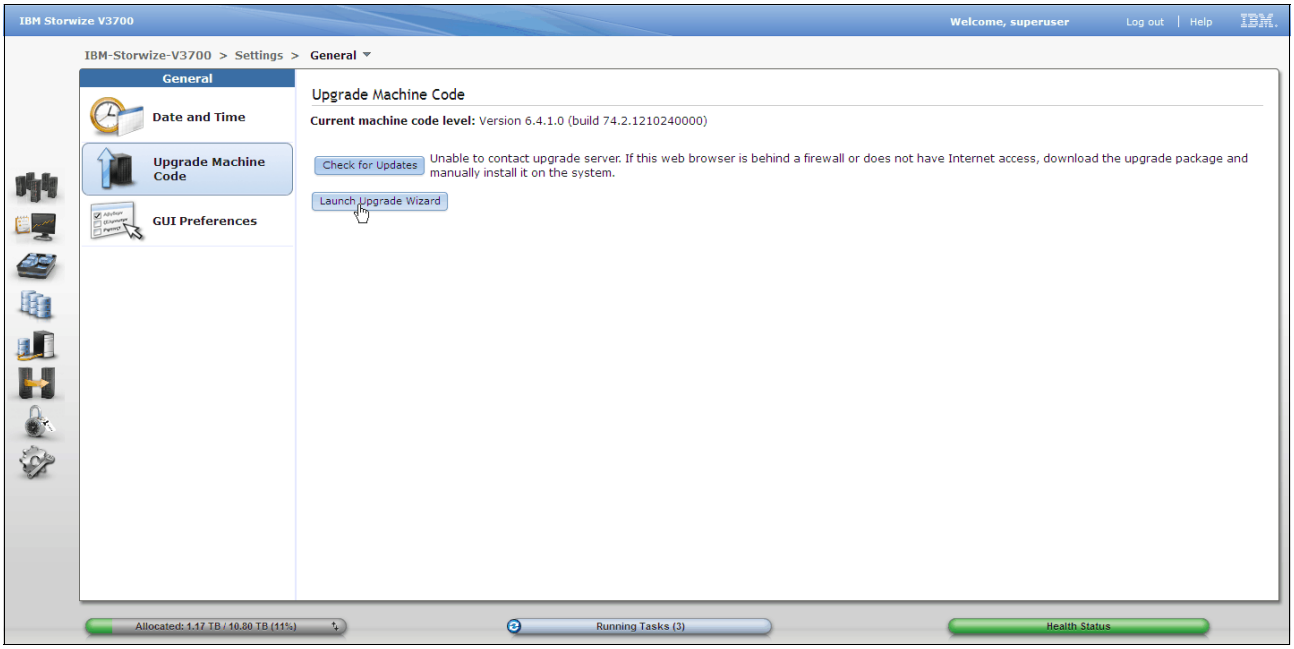

Figure 3-116 Upgrade machine code panel

The first step of the upgrade wizard guides the user to upload the IBM Storwize V3500 Upgrade Test Utility. Figure 3-117 shows step 1of the machine code upgrade wizard.

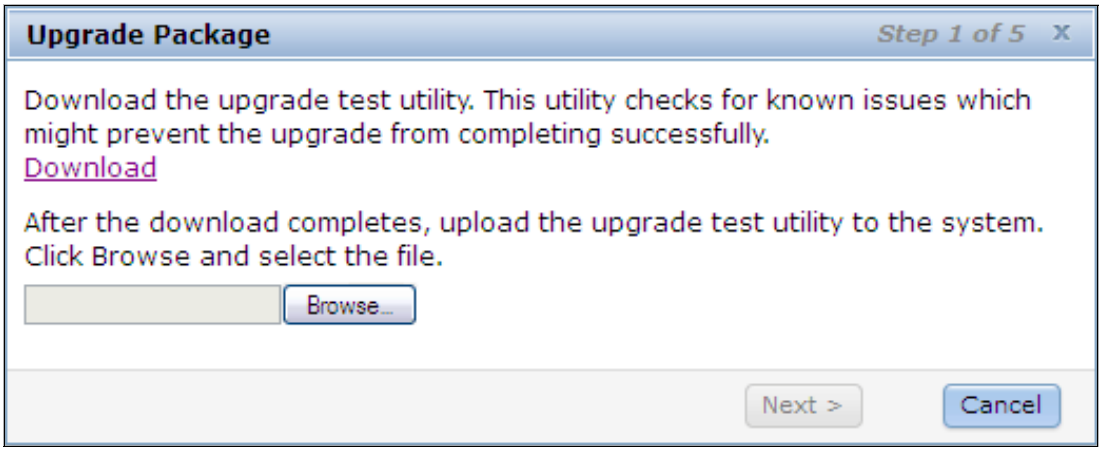

Figure 3-117 Step 1 of the machine code upgrade wizard

GUI Preferences panel

By using the GUI Preferences panel, the user can refresh GUI objects, restore default browser preferences, set table selection policy, and configure the Information Center web address. Figure 3-118 shows the GUI preferences panel.

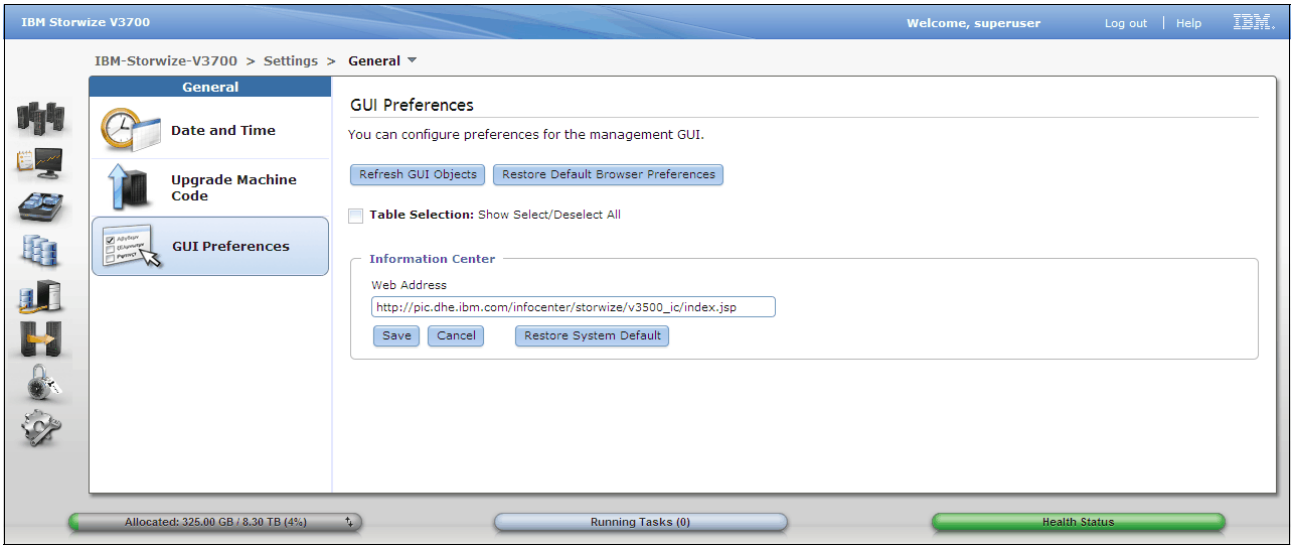

Figure 3-118 GUI preferences panel

3.5 Management GUI help

This section provides information about the various methods that can be used to get help while you are using the Storwize V3500 management GUI. The topics that are covered include the Storwize V3500 Information Center, e-Learning modules, embedded panel help, question mark help, hover help, and IBM endorsed Youtube videos.

3.5.1 Storwize V3500 Information Center

The best source of information for the Storwize V3500 is the Information Center. Click **Visit the Information Center** to access the online version from the Overview panel, as shown in Figure 3-119.

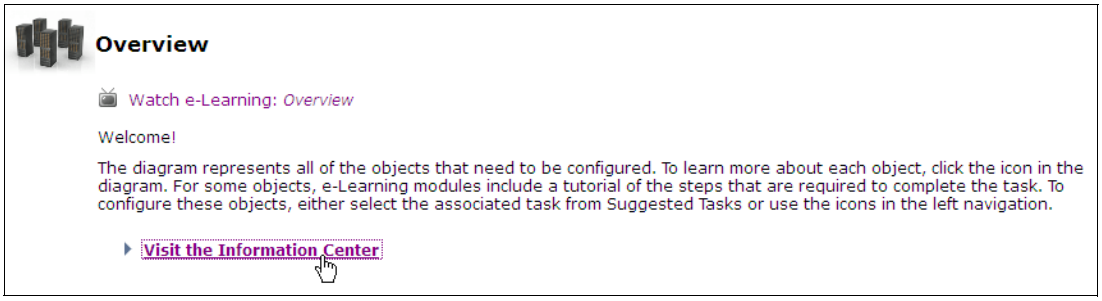

Figure 3-119 Overview panel information center link

3.5.2 Watch e-Learning

The Storwize V3500 provides embedded e-Learning videos. The videos provide easy-to-follow directions to complete various tasks. Click **Watch e-Learning:...** to start a video. Figure 3-120 shows the Watch e-Learning module.

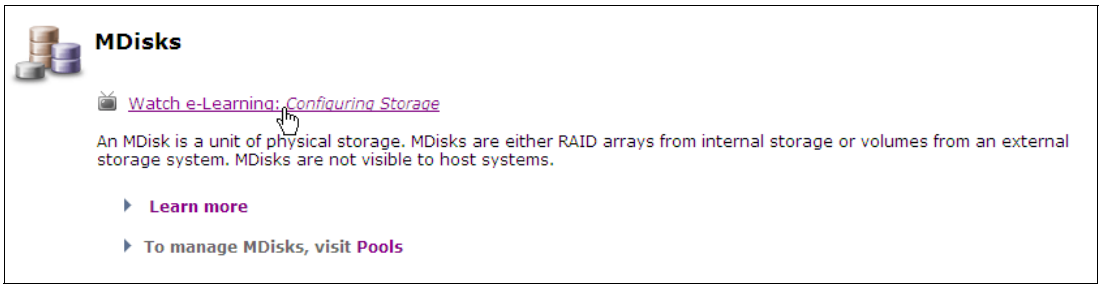

Figure 3-120 Watch e-Learning module

3.5.3 Learn more

The Storwize V3500 provides embedded Learn more links to explain important concepts and panels. Click **Learn more** to open the information panel, as shown in Figure 3-121.

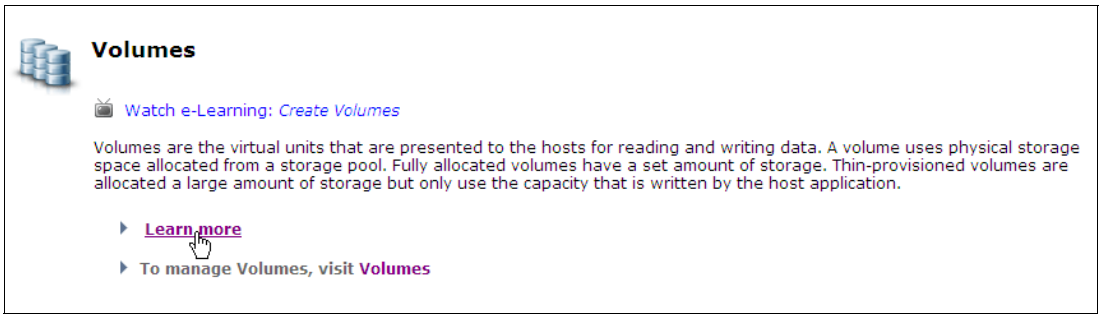

Figure 3-121 Learn more link

Figure 3-122 shows the information panel.

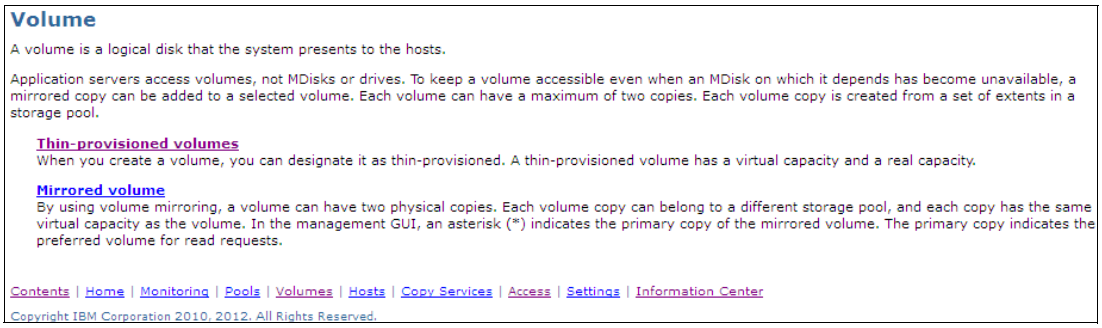

Figure 3-122 Information panel

3.5.4 Embedded panel help

The Storwize V3500 provides embedded help that is available on each panel. Click **Help** to open the information panel, as shown in Figure 3-123.

Figure 3-123 Embedded panel help

Figure 3-124 shows the information panel that is opened from the embedded panel help. The information panel includes hotlinks to various other panels, including the Information Center.

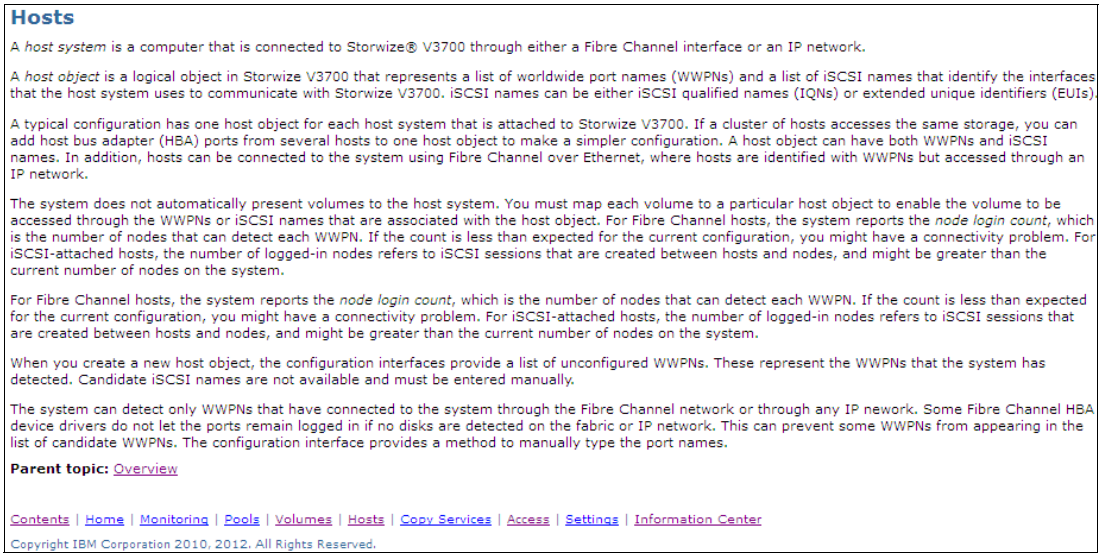

Hosts

A *host system* is a computer that is connected to Storwize® V3700 through either a Fibre Channel interface or an IP network.

A *host object* is a logical object in Storwize V3700 that represents a list of worldwide port names (WWPNs) and a list of iSCSI names that identify the interfaces that the host system uses to communicate with Storwize V3700. iSCSI names can be either iSCSI qualified names (IQNs) or extended unique identifiers (EUIs).

A typical configuration has one host object for each host system that is attached to Storwize V3700. If a cluster of hosts accesses the same storage, you can add host bus adapter (HBA) ports from several hosts to one host object to make a simpler configuration. A host object can have both WWPNs and iSCSI names. In addition, hosts can be connected to the system using Fibre Channel over Ethernet, where hosts are identified with WWPNs but accessed through an IP network.

The system does not automatically present volumes to the host system. You must map each volume to a particular host object to enable the volume to be accessed through the WWPNs or iSCSI names that are associated with the host object. For Fibre Channel hosts, the system reports the *node login count*, which is the number of nodes that can detect each WWPN. If the count is less than expected for the current configuration, you might have a connectivity problem. For iSCSI-attached hosts, the number of logged-in nodes refers to iSCSI sessions that are created between hosts and nodes, and might be greater than the current number of nodes on the system.

For Fibre Channel hosts, the system reports the *node login count*, which is the number of nodes that can detect each WWPN. If the count is less than expected for the current configuration, you might have a connectivity problem. For iSCSI-attached hosts, the number of logged-in nodes refers to iSCSI sessions that are created between hosts and nodes, and might be greater than the current number of nodes on the system.

When you create a new host object, the configuration interfaces provide a list of unconfigured WWPNs. These represent the WWPNs that the system has detected. Candidate iSCSI names are not available and must be entered manually.

The system can detect only WWPNs that have connected to the system through the Fibre Channel network or through any IP nework. Some Fibre Channel HBA device drivers do not let the ports remain logged in if no disks are detected on the fabric or IP network. This can prevent some WWPNs from appearing in the list of candidate WWPNs. The configuration interface provides a method to manually type the port names.

Parent topic: Overview

Contents | Home | Monitoring | Pools | Volumes | Hosts | Copy Services | Access | Settings | Information Center

Figure 3-124 Information panel

3.5.5 Hidden question mark help

The Storwize V3500 provides a hidden question mark help feature for some settings or items that are found in the various configuration panels. This help feature is accessed by hovering the muse pointer adjacent to an item. The question mark is then shown and the help bubble is displayed, as shown in Figure 3-125.

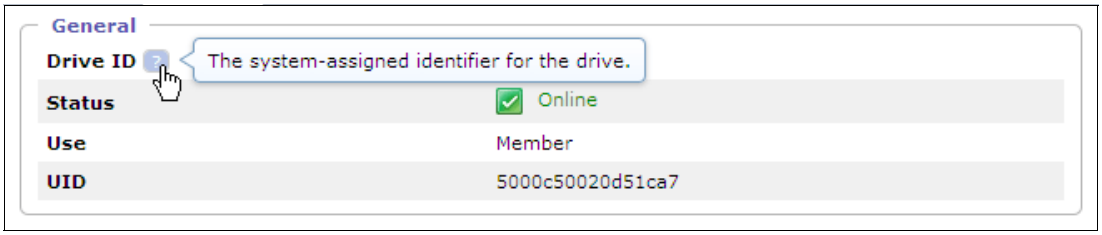

Figure 3-125 Hidden question mark help

3.5.6 Hover help

The Storwize V3500 provides hidden help tags that are shown when you hover the mouse pointer over various functions and items, as shown in Figure 3-126.

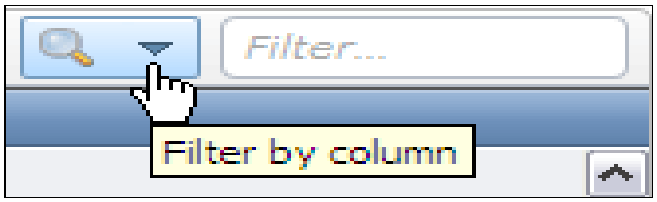

Figure 3-126 Hover help

3.5.7 IBM endorsed YouTube videos

IBM has endorsed a range of YouTube videos for the IBM storage portfolio. Client feedback suggests that this is a good tool to display management GUI navigation and tasks. New videos are available from IBM Storage at this website:

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

4

Host configuration

This chapter describes how to use the IBM Storwize V3500 GUI to create hosts, and how to prepare a host to access the volumes that are created. For more information about volume creation, see Chapter 5, “Basic volume configuration” on page 171.

This chapter includes the following topics:

- Host attachment overview
- Preparing the host operating system
- Creating hosts by using the GUI

4.1 Host attachment overview

A host system is an open-systems computer that is connected to the switch through a Fibre Channel or an iSCSI interface.

This chapter describes the following topics:

- Preparing the host operating system:
  - Windows:
    - Fibre Channel (FC)
    - iSCSI
  - VMware:
    - FC
    - iSCSI
- Creating hosts by using the Storwize V3500 GUI
  - Creating FC hosts
  - Creating iSCSI hosts

In this chapter, we assume that your hosts are connected to your FC or IP network and you completed the steps that are described in Chapter 2, “Initial configuration” on page 27. Follow basic zoning recommendations to 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 is connected to both canisters. This setup assures four paths for failover and failback purposes.

Before mapping the newly created volumes on the host of your choice, a little preparation goes a long way towards ease of use and reliability. Several steps are required to prepare a host for mapping new IBM Storwize V3500 volumes to the host. Use the System Storage Interoperation Center (SSIC) to check which code levels are supported to attach your host to your storage. SSIC is a web tool that checks the interoperation of host, storage, switches, and multipathing drivers. It can be found at this website:

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

The complete support matrix is listed in the IBM Storwize V3500 Supported Hardware List, Device Driver, Firmware, and Recommended Software Levels V6.4 document, which is available at this website:

http://www-01.ibm.com/support/docview.wss?uid=ssg1S1004171

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 V3500 Information Center at this website:

http://pic.dhe.ibm.com/infocenter/storwize/V3500_ic/index.jsp

4.2 Preparing the host operating system

In this section, we describe how to prepare Windows and VMware hosts for attachment to an IBM Storwize V3500 by using FC or iSCSI to connect.

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 V3500 by using FC:

1. Make sure that the latest OS service pack and test fixes are applied to your Microsoft server.
2. Use the latest firmware and driver levels on your host system.
3. Install HBA or HBAs on the Windows server by using the latest BIOS and drivers.
4. Connect the FC Host Adapter ports to the switches.
5. Configure the switches (zoning).
6. Configure the HBA for hosts that are running Windows.
7. Set the Windows timeout value.
8. Install the multipath module.

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 (QLogic), HBAnyware (Emulex), or HBA Software Installer (Brocade) to install the driver. Also, check and update the BIOS (firmware) level of the HBA by using the manufacturer's provided tools. Check the `readme` file to see whether there are Windows registry parameters that should be set for the HBA driver.

The latest supported levels are available at this website:

`http://ibm.com/support/docview.wss?uid=ssg1S1003703#_Win2008`

Configuring Brocade HBAs for Windows

This section applies to Windows hosts that have Brocade HBAs installed. After installing the device driver and firmware, you must configure the HBAs. To perform this task, use the Brocade 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 installing the device driver and firmware, you must configure the HBAs. To perform 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)
- 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 is installed, you must configure the HBAs. To perform 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. To verify this setting, complete the following steps:

1. Click **Start** → **Run**.
2. In the panel, enter `regedit` and press Enter.
3. In the registry editor, locate 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.

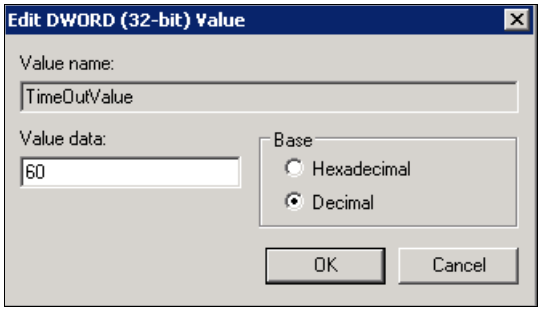

Figure 4-1 Windows timeout value

Installing the multipathing software

Microsoft Multipath Input/Output (MPIO) solutions are designed to work with device-specific modules (DSMs) that are written by vendors. However, the MPIO driver package does not, by itself, form a complete solution. This joint solution allows the storage vendors to design device-specific solutions that are tightly integrated with the Windows operating system. MPIO is not shipped with the Windows operating system; storage vendors must pack the MPIO drivers with their own DSM.

IBM Subsystem Device Driver DSM (SDDDSM) is the IBM multipath I/O solution that is based on Microsoft MPIO technology. It is a device-specific module that is designed to support IBM storage devices on Windows hosts. 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 ensure correct multipathing with IBM Storwize V3500, SDDDSM must be installed on Windows hosts. To install SDDDSM, complete the following steps:

1. Check the SDDDSM download matrix to determine the correct level of SDDDSM to install for Windows 2008 (R2) and download the package from this website:

   `http://ibm.com/support/docview.wss?uid=ssg1S7001350#WindowsSDDDSM`

2. Extract the package to your hard disk drive and run `setup.exe` to install SDDDSM. A command prompt window opens, as shown in Figure 4-2. Confirm the installation by entering `Y`.

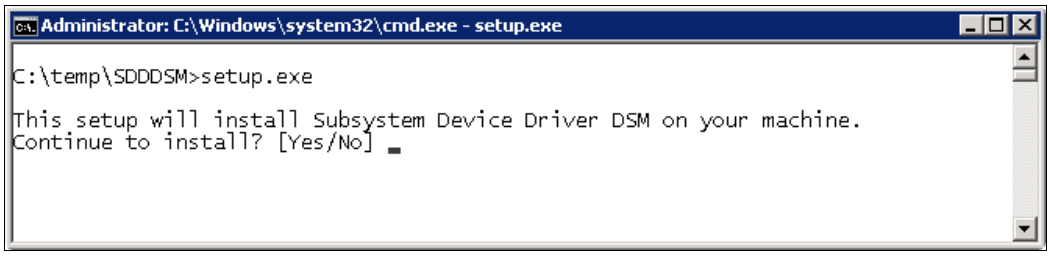

Figure 4-2 SDDDSM setup

3. After the setup completes, you are prompted to restart the system. Confirm this restart by entering `yes` and pressing Enter, as shown in Figure 4-3.

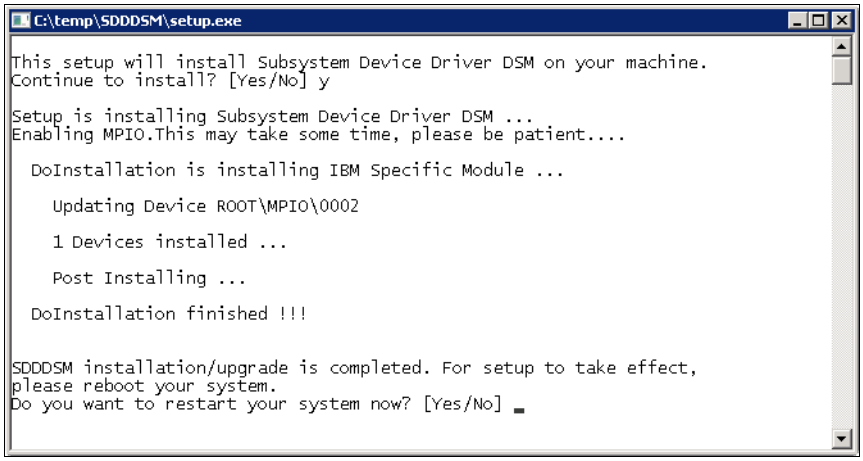

Figure 4-3 Answer yes to restart the host

You successfully installed IBM SDDDSM. You can check the installed driver version if you select **Start** → **All Programs** → **Subsystem Device Driver DSM** → **Subsystem Device Driver DSM**. A command prompt opens. Run the **`datapath query version`** command to determine the version that is installed for this Windows 2008 R2 host, as shown in Example 4-1.

Example 4-1 The datapath query version command

```
C:\Program Files\IBM\SDDDSM>datapath.exe query version
IBM SDDDSM Version 2.4.3.1-2
Microsoft MPIO Version 6.1.7601.17514
```

This command also can be used to determine the WWPNs of the host. Run the **datapath query wwpn** command (as shown in Example 4-2) and note the WWPNs of your host because you need them later.

Example 4-2 The datapath query wwpn command

```
C:\Program Files\IBM\SDDDSM>datapath.exe query wwpn
        Adapter Name    PortWWN
        Scsi Port  7    100000051EC76B89
        Scsi Port  7    100000051EC76B8A
```

For more information about SDDDSM, see *Multipath Subsystem Device Driver User's Guide,* GC52-1309.

The Windows host is prepared to connect to the IBM Storwize V3500 and you know the WWPNs of the host. The next step is to configure a host object for the WWPNs by using the IBM Storwize V3500 GUI. This task is described in 4.3.1, "Creating Fibre Channel hosts" on page 160.

SAN Boot hosts are beyond the intended scope of this book. For more information about these hosts, follow the steps in the Information Center, which is available from the IBM Support Portal.

Windows 2003: The examples focus on Windows 2008 R2, but the procedure for Windows 2003 is similar. If you use Windows 2003, do not forget to install Microsoft Hotfix 908980. If you do not install it before you perform this procedure, preferred pathing is not available. You can download this hotfix at this website:

http://support.microsoft.com/kb/908980

4.2.2 Windows 2008 R2: Preparing for iSCSI attachment

In Windows 2008 R2, the Microsoft iSCSI software initiator is preinstalled. Enter `iscsi` in the search field of the Windows start menu (as shown in Figure 4-4) and click **iSCSI Initiator**.

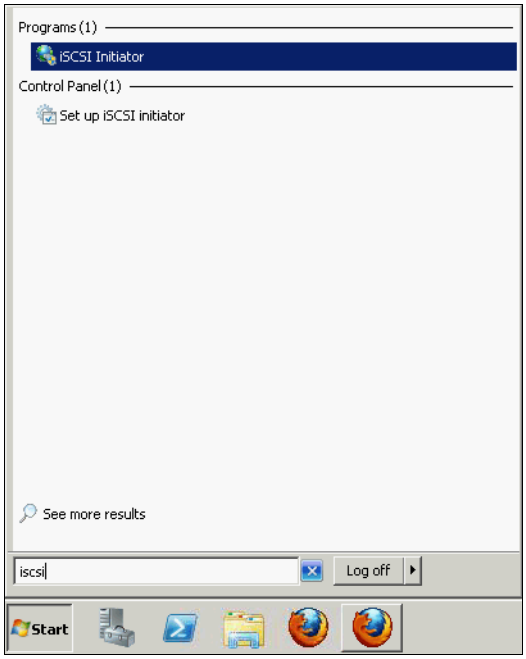

Figure 4-4 Windows iSCSI Initiator

Confirm the automatic startup of the iSCSI Service, as shown in Figure 4-5.

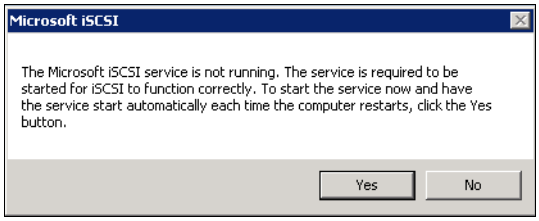

Figure 4-5 Automatic startup of the iSCSI Service

The iSCSI Configuration window opens. Select the **Configuration** tab, as shown in Figure 4-6. Make note of the initiator name of your Windows host because you need it later.

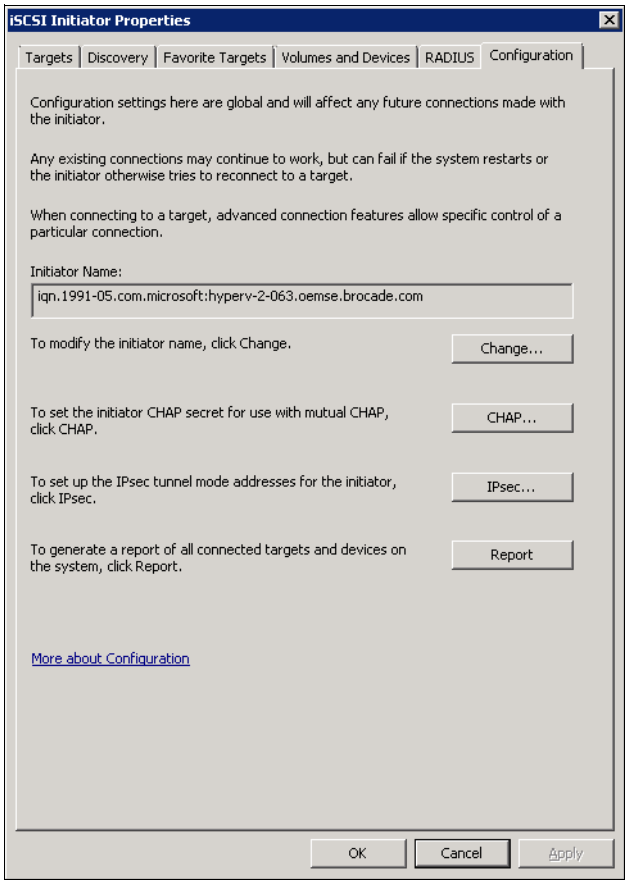

Figure 4-6 iSCSI Initiator Properties window

You can change the initiator name or enable advanced authentication, but these tasks are out of the scope of our basic setup. By default, iSCSI authentication is not enabled. For more information, see the IBM Storwize V3500 Information Center at this website:

http://pic.dhe.ibm.com/infocenter/storwize/v3500_ic/index.jsp?topic=%2Fcom.ibm.storwize.v3500.641.doc%2Fsvc_iscsiwindowsauthen_fu67gt.html

Setting the Windows registry keys

Complete the following steps to change the system registry to make your iSCSI operations more reliable:

1. In the search field of the Windows Start menu, enter `regedit` and click **regedit.exe**.
2. In the registry editor, locate the `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<bus ID>\Parameters\LinkDownTime` key.

   Confirm that the value for the `LinkDownTime` key is 120 (decimal value). If necessary, change the value to 120.

3. In the registry editor, locate the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<bus ID>\Parameters\MaxRequestHoldTime key.

   Confirm that the value for the MaxRequestHoldTime key is 120 (decimal value), and, if necessary, change the value to 120.

4. In the registry editor, locate the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<bus ID>\Parameters\ MaxPendingRequests key.

   Confirm that the value for the MaxPendingRequests key is 2048 (decimal value). If necessary, change the value to 2048.

5. In the registry editor, locate the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Disk\TimeOutValue key.

   Confirm that the value for the TimeOutValue key is 60 (decimal value). If necessary, change the value to 60.

6. Reboot your host for these changes to take effect.

These steps are the basic steps to prepare a Windows 2008 R2 host for iSCSI attachment. For more information about configuring the IBM Storwize V3500 for iSCSI connections, see 4.3.2, “Creating iSCSI hosts” on page 164.

4.2.3 VMware ESX: Preparing for Fibre Channel attachment

Complete the following steps to prepare a VMware ESXi host to connect to an IBM Storwize V3500 by using FC:

1. Install HBA or HBAs on the ESXi server.
2. Make sure that the latest firmware levels are applied on your host system.
3. Update and configure the HBA for hosts that are running ESXi.
4. Connect the FC Host Adapter ports to the switches.
5. Configure the switches (zoning).
6. Install VMware ESXi and load more drivers, if required.

Downloading and installing the supported firmware

Install the latest firmware levels to your host server. For the HBAs, check the “IBM Storwize V3500 Supported Hardware List, Device Driver, Firmware, and Recommended Software Levels V6.4” list for VMware at this website:

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

Download the latest supported HBA firmware for your configuration and apply it to your system. Some HBAs (especially the new CNA Adapters) require another driver to be loaded into ESX. Check the VMware Compatibility Guide to see whether there are any requirements for your configuration, which is available at this website:

http://www.vmware.com/resources/compatibility/search.php

Configuring Brocade HBAs for VMware ESX

This section applies to ESXi hosts that have Brocade HBAs installed. After the firmware is installed, load the default settings of all your adapters that are installed on the host system and make sure that the Adapter BIOS is disabled, unless SAN Boot is used.

Configuring QLogic HBAs for VMware ESX

This section applies to ESXi hosts that have QLogic HBAs installed. After the firmware is installed, you must configure the HBAs. To perform this task, use the QLogic Sansurfer software or the HBA BIOS, load the adapter defaults, and set the following values:

- Host Adapter Settings:
  - Host Adapter BIOS: Disabled (unless the host is configured for SAN Boot)
  - Frame size: 2048
  - Loop Reset Delay: 5 (minimum)
  - Adapter Hard Loop ID: Disabled
  - Hard Loop ID: 0
  - Spinup Delay: Disabled
  - Connection Options 1: Point to point only
  - Fibre Channel Tape Support: Disabled
  - Data Rate: 2
- Advanced Adapter Settings
  - Execution throttle: 100
  - LUNs per Target: 0
  - Enable LIP Reset: No
  - Enable LIP Full Login: Yes
  - Enable Target Reset: Yes
  - Login Retry Count: 8
  - Link Down Timeout: 10
  - Command Timeout: 20
  - Extended event logging: Disabled (enable it only for debugging)
  - RIO Operation Mode: 0
  - Interrupt Delay Timer: 0

Configuring Emulex HBAs for VMware ESXi

This section applies to ESXi hosts that have Emulex HBAs installed. After the firmware is installed, load the default settings of all your adapters that are installed on the host system and make sure that the Adapter BIOS is disabled, unless SAN Boot is used.

VMware ESXi installation

Install your VMware ESXi server and load any other drivers and patches, if required. For more information, see this website:

http://pubs.vmware.com/vsphere-50/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-50-storage-guide.pdf

After you complete your ESXi installation, connect to your ESXi Server by using the vSphere client and browse to the Configuration tab. Click **Storage Adapters** and scroll down to your FC HBAs, as shown in Figure 4-7. Make a note of the WWPNS of the installed adapters for later use.

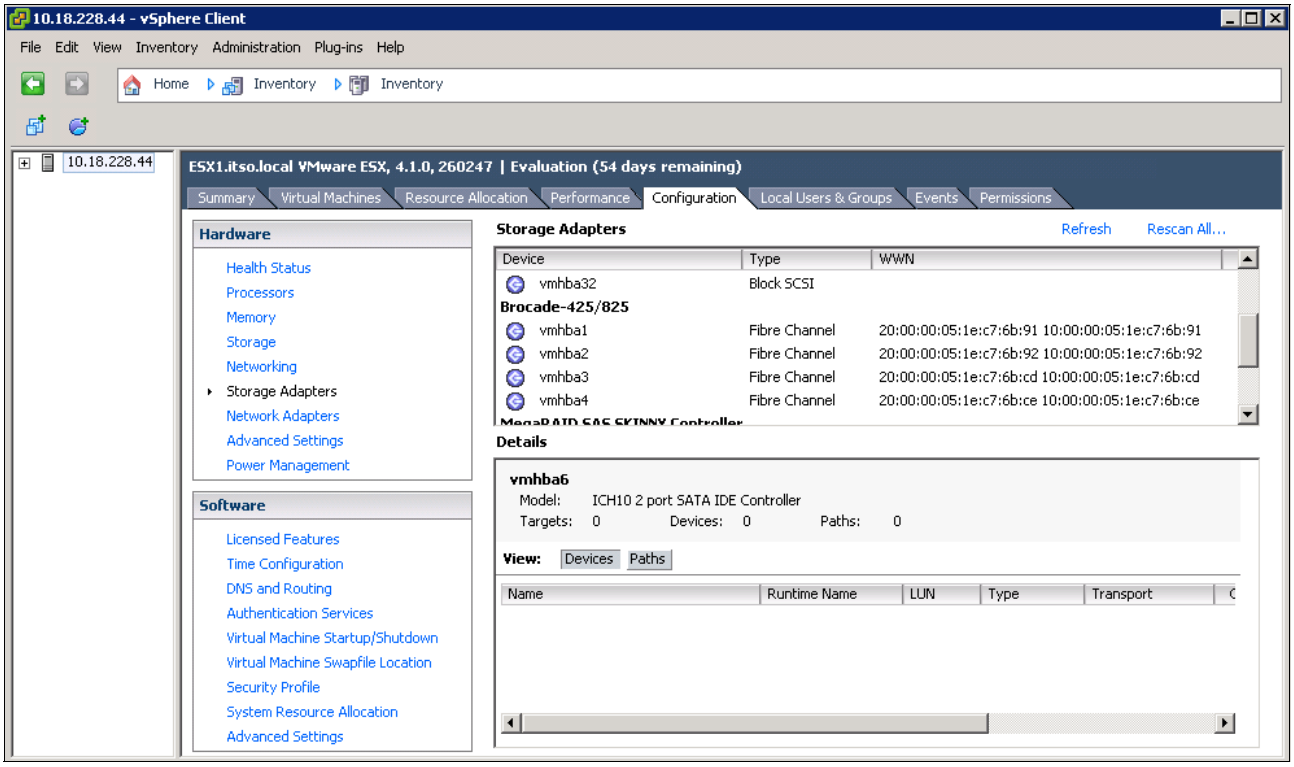

Figure 4-7 Show WWPNs in VMware ESX

VMware ESXi multipathing

The ESXi server has its own multipathing software. You do not need to install a multipathing driver on the ESXi server or on the guest operating systems. The ESXi multipathing policy supports the following operating modes:

- Round Robin
- Fixed
- Most Recently Used (MRU)

The IBM Storwize V3500 is an active/active storage device. Since VMware ESXi 5.0 and later, the suggested multipathing policy is Round Robin. Round Robin performs static load balancing for I/O. If you do not want to have the I/O balanced over all available paths, the Fixed policy also is supported. This policy setting can be selected for every volume. Set this policy after IBM Storwize V3500 LUNs is attached to the ESXi host. For more information, see 5.3.3, “VMware ESX Fibre Channel attachment” on page 202. If you use an older version of VMware ESX (up to Version 3.5), Fixed is the recommended policy setting.

After all of these steps are completed, the ESXi host is prepared to connect to the IBM Storwize V3500. To create the ESX FC host in the IBM Storwize V3500 GUI, see 4.3.1, “Creating Fibre Channel hosts” on page 160.

4.2.4 VMware ESX: Preparing for iSCSI attachment

This section describes how to enable iSCSI on VMware ESX hosts. In this book, we focus on vSphere (ESXi 5.0 and later) because the complete iSCSI stack was updated in this level to offer improved performance. It also supports useful features, such as jumbo frames and increased numbers of switches that are supported per server and doubles the number of port groups that are supported per server. We focus on the basic ESXi iSCSI setup. More information is provided in the VMware *iSCSI SAN Configuration Guide*, which is available at this website:

http://pubs.vmware.com/vsphere-50/topic/com.vmware.icbase/pdf/vsphere-esxi-vcenter-server-50-storage-guide.pdf

Complete the following steps to prepare a VMware ESXi host to connect to an IBM Storwize V3500 by using iSCSI:

1. Make sure that the latest firmware levels are applied on your host system.
2. Install VMware ESXi and load other drivers, if required.
3. Connect the ESXi server to your network. You should use separate network interfaces for iSCSI traffic.
4. Configure your network to fulfill your security and performance requirements.

The iSCSI initiator is installed by default on your ESXi server; you must only enable it. To enable it, complete the following steps:

1. Connect to your ESXi server by using the vSphere Client. Browse to Configuration and select **Networking**, as shown in Figure 4-8.

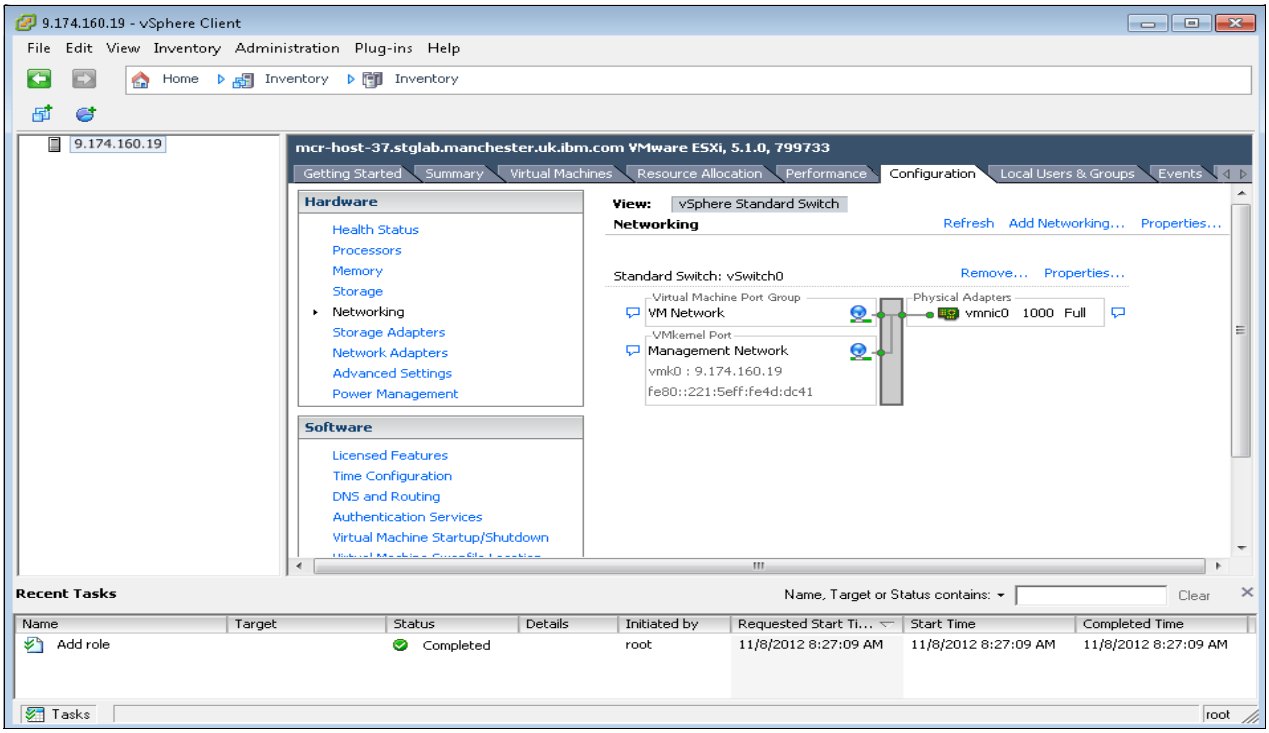

Figure 4-8 Select VMware networking

2. Click **Add Networking** to start the Add Network Wizard, as shown in Figure 4-9. Select **VMkernel** and click **Next**.

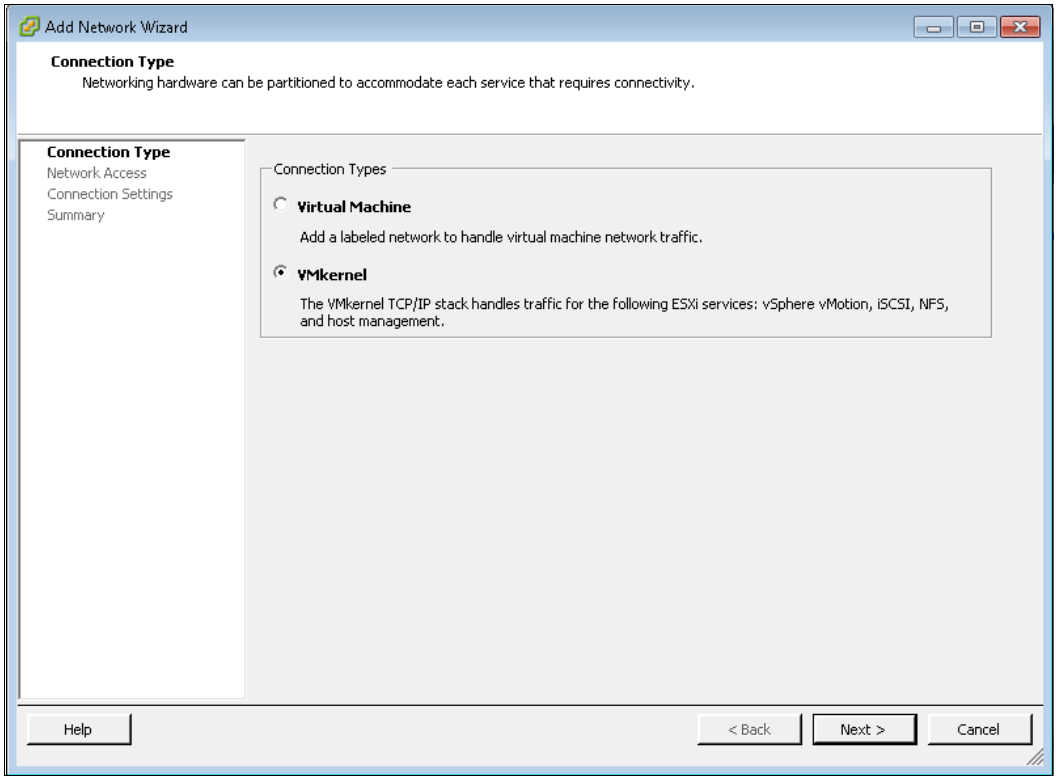

Figure 4-9 Add Network Wizard

Important: The VMkernel performs the virtualization and is referred as the hypervisor.

3. Select one or more network interfaces that you want to use for iSCSI traffic and click **Next**, as shown in Figure 4-10.

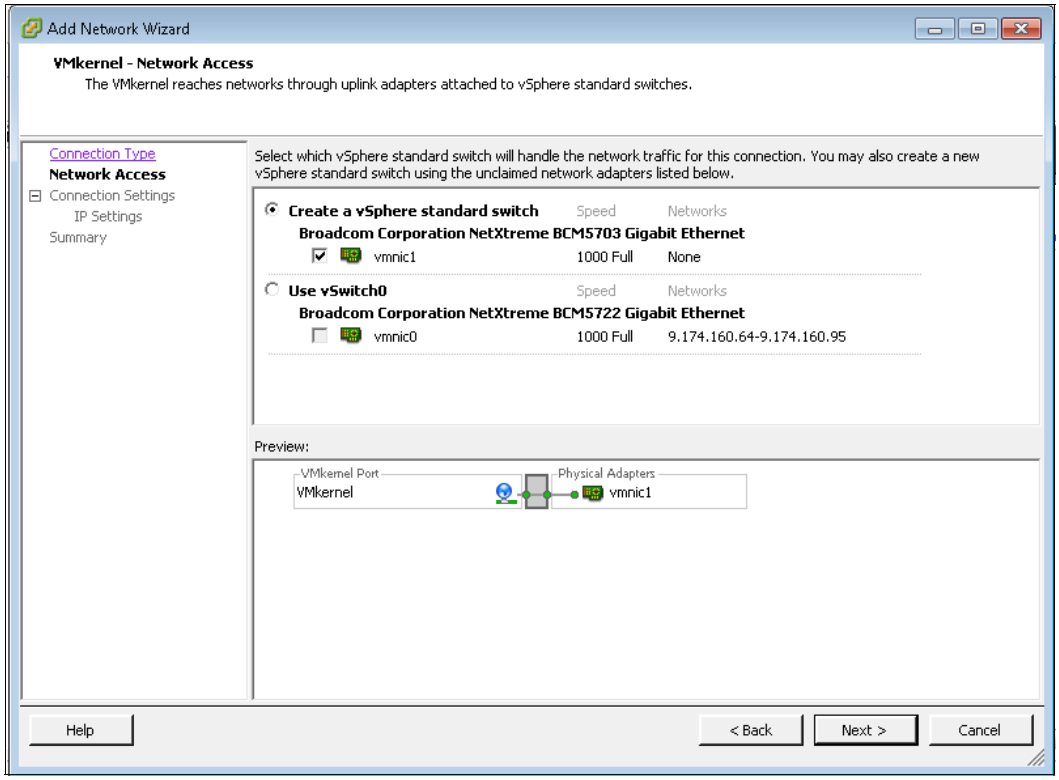

Figure 4-10 Select an iSCSI interface

4. Enter a meaningful Network Label and click **Next**, as shown in Figure 4-11.

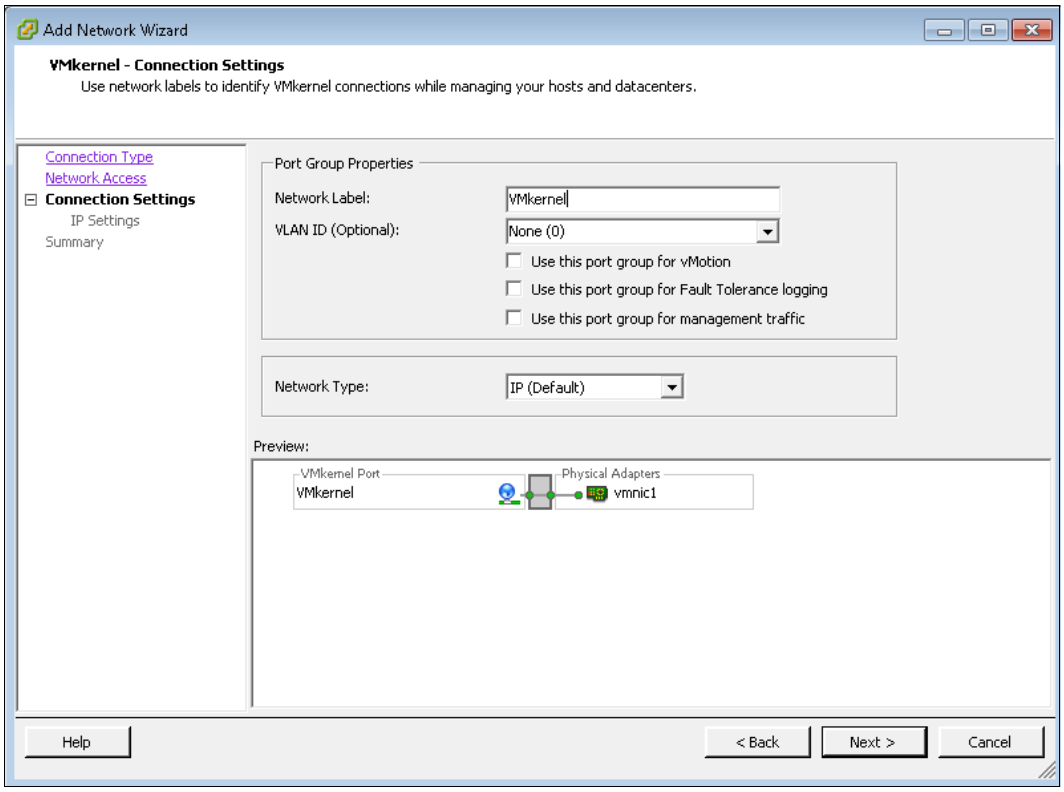

Figure 4-11 Enter a Network Label

5. Enter an IP address for your iSCSI network. You should use a dedicated network for iSCSI traffic, as shown in Figure 4-12.

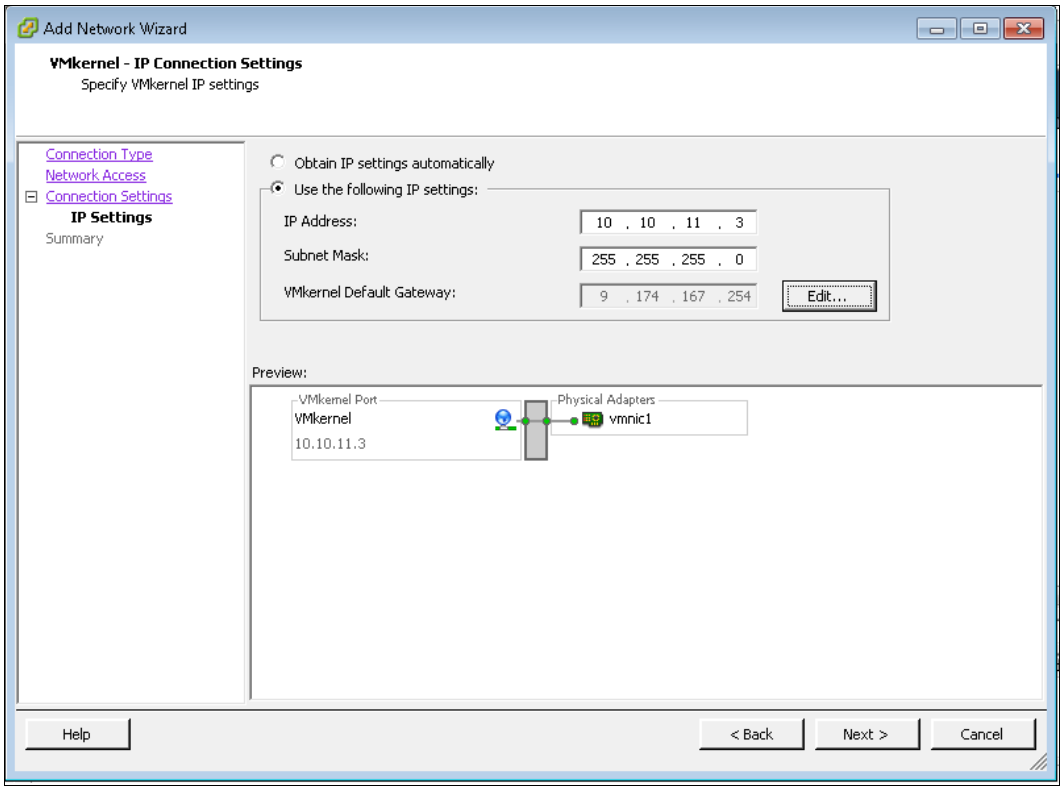

Figure 4-12 Enter iSCSI Network IP

6. Click **Finish** to complete the setup.

7. Select **Storage Adapters** and scroll down to iSCSI Software Adapter, as shown in Figure 4-13. Highlight it and click **Properties**.

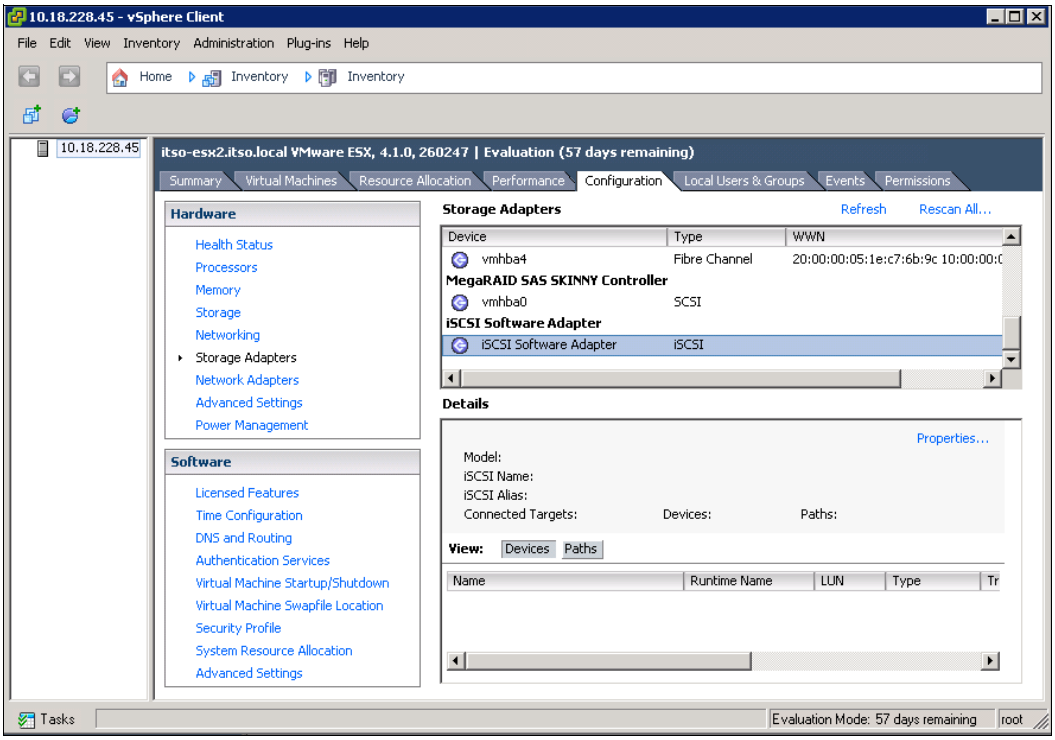

Figure 4-13 iSCSI Software Adapter

8. The iSCSI Software Adapter Properties window opens. As shown in Figure 4-14, the initiator is enabled (changed from VMWare ESX 4.0) by default. If you wanted to change this setting, click **Configure**.

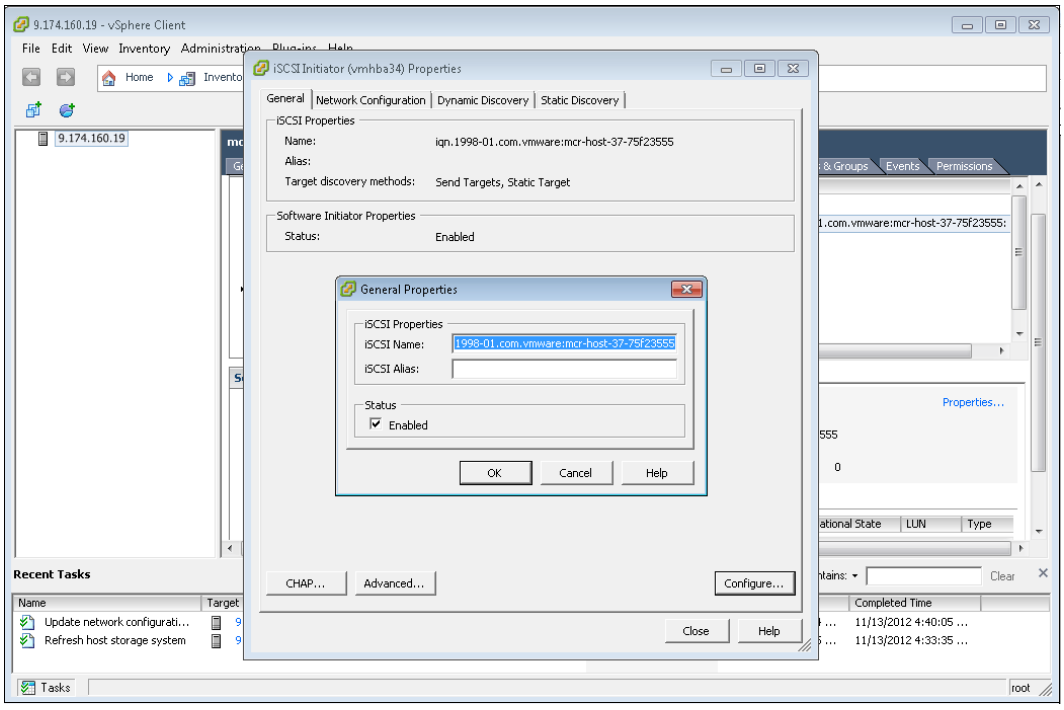

Figure 4-14 iSCSI Software Adapter Properties

The VMware ESX iSCSI initiator is now successfully enabled, as shown in Figure 4-15. Make a note of the initiator name for later use.

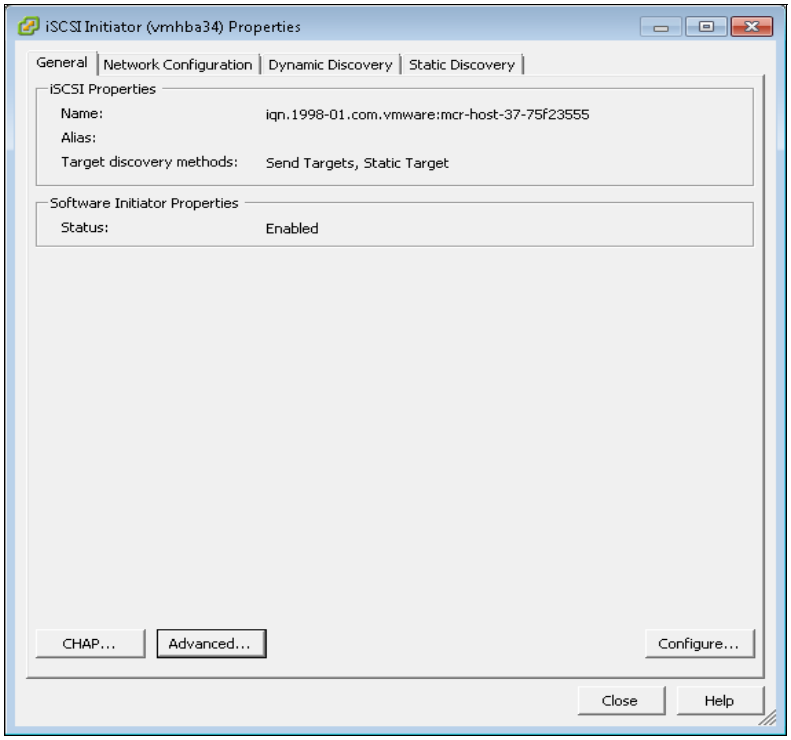

Figure 4-15 VMware iSCSI Initiator enabled

Your VMware ESX host is now prepared to connect to the IBM Storwize V3500. To create the ESX iSCSI host by using the IBM Storwize V3500 GUI, see 4.3.2, “Creating iSCSI hosts” on page 164.

4.3 Creating hosts by using the GUI

This section describes how to create FC and iSCSI hosts by using the IBM Storwize V3500 GUI. We assume that the hosts are prepared for attachment, as described in 4.2, “Preparing the host operating system” on page 143, and that you know the host WWPNs and their iSCSI initiator names.

To create a host, complete the following steps:

1. Open the host configuration window by clicking **Hosts**, as shown in Figure 4-16.

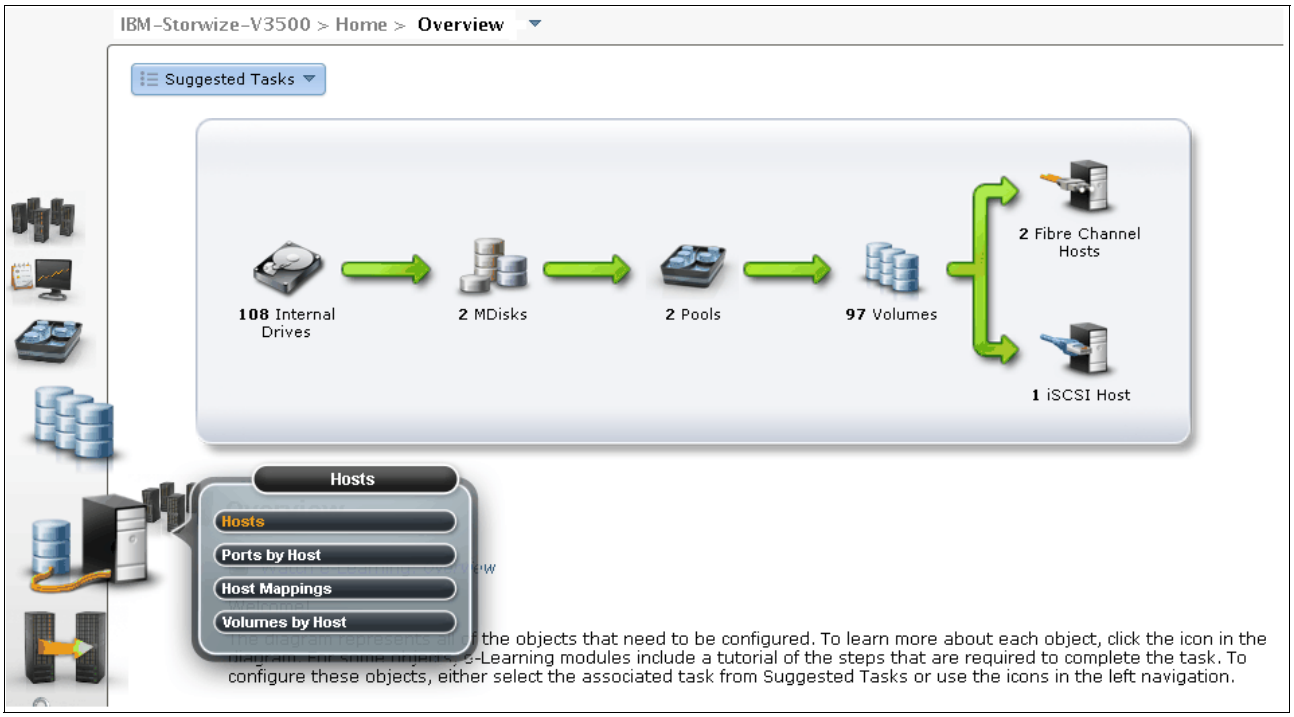

Figure 4-16 Open the host window

2. To create a host, click **New Host** to start the wizard, as shown in Figure 4-17.

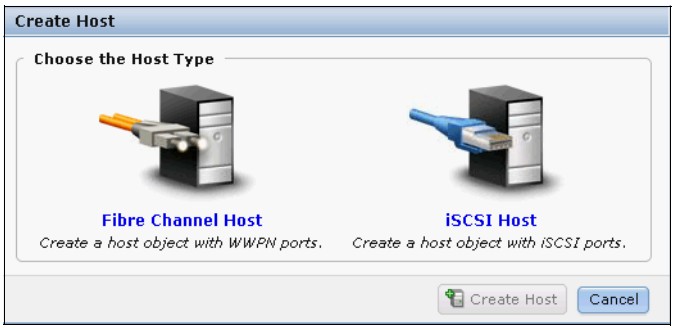

Figure 4-17 Create a host

If you want to create an FC host, continue with 4.3.1, “Creating Fibre Channel hosts” on page 160. To create iSCSI hosts, see 4.3.2, “Creating iSCSI hosts” on page 164.

4.3.1 Creating Fibre Channel hosts

To create FC hosts, complete the following steps:

1. Click **Fibre Channel Host**, as shown in Figure 4-17 on page 159. The Fibre Channel configuration wizard opens, as shown in Figure 4-18.

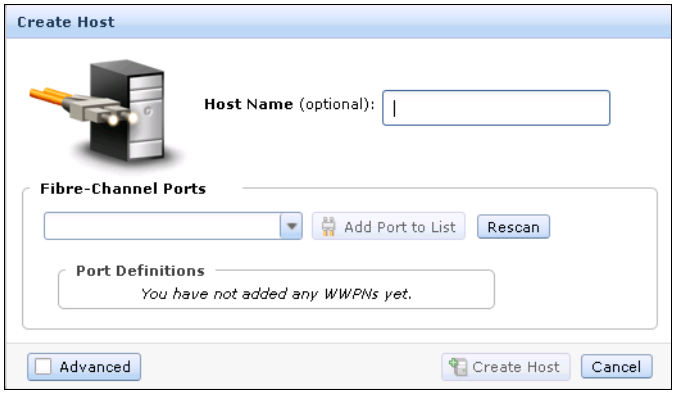

Figure 4-18 Create a Fibre Channel host

2. Enter a host name and click the **Fibre Channel Ports** drop-down menu to see a list of all known WWPNs, as shown in Figure 4-19.

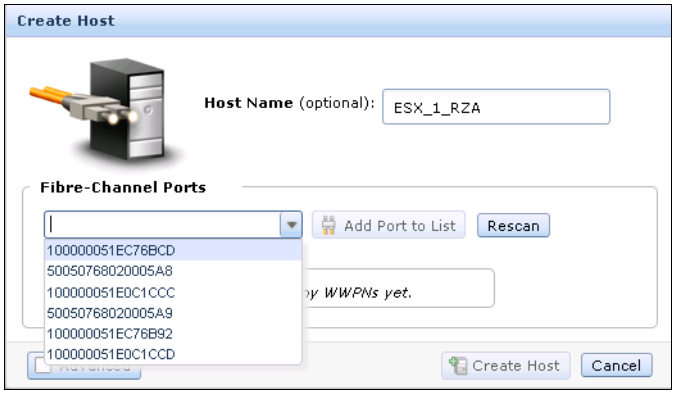

Figure 4-19 Available WWPNs

The IBM Storwize V3500 includes the host port WWPNs that are available if you prepared the hosts as described in 4.2, “Preparing the host operating system” on page 143. 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 SAN zoning and repeat the scanning. For more information about hosts, see this website:

http://pic.dhe.ibm.com/infocenter/storwize/v3500_ic/index.jsp?topic=%2Fcom.ibm.storwize.v3500.641.doc%2Fsvc_fibrechannel_cover.html

AIX HOSTS: AIX hosts WWPNs appear only for a few minutes after logging in to the fabric. You can manually enter the WWPN into the drop-down menu, or you must run the **cfgmgr** command again to allow the host to log in to the fabric and be discovered into the GUI.

3. Select the WWPN for your host and click **Add Port to List**, as shown in Figure 4-20.

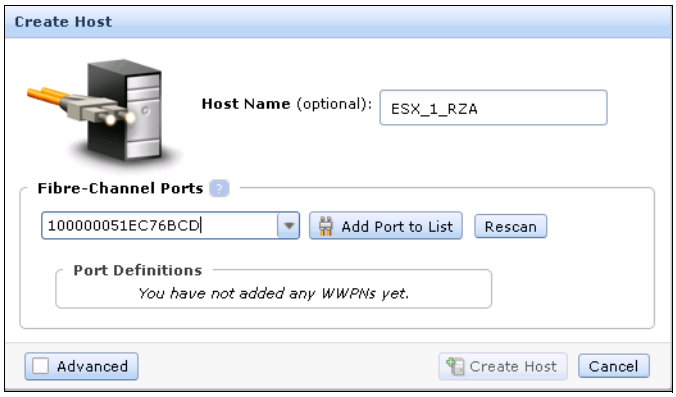

Figure 4-20 Add a port to a list

4. Add all ports that belong to the host, as shown in Figure 4-21.

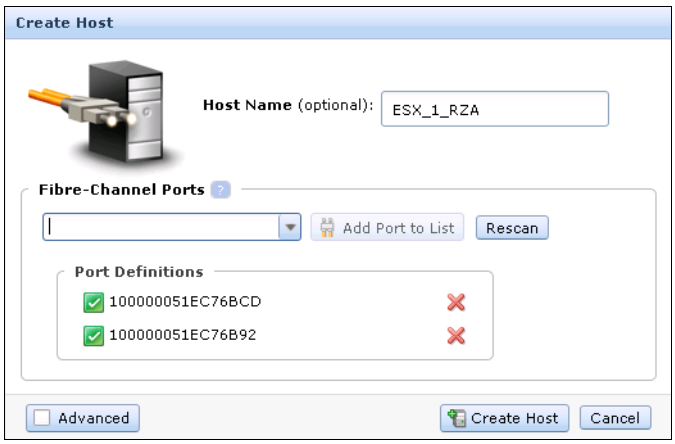

Figure 4-21 Add all WWPNs

Creating offline hosts: If you want to create hosts that are offline or not connected at the moment, it is also possible to enter the WWPNs manually. Enter them into the Fibre Channel Ports panel and add them to the list.

5. If you are creating an HP/UX or TPGS host, select the **Advanced** option and more options appear, as shown in Figure 4-22. Select your host type.

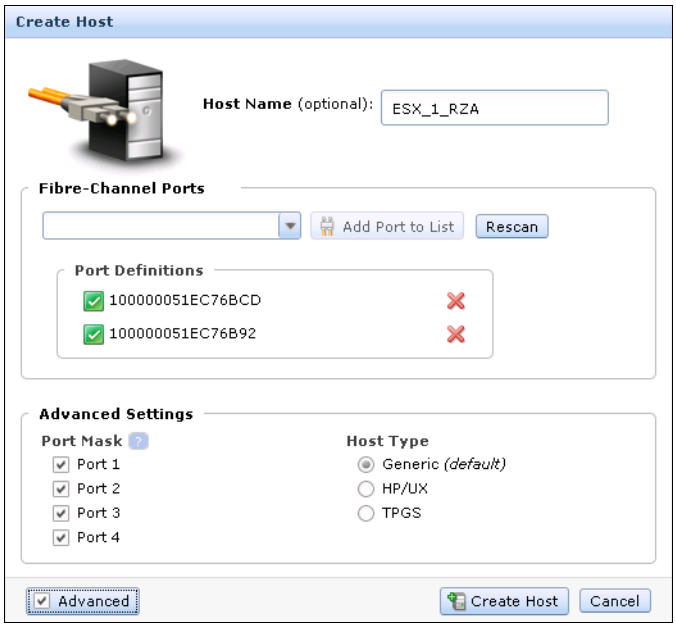

Figure 4-22 Create Host - Advanced Settings

6. Click **Create Host** and the wizard creates the host, as shown in Figure 4-23.

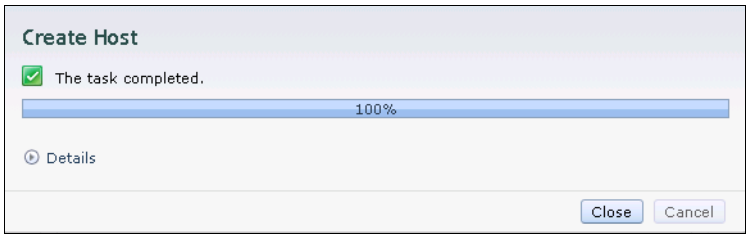

Figure 4-23 Create Host completes

7. Click **Close** to return to the host window, as shown in Figure 4-24.

Figure 4-24 All Hosts

8. Repeat these steps for all of your FC hosts. Figure 4-25 shows the All Hosts window after a second host is created.

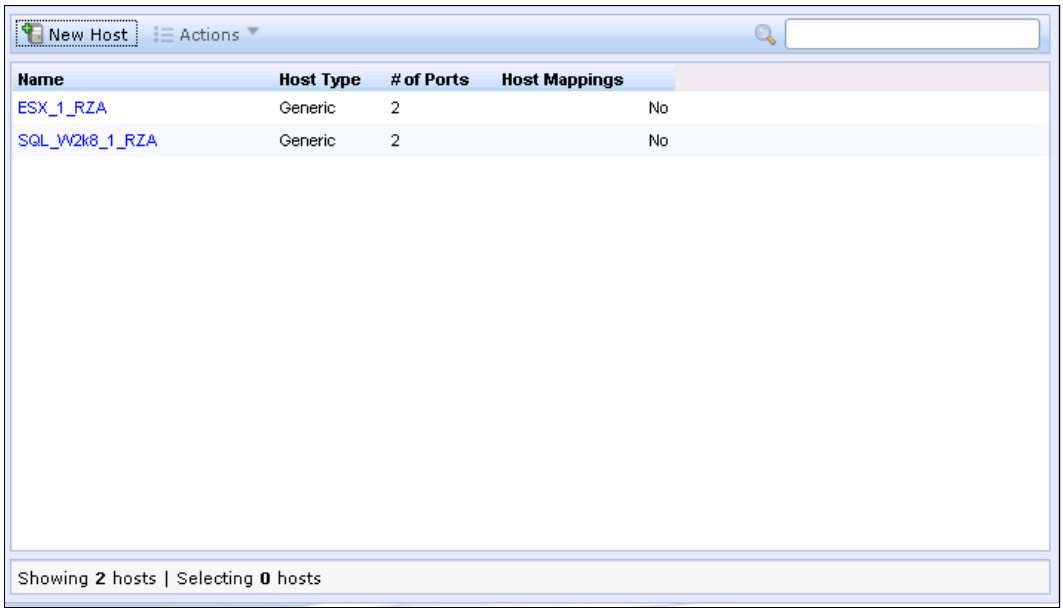

Figure 4-25 All Hosts: After a second host is created

After you create the FC hosts, see Chapter 5, “Basic volume configuration” on page 171 to create volumes and map them to the created hosts.

4.3.2 Creating iSCSI hosts

To create iSCSI hosts, complete the following steps:

1. Click **iSCSI Host** (as shown in Figure 4-17 on page 159) and the iSCSI configuration wizard opens, as shown in Figure 4-26.

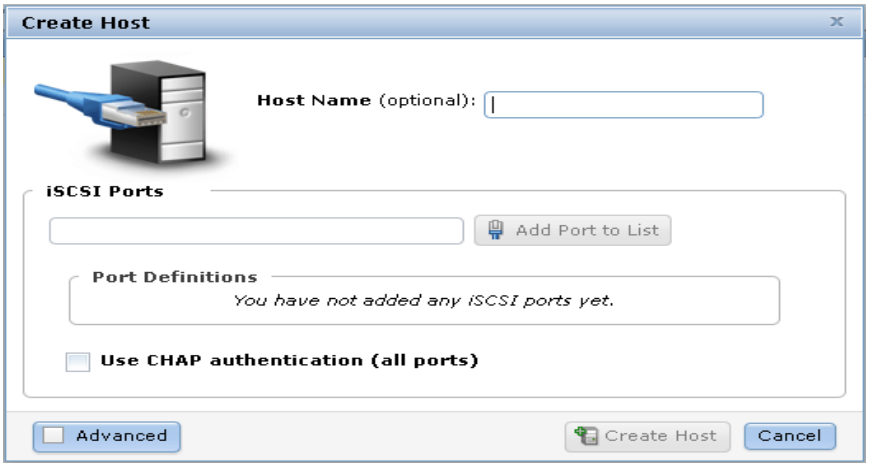

Figure 4-26 Create an iSCSI host

2. Enter a host name. Enter the iSCSI initiator name into the iSCSI Ports panel and click **Add Ports to List**, as shown in Figure 4-27. Repeat this step if you want to add several initiator names to one host.

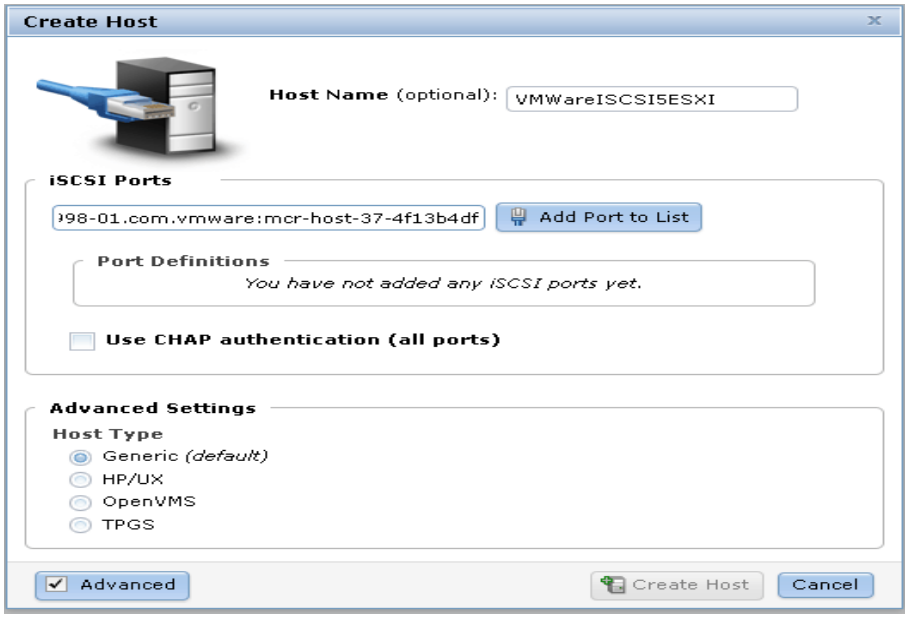

Figure 4-27 Create an iSCSI host - Enter name and iSCSI ports

3. If you are connecting an HP/UX or TPGS host, select the **Advanced** option (as shown in Figure 4-28) and select the correct host type.

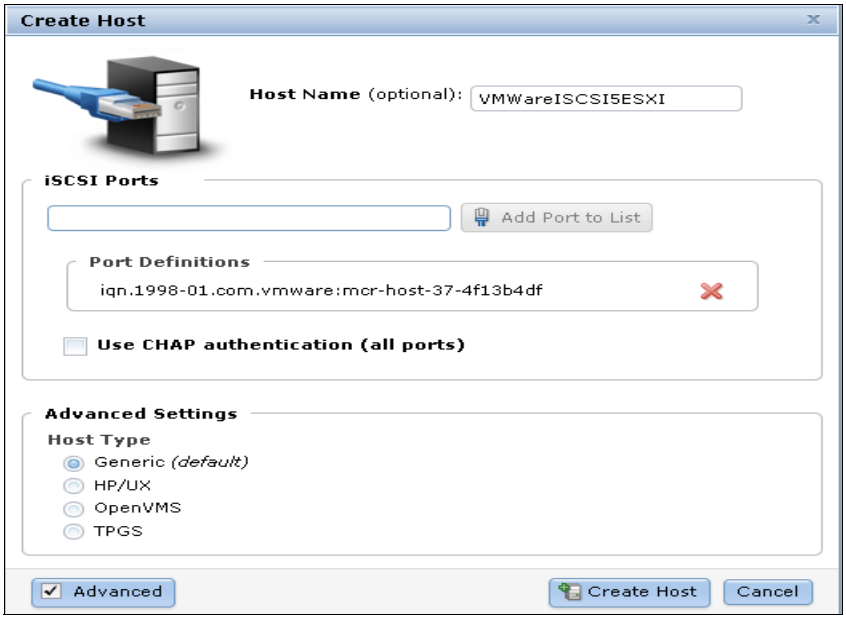

Figure 4-28 Create an iSCSI host: Advanced Settings

4. Click **Create Host** and the wizard completes, as shown in Figure 4-29. Click **Close**.

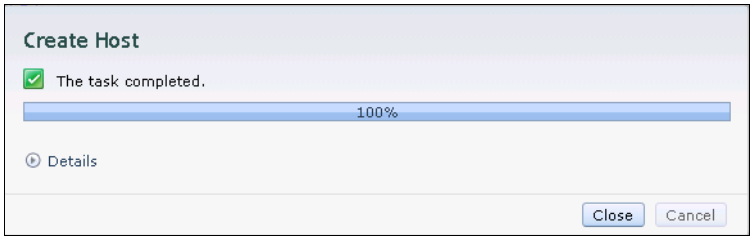

Figure 4-29 Create an iSCSI host: Complete

5. Repeat these steps for every iSCSI host that you want to create.

The iSCSI hosts are now configured on the IBM Storwize V3500. To provide connectivity, the iSCSI Ethernet ports also must be configured. Complete the following steps to enable iSCSI connectivity:

1. Switch to the Configuration window and select **Network**, as shown in Figure 4-30.

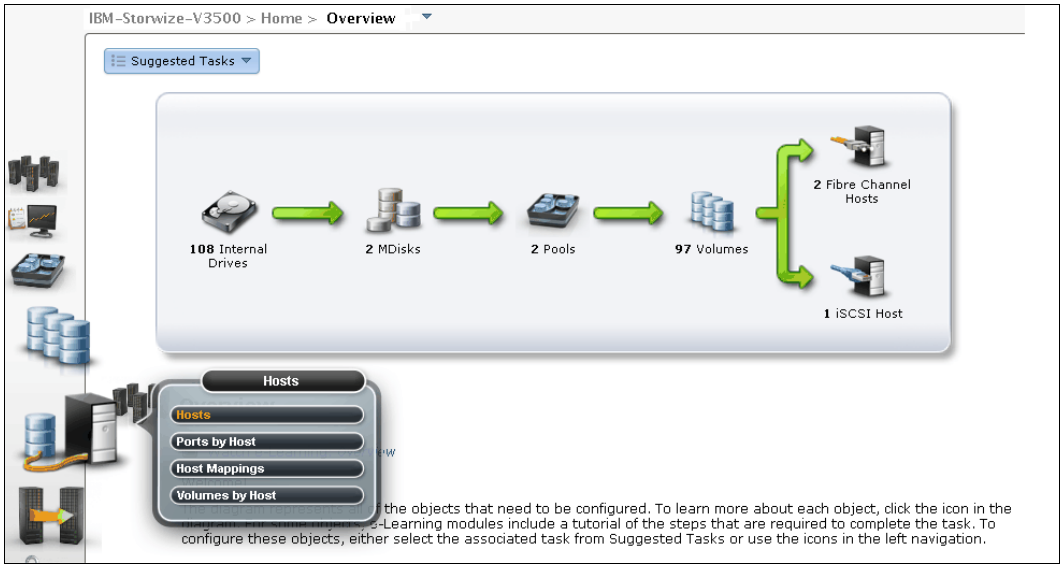

Figure 4-30 Configuration: Network

2. Select **iSCSI** and the iSCSI Configuration window opens, a shown in Figure 4-31.

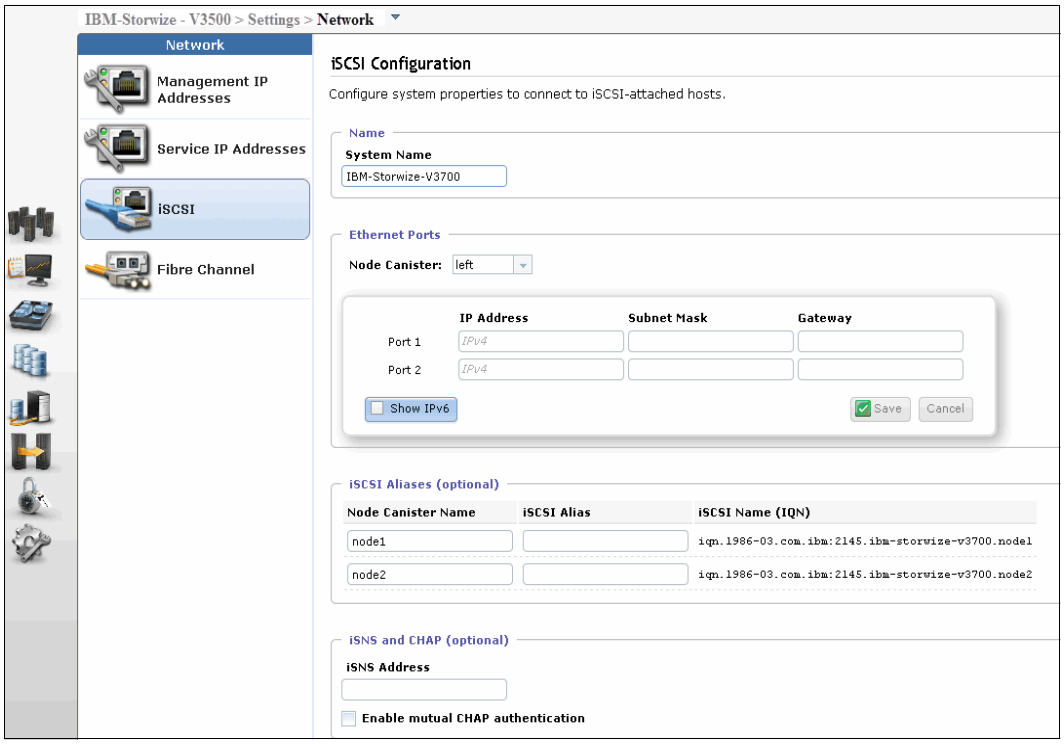

Figure 4-31 iSCSI Configuration window

In the configuration, you see an overview of all the iSCSI settings for the IBM Storwize V3500. You can configure iSCSI Alias, iSNS Addresses, and Chap Authentication Configuration on this window, and the iSCSI IP address, which we edit in the basic setup.

3. Click within the **Ethernet Ports** section and choose left or right to enter the iSCSI IP address, as shown in Figure 4-32. Repeat this step for each port you want to use for iSCSI traffic.

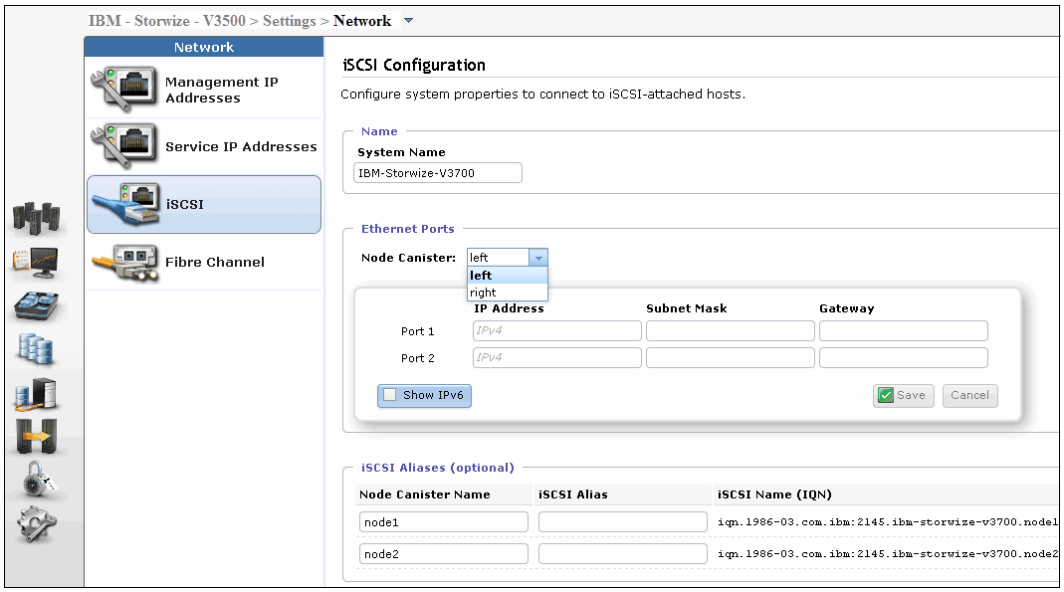

Figure 4-32 Enter an iSCSI IP address

4. After you enter the IP address for each port, click **Save Changes** to enable the configuration, as shown in Figure 4-33.

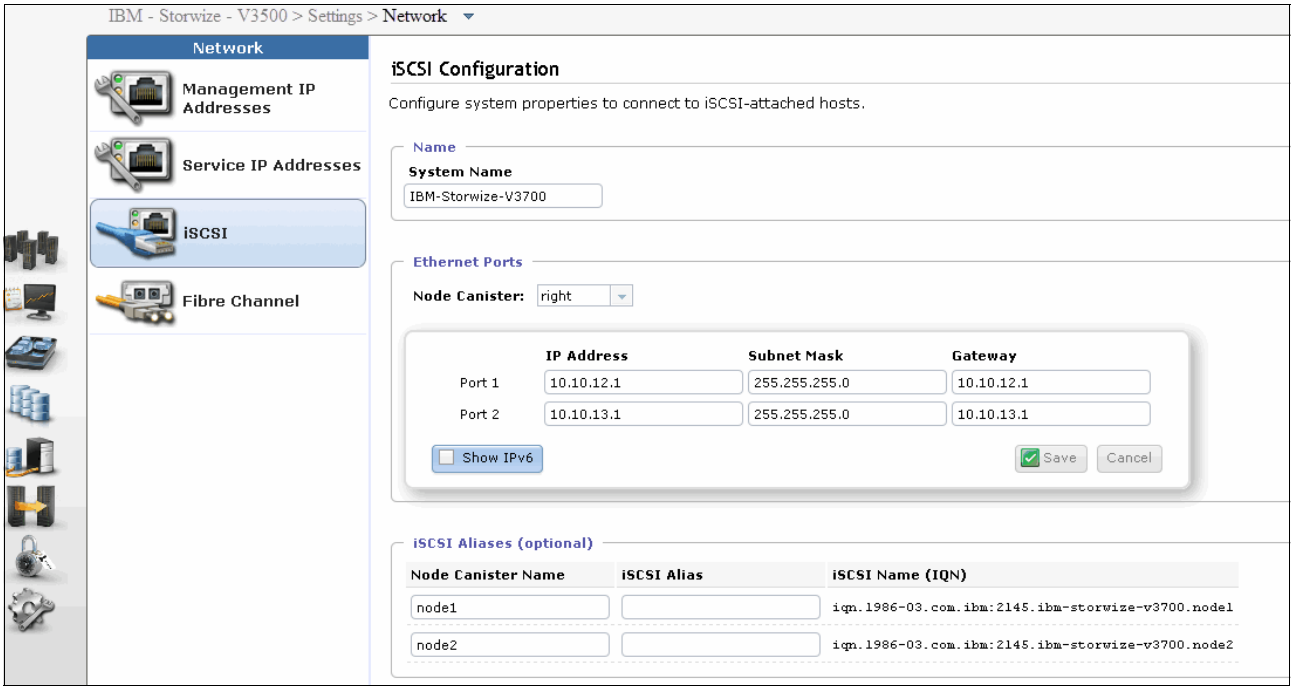

Figure 4-33 iSCSI IP configuration

Before your ESXi host can discover Storwize V3500 storage, the iSCSI initiator must be configured and authentication might have to be done.

You can verify the network configuration by using the `vmkping` utility. If you must authenticate the target, you might need to configure the dynamic discovery address and target name of the Storwize V3500 in vSphere, as shown in Figure 4-34.

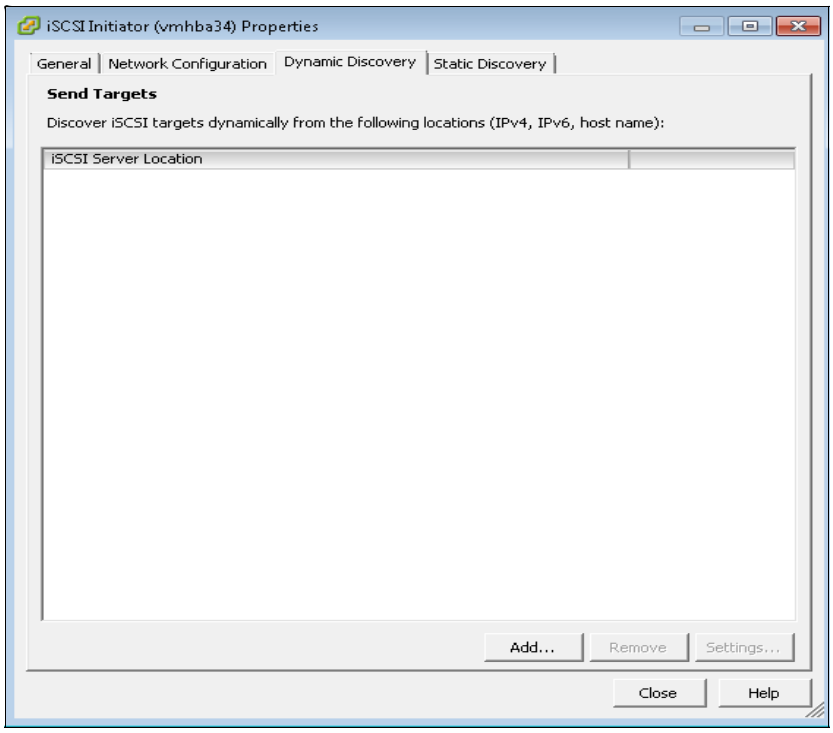

Figure 4-34 iSCSI Dynamic Discovery

Choose to Add Send Target Server, as shown in Figure 4-35.

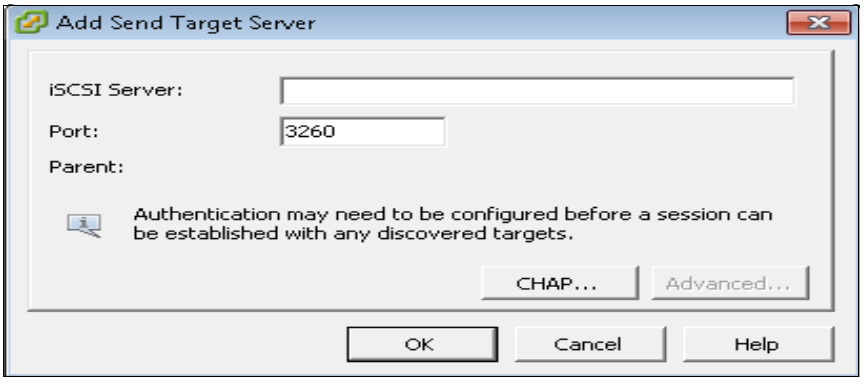

Figure 4-35 iSCSI add Target

Enter the IP address of the target, as shown in Figure 4-36. The target IP address is the IP address of a node in the I/O group from which you are mapping the iSCSI volume.

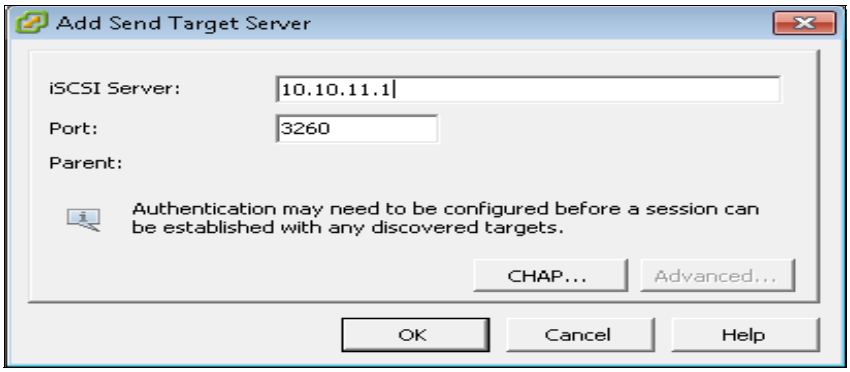

Figure 4-36 Add iSCSI IP of Target

Leave the IP port number at the default value of 3260. Click **OK** to establish the connection between the initiator and target. Repeat this process for each target IP. After you repeat this process and close the window, you are prompted to rescan the host bus adapters, as shown in Figure 4-37. The maximum number of volumes that can be mapped to the VMware iSCSI initiator is 256.

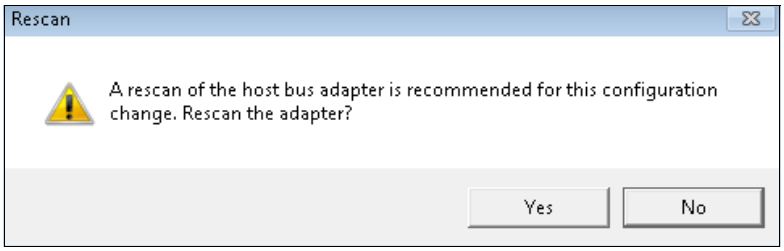

Figure 4-37 Rescan of host bus adapters

Important: iSNS discovery is not supported in VMware ESX 4.

The IBM Storwize V3500 is now configured and ready for iSCSI use. Note the initiator names of your storage canisters (as shown in Figure 4-31 on page 166) because you need them later. To create volumes and map them to a host, see Chapter 5, “Basic volume configuration” on page 171.

5

Basic volume configuration

This chapter describes how to use the IBM Storwize V3500 to create a volume and map a volume to a host. A volume is a logical disk on the IBM Storwize V3500 that is provisioned out of a storage pool. It 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 (see 5.1, “Provisioning storage from IBM Storwize V3500 and making it available to the host” on page 172). After you finish this chapter, your basic configuration is done and you can store data on the IBM Storwize V3500.

Advanced host and volume administration, such as adding and deleting host ports, and creating FlashCopy is described in Chapter 8, “Advanced host and volume administration” on page 317.

This chapter includes the following topics:

- Provisioning storage from IBM Storwize V3500 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 V3500 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. Complete the following steps to complete the basic setup of your environment:

1. Create volumes.
2. Map volumes to the host. For more information, see 5.2, “Mapping a volume to the host” on page 183.
3. Discover the volumes from the host and specify multipath settings. For more information, see 5.3, “Discovering the volumes from the host and specifying multipath settings” on page 187.

Open the All Volumes window of the IBM Storwize V3500 GUI (as shown in Figure 5-1) to start the process of creating volumes.

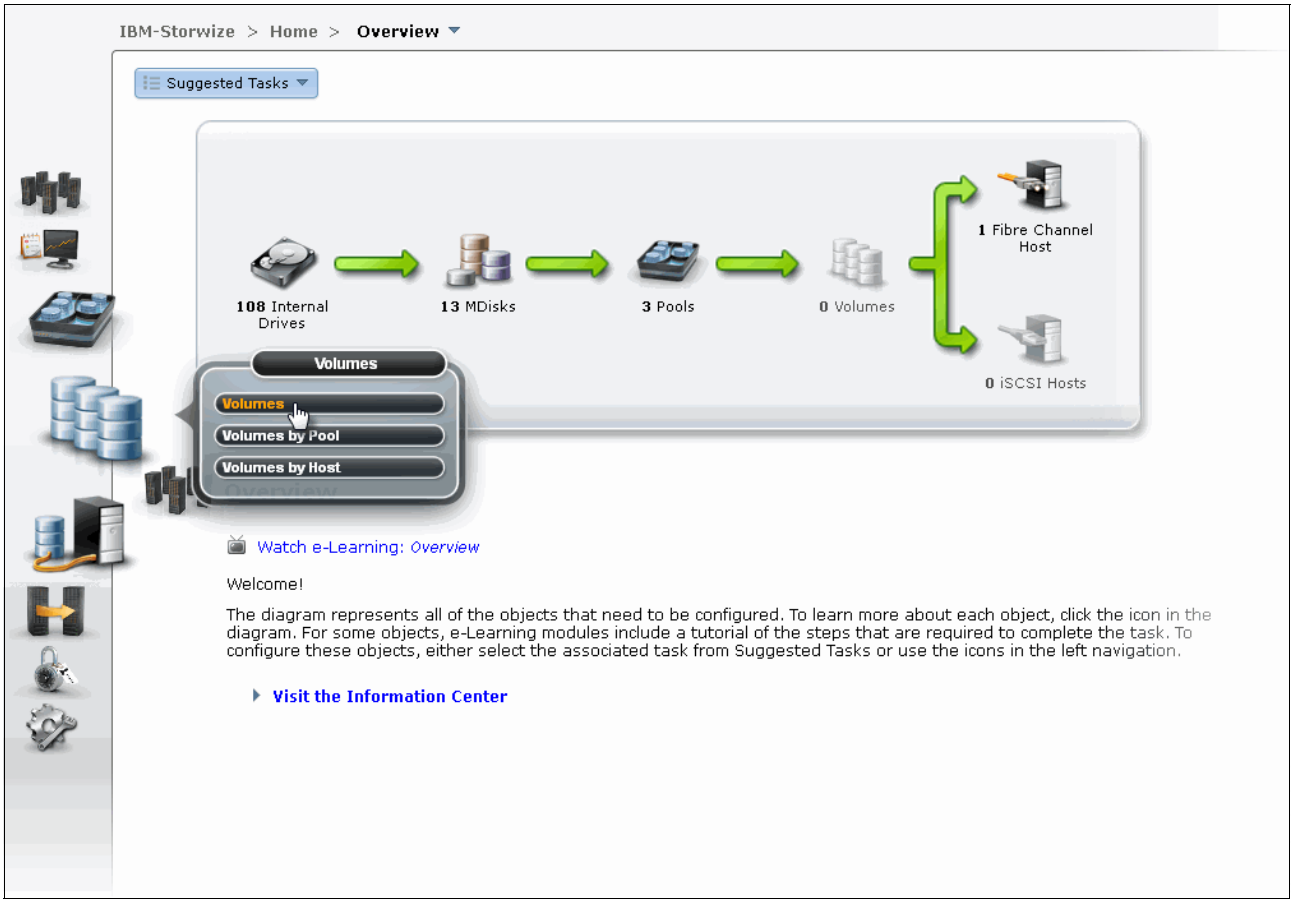

Figure 5-1 Home window

The All Volumes window opens, as shown in Figure 5-2.

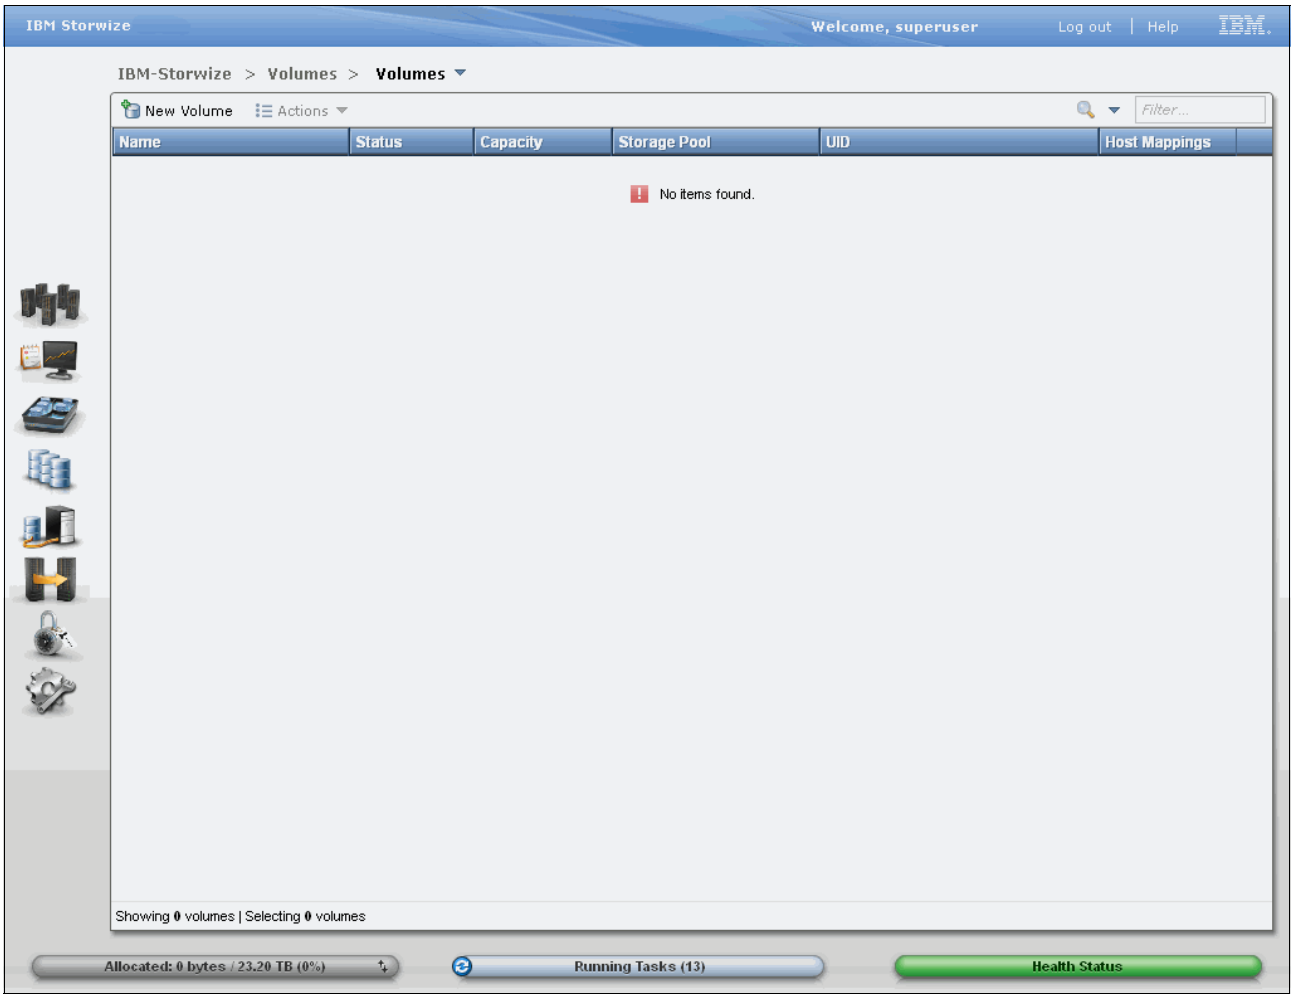

Figure 5-2 All Volumes window

At the moment, no volumes were created, so click **New Volume**, as shown in Figure 5-3.

Figure 5-3 New Volume

The New Volume window opens, as shown in Figure 5-4.

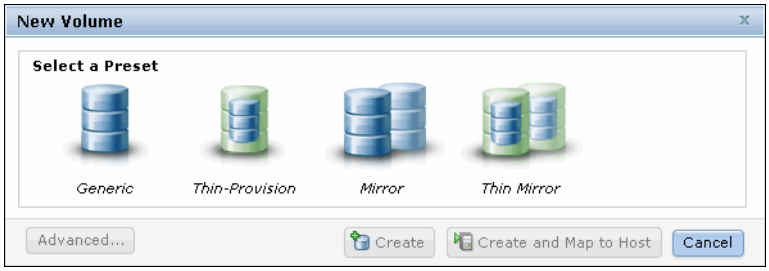

Figure 5-4 New Volume window

By default, all volumes that you create are striped across all available MDisks in one storage pool. The GUI for the IBM Storwize V3500 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 174.
- Thin-provisioned: A striped volume that is space efficient. There are choices under the Advanced option to help determine how much space is fully allocated initially and how large the volume is able to grow, as described in 5.1.2, “Creating a thin-provisioned volume” on page 176.
- 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 178.
- Thin-mirror: Two synchronized copies. Both are thin provisioned, as described in 5.1.4, “Creating a thin-mirror volume” on page 181.

Select the volume that you want to create. For more information, see the relevant section.

5.1.1 Creating a generic volume

The most commonly used type of volume is the generic volume type. This type of volume is fully provisioned, with the entire size of the volume that is dedicated to the volume. The host and the IBM Storwize V3500 see the fully allocated space without a mirror.

Complete the following steps to create a generic volume:

1. We choose a generic volume (as shown in Figure 5-4 on page 174) and then we select the pool in which the volume should be created. Select the pool by clicking it. In our example, we click pool that is called **mdiskgrp0**. The result is shown in Figure 5-5.

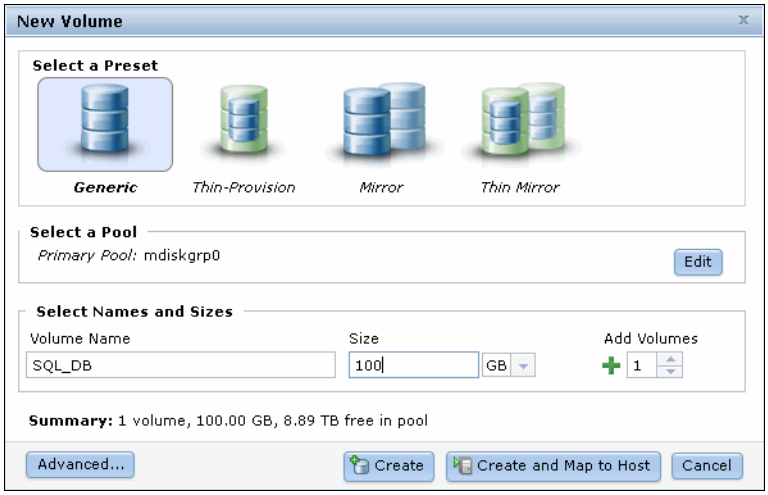

Figure 5-5 Create a generic volume

Important: The “Create and Map to Host” option is disabled if no host is already configured on the IBM Storwize V3500. To create a host, see Chapter 4, “Host configuration” on page 141.

2. Enter a volume name and size and click **Create and Map to Host**. The new Generic Volume is created (as shown in Figure 5-6). Click **Continue** and see 5.2.1, “Mapping newly created volumes to the host by using the wizard” on page 183.

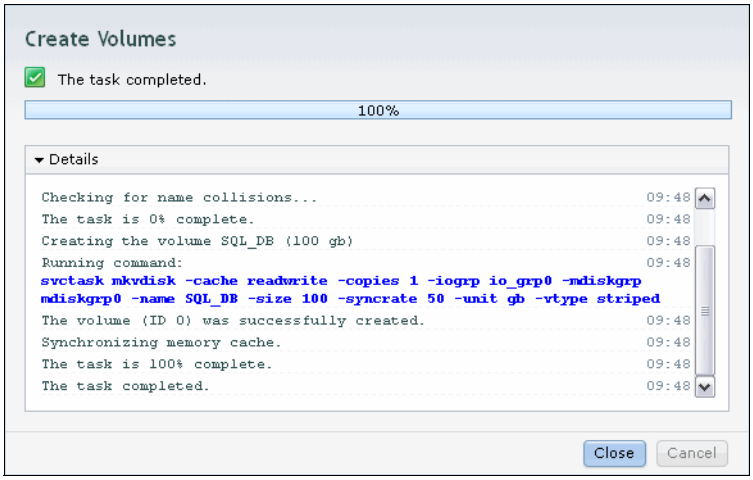

Figure 5-6 Volume created

If you do not want to map the volumes now, click **Create** (as shown in Figure 5-5) to complete the task. Volumes also can be mapped later, as described in 5.2.2, “Manually mapping a volume to the host” on page 184.

5.1.2 Creating a thin-provisioned volume

Volumes can be configured to be thin-provisioned. A thin-provisioned volume behaves concerning application reads and writes as though they were fully allocated. When you are creating a thin-provisioned volume, 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 5-7.

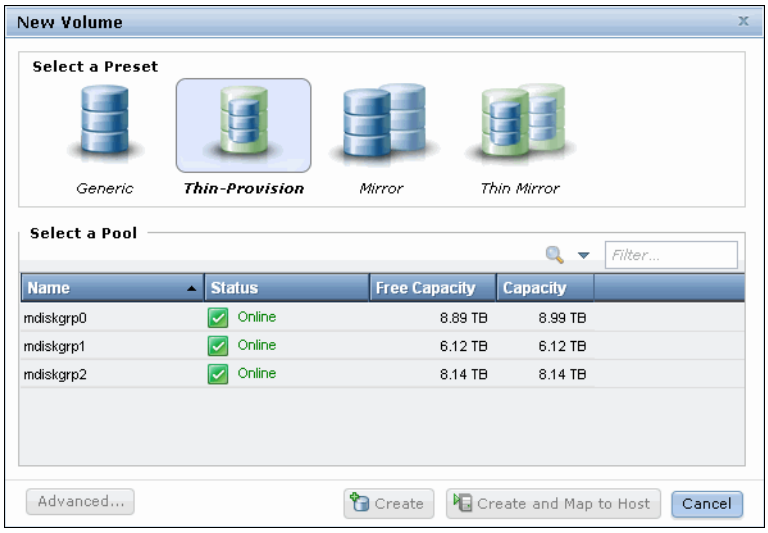

Figure 5-7 Create a thin-provisioned volume

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 click **mdiskgrp0**. The result is shown in Figure 5-8.

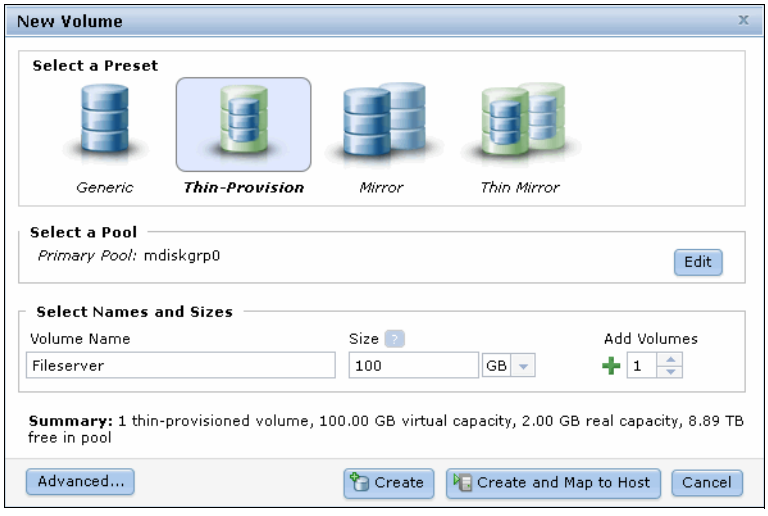

Figure 5-8 Enter the volume name and size

Under the Volume Name field is a summary that shows that you are about to make a thin-provisioned volume, how much virtual space is available, the space that is allocated (real size), and the free capacity in the pool. By default, the real capacity is 2% of the virtual capacity; you can change this setting by clicking **Advanced**.

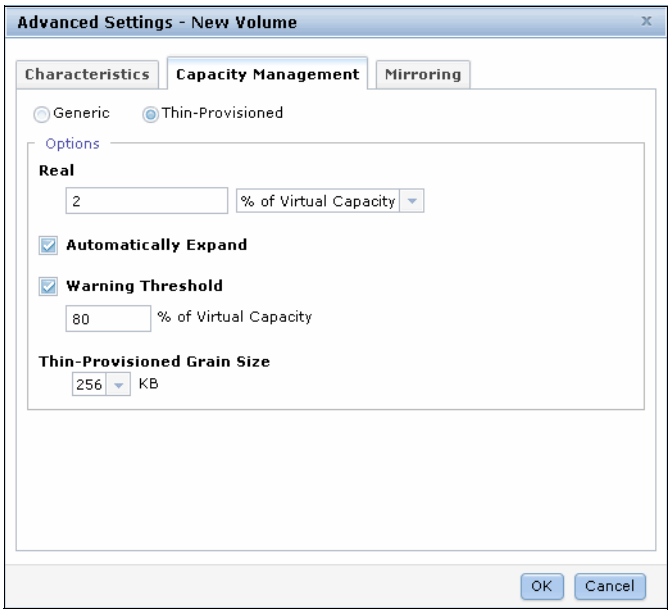

Figure 5-9 Advanced Settings

On the Thin Provisioning tab (as shown in Figure 5-9), the following advanced options are available:

- Real: Specify the size of the real capacity space that is used during creation.
- Automatically Extend: This option enables the automatic expansion of real capacity, if new capacity must be allocated.
- Warning Threshold: Enter a threshold for receiving capacity alerts.
- Thin-Provisioned Grain Size: Specify the grain size for real capacity.

3. Make your choices, if required, and click **OK** to return to the New Volume window.
4. Click **Create and Map to Host** and the creation task completes, as shown in Figure 5-10.

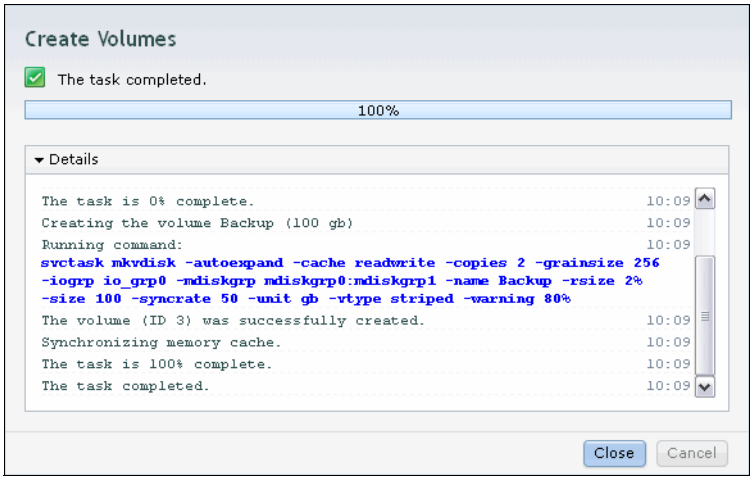

Figure 5-10 Create a thin-provisioned volume

If you do not want to map the volumes now, click **Create** (see Figure 5-8 on page 176) to complete the task. Volumes also can be mapped later, as described in 5.2.2, “Manually mapping a volume to the host” on page 184.

5. Click **Continue** and see 5.2.1, “Mapping newly created volumes to the host by using the wizard” on page 183.

5.1.3 Creating a mirrored volume

IBM Storwize V3500 offers the capability to mirror volumes, which means a single volume is presented to the host. However, 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, such as LVM, 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, you can protect against array failures (for example, multiple disk failures). You also have access to more advanced features, as described in 8.6, “Advanced volume copy functions” on page 380. You also can improve 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.

To create a mirrored volume, complete the following steps:

1. Select **Mirror**, as shown in Figure 5-11.

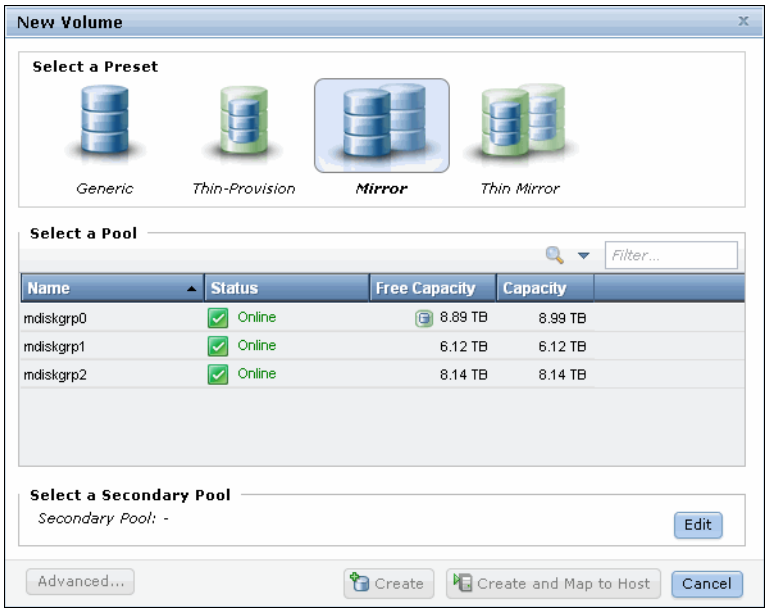

Figure 5-11 Create a mirrored volume

2. Select the primary pool by clicking it and the view changes to the secondary pool, as shown in Figure 5-12.

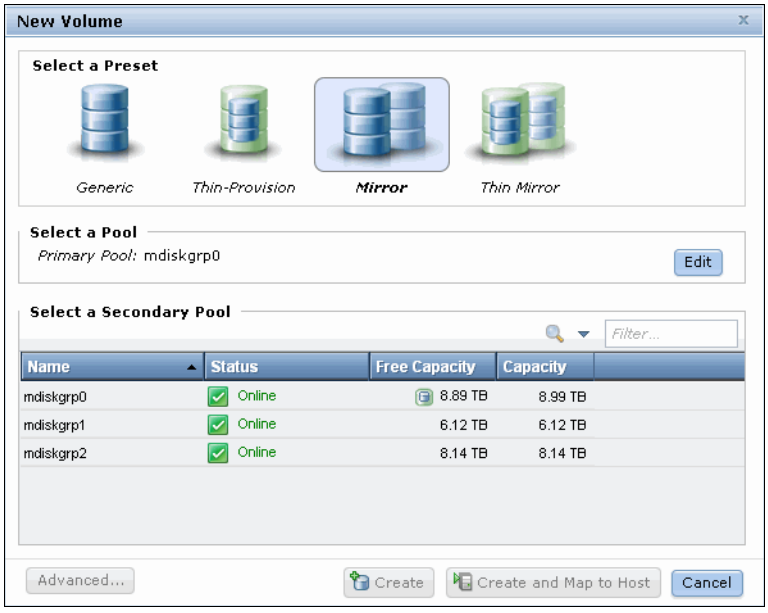

Figure 5-12 Select a pool

3. Select the secondary pool by clicking it and enter a volume name and the required size, as shown in Figure 5-13.

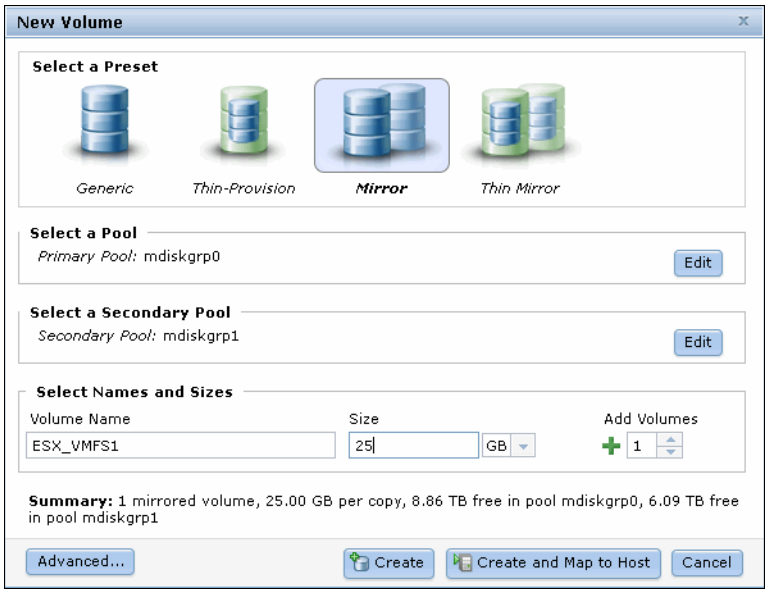

Figure 5-13 Select a pool, name, and size

Important: Before you create a mirrored volume, it is best to have at least two separate storage pools, and then click primary and secondary pool from the graphical user interface and use different pools for each copy. For more information about storage pools, see Chapter 7, “Storage pools” on page 277.

4. The summary shows you capacity information about the pool. If you want to select advanced settings, click **Advanced** and the click the **Mirroring** tab, as shown in Figure 5-14.

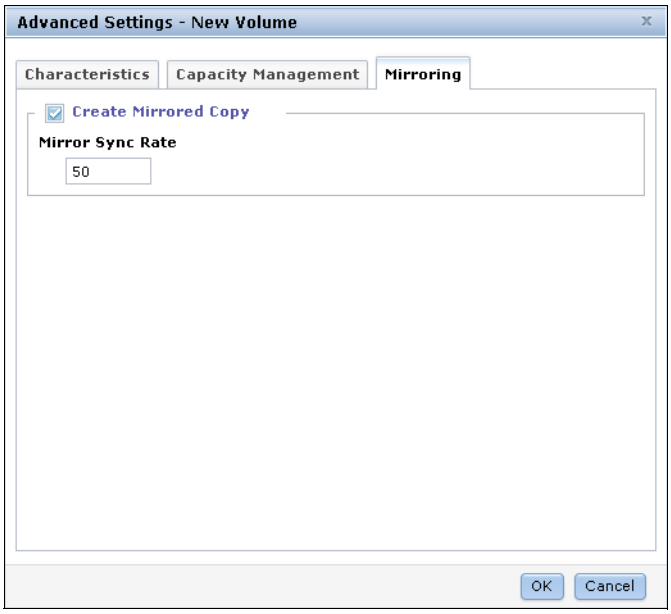

Figure 5-14 Advanced mirroring features

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, with which you can prefer more important volumes for synchronization faster than other mirrored volumes. By default, the rate is set to 50% for all volumes. Click **OK** to return to the New Volume window, as shown in Figure 5-13 on page 179.
6. Click **Create and Map to Host** and the mirrored volume is created, as shown in Figure 5-15.

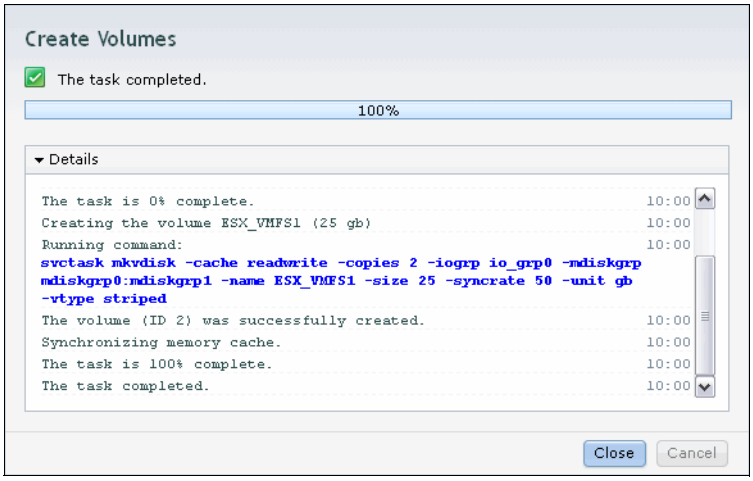

Figure 5-15 Create a mirrored volume

7. Click **Continue** and see 5.2.1, “Mapping newly created volumes to the host by using the wizard” on page 183.

5.1.4 Creating a thin-mirror volume

By using a thin-mirror volume, you can allocate the required space on demand (as described in 5.1.2, “Creating a thin-provisioned volume” on page 176). You also can have several copies of a volume available, as described in 5.1.3, “Creating a mirrored volume” on page 178).

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

1. Select **Thin Mirror**, as shown in Figure 5-16.

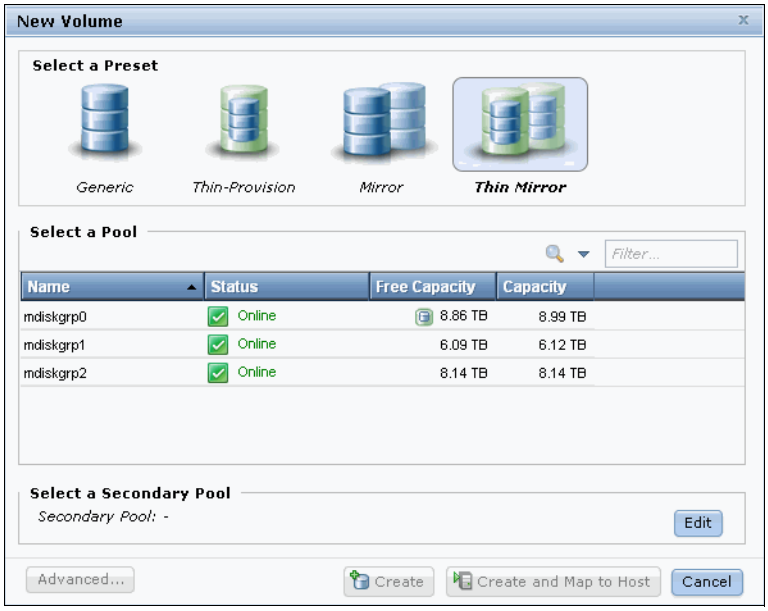

Figure 5-16 Thin Mirror

2. Select the primary pool by clicking it and the view the changes to the secondary pool, as shown in Figure 5-17.

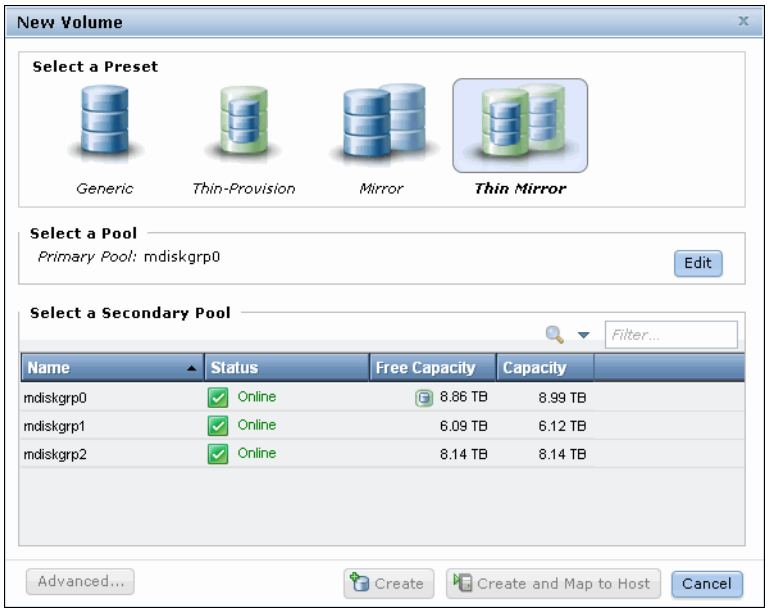

Figure 5-17 Select pools

3. Select the pool for the secondary copy and enter a name and a size for the new volume, as shown in Figure 5-18.

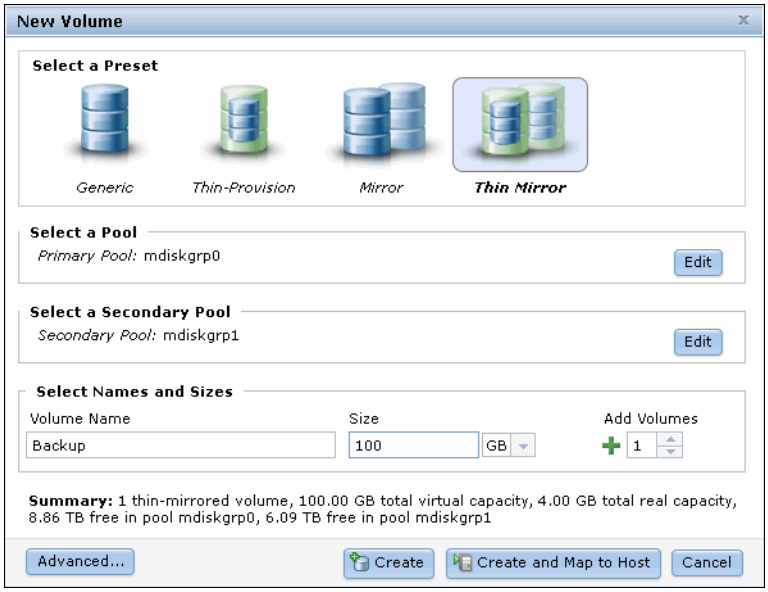

Figure 5-18 Enter a volume name and size

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 5-9 on page 177) or the mirror synchronization rate (as shown in Figure 5-14 on page 180). If you opened the advanced settings, click **OK** to return to the New Volume window, as shown in Figure 5-18.
5. Click **Create and Map to Host** and the mirrored volume is created, as shown in Figure 5-19.

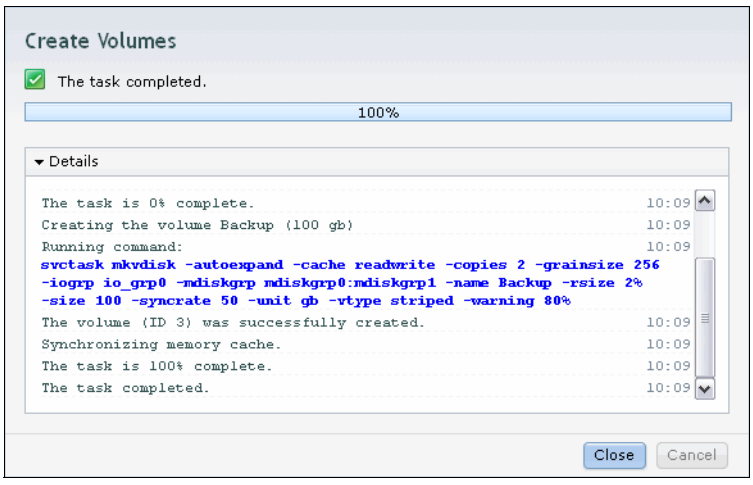

Figure 5-19 Create Thin Mirror Volume

6. Click **Continue** and see 5.2.1, “Mapping newly created volumes to the host by using the wizard” on page 183.

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 the **Create and Map to Host** option. The second part describes the manual host mapping process 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 V3500 and making it available to the host” on page 172. We assume that you followed the procedure and clicked the **Continue** option as shown in Figure 5-5 on page 175.

To map the volumes, complete the following steps:

1. Select a host to which the new volume should be attached, as shown in Figure 5-20.

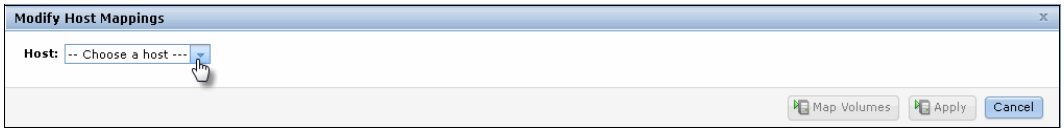

Figure 5-20 Choose a host

2. The **Modify Host Mappings** window opens. Your host and the newly created volume are already selected. Click **Map Volumes** and the volume is mapped to the host, as shown in Figure 5-21.

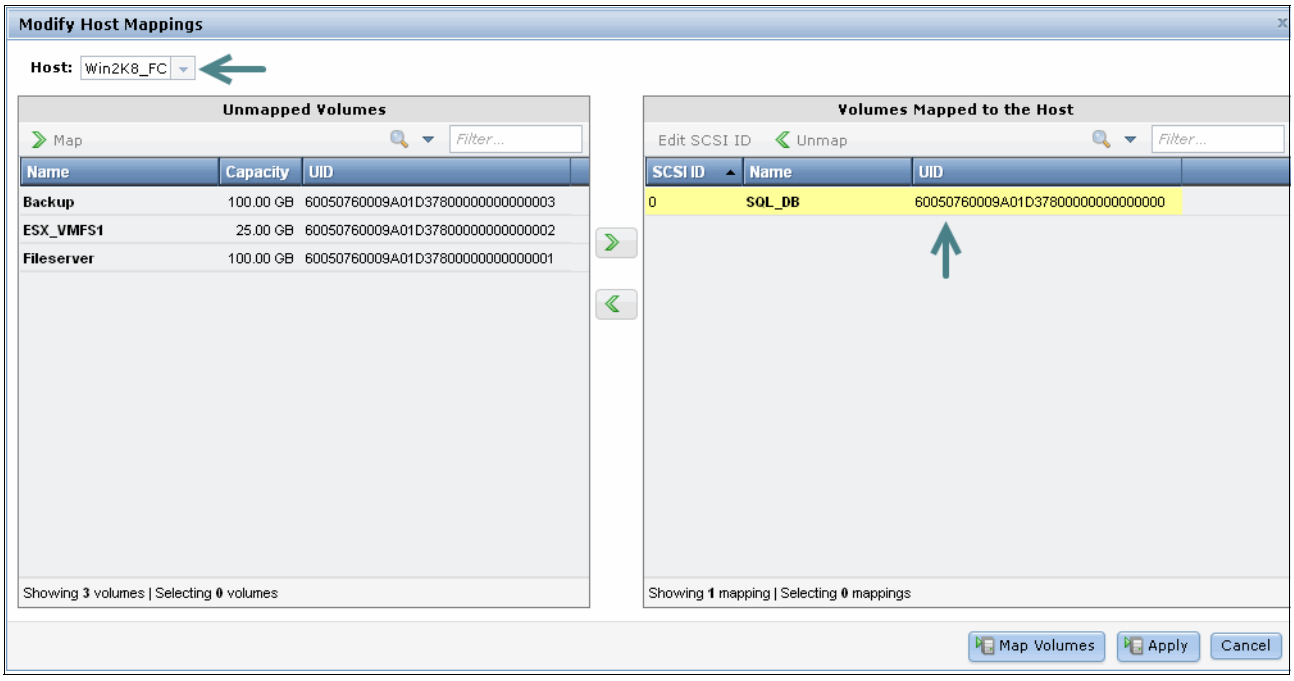

Figure 5-21 Modify host mappings

3. After the task completes, click **Close** (as shown in Figure 5-22), and the wizard returns to the All Volumes window.

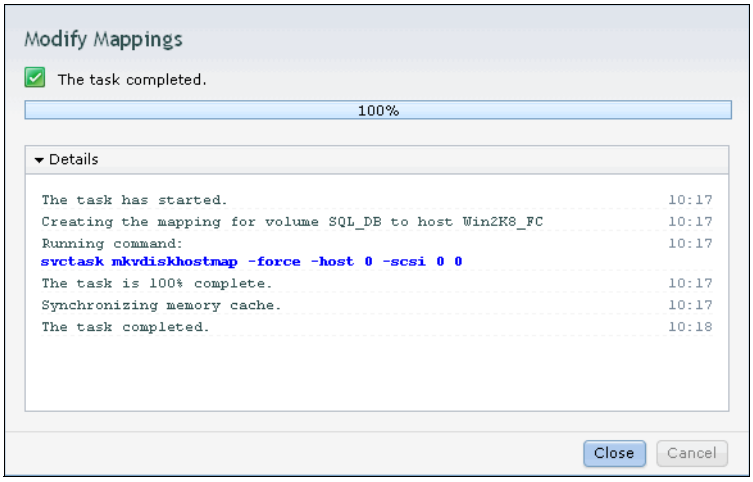

Figure 5-22 Modify Mappings complete

The newly created volume is displayed. We see that it is already mapped to a host, as shown in Figure 5-23.

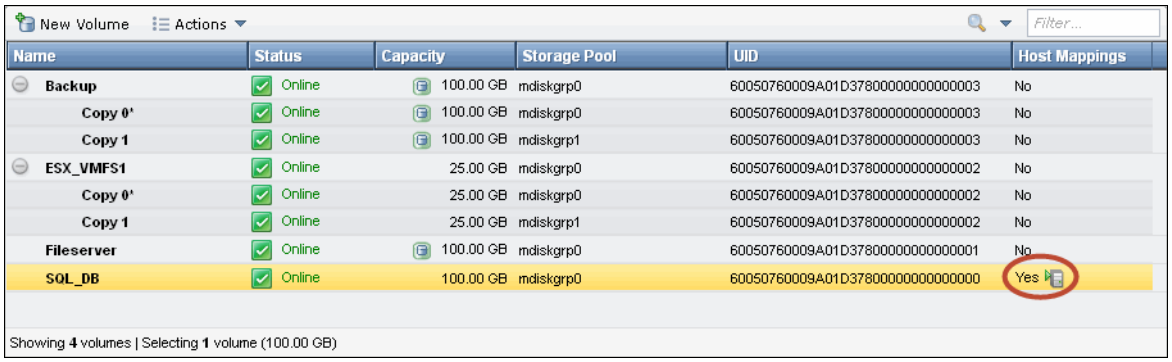

Figure 5-23 New Volume that is mapped to host

The host can access the volumes and store data on it. For more information about discovering the volumes on the host and making other host settings if required, see 5.3, “Discovering the volumes from the host and specifying multipath settings” on page 187.

You also can create multiple volumes in preparation for discovering them later. Mappings can be customized as well. Advanced host configuration is described in 8.1.1, “Modifying Mappings menu” on page 320.

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 V3500 and making it available to the host” on page 172 and clicked **Create** option, as shown in Figure 5-5 on page 175.

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

1. Open the All Hosts window, as shown in Figure 5-24.

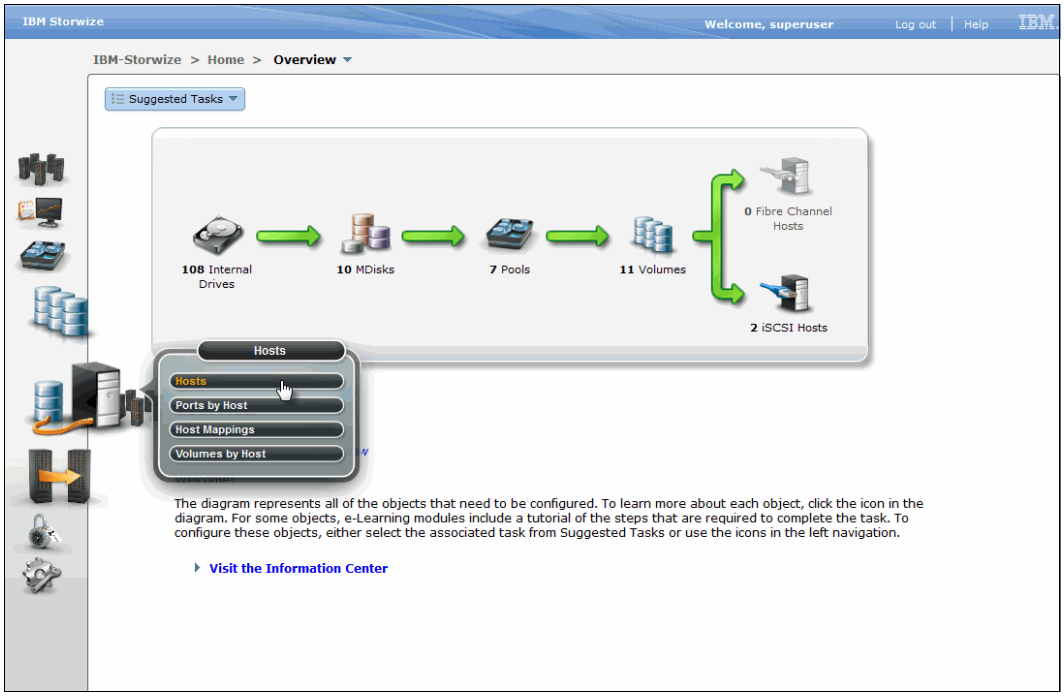

Figure 5-24 All Host window

2. Right-click the host to which a volume is to be mapped and select **Modify Mappings**, as shown in Figure 5-25.

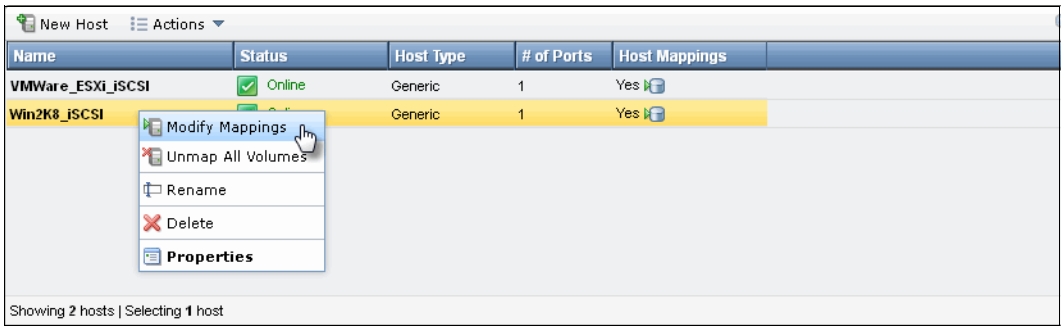

Figure 5-25 All hosts window

3. The Modify Host Mappings window opens. Select the volume that you want to map from the Unmapped Volumes pane on the left, as shown in Figure 5-26.

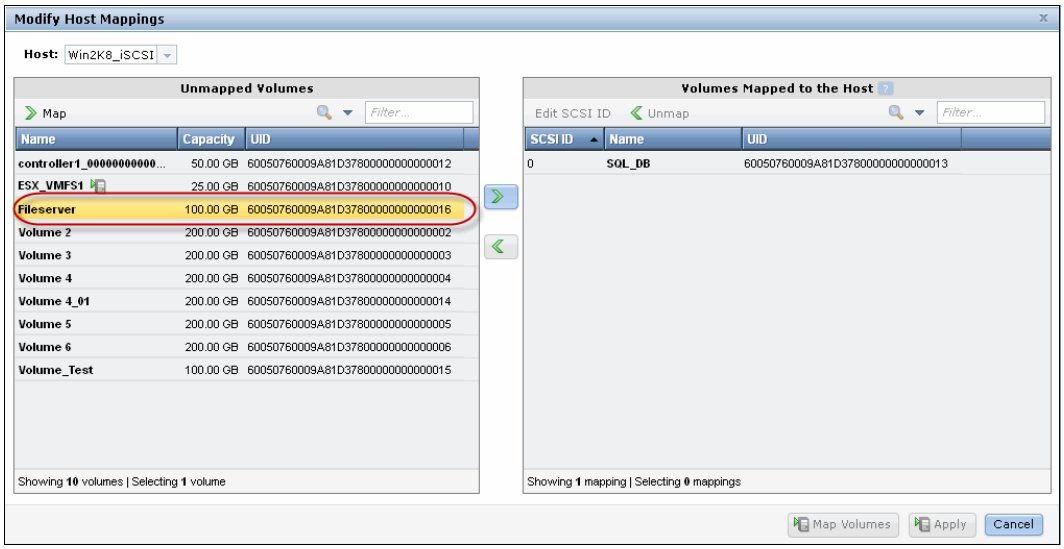

Figure 5-26 Modify host mappings window

4. Click the right-pointing arrow and the volume is moved to Volumes Mapped to the Host pane on the right side, as shown in Figure 5-27. Repeat this step for all of the volumes that you want to map and then click **Map Volumes**.

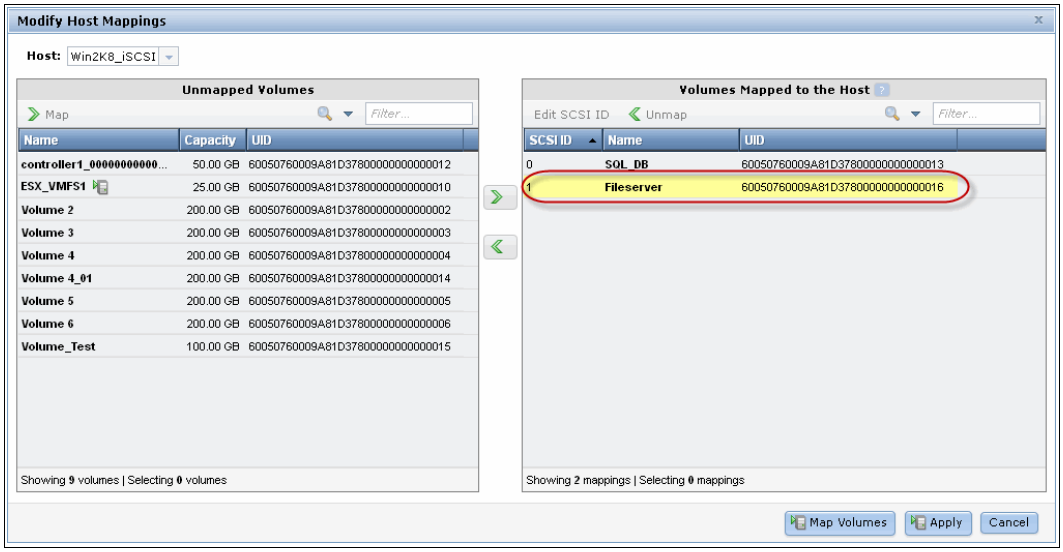

Figure 5-27 Modify host mappings window

5. After the task completes, click **Close**, as shown in Figure 5-28. The wizard returns to the All Hosts window.

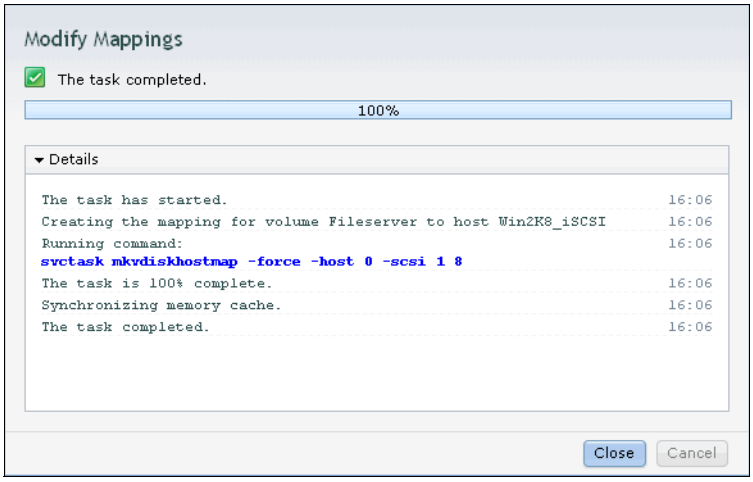

Figure 5-28 Modify mapping complete

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 making other host settings, see 5.3, “Discovering the volumes from the host and specifying multipath settings”.

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 V3500 and making it available to the host” on page 172 and 5.2, “Mapping a volume to the host” on page 183, and set other multipath settings, if required.

We assume that you completed all of the following steps that are described previously in this book so that the hosts and the IBM Storwize V3500 are prepared:

- Prepare your operating systems for attachment (for more information, see Chapter 4, “Host configuration” on page 141).
- Create hosts by using the GUI (for more information, see 4.3, “Creating hosts by using the GUI” on page 159).
- Perform basic volume configuration and host mapping (for more information, see 5.1, “Provisioning storage from IBM Storwize V3500 and making it available to the host” on page 172 and 5.2, “Mapping a volume to the host” on page 183).

This section describes how to discover Fibre Channel and iSCSI volumes from Windows 2008 and VMware ESX 5.x hosts.

In the IBM Storwize V3500 GUI, click **Hosts**, as shown in Figure 5-29.

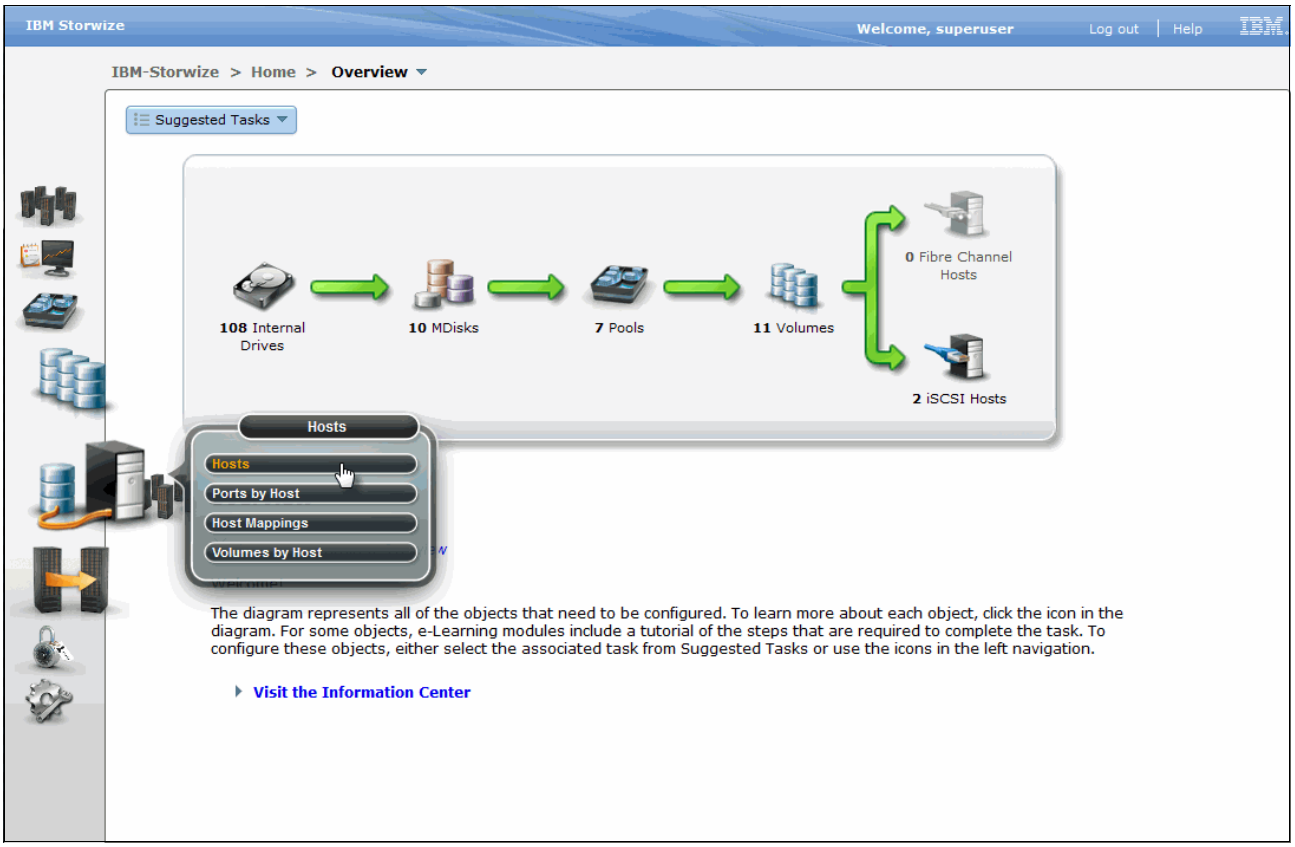

Figure 5-29 Open all hosts

The view that opens gives you an overview of the currently configured and mapped hosts. In our example, we have only one host that is configured, as shown in Figure 5-30.

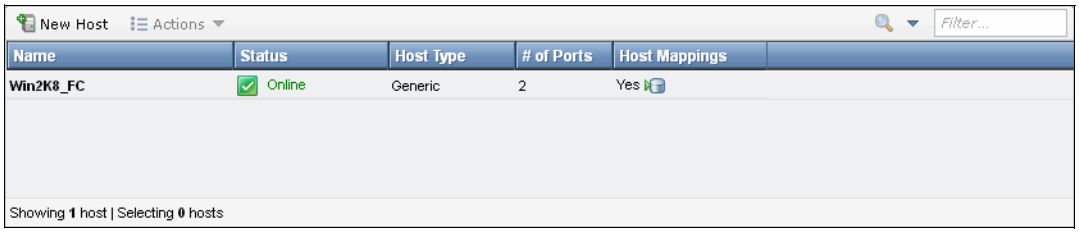

Figure 5-30 All Hosts view

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 All Hosts view (as shown in Figure 5-31) and select **Properties**.

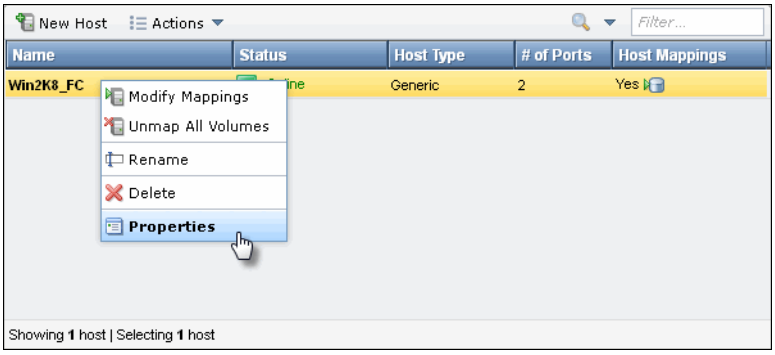

Figure 5-31 Host properties

2. Browse to the Mapped Volumes tab, as shown in Figure 5-32.

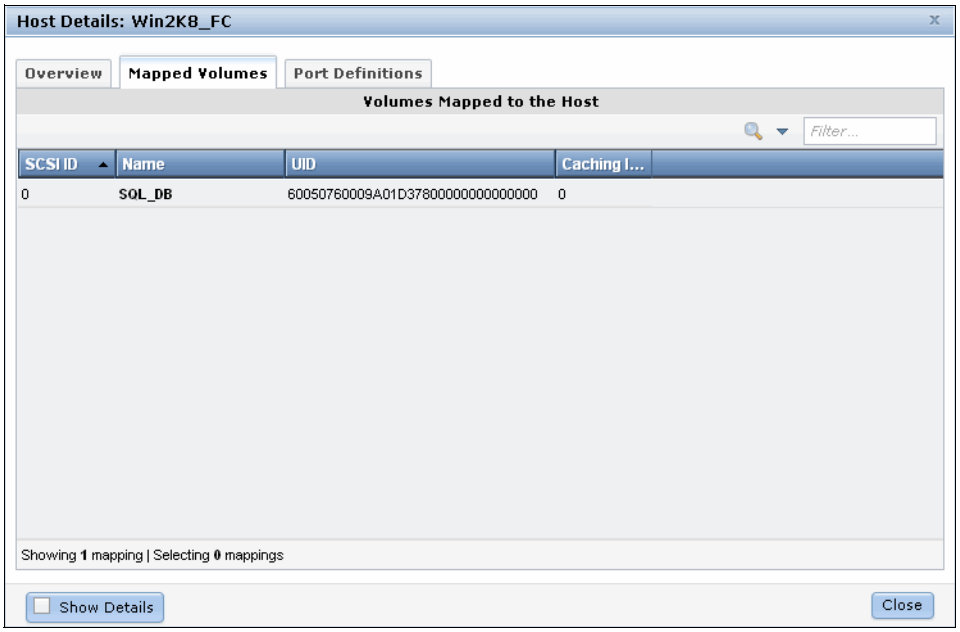

Figure 5-32 Mapped volumes to a host

The host details show you which volumes are currently 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. Log on to your Microsoft host and click **Start** → **All Programs** → **Subsystem Device Driver DSM** → **Subsystem Device Driver DSM**. A command-line interface opens. Enter **`datapath query device`** and press Enter to see whether there are IBM Storwize V3500 disks that are 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 : 1


DEV#:   0  DEVICE NAME: Disk1 Part0  TYPE: 2145       POLICY: OPTIMIZED
SERIAL: 60050760009A01D37800000000000000
============================================================================
Path#            Adapter/Hard Disk        State  Mode       Select     Errors
    0     Scsi Port3 Bus0/Disk1 Part0     OPEN   NORMAL         33          0
    1     Scsi Port4 Bus0/Disk1 Part0     OPEN   NORMAL          0          0

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

4. The output provides information about the connected volumes. In our example, there is one disk that is connected, Disk 1 for the Windows host, and two paths to the disk are available (State = Open). Open the Windows Disk Management window (as shown in Figure 5-33) by clicking **Start** → **Run**. Enter `diskmgmt.msc` and click **OK**.

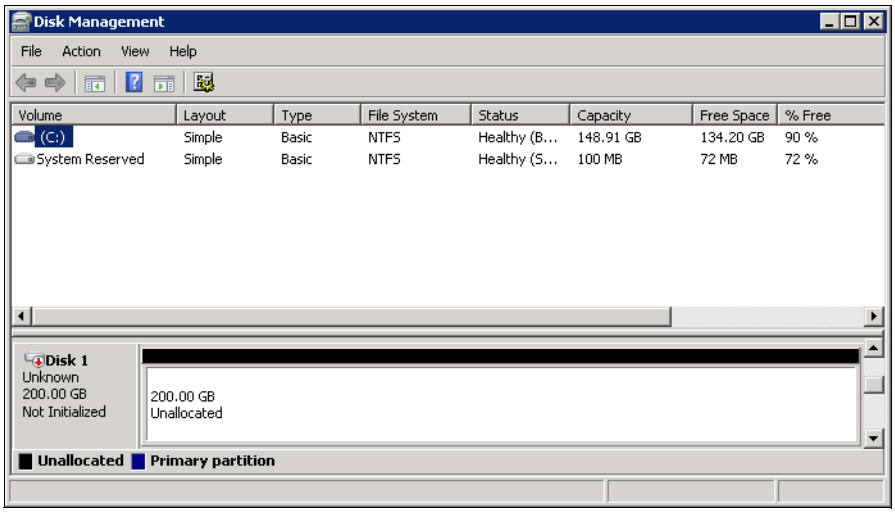

Figure 5-33 Windows Disk Management

Windows device discovery: Windows often automatically discovers new devices, such as disks. If you complete all of the steps and do not see any disks, click **Actions** → **Rescan Disk** in Disk Management to discover the new volumes.

5. Right-click the disk in the left pane and select **Online** if the disk is not online already, as shown in Figure 5-34.

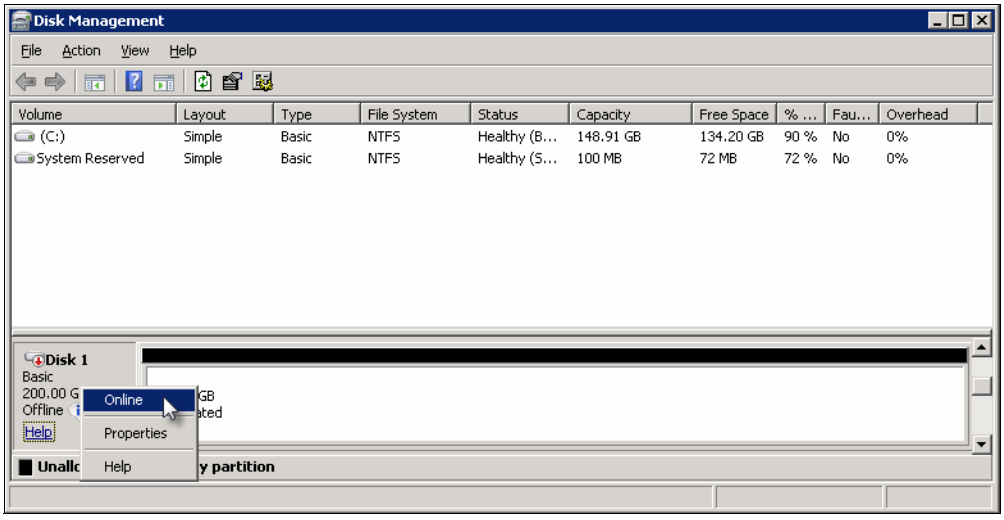

Figure 5-34 Place a disk online

6. Right-click the disk again and then click **Initialize Disk**, as shown in Figure 5-35.

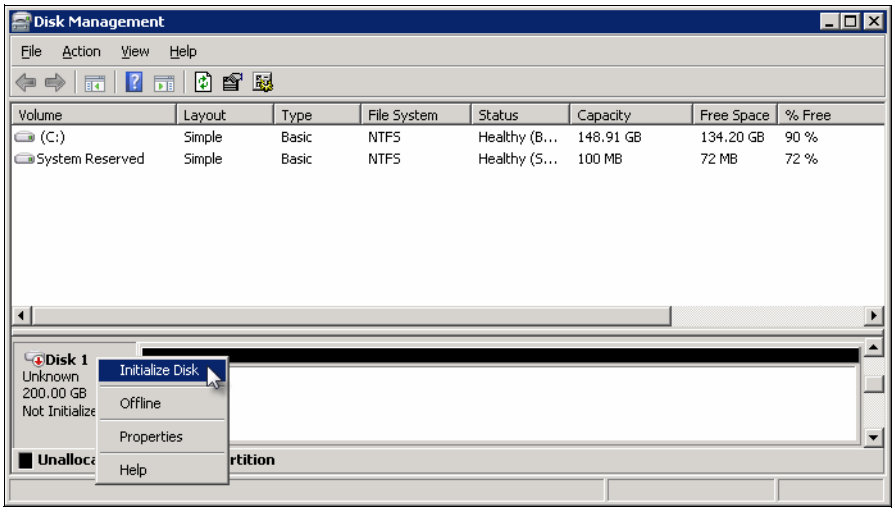

Figure 5-35 Initialize Disk menu

7. Select an initialization option and click **OK**. In our example, we selected MBR, as shown in Figure 5-36.

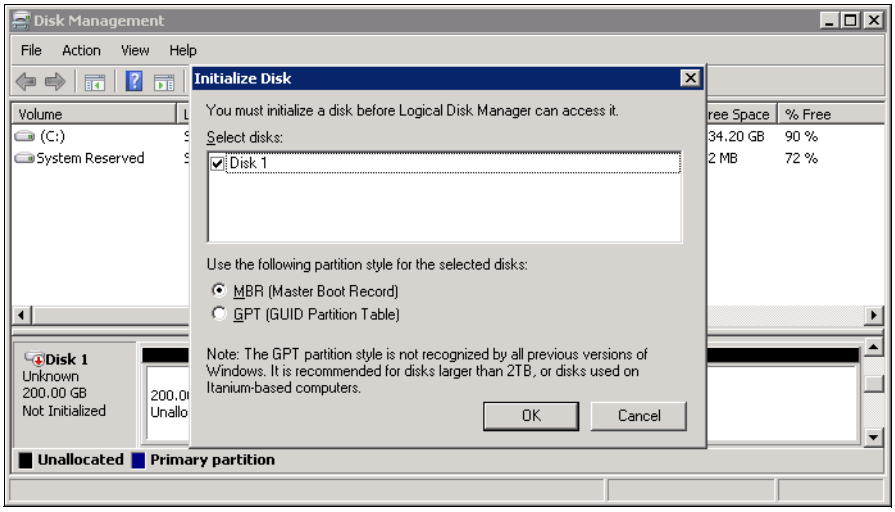

Figure 5-36 Initialize Disk

8. Right-click the right side pane and click **New Simple Volume**, as shown in Figure 5-37.

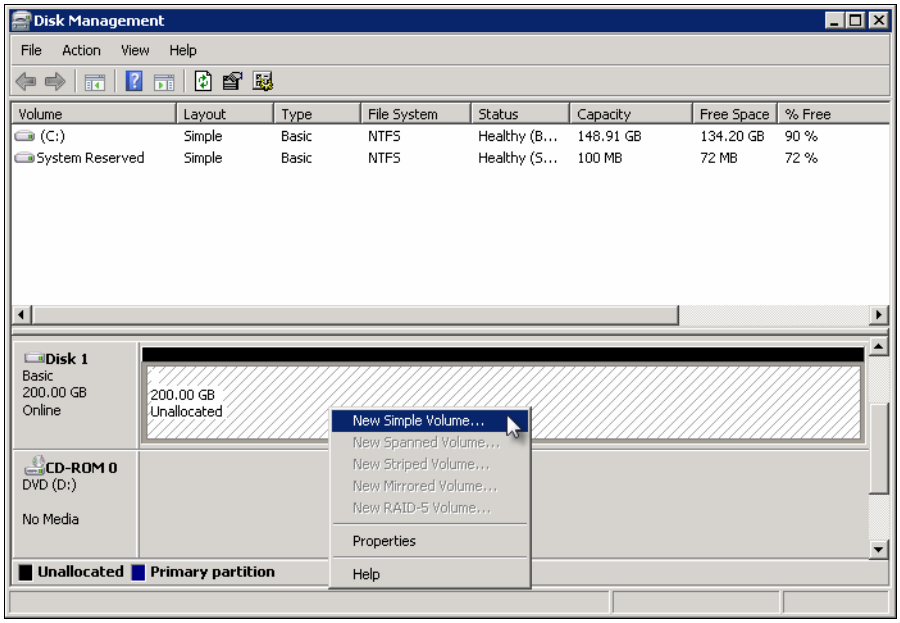

Figure 5-37 New Simple Volume

9. Follow the wizard and the volume is ready to use from your Windows host, as shown in Figure 5-38. In our example, we mapped a 200-GB disk (which is generic on the IBM Storwize V3500) to a Windows 2008 host by using Fibre Channel.

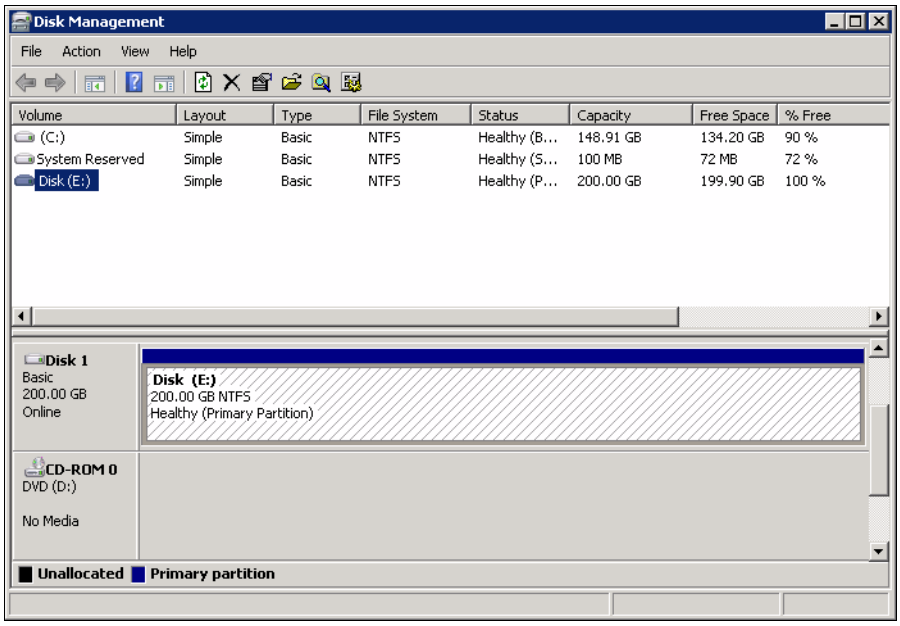

Figure 5-38 Volume is ready to use

The basic setup is now complete. The IBM Storwize V3500 is configured, and the host is prepared to access the volumes over several paths. The hose also can store data on the storage subsystem.

5.3.2 Windows 2008 iSCSI volume attachment

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

1. Right-click your Windows 2008 iSCSI host in the All Hosts view, as shown in Figure 1.

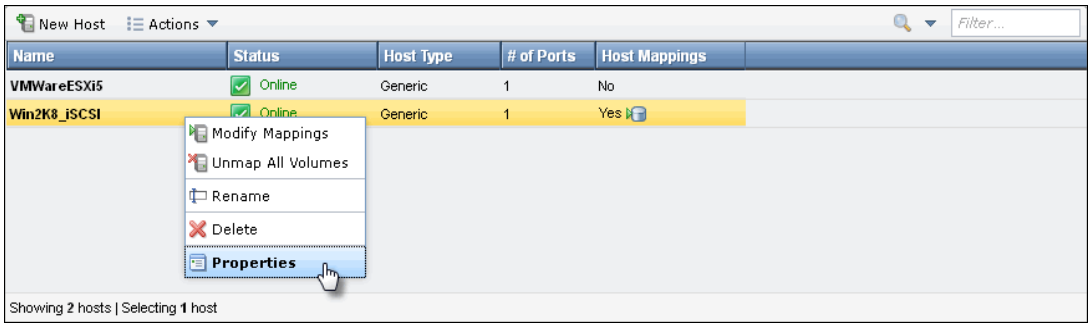

Click **Properties** and then browse to the Mapped Volumes tab, as shown in Figure 5-39.

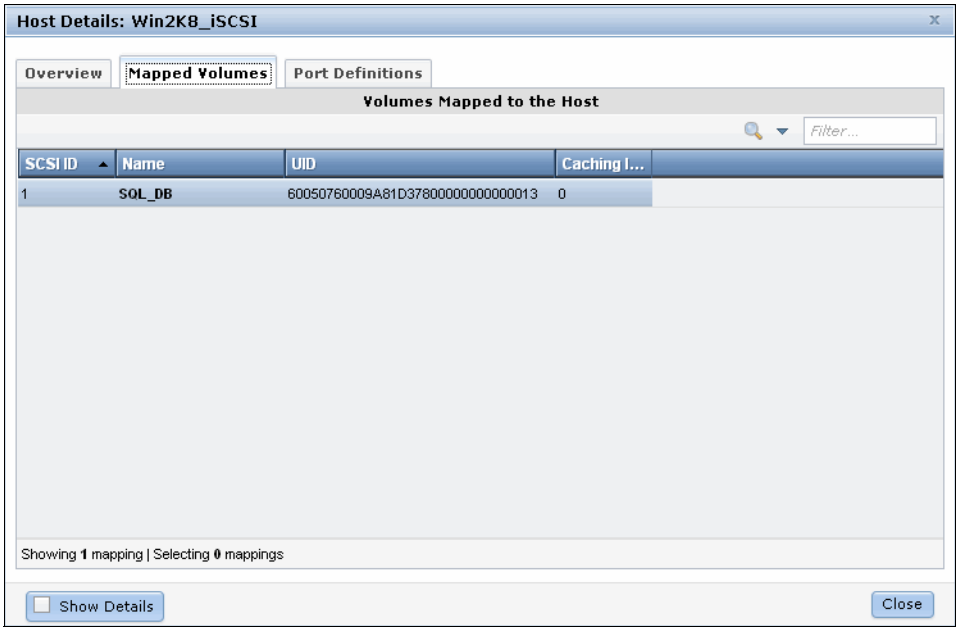

Figure 5-39 Mapped volumes to an iSCSI host

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 1 is mapped to the host.

2. Log on to your Windows 2008 host and click **Start** → **Administrative Tools** → **iSCSI Initiator** to open the iSCSI Configuration tab, as shown in Figure 5-40.

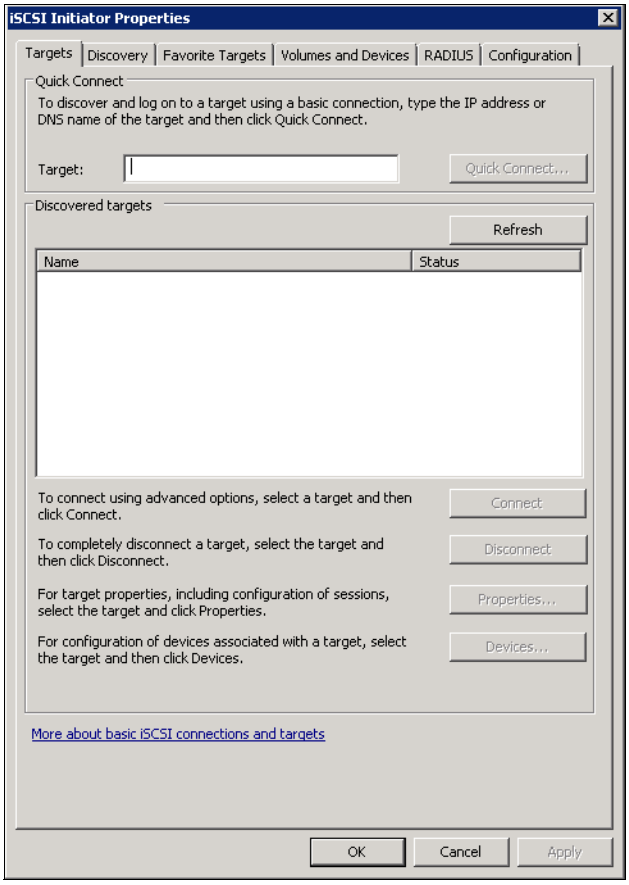

Figure 5-40 Windows iSCSI Configuration tab

3. Enter the IP address of one of the IBM Storwize V3500 iSCSI ports and click **Quick Connect**, as shown in Figure 5-41.

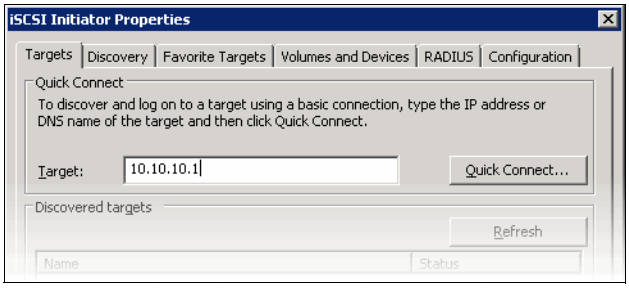

Figure 5-41 iSCSI Quick Connect

iSCSI IP addresses: The iSCSI IP addresses are different for the cluster and canister IP addresses. They were configured in 4.3.2, “Creating iSCSI hosts” on page 164.

The IBM Storwize V3500 initiator is discovered and connected, as shown in Figure 5-42.

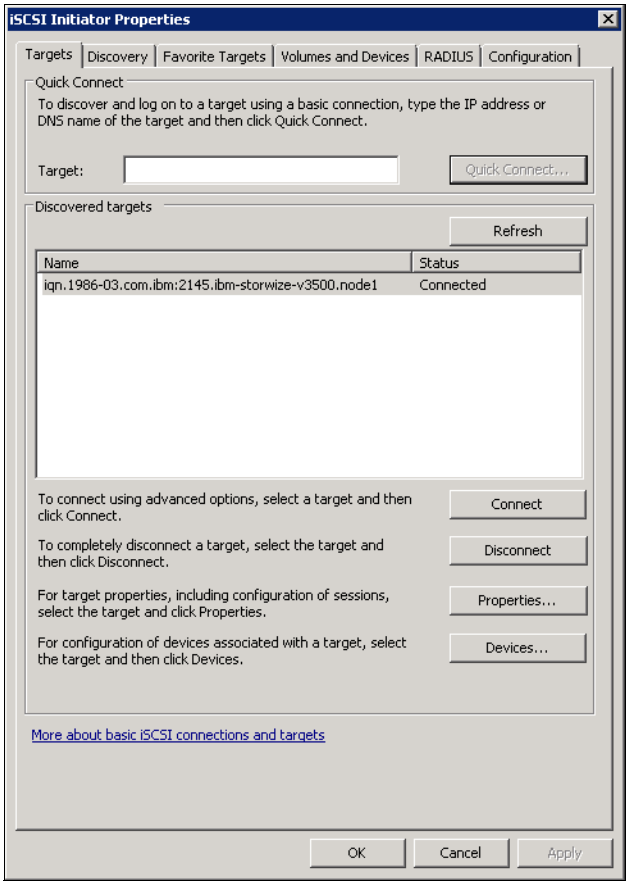

Figure 5-42 iSCSI Initiator target is connected

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

1. Click **Start** → **Run** and enter `cmd` to open a command prompt. Enter **`ServerManagerCMD.exe -install Multipath-IO`** and press Enter, as shown in Example 5-2.

Example 5-2 Installing MPIO

```
C:\Users\Administrator>ServerManagerCmd.exe -Install Multipath-IO

Start Installation...
[Installation] Succeeded: [Multipath I/O] Multipath I/O.
<100/100>

Success: Installation succeeded.
```

2. Click **Start** → **Administrative Tools** → **MPIO**. Click the Discover Multi-Paths tab and then select the **Add support for iSCSI devices** option, as shown in Figure 5-43.

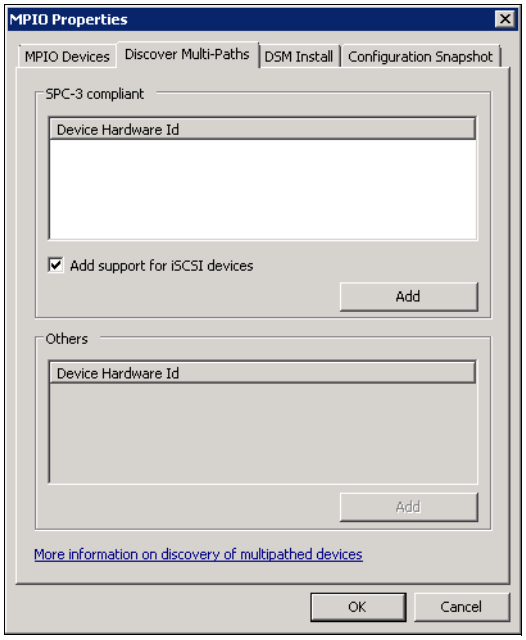

Figure 5-43 Enable iSCSI MPIO

Important: In some cases, the Add support for iSCSI devices option is unavailable. To enable it, you must connect to at least one iSCSI device, as we did in step 3 of this procedure.

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

4. After the reboot, log on again, click **Start** → **Administrative Tools** → **iSCSI Initiator** to open the iSCSI Configuration tab. Browse to the Discovery tab, as shown in Figure 5-44.

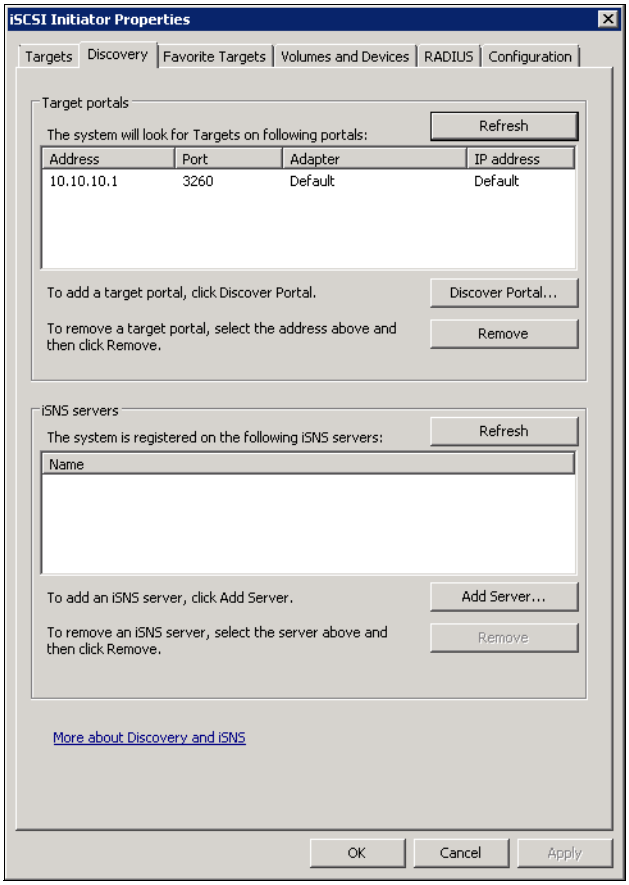

Figure 5-44 iSCSI Properties Discovery tab

5. Click **Discover Portal...**, and enter the IP address of another IBM Storwize V3500 iSCSI port, as shown in Figure 5-45. Click **OK**.

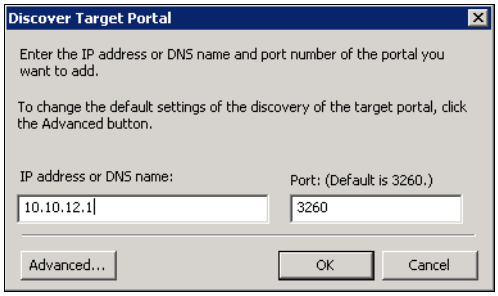

Figure 5-45 Discover Target Portal window

6. Return to the Targets tab (as shown in Figure 5-46) and you find the new connection is listed as Inactive.

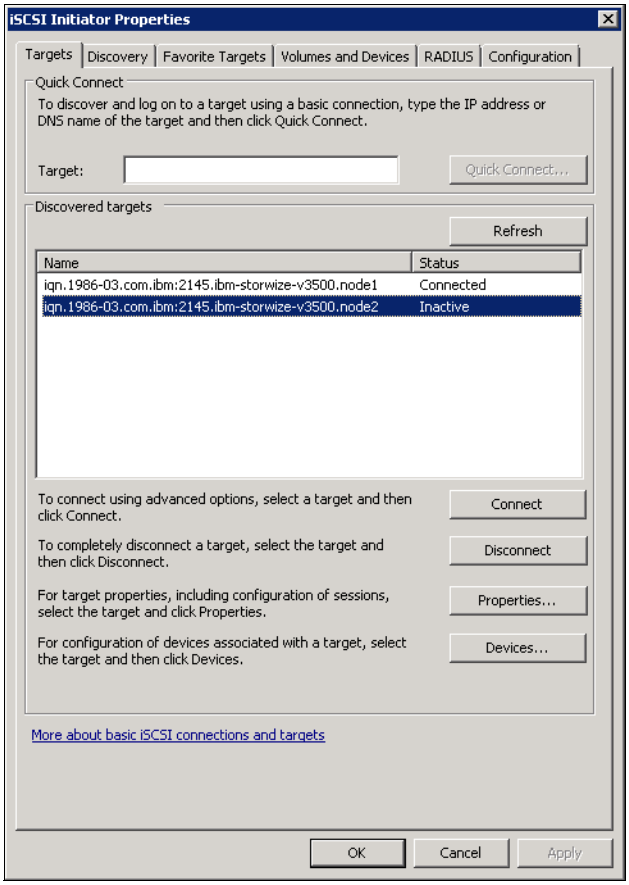

Figure 5-46 Inactive target ports

7. Highlight the inactive port and click **Connect**. The Connect to Target window opens, as shown in Figure 5-47.

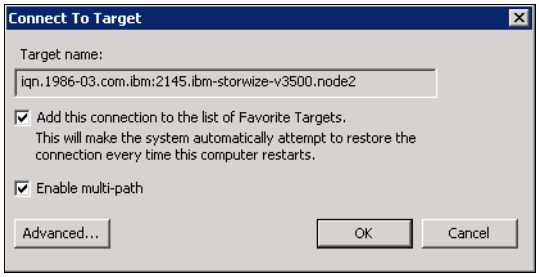

Figure 5-47 Connect to a target

8. Select the **Enable Multipath** option and click **OK**. The second port now appears as Connected, as shown in Figure 5-48.

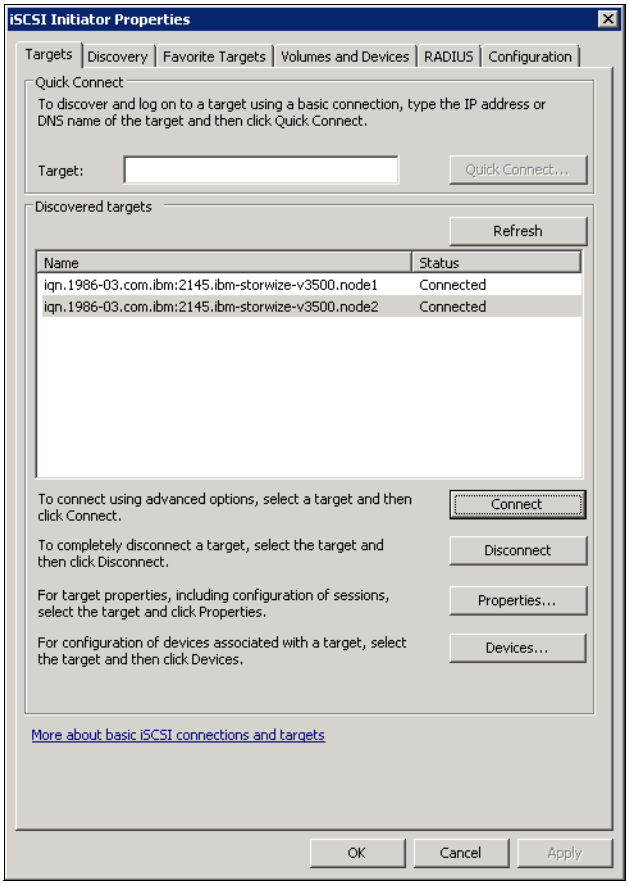

Figure 5-48 Second target port connected

Repeat this step for each IBM Storwize V3500 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 5-49) by clicking **Start** → **Run**. Enter `diskmgmt.msc` and click **OK**.

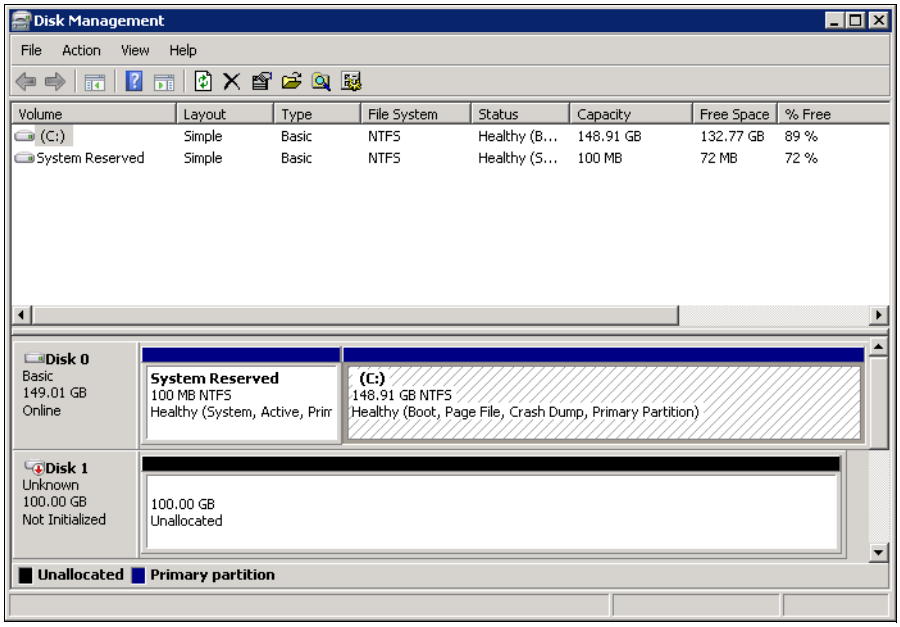

Figure 5-49 Windows Disk Management

10. Set the disk online, initialize it, create a file system on it, and it is ready to use. The steps that are used in this process are the same as the steps that are described in 5.3.1, “Windows 2008 Fibre Channel volume attachment” on page 189.

Now the storage disk is ready for use (as shown in Figure 5-50). In our example, we mapped a 100-GB disk (which is generic on the IBM Storwize V3500) to a Windows 2008 host by using iSCSI.

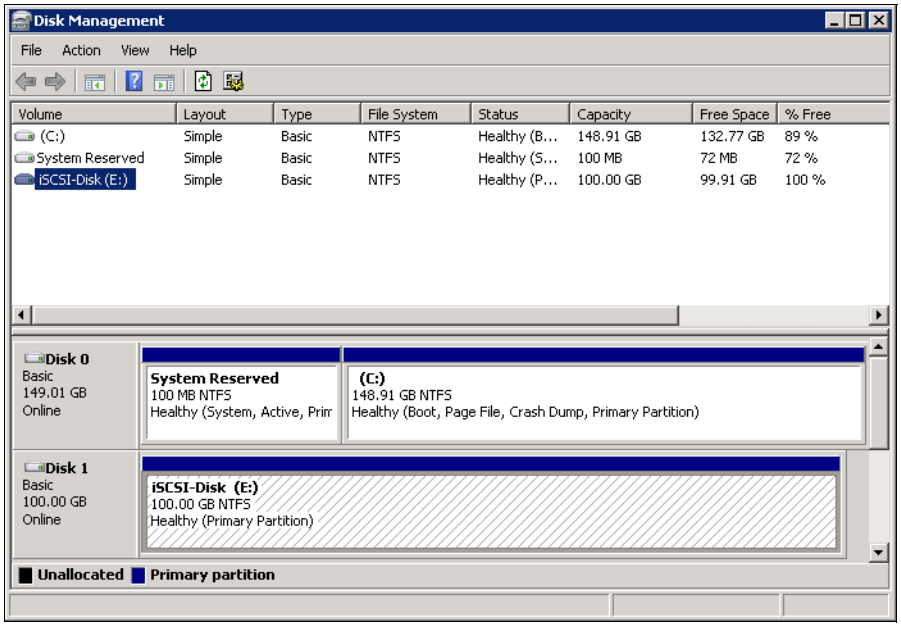

Figure 5-50 Windows Disk Management: Disk is ready to use

5.3.3 VMware ESX Fibre Channel attachment

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

1. Right-click your VMware ESX Fibre Channel host in the All Hosts view (as shown in Figure 5-51) and select **Properties**.

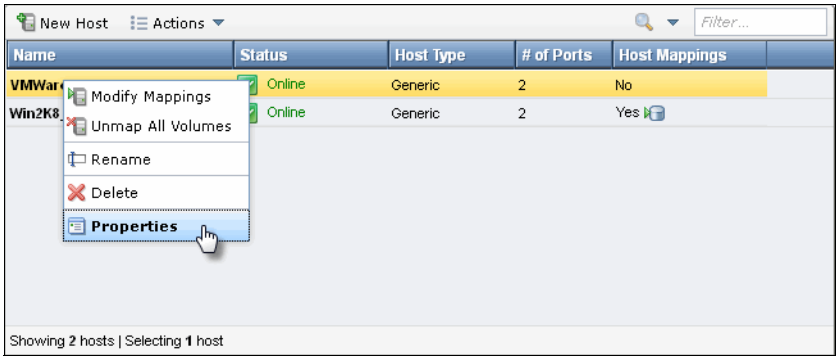

Figure 5-51 Example ESX FC host

2. Browse to the **Mapped Volumes** tab, as shown in Figure 5-52.

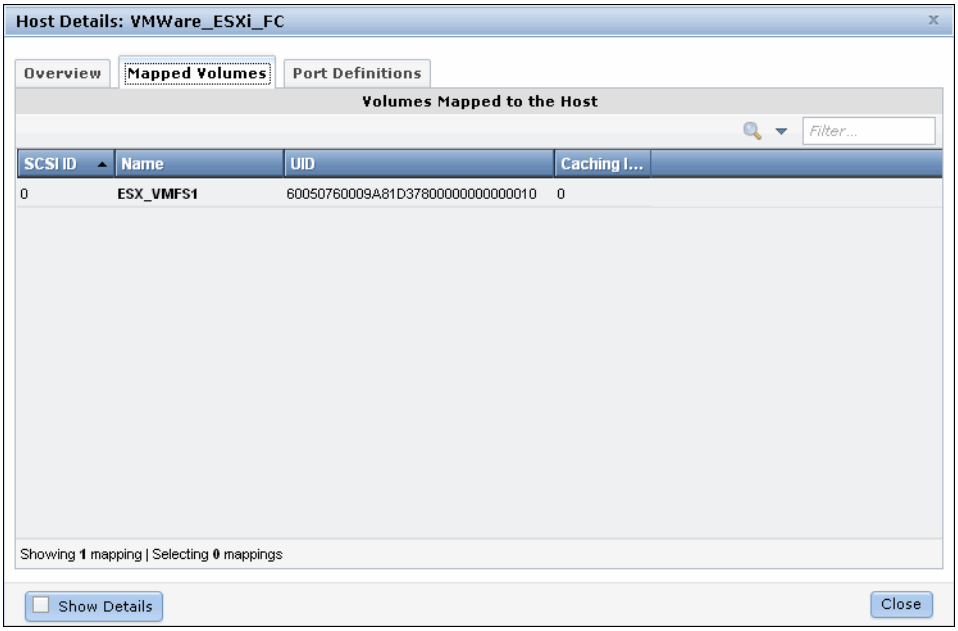

Figure 5-52 Mapped volumes to ESX FC host

In the Host Details window, you see that there is one volume that is connected to the ESX FC host by using 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 5-53.

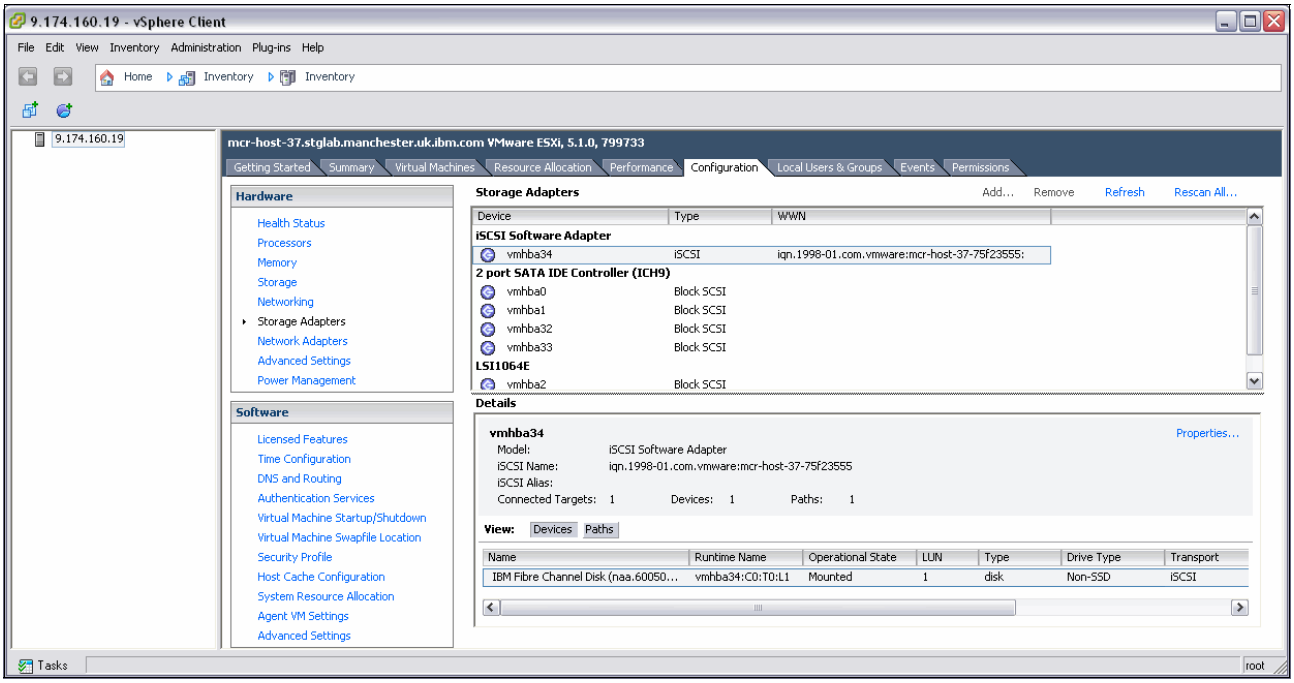

Figure 5-53 vSphere Client: Storage adapters

4. Click **Rescan All...** and click **OK** (as shown in Figure 5-54) to scan for new storage devices.

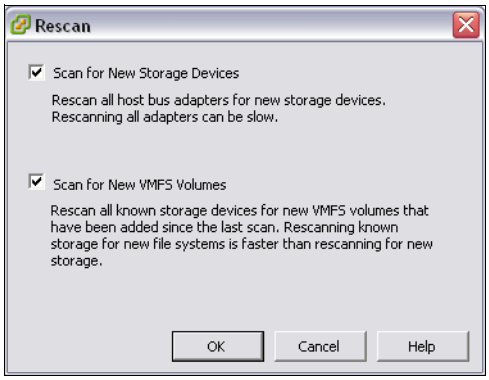

Figure 5-54 Rescan

5. Select **Storage** and click **Add Storage**, as shown in Figure 5-55.

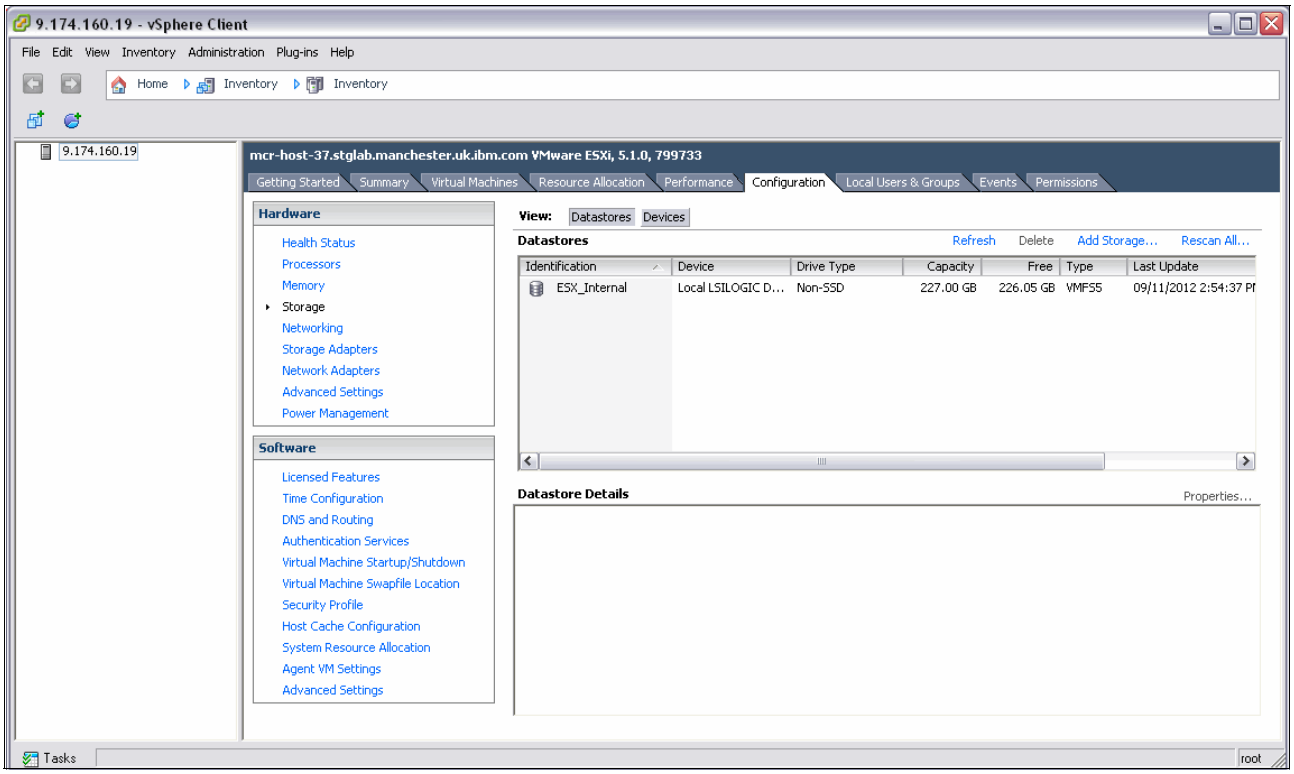

Figure 5-55 vSphere Client: Storage

6. The Add Storage wizard opens. Click **Select Disk/LUN** and then click **Next**. The IBM Storwize V3500 disk opens, as shown in Figure 5-56. Highlight it and click **Next**.

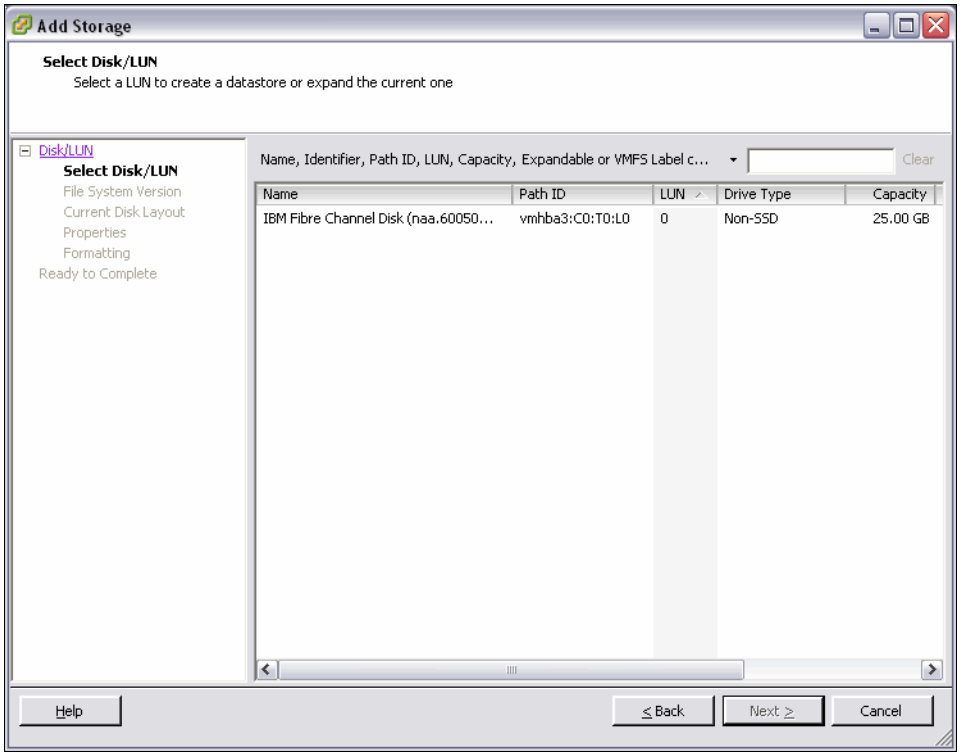

Figure 5-56 Select Disk/LUN menu

7. Follow the wizard to complete the attachment of the disk. After you click **Finish**, the wizard closes and you return to the storage view. In Figure 5-57, you see that the new volume is added to the configuration.

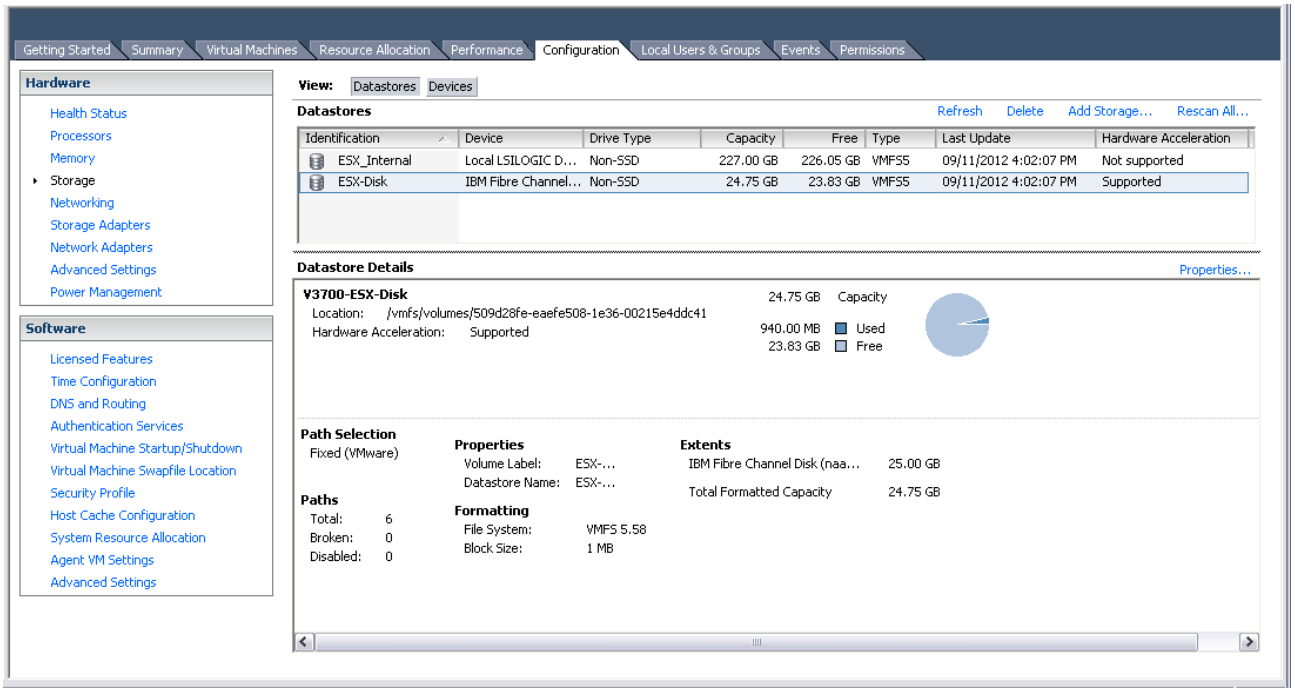

Figure 5-57 Add Storage task complete

8. Highlight the new data store and click **Properties** to see more details, as shown in Figure 5-58.

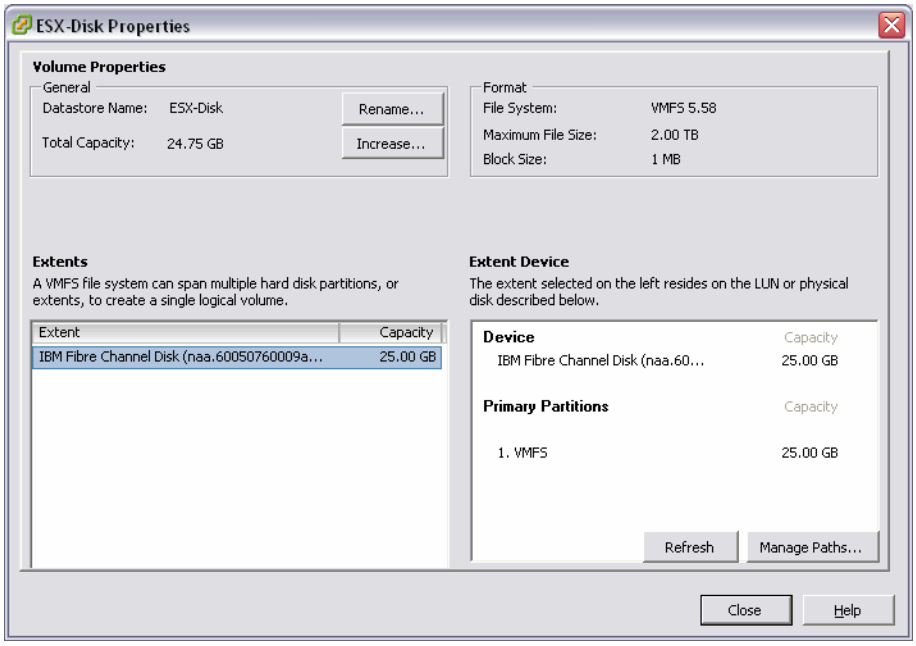

Figure 5-58 Data store properties

9. Click **Manage Paths** to customize the multipath settings. Select **Round Robin** (as shown in Figure 5-59) and click **Change**.

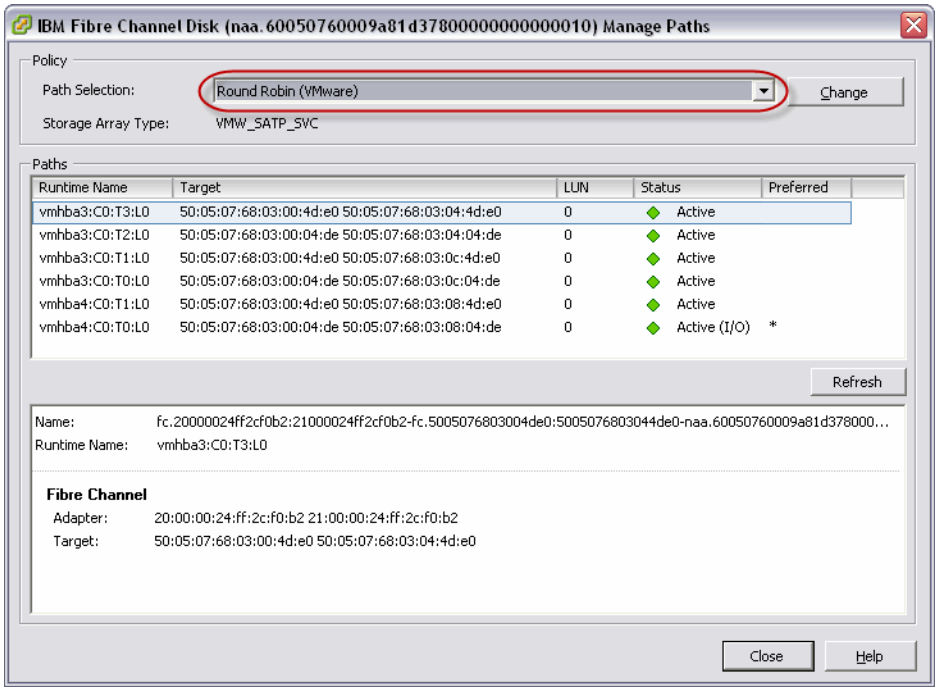

Figure 5-59 Select a data store multipath setting

The storage disk is available and ready to use for your VMware ESX server uses Fibre Channel attachment.

5.3.4 VMware ESX iSCSI attachment

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

1. Right-click your VMware ESX iSCSI host in the All Hosts view (as shown in Figure 5-60) and select **Properties**.

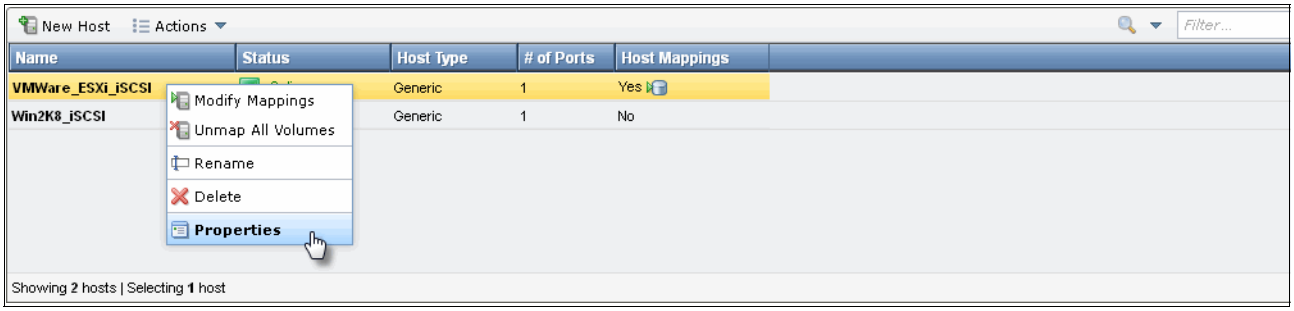

Figure 5-60 Select iSCSI ESX host properties

2. Brows to the Mapped Volumes tab, as shown in Figure 5-61.

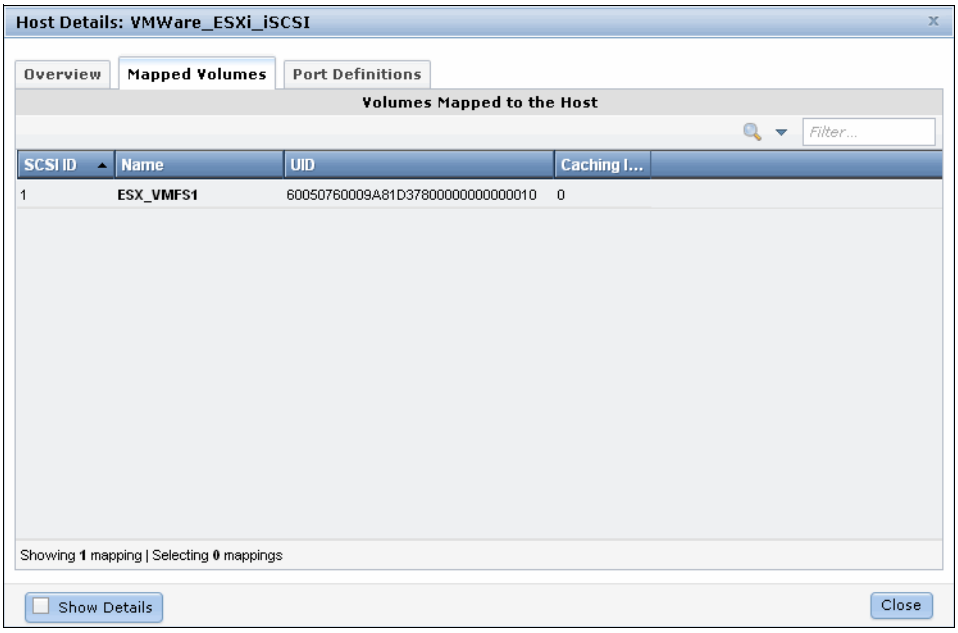

Figure 5-61 iSCSI ESX host properties

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

   Connect to your VMware ESX Server by using the vSphere Client, browse to the Configuration tab, and select **Storage Adapters**, as shown in Figure 5-62.

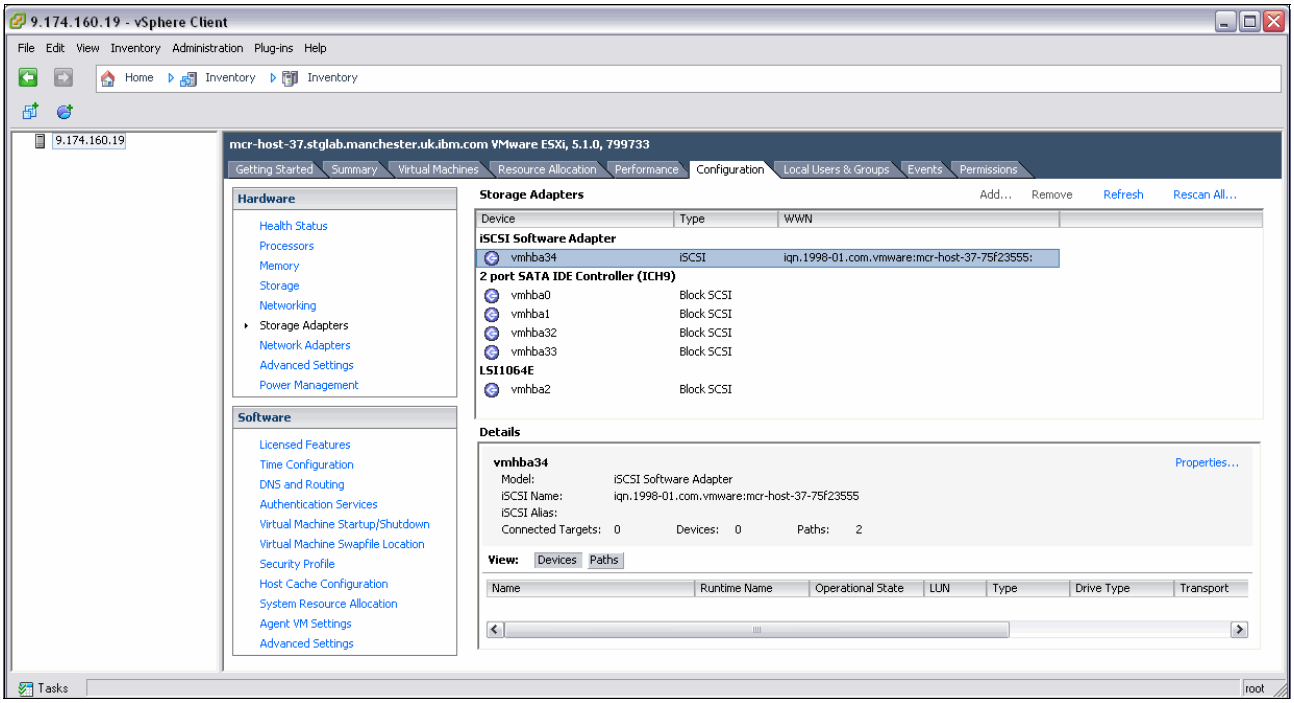

Figure 5-62 vSphere Client: Storage Adapters

4. Highlight the iSCSI Software Adapter and click **Properties**. The iSCSI initiator properties window opens. Select the **Dynamic Discovery** tab (as shown in Figure 5-63) and click **Add**.

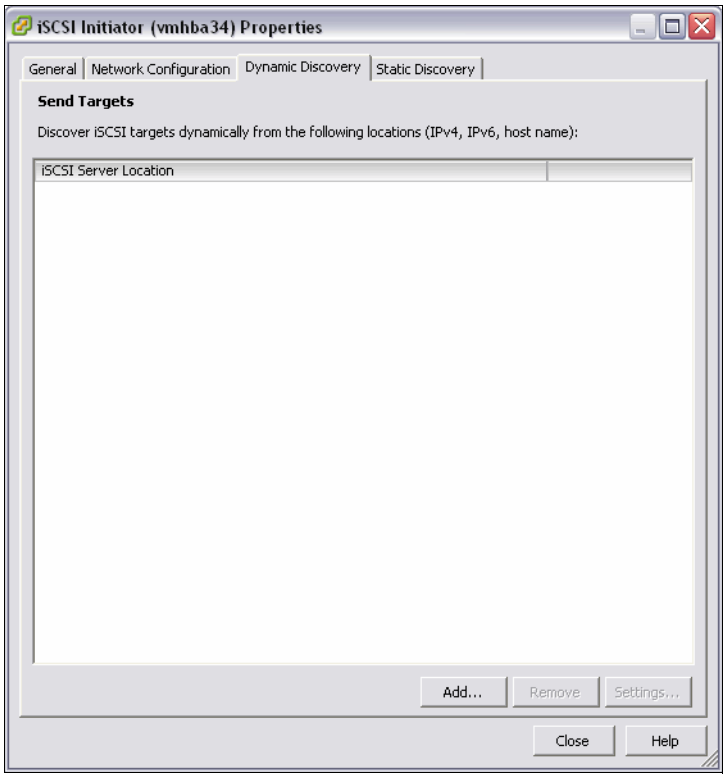

Figure 5-63 iSCSI Initiator properties

5. To add a target, enter the target IP address, as shown in Figure 5-64. The target IP address is the 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.

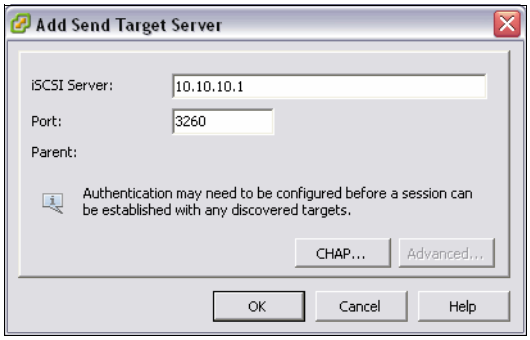

Figure 5-64 Enter a target IP address

Repeat this step for each IBM Storwize V3500 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 were configured in 4.3.2, “Creating iSCSI hosts” on page 164.

6. After you add all of the required ports, close the iSCSI Initiator properties by clicking **Close**, as shown in Figure 5-63 on page 209.

   You are prompted to rescan for new storage devices. Confirm the scan by clicking **Yes**, as shown in Figure 5-65.

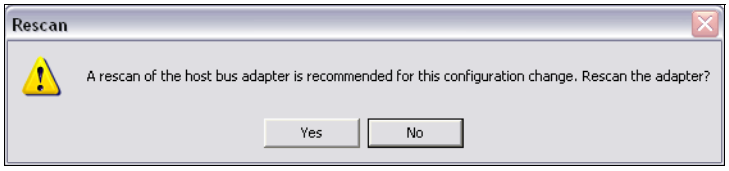

Figure 5-65 Confirm the rescan

7. Go to the storage view that is shown in Figure 5-66 and click **Add Storage**.

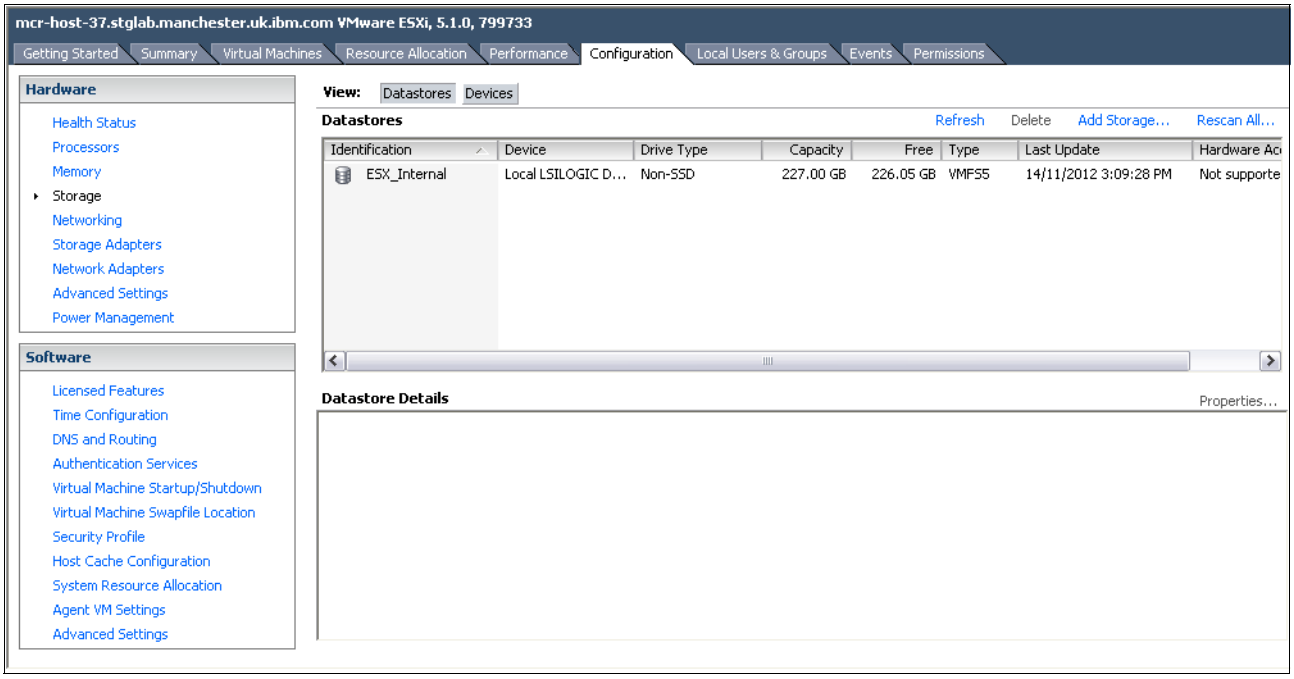

Figure 5-66 Add Storage menu

8. The Add Storage wizard opens, as shown in Figure 5-67. Select **Disk/LUN** and click **Next**.

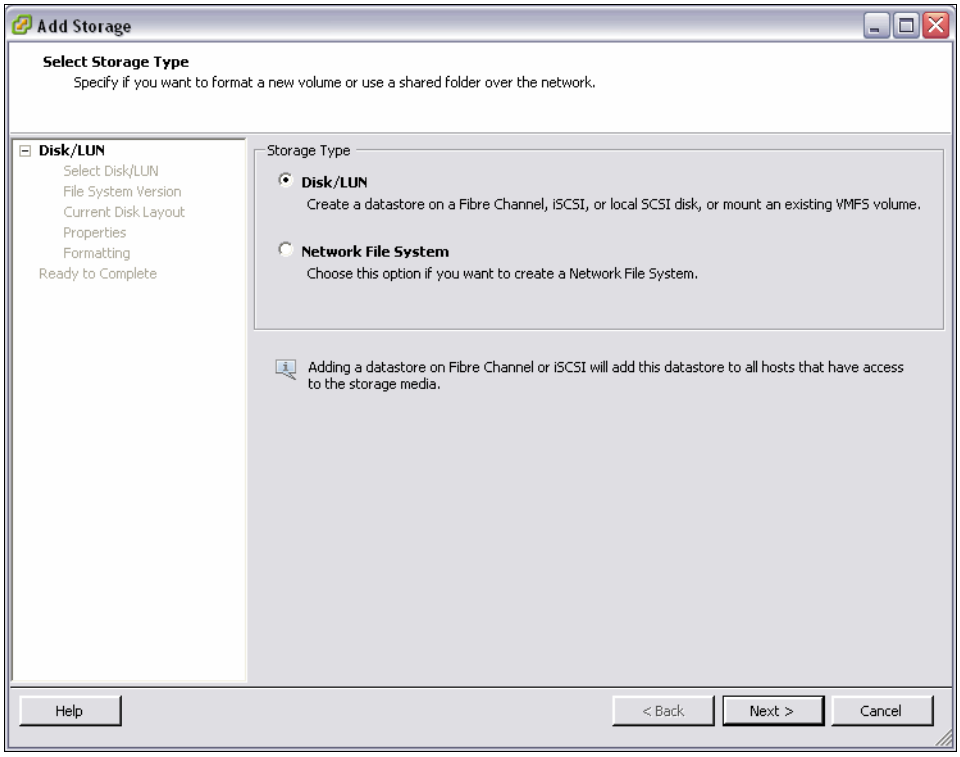

Figure 5-67 Select Disk/LUN menu

9. The new iSCSI LUN displays. Highlight it and click **Next**, as shown in Figure 5-68.

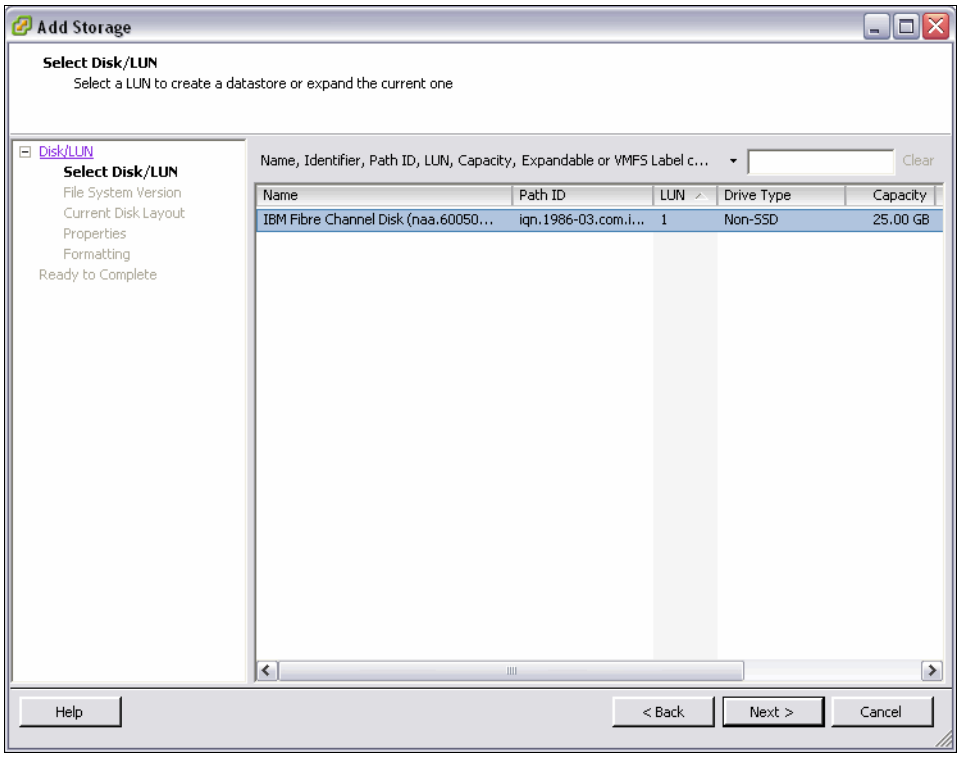

Figure 5-68 Select iSCSI LUN menu

10. Select a File System Version option. In our example, we selected **VMFS-5**, as shown in Figure 5-69.

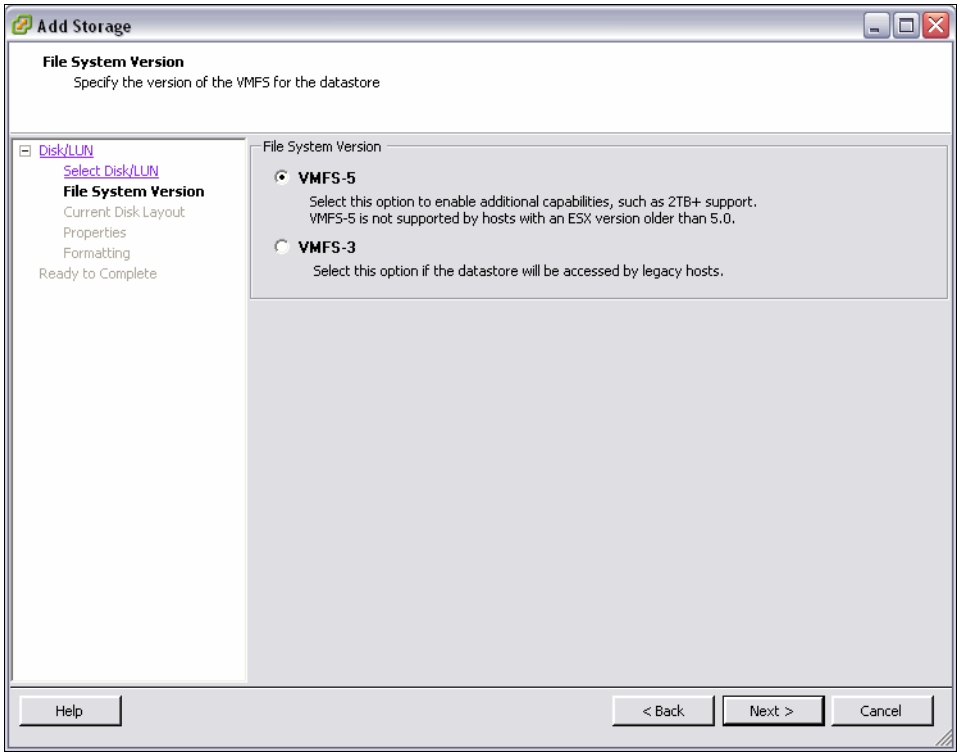

Figure 5-69 Select File System Version

11. Review the disk layout and click **Next**, as shown in Figure 5-70.

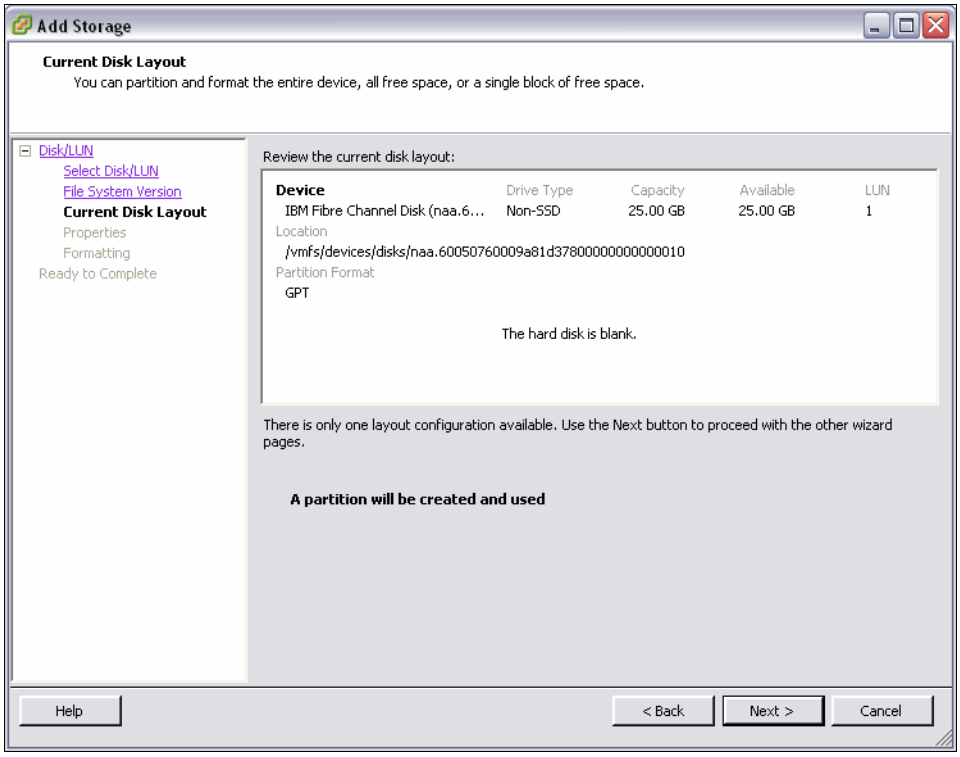

Figure 5-70 Current Disk Layout

12. Enter a name for the data store and click **Next**, as shown in Figure 5-71.

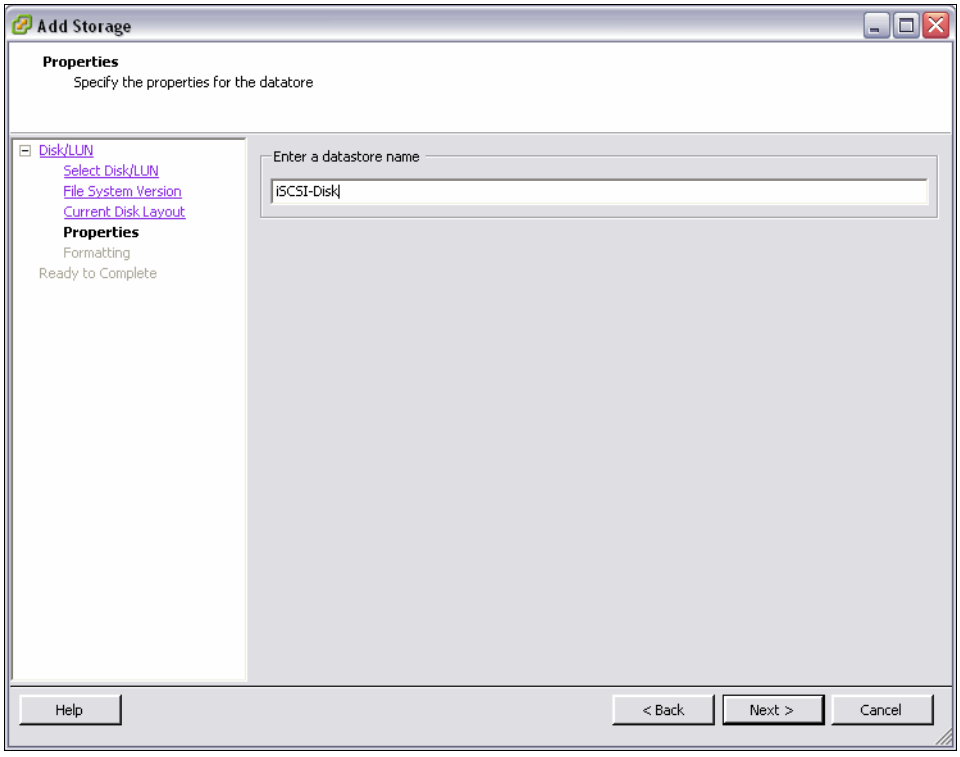

Figure 5-71 Enter a data store name

13. Select the Maximum available space and click **Next**, as shown in Figure 5-72.

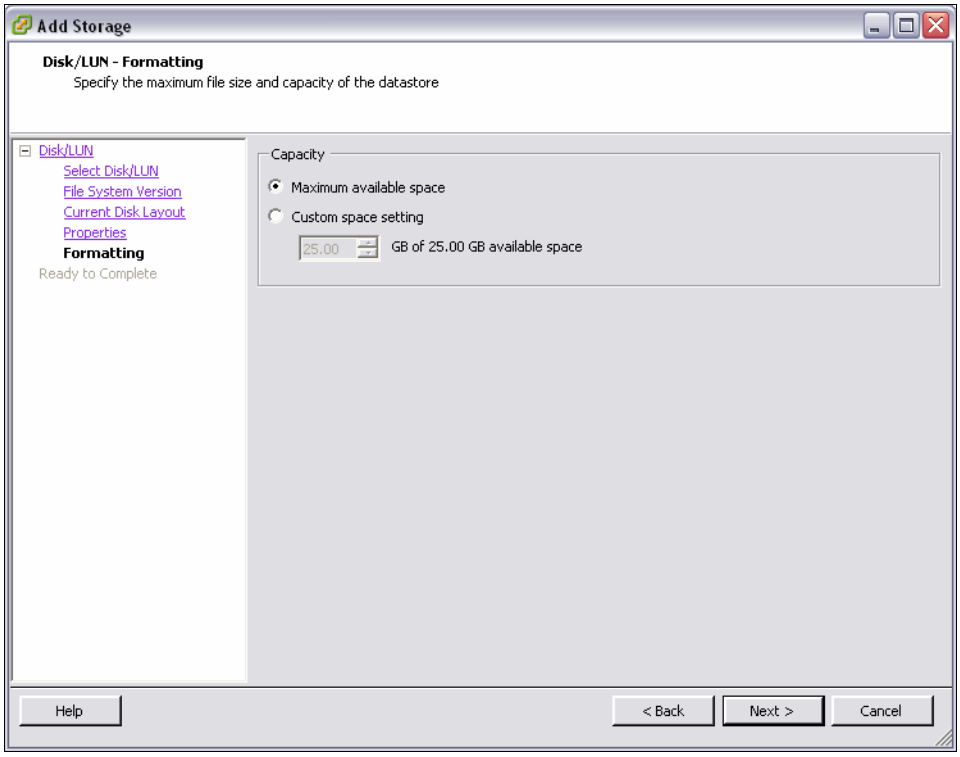

Figure 5-72 Capacity

14. Review your selections and click **Finish**, as shown in Figure 5-73.

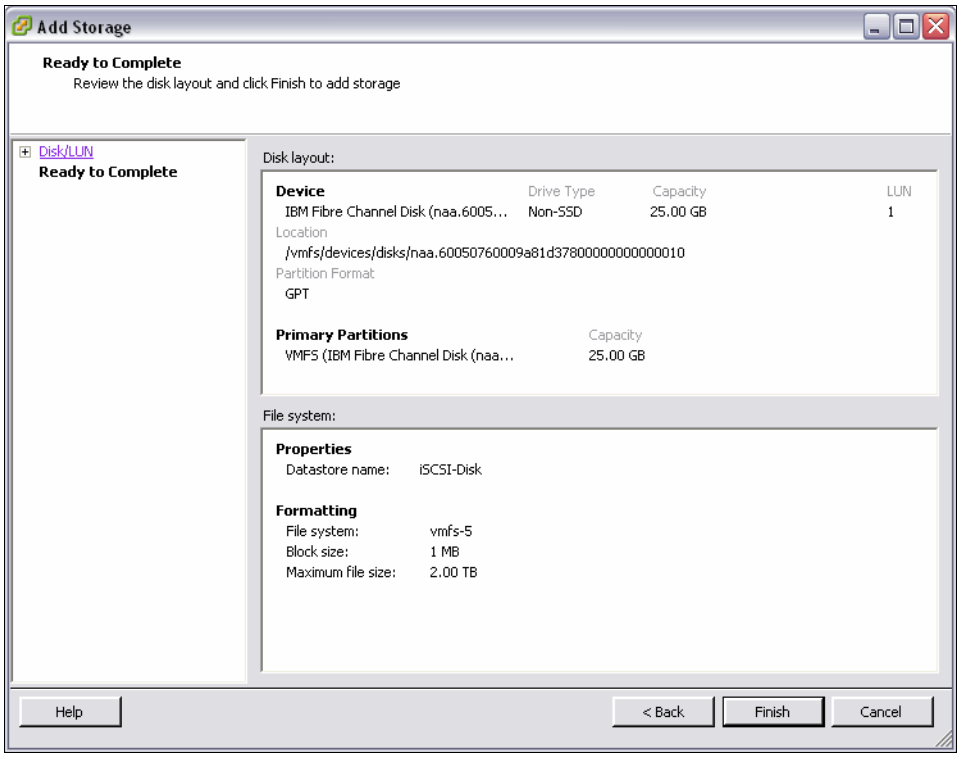

Figure 5-73 Finish the wizard

The process starts to add an iSCSI LUN. This task can take a few minutes. After the tasks complete, the new data store appears in the storage view, as shown in Figure 5-74.

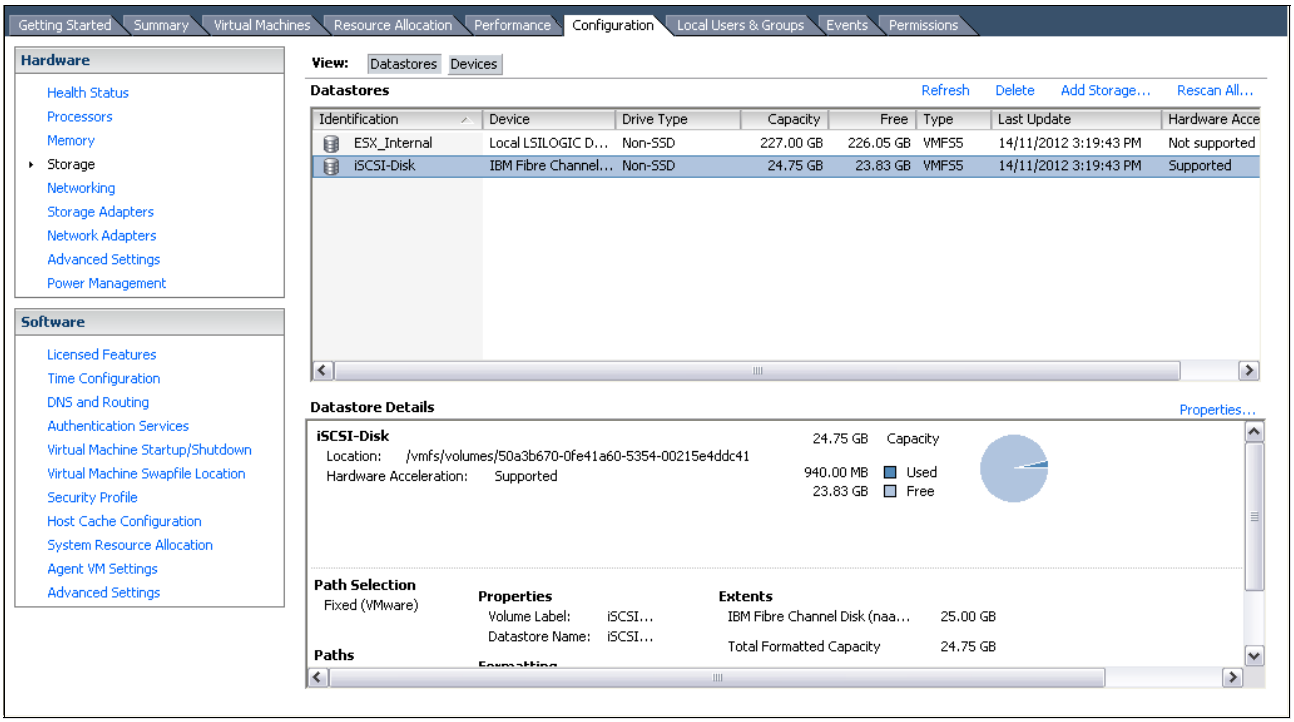

Figure 5-74 New data store available

15. Highlight the new data store and click **Properties** to open and review the data store settings, as shown in Figure 5-75.

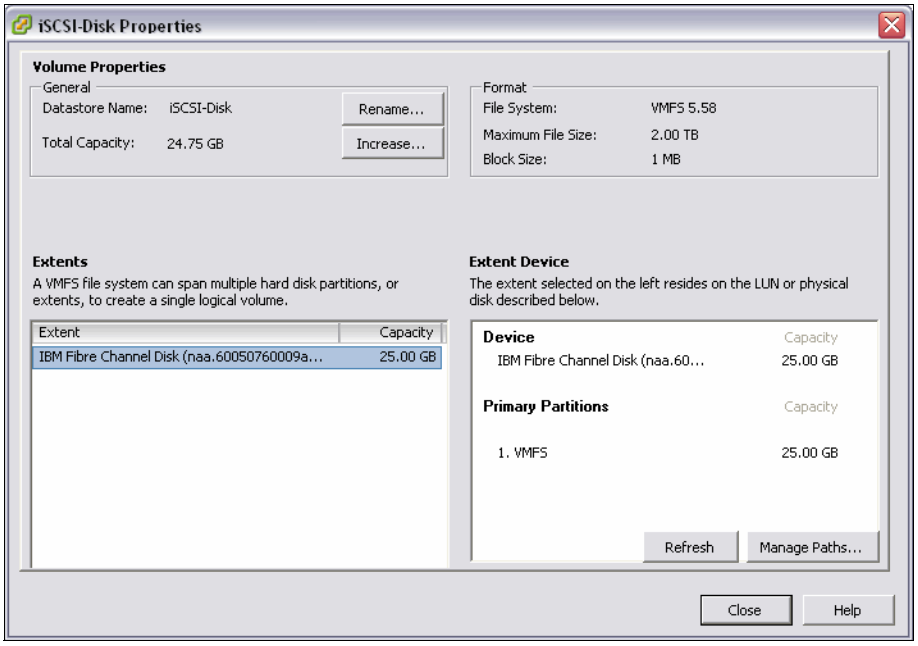

Figure 5-75 Data store properties

16. Click **Manage Paths**, select **Round Robin** as the multipath policy, as shown in Figure 5-76. Click **Change**.

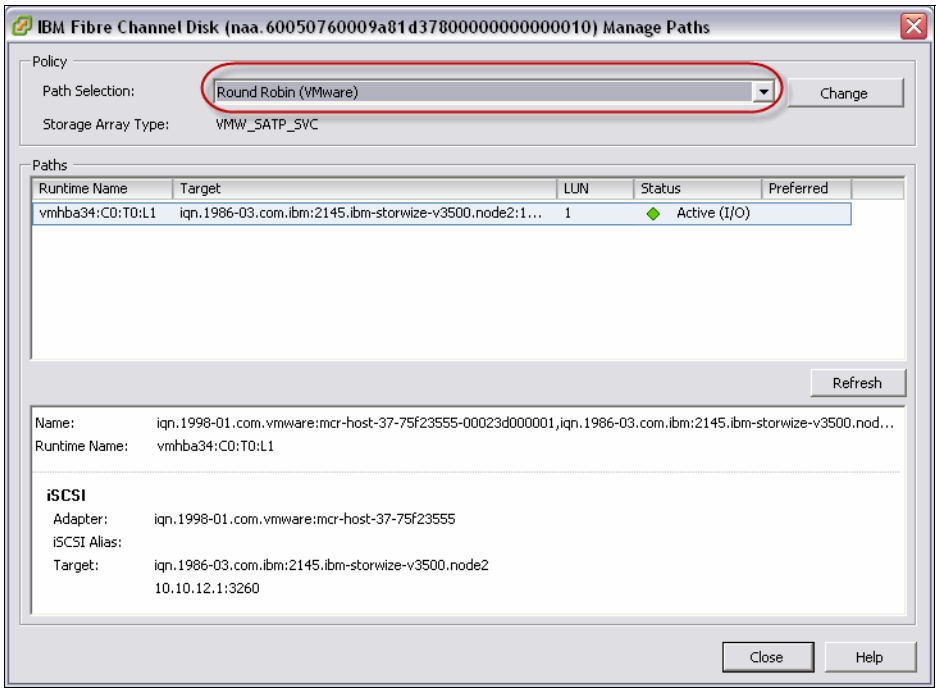

Figure 5-76 Change the multipath policy

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

6

Storage migration wizard

This chapter describes the easy to follow storage migration wizard. The storage migration wizard is used to migrate data from older external storage systems to the internal capacity of the Storwize V3500. Migrating data from older storage systems to the Storwize V3500 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
- The storage migration wizard
- Storage migration wizard example scenario

6.1 Interoperability and compatibility

Interoperability is an important consideration when you set up a new storage system in an environment that contains existing storage infrastructure. In this section, we show how to check the storage environment, the older storage system, and the IBM Storwize V3500 are ready for the data migration process.

Important: The Storwize V3500 node canisters require an appropriate optional Host Interface Card (HIC) to communicate with storage systems that are attached to the same Fibre Channel SAN.

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 the solution is supported and provides recommended hardware and software levels.

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

For more information, see the IBM System Storage Interoperation Center (SSIC) at this website:

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

6.2 The storage migration wizard

The Storwize V3500 storage migration wizard simplifies migration. The wizard is one way with easy to follow panels to guide users through the entire process. The process involves external virtualization of the older storage system, then an online migration is performed. After data migration is complete, the older storage system is removed from Storwize V3500 control and can be retired.

6.2.1 External virtualization capability

To migrate data from an older storage system to the internal capacity of the Storwize V3500, it is necessary to temporarily take advantage of the built-in external virtualization capability. This capability places external Fibre Channel-connected Logical Units (LUs) under the control of the Storwize V3500. Control of the external LUs is established by using and following the storage migration wizard. The built-in external virtualization capability of the Storwize V3500 can be used only for one-way data migration purposes. Any other use of this feature is not supported. If a permanent external virtualization solution is required, the Storwize V7000, Storwize V7000 Unified, or SAN Volume Controller should be considered.

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) logical units (LUs) that are presented on the Fibre Channel SAN to hosts.
- I/O to the LUs is stopped and changes are made to the storage system LUs mappings and to the SAN fabric zoning so that the original LUs are presented directly to the Storwize V3500. The Storwize V3500 discovers the external LUs as unmanaged MDisks.
- The unmanaged MDisks are then imported to the Storwize V3500 as image mode MDisks and placed into 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 with the data remaining unchanged. The Storwize V3500 is simply presenting active images of the SAN-attached LUs.
- The hosts have the older storage system multipath device driver that is removed and then are configured for Storwize V3500 attachment. Further zoning changes are made for host-to-V3500 SAN connections. The Storwize V3500 hosts are defined with WWPNs and the volumes are mapped. After the volumes are mapped, the hosts discover the Storwize V3500 volumes through a host rescan device or reboot operation.
- Storwize V3500 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 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 V3500 LUs control to that 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 V3500.

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 V3500 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 migration progress information, as shown in Figure 6-1.

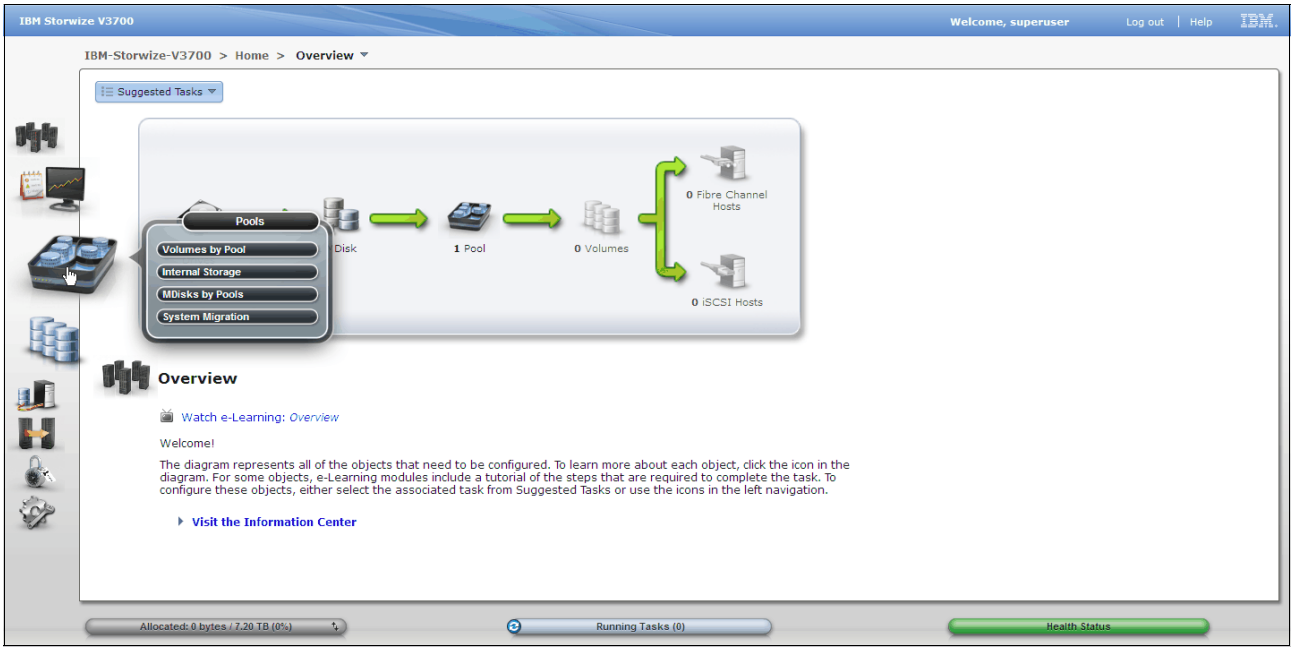

Figure 6-1 Pools menu

Click **Start New Migration** and the storage migration wizard starts, as shown in Figure 6-2.

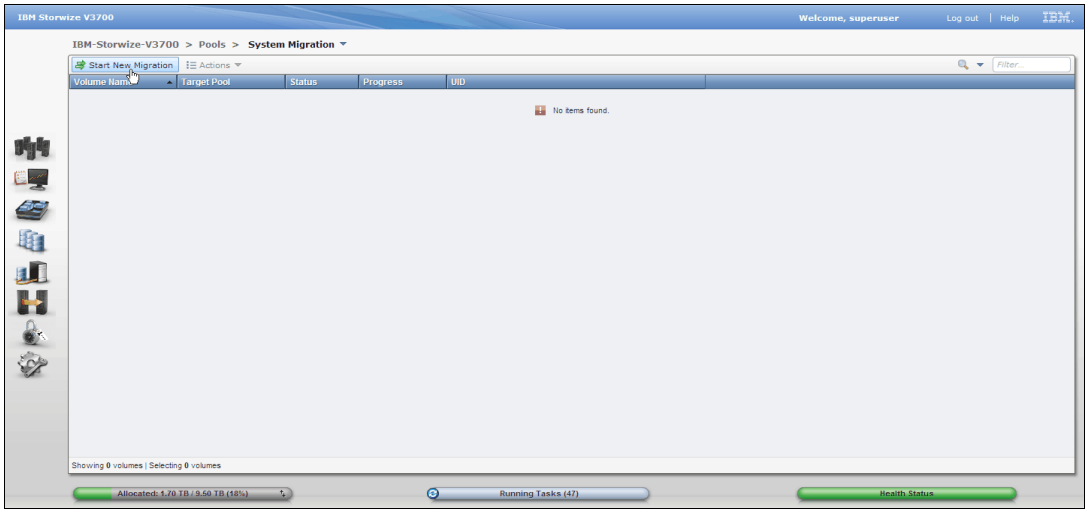

Figure 6-2 System migration panel

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 restrictions and prerequisite that applies to the planned migration, as shown in Figure 6-3.

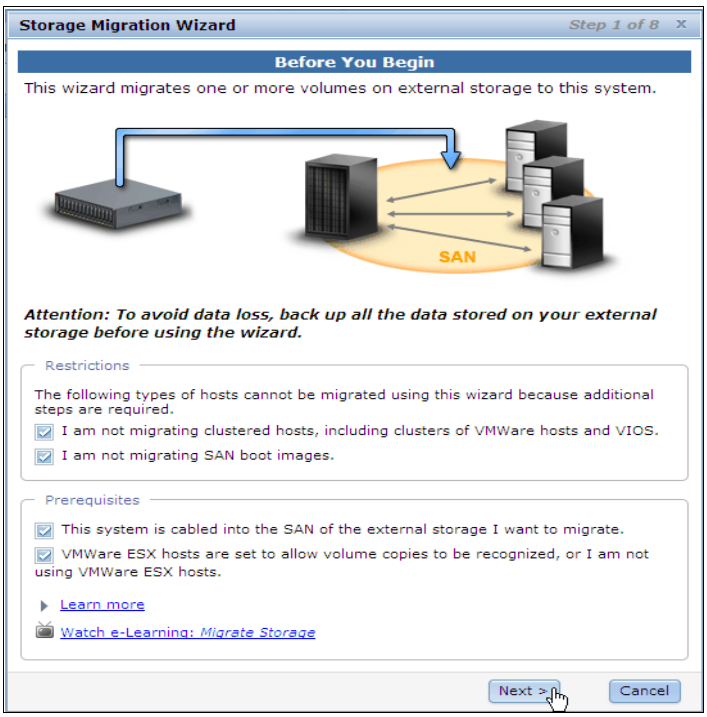

Figure 6-3 Step 1 of the storage migration wizard, with all options selected

Restrictions

Confirm that the following restrictions apply:

- I am not using the storage migration wizard to migrate cluster hosts, including cluster of VMware hosts and VIOS.
- I am not using the storage migration wizard to migrate SAN Boot images.

If the restrictions boxes cannot be selected, the migration must be performed outside of this wizard because more steps are required. For more information about this topic, see the IBM Storwize V3500 Information Center at this website:

`http://pic.dhe.ibm.com/infocenter/storwize/v3500_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 V3500, older storage system, and 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 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 boxes 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 6-3 on page 221 shows step 1 of the storage migration wizard with all restrictions satisfied and prerequisites met.

Step 2: Prepare environment for migration

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

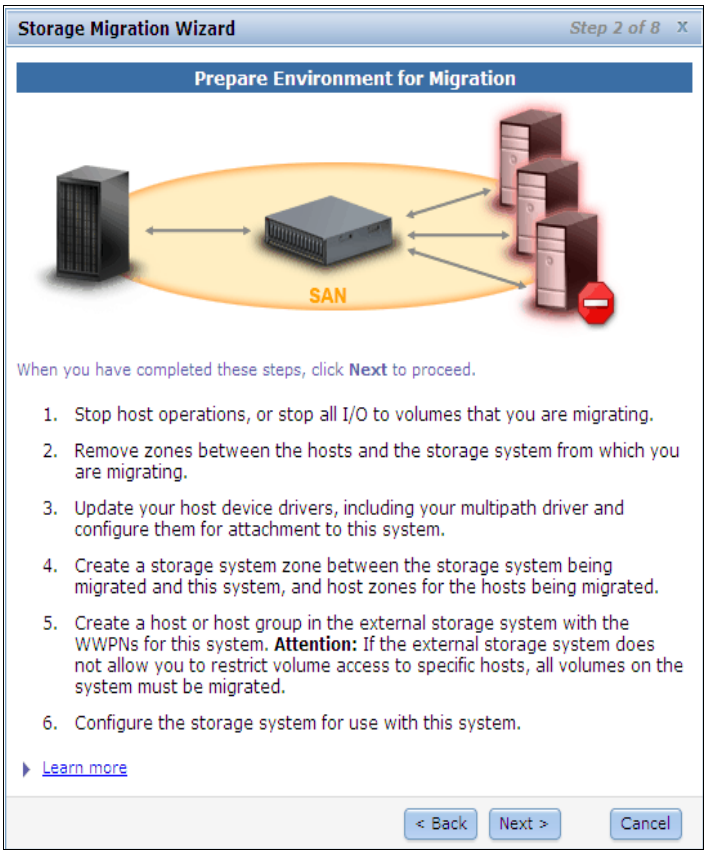

Figure 6-4 Prepare Environment for Migration panel

Step 3: Map storage

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

Table 6-1 Example table for recording 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 that can be used to record host information.

Table 6-2 Example table for recording host information

Host Name/ LU Names	Adapter / Slot / Port	WWPN	HBA F/W	HBA Device Driver	Operating System	V3500 Multipath S/W
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 on page 224 shows the Map Storage panel.

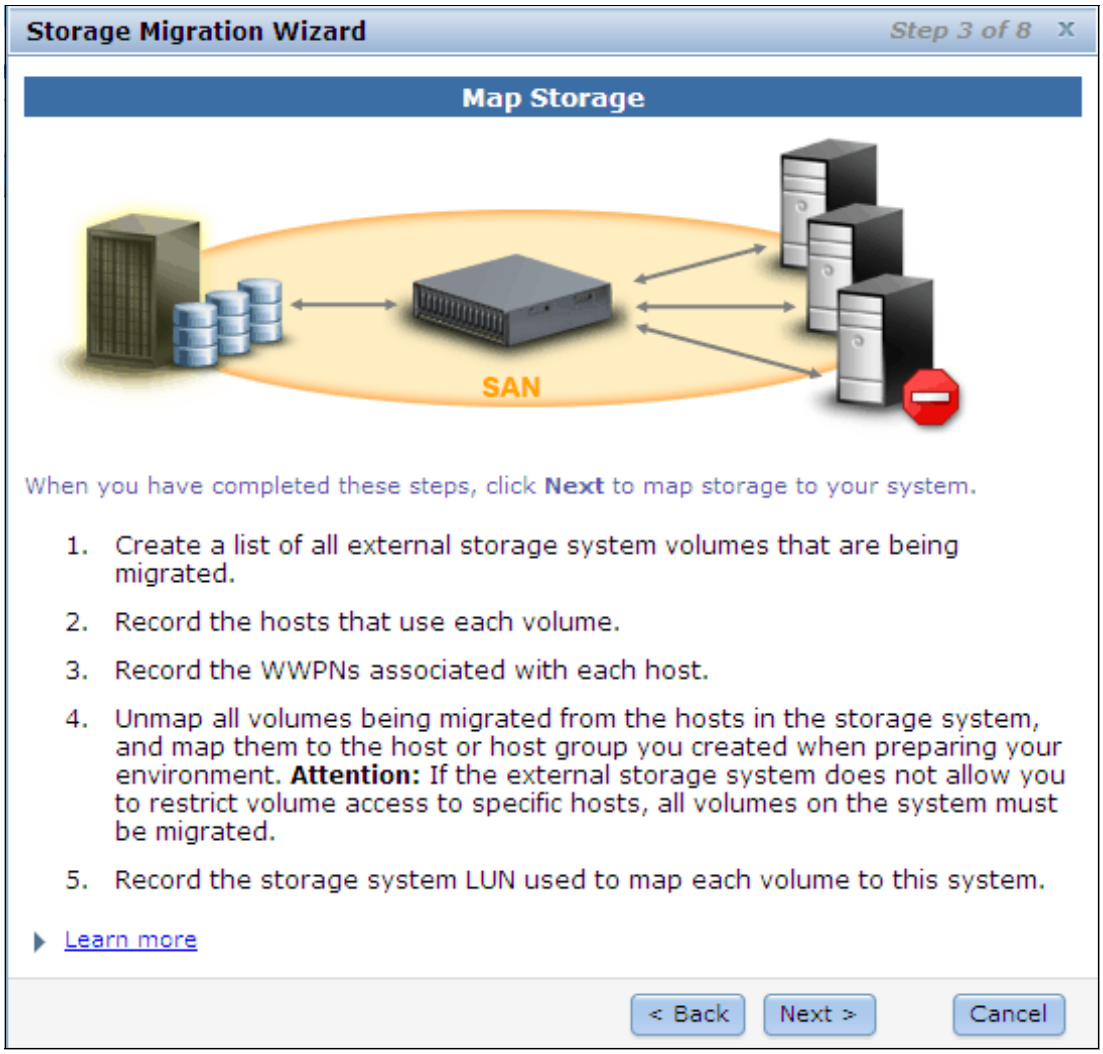

Figure 6-5 Map storage panel

The Storwize V3500 runs the Discover Devices task, as shown in Figure 6-6. After the task completes successfully, click **Close** to continue.

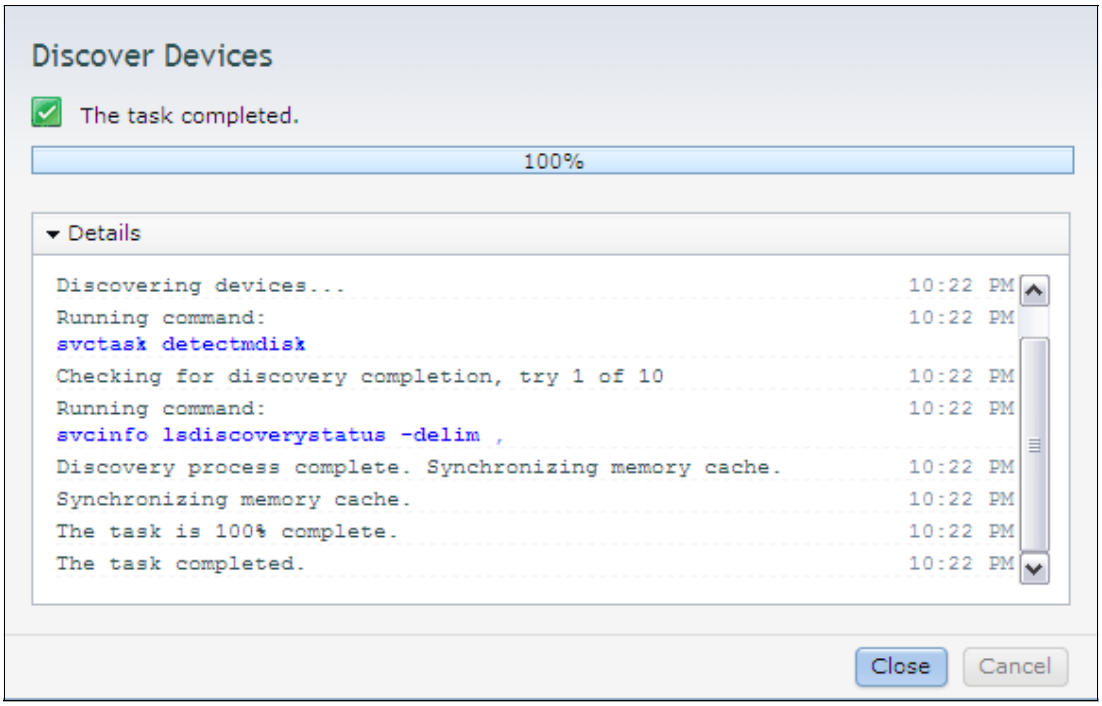

Figure 6-6 Discover Devices task

Step 4: Migrate MDisks

Follow step 4 of the storage migration wizard and select the MDisks that you want to migrate. Click **Next** to continue. Figure 6-7 shows the migrating MDisks panel.

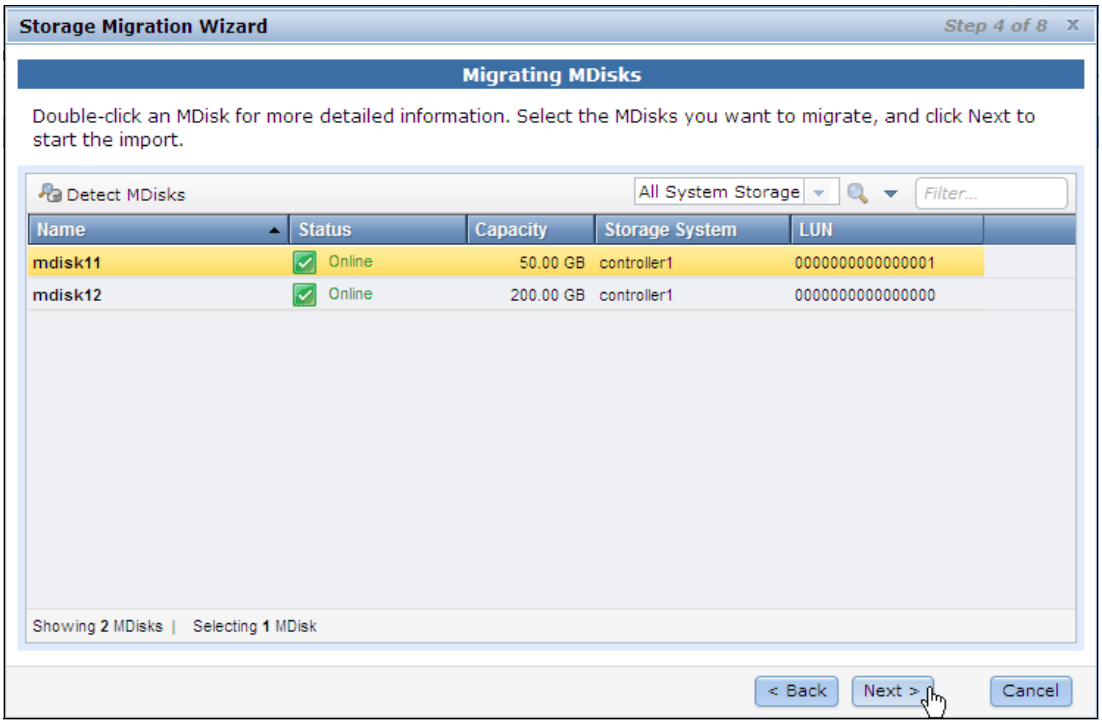

Figure 6-7 Migrating MDisks panel

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

The Storwize V3500 runs the import MDisks task. After the task completes successfully, click **Close** to continue. Figure 6-8 shows the result of the import MDisks task.

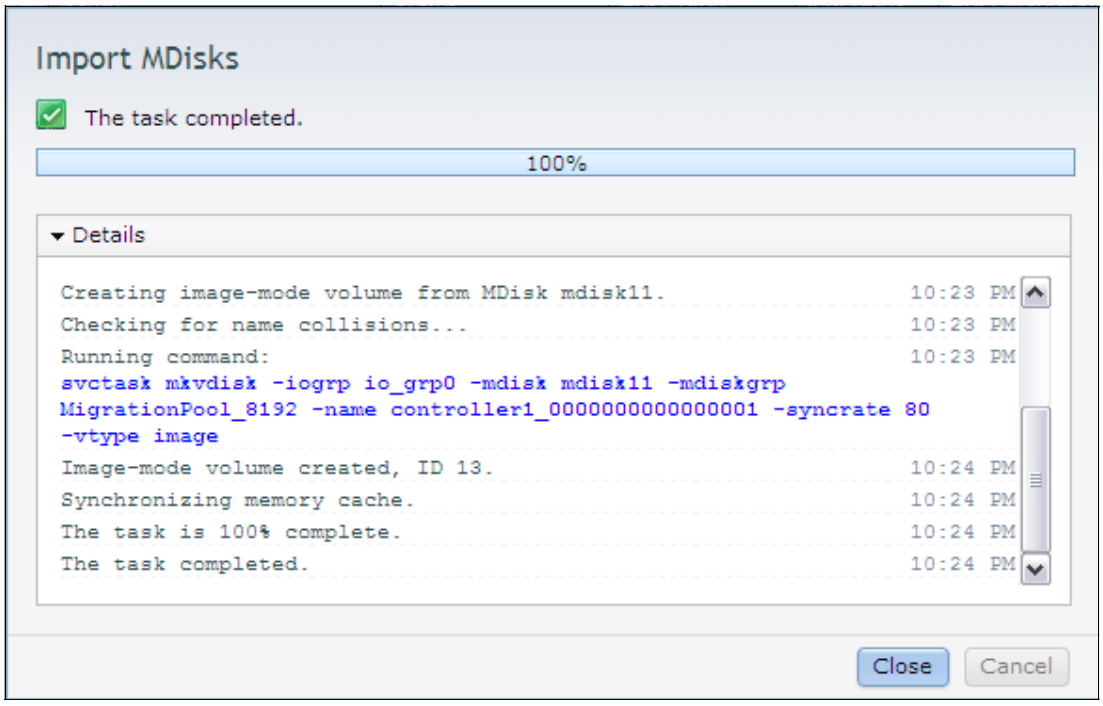

Figure 6-8 Import MDisks task

Step 5: Configure hosts

Follow step 5 of the storage migration wizard. Select or configure new hosts as required and then click **Next** to continue. Figure 6-9 shows the configure hosts panel.

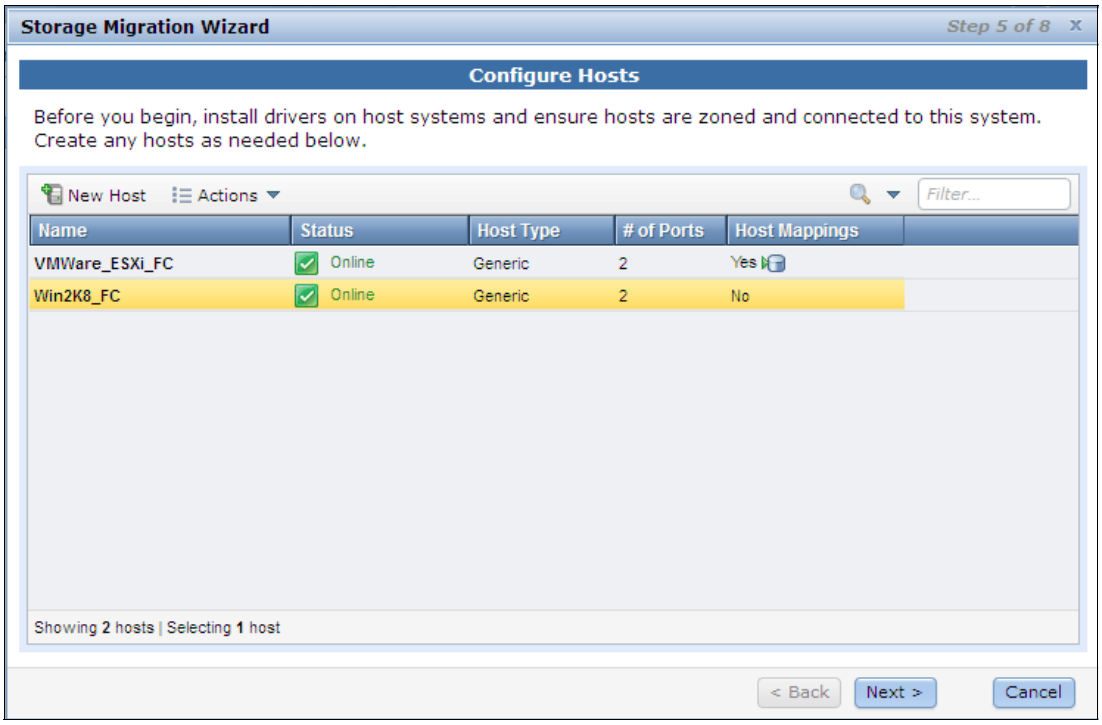

Figure 6-9 Configure hosts panel

Note: It is not mandatory to select the hosts now. Host selection occurs in step 6. However, you might want 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 storage migration wizard. Select the newly migrated volume and click **Map to Host**, as shown in Figure 6-10.

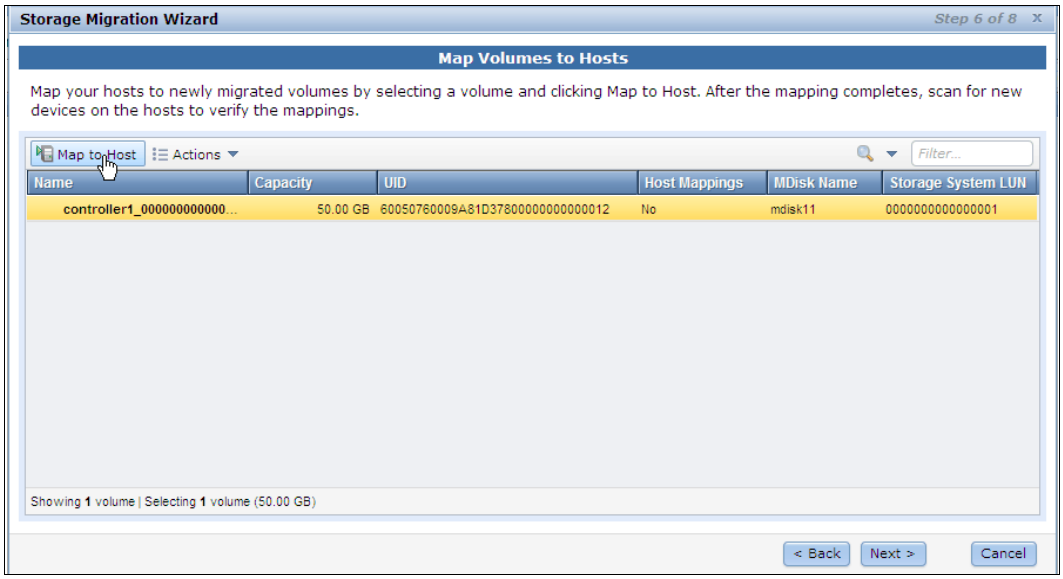

Figure 6-10 Map volumes to hosts panel

The image mode volumes are listed. The names of the image mode volumes are assigned automatically by the Storwize V3500 storage system. The names can be changed to reflect something more meaningful to the user. Select the volume and click **Rename** in the Actions drop-down menu.

Names: Image volume names must begin with a letter. The name can be a maximum of 63 characters. Valid characters are uppercase letters (A-Z), lowercase letters (a-z), digits (0 - 9), underscore (_), period (.), hyphen (-), and space. The names must not begin or end with a space.

A Host drop-down menu is displayed, as shown in Figure 6-11. Select the wanted host and the Modify Host Mappings panel opens.

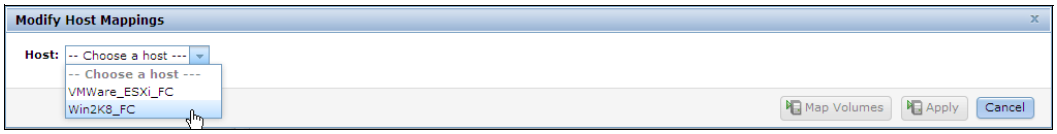

Figure 6-11 Host drop-down menu

The MDisks that were highlighted in step 6 are shown in yellow highlights 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 the volumes as required, which should reflect the same SCSI ID recorded in Step 3. Click **Map Volumes**. Figure 6-12 shows the modify host mapping panel.

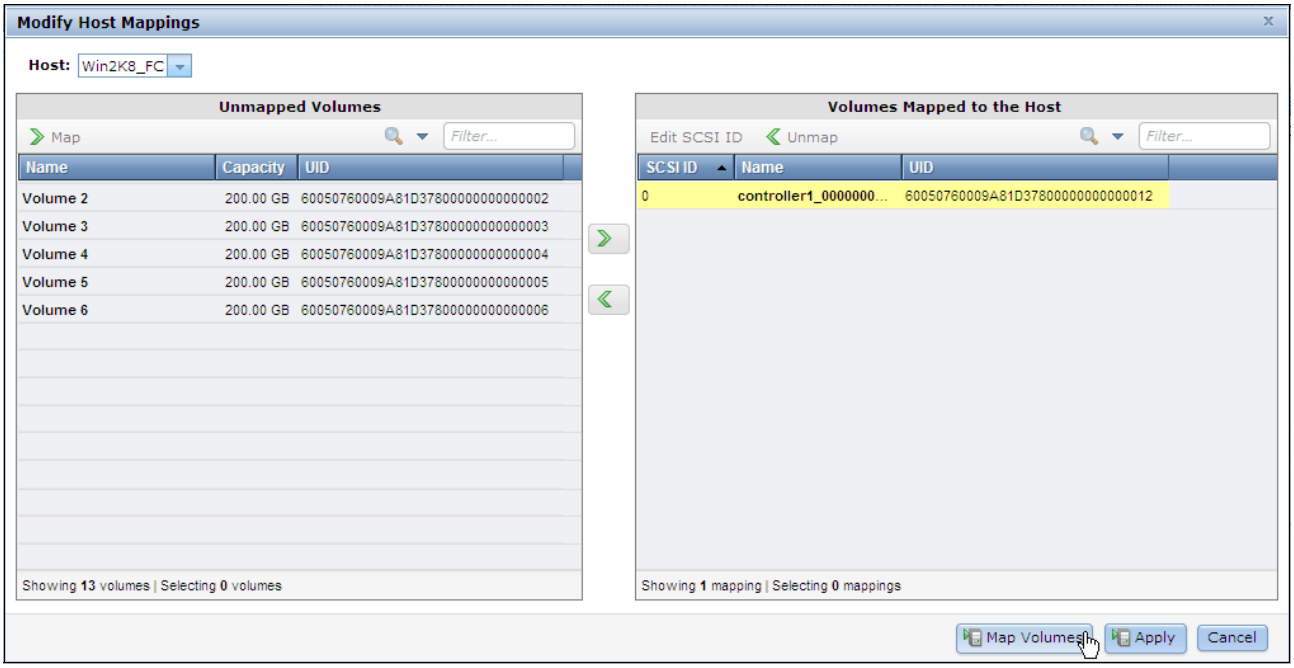

Figure 6-12 Modify hosts mappings panel

The Storwize V3500 runs the Modify Mappings task, as shown in Figure 6-13. After the task is completed successfully, the volume is mapped to the host. Click **Close** to continue.

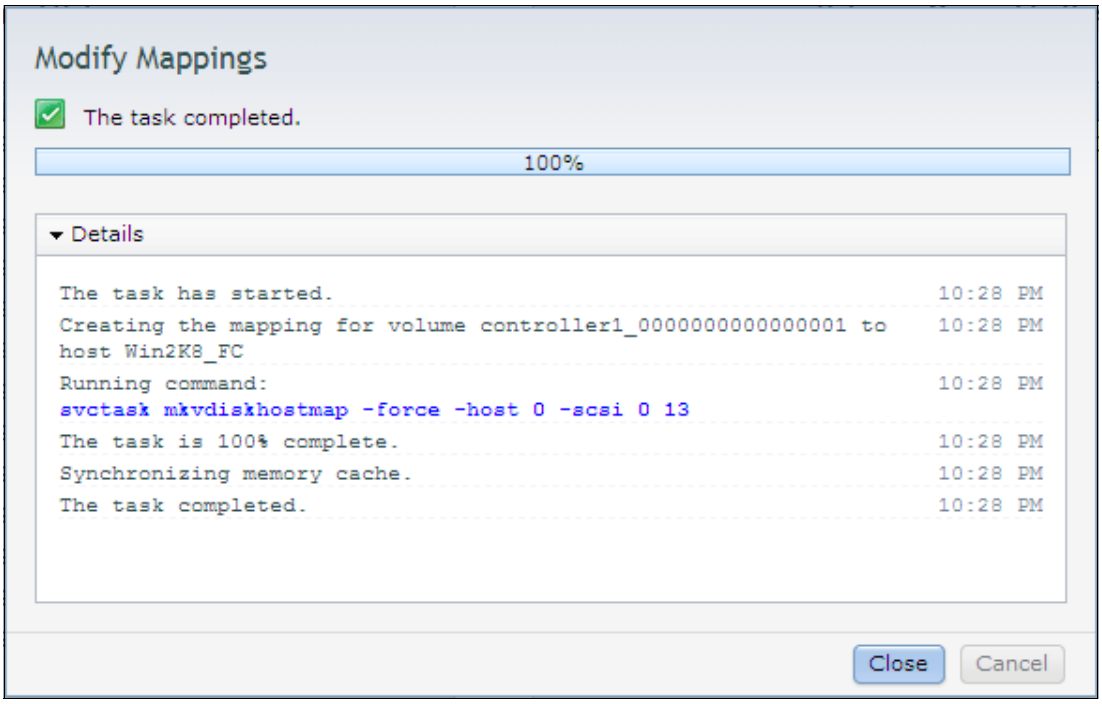

Figure 6-13 Modify mappings tasks

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

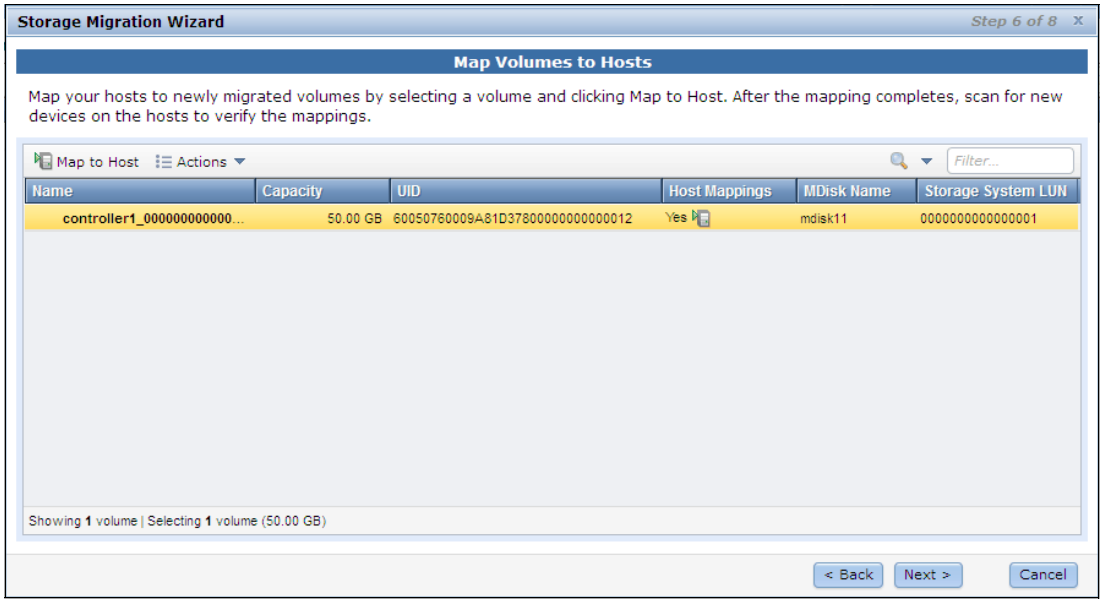

Figure 6-14 Map volumes to hosts panel displaying "yes" in the host mappings column

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 storage migration wizard. Select an internal storage pool and then click **Next** to continue. The destination storage pool of the data migration is an internal storage pool of the Storwize V3500.

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 that is placed on the MDisks in the selected storage pool. The process uses the volume mirroring function that is included with the Storwize V3500. When the process is complete, the volumes have pointers to the new copy that is on the selected internal storage pool and on the older storage system.

Figure 6-15 shows the select storage pool panel.

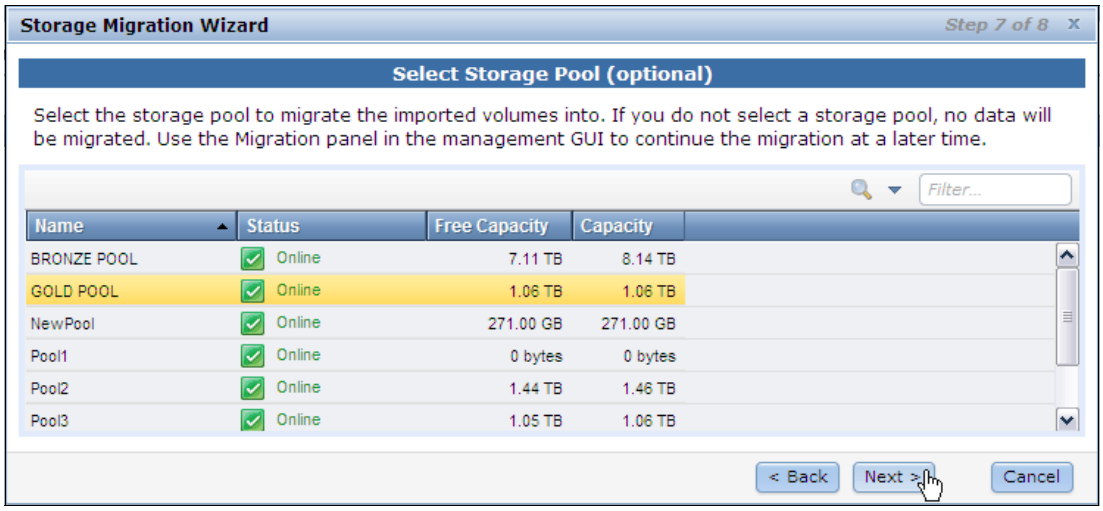

Figure 6-15 Select a storage pool panel

The Storwize V3500 runs the Start Migration task, as shown in Figure 6-16. After the task is completed successfully, click **Close** to continue.

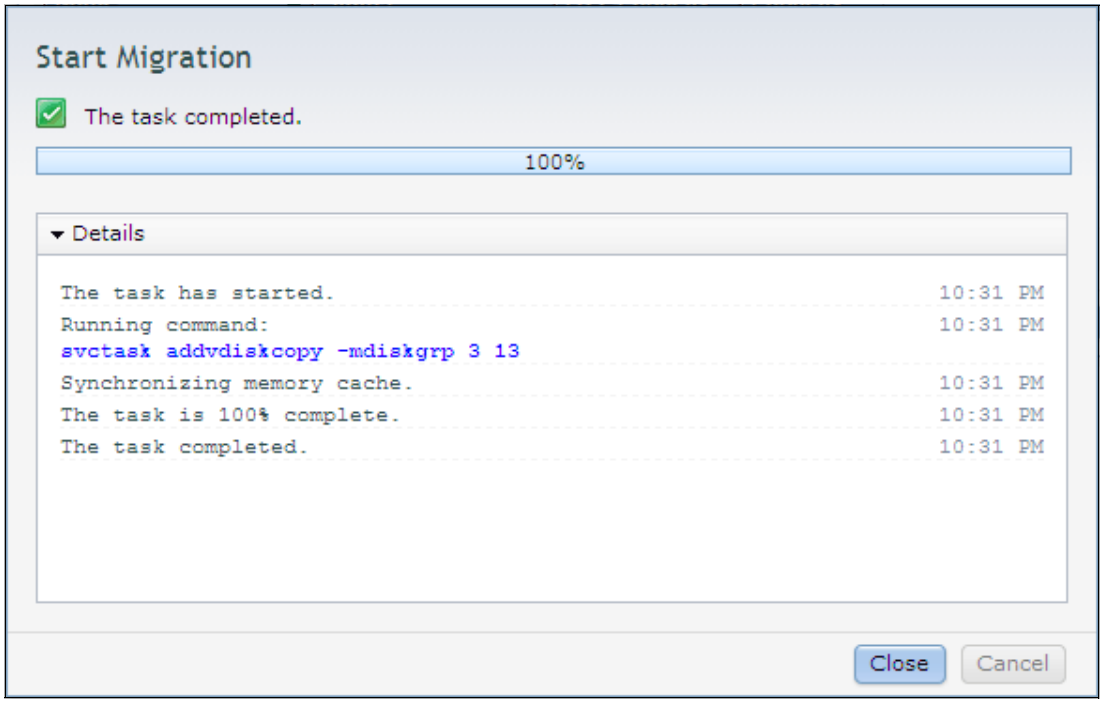

Figure 6-16 Start Migration task

Step 8: Finish the storage migration wizard

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

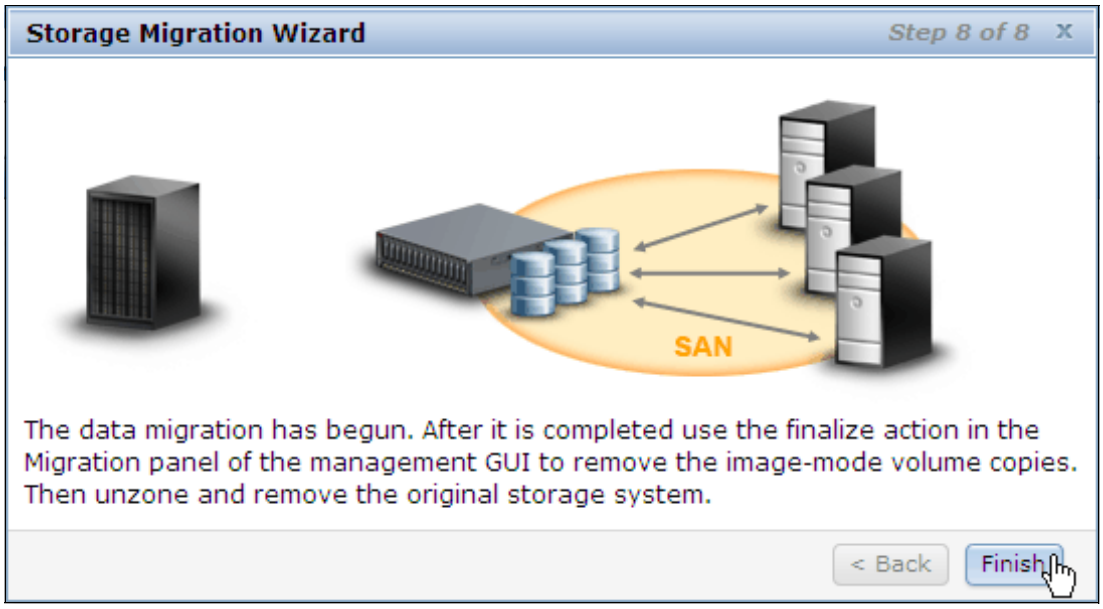

Figure 6-17 Step 8 of 8 of the storage migration wizard

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

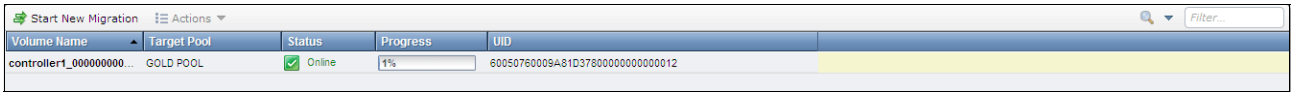

Figure 6-18 Storage Migration panel with a migration in progress

Finalize migrated volumes

When the migration completes, select the Migration and right-click **Finalize.** Verify 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 V3500 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 V3500 to an environment that contains an 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 V3500. The Windows 2008 host has a dual port QLogic Host Bus Adapter (HBA) type QLE2562. Each of the Fibre Channel switches is of type IBM 2498-24B. 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).

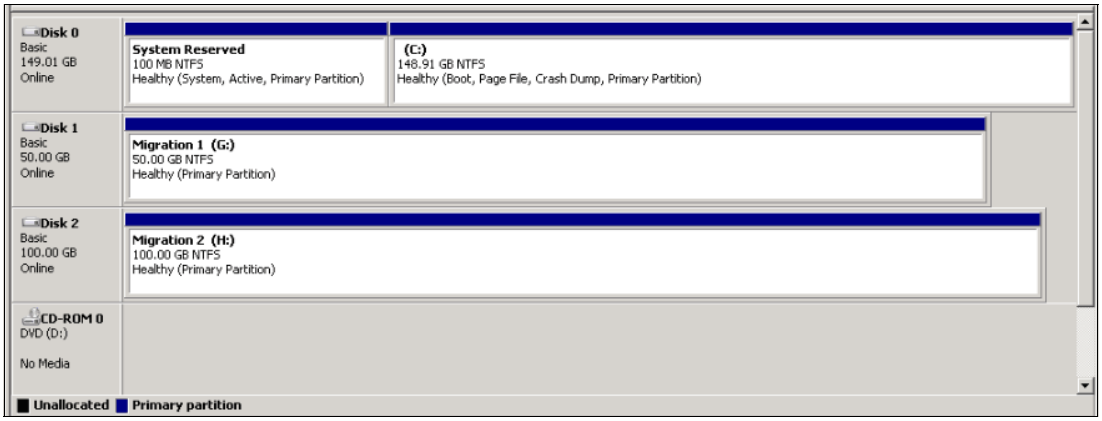

Figure 6-19 Windows 2008 disk management panel

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 open the menu and then select **Properties**, as shown in Figure 6-20.

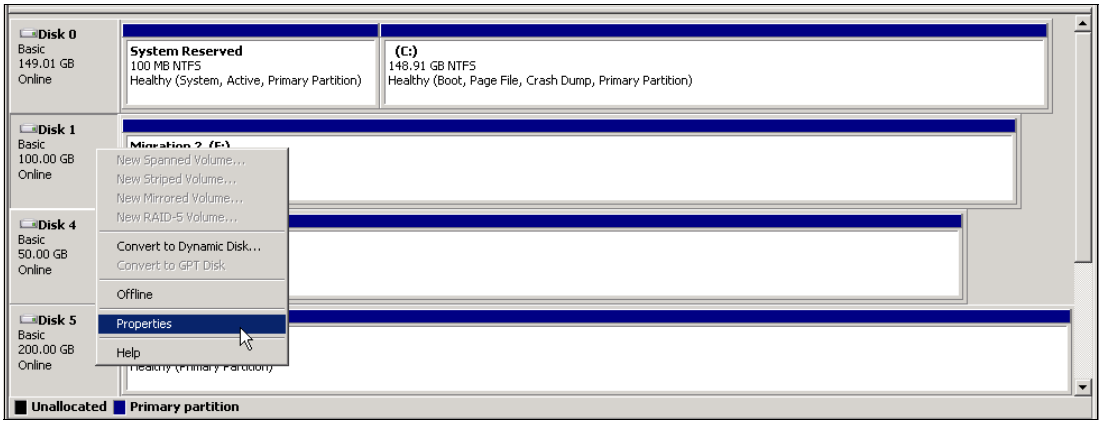

Figure 6-20 Display properties of disk before migration

After the disk properties panel is opened, the General tab shows the disk device type, as shown in Figure 6-21.

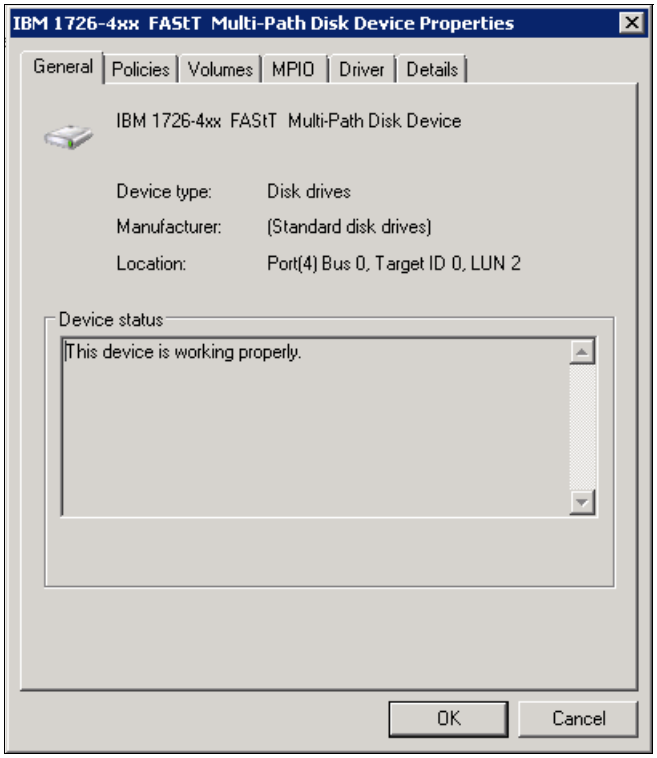

Figure 6-21 Window 2008 disk properties - general tab

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 V3500. If after the Storwize V3500 mapping and host rescan the disk device definitions are changed to IBM 2145 Multi-Path disk device, it confirms that the disks are under Storwize V3500 control.

Example scenario Fibre Channel cabling layout

Figure 6-22 shows the example scenario Fibre Channel cabling layout. The host, IBM DS3400, and Storwize V3500 are cabled into a dual SAN fabric configuration. The connection method that is shown can provide improved availability through fabric and path redundancy. It also can provide improved performance through workload balancing.

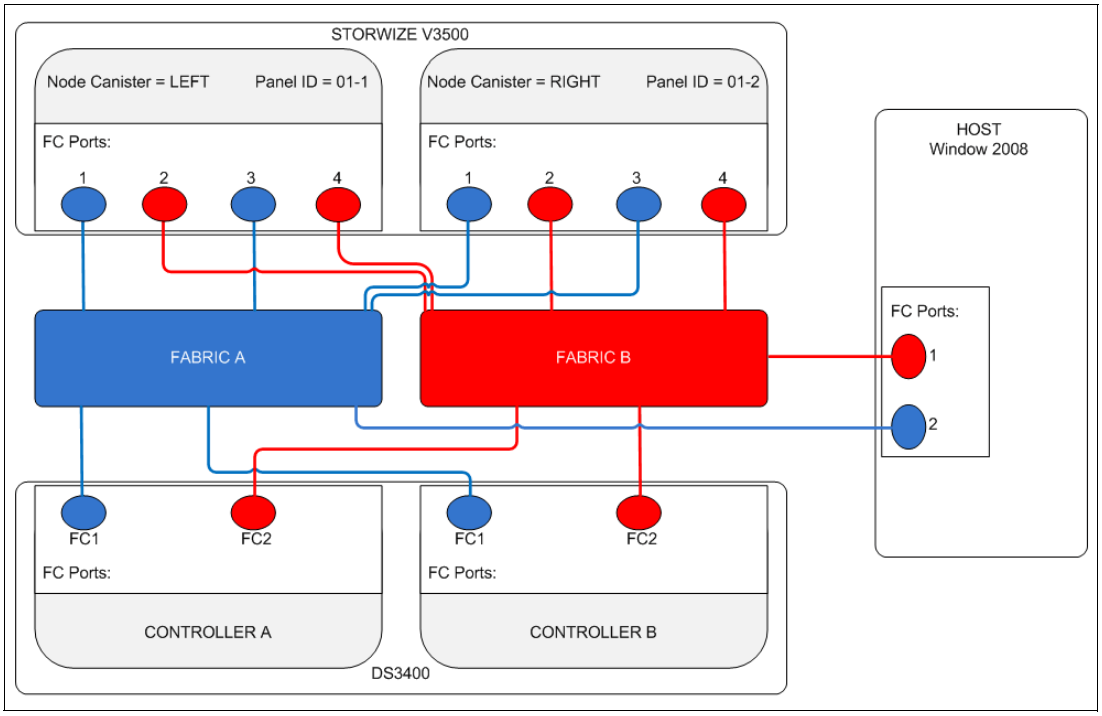

Figure 6-22 Example scenario Fibre Channel cabling layout

6.3.2 Using the storage migration wizard for example scenario

The following 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 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 IBM System Storage Interoperation Center (SSIC) for scenario compatibility.
2. Back up all data that is associated with the host, DS3400, and Storwize V3500.
3. Start New Migration to open the wizard on the Storwize V3500.
4. Follow step 1 of the wizard before you begin.
5. Follow step 2 of the wizard prepare 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 in 4.2.1, “Windows 2008 (R2):

Preparing for FC attachment" on page 143 to prepare the host to connect to Storwize V3500 that uses Fibre Channel.

Pay attention to the following tasks:

- Make sure that the latest OS service pack and test fixes are applied to your Microsoft server.
- Use the latest firmware and driver levels on your host system.
- Install HBA or HBAs on the Windows server by using the latest BIOS and drivers.
- Connect the FC Host Adapter ports to the switches.
- Configure the switches (zoning).
- Configure the HBA for hosts that are running Windows.
- Set the Windows time-out value.
- Install the 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:

- Locate the World Wide Port Numbers (WWPNs) for Host.
- Locate WWPNs for IBM DS3400
- Locate WWPNs for Storwize V3500
- Define port aliases definitions on SAN.
- Add V3500-to-DS3400 zones on SAN.
- Add Host-to-V3500 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 V3500 host group on DS3400.

f. Configure the storage system for use with this system.

Follow the IBM Storwize V3500 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 being migrated.

Create 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 you prepared your environment.

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 V3500 Host Group on DS3400.

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

Update DS3400 LU table.

7. Follow step 4 of the wizard to migrate MDisks. Select discovered MDisk on Storwize V3500.
8. In step 5 of the wizard, configure hosts by completing the following steps:
   a. Create Host on Storwize V3500.
   b. Select Host on Storwize V3500.
9. In step 6 of the wizard, map volumes to hosts by completing the following steps:
   a. Map volumes to Host on Storwize V3500.
   b. Verify disk device type is now 2145 on Host.
   c. SDDDSM datapath query commands on Host.
10. In step 7 of the wizard, select storage pool. Select internal storage pool on Storwize V3500.
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 IBM System Storage Interoperation Center (SSIC) for scenario compatibility.
2. Back up all data that is associated with the host, DS3400, and Storwize V3500.
3. Start New Migration to open the wizard on the Storwize V3500, as shown in Figure 6-23.

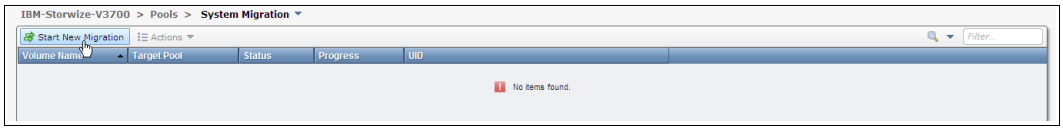

Figure 6-23 Start new migration

4. Follow step 1 of the storage migration wizard and check all of the restrictions and prerequisites, Figure 6-24. Click **Next** to continue.

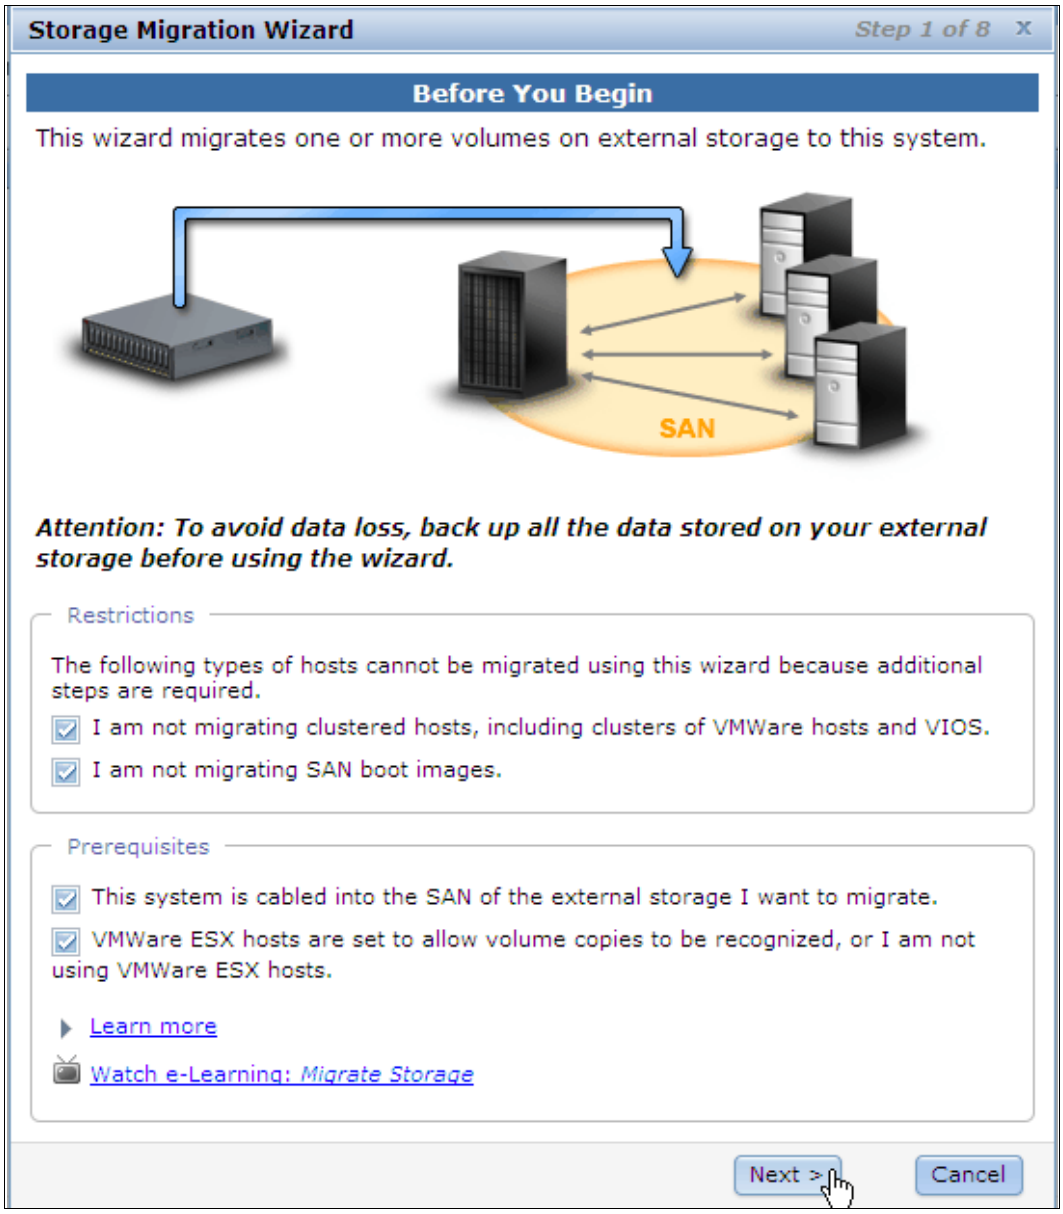

Figure 6-24 Storage migration wizard: Step 1

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

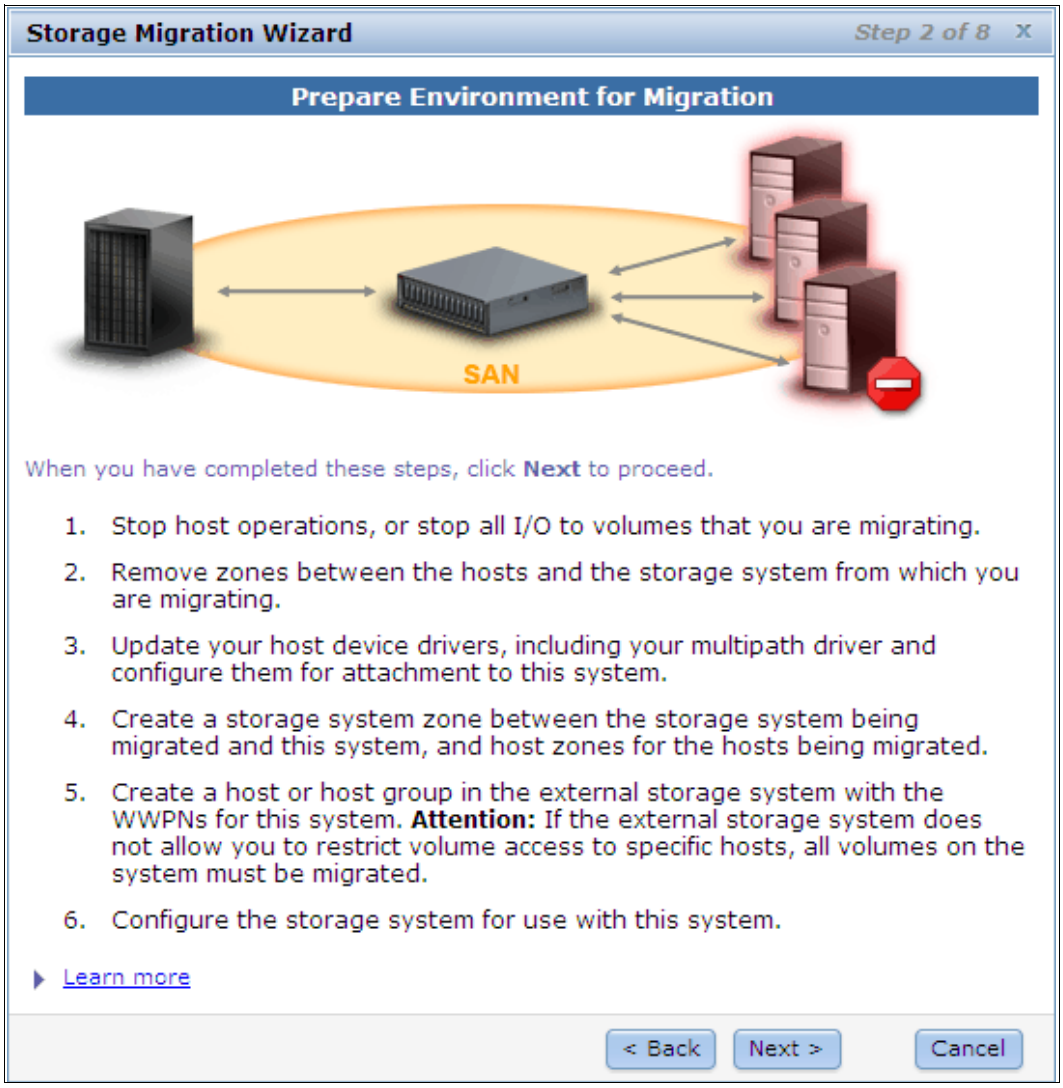

Figure 6-25 Storage migration wizard: Step 2

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. Update your host device drivers (including your multipath driver) and configure them for attachment to this system. Complete the steps in 4.2.1, “Windows 2008 (R2): Preparing for FC attachment” on page 143 to prepare the host to connect to Storwize V3500 by using Fibre Channel.

   Pay attention to the following tasks:

   - Make sure that the latest OS service pack and test fixes are applied to your Microsoft server.
   - Use the latest firmware and driver levels on your host system.

- Install HBA or HBAs on the Windows server that use the latest BIOS and drivers.
- Connect the FC Host Adapter ports to the switches.
- Configure the switches (zoning).
- Configure the HBA for hosts that is running Windows.
- Set the Windows timeout value.
- 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 is migrated.

   To perform this step, locate the WWPNs of the host, IBM DS3400, and Storwize V3500, then create an alias for each port to simplify the zone creation steps.

   Important: A WWPN is the World Wide Port Name. It 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

Refer to the original IBM DS3400 Host definition to locate the WWPNs of the host's dual port QLE2562 Host Bus Adapter (HBA). To perform 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.

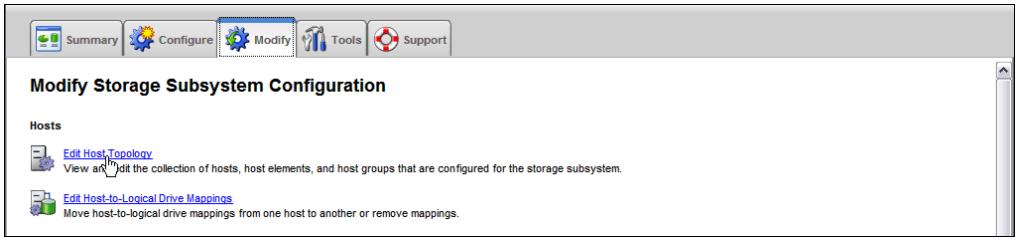

Figure 6-26 IBM DS3400 modify tab: Edit Host Topology

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

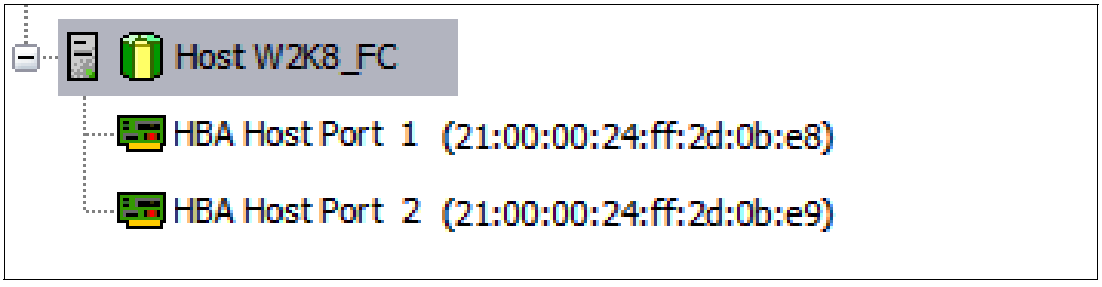

Figure 6-27 IBM DS3400 host definition

Record the WWPNs for alias, zoning, and the Storwize V3500 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 host's WWPNs.

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 and click **Support,** and select **View Storage Subsystem Profile.** Figure 6-28 shows the IBM DS3400 Storage Manager Support tab.

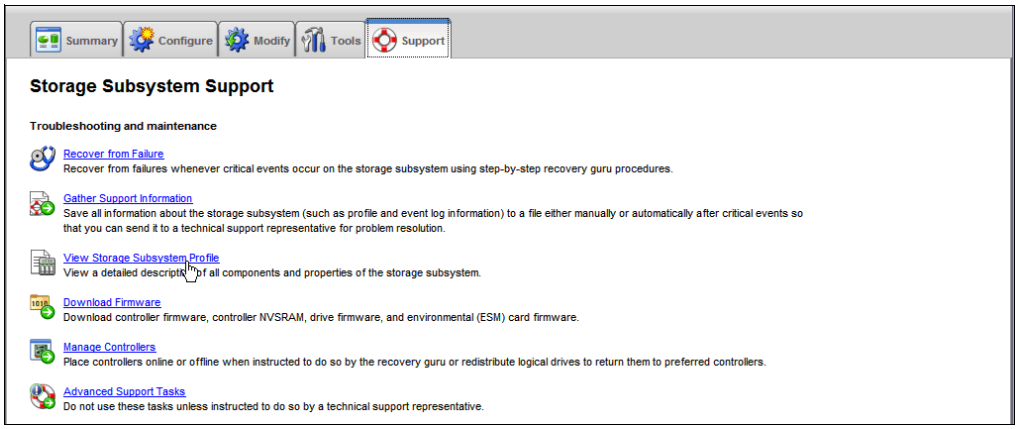

Figure 6-28 IBM Storage Manager: View storage subsystem profile

Click the Controllers tab to show the WWPNs for each controller. Figure 6-29 shows the IBM Ds3400 storage manager storage subsystem profile.

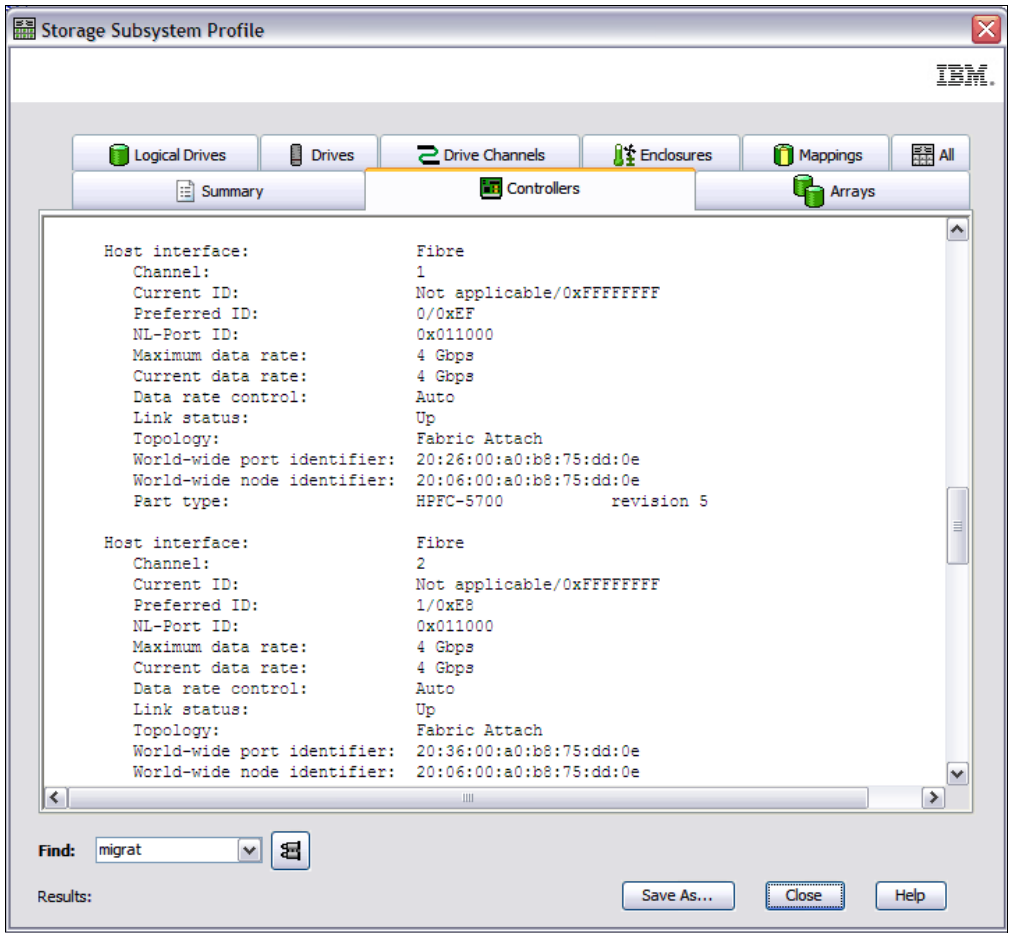

Figure 6-29 IBM DS3400 storage manager: Storage subsystem profile: controller WWPNs

Locating node canister WWPNs on the Storwize V3500

To locate the WWPNs for the Storwize V3500 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 V3500 System Details panel with the WWPNs shown for **IBM-Storwize-V3500** → **Enclosure 1** → **Canister 1**.

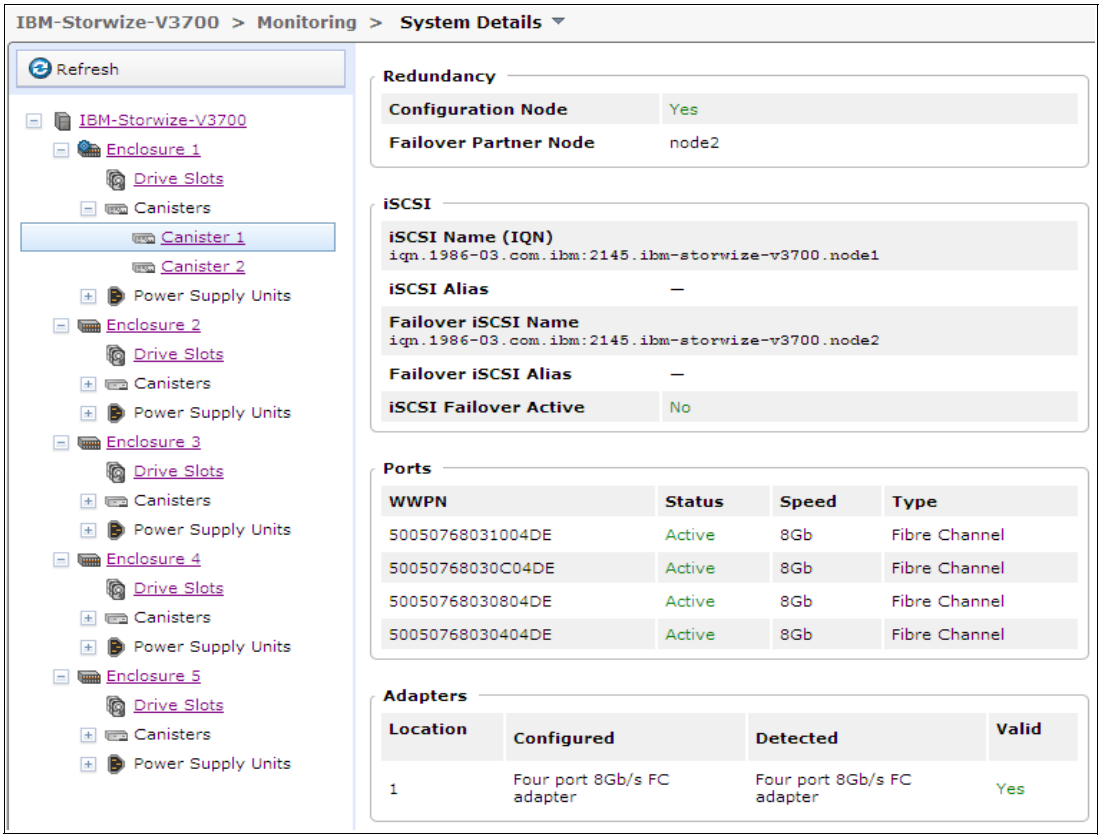

Figure 6-30 Storwize V3500 node canister WWPNs information

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

Example scenario Storwize V3500 and IBM DS3400 WWPN diagram

Each port on the Storwize V3500 and IBM DS3400 has a unique and persistent WWPN. This configuration means that if an HA (Host Adapter) 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 shows the relationship between the device WWPNs and the Fibre Channel ports for the Storwize V3500 and the IBM DS3400 that are used in the example scenario.

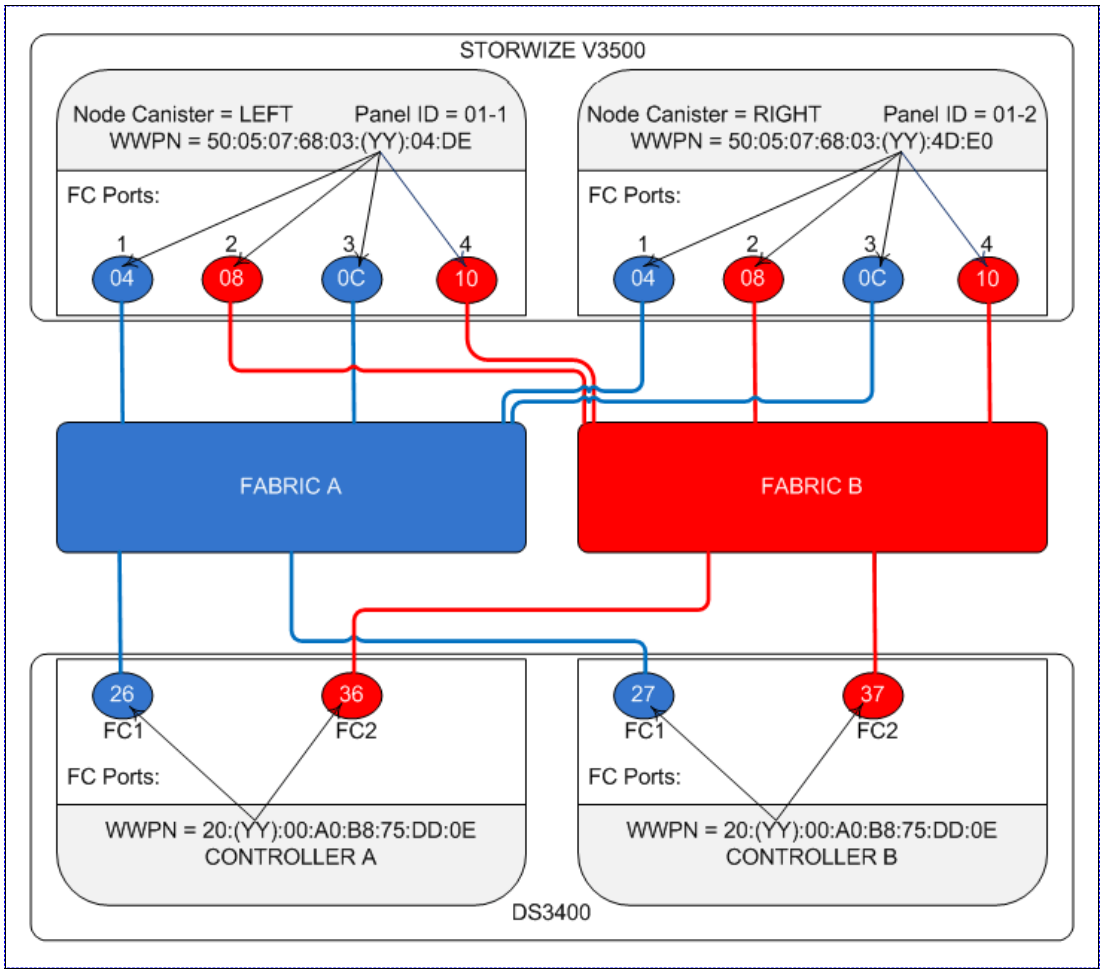

Figure 6-31 Example scenario Storwize V3500 and IBM DS3400 WWPN location diagram

Zoning: Define aliases on the SAN fabrics

Now that the WWPNs for Storwize V3500, IBM DS3400, and Windows 2008 host are located, you can define the WWPN aliases on the SAN fabrics for the Storwize V3500. You also can create aliases for the DS3400 and Windows 2008 host 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 that when host HBA cards are replaced, there can be occasions when a new HBA card contains new WWPNs. As a result, 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 replacement.

See Figure 6-31 on page 243 for the following alias definitions:

```
Storwize V3500 ports connected to SAN Fabric A:
      alias= V3500_Canister_Left_Port1 wwpn= 50:05:07:68:03:04:04:DE
      alias= V3500_Canister_Left_Port3 wwpn= 50:05:07:68:03:0C:04:DE
      alias= V3500_Canister_Right_Port1 wwpn= 50:05:07:68:03:04:4D:E0
      alias= V3500_Canister_Right_Port3 wwpn= 50:05:07:68:03:0C:4D:E0
Storwize V3500 ports connected to SAN Fabric B:
      alias= V3500_Canister_Left_Port2 wwpn= 50:05:07:68:03:08:04:DE
      alias= V3500_Canister_Left_Port4 wwpn= 50:05:07:68:03:10:04:DE
      alias= V3500_Canister_Right_Port2 wwpn= 50:05:07:68:03:08:4D:E0
      alias= V3500_Canister_Right_Port4 wwpn= 50:05:07:68:03:10:4D:E0
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: Define the V3500-to-DS3400 zones on the SAN fabrics

Define the V3500-to-DS3400 zones on the SAN fabrics. The best way to zone DS3400-to-V3500 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 V3500 node canisters must detect the same ports on IBM DS3400 storage system.

See to Figure 6-31 on page 243 and the previously defined SAN aliases for the following zones definitions:

```
FABRIC A
Zone name= ALL_V3500_to_DS3400_CTRLA_FC1:
               DS3400_CTRLA_FC1
               V3500_Canister_Left_Port1
               V3500_Canister_Left_Port3
               V3500_Canister_Right_Port1
               V3500_Canister_Right_Port3
Zone name= ALL_V3500_to_DS3400_CTRLB_FC1:
               DS3400_CTRLB_FC1
               V3500_Canister_Left_Port1
               V3500_Canister_Left_Port3
               V3500_Canister_Right_Port1
               V3500_Canister_Right_Port3
FABRIC B
Zone name= ALL_V3500_to_DS3400_CTRLA_FC2:
               DS3400_CTRLA_FC2
               V3500_Canister_Left_Port2
               V3500_Canister_Left_Port4
```

```
              V3500_Canister_Right_Port2
              V3500_Canister_Right_Port4
Zone name= ALL_V3500_to_DS3400_CTRLB_FC2:
              DS3400_CTRLB_FC2
              V3500_Canister_Left_Port2
              V3500_Canister_Left_Port4
              V3500_Canister_Right_Port2
              V3500_Canister_Right_Port4
```

Zoning: Define the Host-to-V3500 zones on the SAN fabrics

Define the Host-to-V3500 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 235 and the previously defined SAN aliases for the following host zone definitions:

```
FABRIC A
Zone name= W2K8_HOST_P2_to_V3500_Port1s:
              W2K8_HOST_P2
              V3500_Canister_Left_Port1
              V3500_Canister_Right_Port1
FABRIC B
Zone name= W2K8_HOST_P1_to_V3500_Port2s:
              W2K8_HOST_P1
              V3500_Canister_Left_Port2
              V3500_Canister_Right_Port2
```

Important: The configuration of an intra-cluster zone (V3500-to-V3500) on each fabric is recommended. Place all Storwize V3500 port aliases from each node canister into the one zone on each of the fabrics. This 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 V3500, which contains two hosts. Each host is a node canister of the Storwize V3500.

Creating an IBM DS3400 Host Group

To define a new Host Group for the Storwize V3500 by using the DS3400 Storage Manager, click the **Configure** tab and select **Create Host Group** to open the Create Host Group panel, as shown in Figure 6-32.

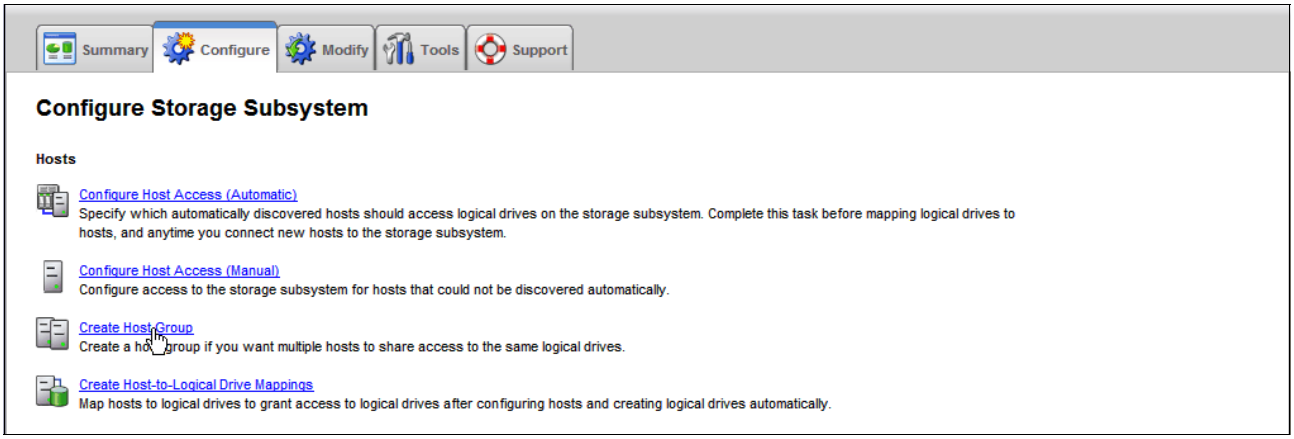

Figure 6-32 IBM DS3400 Configure tab: Create Host Group

By using the IBM DS3400 Storage Manager, create a Host Group named Storwize_V3500, as shown in Figure 6-33.

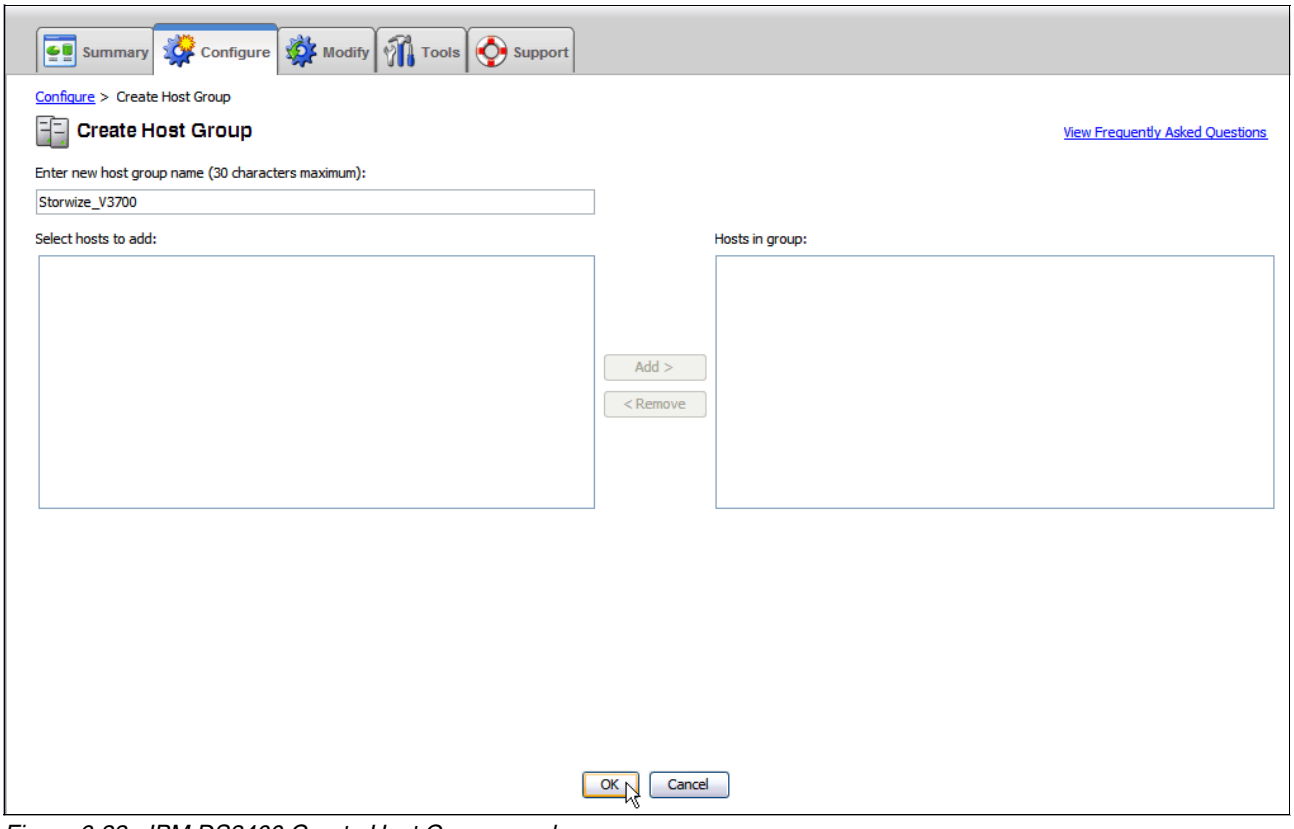

Figure 6-33 IBM DS3400 Create Host Group panel

Creating IBM DS3400 hosts

By using the IBM DS3400 Storage Manager, create a Host for each node canister of the Storwize V3500. To define a new Host by using the DS3400 Storage Manager, click the **Configure** tab and select **Configure Host-Access (Manual)** to open the Configure Host Access panel. Figure 6-34 shows the IBM DS3400 Configure tab.

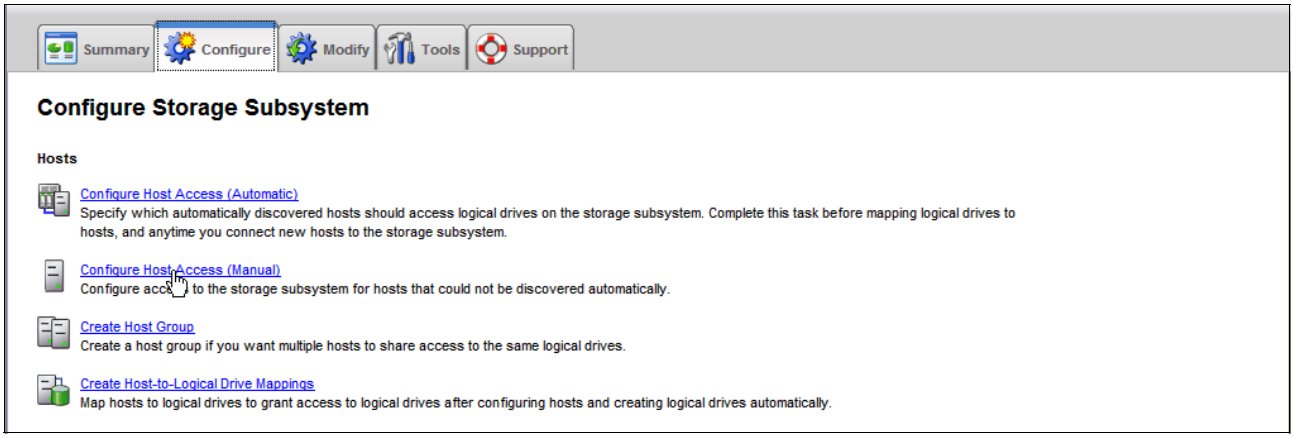

Figure 6-34 IBM DS3400 Storage Manger configure tab: Create host

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_V3500_Canister_Left and Storwize_V3500_Canister_Right. Click **Next** to continue. Figure 6-35 shows the IBM DS3400 Configure Host Access (Manual) panel.

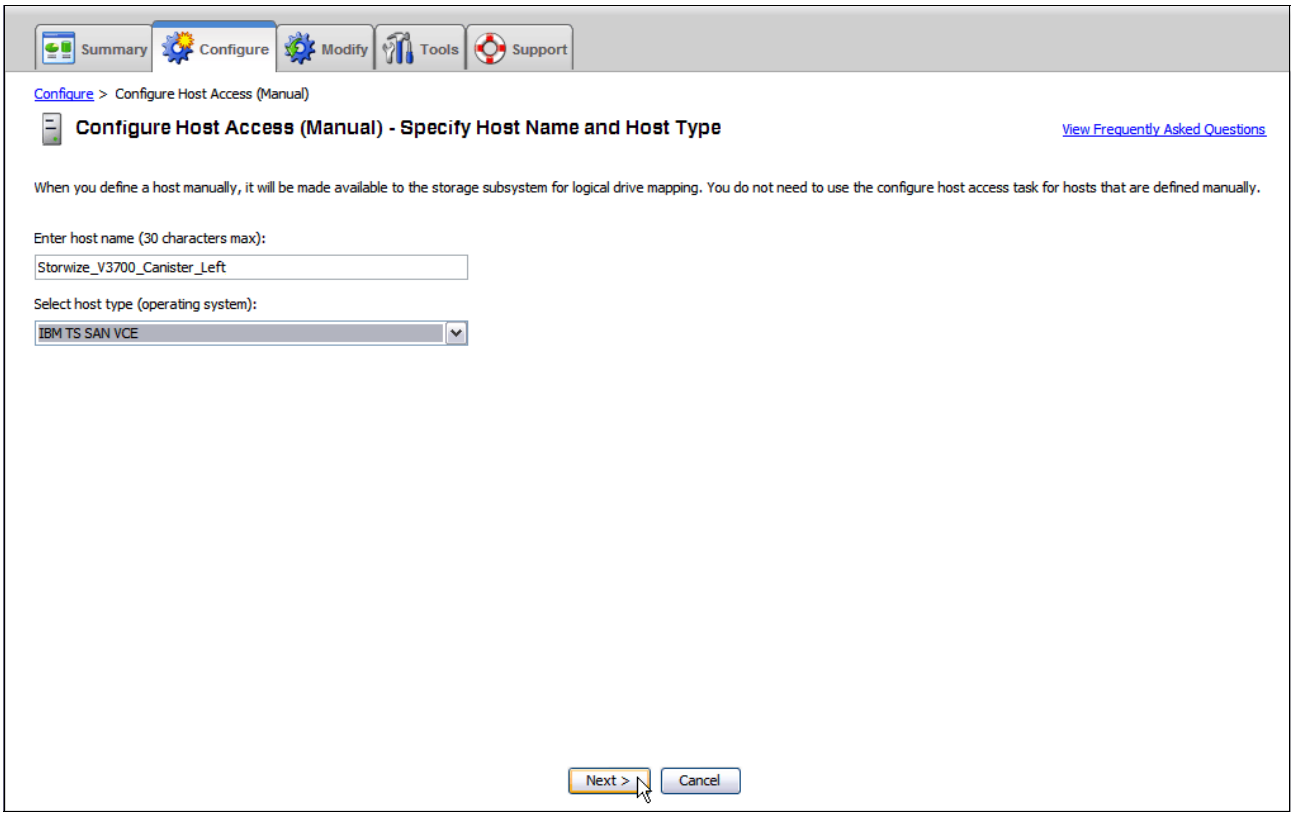

Figure 6-35 IBM DS3400 Storage Manager Configure tab: Configure host

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 move to the right side of the panel. Figure 6-36 shows the IBM DS3400 Specify HBA Host Ports panel.

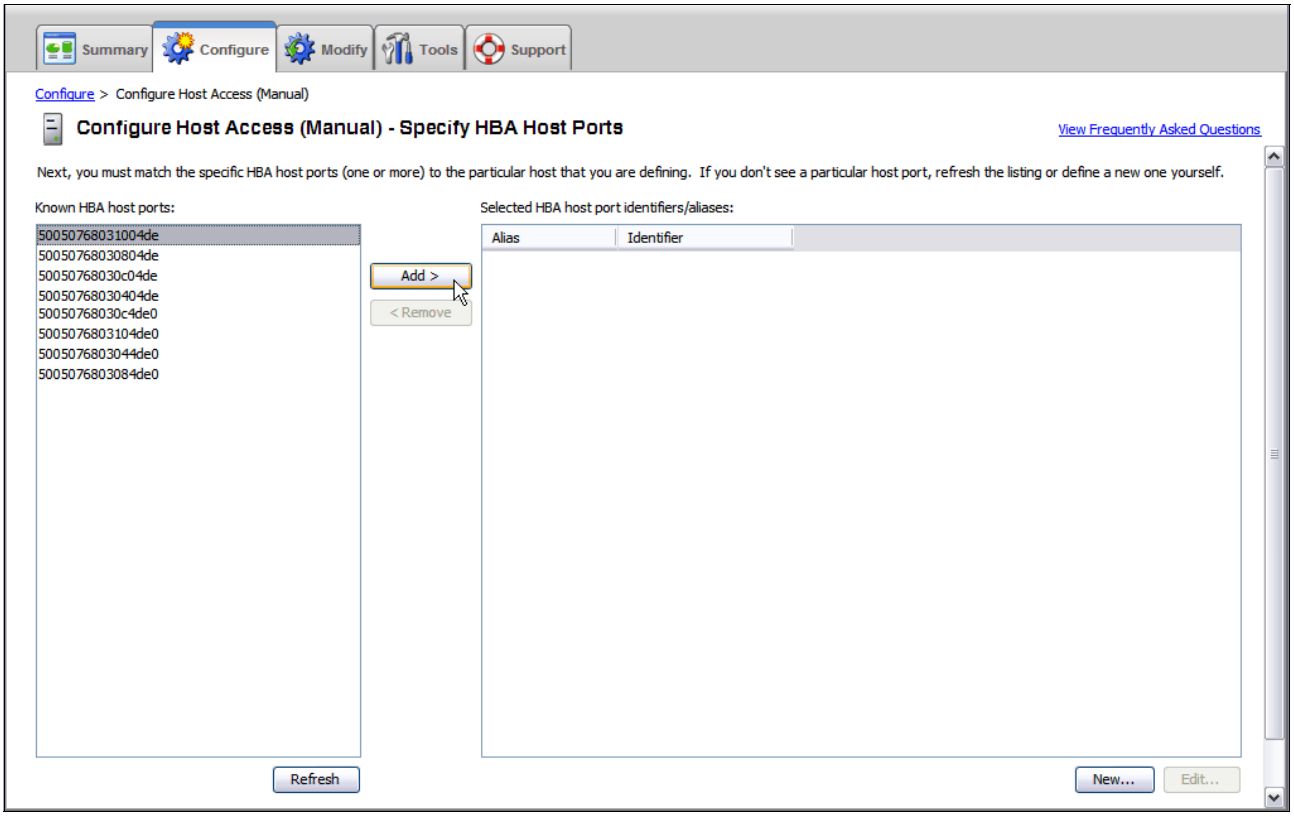

Figure 6-36 IBM DS3400 Specify HBA Host Ports

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

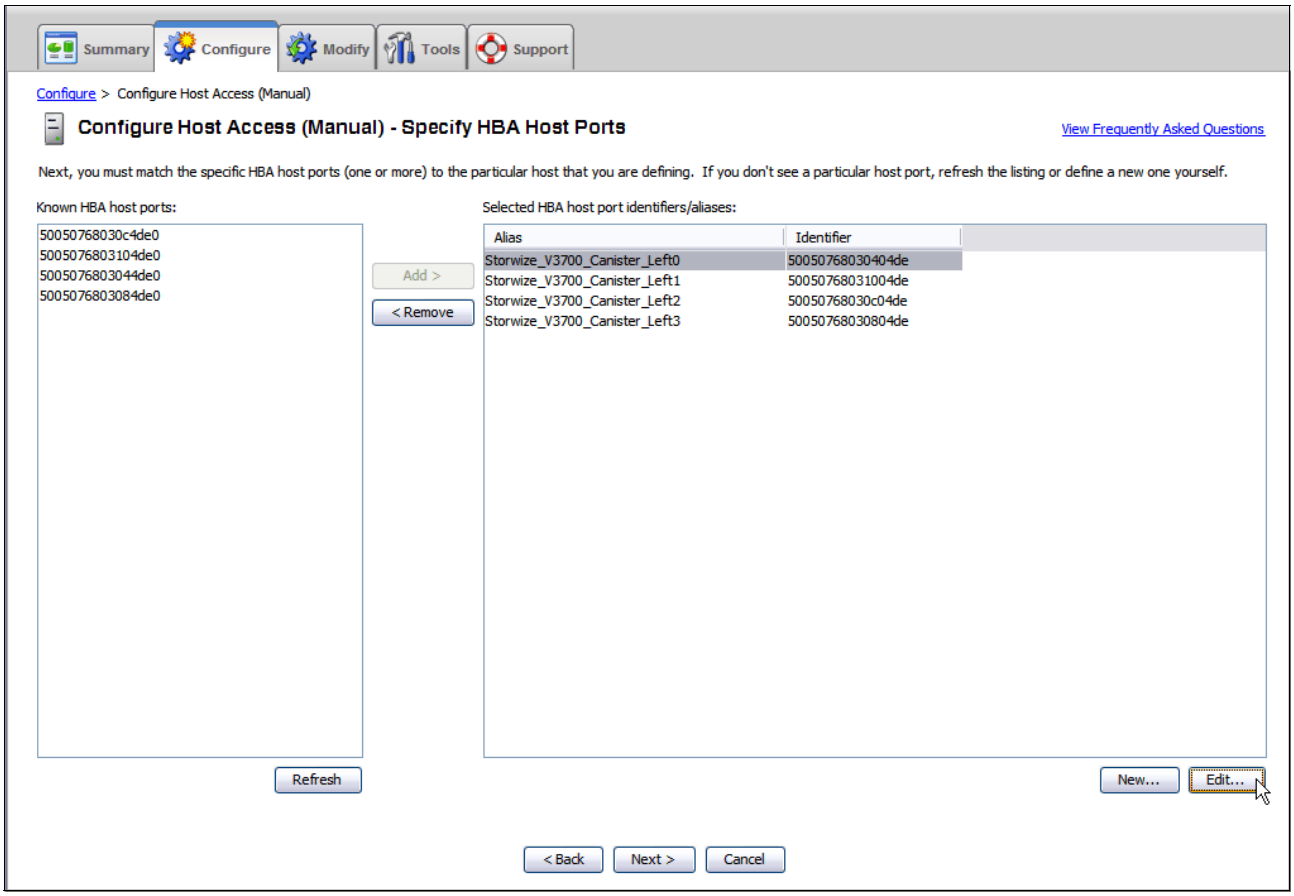

Figure 6-37 IBM DS3400 storage manager specify HBA host ports: edit alias

Figure 6-38 shows the Edit HBA Host Port panel. Enter a meaningful aliases for each of the WWPNs, such as V3500_Canister_Left_P1.

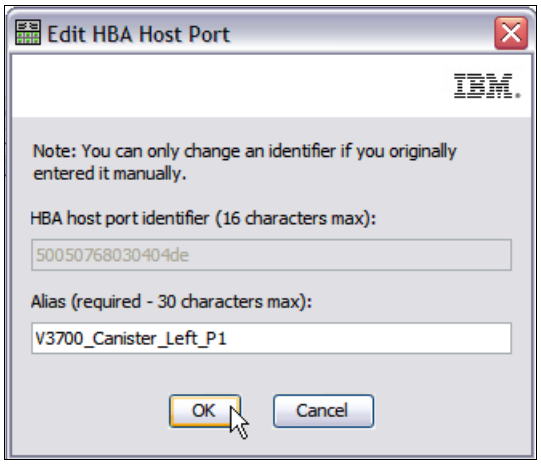

Figure 6-38 IBM DS3400 edit HBA host port panel

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 were added to the host definition on the IBM DS3400 Specify HBA Host Ports panel.

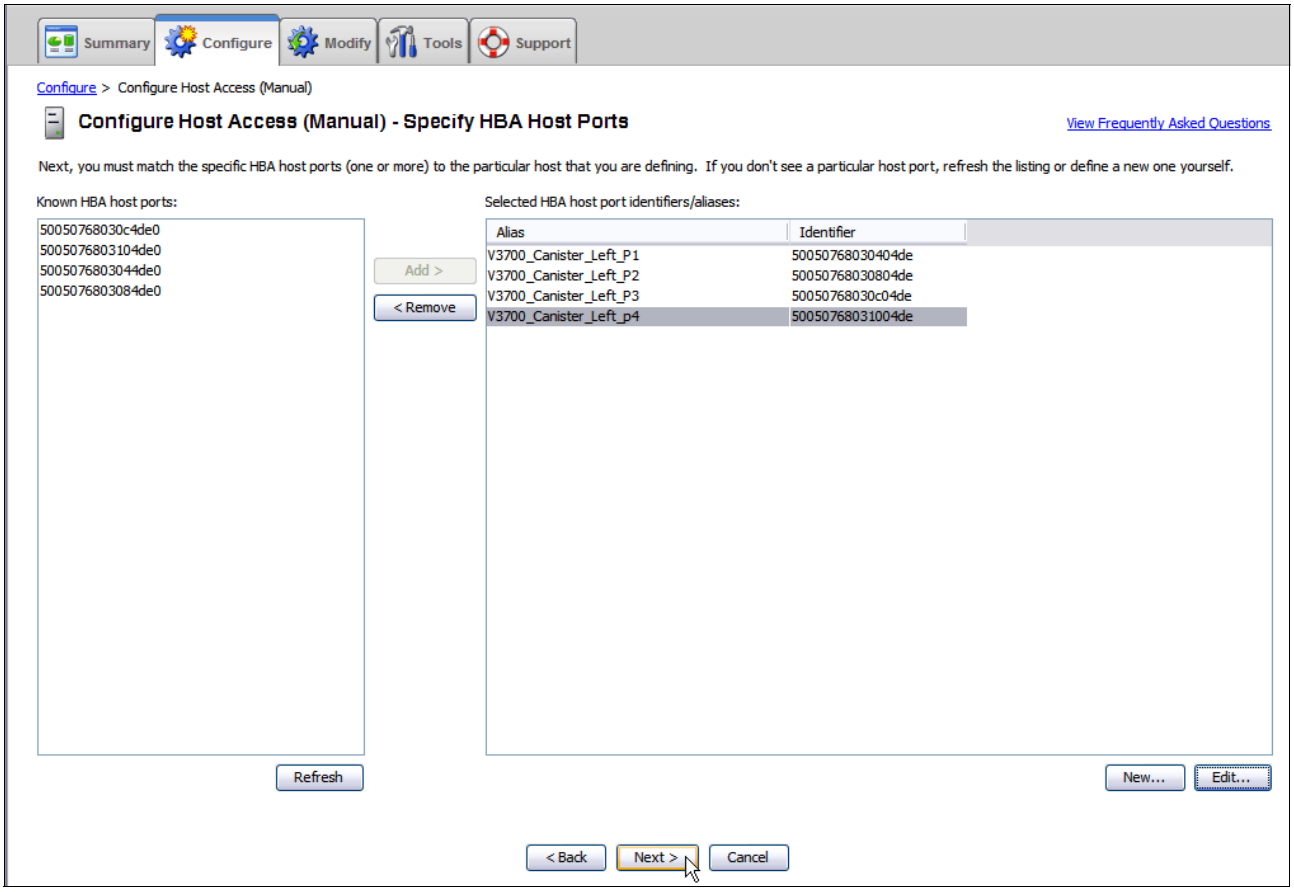

Figure 6-39 IBM DS3400 Specify HBA Host Ports

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_V3500** host group. Select **Next** to continue.

Figure 6-40 shows the IBM DS3400 Specify Host Group panel.

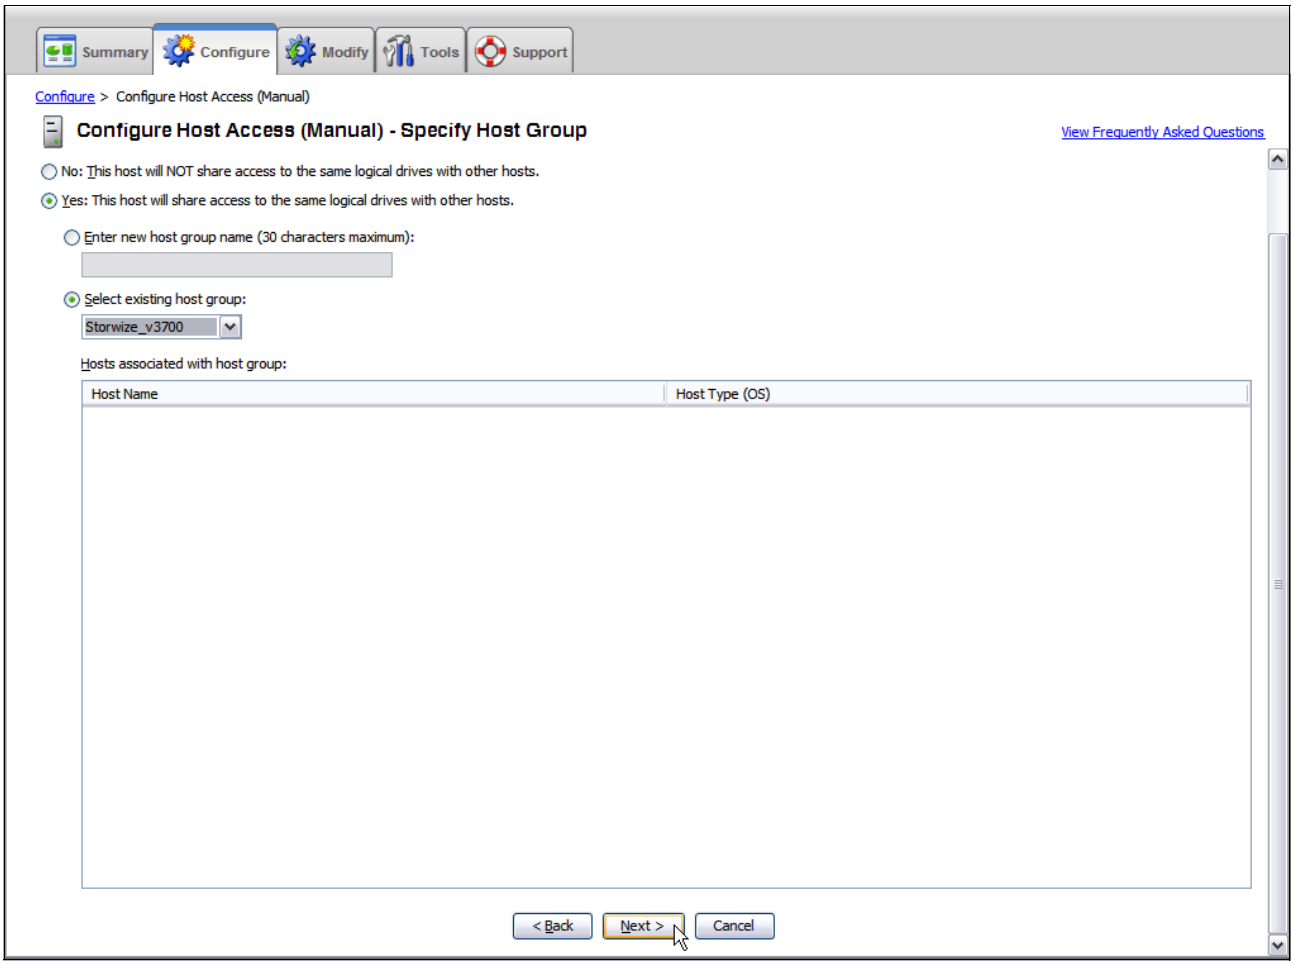

Figure 6-40 IBM DS3400 Specify Host Group panel

A summary panel of the defined host and its associated host group is displayed. Cross-check and confirm the host definition summary, and then select **Finish**. Figure 6-41 shows the IBM DS3400 Confirm Host Definition panel.

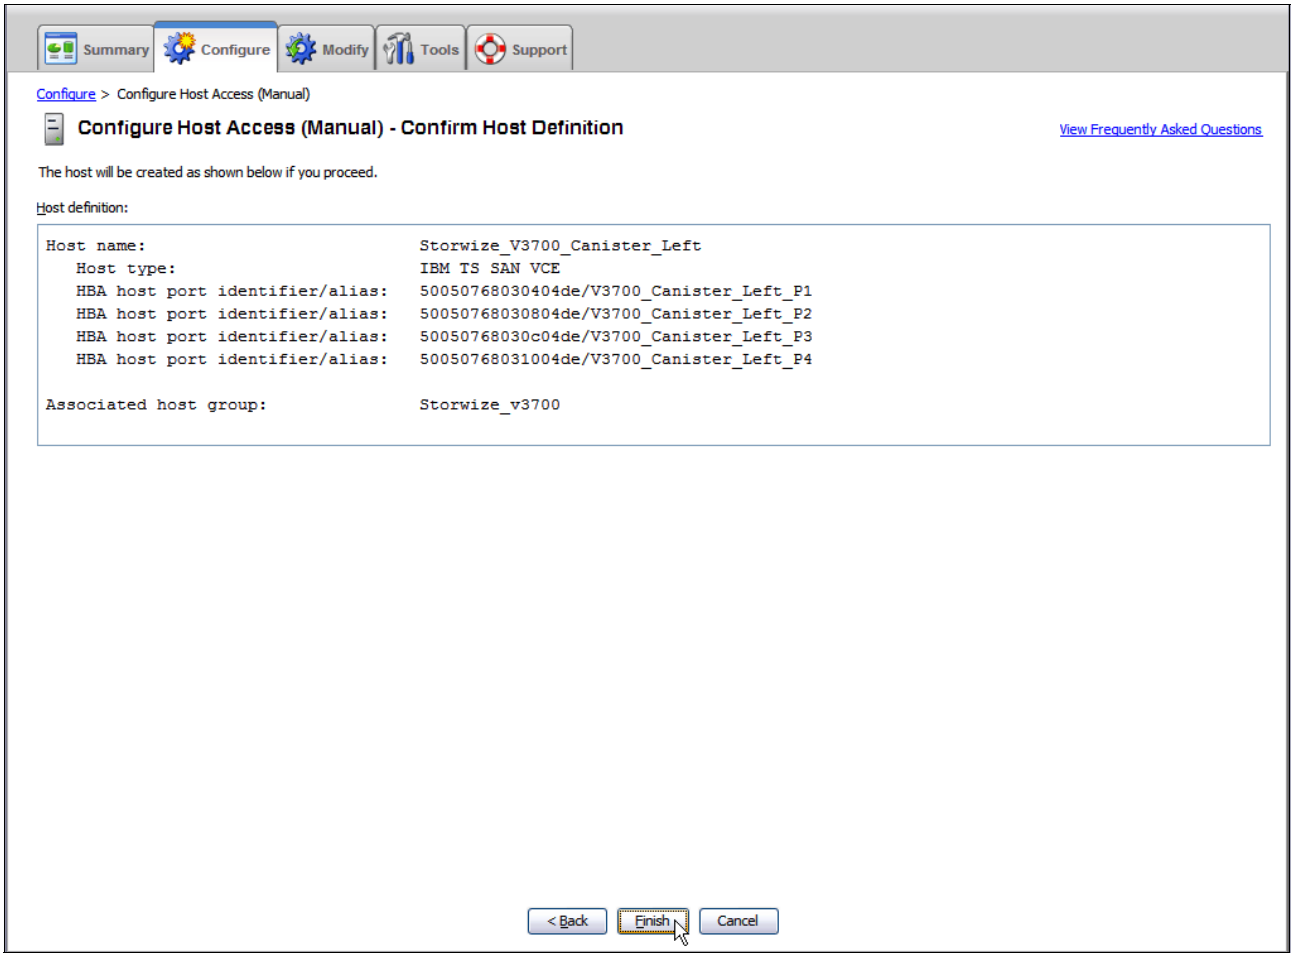

Figure 6-41 IBM DS3400 Confirm Host Definition

A host definition must be created for the other node canister. The host definition must be associated to the Host Group Storwize_V3500. To configure the other node canister, complete the steps that are described in “Creating IBM DS3400 hosts” on page 247.

The node canister Host definitions are logically contained in the Storwize_V3500 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 manage and click the **Modify** tab and select **Edit Host Topology**, as shown in Figure 6-42.

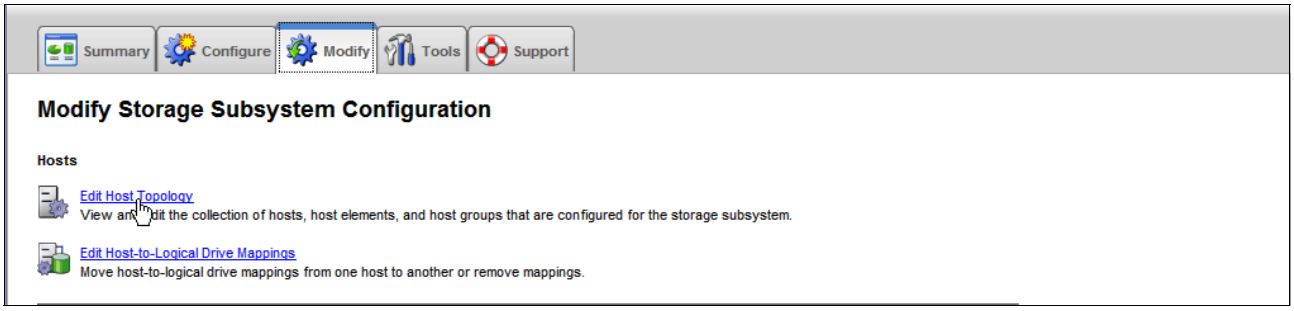

Figure 6-42 IBM DS3400 Modify tab: Edit Host Topology

Figure 6-43 shows the host topology of the defined Storwize_V3500 Host Group with both of the created node canister hosts, as seen through the DS3400 Storage Manager software.

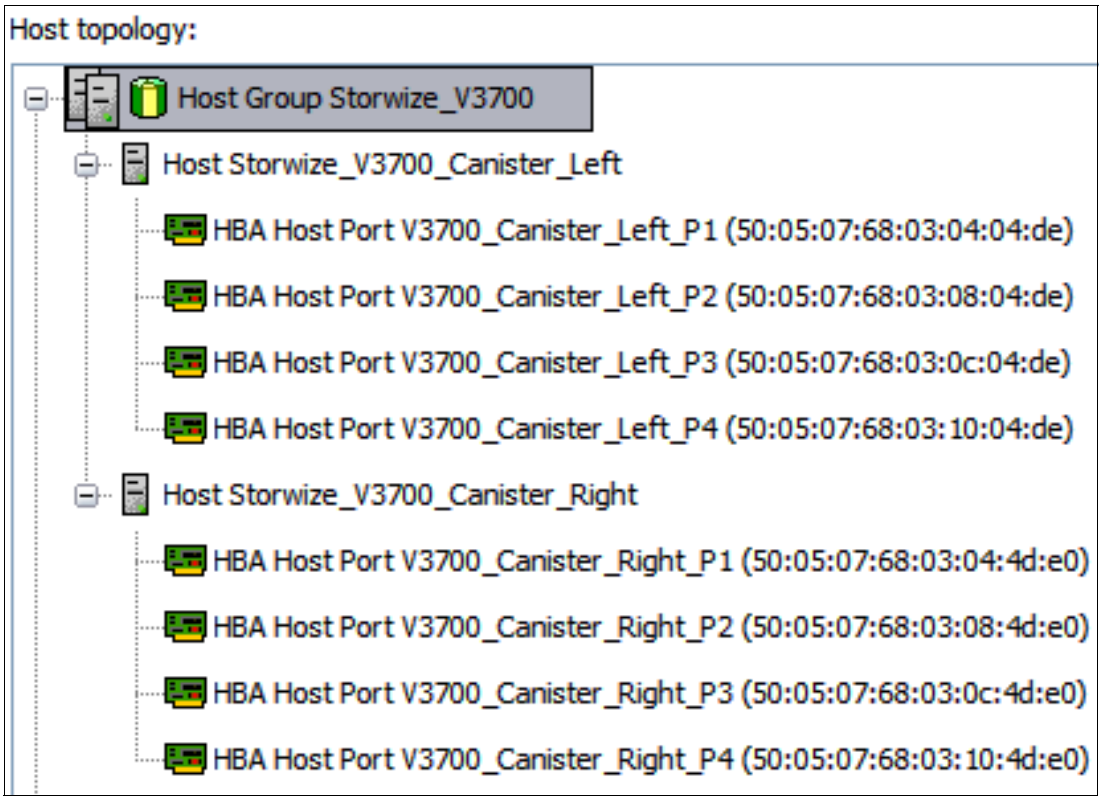

Figure 6-43 IBM DS3400 host group definition for the Storwize V3500

Configure the storage system for use with this system.

Follow the IBM Storwize V3500 Version 6.4.1 Information Center for DS3400 configuration recommendations that are found at this website:

http://pic.dhe.ibm.com/infocenter/storwize/v3500_ic/index.jsp?topic=%2Fcom.ibm.storwize.V3500.641.doc%2Fsvc_configdiskcontrollersovr_22n9uf.html

Now that the environment is prepared, return to step 2 of the storage migration wizard and click **Next** to continue, as shown in Figure 6-44.

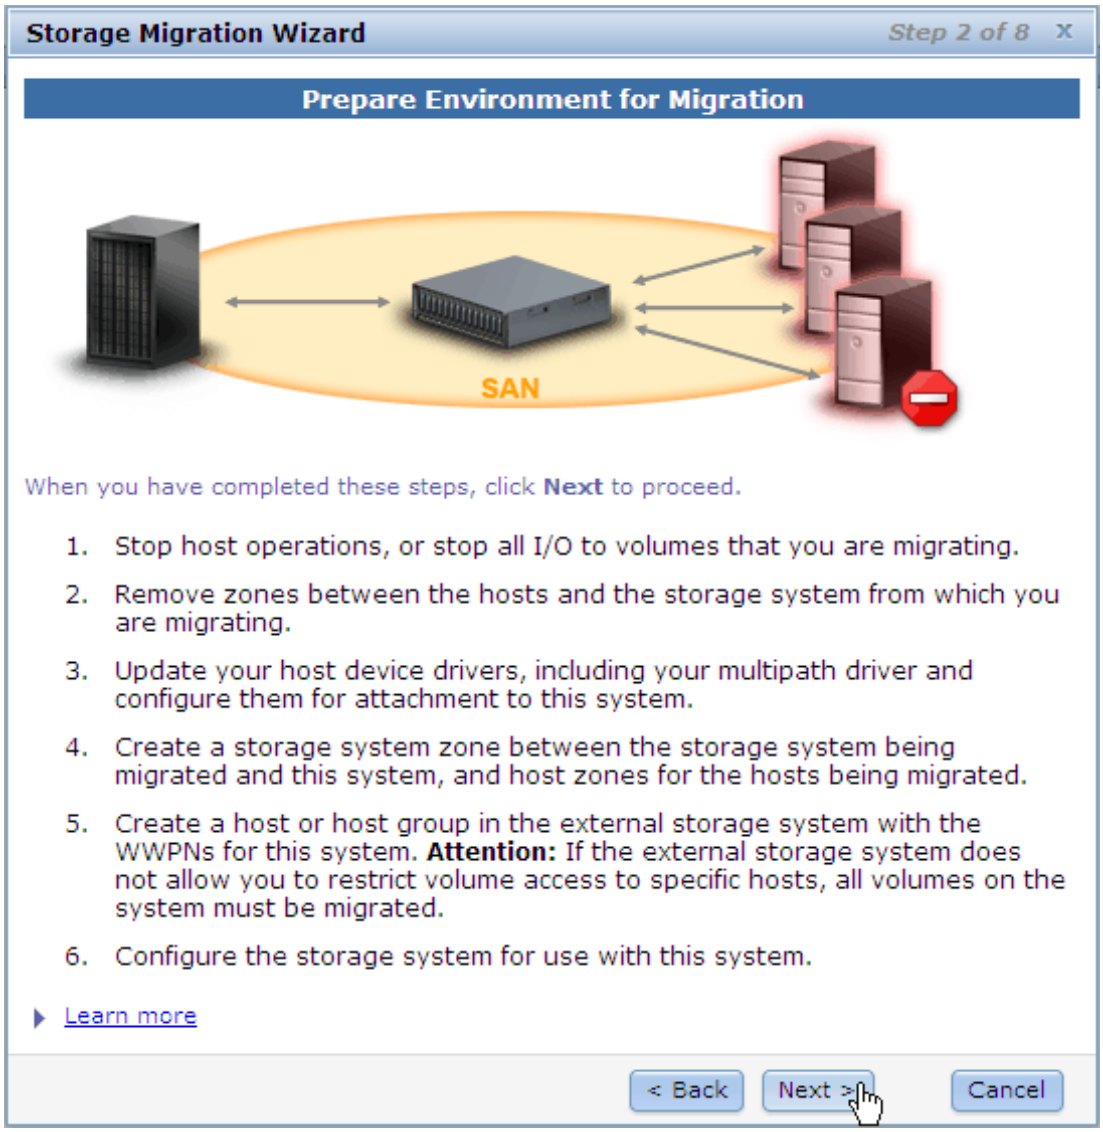

Figure 6-44 Step 2 of the storage migration wizard

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

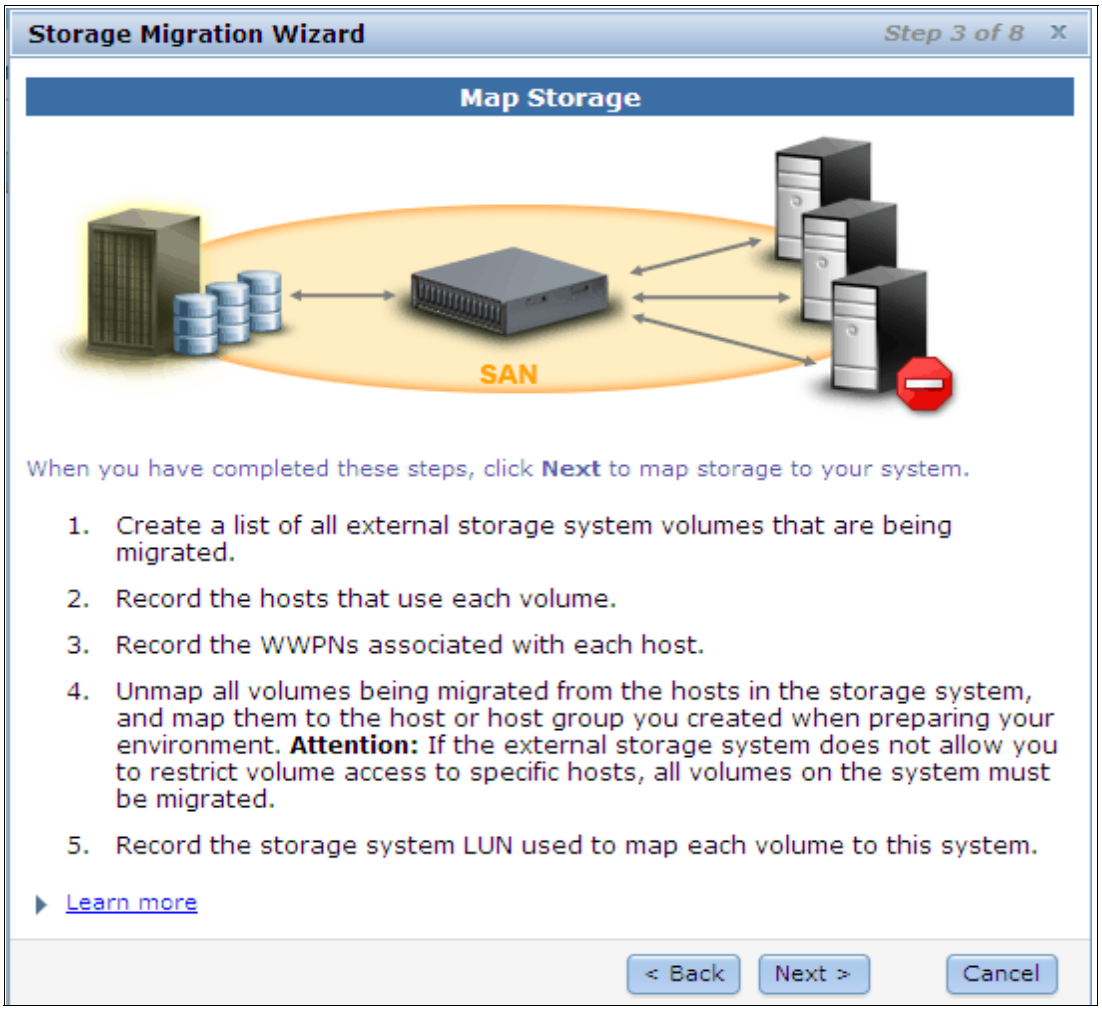

Figure 6-45 Step 3 of the storage migration wizard

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

Table 6-3 shows a list of the IBM DS3400 logical units that are migrated and the host that uses them.

Table 6-3 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 associated to the host can be seen in Table 6-4. It is also recommended to record the HBA firmware, HBA device driver version, adapter information, OS, and V3500 multi-path software version, if possible.

Table 6-4 WWPNs associated with host

Host Name	Adapter/Slot/Port	WWPNs	HBA F/W	HBA Device Driver	Operating System	V3500 Multipath S/W
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, Map them to the host or host group that you created when you were preparing your environment.

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 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 V3500 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**. Figure 6-46 shows the IBM DS3400 Storage Manager Modify tab.

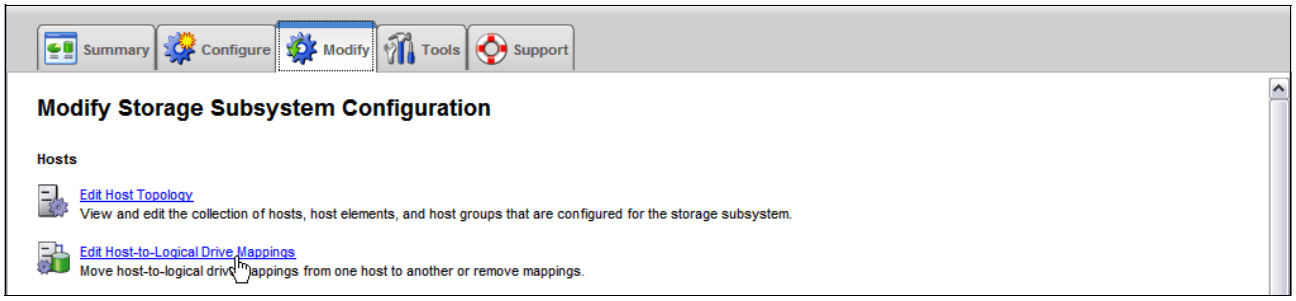

Figure 6-46 IBM DS3400 storage manager modify tab

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.

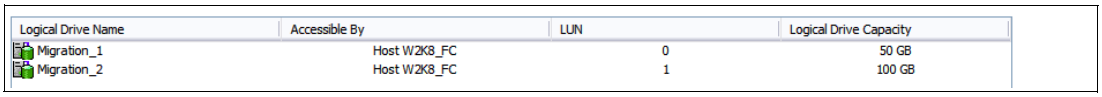

Figure 6-47 IBM DS3400 Logical drives mapping information before changes

To modify the mapping definition so that the LUs are accessible only by the Storwize V3500 Host Group, select **Change...** to open the Change Mapping panel and modify the mapping.

This step ensures that the LU is inaccessible from the Windows 2008 Host. Figure 6-48 shows the **Change...** selection in Modify Mapping panel of the DS3400.

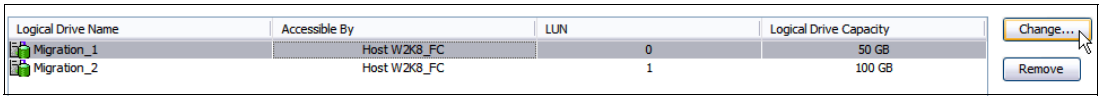

Figure 6-48 IBM DS3400 Modify Mapping panel: Change mapping

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

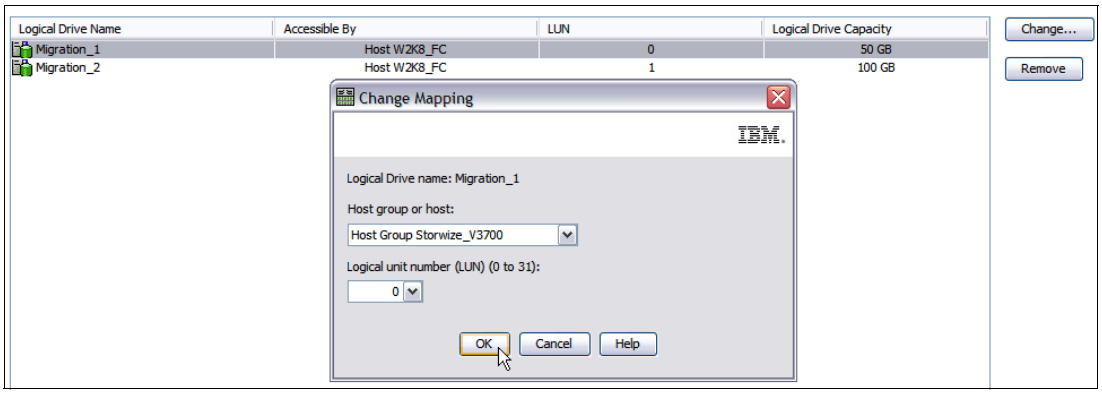

Figure 6-49 IBM DS3400 Change Mapping panel

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

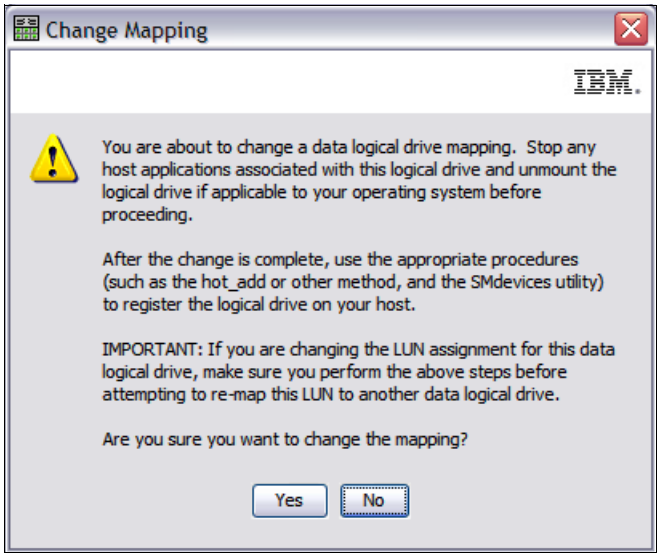

Figure 6-50 Change Mapping confirmation panel

Repeat the steps that are described in “” on page 256 for each of the LUs that are migrated. Confirm that the Accessible By column now reflects the mapping changes. Figure 6-51 shows both logical drives are now accessible by Host Group Storwize_V3500.

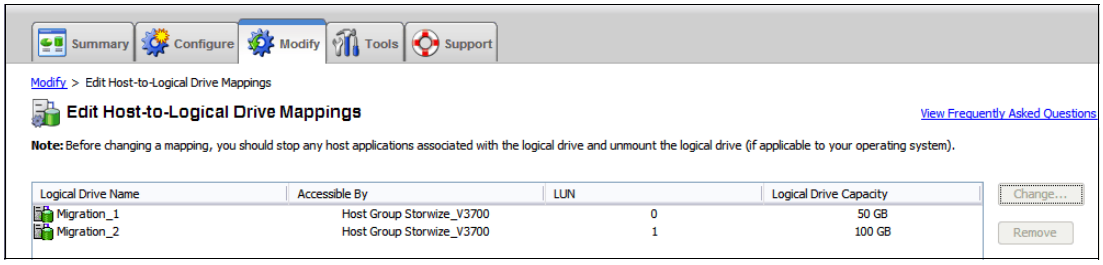

Figure 6-51 iBM DS3400 storage manager modify panel: edit host-to-logical drive mappings

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 255. Now that step 3 of the storage migration wizard is complete, click **Next** to show the Detect MDisks running task. Figure 6-52 shows step 3 of the storage migration wizard.

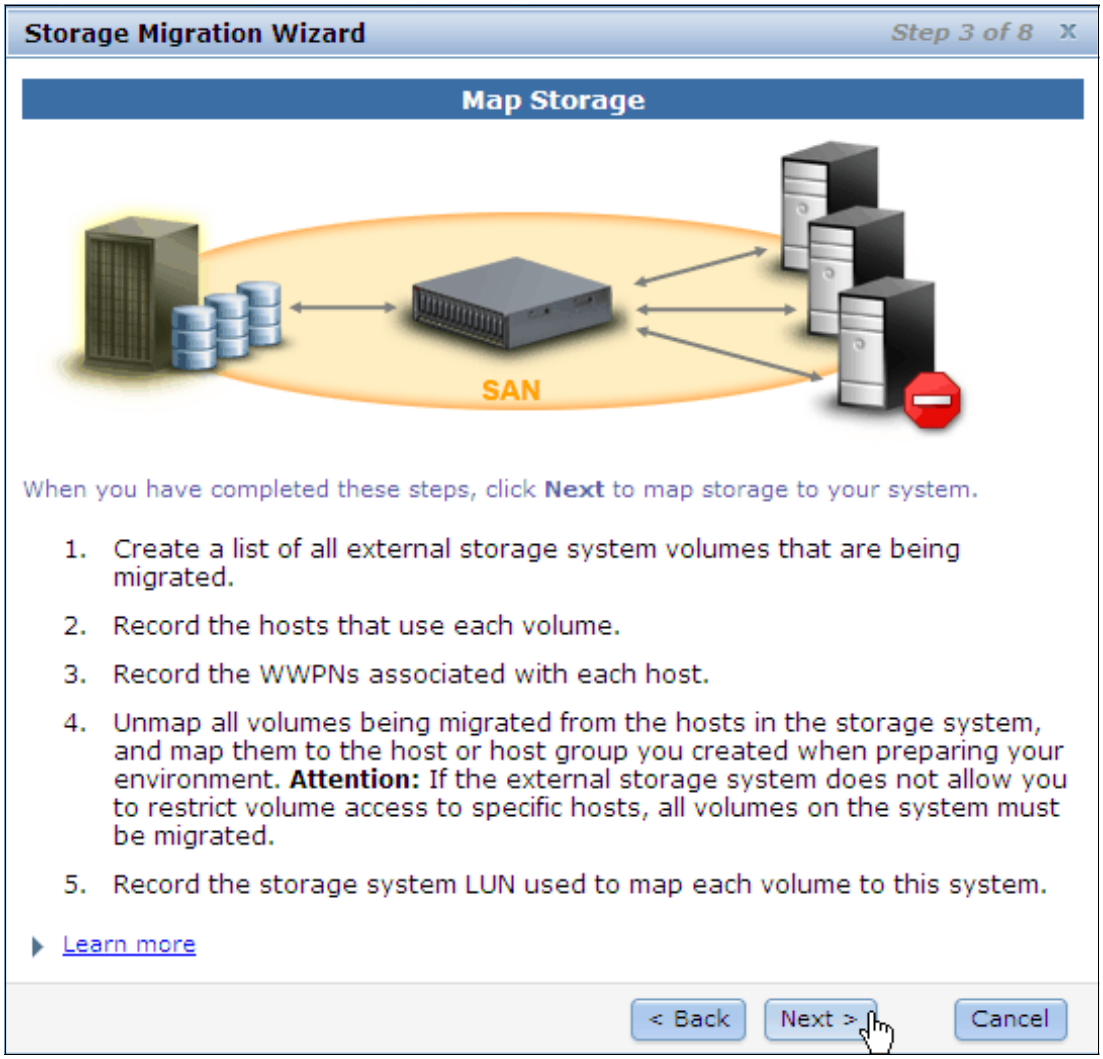

Figure 6-52 Shows step 3 of the storage migration wizard

After the Discover Devices running task is complete, select **Close** to show step 4 of the storage migration wizard. Figure 6-53 shows the Discover Devices panel.

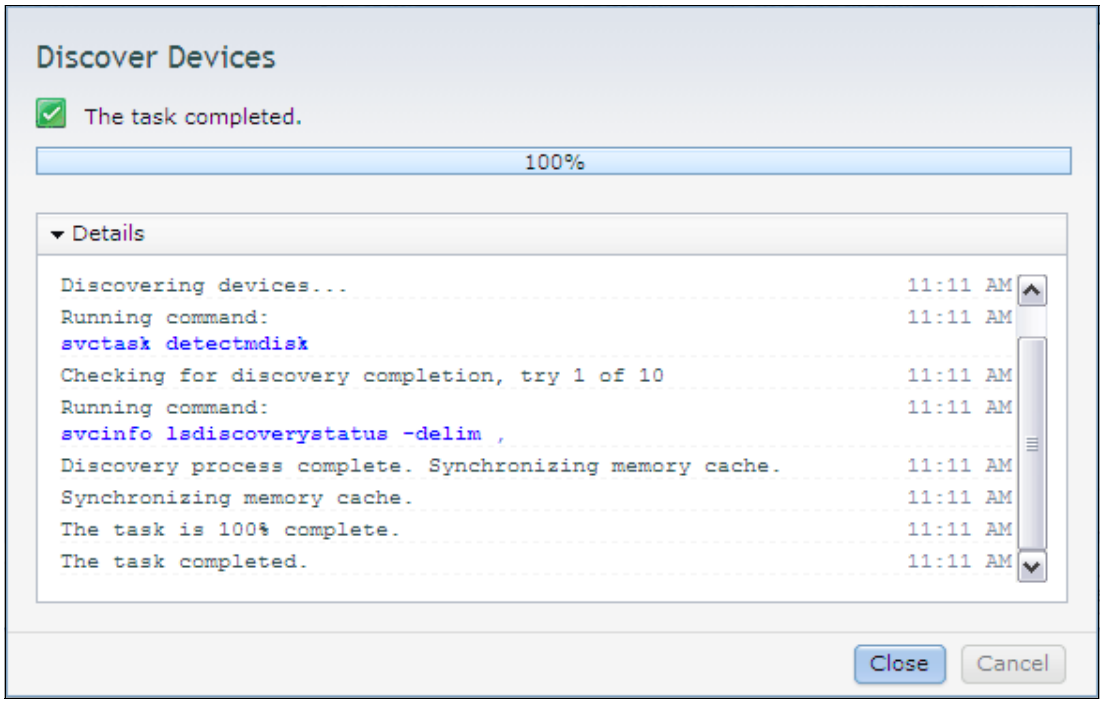

Figure 6-53 Discover devices panel

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.

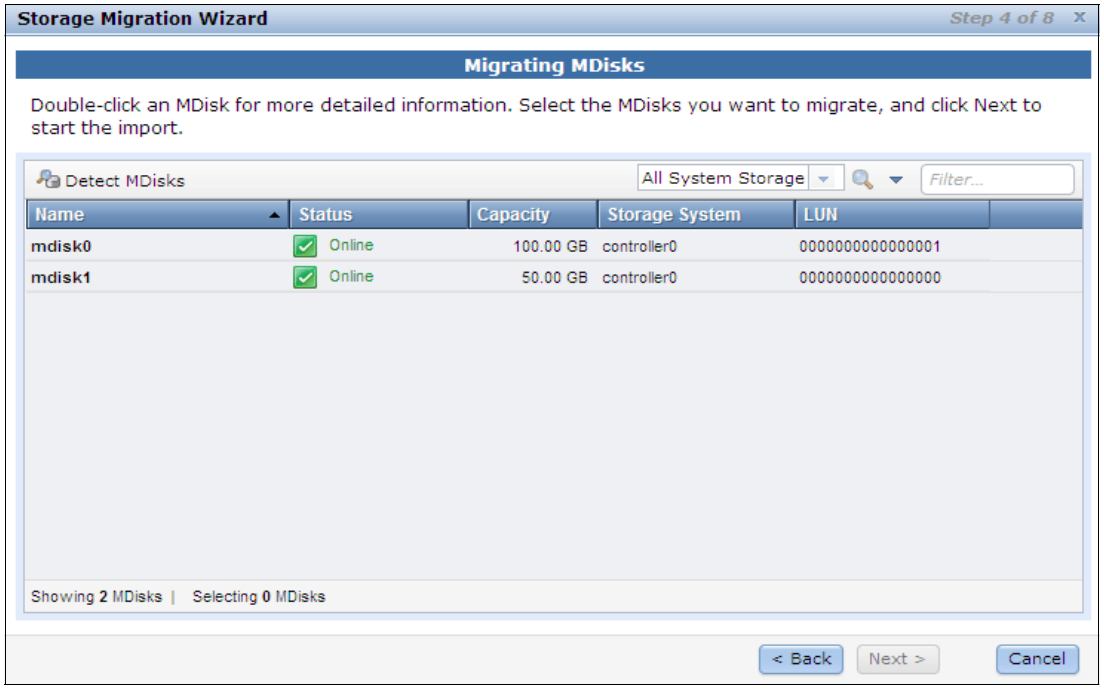

Figure 6-54 Step 4 of the storage migration wizard

Select the discovered MDisks and click **Next** to show the Import MDisks running task panel. Figure 6-55 shows step 4 of the storage migration wizard with the MDisks selected.

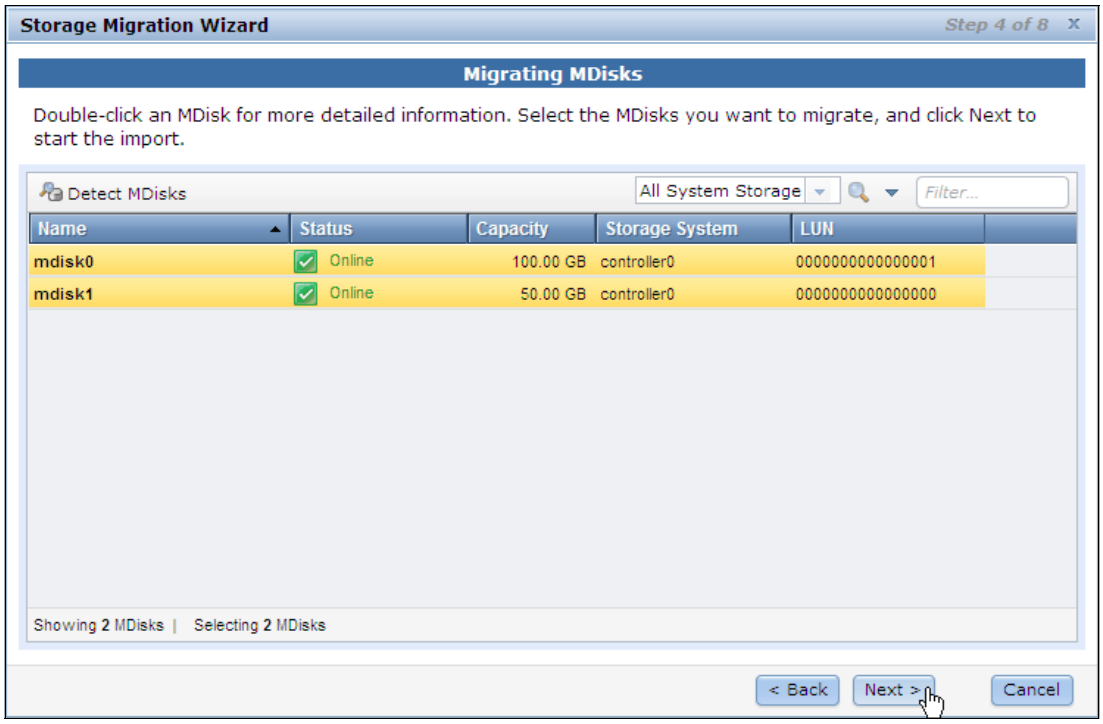

Figure 6-55 Step 4 of the storage migration wizard: Selecting MDisk to be migrated

After the Import MDisks running task is complete, select **Close** to open step 5 of the storage migration wizard. Figure 6-53 shows the Import MDisks panel.

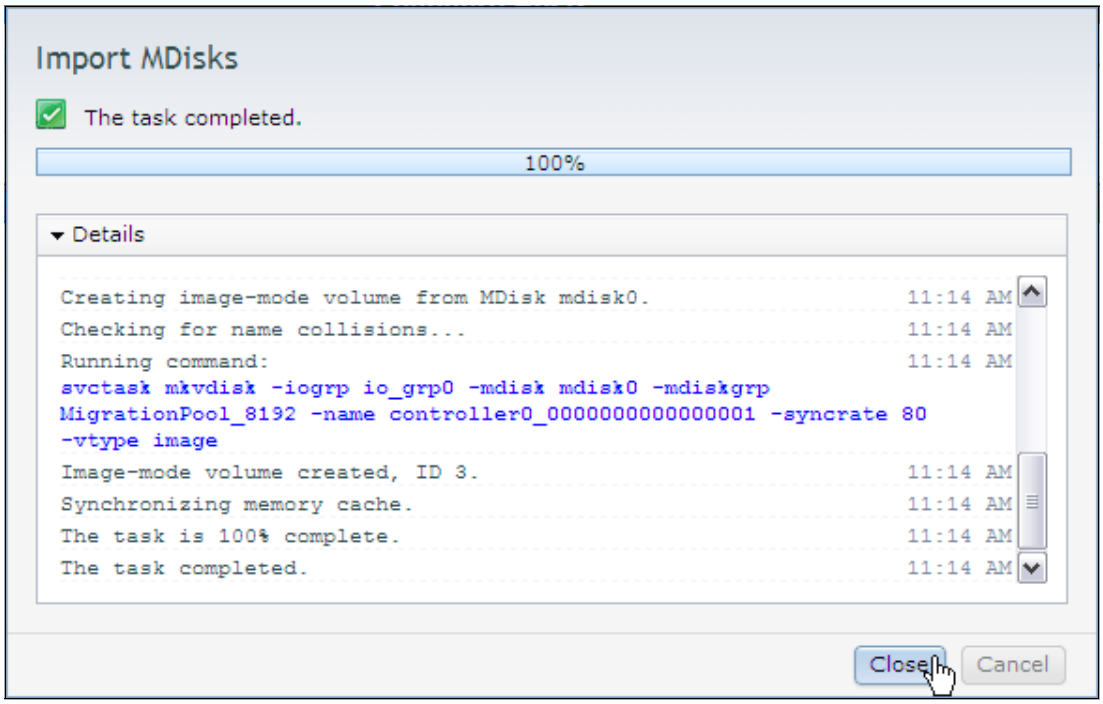

Figure 6-56 Import MDisks panel

Follow step 5 of the storage migration wizard step 5 entitled Configure Hosts, as shown in Figure 6-57. The Windows 2008 host is not yet defined in the Storwize V3500. Select **New Host** to open the create host panel.

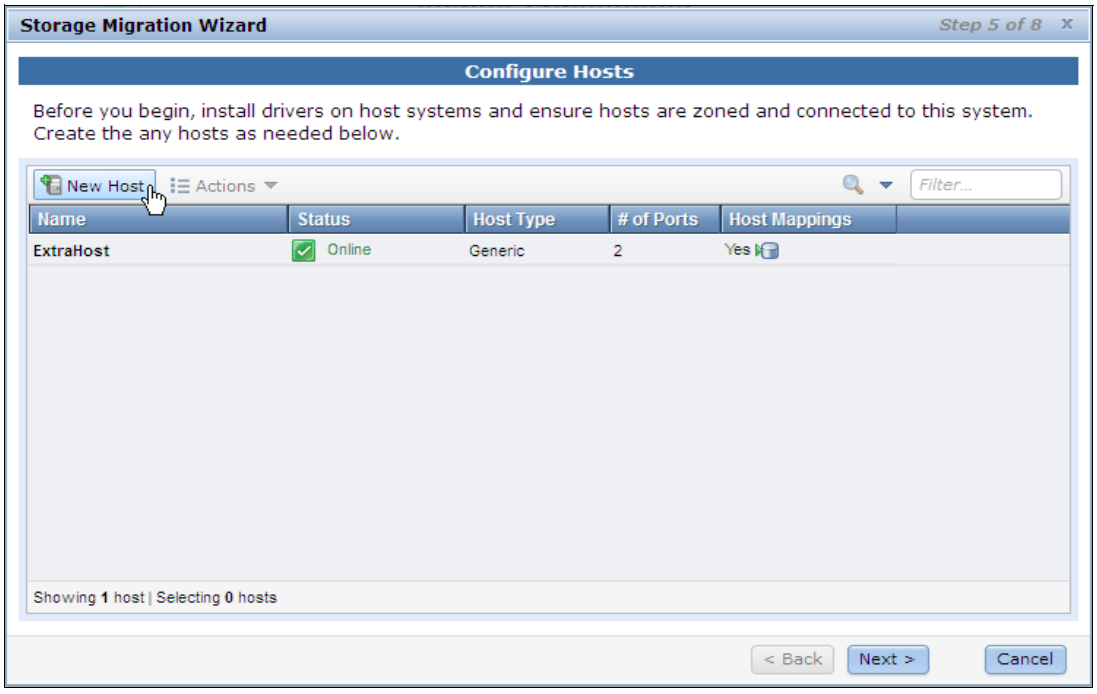

Figure 6-57 Step 5 of the storage migration wizard: Create host

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

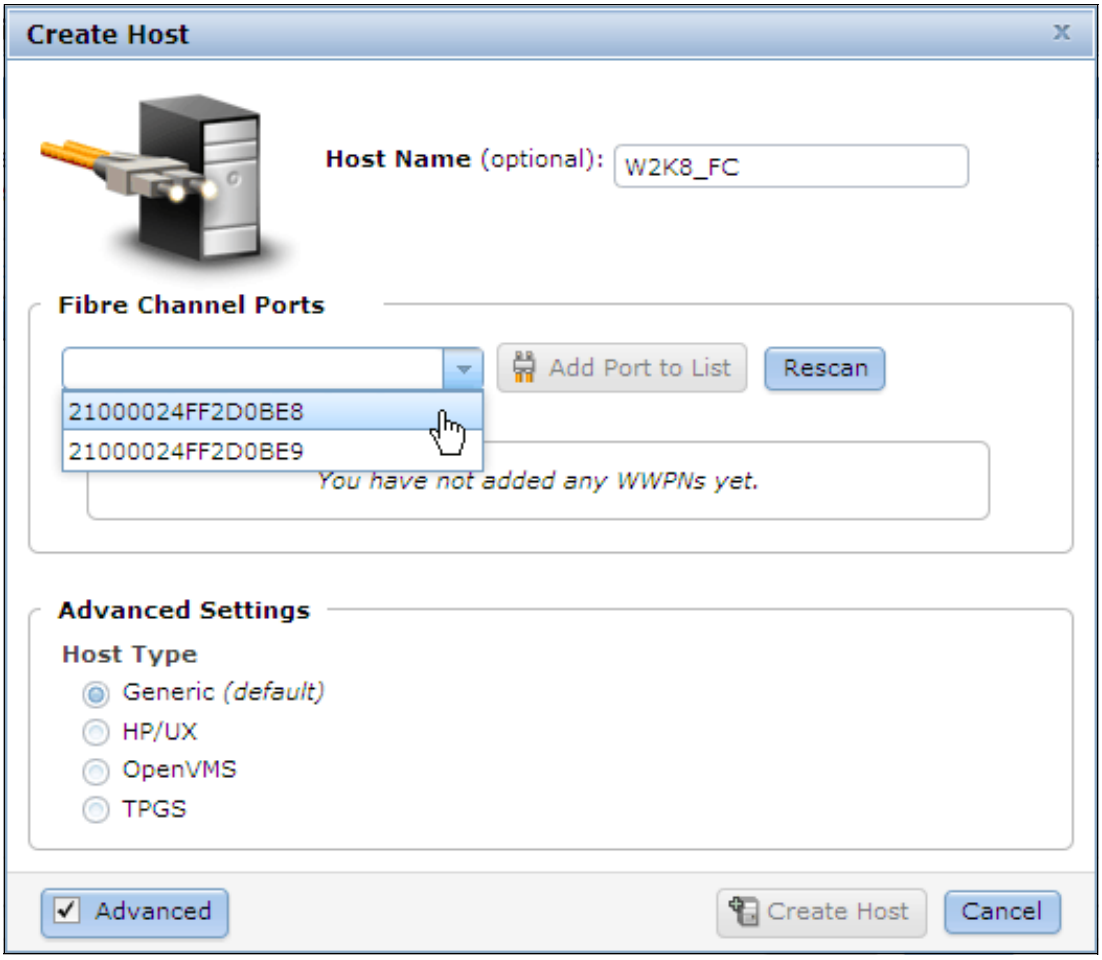

Figure 6-58 Storwize V3500 create host panel: Select WWPNs

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 both of the required port definitions listed.

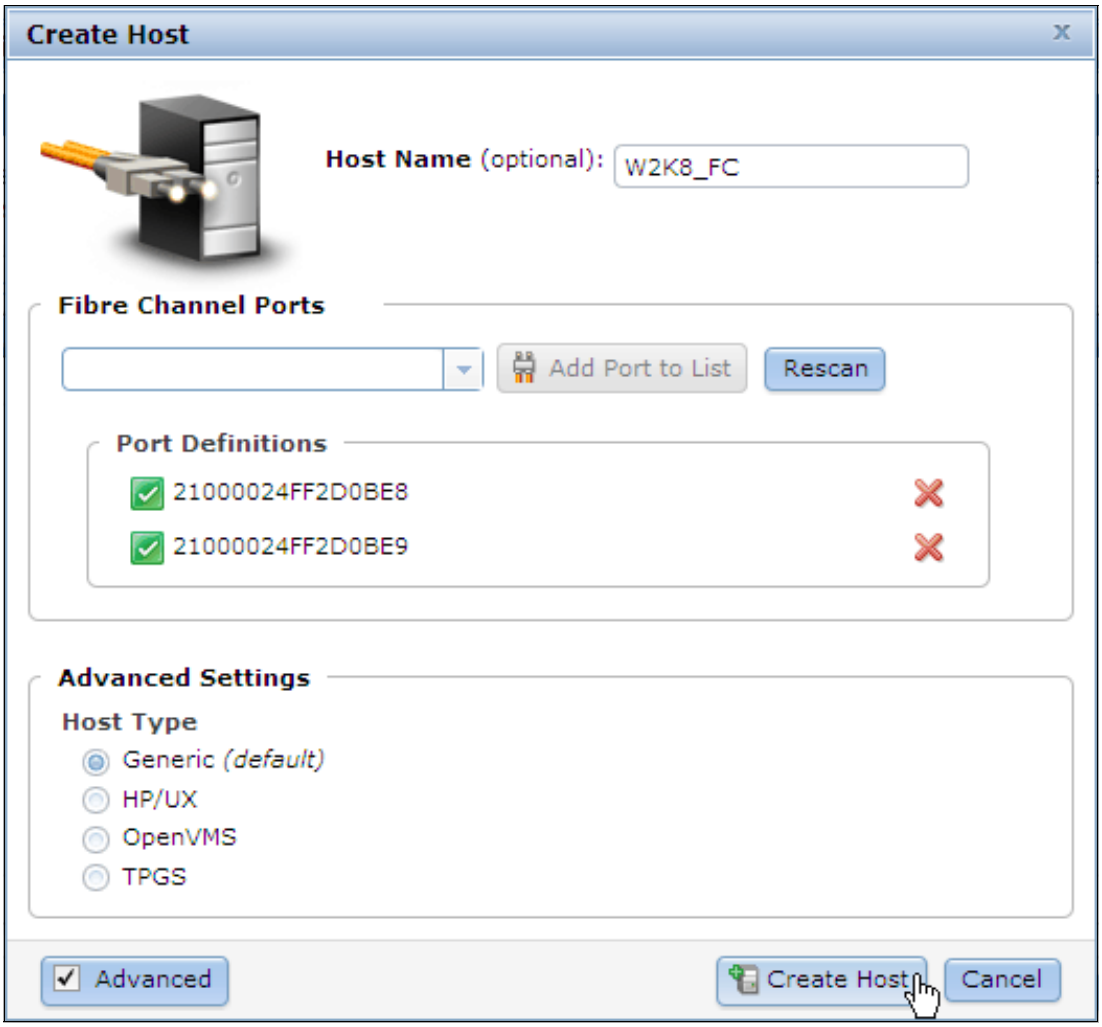

Figure 6-59 Create host panel: Both port definitions listed

After the Create Host running task is complete, select **Close** to open step 5 of the storage migration wizard again. Figure 6-60 the Create Host running task panel.

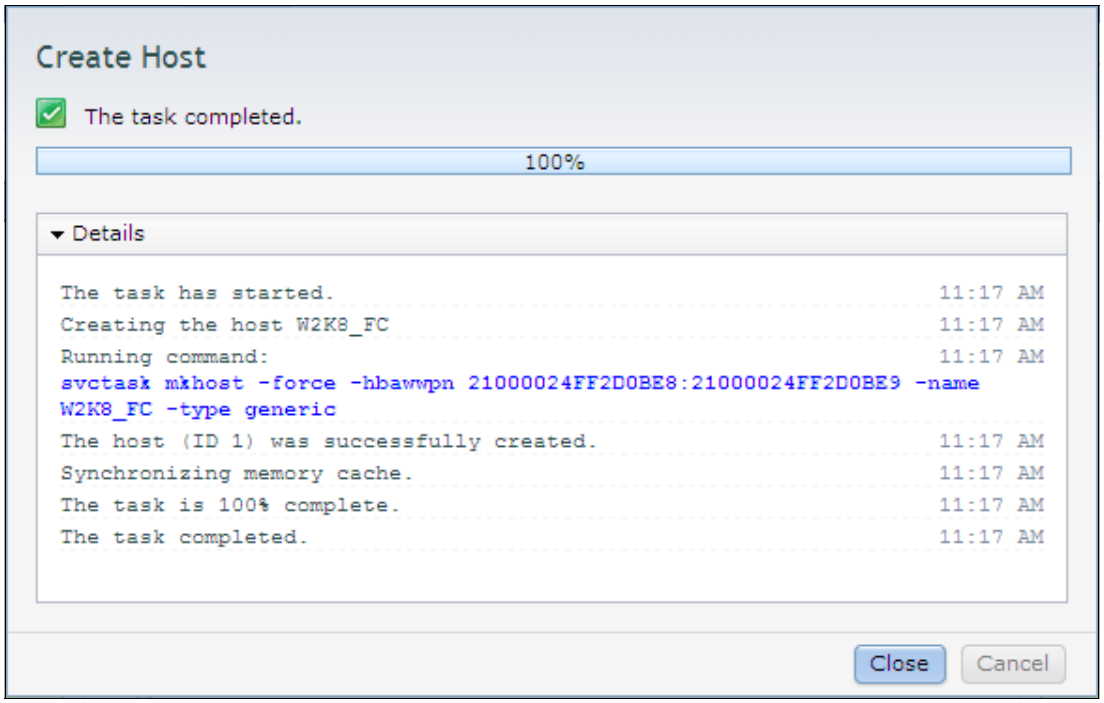

Figure 6-60 Create Host running task panel

From step 5 of the storage migration wizard, select the host that was just configured and click **Next** to open step 6 of the storage migration wizard. Figure 6-61 shows the host was selected in step 5.

Figure 6-61 Step 5 storage migration wizard: 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 storage migration wizard step 6 entitled Map Volumes to Hosts. 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.

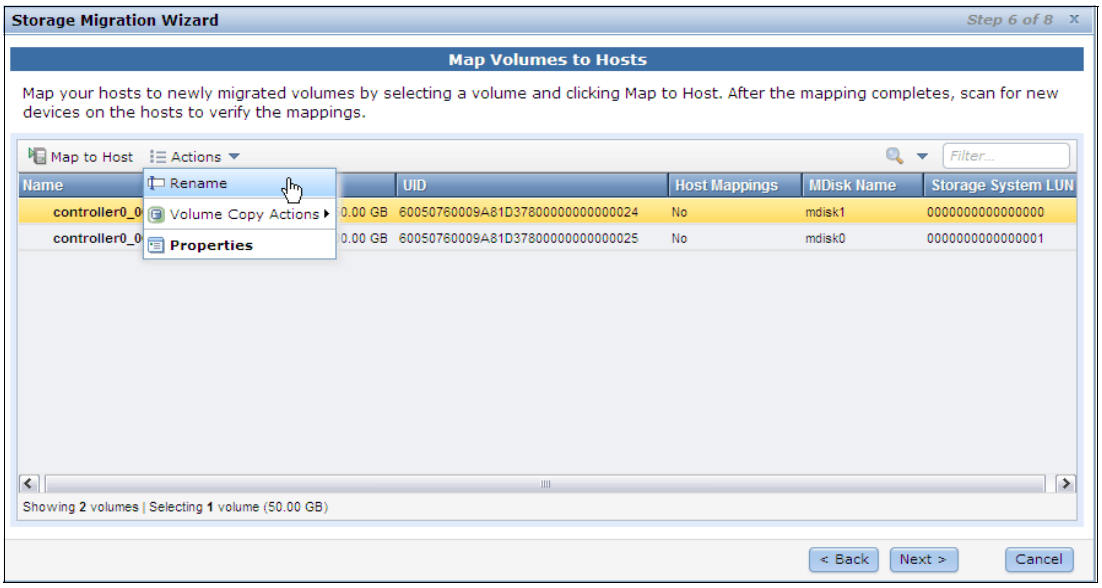

Figure 6-62 Step 6 of the storage migration wizard

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 open the rename running task, as shown in Figure 6-63. Rename both volumes.

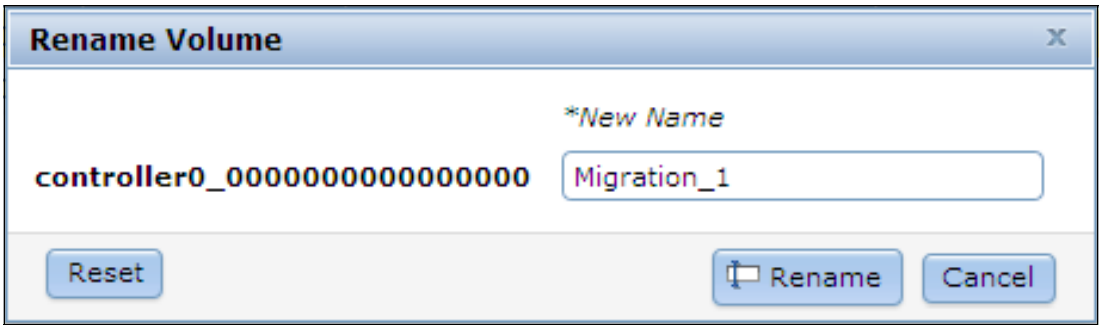

Figure 6-63 Rename Volume panel

After the final rename running task is complete, click **Close** to open step 6 of the storage migration wizard again. Figure 6-64 shows the Rename Volume running task panel.

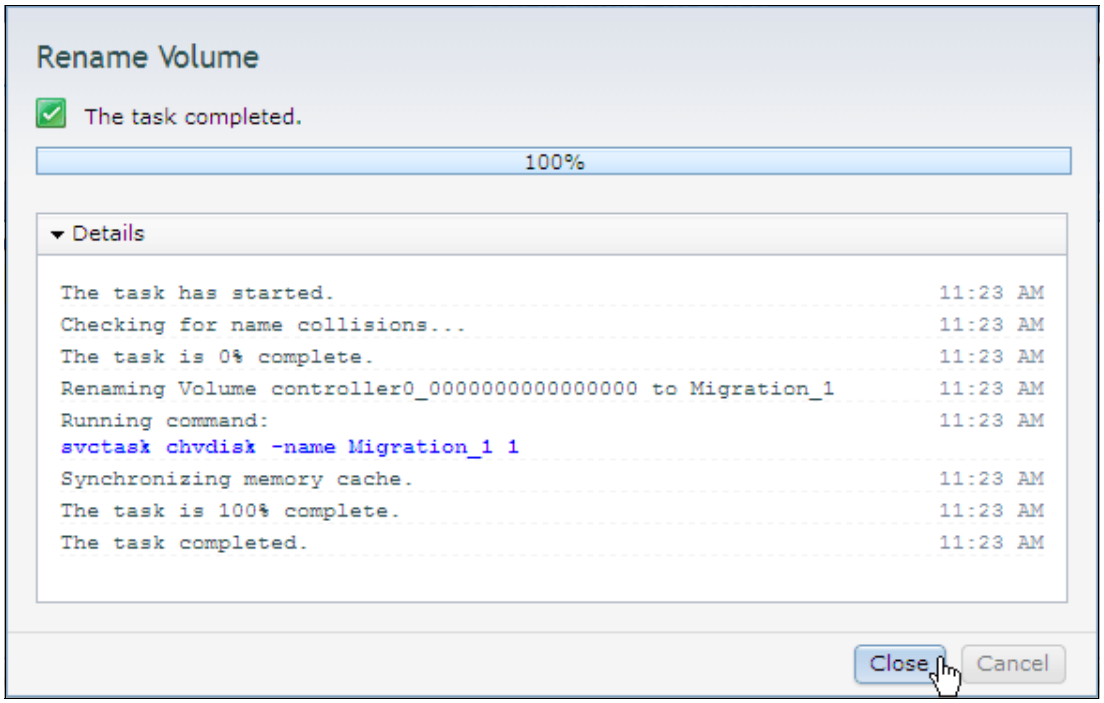

Figure 6-64 Rename Volume running task

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

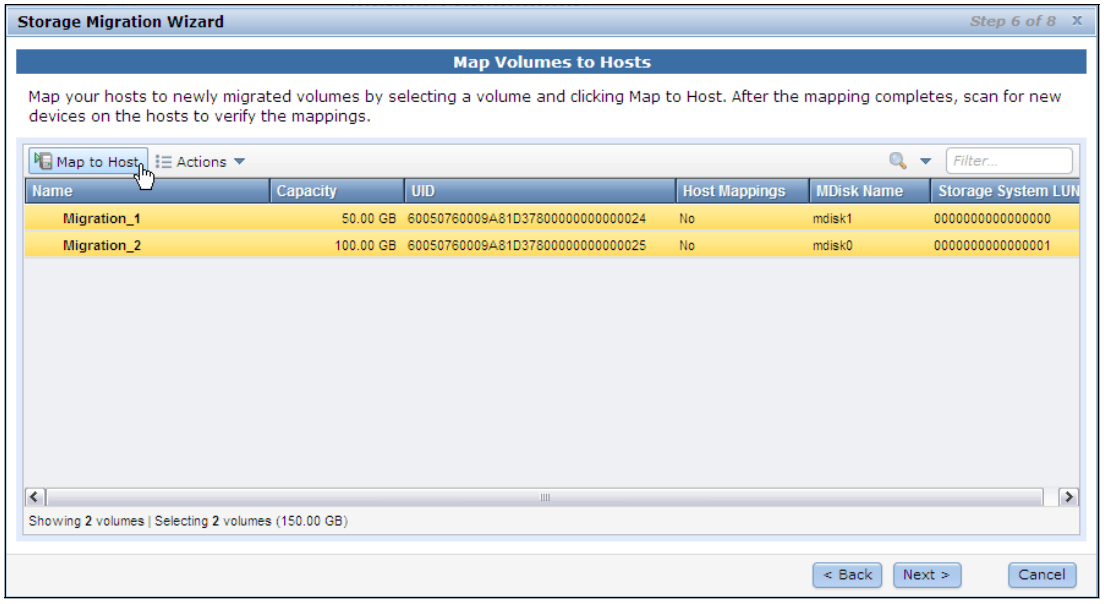

Figure 6-65 Step 6 storage migration wizard: Renamed MDisks highlighted for mapping

Select the host from the menu on the Modify Host Mappings panel. The rest of the Modify Host Mappings opens. Figure 6-66 shows the Modify Host Mappings panel.

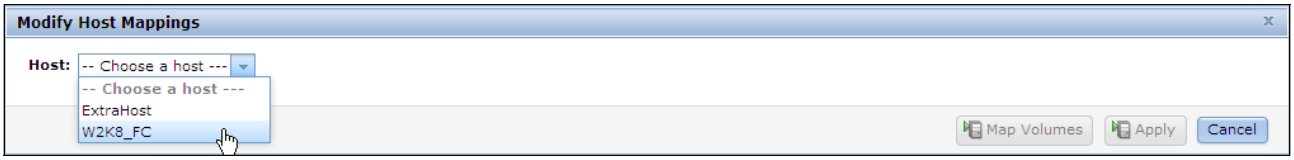

Figure 6-66 Modify Host Mappings panel

The MDisks that were 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. 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. Figure 6-67 shows the Modify Host Mappings panel.

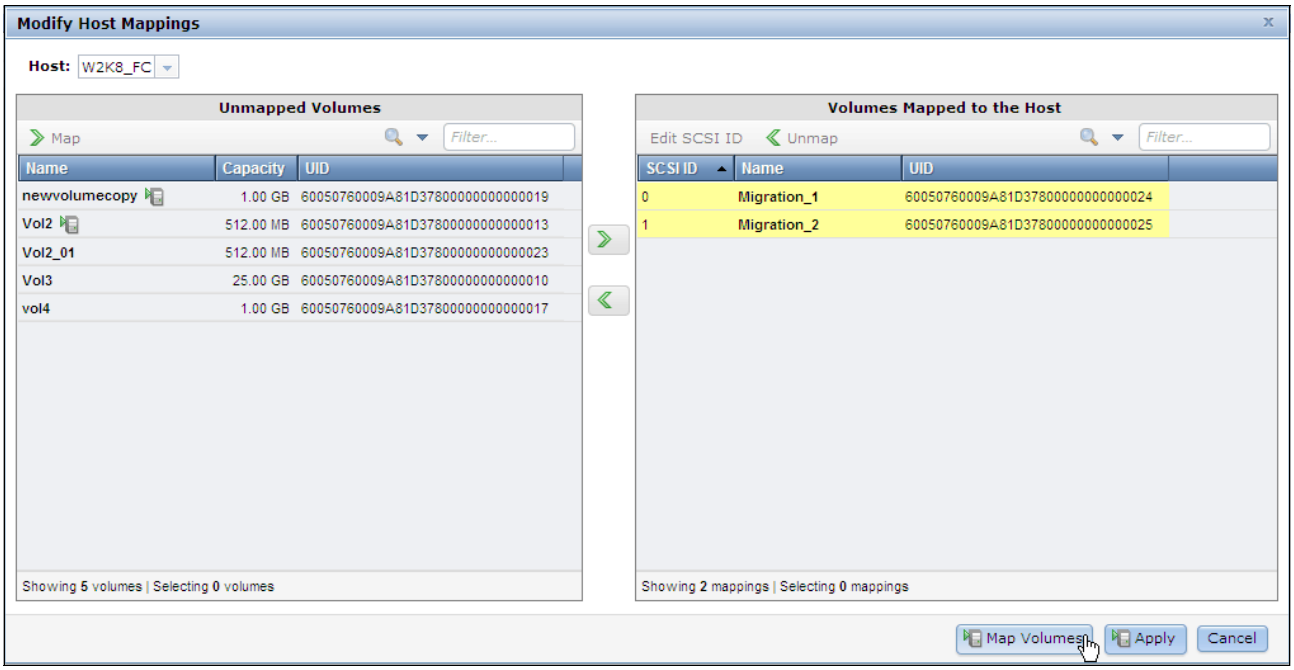

Figure 6-67 Modify Host Mappings panel

After the Modify Mappings running task is complete, select **Close** to show step 6 of the storage migration wizard again. Figure 6-68 shows the Modify Mappings running task.

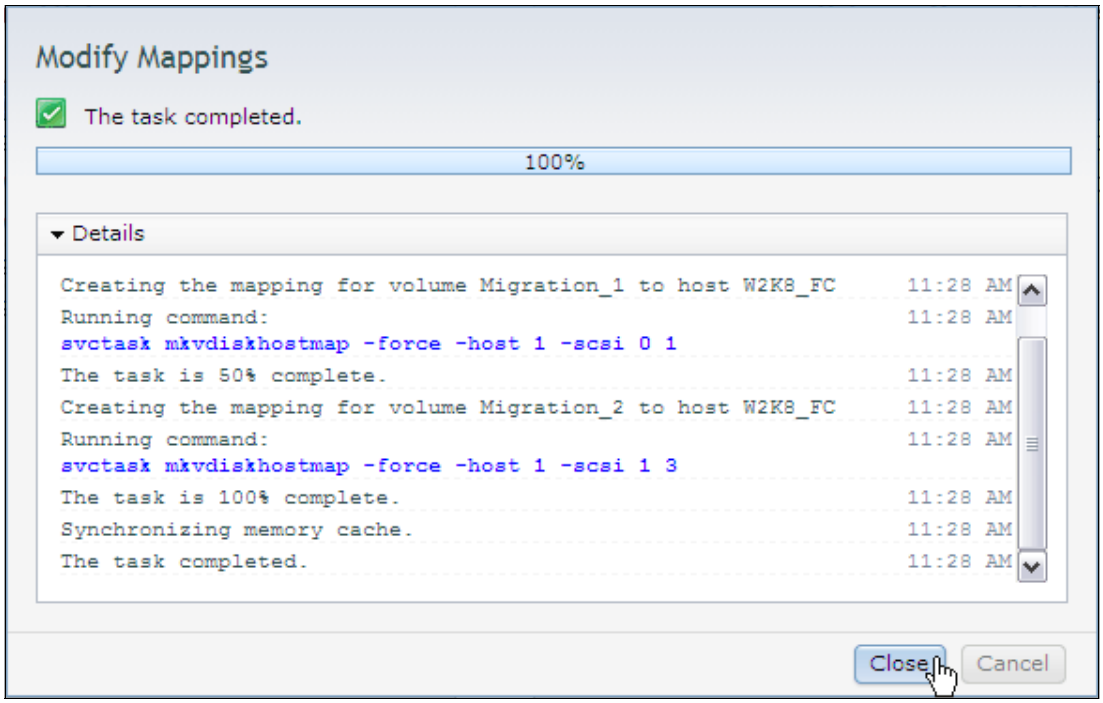

Figure 6-68 Modify Mappings running task

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

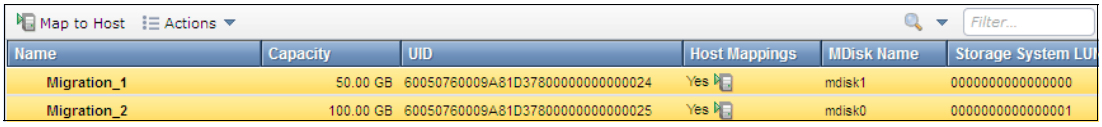

Figure 6-69 Step 6 of the migration wizard: MDisks mapped

Verify migrated disk device type is now 2145 on the host

Now the migrated volumes are mapped to the Storwize V3500 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**. Figure 6-70 shows how to display the properties of the disks.

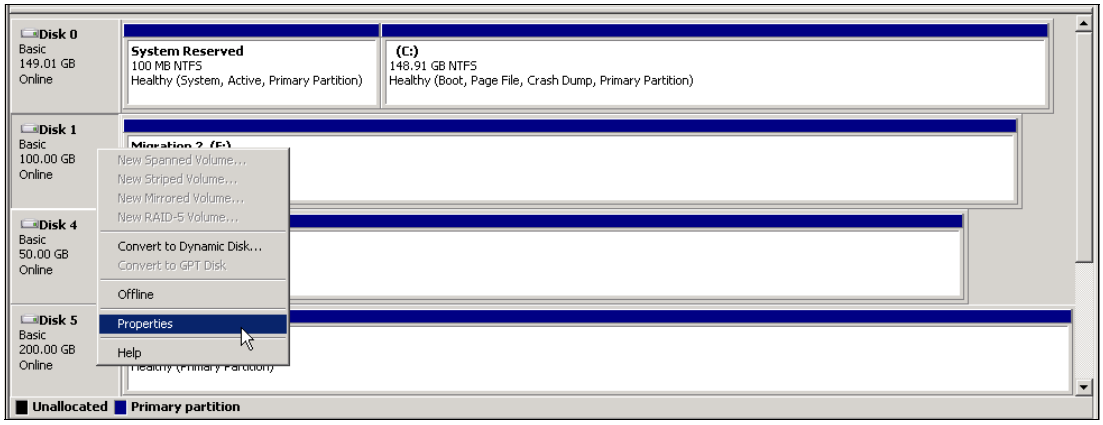

Figure 6-70 Display the disk properties from the Windows 2008 disk migration panel

After the disk properties panel is opened, the General tab shows the disk device type. Figure 6-71 shows the Windows 2008 disk properties General tab.

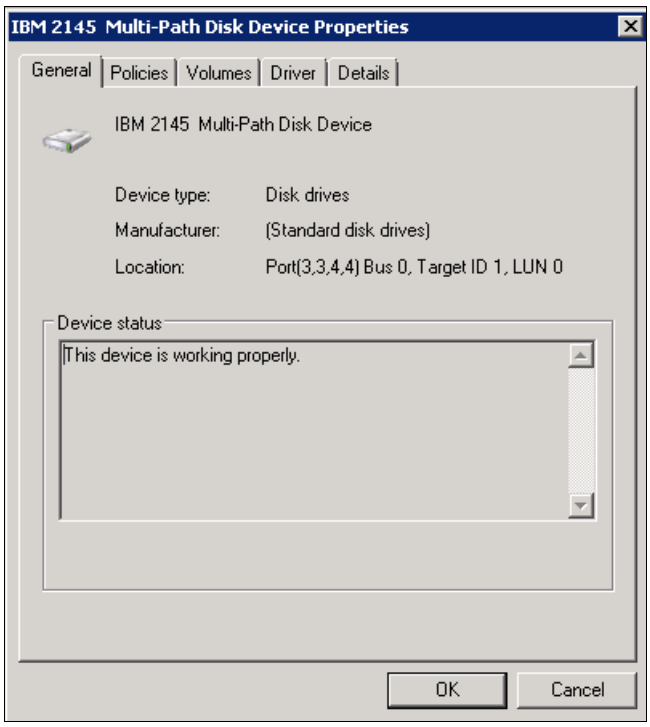

Figure 6-71 Windows 2008 properties General tab

The Storwize V3500 Subsystem Device Driver Device Specific Module (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** (as shown Figure 6-72) to open the SSDDSM CLI window.

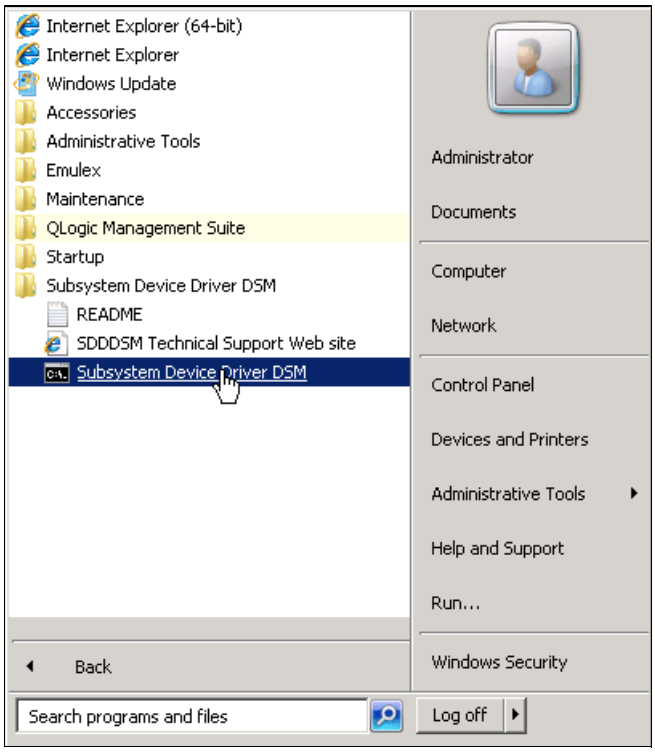

Figure 6-72 Windows 2008 R2 example, open SSDDSM command line

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 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. Figure 6-1 shows that 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.

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

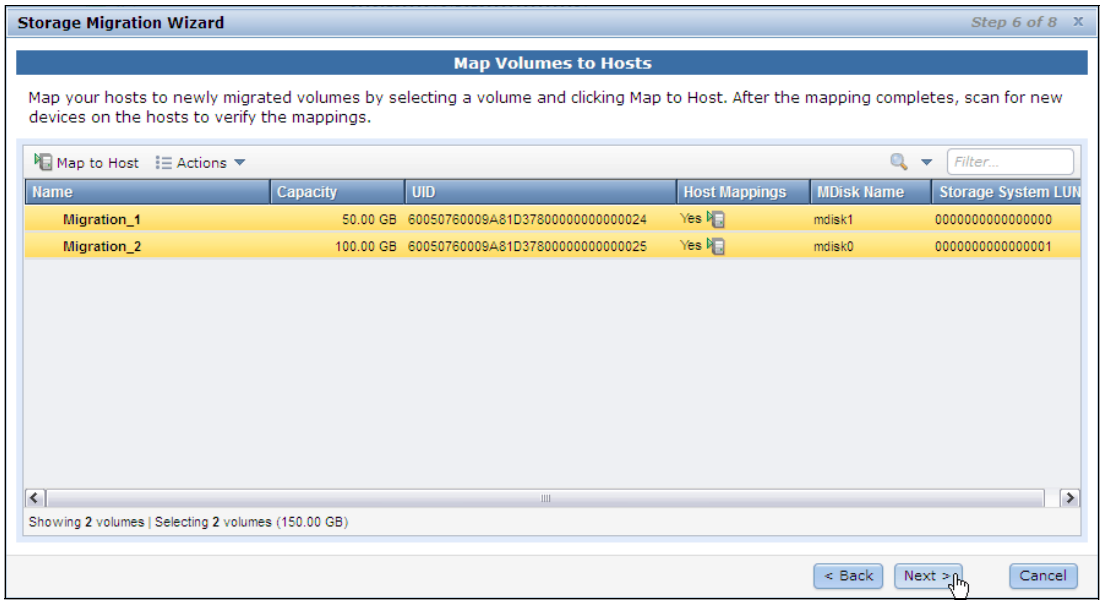

Figure 6-73 Step 6 of the storage migration wizard: mapped volumes and UIDs

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

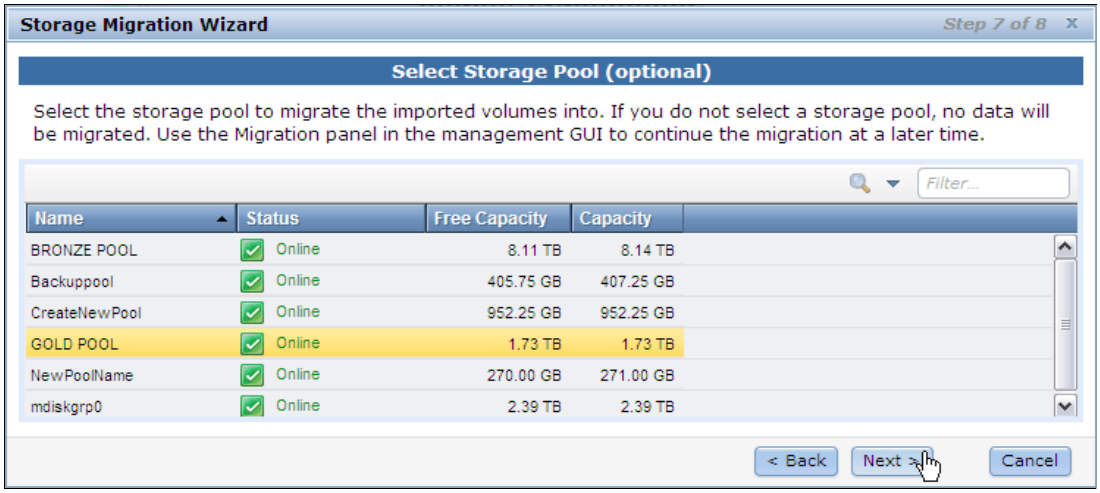

Figure 6-74 Step 7 of the storage migration wizard: Select storage pool

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.

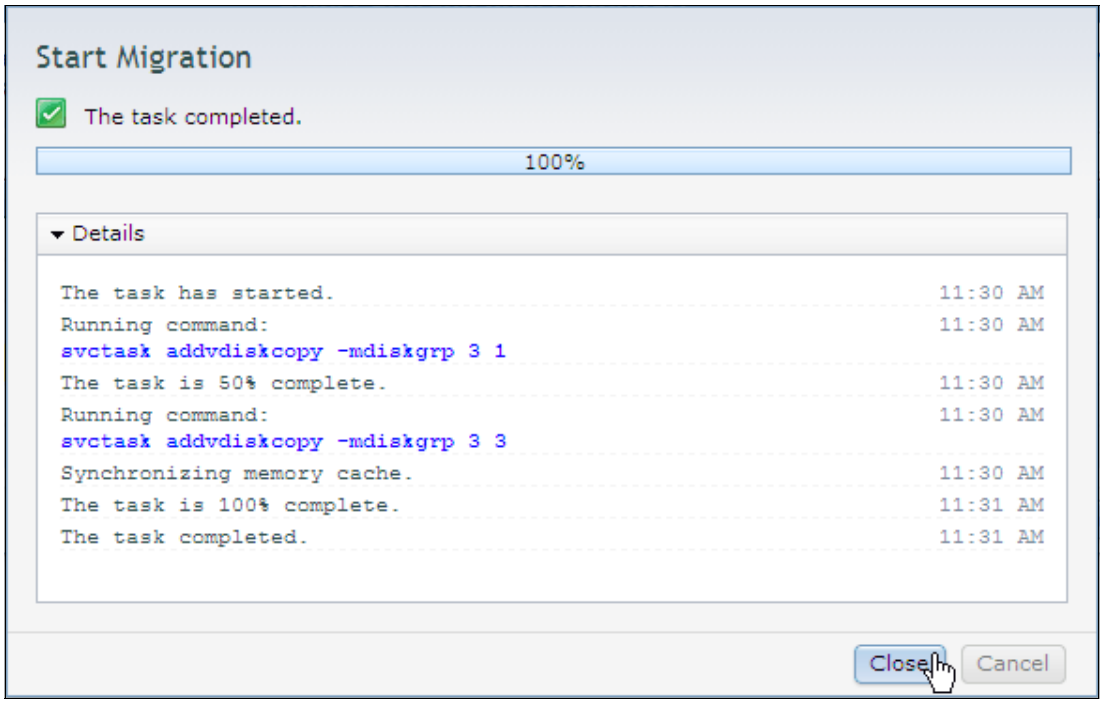

Figure 6-75 Start Migration completed task panel

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

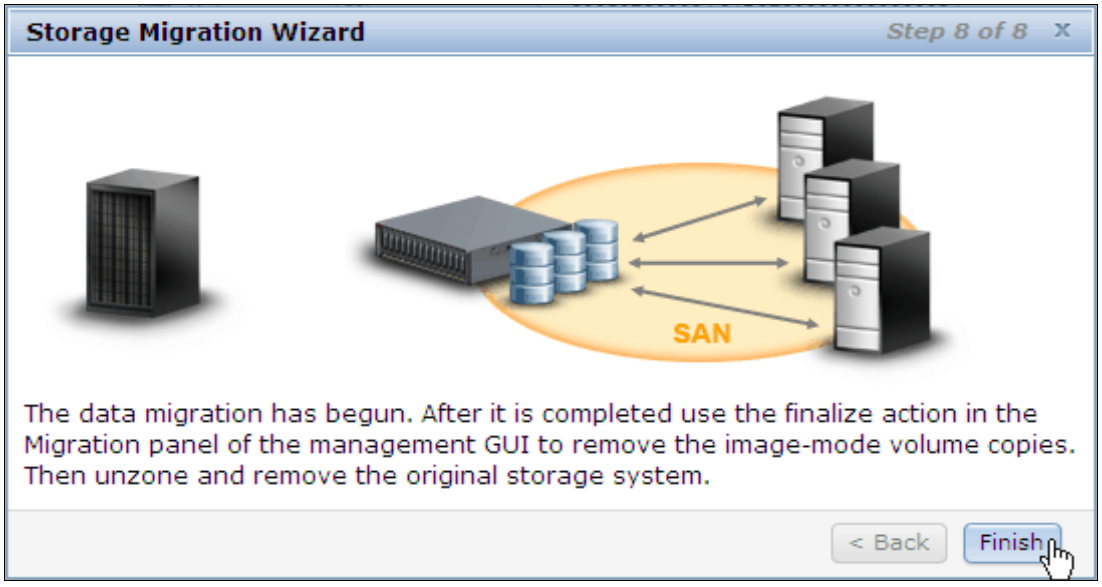

Figure 6-76 Step 8 of the storage migration wizard

The end of the storage migration wizard is not the end of the data migration. The data migration is still in progress. A percentage indication of the migration progress is shown in the System Migration panel. Figure 6-77 shows the System Migration panel migration progress percentage indicator.

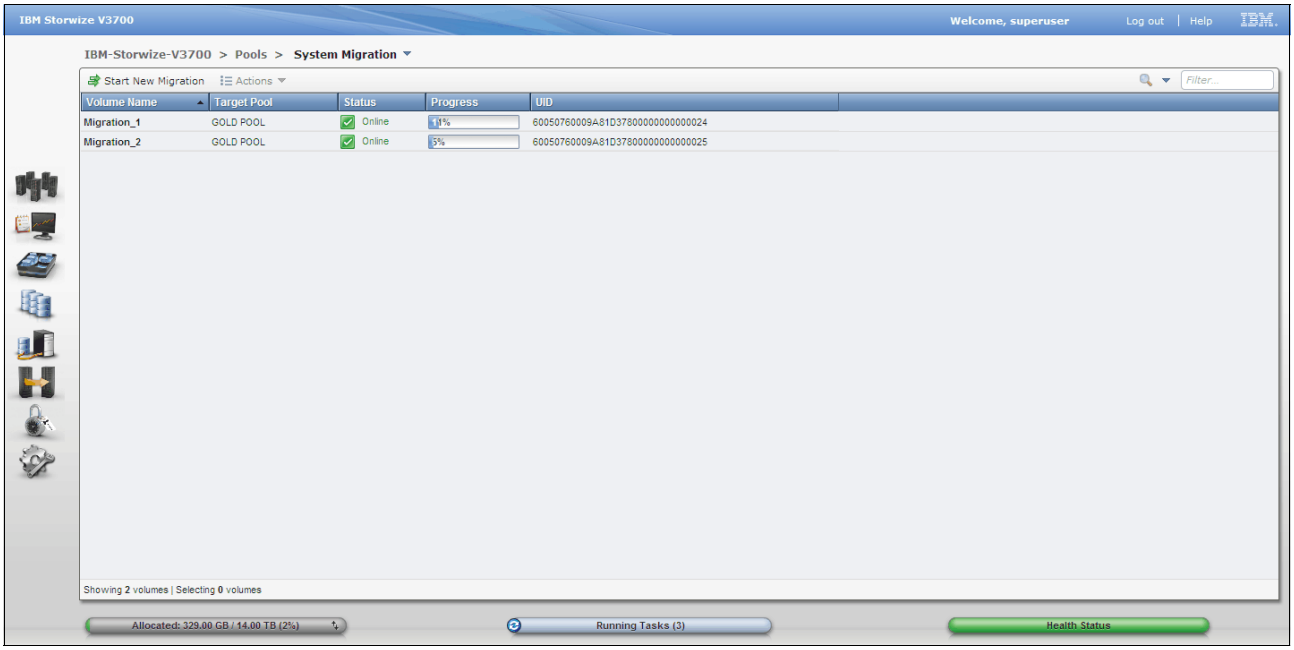

Figure 6-77 System migration panel: Migration progress indicators

Finalize the volume migrations. When the volume migrations complete, select the volume migration instance and right-click **Finalize** to open the Finalize Volume Migrations panel. Figure 6-78 shows the System Migration panel with the completed migrations.

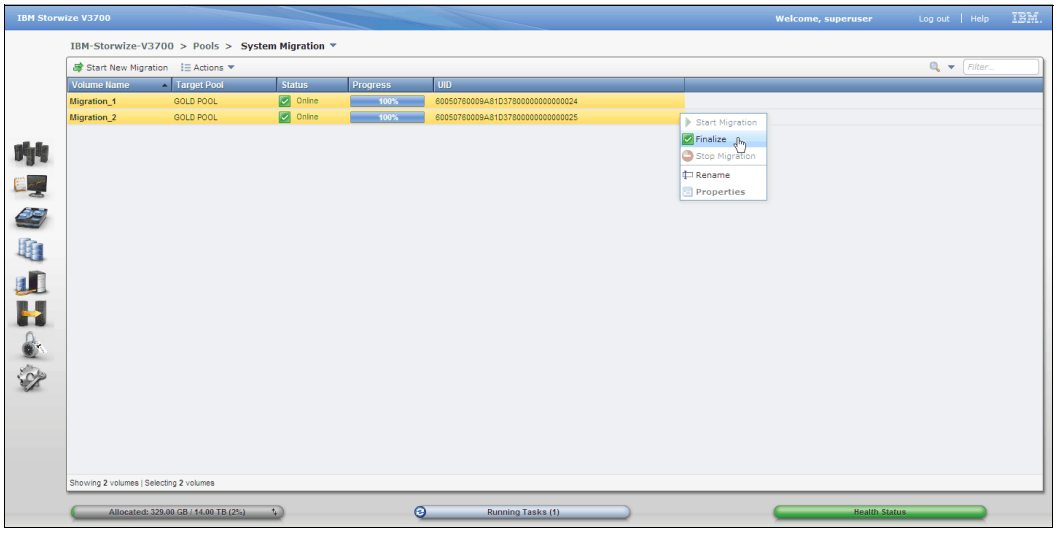

Figure 6-78 System migration panel: finalization

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

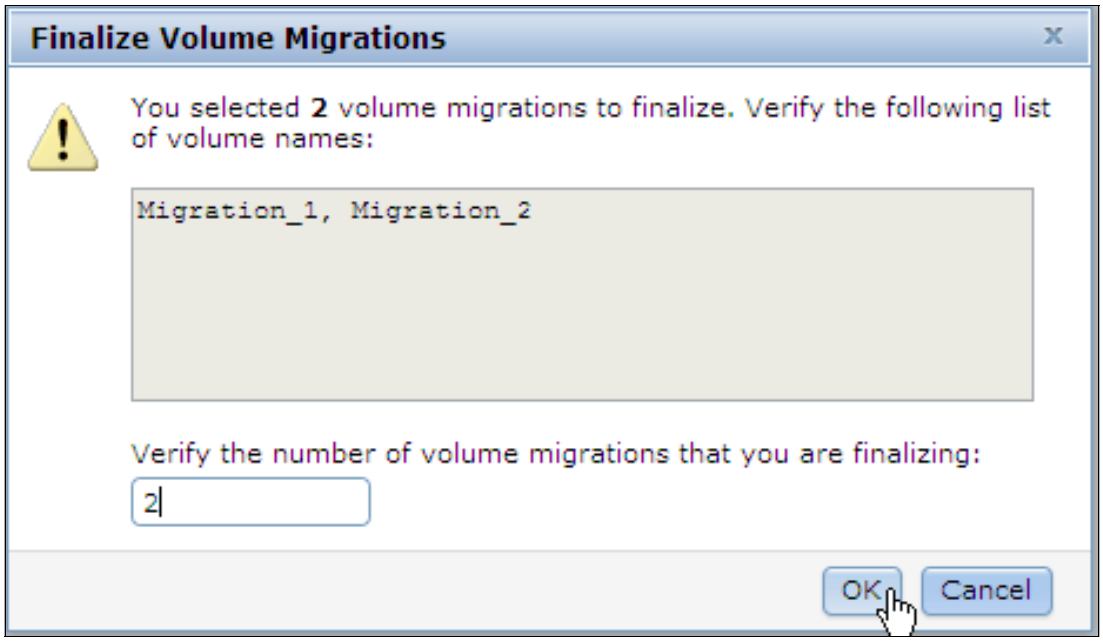

Figure 6-79 Finalize Volume Migrations panel

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 then unmanaged. When the finalization completes, the data migration to the Storwize V3500 is done. Remove the DS3400-to-V3500 zoning and retire the older storage system.

7

Storage pools

This chapter describes how IBM Storwize V3500 manages physical storage resources. All storage resources under IBM Storwize V3500 control are managed by using *storage pools*. Storage pools make it easy to dynamically allocate resources, maximize productivity, and reduce costs.

Storage pools are 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 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, whether to optimize for capacity, and the number of drives to provision.

Tuning the configuration also is described in this chapter.

This chapter includes the following topics:

- Working with internal drives
- Configuring internal storage
- Working with MDisks
- Working with storage pools

7.1 Working with internal drives

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

The IBM Storwize V3500 storage system provides an individual 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.

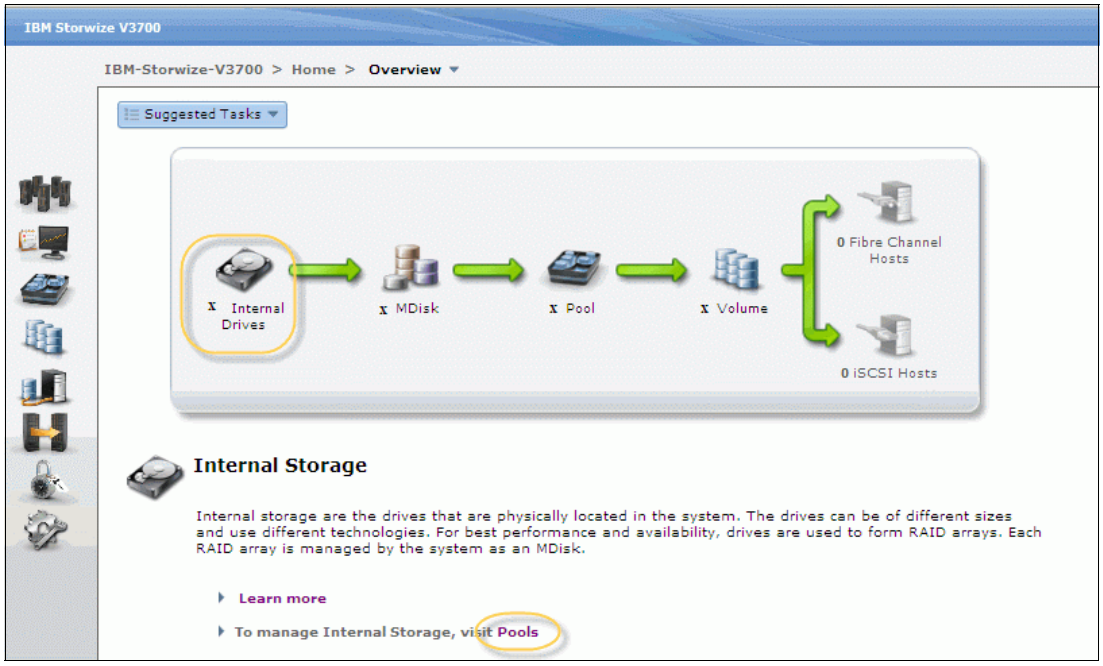

Figure 7-1 Internal Storage via Home Overview

An alternative way to access the Internal Storage window (as shown in Figure 7-2) is by clicking the **Pools** icon on the left side of the window.

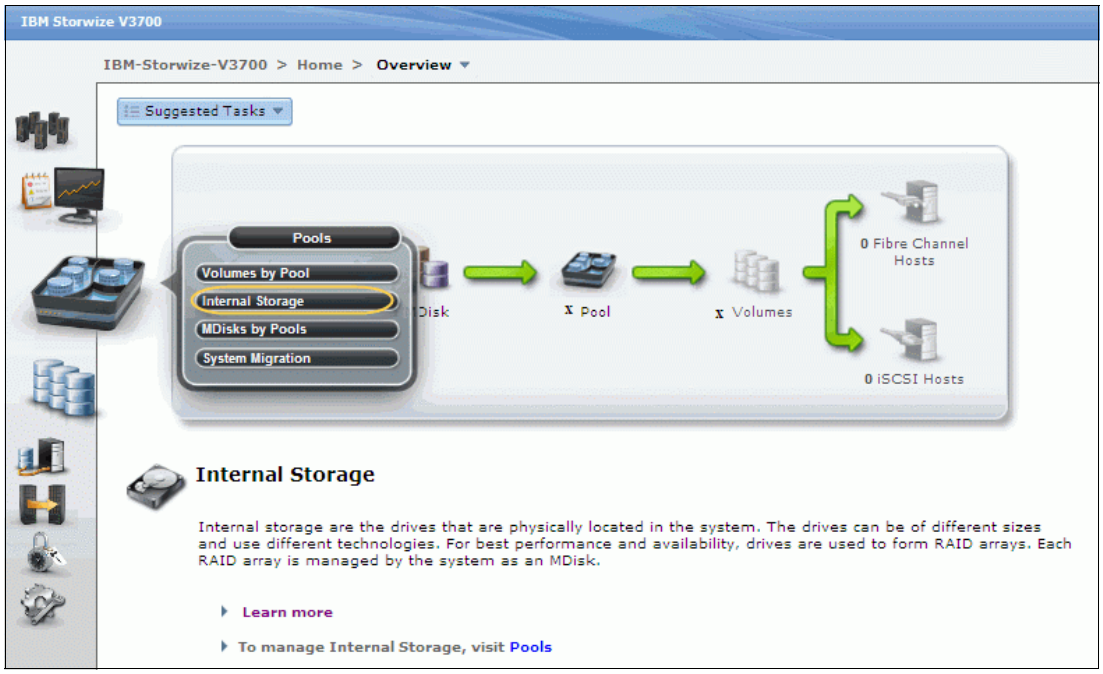

Figure 7-2 Internal Storage Details via Pools icon

7.1.1 Internal storage window

The Internal Storage window (as shown in Figure 7-3) gives an overview of the internal drives that are installed in the IBM Storwize V3500 storage system. Selecting **All Internal** in the Drive Class Filter shows all the drives that are installed in the managed system. Alternatively, you can filter the drives by their type or class; for example, you can choose to show only SAS or SATA drives.

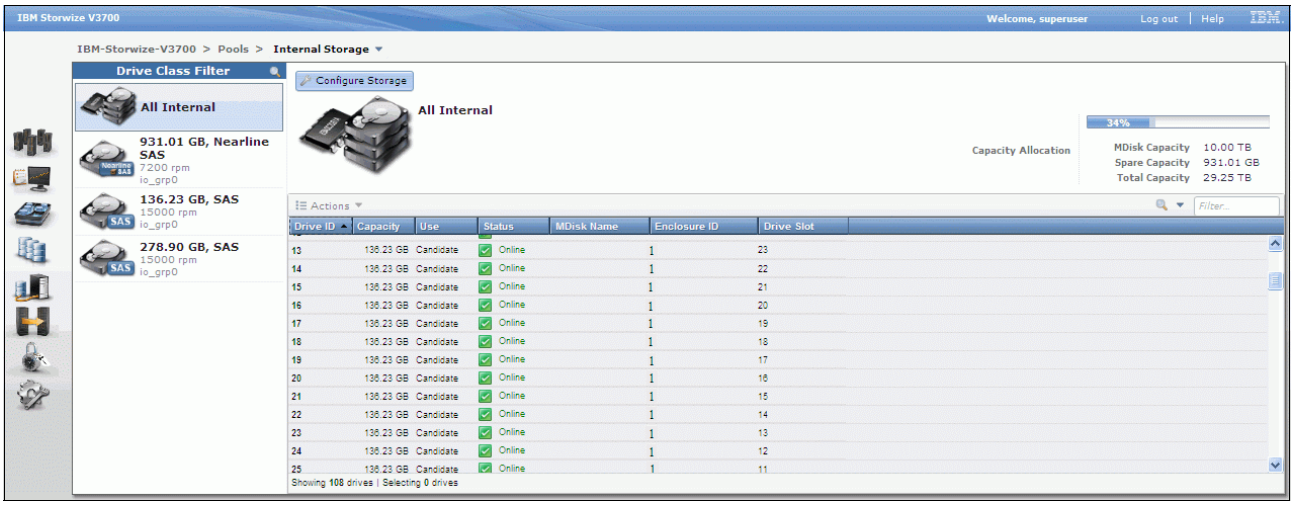

Figure 7-3 Internal storage window

On the right side of the Internal Storage window, the internal disk drives of the selected type are 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 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

Default sort order is by enclosure ID, but can be changed to any other column by left-clicking the column header. To toggle between ascending and descending sort order, the column header can by left-clicked again.

More details can be shown (for example, the drive's RPM speed or its MDisk member ID) by right-clicking the blue header bar of the table, which opens the selection panel, as in Figure 7-4.

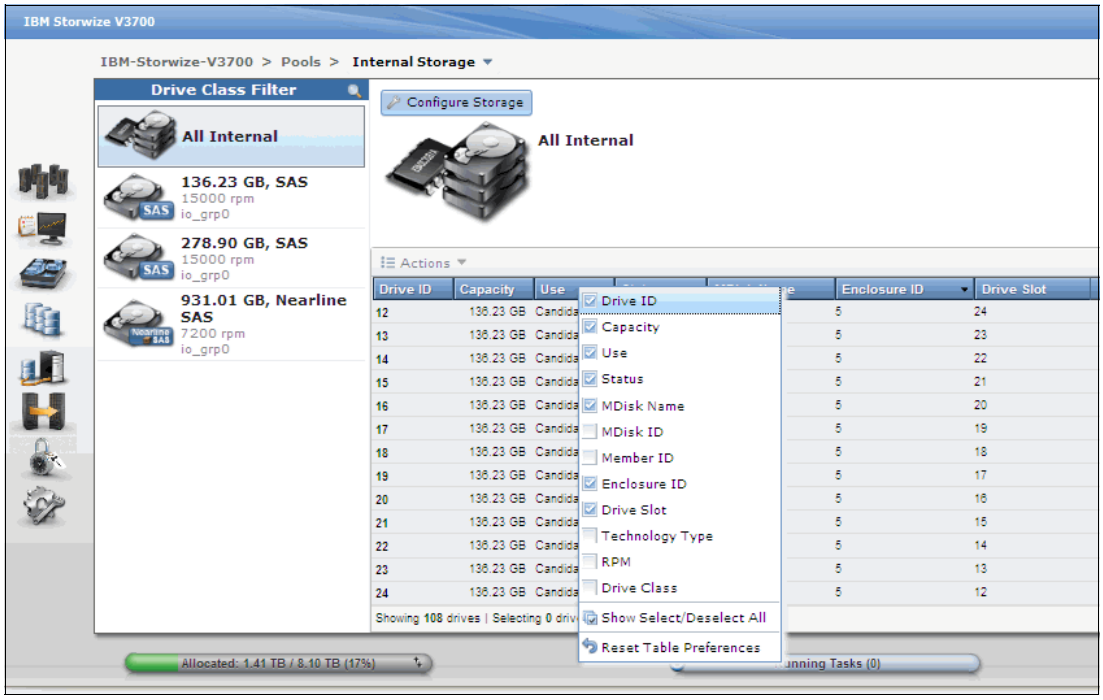

Figure 7-4 Internal storage window details selection

In addition, you can find the internal storage capacity allocation indicator at the upper right. The Total Capacity shows the overall capacity of the internal storage that is installed in this IBM Storwize V3500 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 was allocated.

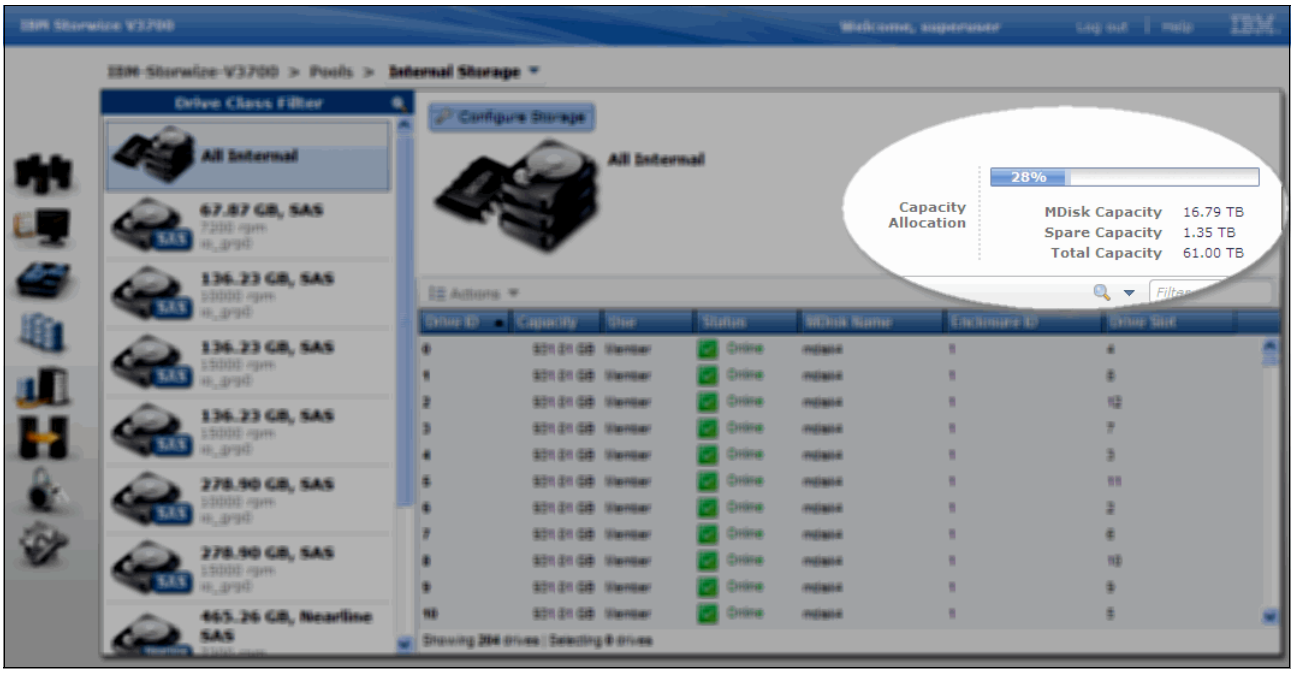

Figure 7-5 Internal storage allocation indicator

Example 7-1 shows how to use the **lsdrive** command display drive configuration and vital product data (VPD).

Example 7-1 The use of the lsdrive command

```
lsdrive
lsdrive drive ID
```

7.1.2 Actions on internal drives

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

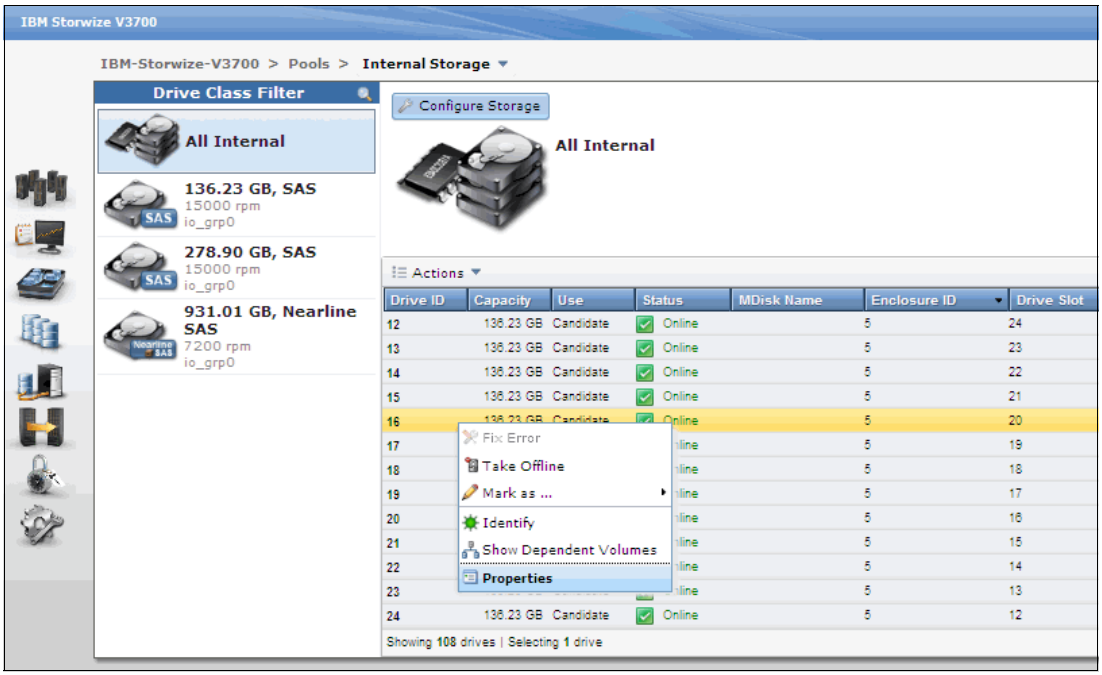

Figure 7-6 Internal drive actions menu

Fix Error

The Fix Error action starts the Directed Maintenance Procedure (DMP) for a defective drive. For more information, see Chapter 9, “RAS, monitoring, and troubleshooting” on page 439.

Take Offline

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

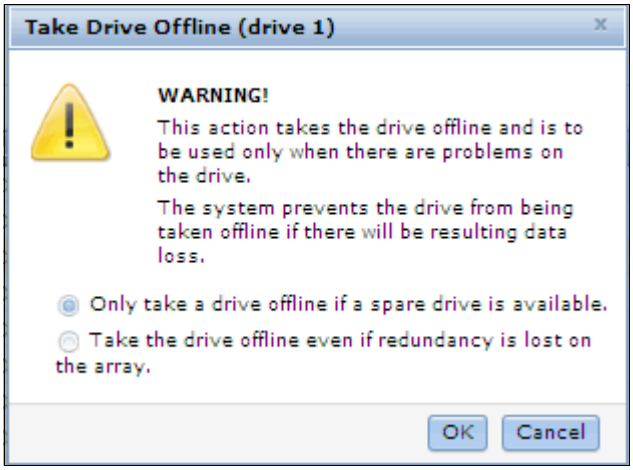

Figure 7-7 Take drive offline warning

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.

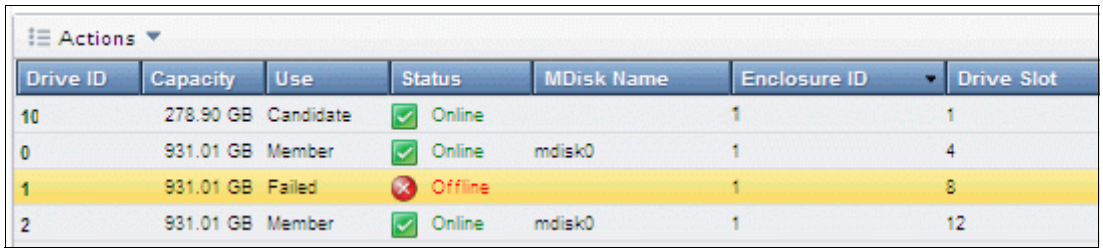

Figure 7-8 Internal drive taken offline

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.

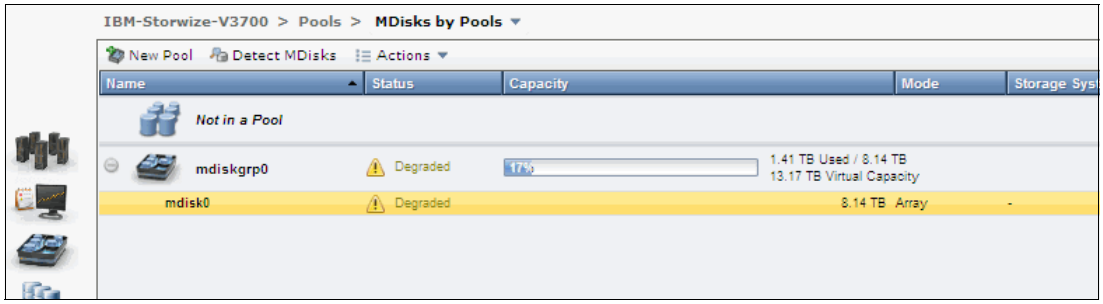

Figure 7-9 Internal drive failed with MDisk degraded

The IBM Storwize V3500 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.

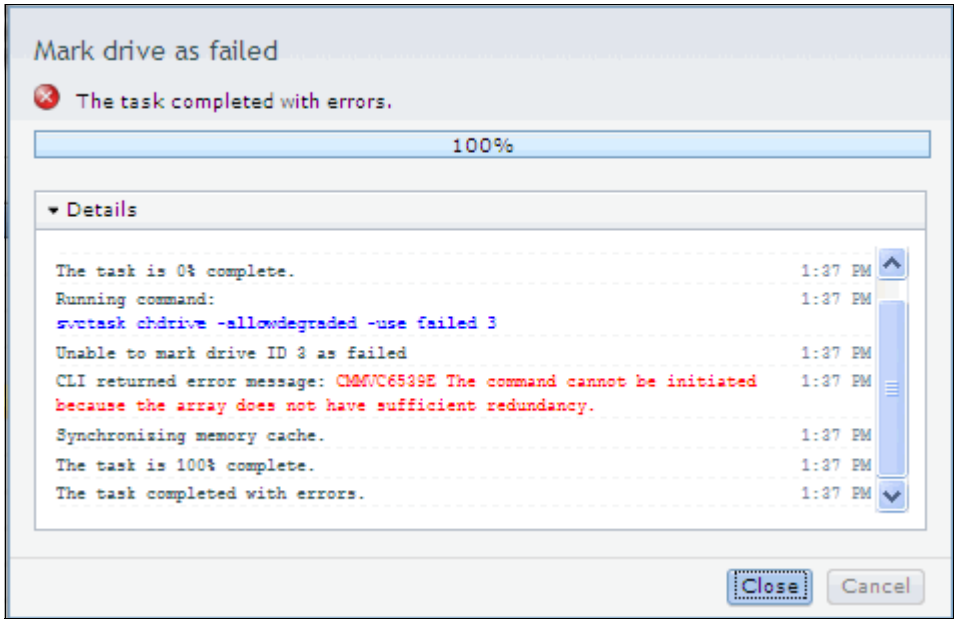

Figure 7-10 Internal drive offline not allowed

Example 7-2 shows how to use the **chdrive** command to set the drive to failed.

Example 7-2 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 V3500 storage system can be assigned to several usage roles, which can be unused, candidate, or spare, as shown in Figure 7-11. The roles have the following meanings:

- Unused: The drive is not in use and is not 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.

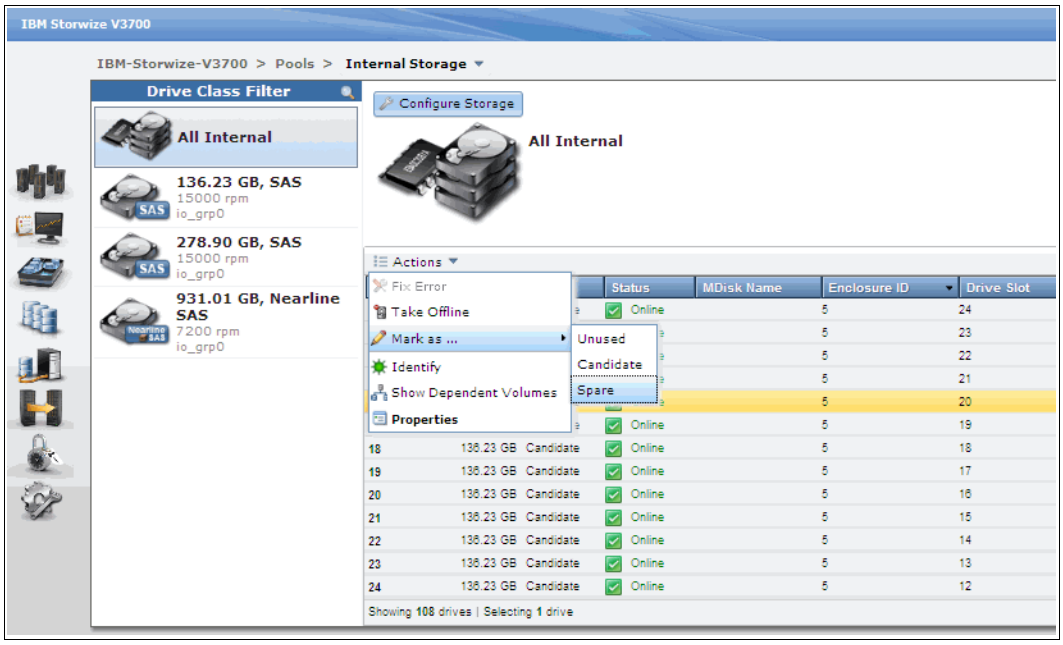

Figure 7-11 Internal drive mark as option

The new role that can be assigned depends on the current drive usage role. Figure 7-12 shows these dependencies.

From \ To	Unused	Candidate	Failed	Member	Spare
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 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 appears when the LED in on.

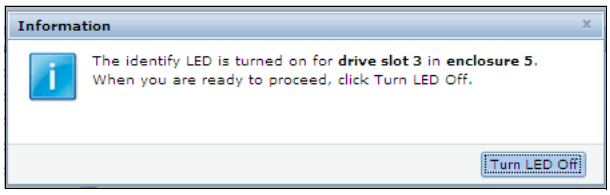

Figure 7-13 Internal drive identification

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

Example 7-3 shows how to use the **chenclosureslot** command to turn on and off the drive LED.

Example 7-3 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 in a case in which a volume is inaccessible because the drive in unavailable. This condition is true for any RAID 0 or if the associated MDisk is degraded already.

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

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

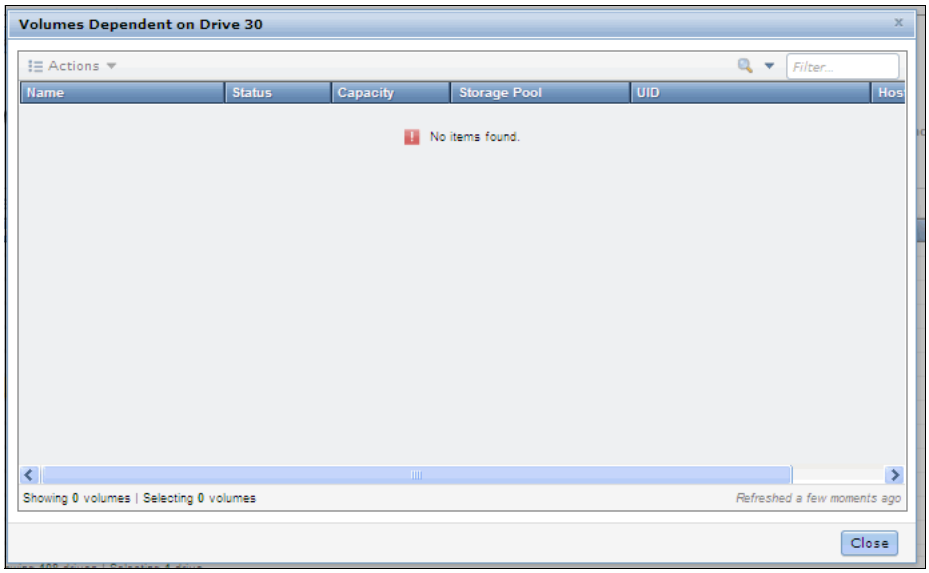

Figure 7-14 Internal drive no dependent volume

Figure 7-15 shows the list of dependent volumes for a drive.

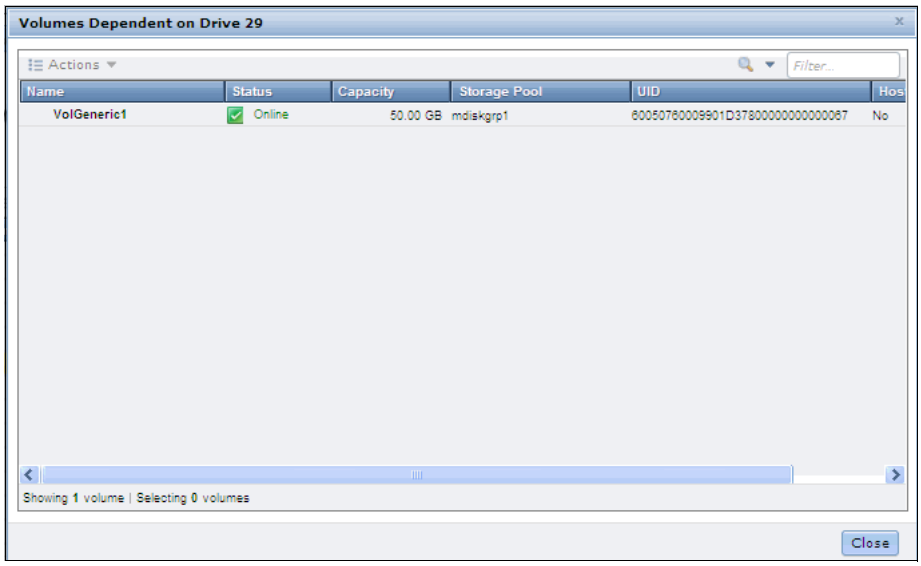

Figure 7-15 Internal drive with dependent volume

Example 7-4 shows how to view dependent volumes for a specific drive by using the CLI.

Example 7-4 Command to view dependent vdisks for a specific drive

```
lsdependentvdisks -drive driveID
```

Properties

Clicking **Properties** (as shown in Figure 7-16) in the Actions drop-down menu or double-clicking the drive provides the vital product data (VPD) and the configuration information.

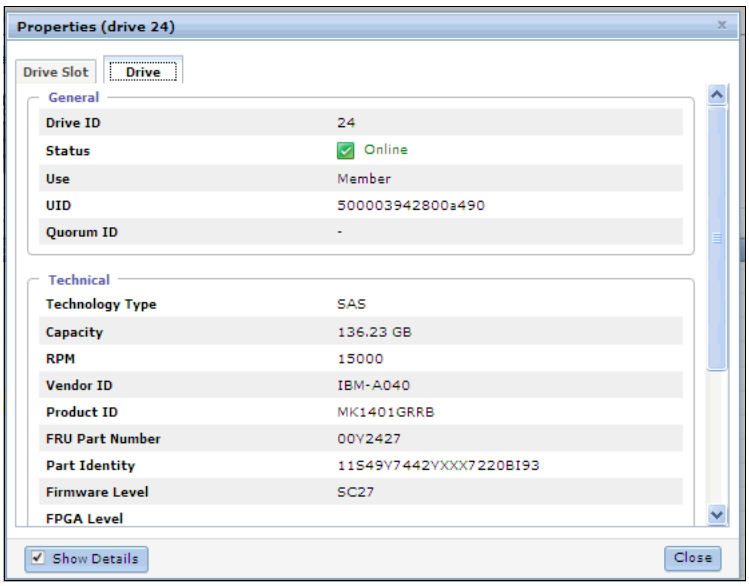

Figure 7-16 Internal drives properties: Part 1

When the Show Details option is selected, more information is displayed, as shown in Figure 7-17.

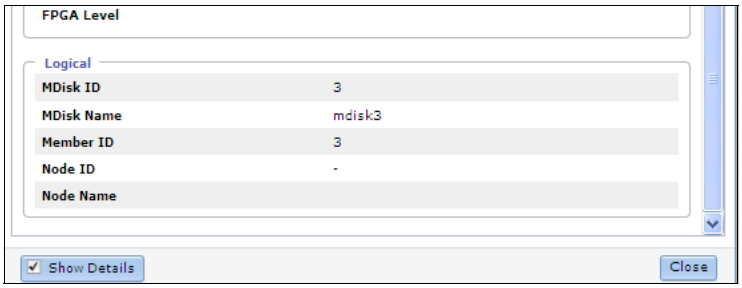

Figure 7-17 Internal drives properties show details

If the Show Details option is not selected, the technical information section is reduced, as shown in Figure 7-18.

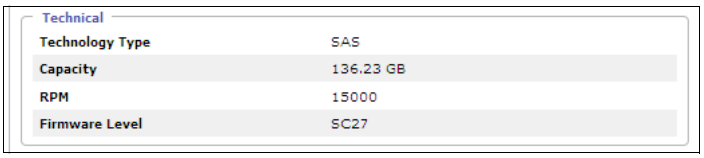

Figure 7-18 Internal drives properties no details

Within the Properties panel, a tag for the drive slot is presented (as shown in Figure 7-19) to get specific information about the slot of the selected drive.

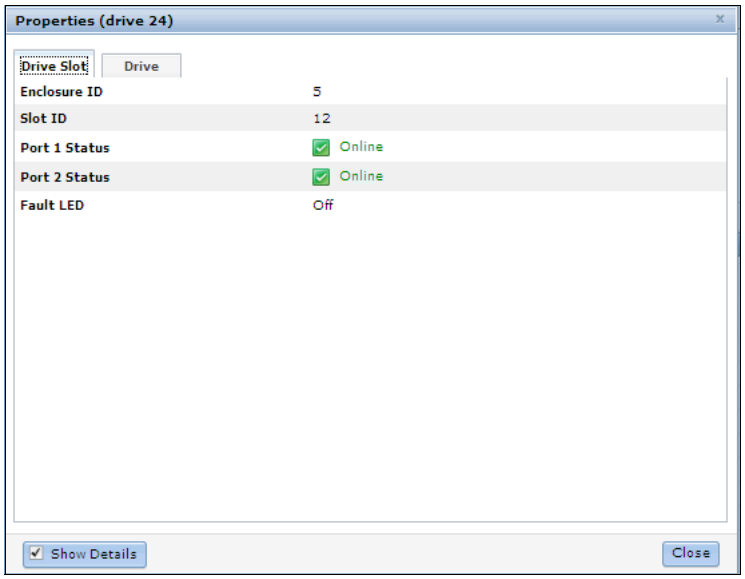

Figure 7-19 Internal drive properties slot

Example 7-5 shows how to use the **lsdrive** command to display configuration information and drive VPD.

Example 7-5 The use of the lsdrive command to display configuration information and drive VPD

```
lsdrive driveID
```

7.2 Configuring internal storage

The complete internal storage of an IBM Storwize V3500 is configured into MDisks and pools by using the system setup wizard during the initial configuration. For more information about this process, see Chapter 2, “Initial configuration” on page 27.

The decision that is show in Figure 7-20 must be made.

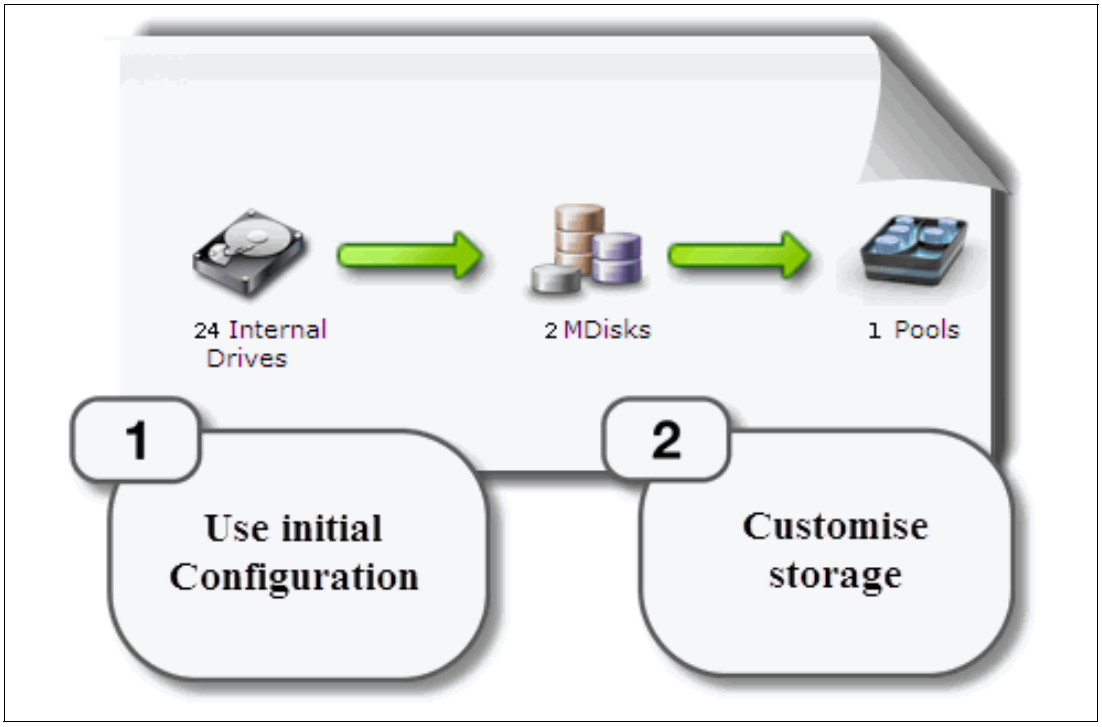

Figure 7-20 Decision to customize storage configuration

The decision choices include the following meanings:

- Use initial configuration

  During system setup all available drives were 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.

- Customise storage configuration

  A storage configuration might be customized for the following reasons:

  – Automated initial configuration does not meet customer requirements.

  – More storage was attached to the IBM Storwize V3500 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 shows an overview of the presets.

Table 7-1 RAID levels that are supported by an IBM Storwize V3500

RAID level	Where data is striped	Drive count (Min - Max)	Spare drive goal
0	Arrays have no redundancy and do not support hot-spare takeover.	1 - 8	0
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	1
5	Arrays stripe data over the member drives with one parity stripe 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	1
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	1
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	1

7.2.2 Customize initial storage configuration

If the initial storage configuration does not meet requirements, pools should be deleted. Select the Pool navigator in the GUI and click **Pools** → **MDisks by Pools**. Select the pool, right-click it, and select **Delete Pool**, as shown in Figure 7-21.

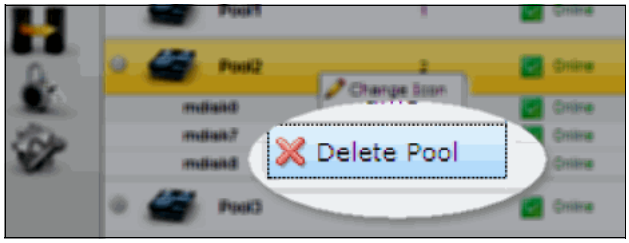

Figure 7-21 Delete selected pool

The option for deleting MDisk must be marked so all associated drives are marked as a candidate for deletion, as shown in Figure 7-22.

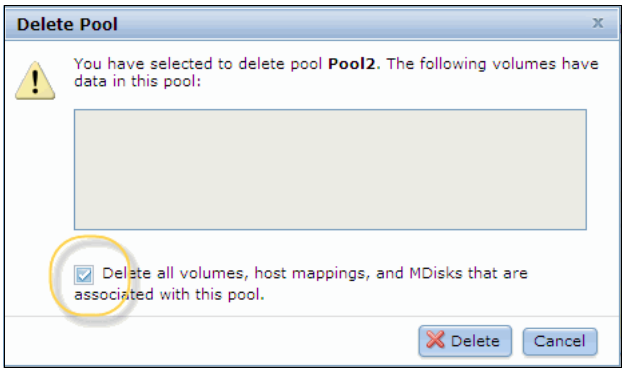

Figure 7-22 Delete pool confirmation

These drives now can be used for a different configuration.

7.2.3 Create MDisks and pools

To configure internal storage for use with hosts, click **Pools** → **Internal Storage** and click **Configure Storage**, as in Figure 7-23.

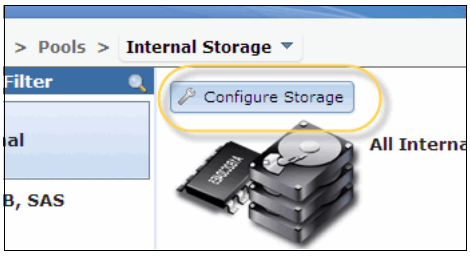

Figure 7-23 Click Configure Storage

A configuration wizard opens and guides you through the process of configuring internal storage. The wizard shows all internal drives with a status of whether it is available for configuration. Figure 7-24 shows an example of which different drives types are available for configuration.

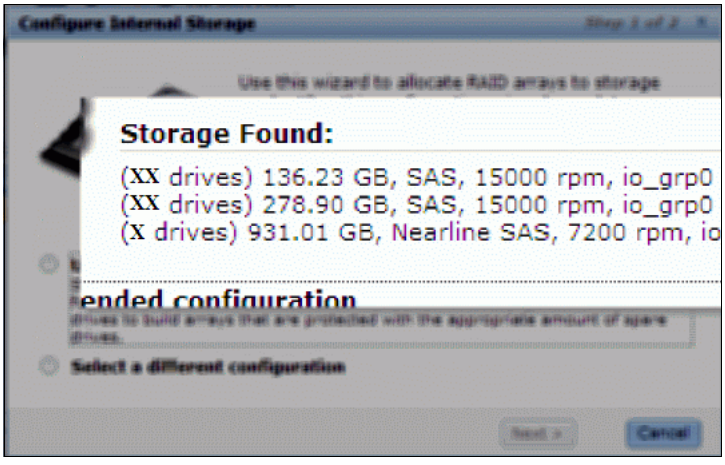

Figure 7-24 Available drives for new MDisk

If there are internal drives with a status of unused, a window opens (as shown in Figure 7-25), which gives the option to include them in the RAID configuration.

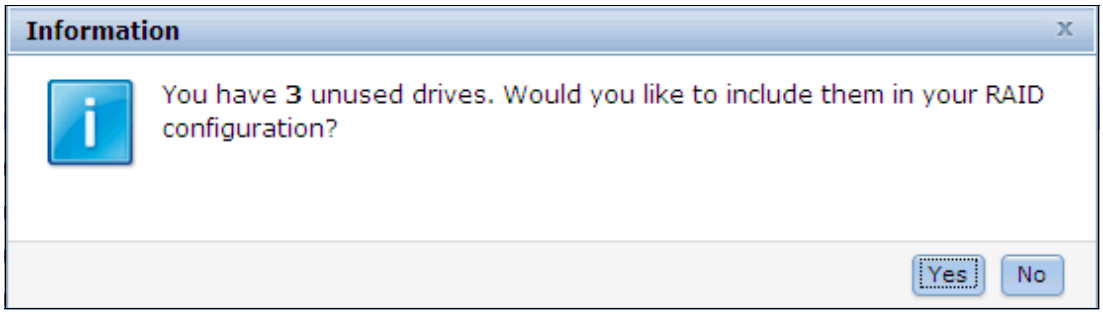

Figure 7-25 Unused drives warning

When the decision was made to include the drives into the RAID, their status is set to candidate, which also makes them available for new MDisks.

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 “Using the recommended configuration” on page 293. Selecting **Select a different configuration** uses the wizard that is described in “Selecting a different configuration” on page 295.

7.2.4 Using the recommended configuration

As shown in Figure 7-26, when you click **Use the recommended configuration**, the wizard offers a recommended storage configuration at the bottom of the window.

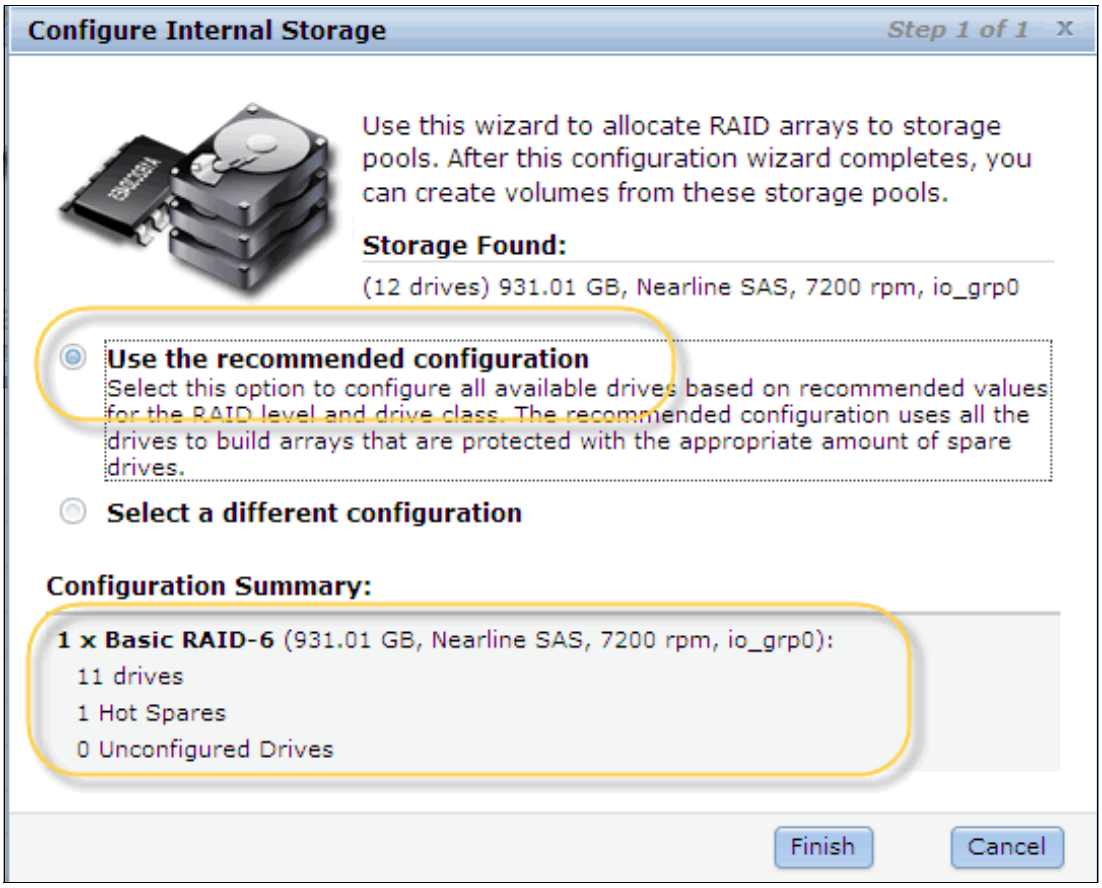

Figure 7-26 The recommended configuration

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

- Basic RAID 5 for SAS drives
- Basic RAID 6 for Nearline SAS drives

By using the recommended configuration, spare drives also are automatically created to meet the spare goals according to the preset that is chosen.

Spare drives in the IBM Storwize V3500 are *global spares*, which means that any spare drive that has at least the same capacity as the drive that must be replaced can be used in any array.

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 option must be omitted, the Select different configuration option is 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 7-27.

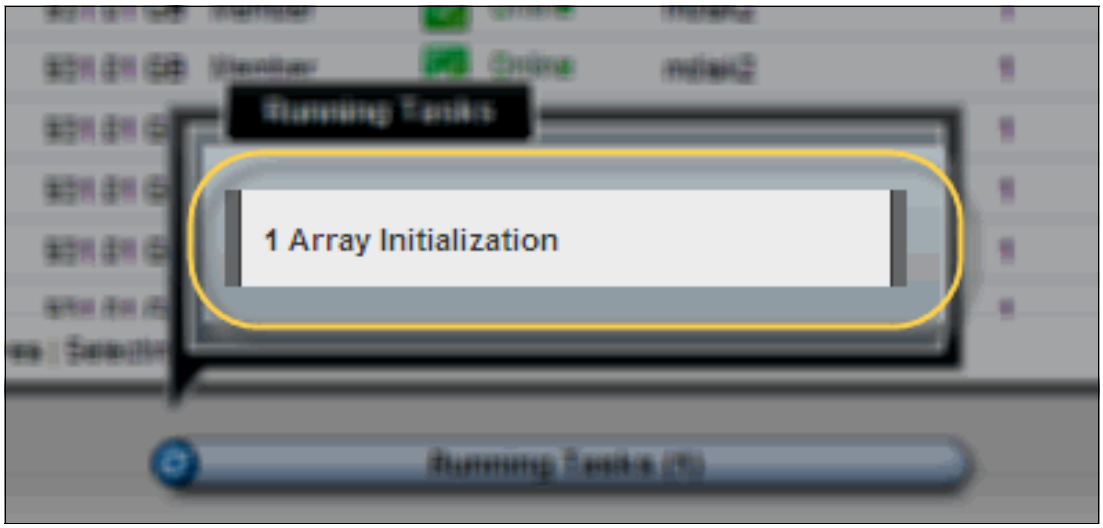

Figure 7-27 Running task panel

Click the task bar to open the progress window, as shown in Figure 7-28. The array is available for I/O during this process. The initialization does not affect the availability because of possible member drive failures.

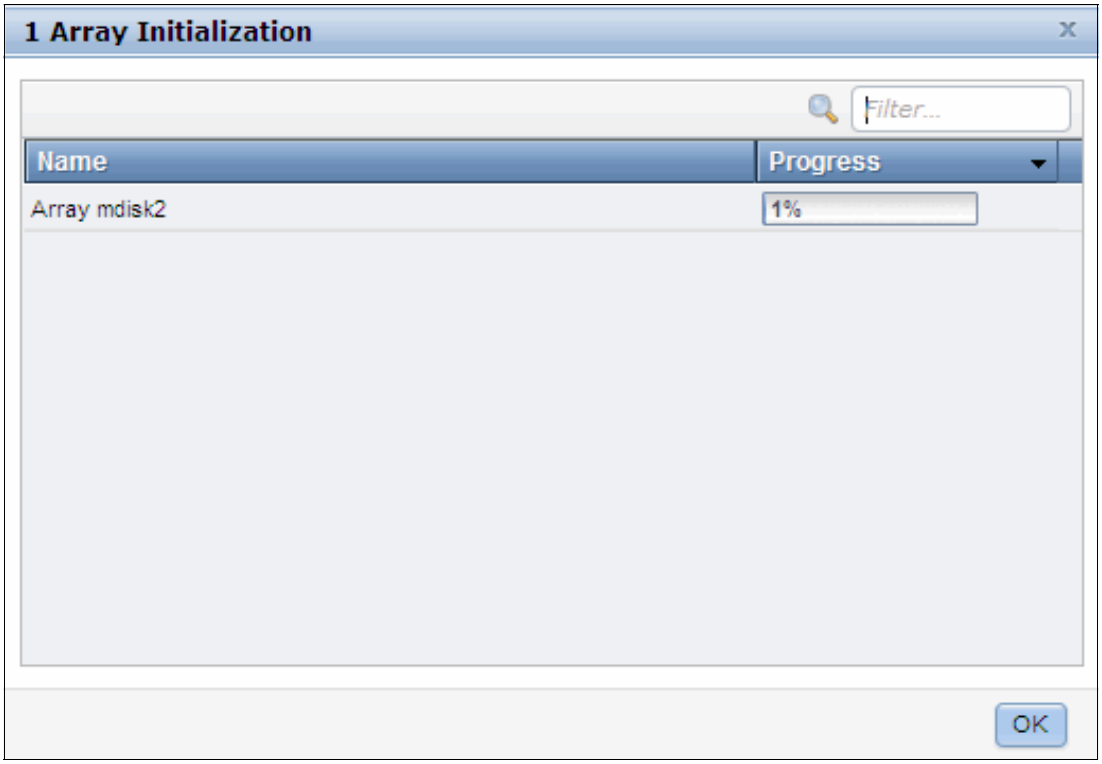

Figure 7-28 Initialization progress view

7.2.5 Selecting a different configuration

The Select a different configuration option offers a more flexible way for the configuration of 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. Select the appropriated RAID preset.

2. Define the RAID attributes.

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

   Each IBM Storwize V3500 preset has a specific goal for the number of drives per array

   For more information, see this website:

   http://pic.dhe.ibm.com/infocenter/storwize/V3500_ic/index.jsp

Table 7-2 shows the RAID goal widths.

Table 7-2 RAID goal width

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

Optimizing for performance creates arrays with the same capacity and performance characteristics. The RAID goal width (as shown in Figure 7-29) must be met for this target.

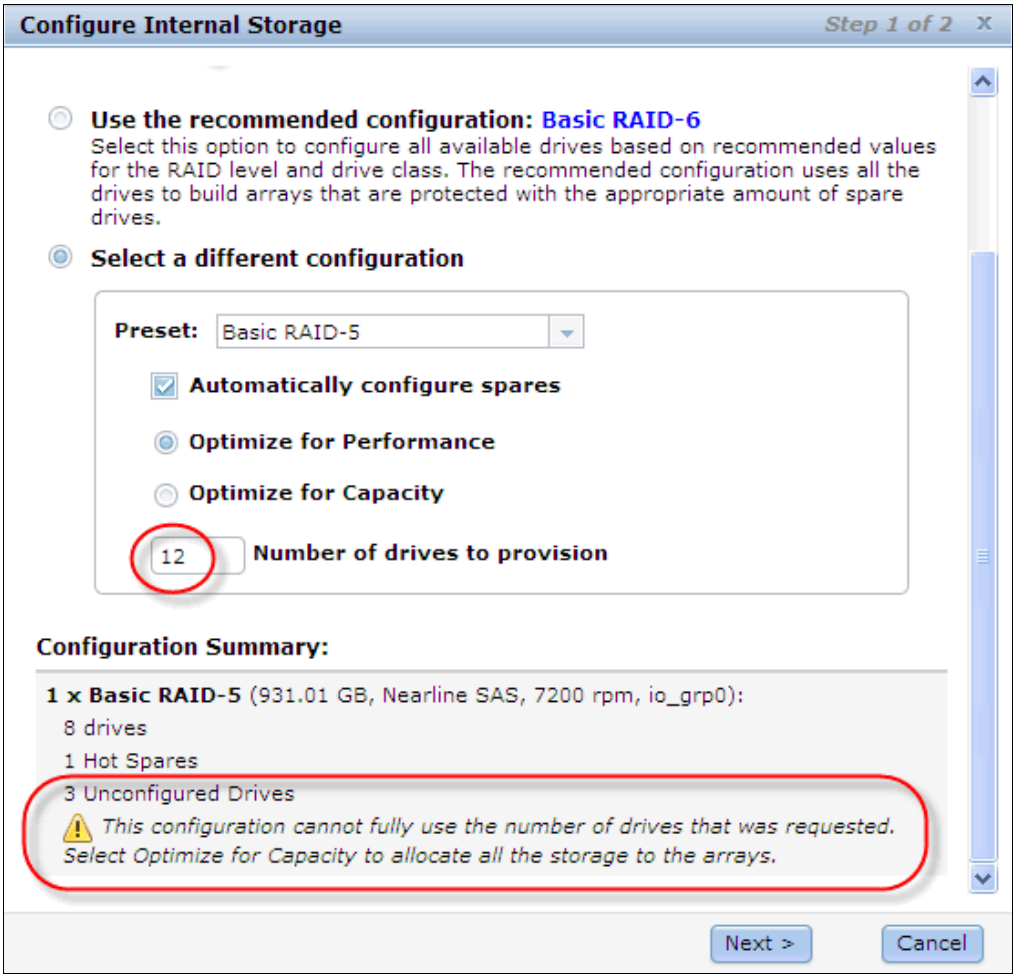

Figure 7-29 Optimization for performance failed

Figure 7-30 shows that not enough drives are available to meet the requirements for the configuration.

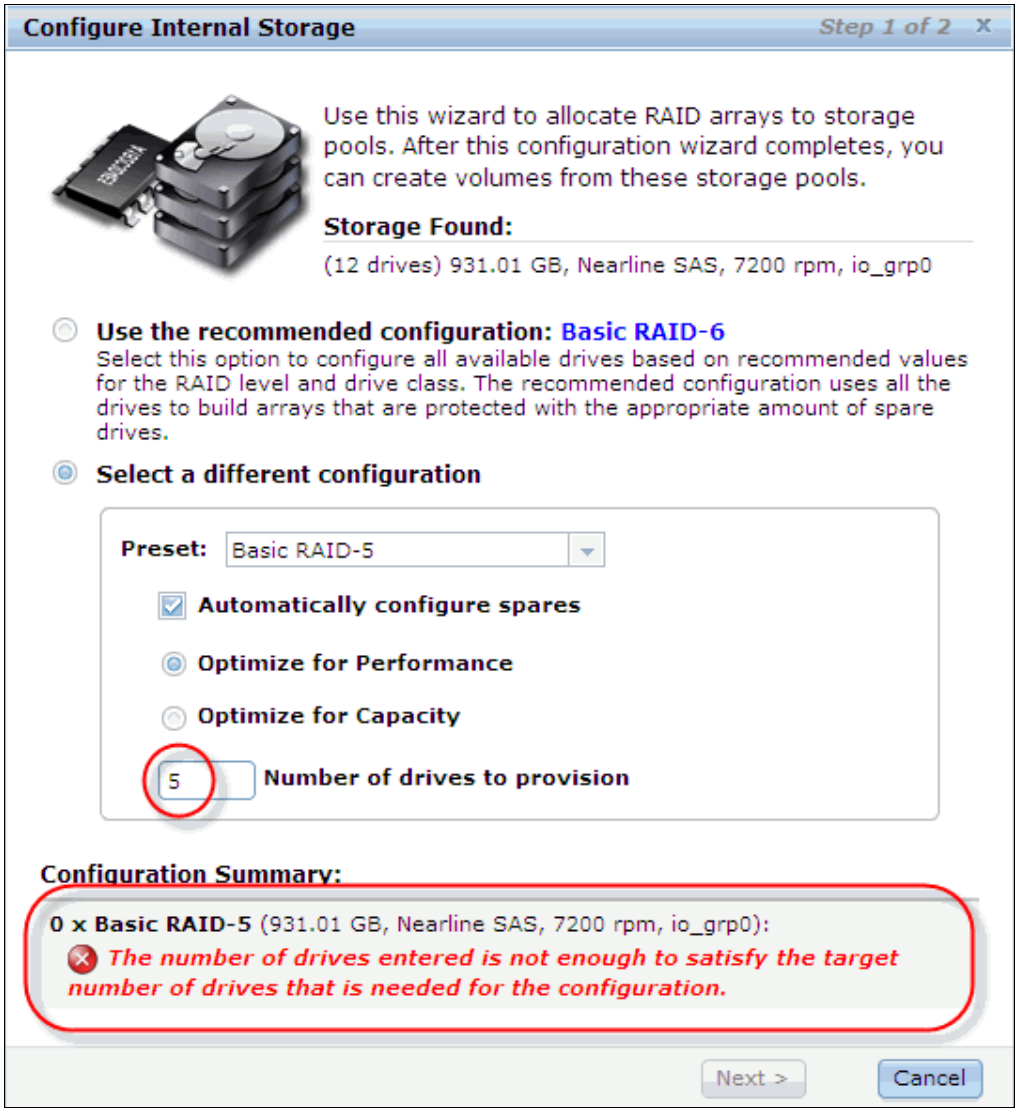

Figure 7-30 Less drive for performance optimization

Figure 7-31 shows that there are a suitable number of drives to configure performance optimized arrays.

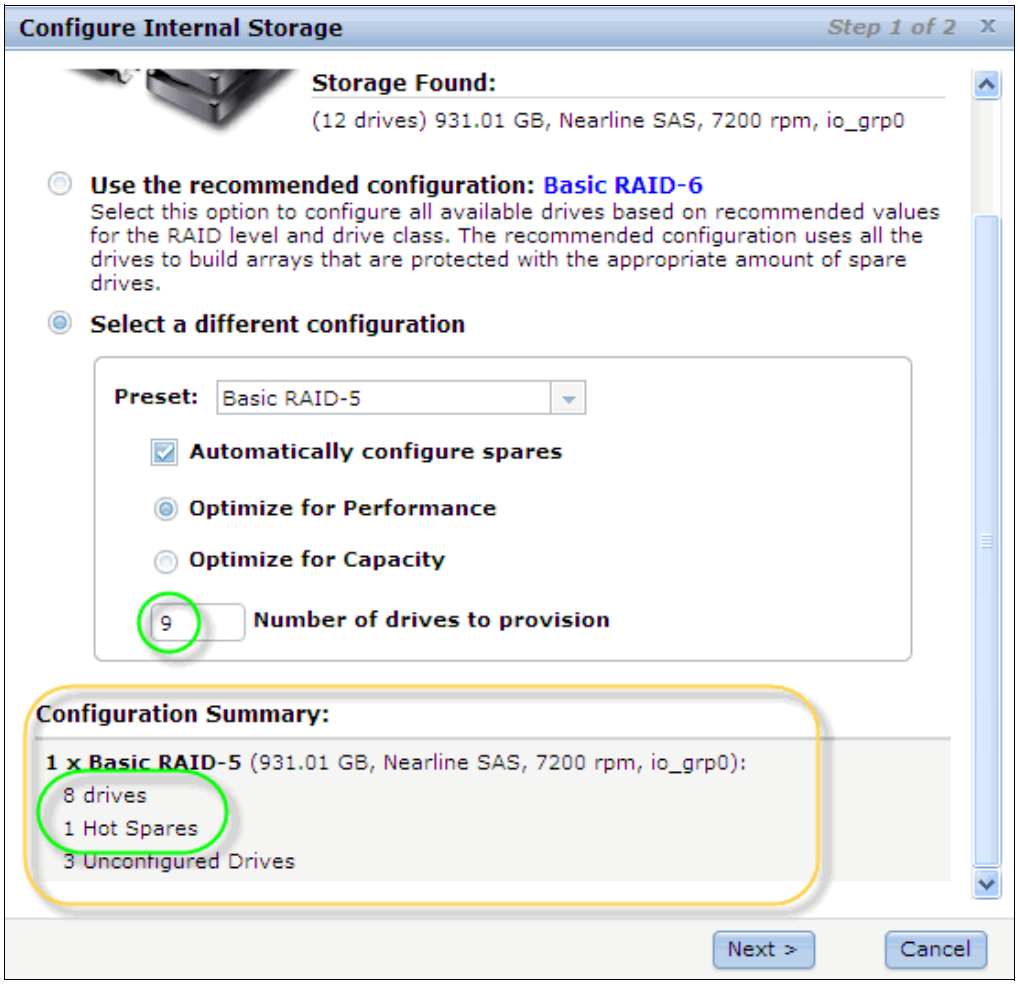

Figure 7-31 Array matching performance goal

One RAID 5 was built and all provisioned drives are used.

Optimizing for capacity creates arrays that allocate whole numbers of drives to provision. It also results in arrays of different capacities and performance, as shown in Figure 7-32.

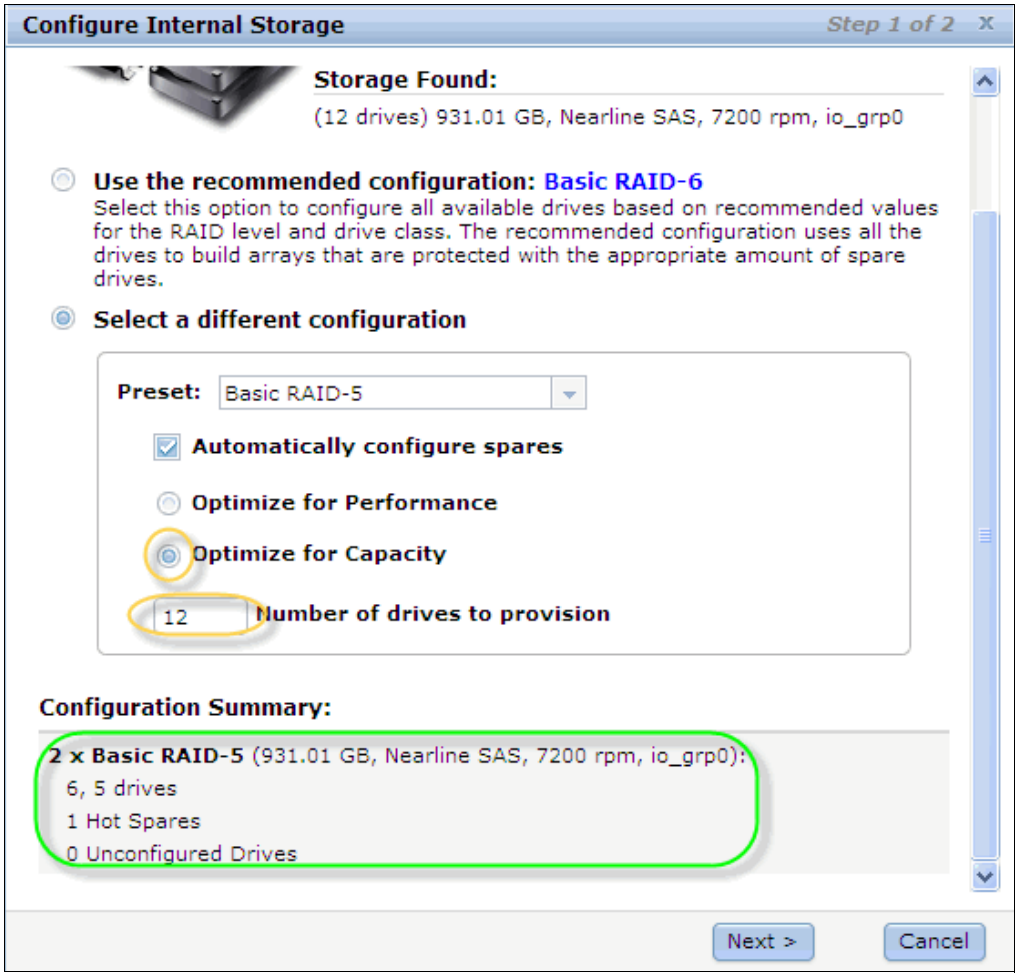

Figure 7-32 Capacity optimized configuration

3. Storage pool assignment

   Choose storage pool assignment if an existing pool should be expanded or if a new pool is to be created for the configuration, as show in Figure 7-33.

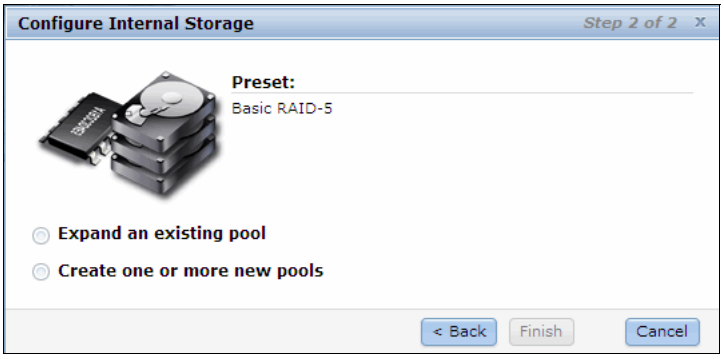

Figure 7-33 Storage pool selection

Complete the following steps to expand or create a pool:

a. Expand existing pool

When an existing pool should be expanded, all available pools without MDisk or those pools that do not match the selected RAID presets are listed, as shown in Figure 7-34.

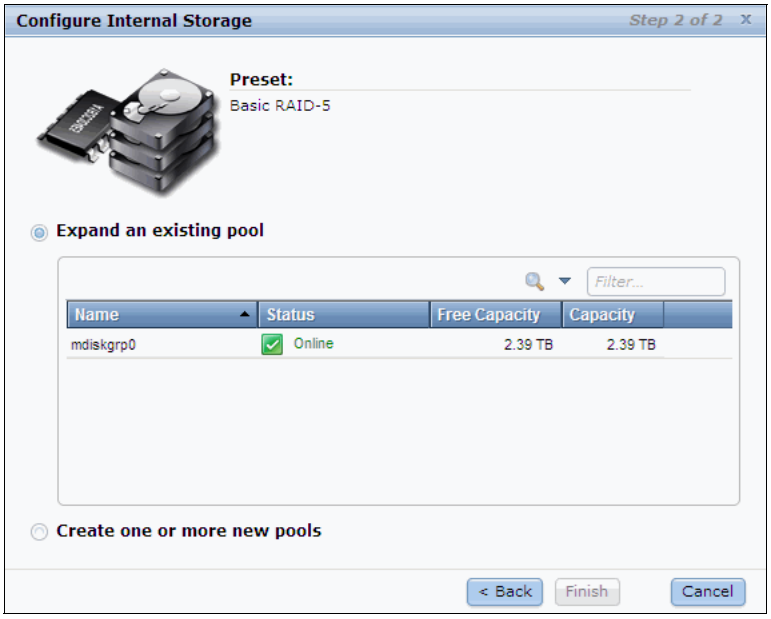

Figure 7-34 List of matching storage pool

b. Create one new pool

Alternatively, a new storage pool is created by using the name that is entered, as shown in Figure 7-35.

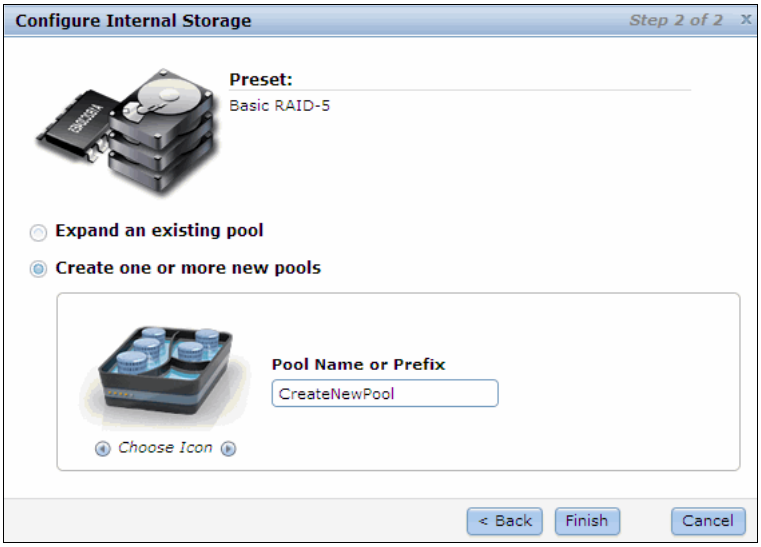

Figure 7-35 Create new pool

All drives are initialized when the configuration wizard is finished.

7.3 Working with MDisks

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

You can access the MDisks window by clicking **Home** → **Overview** and ten clicking the **MDisks** function icon. In the extended help information window, click **Pools**, as shown in Figure 7-36.

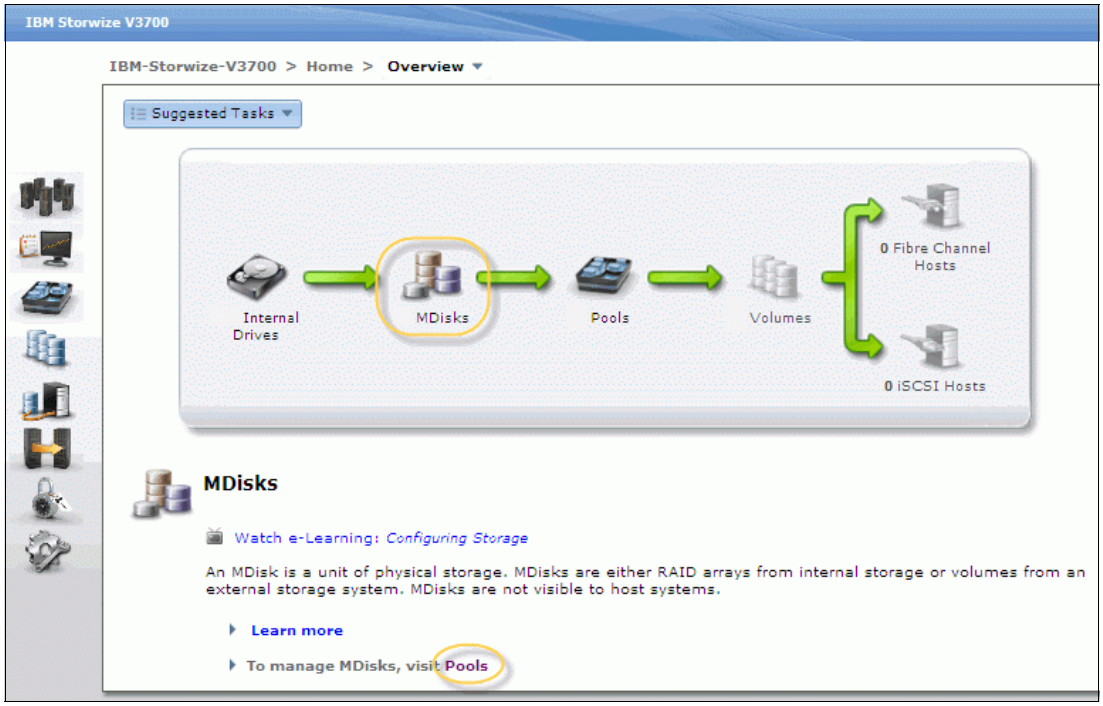

Figure 7-36 MDisk from Overview window

Another way to access the MDisks window is by using the Physical Storage function icon that is on the left, as shown in Figure 7-37.

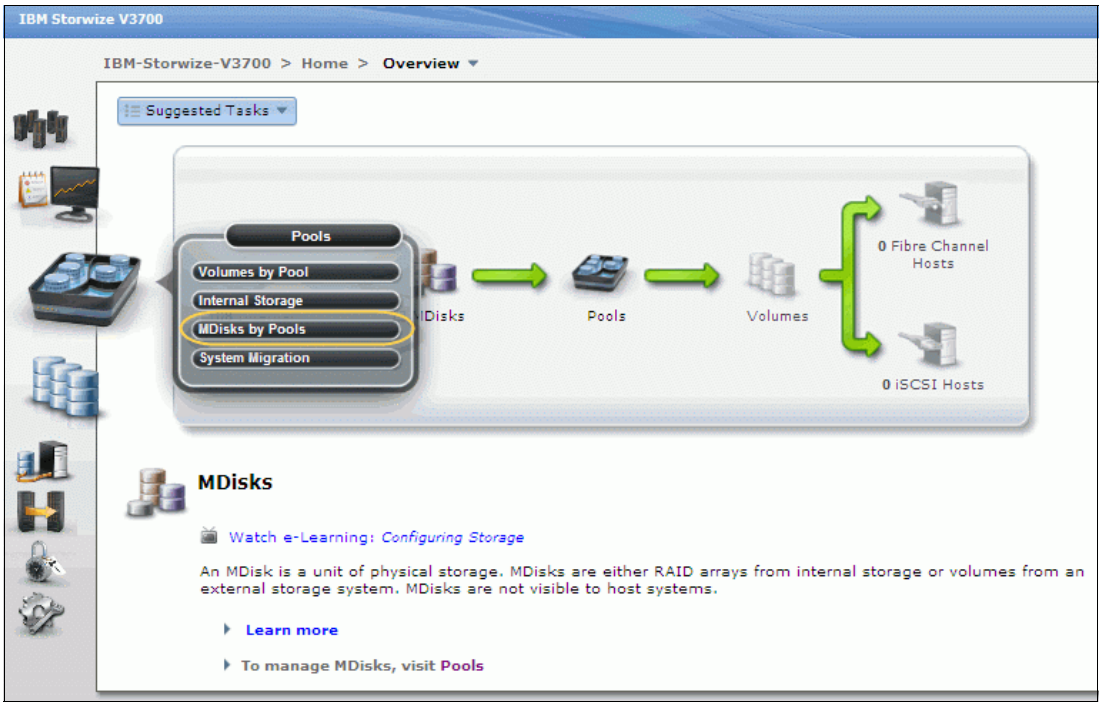

Figure 7-37 MDisk from Physical Storage icon

7.3.1 MDisk by Pools panel

The MDisks by Pools panel (as shown in Figure 7-38) displays information about each of the MDisks. The MDisks are categorized by the pools to which they are attached.

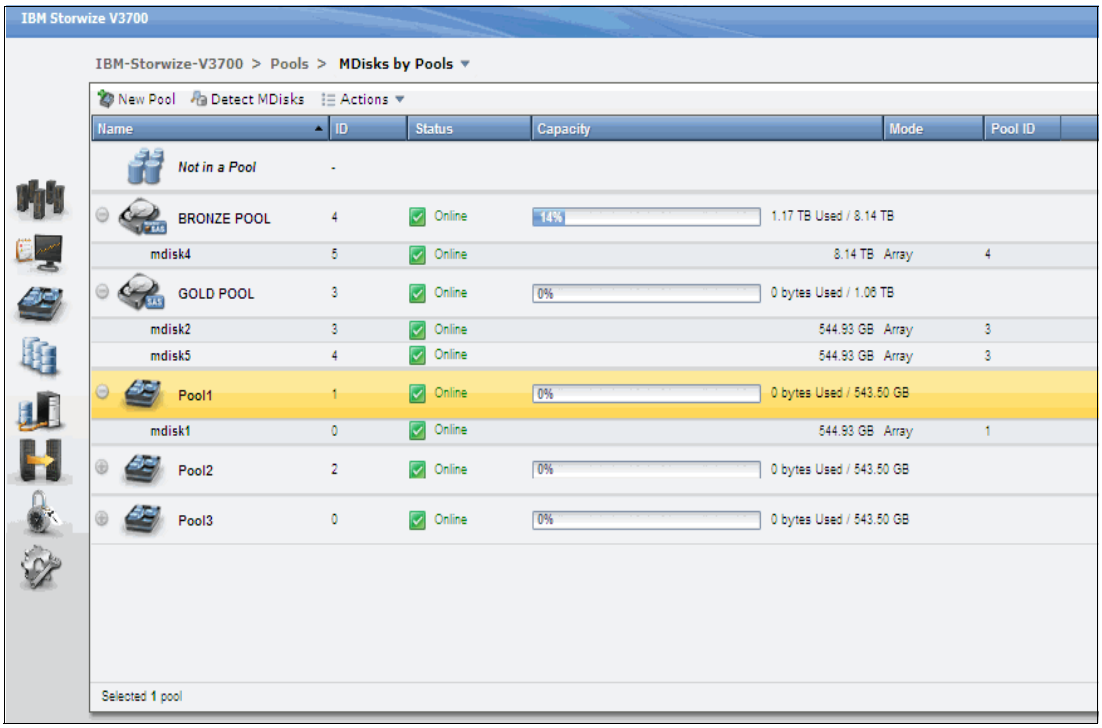

Figure 7-38 MDisk by Pool window

The following default information is provided:

- Name

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

- ID

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

- Status

  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 7-39, the pool also is degraded.

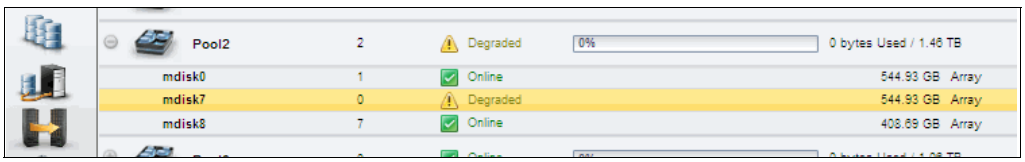

Figure 7-39 One degraded MDisk in pool

– Offline

One or more MDisk in a pool are offline. The pool (Pool3) also changes to offline, as shown in Figure 7-40.

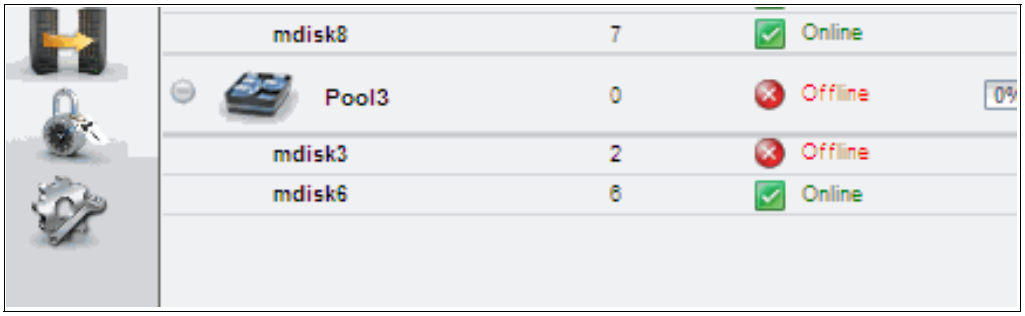

Figure 7-40 Offline MDisk in a pool

- Capacity

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

- Mode

  The mode of the MDisk. The following modes are available in the IBM Storwize V3500:

  – Array

    The MDisk represents a set of drives from internal storage that are managed together by using RAID.

  – Image/unmanaged

    This status is an intermediate status of the migration process and is described in Chapter 6, “Storage migration wizard” on page 217.

- Storage Pool

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

The CLI command l**smdiskgrp** (as shown in Example 7-6) returns a concise list or a detailed view of the storage pools that are visible to the system.

Example 7-6 CLI command lsmdiskgrp

```
lsmdiskgrp
lsmdiskgrp mdiskgrpID
```

7.3.2 RAID action for MDisks

Internal drives in the IBM Storwize V3500 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 7-41.

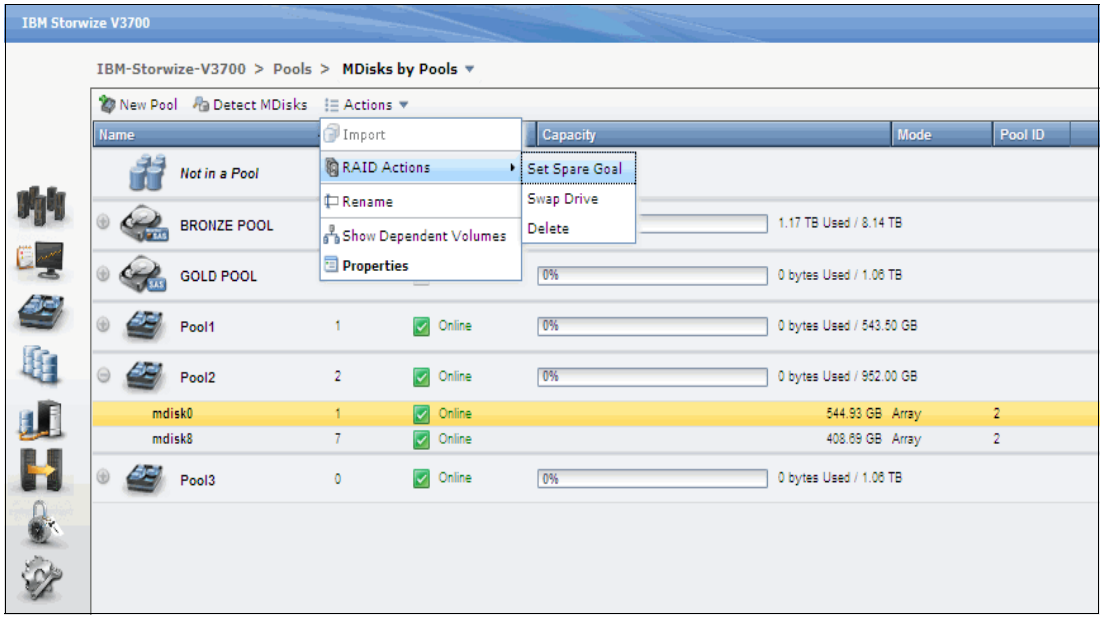

Figure 7-41 MDisk RAID actions

You can choose the following RAID actions:

- Set Spare Goal

  Figure 7-42 shows how to set the number of spare drives that are required to protect the array from drive failures.

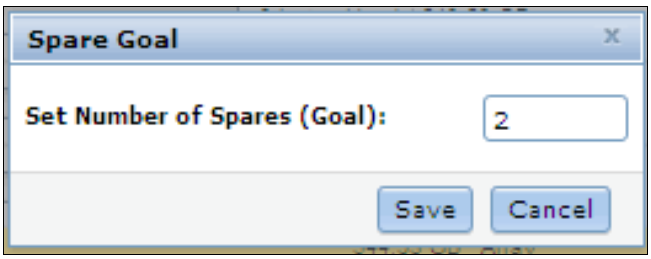

Figure 7-42 MDisk set spare goal

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

Example 7-7 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 drives as spare. 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 289.

- 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 to be replaced (as shown in Figure 7-43) and click **Next**.

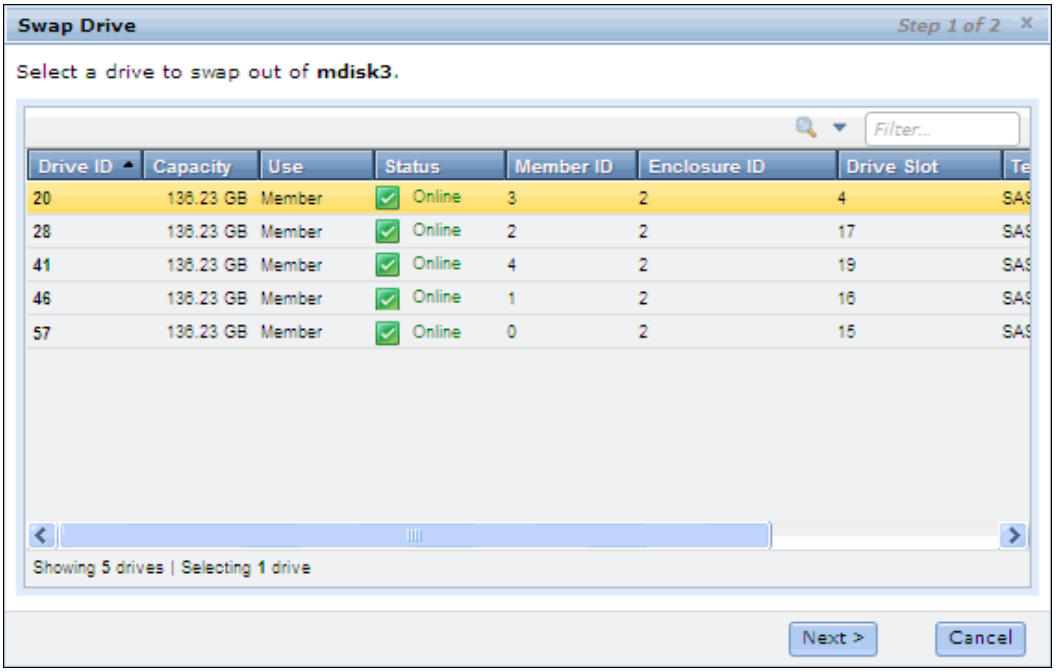

Figure 7-43 MDisk swap drive: Step 1

In step 2 (as shown in Figure 7-44), a list of suitable drives is presented. One drive must be selected to swap into the MDisk. Click **Finish**.

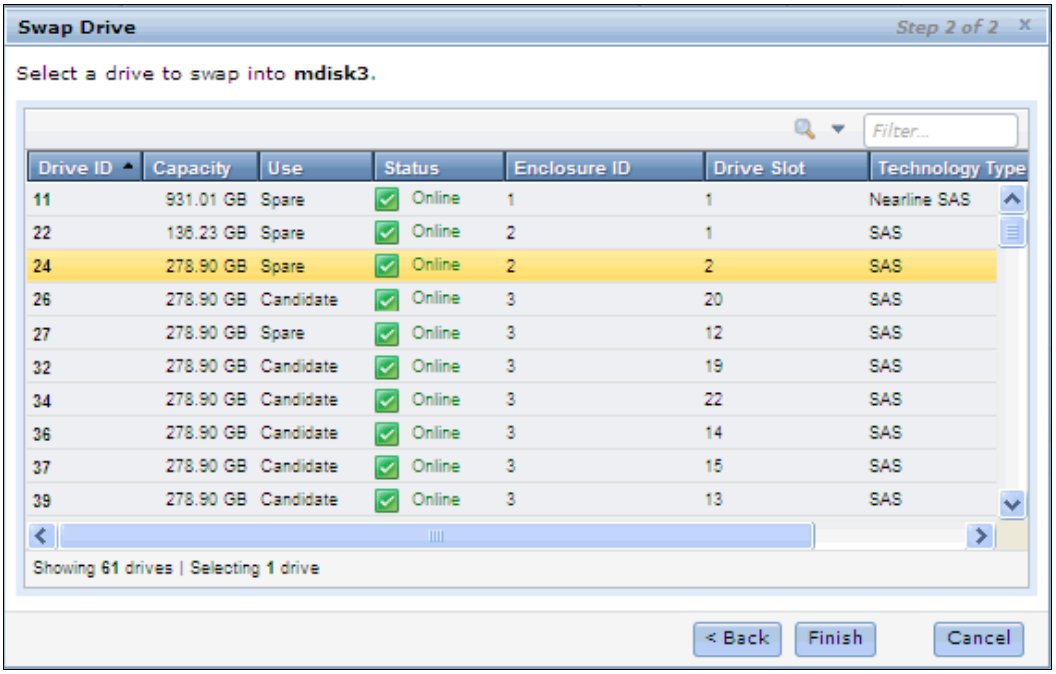

Figure 7-44 MDisk swap drive pool step 2

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

If for any reason the GUI process is not to be used, the CLI command that is shown in Example 7-8 can be run.

Example 7-8 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, use Ctrl+left-mouse click. A confirmation is required (as shown in Figure 7-45) by entering the correct number of MDisks to be deleted. 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.

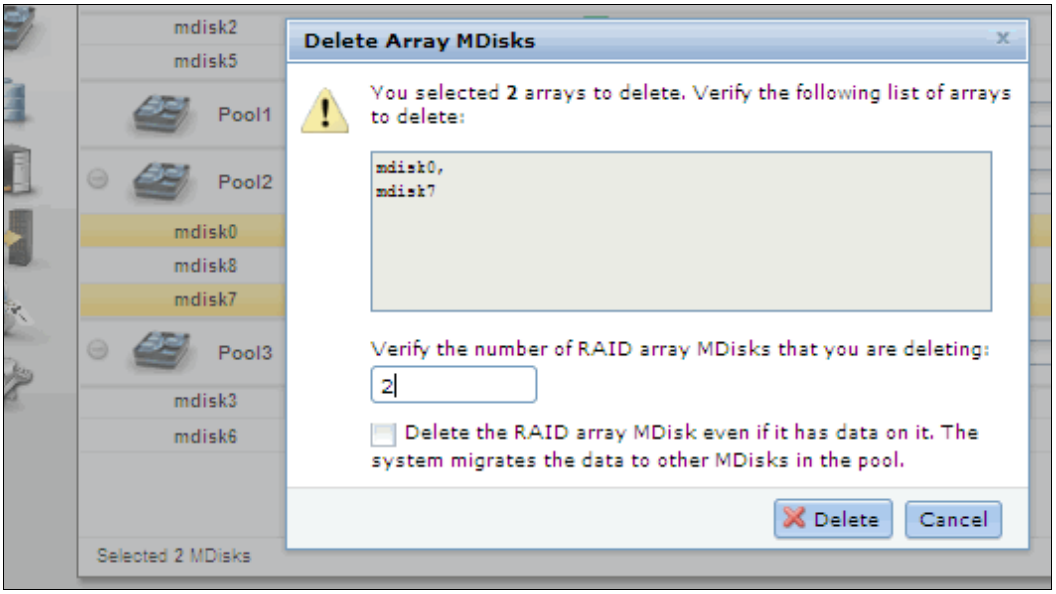

Figure 7-45 MDisk delete confirmation

Data on MDisks are migrated to other MDisks in the pool, assuming 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 7-9.

Example 7-9 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 7-46.

Figure 7-46 Empty storage pool after MDisk deletion

7.3.3 More actions on MDisks

Some of the other actions are available by clicking **MDisk by Pool** → **Action**, as shown in Figure 7-47.

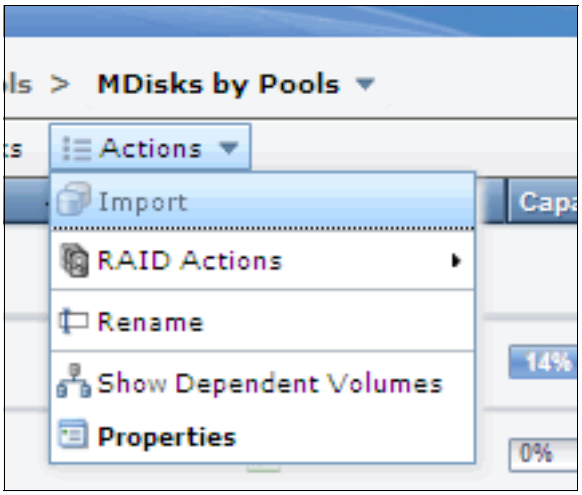

Figure 7-47 MDisk actions

Rename

MDisks can be renamed by selecting the MDisk and clicking **Rename** from the Actions drop-down menu. Enter the new name of your MDisk (as shown in Figure 7-48) and click **Rename**.

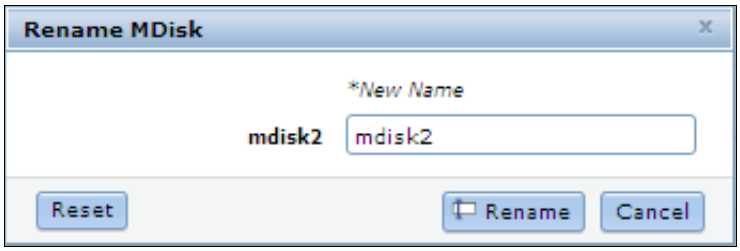

Figure 7-48 Rename MDisk

Show Dependent Volumes

Select the MDisk that should contain the dependent volumes by clicking **Show Dependent Volumes** from the Actions menu. The volumes should be listed with general information.

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) contains information about the MDisk. To show more details, click **Show Details**.

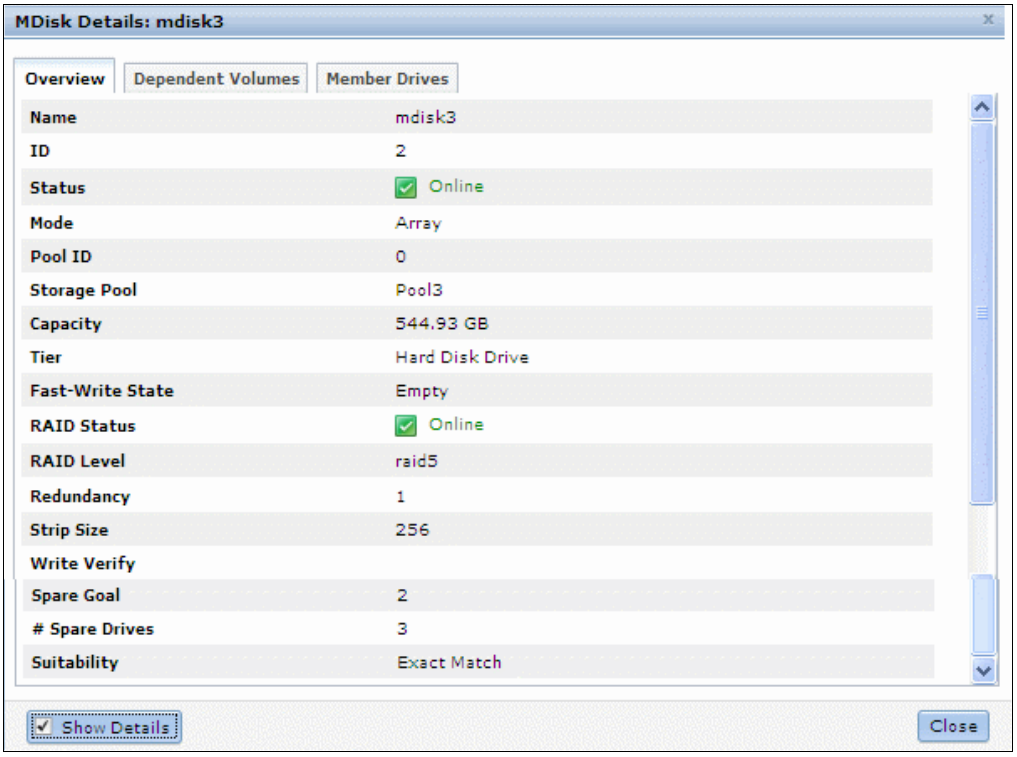

Figure 7-49 MDisk properties overview

- The Dependent Volumes tab lists all of the volumes that use extents on this MDisk.

- In the Member Drives tab (as shown in Figure 7-50), you find all the member drives of this MDisk. Also, all actions that are described in 7.1.2, “Actions on internal drives” on page 282 can be performed on the drives that are listed here.

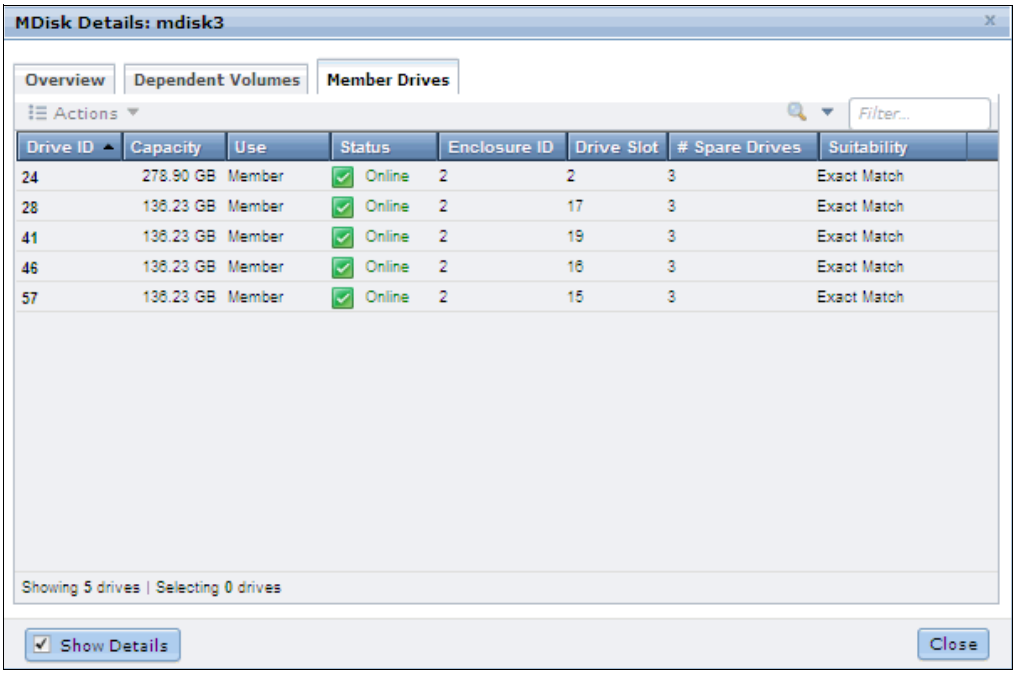

Figure 7-50 MDisk properties member

7.4 Working with storage pools

Storage pools act as a container for MDisks and provision the capacity to volumes. IBM Storwize V3500 organizes storage in storage pools to ease storage management and make it more efficient. Storage pools and MDisks are managed in 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 7-51.

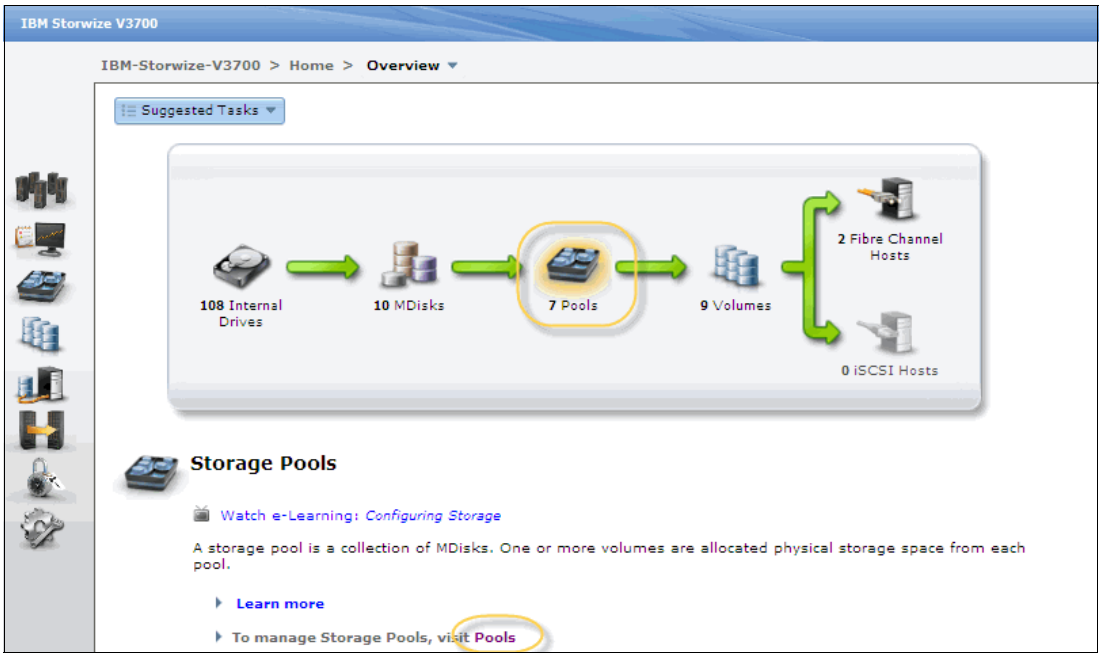

Figure 7-51 Pools from the overview window

An alternative path to the Pools window is to click **Pools** → **MDisks by Pools**, as shown in Figure 7-52.

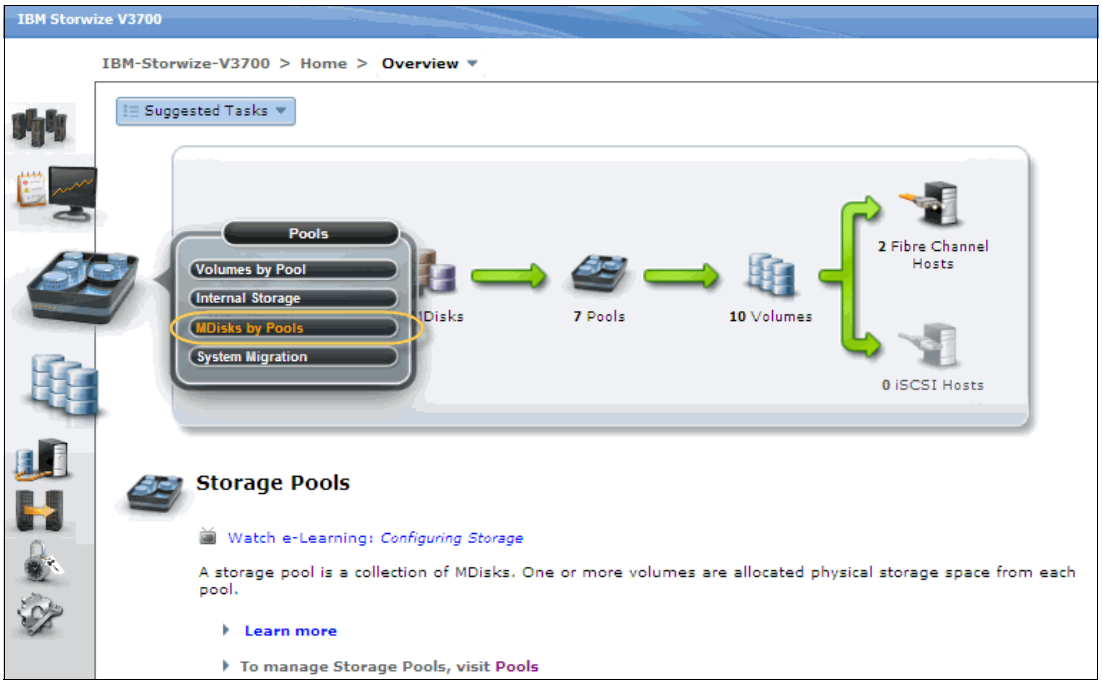

Figure 7-52 Pools from MDisk by pool window

By using the MDisk by Pools window (as shown in Figure 7-53), you can manage storage pools. All existing storage pools are displayed row-by-row. The first row contains the item Not in a Pool, 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 is allocated to volumes.

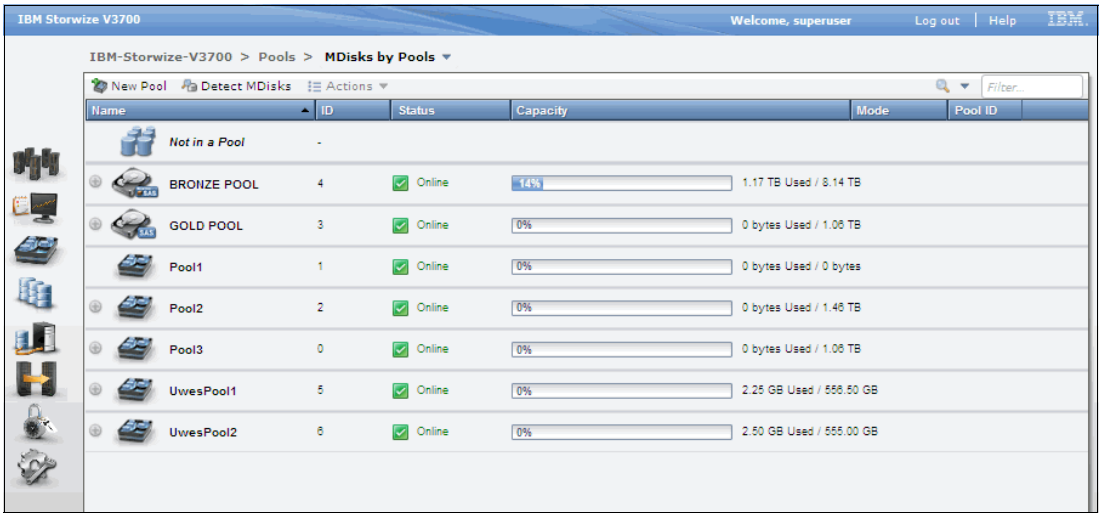

Figure 7-53 Pool window

7.4.1 Creating a pool

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

The only required parameter for the pool is the pool name, which must be entered, as shown in Figure 7-54.

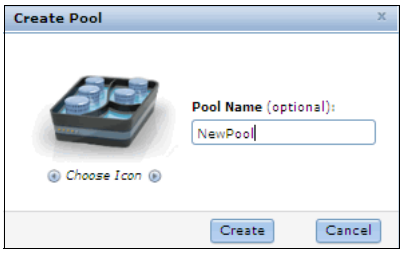

Figure 7-54 Create pool name input

The new pool is listed in the pool list with 0 bytes, as shown in Figure 7-55.

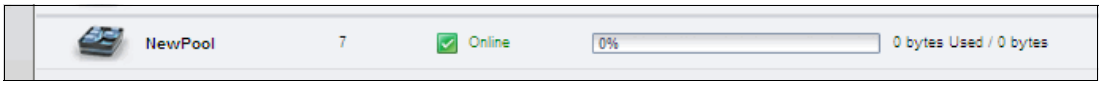

Figure 7-55 Empty pool that is created

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 7-56. A pool can be renamed, its icon can be changed, and it can be deleted.

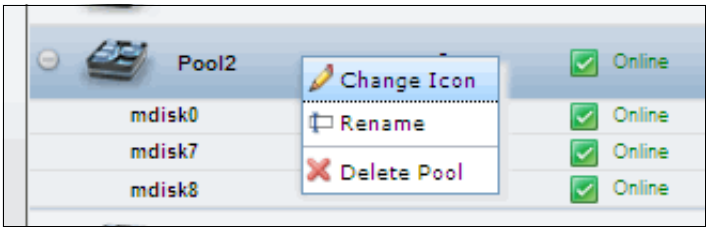

Figure 7-56 Pool action overview

Changing a storage pool icon

There are different storage pool icons available, which can be selected, as shown in Figure 7-57.

Figure 7-57 Change storage pool icon

Renaming a storage pool

The storage pool can be renamed at any time, as shown in Figure 7-58.

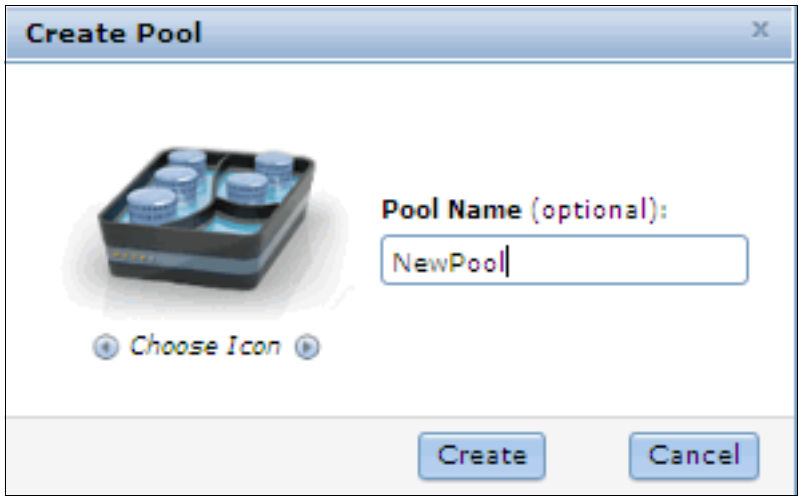

Figure 7-58 Rename Storage Pool

Deleting a storage pool

Pools can be deleted only if there are no MDisks or volumes assigned to it. A confirmation panel (as shown in Figure 7-59) appears to confirm that all associated MDisk and volumes can be deleted with the pool.

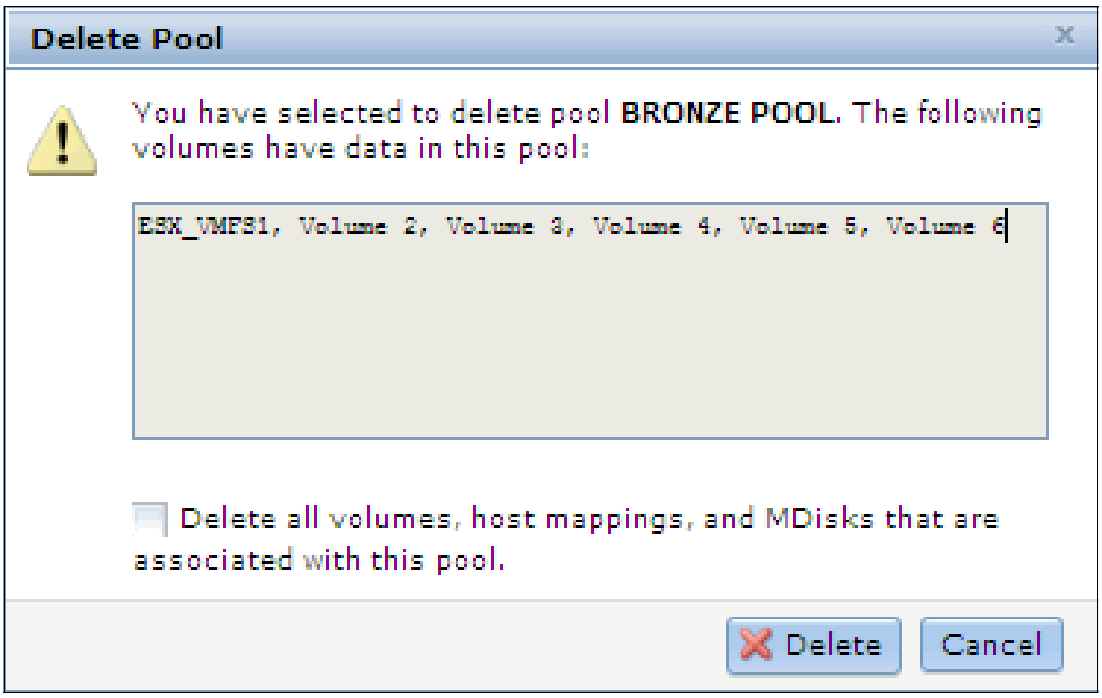

Figure 7-59 Confirmation to delete the storage pool

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.

8

Advanced host and volume administration

The IBM Storwize V3500 offers many functions for volume and hosts creation and configuration. In Chapter 4, “Host configuration” on page 141 and in Chapter 5, “Basic volume configuration” on page 171, the basic host and volume features of IBM Storwize V3500 are covered. Those chapters describe how to create hosts and volumes, and how to map them to a host.

This chapter includes the following topics:

- Advanced host administration
- Advanced volume administration

8.1 Advanced host administration

This section covers host administration, including host modification, host mappings, and deleting hosts. Basic host creation and mapping are described in Chapter 4, “Host configuration” on page 141. We assume here that you already created some hosts and that some volumes are mapped to them.

The following topics are covered in this section:

- All Hosts (8.1.1, “Modifying Mappings menu” on page 320).
- Ports by Host (8.2, “Adding and deleting host ports” on page 338).
- Host Mappings (8.3, “Host mappings overview” on page 348).

For more information, see Chapter 5, “Basic volume configuration” on page 171.

Figure 8-1 shows the IBM Storwize V3500 Hosts menu.

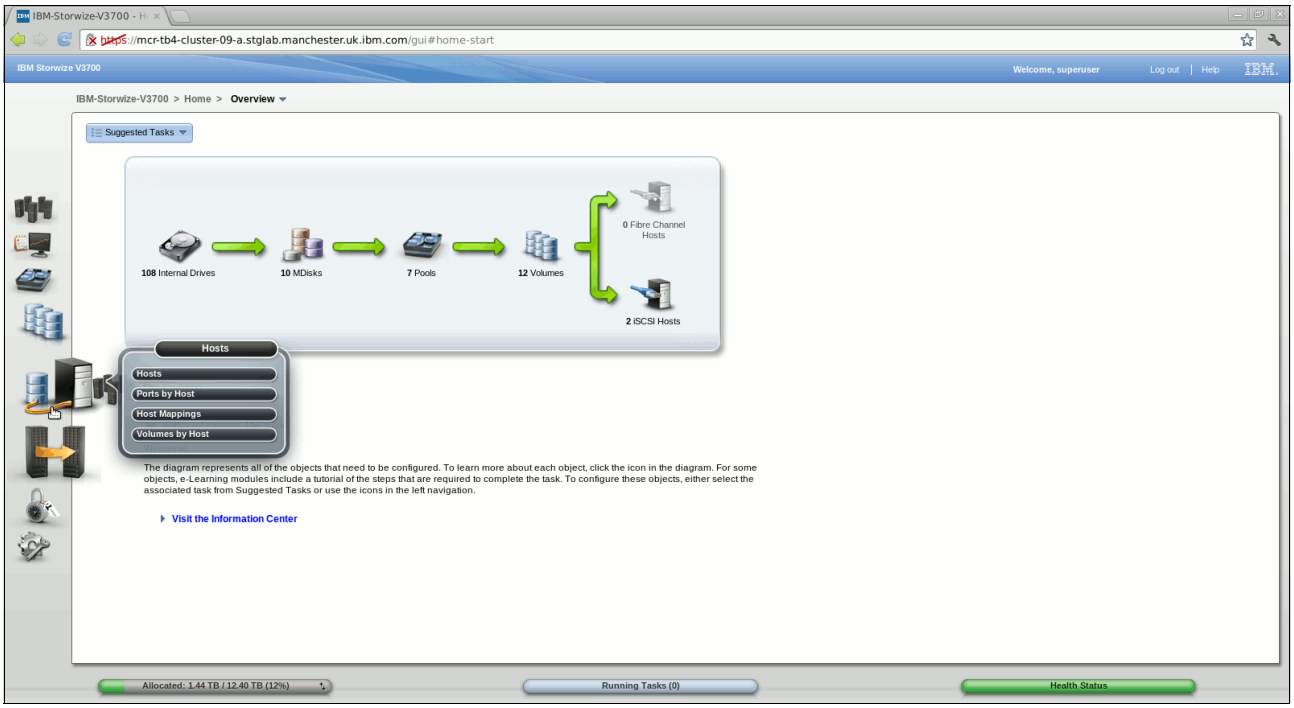

Figure 8-1 IBM Storwize V3500 Hosts menu

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

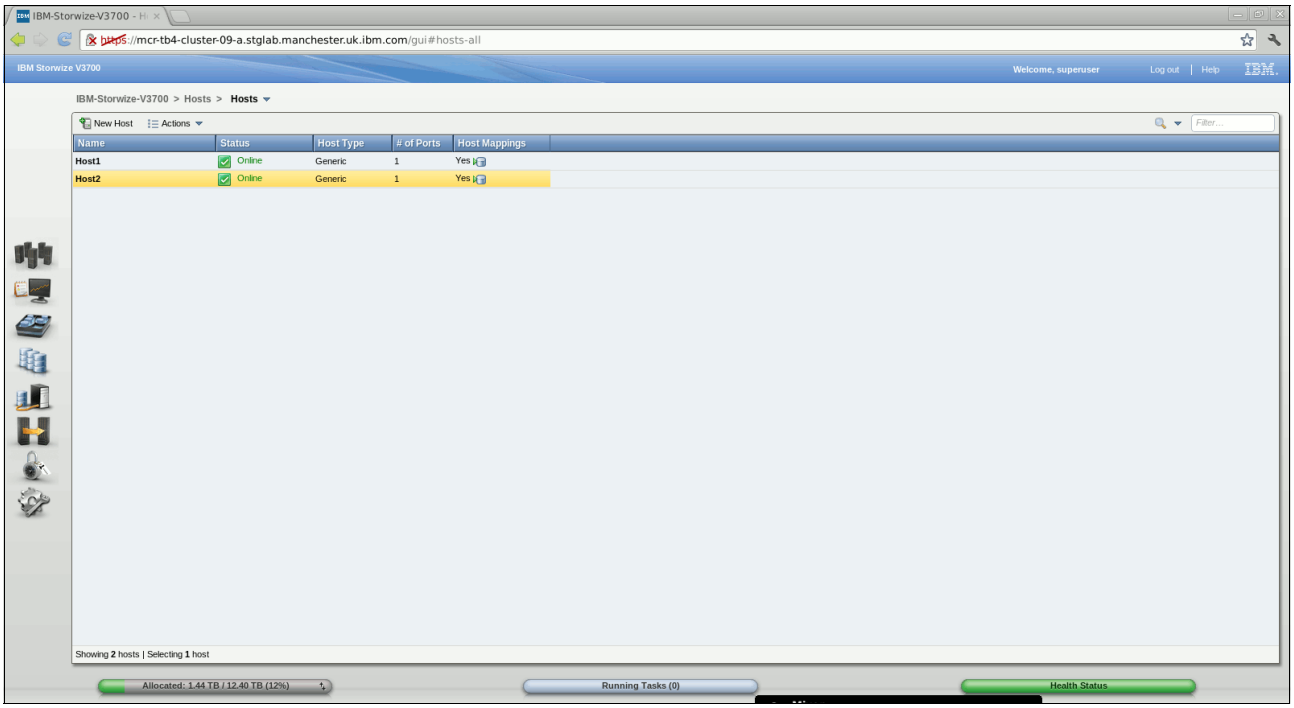

Figure 8-2 Hosts

In our example, two hosts are created and volumes are mapped to one of the hosts. We use these hosts to show the modification possibilities.

If you highlight a host, you can click **Action** (as shown in Figure 8-3) or right-click the host (as shown in Figure 8-4 on page 320) to see all of the available tasks.

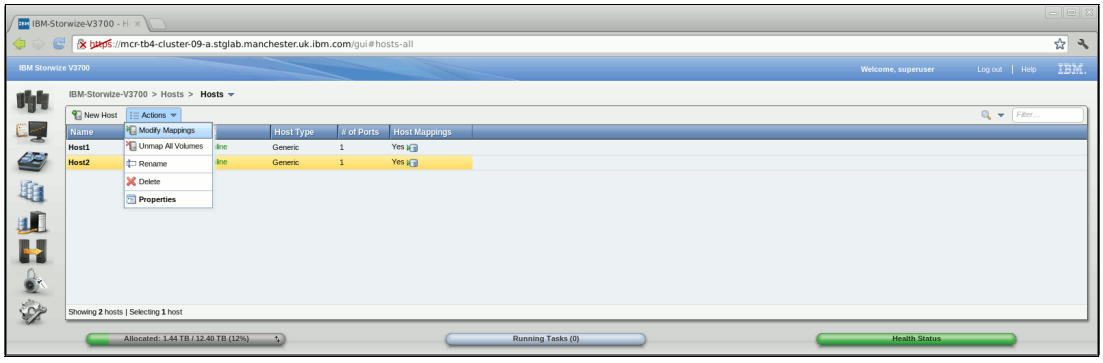

Figure 8-3 Hosts Action window

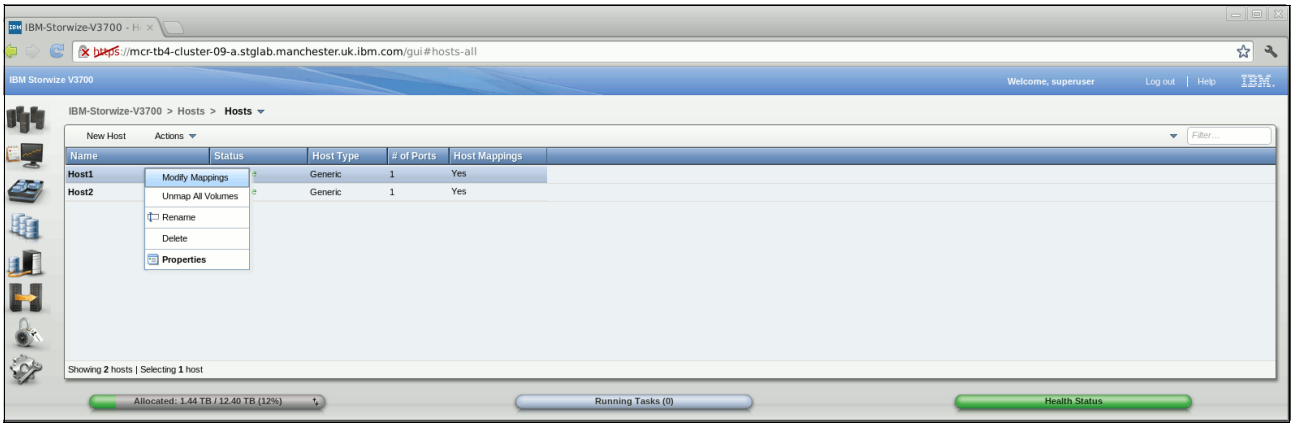

Figure 8-4 Host Menu

8.1.1 Modifying Mappings menu

Highlight a host and select **Modify Mappings** (as shown in Figure 8-3 on page 319) to open the window that is shown in Figure 8-5. In the upper left, you see that the highlighted host is selected. The two panels show all available unmapped and mapped volumes. The left pane shows the volumes that are ready for mapping to this host. The right box includes the volumes that are mapped. In our example, one volume with SCSI ID 0 is mapped to the host Host1 and three more volumes are available.

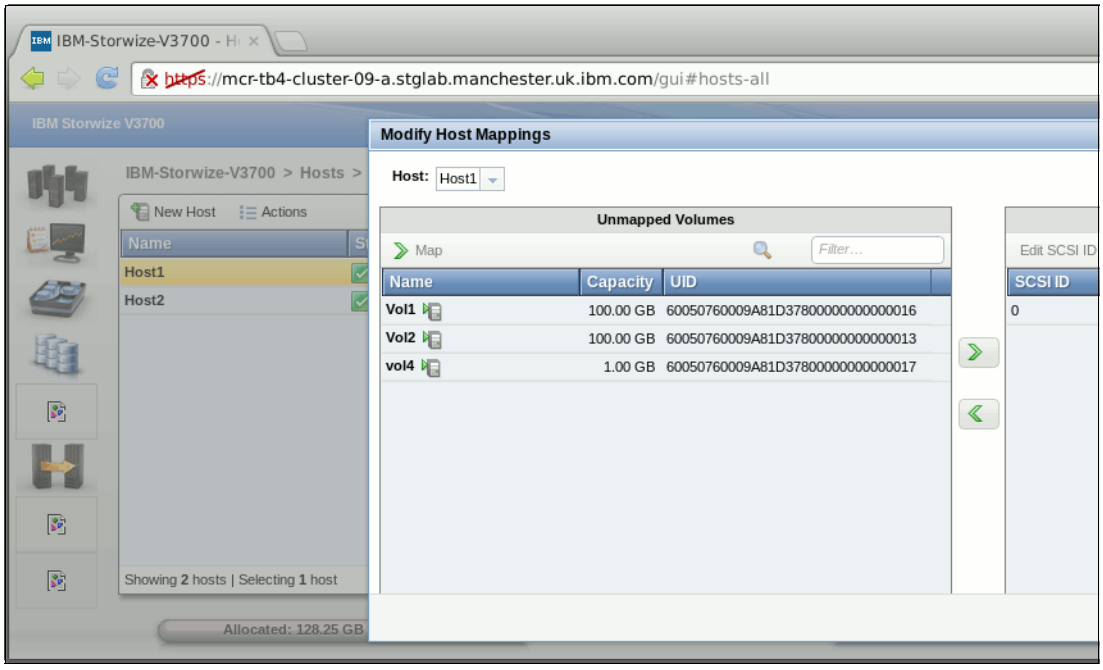

Figure 8-5 Modify Mappings window

To map a volume, highlight the volume that you want to map in the left pane, and select the upper arrow (pointing to the right), as shown in Figure 8-6.

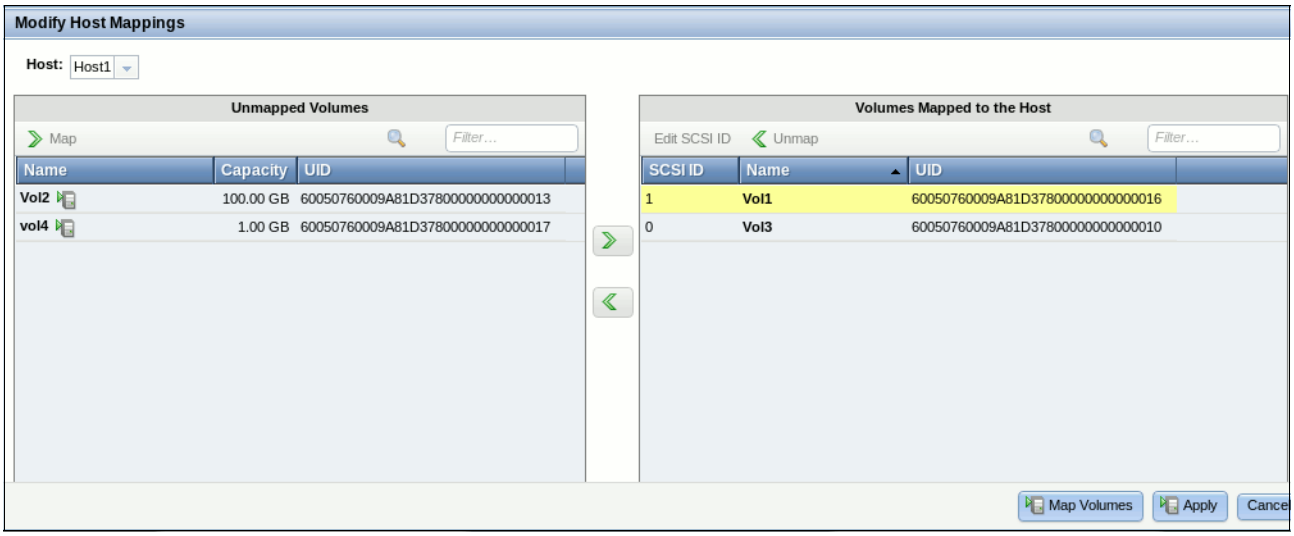

Figure 8-6 Modify Mappings: Add Volume

The changes are marked in yellow and now the Map Volumes and Apply options are enabled.

If you click **Map Volumes**, the changes are applied (as shown in Figure 8-7) and the window closes.

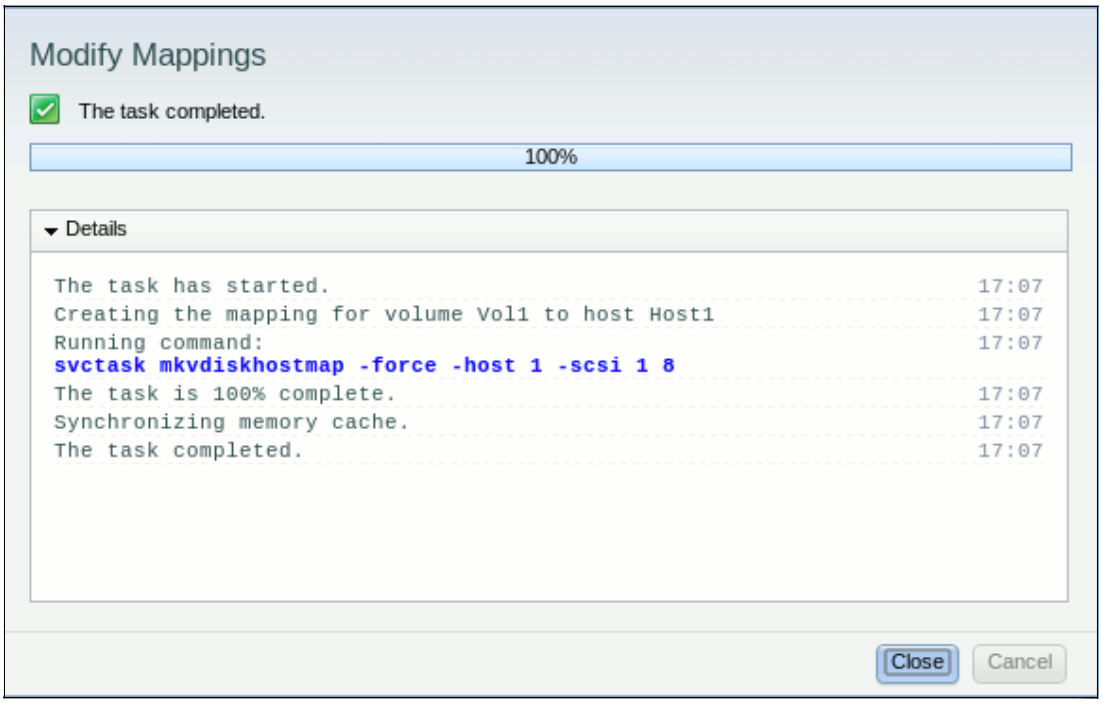

Figure 8-7 Modify Mappings output

If you click **Apply**, the changes are submitted to the system, but the window remains open for further changes, as shown in Figure 8-8.

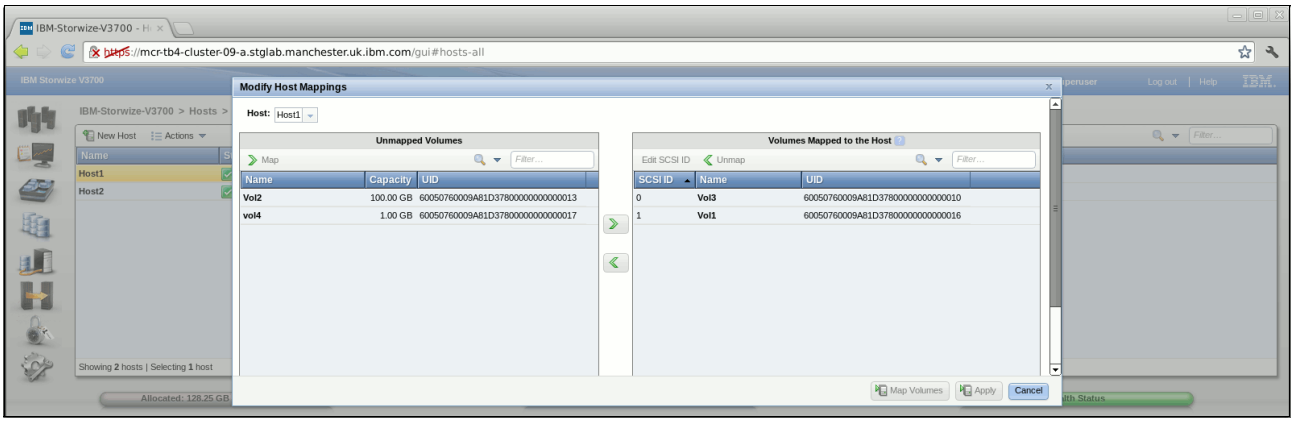

Figure 8-8 Modify Mappings: Applied Changes

You can now choose to modify another host by selecting it from the Hosts drop-down menu (as shown in Figure 8-9) or continue working with the one already selected, as shown in Figure 8-8.

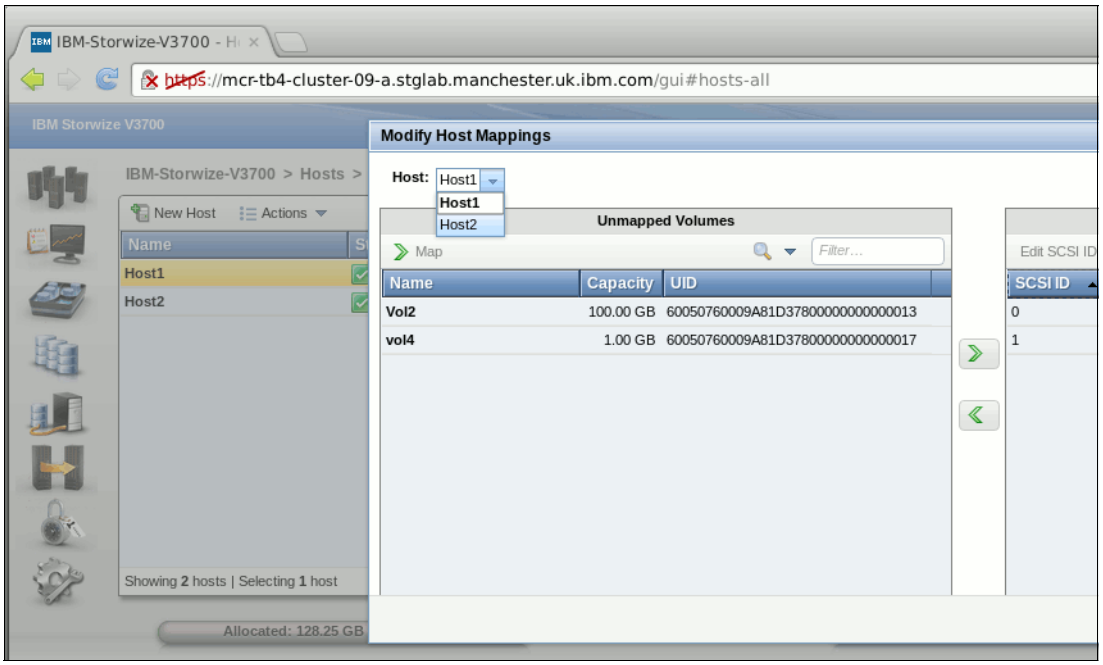

Figure 8-9 Modify another host

Highlight the volume that is to be modified again and click the right arrow button to move it to the right pane. The changes are shown in yellow. If you right-click the yellow volume, you can change the SCSI ID, which is used for the host mapping, as shown in Figure 8-10.

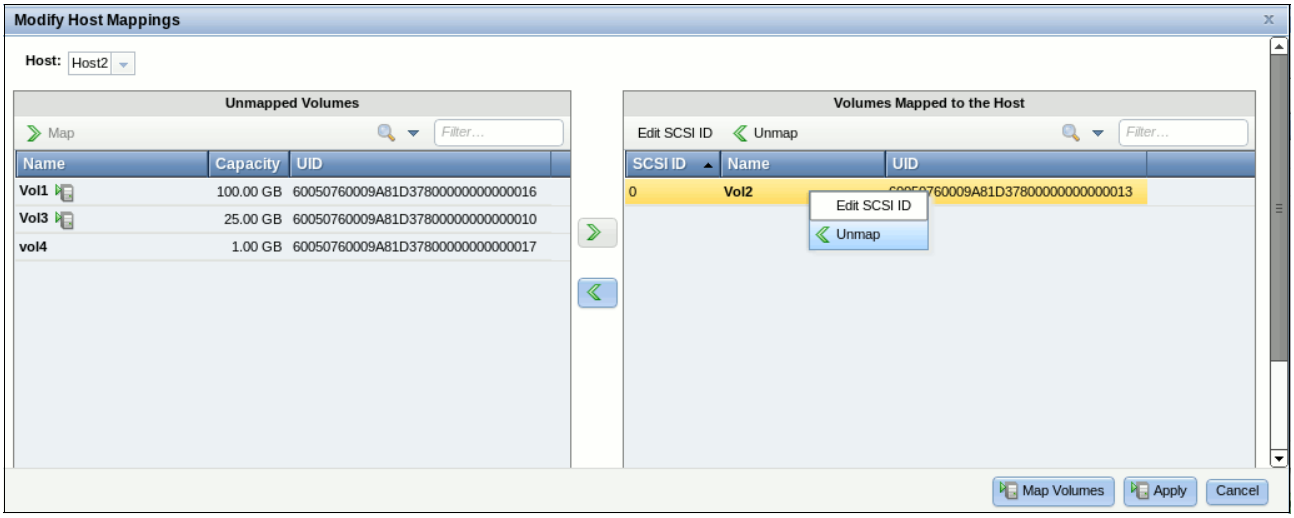

Figure 8-10 Edit SCSI ID

Click **Edit SCSI ID** and then click **OK** to change the ID, as shown in Figure 8-11.

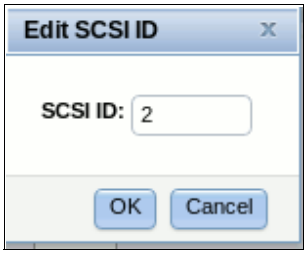

Figure 8-11 Enter a new SCSI ID

The changes are shown in the Modify Host Mappings window, as shown in Figure 8-12.

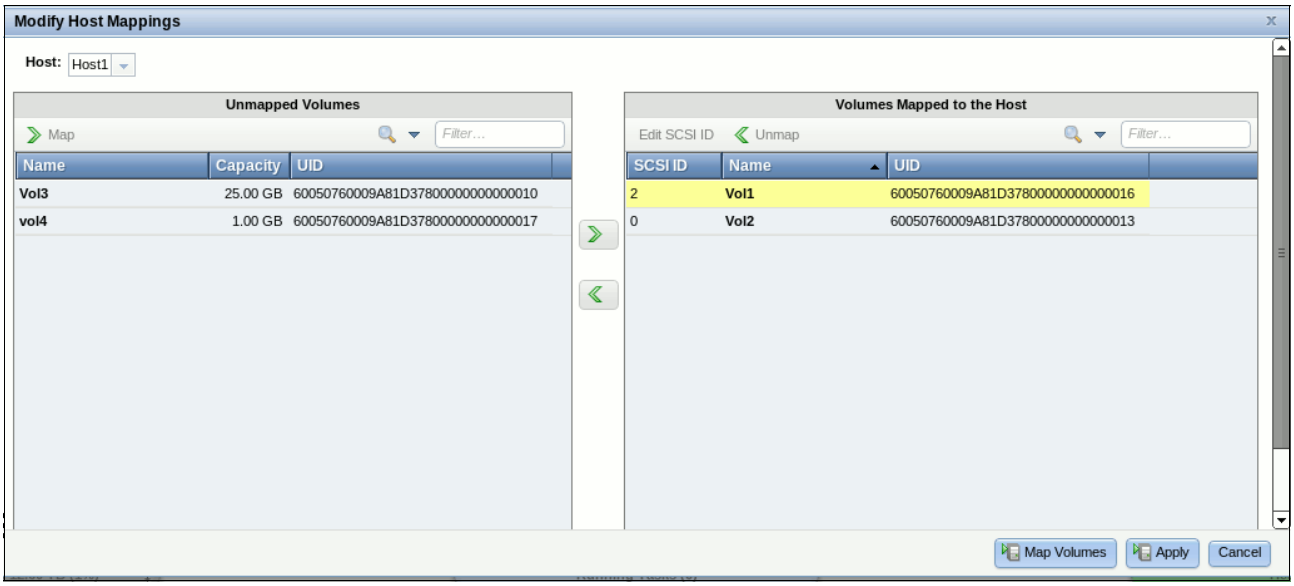

Figure 8-12 Modify Host Mappings window

Click **Apply** to submit the changes. The resulting output is shown in Figure 8-13.

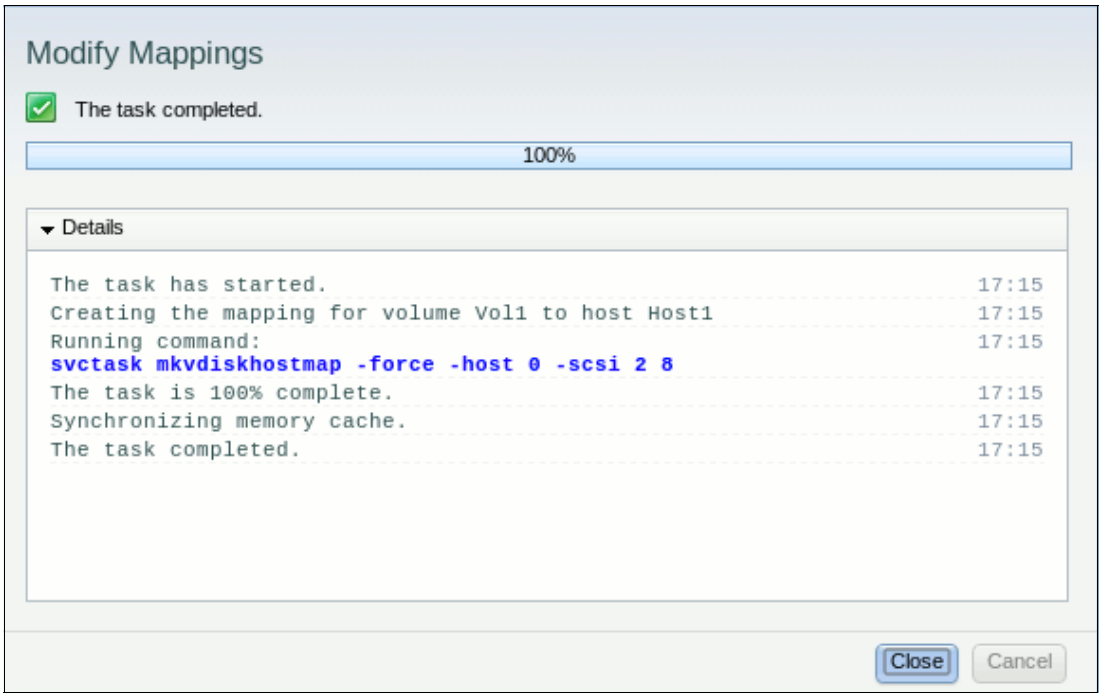

Figure 8-13 Modified mappings

If you want to remove a host mapping, the required steps are much the same, except that you select a volume in the right pane and click the left arrow button to remove the mapping, as shown in Figure 8-14.

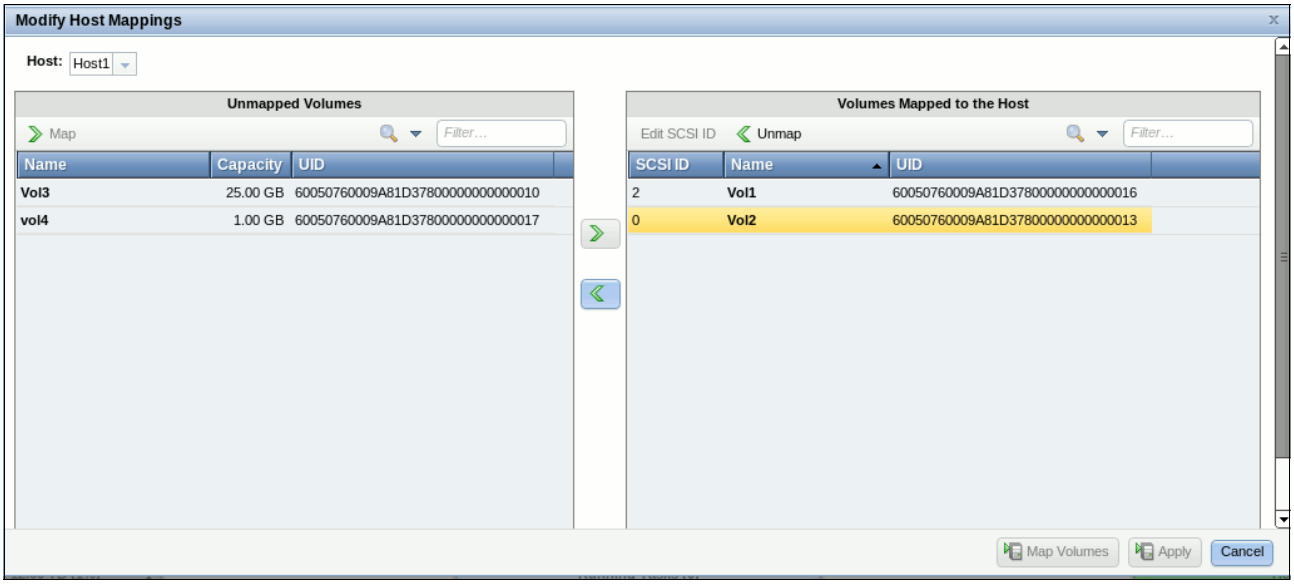

Figure 8-14 Modify Host Mappings: Remove

Figure 8-15 shows that the mapping is removed.

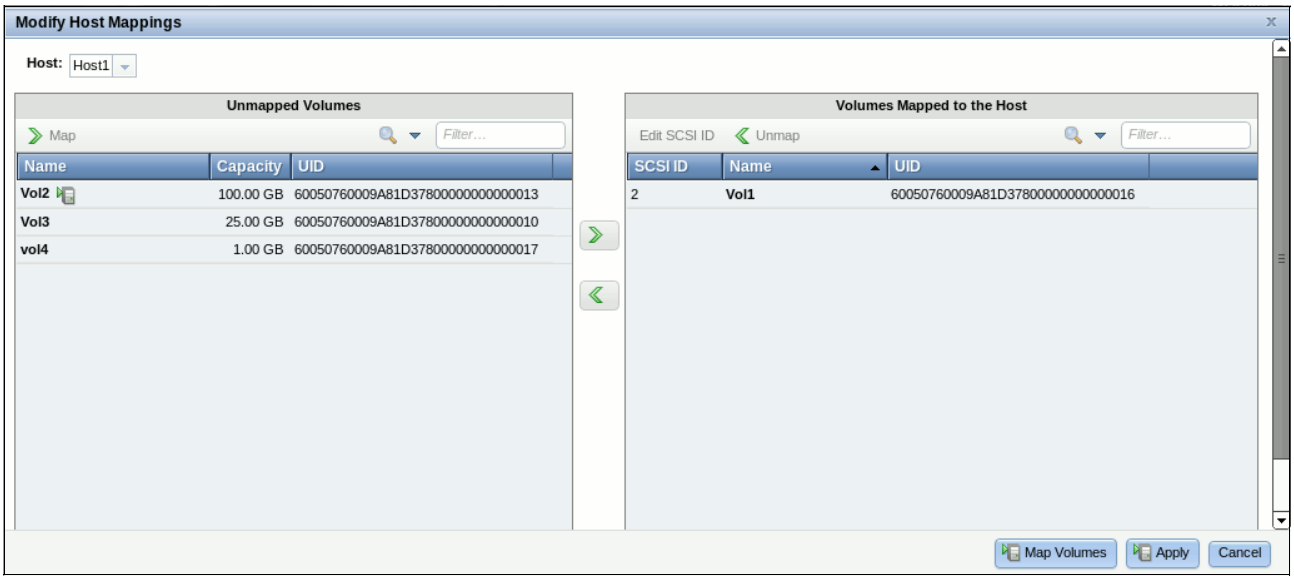

Figure 8-15 Mapping removed

Click **Apply** to submit the changes to the system, as shown in Figure 8-16.

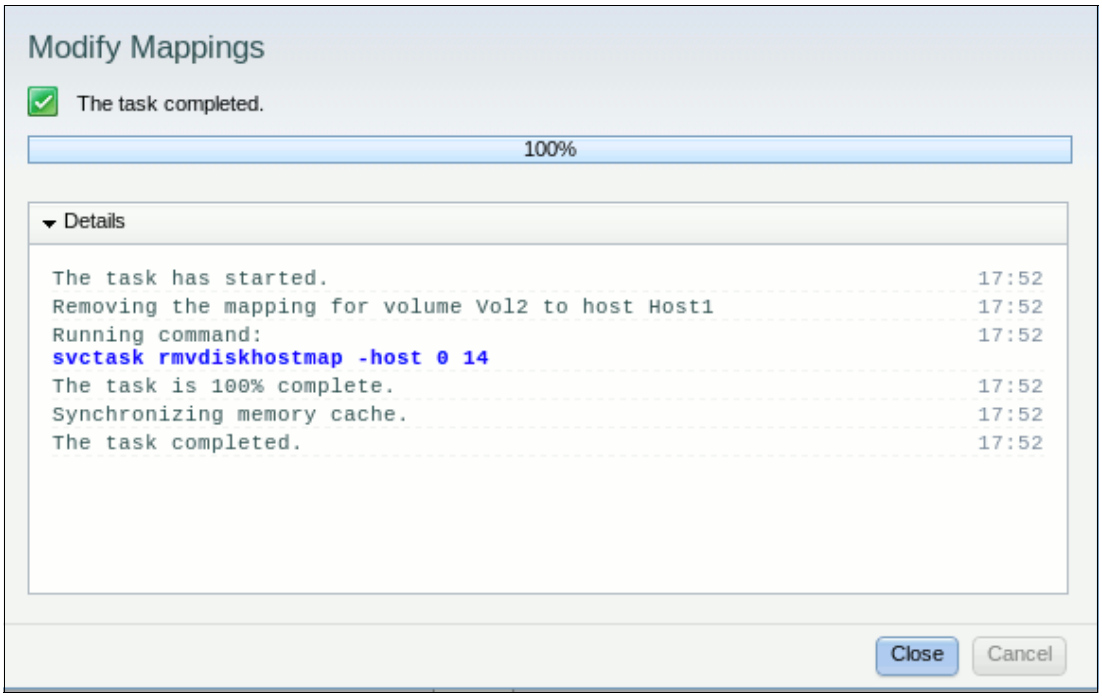

Figure 8-16 Modify Mappings: Removal complete

After you are done with all host mapping modifications, click **Close** to return to the Modify Mappings window, as shown in Figure 8-5 on page 320.

8.1.2 Unmapping all volumes from a host

If you want to remove access to all volumes on your IBM Storwize V3500 from a host, you highlight the host and click **Unmap all volumes**. All volume mappings are removed from the host, as shown in Figure 8-17.

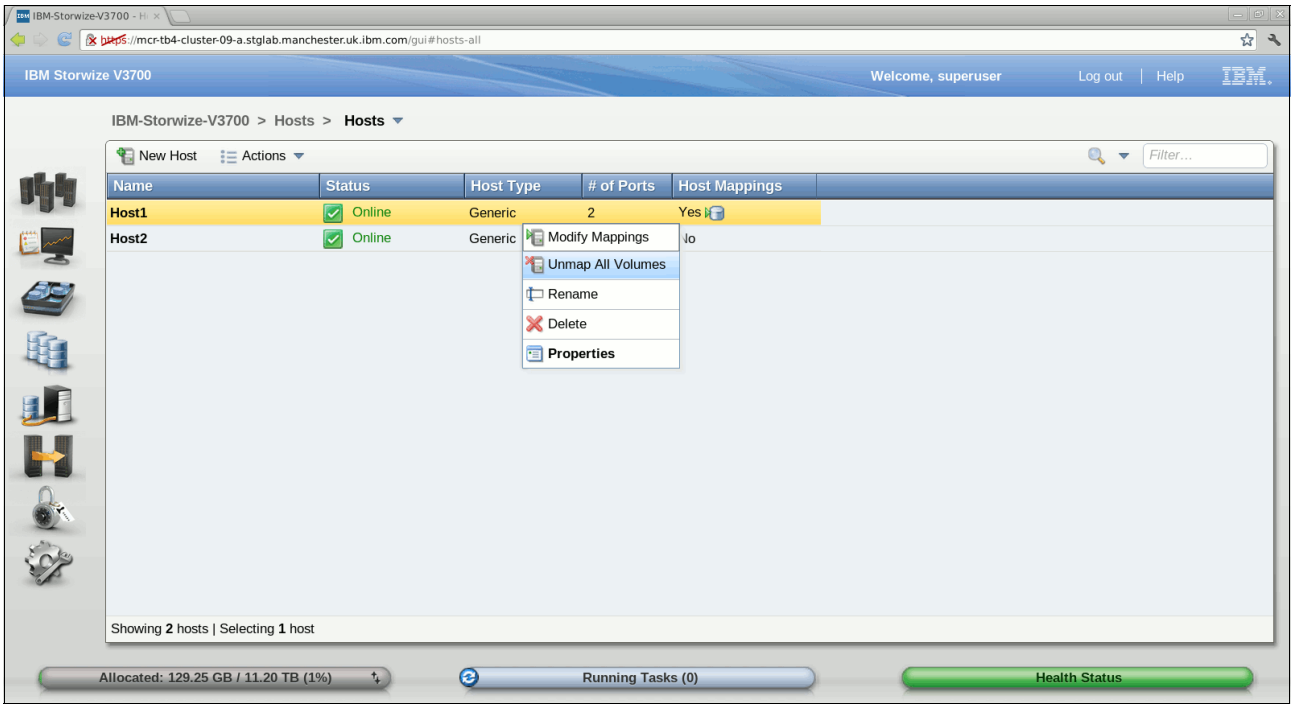

Figure 8-17 Unmap all volumes menu

You are prompted to confirm the number of mappings you want to remove. Enter the number and click **Unmap**, as shown in Figure 8-18. In our example, we remove two mappings.

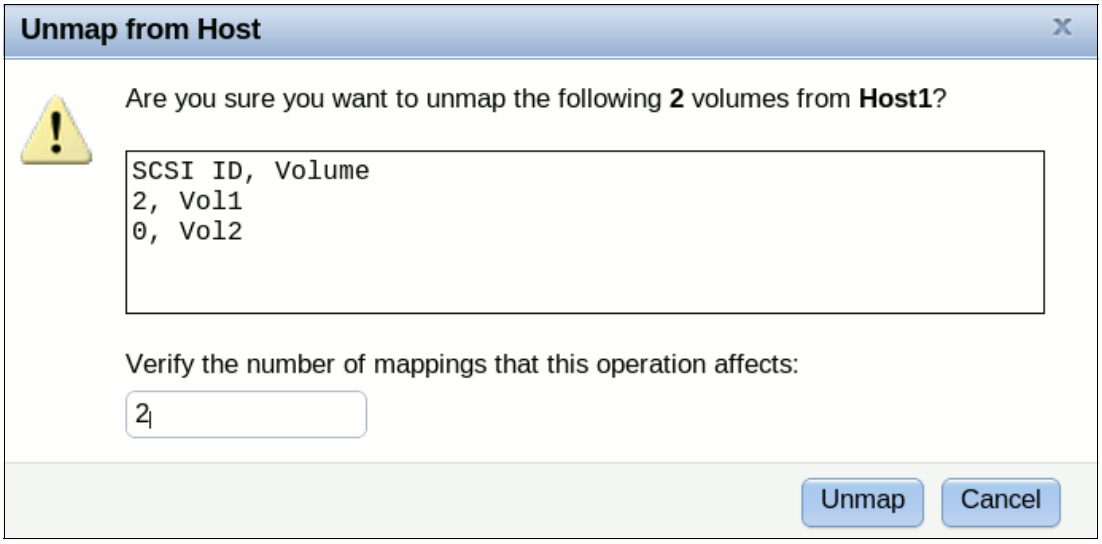

Figure 8-18 Enter the number of mappings to be removed

Unmapping: If you click **Unmap**, all access for this host to volumes that are controlled by IBM Storwize V3700 system is removed. Make sure that you run the required procedures in your host operating system before the volumes are removed from your host.

The changes are applied to the system, as shown in Figure 8-19. Click **Close** after you review the output.

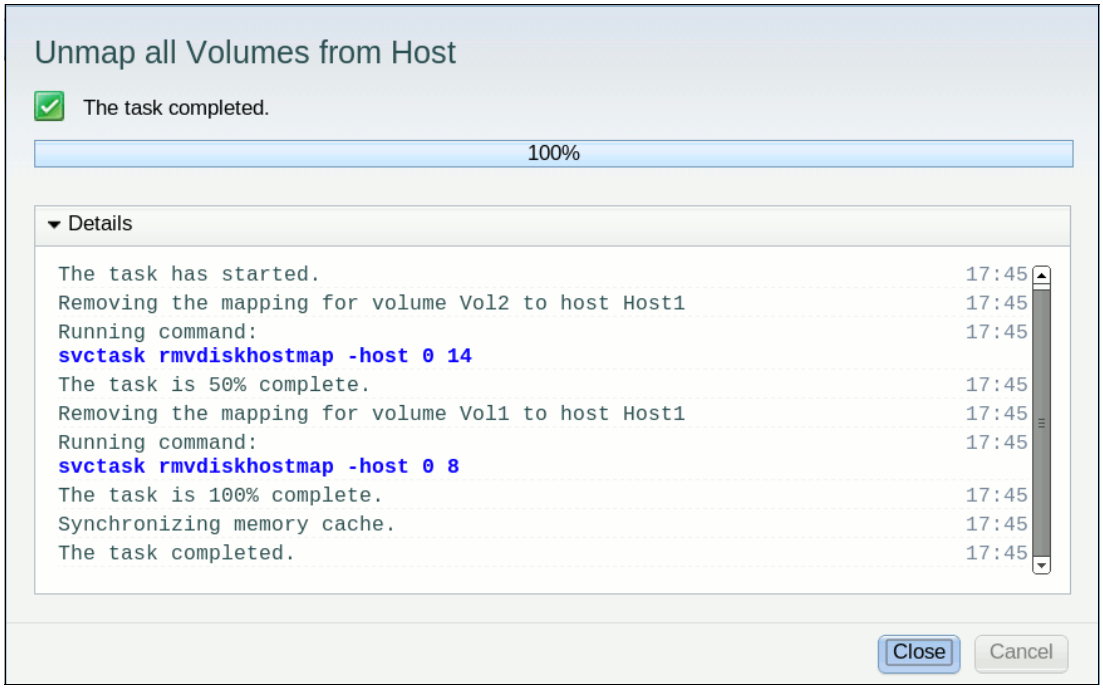

Figure 8-19 Unmap all the volumes from a host: Changes applied

Figure 8-20 shows that the selected host no longer has any host mappings.

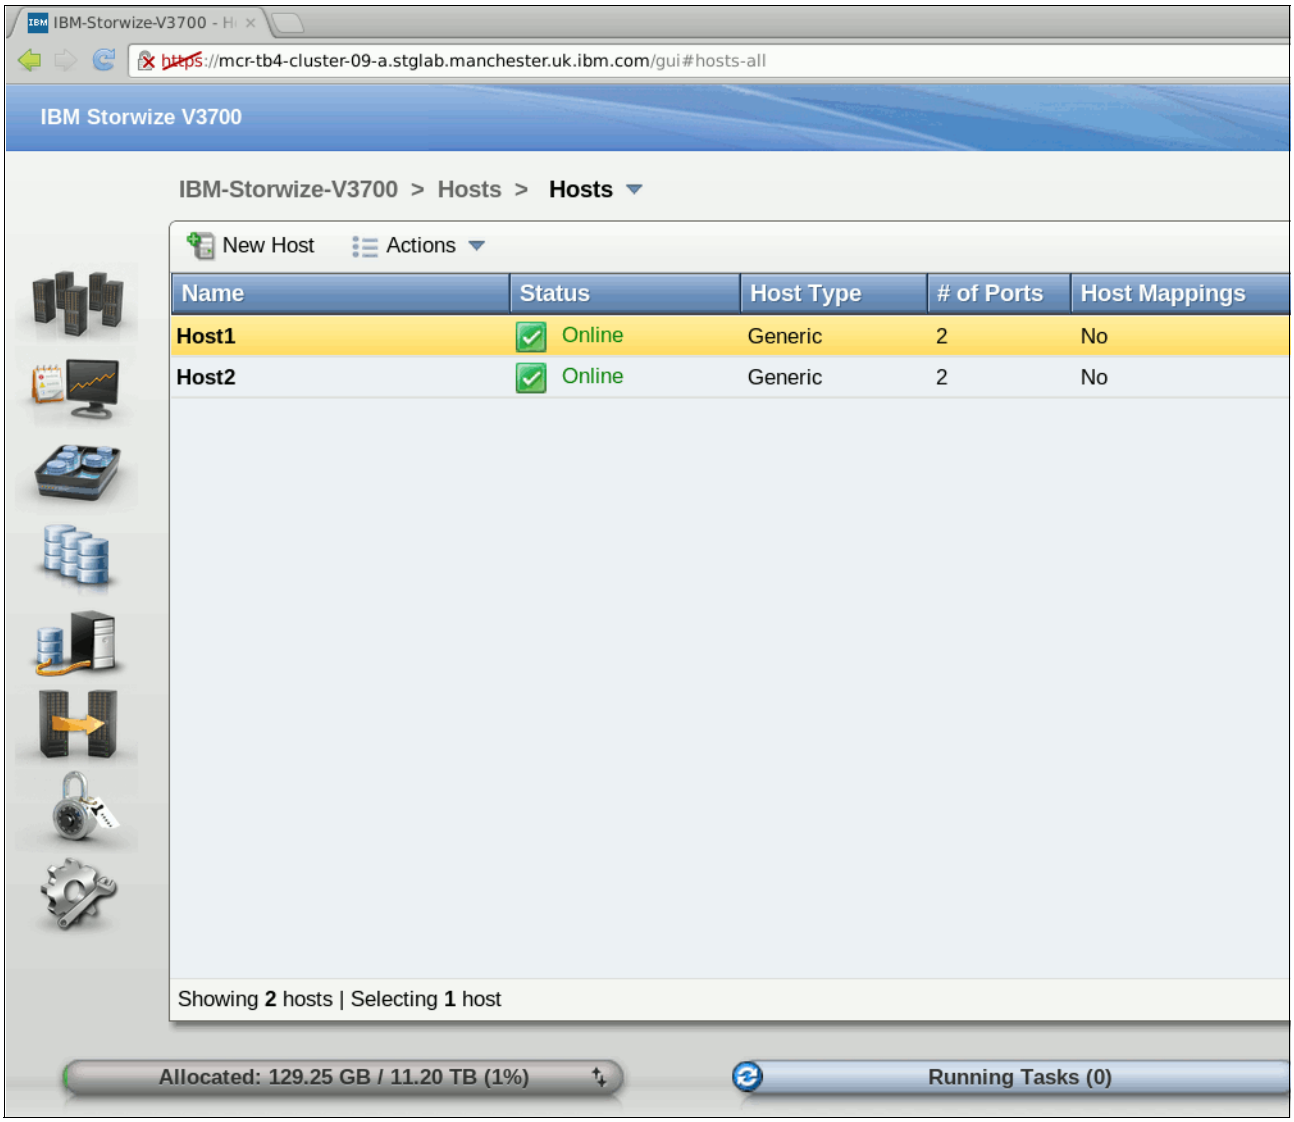

Figure 8-20 Hosts window: No host mappings

8.1.3 Renaming

To rename a host, highlight the host and click **Rename**, as shown in Figure 8-21.

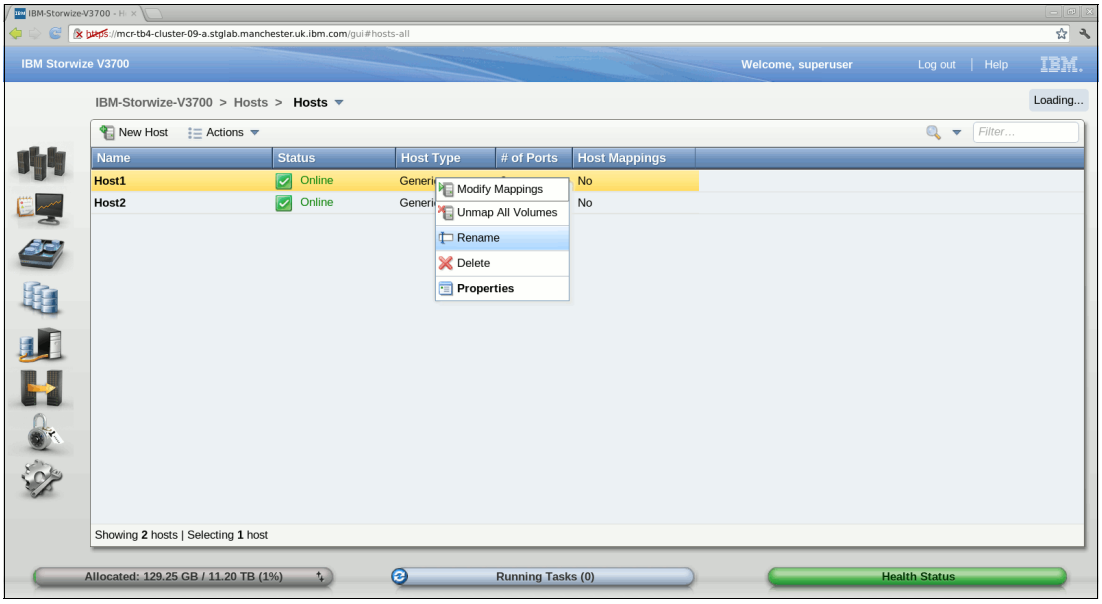

Figure 8-21 Rename a host

Enter a new name and click **Rename**, as shown in Figure 8-22. If you click **Reset**, your changes are not saved and the host retains its original name.

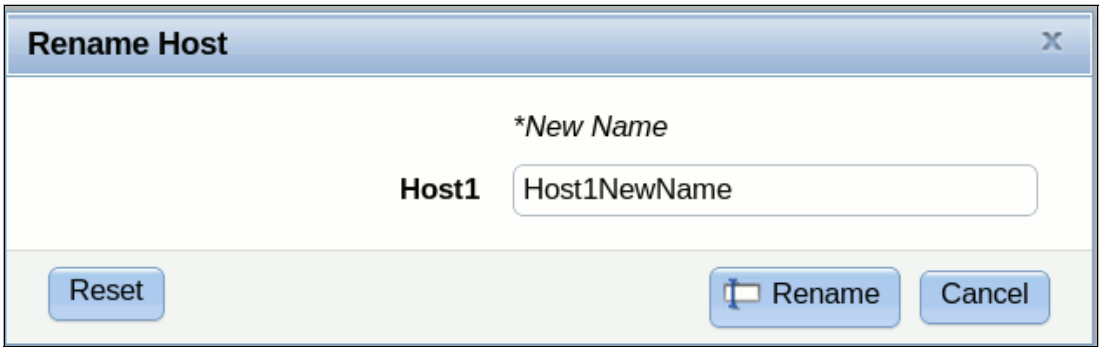

Figure 8-22 Rename Host window

After the changes are applied to the system, click **Close**, as shown in Figure 8-23.

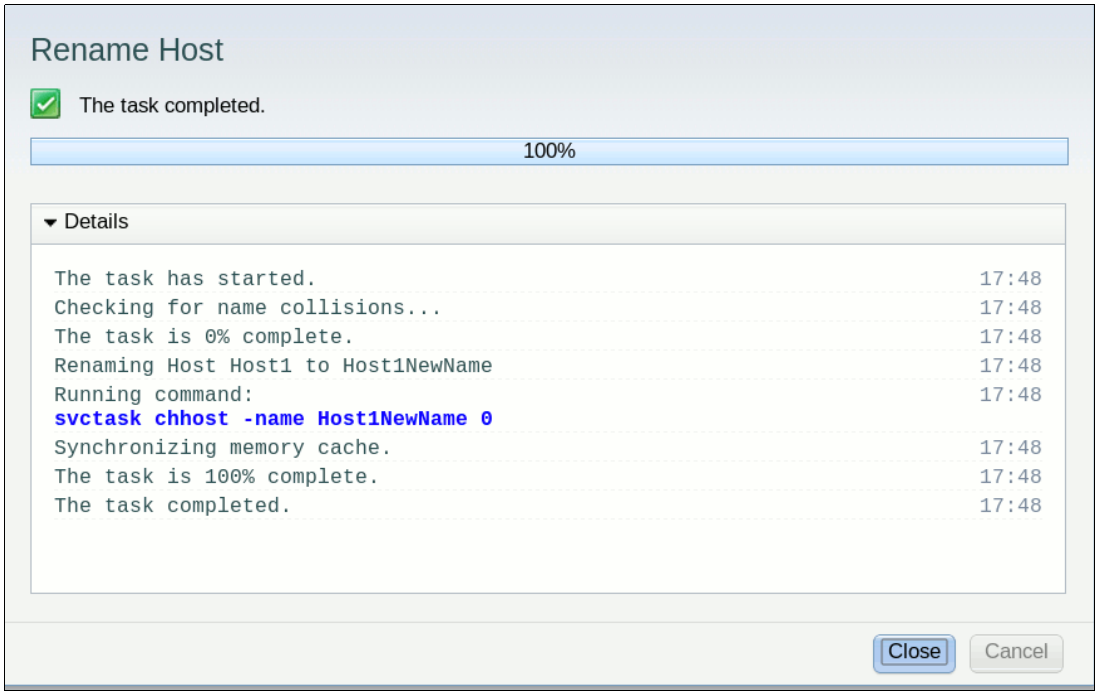

Figure 8-23 Rename Host: Applied changes

8.1.4 Deleting a host

To delete a host, highlight it and select **Delete**, as shown in Figure 8-24.

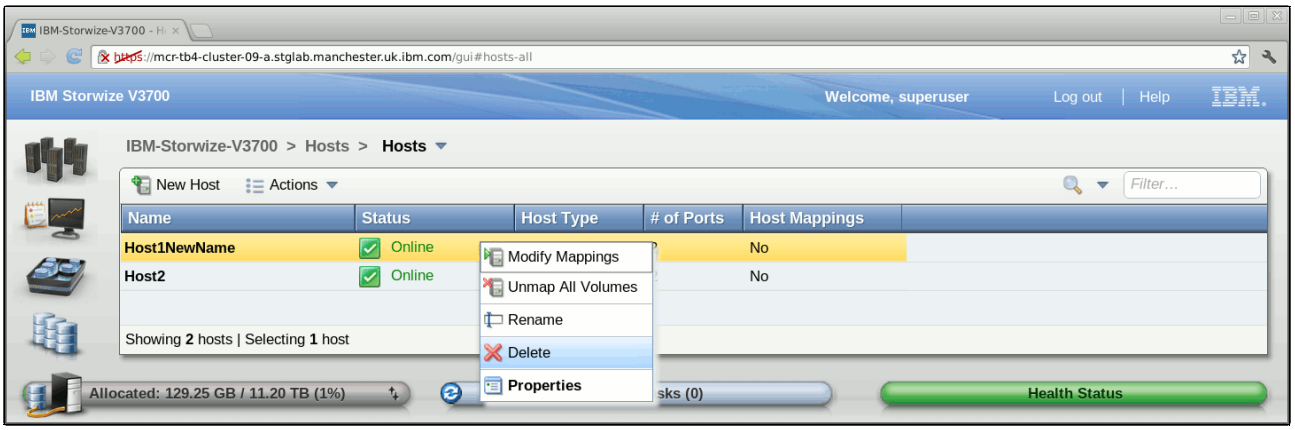

Figure 8-24 Delete a host

Confirm the number of hosts that you want to delete and click **Delete**, as shown in Figure 8-25.

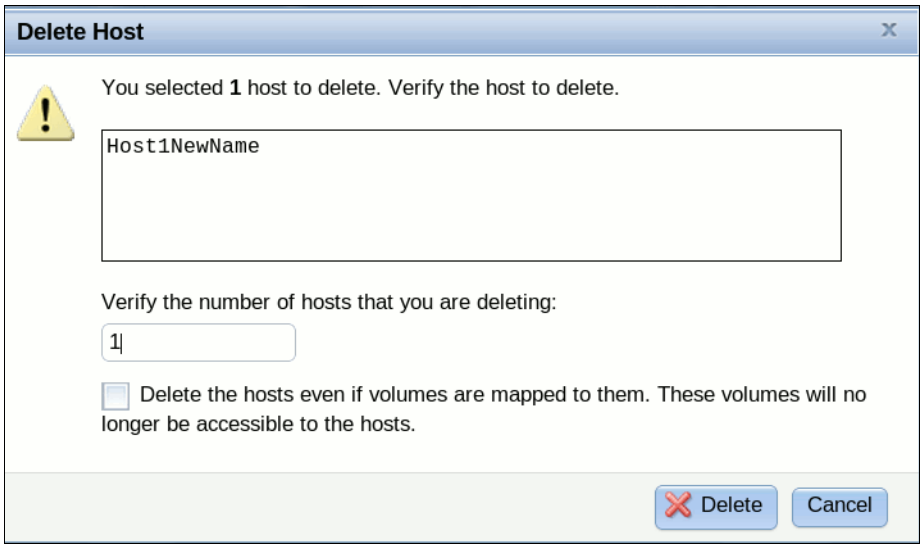

Figure 8-25 Confirm the deletion of the host

If you want to delete a host to which volumes are mapped, you must force the deletion by selecting the option in the lower part of the window. If you select this option, the host is deleted and cannot access the system.

After the task is complete, click **Close** (as shown in Figure 8-26) to return to the mappings window.

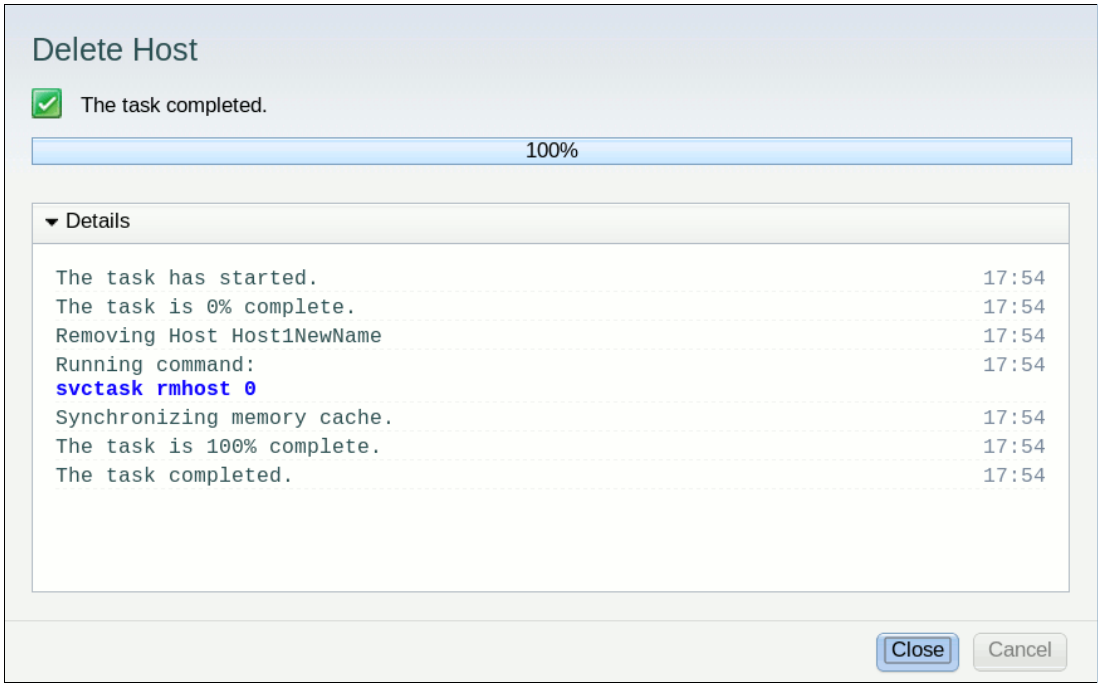

Figure 8-26 Delete host complete

8.1.5 Host properties

Highlight a host and select **Properties** (as shown in Figure 8-27) to open the Host Details window, as shown in Figure 8-28 on page 334). The Host Details window gives you an overview of your host properties. There are three tabs available: Overview, Mapped Volumes, and Port Definitions.

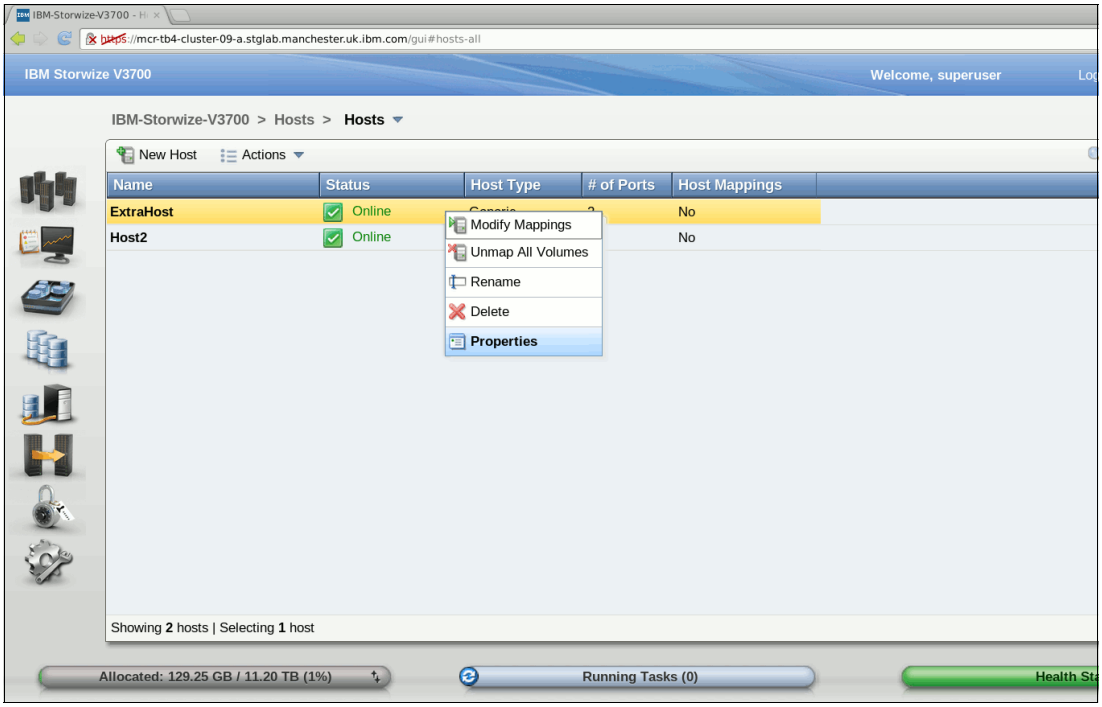

Figure 8-27 Host Properties

Figure 8-28 shows the Host Properties Overview window.

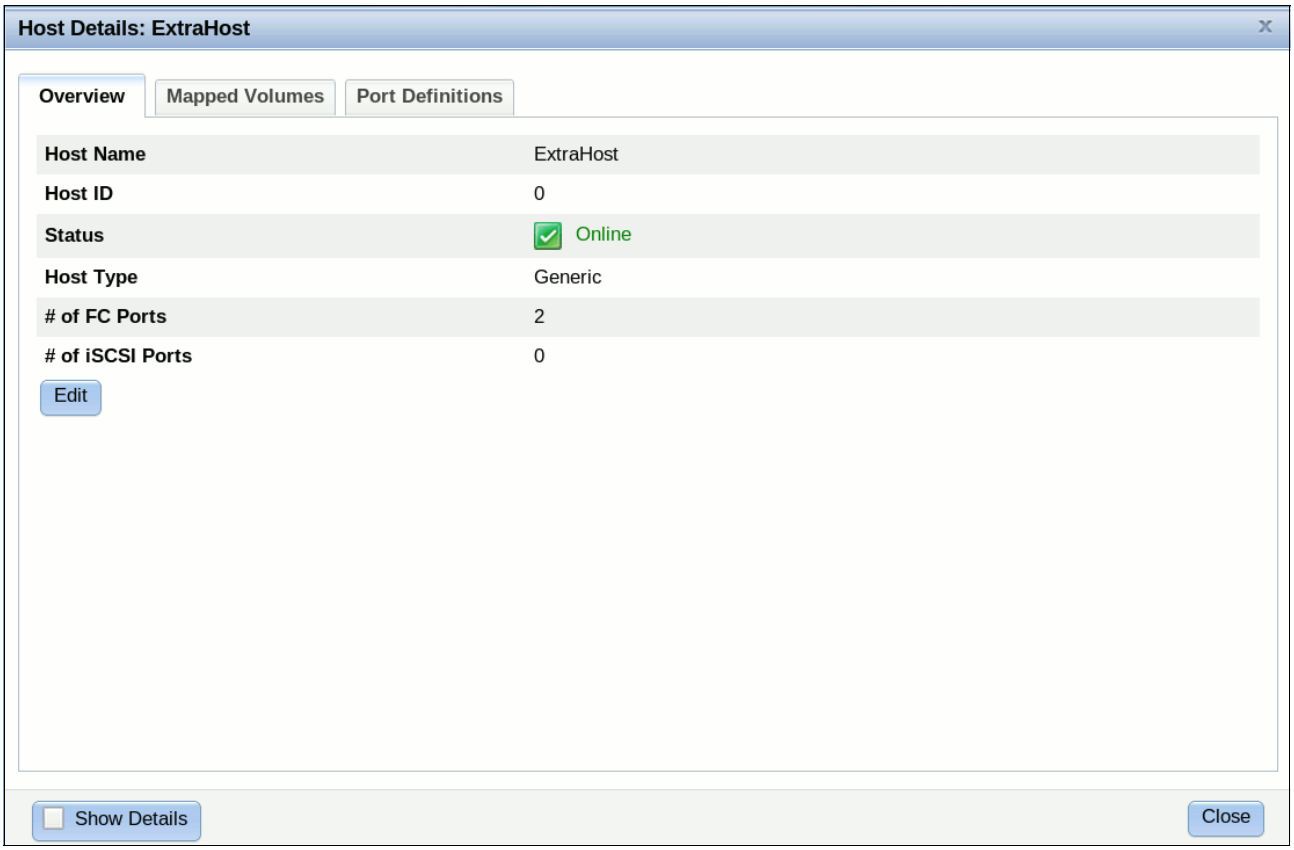

Figure 8-28 Host Properties overview

The Overview tab is shown in Figure 8-29. Select the **Show Details** option to see more information about the host.

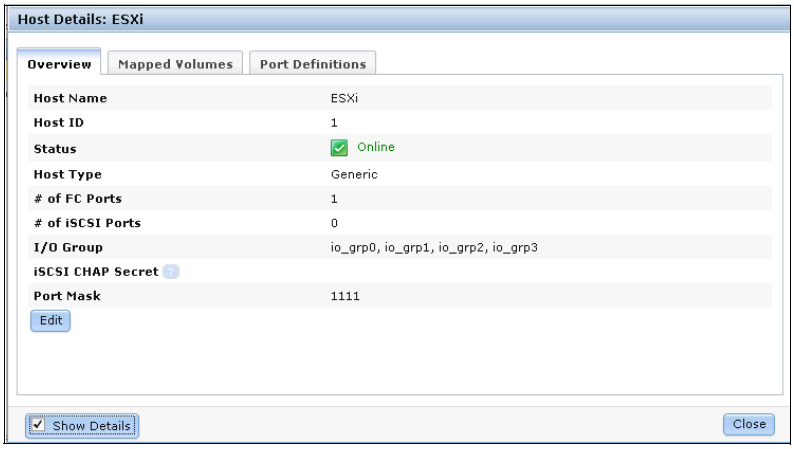

Figure 8-29 Host Properties: Show details

To change the host properties, click **Edit**, as shown in Figure 8-30.

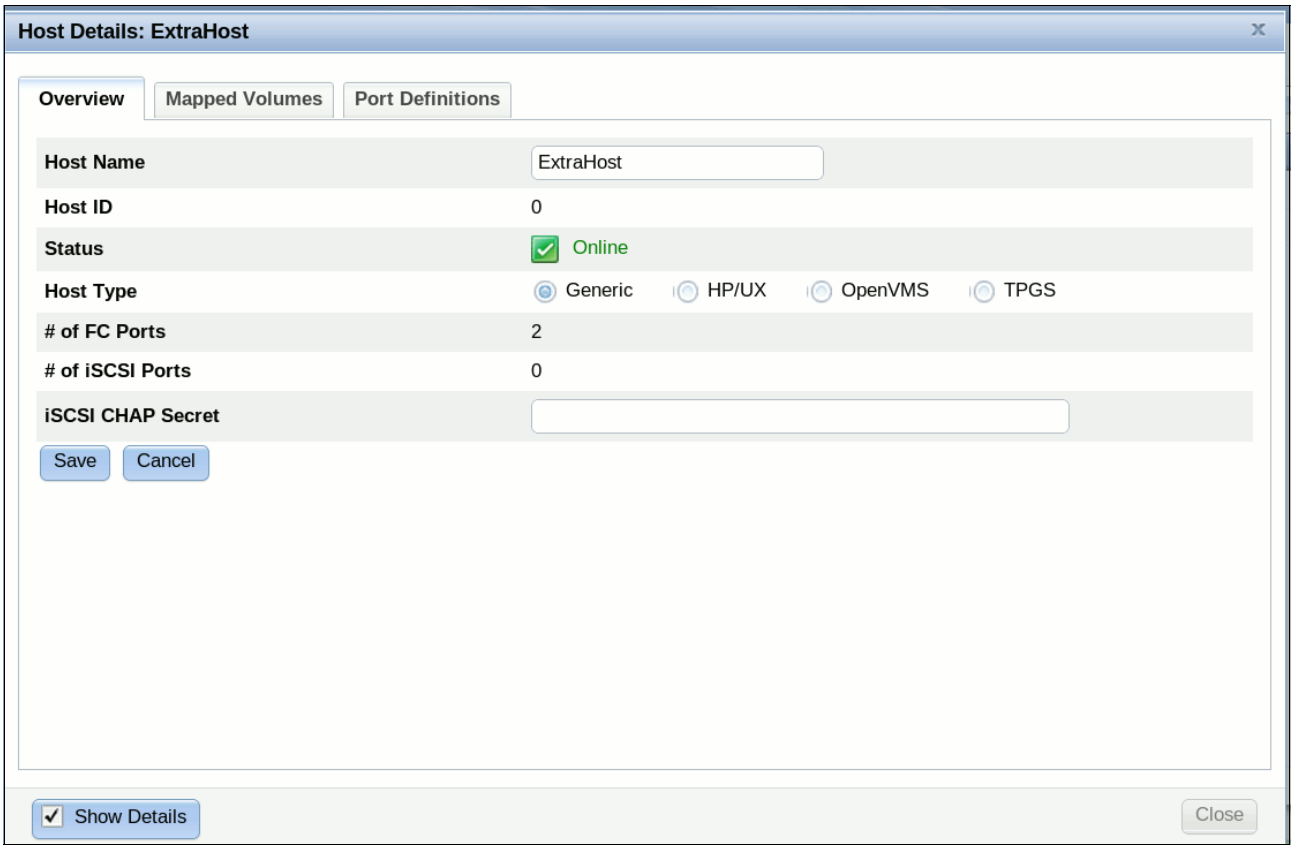

Figure 8-30 Edit Host properties

The following changes can be made:

- Host Name: Change the host name.
- Host Type: Change this setting if you are attaching HP/UX, OpenVMS, or TPGS hosts.
- iSCSI CHAP Secret: Enter or change the iSCSI CHAP secret for this host.
- Port Mask: To change the port mask, disable ports for this host by clearing the check box.

Port masks: You can use port masks to control the node target ports that a host can access. This control can be useful to limit the number of logins with mapped volumes that are visible to a host multipathing driver, instead of changing the SAN zoning. The best practice is to limit each host to four paths.

Make the changes, if required, and click **Save** to apply them, as shown in Figure 8-30 and to close the editing window.

The Mapped Volume tab gives you an overview of which volumes are mapped with which SCSI ID and UID to this host, as shown in Figure 8-31. Selecting the Show Details option does not show any other or different information.

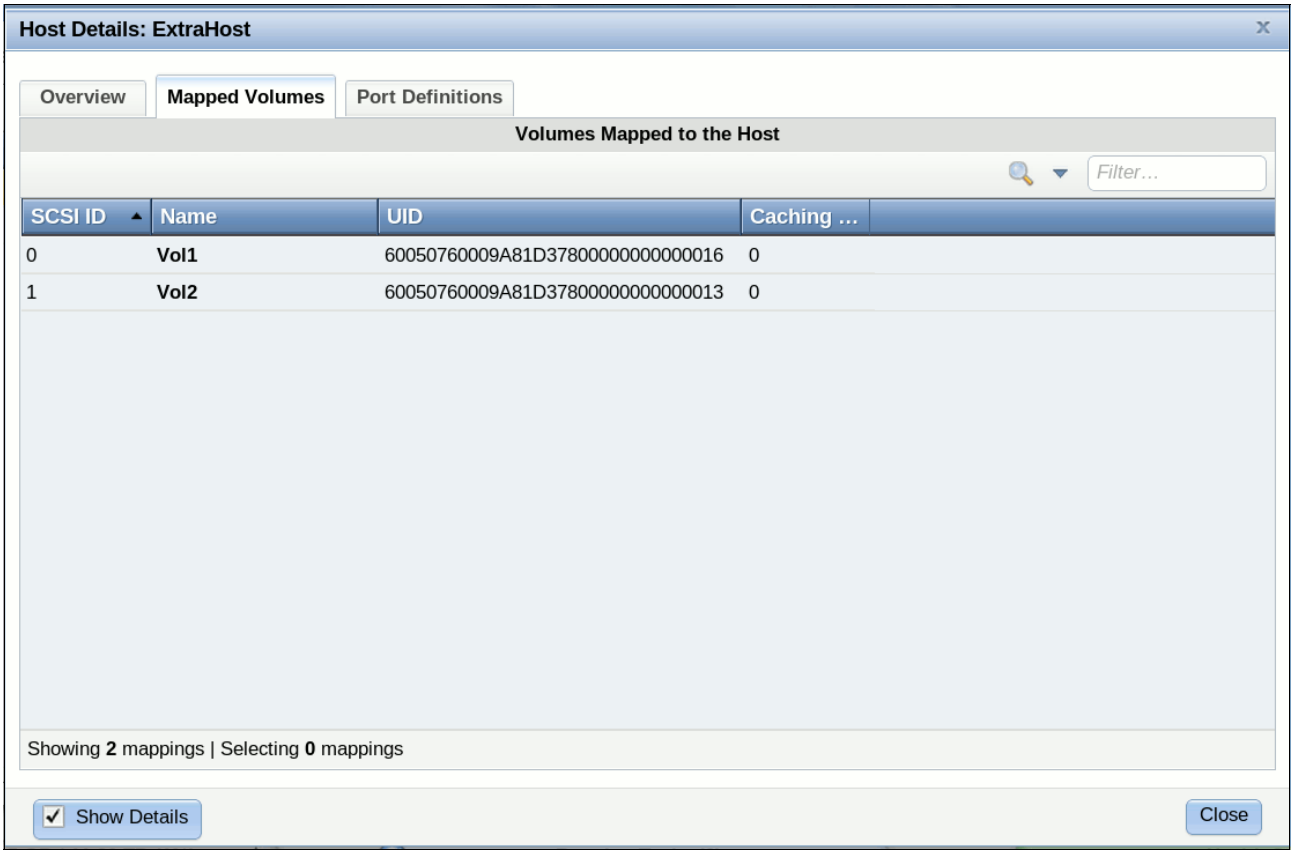

Figure 8-31 Mapped Volumes tab

The Port Definitions tab shows the configured ports of the host and provides status information about them, as shown in Figure 8-32.

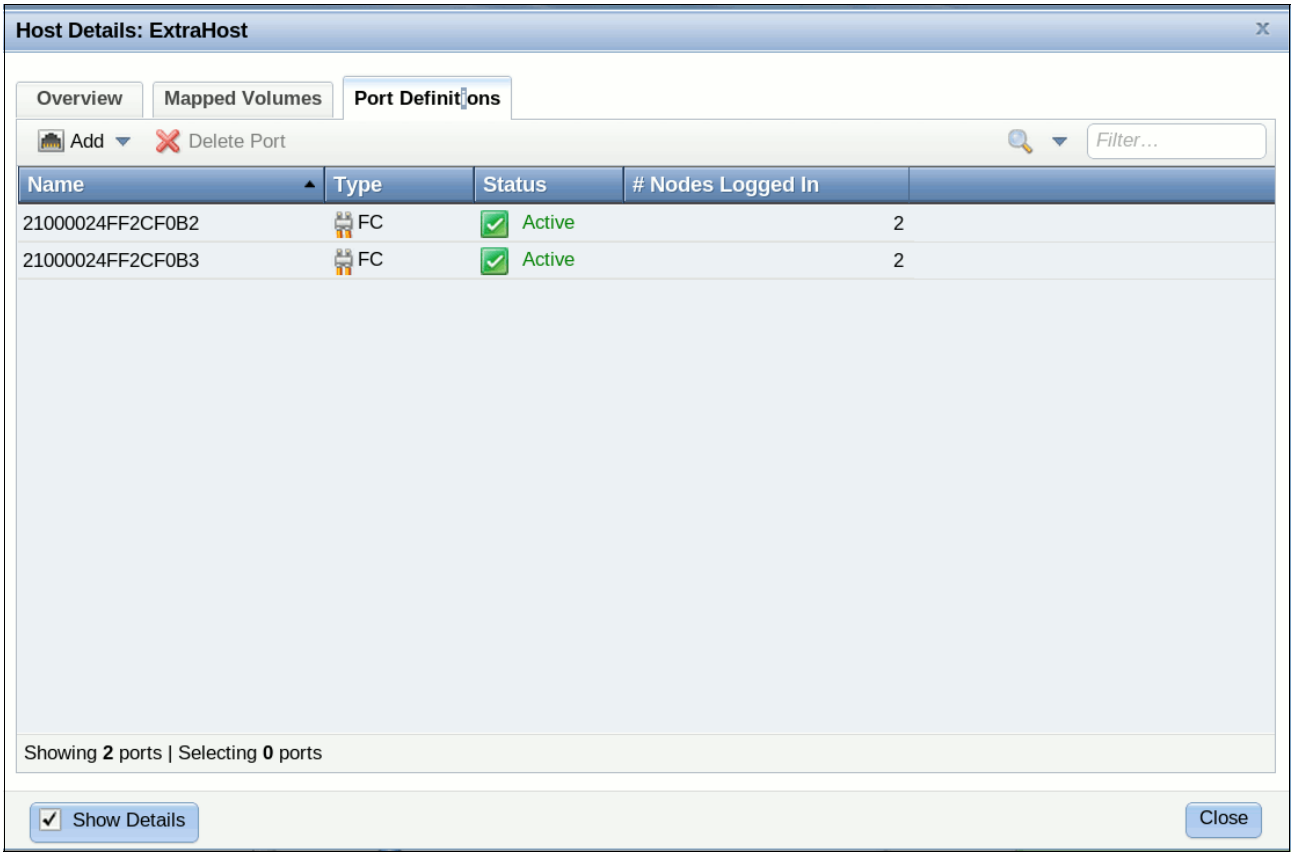

Figure 8-32 Port Definitions tab

This window offers you the option to start Add and Delete Port actions, as described in 8.2, “Adding and deleting host ports”.

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

8.2 Adding and deleting host ports

To configure host ports, click **Ports by Host** (as shown in Figure 8-1 on page 318) to open the associated window, as shown in Figure 8-33.

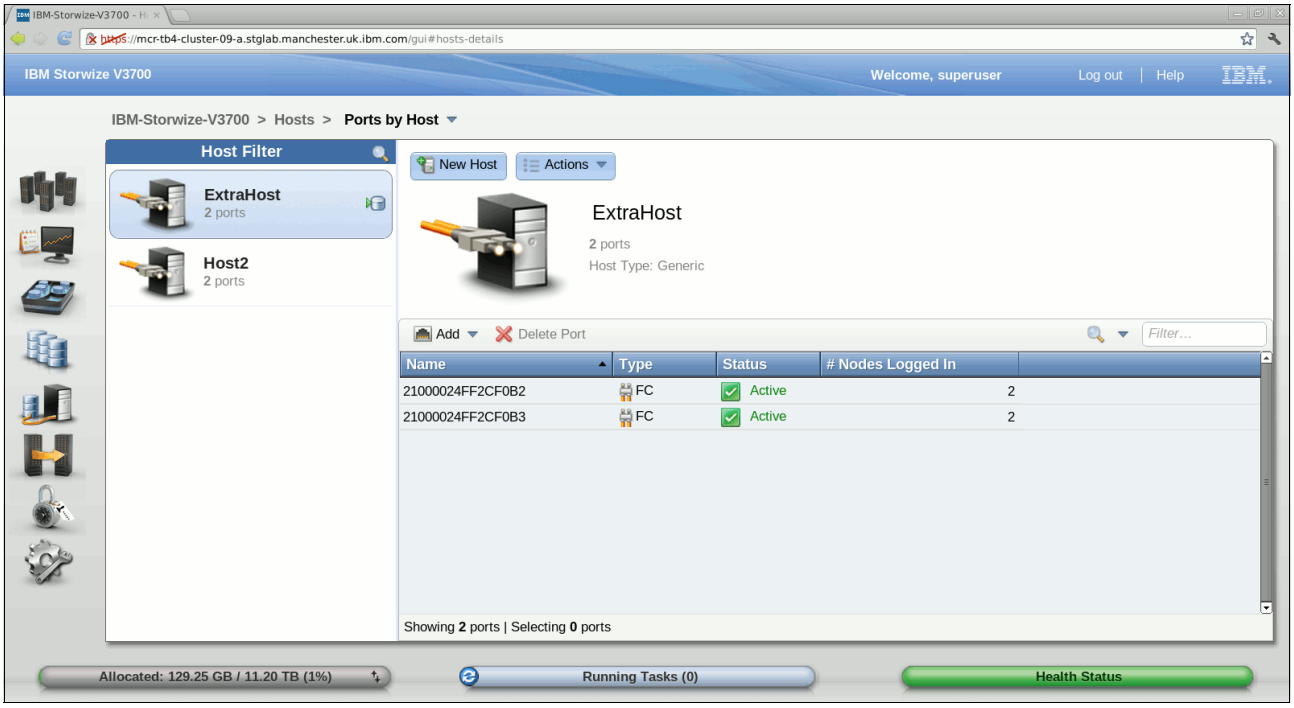

Figure 8-33 Ports By Host

Hosts are listed in the left pane of the window. The function icons show an orange cable for a Fibre Channel host and a blue cable for an iSCSI one. The properties of the highlighted host are shown in the right pane. If you click **New Host**, the wizard that is described in 4.3, “Creating hosts by using the GUI” on page 159 starts.

If you click the Action drop-down menu (as shown in Figure 8-34) the tasks that are described in 8.1.1, “Modifying Mappings menu” on page 320 can be started from this location.

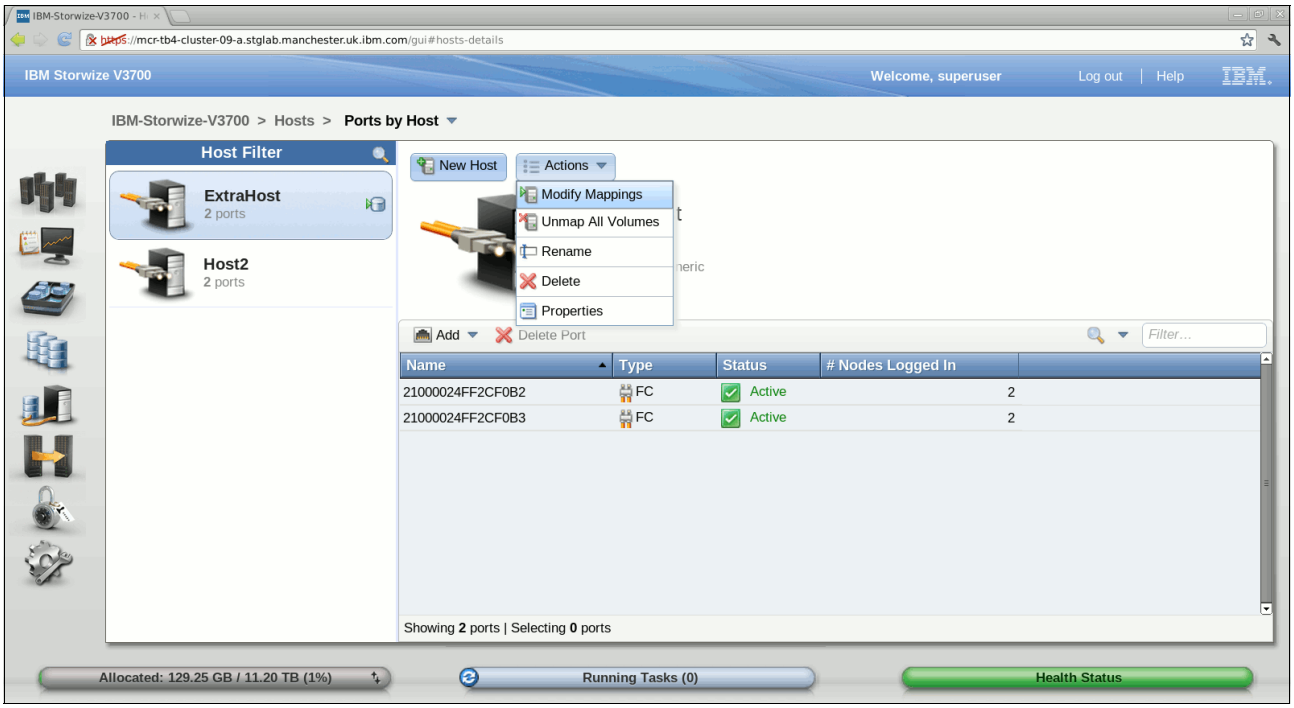

Figure 8-34 Ports By Host Actions

8.2.1 Adding a host port

To add a host port, highlight the host and click **Add** (as shown in Figure 8-33 on page 338), and choose a Fibre Channel or an iSCSI port, as shown in Figure 8-35.

Figure 8-35 shows the Add Host Ports menu options.

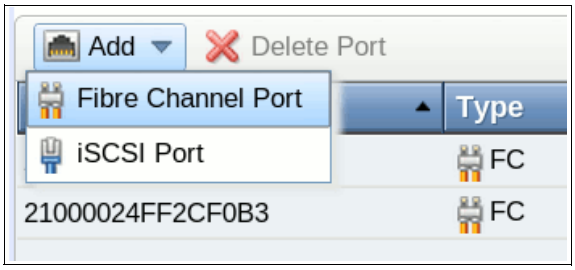

Figure 8-35 Add Host Ports

8.2.2 Adding a Fibre Channel port

In Figure 8-35, click Fibre Channel Port and the Add Fibre Channel Ports window opens, as shown in Figure 8-36.

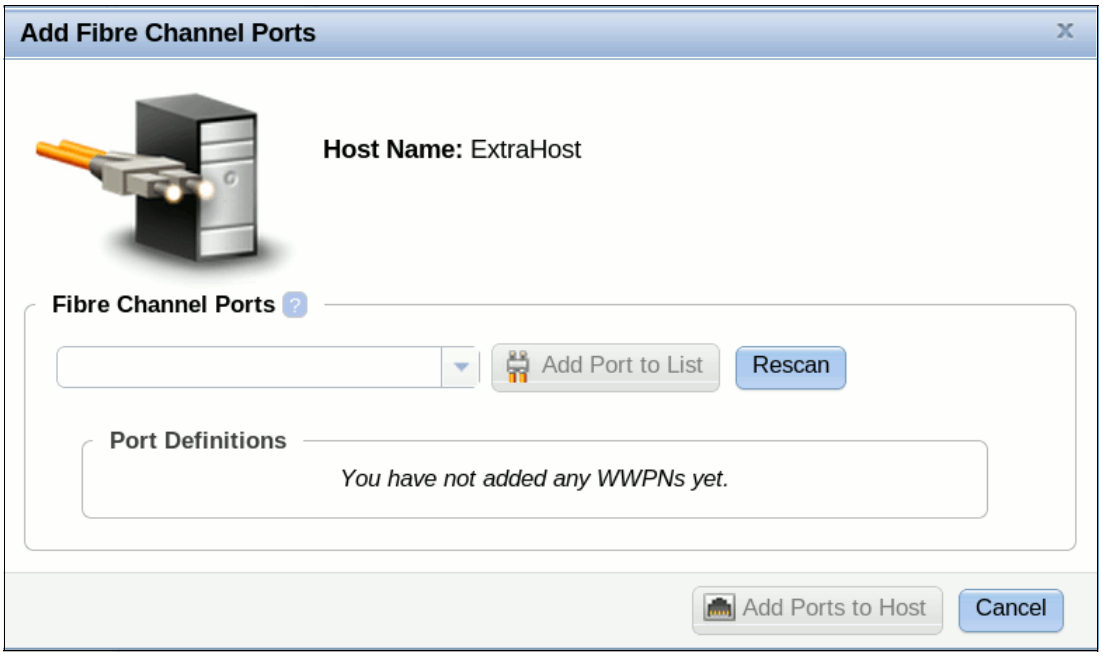

Figure 8-36 Add Fibre Channel Ports

If you click the drop-down menu, you see a list of all known Fibre Channel host ports, as shown in Figure 8-37. If the WWPN of your host is not available in the drop-down menu, check your SAN zoning and rescan the SAN from the host. Click **Rescan** and the new port should now be available in the drop-down menu.

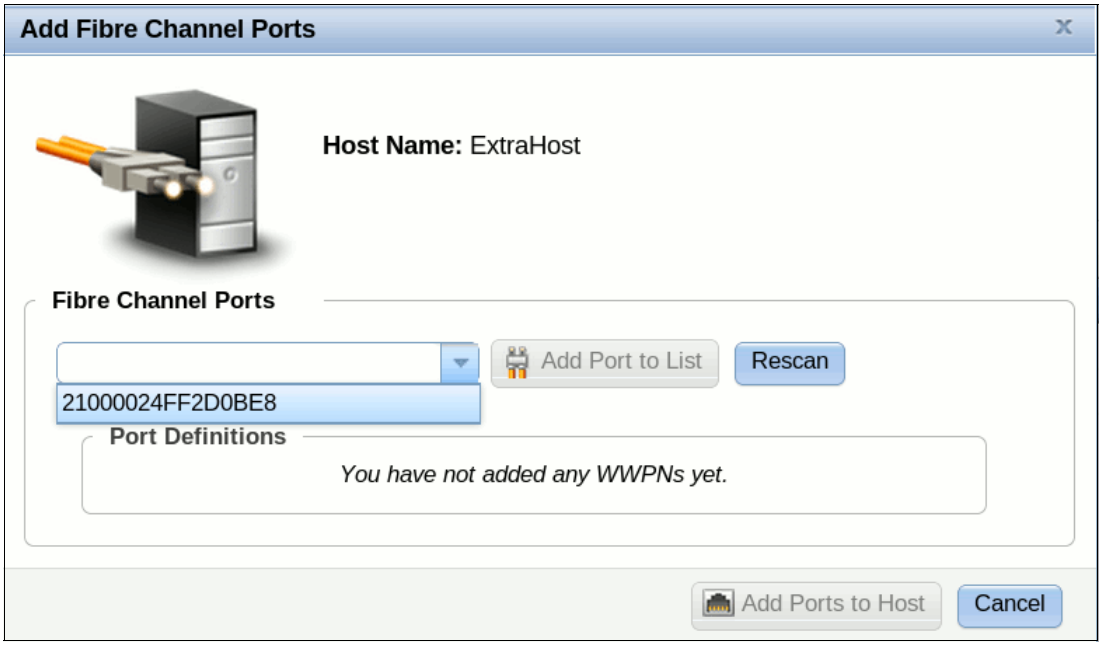

Figure 8-37 Known WWPNs

Select the WWPN to add and click Add Port to List, as shown in Figure 8-38.

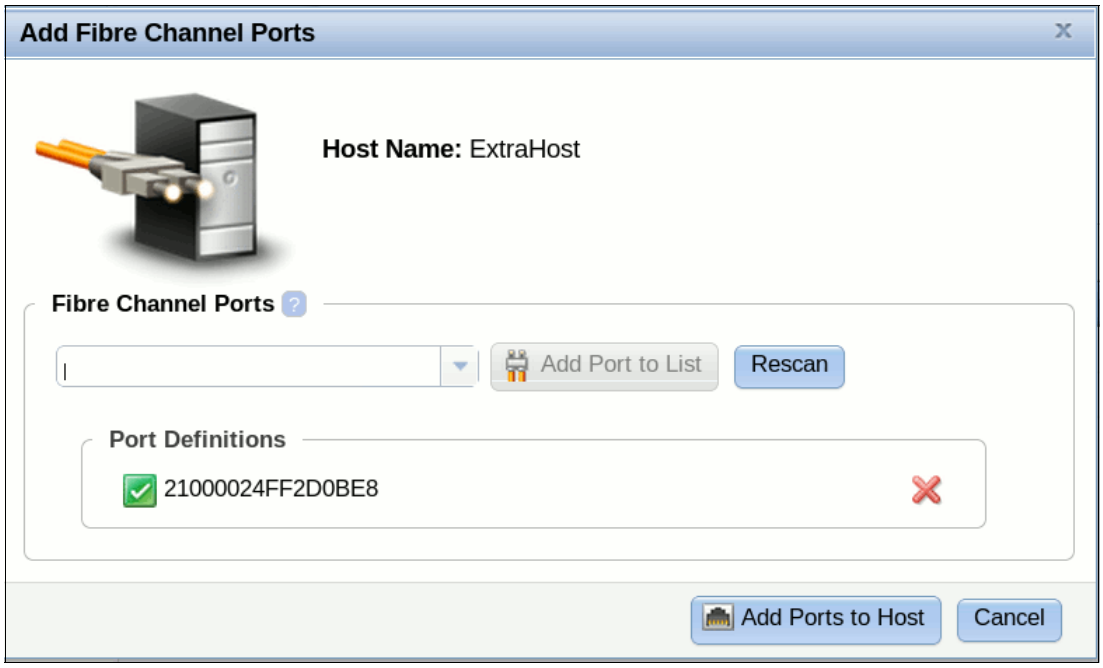

Figure 8-38 Port added to list

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 box (as shown in Figure 8-39) and click **Add Port to List.**

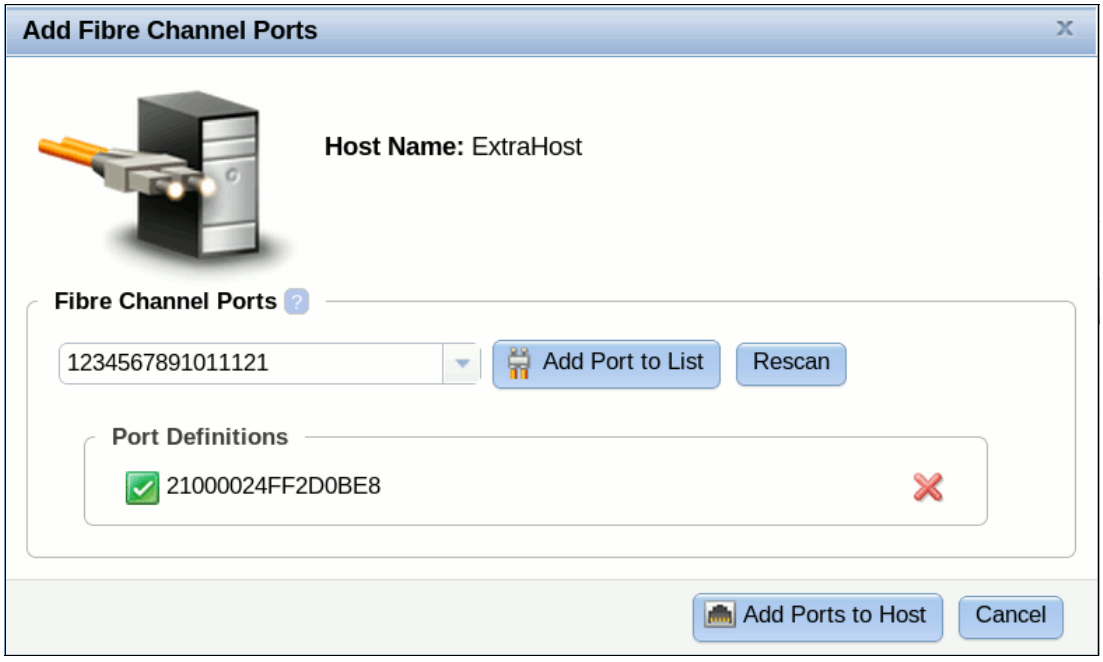

Figure 8-39 Add an offline port

The port appears as unverified (as shown in Figure 8-40) because it is not logged on to the IBM Storwize V3500. The first time that it logs on, it changes its state to online automatically and the mapping is applied to this port.

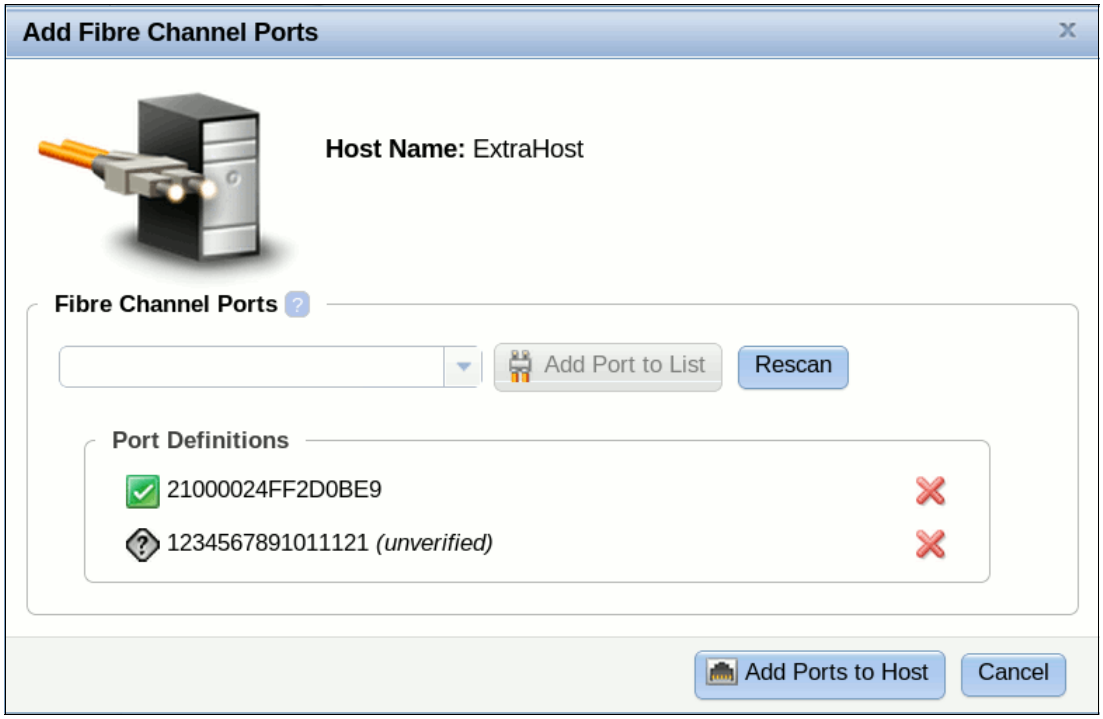

Figure 8-40 Unverified Fibre Channel port

To remove one of the ports from the list, click the red X next to it. In this example, we delete the manually added FC port so only the detected port remains.

Figure 8-41 shows that the manually added Fibre Channel port was deleted.

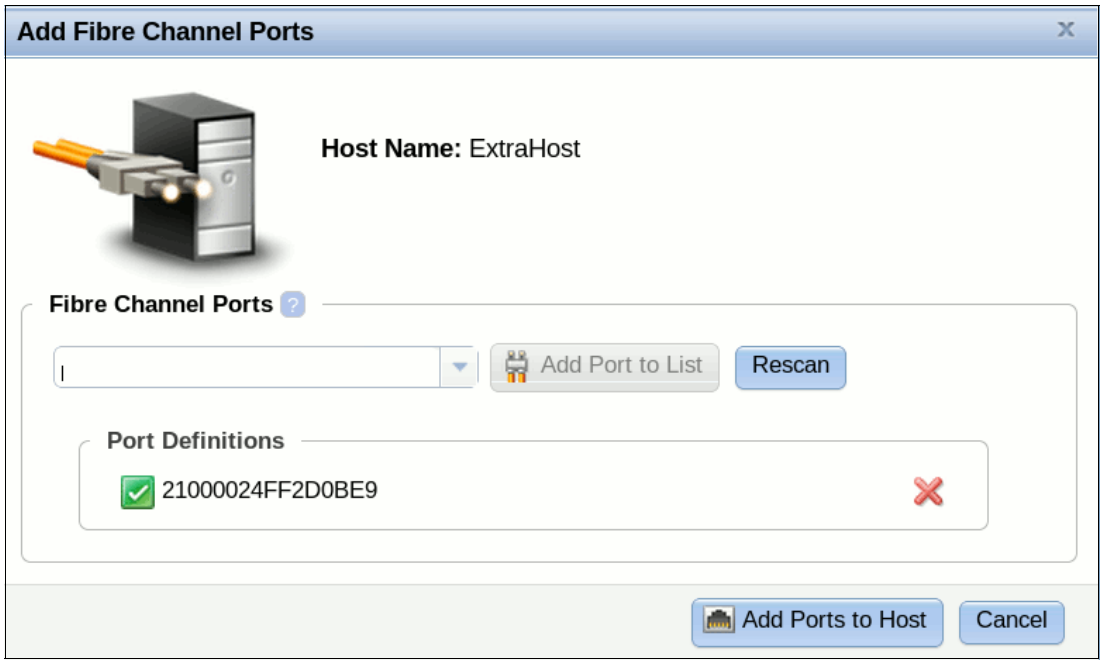

Figure 8-41 Remove a port from a list

Click **Add Ports to Host** and the changes are applied, as shown in Figure 8-42.

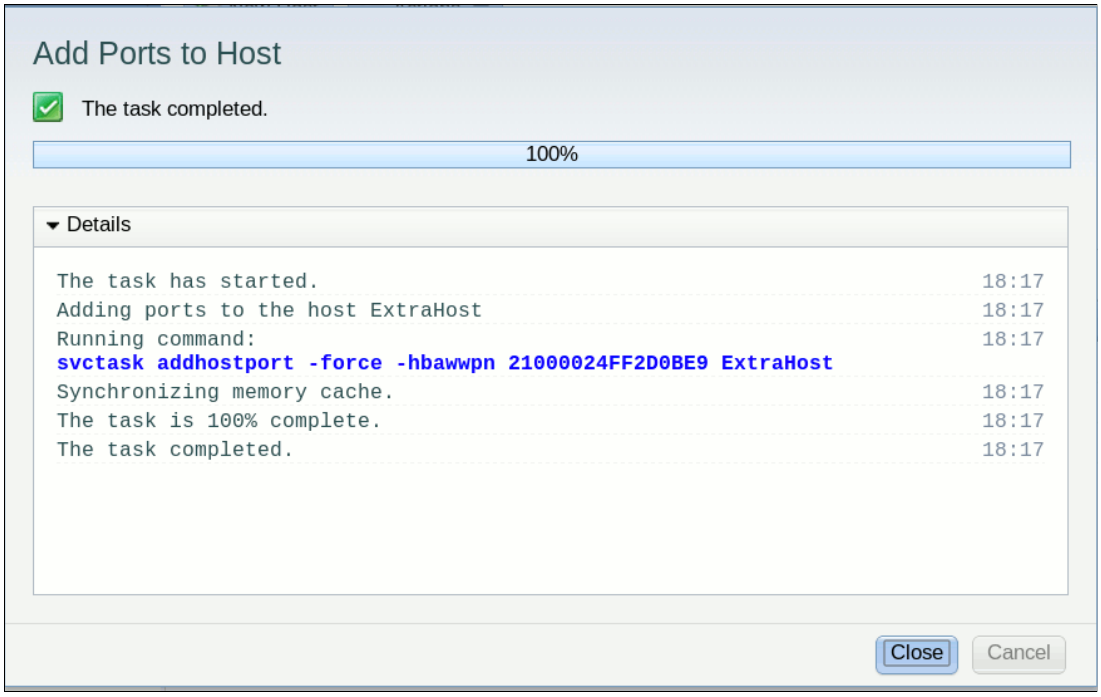

Figure 8-42 Add ports to a host

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

8.2.3 Adding an iSCSI host port

To add an iSCSI host port, click **iSCSI Port** (as shown in Figure 8-35 on page 339) and the Add iSCSI Ports window opens, as shown in Figure 8-43.

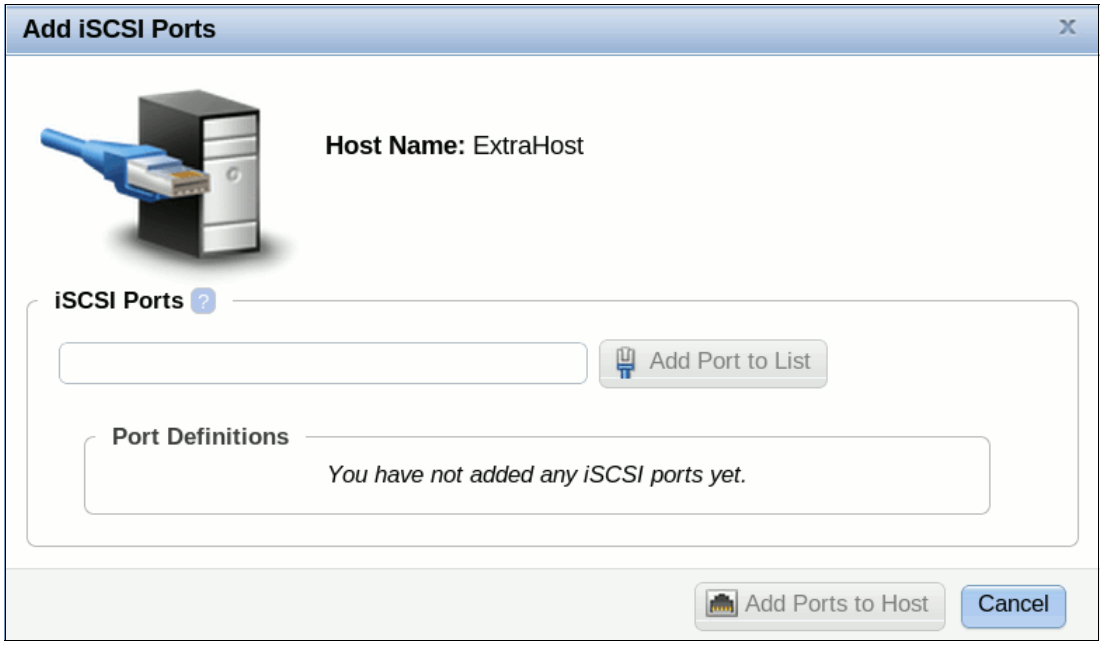

Figure 8-43 Add iSCSI ports

Enter the initiator name of your host, as shown in Figure 8-44.

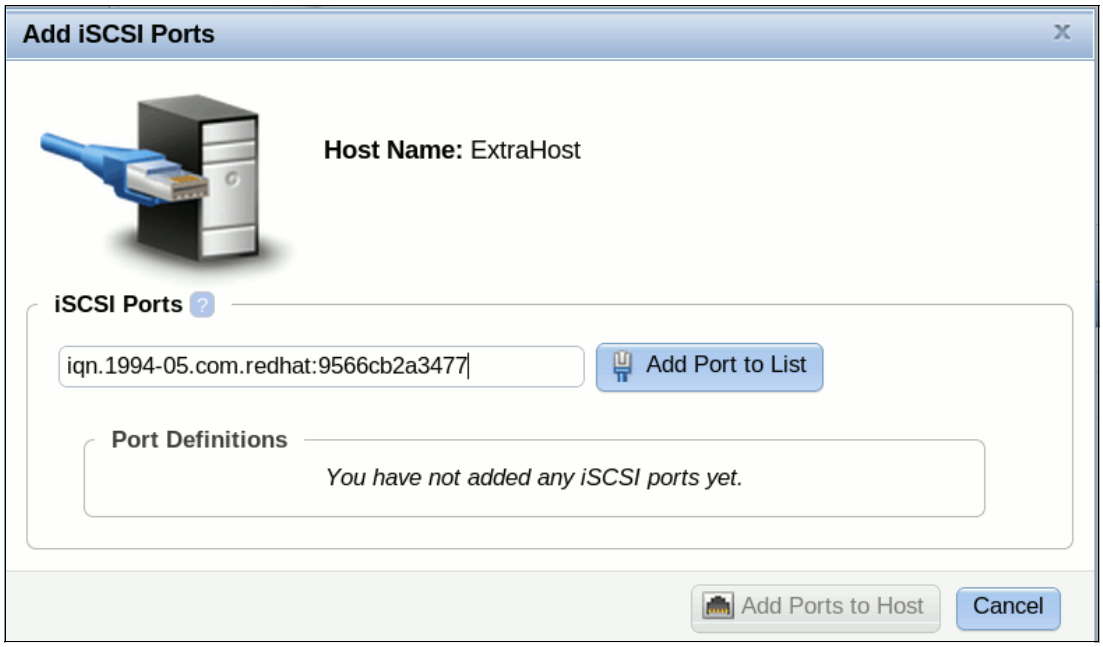

Figure 8-44 Enter initiator name

Click **Add Port to List**, as shown in Figure 8-45.

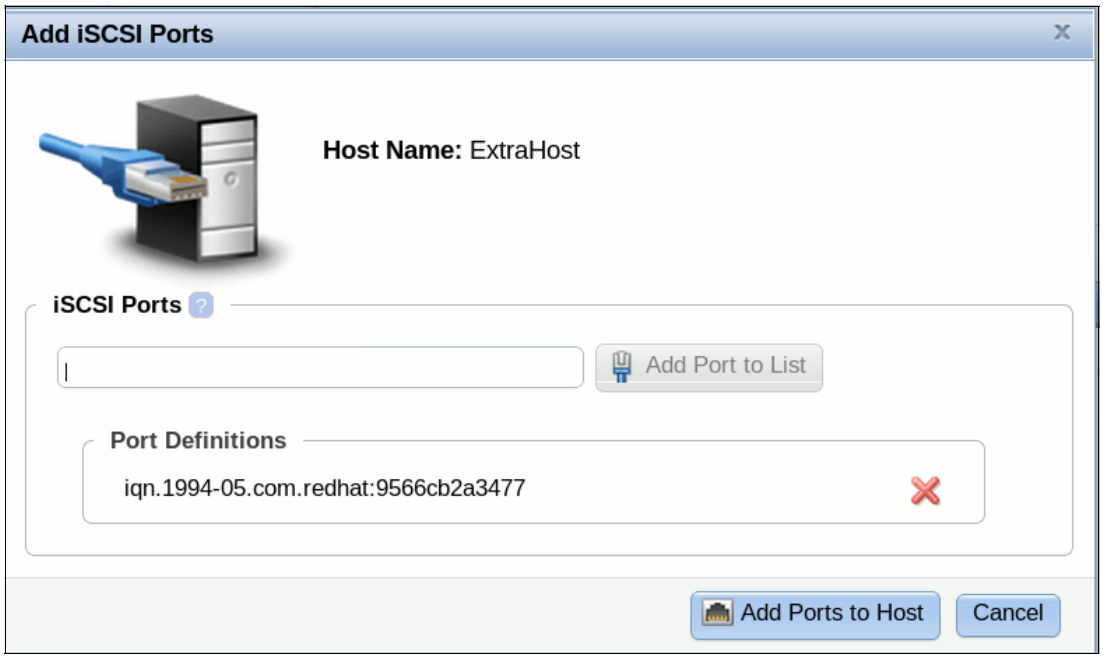

Figure 8-45 Add ports to the port definitions

Click **Add Ports to Host** to complete the tasks and apply the changes to the system, as shown in Figure 8-46. Figure 8-46 shows the output after an iSCSI port is added.

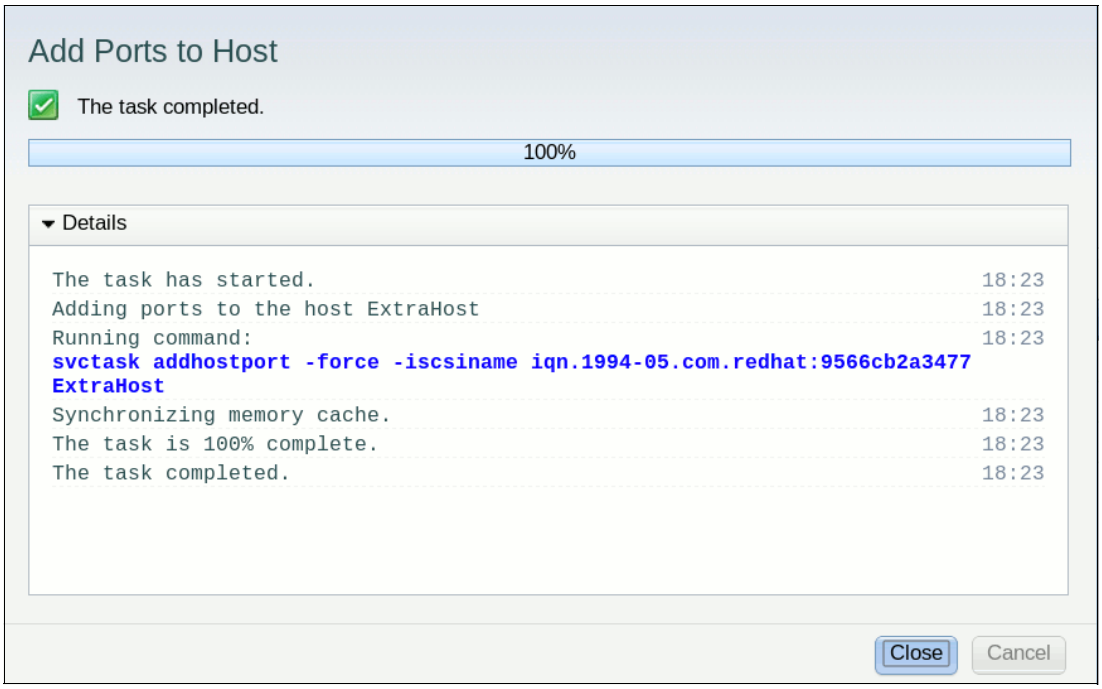

Figure 8-46 Add Host Ports: Apply changes

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

8.2.4 Deleting a host port

To delete a host port, highlight it and click **Delete Port**, as shown in Figure 8-47.

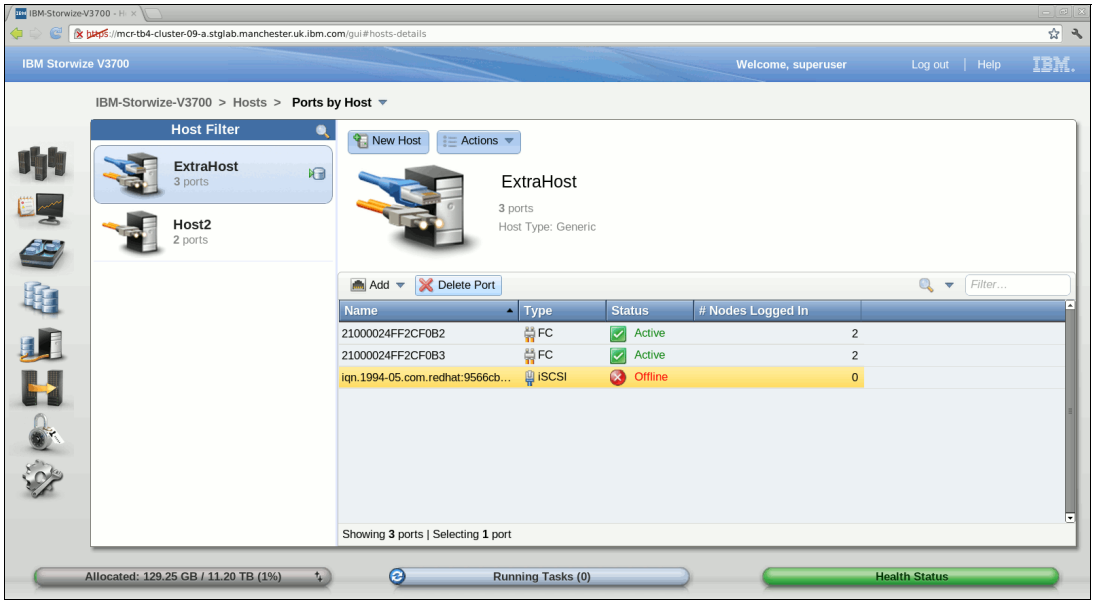

Figure 8-47 Delete port

If you press and hold the Ctrl key, you also can select several host ports to delete, as shown in Figure 8-48.

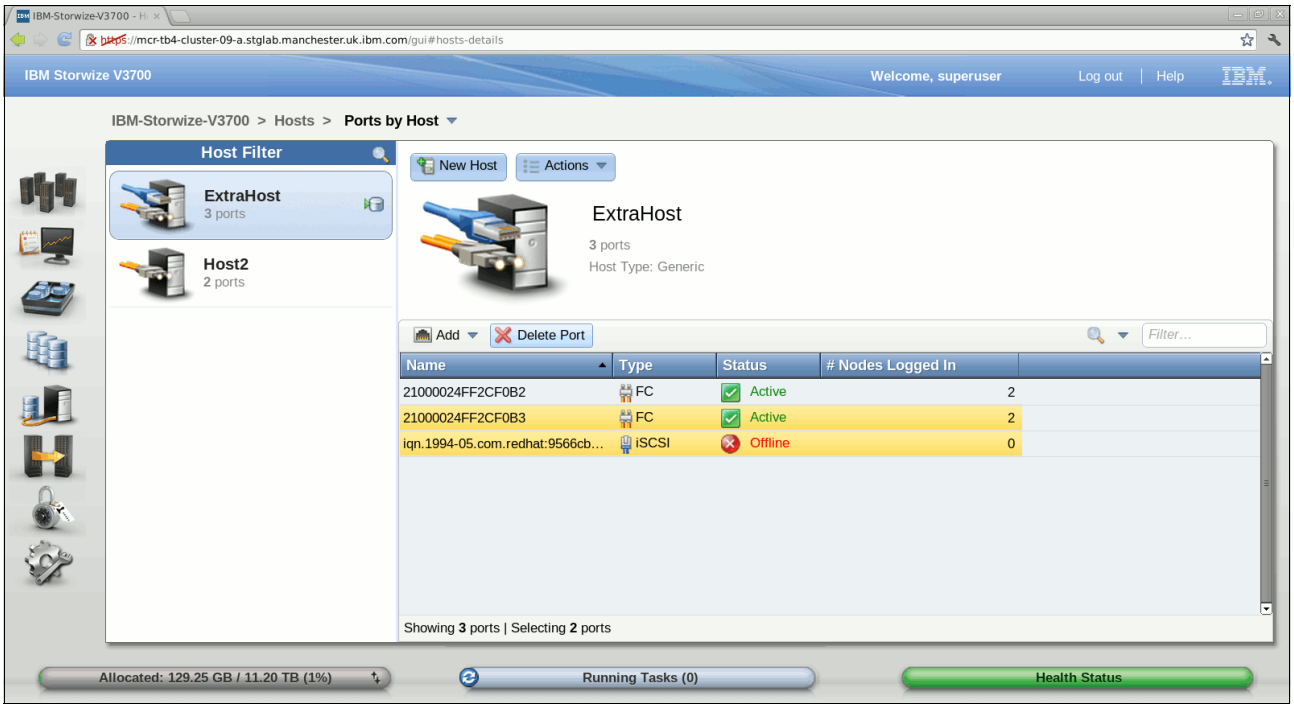

Figure 8-48 Delete multiple host ports

Click **Delete** and enter the number of host ports you want to remove, as shown in Figure 8-49. Figure 8-49 shows the Delete Ports window.

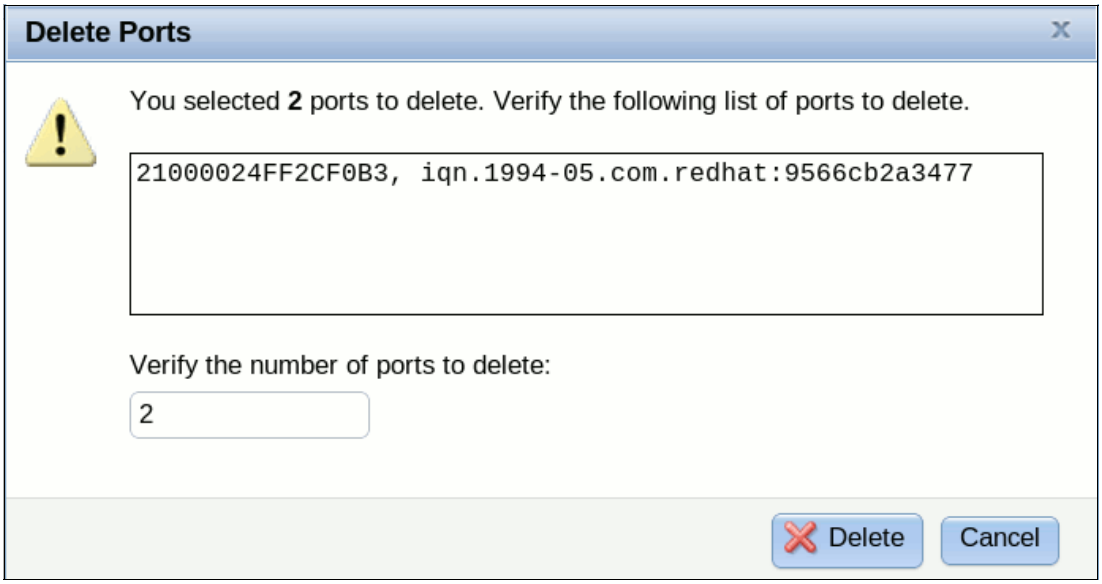

Figure 8-49 Enter the number of ports to delete

Click **Delete** and the changes are applied to the system. Figure 8-50 shows the result of deleting two ports.

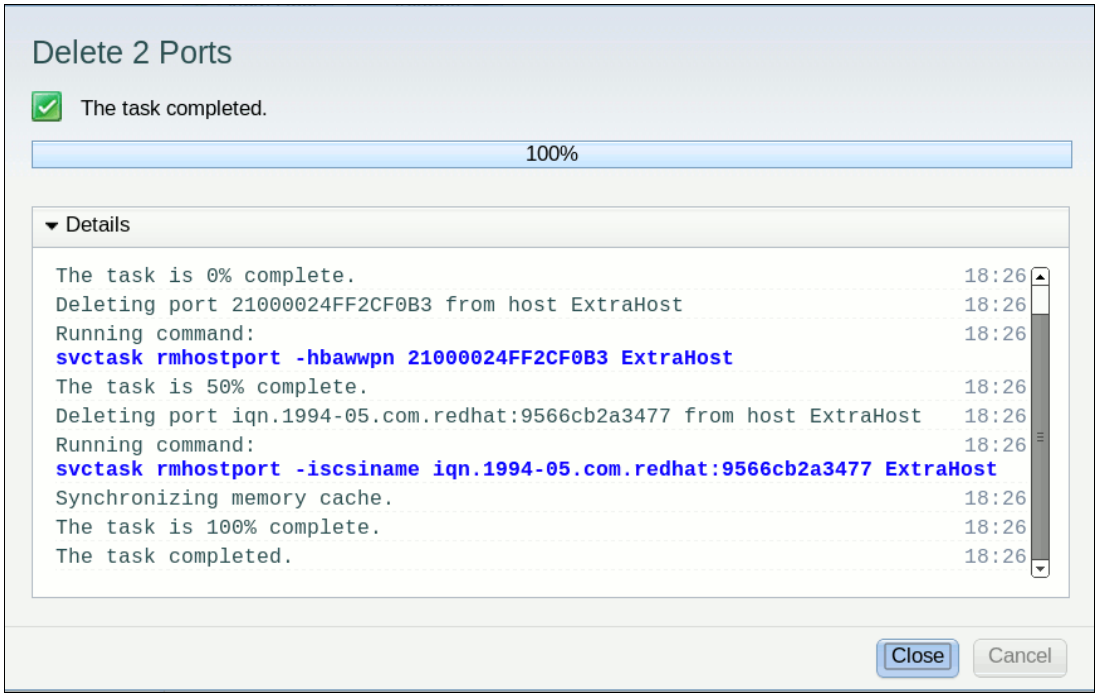

Figure 8-50 Delete host ports

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

8.3 Host mappings overview

Select **Host Mappings** (as shown in Figure 8-1 on page 318) to open the Host Mappings overview, as shown in Figure 8-51.

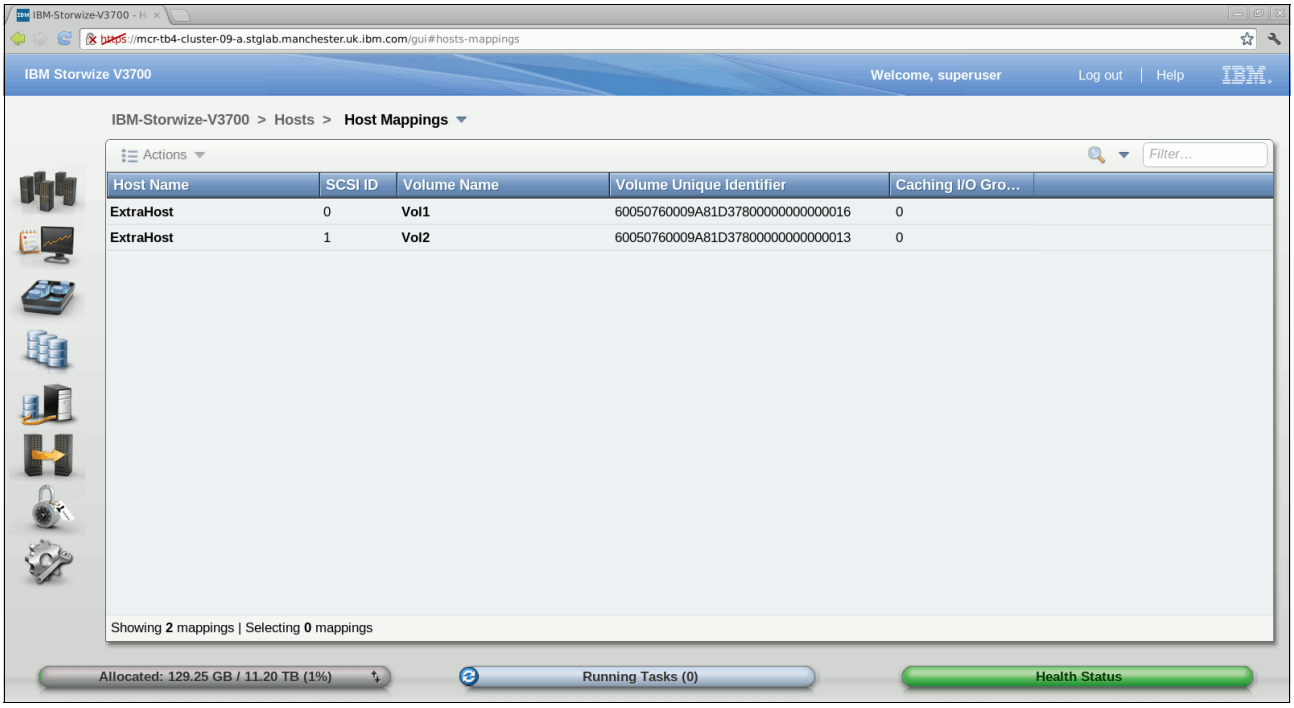

Figure 8-51 Host mappings

This window shows a list of all the host and volumes. In our example, you can see that the host ExtraHost has two volumes that are mapped and the associated SCSI ID, Volume Name, and the Volume Unique Identifier (UID).

If you highlight one line and click **Actions** (as shown in Figure 8-52), the following tasks are available:

- Unmapping a volume
- Properties (Host)
- Properties (Volume)

Figure 8-52 shows the Actions menu.

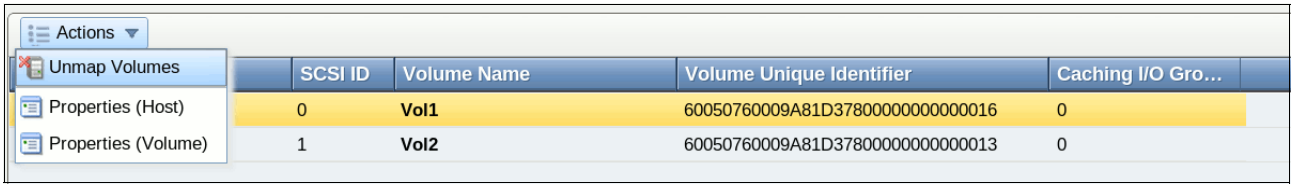

Figure 8-52 Host Mappings Actions drop-down menu

8.3.1 Unmapping a volume

Highlight one or more lines and click **Unmap Volumes**, enter the number of entries to remove (as shown in Figure 8-53), and click **Unmap**. This action removes the mappings for all selected entries.

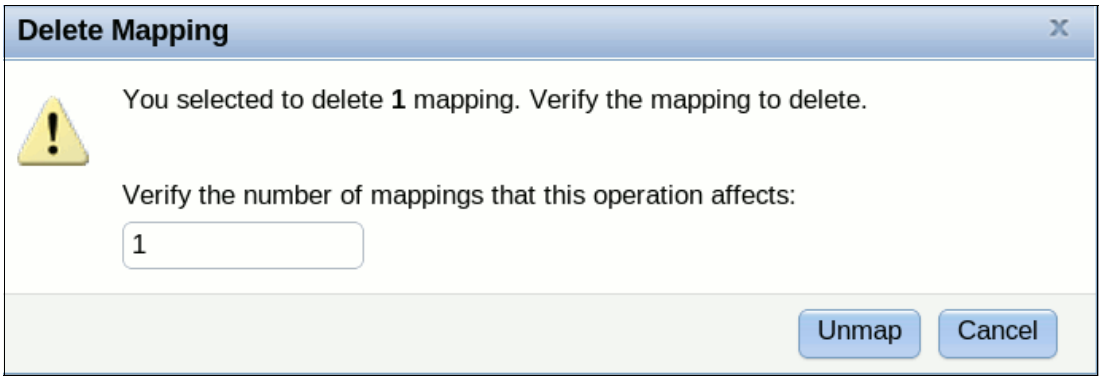

Figure 8-53 Unmap Selected Volumes

Figure 8-54 shows the output from removing a volume mapping.

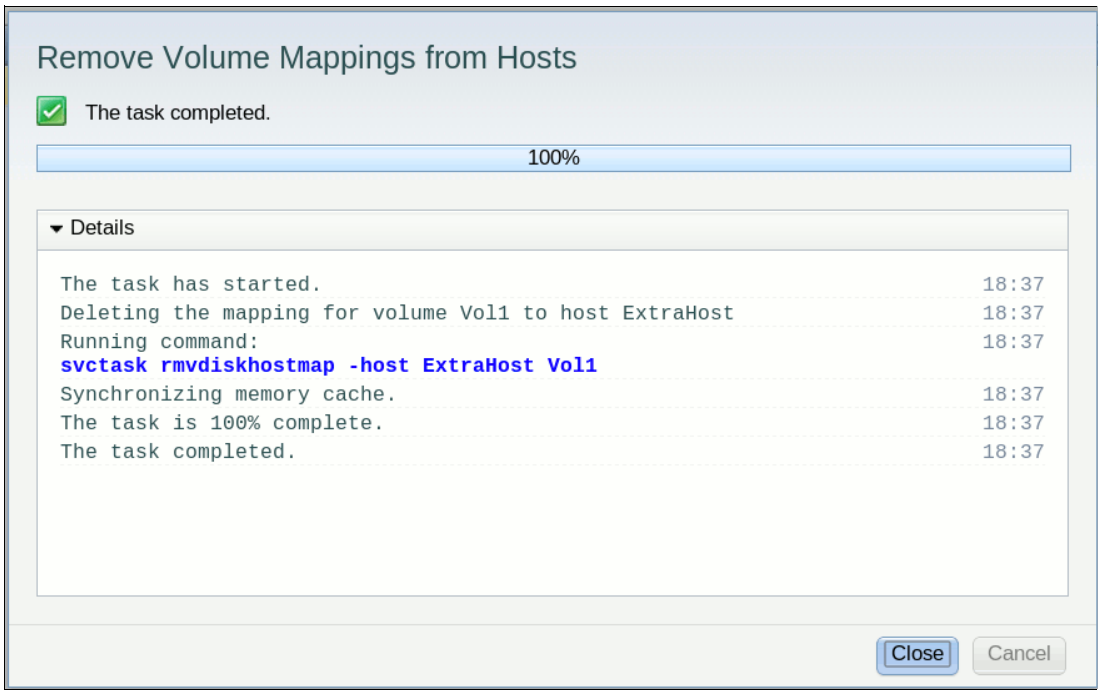

Figure 8-54 Unmapping complete

8.3.2 Properties (Host)

Selecting an entry and clicking **Properties (Host)** (as shown in Figure 8-52 on page 348) opens the Host Properties window. The contents of this window are described in 8.1.5, “Host properties” on page 333.

8.3.3 Properties (Volume)

Selecting an entry and clicking **Properties (Volume)** (as shown in Figure 8-52 on page 348) opens the Volume Properties view. The contents of this window are described in 5.1, “Provisioning storage from IBM Storwize V3500 and making it available to the host” on page 172.

8.4 Advanced volume administration

This section describes volume administration, such as volume modification and the migration or creation of new volume copies. Basic volume creation is covered in 5.1, “Provisioning storage from IBM Storwize V3500 and making it available to the host” on page 172. In this section, we assume that you created some volumes on your IBM Storwize V3500 and that you are familiar with generic, thin-provisioned, and mirrored volumes.

Figure 8-55 shows that there are three sections available within the Volumes menu for the administration of advanced features:

- Volumes, as described in 8.4.1, “Advanced volume functions” and 8.6, “Advanced volume copy functions” on page 380
- Volumes by Host, as described in 8.8, “Volumes by Host” on page 395
- Volumes by Pool, as described in 8.7, “Volumes by storage pool” on page 389

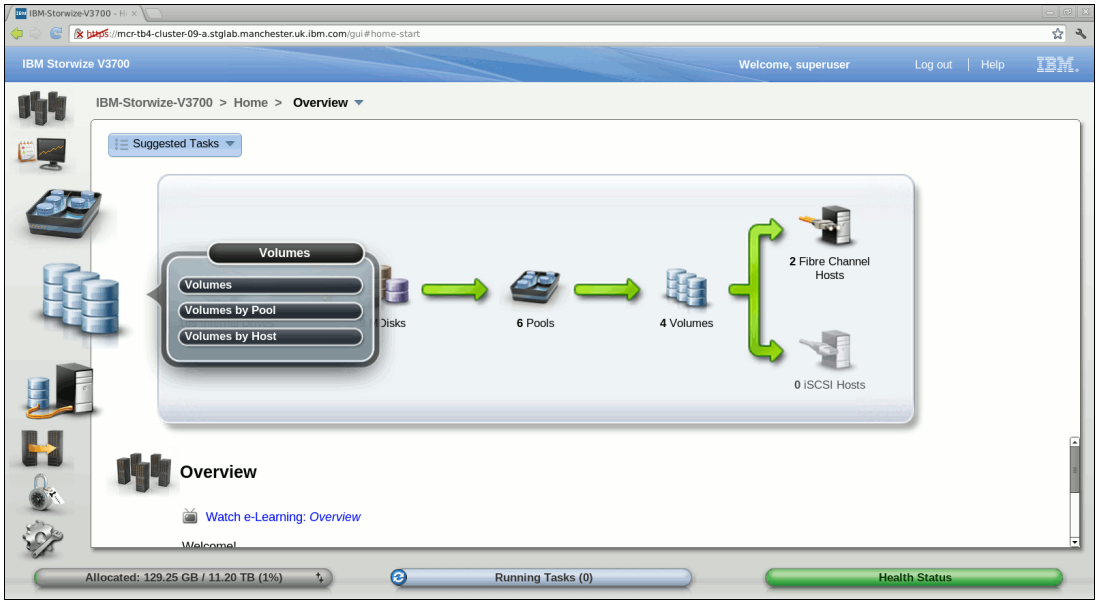

Figure 8-55 Volumes menu

8.4.1 Advanced volume functions

Click **Volumes** (as shown in Figure 8-55) and the Volumes window opens, as shown in Figure 8-56.

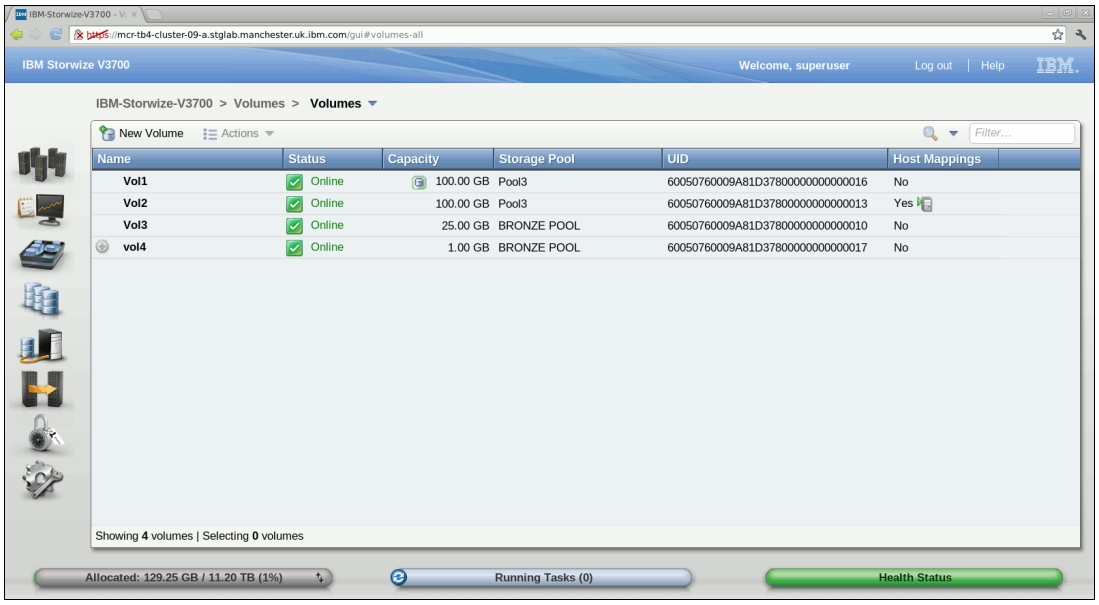

Figure 8-56 Volumes window

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 several copies of this volume. Click it to expand the view and list the copies, as shown in Figure 8-57 on page 352.
- Status: Gives you status information about the volume. It can be online, offline, or degraded.
- Capacity: The disk capacity that is presented to the host is listed here. If there is a blue volume 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. If you have several volume copies, it shows you the pool of the primary copy.
- UID: The volume unique identifier.
- Host Mappings: Shows to which hosts, if any, the volume is mapped.

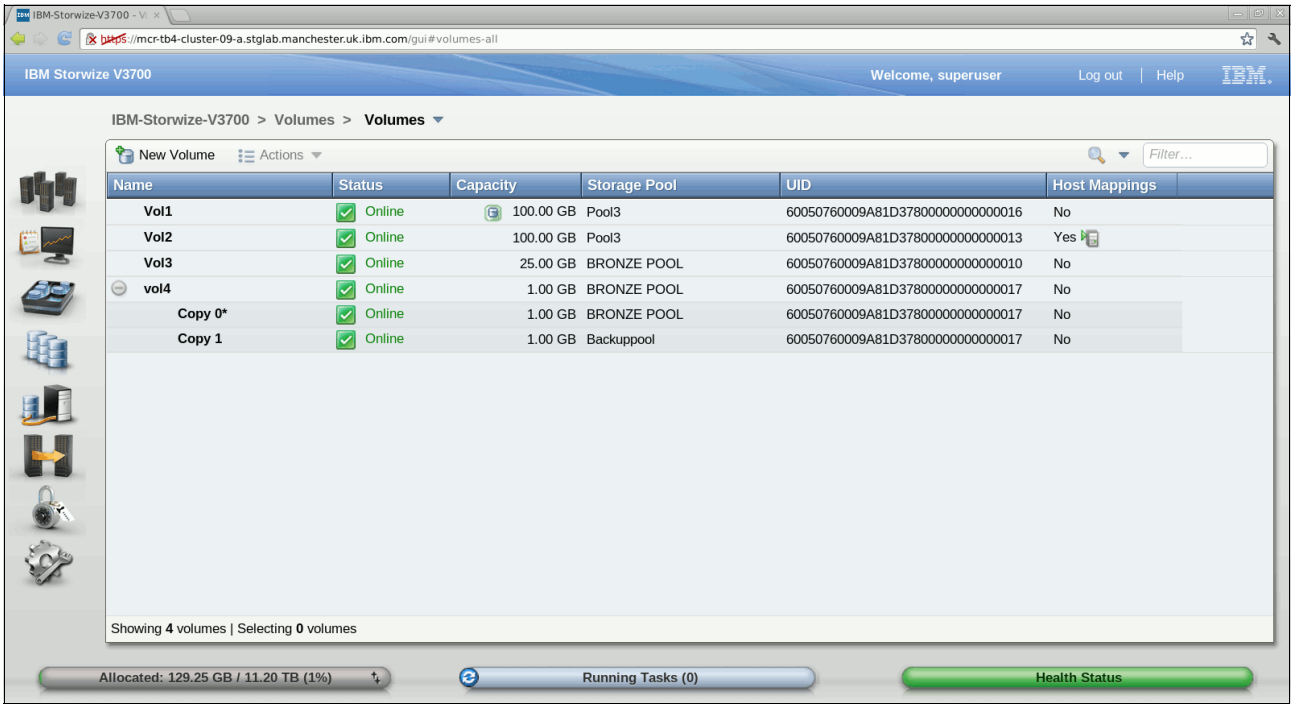

Figure 8-57 Volume copies

To create a volume, click **New Volume** and complete the steps that are described in 8.1, “Advanced host administration” on page 318. Highlight a volume and click **Actions** to see the available actions for a volume, as shown in Figure 8-58.

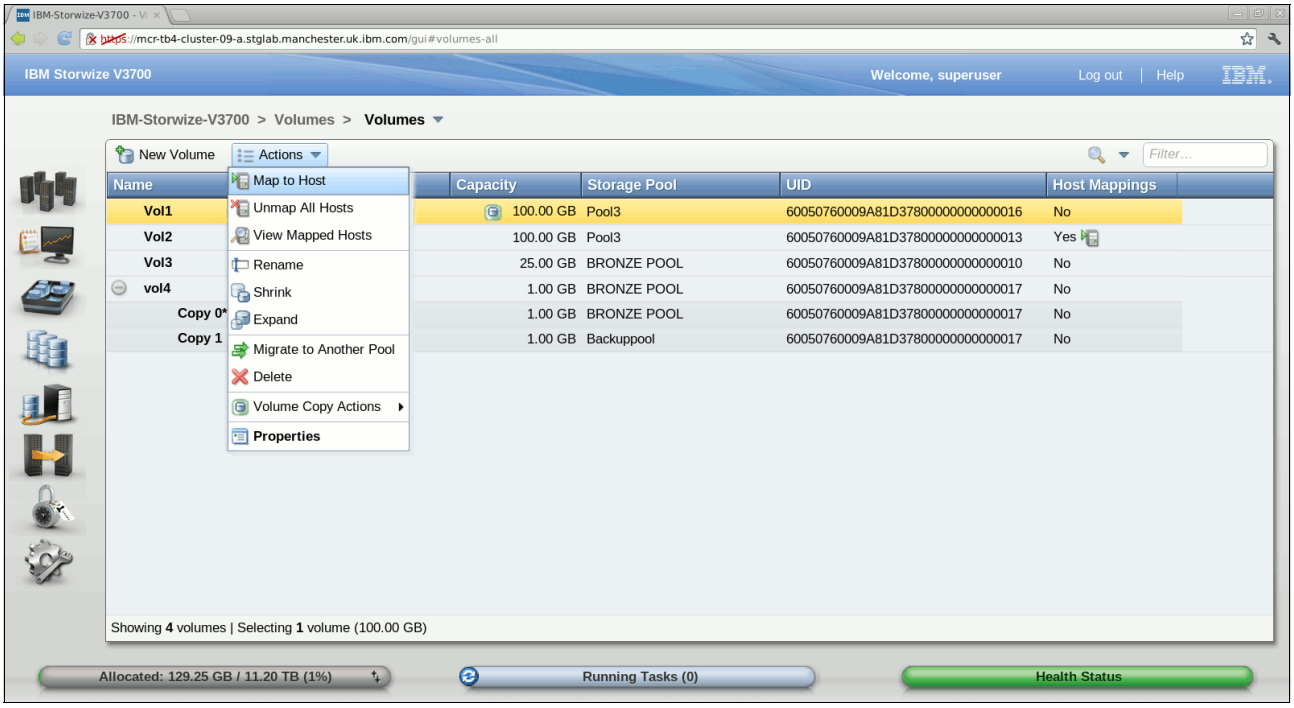

Figure 8-58 Volume Actions

You also can right-click a volume and select the actions in the menu, as shown in Figure 8-59.

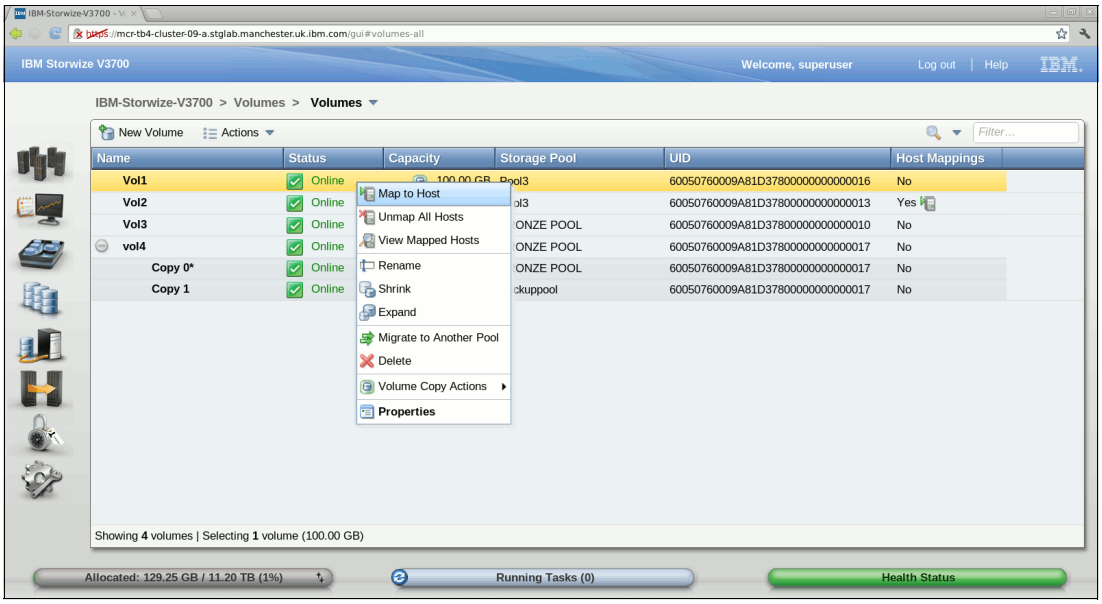

Figure 8-59 Volume actions right-click menu

The following volume options are available:

- Map to Host
- Unmap All Hosts
- View Mapped Host
- Rename
- Shrink
- Expand
- Migration to Another Pool
- Delete
- Properties

Depending on which volume you view, the following options might be available for Volume Copy Actions:

- Add Mirror Copy: Only available for generic volumes
- Thin Provisioned: Only available for thin-provisioned volumes

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, as shown in Figure 8-58 on page 352. Select the host to which you want to map the volume and click **Next**, as shown in Figure 8-60.

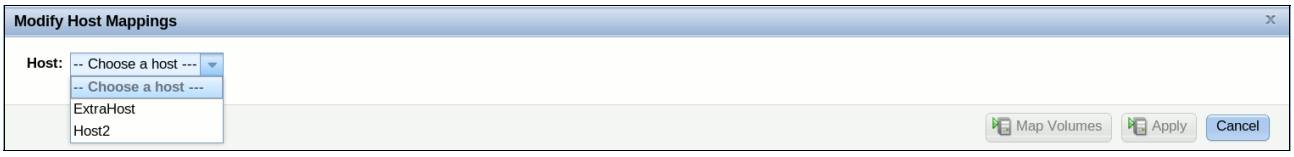

Figure 8-60 Modify Host Mappings menu

The Modify Mappings window opens. In the upper left, you see your selected host and the yellow volume is the selected volume that is mapped, as shown in Figure 8-61. Click **Map Volumes** to apply the changes to the system.

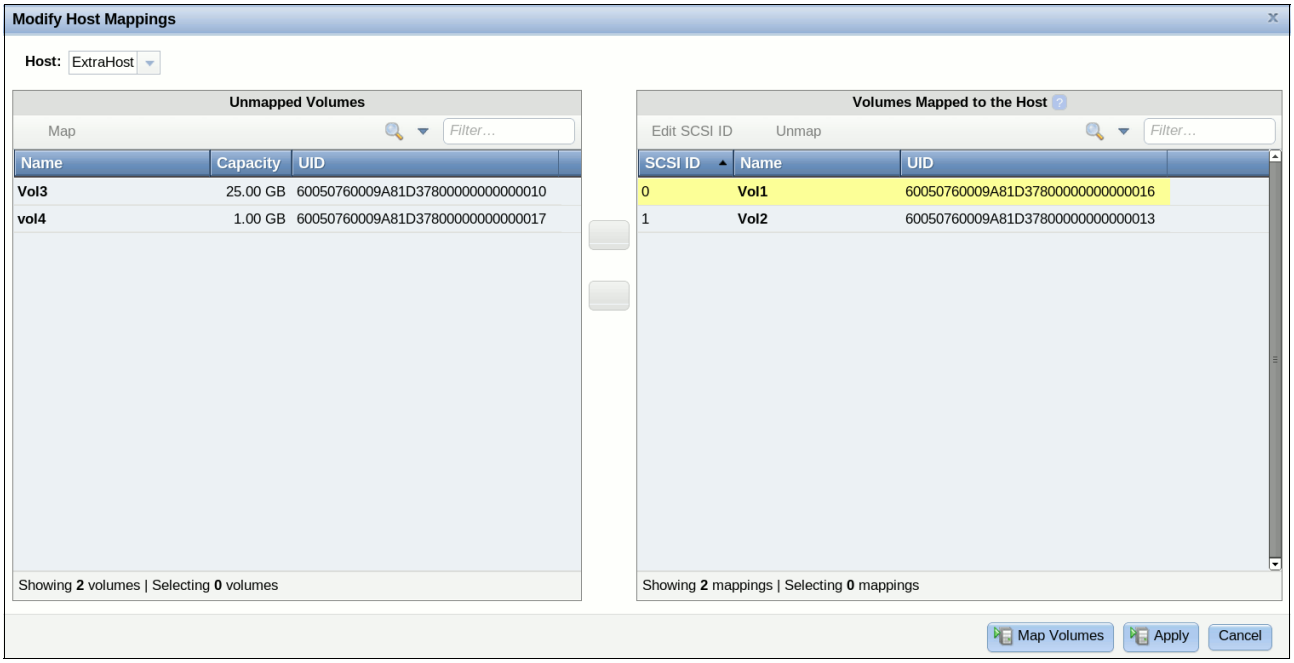

Figure 8-61 Modify Host Mappings

After the changes are complete, click **Close** to return to the All Volumes window. Figure 8-62 shows the output after mappings are modified.

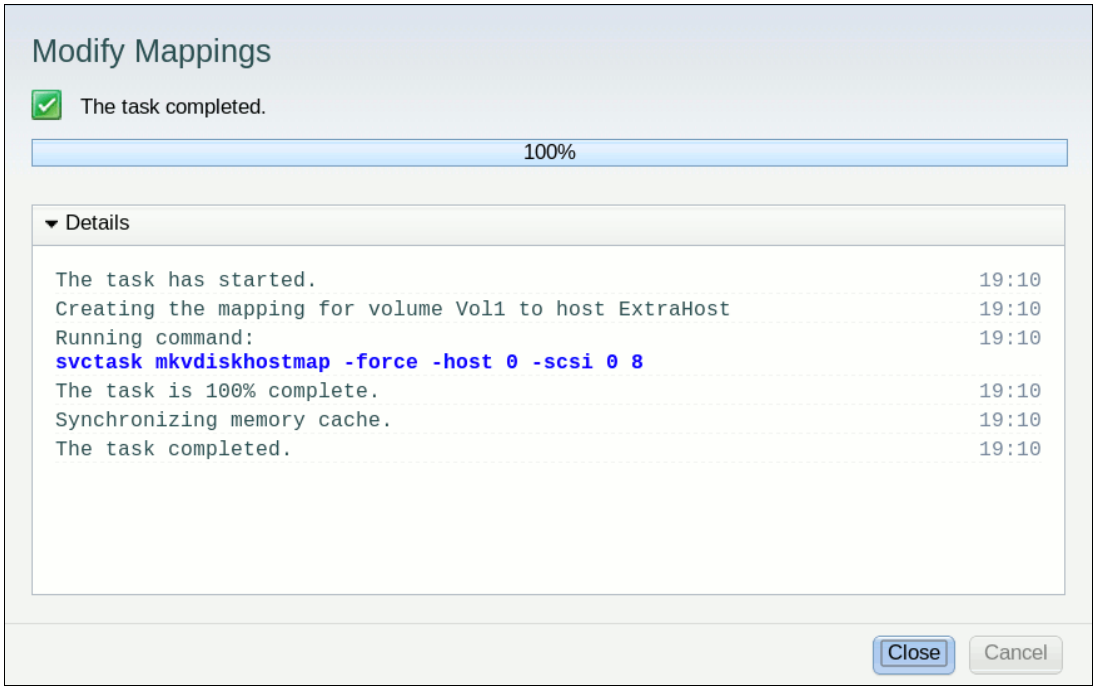

Figure 8-62 Modify Mappings output

Modify Mappings window: The Modify Mappings window is described in 8.1.1, “Modifying Mappings menu” on page 320.

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 8-58 on page 352). All host mappings are removed from this volume, which means that no hosts can access this volume. Enter the number of mappings that are affected and click **Unmap**, as shown in Figure 8-63.

.

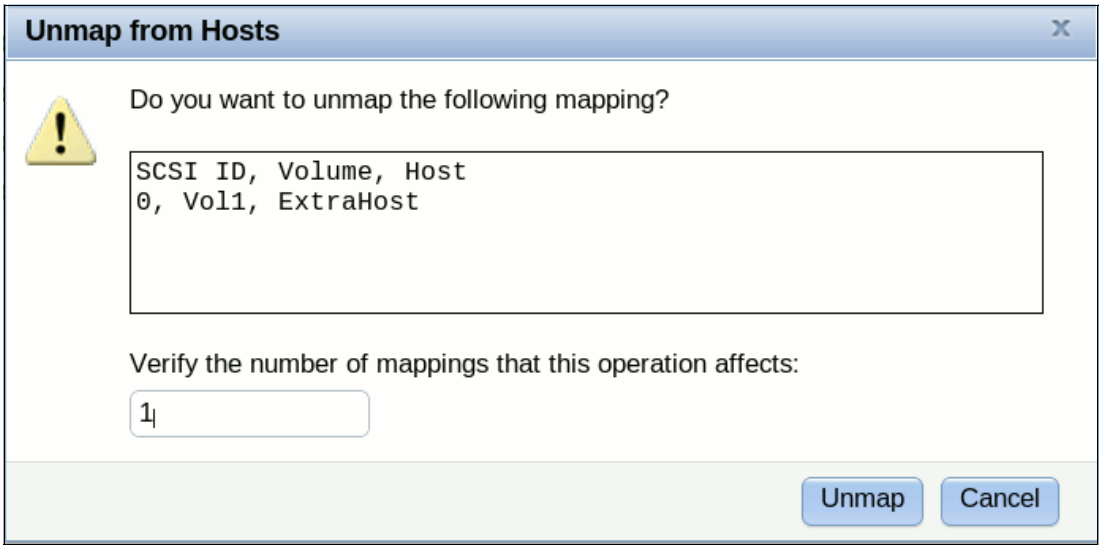

Figure 8-63 Unmap Volumes: Enter number of mappings to be removed

After the task completes, click **Close** to return to the All Volumes window. Figure 8-64 shows the output from unmapping the volume.

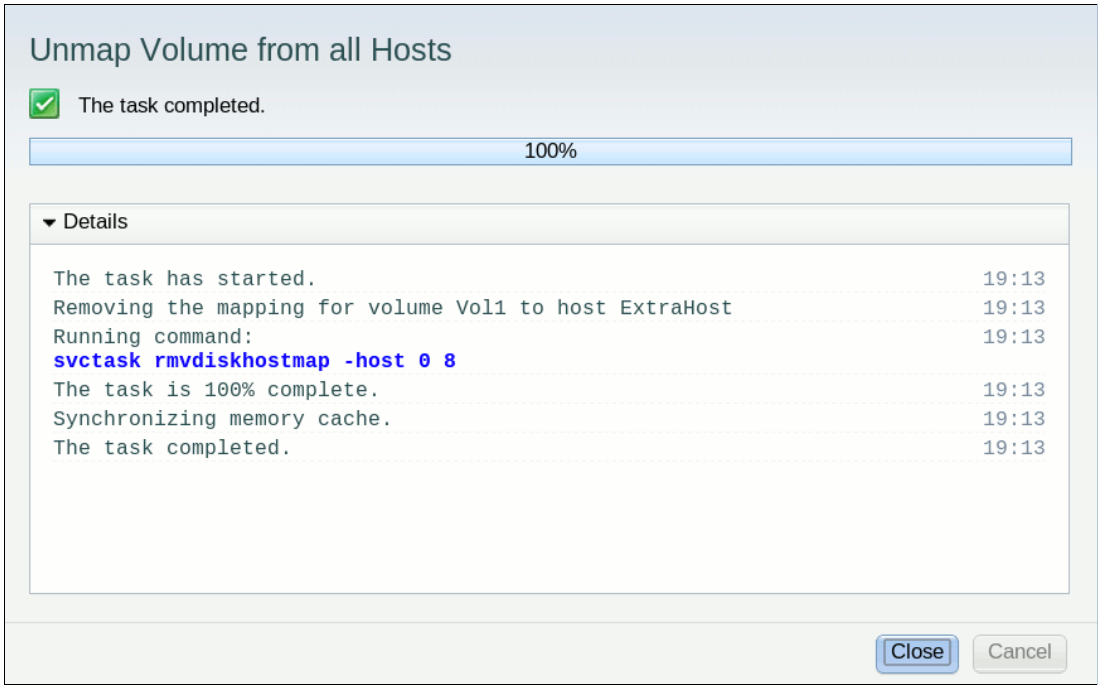

Figure 8-64 Unmap volume complete

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 8-58 on page 352). This action opens the Host Maps tab of the Volume Details window, as shown in Figure 8-65. In this example, you see that there is one existing host mapping for the Vol2 volume.

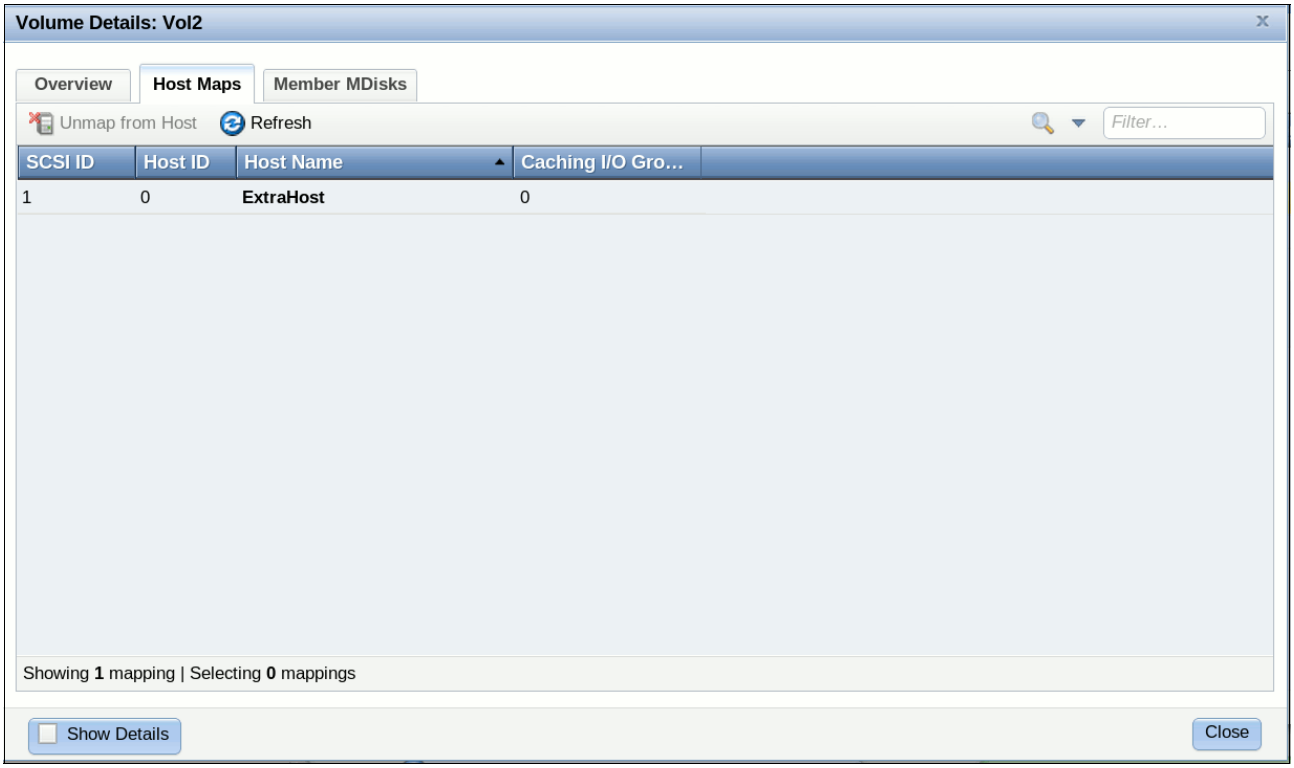

Figure 8-65 Host Maps

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 there are several hosts that 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 8-58 on page 352). A window opens. Enter the new name. Figure 8-66 shows the Rename Volume window.

Figure 8-66 Rename Volume

If you click **Reset**, the name field is always reset to the active name of the volume. Click **Rename** to apply the changes and then click **Close** when you are done, as shown in Figure 8-67.

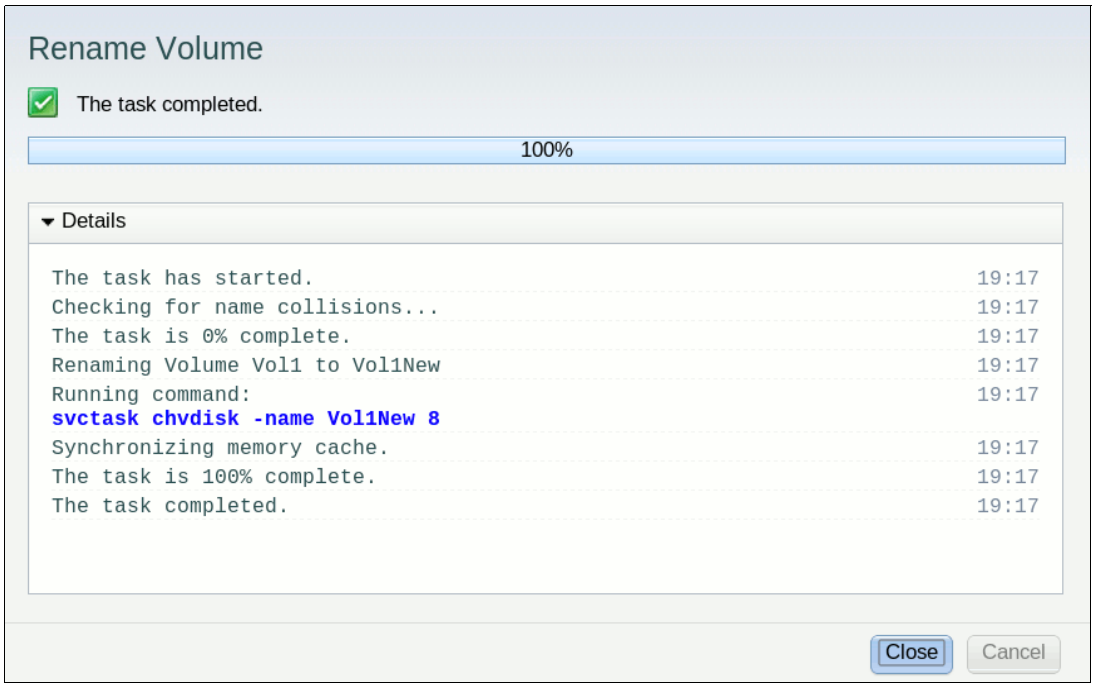

Figure 8-67 Rename Volume completed

8.4.6 Shrinking a volume

By using IBM Storwize V3500, you can shrink volumes. This feature should be used only if your host operating system supports this feature. Before you shrink a volume, complete the preparation that is required in your host operating system. After you prepare your operating system, click **Shrink** (as shown in Figure 8-58 on page 352). You can enter the new size or enter how much the volume should shrink. If you enter a value, the other line updates automatically, as shown in Figure 8-68.

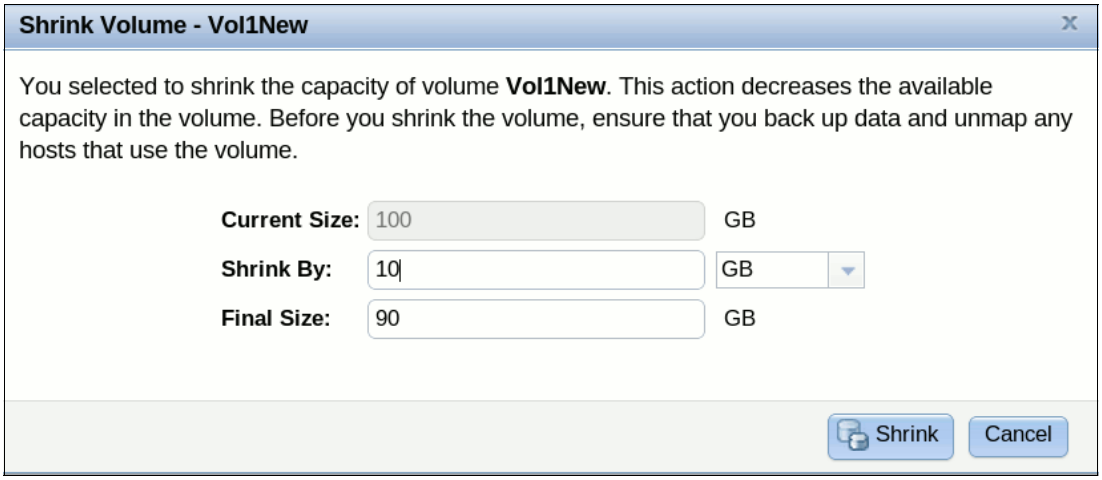

Figure 8-68 Shrink Volume

Click **Shrink** to start the process and click **Close** to return to the All Volumes window.

Figure 8-69 shows the output after the volume is shrunk.

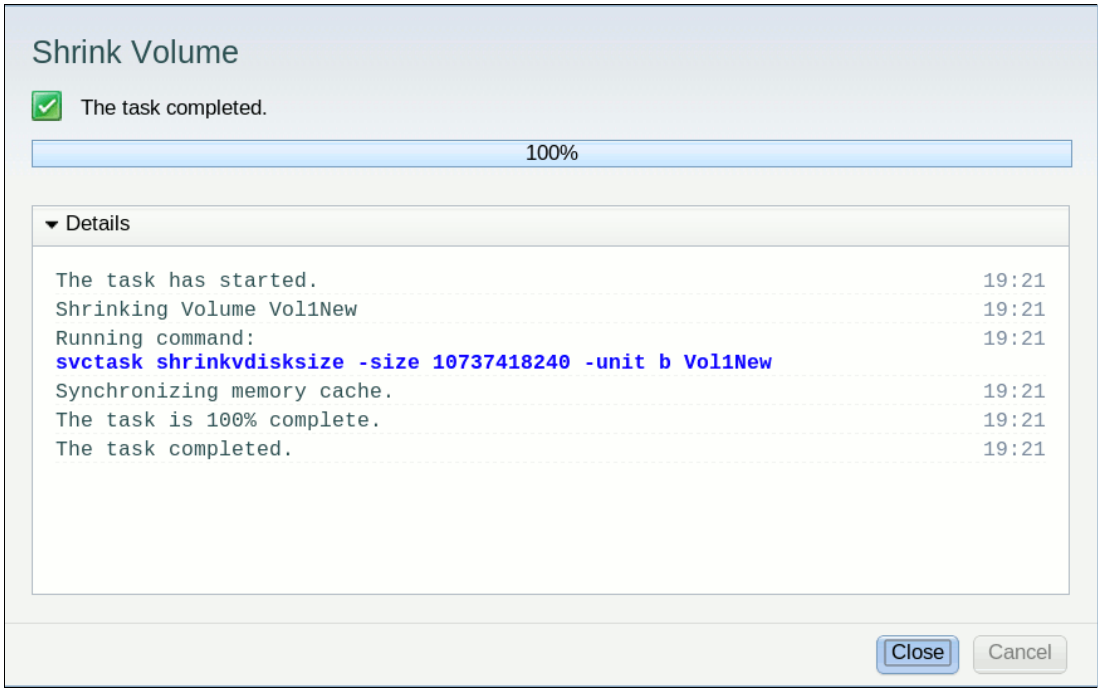

Figure 8-69 Shrink Volume

Run the required procedures on your host to complete the shrinking process.

8.4.7 Expanding a volume

If you want to expand a volume, click **Expand** (as shown in Figure 8-58 on page 352) 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 8-70.

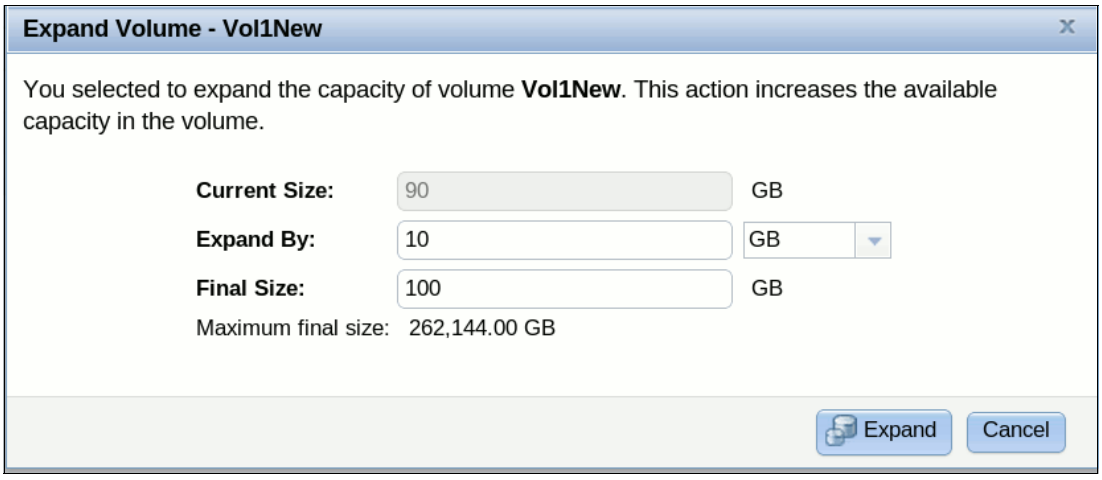

Figure 8-70 Expand volume

After the tasks are complete, click **Close** (as shown in Figure 8-71) to return to the All Volumes window.

Figure 8-71 shows the output after a volume is expanded.

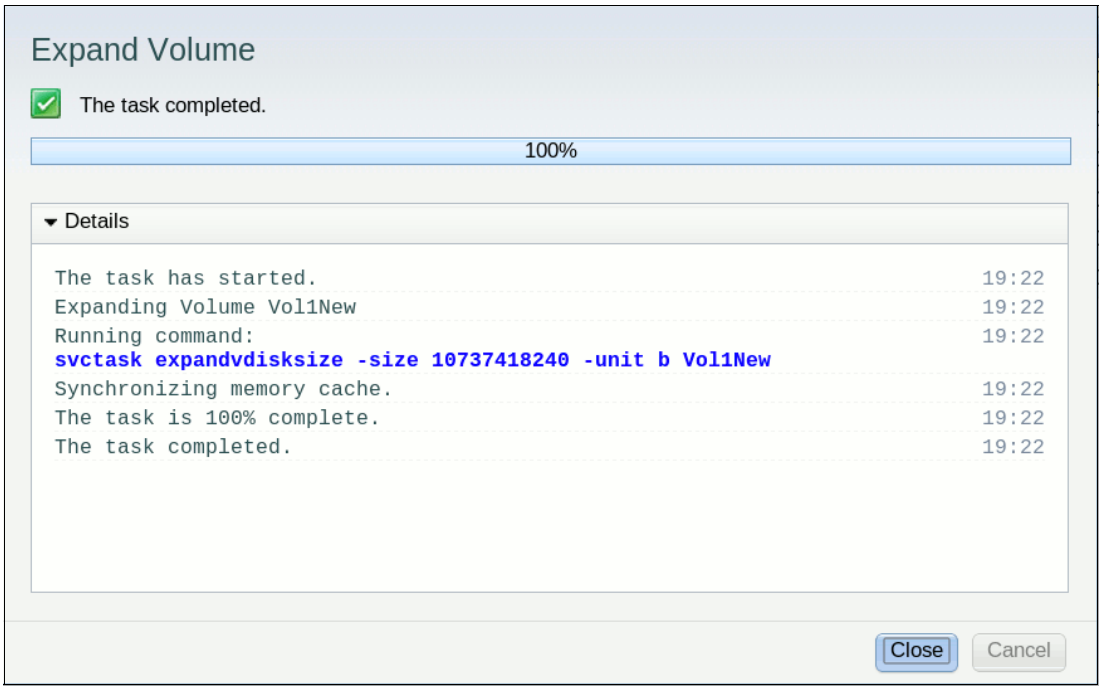

Figure 8-71 Expand Volume complete

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

8.4.8 Migrating a volume to another storage pool

IBM Storwize V3500 provides online volume migration while applications are running. Storage pools are managed disk groups, as described in Chapter 7., “Storage pools” on page 277. By using volume migration, you can move the data between these storage pools, whether the pools are internal or on an external storage system. The migration process is a low priority process that does not affect the performance of the IBM Storwize V3500. However, it moves one extent after another to the new storage pool. However, as the migration process progresses, the performance becomes closer to that of the new pool than the old pool.

To migrate a volume to another storage pool, click **Migrate to Another Pool** (see Figure 8-58 on page 352). 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 8-72). If the selected volume consists of one copy, this selection menu does not appear.

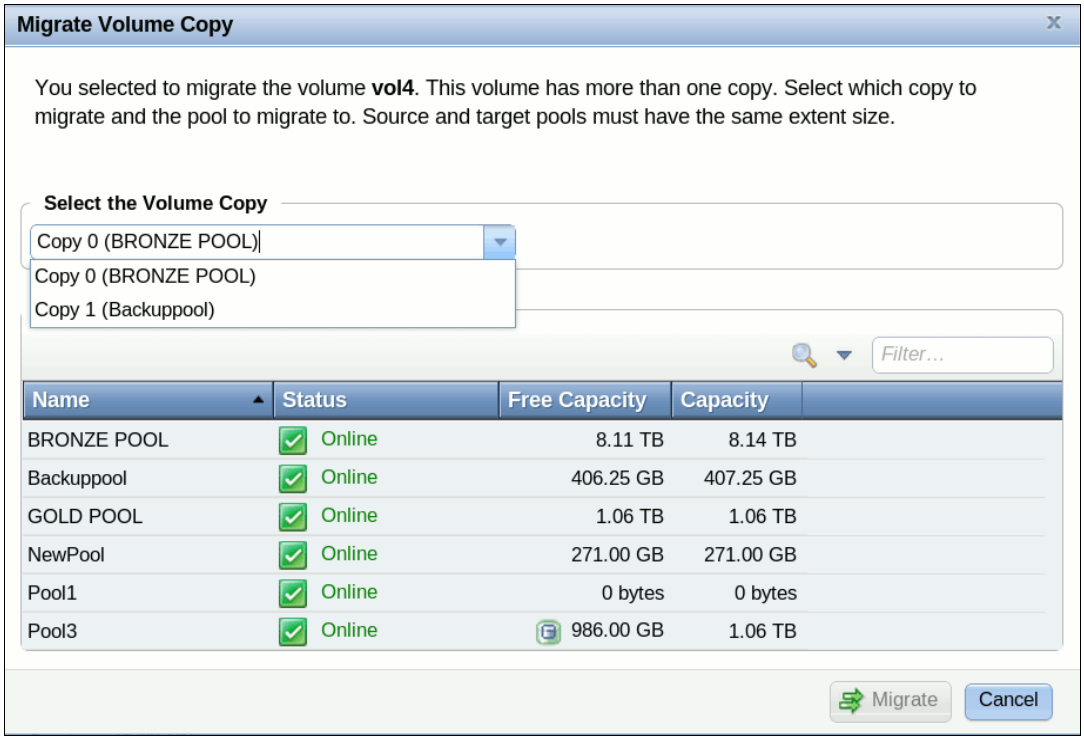

Figure 8-72 Migrate Volume Copy

Select the new target storage pool and click **Migrate**, as shown in Figure 8-73.

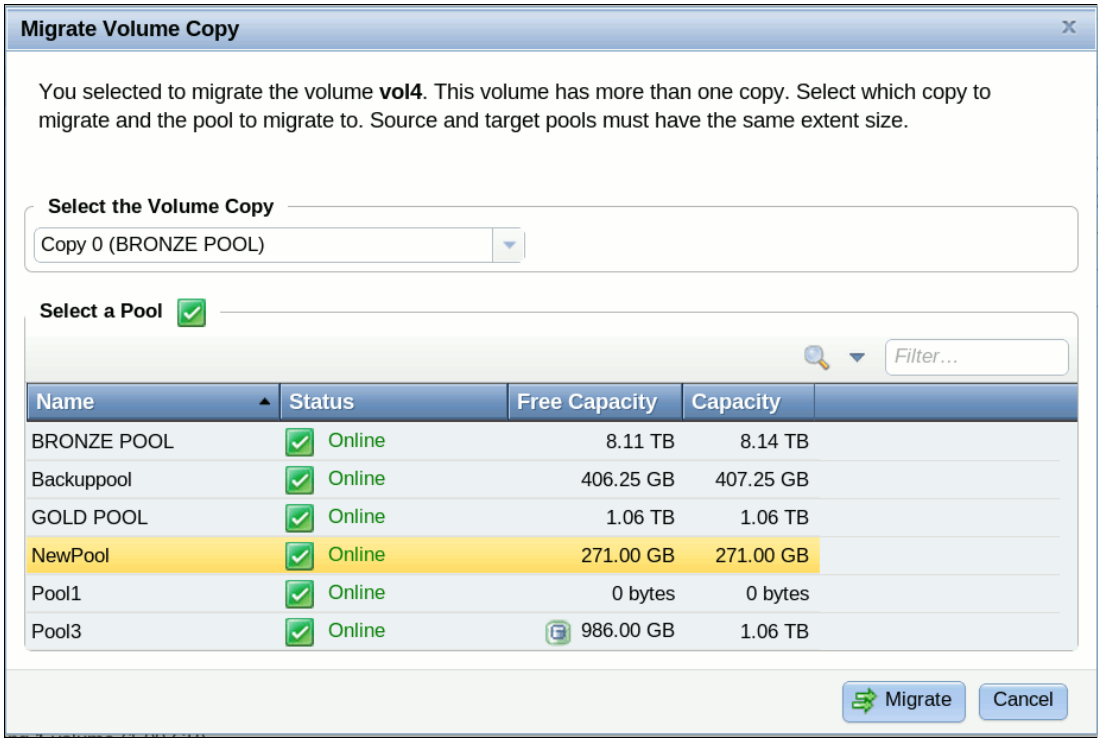

Figure 8-73 Select target storage pool

The migrate volume copy task starts, as shown in Figure 8-74. Click **Close** to return to the All Volumes window.

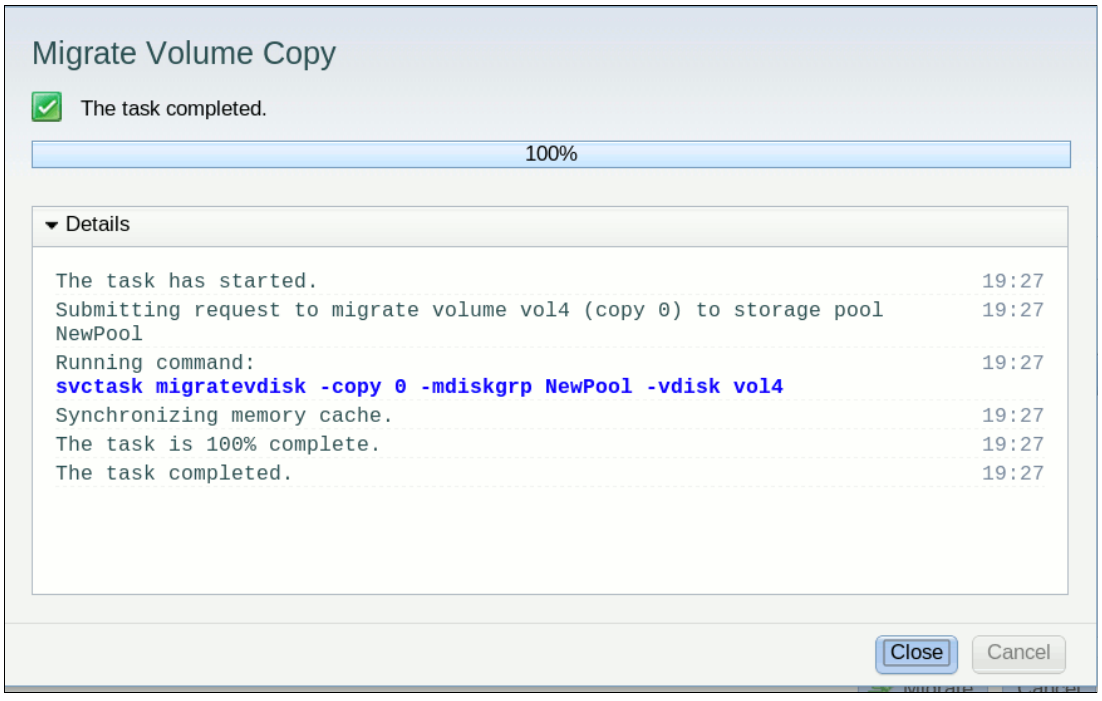

Figure 8-74 Migrate Volume Copy process starts

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, as shown in Figure 8-75.

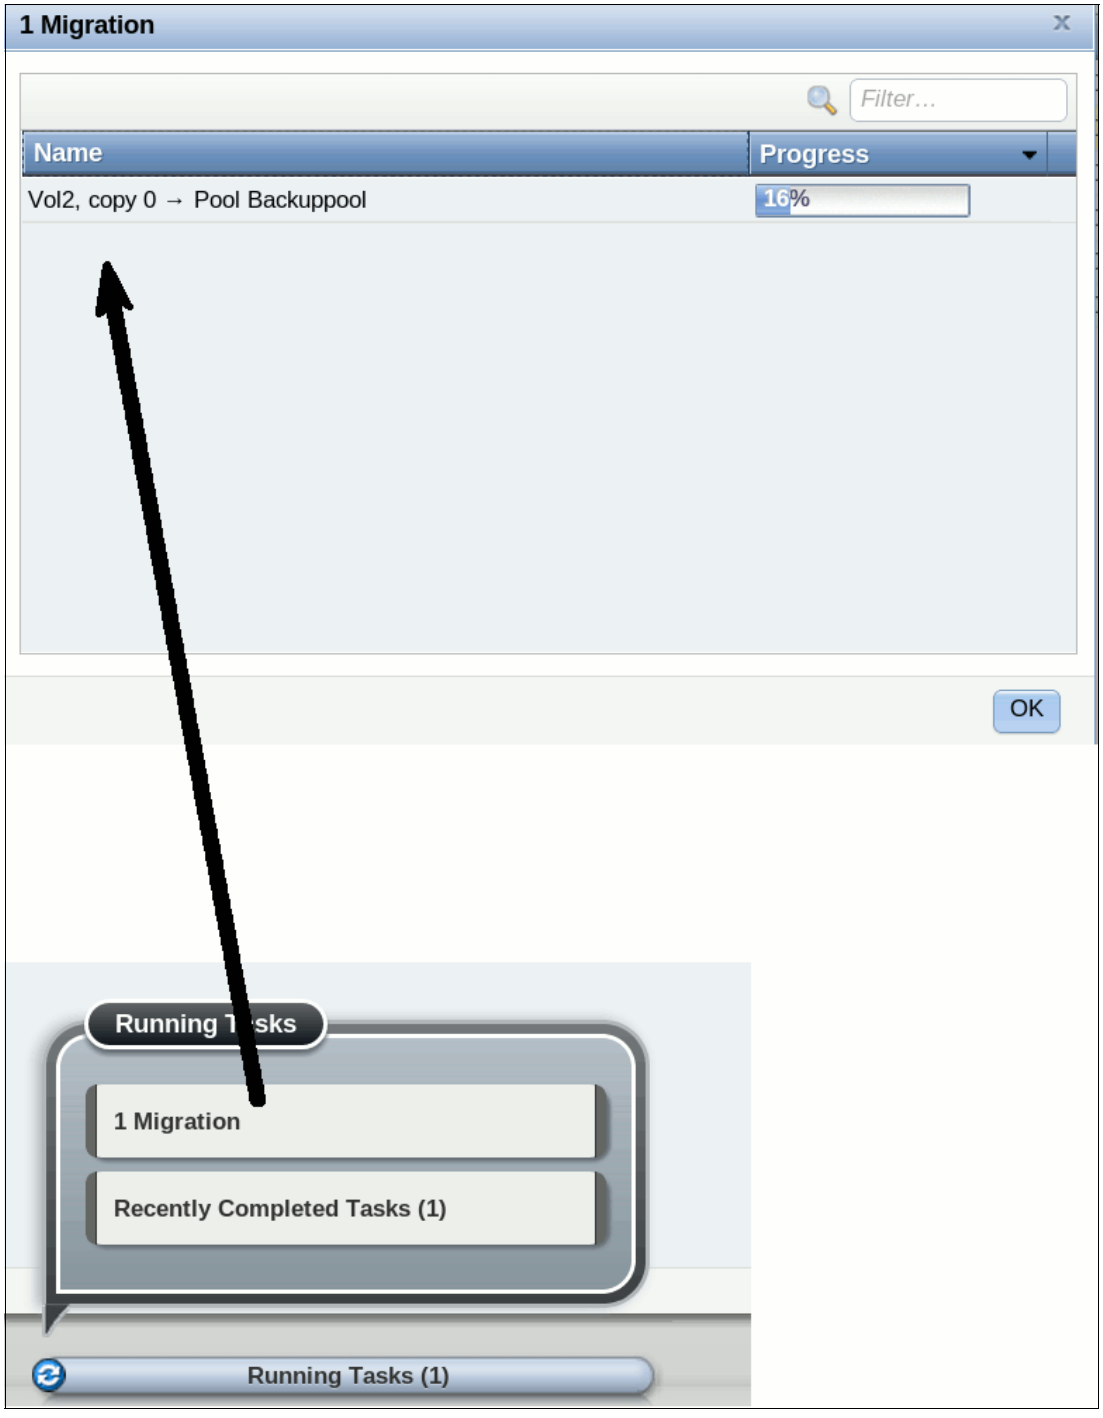

Figure 8-75 Migration progress

After the migration is complete, the volume is shown in the new storage pool. In Figure 8-76, you can see that vol4 moved from the BRONZE POOL storage pool to the NewPool storage pool.

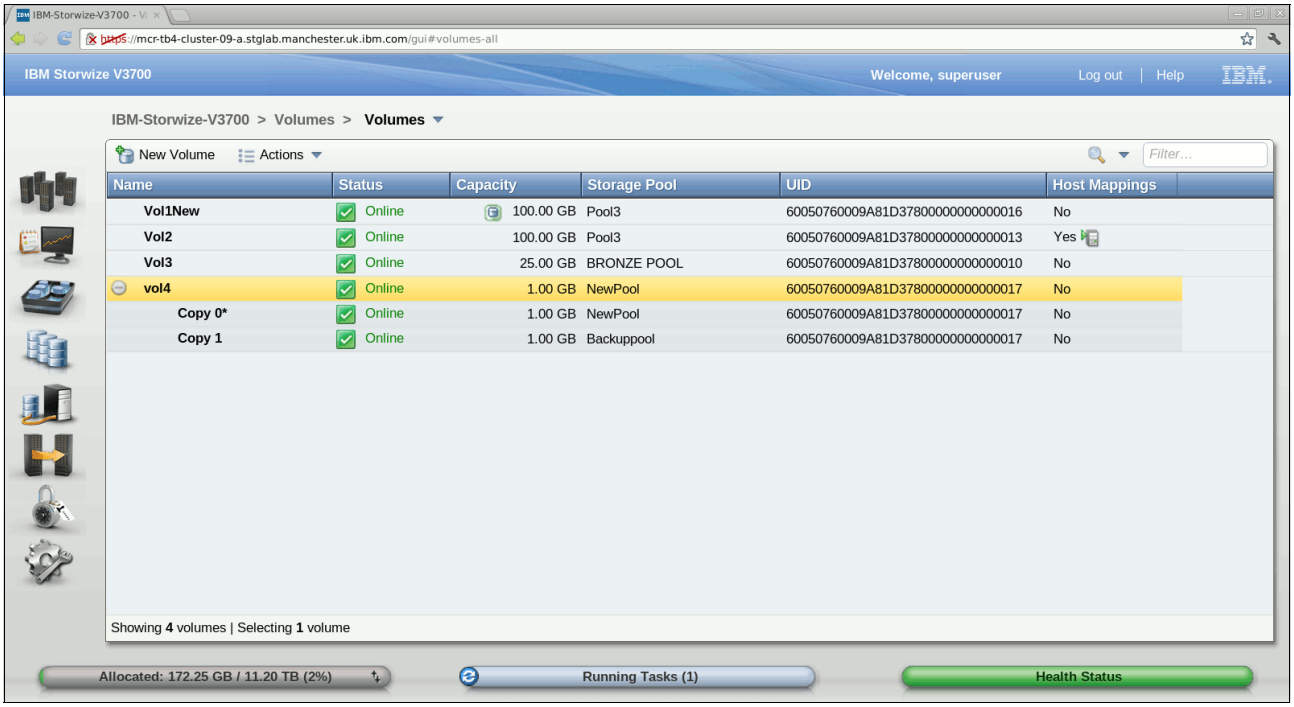

Figure 8-76 Migration complete

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

Another way to migrate volumes to another pool is by using the volume copy feature, as described in 8.6.5, “Migrating volumes by using the volume copy features” on page 388.

8.4.9 Deleting a volume

To delete a volume, select **Delete** (see Figure 8-58 on page 352). Enter the number of volumes that you want to delete and select the option if you want to force the deletion, as shown in Figure 8-77.

Important: You must force the deletion if the volume has host mappings or is used in FlashCopy mappings.

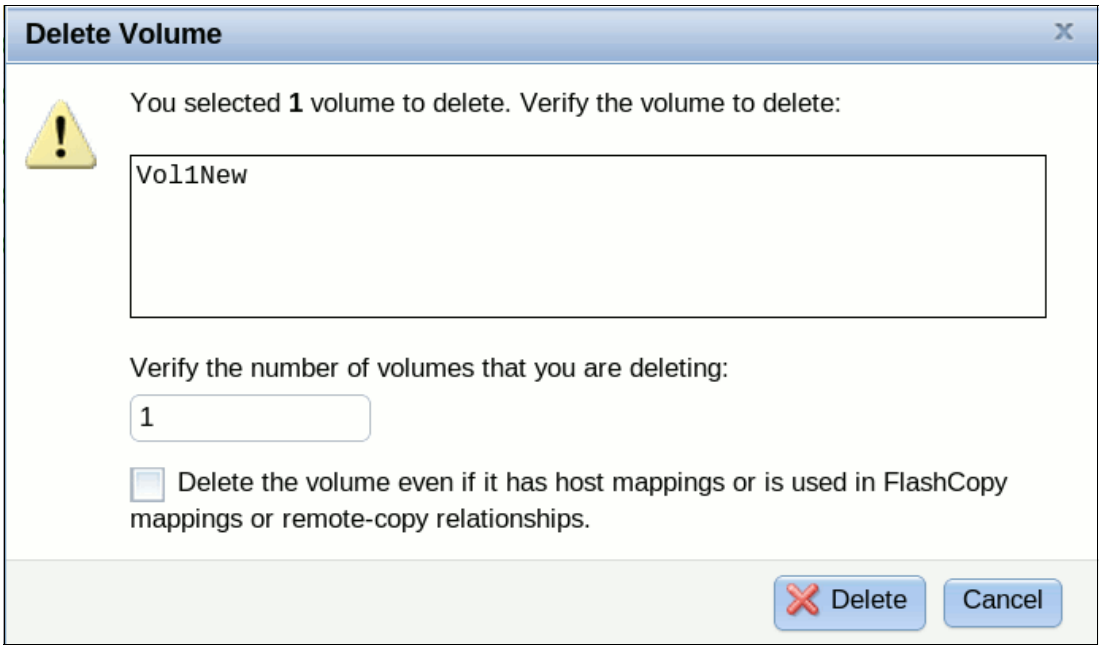

Figure 8-77 Delete volume

Click **Delete** and the volume is removed from the system, as shown in Figure 8-78.

Deleting volumes: This action removes all copies from your storage system and the data on the volume is lost. Before you perform this step, make sure that you do not require the volume or its data.

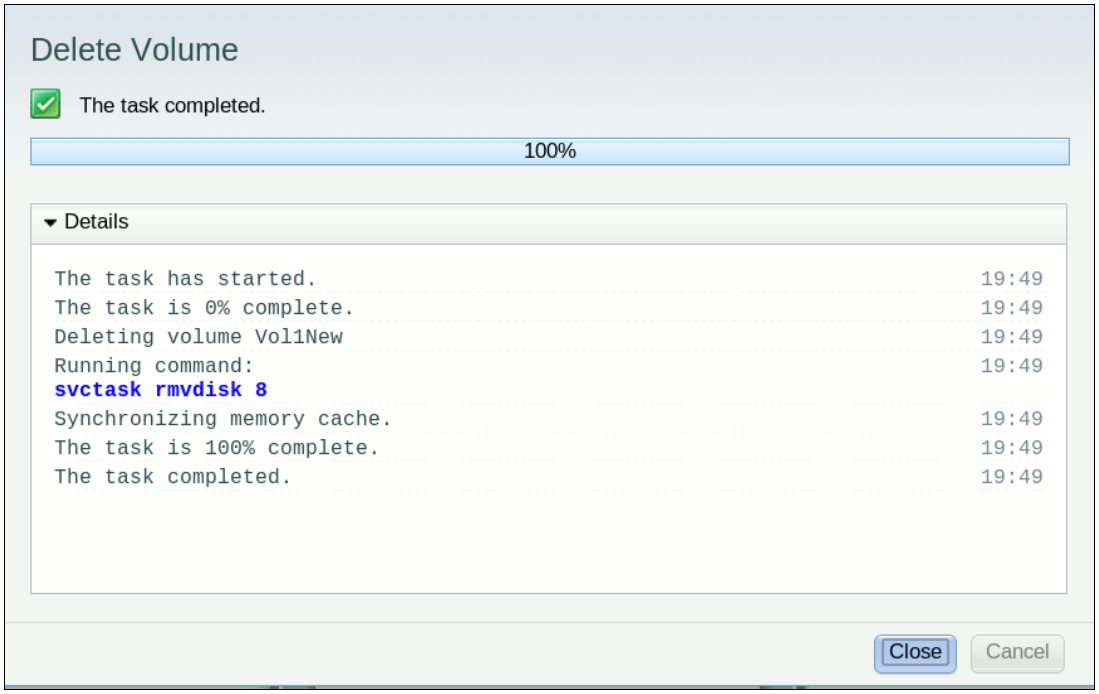

Figure 8-78 Volume deletion complete

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

8.5 Volume properties

To open the advanced view of a volume, select **Properties** (see Figure 8-58 on page 352), and the Volume Details window opens, as shown in Figure 8-79. In this window, the following tabs are available:

- Overview tab
- Host Maps tab
- Member MDisk tab

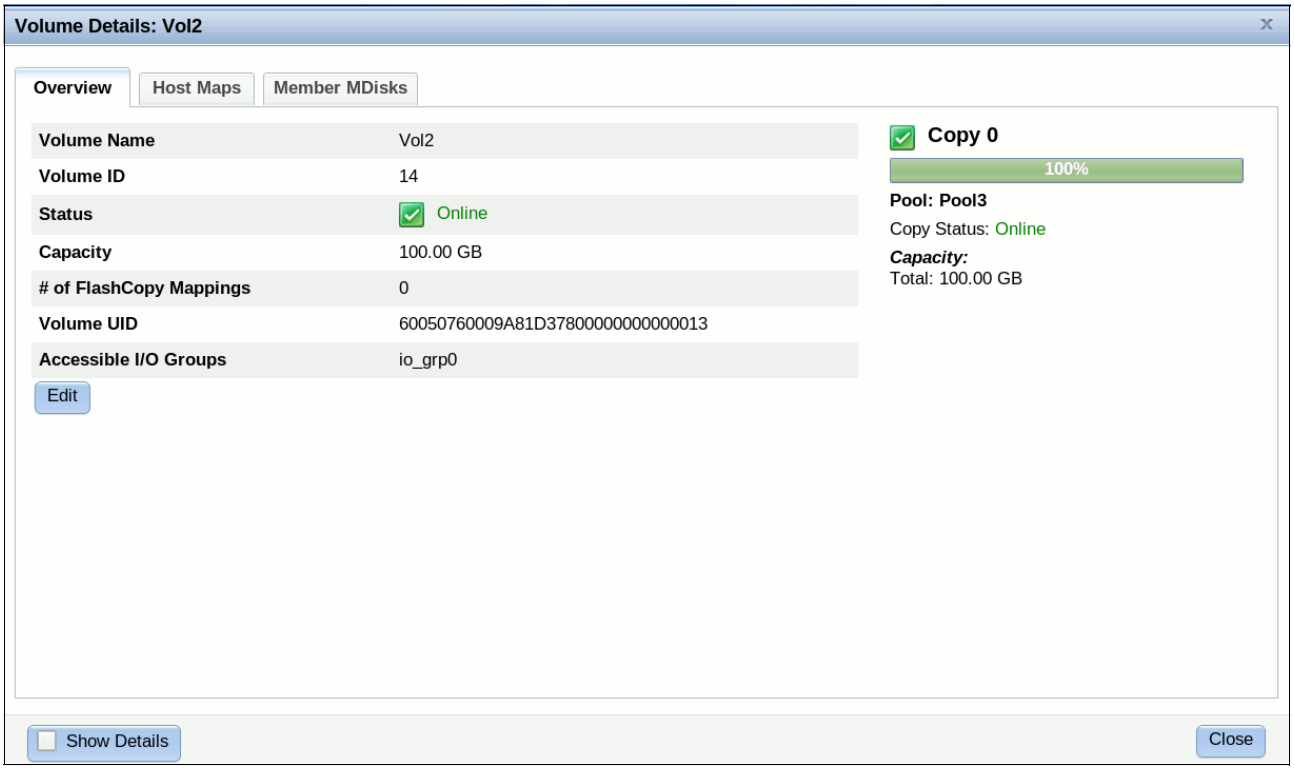

Figure 8-79 Volume Details Overview tab

8.5.1 Overview tab

The Overview tab that is shown in Figure 8-79 on page 367 gives you a basic overview of the volume properties. In the left part of the window, you find common volume properties. In the right part, you see information about the volume copies. To get a more detailed view, select the **Show Details** option in the lower left, as shown in Figure 8-80.

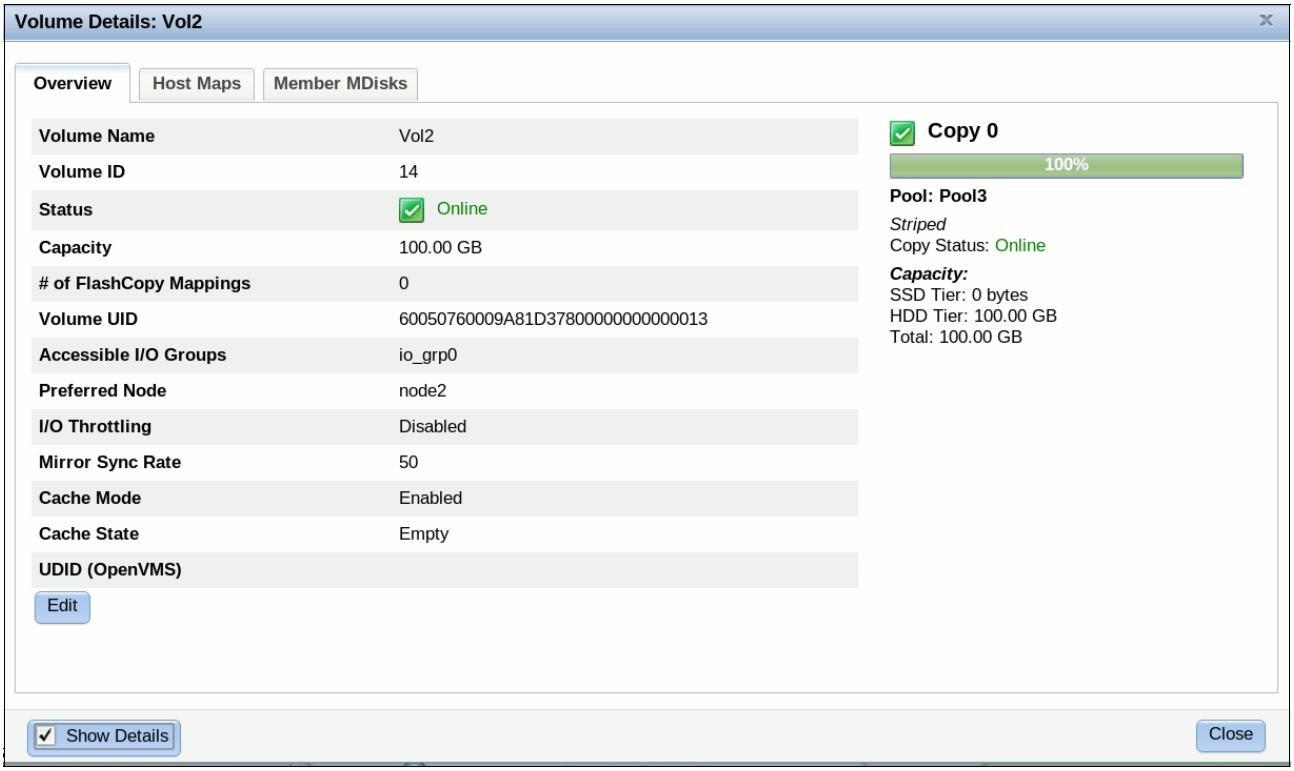

Figure 8-80 Show details

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 that 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: Number of existing FlashCopy relationships. For more information, see 8.9, “FlashCopy” on page 399.
  - Volume UID: The volume unique identifier.
  - I/O Group: Specifies the I/O group to which the volume belongs.
  - 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 483.

- 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 whether the cache is enabled or disabled for this volume.
- Cache State: Gives you feedback if open I/O requests are inside the cache that are 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), and the copy's status.
  - Capacity: Shows the allocated (used) and the virtual (Real) capacity, and the warning threshold and the set grain size.

If you want to edit any of these settings, click **Edit** and the window changes to the modify mode, as shown in Figure 8-81.

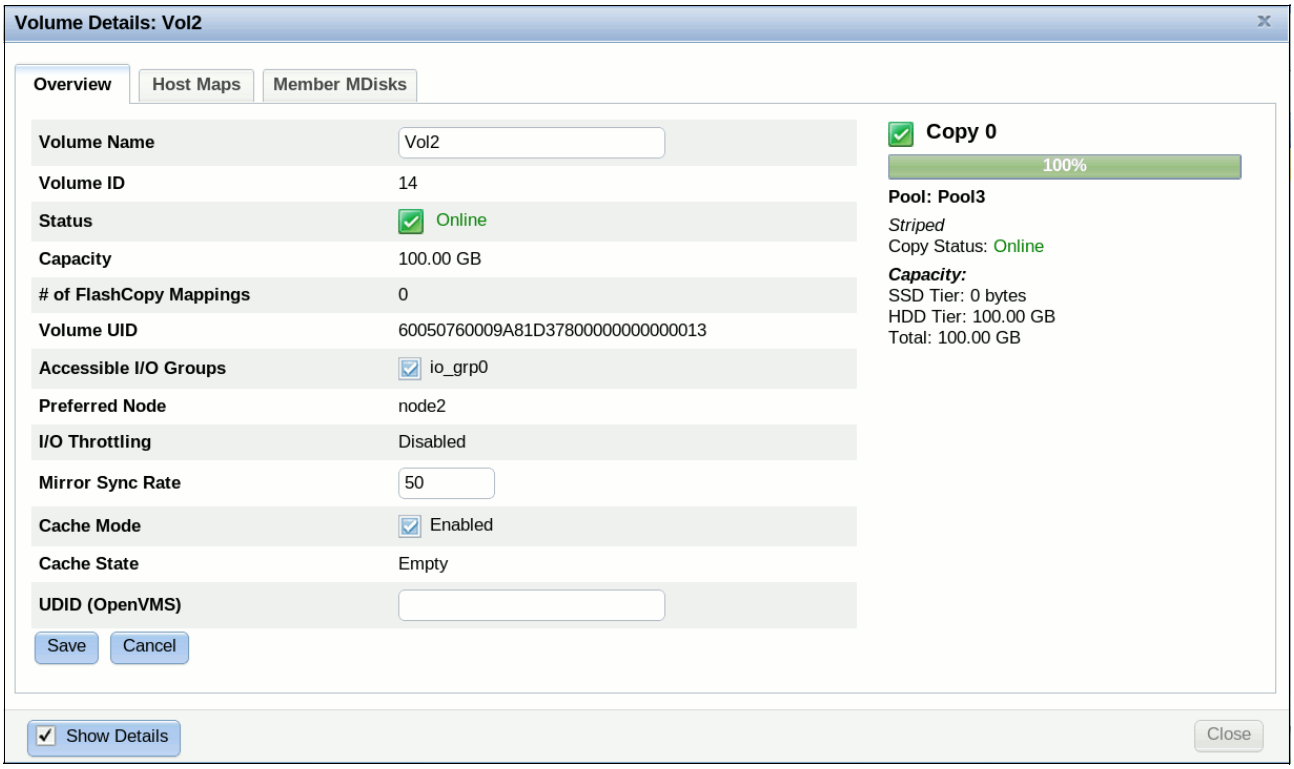

Figure 8-81 Modify Volume Details

Inside the Volume Details window, you can change the following properties:

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

Make any required changes and click **Save**. For this example, we change the mirror sync rate from 50 to 40. Figure 8-82 shows the volume details being edited.

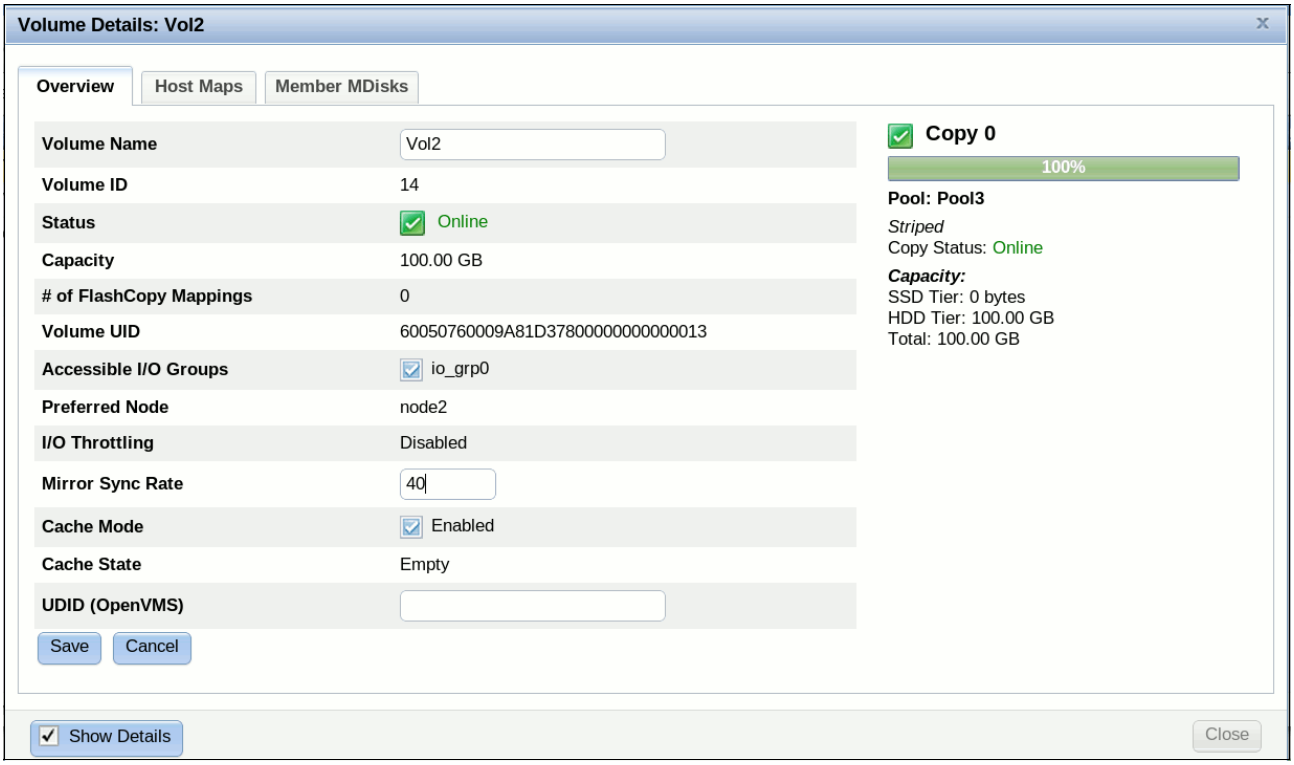

Figure 8-82 Edit volume details

The changes are applied to the system. Figure 8-83 shows the output after the volume properties are edited.

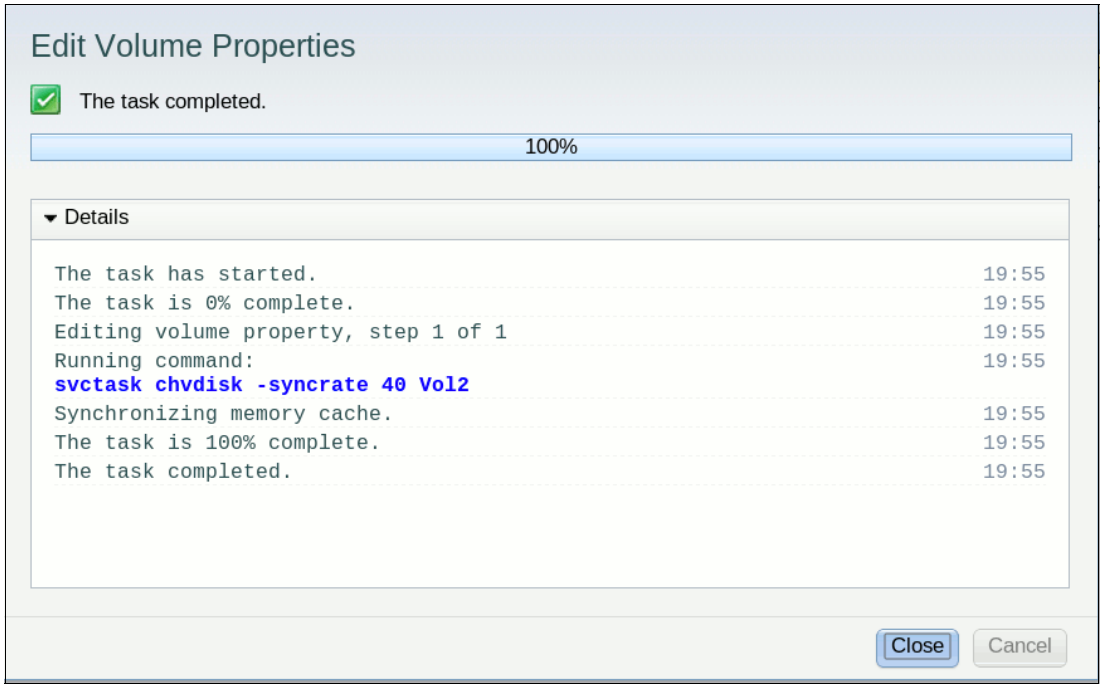

Figure 8-83 Edit Volume Properties completed

Mirror Sync Rate: 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 8-84. All hosts that are mapped to the selected volume are listed in this view.

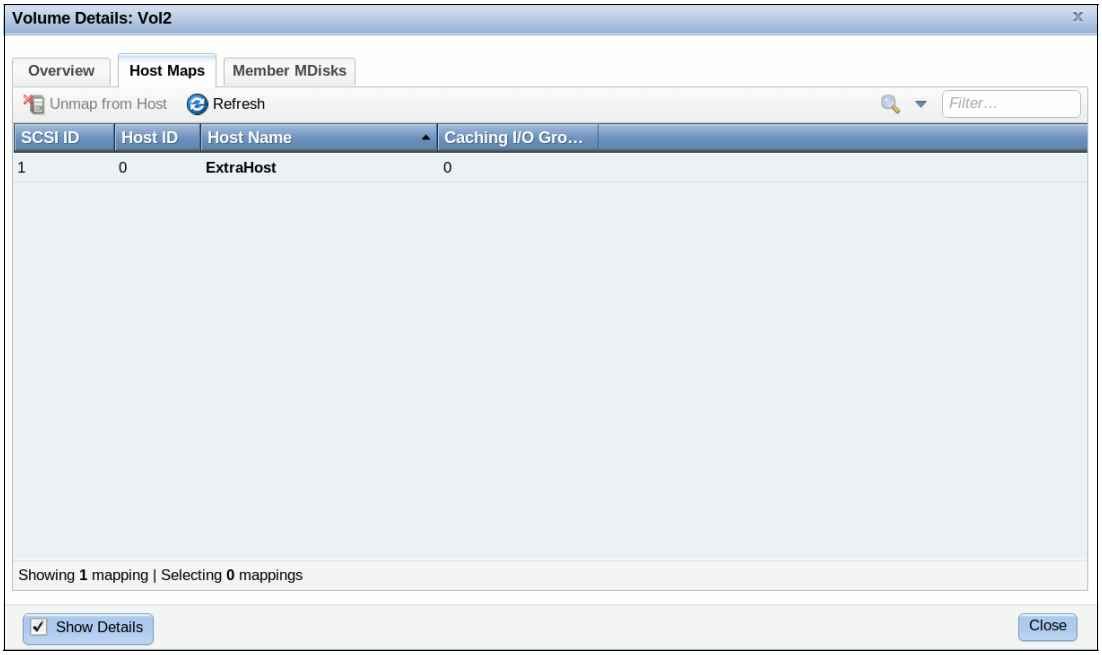

Figure 8-84 Host Maps

To unmap a host from the volume, highlight it and click **Unmap from Host**. Confirm the number of mappings to remove and click **Unmap**, as shown in Figure 8-85.

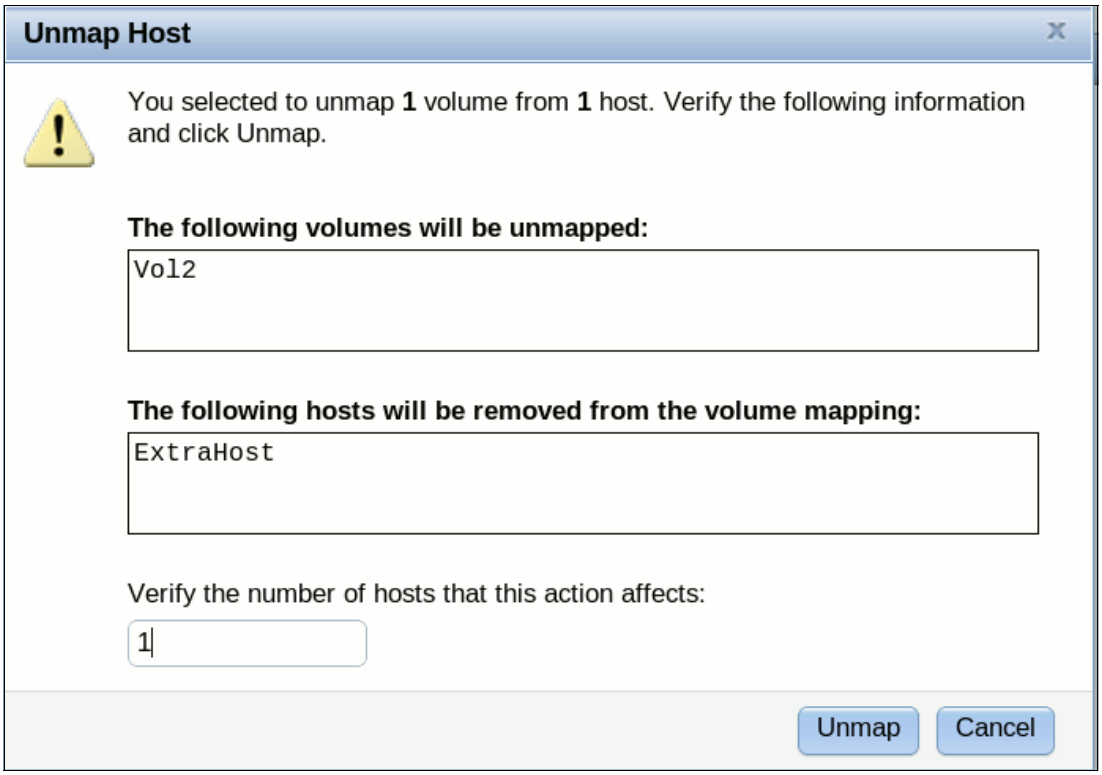

Figure 8-85 Unmap Host window

The changes are applied to the system, as shown in Figure 8-86. The selected host no longer has access to this volume. Click **Close** to return to the Host Maps window.

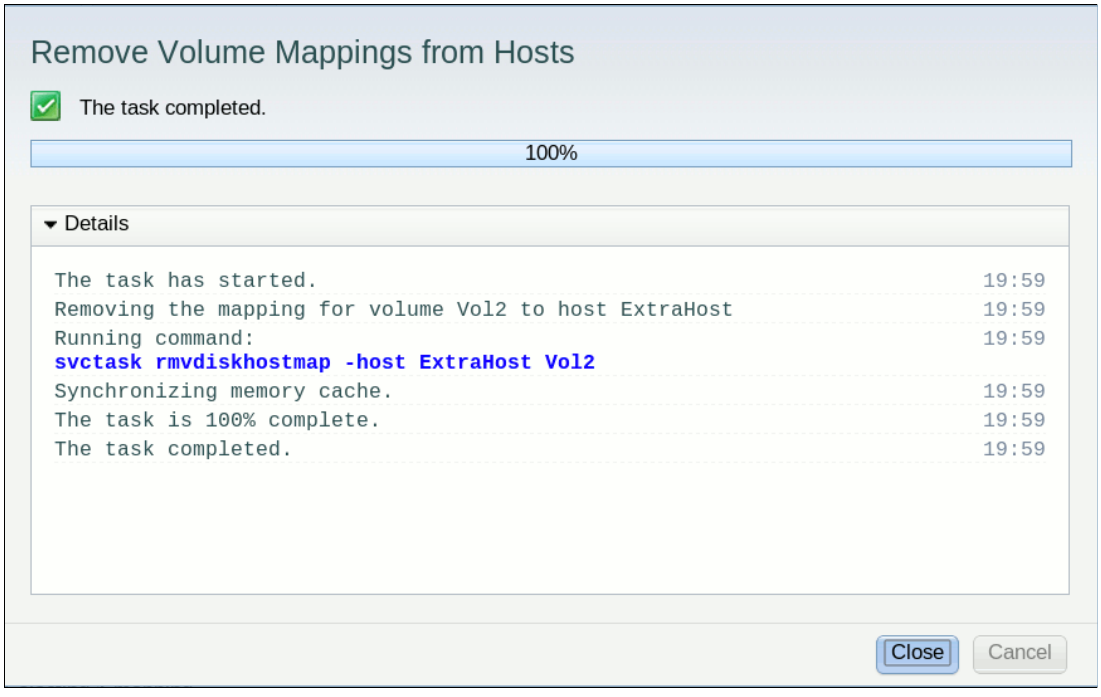

Figure 8-86 Remove Volume Mappings from Hosts window

8.5.3 Member MDisk tab

The third tab is Member MDisk, which lists all MDisks on which the volume is located. Select a copy and the associated MDisks appears in the window, as shown in Figure 8-87.

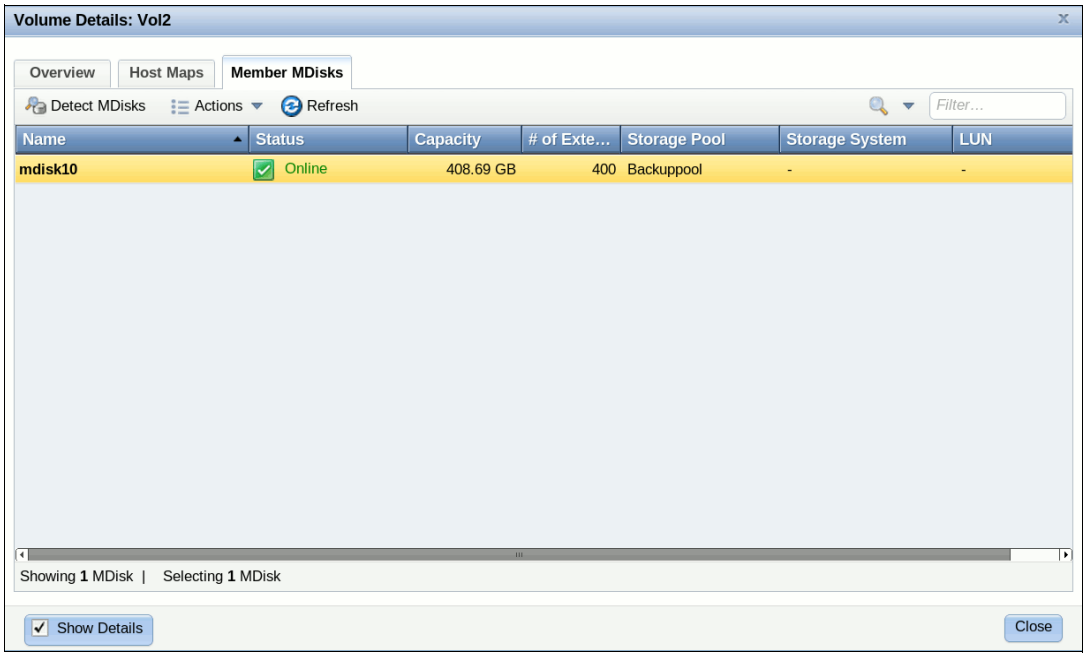

Figure 8-87 Member MDisk tab

Highlight an MDisk and click **Actions** to see a view of the available tasks, as shown in Figure 8-88. The tasks are described in Chapter 7, “Storage pools” on page 277.

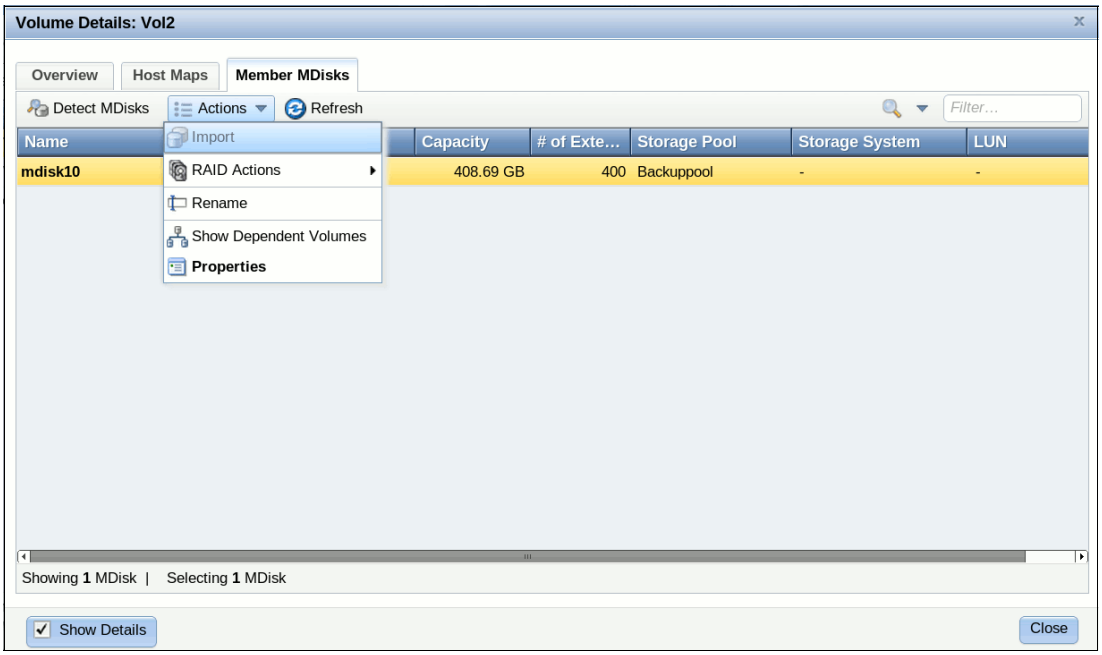

Figure 8-88 Actions

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

8.5.4 Adding a mirrored volume copy

If you have an existing volume that consists of only one copy, you can add a mirrored volume copy to it. This action creates another online copy of your volume. The second copy can be generic or thin-provisioned. It also can be created in any other storage pool.

To add a second copy, click **Volume Copy Actions** → **Add Mirrored Copy**, as shown in Figure 8-89. Select the storage pool to which the new copy should be migrated. If the new copy should be thin-provisioned, select the option and click **Add Copy**.

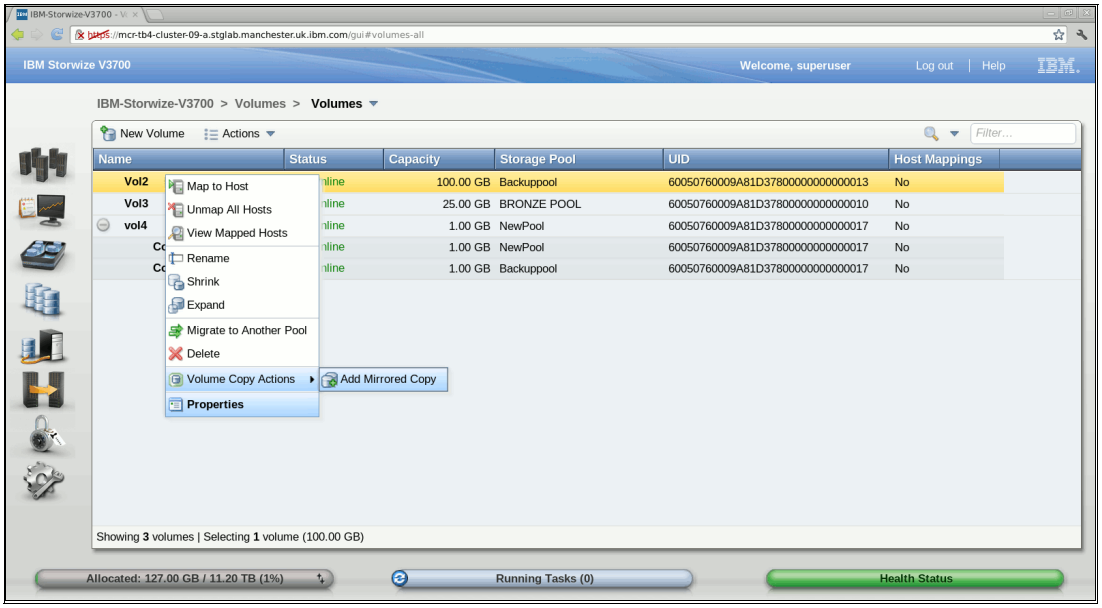

Figure 8-89 Add Mirrored Copy window

Figure 8-90 shows the storage pool selection page.

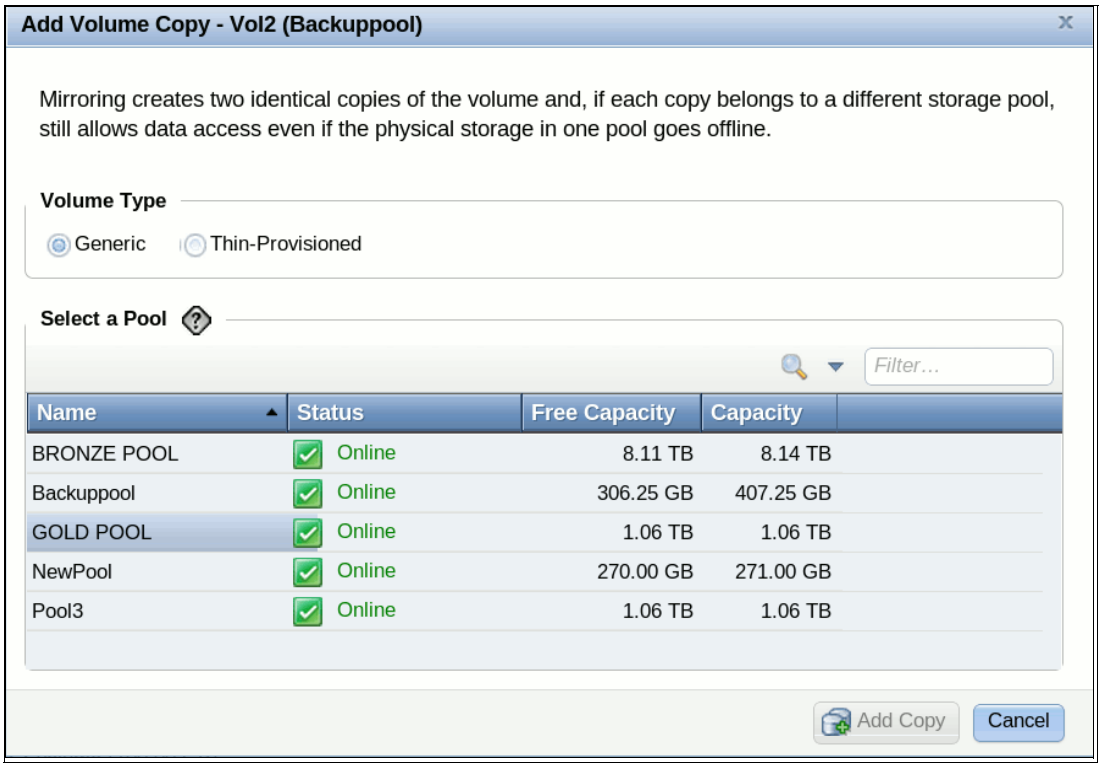

Figure 8-90 Select Storage Pool window

The new copy is created. Figure 8-91 shows the output when the volume copy is created.

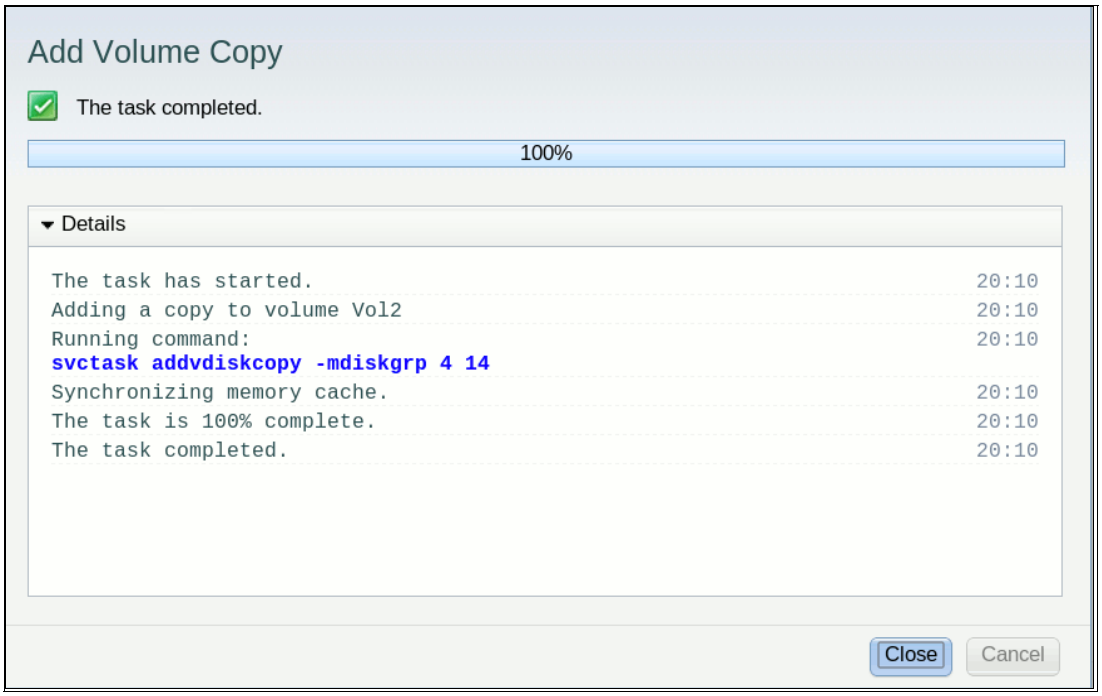

Figure 8-91 Add Volume Copy complete

The new copy is created and the data is synchronized as a background task. Figure 8-92 shows you that the volume that is named `vol2` now has two volume copies and there is one synchronization task running.

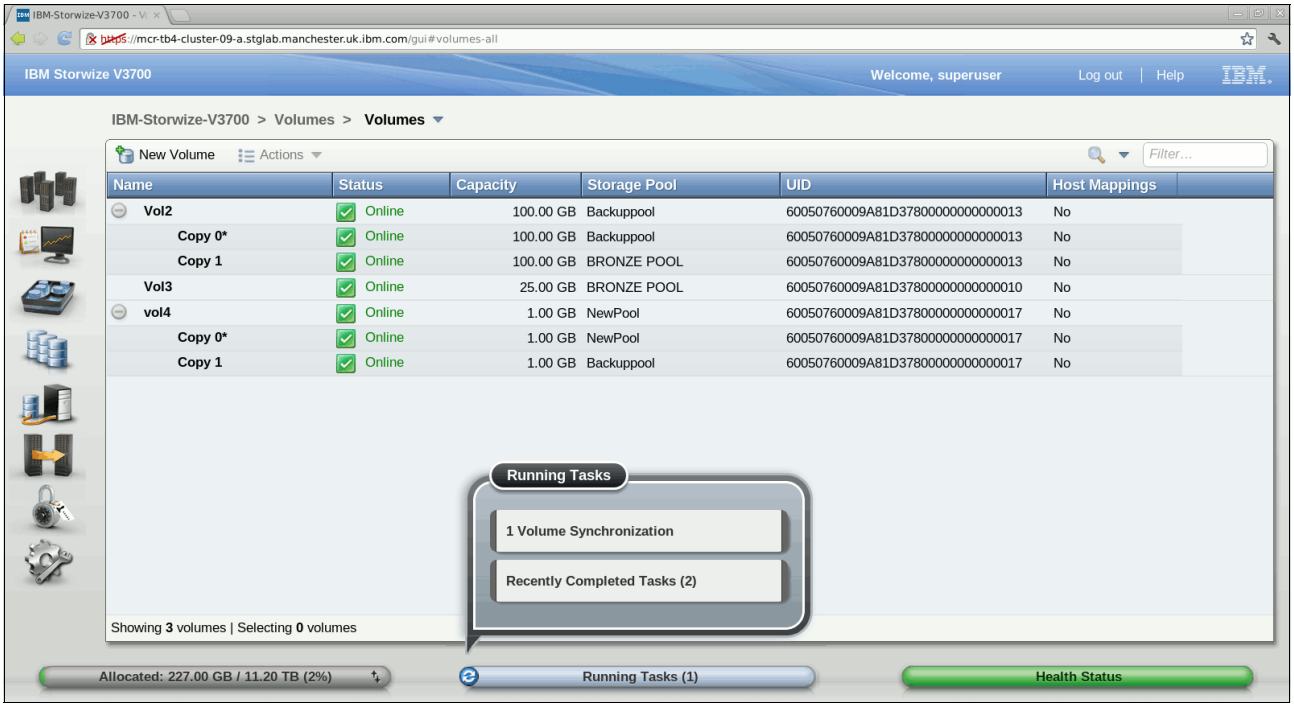

Figure 8-92 Volume copy synchronization

8.5.5 Editing thin-provisioned volume properties

Sections 8.4.6, “Shrinking a volume” on page 358 and 8.4.7, “Expanding a volume” on page 359 described how to modify the volume size that is presented to a host. However, if you have a thin-provisioned volume, you also can edit the allocated size and the warning thresholds. To accomplish this task, select the volume copy, then select **Actions** → **Thin Provisioned** (see Figure 8-59 on page 353). The following options are available (as shown in Figure 8-93):

- Shrink
- Expand
- Edit Properties

This changes are only to the internal storage, so you do not have to change your host.

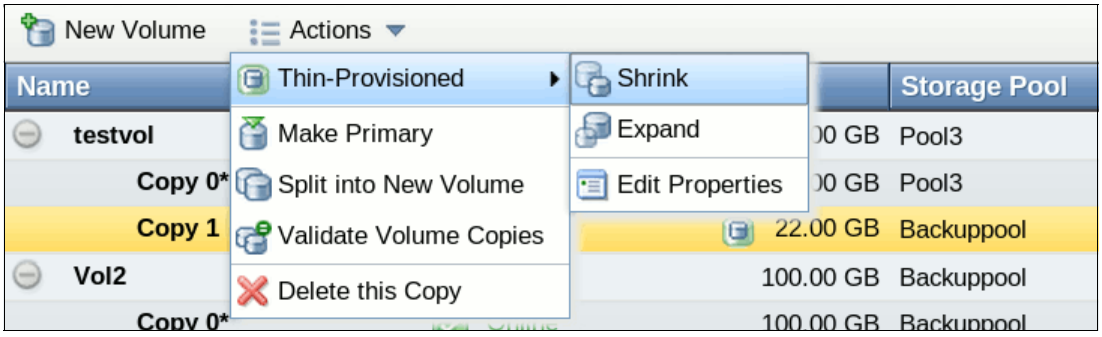

Figure 8-93 Thin-Provisioned menu options

Shrinking thin-provisioned space

Select **Shrink** (as shown in Figure 8-93 on page 376) to reduce the allocated space of a thin-provisioned volume. Enter how much the volume should shrink or the new final size, as shown in Figure 8-94, and click **Shrink**.

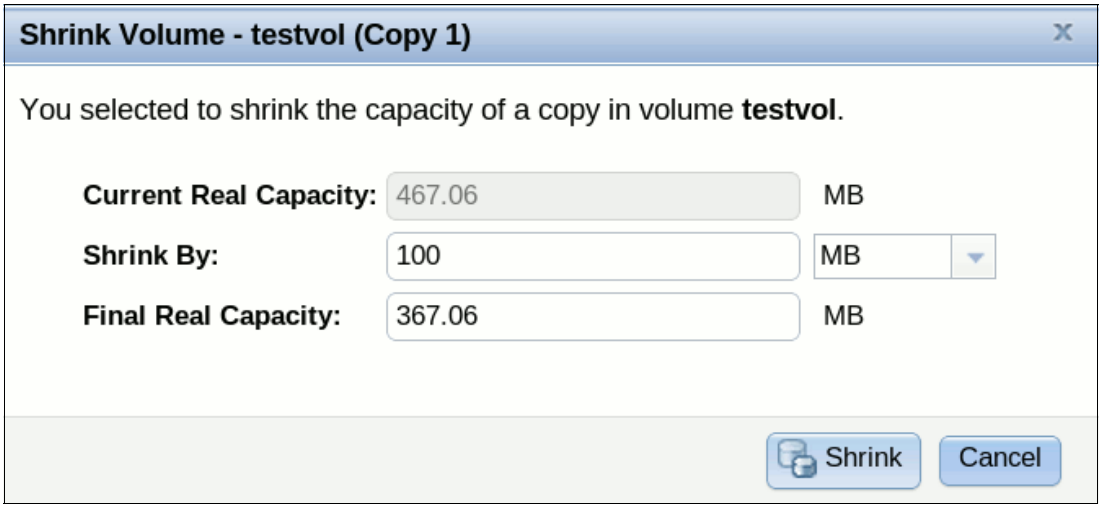

Figure 8-94 Shrink Volume window

After the task completes, click **Close**. Figure 8-95 shows the output from shrinking a volume.

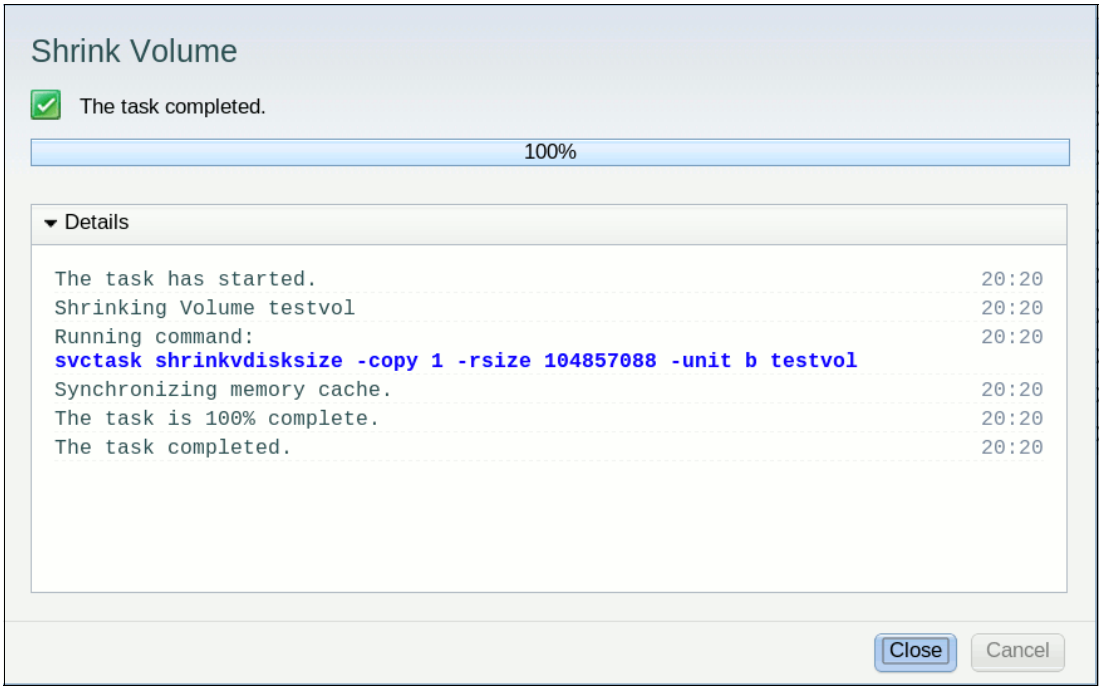

Figure 8-95 Shrink Volume complete

The allocated space of the thin-provisioned volume is now reduced.

Deallocating extents: You can deallocate only extents that do not have stored data on them. If the space is allocated because there is data on them, you cannot shrink the allocated space.

Expanding thin-provisioned space

To expand the allocated space of a thin-provisioned volume, select **Expand** (see Figure 8-93 on page 376). Enter how much space should be allocated or the new final size and click **Expand**, as shown in Figure 8-96.

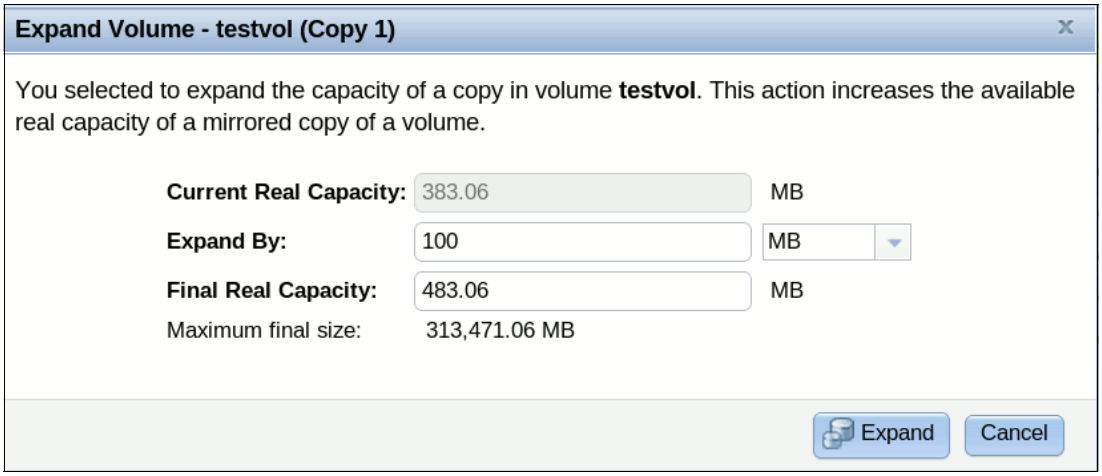

Figure 8-96 Expand Volume window

The new space is now allocated. Click **Close**. Figure 8-97 shows the output when a volume is expanded.

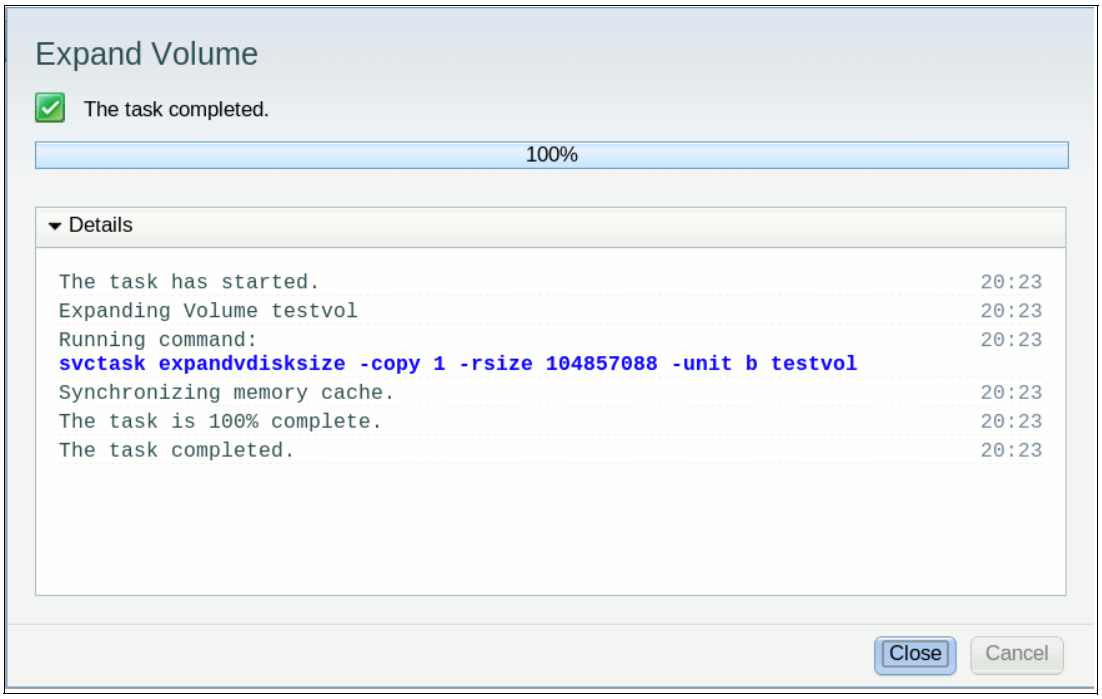

Figure 8-97 Expand Volume complete

Editing thin-provisioned properties

To edit thin-provisioned properties, select **Edit Properties** (see Figure 8-93 on page 376). Edit the settings, if required, as shown in Figure 8-98. Click **OK** to apply the changes.

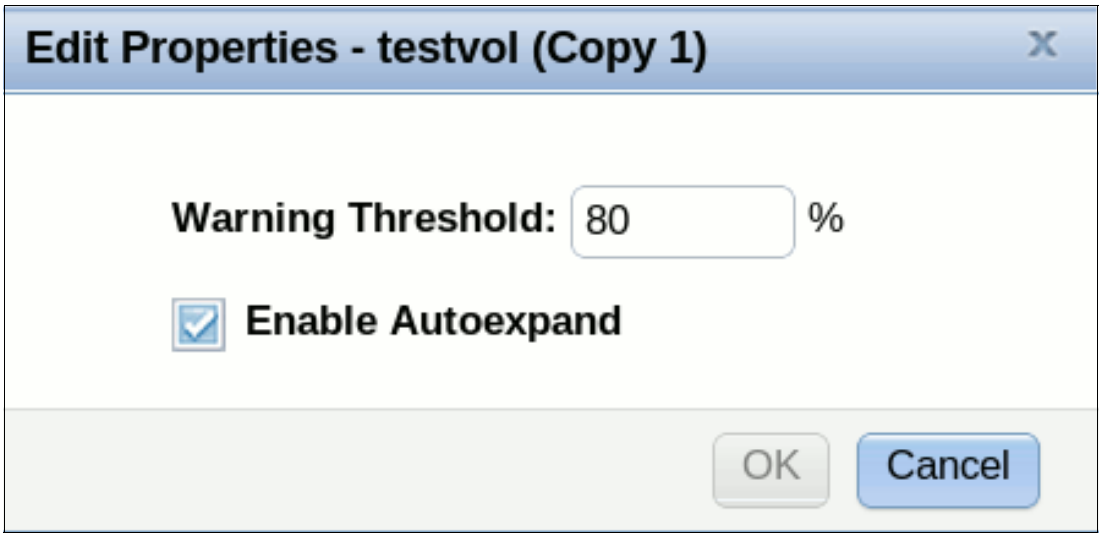

Figure 8-98 Edit Properties

After the task completes, click **Close** (as shown in Figure 8-99) to return to the All Volumes window.

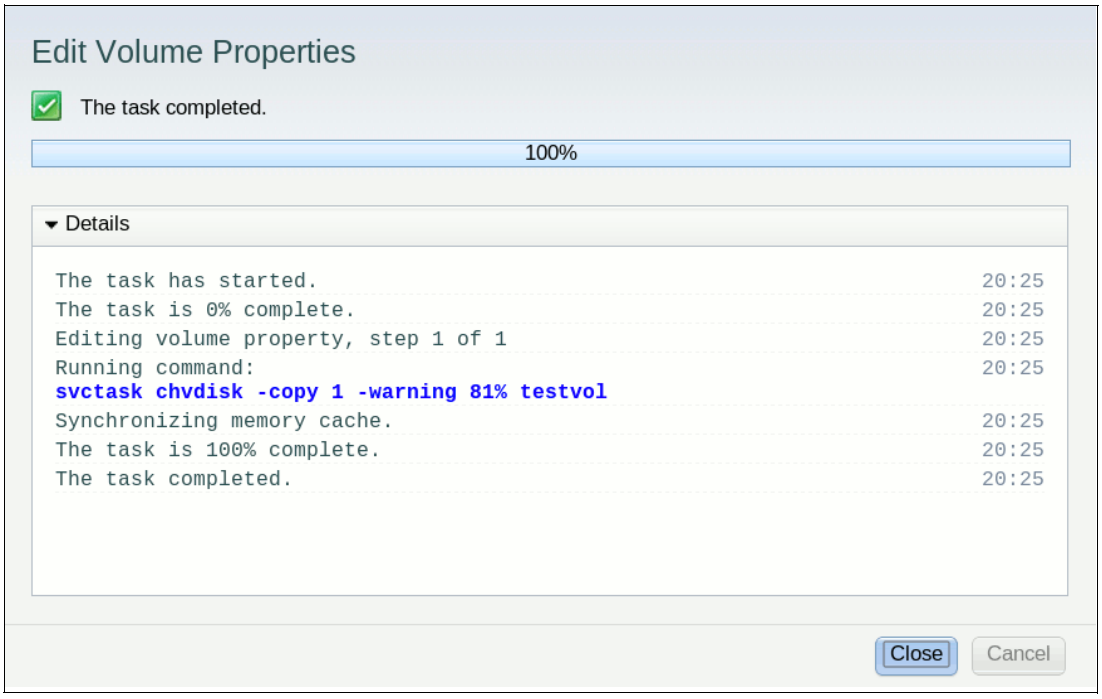

Figure 8-99 Edit Volume Properties completed

8.6 Advanced volume copy functions

In 8.4.1, “Advanced volume functions” on page 351, we described all 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. If you expand the volume and highlight a copy, the following volume copy actions become available, as shown in Figure 8-100:

- Thin-provisioned
- Split into New Volume
- Delete Copy option

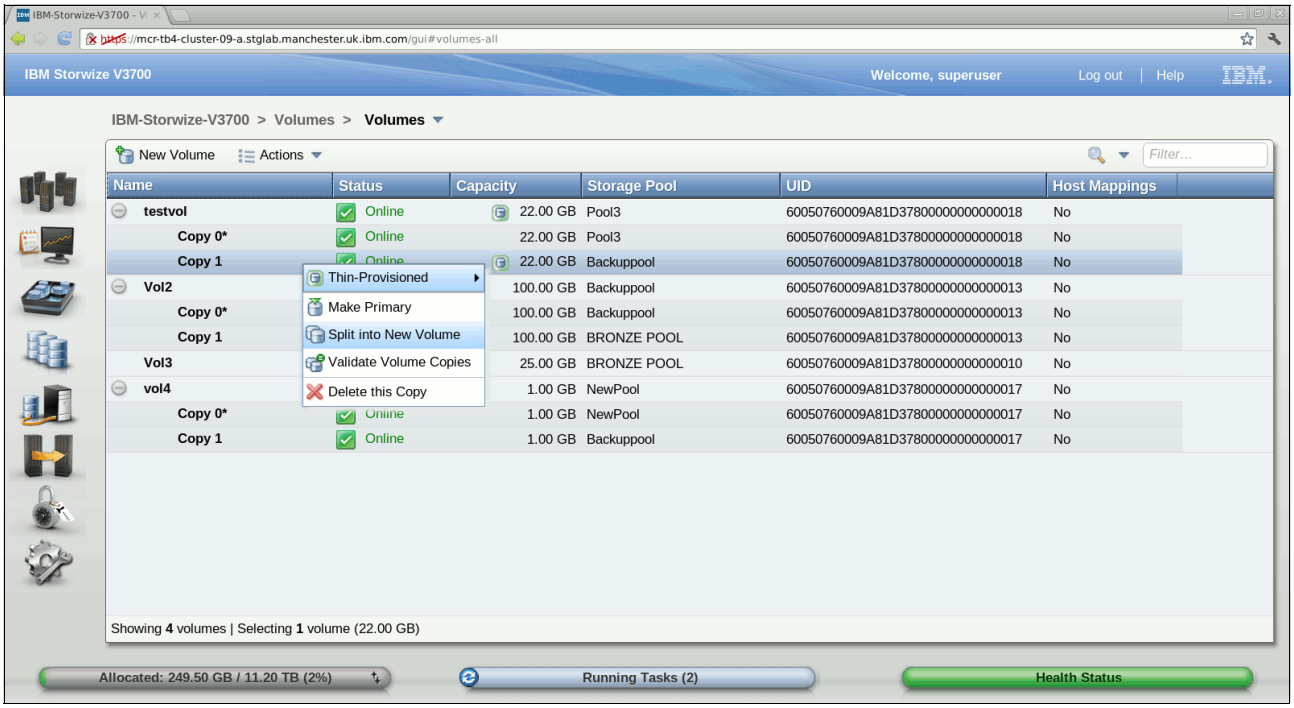

Figure 8-100 Volume copy actions

If you review the volume copies that are shown in Figure 8-100, you notice that one of the copies has a star displayed next to its name, as shown in Figure 8-101.

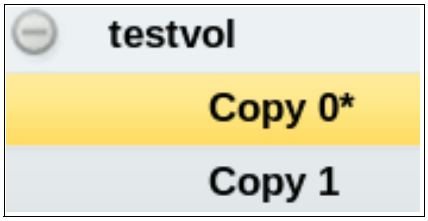

Figure 8-101 Volume copy names

Each volume has a primary and a secondary copy, and the star indicates the primary copy. The two copies always are synchronized, which means that all writes are destaged to both copies, but all reads are always done from the primary copy. By default, the primary and secondary copies always switch between Copy 0 and Copy 1 during creation to balance the reads across your storage pools. However, you should make sure that the I/O to the primary copies is customized to the performance of all your storage pools. Therefore, you can change the roles of your copies.

To accomplish this task, highlight the secondary copy, right-click it, and click **Actions**. The Make Primary task appears. 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. Writes must complete to both copies before the volume can provide acknowledgment to the host that the write completed successfully.

Figure 8-102 shows the secondary copy Actions menu.

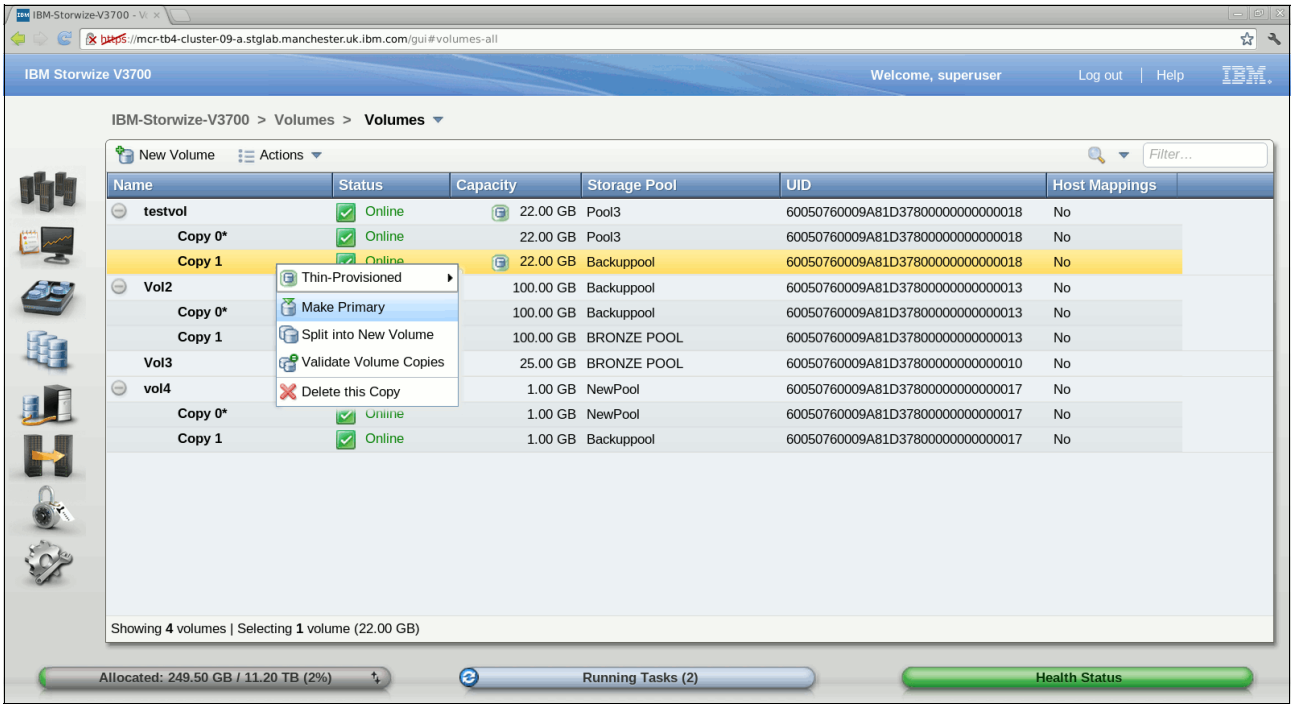

Figure 8-102 Make primary

If you need high read performance only, another possibility is to place the primary copy in an SSD pool and the secondary copy in a normal disk 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 also possible to migrate the copies online 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 360.

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

Figure 8-103 shows the output after volume properties are edited.

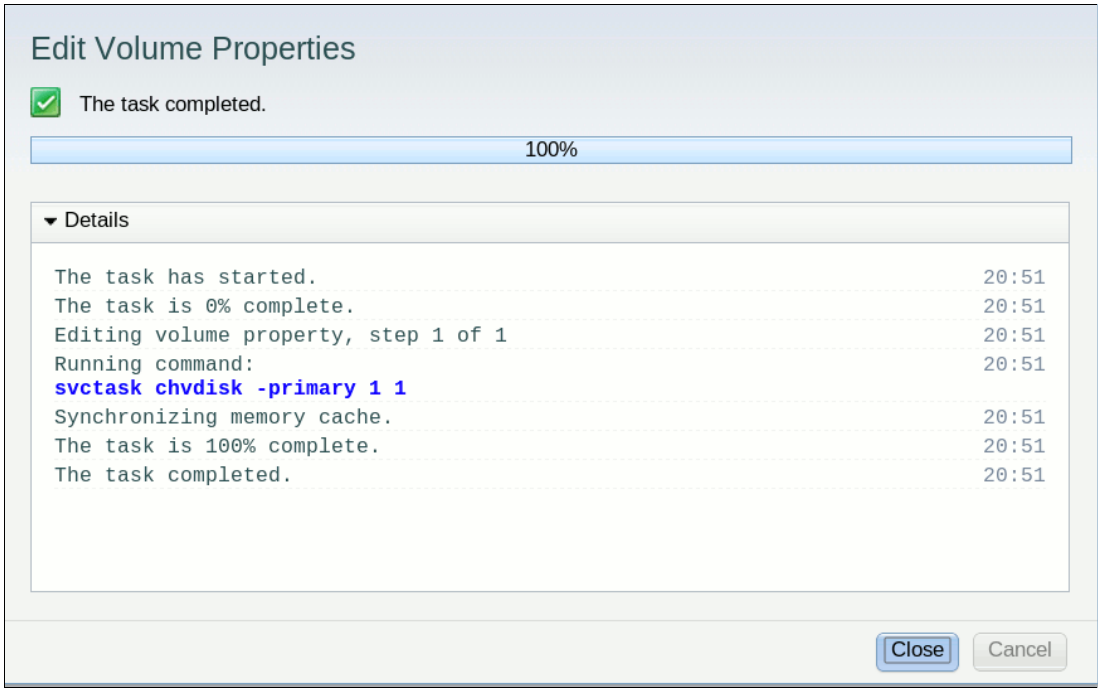

Figure 8-103 Edit Volume Properties complete

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 388.

8.6.1 Thin-provisioned

This menu item includes the same functions that are described in “Shrinking thin-provisioned space” on page 377, “Expanding thin-provisioned space” on page 378, and “Editing thin-provisioned properties” on page 379. You can specify the same settings for each volume copy.

Figure 8-104 shows the Thin-Provisioned menu item.

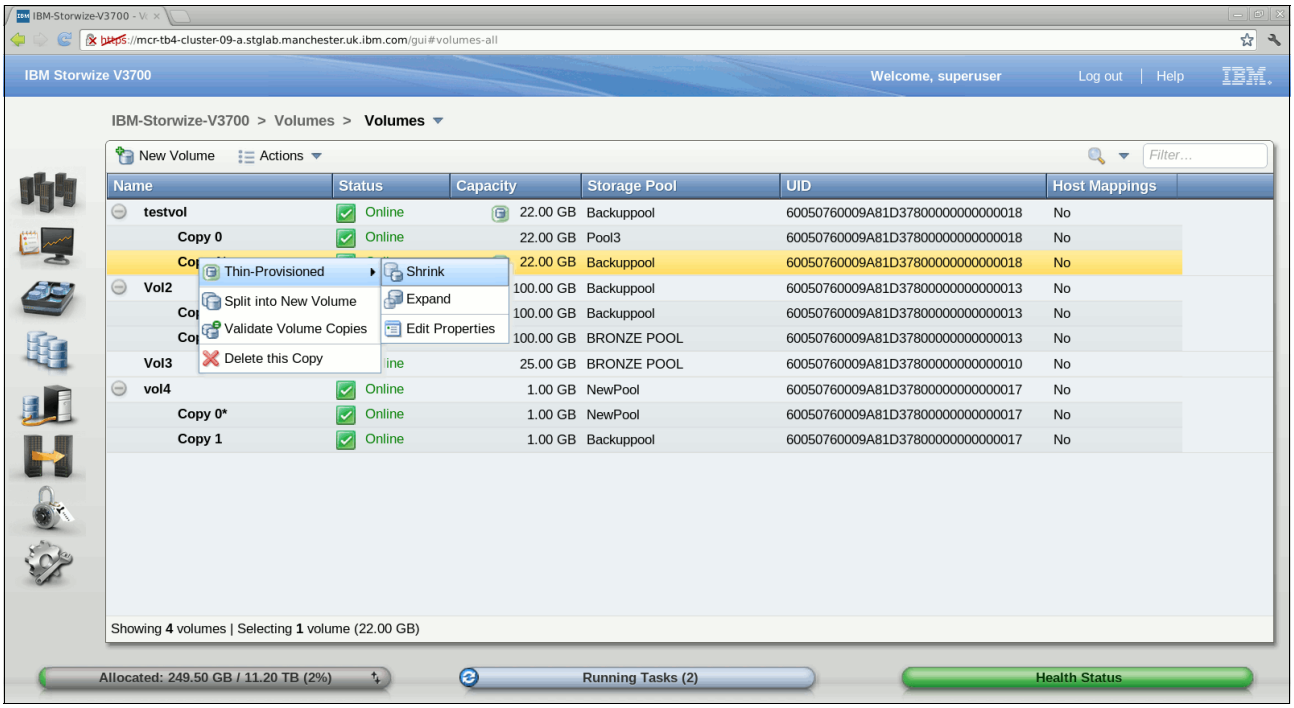

Figure 8-104 Thin-Provisioned menu item

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 could 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 is to shut down the hosts before splitting a copy.

To split a copy, click **Split into New Volume** (see Figure 8-100 on page 380). If you perform this action on the primary copy, the remaining secondary copy automatically becomes the primary for the source volume. Enter a name for the new volume and click **Split Volume Copy**, as shown in Figure 8-105.

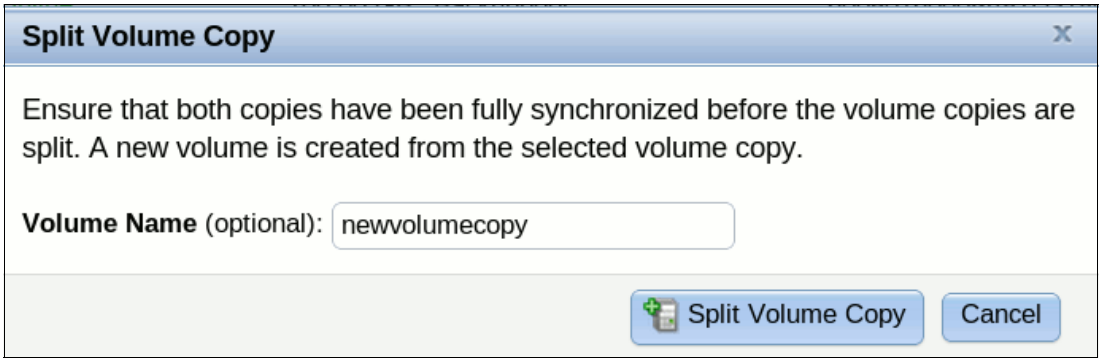

Figure 8-105 Split Volume Copy

After the task completes, click **Close** to return to the All Volumes window, where the copy appears as a new volume that can be mapped to a host, as shown in Figure 8-106.

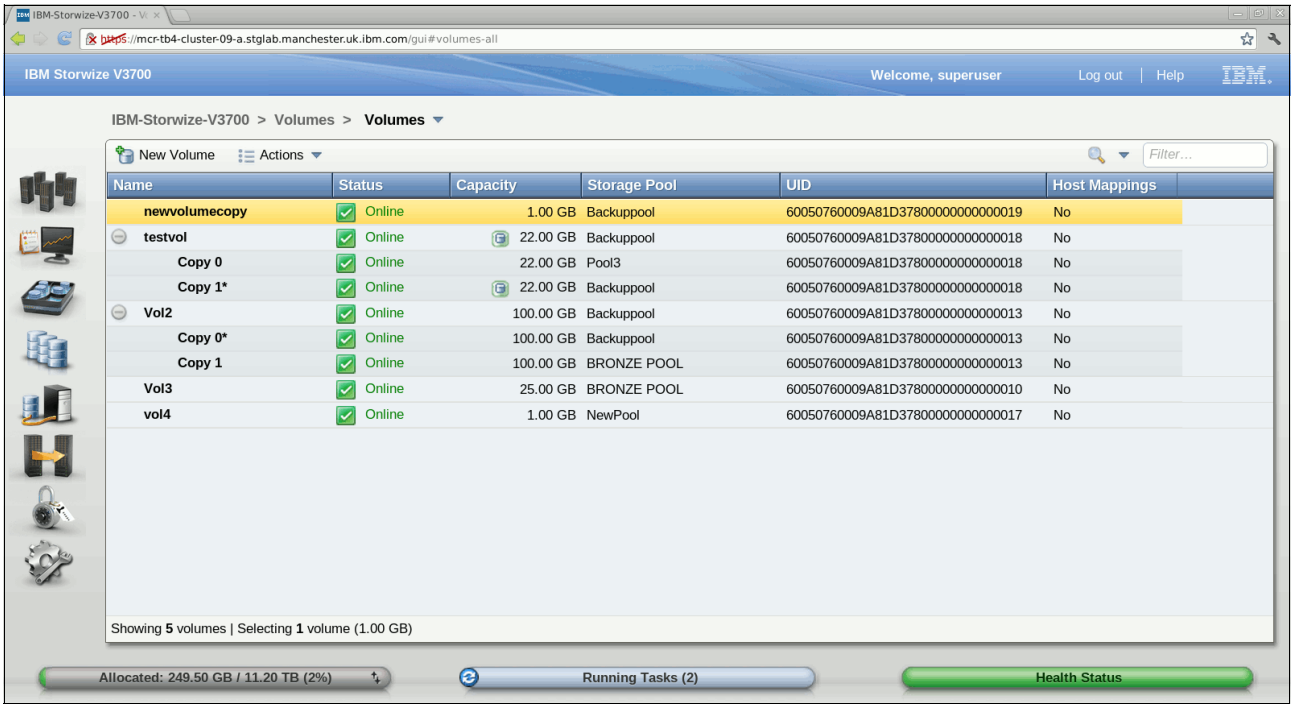

Figure 8-106 All Volumes: New volume from split copy

8.6.3 Validate Volume Copies option

To validate the copies of a mirrored volume, complete the following steps:

1. Select Validate Volume Copies (see Figure 8-100 on page 380). The Validate Volume Copies window opens, as shown in Figure 8-107.

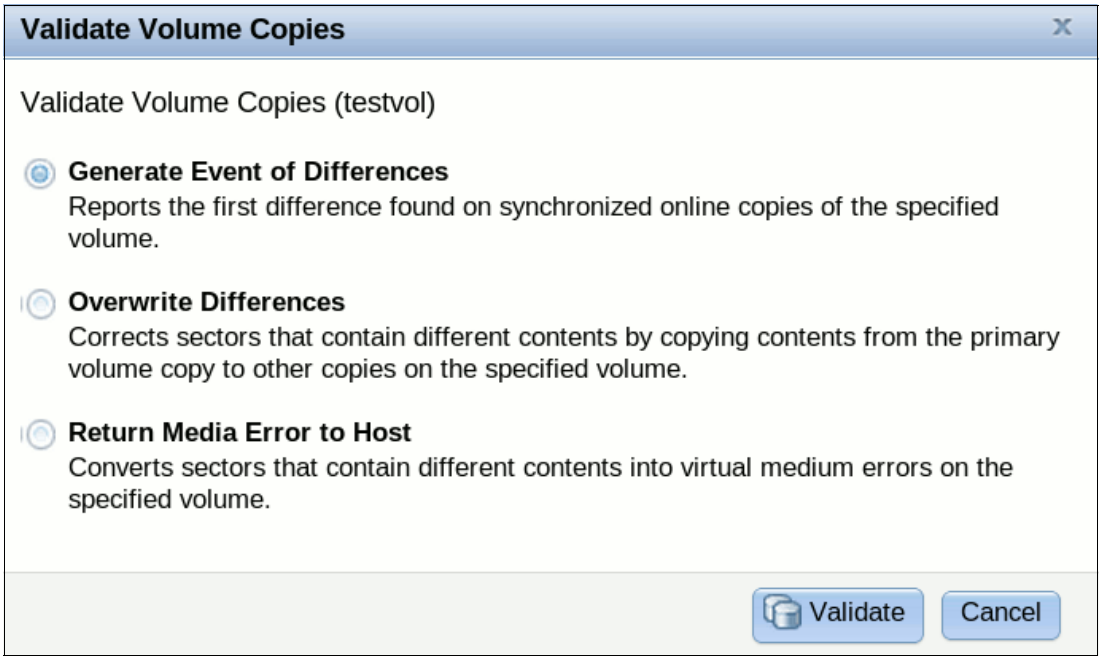

Figure 8-107 Validate Volume Copies

   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, which 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 volumes 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 8-108 shows the output when a volume copy is validated.

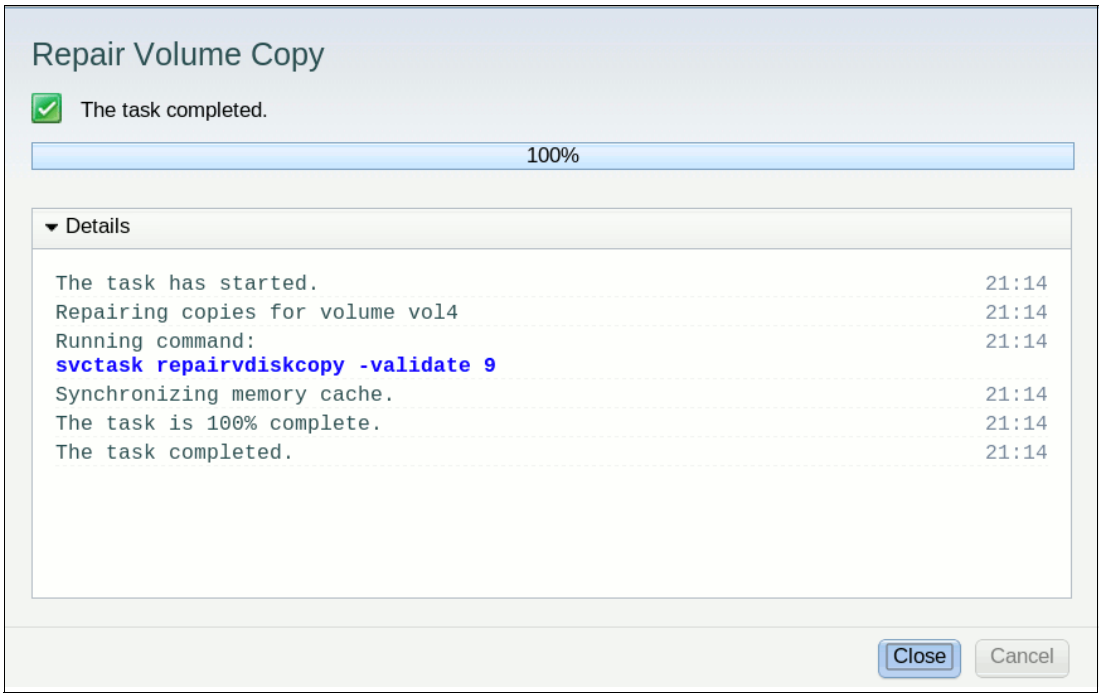

Figure 8-108 Volume copy validation output

The validation process takes some time, depending on the volume size. You can check the status in the Running Tasks window, as shown in Figure 8-109.

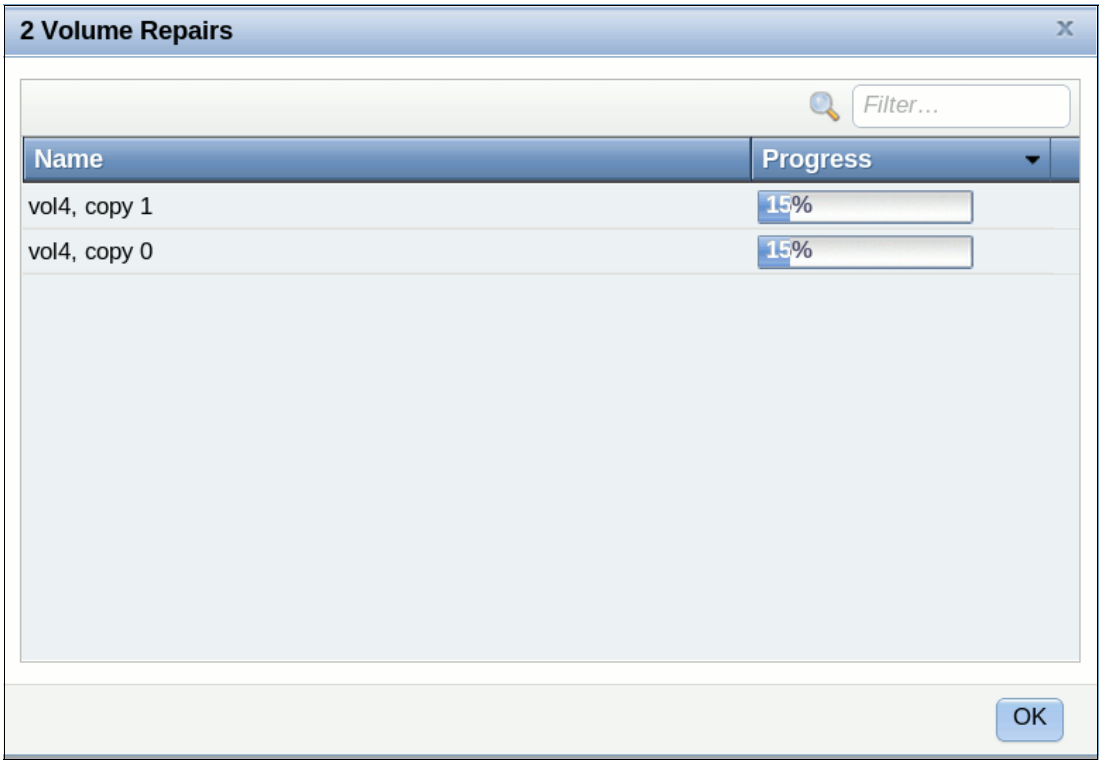

Figure 8-109 Validate Volume Copies: Running Tasks

8.6.4 Delete Copy option

Select **Delete** (see Figure 8-100 on page 380) 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 **OK**, as shown in Figure 8-110.

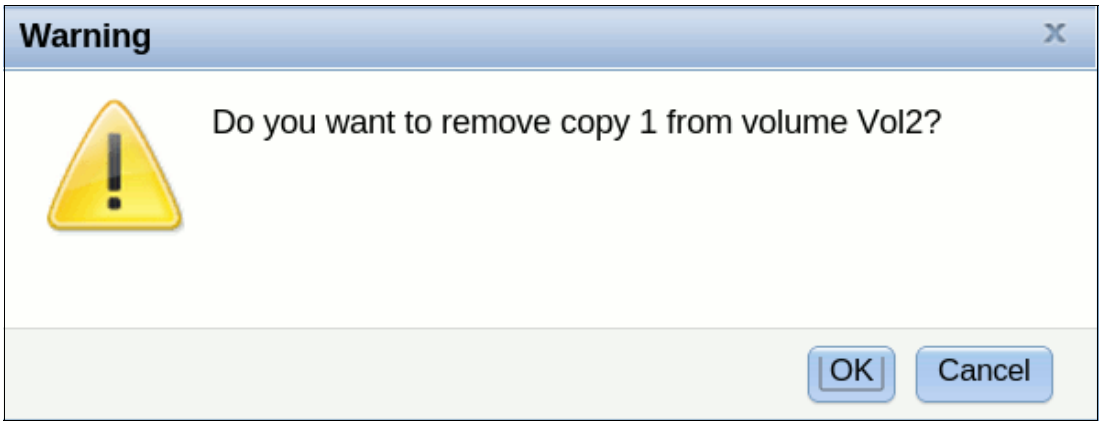

Figure 8-110 Delete a copy

The copy is deleted. Click **Close** to return to the All Volumes window. Figure 8-111 shows the output from deleting a volume copy.

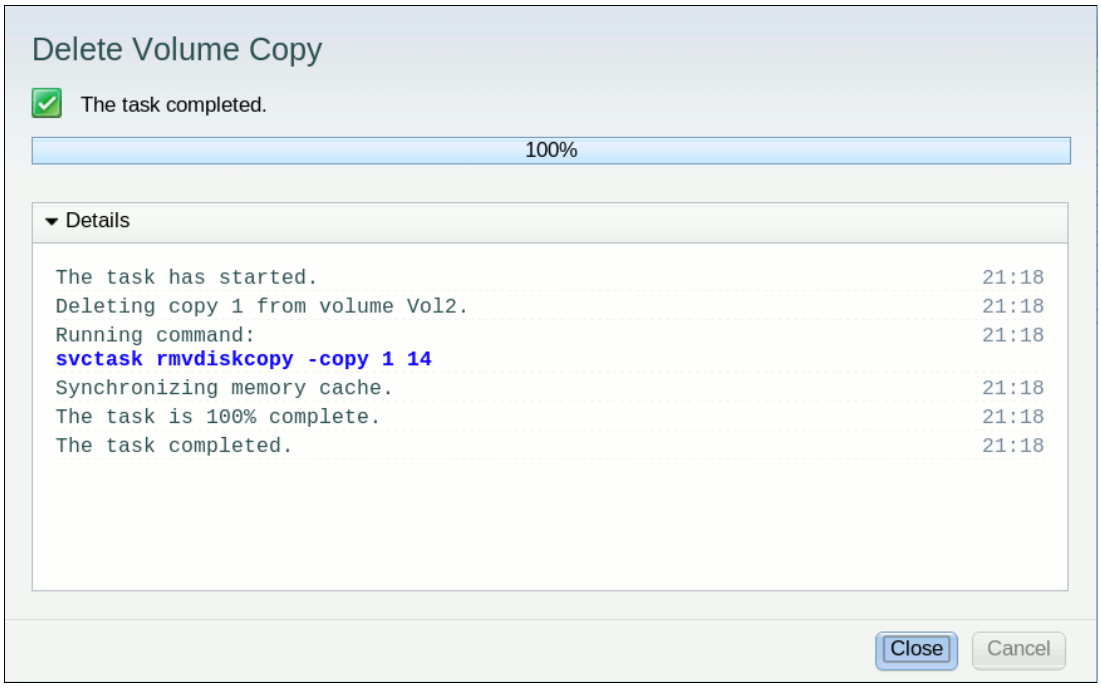

Figure 8-111 Delete volume copy complete

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 360. If you use this feature, one extent after another is migrated to the new storage pool. However, the use of volume copies provides another possible way to migrate volumes. 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 374.
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 380.
4. Split or delete the old copy from the volume. For more information, see 8.6.2, “Splitting into a new volume” on page 383 or 8.6.4, “Delete Copy option” on page 387.

This migration process requires more user interaction, 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 only done from the primary copy.

Now you can switch the role of the copies online (step 3), and test the performance of the new pool. If 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 pool did not meet your requirements.

8.7 Volumes by storage pool

To get an overview of which volumes are on which storage pool, click **Volumes by Pool** (see Figure 8-55 on page 350), and the Volumes by Pool window opens, as shown in Figure 8-112.

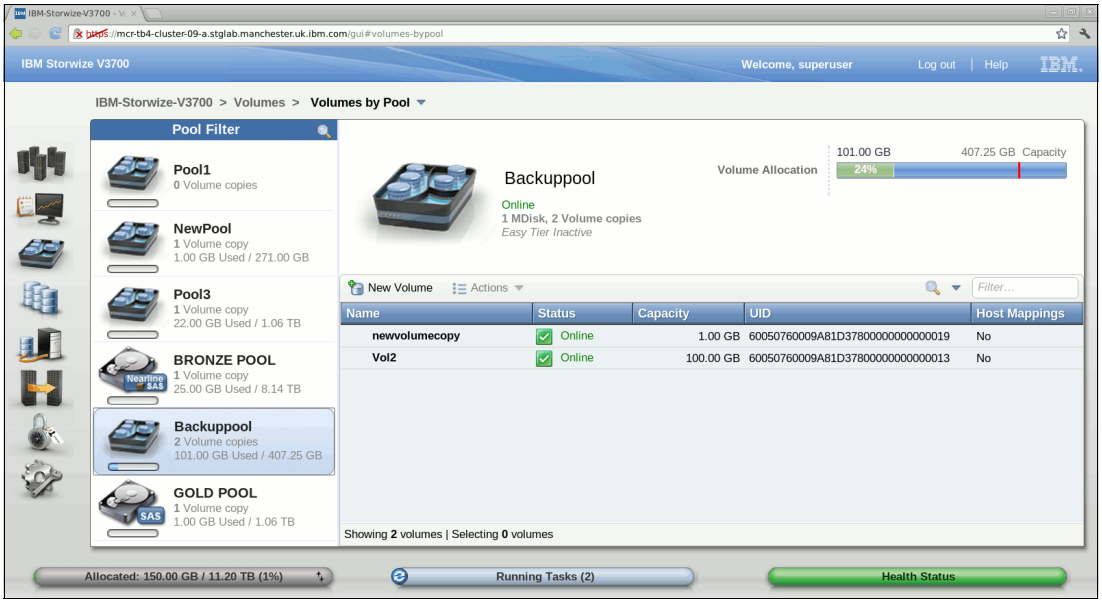

Figure 8-112 Volumes by Pool window

The left pane is named 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 277.

In the upper right, the following information about the selected pool is shown:

- Pool icon: Storage Pools have different characteristics. It is possible to change the pool function icon to identify the pool type. For more information, see 8.7.1, “Changing the Storage Pool function icon” on page 391.
- Pool Name: This name is the name that is given to the storage pool during creation. It is possible to change this name from this window. For more information, see 8.7.2, “Changing the storage pool name” on page 393.
- Pool Details: Provides status information about the pool, such as the number of MDisks and volume copies.
- Volume allocation: Provides details about the available, allocated, and virtual space in this pool.

The lower right lists all volumes that have at least one copy in this storage pool and provides the following information about them:

- Name: Shows the name of the volume.
- Status: Gives status feedback about the volume.
- Capacity: Shows the capacity that is presented to the host of the volume. If it has the green volume sign next to it, this sign means that this volume is thin-provisioned and the capacity that is shown is the virtual size.
- UID: Shows the Volume unique identifier.
- Host Mappings: Shows whether (at least) one host mapping exists.

It is also possible to create a volume in the selected pool. Click **Create Volume** and the wizard that is described in 5.1, “Provisioning storage from IBM Storwize V3500 and making it available to the host” on page 172 starts.

If you highlight a volume and select **Actions** or right-click the volume (as shown in Figure 8-113), the options from the All Volumes window appear. For more information, see 8.4.1, “Advanced volume functions” on page 351.

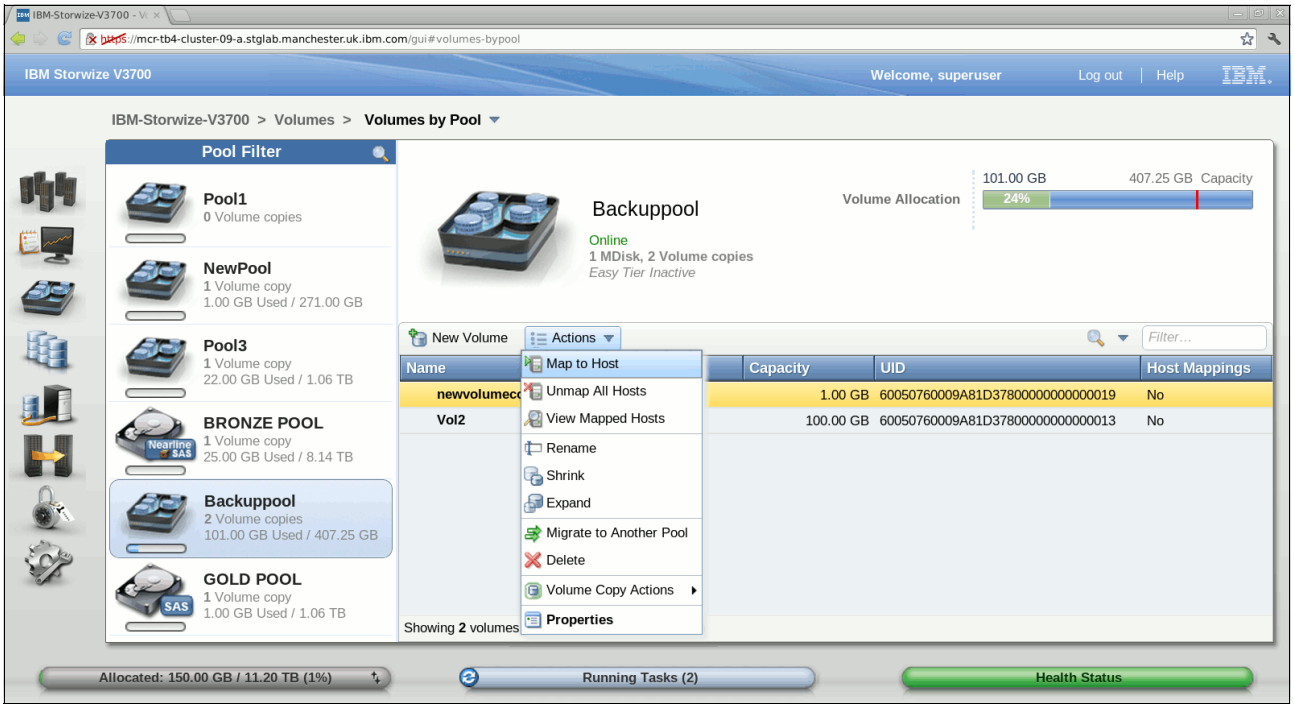

Figure 8-113 Volumes by Pool Actions menu

If you highlight a volume copy and select **Actions** or right-click the volume (as shown in Figure 8-114), the options from the All Volumes window. For more information, see 8.6, “Advanced volume copy functions” on page 380.

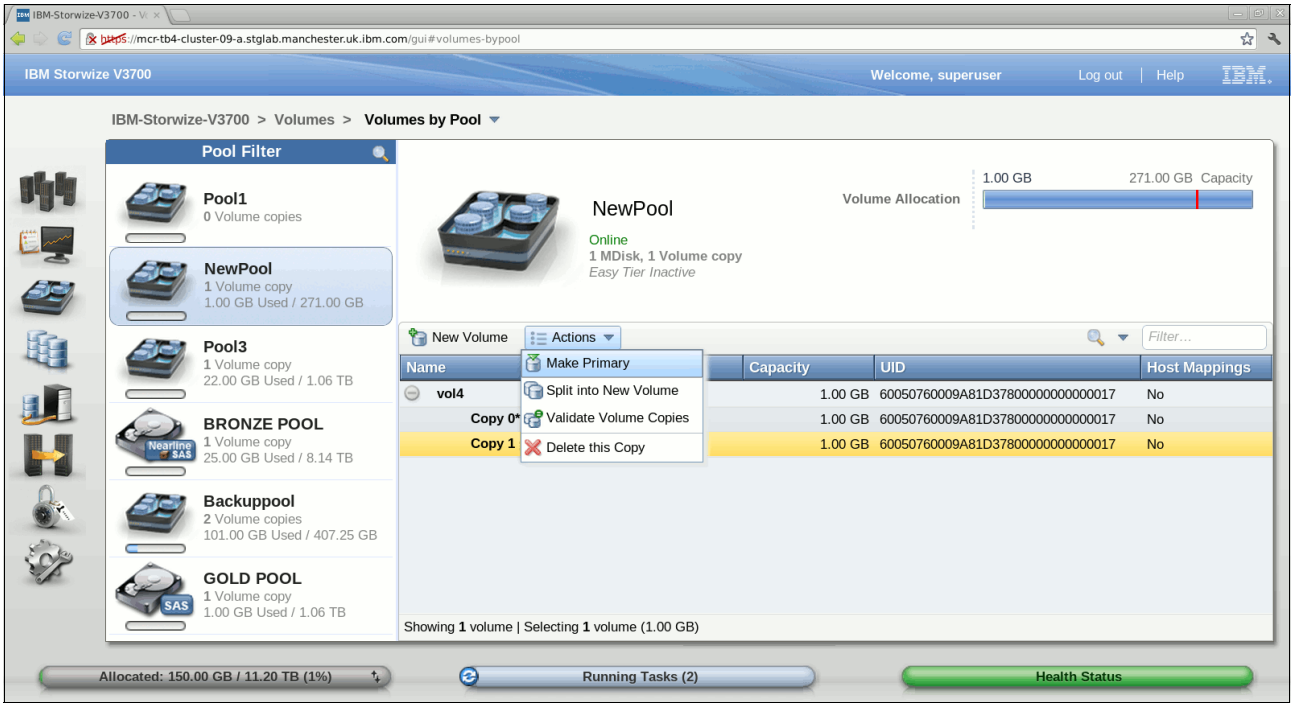

Figure 8-114 Volume Copy Actions

8.7.1 Changing the Storage Pool function icon

Usually, storage pools have different characteristics. You can change the Storage Pool function icon to reflect these differences. Click the pool function icon and the Change Icon window opens, as shown in Figure 8-115.

Figure 8-115 Change Icon window

Use the left and right arrows to select a new function icon, as shown in Figure 8-116. There are several options available.

Figure 8-116 Select new icon

Click **OK** and the changes are applied to the system. Click **Close**.

Figure 8-117 shows the output after the pool icon is changed.

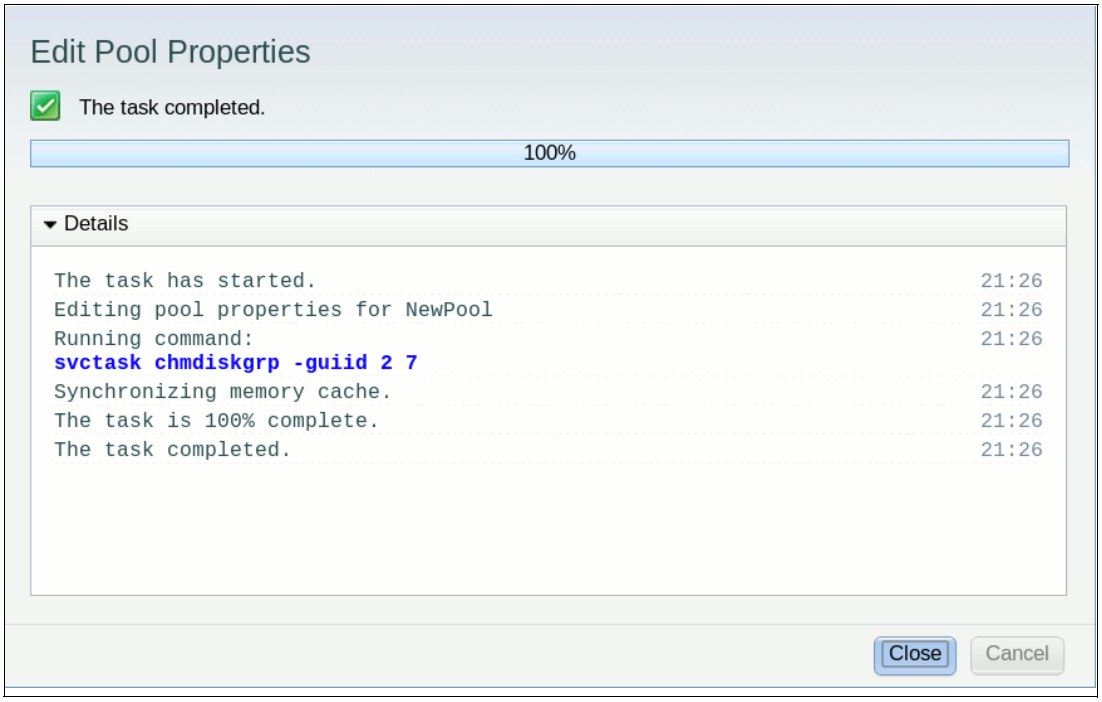

Figure 8-117 Change function icon: Complete

The function icon is changed to make it easier to identify the pool.

Figure 8-118 shows the storage pool with the new icon.

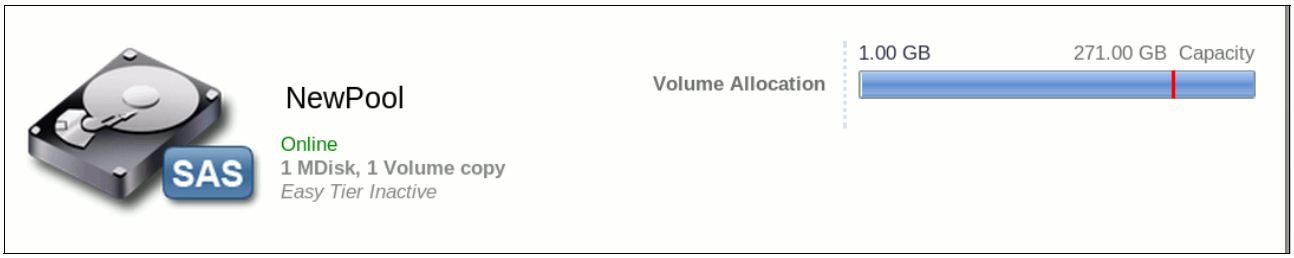

Figure 8-118 Changed function icon

8.7.2 Changing the storage pool name

To change the name of a storage pool, click the name and you can edit it, as shown in Figure 8-119.

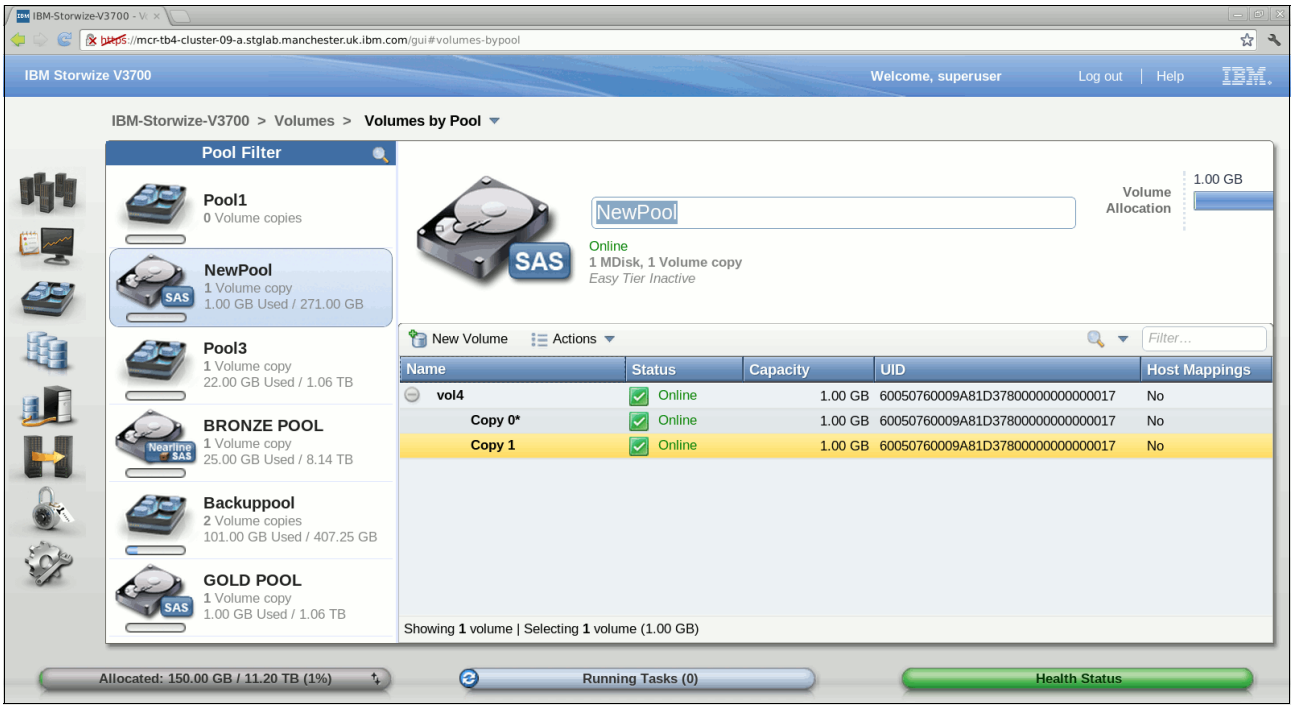

Figure 8-119 Change name

Enter a new name, as shown in Figure 8-120.

Figure 8-120 Enter new name

Press Enter and the changes are applied to the system. Figure 8-121 shows the output after the name of a storage pool is changed.

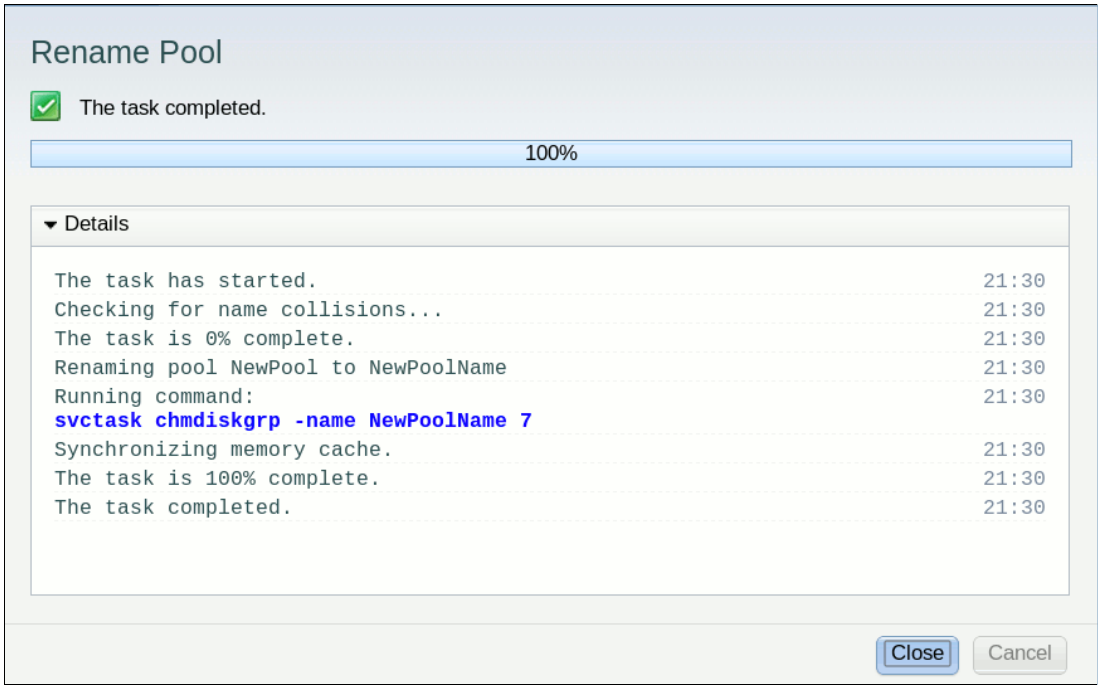

Figure 8-121 Rename storage pool: Complete

The name for the storage pool is changed. Figure 8-122 shows the pool with its new name.

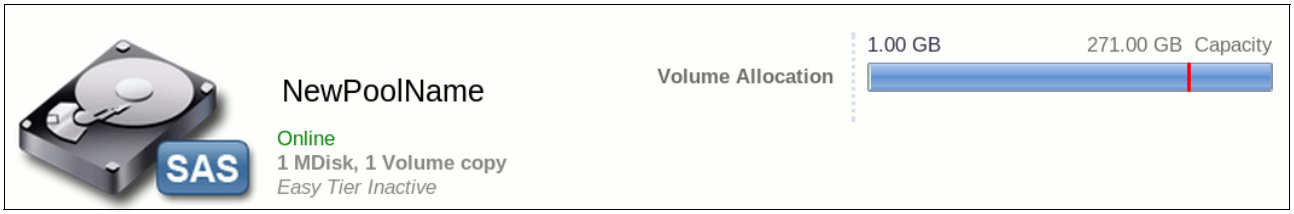

Figure 8-122 Change name: Complete

8.8 Volumes by Host

To see an overview of which volumes a host can access, click **Volumes by Host** (see Figure 8-55 on page 350), and the Volumes by Host window opens, as shown in Figure 8-123.

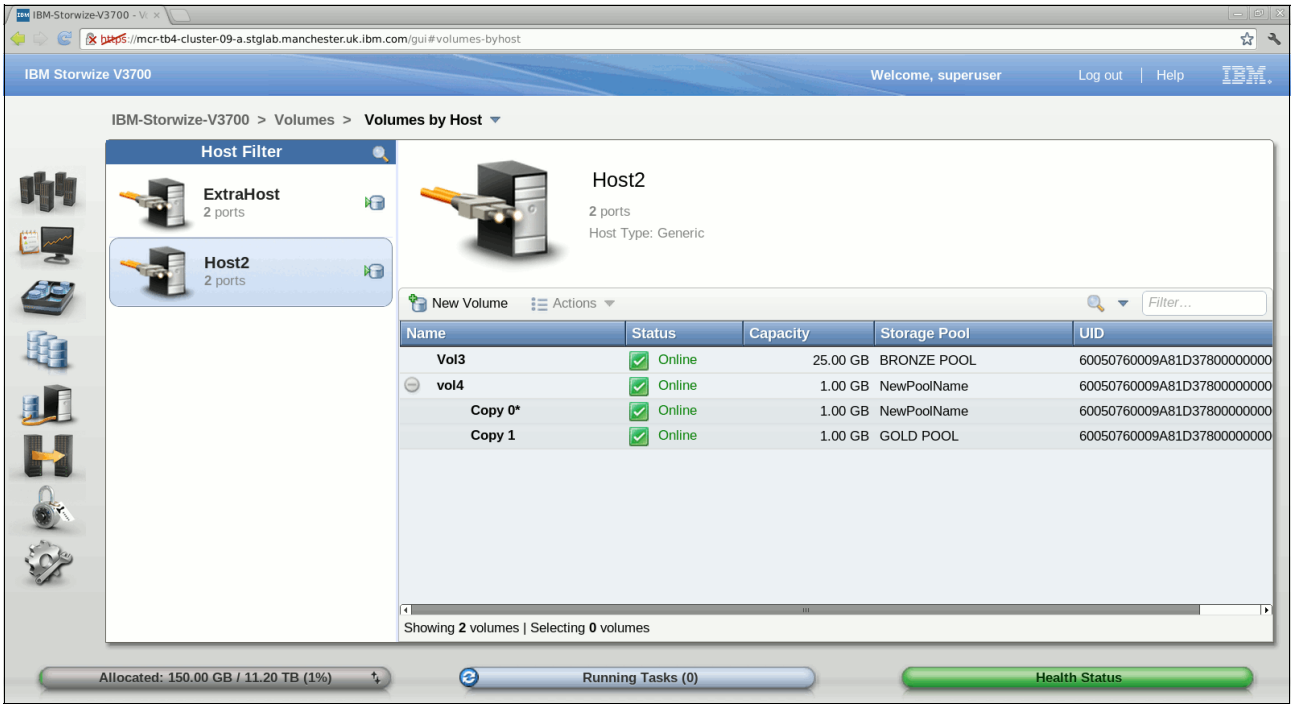

Figure 8-123 Volumes by Host window

In the left pane of the view is the Host Filter. If you select a host, its properties appear in the right pane of the view. The hosts with the orange cable represent Fibre Channel hosts, and the blue cable represents iSCSI hosts. In the upper right side, you see the Host function icon, the host name, the number of host ports, and the host type. The volumes that are mapped to this host are listed in the lower pane.

To create a volume for this host, click **New Volume**, and the wizard that is described in 5.1, “Provisioning storage from IBM Storwize V3500 and making it available to the host” on page 172 starts.

Highlight a volume and select **Actions** or right-click the volume (as shown in Figure 8-124), and the options from the All Volumes window appear. For more information, see 8.4.1, “Advanced volume functions” on page 351).

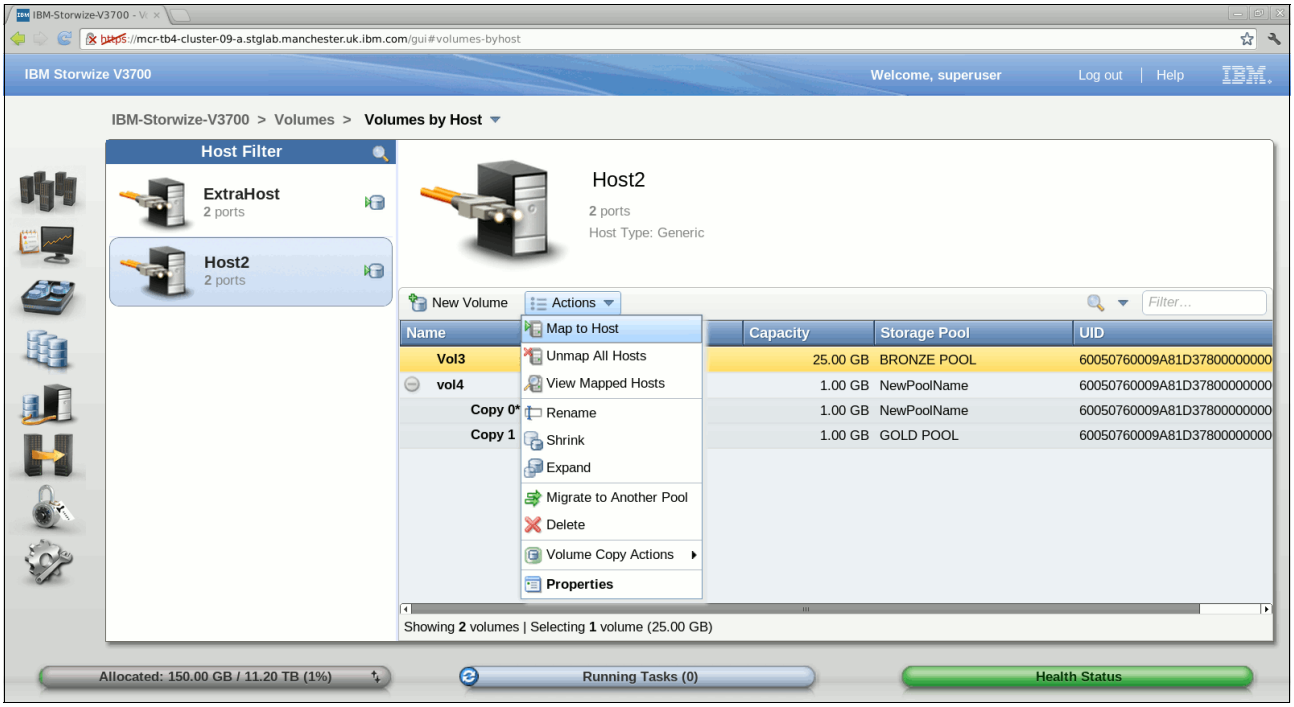

Figure 8-124 Volume Actions

If there are multiple copies of a volume, you can highlight a volume copy and select **Actions** or right-click the copy (as shown in Figure 8-125), and the options from the All Volumes window appear. For more information, see 8.6, “Advanced volume copy functions” on page 380).

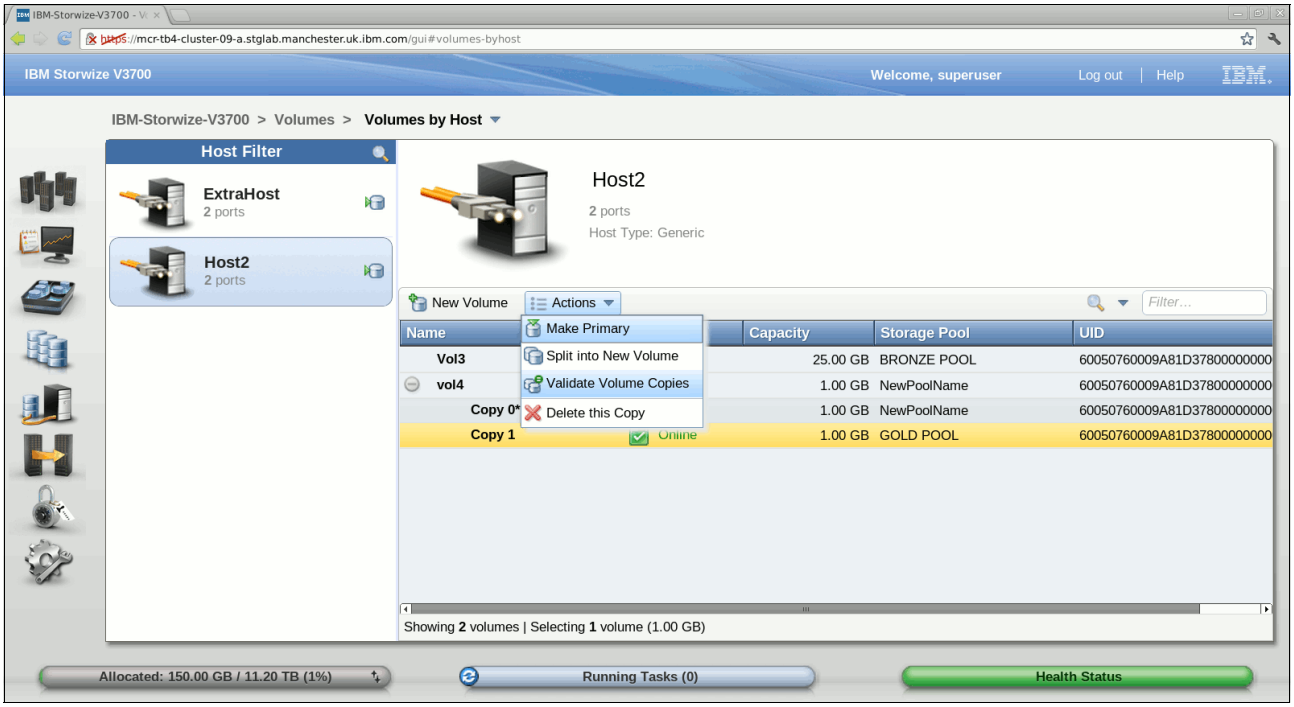

Figure 8-125 Volume Copy Actions

8.8.1 Renaming a host

To rename a host in the Volumes by Host window, click it and you can edit the name, as shown in Figure 8-126.

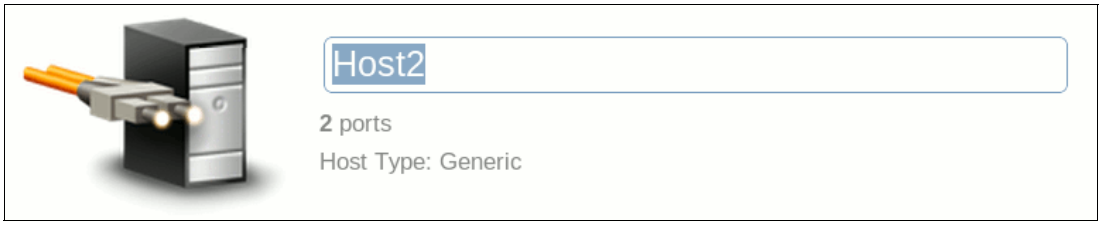

Figure 8-126 Rename a host

Enter a new name, as shown in Figure 8-127.

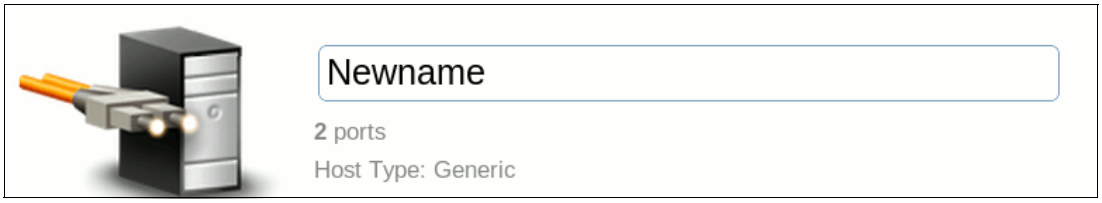

Figure 8-127 Enter a new name

Press Enter to apply the changes to the system. Figure 8-128 shows the output after a host is renamed.

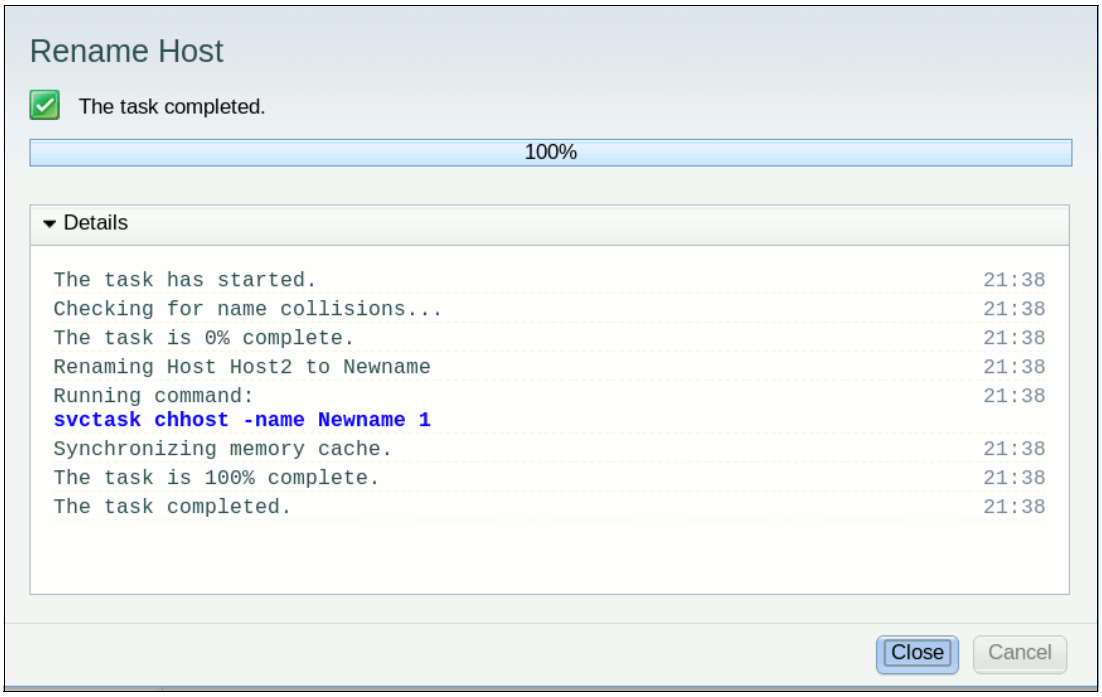

Figure 8-128 Rename host complete

Click **Close** to return to the Volumes by Host window. Figure 8-129 shows the Volumes by Host window, with the renamed host highlighted.

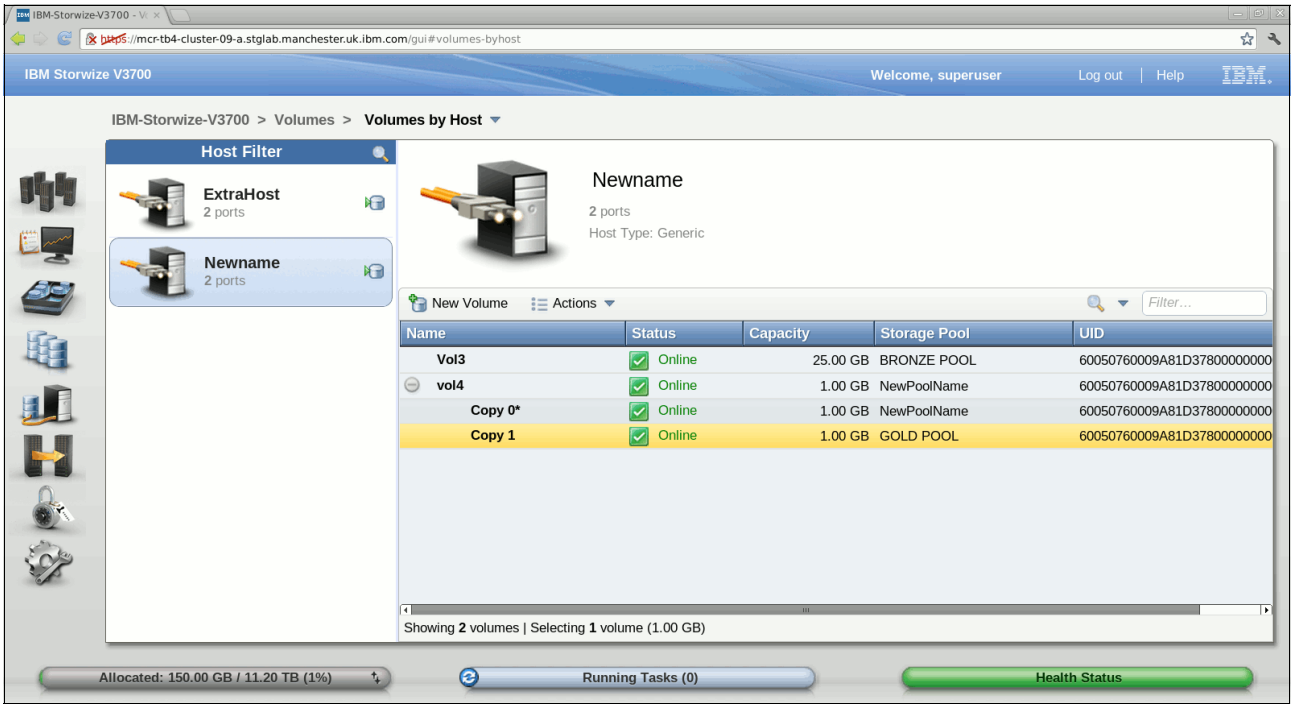

Figure 8-129 Volumes by Host

8.9 FlashCopy

By using the FlashCopy function of the IBM Storwize V3500 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 more information about its configuration and use.

You can use FlashCopy to solve critical and challenging business needs that require 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 and below the host operating system and cache, those levels must be flushed for consistent FlashCopy copies.

While the FlashCopy operation is performed, the source volume is frozen briefly to initialize the FlashCopy bitmap and then I/O 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 initiated, and an indirection layer, which allows data to be read from the source volume transparently.

8.9.1 Business requirements for FlashCopy

With an immediately available copy of the data, FlashCopy could be used in the following business scenarios:

- Rapidly creating consistent backups of dynamically changing data.

  Backups can be created through periodic execution of FlashCopy. 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 be also used for backup-to-tape by attaching them to another server, which offloads the production server. After the copy to tape completes, the target volumes can be discarded (if required) 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 minimizing 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 swifter than other migration methods that are available through the IBM Storwize V3500 because the size or 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 the amount of data and the technology in use, this process could take a day or more. With FlashCopy, a copy

can be created and online for use in a few minutes. The time varies based on the application and the data set size.

- Rapidly creating copies of production data sets for the purposes of auditing and data mining.

  Auditing or data mining normally require the use of 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 situation reduces the load on your production systems and 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 typically is extended. This reduction in the time that is required allows much more frequent refreshes of the quality assurance database.

8.9.2 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. 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 404.

The minimum granularity that IBM Storwize V3500 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 V3500 storage system, but they do not have to be in the same storage pool.

Before a FlashCopy (regardless of the type and options specified) is started, run the **prestartfcmap** or **prestartfcconsistgrp** command from the command line. These commands put the SAN Volume Controller Cache into write-through mode, which flushes the I/O that is bound for your volume. 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.

When data is copied between volumes, it is copied in units of address space that is known as *grains*. Grains are units of data that are grouped to optimize the use of the bitmap and 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 behavior for the source and target volumes:

- Read I/O request to source: Reads are performed from the source volume, which is the same as for non FlashCopy volumes.
- Write I/O request to source: Writes to the source cause the grains to be copied to thetarget if they are not already copied. Then, the write is performed to the source.
- Read I/O request to target: Reads are performed from the target if the grains already are 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 unless the entire grain is being written. Then, the write completes to the target.

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 404.

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 data that was modified only since the initial FlashCopy mapping. This action reduces the amount of time that it takes to re-create an independent FlashCopy image. A FlashCopy mapping can be defined only as incremental at the time of its creation.

Multiple target FlashCopy mappings

It is possible to copy up to 64 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 64 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 64 mappings can exist in a cascade. If cascaded mappings and multiple target mappings are used, a tree of up to 64 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, which provides pointer-based images for limited lifetime uses, is also called the *no-copy option*.

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, 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, the target volume remains a valid copy only 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 `-resize` is correctly set up.)

The relationship of the background copy rate value to the amount of data that is copied per second is shown in Table 8-1.

Table 8-1 Background copy rate

Value	Data 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 with regards to 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 with the backup copy rate (see Table 8-1).

The cleaning rates control the rate at which the cleaning process operates. The cleaning process' purpose 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 was added to prevent data access interruption for dependent mappings, when their source is issued a stop, as was the distinction between stopping and stopped states.

FlashCopy mapping states

At any point, a mapping is in one of the following states:

- 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. If the connection to both nodes in the IBM Storwize V3500 storage system to which the mapping is assigned is lost, the source and target volumes go offline.

- 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 is completed, the mapping records only the differences between the source and target volumes. If the connection to both nodes in the IBM Storwize V3500 storage system to which the mapping is assigned is lost, the source and target volumes go offline.

- 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 is completed, the mapping records only the differences between the source and target volumes. If the connection to both nodes in the IBM Storwize V3500 storage system to which the mapping is assigned is lost, the source and target volumes go offline.

- Stopped

  The mapping is stopped because a stop command was issued 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. If the connection to both nodes in the IBM Storwize V3500 storage system that the mapping is assigned to is lost, the source and target volumes go offline.

- 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 is incomplete, 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 (which are described in “Dependent writes” on page 405), are run in the application’s intended sequence.

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 8-130 shows a consistency group that consists of two FlashCopy mappings.

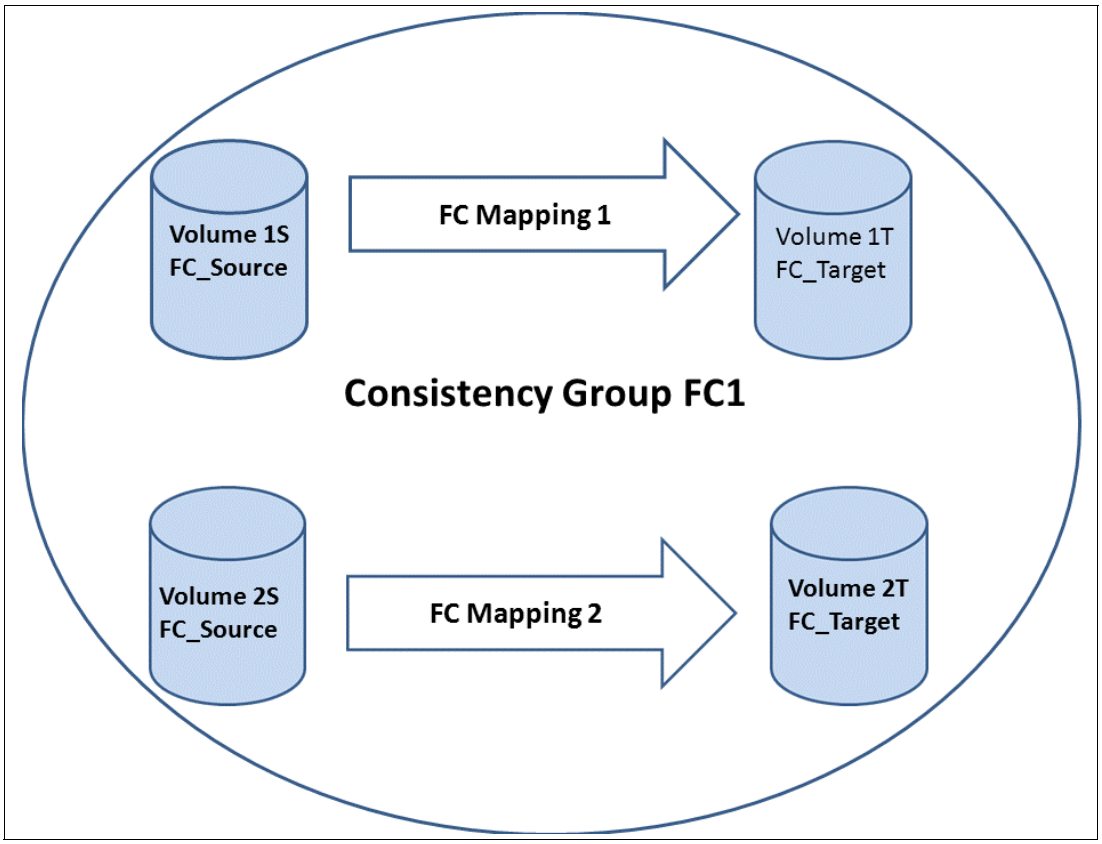

Figure 8-130 Consistency group of two mappings

FlashCopy mapping management: After an individual FlashCopy mapping is 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 the next step is started. 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 that 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 did 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 that are issued to a consistency group are issued simultaneously to all members of the 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. This action 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. It also is possible to add the FlashCopy mapping to a consistency group or change the consistency group of a FlashCopy mapping later. Do not place stand-alone mappings into a consistency group because they become controlled as part of that consistency group.

FlashCopy consistency group states

At any point, a FlashCopy consistency group is in one of the following states:

- 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 inaccessible.

- 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 a command was issued 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.

It also is possible to create an optional copy of the source volume before the reverse copy operation is started. This copy restores the original source data, which can be useful for diagnostic purposes.

Figure 8-131 shows an example of the reverse FlashCopy scenario.

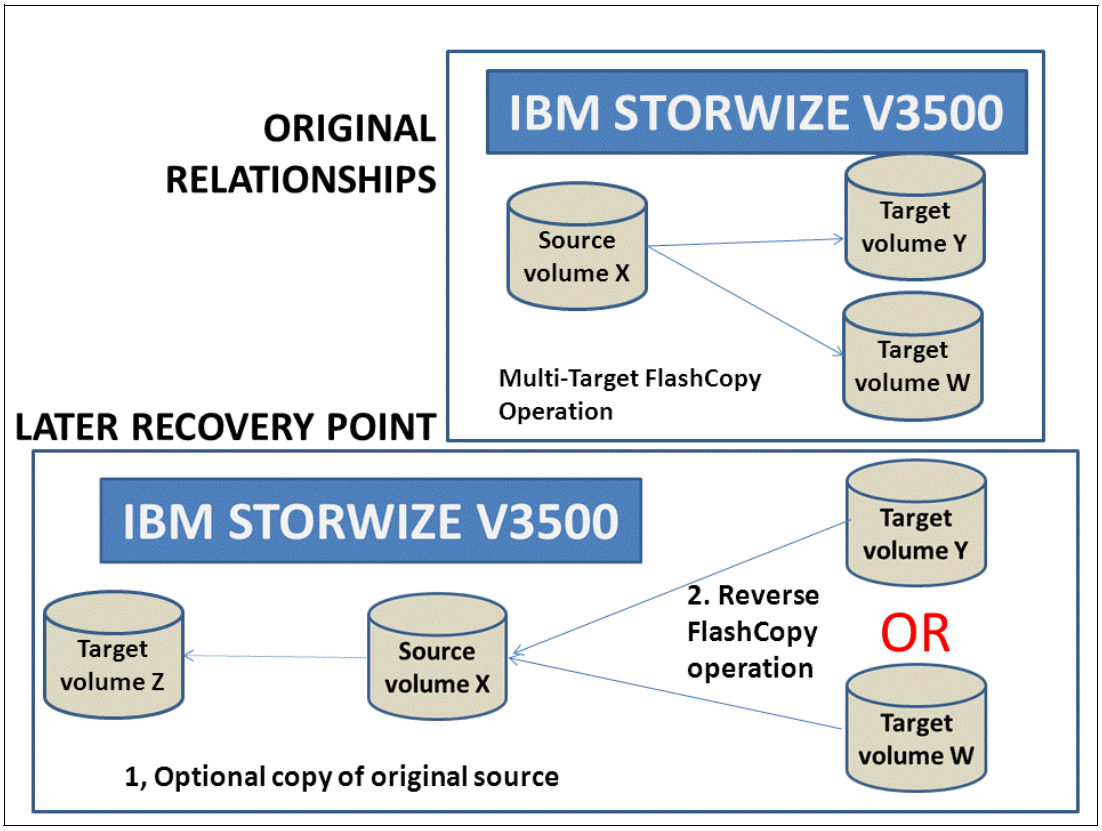

Figure 8-131 Reverse FlashCopy

To restore from a FlashCopy backup, 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, if this map does not exist.
3. Start the FlashCopy map (volume Y -> volume X) with the **`-restore`** option to copy the backup data onto the production disk.

The -restore option: In the GUI, the `-restore` option is applied automatically when you start the FlashCopy mapping from volume Y to volume X. In the CLI, you must add the `-restore` option to the command manually. For more information about the use of the CLI, see Appendix A, “Command-line interface setup and SAN Boot” on page 483.

The production disk is instantly available with the backup data.

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 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.

8.9.3 Planning for FlashCopy

There are several items that must be considered before a FlashCopy is performed, which we describe 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 V3500 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 8-2 shows some of the FlashCopy maximum configurations.

Table 8-2 FlashCopy maximum configurations

FlashCopy property	Maximum
FlashCopy targets per source	64
FlashCopy mappings per cluster	64

FlashCopy presets

The IBM Storwize V3500 storage system provides three FlashCopy presets, named Snapshot, Clone, and Backup, to simplify the more common FlashCopy operations. The presets are described in Table 8-3.

Table 8-3 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. 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 0% of the capacity that is allocated at the time of creation. The preset uses a FlashCopy mapping with 0% 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 itself 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.

8.9.4 Managing FlashCopy by using the GUI

The IBM Storwize V3500 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

Figure 8-132 shows the Copy Services menu.

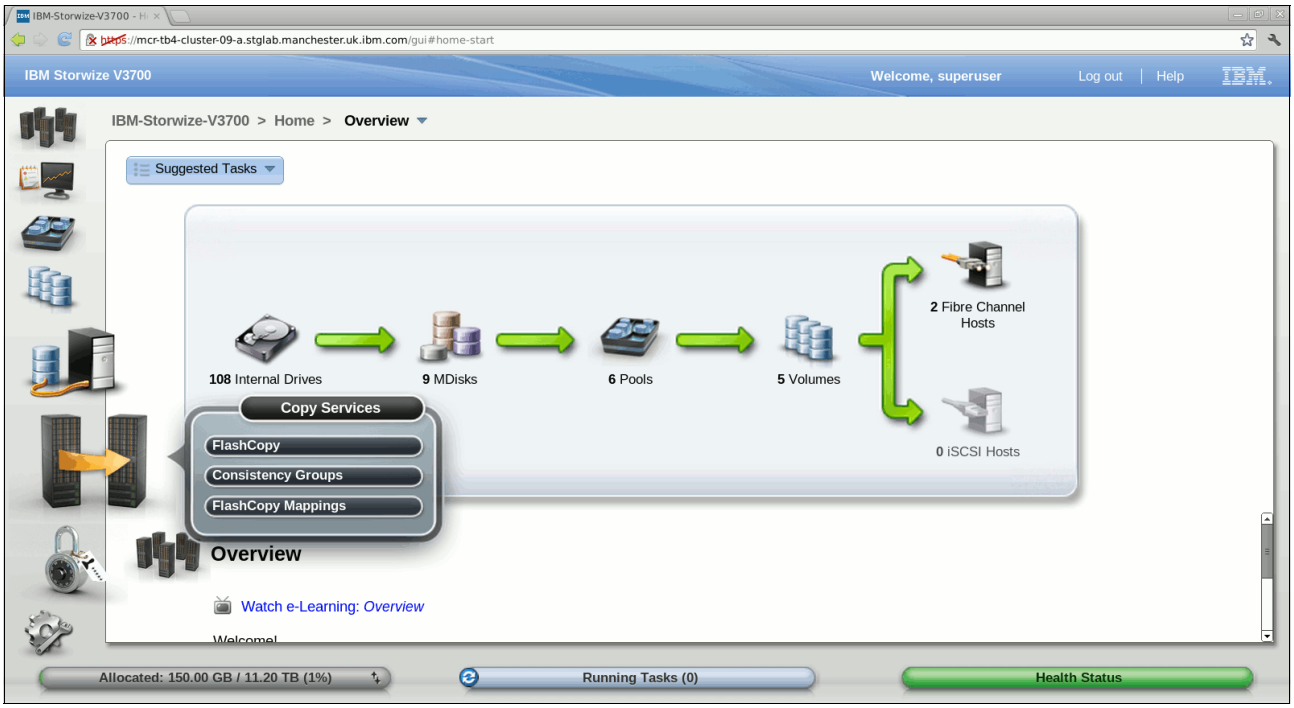

Figure 8-132 Copy Services

Most of the actions to manage the FlashCopy mapping could be done both 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 8-133. In the FlashCopy window, the FlashCopy mappings are organized by volumes.

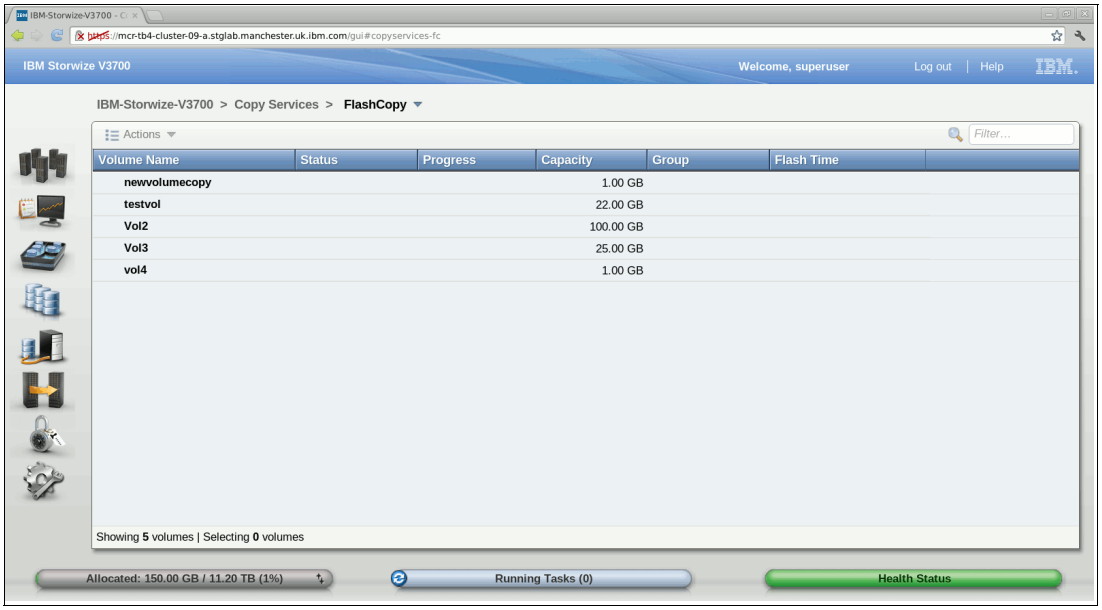

Figure 8-133 FlashCopy window

Click **FlashCopy Mappings** in the Copy Services function icon menu and the FlashCopy Mappings window opens, as shown in Figure 8-134. In the FlashCopy Mappings window, the FlashCopy mappings are listed individually.

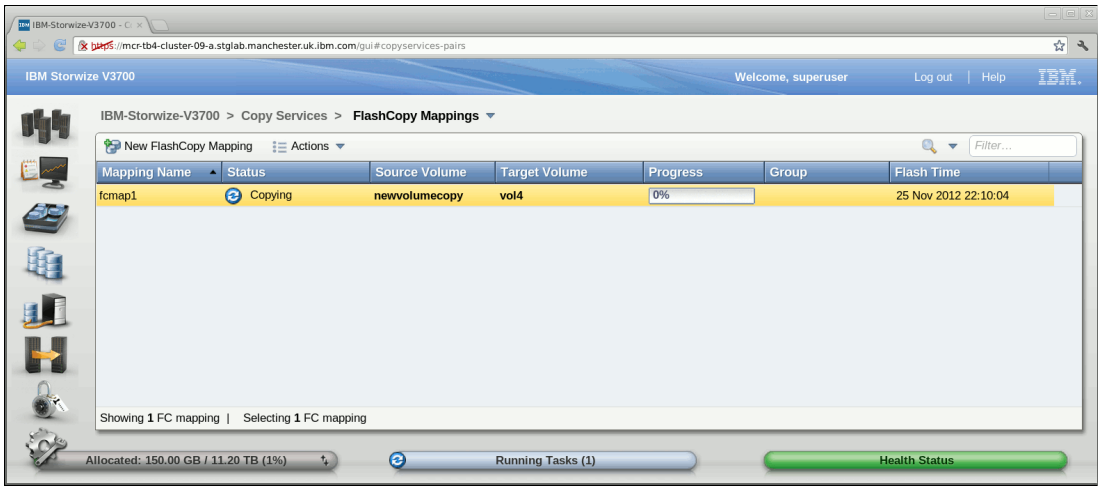

Figure 8-134 FlashCopy mappings

The Consistency Groups window can be 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 8-135.

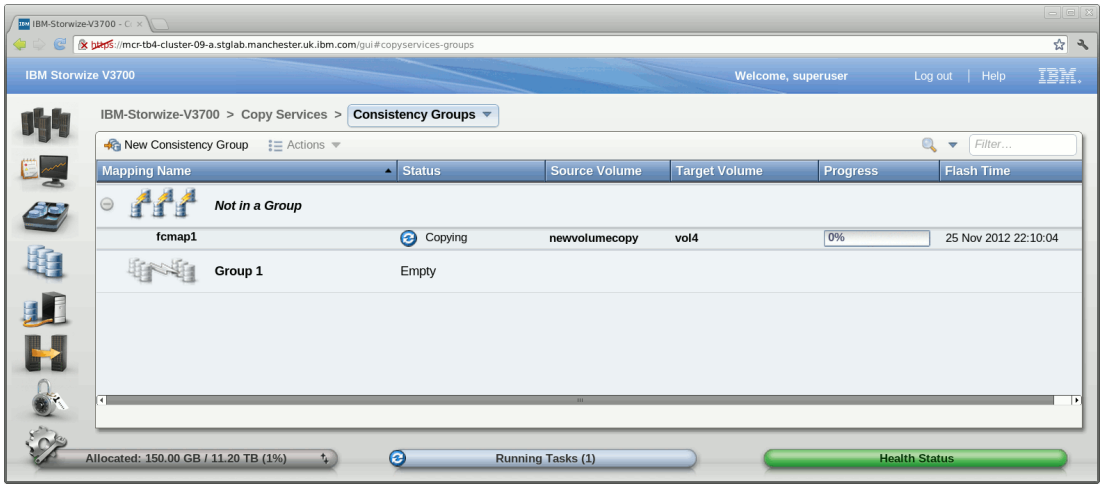

Figure 8-135 Consistency Groups

8.9.5 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 8-136.

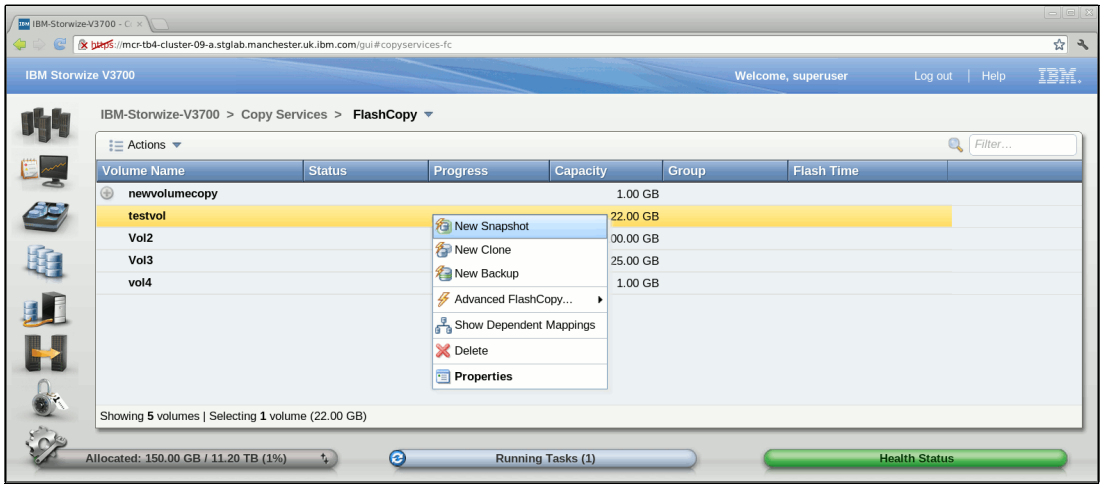

Figure 8-136 New Snapshot

A snapshot of the selected volume is created.

Creating a clone

In the FlashCopy window, choose a volume and click **New Clone** from the Actions drop-down menu, as shown in Figure 8-137.

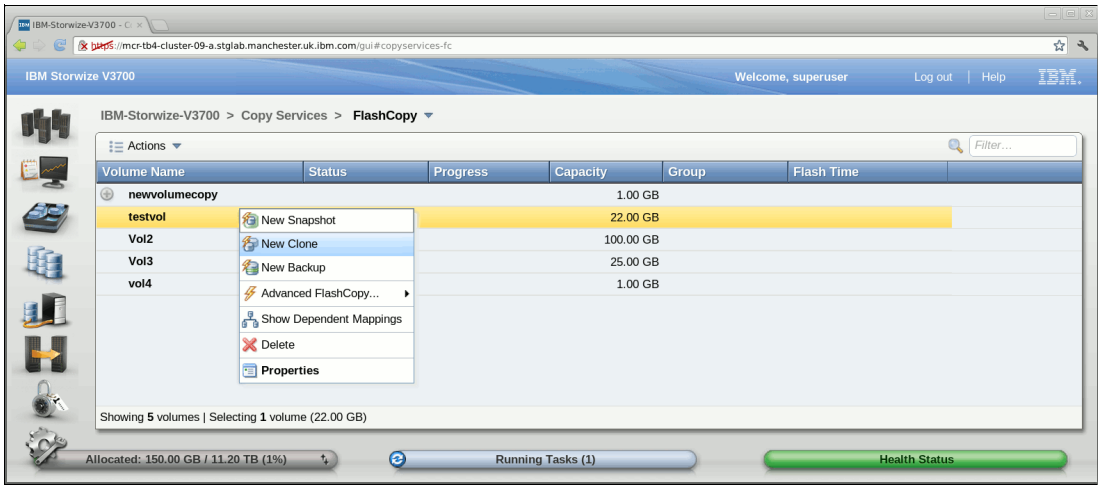

Figure 8-137 New Clone

A clone of the selected volume is created.

Creating a backup

In the FlashCopy window, choose a volume and click **New Backup** from the Actions drop-down menu, as shown in Figure 8-138.

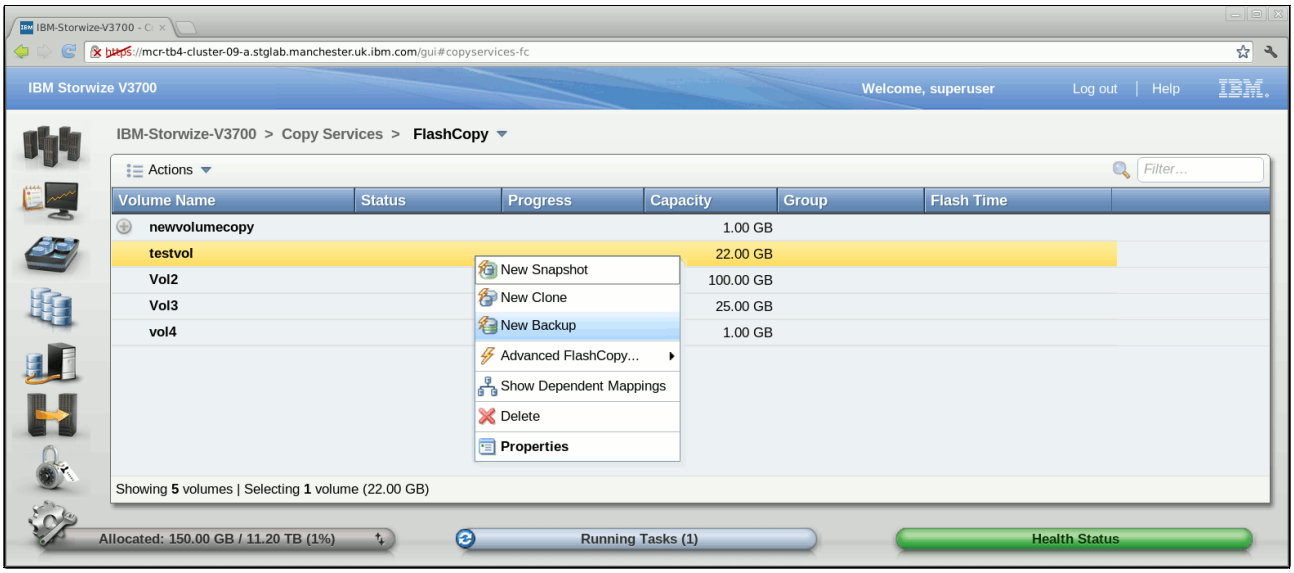

Figure 8-138 New Backup

In the FlashCopy window, three FlashCopy target volumes are visible under the source volume, as shown in Figure 8-139. The progress bars behind the target volumes indicate the copy progress as a percentage. The copy progress remains 0% for the snapshot; there is no change happening at the moment because only the changed data is being copied. At the same time, the copy progresses for clone and backup keep increasing.

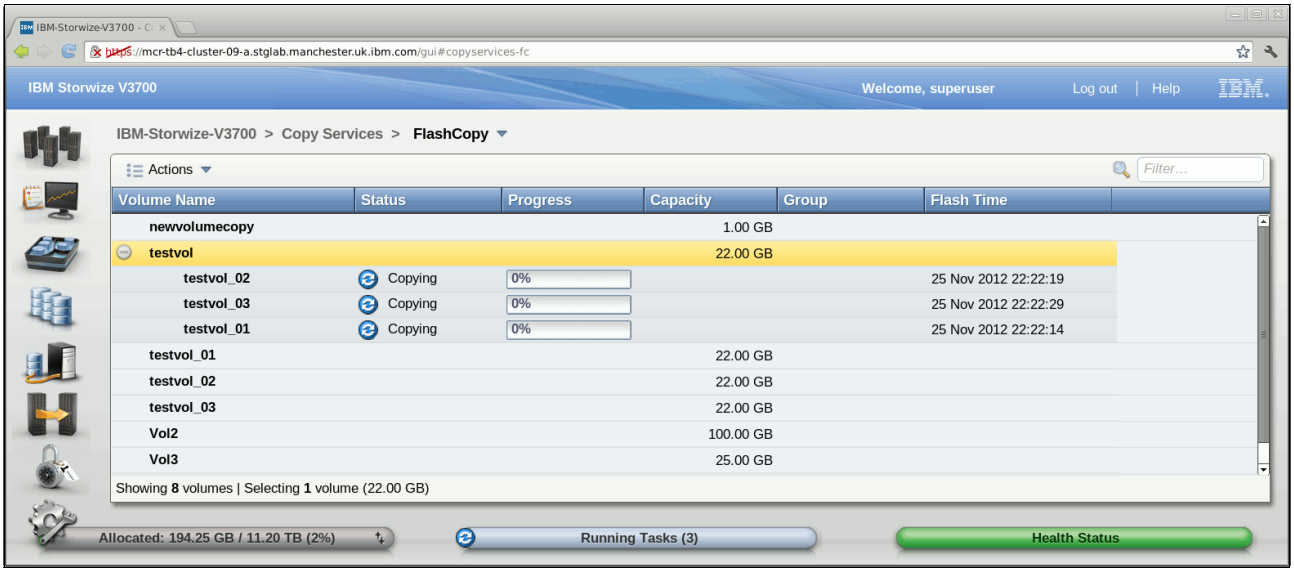

Figure 8-139 FlashCopy progress

The copy progress also can be found under the Running Tasks status indicator, as shown in Figure 8-140. This view is slightly different than the Remote Copy Running Tasks status indicator.

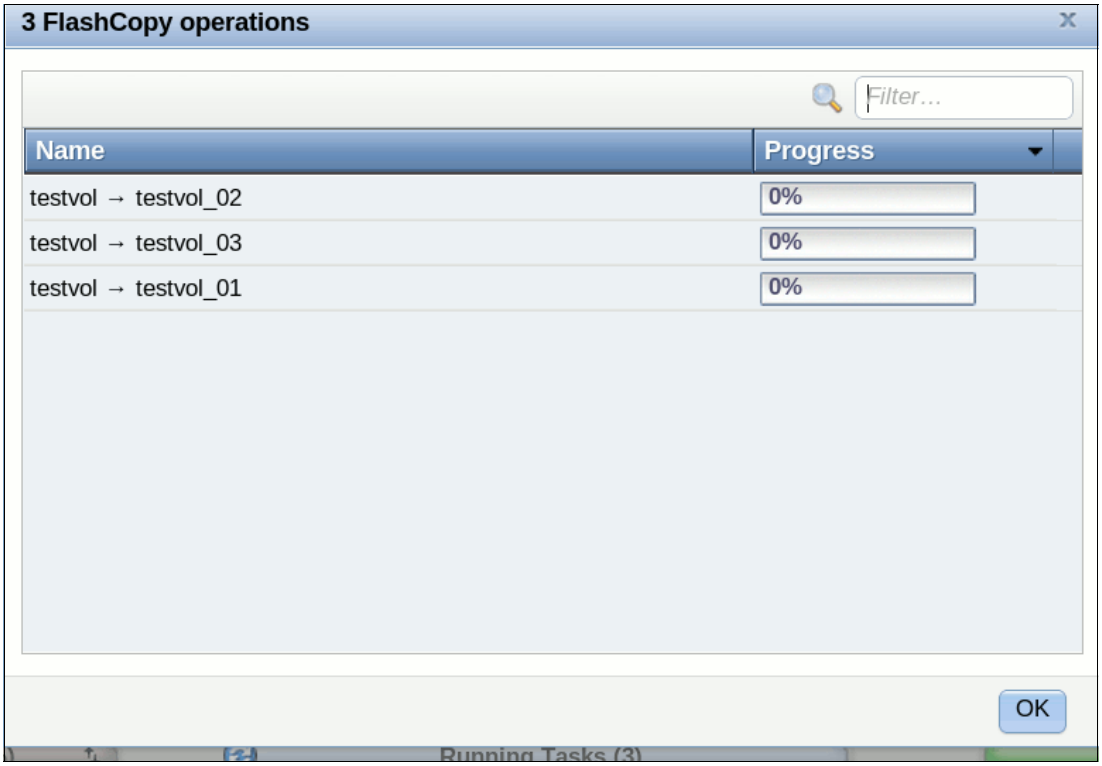

Figure 8-140 FlashCopy progress

After the copying completes, the FlashCopy with the clone preset is automatically deleted, as shown in Figure 8-141.

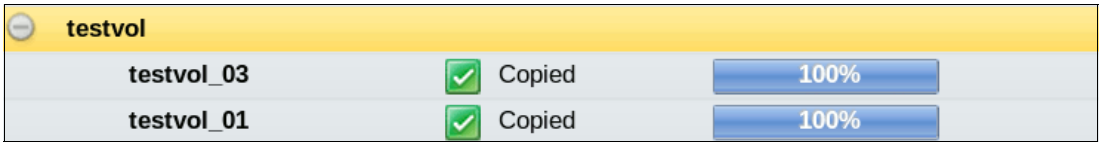

Figure 8-141 FlashCopy completed

8.9.6 Managing FlashCopy mapping

The FlashCopy presets cover the most used FlashCopy configurations for general situations. However, customized FlashCopy mappings are still necessary in some complicated scenarios.

Creating FlashCopy mappings

It is possible to create FlashCopy mappings through the FlashCopy window. Select the source volume for the FlashCopy mapping and click **Advanced FlashCopy...** from the Actions drop-down menu, as shown in Figure 8-142.

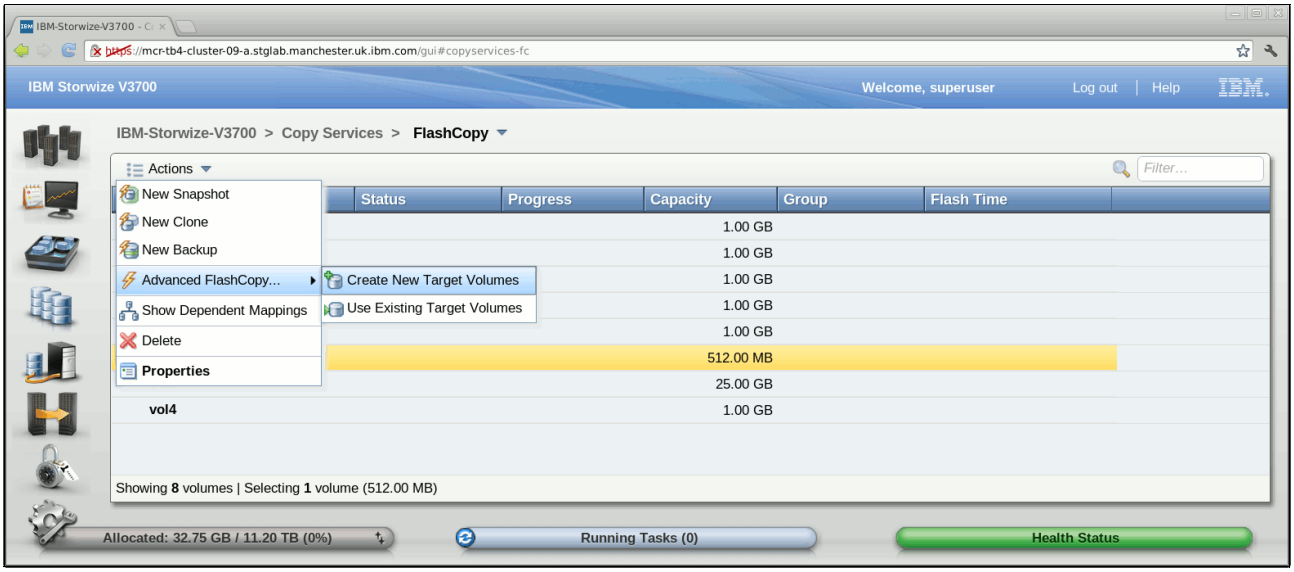

Figure 8-142 Advanced FlashCopy

The following options are available:

- If you did not create the target volume, click **Create New Target Volumes**.

  By using the wizard, you can choose the preset. However, you can modify the setting of the FlashCopy mapping. Choose the preset that has the most similar configuration to the one that is required and click **Advanced Settings** to make any appropriate adjustments to the properties, as shown in Figure 8-143.

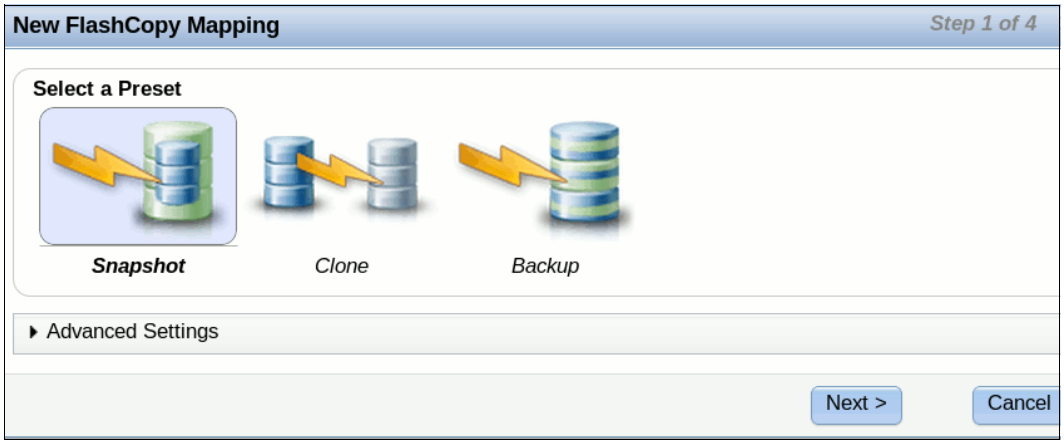

Figure 8-143 Select Preset

As shown in Figure 8-144, if the Snapshot preset is chosen, the following default settings are shown in the Advanced Settings tab:

- Background Copy: 0
- Incremental: No
- Auto Delete after completion: No
- Cleaning Rate: 0

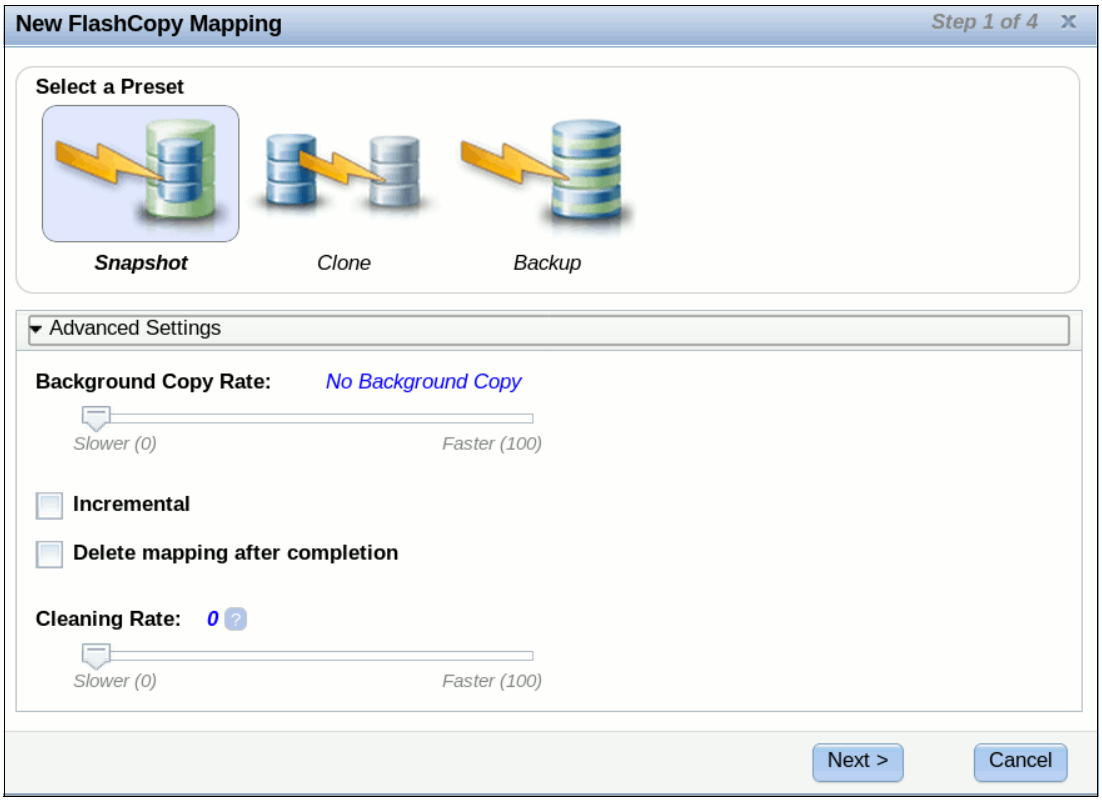

Figure 8-144 Default setting for the Snapshot preset

Similarly, if the Clone preset is chosen, the following default settings are shown in the Advanced Settings tab:

- Background Copy: 50
- Incremental: No
- Auto Delete after completion: Yes
- Cleaning Rate: 50

Figure 8-145 shows the default settings for the Clone preset.

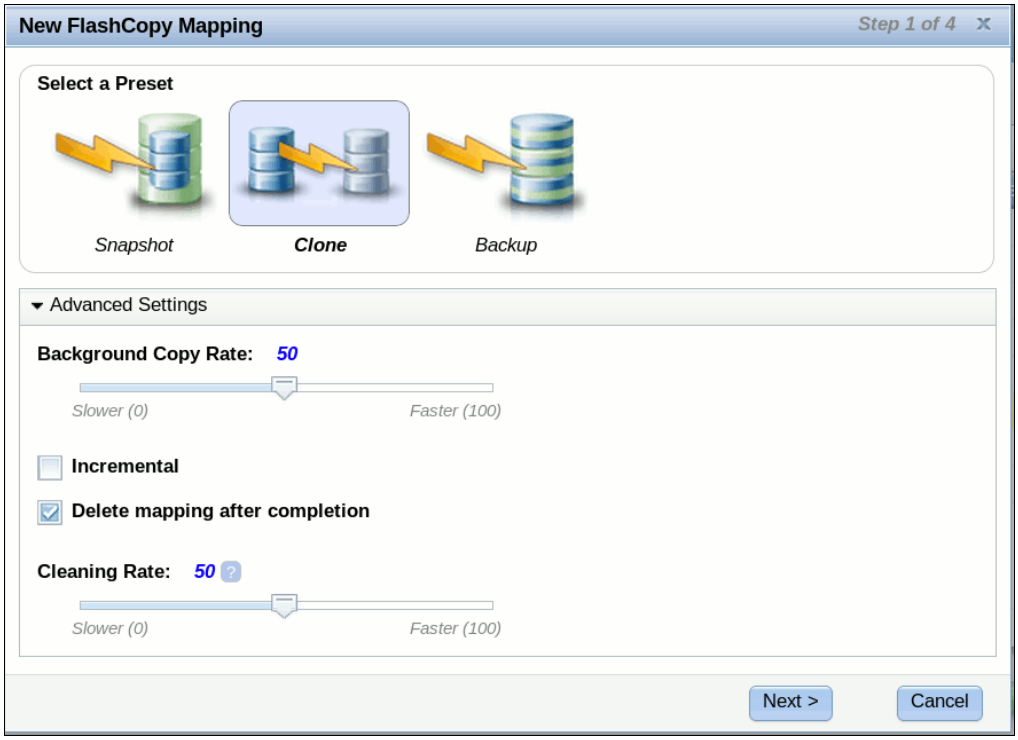

Figure 8-145 Clone preset

Similarly, the following default settings of the Backup preset can be found by clicking **Backup** → **Advanced Settings**, as shown in Figure 8-146:

- Background Copy: 50
- Incremental: Yes
- Auto Delete after completion: No
- Cleaning Rate: 50

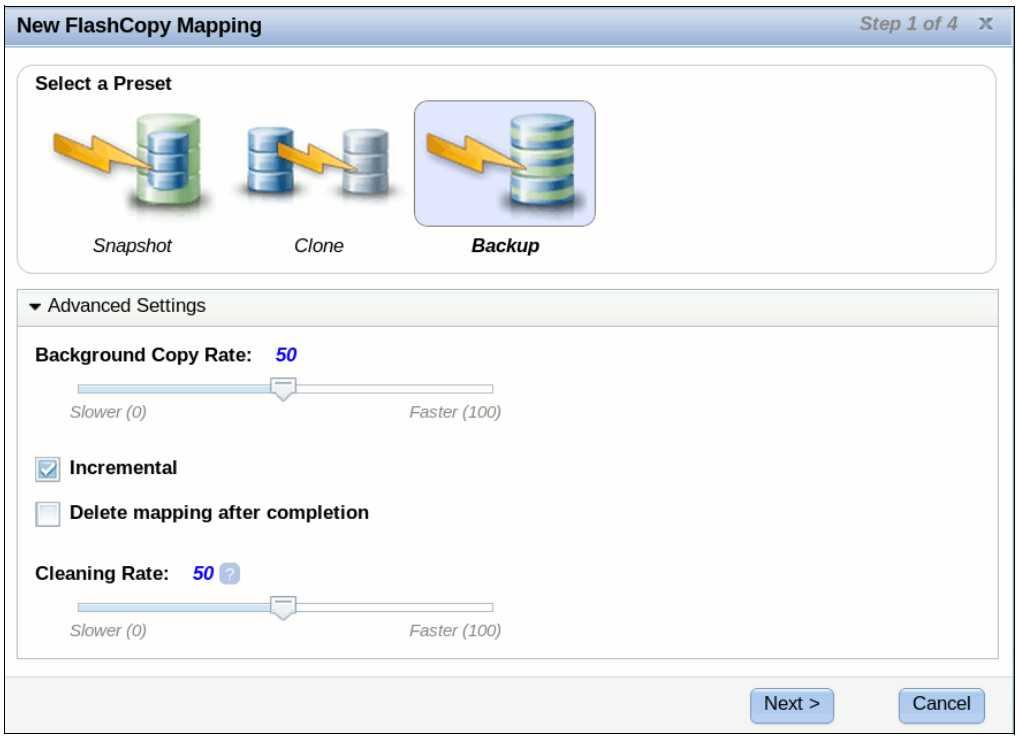

Figure 8-146 Backup default settings

Change the settings of the FlashCopy mapping according to your requirements and click **Next.**

In the next step, it is possible to add the FlashCopy mapping to a consistency group, as shown in Figure 8-147. If the consistency group is not ready, the FlashCopy mapping can be added to the consistency group later. Click **Next** to continue.

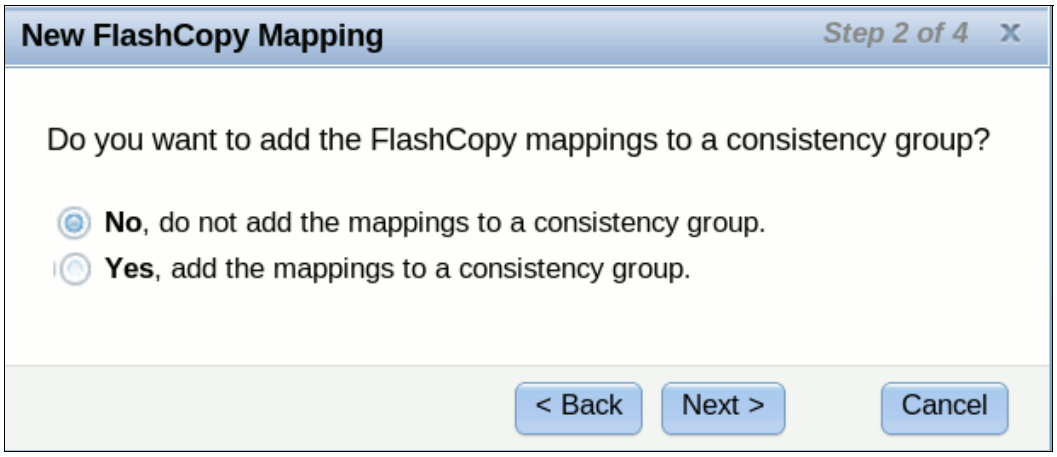

Figure 8-147 Add FlashCopy mapping to consistency group

You can choose the storage pool in which you want to create your volume. As shown in Figure 8-148, you can select the same storage pool that is used by the source volume. Click **Next** to continue.

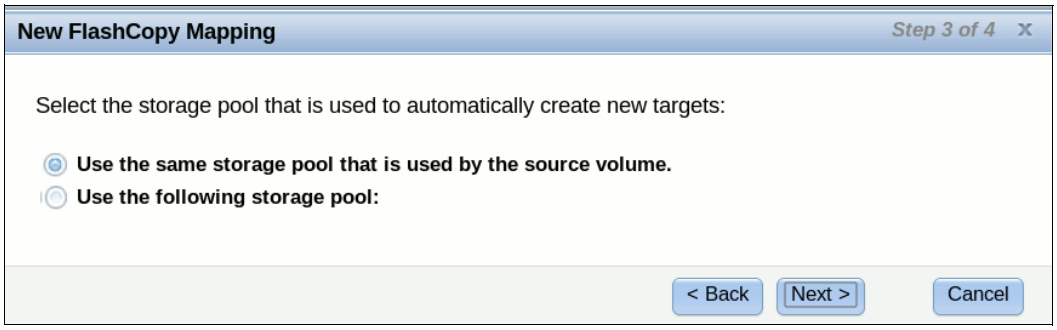

Figure 8-148 Storage pool selection

It also is possible to specify another storage pool for the new volume, as shown in Figure 8-149. Click **Next** to continue.

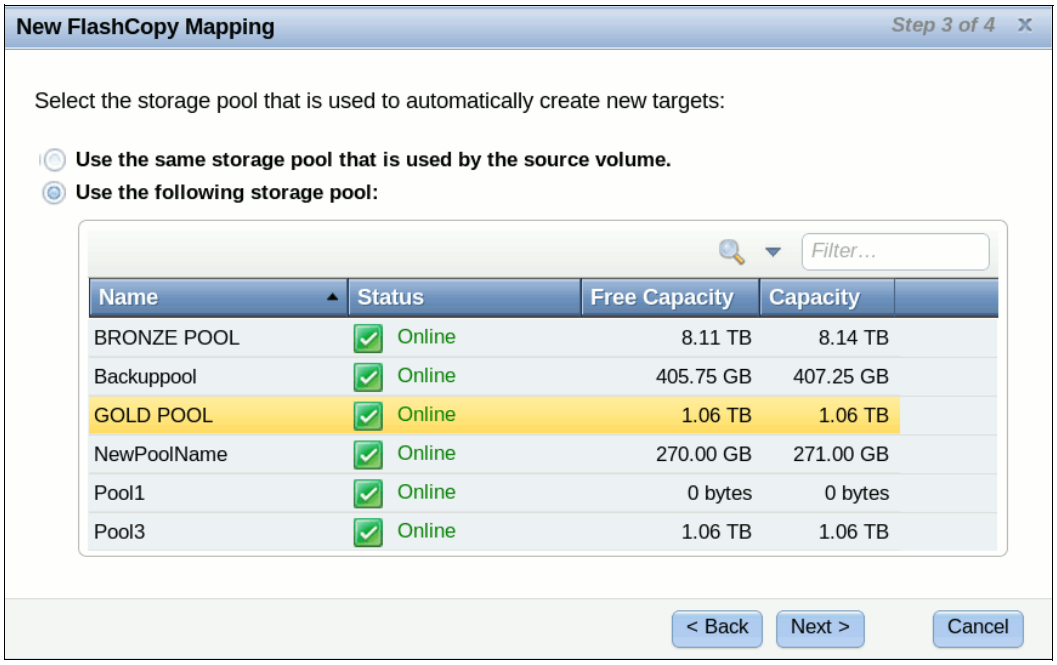

Figure 8-149 Choose another storage pool to create the volume

You are prompted to select a way to manage volume capacity. The default choice is to create a generic volume, as shown in Figure 8-150.

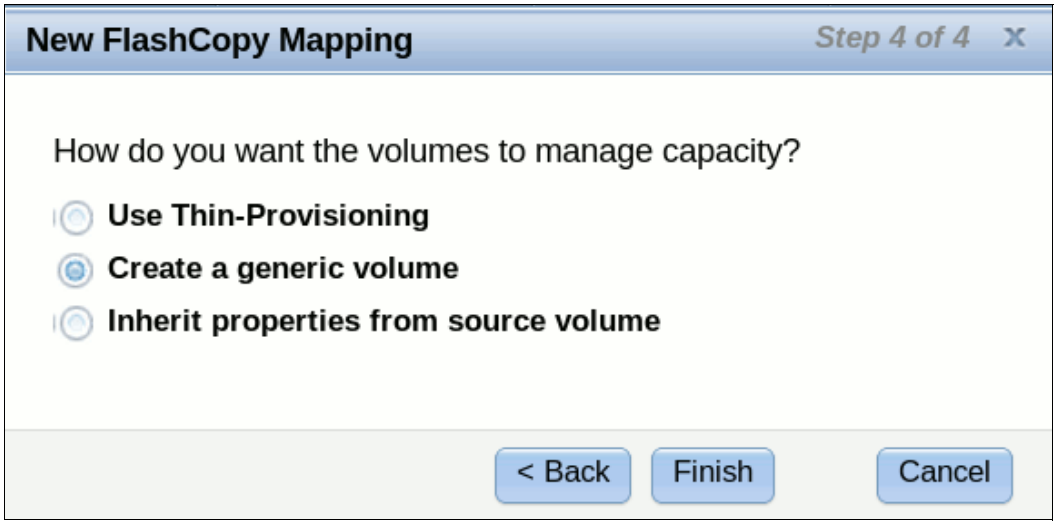

Figure 8-150 Create a generic volume

If you select a thin-provisioned volume, you see more options, as shown in Figure 8-151.

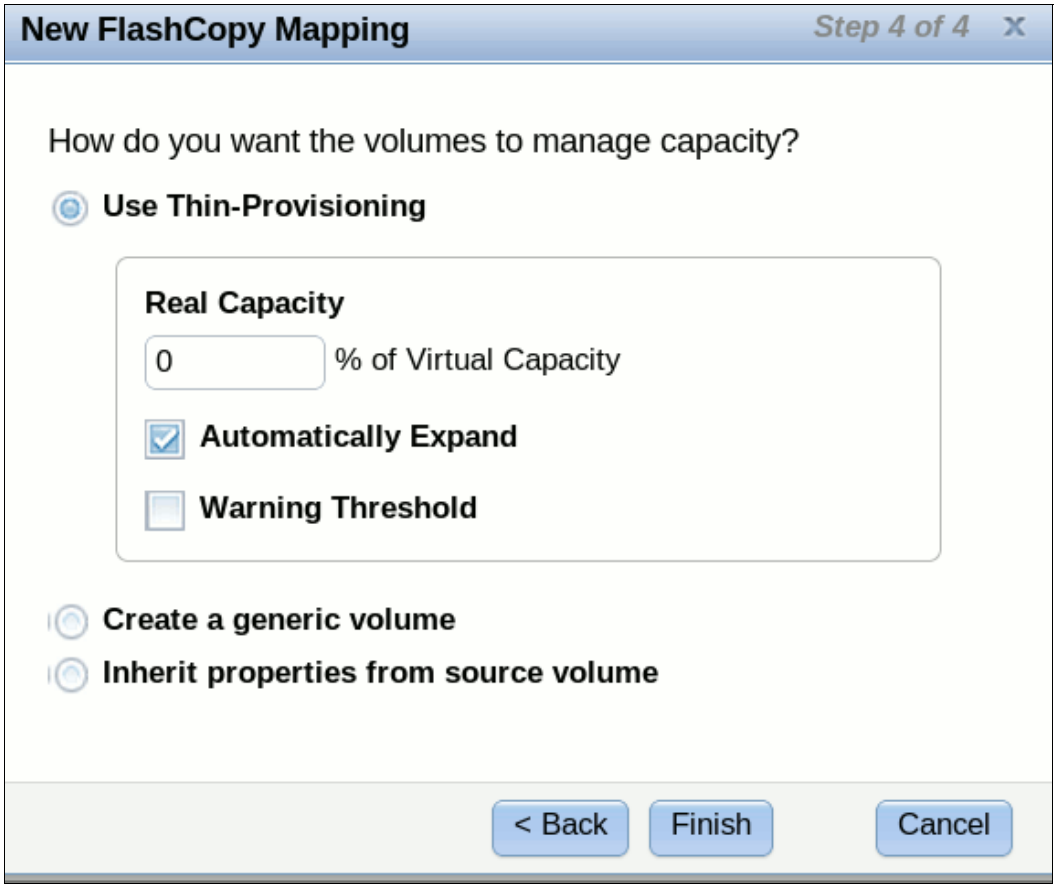

Figure 8-151 Thin-provisioning options

Click **Finish** after you select the options. The FlashCopy mapping is created on your volume with a new target, as shown in Figure 8-152. The status of the newly created FlashCopy mapping is Idle. To start it, see “Starting a FlashCopy mapping” on page 424

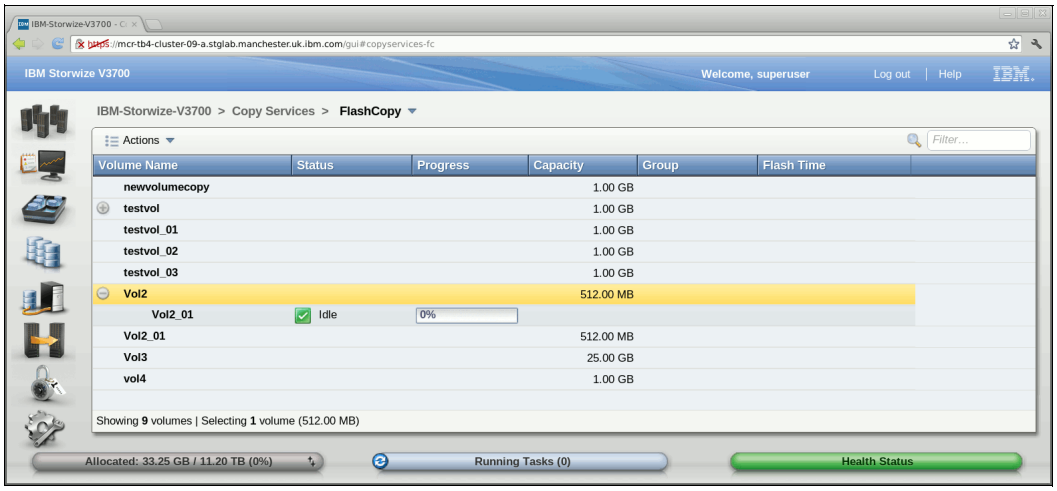

Figure 8-152 FlashCopy window

In the Advanced FlashCopy menu, if you already have candidate target volumes, select **Use existing target volumes**, as shown in Figure 8-153.

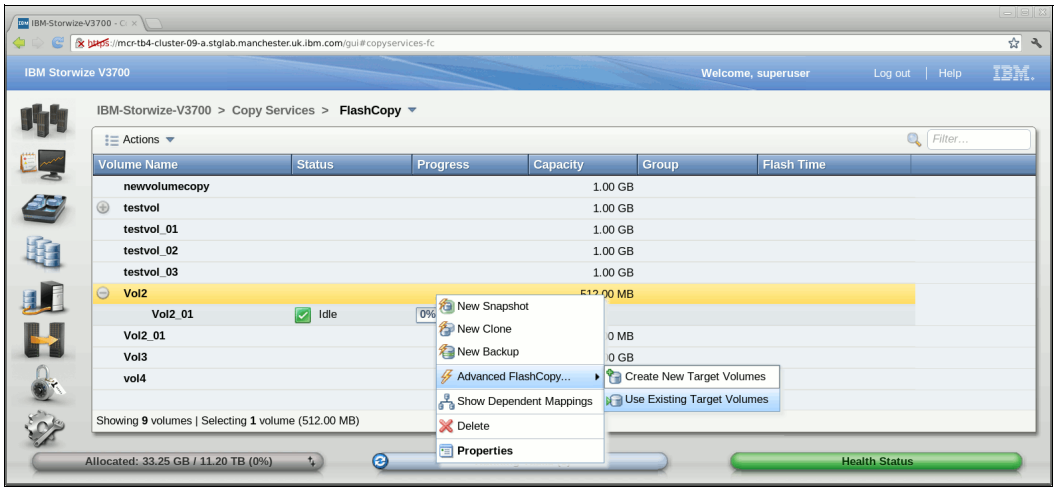

Figure 8-153 Use existing target volumes

You must choose the target volume for the source volume that you selected. Select the target volume from the menu in the right pane of the window and click **Add**, as shown in Figure 8-154.

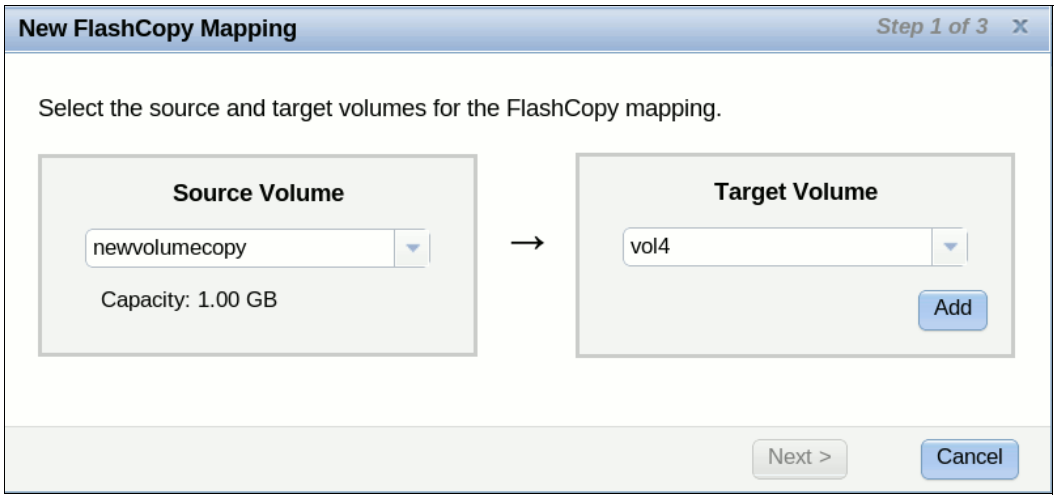

Figure 8-154 Select the target volume

After you click **Add**, the FlashCopy mapping is listed, as shown in Figure 8-155. 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.

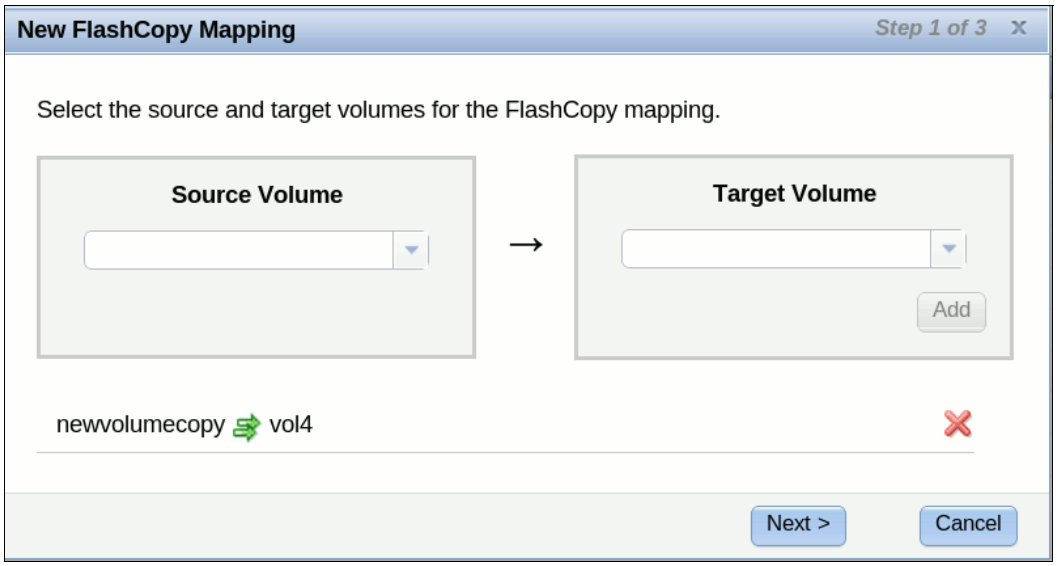

Figure 8-155 Add FlashCopy mapping

Select the preset and adjust the settings, as shown in Figure 8-156. Make sure that the settings meet your requirements. Click **Next**.

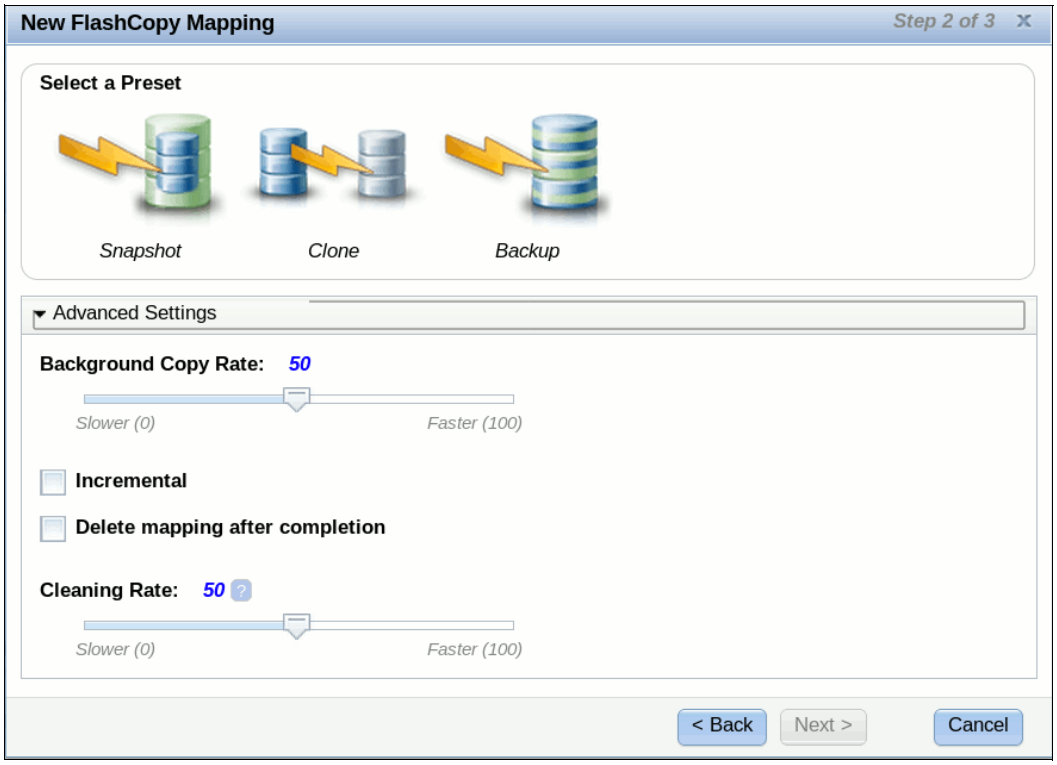

Figure 8-156 Select a preset

Now you can add the FlashCopy mapping to a consistency group, as shown in Figure 8-157. Click **Finish** and the FlashCopy mapping is created with the status of Idle.

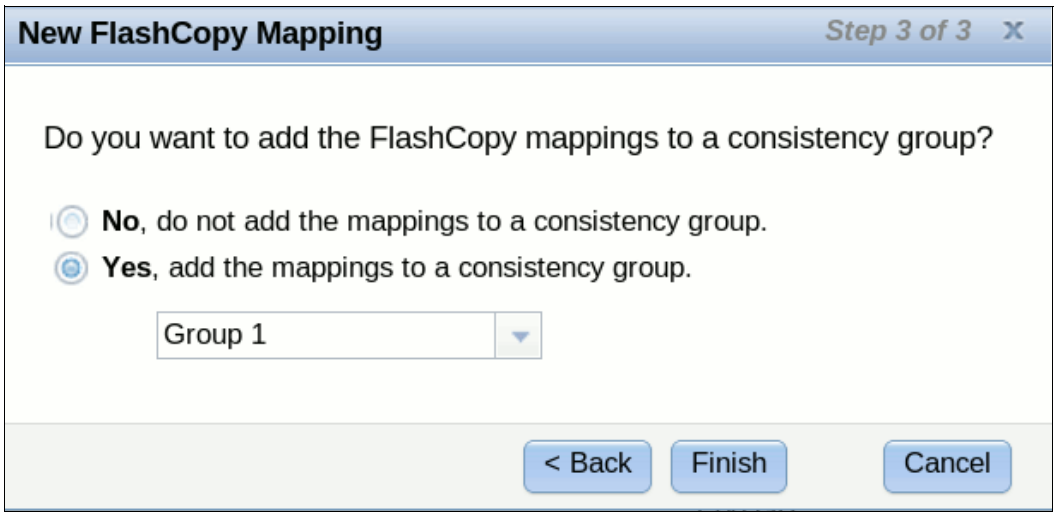

Figure 8-157 Select a consistency group

You also can create the FlashCopy mappings in the FlashCopy Mapping window by clicking **New FlashCopy Mapping** at the upper left, as shown in Figure 8-158.

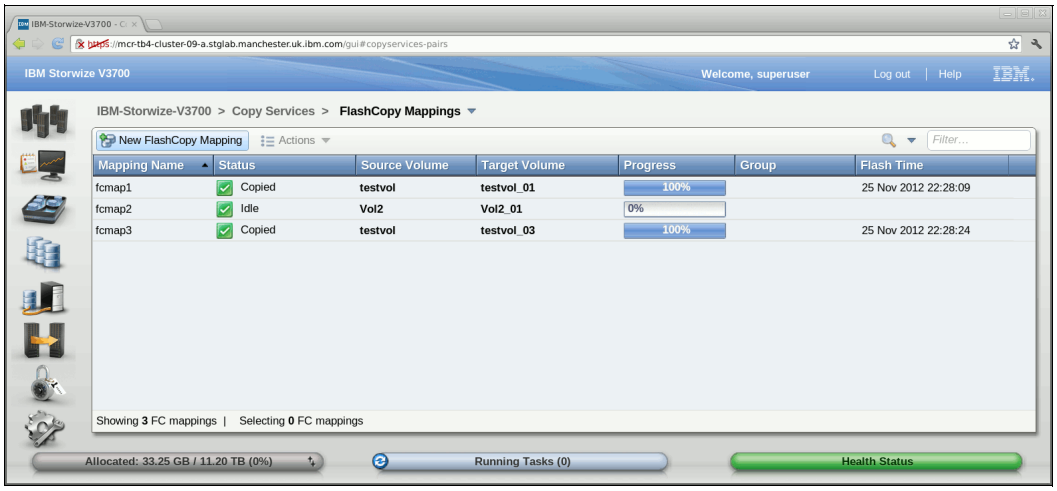

Figure 8-158 New FlashCopy mapping

A wizard opens to guide you through the creation of a FlashCopy mapping. The steps are the same as creating an advanced FlashCopy mapping by using existing target volumes in the FlashCopy window.

Starting a FlashCopy mapping

Most of the actions of FlashCopy mapping can be done 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. However, the steps are the same if you were to use the FlashCopy Mapping window.

The mapping can be started 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 8-159. The status of the FlashCopy mapping changes from Idle to Copying.

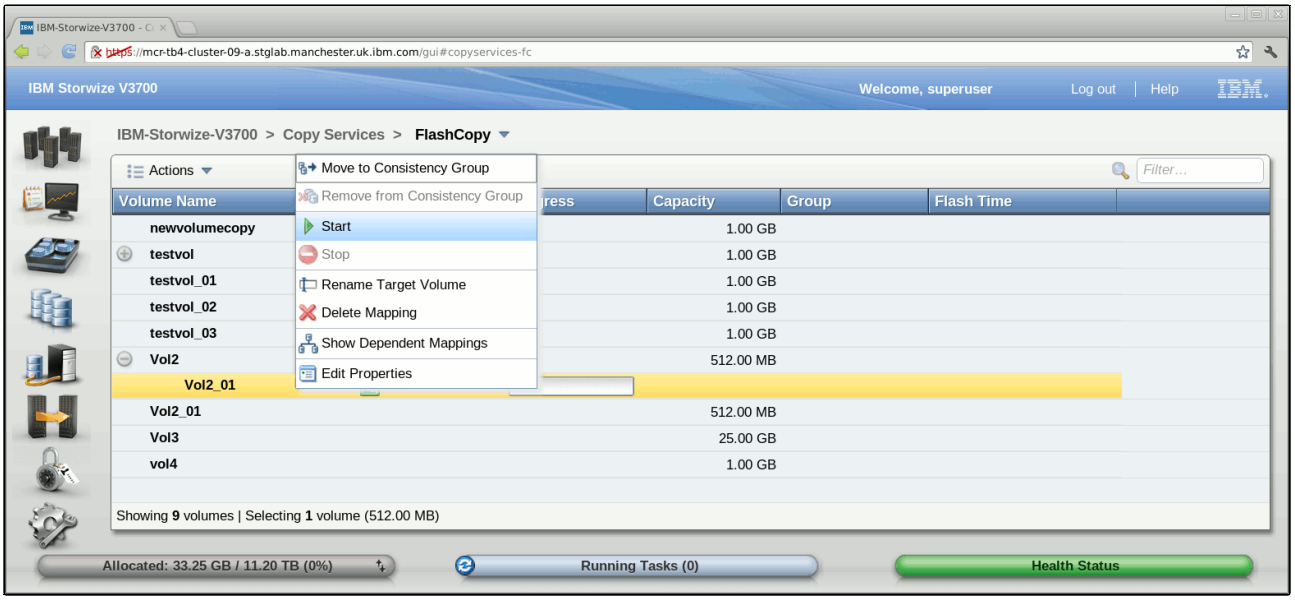

Figure 8-159 Start FlashCopy mapping

Stopping a FlashCopy mapping

The FlashCopy mapping can be stopped by selecting the FlashCopy target volume in the FlashCopy window and clicking **Stop** from the Actions drop-down menu, as shown in Figure 8-160. After the process stops, the status of the FlashCopy mapping is changed to Stopped.

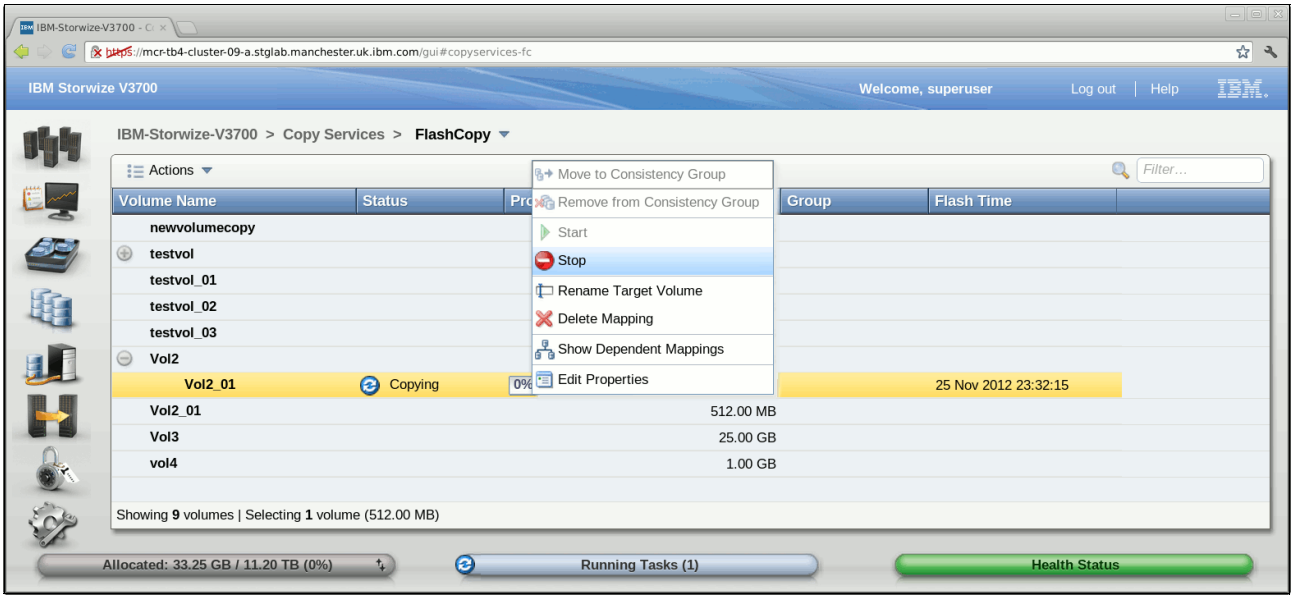

Figure 8-160 Stop a FlashCopy mapping

Renaming the target volume

If the FlashCopy target volumes were created automatically by the IBM Storwize V3500 storage system, the names of the target volumes are the source volume name plus a numerical suffix. It is possible to change the names of the target volumes to be more meaningful in your environment.

To change the name of a target volume, select it in the FlashCopy window and click **Rename Target Volume** from the Actions drop-down menu, as shown in Figure 8-161.

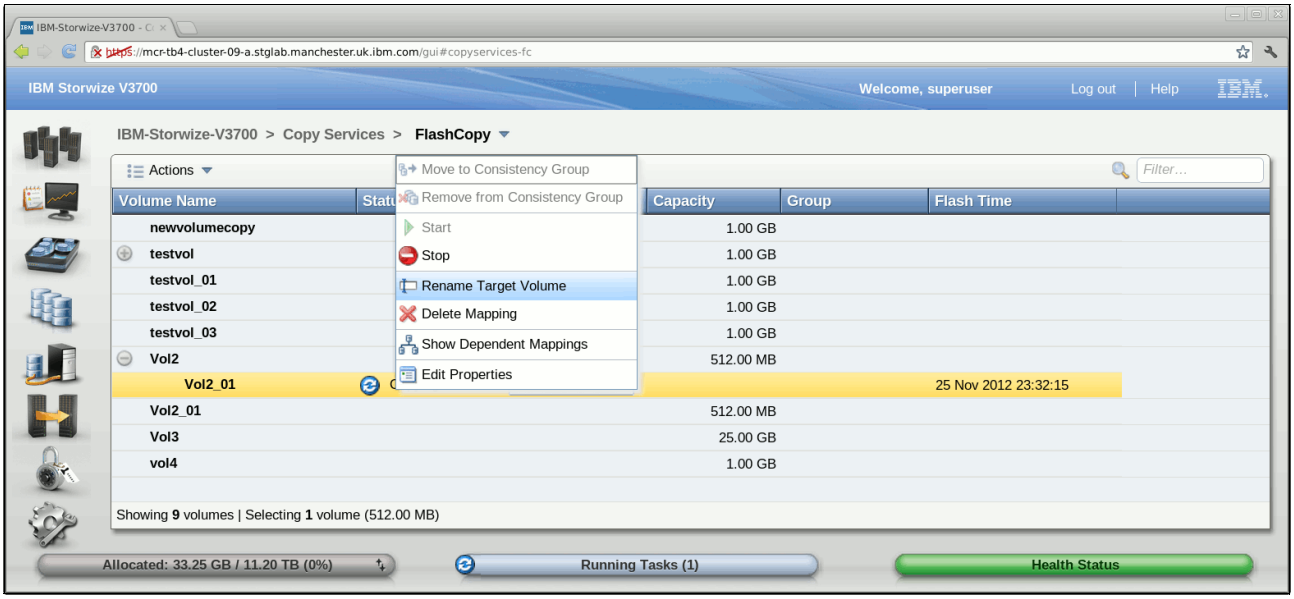

Figure 8-161 Rename Target Volume

Enter the new name for the target volume, as shown in Figure 8-162. Click **Rename** to finish.

Figure 8-162 Rename Volume window

Renaming a FlashCopy mapping

FlashCopy mappings are created with names that begin with `fcmap`. It is possible to replace these names with more meaningful ones.

To change the name of a FlashCopy mapping, select it in the FlashCopy Mappings window and click **Rename Mapping** from the Actions drop-down menu, as shown in Figure 8-163.

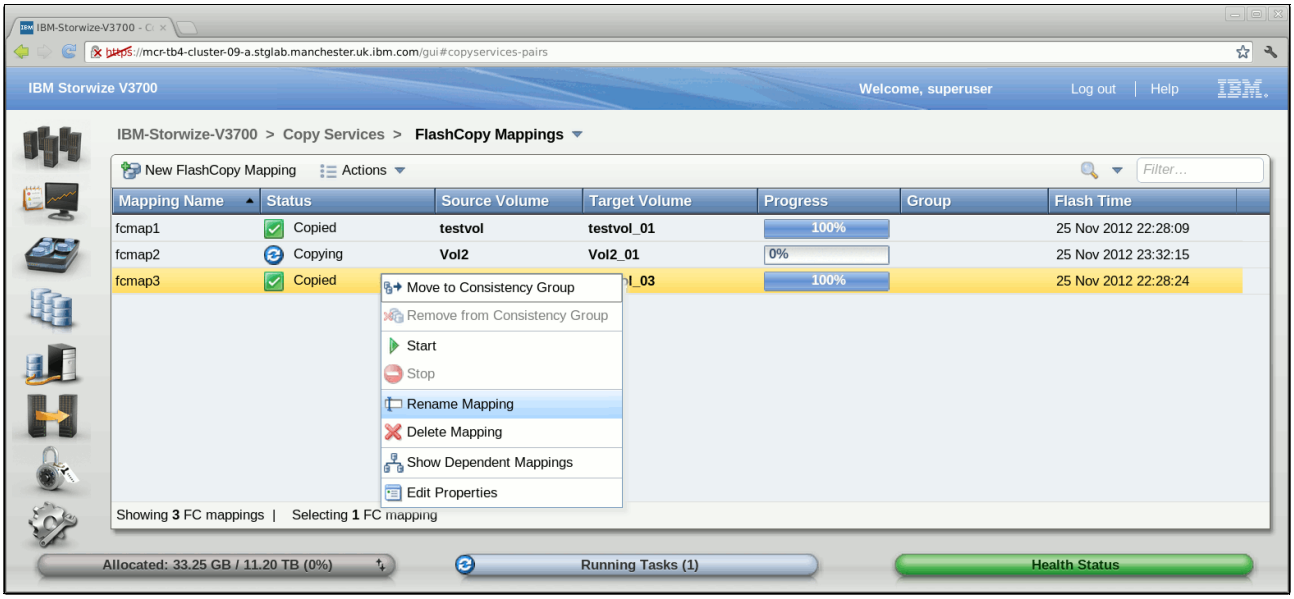

Figure 8-163 Rename Mapping

You must enter the new name for the FlashCopy mapping, as shown in Figure 8-164. Click **Rename** to finish.

Figure 8-164 Rename FlashCopy Mapping window

Deleting a FlashCopy mapping

The FlashCopy mapping can be deleted by selecting the target volume in the FlashCopy window and clicking **Delete Mapping** from the Actions drop-down menu, as shown in Figure 8-165.

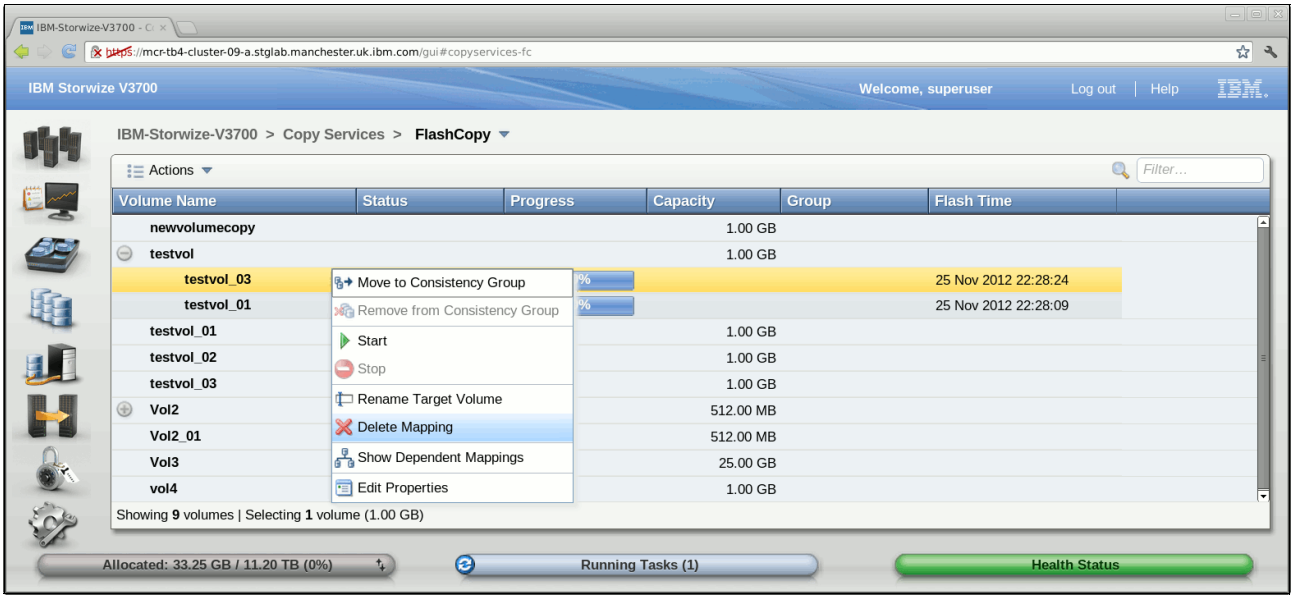

Figure 8-165 Delete Mapping

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 8-166. Verify the number of FlashCopy mappings to delete, and, if you want to delete the FlashCopy mappings while the data on the target volume is inconsistent with the source volume, select the option. Click **Delete** and your FlashCopy mapping is removed.

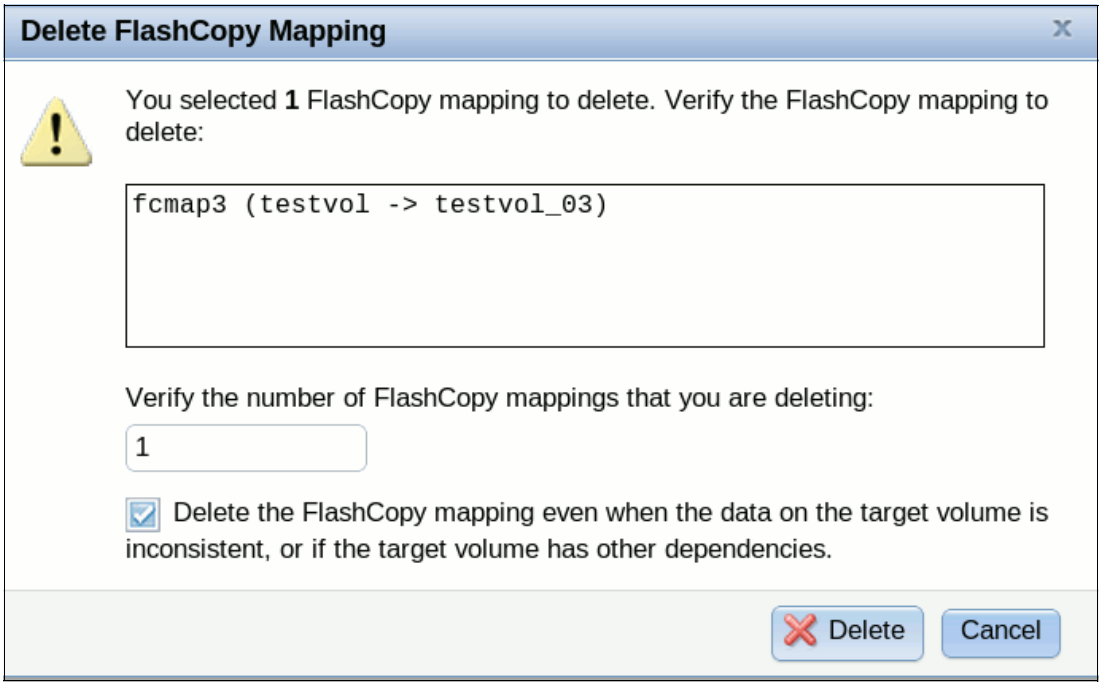

Figure 8-166 Delete FlashCopy Mapping window

Deleting FlashCopy mapping: Deleting the FlashCopy mapping does not delete the target volume. If you must reclaim the storage space that was occupied by the target volume, you must delete the target volume manually.

Showing dependent mapping

It is possible to show the FlashCopy mapping dependency by selecting the source volume in the FlashCopy window and clicking **Show Dependent Mappings** from the Actions drop-down menu, as shown in Figure 8-167.

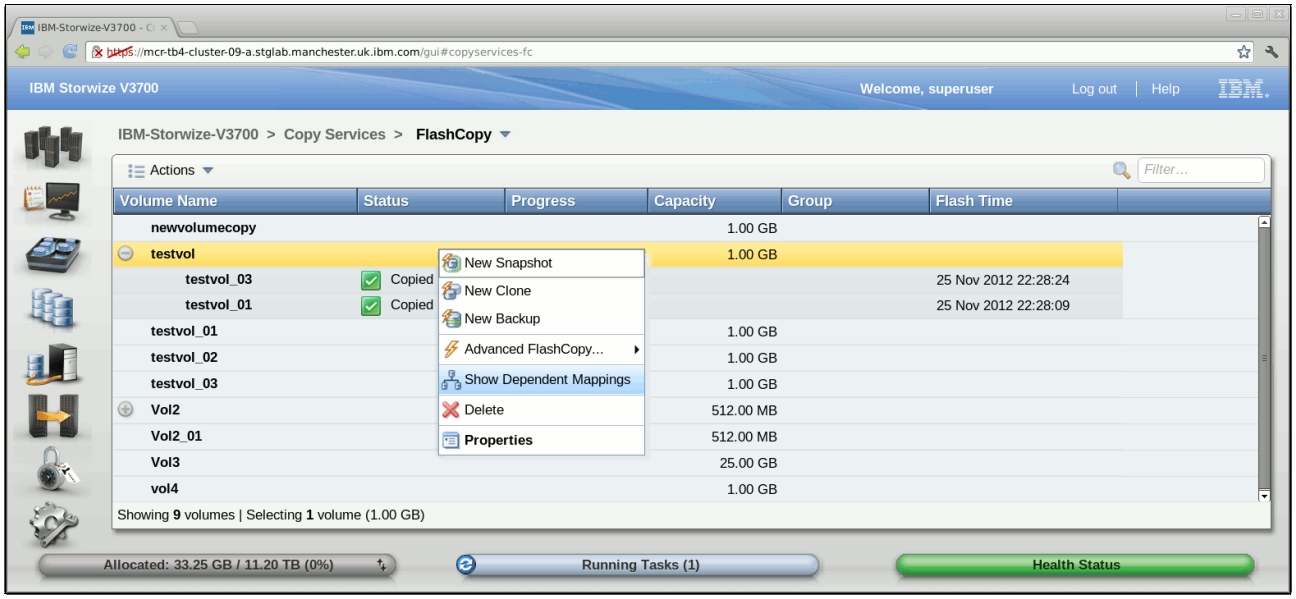

Figure 8-167 Show Dependent Mappings

The FlashCopy mapping dependency tree opens, as shown in Figure 8-168.

Figure 8-168 Dependent Mappings

Editing properties

After the FlashCopy mapping is created, the background copy rate and cleaning rate can be changed by selecting the target volume in the FlashCopy window and clicking **Edit Properties** from the Actions drop-down menu, as shown in Figure 8-169.

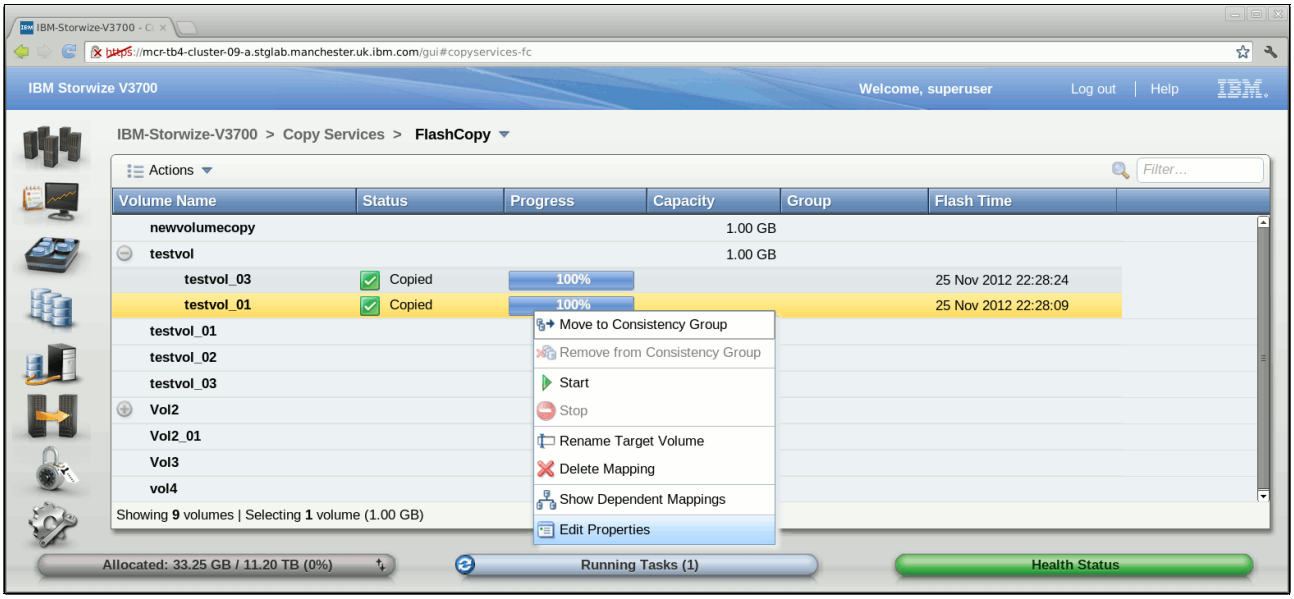

Figure 8-169 Edit properties

It is possible to modify the value of the background copy rate and cleaning rate by moving the pointers on the bars, as shown in Figure 8-170. Click **Save** to save changes.

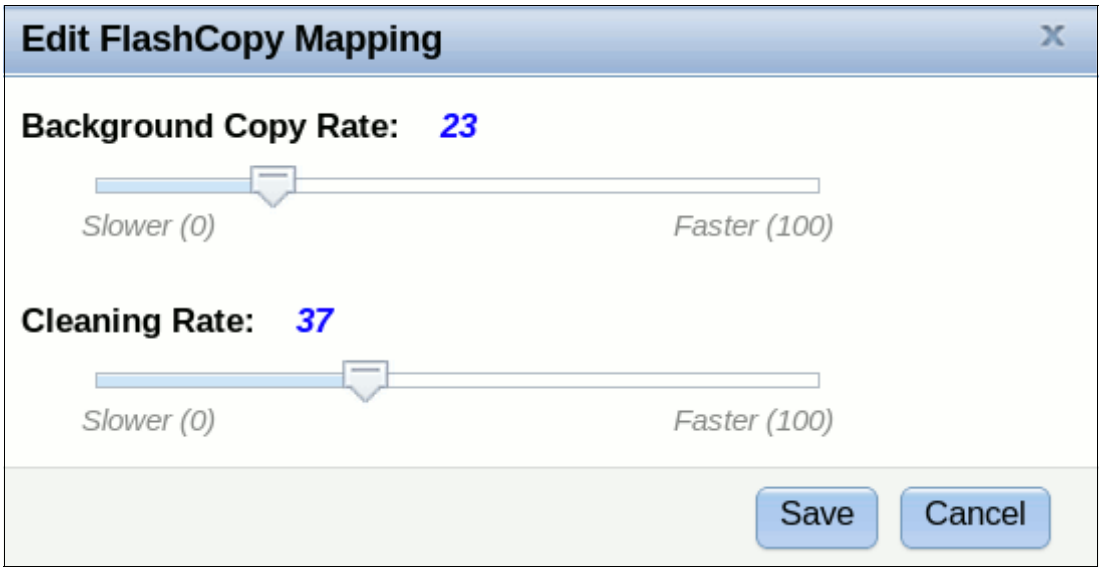

Figure 8-170 Edit FlashCopy Mapping

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 8-171.

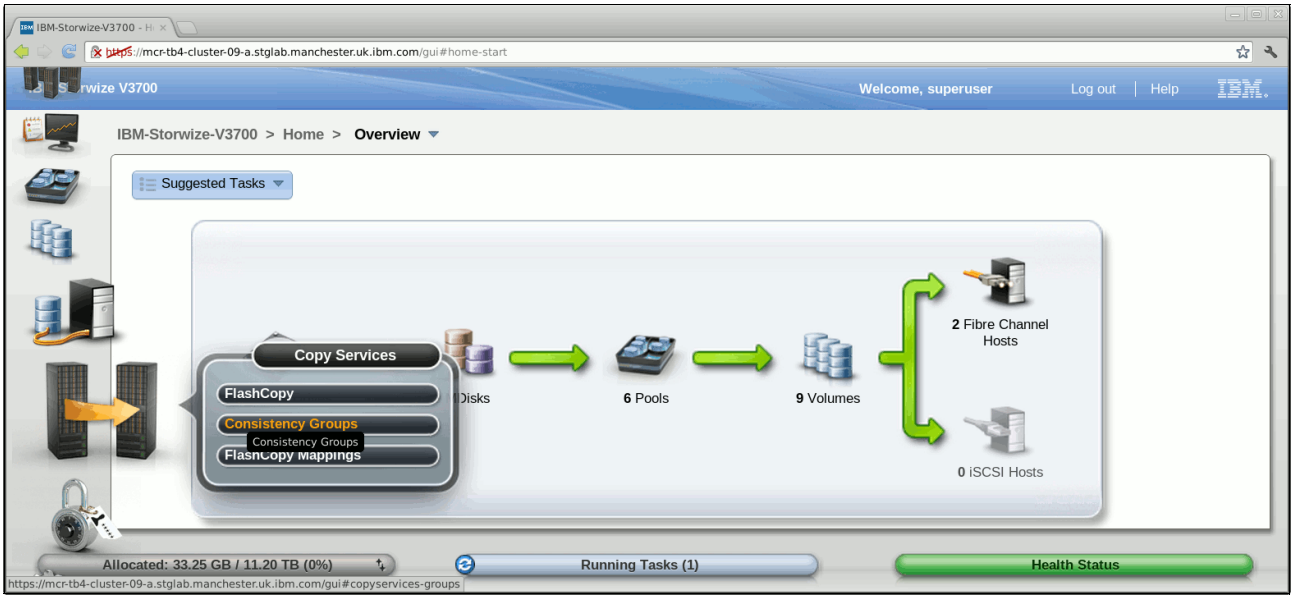

Figure 8-171 Consistency Groups

Consistency groups and FlashCopy mappings can be managed in the Consistency Groups window, as shown in Figure 8-172.

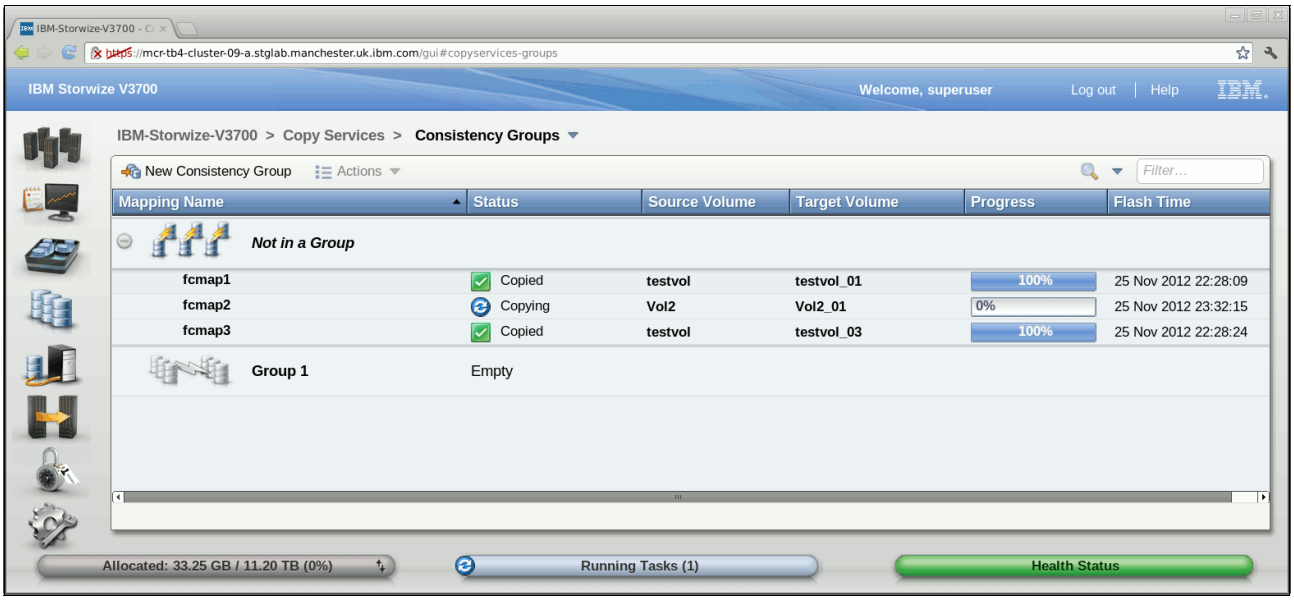

Figure 8-172 Consistency Groups

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 **Not in a Group**. All the FlashCopy mappings that are not in any consistency groups are displayed in the lower pane.

In the lower pane of the Consistency Groups window, you can see the properties of a consistency group and the FlashCopy mappings in it. You also can take actions on any consistency groups and FlashCopy mappings within the Consistency Groups window, as allowed by their state. All the actions that are allowed for the FlashCopy mapping are described in 8.9.6, “Managing FlashCopy mapping” on page 414.

Creating a FlashCopy consistency group

To create a FlashCopy consistency group, click **New Consistency Group** at the top of the Consistency Groups window (see Figure 8-172 on page 432). You are prompted to enter the name of the new consistency group, as shown in Figure 8-173. Following your naming conventions, enter the name of the new consistency group and click **Create**.

Figure 8-173 shows the New Consistency Group window.

Figure 8-173 New Consistency Group

After the creation process completes, your new consistency group is displayed in the consistency groups window, as shown in Figure 8-174.

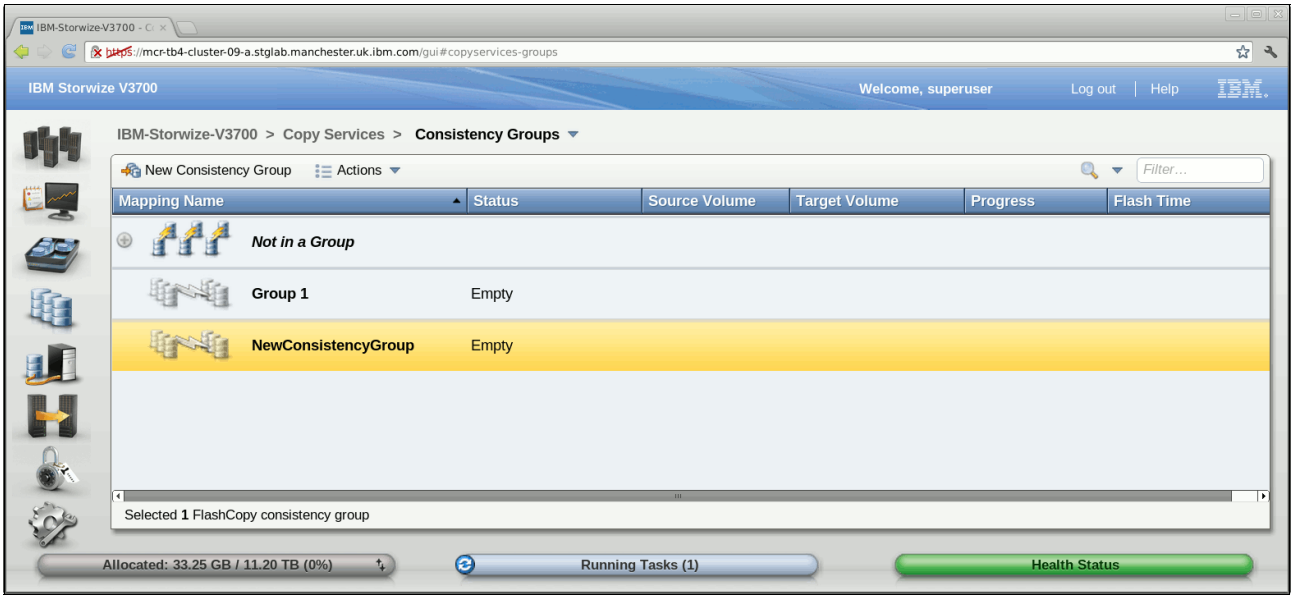

Figure 8-174 The new consistency group

It is possible to rename the consistency group by clicking the name of the consistency group in the main pane and entering a name for it (following your naming convention). 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.

Adding FlashCopy mappings to a consistency group

Click **Not in a Group** to list all the FlashCopy mappings with no consistency group. It is possible to 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 8-175. To select a range, highlight a mapping, press and hold the Shift key, and click the last item in the range. Multiple selections can be made by pressing and holding the Ctrl key and clicking each mapping individually.

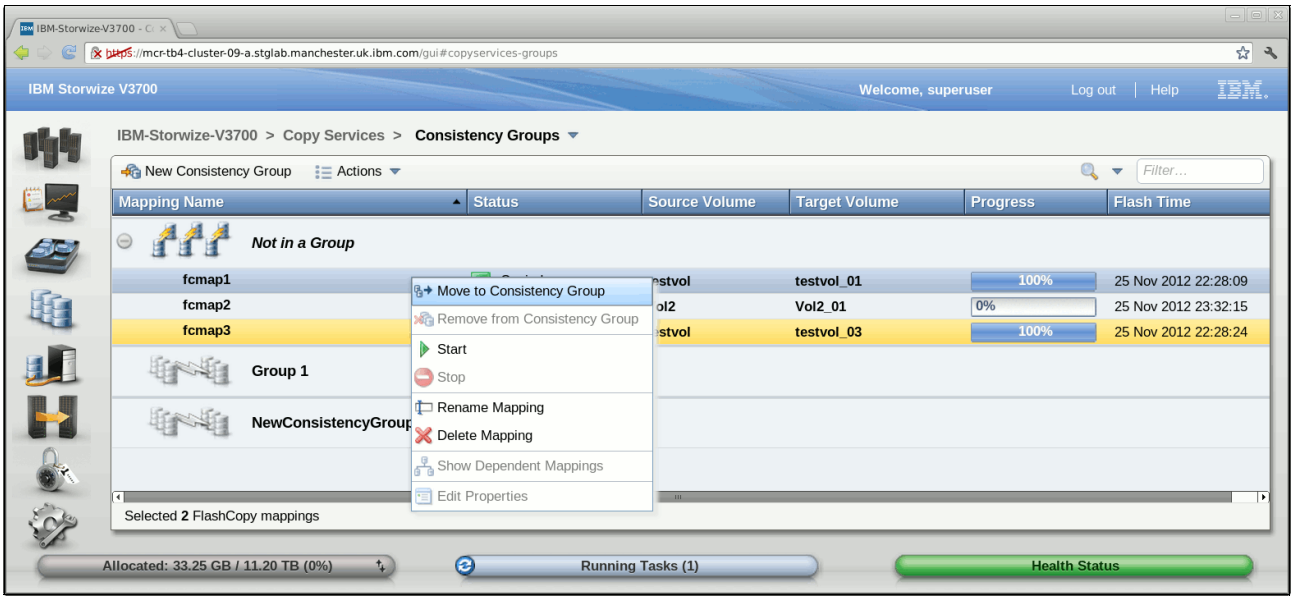

Figure 8-175 Multiple selected consistency groups

You are prompted to specify the consistency group into which you want to move the FlashCopy mapping, as shown in Figure 8-176. Click **Move to Consistency Group** to continue.

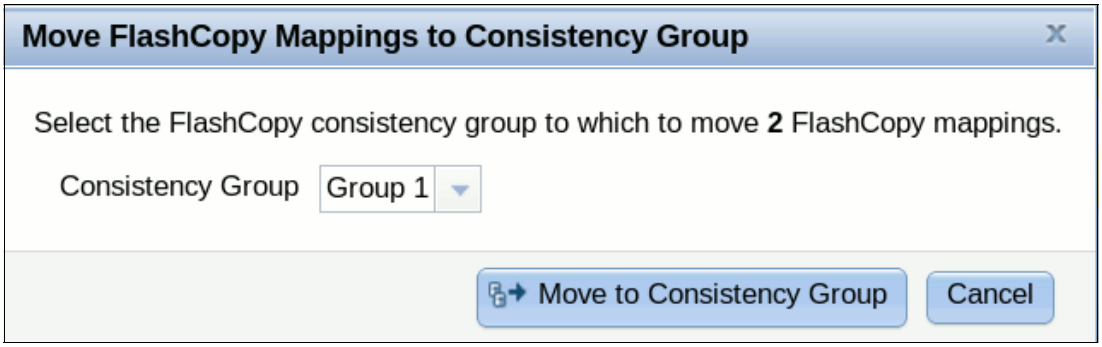

Figure 8-176 Move FlashCopy mappings to consistency group

After the action completes, the FlashCopy mappings you selected is moved from the Not In a Group list to the consistency group that you chose.

Starting a consistency group

To start a consistency group, click **Start** from the Actions drop-down menu, as shown in Figure 8-177.

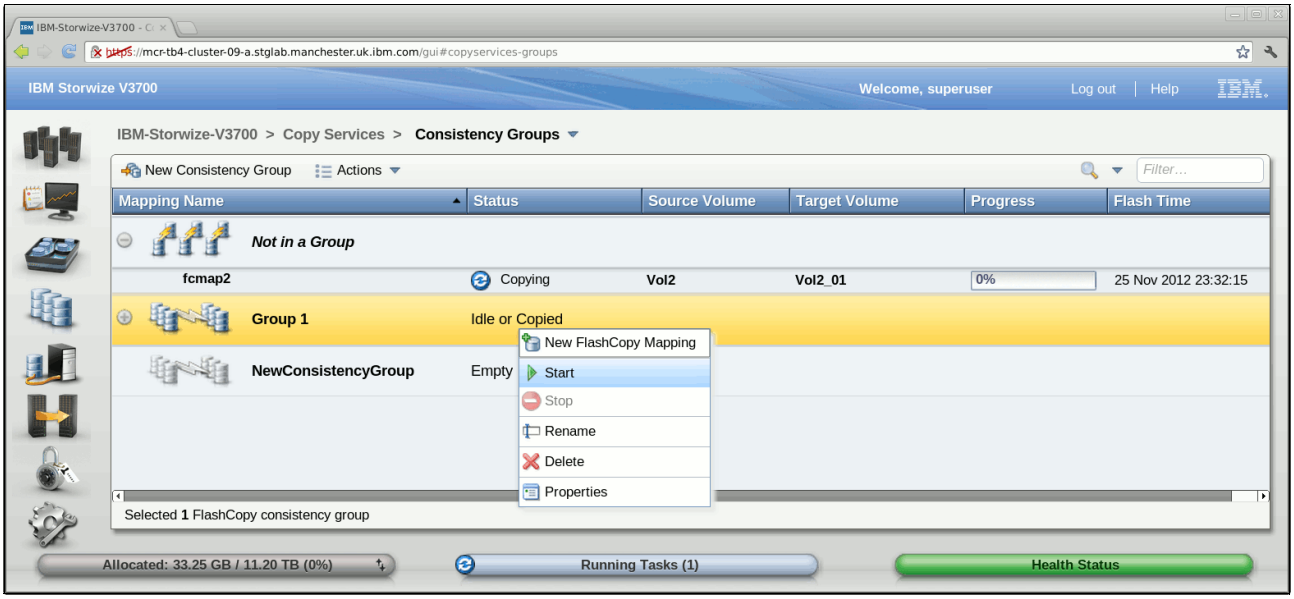

Figure 8-177 Start consistency group

After you start the consistency group, all the FlashCopy mappings in the consistency group are started. The state of the FlashCopy changes to Copying, as shown in Figure 8-178.

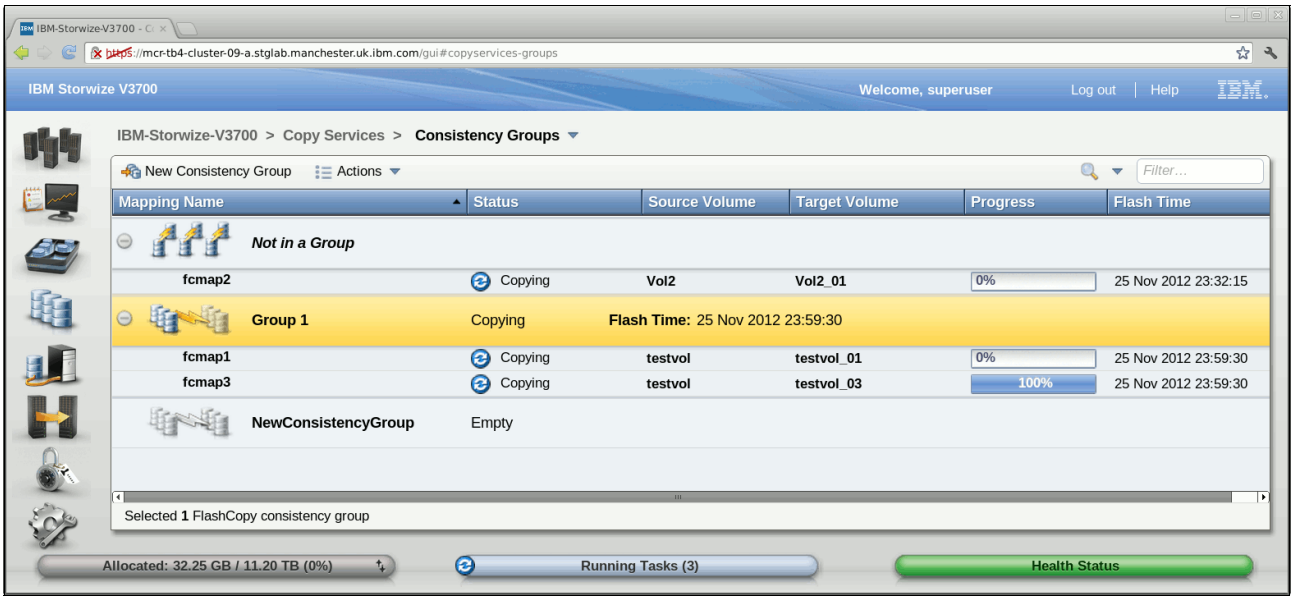

Figure 8-178 Consistency group copying

Stopping a consistency group

The consistency group can be stopped by selecting **Stop** from the Actions drop-down window, as shown in Figure 8-179.

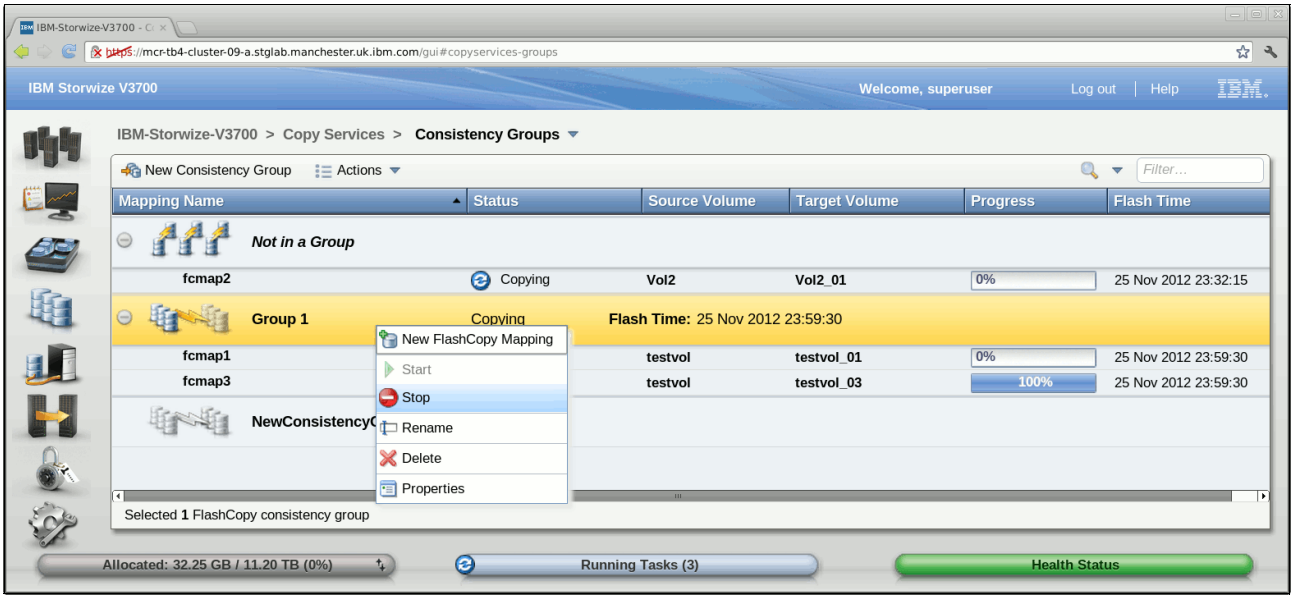

Figure 8-179 Stop consistency group

After the stop process completes, the FlashCopy mappings in the consistency group are in the Stopped state. A red “X” icon appears on the function icon of this consistency group to indicate an alert, as shown in Figure 8-180.

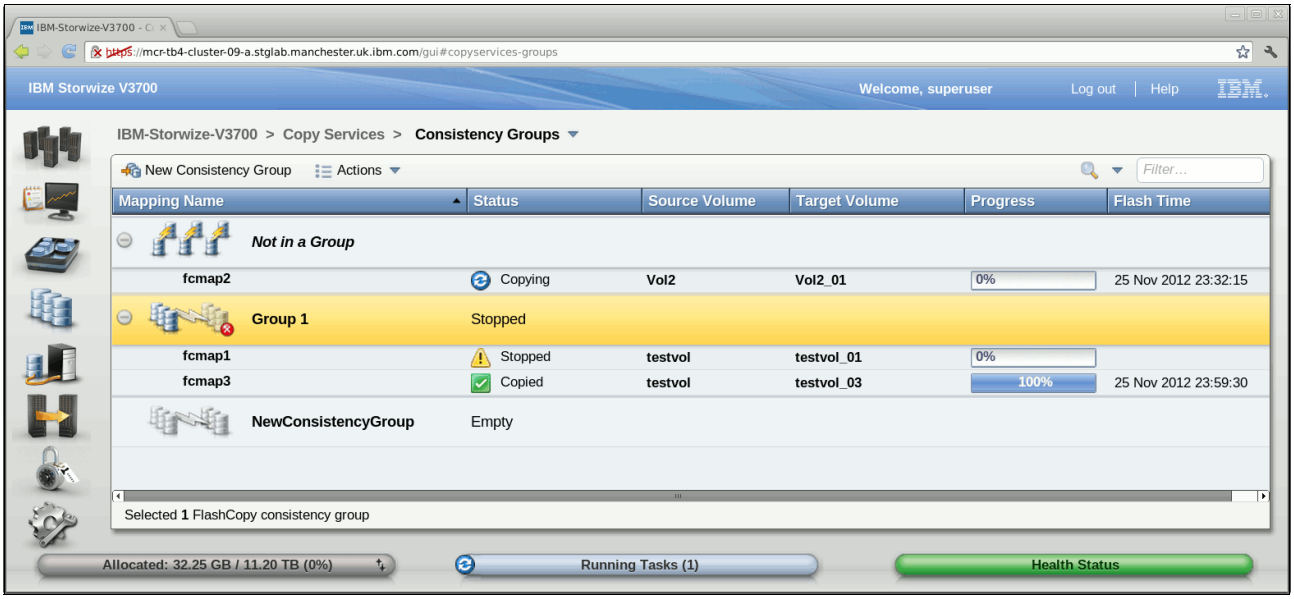

Figure 8-180 Consistency group stop completes

Previously copied relationships that were added to a consistency group that was later stopped before all members of the consistency group completed synchronization do not go out of the Copied state.

Removing FlashCopy mappings from a consistency group

FlashCopy mappings can be removed from a consistency group by selecting the mappings and clicking **Remove from Consistency Group** from the Actions drop-down menu of the FlashCopy mapping, as shown in Figure 8-181.

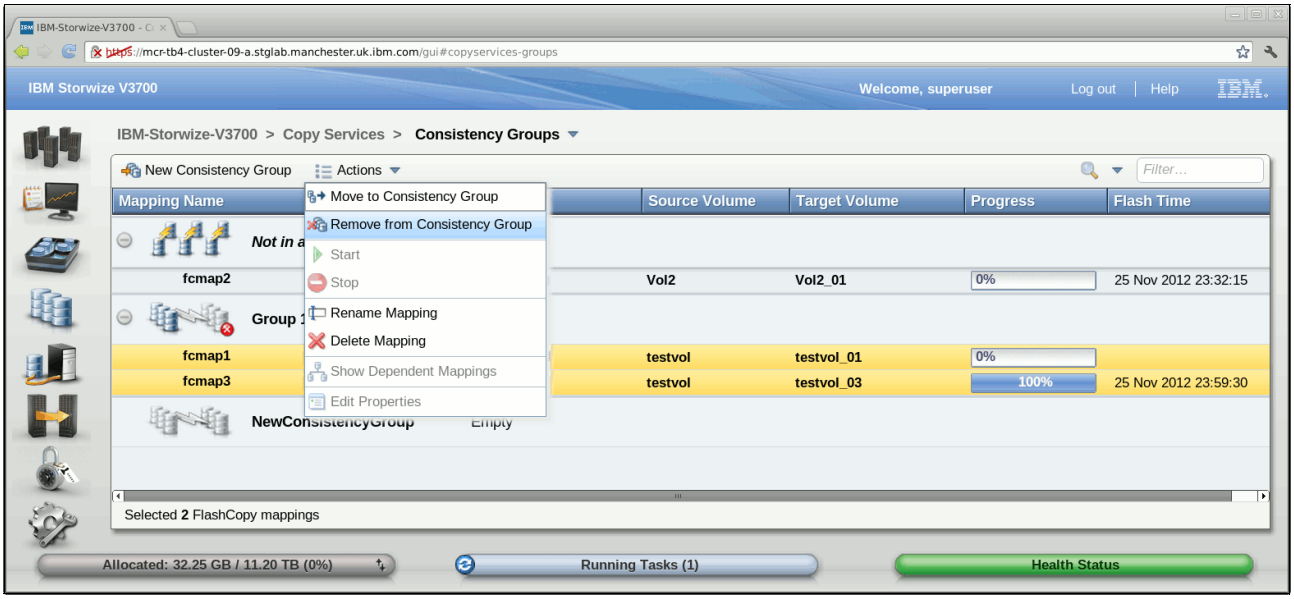

Figure 8-181 Remove from consistency group

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, as shown in Figure 8-182.

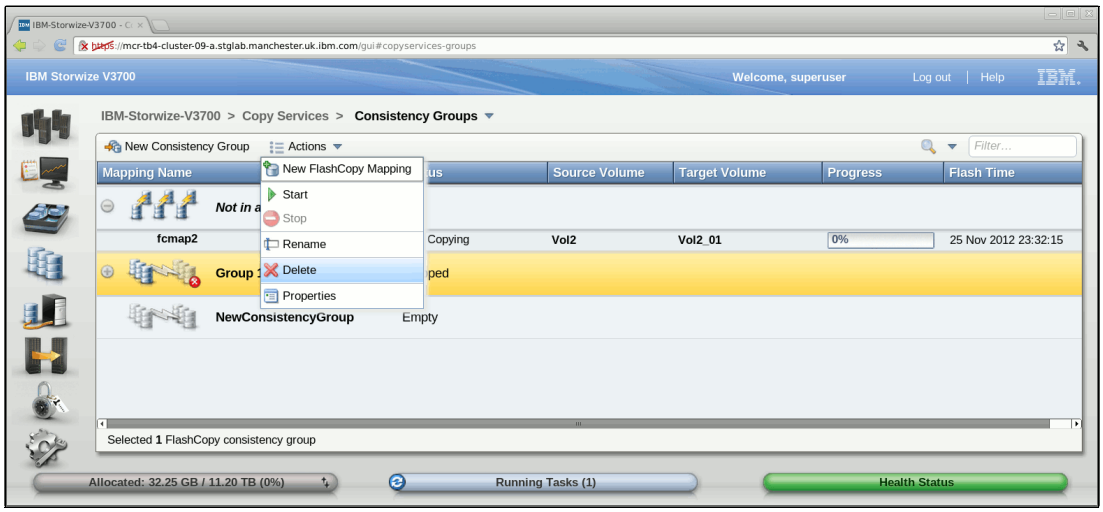

Figure 8-182 Delete consistency group

9

RAS, monitoring, and troubleshooting

There are many ways to monitor and troubleshoot the IBM Storwize V3500. In this chapter, we show the ways in which the Storwize V3500 can be administered from a monitoring and troubleshooting point of view.

This chapter includes the following sections:

- Reliability, availability, and serviceability (RAS)
- Hardware and LEDs descriptions
- Monitoring from a host
- Monitoring from the IBM Storwize V3500
- Backup procedure
- Software upgrade procedure
- Troubleshooting
- Recommended actions
- Event log navigation
- Audit log navigation
- Shutting down the IBM Storwize V3500

9.1 Reliability, availability, and serviceability on the IBM Storwize V3500

This section describes the Reliability, availability, and serviceability (RAS) features of the IBM Storwize V3500, monitoring, and troubleshooting. RAS features are important concepts in the design of the IBM Storwize V3500. Hardware features, software features, design considerations, and operational guidelines all contribute to make the IBM Storwize V3500 reliable.

Fault tolerance and high levels of availability are achieved by the following features:

- The RAID capabilities of the underlying disk subsystems
- The compass architecture that is employed by the Storwize V3500 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 functions:

- 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 V3500 components
- Concurrent upgrade of IBM Storwize V3500 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 the case of a network connection problem

At the heart of the IBM Storwize V3500 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.

9.2 Storwize V3500 components

This section describes each of the components that make up the Storwize V3500 system. Components are examined in terms of location, function, and serviceability.

9.2.1 Enclosure midplane assembly

The enclosure midplane assembly is the unit that contains the node canisters and the power supply units. During the basic system configuration, Vital Product Data (VPD) is written to the enclosure midplane assembly.

Control enclosure midplane

The *control* enclosure midplane holds node canisters and the power supply units. The control enclosure midplane assembly includes 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 by a trained service provider only.

Figure 9-1 shows the Enclosure Midplane Assembly back side.

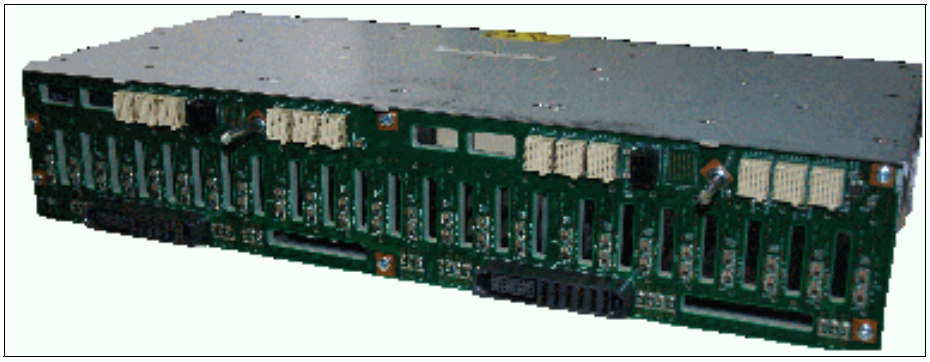

Figure 9-1 Enclosure Midplane Assembly back side

For more information about replacing the control enclosure midplane, see the IBM Storwize V3500 Information Center at this website:

http://pic.dhe.ibm.com/infocenter/storwize/V3500_ic/index.jsp

9.2.2 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 9-2 shows the back of a fully equipped node enclosure.

Figure 9-2 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 9-3 shows the USB ports.

Figure 9-3 Node Canister USB ports

Ethernet ports

There are two 10/100/1000 Mbps Ethernet ports side-by-side on the canister. They are numbered 1 on the left (required) and 2 on the right (optional). The ports are shown in Figure 9-4.

Figure 9-4 Node canister Ethernet ports

Each port has two LEDs and their status are shown in Table 9-1.

Table 9-1 Ethernet LEDs status

LED	Color	Meaning
Link state	Green	It is on when there is an Ethernet link.
Activity	Yellow	It is flashing when there is activity on the link.

SAS ports

There are three 6-Gbps Serial Attached SCSI (SAS) ports side-by-side on the canister. They are numbered 1 on the left and 3 on the right and are not used now. They are shown in Figure 9-5.

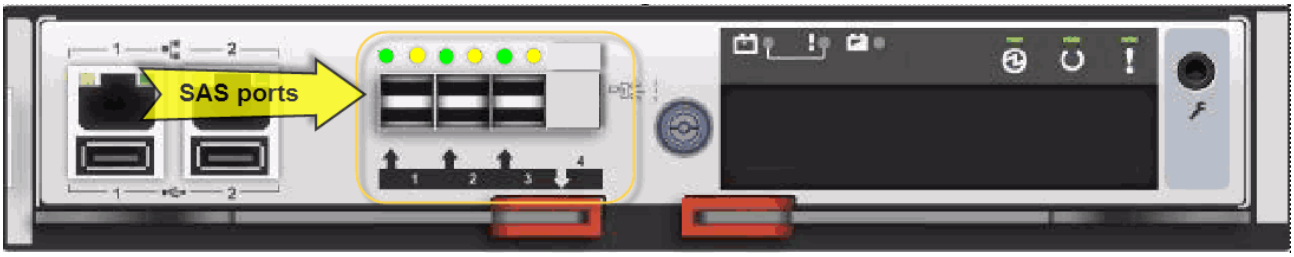

Figure 9-5 Canister SAS ports

Battery status

Each node canister stores a battery the status of which is displayed on three LEDs on the back of the unit. Figure 9-6 is showing the battery indicator LEDs.

Figure 9-6 Node canister battery status

The battery indicator status meanings are shown in Table 9-2.

Table 9-2 Battery indicator on Node canister

Color	Name	Definition
Green (left)	Battery Status	► Fast flash: Indicates the battery is charging and has insufficient charge to complete a single dump. ► Flash: Indicates the battery has sufficient charge to complete a single dump only. ► Solid: Indicates the battery is fully charged and thus has sufficient charge to complete two dumps.
Amber	Fault	Indicates a fault with the battery.
Green (right)	Battery in use	Indicates that hardened/critical data is being written to disk.

Canister status

The status of each canister is displayed by three LEDs, as shown in Figure 9-7.

Figure 9-7 System status indicator

The system status LED meanings are shown in Table 9-3.

Table 9-3 System status indicator

Color	Name	Definition
Green (left)	System Power	► Flash: The canister is in standby mode, which means that V3500 is not running. ► Fast flash: Running a self test. ► On: Powered up and the V3500 code is running.
Green (mid)	System Status	► Off: There is no power to the canister, the canister is in standby mode, Power On Self Test (POST) is running on the canister, or the operating system is loading. ► Flash: The node is in candidate or service state, which means it cannot perform I/O. It is safe to remove the node. ► Fast flash: A code upgrade is running. ► On: Node is part of a cluster.
Amber	Fault	► Off: The node is in candidate or active state. This status does not mean there is no hardware error on the node. Any error that was detected is not severe enough to stop the node from participating in a cluster (or there is no power). ► Flash: Identifies the canister. ► On: The node is in service state, or there is an error that is stopping the software from starting.

9.2.3 Node canister replaceable hardware components

The IBM Storwize V3500 node canister contains the following replaceable components:

- Host Interface Card
- Memory
- Battery

All of the components are customer-replaceable parts. Figure 9-8 shows the location of these parts within the node canister.

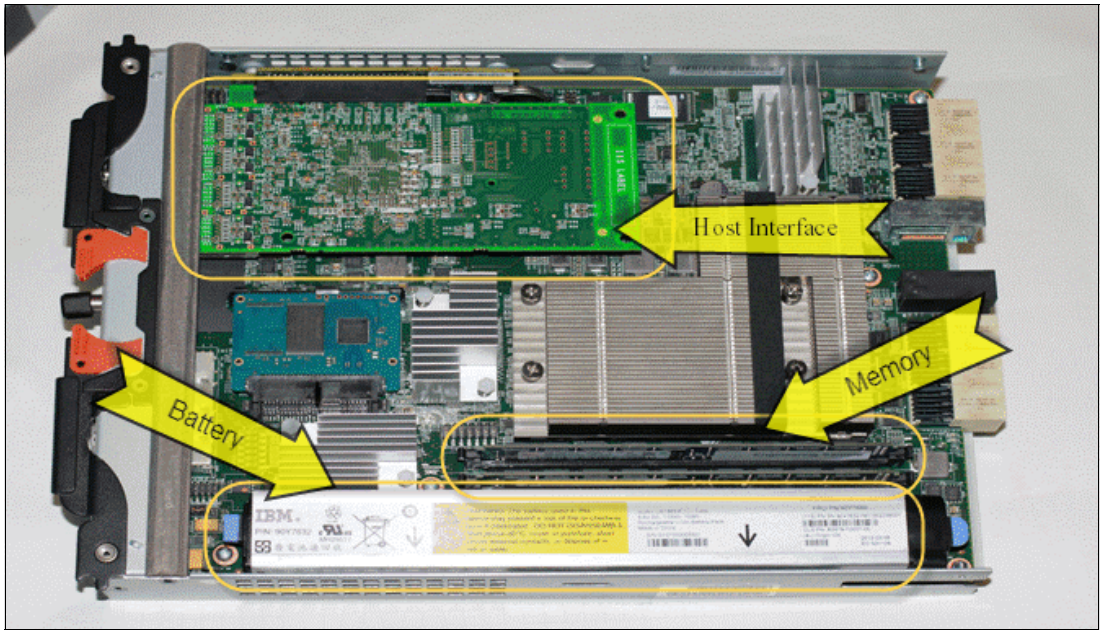

Figure 9-8 Node canister customer replaceable parts

Host Interface card replacement

For more information about the replacement process, see the IBM Storwize V3500 Information Center at this website:

http://pic.dhe.ibm.com/infocenter/storwize/V3500_ic/topic/com.ibm.storwize.V3500.641.doc/V3500_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 9-9.

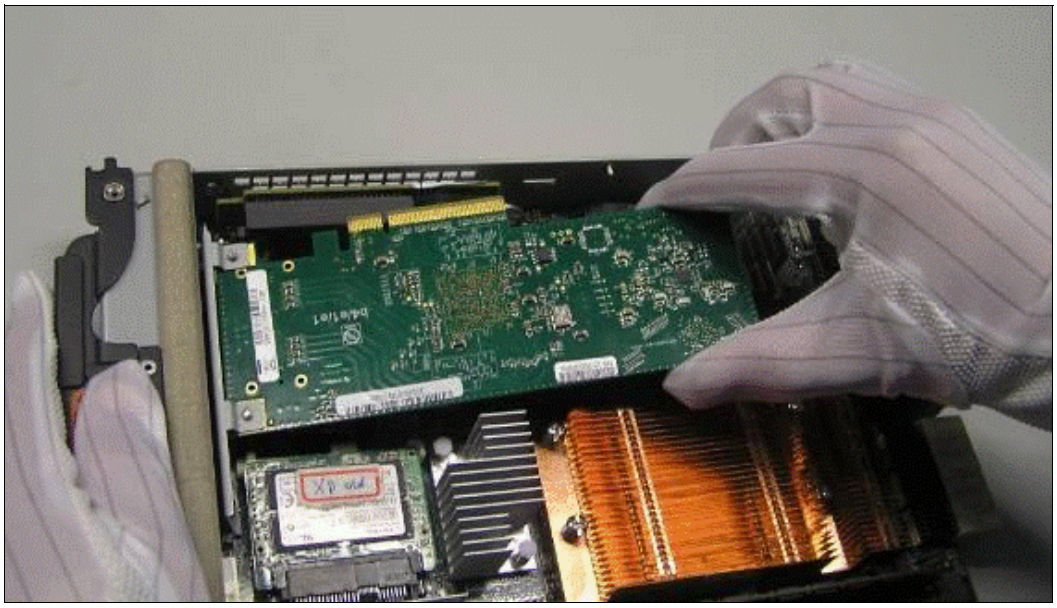

Figure 9-9 Host Interface card replacement

Memory replacement

For more information about the memory replacement process, see the IBM Storwize V3500 Information Center at this website:

http://pic.dhe.ibm.com/infocenter/storwize/V3500_ic/topic/com.ibm.storwize.V3500.641.doc/V3500_rplc_nodecan_dimm.html

At the website, browse to **Troubleshooting Removing and replacing parts** → **Replacing the node canister memory (4 GB DIMM)**.

Figure 9-10 shows the memory location.

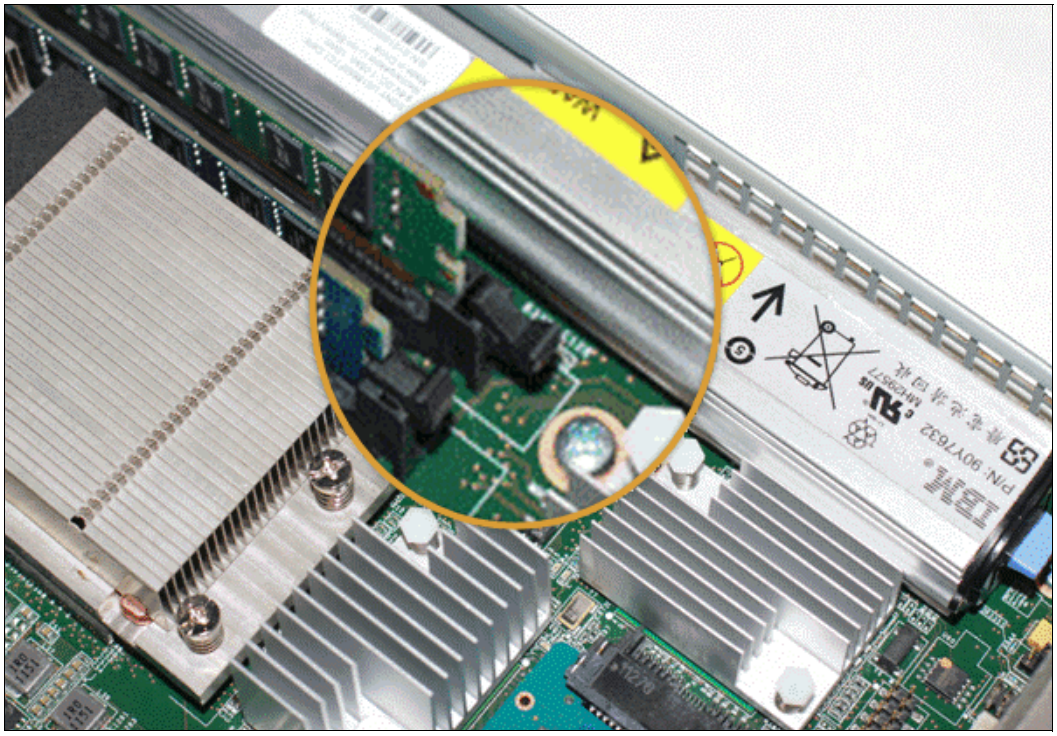

Figure 9-10 Memory replacement

Battery Backup Unit replacement

Caution: The battery is a lithium ion battery. To avoid possible explosion, do not incinerate the battery. Exchange the battery with the IBM approved part only.

Because the Battery Backup Unit (BBU) is seated in the node canister, the BBU replacement leads to a redundancy loss until the replacement is finished.

Therefore, it is recommended to replace a BBU only when it is 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/V3500_ic/topic/com.ibm.storwize.V3500.641.doc/V3500_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 9-11.

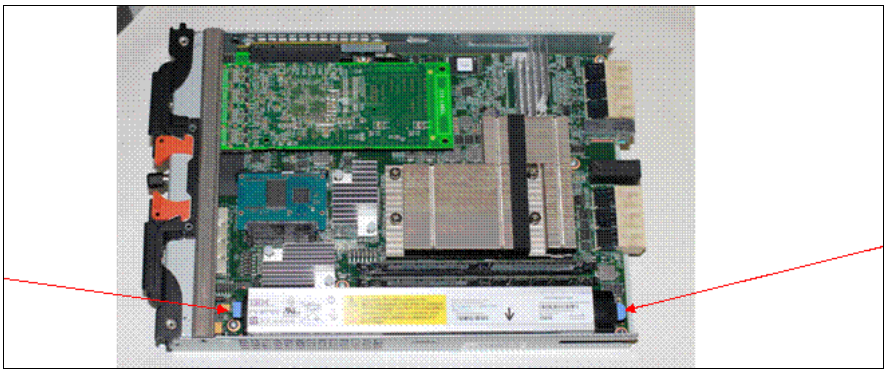

Figure 9-11 BBU replacement: Step 1

2. Lift the battery vertically upwards until the connectors disconnect.

Important: The battery must be kept parallel to the canister system board while it is removed or replaced (as shown in Figure 9-12). Keep equal force, or pressure, on each end.

Figure 9-12 BBU replacement: Step 2

9.2.4 Disk subsystem

The disk subsystem of the IBM Storwize V3500 system is made up of one node enclosure that contains the drives that are based on the enclosure type.

This section describes the all of the parts of the disk subsystem.

Slot numbers in enclosures

The IBM Storwize V3500 is made up of two types of enclosures, as described in Table 9-4.

Table 9-4 Enclosure slot numbering

Enclosure type	Number of slots
Enclosure 12x 3.5-inch drives Control enclosure 2072-12C	Enclosure with 12 slots. 
Enclosure 24x 2.5-inch drives Control enclosure 2072-24C	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 V3500 supports the RAID 0, RAID 1, RAID 5, RAID 6 or RAID 10. Each RAID level is described in Table 9-5.

Table 9-5 RAID levels that are supported by an IBM Storwize V3500

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
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 LBAs are readable and array parity is in synchronization. Arrays are scrubbed independently and each array is entirely scrubbed every seven days.

9.2.5 Power Supply Unit

The enclosure requires two power supply units (PSUs) for normal operation. A single PSU can power the entire enclosure for redundancy.

Figure 9-13 shows the power supplies.

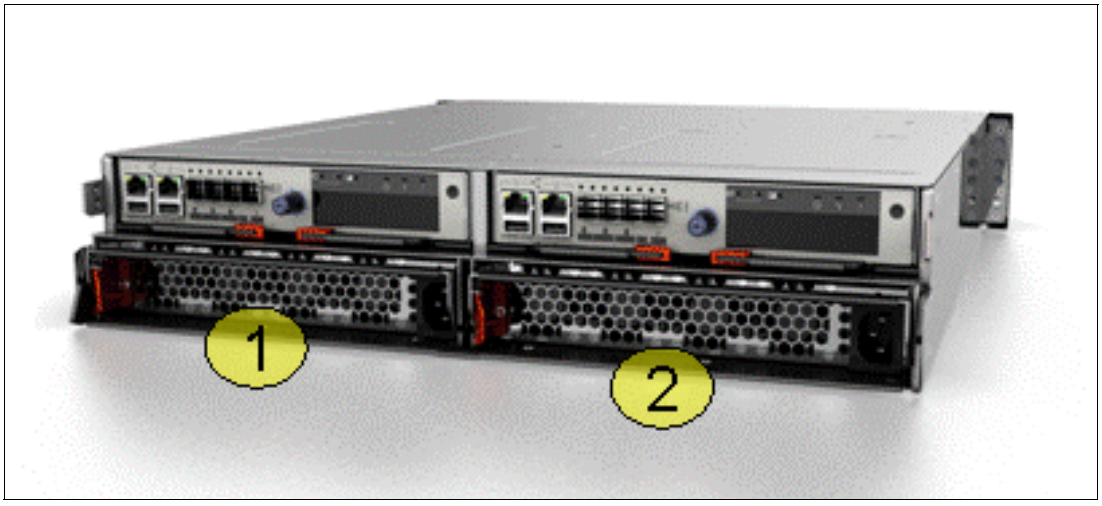

Figure 9-13 Power supplies

The left side PSU is numbered 1 and the right side PSU is numbered 2.

PSU LED indicator

Figure 9-14 shows the PSU LED indicators.

Figure 9-14 PSU LED indicators

Table 9-6 describes the meaning of the different LEDs.

Table 9-6 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 allowed	N/A

9.3 Configuration backup procedure

If there is a serious failure that requires that the system configuration must be restored, the configuration backup file must be used. It contains configuration data such as arrays, pools, and volumes (but no customer applications data). The backup file is updated by the cluster daily.

Even so, it is important to save the file after any changes are made to your system configuration. This save requires command-line interface (CLI) connection to manually start a new backup.

Regularly saving the configuration backup file of the IBM Storwize V3500 is important but 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 by using a script and save it daily).

9.3.1 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 9-1.

Example 9-1 Example of backup CLI command

```
svcconfig backup
```

The progress of the command is detected by dots, as shown in Figure 9-15.

```
IBM_2072:IBM-Storwize-V3700:superuser>svcconfig backup
.....
.....
.....
.....
.....
CMMVC6155I SVCCONFIG processing completed successfully
IBM_2072:IBM-Storwize-V3700:superuser>
```

Figure 9-15 Configuration backup CLI command

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 of the configuration node and are listed on the support view.

Table 9-7 describes the three files that are created by the backup process.

Table 9-7 Files 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.

9.3.2 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 click **Support**, as shown in Figure 9-16.

Figure 9-16 Configuration backup open support view

2. Select the configuration node on the support view, as shown in Figure 9-17.

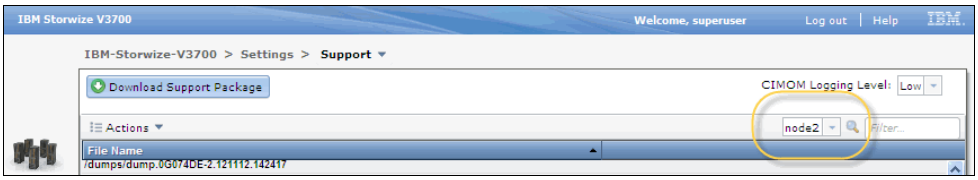

Figure 9-17 Configuration backup select configuration node

3. Select **Show full log listing** (as shown in Figure 9-18) to list all available log files that are stored on the configuration node.

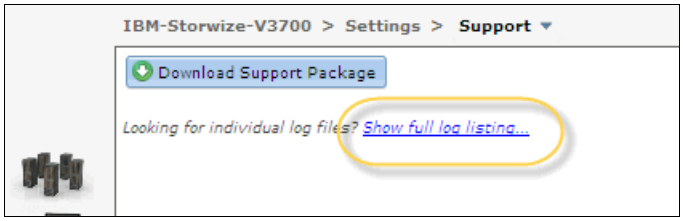

Figure 9-18 Configuration backup show full log list

4. Search for a file named `/dumps/svc.config.backup.xml_*`, as shown in Figure 9-19. Select the file, right-click it, and the select **Download**.

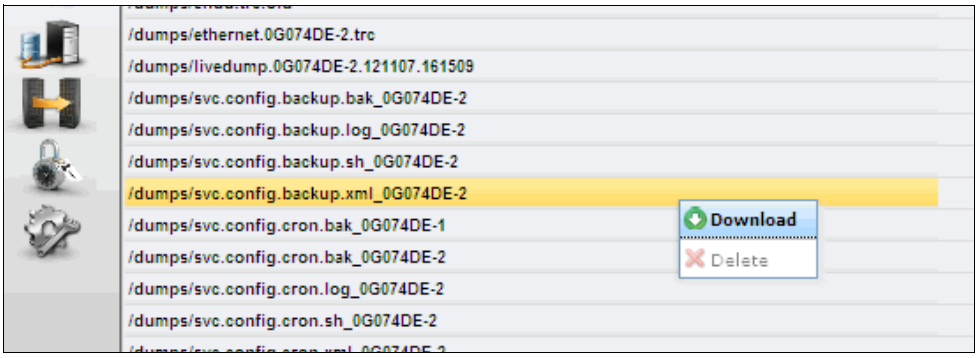

Figure 9-19 Configuration backup start download

5. Save the configuration backup file on your management workstation (as shown in Figure 9-20) where it can be found easily.

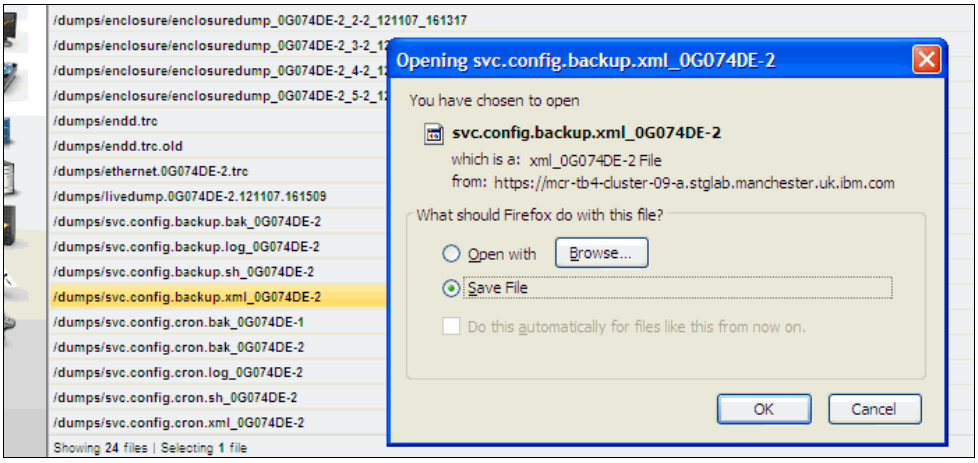

Figure 9-20 Configuration backup save file

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 9-21.

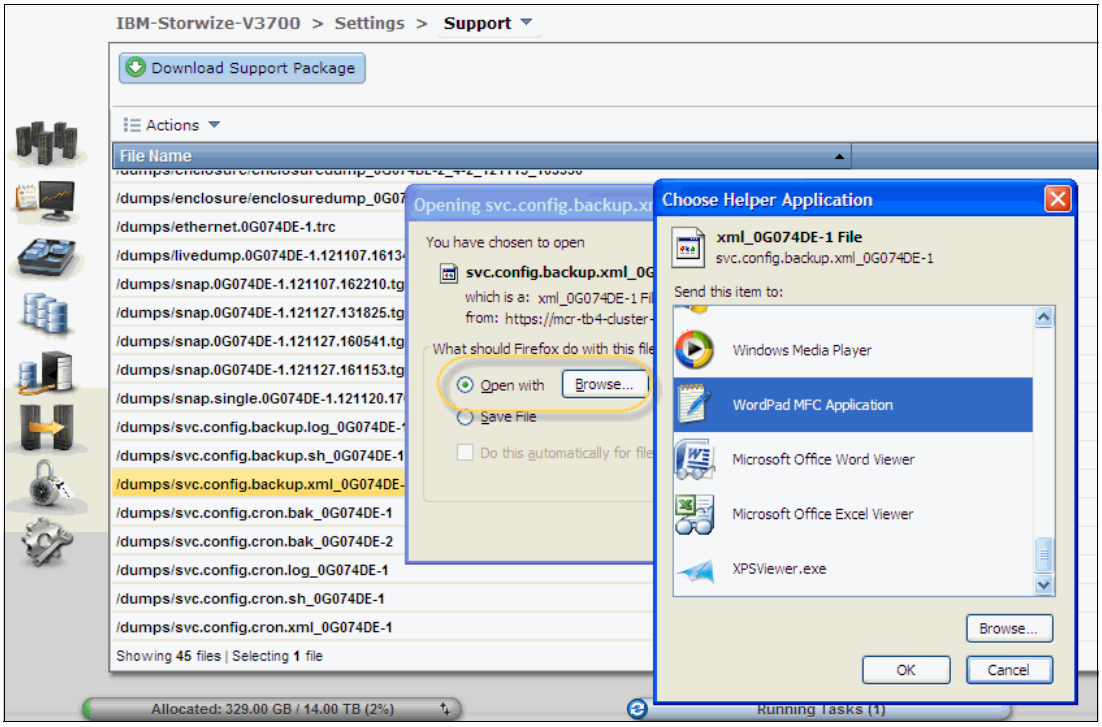

Figure 9-21 Open backup .xml file with WordPad

When the `/dumps/svc.config.backup.xml_xx` file is opened by using and editor, search for the string `timestamp=`. This string is found near of the top of the file. Figure 9-22 shows the file opened with WordPad and the time stamp information in it.

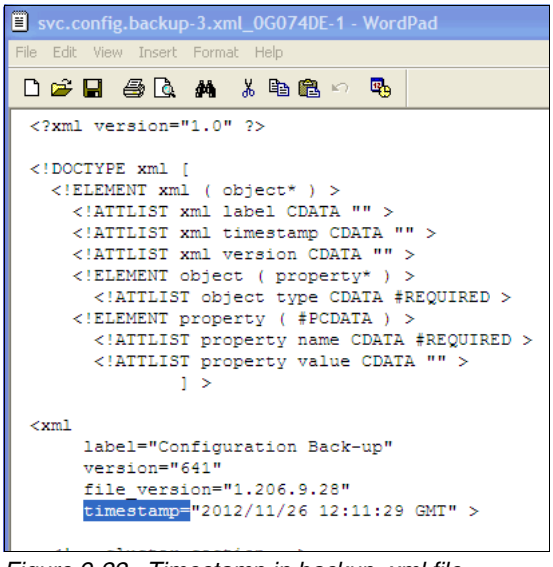

Figure 9-22 Timestamp in backup .xml file

9.4 Software upgrade

The system upgrade process involves upgrading your entire Storwize V3500 environment.

Allow up to one week to plan your tasks, review your preparatory upgrade tasks, and complete the upgrade of the Storwize V3500 environment. The upgrade procedures can be divided into the general processes that are described in Table 9-8.

Table 9-8 Software upgrade tasks

Step	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 also 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 Storwize V3500 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 your Storwize V3500.
5	Upgrade other devices in the Storwize V3500 environment. Examples might include upgrading hosts and switches to the correct levels.

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 over 1 hour for an upgrade.

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 fix any outstanding errors before you attempt to concurrently upgrade the code.

9.4.1 Upgrading software automatically

During the automatic upgrade process, each node in a system is upgraded 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 both the nodes in the system are successfully restarted with the new code level, the new level is automatically committed.

The node that is being upgraded is temporarily unavailable and all I/O operations to that node fail. 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 Storwize V3500 system.

The upgrade normally can be performed concurrently with normal user I/O operations. However, there is a possibility that performance might be impacted.

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 that are going away (fail over) and the error count increasing during the upgrade. When the paths to the nodes are back, the nodes fall back to become a fully redundant system. After the 30-minute delay, the paths to the other node go down.

9.4.2 GUI upgrade process

The automatic upgrade process begins in GUI with starting the upgrade wizard as shown in Figure 9-23. Browse to **Settings** → **General** → **Upgrade Machine Code** → **Launch Upgrade Wizard**.

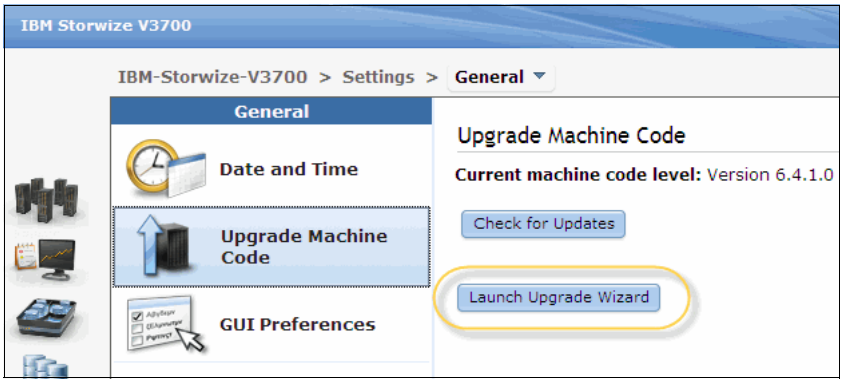

Figure 9-23 Start upgrade wizard

As a first step, the upgrade test utility must be downloaded from the internet. The correct link is provided within the panel. If the tool is already downloaded and stored on the management station, it can be uploaded, as shown in Figure 9-24.

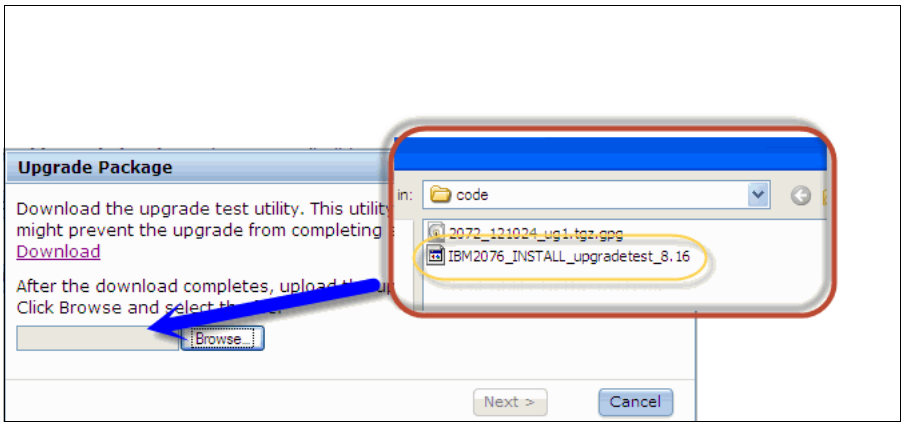

Figure 9-24 Download upgrade test utility

A confirmation panel opens, as shown in Figure 9-25.

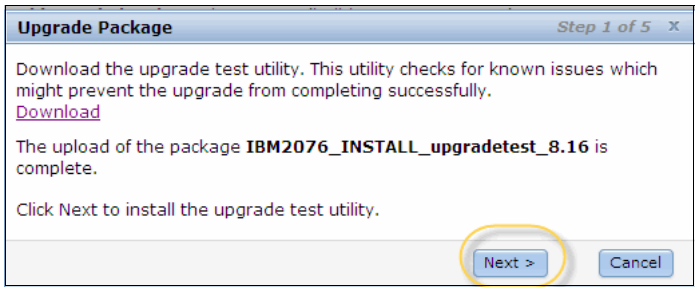

Figure 9-25 Upload test utility complete

The version to which the system should be upgraded must be entered in step 2, as shown in Figure 9-26. By default, the latest code level is displayed.

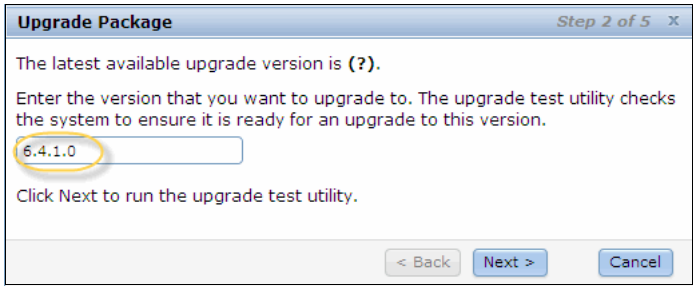

Figure 9-26 Enter version to be checked by tool

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 9-27 shows the panel that indicates that the background test task is running.

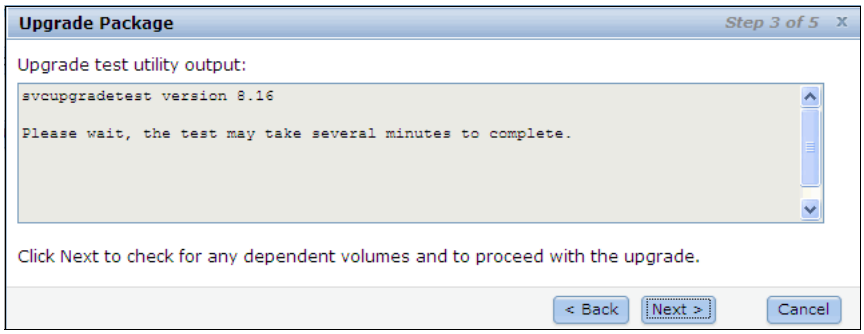

Figure 9-27 Wait utility to complete

Next, the code must be downloaded. If the code is already downloaded to the management station, it can be directly uploaded, as shown in Figure 9-28. Verify that the correct code file is used.

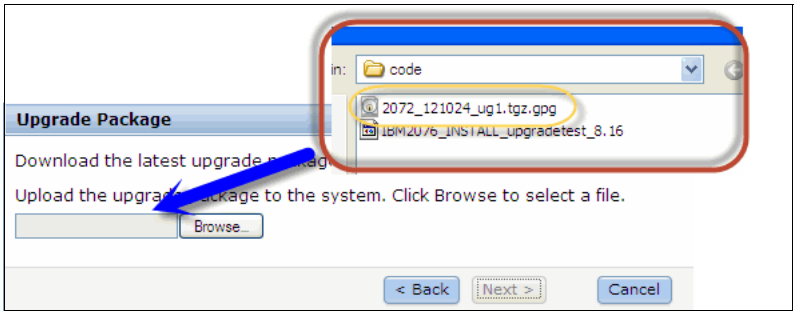

Figure 9-28 Download code

A confirmation panel opens, as shown in Figure 9-29.

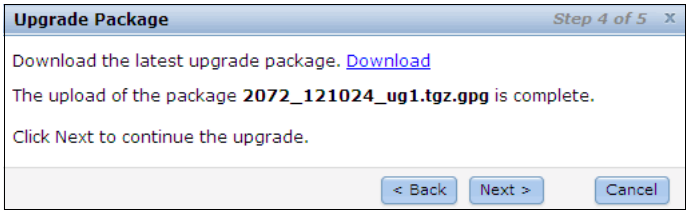

Figure 9-29 Code upload completed

The automated code upgrade can be started by selecting the **Automatic upgrade** option on the decision panel (as show in Figure 9-30), which is the default. If the upgrade is done manually for any reason, the selection must be made accordingly (an automatic upgrade is recommended).

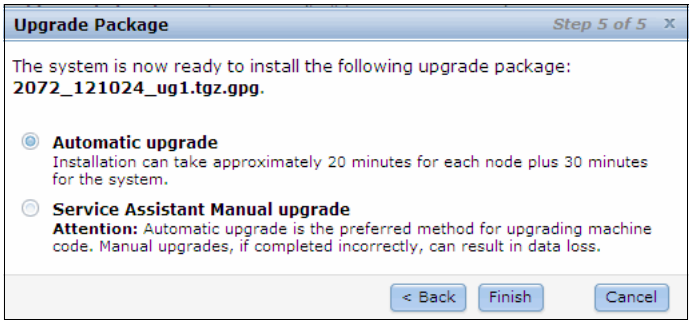

Figure 9-30 Upload mode decision

If you choose to select the **Service Assistant Manual** upgrade option, see 9.4.3, “Upgrading software manually” on page 460.

Selecting **Finish** starts the upgrade process on the nodes. Messages inform you that first one node, then the other, is upgraded. When both nodes are rebooted, the upgrade process is complete. It can take more than an hour to finish this process.

9.4.3 Upgrading software manually

Important: It is highly recommended to upgrade the Storwize V3500 automatically by using the upgrade wizard. If a manual upgrade is used, make sure that you do not skip any step.

The follow steps are for a manual upgrade and are shown on the Service Assistant Upgrade Manually panel.

Complete the following steps to manually upgrade the software:

1. In the management GUI, click **Settings** → **General** → **Upgrade Machine Code** and run the Launch Upgrade wizard. In step 5 of the wizard, select **Service Assistant Manual upgrade**, as shown in Figure 9-31.

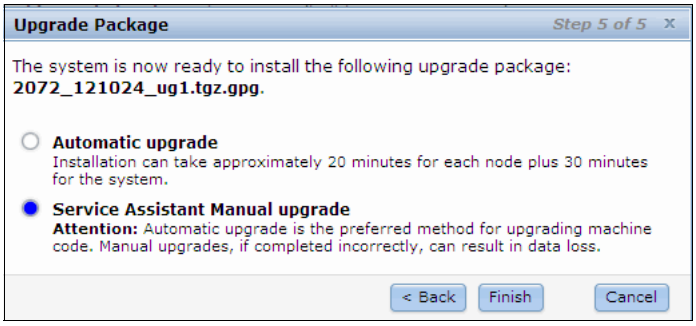

Figure 9-31 Select manual upgrade mode

2. Both nodes are set to Waiting for Upgrade status in the Upgrade Machine Code panel, as shown in Figure 9-32.

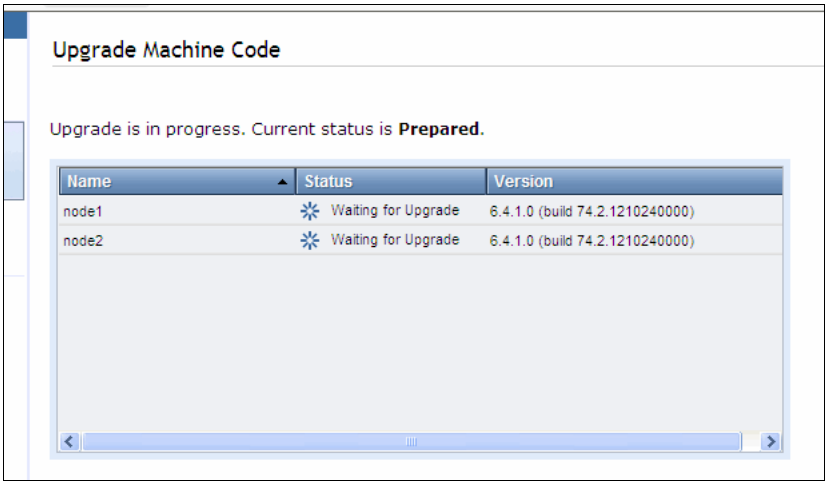

Figure 9-32 Node status to Waiting for Upgrade

3. In the management GUI, select **System Details** and then select the canister that contains the node you want to upgrade next. As shown in Figure 9-33, select **Remove Node** on the action menu.

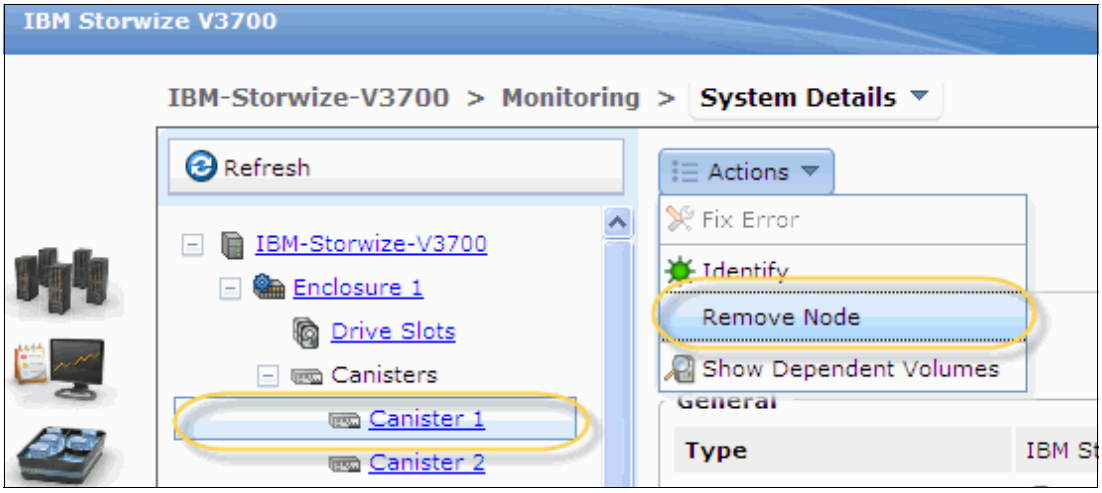

Figure 9-33 Remove the non-config node from cluster

Important: Select the non-config node first.

4. A warning message is displayed, as shown in Figure 9-34.

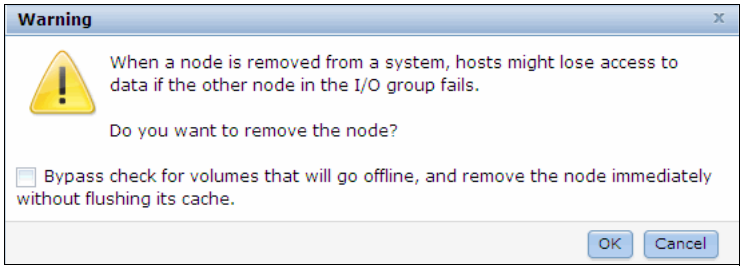

Figure 9-34 Remove node warning message

5. The non-config node is removed from the Upgrade Machine Code panel, as shown in Figure 9-35.

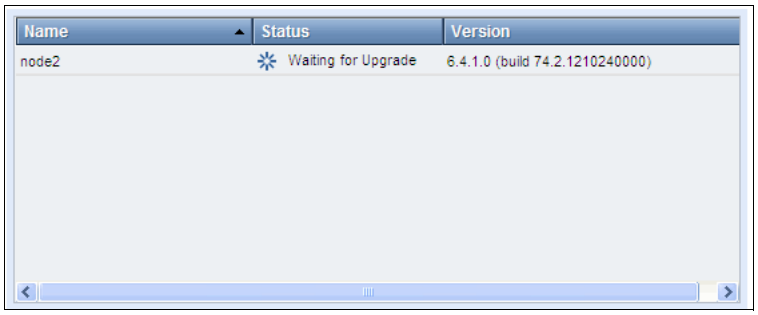

Figure 9-35 Non-config node is removed

6. In System Details, the node is shown as not configured, as shown in Figure 9-36.

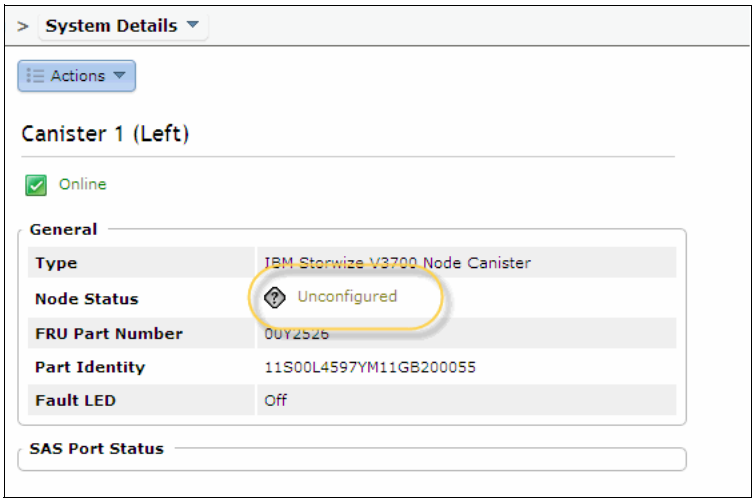

Figure 9-36 Node get status unconfigured

In the Service Assistant panel, the node that is ready for upgrade must be selected. It is the node that is in service mode and has no available cluster information, as shown in Figure 9-37.

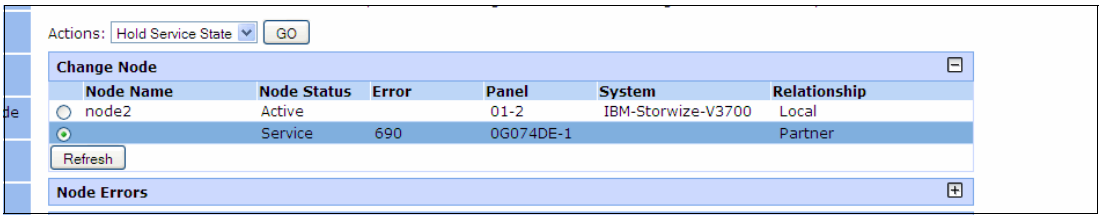

Figure 9-37 Select node in service mode for upgrade

7. On Service Assistant panel, select **Upgrade Manually** and then select the machine code version to upgrade this node, as shown in Figure 9-38.

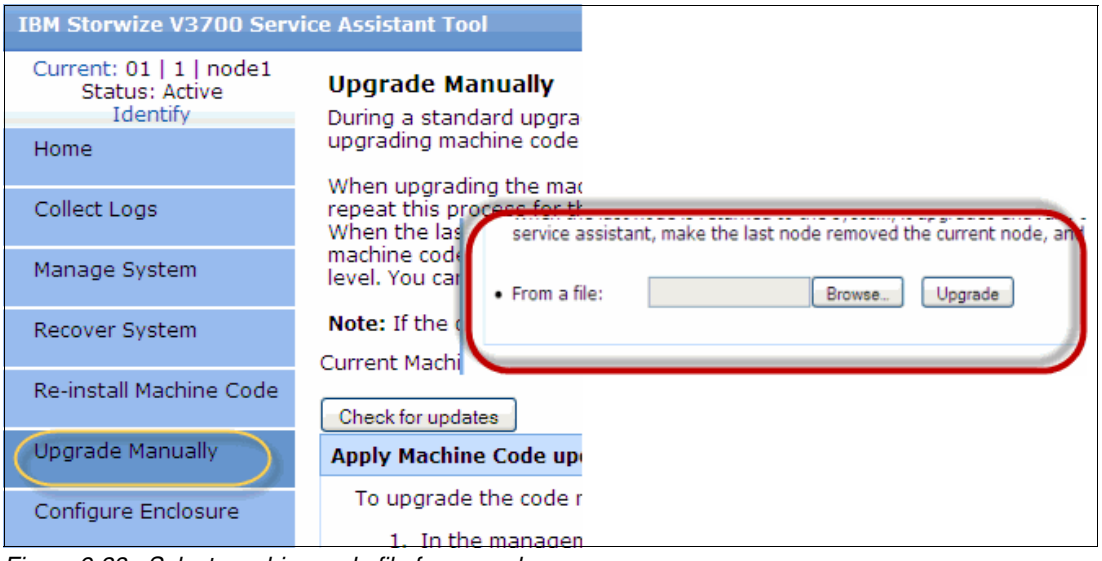

Figure 9-38 Select machine code file for upgrade

8. Click **Upgrade** to start the upgrade process on the first node.

   The node automatically is added again to the system.

   In some circumstances, the node cannot automatically be added again to the system. The upgrade and adding process can take up to 30 minutes. When finished, the first node is completed and the upgrade information panel in the GUI is updated, as shown in Figure 9-39.

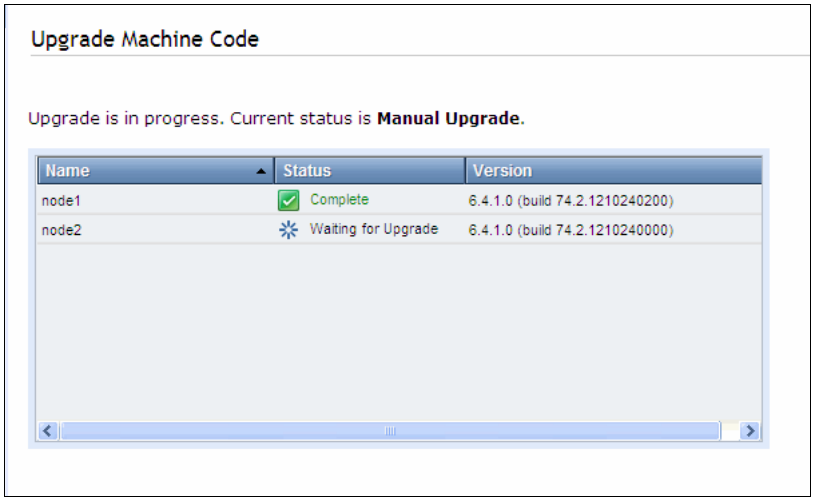

Figure 9-39 Non-config node completed upgrade

9. Repeat steps 2 - 4 for the remaining config node.

   Because the config node is removed from cluster, a warning appears, as shown in Figure 9-40.

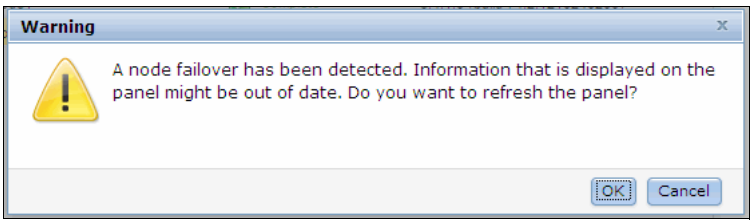

Figure 9-40 Config node failover warning

Note: The config node remains in the Service State when it is added again to the cluster. Exit the Service State manually.

10. To exit the service state, browse to the home panel of the Service Assistant and open the Action menu. Select **Exit Service State**, as shown in Figure 9-41.

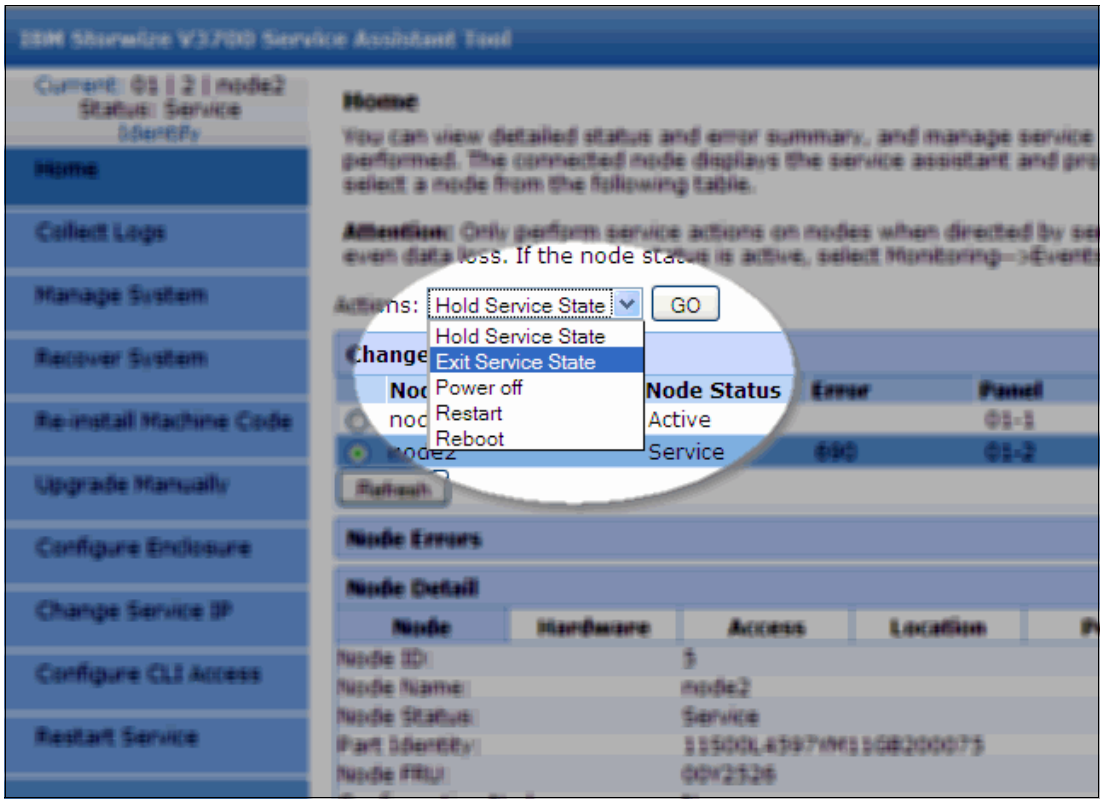

Figure 9-41 Exit service state

Both nodes are now back in the cluster (as shown in Figure 9-42) and the system is running on the new code level.

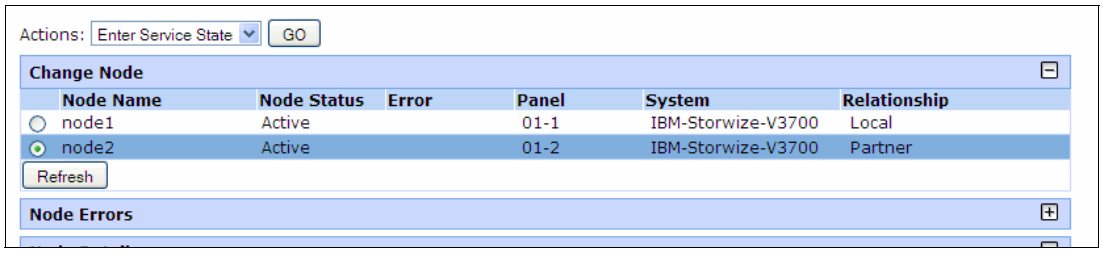

Figure 9-42 Cluster is active again and running new code level

9.5 Event log

Whenever a significant change in status of IBM Storwize V3500 is detected, an event is logged in the event log.

All events are classified as either 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 as soon as 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 9-43.

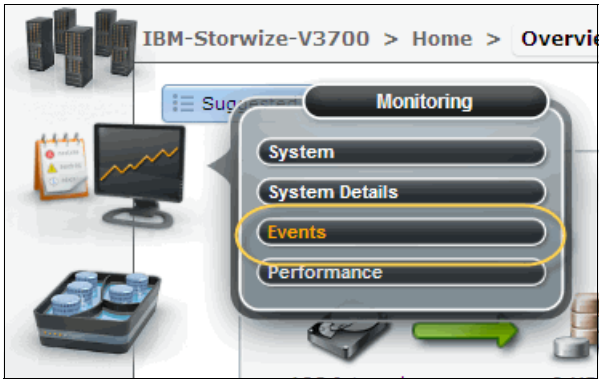

Figure 9-43 Open event log panel

When opened, the event log looks as shown in Figure 9-44.

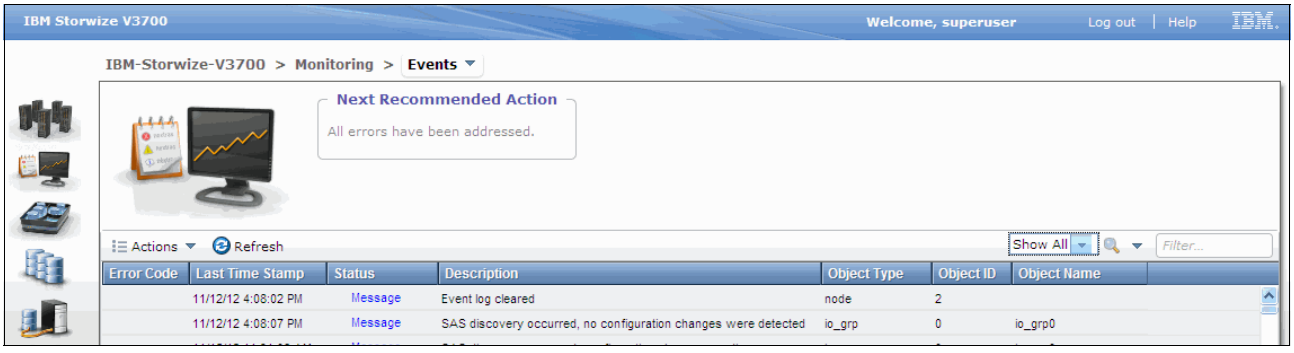

Figure 9-44 The event log view

9.5.1 Managing the event log

The event log has a limited size. After it is full, newer entries replace entries that are not required.

To avoid having 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 any column header opens the option menu to select columns to be shown or hidden. Figure 9-45 shows the columns that can be displayed in the error log view.

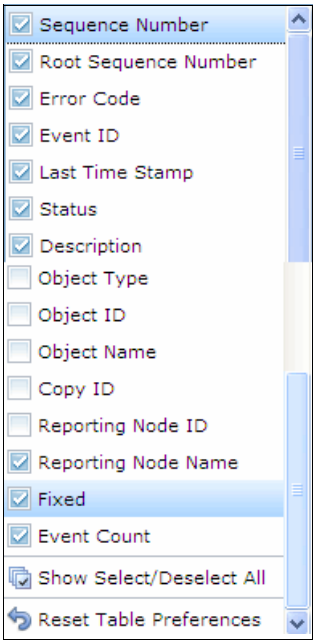

Figure 9-45 Event log columns

The following available fields that are recommended are described to assist you in diagnosing problems:

- Event ID

  This number precisely identifies 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 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 as 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 9-46 shows the main Event log panel options that should be used to handle system events.

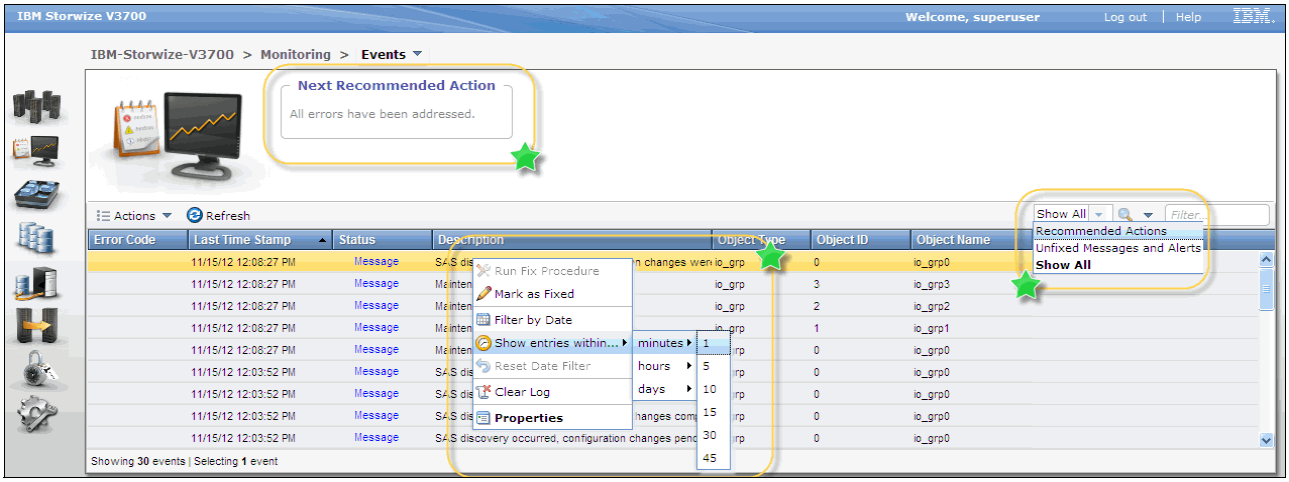

Figure 9-46 Event log panel

Event log filter options

The following log filter options are available:

- Show all

  This option list all available events.

- Unfixed Messages and Alerts

  This option lists unfixed events. It 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.

> **Warning:** Check for this filter option if there is no event is listed. There might be events that are not associated to recommended actions.

Figure 9-47 shows an event log with no items found, which does not necessarily mean that the event log is clear. We check whether the log is clear by using the filter option Show all.

Figure 9-47 No items are found in event log

Actions on single event

Right-clicking a single event gives the following option that might be used for that specific event:

- Mark as Fixed

  It is possible to start the Fix Procedure on this specific event, even 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 limits 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.

  Warning: This action cannot be undone and might prevent the system from being analyzed in case of severe trouble.

- Properties

  This option provides more sense data for the selected event as are shown in the list.

Recommended Action

A fix procedure is a wizard that helps you troubleshoot and correct the cause of an error. Some fix procedures reconfigure the system 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 must always 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 9.5.2, “Alert handling and Recommended Actions” on page 469.

9.5.2 Alert handling and Recommended Actions

All events in an Alert status require attention. Alerts are listed in priority order and should be fixed sequentially by using the available fix procedures.

Example: Power supply failed

For this example, the power cord of the right-side power supply was removed.

The example gives an idea as to how faults are represented in the error log. Information about the fault can be gathered and the DMP can be used to fix the error.

1. Detect an alert

   The Health Status indicator, which is permanently present on most of the GUI panel (see Chapter 3, “Graphical user interface overview” on page 53), is alerting yellow. Click the indicator to receive the specific information, as shown in Figure 9-48.

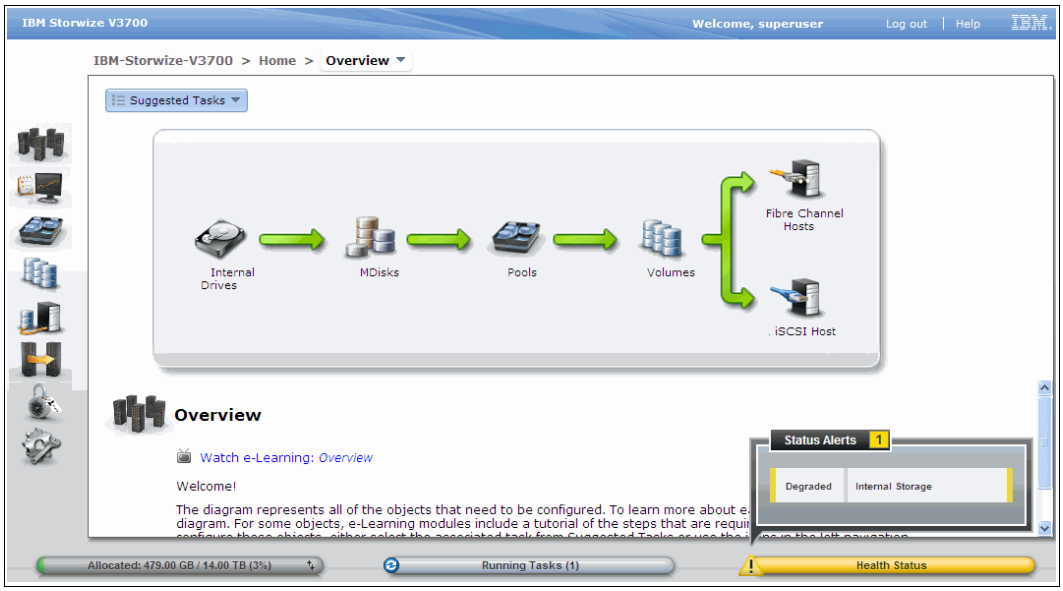

Figure 9-48 Health check shows degraded system status

   Review the event log for more information.

2. Find alert in event log

   The default filter in the error log view (as show in Figure 9-49) is Recommended actions (1). This option lists the alert event only (2) and shows the recommended action to be run first (3).

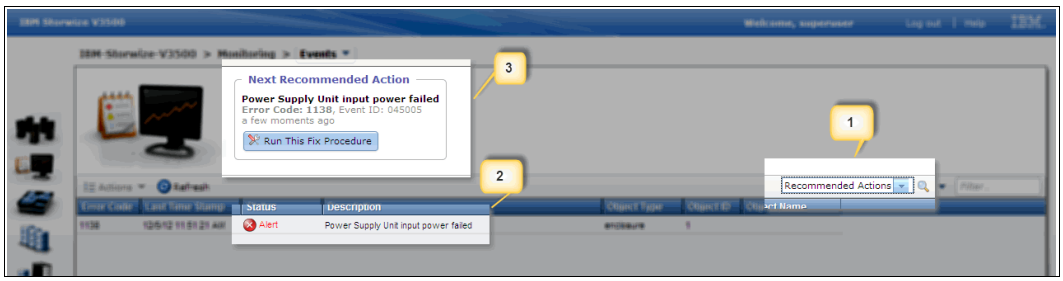

Figure 9-49 Alert in event log

3. Run recommended action

   It is highly recommended to fix alerts under the guidance of the recommended action, or DMP. There are tasks that are running 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 view.

   The steps and the DMP panels are specific to the error that must be fixed. The following figures represent the DMP for the power supply failure event example.

   Figure 9-50 shows the first DMP panel.

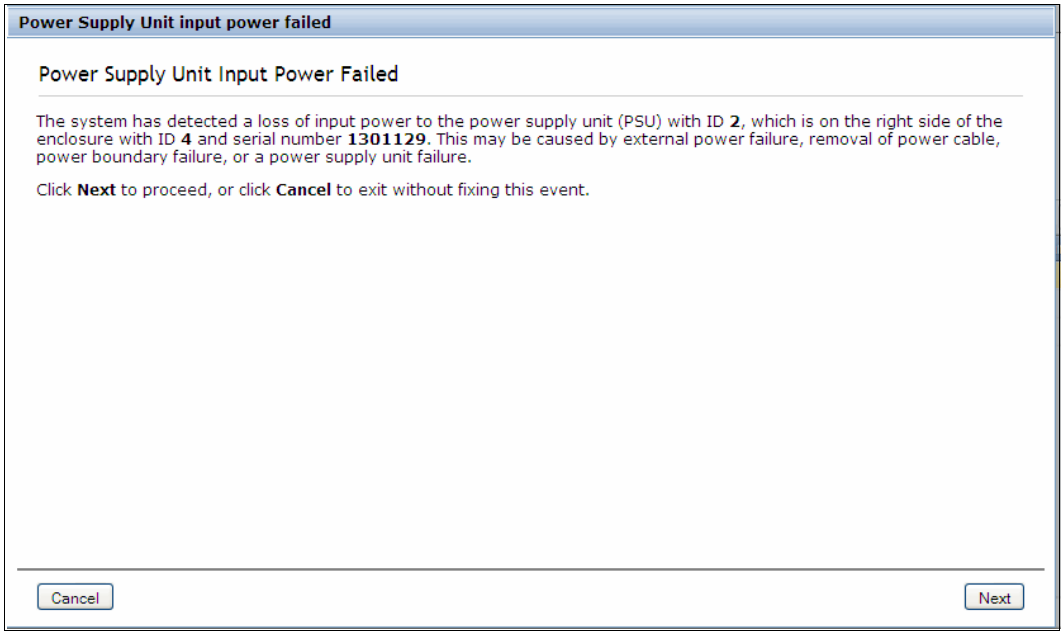

Figure 9-50 Power Supply Unit input power failed

Figure 9-51 shows second DMP panel.

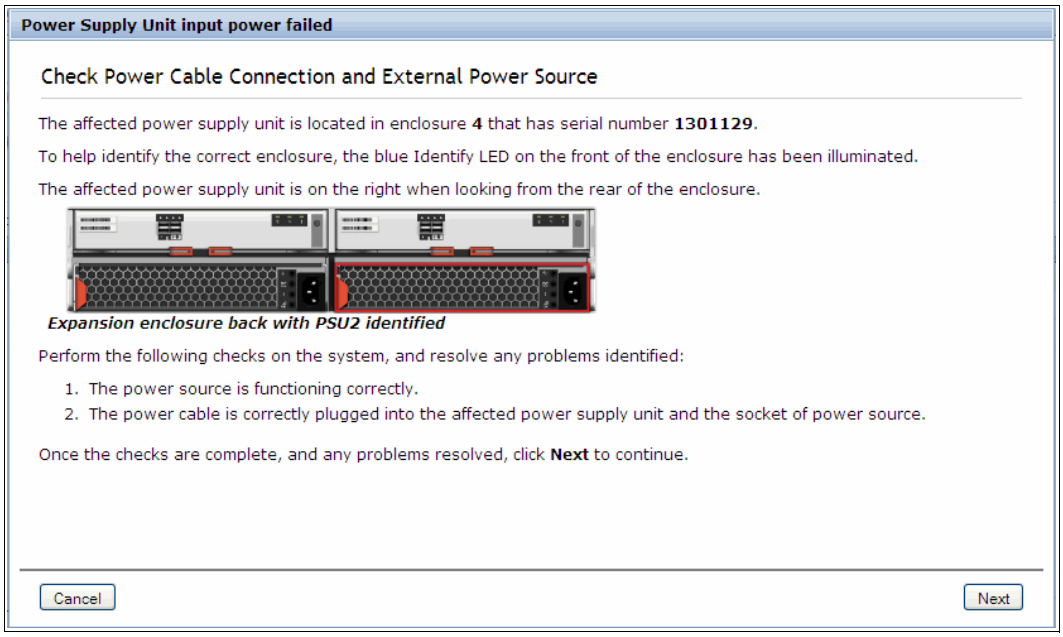

Figure 9-51 Second DMP panel for power failure

Figure 9-52 shows the required action.

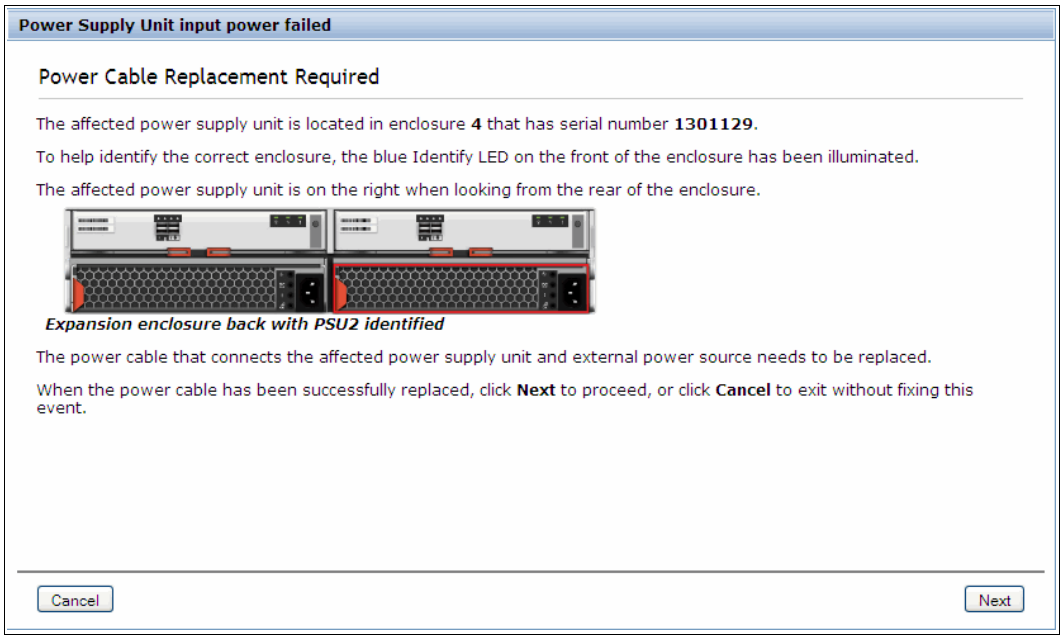

Figure 9-52 Third DMP panel for power failure

Figure 9-53 shows the fourth DMP panel.

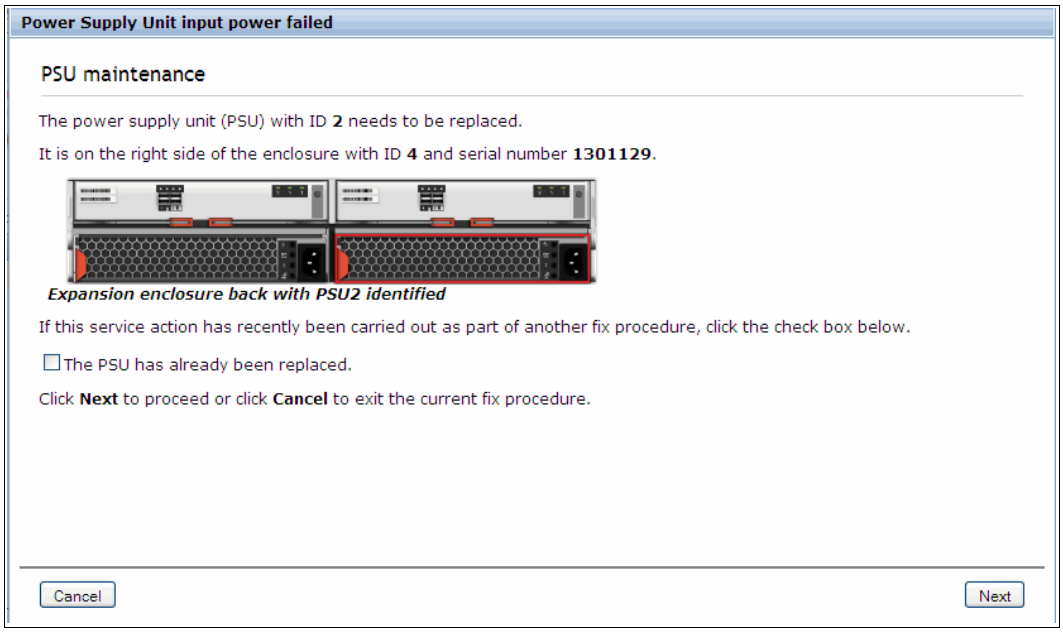

Figure 9-53 Fourth DMP panel for power failure

The next panel that is shown in Figure 9-54 appears if the failure is still present. Therefore, the power supply must be replaced.

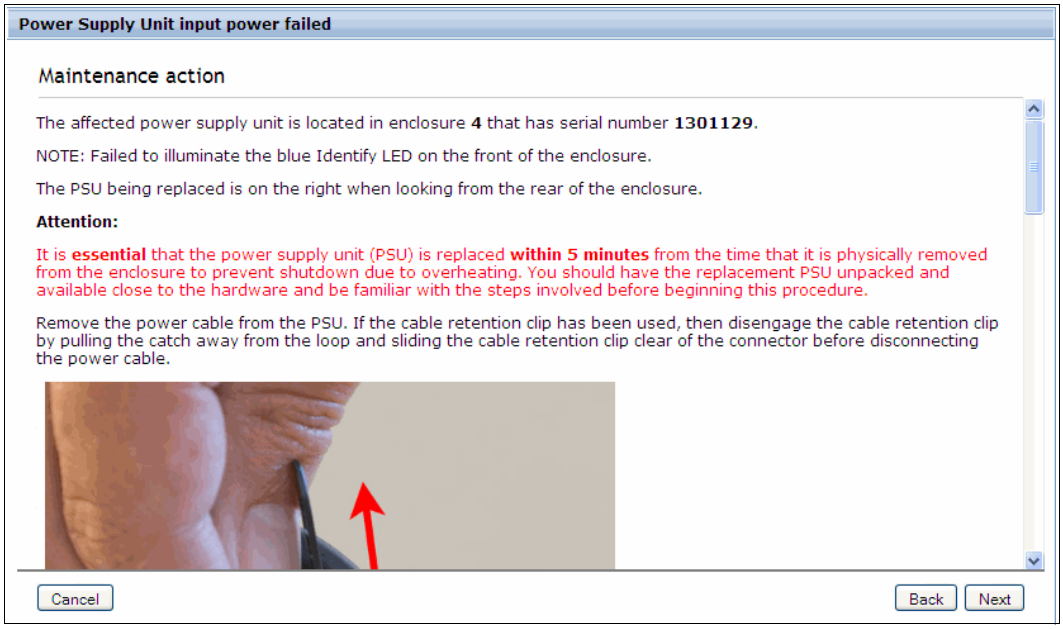

Figure 9-54 Fifth DMP panel for power failure

Figure 9-55 shows the DMP panel that indicates that the failure was fixed successfully.

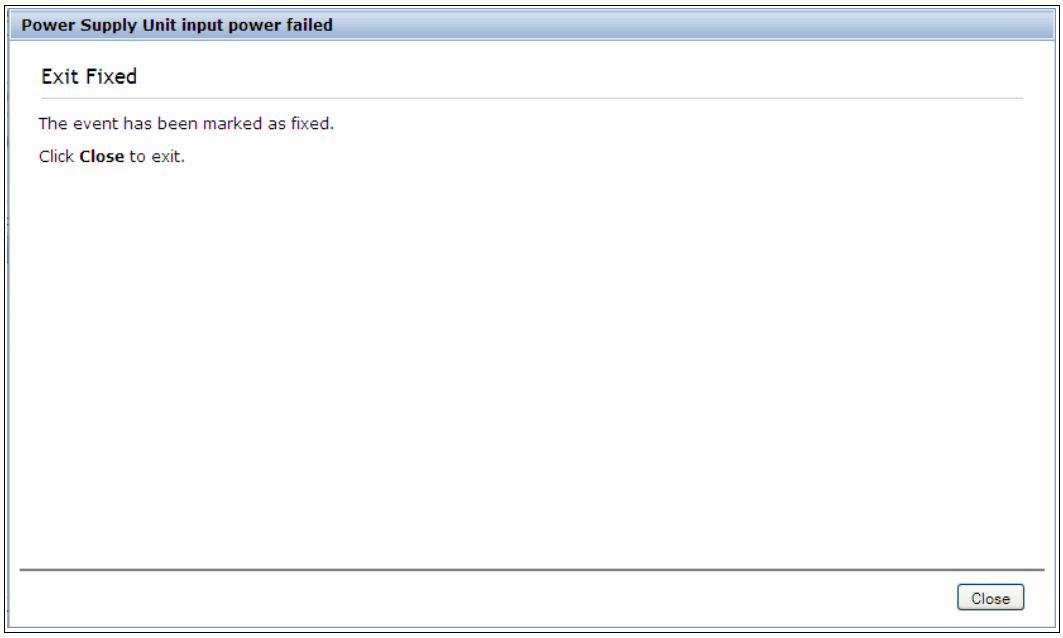

Figure 9-55 Final DMP panel when failure was fixed successful

When all steps of the DMP are completed successfully, the problem should be fixed. Figure 9-56 shows that the red color of the alert event status changed to green, the system health status is green, and there are no errors that require a recommended action. Therefore, the problem is fixed.

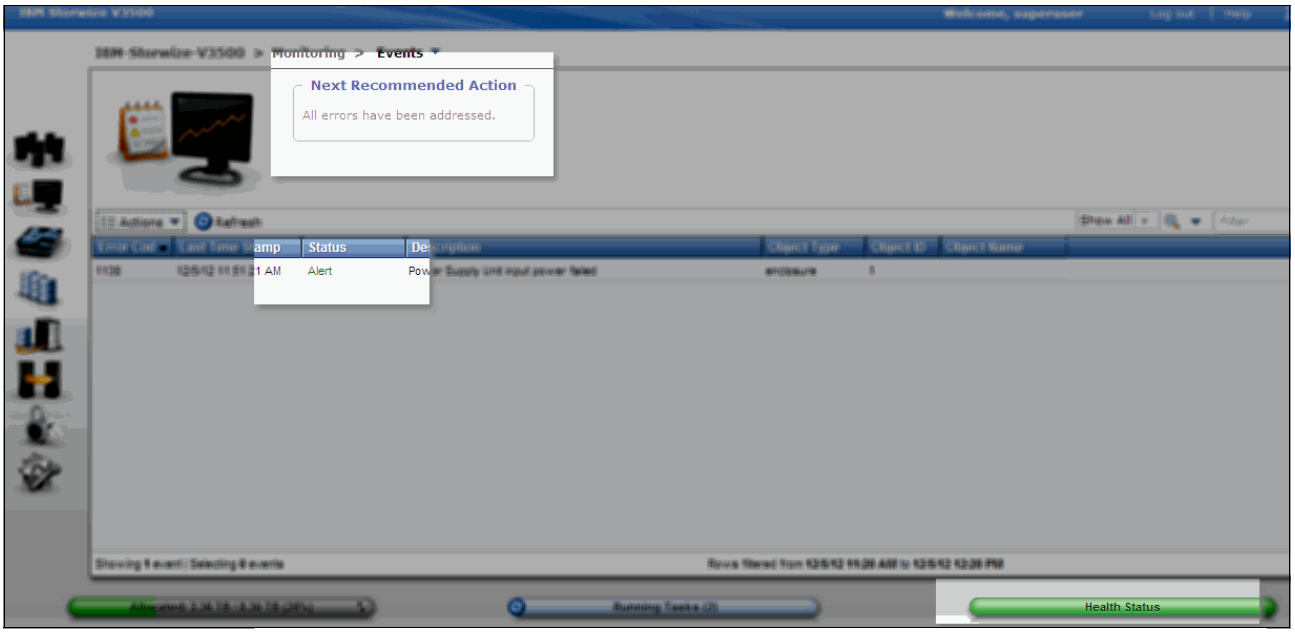

Figure 9-56 Successfully completed DMP

Handling multiple alerts

If there are multiple logged alerts, the IBM Storwize V3500 recommends an action to fix the problem.

Figure 9-57 shows the event log displaying multiple alert events.

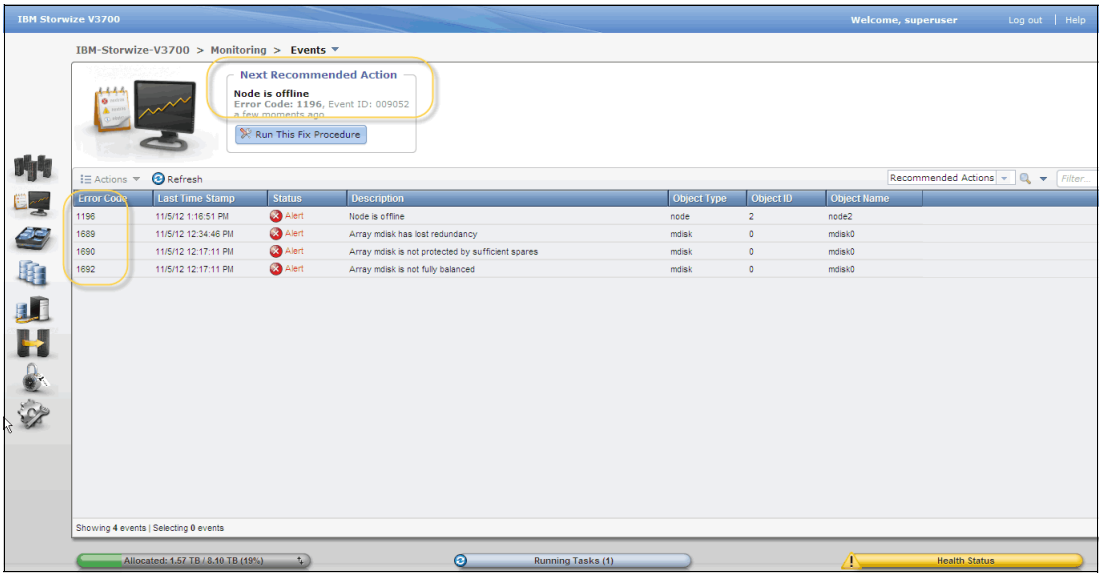

Figure 9-57 Multiple alert events are displayed in the event log

The Next Recommended Action function orders alerts by severity and displays and handles the events with the highest severity first. If multiple events have the same severity, they are ordered by date, with the oldest event 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 often are fixed with the fixture of the most severe fault.

9.6 Collecting support information

If you have a problem and call the IBM Support Center, you might be asked to provide support data.

9.6.1 Support Information via GUI

Click **Settings** and then the **Support** tab (as shown in Figure 9-58) to begin the process of collecting support data.

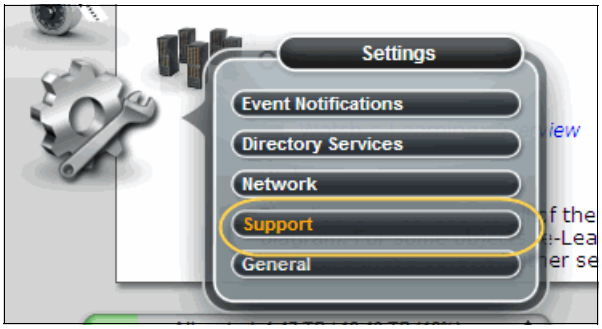

Figure 9-58 Support files via GUI

Click **Download Support Package**, as shown in Figure 9-59.

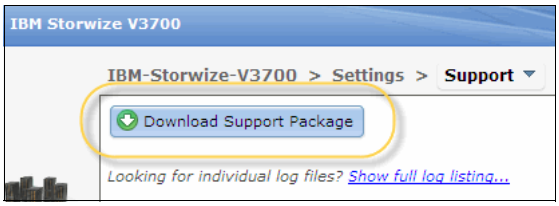

Figure 9-59 Download support package

The panel that is shown in Figure 9-60 opens and you can select one of four different versions of svc_snap support package.

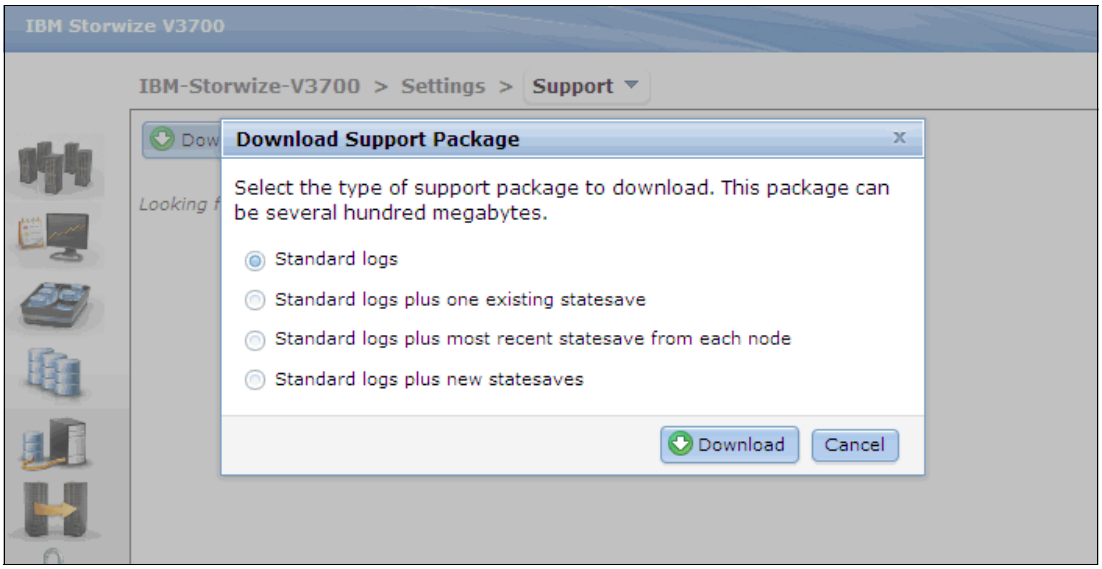

Figure 9-60 Support Package selection

The version that you download depends on the event that is investigated. For example, if you notice that a node was restarted in the event log, capture the snap with the latest statesaves.

The following components are included in the support package:

- Standard logs

  Contains the most recent logs that were collected for 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 for the system and the most recent statesave from any of the nodes in the system. Statesaves also are known as dumps or live dumps.

- Standard logs plus most recent statesave from each node

  This option is mostly used by the Support team for problem analysis.

  They contain the standard logs for the 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 analysis. It generates a new statesave (live dump) for all the nodes and packages them with the most recent logs.

Save the resulting snap file in a directory for later usage or upload it to IBM Support.

9.6.2 Support information via Service Assistant

The management GUI collects information from all of the components in the system. The Service Assistant collects information from a single node canister. When the information that is collected is packaged in a single file, the file is called a snap.

If the package is collected by using the Service Assistant, ensure that the node the logs are collected from is the current node, as shown in Figure 9-61.

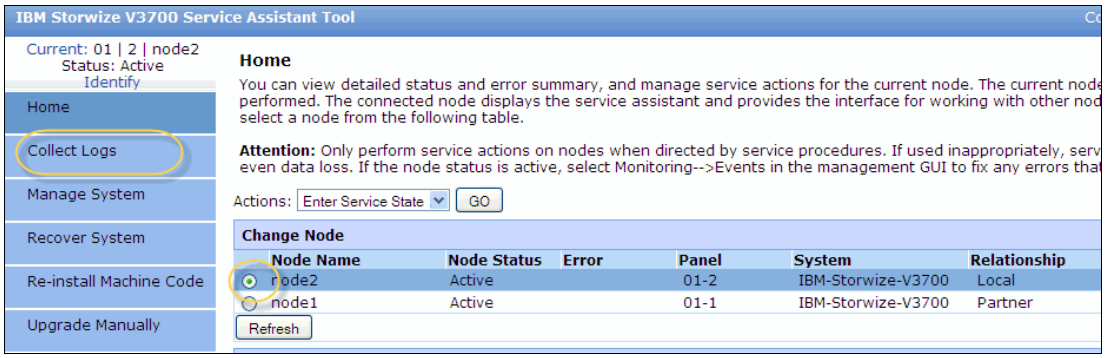

Figure 9-61 Collect logs with Service Assistant

Support information can be downloaded with or without the latest statesave, as shown in Figure 9-62.

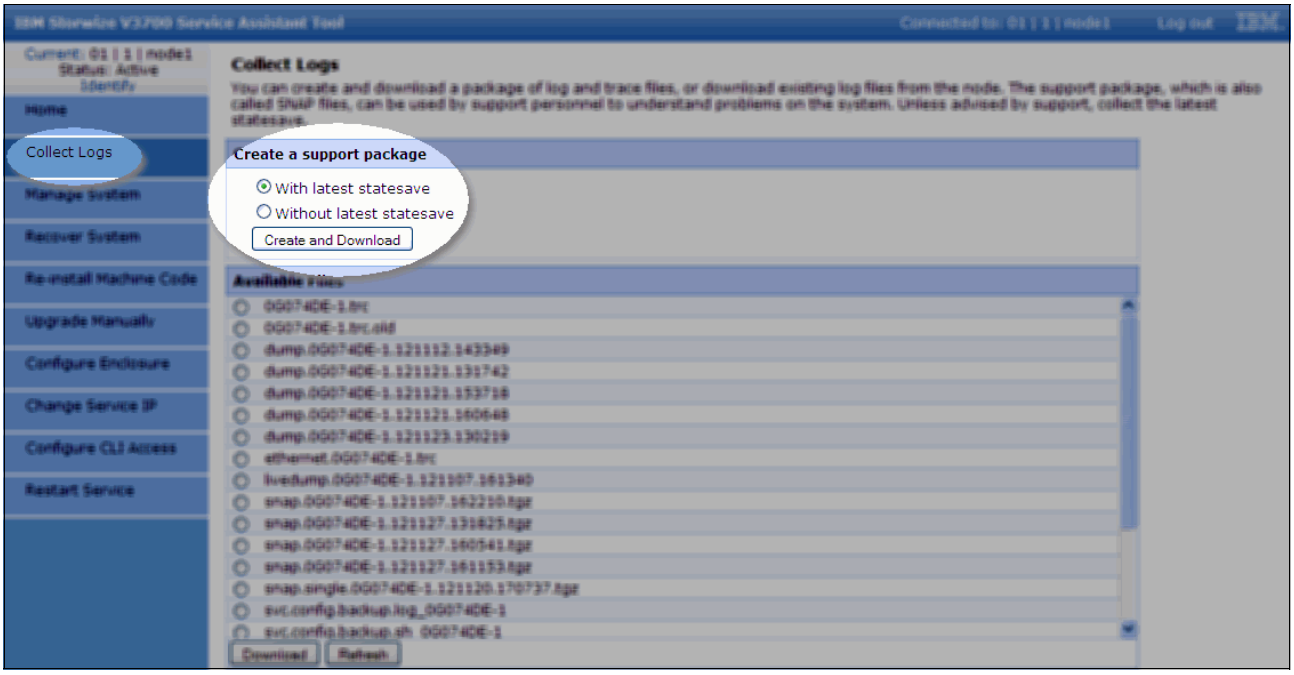

Figure 9-62 Download support file via Service Assistant

9.6.3 Support Information onto USB stick

Whenever GUI, Service Assistant, or remote connection are not available, 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 result. The file that is shown in Figure 9-63 should be stored on the stick in its root directory.

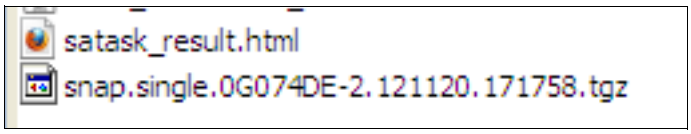

Figure 9-63 Single snap result files on USB stick

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 noted here. Otherwise, any general system information is stored here, as shown in Figure 9-64.

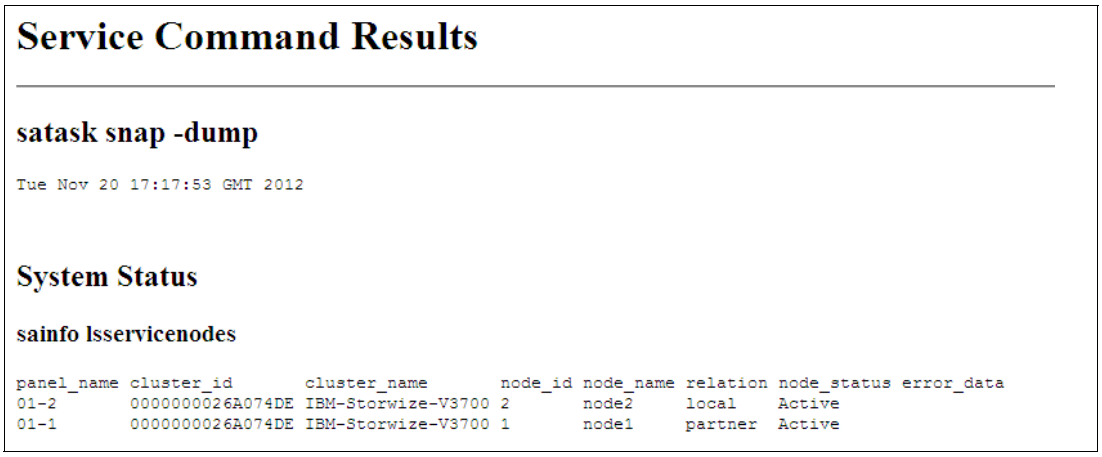

Figure 9-64 satask_result.txt on USB stick (header only)

Snap dump on USB

A complete statesave of the node the USB was attached to is stored in a `.zip` file, as shown in Figure 9-65. The name of the file includes the node name and the time stamp.

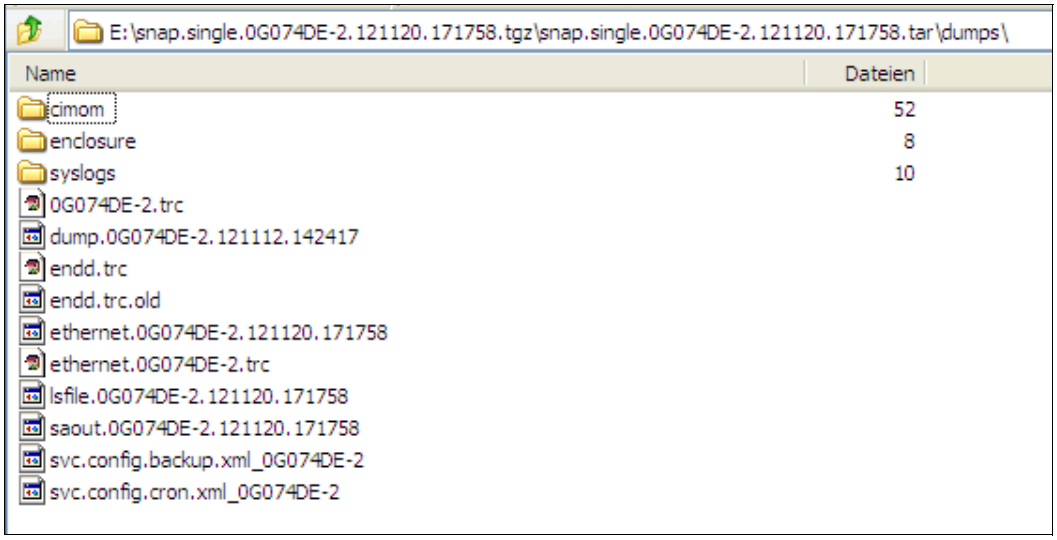

Figure 9-65 Single snap dump on USB stick

9.7 Powering on and shutting down Storwize V3500

In the following sections, we describe the process that is used to power on and shut down the Storwize V3500 system by using the GUI and the CLI.

9.7.1 Shutting down system

In this section, we describe how to shut down a Storwize V3500 system by using the GUI and the CLI.

Important: You should never shut down your IBM Storwize V3500 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 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. Point to the Monitoring function icon (as shown in Figure 9-66) and click **System Details**.

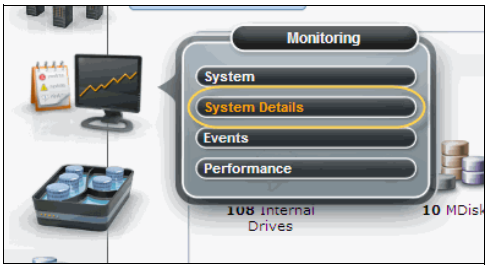

Figure 9-66 Power down via system details

2. Select the root level of the system detail tree. Click **Action** and select **Shut Down System**, as shown in Figure 9-67.

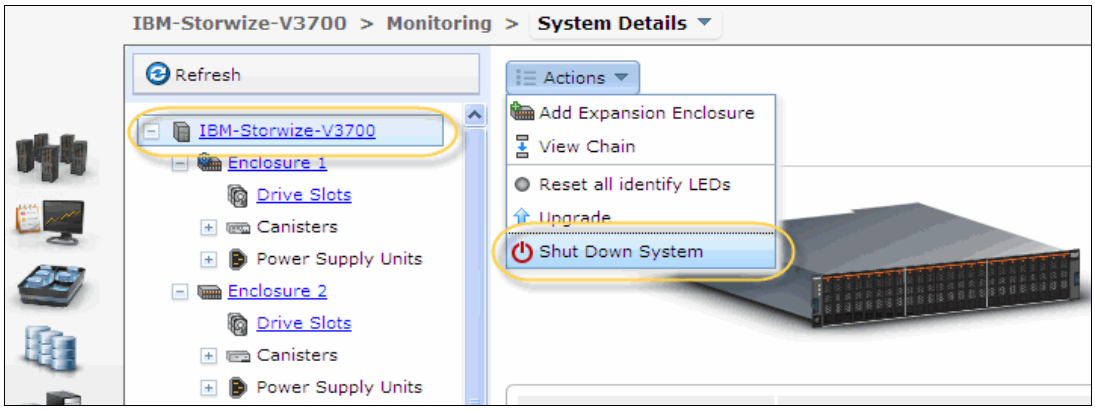

Figure 9-67 Power Down System option

The following process can be used as an alternative to step 1 and step 2 is shown in Figure 9-68:

a. Go to Monitoring navigator and open the system view (1) and (2).
b. Underneath the system display, the system name is shown. Click the name (3).
c. An information panel opens. Click the **Manage** tab (4).
d. Click **Shut down System** (5).

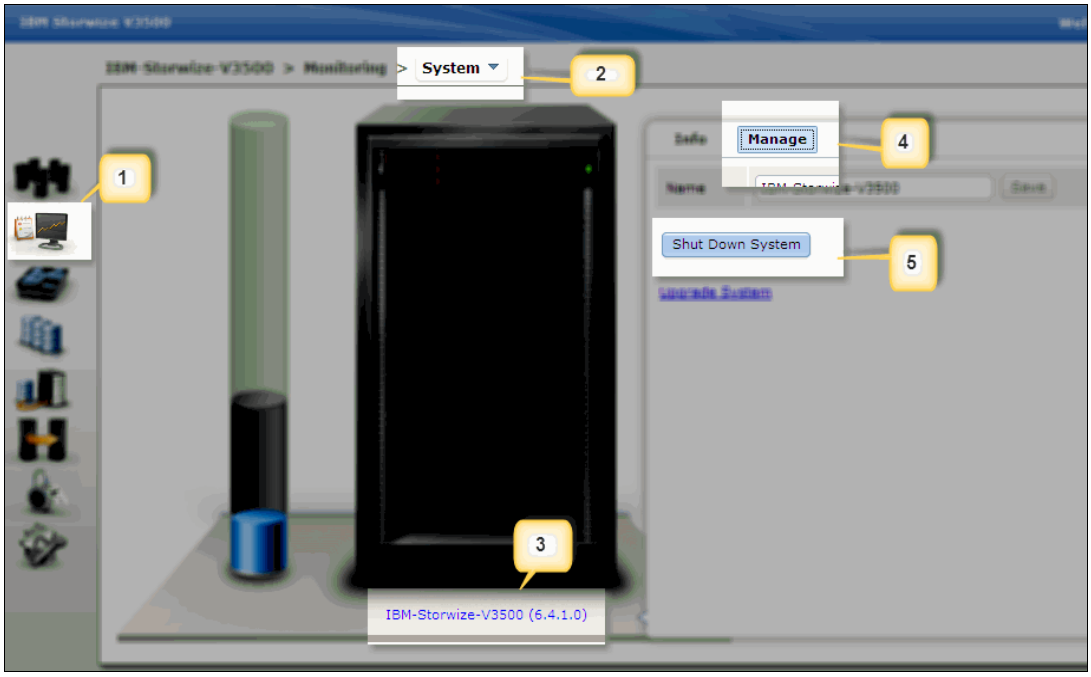

Figure 9-68 Shut down system via Monitoring system GUI

3. The Confirm IBM Storwize V3500 Shutdown IBM Storwize V3500 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.

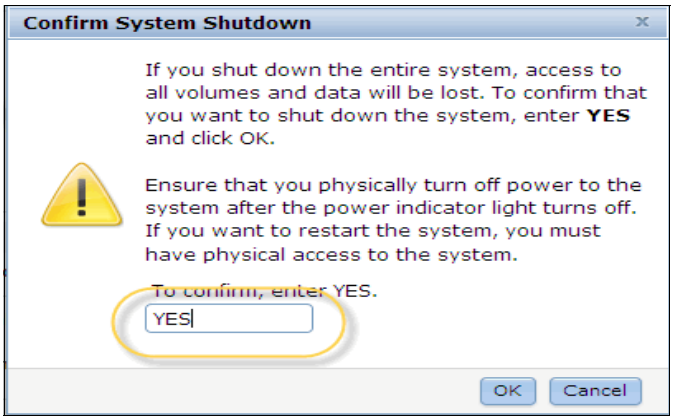

Figure 9-69 Shutdown confirmation

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 V3500, it does not automatically restart. You must manually restart the IBM Storwize V3500.

Shutting down by using the CLI

You also can shut down an IBM Storwize V3500 by using the CLI and 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 Storwize V3500 nodes. This shutdown causes the hardened data to be dumped to the internal hard disk drive.

Run the **stopsystem** command to shut down a clustered system, as shown in Example 9-2.

Example 9-2 Shut down command

```
stopsystem

Are you sure that you want to continue with the shut down?

# Type y to shut down the entire clustered system.
```

9.7.2 Powering on

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 flashes 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.

Complete the following steps to power on the system:

1. Ensure any network switches that the system is connect to are powered on.
2. Power on the control enclosure by connecting the power cords to both power supplies in the rear of the enclosure and switching on the power circuits.
3. The system starts. The system started successfully when all node canisters in the control enclosure have their status LED on permanently, 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 topics:

- Command-line interface
- SAN Boot

Command-line interface

The IBM Storwize V3500 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. Instead, the basic configuration of the IBM Storwize V3500 CLI and some example commands are covered. However, the CLI commands are the same as in the SAN Volume Controller. There also are more commands that are available to manage internal storage. If a task completes in the GUI, the CLI command is always displayed in the details, as shown throughout this book.

For more information about the CLI, see the IBM Storwize V3500 Information Center, which can be found at this website:

http://pic.dhe.ibm.com/infocenter/storwize/v3500_ic/index.jsp?topic=%2Fcom.ibm.storwize.v3500.641.doc%2Fsvc_clicommandscontainer_229g0r.html

Implementing the IBM Storwize V7000 V6.3, SG24-7938, also has much information about using the CLI. The commands in that book also apply to the IBM Storwize V3500 system.

Basic setup

In the IBM Storwize V3500 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 V3500 system. As of IBM Storwize V3500 V6.4, a private and a 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 a private key are generated as a pair.
2. A public key is uploaded to the IBM Storwize V3500 system through the GUI.
3. A client SSH tool must be configured to authenticate with the private key.
4. A secure connection can be established between the client and IBM Storwize V3500 system.

Secure Shell is the communication vehicle that is between the management workstation and the IBM Storwize V3500 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 V3500 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 this website:

http://www.chiark.greenend.org.uk

Download the following tools:

- 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:

- 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.

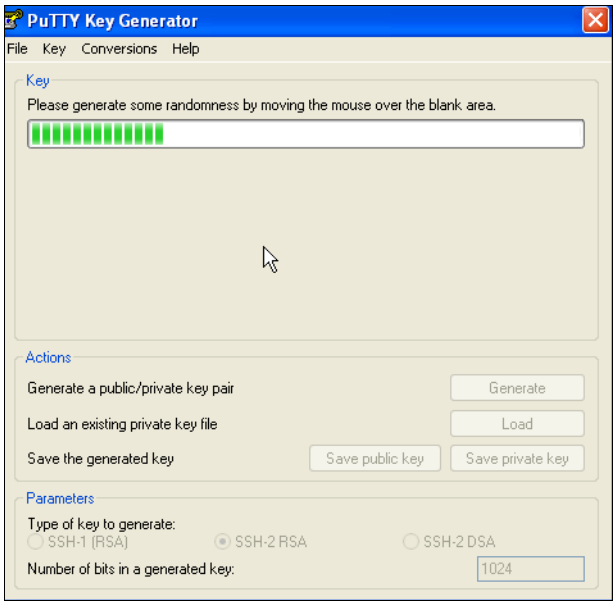

Figure A-2 Generate keys

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. 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.

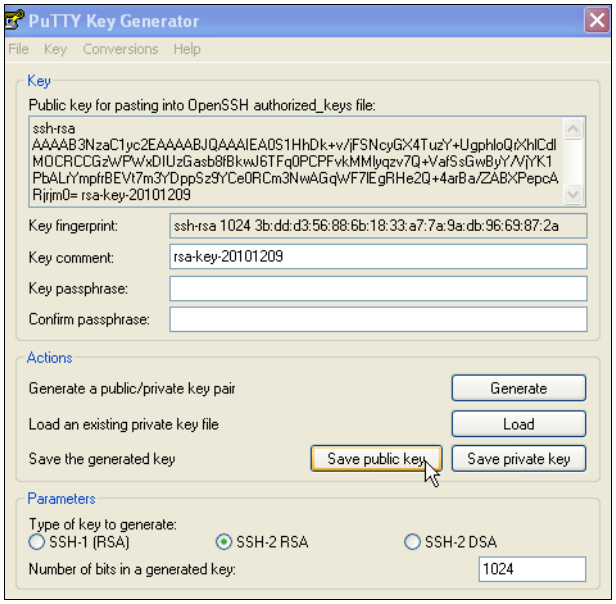

Figure A-3 Save public key

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**.

   Ensure that you record the name and location because the name and location of this SSH public key 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 easily differentiate the SSH public key from the SSH private key.

5. Click **Save private key**, as shown in Figure A-4.

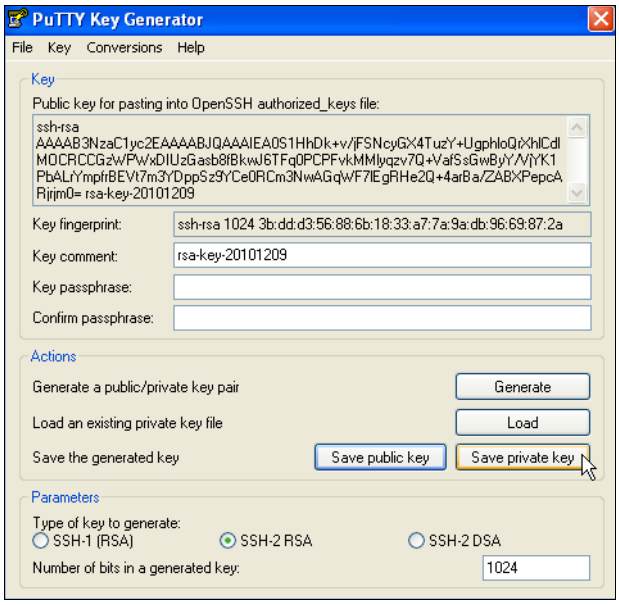

Figure A-4 Save private key

6. You are prompted with a warning message (as shown in Figure A-5). Click **Yes** to save the private key without a passphrase.

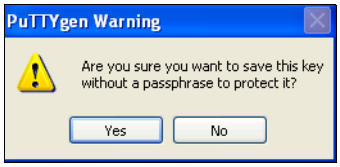

Figure A-5 Confirm the security warning

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 V3500

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.

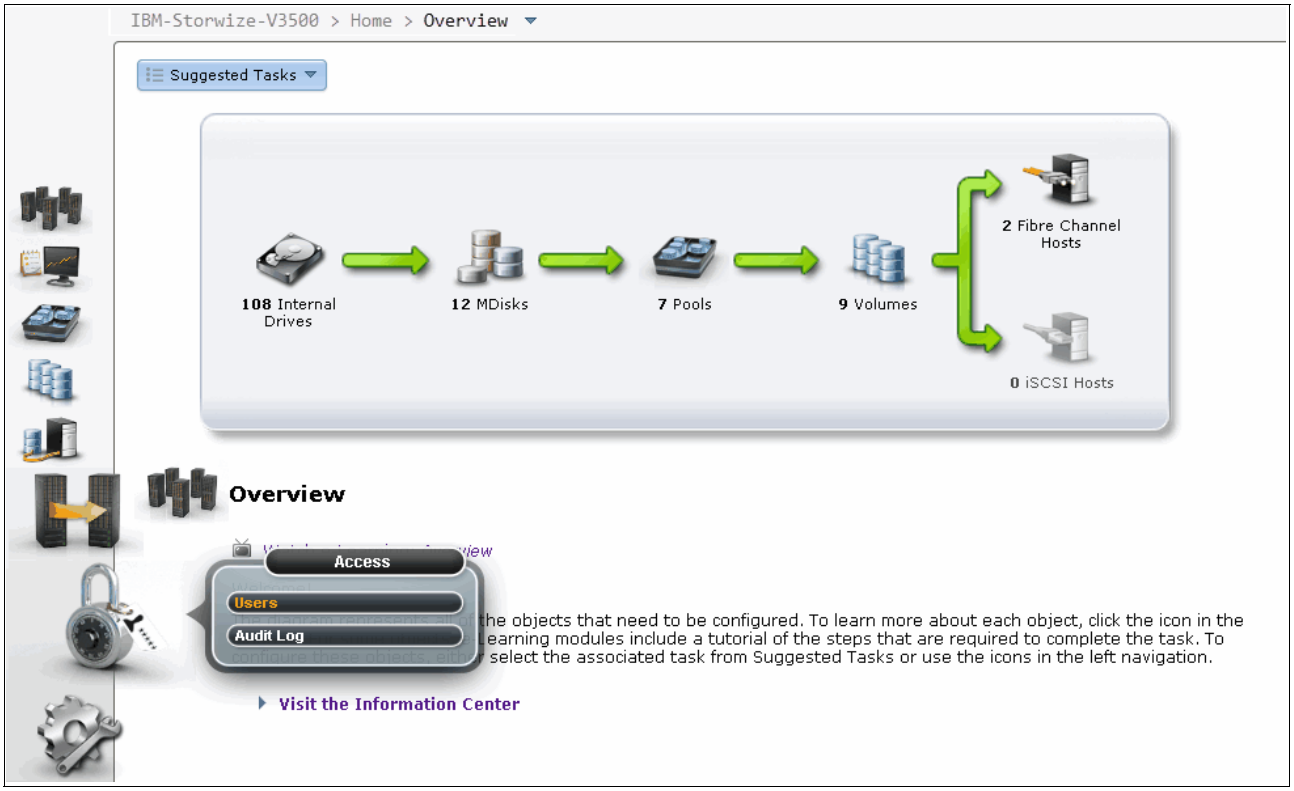

Figure A-6 Open user section

2. Right-click the user for which you want to upload the key and click **Properties**, as shown in Figure A-7.

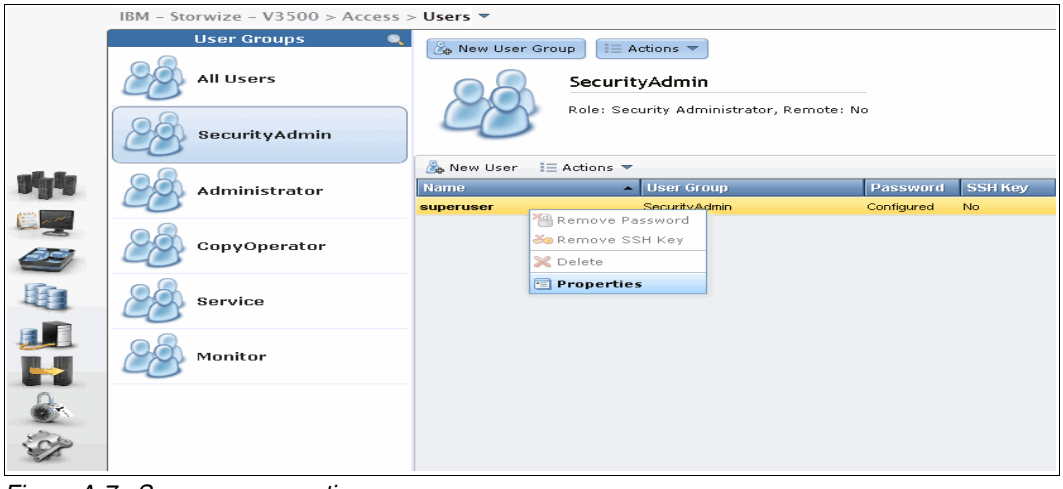

Figure A-7 Superuser properties

3. To upload the public key, click **Browse**, select your public key, and click **OK**, as shown in Figure A-8.

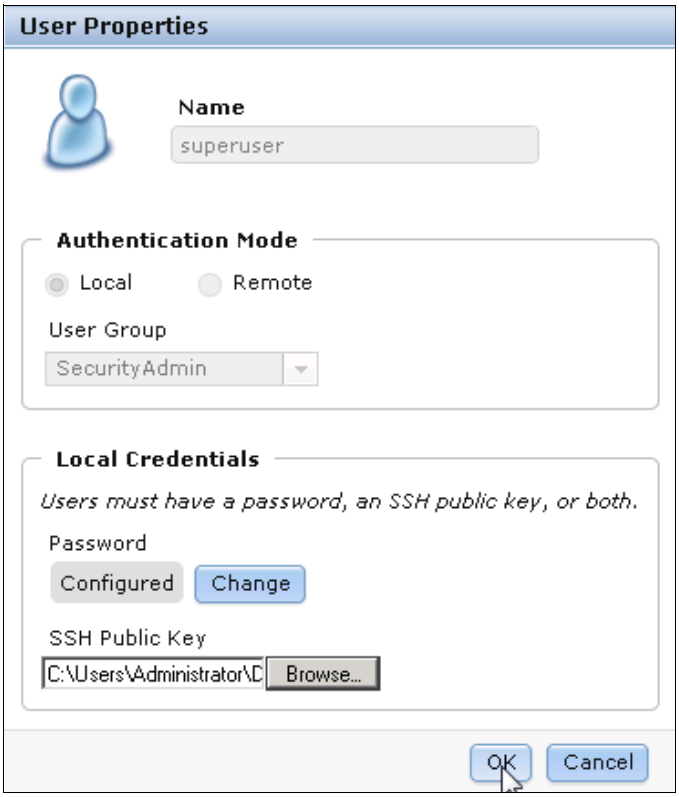

Figure A-8 Select public key

4. Click **OK** and the key is uploaded, as shown in Figure A-9.

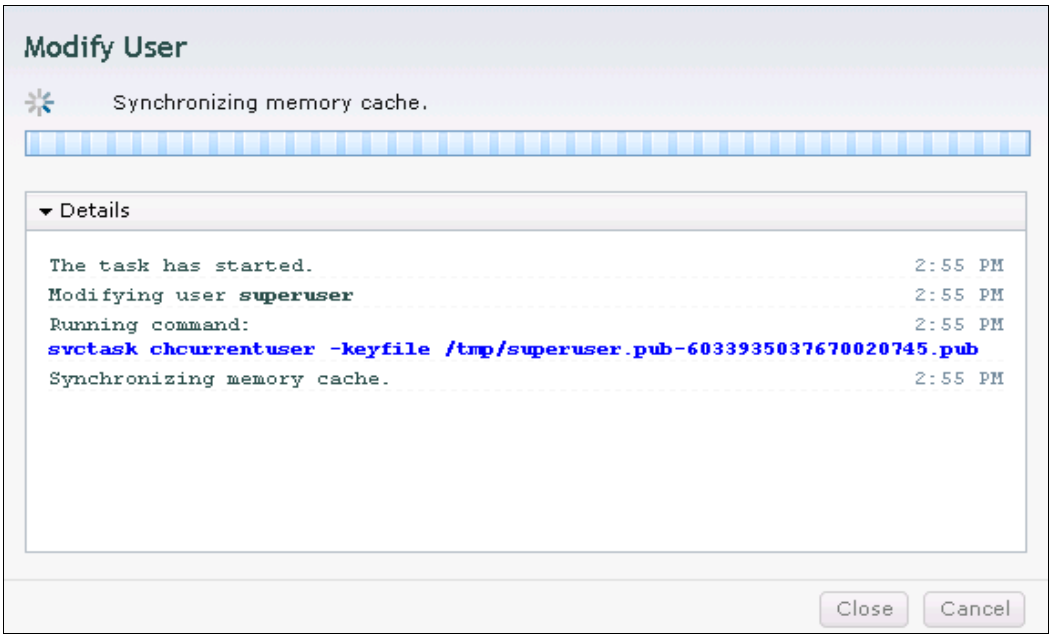

Figure A-9 Public key upload complete

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 SSH client:

1. Start PuTTY, as shown in Figure A-10.

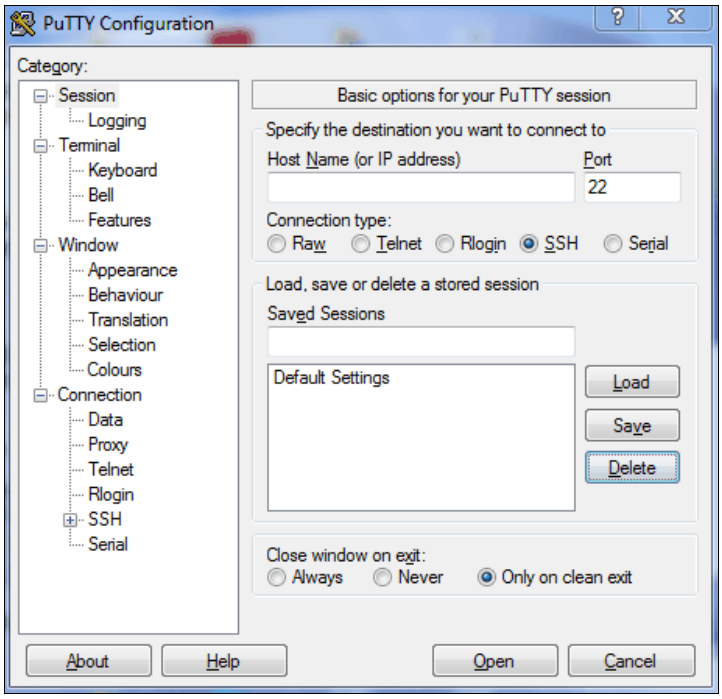

Figure A-10 PuTTY

In the right side pane under the “Specify the destination that 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 on 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.

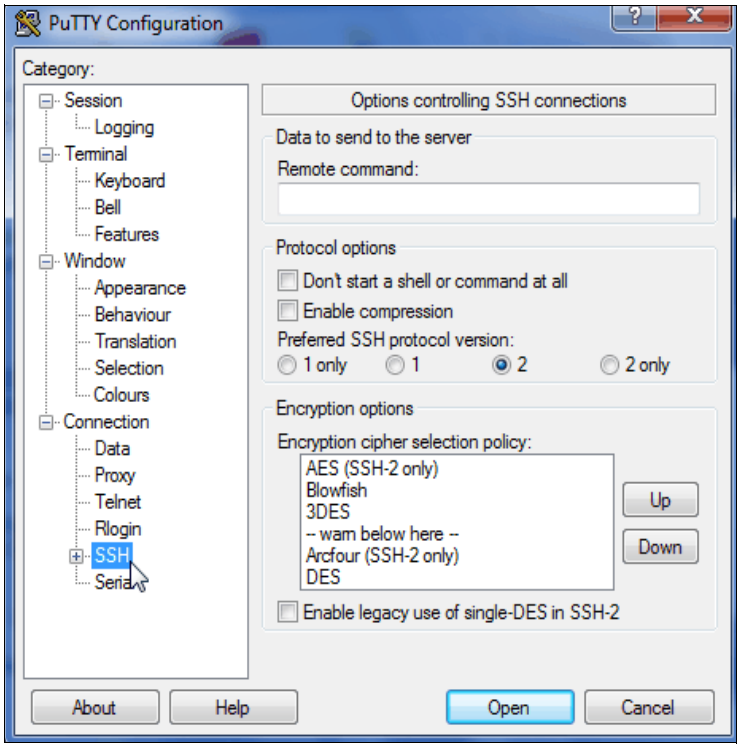

Figure A-11 SSH protocol version 2

3. In the “Preferred SSH protocol version” section of the right side pane, 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 “Private key file for authentication:” field under the Authentication Parameters section in the right side pane, browse to or 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`).

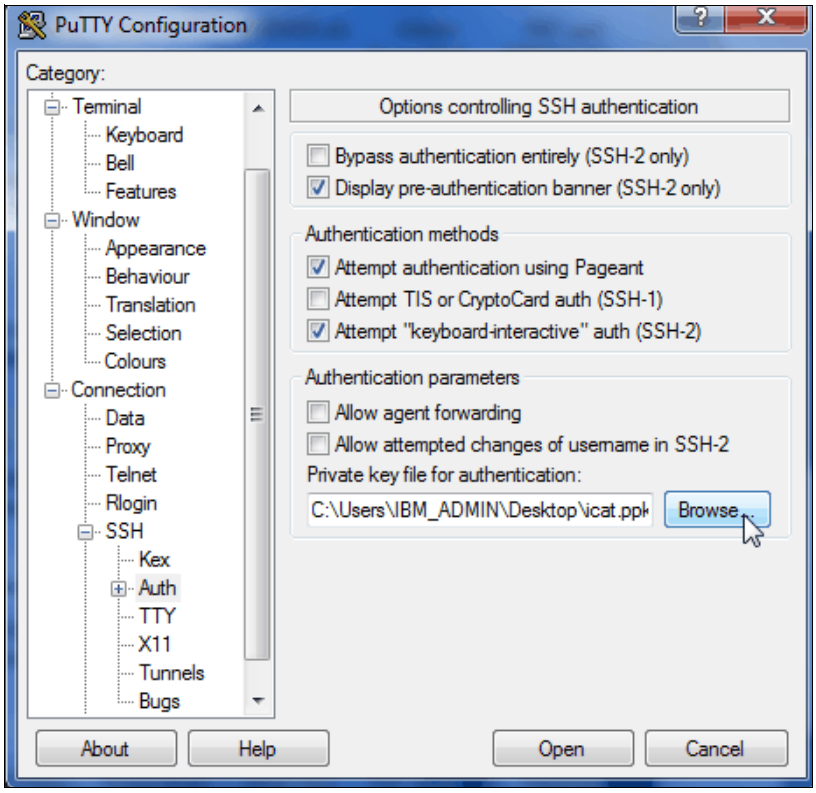

Figure A-12 SSH authentication

5. From the Category pane on the left side of the PuTTY Configuration window, click **Session** to return to the Session view (see Figure A-10 on page 490).

6. In the right side pane, enter the host name or system IP address of the IBM Storwize V3500 clustered system in the Host Name field. Enter a session name in the Saved Sessions field, as shown in Figure A-13.

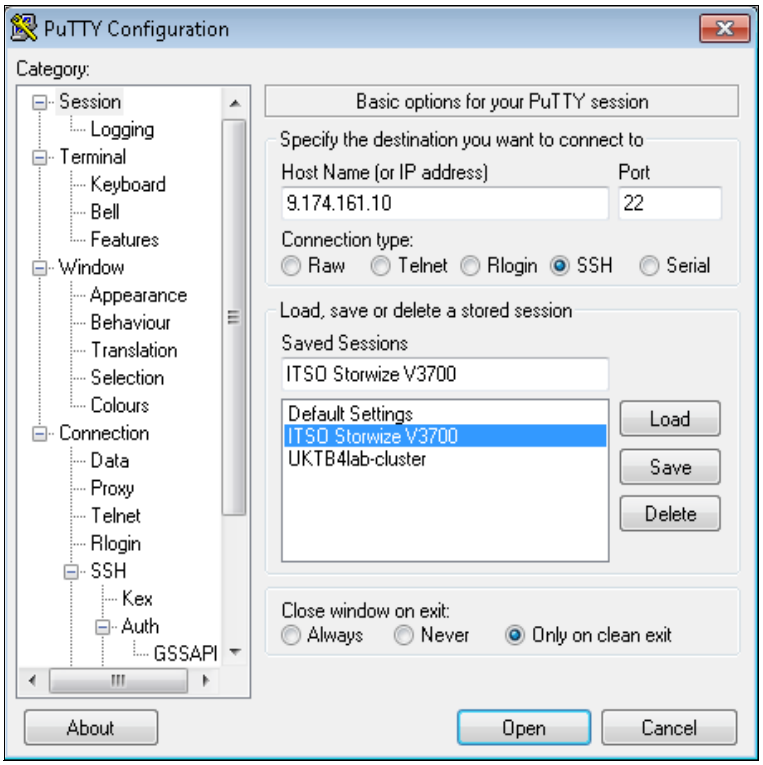

Figure A-13 Enter session information

7. Click **Save** to save the new session, as shown in Figure A-14.

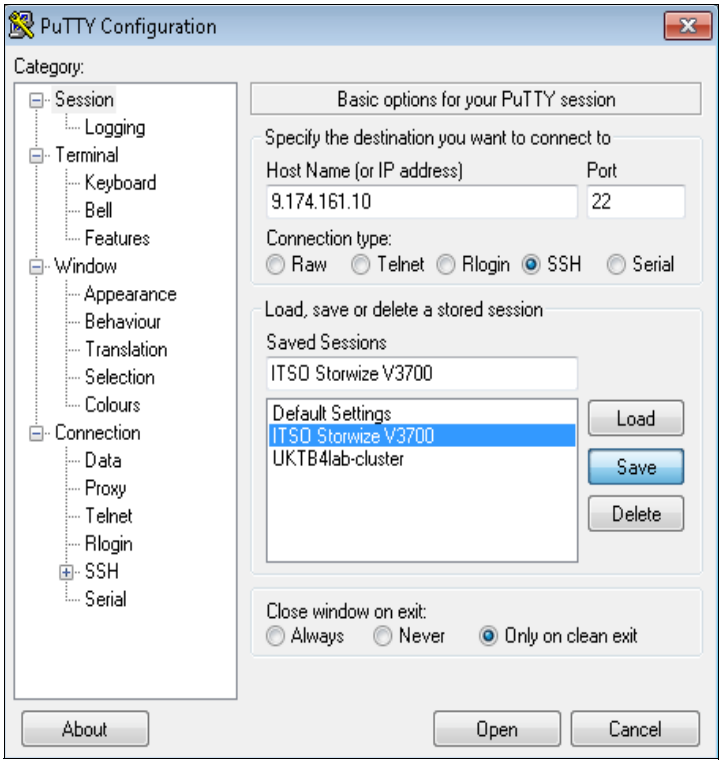

Figure A-14 Save Session

8. Highlight the new session and click **Open** to connect to the IBM Storwize V3500 system.
9. PuTTY now connects to the system and prompts you for a user name. Enter `admin` as the user name and press Enter, as shown in Example A-1.

Example: A-1 Enter user name

```
login as: admin
admin@9.174.161.10's password:
Last login: Thu Nov 29 11:13:14 2012 from 9.174.219.167
IBM_2071:ITSO-Storwize-V3500:admin>
```

The CLI is now configured for IBM Storwize V3500 administration.

Example commands

A detailed description of all the available commands is beyond the intended scope of this book. In this section are sample commands that we referenced in this book.

The **`svcinfo`** and **`svctask`** prefixes are no longer needed in IBM Storwize V3500 V6.3. 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 the **lsvdisk** command, as shown in Example A-2, to list all configured volumes on the system. The example shows that three volumes are configured.

Example: A-2 List all volumes

```
IBM_2071:ITSO-Storwize-V3500-1:admin>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                       fc_map_count copy_count fast_write_state se_copy_count RC_change
0  Hyper-V_PCO_Node1 0           io_grp0       online many         many           100.00GB  many
600507680185853FF000000000000000F 0            2           empty            0             no
1  ESXi-PCO          0           io_grp0       online many         many           100.00GB  many
600507680185853FF0000000000000010 0            2           empty            0             no
2  ESXi-Redbooks     0           io_grp0       online many         many           400.00GB  many
600507680185853FF0000000000000011 0            2           not_empty        0             no
3  Hyper-V_PCO_Node2 0           io_grp0       online many         many           101.00GB  many
600507680185853FF0000000000000012 0            2           empty            0             no
4  ESXi-Demo         0           io_grp0       online many         many           1000.00GB many
600507680185853FF0000000000000013 0            2           not_empty        0             no
```

Enter the **lshost** command to see a list of all configured hosts on the system, as shown in Example A-3.

Example: A-3 List hosts

```
IBM_2071:ITSO-Storwize-V3500-1:admin>lshost
id name          port_count iogrp_count status
2  HyperV-1_FCoE 2          4           online
3  HyperV-2_FCoE 2          4           online
4  ESXi-1        2          4           online
5  ESXi-2        2          4           online
```

To map the volume to the hosts, enter the **mkvdiskhostmap** command, as shown in Example A-4.

Example: A-4 Map volumes to host

```
IBM_2071:ITSO-Storwize-V3500-1:admin>mkvdiskhostmap -host ESXi-1 -scsi 0   -force
ESXi-Redbooks
Virtual Disk to Host map, id [0], successfully created
```

To verify the host mapping, enter the **lsvdiskhostmap** command, as shown in Example A-5.

Example: A-5 List all hosts that are mapped to a volume

```
IBM_2071:ITSO-Storwize-V3500-1:admin>lshostvdiskmap ESXi-1
id name   SCSI_id vdisk_id vdisk_name    vdisk_UID
4  ESXi-1 0       2        ESXi-Redbooks 600507680185853FF0000000000000011
```

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.

Command output: The **lsvdisk** command lists all available properties of a volume and its copies. To make it easier to read, lines in the example output were deleted.

Example: A-6 Enable advanced properties: I/O throttling

```
IBM_2071:ITSO-Storwize-V3500-1:admin>chvdisk -rate 1200 -unit mb ESXi-Redbooks
IBM_2071:ITSO-Storwize-V3500-1:admin>
IBM_2071:ITSO-Storwize-V3500-1:admin>lsvdisk ESXi-Redbooks

id 2
name ESXi-Redbooks
.
.
vdisk_UID 600507680185853FF000000000000011
virtual_disk_throttling (MB) 1200
preferred_node_id 2
.
.
IBM_2071:ITSO-Storwize-V3500-1:admin>
```

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_2071:ITSO-Storwize-V3500-1:admin>chvdisk -rate 4000 ESXi-Redbooks
IBM_2071:ITSO-Storwize-V3500-1:admin>lsvdisk ESXi-Redbooks
id 2
name ESXi-Redbooks
.
.
vdisk_UID 600507680185853FF000000000000011
throttling 4000
preferred_node_id 2
.
.
IBM_2071:ITSO-Storwize-V3500-1:admin>
```

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_2071:ITSO-Storwize-V3500-1:admin>chvdisk -rate 0 ESXi-Redbooks
IBM_2071:ITSO-Storwize-V3500-1:admin>lsvdisk ESXi-Redbooks
id 2
.
.
vdisk_UID 600507680185853FF000000000000011
throttling 0
preferred_node_id 2
IBM_2071:ITSO-Storwize-V3500-1:admin>
```

SAN Boot

IBM Storwize V3500 supports SAN Boot for Windows, VMware, and many other operating systems. SAN Boot support can change, so regularly check the IBM Storwize V3500 interoperability matrix at this website:

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

The IBM Storwize V3500 Information Center has a great deal of information about SAN Boot in combination with different operating systems. For more information, see this website:

http://pic.dhe.ibm.com/infocenter/storwize/v3500_ic/index.jsp?topic=%2Fcom.ibm.storwize.v3500.641.doc%2Fsvc_hostattachmentmain.html

More information about SAN Boot is available in the publication *IBM System Storage. Multipath Subsystem Device Driver User's. Guide*, GC52- 1309-03, which can be found 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 Windows host by using SAN Boot:

1. Configure the IBM Storwize V3500 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 V3500 system node port. The use of multiple paths during installation is not supported.
3. Configure and enable the HBA BIOS.
4. Install the operating system by using the normal procedure. 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 is complete.
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 V3500 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 V3500 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. 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 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 V3500 system.

SAN Boot procedures: For more information about SAN Boot procedures for other operating systems, check the IBM Storwize V3500 Information Center at this website:

http://pic.dhe.ibm.com/infocenter/storwize/v3500_ic/index.jsp?topic=%2Fcom.ibm.storwize.v3500.641.doc%2Fsvc_hostattachmentmain.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 Subsystem Device Driver (SDD) matrix to download packages that are available 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 the image-to-host mappings from the storage controller.
   c. Map the existing SAN Boot image and any other disks to the IBM Storwize V3500 system.

4. Change the zoning so that the host can see the IBM Storwize V3500 I/O group for the target image mode volume.
5. Complete the following configuration changes on the IBM Storwize V3500 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 **bootdiskmigrate** in step 1, 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 V3500.
9. Boot the host in single-path mode.
10. Uninstall any multipathing driver that is not supported for IBM Storwize V3500 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 the 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 that you 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

The publications that are 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 publication. Some publications that are referenced in this 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 V3500 publications

Storwize V3500 publications are available at this website:

`https://ibm.biz/BdxyDT`

IBM Storwize V3500 support

Storwize V3500 support is available at this website:

`https://ibm.biz/BdxyDk`

Help from IBM

IBM Support and downloads:

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

IBM Global Services:

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

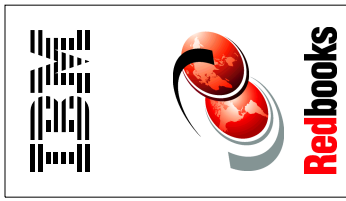

Implementing the Storwize V3500

(1.0” spine)
0.875”<->1.498”
460 <-> 788 pages

Implementing the Storwize V3500

Consolidate and share data with low-cost iSCSI

Deploy storage in minutes and perform storage management

Experience peace of mind

Businesses of all sizes are faced with the challenge of managing huge volumes of data that are becoming increasingly valuable. But storing this data can be costly, and extracting value from the data is becoming more and more difficult. IT organizations have limited resources and cannot afford to make investment mistakes. The IBM Storwize V3500 system provides a smarter solution that is affordable, simple, and efficient, which enables businesses to overcome their storage challenges.

IBM Storwize V3500 is the most recent addition to the IBM Storwize family of disk systems. It delivers easy-to-use, entry-level configurations that are specifically designed to meet the modest budgets of small and medium-sized businesses.

IBM Storwize V3500 features the following highlights:

- Consolidate and share data with low cost iSCSI storage networking.
- Deploy storage in minutes and perform storage management tasks quickly and easily through a breakthrough graphical user interface.
- Experience peace of mind with proven IBM Storwize family high-availability data protection with snapshot technology and IBM warranty support.
- Optimize efficiency by allocating only the amount of disk space needed at the time it is required with high performance, thin-provisioning capabilities.

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

SG24-8125-00 ISBN 0738438154