


Implementing the IBM Storwize V5000

Easily manage and deploy systems with embedded GUI

Experience rapid and flexible provisioning

Protect data with remote mirroring



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

Redbooks



International Technical Support Organization

Implementing the IBM Storwize V5000

October 2013

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

First Edition (October 2013)

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

© Copyright International Business Machines Corporation 2013. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	ix
Trademarks	x
Preface	xi
Authors	xi
Now you can become a published author, too!	xiii
Comments welcome	xiii
Stay connected to IBM Redbooks	xiv
Chapter 1. Overview of the IBM Storwize V5000 system.	1
1.1 IBM Storwize V5000 overview	2
1.2 IBM Storwize V5000 terminology	3
1.3 IBM Storwize V5000 models	5
1.4 IBM Storwize V5000 hardware	7
1.4.1 Control enclosure	8
1.4.2 Expansion enclosure	9
1.4.3 Host connectivity	10
1.4.4 Disk drive types	10
1.5 IBM Storwize V5000 terms	11
1.5.1 Hosts	11
1.5.2 Node canister	11
1.5.3 I/O groups	12
1.5.4 Clustered system	12
1.5.5 RAID	13
1.5.6 Managed disks	14
1.5.7 Quorum disks	14
1.5.8 Storage pools	14
1.5.9 Volumes	16
1.5.10 iSCSI	19
1.5.11 SAS	19
1.6 IBM Storwize V5000 features	19
1.6.1 Mirrored volumes	19
1.6.2 Thin provisioning	20
1.6.3 Easy Tier	21
1.6.4 Storage Migration	21
1.6.5 FlashCopy	21
1.6.6 Remote Copy	22
1.6.7 External virtualization	23
1.7 Problem management and support	23
1.7.1 IBM Assist On-site and remote service	23
1.7.2 Event notifications	24
1.7.3 SNMP traps	24
1.7.4 Syslog messages	24
1.7.5 Call Home email	24
1.8 More information resources	25
1.8.1 Useful IBM Storwize V5000 websites	25
1.8.2 IBM Storwize learning videos on YouTube	26
Chapter 2. Initial configuration	27

2.1	Hardware installation planning	28
2.2	SAN configuration planning	31
2.3	FC Direct-attach planning	33
2.4	SAS Direct-attach planning	33
2.5	LAN configuration planning	35
2.5.1	Management IP address considerations	35
2.5.2	Service IP address considerations	36
2.6	Host configuration planning	37
2.7	Miscellaneous configuration planning	38
2.8	System management	39
2.8.1	GUI	39
2.8.2	CLI	40
2.9	First-time setup	41
2.10	Initial configuration	49
2.10.1	Adding Enclosures after initial configuration	59
2.10.2	Configuring Call Home, email alert, and inventory	69
2.10.3	Service Assistant tool	71
Chapter 3. Graphical user interface overview		75
3.1	Getting started	76
3.1.1	Supported browsers	76
3.1.2	Access the management GUI	76
3.1.3	Overview panel layout	79
3.2	Navigation	80
3.2.1	Function icons navigation	81
3.2.2	Extended help navigation	82
3.2.3	Breadcrumb navigation aid	84
3.2.4	Suggested Tasks feature	85
3.2.5	Presets	86
3.2.6	Access actions	87
3.2.7	Task progress	88
3.2.8	Navigating panels with tables	88
3.3	Status Indicators menus	93
3.3.1	Horizontal bars	93
3.3.2	Allocated status bar menu	93
3.3.3	Running tasks bar menu	95
3.3.4	Health status bar menu	97
3.4	Function icon menus	99
3.4.1	Home menu	100
3.4.2	Monitoring menu	100
3.4.3	Pools menu	110
3.4.4	Volumes menu	121
3.4.5	Hosts menu	125
3.4.6	Copy Services menu	128
3.4.7	Access menu	133
3.4.8	Settings menu	137
3.5	Management GUI help	148
3.5.1	IBM Storwize V5000 Information Center	148
3.5.2	Watching an e-Learning video	148
3.5.3	Learning more	149
3.5.4	Embedded panel help	149
3.5.5	Hidden question mark help	150
3.5.6	Hover help	150

3.5.7 IBM endorsed YouTube videos	151
Chapter 4. Host configuration	153
4.1 Host attachment overview	154
4.2 Preparing the host operating system	155
4.2.1 Windows 2008 R2: Preparing for FC attachment	155
4.2.2 Creating SAS hosts	157
Chapter 5. I/O Group basic volume configuration	161
5.1 Provisioning storage from IBM Storwize V5000 and making it available to the host.	162
5.1.1 Creating a generic volume	164
5.1.2 Creating a thin-provisioned volume.	167
5.1.3 Creating a mirrored volume.	169
5.1.4 Creating a thin-mirror volume	174
5.2 Mapping a volume to the host.	177
5.2.1 Mapping newly created volumes to the host by using the wizard	177
5.2.2 Manually mapping a volume to the host	181
5.3 Discovering the volumes from the host and specifying multipath settings	185
5.3.1 Windows 2008 Fibre Channel volume attachment	186
5.3.2 Windows 2008 iSCSI volume attachment.	191
5.3.3 Windows 2008 Direct SAS volume attachment.	203
5.3.4 VMware ESX Fibre Channel volume attachment	207
5.3.5 VMware ESX iSCSI volume attachment	215
5.3.6 VMware ESX Direct SAS volume attachment.	227
Chapter 6. Storage migration wizard	237
6.1 Interoperability and compatibility.	238
6.2 Storage migration wizard	238
6.2.1 External virtualization capability	238
6.2.2 Overview of the storage migration wizard.	238
6.2.3 Storage migration wizard tasks	239
6.3 Storage migration wizard example scenario	253
6.3.1 Storage migration wizard example scenario description.	253
6.3.2 Using the storage migration wizard for example scenario	255
Chapter 7. Storage pools	295
7.1 Working with internal drives	296
7.1.1 Internal Storage window	297
7.1.2 Actions on internal drives	299
7.2 Configuring internal storage	307
7.2.1 RAID configuration presets	307
7.2.2 Customizing initial storage configuration.	309
7.2.3 Creating an MDisk and pool	310
7.2.4 Using the recommended configuration	312
7.2.5 Selecting a different configuration.	314
7.3 Working with MDisk on internal and external storage	320
7.3.1 Adding Externally Virtualized MDisk to storage pools.	322
7.3.2 Importing externally virtualized MDisk to storage pools	326
7.3.3 MDisk by Pools panel	332
7.3.4 RAID action for MDisk	334
7.3.5 Selecting the drive tier for externally virtualized MDisk	338
7.3.6 More actions on MDisk	339
7.4 Working with storage pools	343
7.4.1 Create Pool option	345

7.4.2	Actions on storage pools	346
Chapter 8.	Advanced host and volume administration	349
8.1	Advanced host administration	350
8.1.1	Modifying Mappings menu	352
8.1.2	Unmapping volumes from a host	356
8.1.3	Renaming a host	359
8.1.4	Deleting a host	361
8.1.5	Host properties	362
8.2	Adding and deleting host ports	367
8.2.1	Adding a host port	368
8.2.2	Adding a Fibre Channel port	369
8.2.3	Adding a SAS host port	370
8.2.4	Adding an iSCSI host port	371
8.2.5	Deleting a host port	372
8.3	Host mappings overview	373
8.3.1	Unmap Volumes	374
8.3.2	Properties (Host)	374
8.3.3	Properties (Volume)	374
8.4	Advanced volume administration	375
8.4.1	Advanced volume functions	376
8.4.2	Mapping a volume to a host	379
8.4.3	Unmapping volumes from all hosts	380
8.4.4	Viewing a host that is mapped to a volume	380
8.4.5	Renaming a volume	381
8.4.6	Shrinking a volume	382
8.4.7	Expanding a volume	382
8.4.8	Migrating a volume to another storage pool	383
8.4.9	Exporting to an image mode volume	385
8.4.10	Deleting a volume	387
8.5	Volume properties	388
8.5.1	Overview tab	388
8.5.2	Host Maps tab	391
8.5.3	Member MDisk tab	392
8.5.4	Adding a mirrored volume copy	394
8.5.5	Editing thin-provisioned volume properties	395
8.6	Advanced volume copy functions	398
8.6.1	Thin-provisioned menu	399
8.6.2	Splitting into a new volume	400
8.6.3	Validate Volume Copies option	402
8.6.4	Delete Volume Copy option	404
8.6.5	Migrating volumes by using the volume copy features	404
8.7	Volumes by Storage Pool	406
8.8	Volumes by Host	408
Chapter 9.	Easy Tier	411
9.1	Easy Tier overview	412
9.2	Easy Tier for IBM Storwize V5000	413
9.2.1	Disk tiers	413
9.2.2	Tiered storage pools	414
9.3	Easy Tier process	415
9.3.1	I/O Monitoring	415
9.3.2	Data Placement Advisor	415

9.3.3	Data Migration Planner	415
9.3.4	Data Migrator	416
9.3.5	Easy Tier operating modes	416
9.3.6	Easy Tier rules	417
9.4	Easy Tier configuration by using the GUI	419
9.4.1	Creating multitiered pools: Enable Easy Tier	419
9.4.2	Downloading Easy Tier I/O measurements.	427
9.5	Easy Tier configuration by using the command-line interface.	429
9.5.1	Enabling Easy Tier evaluation mode.	429
9.5.2	Enabling or disabling Easy Tier on single volumes.	432
9.6	IBM Storage Tier Advisor Tool	434
9.6.1	Creating graphical reports.	434
9.6.2	STAT reports.	435
9.7	Tivoli Storage Productivity Center.	436
9.7.1	Tivoli Storage Productivity Center benefits.	437
9.7.2	Adding IBM Storwize V5000 in Tivoli Storage Productivity Center	437
9.8	Administering and reporting an IBM Storwize V5000 system through Tivoli Storage Productivity Center.	439
9.8.1	Basic configuration and administration	439
9.8.2	Generating reports by using Java GUI	441
9.8.3	Generating reports by using Tivoli Storage Productivity Center web console	444
Chapter 10.	Copy services	449
10.1	FlashCopy	450
10.1.1	Business requirements for FlashCopy	450
10.1.2	FlashCopy functional overview	451
10.1.3	Planning for FlashCopy.	459
10.1.4	Managing FlashCopy by using the GUI	461
10.1.5	Managing FlashCopy mappings	467
10.1.6	Managing a FlashCopy consistency group.	490
10.2	Remote Copy	500
10.2.1	Remote Copy concepts.	500
10.2.2	Global Mirror with Change Volumes	507
10.2.3	Remote Copy planning	512
10.3	Troubleshooting Remote Copy	515
10.3.1	1920 error	515
10.3.2	1720 error	517
10.4	Managing Remote Copy by using the GUI	517
10.4.1	Managing cluster partnerships	518
10.4.2	Managing stand-alone Remote Copy relationships	522
10.4.3	Managing a Remote Copy consistency group	534
Chapter 11.	External storage virtualization	547
11.1	Planning for external storage virtualization	548
11.1.1	License for external storage virtualization.	548
11.1.2	SAN configuration planning	550
11.1.3	External storage configuration planning	551
11.1.4	Guidelines for virtualizing external storage.	552
11.2	Working with external storage.	552
11.2.1	Adding external storage	553
11.2.2	Managing external storage	553
11.2.3	Removing external storage.	557
Chapter 12.	RAS, monitoring, and troubleshooting.	559

12.1 Reliability, availability, and serviceability on the IBM Storwize V5000	560
12.2 IBM Storwize V5000 components	561
12.2.1 Enclosure midplane assembly	561
12.2.2 Node canisters: Ports and LED	562
12.2.3 Node canister replaceable hardware components	566
12.2.4 Expansion canister: Ports and LED	570
12.2.5 Disk subsystem	571
12.2.6 Power supply unit	574
12.3 Configuration backup procedure	576
12.3.1 Generating a configuration backup by using the CLI	576
12.3.2 Downloading a configuration backup by using the GUI	577
12.4 Upgrading software	580
12.4.1 Upgrading software automatically	581
12.4.2 GUI upgrade process	581
12.4.3 Upgrading software manually	584
12.5 Event log	589
12.5.1 Managing the event log	590
12.5.2 Alert handling and recommended actions	593
12.6 Collecting support information	601
12.6.1 Support information via GUI	601
12.6.2 Support information via Service Assistant	602
12.6.3 Support Information onto USB stick	603
12.7 Powering on and shutting down IBM Storwize V5000	605
12.7.1 Shutting down the system	605
12.7.2 Powering on	608
Appendix A. Command-line interface setup and SAN Boot	609
Command-line interface	610
Basic setup	610
Example commands	620
SAN Boot	623
Enabling SAN Boot for Windows	623
Enabling SAN Boot for VMware	624
Windows SAN Boot migration	624
Related publications and information	627
IBM Redbooks	627
IBM Storwize V5000 publications	627
IBM Storwize V5000 support	627
Help from IBM	627

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

The following terms are trademarks of other companies:

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

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

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

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

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

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

Preface

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

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

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

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

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

Authors

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



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



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



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



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

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



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

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

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

IBM Manchester Lab

- ▶ John Fairhurst
- ▶ Paul Marris
- ▶ Paul Merrison

IBM Hursley

- ▶ Mary Connell

IBM STG

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

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

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience by using leading-edge technologies. Your efforts 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 that is found at this website:

<http://www.ibm.com/redbooks>

- ▶ Send your comments in an email to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

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

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>



Overview of the IBM Storwize V5000 system

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

This chapter includes the following topics:

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

1.1 IBM Storwize V5000 overview

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

IBM Storwize V5000 features the following benefits:

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

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

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

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

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

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

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

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

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

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

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

Information also is available at this website:

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

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

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

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

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

1.2 IBM Storwize V5000 terminology

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

Table 1-1 IBM Storwize V5000 terminology

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

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

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

1.3 IBM Storwize V5000 models

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

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

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

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

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

Table 1-2 IBM Storwize V5000 models

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

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



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

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

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



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

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

1.4 IBM Storwize V5000 hardware

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

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

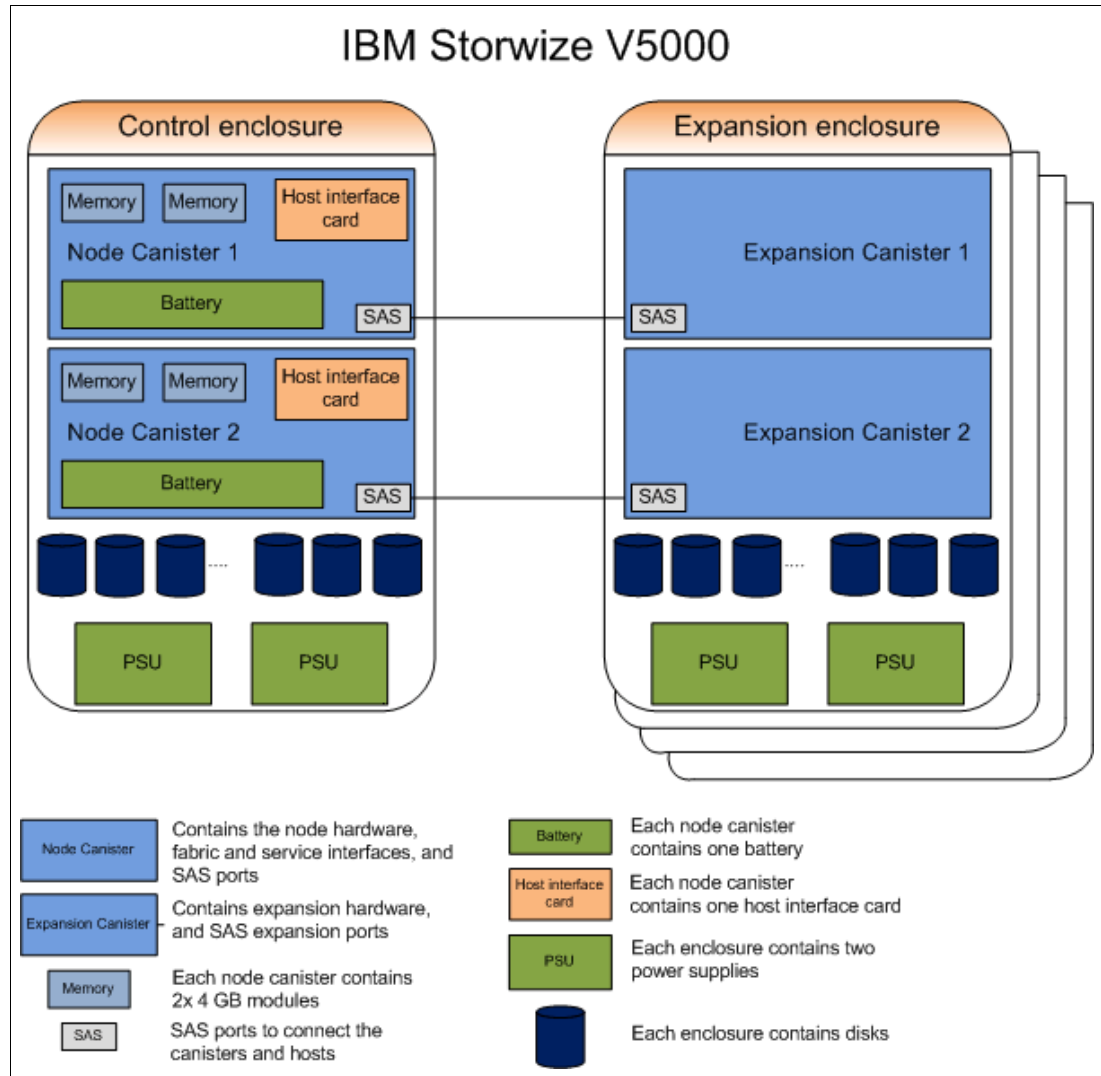


Figure 1-3 IBM Storwize V5000 hardware components

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



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

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

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



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

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

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

1.4.1 Control enclosure

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

Figure 1-6 shows a single node canister.




Figure 1-6 IBM Storwize V5000 node canister

Each node canister contains the following hardware:

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

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

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

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

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

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

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

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

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

1.4.2 Expansion enclosure

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

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



Figure 1-7 Expansion enclosure of the IBM Storwize V5000

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

Figure 1-8 shows the expansion canister ports.




Figure 1-8 Expansion canister ports

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

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

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

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

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

1.4.3 Host connectivity

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

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

1.4.4 Disk drive types

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

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

Table 1-3 IBM Storwize V5000 Disk Drive types

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

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

1.5 IBM Storwize V5000 terms

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

1.5.1 Hosts

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

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

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

1.5.2 Node canister

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

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

1.5.3 I/O groups

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

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

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

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

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

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

1.5.4 Clustered system

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

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

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

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

1.5.5 RAID

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

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

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

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

The following RAID levels are available:

- ▶ RAID 0 (striping, no redundancy)

RAID 0 arrays stripe data across the drives. The system supports RAID 0 arrays with one member, which is similar to traditional JBOD attach. RAID 0 arrays have no redundancy, so they do not support hot spare takeover or immediate exchange. A RAID 0 array can be formed by one to eight drives.
- ▶ RAID 1 (mirroring between two drives, which is implemented as RAID 10 with two drives)

RAID 1 arrays stripe data over mirrored pairs of drives. A RAID 1 array mirrored pair is rebuilt independently. A RAID 1 array can be formed by two drives only.
- ▶ RAID 5 (striping, can survive one drive fault, with parity)

RAID 5 arrays stripe data over the member drives with one parity strip on every stripe. RAID 5 arrays have single redundancy. The parity algorithm means that an array can tolerate no more than one member drive failure. A RAID 5 array can be formed by 3 - 16 drives.
- ▶ RAID 6 (striping, can survive two drive faults, with parity)

RAID 6 arrays stripe data over the member drives with two parity stripes (which is known as the P-parity and the Q-parity) on every stripe. The two parity strips are calculated by using different algorithms, which give the array double redundancy. A RAID 6 array can be formed by 5 - 16 drives.
- ▶ RAID 10 (RAID 0 on top of RAID 1)

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

1.5.6 Managed disks

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

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

An MDisk features the following modes:

- ▶ Array

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

- ▶ Unmanaged

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

- ▶ Managed

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

- ▶ Image

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

1.5.7 Quorum disks

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

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

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

1.5.8 Storage pools

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

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

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

Default extent size: The GUI of IBM Storwize V5000 has a default extent size value of 1 GB when you define a new storage pool. This is a change in the IBM Storwize code v7.1. The GUI cannot change the extent size. Therefore, if you want to create storage pools with a different extent size, this must be done via the CLI by using the `mkmdiskgrp` and `mkarray` commands.

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

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

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

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

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

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

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

Single-tiered storage pool

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

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

Multitiered storage pool

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

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

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

1.5.9 Volumes

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

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

The following types of volumes are available:

- ▶ Striped

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

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

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

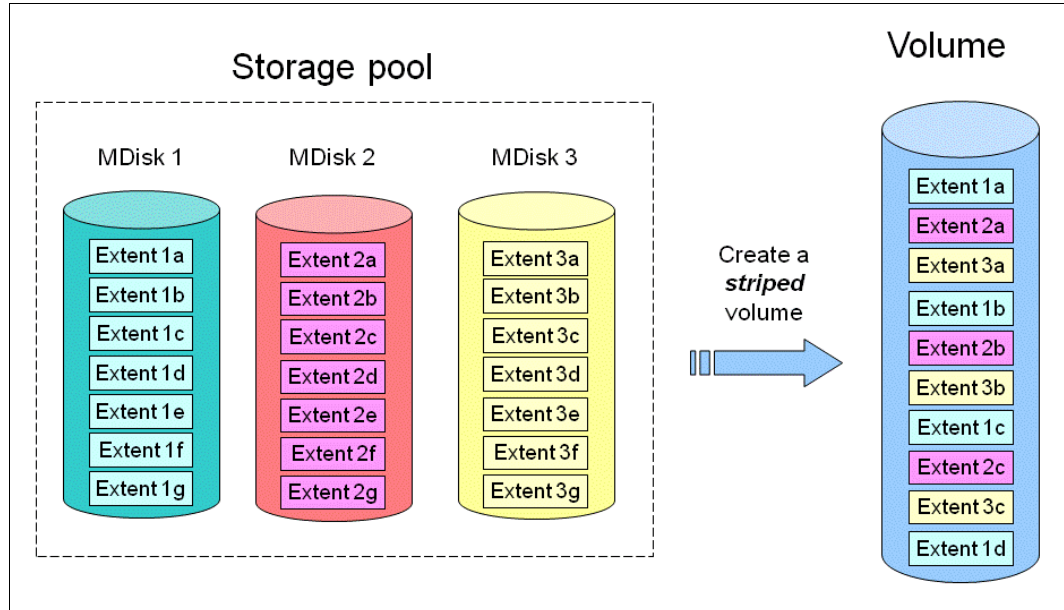


Figure 1-9 Striped volume

► Sequential

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

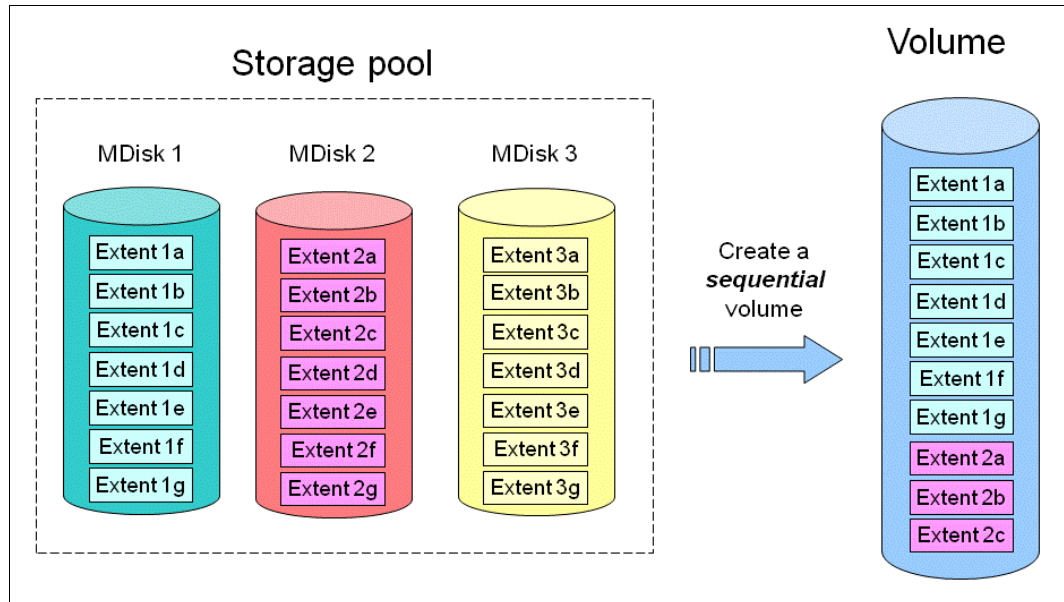


Figure 1-10 Sequential volume

► Image mode

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

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

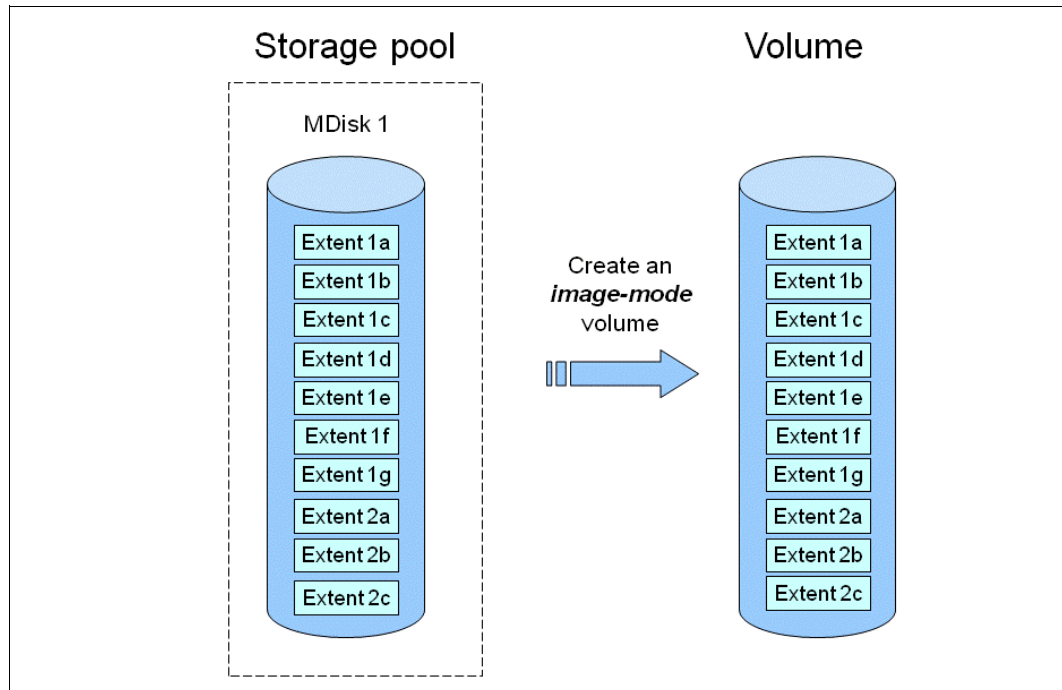


Figure 1-11 Image mode volume

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

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

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

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

1.5.10 iSCSI

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

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

Concepts of names and addresses are carefully separated in iSCSI.

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

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

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

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

1.5.11 SAS

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

1.6 IBM Storwize V5000 features

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

1.6.1 Mirrored volumes

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

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

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

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

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

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

1.6.2 Thin provisioning

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

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

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

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

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

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

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

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

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

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

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

1.6.3 Easy Tier

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

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

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

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

The IBM Easy Tier feature is licensed per enclosure.

1.6.4 Storage Migration

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

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

1.6.5 FlashCopy

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

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

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

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

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

The FlashCopy feature is licensed per enclosure.

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

1.6.6 Remote Copy

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

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

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

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

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

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

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

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

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

The IBM Remote Copy feature is licensed per enclosure.

Copy Services configuration limits

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

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

1.6.7 External virtualization

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

The External Virtualization feature is licensed per disk enclosure.

1.7 Problem management and support

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

1.7.1 IBM Assist On-site and remote service

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

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

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

<http://www.ibm.com/support/assistsite/>

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

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

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

1.7.2 Event notifications

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

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

1.7.3 SNMP traps

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

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

1.7.4 Syslog messages

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

1.7.5 Call Home email

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

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

1.8 More information resources

This section describes resources that are available for more information.

1.8.1 Useful IBM Storwize V5000 websites

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

- ▶ The IBM Storwize V5000 home page:
<http://www.ibm.com/storage/support/storwize/v5000>
- ▶ IBM Storwize V5000 Online Information Center:
http://pic.dhe.ibm.com/infocenter/storwize/v5000_ic/index.jsp

1.8.2 IBM Storwize learning videos on YouTube

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

Table 1-5 Videos available on YouTube

Video description	URL
IBM Storwize V7000 Storage Virtualization Terminology Overview	http://www.youtube.com/watch?v=I2rzt3m2gP0
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=vec0ap-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=BfGbKwCsR4
Introducing IBM Storwize V3700	http://www.youtube.com/watch?v=AePPKiXE4xM
IBM Storwize V3700 Initial Setup	http://www.youtube.com/watch?v=oj9uhTYe6gg
Storwize V7000 Installation	http://www.youtube.com/watch?v=kCCFxm5ZMV4

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



Initial configuration

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

This chapter includes the following topics:

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

2.1 Hardware installation planning

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

- Install the hardware as described in *IBM Storwize V5000 Quick Installation Guide Version 6.4.1*, GC27-4219
- For more information about planning the IBM Storwize V5000 environment, see this website:
http://pic.dhe.ibm.com/infocenter/storwize/V5000_ic/index.jsp?topic=%2Fcom.ibm.storwize.V5000.641.doc%2Fsvc_webplanning_21pb8b.html
- An appropriate 19-inch rack with 2 - 14 U of space should be available, depending on the number of enclosures to install. Each enclosure measures 2 U and a single control enclosure with up to six expansion enclosures constitutes an IBM Storwize V5000 system.
- There should be redundant power outlets in the rack for each of the two power cords that are included per enclosure. In all, 2 - 14 outlets are required, depending on the number of enclosures to install. The power cords conform to the IEC320 C13/C14 standards.
- A minimum of four Fibre Channel ports that are attached to the fabric are required. However, it is a best practice to use eight 2-Gbps, 4-Gbps, or 8-Gbps Fibre Channel ports.

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

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

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

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

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

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




Figure 2-1 Connecting the SAS expansion cables example

The following connections must be made:

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

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

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

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

2.2 SAN configuration planning

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

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

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

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

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

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

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




Figure 2-2 SAN cabling and zoning diagram

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

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

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

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

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

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

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

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

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

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

2.3 FC Direct-attach planning

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

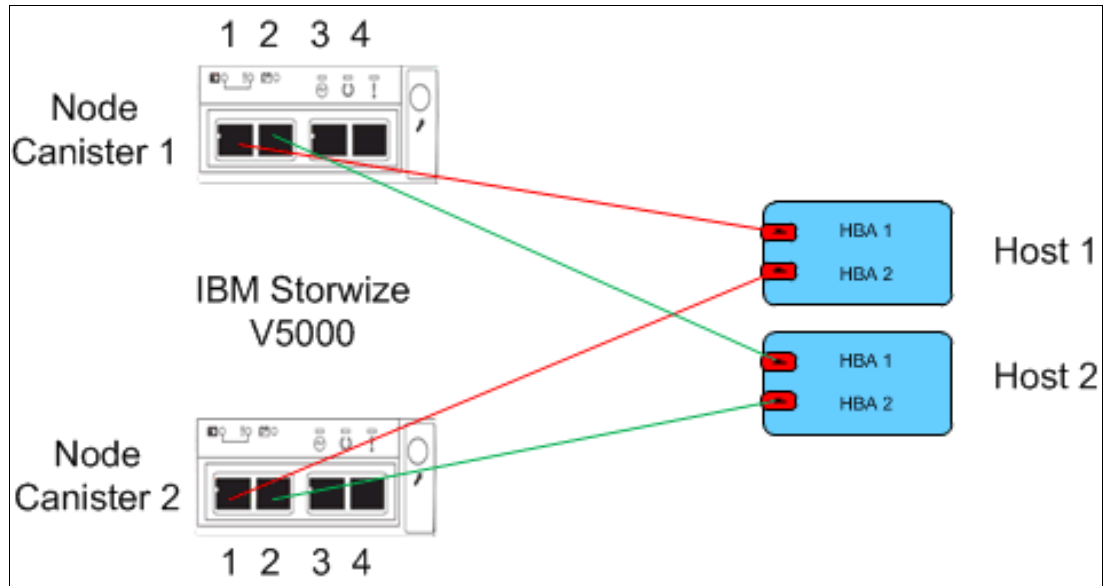


Figure 2-3 FC Direct-attach host configuration

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

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

2.4 SAS Direct-attach planning

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

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




Figure 2-4 SAS port identification

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




Figure 2-5 SAS host direct-attach

2.5 LAN configuration planning

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

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

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

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

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

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

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

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

2.5.1 Management IP address considerations

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

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

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

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

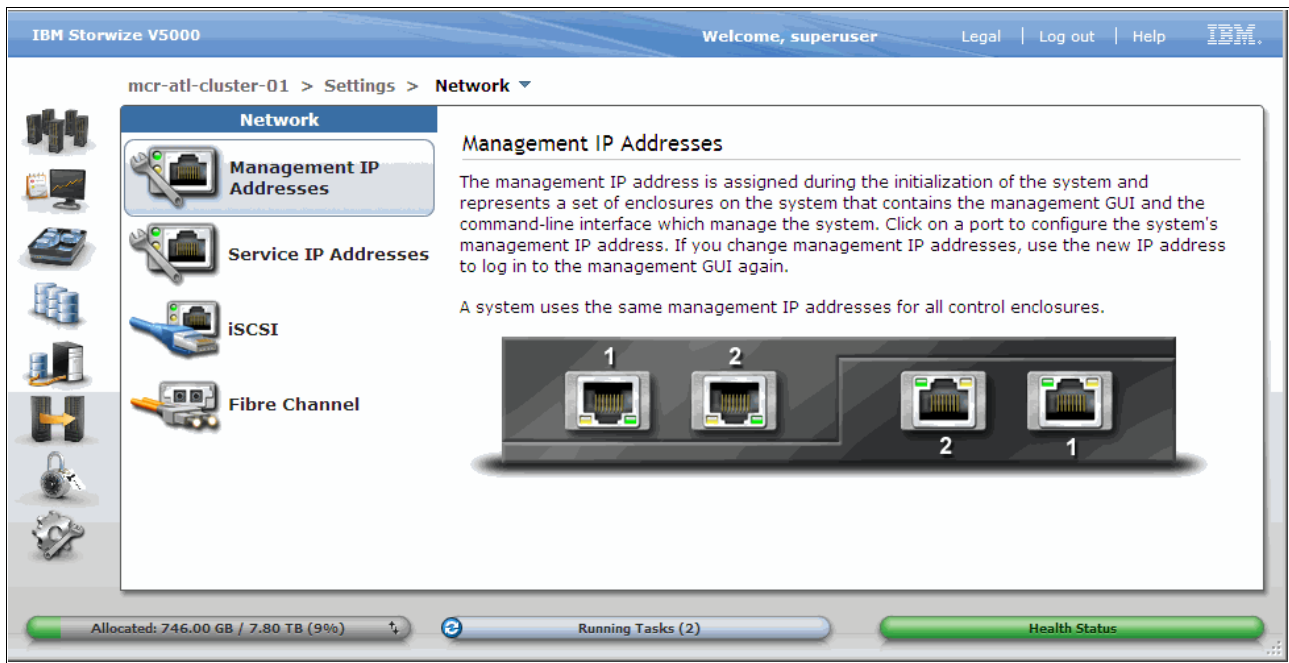


Figure 2-6 Ethernet ports available for configuration

2.5.2 Service IP address considerations

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

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



Figure 2-7 Service IP addresses available for configuration

2.6 Host configuration planning

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

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

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

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

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

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

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

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

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

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

2.7 Miscellaneous configuration planning

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

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

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

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

- ▶ Security Admin
- ▶ Administrator
- ▶ CopyOperator
- ▶ Service
- ▶ Monitor

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

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

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

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

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

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

2.8 System management

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

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

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

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

2.8.1 GUI

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

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

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

1. Browse to one of the following locations:
 - a. `http(s)://host name of your cluster/`
 - b. `http(s)://cluster IP address of your cluster/` Example: `https://192.168.70.120`
2. Use the following default login information:
 - User ID: superuser
 - Password: passw0rd

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

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

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

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

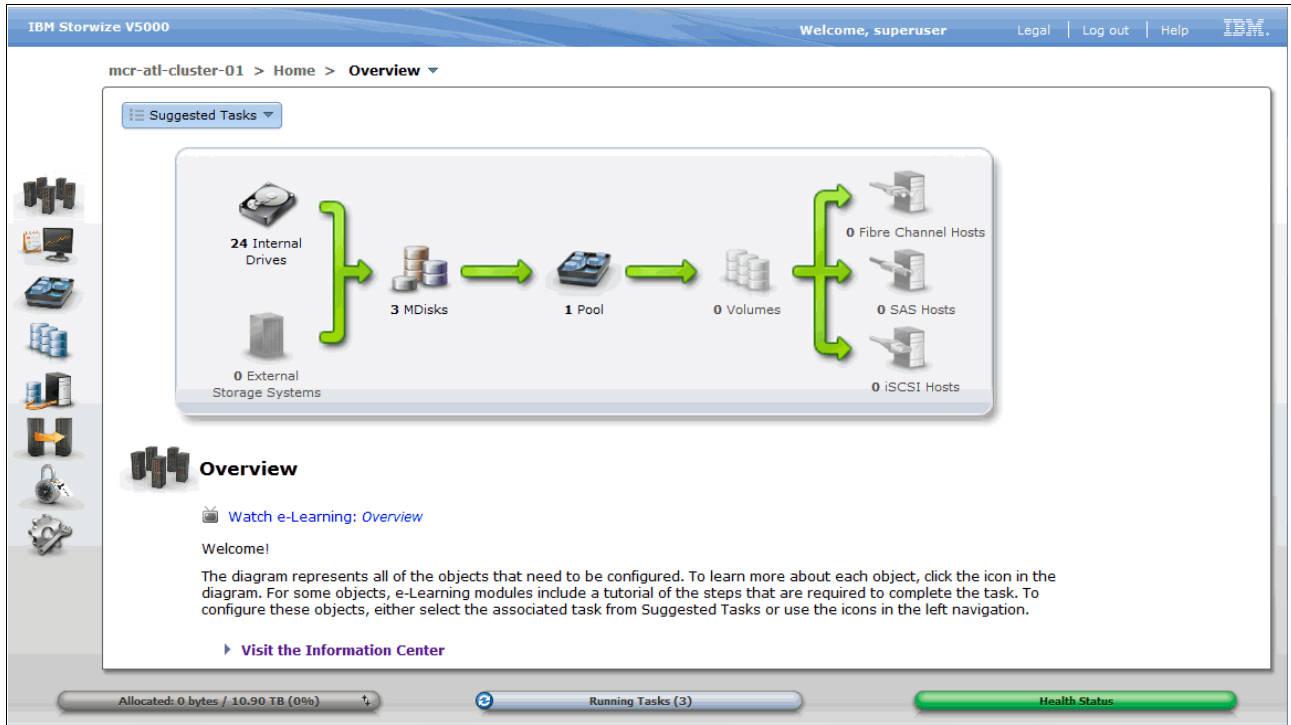


Figure 2-8 Setup wizard: Welcome window

2.8.2 CLI

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

Example 2-1 System management by using the CLI

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


1	10	online	online	yes	11
1	11	online	online	yes	8
1	12	online	online	yes	14
1	13	online	online	yes	15
1	14	online	online	yes	13
1	15	online	online	yes	16
1	16	online	online	yes	19
1	17	online	online	yes	1
1	18	online	online	yes	3
1	19	online	online	yes	6
1	20	online	online	yes	0
1	21	online	online	yes	4
1	22	online	online	yes	7
1	23	online	online	yes	2
1	24	online	online	yes	5

IBM_Storwize:mcr-at1-cluster-01:superuser>

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

2.9 First-time setup

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

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

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

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

<http://www.ibm.com/storage/support/Storwize/V5000>

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

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

We use Windows in the following examples.

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

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

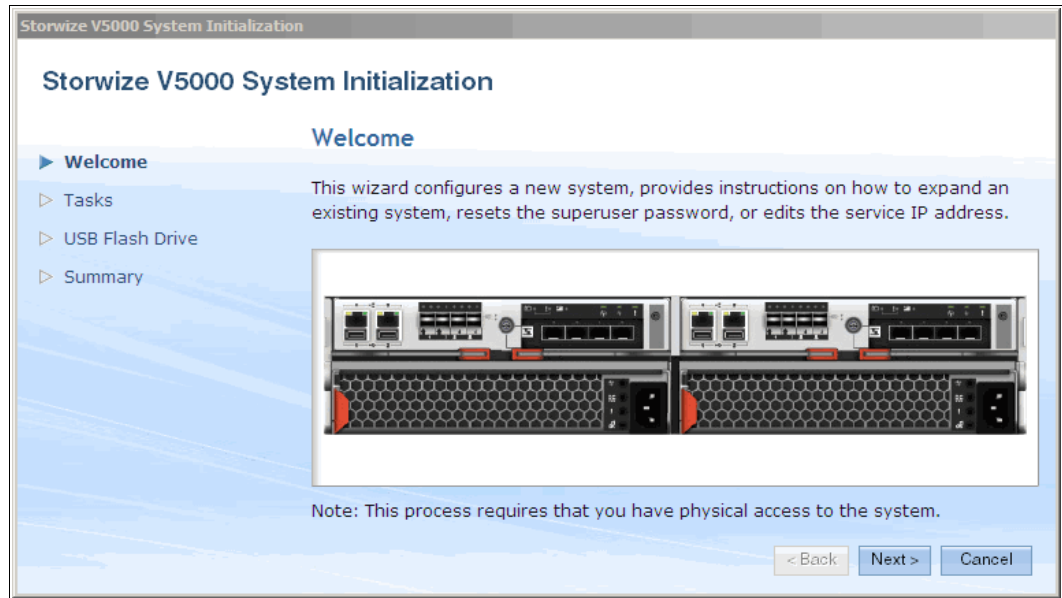


Figure 2-9 System Initialization: Welcome window

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

- a. Open a terminal window.
- b. Locate the root directory of the USB flash drive:
 - For Mac systems, the root directory is often in the `/Volumes/` directory.
 - For Linux systems, the root directory is often in the `/media/` directory.
 - If an automatic mount system is used, the root directory can be located by entering the mount command.
- c. Change the directory to the root directory of the flash drive.
- d. Enter: `sh InitTool.sh`

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

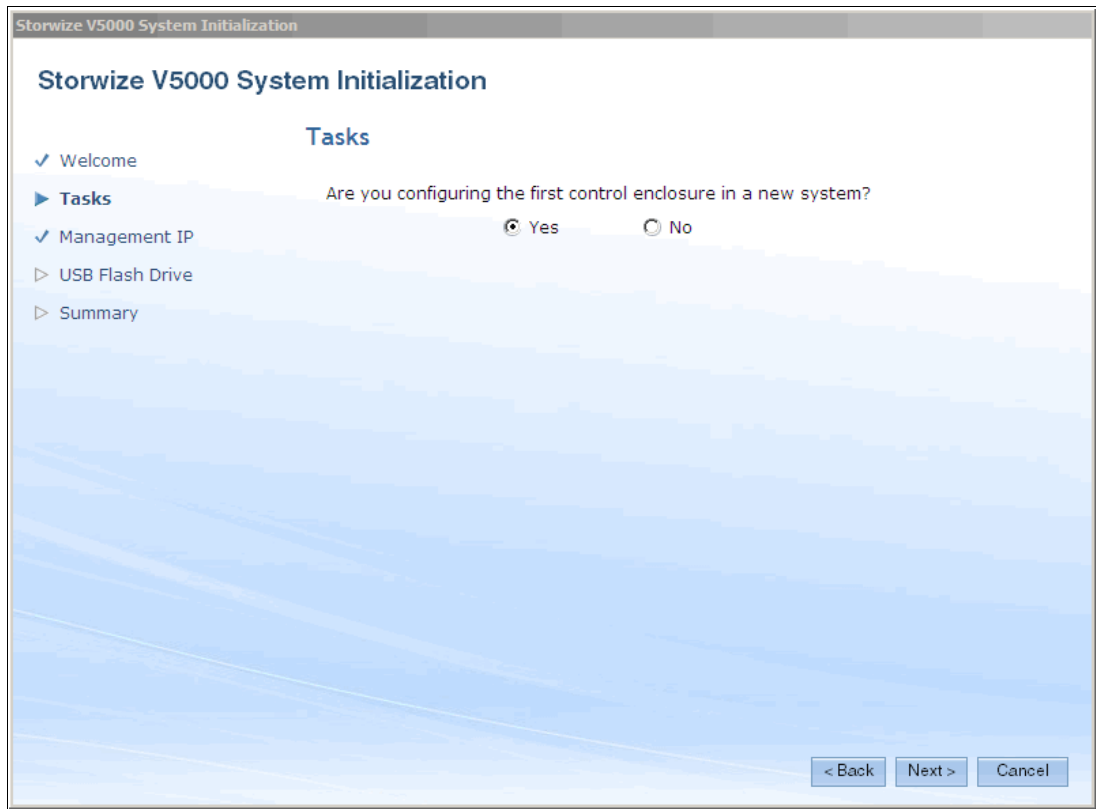


Figure 2-10 System Initialization: Create a system

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

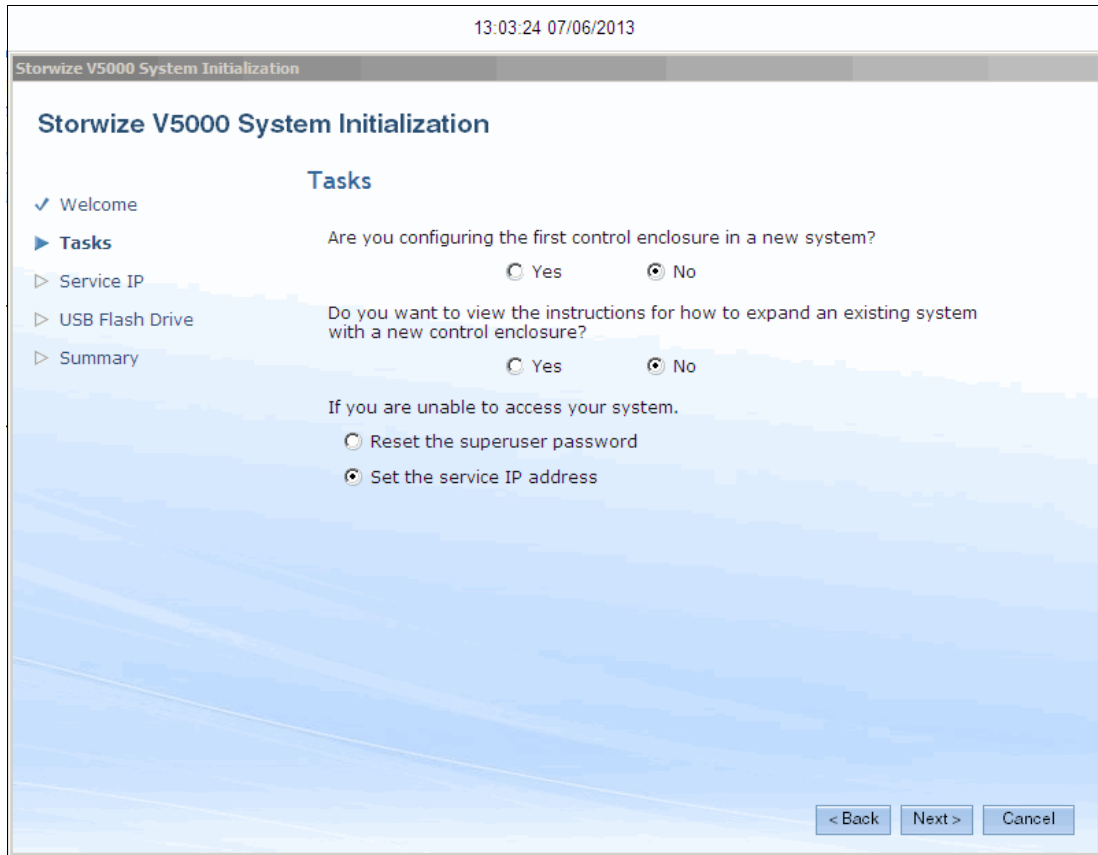


Figure 2-11 Inittool task options

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

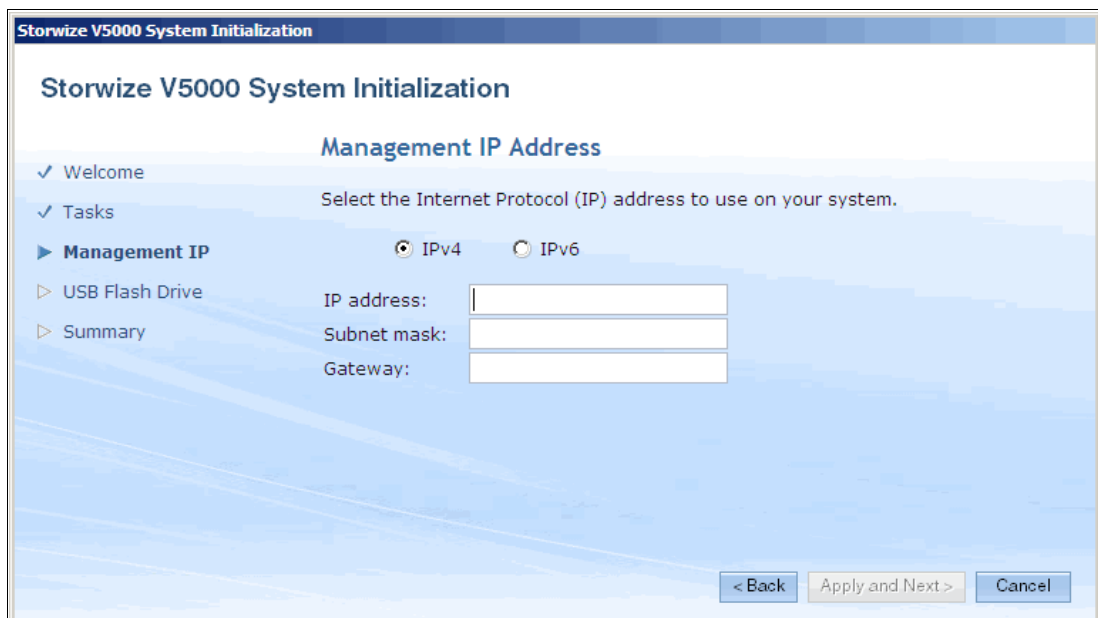


Figure 2-12 System Initialization: Management IP

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

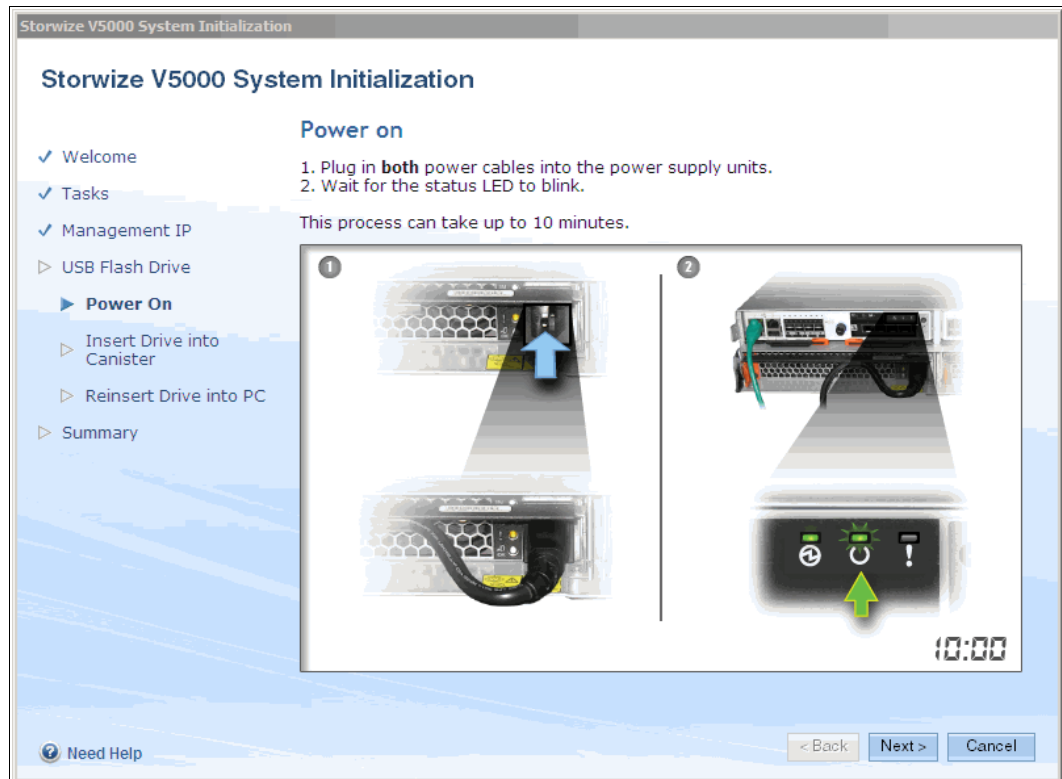


Figure 2-13 Initialization application: V5000 Power up

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

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

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

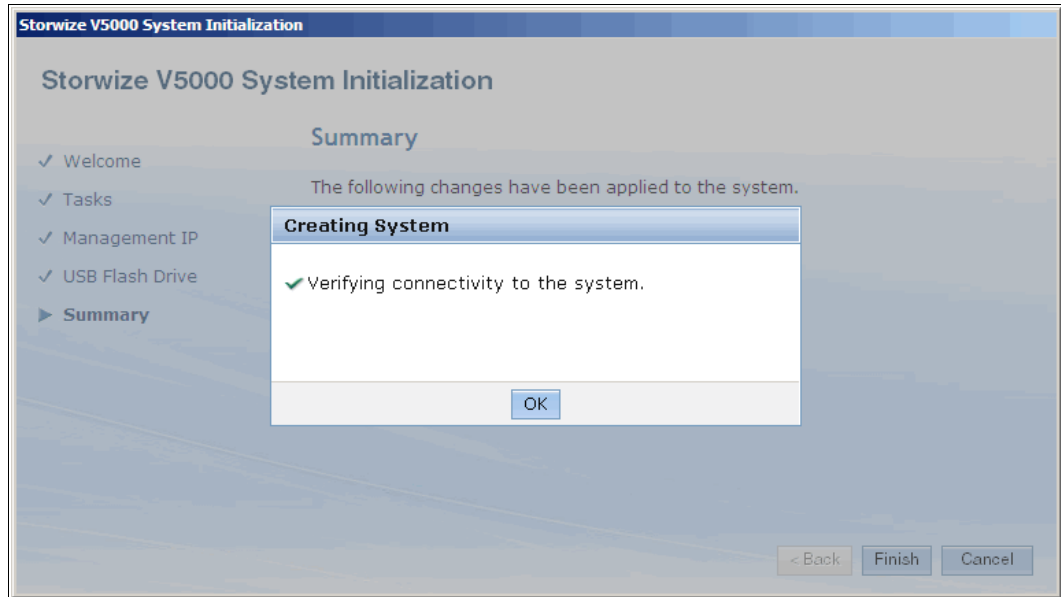


Figure 2-14 Verify system connectivity

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




Figure 2-15 Initialization Summary

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

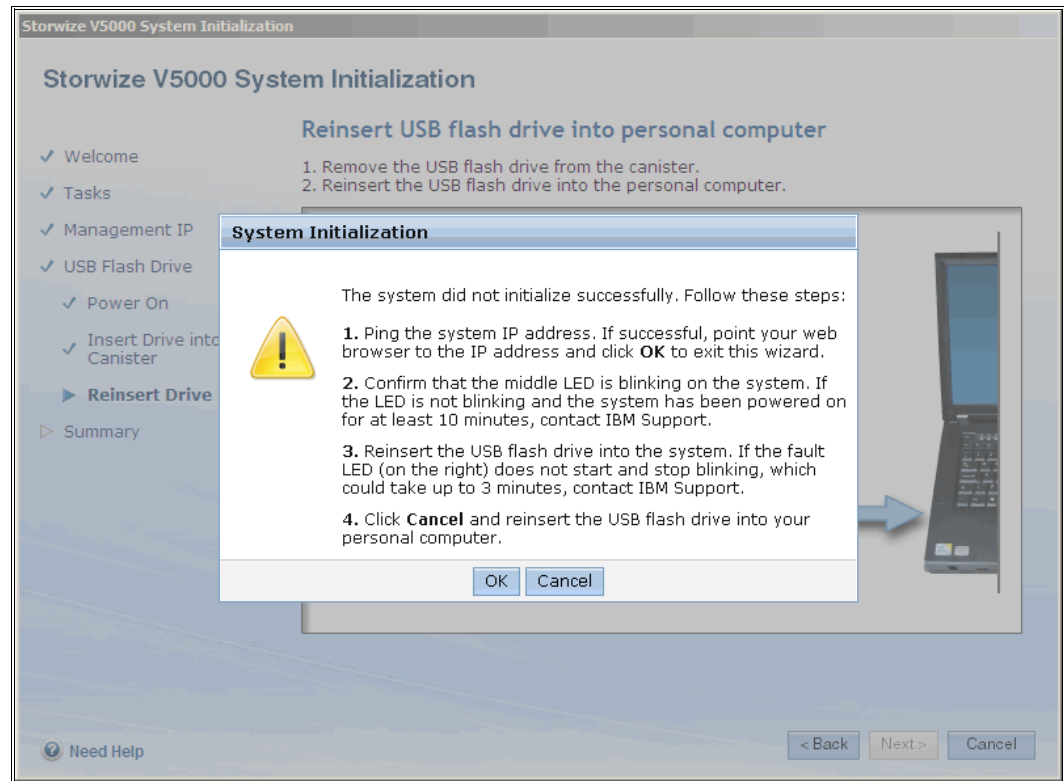


Figure 2-16 Initialization Failure

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

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

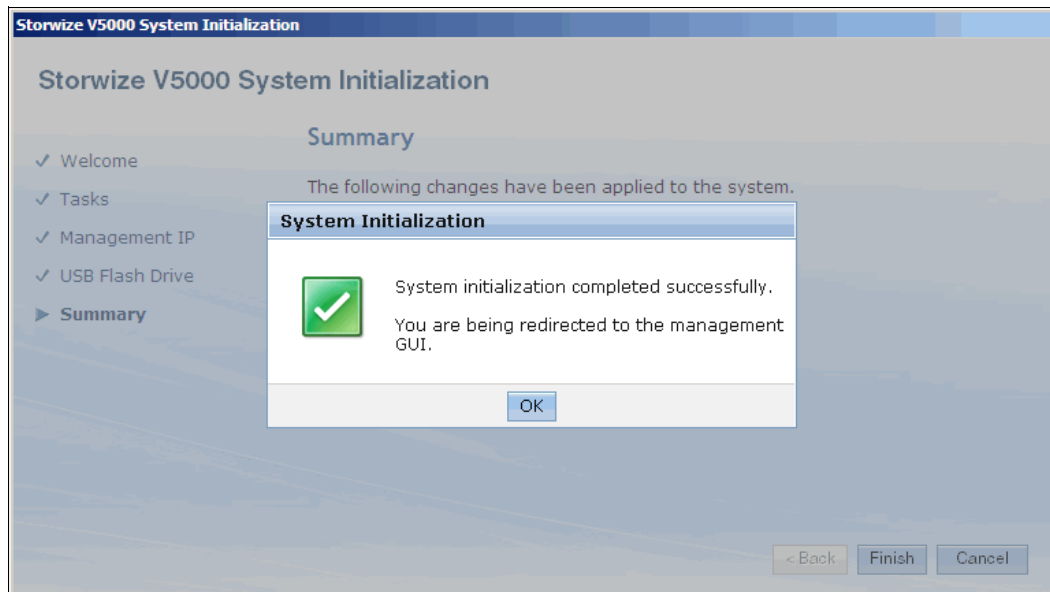


Figure 2-17 System Initialization complete

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

2.10 Initial configuration

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

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

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

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



Figure 2-18 Setup wizard: Login

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

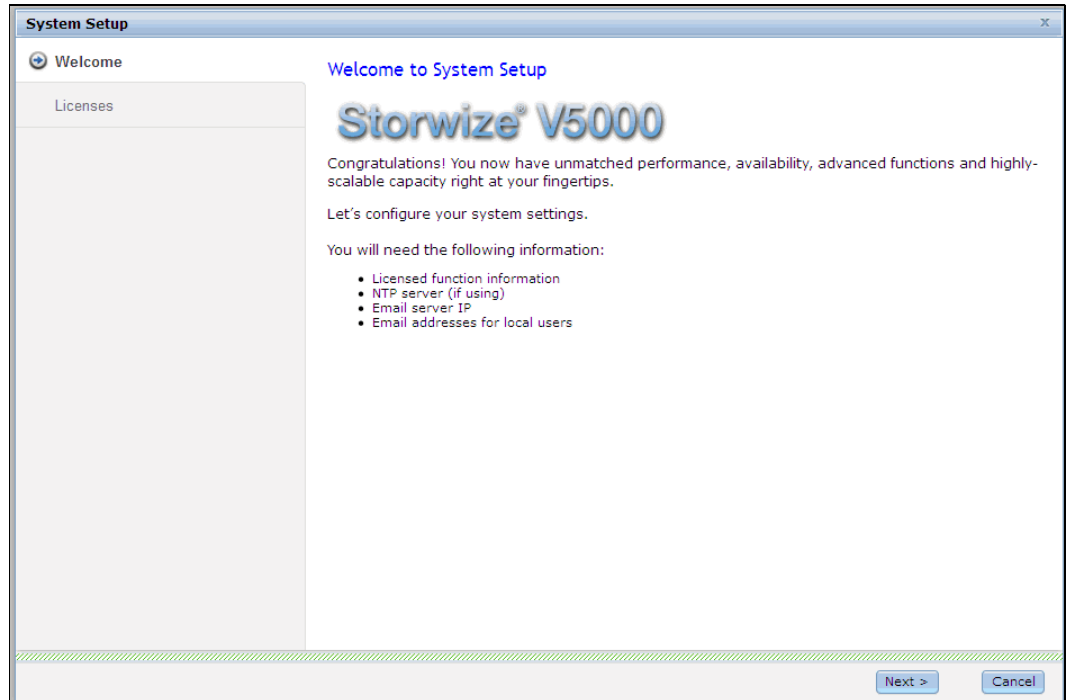


Figure 2-19 Welcome window

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

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

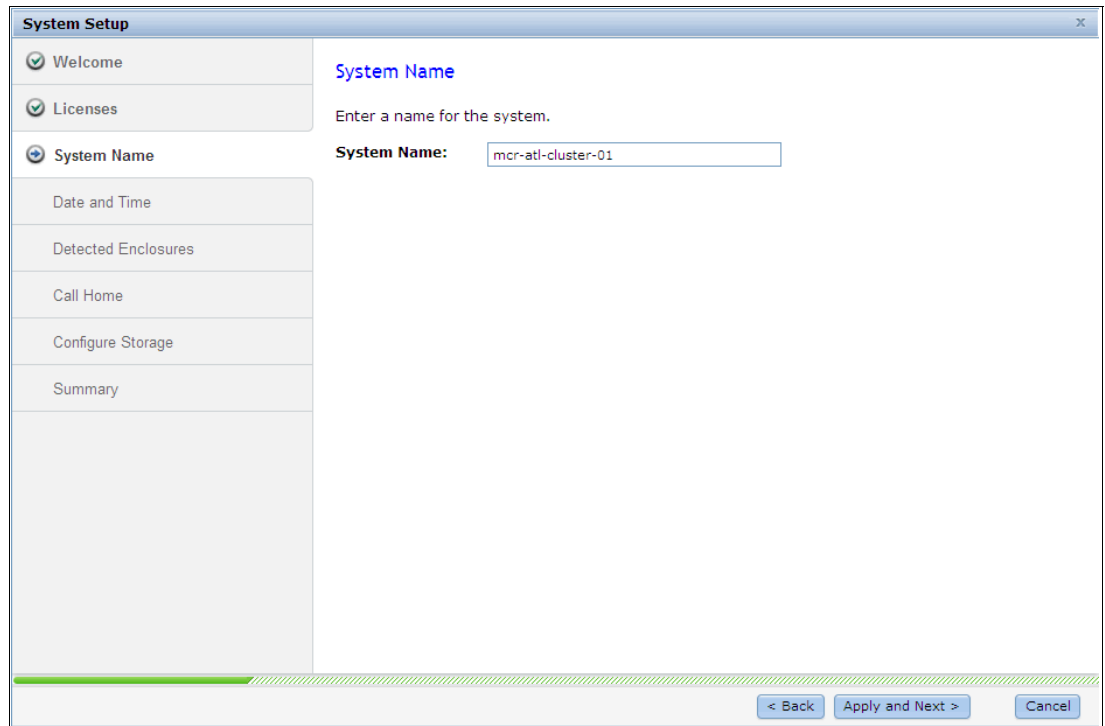


Figure 2-20 Setup wizard: Insert system name

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

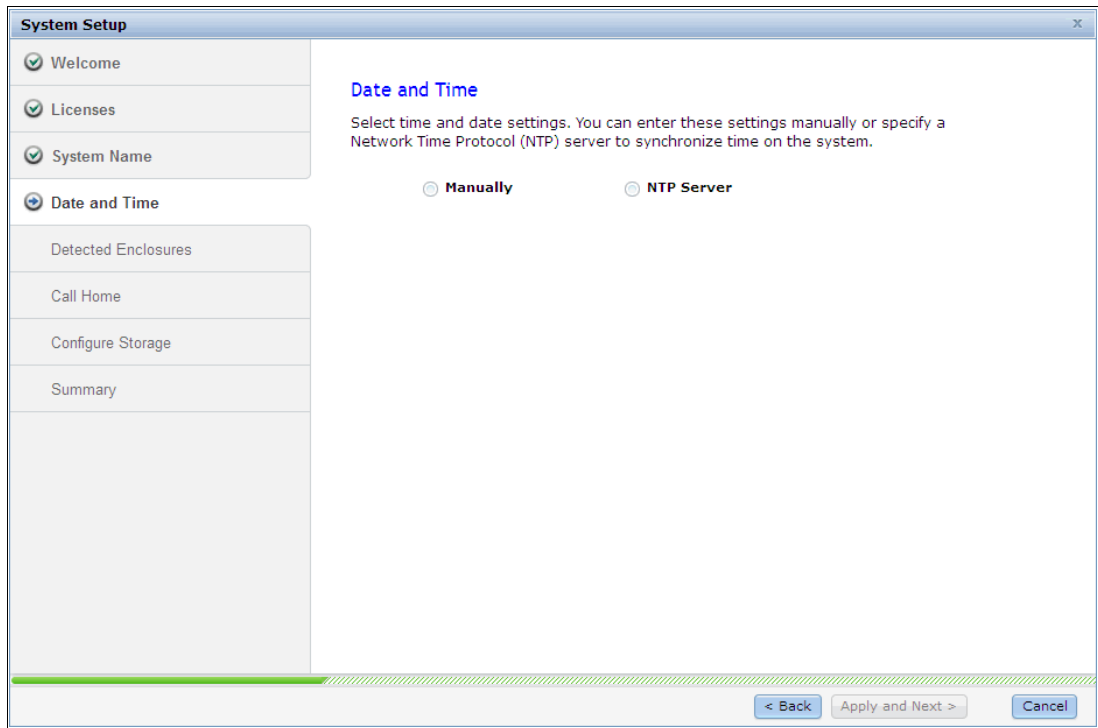


Figure 2-21 Setup wizard: Date and time

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

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

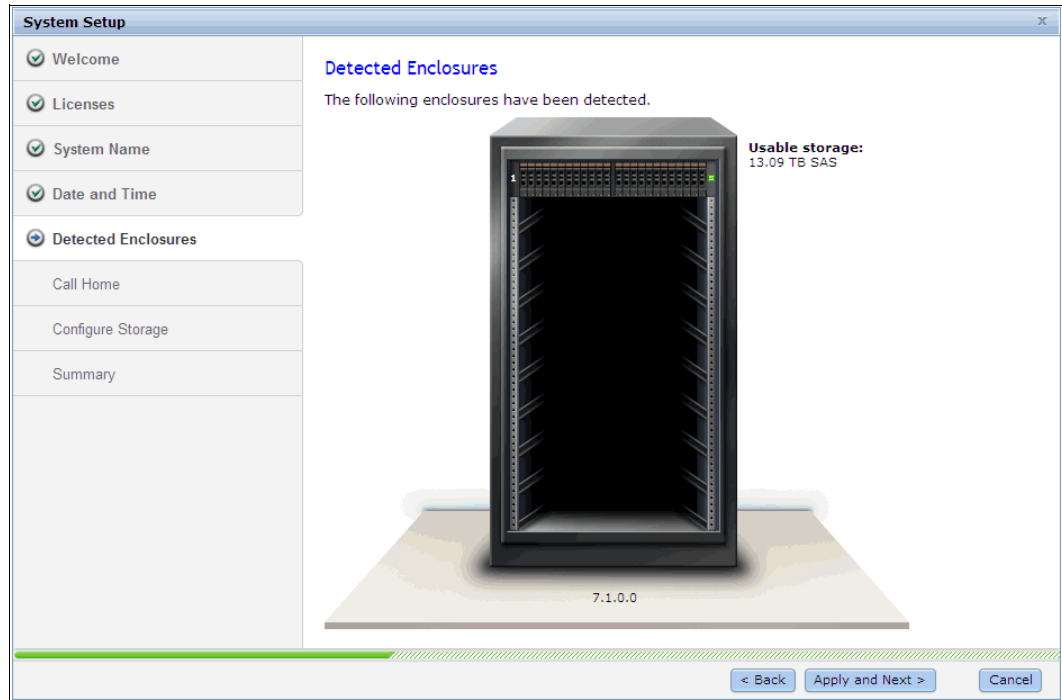


Figure 2-22 Setup wizard: Verify the detected hardware

Click **Apply and Next**.

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

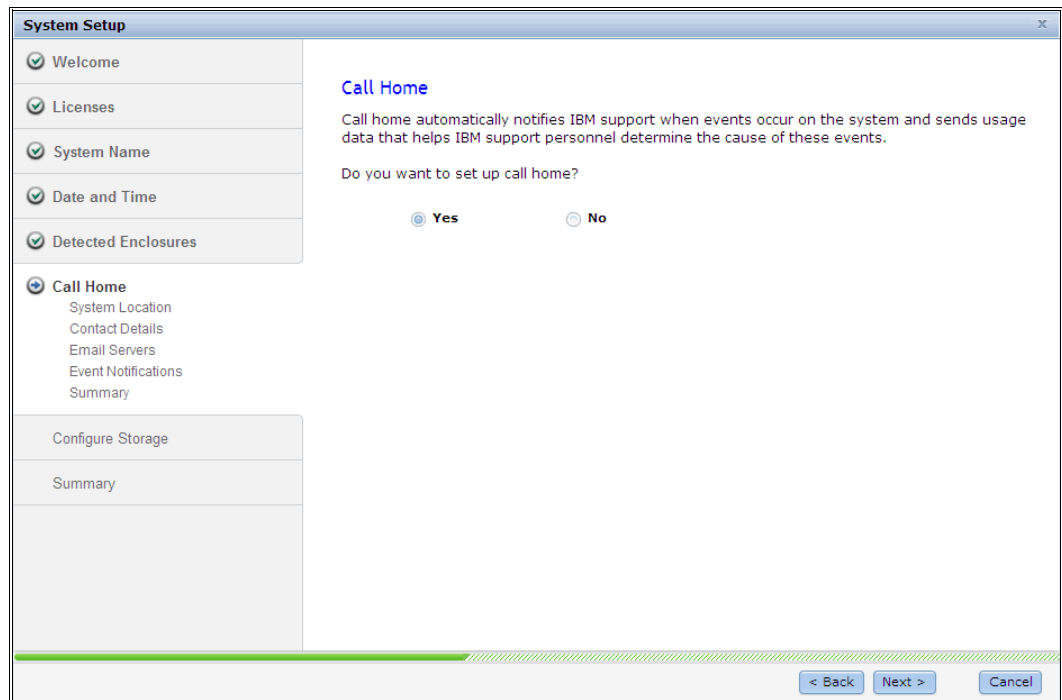


Figure 2-23 Call Home setup

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

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

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

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

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

The screenshot shows a window titled "System Setup" with a sidebar on the left containing a list of configuration steps: Welcome, Licenses, System Name, Date and Time, Detected Enclosures, Call Home (expanded to show System Location, Contact Details, Email Servers, Event Notifications, and Summary), Configure Storage, and Summary. The main area is titled "System Location" and contains the instruction "Enter the company name and address to ship parts." Below this are six input fields: * Company name, * Street address, * City, * State or province (with "XX" as a placeholder), * Postal code, and * Country or region (with a dropdown arrow). At the bottom right, there are three buttons: "< Back", "Next >", and "Cancel".

Figure 2-24 Location details

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

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

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

The screenshot shows a web-based configuration window titled "System Setup". On the left is a navigation pane with the following items: Welcome, Licenses, System Name, Date and Time, Detected Enclosures, Call Home (expanded), System Location, Contact Details (selected), Email Servers, Event Notifications, Summary, Configure Storage, and Summary. The main content area is titled "Contact Details" and includes the instruction: "Enter the name and contact information for the person in your organization that support personnel can contact to help resolve problems on the system." Below this are several input fields: "*Contact name:" with the value "Sid James"; "*Email address:" with the value "Sjames@ibm.com"; "*Telephone (primary):" with the value "1234456789"; "Telephone (alternate):" which is empty; and "Comment:" which is empty. At the bottom right of the window are three buttons: "< Back", "Apply and Next >", and "Cancel".

Figure 2-25 Contact details

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

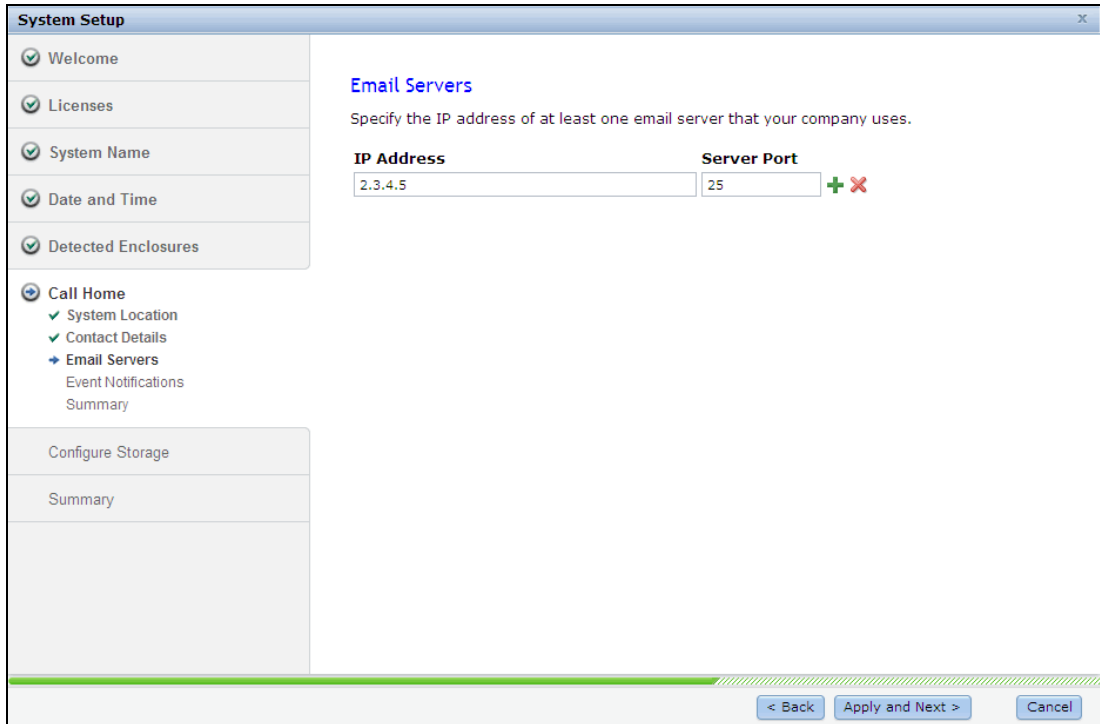


Figure 2-26 Email server details

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

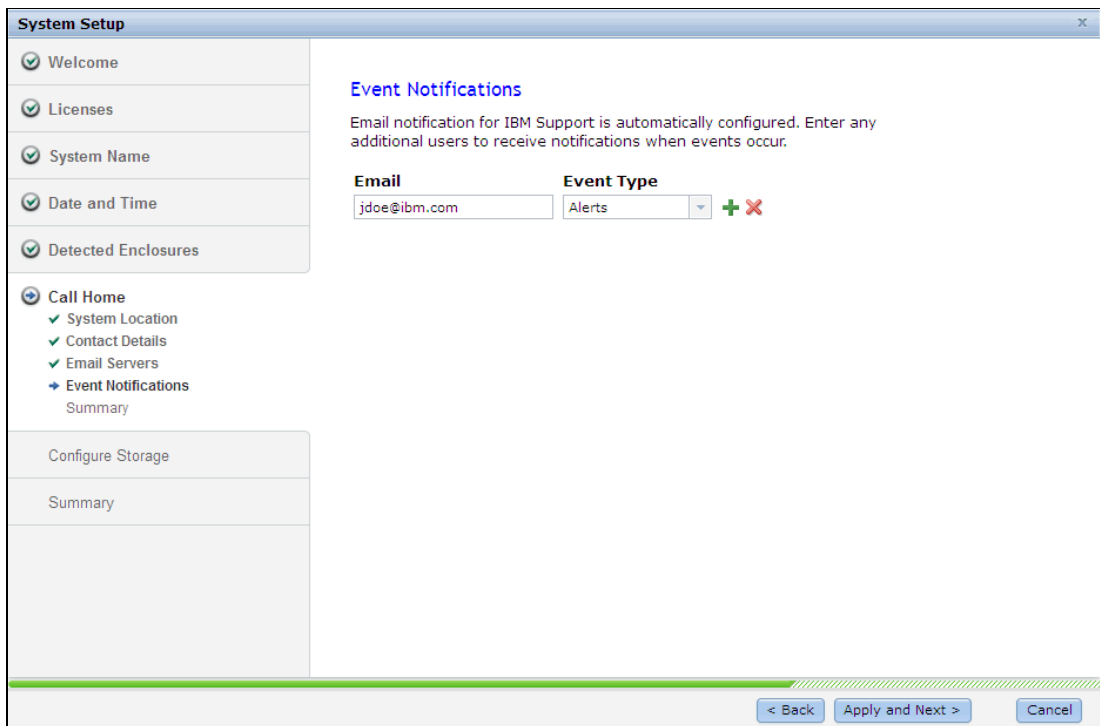


Figure 2-27 Event notification

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

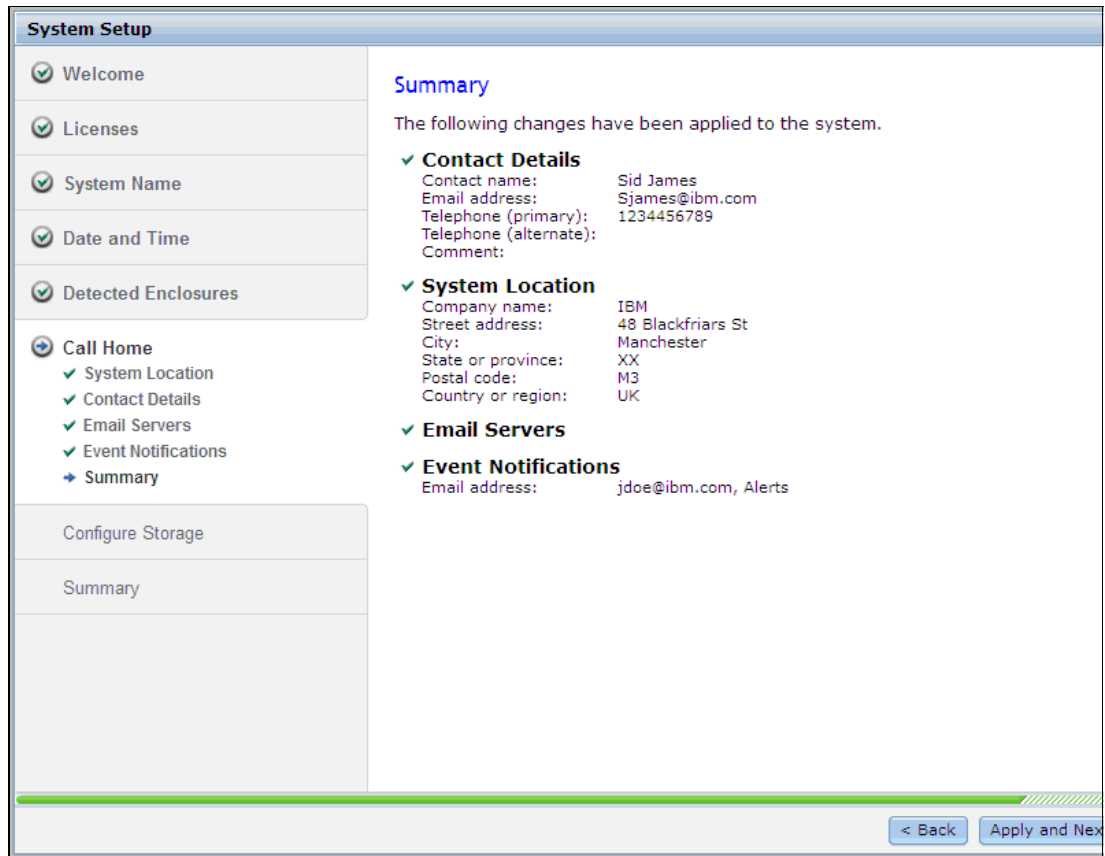


Figure 2-28 Call Home summary

Click **Apply and Next**.

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

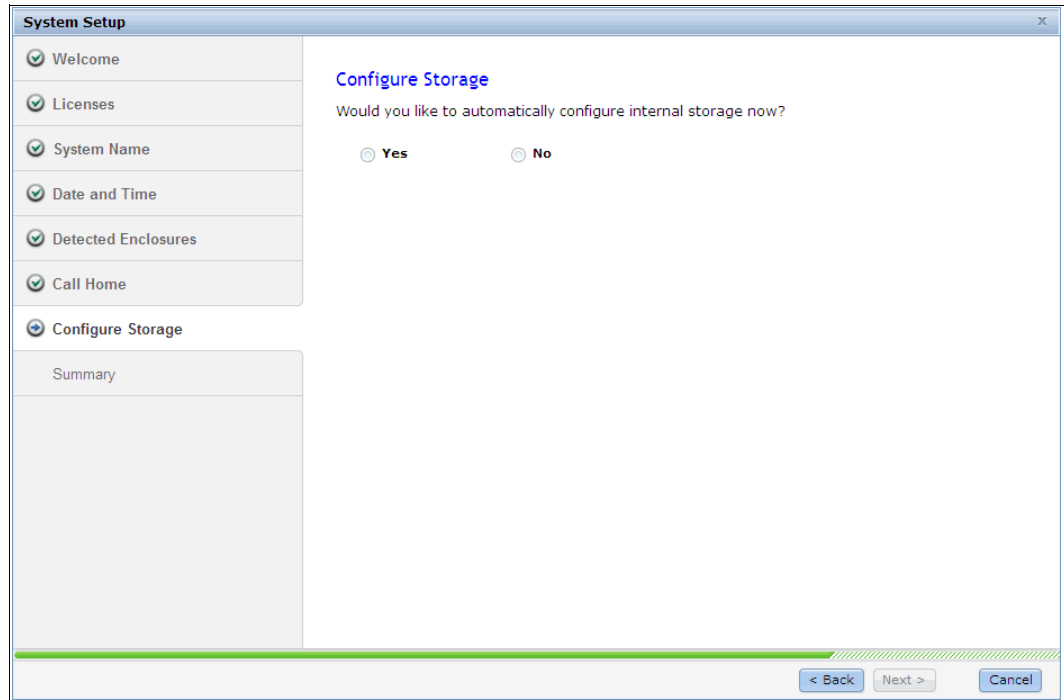


Figure 2-29 Configure Storage option

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

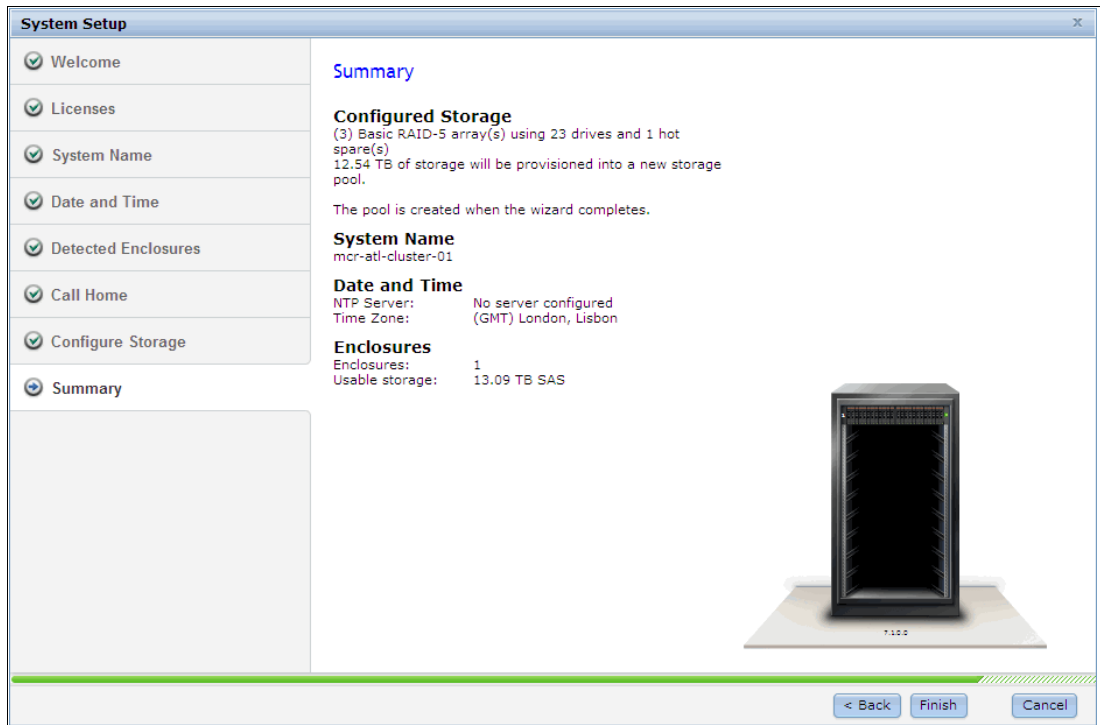


Figure 2-30 Storage Configuration Summary

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

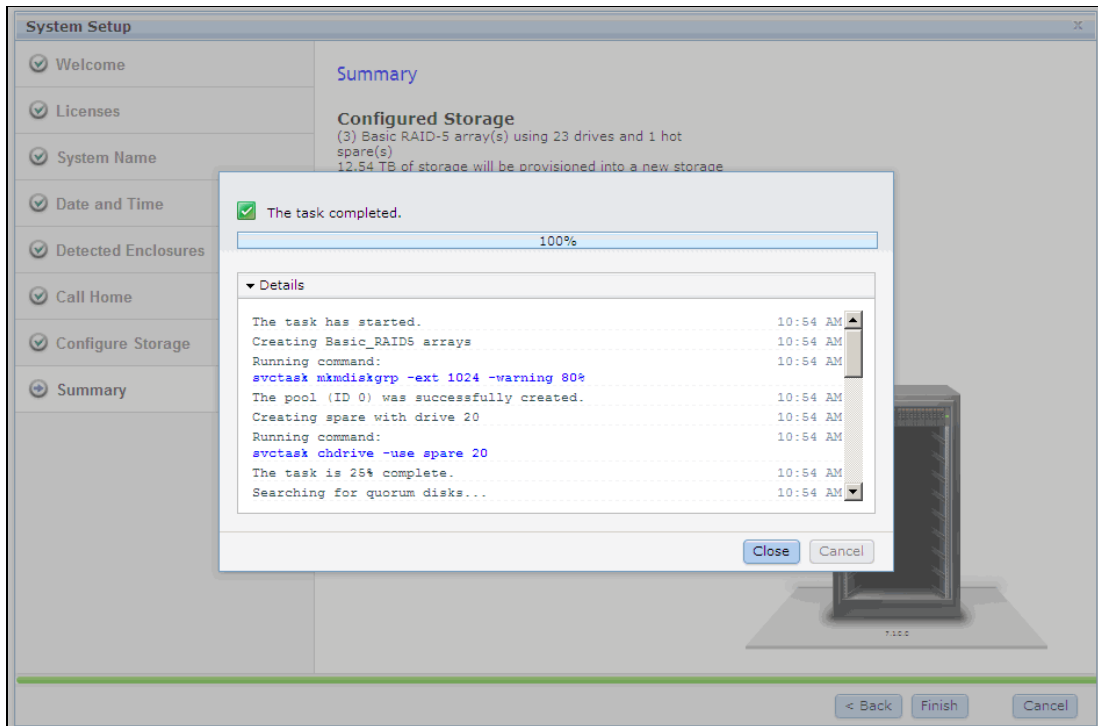


Figure 2-31 Storage array creation

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

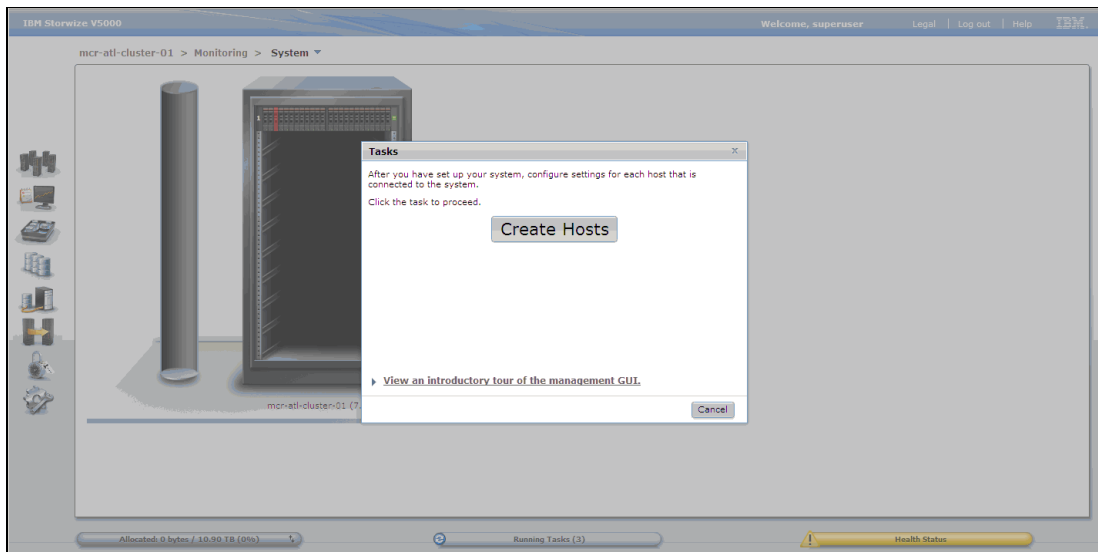


Figure 2-32 Create hosts

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

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

2.10.1 Adding Enclosures after initial configuration

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

Adding a control enclosure

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

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

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

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

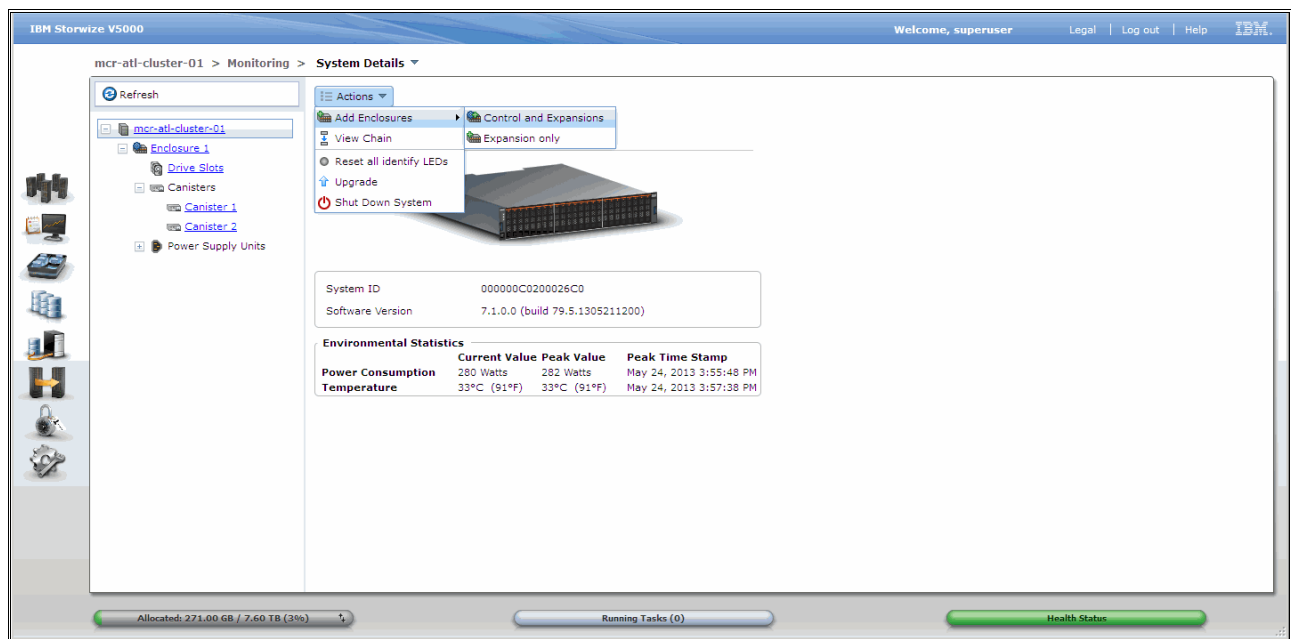


Figure 2-33 Option to add a control enclosure

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

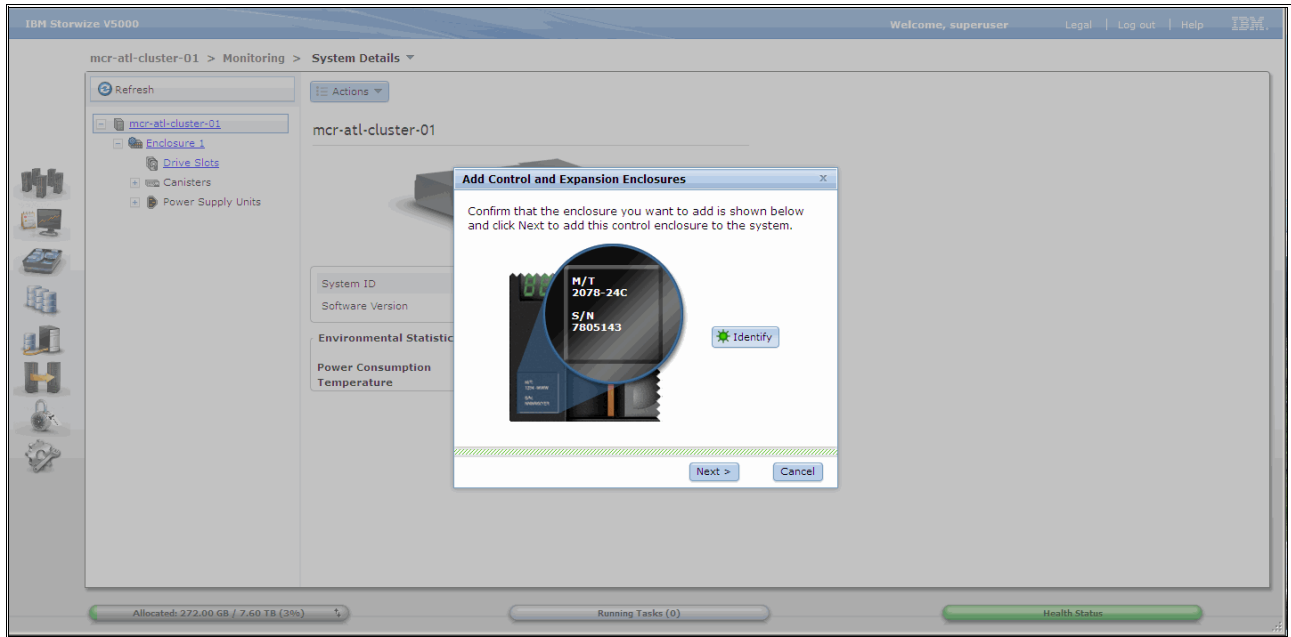


Figure 2-34 New control enclosure identification

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

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

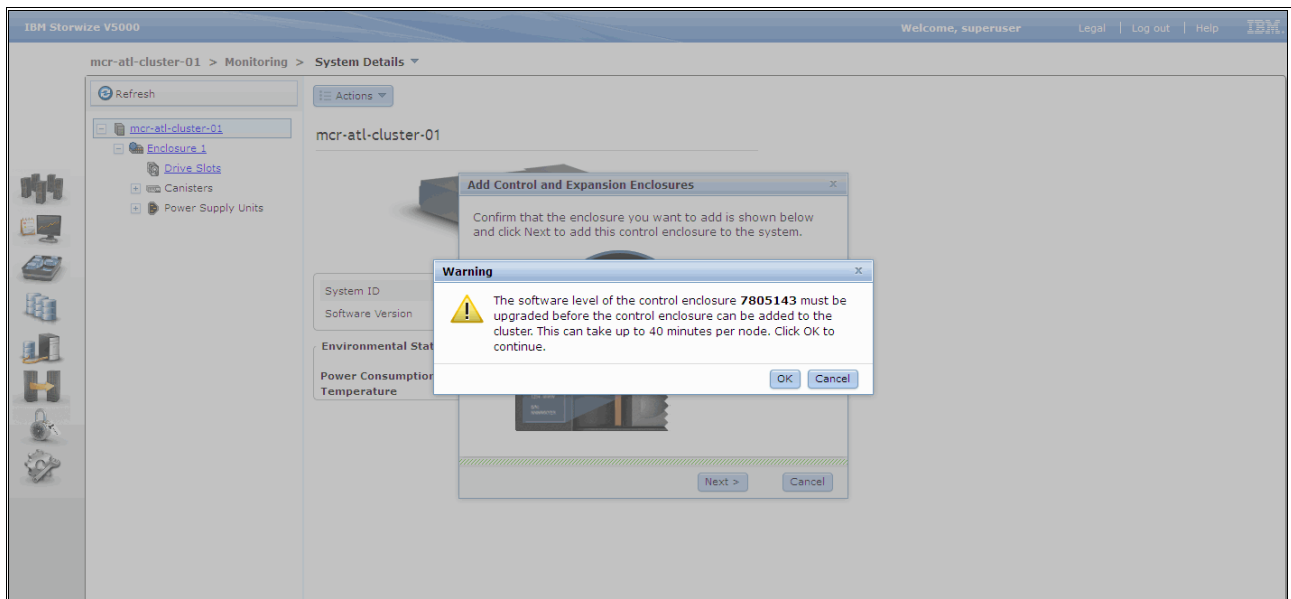


Figure 2-35 New control enclosure software upgrade warning

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

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

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

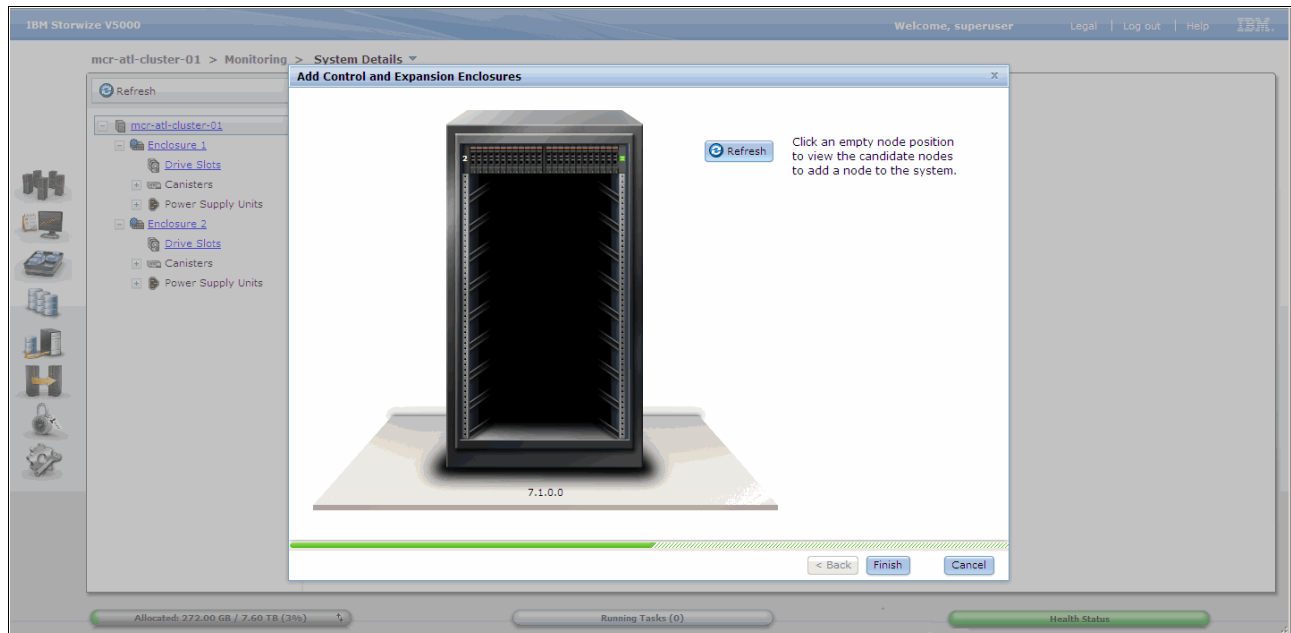


Figure 2-36 Add enclosure complete

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

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

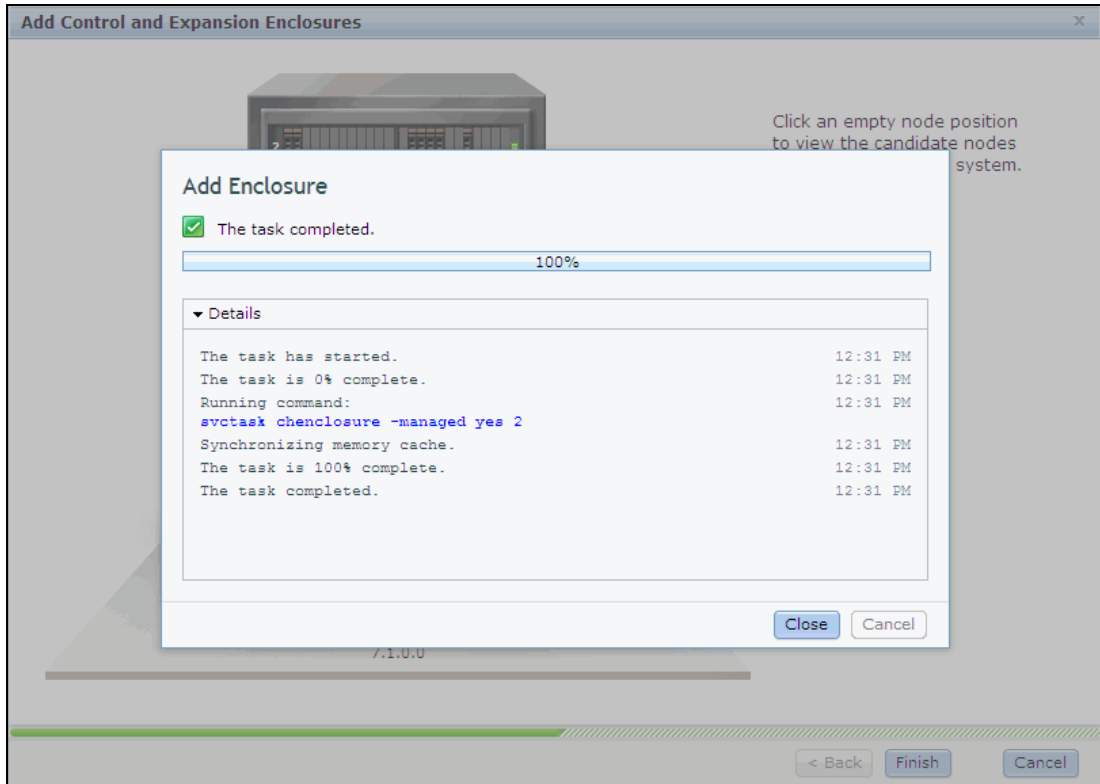


Figure 2-37 Add control enclosure task completion

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

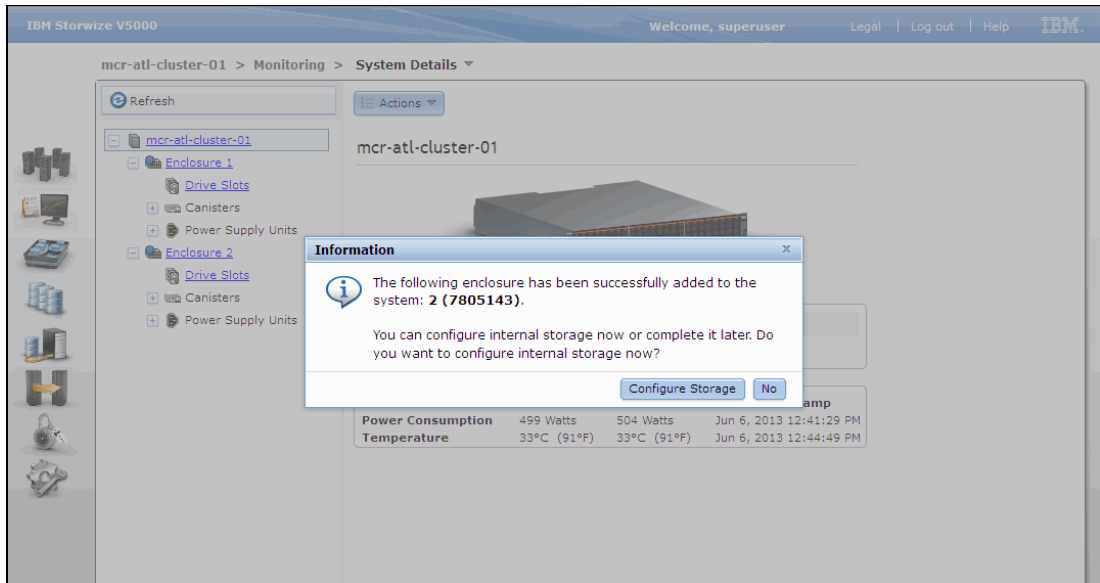


Figure 2-38 New storage configuration prompt

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

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

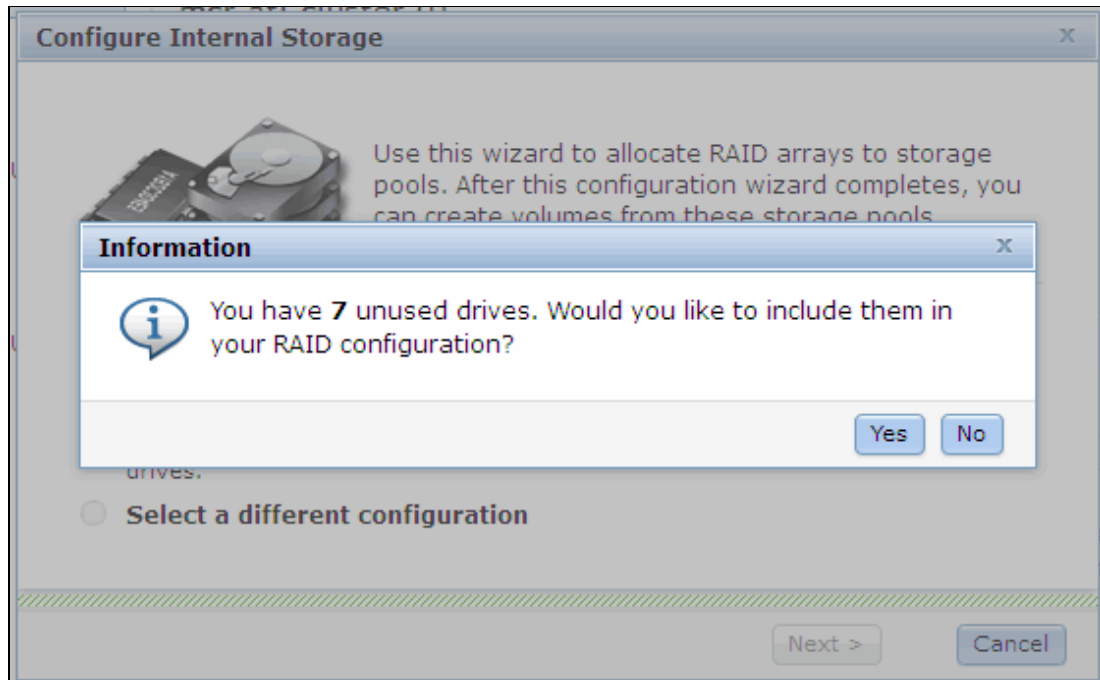


Figure 2-39 Configure new enclosure storage

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

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

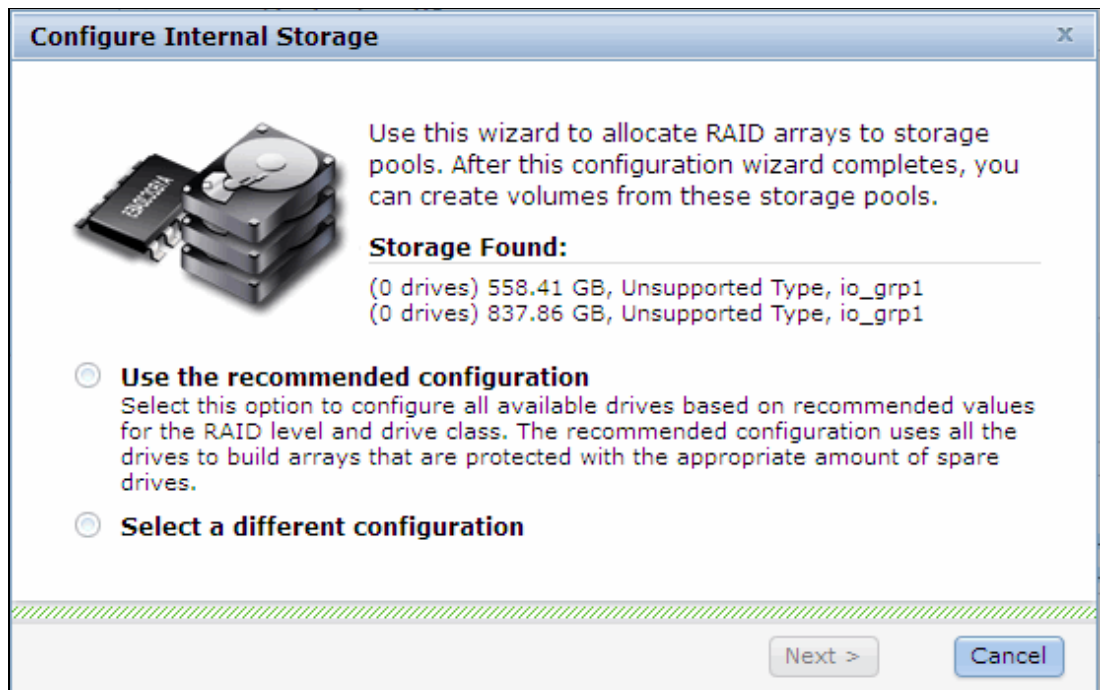


Figure 2-40 Configure new Internal Storage

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

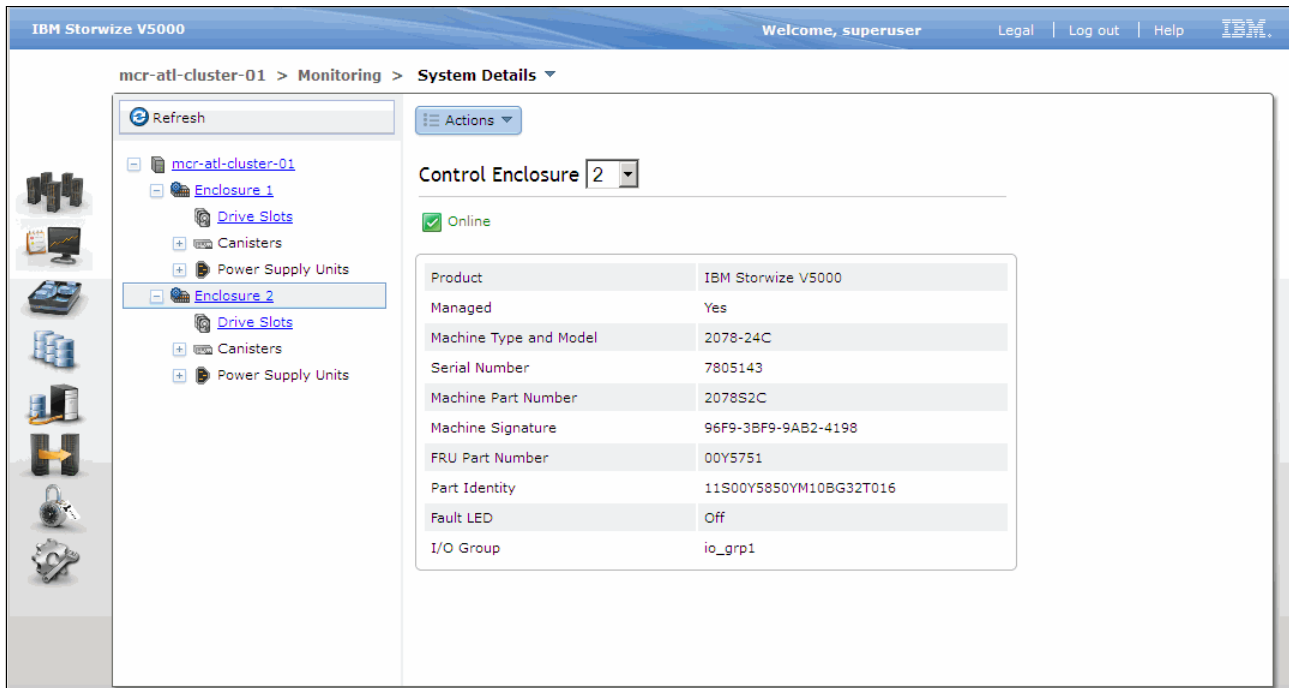


Figure 2-41 New enclosure that is shown as part of existing cluster

Adding a new expansion enclosure

Complete the following steps to add a new expansion controller:

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

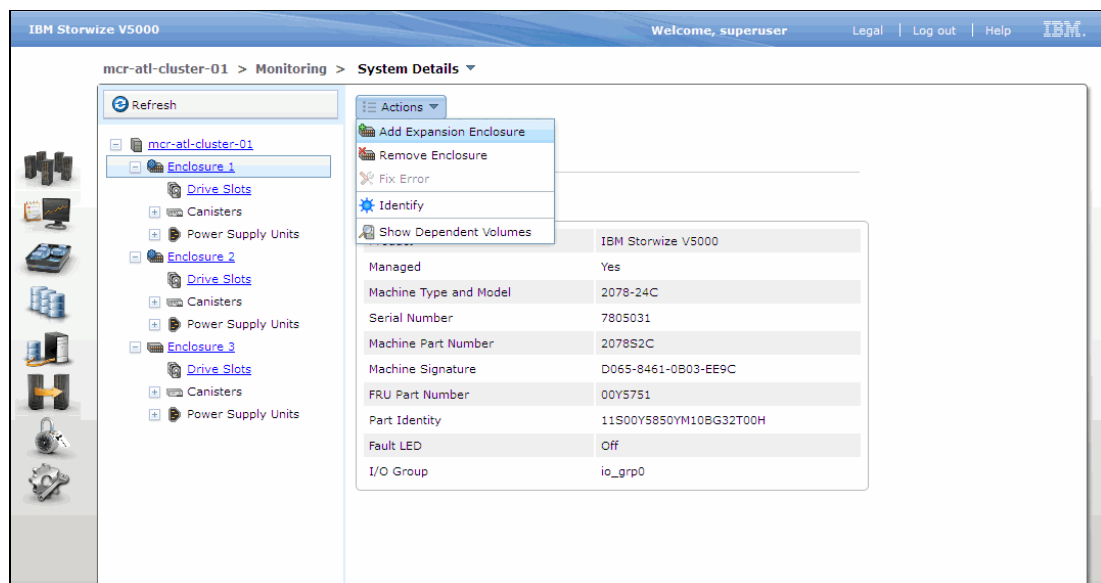


Figure 2-42 Adding an expansion enclosure

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

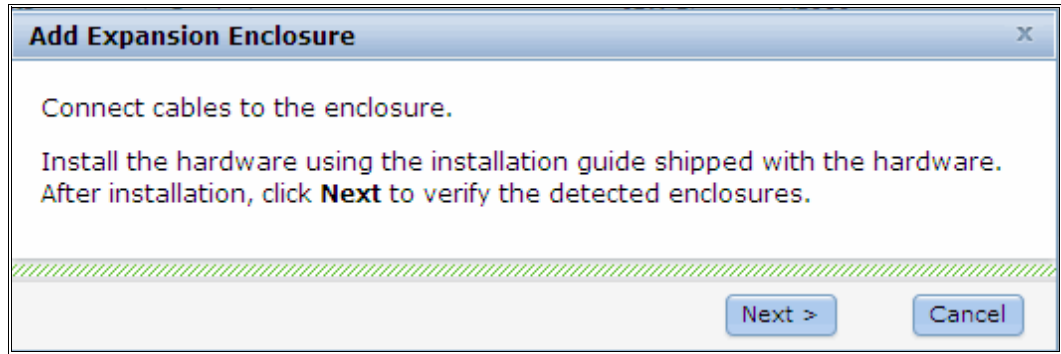


Figure 2-43 Expansion enclosure cable check

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

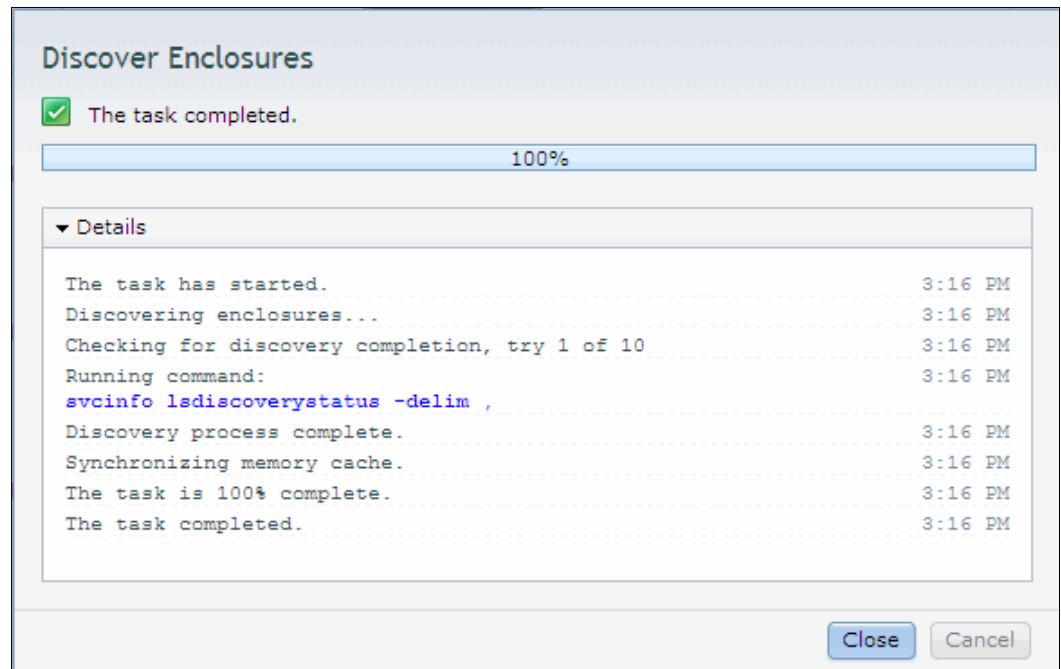


Figure 2-44 New enclosure discovery task

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

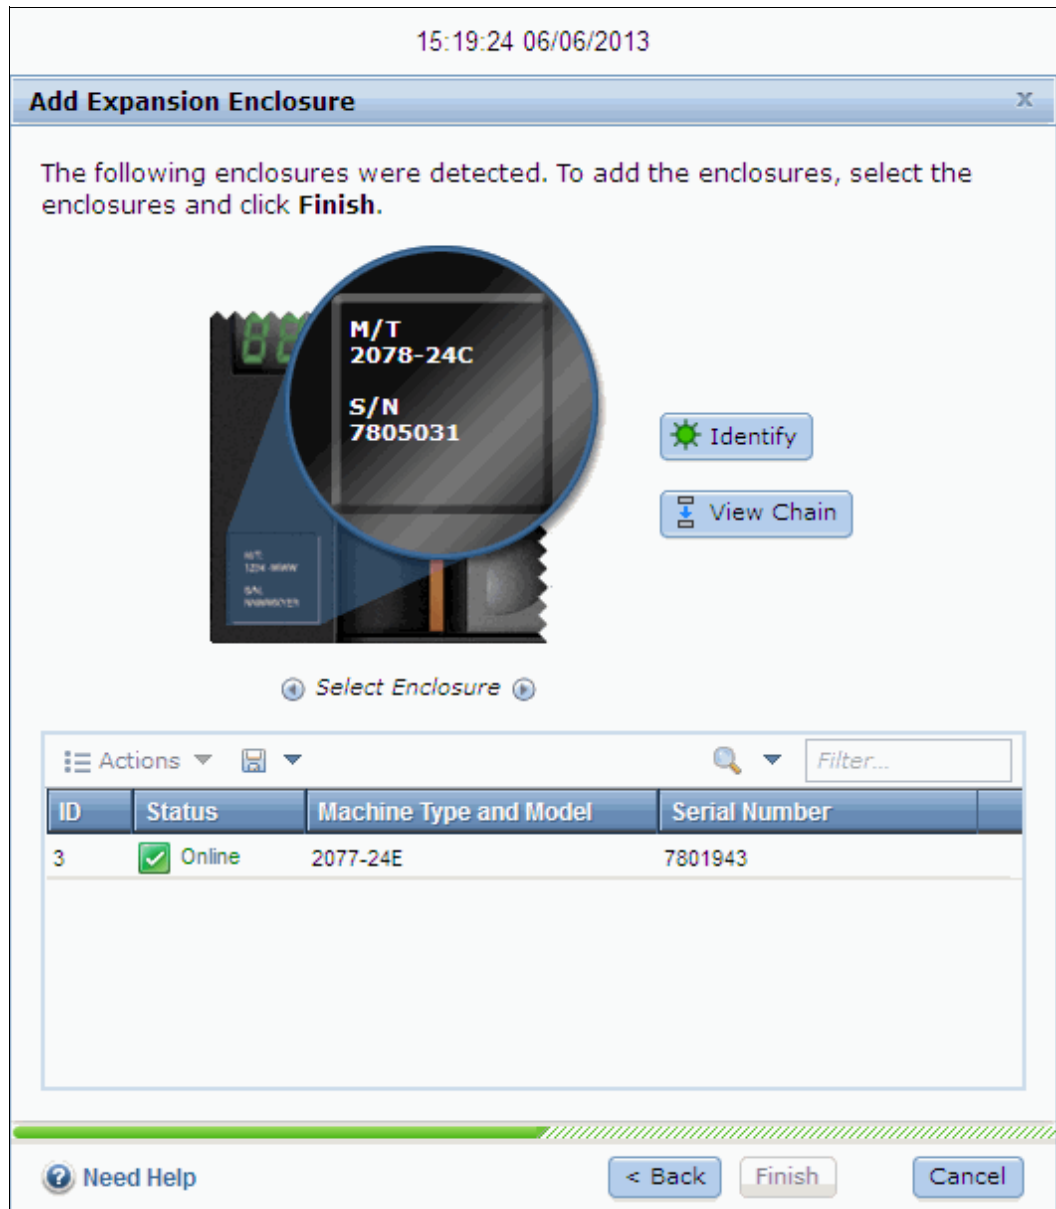


Figure 2-45 New hardware to be added

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

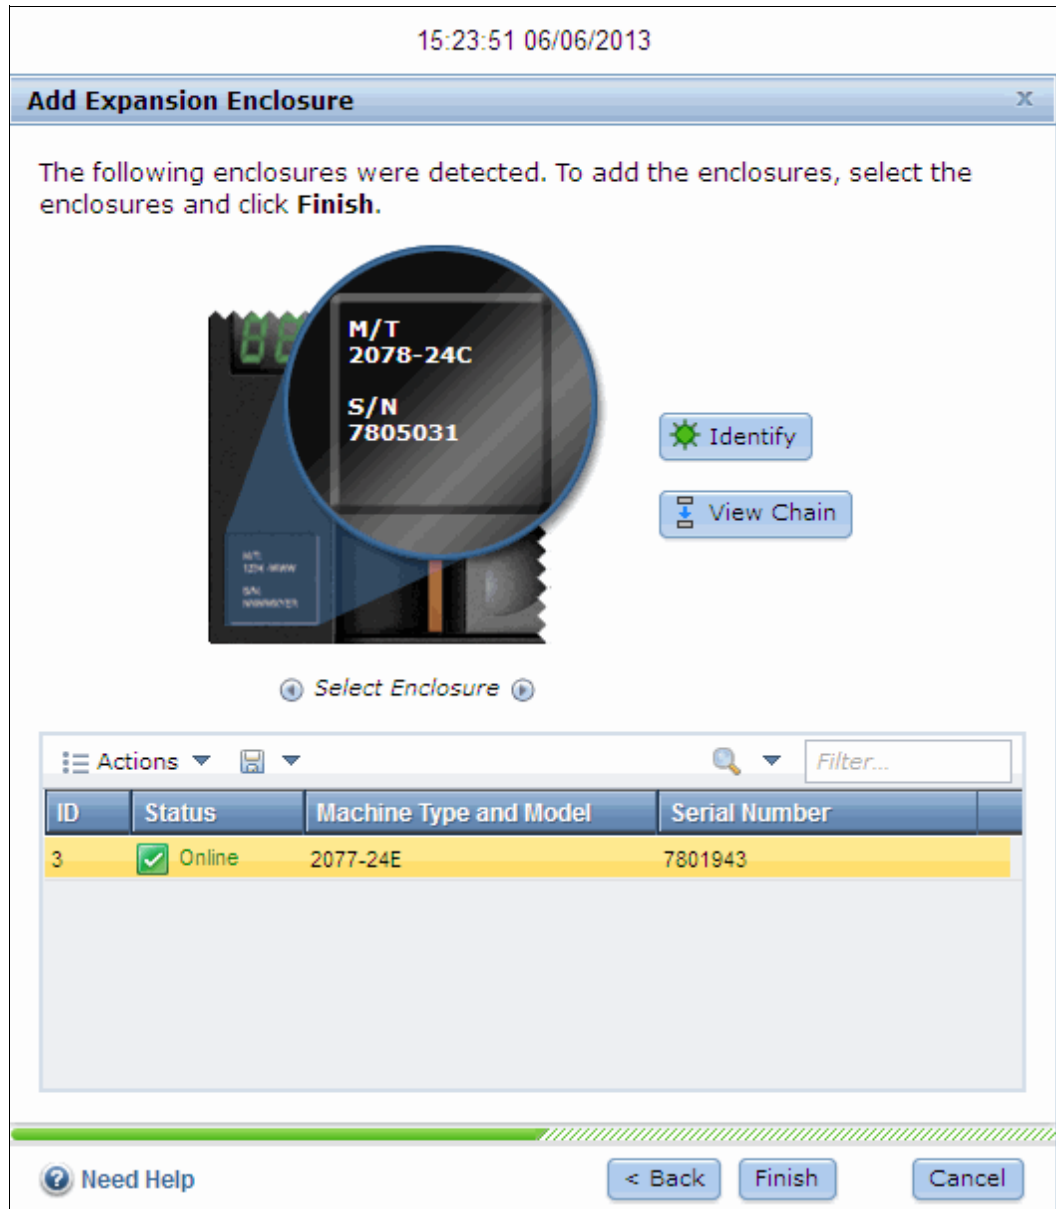


Figure 2-46 Selecting new hardware to be added

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

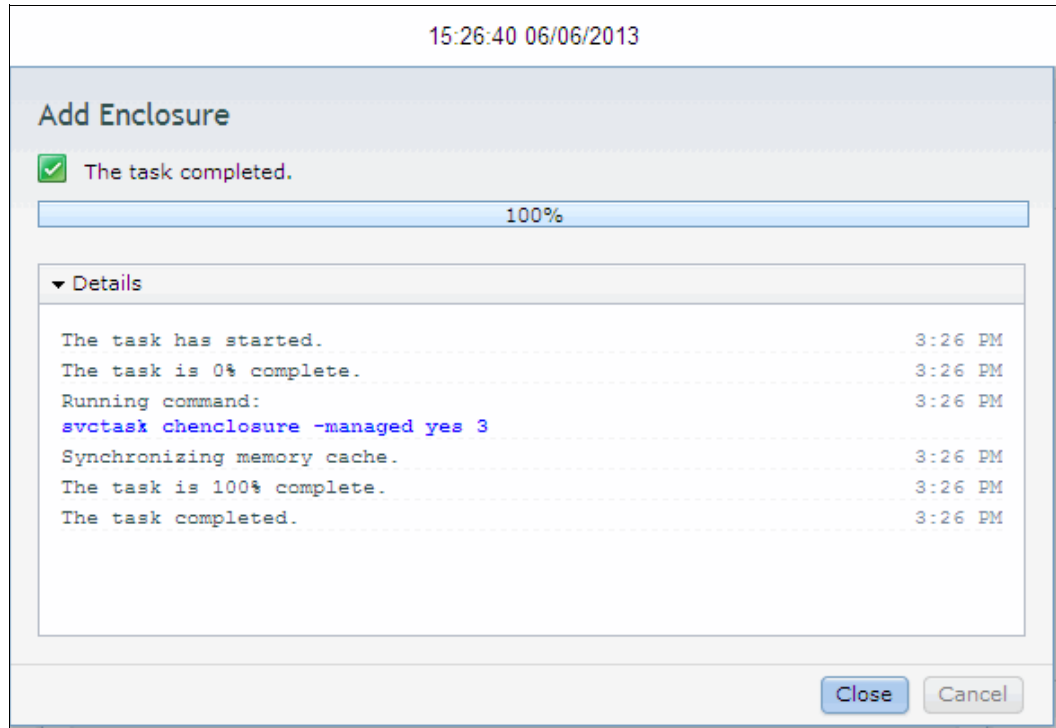


Figure 2-47 Add new enclosure task completion

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

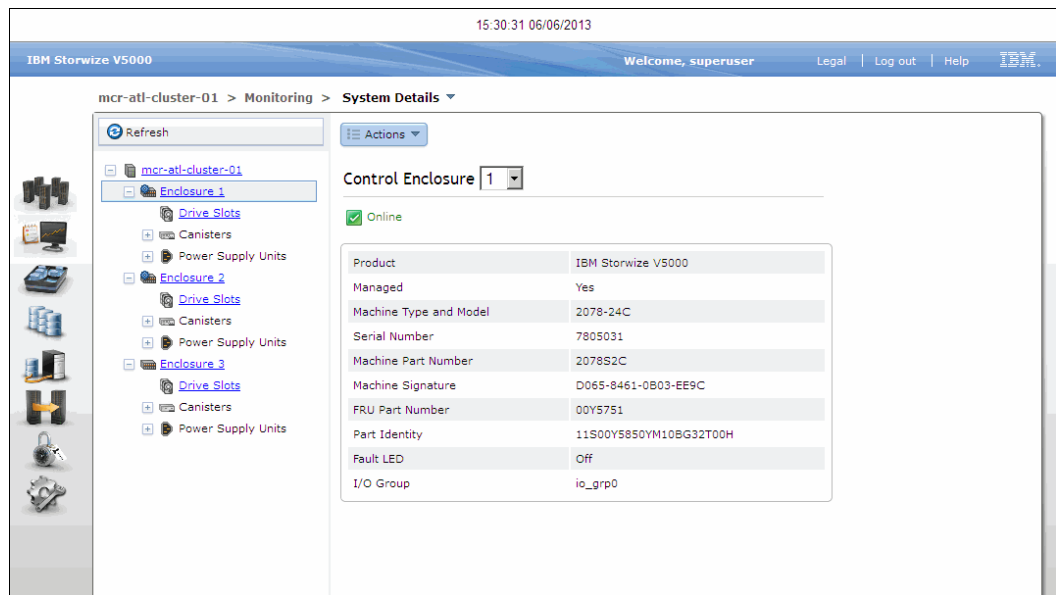


Figure 2-48 New expansion enclosure as part of the cluster

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

2.10.2 Configuring Call Home, email alert, and inventory

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

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

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

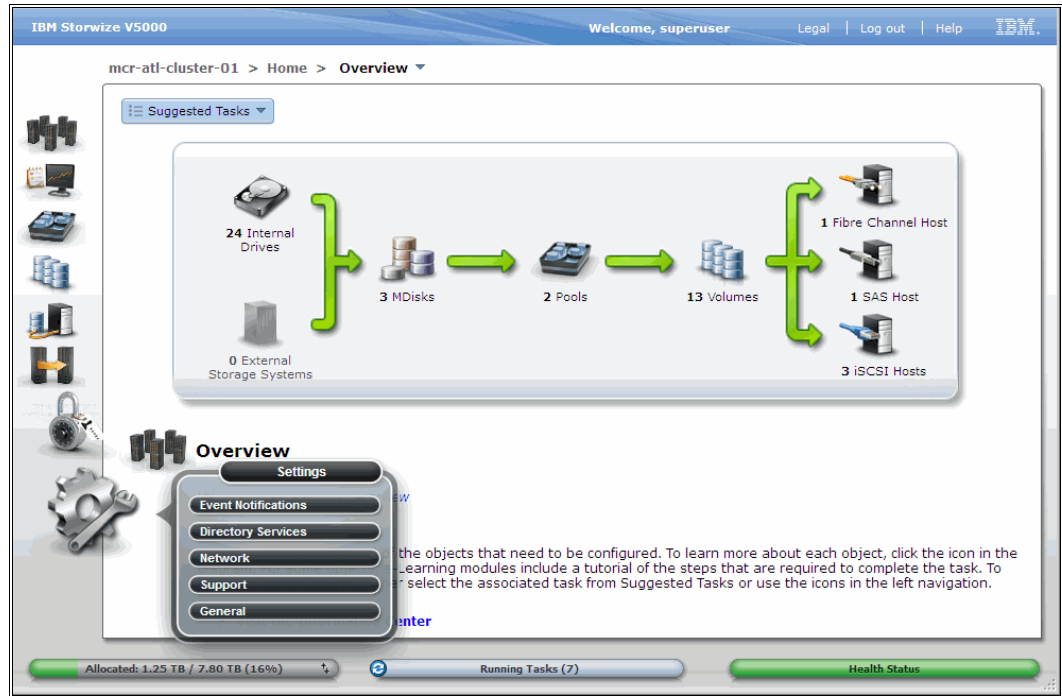


Figure 2-49 Enabling Call Home

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

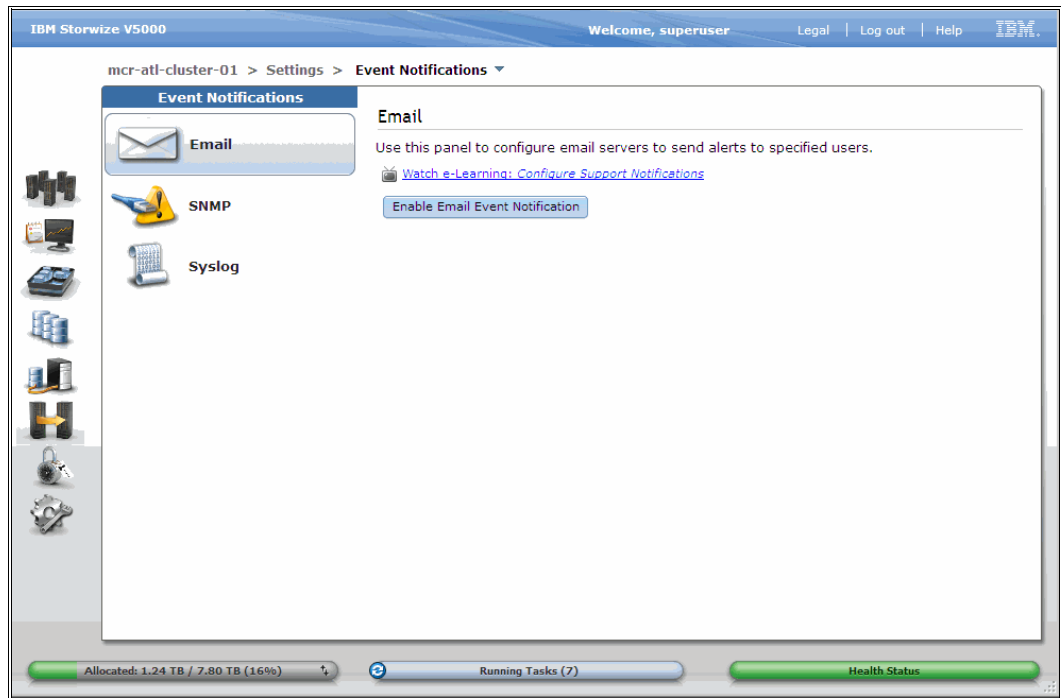


Figure 2-50 Selecting Event Notification

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

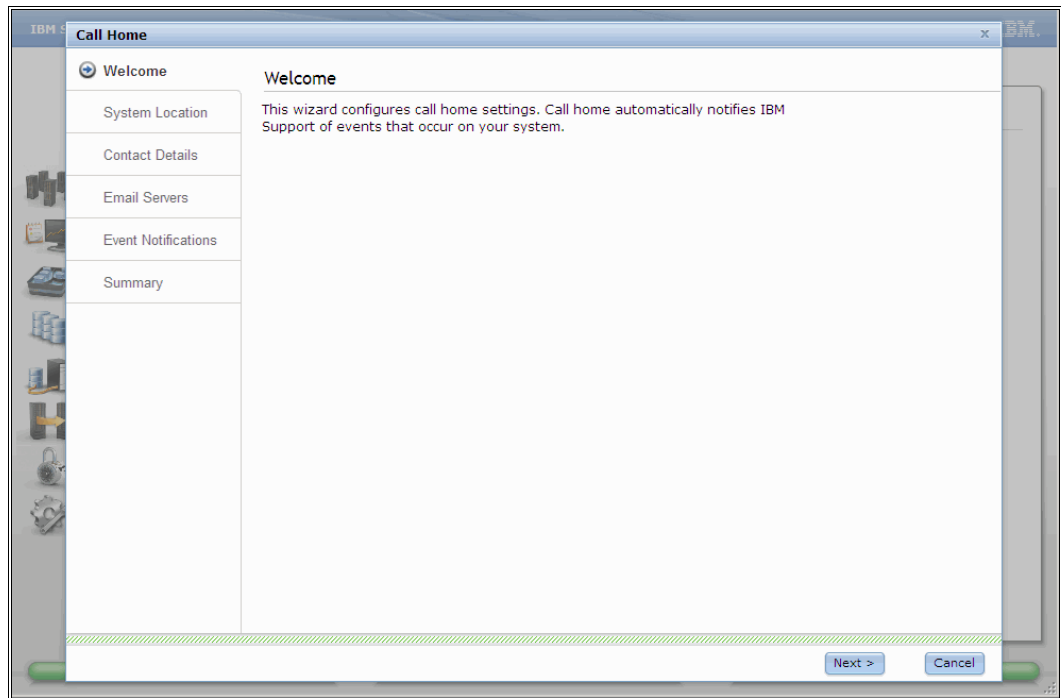


Figure 2-51 Call Home wizard

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

2.10.3 Service Assistant tool

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

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

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

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

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

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

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

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

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

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

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

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

Example:

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

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

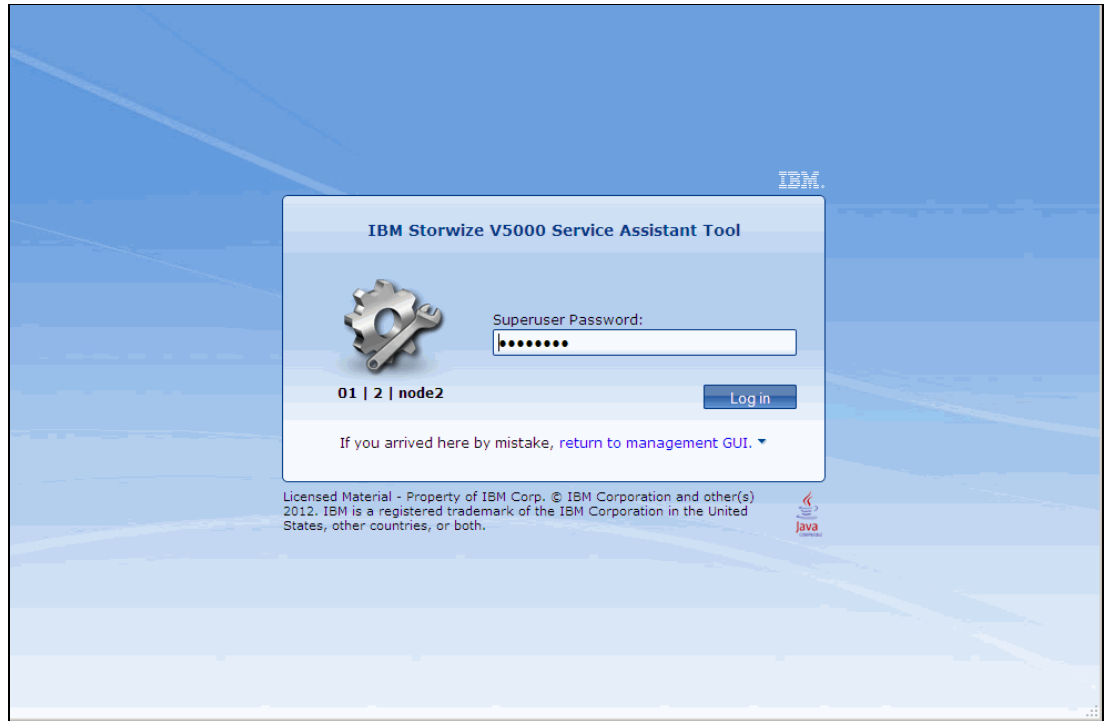


Figure 2-52 Service Assistant Login

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

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

IBM Storwize V5000 Service Assistant Tool Connected to: 01 | 2 | node2 Log out IBM

Current: 01 | 2 | node2
Status: Active
Identify

Home
You can view detailed status and error summary, and manage service actions for the current node. The current node is the node on which service-related actions are performed. The connected node displays the service assistant and provides the interface for working with other nodes on the system. To manage a different node, select a node from the following table.

Attention: Only perform service actions on nodes when directed by service procedures. If used inappropriately, service actions can cause a loss of access to data, or even data loss. If the node status is active, select Monitoring-->Events in the management GUI to fix any errors that are related to the active node.

Actions: Enter Service State GO

Change Node

Node Name	Node Status	Error	Panel	System	Relationship
<input checked="" type="radio"/> node2	Active		01-2	mcr-atl-cluster-01	Local
<input type="radio"/> node1	Active		01-1	mcr-atl-cluster-01	Partner
<input type="radio"/> node3	Active		02-1	mcr-atl-cluster-01	System
<input type="radio"/> node4	Active		02-2	mcr-atl-cluster-01	System

Refresh

Node Errors

Node Detail

Node	Hardware	Access	Location	Ports
Node ID:	8			
Node Name:	node2			
Node Status:	Active			
Part Identity:	11S00Y5764YM10BG32V064			
Node FRU:	00Y5860			
Configuration Node:	Yes			
Model:	AL1			
System:	mcr-atl-cluster-01			
System Software Build:	79.5.1305211200			
Software Version:	7.1.0.0			
Software Build:	79.5.1305211200			
Console IP:	9.174.161.119:443			
Has File Module Key:	No			

Figure 2-53 Service Assistant Home Window

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

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

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

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



Graphical user interface overview

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

This chapter includes the following topics:

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

3.1 Getting started

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

3.1.1 Supported browsers

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

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

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

3.1.2 Access the management GUI

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



Figure 3-1 IBM Storwize V5000 login panel

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

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

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

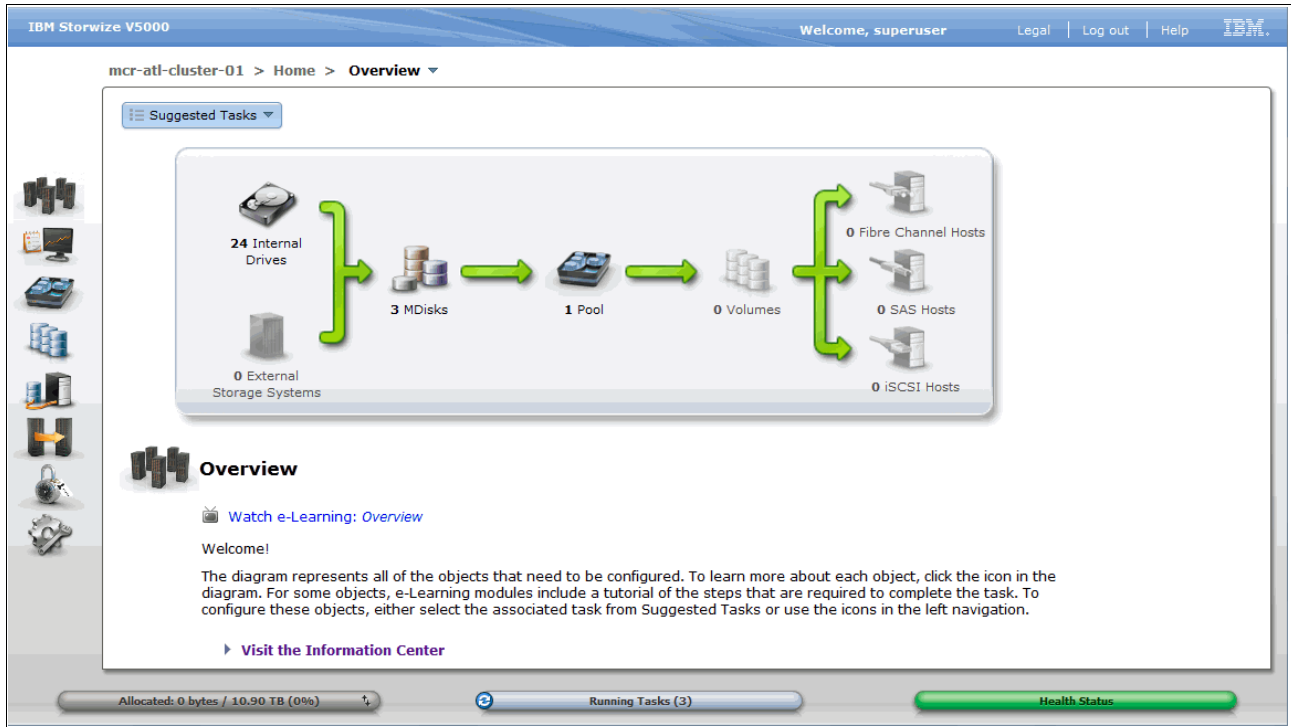


Figure 3-2 IBM Storwize V5000 overview panel

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

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

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

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

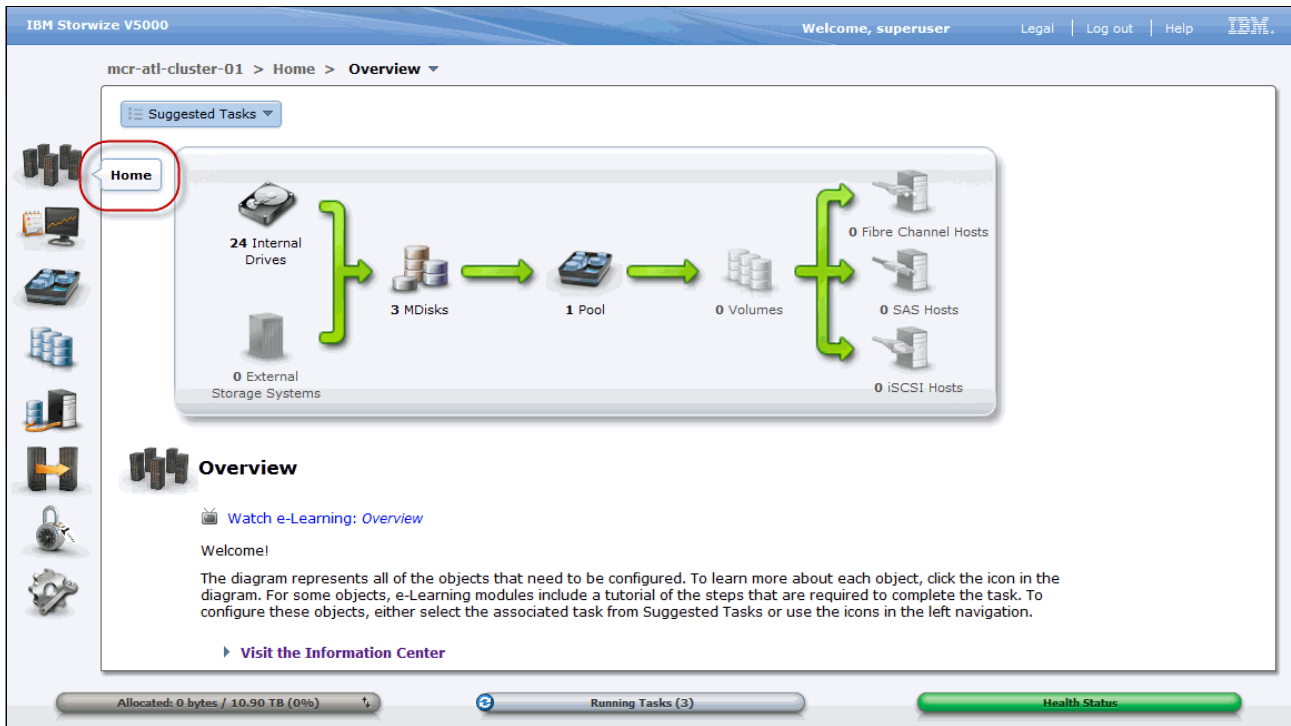


Figure 3-3 Management GUI low graphics mode

3.1.3 Overview panel layout

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

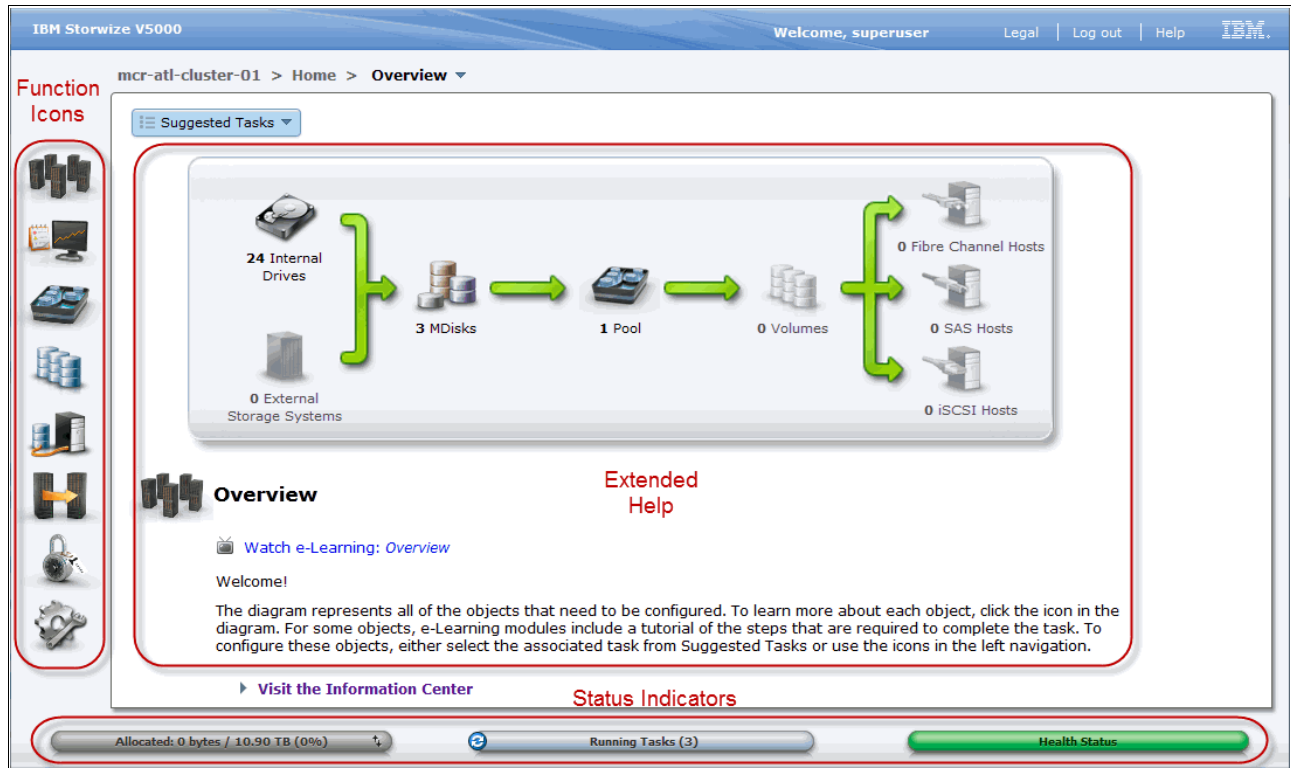


Figure 3-4 Three main sections of the IBM Storwize V5000 overview panel

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

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

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

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

The Status Indicators section shows the following horizontal status bars:

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

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

3.2 Navigation

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

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

3.2.1 Function icons navigation

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

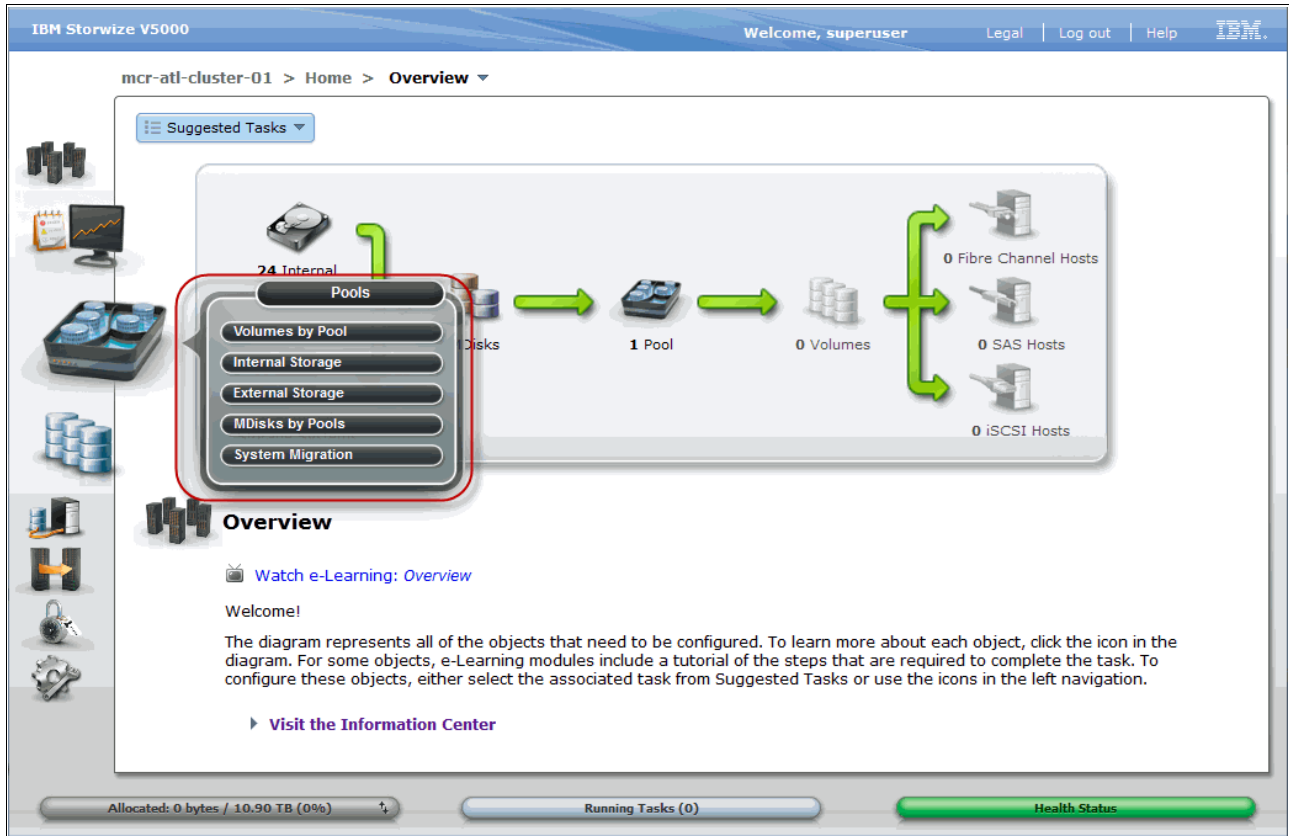


Figure 3-5 Hovering over a function icon

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

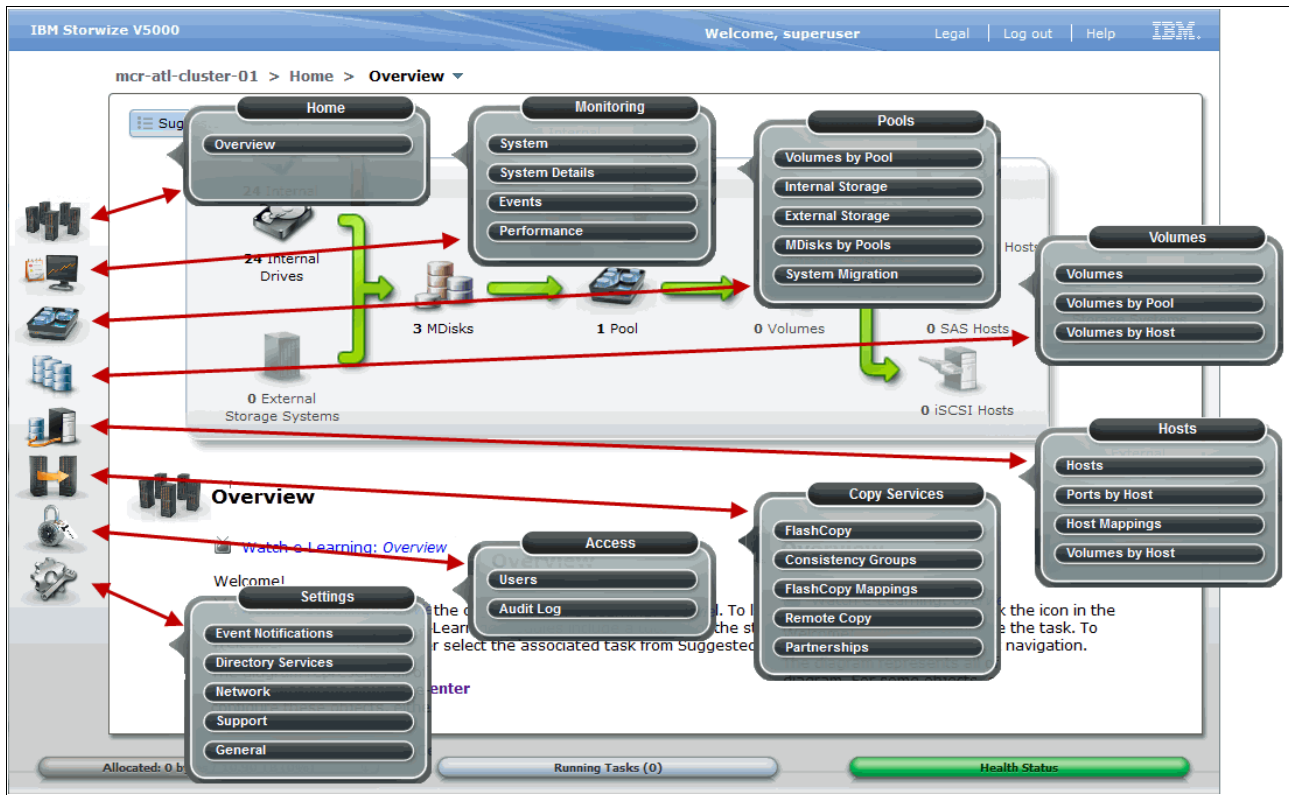


Figure 3-6 Options that are listed under Function Icons section

3.2.2 Extended help navigation

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

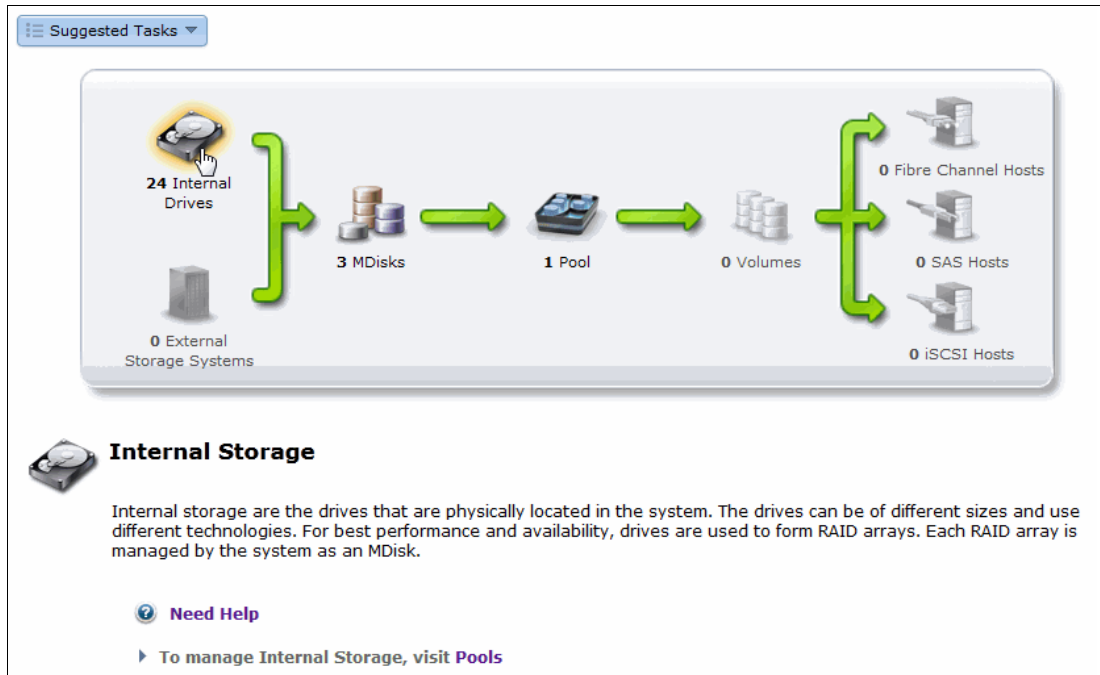


Figure 3-7 Navigating GUI with the extended help section

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

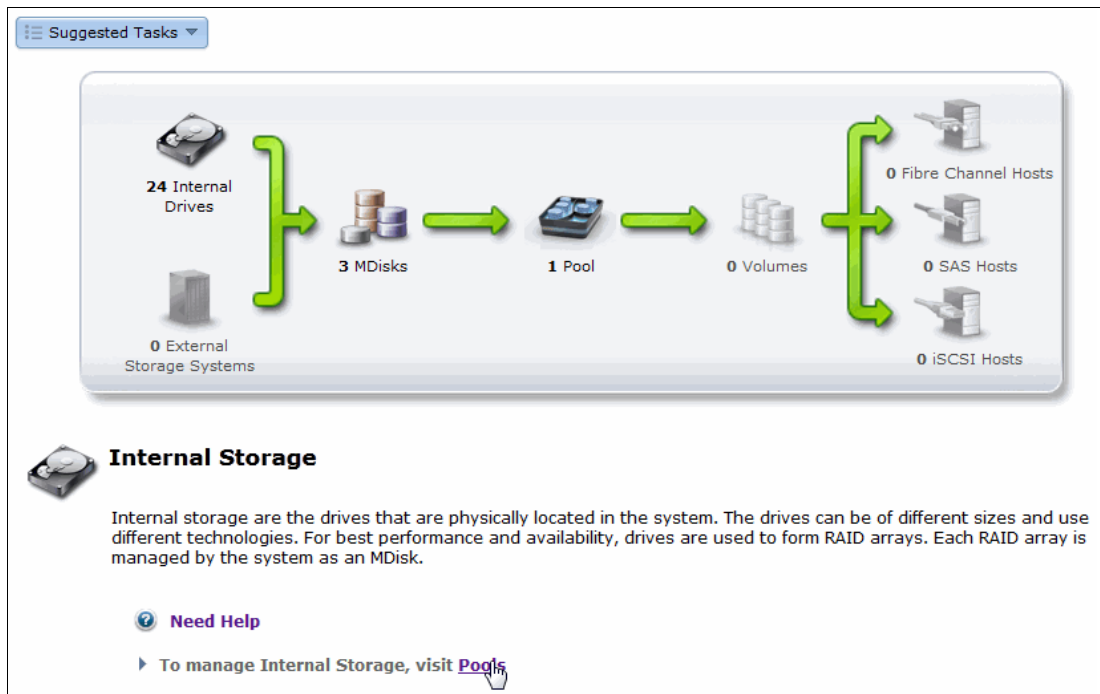


Figure 3-8 Using the extended help section

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

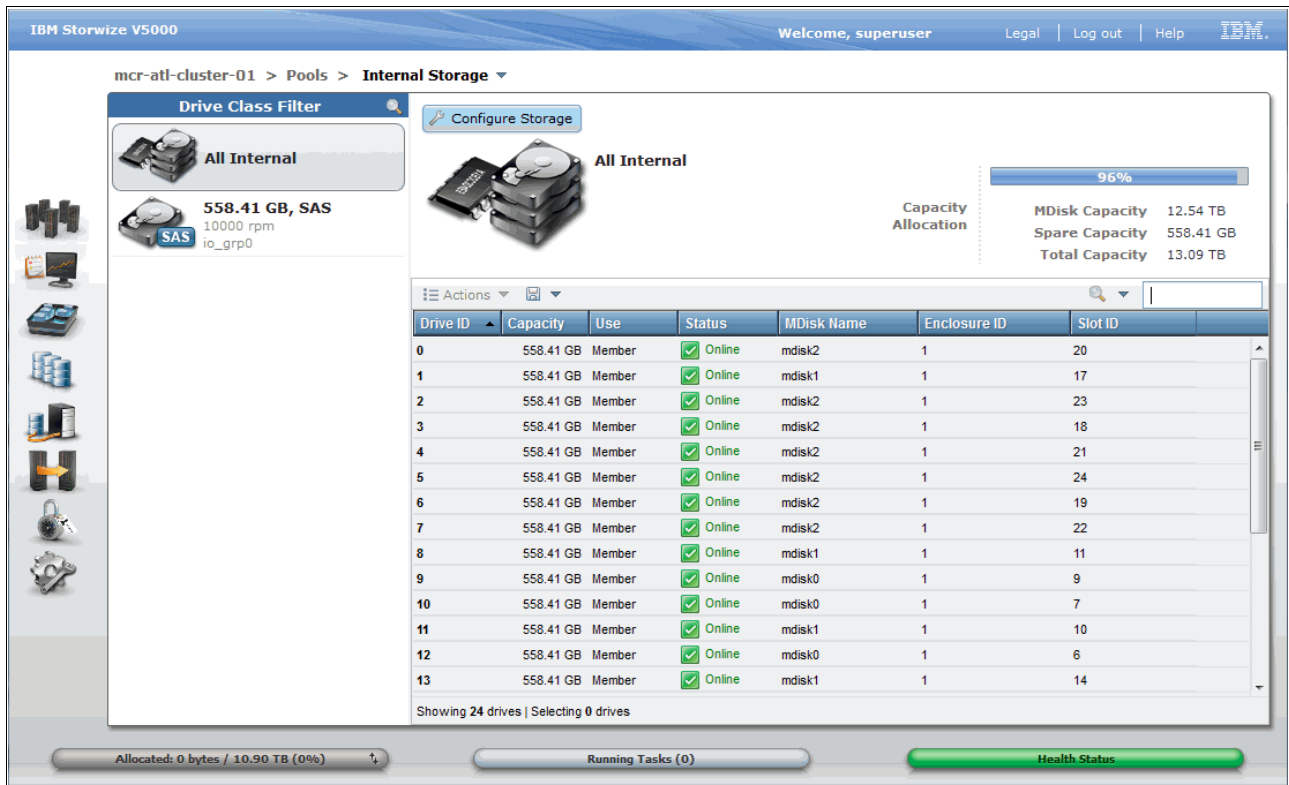


Figure 3-9 The internal storage configuration panel

3.2.3 Breadcrumb navigation aid

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

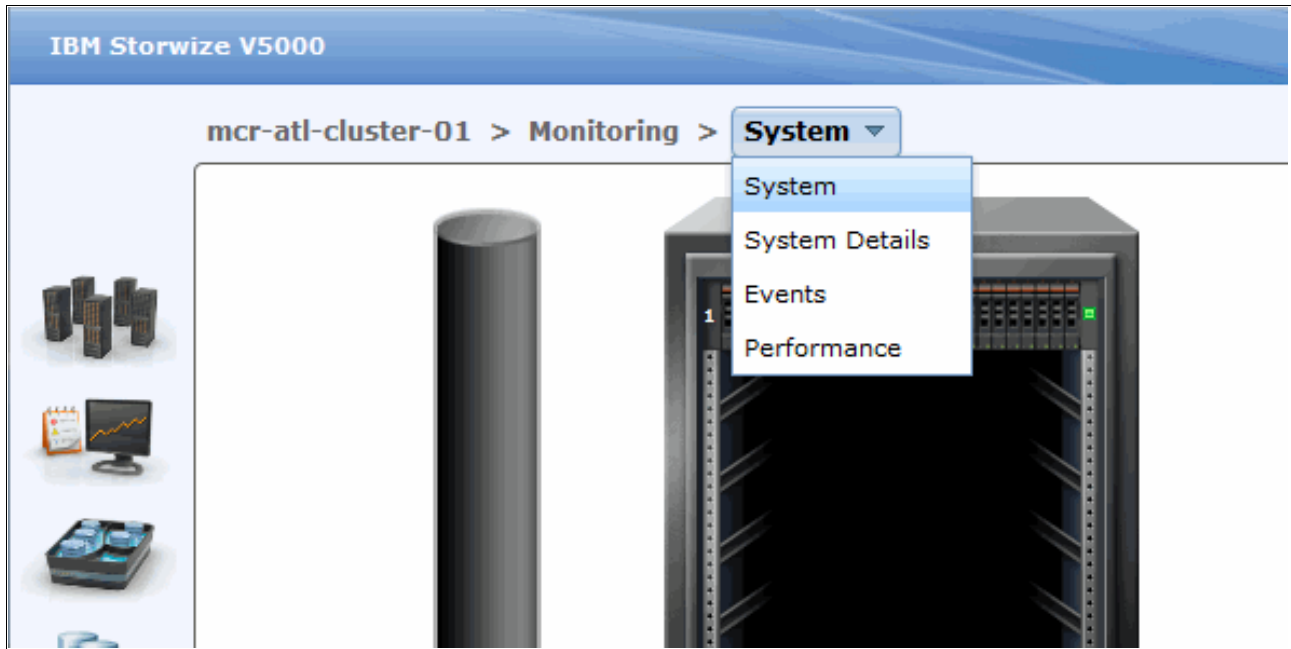


Figure 3-10 Breadcrumb navigation aid

3.2.4 Suggested Tasks feature

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

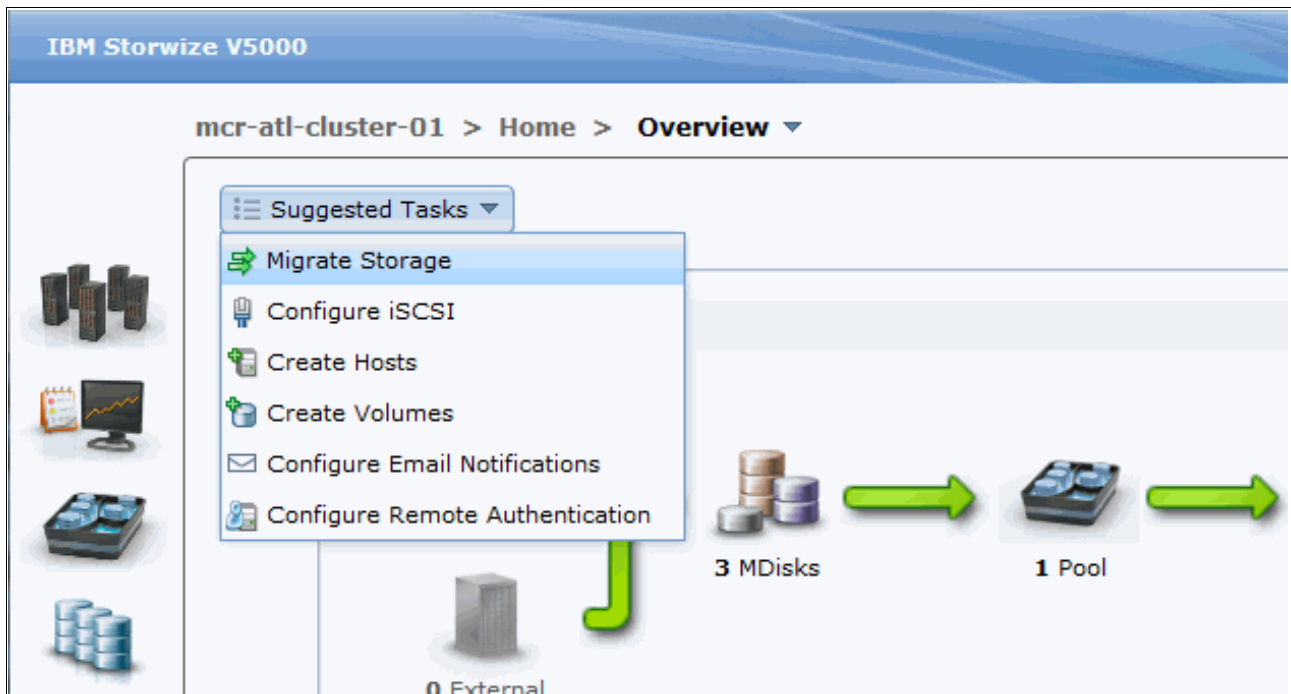


Figure 3-11 Suggested Tasks navigation and configuration aid

3.2.5 Presets

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

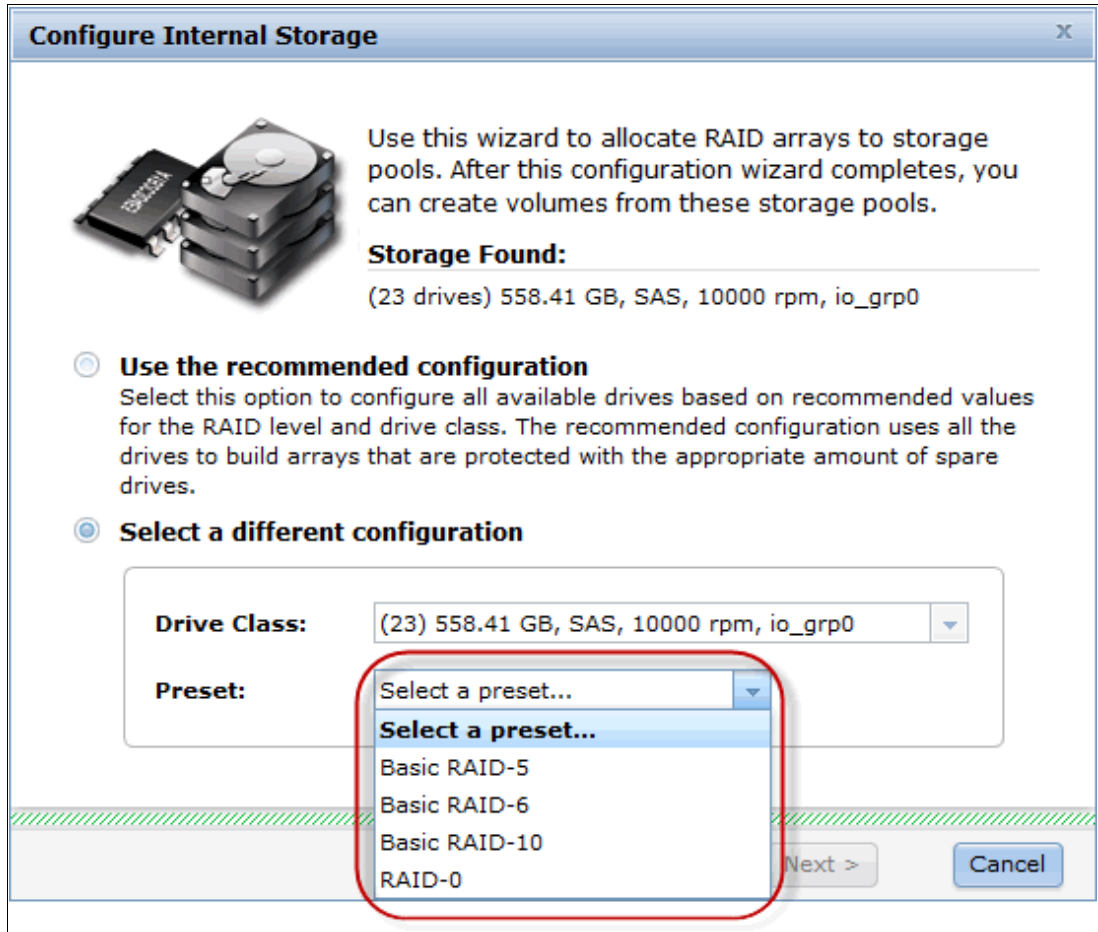


Figure 3-12 Internal storage preset selection

3.2.6 Access actions

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

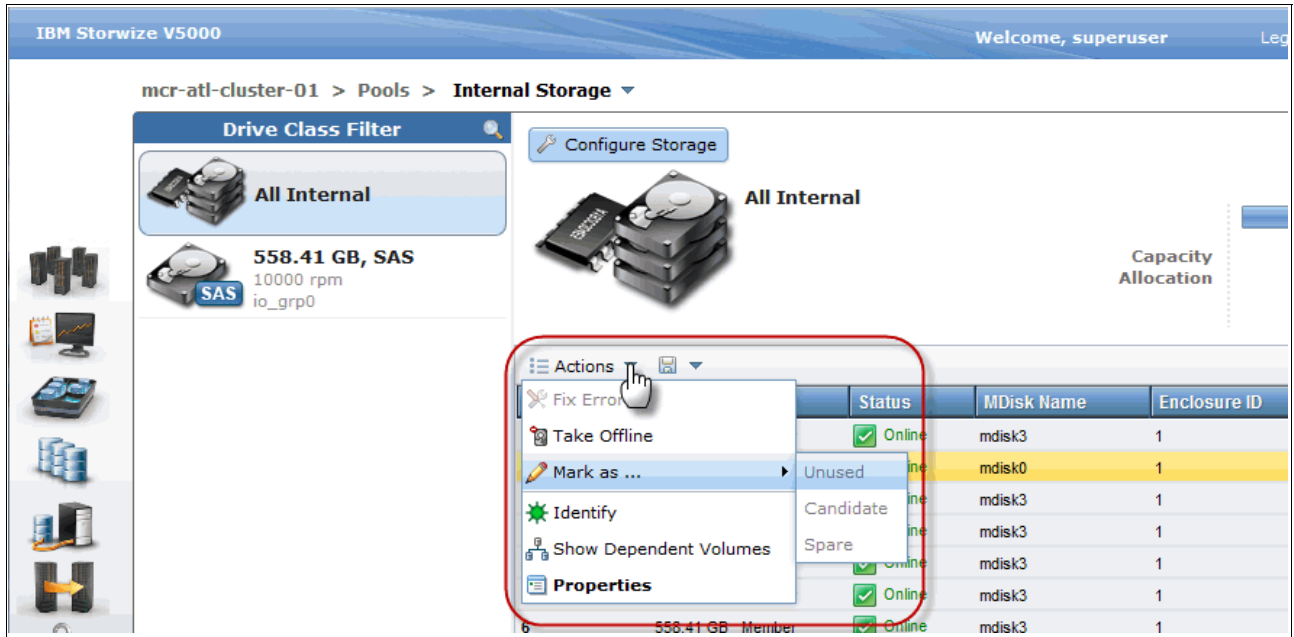


Figure 3-13 Actions menu

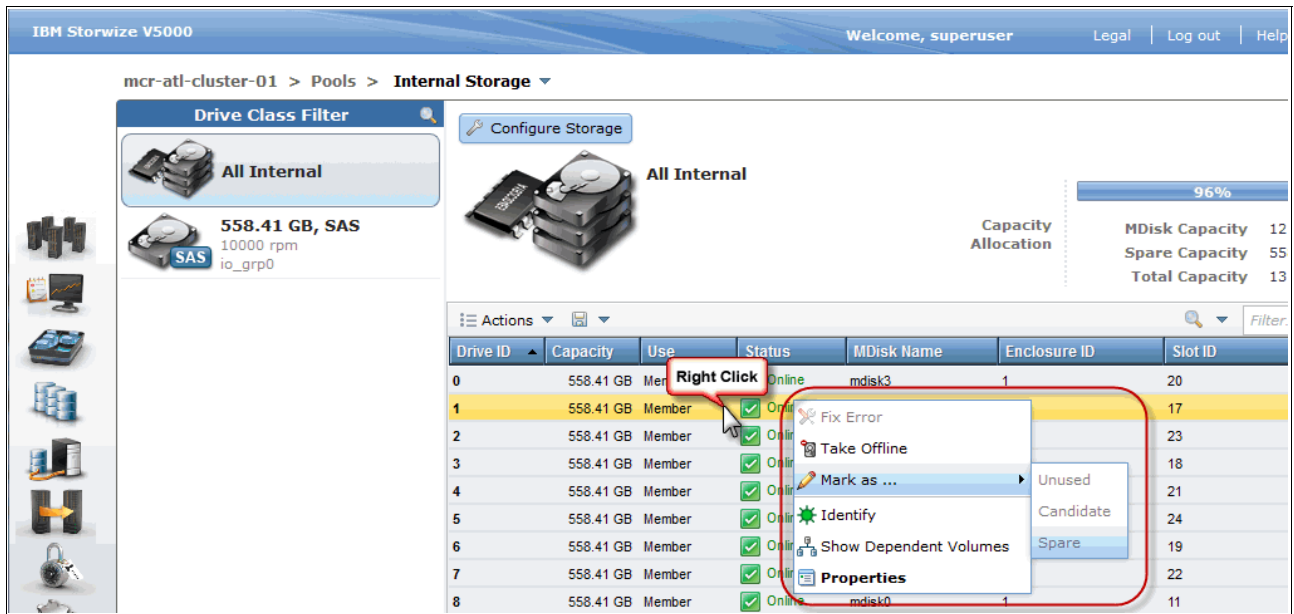


Figure 3-14 Right-clicking the Actions menu

3.2.7 Task progress

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

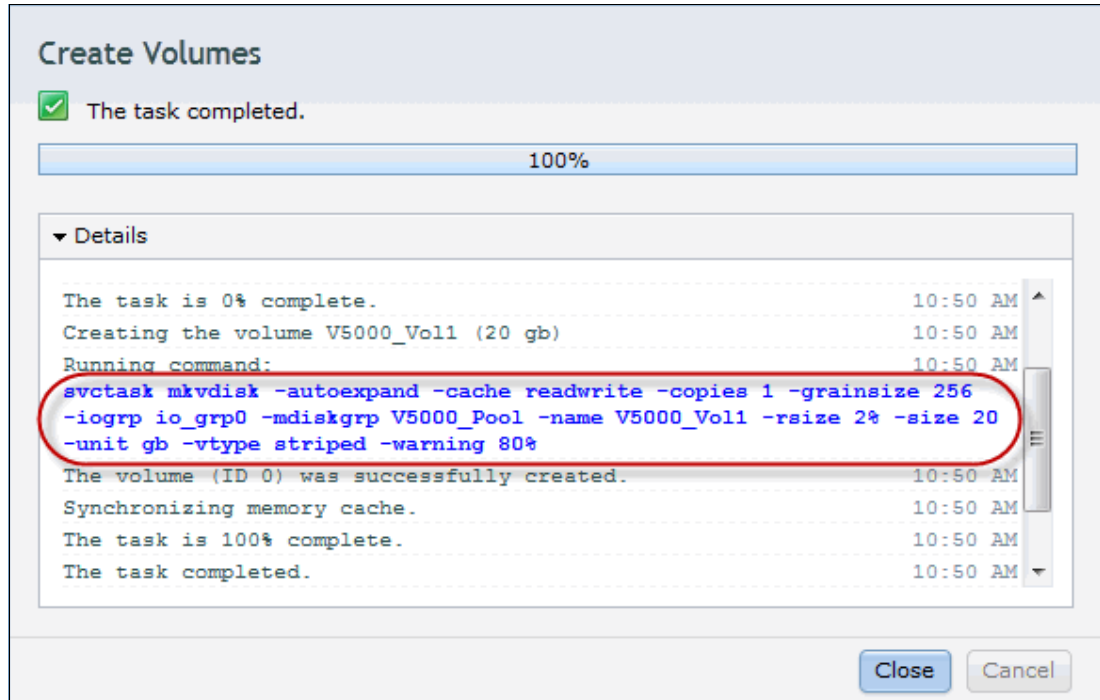


Figure 3-15 Task progress panel

3.2.8 Navigating panels with tables

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

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

Sorting columns

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

IBM Storwize V5000 Welcome, superuser Legal

mcr-atl-cluster-01 > Volumes > Volumes ▾

New Volume Actions ▾

Name	Status	Capacity	Storage Pool	Host Mappings
V5000_Vol10	Online	1.00 GB	V5000_Pool	No
V5000_Vol2	Online	2.00 GB	V5000_Pool	No
V5000_Vol3	Online	2.00 GB	V5000_Pool	No
V5000_Vol4	Online	2.00 GB	V5000_Pool	No
V5000_Vol5	Online	2.00 GB	V5000_Pool	No
V5000_Vol6	Online	2.00 GB	V5000_Pool	No
V5000_Vol7	Online	2.00 GB	V5000_Pool	No
V5000_Vol8	Online	2.00 GB	V5000_Pool	No
V5000_Vol9	Online	2.00 GB	V5000_Pool	No
V5000_Vol11	Online	20.00 GB	V5000_Pool	No

Figure 3-16 Sorting columns by clicking the column heading

Reordering columns

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

New Volume Actions ▾

Name	Status	Capacity	Storage Pool	Host Mappings
V5000_Vol10	Online	1.00 GB	V5000_Pool	No
V5000_Vol2	Online	2.00 GB	V5000_Pool	No
V5000_Vol3	Online	2.00 GB	V5000_Pool	No
V5000_Vol4	Online	2.00 GB	V5000_Pool	No
V5000_Vol5	Online	2.00 GB	V5000_Pool	No
V5000_Vol6	Online	2.00 GB	V5000_Pool	No
V5000_Vol7	Online	2.00 GB	V5000_Pool	No
V5000_Vol8	Online	2.00 GB	V5000_Pool	No
V5000_Vol9	Online	2.00 GB	V5000_Pool	No
V5000_Vol11	Online	20.00 GB	V5000_Pool	No

Figure 3-17 Reordering columns

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

The screenshot shows a table with the following columns: Name, Status, Capacity, Storage Pool, and Host Mappings. A red circle highlights the 'Host Mappings' column header, and a mouse cursor is shown dragging it to the position between 'Name' and 'Status'. A tooltip 'Host Mappi...' is visible above the new position.

Name	Status	Capacity	Storage Pool	Host Mappings
V5000_Vol1	Online	20.00 GB	V5000_Pool	No
V5000_Vol2	Online	2.00 GB	V5000_Pool	No
V5000_Vol3	Online	2.00 GB	V5000_Pool	No
V5000_Vol4	Online	2.00 GB	V5000_Pool	No
V5000_Vol5	Online	2.00 GB	V5000_Pool	No
V5000_Vol6	Online	2.00 GB	V5000_Pool	No
V5000_Vol7	Online	2.00 GB	V5000_Pool	No
V5000_Vol8	Online	2.00 GB	V5000_Pool	No
V5000_Vol9	Online	2.00 GB	V5000_Pool	No
V5000_Vol10	Online	1.00 GB	V5000_Pool	No

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

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

The screenshot shows the same table as in Figure 3-18, but the columns are now ordered: Name, Host Mappings, Status, Capacity, and Storage Pool. The 'Host Mappings' column is now the second column, and the 'Status' column is the third.

Name	Host Mappings	Status	Capacity	Storage Pool
V5000_Vol1	No	Online	20.00 GB	V5000_Pool
V5000_Vol2	No	Online	2.00 GB	V5000_Pool
V5000_Vol3	No	Online	2.00 GB	V5000_Pool
V5000_Vol4	No	Online	2.00 GB	V5000_Pool
V5000_Vol5	No	Online	2.00 GB	V5000_Pool
V5000_Vol6	No	Online	2.00 GB	V5000_Pool
V5000_Vol7	No	Online	2.00 GB	V5000_Pool
V5000_Vol8	No	Online	2.00 GB	V5000_Pool
V5000_Vol9	No	Online	2.00 GB	V5000_Pool
V5000_Vol10	No	Online	1.00 GB	V5000_Pool

Figure 3-19 Reordering column headings

Adding or removing columns

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

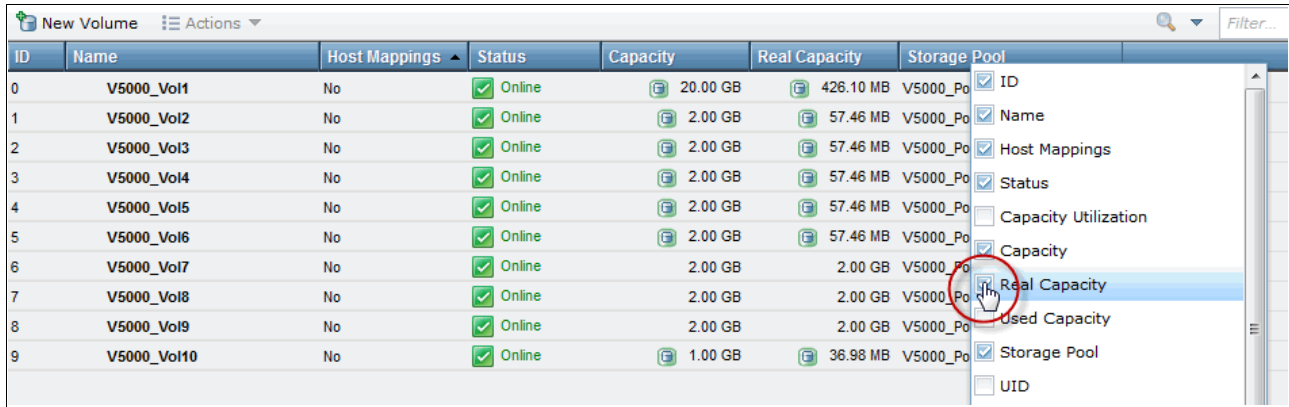


Figure 3-20 Adding column heading Real Capacity

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

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

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

Multiple selections

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

Using the Shift key

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




Figure 3-21 Selection of three sequential items

Using the Ctrl key

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

Name	Status	Capacity	Storage Pool	UID	Host Mappings
Volume 1	Online	200.00 GB	mdiskgrp0	60050760009901D37800000000000003	No
Volume 2	Online	200.00 GB	mdiskgrp0	60050760009901D37800000000000004	No
Volume 3	Online	200.00 GB	mdiskgrp0	60050760009901D37800000000000005	No
Volume 4	Online	200.00 GB	mdiskgrp0	60050760009901D37800000000000006	No
Volume 5	Online	200.00 GB	mdiskgrp0	60050760009901D37800000000000007	No

Figure 3-22 Selecting two non-sequential items

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

Name	Status	Capacity	Storage Pool	UID	Host Mappings
Volume 1	Online	200.00 GB	mdiskgrp0	60050760009901D37800000000000003	No
Volume 2	Online	200.00 GB	mdiskgrp0	60050760009901D37800000000000004	No
Volume 3	Online	200.00 GB	mdiskgrp0	60050760009901D37800000000000005	No
Volume 4	Online	200.00 GB	mdiskgrp0	60050760009901D37800000000000006	No
Volume 5	Online	200.00 GB	mdiskgrp0	60050760009901D37800000000000007	No

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

Filtering objects

To focus on a subset of the listed items that are shown in a panel with columns, use the filter field that is found at the upper right side of the table. This tool shows items that match the value that is entered. Figure 3-24 shows the text **Vol11** was entered into the filter field. Now, only volumes with the text **Vol11** in any column are listed and the filter word also is highlighted.

ID	Name	Host Mappings	Status	Capacity	Real Capacity	Storage Pool
0	V5000_Vol11	No	Online	20.00 GB	426.10 MB	V5000_Pool
9	V5000_Vol10	No	Online	1.00 GB	36.98 MB	V5000_Pool

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

Filter by column

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

ID	Name	Status	Capacity	Real Capacity	Storage P...
0	V5000_Vol11	Online	20.00 ...	426.10 MB	V5000_Pool
1	V5000_Vol2	Online	2.00 GB	57.46 MB	V5000_Pool
10	ESXi-Redbooks	Online	10.00 GB	10.00 GB	V5000_Pool
11	windows2008r2	Online	10.00 GB	10.00 GB	V5000_Pool
12	windows2008_0	Online	10.00 GB	10.00 GB	V5000_Pool
13	windows2008_1	Online	10.00 GB	10.00 GB	V5000_Pool
2	V5000_Vol3	Online	2.00 GB	57.46 MB	V5000_Pool
3	V5000_Vol4	Online	2.00 GB	57.46 MB	V5000_Pool
4	V5000_Vol5	Online	2.00 GB	57.46 MB	V5000_Pool

Figure 3-25 Filter by column

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

The screenshot shows a table with columns: ID, Name, Status, Capacity, Real Capacity, Storage P..., and Host Mappings. The Host Mappings column is filtered to show 'Yes' values. A dropdown menu is open over the Host Mappings column, showing options: -- Choose Value --, Yes, No, and Mapped. The 'Yes' option is selected.

ID	Name	Status	Capacity	Real Capacity	Storage P...	Host Mappings
10	ESXi-Redbooks	Online	10.00 GB	10.00 GB	V5000_Pool	Yes
11	windows2008r2	Online	10.00 GB	10.00 GB	V5000_Pool	Yes
12	windows2008_0	Online	10.00 GB	10.00 GB	V5000_Pool	Yes
13	windows2008_1	Online	10.00 GB	10.00 GB	V5000_Pool	Yes

Figure 3-26 Choosing filter value

3.3 Status Indicators menus

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

3.3.1 Horizontal bars

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

3.3.2 Allocated status bar menu

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

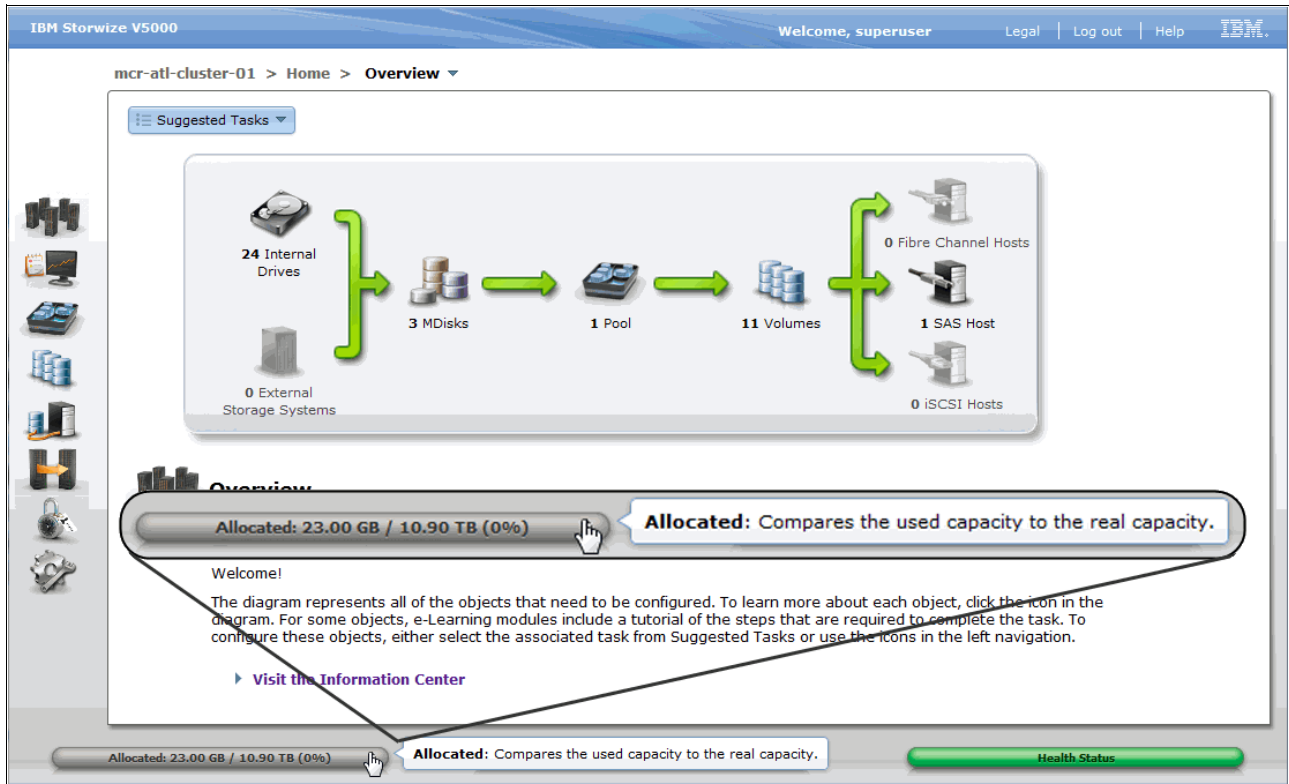


Figure 3-27 Allocated bar that compares used capacity to real capacity

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

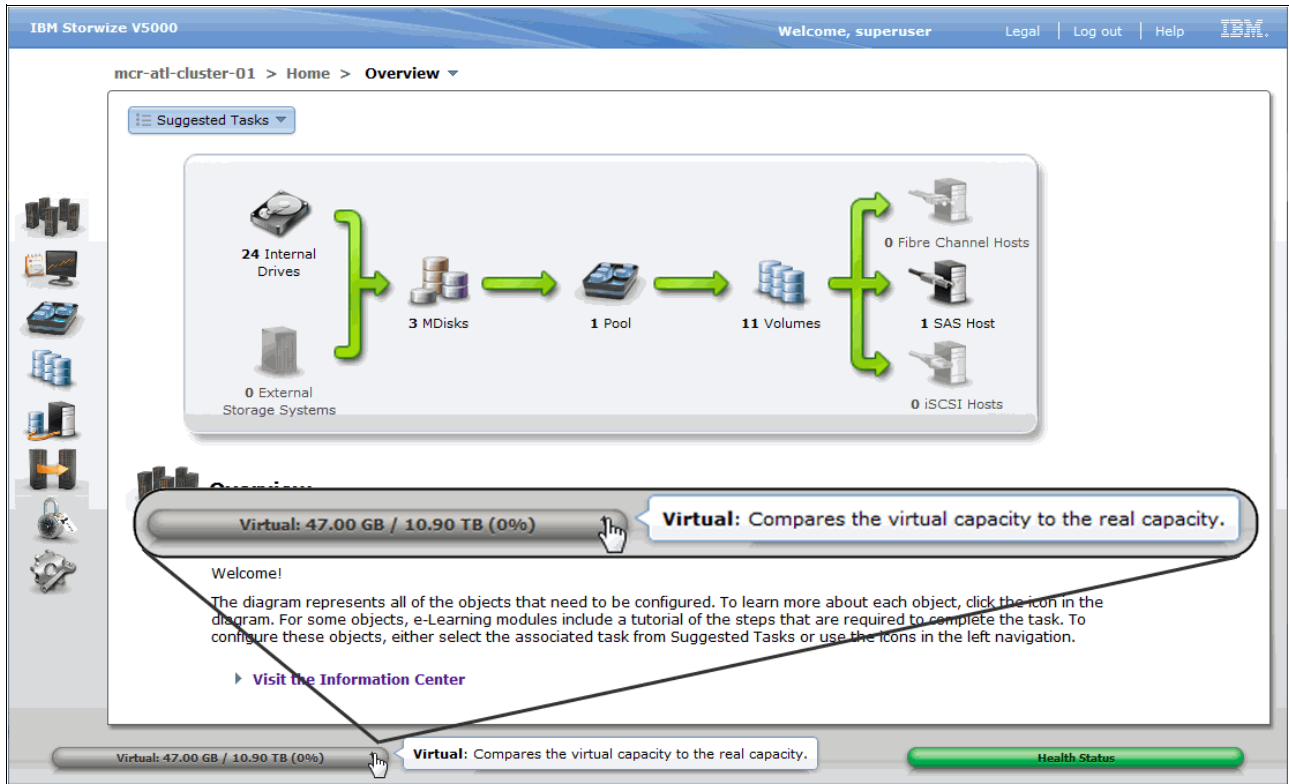


Figure 3-28 Changing the allocated menu comparison, virtual capacity to real capacity

3.3.3 Running tasks bar menu

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

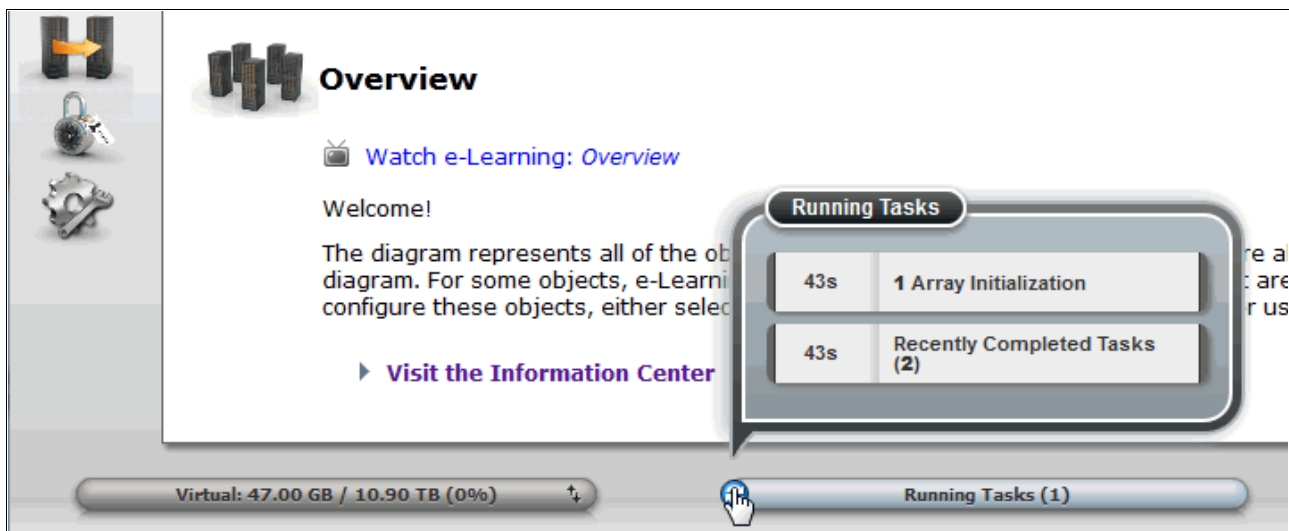


Figure 3-29 Running Tasks menu

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

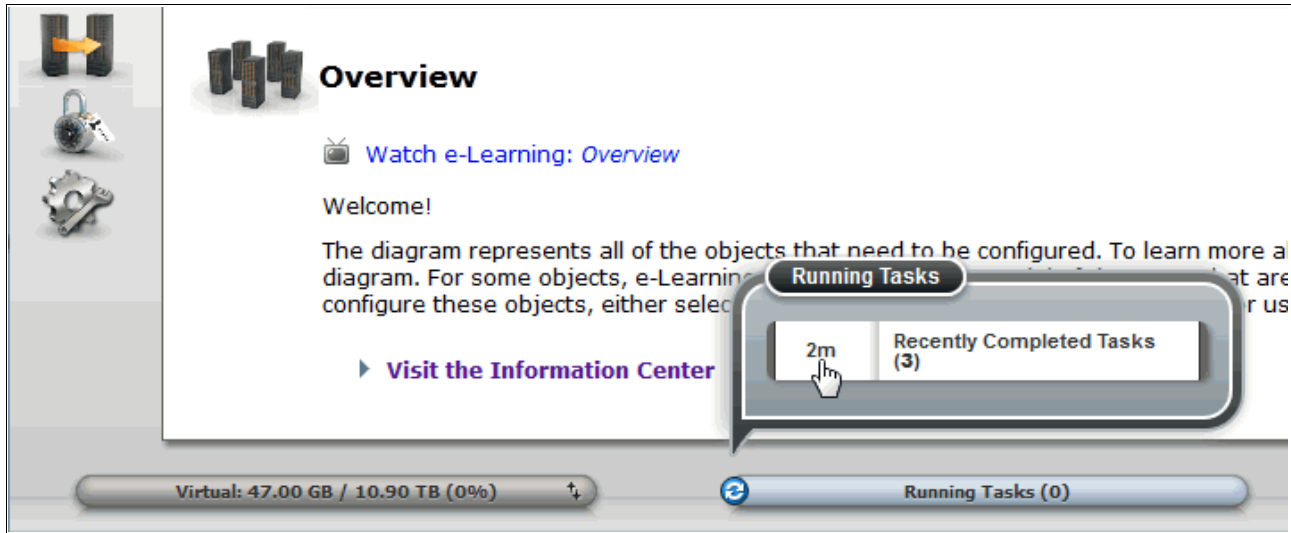


Figure 3-30 Selecting a task from the Running Task menu

Figure 3-31 shows the Recently Completed tasks panel.

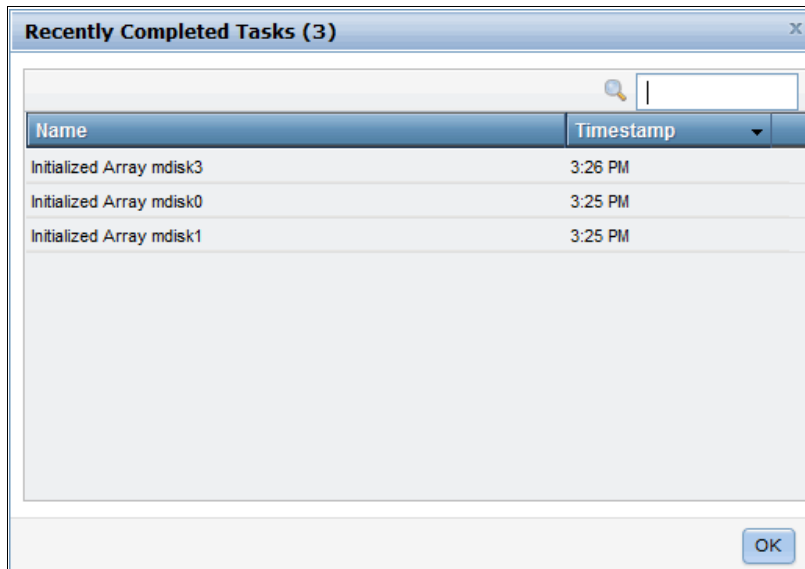


Figure 3-31 Recently Completed tasks panel

3.3.4 Health status bar menu

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

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

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

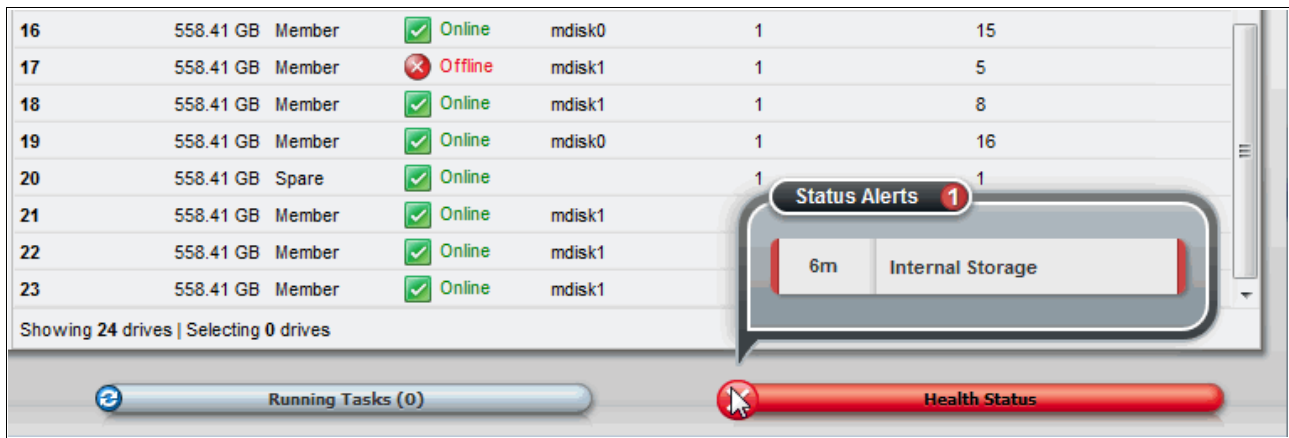


Figure 3-32 Health status menu

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

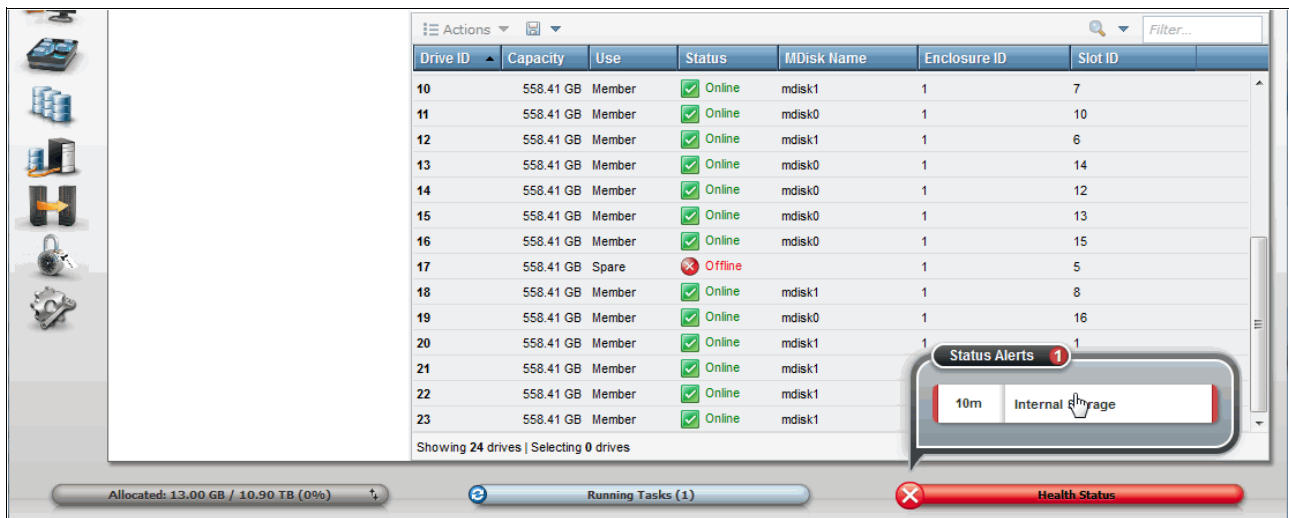


Figure 3-33 Status and description of an alert via the health status menu

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

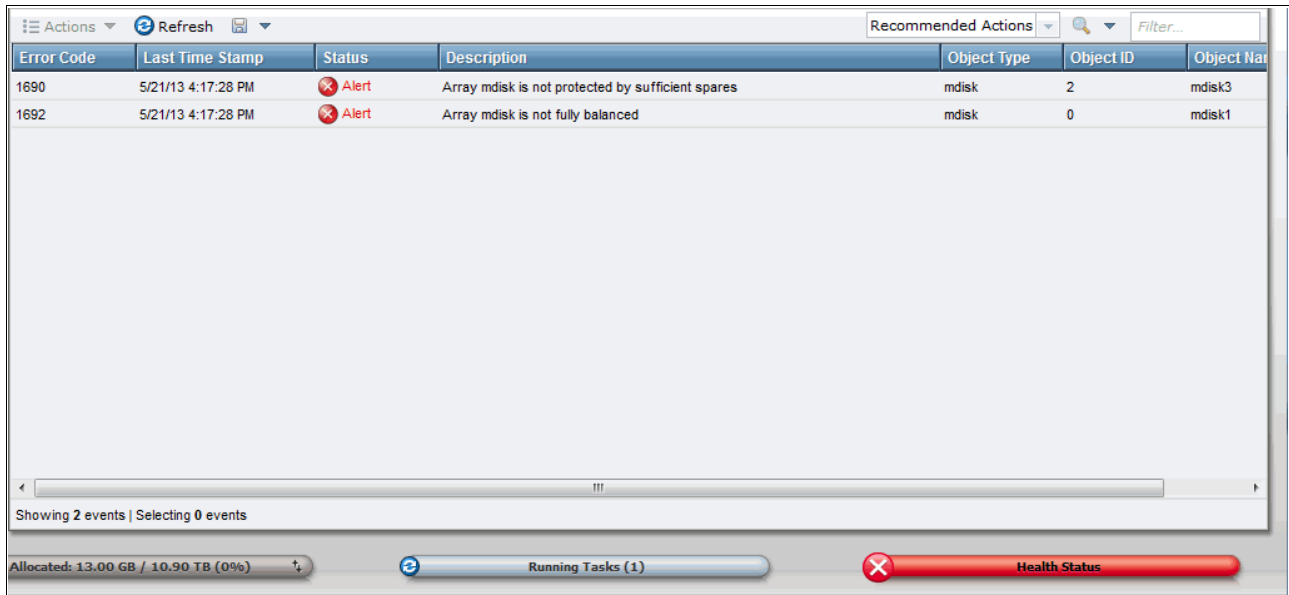


Figure 3-34 Events panel via health status menu

3.4 Function icon menus

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

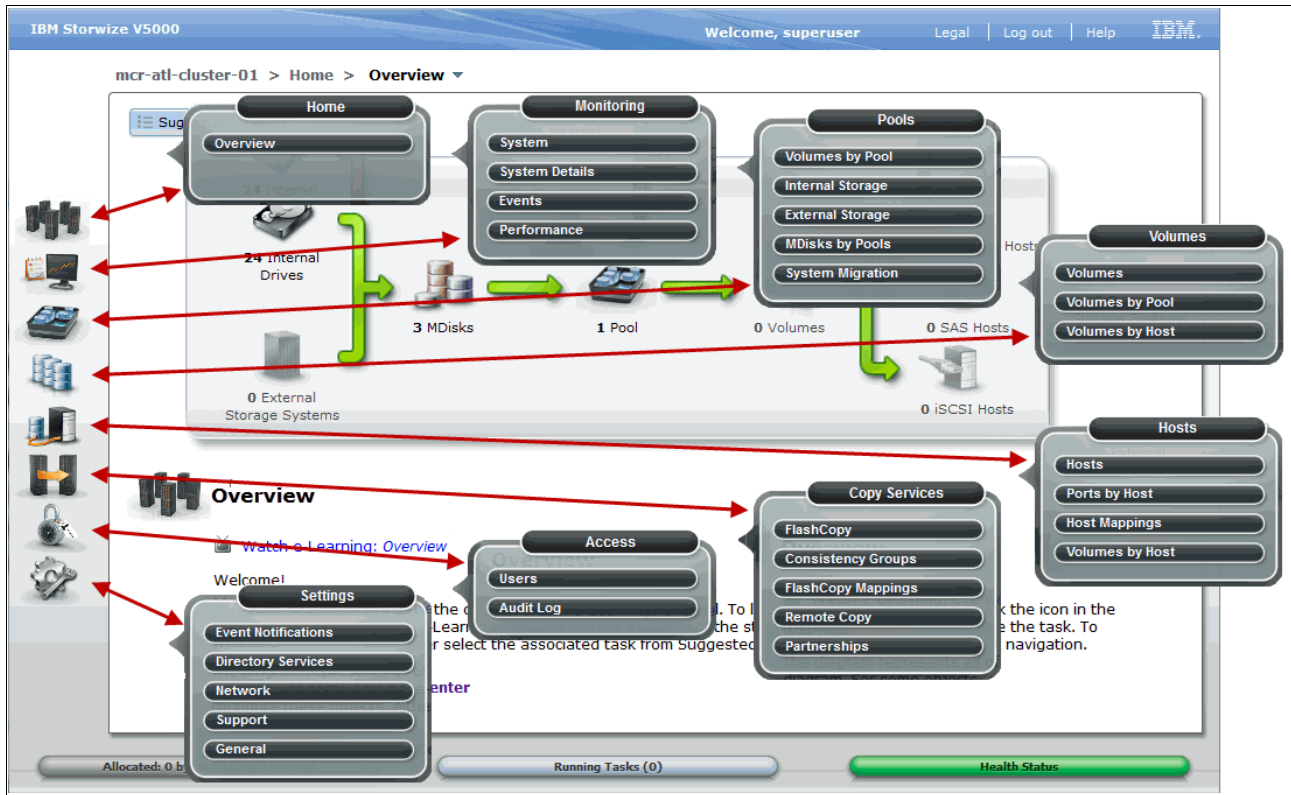


Figure 3-35 All Function Icon menus

3.4.1 Home menu

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

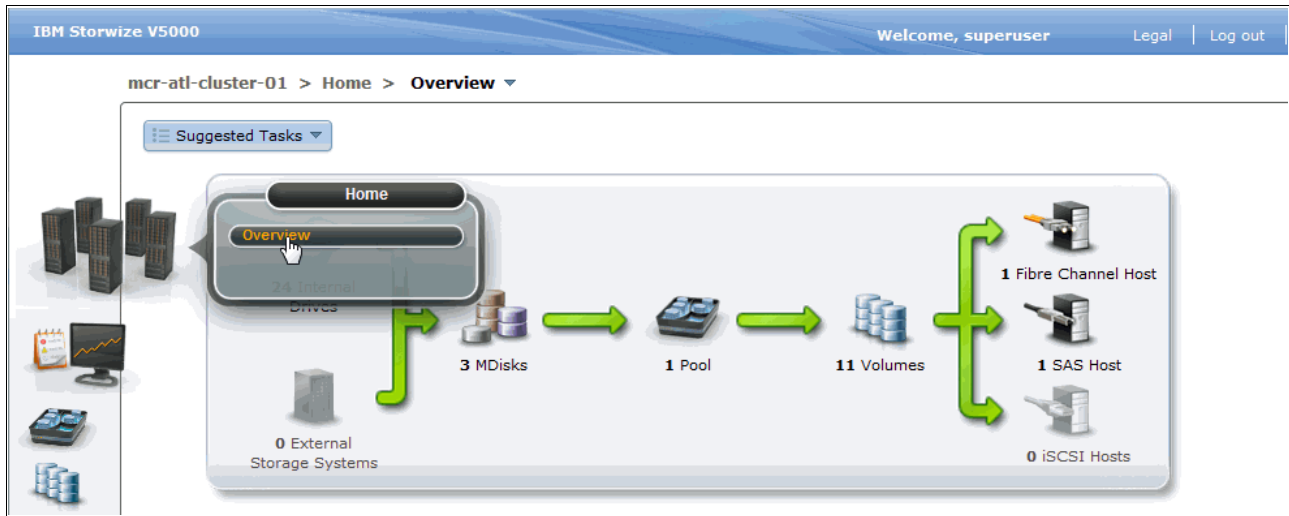


Figure 3-36 Home menu

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

3.4.2 Monitoring menu

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

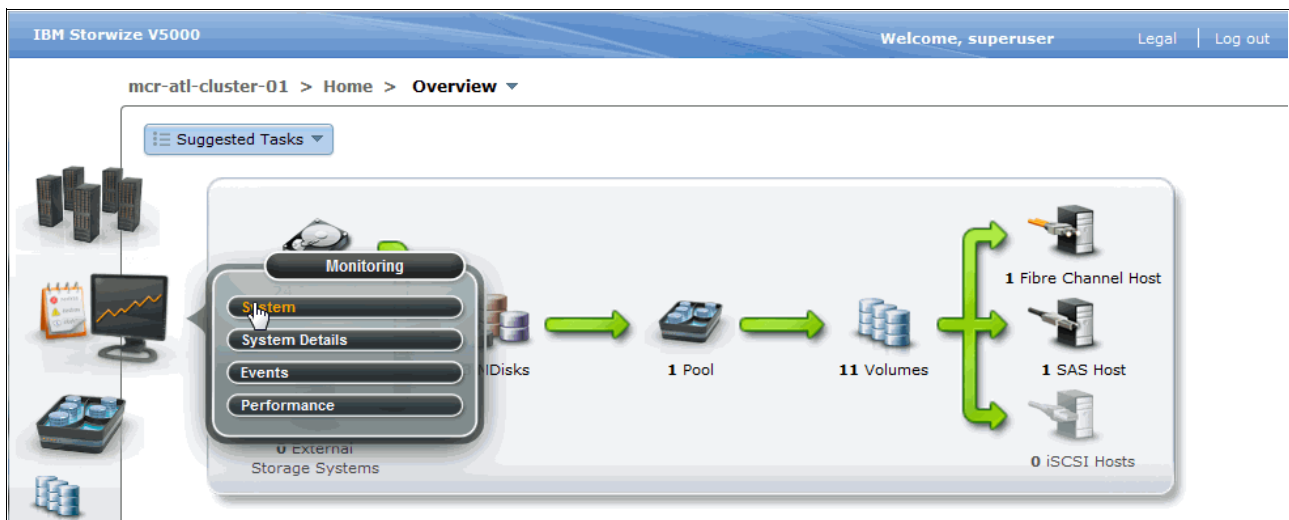


Figure 3-37 Monitoring menu

System panel

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

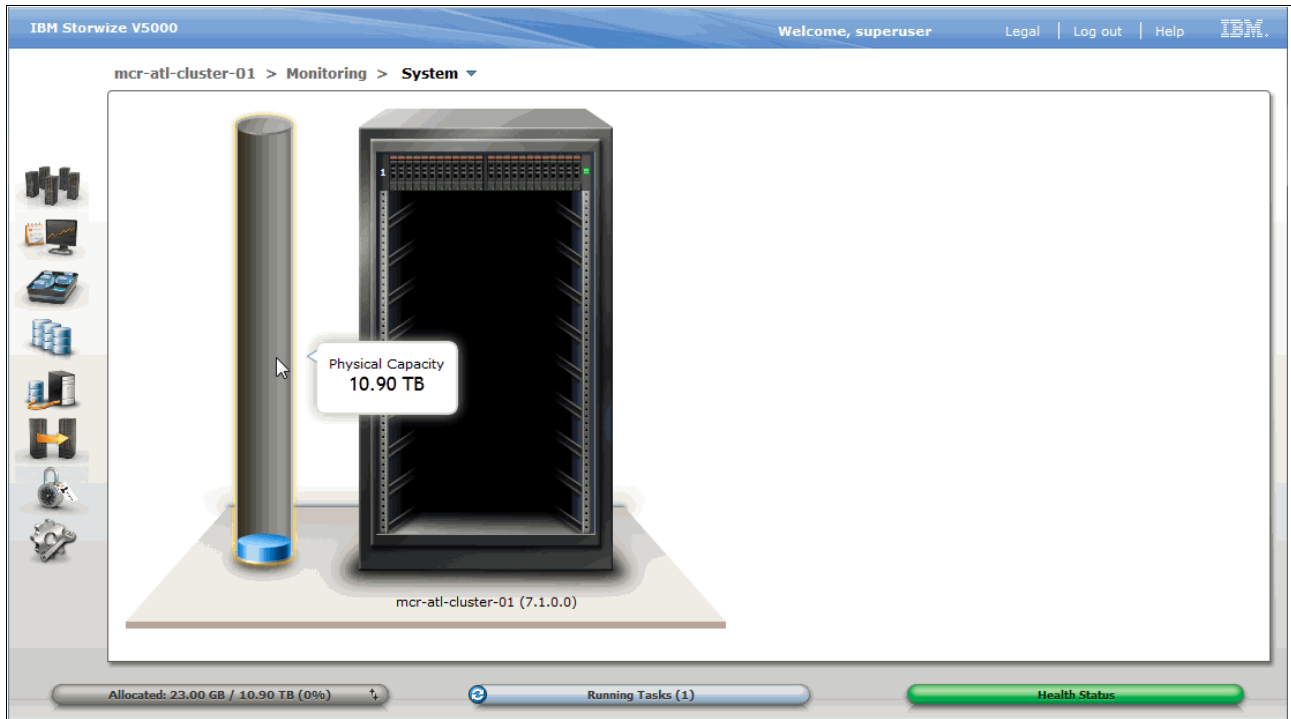


Figure 3-38 The System panel

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

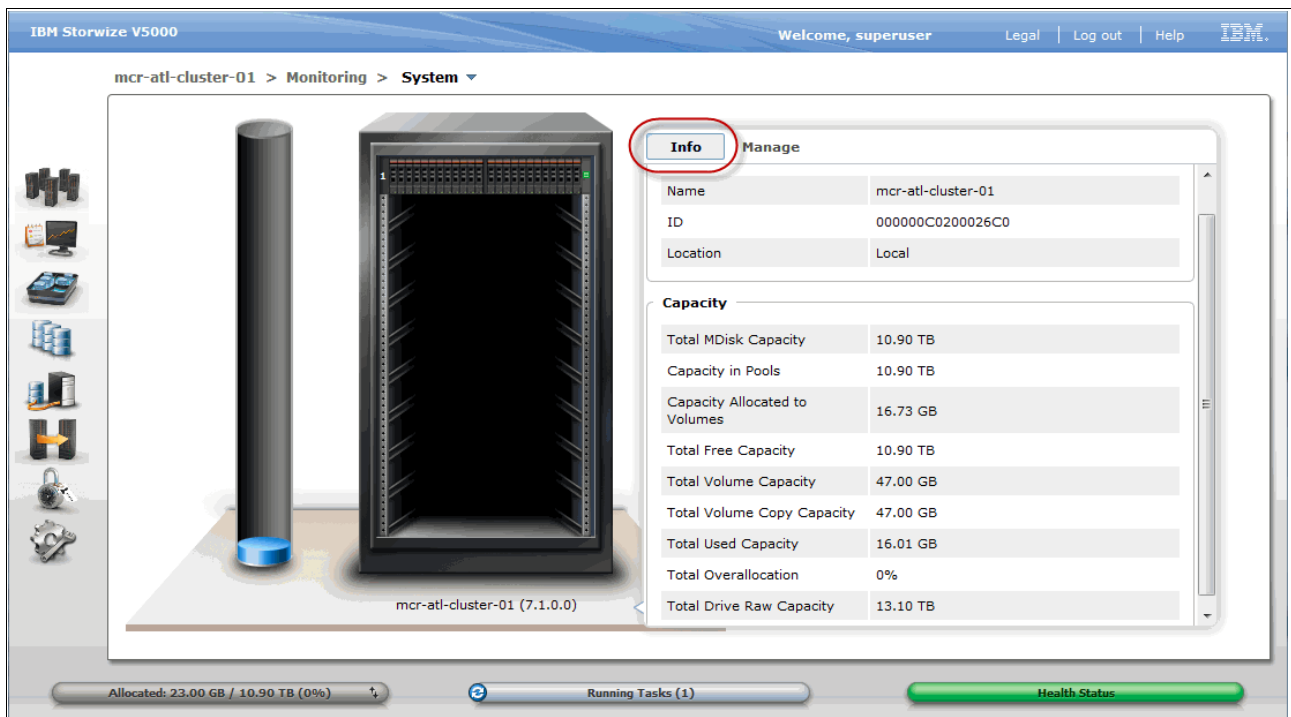


Figure 3-39 System panel Info tab

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

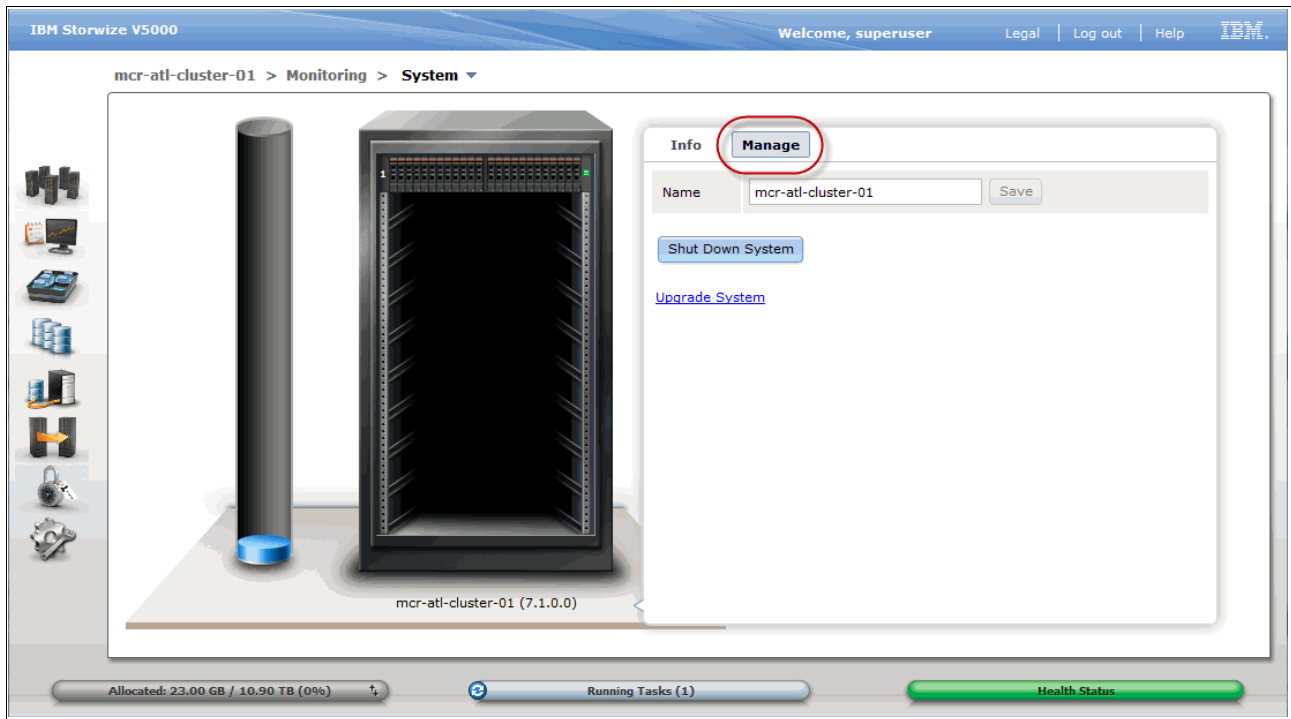


Figure 3-40 System panel Manage tab

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

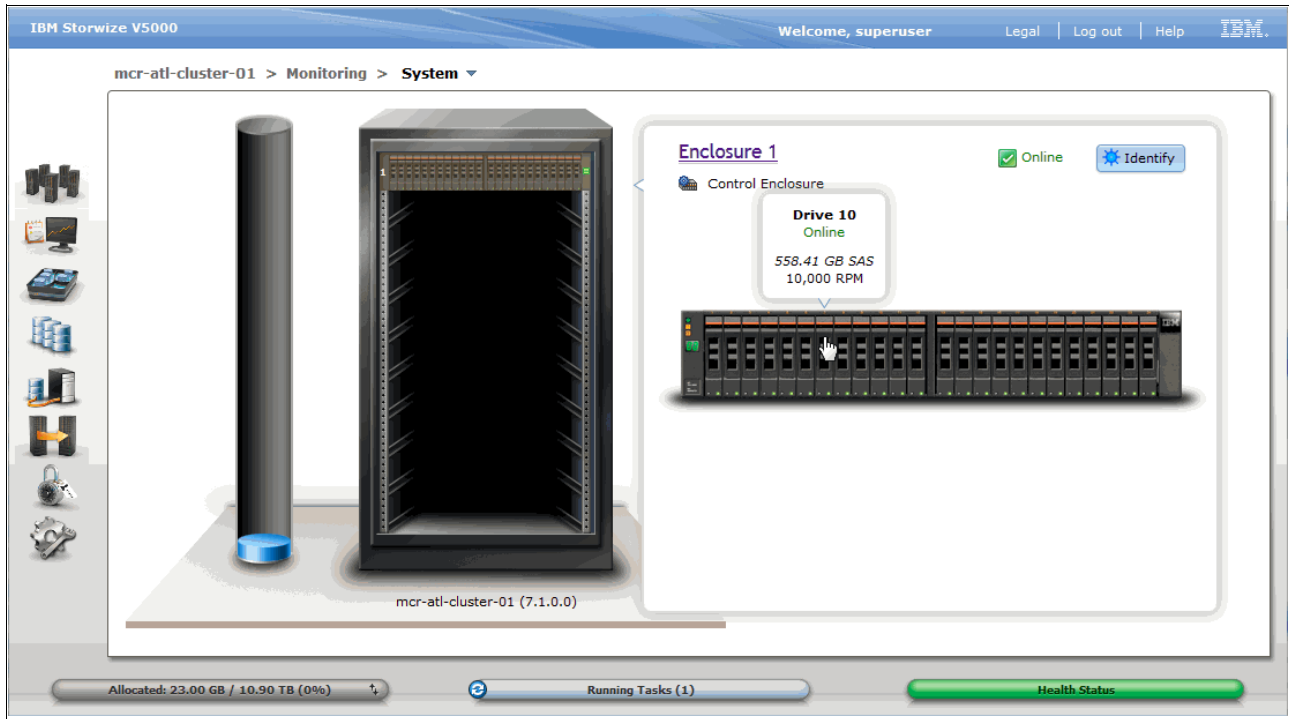


Figure 3-41 System panel enclosure view

System Details panel

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

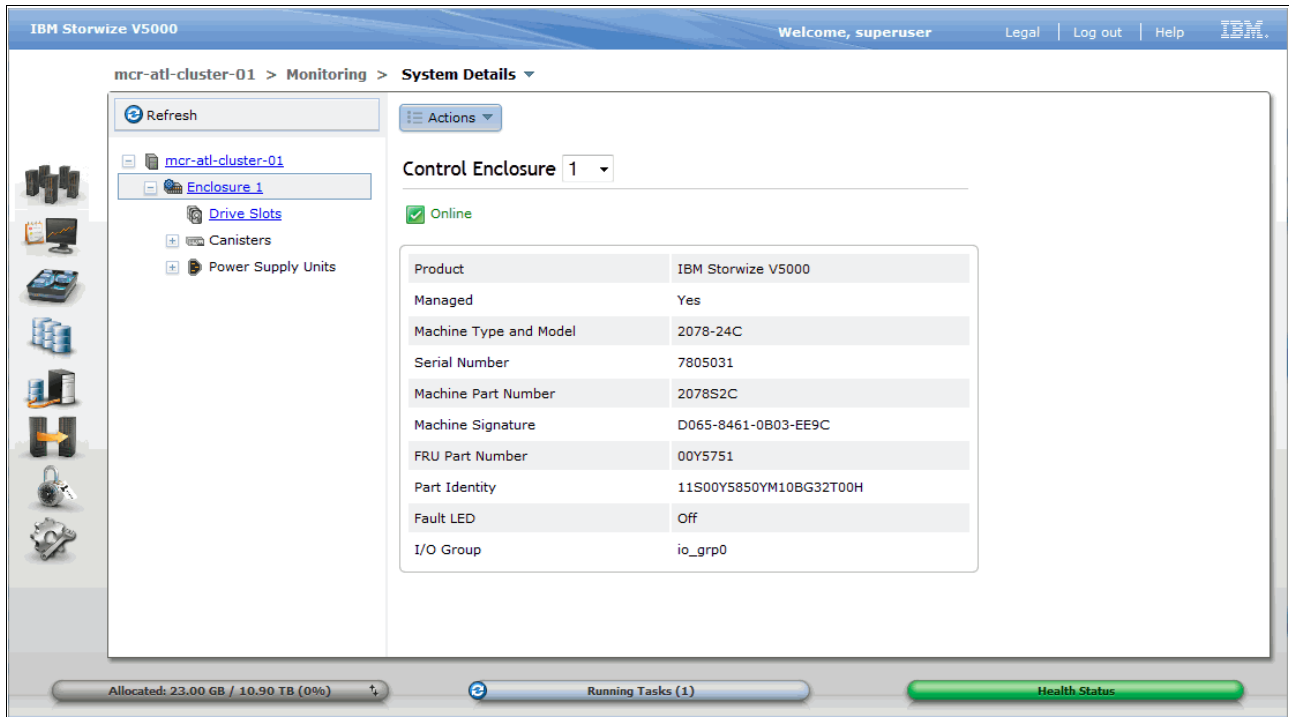


Figure 3-42 System Details panel

Actions and environmental statistics

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

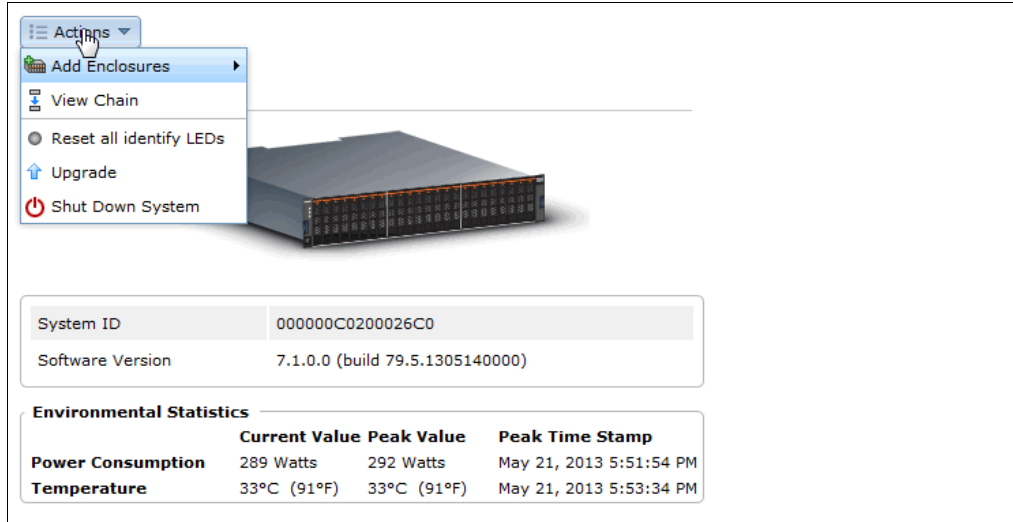


Figure 3-43 System details actions and environmental statistics

Node canister information

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

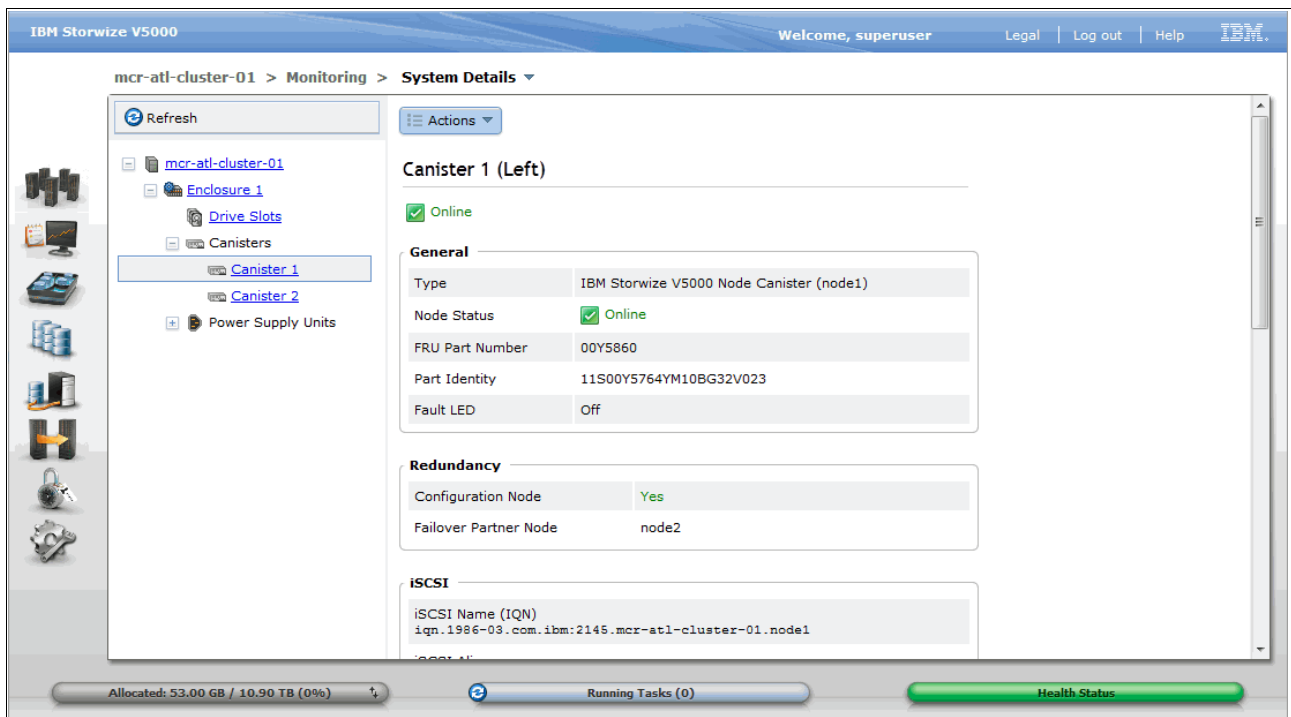


Figure 3-44 Node canister information via system details panel

Events panel

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

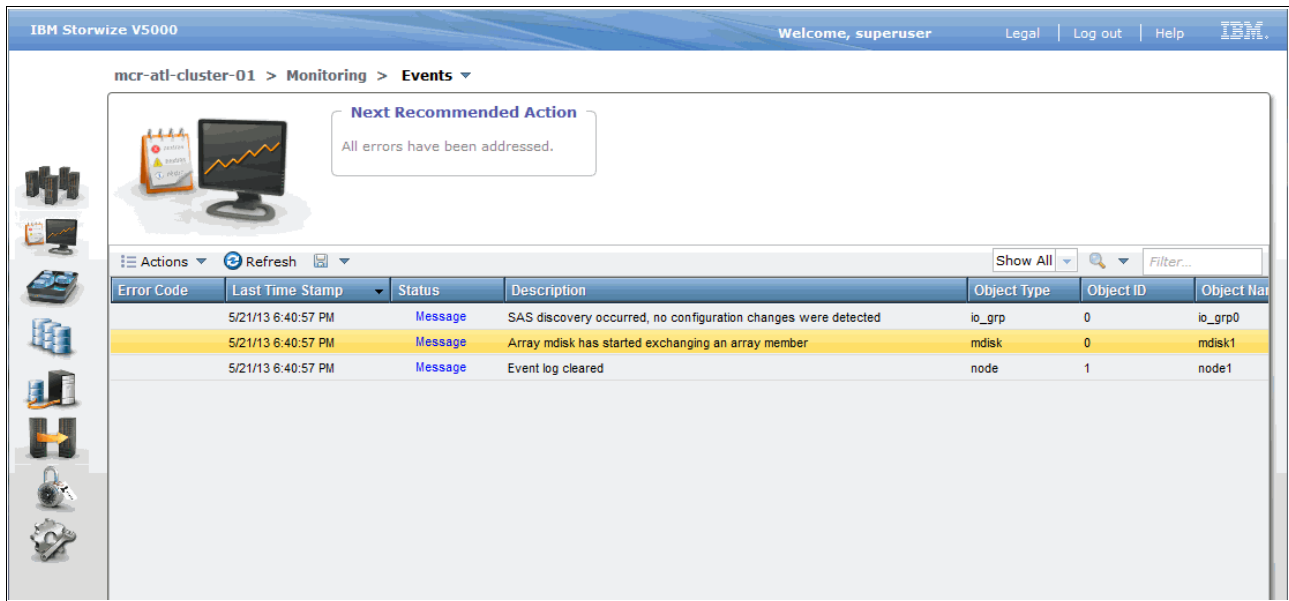


Figure 3-45 Events panel with all errors addressed

Filtering events view

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

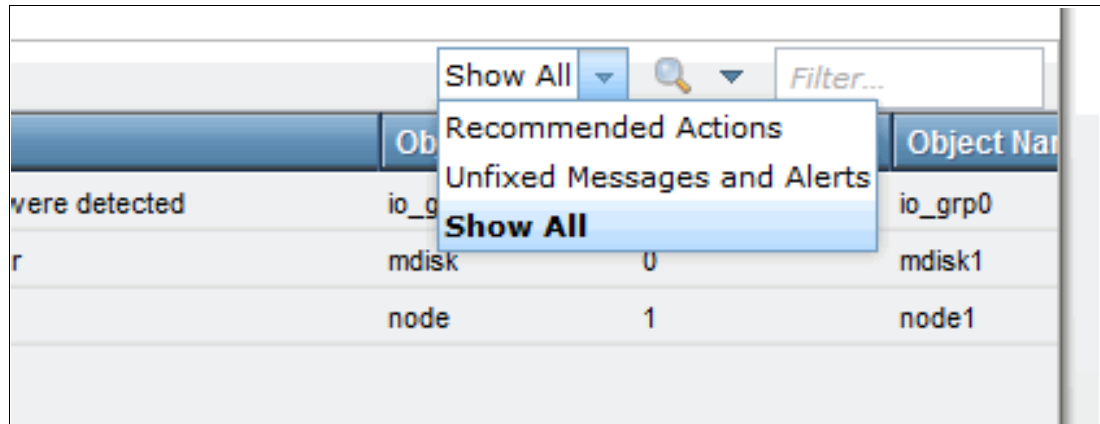


Figure 3-46 Unfixed messages and alerts in the events panel

Event properties

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

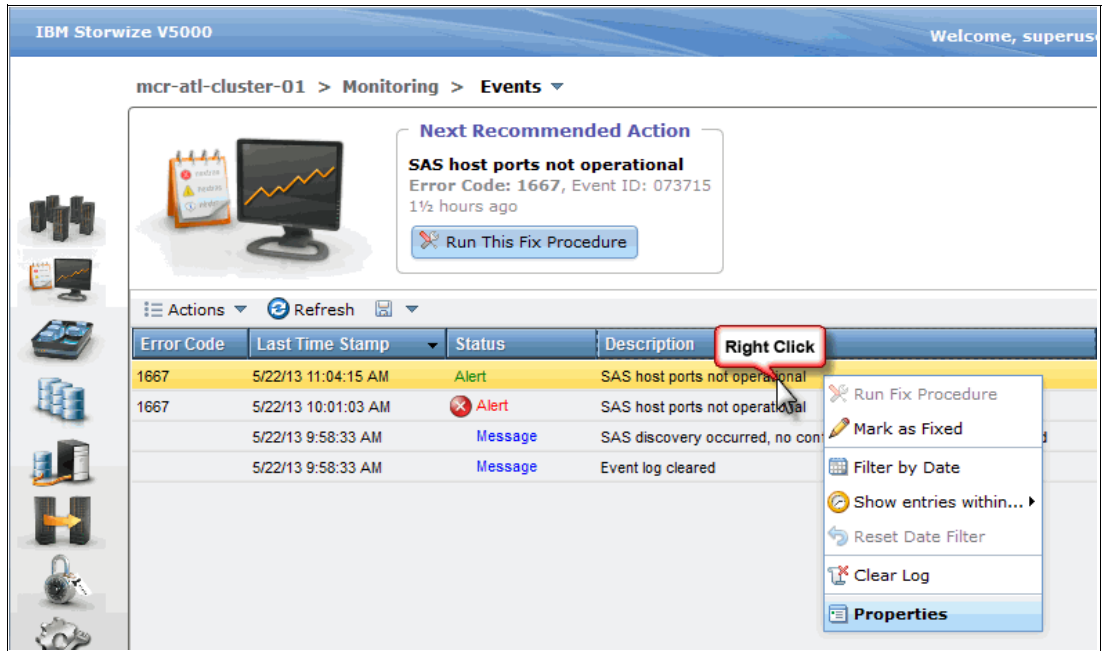


Figure 3-47 Selecting event properties

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

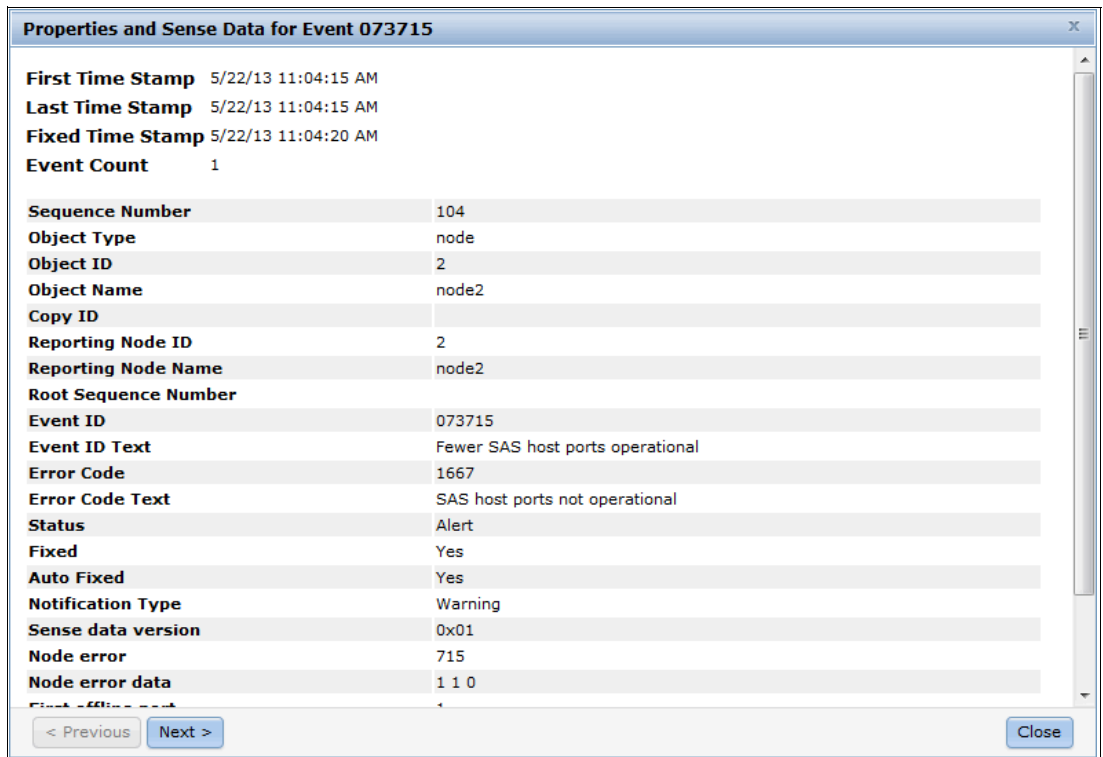


Figure 3-48 Event properties and sense data

Show events entries within

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

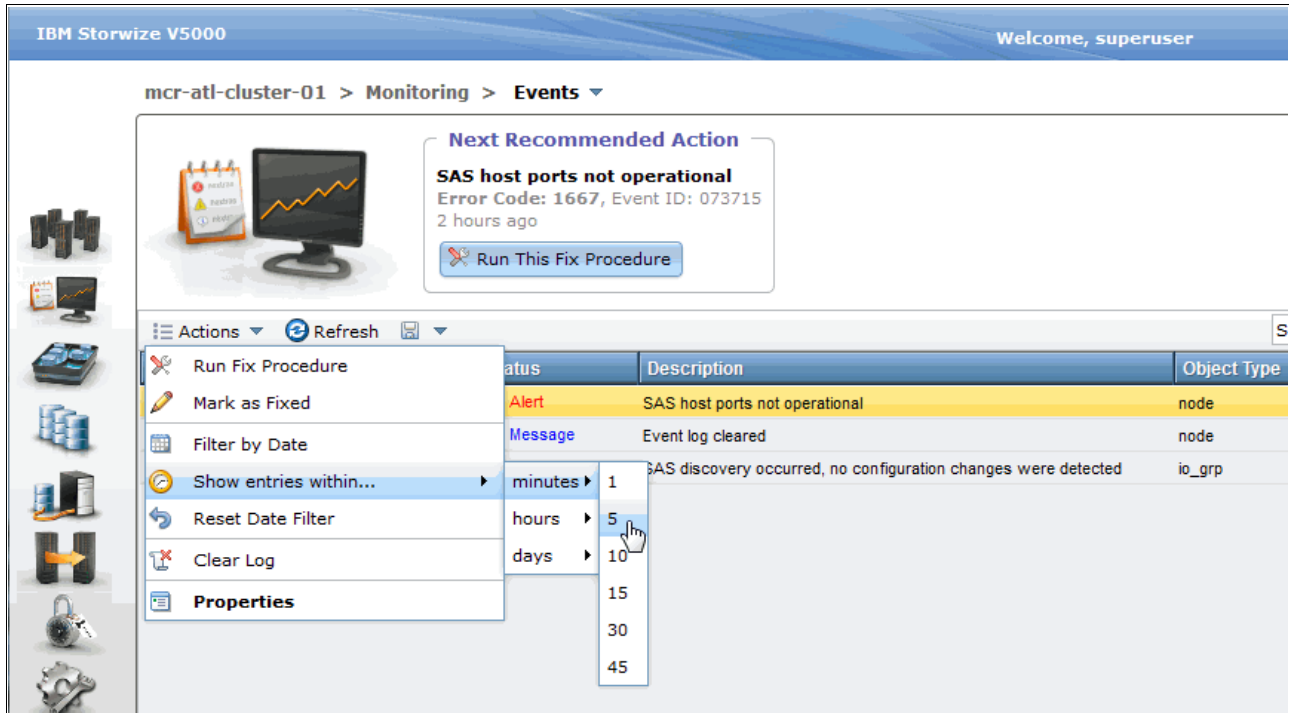


Figure 3-49 Showing events within a set time

Saving events to a file

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

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

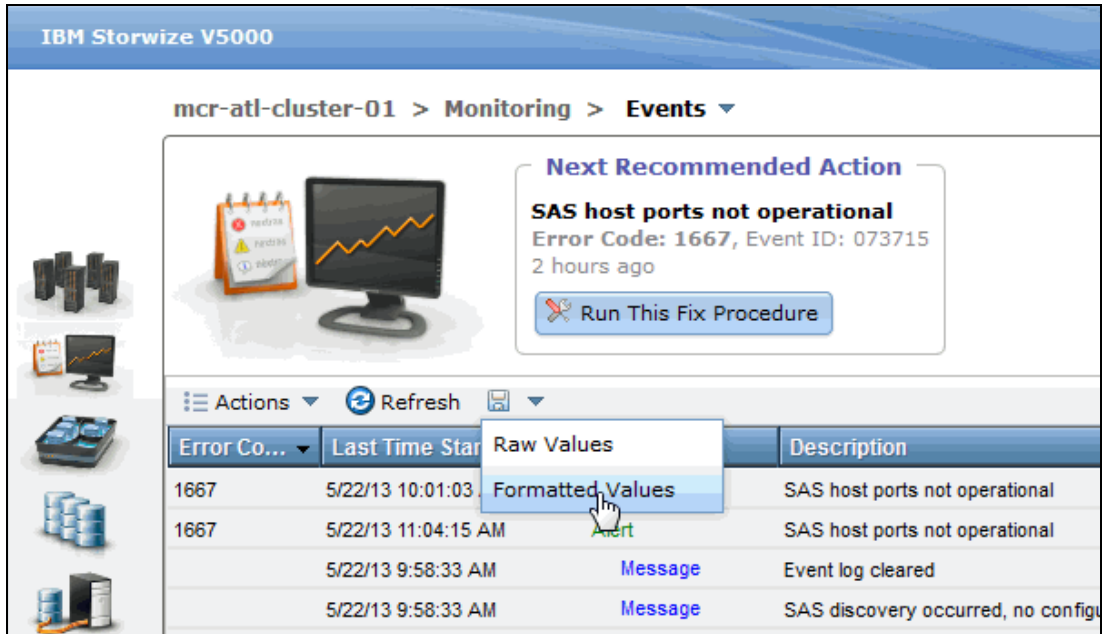


Figure 3-50 Saving events as formatted values

Performance panel

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

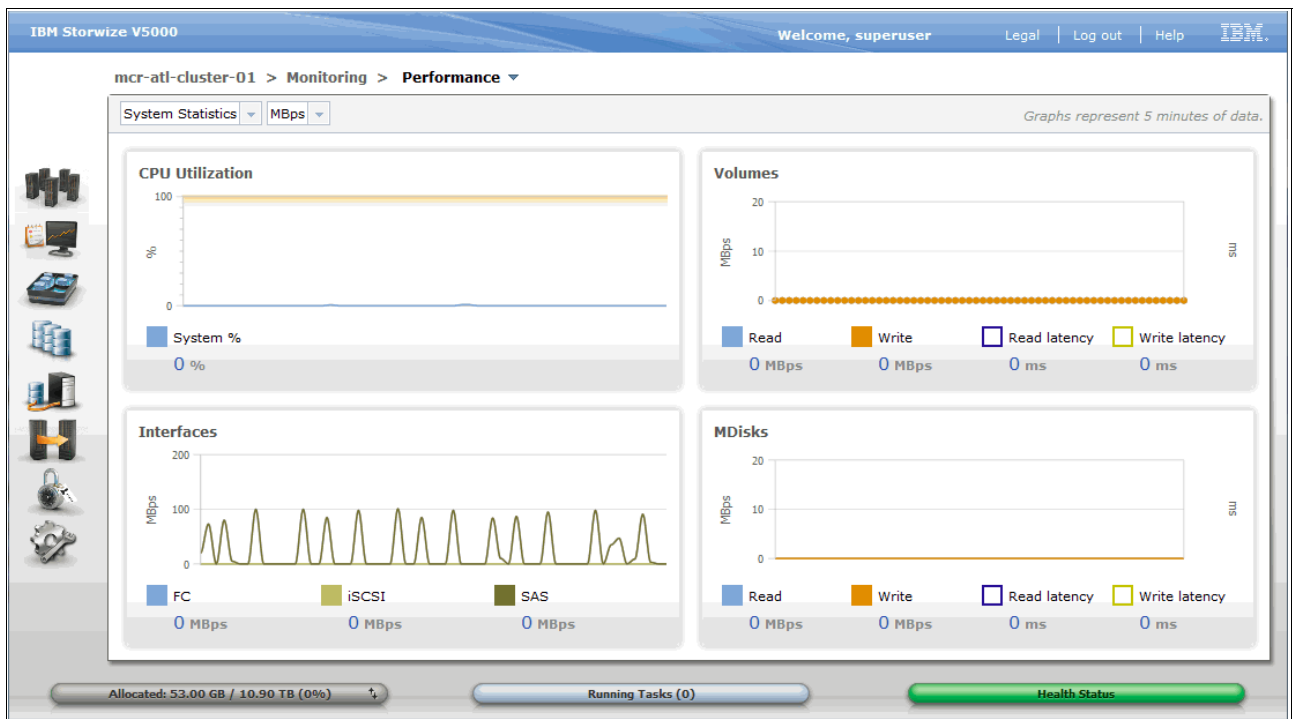


Figure 3-51 Performance panel

Custom-tailoring performance graphs

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

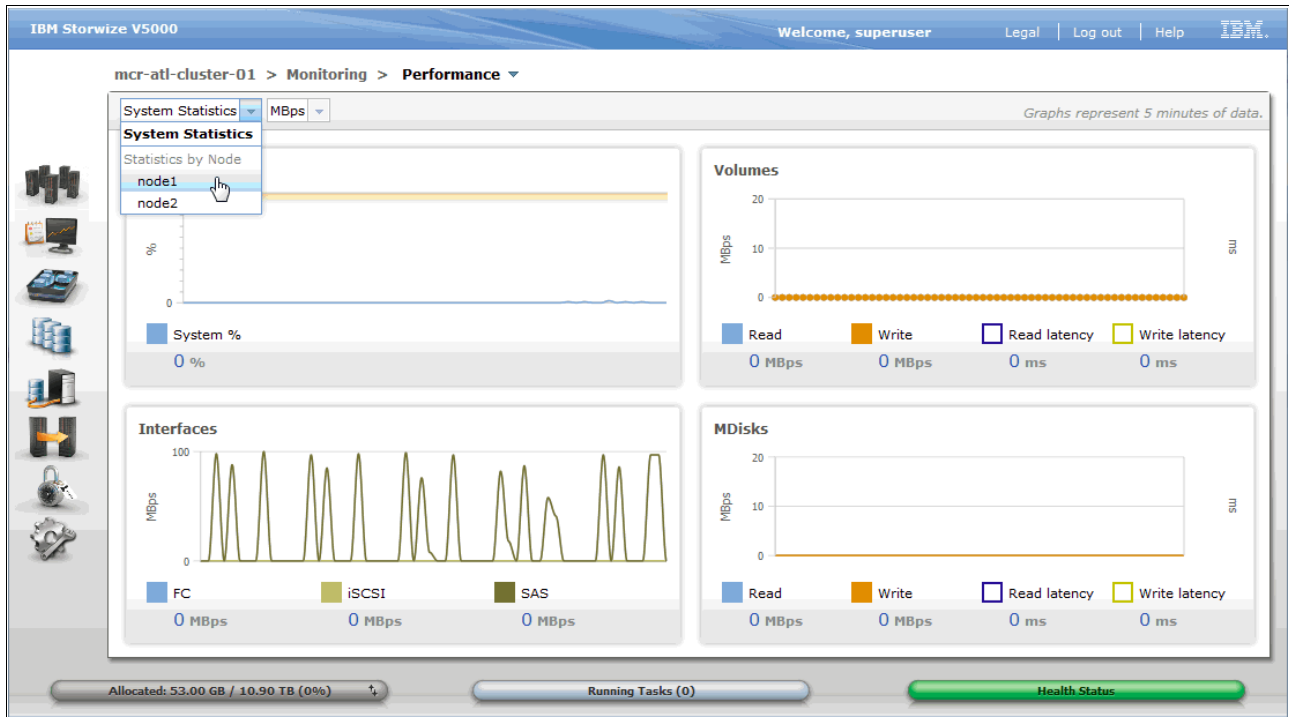


Figure 3-52 Graphs representing performance statistics of a single node

Performance peak value

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

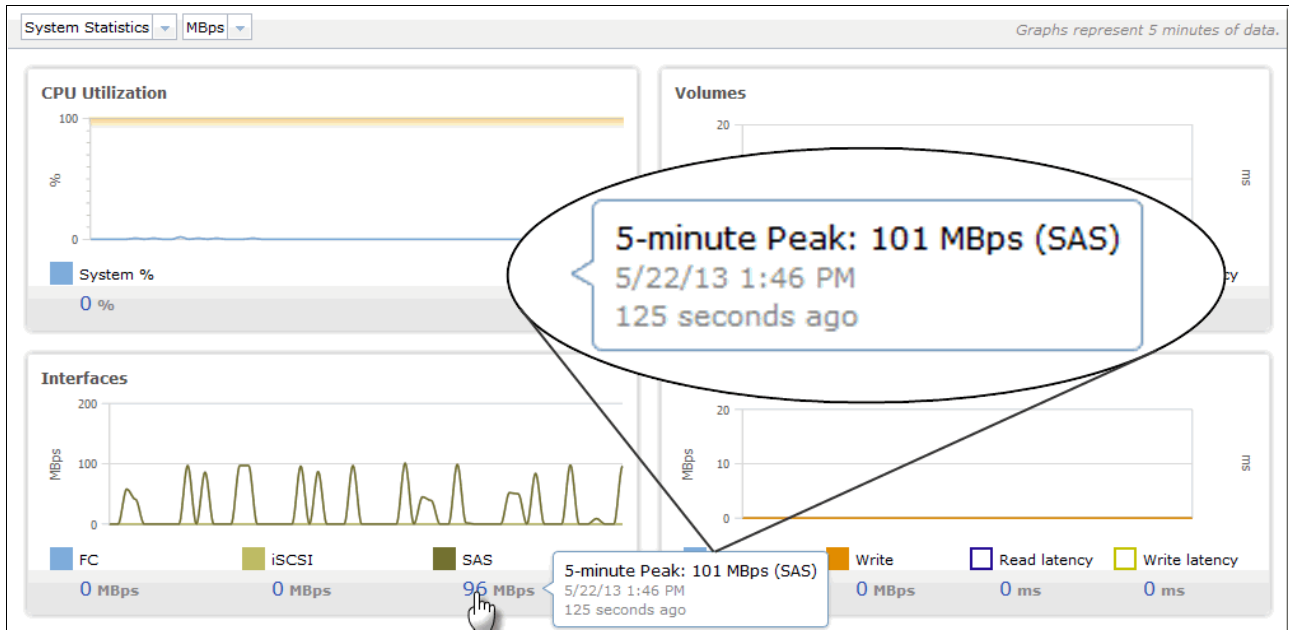


Figure 3-53 Peak SAS Interface usage value over the last 5 minutes

3.4.3 Pools menu

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

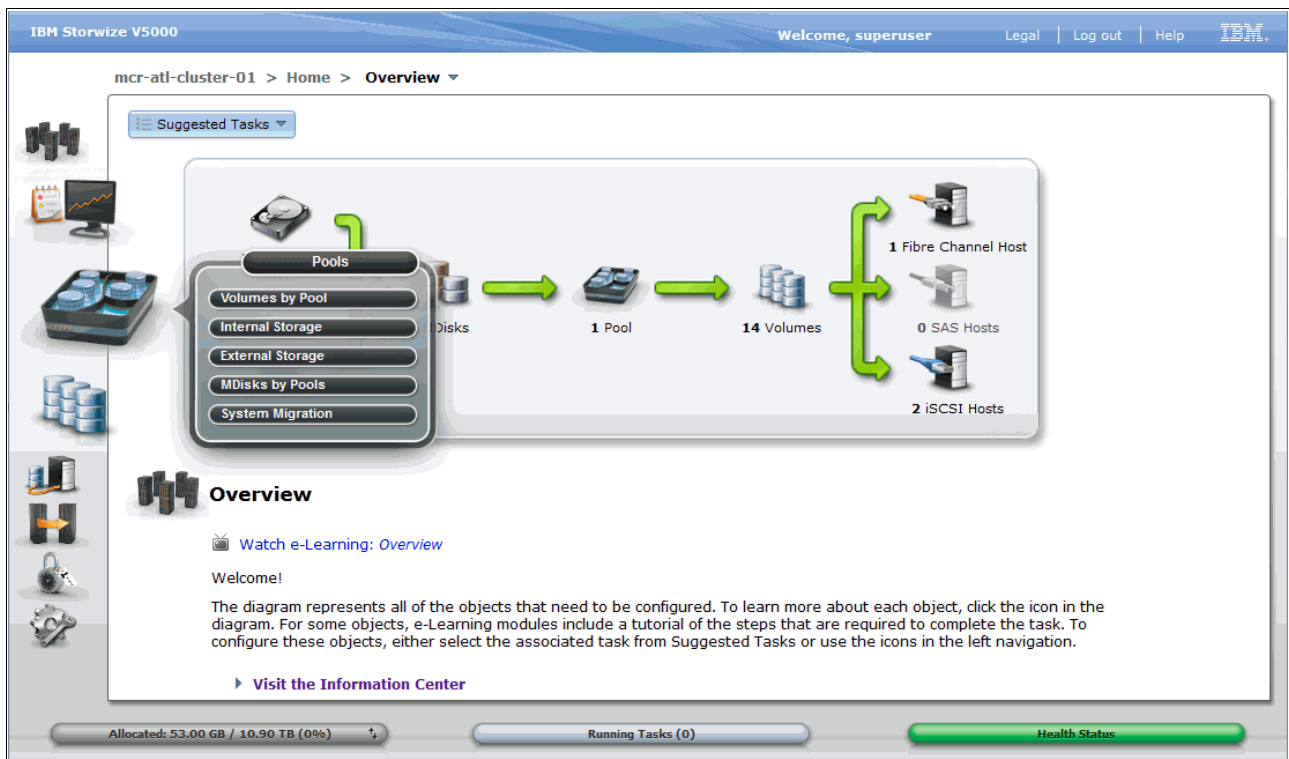


Figure 3-54 Pools menu

Volumes by Pool panel

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

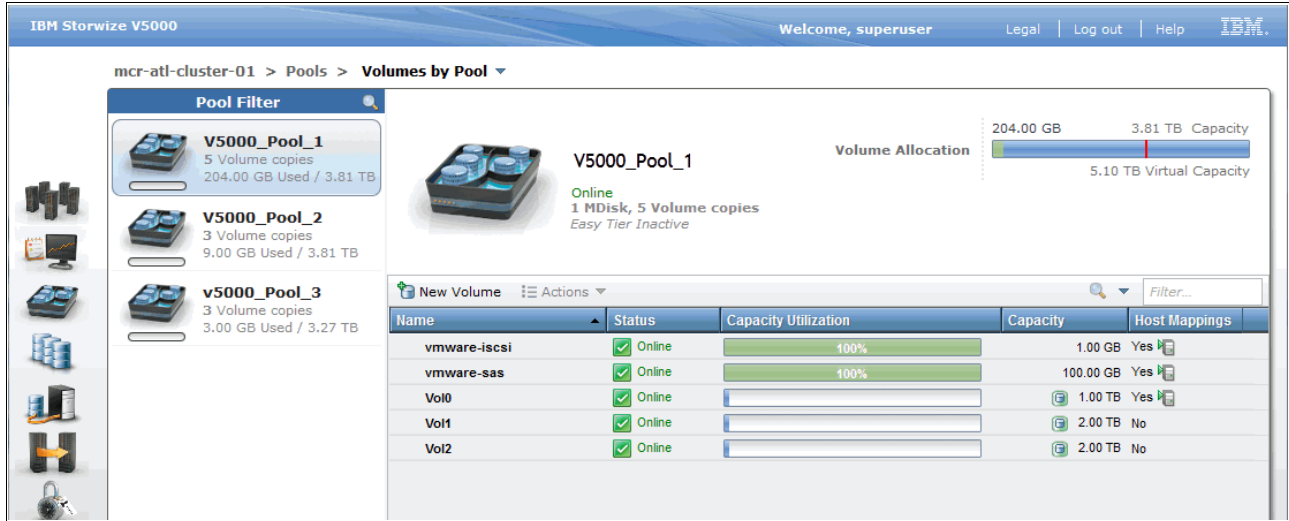


Figure 3-55 Volumes by Pools panel

Volume Allocation

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




Figure 3-56 Volume Allocation

Renaming pools

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

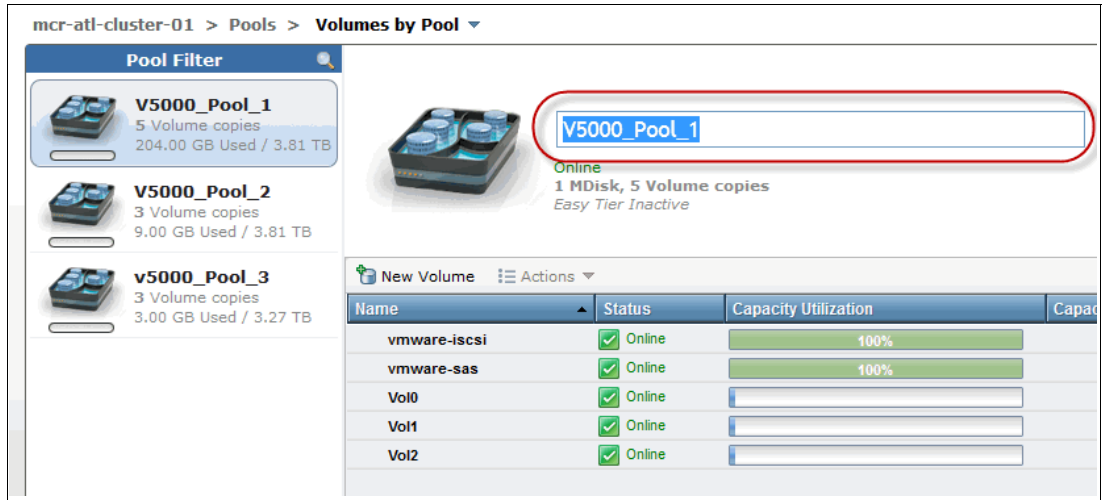


Figure 3-57 Renaming a pool

Changing pool icons

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

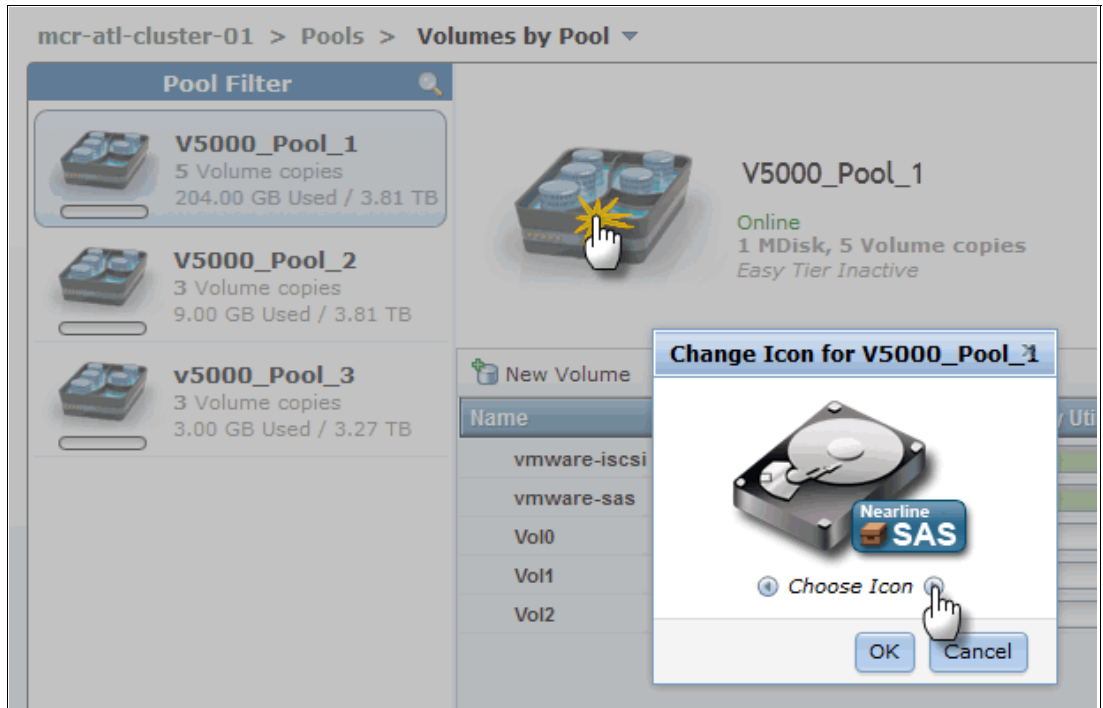


Figure 3-58 Changing a pool icon

Volume functions

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

The screenshot displays the 'Volumes by Pool' interface for 'mcr-atl-cluster-01'. It features a 'Pool Filter' on the left with three pools: V5000_Pool_1 (5 copies, 204.00 GB Used / 3.81 TB), V5000_Pool_2 (3 copies, 9.00 GB Used / 3.81 TB), and V5000_Pool_3 (3 copies, 3.00 GB Used / 3.27 TB). A 'Volume Allocation' bar shows 204.00 GB used and 3.81 TB capacity, with a total of 5.10 TB Virtual Capacity. A table lists five volumes: vmware-iscsi (1.00 GB, Yes), vmware-sas (100.00 GB, No), Vol0 (1.00 TB, No), Vol1 (2.00 TB, No), and Vol2 (2.00 TB, No). A context menu is open over the 'vmware-iscsi' volume, showing actions like 'Map to Host', 'Unmap All Hosts', 'View Mapped Hosts', 'Duplicate Volume', 'Move to Another I/O Group', 'Rename', 'Shrink', 'Expand', 'Migrate to Another Pool', 'Export to Image Mode', 'Delete', 'Volume Copy Actions', and 'Properties'. The bottom status bar shows 'Allocated: 216.00 GB / 10.90 TB (2%)', 'Running Tasks', and 'Health Status'.

Name	Status	Capacity Utilization	Capacity	Host Mappings
vmware-iscsi	Online		1.00 GB	Yes
vmware-sas	Online		100.00 GB	No
Vol0	Online		1.00 TB	No
Vol1	Online		2.00 TB	No
Vol2	Online		2.00 TB	No

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

Internal Storage panel

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




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

Drive actions

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

Drive properties

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




Figure 3-61 Drive properties

Configure internal storage wizard

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

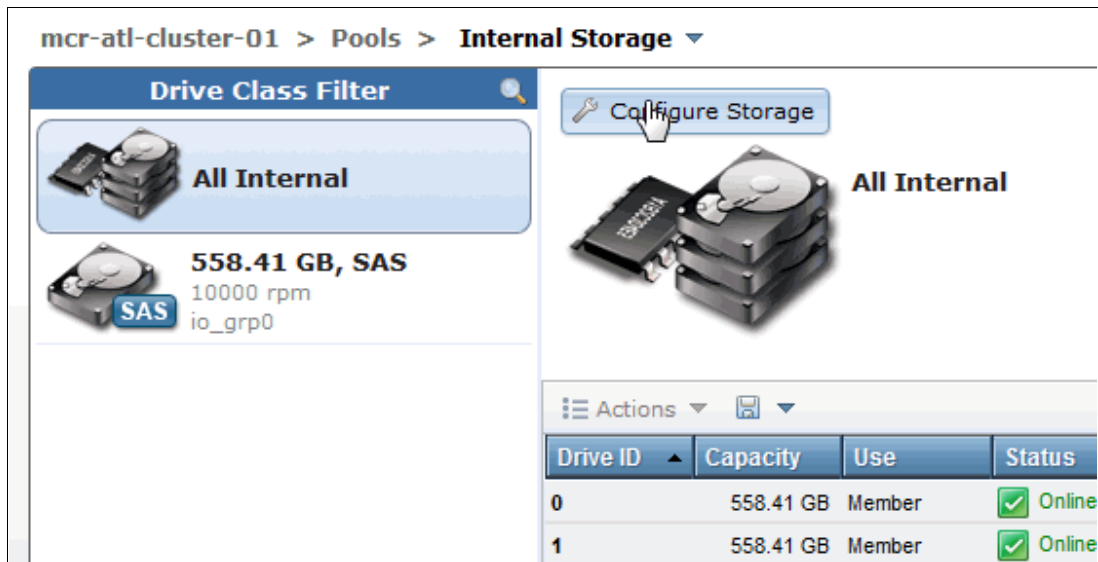


Figure 3-62 Internal Storage panel

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

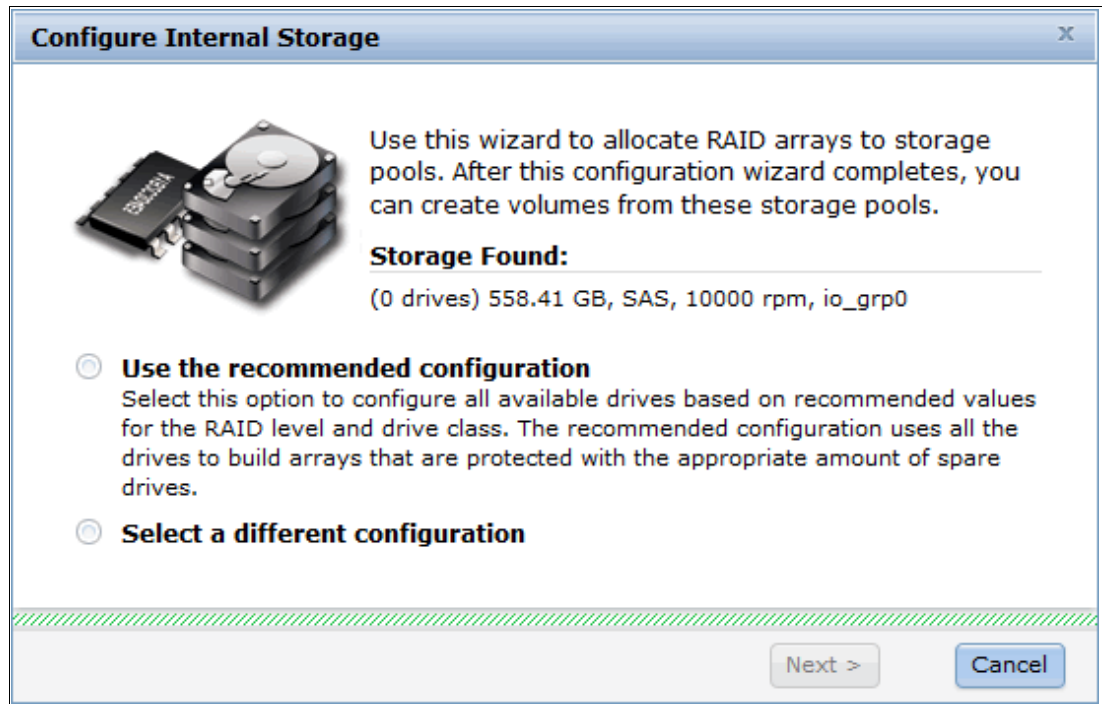


Figure 3-63 Configure Internal Storage wizard: Step 1

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

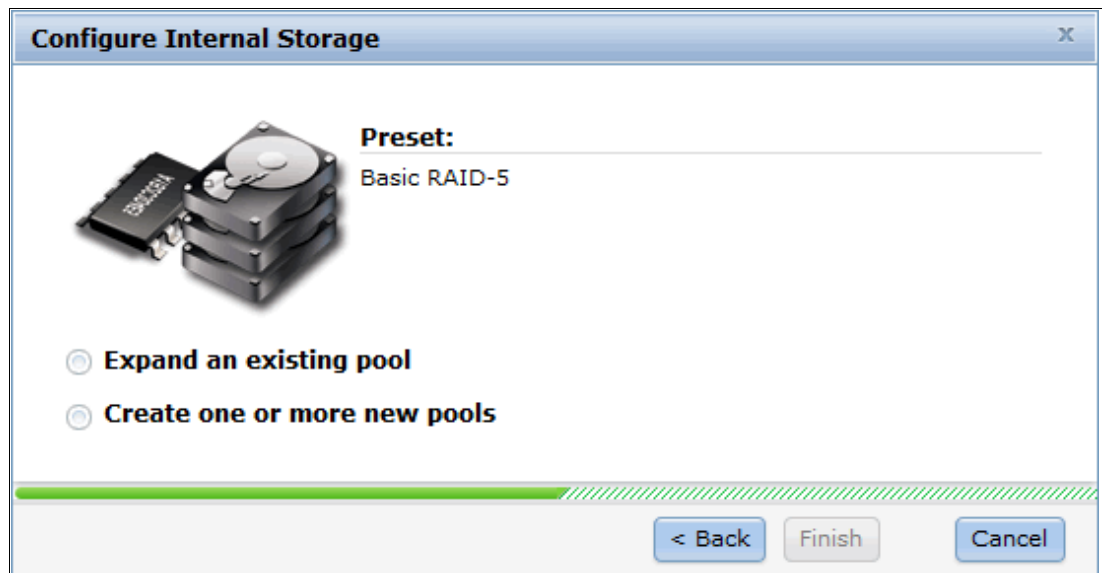


Figure 3-64 Configuring Internal Storage wizard: Step 2

MDisks by Pool panel

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

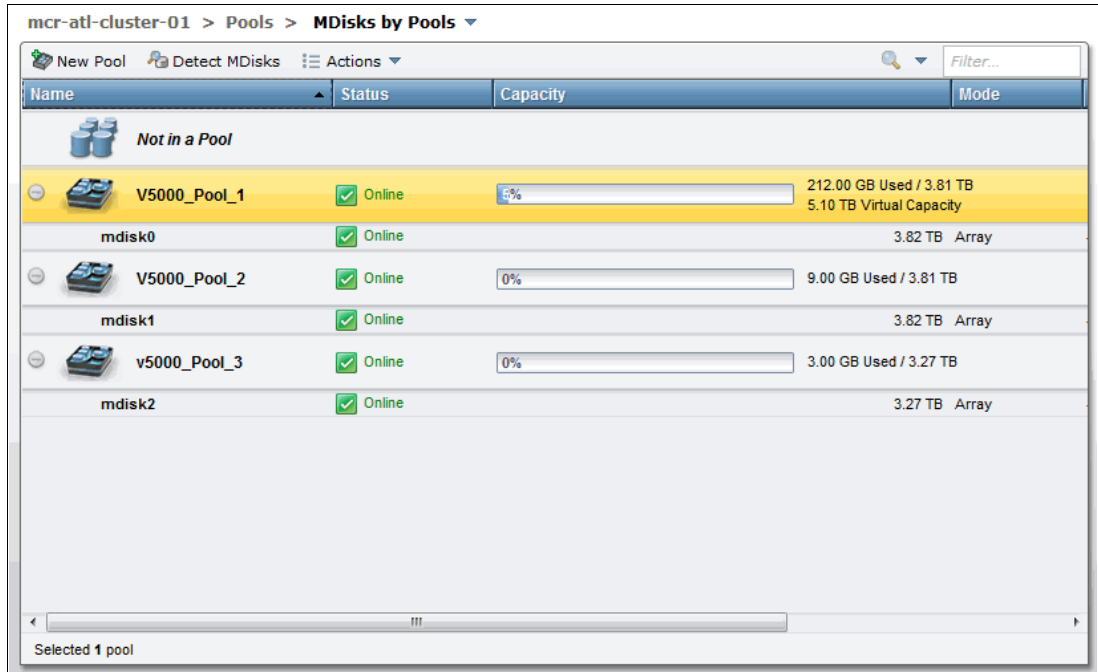


Figure 3-65 MDisks by Pool panel

Pool actions

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




Figure 3-66 Pool actions

RAID actions

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

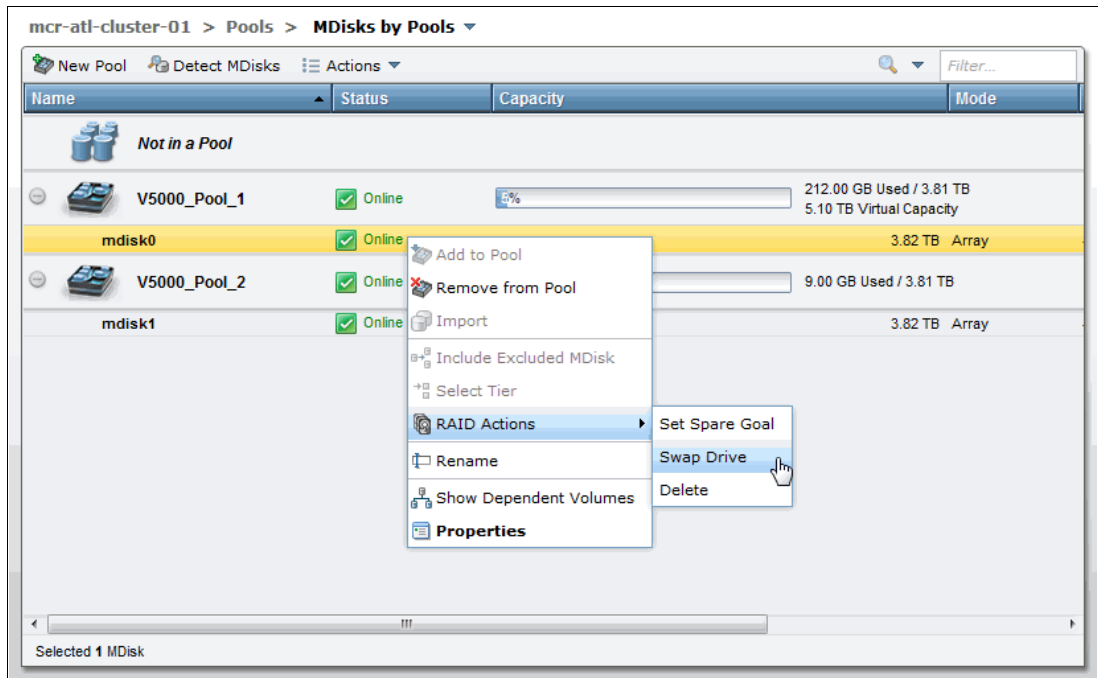


Figure 3-67 RAID actions menu

System Migration panel

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

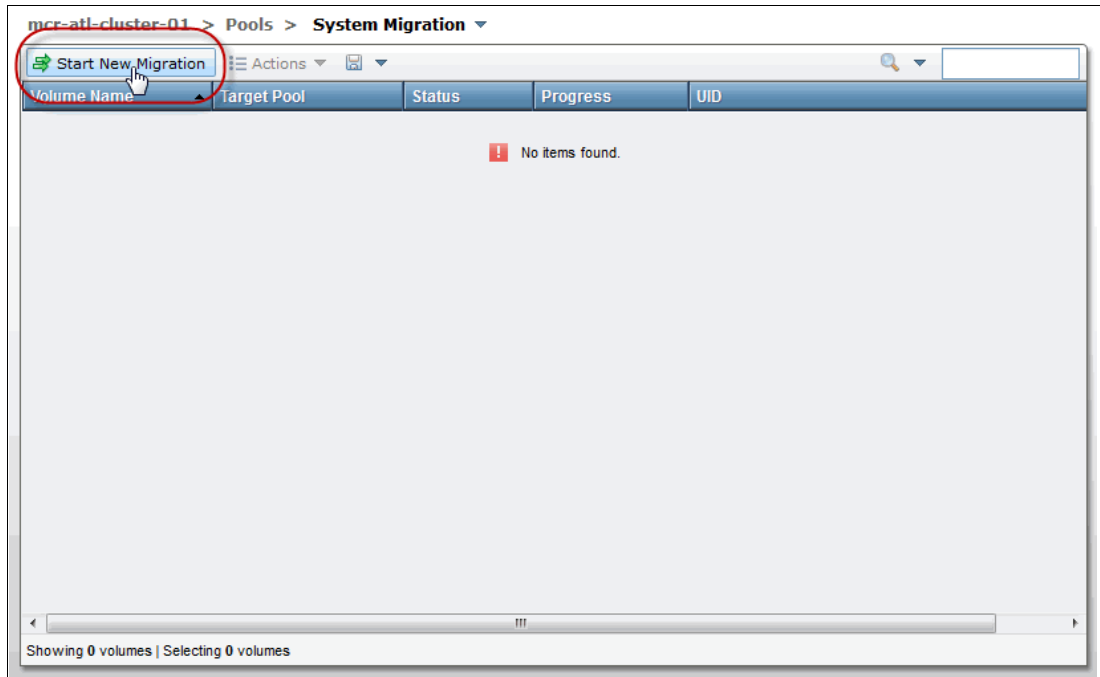


Figure 3-68 System Migration panel

Storage Migration wizard

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

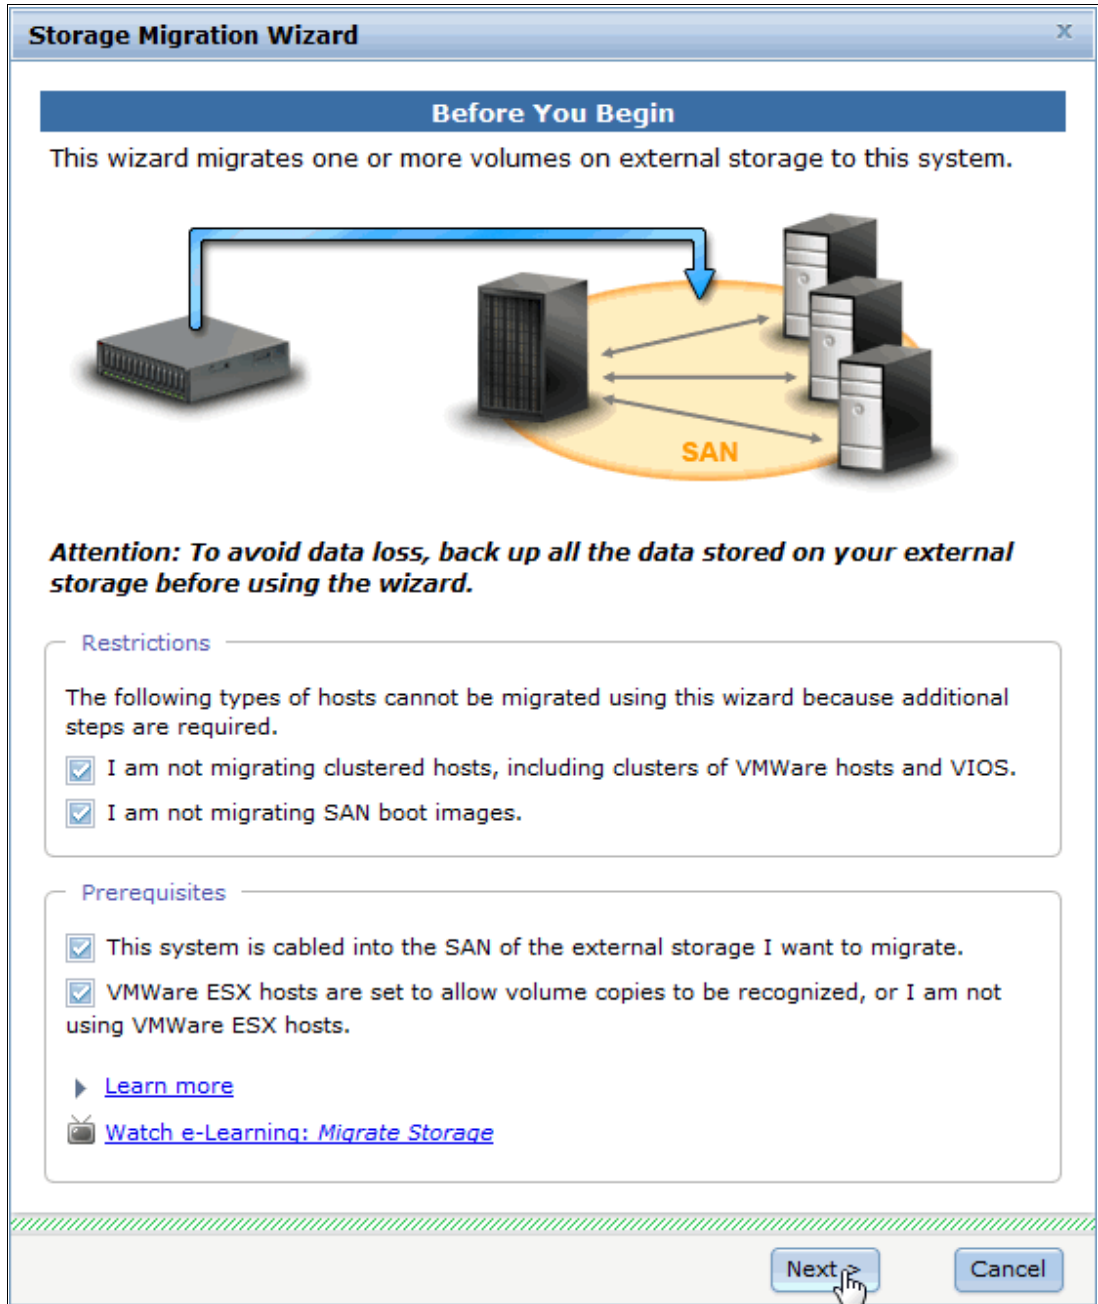


Figure 3-69 Storage Migration wizard

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

3.4.4 Volumes menu

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

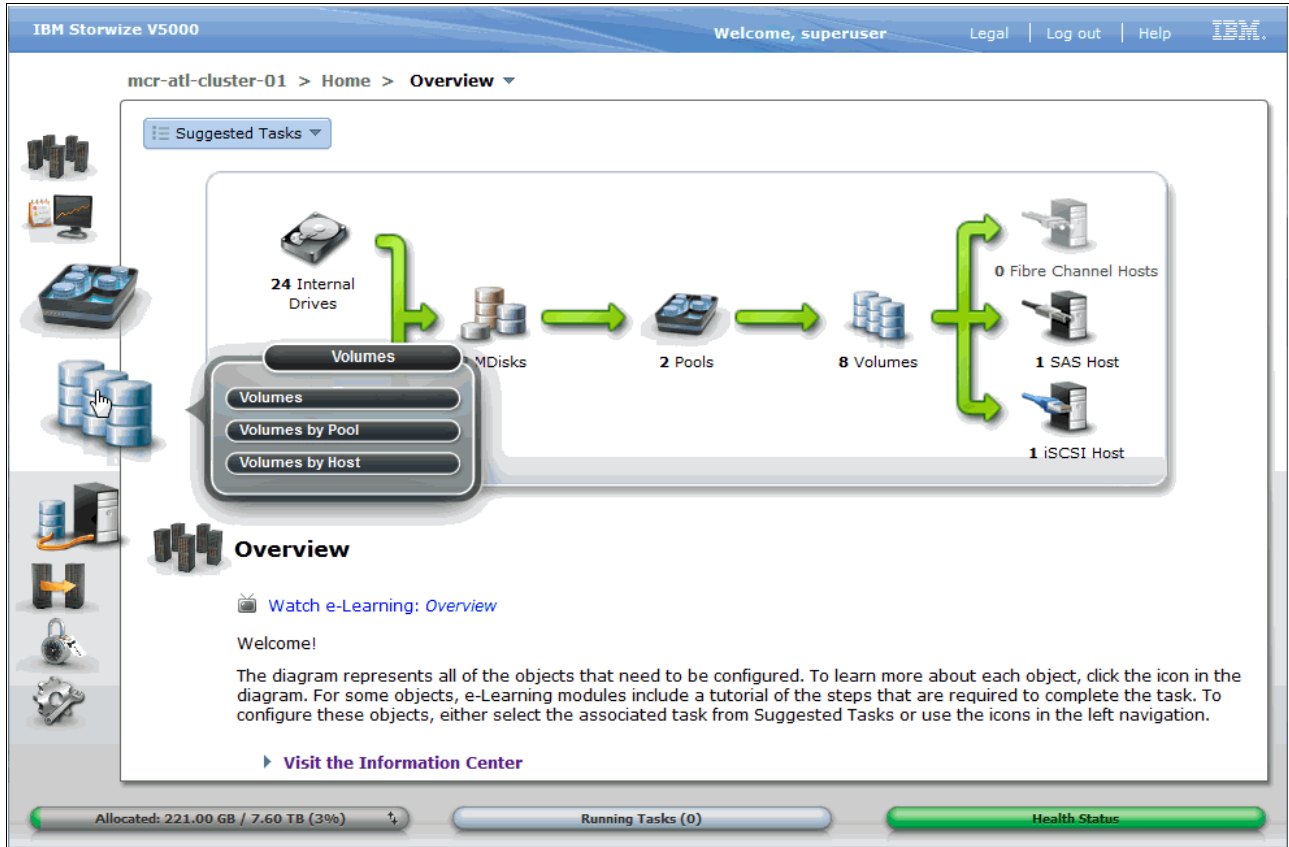


Figure 3-70 Selecting the Volumes menu

Volumes panel

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

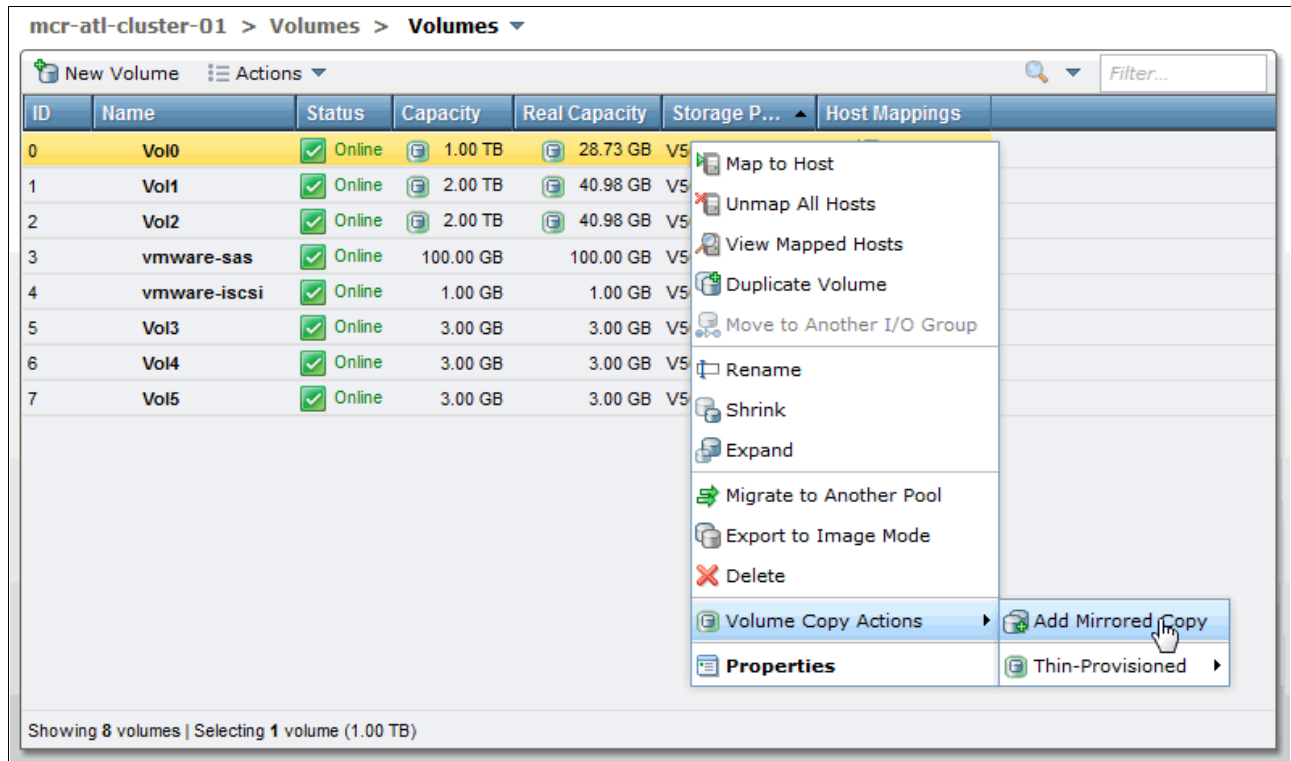


Figure 3-71 Volumes panel

Volume actions

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

Create new volumes

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

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




Figure 3-72 New Volume panel

Creating multiple volumes

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

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

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

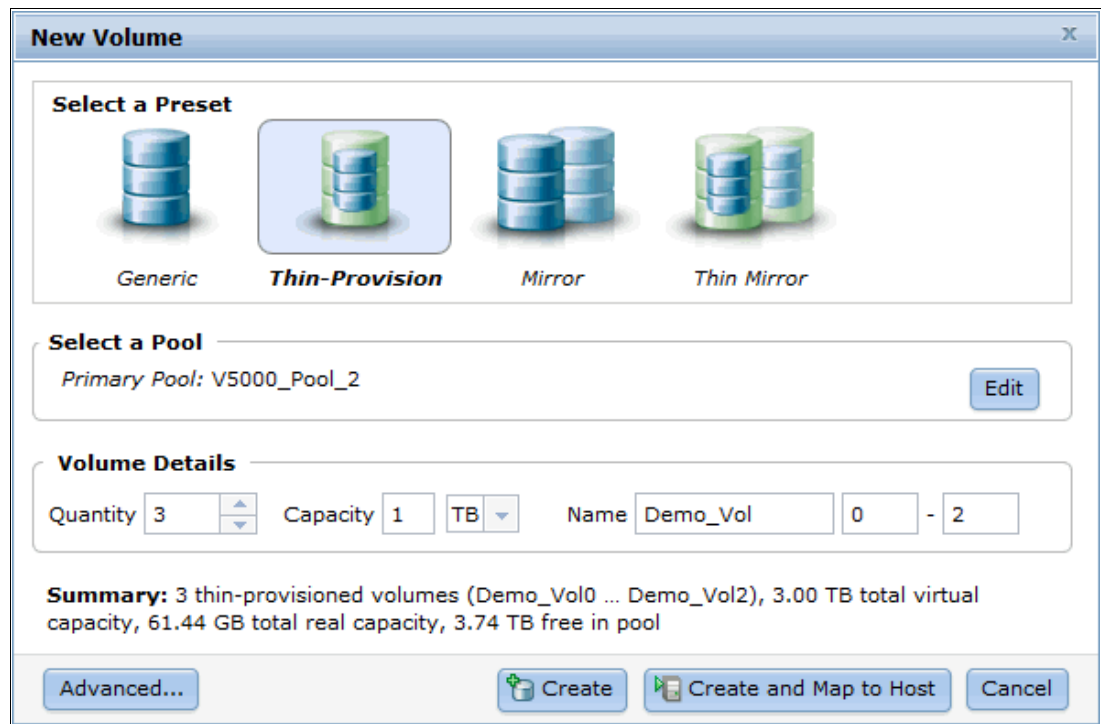


Figure 3-73 Creating multiple volumes quickly

Volume advanced settings

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

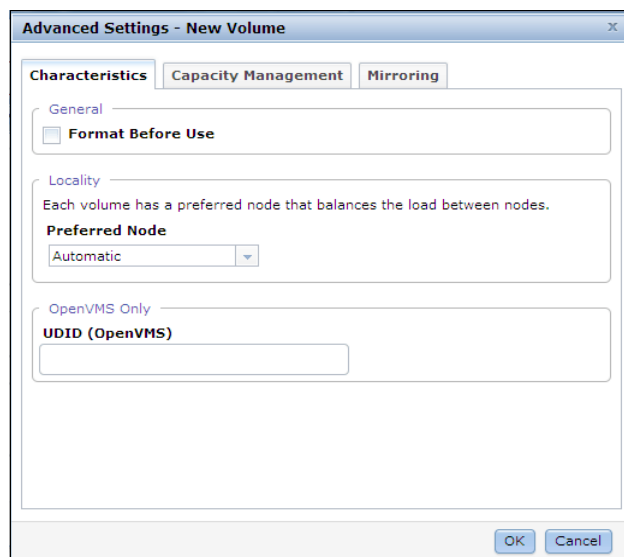


Figure 3-74 Advanced Settings panel

Volumes by Pool panel

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

Volumes by Host panel

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

3.4.5 Hosts menu

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

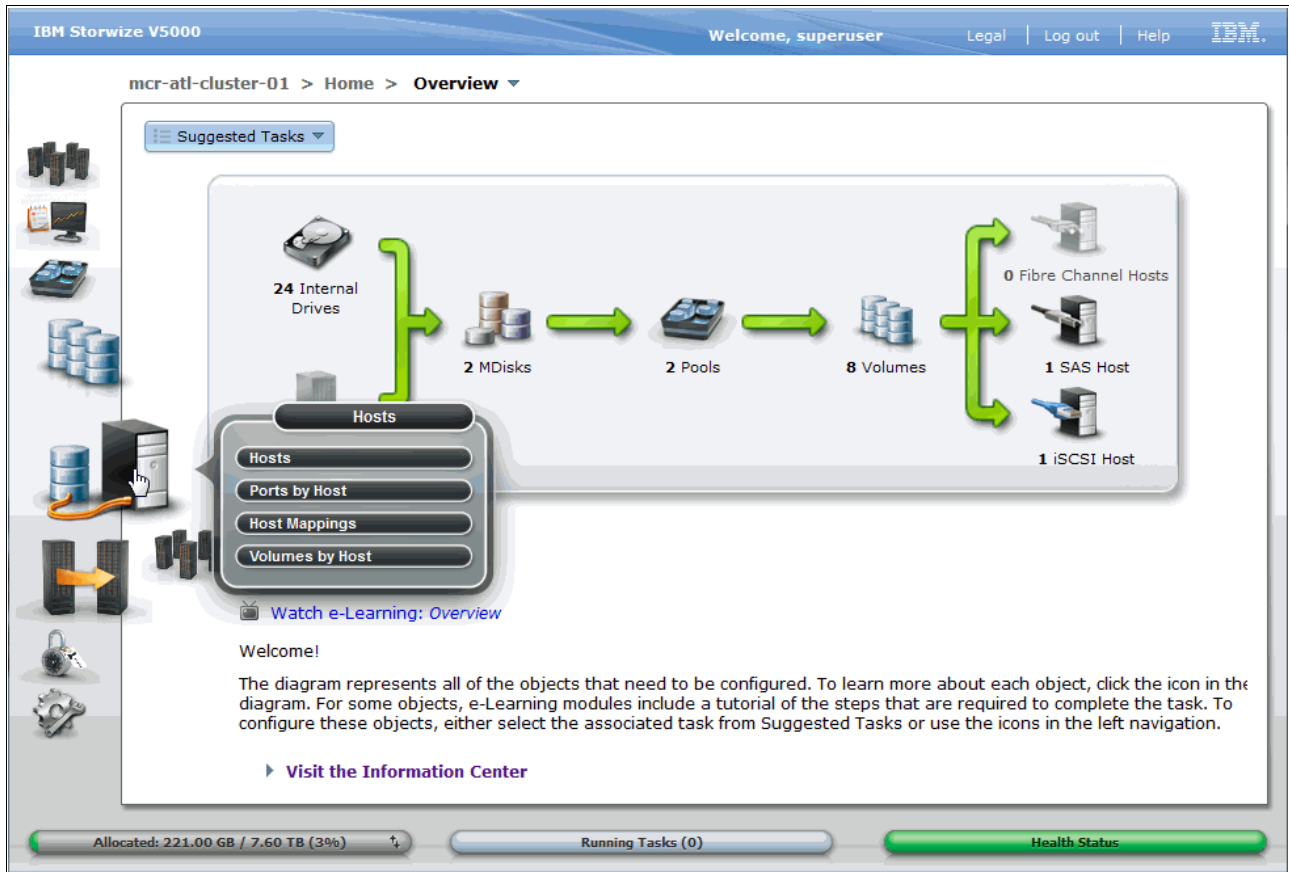


Figure 3-75 Selecting the Hosts menu

Hosts panel

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

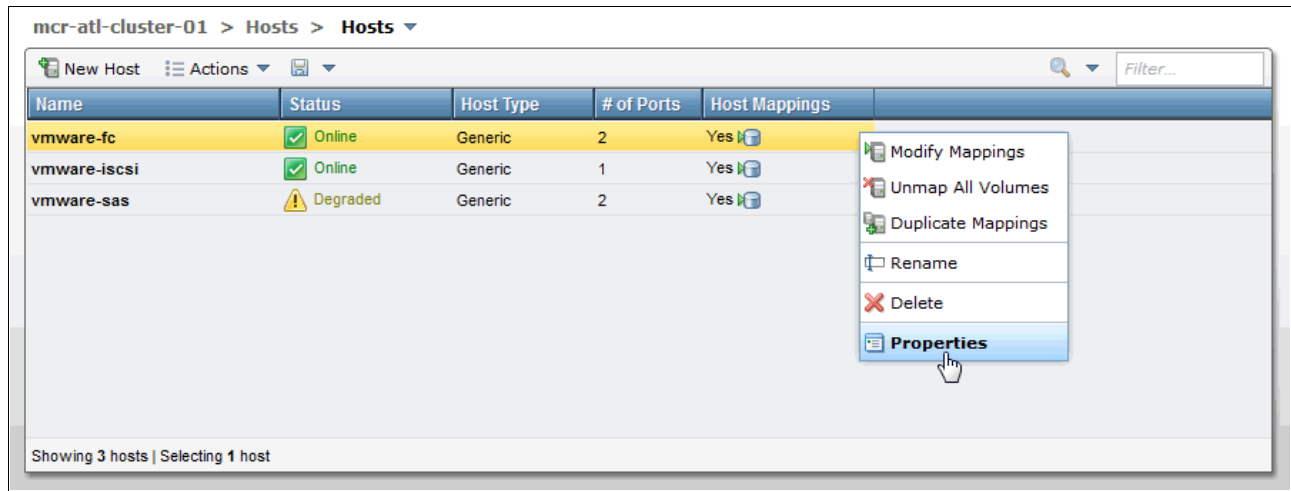


Figure 3-76 Hosts panel

Host Actions

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

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

Creating a host

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

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



Figure 3-77 Choose the Host Type panel

Ports by Host panel

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

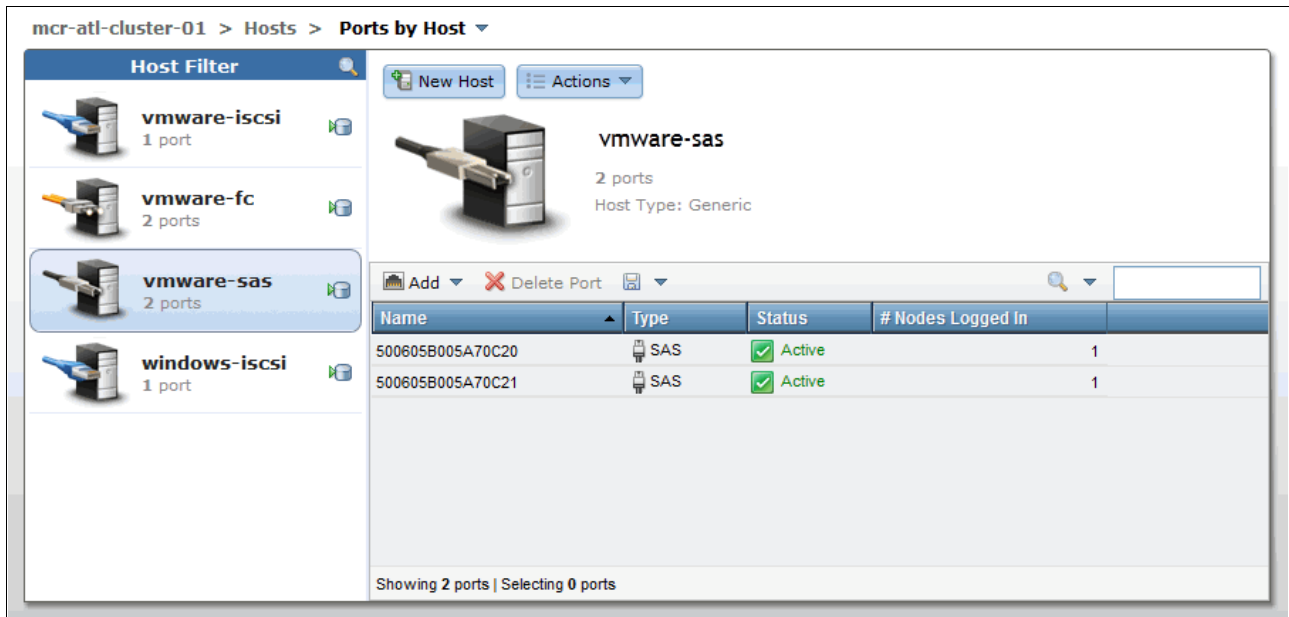


Figure 3-78 Ports by Host panel

Host Mappings panel

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

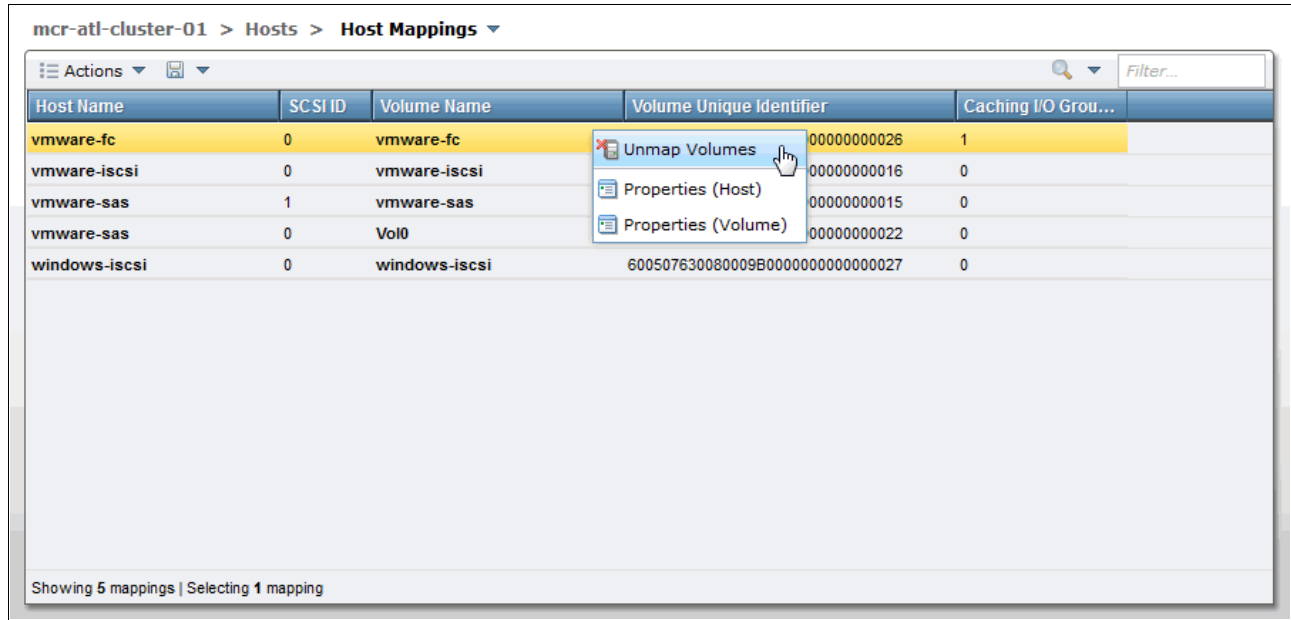


Figure 3-79 Host Mappings panel

Volumes by Host panel

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

3.4.6 Copy Services menu

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

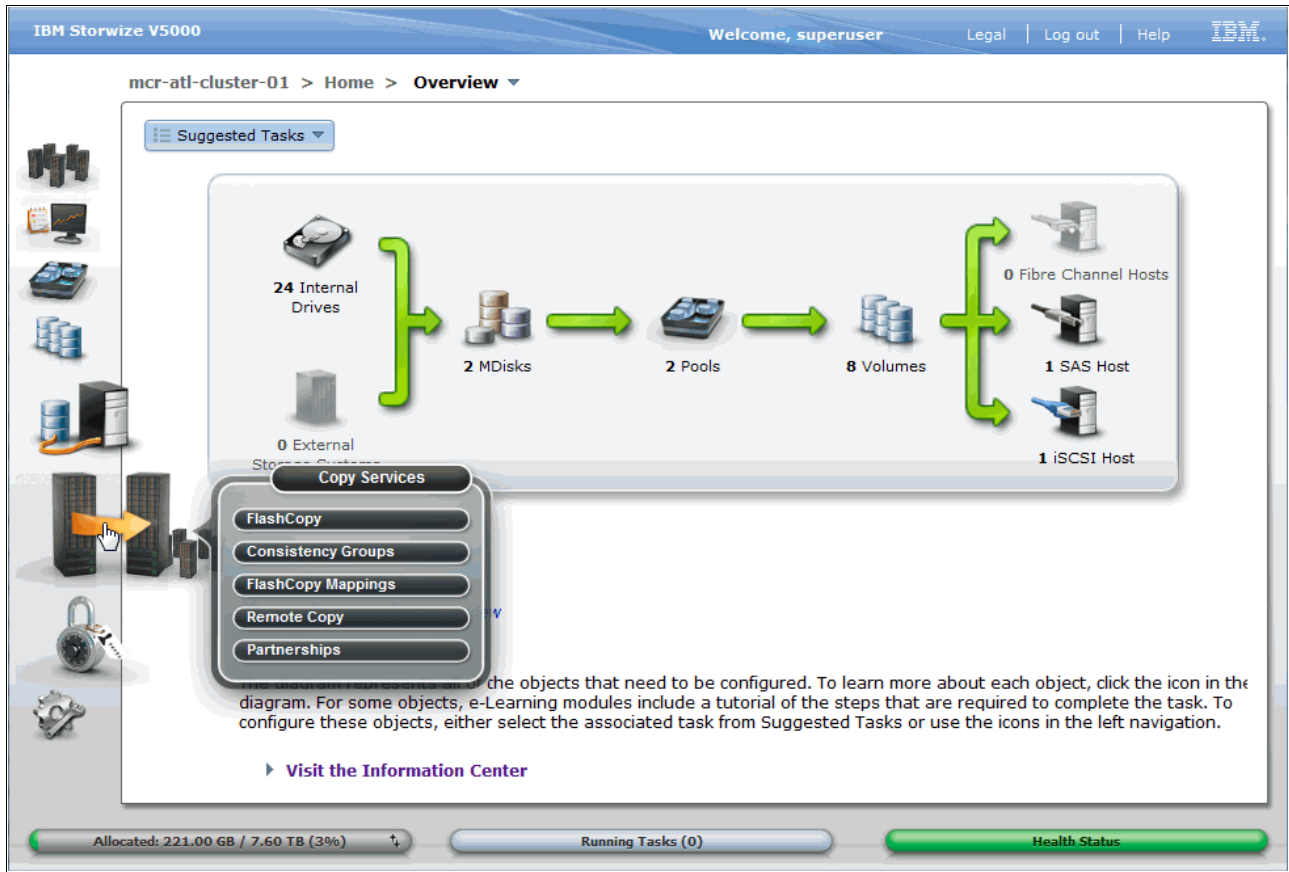


Figure 3-80 Copy Services menu

FlashCopy panel

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

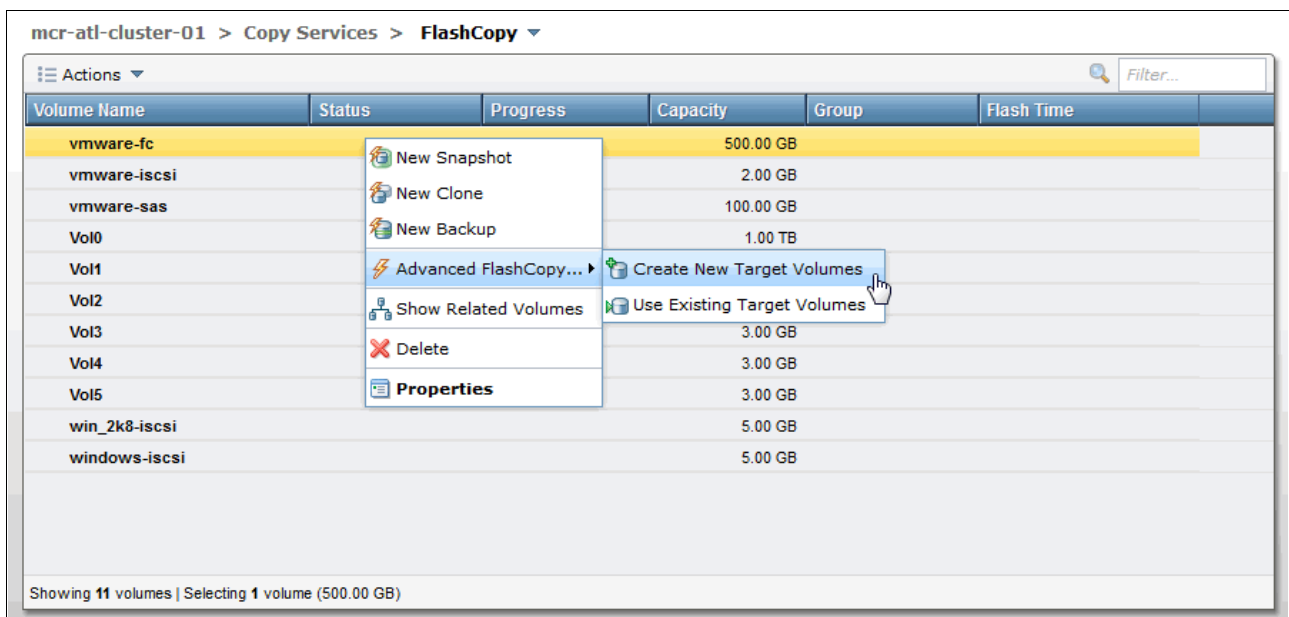


Figure 3-81 FlashCopy panel

FlashCopy actions

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

Consistency Groups panel

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

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

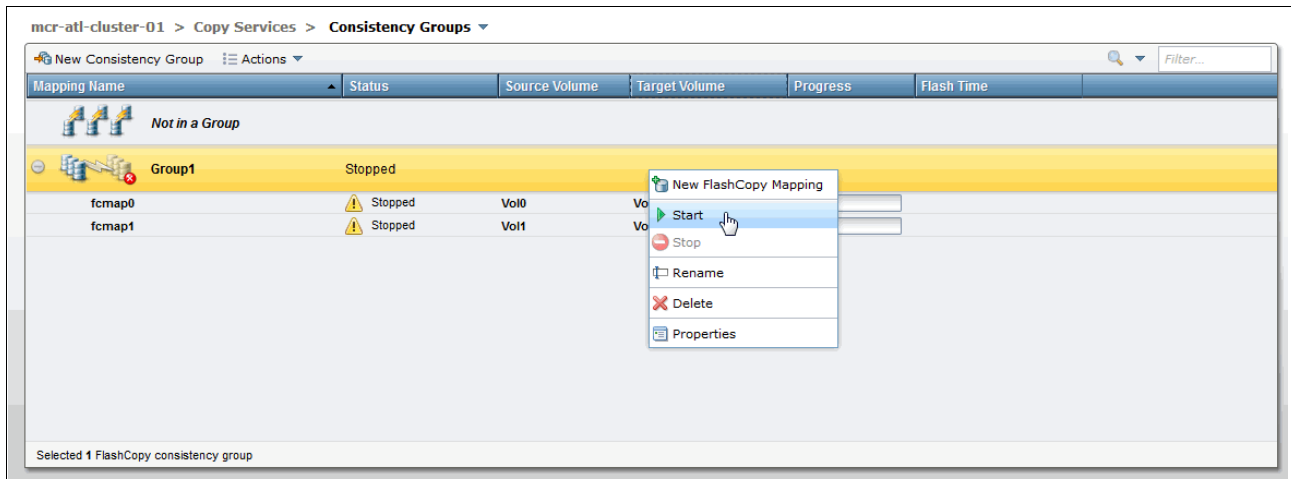


Figure 3-82 Consistency Groups panel

FlashCopy mappings panel

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

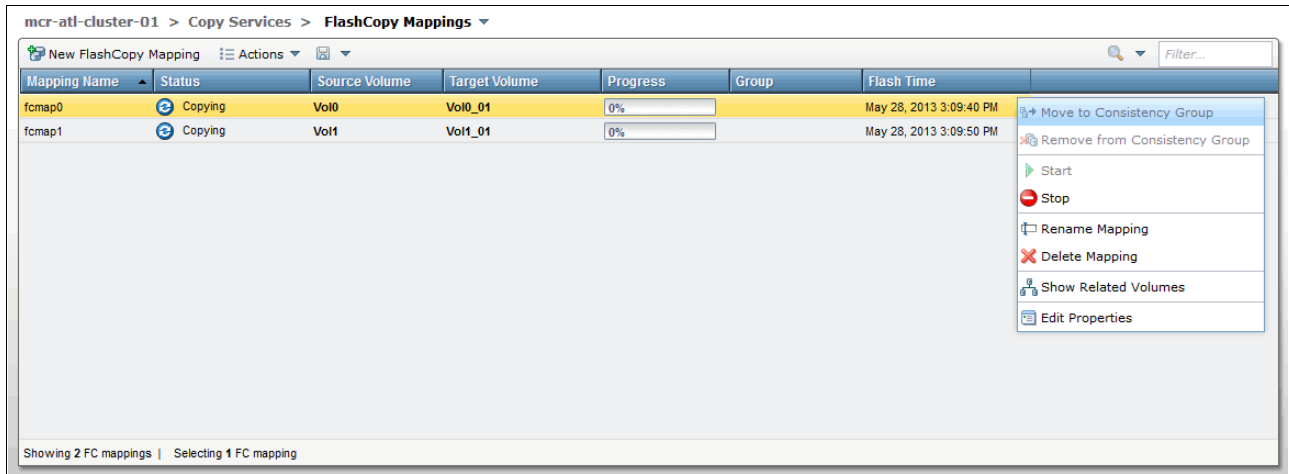


Figure 3-83 FlashCopy Mappings panel

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

Remote Copy panel

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

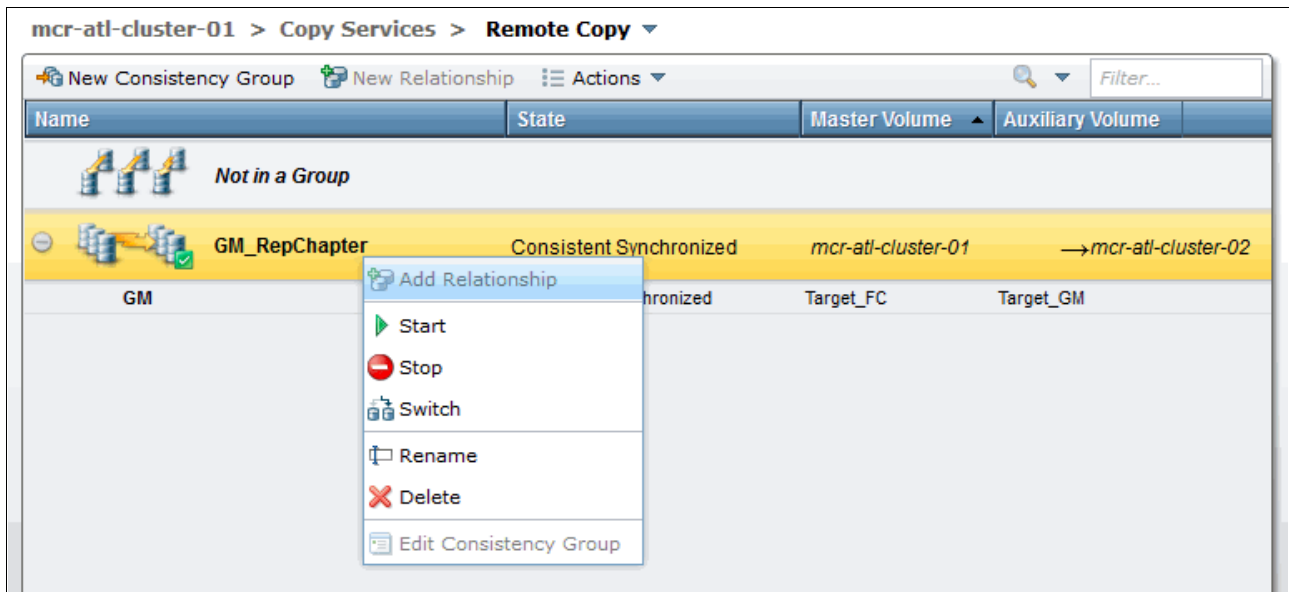


Figure 3-84 Remote Copy window

Partnerships panel

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

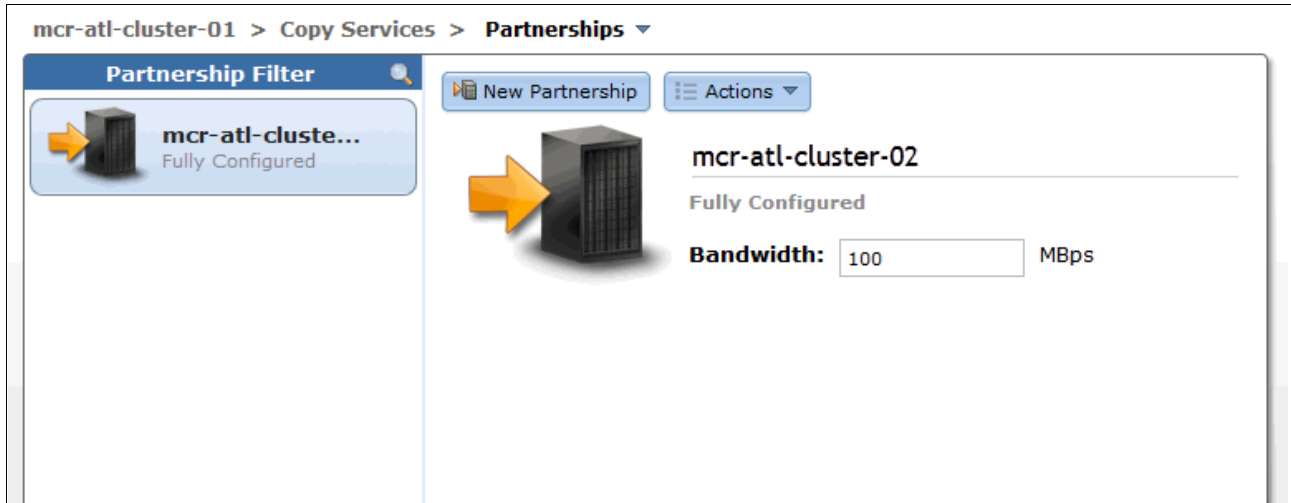


Figure 3-85 Partnerships window

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

3.4.7 Access menu

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

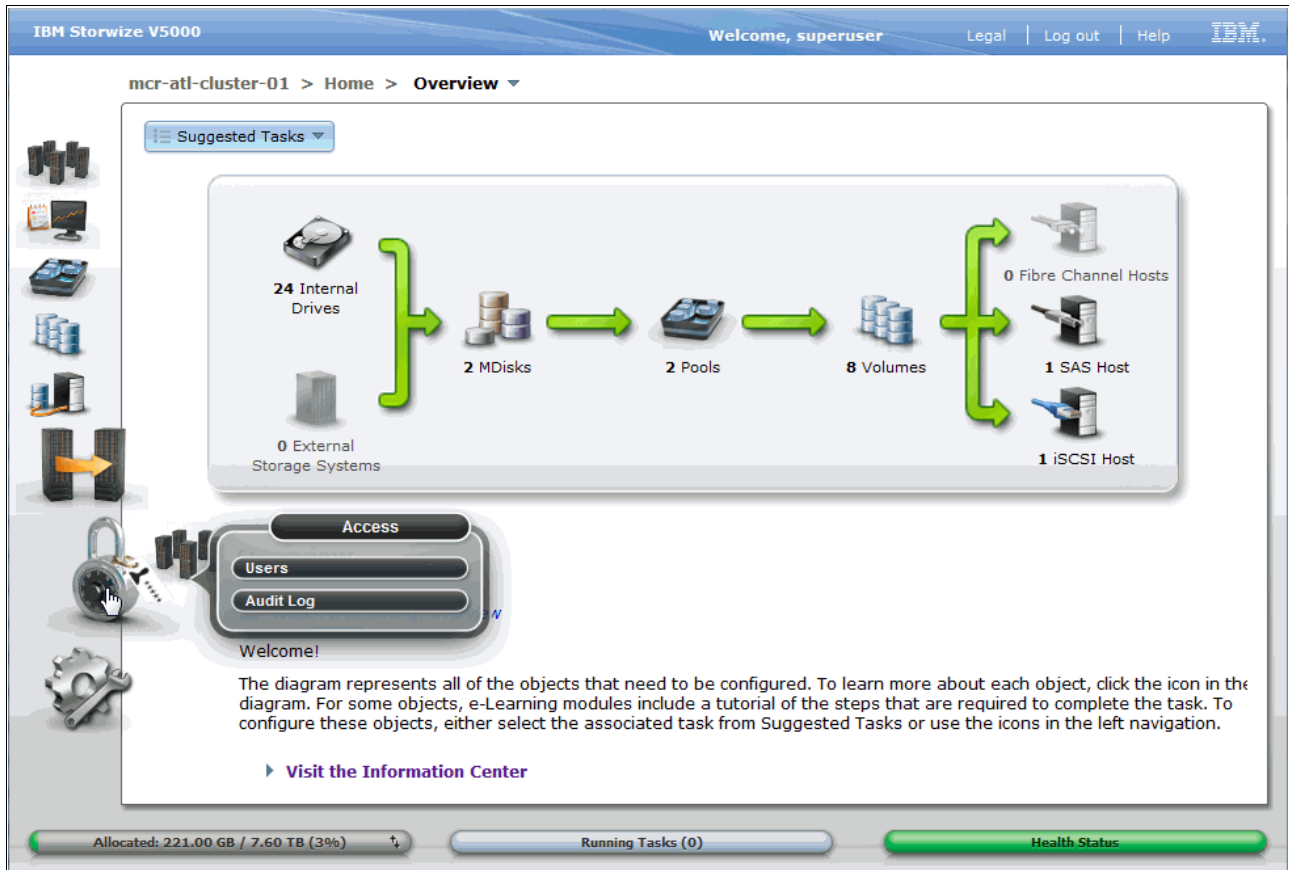


Figure 3-86 Access menu

Users panel

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

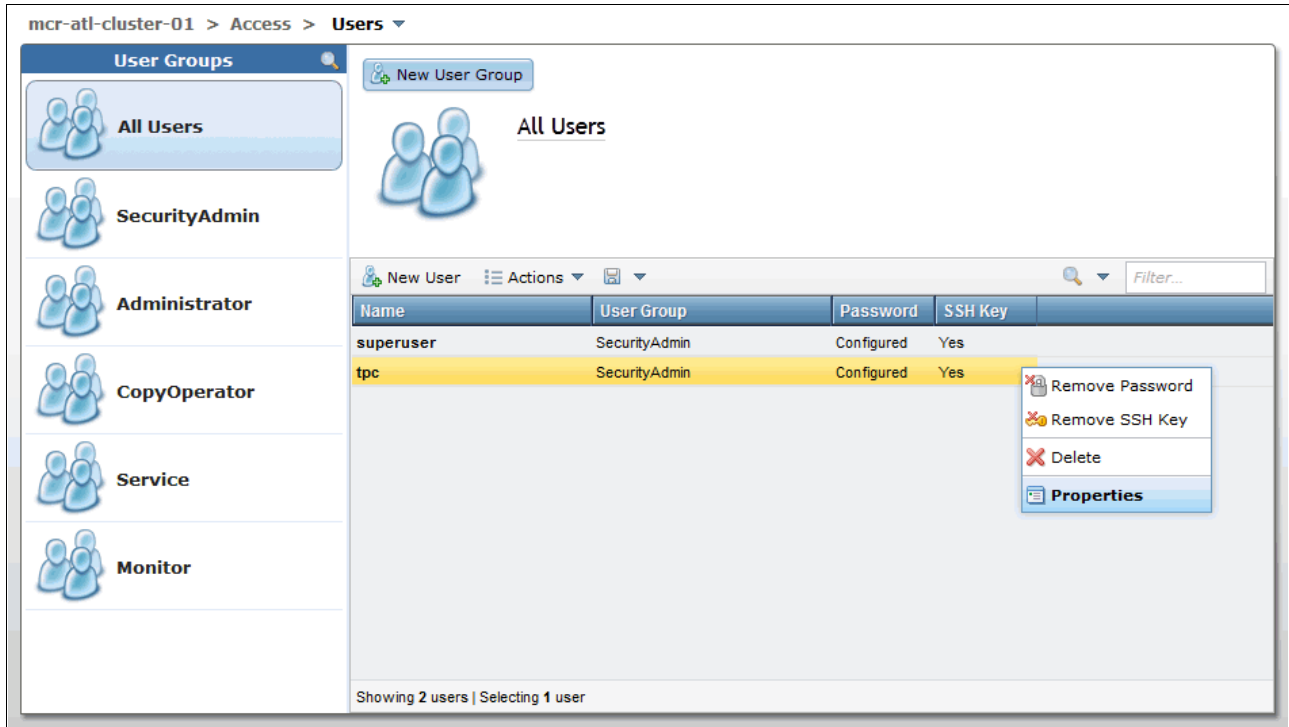


Figure 3-87 Users panel

Creating a user group

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

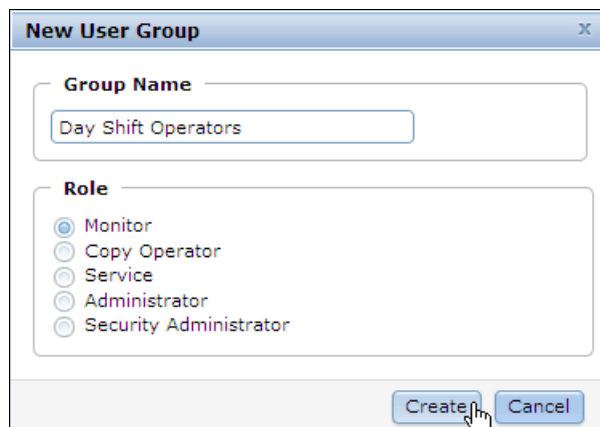


Figure 3-88 New User Group panel

Creating a user

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

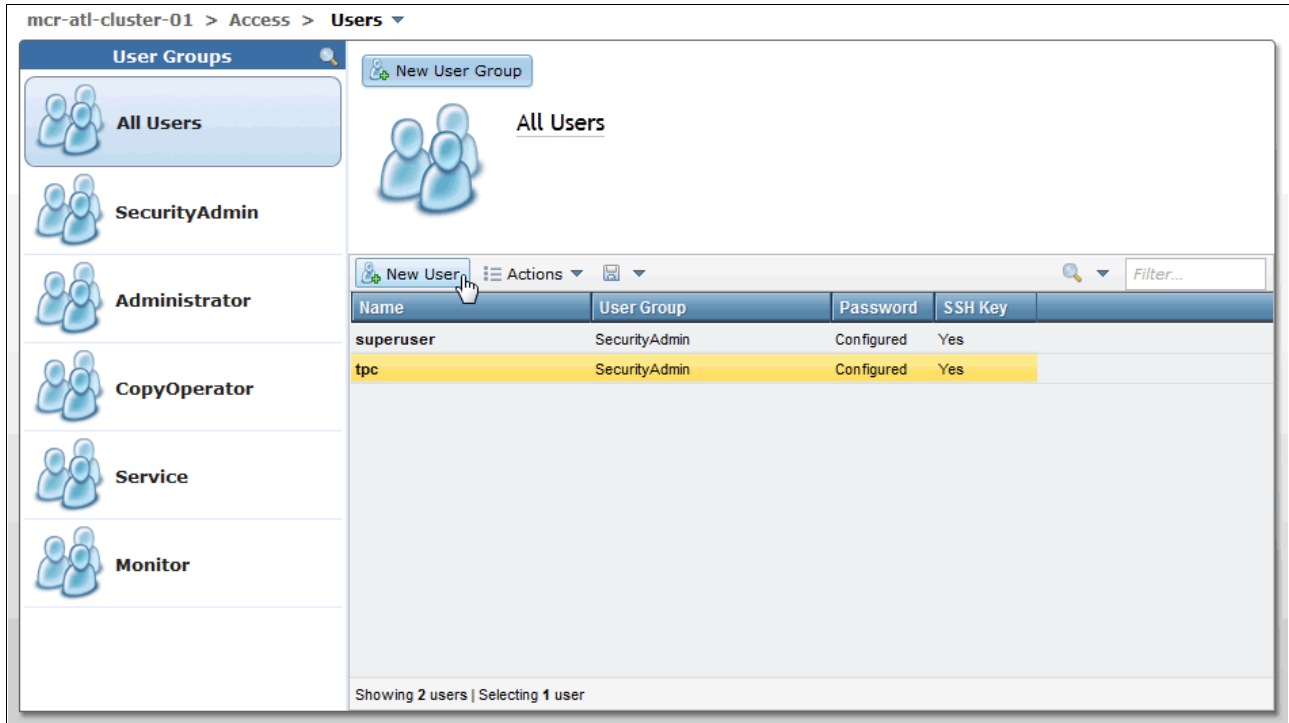


Figure 3-89 Users panel and the New User option

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

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

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

Figure 3-90 shows the New User panel.

Figure 3-90 New User panel

Audit Log panel

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

Date and Time	User Name	Command	Object ID
5/29/13 1:00:03 AM	superuser	svctask detectmdisk	
5/28/13 5:30:16 PM	superuser	svctask mkuser -name Test -usergrp 0 -password ##### -keyfile...	2
5/28/13 5:22:17 PM	superuser	svctask addvdiskcopy -mdiskgrp 0 5	
5/28/13 5:04:59 PM	superuser	svctask chmdiskgrp -name MultiTierPool 1	
5/28/13 4:54:31 PM	superuser	svctask mkarray -level raid0 -drive 24 -sparegoal 0 1	2
5/28/13 4:52:14 PM	superuser	svctask migratevdisk -vdisk Vol3 -mdiskgrp V5000_Pool_2 -copy 0	
5/28/13 4:52:00 PM	superuser	svctask migratevdisk -vdisk Vol2 -mdiskgrp V5000_Pool_2 -copy 0	
5/28/13 4:51:49 PM	superuser	svctask migratevdisk -vdisk Vol1 -mdiskgrp V5000_Pool_2 -copy 0	
5/28/13 4:45:56 PM	superuser	svctask rmdisk -mdisk mdisk3 -force 1	
5/28/13 4:44:12 PM	superuser	svctask migratevdisk -vdisk Vol5 -mdiskgrp V5000_Pool_1 -copy 0	
5/28/13 4:43:14 PM	superuser	svctask migratevdisk -vdisk Vol4 -mdiskgrp V5000_Pool_1 -copy 0	

Figure 3-91 Audit Log panel

3.4.8 Settings menu

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

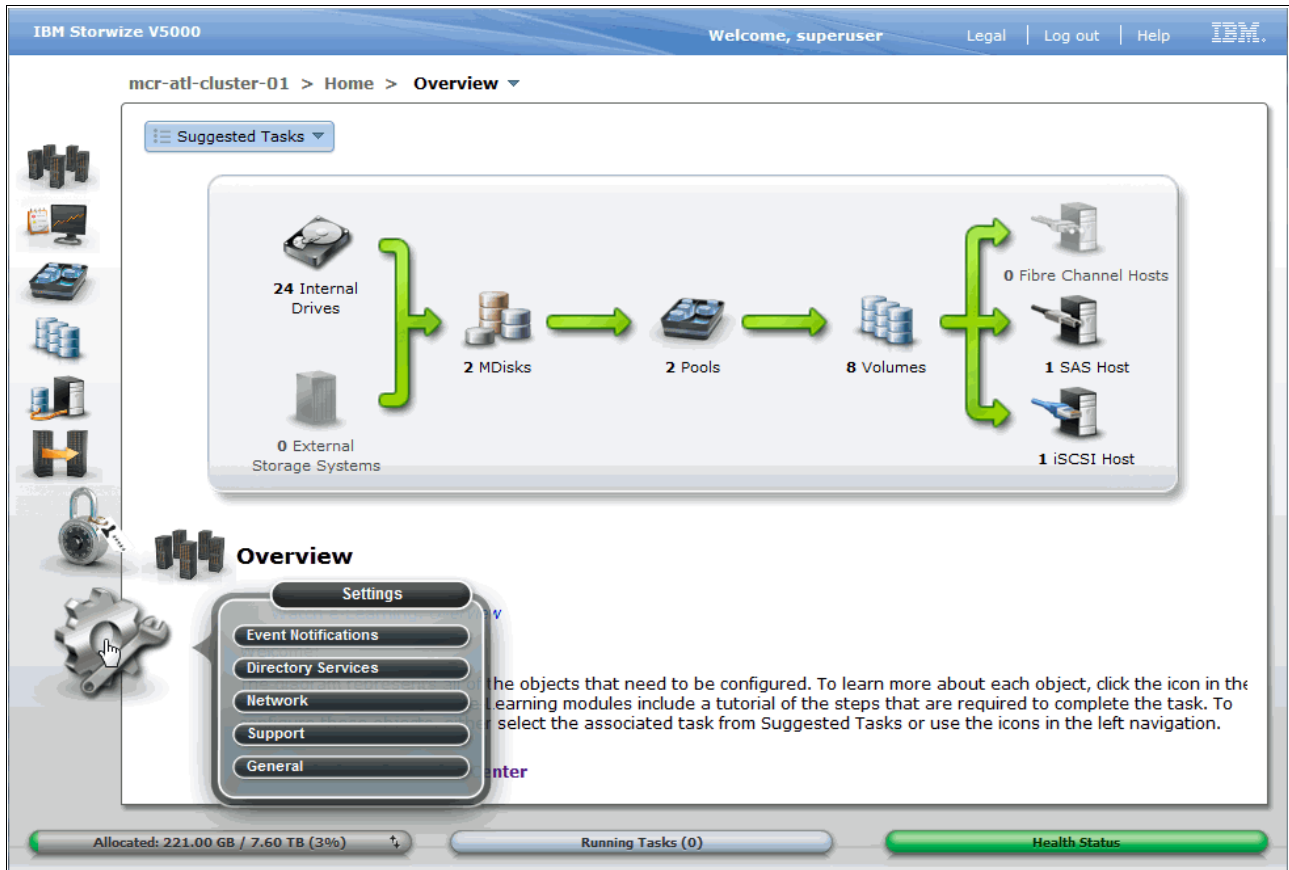


Figure 3-92 Settings menu

Event Notifications panel

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

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

Enabling Email Event Notification option

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

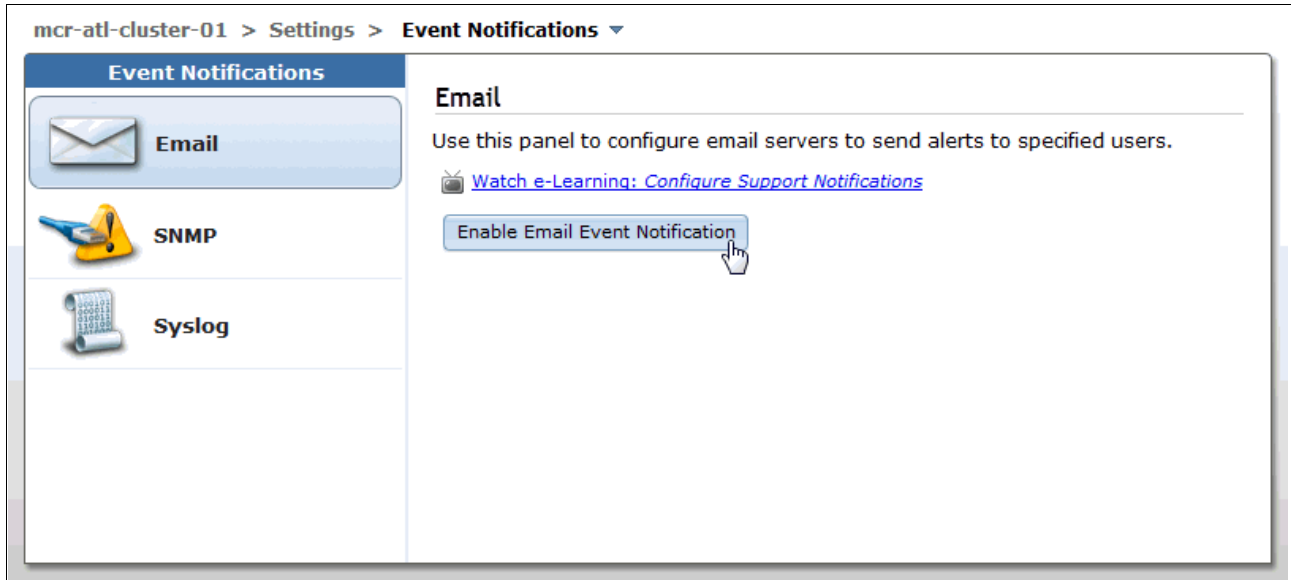


Figure 3-93 Event Notifications panel: Email

Call Home wizard

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

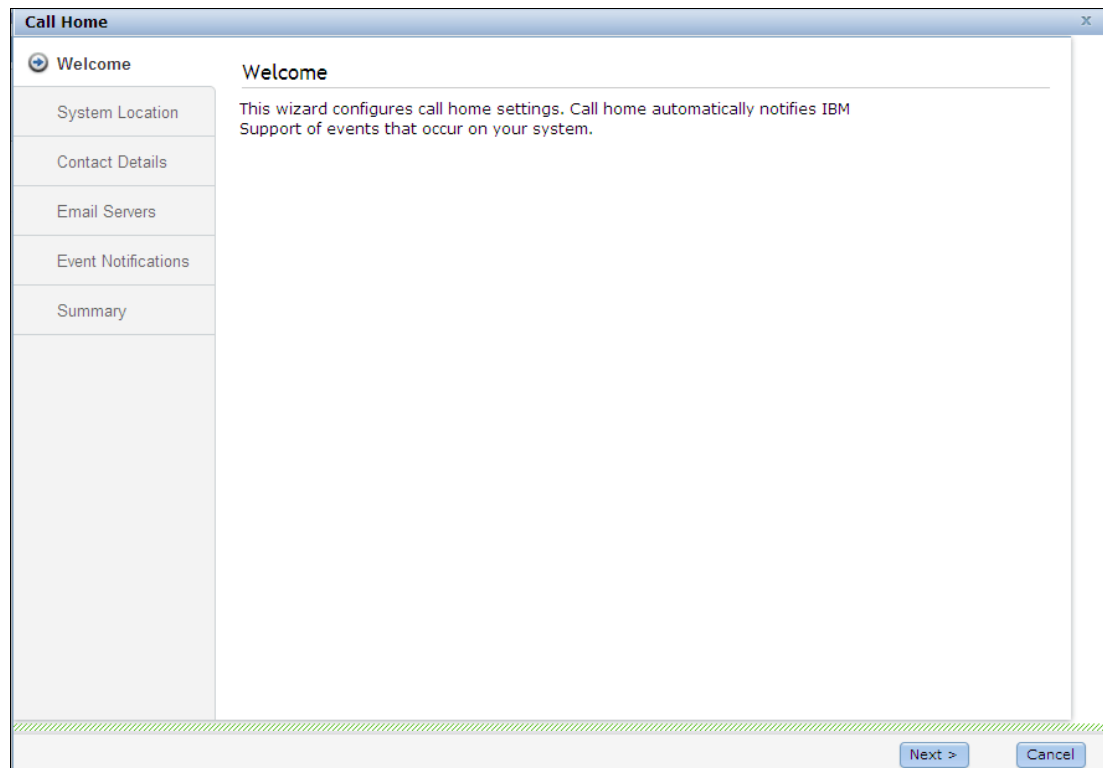


Figure 3-94 Call home wizard

SNMP event notification

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

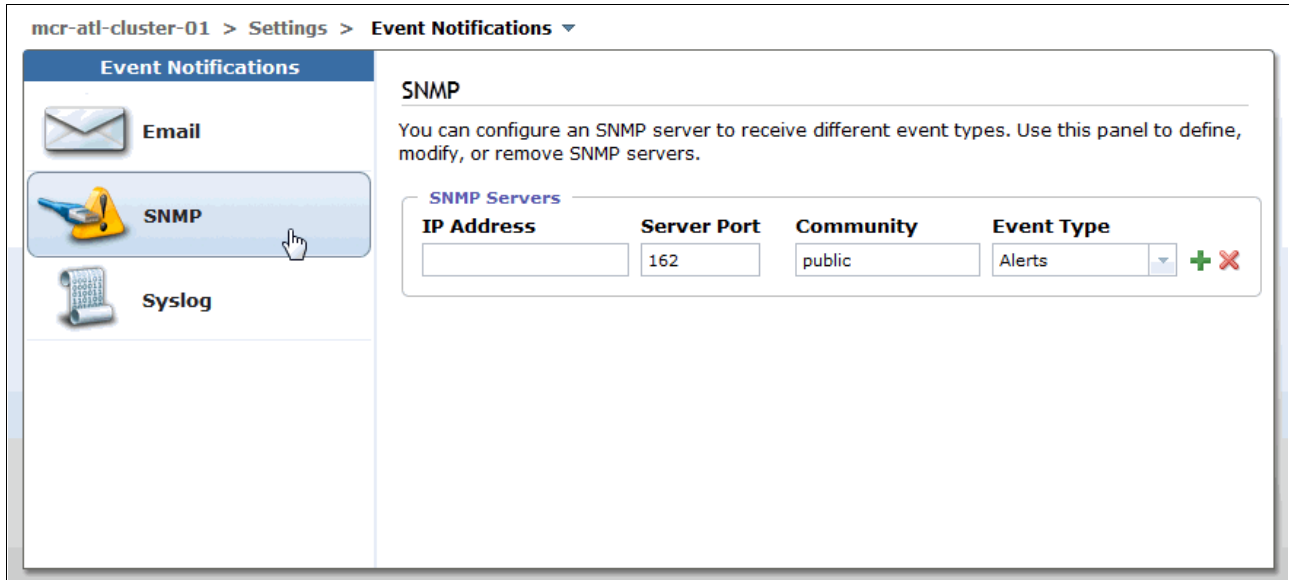


Figure 3-95 SNMP configuration panel

Syslog event notification

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

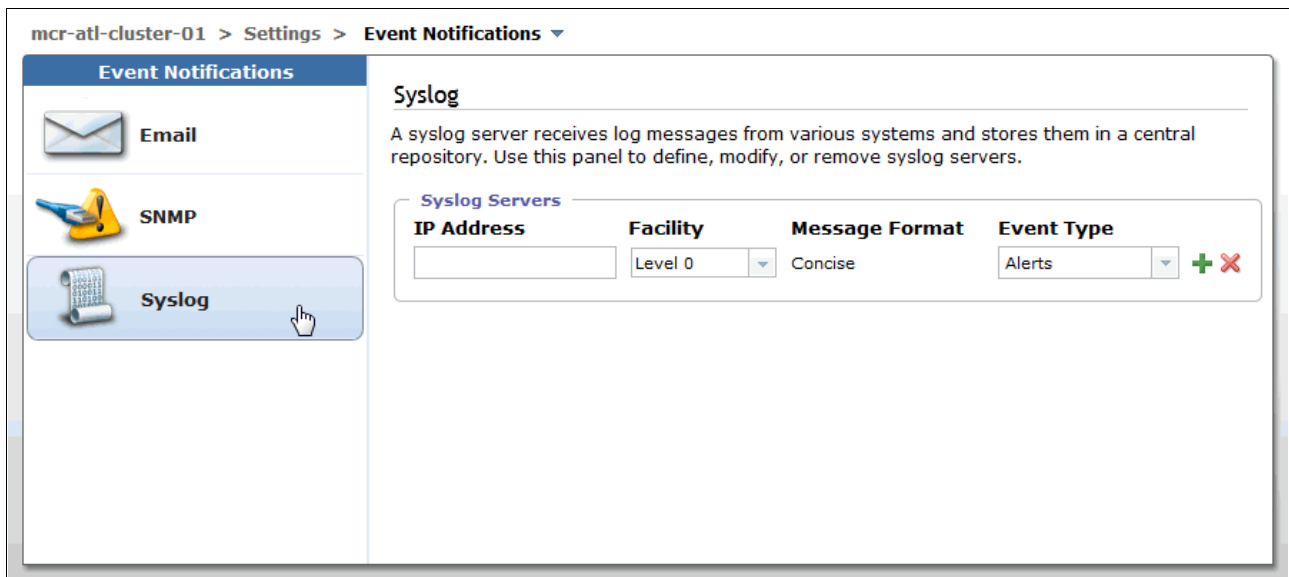


Figure 3-96 Syslog configuration panel

Directory Services panel

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

Enabling Remote Authentication

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

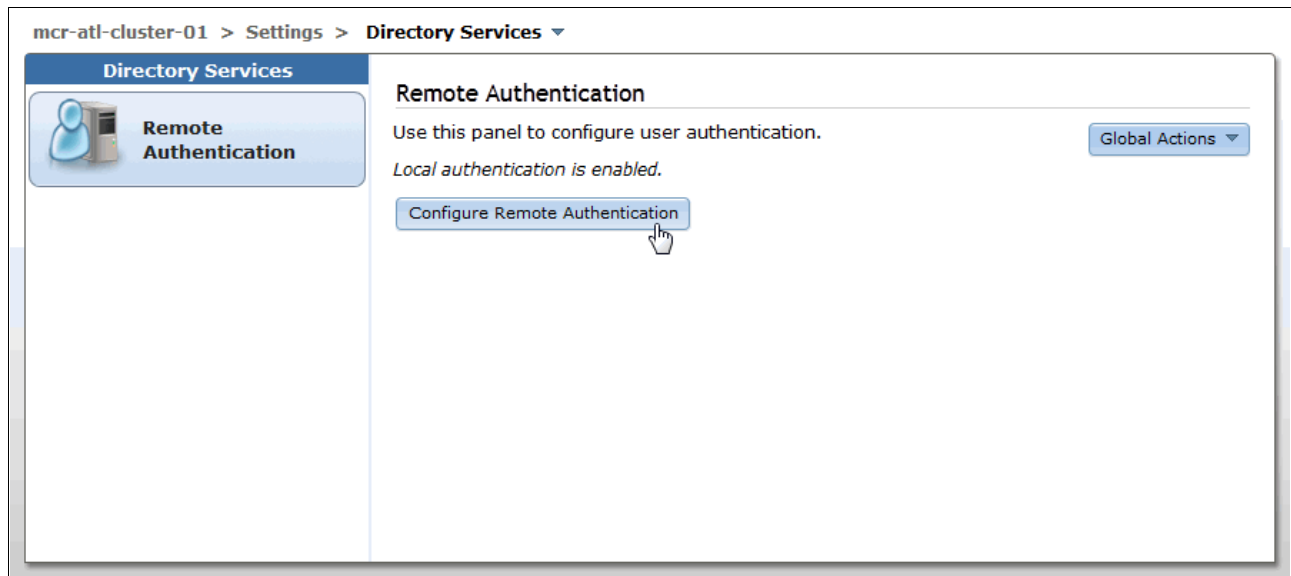


Figure 3-97 Directory Services panel

Network panel

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

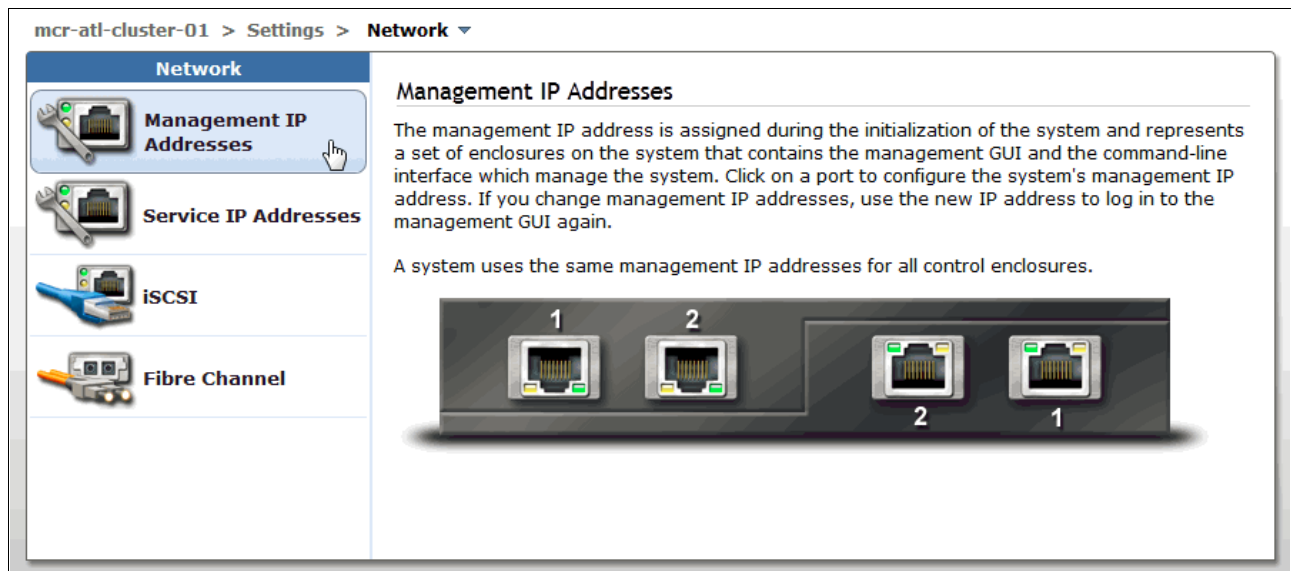


Figure 3-98 Network panel

Management IP addresses

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

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

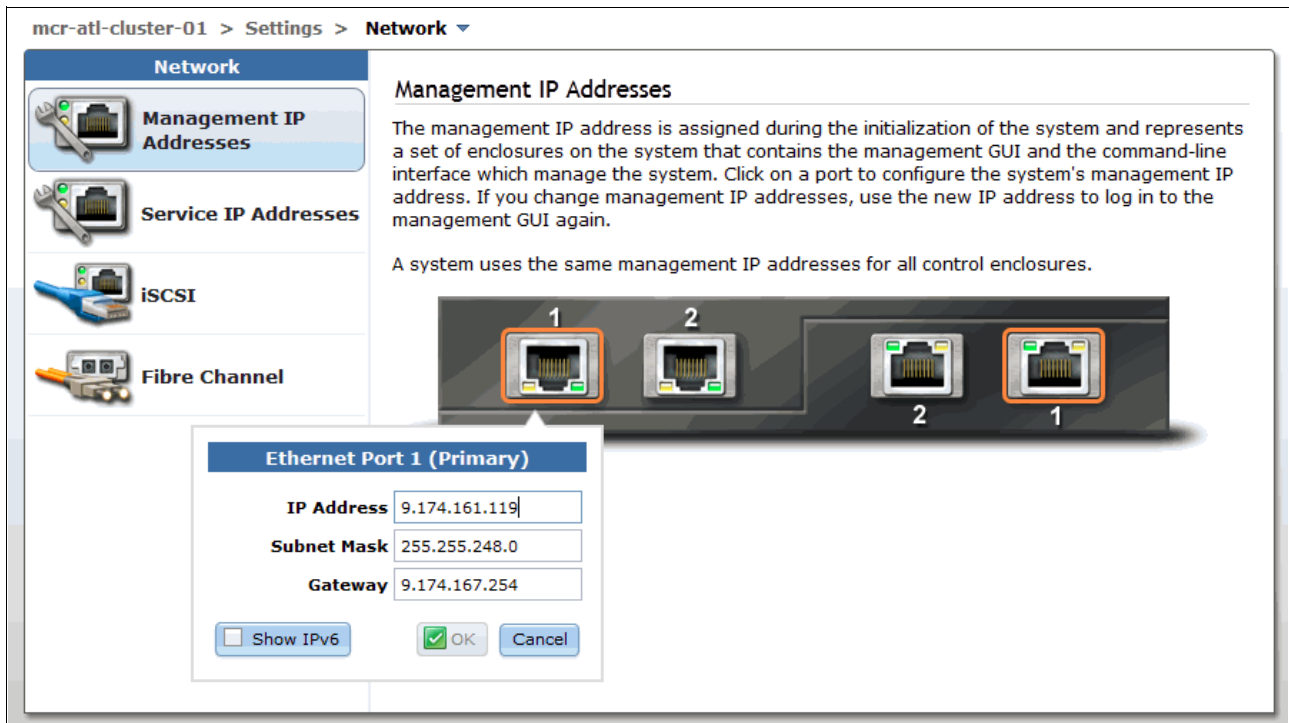


Figure 3-99 Management IP Addresses configuration panel

Service IP Addresses

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

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



Figure 3-100 Service IP Addresses configuration panel

iSCSI connectivity

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

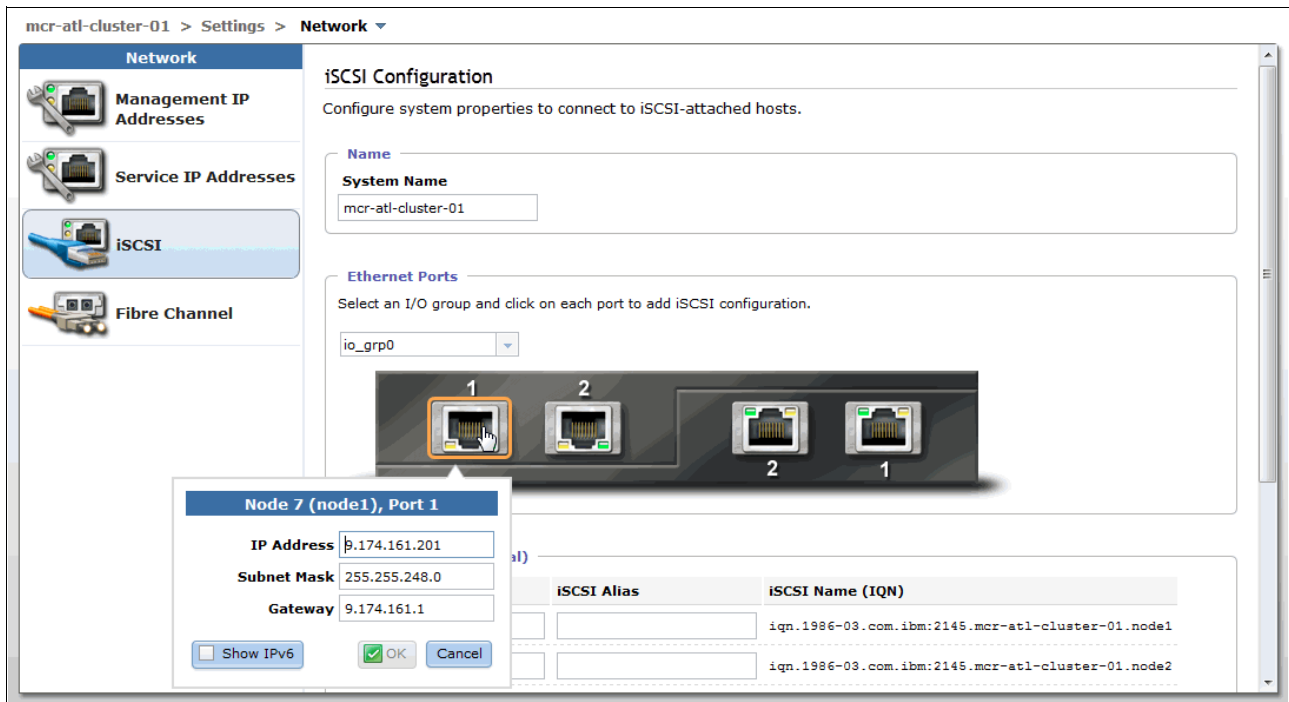


Figure 3-101 iSCSI Configuration panel

Fibre Channel connectivity

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

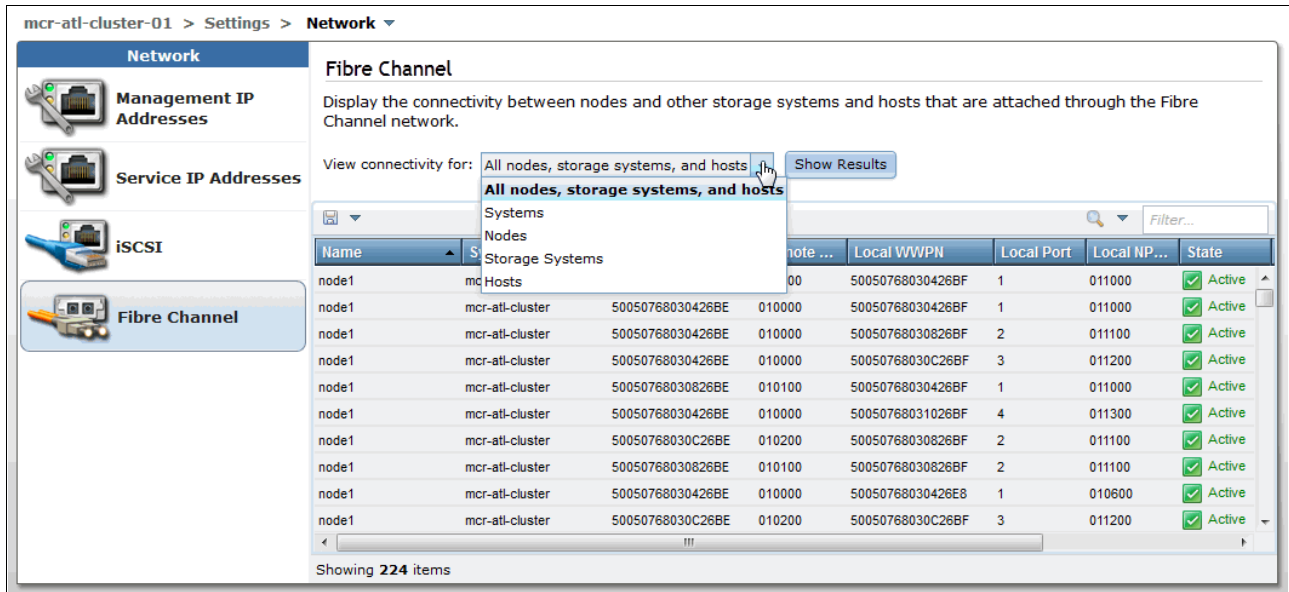


Figure 3-102 Fibre Channel panel

Support panel

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

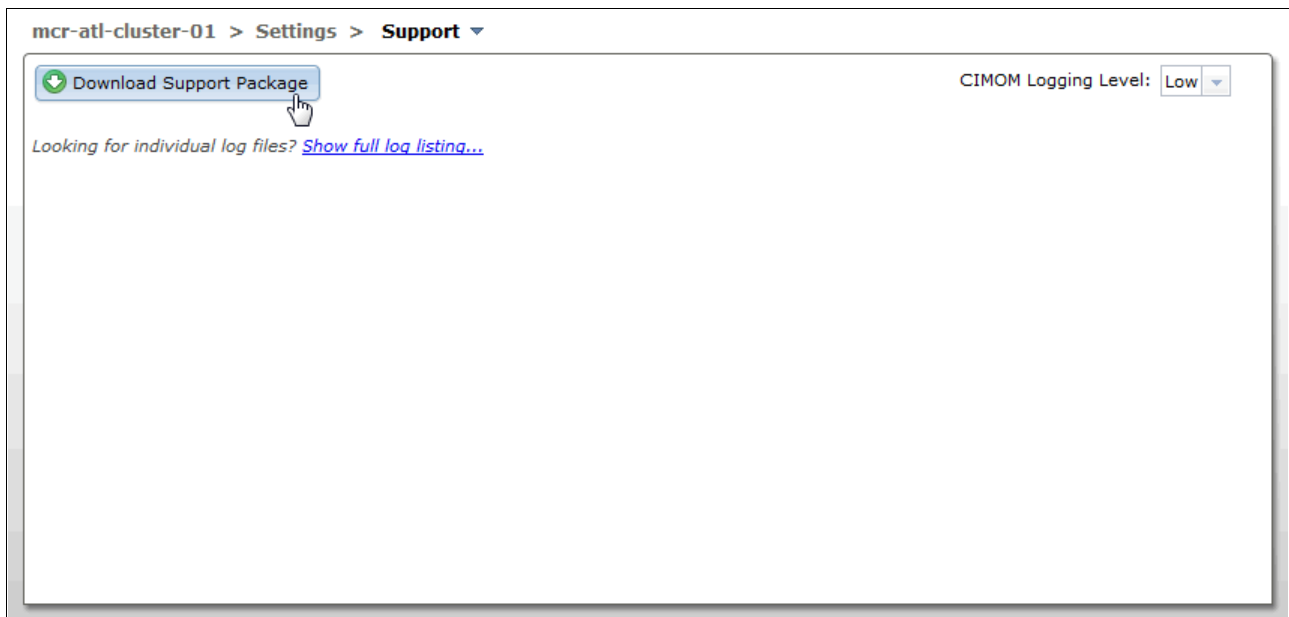


Figure 3-103 Support panel

Download Support Package wizard

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



Figure 3-104 Download Support Package wizard

Show full log listing

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

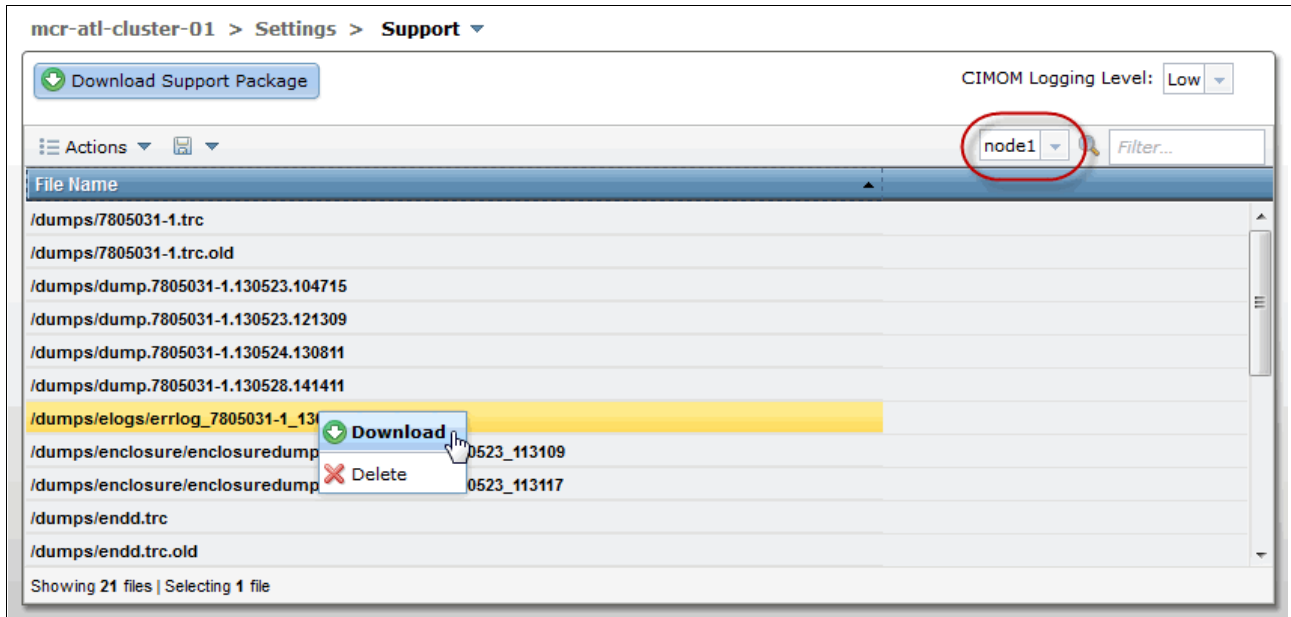


Figure 3-105 Full log listing

General panel

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

Date and Time

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

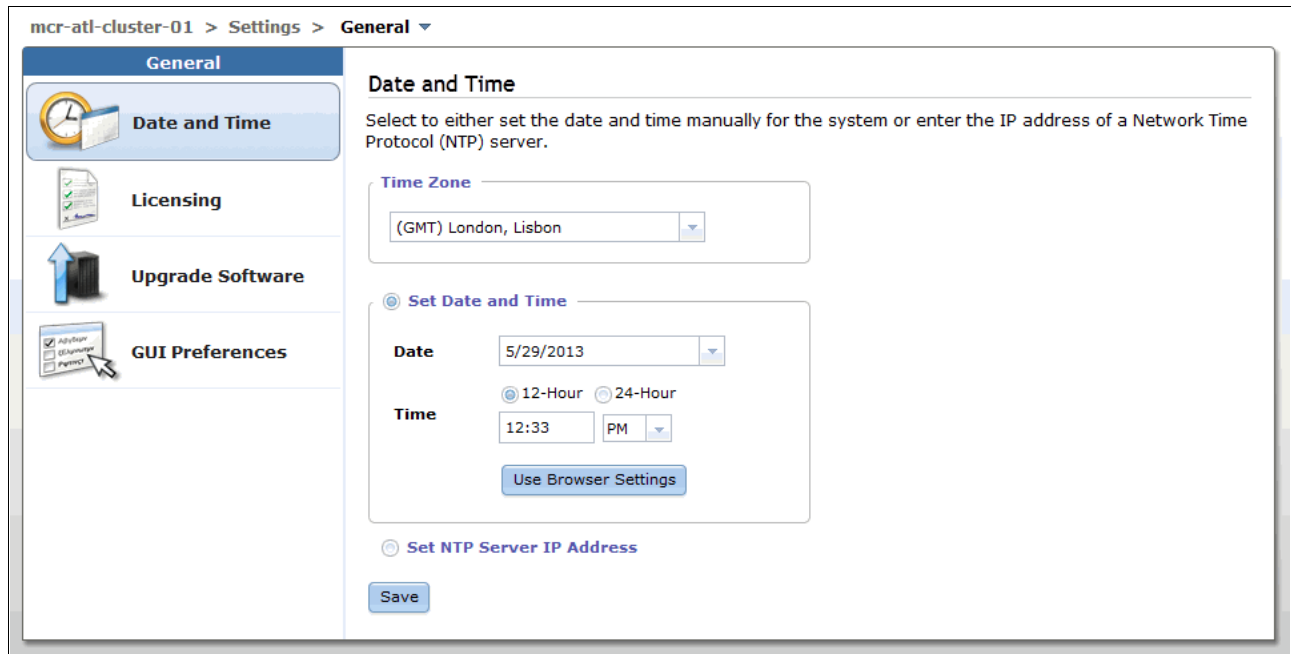


Figure 3-106 General panel

Licensing

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

The following optional licenses are available:

- ▶ FlashCopy
- ▶ Remote Copy
- ▶ Easy Tier
- ▶ External Virtualization

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

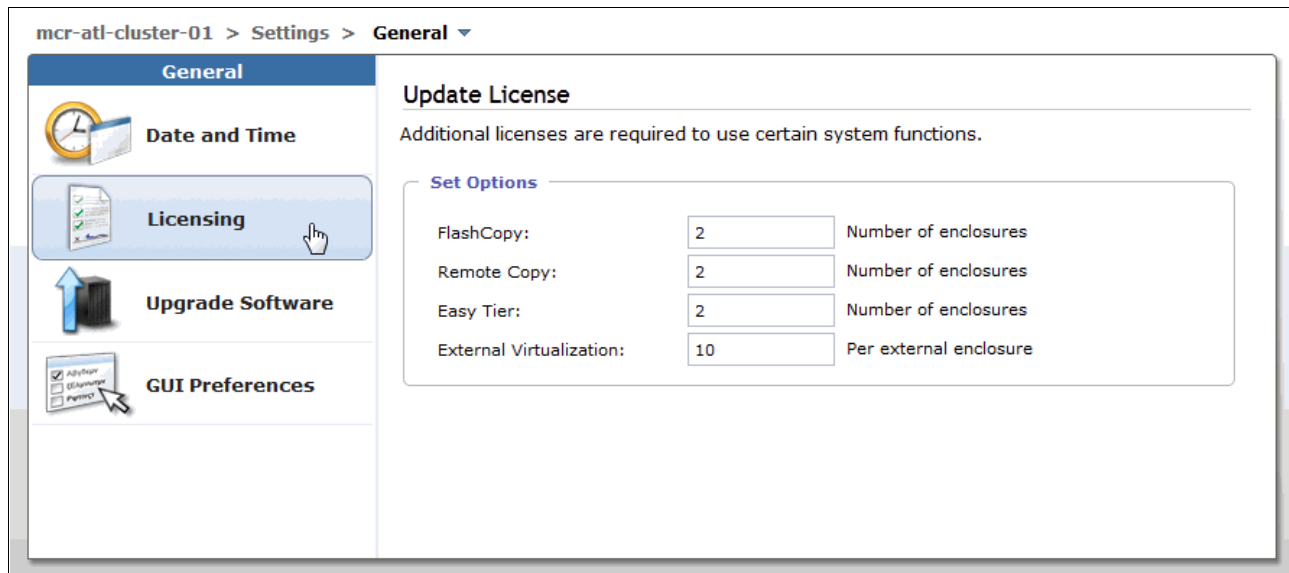


Figure 3-107 Update License panel

Upgrade Software panel

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

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

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

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

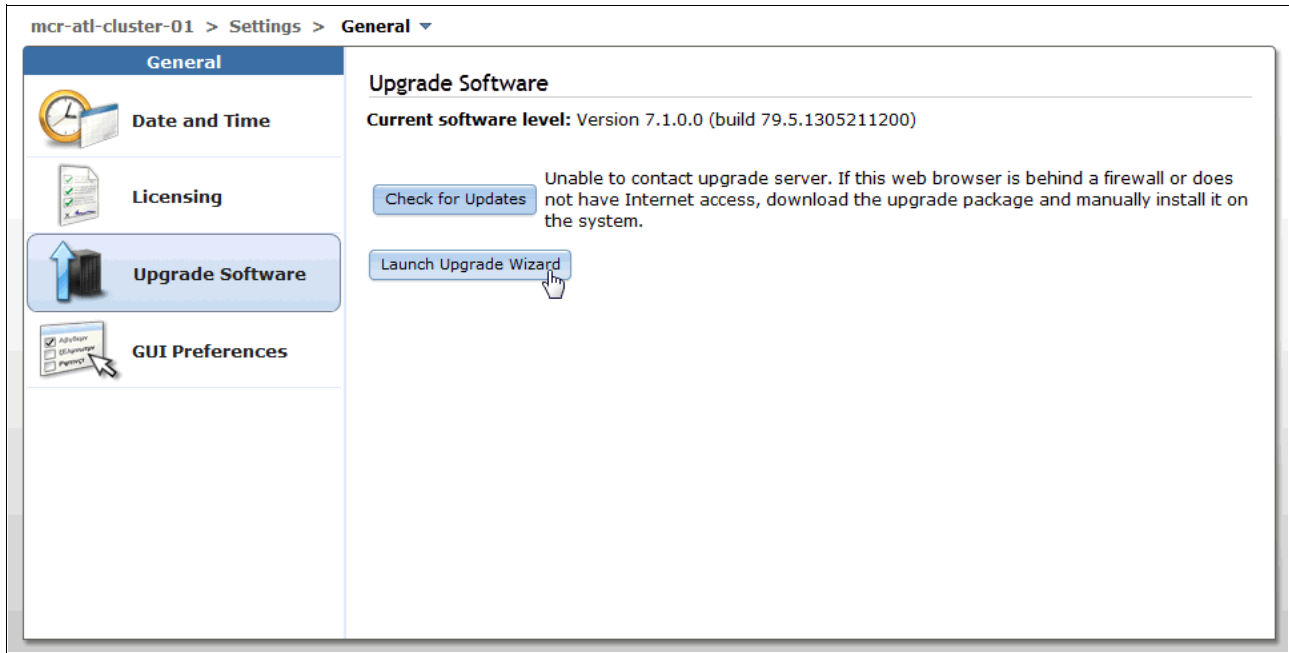


Figure 3-108 Upgrade machine code panel

GUI Preferences panel

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

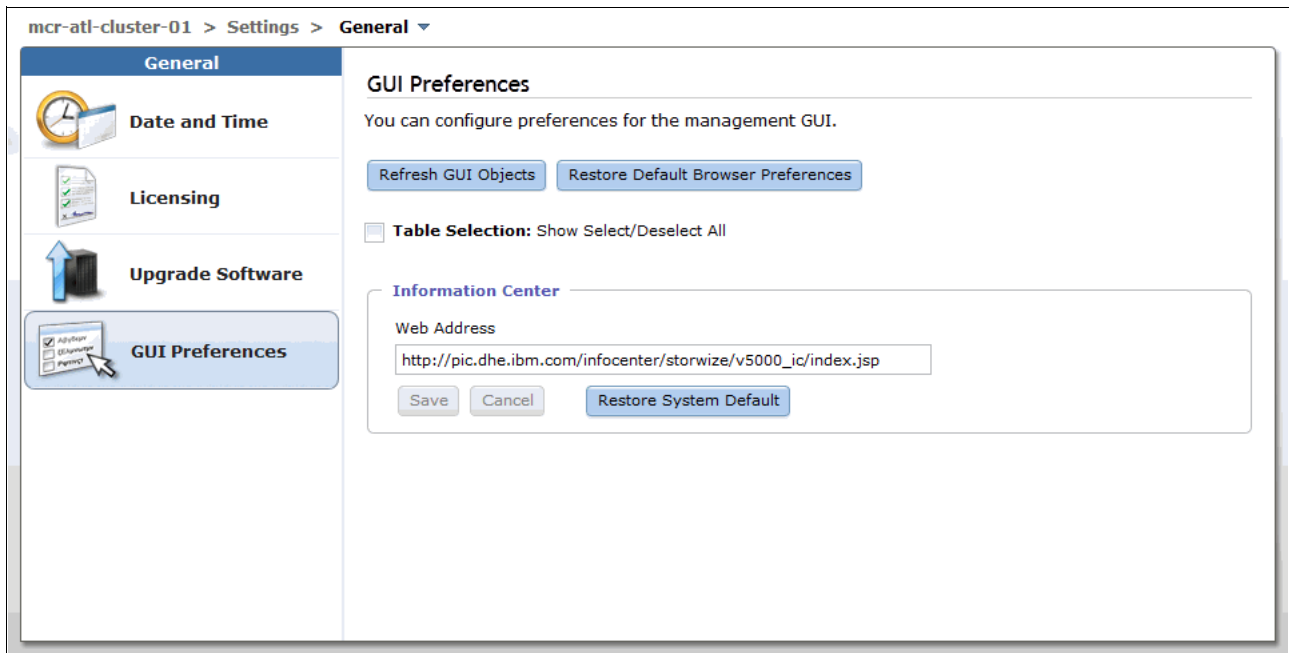


Figure 3-109 GUI Preferences panel

3.5 Management GUI help

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

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

3.5.1 IBM Storwize V5000 Information Center

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

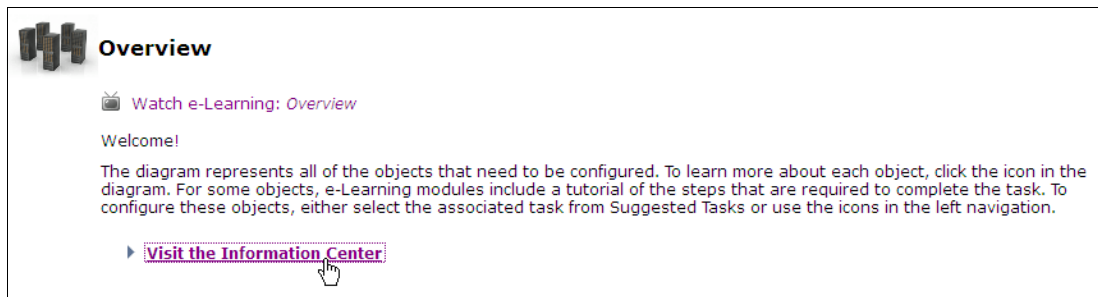


Figure 3-110 Overview panel

3.5.2 Watching an e-Learning video

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

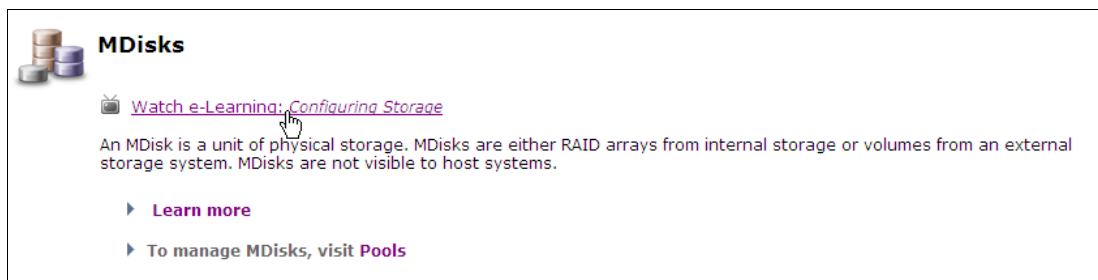


Figure 3-111 Watch e-Learning module

3.5.3 Learning more

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

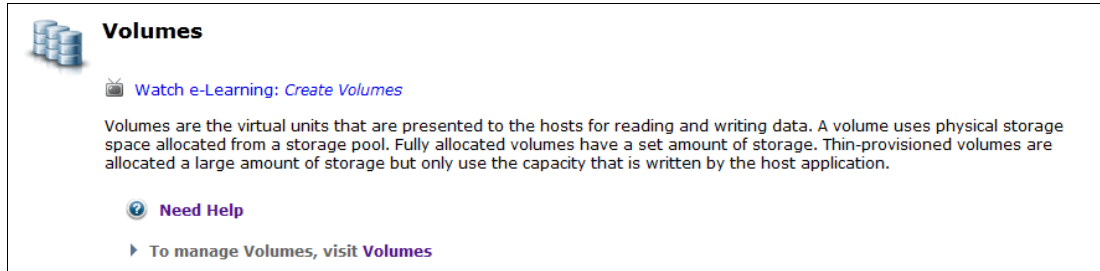


Figure 3-112 Learn more link

Figure 3-113 shows the information panel.

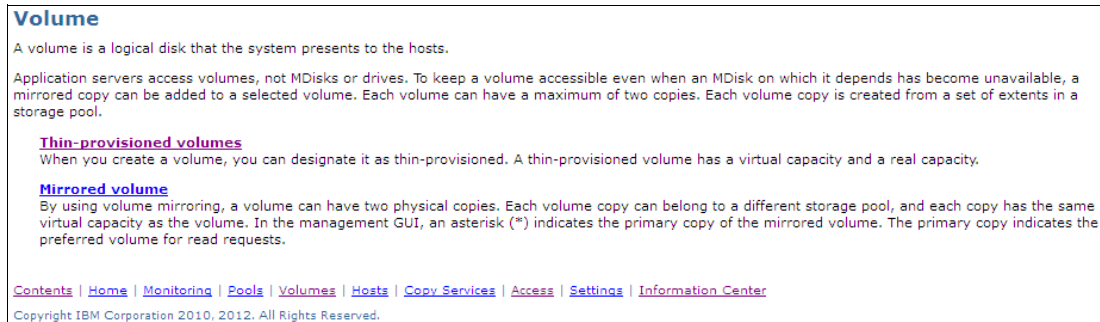


Figure 3-113 Information panel

3.5.4 Embedded panel help

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




Figure 3-114 Embedded panel help

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

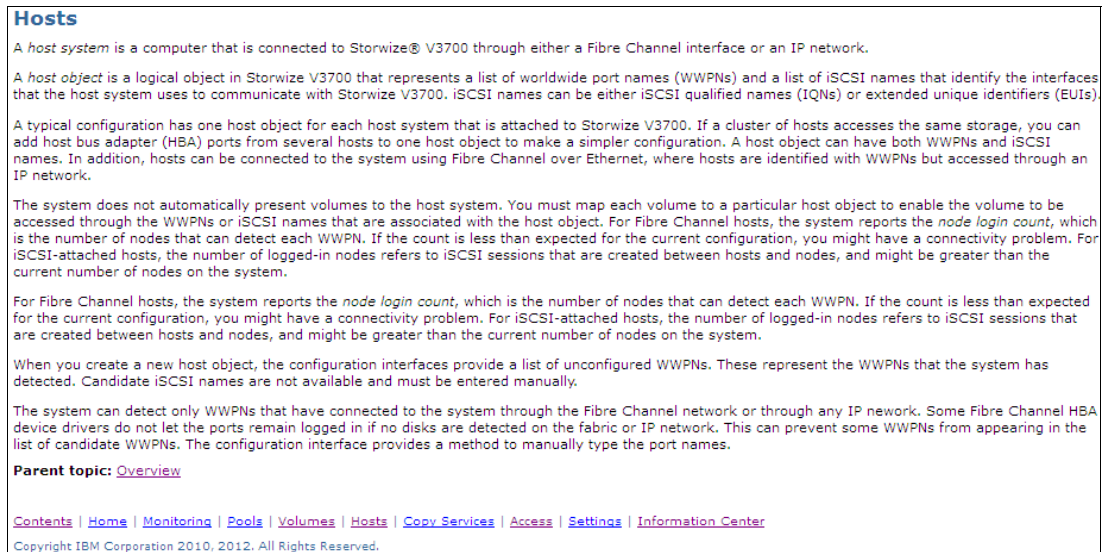


Figure 3-115 Information panel

3.5.5 Hidden question mark help

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




Figure 3-116 Hidden question mark help

3.5.6 Hover help

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



Figure 3-117 Hover help

3.5.7 IBM endorsed YouTube videos

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

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



Host configuration

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

This chapter includes the following topics:

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

4.1 Host attachment overview

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

IBM Storwize V5000 supports the following host attachment protocols:

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

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

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

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

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

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

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

4.2 Preparing the host operating system

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

4.2.1 Windows 2008 R2: Preparing for FC attachment

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

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

Downloading and installing the supported drivers and firmware

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

Configuring Brocade HBAs for Windows

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

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

Configuring QLogic HBAs for Windows

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

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

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

Configuring Emulex HBAs for Windows

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

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

Setting the Windows timeout value

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

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




Figure 4-1 Windows timeout value

Installing Microsoft MPIO multipathing software

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

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

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

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

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



Figure 4-2 iSCSI IP Configuration

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

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

4.2.2 Creating SAS hosts

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

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

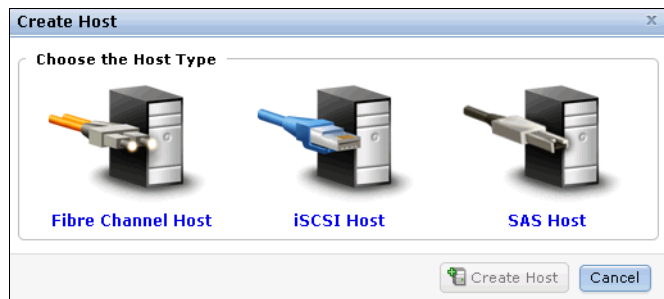


Figure 4-3 Create SAS host

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

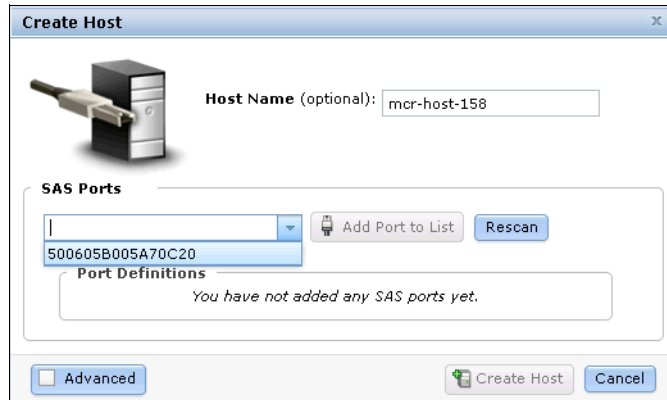


Figure 4-4 Available SAS WWPN or WWPNs

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

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

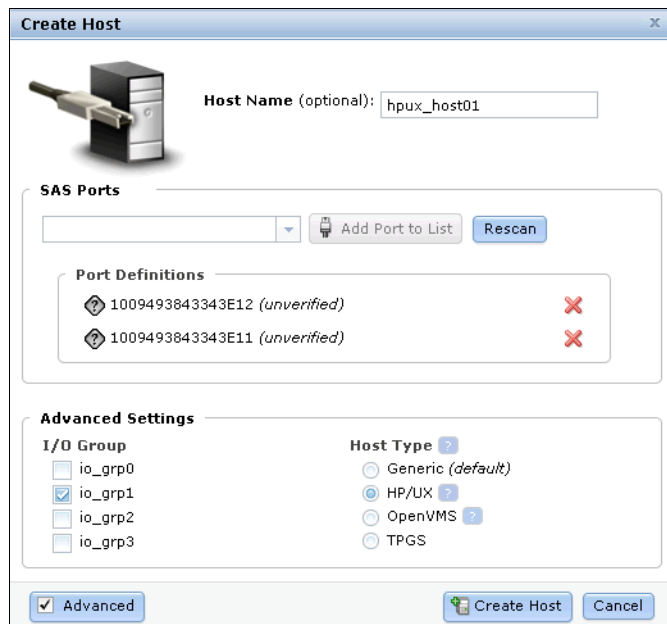


Figure 4-5 Creating HP/UX SAS host

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

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

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

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



I/O Group basic volume configuration

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

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

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

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

This chapter includes the following topics:

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