


Implementing the IBM Storwize V3700

Easily managed and deployed system with embedded GUI

Experience rapid and flexible provisioning

Protect data with remote mirroring



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

Redbooks



International Technical Support Organization

Storwize V3700 Implementation

October 2012

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

First Edition (October 2012)

This edition applies to Version 6 Release 4 of IBM Storwize V3700 machine code.

This document was created or updated on February 27, 2013.

© Copyright International Business Machines Corporation 2012. 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	xi
Trademarks	xii
Preface	xiii
The team who wrote this book	xiii
Now you can become a published author, too!	xv
Comments welcome	xv
Stay connected to IBM Redbooks	xv
Chapter 1. Overview of the Storwize V3700 system	1
1.1 Storwize V3700 overview	2
1.2 Storwize V3700 terminology	5
1.2.1 Storwize V3700 models	6
1.2.2 Storwize V3700 features	8
1.3 Storwize V3700 hardware	9
1.3.1 Control enclosure	11
1.3.2 Expansion enclosure	13
1.3.3 Disk drive types	14
1.4 Storwize V3700 terms	14
1.4.1 Hosts	14
1.4.2 Node canister	14
1.4.3 I/O groups	14
1.4.4 Clustered system	15
1.4.5 RAID	15
1.4.6 Managed disks	16
1.4.7 Quorum disks	17
1.4.8 Storage pools	17
1.4.9 Volumes	18
1.4.10 Thin-provisioned volumes	20
1.4.11 Mirrored volumes	21
1.4.12 iSCSI	21
1.5 FlashCopy	23
1.6 Problem management and support	24
1.6.1 IBM Assist On-site and remote service	24
1.6.2 Event notifications	24
1.6.3 SNMP traps	24
1.6.4 Syslog messages	25
1.6.5 Call Home email	25
1.7 Useful Storwize V3700 websites	25
1.8 IBM Storwize V3700 learning videos on YouTube	25
Chapter 2. Initial configuration	27
2.1 Hardware installation planning	28
2.2 SAN configuration planning	30
2.3 Direct attach planning	31
2.4 LAN configuration planning	32
2.4.1 Management IP address considerations	32
2.4.2 Service IP address considerations	33
2.5 Host configuration planning	34

2.6	Miscellaneous configuration planning	34
2.7	System management	35
2.7.1	Graphical user interface (GUI)	36
2.7.2	Command-line interface	36
2.8	First-time setup	37
2.9	Initial configuration	42
2.9.1	Adding Enclosures	46
2.9.2	Configure Call Home, email alert, and inventory.	51
2.9.3	Configure Event Notifications	53
2.9.4	Service Assistant Tool	56
	Chapter 3. Graphical user interface overview	59
3.1	Getting started.	60
3.1.1	Supported browsers	60
3.1.2	Access the management GUI	60
3.1.3	Overview panel layout.	62
3.2	Navigation	63
3.2.1	Function icons navigation	63
3.2.2	Extended help navigation	64
3.2.3	Breadcrumb navigation aid	66
3.2.4	Suggested tasks aid	67
3.2.5	Presets	67
3.2.6	Access actions	68
3.2.7	Task progress	69
3.2.8	Navigating panels with tables	70
3.3	Status Indicators menus	75
3.3.1	Horizontal bars	75
3.3.2	Allocated status bar menu	75
3.3.3	Running tasks bar menu.	76
3.3.4	Health status bar menu.	77
3.4	Function icon menus	80
3.4.1	Home menu	80
3.4.2	Monitoring menu	81
3.4.3	Pools menu	91
3.4.4	Volumes menu	99
3.4.5	Hosts menu	104
3.4.6	Copy Services menu.	110
3.4.7	Access menu	115
3.4.8	Settings menu	119
3.5	Management GUI help	130
3.5.1	Storwize V3700 Information Center	130
3.5.2	Watch e-Learning	130
3.5.3	Learn more	131
3.5.4	Embedded panel help.	131
3.5.5	Hidden question mark help	132
3.5.6	Hover help.	132
3.5.7	IBM endorsed Youtube videos	133
	Chapter 4. Host configuration	135
4.1	Host attachment overview.	136
4.2	Preparing the host operating system	137
4.2.1	Windows 2008 (R2): Preparing for Fibre Channel attachment	137
4.2.2	Windows 2008 R2: Preparing for iSCSI attachment	140

4.2.3	VMware ESX: Preparing for Fibre Channel attachment	142
4.2.4	VMware ESX: Preparing for iSCSI attachment	145
4.3	Creating hosts using the GUI	150
4.3.1	Creating Fibre Channel hosts	151
4.3.2	Creating iSCSI hosts	155
Chapter 5. Basic volume configuration		163
5.1	Provisioning storage from IBM Storwize V3700 and making it available to the host.	164
5.1.1	Creating a generic volume	166
5.1.2	Creating a thin-provisioned volume.	168
5.1.3	Creating a mirrored volume.	171
5.1.4	Creating a thin-mirror volume	174
5.2	Mapping a volume to the host.	176
5.2.1	Mapping newly created volumes to the host using the wizard	176
5.2.2	Manually mapping a volume to the host	178
5.3	Discovering the volumes from the host and specifying multipath settings	180
5.3.1	Windows 2008 Fibre Channel volume attachment	182
5.3.2	Windows 2008 iSCSI volume attachment.	186
5.3.3	VMware ESX Fibre Channel attachment	195
5.3.4	VMware ESX iSCSI attachment	199
Chapter 6. Storage migration wizard		209
6.1	Interoperability and compatibility.	210
6.2	.The storage migration wizard.	210
6.2.1	External virtualization capability	210
6.2.2	Overview of the storage migration wizard.	210
6.2.3	Storage migration wizard tasks.	211
6.3	Storage migration wizard example scenario	223
6.3.1	Storage migration wizard example scenario description.	223
6.3.2	Using the storage migration wizard for example scenario	226
Chapter 7. Storage pools		267
7.1	Working with internal drives	268
7.1.1	Internal storage window	269
7.1.2	Actions on internal drives	271
7.2	Configuring internal storage	278
7.2.1	RAID configuration presets	279
7.2.2	Customise initial storage configuration	279
7.2.3	Create new MDisk and Pool	280
7.2.4	Using the recommended configuration	281
7.2.5	Selecting a different configuration.	283
7.3	Working with MDisks.	288
7.3.1	MDisk by Pools panel	290
7.3.2	RAID action for MDisks.	291
7.3.3	Additional actions on MDisks	295
7.4	Working with storage pools.	297
7.4.1	Create new Pool	298
7.4.2	Actions on Storage Pools	299
Chapter 8. Advanced host and volume administration		303
8.1	Advanced host administration.	304
8.1.1	Modifying Mappings menu	306
8.1.2	Unmapping all volumes from a host	312
8.1.3	Renaming	315

8.1.4	Deleting a host	317
8.1.5	Host properties	319
8.2	Adding and deleting host ports	323
8.2.1	Adding a host port	325
8.2.2	Adding a Fibre Channel port	325
8.2.3	Adding an iSCSI host port	330
8.2.4	Deleting a host port	332
8.3	Host mappings overview	334
8.3.1	Unmapping a volume	335
8.3.2	Properties (Host)	335
8.3.3	Properties (Volume)	336
8.4	Advanced volume administration	336
8.4.1	Advanced volume functions	336
8.4.2	Mapping a volume to a host	339
8.4.3	Unmapping volumes from all hosts	341
8.4.4	Viewing a host mapped to a volume	342
8.4.5	Renaming a volume	343
8.4.6	Shrinking a volume	344
8.4.7	Expanding a volume	345
8.4.8	Migrating a volume to another storage pool	346
8.4.9	Deleting a volume	350
8.5	Volume properties	353
8.5.1	Overview tab	353
8.5.2	Host Maps tab	357
8.5.3	Member MDisk tab	359
8.5.4	Adding a mirrored volume copy	360
8.5.5	Editing thin-provisioned volume properties	363
8.6	Advanced volume copy functions	367
8.6.1	Thin-provisioned	369
8.6.2	Splitting into a new volume	370
8.6.3	Validate Volume Copies option	371
8.6.4	Delete Copy option	374
8.6.5	Migrating volumes using the volume copy features	374
8.7	Volumes by storage pool	375
8.7.1	Changing the Storage Pool function icon	378
8.7.2	Changing the storage pool name	379
8.8	Volumes by host	382
8.8.1	Renaming a host	384
8.9	FlashCopy	386
8.9.1	Business requirements for FlashCopy	386
8.9.2	FlashCopy functional overview	387
8.9.3	Planning for FlashCopy	395
8.9.4	Managing FlashCopy using the GUI	396
8.9.5	Quick path to create FlashCopy presets	398
8.9.6	Managing FlashCopy mapping	401
Chapter 9. RAS, monitoring, and troubleshooting.		425
9.1	Storwize V3700 components	427
9.1.1	Enclosure midplane assembly	427
9.1.2	Node canisters - Ports and LED	428
9.1.3	Node canister replaceable hardware components	431
9.1.4	Expansion canister - Ports and LED	434
9.1.5	Disk subsystem	436

9.1.6 Power Supply Unit	439
9.2 Configuration backup procedure	441
9.2.1 Generating a configuration backup using the CLI	441
9.2.2 Downloading a configuration backup using the GUI	442
9.3 Software upgrade	444
9.3.1 Upgrading software automatically	445
9.3.2 GUI upgrade process	446
9.3.3 Upgrading software manually	449
9.4 Event log	454
9.4.1 Managing the event log.	455
9.4.2 Alert handling and Recommended Actions.	457
9.5 Collecting support information	464
9.5.1 Support Information via GUI	464
9.5.2 Support information via Service Assistant	465
9.5.3 Support Information onto USB stick	466
9.6 Power on and shut down Storwize V3700	468
9.6.1 Shut down system.	468
9.6.2 Powering on	470
Appendix A. CLI setup and SAN Boot	471
Command-line interface	472
Basic setup	472
Example commands	482
SAN Boot	484
Enabling SAN Boot for Windows.	484
Enabling SAN Boot for VMware	485
Windows SAN Boot migration.	485
Related publications	487
IBM Redbooks	487
IBM Storwize V3700 publications.	487
IBM Storwize V3700 support	487
Help from IBM	487

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®

FlashCopy®

IBM®

Redbooks®

Redbooks (logo) ®

Storwize®

System Storage®

Tivoli®

The following terms are trademarks of other companies:

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

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

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 and more difficult. IT organizations have limited resources but need to stay responsive to dynamic environments and act quickly to consolidate, simplify and optimize their IT infrastructures. The IBM Storwize V3700 system provides a smarter solution that is affordable, easy to use and self-optimizing, enabling organizations to overcome these storage challenges.

Storwize V3700, the most recent addition to the IBM Storwize family of disk systems, delivers efficient, entry-level configurations 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 V3700 offers advanced software capabilities usually found in more expensive systems.

Built upon innovative IBM technology, Storwize V3700 addresses the block storage requirements of small and midsize organizations. Providing up to 120 drives of capacity packaged in a compact 2U, Storwize V3700 is designed to accommodate the most common storage network technologies to enable easy implementation and management.

The team who wrote this book

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



Jon Tate is a Project Manager for IBM System Storage SAN Solutions at the International Technical Support Organization, San Jose Center. Before joining the ITSO in 1999, he worked in the IBM Technical Support Center, providing Level 2 support for IBM storage products. Jon has 26 years of experience in storage software and management, services, and support, and is both an IBM Certified IT Specialist and an IBM SAN Certified Specialist. He is also the UK Chairman of the Storage Networking Industry Association.



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



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




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



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



Nancy Kinney worked as an IBM Remote Technical Support Engineer for IBM Global Technology Services in the Austin AIX/IBM System Storage center where she acquired end-to-end experience in troubleshooting FC technologies. She is currently an SME (Subject Matter Expert) in the Columbia Missouri Storage GDC, responsible for supporting SSO Storage clients. She holds an IBM Midrange Specialist Certification as well as NETAPP/NSeries NCD A Certification. She has a wide range of experience working with FC Storage technologies from multiple vendors as well as working with multipathing drivers for OS's attaching to storage.



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

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

Martyn Spink
 Djihed Afifi
 David Turnbull
 IBM Manchester Lab

Paul Marris
Simon Swindells
IBM Hursley

John Bynum
IBM San Jose

Now you can become a published author, too!

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

- ▶ Use the online **Contact us** review Redbooks form found at:

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 Storwize V3700 system

This chapter provides an overview of the Storwize V3700 architecture and includes a brief explanation of storage virtualization.

1.1 Storwize V3700 overview

The Storwize V3700 solution provides a modular storage system that includes the capability to virtualize its own internal storage. The IBM Storwize® V3700 system is a virtualizing RAID entry / midrange storage system.

Storwize V3700 has the following benefits:

- ▶ Brings enterprise technology to entry / midrange storage
- ▶ Specialty 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 Storwize V3700 subsystem consists of a set of drive enclosures. Control enclosures contain disk drives and two nodes which can be attached to SAN fabric or iSCSI infrastructure. Expansion enclosures contain drives and are attached to control enclosures.

The simplest use of the Storwize V3700 is as a traditional RAID subsystem. The internal drives are configured into RAID arrays and virtual disks volumes created from those arrays.

The Storwize V3700 supports regular and solid-state drives

The Storwize V3700 system provides a number of configuration options that are aimed at simplifying the implementation process. It also provides configuration presets and ,automated wizards, called *Directed Maintenance Procedures (DMP)*, to assist in resolving any events that might occur. The Storwize V3700 system is a clustered, and scalable midrange storage system.

Figure 1-1 shows the Storwize V3700 high-level overview attached via SAN.

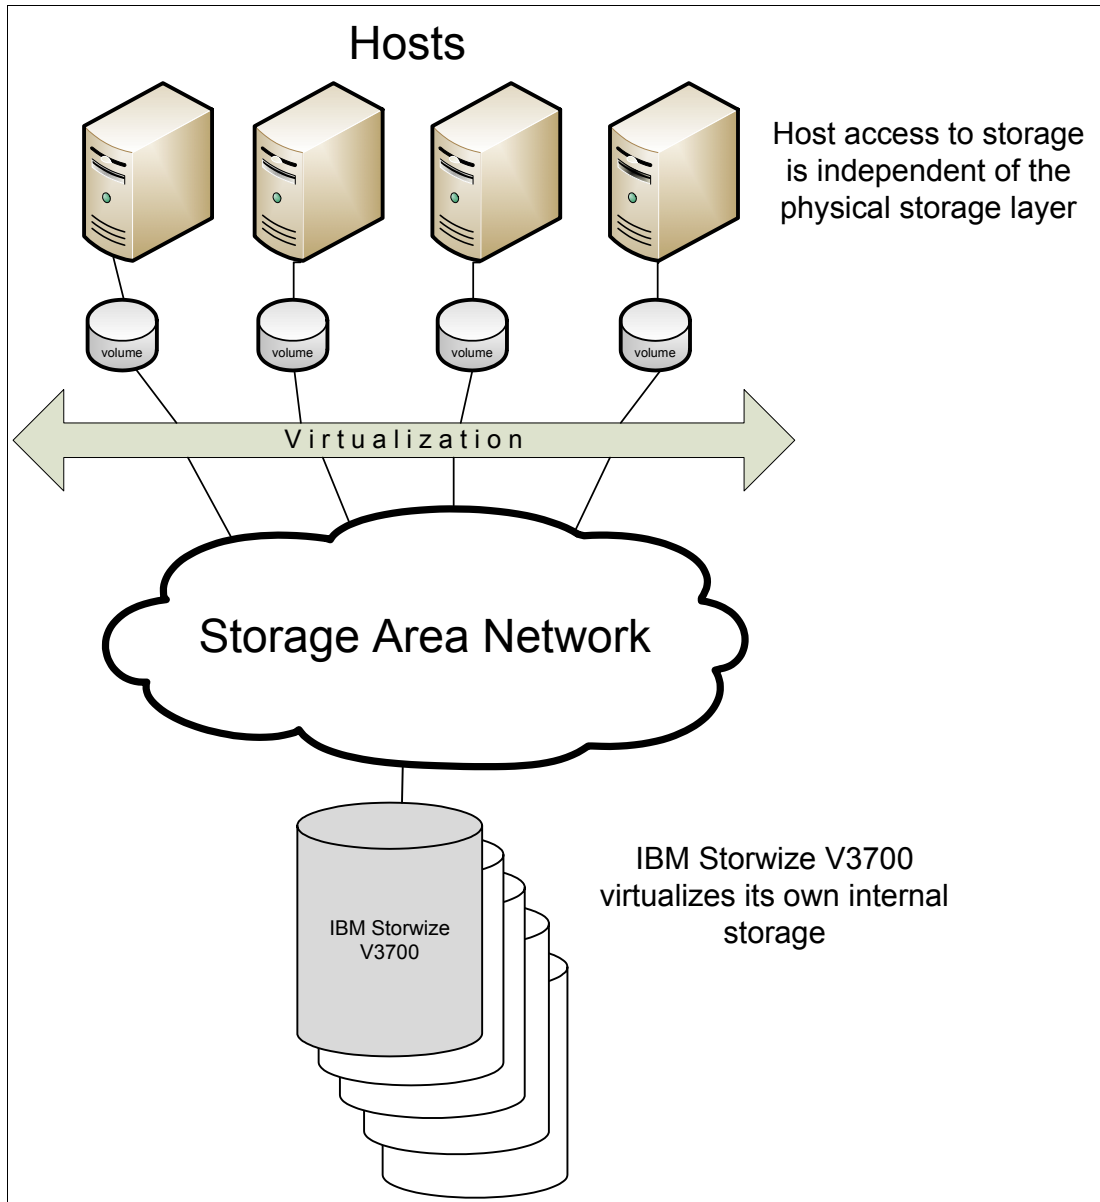


Figure 1-1 Storwize V3700 overview attached with SAN

Note: A big benefit of the Storwize V3700 is that you can direct attach it to a host as shown in Figure 1-2. For any restrictions refer to the IBM System Storage® Interoperation Center (SSIC) <http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss> or to the following restrictions website : <http://www-01.ibm.com/support/docview.wss?uid=ssg1S1004233>

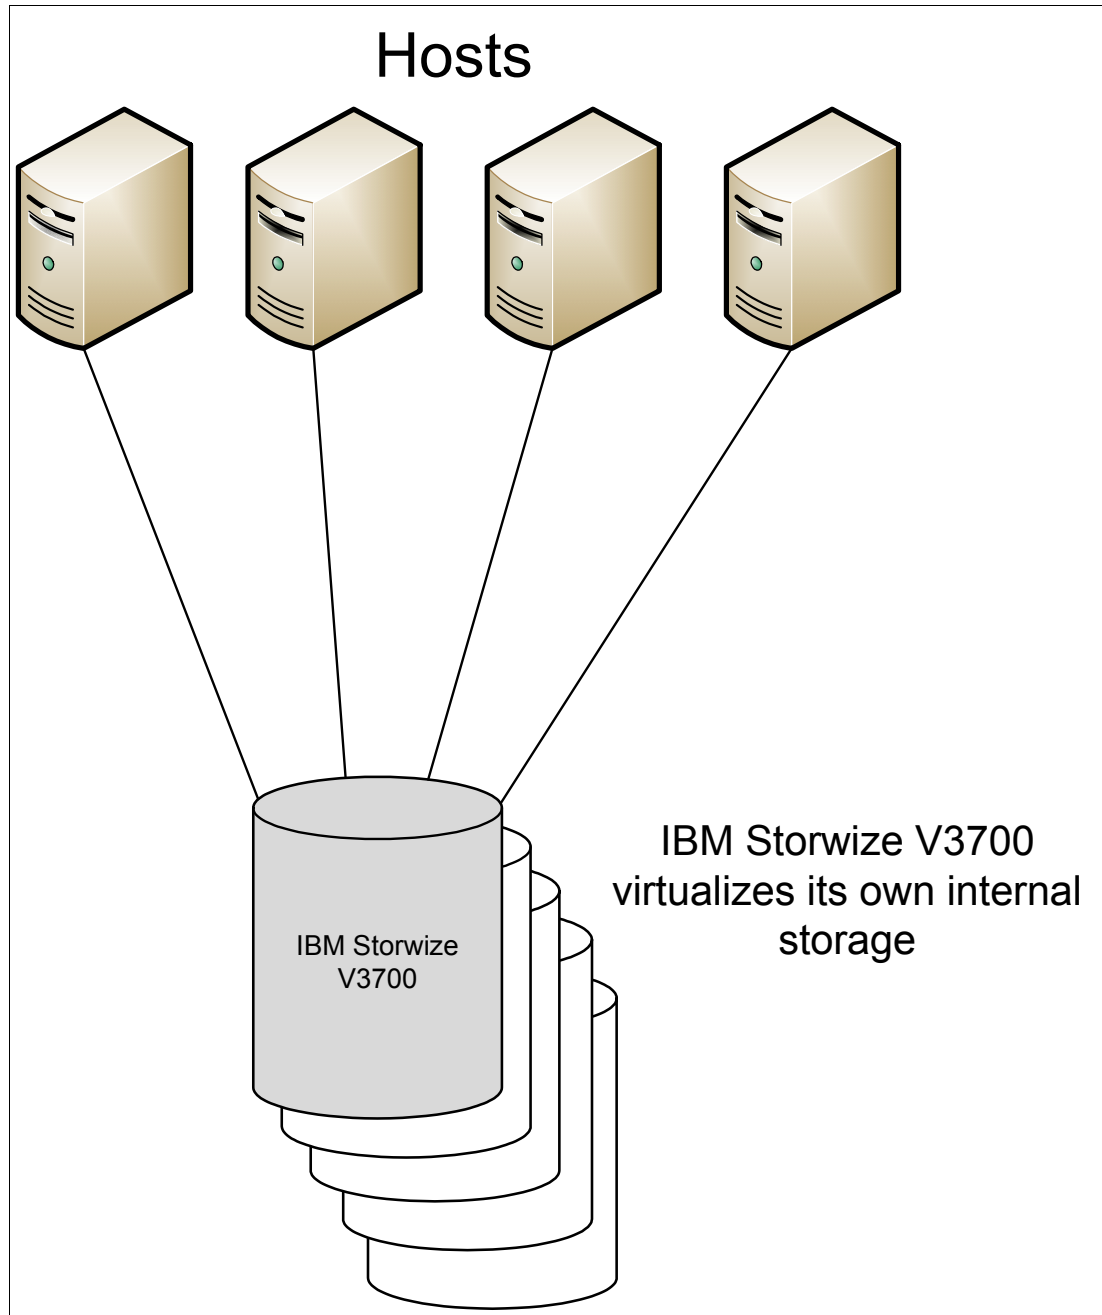


Figure 1-2 Storwize V3700 direct attached

Included with an IBM Storwize V3700 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 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 CLI to setup or control the system.

The IBM Storwize V3700 solution provides a choice of up to 60 x 3.5-inch or 120 x 2.5-inch or a combination of both drive types for the internal storage.

The Storwize V3700 uses SAS cables and connectors to attach to the optional expansion enclosures.

The IBM Storwize V3700 solution consists of one control enclosure and, optionally, up to 4 expansion enclosures (and supports the intermixing of the different expansion enclosures). Within each enclosure are two canisters. Control enclosures contain two node canisters, and expansion enclosures contain two expansion canisters.

There can be up to 180 TB of storage capacity in one complete subsystem.

1.2 Storwize V3700 terminology

The Storwize V3700 system has introduced some terminology, which is consistent with the whole IBM Storwize family and SAN Volume Controller. The terms are defined in Table 1-1.

Table 1-1 Storwize V3700 terminology

Storwize V3700 term	Definition
Battery	Each control enclosure node canister in a Storwize V3700 contains a battery.
Canister	Canisters are hardware units that are subcomponents of a Storwize V3700 enclosures. Each enclosure contains 2 canisters.
Chain	A set of enclosures that are attached to provide redundant access to the drives that are inside the enclosures. Each control enclosure have one chain.
Clone	A copy of a volume on a server at a particular point in time. The contents of the copy can be customized while the contents of the original volume are preserved.
Control enclosure	A hardware unit that includes the chassis, node canisters, drives, and power sources.
Data migration	Storwize V3700 allows you to migrate data from existing external storage to its internal volumes.
Drive	Storwize V3700 supports a range of hard drives (HDDs) and solid-state drives (SSDs)
Enclosure	An enclosure is the basic housing unit for the Storwize V3700. 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 serial-attached SCSI (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.
Fibre Channel port	Fibre Channel ports are connections for the hosts to get access to the Storwize V3700.

Storwize V3700 term	Definition
Host mapping	The process of controlling which hosts have access to specific volumes within a Storwize V3700.
Internal storage	Array managed disks (MDisks) and drives that are held in enclosures and nodes that are part of the Storwize V3700.
Managed disk (MDisk)	A component of a storage pool that is managed by a clustered system. An MDisk is either part of a RAID array of internal storage or a Small Computer System Interface (SCSI) logical unit (LU) for external storage. An MDisk is not visible to a host system on the storage area network.
Node canister	A hardware unit that includes the node hardware, fabric and service interfaces, serial-attached SCSI (SAS), expansion ports, and battery.
PHY	A single SAS lane. There are four PHYs in each SAS cable.
Power Supply Unit	Each enclosure has 2 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..
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 serial-attached SCSI (SAS) connectivity of a set of drives within multiple enclosures. The enclosures can be either control enclosures or expansion enclosures.
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 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 by a worldwide port name (WWPN).

1.2.1 Storwize V3700 models

The Storwize V3700 platform consists of enclosures and drives. Each enclosure contains two canisters that, although they can be replaced independently, are seen as part of the enclosure.

Additional Storwize V3700 information: For the most up-to-date information about the features, benefits, and specifications of Storwize V3700 models, see the following address:

http://www.ibm.com/systems/storage/disk/storwize_v3700/index.html

The information in this book is valid at the time of writing, but as the Storwize V3700 matures, expect to see new features and enhanced specifications.

The Storwize V3700 models are described in Table 1-2.

Table 1-2 Storwize V3700 models

Model	Cache	Drive slots	Power supply
2072-12C (with two node canisters)	8/16 GB	12 x 3.5-inch	Integrated dual power supplies
2072-24C (with two node canisters)	8/16 GB	24 x 2.5-inch	Integrated dual power supplies
2072-12E (with two node canisters)	N/A	12 x 3.5-inch	Integrated dual power supplies
2072-24E (with two node canisters)	N/A	24 x 2.5-inch	Integrated dual power supplies

iSCSI, FC, and FCoE connectivity with intermix flexibility

With 1Gb iSCSI host interface support, along with optional 8Gb FC or 10Gb iSCSI/FCoE support, Storwize V3700 is designed to accommodate the most common storage network. This broad networking support enables deployment of Storwize V3700 in the existing storage network infrastructure with confidence that the system will support additional network technologies should needs change in future.

1Gb iSCSI host interface ports are standard on all Storwize V3700 control enclosure models. 1Gb iSCSI ports, 8Gb FC ports, or 10Gb iSCSI/FCoE ports may also be added as an option. Table 1-3 lists the configurations available

Table 1-3 Host interface port configurations

Standard	Optional		
1Gb iSCSI	1Gb iSCSI	8Gb FC	10Gb iSCSI/FCoE
4 ports	--	--	--
4 ports	8 ports	--	--
4 ports	--	8 ports	--
4 ports	--	--	4 ports

Figure 1-3 shows the front view of the 2072-12C, and 12E enclosures.



Figure 1-3 IBM Storwize V3700 front view for 2072-12C, and 12E enclosures

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

Figure 1-4 on page 8 shows the front view of the 2072-24C, and 24E enclosures.



Figure 1-4 Storwize V3700 front view for 2072-24C, and 24E enclosure

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

1.2.2 Storwize V3700 features

The following features are available with the Storwize V3700:

- ▶ Thin provisioning (Space efficient disks)

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

- ▶ Volume mirroring

Provides a single volume image to the attached host systems while maintaining pointers to two copies of data in separate storage pools. Copies can be on separate disk in a other enclosure. If one copy failing, Storwize V3700 provides continuous data access by redirecting I/O to the remaining copy. When the copy becomes available, automatic resynchronization occurs.

- ▶ FlashCopy

Storwize V3700 provides a volume level point-in-time copy function to its internal virtualized volumes. This function is designed to create copies for backup, parallel processing, testing, and development, and have the copies available almost immediately.

Storwize V3700 includes the following FlashCopy functions:

- Full / Incremental copy

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

- Multitarget FlashCopy

Storwize V3700 supports multiple copies from a single source. Each copy is managed by a unique mapping and, in general, each mapping acts independently and is not affected by other mappings sharing the source volume.

- Cascaded FlashCopy

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

- Reverse FlashCopy

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

– FlashCopy nocopy with thin provisioning

This function provides a combination of using thin-provisioned volumes and FlashCopy together to help reduce disk space requirements when making copies. There are two variations of this option:

- Space-efficient source and target with background copy

Copies only the allocated space.

- Space-efficient target with no background copy

Copies only the space used for changes between the source and target and is generally referred to as “snapshots”.

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

► Consistency groups

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

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

► One way Data Migration (no charge for temporary usage)

Storwize V3700 provides a data migration function that can be used to import external storage systems into the Storwize V3700 system.

It allows you to:

- Move volumes nondisruptively onto a newly installed storage system
- Migrate data from other storage systems to Storwize V3700 managed storage

► The Storwize V3700 can be virtualized from systems such as the IBM SAN Volume Controller (SVC), IBM Storwize V7000, or IBM Flex System V7000 Storage Node.

1.3 Storwize V3700 hardware

The Storwize V3700 solution is a modular storage system that is built on a common enclosure (control enclosure and expansion enclosure).

Figure 1-5 shows an overview of the hardware components of the Storwize V3700 solution.

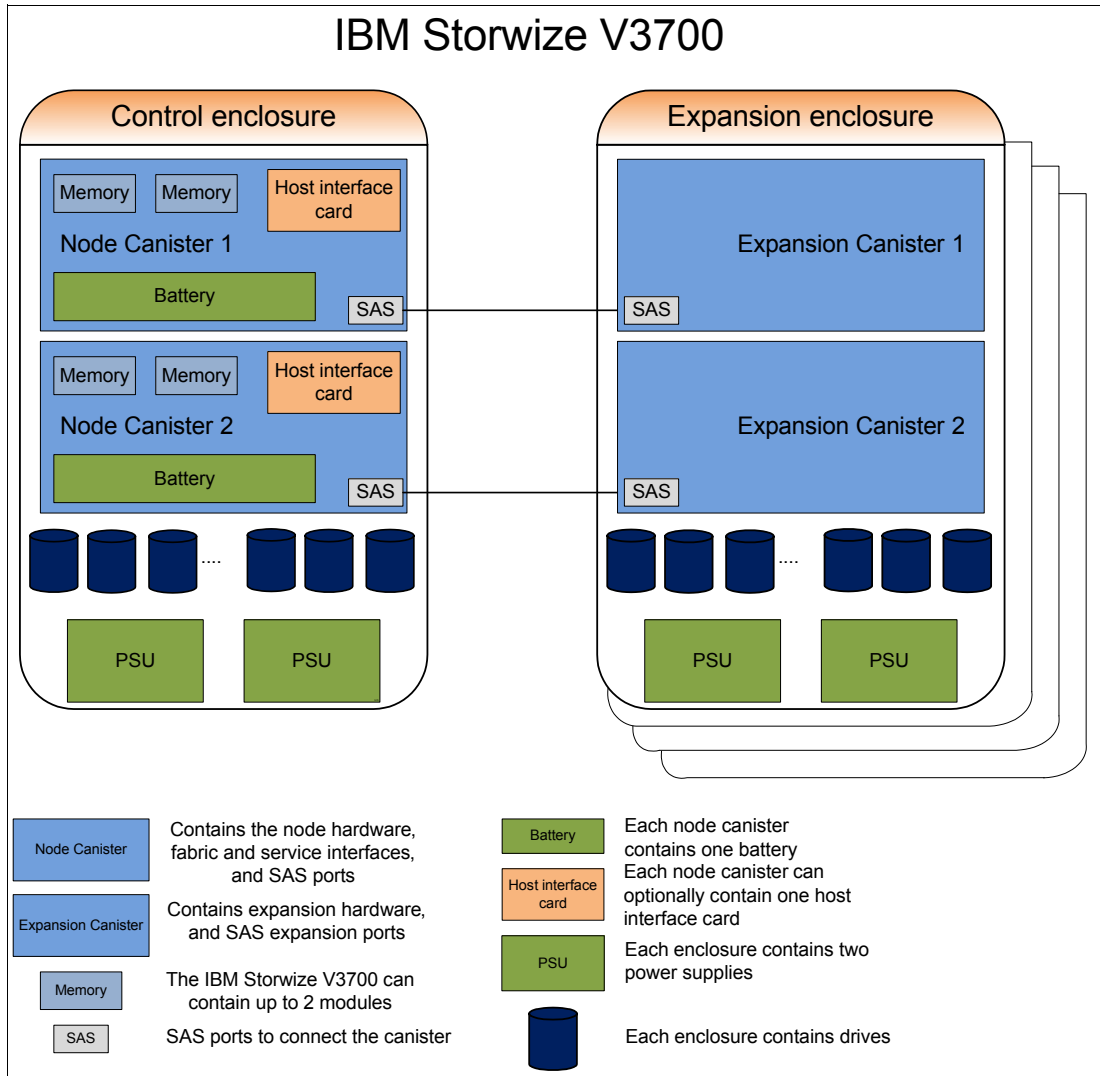


Figure 1-5 Storwize V3700 hardware components

Figure 1-6 shows the controller rear view of IBM Storwize V3700 models 12C and 24C.



Figure 1-6 IBM Storwize V3700 controller rear view - models 12C and 24C

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

There are two *canister* slots at the top of the chassis. In Figure 1-7 you can see the rear of an IBM Storwize V3700 expansion enclosure.



Figure 1-7 IBM Storwize V3700 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 enclosure has only the two SAS ports.

For more details of the expansion enclosure see “Expansion enclosure” on page 13.

1.3.1 Control enclosure

Each Storwize V3700 system has one control enclosure which contains two *node canisters*, disk drives, and two power supplies. There are two models of the control enclosure with providing 12 3.5-inch disk slots (model 12C) and the other providing 24 2.5-inch disk slots (model 24C).

Figure 1-8 shows a single node canister.



Figure 1-8 Storwize V3700 node canister

Each node canister contains the following hardware:

- ▶ Battery
- ▶ Memory - 4 or 8 GB memory that is configurable
- ▶ Host interface card - different options are possible
- ▶ Four 6 Gb/s SAS ports – only port 4 is used for additional enclosure connections
- ▶ 10/100/1000 Mb/s Ethernet ports
- ▶ Two USB 2.0 ports (one port used during install)
- ▶ System flash

The battery is used in case of power loss. The Storwize V3700 system needs these batteries to write the cache data to the internal system flash. This dump is called a fire hose dump. After the system is up again this data is loaded back to the cache to destage it to the disks.

With the Storwize V3700 you have the possibility to order a system with 8 or 16 GB in total. This means that each node canister can have 4 or 8 GB memory.

In Figure 1-9 you see additionally the different ports at the rear which are provided by the Storwize V3700 node canister.




Figure 1-9 Storwize V3700 node ports

The Storwize V3700 offers you the opportunity to select different host interface cards. For the position of the card see number “9” in Figure 1-9 on page 12.

You can choose:

- ▶ 4 x 1Gb/s ethernet ports (iSCSI)
- ▶ 2 x 10Gb/s ethernet ports (iSCSI/FOC)
- ▶ 4 x 8Gb/s Fibre Channel ports (FC)

The Fibre Channel ports operate at 2, 4, or 8 Gbps. The ports of the host interface cards are always numbered from the left to the right starting with 1. For additional info see infocenter: http://pic.dhe.ibm.com/infocenter/storwize/v3700_ic/index.jsp under: **Fibre Channel host interface card ports and indicators**.

The Storwize V3700 is delivered with 4 SAS (serial attached SCSI) ports. They are numbered 1 on the left to 4 on the right. The Storwize V3700 only uses port 4 (the extreme right one) to connect to the next optional expansion enclosure. See number “3” and “4” in Figure 1-9 on page 12.

The two 10 / 100 / 1000 Mbps ethernet ports shown in Figure 1-9 on page 12 number “1” and “2” are used for management. Port 1 (left port) has to be configure. The second port is optional to use for management. Furthermore, both ports can be used for iSCSI traffic. Refer to Chapter 4, “Host configuration” on page 135 for more information.

There are two USB ports numbered with “6” and ”7” in Figure 1-9 on page 12. One is used during the initial configuration or in case of problem. They are numbered 1 on the left and 2 on the right. For more details about usage refer to Chapter 2, “Initial configuration” on page 27.

As already mentioned before the system flash is used to store the dump files in case of power loss and of course for the operating system itself.

Service port: Do *not* use the port marked with “5” in Figure 1-9 on page 12. This port is a service port only.

The two nodes act as a single processing unit and form an I/O group that is attached to the SAN fabric. The pair of nodes is responsible for serving I/O to a given 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 called *I/O groups*.

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

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

1.3.2 Expansion enclosure

The optional expansion enclosure of the Storwize V3700 contains two *expansion canisters*, disk drives, and two power supplies. There are two models of the expansion enclosure with one model providing 12 disk slots and the other providing 24 disk slots.

Figure 1-10 gives an overview of the expansion enclosure.



Figure 1-10 Expansion enclosure of the IBM Storwize V3700

The expansion enclosure power supplies are similar to the control enclosure. There is a single power lead connector on the power supply unit.

Figure 1-11 shows expansion canister ports.




Figure 1-11 expansion canister ports

Each expansion canister showing in Figure 1-11 on page 13 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, as the expansion enclosure must be attached to either a control enclosure or another expansion enclosure. SAS connector 2 is optional, as it is used to attach to additional expansion enclosures.

Each port has 2 LEDs to show the status. The first LED described the link status and the second LED indicates the fault status.

For more details about LED or ports see the infocenter:

http://pic.dhe.ibm.com/infocenter/storwize/v3700_ic/index.jsp

1.3.3 Disk drive types

Storwize V3700 enclosures currently support SSD, SAS, and Nearline SAS drive types. Each SAS drive has two ports (two PHYs) and I/O can be issued down both paths simultaneously.

Table 1-4 shows the Storwize V3700 Disk Drive types that are available at the time of writing.

Table 1-4 Storwize V3700 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	300,600, and 900 GB
2.5-inch form factor	SAS	15,000 rpm	146 and 300 GB
2.5-inch form factor	Nearline SAS	7,200 rpm	500 GB and 1 TB
3.5-inch form factor	Nearline SAS	7,200 rpm	2 and 3 TB

1.4 Storwize V3700 terms

1.4.1 Hosts

A host system is a server that is connected to Storwize V3700 through a Fibre Channel connection or through an iSCSI connection.

Hosts are defined to Storwize V3700 by identifying their worldwide port names (WWPNs) for Fibre Channel hosts. For iSCSI hosts, they are identified by using their iSCSI names. The iSCSI names can either be iSCSI qualified names (IQNs) or extended unique identifiers (EUIs). See Chapter 4, “Host configuration” on page 135 for more information.

A big benefit is that you can attach the host direct to the Storwize V3700. So if you do not have any FC switches you can connect hosts without any FC network directly.

1.4.2 Node canister

Storwize V3700 has a hardware component called the node canister. It is the node, which is a logical object within the node canister, that provides the virtualization of internal volumes, and cache and a copy service (FlashCopy) function. A clustered system consists of a one node pair.

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

1.4.3 I/O groups

Within Storwize V3700, there are two node canisters known as *I/O group*. The Storwize V3700 supports two node canisters in the clustered system, which provides one I/O group.

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 512 volumes per hosts). However, both nodes also act as a failover node for its partner node within the I/O group, so a node takes over the I/O workload from its partner node, if required, with no impact to the server's application.

In a Storwize V3700 environment, using active-active architecture, the I/O handling for a volume can be managed by both nodes of the I/O group. Therefore, it is mandatory for servers that are connected through Fibre Channel connectors to use multipath device drivers to be able to handle this capability.

The Storwize V3700 I/O groups are connected to the SAN so that all application servers accessing volumes from the I/O group have access to them. Up to 256 host server objects can be defined to one Storwize V3700.

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 typical of an active / passive architecture.

1.4.4 Clustered system

All configuration, monitoring, and service tasks are performed at the system level and the configuration settings are replicated across both node canisters in the clustered system. To facilitate these tasks, one or two management IP addresses are set for the clustered system. This allows you to manage the clustered system as a single entity.

There is a process provided to back up the system configuration data on to disk so that the clustered system can be restored in the event of a disaster. This method does not back up application data, only Storwize V3700 system configuration information.

System configuration backup: After backing up the system configuration, save the backup data on your hard disk (or at the least outside of the SAN). If you are unable to access the Storwize V3700, you do not have access to the backup data if it is on the SAN. Do this after each configuration change to be on the save side.

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

The system can be configured using either the Storwize V3700 management software (GUI) or the command-line interface (CLI) or the USB key.

1.4.5 RAID

The Storwize V3700 setup contains a number of internal drives, but these drives cannot be directly added to storage pools. The drives need to be included in a Redundant Array of Independent Disks (RAID) to provide protection against the failure of individual drives.

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

Storwize V3700 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 desired 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 available RAID levels are:

- ▶ RAID 0 (striping, no redundancy)
- ▶ RAID 1 (mirroring between two drives)
- ▶ RAID 5 (striping, can survive one drive fault, with parity)
- ▶ RAID 6 (striping, can survive two drive faults, with parity)
- ▶ RAID 10 (RAID 0 on top of RAID 1)

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

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

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

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

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

1.4.6 Managed disks

A managed disk (MDisk) refers to the unit of storage that Storwize V3700 virtualizes. This unit is a RAID array consisting of internal drives. The Storwize V3700 can then allocate these MDisks into various storage pools. The Storwize V3700 can also pick up MDisk from external storage but only during the migration process.

An MDisk is not visible to a host system on the storage area network, as it is internal in the Storwize V3700 system.

An MDisk has two modes:

- ▶ Array
 - Array mode MDisks are constructed from drives using the RAID function. Array MDisks are always associated with storage pools.
- ▶ Image

Image MDisks are assigned directly to a volume with a one-to-one mapping of extents between the MDisk and the volume. For more detail refer to Chapter 6, “Storage migration wizard” on page 209.

1.4.7 Quorum disks

The clustered system uses quorum disks to break a tie when exactly half the nodes in the system remain after a SAN failure.

The clustered system automatically forms the quorum disk by taking a small amount of space from a local disk. It allocates space from up to three different disks for redundancy, although only one quorum disk is active. It is possible to manage the quorum disks by using the CLI.

1.4.8 Storage pools

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

MDisks can be added to a storage pool at any time to increase the capacity of the storage pool. MDisks can belong in only one storage pool. For more details refer to Chapter 7, “Storage pools” on page 267.

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

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

The effect of extent size on the maximum volume size is shown in Table 1-5, and lists the extent size and the corresponding maximum clustered system size.

Table 1-5 Maximum volume capacity by extent size

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

Use the same extent size for all storage pools in a clustered system, which is a prerequisite if you want to migration 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 267.

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

Default extent size: The GUI of IBM Storwize V3700 has a default extent size value of 256 MB when you define a new storage pool.

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

1.4.9 Volumes

A volume is a logical disk that is presented to a host system by the clustered system. In our virtualized environment, the host system has a volume mapped to it by Storwize V3700. Storwize V3700 translates this volume into a number of extents, which are allocated across MDisks. The advantage with storage virtualization is that the host is “decoupled” from the underlying storage, so the virtualization appliance can move the extents around without impacting the host system.

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

There are three types of volumes:

- ▶ Striped

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

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

Figure 1-12 shows how a striped volume is allocated, assuming 10 extents are required,

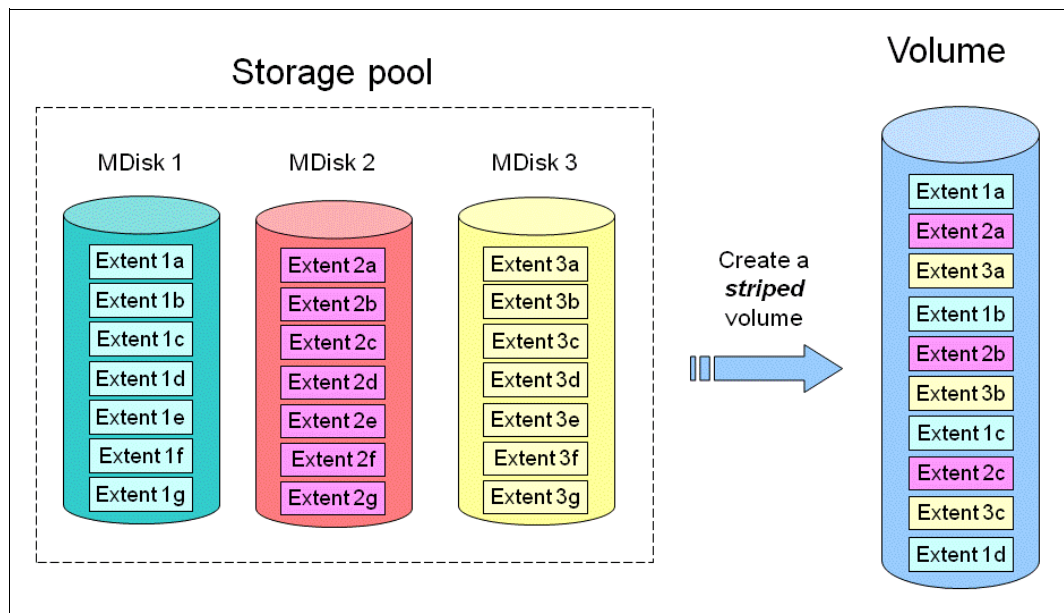


Figure 1-12 Striped volume

► Sequential

A sequential volume is where the extents are allocated one after the other from one MDisk to the next MDisk (Figure 1-13).

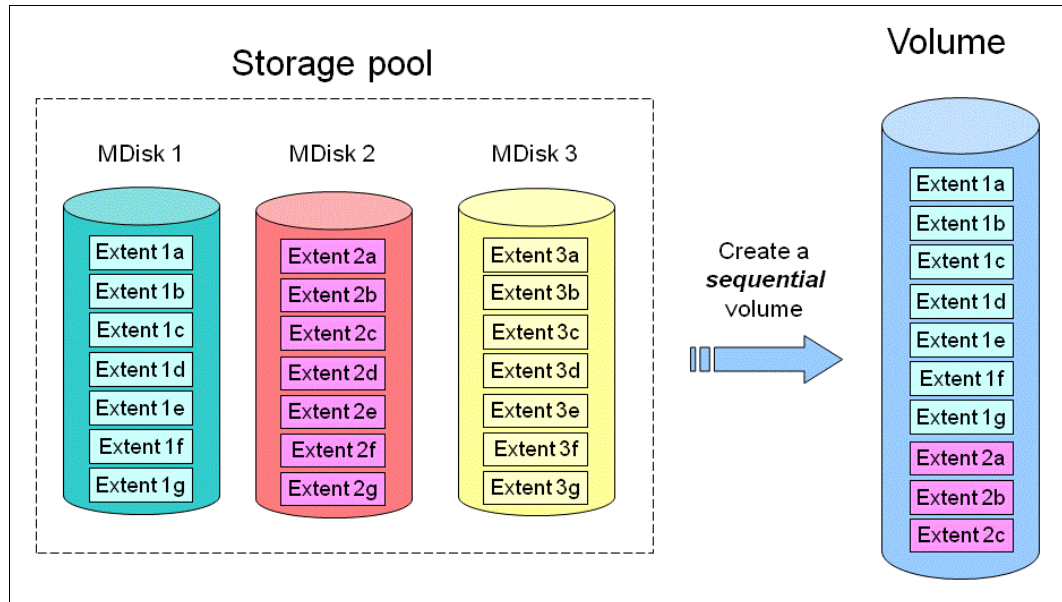


Figure 1-13 Sequential volume

► Image mode

Image mode volumes are special volumes that have a direct relationship with one MDisk. They are used to migrate existing data into and out of the clustered system.

When the image mode volume is created, a direct mapping is made between extents that are on the MDisk and the extents that are on the volume. The logical block address (LBA) x on the MDisk is the same as the LBA x on the volume, which ensures that the data on the MDisk is preserved as it is brought into the clustered system (Figure 1-14).




Figure 1-14 Image mode volume

On the Storwize V3700 the image mode volumes is only use to migrate the volume into the Storwize V3700. After it is migrated you can disconnected the migrated subsystem.

If want to migrate data from an existing storage subsystem use the “Storage Migration Wizard” which will guide you those all steps carefully.

For more details see Chapter 6, “Storage migration wizard” on page 209.

1.4.10 Thin-provisioned volumes

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

The real capacity determines the quantity of MDisk extents that are allocated for the volume. The virtual capacity is the capacity of the volume reported to Storwize V3700 and to the host servers.

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

The thin provisioning feature can be used on its own to create overallocated volumes, 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 Storwize V3700 to automatically expand the real capacity of a thin-provisioned volume as its real capacity is used. This parameter will prevent 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 is initially set to the real capacity that is assigned when the volume is created. If the user modifies the real capacity, the contingency capacity is reset to be the difference between the used capacity and real capacity.

A volume that is created with a zero contingency capacity goes offline as soon as it needs to expand. A volume with a non-zero contingency capacity stays online until it has been 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% has been specified, the warning is logged when 20% of the free capacity remains.

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

1.4.11 Mirrored volumes

Storwize V3700 provides a function called *Storage volume mirroring*, which enables a volume to have two physical copies. Each volume copy can belong to a different storage pool.

When a host system issues a write to a mirrored volume, Storwize V3700 writes the data to both copies. When a host system issues a read to a mirrored volume, Storwize V3700 places it into the primary copy. If one of the mirrored volume copies is temporarily unavailable, the Storwize V3700 automatically uses the alternative copy without any outage for the host system. When the mirrored volume copy is repaired, Storwize V3700 resynchronizes the data.

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

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

1.4.12 iSCSI

iSCSI is an alternative means of attaching hosts to the Storwize V3700. The iSCSI function is a *software function* that is provided by the Storwize V3700 code, not hardware.

In the simplest terms, iSCSI allows the transport of SCSI commands and data over a TCP/IP network, 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.

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

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

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

The concepts of names and addresses have been 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 not only the iSCSI name of an iSCSI node, but also a location of that node. The address consists of a host name or IP address, a TCP port number (for the target), and the iSCSI name of the node. An iSCSI node can have any number of addresses, which can change at any time, particularly if they are assigned by way of Dynamic Host Configuration Protocol (DHCP). An IBM Storwize V3700 node represents an iSCSI node and provides statically allocated IP addresses.

Each iSCSI node, that is, an initiator or target, has a unique iSCSI Qualified Name (IQN), which can have a size of up to 255 bytes. The IQN is formed according to the rules adopted for Internet nodes.

The iSCSI qualified name format is defined in RFC3720 and contains (in order) these elements:

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

For IBM Storwize V3700, the IQN for its iSCSI target is specified as:

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

On a Windows server, the IQN, that is, the name for the iSCSI initiator, can be defined as:

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

The IQNs can be abbreviated by using a descriptive name, known as an *alias*. An alias can be assigned to an initiator or a target. The alias is independent of the name and does not need to be unique. Because it is not unique, the alias must be used in a purely informational way. It cannot be used to specify a target at login or used during authentication. Both targets and initiators can have aliases.

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

Changing names: Before changing system or node names for an IBM Storwize V3700 clustered system that has servers connected to it using SCSI, be aware that because the system and node name are part of the IQN for the IBM Storwize V3700, you can lose access to your data by changing these names. The IBM Storwize V3700 GUI shows a specific warning, but the CLI does *not*.

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

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

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

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

For further details about configuring iSCSI, refer to Chapter 4, “Host configuration” on page 135.

1.5 FlashCopy

FlashCopy makes a copy of a source volume on a target volume. The original content of the target volume is lost. After the copy operation has started, the target volume has the contents of the source volume as they 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.

Storwize V3700 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. Storwize V3700 supports multiple targets and thus multiple rollback points.

Note: There is a limitation of only 64 copies in total. For more information refer to Chapter 8, “Advanced host and volume administration” on page 303.

1.6 Problem management and support

1.6.1 IBM Assist On-site and remote service

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

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

To use the IBM Assist On-site tool, the management PC which is used to manage the Storwize V3700 must be able to access the Internet. The following website provides further information about this tool:

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

Your IBM service representative can provide you with more detailed instructions for using the tool if required.

1.6.2 Event notifications

Storwize V3700 can use Simple Network Management Protocol (SNMP) traps, syslog messages, and a Call Home email to notify you and the IBM Support Center when significant events are detected. Any combination of these notification methods can be used simultaneously.

Each event that Storwize V3700 detects will be sent to different recipient. You can configure Storwize V3700 to send each type of notification to specific recipients or only the alerts which are important to the system.

1.6.3 SNMP traps

SNMP is a standard protocol for managing networks and exchanging messages. Storwize V3700 can send SNMP messages that notify personnel about an event. You can use an SNMP manager to view the SNMP messages that Storwize V3700 sends. You can use the management GUI or the IBM Storwize V3700 command-line interface to configure and modify your SNMP settings.

You can use the Management Information Base (MIB) file for SNMP to configure a network management program to receive SNMP messages that are sent by the Storwize V3700. This file can be used with SNMP messages from all versions of Storwize V3700 Software.

1.6.4 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 either IPv4 or IPv6. Storwize V3700 can send syslog messages that notify personnel about an event. Storwize V3700 can transmit syslog messages in either expanded or concise format. You can use a syslog manager to view the syslog messages that Storwize V3700 sends. Storwize V3700 uses the User Datagram Protocol (UDP) to transmit the syslog message. You can use the management GUI or the Storwize V3700 command-line interface to configure and modify your syslog settings.

1.6.5 Call Home email

The Call Home feature transmits operational and error-related data to you and IBM through a Simple Mail Transfer Protocol (SMTP) server connection in the form of an event notification email. When configured, this function alerts IBM service personnel about hardware failures and potentially serious configuration or environmental issues. You can use the call home function if you have a maintenance contract with IBM or if the Storwize V3700 is within the warranty period.

To send email, you must configure at least one SMTP server. You can specify as many as five additional SMTP servers for backup purposes. The SMTP server must accept the relaying of email from the Storwize V3700 clustered system IP address. You can then use the management GUI or the Storwize V3700 command-line interface 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 using the management GUI or the Storwize V3700 command-line interface.

1.7 Useful Storwize V3700 websites

The Storwize V3700 home page is at the following website:

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

The Storwize V3700 Configuration Limit and Restrictions are at the following website:

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

The Storwize V3700 Online Information Center is at the following website:

http://pic.dhe.ibm.com/infocenter/storwize/v3700_ic/index.jsp

You can see IBM Redbooks publications about Storwize V3700 at the following website:

<http://www.redbooks.ibm.com/cgi-bin/searchsite.cgi?query=V3700&SearchOrder=1&SearchFuzzy=>

1.8 IBM Storwize V3700 learning videos on YouTube

Some interesting videos are available on YouTube that describe the IBM Storwize GUI and are available at the URLs shown in Table 1-6.

Be aware that these videos are not only for Storwize V3700; as the interface to the Storwize V7000 is exceptionally similar, they give you an overview of what is possible, and introduce the look and feel.

Table 1-6 Videos available on YouTube

Description of the video	Link WEB
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



Initial configuration

This chapter provides a description of the initial configuration steps for the IBM Storwize V3700.

In this chapter, we describe the following topics:

- ▶ Planning for IBM Storwize V3700 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. Here is a checklist of requirements to use to plan your installation:

- Install the hardware as described in *IBM Storwize V3700 Quick Installation Guide Version 6.4.1, GC27-4219*
- For information about planning the IBM Storwize V3700 environment, see the following address:

http://pic.dhe.ibm.com/infocenter/storwize/v3700_ic/index.jsp?topic=%2Fcom.ibm.storwize.v3700.641.doc%2Fsvc_webplanning_21pb8b.html

- An appropriate 19-inch rack with 2 - 10U of space should be available depending on the number of enclosures to install. Each enclosure measures 2U and two control enclosure with up to four expansion enclosures constitutes an IBM Storwize V3700 system.
- There should be redundant power outlets in the rack for each of the two power cords included per enclosure. 2 - 10 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 (FC) ports attached to the fabric are required, but it is a best practice to use eight 2-, 4-, or 8- Gbps Fibre Channel ports

FC ports: FC ports are required only if you are using FC hosts. You can use the Storwize V3700 with Ethernet-only cabling for iSCSI hosts.

- You should have eight 2-, 4-, or 8- Gbps compatible Fibre Channel cable drops.
- You should have a minimum of two Ethernet ports on the LAN, with four preferred for additional configuration access redundancy or iSCSI host access.
- You should have a minimum of two Ethernet cable drops, with four preferred for additional configuration access redundancy or iSCSI host access. Ethernet port one on each node canister must be connected to the LAN with port two being 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 address 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 used with these IP addresses is 255.255.255.0 and the default gateway address 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 additional 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 V3700 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 V3700.

- One single 1-, 3-, or 6-meter SAS cable per expansion enclosure is required. The length of the cables depends on the physical rack location they are installed in relative to the control enclosure or other expansion enclosures. Locate the control enclosure so that up to four enclosures can be located below, as shown in Figure 2-1.




Figure 2-1 Connecting the SAS cables example

- 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.

Disk drives: The disk drives included with the control enclosure, model 2076-224 or 2076-212, are part of the single SAS chain. The expansion enclosure should be connected to the SAS chain as shown so that they can take advantage of the full bandwidth of the system.

2.2 SAN configuration planning

The recommended SAN configuration is composed of a minimum of two fabrics with all host ports, and any ports on external storage systems to be virtualized by IBM Storwize V3700. The IBM Storwize V3700 ports themselves are evenly split between the two fabrics to provide redundancy in the event one of the fabrics goes offline (either planned or unplanned).

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

After the IBM Storwize V3700, hosts, and optional external storage systems are connected to the SAN fabrics, zoning needs to be implemented.

In each fabric, create a zone with just the four IBM Storwize V3700 WWPNs, two from each node canister. If there is an external storage system to be virtualized, then in each fabric, create a zone with the four IBM Storwize V3700 WWPNs, two from each node canister, along with up to a maximum of eight WWPNs from the external storage system. Assuming every host has a Fibre Channel connection to each fabric, then in each fabric, create a zone with the host WWPN and one WWPN from each node canister in the IBM Storwize V3700 system.

Maximum ports or WWPNs: IBM Storwize V3700 supports a maximum of 16 ports or WWPNs from a given virtualized external storage system that is virtualized.

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




Figure 2-2 SAN cabling and zoning diagram

Create a host / IBM Storwize V3700 zone for each server to which volumes are mapped to from the clustered system. For example:

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

Verify interoperability with which the IBM Storwize V3700 connects to SAN switches or directors by following the requirements provided in the following link:

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

Switches or directors are at the firmware levels supported by the IBM Storwize V3700.

- ▶ IBM Storwize V3700 port login maximum listed in restriction document must not be exceeded.

Connectivity issues: If you have any connectivity issues between IBM Storwize V3700 ports and Brocade SAN Switches or Directors at 8 Gbps, see <http://www-01.ibm.com/support/docview.wss?rs=591&uid=ssg1S1003699> for the correct setting of the fillword port config parameter in the Brocade operating system.

2.3 Direct attach planning

IBM Storwize V3700 can be used with direct attach Fibre Channel (FC) server configuration. The recommended configuration for direct attachment is to have at least one FC cable from the server connected to each node of the IBM Storwize V3700 to provide redundancy in the event one of the nodes goes offline (Figure 2-3).

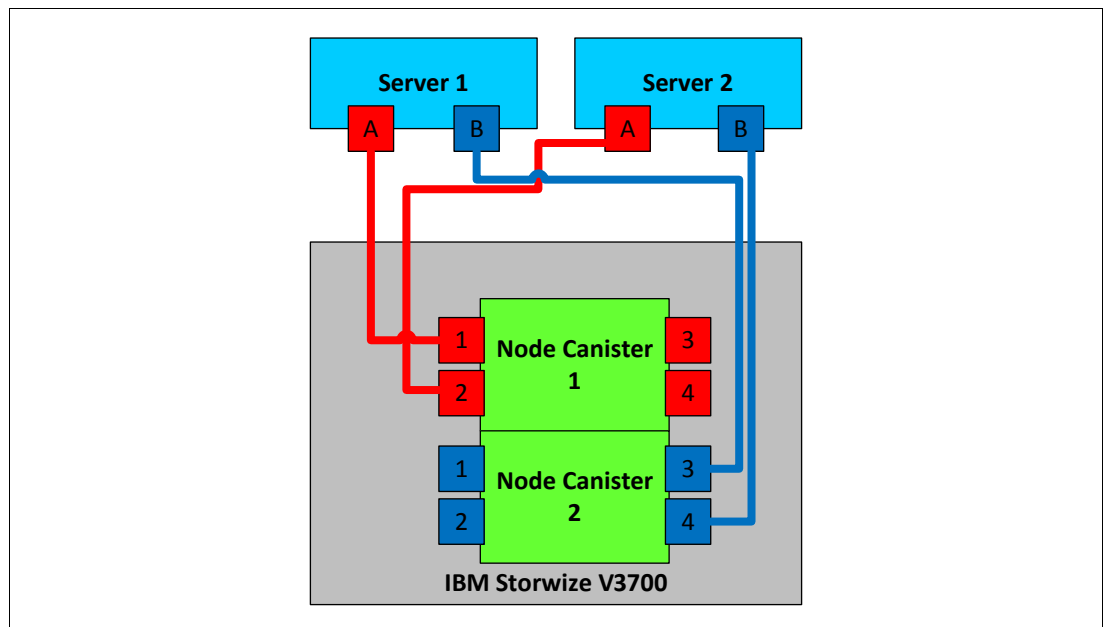


Figure 2-3 Direct attach server configuration

Verify direct attach interoperability with the IBM Storwize V3700, and the supported server operating systems by following the requirements provided in the following link:

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

2.4 LAN configuration planning

There are two Ethernet ports per node canister available for connection to the LAN on an IBM Storwize V3700 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 a given Ethernet port or using the same Ethernet port for management and iSCSI access. However, you cannot use the same IP address for both management and iSCSI host use.

Figure 2-4 shows a possible IP configuration of the Ethernet ports on the IBM Storwize V3700 system.

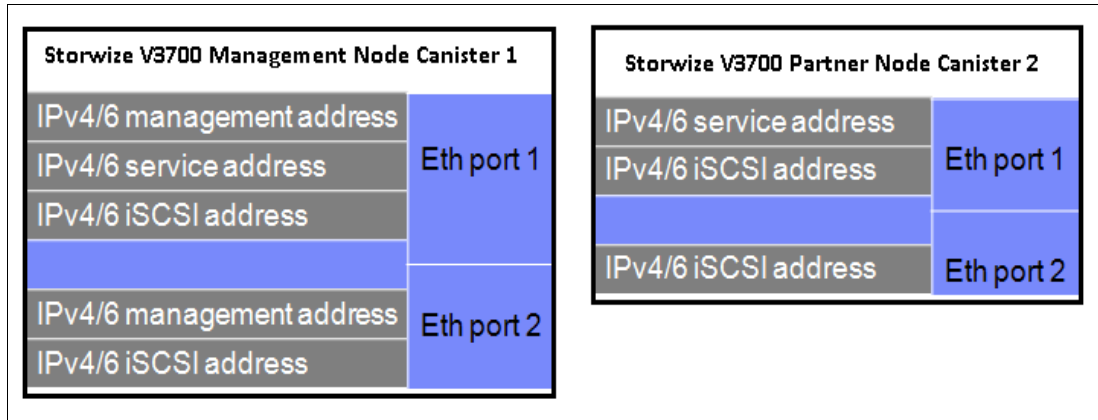


Figure 2-4 Example of IP address configuration options per node canister

IP management addresses: The IP management address shown on node canister 1 is an address on the configuration node; in case of failover, 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).

2.4.1 Management IP address considerations

Because Ethernet port 1 from each node canister must be connected to the LAN, a single management IP address for the clustered system is configured as part of the initial setup of the IBM Storwize V3700 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, either planned or unplanned, the management IP address fails over to the other node's Ethernet port 1.

For additional 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 2-5 on page 33 shows a logical view of the Ethernet ports available for configuration of the one or two management IP addresses. These IP addresses are for the clustered system itself and therefore only associated with one node, which is then considered the configuration node.

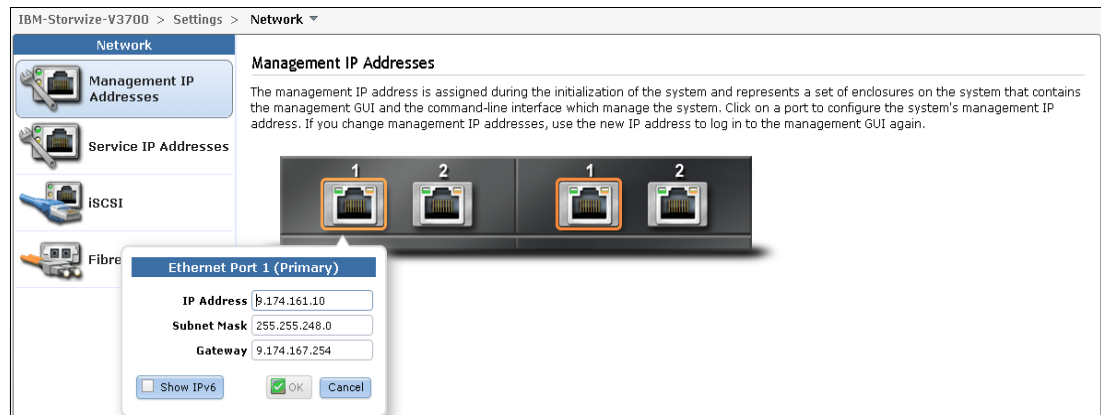


Figure 2-5 Ethernet ports available for configuration

2.4.2 Service IP address considerations

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

Figure 2-6 shows a logical view of the Ethernet ports available for configuration of the service IP addresses. Only port one on each node can be configured with a service IP address.

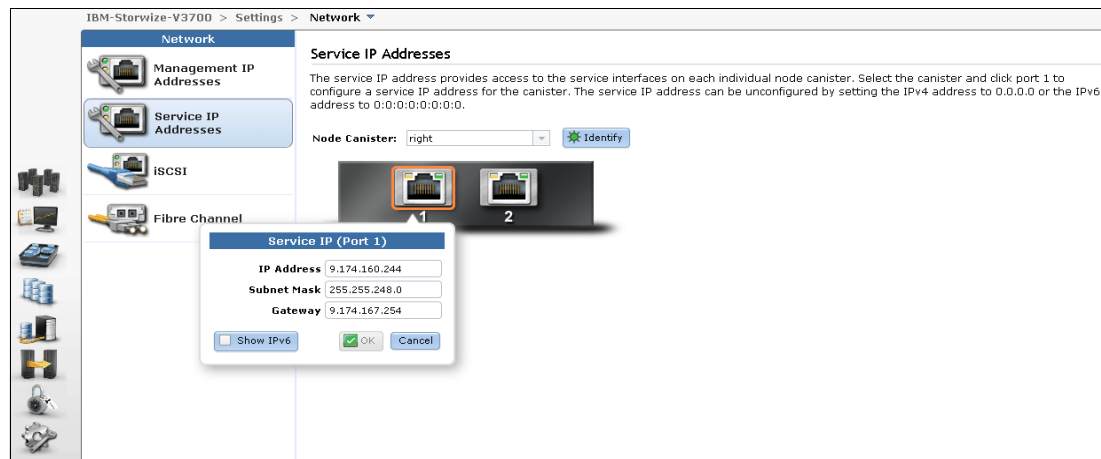


Figure 2-6 Service IP addresses available for configuration

2.5 Host configuration planning

Hosts should have two Fibre Channel connections for redundancy, but IBM Storwize V3700 supports hosts with a single HBA port connection. However, if that HBA, its link to the SAN fabric or the fabric itself fails, then the host loses access to its volumes. Even with just a single connection to the SAN, the host has multiple paths to the IBM Storwize V3700 volumes, because that single connection must be zoned with at least one Fibre Channel port per node. Therefore, a multipath driver is required.

SAN Boot is supported by IBM Storwize V3700. The requirements are listed on the IBM Storwize V3700 support matrix and configuration instructions are provided in the *IBM Storwize V3700 Host Attachment Guide*, which can be found at the following address:

http://pic.dhe.ibm.com/infocenter/storwize/v3700_ic/index.jsp?topic=%2Fcom.ibm.storwize.v3700.641.doc%2Fsvc_hostattachmentmain.html

Verify that the hosts that access volumes from the IBM Storwize V3700 meet the requirements found at the following website:

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

Multiple OS will be supported with IBM Storwize V3700. Refer to the SSIC for details on HBA/Driver/multipath combinations:

<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 supported by the IBM Storwize V3700.
- ▶ HBA BIOS, device drivers, and firmware, along with the multipathing drivers, are at the levels supported by IBM Storwize V3700.
- ▶ If boot from SAN is required, ensure that it is supported for the operating systems to be deployed.
- ▶ If host clustering is required, ensure that it is supported for the operating systems to be deployed.
- ▶ All direct connect hosts should have the HBA set to "point to point".

For more details, see Chapter 4, "Host configuration" on page 135.

2.6 Miscellaneous configuration planning

During the initial setup of the IBM Storwize V3700 system, the installation wizard asks for various information that you should have available during the installation process. Several of these fields are mandatory in order for you to continue the initial configuration.

Here is a checklist of information to gather before performing the initial setup. The date and time can be manually entered, but to keep the clock synchronized, use a network time protocol (NTP) service.

- Document the current LAN NTP server IP address 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 V3700 system for IBM service (for example, Building 22, first floor).
- SMTP or email server address to direct alerts to and from the IBM Storwize V3700.
For the Call Home service to work, the IBM Storwize V3700 system must have access to an SMTP server on the LAN that can forward emails to the default IBM service address `callhome0@de.ibm.com`.
- Email address of local administrators needing to be notified of alerts.
- IP address of SNMP server to direct alerts to if wanted (for example, operations or help desk).

After the IBM Storwize V3700 initial configuration, you might want to add additional users who can manage the system. You can create as many users as you need, but currently there are only five roles generally configured for users:

- ▶ Security Admin
- ▶ Administrator
- ▶ Copyoperator
- ▶ Service
- ▶ Monitor

The Security Admin role allows the user to perform any function on the IBM Storwize V3700.

The Administrator role allows the user to perform any function on the IBM Storwize V3700 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 Copyoperator role allows the user to view anything in the system, but the user can only configure and manage copy functions of the FlashCopy capabilities.

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

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

2.7 System management

The graphical user interface (GUI) is used to configure, manage, and troubleshoot the IBM Storwize V3700 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.

Finally, it offers advanced functions, such as FlashCopy. A command-line interface (CLI) for the IBM Storwize V3700 system is available as well.

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

2.7.1 Graphical user interface (GUI)

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

http://pic.dhe.ibm.com/infocenter/storwize/v3700_ic/index.jsp?topic=%2Fcom.ibm.storwize.v3700.641.doc%2Fsvc_configuringbrowser_1obg15.html

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

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

For more information about how to use this interface, go to the V3700 InfoCenter:

http://pic.dhe.ibm.com/infocenter/storwize/v3700_ic/index.jsp?topic=%2Fcom.ibm.storwize.v3700.641.doc%2Ftbrd_usbgui_1936tw.html

After the initial configuration described in 2.9, “Initial configuration” on page 42 is completed, the IBM Storwize V3700 Welcome window opens as shown in Figure 2-7.

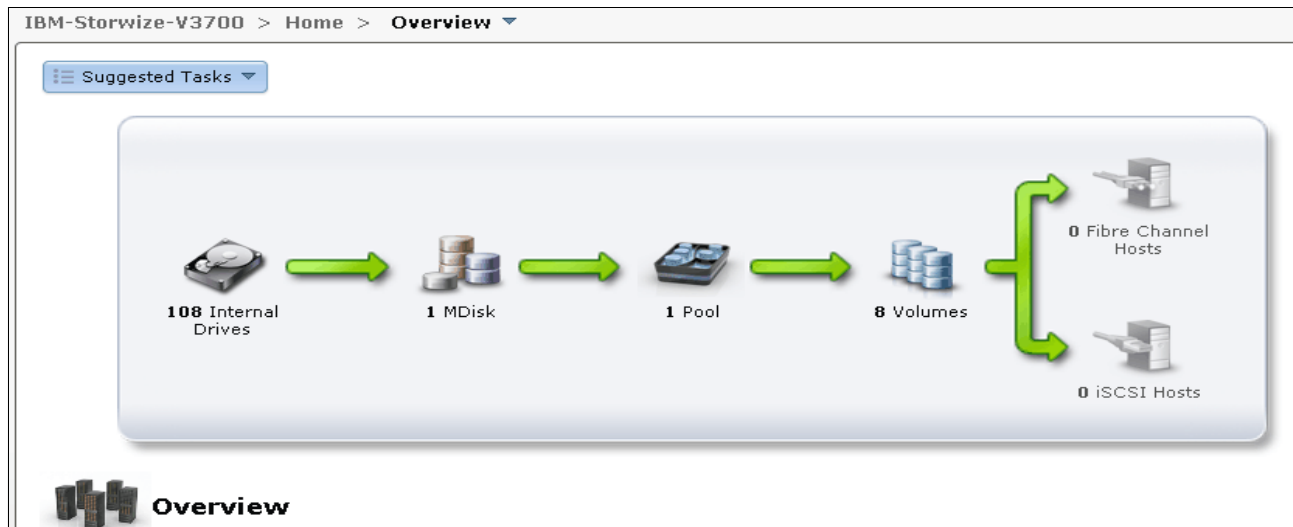


Figure 2-7 Setup Wizard - Welcome window.

2.7.2 Command-line interface

The command-line interface (CLI) is a flexible tool for system management and uses the SSH protocol. A public / private SSH key pair is optional for SSH access. Appendix A, “CLI setup and SAN Boot” on page 471 describes how to set up SSH access for Windows, Linux, or UNIX systems. The storage system can be managed using the CLI (Example 2-1).

Example 2-1 System management using the command-line interface (CLI)

```

IBM_2072:ITS0-Storwize-v3700-2:admin>lsuser
id name      password ssh_key remote usergrp_id usergrp_name
0  superuser  yes      no      no      0          SecurityAdmin
1  Justin     yes      yes     no      0          SecurityAdmin
2  Andrew     yes      no      no      2          CopyOperator
3  Deiter     yes      no      no      3          Service
4  Nancy      yes      no      no      1          Administrator
5  Uwe        yes      no      no      1          Administrator
6  jon        yes      no      no      4          Monitor
IBM_2072:ITS0-Storwize-v3700-2:admin>

```

The initial IBM Storwize V3700 system setup should be done using the process and tools we describe in 2.8, “First-time setup” on page 37.

2.8 First-time setup

This section describes how to perform a first-time IBM Storwize V3700 system setup.

IBM Storwize V3700 uses an easy to use initial setup that is contained within a USB key. The USB key is delivered with each storage system and contains the initialization application called `InitTool.exe`. A system management IP address, the subnet mask, and the network gateway address are required. The initialization application creates a configuration file on the USB key.

The IBM Storwize V3700 starts the initial setup as soon as you plug in the USB key with the newly created file in to the storage system.

USB key: If you are unable to find the official USB key supplied with the IBM Storwize V3700, you can use any USB key you have and download and copy the `initTool.exe` application from the IBM Storwize V3700 Support website at the following address:

<http://www.ibm.com/storage/support/Storwize/v3700>

The USB key contains the `initTool.exe` file (Figure 2-8).


Name ^	Date modified	Type	Size
 InitTool	9/20/2011 10:27 AM	Application	552 KB

Figure 2-8 *InitTool program in the USB key*

To complete the initial setup using the USB key, complete the following steps:

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

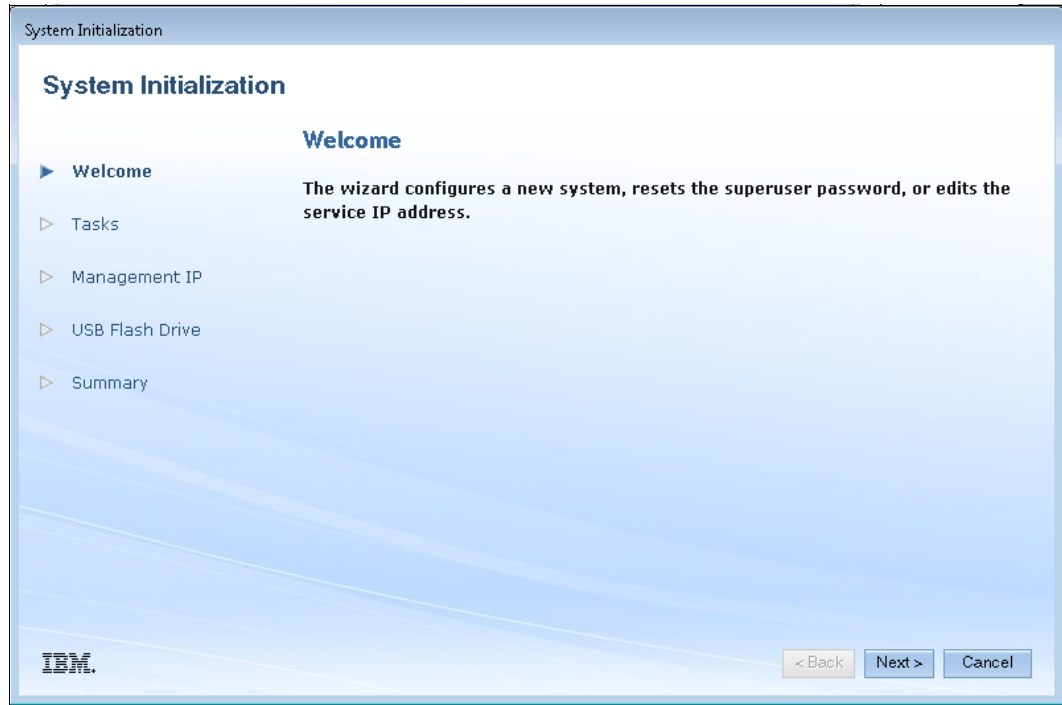


Figure 2-9 System Initialization - Welcome window

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

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

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

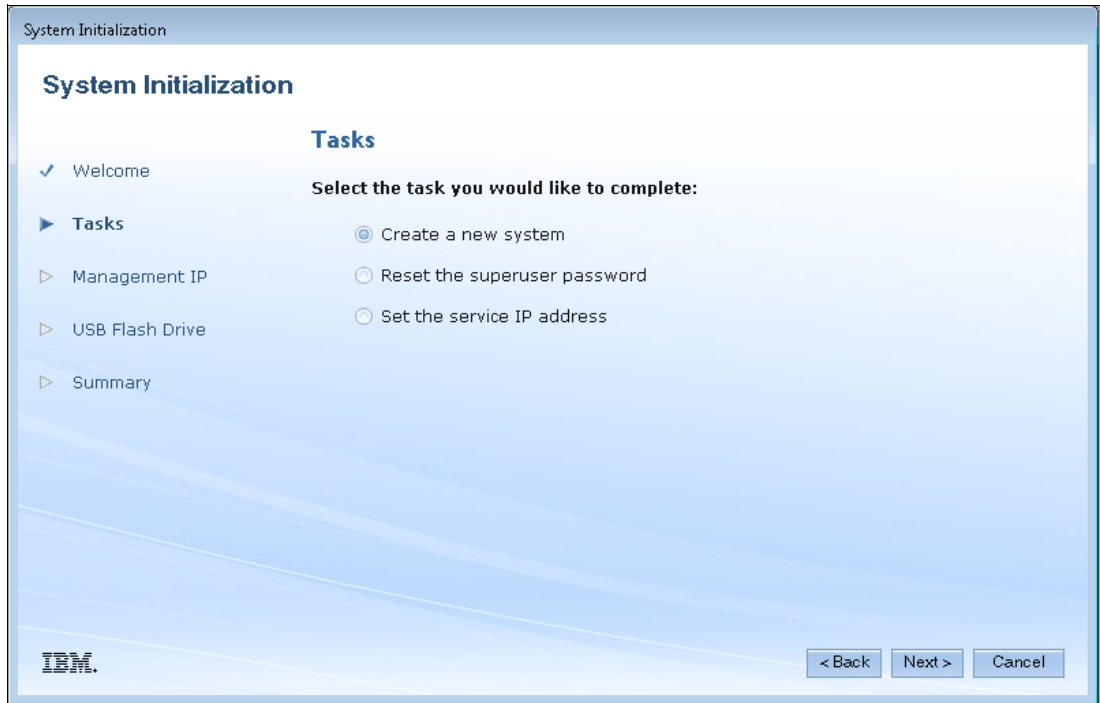


Figure 2-10 System Initialization - Create a new system

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

3. Next is setting the Management IP Address as shown in Figure 2-11.



Figure 2-11 System Initialization - Management IP

4. Click **Apply and Next** to display the USB Flash Drive instructions (Figure 2-12).




Figure 2-12 Initialization application - Finalizing the initial setup

Follow the instructions through the summary portion. Once the initialization completes successfully, click Finish. The application creates a file called `satask.txt` on the USB key (Figure 2-13).

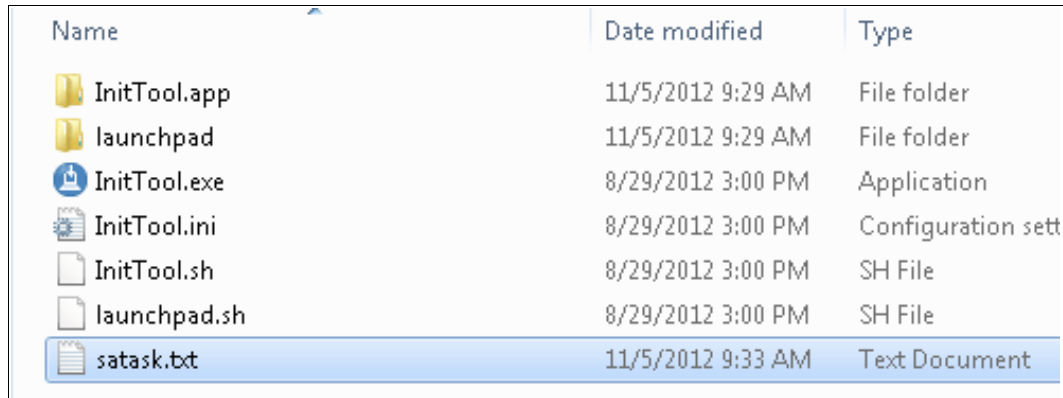


Figure 2-13 The `satask.txt` file created during the `InitTool` procedure

The contents of this file are similar to the following command: `satask mkcluster -clusterip 4.5.6.7 -gw 4.5.6.1 -mask 255.255.255.0`

5. Unplug the USB key from your Windows system and plug it into the IBM Storwize V3700 USB connector slot and the storage system will automatically:
 - a. Detect the key.
 - b. Read the `satask.txt` file and run its contents.
 - c. Delete the `satask.txt` file.
 - d. Write the response to a new file called `satask_result.html`

Clustered system creation: While the clustered system is being created, the amber fault LED on the node canister flashes. When the amber fault LED stops flashing, remove the USB key from IBM Storwize V3700 and insert it in your system to check the results.

6. After this task has completed successfully, the initial setup is done. If you have a network connection to the Storwize system, the system Management GUI is displayed. The IBM Storwize V3700 is available for further configuration changes using the newly defined configuration address.

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

Each IBM Storwize V3700 clustered system has one or two system IP addresses. If the configuration node fails, the system IP addresses are transferred to another node in the same clustered system.

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

2.9 Initial configuration

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

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

Complete the following steps:

1. Start the configuration wizard using a web browser on a workstation and point it to the system management IP address defined in Figure 2-11 on page 39. Type in the default superuser password <passw0rd> (where 0 = zero) Figure 2-14 on page 42



Figure 2-14 Setup Wizard - Login

2. Set up the system name shown in Figure 2-15

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

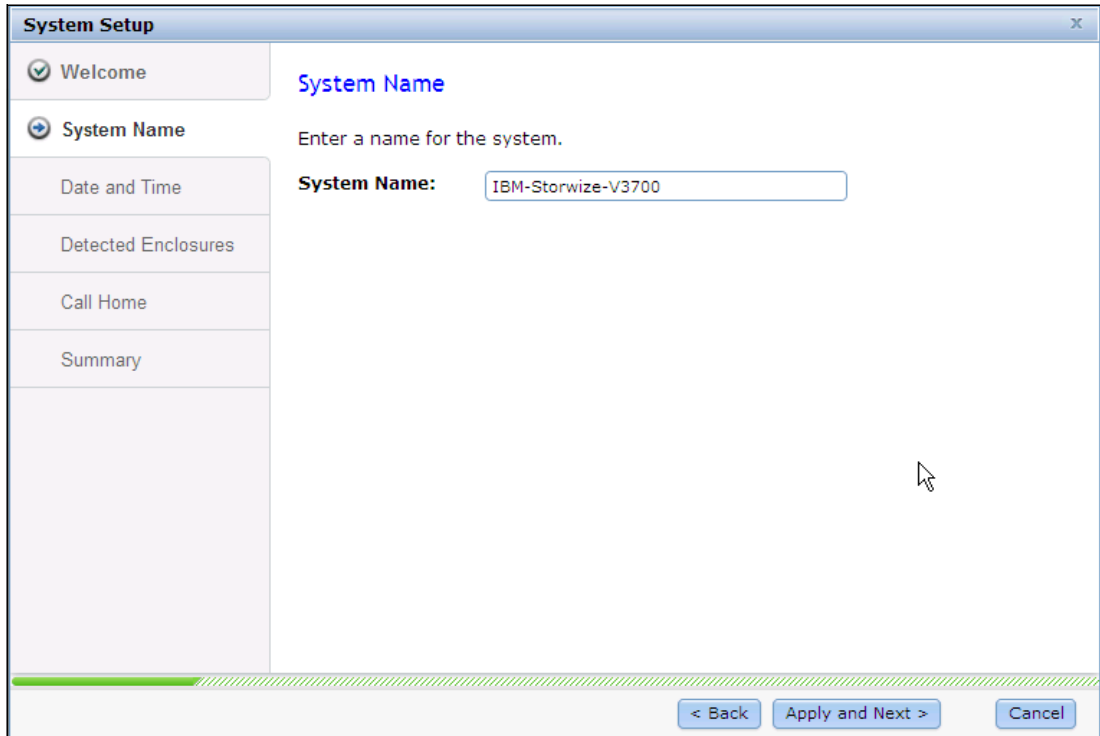


Figure 2-15 Setup Wizard - Insert system name

3. Setup date and time as shown in Figure 2-16 on page 43.

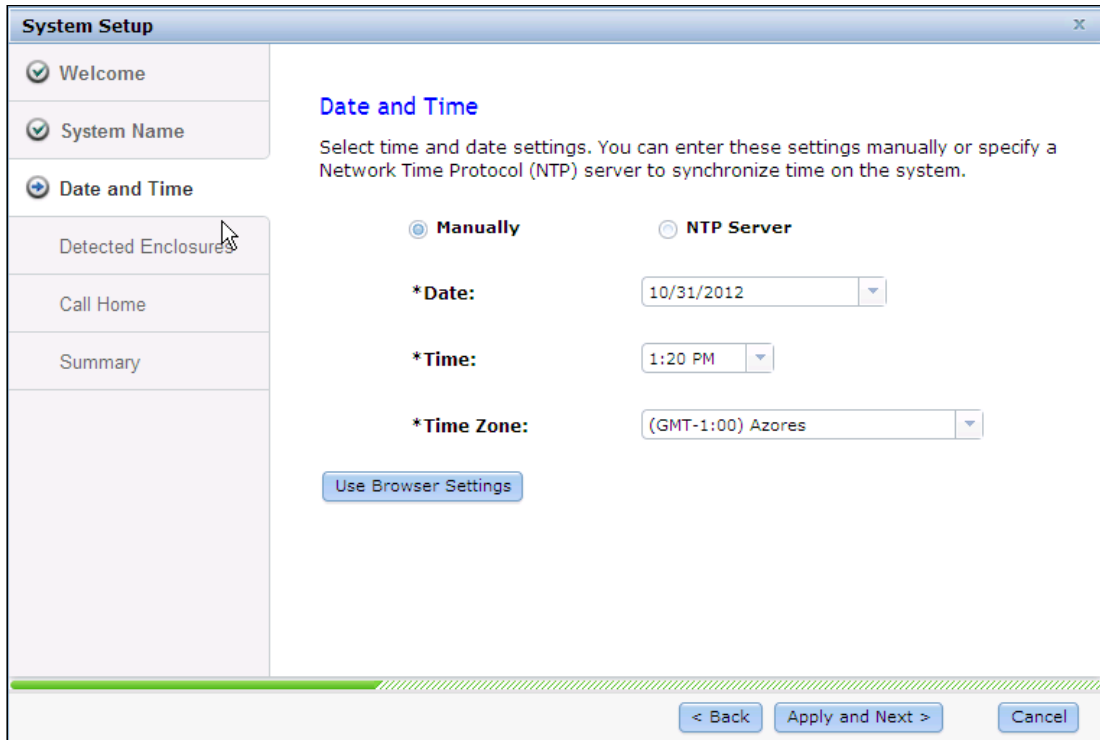


Figure 2-16 Setup Wizard - date and time

4. The configuration wizard continues with the hardware configuration. Verify the hardware (Figure 2-17).

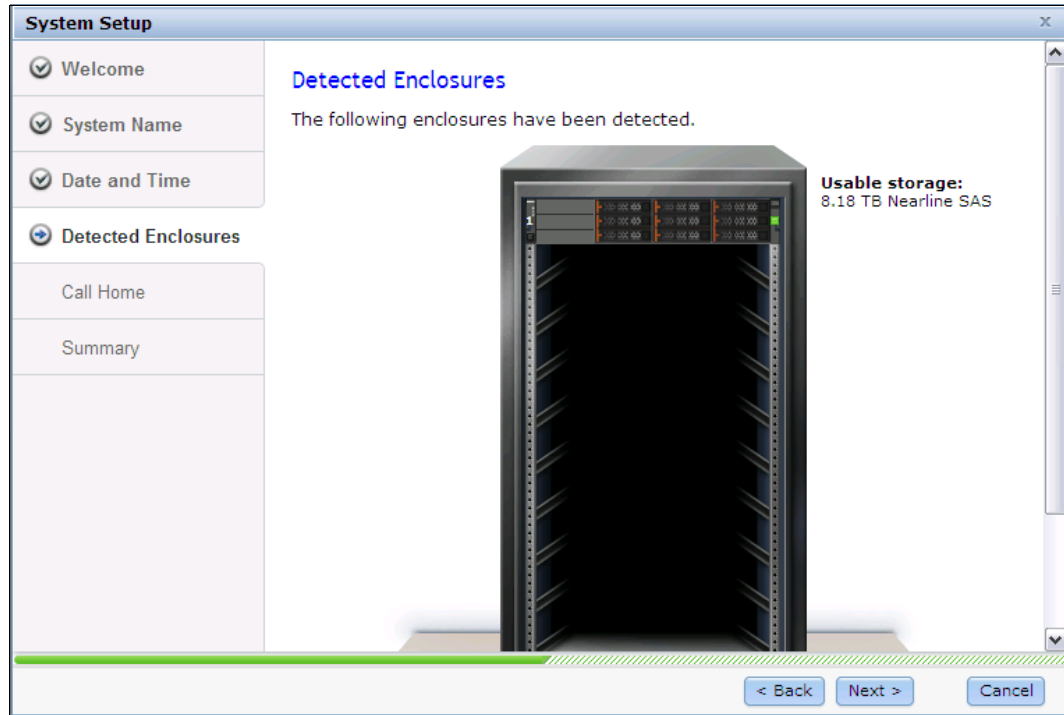


Figure 2-17 Setup wizard - Verify the detected hardware

5. When an array is created, the array members are synchronized with each other by a background initialization process shown in Figure 2-18.

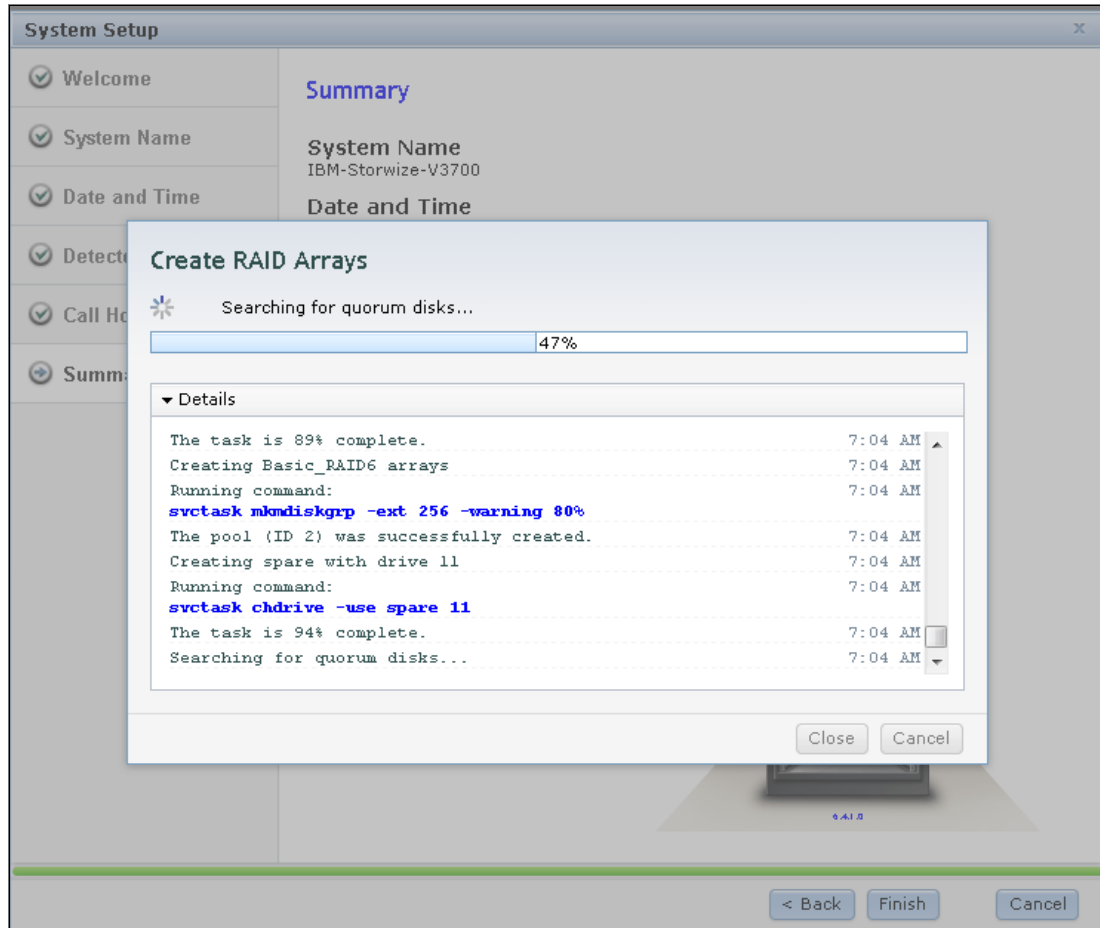


Figure 2-18 Preset RAID configuration

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

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

For a complete description of the management GUI, launch the e-Learning tutorial module by selecting *Watch e-Learning: Overview* once the system has finished configuration Figure 2-19 or refer to Chapter 7, “Storage pools” on page 267.

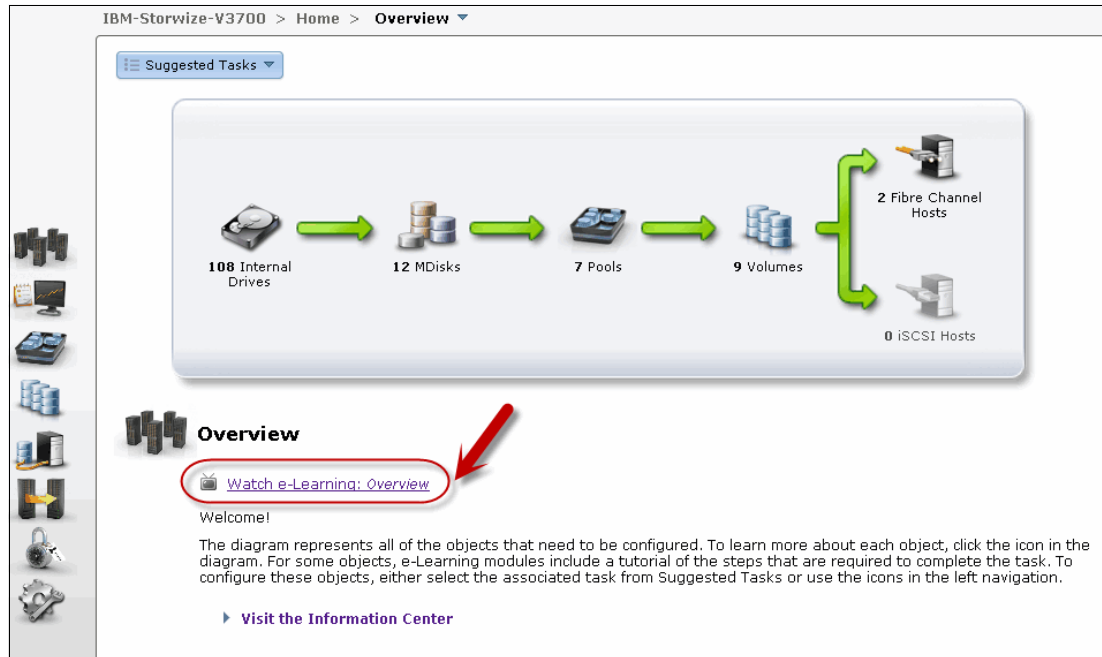


Figure 2-19 Watch e-Learning: Overview

For more details on the presets refer to the infocenter:

http://pic.dhe.ibm.com/infocenter/storwize/v3700_ic/index.jsp?topic=%2Fcom.ibm.storwize.v3700.641.doc%2Fsvc_presets_09271645.html

2.9.1 Adding Enclosures

Install all the expansion enclosures and drives you are adding before configuring the system. This enables the system to make the best use of the enclosures and drives you have added.

If you are installing more than one expansion enclosure, do not configure the storage until all the enclosures and drives have been discovered. This will give you the best RAID configuration. If you add an enclosure after initial configuration this is the procedure

Use the management GUI to configure the new enclosure.

1. In the Monitoring tab, view the system details as shown in Figure 2-20.

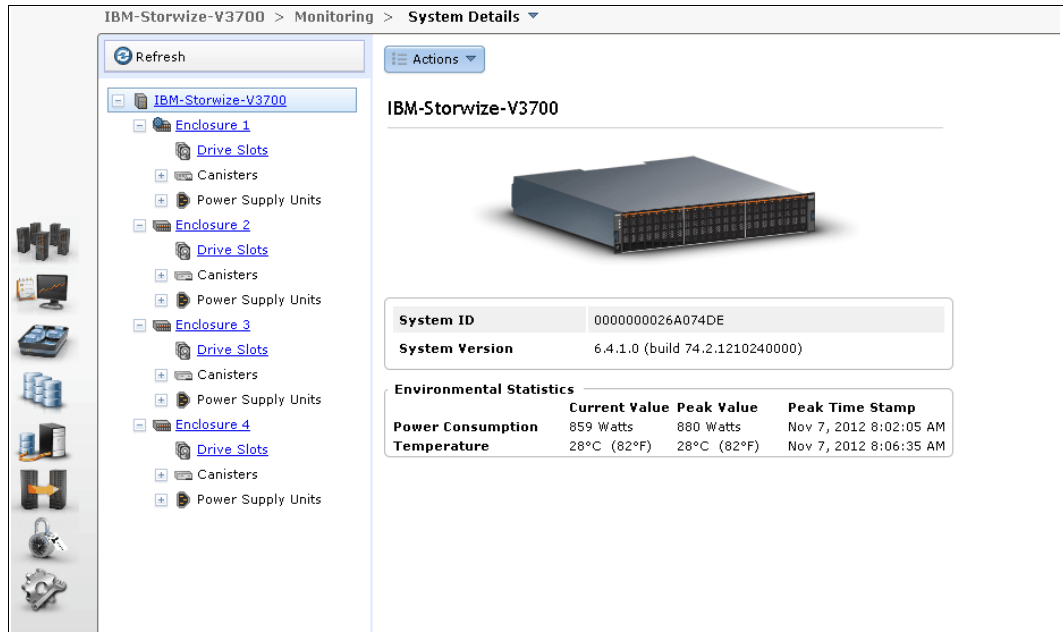


Figure 2-20 System Details

2. Select the option to add a new expansion enclosure from the Actions menu. This wizard guides you through detecting and configuring the new enclosures as shown in Figure 2-21 on page 47.

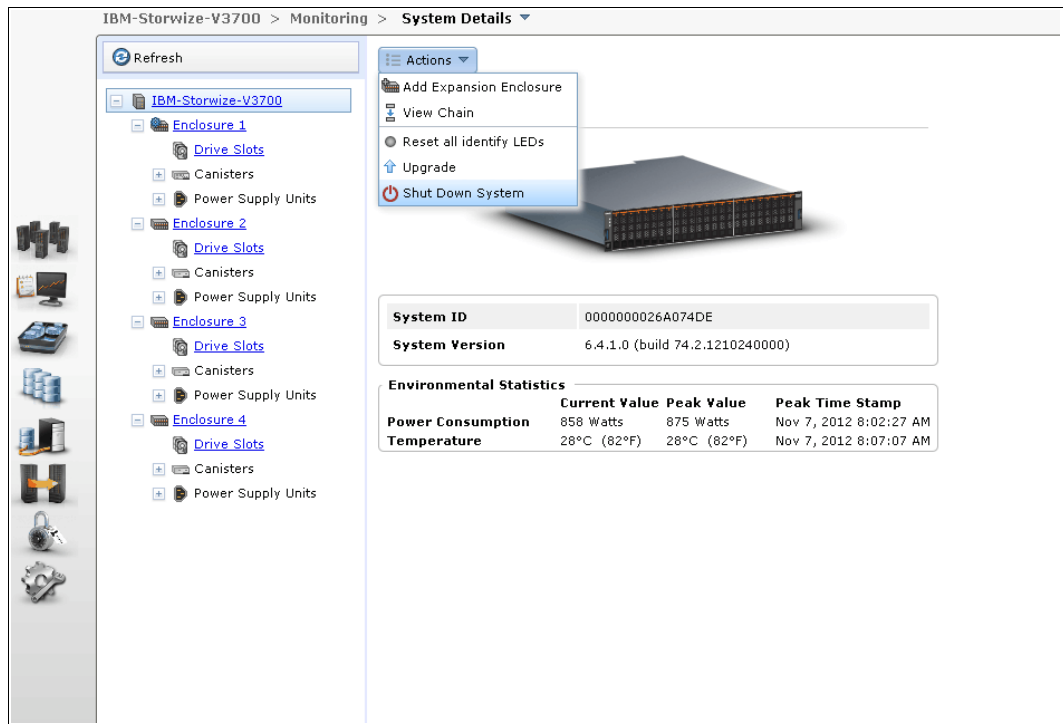


Figure 2-21 Add Expansion Enclosure

3. You will receive a message to verify you have connected cables as shown in Figure 2-22.

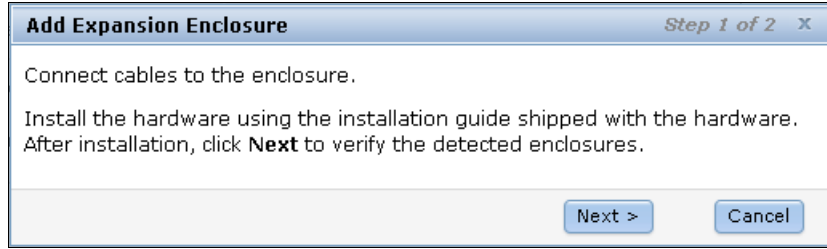


Figure 2-22 Connect Cables

4. Preset discovery will detect the enclosure as shown in Figure 2-23 on page 48.

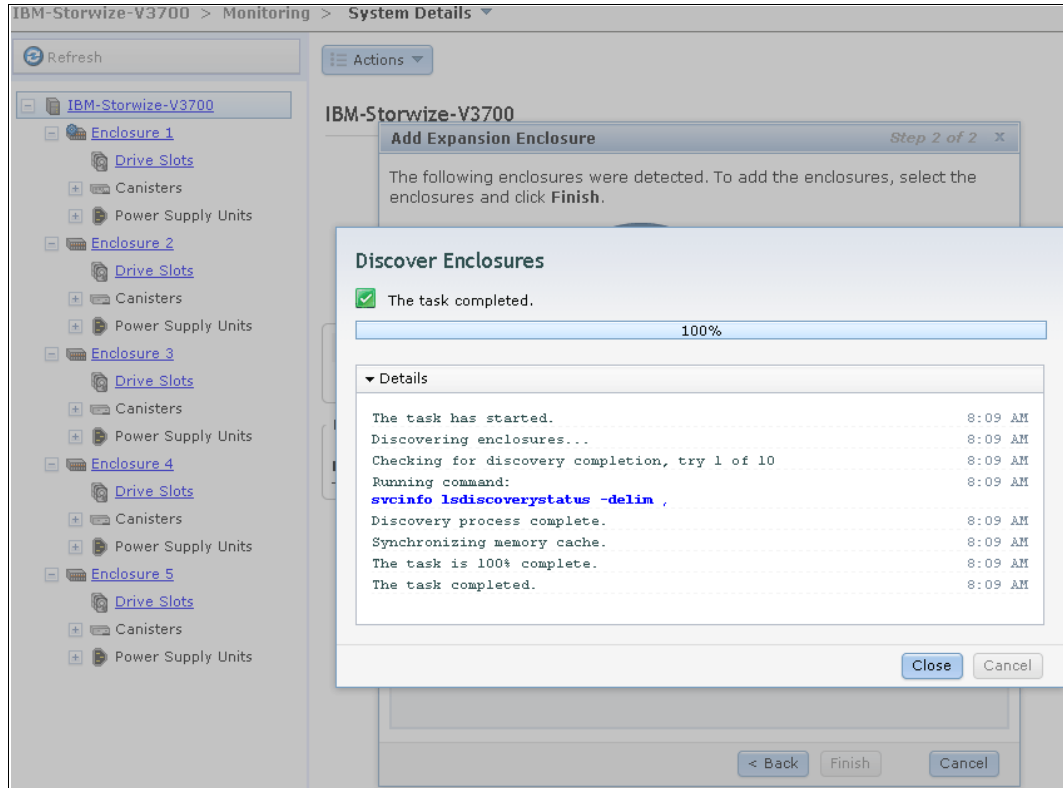


Figure 2-23 Discovery Enclosure

5. To add the enclosure select the new enclosure and choose finish as shown in Figure 2-24 on page 49.




Figure 2-24 Finish enclosure add

- The background "Add Enclosure" to manage new enclosure will display the screen shown in Figure 2-25 on page 50.

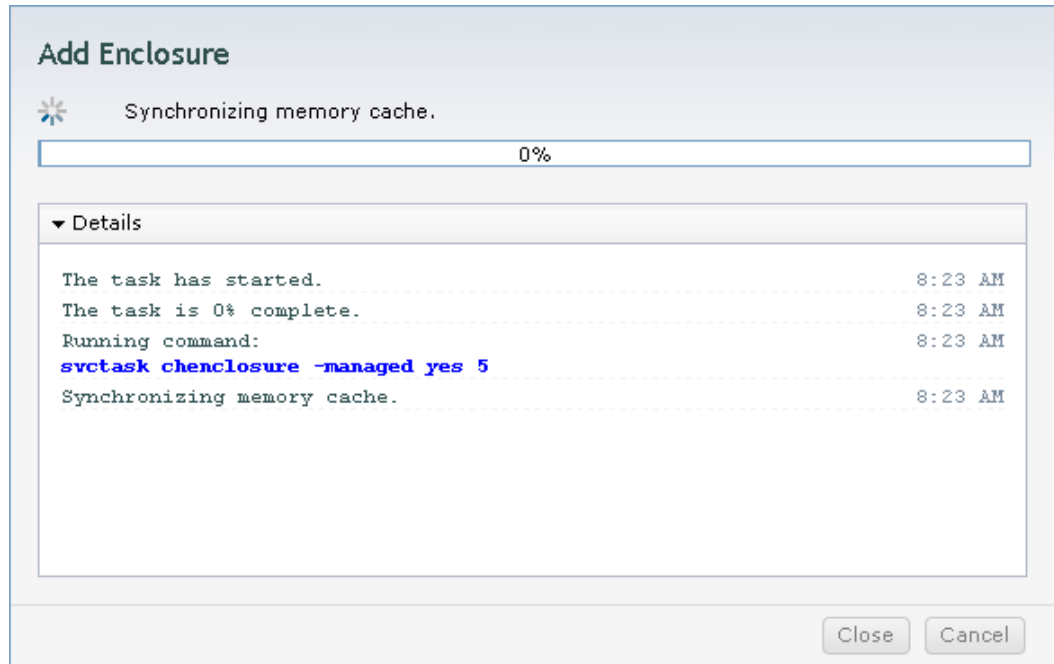


Figure 2-25 Managing new enclosure

- System details will now show the new enclosure as Managed Yes as shown in Figure 2-26.

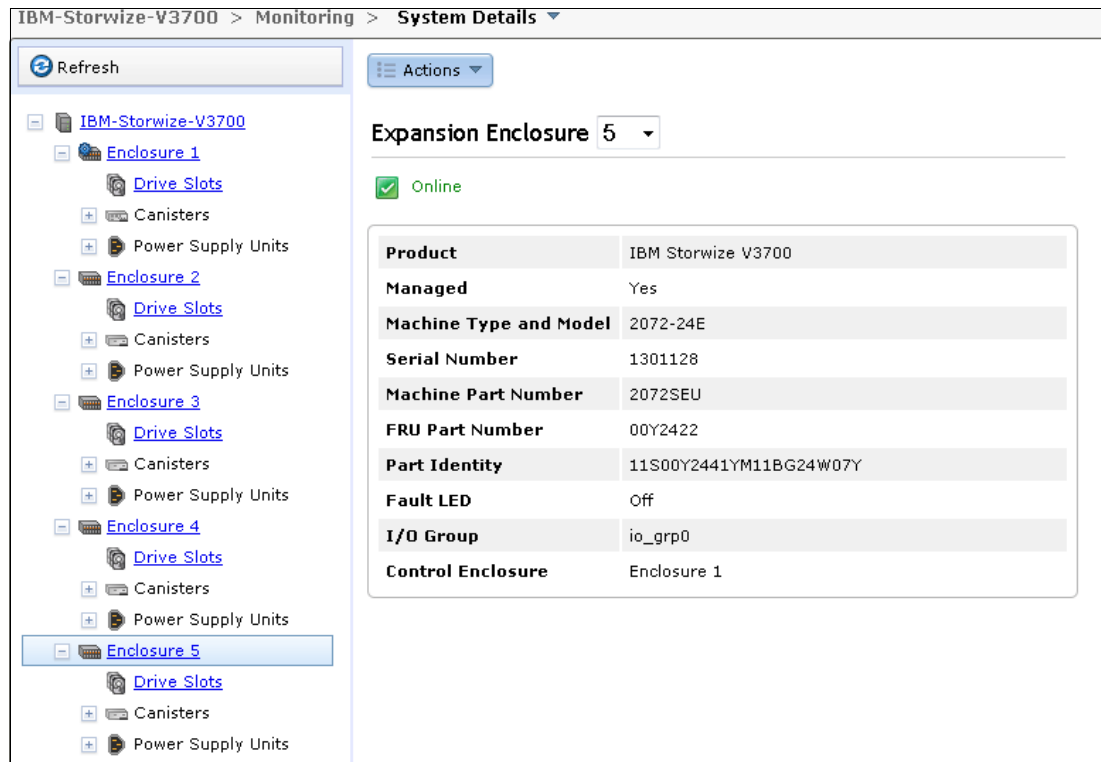


Figure 2-26 New enclosure managed

2.9.2 Configure Call Home, email alert, and inventory

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

Call Home: When configuring Call Home, use one of the following email addresses, depending on country or region of installation:

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

IBM Storwize V3700 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.

To configure the Call Home and email alert event notification in IBM Storwize V3700, complete the following steps.

1. Click **Call Home** as shown in Figure 2-27 on page 51.




Figure 2-27 System setup - Call Home

2. Follow the prompts for callhome system location as shown in Figure 2-28 (all fields are mandatory).

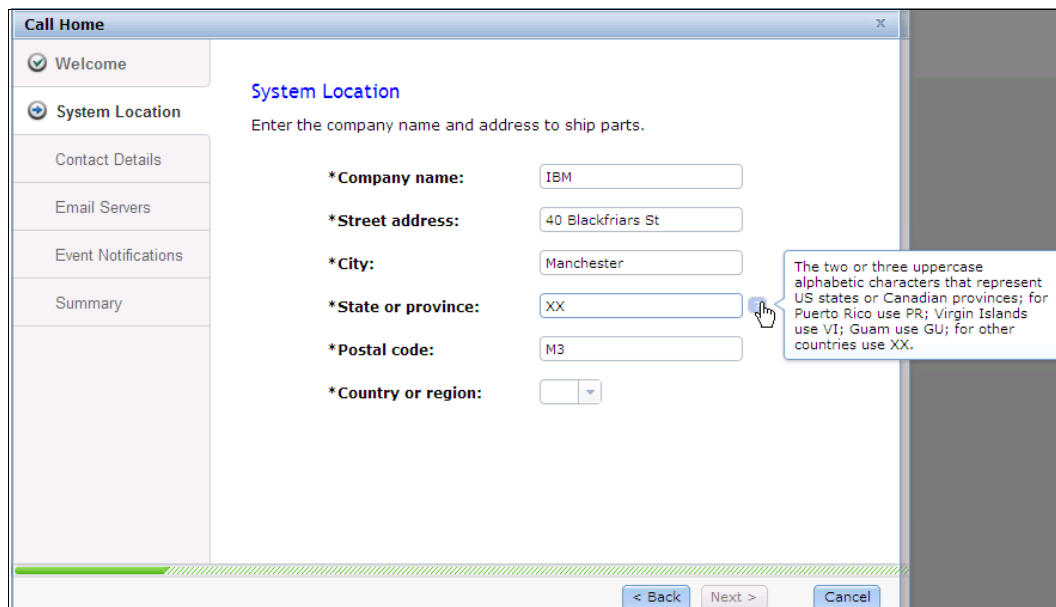


Figure 2-28 Call Home - System location

3. You will need to enter in the contact details as shown in Figure 2-29.

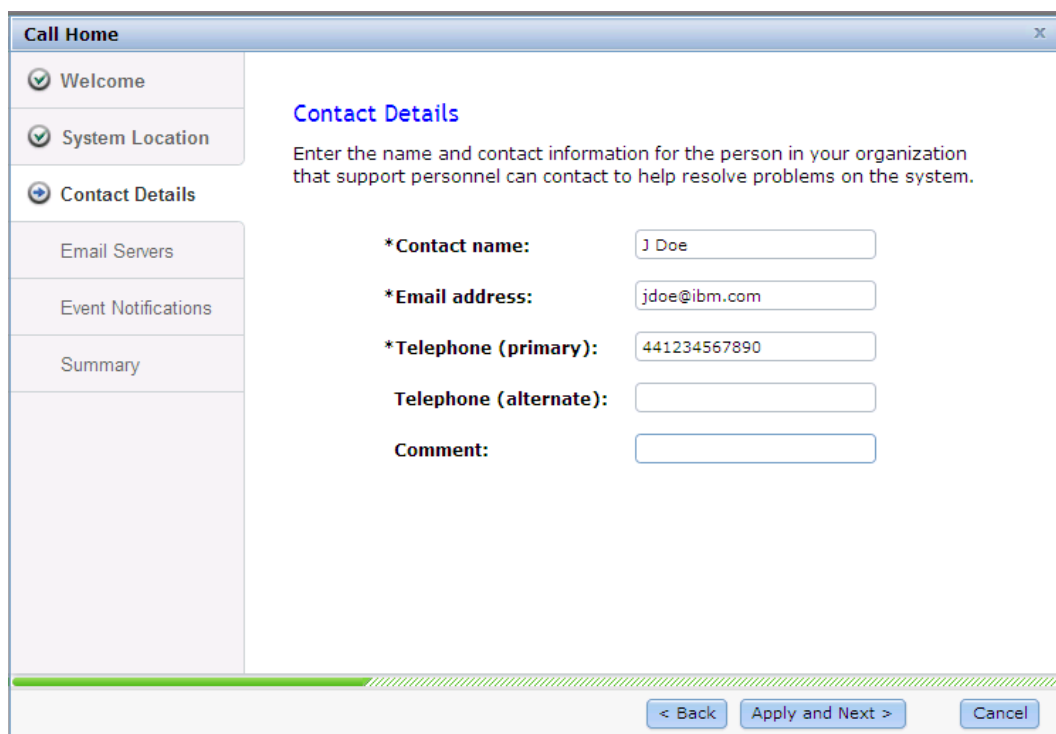


Figure 2-29 Call Home - contact details

4. You will need to configure youemail server IP settings as shown in Figure 2-30 on page 53.

IP Address	Server Port	
1.2.3.4	25	+ -
6.7.8.9	25	+ -

Figure 2-30 Call Home - Email Servers IP address

2.9.3 Configure Event Notifications

Each event that IBM Storwize V3700 detects is assigned a notification type of Error, Warning, or Information. You can configure IBM Storwize V7000 to send each type of notification to specific recipients.

1. You will need to setup to whom you wish to receive these notifications. Figure 2-31 on page 54 shows you that you can configure alerts or all events to be sent to specific users.

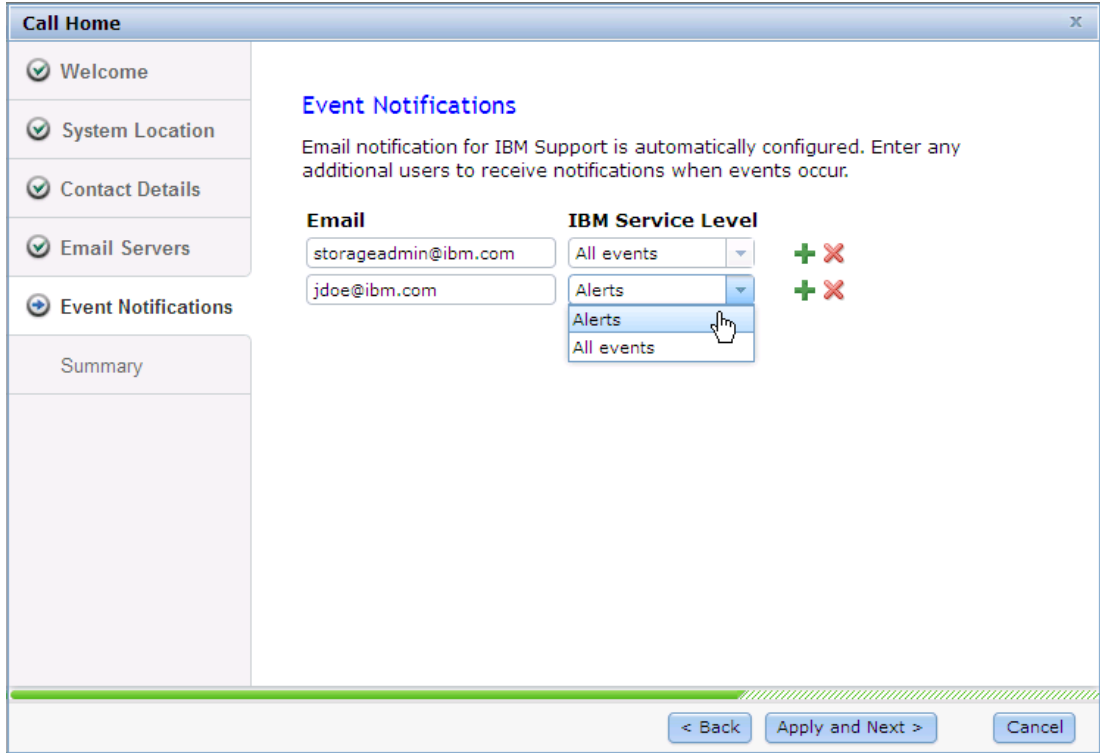


Figure 2-31 Event notifications

2. Once you have specified your event notification, select Apply and Next. This will give you the screen shown in Figure 2-32.

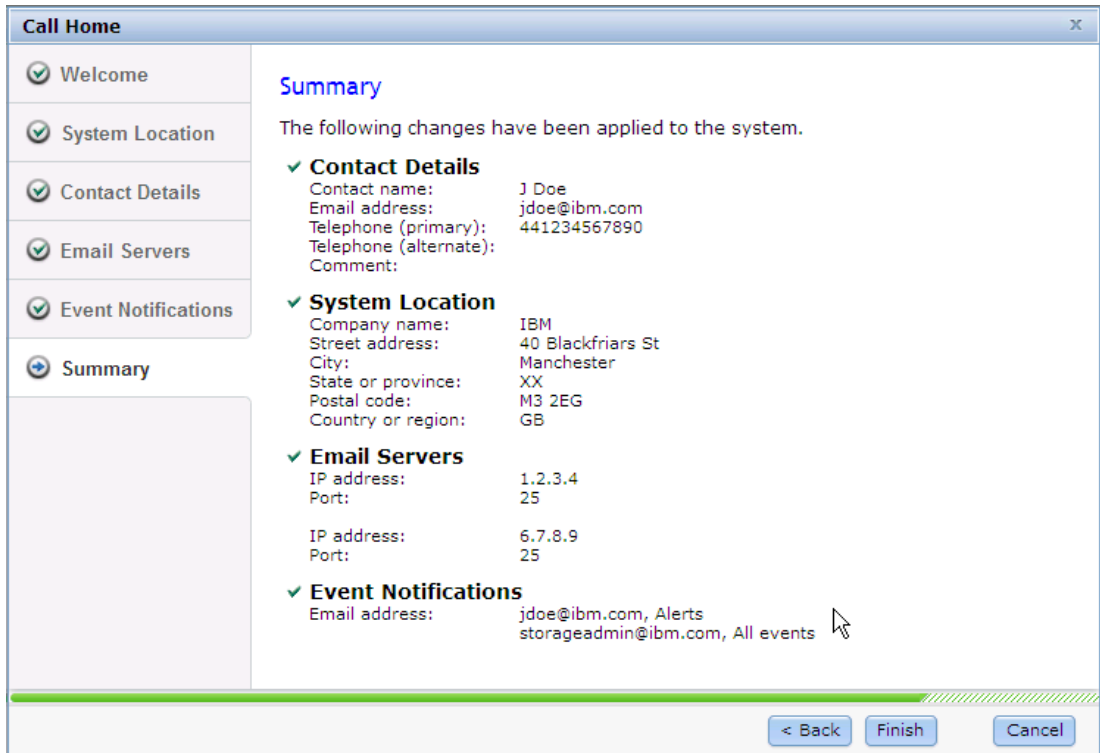


Figure 2-32 Event Notification summary

3. Choosing finish here will setup the email settings for the event notifications and you will see the screen shown in Figure 2-33.

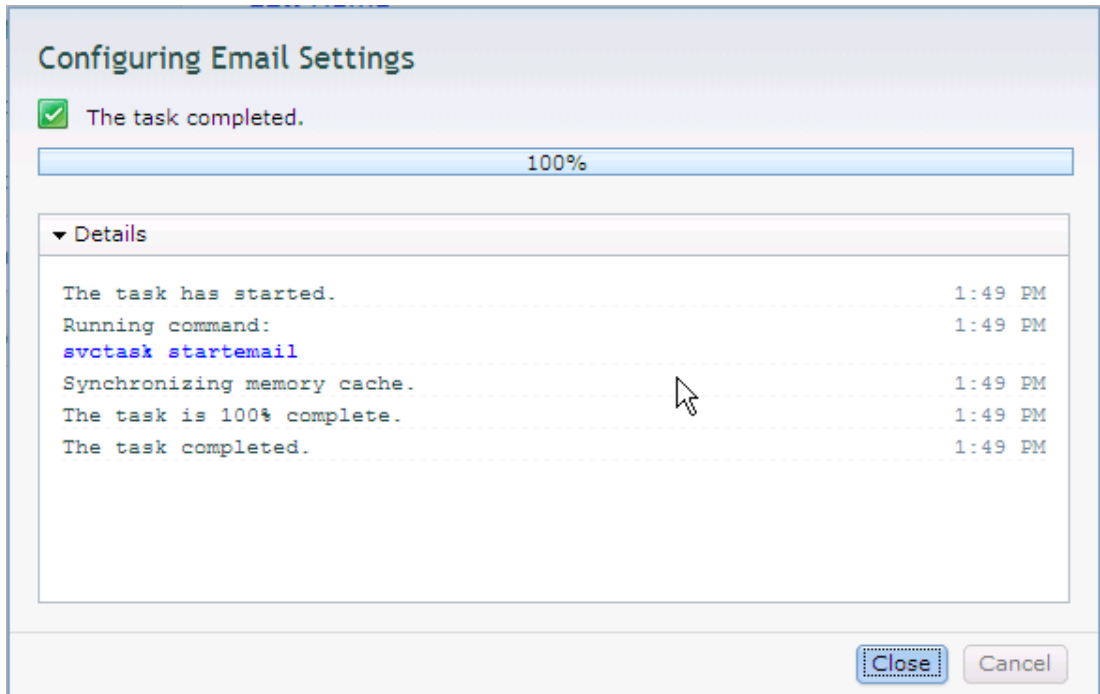


Figure 2-33 Configuring Email Settings

4. Once this is complete you will get the screen shown Figure 2-34.

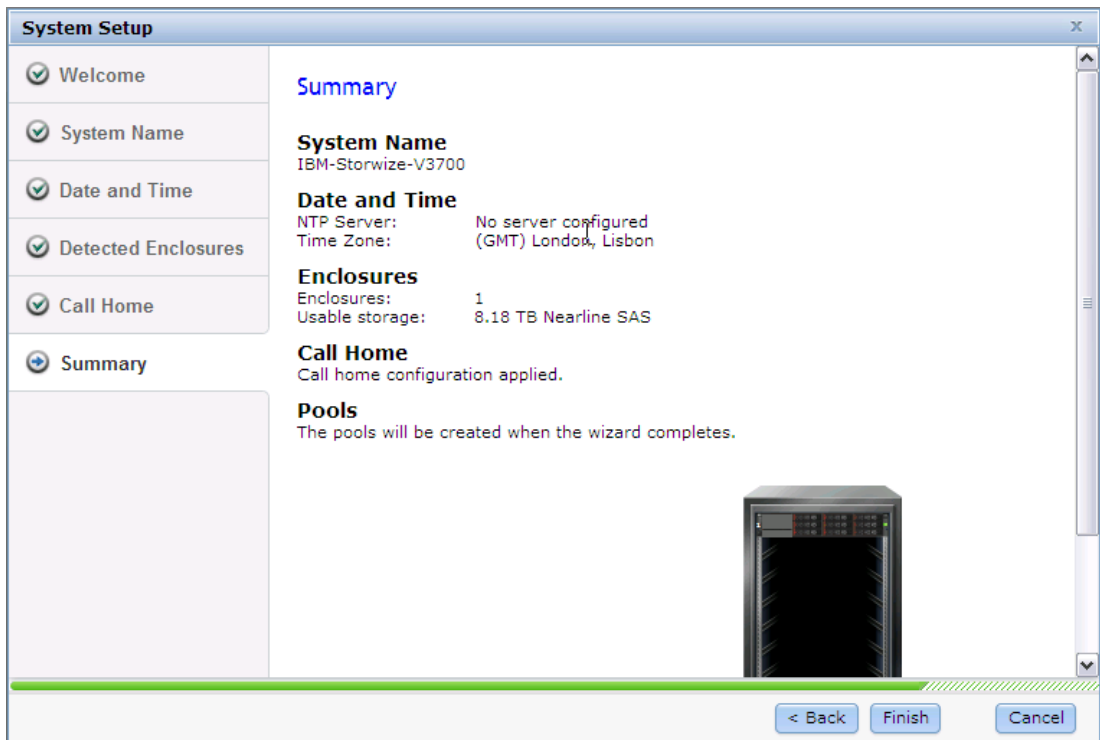


Figure 2-34 Summary

2.9.4 Service Assistant Tool

The V3700 is initially configured with three IP addresses: one service IP address for each node canister and a management IP address, set when the cluster is initialized.

Two methods are available to initially configure the Storwize V3700 system:

- ▶ The Inittool Program 2.8, “First-time setup” on page 37
- ▶ The Service Assistant Tool

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

The SA GUI and CLI are used to service a node even if it is not able to perform I/O. This interface provides full remote access to the information and tasks available in a standard SAN Volume Controller front panel interface. The information and tasks include:

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

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

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

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

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

Example:

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

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

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

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

The IBM Storwize V3700 Login Assistant Login Screen will look as shown in Figure 2-35.



Figure 2-35 Service Assistant Login Screen

Once logged in you will be given the Service Assistant Home Screen. The “Current” canister node is displayed in the top left-hand corner of the GUI. Figure 2-36 shows the Service Assistant Home Screen.

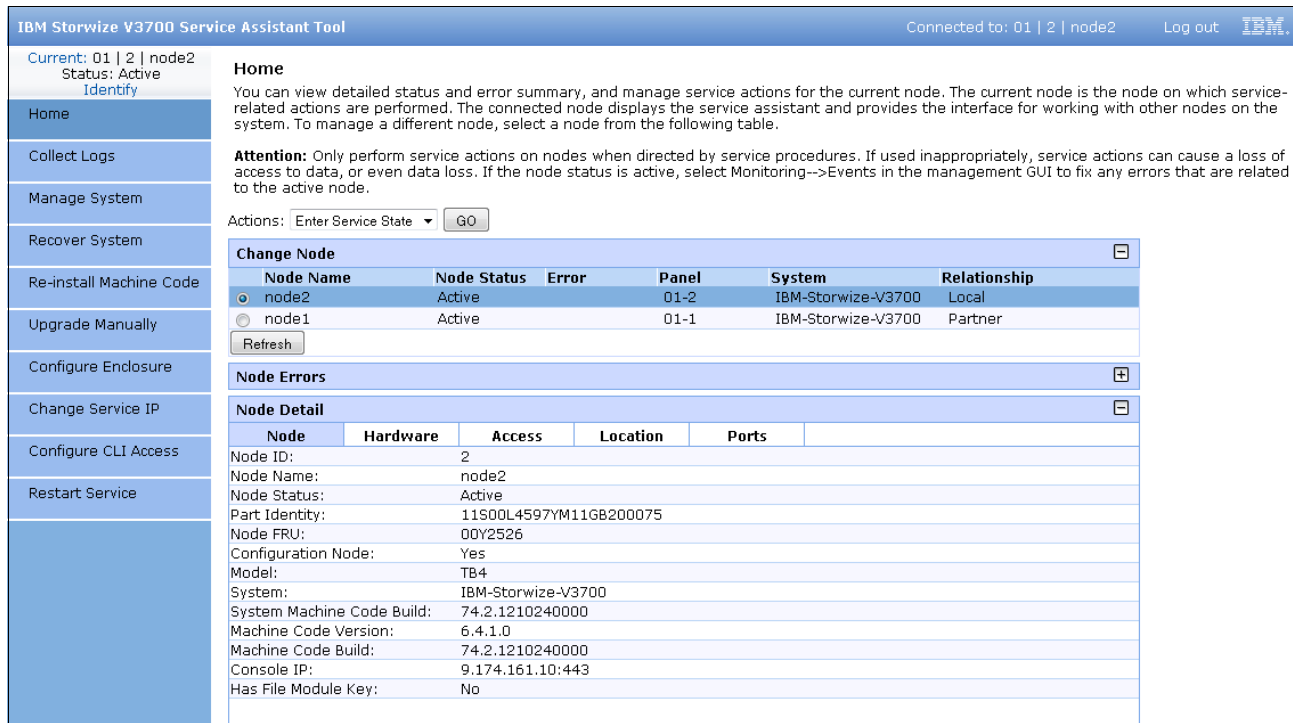


Figure 2-36 Service Assistant Home Screen

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

It is strongly recommended that these IP addresses are configured on all Storwize V3700 node canisters.

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

For more information on how to use the SA Tool please refer to the infocenter at this link:

http://pic.dhe.ibm.com/infocenter/storwize/v3700_ic/index.jsp?topic=%2Fcom.ibm.storwize.v3700.641.doc%2Ftbrd_sagui_1938wd.html



3

Graphical user interface overview

This chapter provides an overview of the Storwize V3700 graphical user interface (GUI) and shows how to navigate the configuration panels.

3.1 Getting started

This section provides information about accessing the Storwize V3700 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 Storwize V3700 management software is a browser-based GUI. It is designed to simplify storage management by providing a single point of control for monitoring, configuration, and management. Table 3-1 lists the web browsers supported by the management GUI. For the latest version of this table see the *IBM Storwize V3700 Quick Installation Guide*, GC27-4219-00.

Table 3-1 Supported web browsers

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

Note: The Web browser requirements, and recommended configuration settings to access the Storwize V3700 management GUI, can be found in the Storwize V3700 Information Center.

http://pic.dhe.ibm.com/infocenter/storwize/v3700_ic/index.jsp

3.1.2 Access the management GUI

To log in to the management GUI, open a supported web browser and enter the management IP address or Hostname of the Storwize V3700. This will display the log in panel. Figure 3-1 shows the Storwize V3700 log in panel.

Default username and password:

username = superuser

password = passw0rd (i.e, the word password with a zero in place of letter o)



Figure 3-1 Storwize V3700 log in panel

A successful log in shows the Overview panel by default. Alternatively, the last opened window from the previous session will be displayed. Figure 3-2 shows the Overview panel.

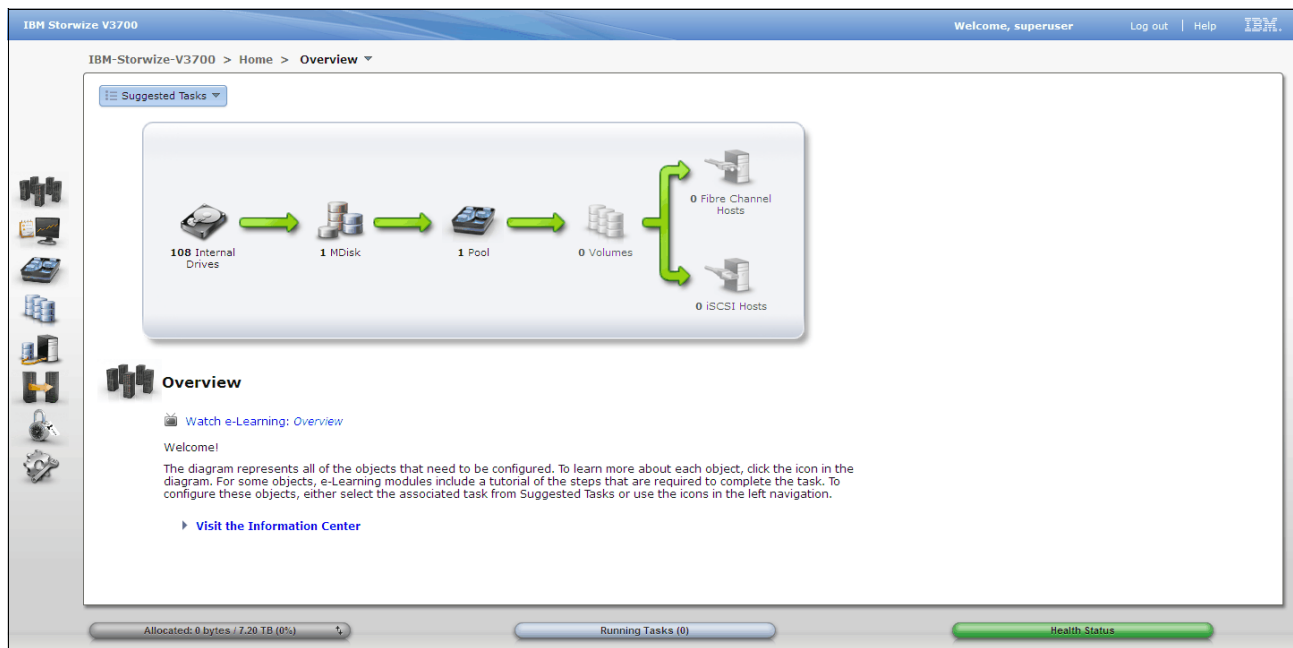


Figure 3-2 Storwize V3700 overview panel

Low graphics mode

Figure 3-1 on page 61 shows the Storwize V3700 log in panel and the check box to enable low graphics mode. This feature can be useful for remote access over narrow band links. The *Function Icons* will no longer enlarge and list available functions, however navigation is achieved by clicking a *Function Icon* and using the breadcrumb navigation aid. The *Function Icons* are described in detail in 3.1.3, “Overview panel layout” on page 62. The breadcrumb navigation aid is described in 3.2.3, “Breadcrumb navigation aid” on page 66. 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

The Overview panel has three main sections: *Function Icons*, *Extended Help*, and *Status Indicators*. Each section can be used to navigate the management GUI. Figure 3-4 shows the three main sections of the Overview panel.




Figure 3-4 The three main sections of the Storwize V3700 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 reveal the menus listed.

- ▶ The Home menu

- ▶ The Monitoring menu
- ▶ The Pools menu
- ▶ The Volumes menu
- ▶ The Hosts menu
- ▶ The Copy Services menu
- ▶ The Access menu
- ▶ The Settings menu

The *Extended Help* section has a flow diagram illustrating the available system resources. The flow diagram is comprised of system resource images and green arrows. The images represent the physical and logical elements of the system. The green arrows in the flow diagram show the order to perform storage allocation tasks and also highlight the various logical layers between the physical internal disks and the logical volumes. Clicking the objects in this area will reveal further information. This information provides *Extended Help* references such as the online version of the Information Center and e-Learning modules. Additionally, this information provides direct links to the various configuration panels that relate to the highlighted image.

The *Status Indicators* section shows three horizontal status bars which are listed.

Allocated indicates status related to the storage capacity of the system.

Running Tasks indicates status of currently running tasks and the recently completed tasks.

Health Status indicates status relating to system health, it is color coded for quick recognition.

Green → **Healthy**, **Yellow** → **Degraded**, **Red** → **Unhealthy**.

Hovering the pointer and clicking on the horizontal bars provides additional information and menus this aspect is discussed in detail in 3.3, “Status Indicators menus” on page 75.

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 either the *Function Icons* section, or the *Extended Help* section of the Overview panel, both sections are described in 3.1.3, “Overview panel layout” on page 62. This section describes the two main navigation methods and introduces the well known breadcrumb navigation aid, and the suggested tasks aid. Finally, information regarding the navigation of panels with tables is provided.

3.2.1 Function icons navigation

Hovering the pointer over one of the eight *Function Icons* on the left of the panel will enlarge the icon and provide a menu to access various functions. Move the pointer to the desired function and click. Figure 3-5 shows the results of hovering the pointer over a *Function Icon*.




Figure 3-5 Hovering the pointer over a function icon

Figure 3-6 shows all the menus under the *Function Icons*.

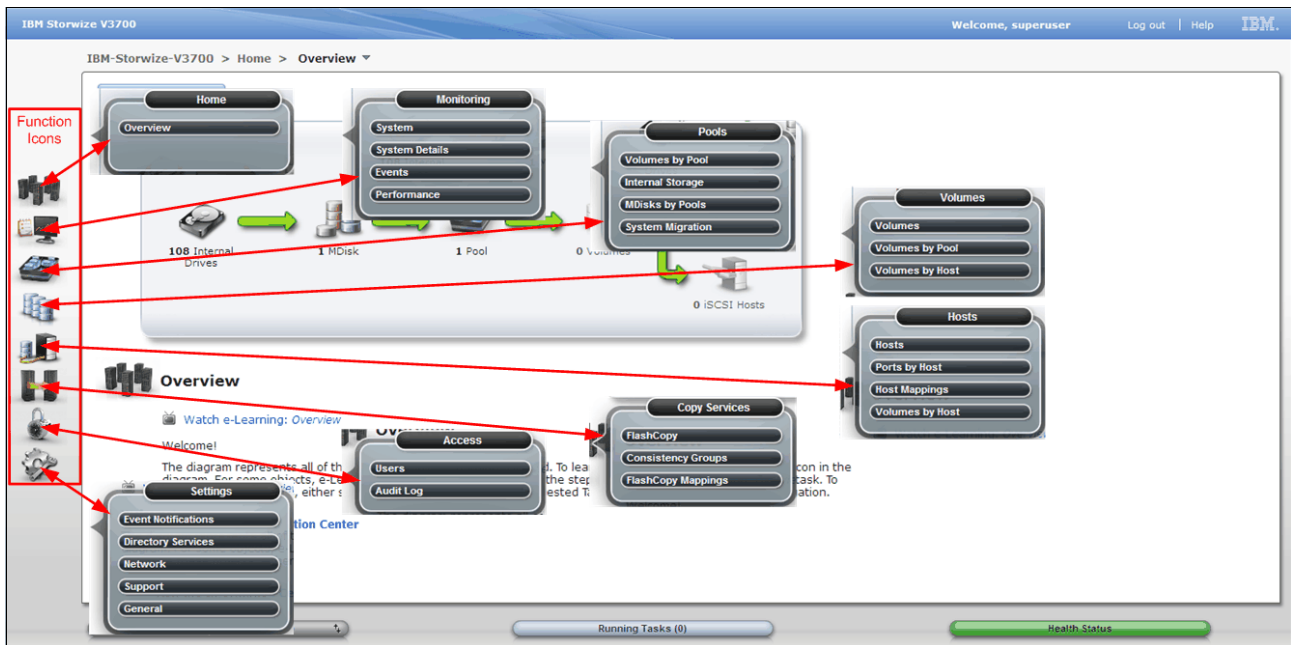


Figure 3-6 Storwize V3700 all options listed under functional icons

3.2.2 Extended help navigation

Selecting an image in the flow diagram of the *Extended Help* section of the Overview panel will reveal information beneath the flow diagram. This information contains links to e-Learning modules and configuration panels related to the selected image. This is convenient when implementing the system, for it is possible to work from left to right, following the flow, and

selecting each object in order. Figure 3-7 on page 65 shows the selection of **Internal Drives** in the flow diagram. The information below the flow diagram now relates to the internal storage.




Figure 3-7 Navigating GUI with the extended help section

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



Figure 3-8 Using the extended help section

Figure 3-9 shows the Internal Storage panel, this has been displayed as a result of selecting and clicking the **Pools** in the information area of the *Extended Help* section.




Figure 3-9 The internal storage configuration panel

3.2.3 Breadcrumb navigation aid

The Storwize V3700 panels uses the breadcrumb navigation aid to show the trail that has been navigated. This breadcrumb navigation aid is located at the top area of the panel and has a drop down selection menu on the last breadcrumb. Figure 3-10 shows the breadcrumb navigation aid for the System panel.




Figure 3-10 Breadcrumb navigation aid

3.2.4 Suggested tasks aid

The suggested tasks feature is a navigation and configuration aid, it is located in the top area of the Overview panel. The list of suggested tasks will change 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.

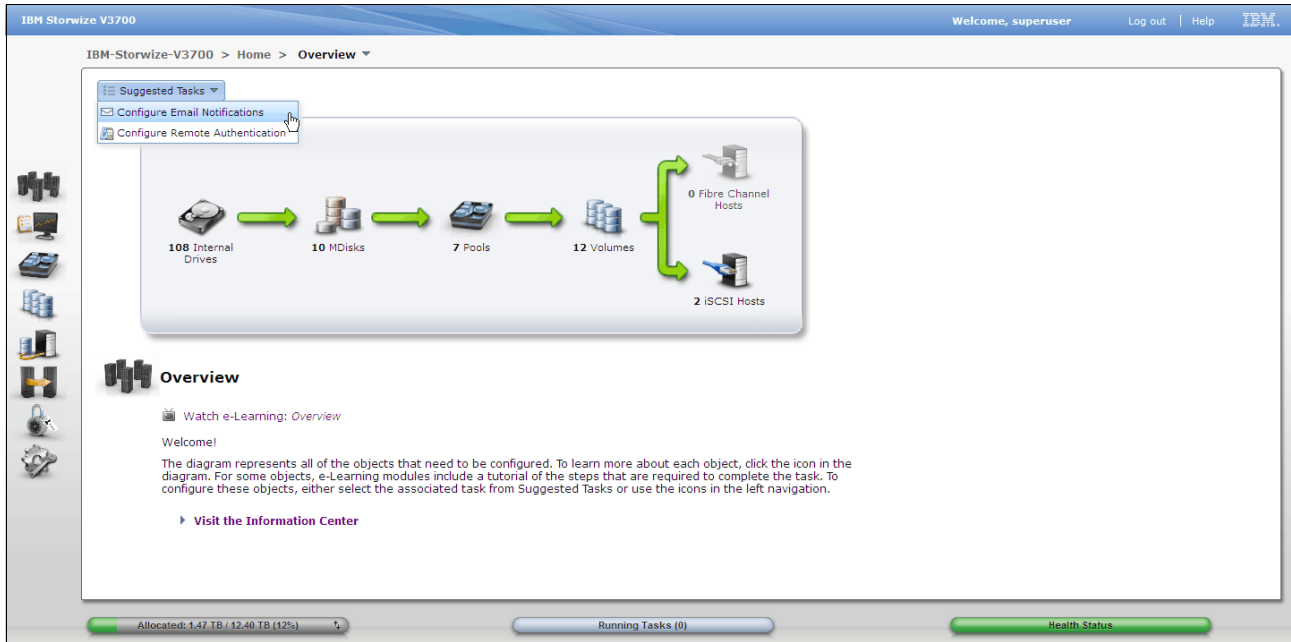


Figure 3-11 Suggested tasks navigation and configuration aid

3.2.5 Presets

The management GUI contains a series of preestablished configuration options 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 on page 68 shows the available internal storage presets.

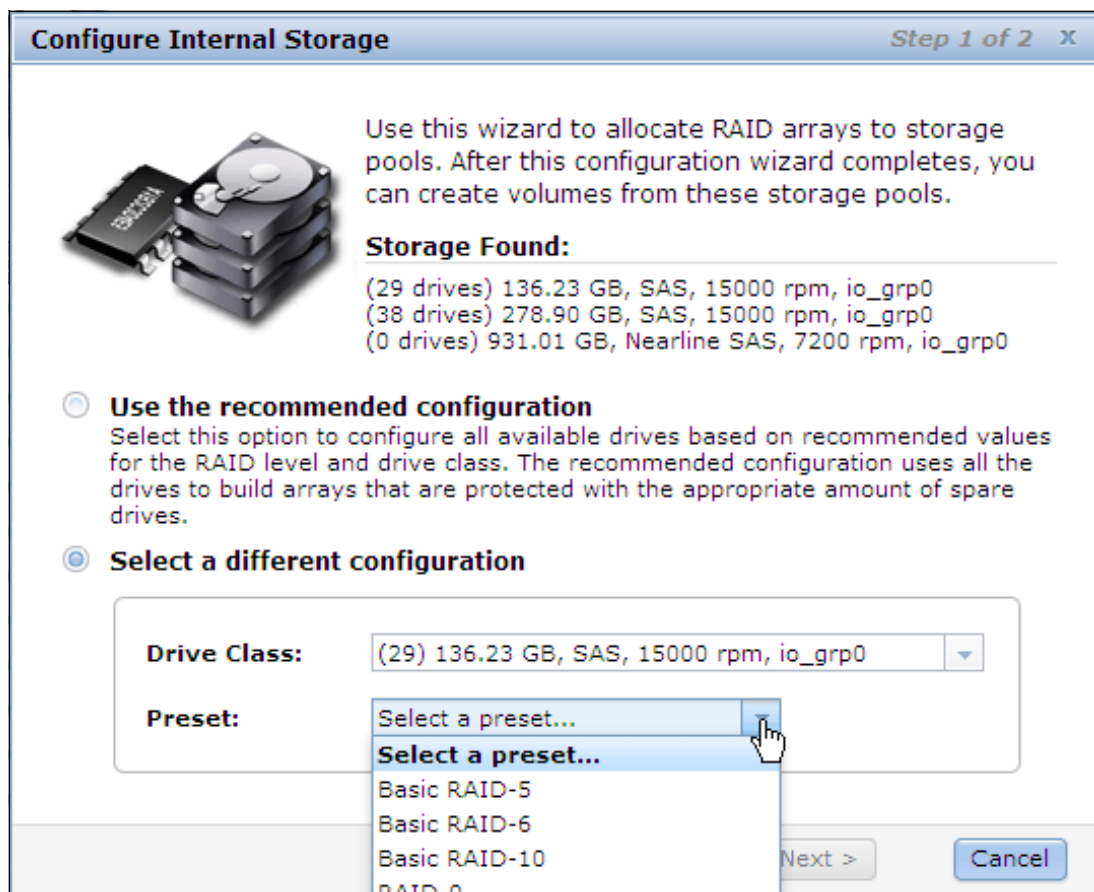


Figure 3-12 Internal storage preset selection

3.2.6 Access actions

The Storwize V3700 functional panels provide access to a range of actions that can be performed, such as: modify attributes, and rename, add, or delete objects. The available actions menus can be accessed by two main methods. Either highlight the resource and use the actions drop box as shown in Figure 3-13 on page 69, or right click on the resources as shown in Figure 3-14 on page 69.

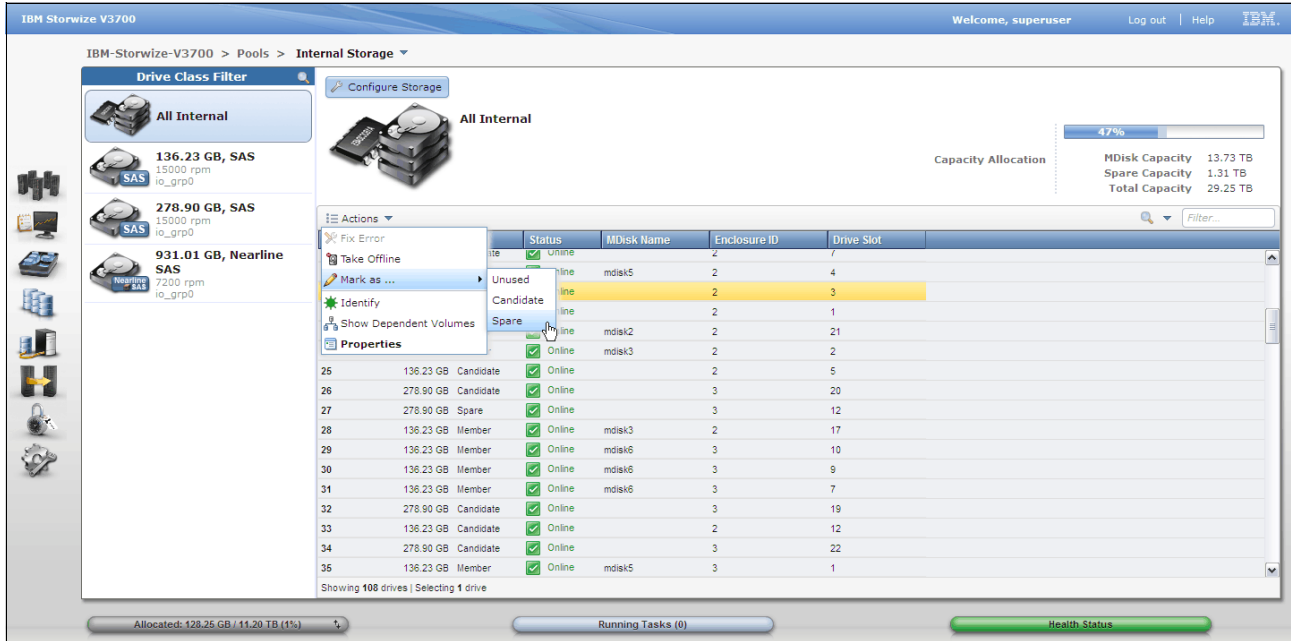


Figure 3-13 Actions drop down menu

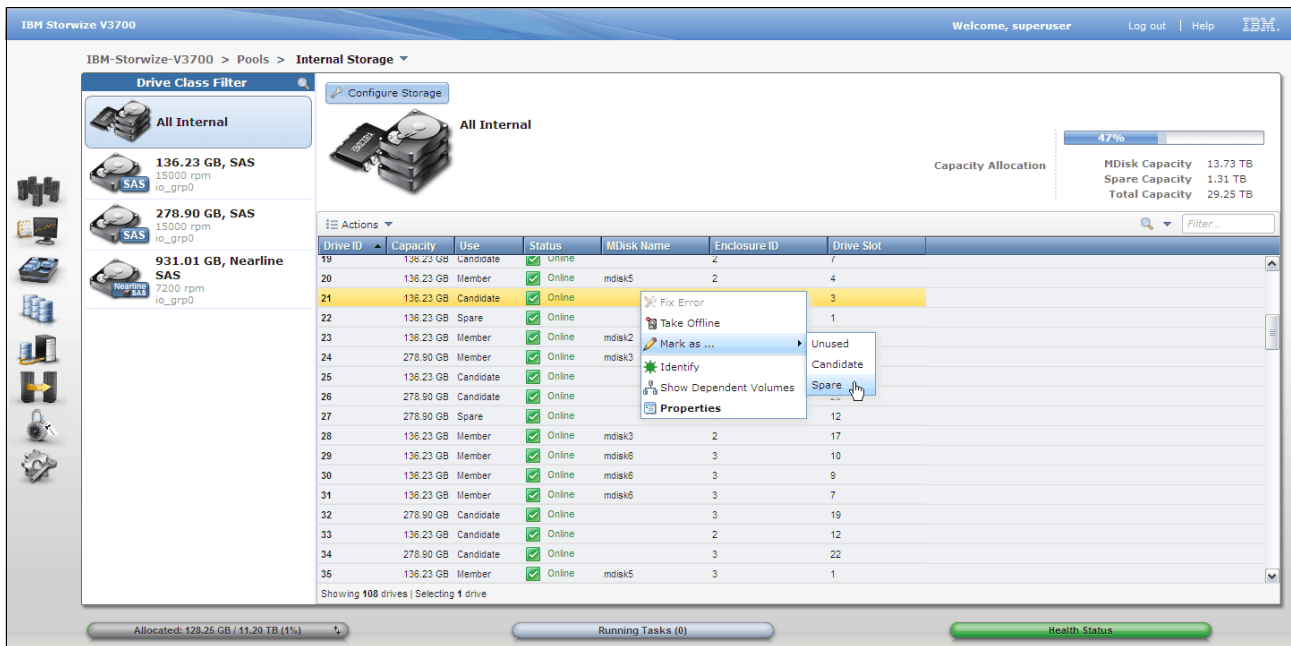


Figure 3-14 Right click actions menu

3.2.7 Task progress

An action will initiate a running task and reveal a task progress box Figure 3-15 on page 70. Click **Details** to reveal the underlying Command Line Interface (CLI) commands. The commands are highlighted in blue and can be cut and then pasted into a configured Storwize V3700 SSH terminal session if required. Figure 3-15 on page 70 shows a task progress box.

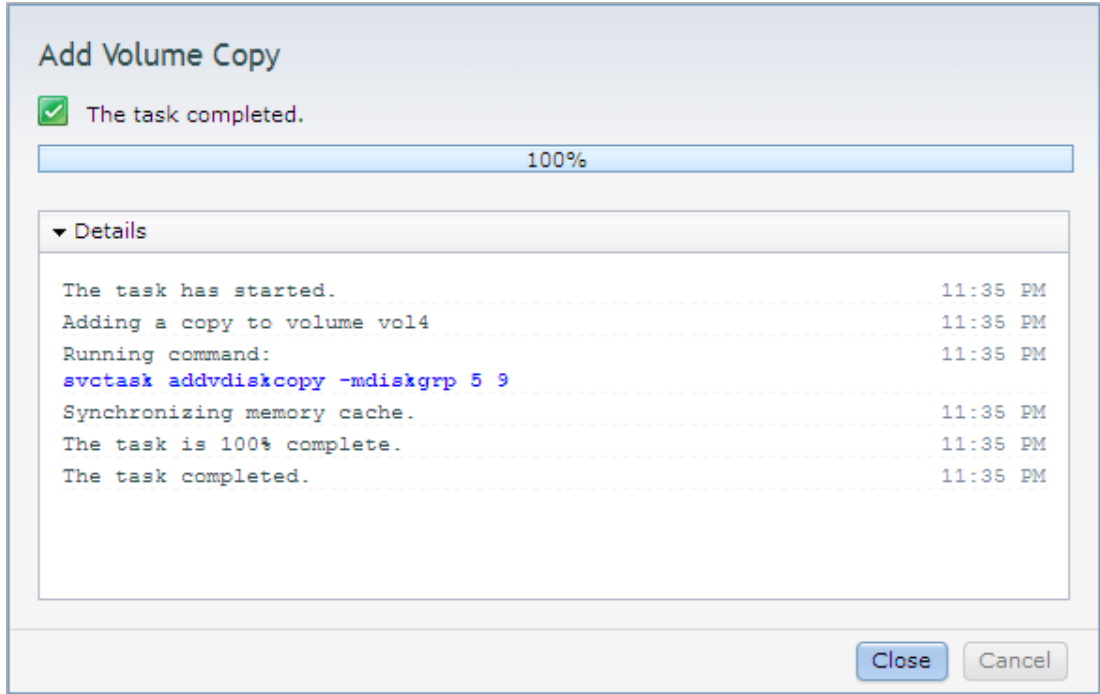


Figure 3-15 Task progress box

3.2.8 Navigating panels with tables

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

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

Sorting columns

Columns can be sorted by clicking on the column heading. Figure 3-16 shows the result of clicking on the heading of the "Capacity Utilization" column. The table is now sorted and has listed volumes with the least amount of capacity utilization at the top of the table.

The screenshot shows the IBM Storwize V3700 GUI interface. At the top, the breadcrumb navigation reads "IBM-Storwize-V3700 > Volumes > Volumes". Below this is a table with the following columns: Name, Status, Capacity Utilization, Capacity, Storage Pool, UID, and Host Mappings. The "Capacity Utilization" column is currently selected, and the table is sorted by this column in ascending order. The first three rows show "Thin Volume 1", "Thin Volume 2", and "Thin Volume 3", all with 0% utilization. The next three rows show "Volume 5", "Copy 1", and "Copy 0*", all with 100% utilization. The final three rows show "Volume 1", "Volume 2", and "Volume 3", all with 100% utilization. A mouse cursor is hovering over the "Capacity Utilization" column header.

Name	Status	Capacity Utilization	Capacity	Storage Pool	UID	Host Mappings
Thin Volume 1	Online	0%	4.00 TB	mdiskgrp0	60050760009901D37800000000000008	No
Thin Volume 2	Online	0%	4.00 TB	mdiskgrp0	60050760009901D37800000000000009	No
Thin Volume 3	Online	0%	4.00 TB	mdiskgrp0	60050760009901D3780000000000000A	No
Volume 5	Online	100%	200.00 GB	mdiskgrp0	60050760009901D37800000000000007	No
Copy 1	Online	100%	200.00 GB	mdiskgrp0	60050760009901D37800000000000007	No
Copy 0*	Online	100%	200.00 GB	mdiskgrp0	60050760009901D37800000000000007	No
Volume 1	Online	100%	200.00 GB	mdiskgrp0	60050760009901D37800000000000003	No
Volume 2	Online	100%	200.00 GB	mdiskgrp0	60050760009901D37800000000000004	No
Volume 3	Online	100%	200.00 GB	mdiskgrp0	60050760009901D37800000000000005	No
Volume 4	Online	100%	200.00 GB	mdiskgrp0	60050760009901D37800000000000006	No

Figure 3-16 Sorting columns by clicking the column heading

Reordering Columns

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

Name	Status	Capacity Utilization	Capacity	Storage Pool	UID	Host Mappings
Volume 1	Online	100%	200.0 GB	mdiskgrp0	60050760009901D37800000000000003	No
Volume 2	Online	100%	200.0 GB	mdiskgrp0	60050760009901D37800000000000004	No
Volume 3	Online	100%	200.0 GB	mdiskgrp0	60050760009901D37800000000000005	No
Volume 4	Online	100%	200.0 GB	mdiskgrp0	60050760009901D37800000000000006	No
Volume 5	Online		200.0 GB	mdiskgrp0	60050760009901D37800000000000007	No
Copy 0*	Online	100%	200.0 GB	mdiskgrp0	60050760009901D37800000000000007	No
Copy 1	Online	100%	200.0 GB	mdiskgrp0	60050760009901D37800000000000007	No
Thin Volume 1	Online		4.00 TB	mdiskgrp0	60050760009901D37800000000000008	No
Thin Volume 2	Online		4.00 TB	mdiskgrp0	60050760009901D37800000000000009	No
Thin Volume 3	Online		4.00 TB	mdiskgrp0	60050760009901D3780000000000000A	No
Vmware 1 Thin	Online		2.00 TB	mdiskgrp0	60050760009901D3780000000000000B	No
Vmware 2 Thin	Online		2.00 TB	mdiskgrp0	60050760009901D3780000000000000C	No
Vmware 3 Thin	Online		2.00 TB	mdiskgrp0	60050760009901D3780000000000000D	No
Vmware 4 Thin	Online		2.00 TB	mdiskgrp0	60050760009901D3780000000000000E	No

Figure 3-17 Reordering columns part 1

Figure 3-18 on page 71 shows the column heading “UID” has been dragged to the desired location.

Name	Status	UID	Capacity Utilization	Capacity	Storage Pool	Host Mappings
Volume 1	Online	60050760009901D37800000000000003	100%	200.0 GB	mdiskgrp0	No
Volume 2	Online	60050760009901D37800000000000004	100%	200.0 GB	mdiskgrp0	No
Volume 3	Online	60050760009901D37800000000000005	100%	200.0 GB	mdiskgrp0	No
Volume 4	Online	60050760009901D37800000000000006	100%	200.0 GB	mdiskgrp0	No
Volume 5	Online	60050760009901D37800000000000007		200.0 GB	mdiskgrp0	No
Copy 0*	Online	60050760009901D37800000000000007	100%	200.0 GB	mdiskgrp0	No
Copy 1	Online	60050760009901D37800000000000007	100%	200.0 GB	mdiskgrp0	No
Thin Volume 1	Online	60050760009901D37800000000000008		4.00 TB	mdiskgrp0	No
Thin Volume 2	Online	60050760009901D37800000000000009		4.00 TB	mdiskgrp0	No
Thin Volume 3	Online	60050760009901D3780000000000000A		4.00 TB	mdiskgrp0	No
Vmware 1 Thin	Online	60050760009901D3780000000000000B		2.00 TB	mdiskgrp0	No
Vmware 2 Thin	Online	60050760009901D3780000000000000C		2.00 TB	mdiskgrp0	No
Vmware 3 Thin	Online	60050760009901D3780000000000000D		2.00 TB	mdiskgrp0	No
Vmware 4 Thin	Online	60050760009901D3780000000000000E		2.00 TB	mdiskgrp0	No
MySQL	Online	60050760009901D3780000000000000F		20.00 GB	mdiskgrp0	No
Copy 0*	Online	60050760009901D3780000000000000F	100%	20.00 GB	mdiskgrp0	No
Copy 1	Online	60050760009901D3780000000000000F	100%	20.00 GB	mdiskgrp0	No

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

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

Name	Status	UID	Capacity Utilization	Capacity	Storage Pool	Host Mappings
Volume 1	Online	60050760009901D37800000000000003	100%	200.0 GB	mdiskgrp0	No
Volume 2	Online	60050760009901D37800000000000004	100%	200.0 GB	mdiskgrp0	No
Volume 3	Online	60050760009901D37800000000000005	100%	200.0 GB	mdiskgrp0	No
Volume 4	Online	60050760009901D37800000000000006	100%	200.0 GB	mdiskgrp0	No
Volume 5	Online	60050760009901D37800000000000007		200.0 GB	mdiskgrp0	No
Copy 0*	Online	60050760009901D37800000000000007	100%	200.0 GB	mdiskgrp0	No
Copy 1	Online	60050760009901D37800000000000007	100%	200.0 GB	mdiskgrp0	No
Thin Volume 1	Online	60050760009901D37800000000000008		4.00 TB	mdiskgrp0	No
Thin Volume 2	Online	60050760009901D37800000000000009		4.00 TB	mdiskgrp0	No
Thin Volume 3	Online	60050760009901D3780000000000000A		4.00 TB	mdiskgrp0	No
Vmware 1 Thin	Online	60050760009901D3780000000000000B		2.00 TB	mdiskgrp0	No
Vmware 2 Thin	Online	60050760009901D3780000000000000C		2.00 TB	mdiskgrp0	No
Vmware 3 Thin	Online	60050760009901D3780000000000000D		2.00 TB	mdiskgrp0	No
Vmware 4 Thin	Online	60050760009901D3780000000000000E		2.00 TB	mdiskgrp0	No
MySQL	Online	60050760009901D3780000000000000F		20.00 GB	mdiskgrp0	No
Copy 0*	Online	60050760009901D3780000000000000F	100%	20.00 GB	mdiskgrp0	No
Copy 1	Online	60050760009901D3780000000000000F	100%	20.00 GB	mdiskgrp0	No

Figure 3-19 Reordering column headings

Adding or removing columns

To add or remove a column, right click on the heading bar and select the desired column headings by checking the box adjacent to the heading name. Figure 3-20 on page 72 shows the addition of the column heading “Mdisk ID”.

Note: Some users may run into a problem, a context menu from the firefox browser is revealed by right clicking to make changes to the column heading. This can be fixed by:

Firefox: **Tools** → **Options** → **Content** → **Advanced (for java setting)** → **Select: Display or replace context menus**

Note: The Web browser requirements, and recommended configuration settings to access the Storwize V3700 management GUI, can be found in the Storwize V3700 Information Center.

http://pic.dhe.ibm.com/infocenter/storwize/v3700_ic/index.jsp

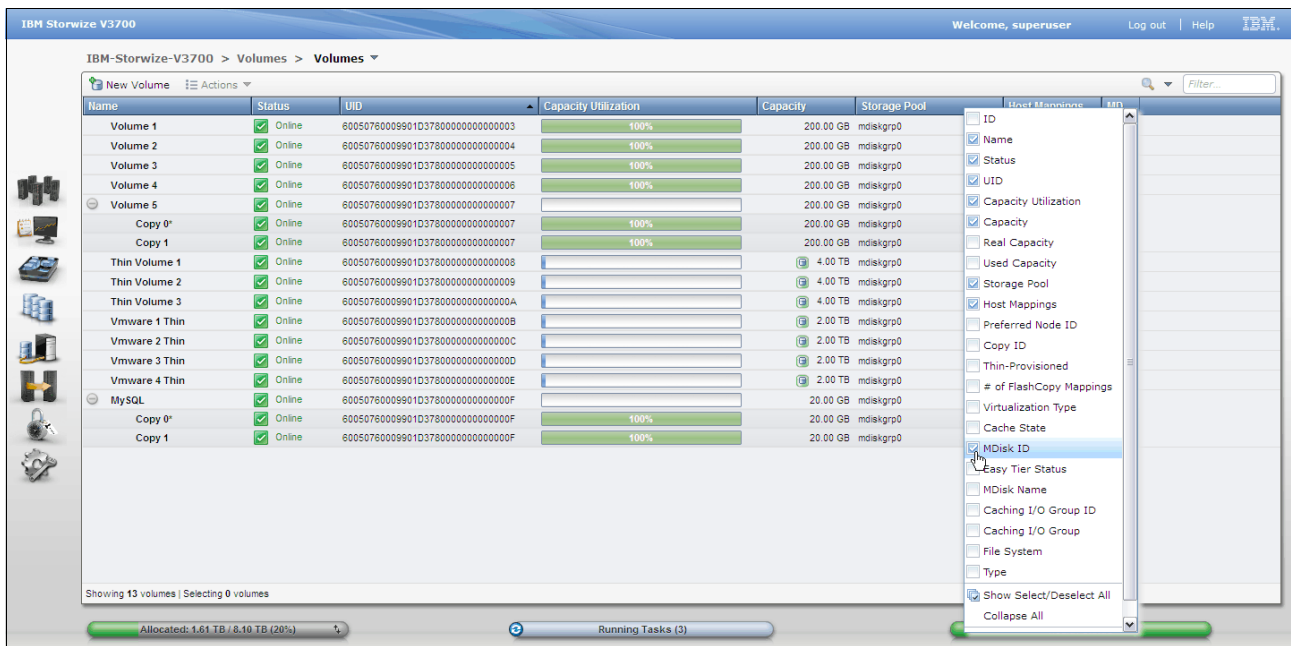


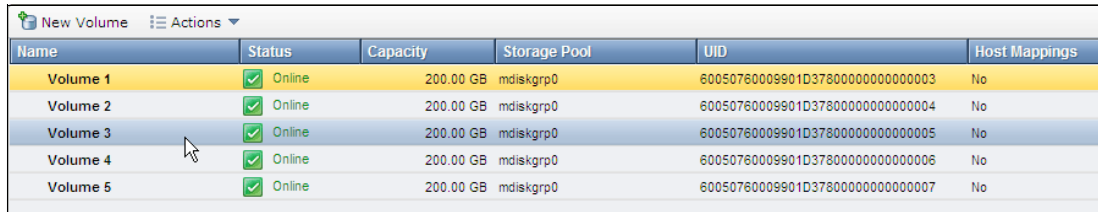
Figure 3-20 Adding column heading “Mdisk ID”

Multiple selections

The management tool also allows selection of multiple items in a list by using a combination of the Shift or Ctrl keys.

Using the Shift key

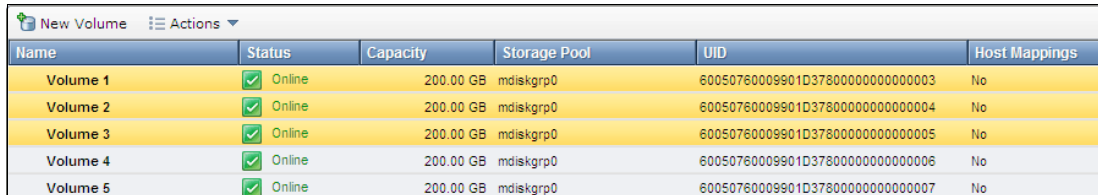
To select multiple items in a sequential order, click the first item listed, press and hold the Shift key, and click the last item in the list. All the items in between those two items are then selected. Figure 3-21 shows the selection of three sequential items.



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 3-21 Selection of three sequential items part 1

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

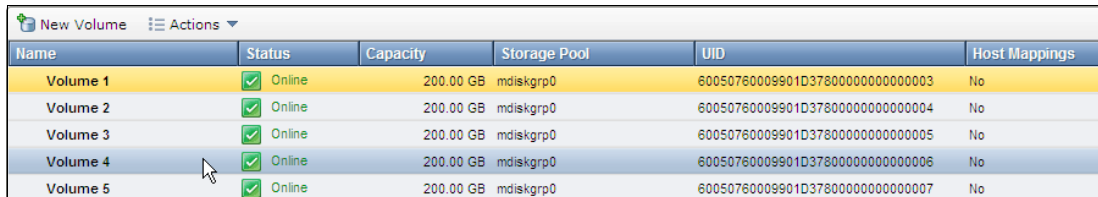


Name	Status	Capacity	Storage Pool	UID	Host Mappings
Volume 1	✓ Online	200.00 GB	mdiskgrp0	60050760009901D37800000000000003	No
Volume 2	✓ Online	200.00 GB	mdiskgrp0	60050760009901D37800000000000004	No
Volume 3	✓ Online	200.00 GB	mdiskgrp0	60050760009901D37800000000000005	No
Volume 4	✓ Online	200.00 GB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

Figure 3-22 Selection of three sequential items part 2

Using the Ctrl key

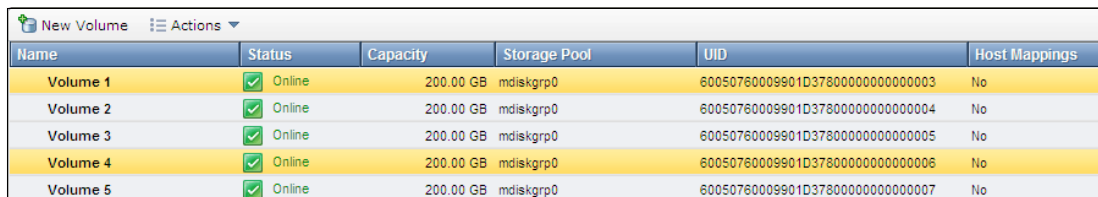
To select multiple items that are not in sequential order, click the first item, press and hold the Ctrl key, and click the other items you require. Figure 3-23 shows the selection of two non-sequential items.



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 3-23 Selection of two non-sequential items part 1

Figure 3-24 shows the result of using 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 3-24 Selection of two non-sequential items part 2

Filtering objects

To focus on a subset of the listed items shown in a panel with columns, use the filter field found at the top right hand side of the table. This tool will display items matching the value entered. Figure 3-25 on page 74 shows the text “Thin” has been entered into the filter field, now only volumes with the text “Thin” in any column are listed and the filter word is also highlighted.

Name	Status	Capacity Utilization	Capacity	Storage Pool	UID	Host Mappings
Thin Volume 1	Online		4.00 TB	mdiskgrp0	60050760009901D37800000000000008	No
Thin Volume 2	Online		4.00 TB	mdiskgrp0	60050760009901D37800000000000009	No
Thin Volume 3	Online		4.00 TB	mdiskgrp0	60050760009901D3780000000000000A	No
Vmware 1 Thin	Online		2.00 TB	mdiskgrp0	60050760009901D3780000000000000B	No
Vmware 2 Thin	Online		2.00 TB	mdiskgrp0	60050760009901D3780000000000000C	No
Vmware 3 Thin	Online		2.00 TB	mdiskgrp0	60050760009901D3780000000000000D	No
Vmware 4 Thin	Online		2.00 TB	mdiskgrp0	60050760009901D3780000000000000E	No

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

Filter by column

Click the magnifying glass adjacent to the filter field to activate the filter by column feature. This feature allows the filter field value to be matched to the values of a specific column. Figure 3-26 on page 74 shows the column filter has been set to “Host Mappings” and the filter value can be selected.

Name	Status	UID	Capacity Utilization	Capacity	Storage Pool	Host Mappings	MDisk ID
esxi_vof1	Online	60050760009A81D37800000000000017	100%	1.00 GB	BRONZE POOL	Yes	
ESX_VMFs1	Online	60050760009A81D37800000000000010		25.00 GB	BRONZE POOL	Yes	many
Copy 0*	Online	60050760009A81D37800000000000010	100%	25.00 GB	BRONZE POOL	Yes	many
Copy 1	Online	60050760009A81D37800000000000010	100%	25.00 GB	BRONZE POOL	Yes	many
controller1_000000000000...	Online	60050760009A81D37800000000000012	100%	50.00 GB	GOLD POOL	No	
Volume_Test	Online	60050760009A81D37800000000000015	100%	100.00 GB	Pool3	No	
Fileserver	Online	60050760009A81D37800000000000016		100.00 GB	Pool3	Yes	
SQL_DB	Online	60050760009A81D37800000000000013	100%	100.00 GB	Pool3	Yes	
Volume 5	Online	60050760009A81D37800000000000008	100%	200.00 GB	BRONZE POOL	No	
Volume 6	Online	60050760009A81D37800000000000006	100%	200.00 GB	BRONZE POOL	No	
Volume 4	Online	60050760009A81D37800000000000004	100%	200.00 GB	BRONZE POOL	No	
Volume 3	Online	60050760009A81D37800000000000003	100%	200.00 GB	BRONZE POOL	No	
Volume 2	Online	60050760009A81D37800000000000002	100%	200.00 GB	BRONZE POOL	No	
Volume 4_01	Online	60050760009A81D37800000000000014	100%	200.00 GB	Pool2	No	

Figure 3-26 Filter by column

3.3 Status Indicators menus

This section provides further information on the horizontal bars displayed at the bottom of the management GUI panels. The bars are *status indicators*, and have associated bar menus. This section covers the Allocated, Running Tasks, and Health Status bar menus.

3.3.1 Horizontal bars

Shown in 3.1.3, “Overview panel layout” on page 62 the *status indicators* include the *allocated*, *running tasks*, and *health status* horizontal bars and appear at the bottom of the panel. The status indicators are colored and draw attention to alerts, events and errors. Hovering and clicking on the bars will reveal further menus.

3.3.2 Allocated status bar menu

The allocated status bar shows capacity status. Hovering over the image of two arrows on the right hand side of the allocated status bar reveals a description of the allocated menu comparison currently in use. Figure 3-27 shows the comparison of the *used* capacity to the *real* capacity.

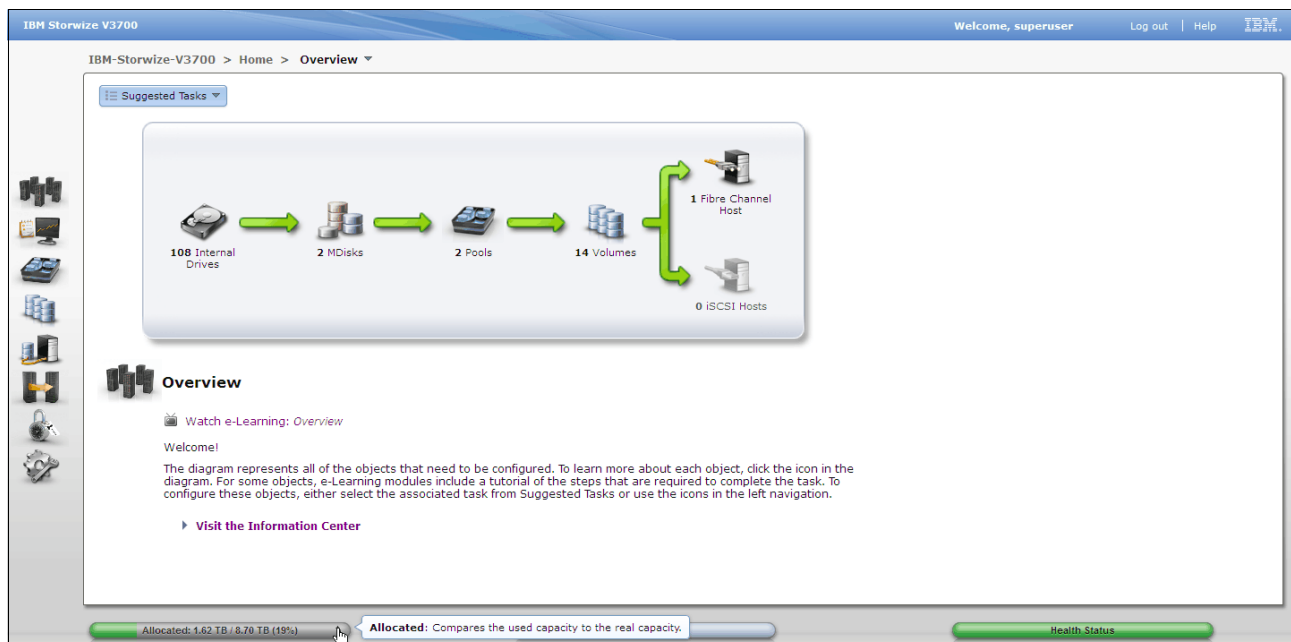


Figure 3-27 Allocated bar, comparing used capacity to real capacity

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

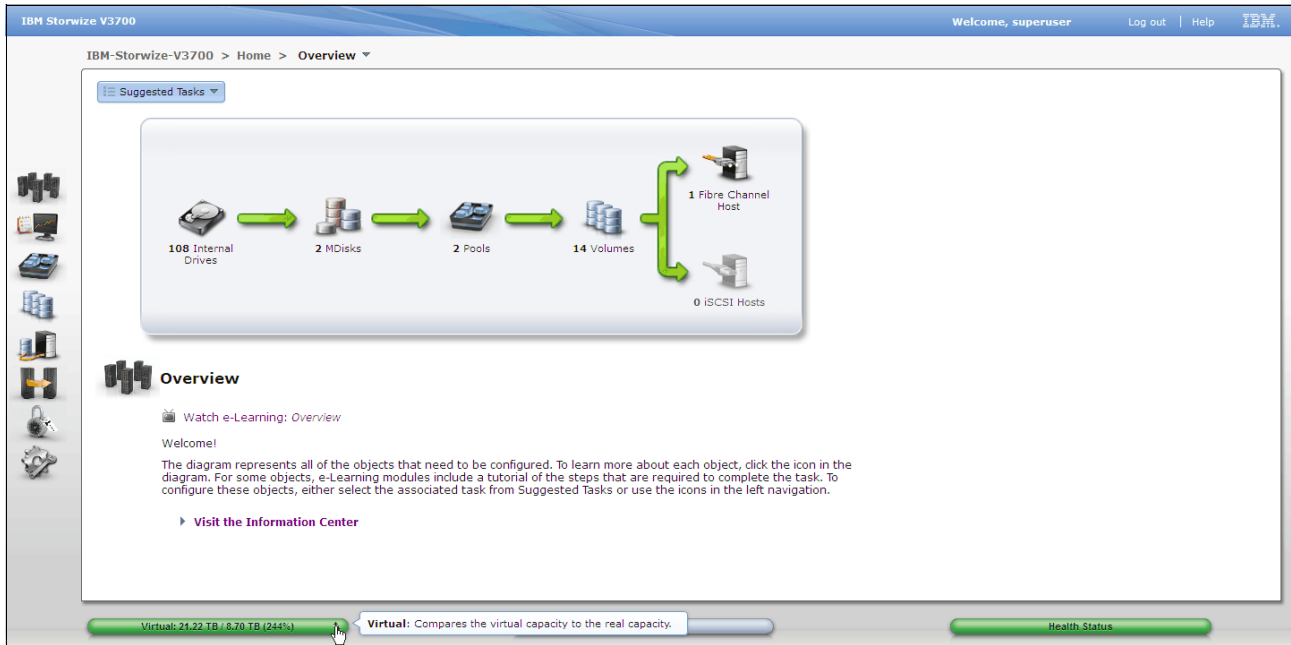


Figure 3-28 Changing the allocated menu comparison, virtual capacity to real capacity

3.3.3 Running tasks bar menu

To reveal the Running Tasks bar menu click the circular image at left end of the running tasks status bar. This menu lists running and recently completed tasks, it will also group like tasks. Figure 3-29 shows Running Tasks bar menu.

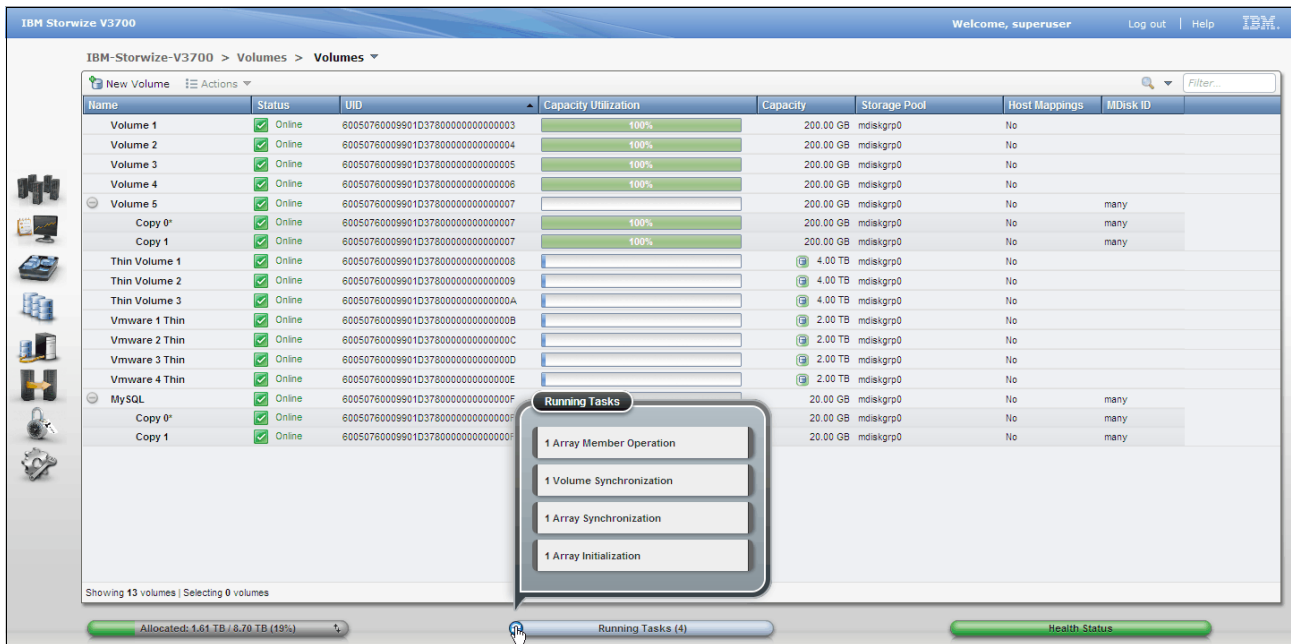


Figure 3-29 Revealing the running tasks bar menu

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

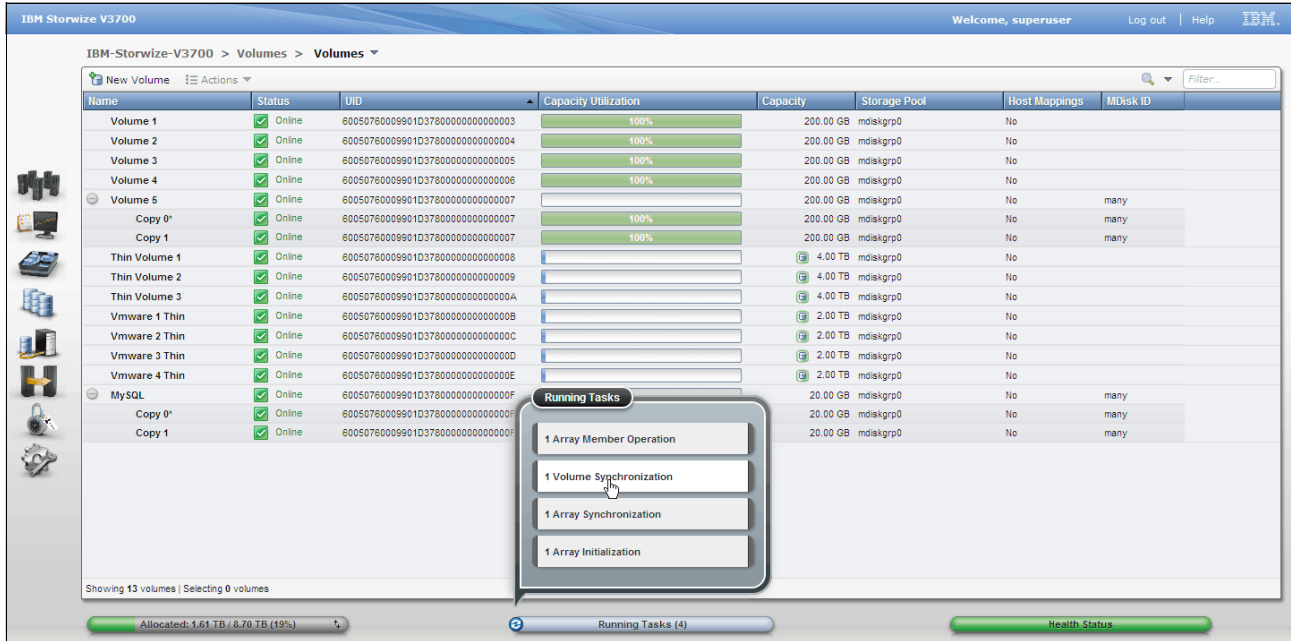


Figure 3-30 Selecting a task from the running task menu for and indication of task progress

Figure 3-31 on page 77 shows the progress of all running “Volume Synchronization” tasks.

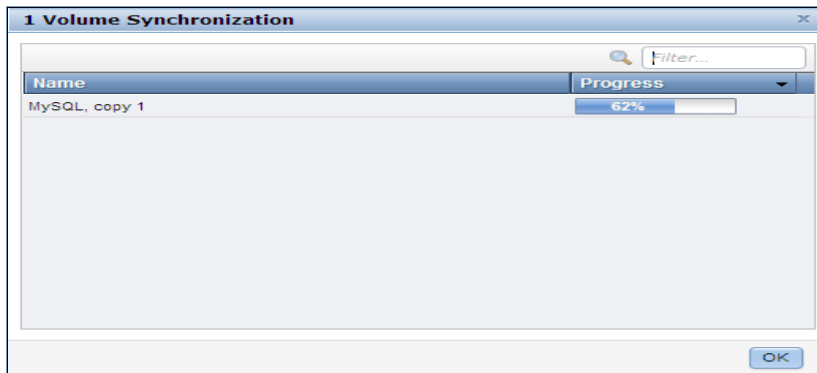


Figure 3-31 Progress of all running “Volume Synchronization” tasks

3.3.4 Health status bar menu

The health status bar provides an indication of the overall health of the system. The color of the status bar denotes the state of Storwize V3700. Green is healthy, yellow is degraded, and red is unhealthy. If a status alert occurs, the health status bar can turn from green to yellow, or red. To reveal the health status menu, simply click the triangular attention image on the left hand side of the health status bar. Figure 3-32 shows the revealed health status menu.

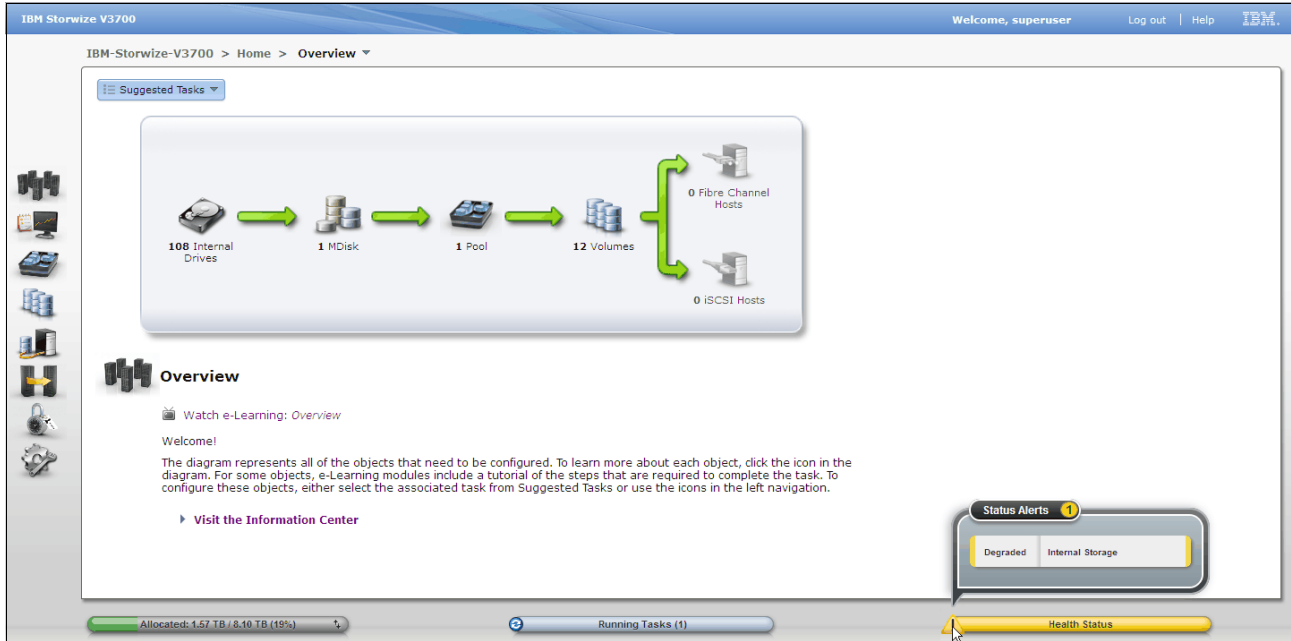


Figure 3-32 Revealing the health status menu

In Figure 3-33 on page 78 the health status bar menu reports the system as *degraded* and provides a description of “Internal Storage” for the type of event that has occurred. To investigate the event, open the health status bar menu and click on the description of the event. Figure 3-33 shows an event in the health status bar menu.




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

A single click on the description of the event in the health status menu reveals the Events panel (**Monitoring** → **Events**). This panel lists all events and provides directed maintenance procedures to help resolve errors. Figure 3-34 shows the Events panel.

Note: Refer to “Events panel” on page 86 for further information.

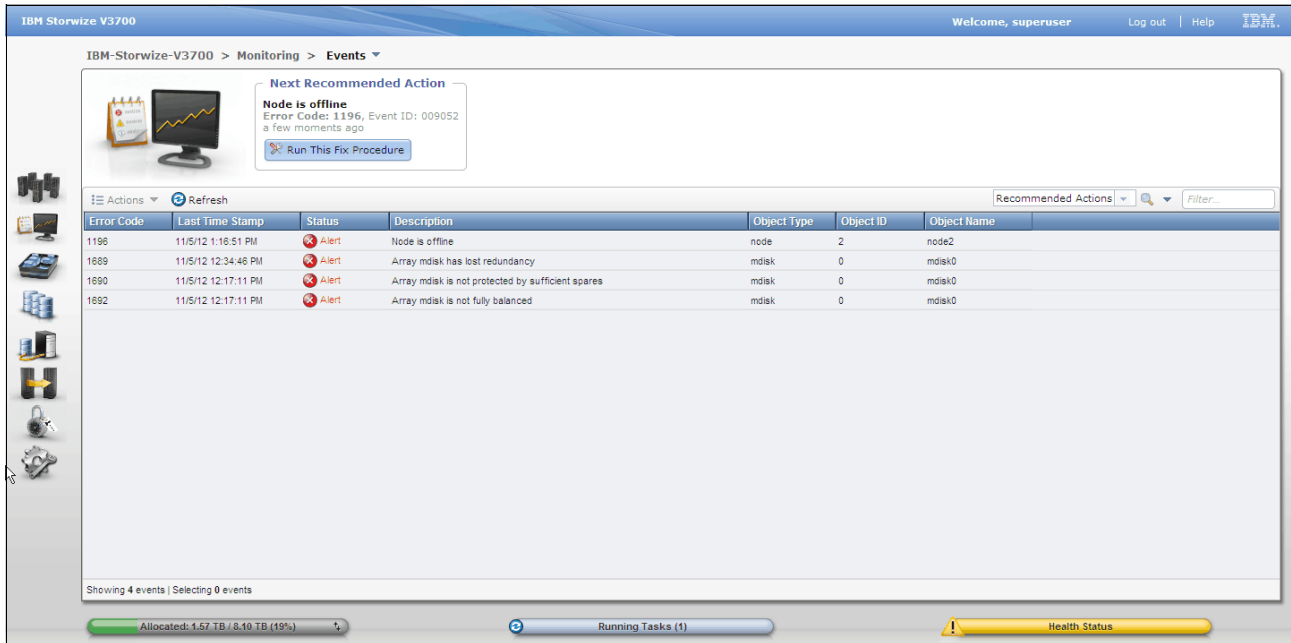


Figure 3-34 Events panel via health status menu

3.4 Function icon menus

The Storwize V3700 management GUI provides *function icons* that are an efficient and quick mechanism for navigation. As discussed in section 3.1.3, “Overview panel layout” on page 62, each small graphic on the left hand side of the panel is a *function icon* and represents a group of interface functions. Hovering over one of the eight *function icons* reveals a menu that lists the functions. Figure 3-35 shows all function icon menus.




Figure 3-35 All function icon menus

3.4.1 Home menu

The Home menu provides access to the Overview panel. Figure 3-36 shows the Home menu.

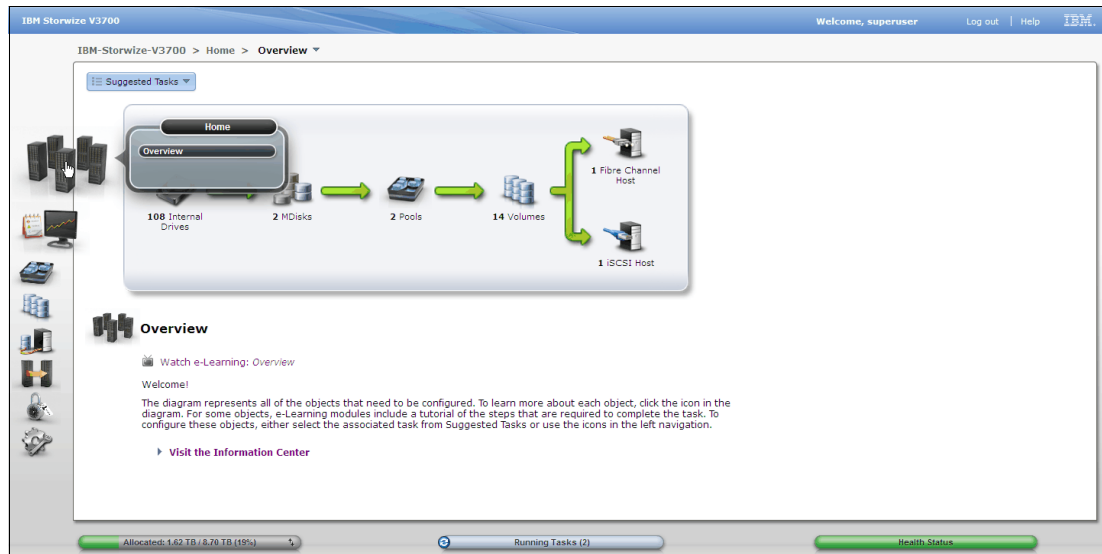


Figure 3-36 Home menu

Overview panel

Select **Overview** in the Home menu to reveal the panel. Detailed information on the Overview panel can be found in section 3.1.3, “Overview panel layout” on page 62.

3.4.2 Monitoring menu

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




Figure 3-37 Monitoring menu

System panel

Select **System** in the Monitoring menu to reveal the panel. The System panel shows capacity usage, the enclosures, and all drives in the system. Figure 3-38 on page 82 shows the System panel.

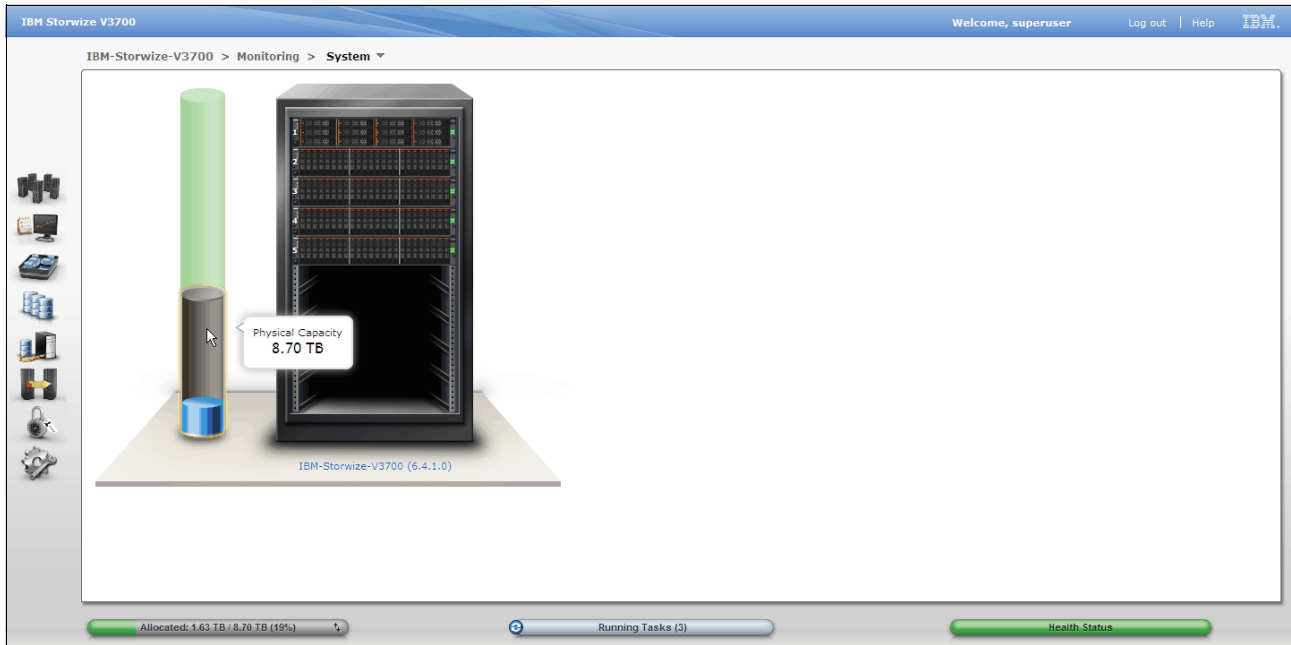


Figure 3-38 The system panel

Selecting the hotlink which includes the name and version of the system reveals additional information related to storage allocation. The information is presented under two tabs, the **Info** tab, and the **Manage** tab. Figure 3-39 on page 82 shows the System panel Info tab.

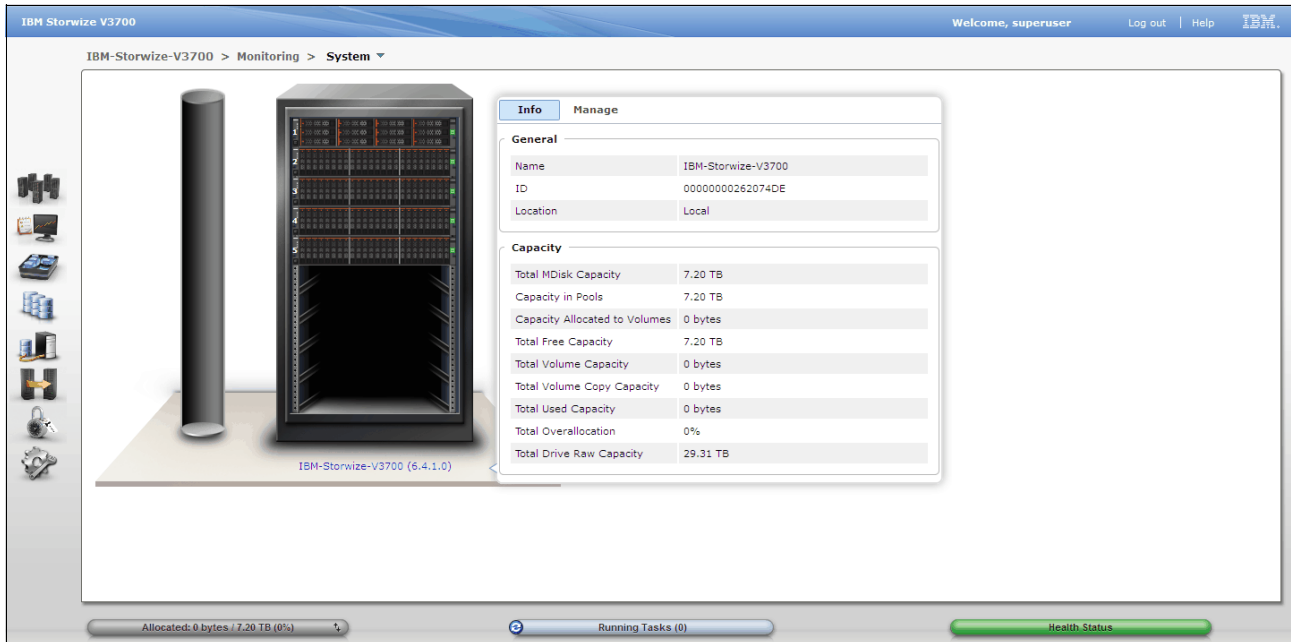


Figure 3-39 System panel info tab

Selecting the **Manage** tab will reveal the name of the system, and shutdown and upgrade actions. Figure 3-40 shows the System panel Manage tab.

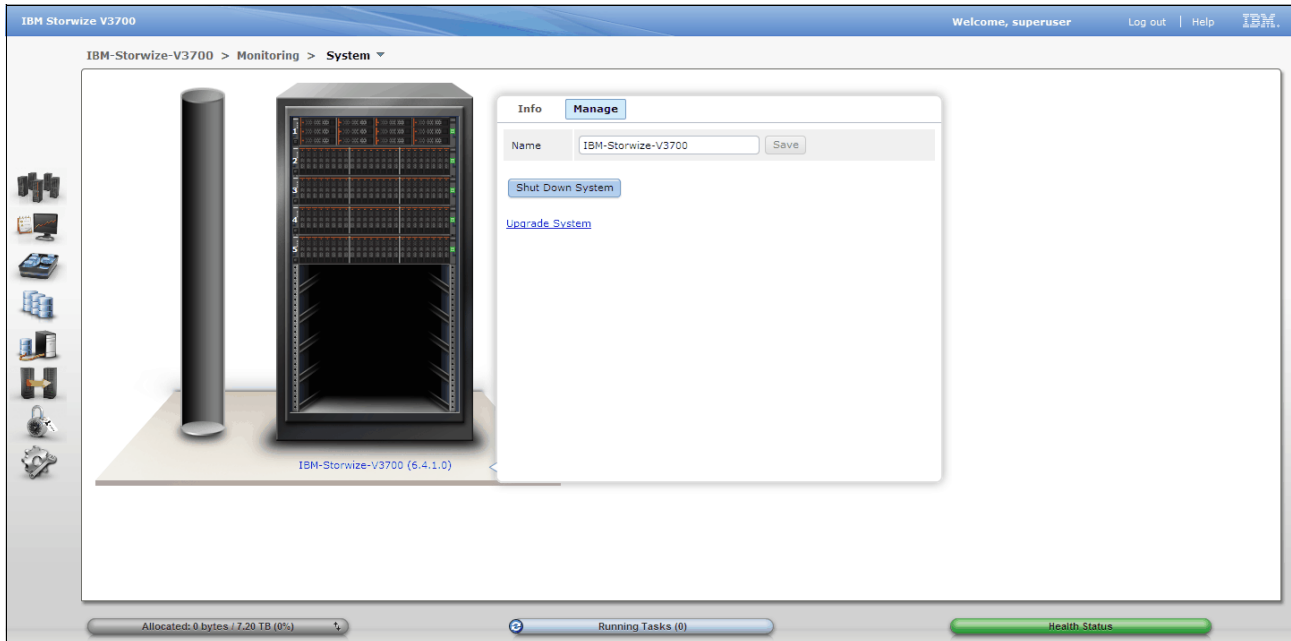


Figure 3-40 System panel manage tab

Selecting a rack mounted enclosure reveals additional information. Hovering over a drive will provide the drive status, size, and speed details. **Identify** will activate the blue identification led on the front of the enclosure. The **Enclosure 1** hotlink will reveal the System Details panel, see "System details panel" on page 84 for further information on this function. Figure 3-41 on page 83 shows the System panel enclosure view.

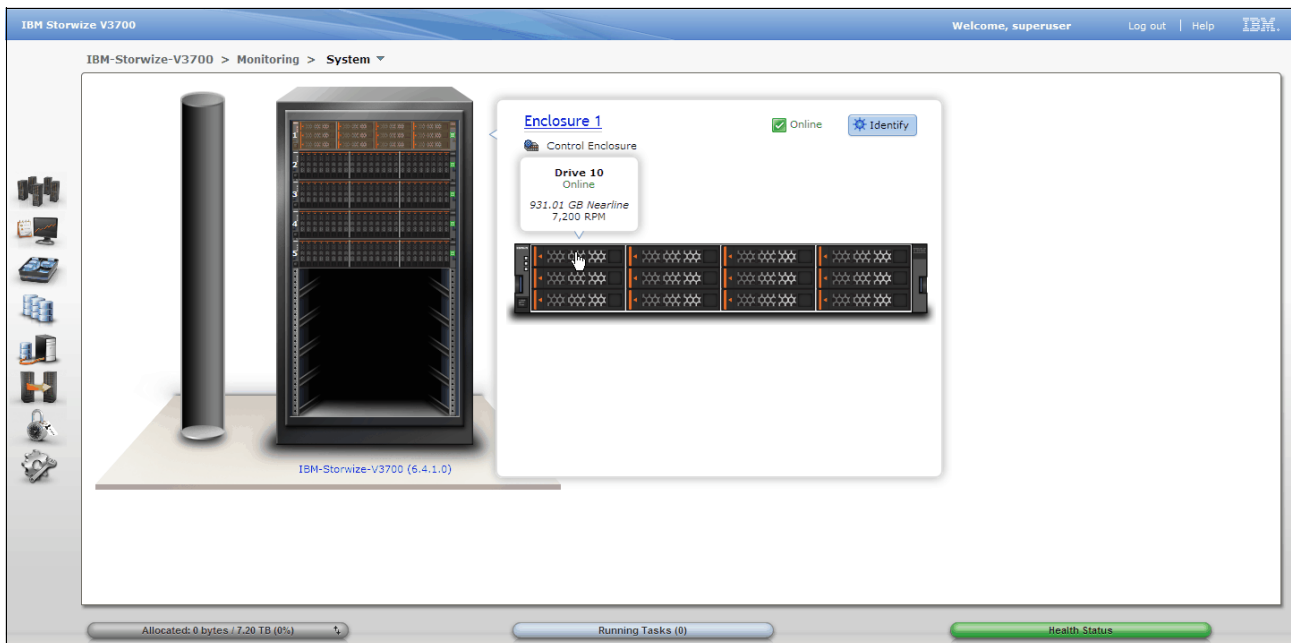


Figure 3-41 System panel enclosure view

System details panel

Select **System Details** in the Monitoring menu to access the panel. The System Details panel provides the status and details of the components that make up the system. Figure 3-42 on page 84 shows the System Details panel.

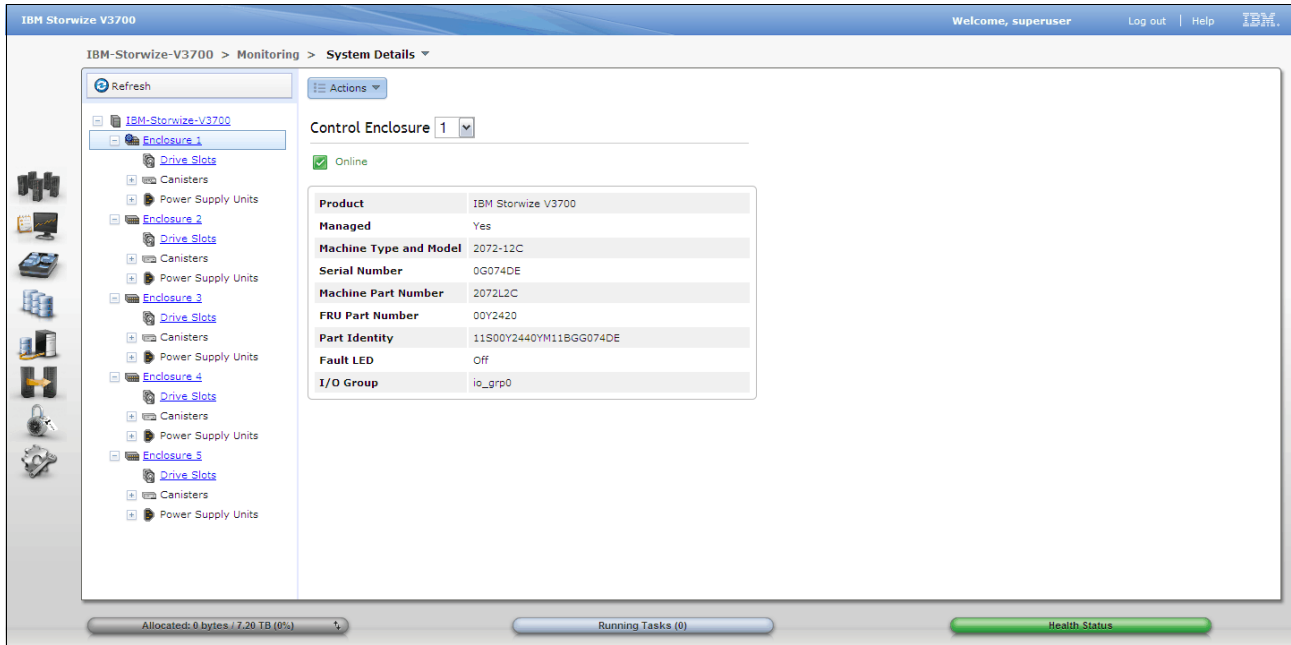


Figure 3-42 System details panel

Actions and environmental statistics

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

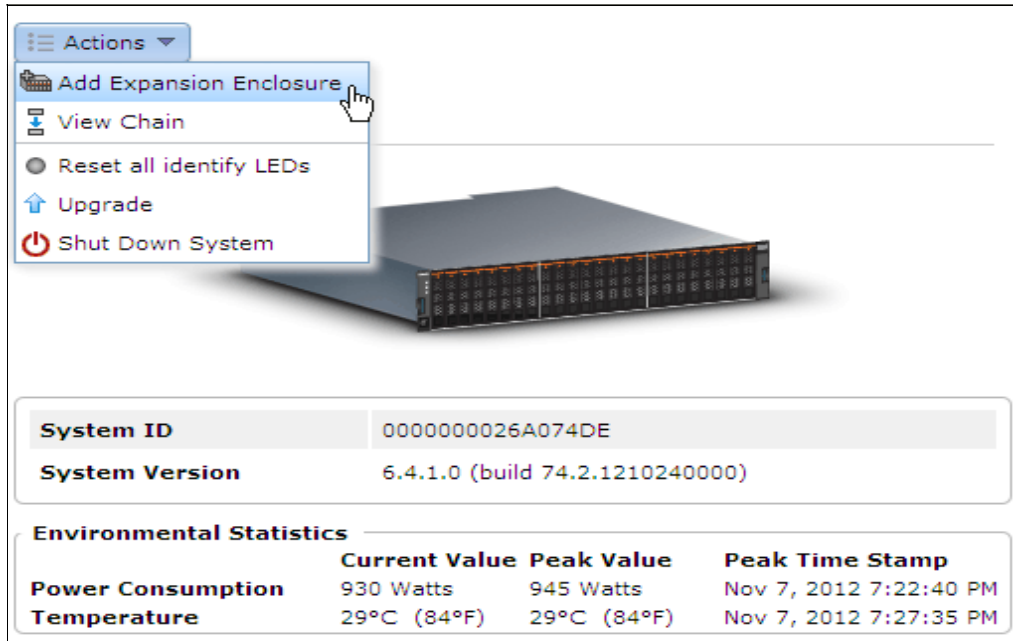


Figure 3-43 System details actions and environmental statistics

Node canister information

Additionally, node canister information such as FC WWPNs and iSCSI IQNs is useful for host attachment purposes, this detail is revealed by clicking on the control enclosure node canister in the System Details panel. Figure 3-44 shows node canister information.

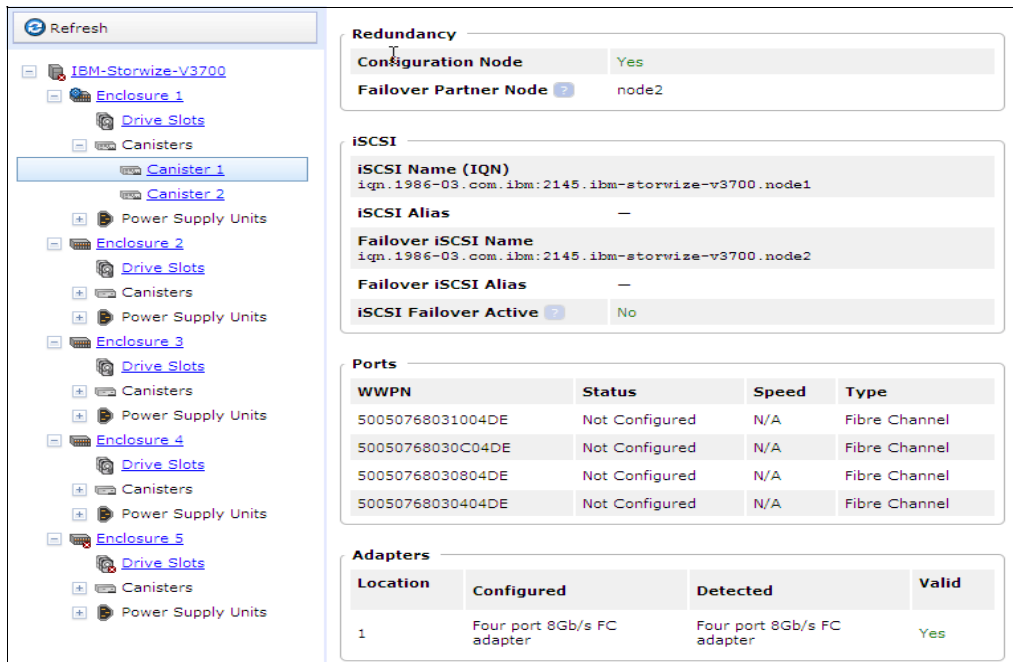


Figure 3-44 Node canister information via system details panel

Events panel

Select **Events** in the Monitoring menu to reveal the panel. The machine is optimal when all errors are addressed and there are no items found in the Events panel. Figure 3-45 on page 86 shows the Events panel with no recommended actions.

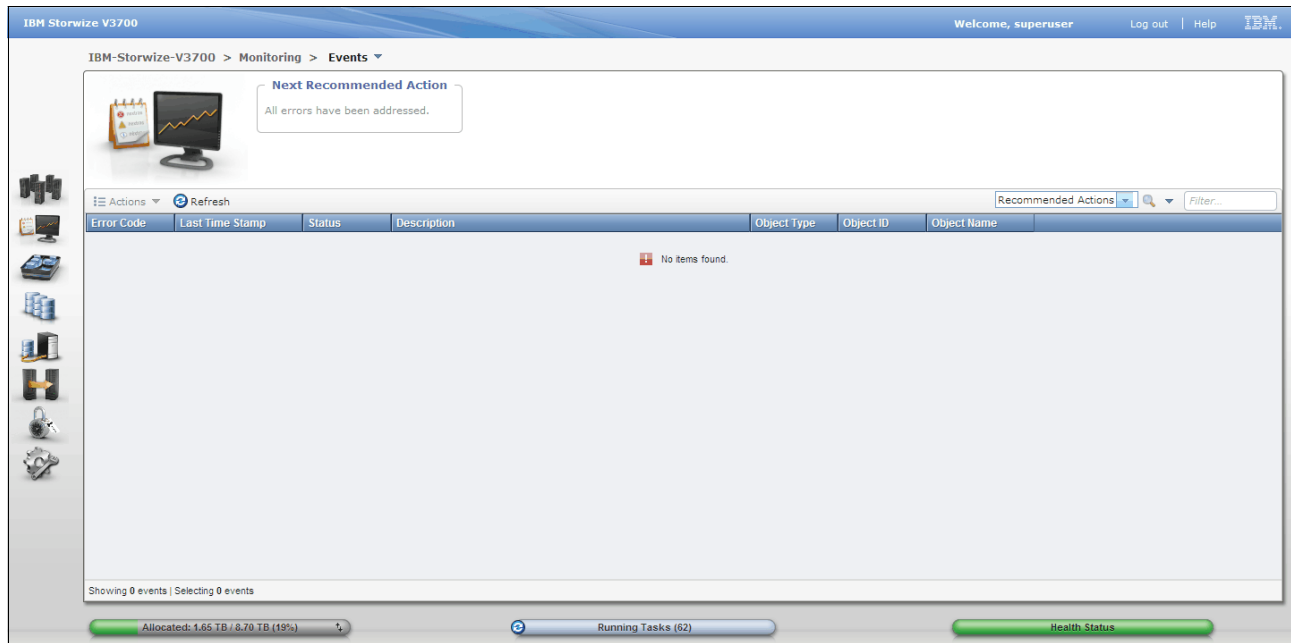


Figure 3-45 Events panel with all errors addressed

Filtering events view

To view unfixed messages and alerts select **Unfixed Messages and Alerts** option from the drop down box adjacent to the filter field. Figure 3-46 shows unfixed messages and alerts.

Note: Using the filter field is described in detail in section “Filtering objects” on page 73.

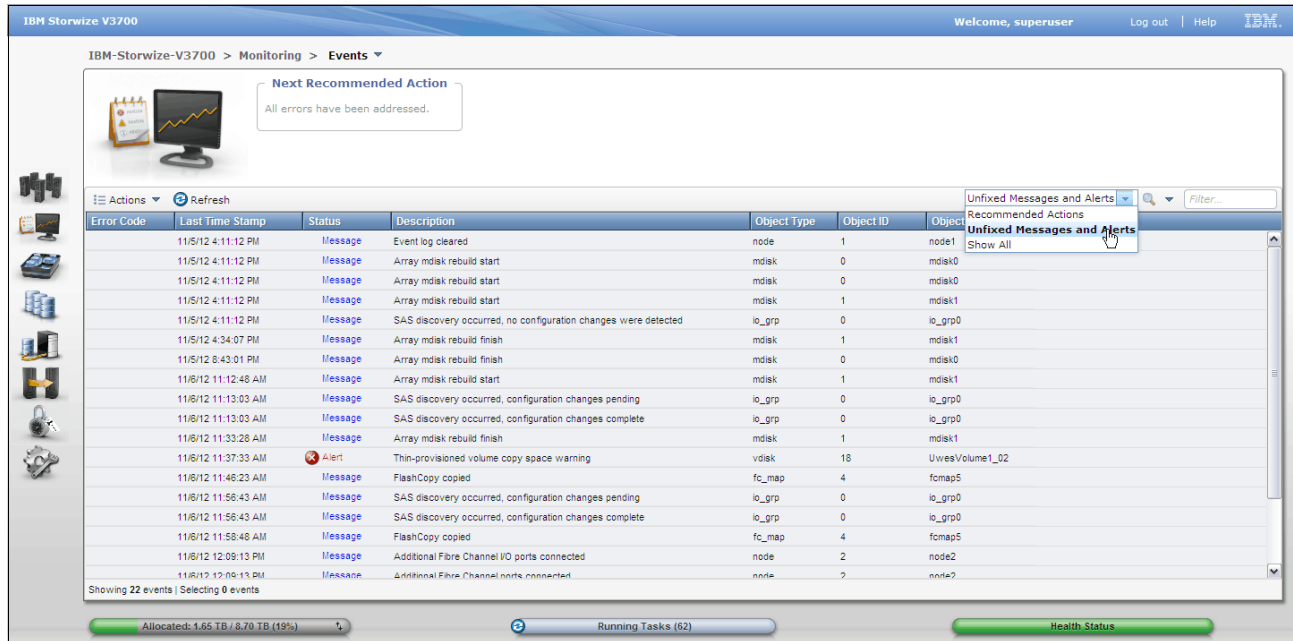


Figure 3-46 Unfixed messages and alerts in the events panel

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

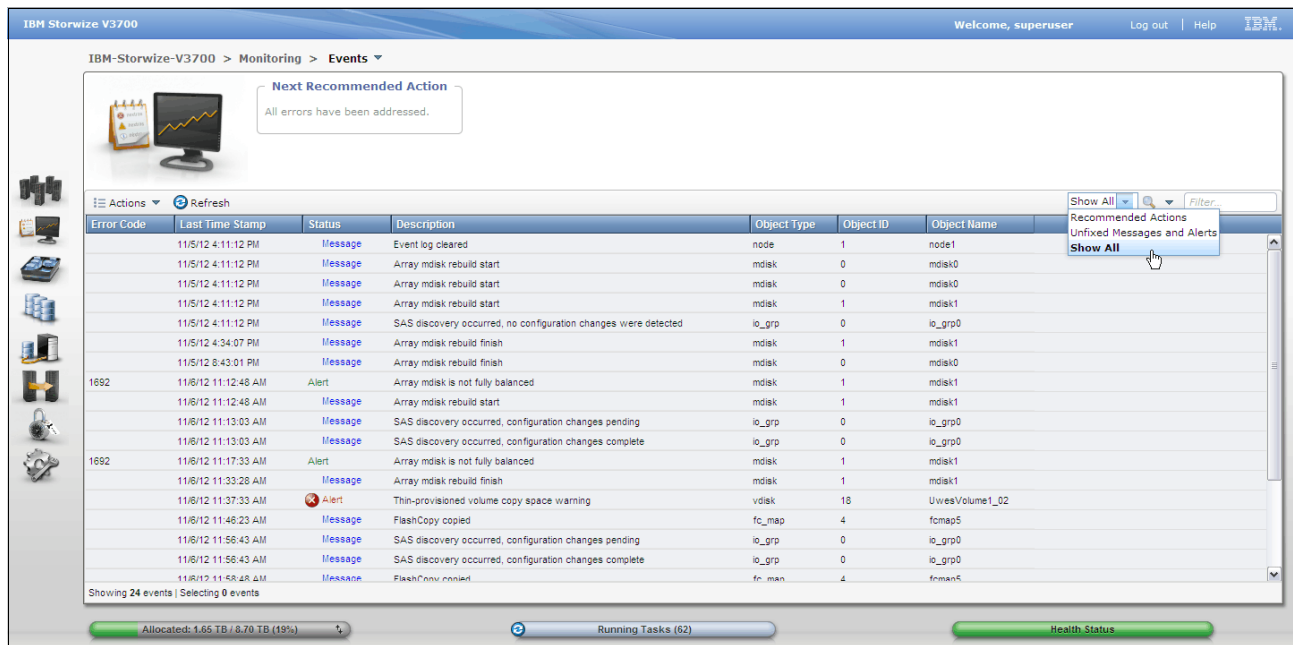


Figure 3-47 Events panel with show all selected

Event properties

To reveal actions and properties related to an event, or to repair an event that is not the “Next Recommended Action”, simply right click on the event to reveal further options. Figure 3-48 shows the selection of **Properties** after a right click on the selected event.

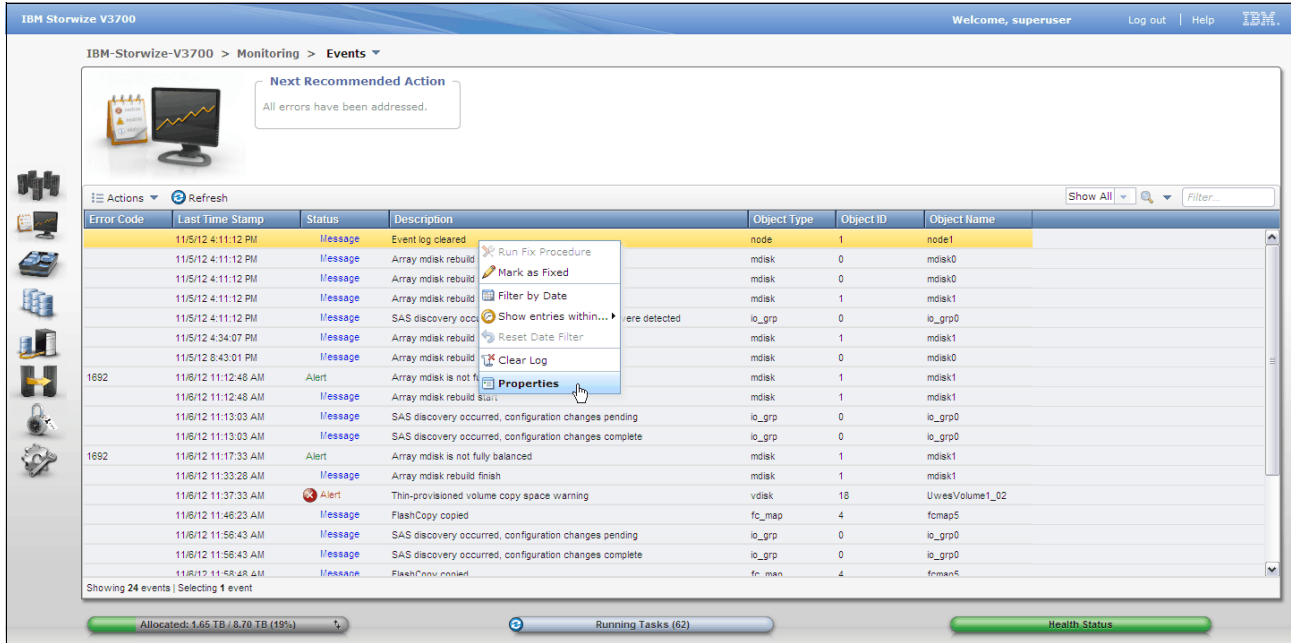


Figure 3-48 Selecting event properties

Figure 3-49 on page 88 shows the properties of an event.

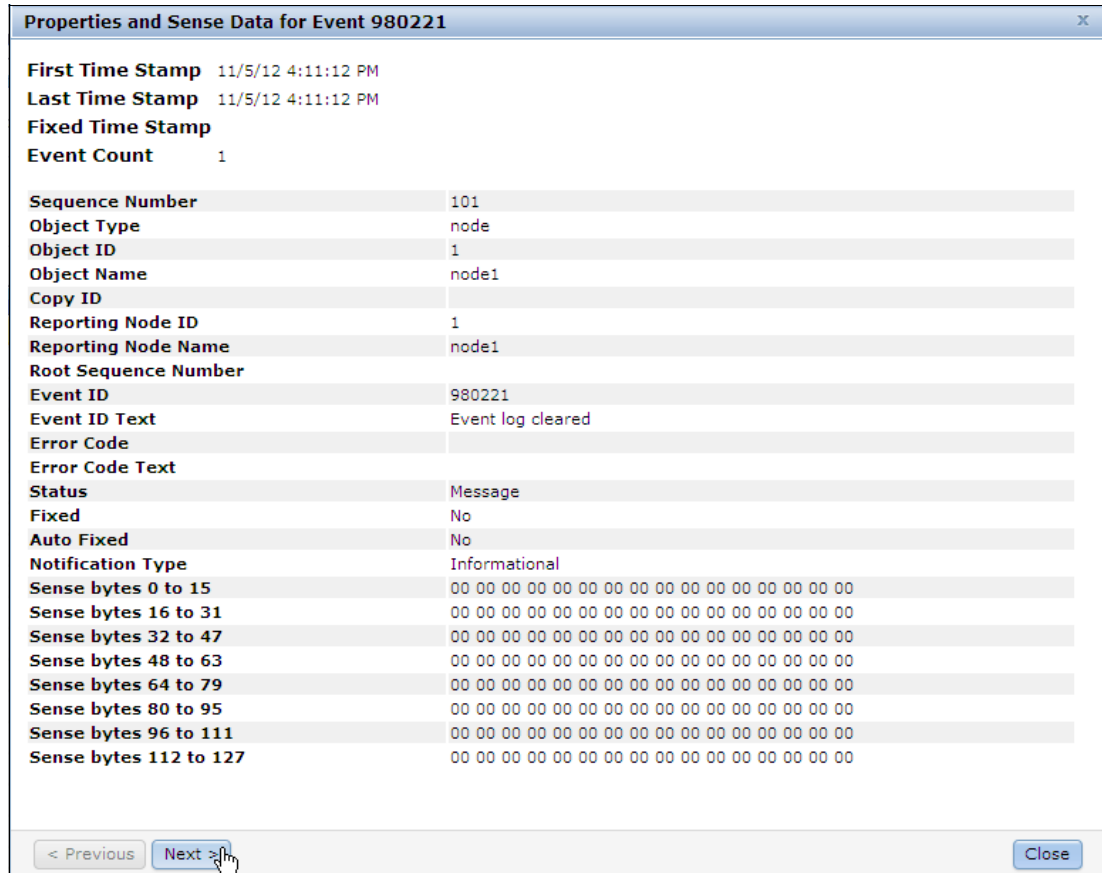


Figure 3-49 Event properties

Show events entries within...

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

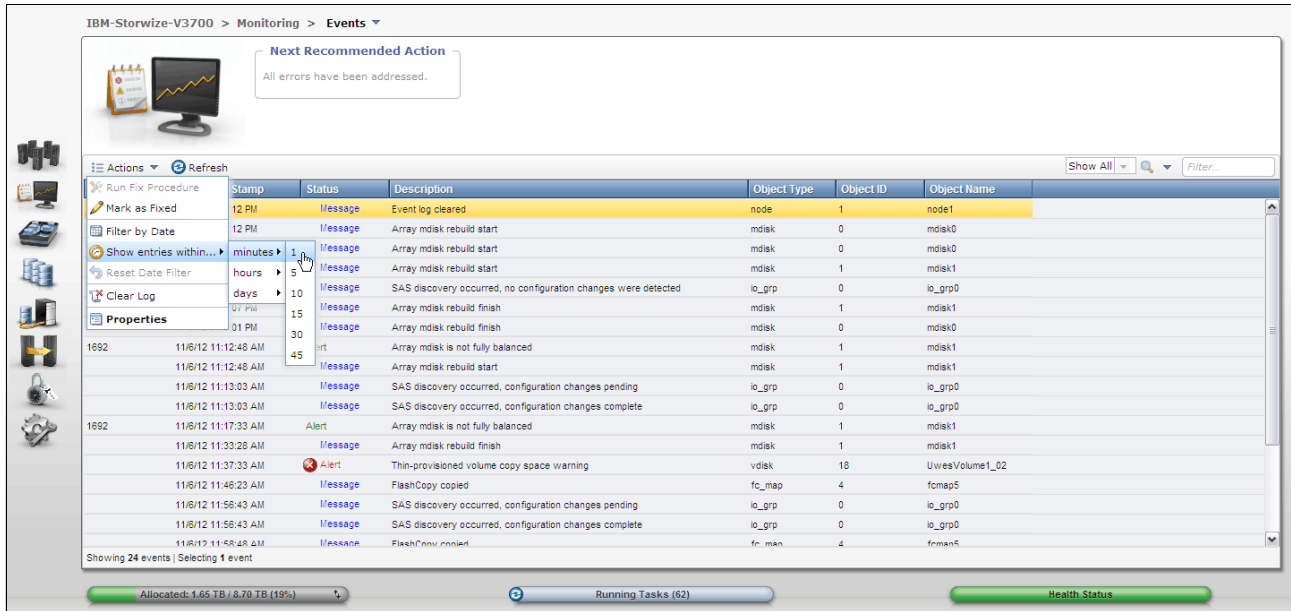


Figure 3-50 Showing events within a set time

Performance panel

Select **Performance** in the Monitoring menu to reveal the panel. This panel reveals graphs that represent the last 5 minutes of performance statistics. The performance graphs include statistics related to CPU utilization, volumes, mdisks, and interfaces. Figure 3-51 shows the Performance panel.

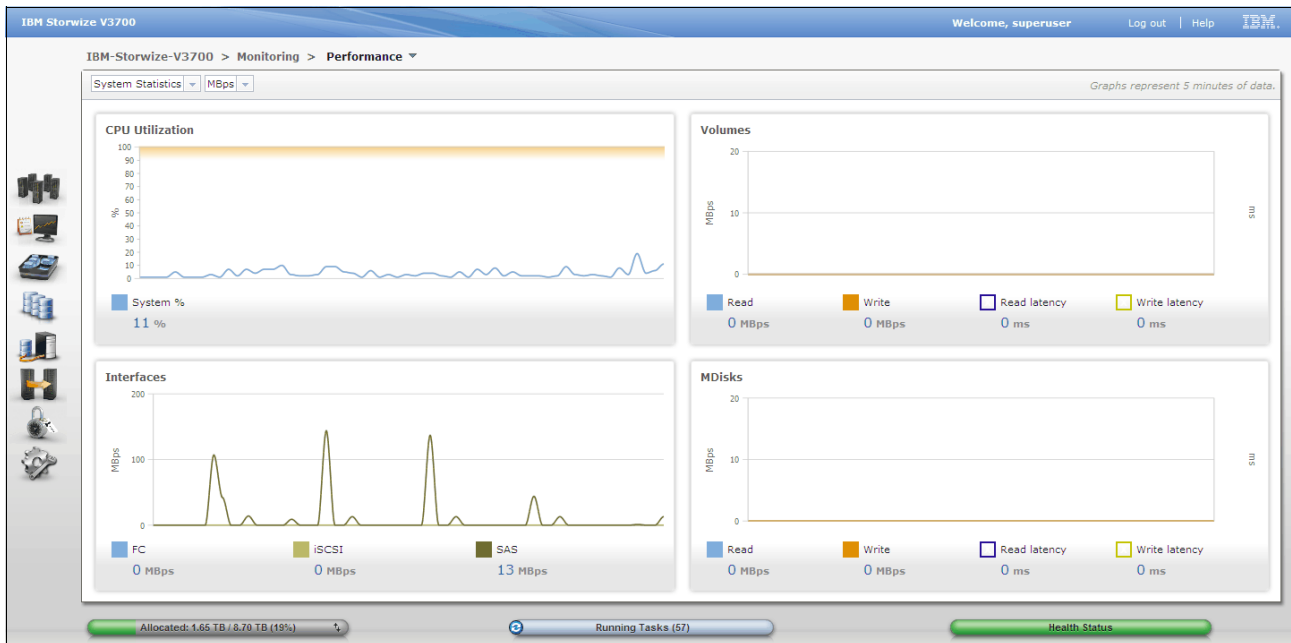


Figure 3-51 Performance panel

Custom-tailoring performance graphs

The Performance panel can be customized to show the workload of a single node. This is useful to help determine if the system is working in a balanced manner. Figure 3-52 on page 90 shows the custom-tailoring of the performance graphs by selecting node 1 from the **System Statistics** drop down box.

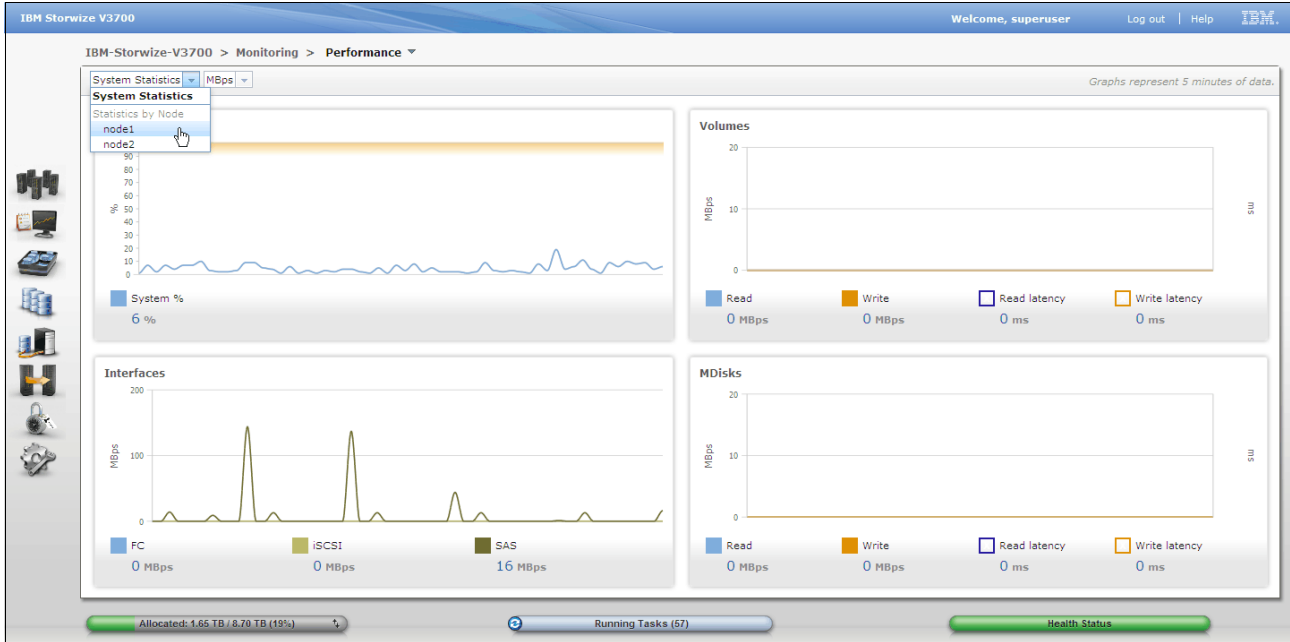


Figure 3-52 Graphs representing performance statistics of a single node

Performance peak value

Peak values over the last 5 minute period can be seen, hovering over the current value will reveal the peak value. Figure 3-53 shows the 5 minute peak for CPU utilization.

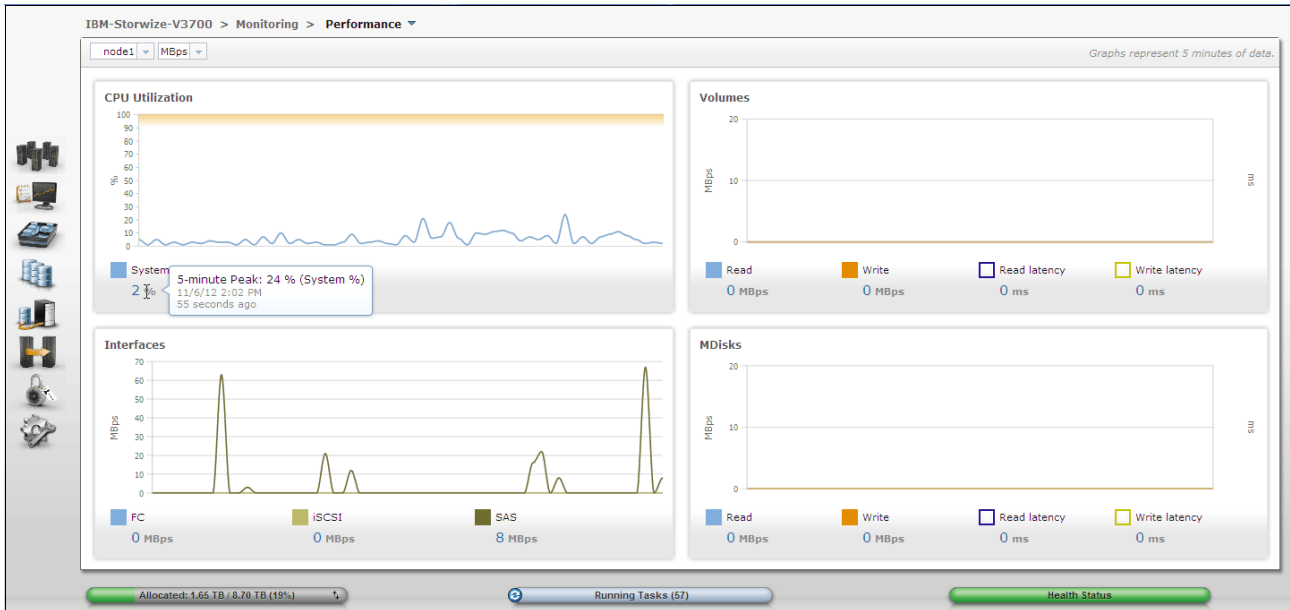


Figure 3-53 Peak CPU utilization value over the last 5 minutes

3.4.3 Pools menu

The Pools menu provides access to the Volumes by Pools, Internal Storage, Mdisks by Pools, and System Migration functions. Figure 3-54 on page 91 shows selection of the Pools menu.

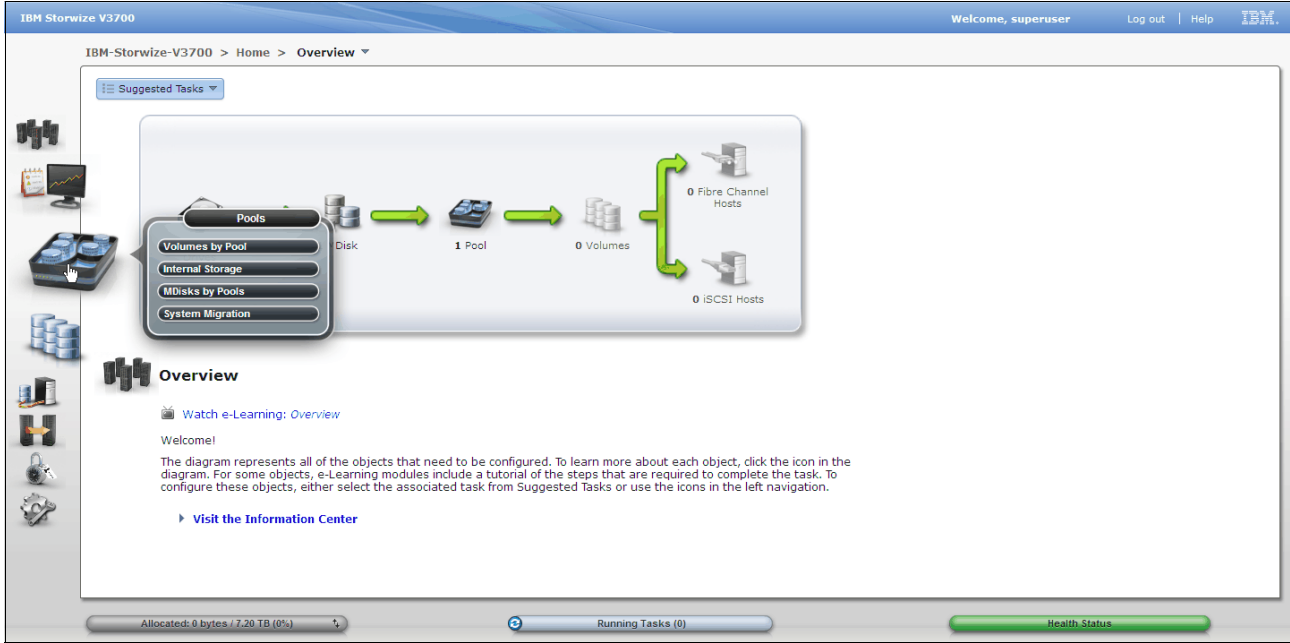


Figure 3-54 Pools menu

Volumes by pools panel

Select **Volumes by Pool** item in the Pools menu to reveal the panel. The Volumes by Pool panel provides the ability to display volumes using the “Pool Filter” function. This view makes it easy to manage volumes and determine the amount of *real* capacity that is available for further allocations. Figure 3-55 shows the Volumes by Pool panel.

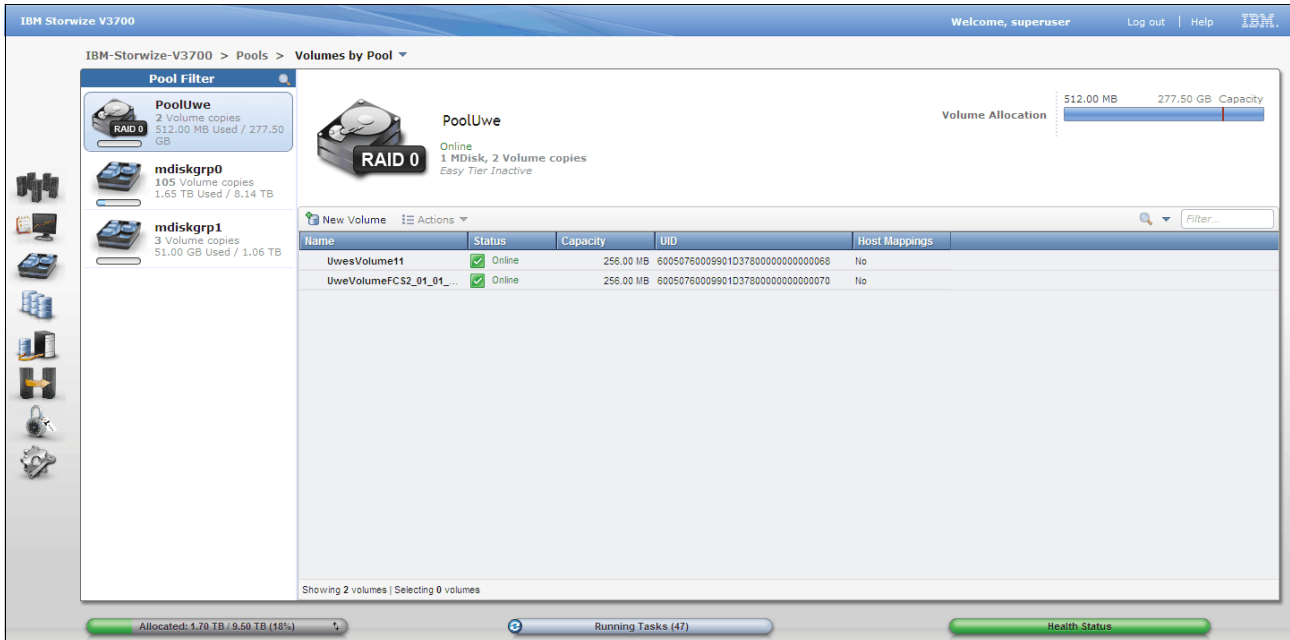


Figure 3-55 Volumes by pools panel

Volume allocation

The top right hand corner of the Volumes by Pool panel shows the actual volume allocation which is the amount of capacity currently allocated to volumes in this pool. Figure 3-56 highlights the volume allocation information that is displayed in the Volumes by Pool panel.




Figure 3-56 Volume allocation

Renaming pools

To change the name of a pool, simply select the pool from the pool filter and click on the name of the pool. Figure 3-57 on page 92 shows that pool “mdiskgrp0” has been renamed to “Gold Pool”.

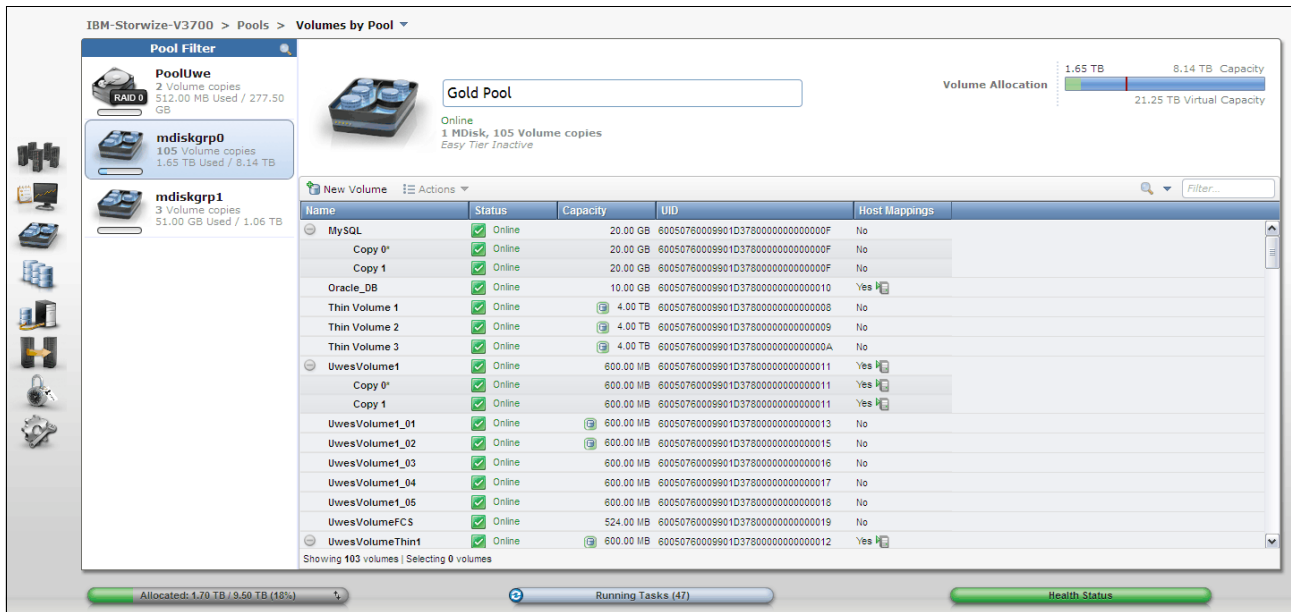


Figure 3-57 Renaming a pool

Changing pool icons

To change the icon associated with a pool, simply select the pool in the pool filter and click on the large pool icon above **New Volume** and **Actions**. This helps to simply management and differentiate between the different classes of drive or the tier of the storage pool. Figure 3-58 shows the pool change icon panel.



Figure 3-58 Changing a pool icon

Volume functions

The Volumes by Pool panel also provides access to the volume functions via the **Actions** drop box, the **New Volume** button, and via a right click on a listed volume. For further information on navigating the Volume panel see 3.4.4, “Volumes menu” on page 99.

Figure 3-59 on page 93 shows the volume functions available via the Volumes by Pool panel.

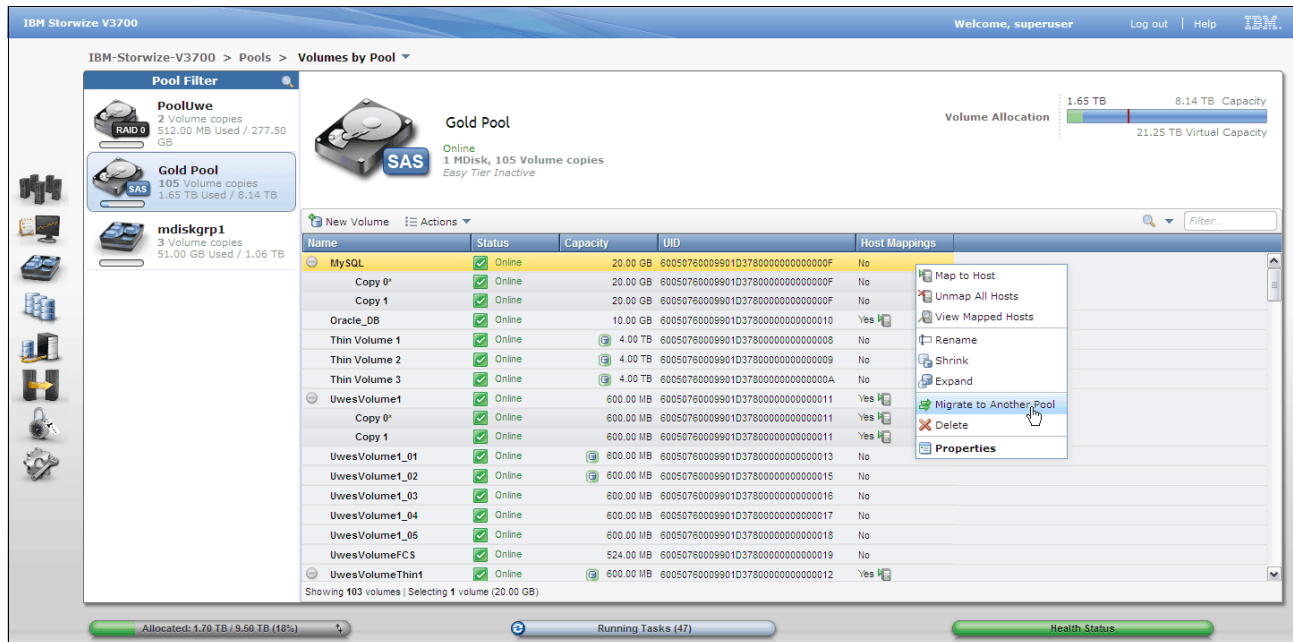


Figure 3-59 Volume functions are available via volume by pools panel

Internal storage panel

Select **Internal Storage** in the pools menu to reveal the Internal Storage panel. The internal storage comprises of the drives contained in the Storwize V3700 control enclosure and the SAS attached Storwize V3700 expansion enclosures. The Internal Storage panel provides the ability to configure the internal storage into RAID-protected storage. This panel also allows the displayed drive list to be filtered by drive class. Figure 3-60 on page 94 show the Internal Storage panel.

Drive actions

Drive level functions such as identifying a drive, and marking a drive as offline, unused, candidate, or spare. Right clicking on listed drives to reveal actions menu. Alternatively, the drives can be selected and the action drop box is used. Multiple drives can be selected, see “Multiple selections” on page 72. Figure 3-60 on page 94 shows the drive actions menu.




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

Drive properties

Drive properties and dependent volumes can be displayed from the Internal Storage panel. Figure 3-61 shows the selection of drive properties from the action drop down box.

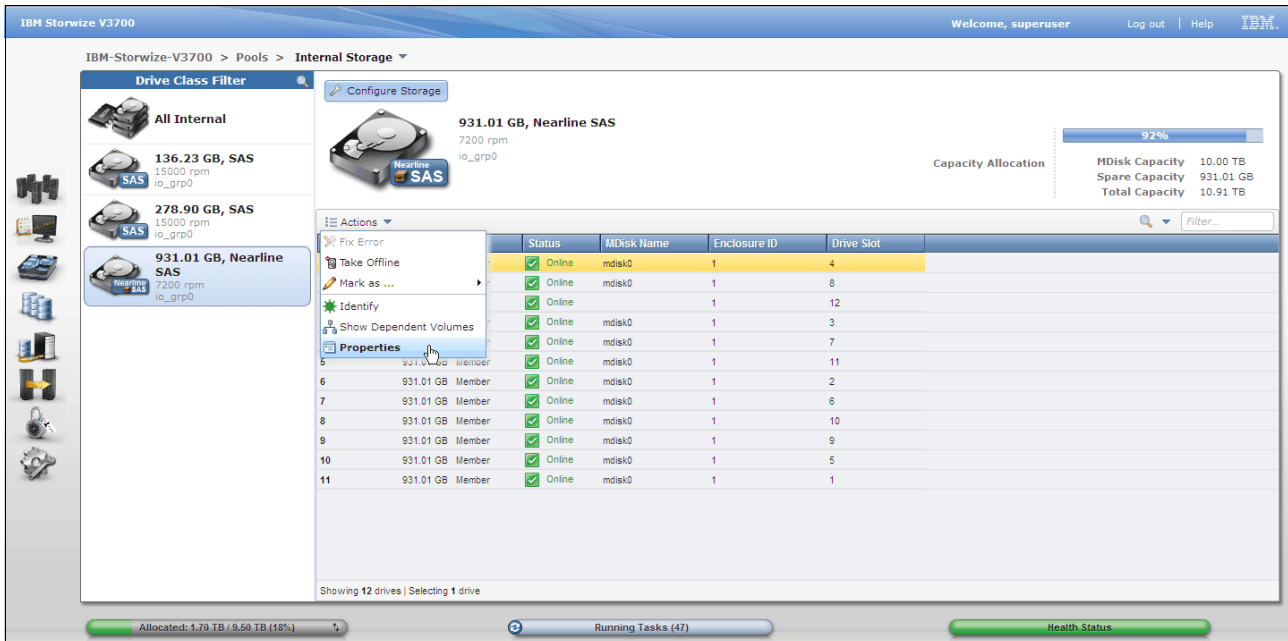


Figure 3-61 Internal storage panels actions

The drive properties panel shows drive attributes and the drive slot SAS port status. Figure 3-62 on page 95 shows drive properties.




Figure 3-62 Drive properties

Configure internal storage wizard

Click **Configure Storage** to reveal the Configure Internal Storage wizard. Figure 3-63 shows selection of the Configure Internal Storage wizard.

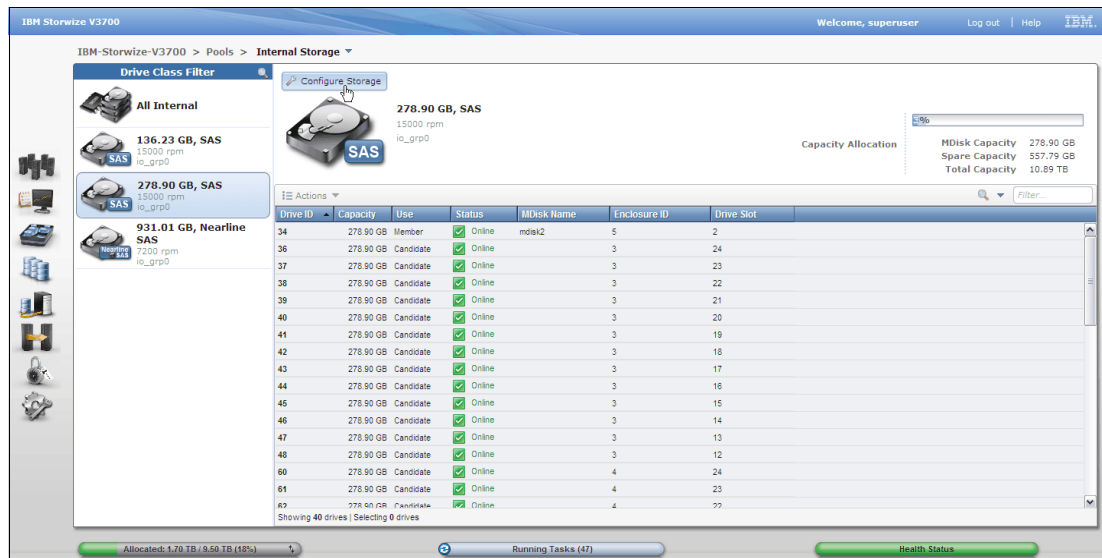


Figure 3-63 Internal Storage panel

This wizard allows the configuration of the RAID-properties and subsequent pool allocation of the internal storage. Figure 3-64 shows step 1 of the Configure Internal Storage wizard.

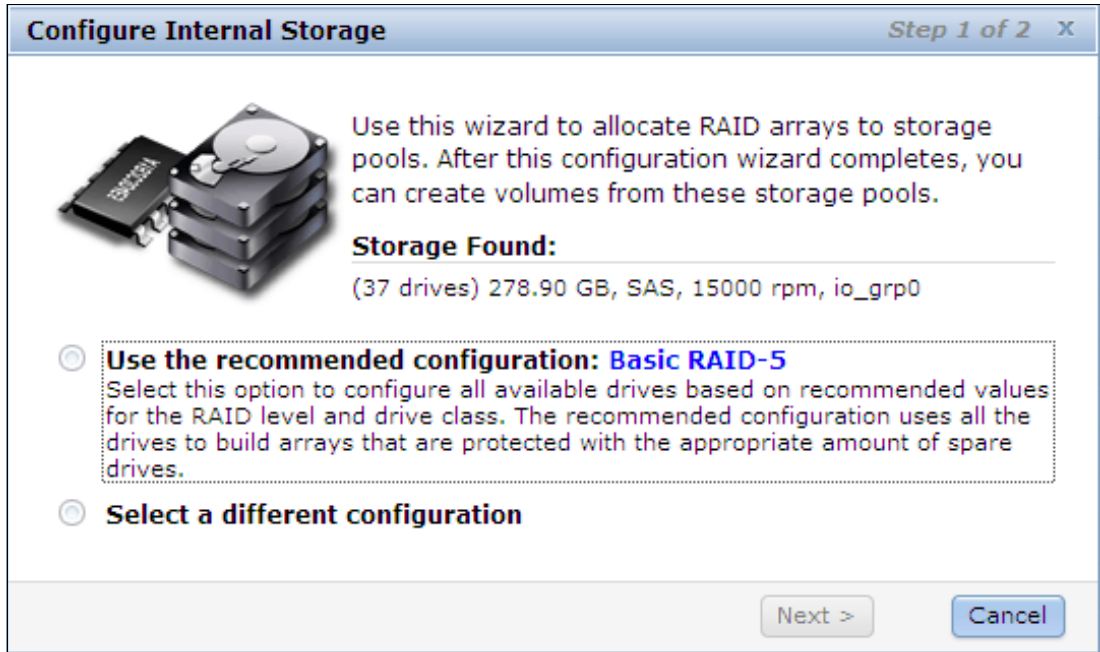


Figure 3-64 Configure internal storage wizard step 1

Figure 3-65 on page 96 shows step 2 of the Configure Internal Storage wizard.

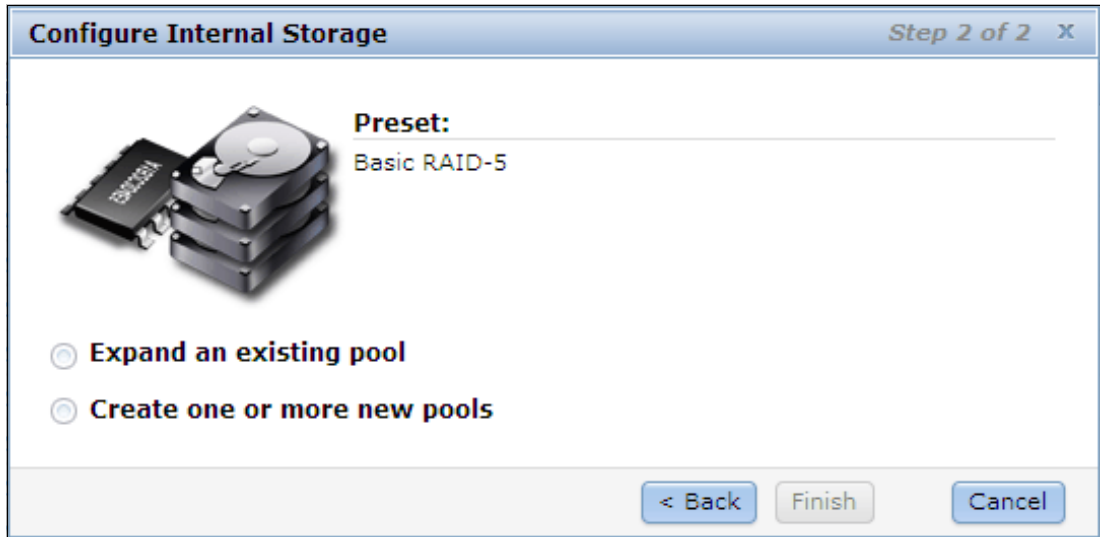


Figure 3-65 Configuring internal storage wizard step 2

Mdisks by pool panel

Select **Mdisks by Pool** in the pools menu to reveal the mdisks by pool panel. The mdisks by pool panel provides the ability to perform tasks such as: display mdisks in each pool, create new pools, delete pools, and detect externally virtualized storage. Figure 3-66 on page 97 shows the mdisks by pool panel.




Figure 3-66 Mdisks by pool panel

Pool actions

To delete a pool, or change the pool name or icon, simply right click on the listed pool. Alternatively, the actions drop down box can be used. Figure 3-67 shows the pool actions.

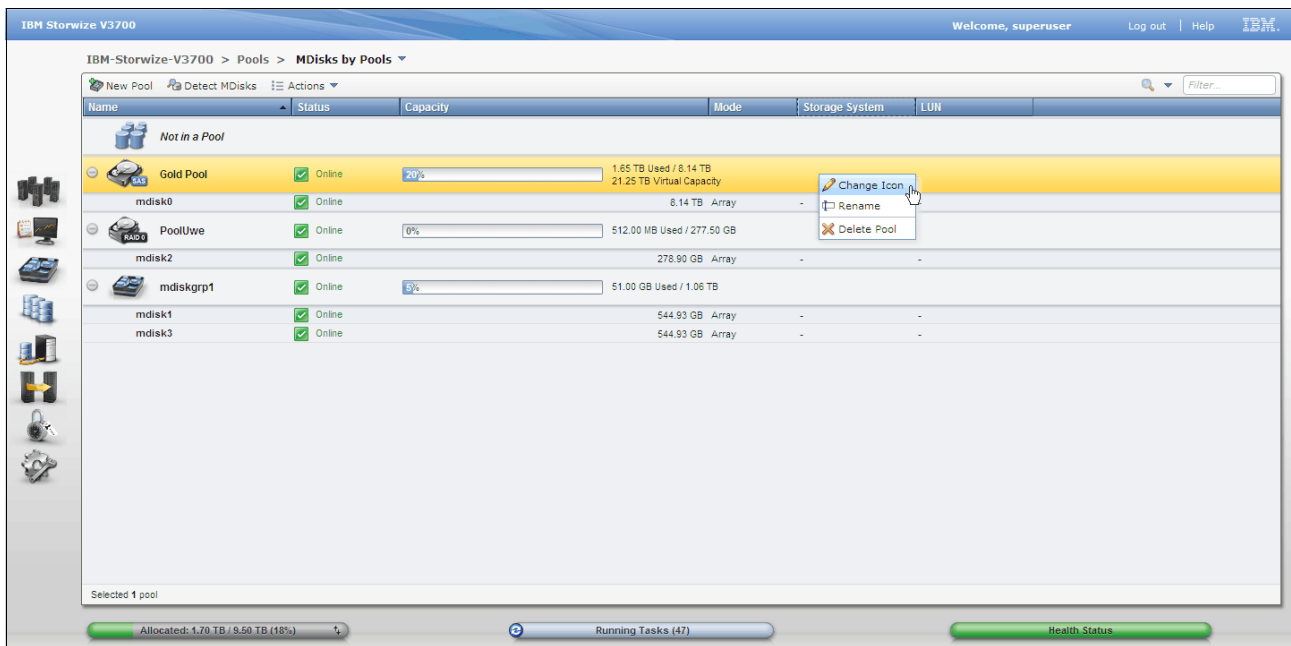


Figure 3-67 Pool actions

RAID actions

The mdisks by pool panel allows mdisk RAID actions such as: set spare goal, swap drive, and delete. To access these functions select the mdisk and right click. Figure 3-68 on page 98 the RAID actions menu.




Figure 3-68 RAID actions menu

System migration panel

Select **System Migration** in the pools menu to reveal the system migration panel. The system migration panel is used to migrate data from externally virtualized storage systems to the internal storage of the Storwize V3700. The panel will display image mode volume information. To initiate a migration, click **Start New Migration** and the Start Migration wizard will be revealed. Figure 3-69 shows the System Migration panel.




Figure 3-69 System migration panel

Storage Migration wizard

The Storage Migration wizard is used for data migration from older storage systems to the Storwize v3700. Figure 3-70 shows the storage migration wizard step 1 of 8 panel.

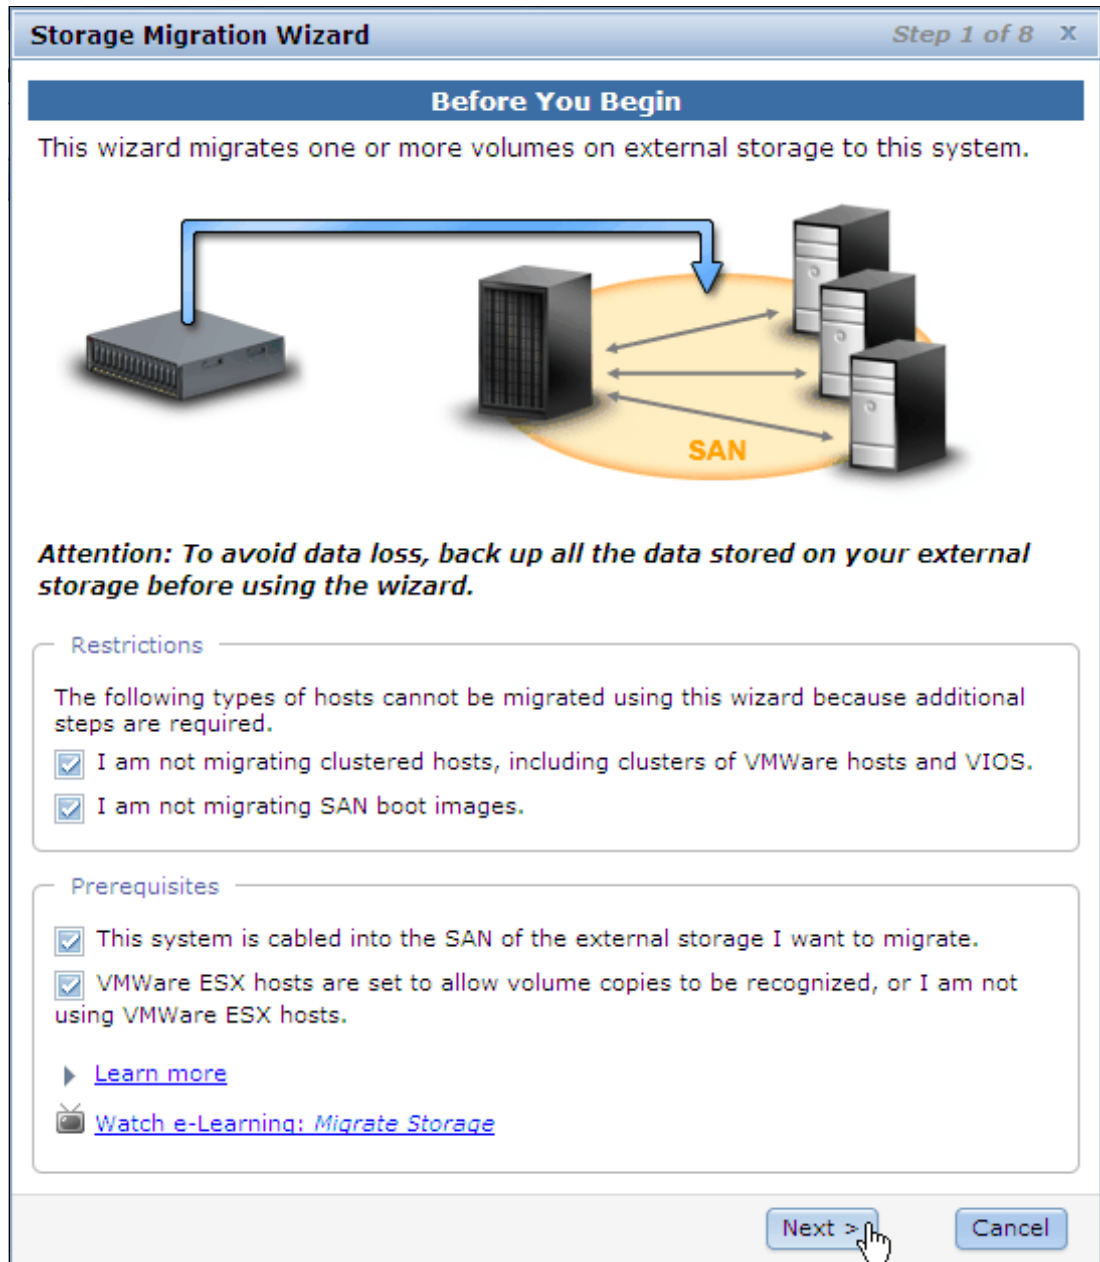


Figure 3-70 Storage migration wizard step 1

3.4.4 Volumes menu

The Volumes menu provides access to the Volumes, Volumes by Pool, and Volumes by host functions. Figure 3-71 shows selection of the volumes menu.

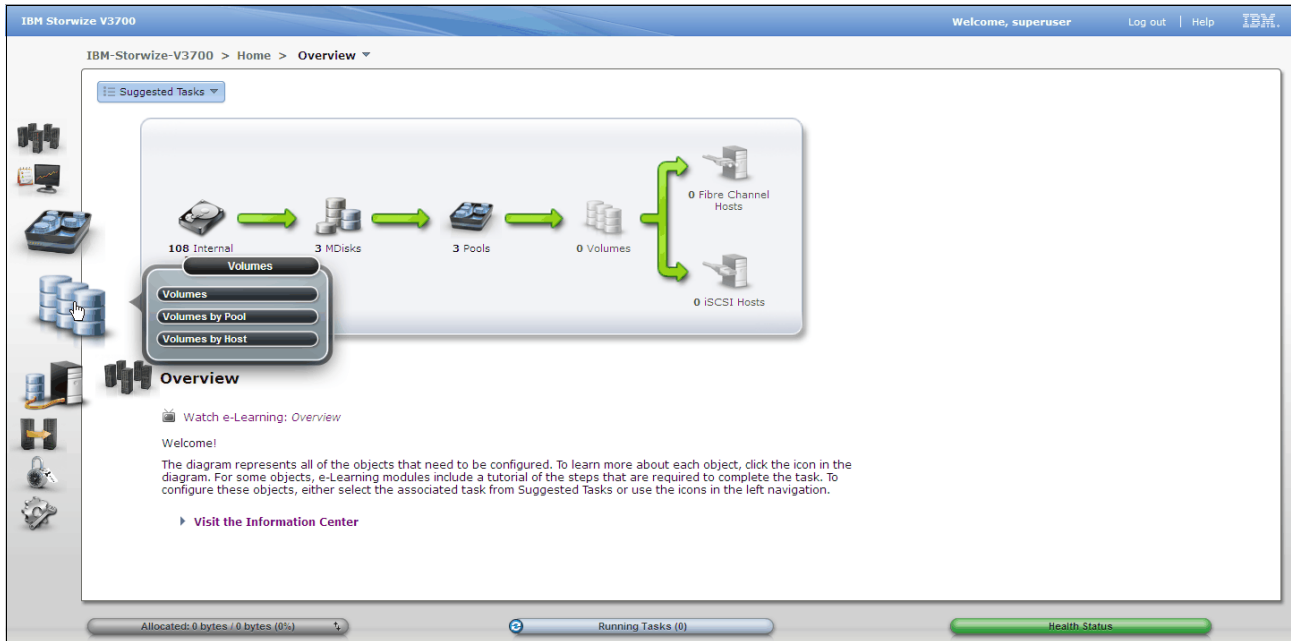


Figure 3-71 Selection of the volumes menu

Volumes panel

Select **Volumes** in the Volumes menu to reveal the panel. The Volumes panel shows all of the volumes in the system. The information displayed is dependant on the columns selected. Figure 3-72 shows the Volumes panel.

Volume actions

Volume actions such as: map, unmap, rename, shrink, expand, migrate to another pool, delete and mirror can be performed from this panel. Figure 3-72 shows the volumes panel.




Figure 3-72 The volumes panel

Create new volumes

Click **New Volume** to reveal the new volume panel. The New Volume panel allows the selection of a preset when creating a new volume. The presets are designed to accommodate the majority of cases. The presets are generic, thin-provisioned, mirror, or thin mirror. Once a preset has been determined, then select the storage pool from which the volumes will be allocated, an area to name and size the volumes is then immediately revealed. Figure 3-73 shows the New Volume panel.

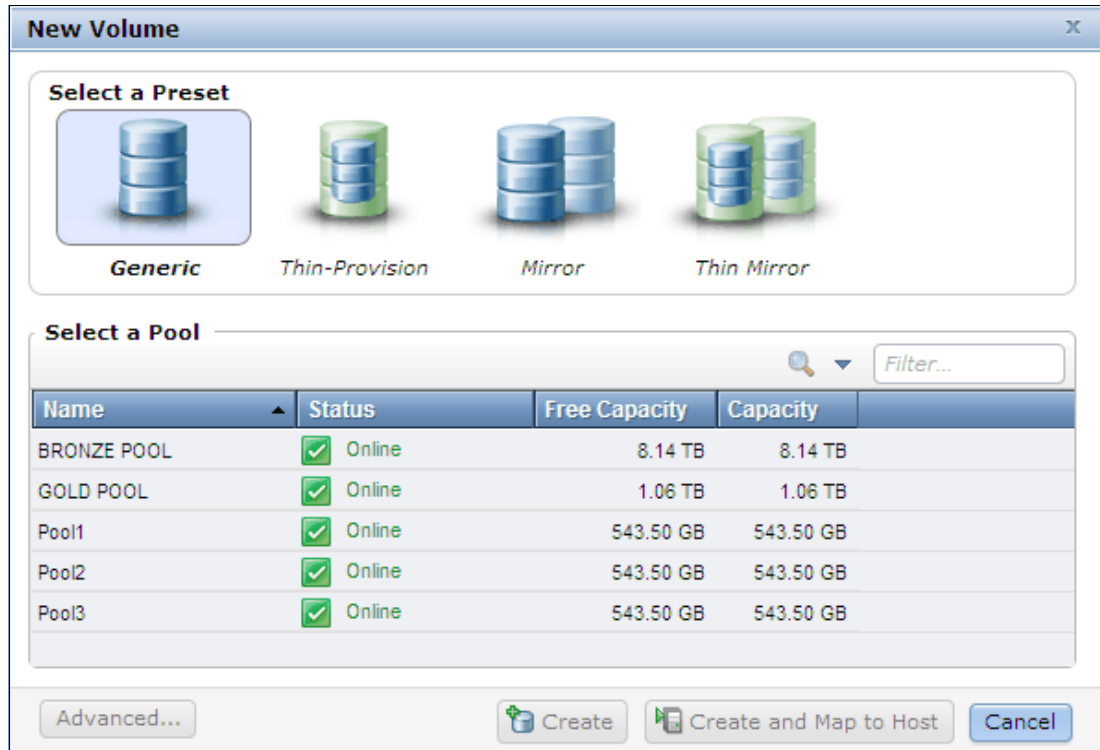


Figure 3-73 New Volume panel

Creating multiple volumes

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

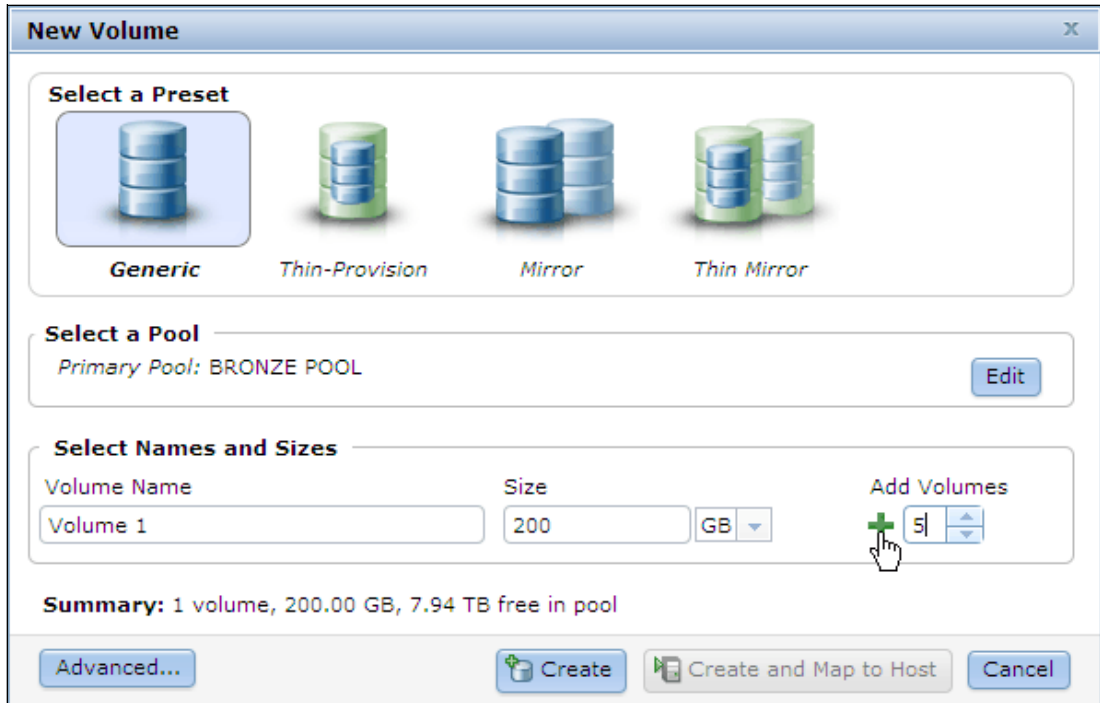


Figure 3-74 Creating multiple volumes quickly

When adding volumes, any number value in the volume name will automatically increment to ensure that each volume is unique. To remove a volume from the list simply click **X**.

Volume advanced settings

Click **Advanced** to reveal additional volume configuration options, use this feature when the preset does not meet requirements. Once the advanced settings have been configured, click **OK** to return to the New Volumes panel. Figure 3-75 shows the Advanced Settings panel.

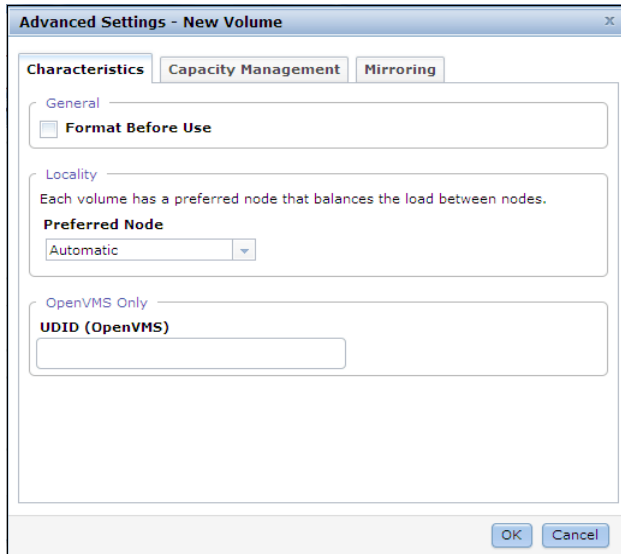


Figure 3-75 New volume advanced settings

Capacity Summary

The New Volumes panel also has a summary which shows the real capacity that will be used if the proposed volumes are created. Once all volumes in the list meet the volumes requirements, click **Create** or **Create and Map to Host** to continue. Figure 3-76 on page 103 shows the New Volume panel and the result of adding 5 volumes by using the add volumes feature.




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

Volumes by pool panel

This is an alternate method to access the volumes by pool panel. For further information see "Volumes by pools panel" on page 91.

Volumes by host panel

Select **Volumes by Host** in the Volumes menu to reveal the panel. The Volume by Hosts panel provides the ability to focus volumes that have been allocated to a particular host by using the host selection filter.

Volume actions

Once a host has been selected all of the volume **Actions** such as: map, unmap, rename, shrink, expand, migrate to another pool, delete and mirror can be performed. Figure 3-77 on page 104 shows the volume by hosts panel.




Figure 3-77

3.4.5 Hosts menu

The Hosts menu provides access to the Hosts, Ports by Host, Host Mappings, Volumes by Host functions. Figure 3-78 shows selection of the Hosts menu.

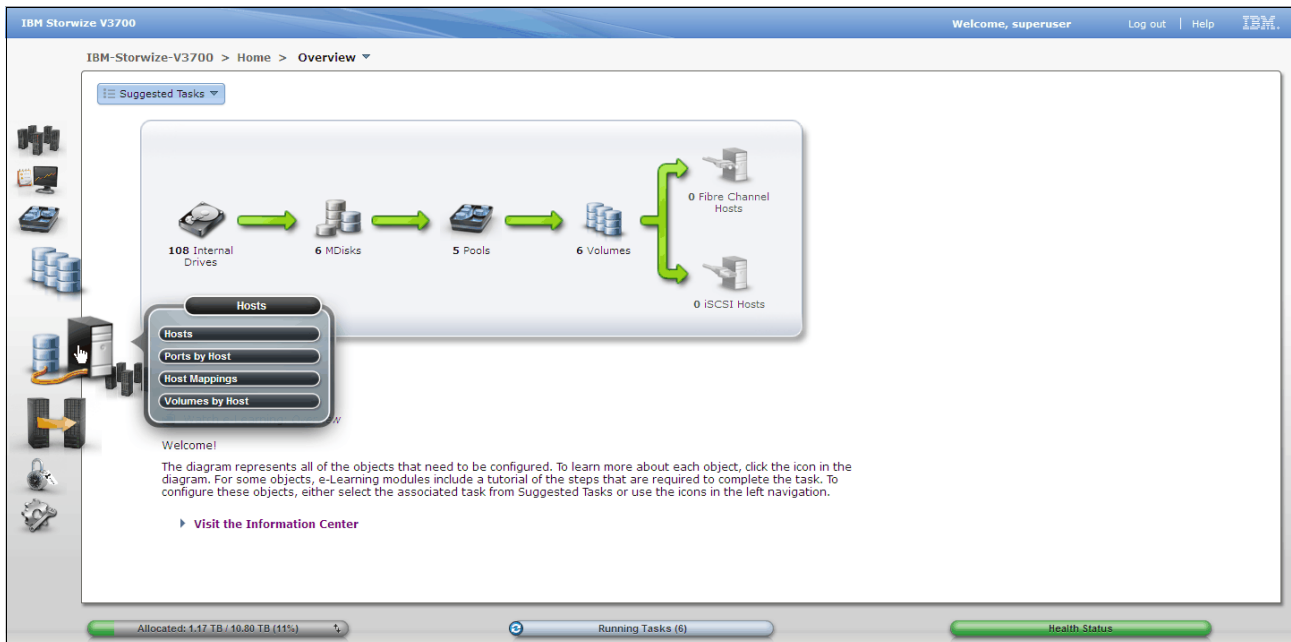


Figure 3-78 Selection of the hosts menu

Hosts panel

Select the **Hosts** item in the Hosts menu to reveal the panel. The Hosts panel shows all of the hosts that have been defined in the system. Figure 3-79 on page 105 shows the Hosts panel.

Host Actions

Host **Actions** such as: map, unmap, rename, and delete can be performed from the Hosts panel. Figure 3-79 shows the actions available from the Hosts panel.




Figure 3-79 The hosts panel

Create host

Click **New Host** and the Create Host panel is revealed. Choose the host type from either Fibre Channel (FC) or iSCSI host and the applicable host configuration panel will be display. Once the host type has been determined the host name and port definitions can be configured. Figure 3-80 shows the choose host type panel.

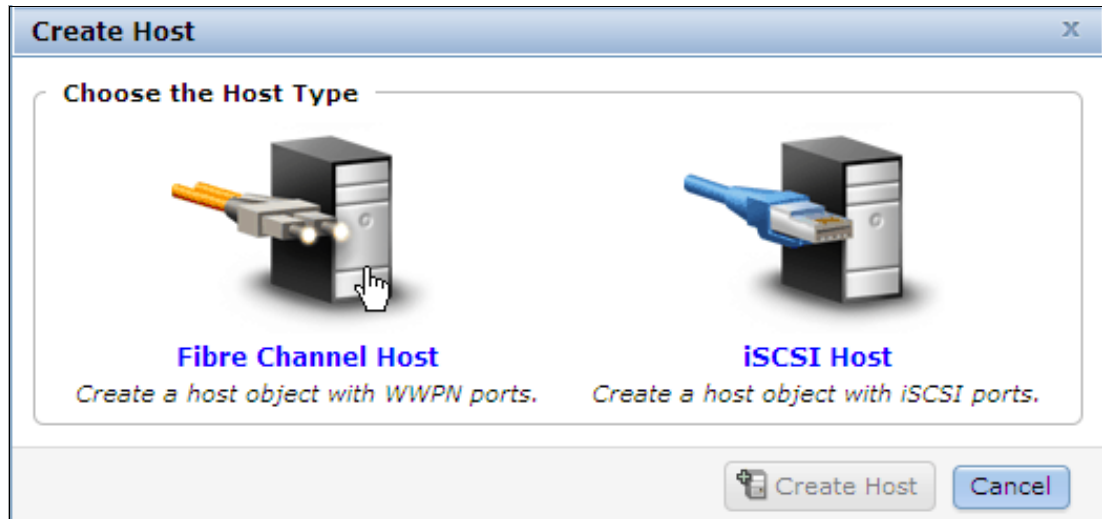


Figure 3-80 Choose the host type panel

Create FC host

Enter a name for the host in the field provided. The FC host port definitions are automatically discovered after the activation of the hosts zoning. The discovered port definitions can be

selected and added from the **Add Port to List** drop down box. Select the port definition from the list and click **Add Port to List** to assign the port to the host definition. If the drop down box has “No candidate HBA ports were found”, or an expected port definition has not been automatically discovered and displayed, check zoning and try **Rescan**.

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

This panel also allows for the manual entry of the port definitions. To do this enter the WWPN in the field provided, it is not case sensitive, and will also accept the bytes separated by colon, however do not put a space character between each byte. Ensure that the port definition entered contains exactly 16 characters and click **Add Port to List** to assign the port to the host definition. This function is useful when the zoning is yet to be activated and the WWPNs are known. In this case the port will have a question mark symbol and “(unverified)” adjacent to the listed port definition. Once the zoning is completed the hosts port status can be seen through the Ports by Host panel, see “Ports by host panel” on page 109. **Advanced Settings** should be checked to allow for the configuration of HP/UX, OpenVMS, and TPGS hosts. At least 1 port definition must be assigned to make **Create Host** become active. Click **Create Host** and the host creation task will run. Figure 3-81 on page 106 shows the FC create host panel.

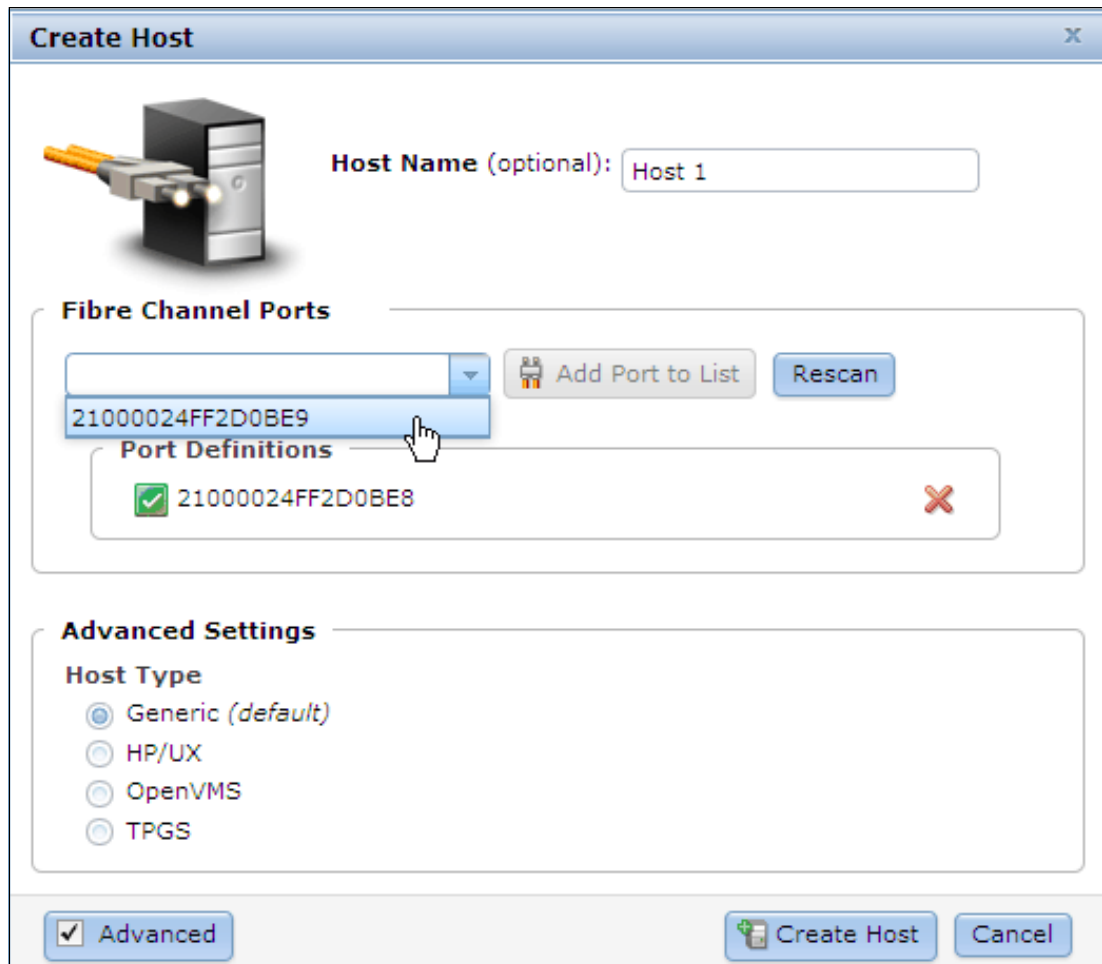


Figure 3-81 FC create host panel with advanced settings checked

Create iSCSI host

Enter a name for the host in the field provided. This panel allows for the entry of the host port definitions. Enter the iSCSI IQN in the field provided and click **Add Port to List** to assign the port to the host definition. To remove a port definition from the list simply click the red **X**.

Advanced Settings should be checked to allow for the configuration of HP/UX, OpenVMS, and TPGS hosts. **Use CHAP authentication** should be checked if it is required, however, to help in problem determination, this step can be delayed until after the first one or two hosts have been configured and their connectivity has been tested without authentication configured. At least 1 port definition must be assigned to make **Create Host** become active. Click **Create Host** and the host creation task will run. Figure 3-82 on page 107 shows the iSCSI create host panel.

Figure 3-82 The iSCSI create host panel with advanced settings checked

Modify host mappings panel

Host mapping is a method of controlling which hosts can access the volumes. A host can only access volumes that have been mapped to it, the Modify Host Mappings panel is used to configure the host mappings. Right click on a host listed in the Hosts panel and select **Modify Mappings** to reveal the modify hosts mappings panel. Figure 3-83 on page 108 shows the method for revealing the Modify Host Mappings panel.




Figure 3-83 Method for revealing the modify hosts panel

Mapping volumes

The Modify Host Mappings panel allows for unmapped volumes to be mapped to the selected host. Verify the correct host has been selected in the top left corner of the panel. To map a volume, select it from the “Unmapped Volumes” section on the left hand side of the panel, use the green **>** to “push” the volume to “Volumes Mapped to the Host” section on the right hand side of the panel. The volume will be highlighted in yellow, at this time the volume is not mapped to the host and hence not accessible by the host. Take this opportunity to **Edit SCSI ID** if it is required. The SCSI ID controls the order in which the volumes are presented to the host. Some host software requires the SCSI IDs to be in sequence. Once the volumes highlighted in yellow are as required then click **Apply** to complete the mapping operation. The mapped volumes will no longer be highlighted in yellow and will now be accessible by the host. Continue to select other hosts and perform other host mappings as required. Once all mapping operations are complete, click **Cancel** to close the Modify Host Mappings panel. **Map Volumes** can also be used, it is best to use this option if only a single mapping operation is required, for it will perform the apply and then close the panel immediately. Figure 3-84 shows the Modify Host Mappings panel.

Unmapping volumes

The Modify Host Mappings panel also allows volumes to be unmapped from hosts. Stop I/O to the volume and verify the correct host has been selected in the top left corner of the panel. To unmap a volume, select it from “Volumes Mapped to the Host” section on the right hand side of the panel, use the green **<** to “pull” the volume back to the “Unmapped Volumes” section on the left hand side, click **Apply** to complete the unmapping operation. Once all unmapping operations are complete, click **Cancel** to close the Modify Host Mappings panel. Figure 3-84 shows the Modify Host Mappings panel.

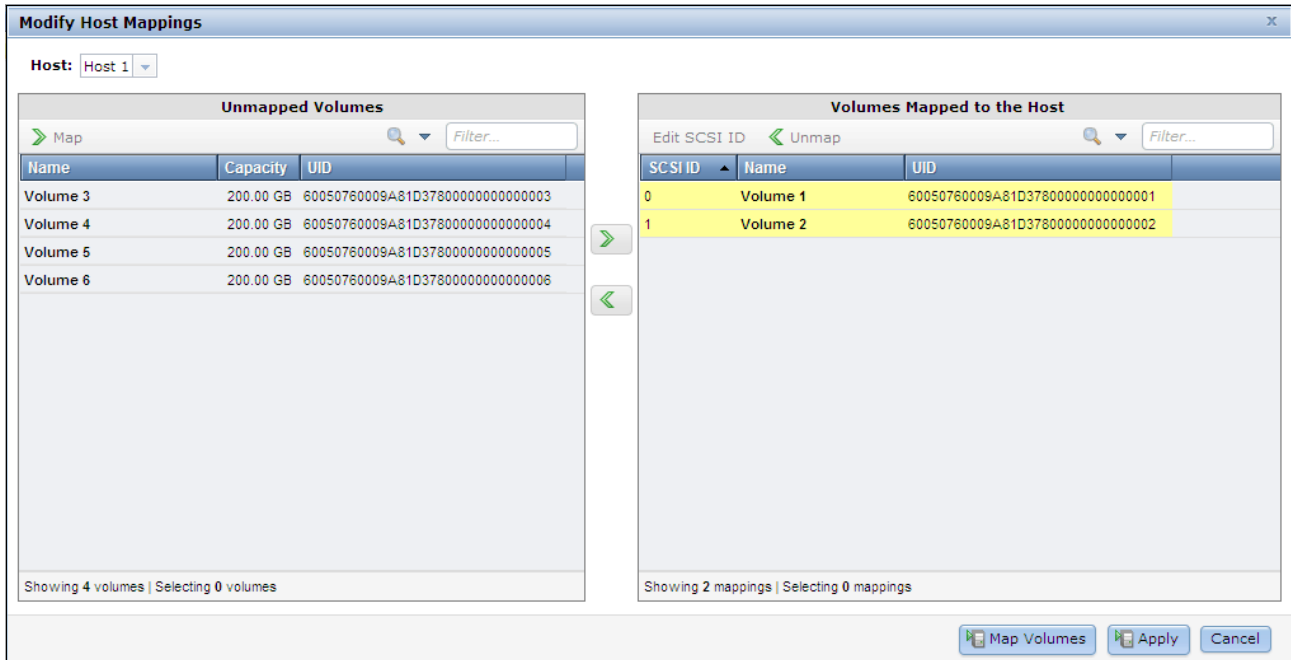


Figure 3-84 Modify host mappings panel

Ports by host panel

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

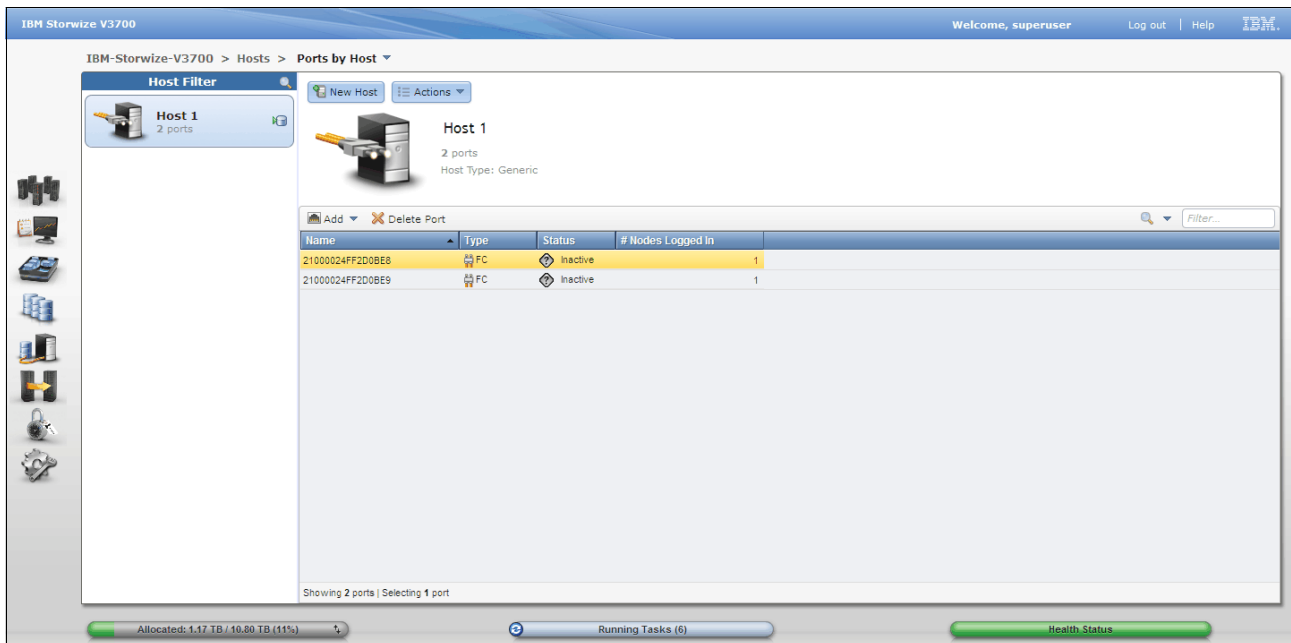


Figure 3-85

Host mappings panel

Select **Host Mappings** in the Hosts menu to reveal the panel. This panel shows the volumes each host can access with the corresponding SCSI ID. **Unmap Volume** action can be performed from this panel. Figure 3-86 on page 110 shows the Host Mappings panel.




Figure 3-86 Host mappings panel

Volumes by host panel

This is an alternate method to access the volumes by host panel. For further information see “Volumes by host panel” on page 103.

3.4.6 Copy Services menu

The Copy Services menu provides access to the FlashCopy, Consistency Groups, and FlashCopy Mappings functions. Figure 3-87 shows the Copy Services menu.

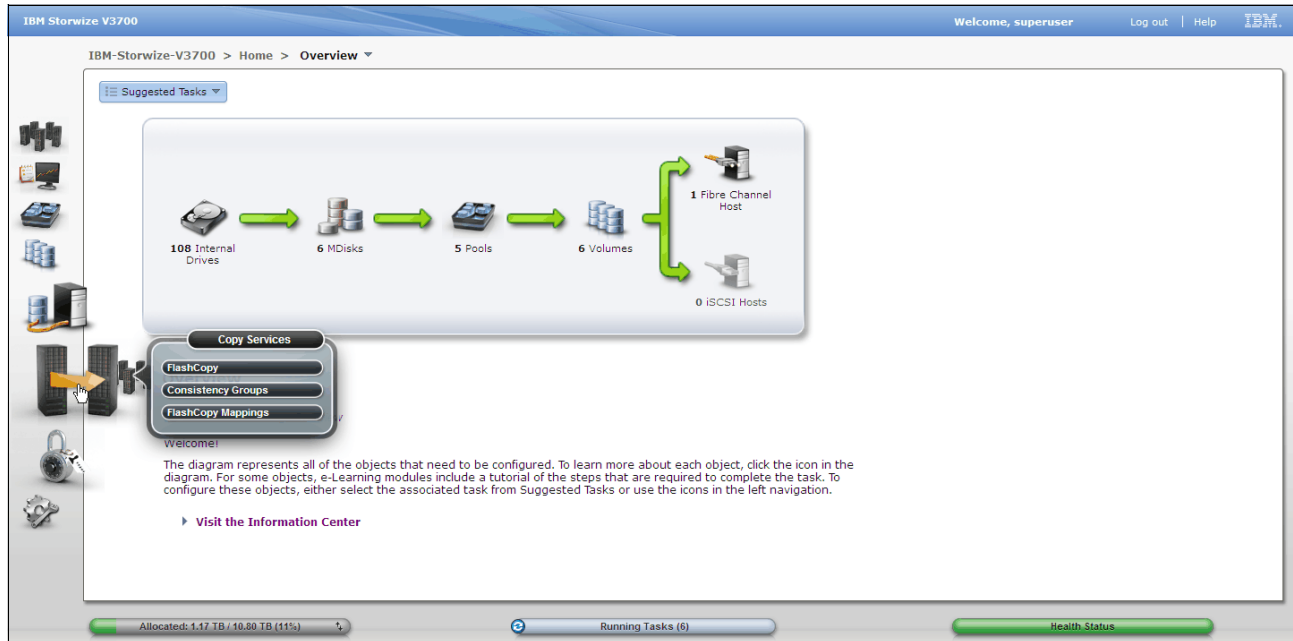


Figure 3-87 Copy services menu

FlashCopy panel

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

FlashCopy actions

FlashCopy **Actions** such as: snapshot, clone, backup, target assignment, and deletion can be performed from this panel. Figure 3-88 shows the actions available from the FlashCopy panel.




Figure 3-88 FlashCopy panel

Consistency groups panel

Select **Consistency Groups** in the Copy Services menu to reveal 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 a simultaneous point in time for the group instead of co-coordinating actions at an individual basis. This can help ensure the groups target volumes are consistent to exactly the same point in time, and remove a number of 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** can also be selected from this panel, this panel is explained in “FlashCopy mappings panel”. Figure 3-89 shows the Consistency Group panel.




Figure 3-89 Consistency groups panel

FlashCopy mappings panel

Select **FlashCopy Mappings** in the Copy Services menu to reveal the panel. FlashCopy mappings define the relationship between a source volumes and a target volumes. The FlashCopy Mappings panel shows information relating to each mapping such as: status, progress, source and target volumes, and flash time. Select **New FlashCopy Mapping** to configure a new mapping. Figure 3-90 on page 113 shows the FlashCopy Mappings panel.




Figure 3-90 FlashCopy mappings panel

Create new flashcopy mapping

To create a new FlashCopy mapping enter the select the source and target volumes from the respective drop down boxes and click **Add**. The mapping definition will be added to the list. To remove a mapping definition from the list simply click the red **X**. Click **Next** to continue. Figure 3-91 shows step 1 of the New FlashCopy Mapping wizard.

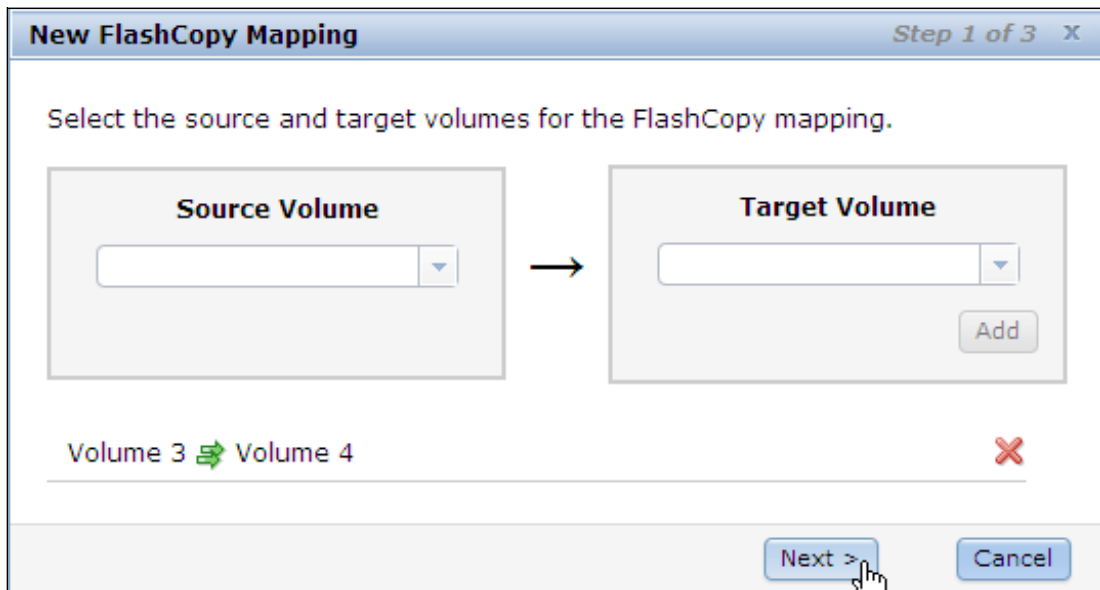


Figure 3-91 Step 1 of the new flashcopy mapping wizard

Select a FlashCopy preset and modify the advance settings as required. Select **Next** to continue. Figure 3-92 on page 114 shows step 2 of the New FlashCopy Mapping wizard.

Note: For detailed information regarding FlashCopy Mappings and Consistency Groups refer to the following section of the IBM Storwize V3700 Information Center.

Product overview → **Technical overview** → **Copy Services features**

http://pic.dhe.ibm.com/infocenter/storwize/v3700_ic/index.jsp

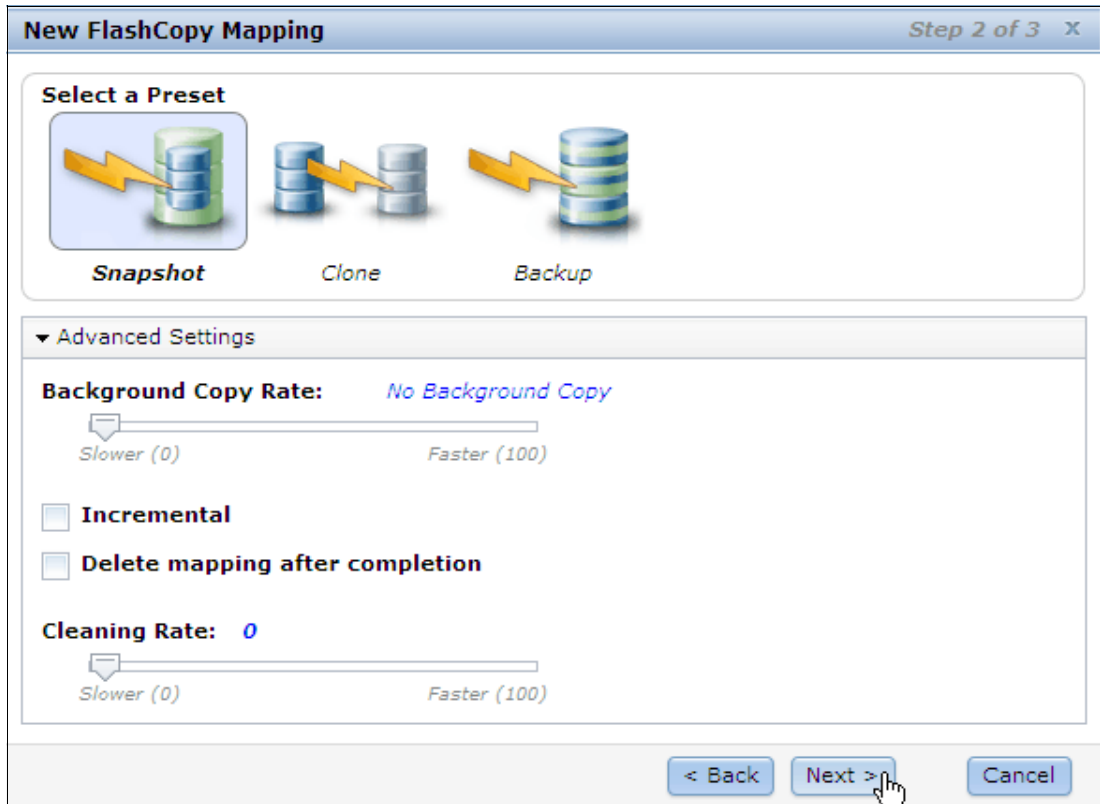


Figure 3-92 Step 2 of the new flashcopy mapping wizard

The newly created FlashCopy mapping can be added to a consistency group. Choose if this mapping will be grouped and click **Finish**. Figure 3-93 shows step 3 of the new Flashcopy Mapping wizard.

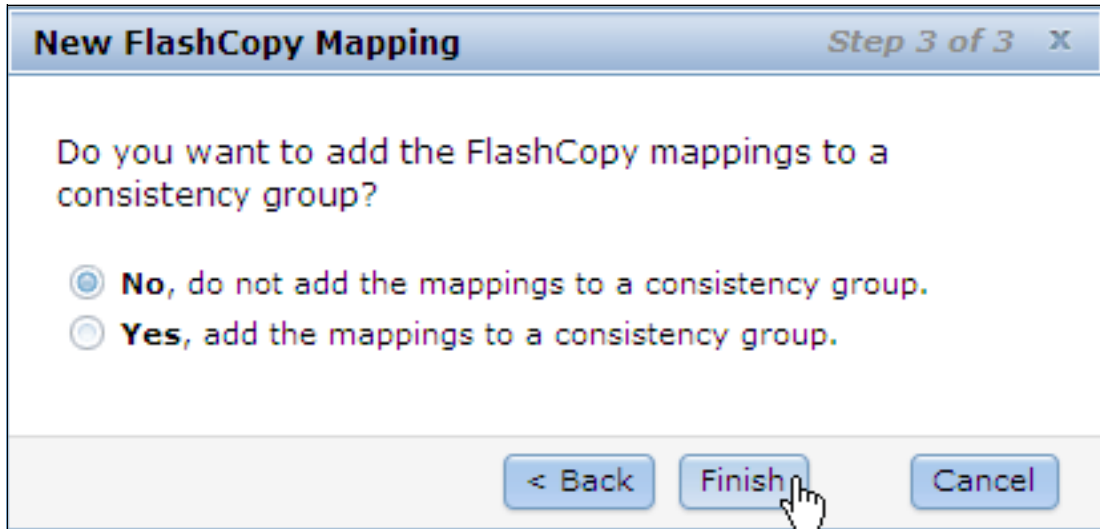


Figure 3-93 Step 3 of the new flashcopy mapping wizard

3.4.7 Access menu

The Access menu provides access to the Users, and Audit Log functions. Figure 3-94 on page 115 shows the Access menu.




Figure 3-94 Access menu

Users panel

Select **Users** in the Access menu to reveal the panel. The Users panel shows the defined user groups and users for the system. The users listed can be filtered by user group. Click **New User Group** to reveal the create a new group panel. Figure 3-95 shows the Users panel.




Figure 3-95 Users panel

Create new user group

The New User Group panel provides a means for user group configuration. Enter the group name and select the role then click **Create** to make a new group. Figure 3-96 shows the New User Group panel.

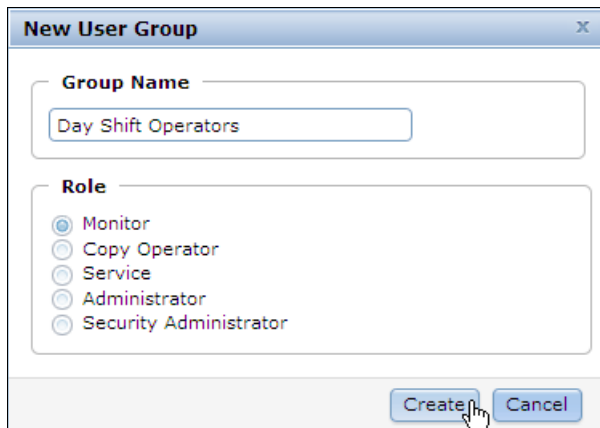


Figure 3-96 New user group panel

Create new user

Click **New User** to define a user to the system. Figure 3-97 on page 117 shows the Users panel and the New User button.

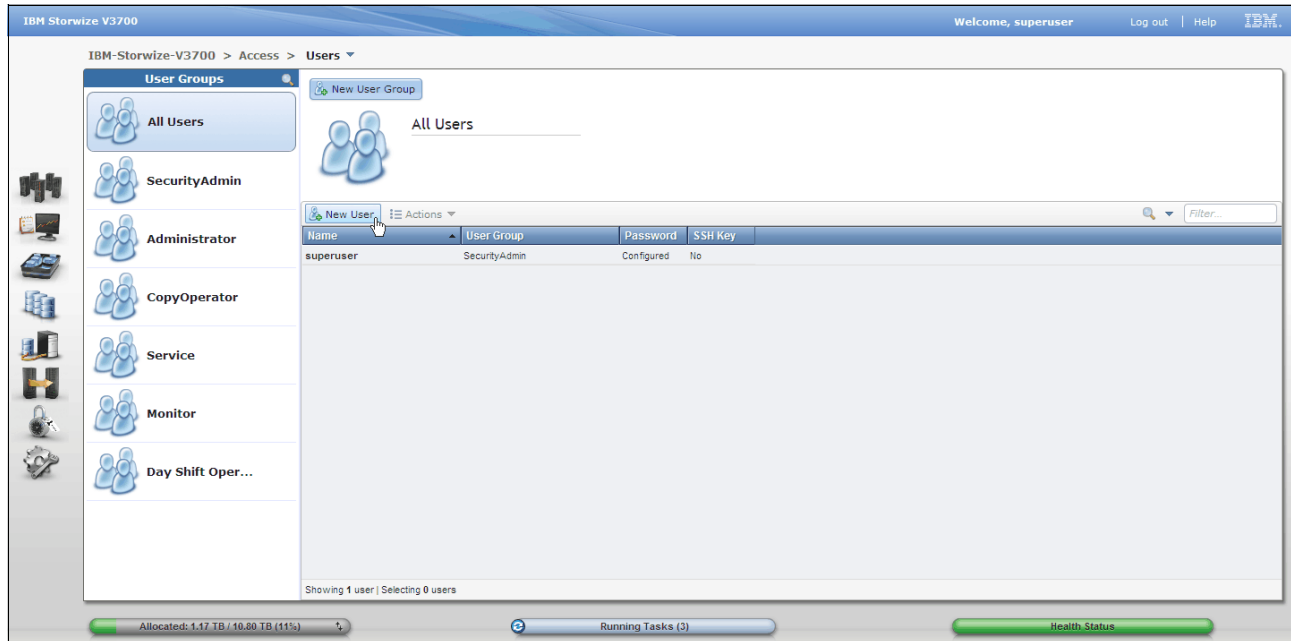


Figure 3-97 Users panel and the new user button

The New User panel provides a means for user name, password and authentication mode configuration. It is essential to enter the user name, password, and group, and authentication mode. The public secure shell (SSH) key is optional. Once the user has been defined click **Create**. The authentication mode can be either set to local or remote. Select local if the Storwize V3700 will perform the authentication locally, select remote if a remote service such as an LDAP server will authenticate the connection. If remote is selected then the remote authentication server will need to be configured in the Storwize V3700, go to the **Settings menu** → **Directory Services panel**. The SSH configuration can be used to establish a more secure connection to the command line interface, it is important to also enter the associated SSH private key into terminal session to make this connection work. Figure 3-98 on page 118 shows the New User panel.

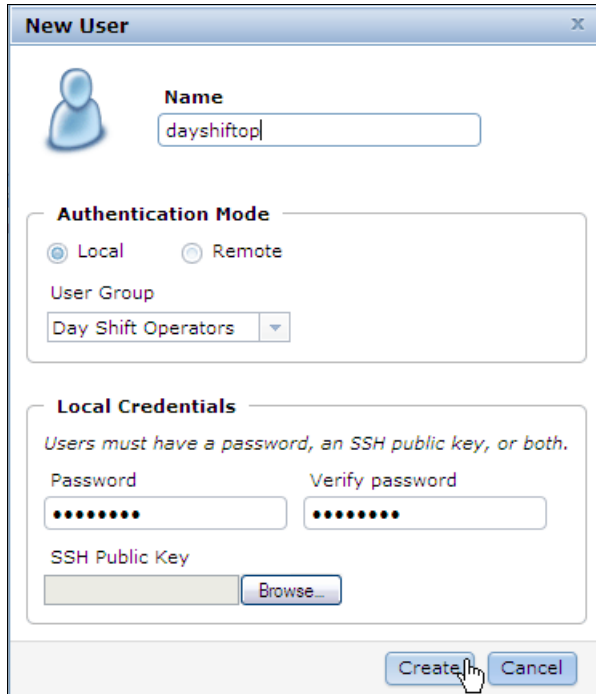


Figure 3-98 New user panel

Audit log panel

Select **Audit Log** in the Access menu to reveal the panel. The audit log keeps track of action commands that are issued through a CLI session or through the management GUI. The Audit Log panel displays information about the actionable command such as: the user, the time stamp and any associated command parameters. The log can be filtered to reduce the number of items listed. Figure 3-99 on page 118 shows the Audit Log panel.

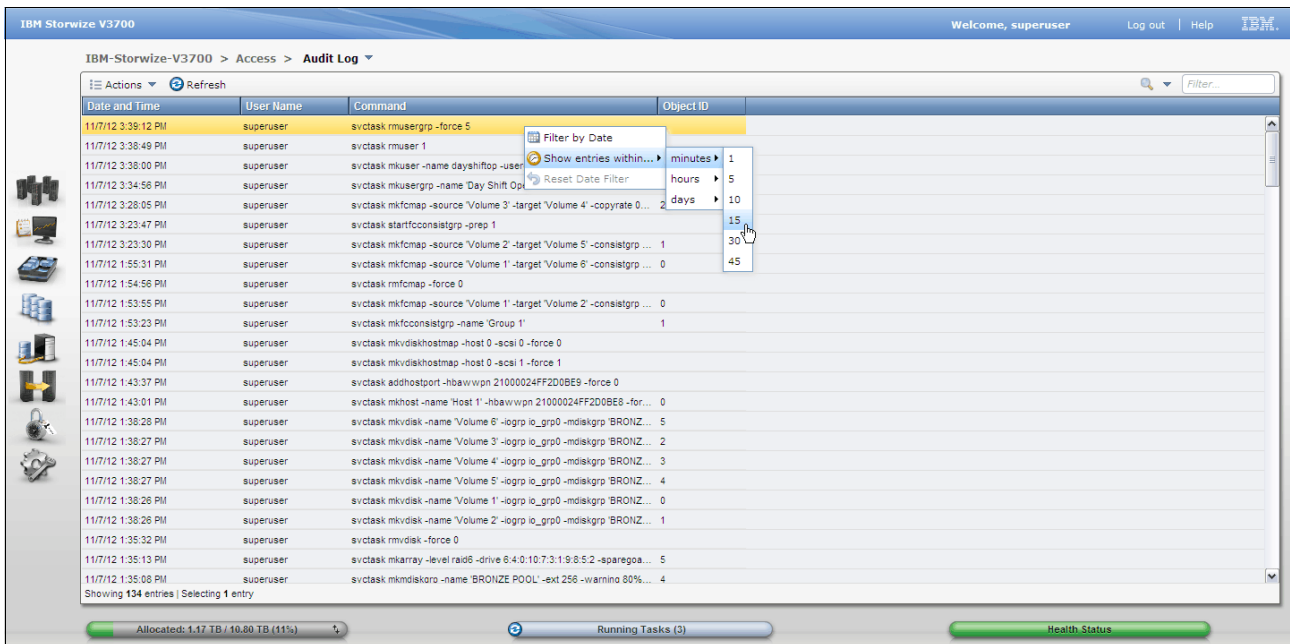


Figure 3-99 Audit log panel

3.4.8 Settings menu

The Setting menu provides access to the Event Notifications, Directory Services, Network, Support, and General functions. Figure 3-100 on page 119 shows the Settings menu.

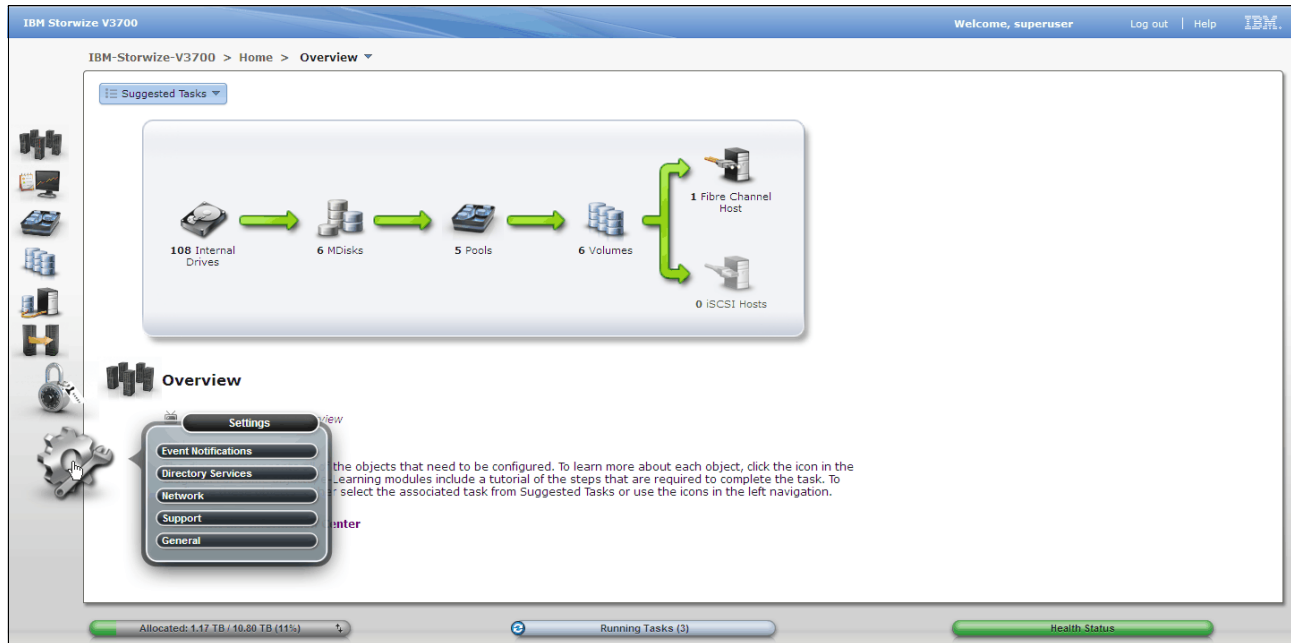


Figure 3-100 Settings menu

Event notifications panel

Select **Event Notifications** in the Settings menu to reveal the panel. The Storwize V3700 can use Simple Network Management Protocol (SNMP) traps, syslog messages, emails and IBM Call Homes to notify parties when events are detected. Each event notification method can be configured to report all events or alerts. Alerts are the significant events and may 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 will be enabled when Email event notification is enabled by following the Call Home wizard.

Enable email event notification

Click **Enable Email Event Notification** to reveal the Call Home wizard. Figure 3-101 on page 120 shows the Event Notifications Email configuration panel.

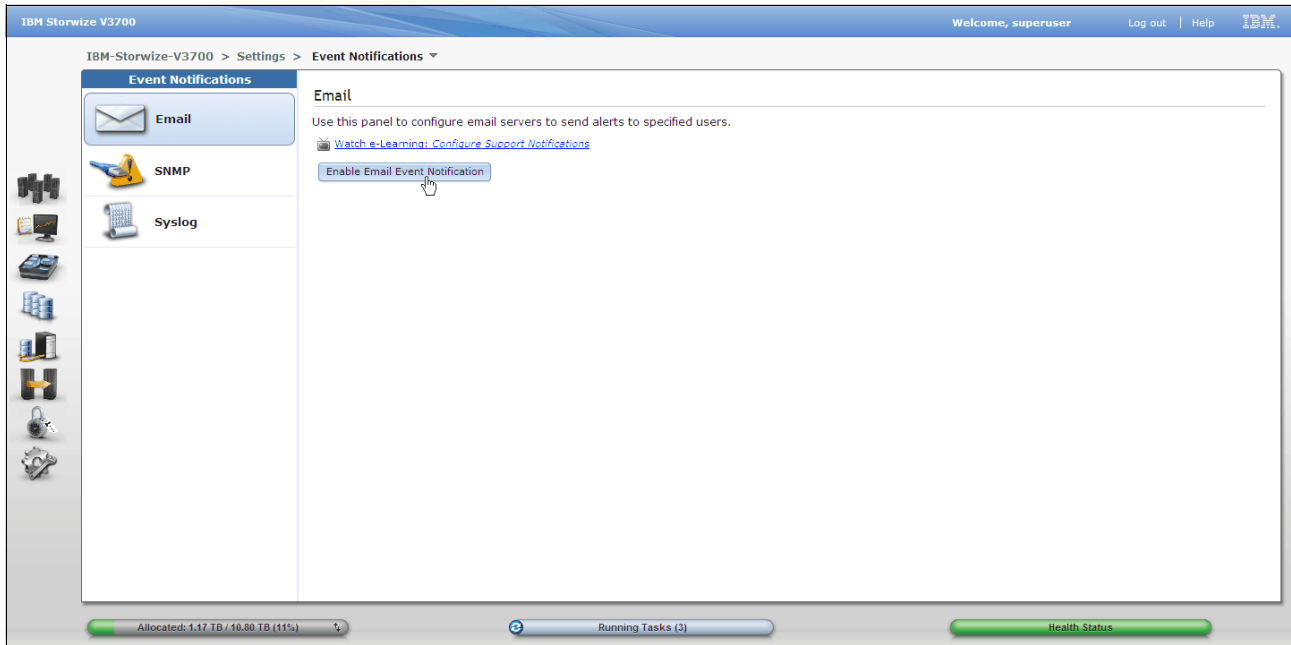


Figure 3-101 Event notification panel

Call home wizard

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

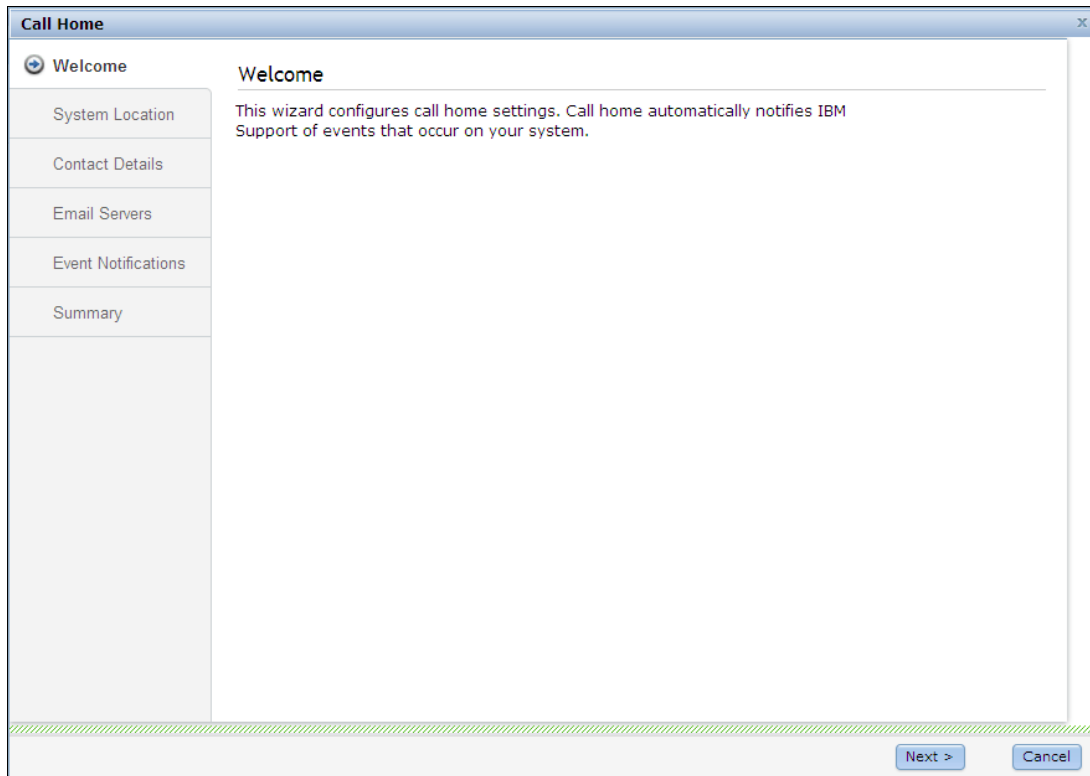


Figure 3-102 Call home wizard

SNMP event notification

The Event Notifications panel provides access to SNMP configuration panel. Click **SNMP** to reveal the panel, then enter the relevant server details. Multiple servers can be configured, click **+** to add more servers. Figure 3-103 on page 121 shows the SNMP configuration panel.

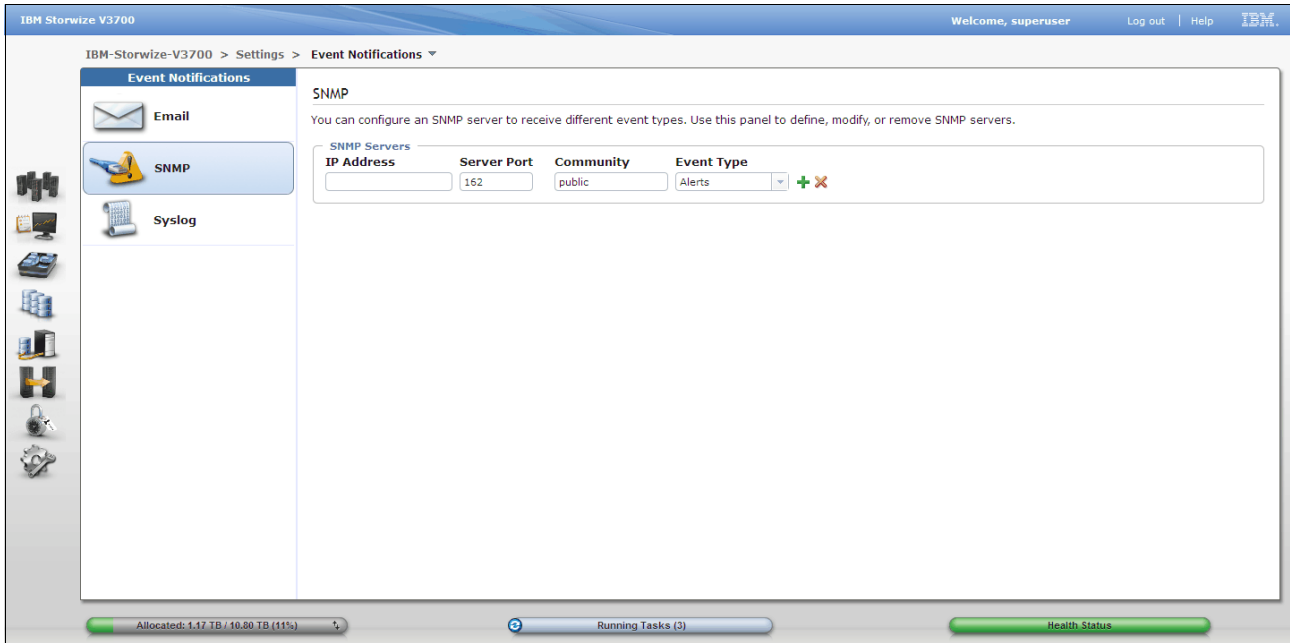


Figure 3-103 SNMP configuration panel

Syslog event notification

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

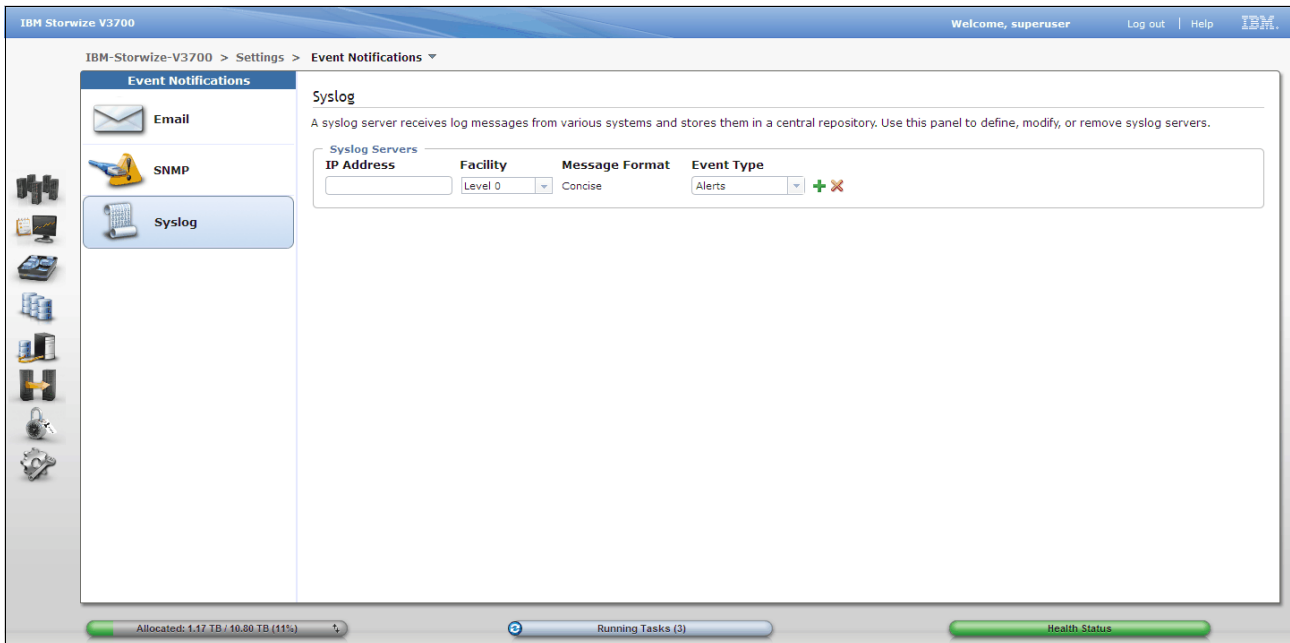


Figure 3-104 Syslog configuration panel

Directory services panel

Select **Directory Services** in the Settings menu to reveal the panel. The Directory Services panel provides access to the Remote Authentication wizard. Remote authentication must be configured to allow remote users to be created on the Storwize v3700. A remote user is authenticated on a remote service such as IBM Tivoli® Integrated Portal or a Lightweight Directory Access Protocol (LDAP) provider.

Enable remote authentication

Click **Configure Remote Authentication** to reveal the wizard. Figure 3-105 on page 122 shows the Directory Services panel.

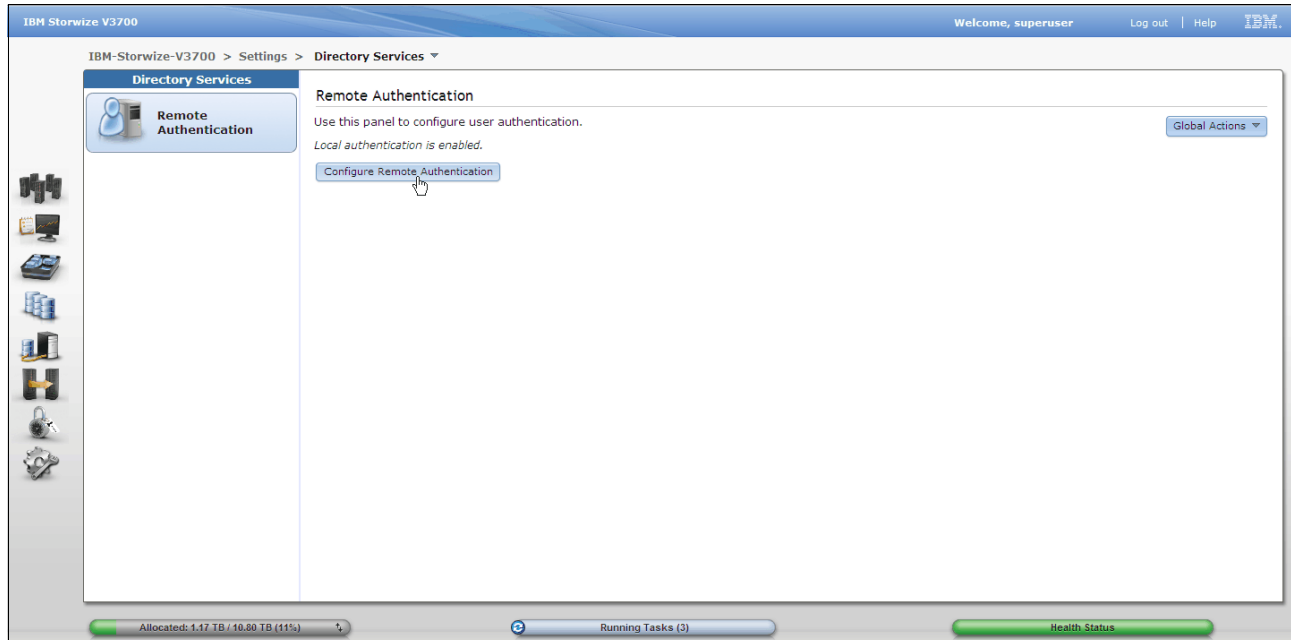


Figure 3-105 Directory services panel

Network panel

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

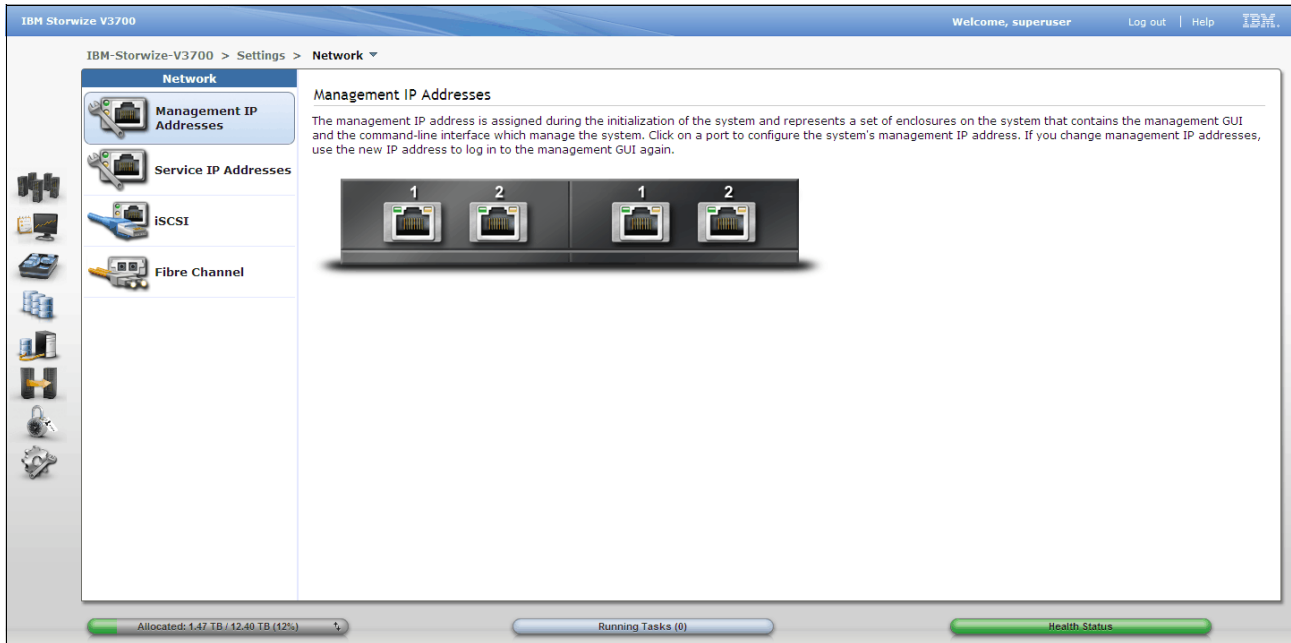


Figure 3-106 Network panel

Management IP addresses

The Management IP Address is the IP address of the system and was configured during initial setup, it can be an IPv4 address, an IPv6 address, or both. The Management IP address is logically assigned to Ethernet port 1 of each node canister, allowing for node canister failover. Furthermore, another Management IP Address can be logically assigned to Ethernet port 2 of each node canister for additional 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 desired **Port** to configure, note the corresponding port on the partner node canister is also highlighted. Figure 3-107 on page 123 shows Management IP Addresses configuration panel.

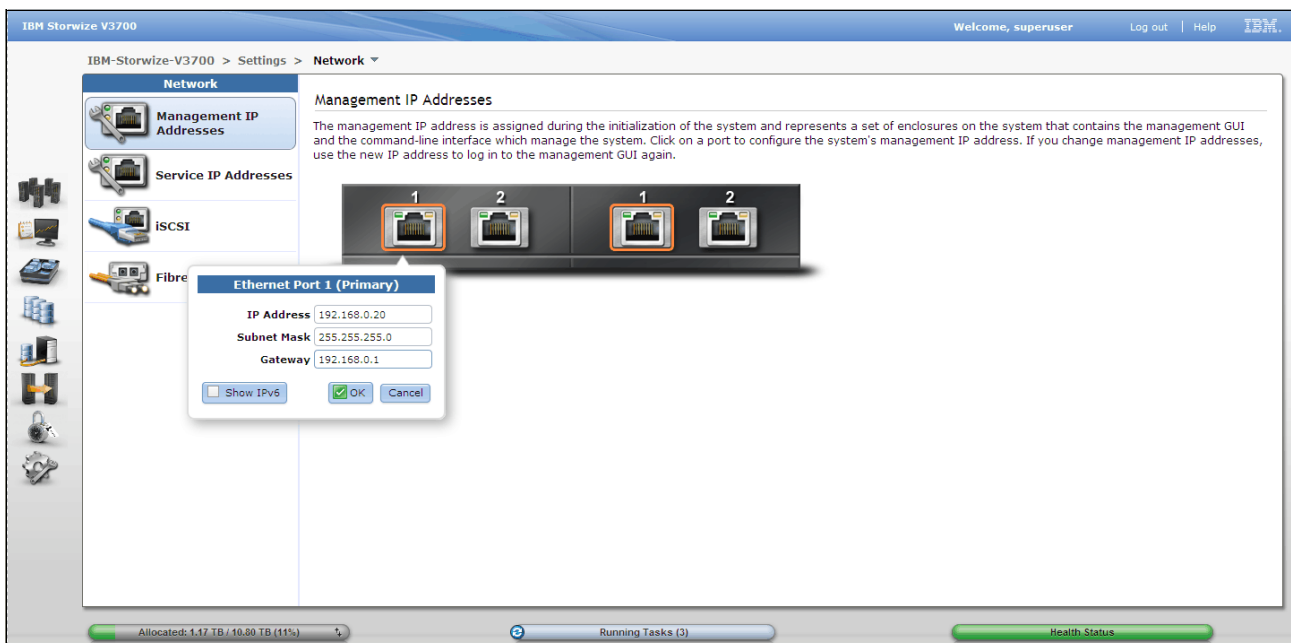


Figure 3-107 Management IP addresses configuration panel

Service IP Addresses

Service IP addresses are used to gain access to the service assistant when a node is not operating correctly as part of the system, it can be an IPv4 address, an IPv6 address, or both. The Service IP addresses are configured on Ethernet port 1 of each node canister. Click **Service IP addresses** and then select the node canister to configure. Figure 3-108 on page 124 shows the Service IP addresses configuration panel.



Figure 3-108 Service IP addresses configuration panel

iSCSI connectivity

The Storwize V3700 supports iSCSI connections for hosts. Click **iSCSI** and select the node canister to configure the iSCSI IP addresses. Figure 3-109 shows the iSCSI configuration panel.




Figure 3-109 The iSCSI configuration panel

Fibre channel connectivity

The Fibre Channel panel displays established fibre channel connections between the Storwize V3700 node canisters and other storage systems, and hosts. Click **Fibre Channel** and select desired view from the **View connectivity for:** drop down box. Figure 3-110 on page 125 shows the Fibre Channel panel.




Figure 3-110 Fibre channel panel

Figure 3-111 shows the results of the selected fibre channel connectivity view.




Figure 3-111 Results of the selected fibre channel connectivity view

Support panel

Select **Support** in the Settings menu to reveal the panel. The Support panel provides access to the IBM support package, which is used by IBM to assist with problem determination. Click **Download Support Package** to access the wizard. Figure 3-112 shows the Support panel.

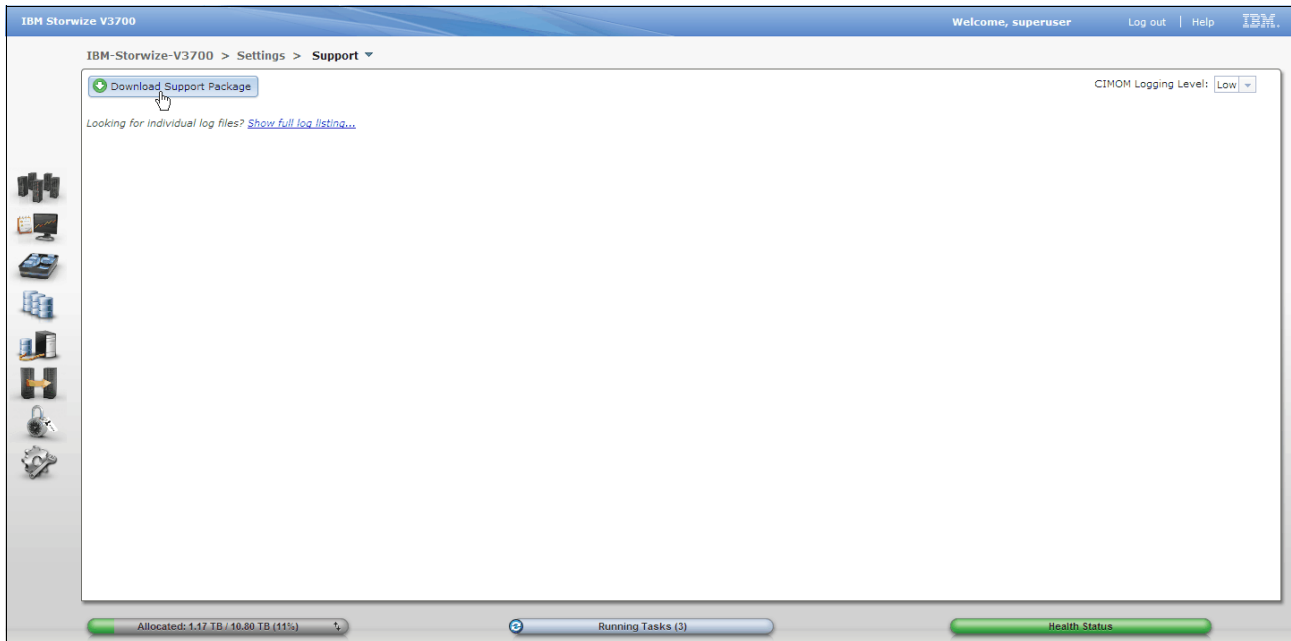


Figure 3-112 Support panel

Download support package wizard

The Download Support Package wizard provides a range of package types. IBM support will provide direction on package type selection as required. To download the package, select

type and click **Download**, the output file can be saved to the user's workstation. Figure 3-113 on page 127 shows the Download Support Package wizard.



Figure 3-113 Download support package wizard

Show full log listing...

The Support panel also provides access to the files currently residing on the node canisters, see Figure 3-112. Click **Show full log listing...** to access the node canister files. To save a file to the user's workstation, select a file, right click and select **Download**. To change to the file listing to show the files residing on a partner node canister, select the node canister from the drop down box adjacent to the panel filter. Figure 3-114 shows the full log listing.




Figure 3-114 Full log listing

General panel

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

Date and time

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




Figure 3-115 General panel

Upgrade machine code

IBM recommends that you use the latest version of machine code, the Upgrade Machine Code panel displays the current machine code level, if the system has connection to the internet it will make contact with the IBM upgrade server to check if the current level is the latest level. If an update is available, a direct link to the code will be provided to make code download process easier. To upgrade the code, the IBM Storwize V3700 Code and the IBM Storwize V3700 Upgrade Test Utility will need to be downloaded. Once the files are downloaded it is best to check the MD5 checksum to ensure the files are sound. Read the release notes, verify compatibility, and follow all IBM recommendations and prerequisites. To upgrade the machine code of the Storwize V3700, click **Launch Upgrade Wizard**. Once the upgrade has started an **Abort** button will be revealed, this can be used to abort the upgrade process. Figure 3-116 on page 129 shows the Upgrade Machine Code panel.

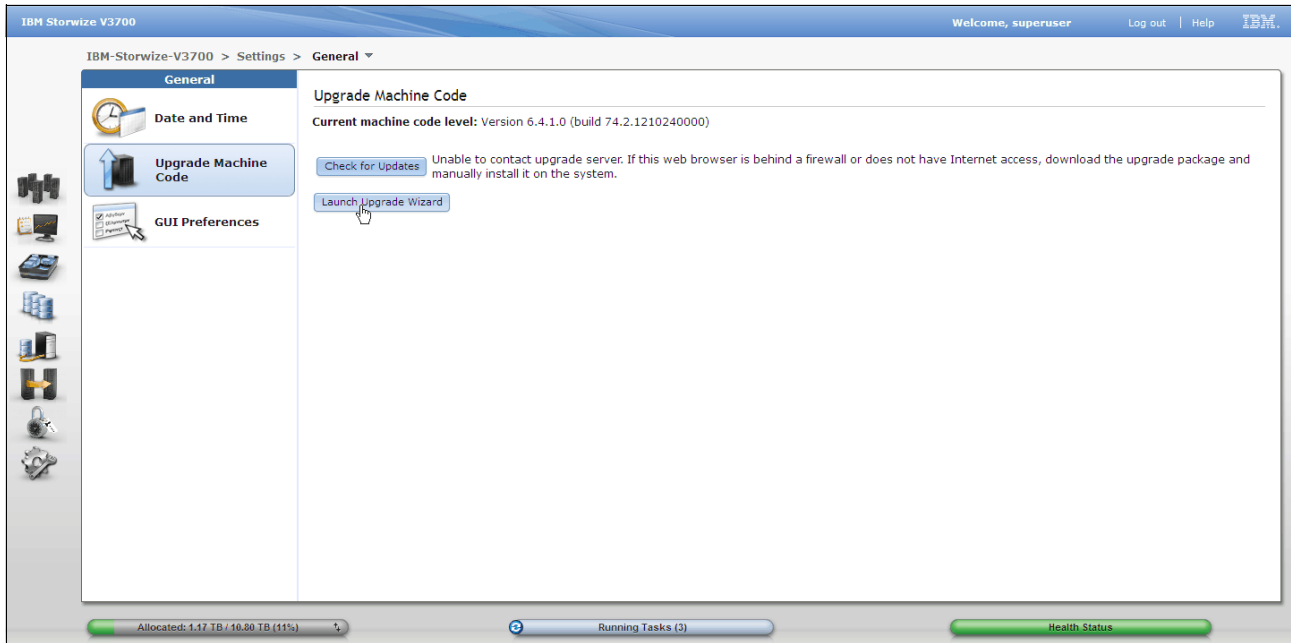


Figure 3-116 Upgrade machine code panel

The first step of the upgrade wizard will guide the user to upload the IBM Storwize V3700 Upgrade Test Utility. Figure 3-117 shows step 1 of the machine code upgrade wizard.

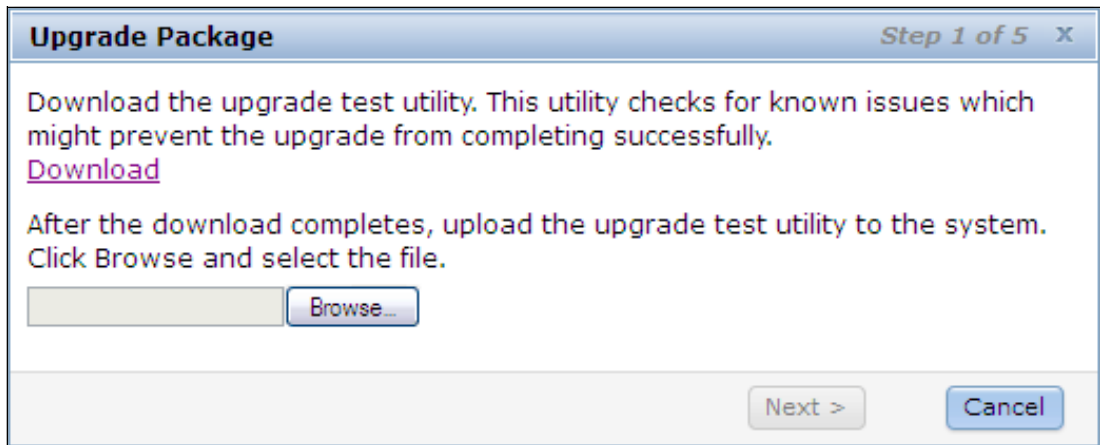


Figure 3-117 Step 1 of the machine code upgrade wizard.

GUI Preferences

The GUI Preferences panel allows the user to refresh GUI objects, restore default browser preferences, set table selection policy, and configure the Information Center web address. Figure 3-118 on page 130 shows the GUI preferences panel.

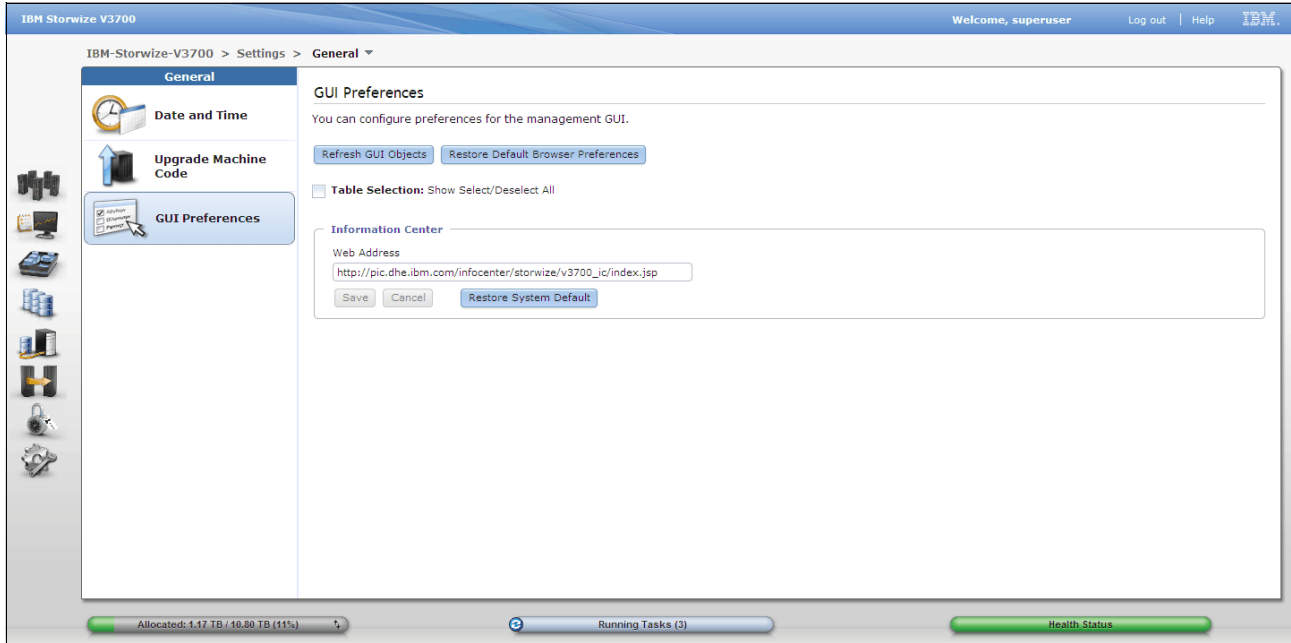


Figure 3-118 GUI preferences panel

3.5 Management GUI help

This section provide information about the various method to get help while using the Storwize V3700 management GUI. The topics covered include the Storwize V3700 Information Center, e-Learning modules, embedded panel help, question mark help, hover help, and IBM endorsed Youtube videos.

3.5.1 Storwize V3700 Information Center

The best source of information for the Storwize V3700 is the Information Center. Click **Visit the Information Center** for direct access to the online version from the Overview panel. Figure 3-119 shows the Overview panel information center link.

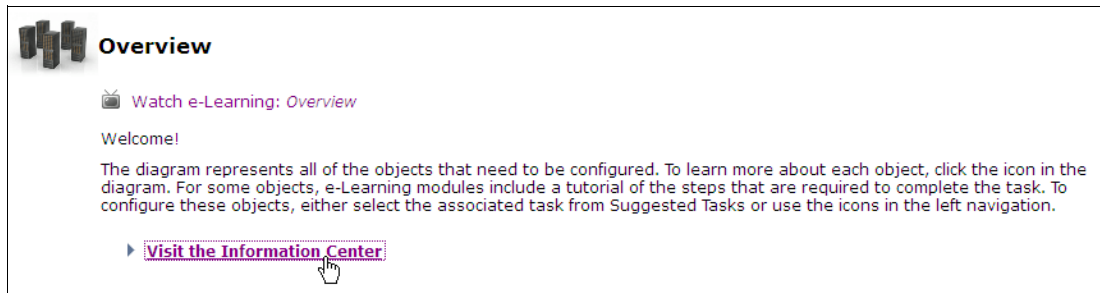


Figure 3-119 Overview panel information center link

3.5.2 Watch e-Learning

The Storwize V3700 provides embedded e-Learning video to watch. The video provide easy to follow directions to complete various tasks. Click **Watch e-Learning:...** to reveal the video. Figure 3-120 show the Watch e-Learning module.

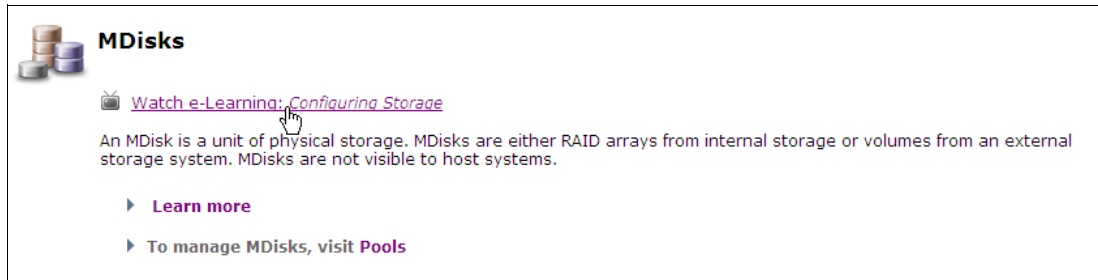


Figure 3-120 Watch e-Learning module

3.5.3 Learn more

The Storwize V3700 provides embedded Learn more links to explain important concepts and panels. Click **Learn more** to reveal the information panel. Figure 3-121 show the Learn more link.

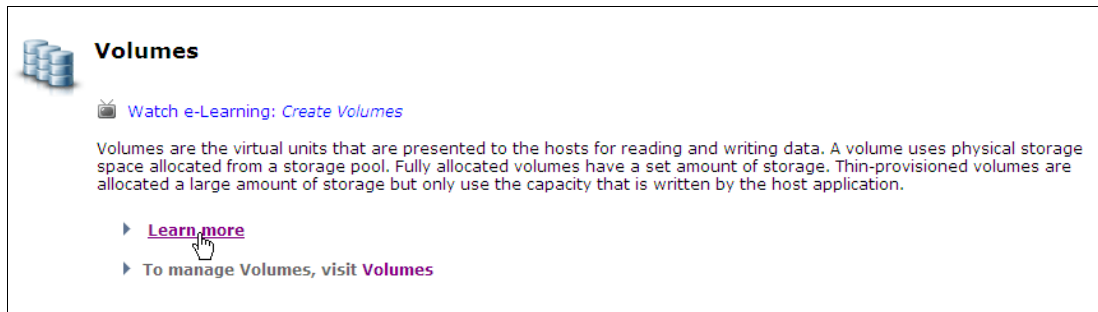


Figure 3-121 Learn more link

Figure 3-122 shows the information panel.

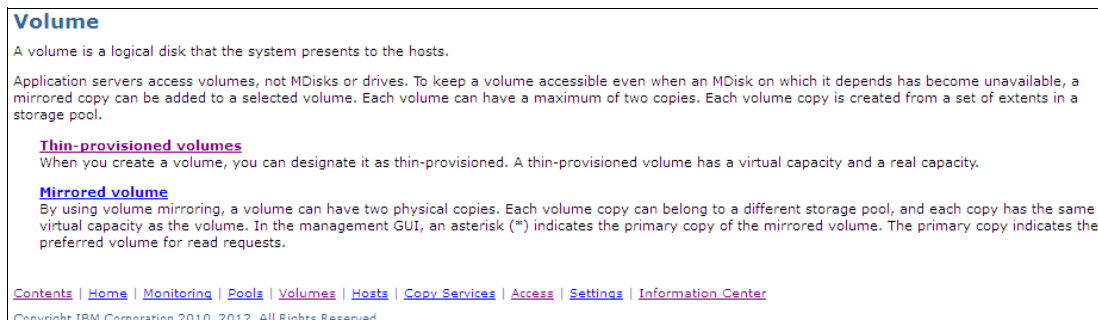


Figure 3-122 Information panel

3.5.4 Embedded panel help

The Storwize V3700 provides embedded help that is available on each panel. Click **Help** to reveal the information panel. Figure 3-123 on page 132 shows the embedded panel help.




Figure 3-123 Embedded panel help

Figure 3-124 shows the information panel revealed from the embedded panel help. The information panel has hotlinks to a range of other panels, including the Information Center.

Hosts

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

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

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

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

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

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

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

Parent topic: [Overview](#)

[Contents](#) | [Home](#) | [Monitoring](#) | [Pools](#) | [Volumes](#) | [Hosts](#) | [Copy Services](#) | [Access](#) | [Settings](#) | [Information Center](#)

Copyright IBM Corporation 2010, 2012. All Rights Reserved.

Figure 3-124 Information panel

3.5.5 Hidden question mark help

The Storwize V3700 provides a hidden question mark help feature for some settings or items found in the various configuration panels. This help feature is provided by hovering adjacent to an item, the question mark is then revealed and the help bubble is displayed. Figure 3-125 show the revealed question mark and help bubble.

General

Drive ID	The system-assigned identifier for the drive.
Status	✔ Online
Use	Member
UID	5000c50020d51ca7

Figure 3-125 Hidden question mark help

3.5.6 Hover help

The Storwize V3700 provides a hidden help tags that are revealed when hovering over various functions and items. Figure 3-126 on page 133 shows the hover help feature.




Figure 3-126 Hover help

3.5.7 IBM endorsed Youtube videos

IBM has endorsed a range of Youtube videos for the IBM storage portfolio. Client feedback suggests that this is a good tool to display management GUI navigation and tasks. Check for new videos from IBM Storage from time to time to find useful information.

IBM System Storage Channel:

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

