


Implementing the IBM Storwize V3700

Easily manage and deploy systems
with embedded GUI

Experience rapid and flexible
provisioning

Protect data with remote
mirroring



Jon Tate
Adam Lyon-Jones
Lee Sirett
Chris Tapsell
Paulo Tomiyoshi Takeda

Redbooks



International Technical Support Organization

Implementing the IBM Storwize V3700

March 2015

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

Third Edition (March 2015)

This edition applies to IBM Storwize V3700 running software version 7.4.

© Copyright International Business Machines Corporation 2014, 2015. All rights reserved.

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

Contents

Notices	ix
Trademarks	x
IBM Redbooks promotions	xi
Preface	xiii
Authors	xiv
Now you can become a published author, too	xvi
Comments welcome	xvi
Stay connected to IBM Redbooks	xvi
Summary of changes	xvii
February 2015, Third Edition	xvii
Chapter 1. Overview of the IBM Storwize V3700 system	1
1.1 IBM Storwize V3700 overview	2
1.2 IBM Storwize V3700 terminology	3
1.3 IBM Storwize V3700 models	5
1.4 IBM Storwize V3700 hardware	7
1.4.1 Control enclosure	8
1.4.2 Expansion enclosure	10
1.4.3 Host interface cards	11
1.4.4 Disk drive types	11
1.5 IBM Storwize V3700 terms	12
1.5.1 Hosts	12
1.5.2 Node canister	12
1.5.3 I/O Group	12
1.5.4 Clustered system	13
1.5.5 RAID	13
1.5.6 Managed disks	14
1.5.7 Quorum disks	14
1.5.8 Storage Pools	15
1.5.9 Child Pools	17
1.5.10 Volumes	17
1.5.11 iSCSI	19
1.5.12 SAS	20
1.6 IBM Storwize V3700 features	20
1.6.1 Volume mirroring	20
1.6.2 Thin provisioning	21
1.6.3 Easy Tier	21
1.6.4 Turbo Performance	22
1.6.5 Storage Migration	22
1.6.6 FlashCopy	22
1.6.7 Remote Copy	23
1.7 Problem management and support	24
1.7.1 IBM Assist On-site and remote service	24
1.7.2 Event notifications	24
1.7.3 Useful IBM Storwize V3700 websites	25

Chapter 2. Initial configuration	27
2.1 Hardware installation planning	28
2.2 Storage area network configuration planning	30
2.3 SAS direct attach planning	32
2.4 LAN configuration planning	34
2.4.1 Management IP address considerations	34
2.4.2 Service IP address considerations	35
2.5 Host configuration planning	36
2.6 Miscellaneous configuration planning	37
2.7 System management	38
2.7.1 Graphical user interface	38
2.7.2 Command-line interface	39
2.8 First-time setup	40
2.9 Initial configuration	49
2.9.1 Adding expansion enclosures after initial configuration	59
2.9.2 Configure Call Home, email alert, and inventory	65
2.9.3 Service Assistant tool	67
Chapter 3. Graphical user interface overview	71
3.1 Getting started	72
3.1.1 Supported browsers	72
3.1.2 Access the management GUI	72
3.1.3 System pane layout	74
3.2 Navigation	76
3.2.1 Function icons navigation	76
3.2.2 Breadcrumb navigation aid	78
3.2.3 Suggested Tasks feature	78
3.2.4 Presets	79
3.2.5 Actions	79
3.2.6 Task progress	81
3.2.7 Navigating panes with tables	81
3.3 Status Indicators menus	89
3.3.1 Allocated / Virtual Capacity status bar menu	89
3.3.2 Running Tasks bar menu	90
3.3.3 Health Status bar menu	91
3.4 Function icon menus	92
3.4.1 Monitoring menu	93
3.4.2 Pools menu	104
3.4.3 Volumes menu	114
3.4.4 Hosts menu	118
3.4.5 Copy Services menu	121
3.4.6 Access menu	126
3.4.7 Settings menu	129
3.5 Management GUI help	144
3.5.1 IBM Storwize V3700 Information Center	145
3.5.2 Watching an e-Learning video	145
3.5.3 Embedded pane help	145
3.5.4 Hidden question mark help	146
3.5.5 Hover help	147
3.5.6 IBM-endorsed YouTube videos	147

Chapter 4. Host configuration	149
4.1 Host attachment overview	150
4.2 Preparing the host operating system	151
4.2.1 Windows 2012 R2: Preparing for FC attachment	151
4.2.2 Windows 2012 R2: Preparing for iSCSI attachment	155
4.2.3 Windows 2012 R2: Preparing for SAS attachment	162
4.2.4 VMware ESXi 5.5: Preparing for FC attachment	163
4.2.5 VMware ESXi 5.5: Preparing for iSCSI attachment	166
4.2.6 VMware ESXi 5.5: Preparing for SAS attachment	178
4.3 Configuring hosts in IBM Storwize V3700	180
4.3.1 Creating FC hosts	182
4.3.2 Configuring IBM Storwize V3700 for FC connectivity	185
4.3.3 Creating iSCSI hosts	188
4.3.4 Configuring IBM Storwize V3700 for iSCSI host connectivity	191
4.3.5 Creating SAS hosts	195
Chapter 5. Volume configuration	199
5.1 Creating volumes in IBM Storwize V3700	200
5.1.1 Creating generic volumes	202
5.1.2 Creating thin-provisioned volumes	205
5.1.3 Creating mirrored volumes	208
5.1.4 Creating thin-mirrored volumes	211
5.2 Mapping volumes to hosts	213
5.2.1 Mapping newly created volumes using Create and Map to Host	213
5.2.2 Manually mapping volumes to hosts	215
5.3 Discovering mapped volumes from host systems	216
5.3.1 Discovering Fibre Channel volumes from Windows 2012 R2	218
5.3.2 Discovering iSCSI volumes from Windows 2012 R2	227
5.3.3 Discovering SAS volumes from Windows 2012 R2	233
5.3.4 Discovering Fibre Channel volumes from VMware ESXi 5.5	234
5.3.5 Discovering iSCSI volumes from VMware ESXi 5.5	235
5.3.6 Discovering SAS volumes from VMware ESXi 5.5	247
Chapter 6. Storage migration	249
6.1 Interoperability and compatibility	250
6.2 Storage migration considerations	250
6.2.1 External virtualization capability	250
6.2.2 Overview of the Storage Migration wizard	250
6.2.3 Storage Migration wizard tasks	251
6.2.4 Storage migration example configuration	252
6.2.5 Recording migration information	252
6.3 Storage migration example	259
Chapter 7. Storage Pools	307
7.1 Working with internal drives	308
7.1.1 Internal Storage window	309
7.1.2 Actions on internal drives	311
7.2 Configuring internal storage	320
7.2.1 RAID configuration presets	321
7.2.2 Customizing initial storage configuration	322
7.2.3 Creating an MDisk and Pool	324
7.2.4 Using the recommended configuration	326
7.2.5 Selecting a different configuration	329
7.2.6 MDisk by Pools pane	336

7.2.7 RAID action for MDisks	339
7.2.8 More actions on MDisks	343
7.3 Working with Storage Pools	348
7.3.1 Create Pool option	351
7.3.2 Actions on Storage Pools	352
7.4 Working with Child Pools	355
7.4.1 Create Child Pools	356
7.4.2 Actions on Child Pools	357
7.4.3 Resizing a Child Pool	357
Chapter 8. Advanced host and volume administration	359
8.1 Advanced host administration	360
8.1.1 Modifying Mappings menu	362
8.1.2 Unmapping volumes from a host	367
8.1.3 Renaming a host	370
8.1.4 Removing a host	371
8.1.5 Host properties	373
8.2 Adding and deleting host ports	377
8.2.1 Adding a host port	378
8.2.2 Adding a Fibre Channel port	378
8.2.3 Adding an SAS host port	380
8.2.4 Adding an iSCSI host port	381
8.2.5 Deleting a host port	382
8.3 Host mappings overview	383
8.3.1 Unmap Volumes	384
8.3.2 Properties (Host)	384
8.3.3 Properties (Volume)	384
8.4 Advanced volume administration	384
8.4.1 Advanced volume functions	385
8.4.2 Mapping a volume to a host	387
8.4.3 Unmapping volumes from all hosts	389
8.4.4 Viewing which host is mapped to a volume	390
8.4.5 Duplicating a volume	390
8.4.6 Renaming a volume	391
8.4.7 Shrinking a volume	392
8.4.8 Expanding a volume	393
8.4.9 Migrating a volume to another Storage Pool	393
8.4.10 Deleting a volume	395
8.5 Volume properties	396
8.5.1 Overview tab	397
8.5.2 Host Maps tab	400
8.5.3 Member MDisks tab	401
8.5.4 Adding a mirrored volume copy	402
8.5.5 Editing thin-provisioned volume properties	403
8.6 Advanced volume copy functions	406
8.6.1 Thin-provisioned menu	408
8.6.2 Splitting into a new volume	408
8.6.3 Validate Volume Copies option	410
8.6.4 Duplicate Volume option	412
8.6.5 Delete Volume Copy option	412
8.6.6 Migrating volumes by using the volume copy features	413
8.7 Volumes by Storage Pool	414
8.8 Volumes by Host	416

Chapter 9. IBM Easy Tier	417
9.1 Generations of IBM Easy Tier	418
9.2 New Features in Easy Tier 3	418
9.3 Easy Tier overview	419
9.3.1 Easy Tier licensing	420
9.3.2 Tiered Storage Pools	423
9.4 Easy Tier process	425
9.4.1 I/O Monitoring	426
9.4.2 Data Placement Advisor	426
9.4.3 Data Migration Planner	426
9.4.4 Data Migrator	426
9.4.5 Easy Tier operating modes	427
9.4.6 Easy Tier status	428
9.4.7 Storage Pool Balancing	428
9.4.8 Easy Tier rules	429
9.5 Easy Tier configuration using the GUI	431
9.5.1 Creating multitiered pools: Enabling Easy Tier	431
9.5.2 Downloading Easy Tier I/O measurements	438
9.6 Easy Tier configuration using the CLI	440
9.6.1 Enabling Easy Tier measured mode	441
9.6.2 Enabling or disabling Easy Tier on single volumes	444
9.7 IBM Storage Tier Advisor Tool	446
9.7.1 Processing heat log files	447
9.7.2 Storage Tier Advisor Tool reports	448
Chapter 10. Copy services	453
10.1 FlashCopy	454
10.1.1 Business requirements for FlashCopy	454
10.1.2 FlashCopy functional overview	455
10.1.3 Planning for FlashCopy	463
10.1.4 Managing FlashCopy using the GUI	465
10.1.5 Managing FlashCopy mappings	474
10.1.6 Managing a FlashCopy consistency group	500
10.2 Remote Copy	514
10.2.1 Remote Copy concepts	515
10.2.2 Global Mirror with Change Volumes	521
10.2.3 Remote Copy planning	525
10.3 Troubleshooting Remote Copy	529
10.3.1 1920 error	529
10.3.2 1720 error	531
10.4 Managing Remote Copy using the GUI	532
10.4.1 Managing cluster partnerships	532
10.4.2 Managing stand-alone Remote Copy relationships	540
10.4.3 Managing a Remote Copy consistency group	556
Chapter 11. RAS, monitoring, and troubleshooting	575
11.1 Reliability, availability, and serviceability	576
11.2 IBM Storwize V3700 components	577
11.2.1 Enclosure midplane assembly	577
11.2.2 Node canisters: Ports and light-emitting diode (LED)	577
11.2.3 Node canister replaceable hardware components	581
11.2.4 Expansion canister: Ports and LED	585
11.2.5 Disk subsystem	586
11.2.6 Power supply unit	590

11.3 Configuration backup procedure	591
11.3.1 Generating a manual configuration backup by using the CLI	591
11.3.2 Downloading a configuration backup by using the graphical user interface	592
11.4 Updating software	595
11.4.1 Node canister software	595
11.4.2 Upgrading drive firmware	608
11.5 Event log	612
11.5.1 Managing the event log.	613
11.5.2 Alert handling and recommended actions.	617
11.6 Collecting support information	624
11.6.1 Support information from the GUI	624
11.6.2 Support information from the Service Assistant	627
11.6.3 Support information to USB stick	628
11.7 Powering on and powering off the IBM Storwize V3700	629
11.7.1 Powering off the system	629
11.7.2 Powering on	632
11.8 Tivoli Storage Productivity Center	632
11.8.1 Tivoli Storage Productivity Center benefits	632
11.8.2 Adding IBM Storwize V3700 in Tivoli Storage Productivity Center	633
11.9 Using Tivoli Storage Productivity Center to administer and generate reports for an IBM Storwize V3700	635
11.9.1 Basic configuration and administration	635
11.9.2 Generating reports by using Java GUI	637
11.9.3 Generating reports using Tivoli Storage Productivity Center web console	641
Appendix A. Command-line interface setup and SAN Boot	647
Command-line interface	648
Basic setup	648
Example commands	659
Upgrading Drive Firmware using the CLI	661
Copying, installing, and running the Upgrade Test Utility on the Storwize unit.	662
SAN Boot	666
Related publications	667
IBM Redbooks	667
IBM Storwize V3700 publications.	667
IBM Storwize V3700 support	667
Help from IBM	667

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 about 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](http://www.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:

AS/400®	InfoSphere®	System i®
DS8000®	Power Systems™	System Storage®
Easy Tier®	Redbooks®	System Storage DS®
FlashCopy®	Redbooks (logo)  ®	Tivoli®
IBM®	Storwize®	XIV®

The following terms are trademarks of other companies:

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

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

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

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

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

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

Find and read thousands of IBM Redbooks publications

- ▶ Search, bookmark, save and organize favorites
- ▶ Get up-to-the-minute Redbooks news and announcements
- ▶ Link to the latest Redbooks blogs and videos


Get the latest version of the Redbooks Mobile App



iOS

Download
Now

Android



Promote your business in an IBM Redbooks publication

Place a Sponsorship Promotion in an IBM® Redbooks® publication, featuring your business or solution with a link to your web site.

Qualified IBM Business Partners may place a full page promotion in the most popular Redbooks publications. Imagine the power of being seen by users who download millions of Redbooks publications each year!



ibm.com/Redbooks

About Redbooks → Business Partner Programs

THIS PAGE INTENTIONALLY LEFT BLANK

Preface

Organizations of all sizes are faced with the challenge of managing massive volumes of increasingly valuable data. However, storing this data can be costly, and extracting value from the data is becoming more and more difficult.

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

Storwize V3700 delivers efficient, entry-level configurations that are specifically designed to meet the needs of small and midsize businesses. Designed to provide organizations with the ability to consolidate and share data at an affordable price, Storwize V3700 offers advanced software capabilities that are usually found in more expensive systems.

Built on innovative IBM technology, Storwize V3700 addresses the block storage requirements of small and midsize organizations, Storwize V3700 is designed to accommodate the most common storage network technologies. This design enables easy implementation and management.

Storwize V3700 includes the following features:

- ▶ Web-based GUI provides point-and-click management capabilities.
- ▶ Internal disk storage virtualization enables rapid, flexible provisioning and simple configuration changes.
- ▶ Thin provisioning enables applications to grow dynamically, but only use space they actually need.
- ▶ Enables simple data migration from external storage to Storwize V3700 storage (one-way from another storage device).
- ▶ Remote Mirror creates copies of data at remote locations for disaster recovery.
- ▶ IBM FlashCopy® creates instant application copies for backup or application testing.

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

The concepts in this book also relate to the IBM Storwize V3500.

This book was written at a software level of version 7 release 4.

Authors

This book was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO), Manchester Labs, United Kingdom (UK).



Jon Tate is a Project Manager for IBM System Storage® storage area network (SAN) Solutions at the ITSO, 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 28 years of experience in storage software and management, services, and support, and is both an IBM Certified information technology (IT) Specialist and an IBM SAN Certified Specialist. He is also the UK Chairman of the Storage Networking Industry Association (SNIA).



Adam Lyon-Jones is a Test Engineer in IBM Storage Systems at the IBM Manchester Lab as part of a global team developing, testing, and supporting IBM Storage products. He has worked for two years performing system and integration testing for IBM SAN Volume Controller and IBM Storwize products. His areas of expertise include storage hardware, storage virtualization, and SANs. Before joining IBM, he studied a Masters degree in Physics at Durham University.



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



Chris Tapsell is a Certified Storage Client Technical Specialist for IBM Systems and Technology Group (STG) in the UK. In his 25+ years at IBM, he has worked as an IBM Service Representative covering products such as Office Products to IBM AS/400® (IBM System i®), as a Support Specialist for all of the IBM Intel server products (PC Server, Netfinity, xSeries, and IBM System x), PCs, and notebooks, and as a Client Technical Specialist for System x. Chris holds several IBM Certifications covering System x and Systems Storage products.



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

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

- ▶ Martyn Spink
- ▶ Djihed Afifi
- ▶ Arthur Wellesley
- ▶ James Whitaker
- ▶ Imran Imtiaz
- ▶ Tobias Fleming
- IBM Manchester Lab
- ▶ John Fairhurst
- ▶ Paul Marris
- ▶ Paul Merrison
- IBM Hursley
- ▶ Samrat Dutta
- IBM Pune

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

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

Now you can become a published author, too

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time. Join an ITSO residency project and help write a book in your area of expertise, while honing your experience 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 two - six weeks, and you can participate either in person or as a remote resident working from your home base.

Learn more about the residency program, browse the residency index, and apply online:

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:

ibm.com/redbooks

- ▶ Send your comments in an email:

redbooks@us.ibm.com

- ▶ Mail your comments:

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>

Summary of changes

This section describes the technical changes made in this edition of the book and in previous editions. This edition might also include minor corrections and editorial changes that are not identified.

Summary of Changes
for SG24-8107-02
for Implementing the IBM Storwize V3700
as created or updated on February 19, 2016.

March 2015, Third Edition

This revision includes the following new and changed information.

New information

- ▶ Scalability data
- ▶ Drive support
- ▶ New graphical user interface (GUI) look and feel
- ▶ New command-line interface (CLI) commands

Changed information

- ▶ GUI screen captures
- ▶ CLI commands and output



Overview of the IBM Storwize V3700 system

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

This chapter includes the following topics:

- ▶ IBM Storwize V3700 overview
- ▶ IBM Storwize V3700 terminology
- ▶ IBM Storwize V3700 models
- ▶ IBM Storwize V3700 hardware
- ▶ IBM Storwize V3700 terms
- ▶ IBM Storwize V3700 features
- ▶ Problem management and support

1.1 IBM Storwize V3700 overview

The IBM 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 Redundant Array of Independent Disks (RAID) entry-level storage system.

IBM Storwize V3700 features the following benefits:

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

The IBM Storwize V3700 addresses the block storage requirements of small and midsize organizations, and consists of one 2U *control enclosure*. Optionally, you can add up to nine 2U *expansion enclosures* that are connected with serial-attached SCSI (SAS) cables, which make up one system that is called an input/output group (*I/O Group*).

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

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

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

The IBM Storwize V3700 scales up to 120 x 3.5-inch, 240 x 2.5-inch, or a combination of both, drive form factors.

The SAS, Nearline SAS (NL-SAS), and solid-state drives (SSDs, also known as flash) are the supported drive types.

IBM Storwize V3700 provides host attachment through onboard 6 Gb SAS and 1 Gb Internet Small Computer System Interface (iSCSI) host interfaces. In addition, there are optional four-port 8 Gb Fibre Channel (FC), two-port 10 Gb iSCSI/Fibre Channel over Ethernet (FCoE) or four-port 6 Gb SAS host interface cards (HIC).

Important: IBM Storwize V3700 can be direct-attached to a host. For more information about restrictions, see the IBM System Storage Interoperation Center (SSIC) on the following website:

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

Host attachment information also is available on the following website:

http://www.ibm.com/support/knowledgecenter/STLM5A_7.4.0/com.ibm.storwize.v3700.740.doc/v3700_ichome_740.html

The IBM Storwize V3700 is a virtualized storage solution that groups its internal drives into RAID arrays, which are called *managed disks* (MDisks). These MDisks are then grouped into *Storage Pools*, from which *volumes* are created and provisioned out to *hosts*. Storage Pools are normally created with MDisks of the same type and capacity of drive.

Volumes can be moved non-disruptively between Storage Pools with differing performance characteristics. For example, a volume can be moved between a Storage Pool made up of NL-SAS drives to a Storage Pool made up of SAS drives.

The IBM Storwize V3700 system also provides several configuration options that are aimed at simplifying the implementation process. It also provides configuration presets and automated wizards, called *Directed Maintenance Procedures* (DMP) to help resolve any events that might occur.

IBM Storwize V3700 system provides a simple 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 command-line interface (CLI) to set up or control the system.

1.2 IBM Storwize V3700 terminology

The IBM Storwize V3700 system introduced some terminology, which is consistent with the entire IBM Storwize family and IBM SAN Volume Controller, as shown in Table 1-1.

Table 1-1 IBM Storwize V3700 terminology

IBM Storwize V3700 term	Definition
Battery	Each control enclosure node canister in a IBM Storwize V3700 contains a battery.
Canister	Canisters are hardware units that are subcomponents of an IBM Storwize V3700 enclosure. Each enclosure contains two canisters.
Chain	A set of enclosures that is attached to provide redundant access to the drives that are inside the enclosures. Each control enclosure has two chains.
Clone	A copy of a volume on a server at a particular point. The contents of the copy can be customized while the contents of the original volume are preserved.
Control enclosure	A hardware unit that includes the chassis, node canisters, drives, and power sources.
Data migration	By using IBM Storwize V3700, you can migrate data from existing external storage to its internal volumes.
Drive	IBM Storwize V3700 supports a range of hard disk drives (HDDs) and SSDs.
Enclosure	An enclosure is the basic housing unit for the IBM Storwize V3700. It is the rack-mounted hardware that contains all the main components of the system: <ul style="list-style-type: none"> ▶ Canisters ▶ Drives ▶ Power supplies

IBM Storwize V3700 term	Definition
Event	An occurrence that is significant to a task or system. Events can include completion or failure of an operation, a user action, or the change in the state of a process.
Expansion canister	A hardware unit that includes the SAS interface hardware that enables the node hardware to use the drives of the expansion enclosure.
Expansion enclosure	A hardware unit that includes expansion canisters, drives, and power supply units.
Fibre Channel port	FC ports are connections for the hosts to get access to the IBM Storwize V3700.
Host mapping	The process of controlling which hosts can access specific volumes in a IBM Storwize V3700.
iSCSI	Internet Protocol (IP)-based storage networking standard for linking data storage facilities.
Internal storage	Array MDisks and drives that are held in enclosures and nodes that are part of the IBM Storwize V3700.
Managed disk	A component of a Storage Pool that is managed by a clustered system. An MDisk is part of a RAID array of internal storage. An MDisk is not visible to a host system on the storage area network (SAN).
Node canister	A hardware unit that includes the node hardware, fabric, and service interfaces, SAS, expansion ports, and battery.
Physical layer (PHY)	A single SAS lane. There are four PHYs in each SAS cable.
Power supply unit (PSU)	Each enclosure has two PSUs.
Quorum disk	A disk that contains a reserved area that is used exclusively for cluster management. The quorum disk is accessed when it is necessary to determine which half of the cluster continues to read and write data.
Serial-Attached SCSI ports	SAS ports are connections for the host to get direct-attached access to the IBM Storwize V3700 and expansion enclosure.
Snapshot	An image backup type that consists of a point-in-time view of a volume.
Storage Pool	A collection of storage capacity that provides the capacity requirements for a volume.
Child Pool	A subset of a Storage Pool that provides the capacity for a volume.
Strand	The SAS connectivity of a set of drives within multiple enclosures. The enclosures can be 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 that is assigned to that storage unit.
Turbo Performance	Increases system maximum input/output (I/O) operations per second (IOPS) and maximum throughput.

IBM Storwize V3700 term	Definition
Volume	A discrete unit of storage on disk, tape, or other data recording medium that supports some form of identifier and parameter list, such as a volume label or I/O control.
Worldwide port name (WWPN)	Each FC or SAS port is identified by their physical port number and worldwide port name (WWPN).

1.3 IBM Storwize V3700 models

The IBM Storwize V3700 platform consists of several different models.

More information: For more information about the features, benefits, and specifications of IBM Storwize V3700 models, see the following website:

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

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

The IBM Storwize V3700 models are described in Table 1-2. All models have two node canisters. C models are control enclosures and E models are expansion enclosures.

Table 1-2 IBM Storwize V3700 models

Model	Total System Cache	Drive slots
2072-12C (with two node canisters)	8 gigabytes (GB) upgradeable to 16 GB	12x 3.5-inch per enclosure
2072-24C (with two node canisters)	8 GB upgradeable to 16 GB	24x 2.5-inch per enclosure
2072-2DC (direct current (dc) powered, with two node canisters)	8 GB upgradeable to 16 GB	24x 2.5-inch per enclosure
2072-12E (one expansion enclosure)	N/A	12x 3.5-inch
2072-24E (one expansion enclosure)	N/A	24x 2.5-inch
2072-2DE (dc powered, with one expansion enclosure)	N/A	24x 2.5-inch

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



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

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

Figure 1-2 shows the front view of the 2072-24C and 24E enclosures.




Figure 1-2 IBM Storwize V3700 front view for 2072-24C and 24E enclosure

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

1.4 IBM Storwize V3700 hardware

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

Figure 1-3 shows an overview of the hardware components of the IBM Storwize V3700 solution.

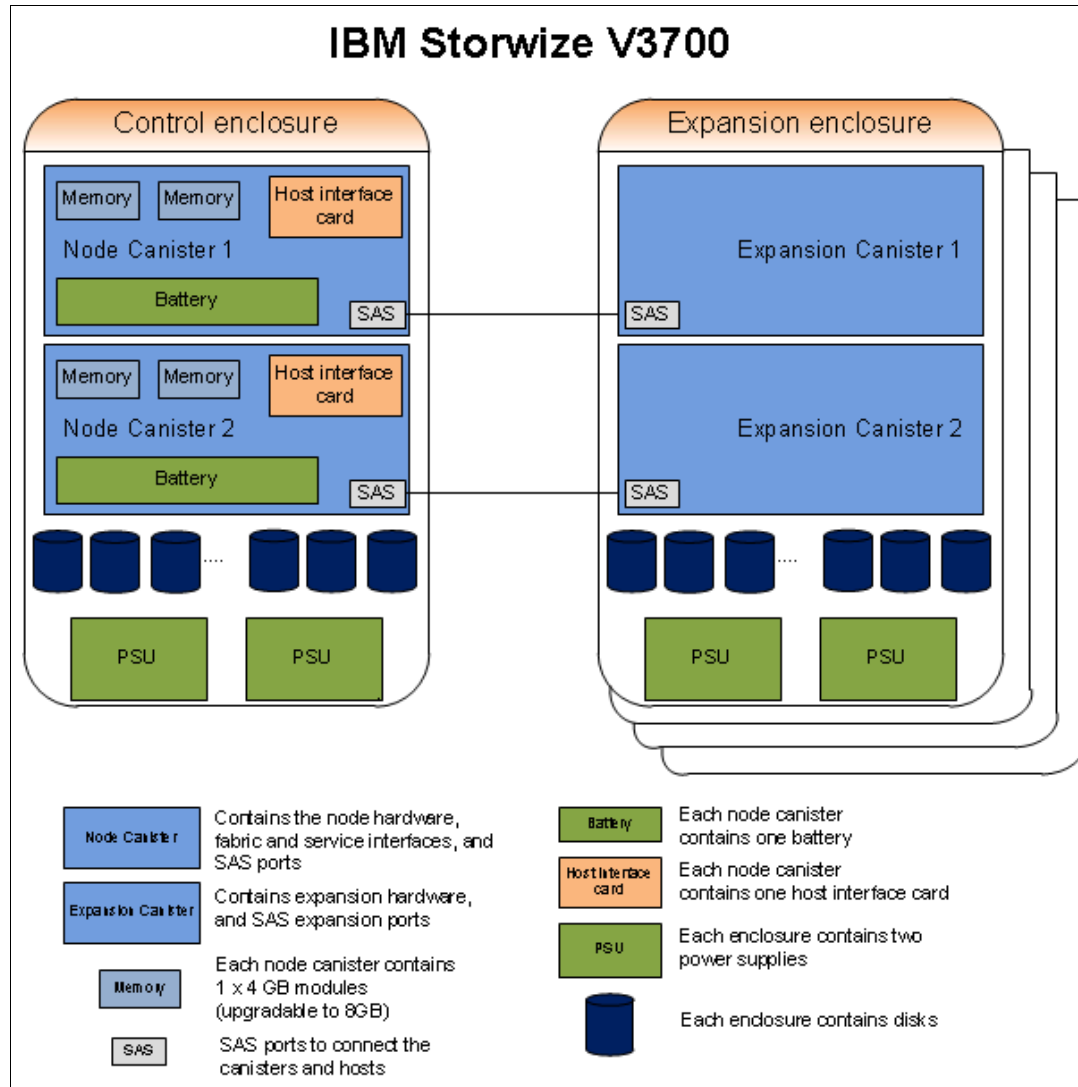


Figure 1-3 IBM Storwize V3700 hardware components

Figure 1-4 shows the controller rear view of IBM Storwize V3700 models 12C and 24C.



Figure 1-4 IBM Storwize V3700 controller rear view: Models 12C and 24C

Figure 1-4 shows the two power supply slots at the bottom of the enclosure. The PSUs are identical and exchangeable. There are two canister slots at the top of the chassis.

Figure 1-5 shows the rear view of an IBM Storwize V3700 expansion enclosure.



Figure 1-5 IBM Storwize V3700 expansion enclosure rear view: Models 12E and 24E

The only difference between the control enclosure and the expansion enclosure are the canisters and vital product data (VPD) written to the backplane. The canisters of the expansion have only the two SAS ports.

For more information about the expansion enclosure, see 1.4.2, “Expansion enclosure” on page 10.

1.4.1 Control enclosure

Each IBM Storwize V3700 system has one control enclosure that contains two node canisters, disk drives, and two power supplies.

Figure 1-6 shows a single node canister.




Figure 1-6 IBM Storwize V3700 node canister

Each node canister contains the following hardware:

- ▶ Battery
- ▶ Memory: 4 GB memory (upgradable to 8 GB)
- ▶ HIC slot (different options are possible)
- ▶ Four 6 gigabits per second (Gbps) SAS ports
- ▶ Two 10/100/1000 megabits per second (Mbps) Ethernet ports
- ▶ Two Universal Serial Bus (USB) 2.0 ports (one port is used during installation)
- ▶ System flash

The battery is used in case of power loss. The IBM Storwize V3700 system uses this battery to power the canister while the cache data is written to the internal system flash. This memory dump is called a *Fire Hose Dump* (FHD). After the system is up again, this data is loaded back to the cache for destage to the disks.

Figure 1-6 on page 8 also shows the following ports, which are provided by the IBM Storwize V3700 node canister:

- ▶ Two 10/100/1000 Mbps Ethernet ports
Port 1 (left port) must be configured. The second port is optional and can be used for management. Both ports can be used for iSCSI traffic. For more information, see Chapter 4, “Host configuration” on page 149.
- ▶ Two USB ports
One port is used during the initial configuration or when there is a problem. They are numbered 1 on the left and 2 on the right. For more information about usage, see Chapter 2, “Initial configuration” on page 27.
- ▶ Four SAS ports
These ports are numbered 1 on the left to 4 on the right. The IBM Storwize V3700 uses ports 1, 2, and 3 for host connectivity, and port 4 to connect to the optional expansion enclosures. The IBM Storwize V3700 incorporates one SAS chain, and nine expansion enclosures can be connected to each chain.

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

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

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 IBM Storwize V3700 only supports one I/O group in a clustered system.

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

1.4.2 Expansion enclosure

The optional IBM Storwize V3700 expansion enclosure contains two expansion canisters, disk drives, and two power supplies.

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



Figure 1-7 Expansion enclosure of the IBM Storwize V3700

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

Figure 1-8 shows the expansion canister ports.




Figure 1-8 Expansion canister ports

Each expansion canister that is shown in Figure 1-8 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.

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

Each port includes two light-emitting diodes (LEDs) to show the status. The first LED indicates the link status, and the second LED indicates the fault status.

For more information about LED or ports, see the following website:

http://www.ibm.com/support/knowledgecenter/STLM5A_7.4.0/com.ibm.storwize.v3700.740.doc/tbrd4_expcanindi.html

1.4.3 Host interface cards

In addition to onboard 1 Gb iSCSI and 6 Gb SAS, there are options to purchase HICs for the control enclosure. Table 1-3 lists the IBM Storwize V3700 configurations available.

Table 1-3 IBM Storwize V3700 configurations available

Standard		Optional			
1 Gb iSCSI	6 Gb SAS	1 Gb iSCSI	6 Gb SAS	8 Gb FC	10 Gb iSCSI/FCoE
4 ports	6 ports	-	-	-	-
4 ports	6 ports	8 ports	-	-	-
4 ports	6 ports	-	8 ports	-	-
4 ports	6 ports	-	-	8 ports	-
4 ports	6 ports	-	-	-	4 ports

1.4.4 Disk drive types

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

Table 1-4 shows the IBM Storwize V3700 disk drive types that are available at the time of writing.

Table 1-4 IBM Storwize V3700 Disk Drive types

Drive type		Speed	Size
2.5-inch form factor	flash	N/A	200 GB, 400 GB, and 800 GB
2.5-inch form factor	SAS	10,000 revolutions per minute (RPM)	600 GB, 900 GB, 1.2 terabytes (TB) and 1.8 TB
2.5-inch form factor	SAS	15,000 RPM	146 GB, 300 GB, and 600 GB
2.5-inch form factor	Nearline SAS	7,200 RPM	1 TB
3.5-inch form factor	SAS	10,000 RPM	900 GB, 1.2 TB, and 1.8 TB ^a
3.5-inch form factor	SAS	15,000 RPM	300 GB and 600 GB ^a
3.5-inch form factor	Nearline SAS	7,200 RPM	2 TB, 3 TB, 4 TB, and 6 TB

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

1.5 IBM Storwize V3700 terms

In this section, we introduce the terms that are used for the IBM Storwize V3700 throughout this book.

1.5.1 Hosts

A host system is a server that is connected to IBM Storwize V3700 through an FC connection, an iSCSI connection, or through an SAS connection.

Hosts are defined on IBM Storwize V3700 by identifying their worldwide port names (WWPNs) for FC and SAS hosts. iSCSI hosts are identified by using their iSCSI names. The iSCSI names can be iSCSI qualified names (IQNs) or extended unique identifiers (EUIs). For more information, see Chapter 4, “Host configuration” on page 149.

Hosts can be FC attached using an existing FC network infrastructure or direct attached, iSCSI/FCoE attached using an existing IP network, or directly attached using SAS. A significant benefit of having direct attachment is that you can attach the host directly to the IBM Storwize V3700 without the need for an FC or IP network.

1.5.2 Node canister

Each single processing unit is a *node canister*. A node canister provides host interfaces, management interfaces, and SAS interfaces to the control enclosure. A node canister has the cache memory, the internal storage to store software and logs, and the processing power to run the IBM Storwize V3700 virtualizing and management software. A clustered system consists of a pair of nodes.

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

1.5.3 I/O Group

Within IBM Storwize V3700, there is one pair of node canisters known as an *I/O group*. The IBM Storwize V3700 supports a maximum of the two node canisters in the clustered system, which provides a single I/O Group.

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

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

In an IBM Storwize V3700 environment that uses active/active architecture, the I/O handling for a volume can be managed by both nodes. Therefore, servers that are connected through FC connectors must use multipath device drivers to handle this capability.

The IBM Storwize V3700 I/O Group can be connected to the SAN so that all application servers can access the volumes from the I/O Group. Up to 256 host server objects can be defined to the IBM Storwize V3700.

Important: The active/active architecture provides availability to process I/Os for both control nodes, and enables 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/logical unit number (LUN) thrashing typical of an active/passive architecture.

1.5.4 Clustered system

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

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

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

The system can be configured by using the IBM Storwize V3700 management software (GUI), the CLI, or the USB key.

1.5.5 RAID

The IBM Storwize V3700 contains several internal drives, but these drives cannot be directly added to Storage Pools. The drives must be included in a 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.

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

The following RAID levels are available:

- ▶ RAID 0 (striping, no redundancy)
- ▶ RAID 1 (mirroring between two drives, implemented as RAID 10 with 2 drives)
- ▶ RAID 5 (striping, can survive one drive fault, with parity)
- ▶ RAID 6 (striping, can survive two drive faults, with parity)
- ▶ RAID 10 (RAID 0 on top of RAID 1)

RAID 0 arrays stripe data across the drives. The system supports RAID 0 arrays with one member, which is similar to traditional “just a bunch of disks” (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 - eight drives.

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

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

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

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

1.5.6 Managed disks

An MDisk refers to the unit of storage that IBM Storwize V3700 virtualizes. This unit is a RAID array consisting of internal drives. The IBM Storwize V3700 then can allocate these MDisks into Storage Pools.

An MDisk is not visible to a host system on the SAN because it is internal in the IBM Storwize V3700 system.

An MDisk features the following modes:

- ▶ Array
Array mode MDisks are constructed from internal drives by using the RAID functionality. Array MDisks are always associated with Storage Pools.
- ▶ Image
Image MDisks are LUNs presented by external storage systems to an IBM Storwize V3700, and assigned directly to a volume with a one-to-one mapping of extents between the MDisk and the volume. For more information, see Chapter 6, “Storage migration” on page 249.

1.5.7 Quorum disks

A quorum disk is an MDisk that contains a reserved area for use exclusively by the system. In the IBM Storwize V3700, internal drives can be considered as *quorum candidates*. The clustered system uses quorum disks to break a tie when exactly half the nodes in the system remain after a SAN failure.

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

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

1.5.8 Storage Pools

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

MDisks can be added to a Storage Pool at any time to increase the capacity of the pool. MDisks can belong in only one Storage Pool. For more information, see Chapter 7, “Storage Pools” on page 307.

Each MDisk in the Storage Pool is divided into several extents. The size of the extent ranges from 16 mebibytes (MiB) - 8 gibibytes (GiB).

Default extent size: The GUI of IBM Storwize V3700 has a default extent size value of 1 GiB when you define a new Storage Pool. This is a change in the IBM Storwize code V7.1.

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

A system can manage 2^{22} (4,194,304) extents. For example, with a 16 MB extent size, the system can manage up to $16 \text{ MiB} \times 4,194,304 = 64 \text{ tibibytes (TiB)}$ of storage.

The effect of extent size on the maximum volume and cluster size is shown in Table 1-5.

Table 1-5 Maximum volume and cluster capacity by extent size

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

Use the same extent size for all Storage Pools in a clustered system, which is a prerequisite if you want to migrate a volume between two Storage Pools. If the Storage Pool extent sizes are not the same, you must use volume mirroring to copy volumes between Storage Pools, as described in Chapter 1, “Overview of the IBM Storwize V3700 system” on page 1.

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

Single-tiered Storage Pool

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

- ▶ They should have the same hardware characteristics, for example, the same RAID type, RAID array size, disk type, and disk RPMs.
- ▶ The disk subsystems providing the MDisks must have similar characteristics, for example, maximum IOPS, response time, cache, and throughput.
- ▶ Use MDisks of the same size, and ensure that the MDisks provide the same number of extents. If this configuration is not feasible, you must check the distribution of the volumes’ extents in that Storage Pool.

Multitiered Storage Pool

A multitiered Storage Pool has a mix of MDisks with more than one type of disk tier attribute; for example, a Storage Pool that contains a mix of flash, enterprise, and Nearline MDisks.

A multitiered Storage Pool contains MDisks with different characteristics, as opposed to the single-tiered Storage Pool. However, each tier should have MDisks of the same size and MDisks that provide the same number of extents.

A multitiered Storage Pool is used to enable automatic migration of extents between disk tiers using the IBM Storwize V3700 Easy Tier® function, as described in Chapter 9, “IBM Easy Tier” on page 417.

1.5.9 Child Pools

A Child Pool is a logical subset object similar to a Storage Pool. A Child Pool is in, and gets capacity exclusively from, one Storage Pool. A user can specify the Child Pool capacity at creation, and it can grow and decrease nondisruptively. A Child Pool inherits the properties of the parent Storage Pool, and provides most of the functions that Storage Pools have. For more information, see 7.4, “Working with Child Pools” on page 355.

1.5.10 Volumes

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

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

The following types of volumes are available:

- ▶ Striped

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

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

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

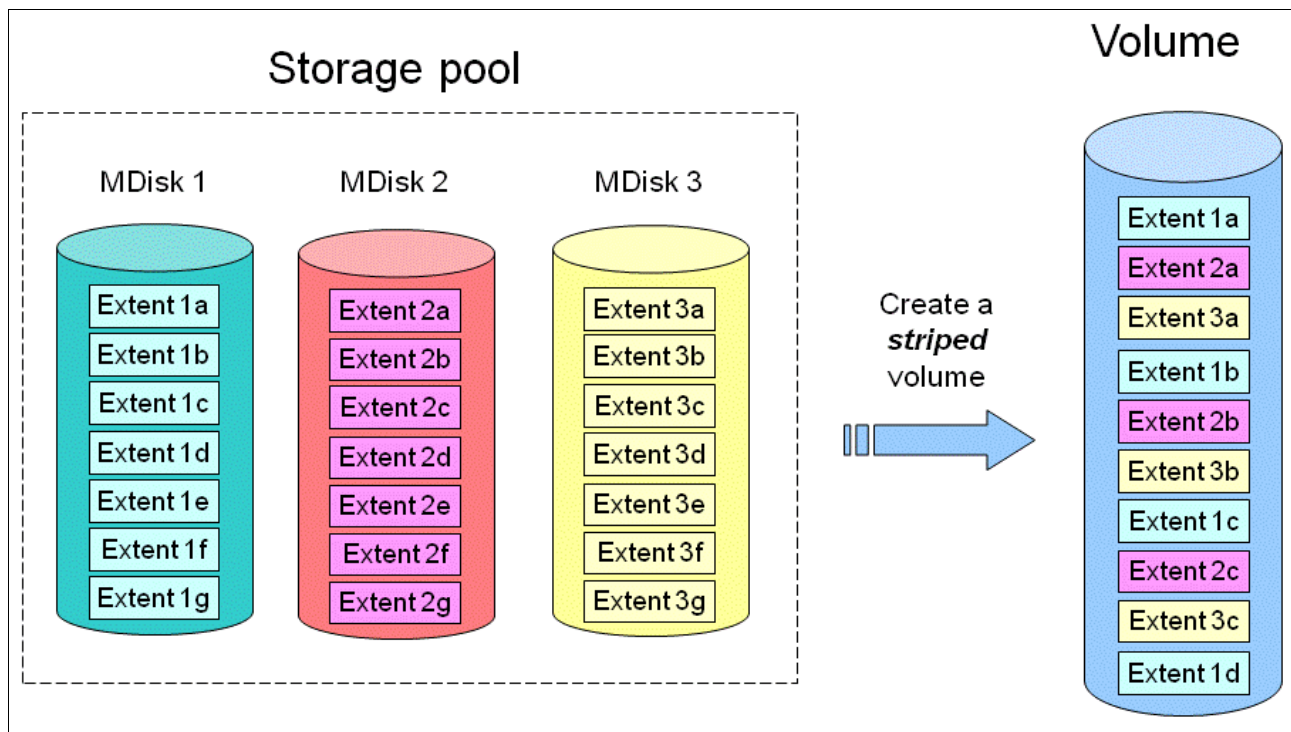


Figure 1-9 IBM Storwize V3700 striped volume

► Sequential

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

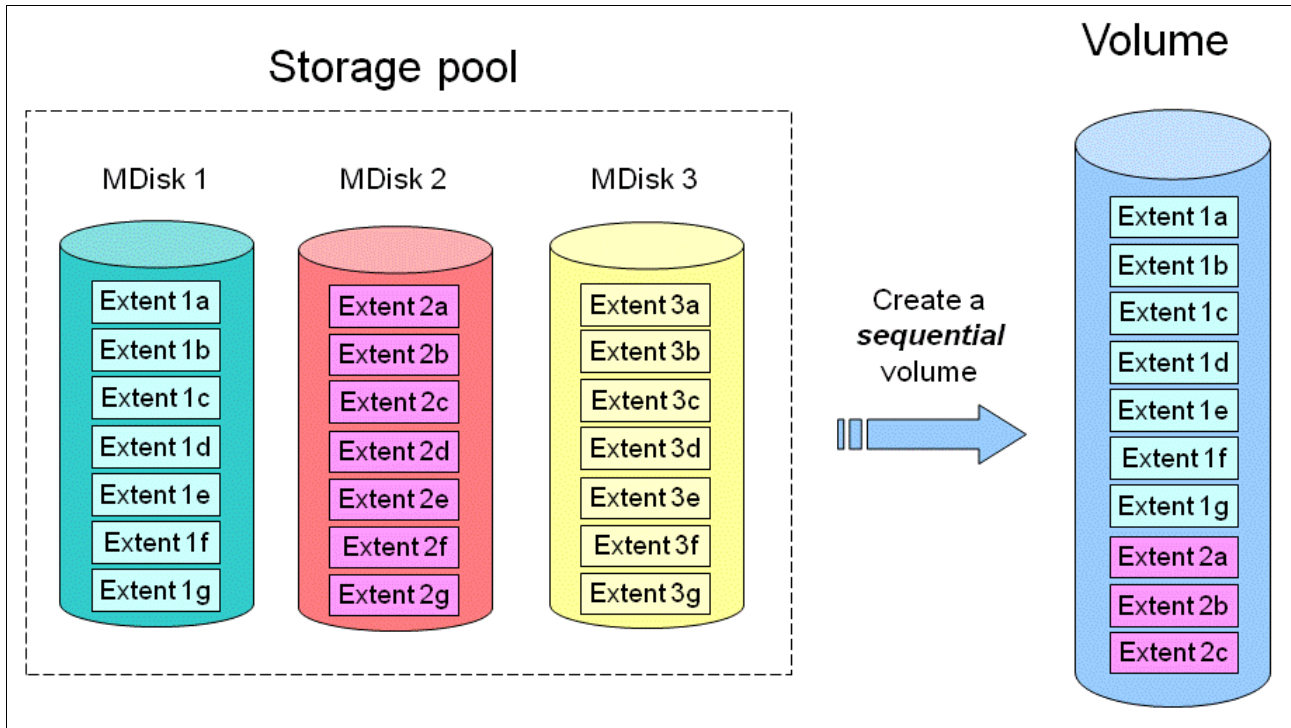


Figure 1-10 Sequential volume

► Image mode

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

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

Figure 1-11 illustrates creating in image mode volume.




Figure 1-11 Image mode volume

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

If you want to migrate data from an existing storage subsystem, use the Storage Migration wizard, which guides you through the process. For more information, see Chapter 6, “Storage migration” on page 249.

If you add an MDisk containing data to a Storage Pool, any data on the MDisk is lost. If you are presenting externally virtualized LUNs that contain data to an IBM Storwize V3700, import them using the migration wizard to ensure data integrity.

1.5.11 iSCSI

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

In the simplest terms, iSCSI enables the transport of SCSI commands and data over an IP network that is based on IP routers and Ethernet switches. iSCSI is a block-level protocol that encapsulates SCSI commands into Transmission Control Protocol/Internet Protocol (TCP/IP) packets and uses an existing IP network, rather than requiring FC host bus adapters (HBAs) and a SAN fabric infrastructure.

The following concepts of names and addresses are carefully separated in iSCSI:

- ▶ An *iSCSI name* is a location-independent, permanent identifier for an iSCSI node. An iSCSI node has one iSCSI name, which stays constant for the life of the node. The terms *initiator name* and *target name* also refer to an iSCSI name.
- ▶ An *iSCSI address* specifies the iSCSI name of an iSCSI node, and a location of that node. The address consists of a host name or IP address, a TCP port number (for the target), and the iSCSI name of the node. An iSCSI node can have any number of addresses, which can change at any time, particularly if they are assigned by way of Dynamic Host Configuration Protocol (DHCP). An IBM Storwize V3700 node represents an iSCSI node, and provides statically allocated IP addresses.

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

The IBM Storwize V3700 also supports the use of the FCoE protocol, which encapsulates the native FC frames into Ethernet frames.

1.5.12 SAS

SAS standard is an alternative method of attaching hosts to the IBM Storwize V3700. The IBM Storwize V3700 supports direct SAS host attachment that provides easy-to-use, affordable storage needs. Each SAS port device has a worldwide unique 64-bit SAS address.

1.6 IBM Storwize V3700 features

The features that are available with the IBM Storwize V3700 are described in this section.

1.6.1 Volume mirroring

IBM Storwize V3700 provides a function that is called *storage volume mirroring*, which enables a volume to have two physical copies. Each volume copy can belong to a different Storage Pool, and can be generic or thin-provisioned, which provides a high-availability (HA) solution.

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

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

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

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

1.6.2 Thin provisioning

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

- ▶ The real capacity of the volume
- ▶ The virtual capacity of the volume

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

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

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

A thin-provisioned volume can be configured to *autoexpand*, which causes the IBM Storwize V3700 to automatically expand the real capacity of a thin-provisioned volume as its real capacity is used. This parameter prevents the volume from going offline. Autoexpand attempts to maintain a fixed amount of unused real capacity on the volume. This amount is known as the *contingency capacity*. The default setting is 2%.

The contingency capacity 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 nonzero contingency capacity stays online until it is used up.

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

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

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

The thin provisioning feature is included as part of the base machine code, and no license is required.

1.6.3 Easy Tier

IBM Easy Tier provides a mechanism to seamlessly migrate hot extents to the most appropriate tier in the IBM Storwize V3700 solution. This migration might be to different tiers of internal drive within IBM Storwize V3700.

The Easy Tier function can be turned on or off at the Storage Pool and volume level. It is possible to demonstrate the potential benefit of Easy Tier in your environment before installing flash drivers by using the IBM Storage Tier Advisor Tool.

Easy Tier is described in more detail in Chapter 9, “IBM Easy Tier” on page 417.

The IBM Easy Tier feature is licensed per IBM Storwize V3700 Storage system.

1.6.4 Turbo Performance

Turbo Performance is a licensed function that provides enhanced performance for the system. A 90-day trial version of this function is available. If a license is not purchased and activated before the trial period expires, the system reverts to non-turbo performance.

The Turbo Performance feature does not require hardware changes, and is a nondisruptive upgrade.

The Turbo Performance feature is licensed per IBM Storwize V3700 Storage system.

1.6.5 Storage Migration

The IBM Storwize V3700 Storage Migration feature enables you to easily move data from other FC or SAS attached external storage to the internal capacity of the IBM Storwize V3700. Migrating data from other storage to the IBM Storwize V3700 storage system provides benefits from more functionality, such as the easy-to-use GUI, internal virtualization, thin provisioning, and Copy Services.

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

1.6.6 FlashCopy

The FlashCopy function copies a source volume to a target volume. There are three types of FlashCopy preset:

- ▶ Snapshot

The snapshot preset creates a point-in-time view of the production data. It is not intended to be an independent copy, but is used to maintain a view of the production data at the time that the snapshot is created. Therefore, it holds only the data from regions of the production volume that changed since the snapshot was created. Because the snapshot preset uses thin provisioning, only the capacity that is required for the changes is used.

- ▶ Clone

The clone preset creates an exact replica of the volume, which can be changed without affecting the original volume. After the copy completes, the mapping that was created by the preset is automatically deleted.

- ▶ Backup

The backup preset creates a point-in-time replica of the production data. After the copy completes, the backup view can be refreshed from the production data, with minimal copying of data from the production volume to the backup volume.

Reverse FlashCopy enables target volumes to become restore points for the source volume without breaking the FlashCopy relationship, and without waiting for the original copy operation to complete. IBM Storwize V3700 supports multiple targets, and therefore multiple rollback points.

The base FlashCopy feature, which requires no license, provides up to 64 mappings. An optional license is available to upgrade FlashCopy mappings up to 4096 mappings.

1.6.7 Remote Copy

Remote Copy can be maintained in one of two modes:

- ▶ Synchronous
- ▶ Asynchronous

With the IBM Storwize V3700, Metro Mirror and Global Mirror are the IBM branded terms for the functions that are synchronous Remote Copy and asynchronous Remote Copy.

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

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

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

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

Global Mirror can operate with or without cycling. When operating without cycling, write operations are applied to the secondary volume as soon as possible after they are applied to the primary volume. The secondary volume is generally less than 1 second behind the primary volume, which minimizes the amount of data that must be recovered in the event of a failover. However, this requires that a high-bandwidth link be provisioned between the two sites.

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

For more information about the IBM Storwize V3700 Copy Services, see Chapter 10, “Copy services” on page 455.

IBM Remote Copy is licensed per IBM Storwize V3700 Storage system.

Copy Services configuration limits

For the current list of these limits, see the following website:

<https://ibm.biz/BdEqMG>

1.7 Problem management and support

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

1.7.1 IBM Assist On-site and remote service

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

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

To use the IBM Assist On-site tool, the management PC that is used to manage the IBM Storwize V3700 must be able to access the Internet. For more information about this tool, see the following website:

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

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

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

1.7.2 Event notifications

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.

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

SNMP traps

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

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

Syslog messages

The syslog protocol is a standard protocol for forwarding log messages from a sender to a receiver on an IP network. The IBM Storwize V3700 can send syslog messages that notify personnel about an event. IBM Storwize V3700 can also transmit syslog messages in expanded or concise format.

You can use a syslog manager to view the syslog messages that IBM Storwize V3700 sends. IBM Storwize V3700 uses the User Datagram Protocol (UDP) to transmit the syslog message. You can use the management GUI or the CLI to configure and modify your syslog settings.

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 it is configured, this function alerts IBM service personnel about hardware failures and potentially serious configuration or environmental issues. You can use the call home function if you have a maintenance contract with IBM, or if the IBM Storwize V3700 is within the warranty period.

To send email, you must configure at least one SMTP server. You can specify as many as five other SMTP servers for backup purposes. The SMTP server must accept the relaying of email from the IBM Storwize V3700 clustered system IP address. You can then use the management GUI or CLI to configure the email settings, including contact information and email recipients.

Set the reply address to a valid email address. Send a test email to check that all connections and infrastructure are set up correctly. You can disable the Call Home function at any time by using the management GUI or CLI.

1.7.3 Useful IBM Storwize V3700 websites

For more information about the IBM Storwize V3700, see the following websites:

- ▶ The IBM Storwize V3700 home page
<https://ibm.biz/BdEqMW>
- ▶ IBM Storwize V3700 Configuration Limit and Restrictions
<https://ibm.biz/BdEqMG>
- ▶ IBM Storwize V3700 Knowledge Center
<https://ibm.biz/BdEqPS>
- ▶ IBM Redbooks publications about IBM Storwize V3700
<https://ibm.biz/BdEqSB>
- ▶ IBM e-learning tutorials about IBM Storwize V3700
<https://ibm.biz/BdEqMQ>

IBM Storwize V3700 learning videos on YouTube

Videos are available on YouTube that describe the Storwize GUI, and are available at the Uniform Resource Locators (URLs) shown in Table 1-6.

Table 1-6 Videos available on YouTube

Video description	URL
Introducing IBM Storwize V3700	http://www.youtube.com/watch?v=AePPKiXE4xM
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
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
IBM Storwize V3700 Initial Setup	http://www.youtube.com/watch?v=oj9uhTYe6gg
IBM Storwize V7000 Installation	http://www.youtube.com/watch?v=kCCFxM5ZMV4

These videos are applicable not only to IBM Storwize V3700 but also to other IBM Storwize products because the GUI interface, functions, and features are similar to all products.



Initial configuration

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

This chapter includes 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. The following checklist of requirements can be used to plan your installation:

- Install the hardware as described in *IBM Storwize V3700 Quick Installation Guide*, GC27-4219.
- For more information about planning the IBM Storwize V3700 environment, see the following website:
<https://ibm.biz/BdEzvZ>
- An appropriate 19-inch rack with 2 - 20 U of space should be available, depending on the number of enclosures to install. Each enclosure measures 2 U, and a single control enclosure with up to nine expansion enclosures constitutes an IBM Storwize V3700 system.
- There should be redundant power outlets in the rack for each of the two power cords that are included per enclosure. In all, 2 - 20 outlets are required, depending on the number of enclosures to install. The power cords conform to the International Electrotechnical Commission (IEC)-60320 (IEC320) C13/C14 standards.
- A minimum of four Fibre Channel (FC) ports that are attached to the fabric are required. However, it is a leading practice to use eight 4 gigabits per second (Gbps) or 8 Gbps FC ports.

Fibre Channel ports: FC ports are required only if you are using FC hosts. You can use the IBM Storwize V3700 with Ethernet-only cabling for Internet Small Computer System Interface (iSCSI) hosts or use serial-attached SCSI (SAS) cabling for direct attach hosts.

- You should have eight 4 Gbps, or 8 Gbps compatible FC cable drops.
- Up to six hosts can be directly connected by using SAS ports 1, 2, and 3 on each node canister, with SFF-8644 mini SAS HD cabling.
- You should have a minimum of two Ethernet ports on the LAN, with four preferred for more configuration access redundancy or iSCSI host access.
- You should have a minimum of two Ethernet cable drops, with four preferred for more configuration access redundancy or iSCSI host access. Ethernet port 1 on each node canister must be connected to the LAN, with port 2 as optional.

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

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

IP addresses: A fourth IP address should be used for backup configuration access. This other IP address enables 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.
- At least two 0.6-meter, 1.5-meter, or 3-meter SAS cables per expansion enclosure are required. The length of the cables depends on the physical rack location of the expansion enclosure relative to the control enclosure or other expansion enclosures. Locate the control enclosure so that up to nine expansion enclosures can be located, as shown in Figure 2-1. The IBM Storwize V3700 supports one external SAS chain using SAS port 4 on the control enclosure node canisters.



Figure 2-1 Connecting the SAS expansion cables example

The following connections must be made:

- Connect SAS port 4 of the left node canister in the control enclosure to SAS port 1 of the left expansion canisters in the first expansion enclosure.
- Connect SAS port 4 of the right node canister in the control enclosure to SAS port 1 of the right expansion canisters in the first expansion enclosure.
- Connect SAS port 2 of the left node canister in the first expansion enclosure to SAS port 1 of the left expansion canister in the second expansion enclosure.
- Connect SAS port 2 of the right node canister in the first expansion enclosure to SAS port 1 of the right expansion canister in the second expansion enclosure.

Continue in this fashion, adding expansion controllers on the SAS chains originating at port 4 on the control enclosure node canisters.

Disk drives: The disk drives that are included with the control enclosure (model 2072-12C or 2072-24C) are part of the single SAS chain. The expansion enclosures should be connected to the SAS chain as shown in Figure 2-1 on page 29 so that they can use the full bandwidth of the system.

2.2 Storage area network configuration planning

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

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

Zoning must be implemented after the IBM Storwize V3700, hosts, and optional external storage systems are connected to the SAN fabrics.

To enable the node canisters to communicate with each other in band, create a zone with only the IBM Storwize V3700 worldwide port names (WWPNs), two from each node canister, on each of the two fabrics. If there is an external storage system that is to be virtualized, create a zone in each fabric with the IBM Storwize V3700 WWPNs (two from each node canister) with up to a maximum of eight WWPNs from the external storage system.

Assuming that every host has a FC 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. The critical point is that there should only ever be one initiator (host bus adapter) in any zone. For load balancing between the node ports on the IBM Storwize V3700, alternate the host FC ports between the ports of the IBM Storwize V3700.

There should be a maximum of eight paths through the SAN from each host to the IBM Storwize V3700. Hosts where this number is exceeded are not supported. The restriction is there to limit the number of paths that the multipathing driver must resolve. A host with only two host bus adapters (HBAs) should not exceed this limit with proper zoning in a dual-fabric SAN.

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

Figure 2-2 shows how to cable devices to the SAN. Refer to this example as 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 and from the clustered system, as shown in the following examples in Figure 2-2:

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

Similar zones should be created for all other hosts with volumes that are on the IBM Storwize V3700.

Verify interoperability with which the IBM Storwize V3700 connects to SAN switches or directors by starting at the IBM System Storage Interoperation Center (SSIC) website:

<https://ibm.biz/BdxQhe>

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

Important: The IBM Storwize V3700 port login maximum that is listed in the restriction document must not be exceeded. The document is available at the following website:

<http://www.ibm.com/support/docview.wss?uid=ssg1S1004921>

Connectivity issues: If you have any connectivity issues between IBM Storwize V3700 ports and Brocade SAN Switches or Directors at 8 Gbps, see the following website for the correct setting of the `fillword` port config parameter in the Brocade operating system:

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

IBM Storwize V3700 can be used with a direct-attach FC host configuration. The suggested configuration for direct attachment is to have at least one FC cable from the host connected to each node of the IBM Storwize V3700 to provide redundancy in the event that one of the nodes goes offline, as shown in Figure 2-3.




Figure 2-3 FC direct-attach host configuration

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

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

2.3 SAS direct attach planning

There are three SAS ports (ports 1, 2, and 3) per node canister that are available for direct host attach on an IBM Storwize V3700. Do not use port 4, because it is reserved for expansion enclosure connectivity only. See Figure 2-4 on page 33 to correctly identify ports 1, 2, and 3. Also, note the keyway in the top of the SAS connector.

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




Figure 2-4 SAS port identification

Although it is possible to attach six hosts, one to each of the three available SAS ports on the two node canisters, the suggested configuration for direct attachment is to have at least one SAS cable from the host connected to each node of the IBM Storwize V3700. This provides redundancy in the event that one of the nodes goes offline, as shown in Figure 2-5.




Figure 2-5 SAS host direct attach

Note: It is possible to connect further hosts using the optional SAS adapter interface card.

2.4 LAN configuration planning

There are two Ethernet ports per node canister that are available for connection to the LAN on an IBM Storwize V3700 system.

Ethernet port 1 is for accessing the management graphical user interface (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 Internet Protocol (IP)v4 or IPv6 addresses on an Ethernet port, or the use of the same Ethernet port for management and iSCSI access. However, you cannot use the same IP address for management and iSCSI host use.

Table 2-1 shows possible IP configuration of the Ethernet ports on the IBM Storwize V3700 system.

Table 2-1 IBM Storwize V3700 IP address configuration options per node canister

Storwize V3700 Management Node Canister 1		Storwize V3700 Partner Node Canister 2	
IPv4/6 management address	ETH PORT 1	IPv4/6 service address	ETH PORT 1
IPv4/6 service address		IPv4/6 iSCSI address	
IPv4/6 iSCSI address			
IPv4/6 management address	ETH PORT 2	IPv4/6 iSCSI address	ETH PORT 2
IPv4/6 iSCSI address			

IP management addresses: The IP management address that is shown on Node Canister 1 in Table 2-1 is an address on the configuration node. 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 (planned or unplanned), the management IP address fails over to the other node's Ethernet port 1.

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

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



Figure 2-6 Ethernet ports available for configuration

2.4.2 Service IP address considerations

Ethernet port 1 on each node canister is used for system management, and for service access when required. In normal operation, the service IP addresses are not needed. However, 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 2-7 shows a logical view of the Ethernet ports that are available for configuration of the service IP addresses. Only port one on each node can be configured with a service IP address.



Figure 2-7 Service IP addresses available for configuration

2.5 Host configuration planning

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

This is also true for direct-attach SAS hosts. They can be connected by using a single host port that enables up to six hosts to be configured. However, two SAS connections per host are advised for redundancy. If two connections per host are used, a multipath driver is also required on the host.

If iSCSI host is to be used, they also require a Microsoft Multipath I/O (MPIO) driver. Both node canisters should be configured and connected to the network so that any iSCSI hosts see at least two paths to volumes, and an MPIO driver is required to resolve these.

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 website:

<https://ibm.biz/BdEzpm>

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

<https://ibm.biz/BdEzpv>

Multiple operating systems (OSs) are supported by IBM Storwize V3700. For more information about HBA/driver/multipath combinations, see the following website:

<https://ibm.biz/BdxQhe>

As indicated on the SSIC, keep the following items under consideration:

- ▶ Host OSs are at the levels that are supported by the IBM Storwize V3700.
- ▶ HBA basic input/output system (BIOS), device drivers, firmware, and multipathing drivers are at the levels that are supported by IBM Storwize V3700.
- ▶ If a start from SAN is required, ensure that it is supported for the OSs that are deployed.
- ▶ If host clustering is required, ensure that it is supported for the OSs that are deployed.
- ▶ All direct connect hosts should have the HBA set to point-to-point.

For more information about host configuration, see Chapter 4, “Host configuration” on page 149.

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 to complete the initial configuration.

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

- Document the LAN NTP server IP address that is used for synchronization of devices.
- For alerts to be sent to storage administrators, and to set up Call Home to IBM for service and support, you need the following information:
 - Name of primary storage administrator for IBM to contact, if necessary.
 - Email address of the storage administrator for IBM to contact, if necessary.
 - Phone number of the storage administrator for IBM to contact, if necessary.
 - Physical location of the IBM Storwize V3700 system for IBM service (for example, Building 22, first floor).
 - Simple Mail Transfer Protocol (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: `callhome1@de.ibm.com` for Americas-based systems and `callhome0@de.ibm.com` for the rest of the world.
 - Email address of local administrators that must be notified of alerts.
 - IP address of Simple Network Management Protocol (SNMP) server to direct alerts to, if required (for example, operations or Help desk).

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

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

The user in the Security Admin role can perform any function on the IBM Storwize V3700.

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

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

The only other role that is available is the Service role, which is used if you create a user ID for the IBM service representative. This user role enables IBM service personnel to view anything on the system (as with the Monitor role). It also enables them to perform service-related commands, such as adding a node back to the system after it is serviced, or including disks that were excluded. For more information about creating users, see Chapter 3, “Graphical user interface overview” on page 71.

2.7 System management

The GUI is used to configure, manage, and troubleshoot the IBM Storwize V3700 system. It is used primarily to configure Redundant Array of Independent Disks (RAID) and logical drives, assign logical drives to hosts, replace and rebuild failed disk drives, and expand the logical drives.

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

The GUI also offers advanced functions, such as FlashCopy, Volume Mirroring, Remote Mirroring, and IBM EasyTier. A command-line interface (CLI) for the IBM Storwize V3700 system is also available.

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

2.7.1 Graphical user interface

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

<https://ibm.biz/BdEzp7>

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

1. Browse to one of the following locations:
 - a. `http(s)://<host name of your cluster>/`
 - b. `http(s)://<cluster IP address of your cluster>/` (for example, `https://192.168.70.120`)
2. Use the following default login information:
 - User ID: superuser
 - Password: passw0rd (with a zero in place of the “o”)

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

<https://ibm.biz/BdEz8c>

After the initial configuration that is described in 2.9, “Initial configuration” on page 49 is completed, the IBM Storwize V3700 Welcome window opens, as shown in Figure 2-8.

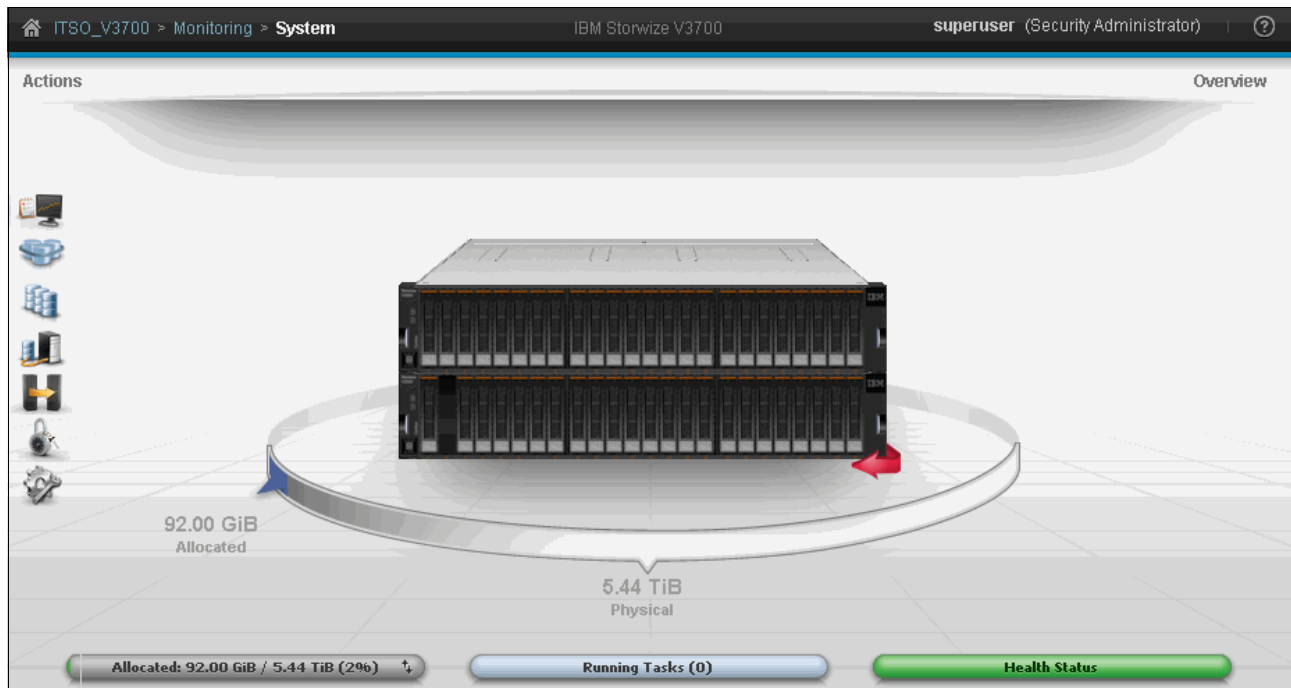


Figure 2-8 Setup wizard: Welcome window

2.7.2 Command-line interface

The CLI is a flexible tool for system management that uses the SSH protocol. A public/private SSH key pair is optional for Secure Shell (SSH) access. For more information about setting up SSH access for Windows, Linux, or UNIX systems, see Appendix A, “Command-line interface setup and SAN Boot” on page 649.

The storage system can be managed by using the CLI, as shown in Example 2-1.

Example 2-1 System management using the CLI

```
IBM_Storwize:ITS0_V3700:superuser>svcinfolenclosurecanister
enclosure_id canister_id status type      node_id node_name
1             1             online node    1       node1
1             2             online node    2       node2
2             1             online expansion
2             2             online expansion
IBM_Storwize:ITS0_V3700:superuser>
```

The initial IBM Storwize V3700 system setup should be done using the process and tools that are described in 2.8, “First-time setup”.

2.8 First-time setup

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

IBM Storwize V3700 has an easy-to-use initial setup process that is contained in a Universal Serial Bus (USB) key. The USB key is delivered with each storage system and contains the initialization application file, `InitTool.exe`. The tool is configured with your IBM Storwize V3700 system management IP address, the subnet mask, and the network gateway address by first plugging the USB stick into a Windows or Linux system.

The IBM Storwize V3700 starts the initial setup when you plug 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 that is supplied with the IBM Storwize V3700, you can use any USB key that you have and download and copy the `initTool.exe` application from the following IBM Storwize V3700 Support website:

<https://ibm.biz/BdEz8s>

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

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

We use Windows in the following examples.

Complete the following steps to complete the initial setup using the USB key:

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



Figure 2-9 System Initialization: Welcome window

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

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

```
sh InitTool.sh
```

2. Select **Create a new system**, as shown in Figure 2-10.




Figure 2-10 System Initialization: Create a new system

There are other options available through the Tasks section. However, these are generally only required after initial configuration. There is the option to **Reset the superuser password** or **Set the service IP address** of a node canister. Selecting **Create a new system** and clicking **Next** (as shown in Figure 2-10) progresses through the initial configuration of the IBM Storwize V3700.

3. Set the Management IP address, as shown in Figure 2-11.




Figure 2-11 System Initialization: Management IP

- Click **Apply and Next** to show the IBM Storwize V3700 power on instructions, as shown in Figure 2-12.

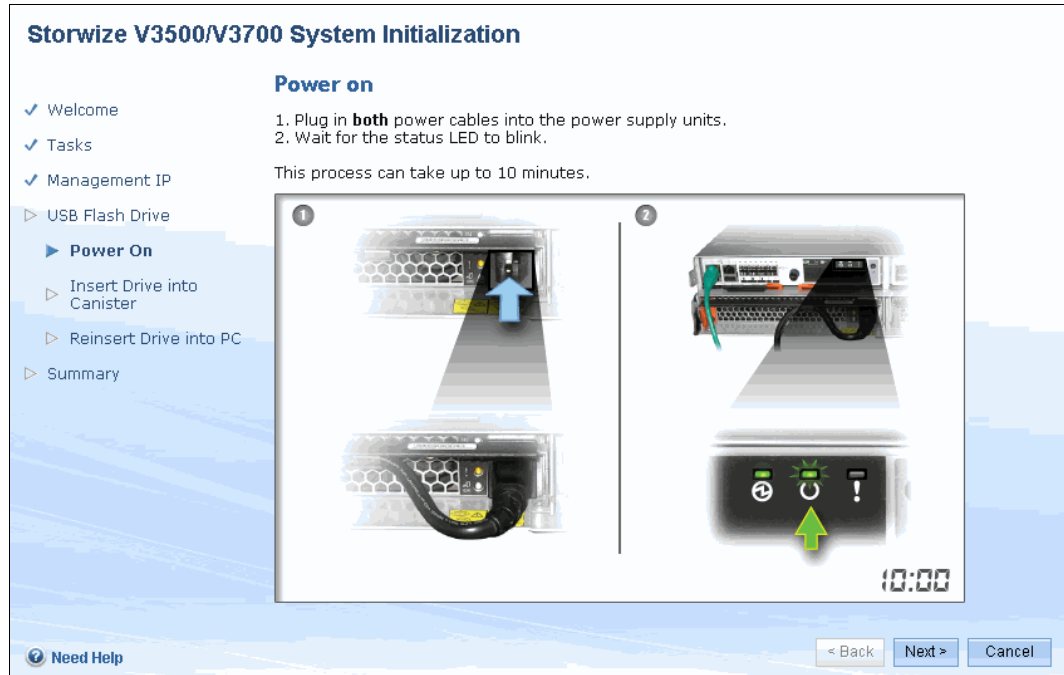


Figure 2-12 Initialization application: Storwize V3700 Power up

- Any expansion enclosures that are part of the system should be powered on and allowed to come ready before the control enclosure. Follow the instructions to power on the IBM Storwize V3700 and wait for the status light-emitting diode (LED) to flash.
- Insert the USB stick in one of the USB ports on the left node canister, as shown in Figure 2-13.



Figure 2-13 Inserting the USB flash drive into the canister

The left node canister becomes the control node, and the other node is the partner node. The fault LED begins to flash.

Clustered system creation: While the clustered system is created, the amber fault LED on the node canister flashes. When this LED stops flashing, remove the USB key from IBM Storwize V3700 and insert it in your system to check the results. This process can take up to 3 minutes to complete.

7. When it stops, return the USB stick to the personal computer (PC), as shown in Figure 2-14.




Figure 2-14 Returning the USB stick to a PC

8. The wizard then attempts to verify connectivity to the IBM Storwize V3700, as shown in Figure 2-15.

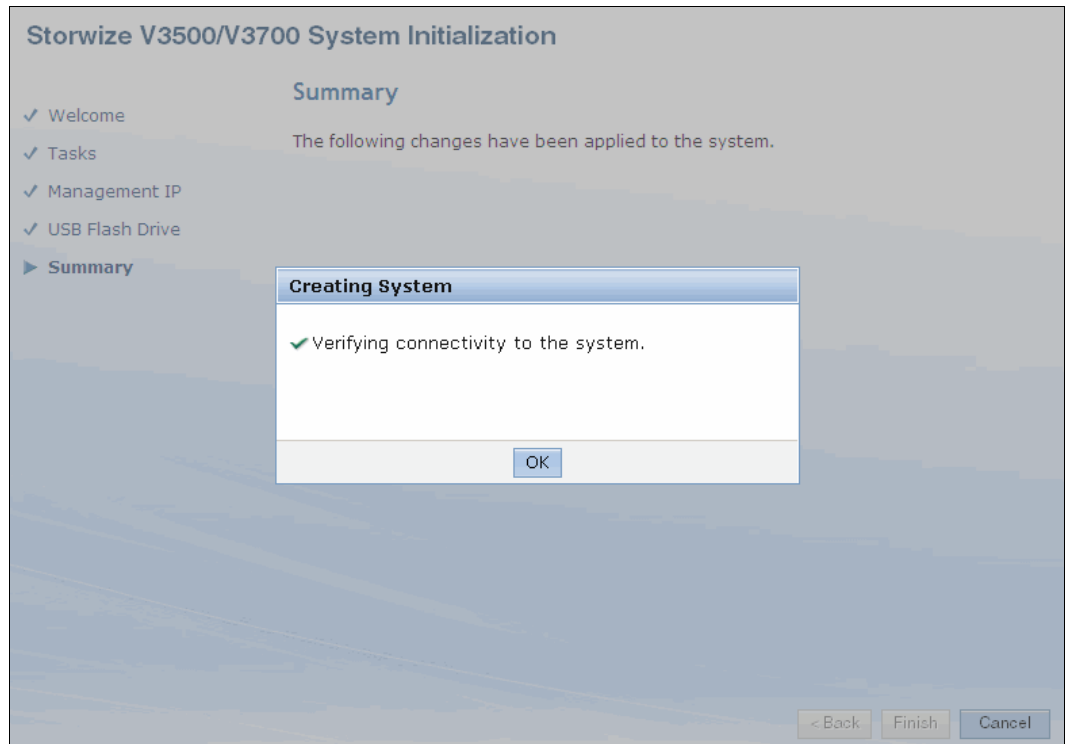


Figure 2-15 Verify system connectivity

9. If successful, a summary page is displayed that shows the settings that were applied to the IBM Storwize V3700, as shown in Figure 2-16.




Figure 2-16 System Initialization Summary

10. If the connectivity to the IBM Storwize V3700 cannot be verified, the wizard displays the error message shown in Figure 2-17.

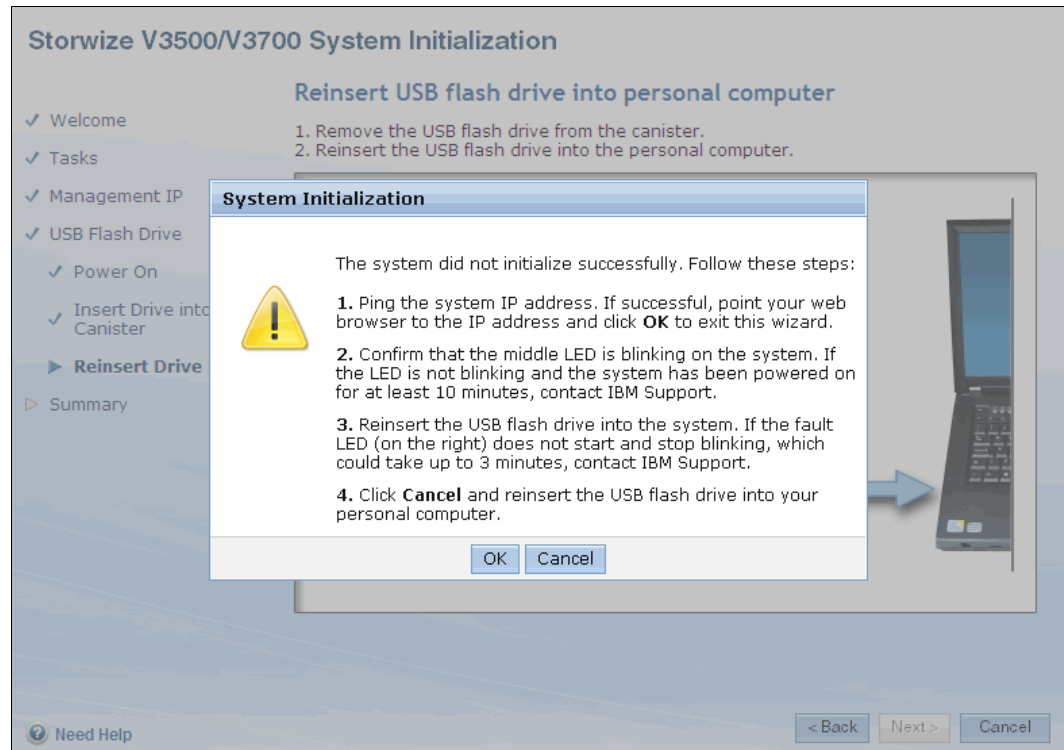


Figure 2-17 Initialization failure

11. Follow the instructions to resolve any issues. The wizard assumes that the system that you are using can connect to the IBM Storwize V3700 through the network. If it is not, you must follow step 1 from a machine that does have network access to the IBM Storwize V3700. After the initialization process completes successfully, click **Finish**.

12. The initial setup is now complete. If you have a network connection to the IBM Storwize V3700, the wizard shows the initialization completed successfully, as shown in Figure 2-18.

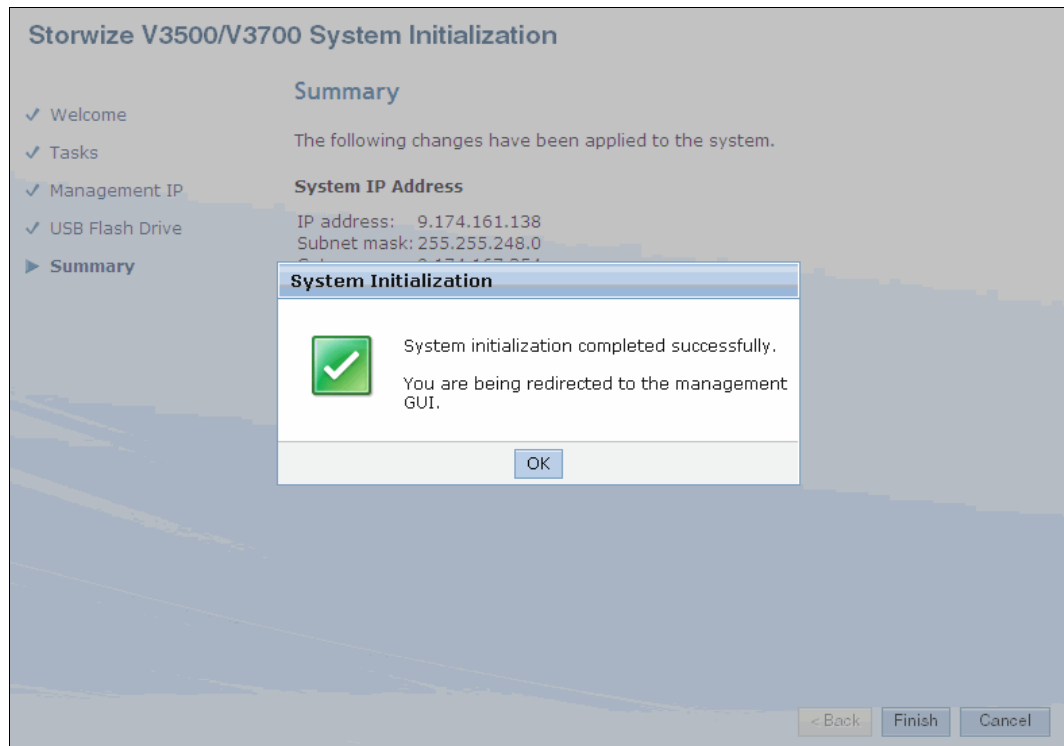


Figure 2-18 System Initialization complete

13. Click **OK** and the wizard automatically redirects you to the IBM Storwize V3700 GUI, and you are prompted to enter a new password for the superuser account, as shown in Figure 2-19. Choose the password and click **Log in** to continue the initial setup of the IBM Storwize V3700. For more information, see 2.9, “Initial configuration” on page 49.



Figure 2-19 First log on

2.9 Initial configuration

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

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

If you just completed the initial setup, you must log on to the IBM Storwize V3700 GUI. Otherwise, complete the following steps to complete the initial configuration process:

1. Start the configuration wizard using a web browser on a workstation and point it to the system management IP address that was defined in step 3 on page 42. Enter the superuser password set in step 13 on page 48, as shown in Figure 2-20.



Figure 2-20 Setup wizard: Log in

2. After you are logged in, a welcome window opens, as shown in Figure 2-21.

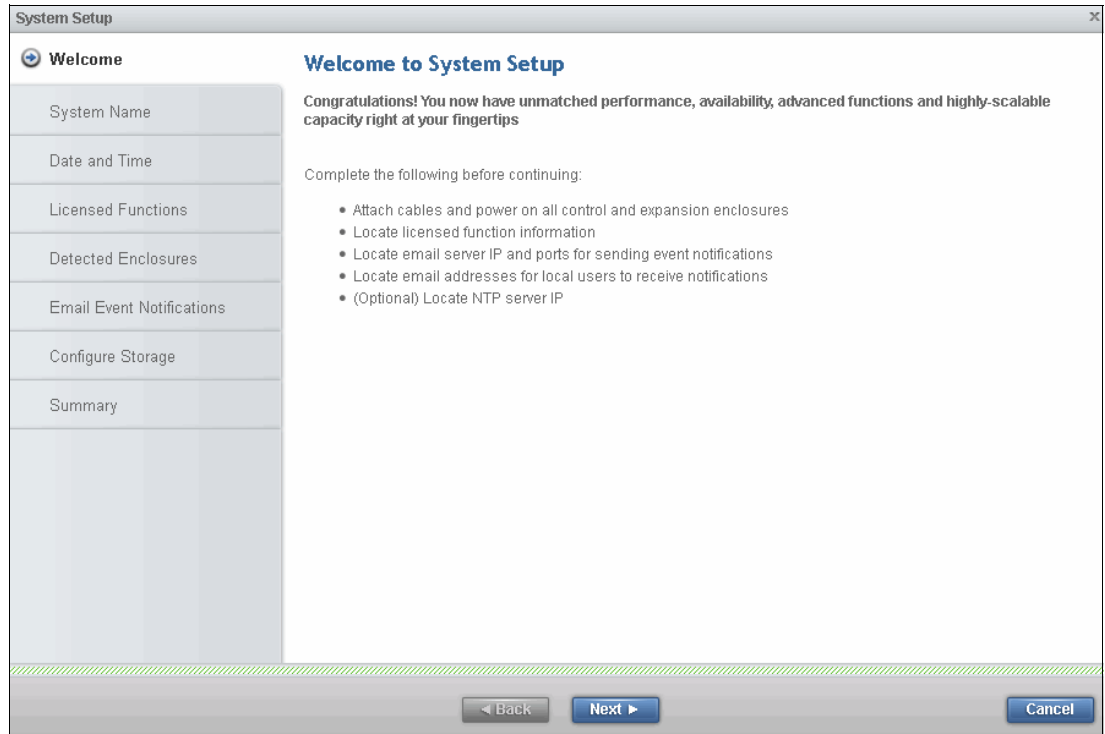


Figure 2-21 GUI Welcome window

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

4. The next window that is opened is the System Name. Set up the system name and click **Apply and Next**, as shown in Figure 2-22.

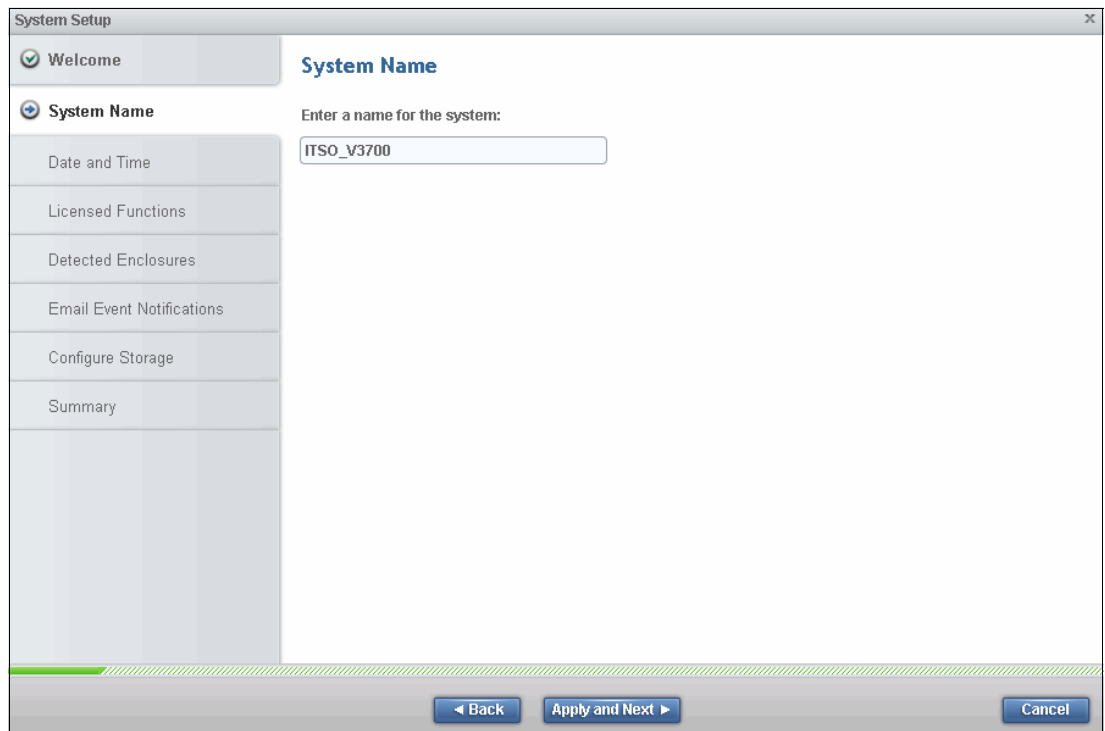


Figure 2-22 IBM Storwize V3700 System Name setup window

- The IBM Storwize V3700 shows the CLI as you go through the configuration steps. Click **Close** when the tasks complete, as shown in Figure 2-23.




Figure 2-23 IBM Storwize V3700 tasks

- There are two options for configuring the date and time. Select the required method and enter the date and time manually, or specify a network address for an NTP server. After this is done, the **Apply and Next** option becomes active, as shown in Figure 2-24.




Figure 2-24 Initial setup: Date and Time

7. The next window of the system setup that is opened is the Licensed Functions window, as shown in Figure 2-25. In our example, the functions Remote Copy and Turbo Performance are licensed. It is possible to activate licensed functions now or later.

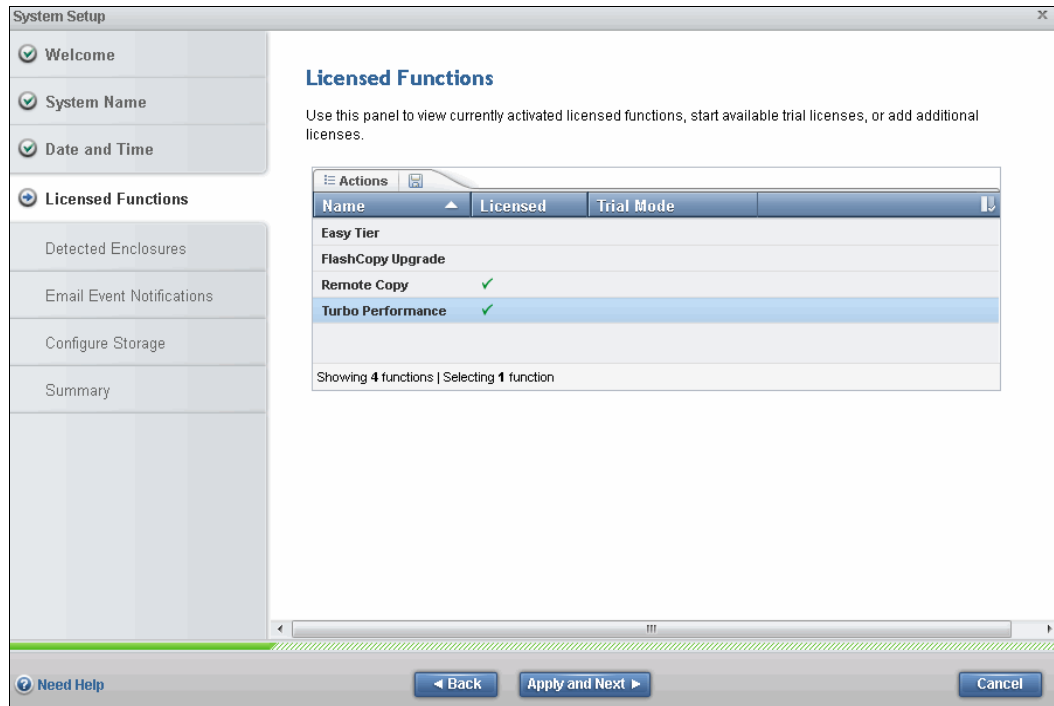


Figure 2-25 Initial setup: Licensed Functions

To learn more about the Licensed Functions of the IBM Storwize V3700 see Chapter 3, “Graphical user interface overview” on page 71.

8. The configuration wizard continues with the hardware configuration. Verify the hardware, as shown in Figure 2-26.

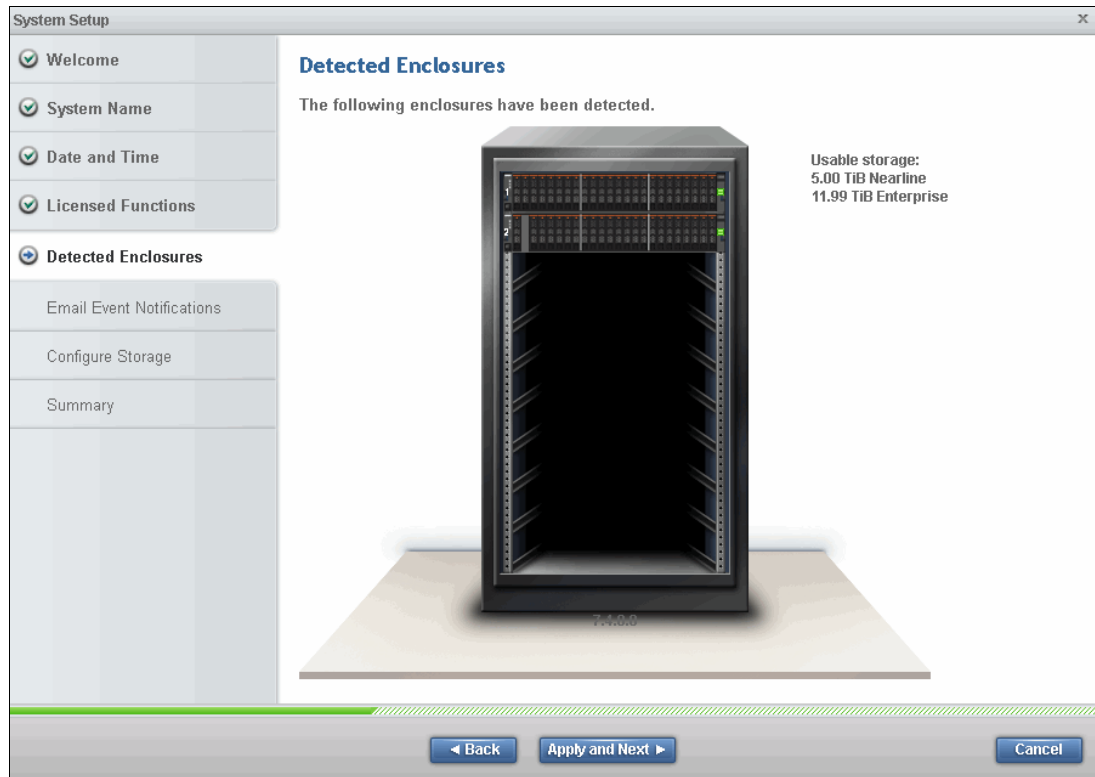


Figure 2-26 Setup wizard: Verify the detected hardware

9. Click **Apply and Next**.

10. The next window in the initial system setup is setting up **Email Event Notifications**. Select **Yes** and click **Next** to set up email notification and call home, as shown in Figure 2-27.

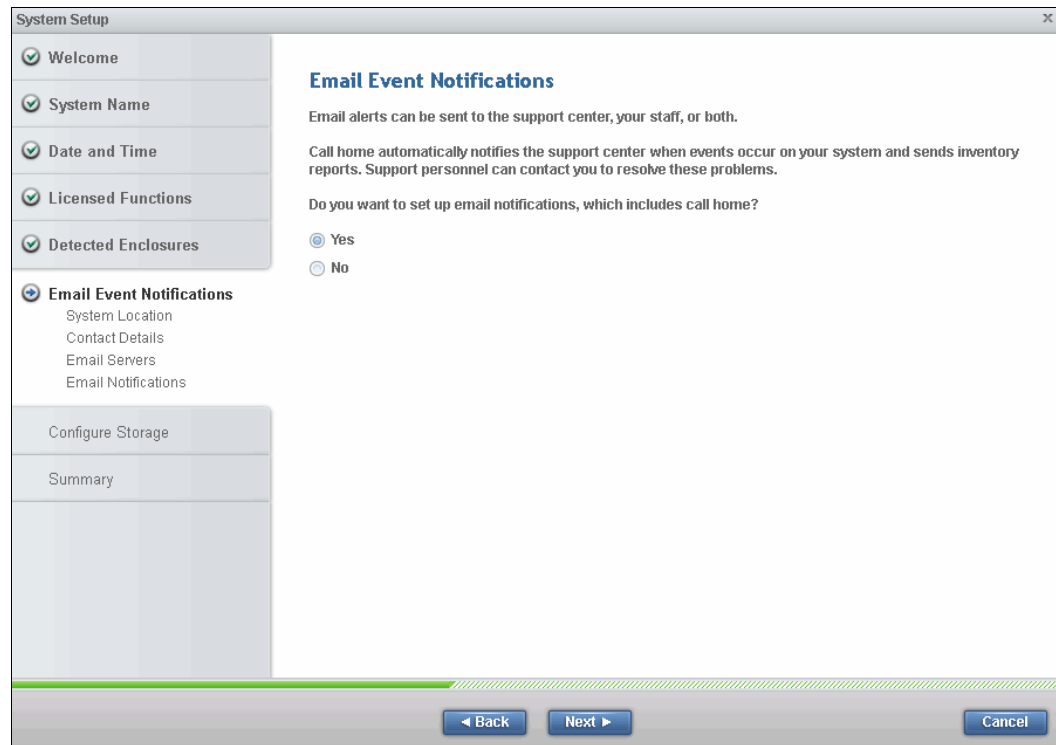


Figure 2-27 Call home setup

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

Call Home: When Call Home is configured, the IBM Storwize V3700 automatically creates a Support Contact with one of the following email addresses, depending on country or region of installation:

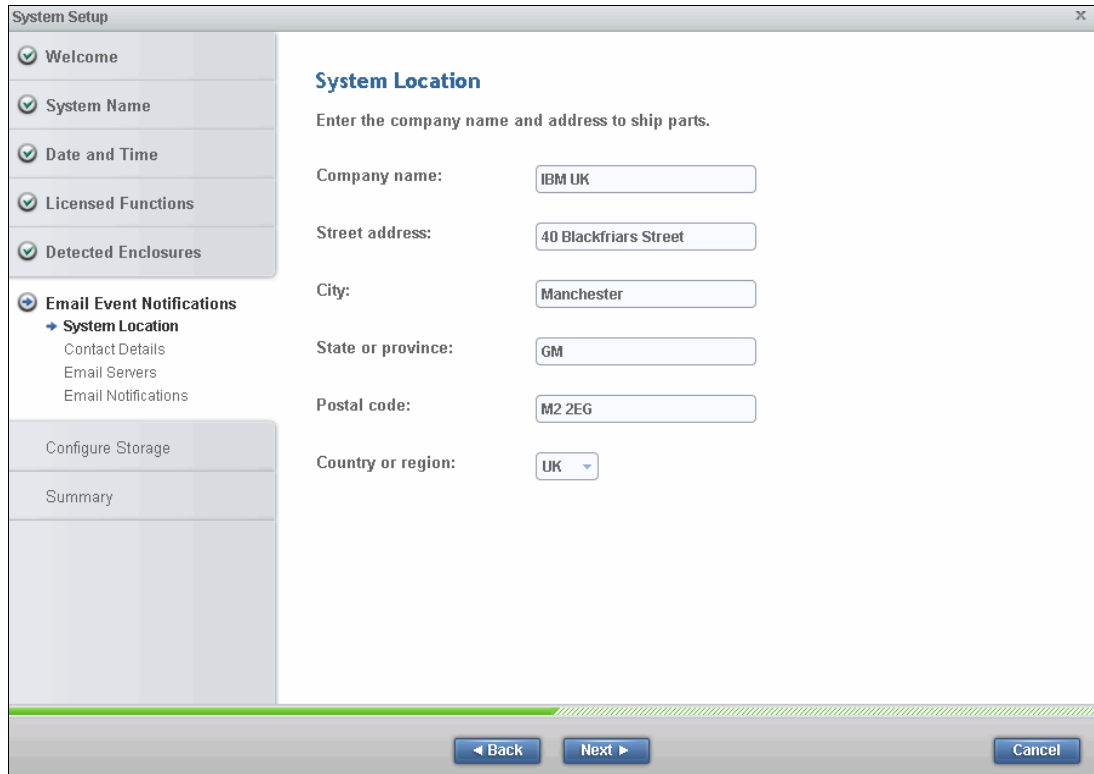
- ▶ US, Canada, Latin America, and Caribbean Islands:
callhome1@de.ibm.com
- ▶ All other countries or regions:
callhome0@de.ibm.com

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

To set up Call Home, you need the location details of the IBM Storwize V3700, Storage Administrators details, and at least one valid SMTP server IP address. If you do not want to configure Call Home, it is possible to do it later by using the GUI option **Settings** → **Event Notification** (for more information, see 2.9.2, “Configure Call Home, email alert, and inventory” on page 65).

If your system is under warranty or you have a hardware maintenance agreement to enable pro-active support of the IBM Storwize V3700, it is advised that Call Home be configured.

11. Select **Yes** and click **Next** to move to the window in which you can enter the location details, as shown in Figure 2-28.



The screenshot shows a 'System Setup' window with a sidebar on the left and a main content area on the right. The sidebar contains a list of steps: Welcome, System Name, Date and Time, Licensed Functions, Detected Enclosures, Email Event Notifications (expanded), System Location (selected), Contact Details, Email Servers, Email Notifications, Configure Storage, and Summary. The main content area is titled 'System Location' and includes the instruction 'Enter the company name and address to ship parts.' Below this are several input fields: 'Company name' (IBM UK), 'Street address' (40 Blackfriars Street), 'City' (Manchester), 'State or province' (GM), 'Postal code' (M2 2EG), and 'Country or region' (UK). At the bottom of the window are three buttons: 'Back', 'Next', and 'Cancel'.

Figure 2-28 Location details

These details are shown on the Call Home data to enable IBM Support to correctly identify where the IBM Storwize V3700 is located.

12. The next window enables you to enter the contact details of the main storage administrator, as shown in Figure 2-29. You can choose to enter the details for a 24-hour operations desk. These details are also sent with any Call Home, which enables IBM Support to contact the correct people quickly to process any issues.

The screenshot shows a 'System Setup' window with a sidebar on the left and a main content area on the right. The sidebar contains a list of setup steps: 'Welcome', 'System Name', 'Date and Time', 'Licensed Functions', 'Detected Enclosures', 'Email Event Notifications' (expanded), 'System Location', 'Contact Details' (selected), 'Email Servers', and 'Email Notifications'. Below these are 'Configure Storage' and 'Summary' buttons. The main content area is titled 'Contact Details' and includes a descriptive text: 'Enter the name and contact information for the person in your organization that the support center can contact to help resolve problems on the system.' Below this text are five input fields: 'Contact name:' with the value 'Oscar Wellington', 'Email address:' with 'owell@ibm.com', 'Telephone (primary):' with '7786100900', 'Telephone (alternate):' (empty), and 'Comment:' (empty). At the bottom of the window are three buttons: 'Back', 'Apply and Next', and 'Cancel'.

Figure 2-29 Contact Details

The next window shows the details of the email server.

13. Click the green plus (+) icon to enter more than one email server, as shown in Figure 2-30.

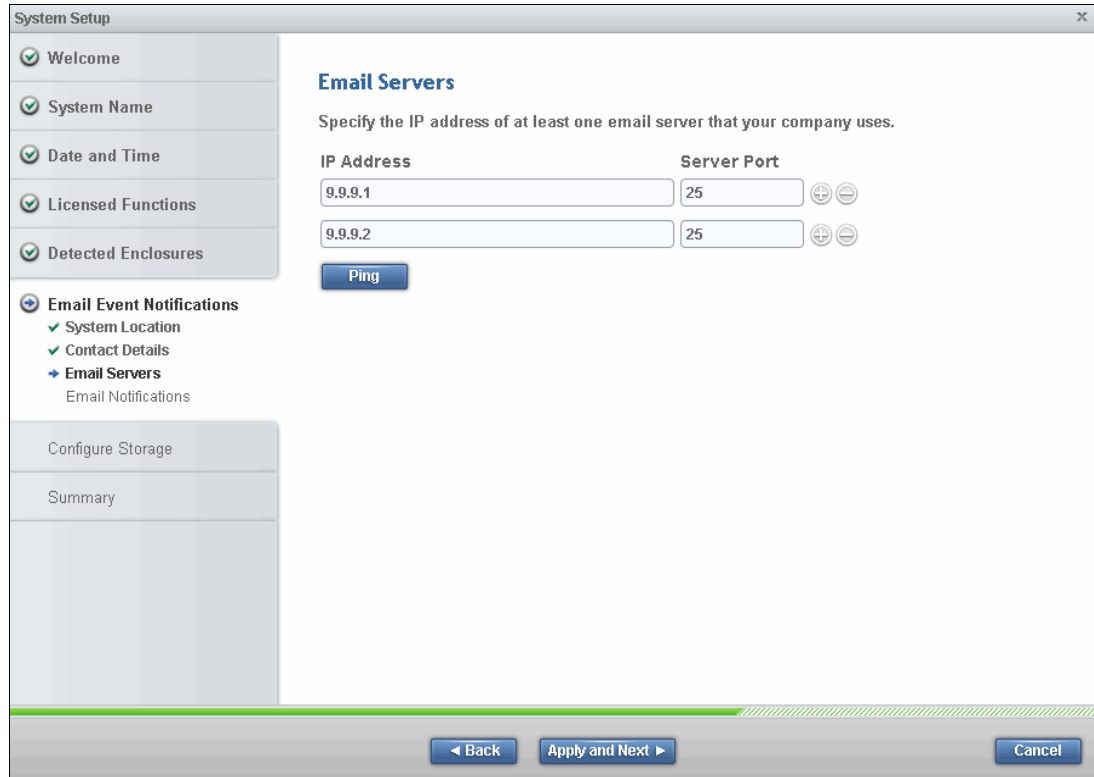


Figure 2-30 Email Servers details

14. The initial setup wizard opens the Configure Storage option next, as shown in Figure 2-31.

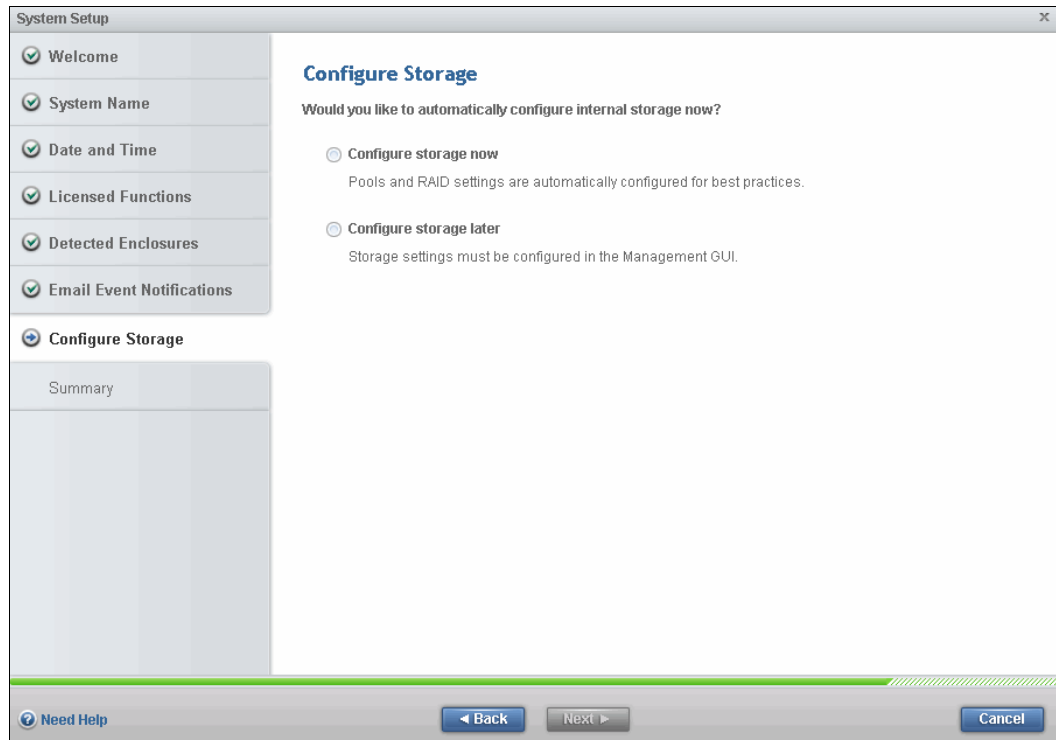


Figure 2-31 System Setup: Configure Storage

This option takes all the disks in the IBM Storwize V3700 and automatically configures them into optimal arrays for use as managed disks (MDisks). If you do not want to automatically configure disks now, select **Configure storage later** and you are taken to the summary window. If you select **Configure storage now**, the system examines the enclosures and provides the best array configuration.

15. Select **Configure storage now**, and then click **Next** to move to the summary window that shows the RAID configuration that the IBM Storwize V3700 implements, as shown in Figure 2-32.

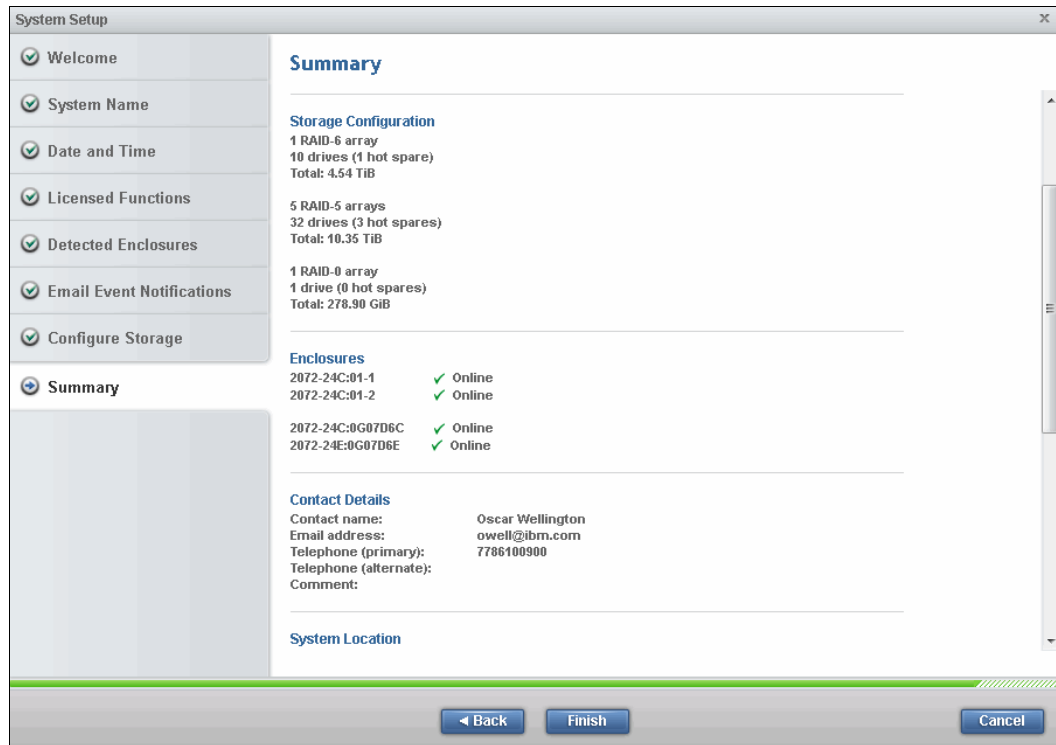


Figure 2-32 System Setup: Summary window

The arrays and the Storage Pool are created when you click **Finish**. Depending on the disks available, this process might take some time to complete as a background task.

16. Closing the task box completes the Initial Configuration wizard, and the IBM Storwize V3700 GUI directs you to the Suggested Tasks, as shown in Figure 2-33.

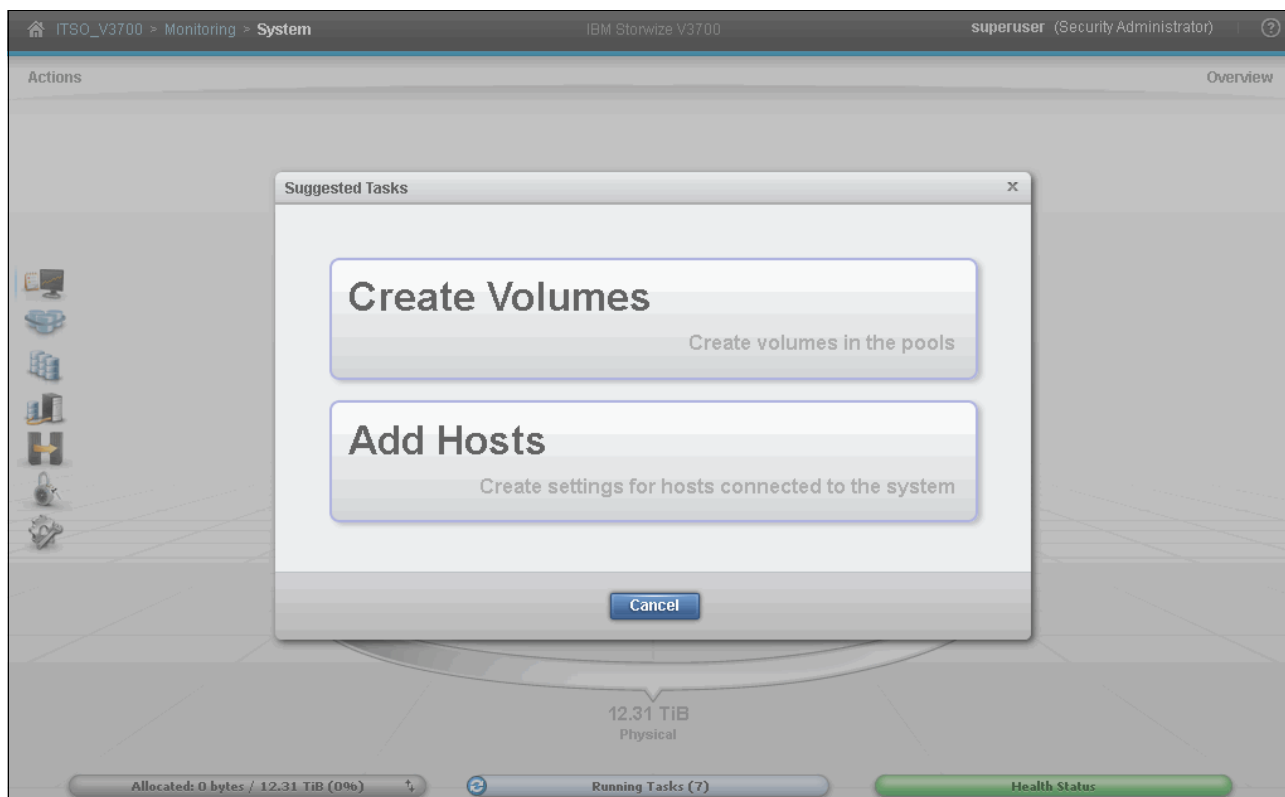


Figure 2-33 IBM Storwize V3700 Suggested Tasks

If you choose to create volumes at this stage, see Chapter 5, “Volume configuration” on page 199.

To work with hosts, see Chapter 4, “Host configuration” on page 149.

17. Selecting **Cancel** exits to the IBM Storwize V3700 GUI. There is also a shortcut link that redirects to a useful web portal about the IBM Storwize V3700.

2.9.1 Adding expansion enclosures after initial configuration

When the initial installation of the IBM Storwize V3700 is complete, all expansion enclosures that were purchased at that time should be installed as part of the initial configuration. This process enables the system to make the best use of the enclosures and drives that you have.

If you are expanding the IBM Storwize V3700 after the initial installation by adding expansion enclosures, follow the physical installation procedures as described in *IBM Storwize V3700 Quick Installation Guide*, GC27-4219. After the hardware is installed, cabled, and powered on, complete the following steps to add an expansion enclosure:

1. Change to the Monitoring tab and you can find two different options to initiate the wizard to add a new enclosure. Each guides the user to the same wizard.

2. Click **Add Enclosure** or **Click to add** to add the new enclosure, as shown in Figure 2-34.

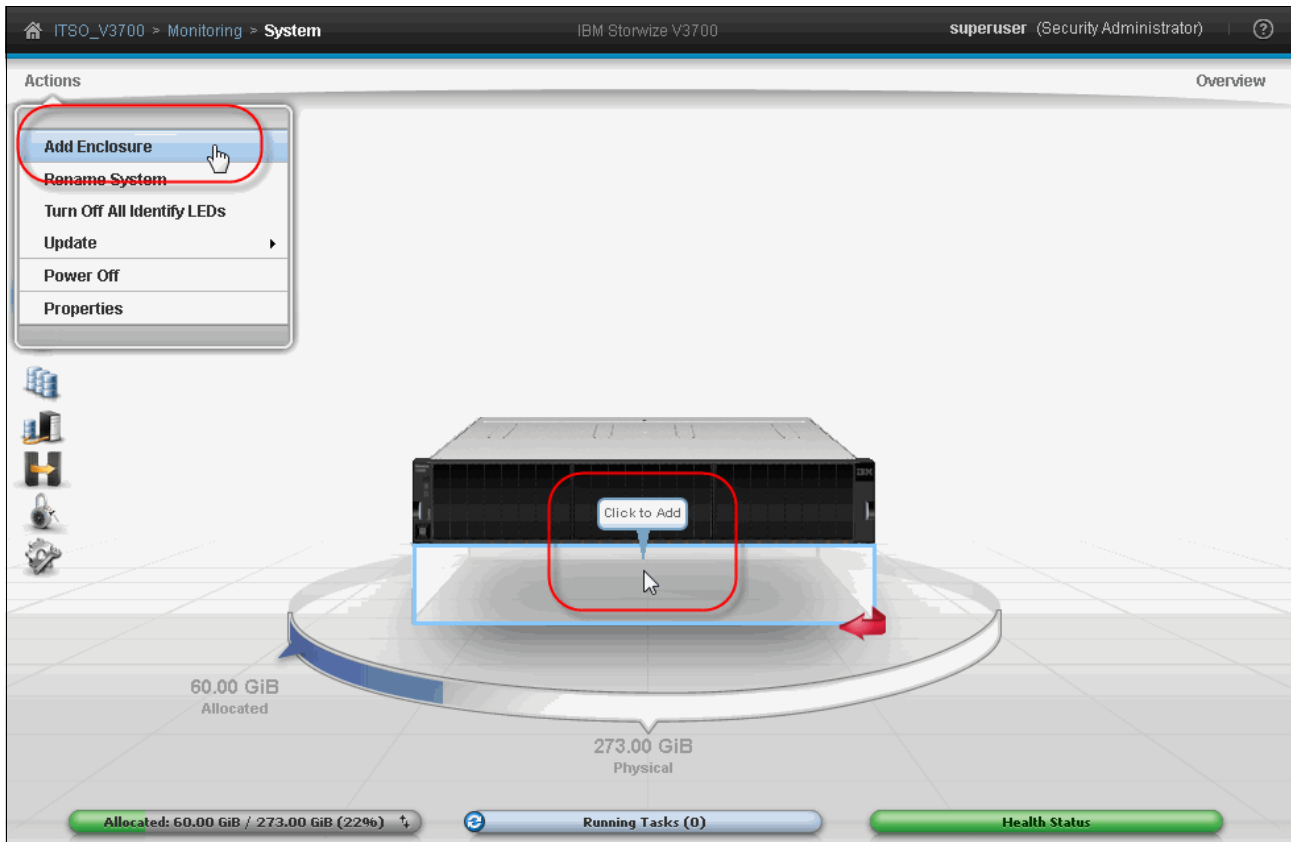


Figure 2-34 Adding new expansion enclosure

3. A message displays that informs you that there are unused drives in the new expansion enclosure. The wizard prompts you to include them into optimal RAID configurations to be used as MDisks. Choose whether you want to configure the new drives. In our example, we choose **NO**. However, if you choose **Yes**, the IBM Storwize V3700 automatically configures the drives as candidates. To learn more about the status of the drives, see Chapter 7, “Storage Pools” on page 307.

4. The wizard shows the details of hardware to be added. Highlight the enclosure that you want to add to your existing IBM Storwize V3700 and click **Next**, as shown in Figure 2-35.

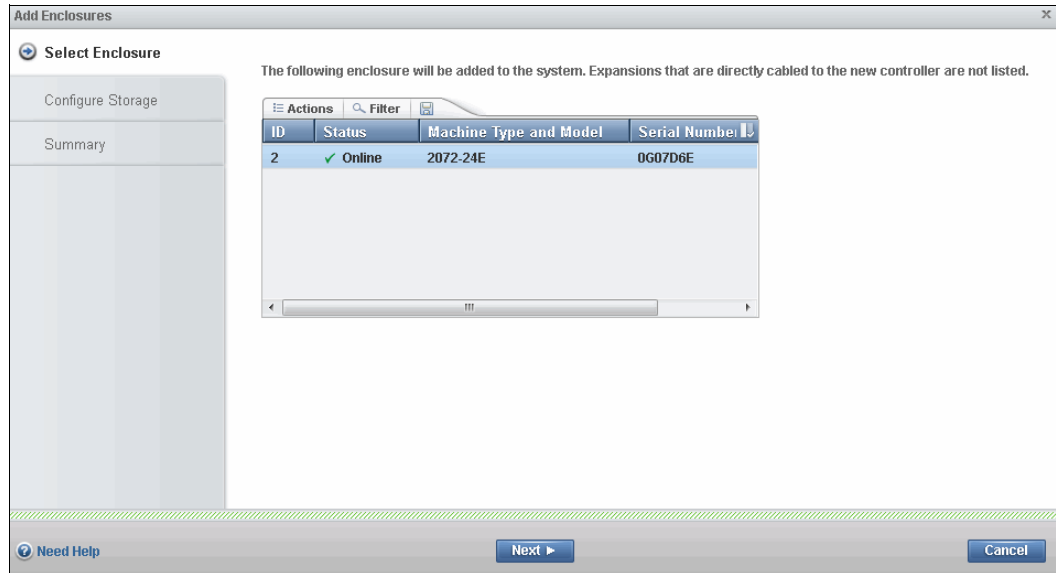


Figure 2-35 Adding a new expansion enclosure

5. As described in step 3 on page 60, you are prompted to configure the new storage as shown in Figure 2-36.

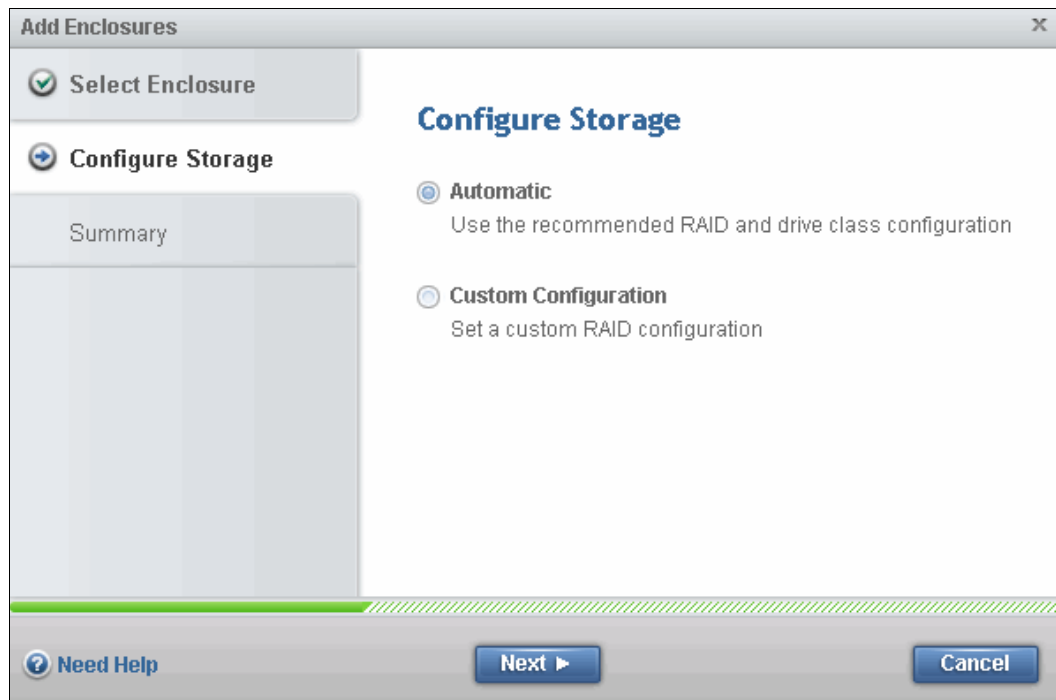


Figure 2-36 The new enclosure as part of the cluster

If you choose **Automatic** configuration, the system automatically configures the storage into optimal arrays for use as MDisks. If you choose **Custom Configuration**, the wizard offers a more flexible way to configure the new storage, as shown in Figure 2-37 and Figure 2-38.

Custom configuration offers a more flexible way to configure the internal storage as compared to the **Use the recommended configuration** preset in terms of drive selection, RAID level, and Storage Pool to be used.

6. As shown in Figure 2-37, select the class of the drivers that you want to include in the RAID arrays and click **Next**.

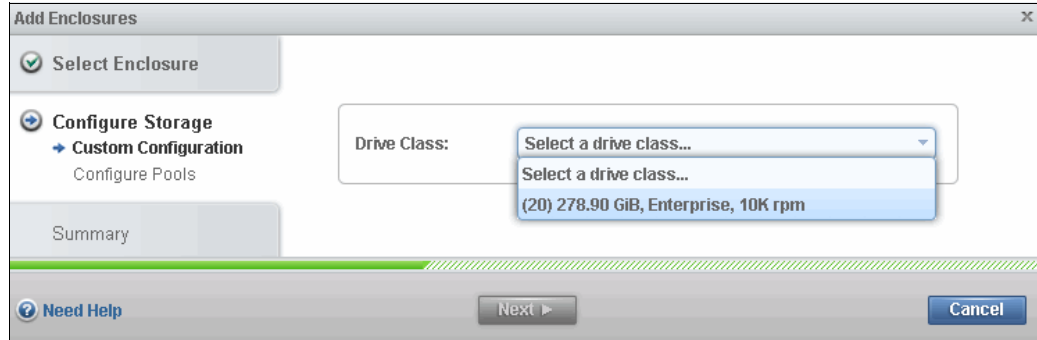


Figure 2-37 Select a drive class

7. Next, select the level of the disk RAID and click **Next**. As shown in Figure 2-38.

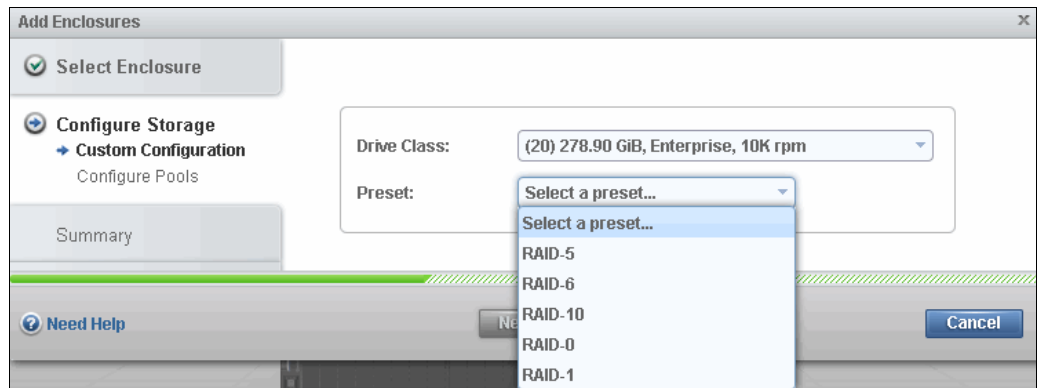


Figure 2-38 Internal storage preset selection

8. As shown in Figure 2-39, the IBM Storwize V3700 offers the suggested internal RAID arrays based on the class of drive and the level of RAID previously chosen. The configuration is shown at the bottom of the window in the summary section. Accept the defaults and click **Next**.

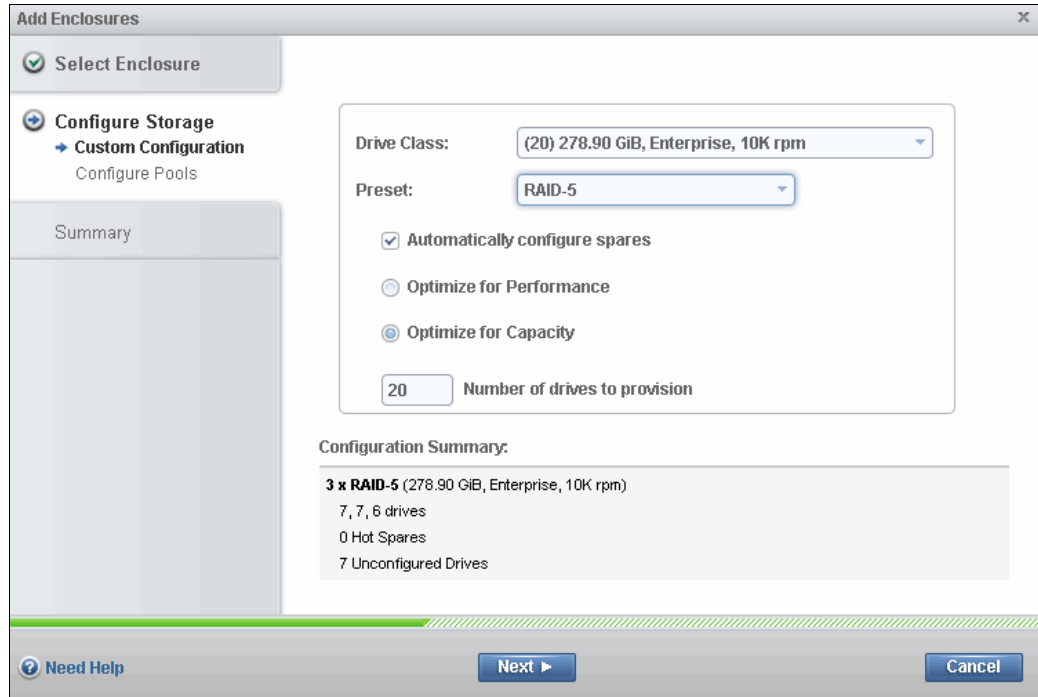


Figure 2-39 The suggested configuration

9. Choose whether an existing pool is expanded or whether a new pool is created for the configuration, as shown in Figure 2-40. In our example, we create a new Storage Pool named V3700_Pool2.

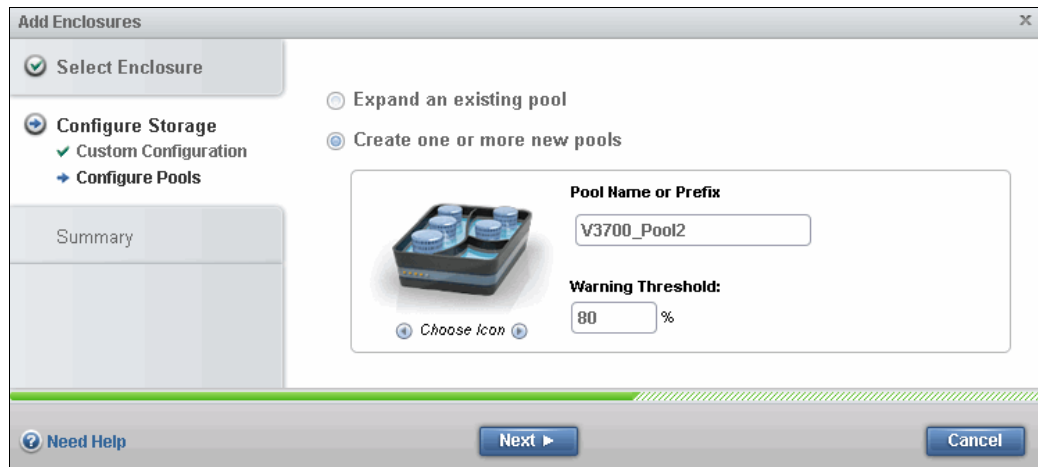


Figure 2-40 Storage Pool selection

10. As shown in Figure 2-41, the wizard shows the internal storage presets. Click **Cancel** if you do not want to continue with the configuration shown, or click **Finish** and the IBM Storwize V3700 GUI creates the arrays shown in the Summary.

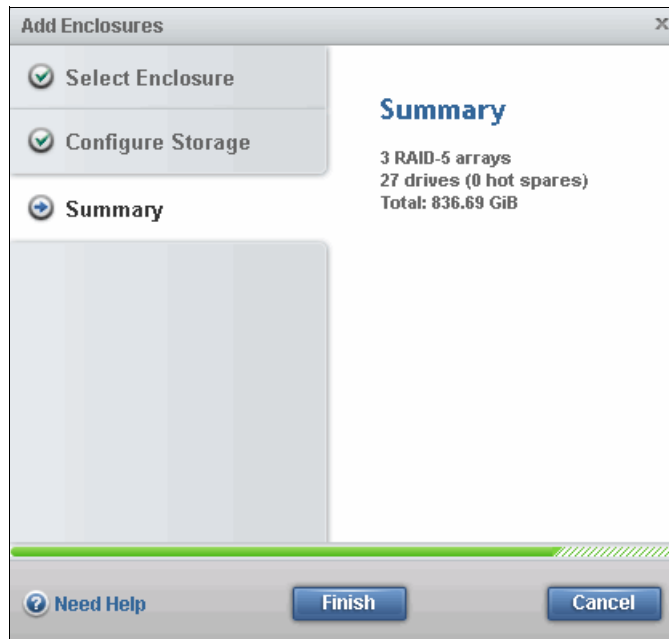


Figure 2-41 Add Enclosures Summary

11. Click **Close** when the task completes, as shown in Figure 2-42.




Figure 2-42 Add enclosure Task completed

The new expansion enclosure is configured within the IBM Storwize V3700, as shown in Figure 2-43.



Figure 2-43 Expansion Enclosure included

For more information about how to provision the new storage in the expansion enclosures, see Chapter 7, “Storage Pools” on page 307.

2.9.2 Configure Call Home, email alert, and inventory

Email and alert notification enables the administrators to receive the events from IBM Storwize V3700. If your system is under warranty or you have a hardware maintenance agreement, it is advised to configure your system to send email reports to IBM if an issue that requires hardware investigation is detected. This feature is known as Call Home, and is typically configured during the Initial Configuration of the system as described in step 10 on page 54.

To configure the Call Home and email alert event notification in IBM Storwize V3700 after the Initial Configuration, complete the following steps:

1. Go to **Settings** → **Event Notifications**, as shown in Figure 2-44.




Figure 2-44 Enabling Call Home

2. If your system has the Email Notification disabled, you can re-enable it by clicking **Email** → **Enable Notifications**, as shown in Figure 2-45.

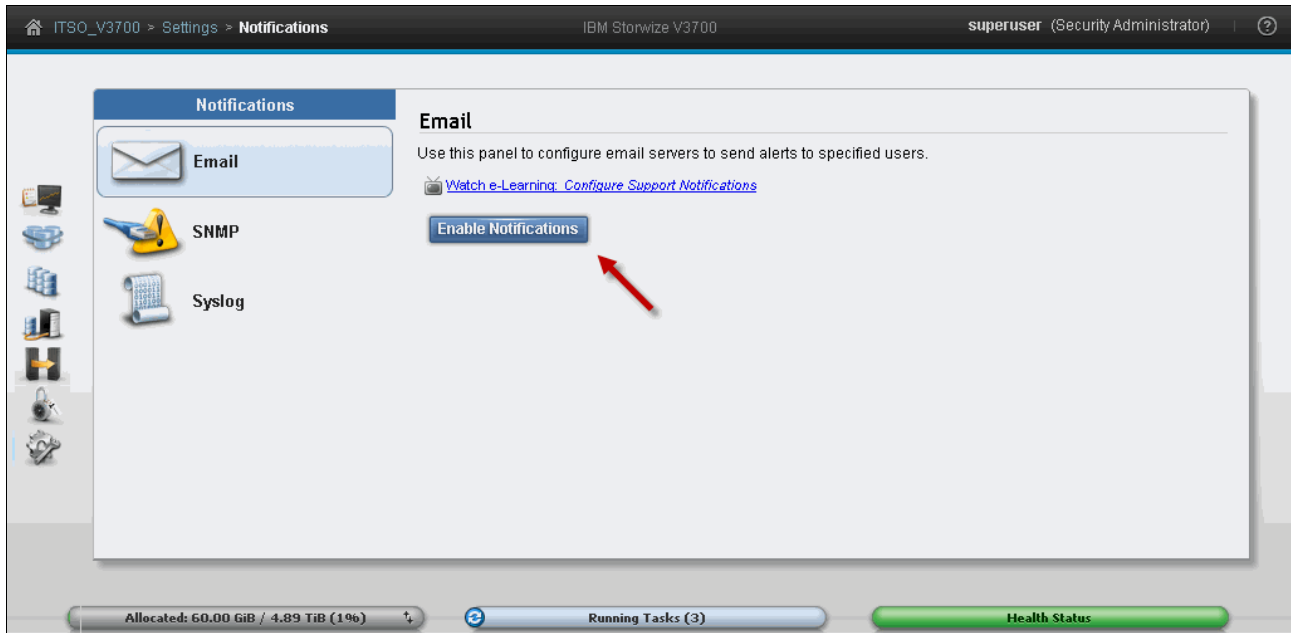


Figure 2-45 Selecting Enable Notifications

3. If the Email and Call Home Notifications were never configured in the IBM Storwize V3700, a window opens and the wizard guides you through the procedure described in step 10 on page 54. However, if the Email and Call Home Notifications were previously configured and you need to edit any of the fields, click **Edit** and the fields to configure Call Home and Email Notification become available, as shown in Figure 2-46.

The screenshot displays the 'Email' configuration page in the IBM Storwize V3700 web interface. The breadcrumb navigation shows 'ITSO_V3700 > Settings > Notifications'. The user is logged in as 'superuser (Security Administrator)'. The left sidebar contains navigation icons for 'Email', 'SNMP', and 'Syslog'. The main content area is titled 'Email' and includes the following sections:

- Email Servers:** A table with columns 'IP Address' (9.9.9.1) and 'Server Port' (25). Buttons for '+', '-', 'Save', and 'Cancel' are present.
- Call Home:** Fields for 'Email Address' (callhome1@de.ibm.com), 'Event Type' (Alerts), and 'Inventory' (checked). A 'Test' button is also visible.
- Email Notifications:** A table with columns 'Email Address' (empty), 'Event Type' (Alerts), and 'Inventory' (unchecked). Buttons for '+', '-', 'Save', and 'Cancel' are present.
- Email Contact:** Fields for '* Contact Name' (Oscar Wellington), '* Email Reply Address' (owell@ibm.com), 'Comment', '* Telephone (Primary)' (7786100900), and 'Telephone (Alternate)' (empty). A '* Required' note is at the bottom.

Figure 2-46 Email notification

4. You must enter accurate and complete information about both the company and the contact. The Email Contact is the person who will be contacted by the support center about any issues, and for further information. You might want to enter a network operations center Email Address here for 24x7 cover. Enter the IP address and the server port for one or more of the email servers that send an email to IBM.
If you would like email alerts to be sent to other recipients in addition to IBM, add their addresses into the Email Notification section. Use the plus (+) symbol to add further recipients and select one or more event types for each user. A maximum of 12 users and a maximum of six email servers are supported.
5. Click **Save** when you finish.

2.9.3 Service Assistant tool

The IBM Storwize V3700 is initially configured with three IP addresses:

- ▶ One service IP address for each of two node canisters
- ▶ One management IP address, which is set when the cluster is started

The following methods are available to configure the IBM Storwize V3700 system:

- ▶ The `InitTool` program, as described in 2.8, “First-time setup” on page 40
- ▶ The Service Assistant (SA) tool, which is described next

Additionally, the Management IP and Service IP addresses can be changed within the GUI, as described in Chapter 3, “Graphical user interface overview” on page 71.

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

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

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

The SA GUI is available using a service assistant IP address on each node. The SA GUI is accessed through the cluster IP addresses by appending `service` to the cluster management Uniform Resource Locator (URL). If the system is down, the only other method of communicating with the node canisters is through the SA IP address directly. Each node can have a single SA IP address on Ethernet port 1. It is advised that these IP addresses are configured on all IBM Storwize V3700 node canisters.

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

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

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

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

For example, you might use the following addresses:

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

Complete the following steps to access the SA:

1. When you are accessing SA by using `<cluster IP address>/service`, the configuration node canister SA GUI login window opens. Enter the Superuser Password, as shown in Figure 2-47.

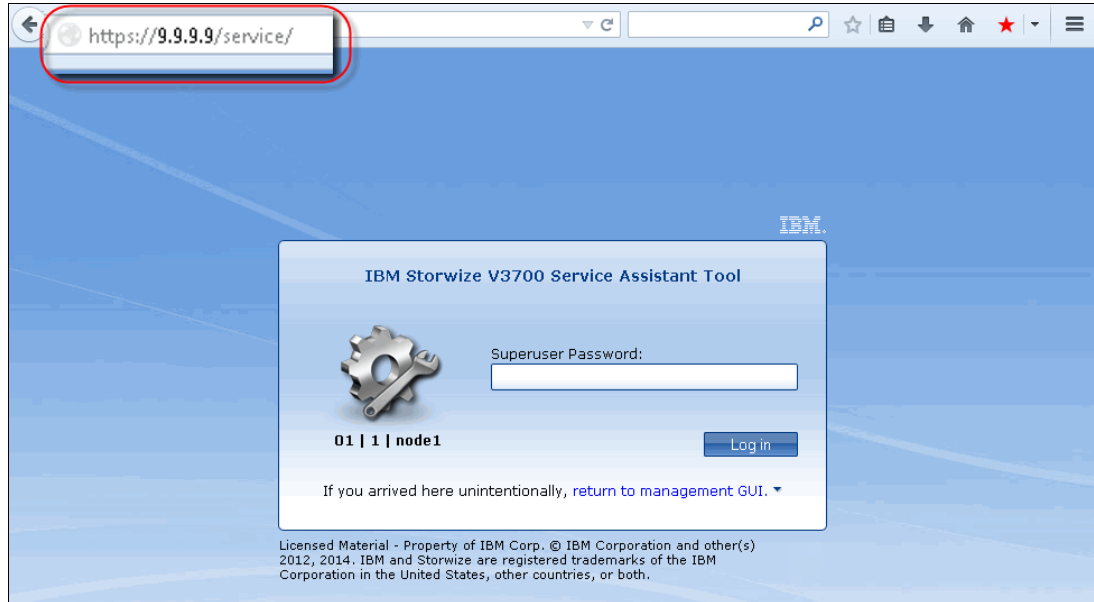


Figure 2-47 Service Assistant Tool login

- After you are logged in, you see the Service Assistant Home window, as shown in Figure 2-48. The SA interfaces can view the status of, and run service actions on, other nodes, in addition to the node where the user is connected.

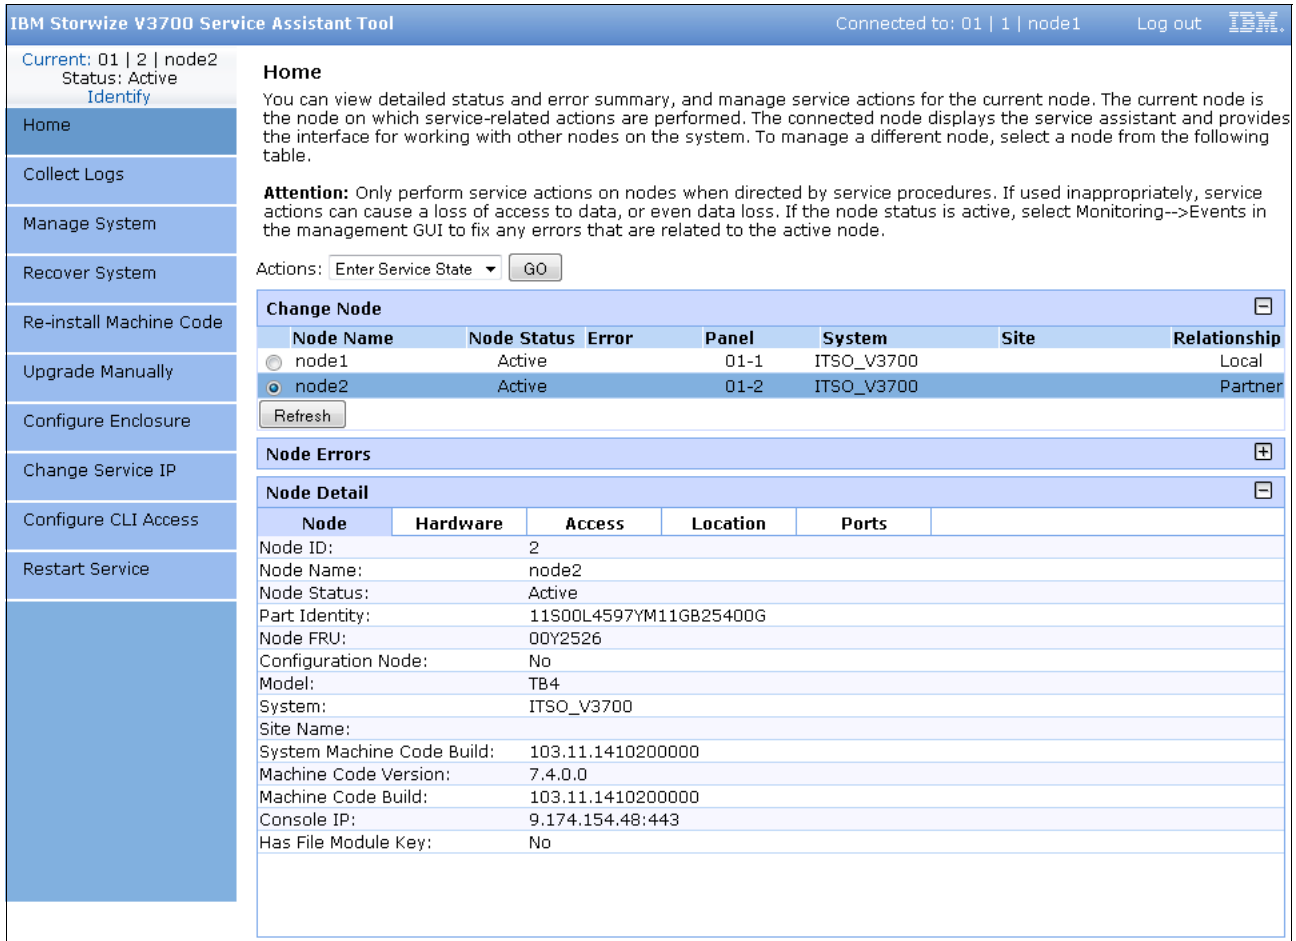


Figure 2-48 Service Assistant Home window

- The Current canister node is displayed in the upper left corner of the GUI. As shown in Figure 2-48, this is node 2. To change the canister, select the relevant node in the Change Node section of the window. You see the details in the upper left change to reflect the new canister.

The SA GUI provides access to service procedures and shows the status of the node canisters. It is advised that these procedures should only be carried out if directed to do so by IBM Support.

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

<https://ibm.biz/BdEz3N>



Graphical user interface overview

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

This chapter includes the following topics:

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

3.1 Getting started

This section provides information about accessing the IBM Storwize V3700 management GUI. It covers topics, such as supported browsers, log in modes, and the layout of the System pane.

3.1.1 Supported browsers

The IBM 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.

For information about supported browsers, see the IBM Storwize V3700 Knowledge Center on the following website:

http://www.ibm.com/support/knowledgecenter/STLM5A_7.4.0/com.ibm.storwize.v3700.740.doc/v3700_ichome_740.html

3.1.2 Access the management GUI

To access the management GUI, complete the following steps:

1. Open a supported web browser and enter the management Internet Protocol (IP) address or host name of the IBM Storwize V3700. The login pane is displayed, as shown in Figure 3-1.



Figure 3-1 IBM Storwize V3700 login pane

Default user name and password:

Change the default password immediately. Use the following information to log in to the IBM Storwize V3700 storage management:

- ▶ User Name: superuser
- ▶ Password: passw0rd (a zero replaces the letter “o”)

Note: Older versions of the software allow low-resolution graphics to be used by selecting a check box on the login pane. Low-resolution graphic mode is now selected using the GUI preferences pane. See “GUI Preferences pane” on page 143.

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

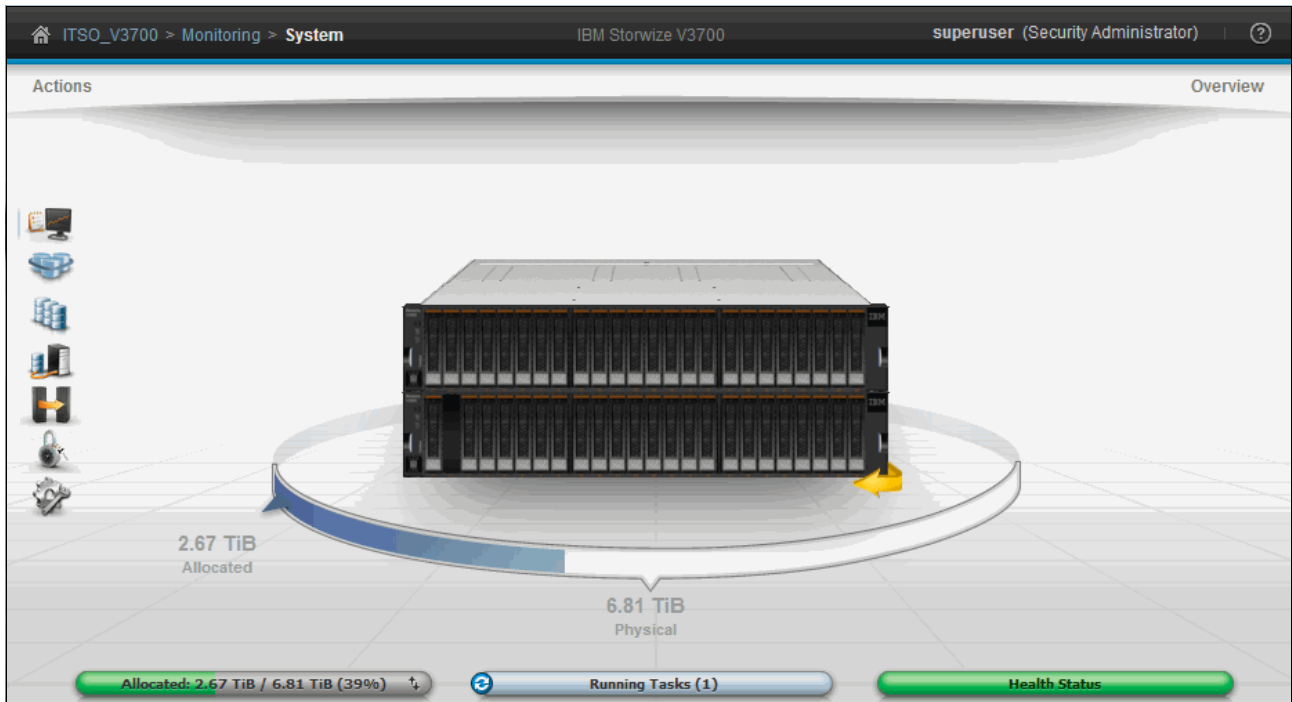


Figure 3-2 IBM Storwize V3700 System pane

3. Click the logged in user name in the top banner, as shown in Figure 3-3. This enables you to Log Out, Modify Password, and Manage SSH Public Key.



Figure 3-3 User menu

3.1.3 System pane layout

As shown in Figure 3-4, the System pane has three main sections:

- ▶ Function Icons
- ▶ System view
- ▶ Status Indicators

The System pane is new in this software release, and replaces the Overview and System Details panes from older software releases.

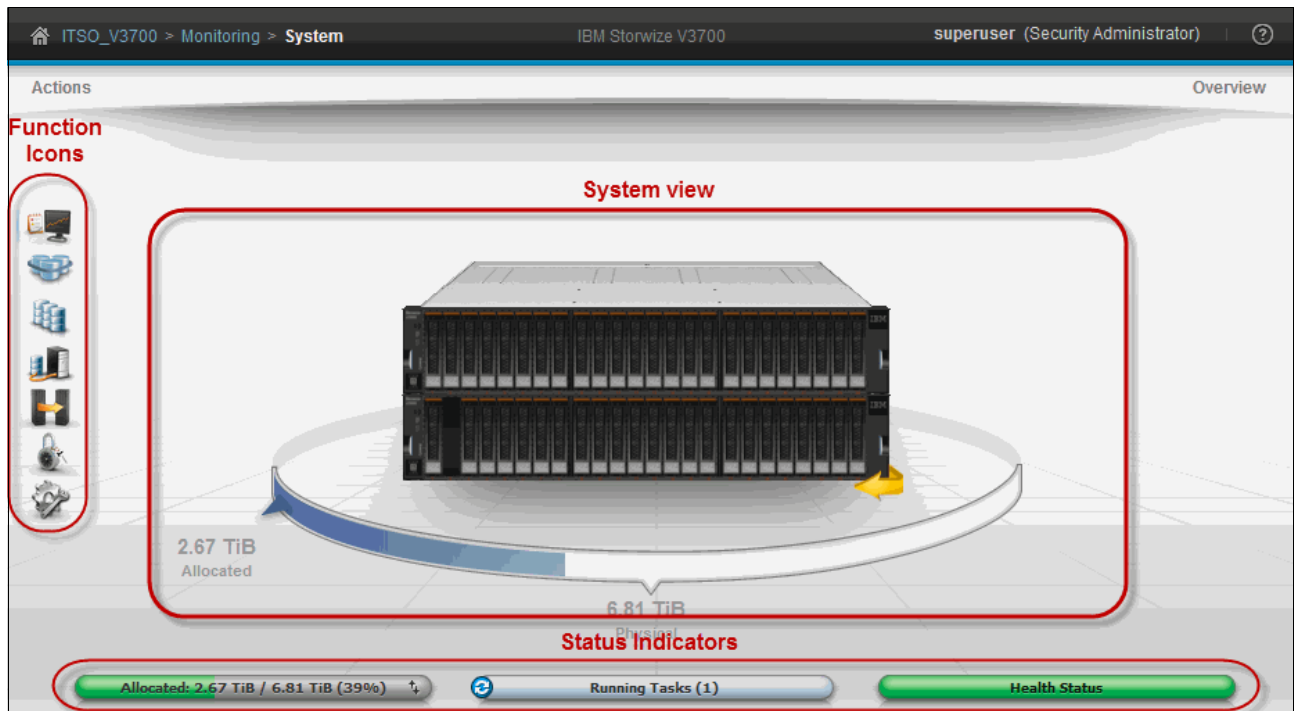


Figure 3-4 Three main sections of the IBM Storwize V3700 System pane

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

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

The System view section shows a 3D representation of the IBM Storwize V3700. Hovering over any of the components shows a component overview, where right-clicking a component shows an Action menu appropriate to that component. Clicking the arrow at the lower right of the graphic rotates it to show the rear of the system. This arrow can be blue, yellow, or red. Blue indicates no issues, yellow indicates a warning, and red indicates an issue of some kind with a component.

The Status Indicators section in Figure 3-4 on page 74 shows the following horizontal status bars:

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

To access more information and other menus, complete the following steps:

1. Hover over or click the horizontal bars to access the menus, which are described in 3.3, “Status Indicators menus” on page 89.
2. There are also two links at the top of the System pane. The Actions link at the upper left opens an Action menu, as shown in Figure 3-5.

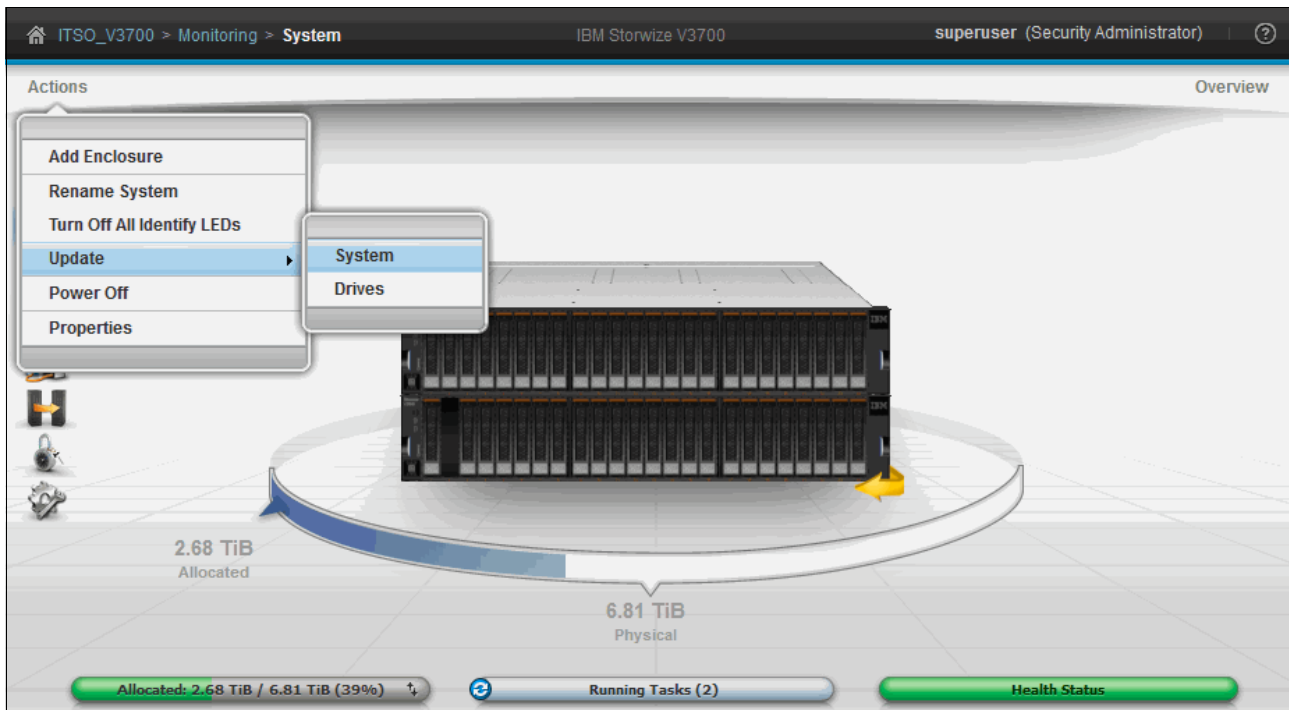


Figure 3-5 Systems pane: Actions menu

3. The Overview link at the upper right toggles an Overview pane, as shown in Figure 3-6.

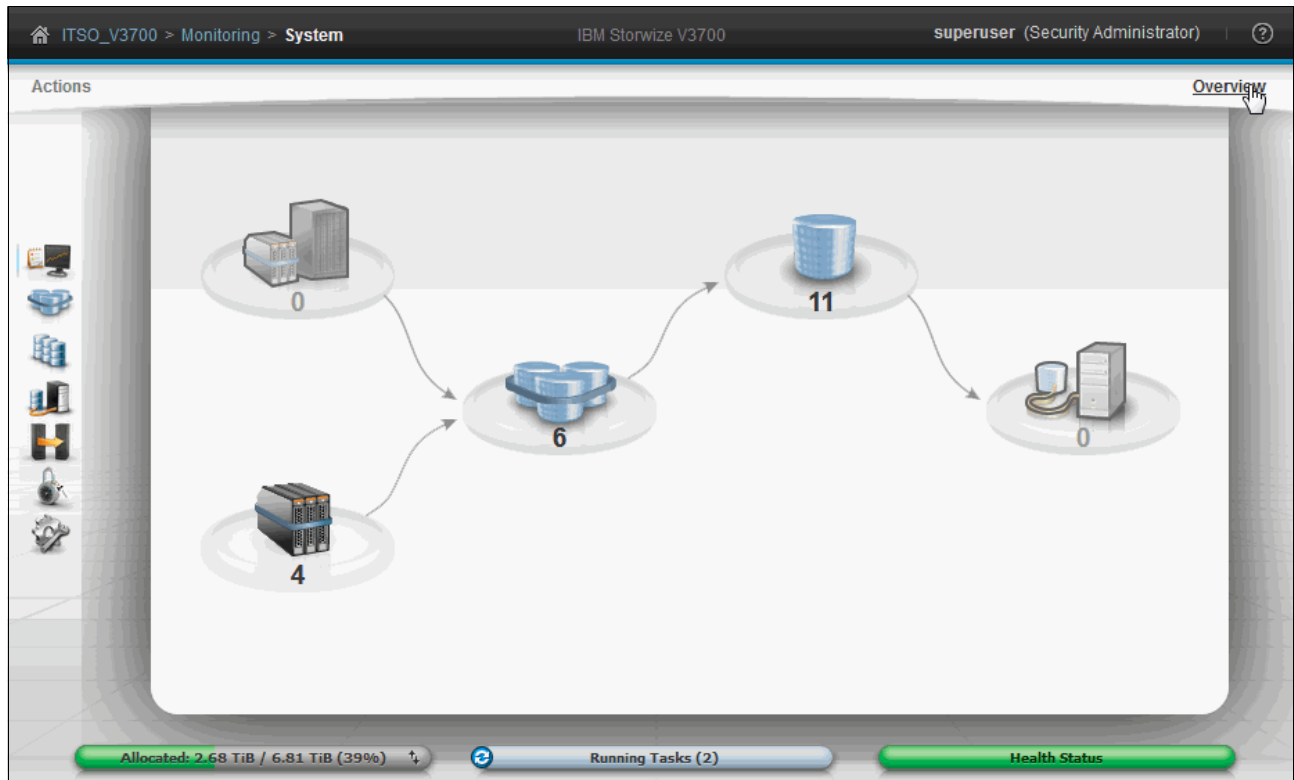


Figure 3-6 System pane: overview

For full details about the System pane, see “System pane” on page 94.

3.2 Navigation

Navigating the management tool is simple and, like most systems, there are many ways to navigate. The two main methods are to use the Function Icons section or the Overview menu.

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

3.2.1 Function icons navigation

To navigate the function icons, complete the following steps:

1. Hover the mouse pointer over one of the seven function icons on the left side of the pane to enlarge the icon and provide a menu with which to access various functions.
2. Move the pointer to the required function and click the function. Figure 3-7 on page 77 shows the results of hovering the mouse pointer over the Pools function icon.




Figure 3-7 Hovering over the Pools function icon

3. Figure 3-8 shows all of the menus with options under the list of function icons.

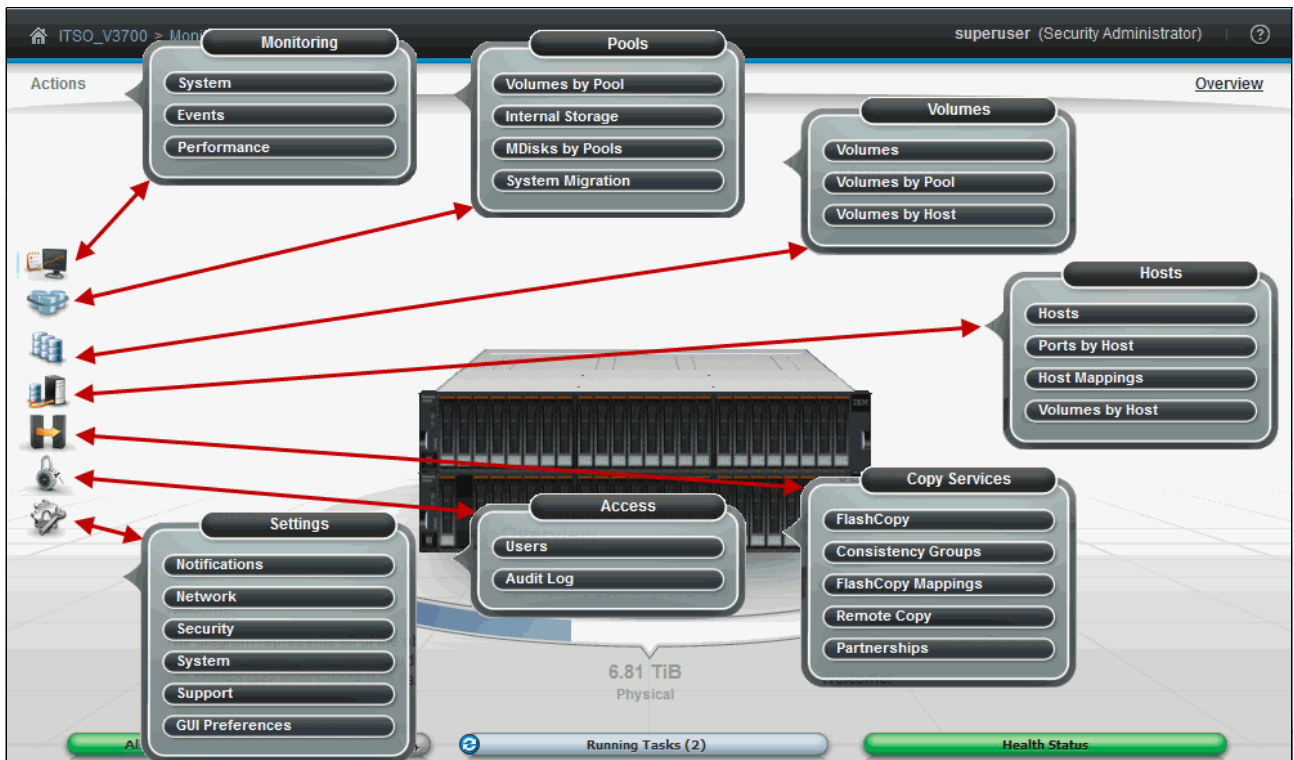


Figure 3-8 Options that are listed under Function Icons section

3.2.2 Breadcrumb navigation aid

The IBM Storwize V3700 panes use the breadcrumb navigation aid to show the path that was browsed. This breadcrumb navigation aid is in the top area of the pane, and hovering over a breadcrumb in the trail shows a menu.

Figure 3-9 shows the breadcrumb navigation aid for the System pane.

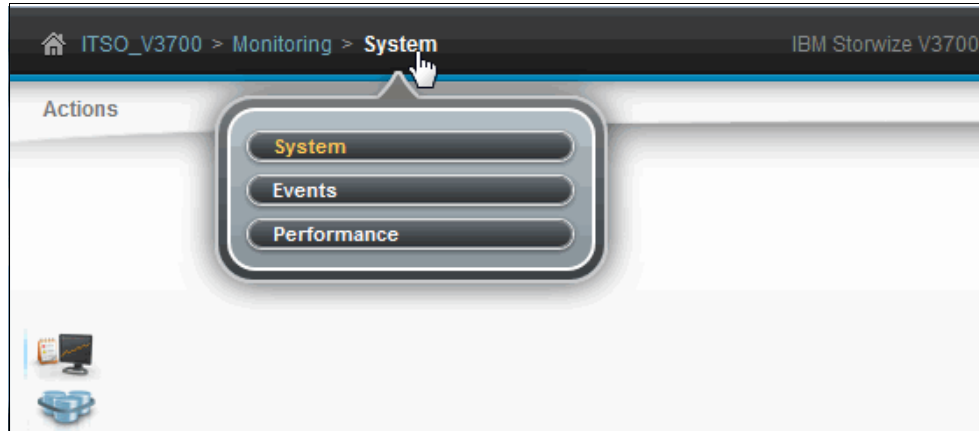


Figure 3-9 Breadcrumb navigation aid

3.2.3 Suggested Tasks feature

The Suggested Tasks feature is a pane that opens at login and displays any outstanding tasks. The number of suggested tasks changes depending on the configuration of the system. Clicking the suggested task opens the corresponding pane to complete the task.

Figure 3-10 shows the Suggested Task pane.

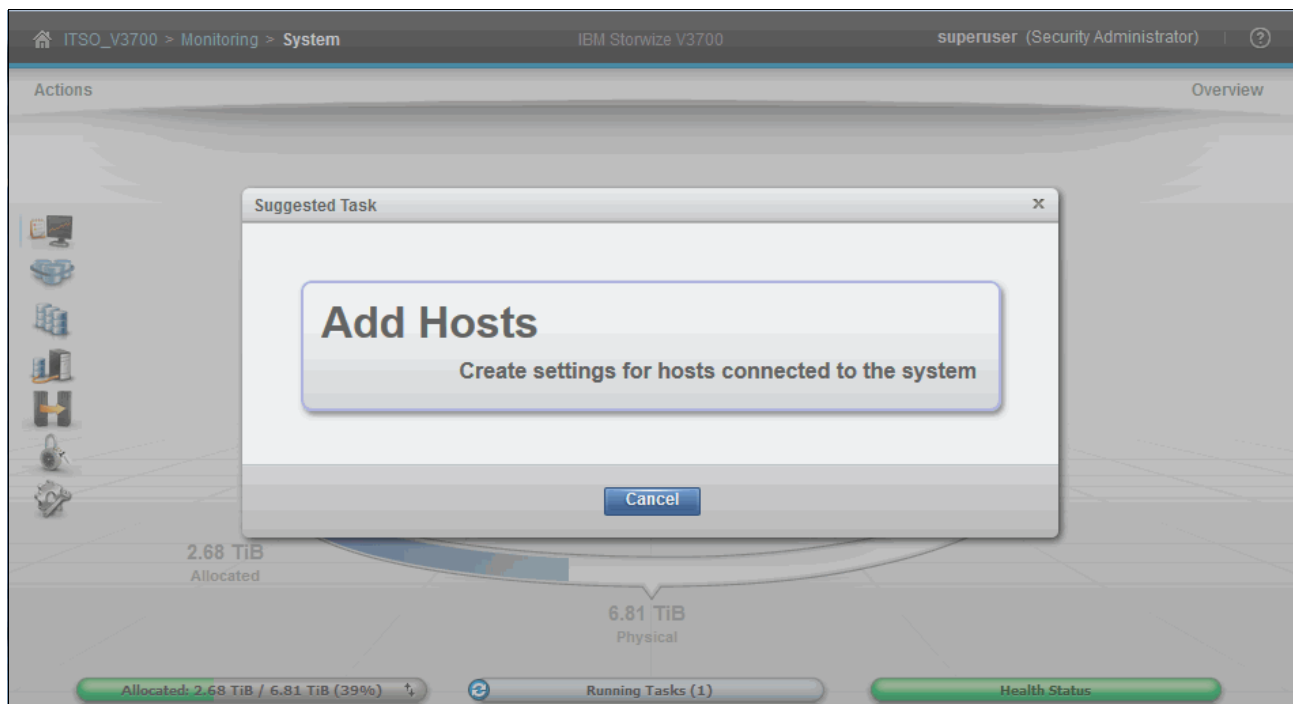


Figure 3-10 Suggested Task pane

3.2.4 Presets

The management GUI contains a series of pre-established configuration options, called *presets*, which are commonly used settings to quickly configure objects on the system. Presets are available for creating volumes, for IBM FlashCopy mappings, and for setting up a Redundant Array of Independent Disks (RAID) configuration.

Figure 3-11 shows the available internal storage presets.

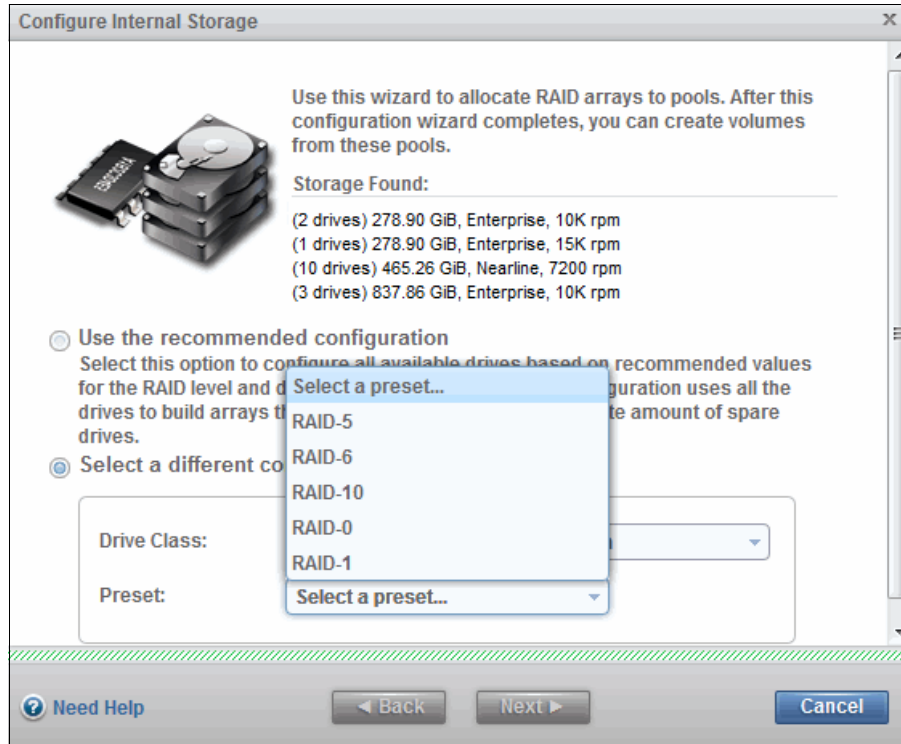


Figure 3-11 Internal storage preset selection

3.2.5 Actions

The IBM Storwize V3700 functional panes provide access to various actions that can be performed, such as modify attributes and rename, add, or delete objects. The Action menu options change depending on the pane accessed.

The available Actions menus can be accessed by using one of two main methods:

- ▶ Highlight the resource and use the Actions drop-down menu, as shown in Figure 3-12.
- ▶ Right-click the resource, as shown in Figure 3-13.

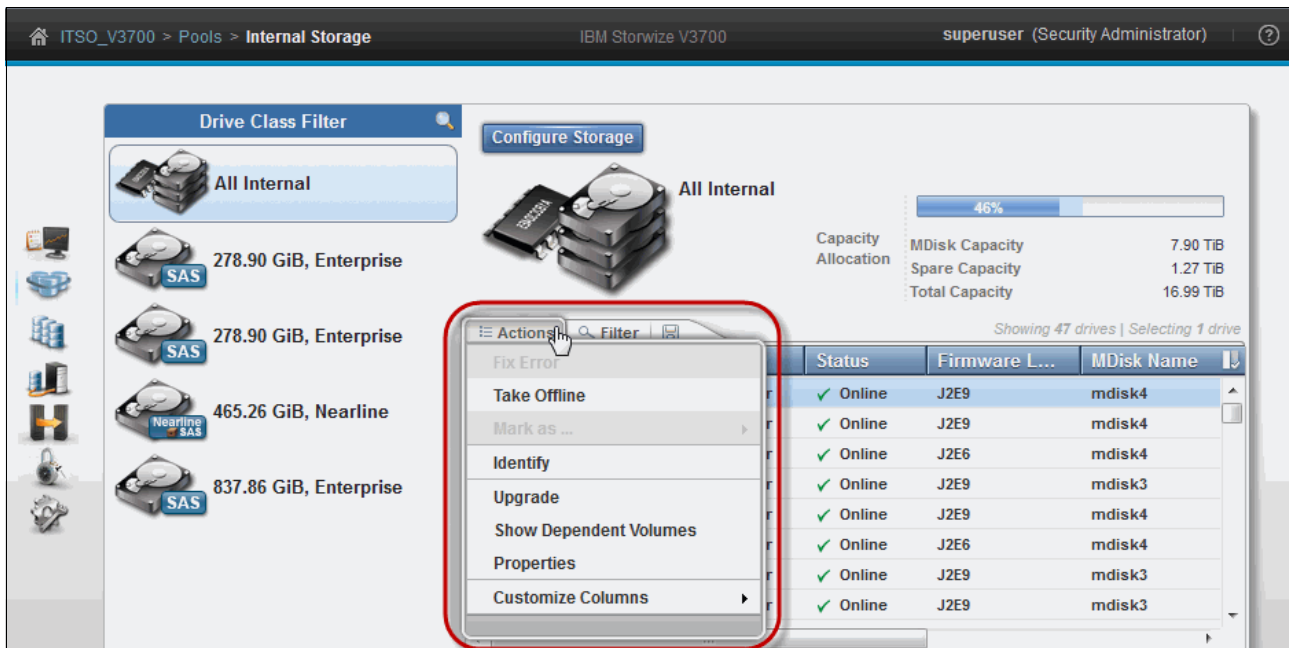


Figure 3-12 Actions menu

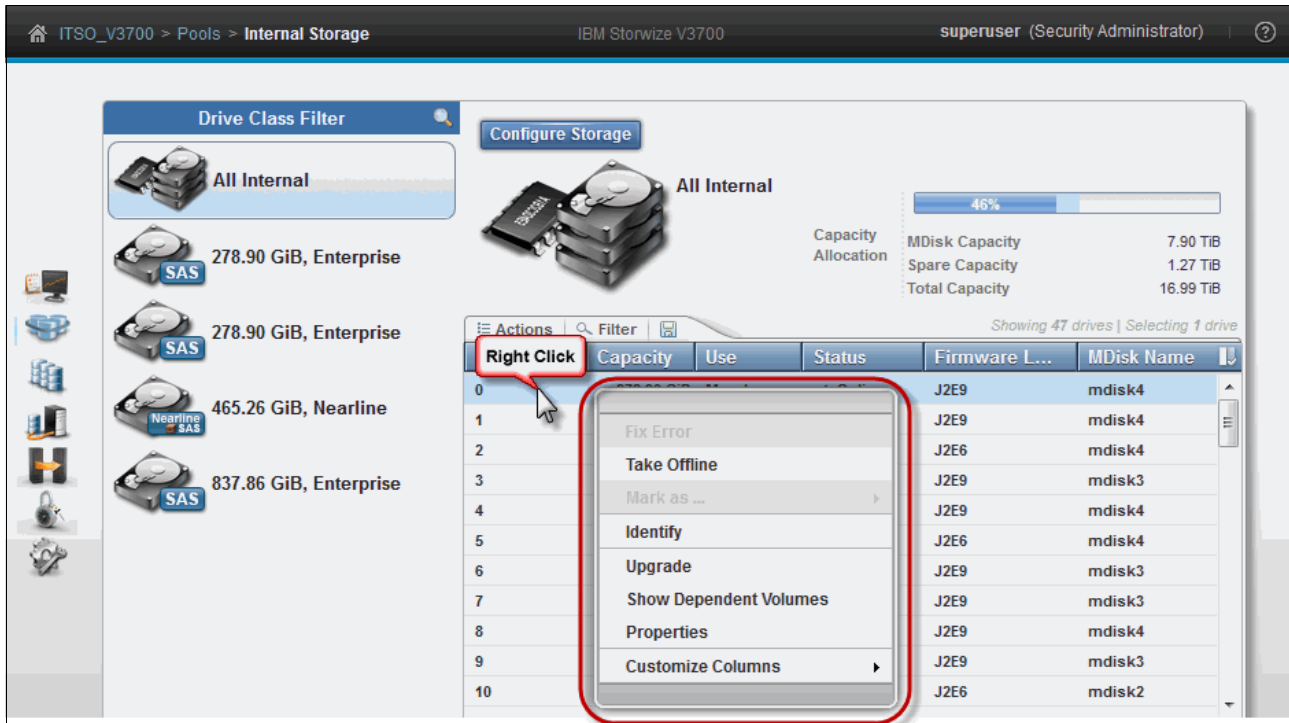


Figure 3-13 Right-clicking an item to access the Actions menu

3.2.6 Task progress

An action starts a running task and shows a task progress pane, as shown in Figure 3-14:

1. Click **View more details** to see the underlying command-line interface (CLI) commands.
2. The commands are highlighted in blue, and can be copied and pasted into a configured IBM Storwize V3700 SSH terminal session, if required. This is useful when developing CLI scripts.

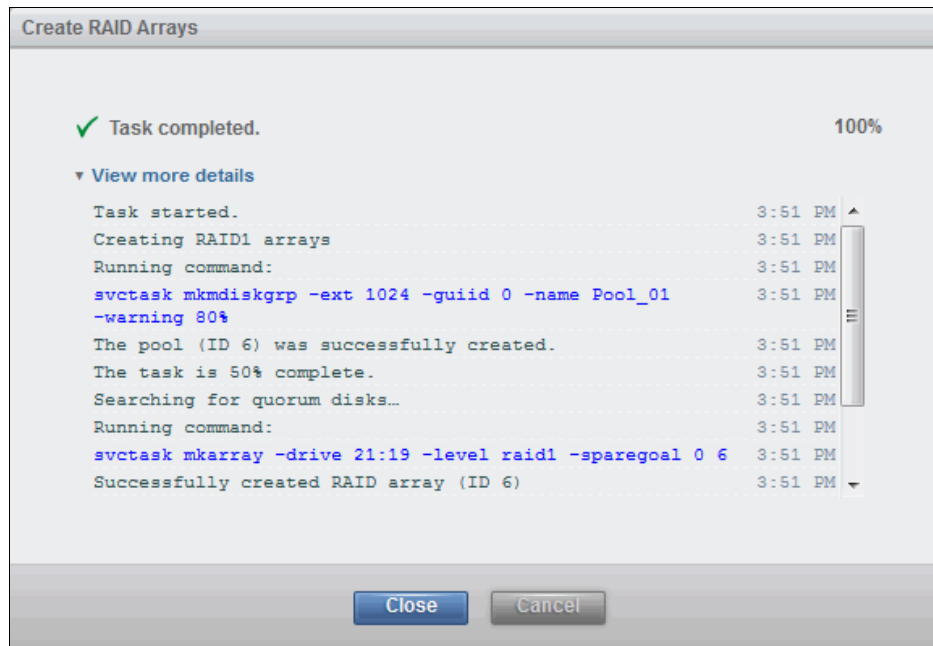


Figure 3-14 Task progress pane

3.2.7 Navigating panes with tables

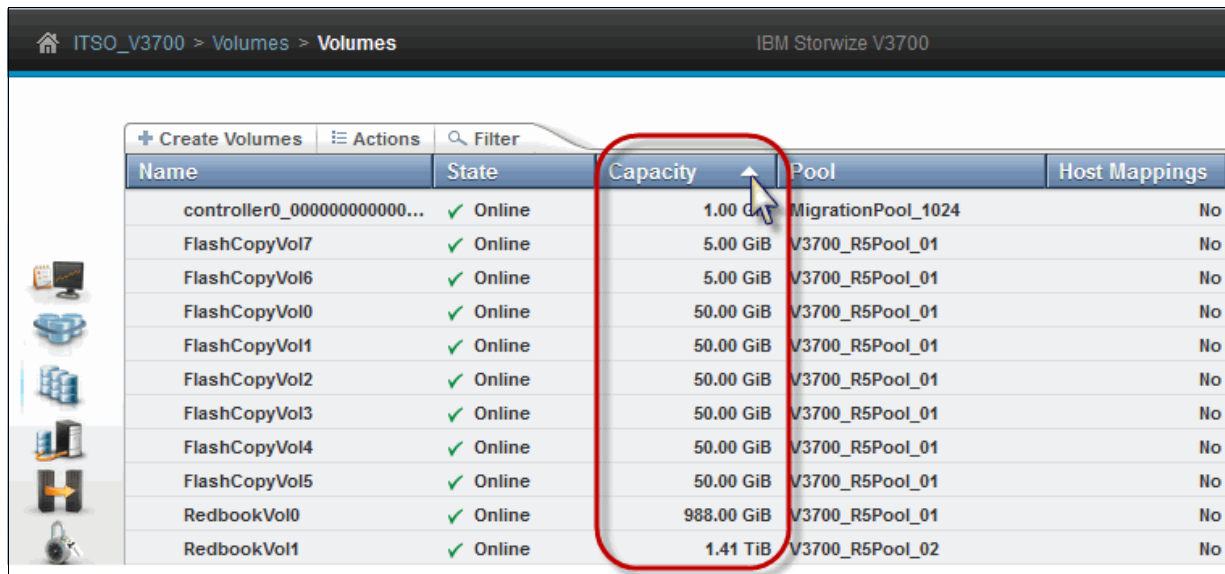
Many of the configuration and status panes show information in a table format. This section describes the following methods to navigate panes with rows and columns:

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

Sorting columns

To sort columns, complete the following steps:

1. Click the column heading. Figure 3-15 shows the result of clicking the heading of the Capacity column. The table is now sorted, and lists volumes with the least amount of capacity at the top of the table. To reverse the sort order, click the heading again.



ITSO_V3700 > Volumes > Volumes IBM Storwize V3700

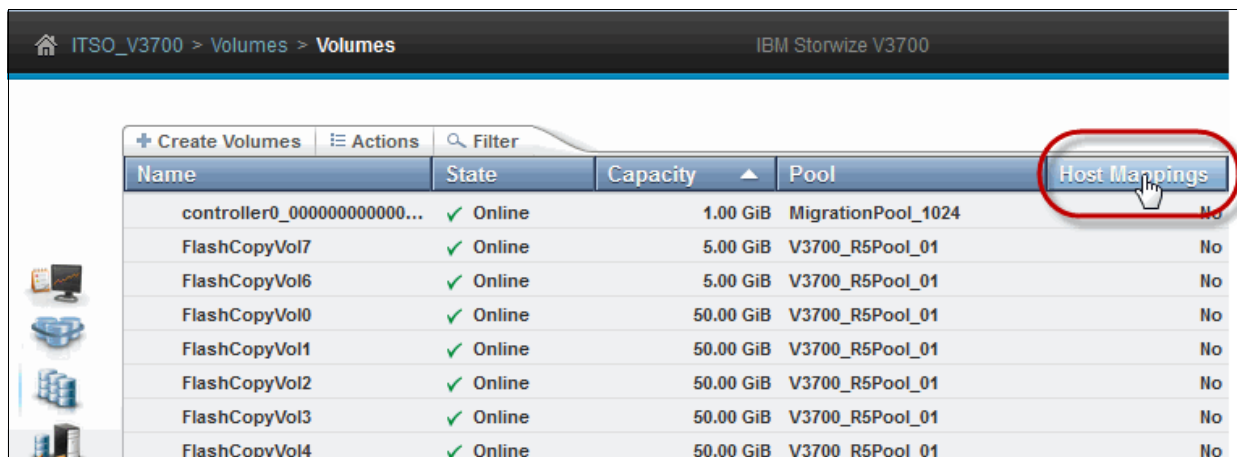
Name	State	Capacity	Pool	Host Mappings
controller0_000000000000...	✓ Online	1.00 GiB	MigrationPool_1024	No
FlashCopyVol7	✓ Online	5.00 GiB	V3700_R5Pool_01	No
FlashCopyVol6	✓ Online	5.00 GiB	V3700_R5Pool_01	No
FlashCopyVol0	✓ Online	50.00 GiB	V3700_R5Pool_01	No
FlashCopyVol1	✓ Online	50.00 GiB	V3700_R5Pool_01	No
FlashCopyVol2	✓ Online	50.00 GiB	V3700_R5Pool_01	No
FlashCopyVol3	✓ Online	50.00 GiB	V3700_R5Pool_01	No
FlashCopyVol4	✓ Online	50.00 GiB	V3700_R5Pool_01	No
FlashCopyVol5	✓ Online	50.00 GiB	V3700_R5Pool_01	No
RedbookVol0	✓ Online	988.00 GiB	V3700_R5Pool_01	No
RedbookVol1	✓ Online	1.41 TiB	V3700_R5Pool_02	No

Figure 3-15 Sorting columns by clicking the column heading

Reordering columns

To reorder columns, complete the following steps:

1. Figure 3-16 shows the location of the column with the Host Mappings heading positioned in the last column. Drag a heading left or right to reorder the columns in the table.



ITSO_V3700 > Volumes > Volumes IBM Storwize V3700

Name	State	Capacity	Pool	Host Mappings
controller0_000000000000...	✓ Online	1.00 GiB	MigrationPool_1024	No
FlashCopyVol7	✓ Online	5.00 GiB	V3700_R5Pool_01	No
FlashCopyVol6	✓ Online	5.00 GiB	V3700_R5Pool_01	No
FlashCopyVol0	✓ Online	50.00 GiB	V3700_R5Pool_01	No
FlashCopyVol1	✓ Online	50.00 GiB	V3700_R5Pool_01	No
FlashCopyVol2	✓ Online	50.00 GiB	V3700_R5Pool_01	No
FlashCopyVol3	✓ Online	50.00 GiB	V3700_R5Pool_01	No
FlashCopyVol4	✓ Online	50.00 GiB	V3700_R5Pool_01	No

Figure 3-16 Reordering columns

- Figure 3-17 shows the column heading Host Mappings as it is dragged to the required location.

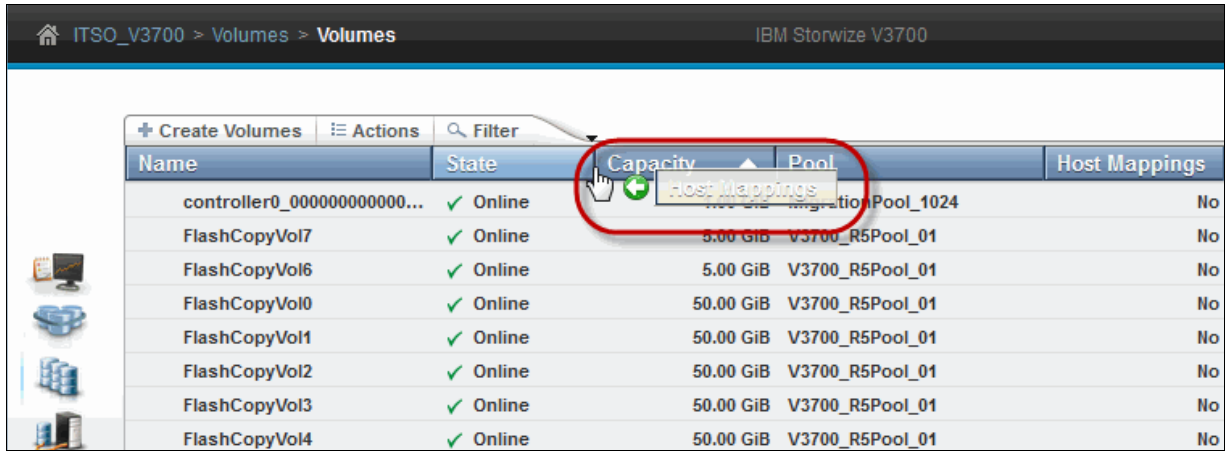


Figure 3-17 Dragging a column heading to the required location

- Figure 3-18 shows the result of dragging the column heading Host Mappings to the new location.

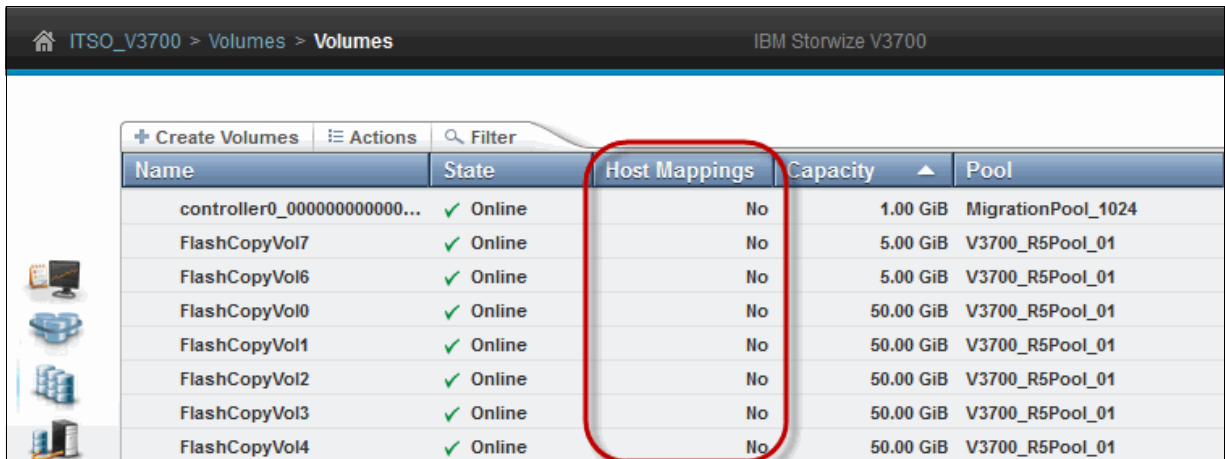


Figure 3-18 Reordering column headings

Adding or removing columns

To add or remove a column, complete the following steps:

1. Right-click the heading bar and select or clear the check box for the required column headings. Figure 3-19 shows the addition of the Real Capacity column.

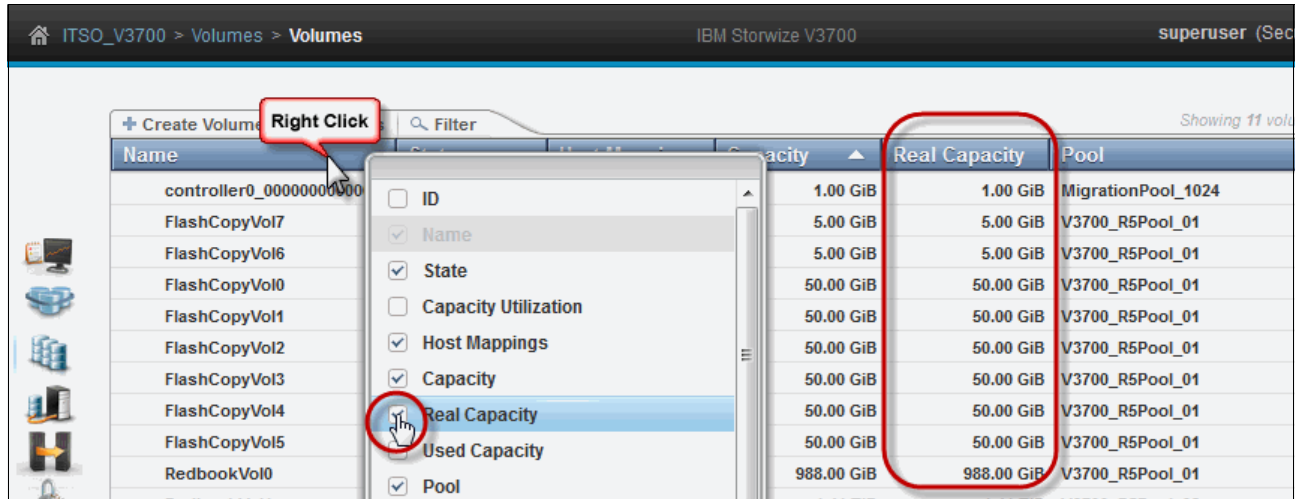


Figure 3-19 Adding column heading Real Capacity

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

In the newest versions of Firefox, this option is not available. If this is the case for you, complete the following steps:

- a. In the address bar, type `about:config` and press Enter.
- b. Click **I'll be careful, I promise**.
- c. In the search bar, search for `"dom.event.contextmenu.enabled"` (with no quotes).
- d. Right-click the result (**dom.event.contextmenu.enabled**) and click **Toggle**.
- e. To revert to the previous setting, repeat these steps.

The web browser requirements and recommended configuration settings to access the IBM Storwize V3700 management GUI can be found in the IBM Storwize V3700 Knowledge Center on the following website:

<http://www.ibm.com/support/knowledgecenter/STLM5A/welcome>

- It is also possible to add or remove columns by clicking the icon at the far right of the heading bar, as shown in Figure 3-20.

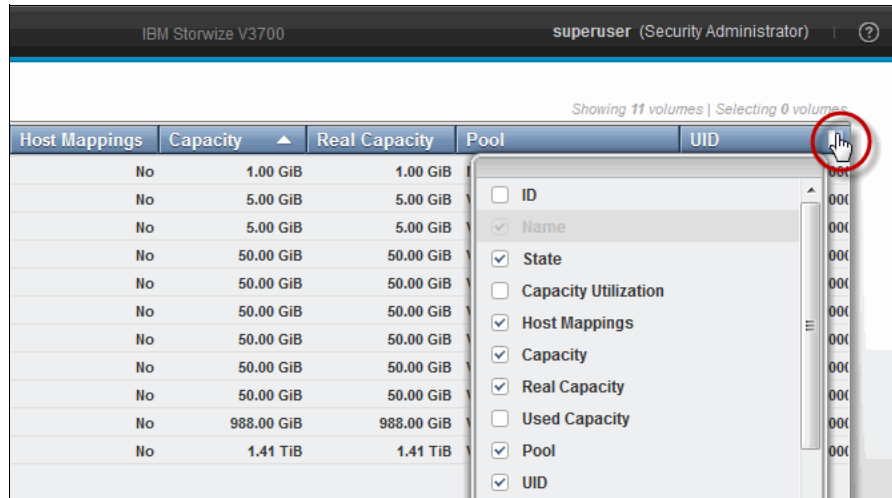


Figure 3-20 Adding columns icon

Multiple selections

It is also possible to select multiple items in a list by using a combination of the Shift or Ctrl keys.

Using the Shift key

To select multiple items in a sequential order, complete the following steps:

- Click the first item that is listed.
- Press and hold the Shift key.
- Click the last item in the list.

All of the items between the first and last items are selected, as shown in Figure 3-21.

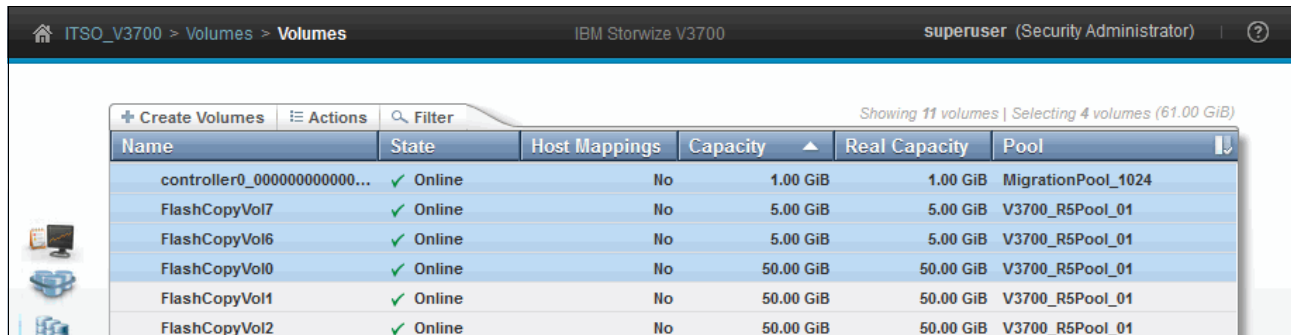


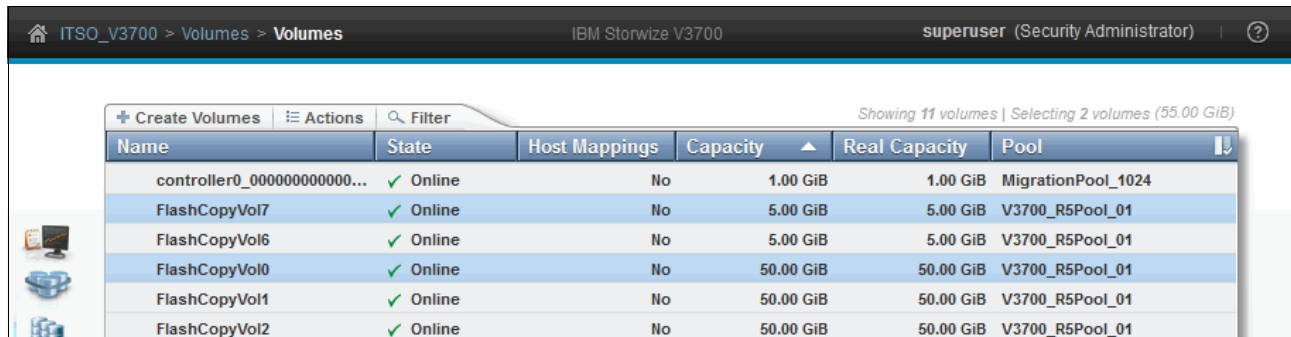
Figure 3-21 Selection of four sequential items

Using the Ctrl key

To select multiple items that are not in sequential order, complete the following steps:

- Click the first item.
- Press and hold the Ctrl key.
- Click the other items that are required.

Figure 3-22 shows the selection of two non-sequential items.



The screenshot shows the IBM Storwize V3700 Volumes page. The breadcrumb navigation is "ITSO_V3700 > Volumes > Volumes". The user is "superuser (Security Administrator)". The page shows a table of volumes with the following data:

Name	State	Host Mappings	Capacity	Real Capacity	Pool
controller0_000000000000...	✓ Online	No	1.00 GiB	1.00 GiB	MigrationPool_1024
FlashCopyVol7	✓ Online	No	5.00 GiB	5.00 GiB	V3700_R5Pool_01
FlashCopyVol6	✓ Online	No	5.00 GiB	5.00 GiB	V3700_R5Pool_01
FlashCopyVol0	✓ Online	No	50.00 GiB	50.00 GiB	V3700_R5Pool_01
FlashCopyVol1	✓ Online	No	50.00 GiB	50.00 GiB	V3700_R5Pool_01
FlashCopyVol2	✓ Online	No	50.00 GiB	50.00 GiB	V3700_R5Pool_01

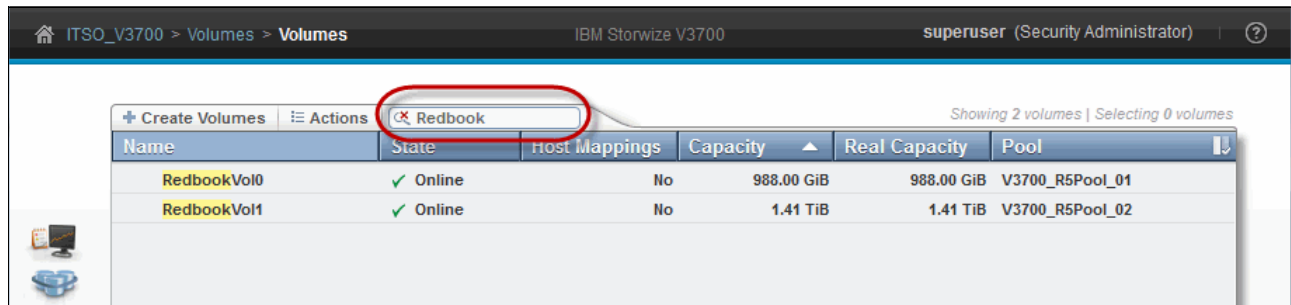
The table indicates "Showing 11 volumes | Selecting 2 volumes (55.00 GiB)".

Figure 3-22 Selecting two non-sequential items

Filtering objects

To focus on a subset of the listed items that are shown in a pane with columns, use the filter field:

1. This tool shows items that match the entered value. Figure 3-23 shows the text Redbook was entered into the filter field.
2. Now, only volumes with the text Redbook in any column are listed, and the filter word is also highlighted.



The screenshot shows the IBM Storwize V3700 Volumes page with a filter applied. The breadcrumb navigation is "ITSO_V3700 > Volumes > Volumes". The user is "superuser (Security Administrator)". The filter field contains "Redbook". The table shows the following data:

Name	State	Host Mappings	Capacity	Real Capacity	Pool
RedbookVol0	✓ Online	No	988.00 GiB	988.00 GiB	V3700_R5Pool_01
RedbookVol1	✓ Online	No	1.41 TiB	1.41 TiB	V3700_R5Pool_02

The table indicates "Showing 2 volumes | Selecting 0 volumes".

Figure 3-23 Filtering objects to display a subset of the volumes

Advanced Filter

To use the Advanced Filter feature, complete the following steps:

1. Click in the Filter field.
2. Click the Advanced Filter icon, as shown in Figure 3-24.



Figure 3-24 Advanced Filter

Figure 3-25 shows the Advanced filter column menu.

3. Select a column name, which enables you to filter by column.
4. Select the appropriate operand, which changes depending on the column being filtered.
5. Finally, enter a filter value. Click **Apply** to display the filtered data.

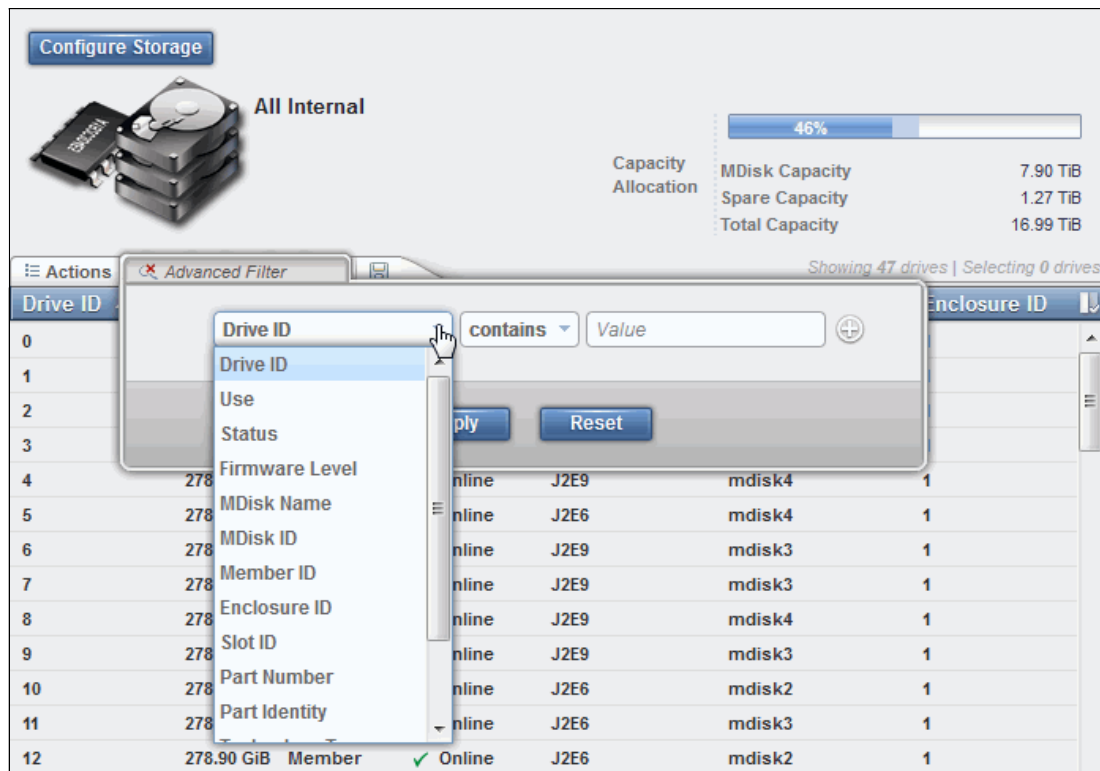


Figure 3-25 Advanced Filter: Choosing values

Figure 3-26 shows the addition of more filters.

6. Click the plus (+) icon to the right of the filter to add further filters.
7. Choose the appropriate logical operand (**AND** or **OR**).
8. Click Apply to display the filtered data.

Filters can also be removed by clicking the minus (-) icon.

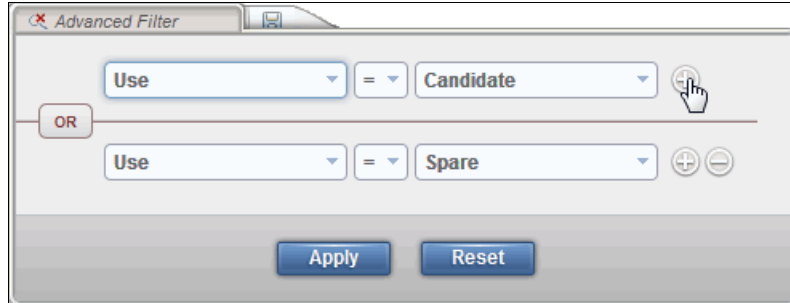


Figure 3-26 Advanced Filter: Additional filters

Figure 3-27 shows the result of the filter used in Figure 3-26. This displays all internal drives that are Candidates or Spares.

9. To reset a filter after the results have been displayed, click the magnifying glass with a red cross through it in the filter field.

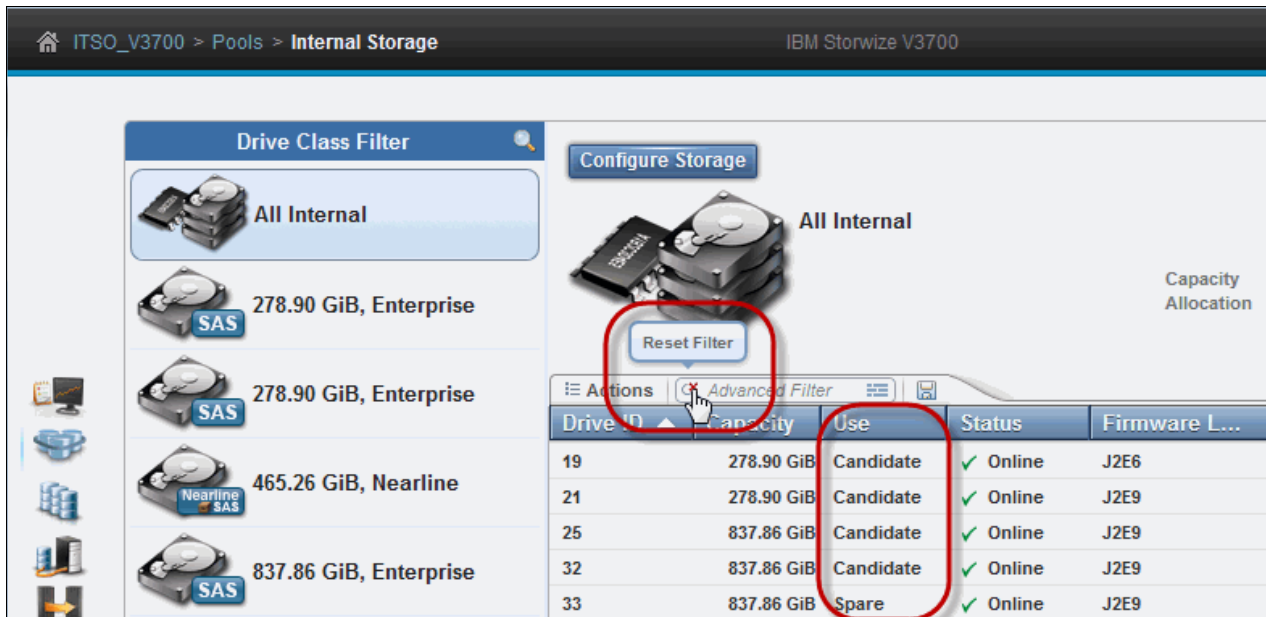


Figure 3-27 Advanced Filter: Results and reset of filter

Saving tabular data to a file

It is possible to save tabular data to a file:

1. Click the diskette icon.
2. Specify the name and location to save the file.

A comma-delimited file (.csv) is created for input to a spreadsheet program, such as Microsoft Excel.

Figure 3-28 shows saving internal storage data.

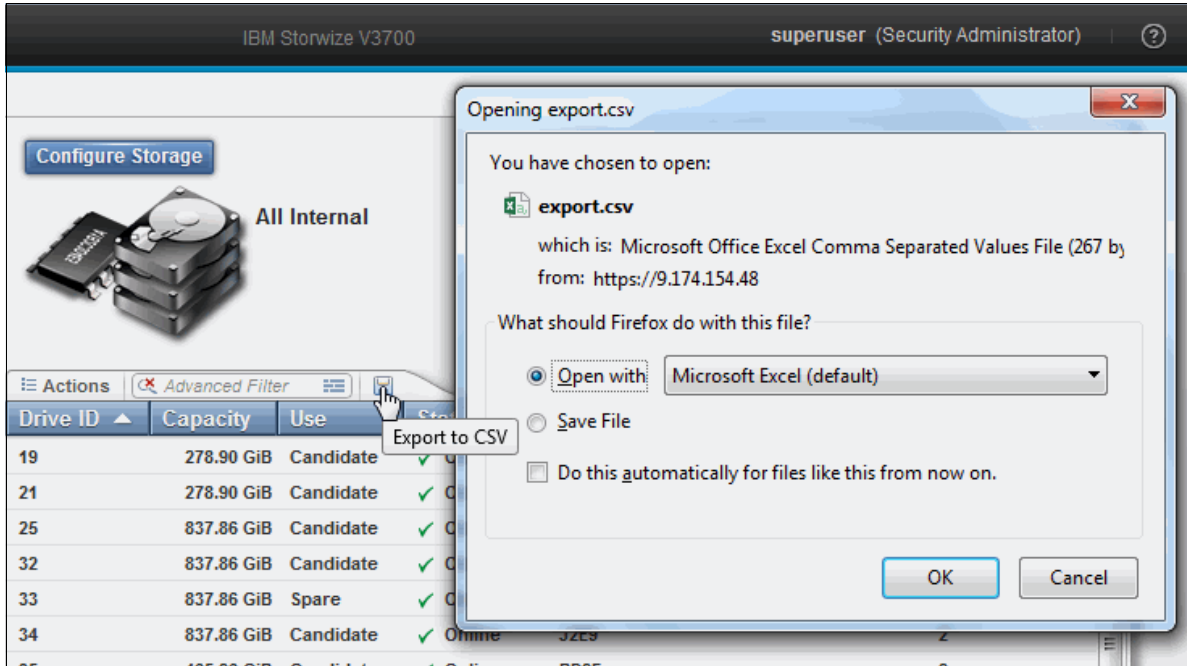


Figure 3-28 Saving internal storage data

3.3 Status Indicators menus

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

3.3.1 Allocated / Virtual Capacity status bar menu

The Allocated / Virtual Capacity status bar shows the capacity status:

1. Hover over the image of two arrows on the right side of this status bar to see a description of the capacity allocation that is in use.

Figure 3-29 shows the comparison of the used capacity to the real capacity.




Figure 3-29 The Allocated / Virtual Capacity bar that compares used capacity to real capacity

2. To change the capacity comparison, click the Allocated / Virtual status bar. This toggles between Allocated and Virtual.

Figure 3-30 shows the new comparison of virtual capacity to real capacity.




Figure 3-30 Changing the Allocated / Virtual comparison, virtual capacity to real capacity

3.3.2 Running Tasks bar menu

To show the Running Tasks menu, complete the following steps:

1. Click the circular image to the left of the running tasks status bar. This menu lists running and recently completed tasks and groups similar tasks.

Figure 3-31 shows the Running Tasks menu.

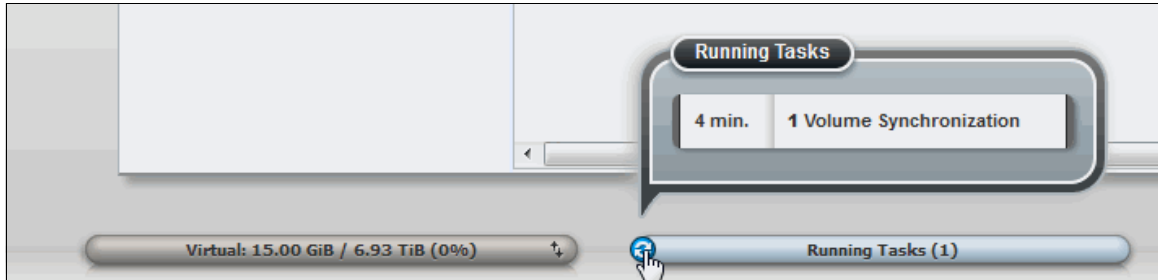


Figure 3-31 Running Tasks menu

2. For an indication of the task progress, browse to the Running Tasks bar menu and click the required task. This opens the task and show its progress, as shown in Figure 3-32.




Figure 3-32 Volume Synchronization task pane

3.3.3 Health Status bar menu

The health status bar provides an indication of the overall health of the system. The color of the status bar indicates the state of IBM Storwize V3700:

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

If a status alert occurs, the Health Status bar can turn from green to yellow or to red. To show the health status menu, complete the following steps:

1. Click in the attention icon on the left side of the health status bar, as shown in Figure 3-33.

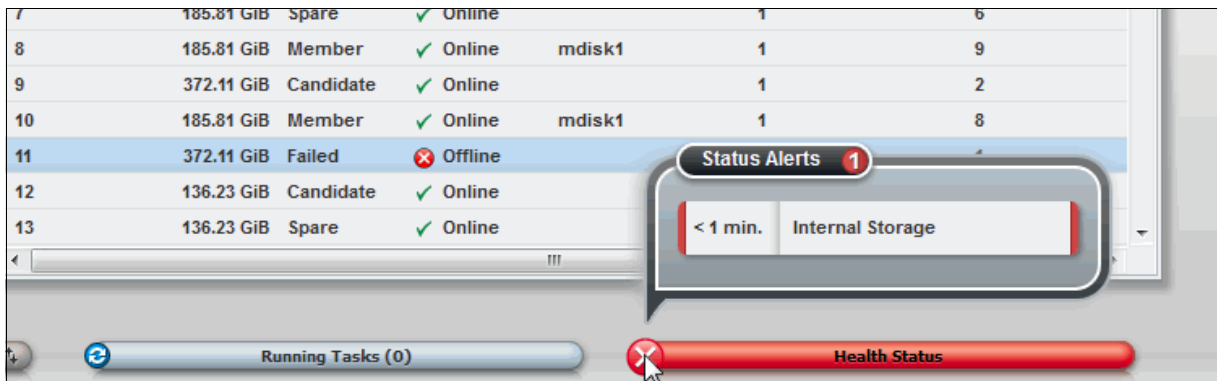


Figure 3-33 Health Status menu

In this example, the Health Status bar menu shows the system as Unhealthy, and provides a description of *Internal Storage* as the type of event that occurred.

2. To investigate the event, open the health status bar menu and click the description of the event to show the Events pane (**Monitoring** → **Events**), as shown in Figure 3-34.

This pane lists all events, and provides Directed Maintenance Procedures (DMPs) to help resolve errors. For more information, see “Events pane” on page 100.

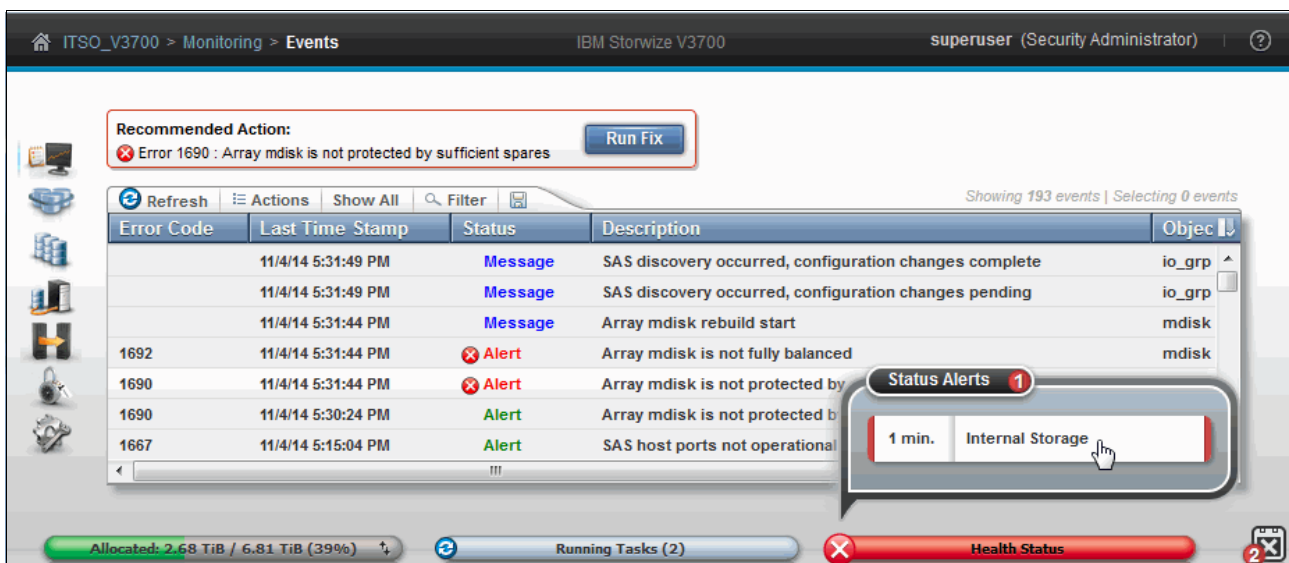


Figure 3-34 Events pane from the Health Status menu

3.4 Function icon menus

The IBM Storwize V3700 management GUI provides function icons that are an efficient and quick mechanism for navigation. As described in section 3.1.3, “System pane layout” on page 74, each graphic on the left side of the pane is a function icon that presents a group of interface functions. Hovering over one of the seven function icons shows a menu that lists the functions. Figure 3-35 shows all of the function icon menus.

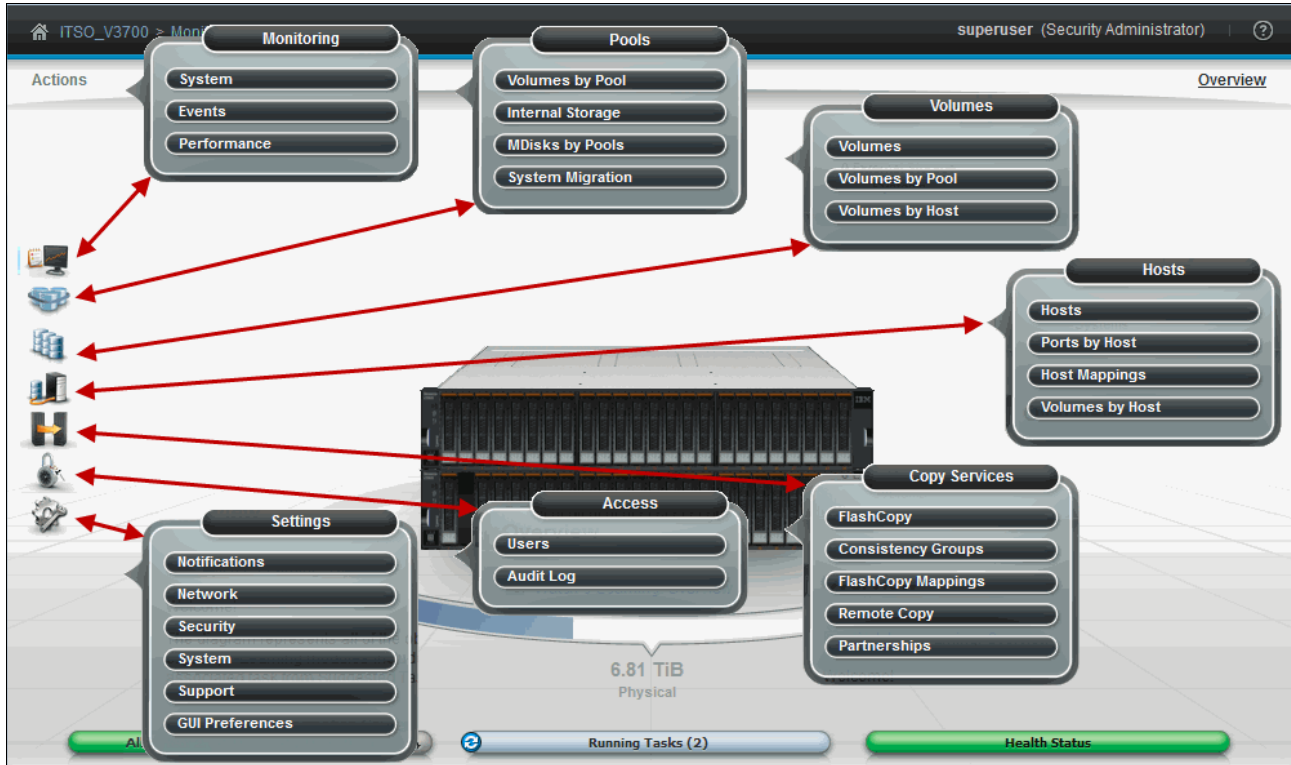


Figure 3-35 All Function Icon menus

3.4.1 Monitoring menu

As shown in Figure 3-36, the Monitoring menu provides access to the System, Events, and Performance panes.




Figure 3-36 Monitoring menu

System pane

To access the System pane, complete the following steps:

1. Select **System** in the Monitoring menu to open the pane. This pane shows a 3D graphical representation of the front of the system, as shown in Figure 3-37.

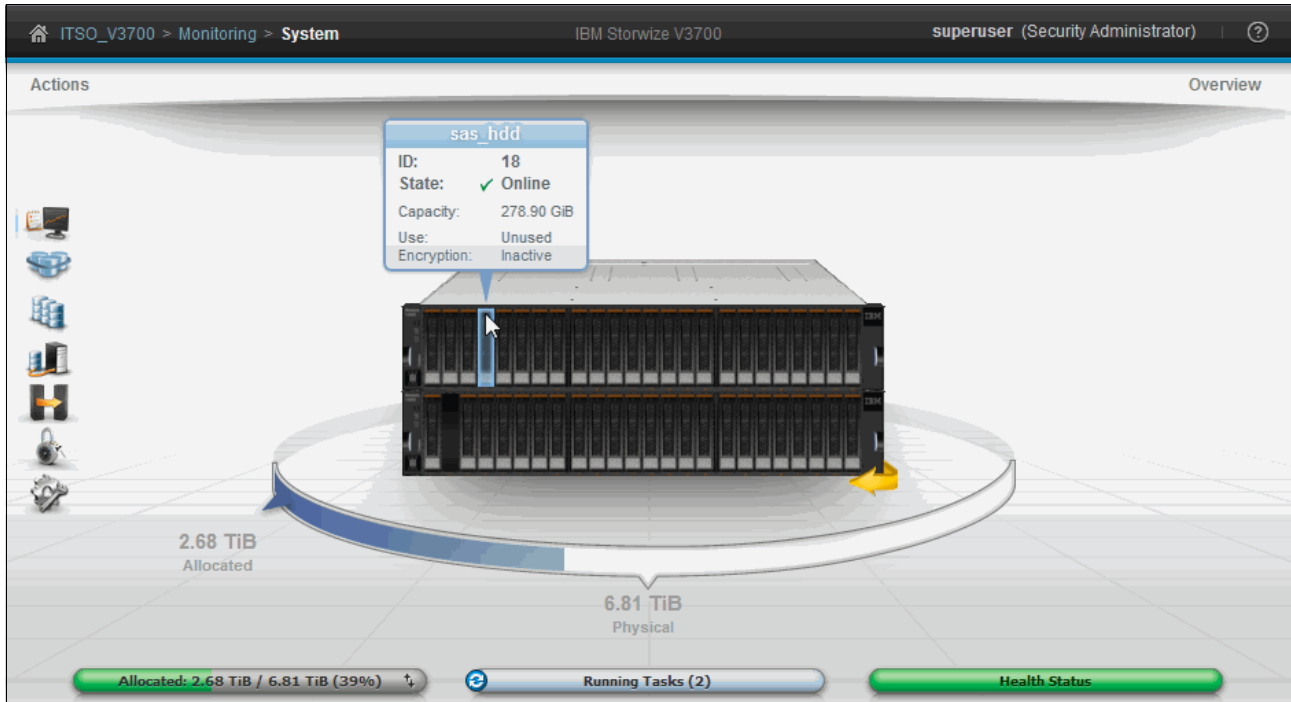


Figure 3-37 System pane: Front view

2. The following actions are possible:
 - Hover over any of the components to show a component overview.
 - Right-click a component to show an Action menu appropriate to that component.
 - Click the arrow at the lower-right of the graphic, as shown in Figure 3-37, to rotate it to show the rear of the system.

This arrow can be blue, yellow, or red. Blue indicates no issues, yellow indicates a warning, and red indicates an issue of some kind with a component.

3. Click the arrow at the lower-right of the graphic again, as shown in Figure 3-38, to rotate it to show the front of the system.

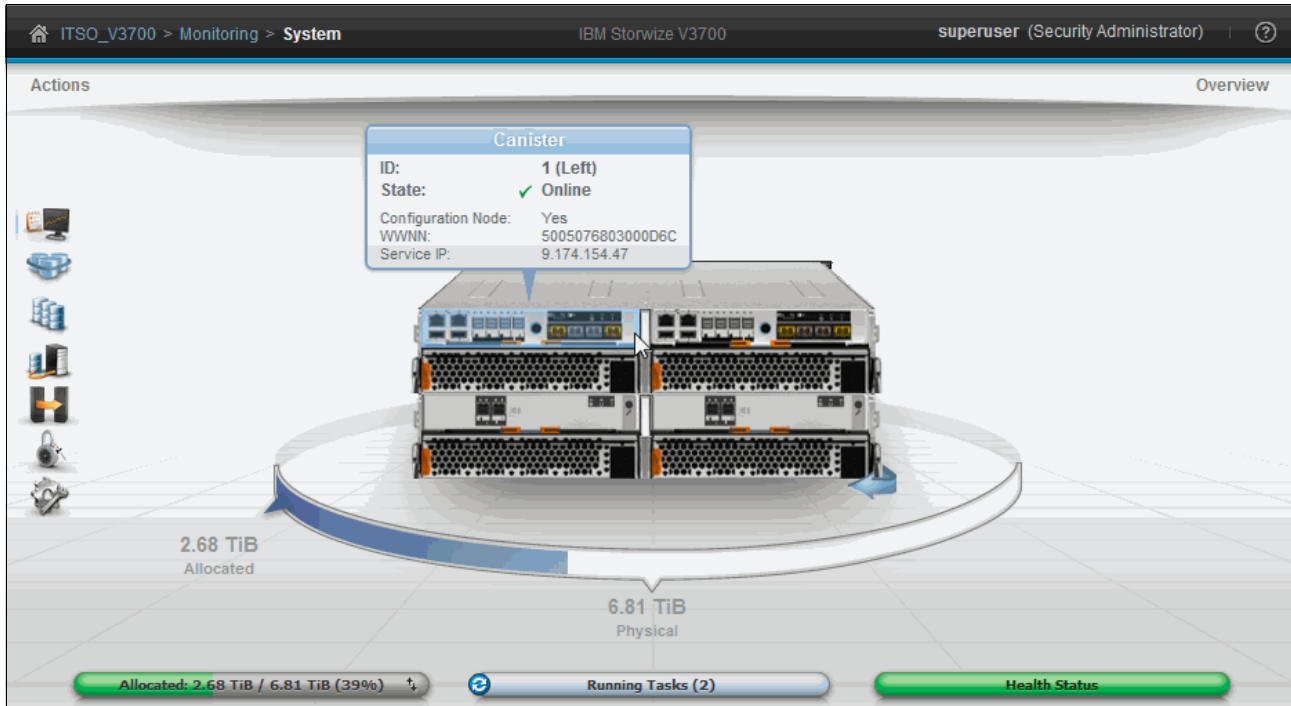


Figure 3-38 System pane: Rear view

The curved band below the graphic of the system, shows capacity usage, including physical installed capacity, the amount of capacity allocated and over-provisioned capacity.

Actions menu

The Actions link at the upper-left of the System pane opens an Actions menu, as shown in Figure 3-39.

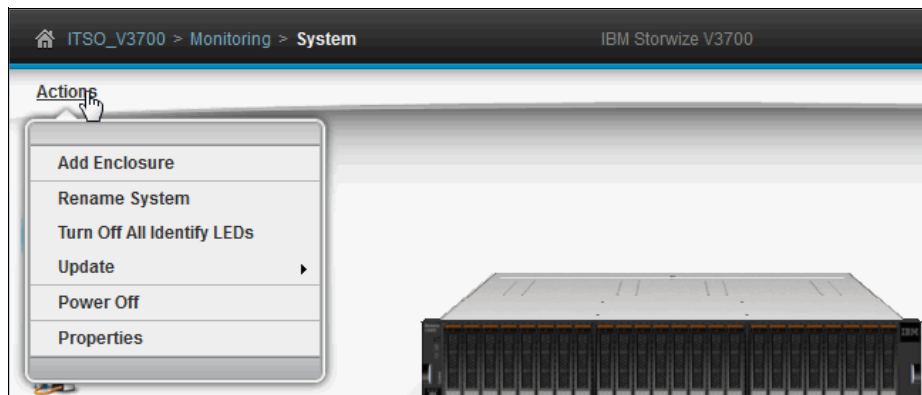


Figure 3-39 System pane Actions menu

The Actions menu enables the user to do the following:

- ▶ Add a new enclosure.
- ▶ Rename the System.
- ▶ Turn off any illuminated Identity LEDs.
- ▶ Update the System and Drive software.
- ▶ Power Off the system.
- ▶ Display System properties, which includes the software version and environment.

Overview

The Overview link at the upper-right of the System pane toggles an Overview pane, as shown in Figure 3-40.

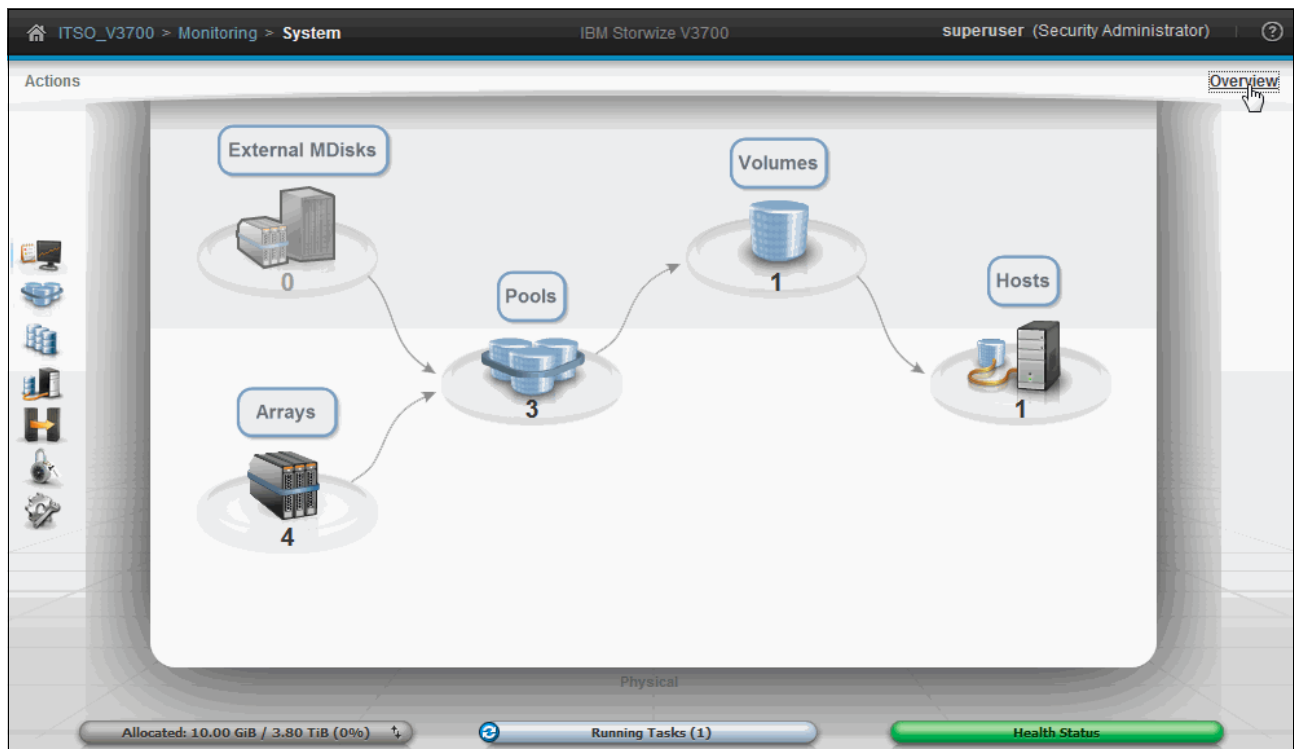


Figure 3-40 Overview pane showing all labels

Clicking any of the Overview icons takes the user to the associated pane. As an example, clicking Arrays takes the user to the **Pools** → **Internal Storage** pane.

System 3D view

On the System pane, hovering over any of the components in the 3D view shows a component overview and right-clicking a component shows an Action menu appropriate to that component.

The following component properties can be displayed:

- ▶ System

The System properties pane shows various information, such as software version, number of enclosures and environmental details.

Figure 3-41 shows the System properties.

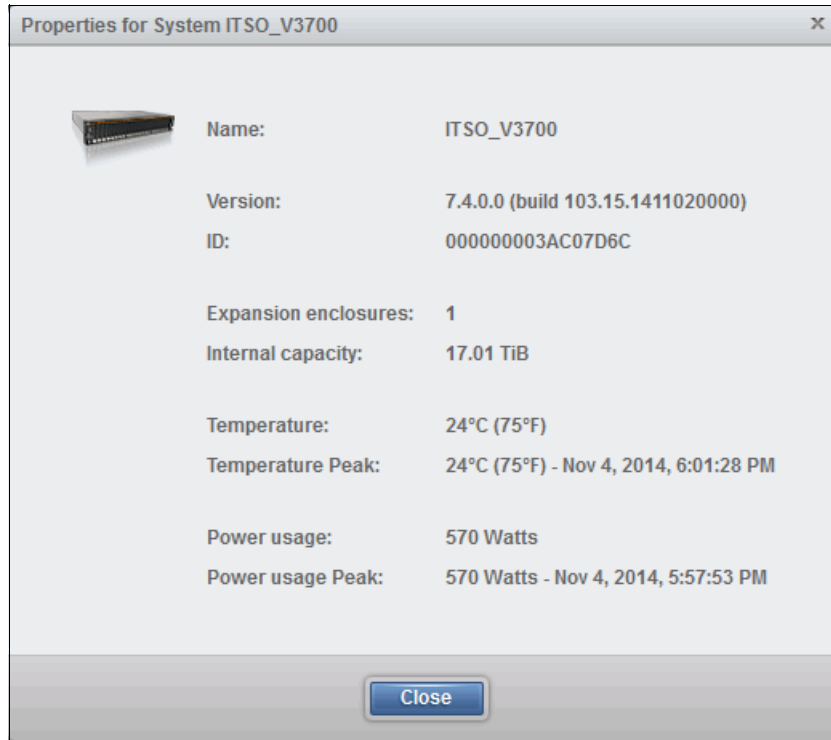


Figure 3-41 Properties for System

► Enclosure

The Enclosure properties pane shows various information, such as the enclosure state, the machine type and serial number, and field-replaceable unit (FRU) part number. Figure 3-42 shows the Enclosure (Control or Expansion) properties.

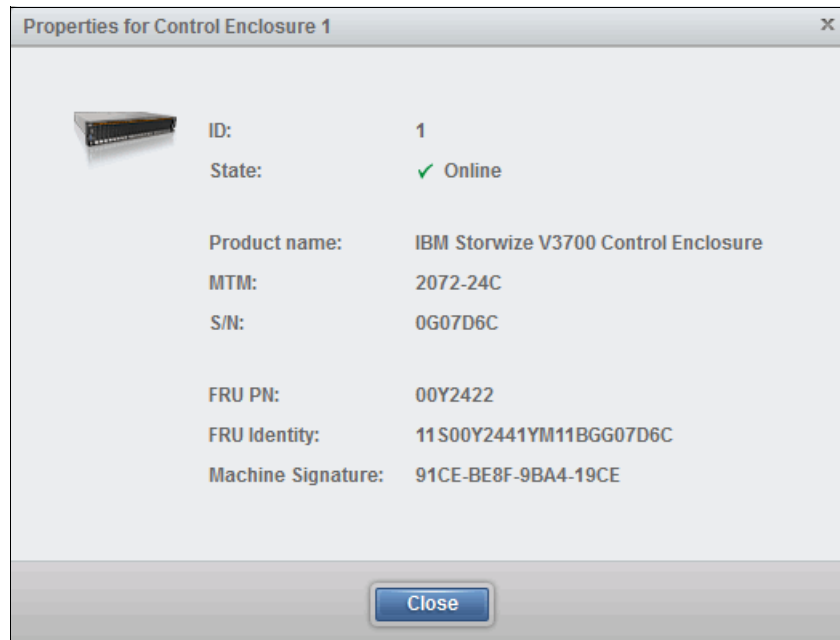


Figure 3-42 Properties for Control Enclosure

► Drive

The Drive properties pane shows various information, such as state, capacity, use (unused, candidate, spare, and member), the drive specification, and its FRU part number. Figure 3-43 shows the Properties for drive 13.

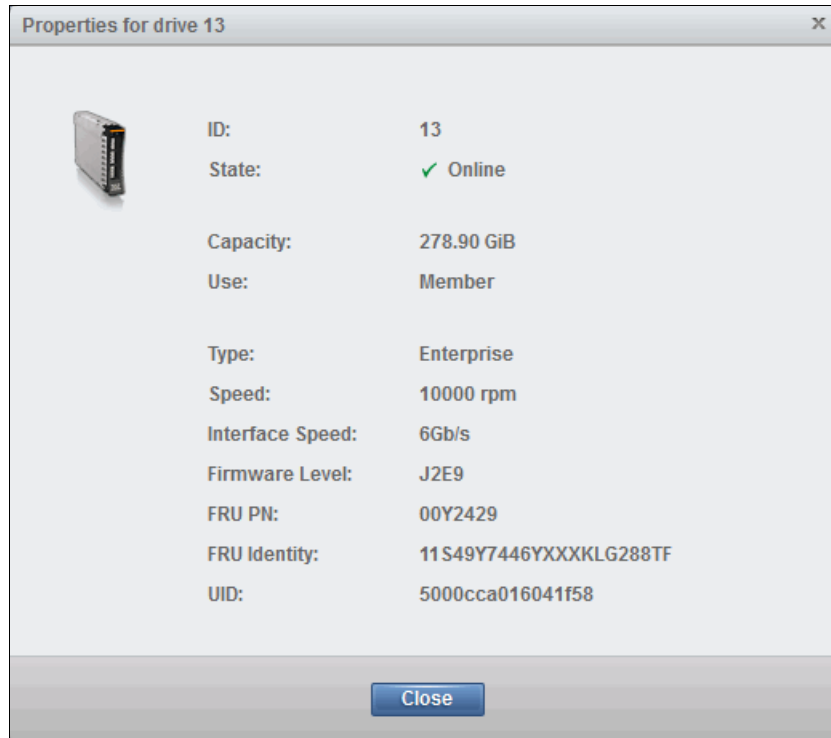


Figure 3-43 Properties for drive

► Canister

The Canister properties pane shows various information, such as Configuration node state, WWNN, memory and central processing unit (CPU) specifications, and FRU part number. Figure 3-44 shows the canister properties.

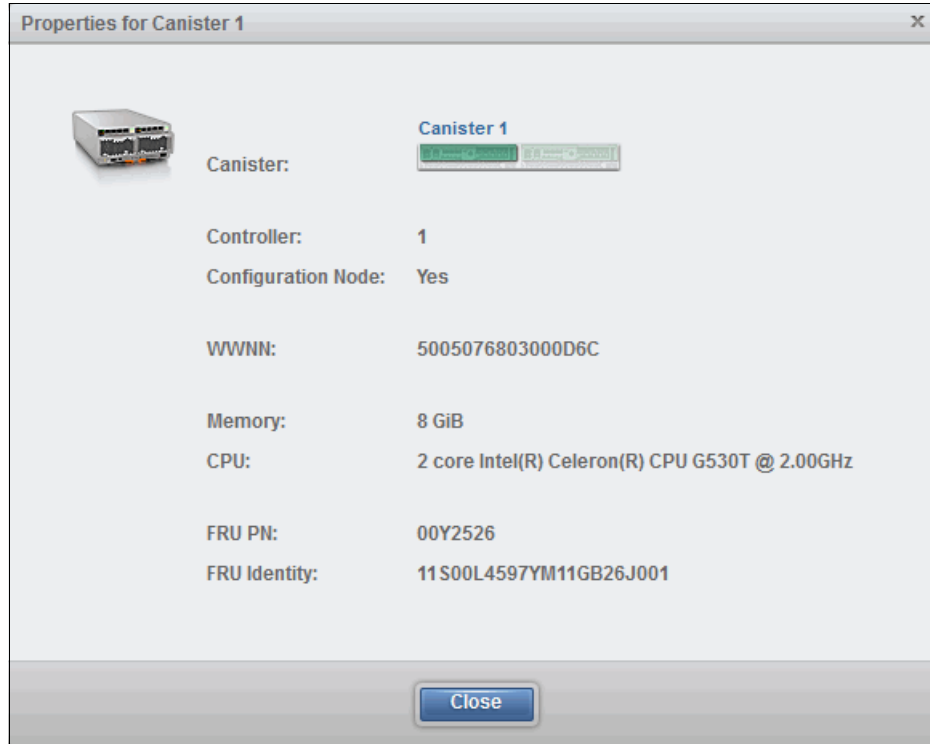


Figure 3-44 Properties for Canister

► Power Supply

Right-clicking the Power Supply shows the Power Supply properties (Figure 3-45) and Battery properties (Figure 3-46 on page 100). The Power Supply properties pane shows various information, such as state and FRU part number.

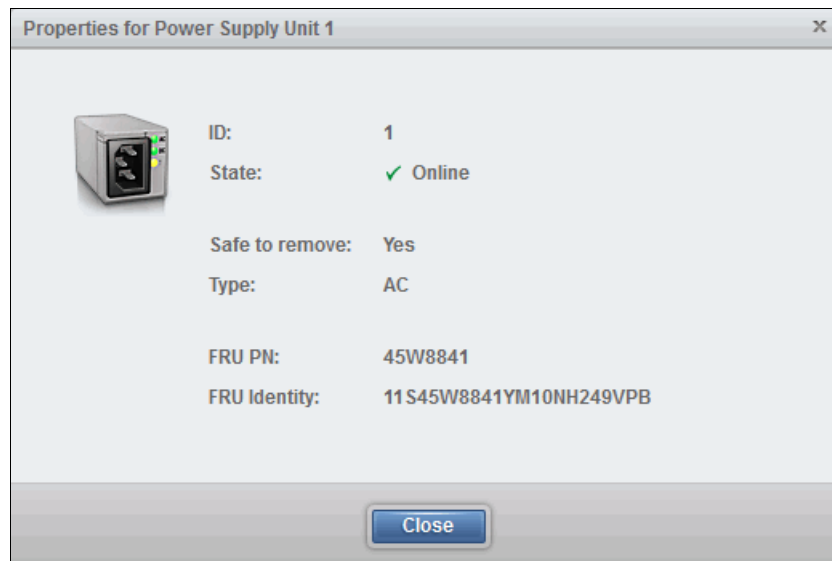


Figure 3-45 Properties for Power Supply Unit

The Battery properties pane shows various information, such as state, charge state, and FRU part number, as shown in Figure 3-46.

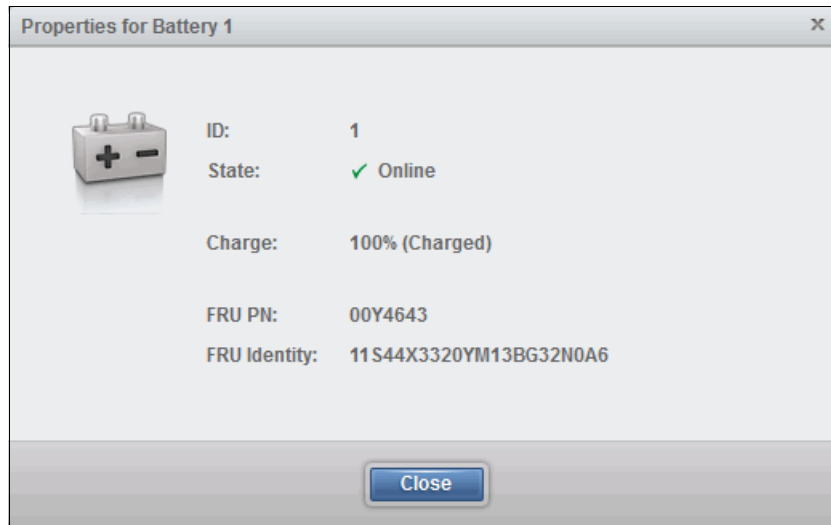


Figure 3-46 Properties for Battery

Events pane

Select **Events** in the Monitoring menu to open the Events pane. The machine is optimal when all errors are addressed and no items are found in this pane, as shown in Figure 3-47.




Figure 3-47 Events pane with all errors addressed

Filtering events view

To view Recommended Actions, Unfixed Messages and Alerts, or to Show All, select the appropriate option from the Recommended Actions menu, as shown in Figure 3-48.

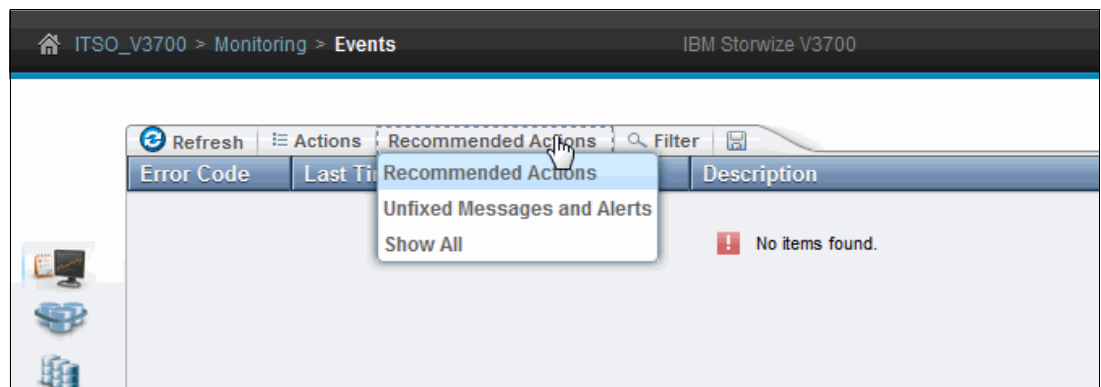


Figure 3-48 Recommended Actions menu

Event properties

To show actions and properties that are related to an event, or to repair an event that is not the Next Recommended Action, right-click the event to show other options.

Figure 3-49 shows the selection of the **Properties** option.

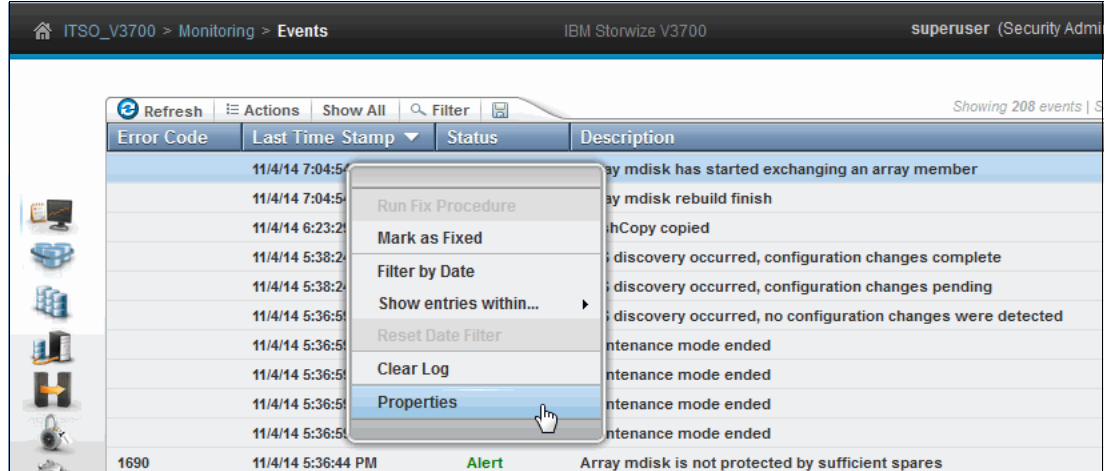


Figure 3-49 Selecting event Properties

Figure 3-50 shows the properties and sense data for an event.

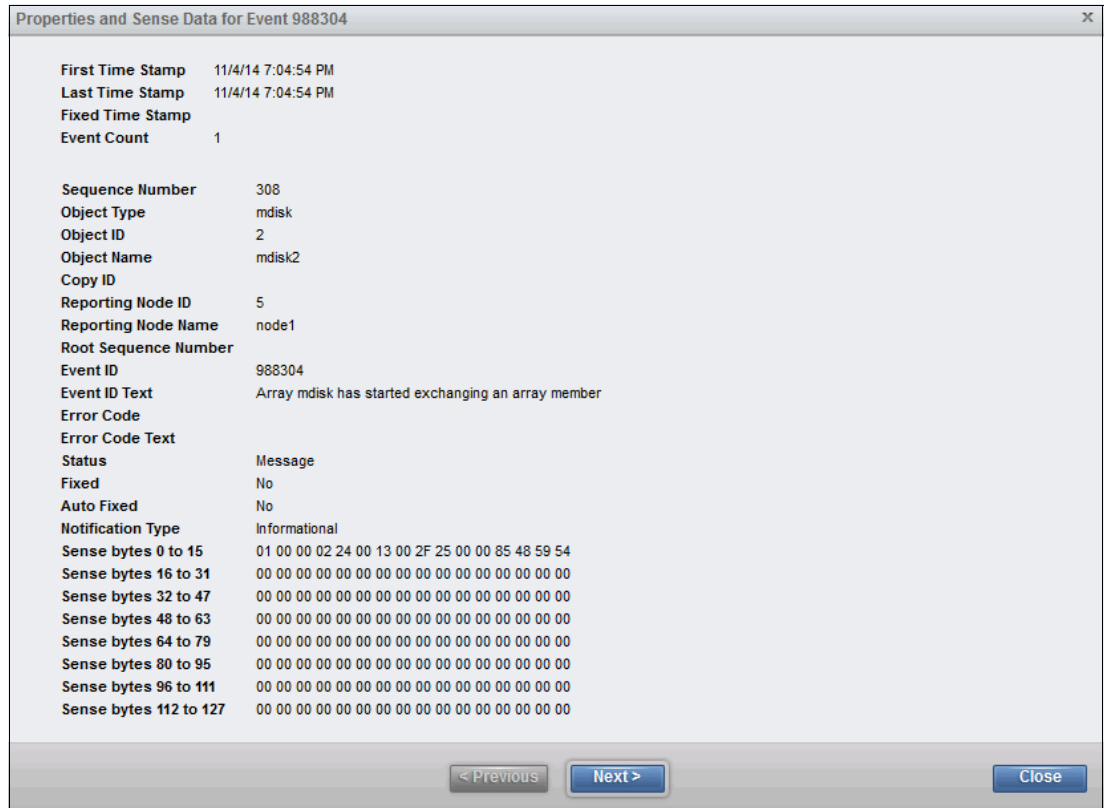


Figure 3-50 Properties And Sense Data for Event

Show entries within

To show events that occurred within a certain time of a particular event, complete the following steps:

1. Select the required event entry.
2. Select **Show entries within** from the Actions menu and set the period value.

Figure 3-51 shows the selection of the **Show entries within** option with a period value of 10 minutes. This shows all events within 10 minutes of the selected event.

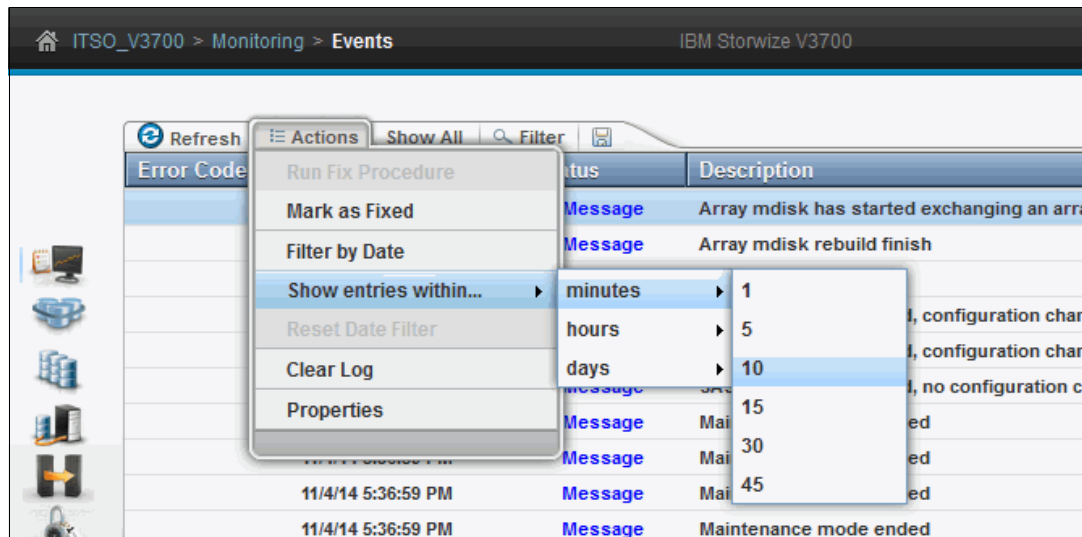


Figure 3-51 Showing Events within a set time

Performance pane

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

Figure 3-52 shows the Performance pane.



Figure 3-52 Performance pane

Custom-tailoring performance graphs

The Performance pane can be customized to show the workload of a single node, which is useful to help determine whether the system is working in a balanced manner. Figure 3-53 shows the custom-tailoring of the performance graphs by selecting **node 1** from the System Statistics menu.

The measurement type can also be changed between throughput in megabytes per second (MBps) or input/output (I/O) operations per second (IOPS) by selecting the relevant value.

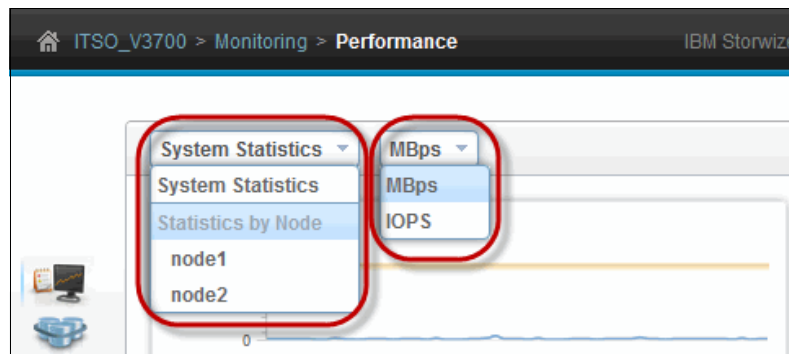


Figure 3-53 Graphs representing Performance System Statistics of a single node

Performance peak value

Peak values over the last 5-minute period can be seen by hovering over the current value, as shown in Figure 3-54 for the serial-attached Small Computer System Interface (SAS).

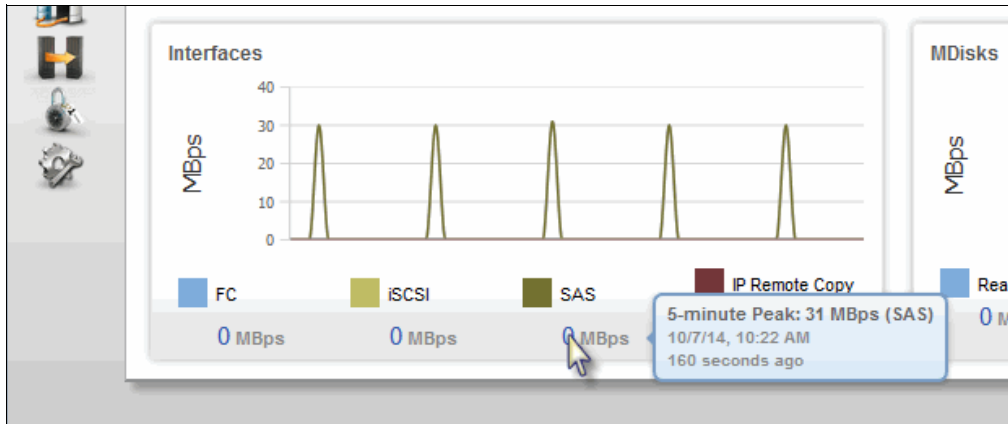


Figure 3-54 Peak SAS Interface usage value over the last 5 minutes

3.4.2 Pools menu

The Pools menu provides access to the Volumes by Pools, Internal Storage, MDisks by Pools, and System Migration functions, as shown in Figure 3-55.




Figure 3-55 Pools menu

Volumes by Pool pane

Select **Volumes by Pool** in the Pools menu to open the pane. The Pool Filter can be used to view volumes associated to a particular Pool. This view makes it easy to manage volumes and determine the amount of real capacity that is available for further use.

Figure 3-56 shows the Volumes by Pool pane.

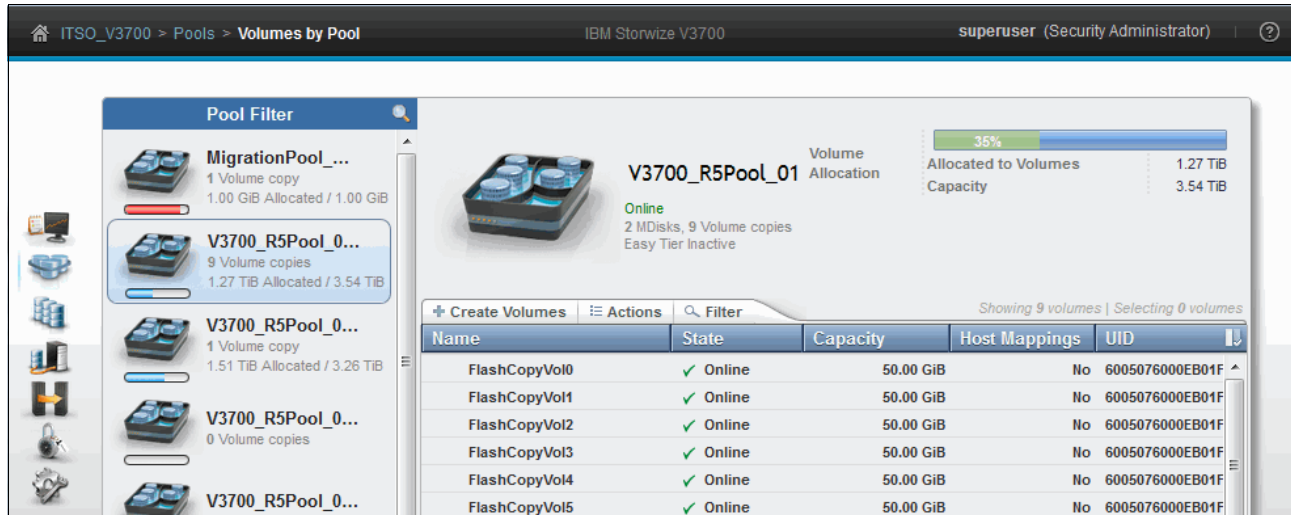


Figure 3-56 Volumes by Pools pane

Volume Allocation

The upper-right corner of the Volumes by Pool pane shows the Volume Allocation, which, in the following example, shows a physical capacity of 738.00 gibibytes (GiB) and a used capacity (52.00 GiB in the green portion). The red bar shows the threshold at which a warning is generated when the used capacity in the pool first exceeds the threshold that is set for its physical capacity. By default, this threshold is set to 80%, but can be altered in the Pool properties.

Figure 3-57 shows the volume allocation information that is displayed in the Volumes by Pool pane.




Figure 3-57 Volume Allocation

Renaming pools

To rename a pool, complete the following steps:

1. Select the pool from the pool filter.
2. Click the name of the pool.

Figure 3-58 shows that pool V3700_R5Pool_01 was selected to be renamed.

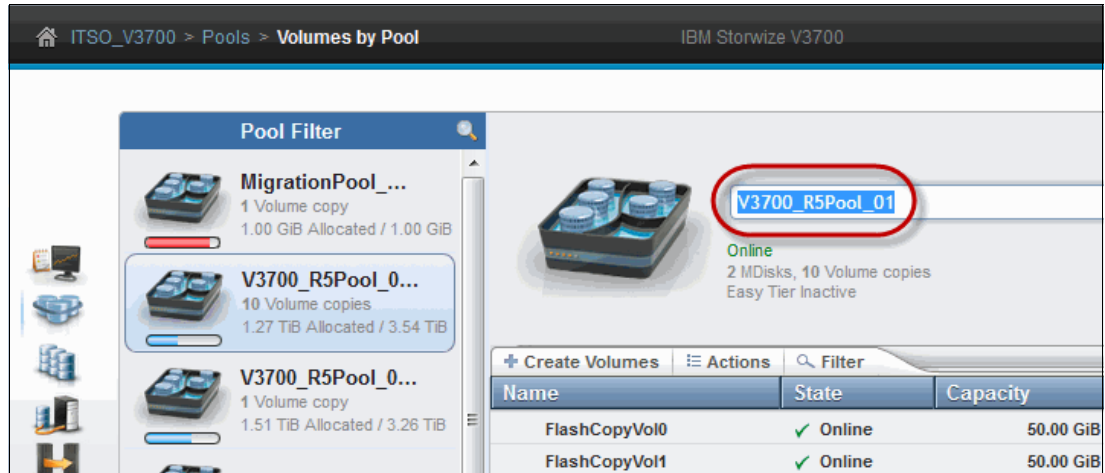


Figure 3-58 Renaming a pool

Changing pool icons

To change the icon that is associated with a pool, complete the following steps:

1. Select the pool in the Pool Filter.
2. Click the large pool icon.
3. Use the Choose Icon buttons to select the required image.

This helps to manage and differentiate between the classes of drive, or the tier of the Storage Pool. Figure 3-59 shows the pool change icon pane.

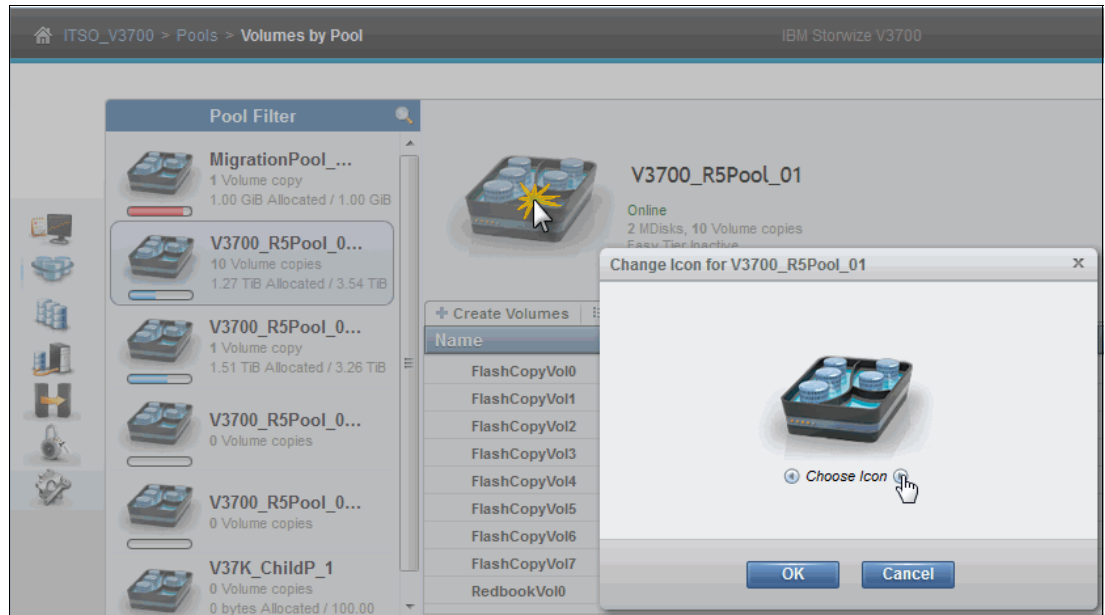


Figure 3-59 Change Icon for a pool

Volume functions

The Volumes by Pool pane also provides access to the volume functions using the Actions menu, the Create Volumes option, and by right-clicking a listed volume. For more information about navigating the Volume pane, see 3.4.3, “Volumes menu” on page 114.

Figure 3-60 shows the volume functions that are available from the Volumes by Pool pane.

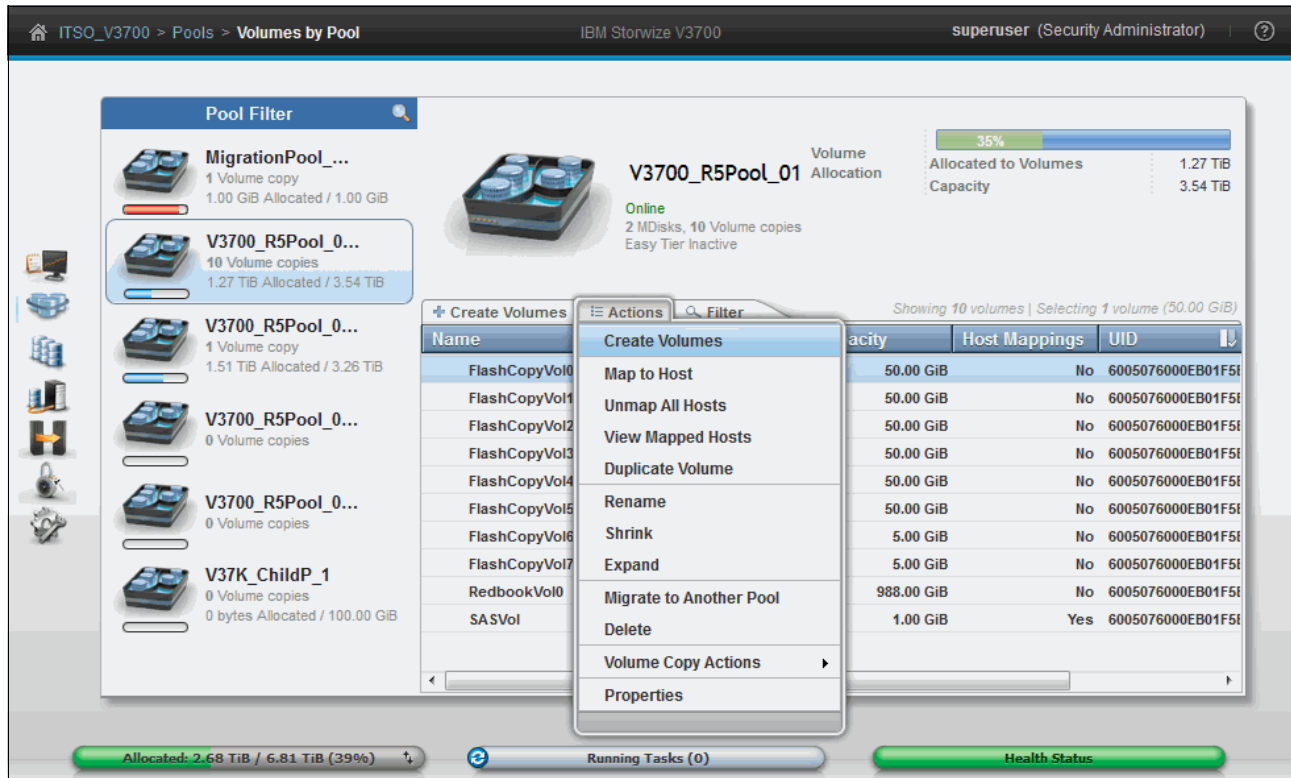


Figure 3-60 Volume functions are available from the Volumes by Pool pane

Internal Storage pane

Select **Internal Storage** in the Pools menu to open the Internal Storage pane, as shown in Figure 3-61. The internal storage consists of the drives that are contained in the IBM Storwize V3700 control enclosure, and any SAS-attached IBM Storwize V3700 expansion enclosures.

Using the Internal Storage pane, it is possible to configure the internal storage into RAID protected storage (MDisks). The Drive Class Filter can be used to display the drive list by drive class.

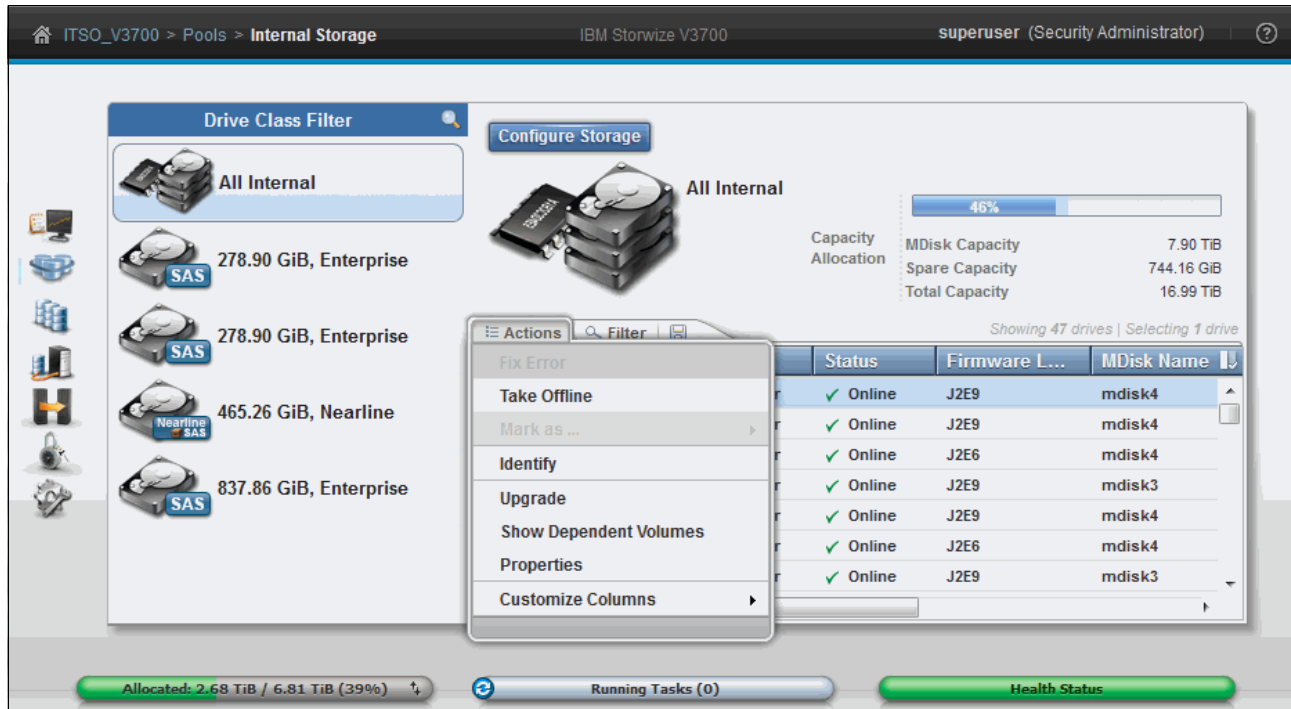


Figure 3-61 Drive Actions menu of the Internal Storage pane

Drive actions

Drive-level functions, such as identifying a drive and marking a drive as offline, unused, candidate, or spare, can be accessed here:

1. Right-click a listed drive to show the Actions menu.
2. Alternatively, the drives can be selected and then the Action menu used.

For more information, see “Multiple selections” on page 85. Figure 3-61 shows the Drive Actions menu.

Drive properties

Drive properties and dependent volumes can be displayed from the Internal Storage pane. Select **Properties** from the Drive Actions menu. The drive Properties pane shows the drive attributes and the drive slot SAS port status. Figure 3-62 on page 109 shows the drive properties with the **Show Details** option selected.

Drive firmware upgrade

In past versions of the IBM Storwize software, the drive firmware could only be upgraded using the CLI. It is now possible to upgrade the firmware using the GUI. The drives can either be upgraded all at once or individually. See Chapter 11, “RAS, monitoring, and troubleshooting” on page 577.

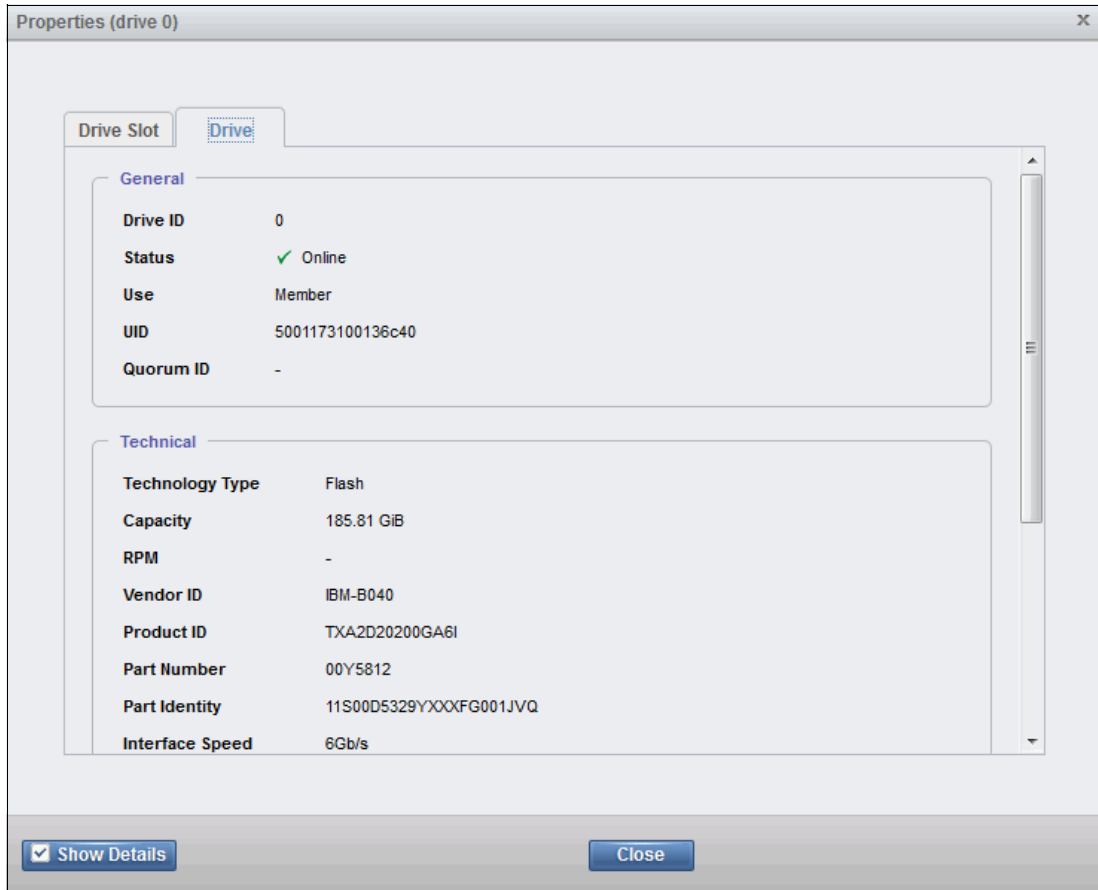


Figure 3-62 Drive Properties

Configure Internal Storage wizard

To show the Configure Internal Storage wizard, perform the following steps:

1. Click **Configure Storage**, as shown in Figure 3-63.



Figure 3-63 Internal Storage pane

2. Using this wizard, is it possible to configure the RAID properties and pool allocation of the internal storage. Figure 3-64 shows the Configure Internal Storage wizard.
- For full details about using the Configure Internal Storage wizard, see Chapter 7, “Storage Pools” on page 307.

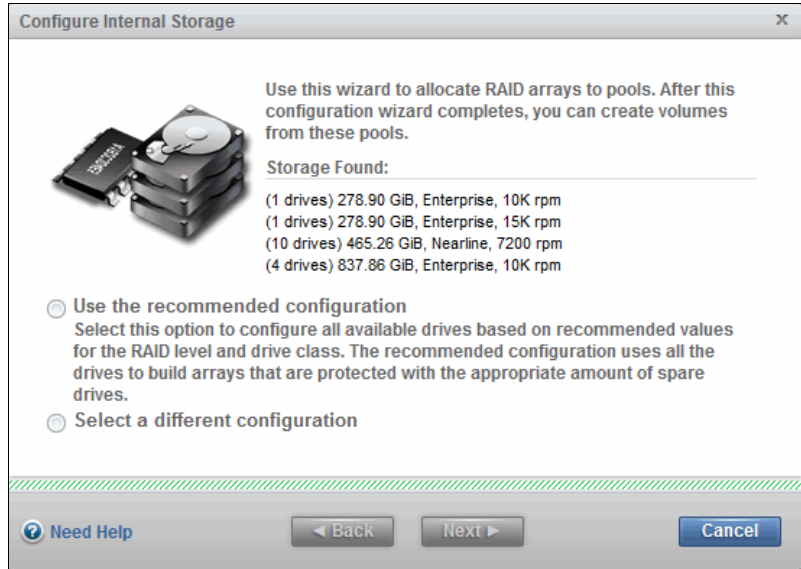


Figure 3-64 Configure Internal Storage wizard

3. Figure 3-65 shows allocating the internal storage to a pool.

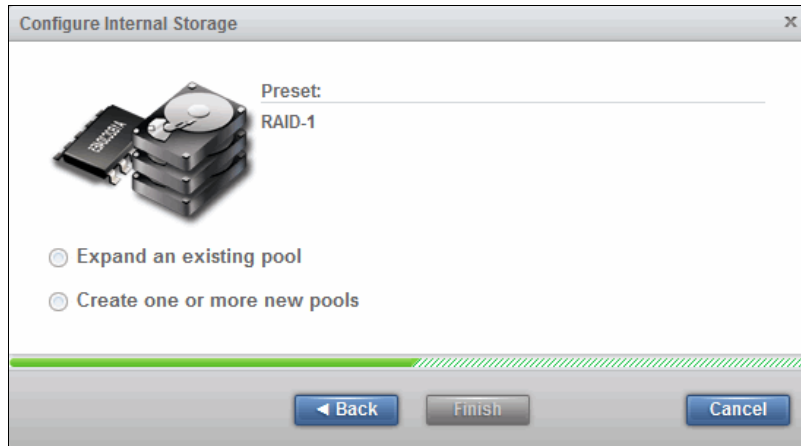


Figure 3-65 Configure Internal Storage wizard: Pool allocation

MDisks by Pool pane

Select **MDisks by Pool** in the Pools menu to open the MDisks by Pool pane. Using this pane, it is possible to perform various tasks, such as display MDisks in each pool, create new storage pools, and delete storage pools.

Figure 3-66 shows the MDisks by Pool pane.

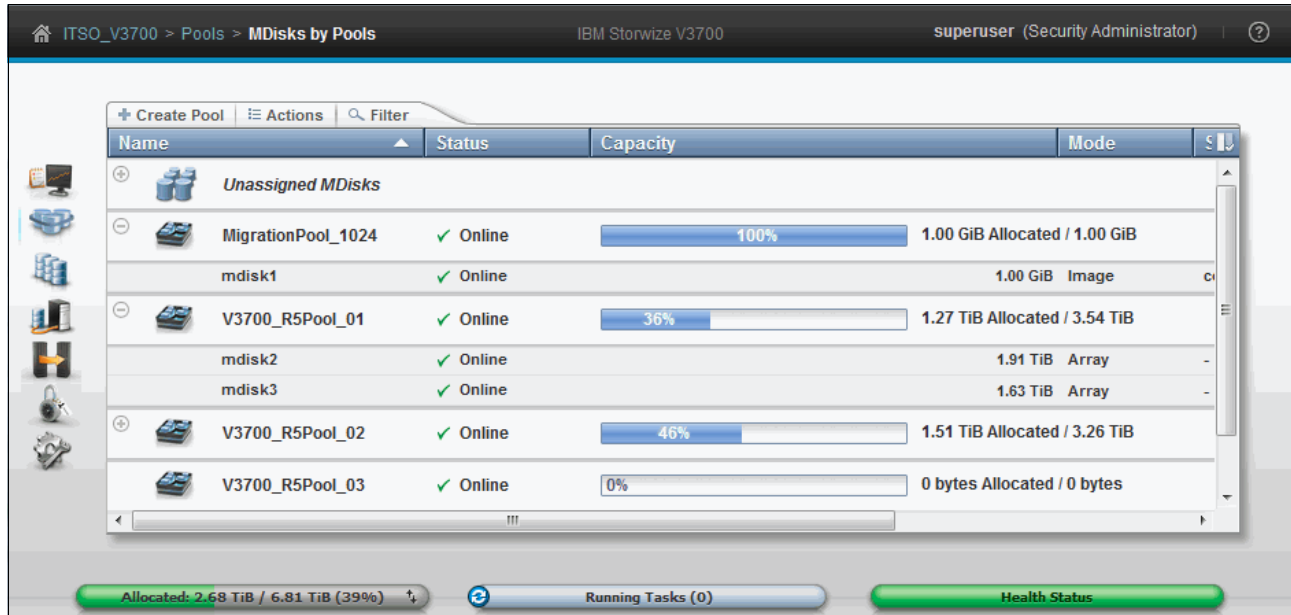


Figure 3-66 MDisks by Pools pane

Pool actions

To create, delete, rename, or change a Pool Icon, view a Child Pool, or view a Pool's properties, complete the following steps:

1. Right-click a pool.
2. Alternatively, the Actions menu can be used.

Figure 3-67 shows the pool actions.

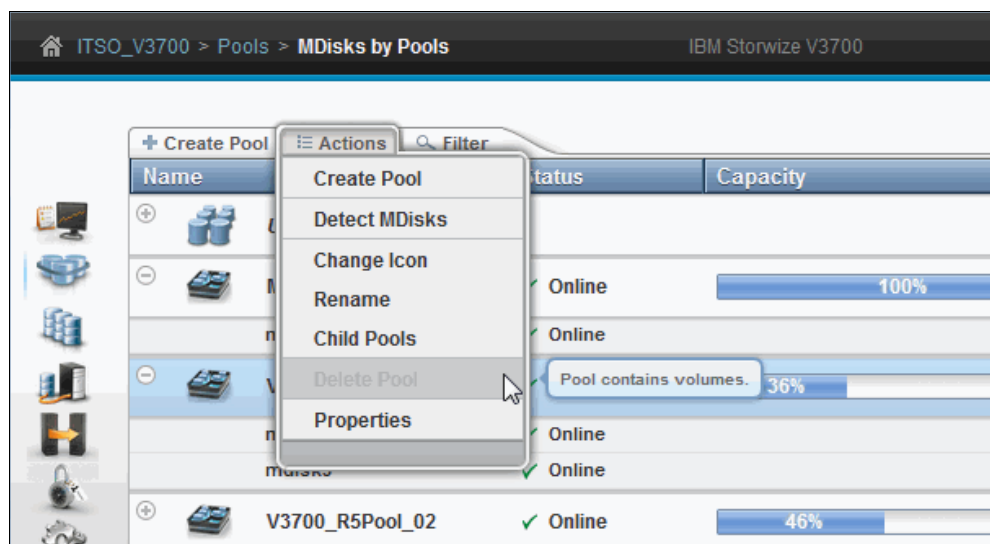


Figure 3-67 Pool Actions

Restriction: In software release 7.4, a pool cannot be deleted if it contains volumes. The volumes must first be deleted using the Volumes by Pool pane. This is a safeguard against accidental volume deletion.

RAID actions

Using the MDisks by Pool pane, it is possible to perform MDisk RAID tasks, such as Set Spare Goal, Swap Drive, and Delete. To access these functions, right-click the MDisk or use the Actions menu, as shown in Figure 3-68.

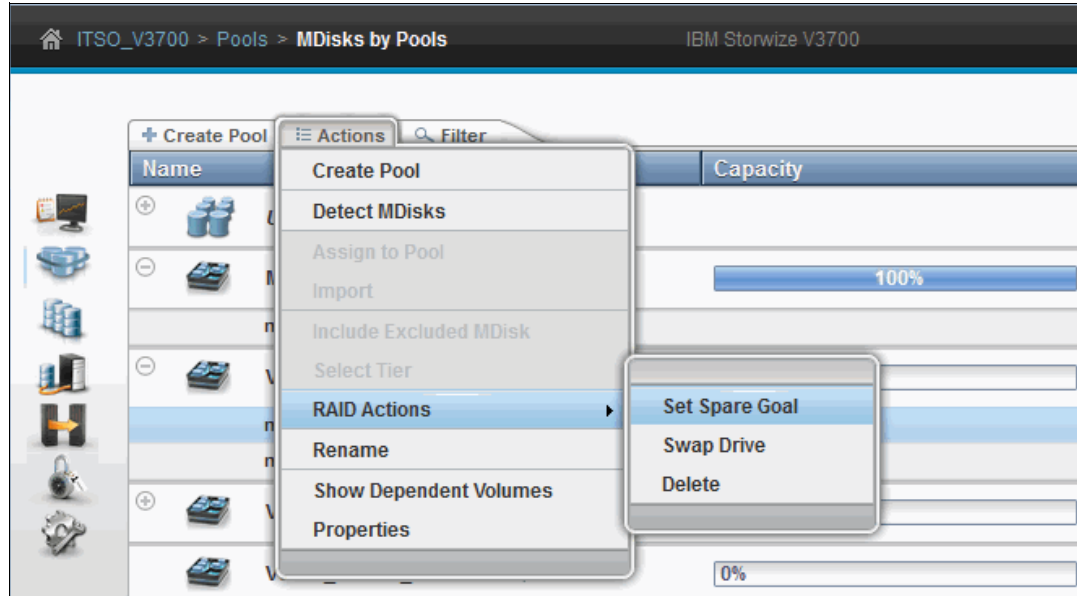


Figure 3-68 RAID Actions menu

System Migration pane

To migrate a system, complete the following steps:

1. Select **System Migration** in the Pools menu to open the System Migration pane, as shown in Figure 3-69. This pane is used to migrate data from externally attached Fibre Channel or SAS-based storage systems to the internal storage of the IBM Storwize V3700. The pane displays image mode volume information.
2. To begin a migration, click **Start New Migration**.

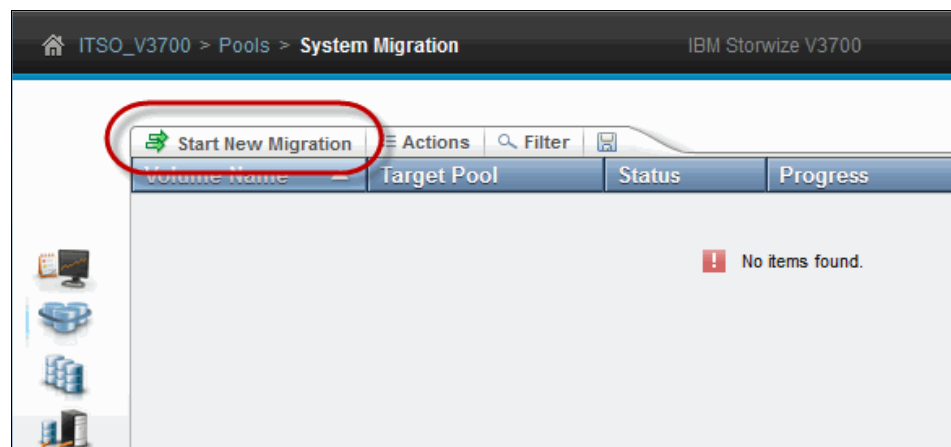


Figure 3-69 System Migration pane

Storage Migration wizard

The Storage Migration wizard is used for data migration from other FC or SAS attached storage systems to the IBM Storwize V3700. Figure 3-70 shows the Storage Migration wizard.

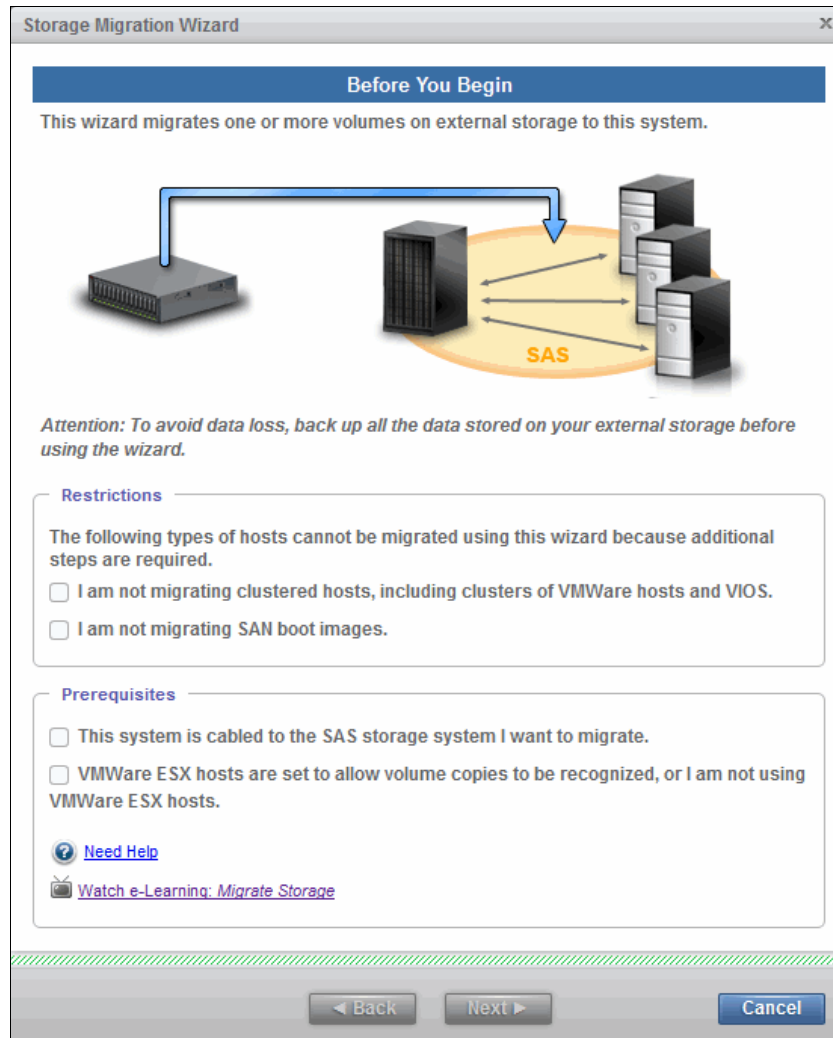


Figure 3-70 Storage Migration wizard

For more information, see Chapter 6, “Storage migration” on page 249.

3.4.3 Volumes menu

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




Figure 3-71 The Volumes menu

Volumes pane

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

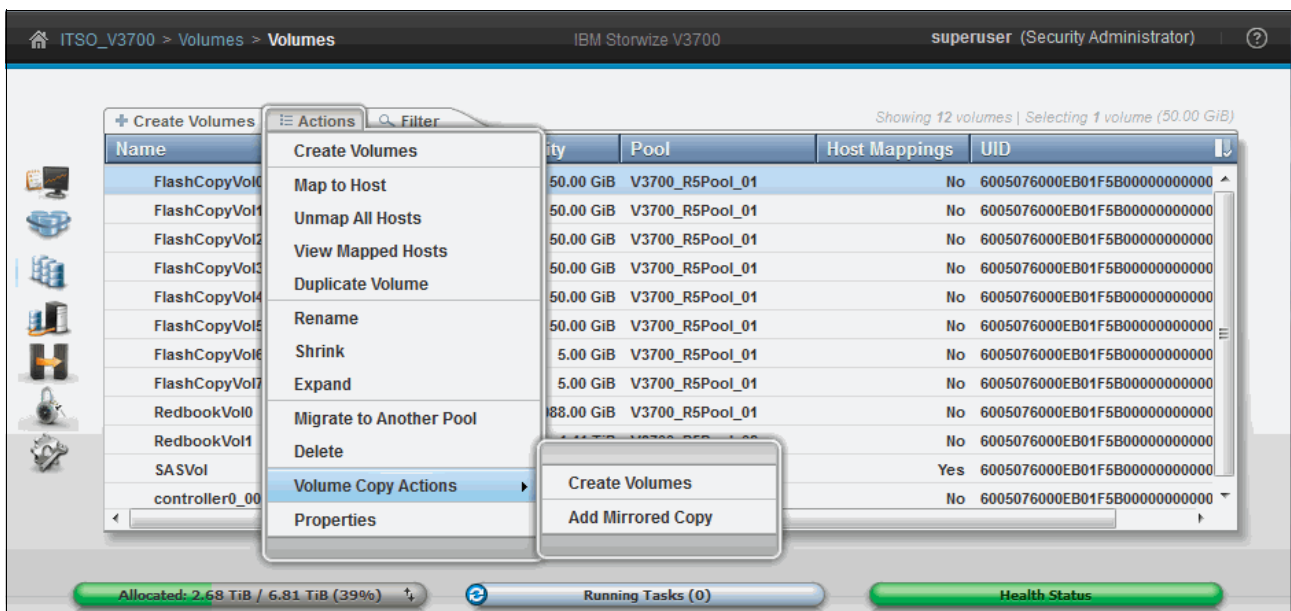


Figure 3-72 Volumes pane

Volume actions

Figure 3-72 on page 114 shows the Volume Action menu and its options.

Create volumes

To create volumes, complete the following steps:

1. Click **Create Volumes** to open the Create Volumes pane, as shown in Figure 3-73.

Using this pane, it is possible to select a preset of the type of volume to create. The presets are designed to accommodate most user cases:

- Generic
- Thin-provisioned
- Mirror
- Thin mirror

2. After a preset is determined, select the Storage Pool from which the volumes are to be allocated. An area to name and size the volumes is shown.

For more information, see Chapter 5, “Volume configuration” on page 199 and Chapter 8, “Advanced host and volume administration” on page 359.




Figure 3-73 Create Volumes pane

Creating multiple volumes

A useful feature is available for quickly creating multiple volumes of the same type and size:

1. Specify the number of volumes required in the Quantity field.
2. Complete the volume capacity and name. A number range can also be specified.

The Create Volumes pane displays a summary that shows the real and virtual capacity that is used if the proposed volumes are created.

3. Click **Create** or **Create and Map to Host** to continue.

Figure 3-74 shows a quantity of 3 in the Quantity field, and this creates three volumes named V3700_Vo1_05, V3700_Vo1_06, and V3700_Vo1_07.

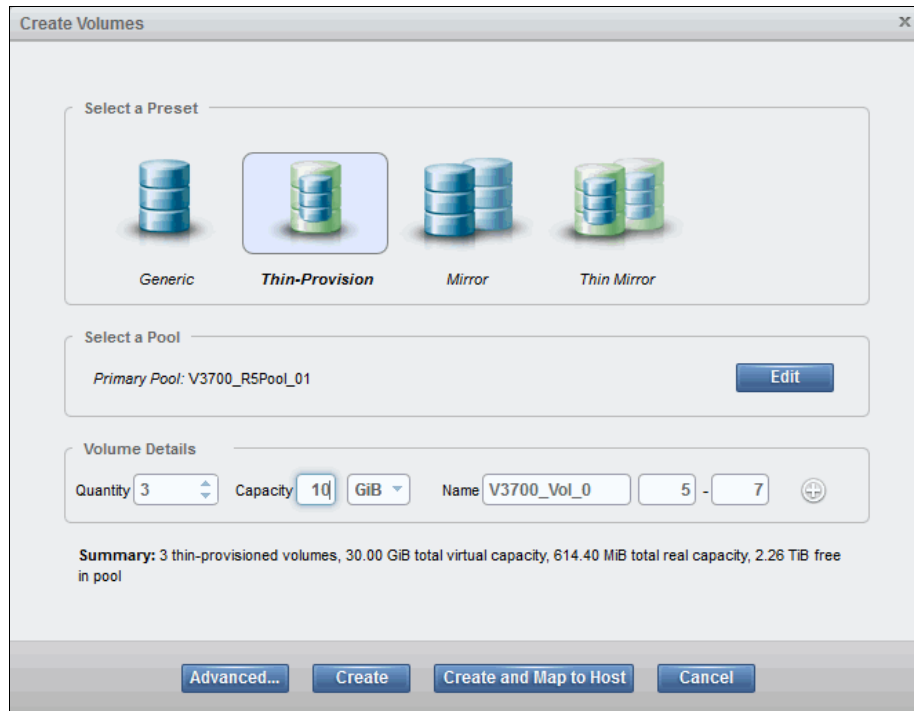


Figure 3-74 Creating multiple volumes quickly

Volume advanced settings

To change other volume settings, complete the following steps:

1. Click **Advanced** to show more volume configuration options. Use this feature when the preset does not meet your requirements.
2. After the advanced settings are configured, click **OK** to return to the Create Volumes pane. Figure 3-75 shows the Advanced Settings pane.

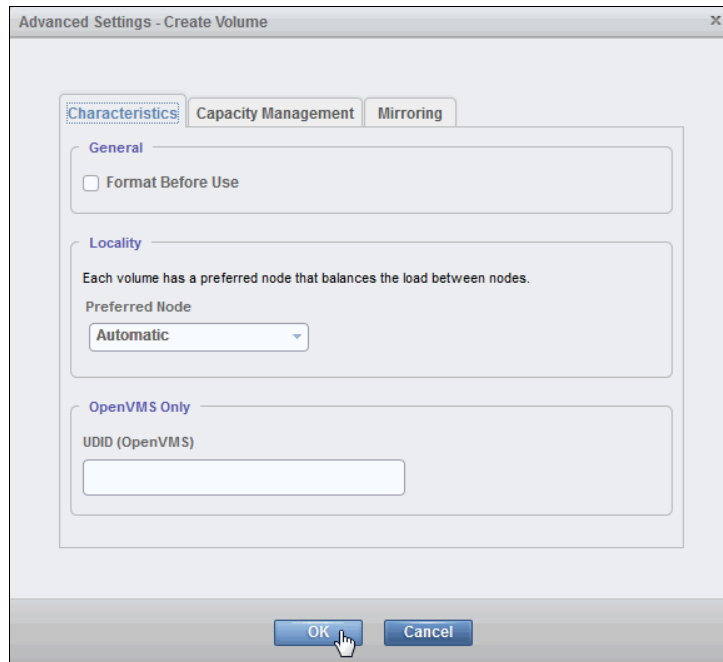


Figure 3-75 Advanced Settings pane

Volumes by Pool pane

For more information, see “Volumes by Pool pane” on page 105.

Volumes by Host pane

Select **Volumes by Host** in the Volumes menu to open the pane. In the Volume by Hosts pane, it is possible to focus on volumes that are allocated to a particular host using the host selection filter.

Figure 3-76 shows the Volumes by Host pane.

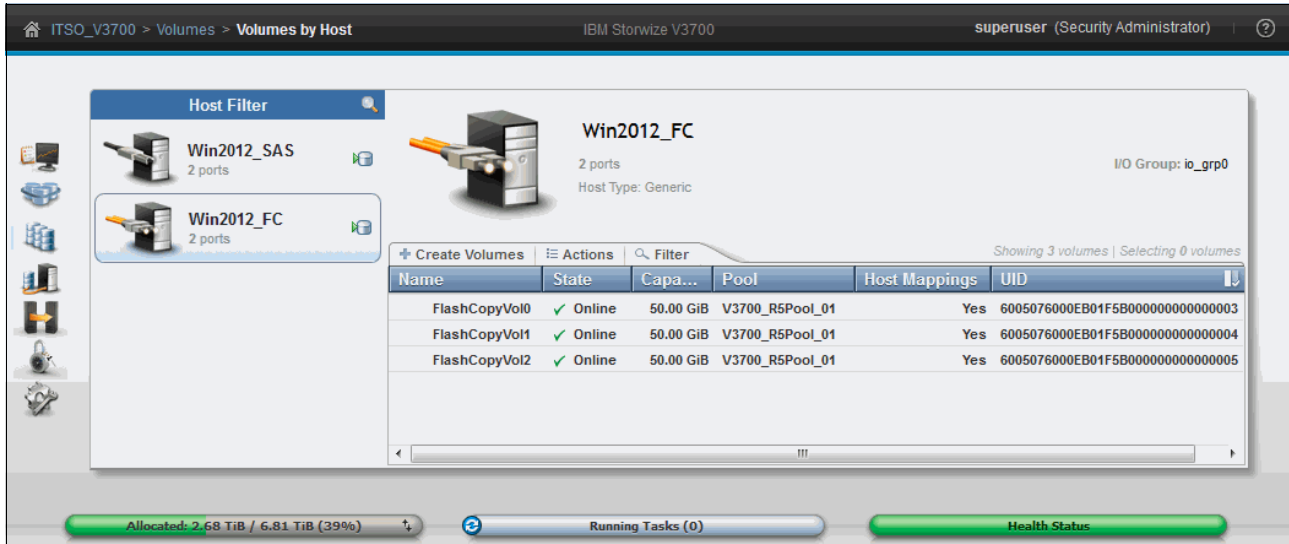


Figure 3-76 Volume by Host pane

3.4.4 Hosts menu

As shown in Figure 3-77, the Hosts menu provides access to the Hosts, Ports by Host, Host Mappings, and Volumes by Host functions.

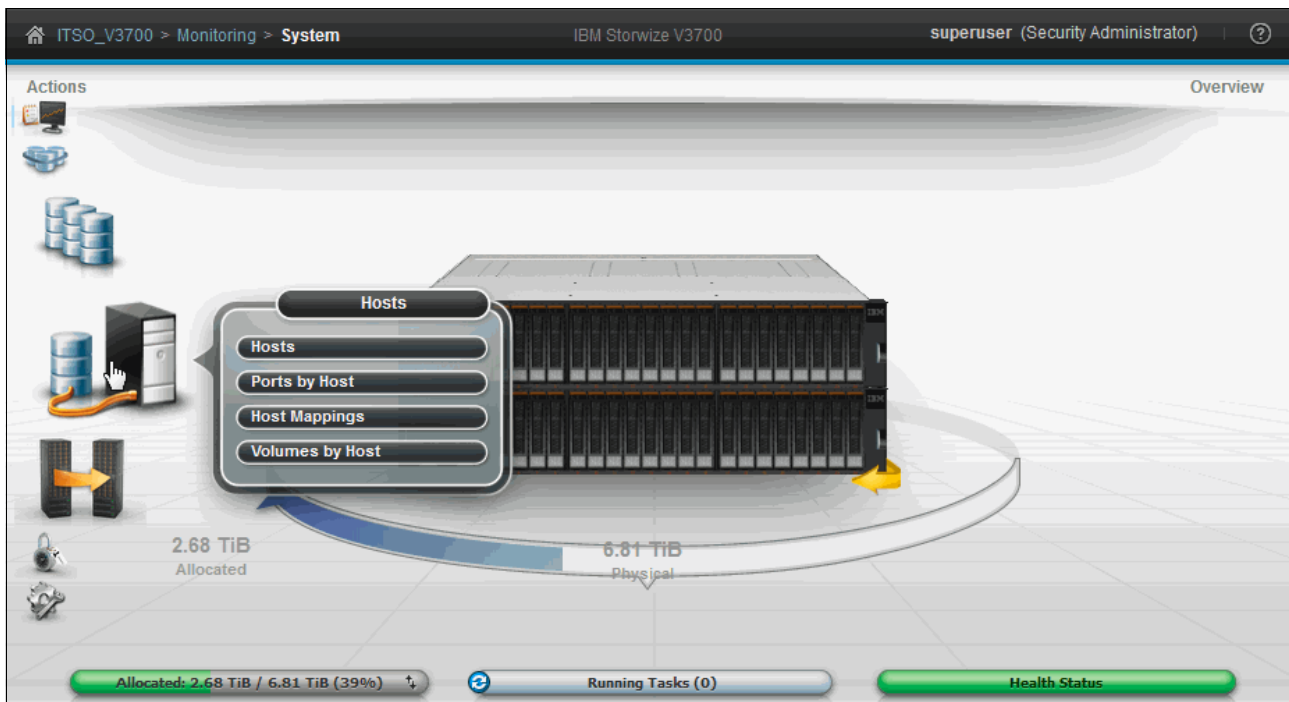


Figure 3-77 The Hosts menu

Hosts pane

Select **Hosts** in the Hosts menu to open the pane, as shown in Figure 3-78. The Hosts pane shows all of the hosts that are defined in the system.

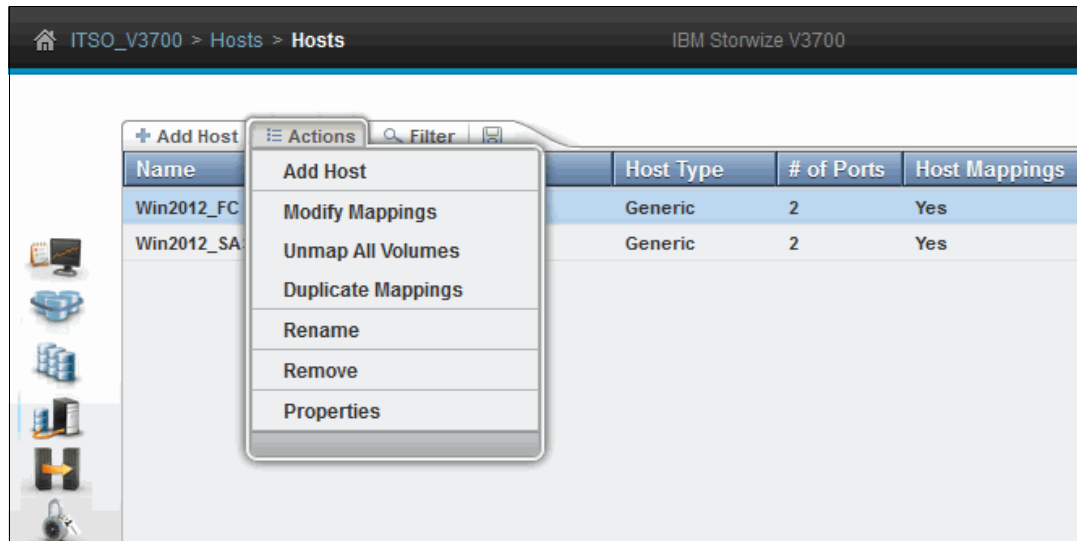


Figure 3-78 Hosts pane

Host Actions

Host Actions, such as Add Host, Modify Mappings, Unmap All Volumes, Duplicate Mappings, Rename, Delete, and Properties can be performed from the Hosts pane. Figure 3-78 shows the actions that are available from the Hosts pane.

For more information about the Hosts Actions menu, see 8.1, "Advanced host administration" on page 360.

Adding a host

To add a host, complete the following steps:

1. Click **Add Host** to open the Add Host window.
2. Choose the type from Fibre Channel Host, Internet Small Computer System Interface (iSCSI) Host, or SAS Host, and the applicable host configuration pane is shown.
3. After the host type is determined, the host name and port definitions can be configured.

Figure 3-79 shows the Choose the Host Type pane.

For more information about how to create hosts, see Chapter 4, “Host configuration” on page 149.

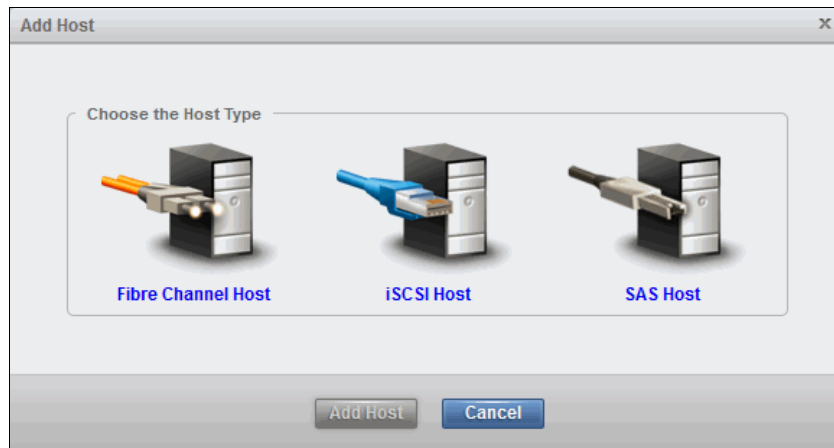


Figure 3-79 Choose the Host Type pane

Ports by Host pane

To perform actions against host ports, complete the following steps:

1. Select **Ports by Host** in the Hosts menu to open the pane, as shown in Figure 3-80. The pane shows the Name, Type, Status, and the # Nodes Logged In for ports that are assigned to the host definition.
2. You can Add and Delete host ports by clicking the respective icon.
3. Actions such as Create Host, Modify mappings, Unmap, Duplicate Mappings, Rename, and Properties can be performed from this pane by clicking the Actions menu.

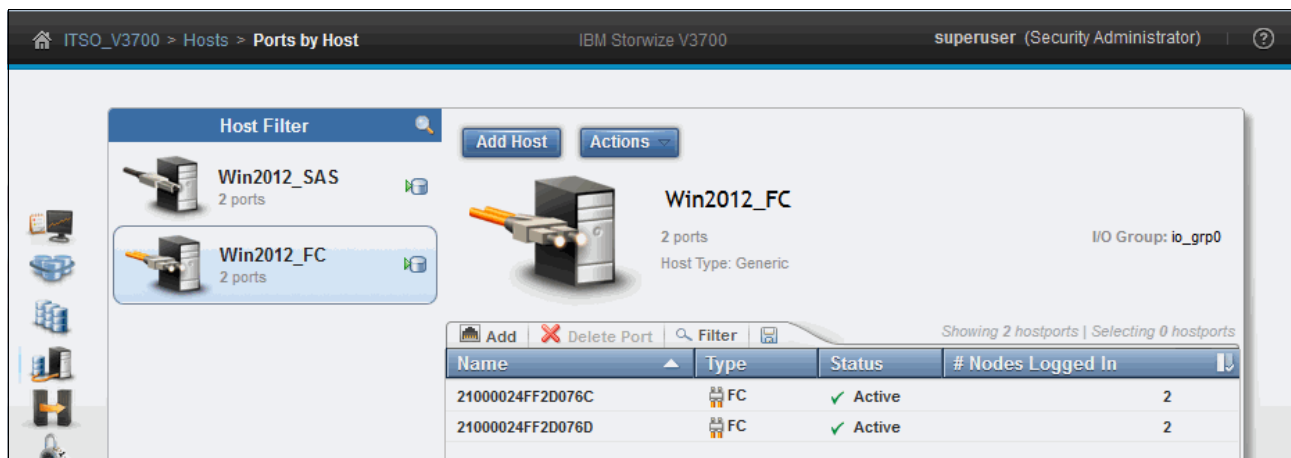


Figure 3-80 Ports by Host pane

Host Mappings pane

To map or unmap volumes, complete the following steps:

1. Select **Host Mappings** in the Hosts menu to open the pane, as shown in Figure 3-81. This pane shows the volumes that each host can access with the corresponding SCSI ID.
2. The Unmap Volume action can also be performed from this pane.



Figure 3-81 Host Mappings pane

Volumes by Host pane

For more information, see “Volumes by Host pane” on page 118.

3.4.5 Copy Services menu

The Copy Services menu provides access to the FlashCopy, Consistency Groups, FlashCopy Mappings, Remote Copy, and Partnership functions. Figure 3-82 shows the Copy Services menu.

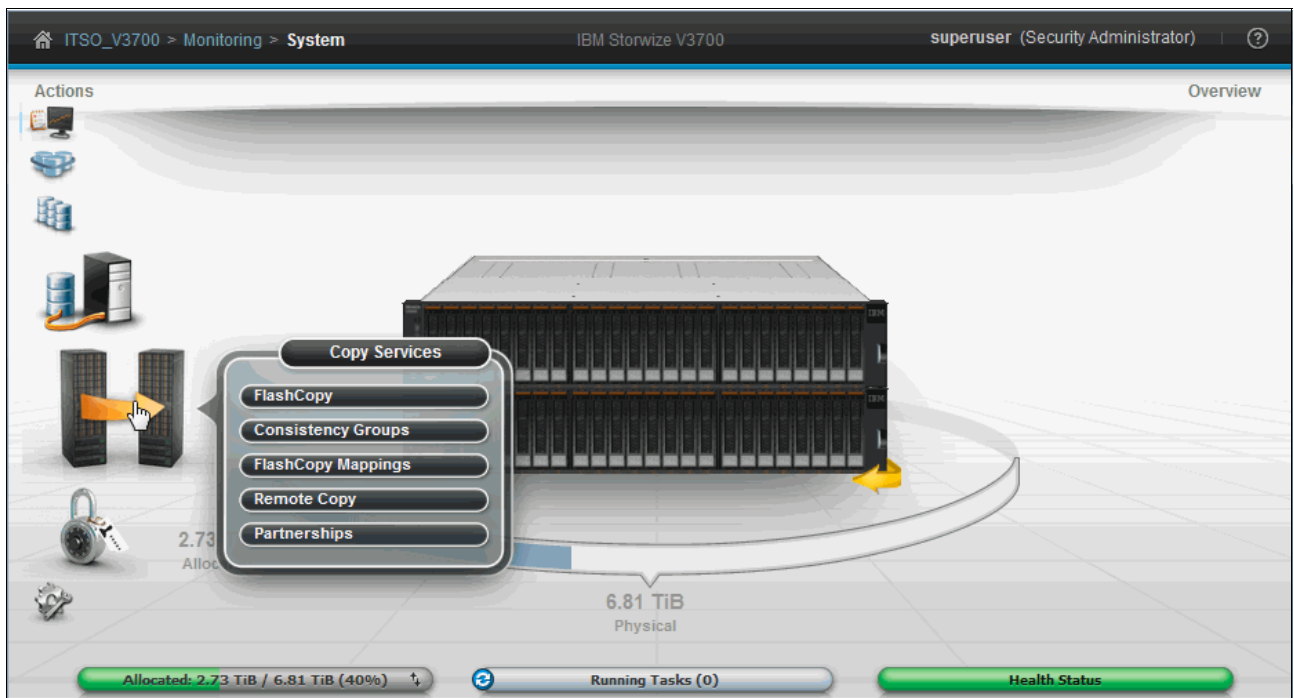


Figure 3-82 Copy Services menu

FlashCopy pane

Select **FlashCopy** in the Copy Services menu to open the pane, as shown in Figure 3-83. The FlashCopy pane displays all of the volumes in the system.

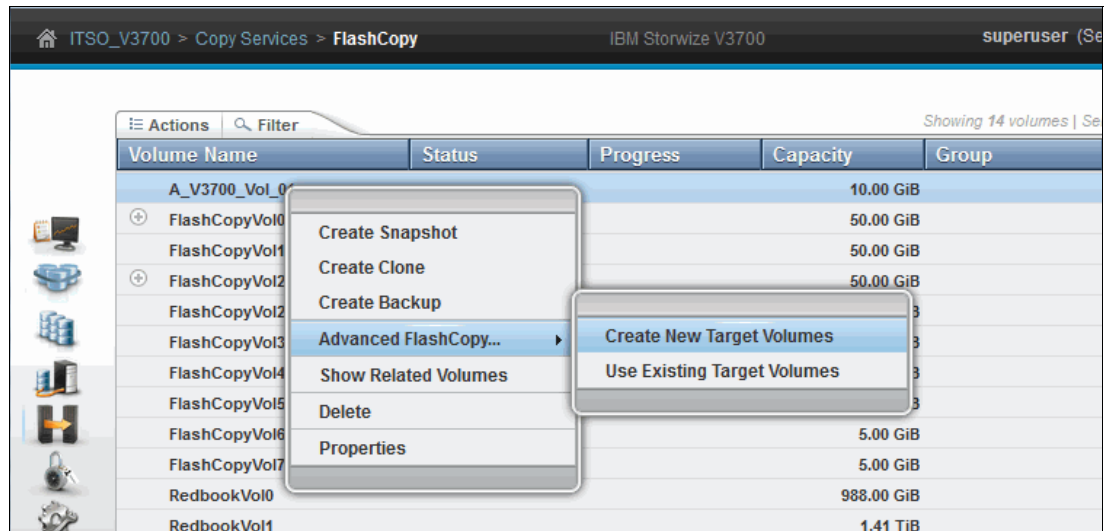


Figure 3-83 FlashCopy pane

FlashCopy actions

FlashCopy actions, such as Create Snapshot, Create Clone, Create Backup, Advanced FlashCopy, and Delete, can be performed from this pane. Figure 3-83 shows all of the actions that are available from the FlashCopy pane.

Consistency Groups pane

A *consistency group* is a container for FlashCopy mappings. Grouping enables FlashCopy mapping actions, such as prepare, start, and stop, to occur at the same time for the group rather than coordinating actions individually. This feature can help to ensure that the group's target volumes are consistent to the same point, and remove several FlashCopy mapping administration tasks.

To modify consistency groups, complete the following steps:

1. Select **Consistency Groups** in the Copy Services menu to open the pane.
The Consistency Group pane shows the defined groups with their associated FlashCopy mappings.
2. Group Actions, such as FlashCopy mapping Start, Stop, and Delete, can be performed from this pane.
3. Create FlashCopy Mapping and Create Consistency Group can also be selected from this pane.

For more information, see “FlashCopy Mappings pane” in the following section.

Figure 3-84 shows the Consistency Groups pane.

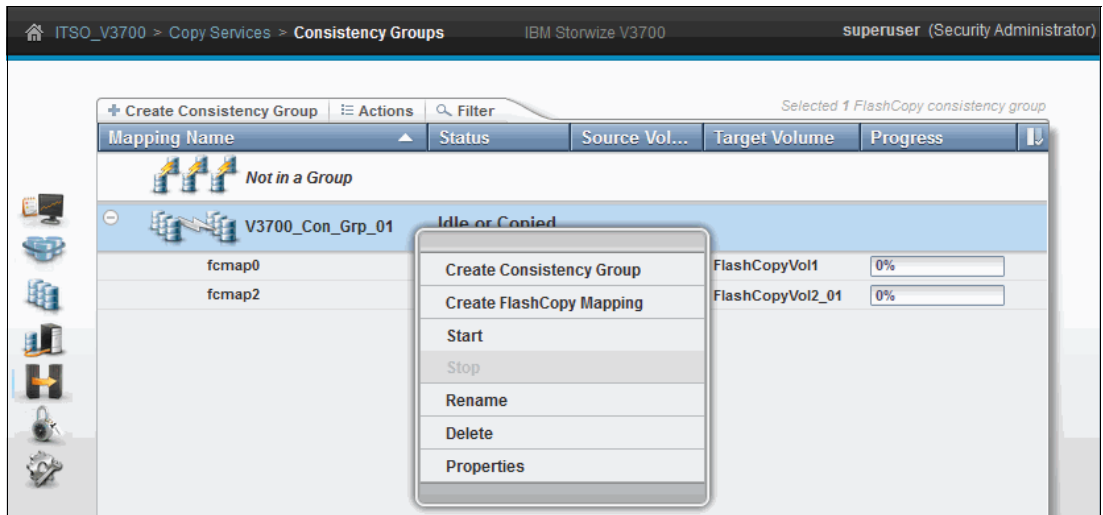


Figure 3-84 Consistency Groups pane

FlashCopy Mappings pane

FlashCopy mappings define the relationship between source volumes and target volumes. The FlashCopy Mappings pane shows information that relates to each mapping, such as status, progress, source and target volumes, and flash time.

To perform FlashCopy mappings, complete the following steps:

1. Select **FlashCopy Mappings** in the Copy Services menu to open the pane.
2. Select **Create FlashCopy Mapping** to configure a new mapping or use the Actions menu to administer the mapping. Figure 3-85 shows the FlashCopy Mappings pane.

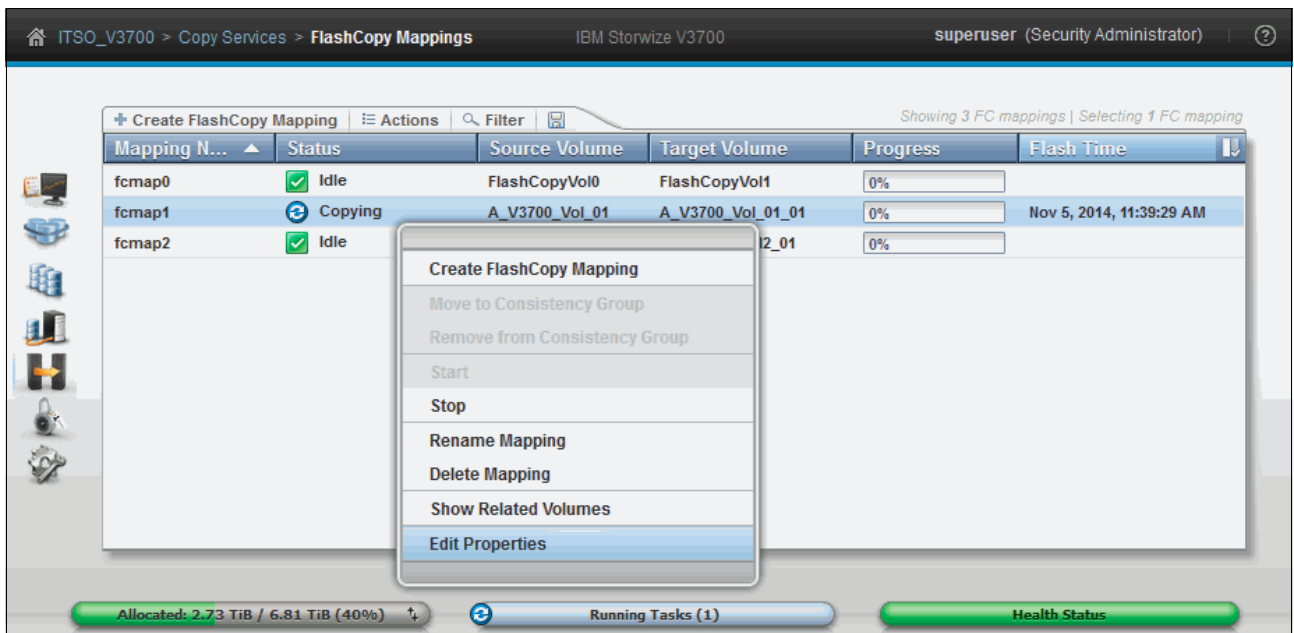


Figure 3-85 FlashCopy Mappings pane

For more information about how to create and administer FlashCopy mappings, see Chapter 10, “Copy services” on page 455.

Remote Copy pane

This pane shows both individual and grouped Remote Copy relationships. To maintain these relationships and groups, complete the following steps:

1. Click **Remote Copy** in the Copy Services menu to open the pane shown in Figure 3-86.
2. The Actions menu provides options to start, stop, switch, rename, and delete relationships.
3. You can also use it to create and modify relationships and consistency groups.

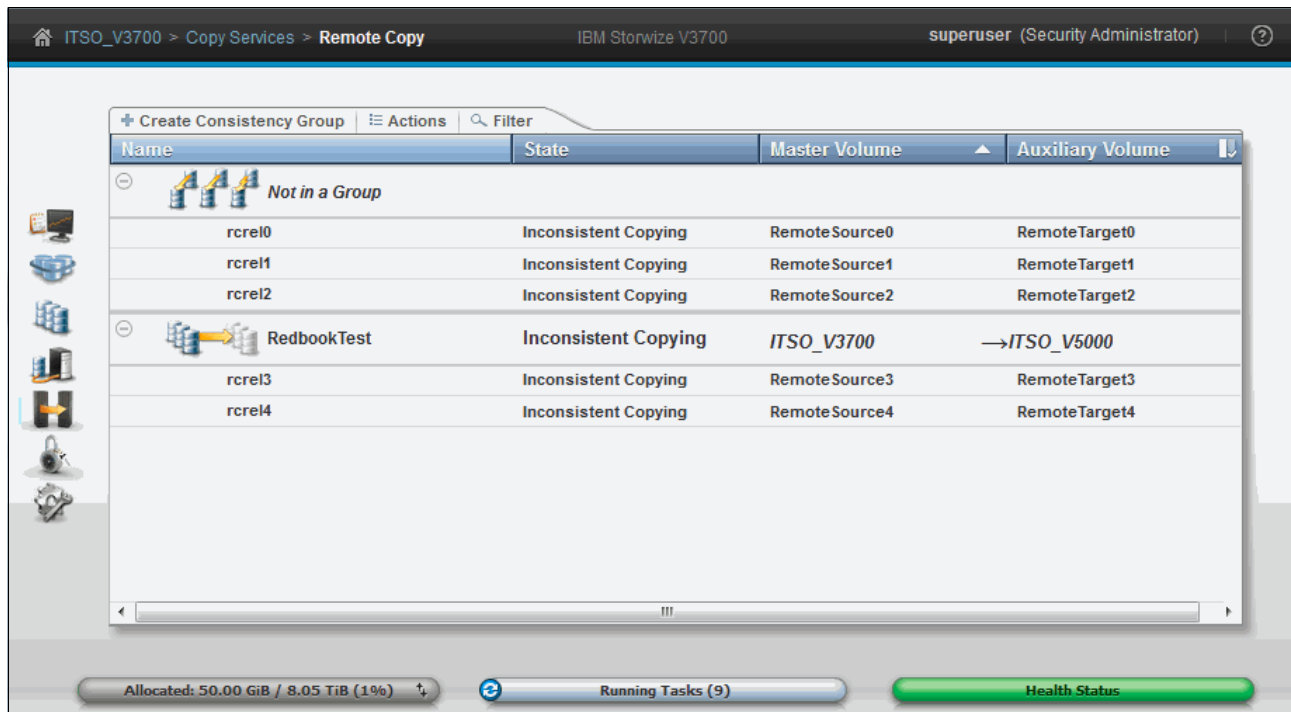


Figure 3-86 Remote Copy pane

Partnerships pane

To manage partnerships, complete the following steps:

1. Click **Partnerships** to open the pane shown in Figure 3-87.
2. From this pane, it is possible to set up a new partnership, or delete an existing partnership, with another IBM Storwize or IBM SAN Volume Controller system for the purposes of Remote Copy.

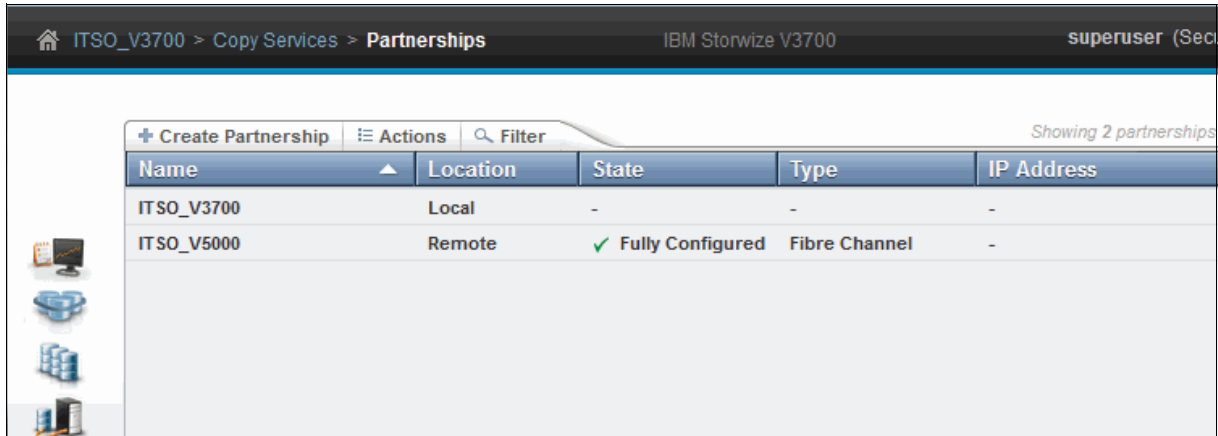


Figure 3-87 Partnerships pane

3. From the Actions menu, access the Partnership Properties pane shown in Figure 3-88.
4. From here, it is possible to change the Link bandwidth and background copy rate.

The Link Bandwidth in Megabits per second (Mbps) specifies the bandwidth that is available for replication between the local and partner systems. The Background copy rate specifies the maximum percentage of the link bandwidth that can be used for background copy operations. This value should be set so that there is enough bandwidth available to satisfy host write operations in addition to the background copy.

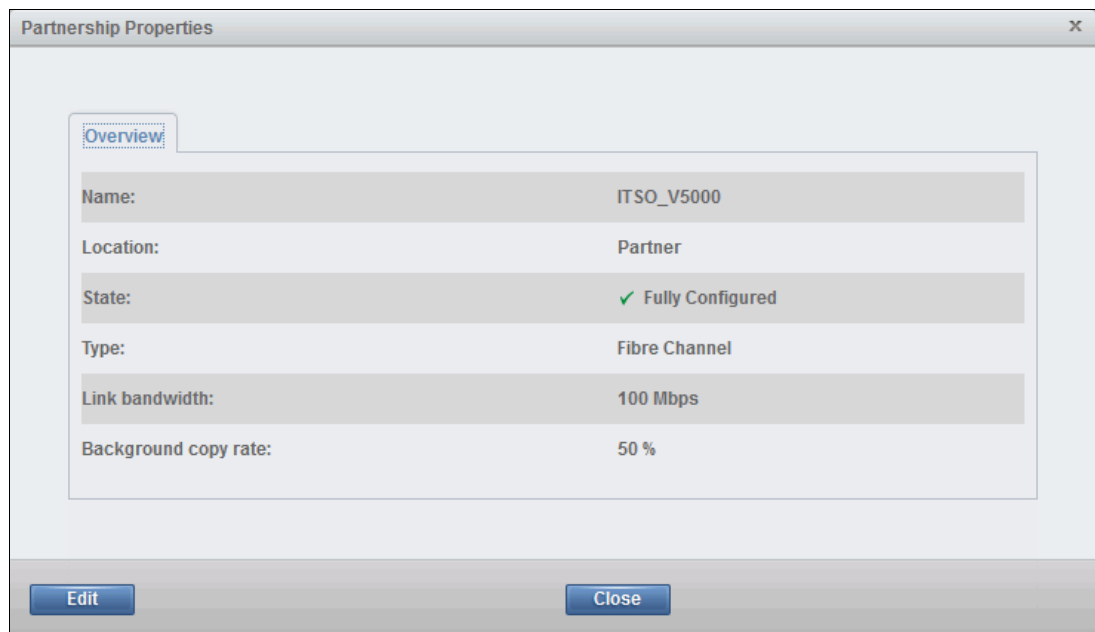


Figure 3-88 Partnership Properties

3.4.6 Access menu

The Access menu provides access to the Users and Audit Log functions, as shown in Figure 3-89.




Figure 3-89 Access menu

Users pane

Select **Users** in the Access menu to open the pane. The Users pane shows the defined user groups and users for the system. The users that are listed can be filtered by user group. Figure 3-90 shows the Users pane and the Users Actions menu.

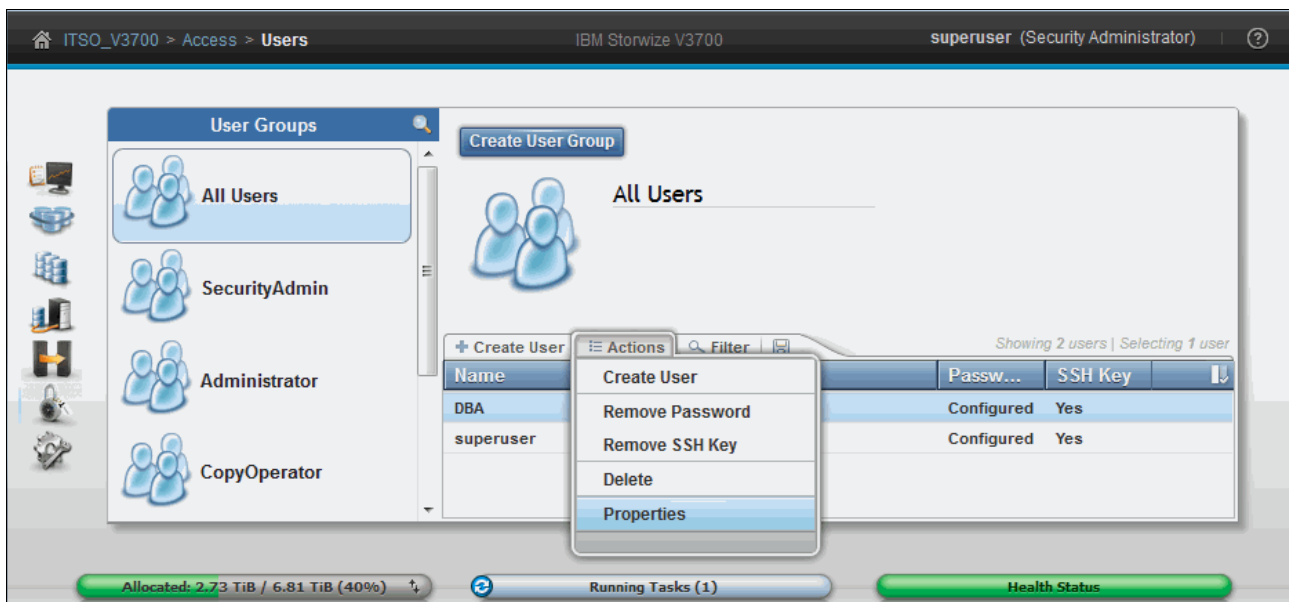


Figure 3-90 Users pane

Creating a user group

To create a user group, complete the following steps:

1. Click **Create User Group** to open the Create User Group pane.
2. Enter the group name, select the role, then click **Create**, as shown in Figure 3-91.
3. Hover over each role to show a help icon that describes the role.

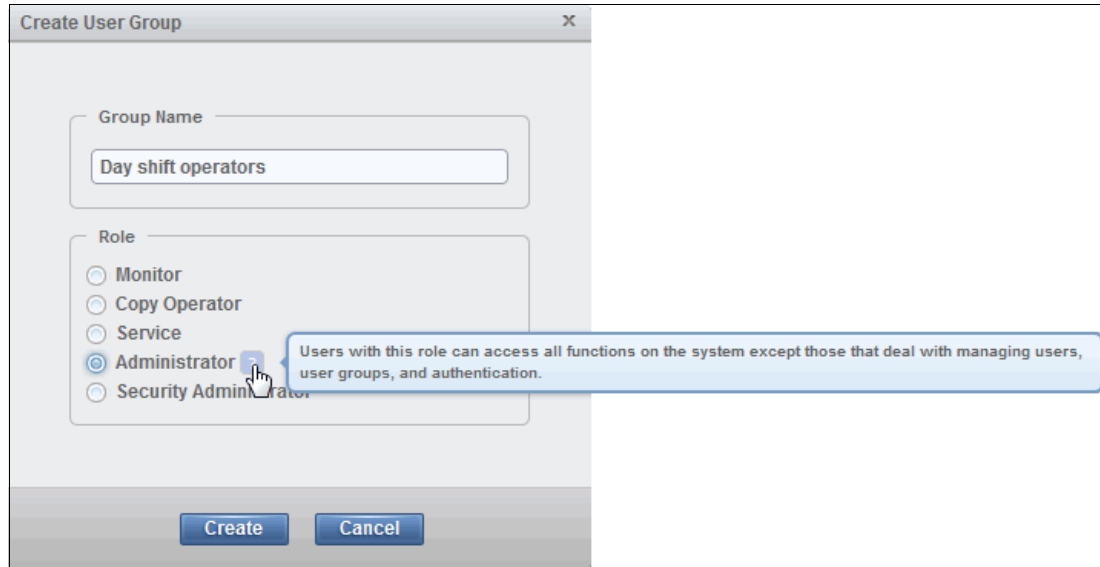


Figure 3-91 Create User Group pane

Creating a user

To create a new user, complete the following steps:

1. Click **New User** to define a user to the system. Figure 3-92 shows the Users pane and the New User option.



Figure 3-92 Create a new user

2. Using the Create User pane, as shown in Figure 3-93 on page 128, it is possible to configure the user name, password, and authentication mode.

The user name, password, group, and authentication mode are mandatory fields but the public Secure Shell (SSH) key is optional.

3. After the user is defined, click **Create**.

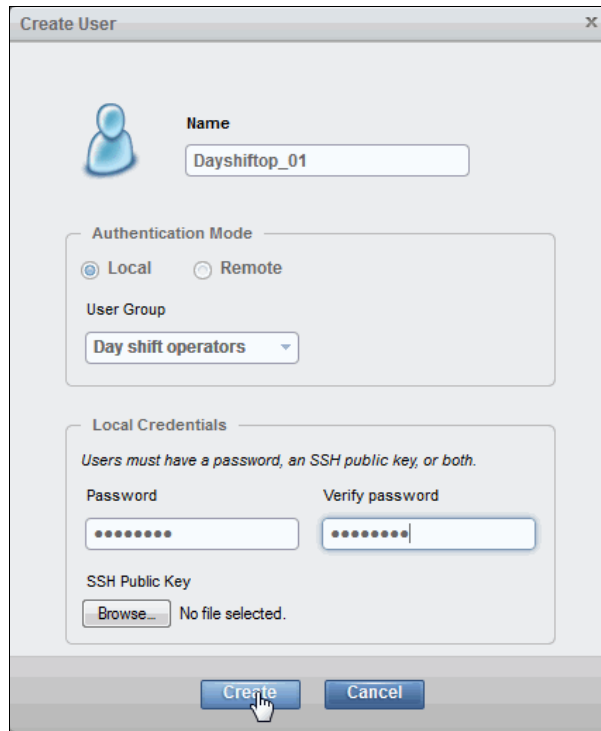


Figure 3-93 Create User pane

4. The authentication mode can be set to local or remote:
 - Select local if the IBM Storwize V3700 authenticates locally.
 - Select remote if a remote service, such as an LDAP server, authenticates the connection. If remote is selected, the remote authentication server must be configured in the IBM Storwize V3700 by using the **Settings menu** → **Security** pane.
5. (Optional) Use an SSH configuration to establish a more secure connection to the CLI. For more information, see Appendix A, “Command-line interface setup and SAN Boot” on page 649.

Audit Log pane

The audit log tracks action commands that are issued through a CLI session or through the management GUI. The Audit Log pane, shown in Figure 3-94, displays information about the command, such as the time stamp, user, and any associated command parameters. It is not possible to delete or alter the Audit log. To open it, complete the following steps:

1. Select **Audit Log** in the Access menu to open the pane.
2. To reduce the number of listed items, select **Filter by Date** or **Show entries within**.

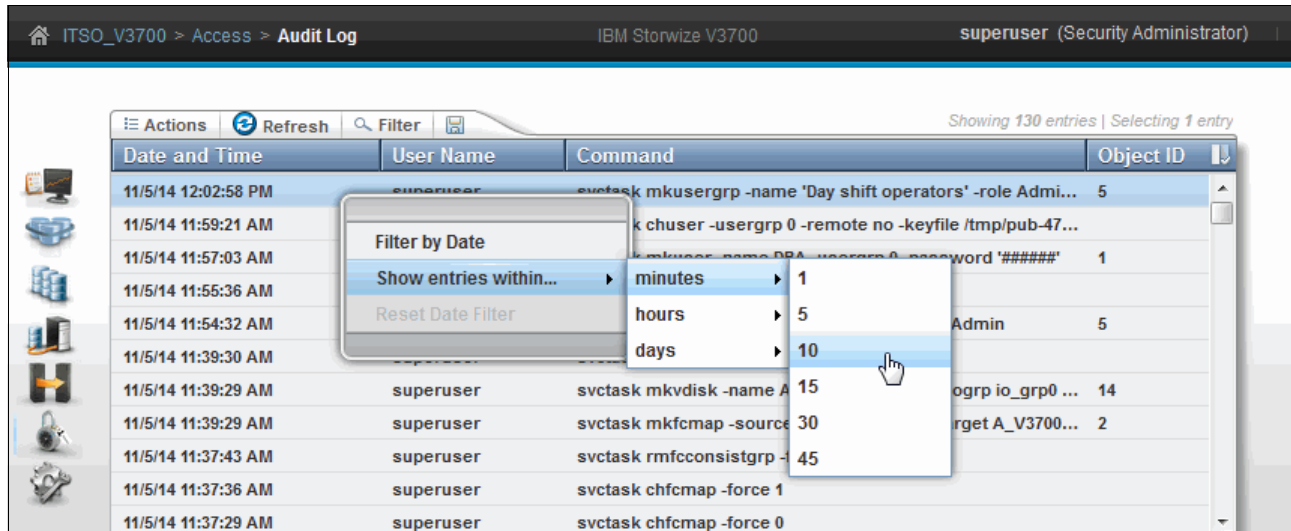


Figure 3-94 Audit Log pane

3.4.7 Settings menu

The Setting menu provides access to the Notifications, Network, Security, System, Support, and GUI Preferences panes. Figure 3-95 shows the Settings menu.

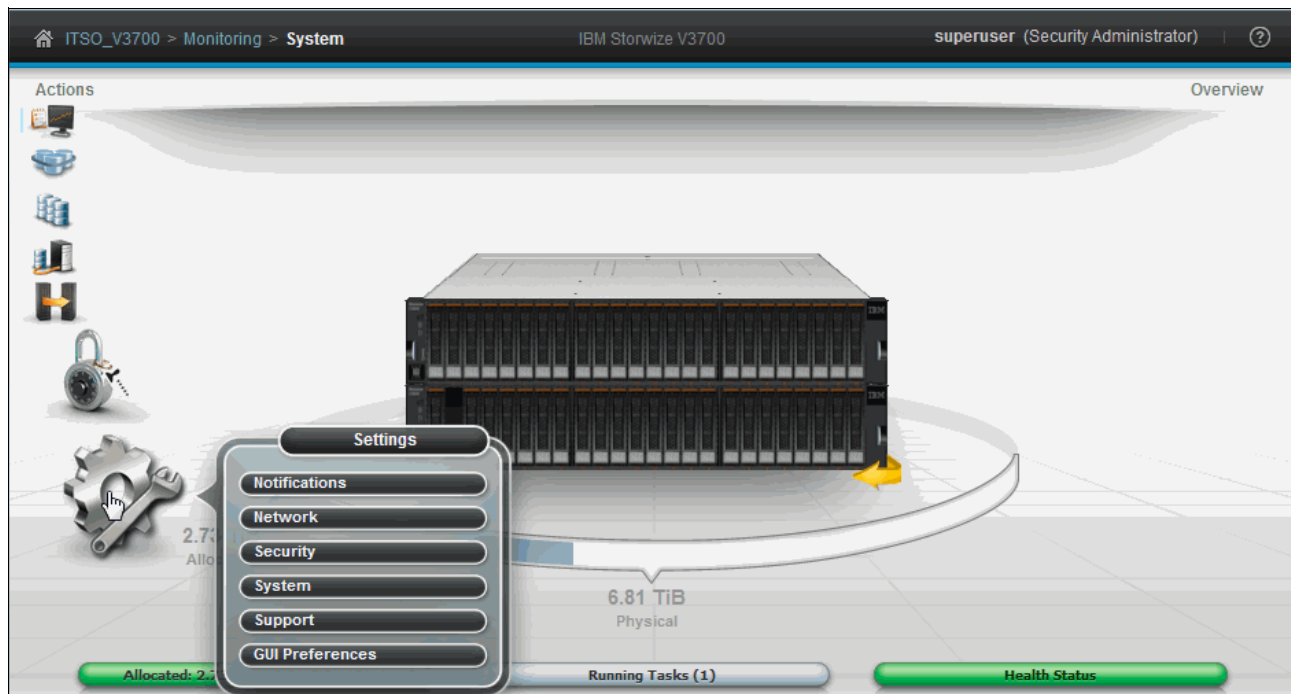


Figure 3-95 Settings menu

Notifications pane

The IBM Storwize V3700 can use Simple Network Management Protocol (SNMP) traps, syslog messages, emails, and IBM Call Home to notify users and IBM when events are detected. Each event notification method can be configured to report all events or alerts. Alerts are the significant events, and might require user intervention. The event notification levels are *critical*, *warning*, and *information*.

Select **Notifications** in the Settings menu to open the Notifications pane, which provides access to the Email, SNMP, and Syslog configuration panes. IBM Call Home is an email notification for IBM Support. It is automatically configured as an email recipient and is enabled when the Email event notification option is enabled by following the Call Home wizard.

Enabling Email Notification option

As shown in Figure 3-96, the Notifications pane provides access to the Email configuration pane:

1. Click **Email** to open the pane.
2. Click **Enable Notifications** to open the Email Event Notifications wizard.

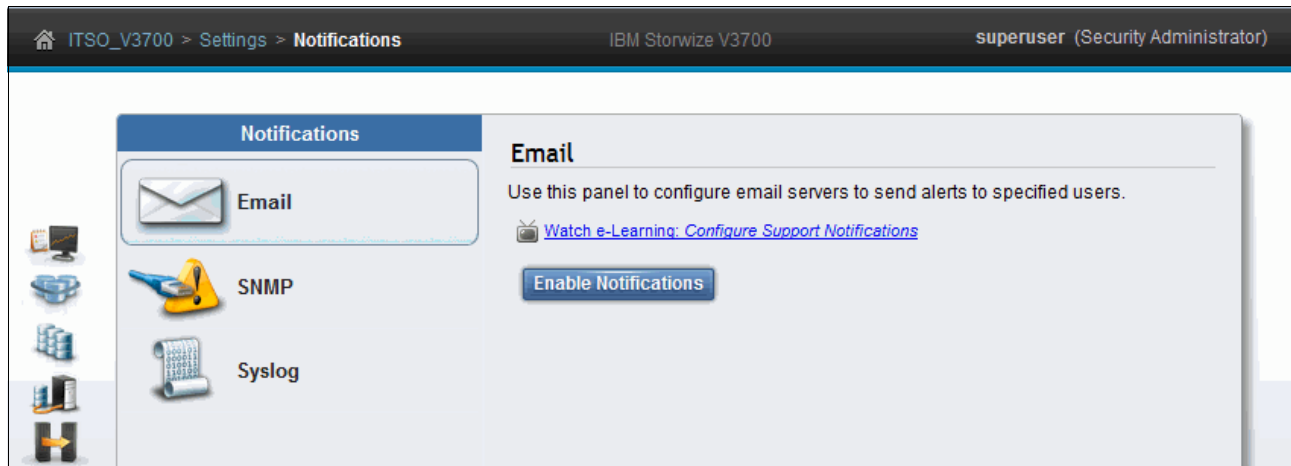


Figure 3-96 Notifications pane: Email

Email Event Notifications wizard

The Email Event Notifications wizard, as shown in Figure 3-97, guides the user through system location, account contact, and email configuration tasks.

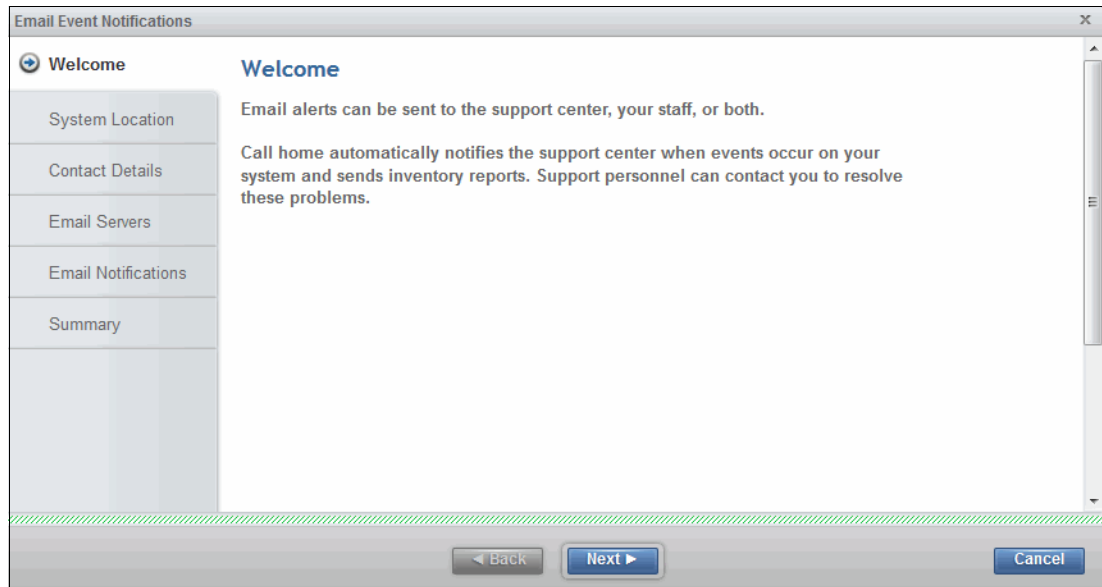


Figure 3-97 Email Event Notifications wizard

SNMP event notification

As shown in Figure 3-98, the Notifications pane provides access to the SNMP pane. To set up SNMP notification, complete the following steps:

1. Click **SNMP** to open the pane and view the configured SNMP servers.

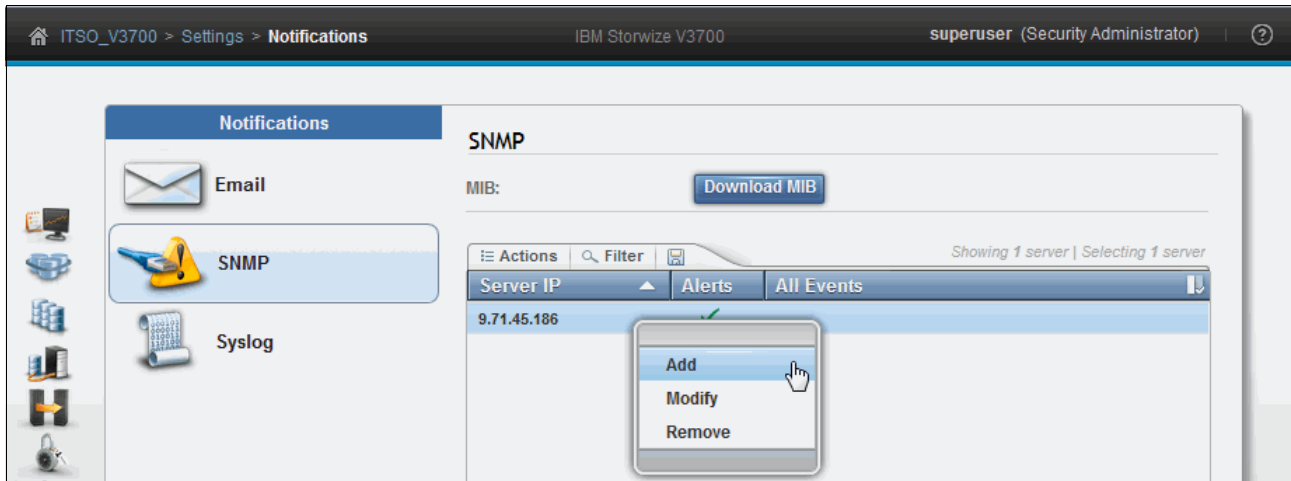
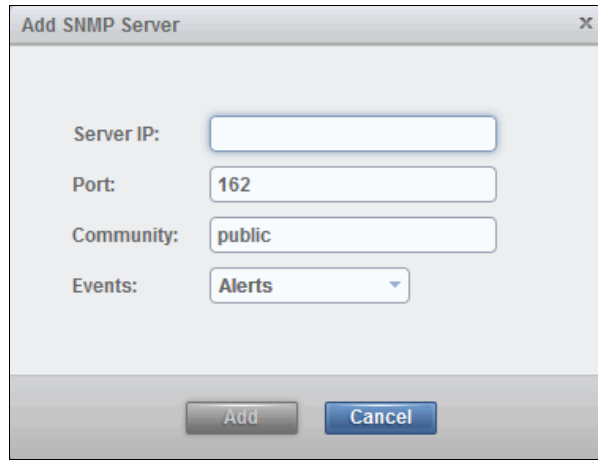


Figure 3-98 SNMP pane

2. A Management Information Base (MIB) file for SNMP can be used to configure a network management program to receive SNMP messages that are sent by the system. This file can be used with SNMP messages from all versions of the software. Click **Download MIB** to download the MIB file.

3. Configure up to six servers using the **Add** option from the Actions menu. Figure 3-99 shows the Add SNMP Server pane.



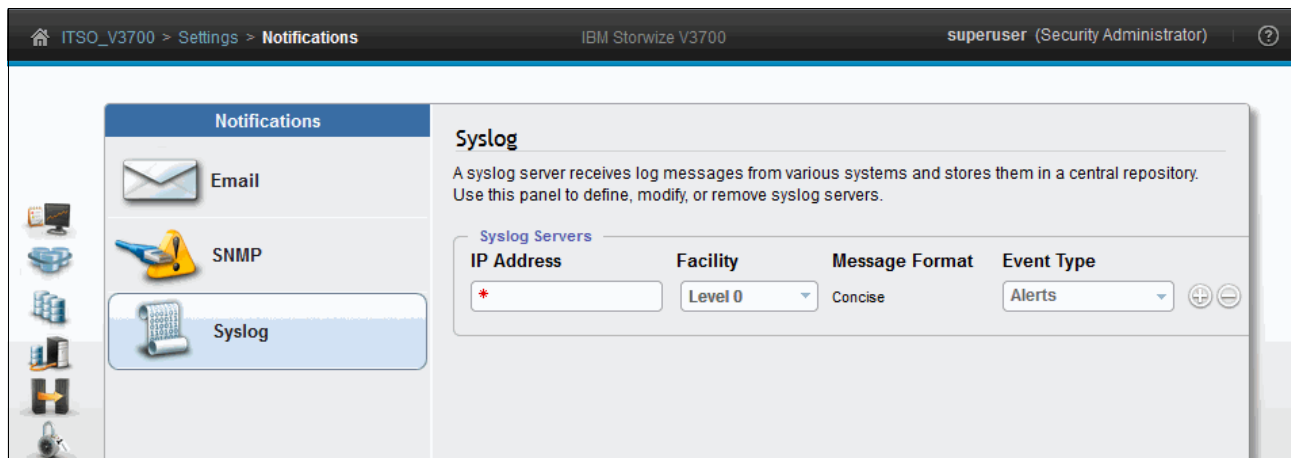
The image shows a dialog box titled "Add SNMP Server" with a close button (X) in the top right corner. It contains four input fields: "Server IP:" (empty), "Port:" (162), "Community:" (public), and "Events:" (Alerts). At the bottom, there are two buttons: "Add" and "Cancel".

Figure 3-99 Add SNMP Server pane

Syslog event notification

As shown in Figure 3-100, the Notifications pane provides access to the Syslog configuration pane:

1. Click **Syslog** to open the pane and enter the server details.
2. Configure multiple servers by clicking plus (+) to add more servers.



The image shows a screenshot of the IBM Storwize V3700 management console. The breadcrumb navigation is "ITSO_V3700 > Settings > Notifications". The user is logged in as "superuser (Security Administrator)". The "Notifications" pane is active, showing options for "Email", "SNMP", and "Syslog". The "Syslog" configuration pane is open, displaying a table of "Syslog Servers".

IP Address	Facility	Message Format	Event Type	
*	Level 0	Concise	Alerts	+ -

Figure 3-100 Syslog configuration pane

Network pane

Select **Network** in the Settings menu to open the pane. As shown in Figure 3-101, the Network pane provides access to the Management IP Addresses, Service IP Addresses, Ethernet Ports, iSCSI, Fibre Channel Connectivity, and Fibre Channel Ports configuration panes.

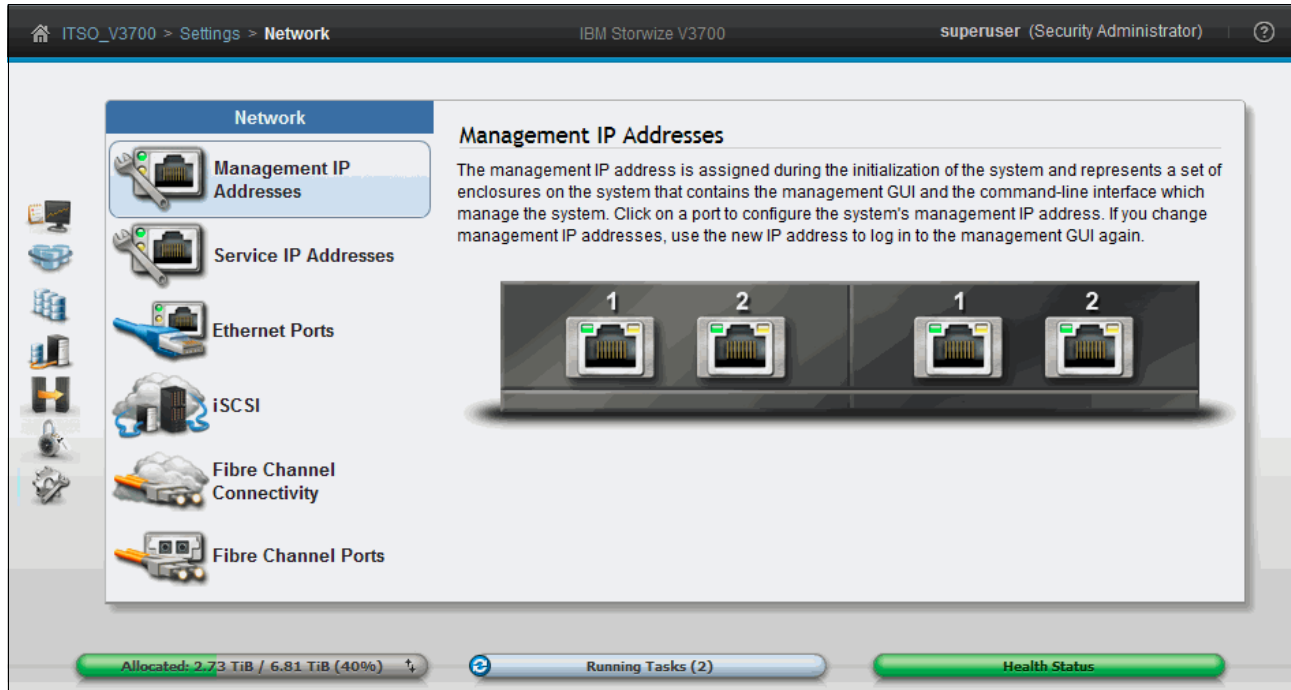


Figure 3-101 Network pane

Management IP addresses

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

Another Management IP address can be logically assigned to Ethernet port 2 of each node canister for fault tolerance:

1. If the Management IP address is changed, use the new IP address to log in to the Management GUI again.
2. Click **Management IP Addresses**.
3. Click the port to configure (the corresponding port on the partner node canister is also highlighted).
4. Provide the required IP, Subnet, and Gateway information.

Figure 3-102 shows the Management IP Addresses configuration pane.

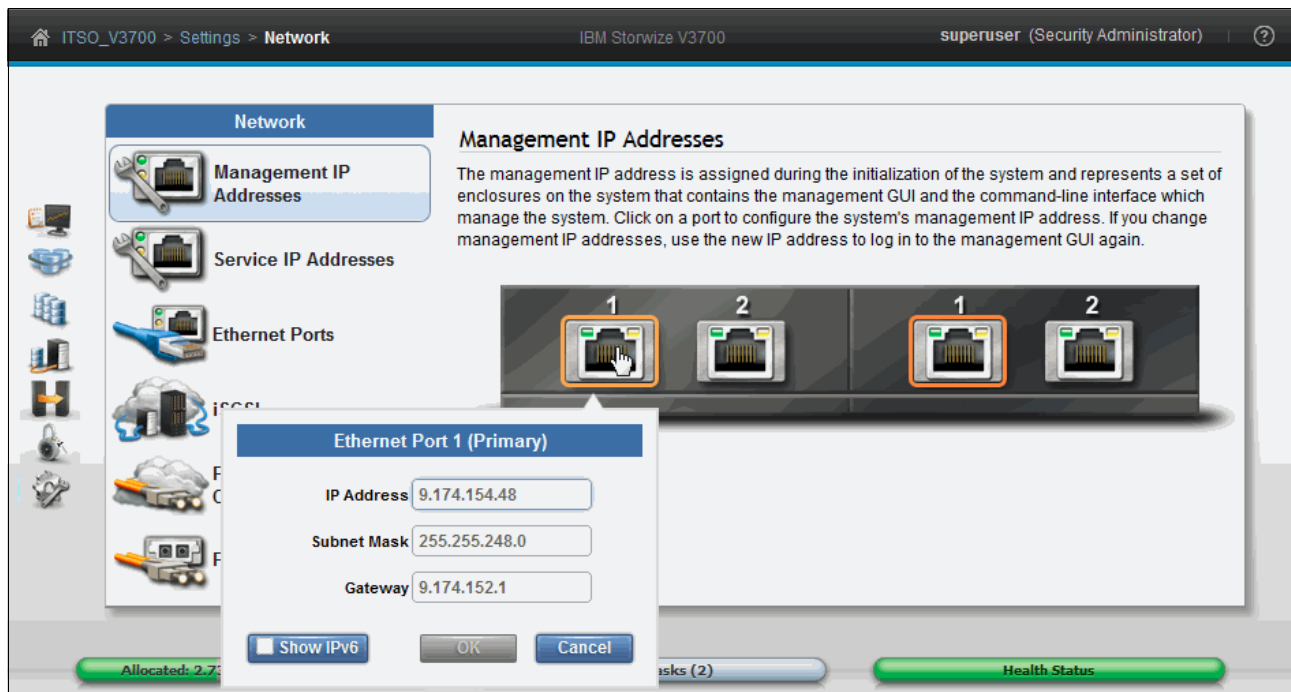


Figure 3-102 Management IP Addresses configuration pane

Service IP addresses

Service IP addresses are used to access the Service Assistant Tool. Each node canister should have its own Service IP address. The address can be an IPv4 address, IPv6 address, or both. For more information, see 2.9.3, “Service Assistant tool” on page 67.

The Service IP addresses are configured on Ethernet port 1 of each node canister:

1. Click **Service IP Addresses**.
2. Select the **Control Enclosure** and **Node Canister** to configure.

Figure 3-103 shows the Service IP addresses configuration pane.



Figure 3-103 Service IP Addresses configuration pane

Ethernet Ports

The Ethernet ports can be used for iSCSI connections, host attachment, and remote copy:

1. Click **Ethernet Ports** to view the available ports.
2. Figure 3-104 shows the Ethernet Ports pane and associated Actions menu. Click **Actions** → **Modify IP Settings** to modify the port settings.

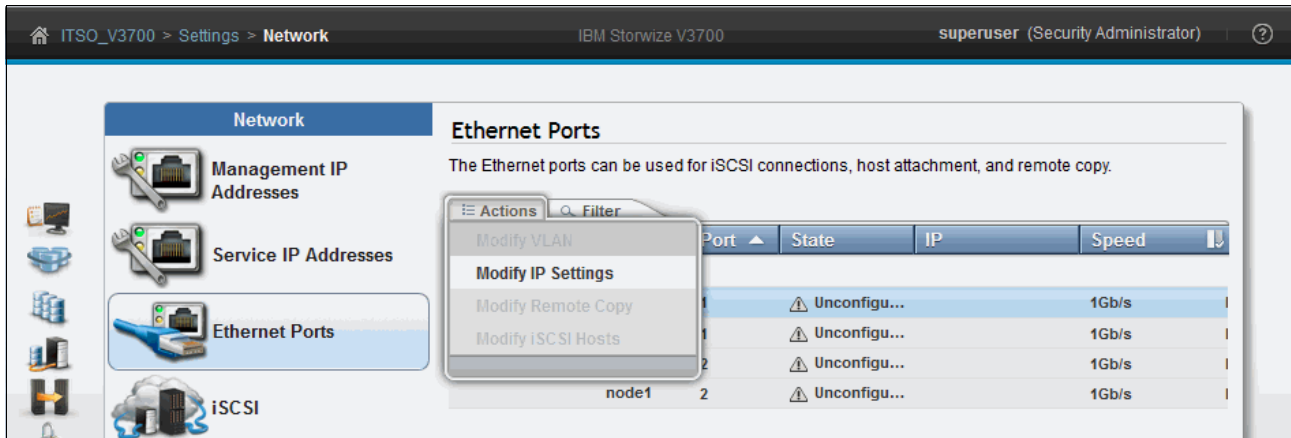


Figure 3-104 Ethernet Ports pane

iSCSI Configuration

The IBM Storwize V3700 supports iSCSI connections for hosts. Click **iSCSI** to configure the System Name, iSCSI Aliases, and iSNS settings. Figure 3-105 shows the iSCSI Configuration pane.

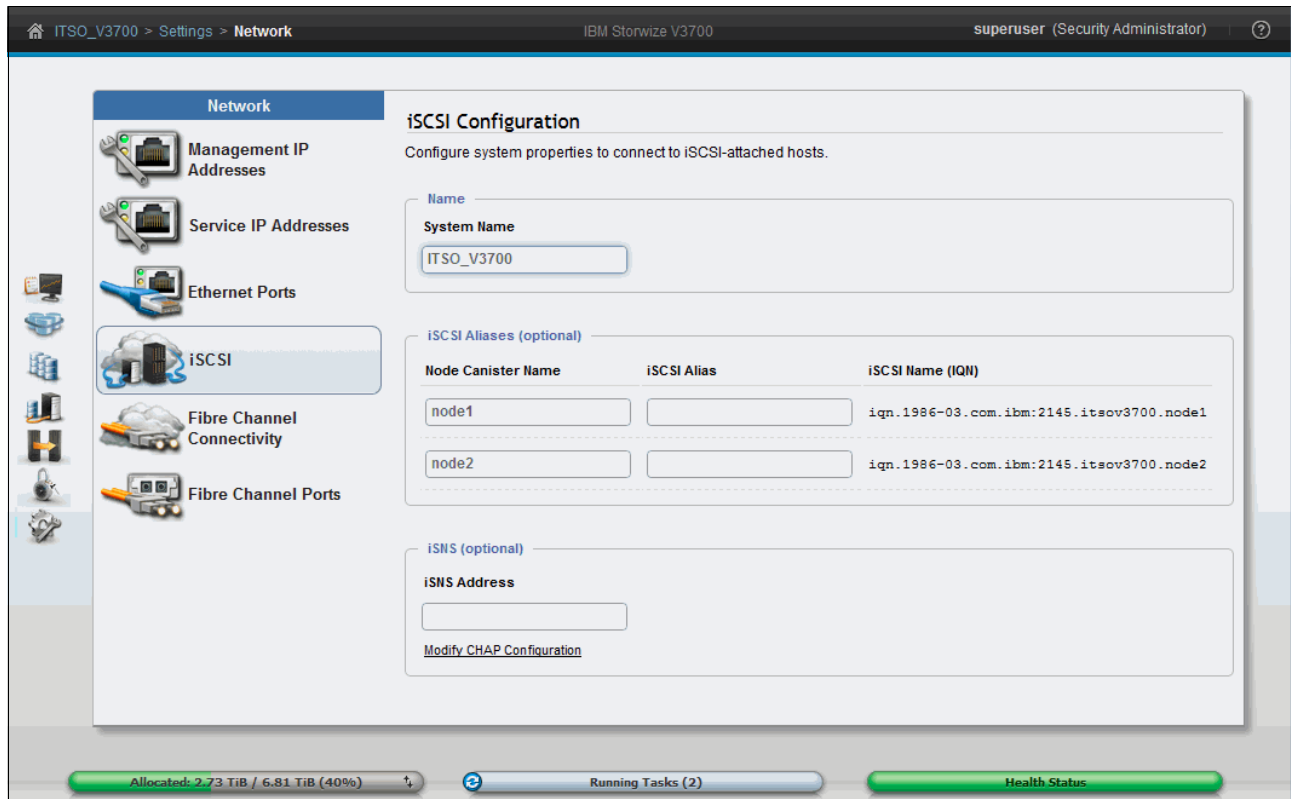


Figure 3-105 iSCSI Configuration pane

Fibre Channel Connectivity

The Fibre Channel Connectivity pane displays the connectivity between nodes and other storage systems and hosts that are attached through the FC network:

1. Click **Fibre Channel Connectivity**.
2. Select the required view from the **View connectivity for** menu.
3. Click the **Show Results** button.

Figure 3-106 shows the Fibre Channel pane with the **All nodes, storage systems, and hosts** option selected.

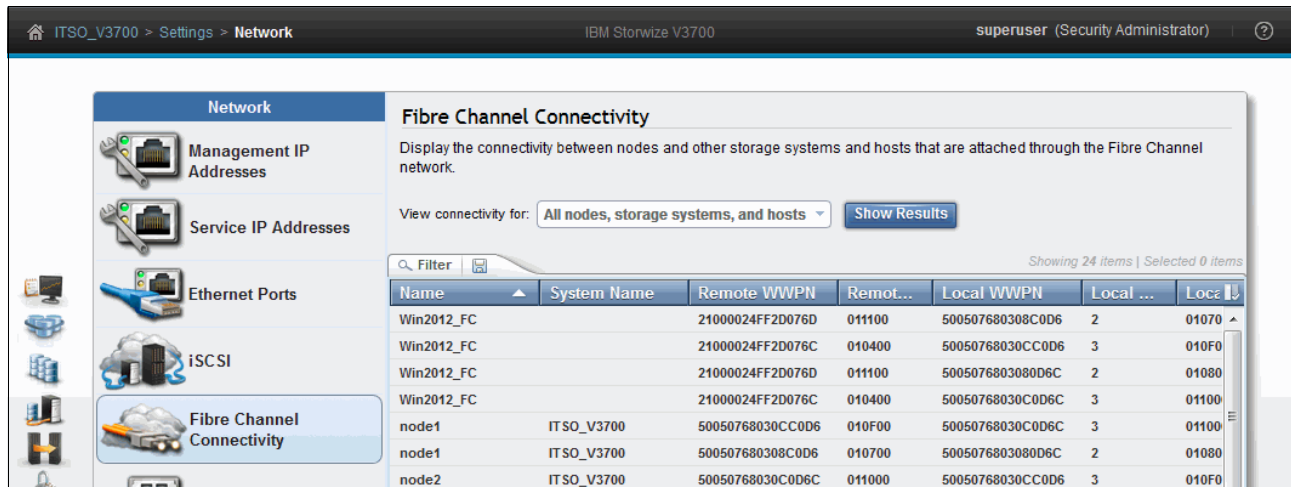


Figure 3-106 Fibre Channel Connectivity pane

Fibre Channel Ports

The Fibre Channel Ports pane displays the Fibre Channel ports used by the system. Each port is allowed to communicate with hosts and storage systems. Figure 3-107 shows the Fibre Channel Ports pane.

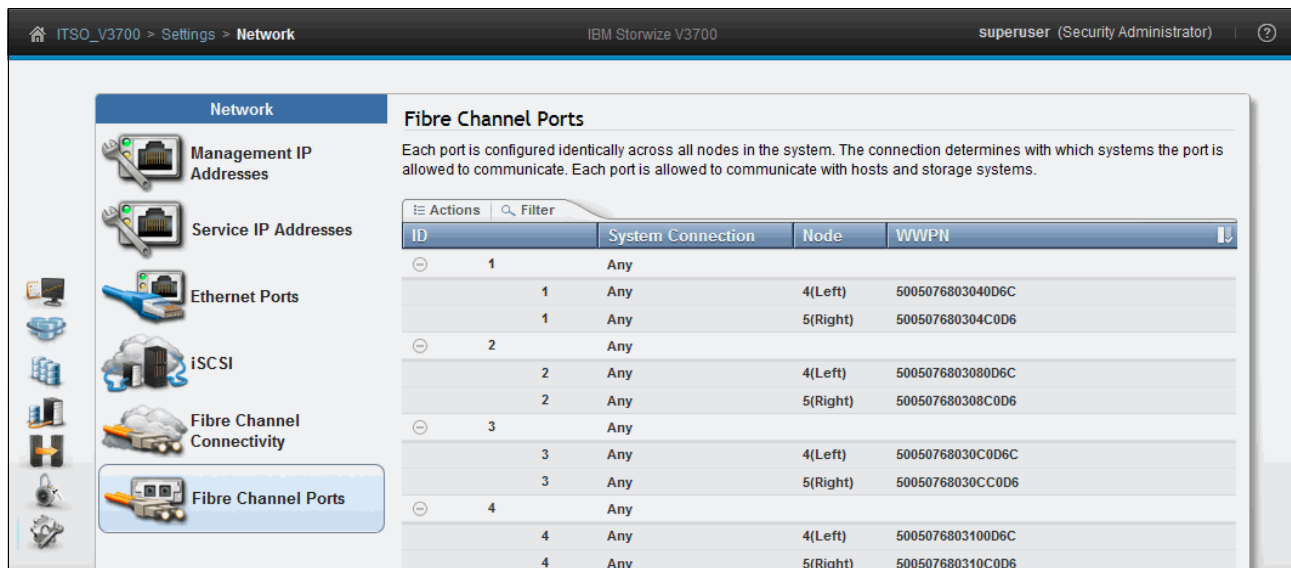


Figure 3-107 Fibre Channel Ports

Security pane

Select **Security** in the Settings menu to open the pane. The Security pane provides access to the Remote Authentication wizard. Remote authentication must be configured to create remote users on the IBM 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.

Enabling Remote Authentication

To enable remote authentication, complete the following steps:

1. Click **Configure Remote Authentication**, as shown in Figure 3-108, to open the wizard.
2. Follow the prompts to configure remote authentication.

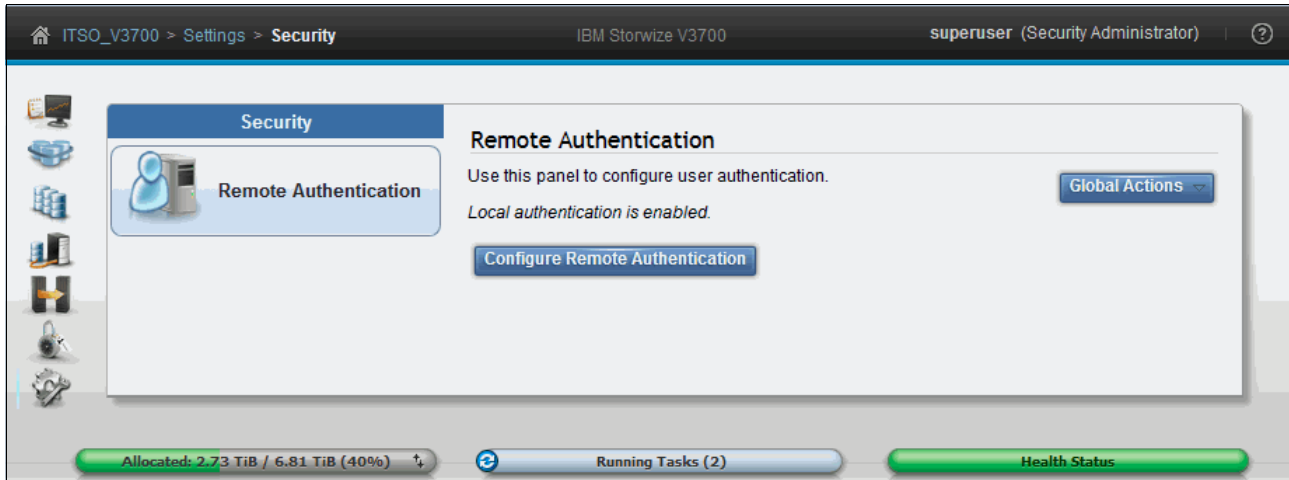


Figure 3-108 Security pane

System pane

Select **System** in the Settings menu to open the pane. This pane provides access to the Date and Time, Licensing, and Upgrade Software configuration panes.

Date and Time

Click **Data and Time** to configure the date and time manually or using a Network Time Protocol (NTP) server. Figure 3-109 shows the Date and Time configuration pane.

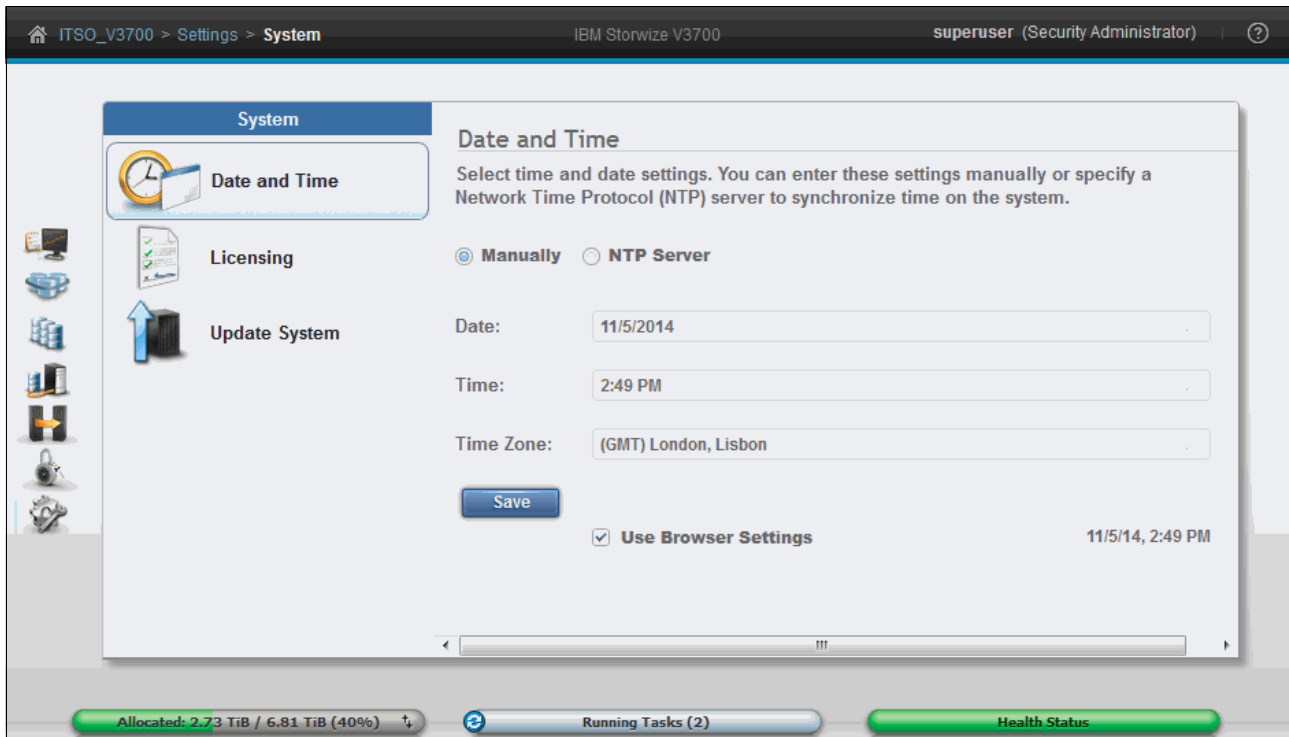


Figure 3-109 System pane

Licensing

The licensing pane shows the current system licensing. The IBM Storwize V3700 is based on per-system licensing, and is not licensed per enclosure as are the other systems in the Storwize family.

The following optional licenses are available:

- ▶ FlashCopy upgrade to 2040 target copies
- ▶ Remote Copy
- ▶ IBM Easy Tier
- ▶ Turbo Performance

One-time trial licenses can be enabled from the GUI for any function except the FlashCopy upgrade. The trial licenses automatically disable after 90 days if they are not fully licensed in the meantime.

IBM sends a printed page with an authorization code when the optional license is ordered.

The following options are available for enabling the licenses:

- ▶ Automatic

The client enters an authorization code into the GUI of an internet-attached IBM Storwize V3700. The system validates the code with IBM and enables functionality.

- ▶ Manual

The client enters an authorization code and machine details into the data storage feature activation (DSFA) website, which validates and creates a license key. The client enters the license key into the IBM Storwize V3700 GUI, which enables the function.

For more information, see the DSFA website:

<https://www.ibm.com/storage/dsfa/storwize/selectMachine.wss>

The authorization code can be used only once, and the license key is specific to one machine. Licenses cannot be transferred between systems and, if the system is sold, license ownership transfers with it.

Figure 3-110 shows the Licensed Functions pane within the Settings pane. The licenses can be activated by right-clicking the function and selecting the required activation method.

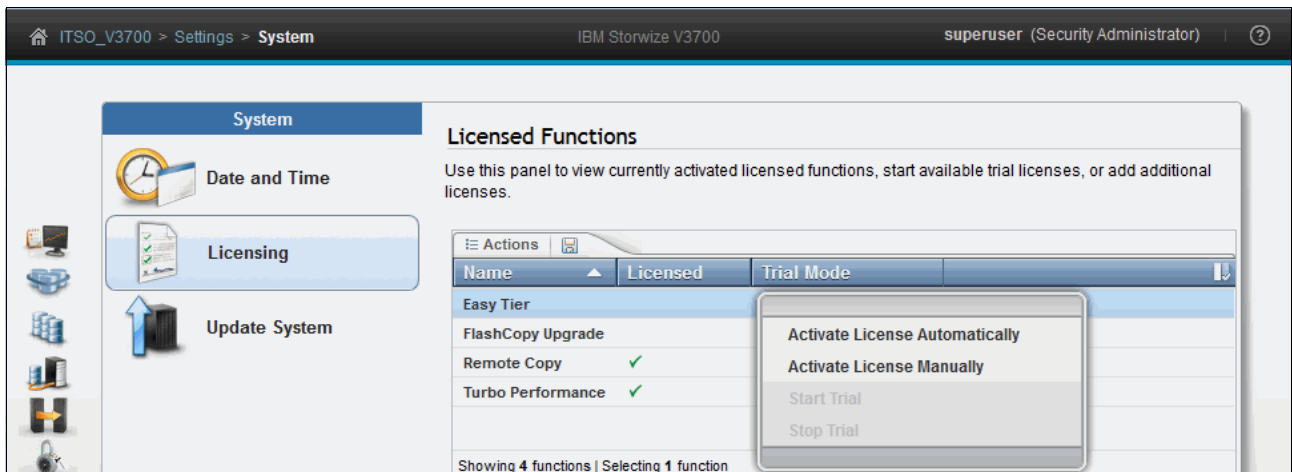


Figure 3-110 Update license

Update System

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

To update the software, the IBM Storwize V3700 software and the IBM Storwize V3700 Upgrade Test Utility must be downloaded. After the files are downloaded, complete the following steps:

1. It is best to validate the checksum to ensure that the files are sound.
2. Read the release notes.
3. Verify compatibility.
4. Follow all IBM recommendations and prerequisites.
5. Finally, to update the software of the IBM Storwize V3700, click **Update**, as shown in Figure 3-111.

For more information, see 11.4, “Updating software” on page 597.

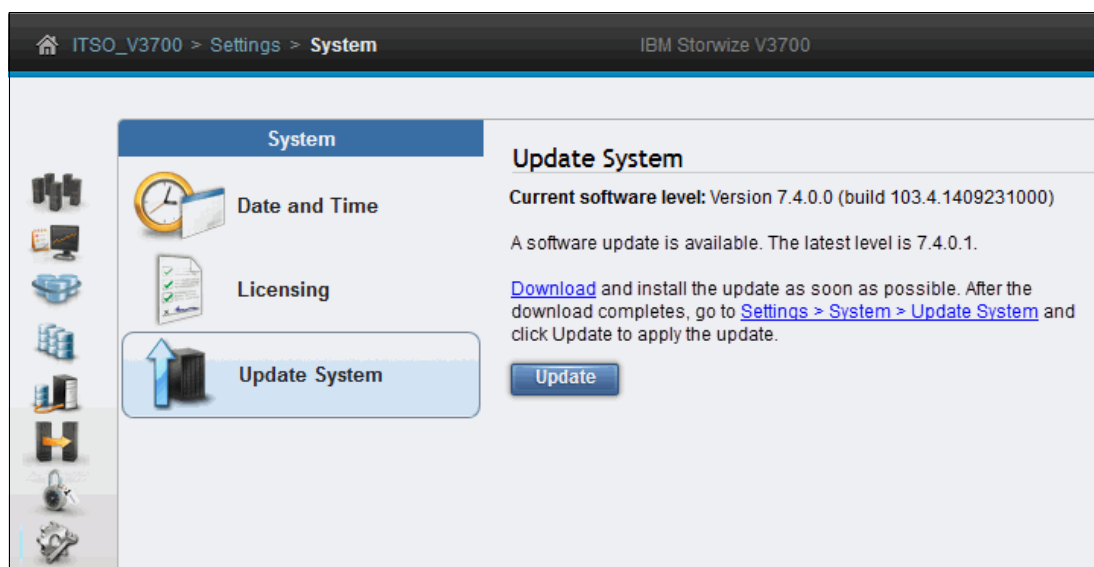


Figure 3-111 Update System pane

Support pane

As shown in Figure 3-112, this pane provides access to the IBM support package, which is used by IBM to assist with problem determination:

1. Select **Support** in the Settings menu to open the Support pane.
2. Click **Download Support Package** to access the wizard.

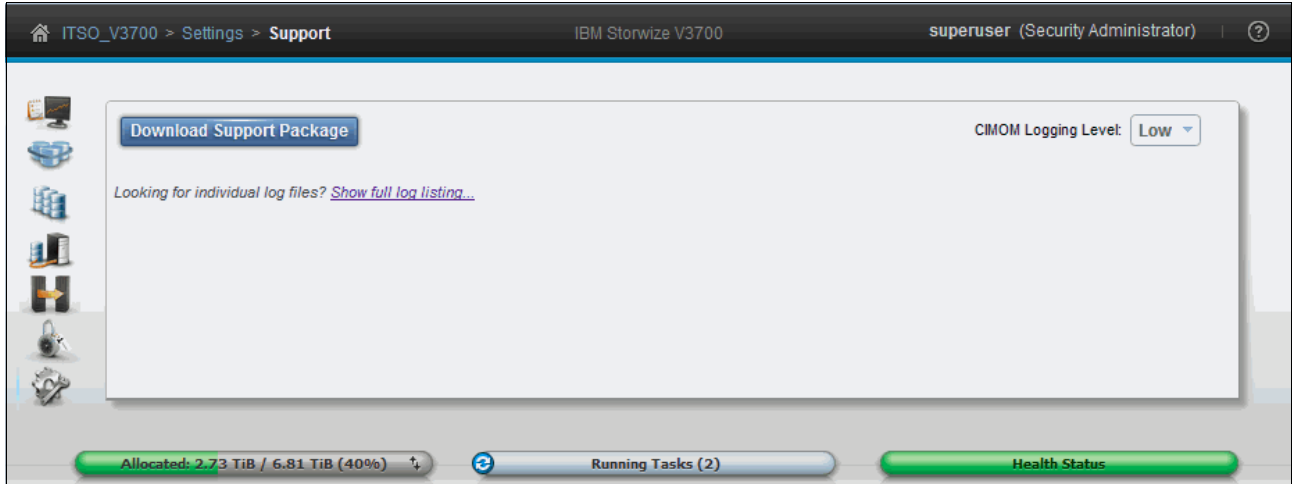


Figure 3-112 Support pane

Download Support Package wizard

The Download Support Package wizard provides a selection of various package types. IBM Support provides direction on package type selection as required. To download the package, complete the following steps:

1. Select the type and click **Download**.
2. The output file can be saved to the user's workstation.

Figure 3-113 shows the Download Support Package wizard.

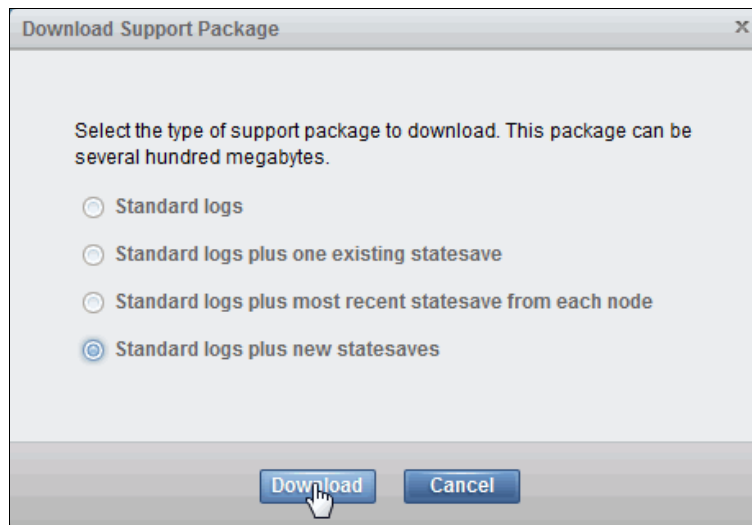


Figure 3-113 Download Support Package wizard

Show full log listing

The Support pane also provides access to the files that are on the node canisters, as shown in Figure 3-114:

1. Click **Show full log listing** to access the node canister files.
2. To save a file to the user's workstation, right-click the file and select **Download**.
3. To change to the file listing to show the files on a partner node canister, select the node canister from the node selection menu.

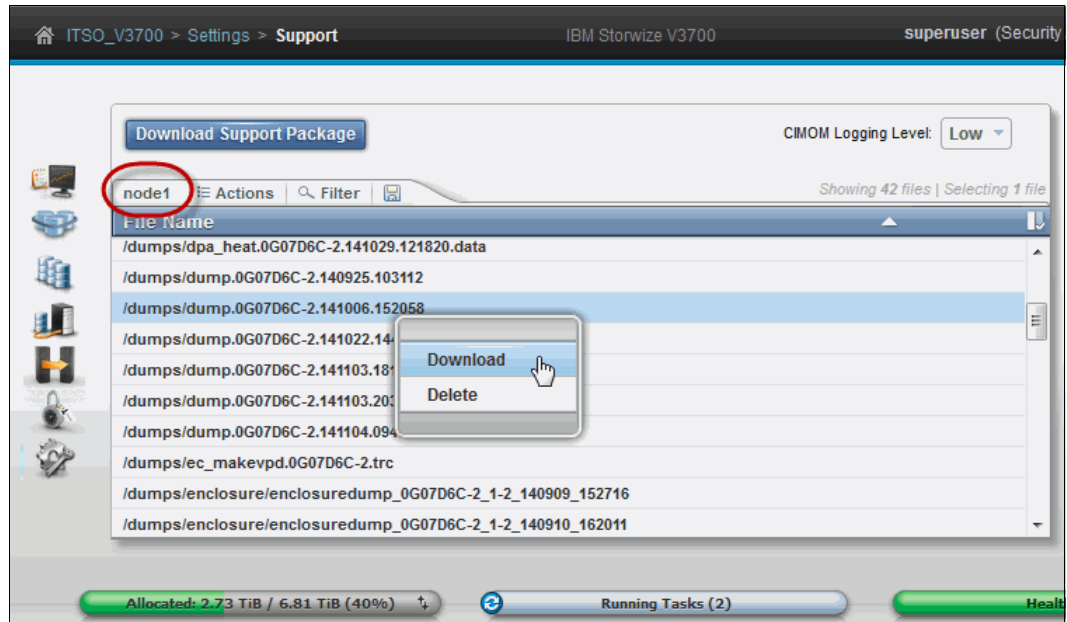


Figure 3-114 Full log listing

GUI Preferences pane

As shown in Figure 3-115, this pane provides options to Refresh GUI Objects, Restore Default Browser Preferences, and configure the Information Center web address. There are also two further check boxes that allow the Extent size to be changed when creating a new Pool and enable the low-resolution graphics mode.

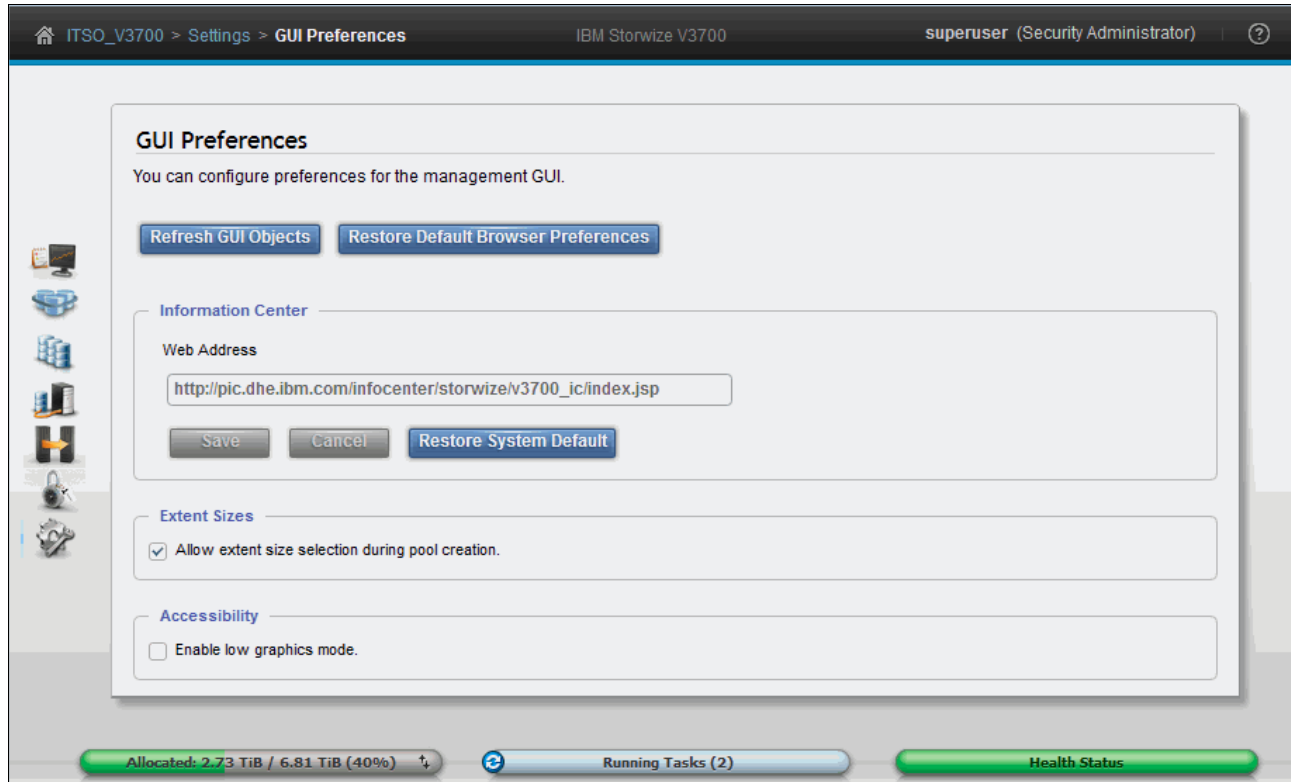


Figure 3-115 GUI Preferences pane

To configure GUI preferences, follow these steps:

1. Select **GUI Preferences** in the Settings menu to open the panel.
2. Select the **Allow extent size selection during pool creation** option to toggle the Advanced Settings button to enable you to specify the size of the extents when creating a new pool, as can be seen in 7.3.1, “Create Pool option” on page 351.
3. Select the **Enable low graphics mode** option if you require remote access over low-bandwidth links. The Function Icons no longer enlarge and list the available functions. However, it is possible to navigate by clicking a Function Icon, and by using the breadcrumb navigation aid.

After selecting the Enable low graphics mode box, the user must log off and log on again for this to take effect.

For more information about the Function Icons, see 3.1.3, “System pane layout” on page 74.

For more information about the breadcrumb navigation aid, see 3.2.2, “Breadcrumb navigation aid” on page 78.

Figure 3-116 shows the management GUI in low graphics mode.

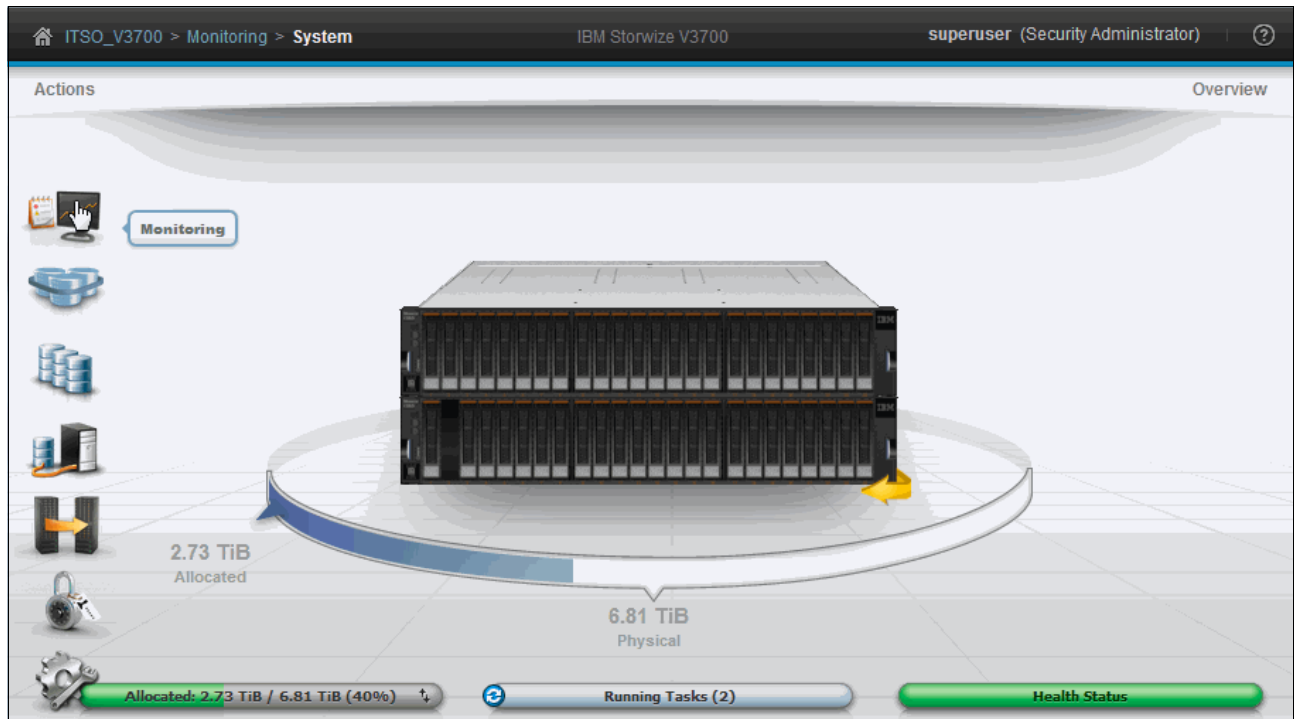


Figure 3-116 Management GUI low graphics mode

3.5 Management GUI help

This section shows how to find help while using the IBM Storwize V3700 management GUI:

- ▶ IBM Storwize V3700 Information Center
- ▶ e-Learning modules
- ▶ Embedded pane help
- ▶ Question mark help
- ▶ Hover help
- ▶ IBM-endorsed YouTube videos

3.5.1 IBM Storwize V3700 Information Center

The best source of information for the IBM Storwize V3700 is the Information Center. For direct access to the online version of the IBM Storwize V3700 Information Center, complete the following steps:

1. Click the help icon in the upper-right corner of any pane.
2. Select **Information Center**, as shown in Figure 3-117.



Figure 3-117 Information Center help

3.5.2 Watching an e-Learning video

IBM e-Learning videos are also available from the **Learning and Tutorials** section of the IBM Storwize V3700 Information Center.

3.5.3 Embedded pane help

The IBM Storwize V3700 provides embedded help for each pane:

1. Click the help icon in the upper-right corner of any pane.
2. Select the first item in the list, which is the pane name.

Figure 3-118 shows the selection of the pane help for the Volumes pane.

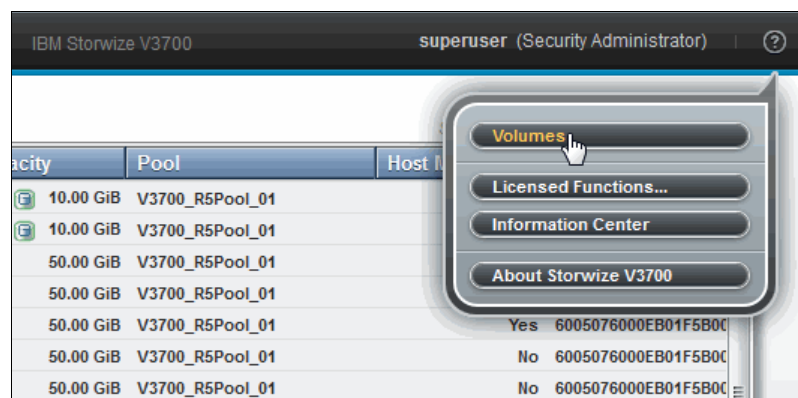


Figure 3-118 Embedded pane help

Figure 3-119 shows the embedded pane help for the Volumes pane, and includes links to various other panes and the Information Center.

Volumes

A volume is a logical disk that the system presents to attached hosts.

Application servers access volumes, not MDisks or drives. To keep a volume accessible even when an MDisk on which it depends has become unavailable, a mirrored copy can be added to a selected volume. Each volume can have a maximum of two copies. Each volume copy is created from a set of extents in a pool.

Thin-provisioned volumes
When you create a volume, you can designate it as thin-provisioned. A thin-provisioned volume has a virtual capacity and a real capacity.

Mirrored volume
By using volume mirroring, a volume can have two physical copies. Each volume copy can belong to a different pool, and each copy has the same virtual capacity as the volume. In the management GUI, an asterisk (*) indicates the primary copy of the mirrored volume. The primary copy indicates the preferred volume for read requests.

Move volumes to another I/O group
You can move a volume to a new I/O group to balance the workload across the canisters in the system without stopping host activity to the volumes. The wizard changes only the I/O group for the volume and does not change where the data resides on the backend storage.

Moving volumes mapped to iSCSI hosts to another I/O group
To move volumes that are mapped to iSCSI hosts to another I/O group, you must unmap the volumes from the hosts before you run the Move Volume to Another I/O Group wizard.

Parent topic: [Pools](#)

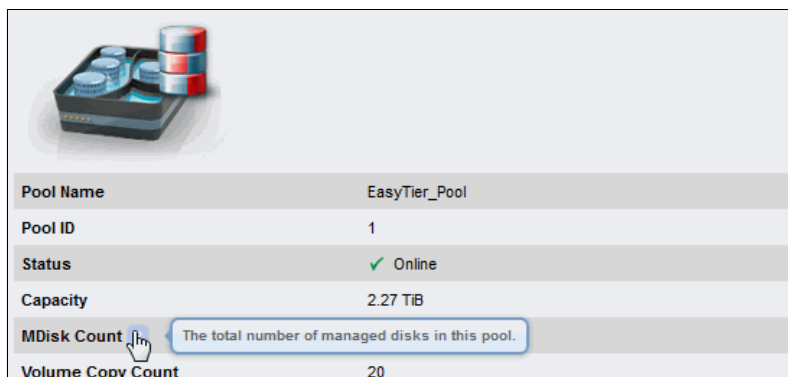
[Contents](#) | [Monitoring](#) | [Pools](#) | [Volumes](#) | [Hosts](#) | [Copy Services](#) | [Access](#) | [Settings](#) | [More Information](#)

Copyright IBM Corporation 2013, 2014. All Rights Reserved.

Figure 3-119 Information pane

3.5.4 Hidden question mark help

The IBM Storwize V3700 provides a *hidden question mark help* feature for some settings and items found in various configuration panes. Access this help feature by hovering next to an item where a question mark is shown. A help bubble is displayed, as shown in Figure 3-120.




Pool Name	EasyTier_Pool
Pool ID	1
Status	✓ Online
Capacity	2.27 TiB
MDisk Count 	The total number of managed disks in this pool.
Volume Copy Count	20

Figure 3-120 Hidden question mark help

3.5.5 Hover help

The IBM Storwize V3700 provides hidden help tags, shown when hovering over various functions and items, as shown in Figure 3-121.

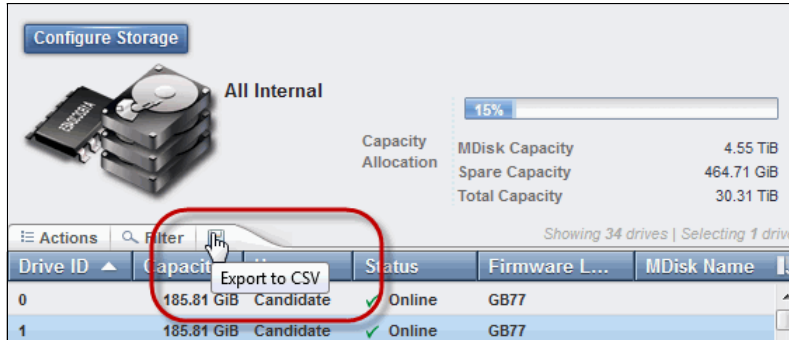


Figure 3-121 Hover help

3.5.6 IBM-endorsed YouTube videos

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

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



Host configuration

This chapter provides an overview of preparing open system hosts for use with IBM Storwize V3700 and using the IBM Storwize V3700 GUI to create logical host objects. This chapter is a prerequisite for creating and mapping Storwize V3700 volumes to hosts, which is described in Chapter 5, “Volume configuration” on page 199.

This chapter includes the following topics:

- ▶ Host attachment overview
- ▶ Preparing the host operating system
- ▶ Configuring hosts in IBM Storwize V3700

4.1 Host attachment overview

A host system is an open system computer that is connected to a switch through a Fibre Channel (FC), Internet Small Computer System interface (iSCSI) or a directly attached serial-attached SCSI (SAS) connection.

In short, IBM Storwize V3700 supports the following host attachment protocols:

- ▶ 8 gigabit (Gb) FC, 10 Gb iSCSI/Fibre Channel over Ethernet (FCoE), or 6 Gb SAS (optional host interface)
- ▶ 6 Gb SAS (standard host interface)
- ▶ 1 Gb iSCSI (standard host interface)

In this chapter, we assume that your hosts are physically ready and attached to your FC and Internet Protocol (IP) network, or directly attached if using SAS host bus adapters (HBAs). We also assume that you have completed the steps that are described in 2.9, “Initial configuration”.

Follow basic switch and zoning recommendations. Also ensure that each host has at least two network adapter ports, that each adapter is on a separate network (or at minimum in a separate zone), and that at least one connection to each node exists per adapter. This setup assures redundant paths for failover and bandwidth aggregation. Further information about storage area network (SAN) configuration is provided in Chapter 2, “Initial configuration” on page 27.

For attachment using SAS, ensure that each host has at least one SAS HBA connection to each IBM Storwize V3700 canister. Further information about configuring SAS hosts is provided in 2.3, “SAS direct attach planning”.

Before new volumes are mapped to the host of your choice, it is advised that preparation is done that will help with ease of use and reliability. There are several steps required on a host system to prepare for mapping new IBM Storwize V3700 volumes.

Use the IBM System Storage Interoperation Center (SSIC) to check which code levels are supported to attach your host to your storage. SSIC is an IBM web tool that checks the interoperation of host, storage, switches, and multipathing drivers. Go to the SSIC web page to get the current IBM Storwize V3700 compatibility information:

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

The complete support matrix is published in the IBM Storwize V3700 Supported Hardware List, Device Driver, Firmware, and Recommended Software Levels V7.4 documents, which are available on the following website:

<http://www.ibm.com/support/docview.wss?uid=ssg1S1004388>

The following sections provide examples of host configuration for each supported attachment type on Microsoft Windows Server 2012 R2 and VMware ESXi 5.5. If you must attach hosts running any other supported operating system (OS), the required information can be found in the IBM Storwize V3700 Knowledge Center:

http://www.ibm.com/support/knowledgecenter/STLM5A/landing/V3700_welcome.html

4.2 Preparing the host operating system

This section describes how to prepare the host attachment that is required to use IBM Storwize V3700 with FC, iSCSI, or SAS connectivity.

In general, the following steps are required to prepare a host to connect to IBM Storwize V3700:

1. Make sure that the current supported system updates are applied to your host operating system.
2. Make sure that the HBAs are physically installed on the host.
3. Make sure that the current, supported HBA firmware and driver levels are installed on your host.
4. Configure the HBA parameters. Although settings are given for each Host OS in the following sections, review the IBM V3700 Information Center to determine the currently supported settings.
5. Configure host input/output I/O parameters (for example, the disk I/O timeout value).
6. Install and configure multipath software.
7. Determine the host worldwide port names (WWPNs).
8. Connect the HBA ports to switches using the appropriate cables, or directly attach to the ports on IBM Storwize V3700.
9. Configure the switches, if applicable.
10. Configure SAN Boot (optional).

Steps 1 - 9 are covered in this chapter. For information about configuring SAN Boot, see Appendix A, “Command-line interface setup and SAN Boot” on page 649.

4.2.1 Windows 2012 R2: Preparing for FC attachment

This section details the steps necessary for preparing for a successful implementation.

Installing and updating supported HBAs

Install a supported HBA model with the current supported firmware and drivers for your configuration:

1. Determine the current supported HBAs and levels for Windows 2012 R2, which are available at the following website:
<http://www.ibm.com/support/docview.wss?uid=ssg1S1004388>
2. Install the driver using Windows Device Manager or vendor tools, such as QLogic QConvergeConsole, Emulex HBAnyware, or Brocade HBA Software Installer.
3. Also, check and update the firmware level of the HBA using the manufacturer’s provided tools.
4. Always check the readme file to see whether there are Windows registry parameters that should be set for the HBA driver.

Important: When using the IBM Subsystem Device Driver Device Specific Module (SDDDSM) for multipath support, the HBA driver must be a Storport (storport.sys) miniport driver.

Configuring Brocade HBAs for Windows

After the device driver and firmware are installed, you must configure the HBAs:

1. Use the Brocade host connectivity manager (HCM) software, or reboot into the HBA basic input/output system (BIOS).
2. Load the adapter defaults.
3. Set the following values:
 - Host Adapter BIOS: Disabled (unless the host is configured for SAN Boot)
 - Queue depth: 4

Configuring QLogic HBAs for Windows

After the device driver and firmware are installed, you must configure the HBAs:

1. Use the QLogic QConvergeConsole software or reboot into the HBA BIOS.
2. Load the adapter defaults.
3. Set the following values:
 - Host Adapter BIOS: Disabled (unless the host is configured for SAN Boot)
 - Adapter Hard Loop ID: Disabled
 - Connection Options: 1 (point-to-point only)
 - LUNs Per Target: 0
 - Port Down Retry Count: 15

Configuring Emulex HBAs for Windows

After the device driver and firmware are installed, you must configure the HBAs:

1. Use the Emulex HBAnyware software or reboot into the HBA BIOS.
2. Load the defaults.
3. Set topology to 1 (10F_Port Fabric).

Setting the Windows disk I/O timeout value

To configure the disk I/O timeout value on Windows 2012 R2, complete the following steps:

1. Press Win+R.
2. In the Run window, enter `regedit` and press Enter. This opens the registry editor.
3. In the registry editor, go to `HKEY_LOCAL_MACHINE\System\CurrentControlSet\services\Disk\TimeOutValue`.
4. Double-click `TimeOutValue`, or right-click it and select **Modify** to view and modify the current timeout value. The window shown in Figure 4-1 opens.

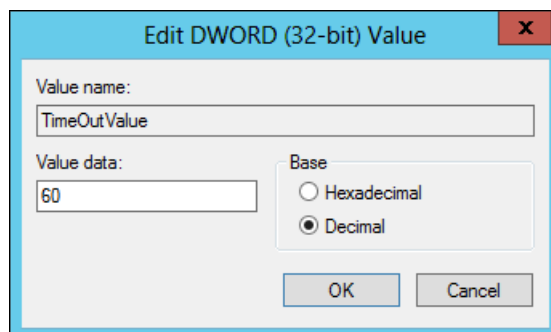


Figure 4-1 Editing the disk I/O timeout value on Windows 2012 R2

The default timeout is 60 seconds, and this is the value suggested by IBM. However, depending on your requirements, you might want to adjust this higher or lower. The following link explains in detail why it might be useful to adjust the timeout:

<http://blogs.msdn.com/b/san/archive/2011/09/01/the-windows-disk-timeout-value-understanding-why-this-should-be-set-to-a-small-value.aspx>

Installing Microsoft MPIO multipathing software

Microsoft Multipath I/O (MPIO) solutions are designed to work with device-specific modules (DSMs) that are written by vendors. The MPIO driver package does not form a complete solution on its own.

By using this joint solution, the storage vendors can design device-specific solutions that are tightly integrated with the Microsoft Windows OS. MPIO in Microsoft Windows 2012 R2 is a DSM designed to work with storage arrays that support the Asymmetric Logical Unit Access (ALUA) control model, such as active/active storage controllers.

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

To enable MPIO on Windows 2012 R2, complete the following steps:

1. Open **Server Manager**.
2. Click **Manage** → **Add Roles and Features**.
3. On the **Installation Type** pane, select **Role-based or feature-based installation** and click **Next**.
4. On the **Server Selection** pane, select the server that you want to configure and click **Next**.
5. On the **Server Roles** pane, click **Next**.
6. On the **Features** pane, select **Multipath I/O** and click **Next**.
7. On the next pane, click **Install** and wait for the installation to complete.
8. When this is complete, restart the server to enable MPIO.

For FC attachment, the installation of SDDDSM is also required. The SDDDSM setup program should enable MPIO automatically. Otherwise, follow the previous steps to enable MPIO manually.

Installing the IBM SDDDSM multipathing software

The SDDDSM is an IBM multipath I/O solution based on Microsoft MPIO technology. This is a device-specific module designed to support IBM storage devices on a range of platforms, including Microsoft Windows. Its purpose is to support a fully redundant multipath configuration environment for IBM Storage Systems.

To ensure correct multipathing with IBM Storwize V3700 when using the FC protocol, SDDDSM must be installed on Microsoft Windows hosts. To install SDDDSM, complete the following steps:

1. Check the SDDDSM download matrix to determine the correct level of SDDDSM to install on Windows 2012 R2 for IBM Storwize V3700:

<http://www.ibm.com/support/docview.wss?uid=ssg1S7001350#WindowsSDDDSM>

2. Download the software package from the following website:

<http://www.ibm.com/support/docview.wss?uid=ssg1S4000350#Storwize>

- Copy the software package to your Windows server and, on the server, run setup.exe to install SDDDSM. A command prompt window opens, as shown in Figure 4-2. Enter Yes to confirm the installation.

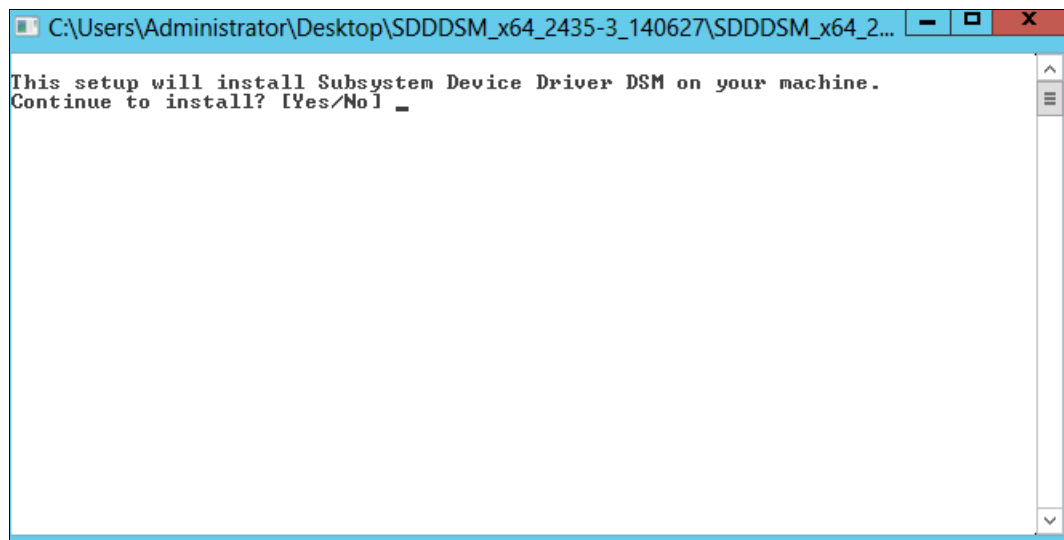


Figure 4-2 Running setup.exe to install SDDDSM on Windows 2012 R2

- After the setup completes, enter Yes to restart the system, as shown in Figure 4-3.

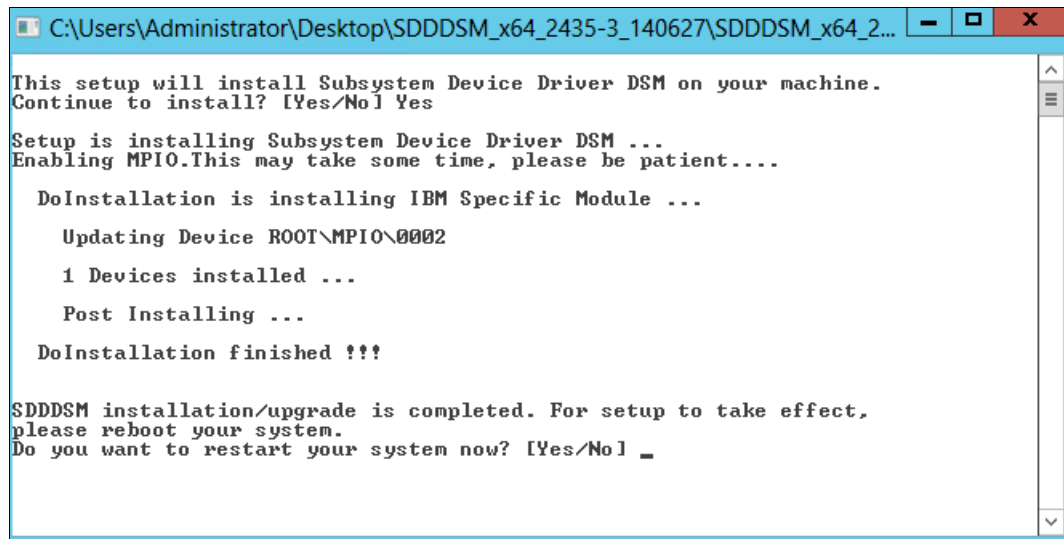


Figure 4-3 Completion of the SDDDSM installation on Windows 2012 R2

- After the server comes back online, the installation should be complete. You can verify the installed driver version by opening a Powershell terminal and running the datapath **query version** command, as shown in Example 4-1.

Example 4-1 Query SDDDSM version details

```
PS C:\Program Files\IBM\SDDDSM>./datapath.exe query version
IBM SDDDSM Version 2.4.3.5-3
Microsoft MPIO Version 6.2.9200.16384
```

For more information about configuring SDDDSM for all supported platforms, see the Multipath Subsystem Device Driver User's Guide. This can be found on the following website:
http://www.ibm.com/support/docview.wss?rs=540&context=ST52G7&q=ssg1*&uid=ssg1S7000303&loc=en_US&cs=utf-8&lang=en%20en

Determining host WWPNS

The worldwide port names (WWPNs) of the FC HBA are required to correctly zone switches and configure host attachment on the IBM Storwize V3700. These can be found using vendor tools, the HBA BIOS, the native Windows command line, or SDDDSM:

1. To determine the host WWPNs using SDDDSM, we can open a Powershell terminal and run the datapath **query wwpn** command, as shown in Example 4-2.

Example 4-2 Query SDDDSM WWPNs

```
PS C:\Program Files\IBM\SDDDSM> .\datapath.exe query wwpn
Adapter Name      PortWWN
-----
Scsi Port2:      21000024FF2D0CC2
Scsi Port3:      21000024FF2D0CC3
```

2. When the WWPNs are known, connect the cables and perform any necessary switch configuration (for example, zoning).

The host is now prepared to connect to the IBM Storwize V3700.

These are the basic steps to prepare a Windows 2012 R2 host for FC attachment. For information about configuring FC attachment on the IBM Storwize V3700 side, see 4.3.2, “Configuring IBM Storwize V3700 for FC connectivity”.

4.2.2 Windows 2012 R2: Preparing for iSCSI attachment

This section details the steps necessary for preparing for a successful implementation.

Installing and updating supported HBAs

Install a supported HBA model with the current supported firmware and drivers for your configuration:

1. Determine the current supported HBAs and levels for Windows 2012 R2, which are available on the following website:
<http://www.ibm.com/support/docview.wss?uid=ssg1S1004388>
2. Install the driver using Windows Device Manager or vendor tools.
3. Also, check and update the firmware level of the HBA using the manufacturer's provided tools.
4. Always check the readme file to see whether there are Windows registry parameters that should be set for the HBA driver.

Important: For Converged Network Adapters (CNAs) that are capable of both FC and iSCSI, it is important to ensure that the Ethernet networking driver is installed in addition to the FCoE driver. This is required for configuring iSCSI.

If you are using a hardware iSCSI HBA, see the manufacturer's documentation and Storwize V3700 Knowledge Center for further details about hardware and host OS configuration:

http://www.ibm.com/support/knowledgecenter/STLM5A/landing/V3700_welcome.html

The following section describes how to configure iSCSI using the software initiator.

Enabling the iSCSI Initiator

On Microsoft Windows systems, the Microsoft iSCSI Initiator can be used to connect to external storage attached with Ethernet. Using the iSCSI protocol enables the use of block storage in a SAN at low cost compared to FC or SAS.

In Windows 2012 R2, the Microsoft iSCSI Initiator is preinstalled. To enable it, complete the following steps:

1. To start the iSCSI Initiator, open **Server Manager** and go to **Tools** → **iSCSI Initiator**, as shown in Figure 4-4.




Figure 4-4 Starting the iSCSI Initiator

2. If this is the first time you have started the iSCSI Initiator, the window shown in Figure 4-5 should open. Click **Yes** to enable the iSCSI Initiator service, which should now start automatically on restart.

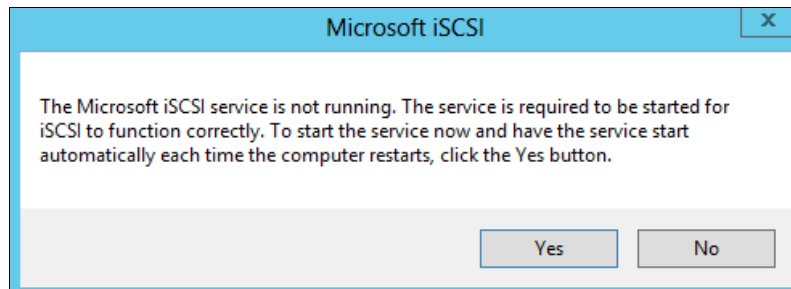


Figure 4-5 Enabling the iSCSI Initiator service on restart

3. The iSCSI Initiator window should open. Select the **Configuration** tab, as shown in Figure 4-6.

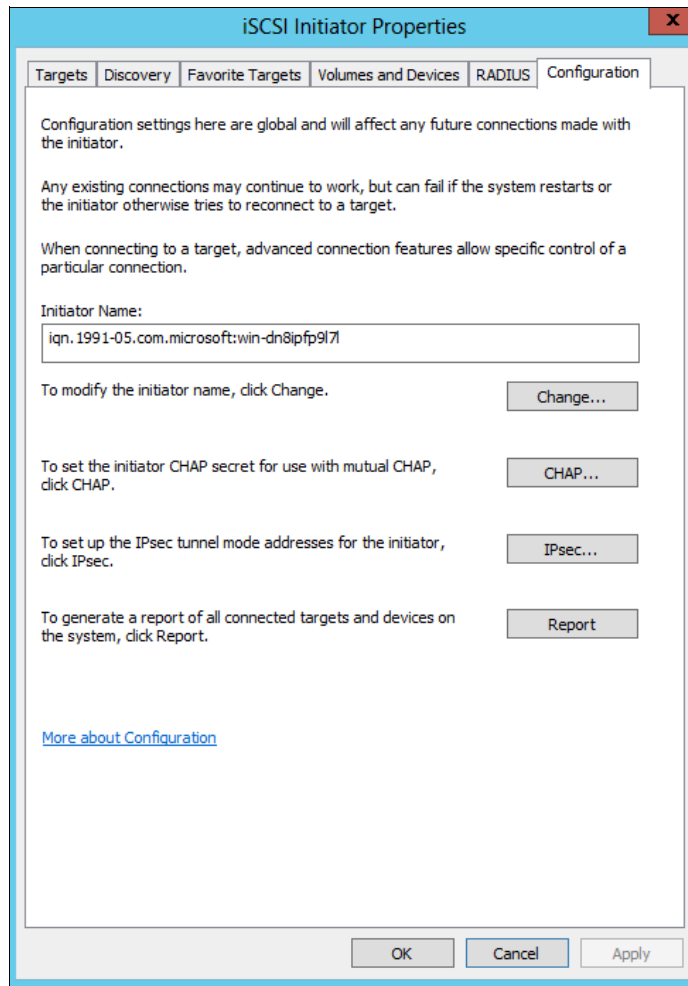


Figure 4-6 Configuring the iSCSI Initiator

4. Make a note of the Initiator Name. This is the iSCSI Qualified Name (IQN) for the host, which is required for configuring iSCSI host attachment on the IBM Storwize V3700. The IQN identifies the host in the same way that WWPNs identify an FC or SAS host.

From the Configuration tab it is possible to change the Initiator Name, enable Challenge Handshake Authentication Protocol (CHAP) authentication and more. However, these tasks are beyond the scope of our basic setup. CHAP authentication is disabled by default. For more information, Microsoft has published a detailed guide about configuring the iSCSI Initiator on the following website:

<http://technet.microsoft.com/en-us/library/ee338476%28v=ws.10%29.aspx>

Configuring Ethernet ports

It is advised to use separate dedicated ports for host management and iSCSI. In this case, we need to configure IPs on the iSCSI ports in the same subnet and VLAN as the external storage to which want to attach.

To configure Ethernet port IPs on Windows 2012 R2, complete the following steps:

1. Go to **Control** → **Network and Internet** → **Network and Sharing Center**. This should take you to the window shown in Figure 4-7.

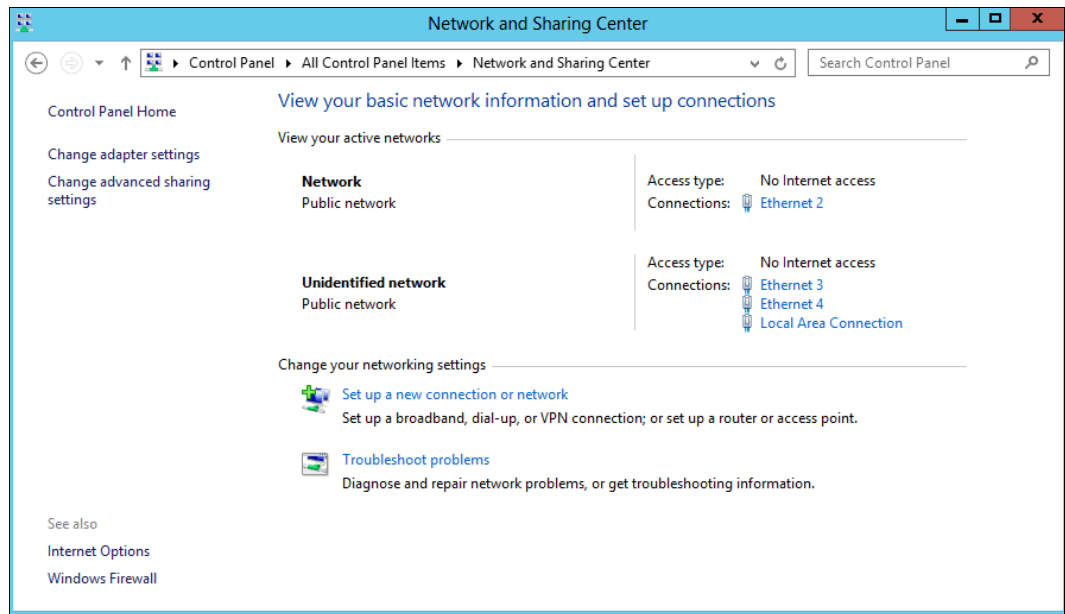


Figure 4-7 Network and Sharing Center in Windows 2012 R2

In this case, we have two networks visible to the system. The first network is the one we are using to connect to the server, consisting of a single dedicated Ethernet port for management. The second network is our iSCSI network, consisting of two dedicated Ethernet ports for iSCSI. It is advised to have at least two dedicated ports for failover purposes.

2. To configure an IP address, click one of the iSCSI Ethernet connections (in this case, Ethernet 3 or Ethernet 4). The window shown in Figure 4-8 should open.




Figure 4-8 Ethernet Status

3. To configure the IP address, click **Properties**. The window shown in Figure 4-9 should open.

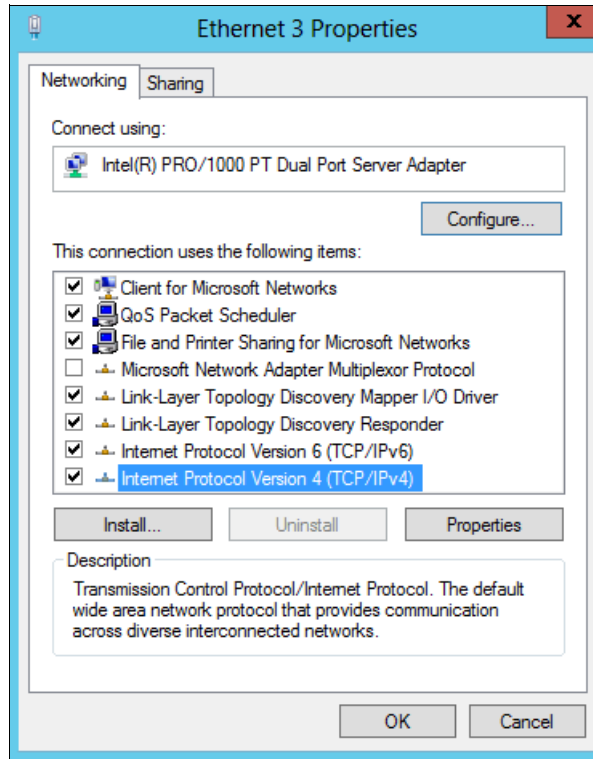


Figure 4-9 Ethernet Properties

4. If you are using IPv6, select **Internet Protocol Version 6 (TCP/IPv6)** and click **Properties**. Otherwise, select **Internet Protocol Version 4 (TCP/IPv4)** and click **Properties** to configure an IPv4 address.
5. For IPv4, the window shown in Figure 4-10 on page 161 should open. To manually set the IP, select **Use the following address** and enter an IP address, Subnet mask, and Default gateway. Set the DNS server address if required. Click **OK** to confirm.

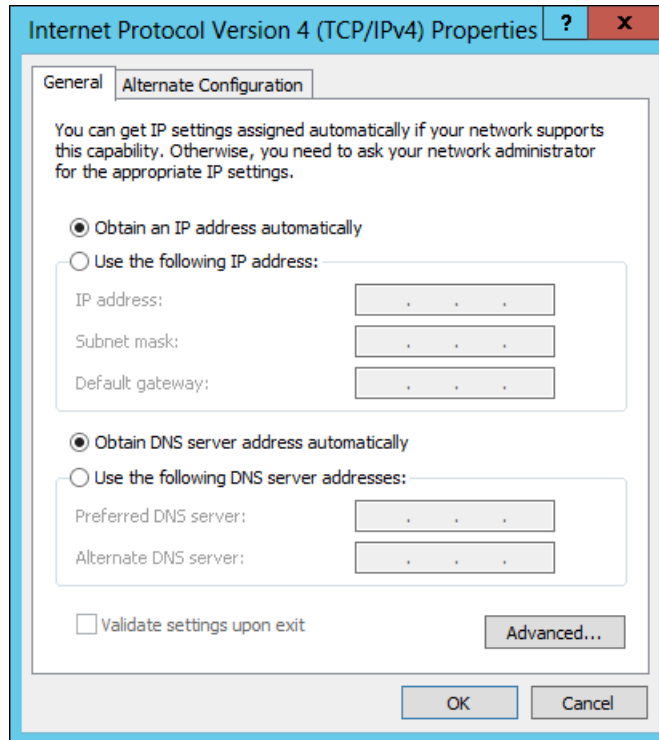


Figure 4-10 Configuring an IPv4 address in Windows 2012 R2

6. Repeat the previous steps for each port that you want to configure for iSCSI attachment.

The Ethernet ports should now be prepared for iSCSI attachment.

Setting the Windows registry keys

You can adjust the system registry to make iSCSI operations more reliable on Windows 2012 R2. To do this, complete the following steps. You must also look at the suggested guidance for your application. To set the registry keys, complete the following steps:

1. Press Win+R.
2. In the dialog box, enter `regedit` and press Enter. This will open the registry editor.
3. In the registry editor, locate the following key:
`HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<bus ID>\Parameters\LinkDownTime`
 If the value for the `LinkDownTime` key is not 120 (decimal value), change the value to 120.
4. In the registry editor, locate the following key:
`HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<bus ID>\Parameters\MaxRequestHoldTime`
 If the value for the `MaxRequestHoldTime` key is not 120 (decimal value) change it to 120.
5. In the registry editor, locate the following key:
`HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<bus ID>\Parameters\MaxPendingRequests`
 If the value for the `MaxPendingRequests` key is not 2048 (decimal value), change it to 2048.

6. In the registry editor, locate the following key:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Disk\TimeOutValue

If the value for the TimeOutValue key is not 60 (decimal value), change the value to 60.

For more details about configuring the disk I/O timeout value, see 4.2.1, “Windows 2012 R2: Preparing for FC attachment”.

7. Restart Windows for these changes to take effect.

Multipath support for iSCSI on Windows

For multipathing with iSCSI, you need to enable MPIO. See 4.2.1, “Windows 2012 R2: Preparing for FC attachment” for instructions on enabling MPIO.

It is important to note that SDDDSM is *not* supported for iSCSI attachment, so do not follow the steps to install this as for FC or SAS.

These are the basic steps to prepare a Windows 2012 R2 host for iSCSI attachment. For information about configuring iSCSI attachment on the IBM Storwize V3700 side, we provide more information later in this book.

4.2.3 Windows 2012 R2: Preparing for SAS attachment

In Windows 2012, complete the following actions to prepare for SAS attachment:

- ▶ Installing and updating supported HBAs
- ▶ Determining host WWPNS
- ▶ Configuring SAS HBAs on Windows
- ▶ Enabling multipath support for SAS on Windows

Installing and updating supported HBAs

Install a supported SAS HBA with the current supported firmware and drivers for your configuration:

1. Determine the current supported HBAs and levels for Windows 2012 R2, which are available on the following website:
<http://www.ibm.com/support/docview.wss?uid=ssg1S1004388>
2. Install the driver using Windows Device Manager or vendor tools.
3. Also, check and update the firmware level of the HBA using the manufacturer’s provided tools.
4. Always check the readme file to see whether there are Windows registry parameters that should be set for the HBA driver.

For LSI cards, the following website provides some guidance on flashing the HBA firmware and BIOS:

http://mycusthelp.info/LSI/_cs/AnswerDetail.aspx?sSessionID=&aid=8352

Determining host WWPNS

The worldwide port names (WWPNs) of the SAS HBA are required to configure host attachment on the IBM Storwize V3700.

The host WWPNS *can* be found using vendor tools or the HBA BIOS.

However, the easiest way is to determine the host WWPNs is to complete the following steps:

1. Connect the SAS cables to the ports on the IBM Storwize V3700.
2. Log on to the Storwize CLI with Secure Shell (SSH).
3. Run `svcinfo lssasportcandidate`, as shown in Example 4-3.

Example 4-3 The svcinfo command

```
IBM_Storwize:ITS0_V3700:superuser>svcinfo lssasportcandidate
sas_wwpn
500062B200556140
500062B200556141
```

For information about connecting to IBM Storwize V3700 with SSH and using the Storwize CLI, see Appendix A, “Command-line interface setup and SAN Boot” on page 649.

Configuring SAS HBAs on Windows

The following settings are suggested:

- ▶ I/O Timeout for Block Devices: 10
- ▶ I/O Timeout for Sequential Devices: 10
- ▶ I/O Timeout for Other Devices: 10
- ▶ LUNs to Scan for Block Devices: All
- ▶ LUNs to Scan for Sequential Devices: All
- ▶ LUNs to Scan for Other Devices: All

Enabling multipath support for SAS on Windows

For multipathing with SAS, you need to enable MPIO and install SDDDSM. See Chapter 4.2.1, “Windows 2012 R2: Preparing for FC attachment” on page 151 for instructions about this.

These are the basic steps to prepare a Windows 2012 R2 host for SAS attachment. For information about configuring SAS attachment on the IBM Storwize V3700 side, see Chapter 4.3.5, “Creating SAS hosts” on page 195. For information about discovering mapped SAS volumes from Windows 2012 R2, see 5.3.3, “Discovering SAS volumes from Windows 2012 R2”.

4.2.4 VMware ESXi 5.5: Preparing for FC attachment

Complete the following actions to prepare for FC attachment:

- ▶ Installing and updating supported HBAs
- ▶ Configuring Brocade HBAs for VMware ESX
- ▶ Configuring QLogic HBAs for VMware ESX
- ▶ Configuring Emulex HBAs for VMware ESXi
- ▶ Configuring multipath support for FC on VMware ESXi
- ▶ Determining host WWPNs

Installing and updating supported HBAs

Install a supported HBA with the current supported firmware and drivers for your configuration:

1. A list of the current supported HBAs and levels for VMware ESXi 5.5 are available at this website:

<http://www.ibm.com/support/docview.wss?uid=ssg1S1004388>

2. Install the driver using VMware vSphere Client, the ESXi CLI, or vendor tools.

3. Also, check and update the firmware level of the HBA using the manufacturer's provided tools.
4. Always check the readme file to see whether there is further configuration required for the HBA driver.

See the VMware Compatibility Guide for more information:

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

Configuring Brocade HBAs for VMware ESX

Follow these steps to configure Brocade HBAs:

1. After the firmware is installed, load the default settings for all of your adapters that are installed on the host system.
2. Ensure that the Adapter BIOS is disabled, unless SAN Boot is used.

Configuring QLogic HBAs for VMware ESX

After the firmware is installed, you must configure the HBAs. To perform this task, use the QLogic QConvergeConsole software or the HBA BIOS:

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

Configuring Emulex HBAs for VMware ESXi

Follow these steps to configure Emulex HBAs:

1. After the firmware is installed, load the default settings for all of your adapters that are installed on the host system.
2. Ensure that the Adapter BIOS is disabled, unless SAN Boot is used.

Configuring multipath support for FC on VMware ESXi

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

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

The IBM Storwize V3700 is an active/active storage device. Since VMware ESXi 5.0 and later, the suggested multipathing policy is *Round Robin*. Round Robin performs dynamic load balancing for I/O.

If you do not want to have the I/O balanced over all available paths, the *Fixed* policy is supported. This policy setting can be configured for every volume. Set this policy after IBM Storwize V3700 volumes have been mapped to the ESXi host. For more information about mapping volumes to hosts, see Chapter 5, “Volume configuration” on page 199.

Determining host WWPNs

The WWPNs of the FC HBA are required to correctly zone switches and configure host attachment on the IBM Storwize V3700. On VMware ESXi 5.5, these can be found using the VMware vSphere Client.

Note: In VMware ESXi 5.5, some new features can only be accessed through the vSphere web client. However, we do not demonstrate any of these features, so this and all of the following examples continue to focus on using the desktop client.

To determine the host WWPNs, complete the following steps:

1. Connect to the ESXi server (or VMware vCenter) using the VMware vSphere Client.
2. Browse to the **Configuration** tab. HBA WWPNs are listed under **Storage Adapters**, as shown in Figure 4-11.

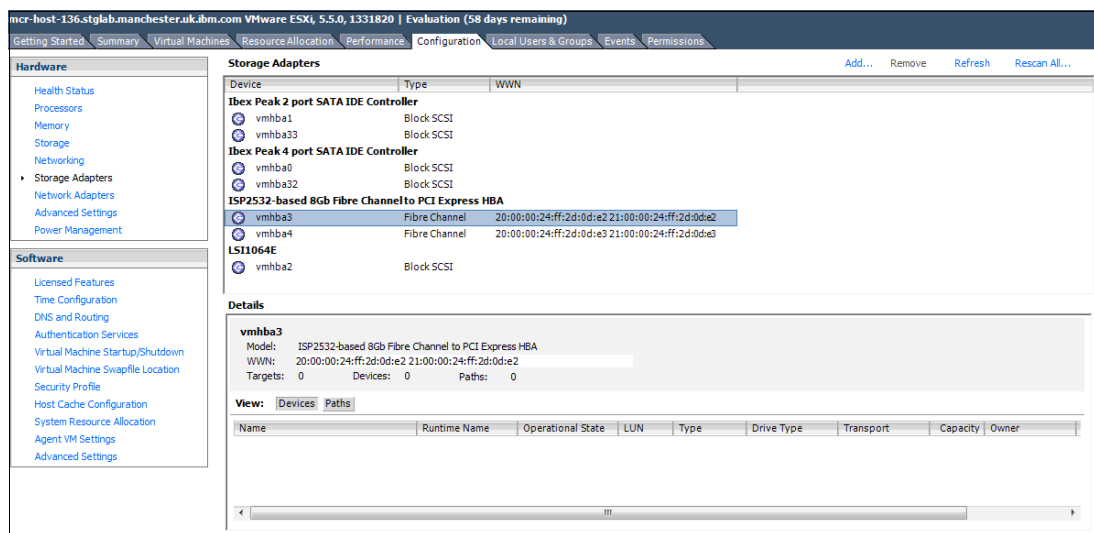


Figure 4-11 FC WWPNs in VMware vSphere Client

These are the basic steps required to prepare a VMware ESXi 5.5 host for FC attachment. We provide more information later in this book.

4.2.5 VMware ESXi 5.5: Preparing for iSCSI attachment

Installing and updating supported HBAs

Install a supported HBA model with the current firmware and drivers for your configuration:

1. Determine the current supported HBAs and levels for VMware ESXi 5.5, which are available on the following website:
<http://www.ibm.com/support/docview.wss?uid=ssg1S1004388>
2. Install the driver using VMware vSphere Client, the ESXi CLI, or vendor tools.
3. Also, check and update the HBA firmware level using the manufacturer's provided tools.
4. Always check the readme file to see whether there is further configuration required for the HBA driver.

See the VMware Compatibility Guide for more information:

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

Important: For Converged Network Adapters (CNAs) that are capable of both FC and iSCSI, it is important to ensure that the Ethernet networking driver is installed in addition to the FCoE driver. This is required for configuring iSCSI.

If you are using a hardware iSCSI HBA, see the manufacturer's documentation and IBM Storwize V3700 Knowledge Center for further details about configuring the hardware and host OS. The following section describes how to configure iSCSI using the software initiator.

Enabling the iSCSI Initiator

The iSCSI initiator is installed by default on VMware ESXi 5.5, and only needs to be enabled. To do this, you need to add and enable the iSCSI Software Adapter. Complete the following steps:

1. Connect to your ESXi server using the VMware vSphere Client.
2. Browse to the **Configuration** tab and select **Storage Adapters**.

This should open the view shown in Figure 4-12.

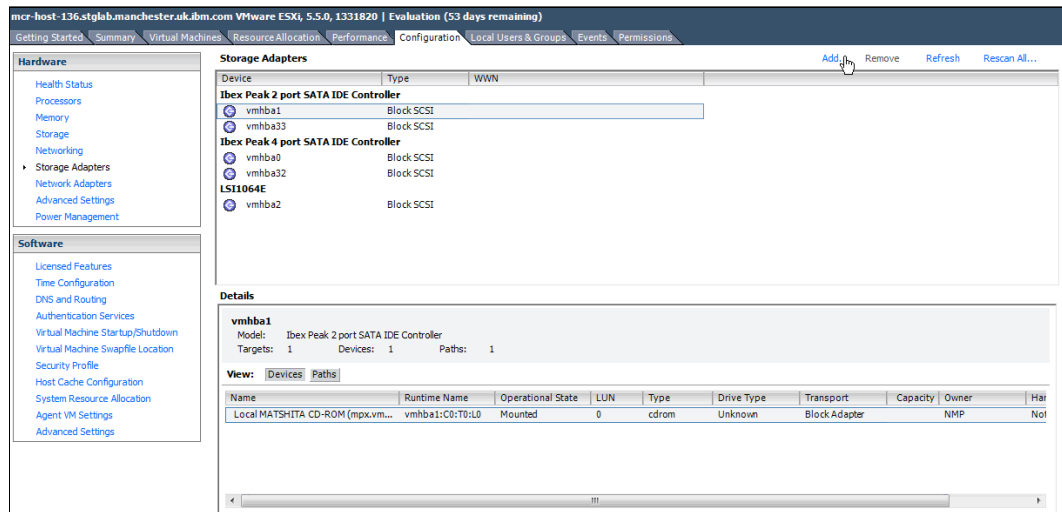


Figure 4-12 Adding a new storage adapter using VMware vSphere Client

3. Click **Add**, and the window shown in Figure 4-13 should open.




Figure 4-13 Add Software iSCSI Adapter

4. Select **Add Software iSCSI Adapter** and click **OK**. The window shown in Figure 4-14 should open.

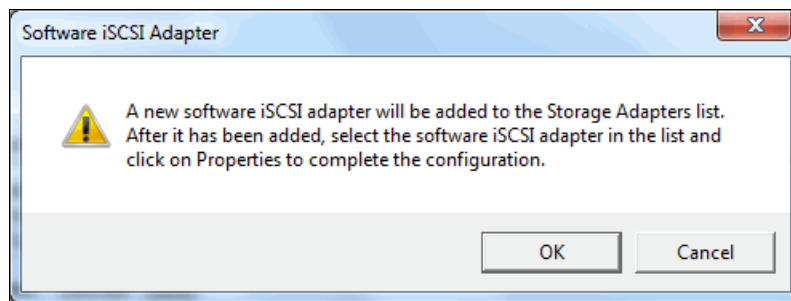


Figure 4-14 Confirm addition of the Software iSCSI Adapter

5. Click **OK** to confirm. The Software iSCSI Adapter should now be listed under **Storage Adapters**, as shown in Figure 4-15.

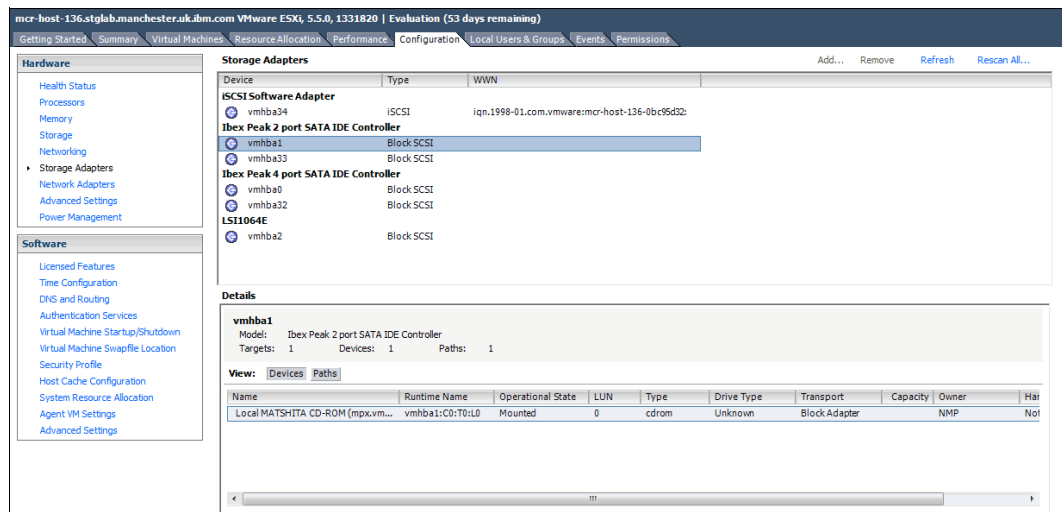


Figure 4-15 The Storage Adapters view showing the newly added iSCSI Software Adapter

- To find the iSCSI initiator name and configure the iSCSI Software Adapter, right-click the adapter and click **Properties**. The window shown in Figure 4-16 should open. Make a note of the initiator name and perform any other necessary configuration.

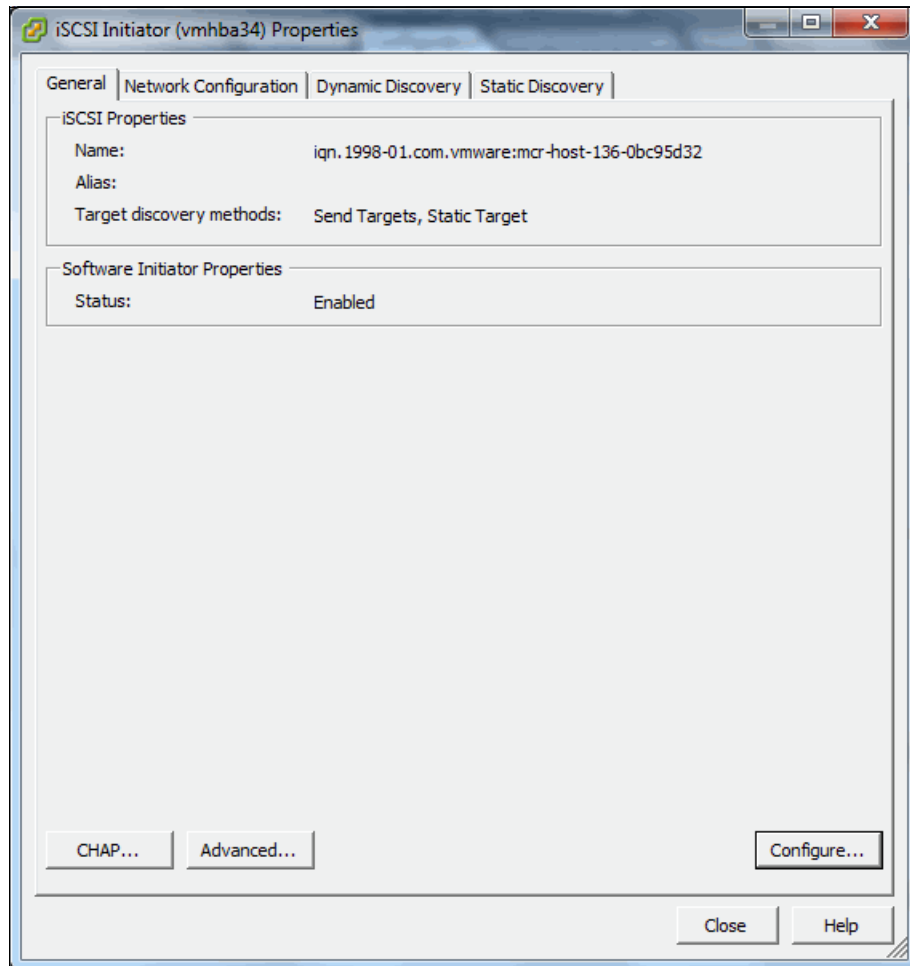


Figure 4-16 Properties of the iSCSI software adapter

The VMware ESXi 5.5 iSCSI initiator is now enabled and ready for use with IBM Storwize V3700. The next step required is to configure the Ethernet adapter ports.

Configuring Ethernet ports

To configure Ethernet adapter ports for iSCSI attachment in VMware ESXi 5.5, complete the following steps:

1. Browse to the **Configuration** tab and select **Networking**, as shown in Figure 4-17.

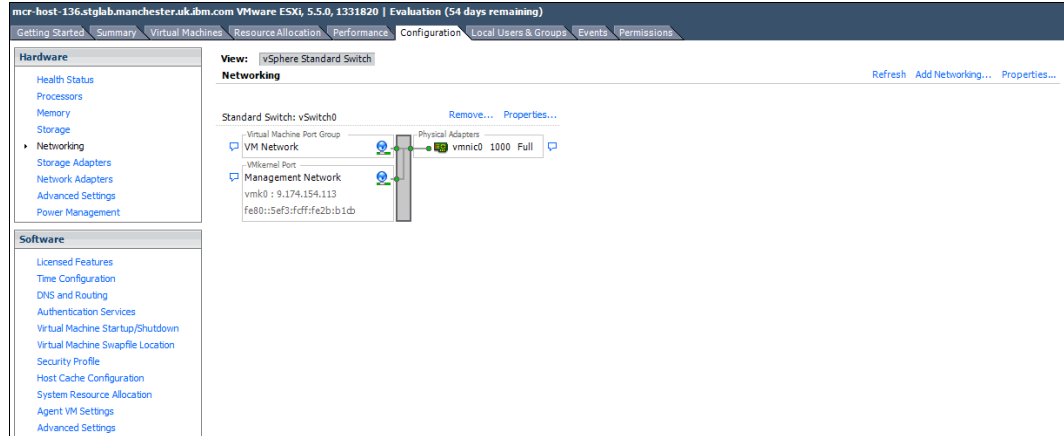


Figure 4-17 The network Configuration tab in VMware vSphere Client

2. Click **Add Networking** to start the Add Network Wizard, as shown in Figure 4-18.

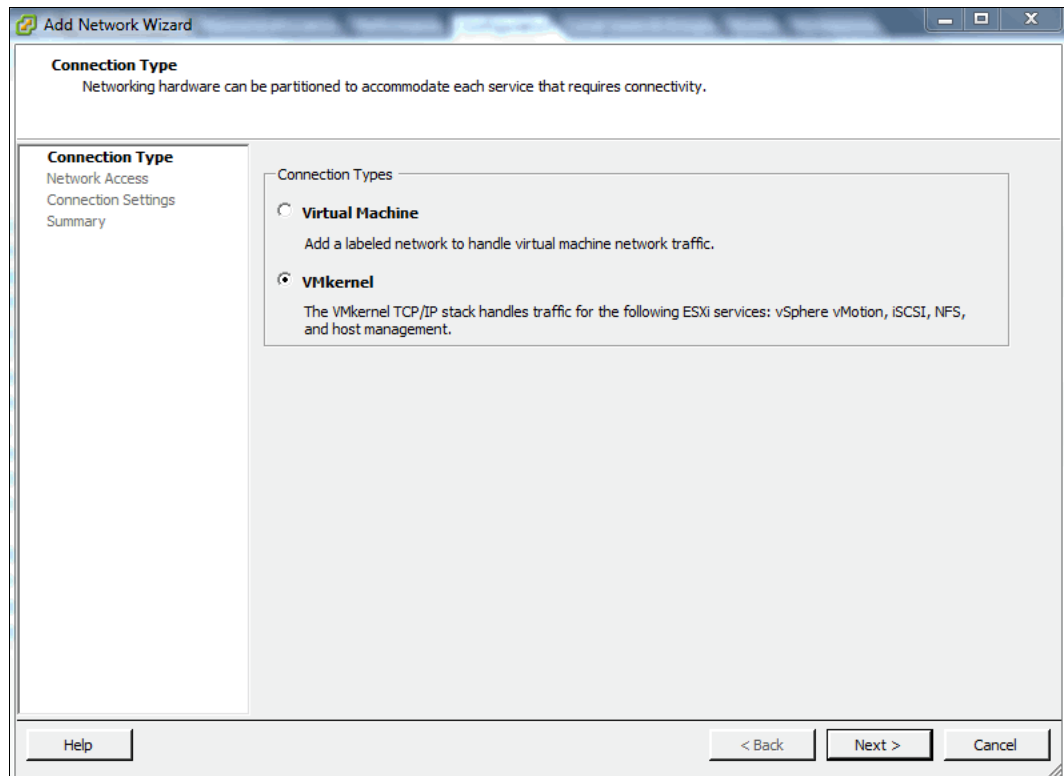


Figure 4-18 The Add Network Wizard in VMware vSphere Client

3. Select **VMkernel** and click **Next**.

4. Create a new vSwitch by selecting one or more network interfaces to use for iSCSI traffic and click **Next**, as shown in Figure 4-19.

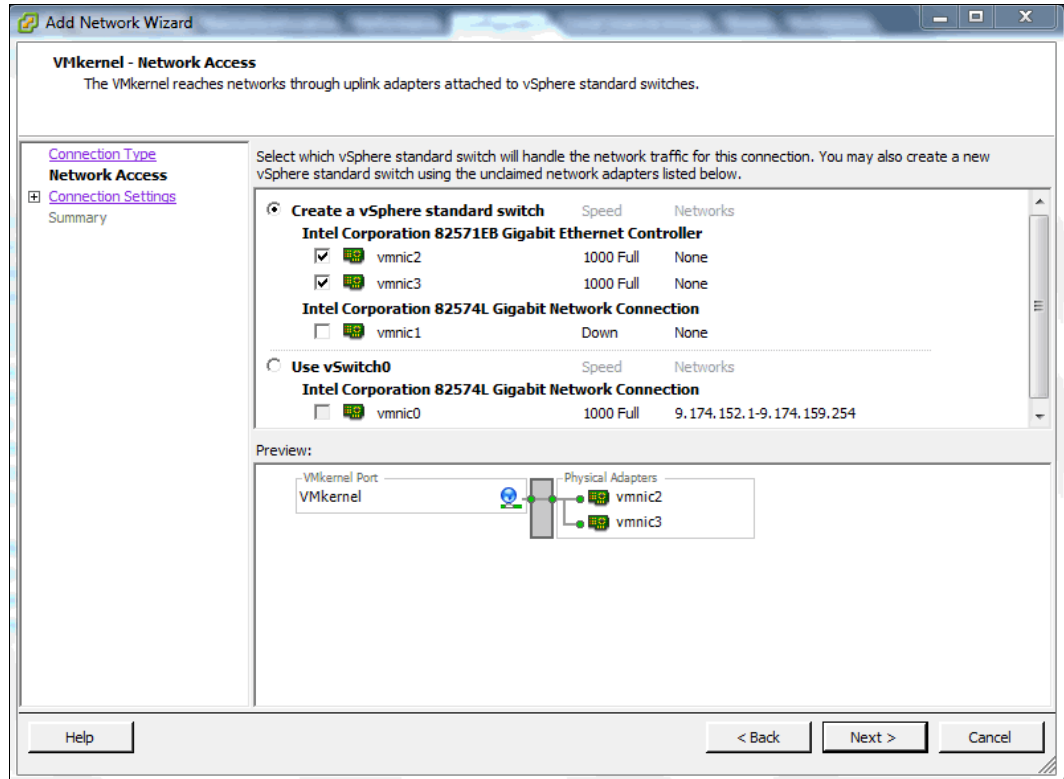


Figure 4-19 Creating a new vSwitch using two physical adapter ports

Note: It is advised to select at least two interfaces for failover purposes.

- The wizard then prompts you to create a single VMkernel port within the new vSwitch, as shown in Figure 4-20.
- Enter a meaningful Network Label and, if using VLAN tagging, enter a VLAN ID or select **All** from the drop-down menu.

There are other options available. See VMware documentation for details about configuring these.

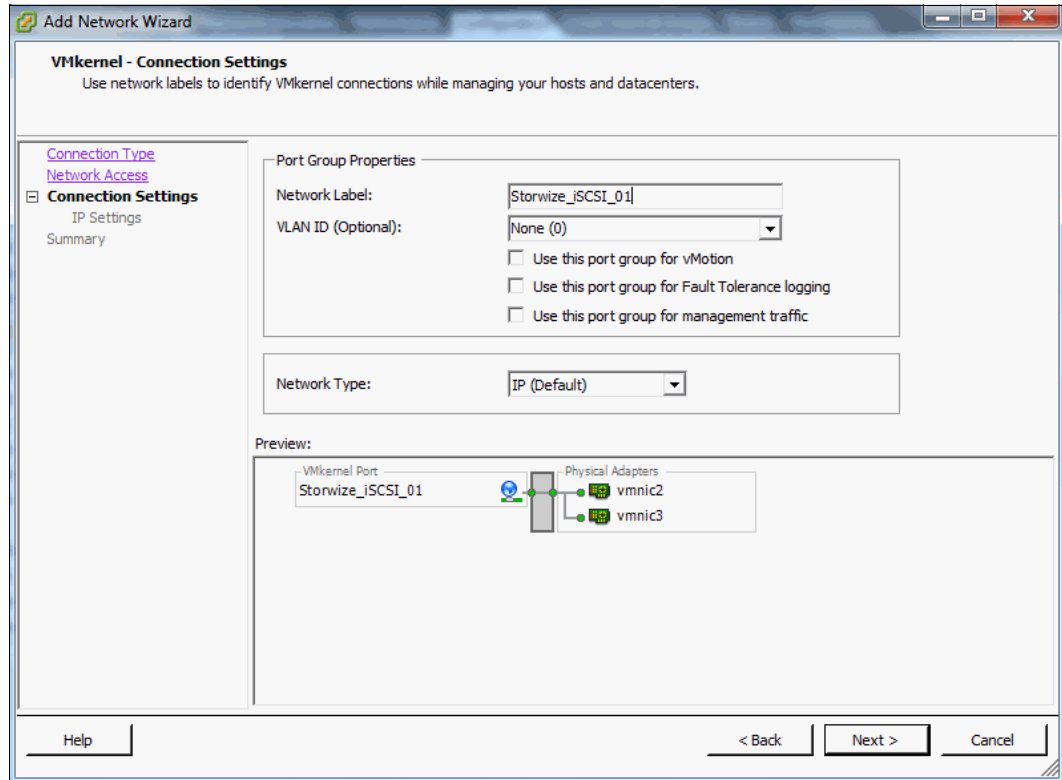


Figure 4-20 Configuring a VMkernel port on the new vSwitch

- Click **Next** to continue.

8. Enter an IP address for your iSCSI network. You should use a dedicated network for iSCSI traffic, as shown in Figure 4-21.




Figure 4-21 Configuring IP settings for the new VMkernel port

9. Click **Next** to see what is added to the host network settings and **Finish** to exit from iSCSI configuration. The network configuration page should now look like Figure 4-22.

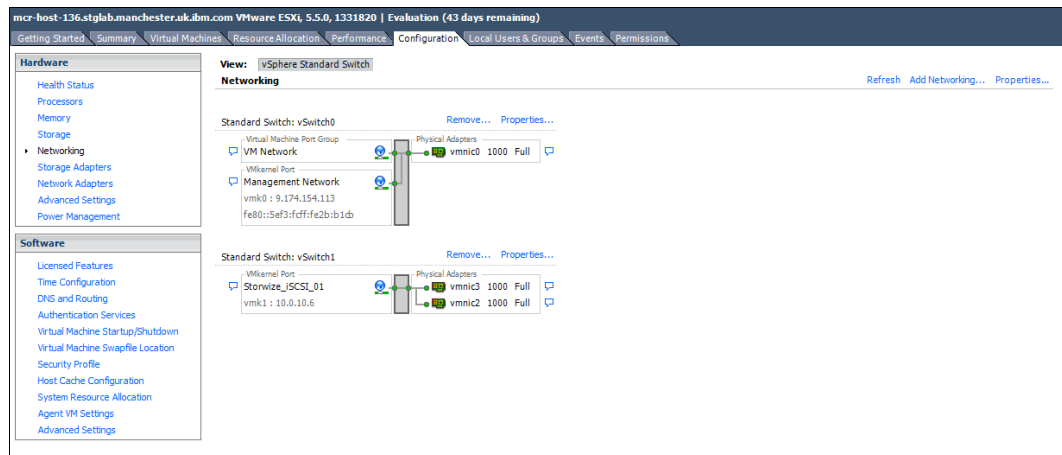


Figure 4-22 The network Configuration page with the newly created vSwitch and VMkernel port

10. To add further VMkernel ports to this switch, click **Properties** to open the window shown in Figure 4-23, then click **Add**. This reopens the same wizard you have just gone through, so follow the same steps to add a second VMkernel port.

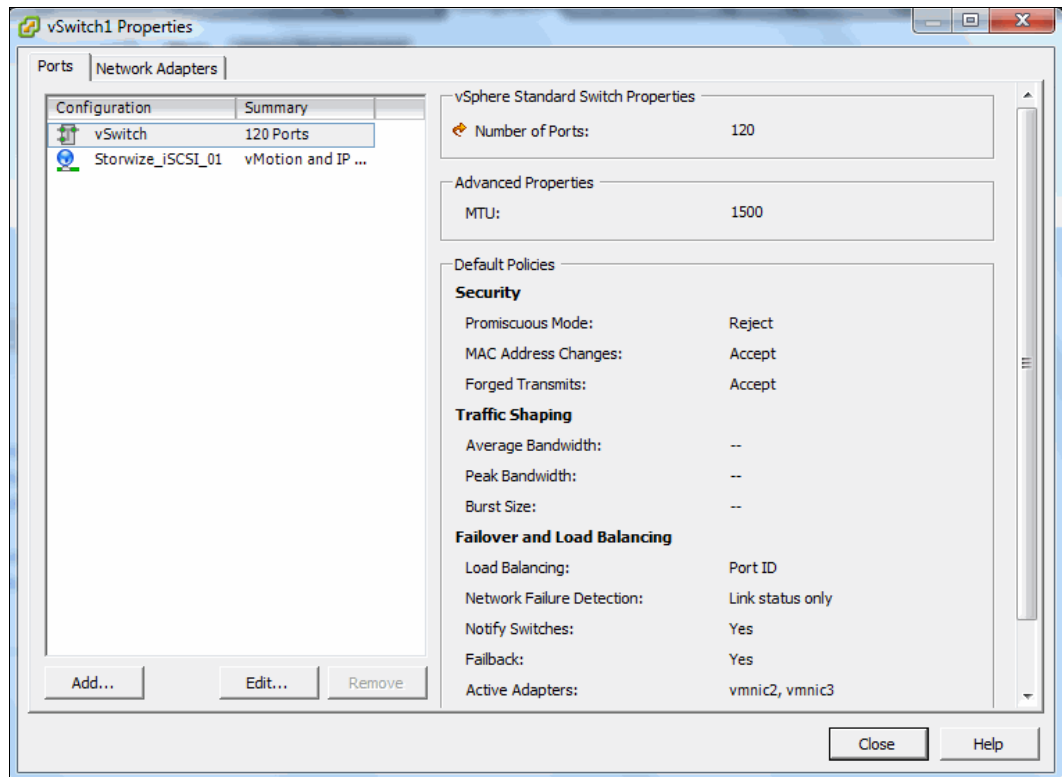


Figure 4-23 Adding another VMkernel port in the vSwitch Properties

11. In this example, you create a second VMkernel port with IP 10.0.10.7. The network configuration page should now look like Figure 4-24.

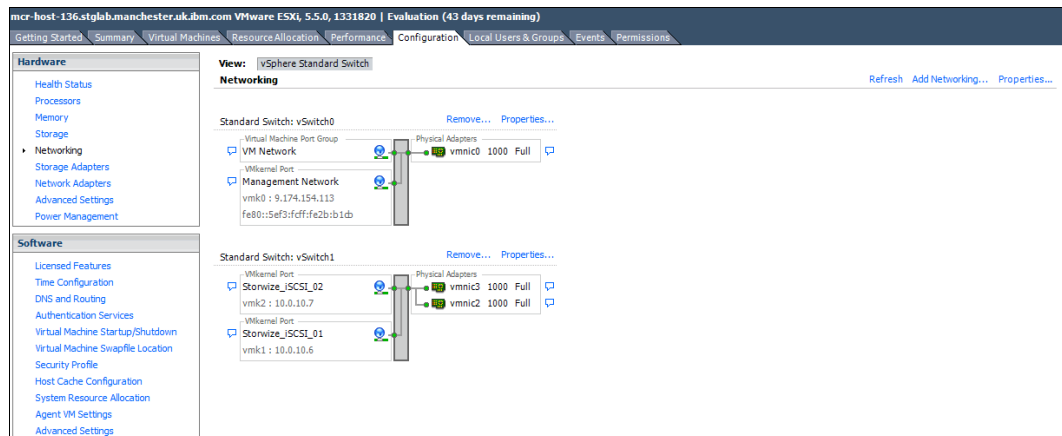


Figure 4-24 The finished network Configuration page

12. You must now bind the VMkernel ports to the iSCSI Software Adapter. To do this, go to **Storage Adapters**, right click the iSCSI Software Adapter and click **Properties**.

13. Then, go to **Network Configuration** and click **Add** to start adding the VMkernel ports.

14. The window shown in Figure 4-25 should open. Click **OK** to add the first port, then repeat for the second.

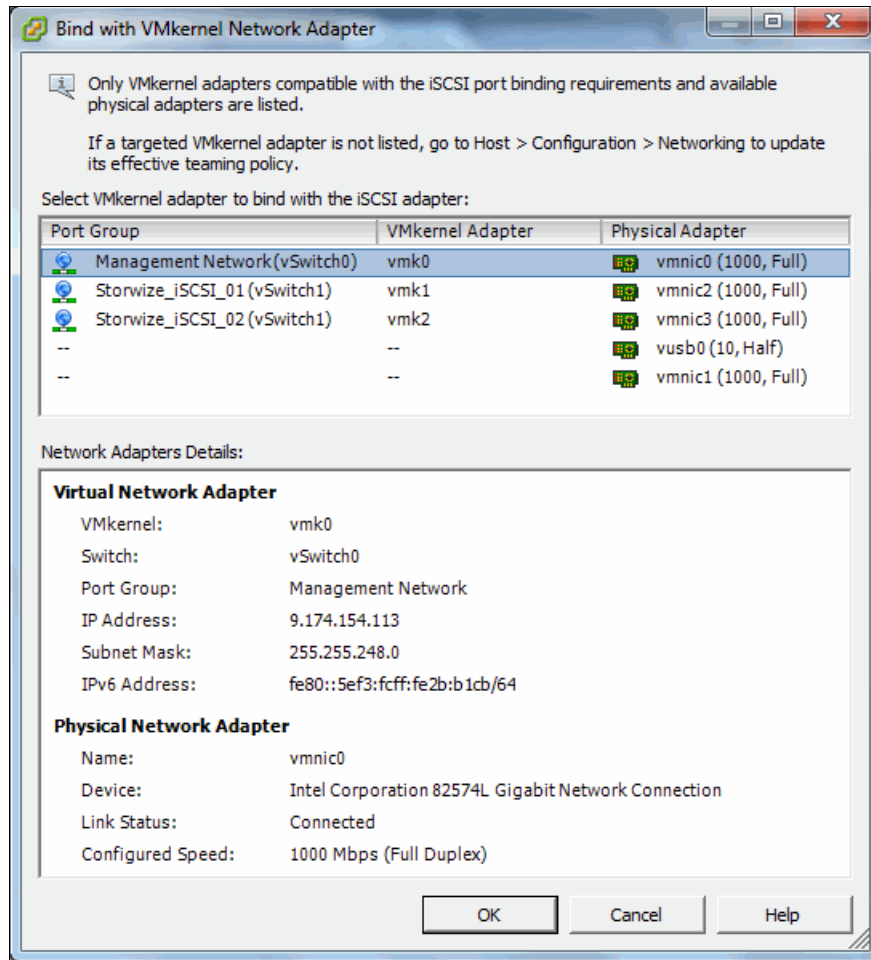


Figure 4-25 Adding VMkernel ports to the iSCSI Adapter

Important: If the VMkernel ports are reported as *Not Compliant*, complete the steps described in “**Configuring multipath support for iSCSI on VMware ESXi**” on **page 176** and try again.

15. The **Network Configuration** tab should now look like Figure 4-26. Click **Close** to finish binding the VMkernel ports.

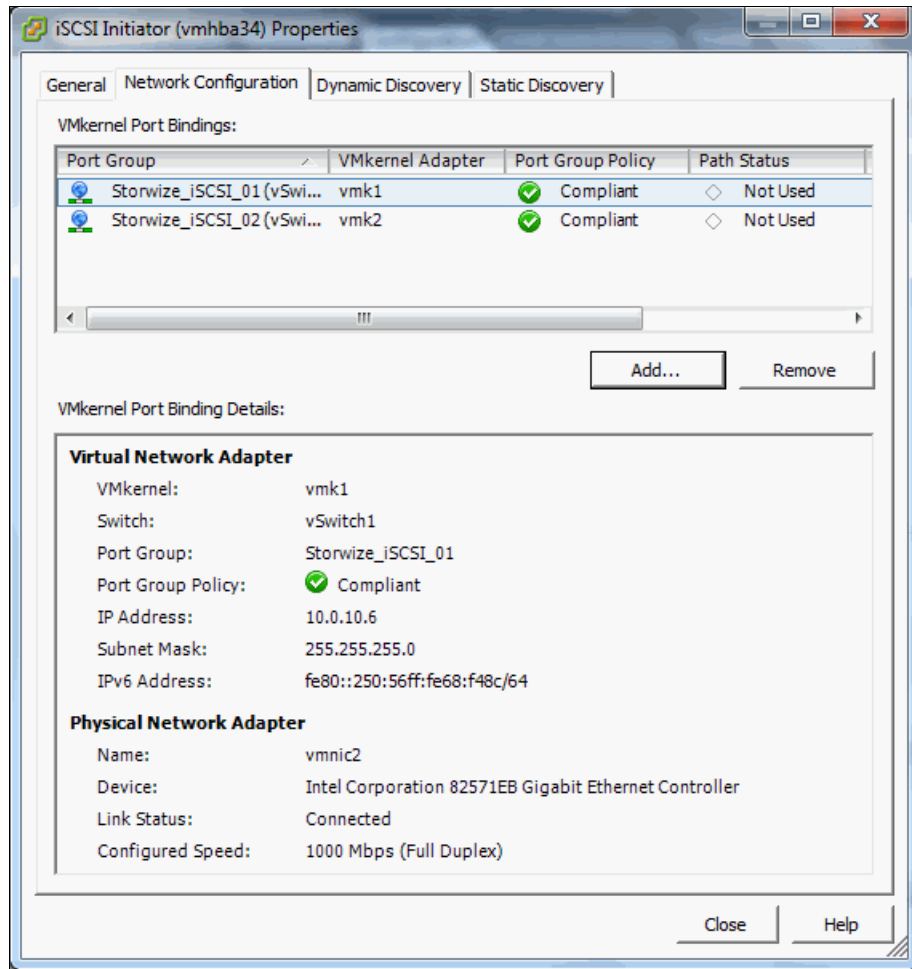


Figure 4-26 The VMkernel ports added to the iSCSI Adapter

16. The window shown in Figure 4-27 should open. Click **Yes** to rescan the adapter and complete this procedure.




Figure 4-27 Rescan the iSCSI Software Adapter

Configuring multipath support for iSCSI on VMware ESXi

As explained in Chapter 4.2.4, “VMware ESXi 5.5: Preparing for FC attachment” on page 163, the ESXi server has its own multipathing software.

For iSCSI, there is some extra configuration required in the VMkernel port properties to enable path failover. Each VMkernel port must map to just one physical adapter port, which is not the default setting. To fix this, complete the following steps:

1. Browse to the **Configuration** tab and select **Networking**.
2. Click **Properties** next to the vSwitch that you configured for iSCSI to open the window shown in Figure 4-28.

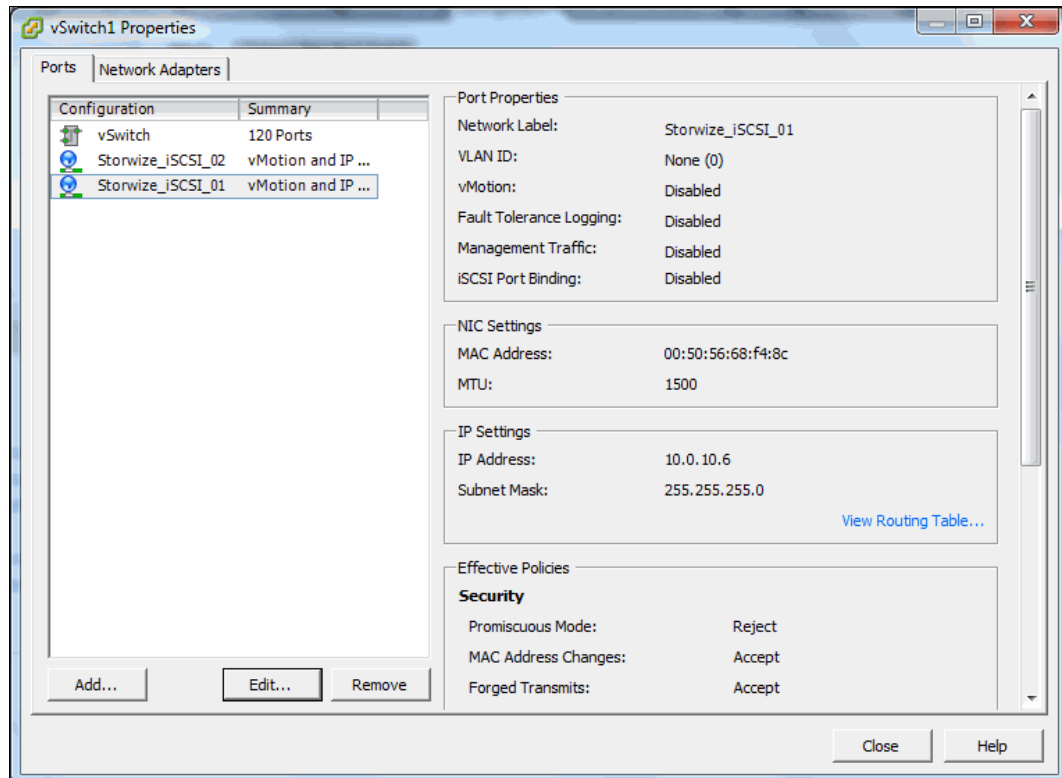


Figure 4-28 vSwitch Properties with VMkernel ports listed

3. Select one of the VMkernel ports and click **Edit**. The window shown in Figure 4-29 should open.

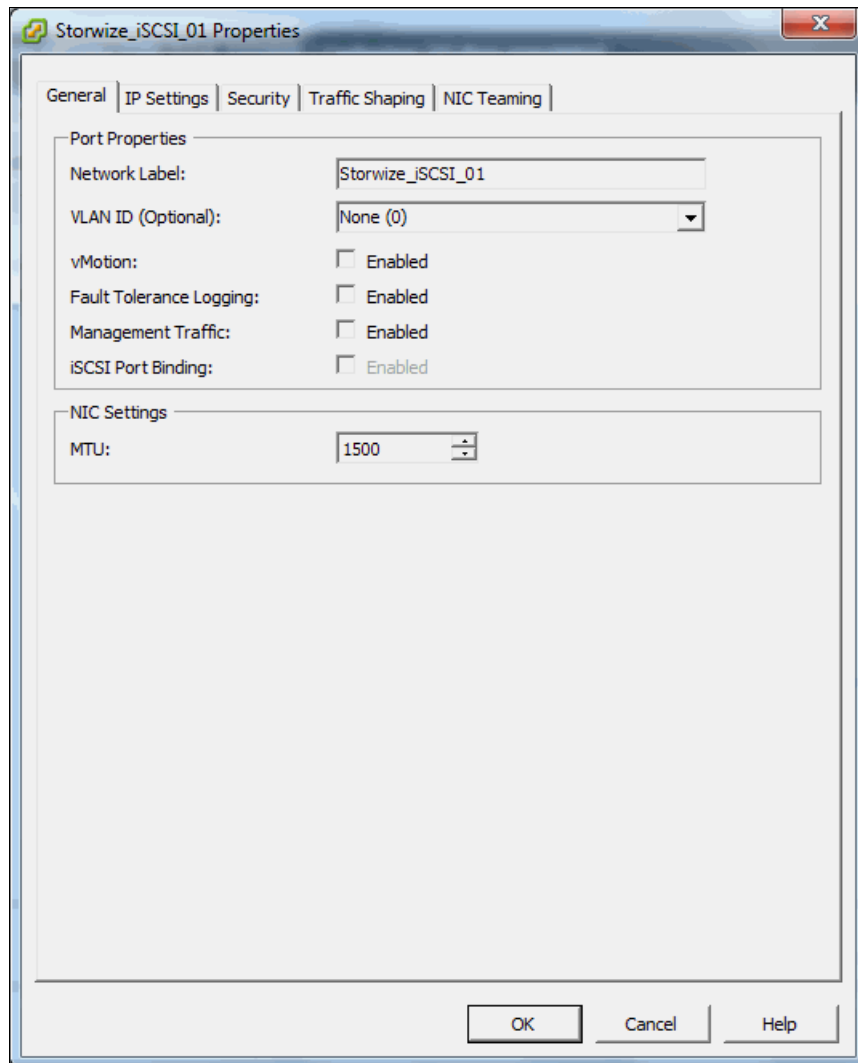


Figure 4-29 Editing a VMkernel port

4. Go to the NIC Teaming tab, select **Override switch failover order**, and then ensure that each port is tied to just one physical adapter port, as shown in Figure 4-30.

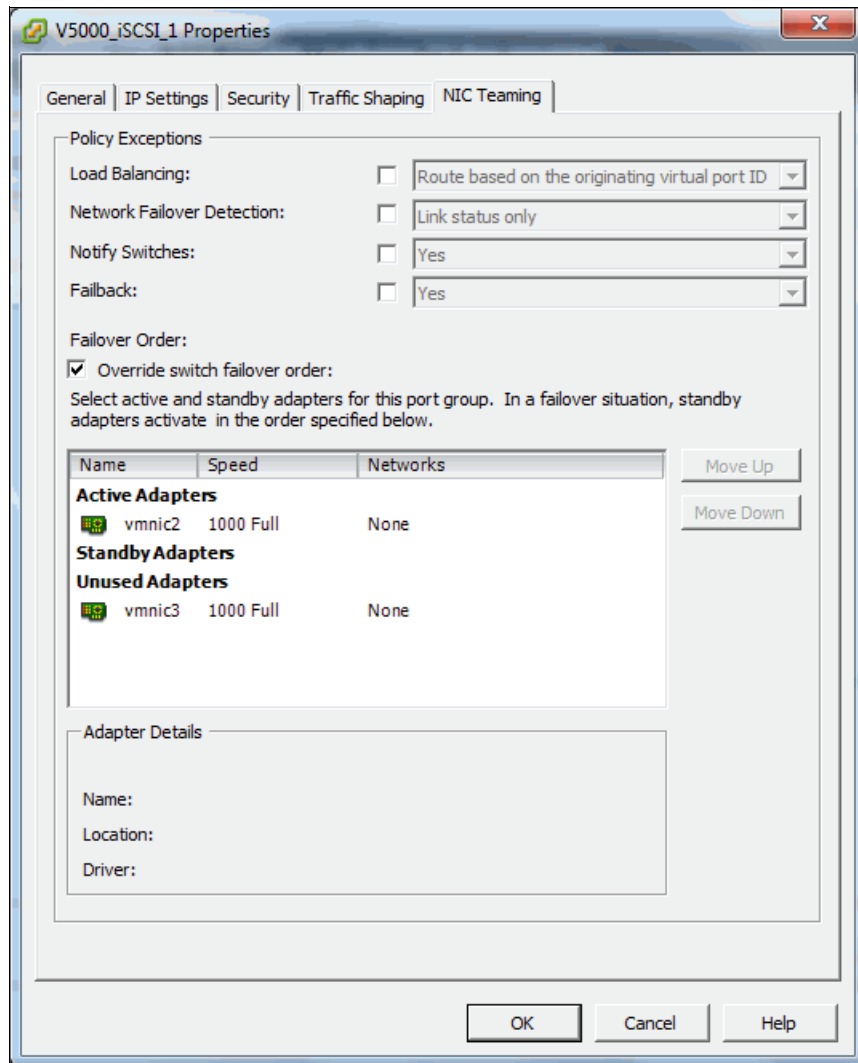


Figure 4-30 Configuring a VMkernel port to bind to a single physical adapter port

These are the basic steps required to prepare a VMware ESXi 5.5 host for iSCSI attachment. We provide more information later in this book.

For more information about configuring iSCSI attachment on the VMware ESXi side, the following white paper published by VMware is a useful resource:

<http://www.vmware.com/files/pdf/techpaper/vmware-multipathing-configuration-software-iSCSI-port-binding.pdf>

4.2.6 VMware ESXi 5.5: Preparing for SAS attachment

To prepare for SAS attachment, complete the following steps:

- ▶ Installing and updating supported HBAs
- ▶ Configuring SAS HBAs on VMware ESXi
- ▶ Confirming multipath support for SAS on VMware ESXi
- ▶ Determining host WWPNS

Installing and updating supported HBAs

Install a supported HBA with the current supported firmware and drivers for your configuration:

1. Determine a list of the current supported HBAs and levels for VMware ESXi 5.5, which are available on the following website:

<http://www.ibm.com/support/docview.wss?uid=ssg1S1004388>

2. Install the driver using VMware vSphere Client, the ESXi CLI, or vendor tools.
3. Also, check and update the firmware level of the HBA using the manufacturer's provided tools.
4. Always check the readme file to see whether there is further configuration required for the HBA driver.

See the VMware Compatibility Guide for more information:

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

For LSI cards, the following website provides some guidance on flashing the HBA firmware and BIOS:

http://mycusthelp.info/LSI/_cs/AnswerDetail.aspx?sSessionID=&aid=8352

Configuring SAS HBAs on VMware ESXi

In this example, we used an LSI 9207-8e card and did not have to configure HBA parameters beyond the default settings. It is advised to check the parameters using the HBA BIOS or vendor tools to confirm that they are suitable for your requirements.

Confirming multipath support for SAS on VMware ESXi

As with FC, we can use native ESXi multipathing for SAS attachment on VMware ESXi 5.5. See Chapter 4.2.4, "VMware ESXi 5.5: Preparing for FC attachment" on page 163 for more details.

Determining host WWPNS

The worldwide port names (WWPNs) of the SAS HBA are required to configure host attachment on the IBM Storwize V3700.

The host WWPNs are not directly available through VMware vSphere Client. However, you can find them using vendor tools or the HBA BIOS. The method described in Chapter 4.2.3, "Windows 2012 R2: Preparing for SAS attachment" on page 162 also works.

These are the basic steps required to prepare a VMware ESXi 5.5 host for SAS attachment. We provide more information later in this book.

For further information and guidance on attaching storage with VMware ESXi 5.5, the following document published by VMware is a useful resource:

<http://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-552-storage-guide.pdf>

4.3 Configuring hosts in IBM Storwize V3700

This section describes how to create FC, iSCSI, and SAS hosts using the IBM Storwize V3700 GUI. We assume that the hosts have been previously prepared, as described in Chapter 4.2, “Preparing the host operating system” on page 151. We also assume that FC addresses, SAS addresses, or iSCSI initiator names were recorded for host creation purposes.

Considerations when configuring hosts in IBM Storwize V3700

When creating a host object in IBM Storwize V3700, it is important to verify the configuration limits and restrictions, which are published on the following website:

<http://www.ibm.com/support/docview.wss?uid=ssg1S1004921>

Creating a host object using the GUI

To create a host object, complete the following steps:

1. From the IBM Storwize V3700 GUI, open the hosts view by navigating to **Hosts** → **Hosts**, as shown in Figure 4-31.




Figure 4-31 Opening the hosts view

This should load the page shown in Figure 4-32.




Figure 4-32 The Hosts view with no host objects configured

2. Click the **Add Host** button to open the Add Host wizard, as shown in Figure 4-33.

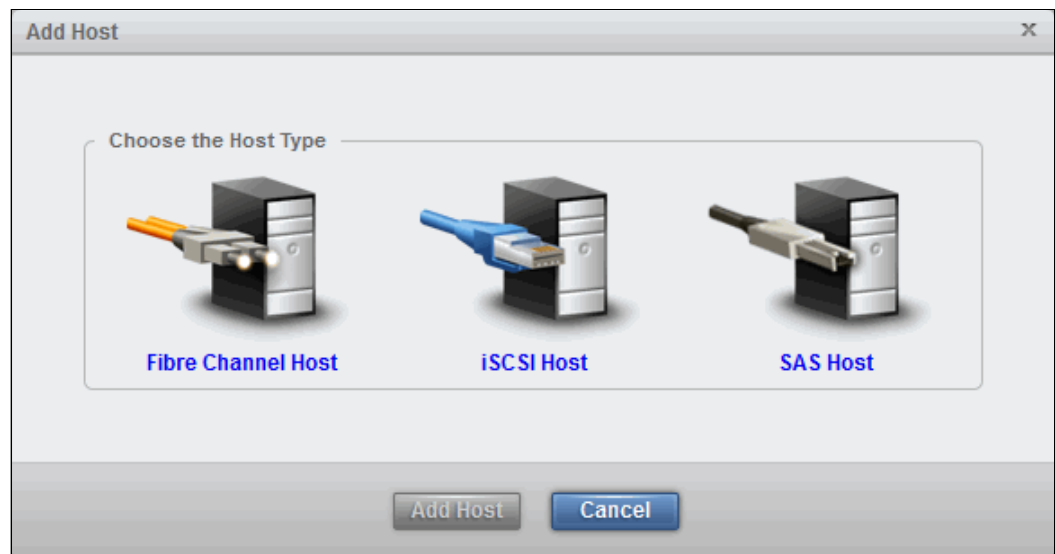


Figure 4-33 Choosing the host type in the Add Host wizard

3. Choose whether to create an FC host, an iSCSI host, or an SAS host, and follow the steps in the wizard. The following sections provide guidance for each attachment type.

4.3.1 Creating FC hosts

To create an FC host, complete the following steps:

1. Click **Fibre Channel Host**. The Fibre Channel host configuration wizard opens, as shown in Figure 4-34.

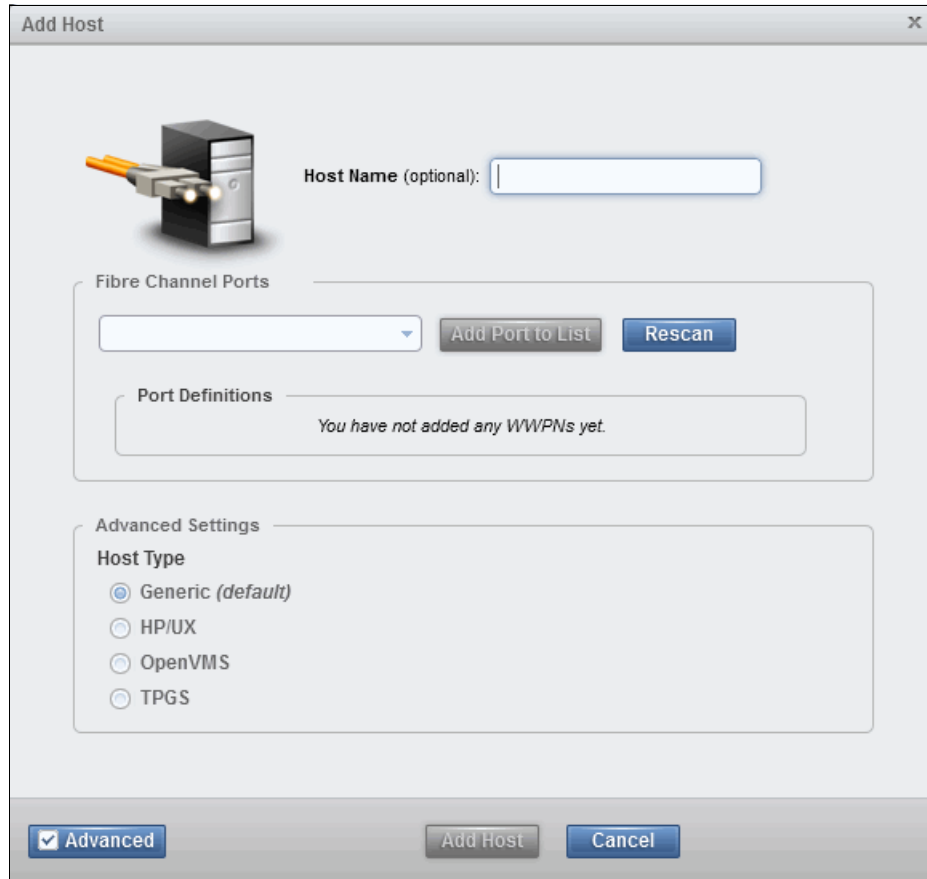


Figure 4-34 The FC host configuration wizard

2. Enter a descriptive **Host Name** and click the **Fibre Channel Ports** drop-down menu to view a list of all FC WWPNs visible to the system, as shown in Figure 4-35.

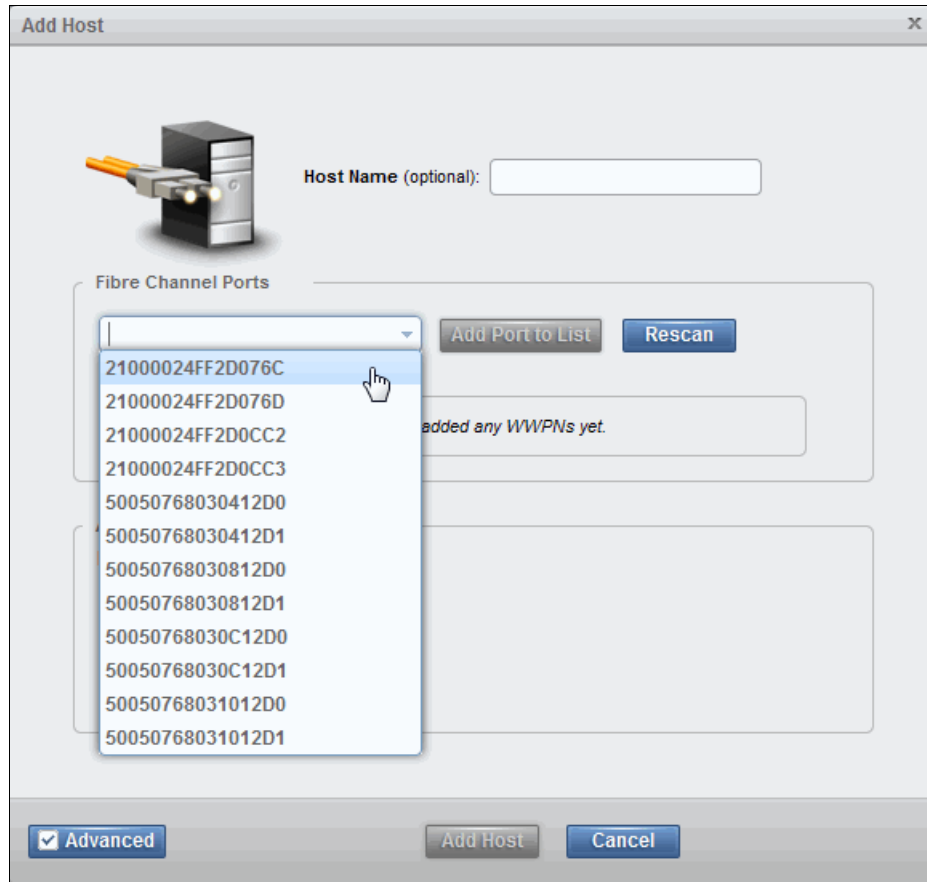


Figure 4-35 FC WWPNs visible to the system

3. If you have previously prepared an FC host, as described in Chapter 4.2, "Preparing the host operating system" on page 151, the WWPNs you recorded in this section should display.
4. If they do not display in the list, verify that you have completed all of the required steps and check your SAN zoning and cabling. Click **Rescan** in the configuration wizard.

Important: It is possible to enter WWPNs manually. However, if these are not visible to IBM Storwize V3700, the host object shows as offline and is unusable for I/O operations until the ports become visible.

5. Select the WWPNs for your host and click **Add Port to List** for each. These will be added to **Port Definitions**, as shown in Figure 4-36.

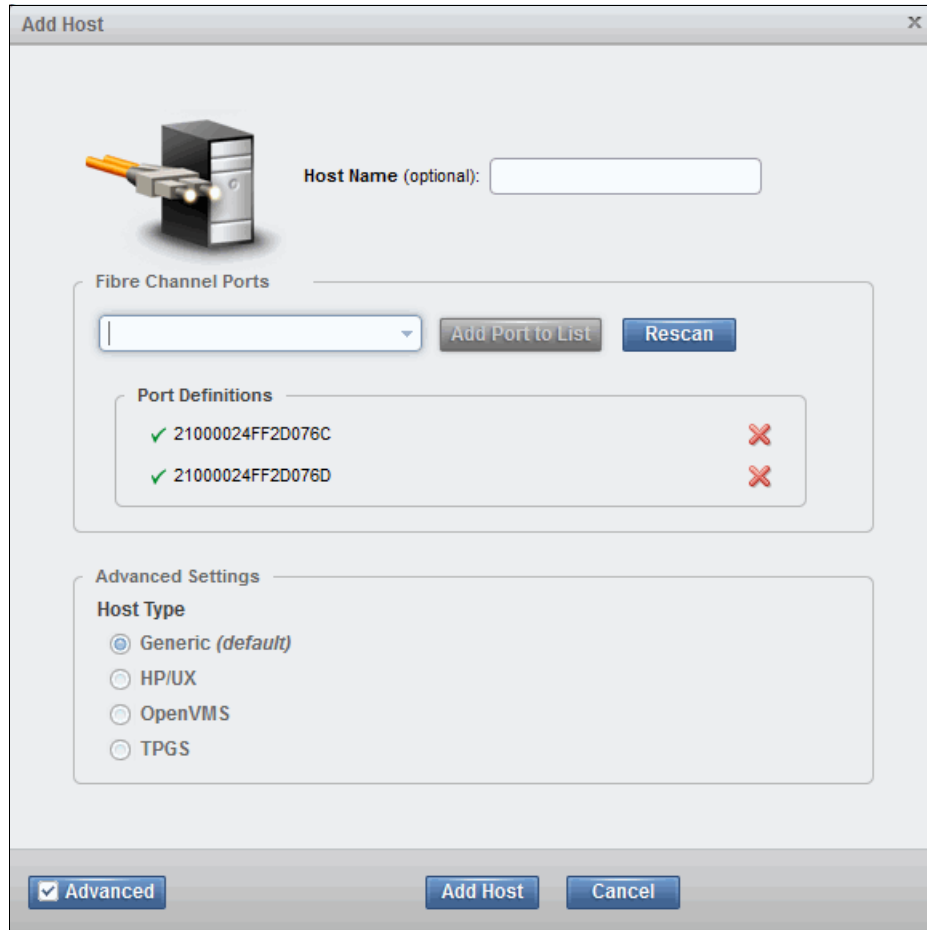


Figure 4-36 Adding the host WWPNs

6. In the **Advanced Settings**, it is possible to choose the Host Type. If using **HP/UX**, **OpenVMS**, or **TPGS**, this needs to be configured. Otherwise, the default option (**Generic**) is fine.
7. Click **Add Host** and the wizard completes, as shown in Figure 4-37.




Figure 4-37 Completion of the Add Host wizard

- Click **Close** to return to the host view, which should now list your newly created host object, as shown in Figure 4-38.




Figure 4-38 The Hosts view with the newly created host object listed

- Repeat these steps for all of the FC hosts.

After the host objects have been created, see Chapter 5, “Volume configuration” on page 199 to create volumes and map them to the hosts.

4.3.2 Configuring IBM Storwize V3700 for FC connectivity

It is possible to configure the FC ports on IBM Storwize V3700 to only be used for certain connections. This is sometimes referred to as *Port Masking*. In a system with multiple I/O groups and remote partnerships, this is a useful tool for ensuring peak performance and availability.

The following options are available per port:

- ▶ Any. Allow local and remote communication between nodes.
- ▶ Local. Allow only local node communication.
- ▶ Remote. Allow only remote node communication.
- ▶ None. Do not allow any node communication.

In all cases, host I/O is still permitted, so the **None** option could be used to exclusively reserve a port for host I/O.

There is a limit of **16** logins per node from another node before an error is logged. A combination of port masking and SAN zoning can help to manage this and support optimum I/O performance with local, remote, and host traffic.

To configure FC ports, complete the following steps:

1. Go to **Settings** → **Network**, as shown in Figure 4-39.




Figure 4-39 Opening the Settings > Network view

2. Select **Fibre Channel Ports** and the FC port configuration view should open, as shown in Figure 4-40.




Figure 4-40 The Fibre Channel Ports pane

- To configure a port, right-click the port, and select **Modify Connection**. This should open the window shown in Figure 4-41.

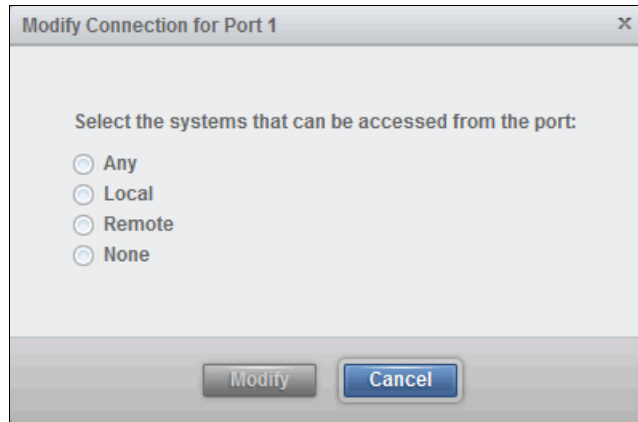


Figure 4-41 Modifying Connection for Port 1

- In this example, select **None** to restrict traffic on this port to host I/O only.
- Click **Modify** to confirm the selection.

Note: In doing this, we are configuring Port 1 for **all** nodes. It is not possible to configure FC ports on a per-node basis.

You can view connections between nodes, storage systems, and hosts by selecting **Fibre Channel Connectivity** while in the network settings view.

- Choose which connections you want to view and click **Show Results**, as shown in Figure 4-42.

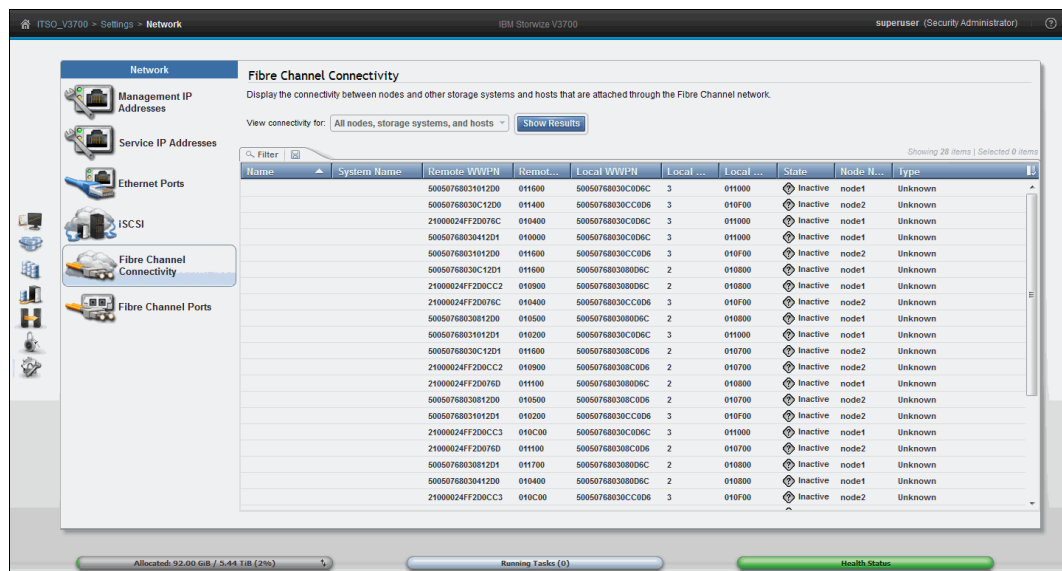


Figure 4-42 Viewing Fibre Channel Connectivity between nodes, storage systems, and hosts

4.3.3 Creating iSCSI hosts

To create an iSCSI host, complete the following steps:

1. Click **iSCSI Host** and the iSCSI host configuration wizard opens, as shown in Figure 4-43.

The screenshot shows the 'Add Host' configuration window. It includes a 'Host Name (optional)' text box, an 'iSCSI Ports' section with an input field and an 'Add Port to List' button, a 'Port Definitions' section with the message 'You have not added any iSCSI ports yet.', a checkbox for 'Use CHAP authentication (all ports)', and an 'Advanced Settings' section with a 'Host Type' dropdown menu. The 'Host Type' options are 'Generic (default)', 'HP/UX', 'OpenVMS', and 'TPGS'. At the bottom, there is a checked 'Advanced' checkbox and 'Add Host' and 'Cancel' buttons.

Figure 4-43 The iSCSI host configuration wizard

2. Enter a descriptive **Host Name** and enter the **iSCSI Initiator Name** into the iSCSI Ports box, then click **Add Ports to List**.

3. This displays under **Port Definitions**, as shown in Figure 4-44. Repeat this step if several initiator names are required.

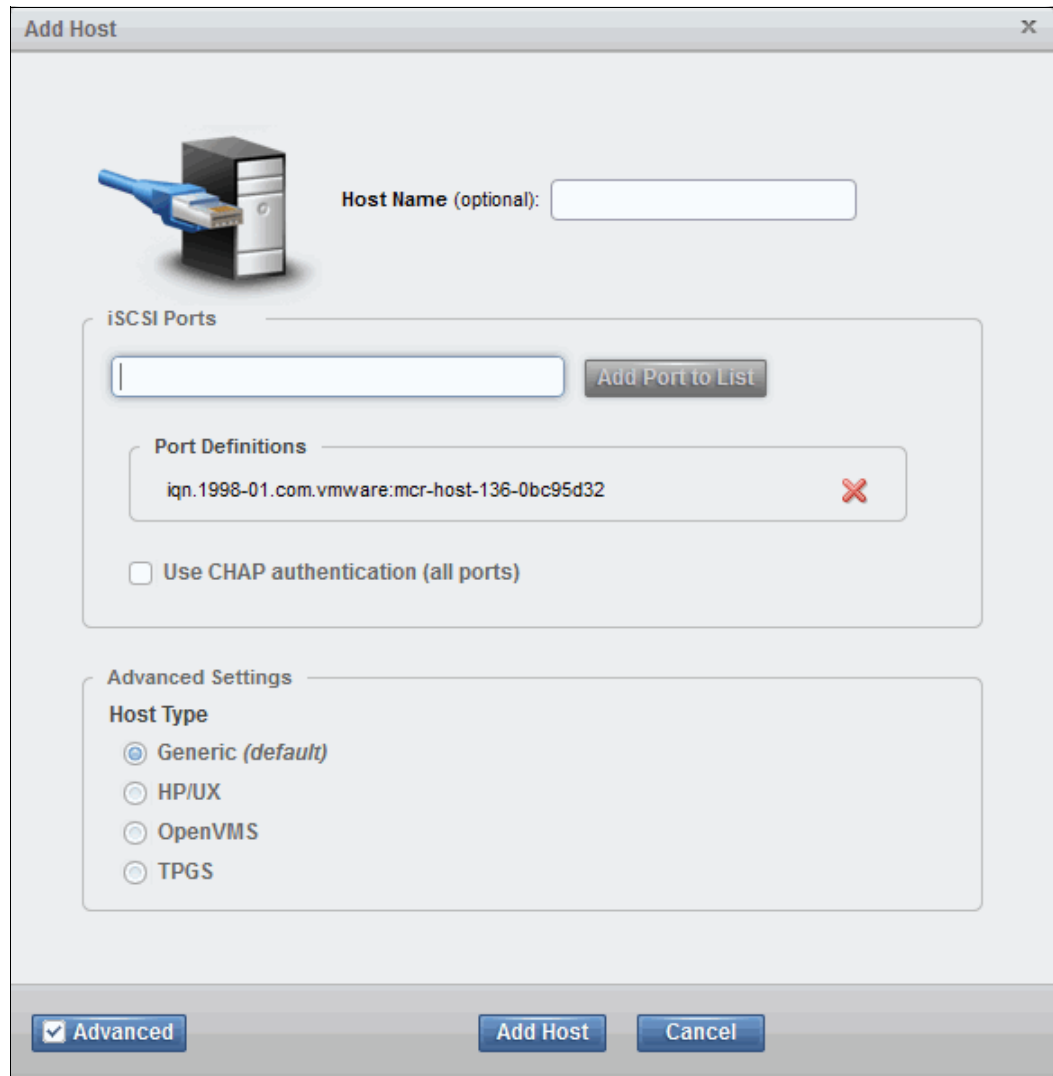


Figure 4-44 Adding the iSCSI initiator Host Name

4. In the **Advanced Settings**, it is possible to choose the Host Type. If using **HP/UX**, **OpenVMS**, or **TPGS**, this needs to be configured. Otherwise, the default option (**Generic**) is fine.

5. Click **Add Host** and the wizard completes, as shown in Figure 4-45.




Figure 4-45 Completion of the Add Host wizard

6. Click **Close** to return to the host view, which should now list your newly created host object, as shown in Figure 4-46.




Figure 4-46 The Hosts view with the newly created host object listed

7. Repeat these steps for all of your iSCSI hosts.

After the host objects have been created, see Chapter 5, “Volume configuration” on page 199 to create volumes and map them to the hosts.

Note: iSCSI hosts might have Degraded status until volumes are mapped. This is a limitation related to the implementation of iSCSI in IBM Storwize V3700, and not necessarily a sign of problems with network connectivity or the host configuration.

4.3.4 Configuring IBM Storwize V3700 for iSCSI host connectivity

The iSCSI hosts are now configured on the IBM Storwize V3700. To provide connectivity, the iSCSI Ethernet ports must also be configured. Complete the following steps to enable iSCSI connectivity:

1. Go to **Settings** → **Network**, as shown in Figure 4-47.




Figure 4-47 Opening the network settings view

2. Select **Ethernet Ports** and the Ethernet port configuration view should open, as shown in Figure 4-48.




Figure 4-48 The Ethernet Ports view

- To configure an IP address on a node port, expand the I/O Group, right-click the wanted port, and click **Modify IP Settings**, as shown in Figure 4-49.

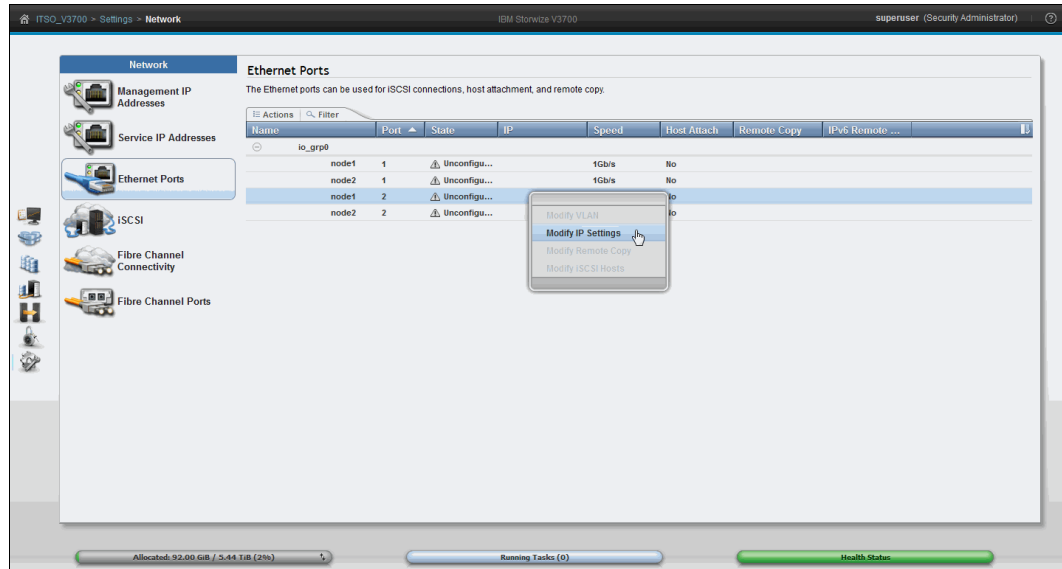


Figure 4-49 Opening the Modify IP Settings window for Node 1, Port 2

Note: It is generally advised that at least one port is reserved for the management IP address. This is normally Port 1, so in this example we configure Port 2 for iSCSI.

- The window shown in Figure 4-50 should open. To configure IPv6, click **IPv6** to reveal these options.
- Enter the wanted IP address, subnet mask, and gateway. It is important to make sure that these exist in the same subnet as the host IP addresses, and that the chosen address is not already in use. Click **Modify** to confirm.

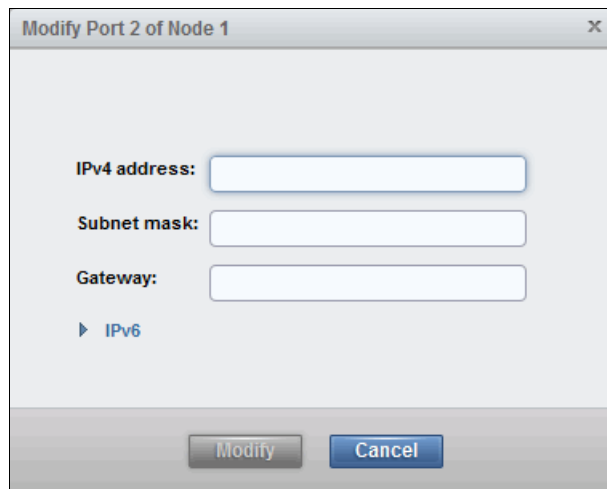


Figure 4-50 Configuring an IP address, subnet mask, and gateway for Port 2 of Node 1

6. The port should now be listed as Configured, as shown in Figure 4-51.




Figure 4-51 Showing configured port

7. Right-click the configured port again. This time, the options that were previously disabled should now be available. To confirm that the port is enabled for iSCSI, click **Modify iSCSI Hosts**.
8. The window shown in Figure 4-52 should open. If the port is not enabled, do so using the drop-down box and click **Modify** to confirm.

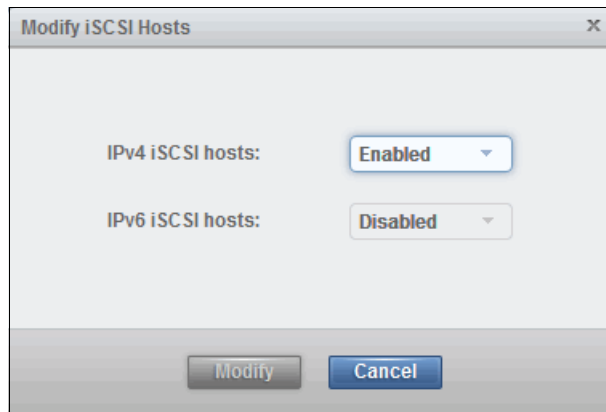


Figure 4-52 Enabling an Ethernet port for iSCSI host attachment

9. To modify VLAN settings, right-click the port again, and click **Modify VLAN**.

10. The window shown in Figure 4-53 should open. Check **Enable** to turn on VLAN tagging and set the tag in the box provided. It is advised that the failover port belong to the same VLAN, so leave this box selected. Click **Modify** to confirm.

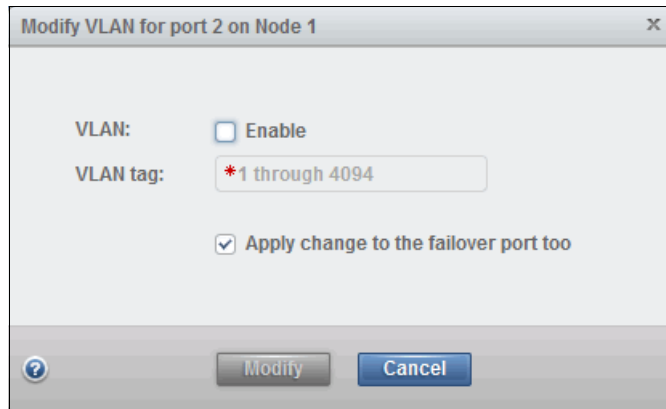


Figure 4-53 Modifying VLAN settings for Port 2 of Node 1

11. Repeat the previous steps for all ports that need to be configured.
12. It is also possible to configure iSCSI aliases, an iSNS name, and CHAP authentication. These options are in the iSCSI view. To access this, click **iSCSI** while in the network settings view, as shown in Figure 4-54.

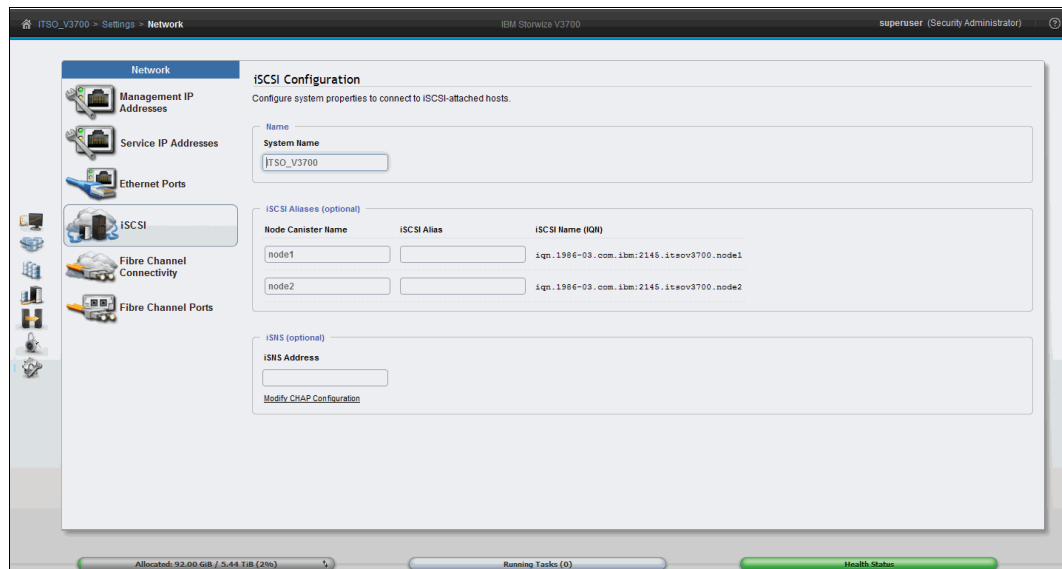


Figure 4-54 Advanced iSCSI configuration

4.3.5 Creating SAS hosts

To create an SAS host, complete the following steps:

1. Click **SAS Host**. The SAS host configuration wizard opens, as shown in Figure 4-55.

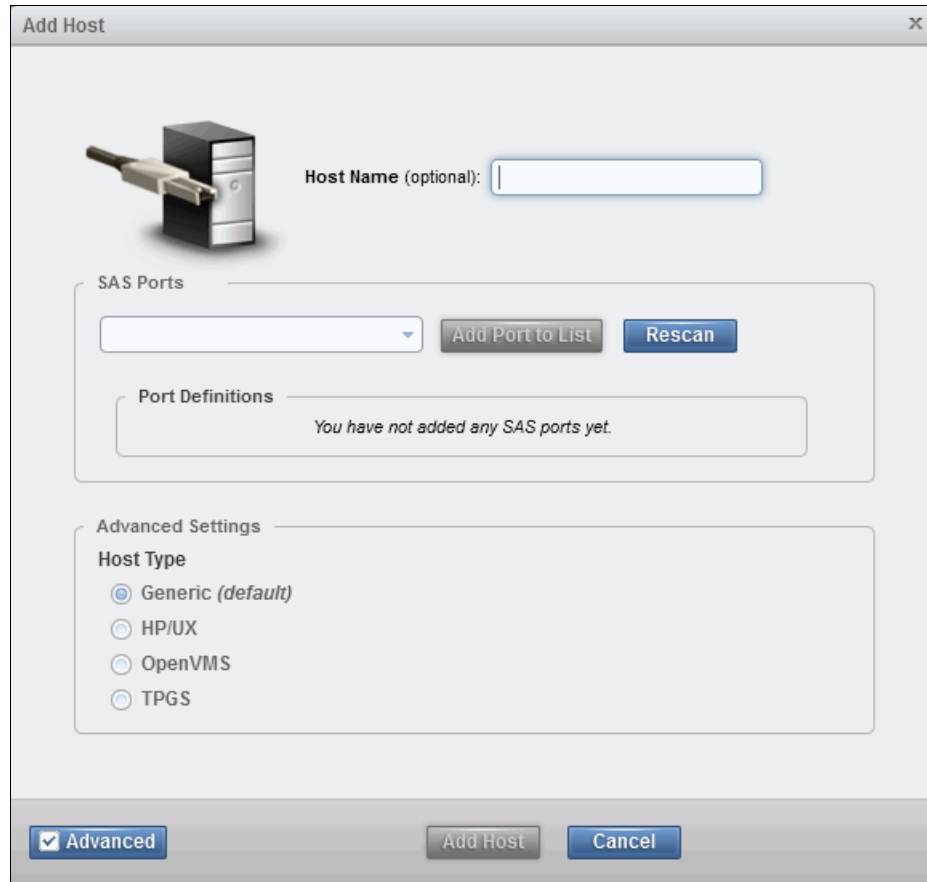


Figure 4-55 The SAS host configuration wizard

2. Enter a descriptive **Host Name** and click the **SAS Ports** drop-down menu to view a list of all FC WWPNs visible to the system, as shown in Figure 4-56.

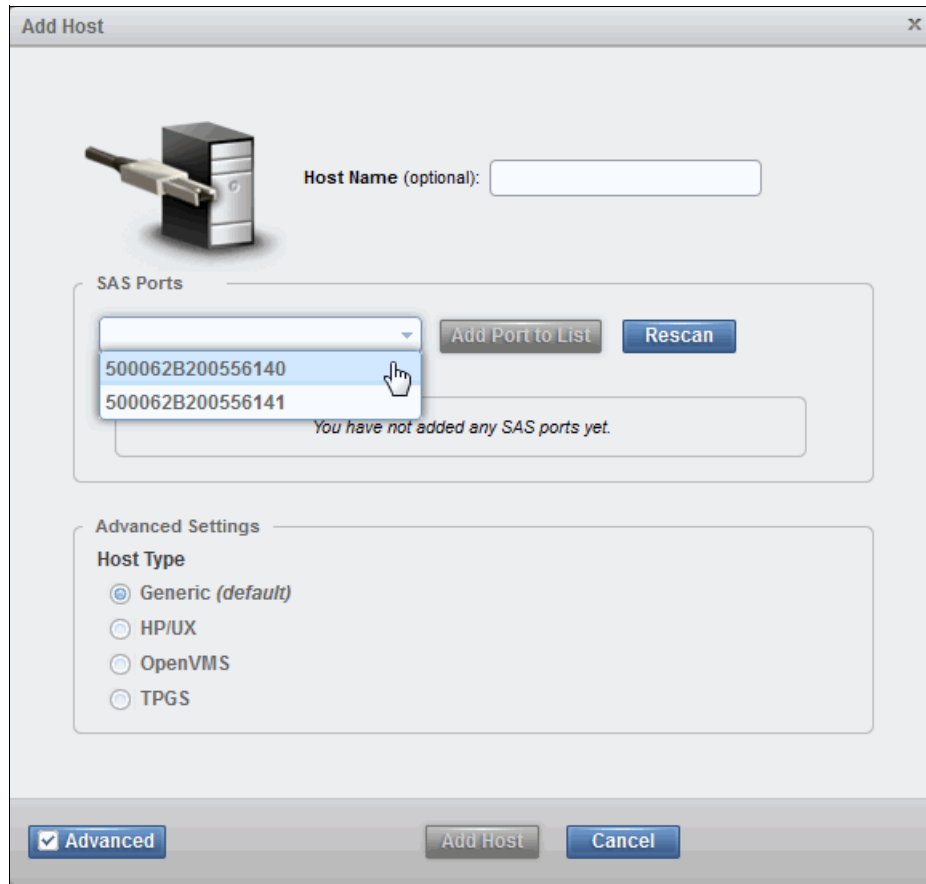


Figure 4-56 SAS WWPNs visible to the system

3. If you have previously prepared a SAS host, as described in Chapter 4.2, “Preparing the host operating system” on page 151, the WWPNs that you recorded in this section should display.
4. If they do not display in the list, verify that you have completed all of the required steps and check your cabling, then click **Rescan** in the configuration wizard. In particular, make sure that the ends of the SAS cables are aligned properly.

Important: It is possible to enter WWPNs manually. However, if these are not visible to IBM Storwize V3700, the host object shows as offline and is unusable for I/O operations until the ports become visible.

5. Select the WWPNs for your host, and click **Add Port to List** for each. These are added to **Port Definitions**, as shown in Figure 4-57.

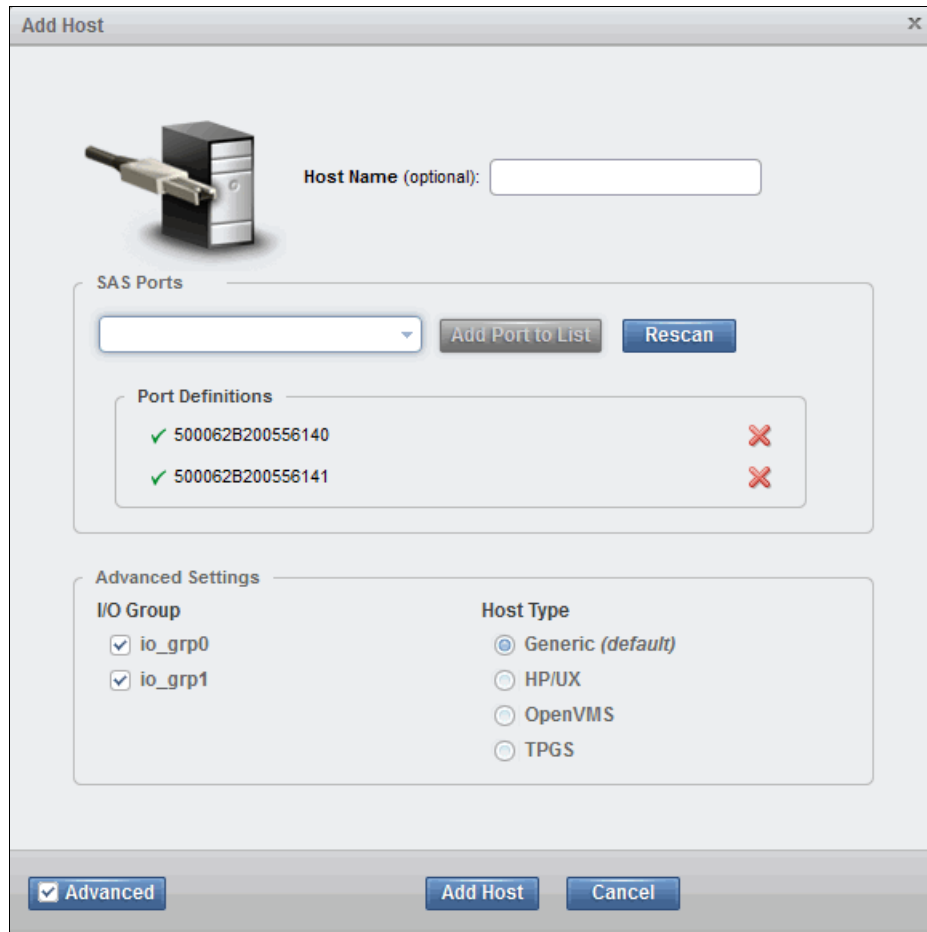


Figure 4-57 Adding the host WWPNs

6. In the **Advanced Settings**, it is possible to choose the Host Type. If using **HP/UX**, **OpenVMS**, or **TPGS**, this needs to be configured. Otherwise, the default option (**Generic**) is fine.
7. Click **Add Host** and the wizard completes, as shown in Figure 4-58.




Figure 4-58 Completion of the Add Host wizard

8. Click **Close** to return to the host view, which should now list your newly created host object, as shown in Figure 4-59.




Figure 4-59 The Hosts view with the newly created host object listed

9. Repeat these steps for all of your SAS hosts.
After the host objects have been created, see Chapter 5, “Volume configuration” on page 199 to create volumes and map them to the hosts.



Volume configuration

This chapter describes how to use the IBM Storwize V3700 to create and map volumes. A volume is a logical disk on the IBM Storwize V3700 that is provisioned out of a Storage Pool. Mappings are defined in IBM Storwize V3700 between volumes and host objects, at which point the storage is ready to be used by the host operating systems (OSs).

For information about preparing host OSs and creating host objects in IBM Storwize V3700, see Chapter 4, “Host configuration” on page 149.

The first part of this chapter describes how to create volumes of different types, and map them to host objects in IBM Storwize V3700. The second part of this chapter describes how to discover these volumes from the host OS.

After completing this chapter, your basic configuration is complete, and you can begin to store and access data on the IBM Storwize V3700.

For information about advanced volume administration, such as reconfiguring settings for existing volumes, see Chapter 8, “Advanced host and volume administration” on page 359.

This chapter includes the following topics:

- ▶ Creating volumes in IBM Storwize V3700
- ▶ Mapping volumes to hosts
- ▶ Discovering mapped volumes from host systems

5.1 Creating volumes in IBM Storwize V3700

This section describes how to create volumes using the IBM Storwize V3700 graphical user interface (GUI). We assume that you have completed the steps described in Chapter 2, “Initial configuration” on page 27, and have storage arrays and pools already configured.

To create new volumes using the IBM Storwize V3700 GUI, complete the following steps:

1. Open the Volumes view by navigating to **Volumes** → **Volumes**, as shown in Figure 5-1.




Figure 5-1 Opening the Volumes view

This loads the page shown in Figure 5-2.




Figure 5-2 The Volumes view with no volumes configured

2. Click **Create Volumes** to open the Create Volumes wizard, as shown in Figure 5-3 on page 201.




Figure 5-3 Selecting a volume preset and pool in the Create Volume wizard

3. Choose whether to create Generic, Thin Provision, Mirror, or Thin Mirror volumes, and which pool to provision this storage from.

By default, volumes created with the Storwize V3700 GUI are striped across all available managed disks (MDisks) in the chosen Storage Pool. The following section provides descriptions of the volume presets, and guidance for creating each volume type:

- ▶ **Generic.** A striped volume that is fully provisioned, meaning that the volume capacity reflects the same size physical disk capacity.

Creating generic volumes is described in 5.1.1, “Creating generic volumes” on page 202.

- ▶ **Thin Provision.** A striped volume that reports a greater capacity than is initially allocated. Thin provisioned volumes are sometimes referred to as *space-efficient*.

Such volumes have two capacities:

- The *real* capacity controls the quantity of extents that is initially allocated to the volume.
- The *virtual* capacity is reported to all of the other components (for example, IBM FlashCopy), and to the hosts.

Thin provisioning is useful for applications where the amount of physical disk space actually used is often much smaller than the amount allocated. Thin-provisioned volumes can be configured to expand automatically when more physical capacity is used. It is also possible to configure the amount of space that is initially allocated. IBM Storwize V3700 includes a configurable alert, which alerts system administrators when thin-provisioned volumes begin to reach their capacity.

Creating thin-provisioned volumes is described in 5.1.2, “Creating thin-provisioned volumes” on page 205.

- ▶ **Mirror.** A striped volume that consists of two striped copies and is synchronized to protect against loss of data if the underlying Storage Pool of one copy is lost.

In this case, a single volume is presented to the host. However, two copies exist on the storage back end, usually in different Storage Pools (all reads are handled by the primary copy). This feature is similar to host-based software mirroring, but provides a single point of management and high availability (HA) to OSs that do not support software mirroring.

Creating mirror copies in different Storage Pools protects against array failures. However, the mirroring feature is *not* a complete disaster recovery solution, because both copies are accessed by the same node pair and are addressable only by a single cluster. If you want to create a robust disaster recovery solution using IBM Storwize V3700, see 10.2, “Remote Copy” on page 516.

Creating mirror volumes is described in 5.1.3, “Creating mirrored volumes” on page 208.

- ▶ **Thin Mirror.** A mirror volume that is also thin provisioned.

Creating thin mirror volumes is described in 5.1.4, “Creating thin-mirrored volumes” on page 211.

5.1.1 Creating generic volumes

To create generic volumes, complete the following steps:

1. Select **Generic** as the preset, as shown in Figure 5-4.
2. From the list shown in Figure 5-4, select a pool to provision storage from.
3. The view shown in Figure 5-4 opens. From here, you can choose how many volumes to create, the capacity of each volume, and what to name the volumes.

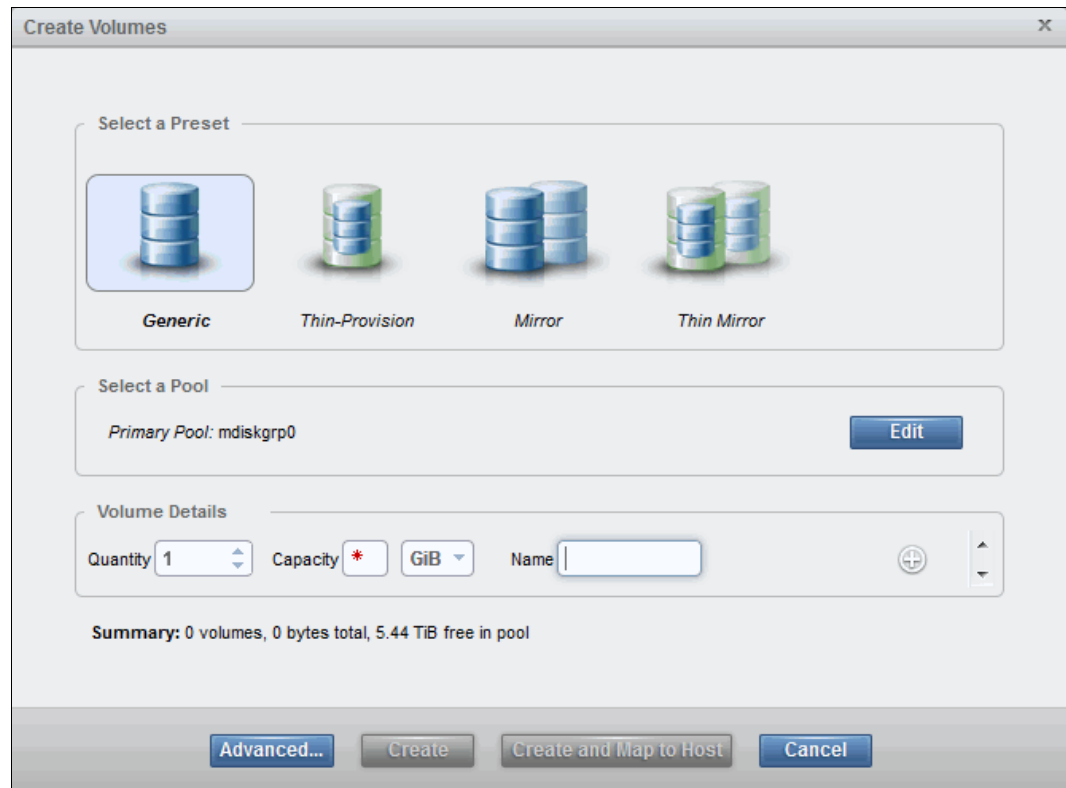


Figure 5-4 Creating generic volumes using the Create Volumes wizard

4. Click **Advanced** to view more options, as shown in Figure 5-5.

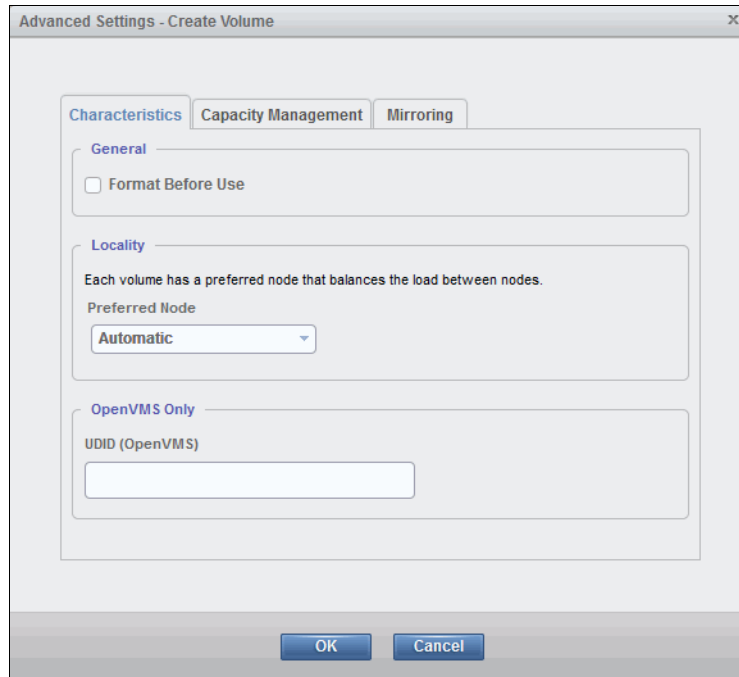


Figure 5-5 Advanced Settings for volume creation

Some important points to note:

- For generic volumes, capacity management and mirroring do not apply.
 - There is an option to set the preferred node within the input/output (I/O) Group that owns the volume. The suggested option here is to set this to **Automatic** and allow the system to balance volume I/O across both node canisters in the I/O Group.
5. When you are done, click **OK** to return to the main wizard.
 6. Click **Create** or **Create and Map to Host** when you are finished. **Create and Map to Host** creates the volumes and then takes you to the Host Mapping wizard, which is described in 5.2, “Mapping volumes to hosts” on page 213.

Important: The **Create and Map to Host** option is disabled if no hosts are configured on the IBM Storwize V3700. For more information about configuring hosts, see Chapter 4, “Host configuration” on page 149.

7. If you clicked **Create**, the wizard completes, as shown in Figure 5-6.




Figure 5-6 Completion of the Create Volumes wizard

8. Click **Close** to return to the Volumes view, which now lists your newly created volume, as shown in Figure 5-7.




Figure 5-7 The Volumes view with the newly created volume listed

9. Repeat these steps for all generic volumes you want to create.

If you did not choose **Create and Map to Host** in Step 5, see 5.2.2, “Manually mapping volumes to hosts” on page 215 for instructions about mapping volumes manually after host creation.

5.1.2 Creating thin-provisioned volumes

To create thin-provisioned volumes, complete the following steps:

1. Select **Thin Provision** as the preset, as shown in Figure 5-8.
2. From the list shown in Figure 5-8, select a pool to provision storage from.
3. The view shown in Figure 5-8 opens. From here, you can choose how many volumes to create, the capacity of each volume, and what to name the volumes.

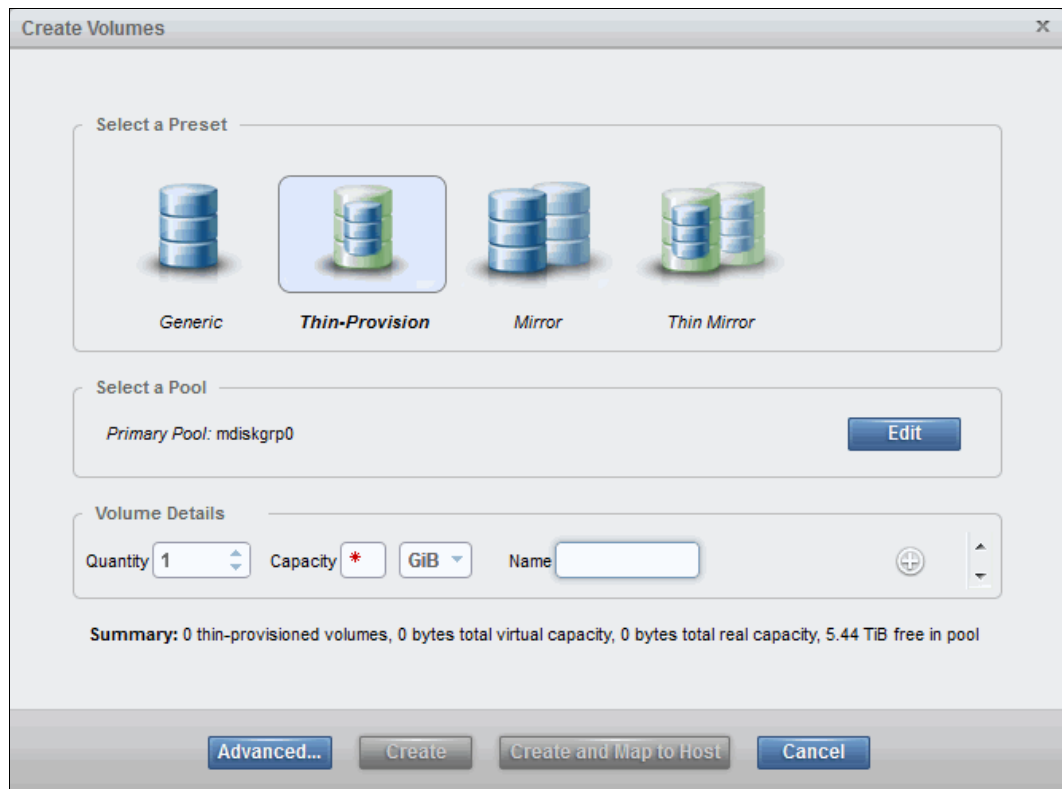


Figure 5-8 Creating thin provisioned volumes using the Create Volumes wizard

Note: The capacity specified is the *virtual* capacity of the volume, where the real capacity is 2% of this value. The real capacity can be modified in the advanced settings.

4. Click **Advanced** to view more options. The **Characteristics** tab is the same as for generic volumes. However, there are more options under **Capacity Management**, as shown in Figure 5-9.

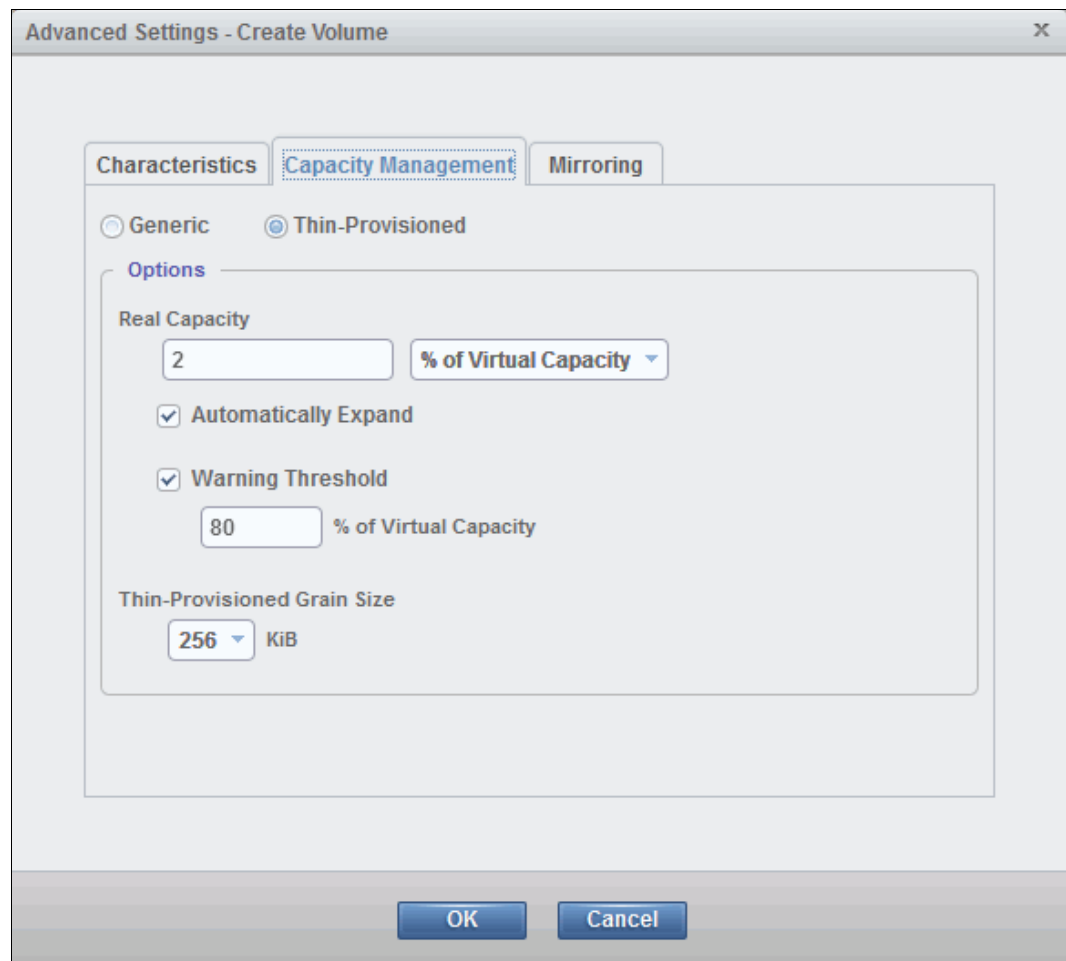


Figure 5-9 Advanced Capacity Management options for Thin-Provisioned volumes

Some important points to note:

- For thin-provisioned volumes, mirroring does not apply.
- The *real capacity* is the space that is actually allocated during creation. This can be specified as a percentage of the virtual capacity, or as a value in gibibytes (GiB). Increasing the real capacity reserves more space for the volumes, ensuring that it does not get used elsewhere. Increasing the real capacity diminishes the advantage of thin-provisioned volumes compared to generic volumes, however, because there is less virtual capacity left to work with and more physical capacity allocated.
- Enabling *automatic expansion* enables the real capacity to grow automatically as the amount of actual data stored on the volumes increases.
- The *warning threshold* is the point at which the system begins to log capacity alerts for the volumes. This gives system administrators advance warning to plan for expanding the real capacity of volumes. If a thin-provisioned volume becomes over-allocated, the system takes the volume offline, so it is important to pay attention to the warnings.
- You can also specify the grain size for the real capacity used. This is relevant when creating FlashCopy mappings from the volumes. For more information about FlashCopy, see Chapter 10, “Copy services” on page 455.

- When you are done, click **OK** to return to the main wizard.
- Click **Create** or **Create and Map to Host** when you are finished. **Create and Map to Host** creates the volumes and then takes you to the Host Mapping wizard, which is described in 5.2, “Mapping volumes to hosts” on page 213.

Important: The **Create and Map to Host** option is disabled if no hosts are configured on the IBM Storwize V3700. For more information about configuring hosts, see Chapter 4, “Host configuration” on page 149.

- If you clicked **Create**, the wizard completes, as shown in Figure 5-10.




Figure 5-10 Completion of the Create Volumes wizard

- Click **Close** to return to the Volumes view, which now lists your newly created volume, as shown in Figure 5-11.




Figure 5-11 The Volumes view with the newly created volume listed

- Repeat these steps for all thin-provisioned volumes that you want to create.

If you did not choose **Create and Map to Host** in Step 5, see 5.2.2, “Manually mapping volumes to hosts” on page 215 for instructions about mapping volumes manually after host creation.

5.1.3 Creating mirrored volumes

To create mirrored volumes, complete the following steps:

1. Select **Mirror** as the preset, as shown in Figure 5-12.
2. From the list shown in Figure 5-12, for each mirrored copy, select a pool to provision storage from. To take full advantage of the redundancy offered by the mirroring feature, each copy should ideally be provisioned from separate pools.
3. The view shown in Figure 5-12 opens. From here, you can choose how many mirrored pairs to create, the capacity of each volume, and what to name the volumes.

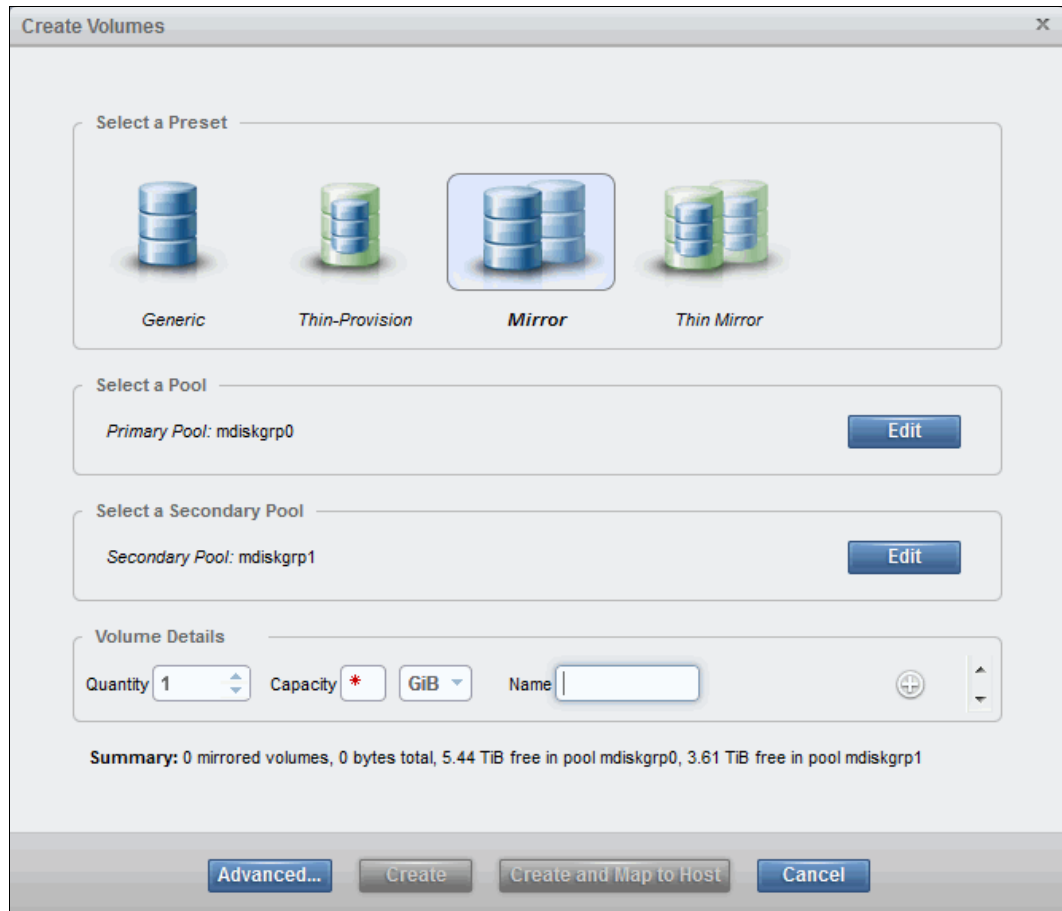


Figure 5-12 Creating mirrored volumes using the Create Volumes wizard

Note: The capacity specified is *per copy*, so the total capacity required is the number of mirrored pairs, multiplied by the capacity, *then multiplied by two*.

4. Click **Advanced** to view more options.

5. The **Characteristics** and **Capacity Management** tabs are the same as for generic volumes. However, there are more options under **Mirroring**, as shown in Figure 5-13.




Figure 5-13 Advanced Mirroring options for mirrored volumes

Here, it is possible to set the mirror sync rate. Higher sync rates provide greater availability at a small cost to performance, and vice versa.

6. When you are done, click **OK** to return to the main wizard.
7. Click **Create** or **Create and Map to Host** when you are finished. **Create and Map to Host** creates the volumes and then takes you to the Host Mapping wizard, which is described in 5.2, "Mapping volumes to hosts" on page 213.

Important: The **Create and Map to Host** option is disabled if no hosts are configured on the IBM Storwize V3700. For more information about configuring hosts, see Chapter 4, "Host configuration" on page 149.

8. If you clicked **Create**, the wizard completes, as shown in Figure 5-14.




Figure 5-14 Completion of the Create Volumes wizard

9. Click **Close** to return to the Volumes view, which now lists your newly created volumes, as shown in Figure 5-15.




Figure 5-15 The Volumes view with the newly created volume listed

10. Repeat these steps for all mirrored volumes that you want to create.

If you did not choose **Create and Map to Host** in Step 5, see 5.2.2, “Manually mapping volumes to hosts” on page 215 for instructions on mapping volumes manually after host creation.

5.1.4 Creating thin-mirrored volumes

To create thin-mirrored volumes, complete the following steps:

1. Select **Thin Mirror** as the preset, as shown in Figure 5-16.
2. From the list shown in Figure 5-16, for each mirrored copy, select a pool to provision storage from. To take full advantage of the redundancy offered by the mirroring feature, each copy should ideally be provisioned from separate pools.
3. The view shown in Figure 5-16 opens. From here, you can choose how many mirrored pairs to create, the capacity of each volume, and what to name the volumes.

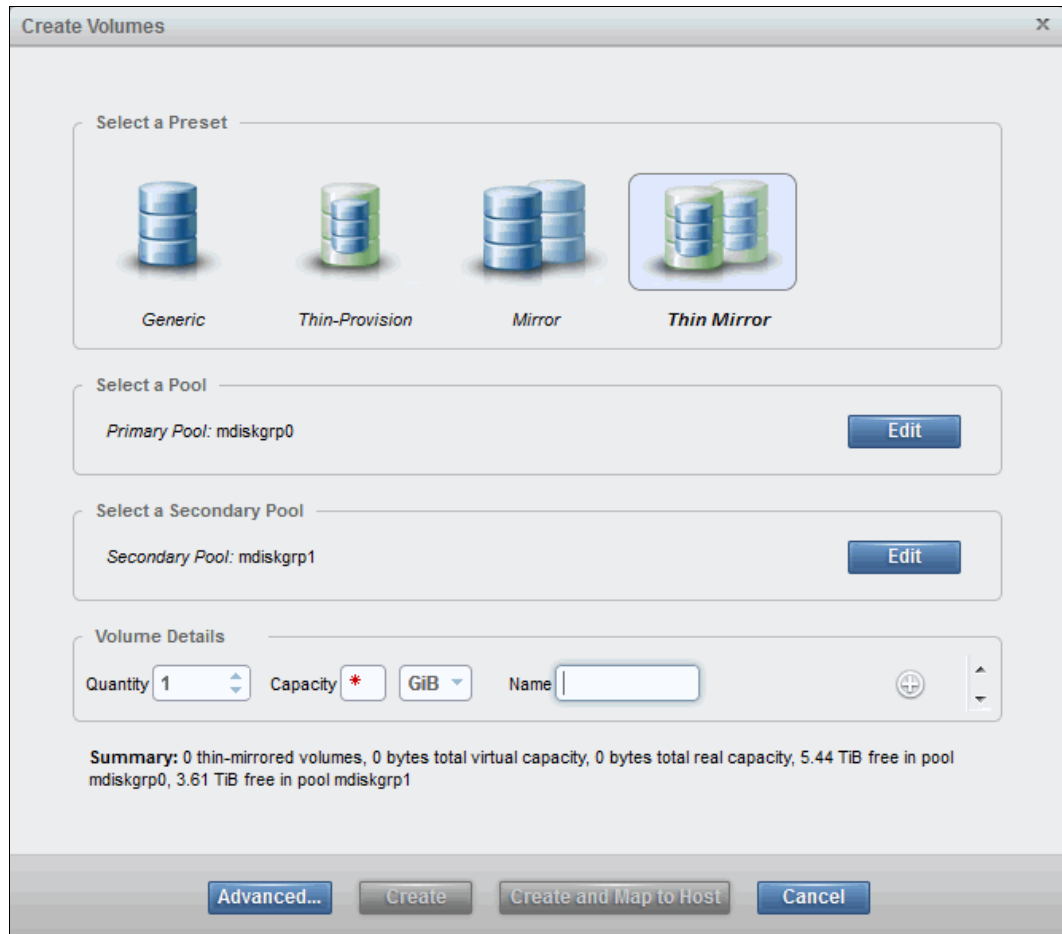


Figure 5-16 Creating thin mirrored volumes using the Create Volumes wizard

Note: The capacity specified is *per copy*, so the total virtual capacity is the number of mirrored pairs, multiplied by the capacity, *then multiplied by two*.

4. Click **Advanced** to view more options.
5. Here, you can configure **Characteristics** as for generic volumes, **Capacity Management** as for thin-provisioned volumes, and **Mirroring** as for mirrored volumes.

When you are done, click **OK** to return to the main wizard.

- Click **Create** or **Create and Map to Host** when you are finished. **Create and Map to Host** creates the volumes and then takes you to the Host Mapping wizard, which is described in 5.2, “Mapping volumes to hosts” on page 213.

Important: The Create and Map to Host option is disabled if no hosts are configured on the IBM Storwize V3700. For more information about configuring hosts, see Chapter 4, “Host configuration” on page 149.

- If you clicked **Create**, the wizard completes, as shown in Figure 5-17.




Figure 5-17 Completion of the Create Volumes wizard

- Click **Close** to return to the Volumes view, which now lists your newly created volumes, as shown in Figure 5-18.

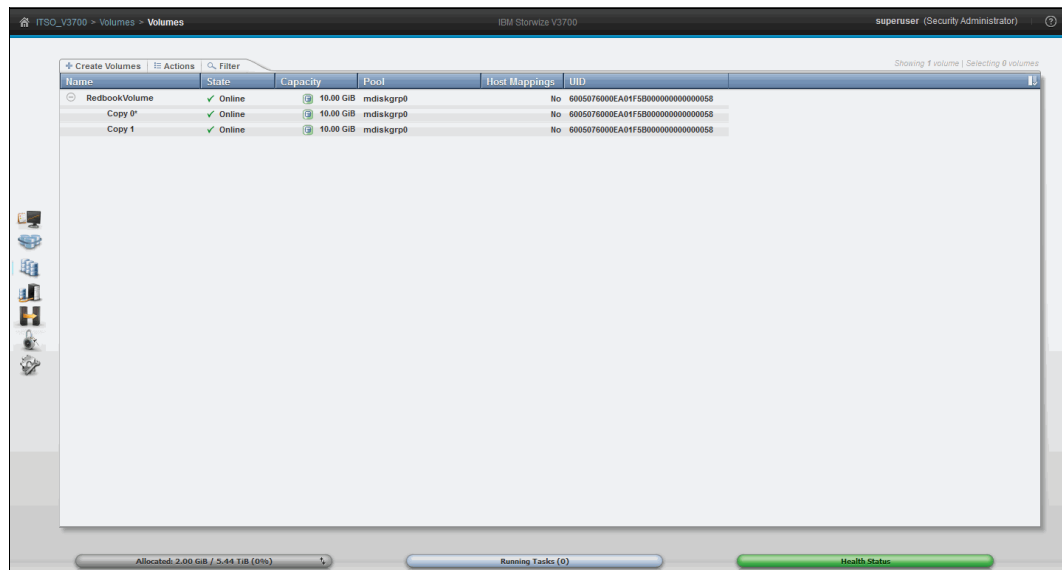


Figure 5-18 The Volumes view with the newly created volume listed

- Repeat these steps for all thin mirrored volumes that you want to create.

If you did not choose **Create and Map to Host** in Step 5, see 5.2.2, “Manually mapping volumes to hosts” on page 215 for instructions about mapping volumes manually after host creation.

5.2 Mapping volumes to hosts

This section describes how to create mappings between volumes and hosts on IBM Storwize V3700. When this task is complete, the mapped storage is available for use by host systems.

The first part of this section describes how to map volumes to hosts if you clicked **Create and Map to Host** in 5.1, “Creating volumes in IBM Storwize V3700” on page 200. The second part of this section describes the manual host mapping process that is used to create customized mappings.

5.2.1 Mapping newly created volumes using Create and Map to Host

In this section, we assume that you are following one of the procedures in 5.1, “Creating volumes in IBM Storwize V3700” on page 200 and clicked **Create and Map to Host**, followed by **Continue** when the Create Volumes task completed, as shown in Figure 5-19.




Figure 5-19 Continuing to map volumes after the Create Volumes task has completed

To map the volume, complete the following steps:

1. After clicking **Continue**, the window shown in Figure 5-20 opens.

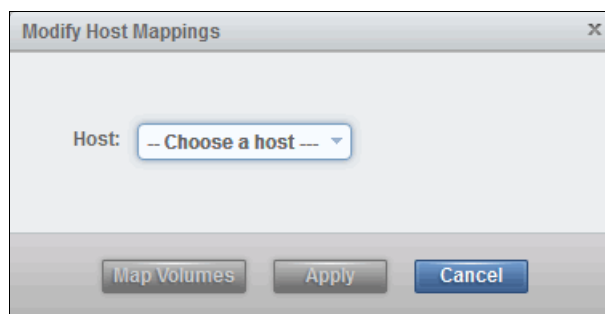


Figure 5-20 Selecting the I/O Group and host to assign the mappings to

- Select the host that you want to assign the mappings to. The window expands to the view shown in Figure 5-21.

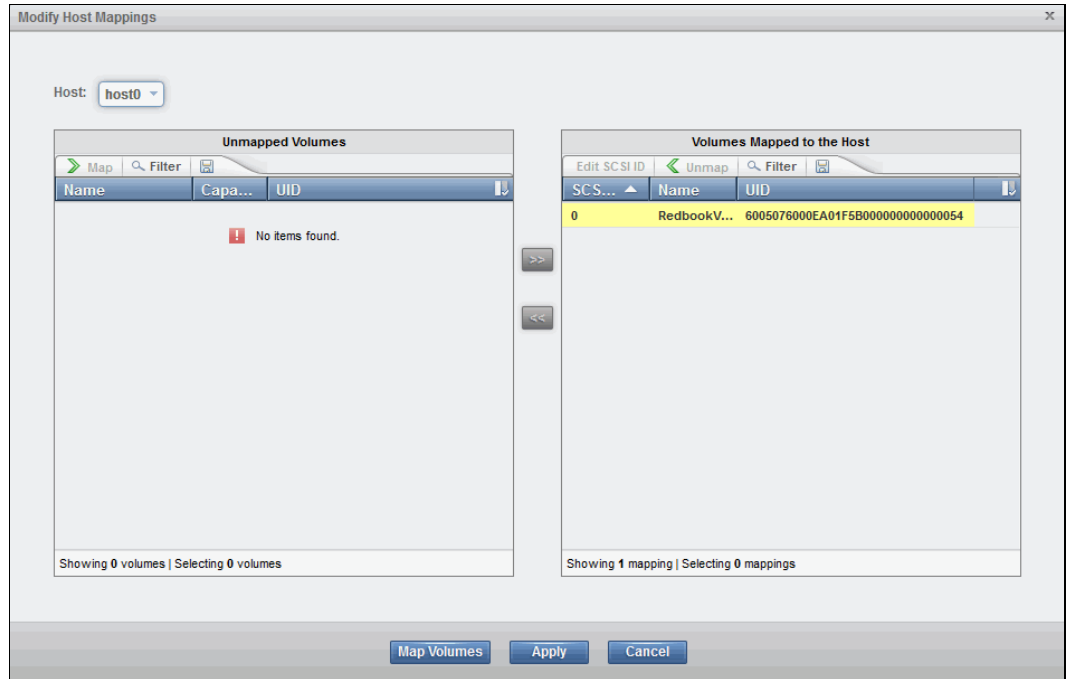


Figure 5-21 The Modify Host Mappings wizard showing Volume Mapped to the Host

- The newly created volume should already be mapped to the host. To confirm the mapping, click **Map Volumes** to complete the operation and close the window.
Alternatively, click **Apply** to complete the operation but leave the Host Mapping wizard open.
- After clicking **Map Volumes**, the wizard completes, as shown in Figure 5-22.




Figure 5-22 Completion of the Modify Mappings wizard

5. Click **Close** to return to the Volumes view, which now lists your newly created volume, as shown in Figure 5-23.




Figure 5-23 The Volumes view with the newly created and mapped volume listed

5.2.2 Manually mapping volumes to hosts

In this section, we assume that you have followed one of the procedures in 5.1, “Creating volumes in IBM Storwize V3700” on page 200 and clicked **Create**. This section describes how to manually map these volumes to hosts after volume creation.

To manually map volumes to hosts, complete the following steps:

1. Go to **Volumes** → **Volumes**, select the volumes that you want to map, and then click **Actions** → **Map to Host**, as shown in Figure 5-24.




Figure 5-24 Manually creating host mappings from the Volumes view

2. This opens the window shown in Step 1 of the procedure in 5.2.1, “Mapping newly created volumes using Create and Map to Host” on page 213. Complete the steps in that section to finish mapping the volumes.

The volumes should now be accessible from the host OS. The following section describes how to discover and use mapped volumes from hosts.

5.3 Discovering mapped volumes from host systems

This section describes how to discover the volumes that were created and mapped in 5.1, “Creating volumes in IBM Storwize V3700” on page 200 and 5.2, “Mapping volumes to hosts” on page 213, from the host operating system.

We assume that you have completed all of the following tasks described in this book so that the hosts and the IBM Storwize V3700 are prepared:

1. Prepare your operating systems for attachment, including installing multipath support. For more information, see 4.2, “Preparing the host operating system” on page 151.
2. Create host objects using the IBM Storwize V3700 GUI. For more information, see 4.3, “Configuring hosts in IBM Storwize V3700” on page 180.
3. Create volumes and map these to the host objects using the IBM Storwize V3700 GUI. For more information, see 5.1, “Creating volumes in IBM Storwize V3700” on page 200, and 5.2, “Mapping volumes to hosts” on page 213.

First, you need to look up the unique identifier (UIDs) of the volumes. Knowing the UIDs of the volumes helps determine which is which from the host point of view. To do this, complete the following steps:

1. Open the volumes by host view by going to **Volumes** → **Volumes by Host**, as shown in Figure 5-25.

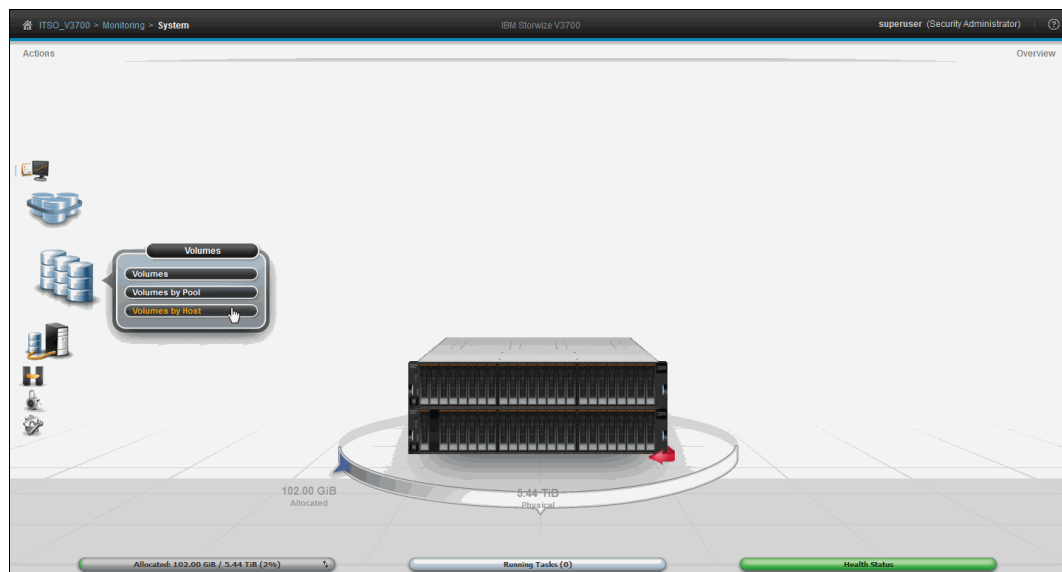


Figure 5-25 Opening the Volumes by Host view

This should load the page shown in Figure 5-26.

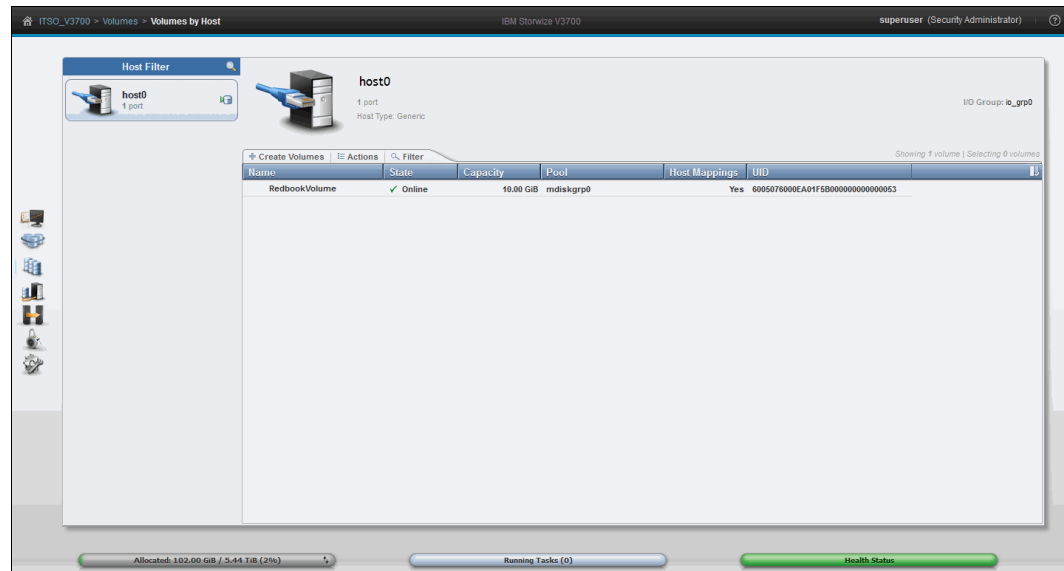


Figure 5-26 The volumes by host view with a single host and volume mapping configured

2. Examine the UID column and note the value for each mapped volume.

Next, you need to discover the mapped volumes from the host. The following sections describe how to discover mapped volumes from Windows 2012 R2 and VMware ESXi 5.5 for each supported attachment type, Fibre Channel (FC), Internet Small Computer System Interface (iSCSI), and serial-attached SCSI (SAS).

Note: For Windows 2012 R2, we describe the use of IBM Subsystem Device Driver Device Specific Module (SDDDSM) for multipath for FC and SAS, where 5.3.2, “Discovering iSCSI volumes from Windows 2012 R2” on page 227 describes how to configure native Microsoft Multipath I/O (MPIO). This is because only native MPIO is supported for iSCSI on Windows 2012 R2.

5.3.1 Discovering Fibre Channel volumes from Windows 2012 R2

This section describes how to discover FC volumes.

Discovering mapped volumes

To discover mapped FC volumes from Windows 2012 R2, complete the following steps:

1. Connect to your Windows server and open **Server Manager**.
2. Go to **File and Storage Services**, as shown in Figure 5-27.




Figure 5-27 Opening File and Storage Services in Server Manager

Note: File and Storage Services should be enabled by default. If not, they can be enabled by going to **Manage** → **Add Roles and Features** and following the wizard.

3. In File and Storage Services, go to **Disks**. Here, all storage devices known to Windows should be listed, as shown in Figure 5-28.




Figure 5-28 The Disks view in File and Storage Services

- If your mapped volumes are not shown, go to **Tasks** → **Rescan Storage**, as shown in Figure 5-29.




Figure 5-29 Rescan storage from File and Storage Services

- A bar displays at the top of the pane to show that Windows is attempting to scan for devices, as shown in Figure 5-30. Wait for this to complete.




Figure 5-30 Waiting for Windows to rescan storage

- The mapped volumes are now listed, as shown in Figure 5-31.




Figure 5-31 The mapped volumes listed in Windows 2012 R2

- You need to bring the disk online. To do this, right-click and select **Bring Online**, as shown in Figure 5-32.




Figure 5-32 Bring a disk online

- The window shown in Figure 5-33 opens. Click **Yes** to confirm that you want to bring the disk online on this server.




Figure 5-33 Confirm that you want to bring the disk online

- The next step is to initialize the disk. To do this, right-click the disk and click **Initialize**, as shown in Figure 5-34.




Figure 5-34 Initializing the disk

- The window shown in Figure 5-35 opens. Click **Yes** to confirm that you want to initialize the disk.

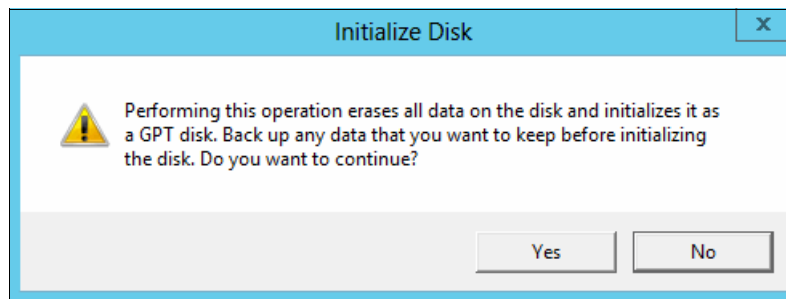


Figure 5-35 Confirm that you want to initialize the disk

- The disk now shows as online and initialized with Globally Unique Identifier (GUID) Partition Table (GPT) partitioning, as shown in Figure 5-36.




Figure 5-36 The mapped volume online and initialized as a GPT disk

- Repeat the previous steps for any other mapped volumes.

The mapped volumes are now ready to be managed by the host OS.

Creating a file system on a mapped volume

To create a file system on a mapped volume in Windows 2012 R2, complete the following steps:

- In the Disks view in **File and Storage Services**, right-click the disk and click **New Volume**, as shown in Figure 5-37.




Figure 5-37 Creating a new file system on a disk

Definition: In this context, *Volume* refers to a Windows concept, and not specifically an IBM Storwize V3700 volume.

2. The New Volume Wizard opens. Click **Next** on the first page to open the page shown in Figure 5-38.
3. Choose which server and disk to provision and click **Next**.

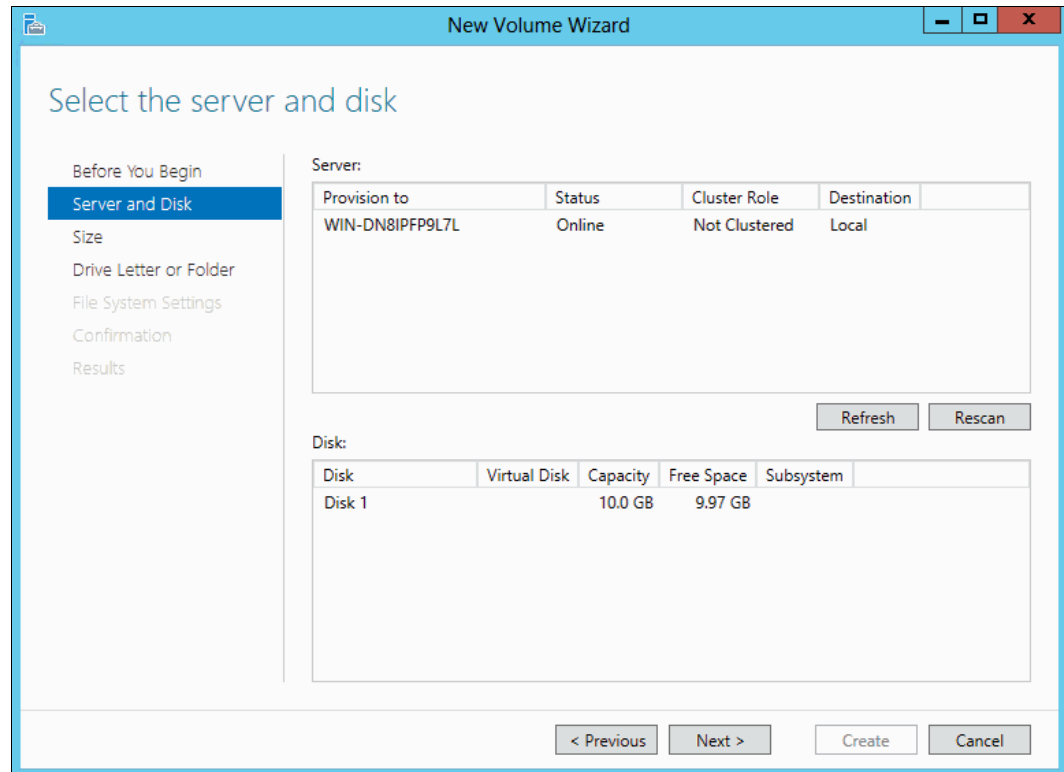


Figure 5-38 Choosing the Server and Disk

- The next step is to choose the size of the volume, as shown in Figure 5-39. Set the volume size and click **Next** to continue.

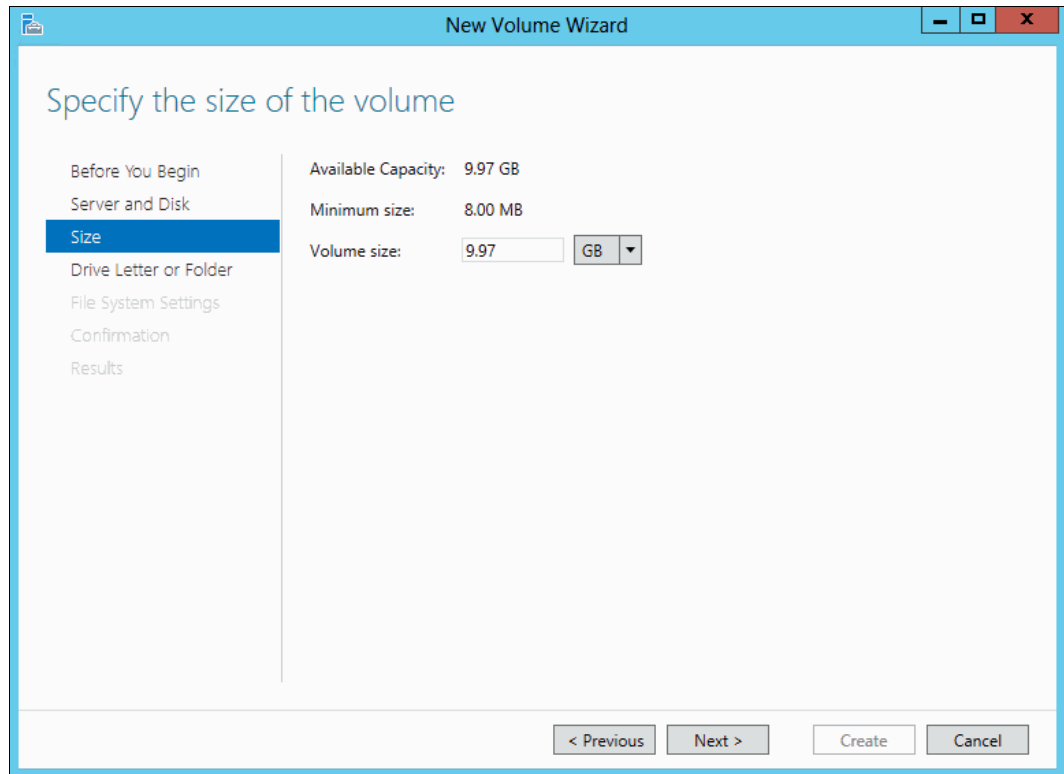


Figure 5-39 Choosing the volume size

- Next, choose whether to assign a Drive letter, as shown in Figure 5-40. Click **Next** to continue.

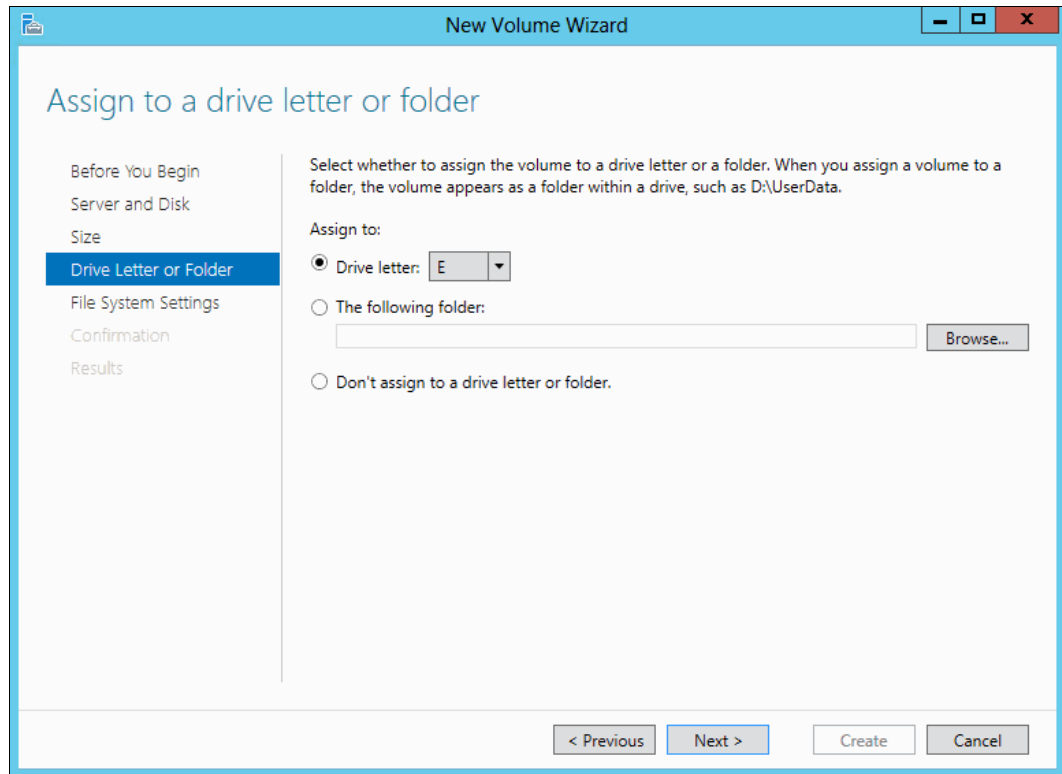


Figure 5-40 Assigning a Drive letter to the new volume

6. Choose a file system type, unit size, and name for the new volume, as shown in Figure 5-41. Click **Next** to continue.

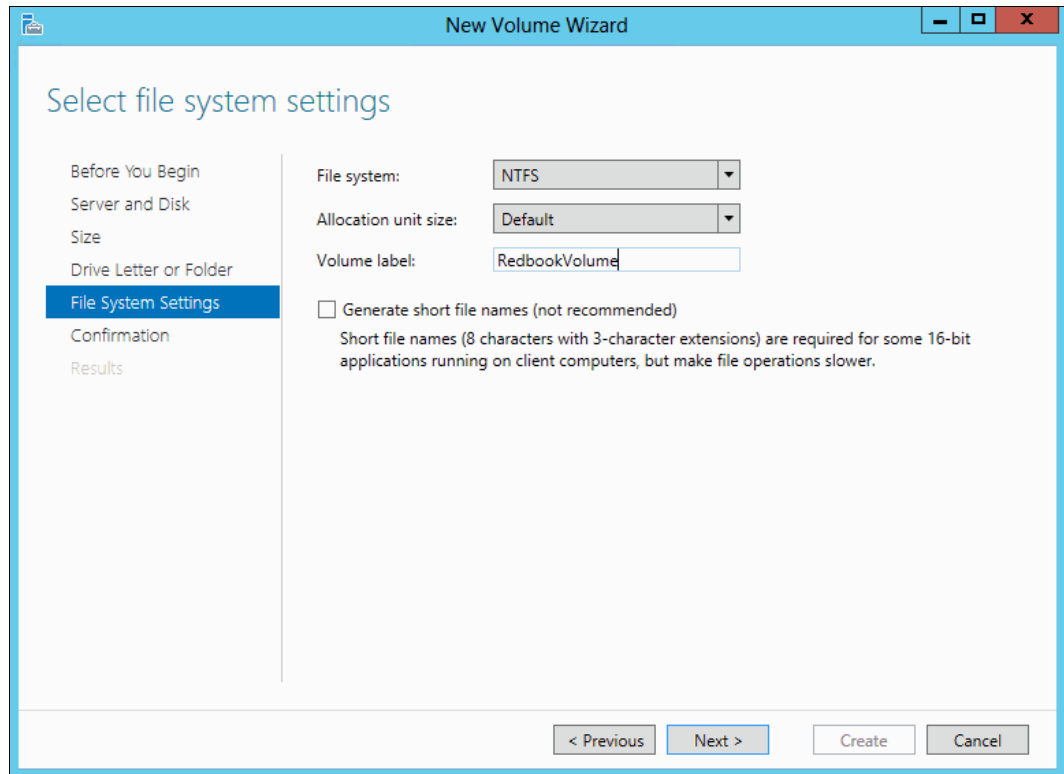


Figure 5-41 Configuring File System Settings

7. Confirm the new volume configuration on the next page, click **Next** and confirm that volume creation is successful, as shown in Figure 5-42.

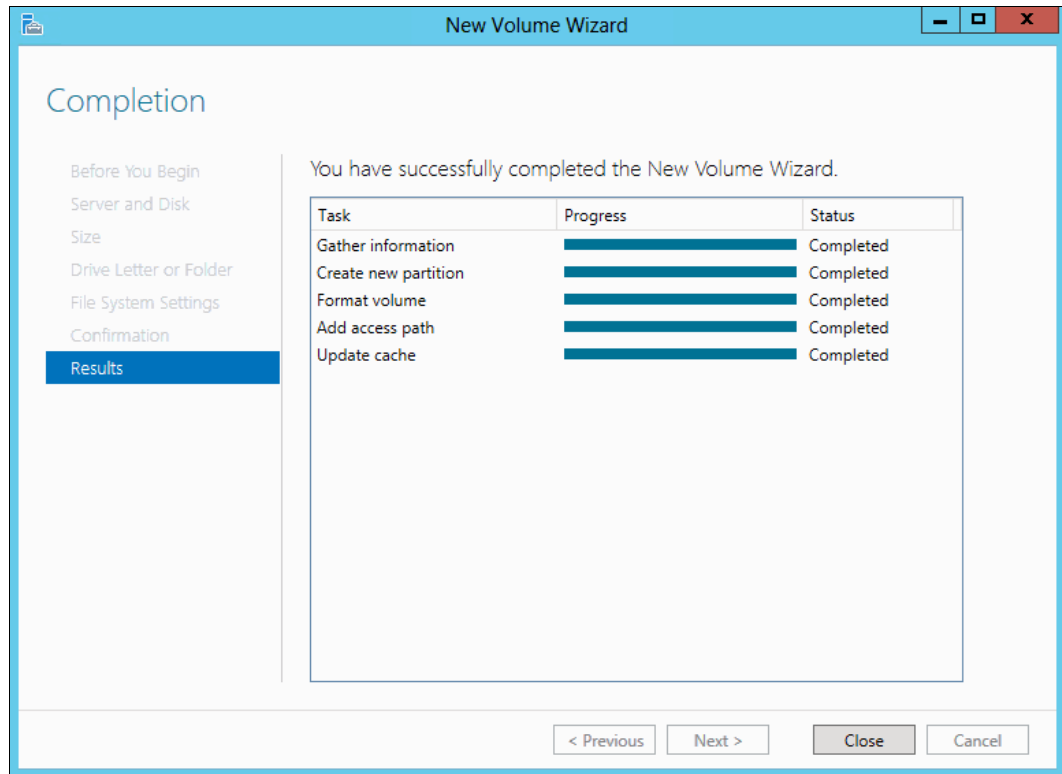


Figure 5-42 New volume success

8. The new volume displays under the Volumes view of **File and Storage Services**, as shown in Figure 5-43.




Figure 5-43 The new volume listed under the Volumes view

Configuring multipath settings for mapped volumes

To configure SDDDSM multipath settings for mapped volumes in Windows 2012 R2, you can use the **datapath.exe** command-line interface (CLI) tool.

To query paths to discovered devices, complete the following steps:

1. Open a Powershell terminal.
2. Change to the SDDDSM directory.

3. Run **datapath.exe query device**, as shown in Example 5-1.

Example 5-1 Datapath query device

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

Total Devices : 1

DEV#: 1 DEVICE NAME: Disk1 Part0 TYPE: 2145 POLICY: OPTIMIZED
SERIAL: 6005076300AF004B4800000000000027 LUN SIZE: 10.0GB
=====
Path# Adapter/Hard Disk State Mode Select Errors
 0 * Scsi Port2 Bus0/Disk1 Part0 OPEN NORMAL 0 0
 1 * Scsi Port2 Bus0/Disk1 Part0 OPEN NORMAL 0 0
 2 Scsi Port2 Bus0/Disk1 Part0 OPEN NORMAL 43 0
 3 Scsi Port2 Bus0/Disk1 Part0 OPEN NORMAL 29 0
 4 * Scsi Port3 Bus0/Disk1 Part0 OPEN NORMAL 0 0
 5 * Scsi Port3 Bus0/Disk1 Part0 OPEN NORMAL 0 0
 6 Scsi Port3 Bus0/Disk1 Part0 OPEN NORMAL 7 0
 7 Scsi Port3 Bus0/Disk1 Part0 OPEN NORMAL 6 0
=====
```

The output provides multipath information about the mapped volumes. In our example, one disk is connected (Disk 1) and eight paths to the disk are available (State = Open). You can also confirm that the serial numbers match the UIDs recorded earlier.

Important: Correct storage area network (SAN) switch zoning must be implemented to allow only eight paths to be visible from the host to any one volume. Volumes with more than eight paths are not supported.

4. To set the multipath policy for a device, run **datapath.exe set device <device ID> policy <rr/fo/lb/df>**, where the variables represent the following actions:

- *rr* is Round Robin
- *fo* is Failover
- *lb* is Load Balancing (reported as *optimized*)
- *df* is Default (which should be Load Balancing)

In Example 5-2, we have changed the policy from Load Balancing to Round Robin.

Example 5-2 Setting a Round Robin multipath policy

```
PS C:\Program Files\IBM\SDDDSM> .\datapath.exe set device 1 policy rr

DEV#: 1 DEVICE NAME: Disk1 Part0 TYPE: 2145 POLICY: ROUND ROBIN
SERIAL: 6005076300AF004B4800000000000027 LUN SIZE: 10.0GB
=====
Path# Adapter/Hard Disk State Mode Select Errors
 0 * Scsi Port2 Bus0/Disk1 Part0 OPEN NORMAL 0 0
 1 * Scsi Port2 Bus0/Disk1 Part0 OPEN NORMAL 0 0
 2 Scsi Port2 Bus0/Disk1 Part0 OPEN NORMAL 43 0
 3 Scsi Port2 Bus0/Disk1 Part0 OPEN NORMAL 29 0
 4 * Scsi Port3 Bus0/Disk1 Part0 OPEN NORMAL 0 0
 5 * Scsi Port3 Bus0/Disk1 Part0 OPEN NORMAL 0 0
 6 Scsi Port3 Bus0/Disk1 Part0 OPEN NORMAL 7 0
 7 Scsi Port3 Bus0/Disk1 Part0 OPEN NORMAL 6 0
=====
```

For more information about configuring SDDDSM, see the Multipath Subsystem Device Driver User's Guide. This can be found on the following website:

http://www.ibm.com/support/docview.wss?rs=540&context=ST52G7&q=ssg1*&uid=ssg1S7000303&loc=en_US&cs=utf-8&lang=en%20en

5.3.2 Discovering iSCSI volumes from Windows 2012 R2

This section describes how to discover iSCSI volumes.

Discovering mapped volumes

To discover mapped iSCSI volumes from Windows 2012 R2, complete the following steps:

1. Connect to your Windows server and open the iSCSI Initiator by opening **Server Manager** and going to **Tools** → **iSCSI Initiator**, as shown in Figure 5-44.




Figure 5-44 Starting the iSCSI Initiator

2. While in the iSCSI Initiator properties, log on to an IBM Storwize V3700 iSCSI port by typing the port Internet Protocol (IP) address into the Target box and clicking **Quick Connect**, as shown in Figure 5-45.




Figure 5-45 Connecting to the IBM Storwize V3700 iSCSI ports

3. If the connection is successful, the window shown in Figure 5-46 opens. Click **Done** to proceed.

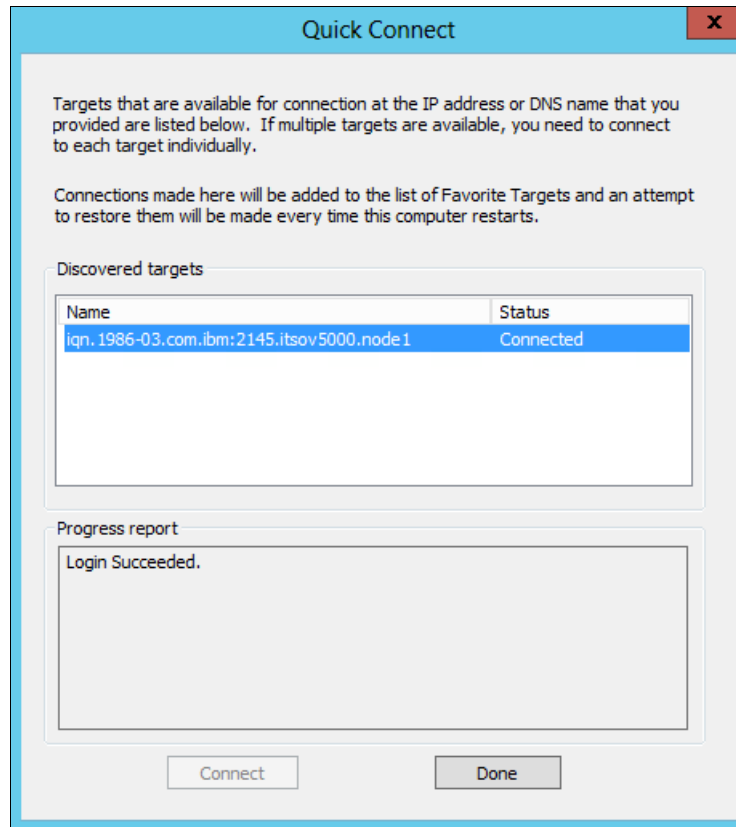


Figure 5-46 Successful iSCSI port connection

4. Repeat the previous two steps for all iSCSI IPs that you want to connect to.
5. In the iSCSI Initiator properties, the connected ports should now be listed under **Discovered targets**, as shown in Figure 5-47.

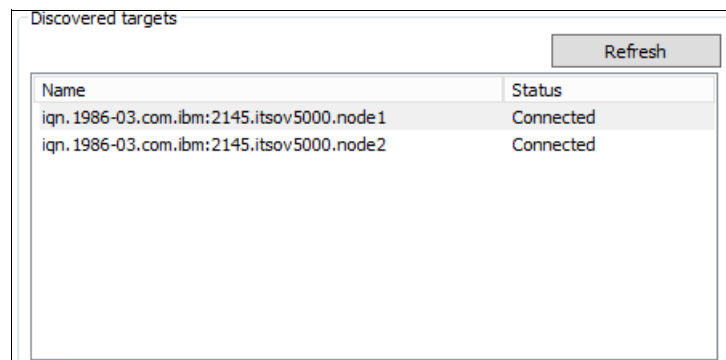


Figure 5-47 Discovered iSCSI targets

6. The next step is to scan for disks. Follow the steps in 5.3.1, “Discovering Fibre Channel volumes from Windows 2012 R2” on page 218 to do this. The number of disks listed should be the number of mapped volumes multiplied by the number of paths, because you have not yet enabled MPIO for iSCSI devices.

7. To enable MPIO for iSCSI devices, first go to **Tools** → **MPIO** in Server Manager, as shown in Figure 5-48.

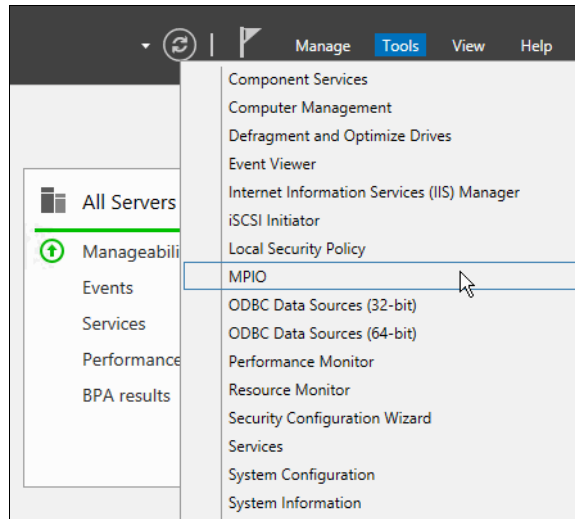


Figure 5-48 Opening the MPIO properties window

This opens the window shown in Figure 5-49.

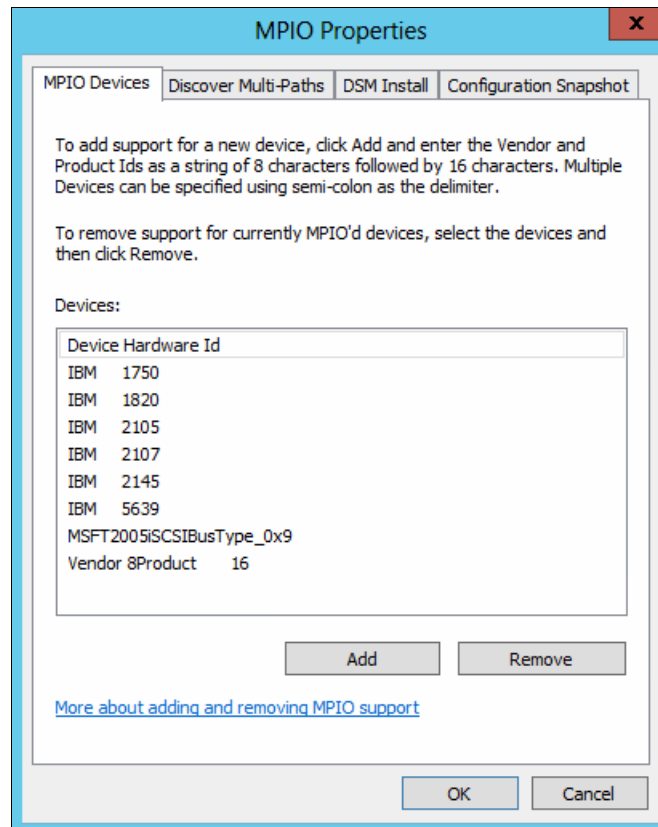


Figure 5-49 MPIO properties

Note: IBM 2145 devices need to be added to the list of supported devices for MPIO to work with IBM Storwize V3700. Click **Add** to do this if they are not already listed.

8. Go to the Discover Multi-Paths tab and select **Add support for iSCSI devices**, as shown in Figure 5-50.

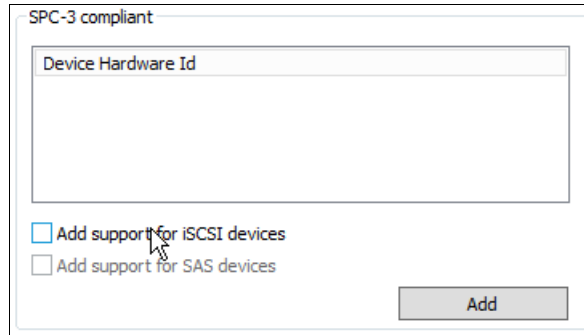


Figure 5-50 Adding support for iSCSI devices in MPIO

9. Click **Add** and then click **Yes** when prompted to restart. Wait for the server to finish restarting.

When the server is back online, the correct number of disks should now be listed under the Disks view in **File and Storage Services**.

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

10. To make iSCSI devices readily available on system boot, reopen the iSCSI Initiator properties and go to **Volumes and Devices**, as shown in Figure 5-51 on page 231.

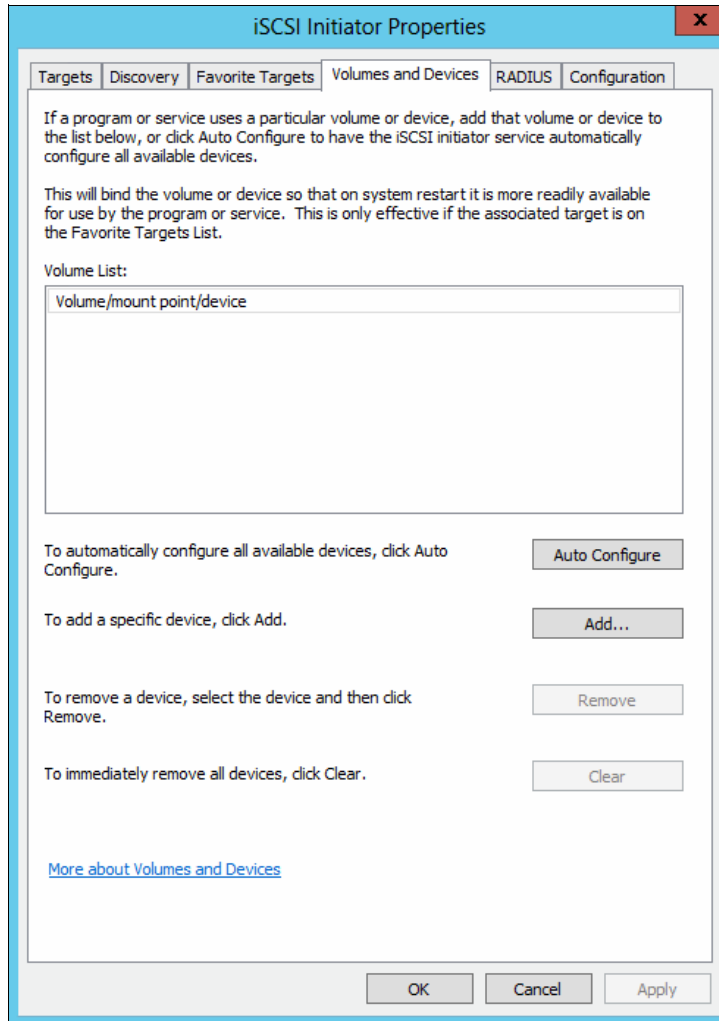


Figure 5-51 iSCSI volumes and devices

11. Click **Auto Configure** and the mapped volumes should display under Volume List, as shown in Figure 5-52. This ensures that, on restart, the volumes are available for use by dependent services and applications before these start.




Figure 5-52 Configured iSCSI device

12. Follow the instructions in 5.3.1, “Discovering Fibre Channel volumes from Windows 2012 R2” on page 218 to bring the disks online and initialize them.

The mapped volumes are now ready to be managed by the host operating system.

Creating a file system from a mapped volume

To create a file system from a mapped volume in Windows 2012 R2, complete the steps described in 5.3.1, “Discovering Fibre Channel volumes from Windows 2012 R2” on page 218.

This procedure is common to all attachment types.

Configuring multipath settings for mapped volumes

To configure MPIO settings for mapped volumes in Windows 2012 R2, you can use the `mpclaim.exe` CLI tool:

1. To save details of the current MPIO configuration, including all connected devices and paths, to a text file, open a Powershell terminal and run `mpclaim -v <directory path\file>`.

For example, you could run `mpclaim -v C:\multipath_config.txt`. Some sample output is shown in Example 5-3.

Example 5-3 MPIO configuration details

MPIO Storage Snapshot on Monday, 27 October 2014, at 14:57:03.169

Registered DSMs: 1

=====

DSM Name	Version	PRP	RC	RI	PVP	PVE	
Microsoft DSM		006.0002.09200.16384	0020	0016	0001	060	True

Microsoft DSM

=====

MPIO Disk0: 04 Paths, Round Robin with Subset, Implicit Only

SN: 6057630AF04B480000002B

Supported Load Balance Policies: F00 RRWS LQD WP LB

Path ID	State	SCSI Address	Weight
0000000077040002	Active/Optimized	004 000 002 000	0
TPG_State: Active/Optimized , TPG_Id: 0, TP_Id: 2432			
Adapter: Microsoft iSCSI Initiator... (B D F: 000 000 000)			
Controller: 46616B65436F6E74726F6C6C6572 (State: Active)			
0000000077040003	Active/Optimized	004 000 003 000	0
TPG_State: Active/Optimized , TPG_Id: 0, TP_Id: 2432			
Adapter: Microsoft iSCSI Initiator... (B D F: 000 000 000)			
Controller: 46616B65436F6E74726F6C6C6572 (State: Active)			
0000000077040001	Active/Unoptimized	004 000 001 000	0
TPG_State: Active/Unoptimized, TPG_Id: 1, TP_Id: 384			
Adapter: Microsoft iSCSI Initiator... (B D F: 000 000 000)			
Controller: 46616B65436F6E74726F6C6C6572 (State: Active)			
0000000077040000	Active/Unoptimized	004 000 000 000	0
TPG_State: Active/Unoptimized, TPG_Id: 1, TP_Id: 384			
Adapter: Microsoft iSCSI Initiator... (B D F: 000 000 000)			
Controller: 46616B65436F6E74726F6C6C6572 (State: Active)			

MSDSM-wide default load balance policy: N\A

No target-level default load balance policies have been set.

=====

2. Confirm that all paths are active as expected, and that the serial numbers match the UIDs recorded earlier.
3. To change the load balancing policy for all devices, run `mpclaim -L -M <policy>`, where the allowed policies are shown in Table 5-1.

Table 5-1 MPIO load balancing policies

Parameter	Definition
0	Clear the Policy
1	Failover Only
2	Round Robin
3	Round Robin with Subset
4	Least Queue Depth
5	Weighted Paths
6	Least Blocks
7	Vendor Specific

For detailed explanations of each policy, see the following link:

<http://technet.microsoft.com/en-gb/library/dd851699.aspx>

The following link provides further information and guidance about using `mpclaim.exe`:

<http://technet.microsoft.com/en-us/library/ee619743%28v=ws.10%29.aspx>

5.3.3 Discovering SAS volumes from Windows 2012 R2

This section describes how to discover SAS volumes.

Discovering mapped volumes

To discover mapped SAS volumes from Windows 2012 R2, complete the steps described in 5.3.1, “Discovering Fibre Channel volumes from Windows 2012 R2” on page 218. The steps required are identical for both FC and SAS.

The mapped volumes are now ready to be managed by the host operating system.

Creating a file system from a mapped volume

To create a file system from a mapped volume in Windows 2012 R2, complete the steps described later in this chapter. This procedure is common to all attachment types.

Configuring multipath settings for mapped volumes

To configure multipath settings for mapped volumes in Windows 2012 R2, complete the steps described in “Configuring multipath settings for mapped volumes” on page 225.

5.3.4 Discovering Fibre Channel volumes from VMware ESXi 5.5

This section describes how to discover FC volumes.

Discovering mapped volumes

To discover mapped FC volumes from VMware ESXi 5.5, complete the following steps:

1. Connect to your ESXi Server using the VMware vSphere Client.
2. Browse to the **Configuration** tab and select **Storage Adapters**. This opens the view shown in Figure 5-53.

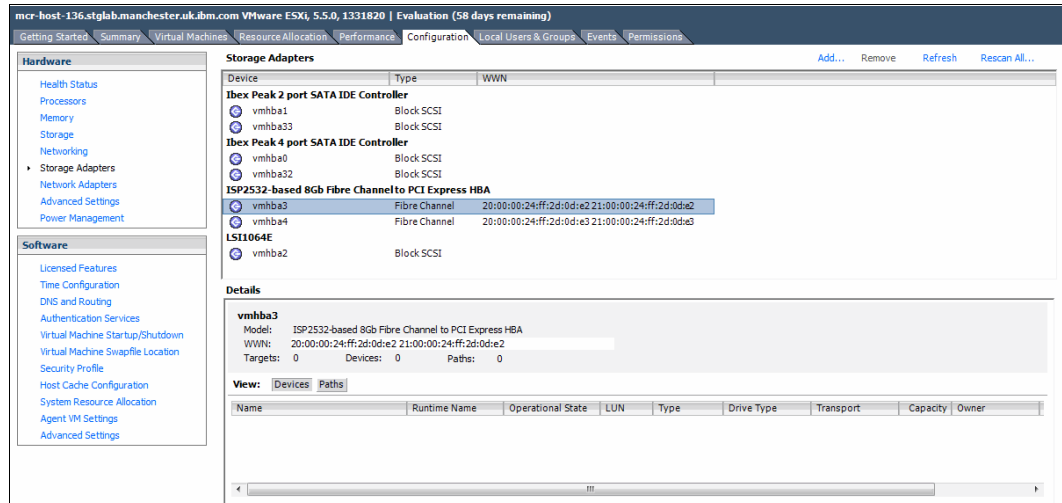


Figure 5-53 The storage adapters view in VMware vSphere Client

3. Click **Rescan all**. The window shown in Figure 5-54 opens. Make sure **Scan for New Storage Devices** is selected. Click **OK** to confirm.




Figure 5-54 Rescan all adapters

- The rescan task starts and displays under **Recent Tasks**, as shown in Figure 5-55. Wait for this to finish.

Name	Target	Status	Details	Initiated by	Requested Start Time	Start Time	Completed Time
Rescan all HBAs	mcr-host-136.s...	In Progress		root	23/10/2014 17:42:10	23/10/2014 17:42:10	

Figure 5-55 The Rescan all HBAs task in progress

- If you now select the FC adapters, the mapped volumes are listed, as shown in Figure 5-56.

Name	Runtime Name	Operational State	LUN	Type	Drive Type	Transport	Capacity	Owner	Hardware
IBM Fibre Channel Disk (naa.60050...	vmhba3:C0:T0:L0	Mounted	0	disk	Non-SSD	Fibre Channel	10.00 GB	NMP	Sup

Figure 5-56 The mapped volumes listed in VMware vSphere Client

- Check the UIDs that you recorded earlier, and confirm that all of the mapped volumes are listed.

The mapped volumes are now ready to be managed by the host operating system.

Creating a data store from a mapped volume

To create a data store from a mapped volume in VMware ESXi 5.5, complete the steps described in 5.3.5, “Discovering iSCSI volumes from VMware ESXi 5.5” on page 235. This procedure is common to all attachment types.

Configuring multipath settings for mapped volumes

To configure multipath settings for a mapped volume in VMware ESXi 5.5, complete the steps described in 5.3.5, “Discovering iSCSI volumes from VMware ESXi 5.5” on page 235. This procedure is common to all attachment types.

5.3.5 Discovering iSCSI volumes from VMware ESXi 5.5

This section describes how to discover iSCSI volumes.

Discovering mapped volumes

To discover mapped iSCSI volumes from VMware ESXi 5.5, complete the following steps:

- Connect to your ESXi Server using the VMware vSphere Client. Browse to the **Configuration** tab and select **Storage Adapters**.

This opens the view shown in Figure 5-57.

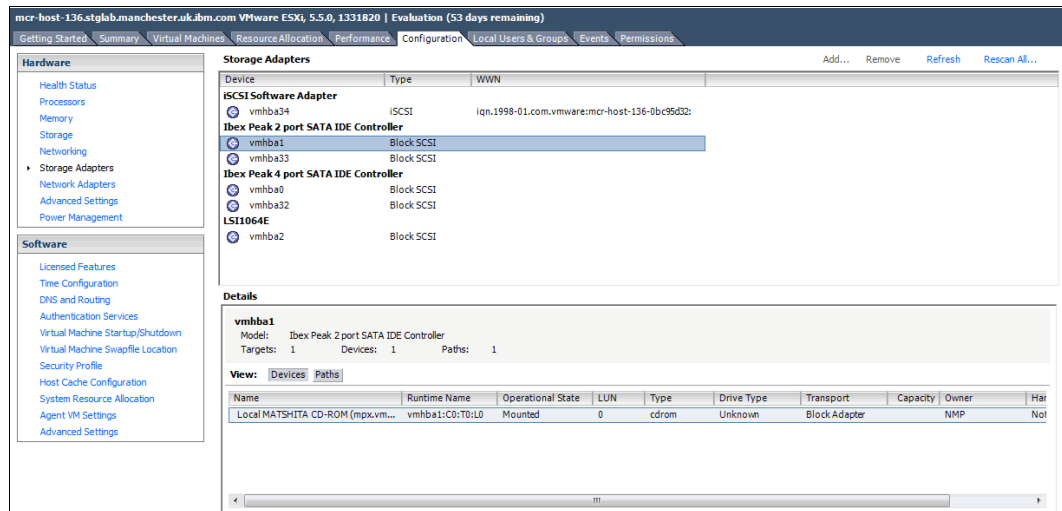


Figure 5-57 The storage adapters view in VMware vSphere Client

2. Right-click the iSCSI Software Adapter and click **Properties**. The iSCSI initiator properties window opens. Go to **Dynamic Discovery** and click **Add**, as shown in Figure 5-58.

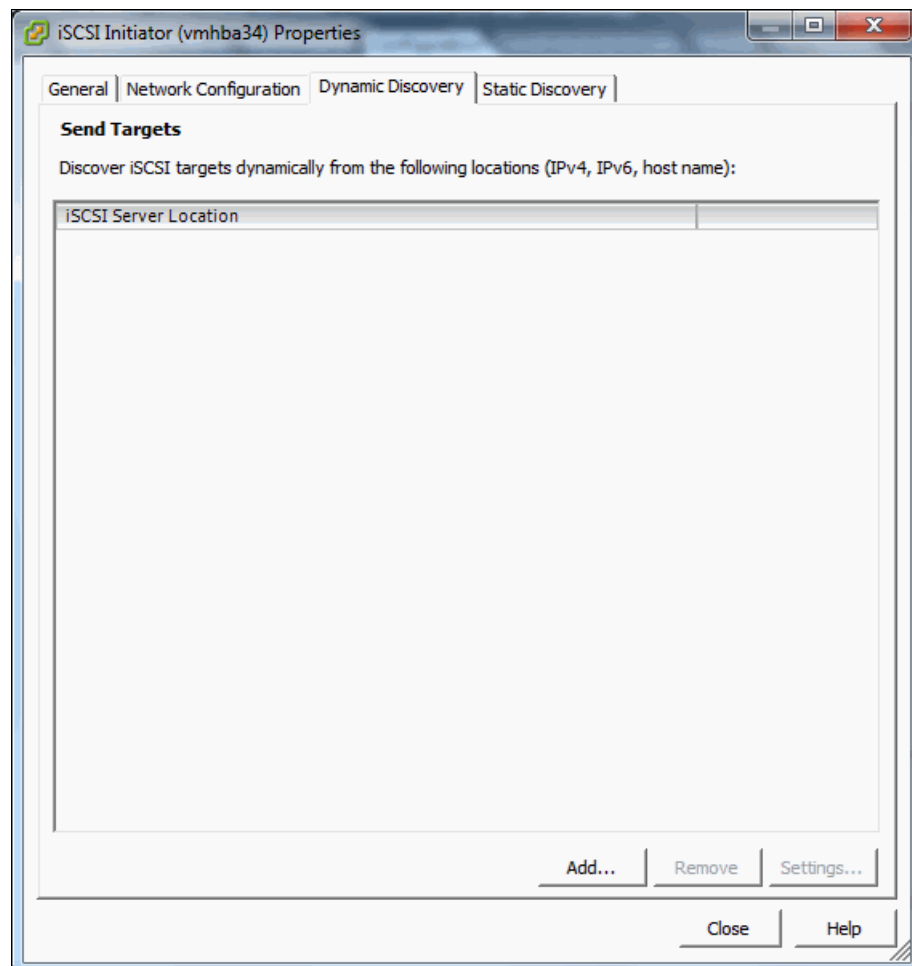


Figure 5-58 Discover iSCSI targets dynamically from the iSCSI Initiator

3. The window shown in Figure 5-59 opens. Enter a target IP address, which should be one of the IPs configured during setup, as described earlier in this book. The target IP address is the iSCSI IP address of a node in the I/O Group from which the iSCSI volumes are mapped. Keep the default value of 3260 for the port and click **OK**.

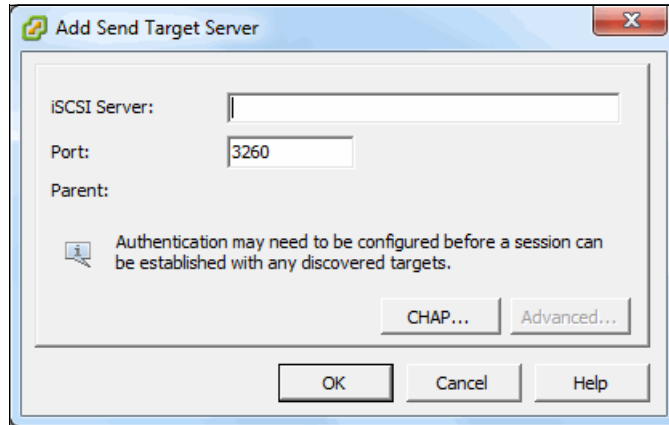


Figure 5-59 Adding a target IP to the iSCSI Initiator

4. Under **Recent Tasks**, a task displays for connecting to the target, as shown in Figure 5-60. Wait for this to complete.

Name	Target	Status	Details	Initiated by	Requested Start Time	Start Time	Completed Time
Add Internet SCSI send...	mcr-host-136.s...	In Progress		root	21/10/2014 17:50:09	21/10/2014 17:50:09	

Figure 5-60 Connecting to iSCSI targets

5. Repeat the previous two steps for each IBM Storwize V3700 iSCSI port that you want to use for iSCSI connections with this host.

iSCSI IP addresses: The iSCSI IP addresses are different from the cluster and canister IP addresses. They are configured as described in Chapter 4, “Host configuration” on page 149.

6. After all of the required ports have been added, close the iSCSI Initiator properties window.
7. The window shown in Figure 5-61 opens. Click **Yes** to rescan the adapter.




Figure 5-61 Rescan the iSCSI Software Adapter

- In **Storage Adapters**, click the iSCSI Software Adapter. When the rescan has completed, the mapped volumes are listed in the adapter details, as shown in Figure 5-62.

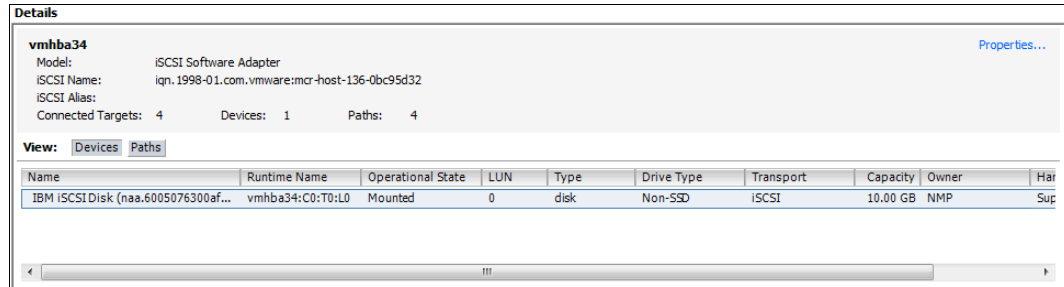


Figure 5-62 The mapped volumes listed in VMware vSphere Client

- Check the UIDs that you recorded earlier and confirm that all of the mapped volumes are listed.

The mapped volumes are now ready to be managed by the host operating system. The next section describes how to create a data store from a mapped volume.

Creating a data store from a mapped volume

To create a data store from a mapped volume in VMware ESXi 5.5, complete the following steps:

- Go to **Configuration** → **Storage**. This opens the window shown in Figure 5-63.

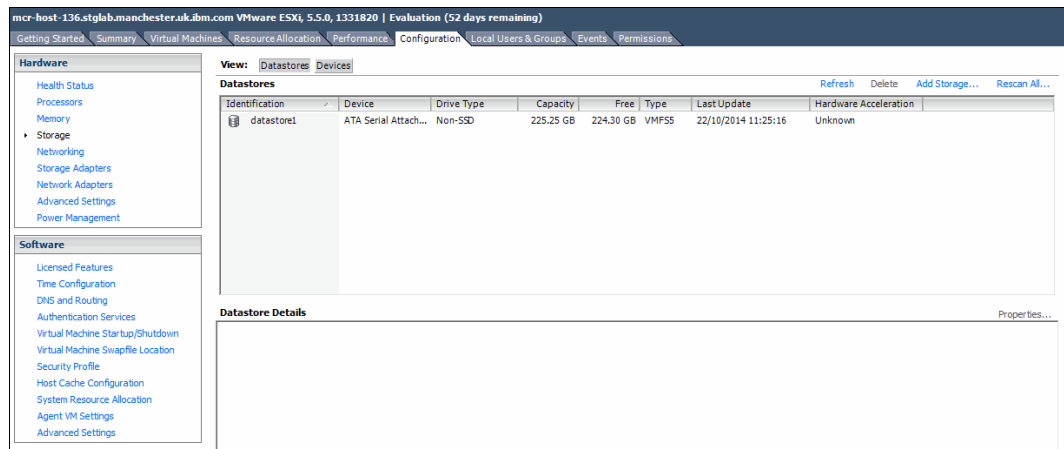


Figure 5-63 The storage view in VMware vSphere Client

2. To create a new data store, click **Add Storage** to open the Add Storage wizard. The window shown in Figure 5-64 opens. Select **Disk/LUN** and click **Next**.

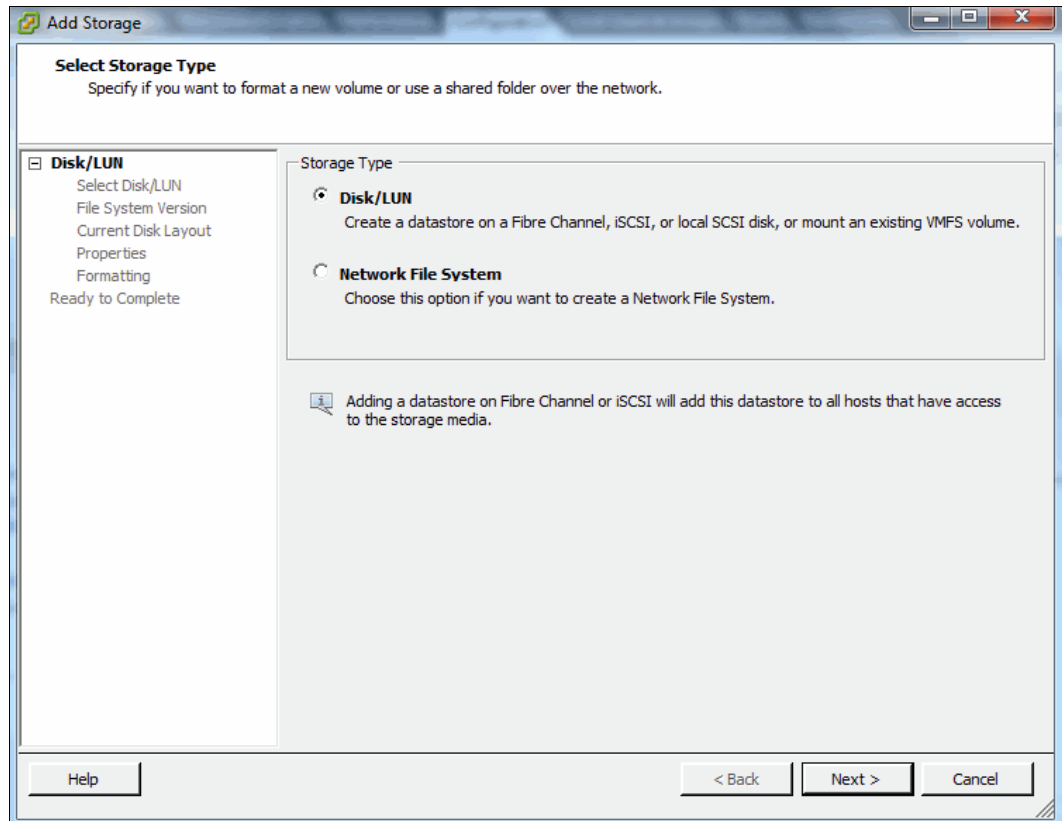


Figure 5-64 The Add Storage wizard

3. The available disks and LUNs (mapped volumes) will be listed on the next page, as shown in Figure 5-65. Select a volume to create the data store from. In this example, we have just one mapped volume, so the choice is easy.




Figure 5-65 Choosing the disk or LUN from which to create the new data store

4. Select an option for the File System Version. In this example, select **VMFS-5**, as shown in Figure 5-66. Click **Next** to continue.

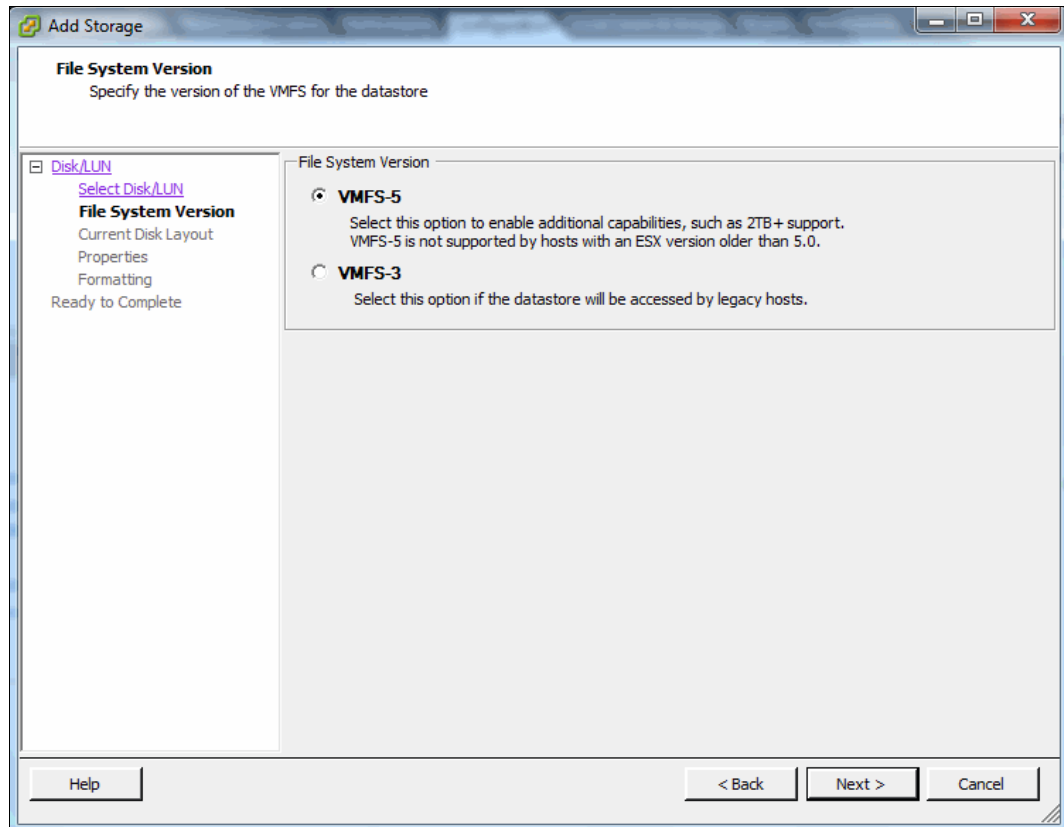


Figure 5-66 Selecting the File System Version for the new data store

5. Review the disk layout and click **Next** to continue, as shown in Figure 5-67.

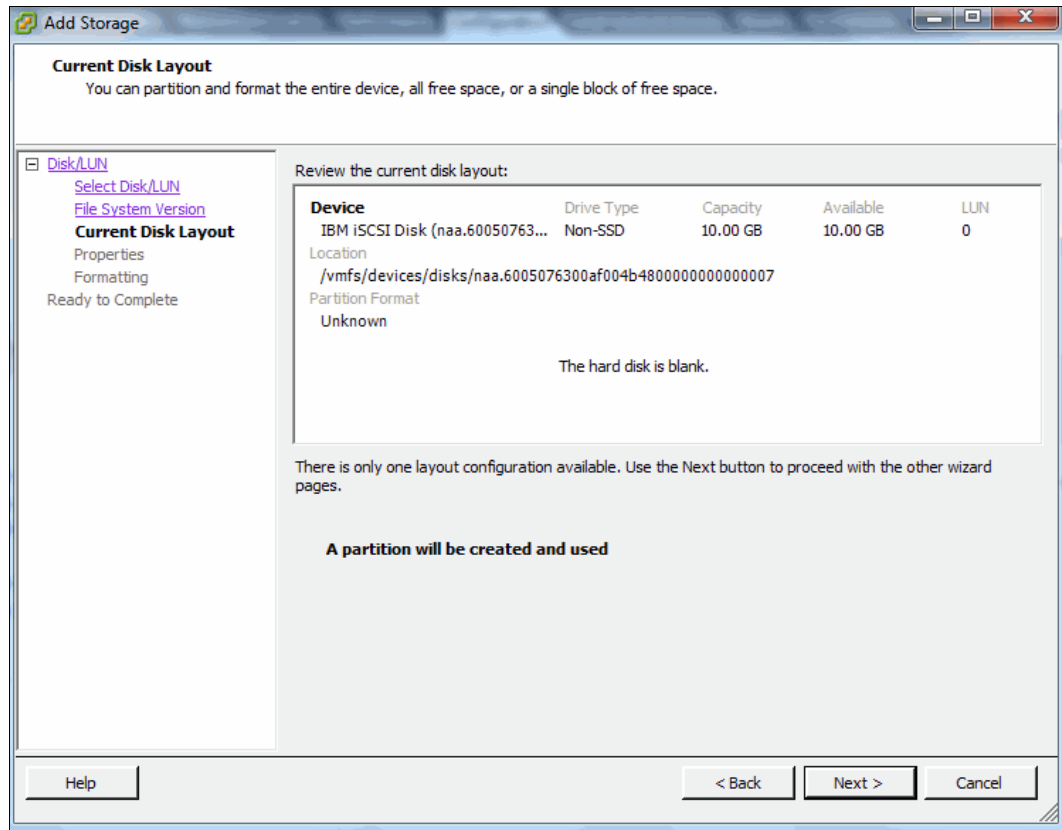


Figure 5-67 Review the current disk layout

6. Enter a name for the data store and click **Next**, as shown in Figure 5-68.




Figure 5-68 Choosing a name for the new data store

7. Select how much of the available disk space to use and click **Next**, as shown in Figure 5-69.

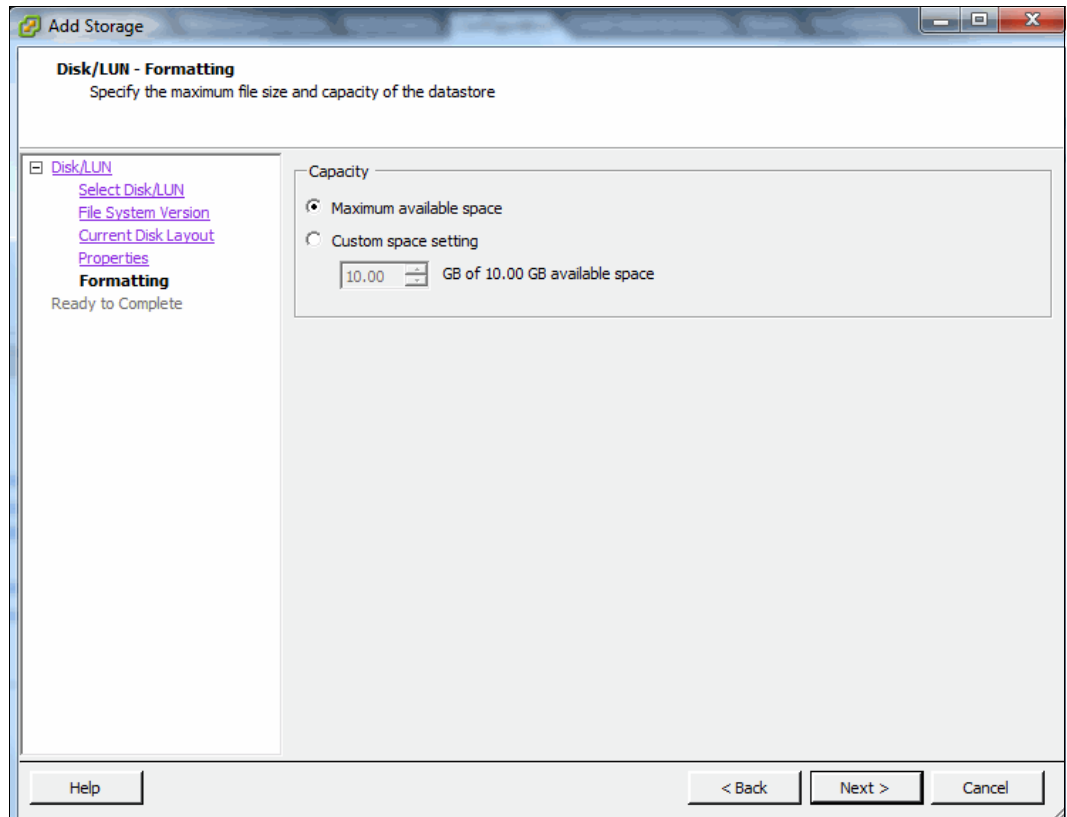


Figure 5-69 Selecting how much of the available disk space to use

- Review your selections and click **Finish**, as shown in Figure 5-70.




Figure 5-70 Completing the Add Storage wizard

- The **Create VMFS datastore** task starts, and displays under **Recent Tasks**, as shown in Figure 5-71. Wait for this to complete.

Name	Target	Status	Details	Initiated by	Requested Start Time	Start Time	Completed Time
Create VMFS datastore	mcr-host-136.s...	Completed		root	22/10/2014 12:03:56	22/10/2014 12:03:56	22/10/2014 12:03:58

Figure 5-71 The Create VMFS data store task under Recent Tasks

- When the task is complete, the new data store displays under **Configuration** → **Storage**, as shown in Figure 5-72.

Identification	Device	Drive Type	Capacity	Free	Type	Last Update	Hardware Acceleration
datastore1	ATA Serial Attach...	Non-SSD	225.25 GB	224.30 GB	VMFS5	22/10/2014 12:04:28	Unknown
RedbookDatastore	IBM iSCSI Disk (n...	Non-SSD	9.75 GB	8.89 GB	VMFS5	22/10/2014 12:04:28	Supported

Figure 5-72 The new data store listed in the storage view

Configuring multipath settings for mapped volumes

To configure multipath settings for a mapped volume managed by VMware as a data store, complete the following steps:

1. Go to **Configuration** → **Storage**.
2. Right-click the data store and click **Properties**. The window shown in Figure 5-73 opens.

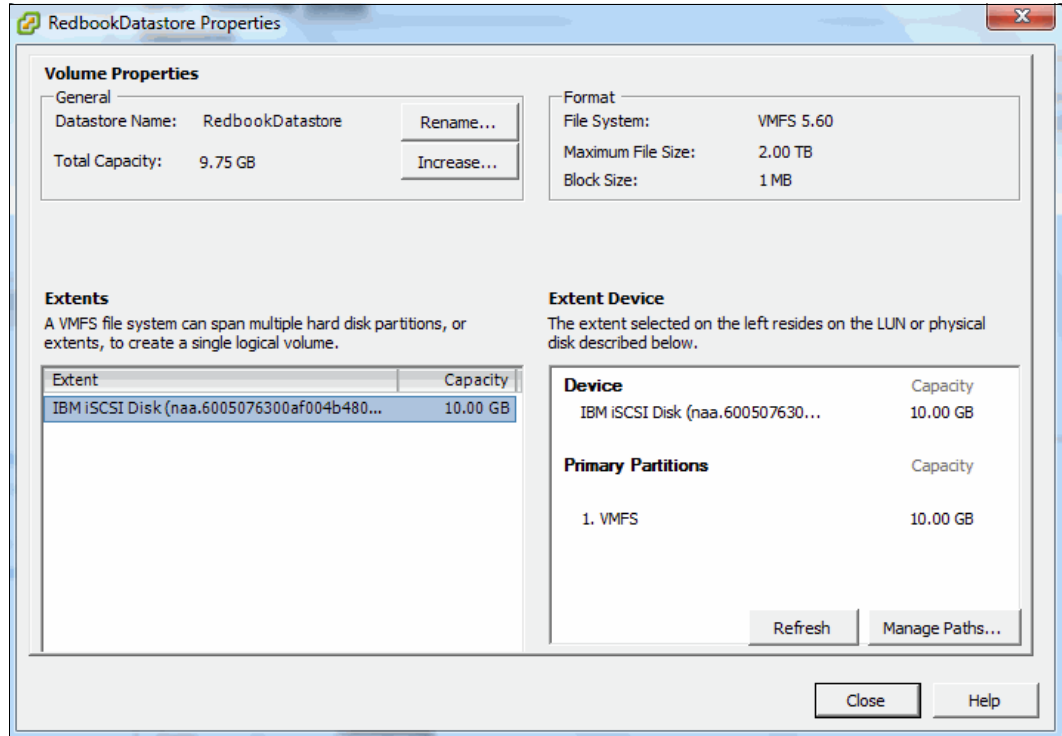


Figure 5-73 Data store properties

3. Click **Manage Paths**. The window shown in Figure 5-74 on page 247 opens. Select either **Most Recently Used**, **Round Robin**, or **Fixed**, and then click **Change** to confirm.
Round Robin is the suggested option, because this enables for I/O load balancing. For more information about the available options, see Chapter 4, “Host configuration” on page 149.




Figure 5-74 Data store multipath settings

4. When you are finished, close both windows.

The mapped volume is now configured for maximum availability and ready to be used.

5.3.6 Discovering SAS volumes from VMware ESXi 5.5

This section describes how to discover SAS volumes.

Discovering mapped volumes

To discover mapped SAS volumes from VMware ESXi 5.5, complete the steps described in 5.3.4, “Discovering Fibre Channel volumes from VMware ESXi 5.5” on page 234. The steps required are identical for both FC and SAS.

The mapped volumes are now ready to be managed by the host operating system.

Creating a data store from a mapped volume

To create a data store from a mapped volume in VMware ESXi 5.5, complete the steps described in 5.3.5, “Discovering iSCSI volumes from VMware ESXi 5.5” on page 235. This procedure is common to all attachment types.

Configuring multipath settings for mapped volumes

To configure multipath settings for mapped volumes in VMware ESXi 5.5, complete the steps described in 5.3.5, “Discovering iSCSI volumes from VMware ESXi 5.5” on page 235. This procedure is common to all attachment types.



Storage migration

This chapter describes the steps involved in using the Storage Migration wizard. The Storage Migration wizard is used to migrate data from existing external storage systems to the internal capacity of the IBM Storwize V3700. Migrating data from other storage systems to the Storwize V3700 storage system consolidates storage and enables the Storwize V3700 functionality.

This includes the graphical user interface (GUI), internal virtualization, thin provisioning, IBM Easy Tier, and IBM FlashCopy, to be realized across all volumes. When the migration is complete, the external system can be retired.

This chapter includes the following topics:

- ▶ Interoperability and compatibility
- ▶ Storage migration considerations
- ▶ Storage migration example configuration

6.1 Interoperability and compatibility

Interoperability is an important consideration when a new storage system is set up in an environment that contains existing storage infrastructure. This section describes how to check that the storage environment, the existing storage system, and IBM Storwize V3700 are ready for the data migration process.

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

If the required configuration is not listed for support in the SSIC, contact your IBM sales representative and ask for a request for price quotation (RPQ) for your specific configuration.

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

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

6.2 Storage migration considerations

The Storwize V3700 Storage Migration wizard simplifies the migration task. The wizard features easy-to-follow panes that guide you through the entire process.

6.2.1 External virtualization capability

To migrate data from an existing storage system to the Storwize V3700, it is necessary to use the built-in external virtualization capability of the Storwize V3700. This capability places external Fibre Channel (FC)-connected logical units (LUs) under the control of the Storwize V3700. External virtualization also enables Storwize V3700 to control serial-attached Small Computer System Interface (serial-attached SCSI, known as SAS)-connected LUs.

When the volumes are virtualized, hosts continue to access them, but do so through the IBM Storwize V3700, which acts as a proxy.

Restriction: The IBM Storwize V3700 only supports external virtualization for the purposes of data migration. Permanently virtualizing external storage for use under the Storwize V3700 is not supported.

6.2.2 Overview of the Storage Migration wizard

The Storage Migration wizard process includes the following considerations and steps:

1. Typically, storage systems divide storage into many SCSI LUs that are presented on a FC SAN to hosts. Storage can also be presented through direct SAS attachment to a host. Broadly speaking, the steps to migrate either one of these storage systems are the same.
2. Input/output (I/O) to the LUNs must be stopped and changes made to the mapping of the storage system LUs and to the SAN fabric zoning or SAS connectivity, so that the original LUs are presented directly to the Storwize V3700 and not to the hosts. The Storwize V3700 discovers the external LUs as *unmanaged* (not a member of any Storage Pools) managed disks (MDisks).

3. The unmanaged MDisks are then imported to the Storwize V3700 as image mode MDisks and placed into a Storage Pool. This Storage Pool is now a logical container for the externally attached LUs.
4. Each existing volume has a one-to-one mapping with an image mode MDisk. From a data perspective, the image mode volume represents the externally attached LUs exactly as they were before the import operation. The image mode volumes are on the same physical drives of the existing storage system, and the data remains unchanged. The Storwize V3700 is presenting active images of the externally attached LUs, and is acting as a proxy.
5. The hosts must have the existing storage system multipath device driver removed, and then be configured for Storwize V3700 attachment. Further zoning or SAS connectivity changes are made for host-to-V3700 SAN connections. The Storwize V3700 hosts are defined with worldwide port names (WWPNs), and the volumes are mapped to the hosts. After the volumes are mapped, the hosts discover the Storwize V3700 volumes through a host rescan or restart operation.
6. Storwize V3700 volume mirror operations are then initiated. The image mode volumes are mirrored to generic volumes. The mirrors are online migration tasks, which means a defined host can still access and use the volumes during the mirror synchronization process.
7. After the mirror operations are complete, the volume mirror relationships and the image mode volumes are removed. The external storage system LUs are now migrated and the now redundant external storage can be retired or reused elsewhere.

Restriction: If you are going to be migrating volumes from another IBM Storwize product, be aware that the target Storwize unit needs to be configured at the *Replication* layer in order for the source to discover the target system as a host. The default layer setting for Storwize V3700 is *Storage*.

If this Replication layer configuration is not done, it will not be possible to add the target system as a host on the source storage system, or to see source volumes on the target system. For more information about layers and how to change them, see Chapter 10, “Copy services” on page 455.

6.2.3 Storage Migration wizard tasks

The Storage Migration wizard is designed to provide easy and nondisruptive migration of data from other storage systems to the internal capacity of the Storwize V3700, taking into account the considerations in 6.2.2, “Overview of the Storage Migration wizard” on page 250. Although this is true, the initial information panes of the migration wizard should all be read first, to prepare and enable a successful migration. Some of the information in these panes needs to be recorded before actions in previous panes can be carried out.

In early releases of V7.4 code, it is not possible to start the migration wizard unless you have already cabled or zoned the external system to the Storwize V3700. Therefore, we list the Storage migration tasks in the same order as the wizard displays them, but it is important to read this chapter fully before beginning *any* migration task and to collate the relevant information before starting the wizard.

The following are the storage migration wizard tasks as presented by the wizard:

- ▶ Accessing the Storage Migration wizard
- ▶ Before you begin
- ▶ Prepare environment for migration
- ▶ Map storage

- ▶ Migrating MDisks
- ▶ Configure hosts
- ▶ Map volumes to hosts
- ▶ Select Storage Pool
- ▶ Finish the Storage Migration wizard
- ▶ Finalize migrated volumes

However, the first task is to document the existing system and record the data needed for the migration.

6.2.4 Storage migration example configuration

The example scenario we use involves migrating existing volumes that are presented from a direct SAS attached IBM System Storage DS3200 to a Microsoft Windows 2012 host. There are three volumes presented from the System Storage DS3200: Redbookvol1, drive letter G: 8 gigabytes (GB); Redbookvol2, drive letter E: 15 GB; and Redbookvol3, drive letter F: 10 GB. These three volumes are to be migrated from the System Storage DS3200 onto the internal storage capacity of the Storwize V3700.

6.2.5 Recording migration information

Before a successful migration can take place, it is necessary to record certain information about the existing system:

1. First, record information about the existing external logical unit numbers (LUNs). In our example, from the Windows 2012 Server Manager Dashboard, select **Tools** → **Computer Management** → **Storage** → **Disk Management**, as shown in Figure 6-1. You can see the three volumes, disk (G:), disk (E:), and disk (F:).




Figure 6-1 Host volume properties

- Highlight either the Disk 1, Disk 2, or Disk 3 labels. Right-click and select **Properties**, as shown in Figure 6-2.

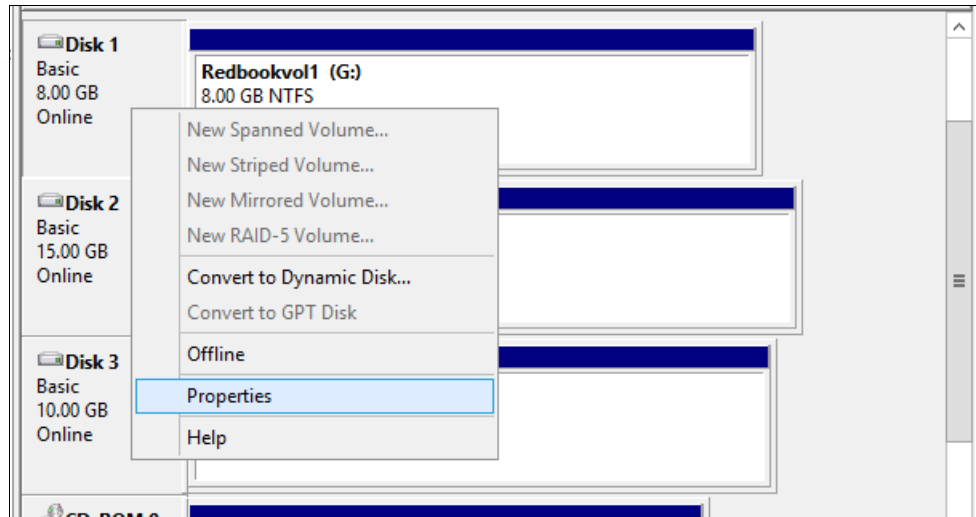


Figure 6-2 Selecting disk properties: Windows 2012

The properties shown indicate that these volumes are presented from a System Storage DS3200, as indicated by the device type 1726 - 2xx shown in Figure 6-3. You can also tell that the volume is using the IBM device-specific module (DSM) for the native Microsoft Multipath I/O (MPIO) feature.




Figure 6-3 Disk properties: Windows 2012

- Using IBM System Storage DS® Storage Manager, connect to the external System Storage DS3200 system, as shown in Figure 6-4, and select **Manage a Storage Subsystem**. Select the system that you want to manage and click **OK**.

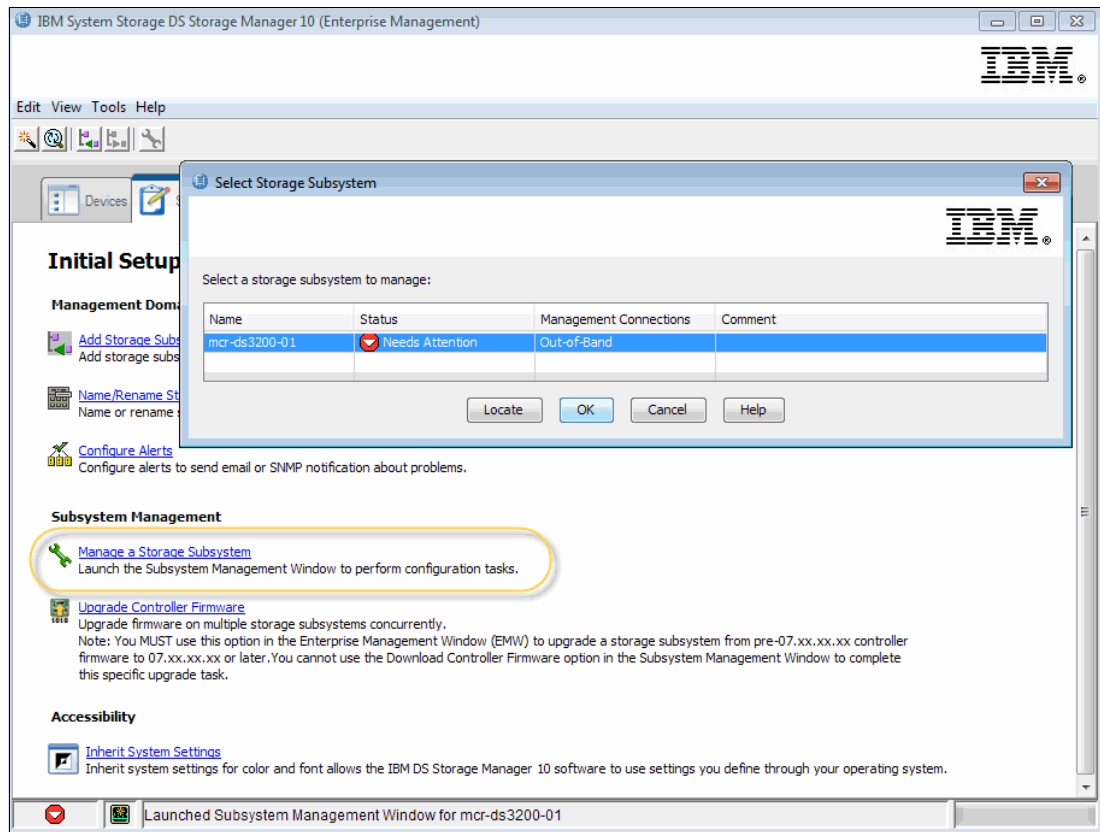


Figure 6-4 Selecting Storage system to manage: DS Storage Manager

- A new window opens to enable management of the IBM DS3200 subsystem. From the summary pane, select the **Storage Subsystem Profile** option, as shown in Figure 6-5.

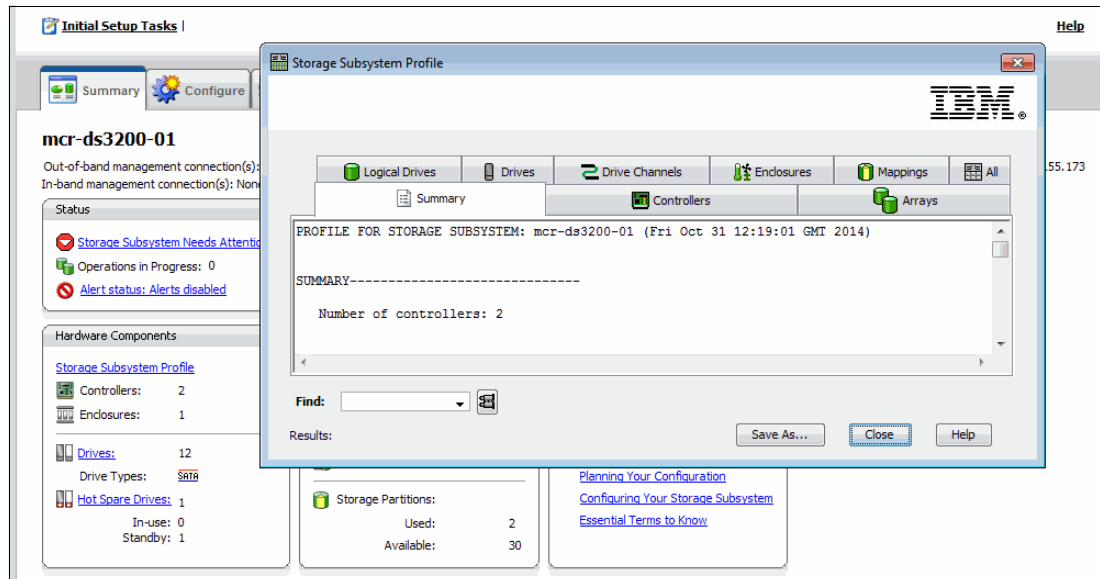


Figure 6-5 Storage Subsystem Profile: System Storage DS Storage Manager

5. Use the information in the tabs to complete the information shown in Table 6-1.

Table 6-1 Example table for capturing external LU information

LU name	Controller	Array	SCSI ID	Host name	Capacity	Host Volume
Win2012vol_01	SlotA	DS3200	1	mcr-host-127	15 GB	Redbookvol2
Win2012vol_02	SlotA	DS3200	2	mcr-host-127	10 GB	Redbookvol3
Win2012vol_03	SlotA	DS3200	0	mcr-host-127	8 GB	Redbookvol1

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

6. Next, record the host information. The easiest way of determining the WWPN of host adapters is to look on the host profile for the existing storage. In this case, from the summary pane of the IBM System Storage DS Storage Manager, select the **Configured Hosts** option to show the host SAS ports, as shown in Figure 6-6.




Figure 6-6 Obtaining host WWPN from the external storage

Remember: Don't forget to record the WWPN of the existing host. It can sometimes be difficult getting this information from the host after it has been removed from the existing storage.

7. Go to the Windows 2012 Server Manager Dashboard and select **Tools** → **Computer Management** → **Device Manager**.
8. Select the storage controller used to connect to the external storage, as shown in Figure 6-7.

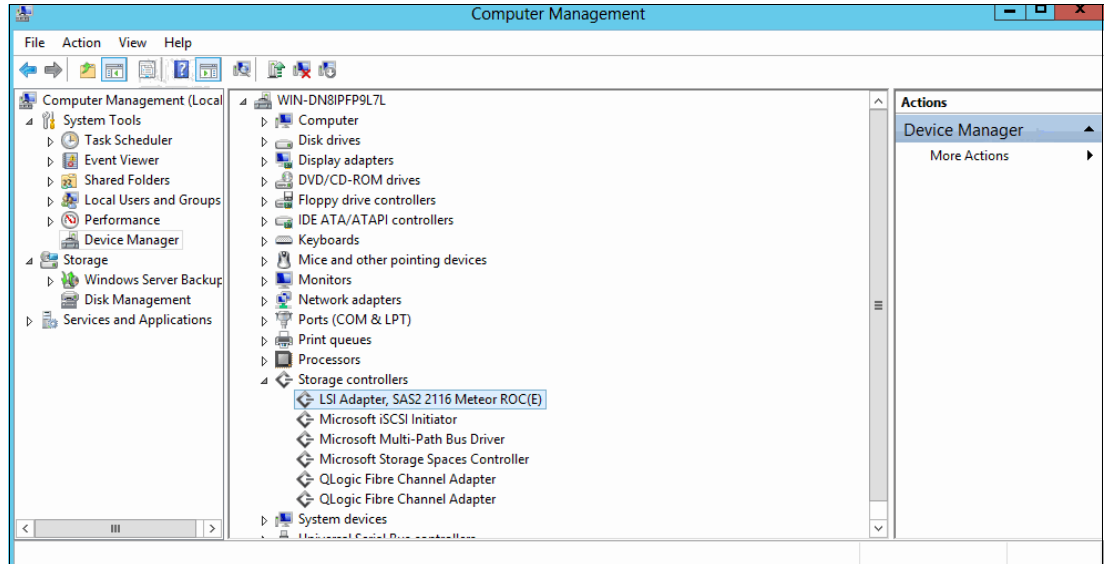


Figure 6-7 Selecting host HBA: Windows 2012

9. From here, use the tabs to record the information about the controller, as shown in Figure 6-8.

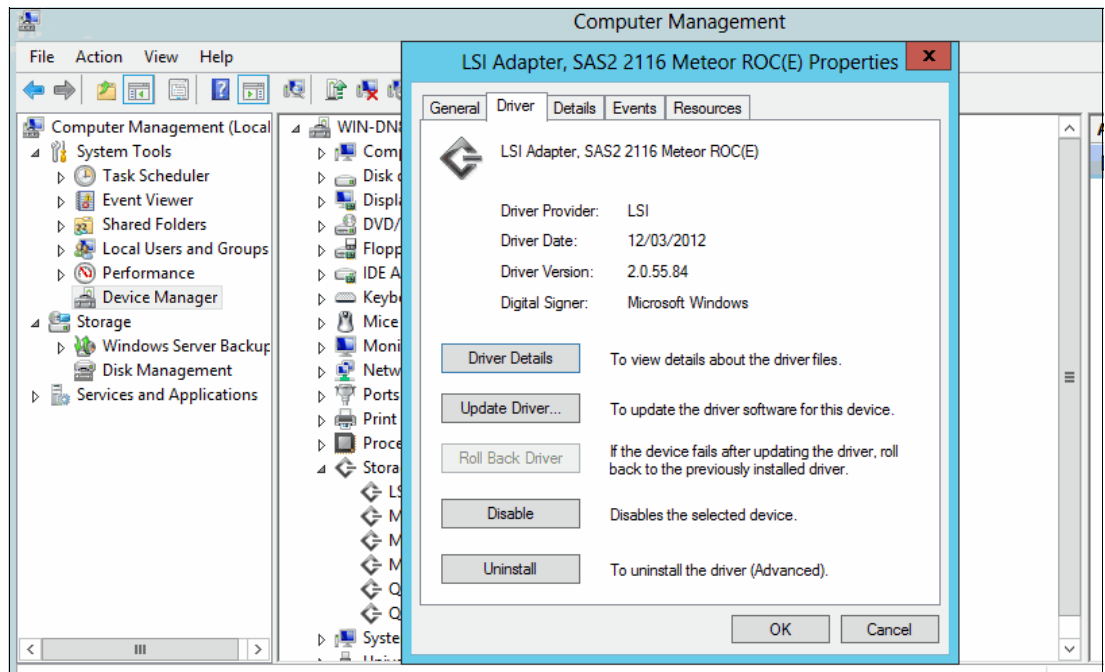


Figure 6-8 Host HBA details: Windows 2012

Table 6-2 shows an example table for capturing host information, including host bus adapter (HBA).

Table 6-2 Example table for capturing host information

Host Name/ LU Names	Adapter / Slot / Port	WWPN	HBA F/W	HBA Device Driver	Operating System	Multipath Software
mcr-host-127	LSI SAS2 2116/ 2 / 1	500062b200556140	16.0	2.0.55.84	Win2012	Native MPIO with DSM
mcr-host-127	LSI SAS2 2116/ 2 / 2	500062b200556141	16.0	2.0.55.84	Win2012	Native MPIO with DSM
mcr-host-127	LSI SAS2 2116/ 2 / 2	500062b200556142	16.0	2.0.55.84	Win2012	Native MPIO with DSM
mcr-host-127	LSI SAS2 2116/ 2 / 2	500062b200556143	16.0	2.0.55.84	Win2012	Native MPIO with DSM

Note: In this example, there are more than the four ports (shown in Table 6-2) listed against the host profile on the System Storage DS3200. For direct-attached SAS hosts, this is sometimes necessary, as described in the following support article:

<https://www.ibm.com/support/docview.wss?uid=ssg1S1004386>

Our table is reduced for clarity.

10. Next, record the port numbers of the external storage controllers and the IBM Storwize V3700. To find the WWPN of the Storwize V3700, complete the following steps:

- a. Go to the System view and switch the 3D image to the rear.
- b. Hover over the point between both canisters and right-click.
- c. Select **View** → **SAS Ports**, as shown in Figure 6-9.

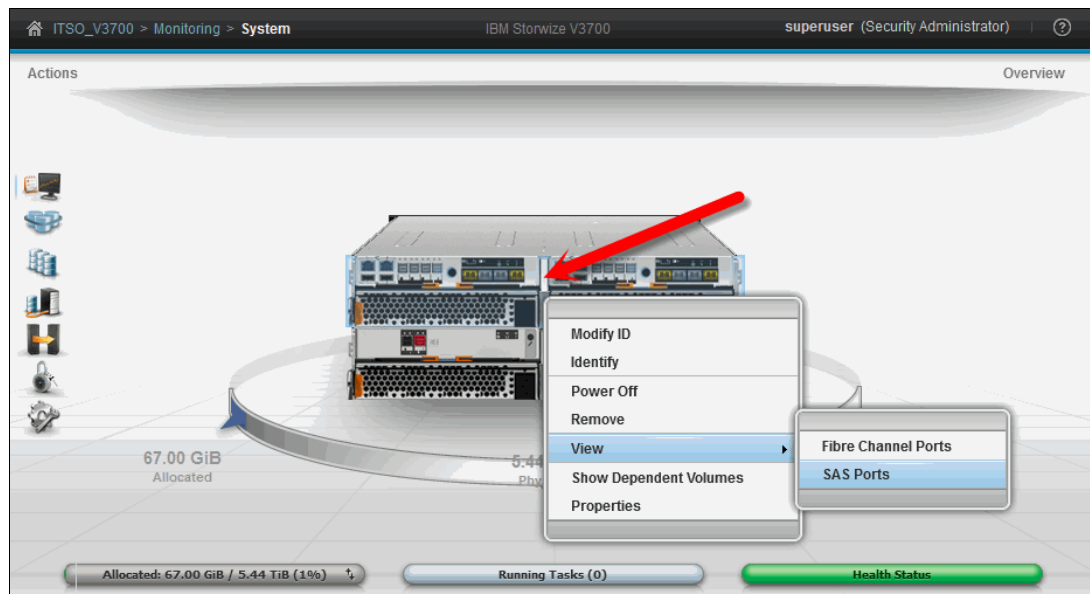


Figure 6-9 Identifying Storwize V3700 SAS port WWPN

11. A list of the SAS port WWPNs is presented, as shown in Figure 6-10. Record the WWPN of the ports that are used to connect to the external System Storage DS3200.

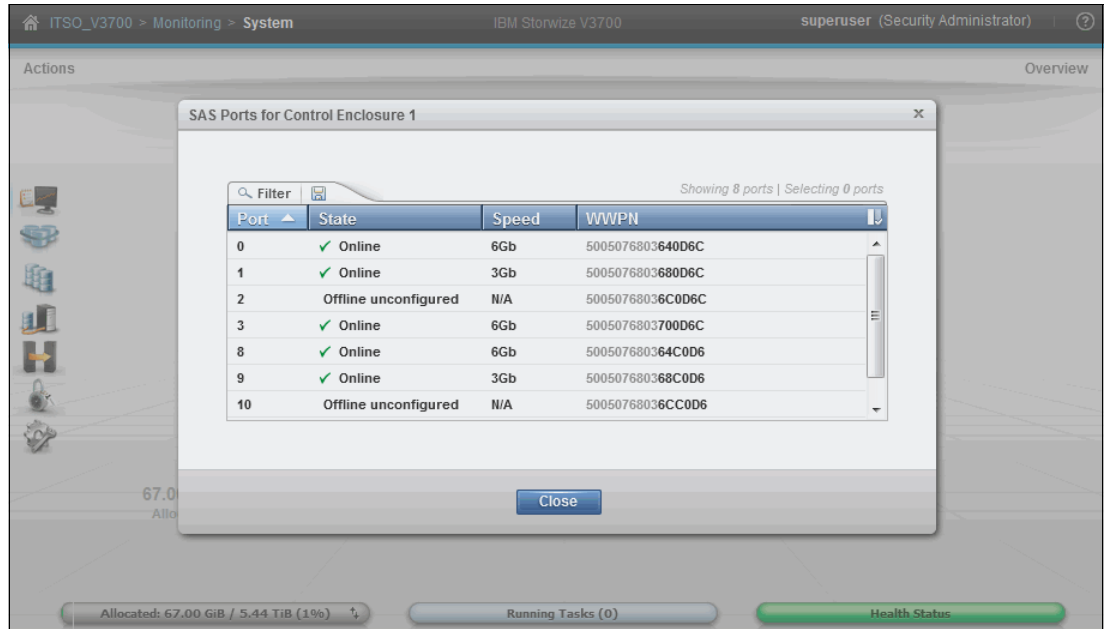


Figure 6-10 Storwize V3700 SAS port list

12. To find the SAS port WWPN from the IBM System Storage DS Storage Manager, go to the Subsystem Management summary window and click **Storage Subsystem Profile**.

13. Scroll through the Controllers tab to find the SAS port WWPN that is used to connect to the Storwize V3700, as shown in Figure 6-11.

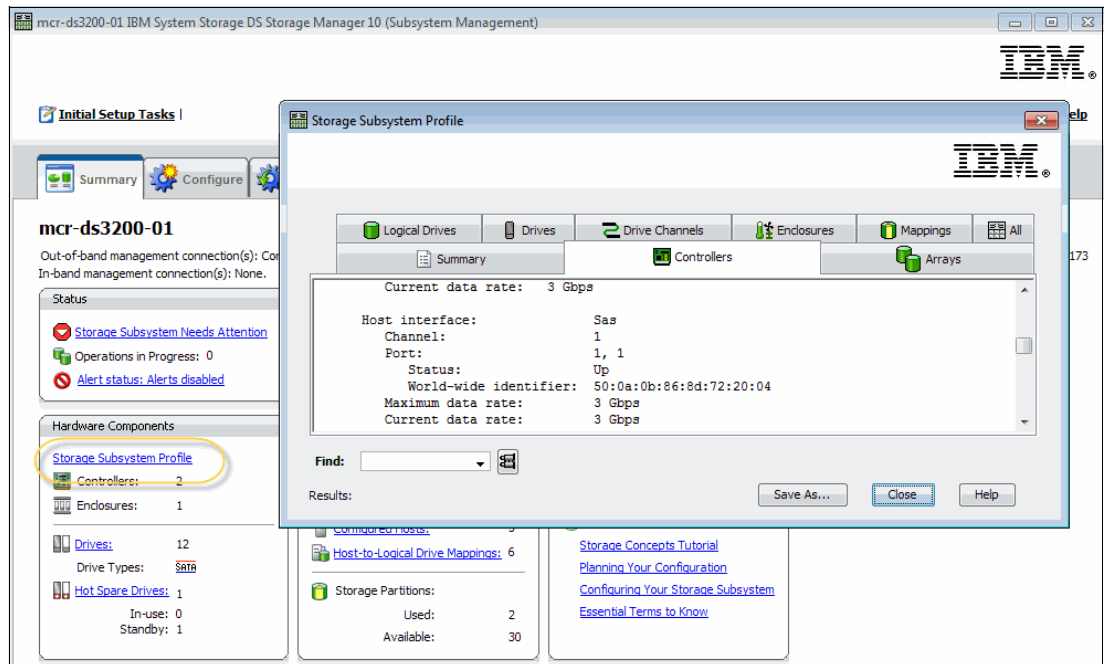


Figure 6-11 Obtaining DS3200 SAS port WWPN: DS Storage Manager GUI

14. Record the information shown in Table 6-3.

Table 6-3 SAS port WWPN for Storage Controllers

Array	Controller	WWPN
DS3200	A	500A0B868D722008
DS3200	B	500A0B868D74000C
V3700	canister 1	5005076803680D6C
V3700	canister 2	500507680368C0D6

15. In addition, make sure that you pay attention to the following tasks:

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

6.3 Storage migration example

Avoiding data loss

The risk of losing data when the Storage Migration wizard is used correctly is low. However, it is prudent to avoid potential data loss by creating a backup of all of the data stored on the hosts, the existing storage systems, and the Storwize V3700 before the wizard is used.

Accessing the Storage Migration wizard

To navigate the Storage Migration wizard, complete the following steps:

1. Select **System Migration** in the **Pools** menu to open the System Migration pane, as shown in Figure 6-12. The System Migration pane provides access to the Storage Migration wizard and displays the migration progress information.




Figure 6-12 Pools menu

2. Click **Start New Migration** to begin the storage migration wizard. Figure 6-13 shows the System Migration pane.




Figure 6-13 System Migration pane

Important: If you have not already configured your zoning correctly or attached your external SAS storage (or have the layer set incorrectly if this is another Storwize system), you might receive a warning message indicating that no external storage controllers could be found, as shown in Figure 6-14 on page 261. Click **Close** and correct the problem before trying to start the migration wizard again. This occurs with early V7.4 release code.

The subsequent panes in the migration wizard (as shown in Figure 6-18 on page 264) direct you to remove host zoning or SAS connectivity to the external storage. You are also prompted to create zones or SAS connections between the Storwize V3700 and the external storage.

However, you must already have done this before you start the wizard. See “Prepare the environment for migration” on page 264 and “Map storage” on page 278 for a list of instructions to complete before starting the data migration wizard.

Figure 6-14 shows the error message indicating that no external controllers could be found.

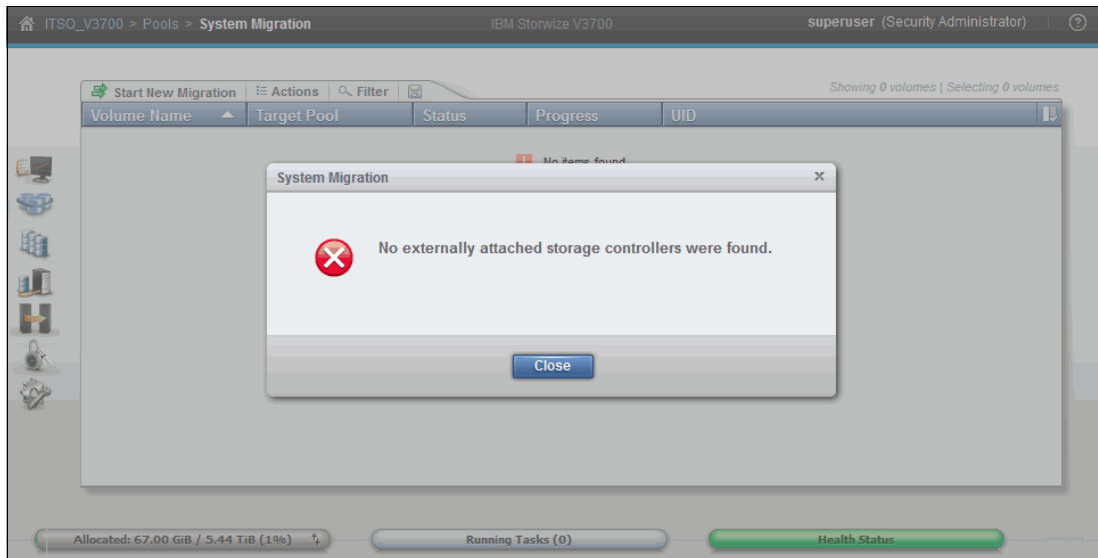


Figure 6-14 Error message displayed when no externally attached storage controllers are found

3. If there are both FC-attached and SAS-attached external storage systems that the IBM Storwize V3700 can see, the wizard progresses with a choice of migrating an Fibre Channel or SAS host, as shown in Figure 6-15.

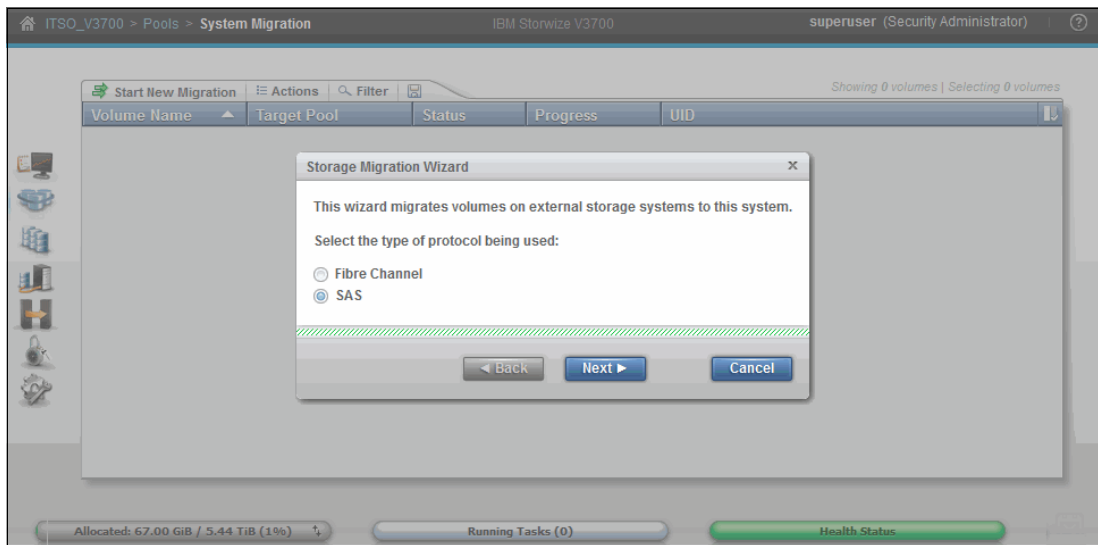


Figure 6-15 Choosing FC or SAS attached storage for migration

We have chosen SAS, because in our example, we are migrating from an external IBM DS3200 system. This choice does not display if there is only one type of attached external storage.

Before you begin

The Before You Begin pane of the Storage Migration wizard describes the restrictions and prerequisites for the wizard to complete successfully, as shown in Figure 6-16.

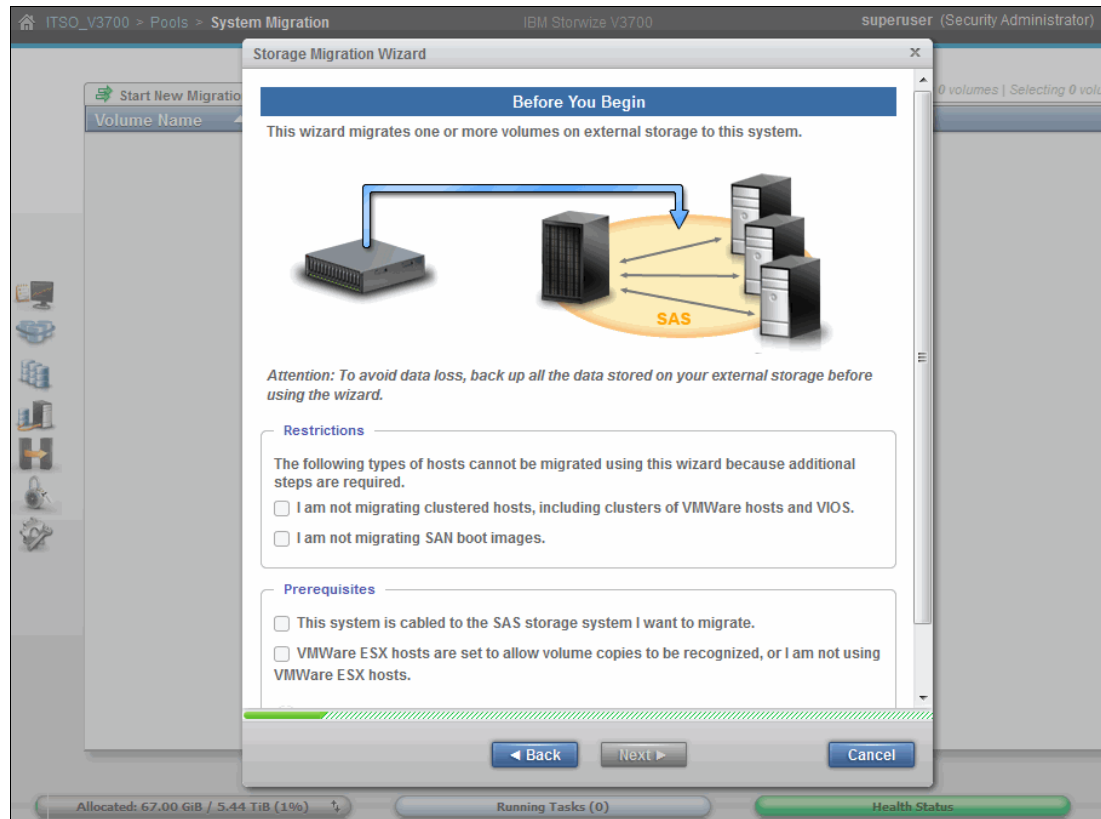


Figure 6-16 Step 1 of the Storage Migration wizard

Restrictions

Confirm that the following conditions are met:

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

If the restrictions options cannot be selected, the migration must be performed outside of this wizard, because more steps are required. For more information, see the IBM Storwize V3700 Knowledge Center on the following website:

http://www.ibm.com/support/knowledgecenter/STLM5A/landing/V3700_welcome.html

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

<http://www.vmware.com/products/vsphere/features/storage-vmotion.html>

For more information about migrating SAN boot images see Appendix A, “Command-line interface setup and SAN Boot” on page 649.

Prerequisites

Confirm that the following prerequisites apply:

- ▶ Make sure that the Storwize V3700, the existing storage system, and host FC ports are physically connected to the SAN fabrics, or that SAS connectivity exists between the Storwize V3700 and the External storage system.
- ▶ If there are VMware ESX hosts involved in the data migration, make sure that the VMware ESX hosts are set to allow volume copies to be recognized. For more information, see the VMware ESX product documentation on the following website:

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

After you have confirmed the restrictions and prerequisites, complete the following steps:

1. If all options *can* be selected, click **Next** to continue.
2. In all other cases, **Next** cannot be selected, and the data must be migrated without use of this wizard.

Figure 6-17 shows the Storage Migration wizard with all restrictions satisfied and prerequisites met. Notice at the bottom of the wizard pane that additional help and an eLearning video are available.

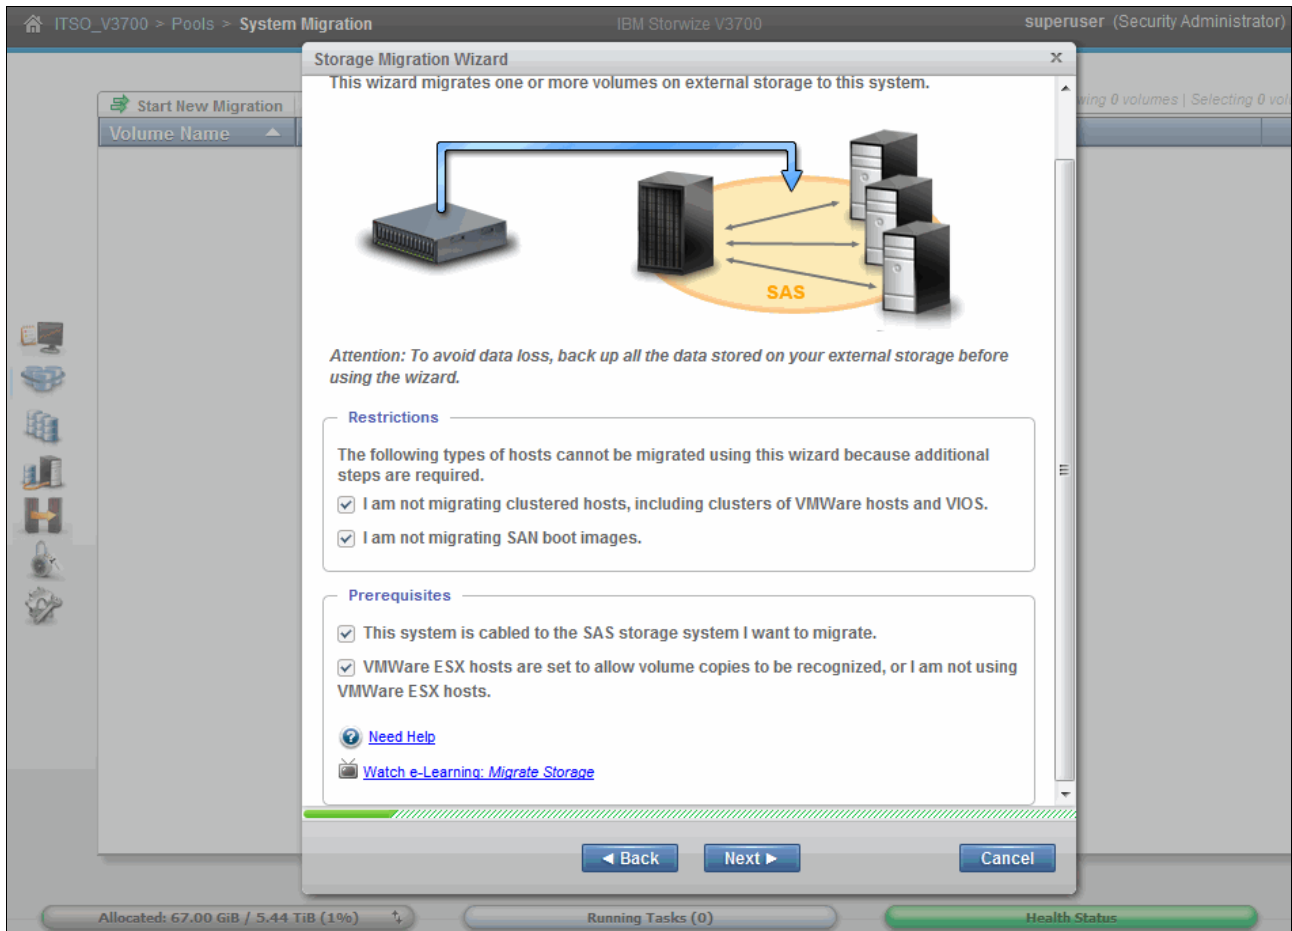


Figure 6-17 Confirming prerequisites

Prepare the environment for migration

Follow the instructions given in this pane carefully:

1. When all of the required tasks are complete, click **Next** to continue. Figure 6-18 shows the Prepare Environment for Migration pane.

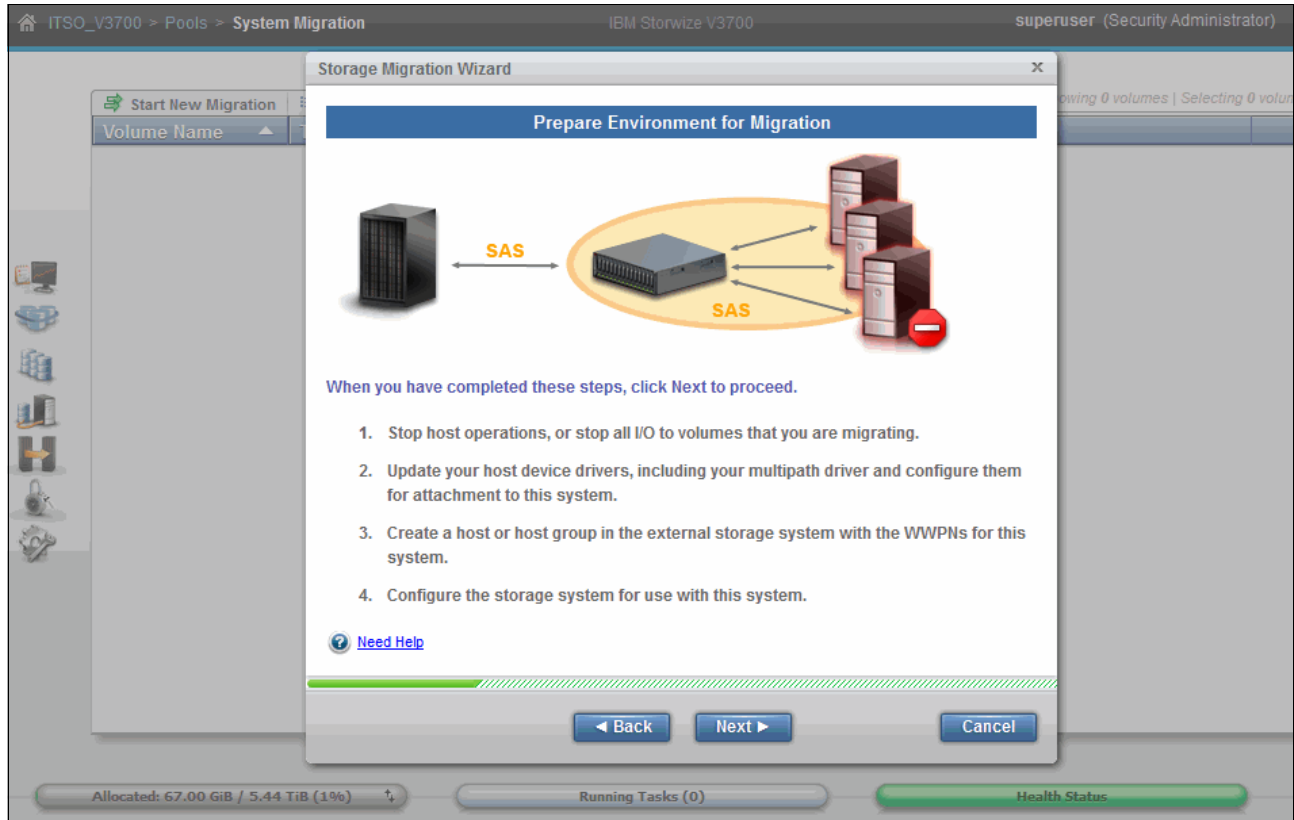


Figure 6-18 Instructions to prepare the migration environment: SAS attached storage

2. Slightly different instructions are given if only FC external storage is attached, or if you have chosen an FC-attached migration, as shown in Figure 6-19.




Figure 6-19 Instructions to prepare the migration environment: FC attached storage

3. For this example, there is SAS-attached storage, so you must complete the following steps:
 - a. Stop host operations, or stop all I/O to volumes that are being migrated.
 - b. Update host device drivers, including the multipath driver, and configure them for attachment to the IBM Storwize V3700.
 - c. Create a host or host group in the external storage system with the WWPNs for the Storwize V3700 system.
 - d. Configure the storage system for use with the Storwize V3700.

Stop host operations or stop all I/O to volumes that are being migrated

To stop host operations or volume I/O, complete the following steps:

1. To stop host I/O, stop all of the applications using the volumes and set them offline. From the Windows 2012 Server Manager Dashboard select **Tools** → **Computer Management** → **Storage** → **Disk Management**, as shown in Figure 6-20.




Figure 6-20 Offlining DS3200 volumes from Windows 2012

Next we must remove the volumes from the Windows 2012 host. To do this, we must first unmap from the host definition in the external storage system. In this case, this is the Storage Manager GUI controlling the DS3200.

- From the Subsystem Management summary window, click the **Modify** tab and select **Edit host-logical-drive mappings**, as shown in Figure 6-21.

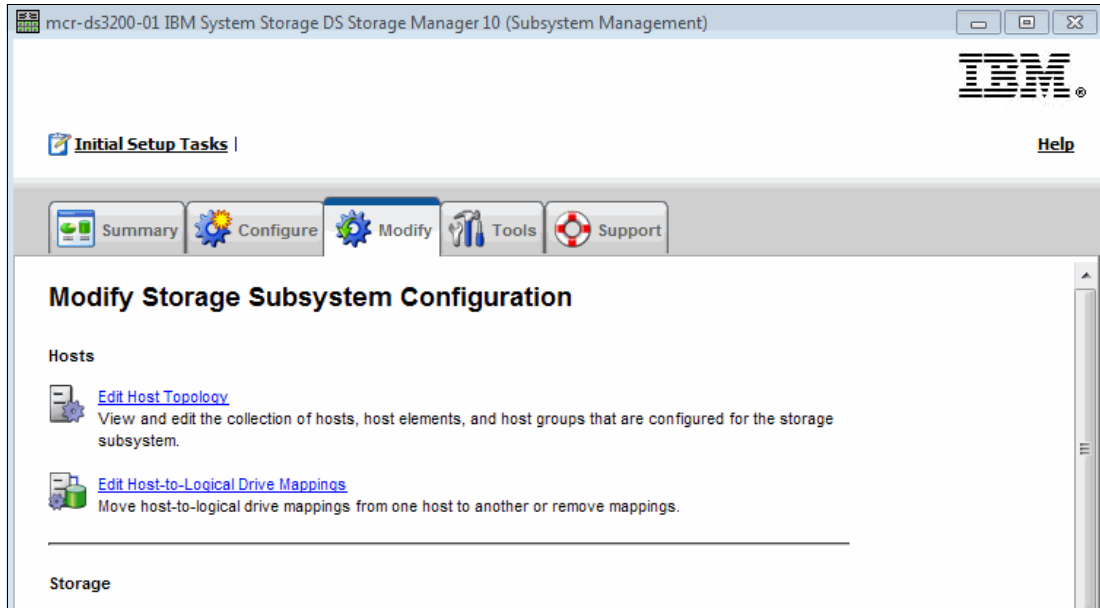


Figure 6-21 Editing host/volume mappings for IBM DS3200

- A list of volume mappings to hosts are displayed. Select the volume that you want to delete the mapping from, then click **Remove**, as shown in Figure 6-22.

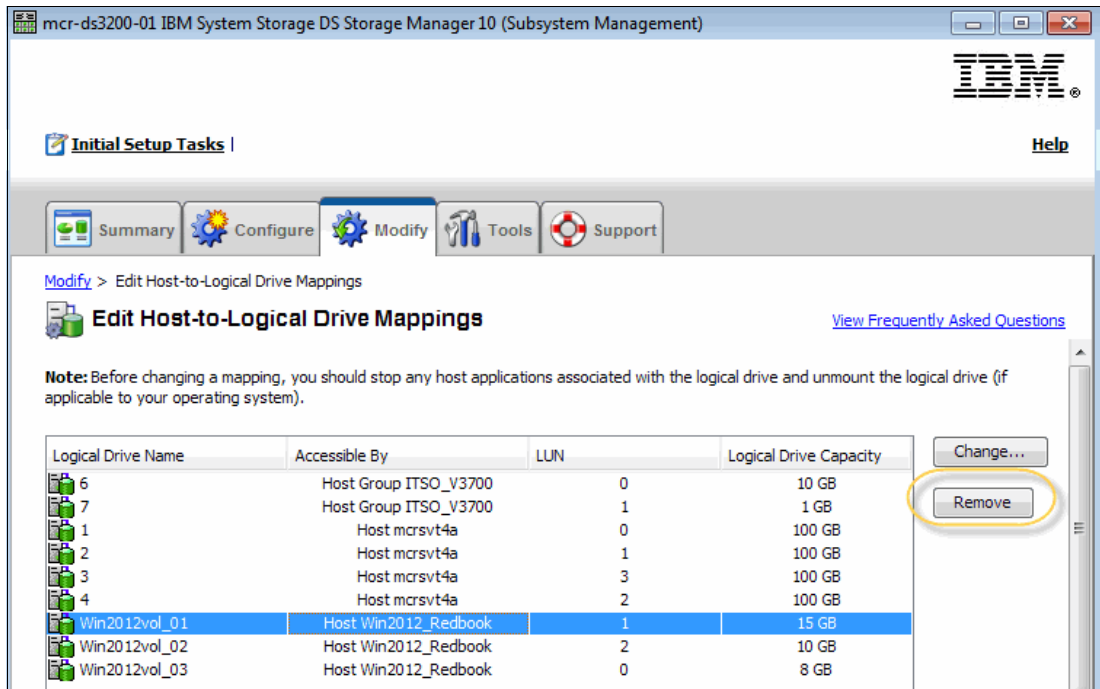


Figure 6-22 Removing a host volume mapping: IBM DS3200

4. A warning message is presented, as shown in Figure 6-23.

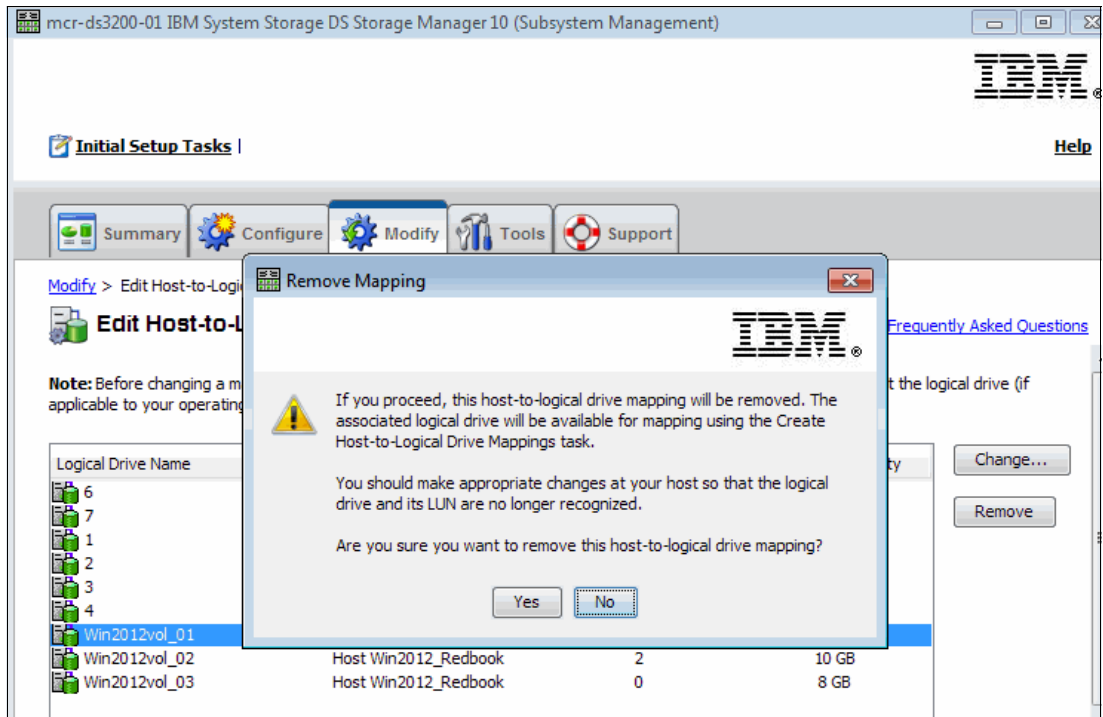


Figure 6-23 Host mapping warning: IBM DS3200

5. Select **Yes** to remove the mapping.

- After all of the volumes have been unmapped, return to the host and perform a drive rescan. For a Windows 2012 host, go to the Windows 2012 Server Manager Dashboard, select **Tools** → **Computer Management** → **Storage** → **Disk Management**, and select **Rescan Disks** from the **Actions** menu, as shown in Figure 6-24.

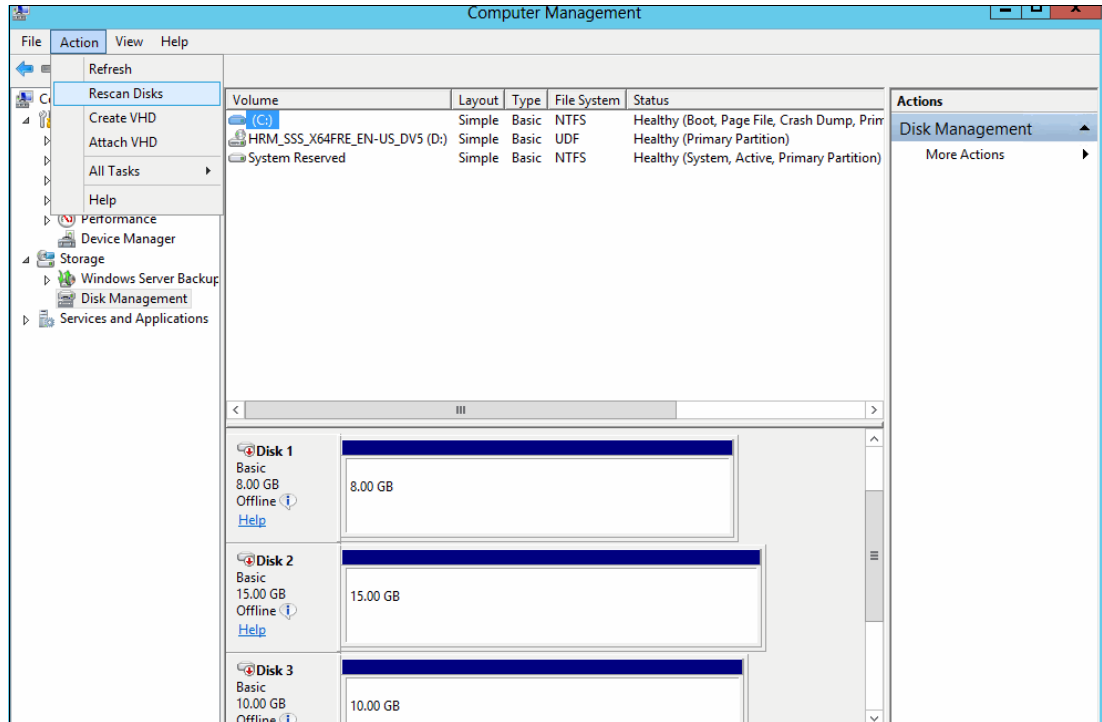


Figure 6-24 Performing a drive rescan: Windows 2012

The volumes no longer display in the windows configuration.

Update host device drivers and configure fhost to access Storwize V3700

The three volumes are currently under the control of the native MPIO driver with the added IBM DSM for System Storage DS3x00. This needs to be removed from the volumes that you are going to migrate. They then need to be put under the control of IBM SDDDSM for use with the Storwize V3700.

To update the drivers and configure fhost, complete the following steps:

1. To remove the existing MPIO from the server Manager Dashboard, select **Tools** → **MPIO**. The MPIO properties pane opens, as shown in Figure 6-25.




Figure 6-25 Windows 2012 MPIO properties

2. Select the device for which you want support to be removed and click **Remove**. In this example, for clarity, you remove all support for all System Storage DS and IBM InfoSphere® FastTrack (FASTT) devices. This is because there is only one System Storage DS3200. We only need to remove support for the 1726-2xx device (System Storage DS3200).

- This action requires a reboot, and the operating system prompts you to perform one, although you can defer it, as shown in Figure 6-26.

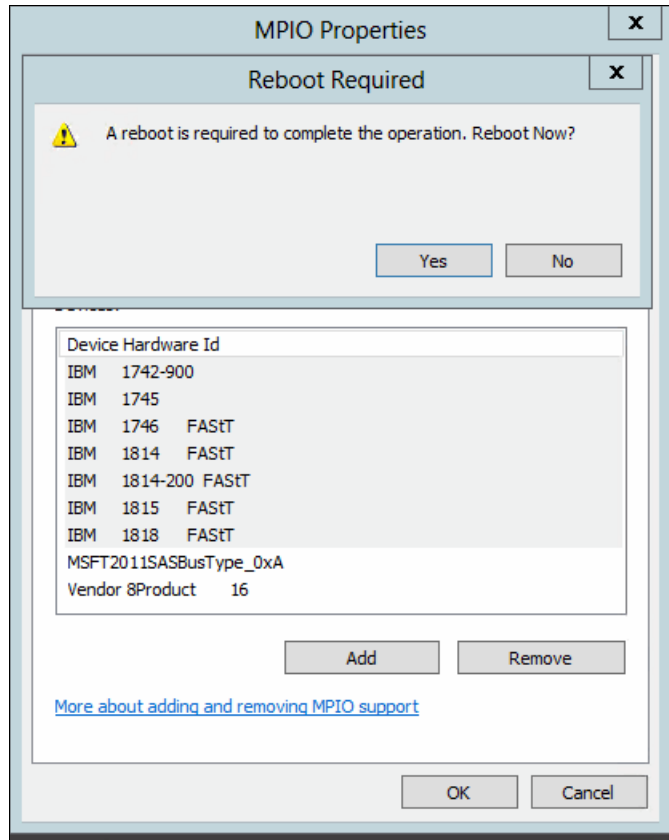


Figure 6-26 Windows 2012 reboot prompt

Next you must install MPIO support for the IBM Storwize V3700. The IBM Storwize family uses the IBM SDDDSM for MPIO support in Windows 2012. This can be obtained from the following website:

http://www.ibm.com/support/docview.wss?rs=540&context=ST52G7&dc=D430&uid=ssg1S4000350&loc=en_US&cs=utf-8&lang=en#Storwize3700

- Unpack the file to your host server and run the **setup** file, as shown in Figure 6-27. When prompted, type yes to install the device driver.




Figure 6-27 Installing IBM SDDDSM: Windows 2012

- In this case MPIO is already installed, so only the DSM is needed. When complete, the installation prompts you to reboot your host to activate the new device driver, as shown in Figure 6-28.

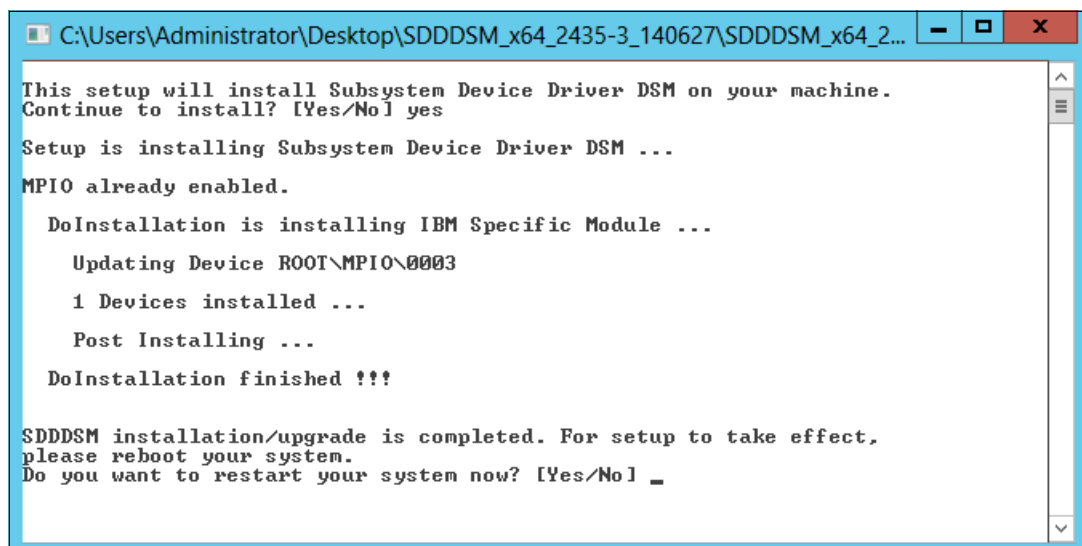


Figure 6-28 IBM SDDDSM installation complete: Windows 2012

6. Type yes to reboot your host and activate the driver. When the system restarts, go into the Server Manager Dashboard, select **Tools** → **MPIO**, and view the MPIO drivers, as shown in Figure 6-29.




Figure 6-29 Checking MPIO support: Windows 2012

As can be seen in the example, the DSM has been loaded to support the Storwize range of Storage subsystems.

Create a host or host group in the external storage system

To create a host or host group, complete the following steps:

1. Using IBM System Storage DS Storage Manager, connect to the external DS3200 system as shown in Figure 6-4 on page 254, and select **Manage a Storage Subsystem**.
2. Select the system that you want to manage and click **OK**.
3. Each canister of the IBM Storwize V3700 needs to be defined on the System Storage DS3200 as a host, and then added to a host group. Click the **Configure** tab and select **Configure Host Access (Manual)**, as shown in Figure 6-30.

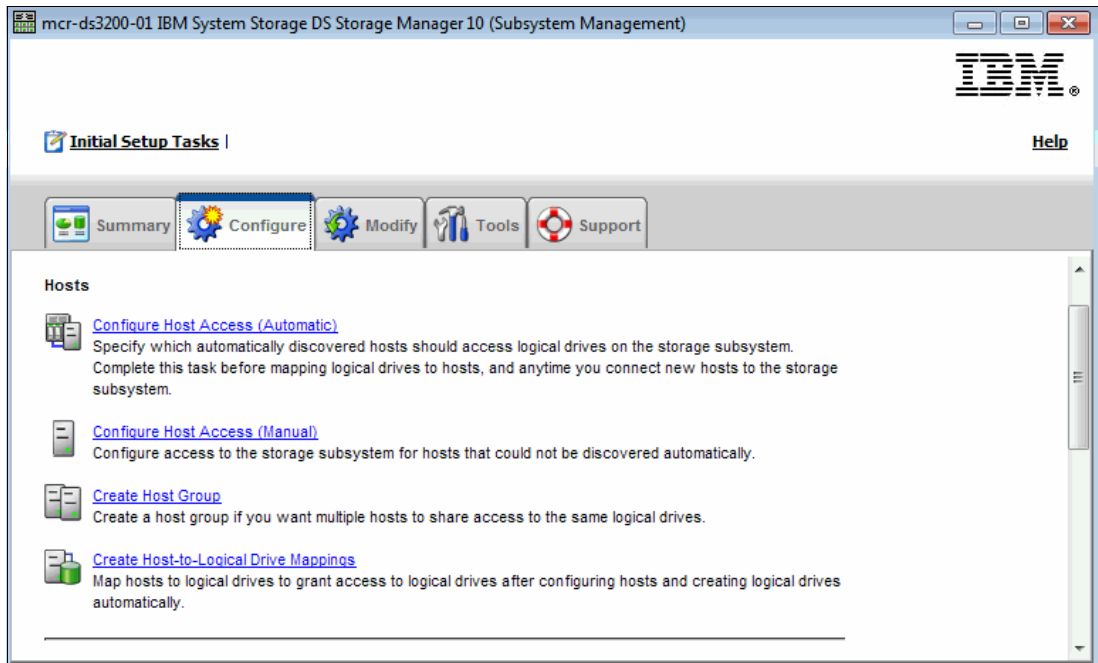


Figure 6-30 Creating hosts: Storage Manager GUI

4. A wizard starts to guide you through the creation of each host. Enter the host name and define the host type. For IBM Storwize V3700, use **IBM TS SAN VCE**, as shown in Figure 6-31.
5. Click **Next**.

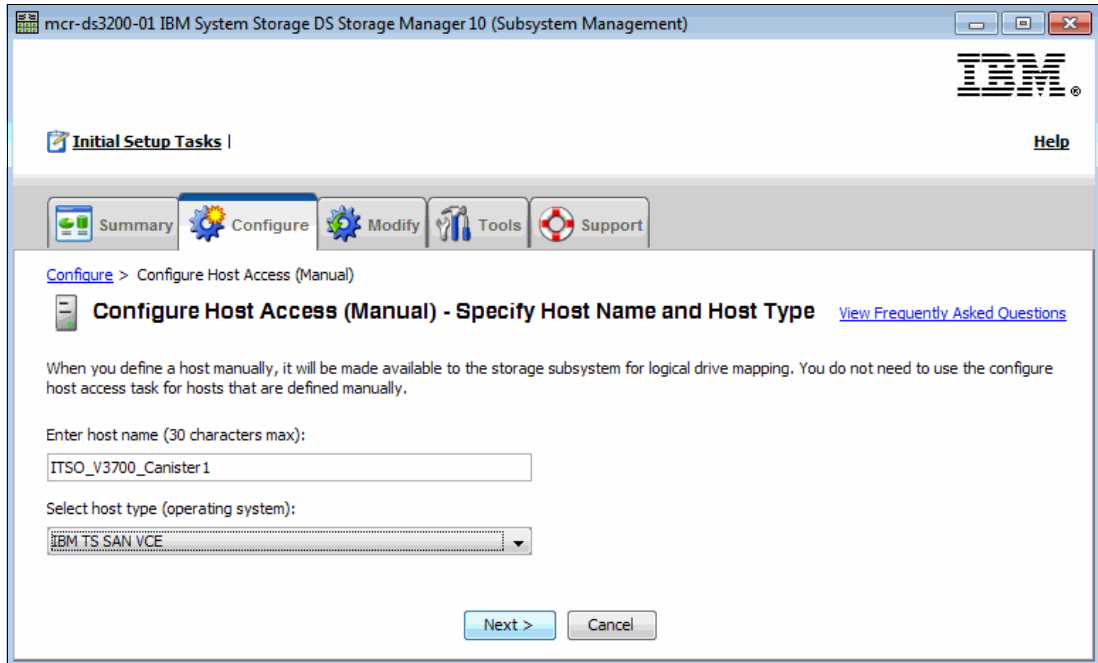


Figure 6-31 Defining a host: Storage Manager GUI

6. In the next pane, assign the V3700 canister WWPN to the host definition, as shown in Figure 6-32. Use the data gathered in 6.2.5, "Recording migration information" on page 252.

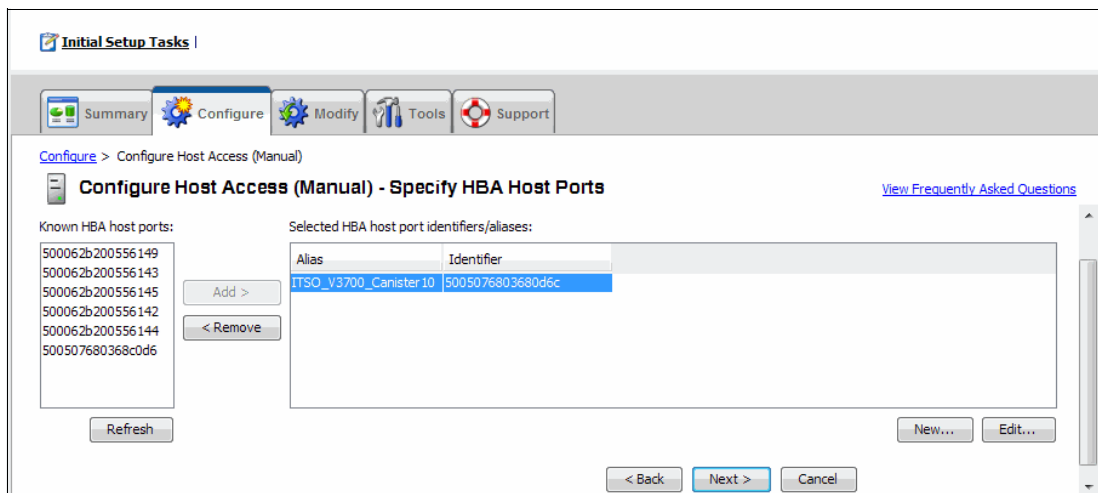


Figure 6-32 Adding WWPN to the host definition

- Return to the Storage Manager GUI and continue creating the Storwize V3700 hosts. Click **Next** to add the host to a new host group, as shown in Figure 6-33. In this example, call the host group ITS0_V3700.

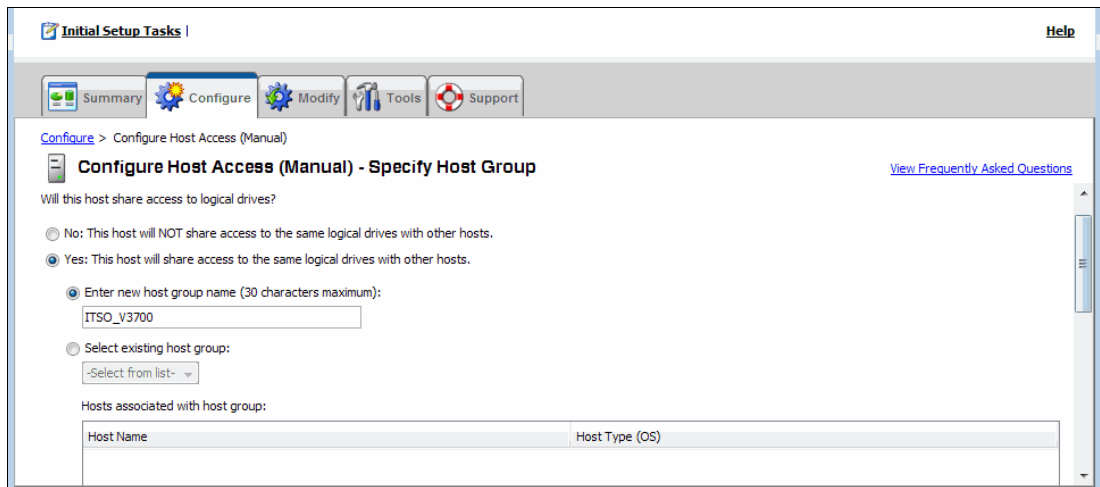


Figure 6-33 Adding host to host group: Storage Manager GUI

- Click **Next** and a confirmation window opens. If everything is correct, click **Finish** to create the host and add it to the host group, as shown in Figure 6-34.

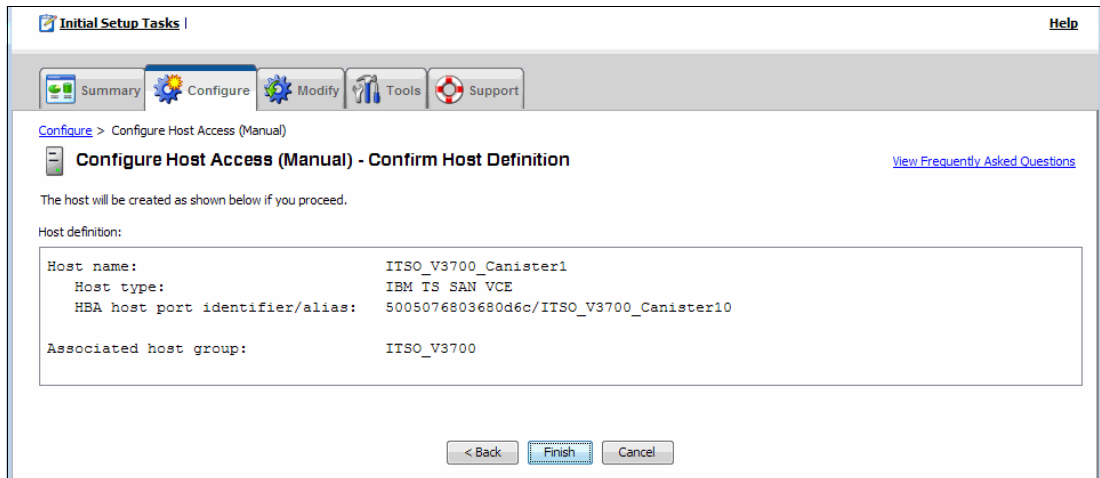


Figure 6-34 Confirming host creation: Storage Manager GUI

The wizard completes and prompts you whether a second host needs to be created.

9. Click **Yes** and follow the steps from Figure 6-31 on page 275 to Figure 6-33 on page 276. When adding the second host to a host group, this time select an existing host group, as shown in Figure 6-35.

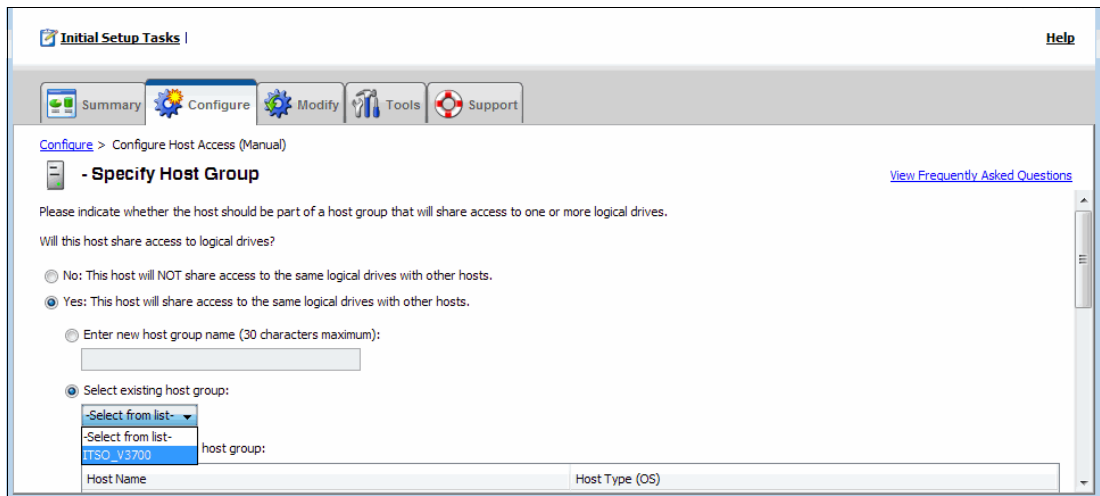


Figure 6-35 Adding second host to existing host group: Storage Manager GUI

10. Click **Next** and **Finish** to complete the task. This time, select **No** when prompted to create further hosts.

Configure the storage system for use with the Storwize V3700

To configure the storage system, complete the following steps:

1. Recable or rezone your external storage system to connect to the IBM StorwizeV3700.
2. Now click **Next**, as shown in Figure 6-18 on page 264, to move to the next stage of the migration.

Map storage

The next pane, shown in Figure 6-36, directs you to record information about the LUNs to be migrated, their hosts, the host adapter WWPNs, and the SCSI IDs. It also prompts you to unmap those volumes to be migrated from their hosts, and to remap them to the new host group definition for the Storwize V3700.

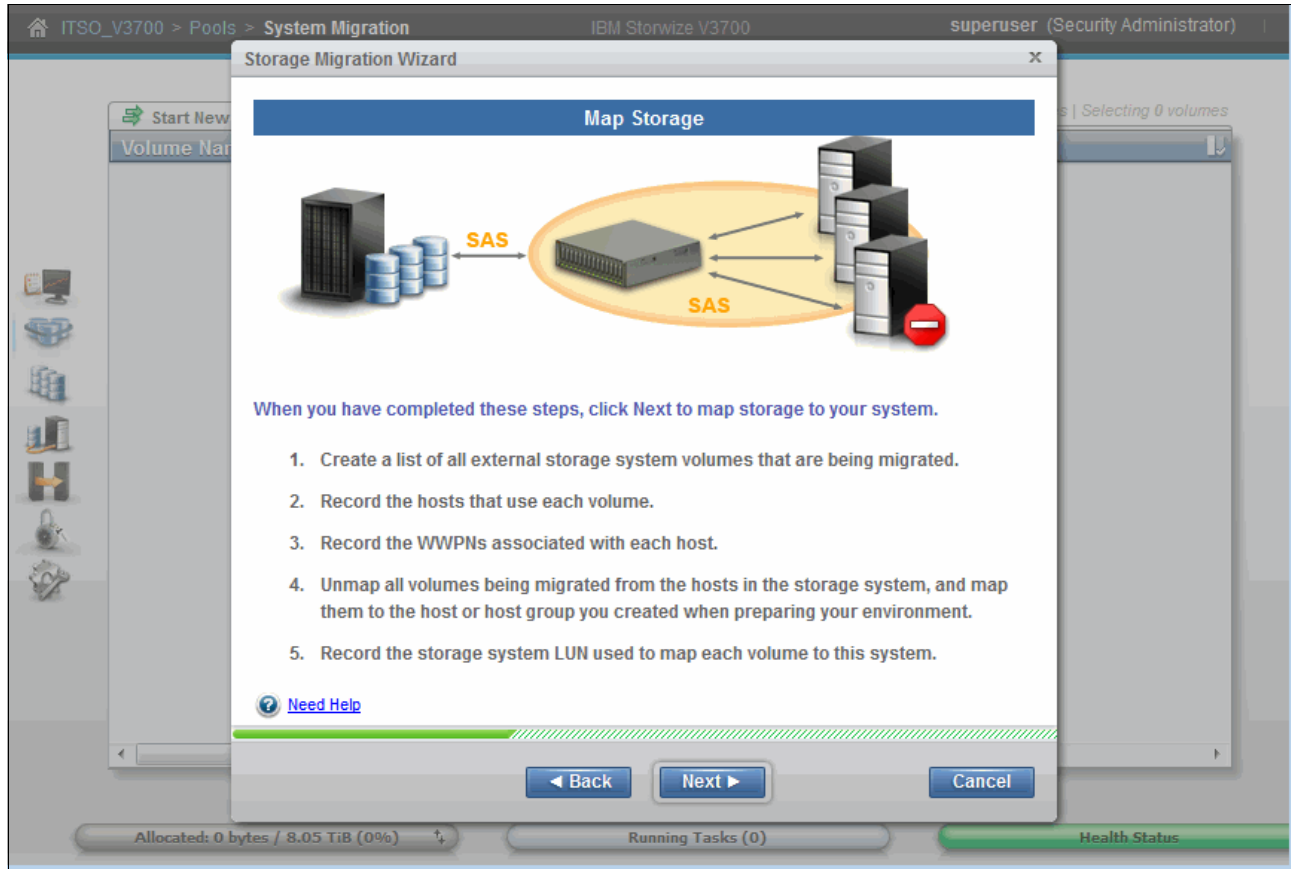


Figure 6-36 Directions to record volume migration data

In this example scenario, you have already collected this data in 6.2.5, “Recording migration information” on page 252. In addition, in the last section, you unmapped the volumes on the external IBM DS3200 storage system from their Windows 2012 host. Therefore, all that remains for you now is to remap those external LUNs to the new Storwize V3700 host definition in the IBM System Storage DS3200 Storage Manager GUI.

To remap the external LUNs, complete the following steps:

1. Connect to the external DS3200 system, as shown in Figure 6-4 on page 254, and select **Manage a Storage Subsystem**.
2. Select the system that you want to manage and click **OK**.
3. Click the **Configure** tab and select **Create Host-to-Logical Drive Mappings**, as shown in Figure 6-37.

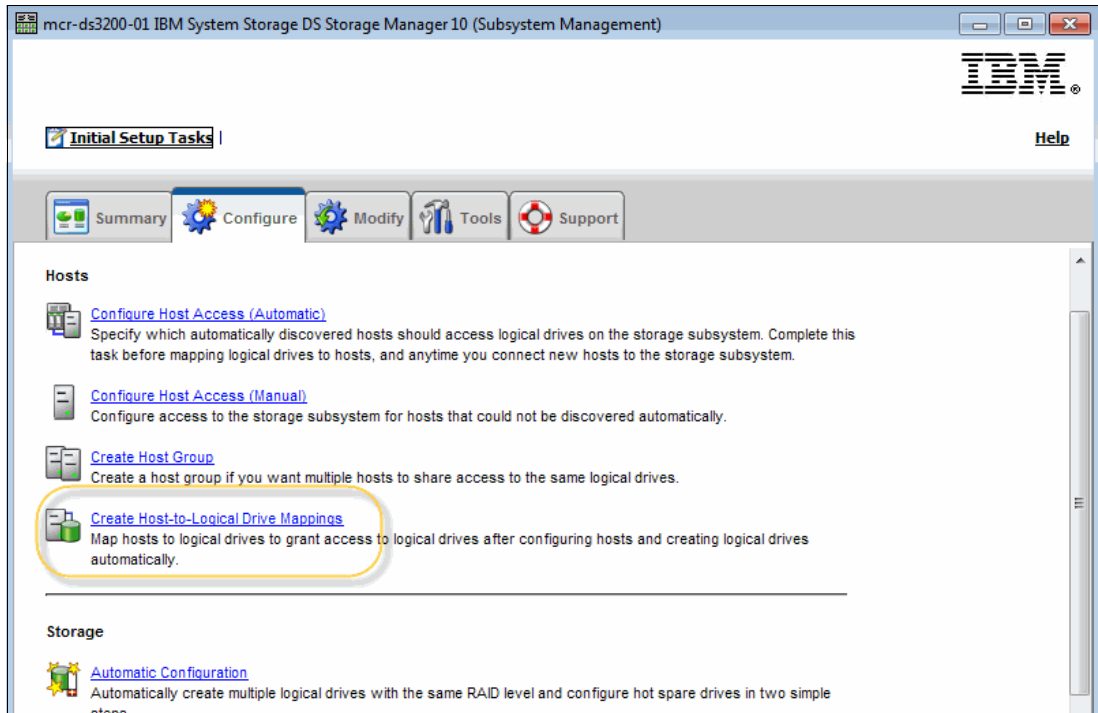


Figure 6-37 Create host mappings: IBM DS3200

4. A wizard starts, to guide you through the steps of creating the new host mappings. Select the new Storwize V3700 host group from those displayed and click **Next**, as shown in Figure 6-38.

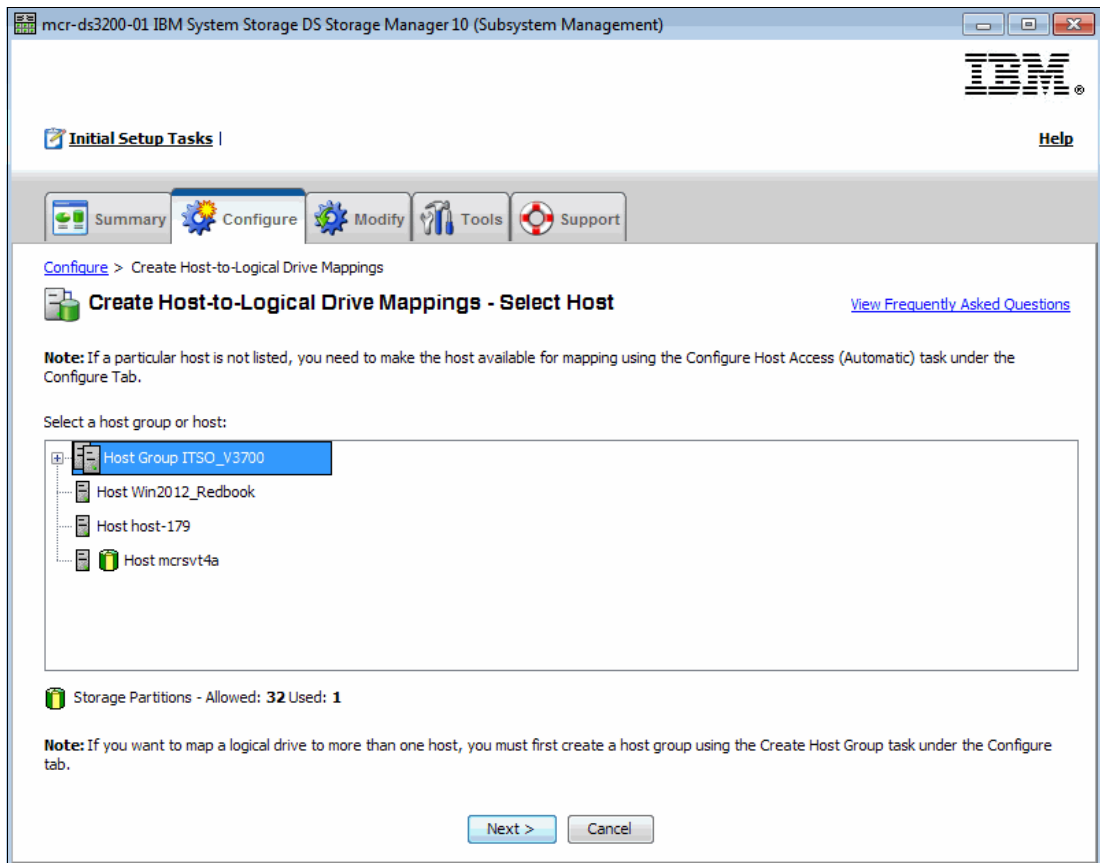


Figure 6-38 Selecting new Storwize V3700 host group: IBM DS3200

- Select the volume to be mapped to the host and, using the data recorded in 6.2.5, “Recording migration information” on page 252, configure the correct SCSI ID for the volume. Click **Finish** to make the mapping, as shown in Figure 6-39.

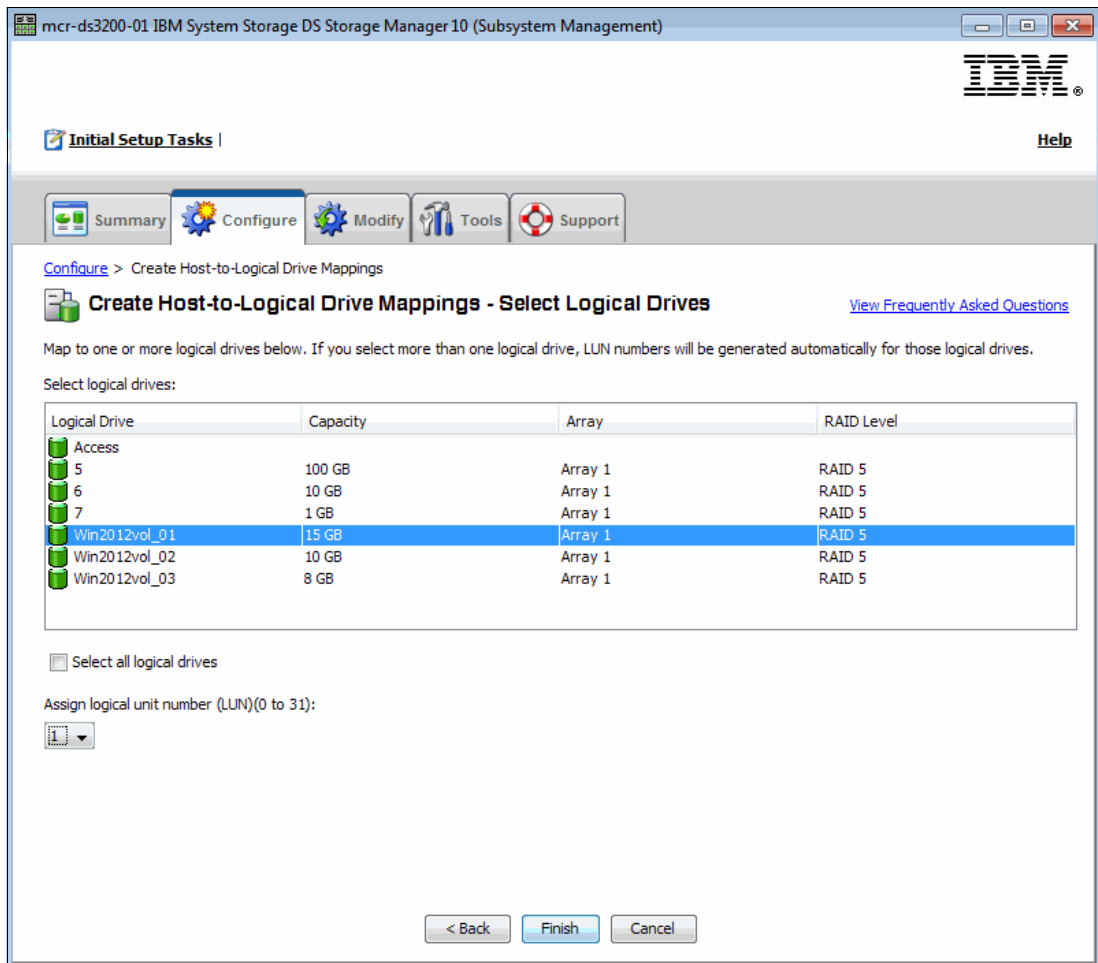


Figure 6-39 Selecting volume to map to Storwize V3700 host: IBM DS3200

6. A confirmation box opens to show that the mapping has been created, as shown in Figure 6-40.




Figure 6-40 Mapping complete: IBM DS3200

7. Click **OK**, and you are then prompted to create another mapping, as shown in Figure 6-41.

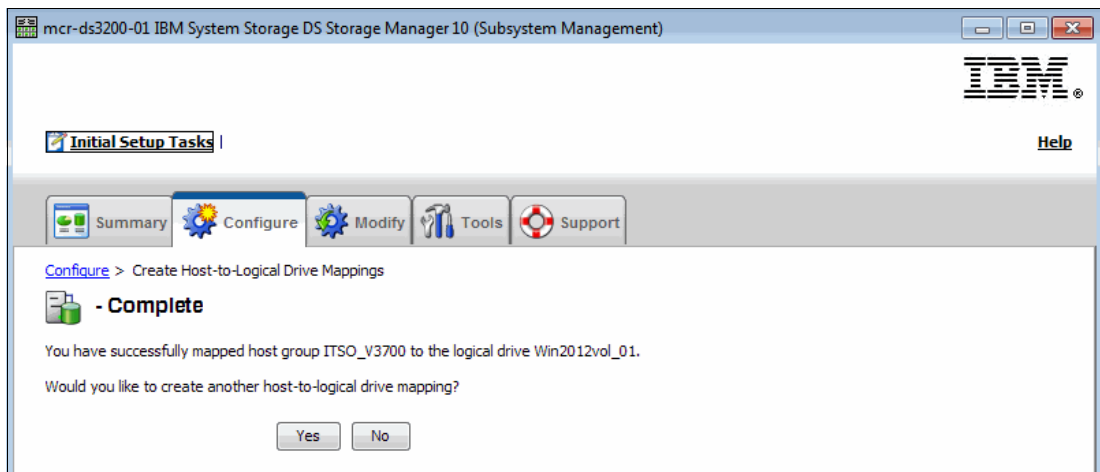


Figure 6-41 Creating further host-volume mappings: IBM DS3200

- Click **Yes** if you have more volumes to map to the new Storwize V3700 host definition, and repeat the steps from Figure 6-38 on page 280. Click **No** if there are no more volumes to be mapped to exit the wizard.

The following tasks have been completed:

- All of the existing external volumes have now been mapped to your new Storwize V3700 host definition.
 - Cabling or zoning has been changed to ensure that the external storage system can connect to the Storwize V3700, and that hosts can connect to the Storwize V3700.
 - Volumes have been removed from the host configuration.
 - MPIO drivers have been updated to reflect the impending connection to the Storwize V3700.
- We are now in a position to move to the next stage of the Storwize V3700 Data Migration wizard. Click **Next** to continue, as shown in Figure 6-36 on page 278.
 - The Storwize V3700 runs the Discover Devices task. After the task is complete, click **Close** to continue. Figure 6-42 shows the results of the Discover Devices task. You can see in the background that it has successfully detected MDisks 9, 10, and 11. These are the three volumes that were presented from the System Storage DS3200 to the Storwize V3700 host group definition.

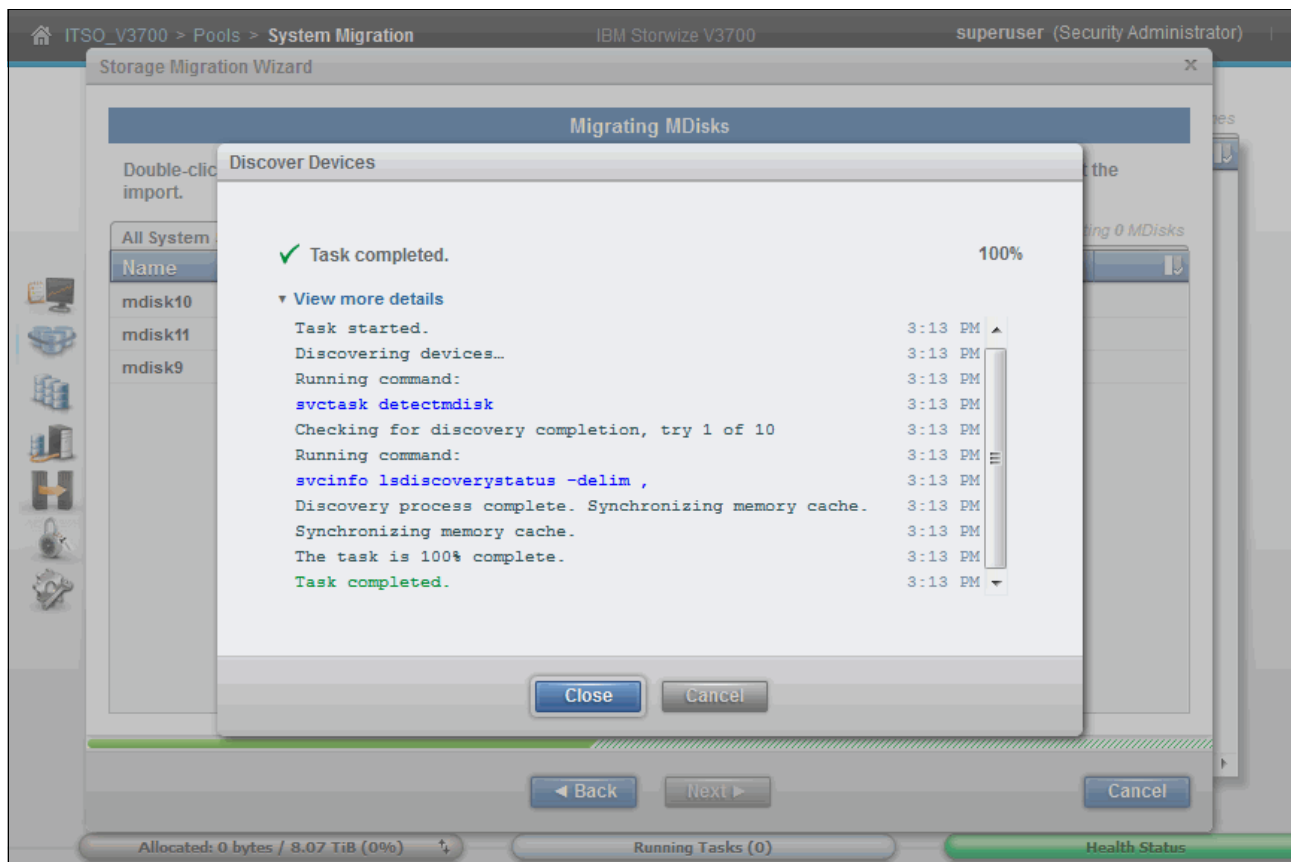


Figure 6-42 Discover Devices task

Migrating MDisks

To migrate the MDisks, perform the following steps:

1. In the next step of the wizard, select the MDisks that are to be migrated and then click **Next** to continue. Hold the Ctrl key down while clicking MDisks to highlight more than one MDisk for migration. Figure 6-43 shows the Migrating MDisks pane.

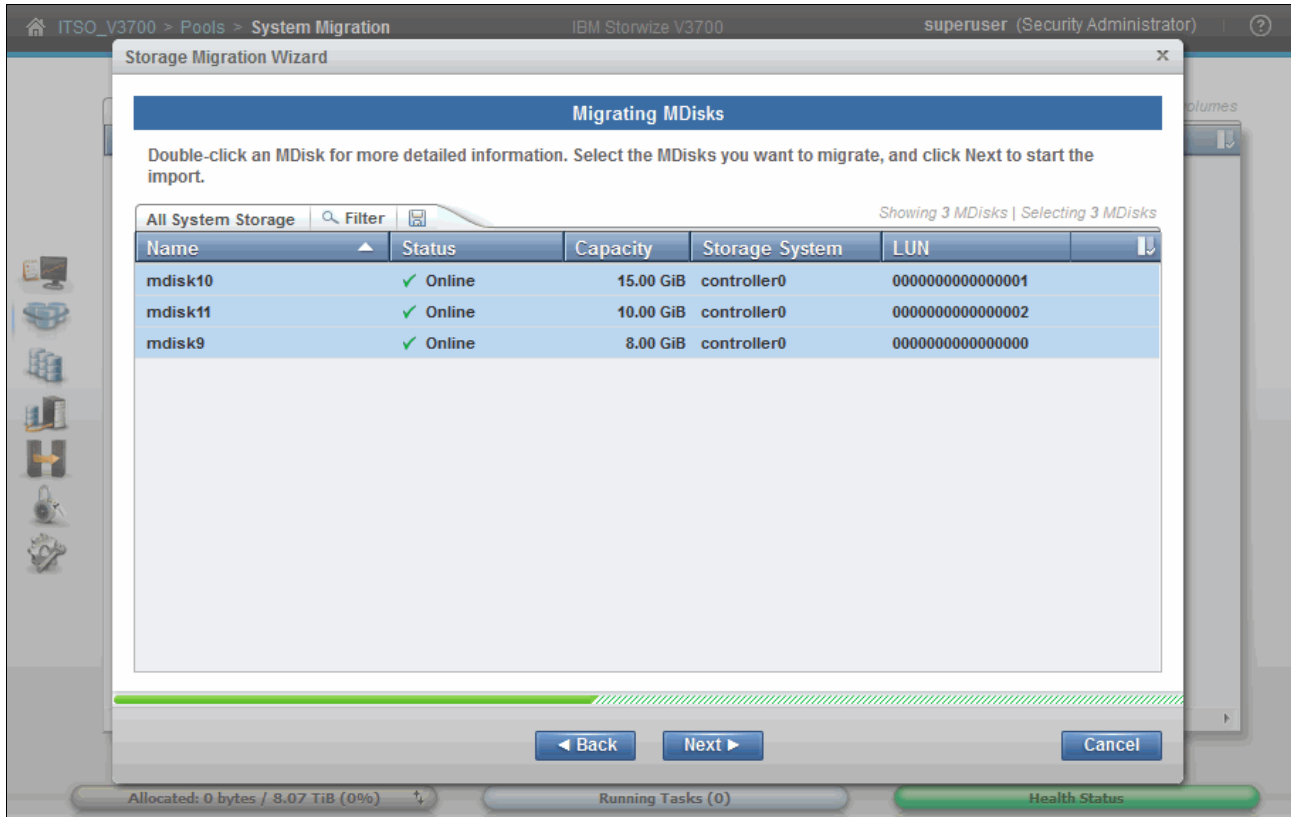


Figure 6-43 Migrating MDisks pane

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

2. The Storwize V3700 runs the Import MDisks task. After the task is complete, click **Close** to continue. Figure 6-44 shows the result of the Import MDisks task.

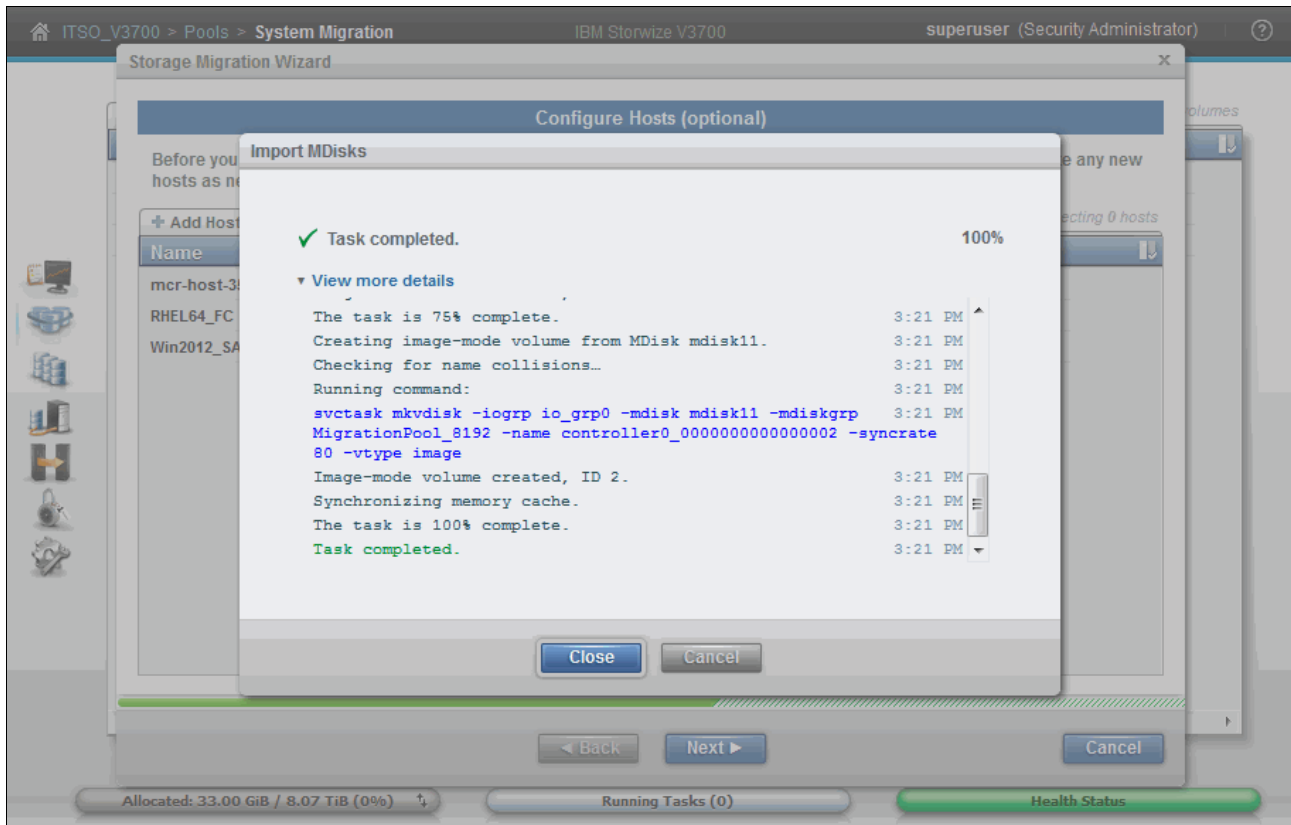


Figure 6-44 Import MDisks task

Configure hosts

Optional: This step is optional. You can bypass it by selecting **Next** and moving to “Map volumes to hosts” on page 291.

Follow this part of the wizard to select or configure new hosts as required:

1. Figure 6-45 shows the Configure Hosts pane. In this example scenario, there is no host definition for the Windows 2012 host, so define one at this stage.

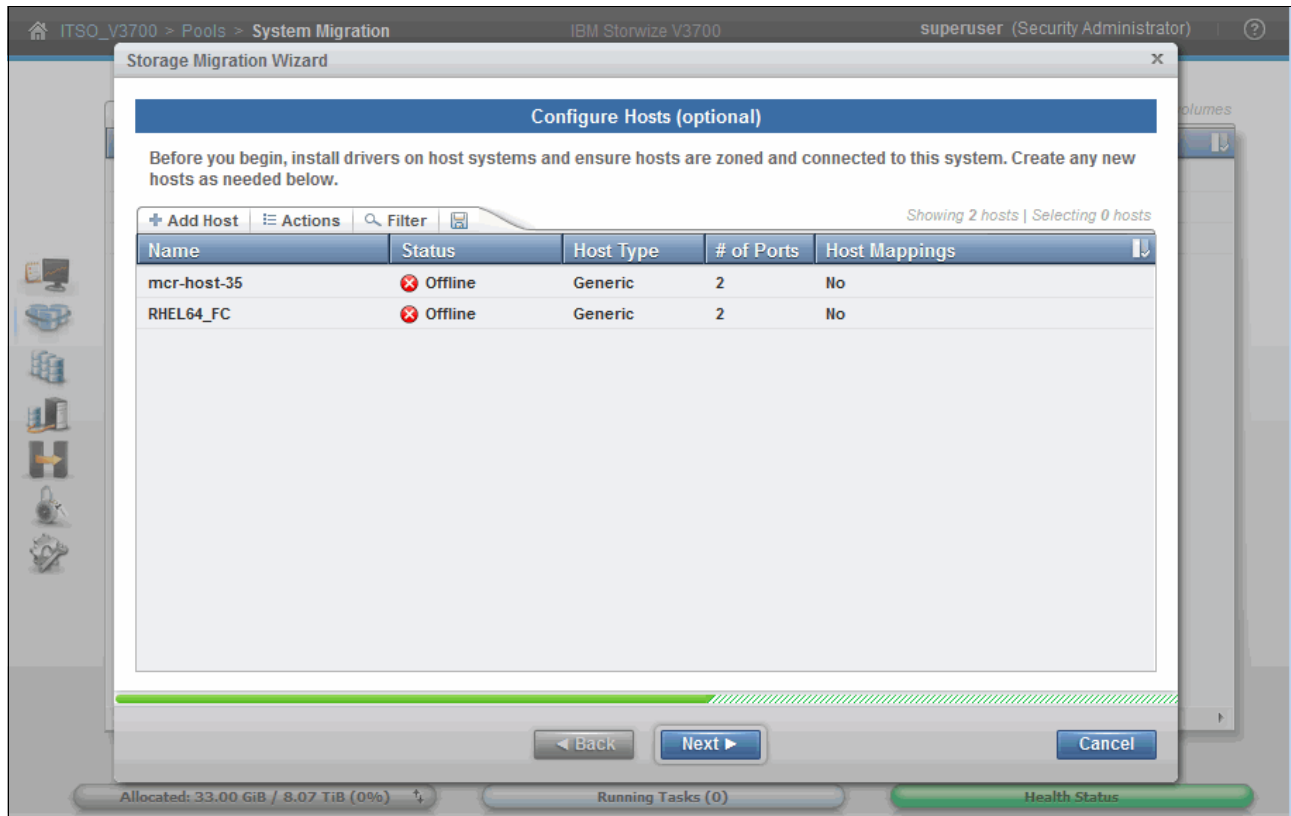


Figure 6-45 Configure hosts pane

If hosts are already defined they show up listed in the pane, as shown in Figure 6-45.

2. In this case, create the host by selecting the **Add Host** option. This begins the process of defining a host, with the choice of connectivity shown in Figure 6-46.




Figure 6-46 Add host connection type

3. Select your connection type, and the wizard moves to the next pane. Here, you can name the host, assign ports (in this case, SAS WWPNs) from the drop-down selection and, in the advanced settings, assign a host type, as shown in Figure 6-47.

In this example, there are two ports listed that the Storwize V3700 can see as host connections.

If there are no connections displayed, check your cabling and zoning to ensure that your host is properly connected to the Storwize V3700 (or any SAN fabric), and click **Rescan**.

4. Select each of the port WWPNs and click **Add port to List** to add them to the host definition.

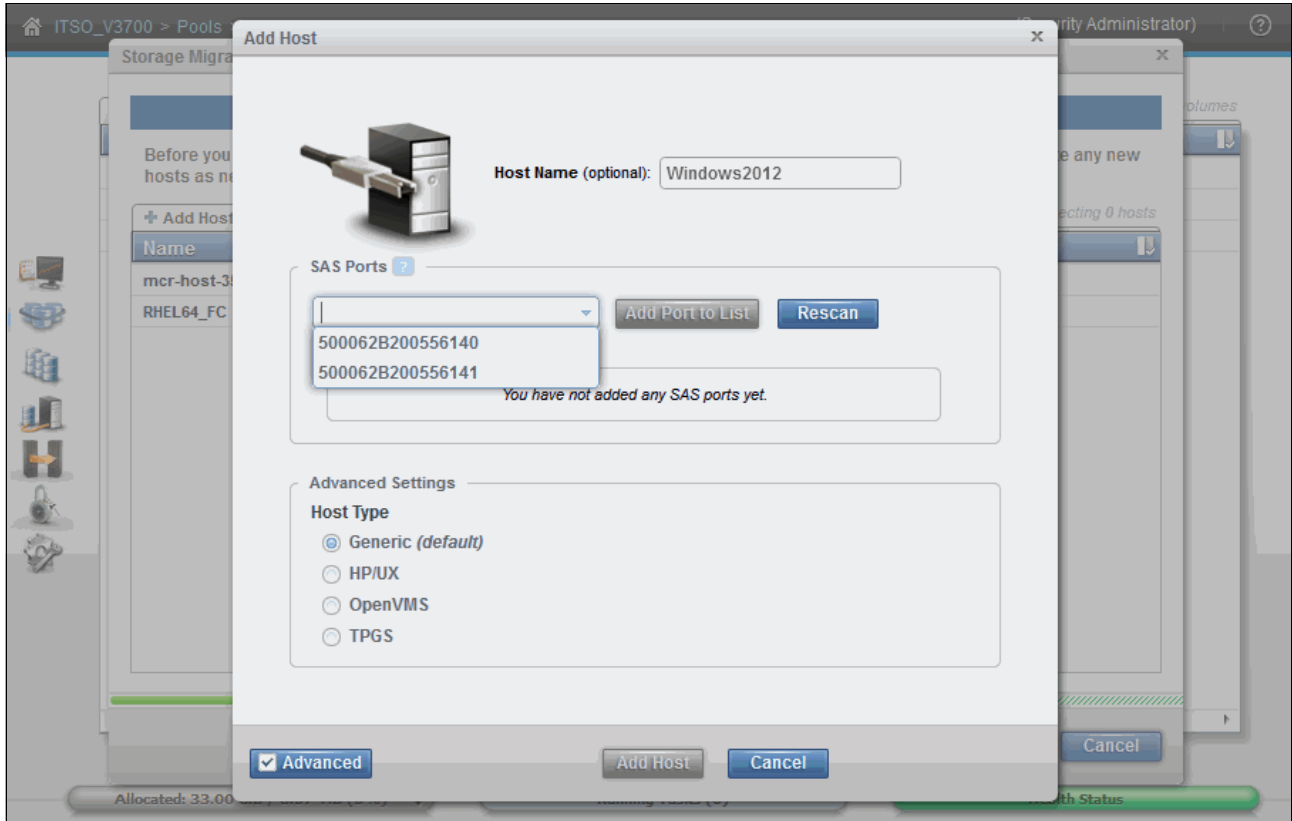


Figure 6-47 Naming host, selecting ports, and setting advanced options

5. Click **Add Host** to complete the task, as shown in Figure 6-48. Notice that the new host now displays in the background list.

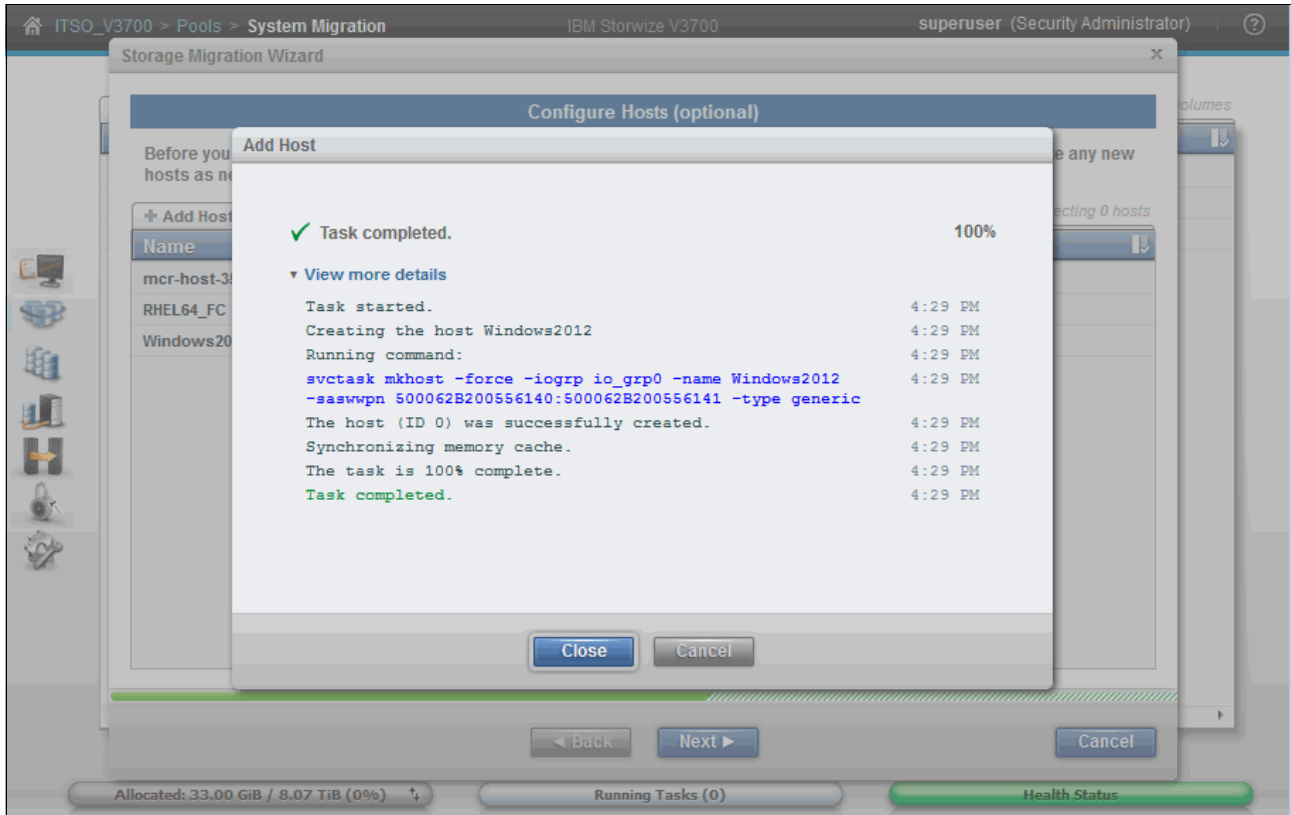


Figure 6-48 Create host task complete

6. Click **Close** and select the new host definition, as shown in Figure 6-49. Click **Next** to move to the next phase of the data migration.

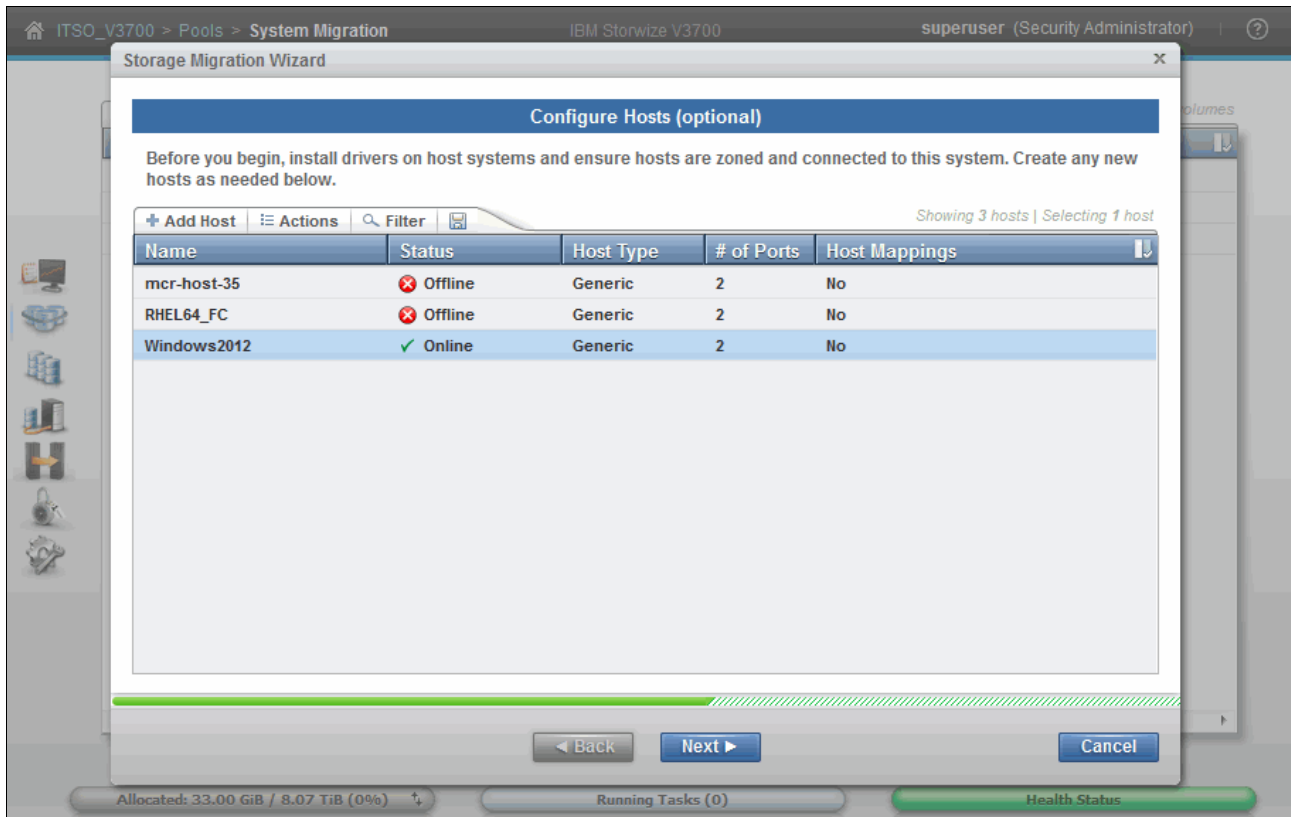


Figure 6-49 Configure host with host listed

Map volumes to hosts

This step of the wizard enables you to select the newly migrated volumes, rename them, and map them to hosts:

1. Start in the Map Volumes to Hosts pane, as shown in Figure 6-50.

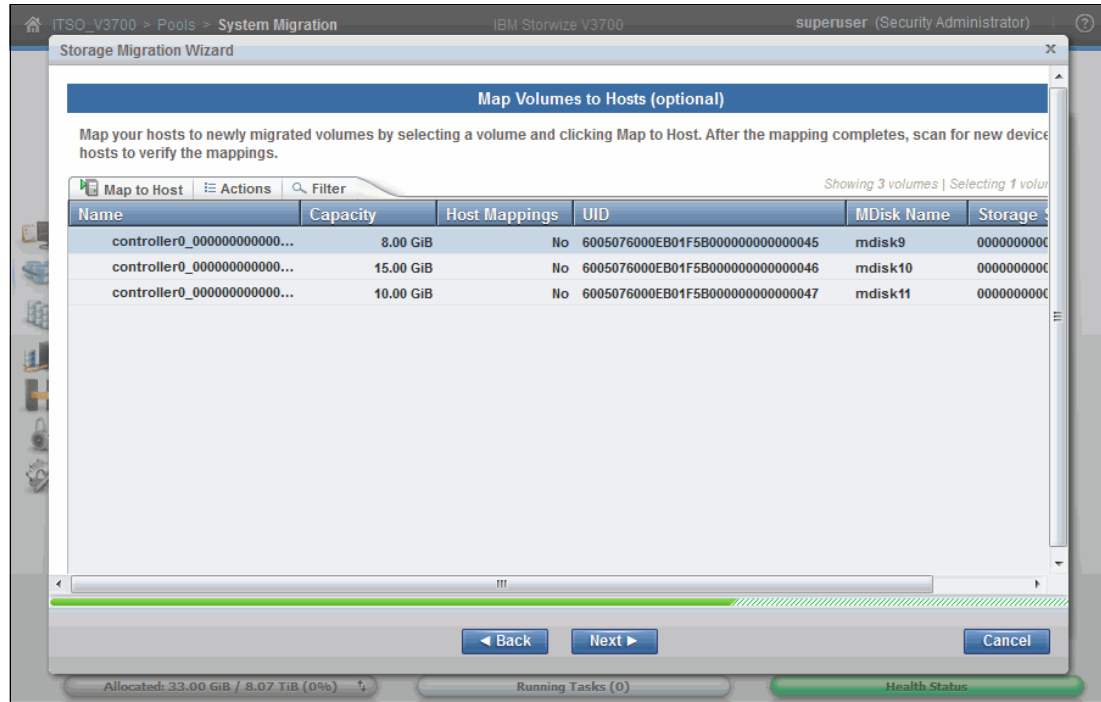


Figure 6-50 Map Volumes to Hosts pane

2. The image mode volumes are listed, the names of which are assigned automatically by the Storwize V3700 storage system. The names can be changed to reflect something more meaningful to the user by selecting the volume and clicking **Rename** in the Actions menu, or by right-clicking the volume and selecting **Rename**.

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

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

The names must not begin or end with a space.

- Use the information collected in 6.2.5, “Recording migration information” on page 252 to rename the volumes to their original names in the external IBM DS3200 Storage array, as shown in Figure 6-51.

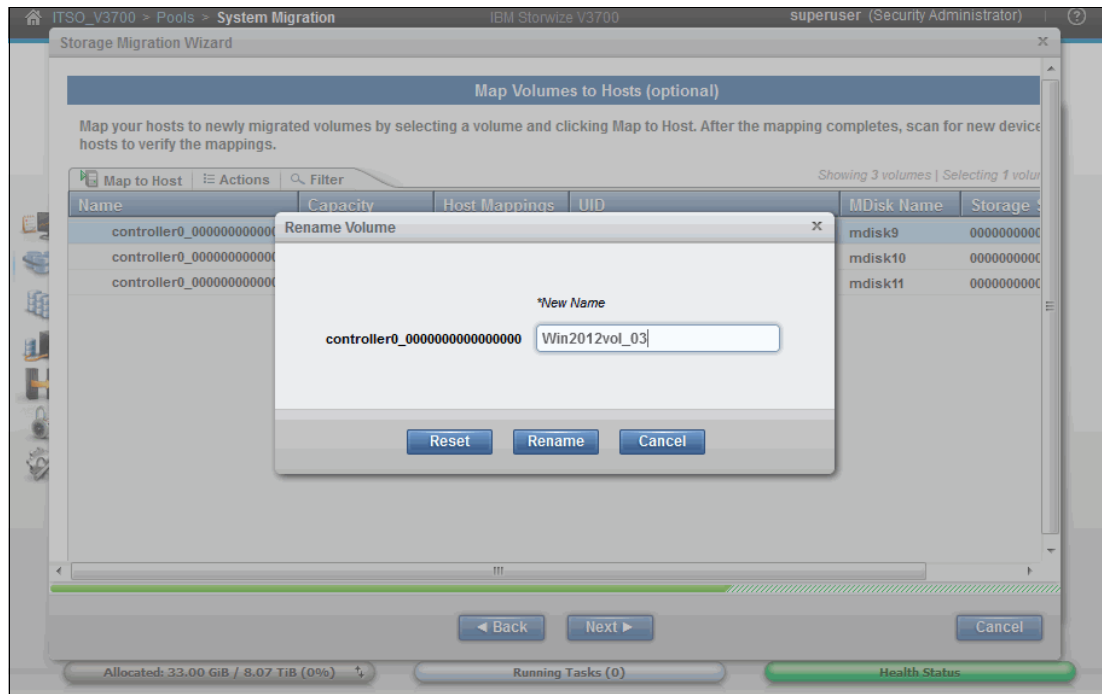


Figure 6-51 Renaming an image mode volume

- Enter the new name and click **Rename** to complete the operation. The task completes, as shown in Figure 6-52.

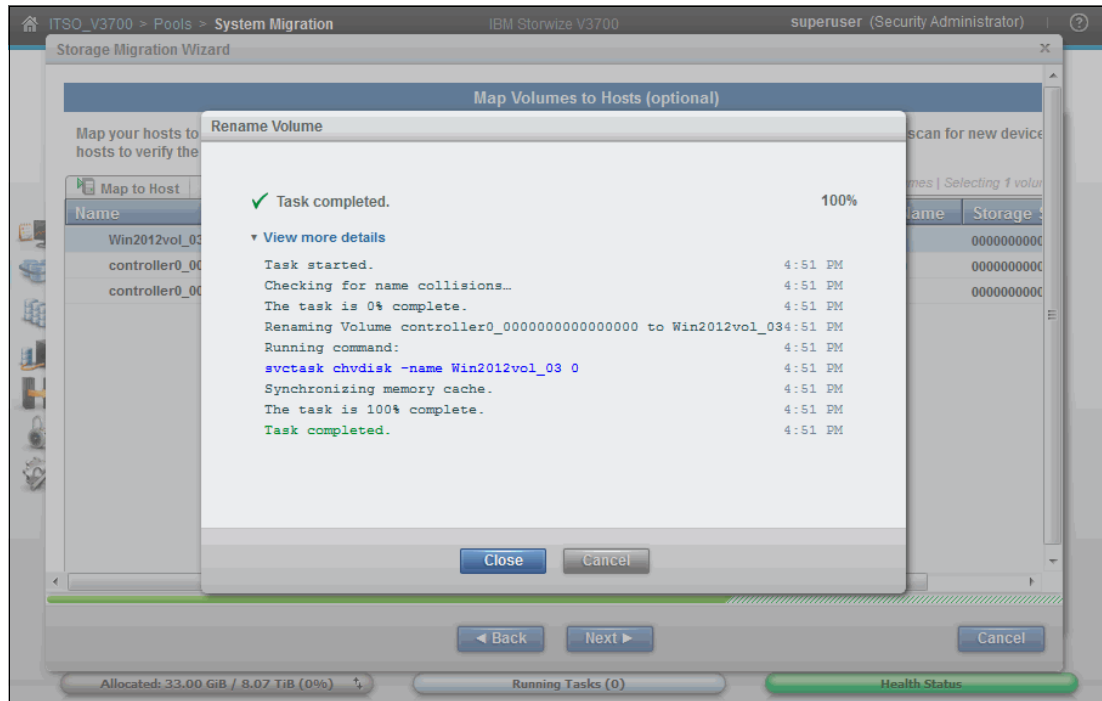


Figure 6-52 Rename task complete

5. Click **Close** to return to the Map Volumes pane as shown in Figure 6-53.




Figure 6-53 Renamed volume

6. Notice that the volume now has a new name. Follow the process from Figure 6-50 on page 291 to rename the rest of the volumes.
7. After all of the volumes have been renamed, select the volumes that you want to map to your host. Hold Ctrl and click to select multiple volumes, and then click **Map to Host**, as shown in Figure 6-54.

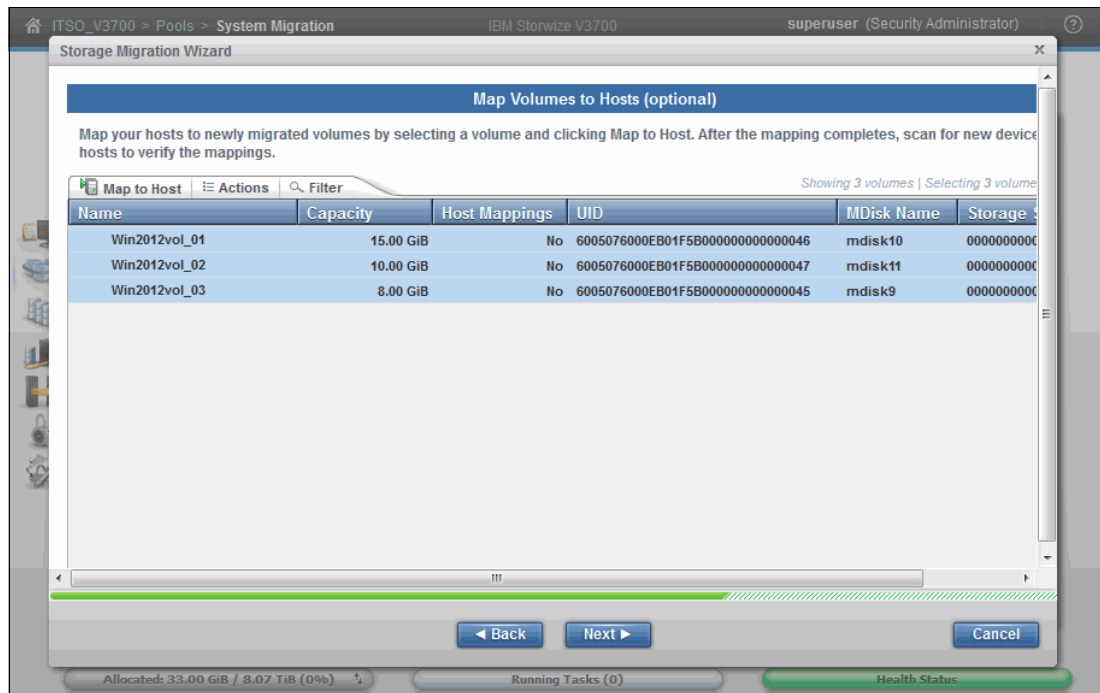


Figure 6-54 Selecting volumes to map

- A box opens to allow you to choose the host to which the volumes are to be mapped, as shown in Figure 6-55. In this example, you have chosen the previously defined Windows 2012 host.




Figure 6-55 Modify host mappings

The Modify Host Mappings pane is now displayed, listing the volumes that you have chosen to be mapped to your particular host, as shown in Figure 6-56 on page 295. The volumes are shown in yellow in the Modify Host Mappings pane. The yellow highlighting means that they are not yet mapped to the host.

- If the SCSI ID was not set correctly when mapping the external volume to the Storwize V3770, as shown in Figure 6-38 on page 280, it can be changed now. Click **Edit SCSI ID** and modify as required. The SCSI ID should reflect the same SCSI ID that was recorded in 6.2.5, “Recording migration information” on page 252.

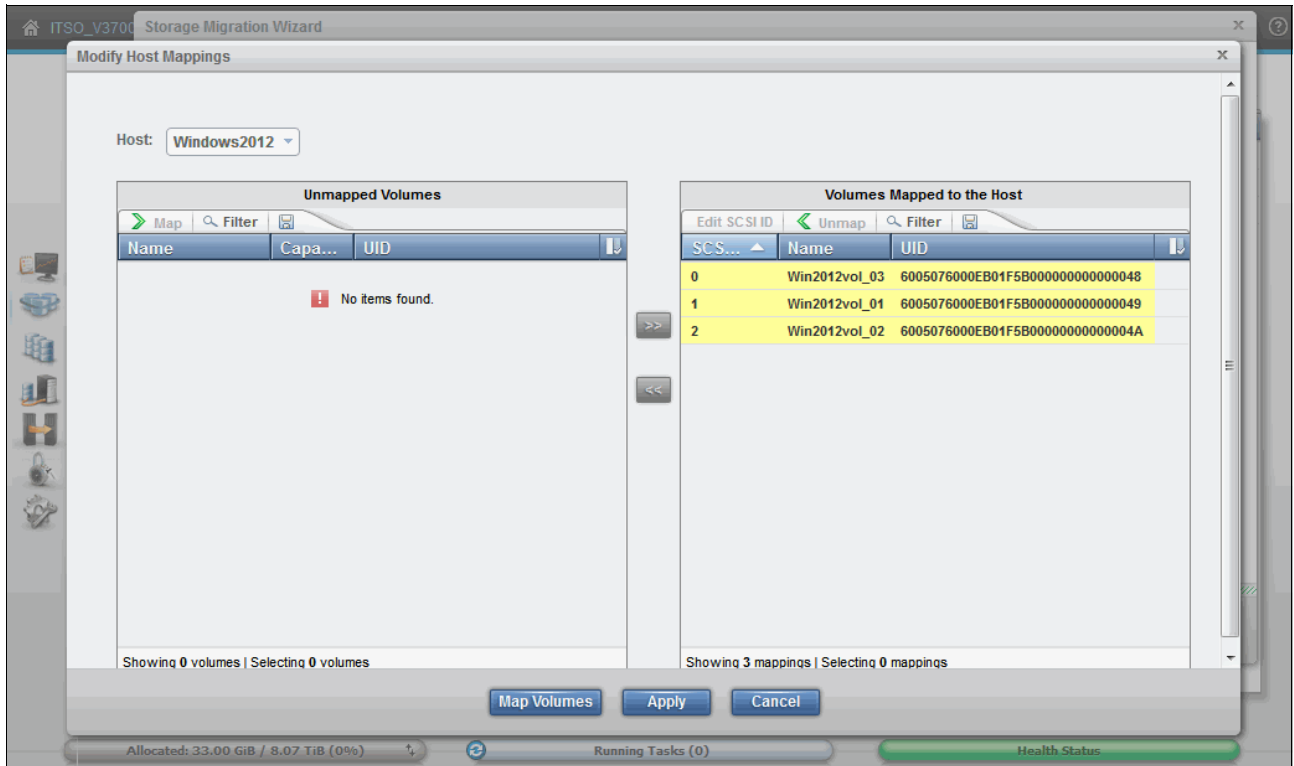


Figure 6-56 Mapping volumes to host

10. If the selections are correct, click **Apply** to complete the task.

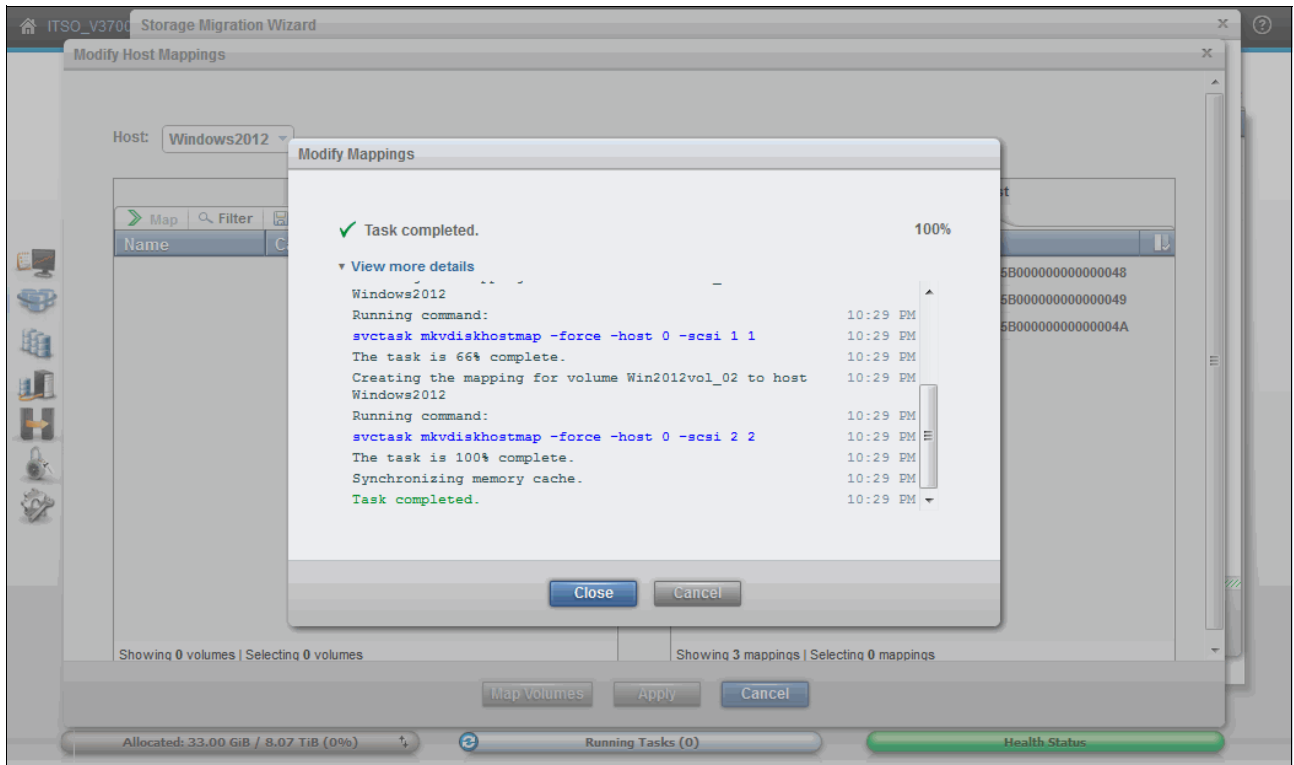


Figure 6-57 Mapping task complete

11. When you click **Close**, the program returns to the Modify Host Mappings pane. If any errors were made, you can unmap volumes at this stage.
12. If everything is correct, use the **Cancel** option, as shown in Figure 6-58. This is normal and is not an indication that the wizard is being canceled. The volumes have been successfully mapped, indicated by the fact that there are no volumes highlighted in yellow, and the only option available is **Cancel**.

If you use the **Map Volumes** button rather than the **Apply** button, the Modify Host Mappings closes when the task window closes.

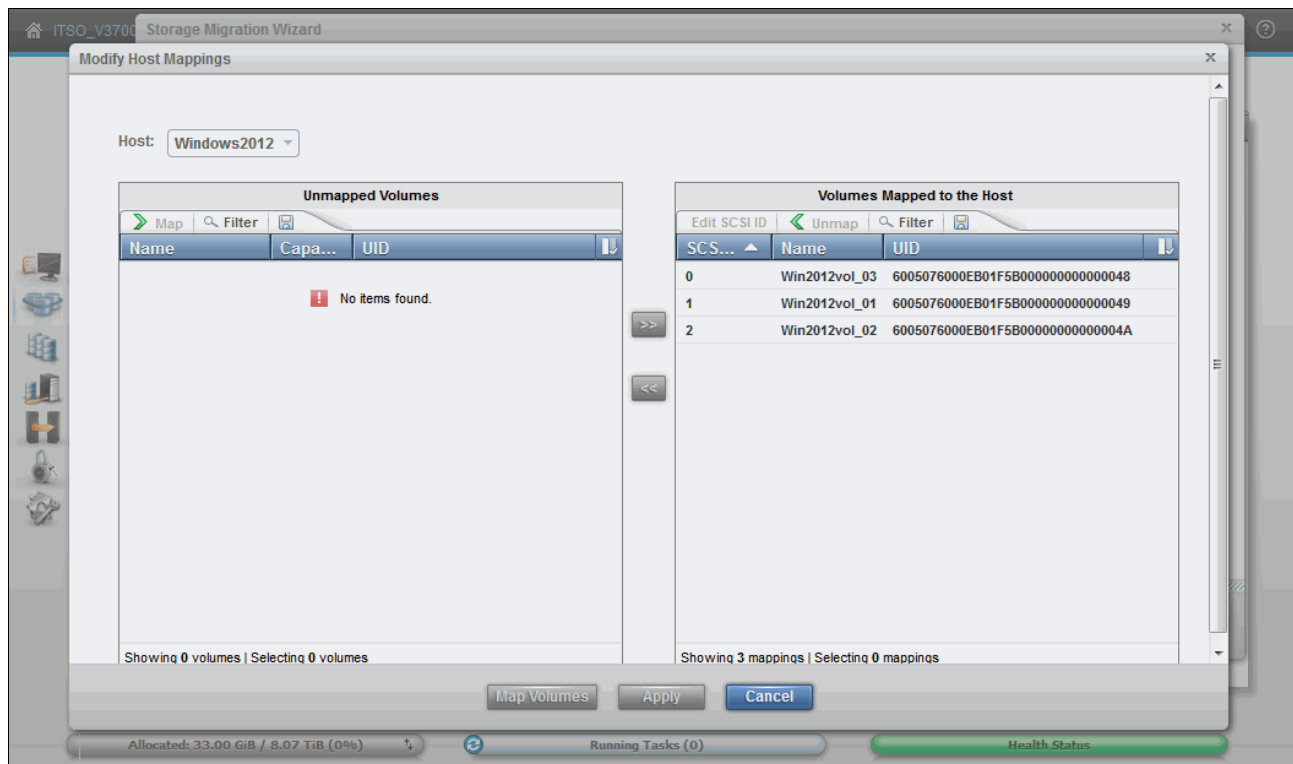


Figure 6-58 Canceling out of Modify Host Mappings

Click **Cancel** to return to the data migration wizard and click **Next** to continue.

Select Storage Pool

To continue with the Storage Migration wizard, complete the following steps:

1. You must now select a Storage Pool to migrate the imported volumes to, as shown in Figure 6-59. The destination Storage Pool must be an internal Storage Pool on the Storwize V3700.

Ensure that there is enough space in the selected Storage Pool to accommodate the migrated volume. The migration task runs in the background and results in a copy of the data being placed on the MDisks in the selected Storage Pool.

The process uses the volume mirroring function that is included with the Storwize V3700.

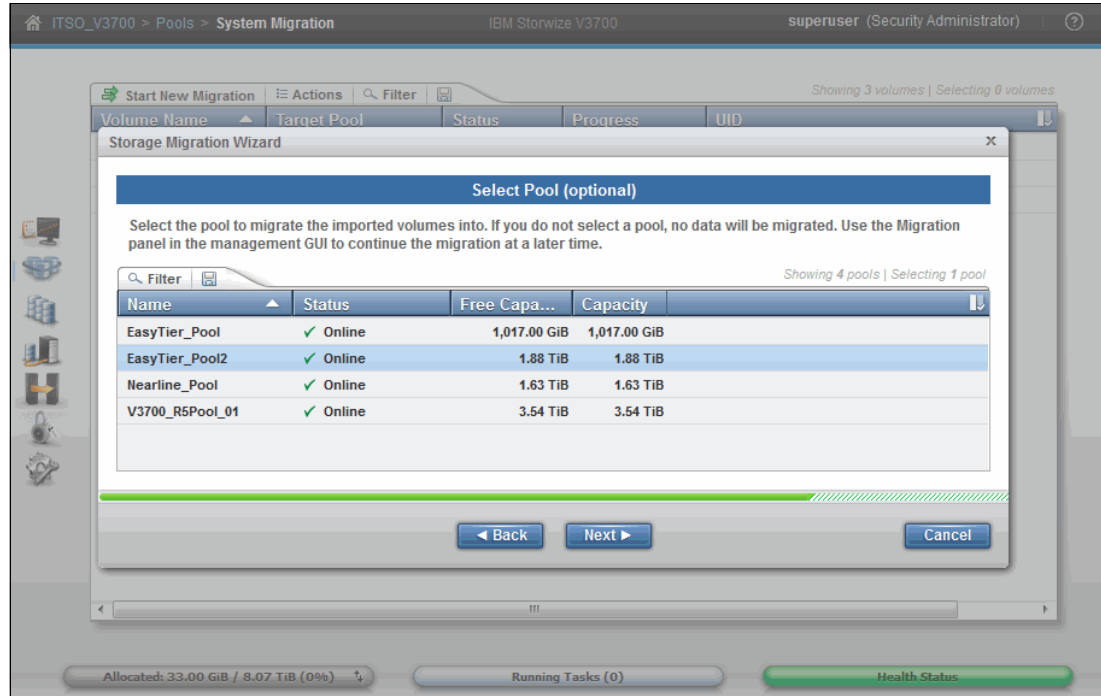


Figure 6-59 Storage Pool selection

2. In this example, you have chosen the EasyTier_Pool2 Storage Pool. Click **Next** to continue.

3. The task to create volumes in the internal Storage Pool completes, as shown in Figure 6-60. Click **Close** to continue.

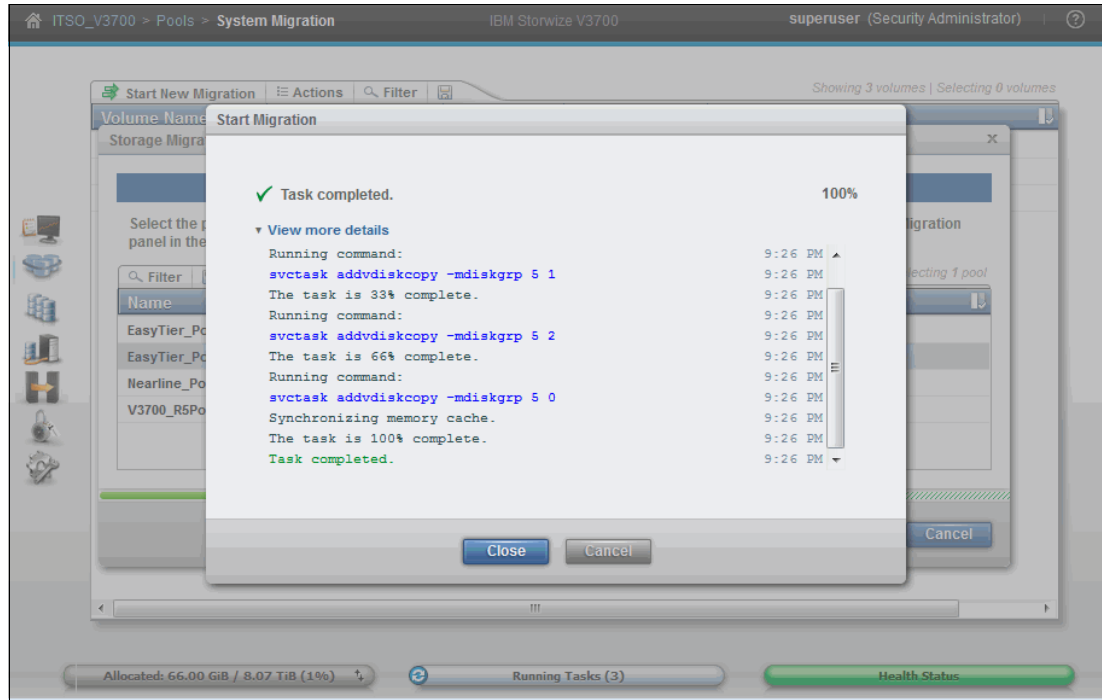


Figure 6-60 Adding volume copies to target pool

Finish the Storage Migration wizard

We are now at the end of the Storage Migration wizard:

1. Click **Finish** to end, as shown in Figure 6-61.

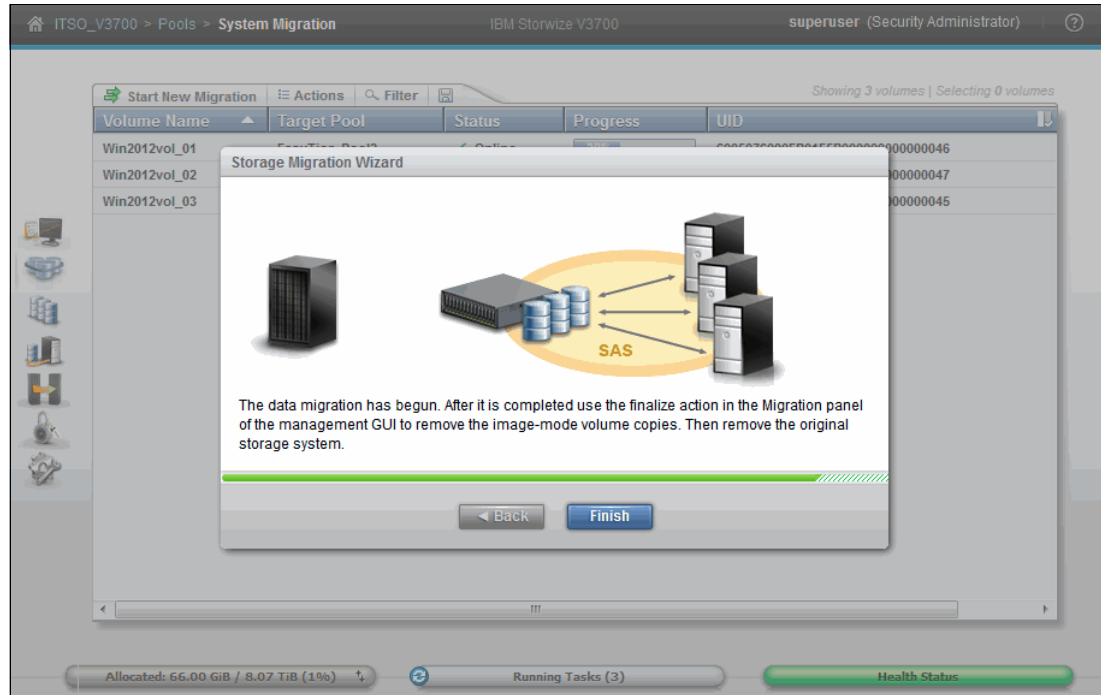


Figure 6-61 Finishing the data migration wizard

2. The end of the Storage Migration wizard is not the end of the data migration task. It is still in progress. A percentage indicator is displayed in the Storage Migration pane, as shown in Figure 6-62.




Figure 6-62 Storage Migration progress

Finalize migrated volumes

To finalize the migration process, complete the following steps:

1. When the migration completes with all of the progress indicators at 100%, select all of the volume migrations you want to finalize by holding the Ctrl key and clicking to select each one.
2. Select the **Actions** box and the **Finalize** option, as shown in Figure 6-63. Alternatively, right-click the selected volumes and click **Finalize**.

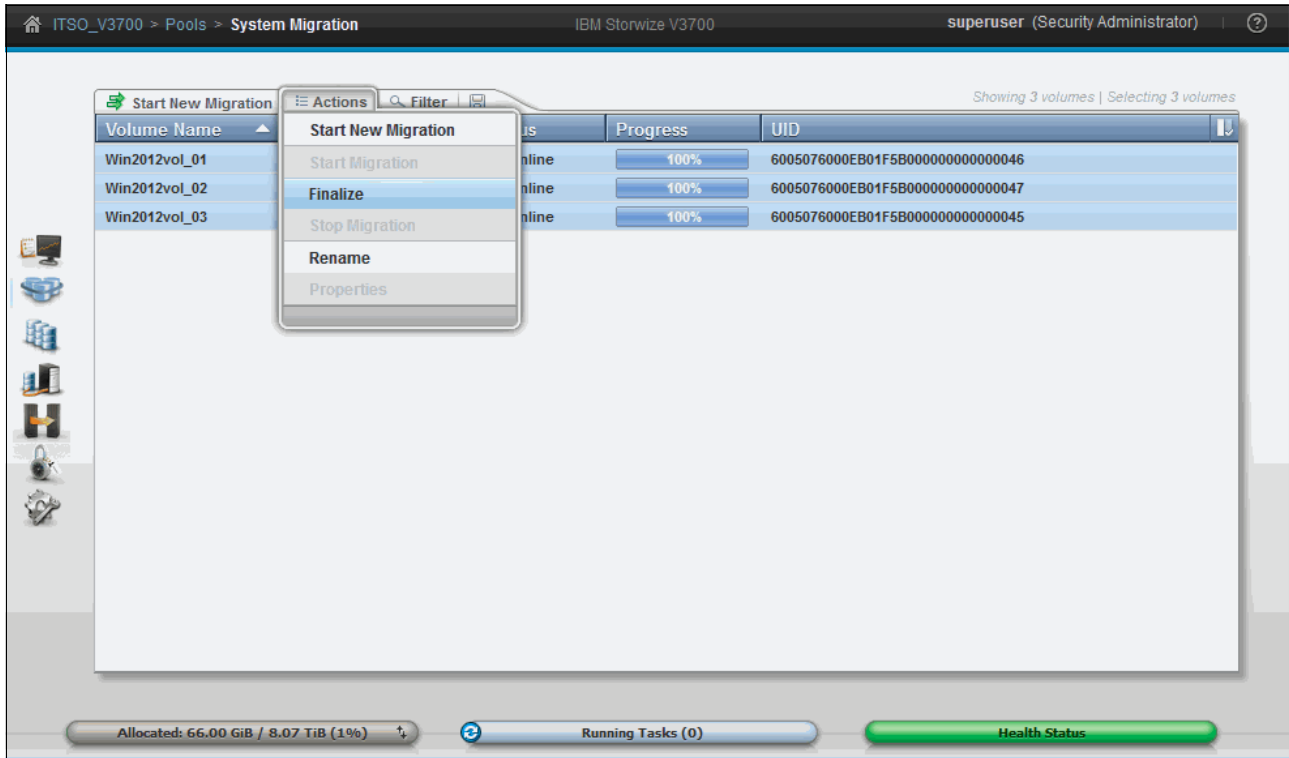


Figure 6-63 Finalize storage migration

3. You are asked to confirm the number of volumes that you want to finalize, as shown in Figure 6-64.

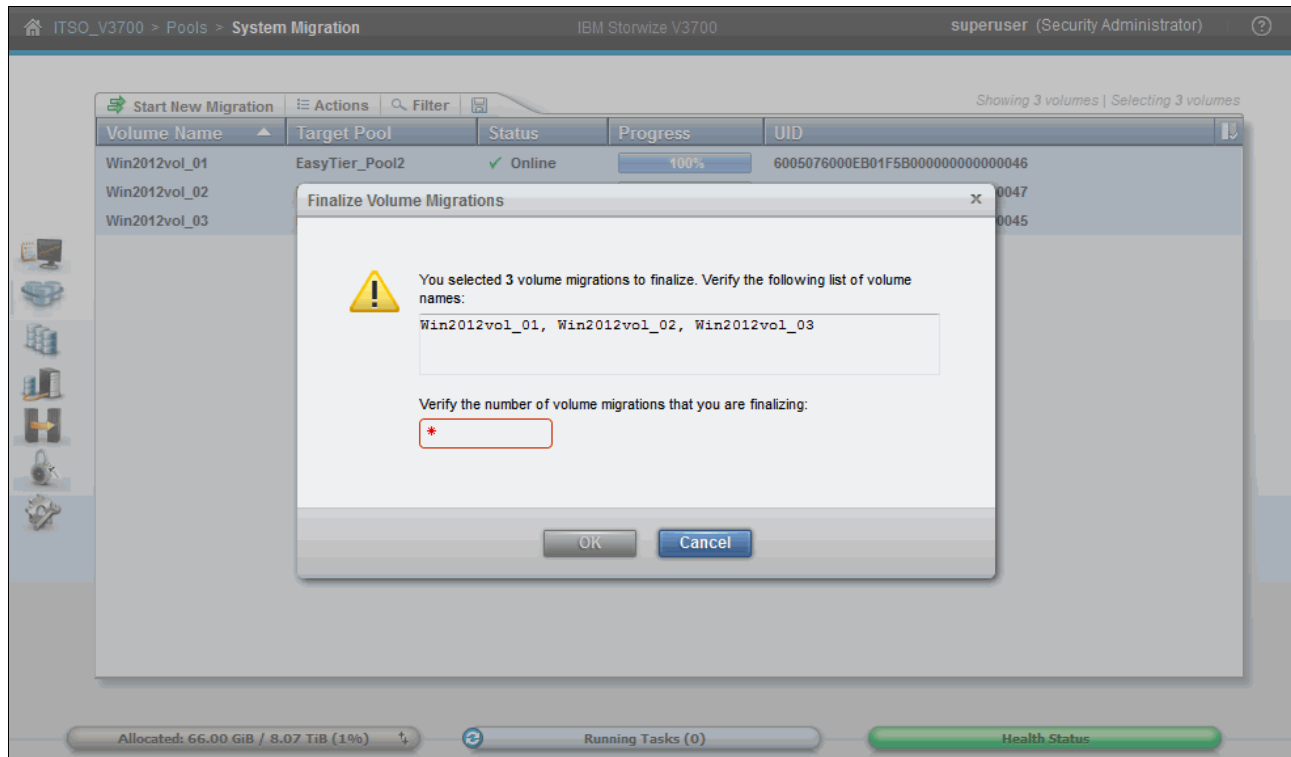


Figure 6-64 Confirm volumes to finalize

4. Verify the volume names, enter the number of migrations to finalize, and click **OK**. The image mode volumes are deleted and the associated image mode MDisks from the migration Storage Pool are removed. Figure 6-65 shows the tasks completing.

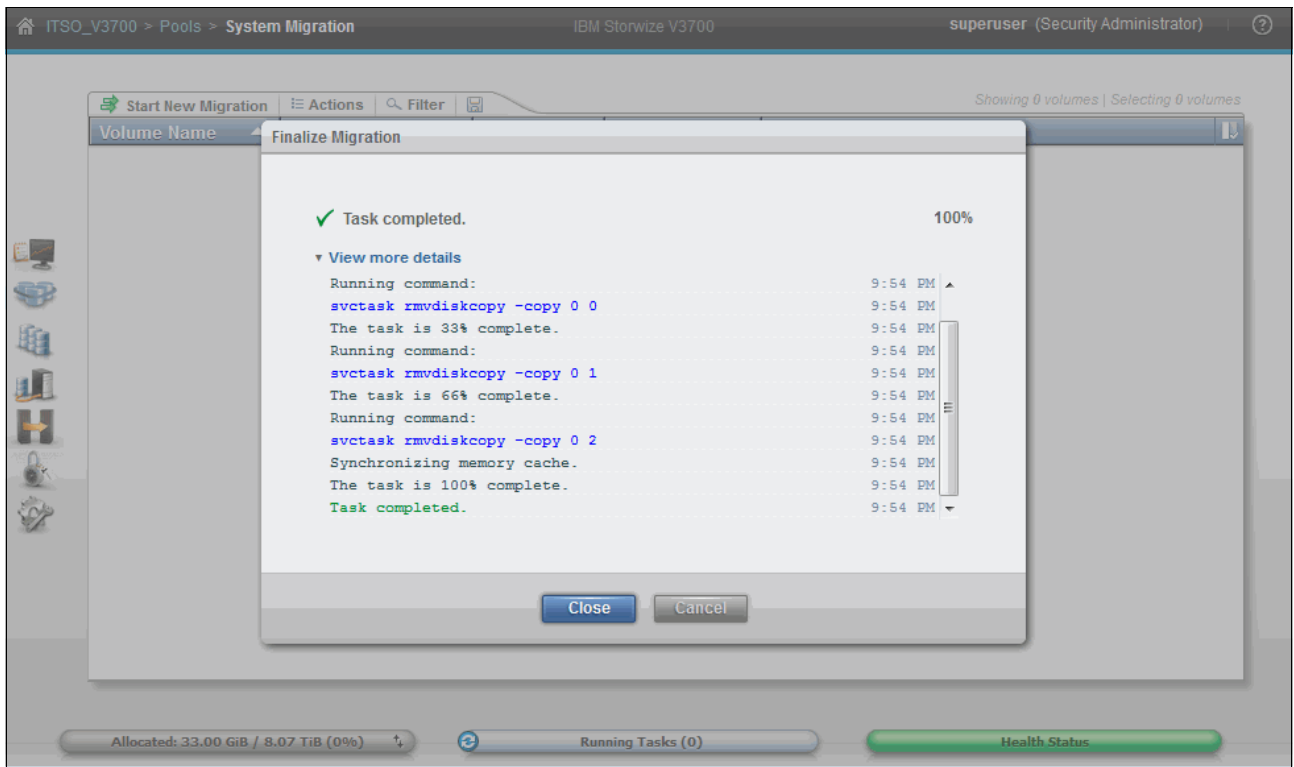


Figure 6-65 Removal of image mode volumes

5. The status of those image mode MDisks is now Unmanaged, as shown in the MDisks by Pools view shown in Figure 6-66.
6. When the finalization completes, the data migration to the IBM Storwize V3700 is done. Remove zoning and connectivity, and then retire the external storage system. In this example, because there is SAS connection, there is no zoning to remove.

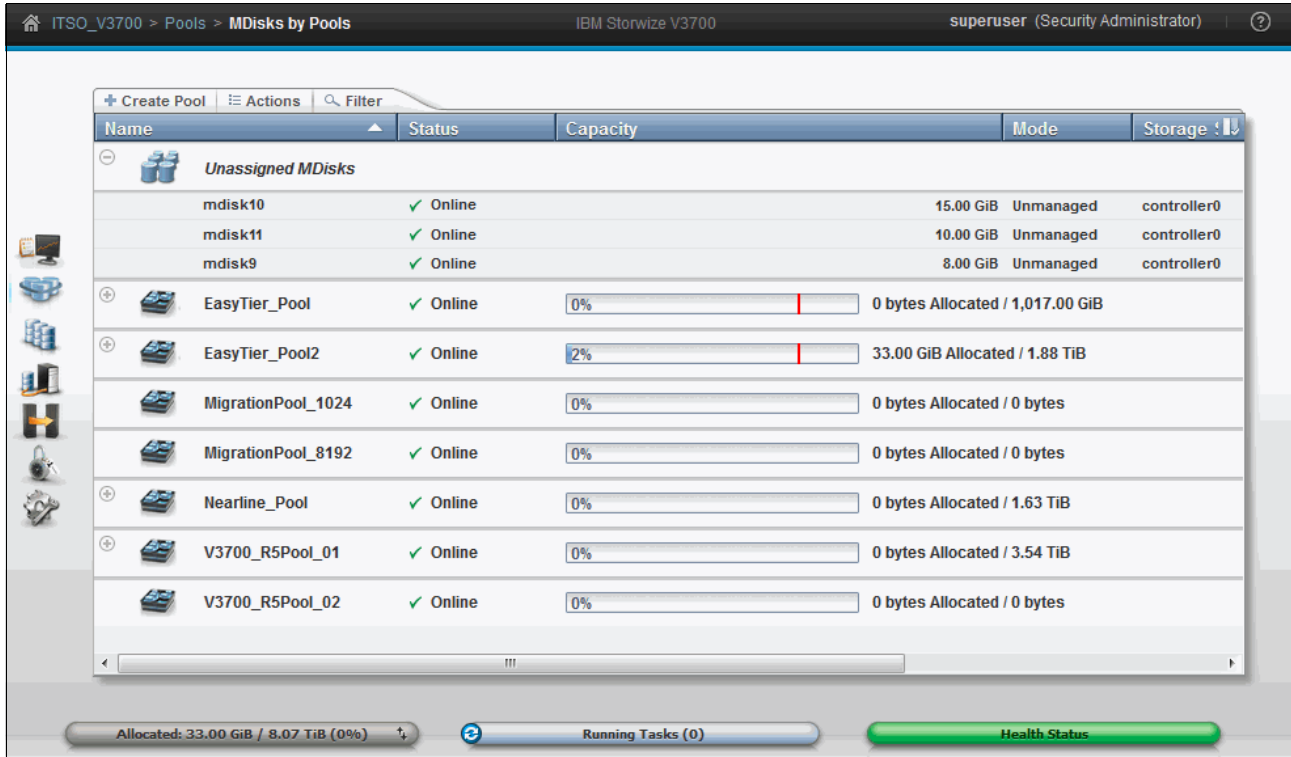


Figure 6-66 Unmanaged external MDisks

Host rescan

After the host is mapped to the Storwize V3700, it is possible to perform a drive rescan and rediscover the volumes on the host side. The data migrations do not have to complete for the host to access the volumes.

To rescan the drive to locate the new Storwize V3700 volumes, complete the following steps:

1. In this example, from the Windows 2012 host, go to the Server Manager Dashboard and select **Tools** → **Computer Management** → **Storage** → **Disk Management**, as shown in Figure 6-24 on page 269. The rescan should find the previous volumes, as shown in Figure 6-67.

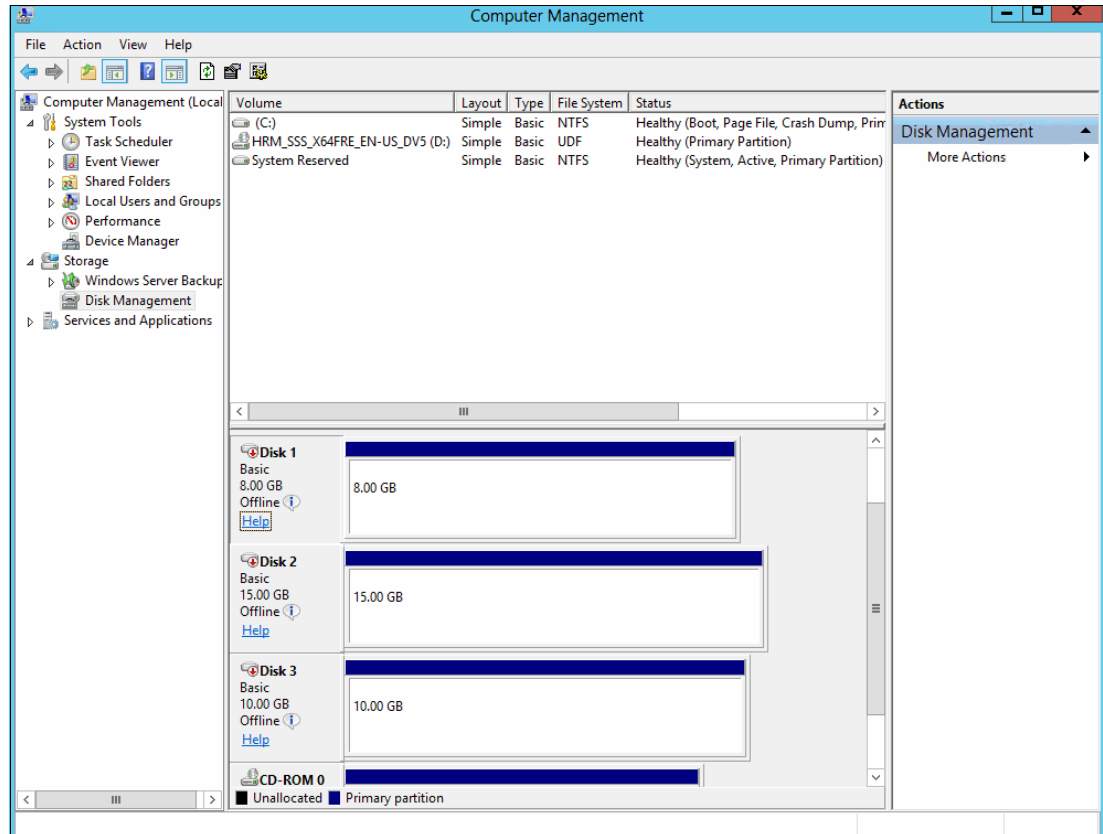


Figure 6-67 Rescanning for new Storwize V3700 volumes: Windows 2012

2. As can be seen in this example, three volumes were found of the same size as before the migration took place. If we put these volumes online, they are identical to those that we removed in Figure 6-24 on page 269. This is shown in Figure 6-68.

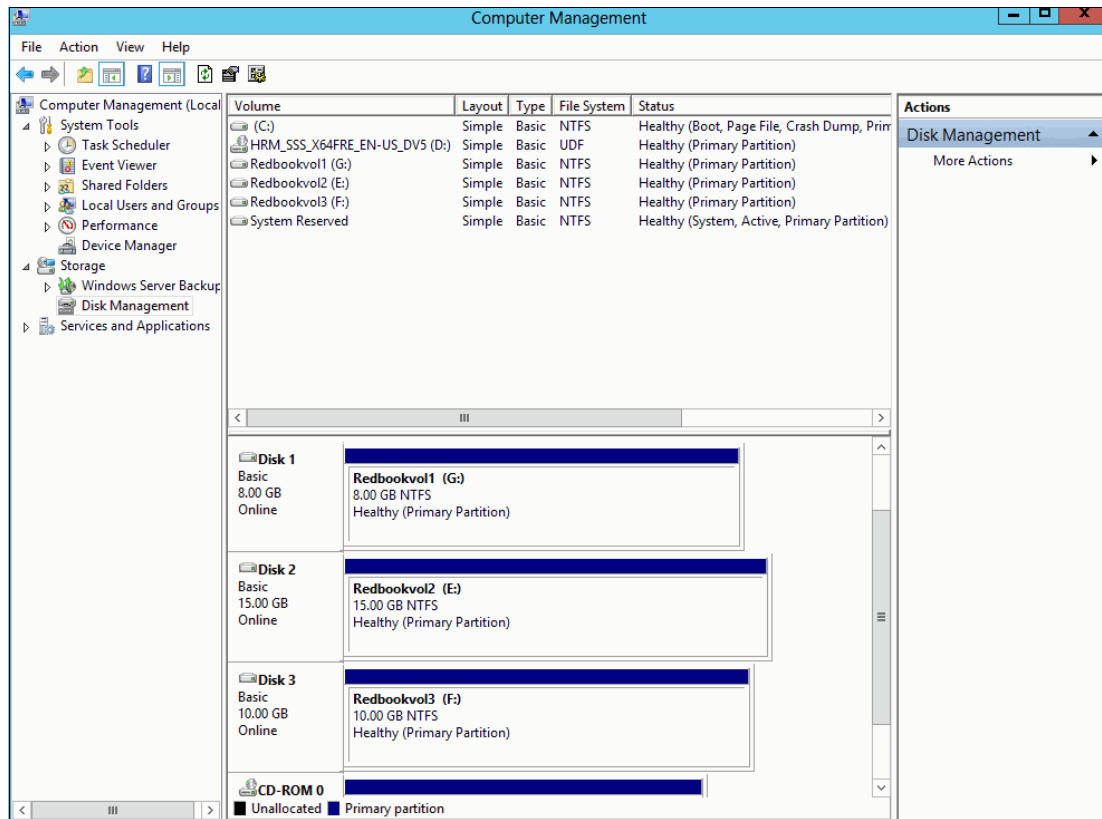


Figure 6-68 Bringing volumes online: Windows 2012

3. The only difference now is that these volumes are presented from the Storwize V3700 using the IBM SDDDSM MPIO driver. You can see this by right-clicking any disk label (Disk 1, Disk 2, or Disk 3) and clicking the properties, as shown in Figure 6-69.
4. As you can see, these disks are now reported as IBM 2145 device type, indicating that they are now presented by a Storwize device, and not the previous external IBM DS3200 device.




Figure 6-69 Device properties



Storage Pools

This chapter describes how IBM Storwize V3700 manages physical storage resources. All storage resources that are under IBM Storwize V3700 control are managed using *Storage Pools*. Storage Pools make it easy to dynamically allocate resources, maximize productivity, and reduce costs. Advanced internal storage, managed disks (MDisks), and Storage Pool management are covered in this chapter.

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

They can also be configured after the initial setup through the management graphical user interface (GUI), which provides a set of presets to help you configure different Redundant Array of Independent Disks (RAID) types.

The suggested configuration presets configure all drives into RAID types based on drive class, and protect them with the appropriate number of spare drives. Alternatively, you can configure the storage to your own requirements. Selections include the drive class, the number of drives to configure, whether to configure spare drives, and optimizing for performance or capacity.

This chapter includes the following topics:

- ▶ Working with internal drives
- ▶ Configuring internal storage
- ▶ Working with Storage Pools
- ▶ Working with child pools

For information about using the IBM Storwize V3700 command-line interface (CLI), see Appendix A, “Command-line interface setup and SAN Boot” on page 649.

7.1 Working with internal drives

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

The IBM Storwize V3700 storage system provides an Internal Storage window for managing all internal drives. The Internal Storage window can be accessed by completing the following steps:

1. Open the System window.
2. Click the **Pools** function icon.
3. Click **Internal Storage**, as shown in Figure 7-1.

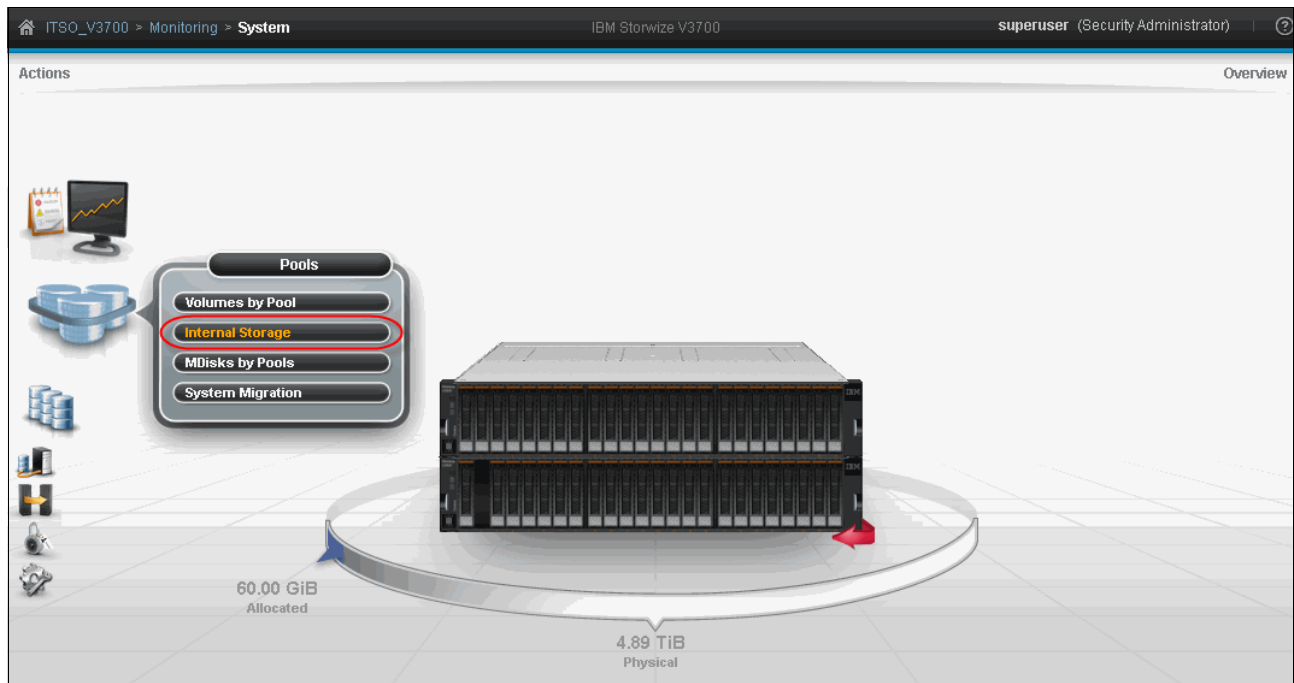


Figure 7-1 Internal Storage Details accessed from the Pools icon

7.1.1 Internal Storage window

The Internal Storage window (as shown in Figure 7-2) provides an overview of the internal drives that are installed in the IBM Storwize V3700 storage system:

1. Selecting **All Internal** under the Drive Class Filter shows all of the drives that are installed in the managed system, including attached expansion enclosures.
2. Alternatively, you can filter the drives by their type or class. For example, you can choose to show only Enterprise drive class (serial-attached Small Computer System Interface (SCSI), called *SAS*), Nearline SAS (NL-SAS), or flash drives.

Drive ID	Capacity	Use	Status	MDisk Name	Enclosure ID	Slot ID
0	278.90 GiB	Member	✓ Online	mdisk4	1	24
1	278.90 GiB	Member	✓ Online	mdisk0	1	13
2	278.90 GiB	Member	✓ Online	mdisk4	1	20
3	278.90 GiB	Member	✓ Online	mdisk4	1	21
4	278.90 GiB	Member	✓ Online	mdisk1	1	9
5	278.90 GiB	Member	✓ Online	mdisk0	1	18
6	278.90 GiB	Member	✓ Online	mdisk4	1	23
7	278.90 GiB	Member	✓ Online	mdisk4	1	22
8	278.90 GiB	Member	✓ Online	mdisk0	1	17
9	278.90 GiB	Spare	✓ Online		1	2
10	278.90 GiB	Member	✓ Online	mdisk1	1	8
11	278.90 GiB	Member	✓ Online	mdisk0	1	16
12	278.90 GiB	Member	✓ Online	mdisk0	1	12
13	278.90 GiB	Member	✓ Online	mdisk0	1	15
14	278.90 GiB	Member	✓ Online	mdisk1	1	6

Figure 7-2 Internal storage window

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

- Logical Drive ID
 - Drive capacity
 - Current type of use (unused, candidate, member, spare, or failed)
 - Status (online, offline, or degraded)
 - MDisk name that the drive is a member of
 - Enclosure ID that it is installed in
 - Slot ID of the enclosure in which drive is installed
3. The default sort order is by enclosure ID. This default can be changed to any other column by left-clicking the column header. To toggle between ascending and descending sort order, left-click the column header again.

- You can see more details, for example the drive's Technology Type, by right-clicking the blue header bar of the table, which opens the selection pane, as shown in Figure 7-3.

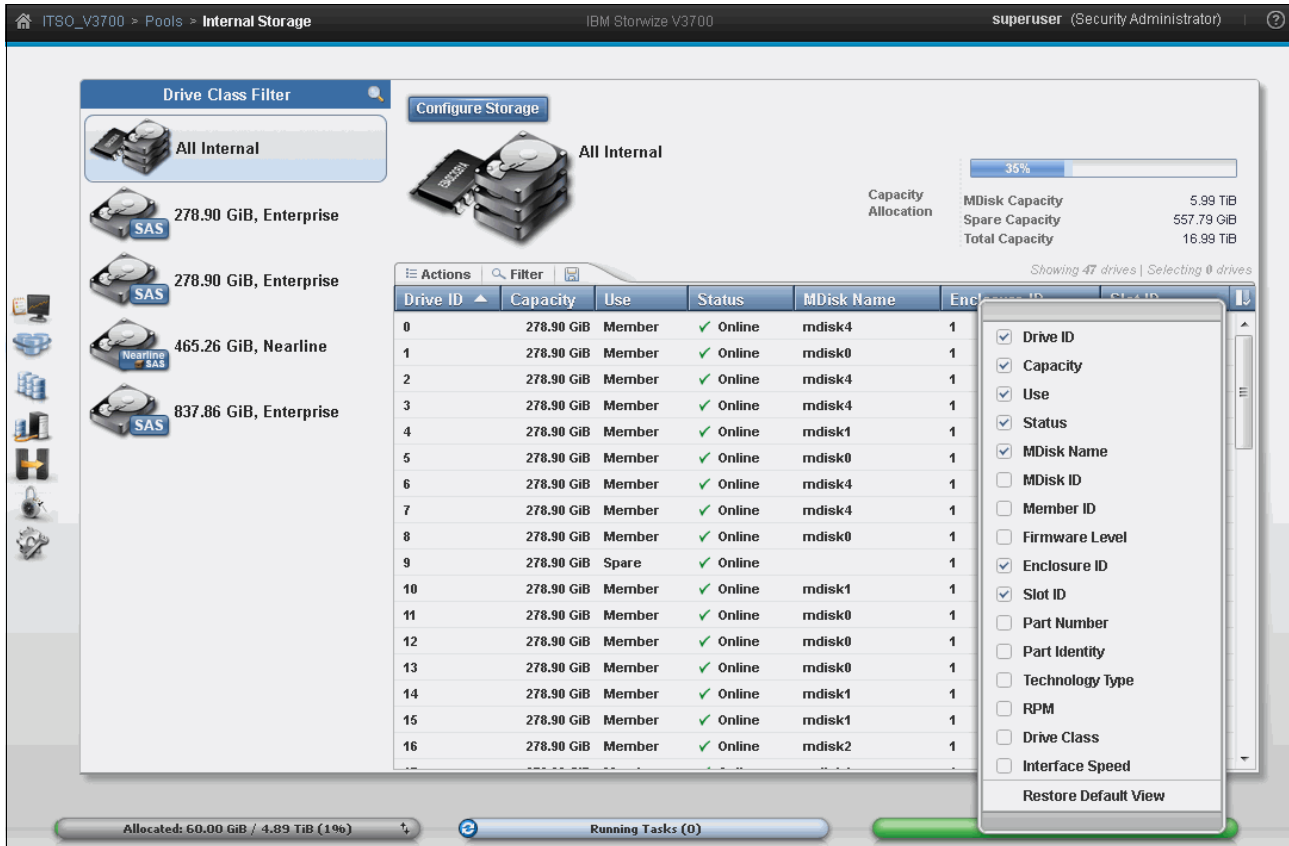


Figure 7-3 Internal storage window details selection

- You can also find the overall internal storage capacity allocation indicator in the upper right corner. The **Total Capacity** shows the overall capacity of the internal storage installed in the IBM Storwize V3700 storage system. The **MDisk Capacity** shows the internal storage capacity that is assigned to the MDisks. The **Spare Capacity** shows the internal storage capacity that is used for hot spare disks.

The percentage bar that is shown in Figure 7-4 indicates how much capacity is allocated.



Figure 7-4 Internal storage allocation indicator

7.1.2 Actions on internal drives

There are several actions that can be performed on the internal drives when you select them and right-click, or click the **Actions** menu, as shown in Figure 7-5.

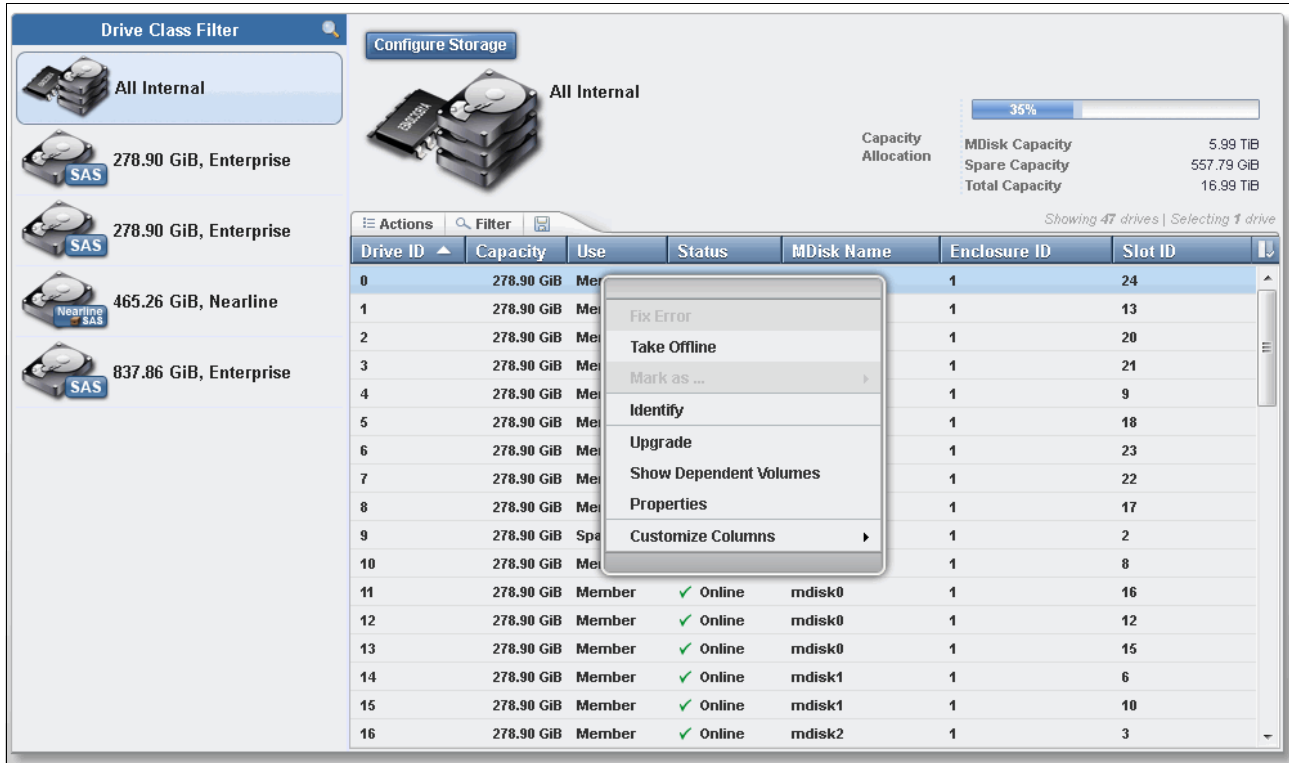


Figure 7-5 Internal drive Actions menu

Depending on the status of the drive selected, the actions available are:

Take Offline

The internal drives can be taken offline if there are problems with them:

1. When you attempt to take a drive offline, a confirmation window opens, as shown in Figure 7-6.

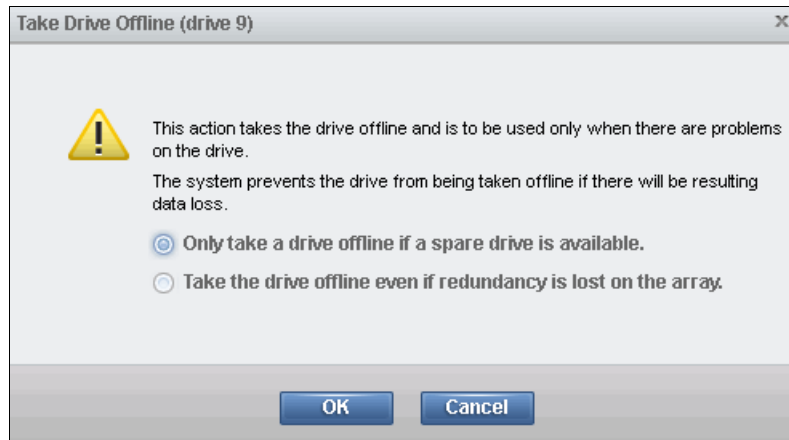


Figure 7-6 Take internal drive offline warning

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

Drive ID	Capacity	Use	Status	MDisk Name	Enclosure ID	Slot ID
8	185.81 GiB	Member	✓ Online	mdisk1	1	9
10	185.81 GiB	Member	✓ Online	mdisk1	1	8
4	185.81 GiB	Failed	✗ Offline		1	17
6	185.81 GiB	Candidate	✓ Online		1	16
11	372.11 GiB	Spare	✓ Online		1	1

Figure 7-7 Internal drive taken offline

3. If no sufficient spare drives are available and one drive must be taken offline, the second option for no redundancy must be selected. This option results in a degraded Storage Pool because of the degraded MDisk, as shown in Figure 7-8.

MDisk	Status	Capacity	Type
V3700_Pool1	Degraded	100.00 GiB Allocated / 4.62 TiB	Storage Pool
mdisk0	Degraded	1.63 TiB	Array
mdisk1	Online	1.63 TiB	Array
mdisk4	Online	1.36 TiB	Array

Figure 7-8 Internal drive that is failed with MDisk degraded

The IBM Storwize V3700 storage system prevents the drive from being taken offline if there might be data loss as a result. A drive cannot be taken offline (as shown in Figure 7-9) if no suitable spare drives are available and, based on the RAID level of the MDisk, drives are already offline.

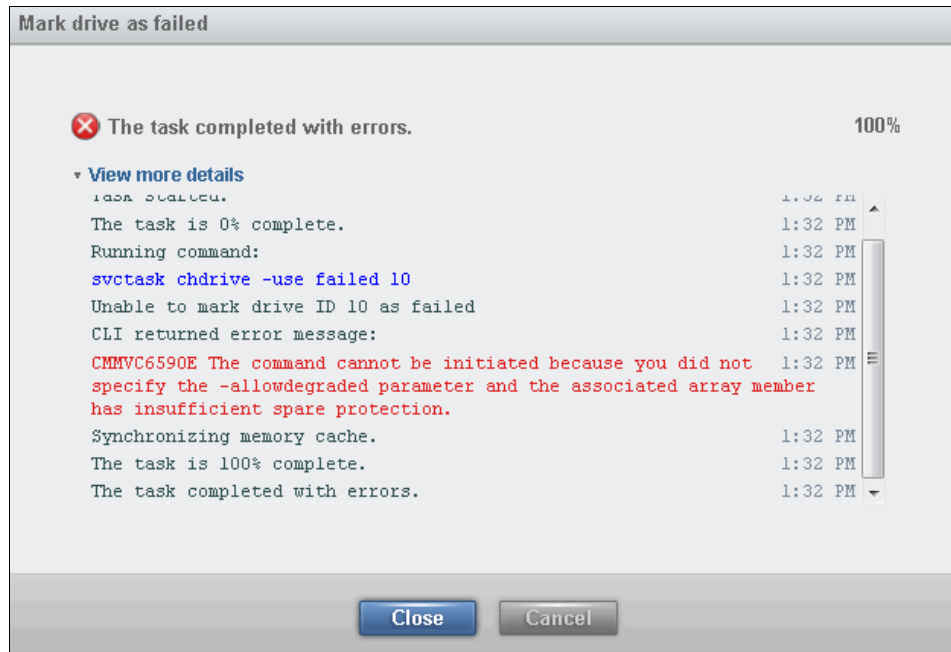


Figure 7-9 Internal drive offline not allowed because of insufficient redundancy

- Example 7-1 shows how to use the `chdrive` CLI command to bypass this check and set the drive to failed anyway.

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

```
IBM_Storwize:ITS0_V3700:superuser>chdrive -use failed -allowdegraded driveID
```

The `-allowdegraded` flag enables you to bypass the check for suitable spare drives.

For information about using the IBM Storwize V3700 CLI, see Appendix A, “Command-line interface setup and SAN Boot” on page 649.

Mark as

The internal drives in the IBM Storwize V3700 storage system can be assigned the following usage roles, as shown in Figure 7-10:

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

The screenshot shows a table of 47 drives. The table has columns for Drive ID, Capacity, Use, Status, MDisk Name, Enclosure ID, and Slot ID. Drive 27 is selected, and a context menu is open over it. The 'Mark as...' option is highlighted, and a sub-menu is displayed with 'Unused', 'Candidate', and 'Spare' options. The 'Unused' option is circled in red.

Drive ID	Capacity	Use	Status	MDisk Name	Enclosure ID	Slot ID
9	278.90 GiB	Spare	✓ Online		1	2
23	278.90 GiB	Candidate	✓ Online		1	1
24	837.86 GiB				2	5
25	278.90 GiB				2	12
26	278.90 GiB					13
27	278.90 GiB					8
28	278.90 GiB					7
29	278.90 GiB					10
30	837.86 GiB					6
31	837.86 GiB				2	1
32	278.90 GiB				2	11
33	837.86 GiB				2	4
34	278.90 GiB	Candidate	✓ Online		2	9
35	465.26 GiB	Candidate	✓ Online		2	22
36	465.26 GiB	Candidate	✓ Online		2	14
37	465.26 GiB	Candidate	✓ Online		2	19
38	278.90 GiB	Candidate	✓ Online		2	3

Figure 7-10 Internal drive Mark as... option

The new role that can be assigned depends on the current drive usage role. These dependencies are shown in Figure 7-11.

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

Figure 7-11 Internal drive usage role table

Identify

Use the Identify action to turn on the LED light so that you can easily identify a drive that must be replaced or that you want to troubleshoot:

1. The pane that is shown in Figure 7-12 opens when the LED is on.




Figure 7-12 Internal drive identification

2. Click **Turn LED Off** when you are finished.
3. Example 7-2 shows how to use the `chenclosureslot` CLI command to turn the drive LED on and off.

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

```
IBM_Storwize:ITS0_V3700:superuser>chenclosureslot -identify yes/no -slot slot enclosureID
```

For information about using the IBM Storwize V3700 CLI, see Appendix A, “Command-line interface setup and SAN Boot” on page 649.

Upgrade

Using this option, you can easily upgrade the drive firmware. The GUI enables you to upgrade individual drives or upgrade all drives that have available updates. For more information about upgrading drive firmware, see Chapter 11, “RAS, monitoring, and troubleshooting” on page 577.

Show Dependent Volumes

To show dependent volumes, complete the following steps:

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

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

Important: A lack of listed dependent volumes does not imply that there are no volumes using the drive.

2. Figure 7-13 shows an example if no dependent volumes are detected for a specific drive.

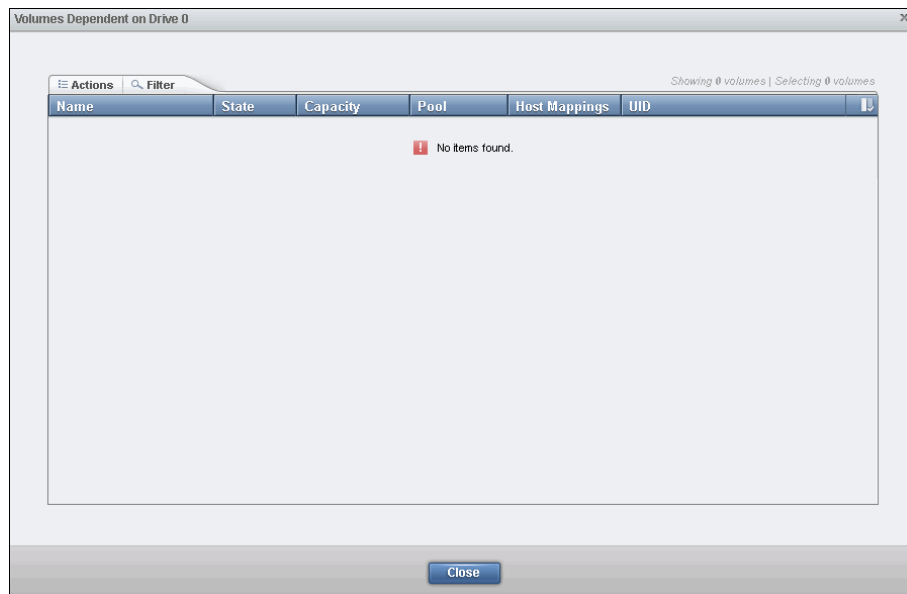


Figure 7-13 Internal drive no dependent volume

3. Figure 7-14 shows the list of dependent volumes for a drive when its MDisk is in a degraded state.




Figure 7-14 Internal drive with dependent volume

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

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

```
IBM_Storwize:ITS0_V3700:superuser>lsdependentvdisks -drive driveID
```

Properties

To view the drive properties, complete the following steps:

1. Click **Properties** (as shown in Figure 7-15) in the Actions menu, or double-click the drive, to see the vital product data (VPD) and the configuration information. The Show Details option was selected to show more information.

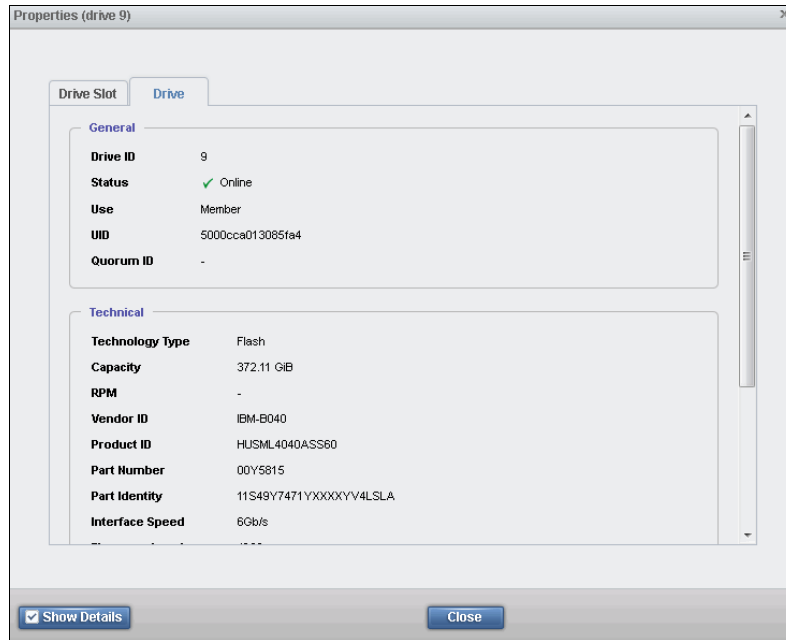


Figure 7-15 Internal drives properties: Part1

2. If the Show Details option is not selected, the technical information section is reduced, as shown in Figure 7-16.




Figure 7-16 Internal drives properties no details

3. A tab for the Drive Slot is available in the Properties pane (as shown in Figure 7-17) to get specific information about the slot of the selected drive.

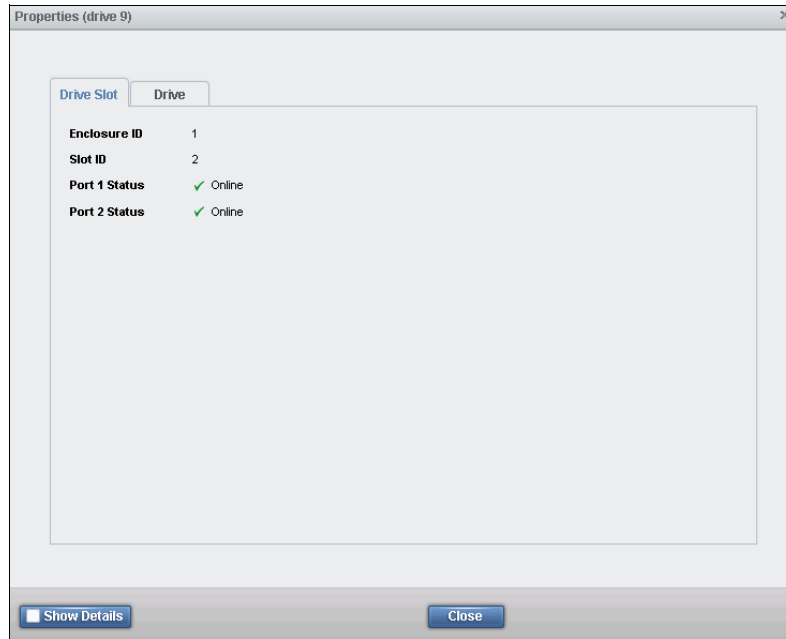


Figure 7-17 Internal drive properties slot

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

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

```
IBM_Storwize:ITS0_V3700:superuser>svcinfo lsdrive 19
ID 19
status online
error_sequence_number
use member
UID 5000cca01604272c
tech_type sas_hdd
capacity 278.9GB
block_size 512
vendor_id IBM-A040
product_id HUC109030CSS60
FRU_part_number 00Y2429
FRU_identity 11S49Y7446YXXXKLG28TYF
RPM 10000
firmware_level J2E9
FPGA_level
mdisk_id 0
mdisk_name mdisk1
member_id 6
enclosure_id 1
slot_id 11
node_id
node_name
quorum_id
port_1_status online
```

```
port_2_status online
interface_speed 6Gb
protection_enabled no
drive_auto_manage idle
```

Customize Columns

To customize columns, complete the following steps:

1. Click **Customize Columns** in the actions menu. This enables you to add or remove several options available in the **Internal Storage** view.
2. To restore the current view, select **Restore Default View**, as shown in Figure 7-18.

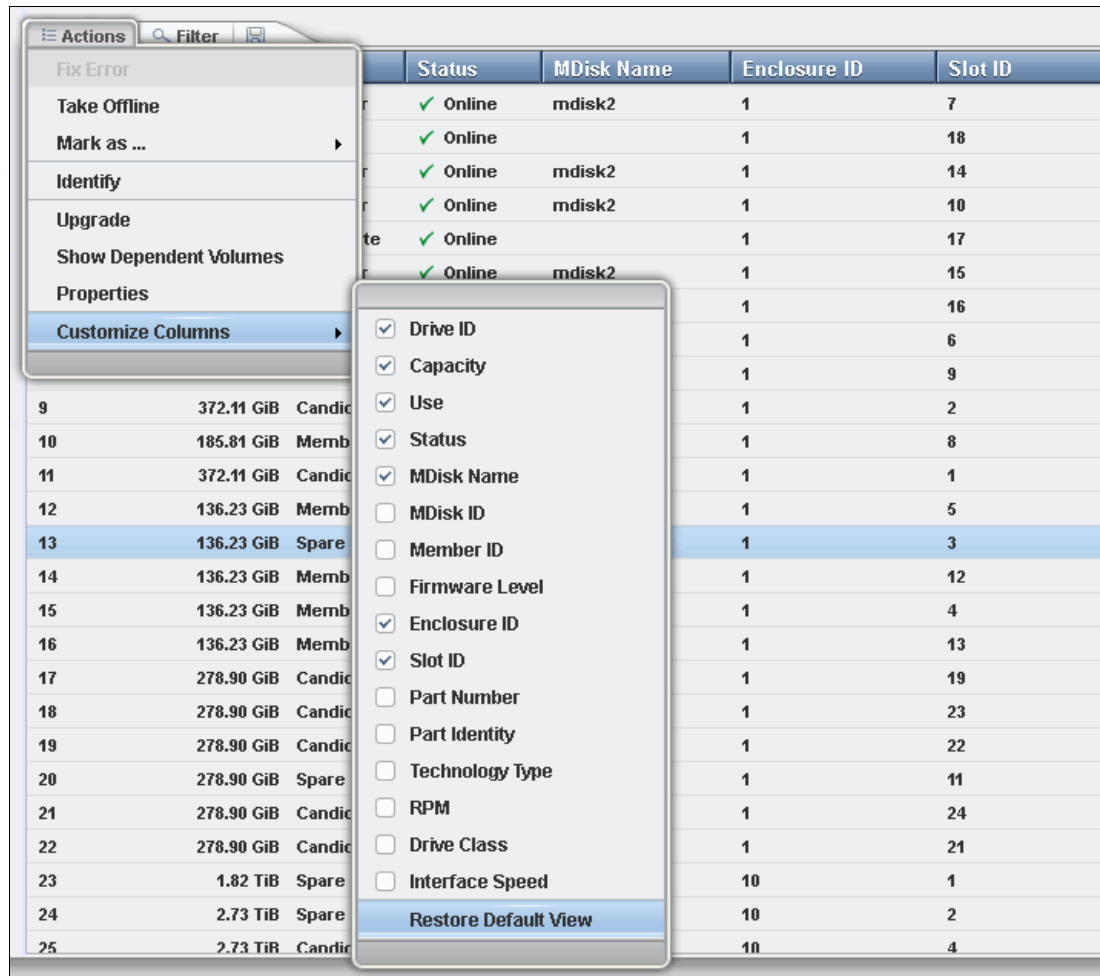


Figure 7-18 Customizing columns

7.2 Configuring internal storage

The internal storage of an IBM Storwize V3700 can be automatically configured into MDisks and pools using the system setup wizard during the initial configuration.

If the internal storage of an IBM Storwize V3700 was not configured during the initial setup, you can manually configure arrays and MDisks:

1. Click **Configure Storage** in the **internal Storage** window as shown in Figure 7-19.

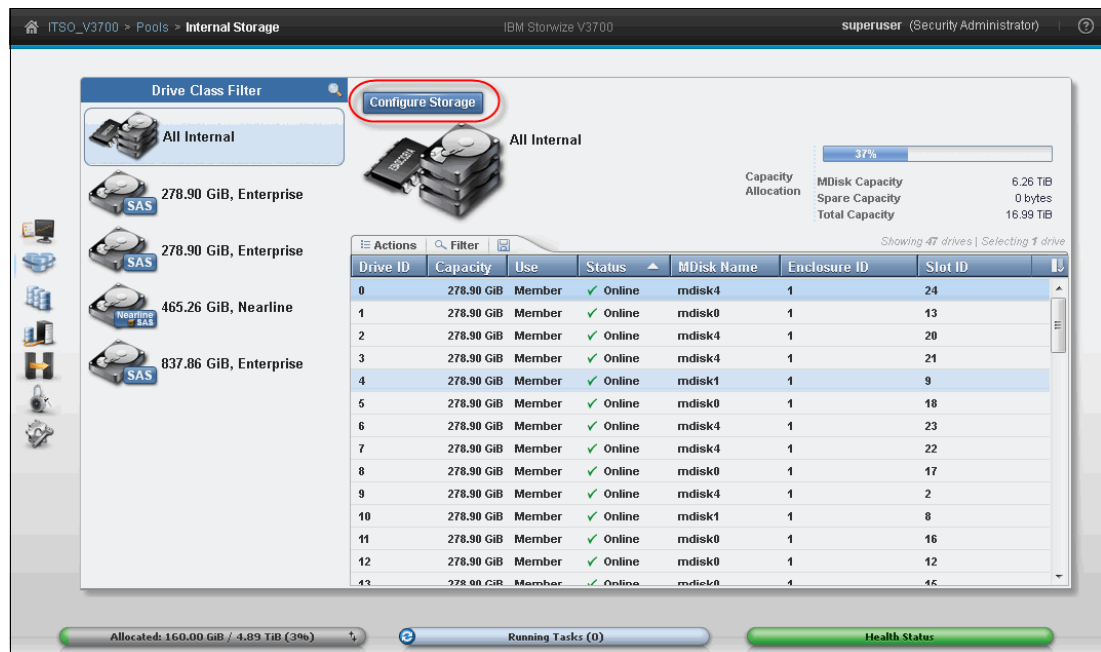


Figure 7-19 Configuring internal storage

2. Choose one of the following options:

- Use the Recommended Configuration

This option will configure all available drives based on the RAID configuration presets. The setup creates MDisks and pools, but does not create volumes.

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

- Select a Different Configuration

This option starts a wizard to customize storage configuration. A storage configuration might be customized for the following reasons:

- The automated initial configuration does not meet customer requirements.
- More storage was attached to the IBM Storwize V3700 and must be integrated into the existing configuration.

7.2.1 RAID configuration presets

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

Table 7-1 describes the presets that are used for flash drives in the IBM Storwize V3700 storage system.

Table 7-1 Flash RAID presets

Preset	Purpose	RAID level	Drives per array goal	Drive count (Min - Max)	Spare drive goal
Flash RAID 5	Protects against a single drive failure. Data and one stripe of parity are striped across all array members.	5	8	3 - 16	1
Flash RAID 6	Protects against two drive failures. Data and two stripes of parity are striped across all array members.	6	12	5 - 16	1
Flash RAID 10	Protects against at least one drive failure. All data is mirrored on two array members.	10	8	4 - 16 (even)	1
Flash RAID 1	Protects against at least one drive failure. All data is mirrored on two array members.	1	2	2	1
Flash RAID 0	Provides no protection against drive failures.	0	8	1 - 8	0
Flash Easy Tier	Mirrors data to protect against drive failure. The mirrored pairs are spread between Storage Pools to be used for the Easy Tier function.	10	2	4 - 16 (even)	1

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

Table 7-2 describes the RAID presets that are used for Enterprise SAS and NL-SAS drives for the IBM Storwize V3700 storage system.

Table 7-2 HDD RAID presets

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

Preset	Purpose	RAID level	Drives per array goal	Drive count (Min - Max)	Spare goal	Chain balance
Basic RAID 10	Protects against at least one drive failure. All data is mirrored on two array members.	10	8	4 - 16 (evens)	1	All drives in the array are from the same chain wherever possible.
Balanced RAID 10	Protects against at least one drive or enclosure failure. All data is mirrored on two array members. The mirrors are balanced across the two enclosure chains.	10	8	4 - 16 (evens)	1	Exactly half of the drives are from each chain.
RAID 0	Provides no protection against drive failures.	0	8	1 - 8	0	All drives in the array are from the same chain wherever possible.

7.2.2 Customizing initial storage configuration

If the initial storage configuration carried out during setup, or the recommended configuration, has been run and does not meet the requirements, the pools created must first be deleted before you can customize the configuration:

1. Select the Pool navigator in the GUI and click **Pools** → **MDisks by Pools**.
2. Right-click the pool and then select **Delete Pool**, as shown in Figure 7-20.

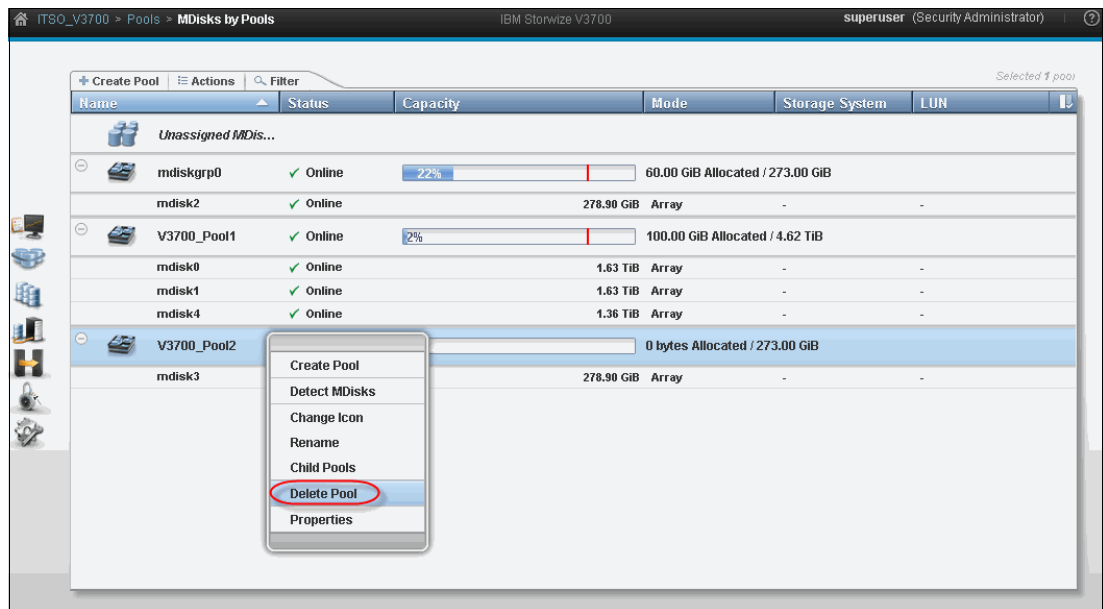


Figure 7-20 Delete selected pool

Note: The **Delete Pool** option only becomes available when the pool is not provisioning volumes.

3. Select **Yes**. The MDisks associated to the pool are removed, and the drives become candidate drives, as shown in Figure 7-21.

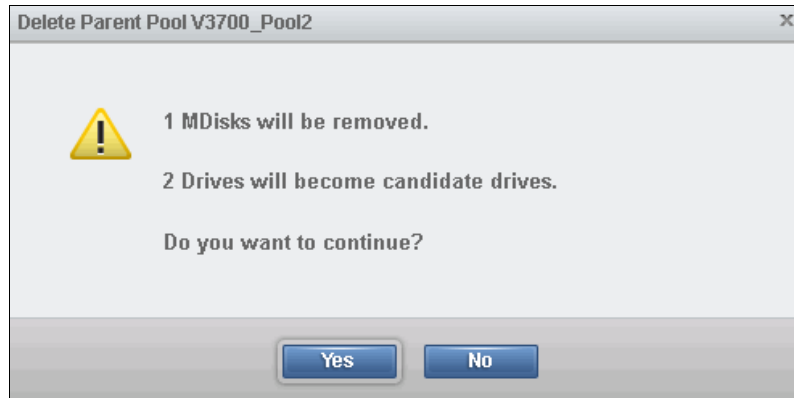


Figure 7-21 Delete pool confirmation

These drives can now be used for a different configuration.

Attention: When deleting a pool using the CLI, it is possible to bypass checks with a **-force** parameter. In this case, any volume provisioned from this pool is deleted, including the data stored on it, *with no option for recovery*.

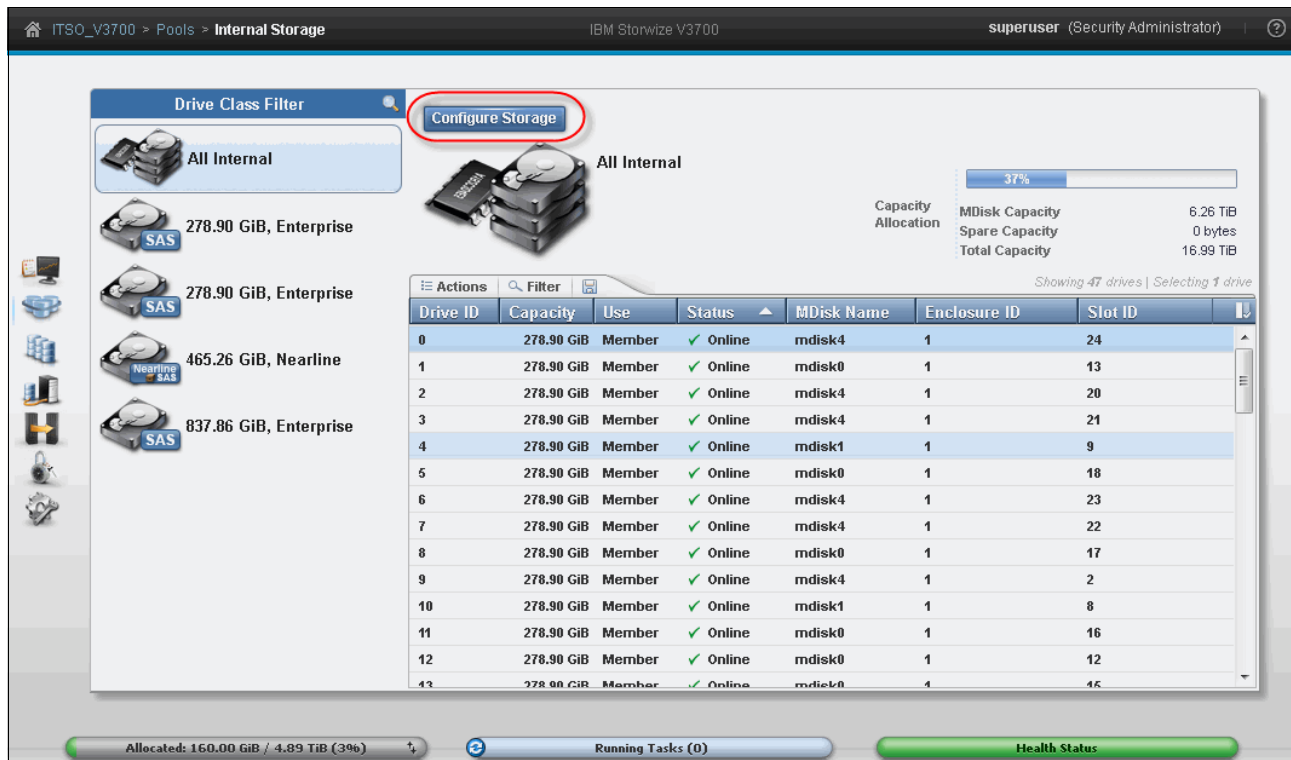
7.2.3 Creating an MDisk and Pool

To configure internal storage for use with hosts, complete the following steps:

1. Click **Pools** → **Internal Storage**. The IBM Storwize V3700 shows all internal drives, their status, and their use. The status shows whether it is Online, Offline, or Degraded.

The Use status shows if a drive is Unused, a Candidate for configuration, a Spare, a Member of a current configuration, or Failed.

2. Click **Configure Storage**, as shown in Figure 7-22.



The screenshot displays the IBM Storwize V3700 management console. The breadcrumb navigation shows 'Internal Storage'. A 'Configure Storage' button is circled in red. On the left, a 'Drive Class Filter' shows 'All Internal' selected. Below it, a list of drives is shown with their capacities and types (SAS, Nearline SAS). On the right, a 'Capacity Allocation' bar shows 37% usage, with a table below it showing MDisk Capacity (6.26 TiB), Spare Capacity (0 bytes), and Total Capacity (16.99 TiB). A table of drive details is shown below, with columns for Drive ID, Capacity, Use, Status, MDisk Name, Enclosure ID, and Slot ID. The table shows 13 drives, all with a capacity of 278.90 GiB and a status of 'Online'. The 'Use' column shows 'Member' for all drives. The 'MDisk Name' column shows 'mdisk4' for drives 0, 1, 2, 3, 4, 6, 7, 8, 9, 10, 11, 12, and 13. The 'Enclosure ID' column shows '1' for all drives. The 'Slot ID' column shows '24', '13', '20', '21', '9', '18', '23', '22', '17', '2', '8', '16', and '12' for drives 0 through 13 respectively.

Drive ID	Capacity	Use	Status	MDisk Name	Enclosure ID	Slot ID
0	278.90 GiB	Member	Online	mdisk4	1	24
1	278.90 GiB	Member	Online	mdisk0	1	13
2	278.90 GiB	Member	Online	mdisk4	1	20
3	278.90 GiB	Member	Online	mdisk4	1	21
4	278.90 GiB	Member	Online	mdisk1	1	9
5	278.90 GiB	Member	Online	mdisk0	1	18
6	278.90 GiB	Member	Online	mdisk4	1	23
7	278.90 GiB	Member	Online	mdisk4	1	22
8	278.90 GiB	Member	Online	mdisk0	1	17
9	278.90 GiB	Member	Online	mdisk4	1	2
10	278.90 GiB	Member	Online	mdisk1	1	8
11	278.90 GiB	Member	Online	mdisk0	1	16
12	278.90 GiB	Member	Online	mdisk0	1	12
13	278.90 GiB	Member	Online	mdisk0	1	15

Figure 7-22 Click Configure Storage

3. A configuration wizard opens and guides you through the process of configuring internal storage. Figure 7-23 shows an example in which 38 drives are available for configuration.

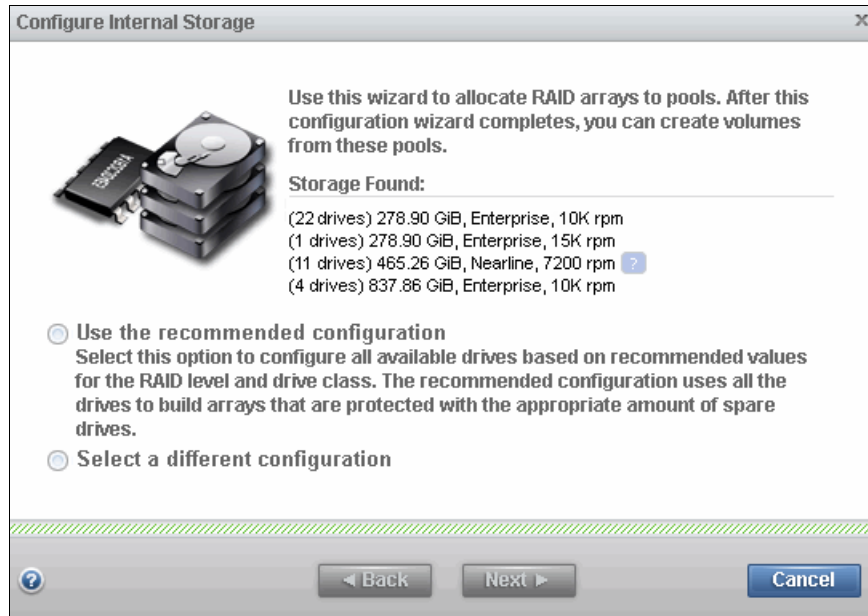


Figure 7-23 Available drives for new MDisk

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

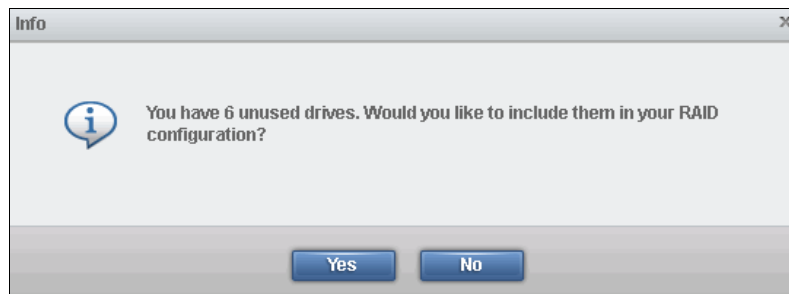


Figure 7-24 Unused drives notification

5. When the decision is made to include the drives into the RAID configuration, their status is set to Candidate, which also makes them available for a new MDisk.
6. The use of the storage configuration wizard simplifies the initial disk drive setup, and offers the following options:
 - Use the recommended configuration. Selecting **Use the recommended configuration** guides you through the wizard that is described in 7.2.4, “Using the recommended configuration” on page 326.
 - Select a different configuration. Selecting **Select a different configuration** uses the wizard that is described in 7.2.5, “Selecting a different configuration” on page 329.

7.2.4 Using the recommended configuration

To use the recommended configuration, complete the following steps:

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

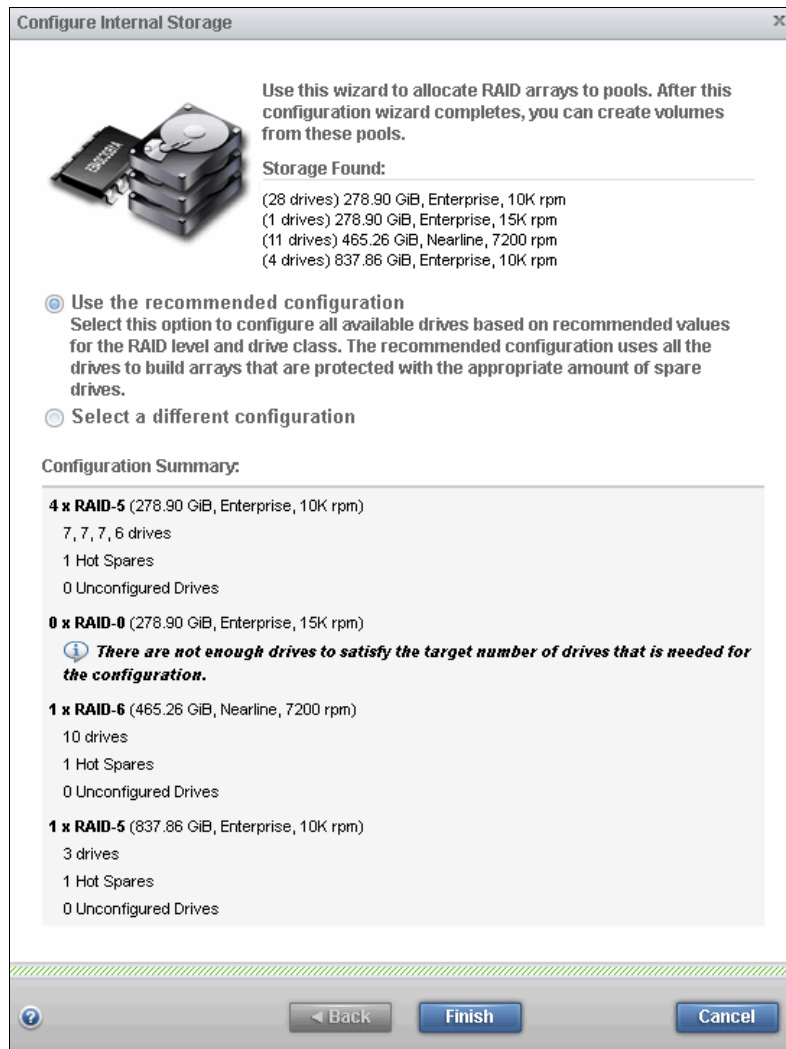


Figure 7-25 The recommended configuration

2. In the Configuration Summary section (shown in Figure 7-26 on page 327), the wizard warns that there are not enough disk drives to satisfy the target number of the drives needed for the configuration, as follows:
 - There are insufficient Enterprise disk classes, 278.9 gibibytes (GiB) 15,000 rpm.

- Therefore, the wizard automatically creates several MDisks, as shown in Figure 7-26.

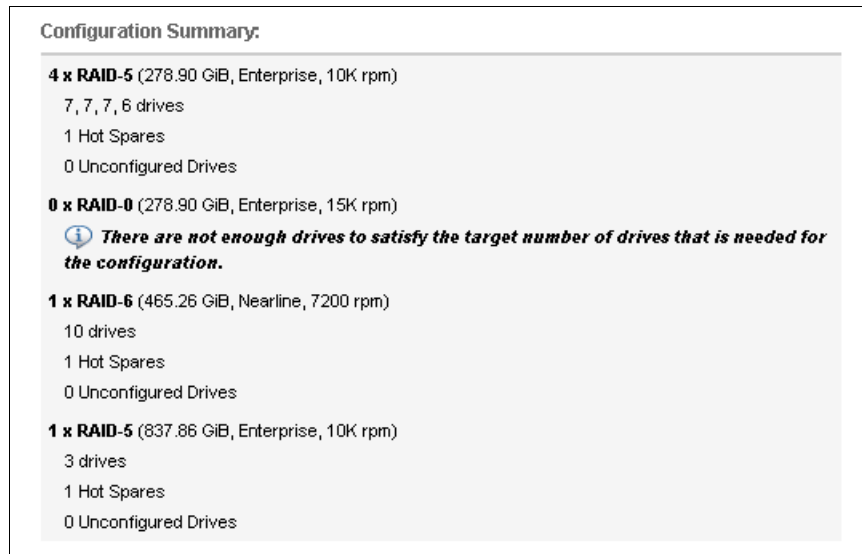


Figure 7-26 The recommended configuration summary

- Spare drives are also automatically created to meet the spare goals according to the chosen preset. One spare drive is created out of every 24 disk drives of the same drive class on a single chain. Spares are not created if sufficient spares are already configured.

Spare drives in the IBM Storwize V3700 are *global spares*, which means that any spare drive that has at least the same capacity as the drive to be replaced can be used in any array. Therefore, a flash array with no flash spare available uses a SAS spare instead.

- If the proposed configuration meets your requirements, click **Finish**, and the system automatically creates the array MDisks.
- Storage Pools are also automatically created to contain the MDisks. MDisks with similar performance characteristics, including RAID level, number of member drives, and drive class, are grouped into the same Storage Pool.

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

- After an array is created, the Array MDisk members are synchronized with each other through a background initialization process. You can monitor the progress of the initialization process by clicking the icon at the left of the Running Tasks status bar and selecting the initialization task to view the status, as shown in Figure 7-27.

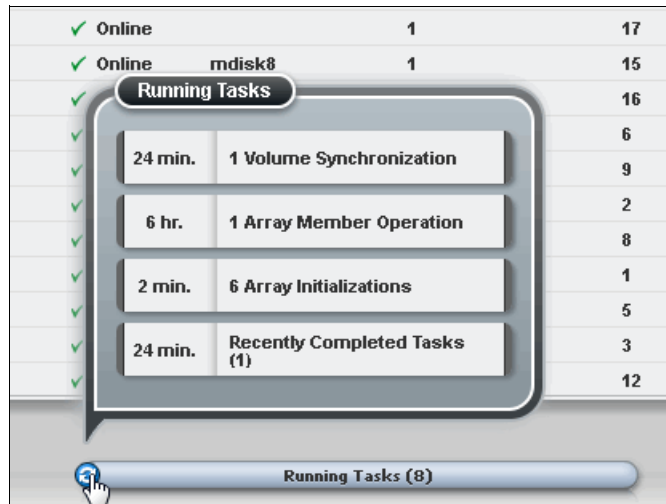


Figure 7-27 Running task pane

- Click the taskbar to open the progress window, as shown in Figure 7-28. The array is available for I/O during this process, however performance is affected until the initialization task has completed.

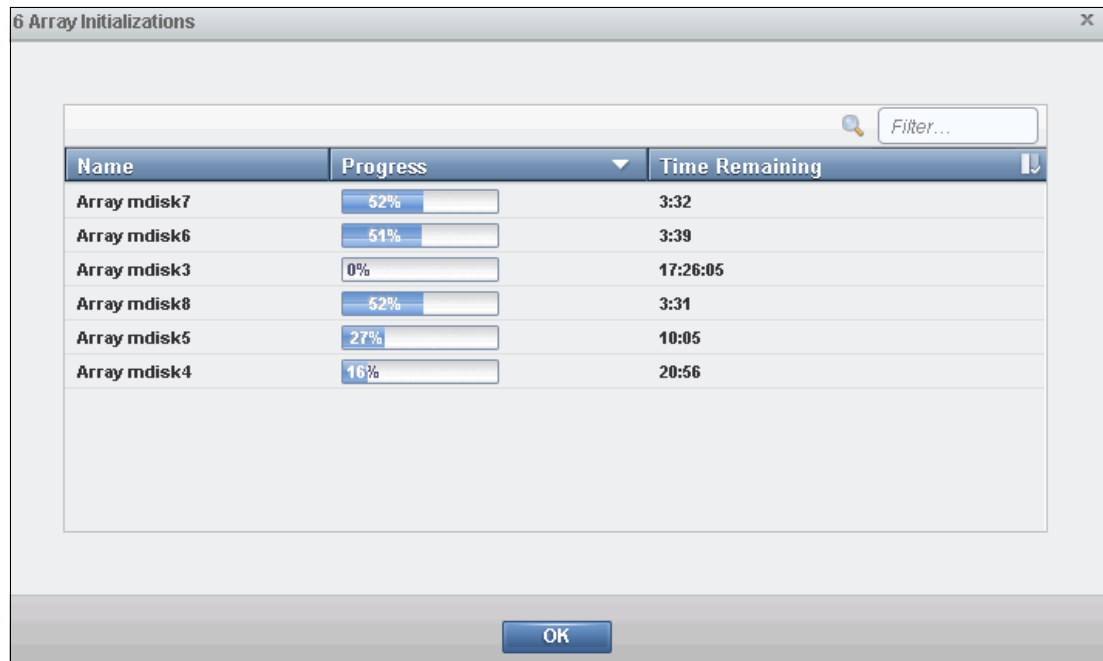


Figure 7-28 Initialization progress view

7.2.5 Selecting a different configuration

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

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

Complete the following steps to select a different configuration:

1. Choose drive class.

The drive class selection list contains each drive class available for configuration, as shown in Figure 7-29.

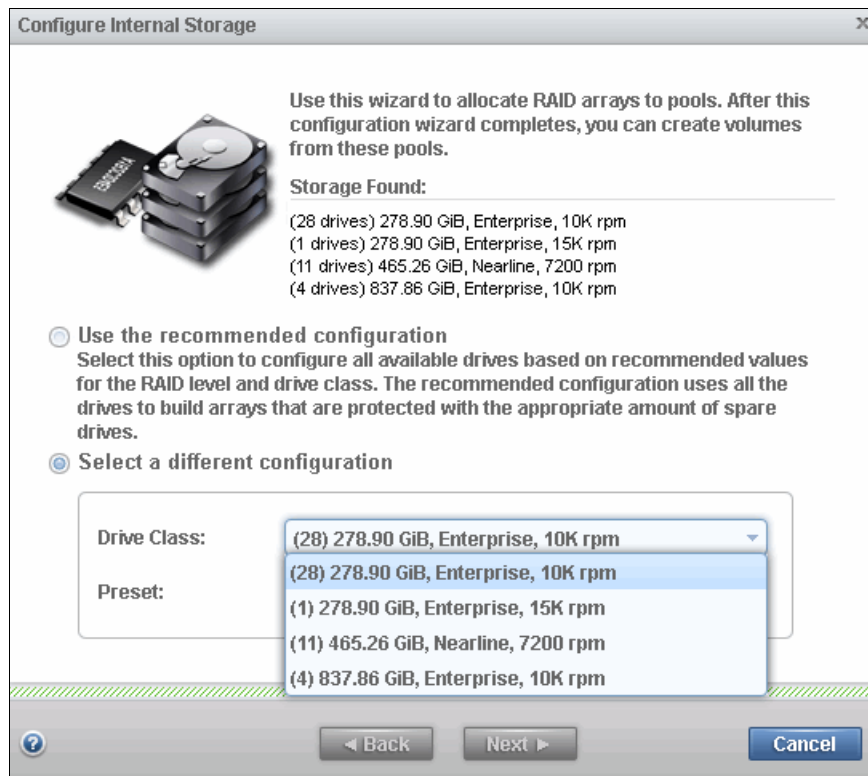


Figure 7-29 Select drive class for new configuration

2. Click **Next** and select the appropriated RAID preset, as shown in Figure 7-30.

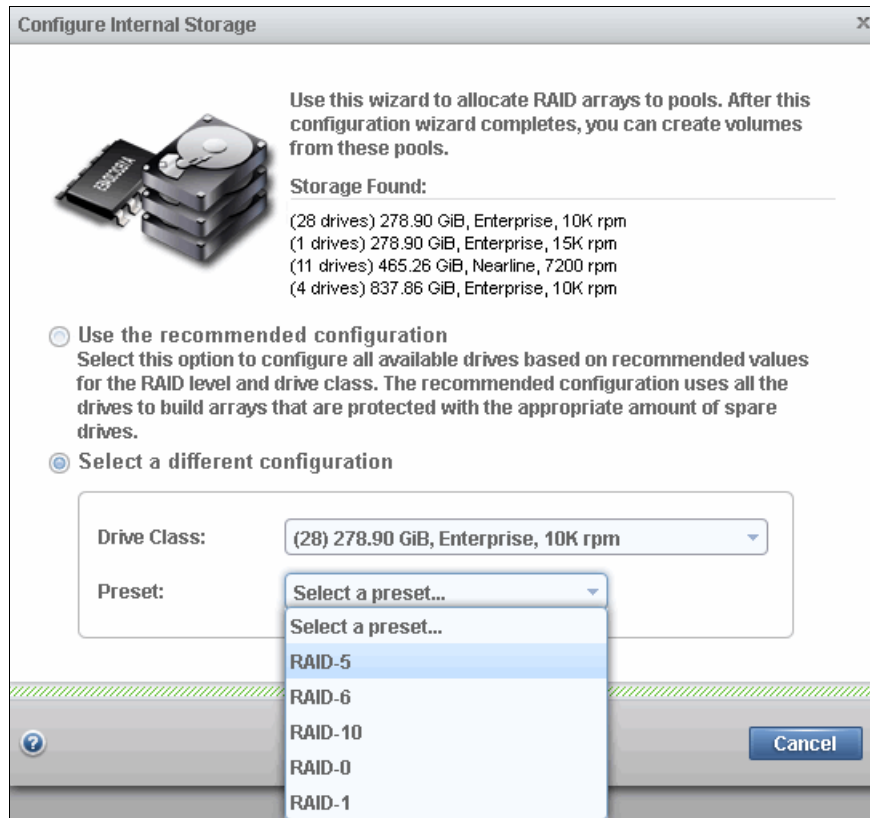


Figure 7-30 Select the RAID preset

3. Define the RAID attributes, including the following selections:

- Configuration of spares
- Optimization for performance
- Optimization for capacity
- Number of drives to provision

Each IBM Storwize V3700 preset has a specific goal for the number of drives per array. For more information, see the IBM Storwize V3700 Information Center on the following website:

<https://ibm.biz/BdEz7J>

Table 7-3 shows the RAID goal widths.

Table 7-3 RAID goal width

RAID level	Enterprise goal width	Flash goal width
0	8	8
5	8	9
6	12	10
10	8	8

The following RAID configurations are available:

– Optimize for Performance

Optimizing for performance creates arrays with the same capacity and performance characteristics. The RAID goal width (as shown in Table 7-3) must be met for this target. In a performance-optimized setup, the IBM Storwize V3700 provisions eight physical disk drives in a single array MDisk, except for the following situations:

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

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

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

Figure 7-31 shows an example in which not all of the provisioned drives can be used in a performance optimized configuration (five drives remain).

The screenshot shows a configuration window titled "Select a different configuration". It contains the following elements:

- Drive Class:** A dropdown menu showing "(31) 278.90 GiB, Enterprise, 10K rpm".
- Preset:** A dropdown menu showing "RAID-5".
- Options:** Three radio buttons: "Automatically configure spares" (checked), "Optimize for Performance ?" (selected), and "Optimize for Capacity".
- Number of drives to provision:** A text input field containing "31", which is circled in red.

Below the configuration options is a "Configuration Summary:" section with a grey background. It lists:

- 3 x RAID-5** (278.90 GiB, Enterprise, 10K rpm)
- 8, 8, 8 drives
- 2 Hot Spares
- 5 Unconfigured Drives

A red circle highlights the "5 Unconfigured Drives" line, which includes a warning icon and the text: "This configuration cannot fully use the number of drives that was requested. Select Optimize for Capacity to allocate all the storage to the arrays."

Figure 7-31 Optimization for performance failed

Figure 7-32 shows that the number of drives is not enough to satisfy the needs of the configuration.

The screenshot shows a configuration window titled "Select a different configuration". It includes a "Drive Class" dropdown set to "(31) 278.90 GiB, Enterprise, 10K rpm" and a "Preset" dropdown set to "RAID-5". There are three radio buttons: "Automatically configure spares" (checked), "Optimize for Performance" (selected), and "Optimize for Capacity". A text input field for "Number of drives to provision" contains the value "5". Below this is a "Configuration Summary" section showing "0 x RAID-5 (278.90 GiB, Enterprise, 10K rpm)" and a red-bordered error message: "The number of drives entered is not enough to satisfy the target number of drives that is needed for the configuration."

Figure 7-32 Not enough drives for performance optimization

Figure 7-33 shows that there are a suitable number of drives to configure performance optimized arrays.

The screenshot shows the same configuration window as Figure 7-32. The "Number of drives to provision" field now contains the value "26". The "Configuration Summary" section shows "3 x RAID-5 (278.90 GiB, Enterprise, 10K rpm)" with sub-items: "8, 8, 8 drives", "2 Hot Spares", and "5 Unconfigured Drives".

Figure 7-33 Arrays match performance goals

Three RAID 5 arrays are possible, and 26 out of 31 provisioned drives are used. Five drives remain.

- Optimize for Capacity

Optimizing for capacity creates arrays that allocate all of the drives that are specified in the Number of drives to provision field. This option results in arrays of different capacities and performance. The number of drives in each MDisk does not vary by more than one drive, as shown in Figure 7-34.

4. After all of these choices are made, click **Next**.

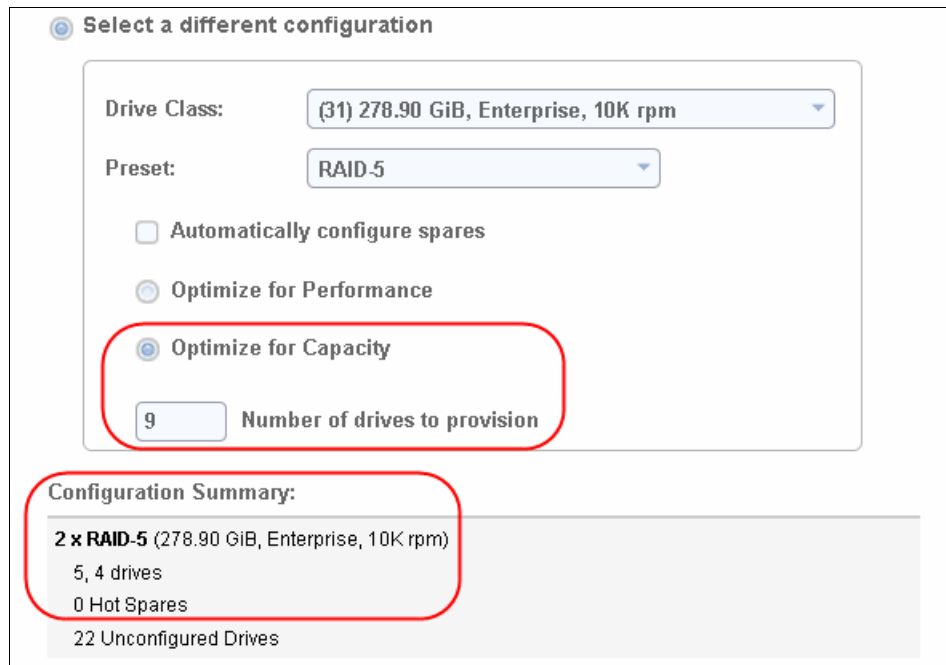


Figure 7-34 Capacity optimized configuration

5. Storage Pool assignment.

Choose whether an existing pool must be expanded or whether a pool is to be created for the configuration, as shown in Figure 7-35.

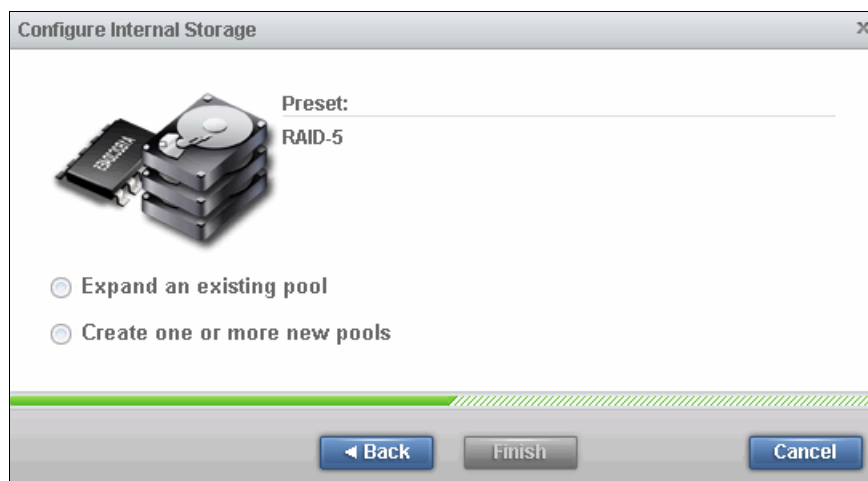


Figure 7-35 Storage Pool selection

Complete the following steps to expand or create a pool:

- a. Expand an existing pool.

When an existing pool is to be expanded, you can select an existing Storage Pool that does not contain MDisks, or a pool that contains MDisks with the same performance characteristics (which is listed automatically), as shown in Figure 7-36.

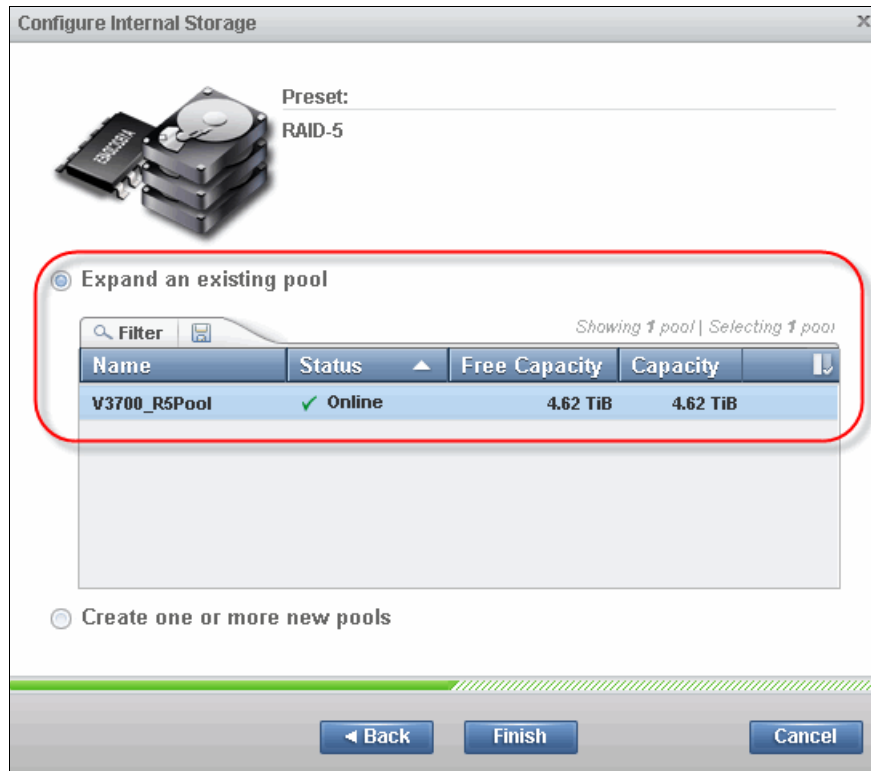


Figure 7-36 List of matching Storage Pool

- b. Create one or more pools.

Alternatively, a Storage Pool is created by entering the required name, as shown in Figure 7-37.

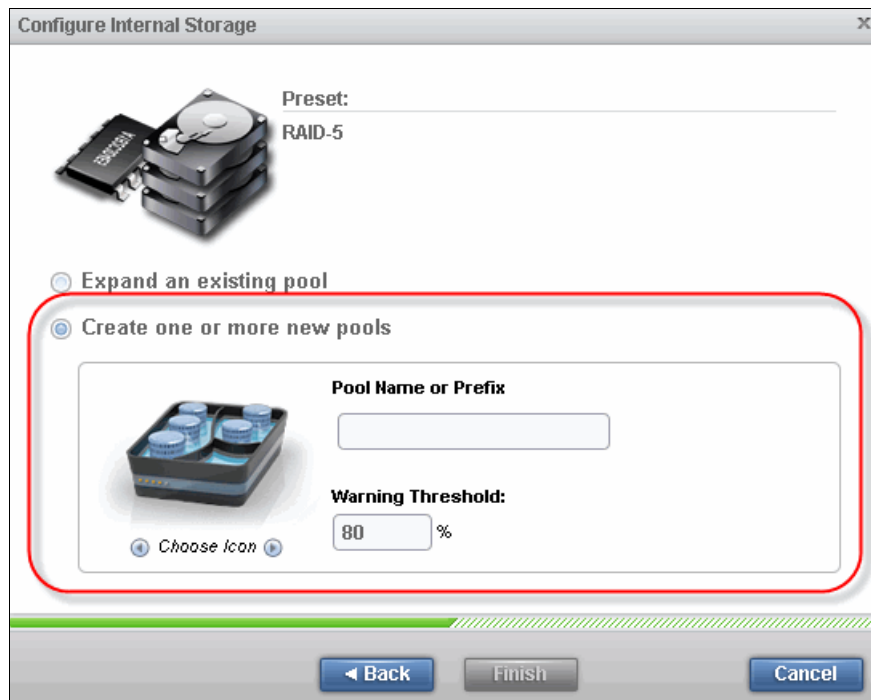


Figure 7-37 Create new pool

All drives are initialized when the Configuration wizard is finished.

7.2.6 MDisk by Pools pane

The MDisks by Pools pane (as shown in Figure 7-38) displays information about all MDisks made of internal and external storage. The MDisks are categorized by the pools to which they are attached.

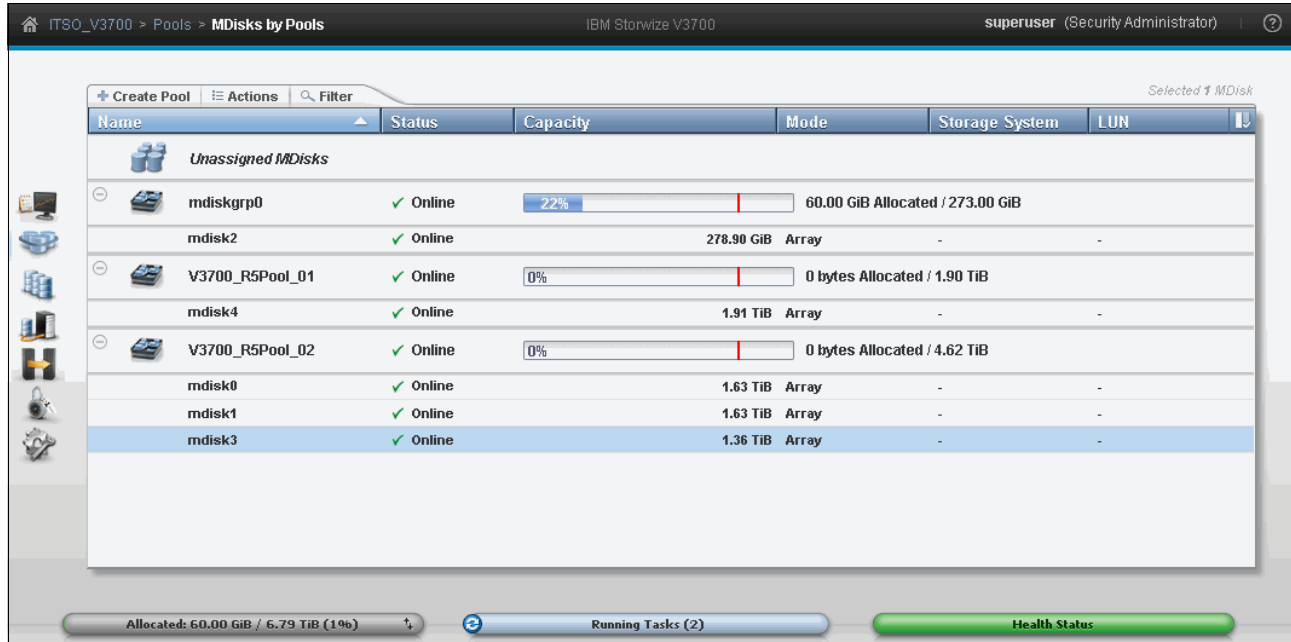


Figure 7-38 MDisk by Pool window

The following default information is provided:

- ▶ Name

The MDisk and the Pool name that is provided during the configuration process.

- ▶ Status

The status of the MDisk and Storage Pool. The following statuses are possible:

- Online

All MDisks are online and performing optimally.

- Degraded

One MDisk is in degraded state for example, missing SAS connection to an enclosure of member drives or a failed drive with no spare available. A node could also be a contributing factor to this state.

Figure 7-39 shows that the pool is degraded.

Name	Status	Capacity	Mode	Storage System
<i>Unassigned MDisks</i>				
mdiskgrp0	Online	22%	60.00 GiB Allocated / 273.00 GiB	
mdisk2	Online		278.90 GiB Array	-
V3700_R5Pool_01	Degraded	0%	0 bytes Allocated / 1.90 TiB	
mdisk4	Degraded		1.91 TiB Array	-
V3700_R5Pool_02	Online	0%	0 bytes Allocated / 4.62 TiB	
mdisk0	Online		1.63 TiB Array	-
mdisk1	Online		1.63 TiB Array	-
mdisk3	Online		1.36 TiB Array	-

Figure 7-39 One degraded MDisk in pool

– Offline

One or more MDisks in a pool are offline. The pool (V3700_R5Pool_02) also changes to offline, as shown in Figure 7-40.

Name	Status	Capacity	Mode	Storage System
<i>Unassigned MDisks</i>				
mdiskgrp0	Online	22%	60.00 GiB Allocated / 273.00 GiB	
mdisk2	Online		278.90 GiB Array	-
V3700_R5Pool_01	Degraded	0%	0 bytes Allocated / 1.90 TiB	
mdisk4	Degraded		1.91 TiB Array	-
V3700_R5Pool_02	Offline	0%	0 bytes Allocated / 4.62 TiB	
mdisk0	Online		1.63 TiB Array	-
mdisk1	Online		1.63 TiB Array	-
mdisk3	Offline		1.36 TiB Array	-

Figure 7-40 Offline MDisk in a pool

► Capacity

The capacity of the individual MDisks. The capacity for the Storage Pool, which is the total of all of the MDisks in this Storage Pool, is also shown. The usage of the Storage Pool is represented by a bar and the number, as shown in Figure 7-41.

Name	Status	Capacity	Mode	Storage System
<i>Unassigned MDisks</i>				
mdiskgrp0	Online	22%	60.00 GiB Allocated / 273.00 GiB	
mdisk2	Online		278.90 GiB Array	-
V3700_R5Pool_01	Online	50%	980.00 GiB Allocated / 1.90 TiB	
mdisk4	Online		1.91 TiB Array	-
V3700_R5Pool_02	Online	30%	1.41 TiB Allocated / 4.62 TiB	
mdisk0	Online		1.63 TiB Array	-
mdisk1	Online		1.63 TiB Array	-
mdisk3	Online		1.36 TiB Array	-

Figure 7-41 Capacity of the Pool

► Mode

An MDisk features the following modes:

- Array

Array mode MDisk are constructed from internal drives by using the RAID functionality. Array MDisks are always associated with Storage Pools.

- Unmanaged

The MDisk is not a member of any Storage Pools, which means it is not used by the IBM Storwize V3700 storage system. This status is an intermediate status of the migration process, and is described in Chapter 6, “Storage migration” on page 249.

- Image

When an unmanaged MDisk is imported to an IBM Storwize V3700 Storage Pool, the MDisk becomes an Image mode MDisk.

► Storage System

The name of the external storage system from which the logical unit number (LUN) is presented when working with storage migration.

► LUN

Represents the logical unit number of the MDisk.

Figure 7-42 shows an example of MDisk features.

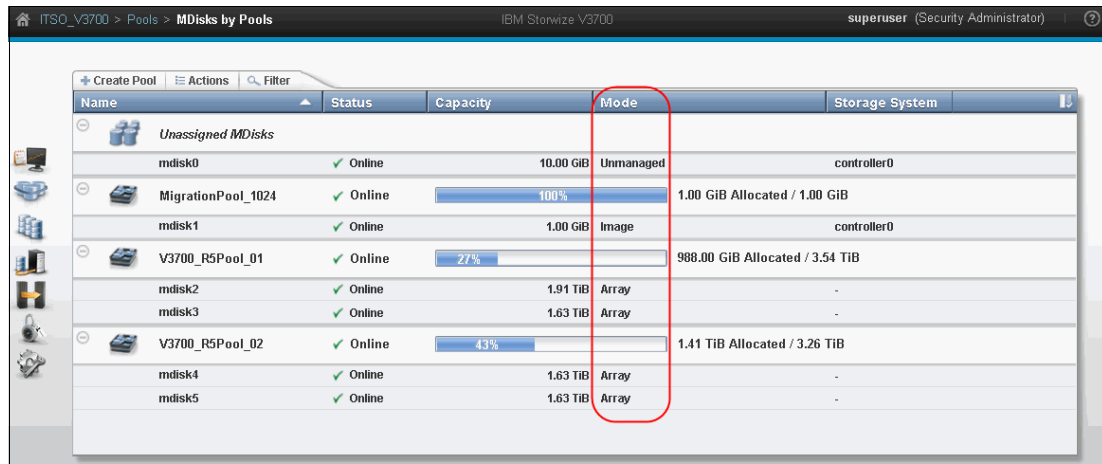


Figure 7-42 Working with MDisks

For more information about how to attach external storage to an IBM Storwize V3700 storage system, see Chapter 6, “Storage migration” on page 249.

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

Example 7-5 CLI `lsmdiskgrp` command

```
IBM_Storwize:ITS0_V3700:superuser>lsmdiskgrp
IBM_Storwize:ITS0_V3700:superuser>lsmdiskgrp mdiskgrpID
```

7.2.7 RAID action for MDisks

Internal drives in the IBM Storwize V3700 are managed as Array mode MDisks, on which several RAID actions can be performed:

1. Select the appropriate Array MDisk by clicking **Pools** → **MDisks by Pools**.
2. Then, click **Actions** → **RAID Actions**, as shown in Figure 7-43.

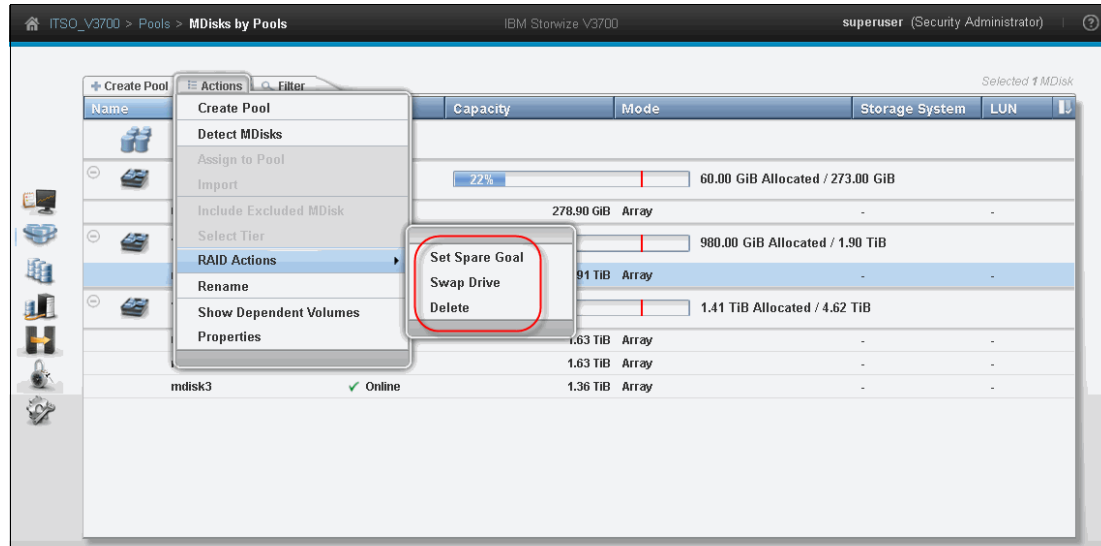


Figure 7-43 MDisk RAID actions

3. You can choose the following RAID actions:
 - a. Set Spare Goal:
 - i. Figure 7-44 shows how to, in the GUI, set the number of spare drives that are required to protect the array from drive failures.

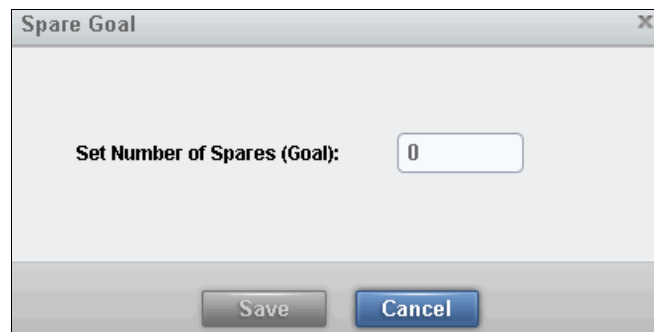


Figure 7-44 MDisk set spare goal

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

Example 7-6 CLI command to set spares

```
IBM_Storwize:ITS0_V3700:superuser>charray -sparegoal mdiskID goal
```

If the number of drives that are assigned as Spare does not meet the configured spare goal, an error is logged in the event log that reads: “Array MDisk is not protected by sufficient spares.” This error can be fixed by adding more drives as spares.

During the internal drive configuration, spare drives are automatically assigned according to the chosen RAID preset's spare goals (see 7.2, "Configuring internal storage" on page 320).

b. Swap Drive:

- i. The Swap Drive action can be used to replace a drive in the array with another drive with the status of Candidate or Spare. This action is used to replace a drive that failed, or is expected to fail soon (for example, as indicated by an error message in the event log). Select an MDisk that contains the drive to be replaced and click **RAID Actions** → **Swap Drive**.
- ii. In the Swap Drive window, select the member drive that is to be replaced (as shown in Figure 7-45) and click **Next**.

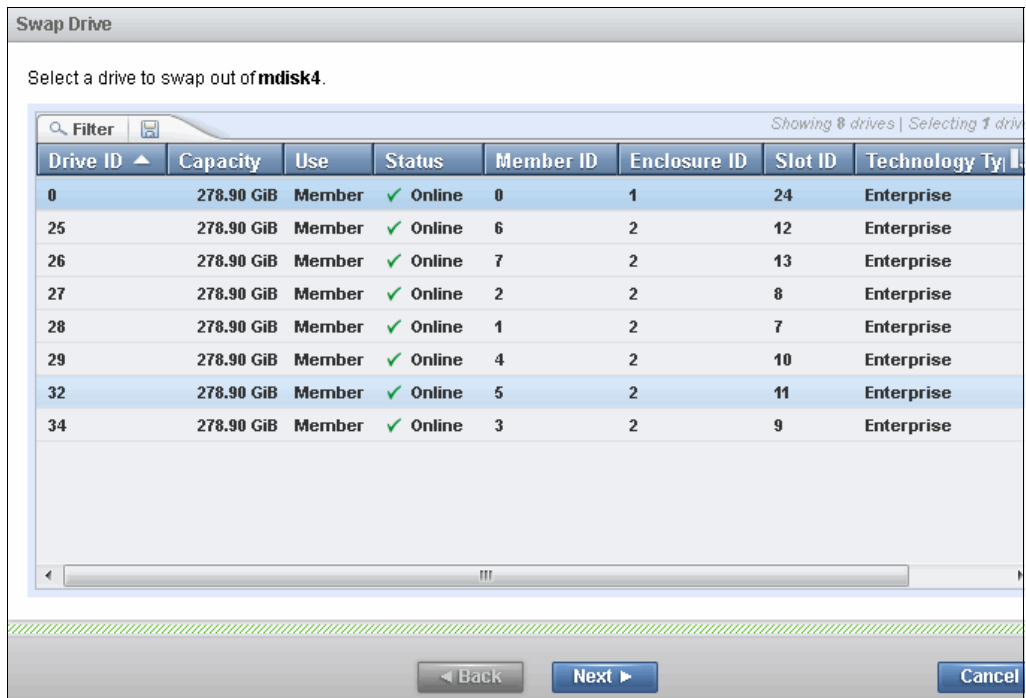


Figure 7-45 MDisk swap drive: Step 1

- iii. In step 2 (as shown as Figure 7-46), a list of suitable drives is presented. One drive must be selected to swap into the MDisk. Click **Finish**.

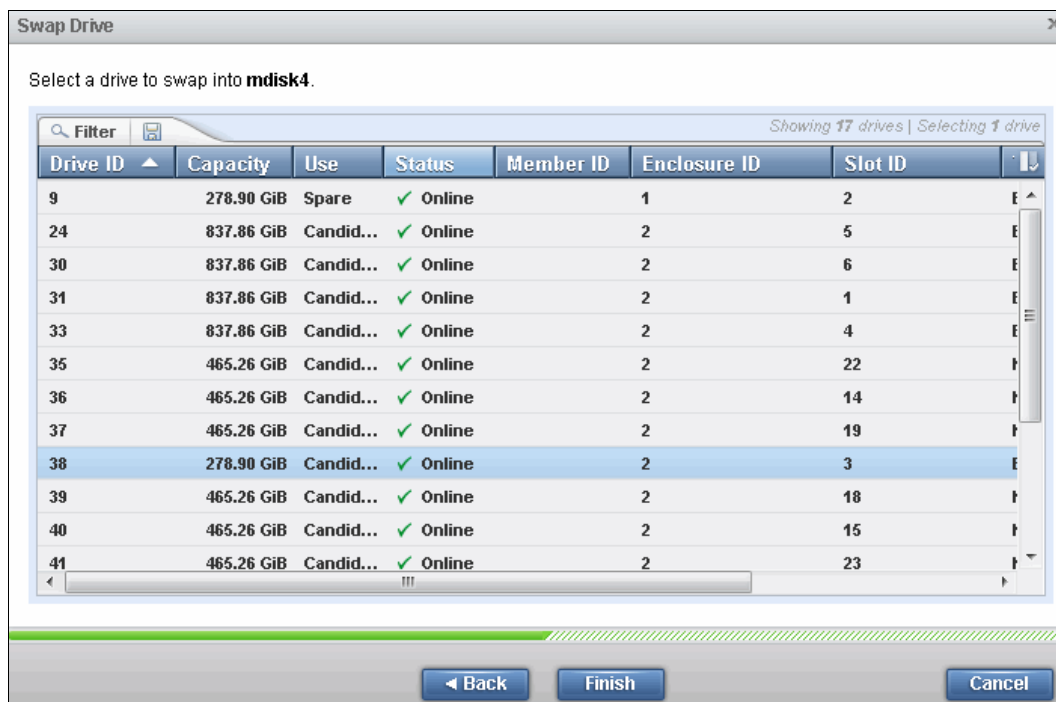


Figure 7-46 MDisk swap drive: Step 2

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

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

Example 7-7 CLI command to swap drives

```
IBM_Storwize:ITS0_V3700:superuser>charraymember -balanced -member
oldDriveID -newdrive newDriveID mdiskID
```

- c. Delete:
 - i. An Array MDisk can be deleted by clicking **RAID Actions** → **Delete**. To select more than one MDisk, press Ctrl+left-mouse click.
 - ii. A confirmation is required by entering the correct number of MDisks to be deleted, as shown in Figure 7-47. You must confirm the number of MDisks that you want to delete. If there is data on the MDisk, it can be deleted only by selecting the **Delete the RAID array MDisk even if it has data on it** option.

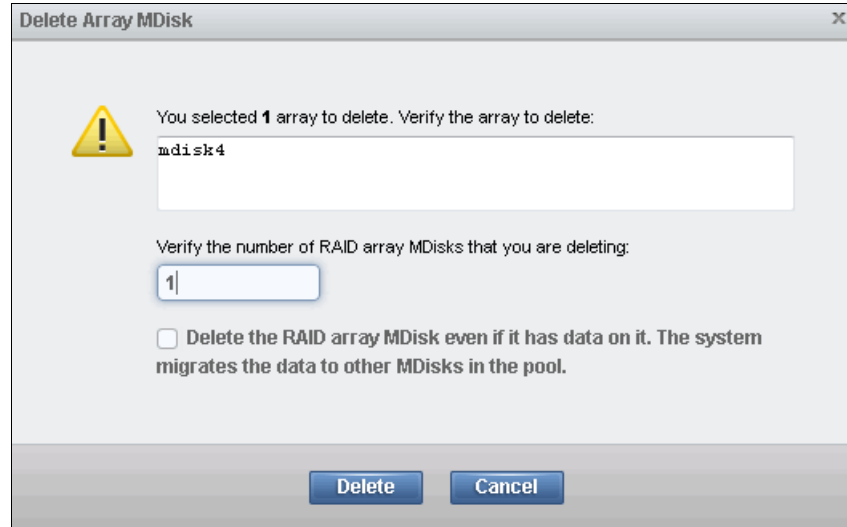


Figure 7-47 MDisk delete confirmation

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

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

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

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

Example 7-8 CLI command to delete MDisk

```
IBM_Storwize:ITS0_V3700:superuser>rmmdisk -mdisk list -force mdiskgrpID
```

If all the MDisks of a Storage Pool were deleted, the pool remains as an empty pool with 0 bytes of capacity, as shown in Figure 7-48.

Name	Status	Capacity	Mode	Storage System
<i>Unassigned MDisks</i>				
mdiskgrp0	✓ Online	22%		60.00 GiB Allocated / 273.00 GiB
V3700_R5Pool_01	✓ Online	0%		0 bytes Allocated / 0 bytes
V3700_R5Pool_02	✓ Online	30%		1.41 TiB Allocated / 4.62 TiB

Figure 7-48 Empty Storage Pool after MDisk deletion

7.2.8 More actions on MDisks

The following actions can be performed on MDisks:

1. Detect MDisks.

The **Detect MDisks** action, available from the **Actions** menu at the upper left of the MDisks by Pools window, is useful if you have external storage controllers in your environment (for more information, see Chapter 6, “Storage migration” on page 249).

The **Detect MDisks** action starts a rescan of the fibre channel (FC) or SAS network. It discovers any new LUNs that were mapped to the IBM Storwize V3700 storage system and rebalances MDisk access across the available controller device ports. This action also detects any loss of controller port availability, and updates the IBM Storwize V3700 configuration to reflect any changes.

When external storage controllers are presented to the IBM Storwize V3700 environment, the IBM Storwize V3700 automatically discovers the controllers, and the LUNs that are presented by those controllers are listed as unmanaged MDisks. However, if you attached new storage and the IBM Storwize V3700 did not detect it, you might need to use the **Detect MDisks** action before the system detects the new LUNs.

If the configuration of the external controllers is modified afterward, the IBM Storwize V3700 might be unaware of these configuration changes. Use the **Detect MDisks** action to rescan the FC network and the SAS ports to update the list of unmanaged MDisks. Figure 7-49 shows the Detect MDisks action.

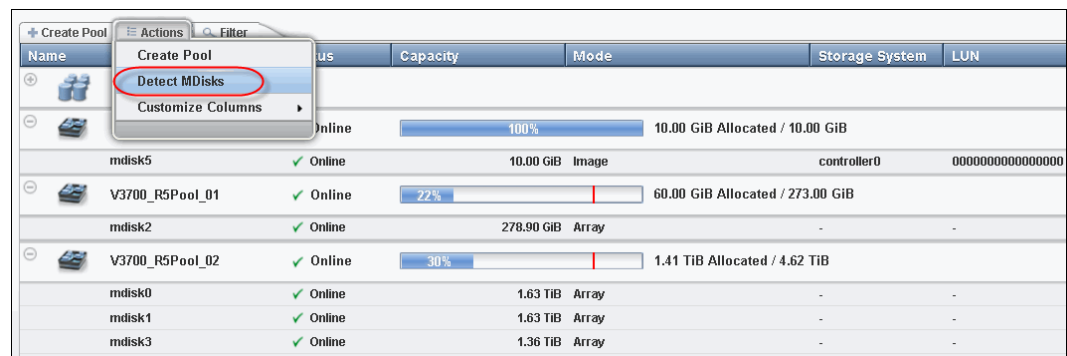


Figure 7-49 Detect MDisks

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

2. Include Excluded MDisks.

An MDisk can be excluded from the IBM Storwize V3700 because of multiple input/output (I/O) failures. These failures might be caused, for example, by link errors. After a fabric-related problem is fixed, the excluded disk can be added back into the IBM Storwize V3700:

- Select the MDisks.
- Click **Include Excluded MDisk** from the Actions menu.

3. Select Tier.

Internal drives and external MDisks have their tier assigned automatically by the IBM Storwize V3700. The tier of an internal MDisk cannot be changed, but if the disk tier of an external MDisk is recognized incorrectly by the IBM Storwize V3700, use this option to manually change the disk tier.

To specify a new tier, select the MDisk and click **Actions** → **Select Tier**, as shown in Figure 7-50.

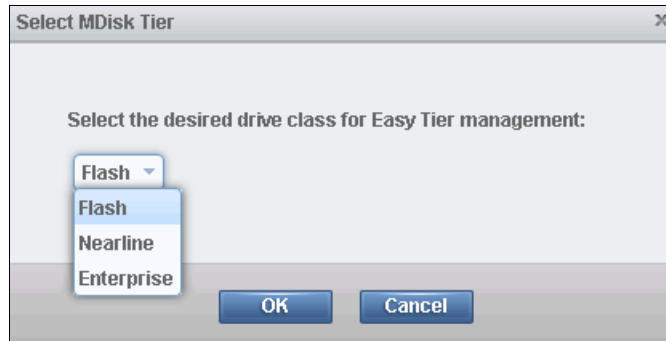


Figure 7-50 Select the wanted disk tier

Some of the other actions are available by clicking **MDisk by Pool** → **Actions**, as shown in Figure 7-51.

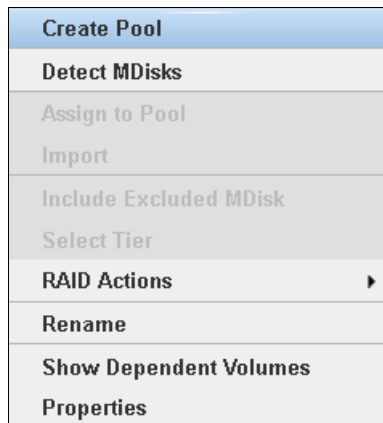


Figure 7-51 MDisk actions on externally virtualized storage

You can perform the following additional actions:

1. Rename:
 - a. Select the MDisk and clicking **Rename** from the Actions menu.
 - b. Enter the new name of your MDisk (as shown in Figure 7-52) and click **Rename**.



Figure 7-52 Rename MDisk

2. Show Dependent Volumes:
 - a. Figure 7-53 shows the volumes that are dependent on an MDisk.
 - b. The volumes can be displayed by selecting the MDisk and clicking **Show Dependent Volumes** from the Actions menu. The volumes are listed with general information.




Figure 7-53 Show dependent volumes

- 3. Display Properties:
 - a. The Properties action for an MDisk shows the information that you need to identify it. In the MDisks by Pools window, select the MDisk and click **Properties** from the Actions menu.
 - b. The following tabs are available in this information window:
 - i. The Overview tab (as shown in Figure 7-54) contains information about the MDisk. To show more details, click **Show Details**.

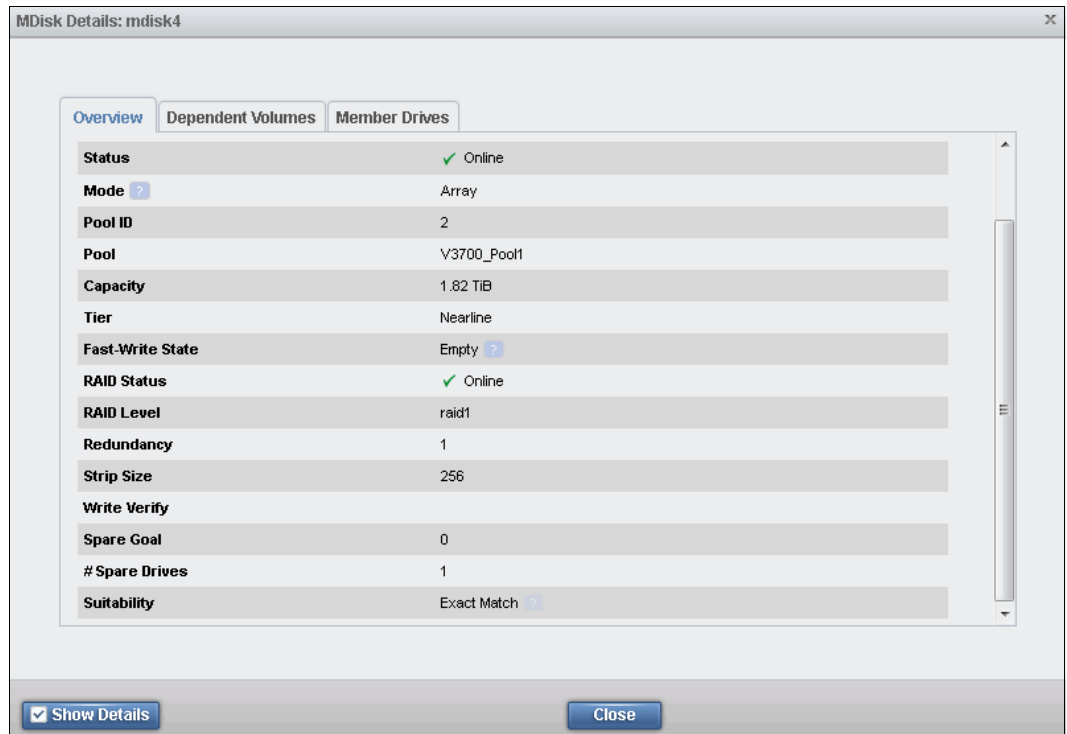


Figure 7-54 MDisk properties overview

- ii. The Dependent Volumes tab (as shown in Figure 7-55) lists all of the volumes that use extents on this MDisk.




Figure 7-55 MDisk dependent volumes

- iii. In the Member Drives tab (as shown in Figure 7-56), you find all of the member drives of this MDisk. Also, all of the actions that are described in 7.1.2, “Actions on internal drives” on page 311, can be performed on the drives that are listed here.




Figure 7-56 MDisk properties member

7.3 Working with Storage Pools

The next sections provides a detailed overview about the concept and use of the Storage Pools and Child Pools.

Storage Pools (or Pools) act as a container for MDisks and provision the capacity to volumes. IBM Storwize V3700 organizes storage in to Storage Pools to ease storage management and make it more efficient. Child Pools act as subcontainers, and they inherit the parent Storage Pool’s properties. They can also be used to provision capacity to volumes. For more information about the Child Pools, see 7.4, “Working with Child Pools” on page 355.

Storage Pools and MDisks are managed from the MDisks by Pools window:

1. You can access the MDisks by Pools window by clicking the upper right option **Overview** → **Pools**, as shown in Figure 7-57.

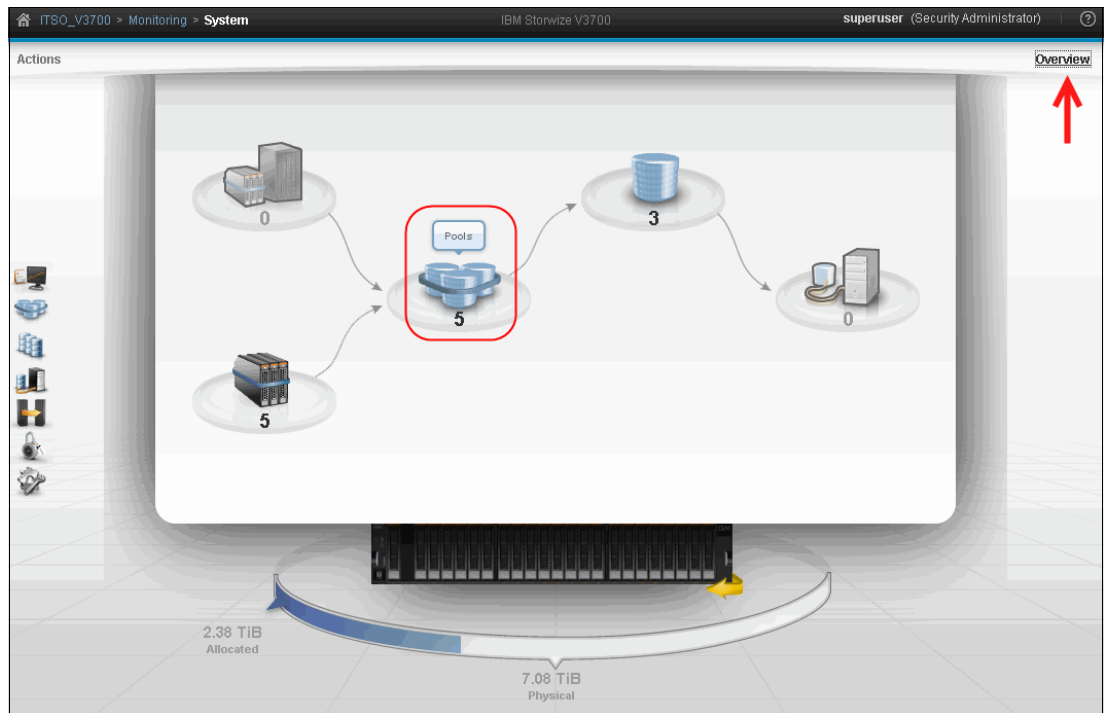


Figure 7-57 Pools from the overview window

2. An alternative path to the Pools window is to click **Pools** → **MDisks by Pools**, as shown in Figure 7-58.




Figure 7-58 Pools from MDisk by Pools window

- Using the MDisk by Pools window (as shown in Figure 7-59), you can manage internal and external Storage Pools. All existing Storage Pools are displayed row-by-row. The first row features Unassigned MDisks, which contains all unmanaged MDisks, if any exist. Each defined Storage Pool is displayed with the following information:
 - Assigned icon and name
 - Numerical ID
 - Status
 - A graphical indicator of pool capacity percentage that is allocated to volumes

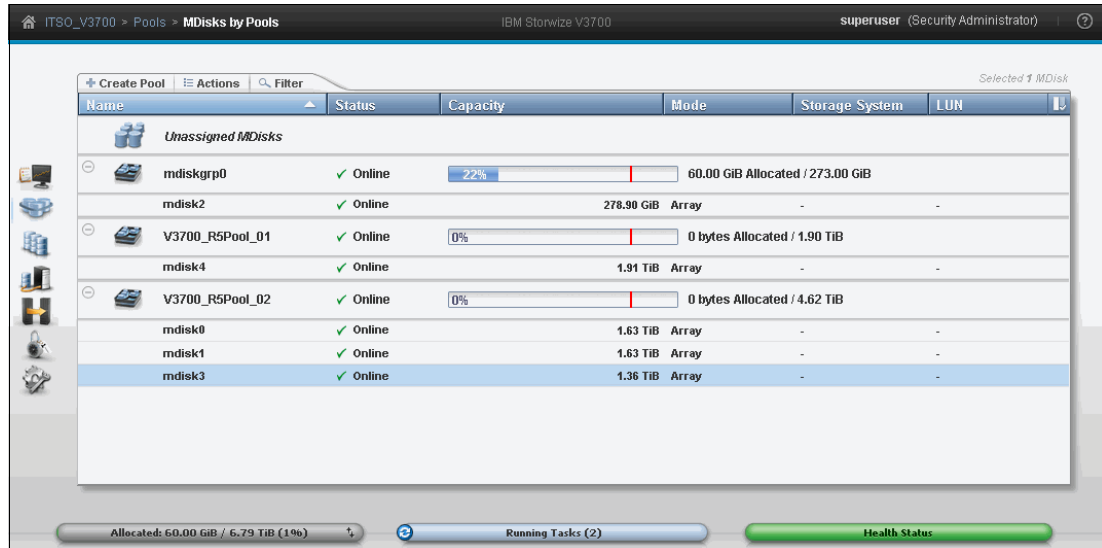


Figure 7-59 Pool window

- When you expand a pool's entry by clicking the plus sign (+) to the left of the pool's icon, you can access the MDisks that are associated with this pool. You can perform actions on these MDisks as described in 7.2, "Configuring internal storage" on page 320.

7.3.1 Create Pool option

New Storage Pools are built when an MDisk is created if this MDisk is not attached to an existing pool.

To create an empty pool, complete the following steps:

1. Click the **Create Pool** option in the pool window. The only required parameter for the pool is the pool name, as shown in Figure 7-60.



Figure 7-60 Create pool name input

2. The IBM Storwize V3700 GUI has a default extent size value of 1 GiB when you define a new Storage Pool. Therefore, if you are creating an empty Pool using the GUI, and you want to specify the size of the extents, you must ensure that the **Allow extent size selection during pool creation** option is selected in the **Settings** → **GUI Preferences**. The size of the extent is selected at creation, and it cannot be changed later. Figure 7-61 shows an example of how to create a Pool choosing an extent size.

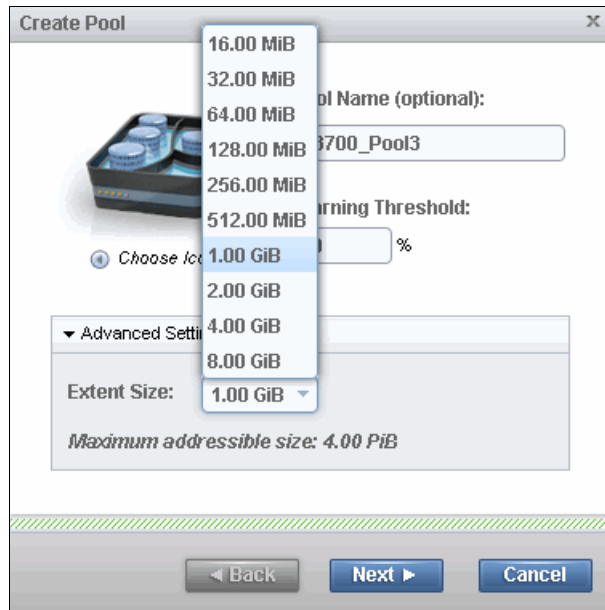
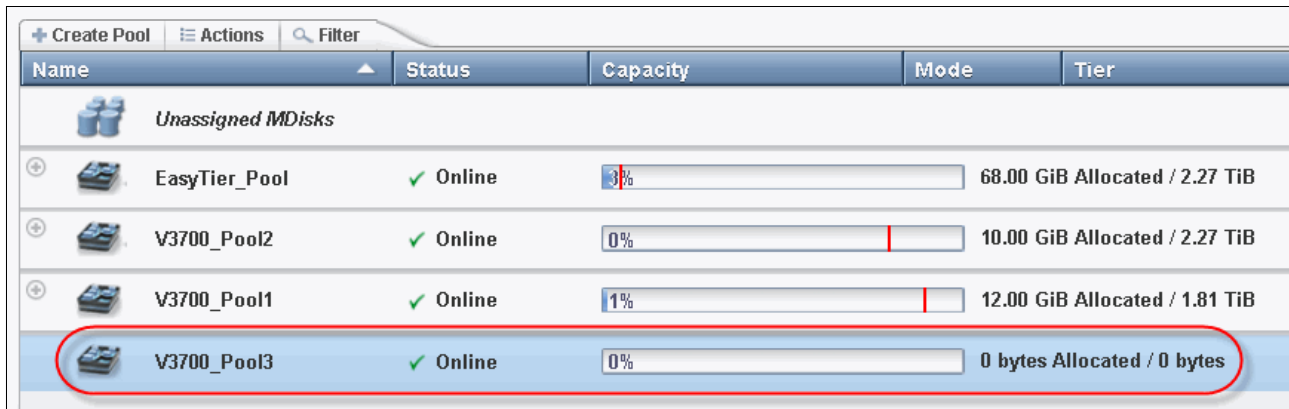


Figure 7-61 Creating a pool choosing the extent size

The new pool is included in the pool list with 0 bytes, as shown in Figure 7-62.

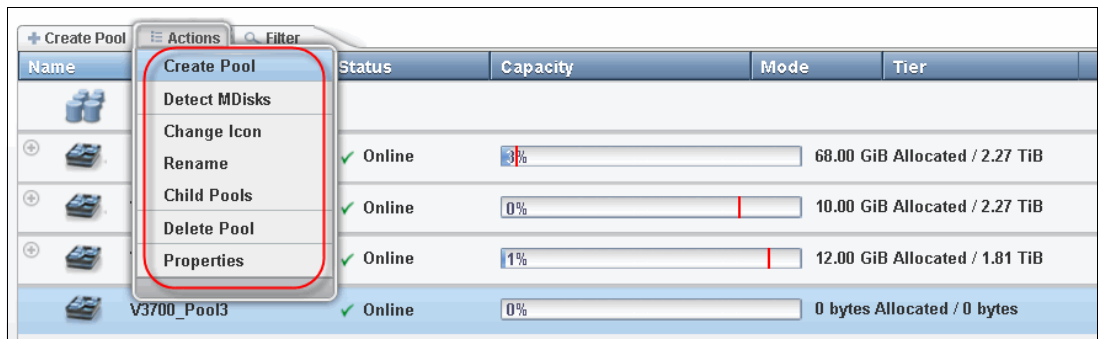


Name	Status	Capacity	Mode	Tier
Unassigned MDisks				
EasyTier_Pool	✓ Online	3%	68.00 GiB Allocated / 2.27 TiB	
V3700_Pool2	✓ Online	0%	10.00 GiB Allocated / 2.27 TiB	
V3700_Pool1	✓ Online	1%	12.00 GiB Allocated / 1.81 TiB	
V3700_Pool3	✓ Online	0%	0 bytes Allocated / 0 bytes	

Figure 7-62 Creating an empty pool

7.3.2 Actions on Storage Pools

There are several actions that can be performed on Storage Pools using the Actions menu, as shown in Figure 7-63.



Name	Status	Capacity	Mode	Tier
Actions Menu: <ul style="list-style-type: none">Create PoolDetect MDisksChange IconRenameChild PoolsDelete PoolProperties				
EasyTier_Pool	✓ Online	3%	68.00 GiB Allocated / 2.27 TiB	
V3700_Pool2	✓ Online	0%	10.00 GiB Allocated / 2.27 TiB	
V3700_Pool1	✓ Online	1%	12.00 GiB Allocated / 1.81 TiB	
V3700_Pool3	✓ Online	0%	0 bytes Allocated / 0 bytes	

Figure 7-63 Pool action overview

Create Pool

As shown in 7.3.1, “Create Pool option” on page 351, this option is used to create an empty Pool.

Detect MDisks

As detailed in 7.2.8, “More actions on MDisks” on page 343, the Detect MDisk option starts a rescan of the FC network and SAS connections. It discovers any new MDisks that were recently mapped to the IBM Storwize V3700.

Change Icon

There are several Storage Pool icons available that can be selected, as shown in Figure 7-64. These icons can be used to differentiate between different storage tier, RAID level, or types of drives.



Figure 7-64 Change Storage Pool icon

Rename

The Storage Pool can be renamed at any time, as shown in Figure 7-65.



Figure 7-65 Rename Storage Pool

Delete

A pool can be deleted using the GUI only if there are no volumes assigned to it. A pane displays to confirm that all associated MDisks will be removed, and the associated drives will become candidate drives, as shown in Figure 7-66.




Figure 7-66 Confirmation to delete the Storage Pool

Through the CLI, you can delete a Storage Pool and all of its contents using the **-force** parameter. However, all of the volumes and host mappings are deleted without being able to recover them.

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

After you delete the pool, all of the managed or image mode MDisks in the pool return to a status of unmanaged.

Child Pools

Figure 7-67 shows the list of Child Pools that are dependent on a storage parent pool. A Storage Pool can be either a *Parent* or *Child* type. In our example, we show the Child Pools underlying on the V3700_Pool2 parent Storage Pool. For more information about Child Pools, see 7.4, “Working with Child Pools”.

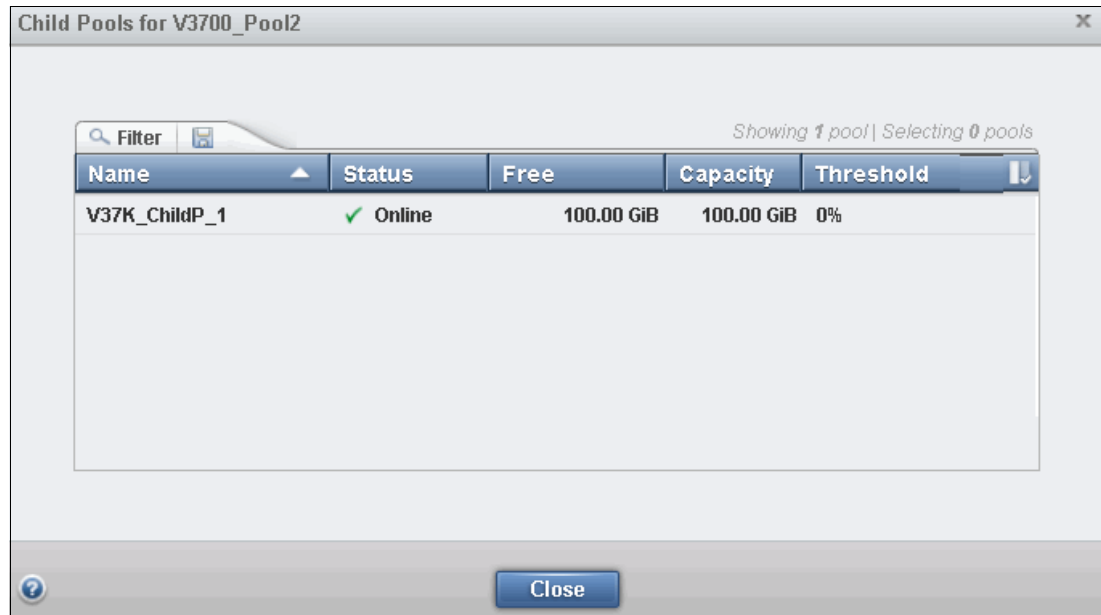


Figure 7-67 Child Pools list

7.4 Working with Child Pools

A Child Pool is an object similar to a Storage Pool, and is a user-configurable object. However, a user can create a Child Pool through the CLI only.

A storage parent pool is a collection of MDisks that provides the pool of storage from which volumes are provisioned. A Child Pool is in, and gets capacity exclusively from, one Storage Pool. A user can specify the Child Pool capacity at creation. The Child Pool can grow and decrease non-disruptively.

The Child Pool shares the properties of the parent Storage Pool, and provides most of the functions that Storage Pools have. The Child Pool has the following common properties and behaviors:

- ▶ ID
- ▶ Name
- ▶ Status
- ▶ Volume Count
- ▶ Warning Thresholds (Child Pools can have their own threshold warning setting)
- ▶ MDisk (this value is always zero “0” for child pools)
- ▶ Space Utilization (in percent)
- ▶ Easy Tier (for a Child Pool, this value must be the same as the parent mdiskgrp’s Easy Tier setting)

Figure 7-68 illustrates the concept of Child Pools.

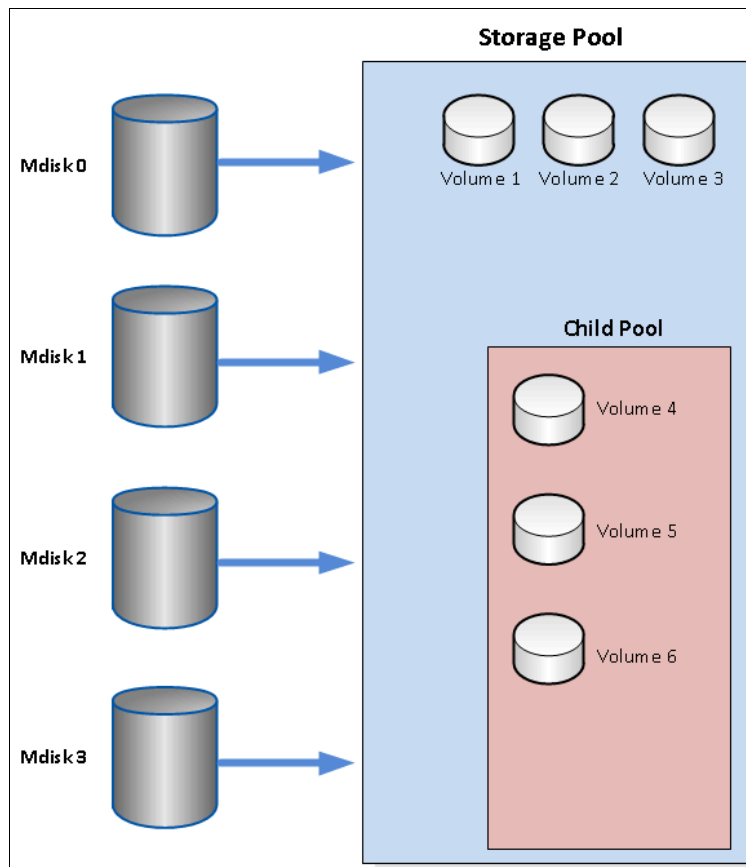


Figure 7-68 Child Pool representation

7.4.1 Create Child Pools

Example 7-9 demonstrates how to create a Child Pool using the CLI.

Example 7-9 Creating a Child Pool

```
IBM_Storwize:ITS0_V3700:superuser>svctask mkmdiskgrp -name V37K_ChildP_1 -warning  
80 -parentmdiskgrp V3700_Pool2 -unit gb -size 100  
MDisk Group, ID [0], successfully created
```

The following parameters must be specified when creating Child Pools:

- ▶ **-name** (Child Pool name)
- ▶ **-warning** (Generates a warning when the used disk capacity exceeds the specified threshold)
- ▶ **-parentmdiskgrp** (The ID or name of the Parent Pool)
- ▶ **-unit** (Specifies the data unit; it must be in multiples of extent size of the parent pool)
- ▶ **-size** (Specifies the capacity of the Child Pool)

Note: The capacity of the Child Pool is always limited by the free capacity of the parent pool.

Because a Child Pool must have a unique Parent Pool, the maximum number of Child Pools per parent pool is 127.

For more information about the limits and restrictions, see the following website:

<http://www-01.ibm.com/support/docview.wss?uid=ssg1S1004923>

7.4.2 Actions on Child Pools

By selecting a Child Pool and clicking the action menu, only the following actions are applicable to Child Pools:

- ▶ Change Icon

As described in “Change Icon” on page 353, this option helps to manage and differentiate between pools and drive classes.

- ▶ Rename

Select this option to rename a Child Pool. Actions are explained in “Rename” on page 353.

- ▶ Delete Pool

If the Child Pool is not provisioning any capacity to volumes, use this option to delete a Child Pool.

- ▶ Properties

The properties of the Child Pool can be seen by choosing this option.

7.4.3 Resizing a Child Pool

A Child Pool can be easily expanded and reduced. However, re-sizing a Child Pool can only be performed through CLI:

1. If you are certain that the Parent Pool has sufficient free extents, you can use the **svctask chmdiskgrp** command with the **-size** parameter to set the new size of the Child Pool.
2. The capacity of the Child Pool can also be reduced. You cannot reduce the size of the Child Pool to less than its used capacity.

Assuming that the Child Pool capacity is 100.00 GB, Example 7-10 shows how to use the CLI to increase the capacity by 20 GB.

Example 7-10 Expanding Child Pool

```
IBM_Storwize:ITS0_V3700:superuser>svctask chmdiskgrp -unit gb -size 120
V37K_ChildP_1
```

3. After you expand the size of a Child Pool, you might need to run `svcinfo lsmdiskgrp` to have a detailed view of Child Pool, as shown in Example 7-11.

Example 7-11 View of the Child Pool

```
IBM_Storwize:ITS0_V3700:superuser>svcinfo lsmdiskgrp V37K_ChildP_1
id 0
name V37K_ChildP_1
status online
mdisk_count 0
vdisk_count 0
capacity 120.00GB
extent_size 1024
free_capacity 120.00GB
virtual_capacity 0.00MB
used_capacity 0.00MB
real_capacity 0.00MB
overallocation 0
warning 80
easy_tier auto
easy_tier_status active
tier ssd
tier_mdisk_count 0
tier_capacity 0.00MB
tier_free_capacity 0.00MB
tier enterprise
tier_mdisk_count 0
tier_capacity 0.00MB
tier_free_capacity 0.00MB
tier nearline
tier_mdisk_count 0
tier_capacity 0.00MB
tier_free_capacity 0.00MB
site_id
site_name
parent_mdisk_grp_id 3
parent_mdisk_grp_name V3700_Pool2
child_mdisk_grp_count 0
child_mdisk_grp_capacity 0.00MB
type child_thick
```

4. Assuming that the child pool capacity is 120 GB and you want to reduce the capacity by 40 GB, use the `svctask chmdiskgrp` command with the `-size` parameter to reduce the capacity to 80 GB, as shown in Example 7-12.

Example 7-12 Shrinking a Child Pool

```
IBM_Storwize:ITS0_V3700:superuser>svctask chmdiskgrp -size 80 -unit gb
V37K_ChildP_1
```



Advanced host and volume administration

The IBM Storwize V3700 offers many functions for volume and host configuration. The basic host and volume features of IBM Storwize V3700 are described in Chapter 4, “Host configuration” on page 149 and Chapter 5, “Volume configuration” on page 199. These chapters also describe how to create hosts and volumes and how to map them to a host.

This chapter includes the following topics:

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

8.1 Advanced host administration

This section describes host administration, including host modification, host mappings, and deleting hosts. Basic host creation and mapping are described in Chapter 4, “Host configuration” on page 149. It is assumed that hosts have already been defined, and that volumes are mapped to them.

The following topics are covered in this section:

- ▶ Modifying Hosts, as described in 8.1.1, “Modifying Mappings menu” on page 362
- ▶ Ports by Host, as described in 8.2, “Adding and deleting host ports” on page 377
- ▶ Host Mappings, as described in 8.3, “Host mappings overview” on page 383

The Hosts menu is shown in Figure 8-1.

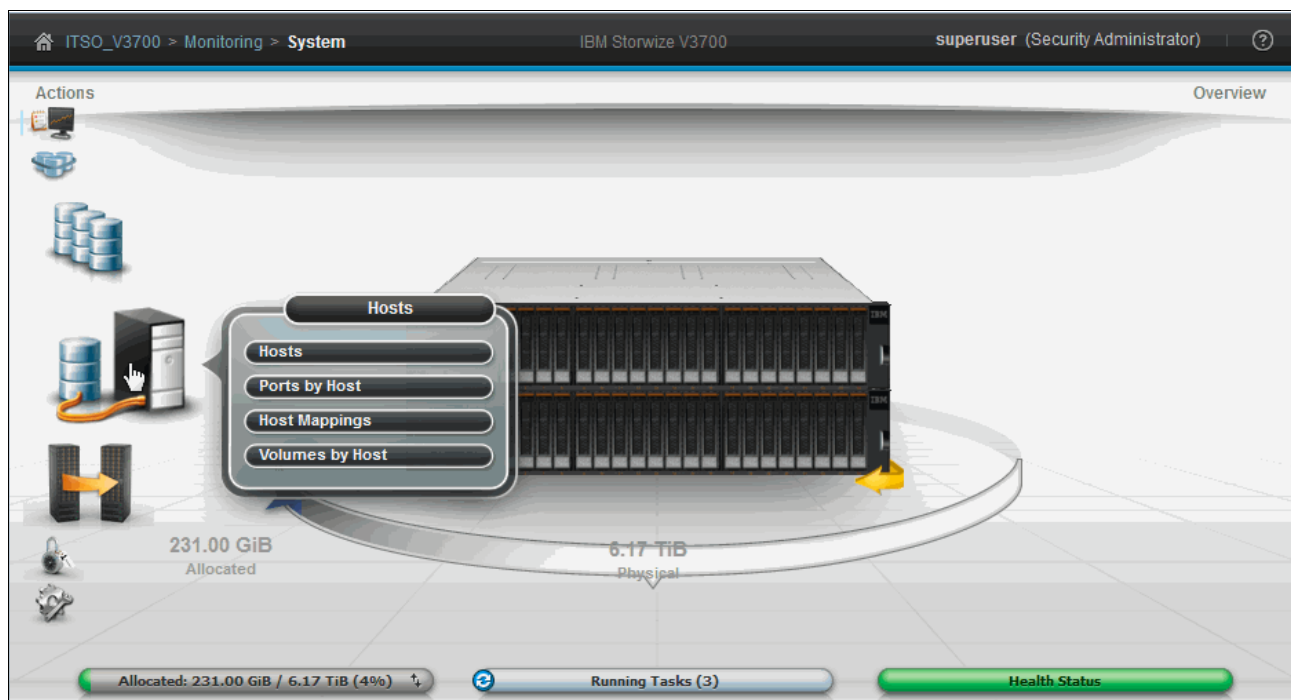


Figure 8-1 Host menu

To access the available host Actions, complete the following steps:

1. Select **Hosts** to open the Hosts pane, as shown in Figure 8-2.

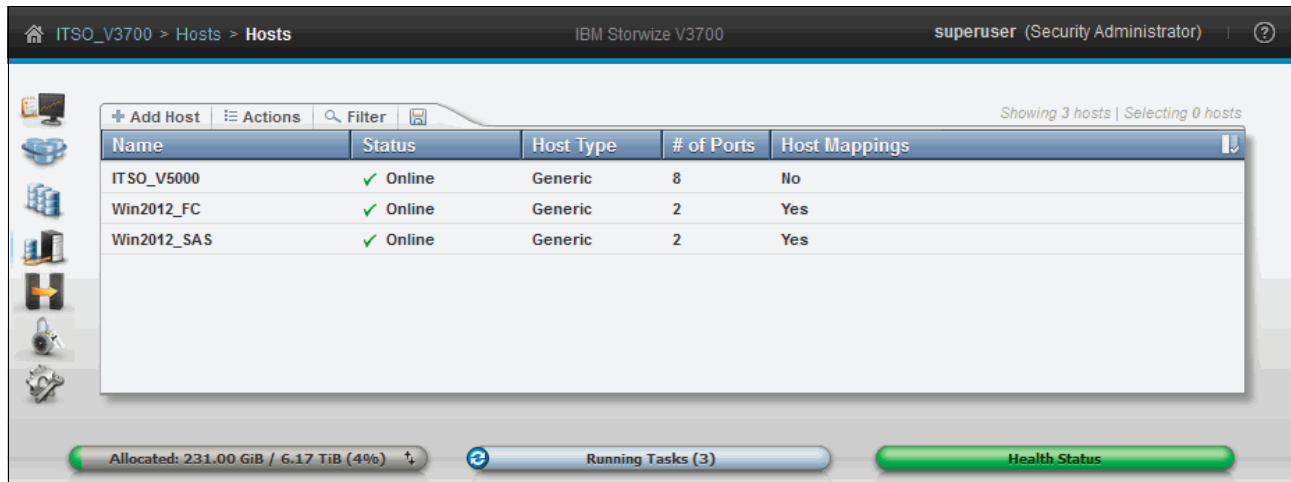


Figure 8-2 Hosts pane

Figure 8-2, shows several hosts, some with volumes mapped to them.

2. Select a host and click **Actions** (as shown in Figure 8-3), or right-click the host, to show the available actions.

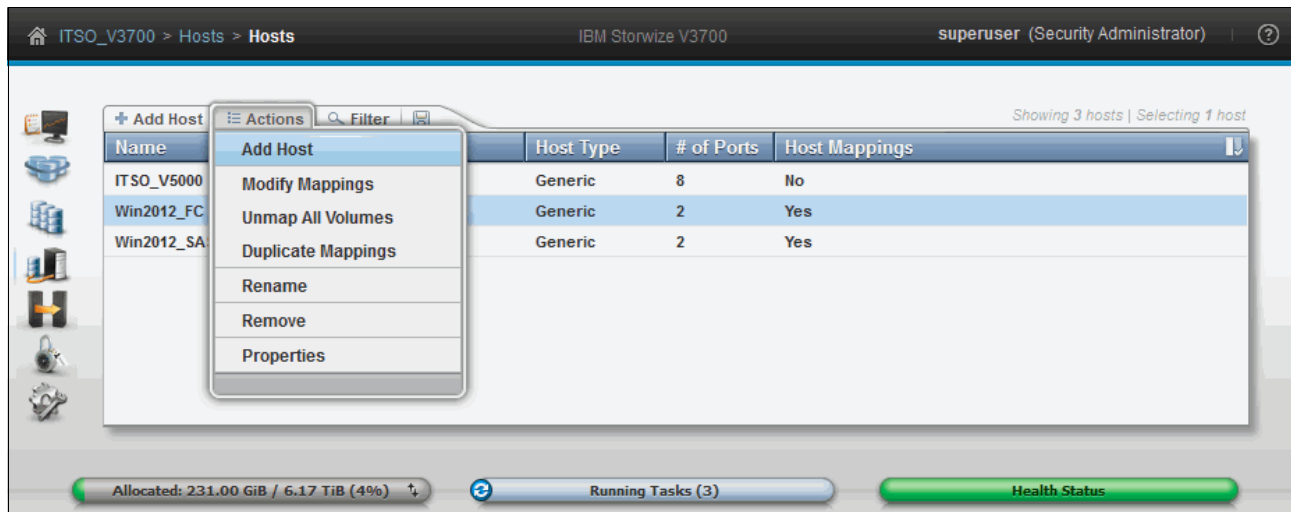


Figure 8-3 Host: Action menu

As shown in Figure 8-3, there are several Actions associated with host mapping. For more information, see 8.1.1, “Modifying Mappings menu” and 8.1.2, “Unmapping volumes from a host” on page 367.

8.1.1 Modifying Mappings menu

To modify mappings, complete the following steps:

1. From the Host pane, after selecting a host and opening the Actions menu, select **Modify Mappings** to open the Modify Host Mappings pane, as shown in Figure 8-4.




Figure 8-4 Host mappings pane

The Host menu lists the hosts that are attached to the IBM Storwize V3700.

The left pane (Unmapped Volumes) shows the volumes that are available for mapping to the chosen host. The right pane (Volumes Mapped to the Host) shows the volumes that are already mapped. Figure 8-4 shows that the volumes named FlashCopyVol0, FlashCopyVol1, FlashCopyVol2, and FlashCopyVol6 with Small Computer System Interface (SCSI) ID's 0, 1, 2, and 3 are mapped to the Win2012_FC host. The left pane shows that a further 13 volumes are available for mapping.

Important: The Unmapped Volumes pane refers to volumes that are not mapped to the chosen host. However, they might be mapped to other hosts.

- To map further volumes to a chosen host, select the volume from the left pane, as shown in Figure 8-5.

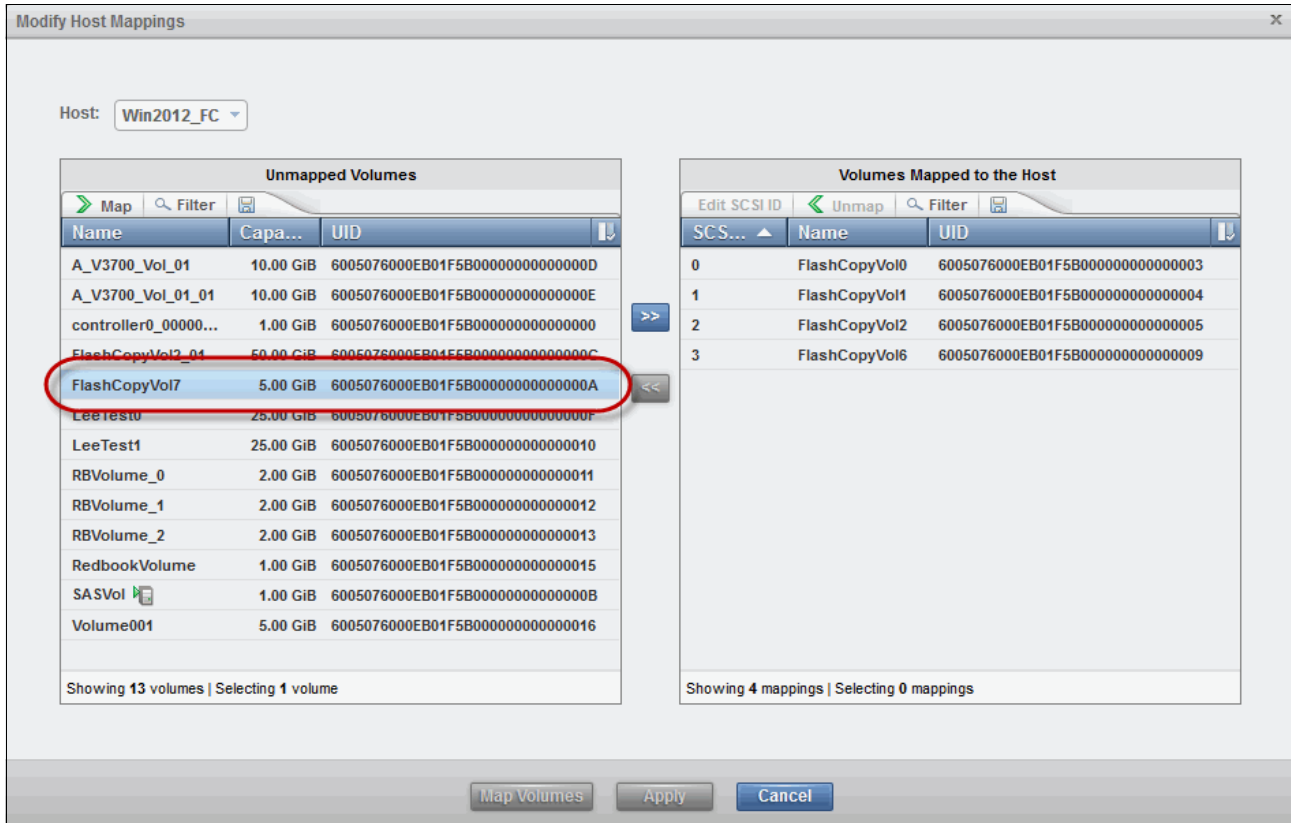


Figure 8-5 Modify Host Mappings

- Click the Right Arrow icon to move the volume to the right pane, as shown in Figure 8-6. The changes are marked in yellow, and the **Map Volumes** and **Apply** buttons are enabled.

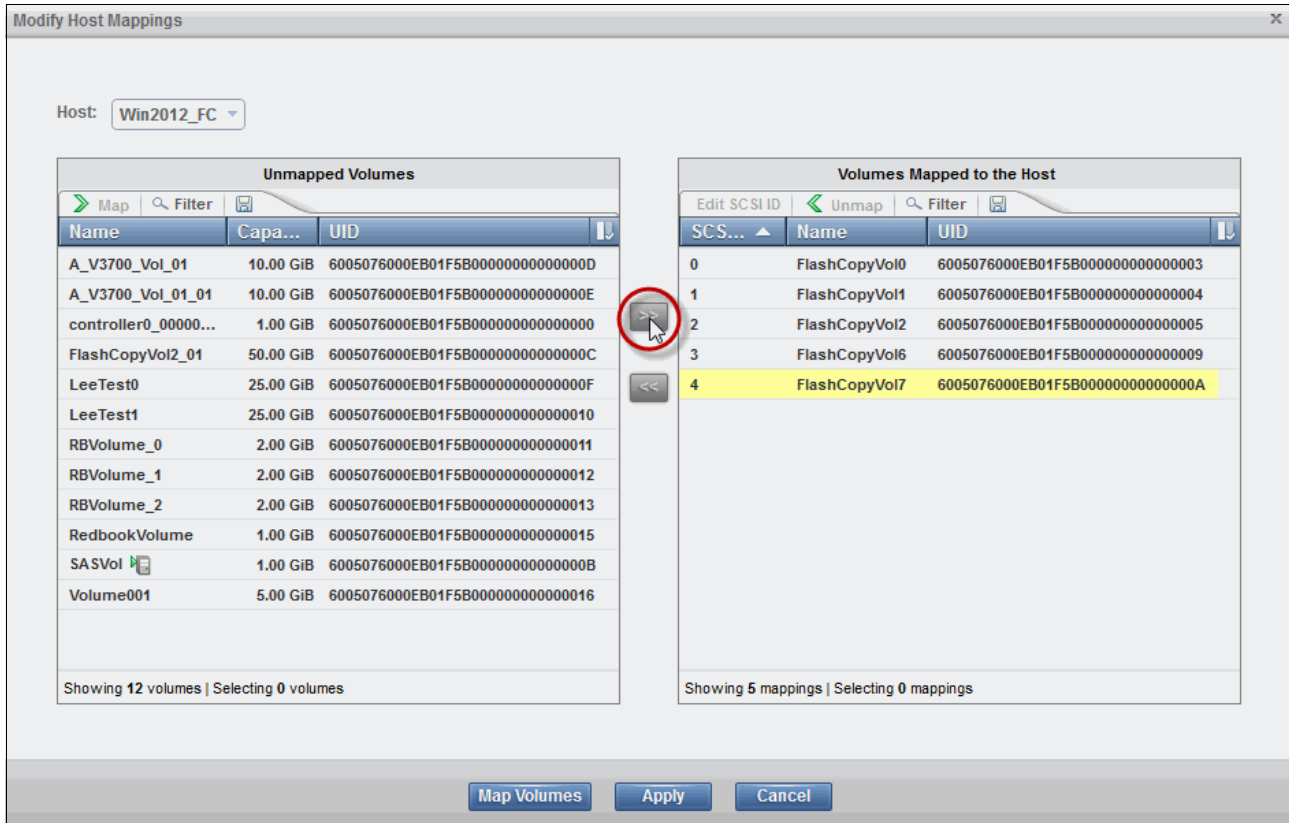


Figure 8-6 Modify Host Mappings

- Click **Map Volumes** to apply the changes, and the Modify Mappings pane shows that the task completed successfully, as shown in Figure 8-7.

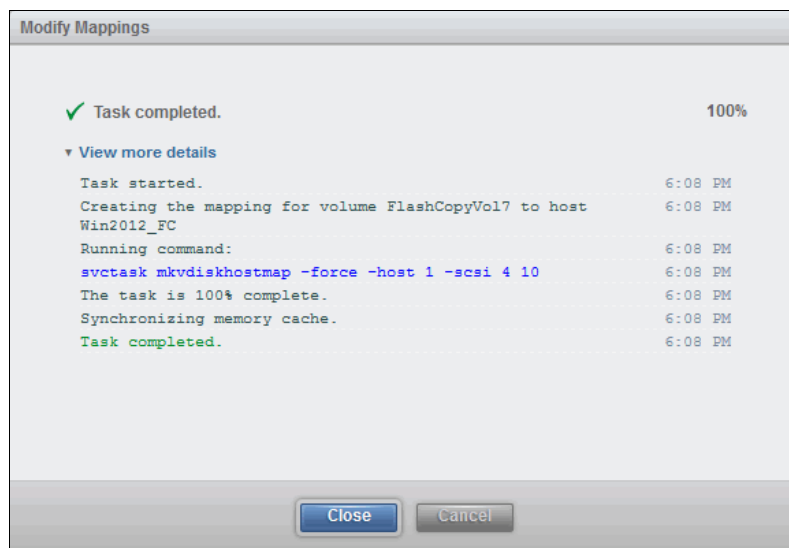


Figure 8-7 Modify Mappings task completed

- Click **Close**, and the Modify Mapping pane closes.

6. If **Apply** had been clicked in the pane shown in Figure 8-6 on page 364, the changes would have been submitted to the system, but the Modify Host Mappings pane remains open for further changes.
7. From the Modify Host Mappings pane, it is possible to modify another host by selecting it from the Hosts menu, as shown in Figure 8-8.

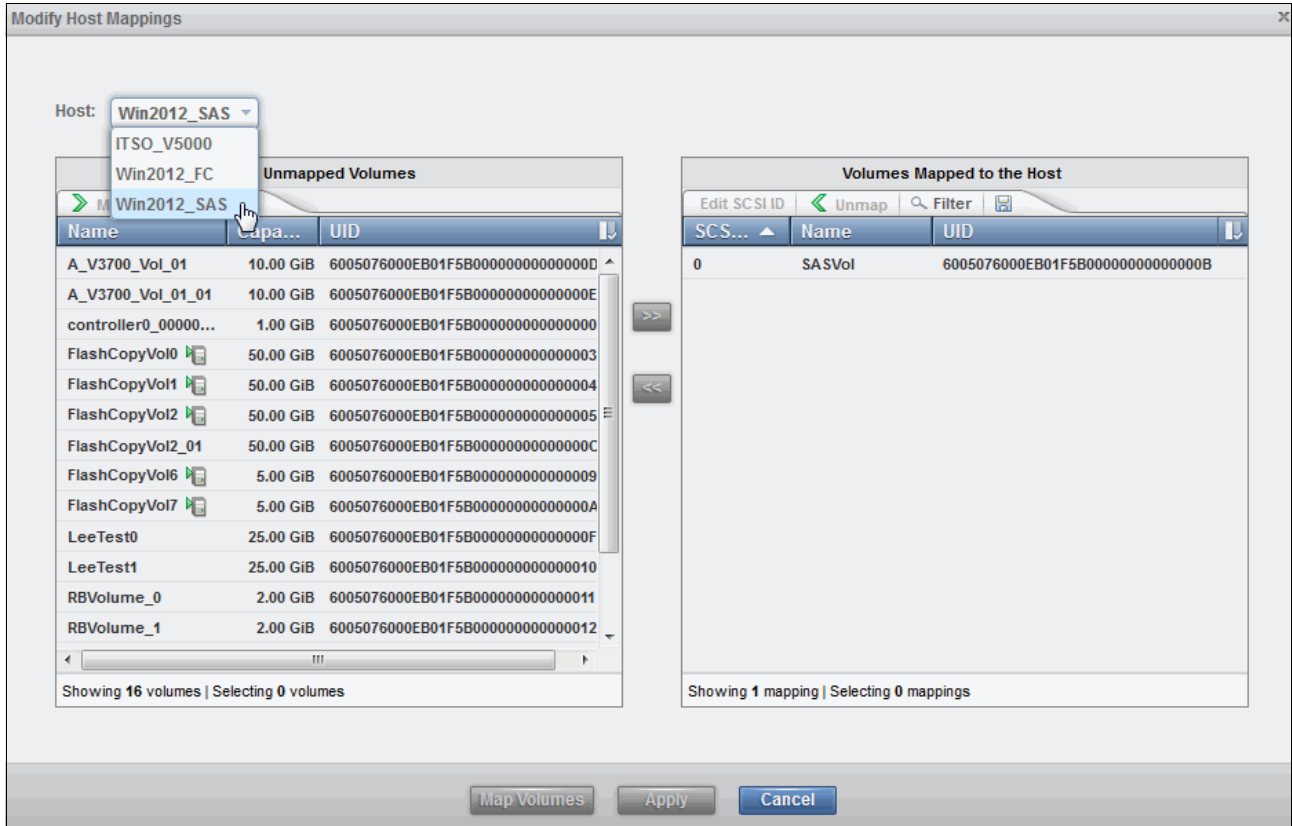


Figure 8-8 Selecting another host to modify

- After moving a volume to the right pane in the Modify Host Mappings pane, it is possible to right-click the yellow unmapped volume and change the SCSI ID used for the host mapping, as shown in Figure 8-9.

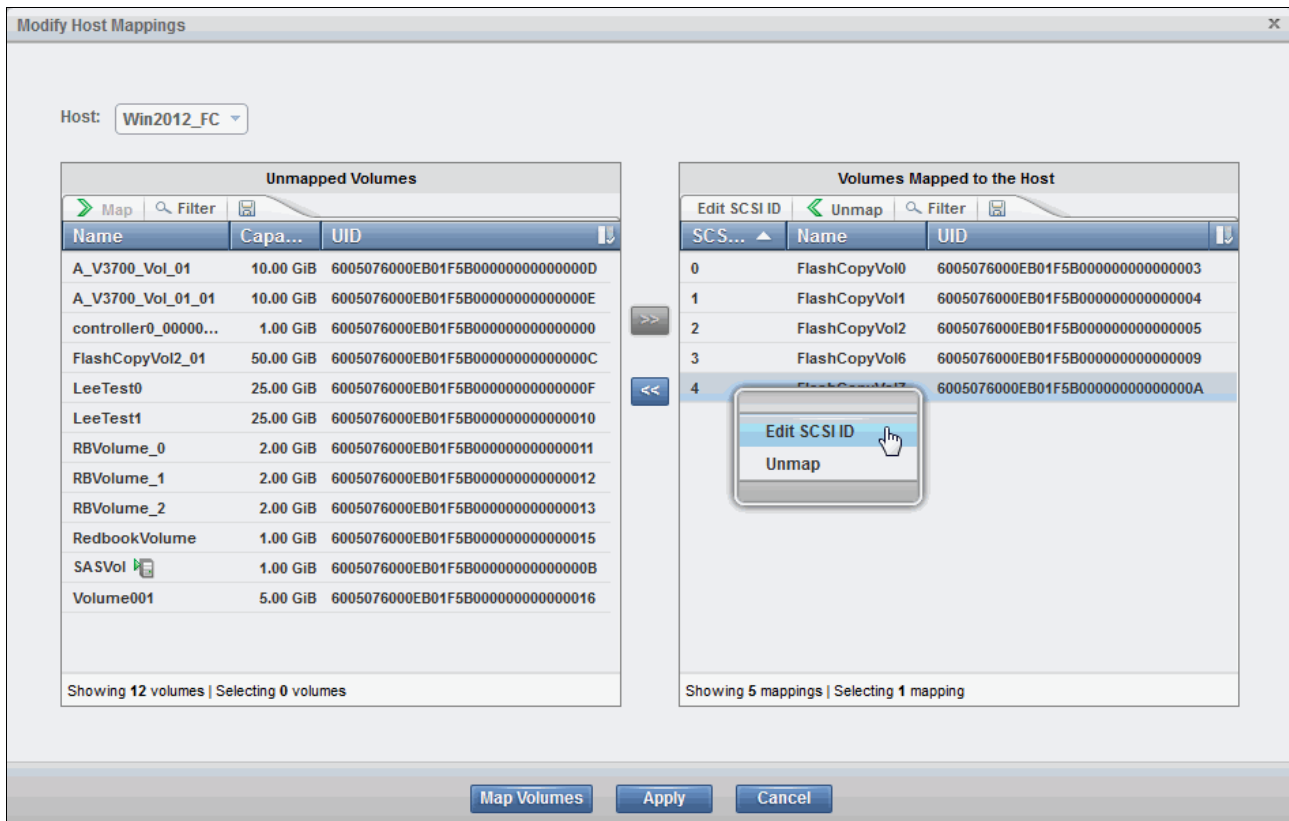


Figure 8-9 Editing SCSI ID

- Click **Edit SCSI ID**, select the required SCSI ID, then click **OK** to save the change.
- Click **Apply** to submit the changes and complete the host volume mapping.

Important: The IBM Storwize V3700 automatically assigns the lowest available SCSI ID if none is specified. However, you can set a SCSI ID for the volume. The SCSI ID cannot be changed while volume is assigned to host.

Attention: It is also possible to map a volume to multiple hosts. However, this would normally only be done in a clustered host environment, because data corruption is possible.

11. If you map a volume to multiple hosts, a warning is displayed, as shown in Figure 8-10.

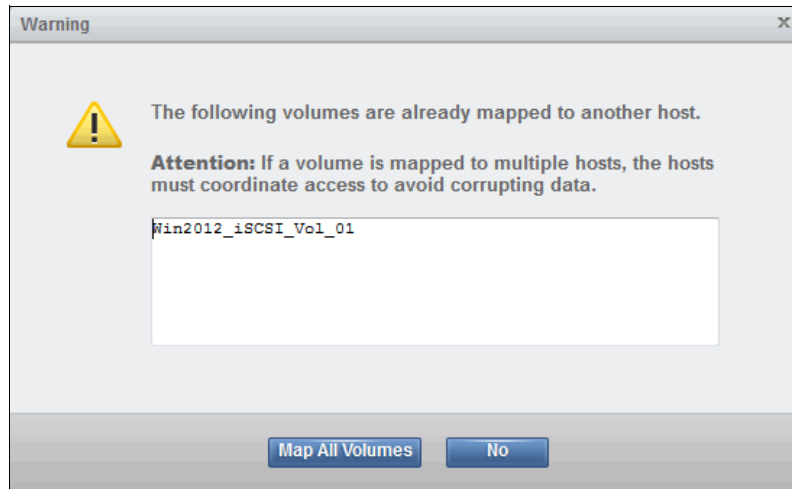


Figure 8-10 Volume mapped to multiple hosts warning

8.1.2 Unmapping volumes from a host

To unmap a volume from a host, complete the following steps:

1. Select the volume and click the Left Arrow icon to move the volume to the left (Unmapped Volumes) pane. Multiple volumes can be unmapped by selecting them with the Ctrl or Shift key, as shown in Figure 8-11.

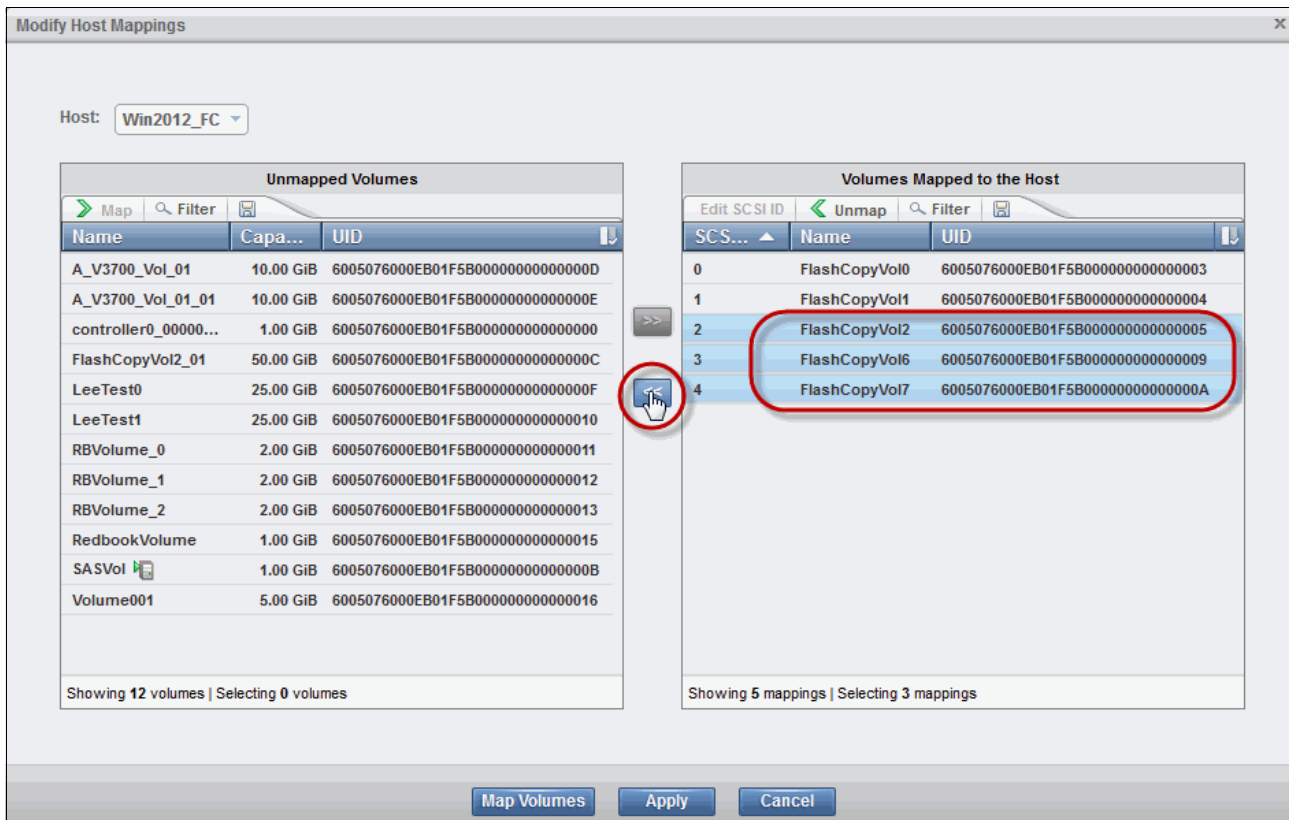


Figure 8-11 Unmapping volumes

- To unmap all volumes from a host, select a host from the Hosts pane and, using the Action menu, click **Unmap all Volumes**, as shown in Figure 8-12.

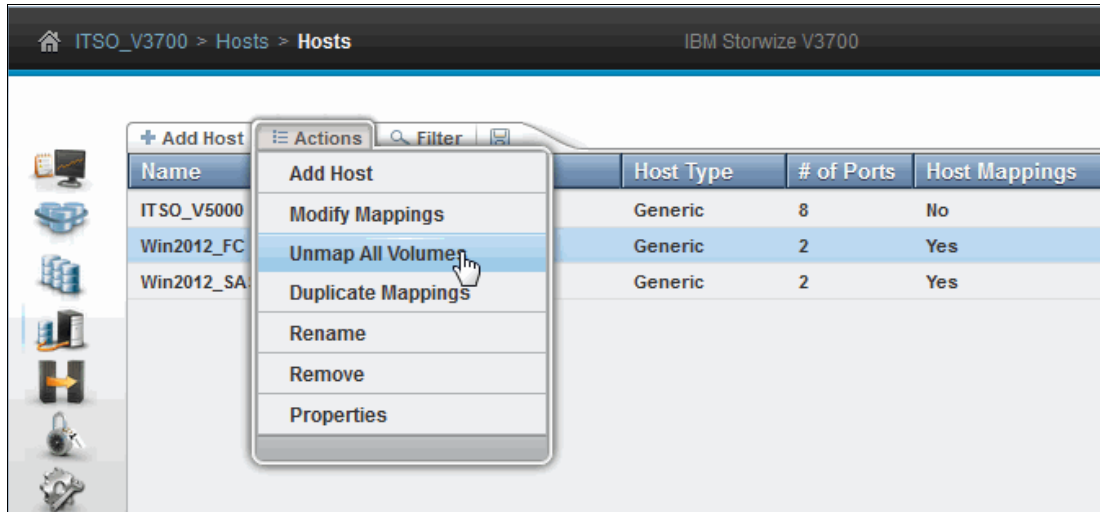


Figure 8-12 Unmap all volumes

- A prompt is displayed requesting confirmation of the number of mappings to be removed. Enter the number of mappings to be removed and click **Unmap**. In the example, five mappings are removed, as shown in Figure 8-13.

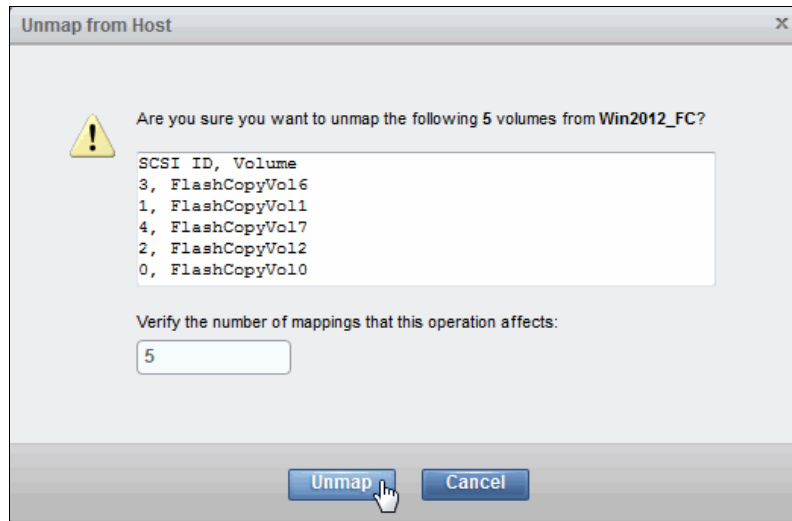


Figure 8-13 Enter the number of mappings to be removed

Unmapping: When you click **Unmap**, access to all of the volumes for this host is removed. Ensure that the required procedures are run on the host operating system (OS) before the unmapping procedure is performed.

- After you click Unmap, the changes are applied to the system, as shown in Figure 8-14. Click **Close** after you review the output.

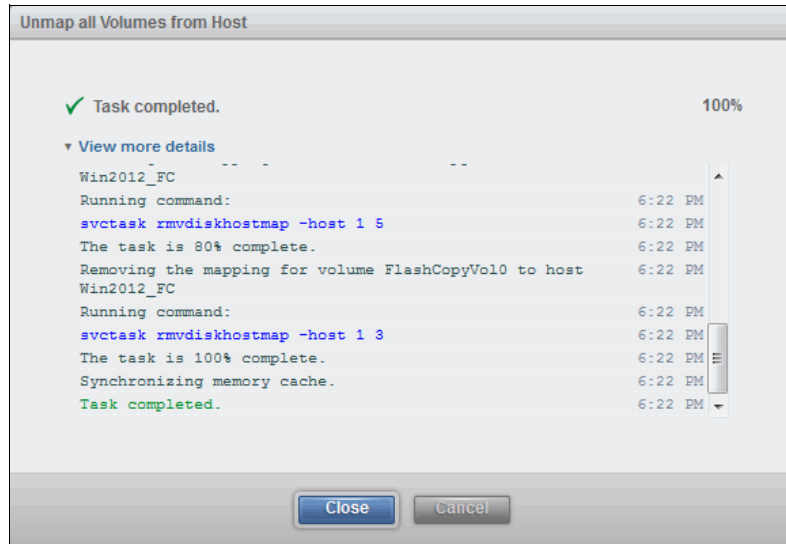


Figure 8-14 Unmapping all volumes from host

Figure 8-15 shows that the selected host no longer has any volume mappings.

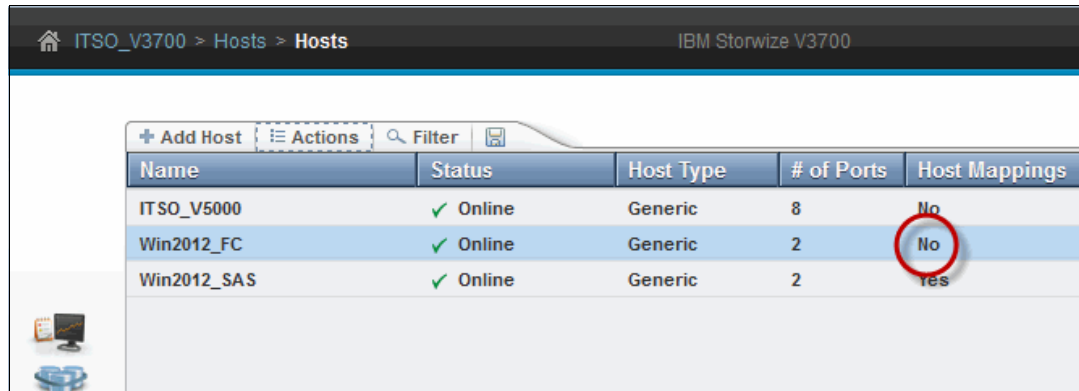


Figure 8-15 Host mapping

8.1.3 Renaming a host

To rename a host, complete the following steps:

1. Select the host and, from the Actions menu, click **Rename**, as shown in Figure 8-16.

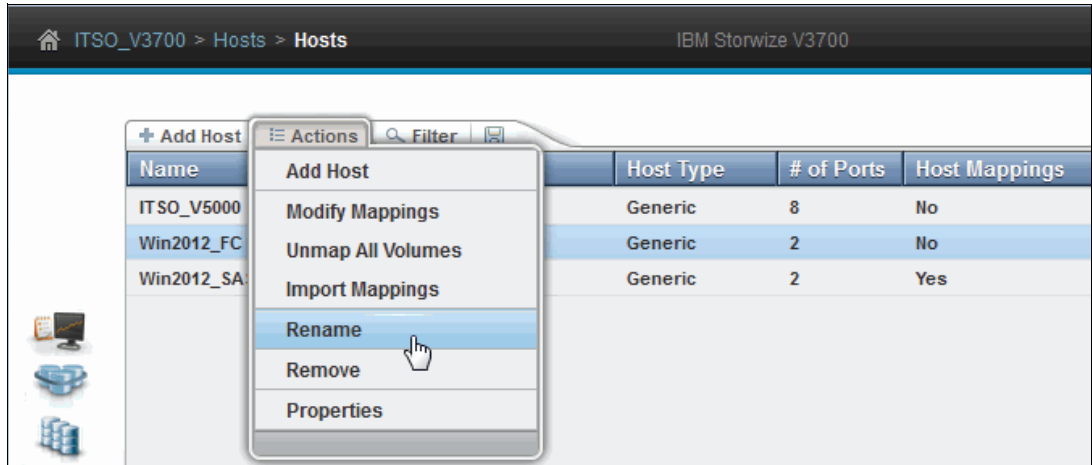


Figure 8-16 Renaming a host

2. Enter a new name and click **Rename**, as shown in Figure 8-17. Clicking **Reset** changes the host name back to its original name, where clicking **Cancel** cancels the name change and closes the Rename Host pane.



Figure 8-17 Rename Host pane

3. After the changes are applied to the system, click **Close**, as shown in Figure 8-18.

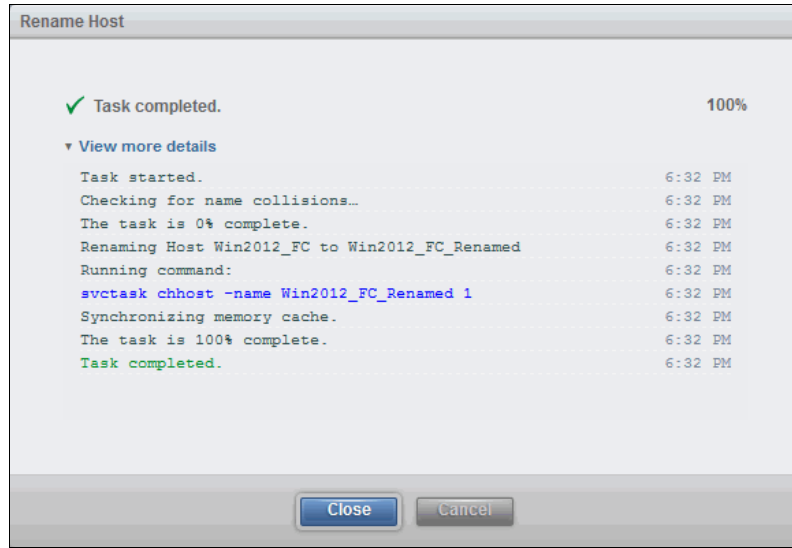


Figure 8-18 Rename a host task completed

8.1.4 Removing a host

To remove a host, complete the following steps:

1. Select the host and from the Actions menu, click **Remove**, as shown in Figure 8-19.

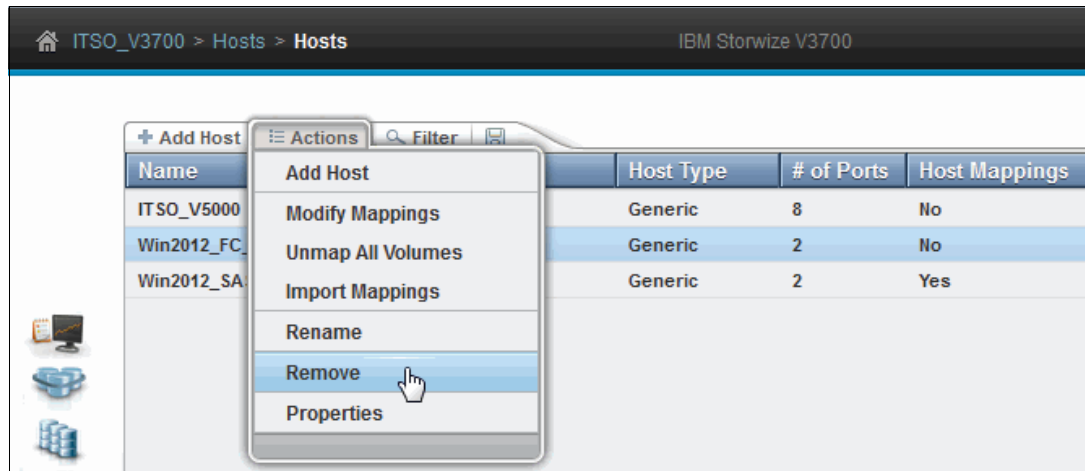


Figure 8-19 Removing a host

2. A prompt is displayed requesting confirmation of the number of hosts to be removed. Enter the number of hosts to be removed and click **Remove**, as shown in Figure 8-20.

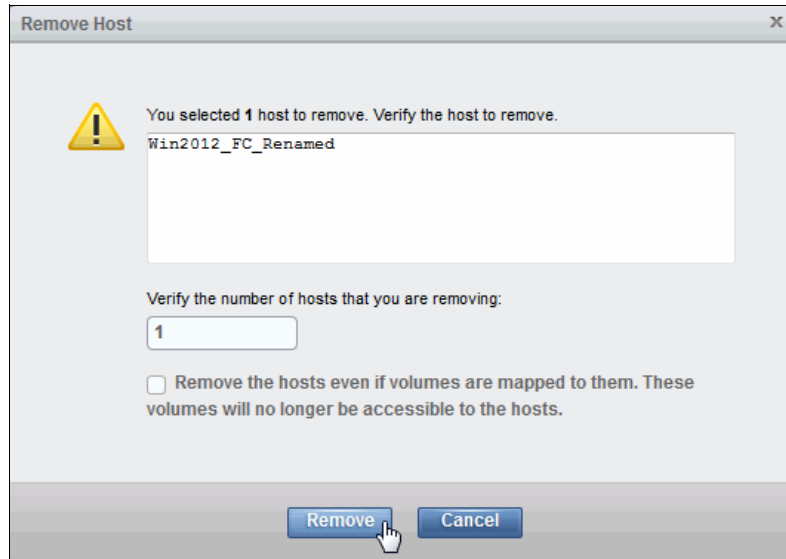


Figure 8-20 Removing a host

If it is necessary to remove a host with volumes assigned, deletion must be forced by selecting the option to **Remove the hosts even if volumes are mapped to them** (see Figure 8-20). This removes both the host and its associated volume mappings.

3. After the task is complete, click **Close** to return to the Hosts pane, as shown in Figure 8-21.

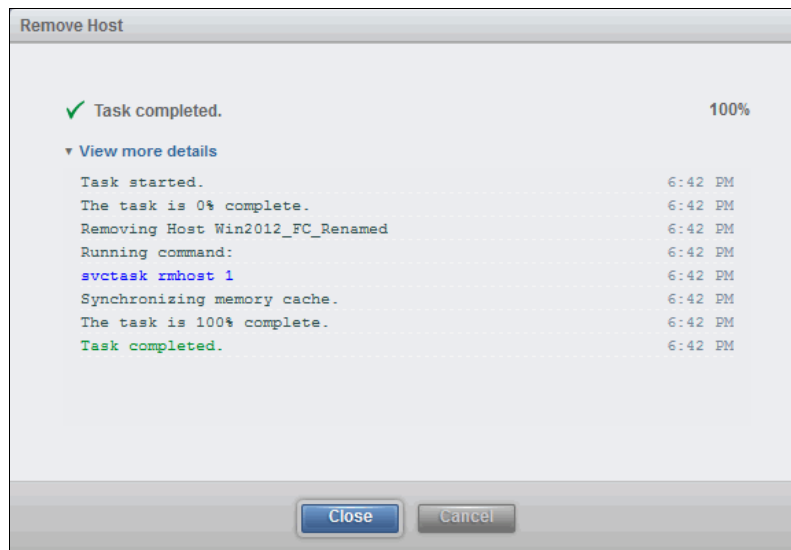


Figure 8-21 Remove host task completed

8.1.5 Host properties

This section describes the host properties, which provide the following information:

- ▶ Overview
- ▶ Mapped Volumes
- ▶ Port Definitions

To open the Host Details pane, complete the following steps:

1. Select the host and, from the Actions menu, click **Properties**. You also can highlight the host and right-click it, as shown in Figure 8-22.

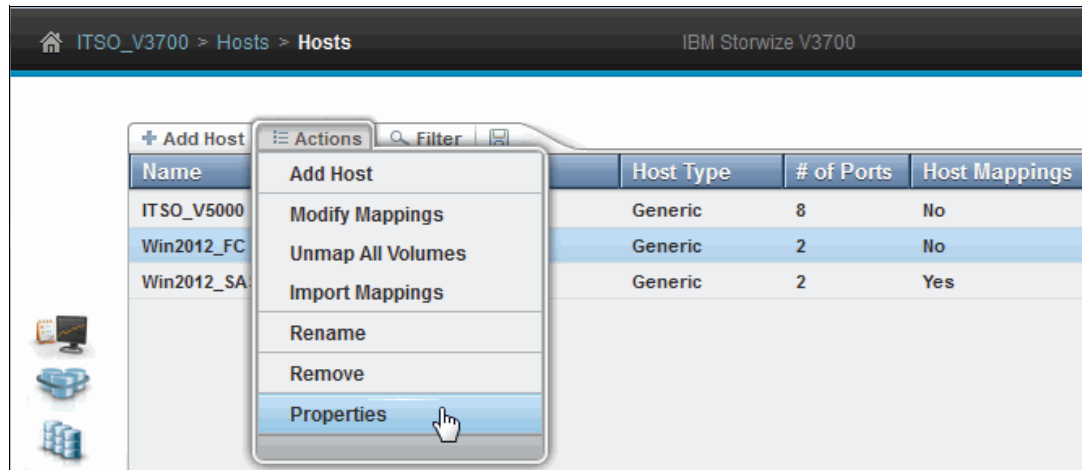


Figure 8-22 Opening host properties

2. Figure 8-23 shows the Overview tab of the Host Details pane. Select the **Show Details** box in the lower left to see more information about the host.

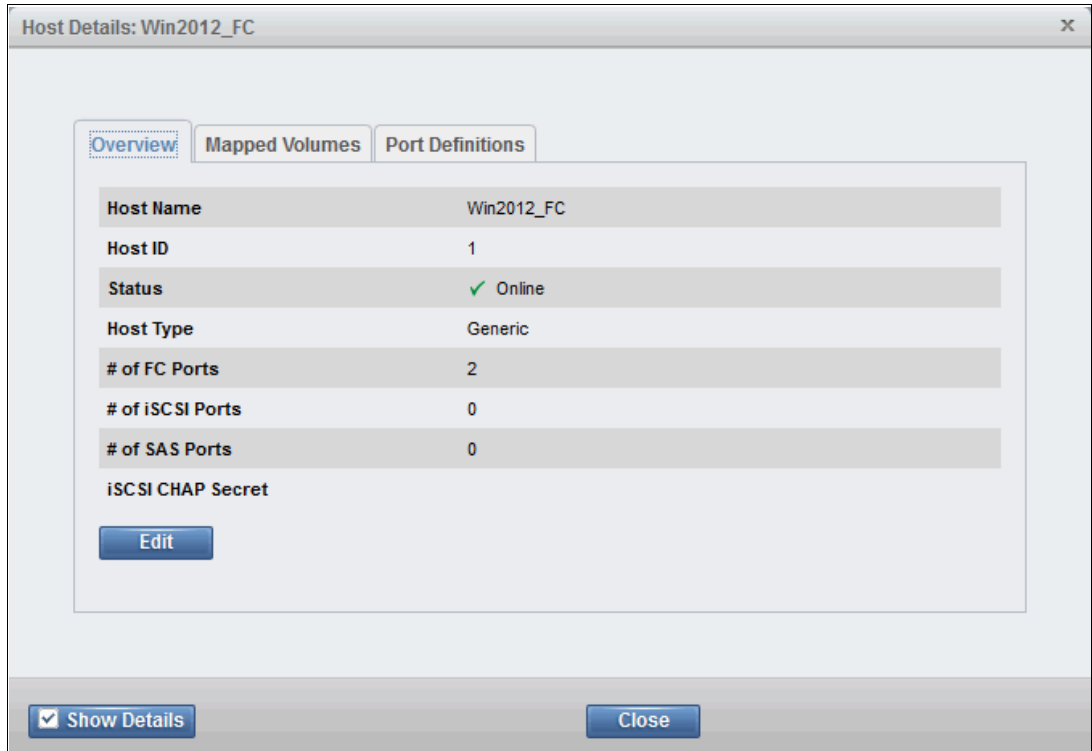


Figure 8-23 Host details

This tab provides the following information:

Host Name	Host object name.
Host ID	Host object identification (ID) number.
Status	The current host object status (Online, Offline, or Degraded).
# of FC Ports	The number of host Fibre Channel (FC) ports.
# of iSCSI Ports	The number of host Internet SCSI (iSCSI) names or host iSCSI qualified name (IQN) IDs.
# of SAS Ports	The number of host serial-attached SCSI (SAS) ports.
iSCSI CHAP Secret	The Challenge Handshake Authentication Protocol (CHAP) information if it exists or is configured.

3. To change the host properties, click **Edit** and several fields can be edited, as shown in Figure 8-24.

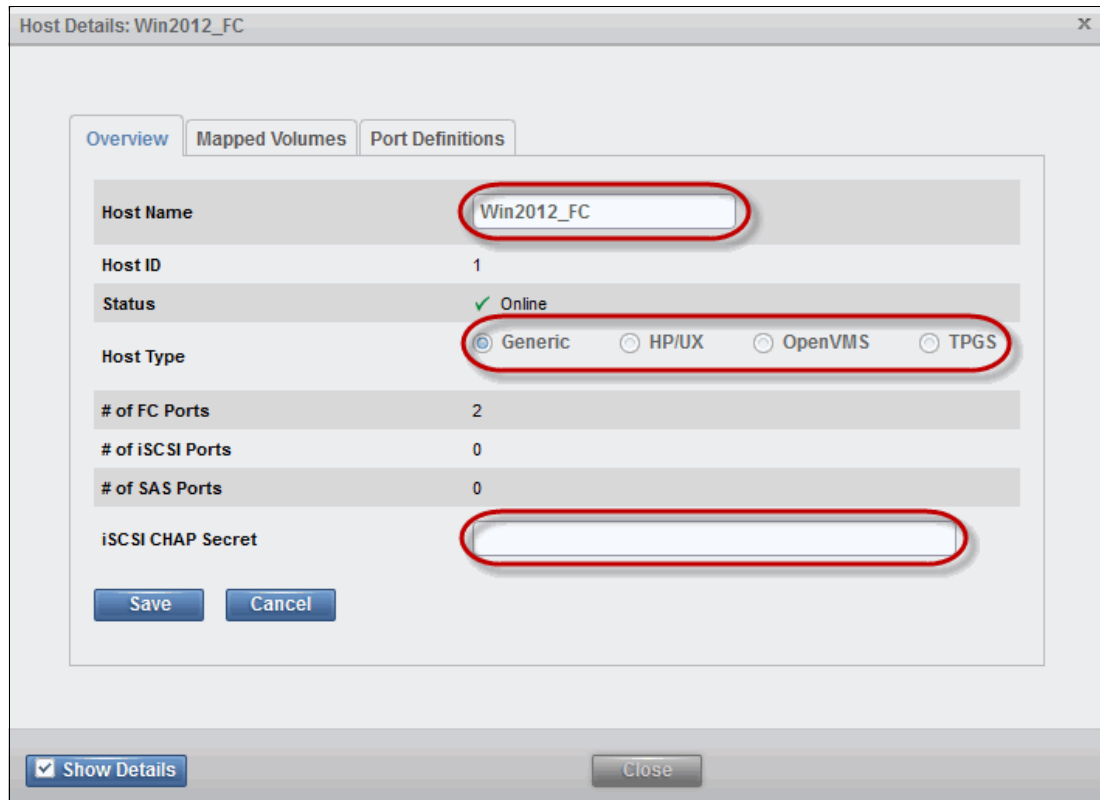


Figure 8-24 Host properties: Editing host information

4. Make any necessary changes and click **Save** to apply them, and then click **Close** to return to the Host Details pane.

Figure 8-25 on page 376 shows the **Mapped Volumes** tab, which provides an overview of the volumes mapped to the host. This tab provides the following information:

- SCSI ID
- Volume name
- UID
- Caching input output (I/O) Group ID

5. Clicking the **Show Details** option does not show any detailed information, as shown in Figure 8-25.

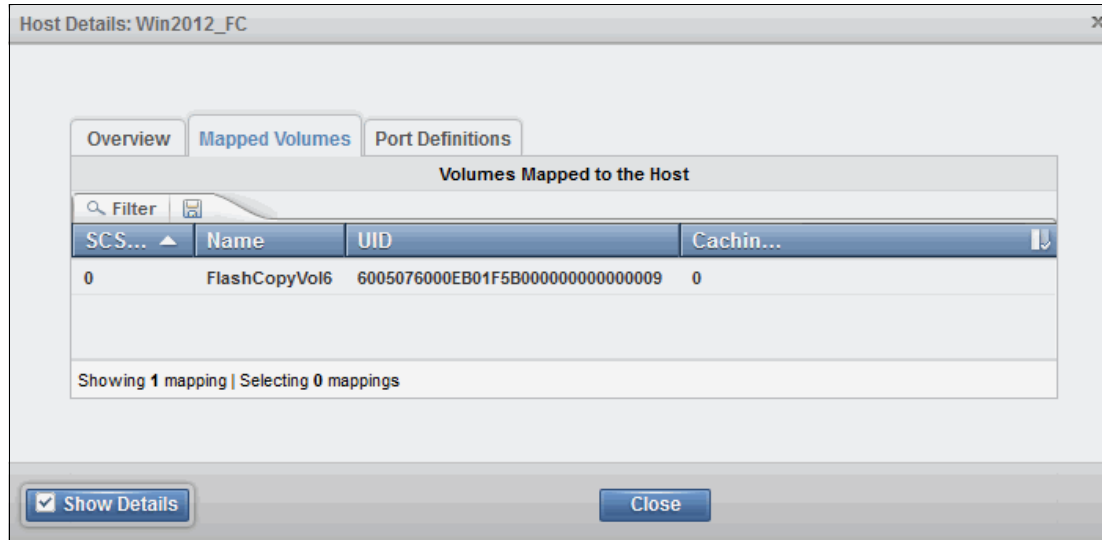


Figure 8-25 Host Details: Mapped Volumes

6. Figure 8-26 shows the Port Definitions tab, which shows the configured host ports and their status. This tab provides the following information:
 - The worldwide port names (WWPNs) for SAS and FC hosts.
 - IQN for iSCSI hosts.
 - The Type column shows the port type information.
 - The Status column states whether a port is Active or Inactive.
 - The # Nodes Logged In column lists the number of IBM Storwize V3700 node canisters that each port (initiator port) is logged in to.

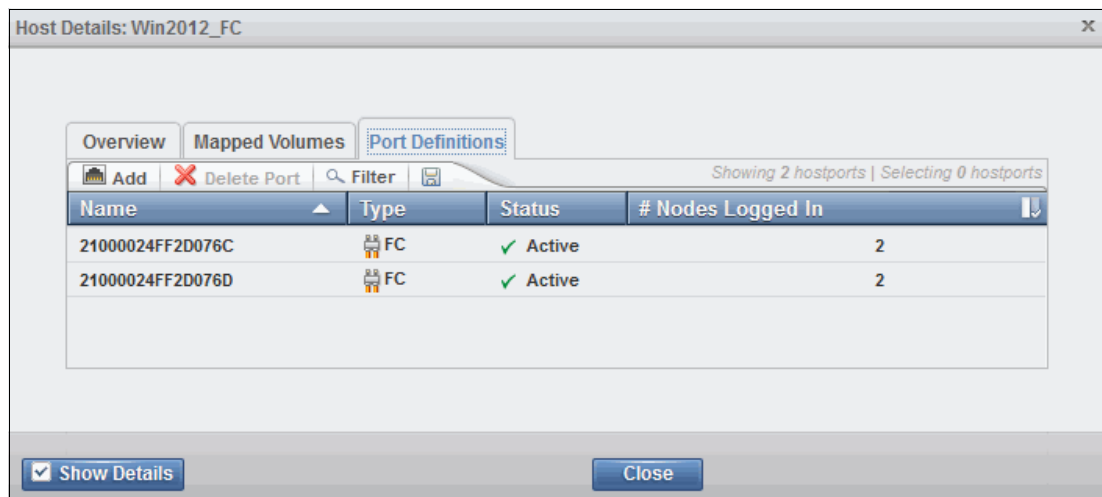


Figure 8-26 Host Details: Port Definitions

7. Using this pane, you can also Add and Delete Host Ports, as described in 8.2, “Adding and deleting host ports” on page 377. Selecting the **Show Details** option does not show any further information.
8. Click **Close** to close the Host Details section.

8.2 Adding and deleting host ports

To configure host ports, complete the following steps:

1. Browse to **Hosts** → **Ports by Host** to open the Ports by Host pane, as shown in Figure 8-27.




Figure 8-27 Ports by Host pane

Hosts are listed in the left pane. The Host icons show an orange cable for a FC host, a black cable for an SAS host, and a blue cable for an iSCSI host.

2. The properties of the selected host are shown in the right pane. Click **Add Host** to start the wizard described in Chapter 4, “Host configuration” on page 149.

Figure 8-28, shows the Action menu and associated tasks for a host.

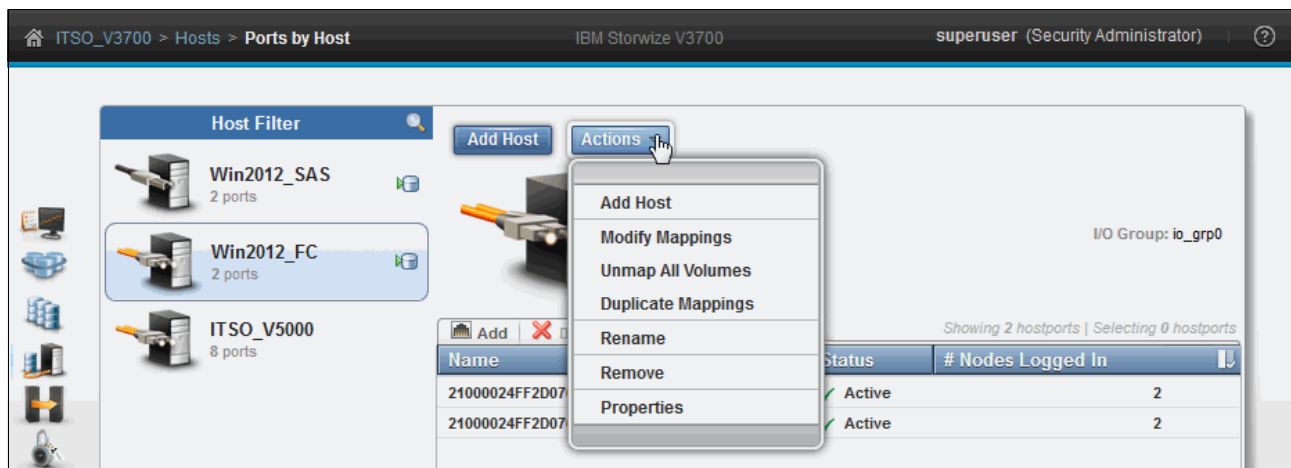


Figure 8-28 Ports by Host: Action menu

8.2.1 Adding a host port

To add a host port, complete the following steps:

1. Select the host from the left pane and click **Add**.
2. Choose a Fibre Channel, SAS, or iSCSI port, as shown in Figure 8-29.

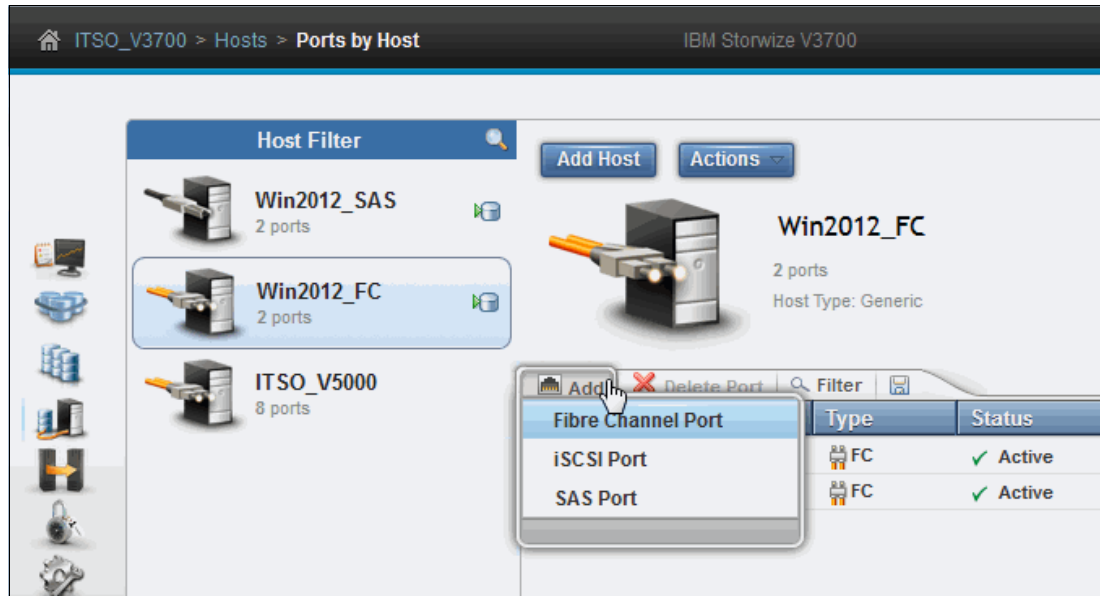


Figure 8-29 Adding a host port

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

8.2.2 Adding a Fibre Channel port

To add an FC port, complete the following steps:

1. After clicking **Add**, as shown in Figure 8-29, click **Fibre Channel Port** to open the Add Fibre Channel Ports pane.
2. Click the **Fibre Channel Ports** menu to display a list of all available FC host ports.
If the FC WWPN of your host is not available in the menu, check the SAN zoning and rescan the SAN from the host or try clicking **Rescan**.

3. Select the required WWPN and click **Add Port to List**. This adds the port to the Port Definitions list, as shown in Figure 8-30. Repeat this step to add more ports to the host.

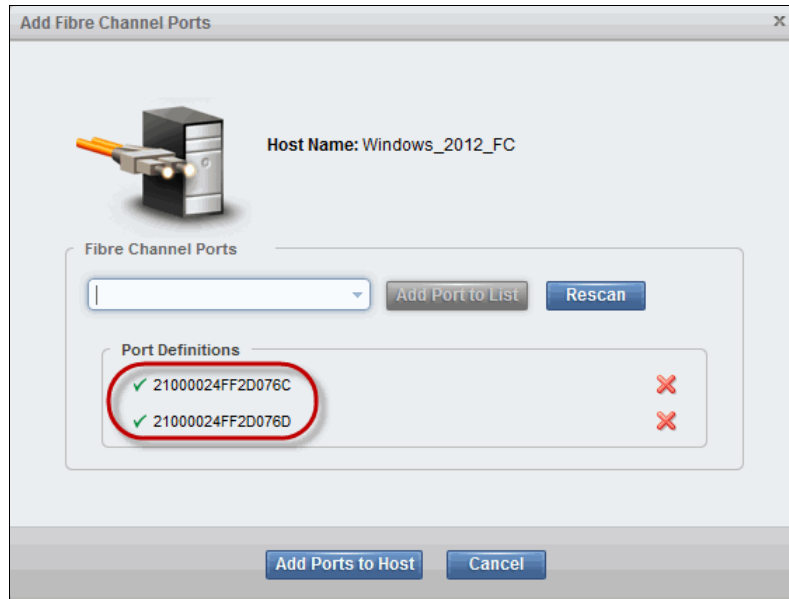


Figure 8-30 Adding online host ports

4. To add an offline port, manually enter the WWPN of the port into the Fibre Channel Ports field and click **Add Port to List**, as shown in Figure 8-31.

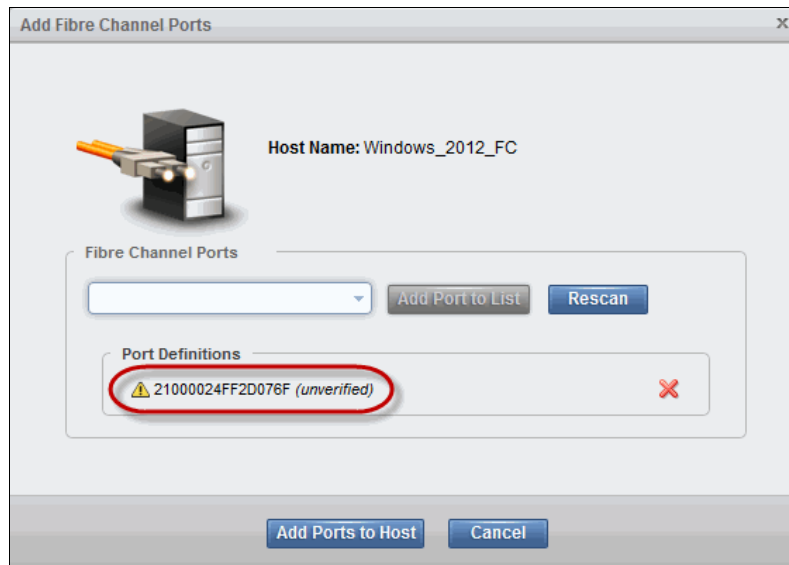


Figure 8-31 Adding offline port

The port shows as unverified because it is not logged in to the SAN fabric or known to the IBM Storwize V3700. The first time the port logs in, the state automatically changes to online, and the mapping is applied to this port.

5. To remove one of the ports from the list, click the red X next to it. In Figure 8-31, we manually added an FC port.

Important: When removing online or offline ports, the IBM Storwize V3700 prompts you to confirm the number of ports to be deleted, but does not warn about mappings. Disk mapping is associated to the host object, and logical unit number (LUN) access is lost if all ports are deleted.

6. Click **Add Ports to Host** to apply the changes.

8.2.3 Adding an SAS host port

To add an SAS host port, complete the following steps:

1. After clicking **Add**, as shown in Figure 8-29 on page 378, click **SAS Port** to open the Add SAS Ports pane.
2. Click the **SAS Ports** menu to see a list of all available SAS ports. If the SAS WWPNs are not available, try the **Rescan** option or check the physical connections.

Important: The IBM Storwize V3700 enables you to add an offline SAS port. To do this, enter the SAS WWPN in the SAS Port field, and then click **Add Port to List**.

3. Select the required SAS WWPN and click **Add Port to List**, as shown in Figure 8-32. Repeat this step to add more ports to a host.

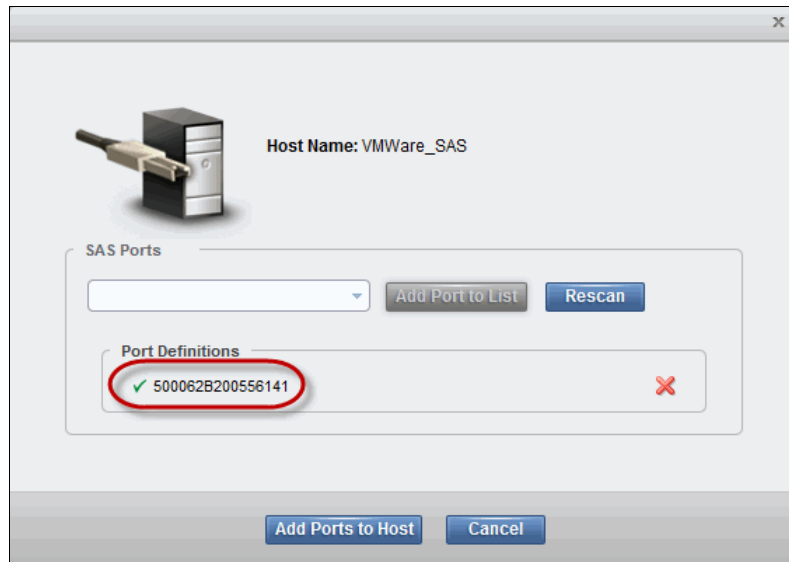


Figure 8-32 Adding an online SAS port

4. Click **Add Ports to Host** to apply the changes.

8.2.4 Adding an iSCSI host port

To add an iSCSI port, complete the following steps:

1. After clicking **Add**, as shown in Figure 8-29 on page 378, select **iSCSI Port** to open the Add iSCSI Ports pane shown in Figure 8-33.

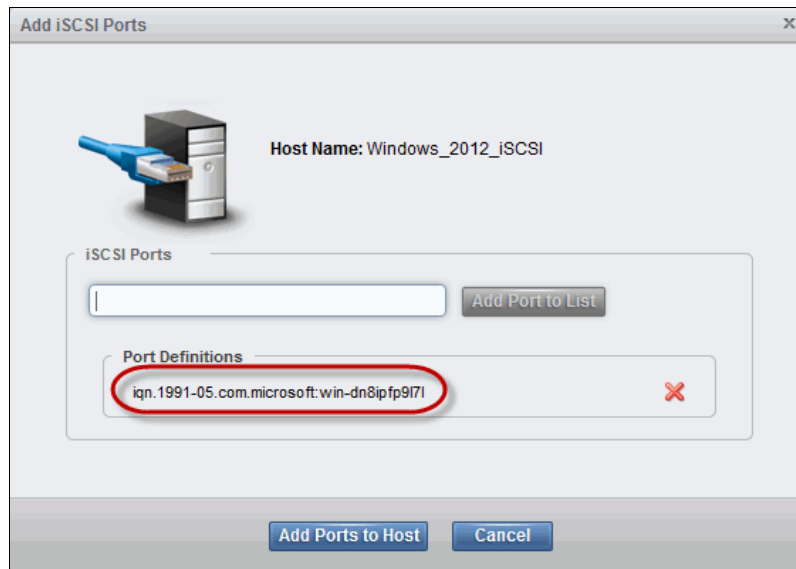


Figure 8-33 Adding iSCSI Host ports

2. Enter the initiator name of your host and click **Add Port to List**, as shown in Figure 8-33.
3. Click **Add Ports to Host** to apply the changes. The iSCSI port status remains unknown until it is added to the host and a host rescan process is completed.
4. Click **Close** to return to the Ports by Host pane.

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

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

8.2.5 Deleting a host port

To delete host ports, complete the following steps:

1. Click **Host** → **Ports by Host** to open the Ports by Host pane.
2. Select the host in the left pane and select the host port to be deleted, then click **Delete Port**, as shown in Figure 8-34.



Figure 8-34 Delete host port

Multiple ports can be selected by holding down the Ctrl or Shift keys while selecting the host ports to delete.

3. Figure 8-35 shows the Delete Ports confirmation pane. Confirm the number of ports to be deleted and click **Delete** to apply the changes.

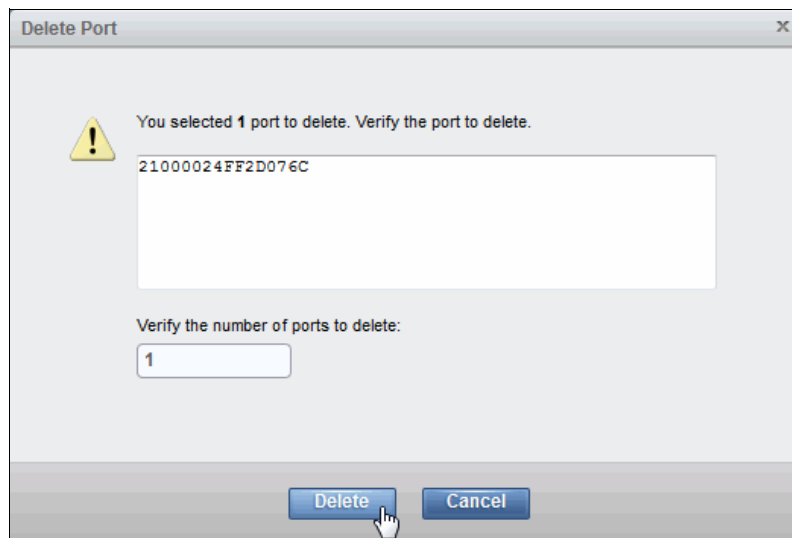


Figure 8-35 Delete host port

4. A task pane opens that shows the results. Click **Close** to return to the Ports by Host pane.

8.3 Host mappings overview

To access host mappings, complete the following steps:

1. Select **Host** → **Host Mappings** to open the Host Mappings pane, as shown in Figure 8-36.

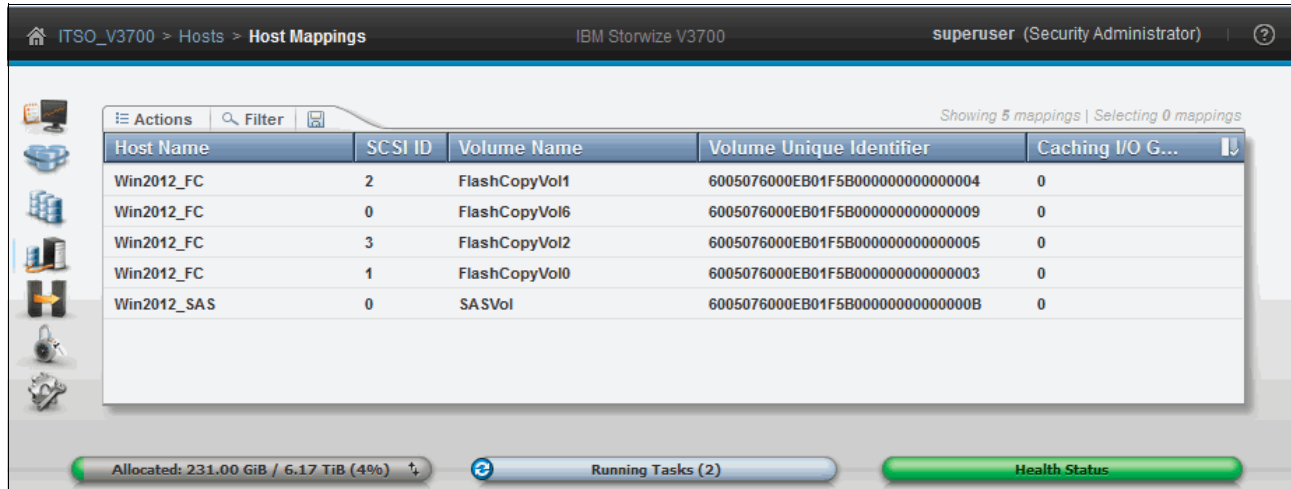


Figure 8-36 Host mappings

The Host Mapping pane shows a list of mapped hosts and volumes, along with their respective SCSI ID and Volume Unique Identifier (UID). Figure 8-36 shows the host Windows2012_FC has four mapped volumes (FlashCopyVol0, 1, 2 and 6), and the associated SCSI ID (1, 2, 3 and 0), Volume Name, UID, and Caching I/O Group ID.

2. As shown in Figure 8-37, select a mapping and open the **Actions** menu to open the following options:
 - Unmap Volumes
 - Properties (Host)
 - Properties (Volume)

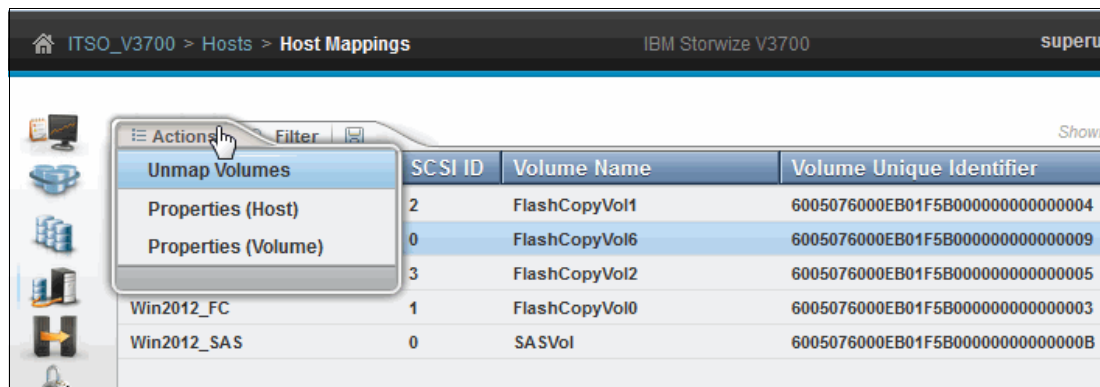


Figure 8-37 Host mapping actions

If multiple mappings are selected (by holding the Ctrl or Shift keys), only the **Unmap Volumes** option is available.

8.3.1 Unmap Volumes

To unmap volumes, complete the following steps:

1. Select one or more mappings.
2. From the Actions menu, Select **Unmap Volumes**.
3. Enter the number of mappings to be removed (as shown in Figure 8-38), and click **Unmap**. The mappings for all selected entries are removed.

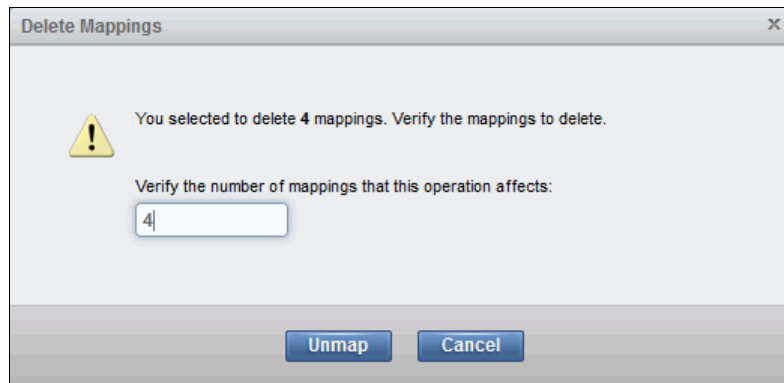


Figure 8-38 Unmapping a volume from host

Warning: Ensure that the required procedures are run on the host OS before unmapping volumes on the IBM Storwize V3700.

8.3.2 Properties (Host)

Selecting a mapping and clicking **Properties (Host)** from the Actions menu (as shown in Figure 8-37 on page 383) opens the Host Properties pane. For more information, see 8.1.5, “Host properties” on page 373.

8.3.3 Properties (Volume)

Selecting a mapping and clicking **Properties (Volume)** from the Actions menu (as shown in Figure 8-37 on page 383) opens the Volume Properties pane. For more information about volume properties, see 8.5, “Volume properties” on page 396.

8.4 Advanced volume administration

This section describes volume administration tasks, such as volume modification, volume migration, and creation of volume copies. We assume that you have completed Chapter 5, “Volume configuration” on page 199, and so are familiar with volume creation and generic, thin-provisioned, mirrored, and thin-mirrored volumes.

Figure 8-39 on page 385 shows the following options that are available in the Volumes menu for advanced features administration:

- ▶ Volumes
- ▶ Volumes by Pool
- ▶ Volumes by Host




Figure 8-39 Volumes menu

8.4.1 Advanced volume functions

To perform advanced volume functions, complete the following steps:

1. Click **Volumes** (as shown in Figure 8-39) and the Volumes pane opens, as shown in Figure 8-40.

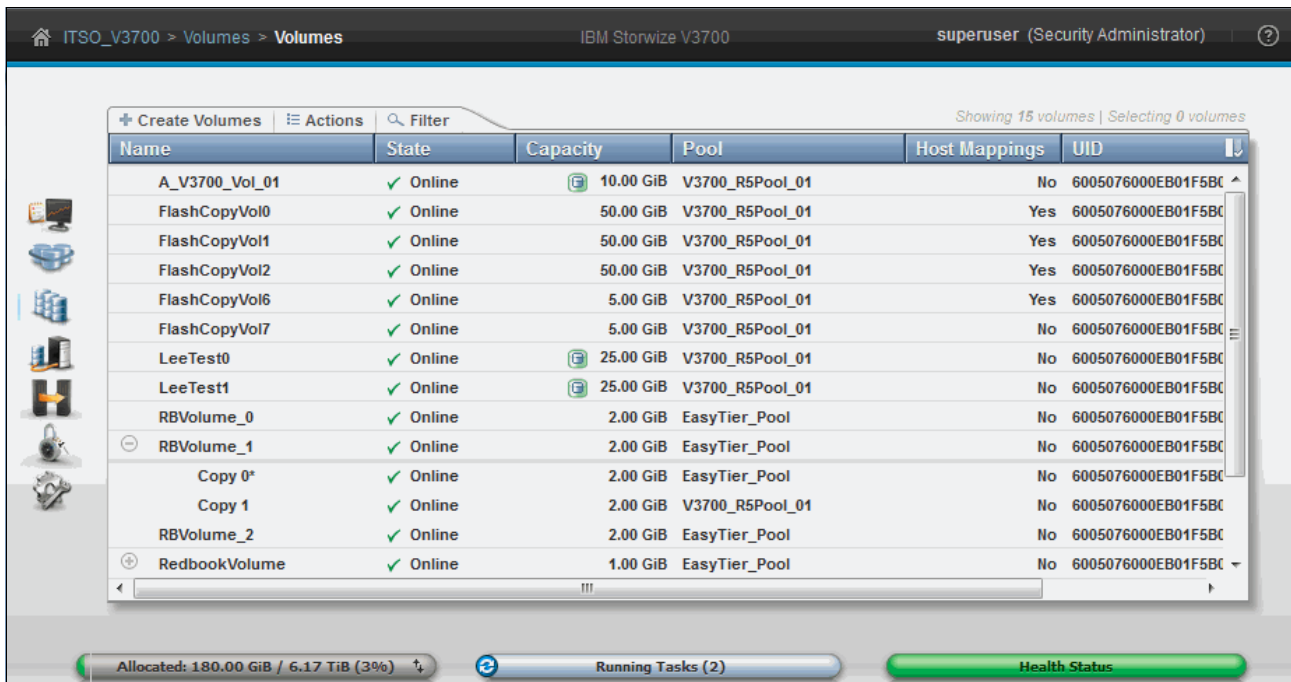


Figure 8-40 Volume pane

This pane lists all configured volumes on the system and provides the following information:

- Name** Shows the name of the volume. If there is a plus sign (+) next to the name, this sign means that there are two copies of this volume. Click the plus sign (+) to expand the view and list the copies, as shown in Figure 8-40 on page 385.
- Status** Provides the status information about the volume, which can be online, offline, or degraded.
- Capacity** The disk capacity that is presented to the host. If a blue volume is listed next to the capacity, this means that this volume is a thin-provisioned volume. Therefore, the listed capacity is the virtual capacity, which might be more than the real capacity on the system.
- Pool** Shows in which Storage Pool the volume is stored. The primary copy is shown unless you expand the volume copies.
- UID** The volume unique identifier (UID).
- Host Mappings** Shows if a volume has host mapping. The value is Yes when host mapping exists, and No when there are no hosting mappings.

Tip: Right-clicking anywhere in the blue title bar enables you to customize the volume attributes that are displayed. You might want to add some useful information, such as Caching I/O Group and Real Capacity.

2. To create a volume, click **Create Volumes** and complete the steps as described in 5.1, “Creating volumes in IBM Storwize V3700” on page 200.
3. Right-click a volume, or select a volume and open the **Actions** menu, to display the available actions for a volume, as shown in Figure 8-41.




Figure 8-41 Action menu for a volume

Depending on the volume selected, the following Volume options are available:

- Create Volumes
- Map to Host
- Unmap All Hosts
- View Mapped Host
- Duplicate Volume
- Rename
- Shrink
- Expand
- Migrate to Another Pool
- Delete
- Volume Copy Actions:
 - Create Volumes
 - Add Mirrored Copy
- Properties

For thin-provisioned volumes, the following Volume Copy Actions are available:

- Create Volumes
- Add Mirror Copy
- Thin Provisioned:
 - Shrink
 - Expand
 - Edit Properties

These options are described in the next sections.

8.4.2 Mapping a volume to a host

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

1. Select **Map to Host** from the Action menu.
2. Select the Host to which you want to map the volume. Figure 8-42 shows the Modify Host Mappings pane.

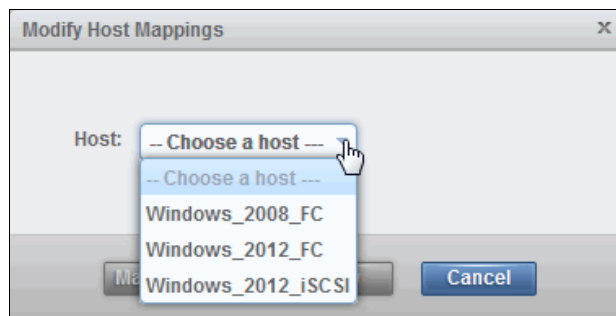


Figure 8-42 Modify Host Mappings pane

- After a host is selected, the Modify Host Mappings pane opens. In the upper left, the selected host is displayed. The volume that is to be mapped is highlighted in yellow, as shown in Figure 8-43. Click **Map Volumes** to apply the changes to the system.

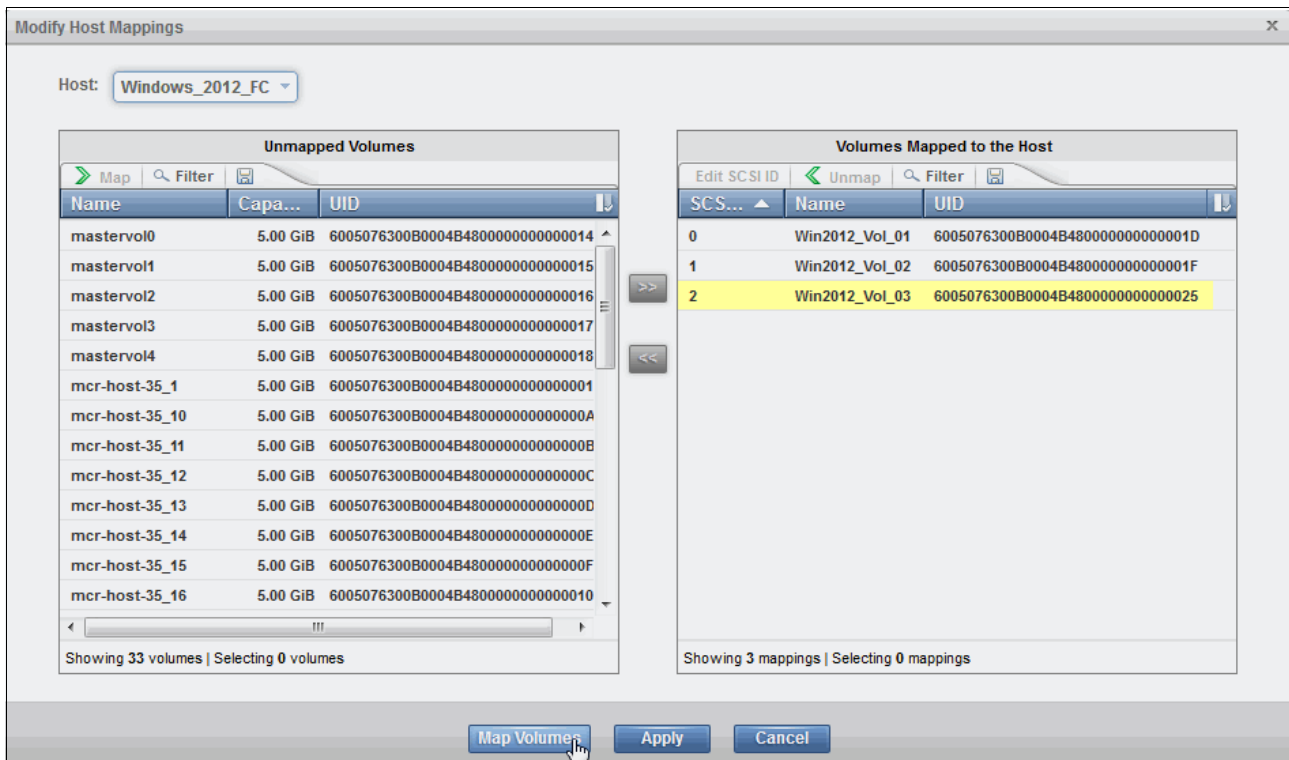


Figure 8-43 Modify Host Mappings

- After the changes are made, click **Close** to return to the Volumes pane.

Modify Host Mappings pane: For more information about the Modify Host Mappings pane, see 8.1.1, “Modifying Mappings menu” on page 362.

8.4.3 Unmapping volumes from all hosts

To remove all host mappings from a volume, complete the following steps:

1. Select **Unmap All Hosts** from the actions menu. This action removes all host mappings, which means that no hosts can access this volume.
2. Confirm the number of mappings that are to be removed and click **Unmap**, as shown in Figure 8-44.

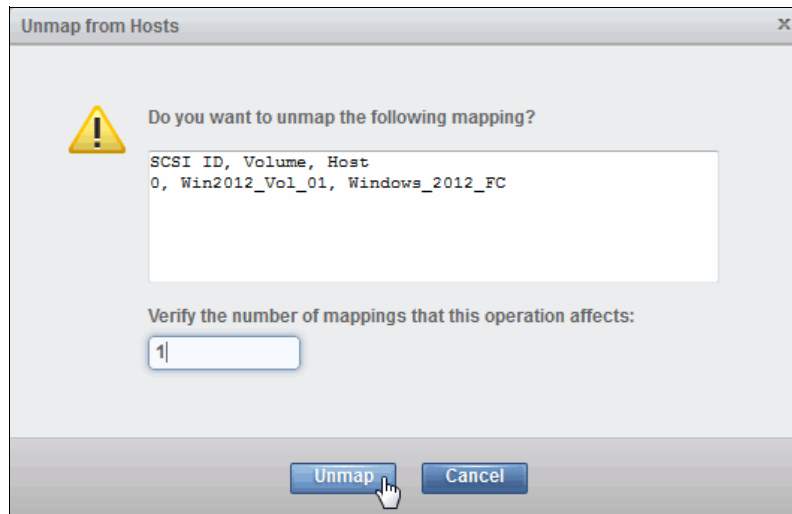


Figure 8-44 Unmapping from host (or hosts)

3. After the task completes, click **Close** to return to the Volumes pane.

Important: Ensure that the required procedures are run on the host OS before the unmapping procedure is performed.

8.4.4 Viewing which host is mapped to a volume

To determine which host mappings are configured, complete the following steps:

1. Highlight a volume and select **View Mapped Host** from the Actions menu. The Host Maps tab of the Volume Details pane opens, as shown in Figure 8-45. In this example, host Windows_2012_FC is mapped to the Win2012_Vol_01 volume.

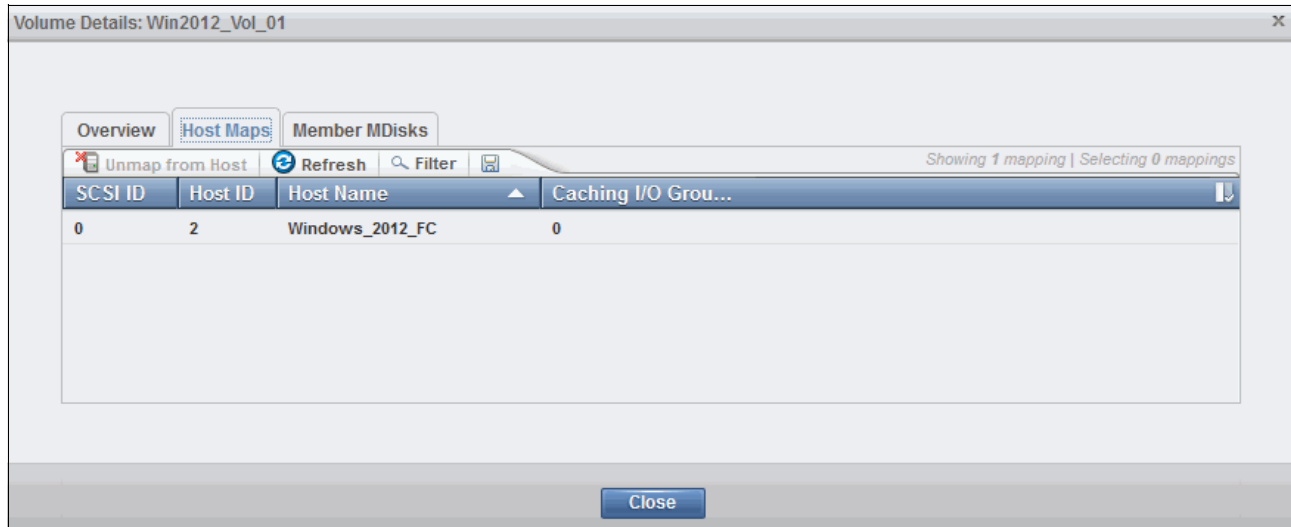


Figure 8-45 Volume to host mapping

2. To remove a mapping, highlight the host and click **Unmap from Host**. If several hosts are mapped to this volume (for example, in a cluster), only the selected host is removed.

8.4.5 Duplicating a volume

To duplicate a volume, complete the following steps:

1. After selecting a volume, click **Duplicate Volume** from the Actions menu to make an exact copy of an existing volume.
2. Enter a name for the duplicate volume and click **Duplicate** to apply the change, as shown in Figure 8-46.




Figure 8-46 Create a duplicate volume

8.4.6 Renaming a volume

To rename a volume, complete the following steps:

1. Select **Rename** from the Action menu. The Rename Volume window opens.
2. Enter the new name, as shown in Figure 8-47.



Figure 8-47 Renaming a volume

3. Click **Reset** to reset the name field to the original name of the volume.
4. Click **Rename** to apply the changes.
5. Click **Close** after the task pane completes.

8.4.7 Shrinking a volume

The IBM Storwize V3700 can shrink volumes. This feature should be used *only if the host OS supports it*. This capability reduces the capacity that is allocated to the particular volume by the amount specified. To shrink a volume, complete the following steps:

1. Select **Shrink** from the Actions menu.
2. Enter the new volume size or by how much the volume should shrink, as shown in Figure 8-48.

Attention: Before you shrink a volume, ensure that the host OS supports this. If the OS does not support shrinking a volume, it is likely that the OS will log disk errors, and data corruption will occur.

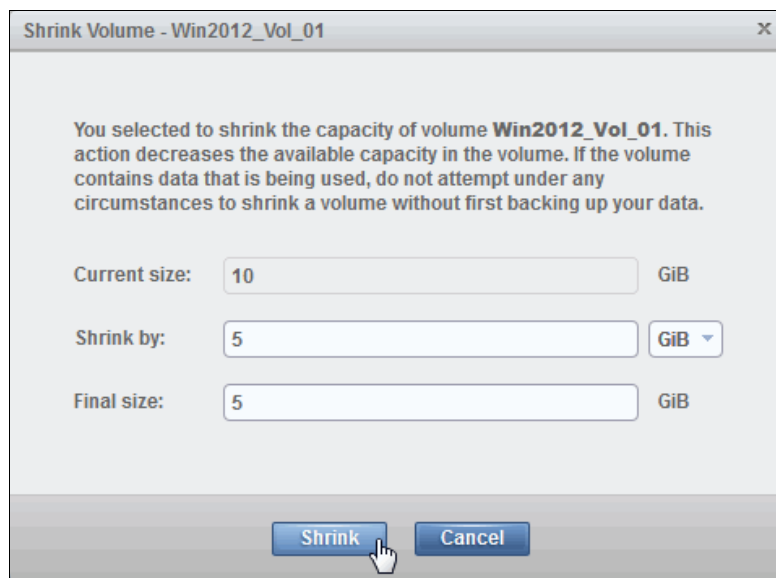


Figure 8-48 Shrink Volume pane

3. Click **Shrink** to start the process.
4. Click **Close** when the task pane completes to return to the Volumes pane.
5. Run the required procedures on the host OS after the shrinking process.

Important: For volumes that contain more than one copy, you might receive a CMMVC6354E error. In this case, complete the following steps:

1. Run the `lsvdisksyncprogress` command from the IBM Storwize V3700 command-line interface (CLI) to view the synchronization status.
2. Wait for the copy to synchronize.
3. If you want the synchronization process to complete more quickly, increase the rate by running the `chvdisk` command.
4. When the copy is synchronized, resubmit the shrink process.

8.4.8 Expanding a volume

The IBM Storwize V3700 can expand volumes. This feature should be used *only if the host OS supports it*. This capability increases the capacity that is allocated to the particular volume by the amount specified. To expand a volume, complete the following steps:

1. Select **Expand** from the Actions menu.
2. Enter the new volume size, or by how much the volume should expand, and click **Expand**, as shown in Figure 8-49.

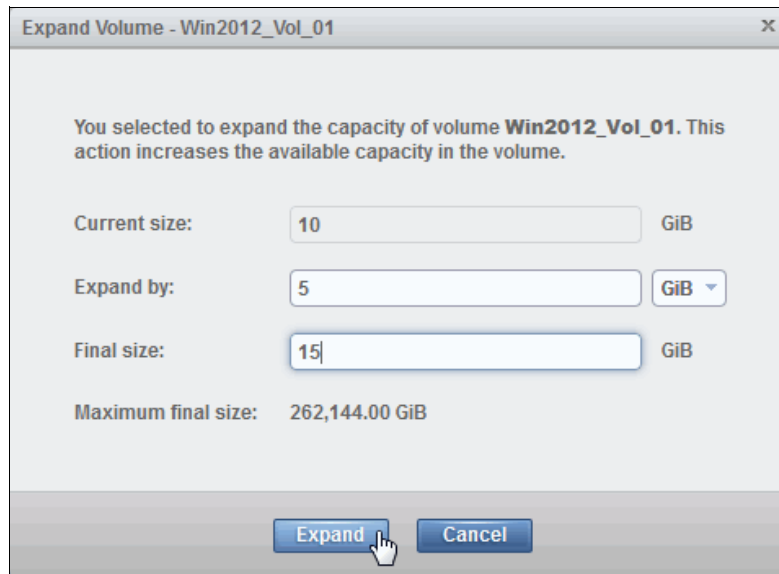


Figure 8-49 Expand Volume pane

3. After the task completes, click **Close** to return to the Volumes pane.
4. Run the required procedures on host OS to use the available space.

8.4.9 Migrating a volume to another Storage Pool

The IBM Storwize V3700 supports online volume migration while applications are running. Using volume migration, volumes can be moved between Storage Pools. The migration process is a low-priority task. One extent is moved at a time, and the migration has a slight effect on performance.

Important: To migrate a volume, the source and target Storage Pool must have the same extent size. For more information about extent size, see Chapter 1, "Overview of the IBM Storwize V3700 system" on page 1.

To migrate a volume to another Storage Pool, complete the following steps:

1. Select **Migrate to Another Pool** from the Actions menu. The Migrate Volume Copy pane opens.
2. If the volume consists of more than one copy, you are asked which copy you want to migrate to another Storage Pool, as shown in Figure 8-50. If the selected volume consists of only one copy, this option does not display.

Notice that the Win2012_Vol_01 volume has two copies stored in two different Storage Pools. The Storage Pools to which they belong are shown in parentheses.

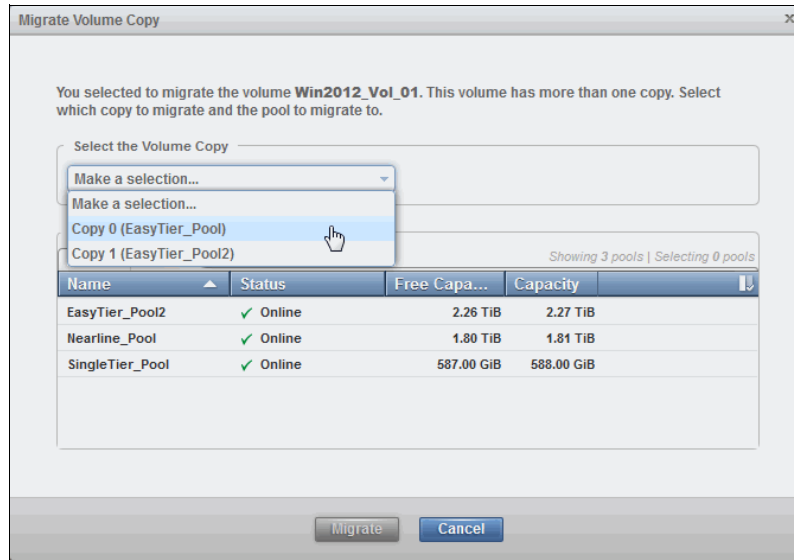


Figure 8-50 Migrate Volume Copy

3. Select the new target Storage Pool and click **Migrate**.
4. The volume migration starts, as shown in Figure 8-51. Click **Close** to return to the Volumes pane.

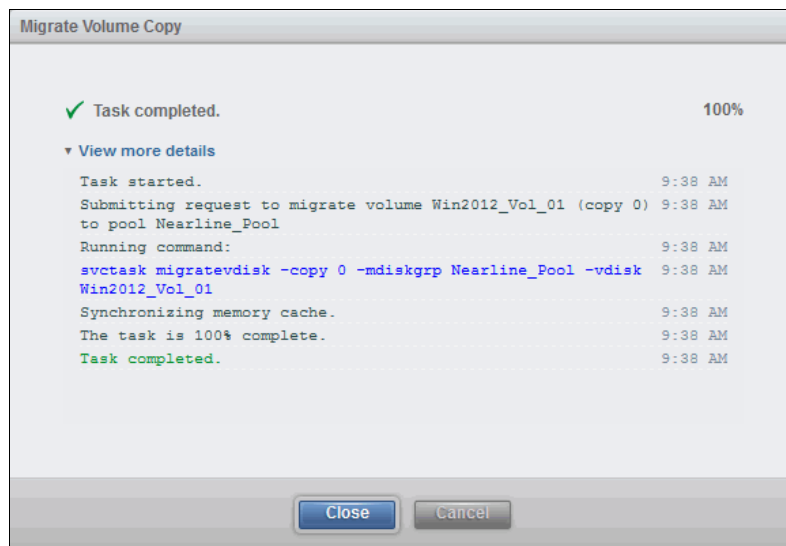


Figure 8-51 Migrate Volume Copy task progress

Depending on the size of the volume, the migration process can take some time. The status of the migration can be monitored in the **Running Tasks** bar at the bottom of the graphical user interface (GUI) pane. Volume migration tasks cannot be interrupted.

5. After the migration completes, the Copy 0* from the Win2012_Vol_01 volume is shown in the new Storage Pool, as shown in Figure 8-52.

Name	State	Capacity	Pool
controller0_000000000000...	✓ Online	1.00 GiB	MigrationPool_1024
Win2012_Vol_01	✓ Online	5.00 GiB	Nearline_Pool
Copy 0*	✓ Online	5.00 GiB	Nearline_Pool
Copy 1	✓ Online	5.00 GiB	EasyTier_Pool2
Volume001	✓ Online	5.00 GiB	EasyTier_Pool
SASVol	✓ Online	1.00 GiB	V3700_R5Pool_01

Figure 8-52 Volume showing in new Storage Pool

The volume copy migrates to the new Storage Pool without any downtime. It is also possible to migrate both volume copies to other Storage Pools.

The volume copy feature can also be used to migrate volumes to a different pool with a different extent size, as described in 8.6.6, “Migrating volumes by using the volume copy features” on page 413.

8.4.10 Deleting a volume

To delete a volume, complete the following steps:

1. Select **Delete** from the Action menu.
2. Confirm the number of volumes to be deleted. If you want to force the deletion, select the **Delete the volume even if it has host mappings** option, as shown in Figure 8-53.




Figure 8-53 Delete Volume pane

3. Click **Delete** and the volumes are removed from the system.
4. After the task completes, click **Close** to return to the Volumes pane

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

8.5 Volume properties

This section provides an overview of all available information that is related to IBM Storwize V3700 volumes.

To open the advanced view of a volume, complete the following steps:

1. Select **Properties** from the Action menu, and the Volume Details pane opens, as shown in Figure 8-54 on page 397.
2. The following tabs are available, and are described in the following sections:
 - Overview
 - Host Maps
 - Member MDisk

8.5.1 Overview tab

The Overview tab gives you a complete overview of the volume properties. The left part of the pane displays common volume properties, and the right part of the pane displays information about the volume copies, as shown in Figure 8-54.

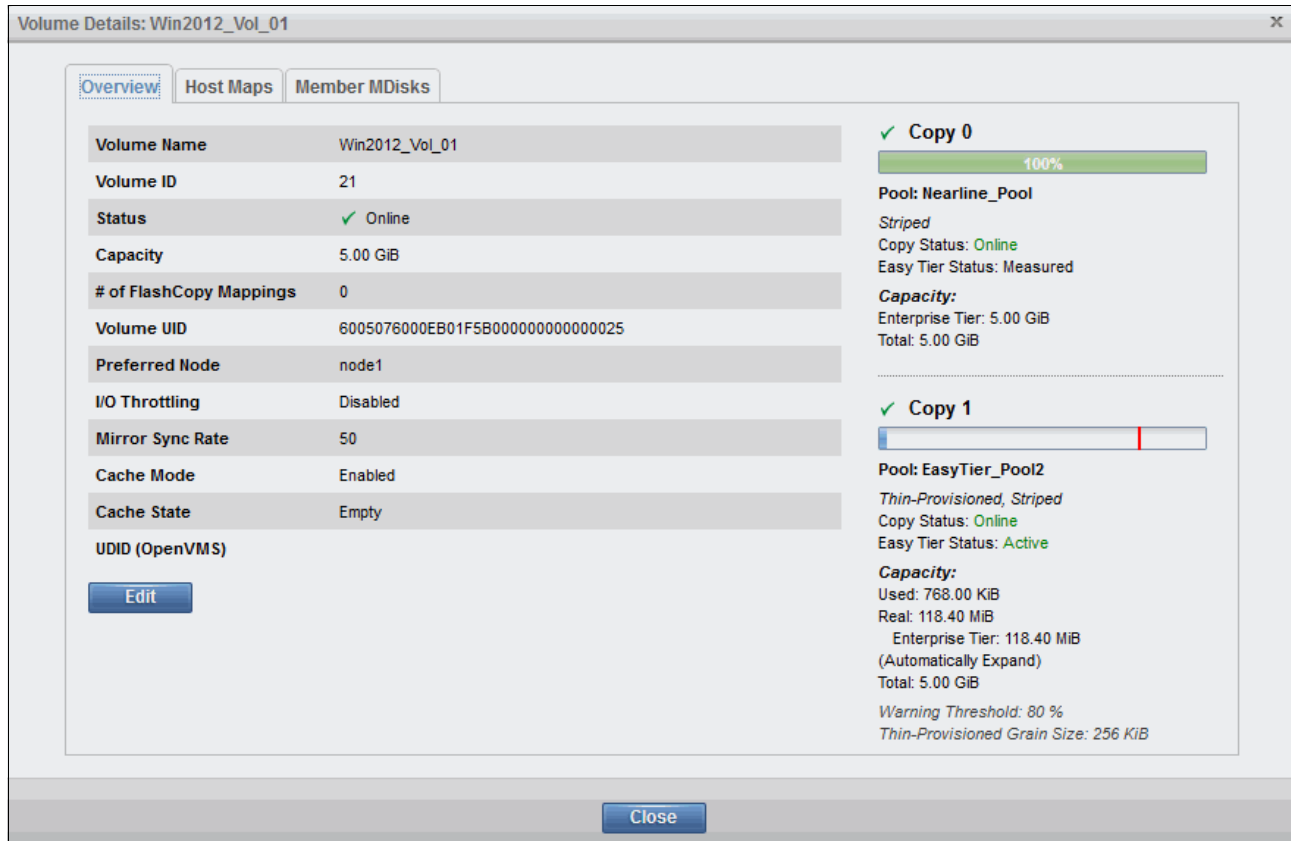


Figure 8-54 Volume Details: Overview

The following details are available:

► Volume Properties:

Volume Name	Shows the name of the volume.
Volume ID	Shows the ID of the volume.
Status	Gives status information about the volume, which can be online, offline, or degraded.
Capacity	Shows the capacity of the volume. If the volume is thin-provisioned, this number is the virtual capacity. The real capacity is displayed for each copy.
# of FlashCopy Mappings	The number of existing FlashCopy relationships. For more information, see Chapter 10, "Copy services" on page 455.
Volume UID	The volume unique identifier.
Preferred Node	Specifies the ID of the preferred node for the volume.

I/O Throttling	It is possible to set a maximum rate at which the volume processes I/O requests. The limit can be set in I/Os to megabytes per second (MBps). This feature is an advanced feature, and it is possible to enable it only through the CLI, as described in Appendix A, “Command-line interface setup and SAN Boot” on page 649.
Mirror Sync Rate	After creation, or if a volume copy is offline, the mirror sync rate weights the synchronization process. Volumes with a high sync rate (100%) complete the synchronization faster than volumes with a lower priority. By default, the rate is set to 50% for all volumes.
Cache Mode	Shows if the cache is enabled or disabled for this volume.
Cache State	Provides feedback if open I/O requests are inside the cache that are not destaged to the disks.
UDID (OpenVMS)	The unit device identifiers (UDIDs) are used by OpenVMS hosts to access the volume.
▶ Copy Properties:	
Storage Pool	Provides information about which pool the copy is in, what type of copy it is (generic or thin-provisioned), the status of the copy, and IBM Easy Tier status.
Capacity	Shows the allocated (used) and the virtual (real) capacity from both (each storage tier and the warning threshold), and the grain size for Thin-Provisioned volumes.

3. It is possible to modify some of these settings. Click **Edit** to change the available settings, as shown in Figure 8-55.

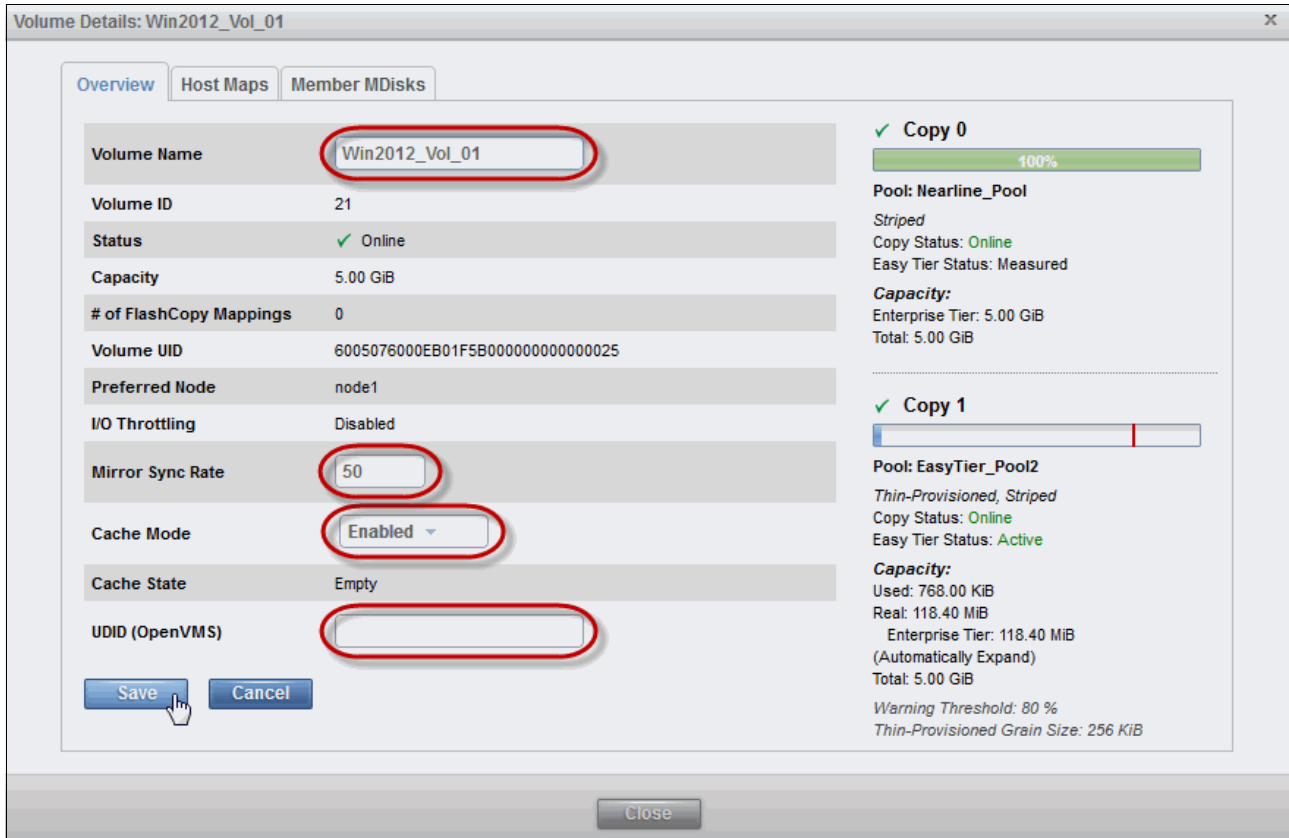


Figure 8-55 Modify Volume Details pane

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

- Volume Name
- Mirror Sync Rate
- Cache Mode
- UDID

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

8.5.2 Host Maps tab

The second tab of the Volume Details pane is Host Maps, as shown in Figure 8-56. All hosts that are mapped to the selected volume are listed in this view.

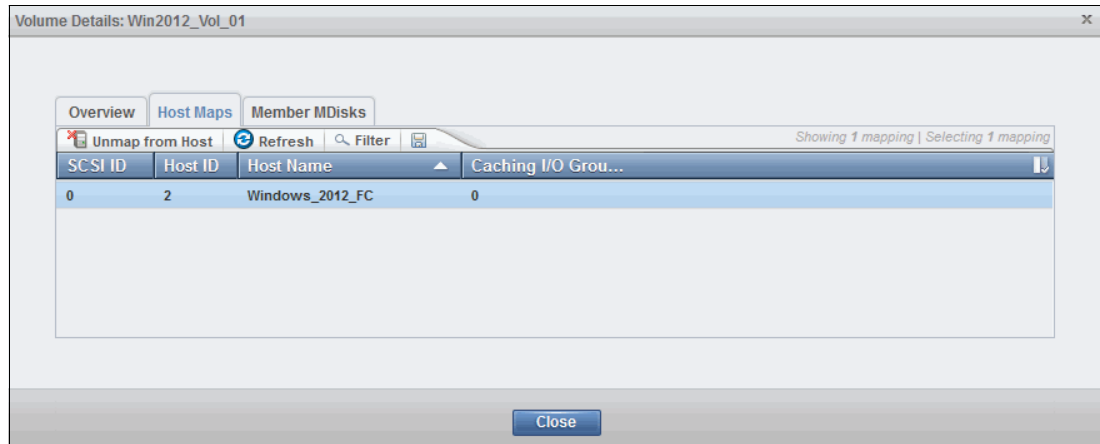


Figure 8-56 Volume Details: Host Maps

To unmap a host from the volume, complete the following steps:

1. Highlight the volume and click **Unmap from Host**.
2. Confirm the number of mappings to remove and click **Unmap**. Figure 8-57 shows the Unmap Host pane.

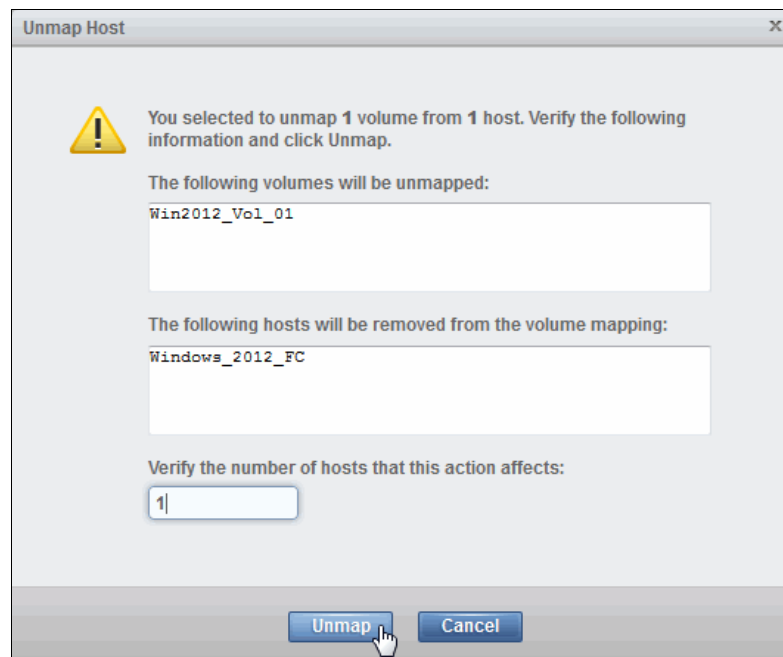


Figure 8-57 Unmap Host pane

3. When the changes are applied to the system, the selected host no longer has access to this volume. Click **Close** to return to the Host Maps window. For more information about host mappings, see 8.3, "Host mappings overview" on page 383.

8.5.3 Member MDisks tab

The third tab is the Member MDisks tab, which lists all MDisks on which the volume is located:

1. If there are multiple copies of the volume, it is possible to select an individual copy, or choose to display **Both Copies**, as shown in Figure 8-58. The associated MDisks are displayed.

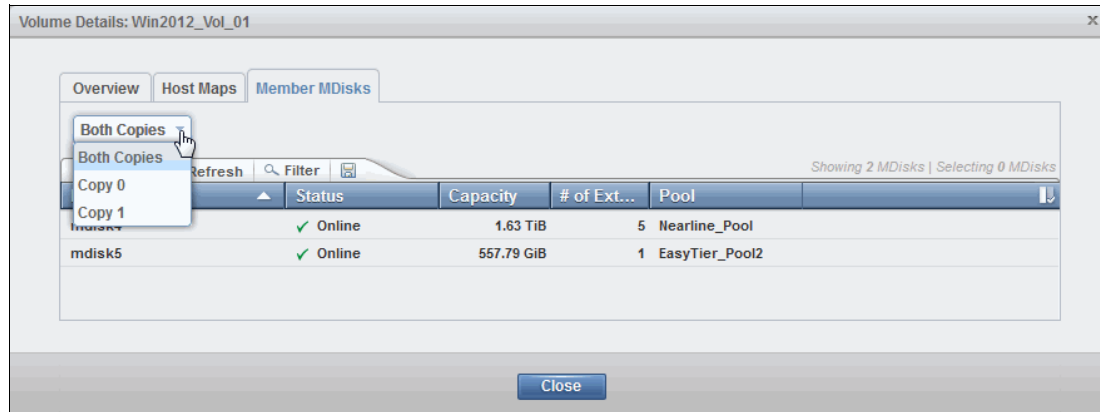


Figure 8-58 Member MDisk tab

2. Select an MDisk and click **Actions** to see the available tasks, as shown in Figure 8-59. For more information about the available tasks, see Chapter 7, “Storage Pools” on page 307.




Figure 8-59 MDisk Actions menu

3. Click **Close** to return to the Volumes pane.

8.5.4 Adding a mirrored volume copy

If a volume consists of only one copy, it is possible to add a second mirrored copy of the volume. This second copy can be either generic or thin-provisioned.

You can also use this method to migrate data across Storage Pools with different extent sizes.

To add a second copy, complete the following steps:

1. Select the volume and click **Actions** → **Volume Copy Actions** → **Add Mirrored Copy**, as shown in Figure 8-60.

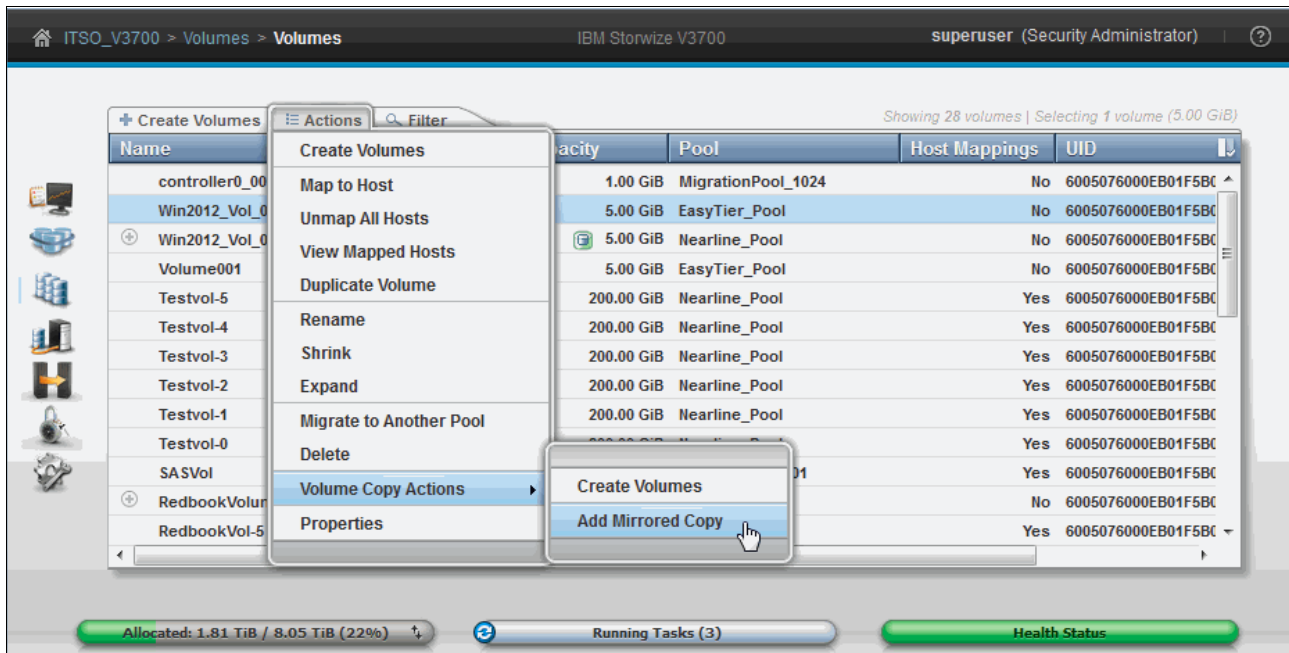


Figure 8-60 Add Mirrored Copy

2. Select if the volume should be Generic or Thin-Provisioned, select the Storage Pool in which the new copy should be created, and click **Add Copy**, as shown in Figure 8-61.

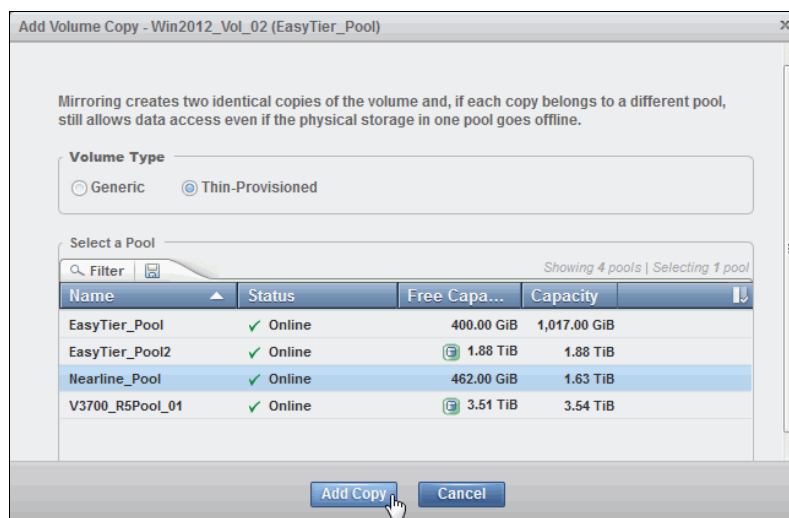


Figure 8-61 Select Storage Pool

The copy is created after clicking **Add Copy**, and data starts to synchronize as a background task. Figure 8-62 shows that the volume named Win2012_Vol_02 now has two volume copies that are stored in two different Storage Pools, and that the new copy is thin-provisioned.

Name	State	Capacity	Pool
controller0_000000000000...	✓ Online	1.00 GiB	MigrationPool_1024
Win2012_Vol_02	✓ Online	5.00 GiB	EasyTier_Pool
Copy 0*	✓ Online	5.00 GiB	EasyTier_Pool
Copy 1	✓ Online	5.00 GiB	Nearline_Pool
Win2012_Vol_01	✓ Online	5.00 GiB	Nearline_Pool
Volume001	✓ Online	5.00 GiB	EasyTier_Pool
Testvol-5	✓ Online	200.00 GiB	Nearline_Pool
Testvol-4	✓ Online	200.00 GiB	Nearline_Pool
Testvol-3	✓ Online	200.00 GiB	Nearline_Pool
Testvol-2	✓ Online	200.00 GiB	Nearline_Pool

Figure 8-62 Volume copies

8.5.5 Editing thin-provisioned volume properties

The processes that are used to modify the volume size that is presented to a host are described in 8.4.7, “Shrinking a volume” on page 392 and 8.4.8, “Expanding a volume” on page 393. However, with a thin-provisioned volume, it is also possible to edit the allocated size and the warning thresholds.

To edit these settings, complete the following steps:

1. Select the volume then select **Actions** → **Volume Copy Actions** → **Thin-Provisioned**, as shown in Figure 8-63 on page 404.

The following options are available, as shown in Figure 8-63:

- Shrink
- Expand
- Edit Properties




Figure 8-63 Working with thin-provisioned volumes

Shrinking thin-provisioned space

To shrink thin-provisioned space, complete the following steps:

1. Select **Shrink**, as shown in Figure 8-63, to reduce the allocated space of a thin-provisioned volume.
2. Enter the amount by which the volume should shrink, or enter the new final size, and then click **Shrink**.

Deallocating extents: Only extents that do not include stored data can be deallocated. If the space is allocated because there is data on the extent, the allocated space cannot be shrunk and an out-of-range warning message displays.

Figure 8-64 shows the Shrink Volume window.

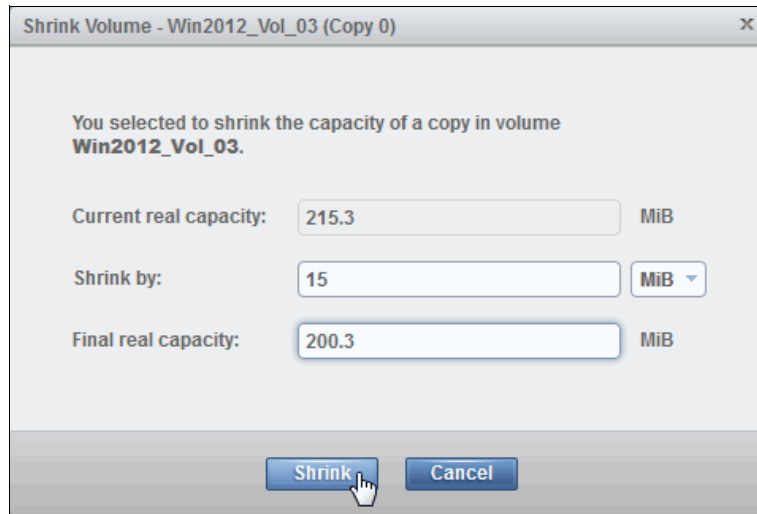


Figure 8-64 Shrink Volume pane

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

Expanding thin-provisioned space

To expand the allocated space of a thin-provisioned volume, complete the following steps:

1. Select **Expand**, as shown in Figure 8-63 on page 404.
2. Enter the amount by which space should be allocated or the new final size and click **Expand**. In this example, shown in Figure 8-65, we are expanding the thin-provisioned space by 15 mebibytes (MiB).

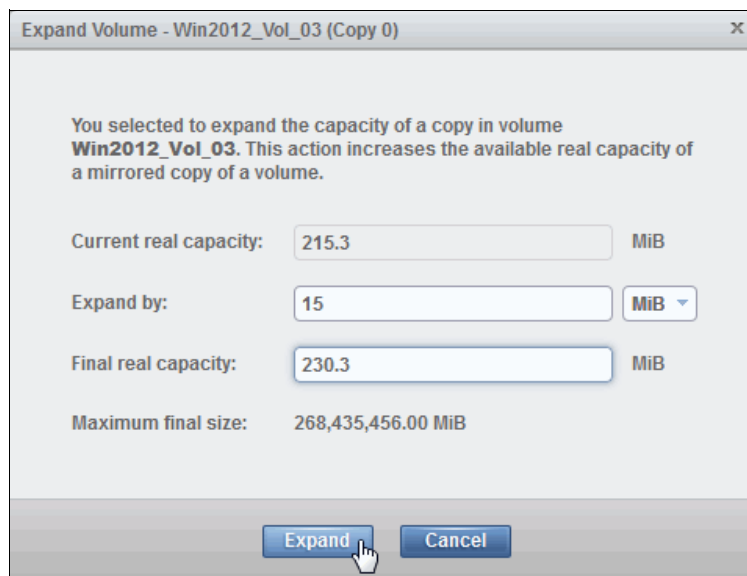


Figure 8-65 Expand Thin-Provisioned Volume pane

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

Editing thin-provisioned properties

To edit thin-provisioned properties, complete the following steps:

1. Select **Edit Properties**, as shown in Figure 8-63 on page 404.
2. Edit the settings (if required) and click **OK** to apply the changes.

Figure 8-66 shows the Edit Properties window.

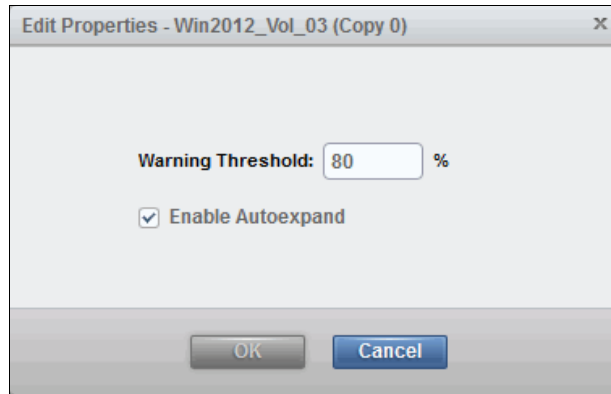


Figure 8-66 Edit Properties pane

3. After the task completes, click **Close** to return to the Volumes pane.

8.6 Advanced volume copy functions

In 8.4.1, “Advanced volume functions” on page 385, we described all of the available actions at a volume level, and how to create a second volume copy. In this section, we focus on volumes that consist of two volume copies, and how to apply the concept of two copies for business continuity and data migration.

Expanding a volume, selecting a copy, and opening the Action menu displays the following volume copy actions, as shown in Figure 8-67 on page 407:

- ▶ Thin-Provisioned (for thin volumes)
- ▶ Make Primary (for non-primary copy)
- ▶ Split into New Volume
- ▶ Validate Volume Copies
- ▶ Duplicate Volume
- ▶ Delete this Copy

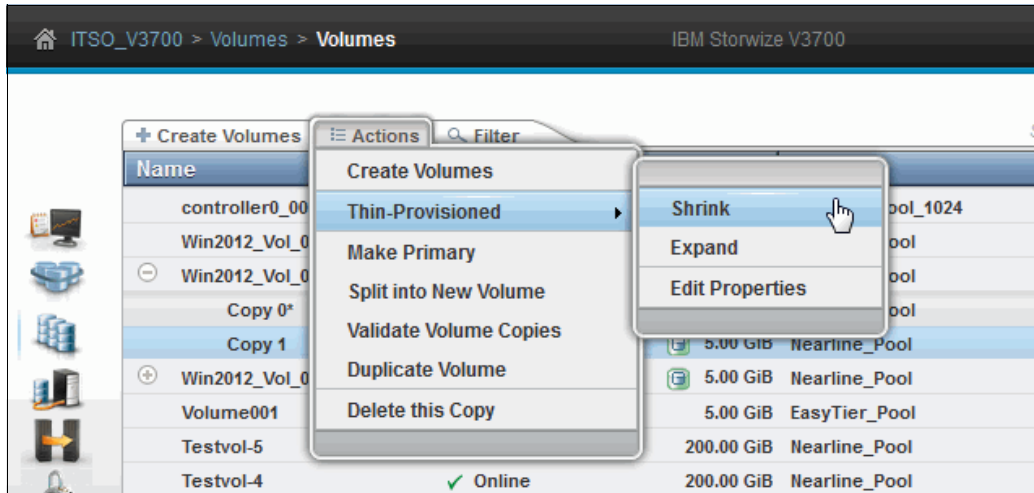


Figure 8-67 Volume copy actions

Looking at the volume copies that are shown in Figure 8-67, it is possible to see that one of the copies has an asterisk (*) displayed next to its name, as also shown in Figure 8-68.

Win2012_Vol_02	✓ Online	10.00 GiB	EasyTier_Pool
Copy 0*	✓ Online	10.00 GiB	EasyTier_Pool
Copy 1	✓ Online	10.00 GiB	Nearline_Pool

Figure 8-68 Volume copy names

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

To accomplish this task, complete the following steps:

1. Select the secondary copy and then click **Actions** → **Make Primary**. Usually, it is a leading practice to place the volume copies on Storage Pools with similar performance, because the write performance is constrained if one copy is on a lower performance pool.

Figure 8-69 shows the secondary copy Actions menu.

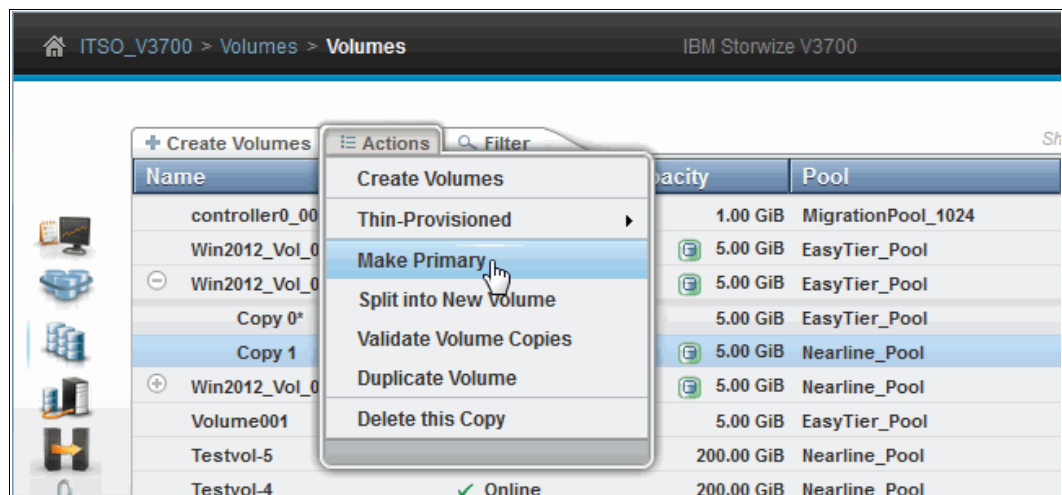


Figure 8-69 Make primary

If high read performance is demanded, another possibility is to place the primary copy in an solid-state drive (SSD) pool and the secondary copy in an hard disk drive (HDD) pool. This action maximizes the read performance of the volume, and makes sure that you have a synchronized second copy in your less expensive disk pool.

It is possible to migrate online copies between Storage Pools. For more information about how to select which copy you want to migrate, see 8.4.9, “Migrating a volume to another Storage Pool” on page 393.

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

The volume copy feature is also a powerful option for migrating volumes, as described in 8.6.6, “Migrating volumes by using the volume copy features” on page 413.

8.6.1 Thin-provisioned menu

The Thin-Provisioned menu includes the same functions that are described in “Shrinking thin-provisioned space” on page 404, “Expanding thin-provisioned space” on page 405, and “Editing thin-provisioned properties” on page 406.

Figure 8-70 shows the Thin-Provisioned menu items.

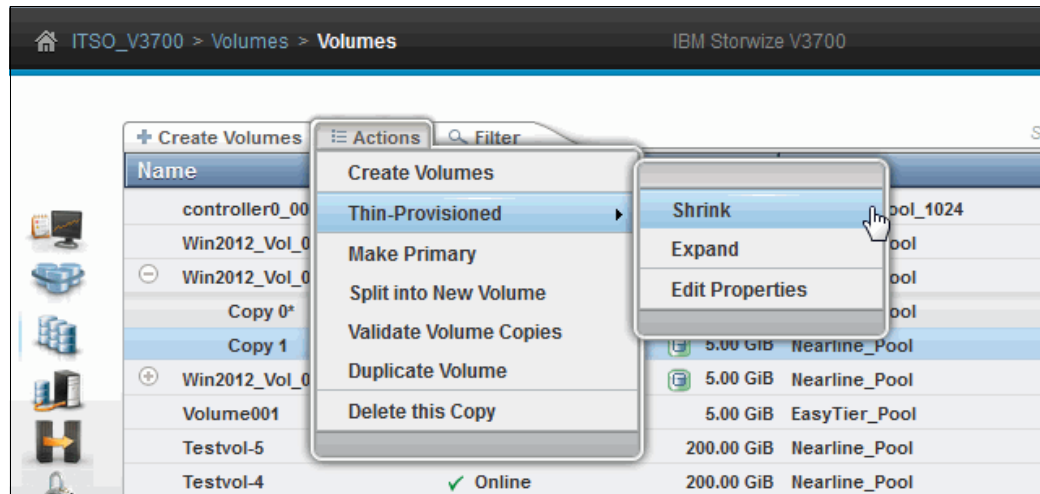


Figure 8-70 Thin-Provisioned menu items

8.6.2 Splitting into a new volume

If the two-volume copies are synchronized, it is possible to split one of the copies to a new volume, and map this volume to another host. From a storage point of view, this procedure can be performed online, which means that you can split one copy from the volume and create a copy from the remaining volume without any host effect.

However, if you want to use the split copy for testing or backup purposes, you must make sure that the data inside the volume is consistent. Therefore, the data must be flushed to storage to make the copies consistent.

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

In this example, volume Win2012_Vol_01 has two copies: Copy 0 as primary and Copy 1 as secondary. To split a copy, complete the following steps:

1. Click **Split into New Volume** (as shown in Figure 8-67 on page 407) on any copy, and the remaining secondary copy automatically becomes the primary for the source volume. Optionally, a name for the new volume can be specified.

Figure 8-71 shows the Split Volume Copy pane.

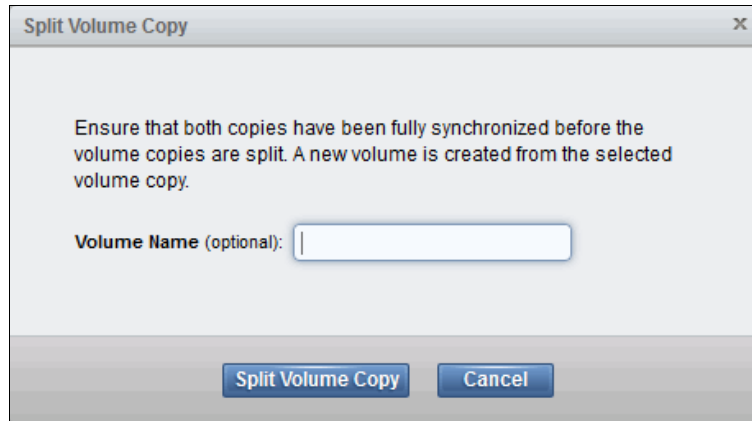


Figure 8-71 Split Volume Copy pane

2. After the task completes, click **Close** to return to the Volumes pane, where the copy displays as a new volume named `vdisk0` (unless a name was specified) that can be mapped to a host, as shown in Figure 8-72.




Figure 8-72 Volumes: New volume from split copy

Important: If you receive error message code CMMVC6357E while you are splitting a volume copy, use the `lsvdisksyncprogress` command to view the synchronization status, or wait for the copy to synchronize.

Example 8-1 shows output of the `lsvdisksyncprogress` command.

Example 8-1 Output of the lsvdisksyncprogress command

```
IBM_Storwize:ITS0_V3700:superuser>lsvdisksyncprogress
vdisk_id vdisk_name      copy_id progress estimated_completion_time
30       Win2012_Vol_02 1       77       141031132545
IBM_Storwize:ITS0_V3700:superuser>
```

8.6.3 Validate Volume Copies option

It is possible to verify that the volume copies are identical, or process the differences between them.

To validate the copies of a mirrored volume, complete the following steps:

1. Select **Validate Volume Copies**, as shown in Figure 8-67 on page 407. The Validate Volume Copies pane opens, as shown in Figure 8-73.

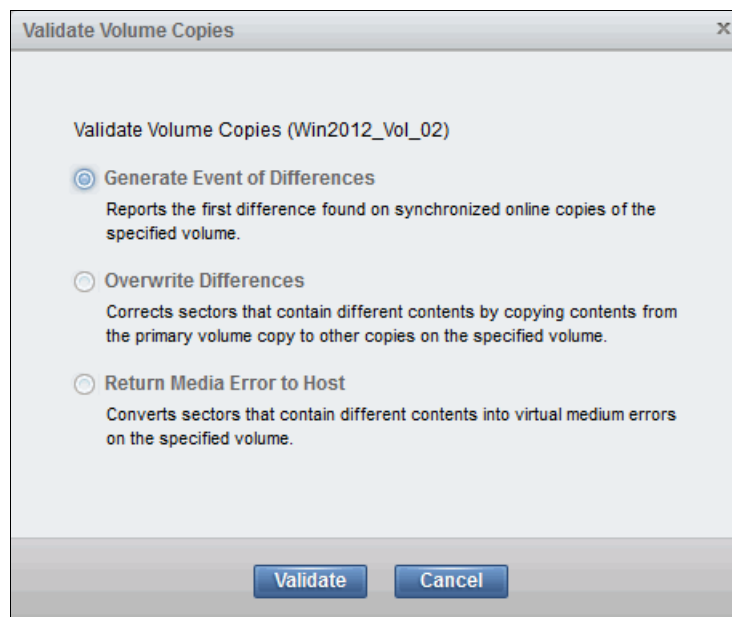


Figure 8-73 Validate Volume Copies pane

The following options are available:

- Generate Event of Differences

Use this option if you want to verify only that the mirrored volume copies are identical. If any difference is found, the command stops and logs an error that includes the logical block address (LBA) and the length of the first difference. You can use this option, starting at a different LBA each time, to count the number of differences on a volume.

- Overwrite Differences

Use this option to overwrite contents from the primary volume copy to the other volume copy. The command corrects any differing sectors by copying the sectors from the primary copy to the copies that are compared. Upon completion, the command process logs an event that indicates the number of differences that were corrected. Use this option if you are sure that the primary volume copy data is correct, or that your host applications can handle incorrect data.

- Return Media Error to Host

Use this option to convert sectors on all volume copies that contain different contents into virtual medium errors. Upon completion, the command logs an event, which indicates the number of differences that were found, the number that were converted into medium errors, and the number that were not converted. Use this option if you are unsure what the correct data is, and you do not want an incorrect version of the data to be used.

2. Select which action to perform and click **Validate** to start the task. The volume is now checked. Click **Close**.

Figure 8-74 shows the output when the volume copy **Generate Event of Differences** option is chosen.

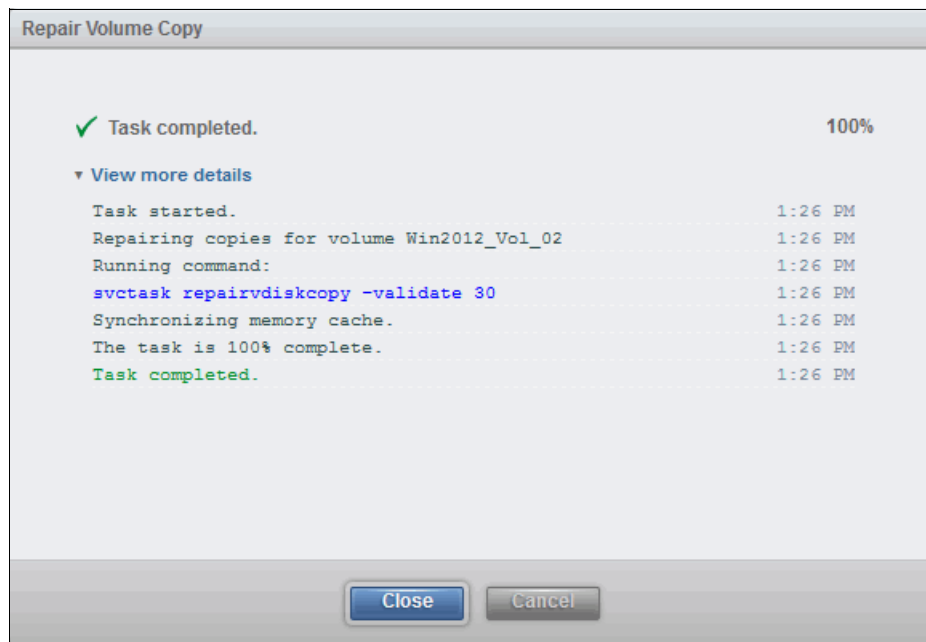


Figure 8-74 Volume copy validation output

The validation process runs as a background process and might take some time, depending on the volume size.

3. You can check the status in the Running Tasks window, as shown in Figure 8-75.

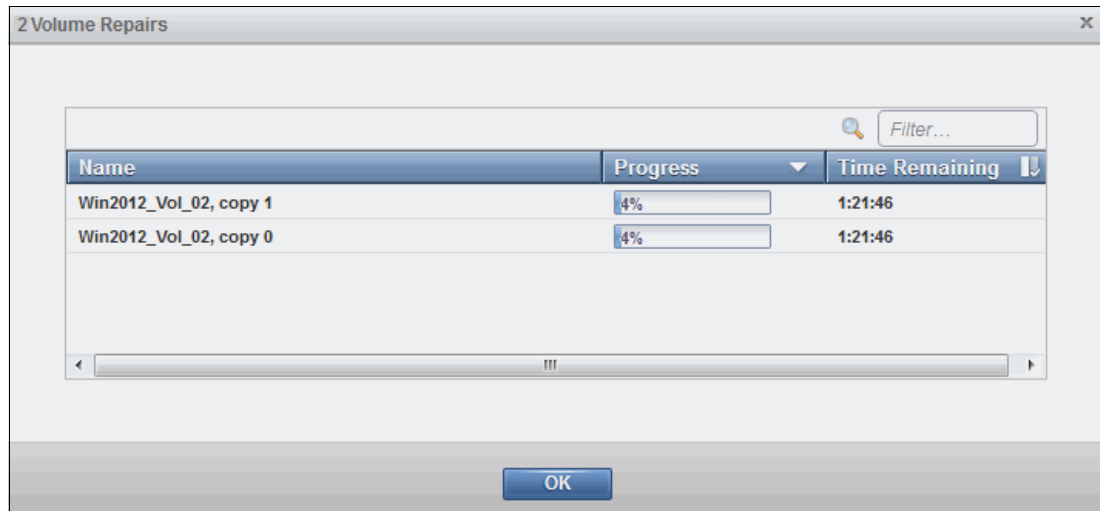


Figure 8-75 Validate Volume Copies: Running Tasks

8.6.4 Duplicate Volume option

To make an exact copy of an existing volume copy, complete the following steps:

1. After selecting a volume copy, click **Duplicate Volume** (as shown in Figure 8-67 on page 407).
2. Enter a name for the duplicate volume and click **Duplicate** to apply the change, as shown in Figure 8-76.

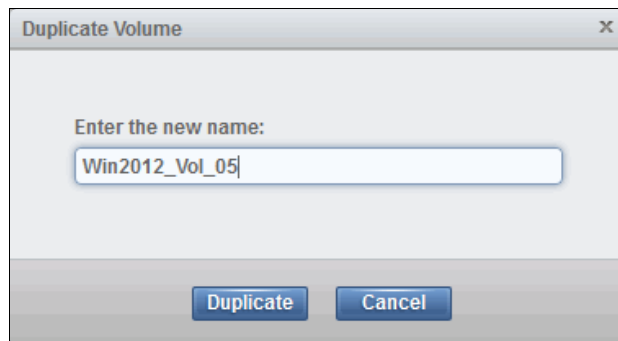


Figure 8-76 Create a duplicate volume

8.6.5 Delete Volume Copy option

To delete a volume copy, complete the following steps:

1. Click **Delete** (as shown in Figure 8-67 on page 407) to delete a volume copy. The copy is deleted, but the volume remains online by using the remaining copy.

2. Confirm the deletion process by clicking **Yes**. Figure 8-77 shows the copy deletion Warning pane.

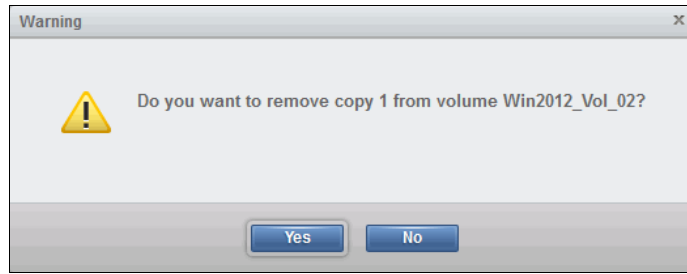


Figure 8-77 Delete a copy

3. After the copy is deleted, click **Close** to return to the Volumes pane.

8.6.6 Migrating volumes by using the volume copy features

In the previous sections, we showed that it is possible to create, synchronize, split, and delete volume copies. A combination of these tasks can be used to migrate volumes to other Storage Pools.

The easiest way to migrate volume copies is to use the migration feature that is described in 8.4.9, “Migrating a volume to another Storage Pool” on page 393. Using this feature, one extent after another is migrated to the new Storage Pool. However, the use of volume copies provides another way to migrate volumes if you have different Storage Pool characteristics in terms of extent size.

To migrate a volume, complete the following steps:

1. Create a second copy of your volume in the target Storage Pool. For more information, see 8.5.4, “Adding a mirrored volume copy” on page 402.
2. Wait until the copies are synchronized.
3. Change the role of the copies and make the new copy the primary copy. For more information, see 8.6, “Advanced volume copy functions” on page 406.
4. Split or delete the old copy from the volume. For more information, see 8.6.2, “Splitting into a new volume” on page 408 or 8.6.5, “Delete Volume Copy option” on page 412.

This migration process requires more user interaction with the IBM Storwize V3700 GUI, but it offers some benefits.

As an example, we look at migrating a volume from a tier 1 Storage Pool to a lower performance tier 2 Storage Pool.

In step 1, you create the copy on the tier 2 pool, while all reads are still performed in the tier 1 pool to the primary copy. After the synchronization, all writes are destaged to both pools, but the reads are still done only from the primary copy.

Because the copies are fully synchronized, you can switch their role online (see step 3 previously), and analyze the performance of the new pool. When you are done testing your lower performance pool, you can split or delete the old copy in tier 1, or switch back to tier 1 in seconds if the tier 2 Storage Pool did not meet your requirements.

8.7 Volumes by Storage Pool

To see an overview of which volumes are on which Storage Pool, perform the following steps:

1. Click **Volumes by Pool**, as shown in Figure 8-78.




Figure 8-78 Volumes by Pool

2. The Volumes by Pool pane opens, as shown in Figure 8-79.

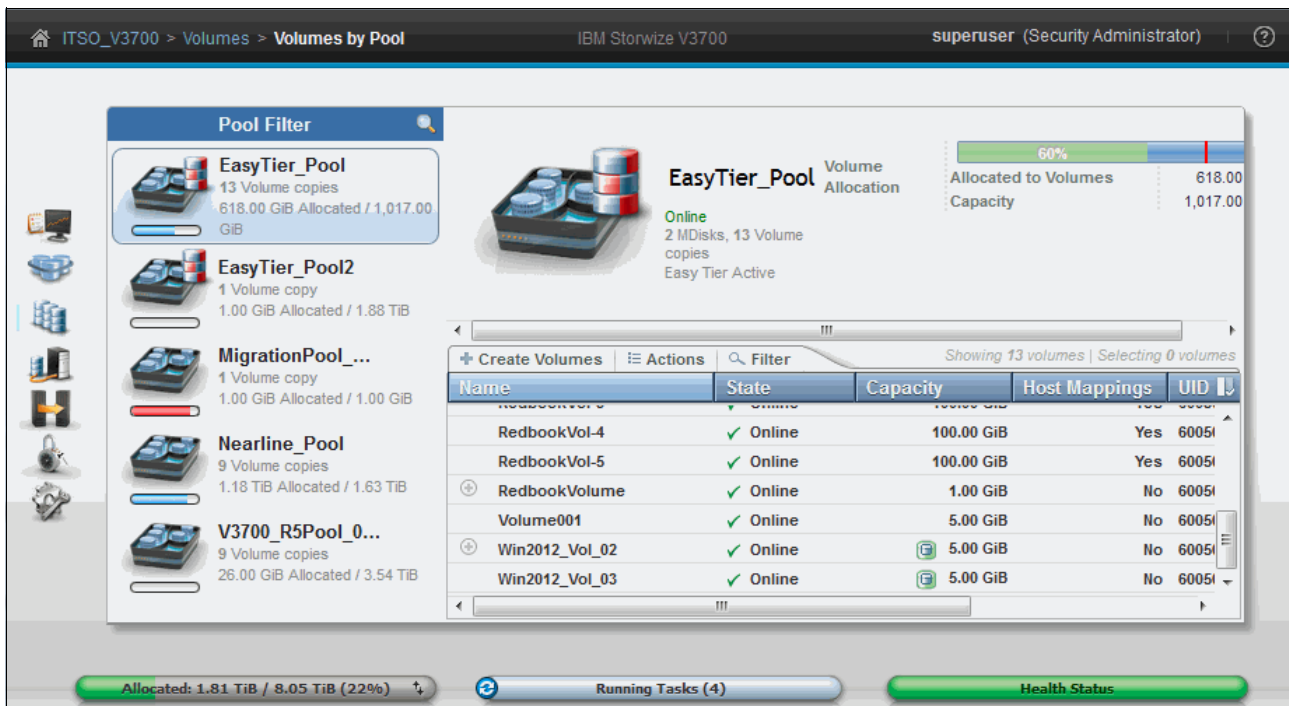


Figure 8-79 Volumes by Pool pane

The left pane is called the Pool Filter and the Storage Pools are displayed there. For more information about Storage Pools, see Chapter 7, “Storage Pools” on page 307.

In the upper right, you see information about the pool that you selected in the pool filter. The following information is also shown:

- Pool icon** Because Storage Pools can have different characteristics, it is possible to change the Storage Pool icon. For more information, see 7.3, “Working with Storage Pools” on page 348.
- Pool Name** The name that is given during the creation of the Storage Pool. For more information about changing the Storage Pool name, see “Rename” on page 353.
- Pool Details** Shows you the information about the Storage Pools, such as status, the number of managed disks, and Easy Tier status.
- Volume allocation** Shows you the amount of capacity that is allocated to volumes from this Storage Pool.

The lower right section (as shown in Figure 8-80) lists all volumes that have at least one copy in the selected Storage Pool. The following information is provided:

- Name** Shows the name of the volume.
- Status** Shows the status of the volume.
- Capacity** Shows the capacity that is presented to hosts.
- UID** Shows the volume unique identifier.
- Host Mappings** Shows if host mappings exist.

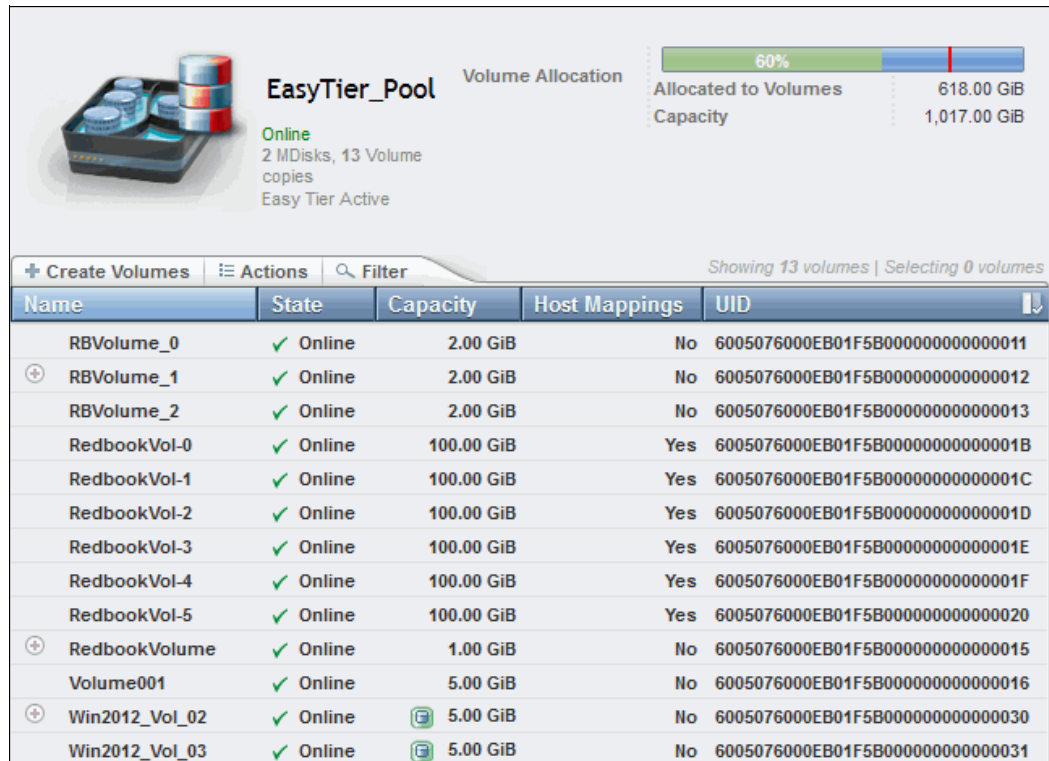


Figure 8-80 Volumes by Storage Pool

3. It is also possible to create volumes from this pane. Click **Create Volume** to start the Volume Creation pane, as described in Chapter 5, “Volume configuration” on page 199.
4. Selecting a volume and opening the **Actions** menu, or right-clicking the volume, shows the same options as described in 8.4, “Advanced volume administration” on page 384.

8.8 Volumes by Host

To see an overview on which volumes a host can access, complete the following steps:

1. Click **Volumes by Host** from the Volumes menu and the Volumes by Host window opens, as shown in Figure 8-81.



Figure 8-81 Volumes by Host

2. In the left pane of the view is the Host Filter. Selecting a host, shows its properties in the right pane, such as the host name, number of ports, and host type. The Host icons show an orange cable for a FC host, a black cable for an SAS host, and a blue cable for an iSCSI host. The volumes that are mapped to this host are listed, as shown in Figure 8-82.



Figure 8-82 Volumes by Host

3. It is also possible to create a volume from this pane. Clicking **Create Volumes**, opens the same wizard as described in 5.1, “Creating volumes in IBM Storwize V3700” on page 200.
4. Selecting a volume and opening the **Actions** menu or right-clicking the volume, shows the same options as described in 8.4, “Advanced volume administration” on page 384.



IBM Easy Tier

In today's storage industry, flash drives are emerging as an attractive alternative to hard disk drives (HDDs). Because of their low response times, high throughput, and input/output (I/O) operations per second (IOPS)-energy-efficient characteristics, flash drives have the potential to allow your storage infrastructure to achieve significant savings in operational costs. However, the current acquisition cost per gigabyte (GB) for flash is higher than for Enterprise serial-attached Small Computer System Interface (SAS) and Nearline SAS (NL-SAS).

Enterprise SAS drives replaced the old Small Computer System Interface (SCSI) drives, and have become common in the storage market. They are offered in a wide variety of capacity, spindle speeds, and form factors. NL-SAS is the low-cost, large-capacity storage drive class, commonly offered at 7200 revolutions per minute (RPM) spindle speed.

It is critical to choose the right mix of drives and the right data placement to achieve optimal performance at low cost. Maximum value can be derived by placing data as follows:

- ▶ “Hot” data with high I/O density and low response time requirements on flash
- ▶ “Warm” data targeted on Enterprise class disks
- ▶ “Cold” data accessed more sequentially and at lower rates on Nearline disks

This chapter describes the function provided by the IBM Easy Tier disk performance optimization function of the IBM Storwize V3700. It also describes how to activate the Easy Tier process for both evaluation purposes and automatic extent migration.

This chapter includes the following topics:

- ▶ Generations of IBM Easy Tier
- ▶ New Features in Easy Tier 3
- ▶ Easy Tier overview
- ▶ Easy Tier process
- ▶ Easy Tier configuration using the graphical user interface (GUI)
- ▶ Easy Tier configuration using the command-line interface (CLI)
- ▶ IBM Storage Tier Advisor Tool (STAT)

9.1 Generations of IBM Easy Tier

IBM Storwize Family software has benefited from software development work for the IBM System Storage DS8000 product, in which there have been six versions of Easy Tier. Of those versions, versions 1 and 3 have been implemented within IBM Storwize Family software.

The first generation of Easy Tier introduced automated storage performance management, efficiently boosting enterprise-class performance with flash drives, and automating storage tiering from enterprise class drives to flash drives.

The second generation of Easy Tier introduced the combination of Nearline drives with the objective of maintaining optimal performance at low cost. The second generation also introduced auto rebalancing of extents, and was implemented within DS8000 products only.

The third generation of Easy Tier introduces further enhancements that provide automated storage performance and storage economics management across all three drive tiers (flash, Enterprise, and Nearline storage tiers). It enables you to consolidate and efficiently manage more workloads on a single IBM Storwize V3700. It also introduces support for Storage Pool Balancing in homogeneous pools.

9.2 New Features in Easy Tier 3

The following enhancements are included in Easy Tier 3:

- ▶ Support for three tiers of disk, or a mixture of any two tiers
- ▶ Storage Pool Balancing
- ▶ Enhancements to the Storage Tier Advisor Tool, including additional graphing from the STAT utility

Figure 9-1 shows the supported Easy Tier pools available in the third generation of Easy Tier.

Tier 0	Tier 1	Tier 2
Three tier Pools:		
SSD	Enterprise	Nearline
Two Tier Pools:		
SSD	Enterprise	
SSD	Nearline	
	Enterprise	Nearline
Single Tier Pools:		
SSD		
	Enterprise	
		Nearline

Figure 9-1 Supported Easy Tier Pools

9.3 Easy Tier overview

Easy Tier is an optional licensed function of IBM Storwize V3700 that brings enterprise storage enhancements to the entry-level segment. It enables automated subvolume data placement throughout different storage tiers to intelligently align the system with current workload requirements, and to optimize the usage of storage.

This function includes the ability to automatically and non-disruptively relocate data (at the extent level) from one tier to another tier in either direction. This helps to achieve the best available storage performance for your workload in your environment.

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

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

Easy Tier is a performance optimization function that overcomes this issue, because it automatically migrates (or moves) extents that belong to a volume between different storage tiers, as shown in Figure 9-2. Because this migration works at the extent level, it is often referred to as *sub-logical unit number (LUN) migration*.

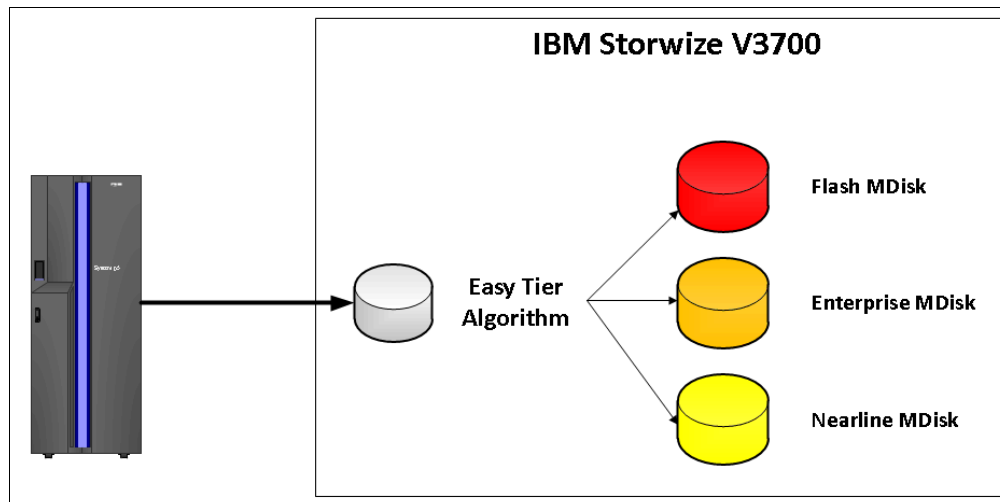


Figure 9-2 Easy Tier

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

To enable this migration between managed disks (MDisks) with different tier levels, the target Storage Pool must consist of MDisks with different characteristics. These pools are named *multitiered Storage Pools*. IBM Storwize V3700 Easy Tier is optimized to boost the performance of Storage Pools that contain flash, Enterprise, and Nearline drives.

To identify the potential benefits of Easy Tier in your environment before actually installing higher MDisk tiers (such as flash), it is possible to enable the Easy Tier monitoring on volumes in single-tiered Storage Pools. Although the Easy Tier extent migration is not possible within a single-tiered pool, the Easy Tier statistical measurement function is possible. Enabling Easy Tier on a single-tiered Storage Pool starts the monitoring process, and logs the activity of the volume extents.

The STAT tool is a no-cost tool that helps you to analyze this data. If you do not have an IBM Storwize V3700, use Disk Magic to get a better understanding of the required number of different drive types that are appropriate for your workload.

Easy Tier is available for IBM Storwize V3700 internal volumes only.

9.3.1 Easy Tier licensing

To use the Easy Tier functions in the IBM Storwize V3700, you must obtain the authorization codes to your system. Do so by connecting to the IBM data storage feature activation (DSFA) website. When you acquire licensed functions for your IBM Storwize V3700, you receive a license authorization document, which contains detailed instructions about how to retrieve the license keys. To generate and download your authorization code for Easy Tier function, go to the following website:

<https://www.ibm.com/storage/dsfa/home.wss>

To apply and activate the Easy Tier function on your IBM Storwize V3700, complete the following steps:

1. Using the **System** page, go to **Settings** → **System** and choose **Licensing**, as shown in Figure 9-3.

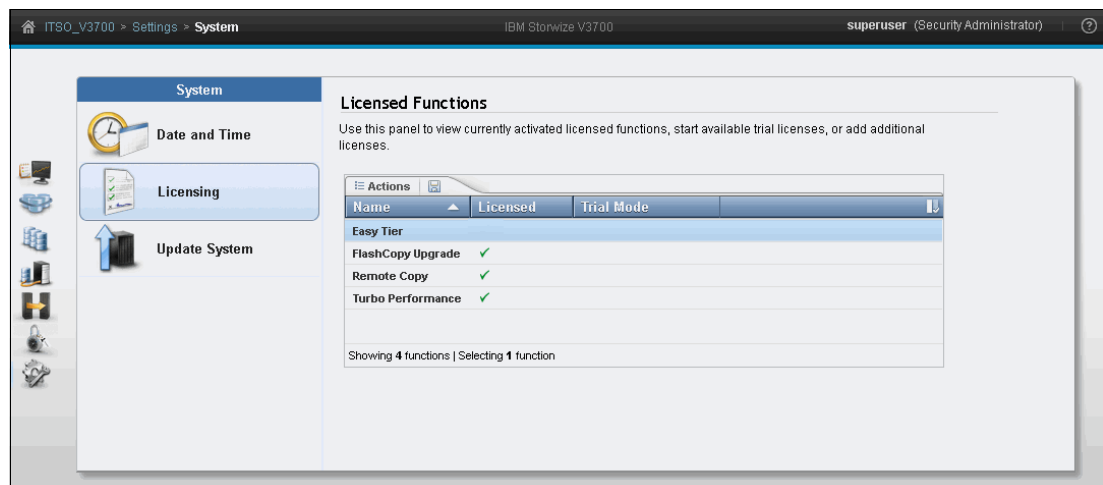


Figure 9-3 Licensed functions window

Alternatively, you can click the question mark in the upper right of the system window as shown in Figure 9-4.

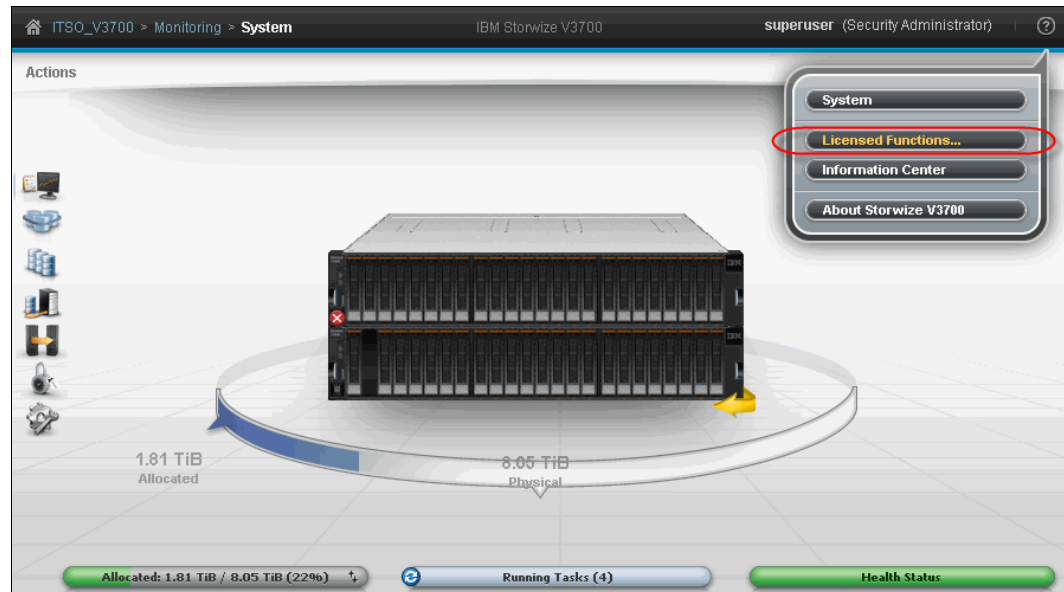


Figure 9-4 Licensed Functions option

2. Select **Easy Tier** → **Actions** and choose **Active License Automatically** or **Activate License Manually**, as shown in Figure 9-5.

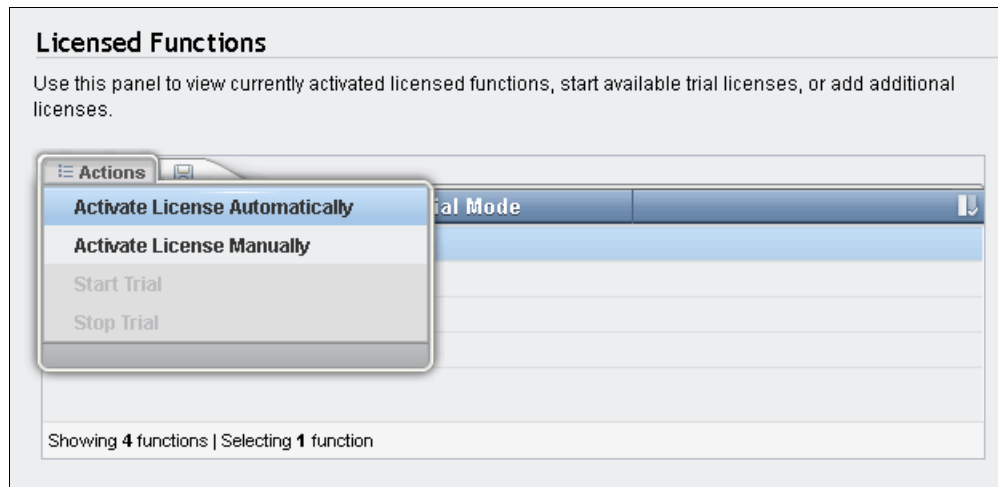


Figure 9-5 Easy Tier activation mode

When activating the Easy Tier function automatically, you are required to apply the authorization code by copying or downloading it from the IBM DSFA website. To activate the function Easy Tier manually, you can either apply the authorization code or click the folder icon to upload the license code as a file.

To activate Easy Tier function automatically, complete the next step. To activate Easy Tier function manually, go to step 4 on page 422.

3. Choose **Activate License Automatically** to activate the Easy Tier function automatically and click **Activate** as shown in Figure 9-6.

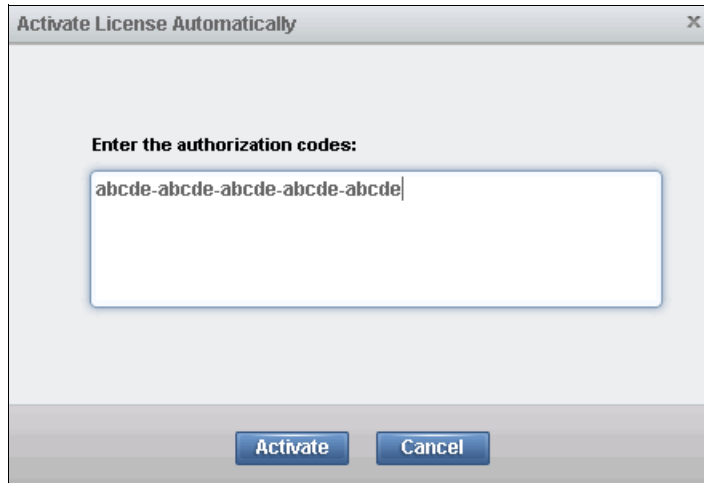


Figure 9-6 Activating license automatically

4. Choose **Activate License Manually** to activate Easy Tier function manually and click in the folder icon to import the activation license file, as shown in Figure 9-7.



Figure 9-7 Activating license manually

5. Change to the directory where the activation license file is saved and click **Open**, as shown in Figure 9-8.

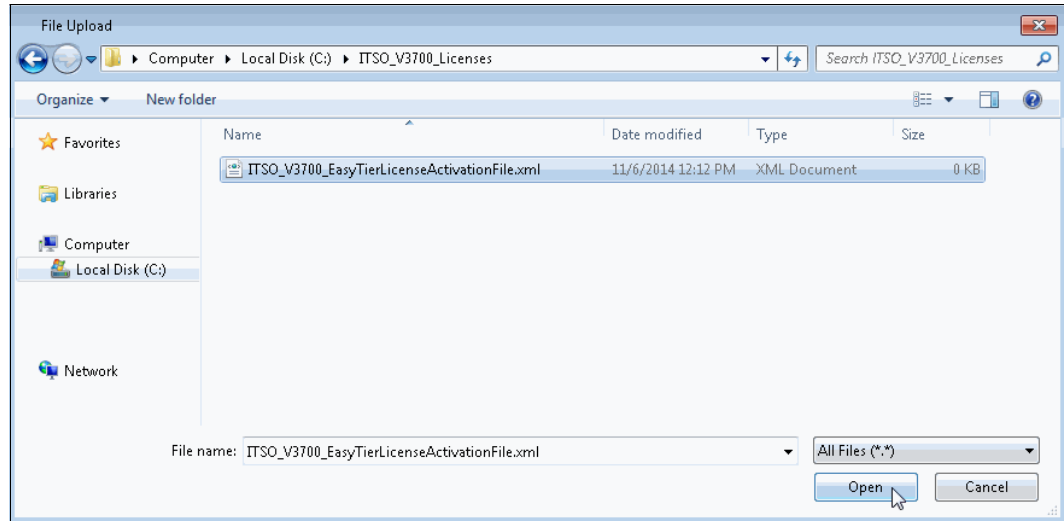


Figure 9-8 Importing the Easy Tier license file

6. Click **Close** when the import of the activation license file completes.
7. As shown in Figure 9-9, the Easy Tier function is licensed on your IBM Storwize V3700.

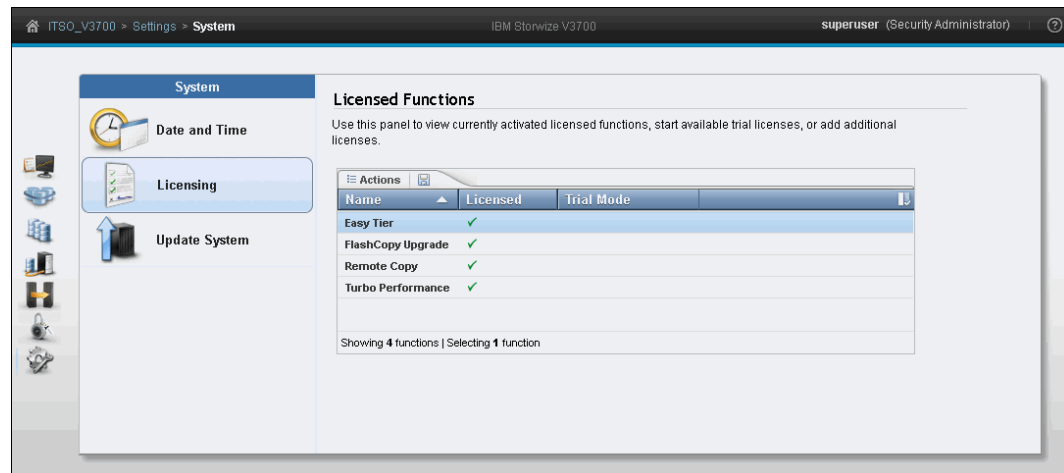


Figure 9-9 Easy Tier function licensed

9.3.2 Tiered Storage Pools

With IBM Storwize V3700, you must differentiate between the following types of Storage Pools:

- ▶ Single-tiered Storage Pools
- ▶ Multitiered Storage Pools

Figure 9-10 on page 424 shows that single-tiered Storage Pools include one type of disk tier attribute. Each disk should have the same size and performance characteristics. Multitiered Storage Pools are populated with two or more different disk tier attributes, high-performance flash drives, enterprise SAS drives, and Nearline drives.

A volume migration is when the complete volume is migrated from one Storage Pool to another Storage Pool. An Easy Tier data migration moves only extents inside the Storage Pool to different performance attributes, as shown in Figure 9-10.

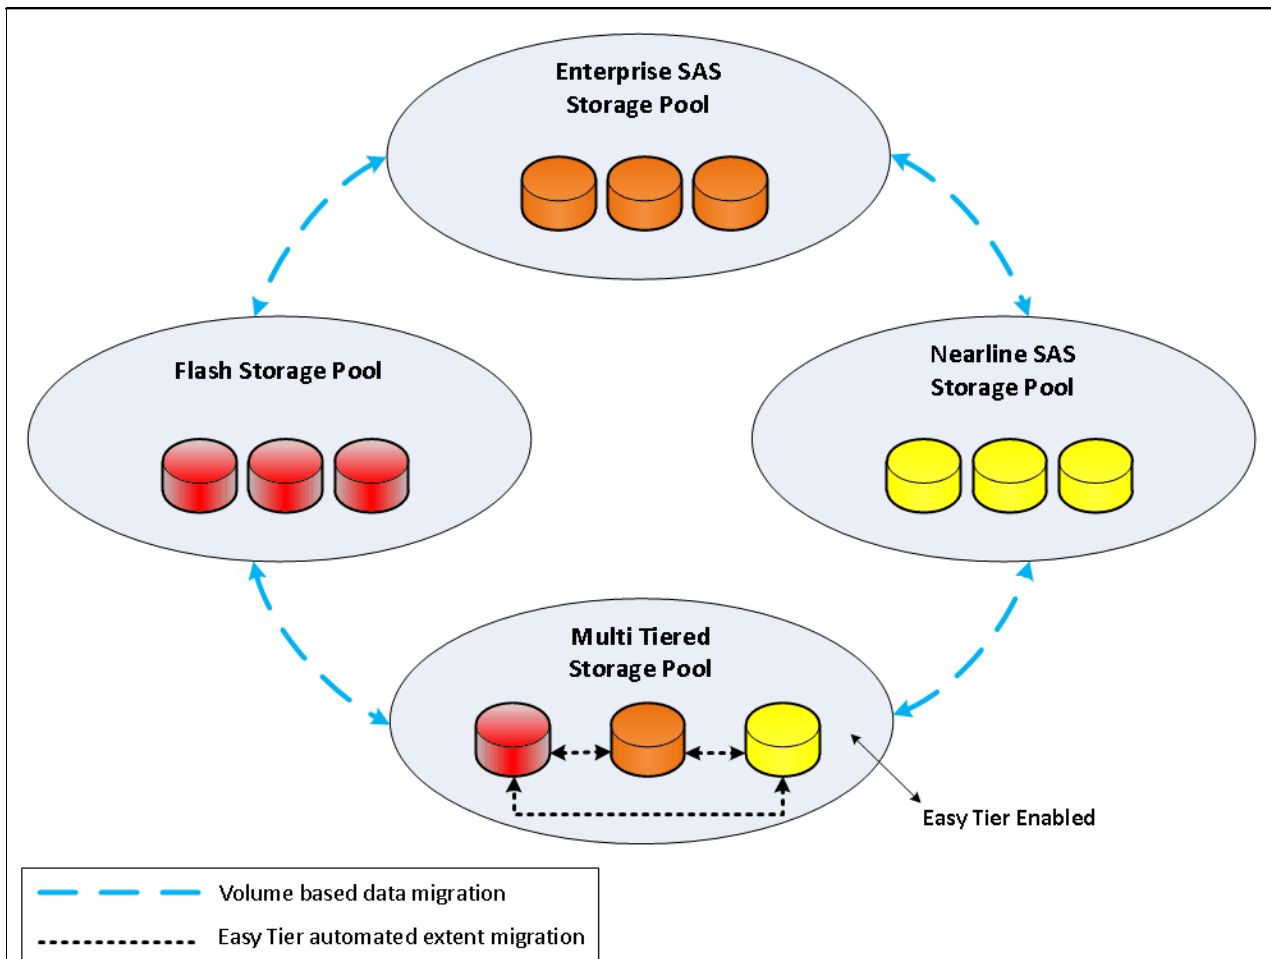


Figure 9-10 Tiered Storage Pools

When licensed, Easy Tier is enabled on any pool that contains two or more classes of disk drive. The Easy Tier function, manages the extent migration as follows:

- ▶ Promote:
 - Moves the candidate hot extent to a higher performance tier.
- ▶ Warm Demote:
 - Prevents performance overload of a tier by demoting a warm extent to a lower tier.
 - Triggered when bandwidth or IOPS exceeds predefined threshold.
- ▶ Cold Demote:
 - Coldest extent moves to a lower tier.
- ▶ Expanded or Cold Demote:
 - Demote appropriate sequential workload to the lowest tier to better use Nearline bandwidth.
- ▶ Swap:
 - This operation exchanges a cold extent in a higher tier with a hot extent in a lower tier, or vice versa.

Note: Extent migrations occur only between adjacent tiers.

Figure 9-11 shows the Easy Tier extent migration.

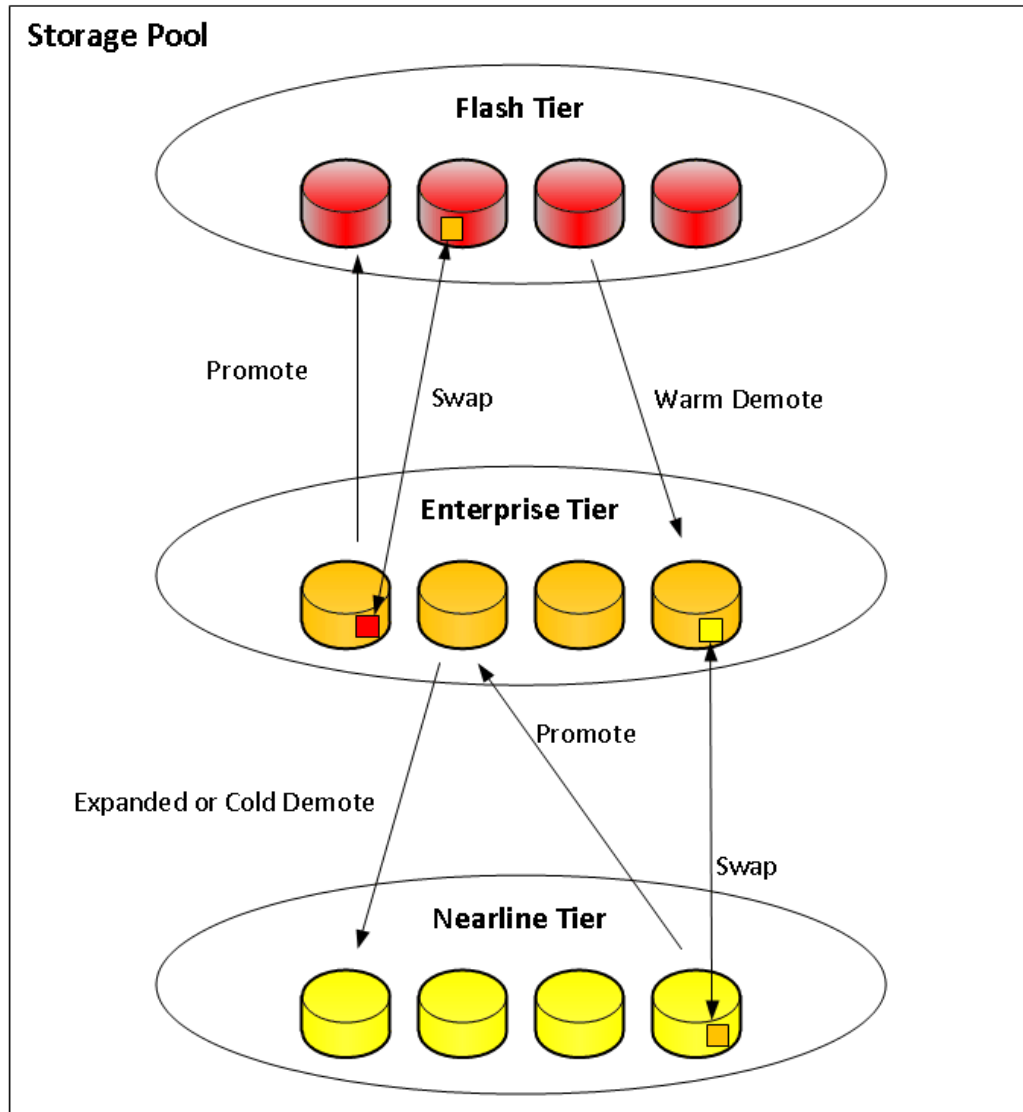


Figure 9-11 Easy Tier process

9.4 Easy Tier process

Easy Tier is based on an algorithm with a threshold to evaluate if an extent is cold, warm, or hot, and consists of four main processes. These processes ensure that the extent allocation in multitiered Storage Pools is optimized for the best performance, based on your workload in the last 24 hours. The processes are described in the following list:

- ▶ I/O Monitoring
- ▶ Data Placement Advisor
- ▶ Data Migration Planner
- ▶ Data Migrator

Figure 9-12 shows the flow between these processes.




Figure 9-12 Easy Tier process flow

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

9.4.1 I/O Monitoring

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

Easy Tier makes allowances for large block I/Os, and therefore considers only I/Os up to 64 kilobytes (KB) as migration candidates.

This is an efficient process, and adds negligible processing effect to the IBM Storwize V3700 node canisters.

9.4.2 Data Placement Advisor

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

This process also identifies extents that must be migrated back to a lower tier.

9.4.3 Data Migration Planner

Using the previously identified extents, the Data Migration Planner (DMP) process builds the extent migration plan for the Storage Pool.

9.4.4 Data Migrator

The Data Migrator (DM) process involves scheduling and the actual movement, or migration, of the volume's extents up to, or down from, the high disk tier. The extent migration rate is capped to a maximum of up to 15 MBps.

This rate equates to around two terabytes (TB) per day being migrated between disk tiers, as shown in Figure 9-13 on page 427.

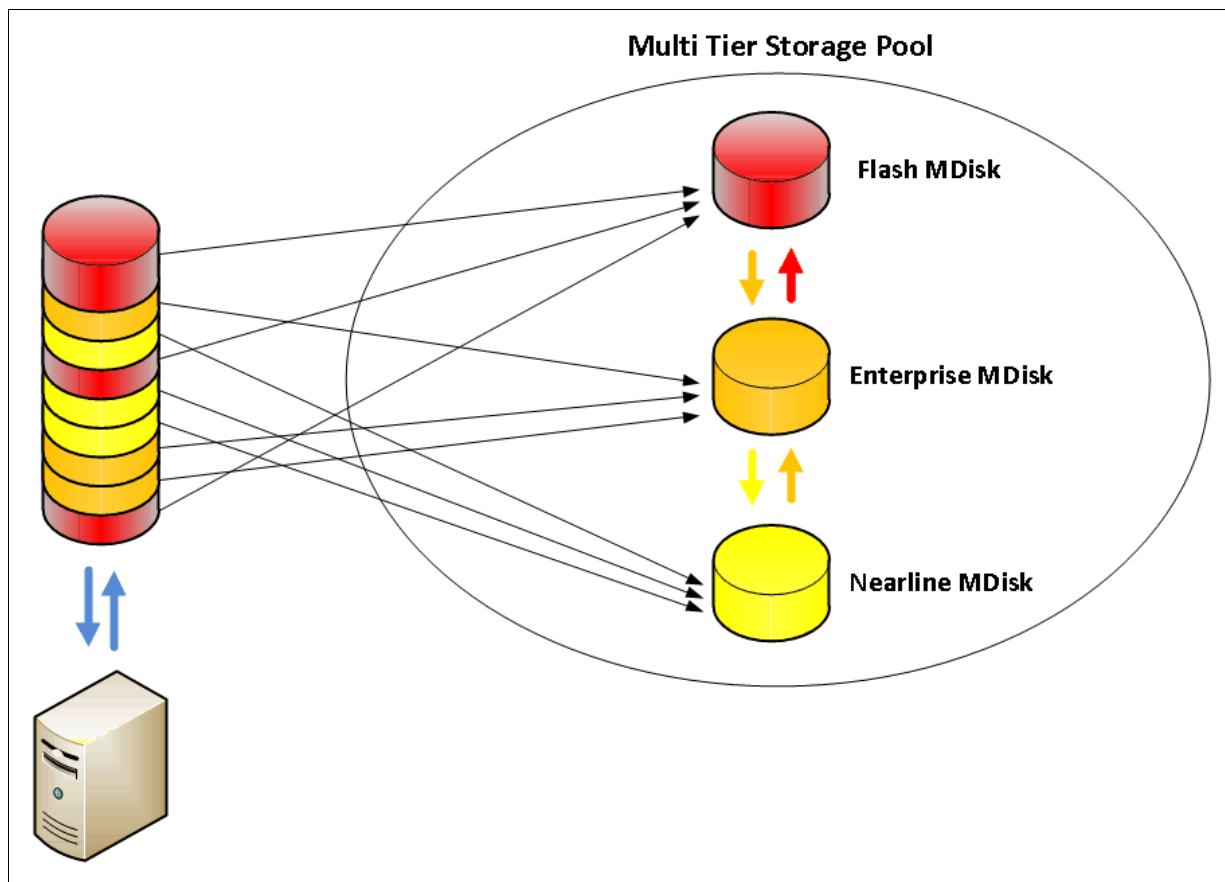


Figure 9-13 Easy Tier Data Migrator

9.4.5 Easy Tier operating modes

IBM Storwize V3700 offers the following operating modes for Easy Tier:

- ▶ Easy Tier: Off

Easy Tier can be turned off. No statistics are recorded, and no extents are moved.

- ▶ Easy Tier: On

When the Easy Tier function is turned on, Easy Tier measures the I/O activity for all extents. When you have a multitiered pool, the extents are migrated dynamically by the Easy Tier processes to achieve the best performance. The movement is not apparent to the host server and applications.

A statistic summary file is created, and can be offloaded and analyzed with the IBM STAT, as described in 9.7, "IBM Storage Tier Advisor Tool" on page 447. Easy Tier can be turned on for any single tier or multitiered pool.

- ▶ Easy Tier: Measured

When Easy Tier is in measured mode, it means Easy Tier measures the I/O activity for all extents, but does not move any extents in the Storage Pool. A statistic summary file is created and can be offloaded from the IBM Storwize V3700. This file can be analyzed with the IBM STAT. This analysis shows the benefits to your workload if you were to add or remove a different drive class to or from your pool before any actual hardware is acquired.

Note: When the Easy Tier trial license 90 days period is finished, running Easy Tier measured mode will require a full license.

► Easy Tier: Auto

This is the operating mode default. When you have Easy Tier set to auto on a single-tiered Storage Pool, Easy Tier is disabled for all volumes inside the Storage Pool, and no extents are moved. If you have Easy Tier set to auto on a multitiered Storage Pool, Easy Tier becomes active for all volumes inside the Storage Pool, and the extents are migrated dynamically by the EasyTier process. This scenario is similar to Easy Tier being set to on, but the extents are not migrated across the tiers if the Easy Tier function is not licensed.

9.4.6 Easy Tier status

Depending on Easy Tier mode attributes, a Storage Pool can have one of the following Easy Tier statuses:

- Active. This status indicates that Easy Tier is actively managing the extents of the Storage Pool.
- Balanced. This status applies to homogeneous Storage Pools, and indicates that Easy Tier is actively managing the extents to provide enhanced performance by rebalancing the extents among the MDisks within the tier. This re-balancing characteristic is called Storage Pool Balancing, which is described in 9.4.7, “Storage Pool Balancing”.
- Measured. This status means that Easy Tier is constantly measuring the I/O activity for all extents to generate an I/O statistics report.
- Inactive. When the Easy Tier status is inactive, no extents are being monitored, and no statistics are being recorded.

9.4.7 Storage Pool Balancing

Storage Pool Balancing is a new function in Storwize software V7.4 or higher that, although associated with Easy Tier, operates independently and does not require a license. Storage Pool Balancing works with Easy Tier when multiple tiers exist in a single pool.

Storage Pool Balancing algorithms assess the extents in a storage tier and balance them automatically across all MDisks within that tier. Storage Pool Balancing moves the extents to achieve a balanced workload distribution and avoid hotspots. Storage Pool Balancing is an algorithm based on MDisk IOPS usage, which means that it is not capacity-based but performance-based. It works on a six-hour performance window.

When a new MDisk is added into an existing Storage Pool, Storage Pool Balancing automatically balances the extents across all MDisks.

Figure 9-14 represents an example of Storage Pool Balancing.

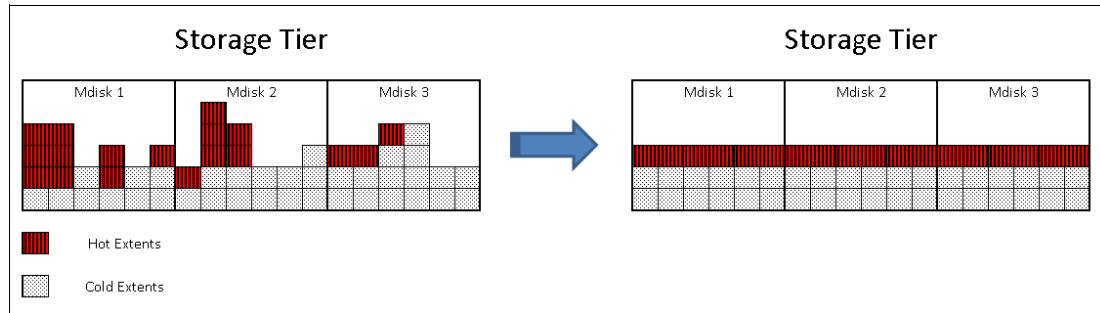


Figure 9-14 Storage Pool Balancing

9.4.8 Easy Tier rules

The following operating rules apply when IBM System Storage Easy Tier is used on the IBM Storwize V3700:

- ▶ Automatic data placement and extent I/O activity monitors are supported on each copy of a mirrored volume. Easy Tier works with each copy independently of each other.

Volume mirroring: Volume mirroring can have different workload characteristics for each copy of the data, because reads are normally directed to the primary copy, and writes occur to both. Therefore, the number of extents that Easy Tier migrates is probably different for each copy.

- ▶ Easy Tier works with all striped volumes, including the following types of volumes:
 - Generic volumes
 - Thin-provisioned volumes
 - Mirrored volumes
 - Thin-mirrored volumes
 - Global Mirror and Metro Mirror sources and targets
- ▶ Easy Tier automatic data placement is not supported for image mode or sequential volumes. I/O monitoring for such volumes is supported, but you cannot migrate extents on such volumes unless you convert image or sequential volume copies to striped volumes.
- ▶ IBM Storwize V3700 creates volumes or volume expansions using extents from MDisks from the Enterprise and Nearline tier. Extents from MDisks in the flash tier are used if no Enterprise or Nearline space is available.
- ▶ When a volume is migrated out of a Storage Pool managed with Easy Tier, automatic data placement mode is no longer active on that volume. Automatic data placement is also turned off while a volume is migrated, even if it is between pools that both have Easy Tier automatic data placement enabled. Automatic data placement for the volume is re-enabled when the migration is complete.
- ▶ Flash drive performance depends on block size (small blocks perform much better than larger blocks). Easy Tier measures I/O blocks smaller than 64 KB, but it migrates the entire extent to the appropriate disk tier.
- ▶ As extents are migrated, the use of smaller extents makes Easy Tier more efficient.
- ▶ The first migration starts about one hour after automatic data placement mode is enabled. It takes up to 24 hours to achieve optimal performance.

- ▶ In the current IBM Storwize V3700 Easy Tier implementation, it takes about two days before hot spots are considered moved from tier to tier, which prevents hot spots from being moved from a fast tier if the workload changes over a weekend.
- ▶ If you run an unusual workload over a longer period, automatic data placement can be turned off and on online to avoid data movement.

Depending on which Storage Pool and which Easy Tier configuration is set, a volume copy can have the Easy Tier states that are shown in Table 9-1.

Table 9-1 Easy Tier states

Storage Pool	Single-tiered or multitiered Storage Pool	Volume copy Easy Tier setting	Easy Tier status on volume copy
Off	Single-tiered	Off	Inactive
Off	Single-tiered	On	Inactive
Off	Multitiered	Off	Inactive
Off	Multitiered	On	Inactive
Auto ^a	Single-tiered	Off	Measured ^b
Auto ^a	Single-tiered	On	Balanced ^e
Auto ^a	Multitiered	Off	Measured ^b
Auto ^a	Multitiered	On	Active ^{c d}
On	Single-tiered	Off	Measured ^b
On	Single-tiered	On	Balanced ^e
On	Multitiered	Off	Measured ^b
On	Multitiered	On	Active ^c
Measure	Single-tiered	On	Measured ^b
Measure	Single-tiered	Off	Measured ^b
Measure	Multitiered	On	Measured ^b
Measure	Multitiered	Off	Measured ^b

- a. The default Easy Tier setting for a Storage Pool is Auto, and the default Easy Tier setting for a volume copy is On. This scenario means that Easy Tier functions are disabled for Storage Pools with a single tier.
- b. When the volume copy status is measured, the Easy Tier function collects usage statistics for the volume, but automatic data placement is not active.
- c. If the volume copy is in image or sequential mode or is being migrated, the Easy Tier status of the volume copy is measured rather than active.
- d. When the volume copy status is active, the Easy Tier function operates in automatic data placement mode for that volume.
- e. When the volume Easy Tier status is Balanced, the Easy Tier is actively managing the extents by rebalancing them among the MDisk within the tier.

9.5 Easy Tier configuration using the GUI

This section describes how to activate Easy Tier to a Storage Pool using the IBM Storwize V3700.

9.5.1 Creating multitiered pools: Enabling Easy Tier

This section describes how to create multitiered Storage Pools using the GUI. When a Storage Pool changes from single-tiered to multitiered, Easy Tier is enabled by default for the pool, and on all volume copies inside this pool.

In this example, you create a pool containing Enterprise and Nearline MDisks.

To create a multitiered pool, complete the following steps:

1. Click **Pools** → **Internal Storage**. Figure 9-15 shows all of the internal drive classes installed in the IBM Storwize V3700.

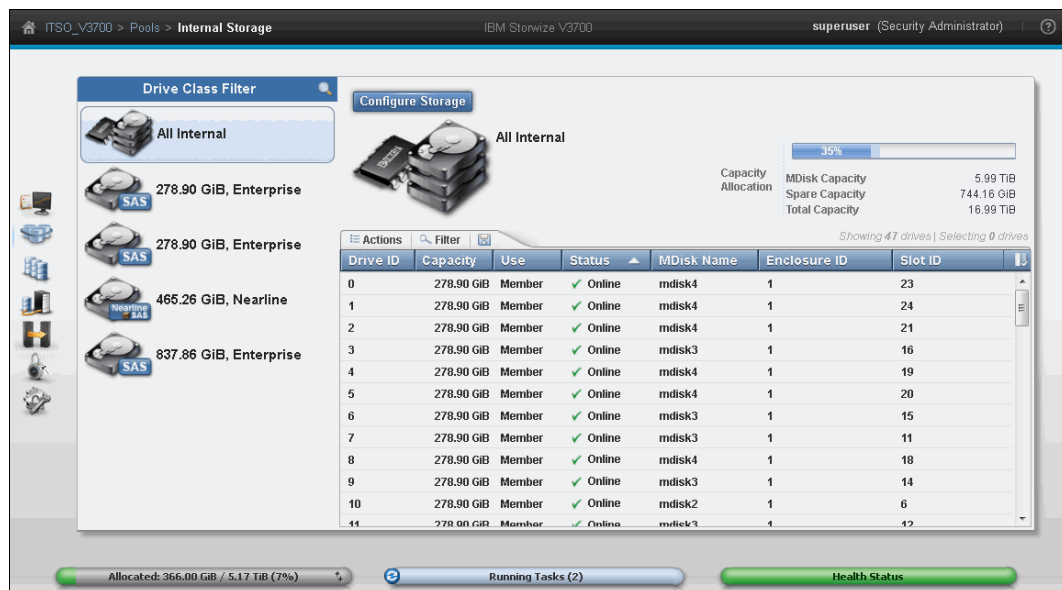


Figure 9-15 Creating Easy Tier Pool

2. Select the class of drive that you want to include in the multitiered pool and click **Configure Storage**. In this scenario, we provision one Redundant Array of Independent Disks (RAID) 10 MDisk array using the Enterprise drive class, as shown in Figure 9-16. Click **Next** to continue.

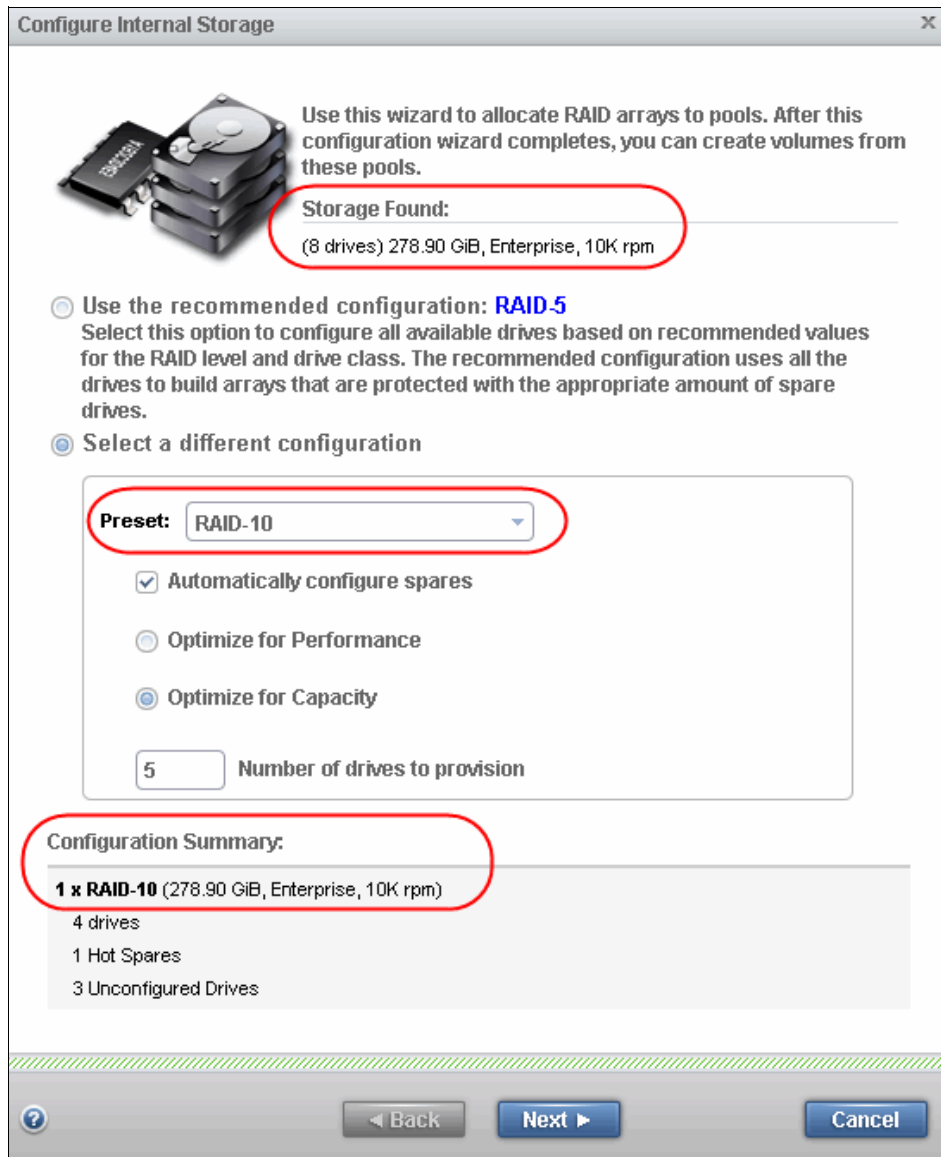


Figure 9-16 Creating a new MDisk

3. Select **Create one or more new pools**, and include the pool name or prefix. In this example scenario, we create a new Storage Pool named **EasyTier_Pool**, as shown in Figure 9-17.

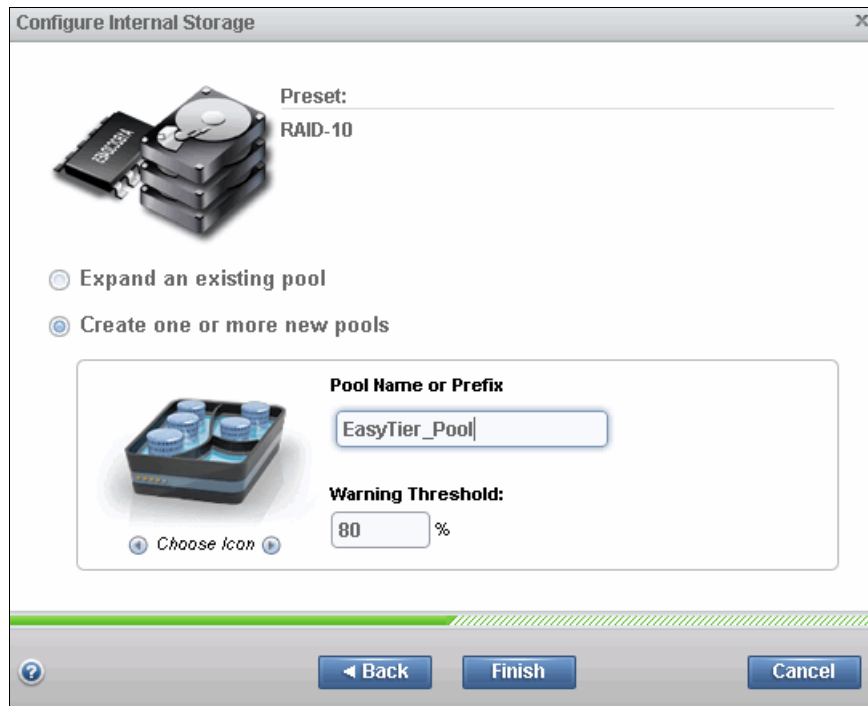


Figure 9-17 Creating a new pool

4. Click **Finish**.

Using the IBM Storwize V3700 GUI, you can check the properties of the **EasyTier_Pool**. Because the Storage Pool contains only one Enterprise MDisk type, the Easy Tier status remains *Inactive*, as shown in Figure 9-18.

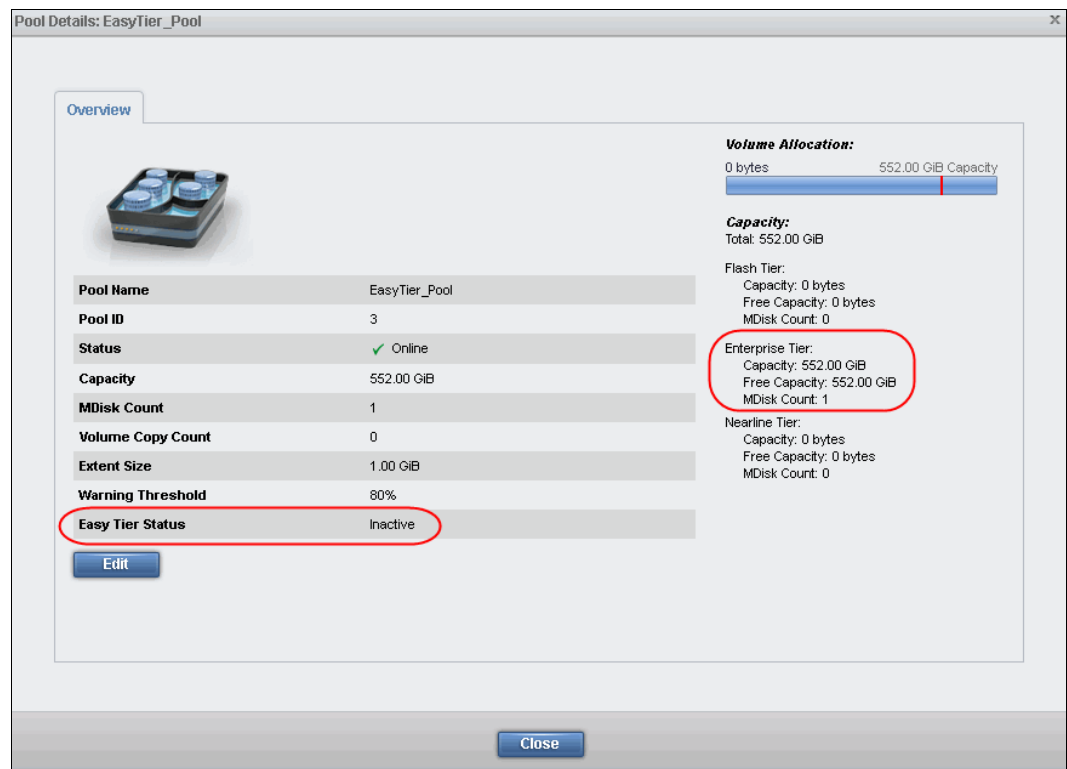


Figure 9-18 Storage Pool properties

5. From the main page, click **Pools** → **Internal Storage** and repeat step 1 on page 432 and step 2 on page 433 to create a further disk class MDisk using the available storage. In this example, we used Nearline SAS.

- As shown in Figure 9-19, select **Expand an existing pool** and select the pool that you want to change to a multitiered pool. In this example, EasyTier_Pool1 is selected.

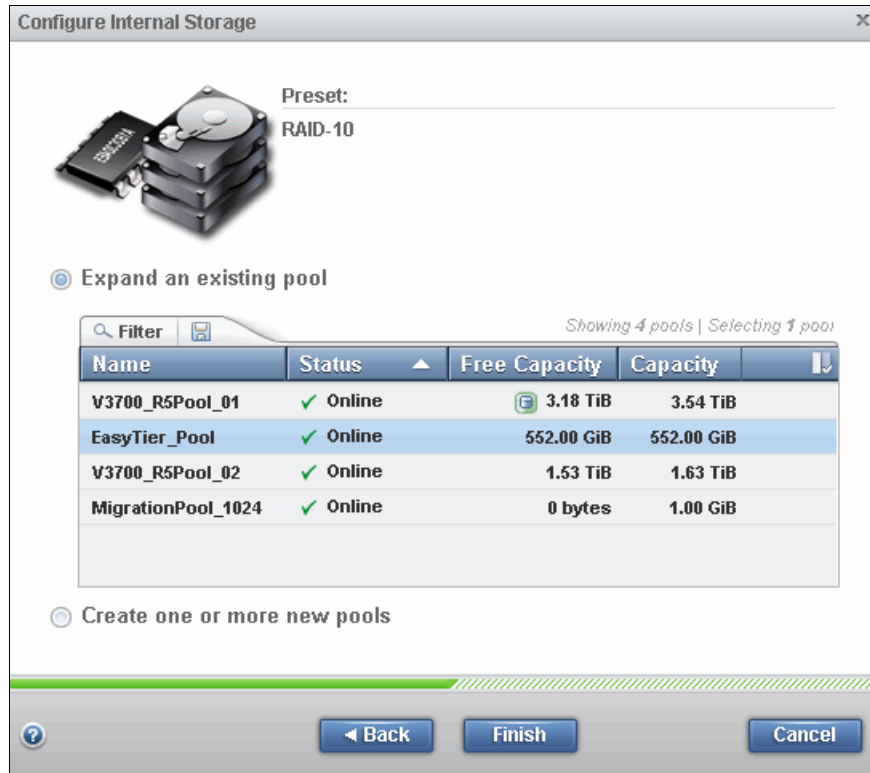


Figure 9-19 Configure Internal Storage window

- Click **Finish** and **Close** when the task completes. Add further MDisk and classes by repeating the previous steps.
- In Figure 9-20, you see that the EasyTier_Pool Storage Pool contains Enterprise MDisk and Nearline MDisk types.




Figure 9-20 Easy Tier active

As soon as a second MDisk of a different disk class is added, the status of Easy Tier changes from Inactive to Active. In this pool, IOM has started and the other Easy Tier processes start to work.

9. Figure 9-21 shows the Storage Pool changed to a multitiered Storage Pool containing Enterprise and Nearline tiers. The Storage Pool icon changes from a single-tiered to a multitiered characteristic.

Note: When Easy Tier functionality is licensed, the IBM Storwize V3700 changes the Easy Tier status of the Storage Pool from Inactive to Active automatically.

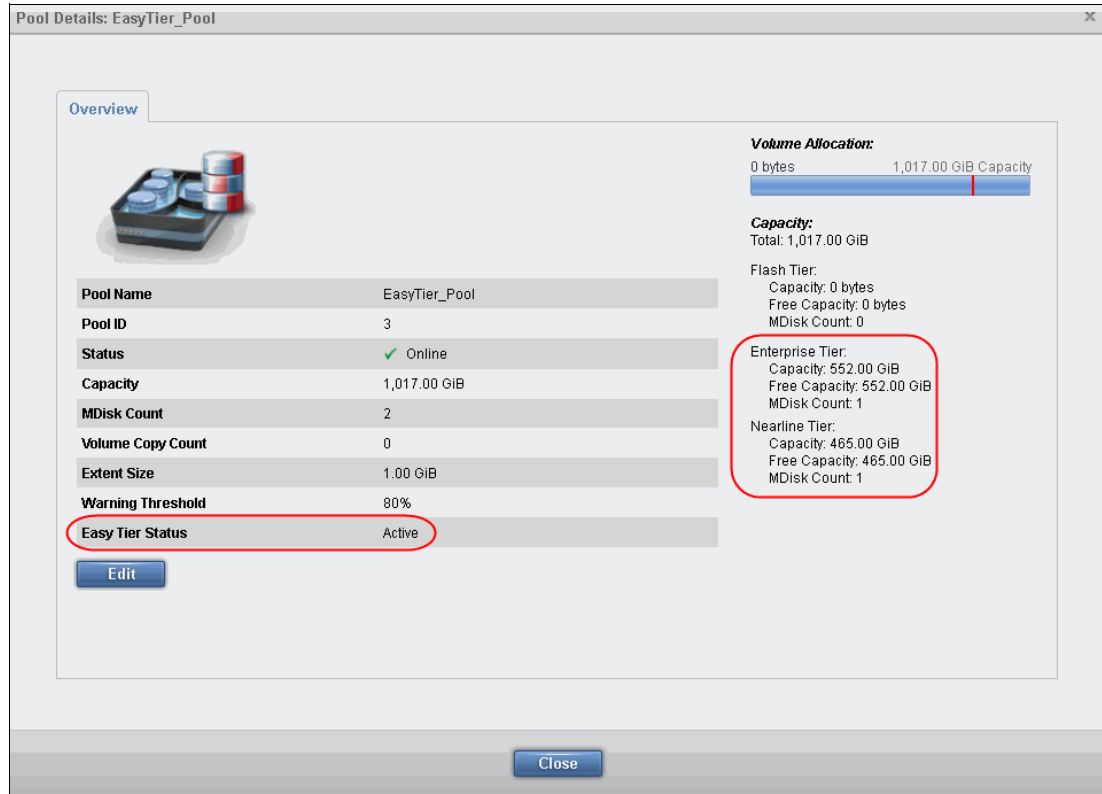


Figure 9-21 Multitiered Storage Pool

Figure 9-22 shows three volumes on the multitiered Storage Pool.

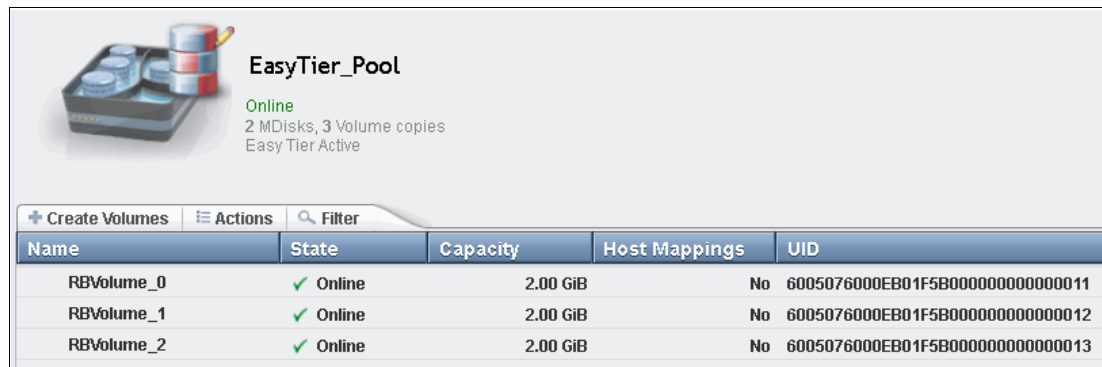


Figure 9-22 Volumes by Pool

10. If you open the properties of a volume by clicking **Actions** → **Properties**, you can also see that Easy Tier is enabled on the volume, as shown in Figure 9-23. Volumes inherit the Easy Tier state of their parent pool.

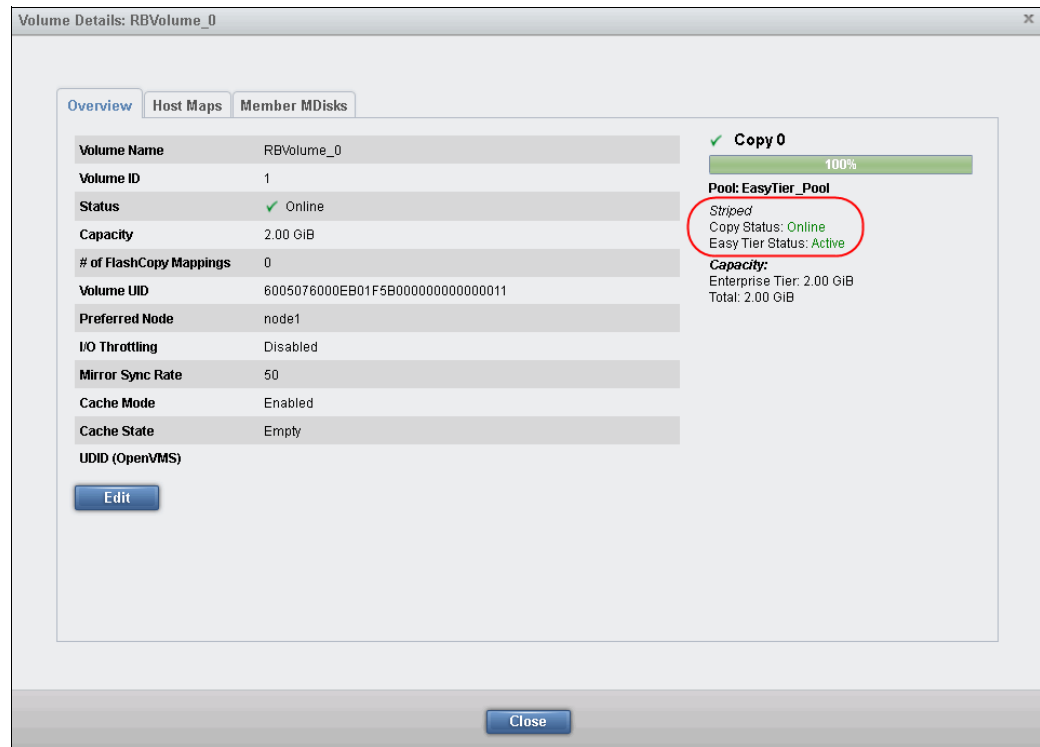


Figure 9-23 Easy Tier enabled volume

11. If a volume has more than one copy, Easy Tier can be enabled and disabled on each copy separately. This action depends on the Storage Pool where the volume copy is defined. You can see a volume with two copies that are stored in two different Storage Pools, as shown in Figure 9-24.




Figure 9-24 Easy Tier by Copy

12. If you want to enable Easy Tier on the second copy, change the Storage Pool of the second copy to a multitiered Storage Pool by repeating the steps in this section.

When external MDisks are presented to an IBM Storwize V3700, you can select the tier manually, and then add it to a Storage Pool for migration purposes, as described in Chapter 6, “Storage migration” on page 249.

9.5.2 Downloading Easy Tier I/O measurements

Easy Tier is now enabled, and automatic data placement mode is active. Extents are automatically migrated to or from disk tiers, and the statistic summary collection is now active. The statistics log file can be downloaded to analyze how the extents were migrated, and to monitor if it makes sense to add more flash to the multitiered Storage Pool.

Heat data files are produced approximately once a day (that is, roughly every 24 hours) when Easy Tier is active on one or more Storage Pools.

To download the statistics file, complete the following steps:

1. Click **Settings** → **Support**, as shown in Figure 9-25.




Figure 9-25 Settings menu

2. Click **Show full log listing**, as shown in Figure 9-26.

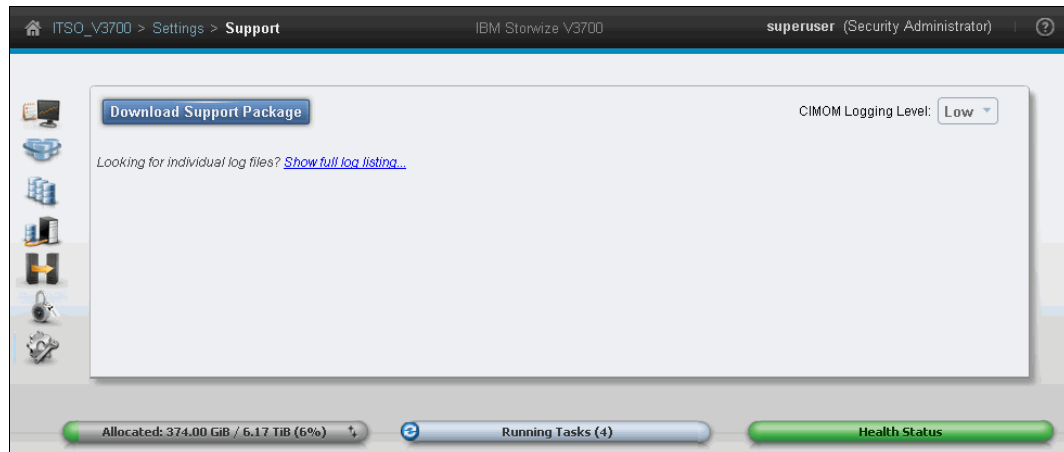


Figure 9-26 Download files menu

This action lists all of the log files available to download. The Easy Tier log files are always named `dpa_heat.<canister_name>_<date>.<time>.data`.

If you run Easy Tier for a longer period, it generates a heat file at least every 24 hours. The time and date that the file was created are included in the file name.

- To download the statistics file, select the file for the most representative period and click **Actions** → **Download** as shown in Figure 9-27. Usually this is the latest file.




Figure 9-27 Download dpa_heat file

Log file creation: Depending on your workload and configuration, it can take up to 24 hours until a new Easy Tier log file is created.

- You can also use the search field on the right to filter your search, as shown in Figure 9-28.




Figure 9-28 Filter your search

- Depending on your browser settings, the file is downloaded to your default location, or you are prompted to save it to your computer. This file can be analyzed as described in 9.7, “IBM Storage Tier Advisor Tool” on page 447.

9.6 Easy Tier configuration using the CLI

The process used to enable IBM Storwize V3700 Easy Tier using the GUI is described in 9.5, “Easy Tier configuration using the GUI” on page 432. Easy Tier can also be configured using the CLI. For the advanced user, this method offers more options for Easy Tier configuration.

Before you use the CLI, you must configure CLI access, as described in Appendix A, “Command-line interface setup and SAN Boot” on page 649.

Readability: In most examples shown in this section, many lines were deleted in the command output or responses so that you can concentrate on the information that is related to Easy Tier.

9.6.1 Enabling Easy Tier measured mode

It is possible to enable Easy Tier in measured mode, on either a single-tiered or multitiered Storage Pool. To do so, complete the following steps:

1. Connect to your IBM Storwize V3700 using the CLI and run the **svcinfo lsmdiskgrp** command, as shown in Example 9-1. This command shows an overview of all configured Storage Pools and their Easy Tier status. In this example, there are two Storage Pools:
 - Nearline_Pool with Easy Tier in auto
 - Multi_Tier_Pool with Easy Tier enabled

Example 9-1 Show all configured Storage Pools

```
IBM_Storwize:ITSO_V3700:superuser>svcinfo lsmdiskgrp
id name                status mdisk_count easy_tier easy_tier_status type
1 Nearline_Pool        online 1             auto    balanced    parent
3 EasyTier_Pool        online 3             on      active      parent
IBM_Storwize:ITSO_V3700:superuser>
```

2. To get a more detailed view of the single-tiered Storage Pool, run the **svcinfo lsmdiskgrp <Storage Pool name>** command, as shown in Example 9-2.

Example 9-2 Storage Pools details: Easy Tier inactive

```
IBM_Storwize:ITSO_V3700:superuser>svcinfo lsmdiskgrp Nearline_Pool
id 1
name Nearline_Pool
status online
mdisk_count 1
.
.
.
easy_tier auto
easy_tier_status balanced
tier ssd
tier_mdisk_count 0
tier_capacity 0.00MB
tier_free_capacity 0.00MB
tier enterprise
tier_mdisk_count 1
tier_capacity 0.00MB
tier_free_capacity 0.00MB
tier nearline
tier_mdisk_count 0
tier_capacity 1.63TB
tier_free_capacity 1.63TB
.
.
.
IBM_Storwize:ITSO_V3700:superuser>
```

- To enable Easy Tier on a single-tiered Storage Pool in measured mode, run the **chmdiskgrp -easytier measure <Storage Pool name>** command, as shown in Example 9-3.

Example 9-3 Enable Easy Tier on a single-tiered Storage Pool

```
IBM_Storwize:ITS0_V3700:superuser>chmdiskgrp -easytier measure Nearline_Pool
IBM_Storwize:ITS0_V3700:superuser>
```

- Check the status of the Storage Pool again by running the **lsmdiskgrp <Storage Pool name>** command again, as shown in Example 9-4.

Example 9-4 Storage Pool details: Easy Tier ON

```
IBM_Storwize:ITS0_V3700:superuser>lsmdiskgrp Nearline_Pool
id 1
name Nearline_Pool
status online
mdisk_count 1
...
easy_tier measure
easy_tier_status measured
tier ssd
tier_mdisk_count 0
tier_capacity 0.00MB
tier_free_capacity 0.00MB
tier enterprise
tier_mdisk_count 1
tier_capacity 0.00MB
tier_free_capacity 0.00MB
tier nearline
tier_mdisk_count 0
tier_capacity 1.63TB
tier_free_capacity 1.63TB
...
IBM_Storwize:ITS0_V3700:superuser>
```

- To get a list of all of the defined volumes, run the **lsvdisk** command, as shown in Example 9-5. For this example, we are only interested in the RedbookVolume volume.

Example 9-5 All volumes list

```
IBM_Storwize:ITS0_V3700:superuser>lsvdisk
id name          IO_group_id IO_group_name status mdisk_grp_id
...
6 RedbookVolume 0          io_grp0      online 1
...
IBM_Storwize:ITS0_V3700:superuser>
```

- To get a more detailed view of a volume, run the **lsvdisk volume_name** command, as shown in Example 9-6 on page 444. This output shows two copies of a volume. Both copies now have Easy Tier turned on.

However, Copy 0 is in a multitiered Storage Pool and automatic data placement is active, as indicated by the `easy_tier_status` line. Copy 1 is in the single-tiered Storage Pool, and Easy Tier mode is measured, as indicated by the `easy_tier_status` line.

Example 9-6 Volume details

```
IBM_Storwize:ITSO_V3700:superuser>lsvdisk RedbookVolume
id 6
name RedbookVolume
status online
capacity 1.00GB
copy_count 2
copy_id 0
status online
mdisk_grp_name EasyTier_Pool
type striped
mdisk_id
mdisk_name
fast_write_state empty
used_capacity 1.00GB
real_capacity 1.00GB
free_capacity 0.00MB
overallocation 100
easy_tier on
easy_tier_status active
tier ssd
tier_capacity 0.00MB
tier enterprise
tier_capacity 1.00GB
tier nearline
tier_capacity 0.00MB
parent_mdisk_grp_name EasyTier_Pool

copy_id 1
status online
mdisk_grp_name Nearline_Pool
type striped
mdisk_id
mdisk_name
fast_write_state not_empty
used_capacity 1.00GB
real_capacity 1.00GB
free_capacity 0.00MB
overallocation 100
easy_tier on
easy_tier_status measured
tier ssd
tier_capacity 0.00MB
tier enterprise
tier_capacity 1.00GB
tier nearline
tier_capacity 0.00MB
parent_mdisk_grp_name Nearline_Pool
IBM_Storwize:ITSO_V3700:superuser>
```

Easy Tier measured mode does not carry out any data placement. It only collects statistics for measurement. For more information about downloading and analyzing the I/O statistics, see 9.5.2, “Downloading Easy Tier I/O measurements” on page 439.

- These changes are also reflected in the GUI, as shown in Figure 9-29. Click **Pools** → **Volumes by Pool**, select Nearline_Pool, and find the volume RedbookVolume.
- Then, click **Actions** → **Properties** to view the details of the Easy Tier for each of the volume copies.

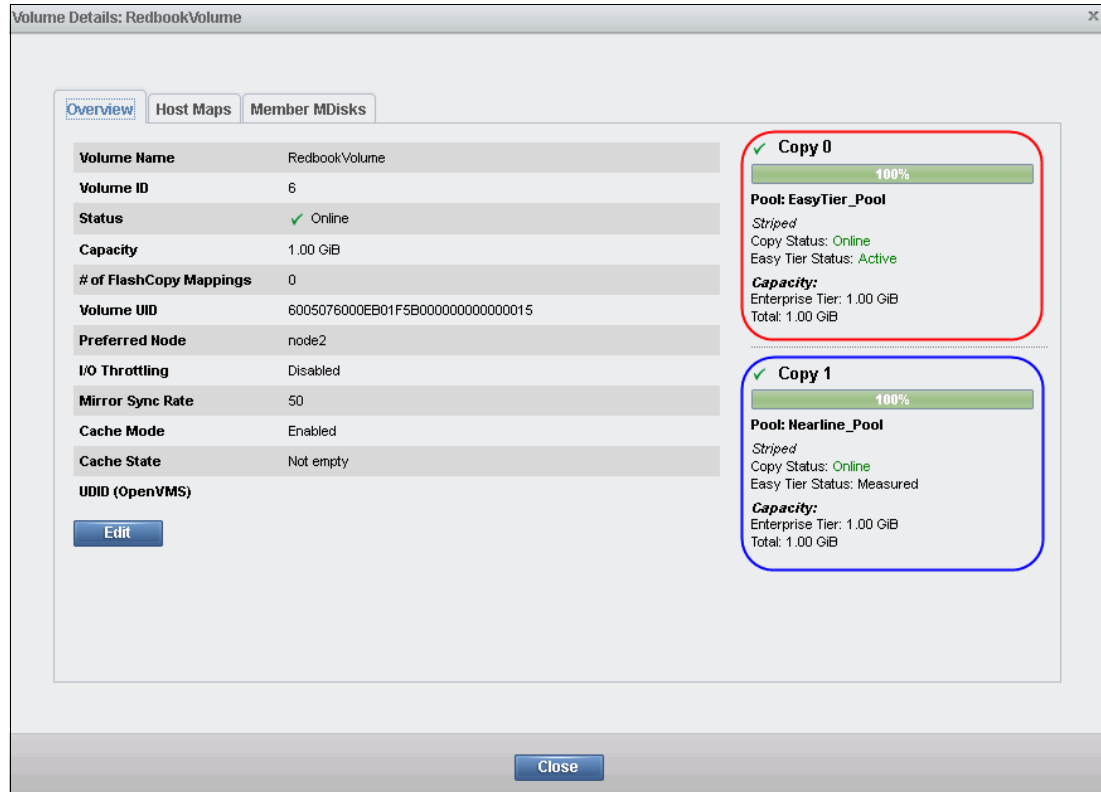


Figure 9-29 Volume properties for Easy Tier

9.6.2 Enabling or disabling Easy Tier on single volumes

If you enable Easy Tier on a Storage Pool, all volume copies inside the Easy Tier pools also have Easy Tier enabled by default. This setting applies to multitiered and single-tiered Storage Pools. It is also possible to turn Easy Tier on and off for single volume copies.

To disable Easy Tier, complete the following steps:

- Before disabling Easy Tier on a volume, run `svcinfo lsmdisgrp <Storage Pool name>` command to show all configured Storage Pools, as shown in Example 9-7. In this example, EasyTier_Pool is the Storage Pool used as a reference.

Example 9-7 Listing the Storage Pool

```
IBM_Storwize:ITSO_V3700:superuser>svcinfo lsmdisgrp
id name                status easy_tier easy_tier_status
.
3 EasyTier_Pool        online on          active
.
.
IBM_Storwize:ITSO_V3700:superuser>
```

2. Run the **svcinfolsvdisk** to show all configured volumes within your IBM Storwize V3700 as shown in Example 9-8. We are interested only in the volume called Volume001.

Example 9-8 Show all configured volumes

```
IBM_Storwize:ITS0_V3700:superuser>svcinfolsvdisk
id name                IO_group_id status mdisk_grp_name  capacity
0  Volume001           0          online EasyTier_Pool   5.00GB
IBM_Storwize:ITS0_V3700:superuser>
```

3. To disable Easy Tier on single volumes, run the **svctask chvdisk -easytier off <volume name>** command, as shown in Example 9-9.

Example 9-9 Disable Easy Tier on a single volume

```
IBM_Storwize:ITS0_V3700:superuser>svctask chvdisk -easytier off Volume001
IBM_Storwize:ITS0_V3700:superuser>
```

This command disables Easy Tier on all copies of the volume. Example 9-10 shows Easy Tier turned off for Copy 0 even if Easy Tier is still enabled on the Storage Pool. Note that on Copy 0, the status changed to measured, because the pool is still actively measuring the I/O on the volume.

Example 9-10 Easy Tier disabled

```
IBM_Storwize:ITS0_V3700:superuser>lsvdisk Volume001
id 0
name Volume001
mdisk_grp_name EasyTier_Pool
capacity 5.00GB
parent_mdisk_grp_id 3
parent_mdisk_grp_name EasyTier_Pool

copy_id 0
status online
mdisk_grp_name EasyTier_Pool
used_capacity 5.00GB
real_capacity 5.00GB
free_capacity 0.00MB
overallocation 100
autoexpand
warning
grainsize
se_copy no
easy_tier off
easy_tier_status measured
tier ssd
tier_capacity 0.00MB
tier enterprise
tier_capacity 5.00GB
tier nearline
tier_capacity 0.00MB
compressed_copy no
parent_mdisk_grp_id 3
parent_mdisk_grp_name EasyTier_Pool
IBM_Storwize:ITS0_V3700:superuser>
```

4. To enable Easy Tier on a volume, run the `svctask chvdisk -easytier on <volume name>` command (as shown in Example 9-11), and Easy Tier changes back to on (as shown in Example 9-12). Notice copy 0 status also changed back to active.

Example 9-11 Easy Tier enabled

```
IBM_Storwize:ITS0_V3700:superuser>svctask chvdisk -easytier on Volume001
IBM_Storwize:ITS0_V3700:superuser>
```

Example 9-12 Easy Tier on single volume enabled

```
IBM_Storwize:ITS0_V3700:superuser>svcinfolsvdisk Volume001
id 0
name Volume001
status online
mdisk_grp_name EasyTier_Pool
capacity 5.00GB
parent_mdisk_grp_id 3
parent_mdisk_grp_name EasyTier_Pool

copy_id 0
status online
mdisk_grp_name EasyTier_Pool
used_capacity 5.00GB
real_capacity 5.00GB
free_capacity 0.00MB
overallocation 100
easy_tier on
easy_tier_status active
tier ssd
tier_capacity 1.00GB
tier enterprise
tier_capacity 4.00GB
tier nearline
tier_capacity 0.00MB
compressed_copy no
parent_mdisk_grp_id 1
parent_mdisk_grp_name EasyTier_Pool
IBM_Storwize:ITS0_V3700:superuser>
```

9.7 IBM Storage Tier Advisor Tool

The IBM STAT is a Windows console tool. If you run Easy Tier in measured mode, the tool analyzes the extents and captures I/O profiles to estimate how much benefit you would derive from implementing Easy Tier automatic data placement with SSD MDisks.

If automatic data placement mode is already active, the analysis also includes an overview of migrated hot data, and recommendations about whether you can derive any benefit from adding more flash drives. The output provides a graphical representation of the performance data collected by Easy Tier over a 24-hour operational cycle.

The tool comes as a .iso file, which needs to be extracted to a temporary folder. The installation is initiated by running `setup.exe`.

The STAT Tool can be downloaded from the following website:

<http://www-01.ibm.com/support/docview.wss?uid=ssg1S4000935>

9.7.1 Processing heat log files

STAT takes input from the dpa_heat log file and produces a Hypertext Markup Language (HTML) file that contains the report. Download the heat_log file, as described in 9.5.2, “Downloading Easy Tier I/O measurements” on page 439, and save it to the HDD of a Windows system.

For more information about the tool and to download it, see the following website:

<http://www-01.ibm.com/support/docview.wss?uid=ssg1S4000935>

To process the files, complete the following steps:

1. Click **Start** → **Run**, enter cmd, and then click **OK** to open a command prompt.
2. Typically, the tool is installed in the C:\Program Files(x86)\IBM\STAT directory. Enter the command to generate the report, as shown in Example 9-13:

```
C:\Program Files(x86)\IBM\STAT>STAT.exe -o c:\<directory_where_you_want_the  
output_to_go> c:\location_of_dpa_heat_data_file
```

If you do not specify **-o c:\<directory_where_you_want_the_output_to_go>**, the output goes to the directory where the STAT.exe file is located.

Example 9-13 Generate HTML file

```
C:\EasyTier>STAT.exe -o C:\EasyTier  
C:\StorwizeV3700_Logs\dpa_heat.01010KV-1.101  
209.141106.data
```

```
CMUA00019I The STAT.exe command has completed.
```

```
C:\EasyTier>
```

3. The Storage Tier Advisor Tool creates a set of HTML files. Browse to the directory where you directed the output file, and find the file named index.html. Open the file using your browser to view the report.

Typically, the STAT tool generates the report as described in 9.7.2, “Storage Tier Advisor Tool reports”.

9.7.2 Storage Tier Advisor Tool reports

When you open the `index.html` file of an IBM Storwize V3700 system, a window opens that gives you a complete system summary report, as shown in Figure 9-30.

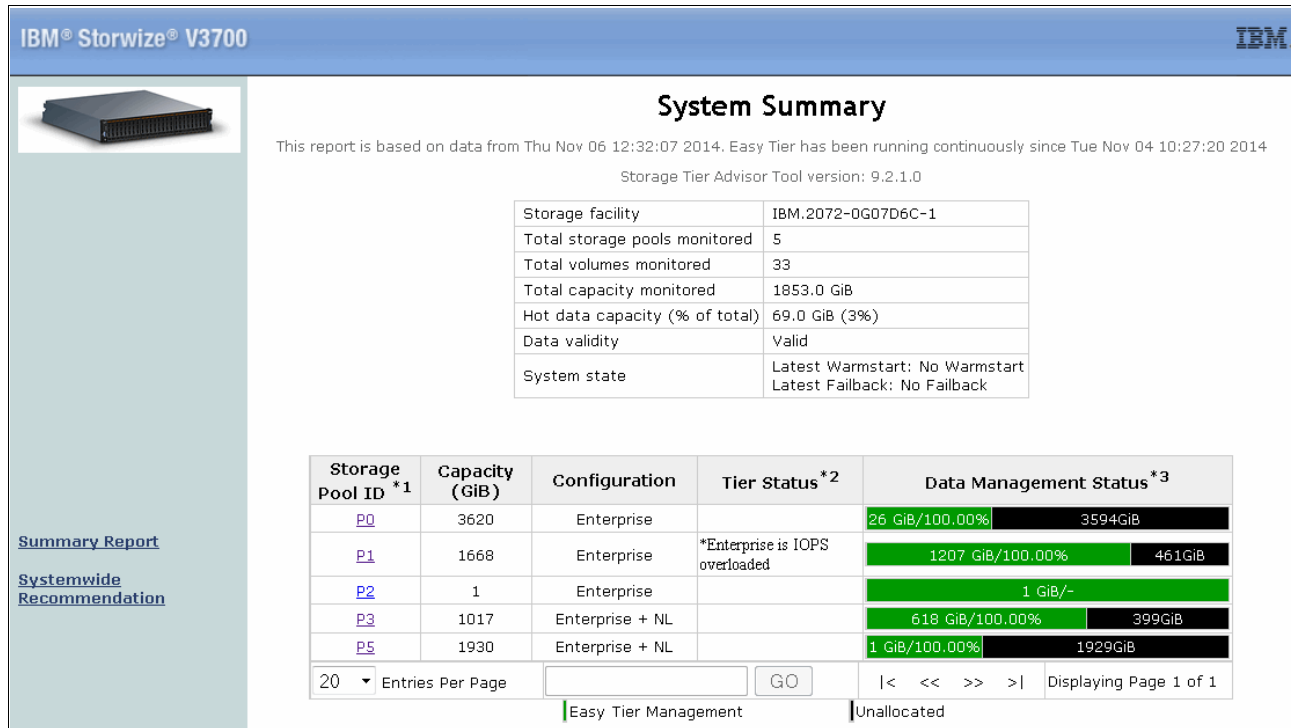


Figure 9-30 STAT report: System Summary

The first page of the report shows the following information:

- ▶ Storage facility
- ▶ Total number of Storage Pools and volumes monitored
- ▶ Total capacity monitored
- ▶ Hot data capacity
- ▶ Data validity
- ▶ System state

As shown in In Figure 9-30, the prediction generated by the STAT is pool-based.

The next table shows how data is managed in the particular Storage Pools, and is denoted by different colors. The total capacity of the Storage Pool is divided into two parts:

- ▶ The data managed by Easy Tier
- ▶ The unallocated space

The green portion of the bar represents the data managed by Easy Tier. The black portion of the bar represents the unallocated space or data that is not being managed by Easy Tier.

In our example, a total of five Storage Pools were monitored with a total capacity of 1853.0 gibibytes (GiB). Therefore, the total of hot data capacity is approximately 3% or 69 GiB, as shown in the system summary page. Click the Storage Pool **P1** as shown in Figure 9-31 on page 450.

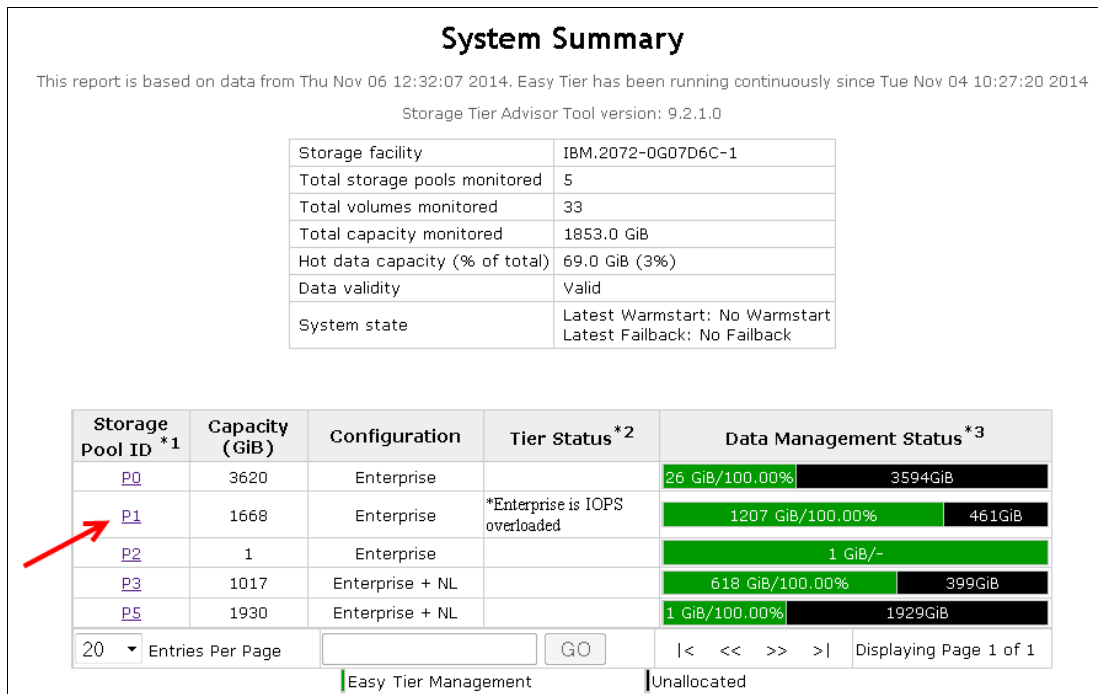


Figure 9-31 Selecting Storage Pool P1

As shown in Figure 9-32, the STAT shows the MDisks that construct the Storage Pool. If the Storage Pool is constructed by more than one MDisk, the STAT shows additional lines for each MDisk.

The tool also shows the performance statistics and the projected IOPS for the MDisks after the Easy Tier processes and rebalancing operations are completed on the extents of the Storage Pool.

For each MDisk, the STAT tool shows the following:

- ▶ The MDisk ID
- ▶ The Storage Pool ID
- ▶ The MDisk type
- ▶ The number of IOPS threshold
- ▶ The utilization of MDisk IOPS
- ▶ The projected utilization of MDisk IOPS

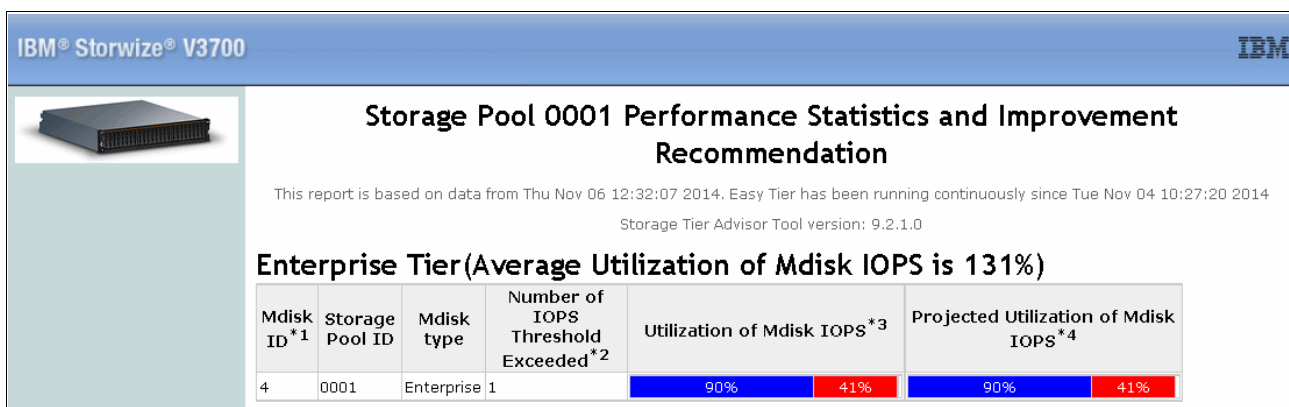


Figure 9-32 Storage Pool performance statistics

The blue portion of the bar represents the percentage range, which is below the tier average utilization of MDisk IOPS.

The red portion of the bar represents percentage range, which is above the maximum allowed IOPS threshold.

In this example, the STAT shows 131% of utilization for the Enterprise Tier, so the IOPS exceeded the threshold by 31%. The blue portion shows 90% out of 131% (approximately 69%) did not exceed the threshold for this particular tier. The red portion shows 41% out of 131% (approximately 31%) of IOPS exceeded the threshold, and the extents are potential candidates to move to a higher tier.

In the **Workload Distribution Across Tier** section, the STAT shows the skew of the workload of the Storage Pool selected. The X-axis (horizontal) denotes the percentage of extents in use. The Y-axis (vertical) denotes the percentage of extents that are busy out of the given percentage from the X-axis. In this graph, for instance, when you look at 10% of the extents in the pool, about 84% of these are determined to be very busy. Figure 9-33 shows an example of this graph.

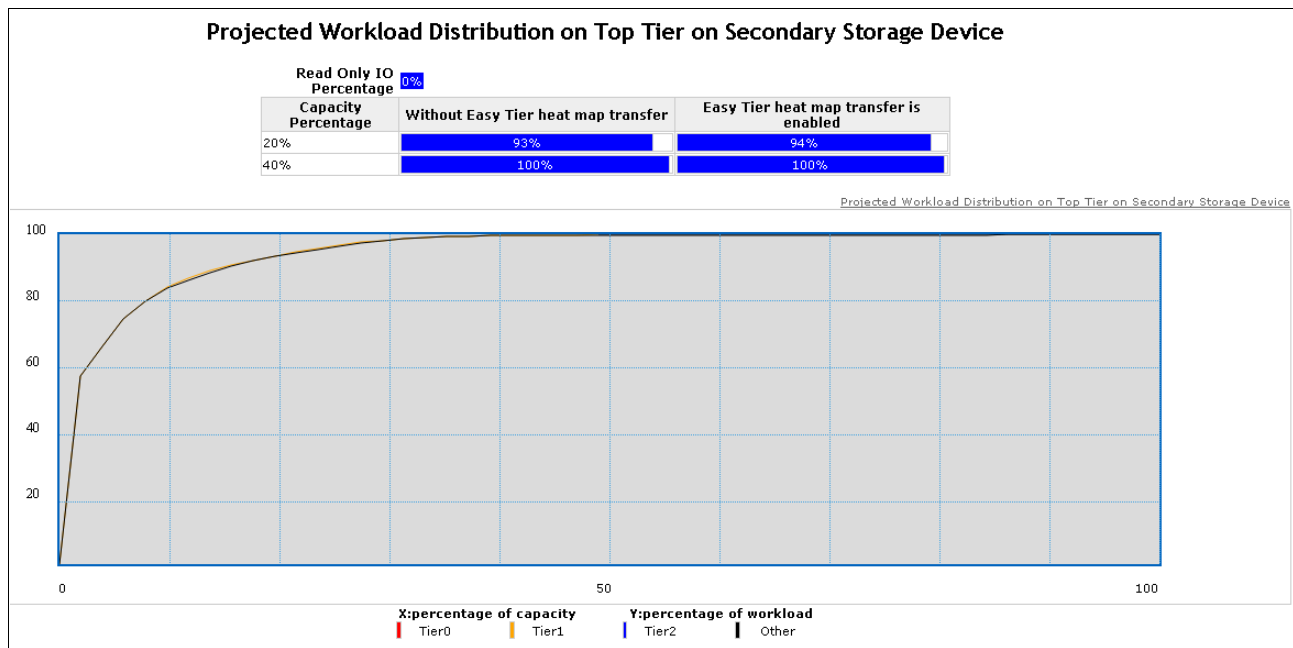


Figure 9-33 STAT Tool project workload distribution

In the **Volume Heat Distribution** section, the STAT shows volume heat distribution of all volumes in the Storage Pool. The following list describes the columns:

- ▶ VDisk ID
- ▶ VDisk Copy ID
- ▶ Configured size or VDisk capacity
- ▶ I/O percentage of extent pool
- ▶ The tiers that the volume is taking capacity from
- ▶ The capacity on each of the tiers
- ▶ The heat distribution of that Storage Pool

The heat distribution of a volume is displayed using a color bar, which represents the following:

- ▶ The blue portion of the bar represents the capacity of *cold extents* in the volume. Extents are considered cold when not used heavily or the IOPS is very low.
- ▶ The orange portion of the bar represents the capacity of *warm extents* in the volume. Data is considered warm when it is either used relatively heavily, or the IOPS is relatively more compared to the cold extents or lower when compared to hot extents.
- ▶ The red portion of the bar represents the capacity of *hot data* in the volume. Data is considered hot when used heavily, or when the IOPS is high.

Figure 9-34 shows a few examples of the volume heat distribution.

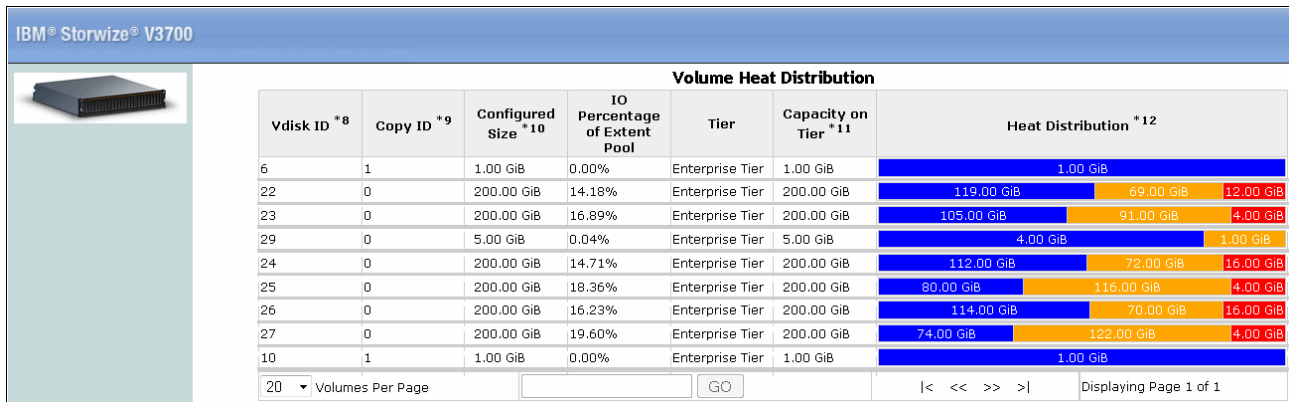



Figure 9-34 Volume heat distribution

The **Systemwide Recommendation** section can be viewed on another page, which shows the recommended configuration for the tiers as applicable for the configuration of the system. Typically, it shows three recommended configurations:

- ▶ Flash (SSD)
- ▶ Enterprise
- ▶ Nearline

Each recommendation displays the Storage Pools, the MDisk recommendation, and the expected improvement. An example is shown in Figure 9-35.

IBM® Storwize® V3700



Systemwide Recommendation

This report is based on data from Thu Nov 06 12:32:07 2014 Easy Tier has been running continuously since Tue Nov 04 10:27:20 2014
Storage Tier Advisor Tool version: 9.2.1.0

Recommended SSD Configuration

Storage Pool ID	SSD Configuration	Predicted Performance Improvement	Total Improvement
0001	Performance Improved by Adding 1 SSD_800G_R10_W8_50K Mdisk(s)	16% ~ 36%	16% ~ 36%
0003	Performance Improved by Adding 1 SSD_800G_R10_W8_50K Mdisk(s)	0% ~ 10%	16% ~ 36%

Recommended Enterprise Configuration

Storage Pool ID	Recommended Enterprise Configuration	Predicted IOPS Improvement	Total Improvement
0003	Existing Enterprise Mdisk(s) Free Capacity (5.0 GiB)	0% ~ 10%	0% ~ 10%

Recommended NL Configuration

Storage Pool ID	Recommended NL Configuration	Cold Data Capacity(GiB)	Total Improvement
0001	Adding 1 7.2K_NL_4000G_R10_W8 Mdisk(s)	610.0	610.0
0003	Adding 1 7.2K_NL_4000G_R10_W8 Mdisk(s)	509.0	1119.0
0000	Adding 1 7.2K_NL_4000G_R10_W8 Mdisk(s)	26.0	1145.0
0002	Adding 1 7.2K_NL_4000G_R10_W8 Mdisk(s)	1.0	1146.0

[Summary Report](#)
[Systemwide Recommendation](#)

Figure 9-35 STAT Tool systemwide recommendation



Copy services

In this chapter, we describe the copy services functions provided by the IBM Storwize V3700 storage system, including IBM FlashCopy and Remote Copy. Copy services functions are useful for making data copies for backup, application test, recovery, and so on. The IBM Storwize V3700 system makes it easy to apply these functions to your environment through its intuitive graphical user interface (GUI).

This chapter includes the following topics:

- ▶ FlashCopy
- ▶ Remote Copy
- ▶ Troubleshooting Remote Copy
- ▶ Managing Remote Copy using the GUI

10.1 FlashCopy

Using the FlashCopy function of the IBM Storwize V3700 storage system, you can create a *point-in-time copy* of one or more volumes. In this section, we describe the structure of FlashCopy and provide details about its configuration and use.

You can use FlashCopy to solve critical and challenging business needs that require the duplication of data on your source volume. Volumes can remain online and active while you create consistent copies of the data sets. Because the copy is performed at the block level, it operates below the host operating system (OS) and cache and, therefore, is not apparent to the host.

Flushing: Because FlashCopy operates at the block level below the host operating system and cache, those levels do need to be flushed for consistent FlashCopy copies.

While the FlashCopy operation is performed, input/output (I/O) to the source volume is disabled briefly to initialize the FlashCopy bitmap, and then is allowed to resume.

Although several FlashCopy options require the data to be copied from the source to the target in the background (which can take time to complete), the resulting data on the target volume copy appears to complete immediately. This task is accomplished by using a bitmap (or bit array) that tracks changes to the data after the FlashCopy is started, and an indirection layer, which enables data to be read from the source volume transparently.

10.1.1 Business requirements for FlashCopy

When you are deciding whether FlashCopy addresses your needs, you must adopt a combined business and technical view of the problems that you must solve. Determine your needs from a business perspective, and then determine whether FlashCopy fulfills the technical needs of those business requirements.

With an immediately available copy of the data, FlashCopy can be used in the following business scenarios:

- ▶ Rapidly creating consistent backups of dynamically changing data

FlashCopy can be used to create backups through periodic running. The FlashCopy target volumes can be used to complete a rapid restore of individual files, or of the entire volume through Reverse FlashCopy (by using the **-restore** option).

The target volumes that are created by FlashCopy can also be used for backup to tape. By attaching them to another server and performing backups from there, it enables the production server to continue largely unaffected. After the copy to tape completes, the target volumes can be discarded, or kept as a rapid restore copy of the data.

- ▶ Rapidly creating consistent copies of production data to facilitate data movement or migration between hosts

FlashCopy can be used to facilitate the movement or migration of data between hosts, while minimizing downtime for applications. FlashCopy enables application data to be copied from source volumes to new target volumes while applications remain online. After the volumes are fully copied and synchronized, the application can be stopped, and then immediately started on the new server accessing the new FlashCopy target volumes.

This mode of migration is faster than other migration methods that are available through the IBM Storwize V3700, because the size and the speed of the migration is not as limited.

- ▶ Rapidly creating copies of production data sets for application development and testing
Under normal circumstances, to perform application development and testing, data must be restored from traditional backup media, such as tape. Depending on the amount of data and the technology in use, this process can easily take a day or more. With FlashCopy, a copy can be created and be online for use in just a few minutes. The time varies based on the application and the data set size.
- ▶ Rapidly creating copies of production data sets for auditing purposes and data mining
Auditing or data mining normally require the usage of the production applications. This situation can cause high loads for databases that track inventories or similar data. With FlashCopy, you can create copies for your reporting and data mining activities. This feature reduces the load on your production systems, which increases their performance.
- ▶ Rapidly creating copies of production data sets for quality assurance
Quality assurance is an interesting case for FlashCopy. Because traditional methods involve so much time and labor, the refresh cycle is typically extended. Because FlashCopy reduces the time that is required, it enables much more frequent refreshes of the quality assurance database.

10.1.2 FlashCopy functional overview

FlashCopy occurs between a source volume and a target volume. The source and target volumes must be the same size. Multiple FlashCopy mappings (source-to-target relationships) can be defined, and point-in-time consistency can be maintained across multiple point-in-time mappings by using *consistency groups*. For more information about FlashCopy consistency groups, see “FlashCopy consistency groups” on page 462.

The minimum granularity that the IBM Storwize V3700 storage system supports for FlashCopy is an entire volume. It is not possible to use FlashCopy to copy only part of a volume. Additionally, the source and target volumes must belong to the same IBM Storwize V3700 storage system, but they do not have to be in the same Storage Pool.

Before you start a FlashCopy (regardless of the type and options that are specified), the IBM Storwize V3700 must put the cache into write-through mode, which flushes the I/O that is bound for the source volume. If you are scripting FlashCopy operations from the command-line interface (CLI), you must run the **prestartfcmap** or **prestartfcconsistgrp** command. However, this step is managed for you and carried out automatically by the GUI.

This is not the same as flushing the host cache, which is not required. After FlashCopy is started, an effective copy of a source volume to a target volume is created. The content of the source volume is immediately presented on the target volume, and the original content of the target volume is lost. This FlashCopy operation is also referred to as a *time-zero copy* (T_0).

Immediately following the FlashCopy operation, the source and target volumes are available for use. The FlashCopy operation creates a bitmap that is referenced and maintained to direct I/O requests within the source and target relationship. This bitmap is updated to reflect the active block locations as data is copied in the background from the source to target, and updates are made to the source.

Figure 10-1 shows the redirection of the host I/O toward the source volume and the target volume.

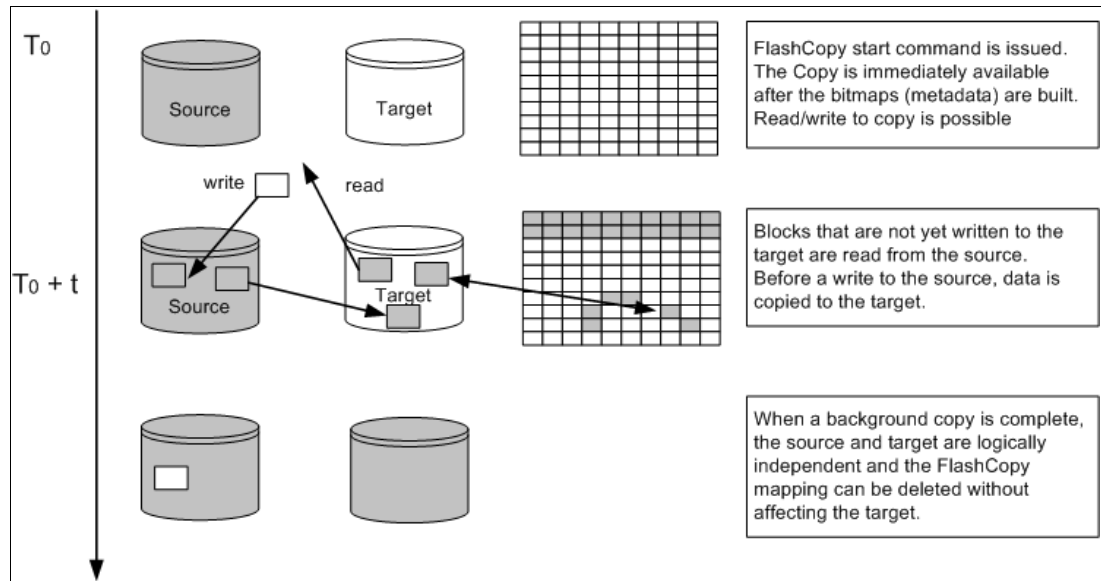


Figure 10-1 Redirection of host I/O

When data is copied between volumes, it is copied in units of address space known as *grains*. Grains are units of data that are grouped to optimize the use of the bitmap that track changes to the data between the source and target volume. You have the option of using 64 kilobytes (KB) or 256 KB grain sizes (256 KB is the default). The FlashCopy bitmap contains 1 bit for each grain, and is used to track whether the source grain was copied to the target. The 64 KB grain size uses bitmap space at a rate of four times the default 256 KB size.

The FlashCopy bitmap dictates the following read and write behaviors for the source and target volumes:

- ▶ Read I/O request to source
Reads are performed from the source volume the same as for non-FlashCopy volumes.
- ▶ Write I/O request to source
Writes to the source cause the grains of the source volume to be copied to the target if they were not already, and then the write is performed to the source.
- ▶ Read I/O request to target
Reads are performed from the target if the grains were already copied. Otherwise, the read is performed from the source.
- ▶ Write I/O request to target
Writes to the target cause the grain to be copied from the source to the target first, unless the entire grain is being written, in which case the write completes to the target only.

FlashCopy mappings

A FlashCopy mapping defines the relationship between a source volume and a target volume. FlashCopy mappings can be stand-alone mappings or a member of a consistency group, as described in “FlashCopy consistency groups” on page 462.

Incremental FlashCopy mappings

In an incremental FlashCopy, the initial mapping copies all of the data from the source volume to the target volume. Subsequent FlashCopy mappings copy only data that was modified since the initial FlashCopy mapping. This action reduces the amount of time that it takes to re-create an independent FlashCopy image. You can define a FlashCopy mapping as incremental *only* when you create the FlashCopy mapping.

Multiple target FlashCopy mappings

You can copy up to 256 target volumes from a single source volume. Each relationship between a source and target volume is managed by a unique mapping such that a single volume can be the source volume for up to 256 mappings.

Each of the mappings from a single source can be started and stopped independently. If multiple mappings from the same source are active (in the copying or stopping states), a dependency exists between these mappings.

If a single source volume has multiple target FlashCopy volumes, the write to the source volume does not cause its data to be copied to all of the targets. Instead, it is copied to the newest target volume only. The older targets refer to new targets first before they refer to the source. A dependency relationship exists between a particular target and all newer targets that share a source until all data is copied to this target and all older targets.

Cascaded FlashCopy mappings

The cascaded FlashCopy function enables a FlashCopy target volume to be the source volume of another FlashCopy mapping. Up to 256 mappings can exist in a cascade. If cascaded mappings and multiple target mappings are used, a tree of up to 256 mappings can be created.

Cascaded mappings differ from multiple target FlashCopy mappings in depth. Cascaded mappings have an association in the manner of $A > B > C$, where multiple target FlashCopy has an association in the manner $A > B1$ and $A > B2$.

Background copy

The background copy rate is a property of a FlashCopy mapping that is defined as a value of 0 - 100. The background copy rate can be defined and dynamically changed for individual FlashCopy mappings. A value of 0 disables background copy. This option is also called the *no-copy option*, which provides pointer-based images for limited lifetime uses.

With FlashCopy background copy, the source volume data is copied to the corresponding target volume in the FlashCopy mapping. If the background copy rate is set to 0 (which means disable the FlashCopy background copy), only data that changed on the source volume is copied to the target volume. The benefit of using a FlashCopy mapping with background copy enabled is that the target volume becomes a real independent clone of the FlashCopy mapping source volume after the copy is complete.

When the background copy is disabled, only the target volume is a valid copy of the source data while the FlashCopy mapping remains in place. Copying only the changed data saves your storage capacity (assuming that it is thin provisioned, and that the `-rsize` option was correctly set up.)

The relationship of the background copy rate value to the amount of data that is copied per second is shown in Table 10-1.

Table 10-1 Background copy rate

Value	Data that is copied per second	Grains per second (256 KB grain)	Grains per second (64 KB grain)
1 - 10	128 KB	0.5	2
11 - 20	256 KB	1	4
21 - 30	512 KB	2	8
31 - 40	1 megabyte (MB)	4	16
41 - 50	2 MB	8	32
51 - 60	4 MB	16	64
61 - 70	8 MB	32	128
71 - 80	16 MB	64	256
81 - 90	32 MB	128	512
91 - 100	64 MB	256	1024

Data copy rate: The data copy rate remains the same, regardless of the FlashCopy grain size. The difference is the number of grains that are copied per second. The grain size can be 64 KB or 256 KB. The smaller size uses more bitmap space, and therefore limits the total amount of FlashCopy space possible. However, it might be more efficient regarding the amount of data that is moved, depending on your environment.

Cleaning rate

The cleaning rate provides a method for FlashCopy copies with dependent mappings (multiple target or cascaded) to complete their background copies before their source goes offline, or is deleted after a stop is issued.

When you create or modify a FlashCopy mapping, you can specify a cleaning rate for the FlashCopy mapping that is independent of the background copy rate. The cleaning rate is also defined as a value of 0 - 100, which has the same relationship to data copied per second as the backup copy rate (see Table 10-1).

The cleaning rate controls the rate at which the cleaning process operates. The purpose of the cleaning process is to copy (or flush) data from FlashCopy source volumes upon which there are dependent mappings.

For cascaded and multiple target FlashCopy, the source might be a target for another FlashCopy or a source for a chain (cascade) of FlashCopy mappings. The cleaning process must complete before the FlashCopy mapping can go to the stopped state. This feature and the distinction between *stopping* and *stopped* states was added to prevent data access interruption for dependent mappings when their source is issued a stop.

FlashCopy mapping states

A mapping is in one of the following states at any point:

- ▶ Idle or Copied

The source and target volumes act as independent volumes even if a mapping exists between the two. Read/write caching is enabled for the source and the target volumes. If the mapping is incremental and the background copy is complete, the mapping records only the differences between the source and target volumes. The source and target volumes go offline if the connection to both nodes in the IBM Storwize V3700 storage system that the mapping is assigned to is lost.

- ▶ Copying

The copy is in progress. Read/write caching is enabled on the source and target volumes.

- ▶ Prepared

The mapping is ready to start. The target volume is online, but is not accessible. The target volume cannot perform read or write caching. Read and write caching is failed by the Small Computer System Interface (SCSI) front end as a hardware error.

If the mapping is incremental and a previous mapping completed, the mapping records only the differences between the source and target volumes. The source and target volumes go offline if the connection to both nodes in the IBM Storwize V3700 storage system that the mapping is assigned to is lost.

- ▶ Preparing

The target volume is online, but not accessible. The target volume cannot perform read or write caching. Read and write caching is failed by the SCSI front end as a hardware error. Any changed write data for the source volume is flushed from the cache. Any read or write data for the target volume is discarded from the cache.

If the mapping is incremental and a previous mapping completed, the mapping records only the differences between the source and target volumes. The source and target volumes go offline if the connection to both nodes in the IBM Storwize V3700 storage system that the mapping is assigned to is lost.

- ▶ Stopped

The mapping is stopped because you issued a **stop** command or an I/O error occurred. The target volume is offline and its data is lost. To access the target volume, you must restart or delete the mapping. The source volume is accessible, and the read and write cache is enabled. If the mapping is incremental, the mapping is recording write operations to the source volume. The source and target volumes go offline if the connection to both nodes in the IBM Storwize V3700 storage system that the mapping is assigned to is lost.

- ▶ Stopping

The mapping is copying data to another mapping. If the background copy process is complete, the target volume is online while the stopping copy process completes. If the background copy process did not complete, data is discarded from the target volume cache. The target volume is offline while the stopping copy process runs. The source volume is accessible for I/O operations.

- ▶ Suspended

The mapping did start, but it did not complete. Access to the metadata is lost, which causes the source and target volume to go offline. When access to the metadata is restored, the mapping returns to the copying or stopping state, and the source and target volumes return online. The background copy process resumes.

Any data that was not flushed, and was written to the source or target volume before the suspension, is in cache until the mapping leaves the suspended state.

FlashCopy consistency groups

Consistency groups address the requirement to preserve point-in-time data consistency across multiple volumes for applications that include related data that spans them. For these volumes, consistency groups maintain the integrity of the FlashCopy by ensuring that dependent writes are run in the application's intended sequence. For more information about dependent writes, see "Dependent writes" on page 462.

When consistency groups are used, the FlashCopy commands are issued to the FlashCopy consistency group, which performs the operation on all FlashCopy mappings that are contained within the consistency group.

Figure 10-2 shows a consistency group that consists of two FlashCopy mappings.

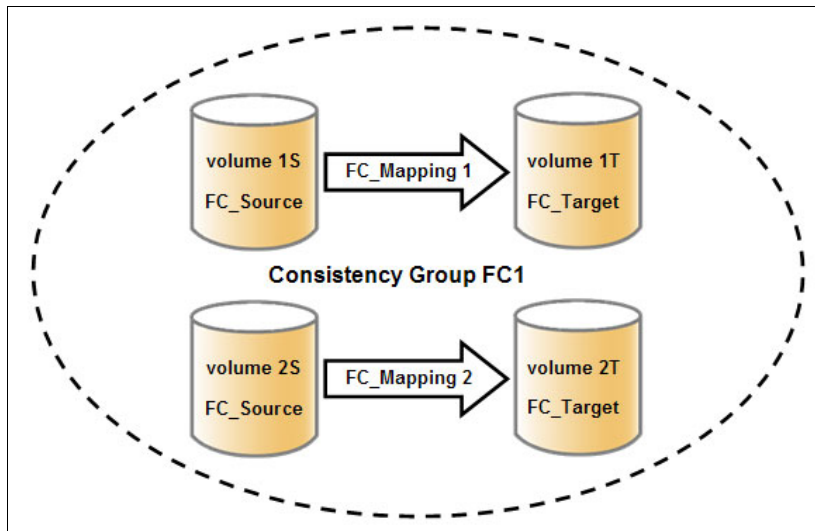


Figure 10-2 FlashCopy consistency group

FlashCopy mapping management: After an individual FlashCopy mapping is added to a consistency group, it can be managed only as part of the group. Operations such as **start** and **stop** are no longer allowed on the individual mapping.

Dependent writes

To show why it is crucial to use consistency groups when a data set spans multiple volumes, consider the following typical sequence of writes for a database update transaction:

1. A write is run to update the database log, which indicates that a database update is about to be performed.
2. A second write is run to complete the actual update to the database.
3. A third write is run to update the database log, which indicates that the database update completed successfully.

The database ensures the correct ordering of these writes by waiting for each step to complete before it starts the next step. However, if the database log (updates 1 and 3) and the database (update 2) are on separate volumes, it is possible for the FlashCopy of the database volume to occur before the FlashCopy of the database log. This situation can result in the target volumes seeing writes (1) and (3) but not (2), because the FlashCopy of the database volume occurred before the write was completed.

In this case, if the database was restarted using the backup that was made from the FlashCopy target volumes, the database log indicates that the transaction completed successfully when, in fact, it had not. This situation occurs because the FlashCopy of the volume with the database file was started (bitmap was created) before the write completed to the volume. Therefore, the transaction is lost and the integrity of the database is in question.

To overcome the issue of dependent writes across volumes, and to create a consistent image of the client data, it is necessary to perform a FlashCopy operation on multiple volumes as an atomic operation using consistency groups. A FlashCopy consistency group can contain up to 512 FlashCopy mappings. The more mappings that you have, the more time it takes to prepare the consistency group.

FlashCopy commands can then be issued to the FlashCopy consistency group and simultaneously for all of the FlashCopy mappings that are defined in the consistency group. For example, when the FlashCopy for the consistency group is started, all FlashCopy mappings in the consistency group are started at the same time, which results in a point-in-time copy that is consistent across all FlashCopy mappings that are contained in the consistency group. A consistency group aggregates FlashCopy mappings, not volumes.

Therefore, where a source volume has multiple FlashCopy mappings, they can be in the same or separate consistency groups. If a particular volume is the source volume for multiple FlashCopy mappings, you might want to create separate consistency groups to separate each mapping of the same source volume. Regardless of whether the source volume with multiple target volumes is in the same consistency group or in separate consistency groups, the resulting FlashCopy produces multiple identical copies of the source data.

The consistency group can be specified when the mapping is created. You can also add the FlashCopy mapping to a consistency group, or change the consistency group of a FlashCopy mapping later.

Important: Do not place stand-alone mappings into a consistency group, because they become controlled as part of that consistency group.

FlashCopy consistency group states

A FlashCopy consistency group is in one of the following states at any point:

- ▶ Idle or Copied. All FlashCopy Mappings in this consistency group are in the Idle or Copied state.
- ▶ Preparing. At least one FlashCopy mapping in this consistency group is in the Preparing state.
- ▶ Prepared. The consistency group is ready to start. While in this state, the target volumes of all FlashCopy mappings in this consistency group are not accessible.
- ▶ Copying. At least one FlashCopy mapping in the consistency group is in the Copying state, and no FlashCopy mappings are in the Suspended state.
- ▶ Stopping. At least one FlashCopy mapping in the consistency group is in the Stopping state, and no FlashCopy mappings are in the Copying or Suspended state.
- ▶ Stopped. The consistency group is stopped because you issued a command or an I/O error occurred.
- ▶ Suspended. At least one FlashCopy mapping in the consistency group is in the Suspended state.
- ▶ Empty. The consistency group does not have any FlashCopy mappings.

Reverse FlashCopy

Reverse FlashCopy enables FlashCopy targets to become restore points for the source without breaking the FlashCopy relationship, and without waiting for the original copy operation to complete. It supports multiple targets and multiple rollback points. A key advantage of Reverse FlashCopy is that it does not delete the original target, therefore enabling processes that use the target, such as a tape backup, to continue uninterrupted.

You can also create an optional copy of the source volume that is made before the reverse copy operation is started. This copy restores the original source data, which can be useful for diagnostic purposes.

Figure 10-3 shows an example of the reverse FlashCopy scenario.

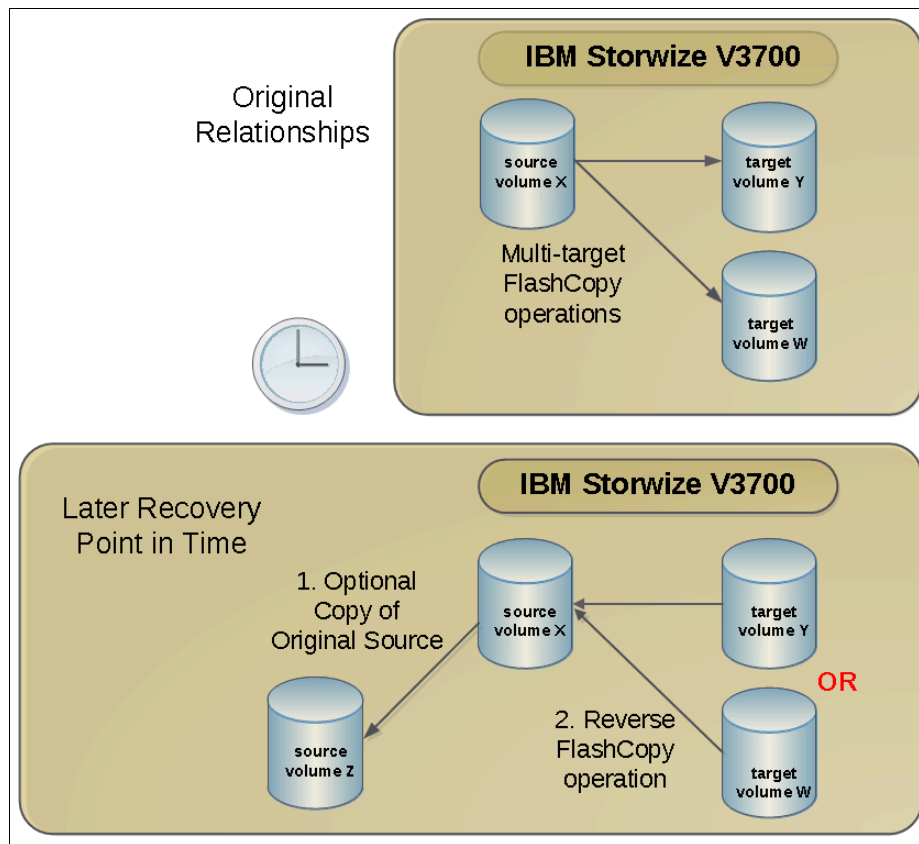


Figure 10-3 Reverse FlashCopy scenario

To restore from a FlashCopy backup using the GUI, complete the following steps:

1. (Optional) Create a target volume (volume Z), and run FlashCopy on the production volume (volume X) to copy data on to the new target for later problem analysis.
2. Create a FlashCopy map with the backup to be restored (volume Y), or (volume W) as the source volume and volume X as the target volume.
3. Start the FlashCopy map (volume Y to volume X).

The -restore option: In the CLI, you must add the **-restore** option to the **svctask startfcmap** command manually. For more information about using the CLI, see Appendix A, "Command-line interface setup and SAN Boot" on page 649. The GUI adds this suffix automatically if it detects that you are mapping from a target back to the source.

Regardless of whether the initial FlashCopy map (volume X to volume Y) is incremental, the Reverse FlashCopy operation copies only the modified data.

Consistency groups are reversed by creating a set of new “reverse” FlashCopy maps and adding them to a new “reverse” consistency group. Consistency groups cannot contain more than one FlashCopy map with the same target volume.

10.1.3 Planning for FlashCopy

There are several items that must be considered before a FlashCopy is performed, which are described in this section.

Guidelines for FlashCopy implementation

Consider the following guidelines for FlashCopy implementation:

- ▶ The source and target volumes must be on the same IBM Storwize V5000 storage system.
- ▶ The source and target volumes do not need to be in the same Storage Pool.
- ▶ The FlashCopy source and target volumes can be thin-provisioned.
- ▶ The source and target volumes must be the same size. The size of the source and target volumes cannot be altered (increased or decreased) while a FlashCopy mapping is defined.
- ▶ FlashCopy operations perform in direct proportion to the performance of the source and target disks. If you have a fast source disk and slow target disk, the performance of the source disk is reduced because it must wait for the write operation to occur at the target before it can write to the source.

Maximum configurations for FlashCopy

Table 10-2 shows some of the FlashCopy maximum configurations. For more information about the most current values, see the IBM Storwize V3700 Knowledge Center on the following website:

http://www.ibm.com/support/knowledgecenter/STLM5A/landing/V3700_welcome.html

Also see the planning and installation documentation on the following website:

https://www.ibm.com/support/entry/myportal/plan_install/system_storage/disk_systems/mid-range_disk_systems/ibm_storwize_v3700?productContext=-2033461677

Table 10-2 FlashCopy maximum configurations

FlashCopy property	Maximum
FlashCopy targets per source	256
FlashCopy mappings per cluster	4096
FlashCopy consistency groups per cluster	255
FlashCopy mappings per consistency group	512

FlashCopy presets

The IBM Storwize V3700 storage system provides three FlashCopy presets (Snapshot, Clone, and Backup) to simplify the more common FlashCopy operations, as shown in Table 10-3.

Table 10-3 FlashCopy presets

Preset	Purpose
Snapshot	<p>Creates a point-in-time view of the production data. The snapshot is not intended to be an independent copy. Instead, it is used to maintain a view of the production data at the time that the snapshot is created.</p> <p>This preset automatically creates a thin-provisioned target volume with none of the capacity that is allocated at the time of creation. The preset uses a FlashCopy mapping with none of the background copy, so that only data written to the source or target is copied to the target volume.</p>
Clone	<p>Creates an exact replica of the volume, which can be changed without affecting the original volume. After the copy operation completes, the mapping that was created by the preset is automatically deleted.</p> <p>This preset automatically creates a volume with the same properties as the source volume, and creates a FlashCopy mapping with a background copy rate of 50. The FlashCopy mapping is configured to automatically delete when the FlashCopy mapping reaches 100% completion.</p>
Backup	<p>Creates a point-in-time replica of the production data. After the copy completes, the backup view can be refreshed from the production data, with minimal copying of data from the production volume to the backup volume.</p> <p>This preset automatically creates a volume with the same properties as the source volume. The preset creates an incremental FlashCopy mapping with a background copy rate of 50.</p>

Presets: All of the presets can be adjusted by using the expandable Advanced Settings section in the GUI.

10.1.4 Managing FlashCopy using the GUI

The IBM Storwize V3700 storage system provides a separate function icon to access copy service management. The following windows are available for managing FlashCopy under the Copy Services function icon:

- ▶ FlashCopy
- ▶ Consistency Groups
- ▶ FlashCopy Mappings

The Copy Services function icon is shown in Figure 10-4.

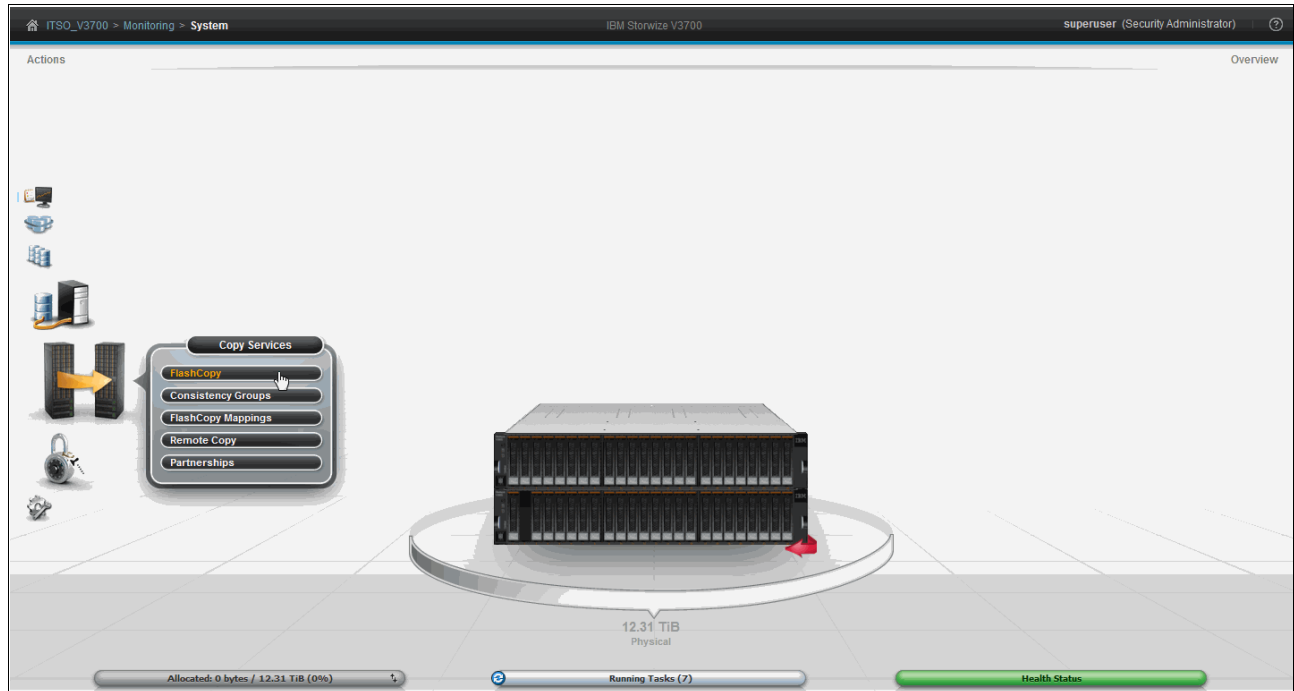


Figure 10-4 Copy Services function icon

Most of the actions to manage the FlashCopy mapping can be done in the FlashCopy window or the FlashCopy Mappings windows, although the quick path to create FlashCopy presets can be found only in the FlashCopy window.

To manage FlashCopy mapping, complete the following steps:

1. Click **FlashCopy** in the Copy Services function icon menu and the FlashCopy window opens, as shown in Figure 10-5. The FlashCopy mappings are organized by volumes.

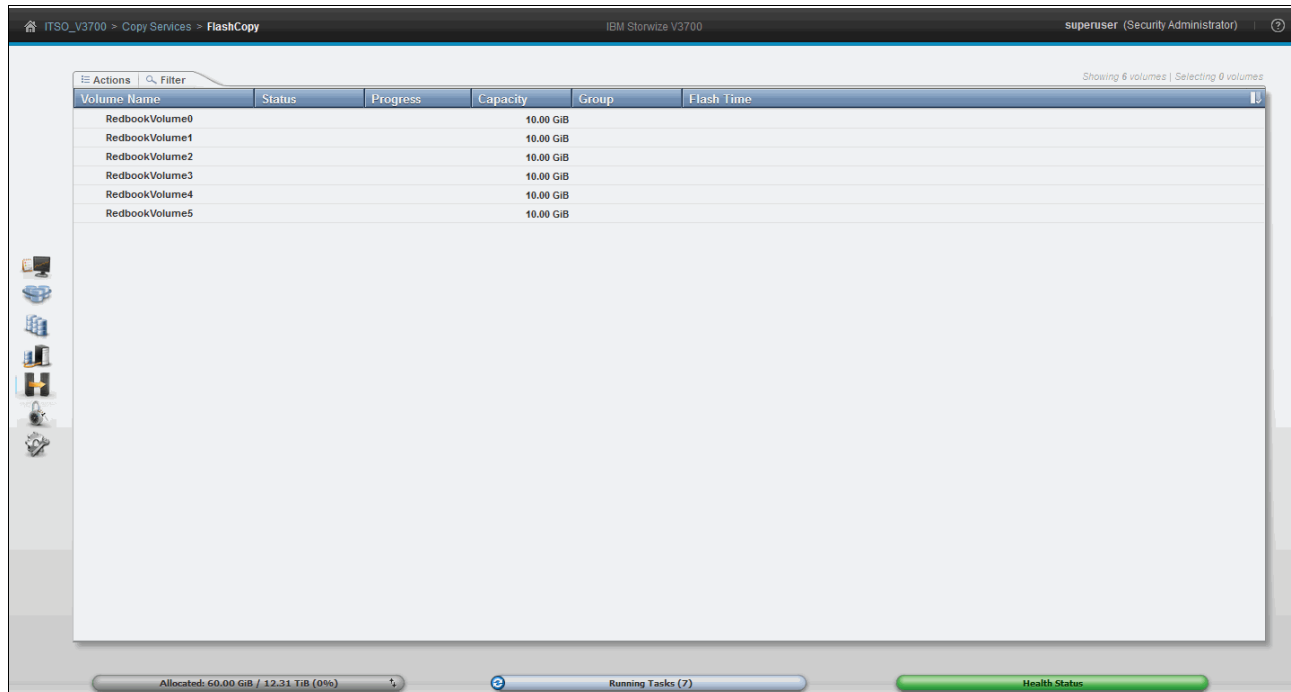


Figure 10-5 FlashCopy window

2. Click **FlashCopy Mappings** in the Copy Services function icon menu and the FlashCopy Mappings window opens, as shown in Figure 10-6. In this window, the FlashCopy mappings are listed individually.




Figure 10-6 FlashCopy Mappings window

3. The Consistency Groups window is used to manage the consistency groups for FlashCopy mappings. Click **Consistency Groups** in the Copy Services function icon menu and the Consistency Groups window opens, as shown in Figure 10-7.




Figure 10-7 Consistency Groups window

Quick path to create FlashCopy presets

It is easy to create a FlashCopy by using the presets in the FlashCopy window.

Creating a snapshot

To create a snapshot, complete the following steps:

1. In the FlashCopy window, choose a volume and click **Create Snapshot** from the **Actions** menu, as shown in Figure 10-8.
2. Alternatively, you can highlight your chosen volume and right-click to access the **Actions** menu.




Figure 10-8 Create a snapshot using the preset

You now have a snapshot volume for the volume you selected.

Creating a clone

To create a clone, complete the following steps:

1. In the FlashCopy window, choose a volume and click **Create Clone** from the **Actions** menu, as shown in Figure 10-9.
2. Alternatively, highlight your chosen volume and right-click to access the **Actions** menu.




Figure 10-9 Create a clone from the preset

You now have a clone volume for the volume you selected.

Creating a backup

To create a backup, complete the following steps:

1. In the FlashCopy window, choose a volume and click **Create Backup** from the **Actions** menu, as shown in Figure 10-10.
2. Alternatively, highlight your chosen volume and right-click to access the **Actions** menu.




Figure 10-10 Create a backup from the preset

You now have a backup volume for the volume you selected.

- In the FlashCopy window and in the FlashCopy Mappings window, you can monitor the progress of the running FlashCopy operations, as shown in Figure 10-11. The progress bars for each target volume indicate the copy progress as a percentage. The copy progress remains 0% for snapshots (there is no change until data is written to the target volume). The copy progresses for clone and backup continues to increase until the copy process completes.




Figure 10-11 FlashCopy in progress viewed in the FlashCopy Mappings window

4. The copy progress can also be found under the Running Tasks status indicator, as shown in Figure 10-12.

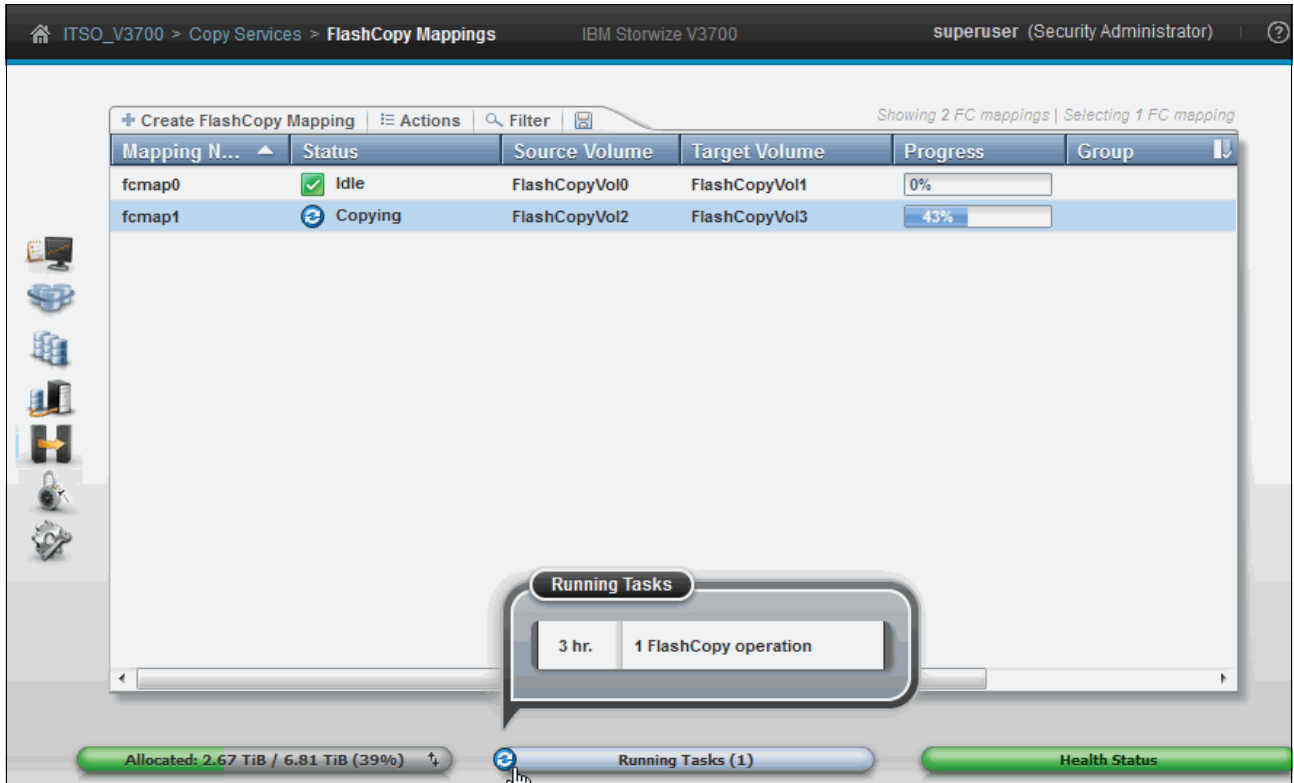


Figure 10-12 Running Tasks bar: FlashCopy operations

5. This view is slightly different from the FlashCopy and FlashCopy Mappings windows, as shown in Figure 10-13.




Figure 10-13 FlashCopy operations shown through Running Tasks

After the copy processes complete, you find that a FlashCopy mapping with the clone preset is deleted automatically. There are now two identical volumes that are independent of each other.

10.1.5 Managing FlashCopy mappings

The FlashCopy presets cover the most frequently used FlashCopy configurations for general situations. However, customized FlashCopy mappings are still necessary in some complicated scenarios.

Creating FlashCopy mappings

You can create FlashCopy mappings through the FlashCopy window:

1. Select the volume that you want to be the source volume for the FlashCopy mapping and click **Advanced FlashCopy** from the **Actions** menu, as shown in Figure 10-14.
2. Alternatively, select the volume and right-click.

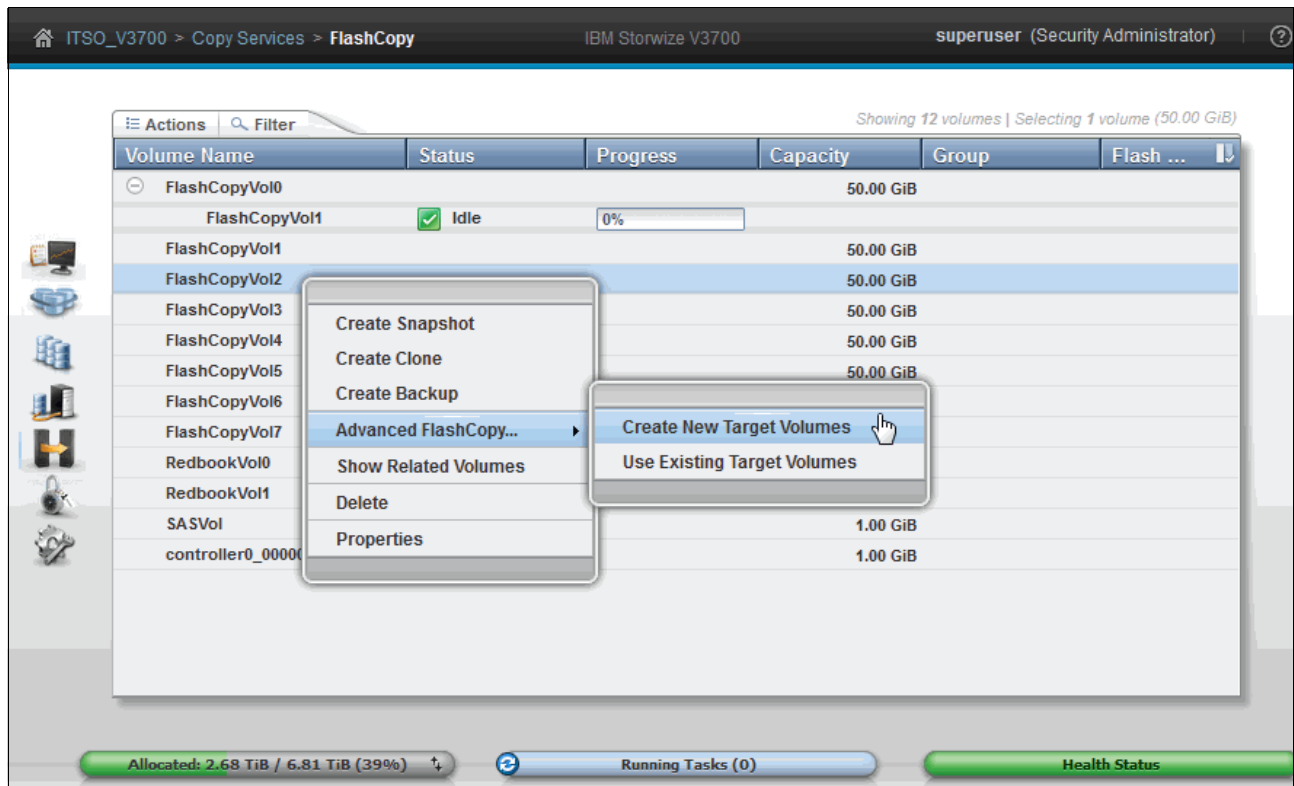


Figure 10-14 Create advanced FlashCopy

3. You can **Create New Target Volumes** as part of the mapping process, or **Use Existing Target Volumes**. We describe creating volumes next. To use existing volumes, see “Using existing target volumes” on page 482.

Creating target volumes

Complete the following steps to create target volumes:

1. Click **Create new target volumes** if you have not yet created the target volume.

- The wizard guides you to choose a preset, as shown in Figure 10-15. Choose one preset that has the most similar configuration to the one that is required.




Figure 10-15 Choose a preset most similar to your requirement

- Click **Advanced Settings** to make any appropriate adjustments to the advanced settings:
 - The following default advanced settings for the *Snapshot* preset are shown in Figure 10-16:
 - Background Copy: 0
 - Incremental: No
 - Auto Delete after completion: No
 - Cleaning Rate: 0

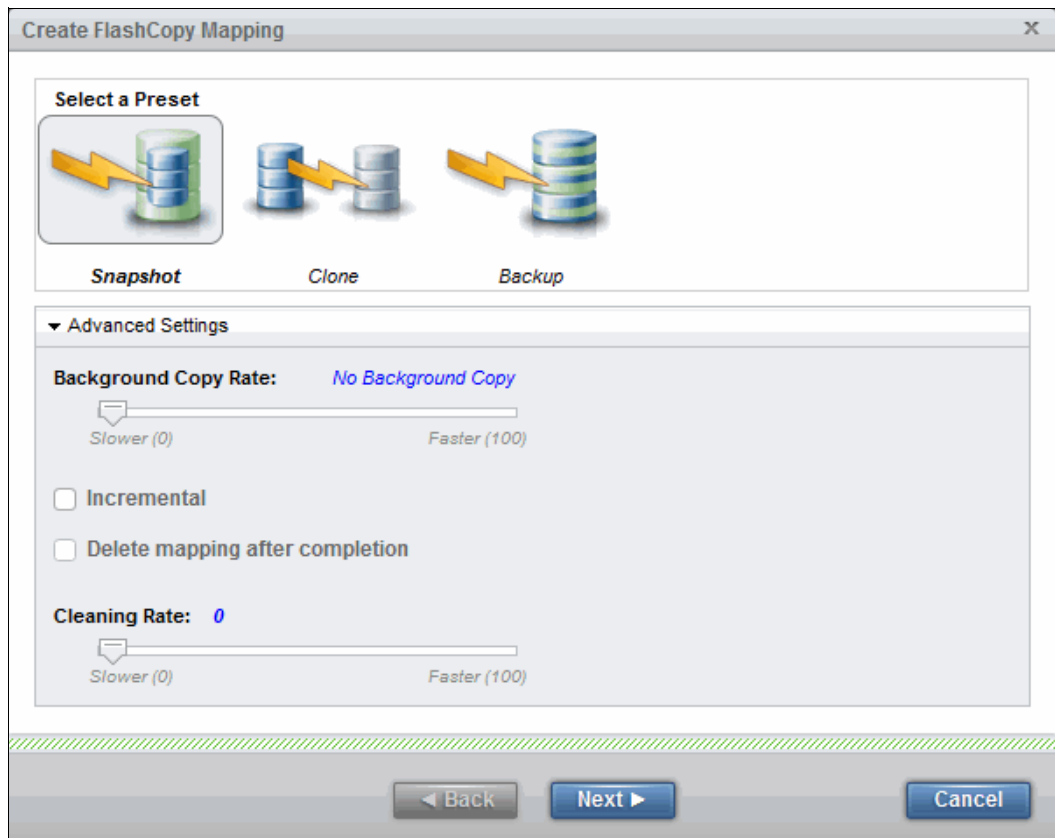


Figure 10-16 Default settings for the snapshot preset

b. The following Advanced Settings for the *Clone* preset are shown in Figure 10-17:

- Background Copy: 50
- Incremental: No
- Auto Delete after completion: Yes
- Cleaning Rate: 50

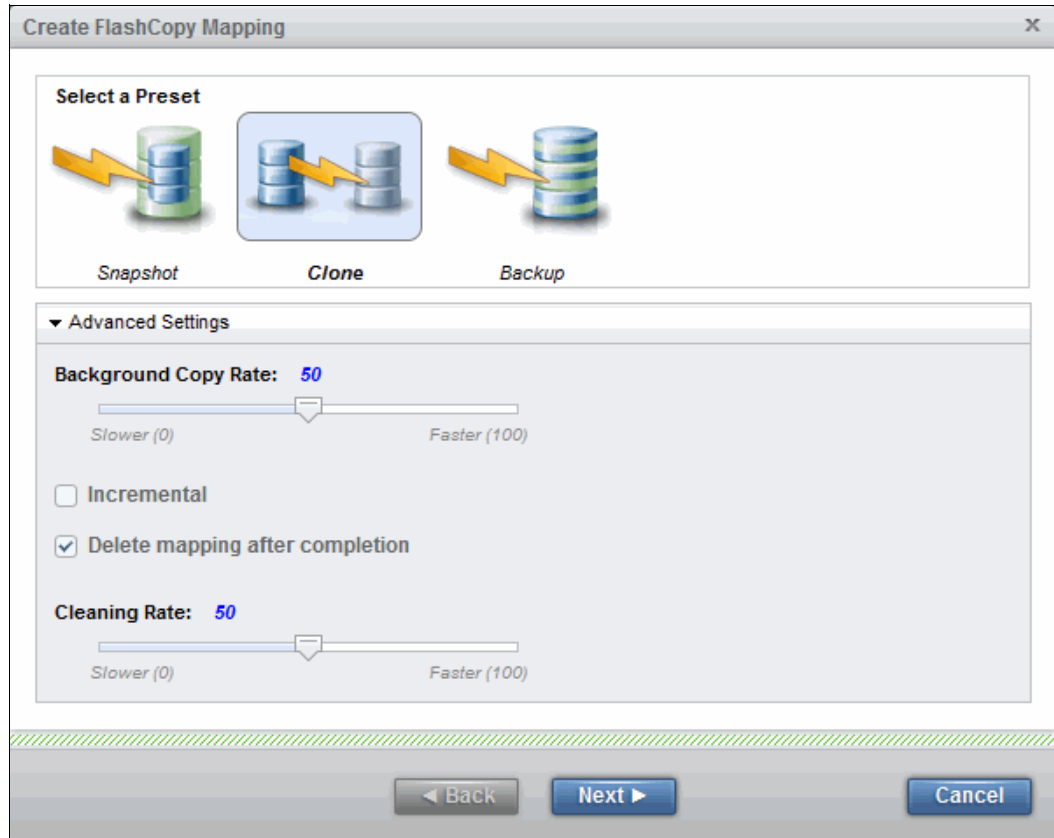


Figure 10-17 Default settings for the clone preset

c. Figure 10-18 shows the following Advanced Settings for the *Backup* preset:

- Background Copy: 50
- Incremental: Yes
- Auto Delete after completion: No
- Cleaning Rate: 50

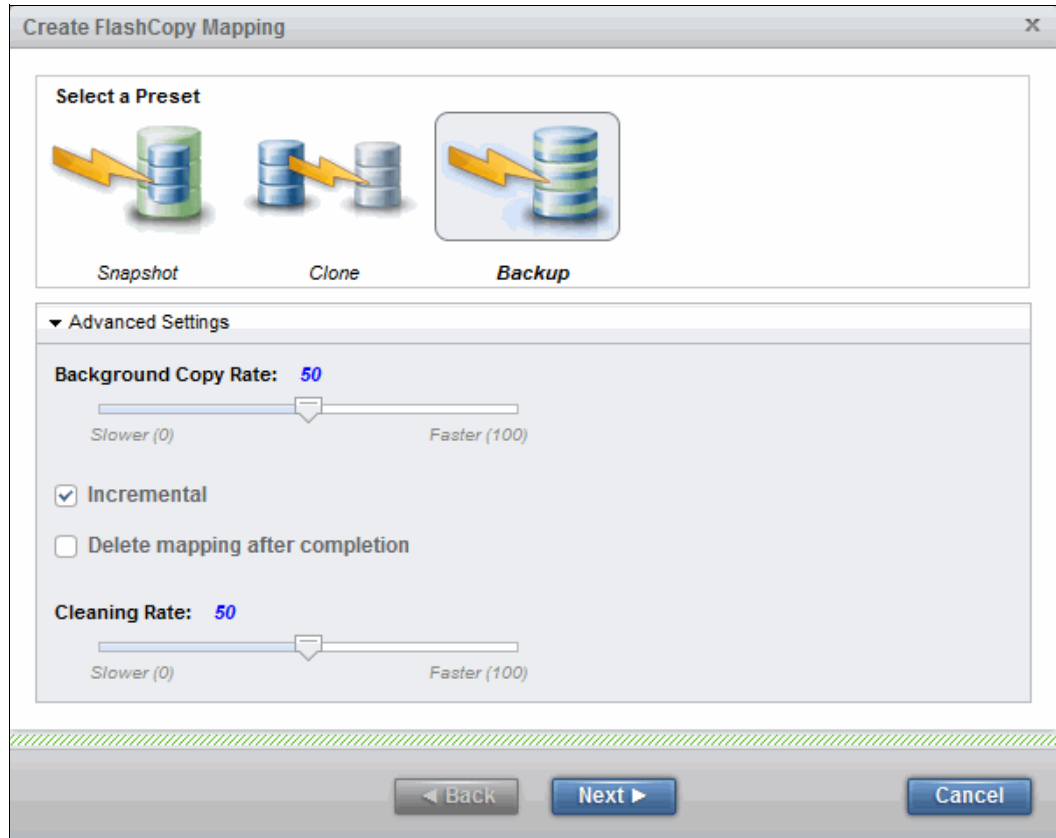


Figure 10-18 Default settings for the backup preset

4. Change the settings of the FlashCopy mapping according to your requirements and click **Next**.
5. In the next step, you can add your FlashCopy mapping to a consistency group, as shown in Figure 10-19. If the consistency group is not ready, the FlashCopy mapping can be added to the consistency group afterward. Click **Next** to continue.

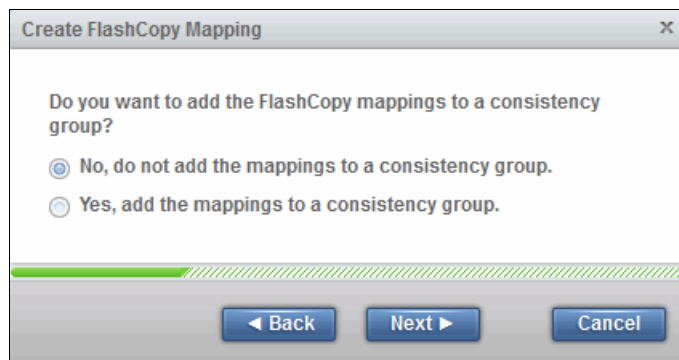


Figure 10-19 Add FlashCopy mapping to a consistency group

6. You can define how the new target volumes manage capacity, as shown in Figure 10-20. The Create a Generic volume option is your default choice if you selected Clone or Backup as your basic preset.

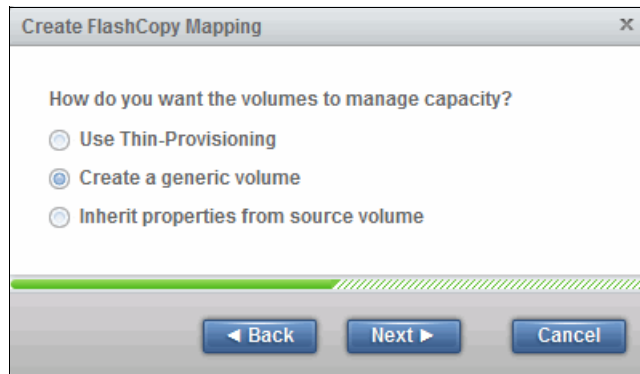


Figure 10-20 Defining how target volumes use capacity

7. If you select a thin-provisioned volume, more options are available, as shown in Figure 10-21. The other option is to inherit the Source Volume properties.

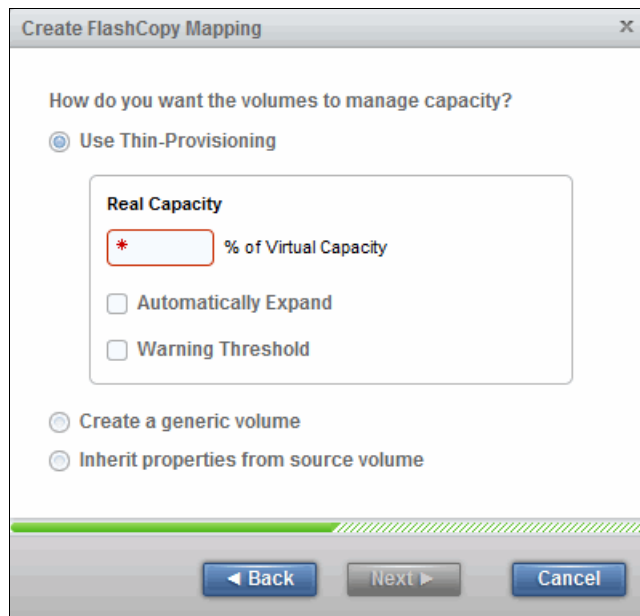


Figure 10-21 Advanced thin provision options

- Click **Next** to continue. You can choose from which Storage Pool you want to create your target volume, as shown in Figure 10-22. You can select the same Storage Pool used by the source volume, or select a different pool.

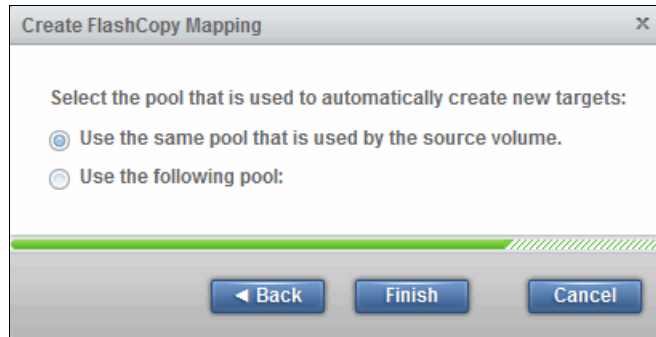


Figure 10-22 Selecting Storage Pool for target volume

- Click **Finish** when you make your decision. The mappings and volume are created.
- Click **Close** to see the FlashCopy mapping created on your volume with a new target, as shown in Figure 10-23. The status of the newly created FlashCopy mapping is Idle. It can be started, as described in “Starting a FlashCopy mapping” on page 486.

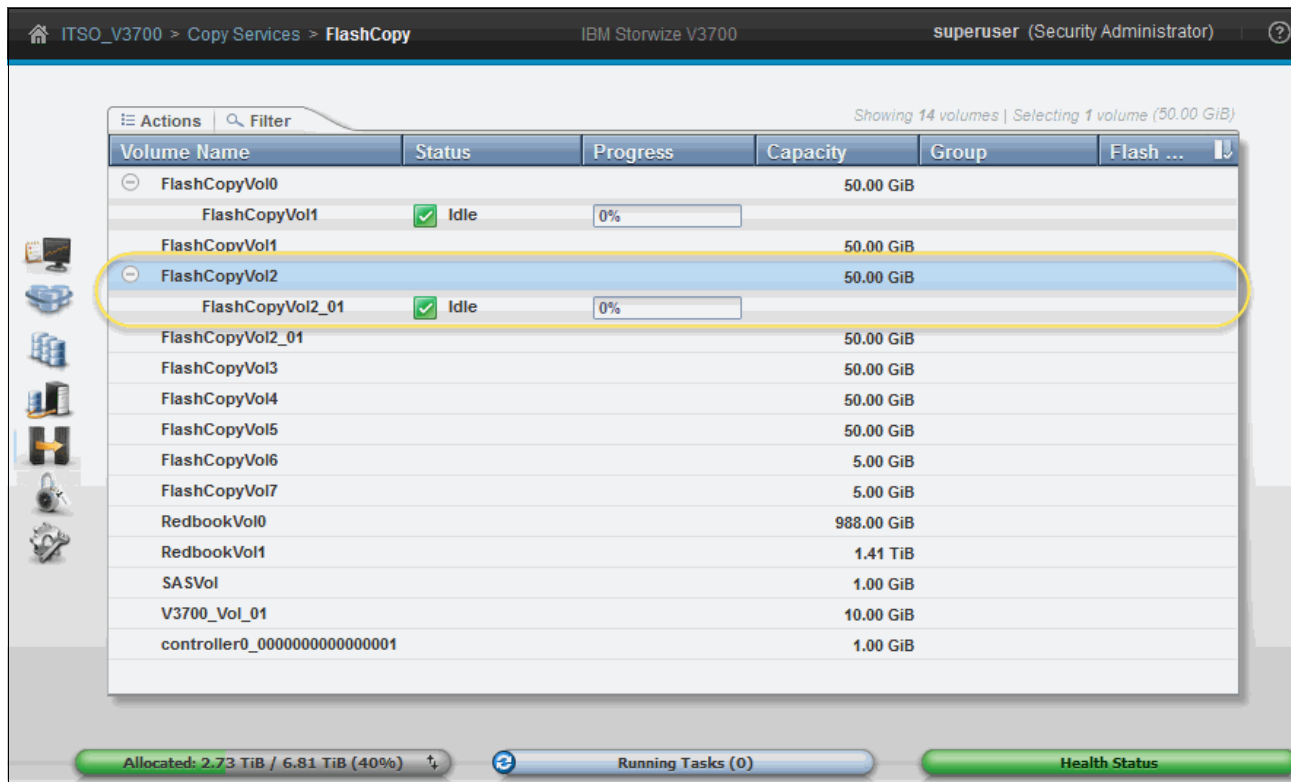


Figure 10-23 New FlashCopy mapping is created with a new target

Using existing target volumes

Complete the following steps to use existing target volumes:

1. If you already have candidate target volumes, select **Use Existing Target Volumes** in the Advanced FlashCopy menu, as shown in Figure 10-24.




Figure 10-24 Create FlashCopy mapping using existing target volume

2. You must choose the target volume for the source volume that you selected. Select the target volume from the drop-down menu in the right pane of the window and click **Add**, as shown in Figure 10-25.

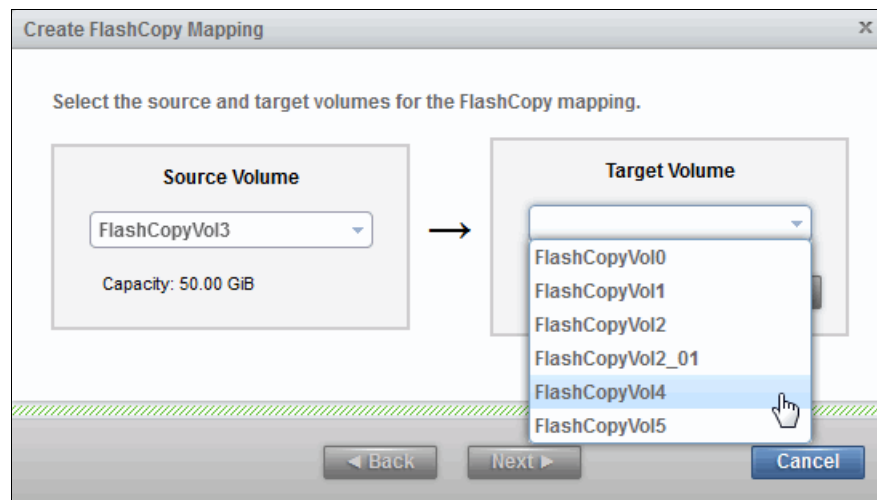


Figure 10-25 Select the target volume

3. After you click **Add**, the FlashCopy mapping is listed, as shown in Figure 10-26. Click the red X to delete it if the FlashCopy mapping is not the one that you want to create.
4. Additional mappings can be made by selecting further source and target volumes from the drop-down boxes and clicking **Add**. When all of the FlashCopy mappings that you want are defined, click **Next** to continue.

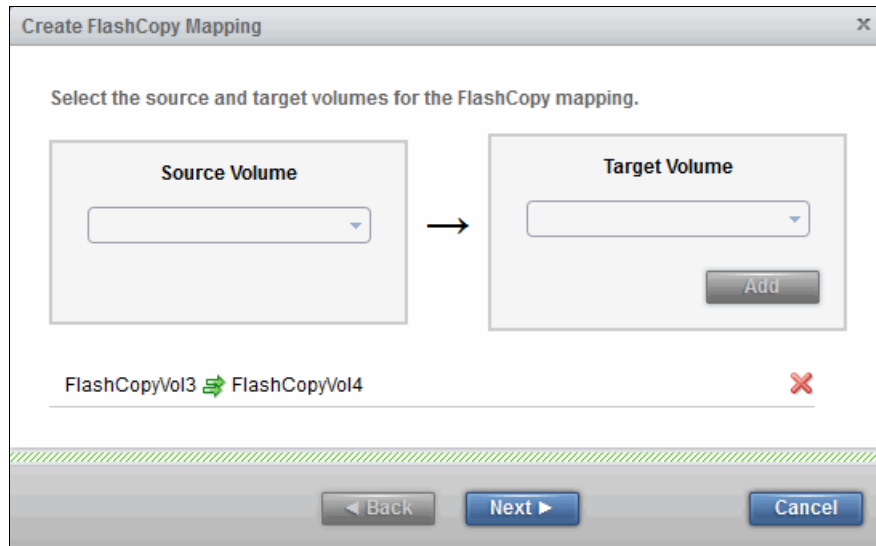


Figure 10-26 Add FlashCopy mapping

5. Select the preset and (if necessary) adjust the settings by using the Advanced Settings section, as shown in Figure 10-27. (For more information about the Advanced Settings, see “Creating target volumes” on page 476.) Confirm that the settings meet your requirements and then click **Next**.

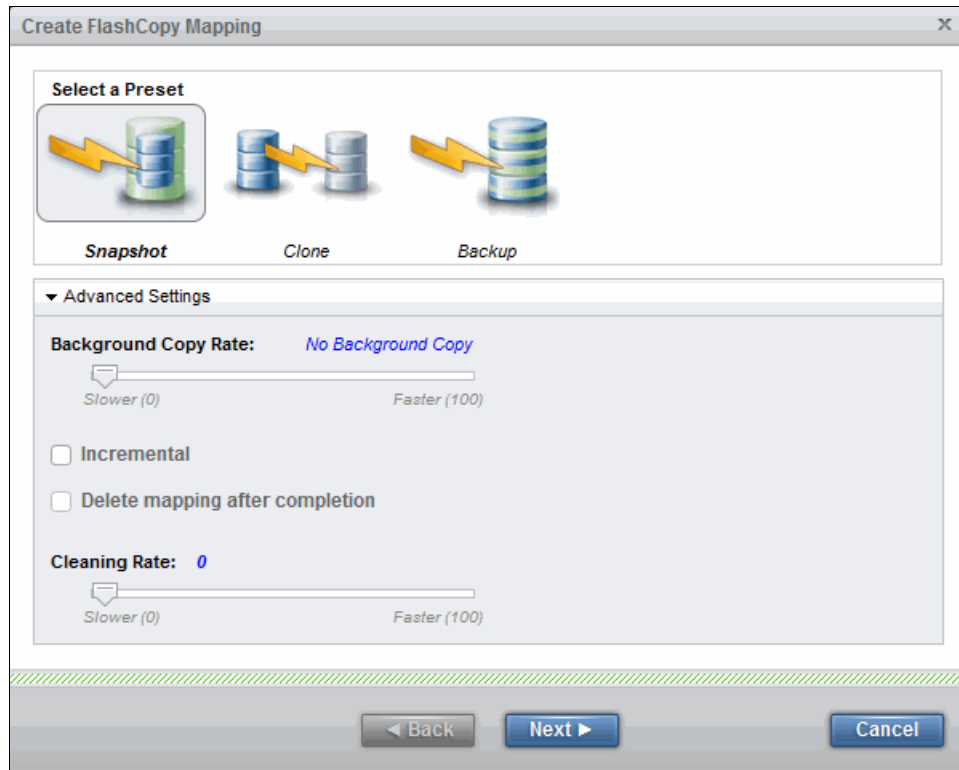


Figure 10-27 Select a preset and make your adjustments

6. You can now add the FlashCopy mapping to a consistency group (if necessary), as shown in Figure 10-28. Selecting **Yes** displays a drop-down menu from which you can select a consistency group. Click **Finish** and the FlashCopy mapping is created with the status of Idle, as shown in Figure 10-28.

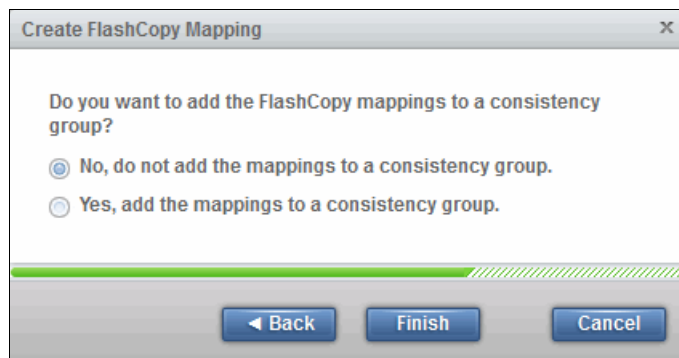


Figure 10-28 Select a preset and make your adjustments

Creating new FlashCopy mappings

To create FlashCopy mappings in the FlashCopy Mappings window, complete the following steps:

1. Click **Create FlashCopy Mapping**, as shown in Figure 10-29.

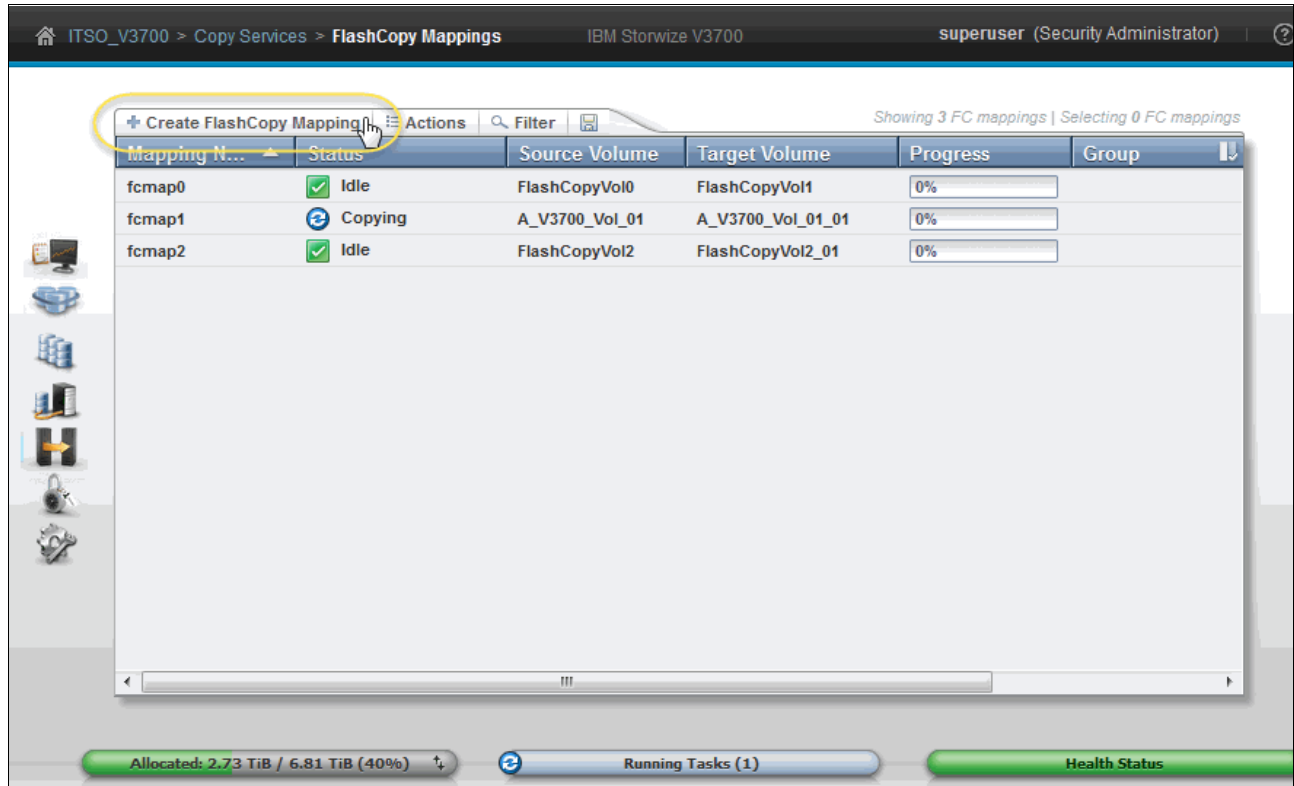


Figure 10-29 Create a FlashCopy mapping in the FlashCopy Mappings window

2. A wizard guides you through the process to create a FlashCopy mapping. The steps are the same as creating an Advanced FlashCopy mapping using Existing Target Volumes, as described in “Using existing target volumes” on page 482.

Starting a FlashCopy mapping

Most of the FlashCopy mapping actions can be performed in the FlashCopy window or the FlashCopy Mapping window. For the actions that are available in both windows, we show in the following sections the steps in the FlashCopy window, although the steps are the same if you were to use the FlashCopy Mapping window. To start mapping, complete the following steps:

1. Select the FlashCopy target volume in the FlashCopy window.
2. Select the **Start** option from the **Actions** menu (as shown in Figure 10-30), or right-click and select **Start**. The status of the FlashCopy mapping changes from Idle to Copying.




Figure 10-30 Start FlashCopy mapping

Stopping a FlashCopy mapping

To stop FlashCopy mapping, complete the following steps:

1. Select the FlashCopy target volume in the FlashCopy window.
2. Click the **Stop** option from the **Actions** menu, as shown in Figure 10-31. After the stopping process completes, the status of the FlashCopy mapping is changed to Stopped.

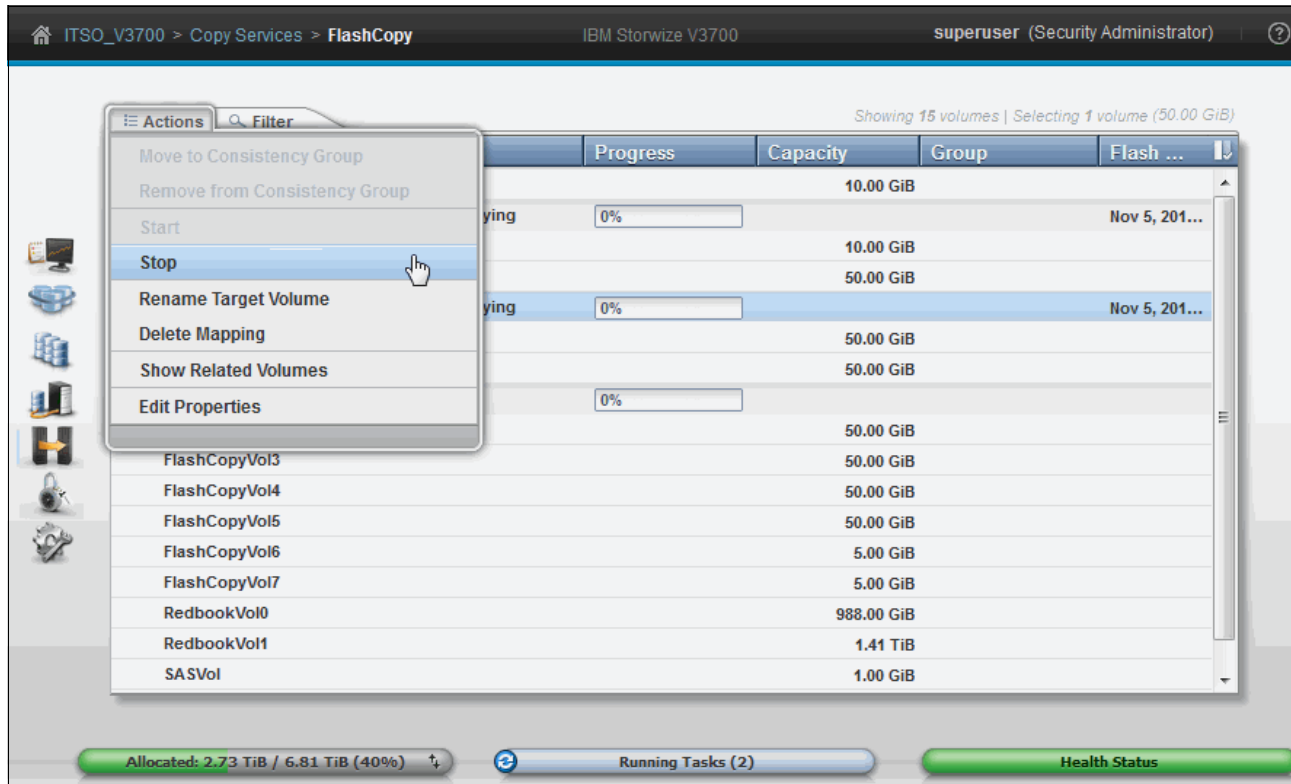


Figure 10-31 Stopping a FlashCopy mapping

Renaming the target volume

If the FlashCopy target volumes were created automatically by IBM Storwize V3700, the name of the target volume is the source volume name plus a suffix that includes numbers. The name of the target volumes can be changed to be more meaningful in your environment:

1. Select the FlashCopy target volume in the FlashCopy window.
2. Click the **Rename Target Volume** option from the **Actions** menu (as shown in Figure 10-32), or right-click the selected volume.

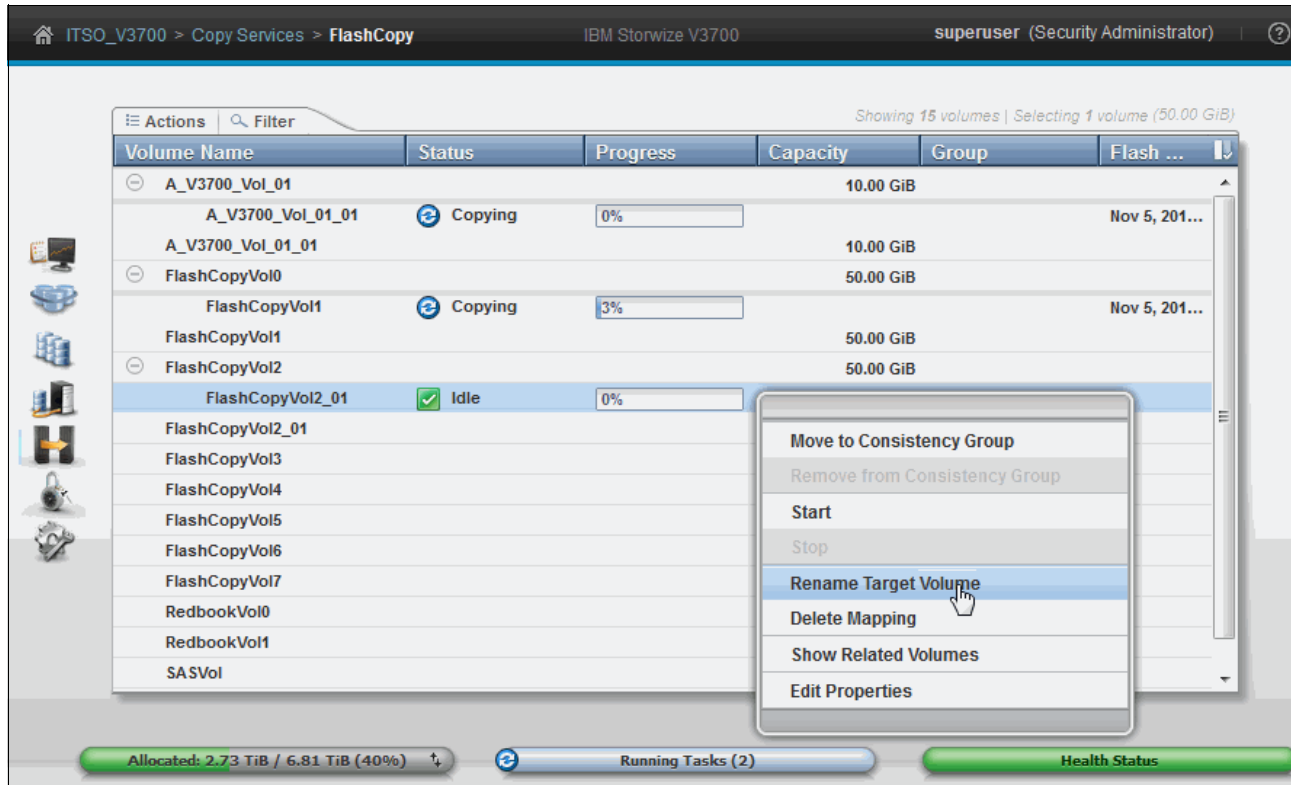


Figure 10-32 Rename a target volume

3. Enter your new name for the target volume, as shown in Figure 10-33. Click **Rename** to finish.

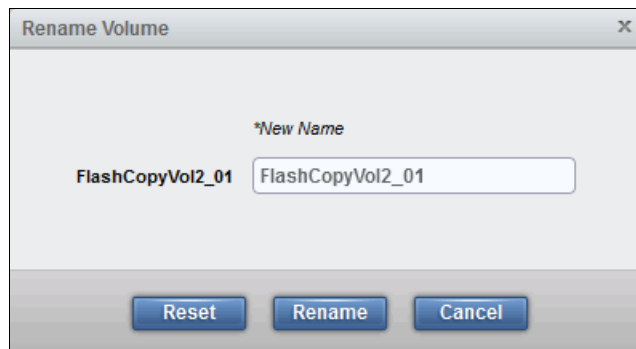


Figure 10-33 Rename a target volume

Renaming a FlashCopy mapping

The FlashCopy mappings are created with names that begin with fcmap. You can change this to be something more meaningful:

1. Select the FlashCopy mapping in the FlashCopy Mappings window.
2. Click **Rename Mapping** in the **Actions** menu, as shown in Figure 10-34.

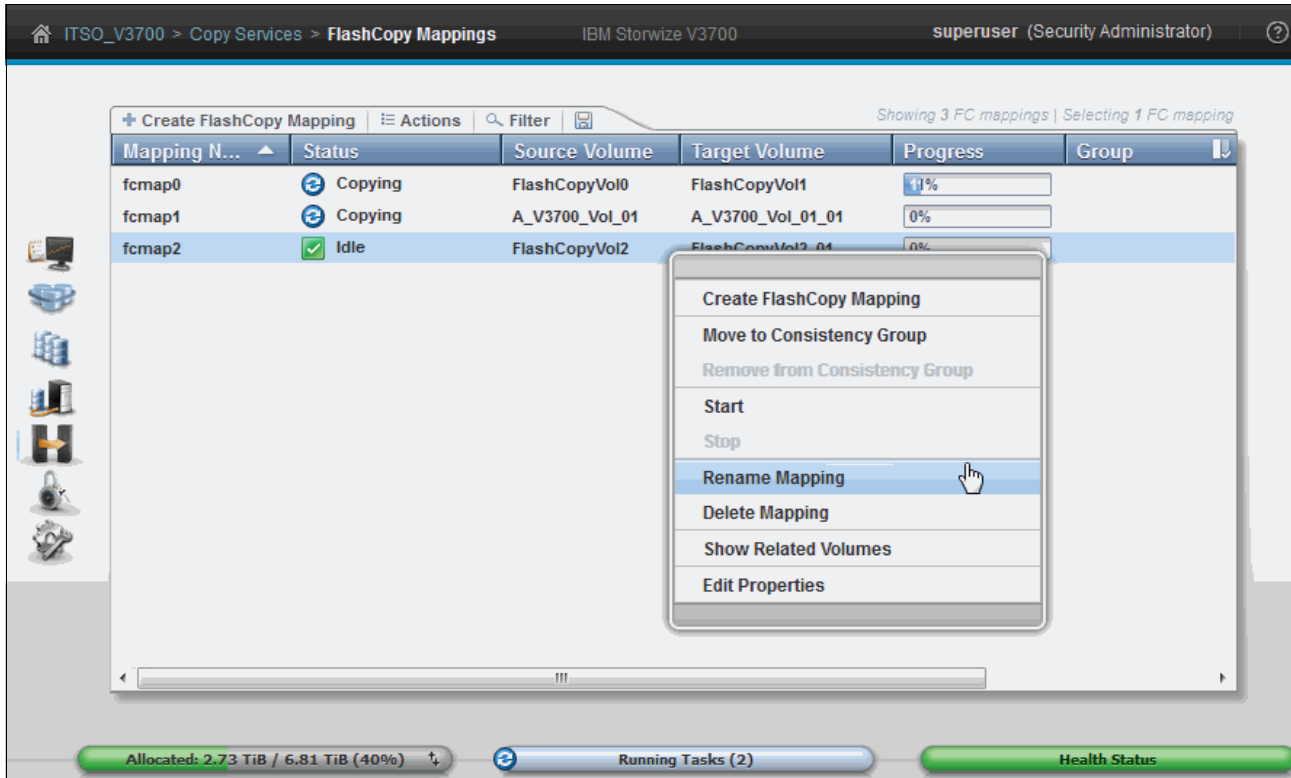


Figure 10-34 Rename a FlashCopy mapping

3. Enter the new name for the FlashCopy mapping, as shown in Figure 10-35. Click **Rename** to finish.

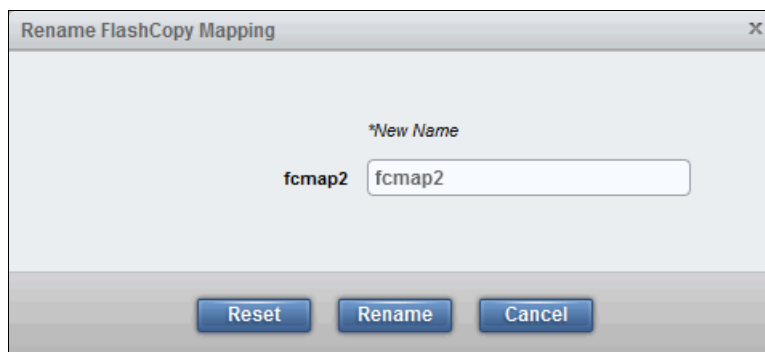


Figure 10-35 Enter a new name for the FlashCopy mapping

Deleting a FlashCopy mapping

To delete FlashCopy mapping, complete the following steps:

1. Select the FlashCopy target volume in the FlashCopy window.
2. Click **Delete Mapping** in the **Actions** menu (as shown in Figure 10-36), or right-click the selected volume.

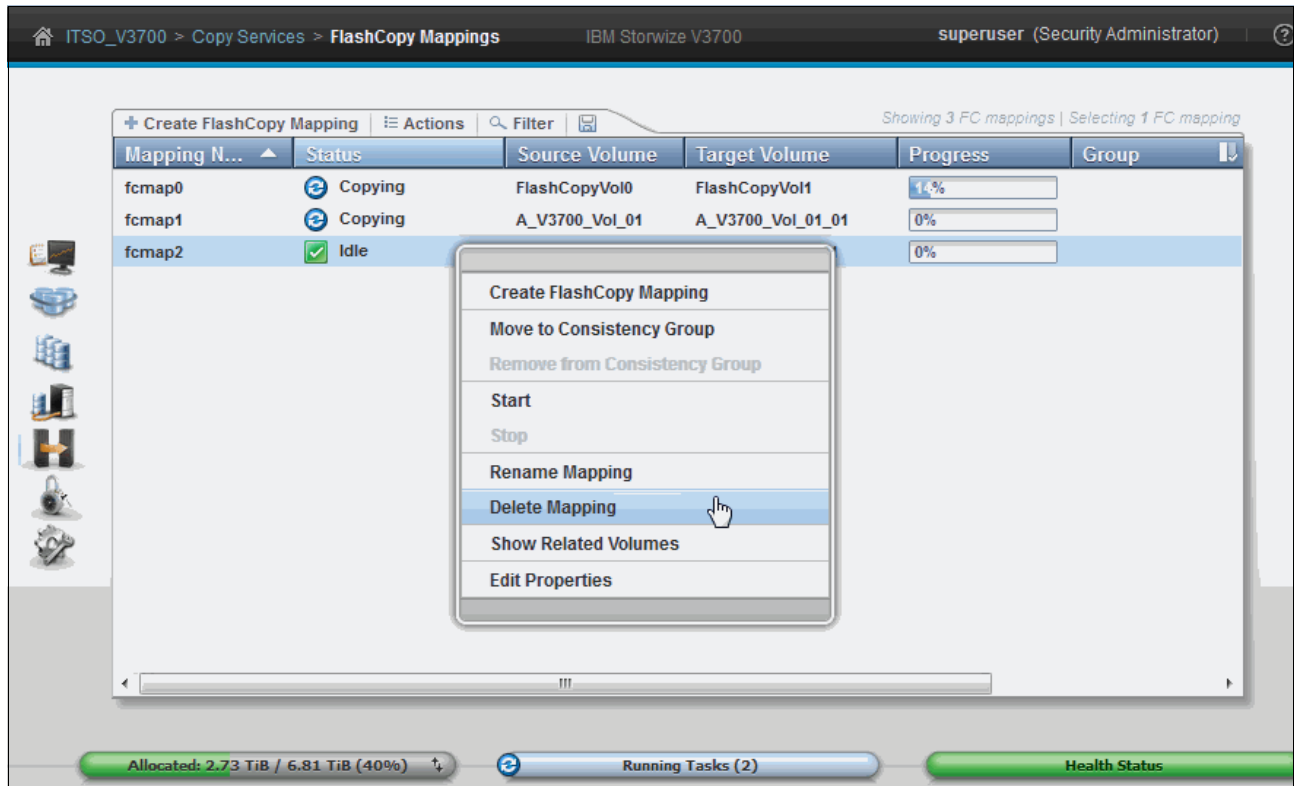


Figure 10-36 Select Delete Mapping

FlashCopy Mapping state: If the FlashCopy mapping is in the Copying state, it must be stopped before it is deleted.

3. You must confirm your action to delete FlashCopy mappings in the window that opens, as shown in Figure 10-37. Verify the number of FlashCopy mappings that you want to delete. If you want to delete the FlashCopy mappings when the data on the target volume is inconsistent with the source volume, select the option to do so.

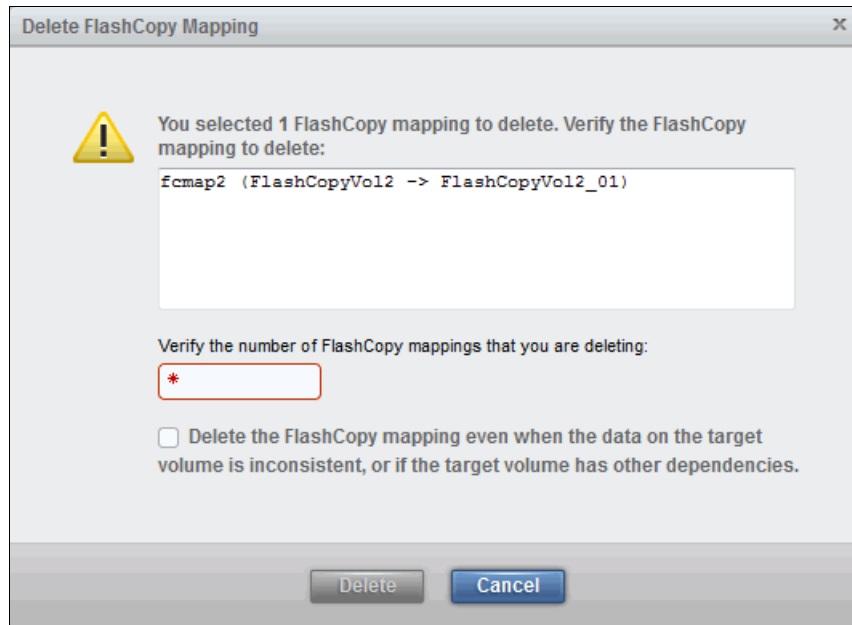


Figure 10-37 Confirm the deletion of FlashCopy mappings

4. Click **Delete** and your FlashCopy mapping is removed.

Important: Deleting the FlashCopy mapping does *not* delete the target volume. If you want to reclaim the storage space occupied by the target volume, you must delete the target volume manually.

Showing related volumes

To show the FlashCopy mapping dependencies, complete the following steps:

1. Select a target or source volume in the FlashCopy window.
2. Click **Show Related Volumes** in the **Actions** menu (as shown in Figure 10-38) or right-click the selected volume.

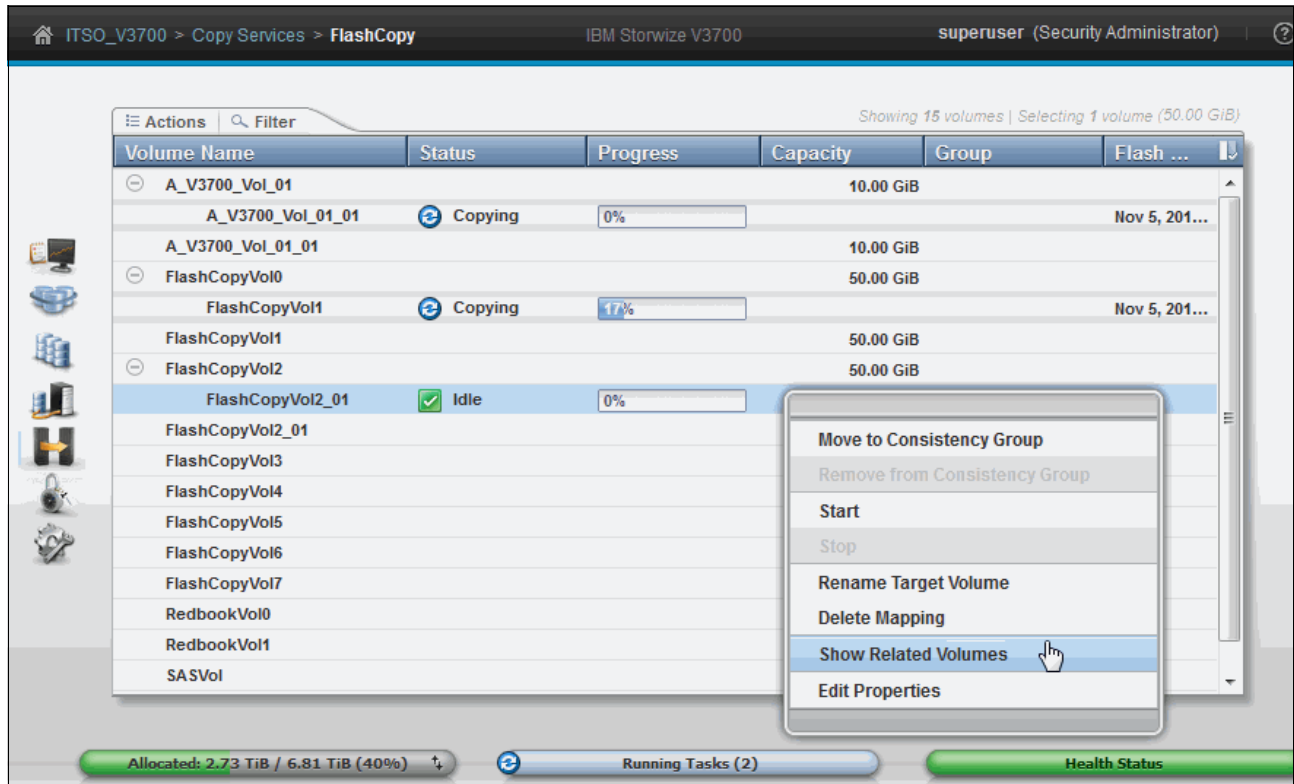


Figure 10-38 Show Related Volumes menu

3. The FlashCopy mapping dependency tree opens, as shown in Figure 10-39.



Figure 10-39 FlashCopy mapping dependency

4. Clicking either volume shows the properties of the volume, as shown in Figure 10-40.

The screenshot shows a window titled "Volume Details: FlashCopyVol2_01" with three tabs: "Overview", "Host Maps", and "Member MDisks". The "Overview" tab is active, displaying a list of volume properties and a copy status section.

Volume Name	FlashCopyVol2_01
Volume ID	12
Status	✓ Online
Capacity	50.00 GiB
# of FlashCopy Mappings	1
Volume UID	6005076000EB01F5B00000000000000C
Preferred Node	node2
I/O Throttling	Disabled
Mirror Sync Rate	50
Cache Mode	Enabled
Cache State	Empty
UDID (OpenVMS)	

Below the table is an "Edit" button. To the right, the "Copy 0" status is shown with a green progress bar at 100%. Below the progress bar, the following information is displayed:

- Pool: V3700_R5Pool_01
- Striped
- Copy Status: Online
- Capacity: Enterprise Tier: 50.00 GiB, Total: 50.00 GiB

At the bottom of the window is a "Close" button.

Figure 10-40 Target FlashCopy Volume details

Editing properties

The background copy rate and cleaning rate can be changed after the FlashCopy mapping is created. To do so, complete the following steps:

1. Select the FlashCopy target mapping in the FlashCopy window.
2. Click **Edit Properties** in the **Actions** menu (as shown in Figure 10-41), or right-click the FlashCopy target mapping.

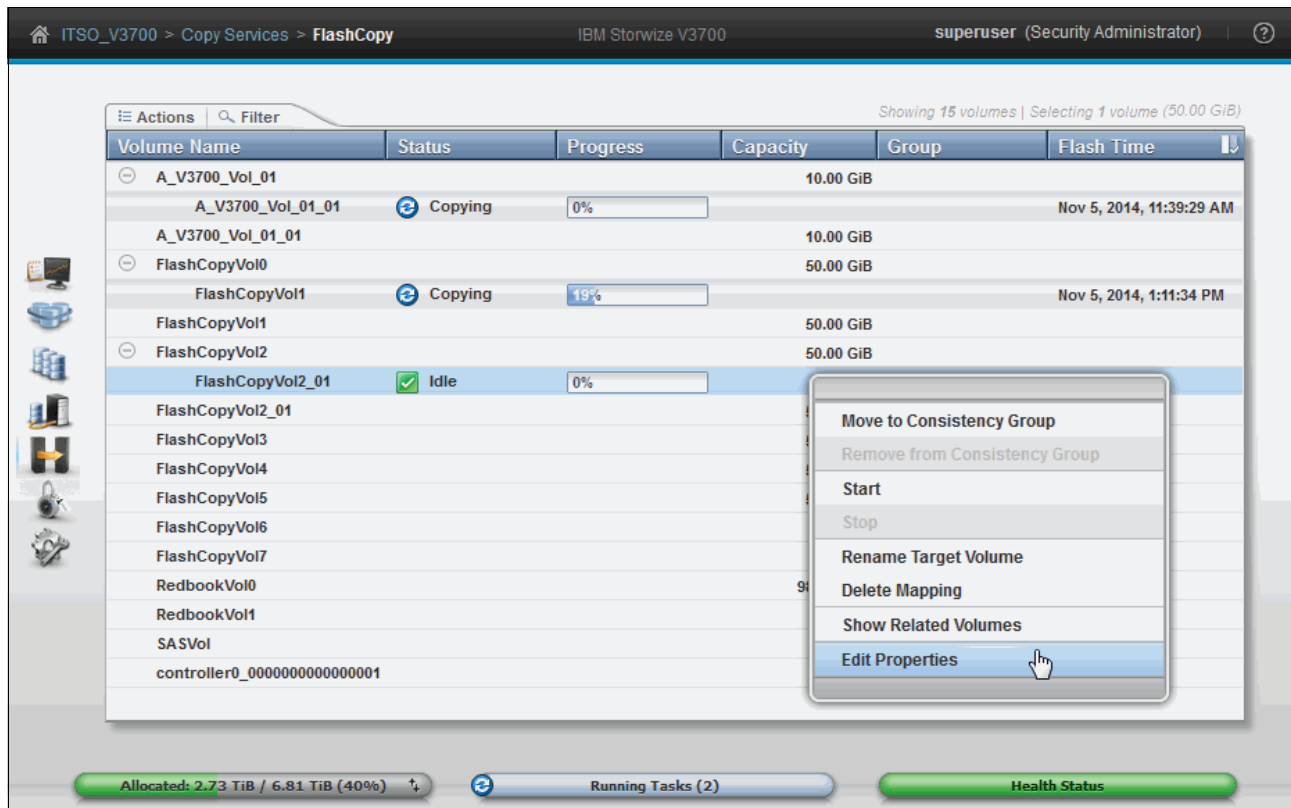


Figure 10-41 Edit Properties menu

3. You can then modify the value of the background copy rate and cleaning rate by moving the pointers on the bars, as shown in Figure 10-42.

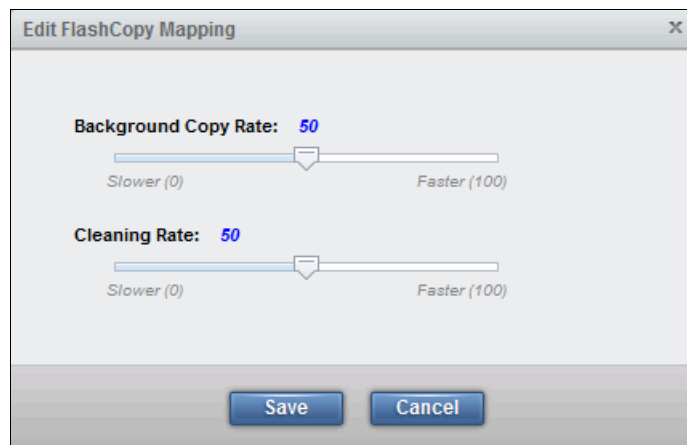


Figure 10-42 Change the copy and cleaning rate

4. Click **Save** to save changes.

Restoring from a FlashCopy

Complete the following steps to manipulate FlashCopy target volumes to restore a source volume to a previous known state from a snapshot:

1. Identify the FlashCopy relationship that you want to restore. In this example, you want to restore FlashCopyVol6, as shown in Figure 10-43. Restore from FlashCopyVol7, which is a snapshot of the data from an earlier point in time.

The screenshot shows the IBM Storwize V3700 FlashCopy management interface. The breadcrumb navigation is "ITSO_V3700 > Copy Services > FlashCopy". The user is "superuser (Security Administrator)". The interface displays a table of 17 volumes. The volume "FlashCopyVol6" is highlighted with a yellow circle, and a restore operation is shown in progress from "FlashCopyVol7". The progress bar for this operation is at 12%. The "Flash Time" for the restore is "Nov 5, 2014, 7:40:02 PM".

Volume Name	Status	Progress	Capacity	Group	Flash Time
A_V3700_Vol_01			10.00 GiB		
A_V3700_Vol_01_01			10.00 GiB		
FlashCopyVol0			50.00 GiB		
FlashCopyVol1			50.00 GiB		
FlashCopyVol2			50.00 GiB		
FlashCopyVol2_01			50.00 GiB		
FlashCopyVol6			5.00 GiB		
FlashCopyVol7	Copying	12%			Nov 5, 2014, 7:40:02 PM
FlashCopyVol7			5.00 GiB		
LeeTest0			25.00 GiB		
LeeTest1			25.00 GiB		
RBVolume_0			2.00 GiB		
RBVolume_1			2.00 GiB		
RBVolume_2			2.00 GiB		
RedbookVolume			1.00 GiB		
SASVol			1.00 GiB		
Volume001			5.00 GiB		
controller0_0000000000000001			1.00 GiB		

Allocated: 231.00 GiB / 6.17 TiB (4%) Running Tasks (2) Health Status

Figure 10-43 Starting FlashCopy restore

Note: The mapping progress in the case of a snapshot indicates the percentage of the original data that has changed since the snapshot was taken. Because this old data no longer exists on the source volume, it has been physically copied to the target volume rather than being bitmap referenced to the source.

The greater the percentage of data changed, the longer the snapshot restore process takes, because more data needs to be physically copied back to the source. Therefore, it is good practice to take regular snapshots so that data can be restored quickly.

2. Create a mapping using the target volume of the mapping to be restored. In this example, it is FlashCopyVol17, as shown in Figure 10-44. Select **Advanced FlashCopy** → **Use Existing Target Volumes**.




Figure 10-44 Create reverse mapping

3. The Source Volume is preselected with the target volume that was selected in the previous step. Select the Target Volume from the drop-down menu (you select the source volume that you want to restore). In this example, select FlashCopyVol16, as shown in Figure 10-45. Because FlashCopyVol16 is the source volume for FlashCopyVol17, you are effectively reversing the copy direction by doing this.

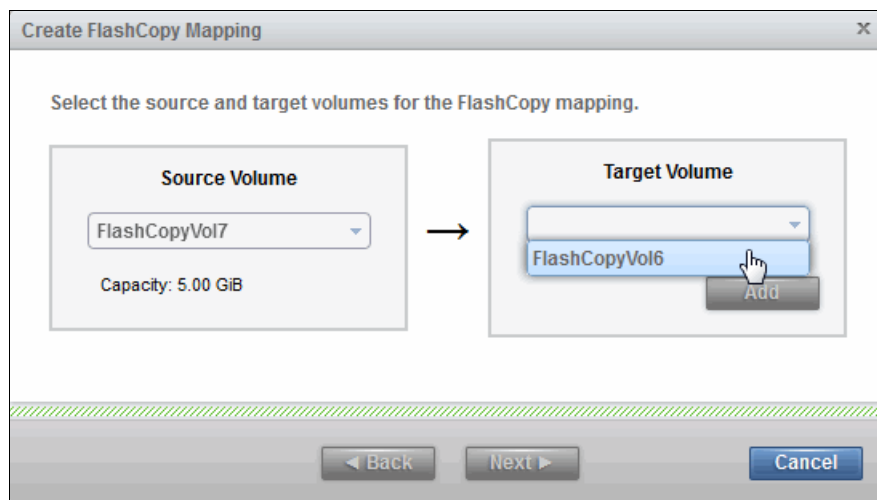


Figure 10-45 Select target volume

- Click **Add**. A warning message displays, as shown in Figure 10-46. Click **Close**. This message is shown because you are using a source as a target.



Figure 10-46 FlashCopy restore warning

- Click **Next** and you see a FlashCopy preset choice, as shown in Figure 10-47.




Figure 10-47 Choose FlashCopy preset

- Select **Snapshot** and click **Next**.
- In the next window, you are asked if the new mapping is to be part of a consistency group, as shown in Figure 10-48. In this example, the new mapping is not part of a consistency group, so click **No** and then **Finish** to create the mapping.

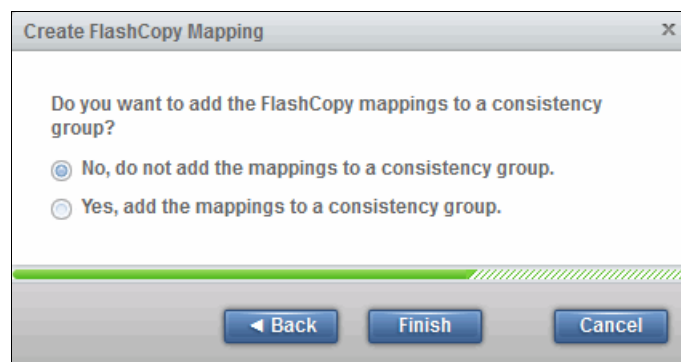


Figure 10-48 Choosing whether to add the mapping to a consistency group

8. The new reverse mapping is now created and shown in the Idle state, as shown in Figure 10-49.

The screenshot shows the IBM Storwize V3700 FlashCopy management interface. The breadcrumb navigation is 'ITSO_V3700 > Copy Services > FlashCopy'. The user is 'superuser (Security Administrator)'. The interface displays a table of 17 volumes. The 'FlashCopyVol7' row is highlighted with a yellow circle, indicating it is the selected volume. The status of 'FlashCopyVol7' is 'Idle', and the status of 'FlashCopyVol6' is 'Copying'.

Volume Name	Status	Progress	Capacity	Group	Flash Time
A_V3700_Vol_01			10.00 GiB		
A_V3700_Vol_01_01			10.00 GiB		
FlashCopyVol0			50.00 GiB		
FlashCopyVol1			50.00 GiB		
FlashCopyVol2			50.00 GiB		
FlashCopyVol2_01			50.00 GiB		
FlashCopyVol6			5.00 GiB		
FlashCopyVol7	Copying	100%			Nov 5, 2014, 7:40:02 PM
FlashCopyVol7	Idle	0%	5.00 GiB		
FlashCopyVol6	Idle	0%			
LeeTest0			25.00 GiB		
LeeTest1			25.00 GiB		
RBVolume_0			2.00 GiB		
RBVolume_1			2.00 GiB		
RBVolume_2			2.00 GiB		
RedbookVolume			1.00 GiB		
SASVol			1.00 GiB		
Volume001			5.00 GiB		
controller0_0000000000000001			1.00 GiB		

At the bottom of the interface, there are three status bars: 'Allocated: 231.00 GiB / 6.17 TiB (4%)', 'Running Tasks (2)', and 'Health Status'.

Figure 10-49 New reverse mapping

- To restore the original source volume FlashCopyVol16 with the snapshot you took (FlashCopyVol17), select the new mapping and right-click to open the Actions menu, as shown in Figure 10-50.

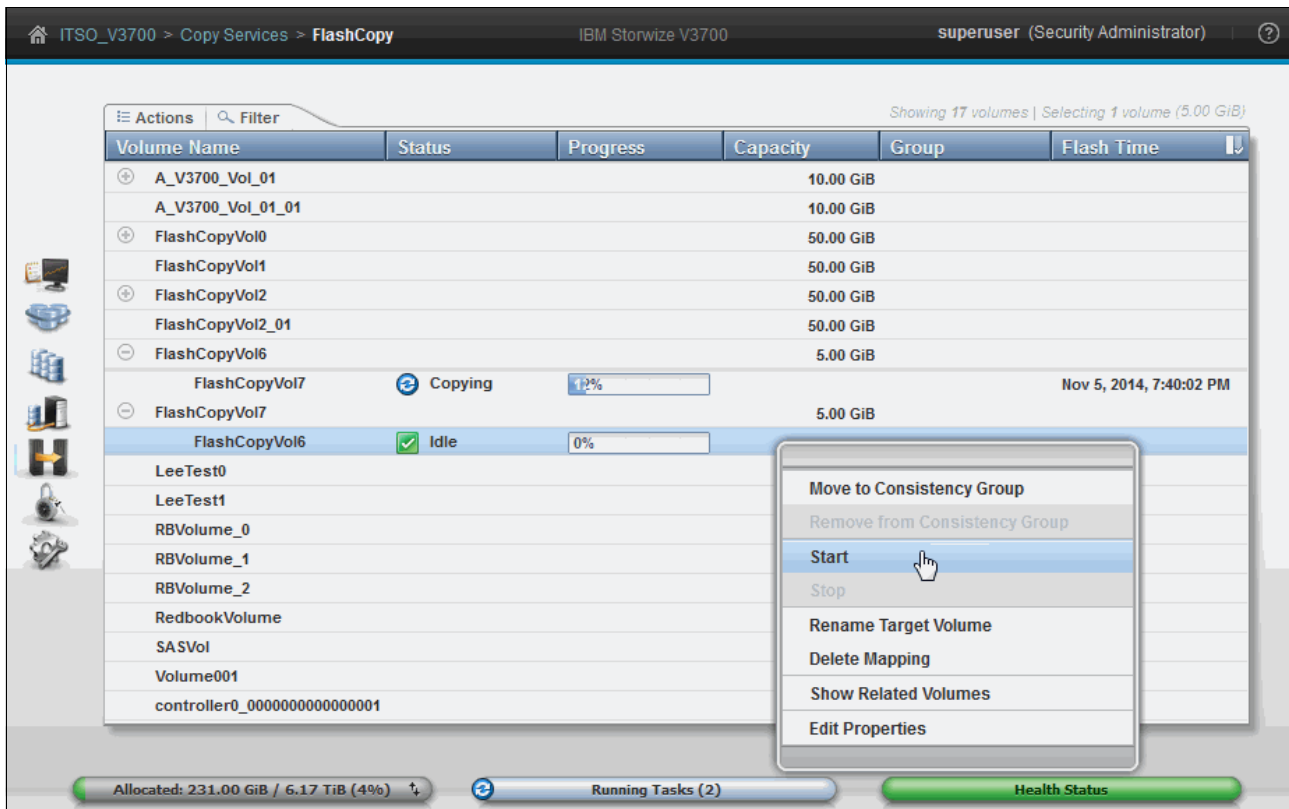


Figure 10-50 Starting the reverse mapping

10. Click **Start** to write over FlashCopyVo16 with the original bitmap data that was saved in the FlashCopy FlashCopyVo17. The command then completes, as shown in Figure 10-51.

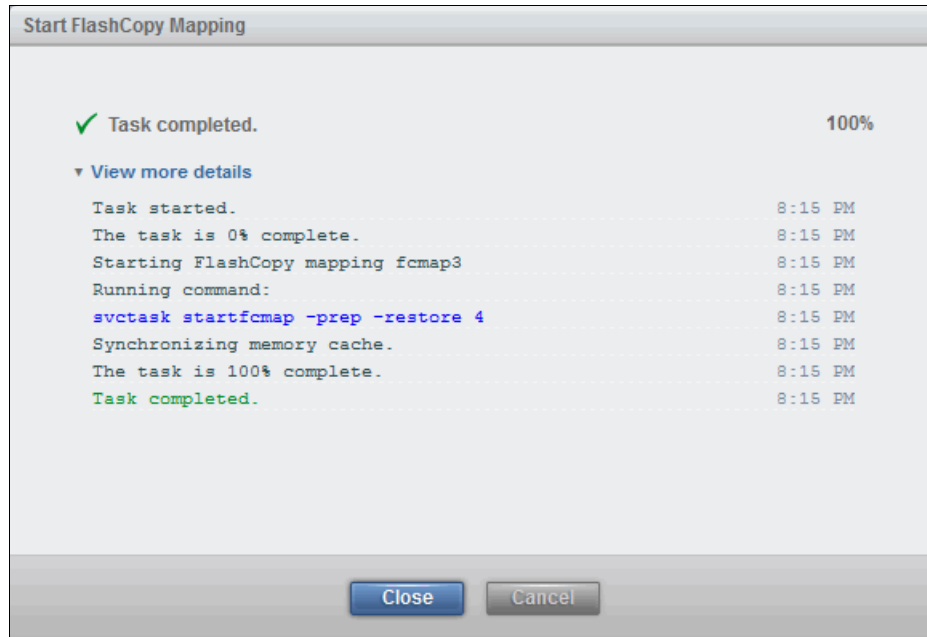


Figure 10-51 FlashCopy restore command

Important: The underlying command that is run by the IBM Storwize V3700 appends the **-restore** option automatically.

11. The reverse mapping is now reported as Copying, as shown in Figure 10-52. The progress reported is the opposite of the original mapping, because the direction is reversed. The restore process is now complete, and the original data from the snapshot is available on the source volume.

The screenshot shows the IBM Storwize V3700 FlashCopy interface. The breadcrumb navigation is 'ITSO_V3700 > Copy Services > FlashCopy'. The user is 'superuser (Security Administrator)'. The interface shows a table of 17 volumes. Two volumes are in the 'Copying' status, with progress bars indicating 12% and 87% completion. The 'Flash Time' column shows the start times for these operations: 'Nov 5, 2014, 7:40:02 PM' and 'Nov 5, 2014, 8:15:17 PM'. At the bottom, there are three status indicators: 'Allocated: 231.00 GiB / 6.17 TiB (4%)', 'Running Tasks (3)', and 'Health Status'.

Volume Name	Status	Progress	Capacity	Group	Flash Time
A_V3700_Vol_01			10.00 GiB		
A_V3700_Vol_01_01			10.00 GiB		
FlashCopyVol0			50.00 GiB		
FlashCopyVol1			50.00 GiB		
FlashCopyVol2			50.00 GiB		
FlashCopyVol2_01			50.00 GiB		
FlashCopyVol6			5.00 GiB		
FlashCopyVol7	Copying	12%			Nov 5, 2014, 7:40:02 PM
FlashCopyVol7			5.00 GiB		
FlashCopyVol6	Copying	87%			Nov 5, 2014, 8:15:17 PM
LeeTest0			25.00 GiB		
LeeTest1			25.00 GiB		
RBVolume_0			2.00 GiB		
RBVolume_1			2.00 GiB		
RBVolume_2			2.00 GiB		
RedbookVolume			1.00 GiB		
SASVol			1.00 GiB		
Volume001			5.00 GiB		
controller0_0000000000000001			1.00 GiB		

Allocated: 231.00 GiB / 6.17 TiB (4%) Running Tasks (3) Health Status

Figure 10-52 Source volume restore complete

10.1.6 Managing a FlashCopy consistency group

FlashCopy consistency groups can be managed by completing the following steps:

1. Click **Consistency Groups** under the Copy Services function icon. This displays the view shown in Figure 10-53.

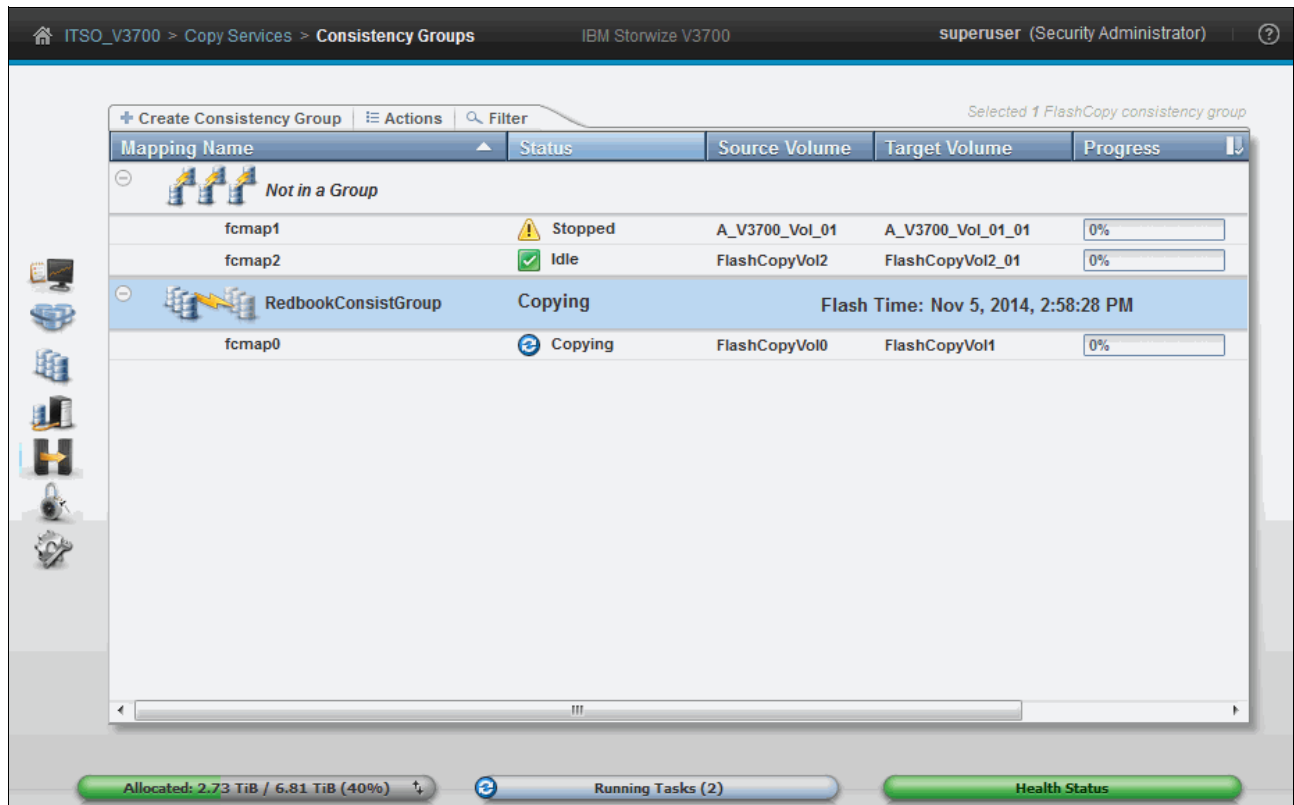


Figure 10-53 Access to the Consistency Groups window

The Consistency Groups window is where you can manage consistency groups and FlashCopy mappings within these groups.

2. All FlashCopy mappings are shown, those in consistency groups and those that are not. Click **Not in a Group**, and then expand your selection by clicking the plus sign (+) icon next to it. All of the FlashCopy mappings that are not in any consistency groups are displayed underneath.
3. Individual consistency groups are displayed underneath those not in groups, and you can discover the properties of a consistency group and the FlashCopy mappings in it by clicking the plus sign (+) to the left of the group name.

You can also take action on any consistency groups and FlashCopy mappings within the Consistency Groups window, as allowed by their state. For more information, see 10.1.5, “Managing FlashCopy mappings” on page 476.

Creating a FlashCopy consistency group

To create a FlashCopy consistency group, complete the following steps:

1. Click **Create Consistency Group** at the top of the Consistency Groups window, as shown in Figure 10-54.




Figure 10-54 New Consistency Group option

2. You are prompted to enter the name of the new consistency group. Following your naming conventions, enter the name of the new consistency group in the name field, and then click **Create**.

3. After the creation process completes, you find a new empty consistency group, as shown in Figure 10-55.

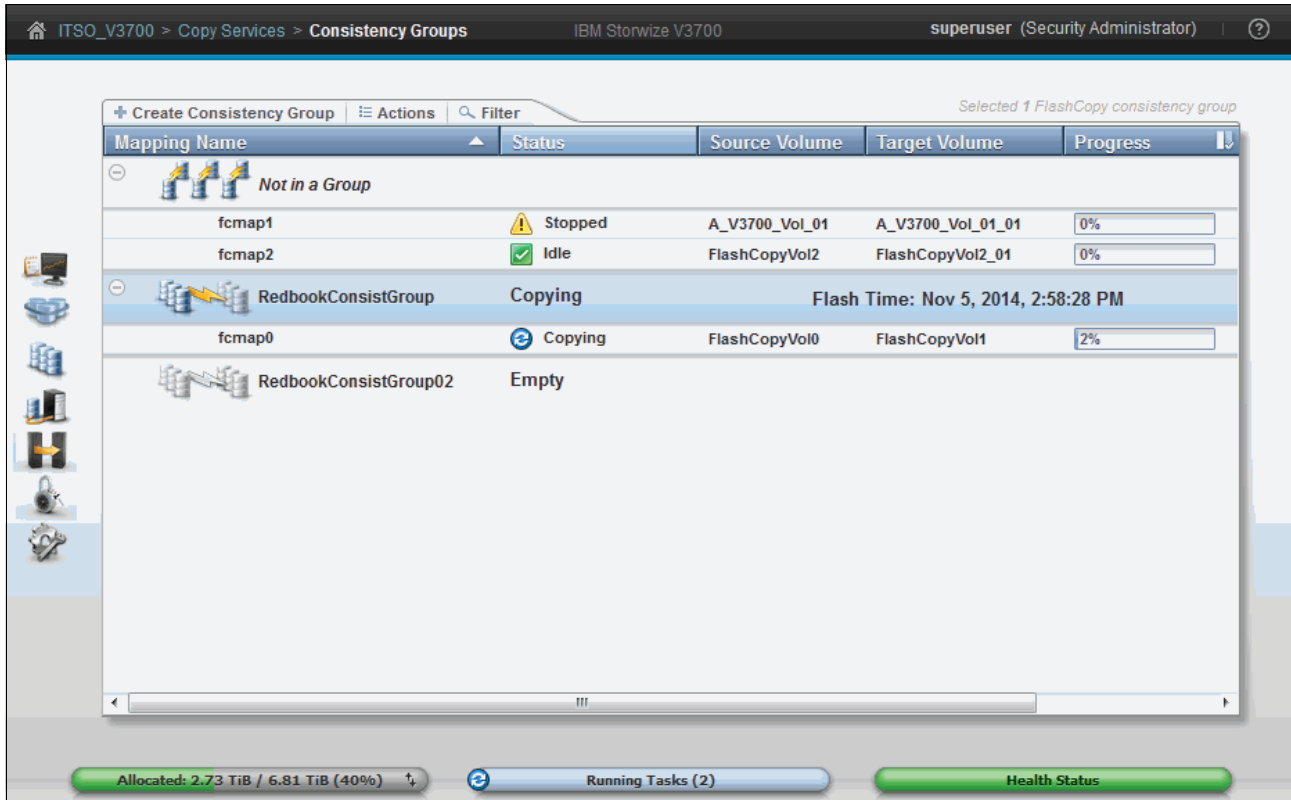


Figure 10-55 New consistency group

4. You can rename the Consistency Group by selecting it and then right-clicking, or by using the **Actions** menu. Select **Rename** and enter the new name, as shown in Figure 10-56. Next to the name of the consistency group, the state shows that it is now an empty consistency group with no FlashCopy mapping in it.

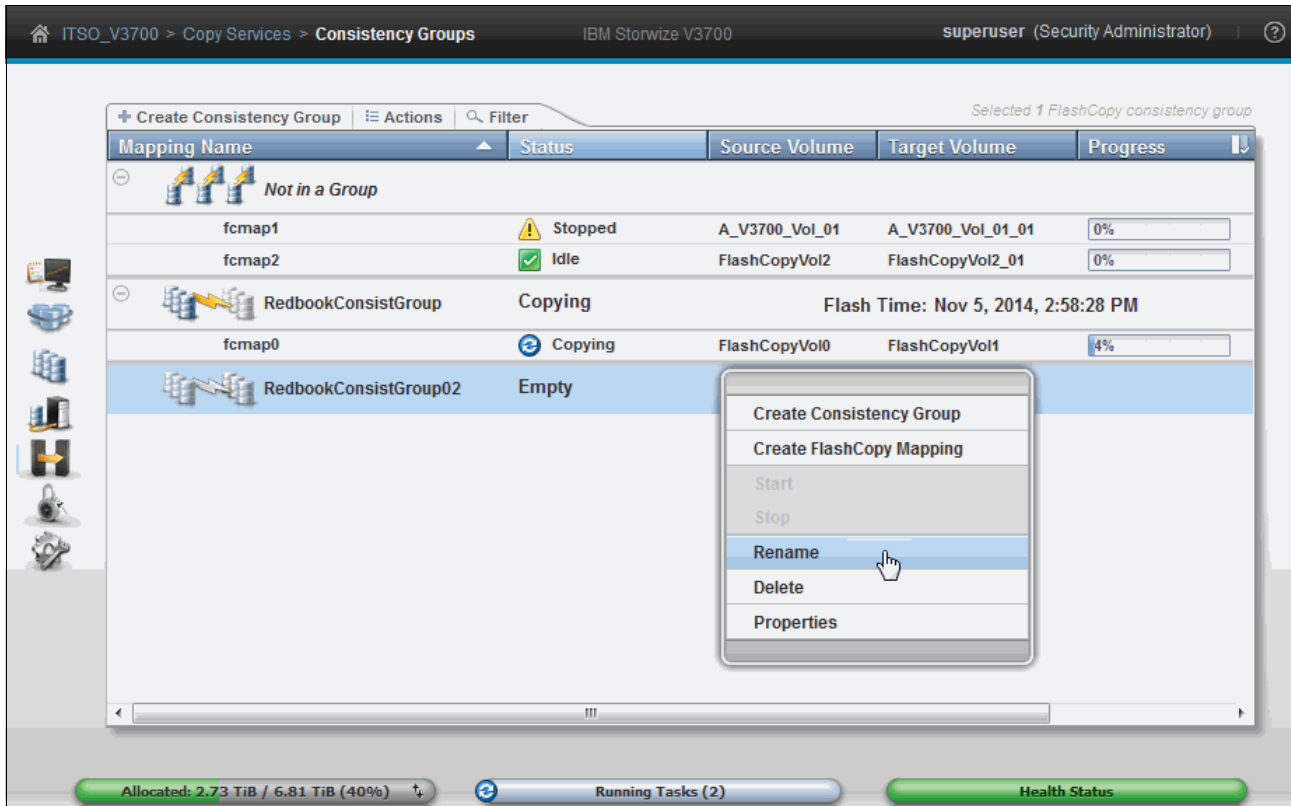


Figure 10-56 Renaming a consistency group

Adding FlashCopy mappings to a consistency group

To add FlashCopy mappings to a consistency group, complete the following steps:

1. Click **Not in a Group** to list all of the FlashCopy mappings with no consistency group.
2. You can add FlashCopy mappings to a consistency group by selecting them and clicking the **Move to Consistency Group** option from the **Actions** menu, as shown in Figure 10-57.

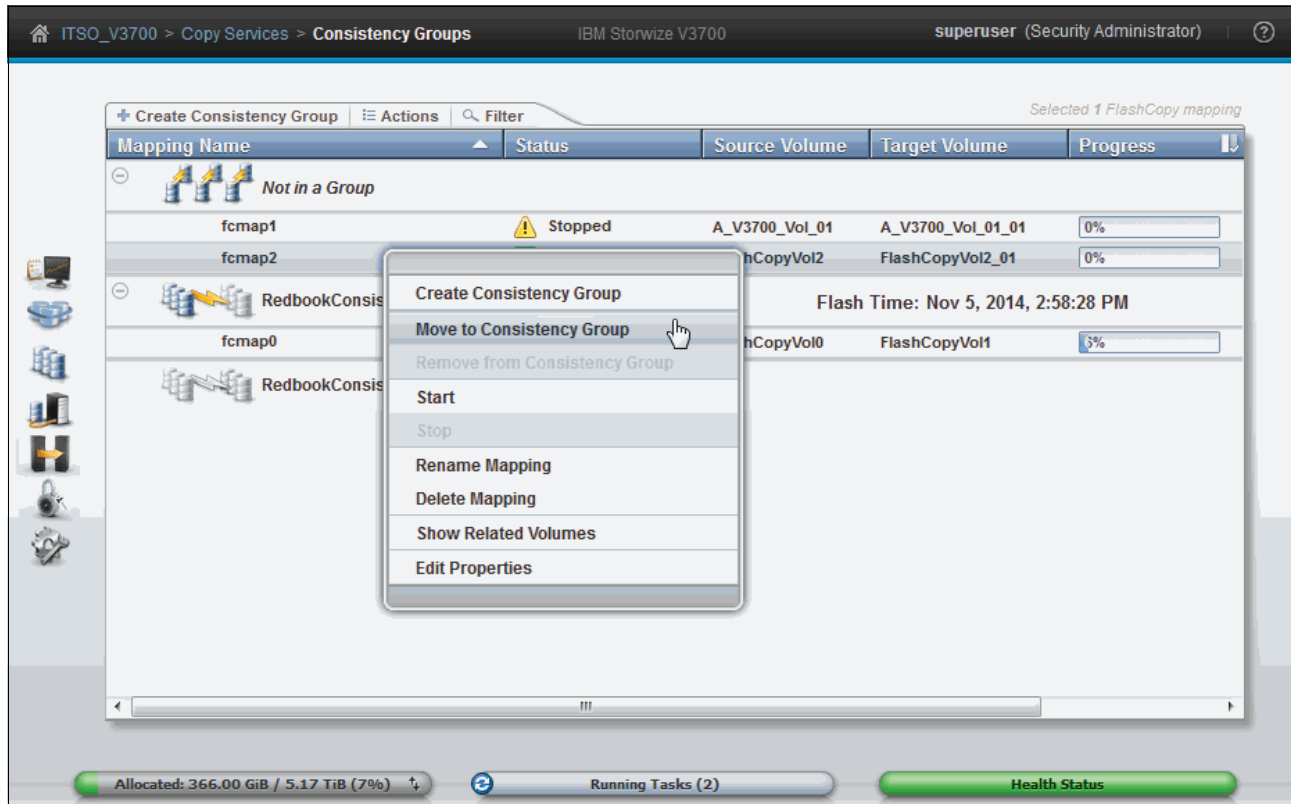


Figure 10-57 Select the FlashCopy mappings to add to a consistency group

Important: You cannot move mappings that are copying. Selecting a FlashCopy that is already running results in the Move to Consistency Group option being disabled.

3. Selections of a range are performed by highlighting a mapping, pressing and holding the Shift key, and clicking the last item in the range. Multiple, nonsequential selections can be made by pressing and holding the Ctrl key and clicking each mapping individually. The option is also available by right-clicking individual mappings.

- You are prompted to specify which consistency group you want to move the FlashCopy mapping into, as shown in Figure 10-58. Choose from the list in the drop-down menu. Click **Move to Consistency Group** to continue.



Figure 10-58 Select consistency group

After the action completes, you find that the FlashCopy mappings you selected were removed from the Not In a Group list to the consistency group that you chose.

Starting a consistency group

To start a consistency group, complete the following steps:

- Highlight the required group and click **Start** from the **Actions** menu, or right-click the consistency group, as shown in Figure 10-59.

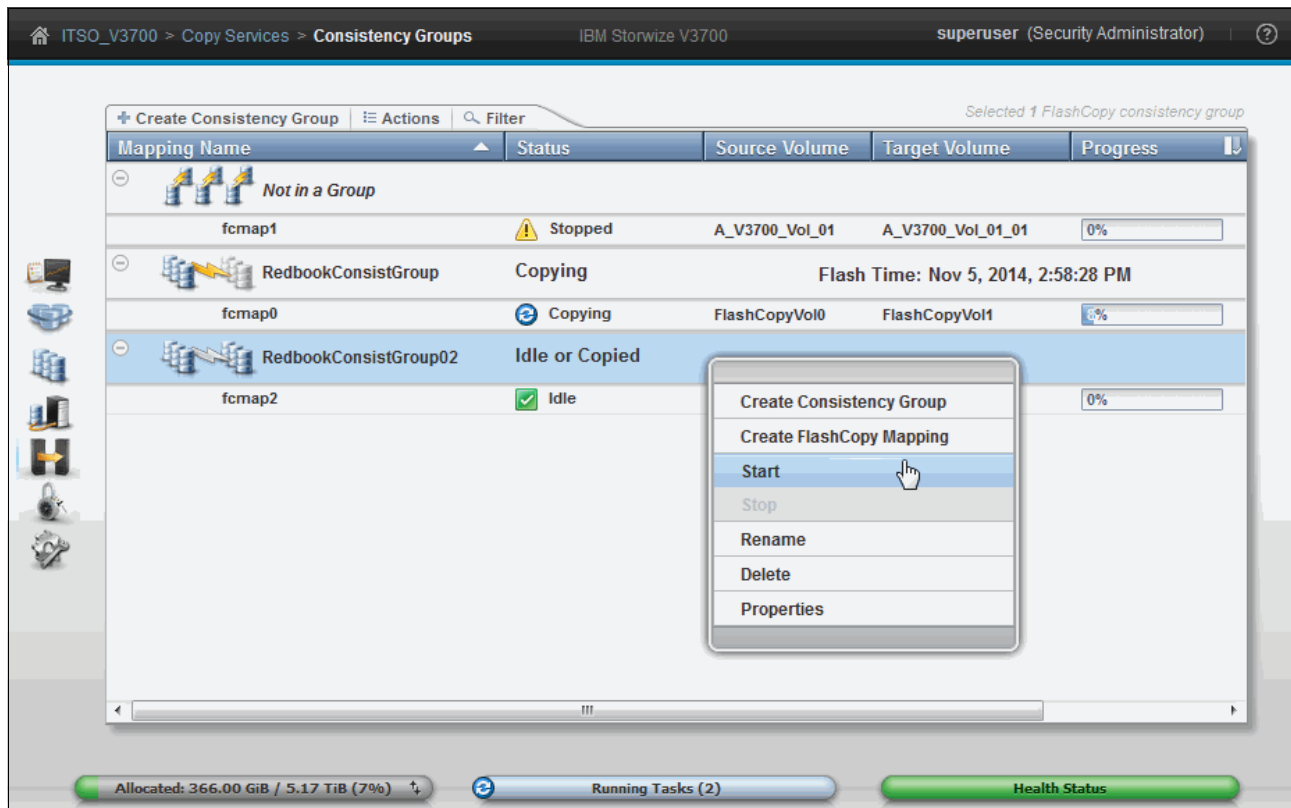


Figure 10-59 Start a consistency group

- After you start the consistency group, all of the FlashCopy mappings in the consistency group start at the same time. The state of the consistency group and all of the underlying mappings changes to Copying, as shown in Figure 10-60.




Figure 10-60 Consistency group started

Stopping a consistency group

The consistency group can be stopped by completing the following steps:

1. Select **Stop** from the **Actions** menu or right-click, as shown in Figure 10-61.

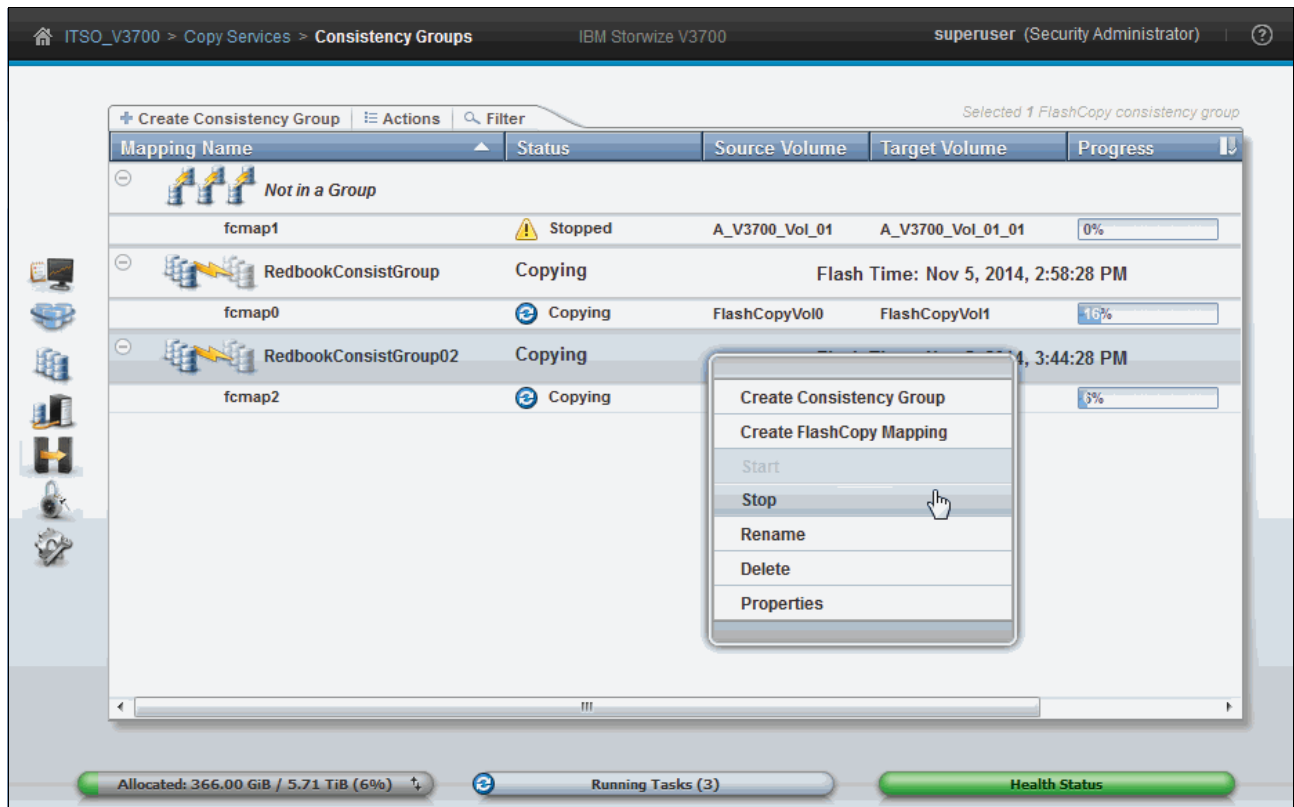


Figure 10-61 Stop a consistency group

- After the stop process completes, the FlashCopy mappings in the consistency group are in the Stopped state, and a red X icon displays on the function icon of this consistency group to indicate an alert, as shown in Figure 10-62.




Figure 10-62 Consistency group stop completes

Any previously copied relationships that were added to a consistency group that was later stopped before all members of the consistency group completed synchronization remain in the Copied state.

Removing FlashCopy mappings from a consistency group

FlashCopy mappings can be removed from a consistency group by completing the following steps:

1. Select the FlashCopy mappings.
2. Click **Remove from Consistency Group** from the **Actions** drop-down menu of the FlashCopy mapping, or right-click, as shown in Figure 10-63.

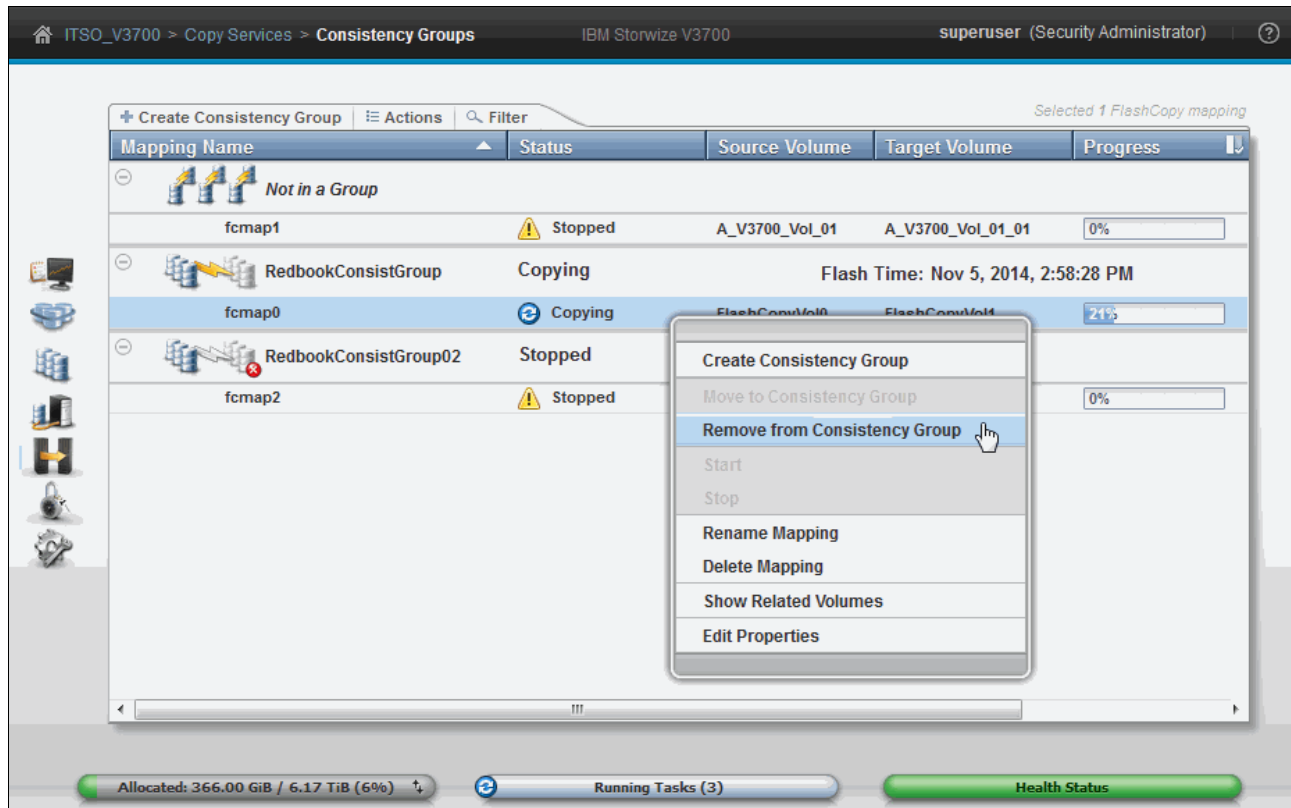


Figure 10-63 Remove from consistency group

The FlashCopy mappings are returned to the Not in a Group list after they are removed from the consistency group.

Deleting a consistency group

A consistency group can be deleted by completing the following steps:

1. Click **Delete** from the **Actions** menu, or right-click the selected group, as shown in Figure 10-64.

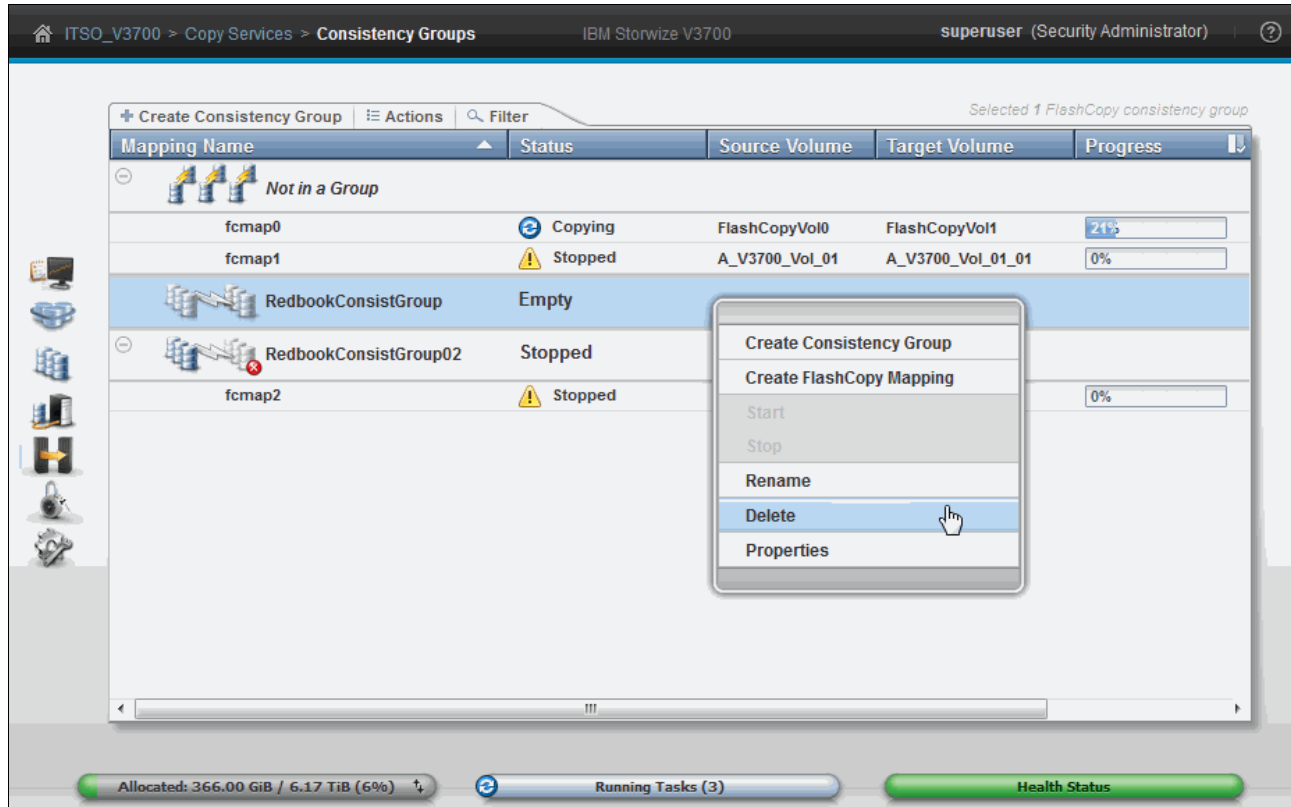


Figure 10-64 Delete a consistency group

2. You are presented with a confirmation pane as shown in Figure 10-65.

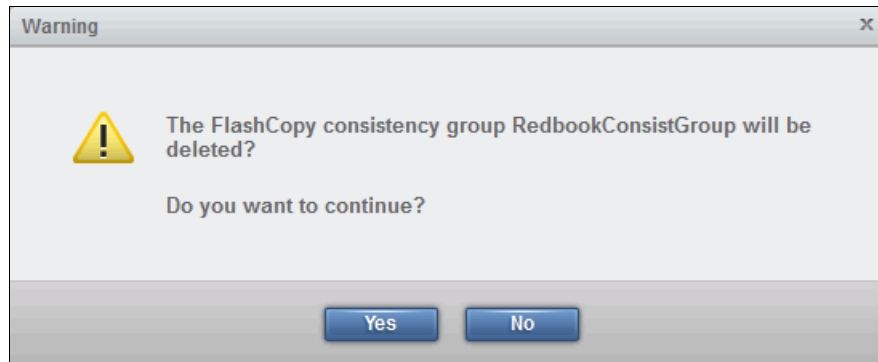


Figure 10-65 Confirming deletion of a consistency group

3. Click **Yes** to delete the group. Consistency groups do not have to be empty to be deleted. Any FlashCopy mappings that exist in them are moved to the Not in a Group mappings list when the consistency group is deleted.

Restoring from a FlashCopy Consistency Group

It is possible to manipulate FlashCopy snapshots that were captured as part of a consistency group to restore the source volumes of those mappings to the state that they were all in at the time the snapshot was taken:

1. To restore a consistency group from a FlashCopy, you must create a reverse mapping of all of the individual volumes that are contained within the original consistency group. In this example, there are two FlashCopy mappings (fcmap0 and fcmap1) in a consistency group called RedbookConsistGroup, as shown in Figure 10-66. These mappings are both snapshots.




Figure 10-66 Creating FlashCopy reverse mapping

2. Click **New Consistency Group** in the upper left corner (as shown in Figure 10-66) and create a consistency group. In this example, we created a group called RedBookReverseConsistGroup.
3. Follow the procedure that is described in “Restoring from a FlashCopy” on page 495 to create reverse mappings for each of the mappings that exist in the source consistency group (RedbookConsistGroup). When prompted to add to a consistency group, select **Yes**.

- From the drop-down menu, select the new “reverse” consistency group that you created in step 3 on page 513. In our example, this group is RedBookReverseConsistencyGroup. The result should be similar to that shown in Figure 10-67.

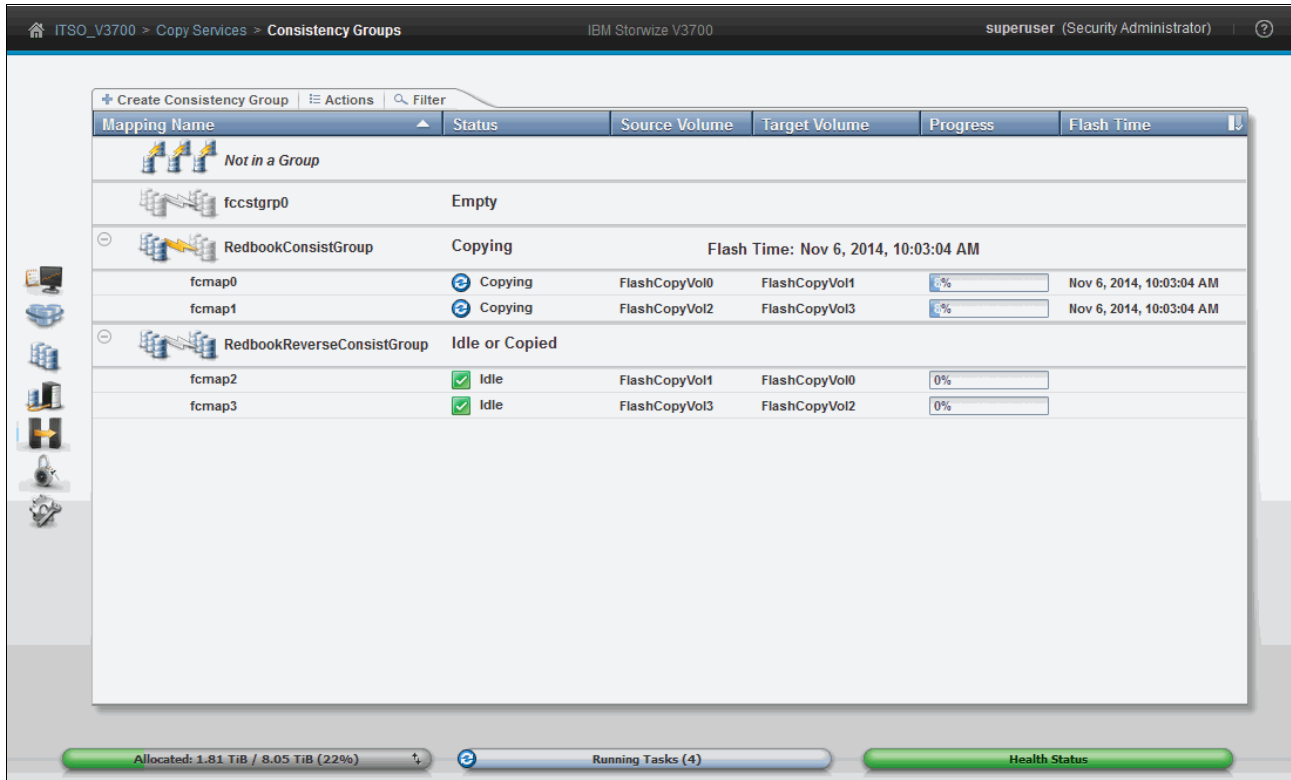


Figure 10-67 Reverse Consistency group populated

- To restore the consistency group, highlight the reverse consistency group and click **Start**, as shown in Figure 10-68.

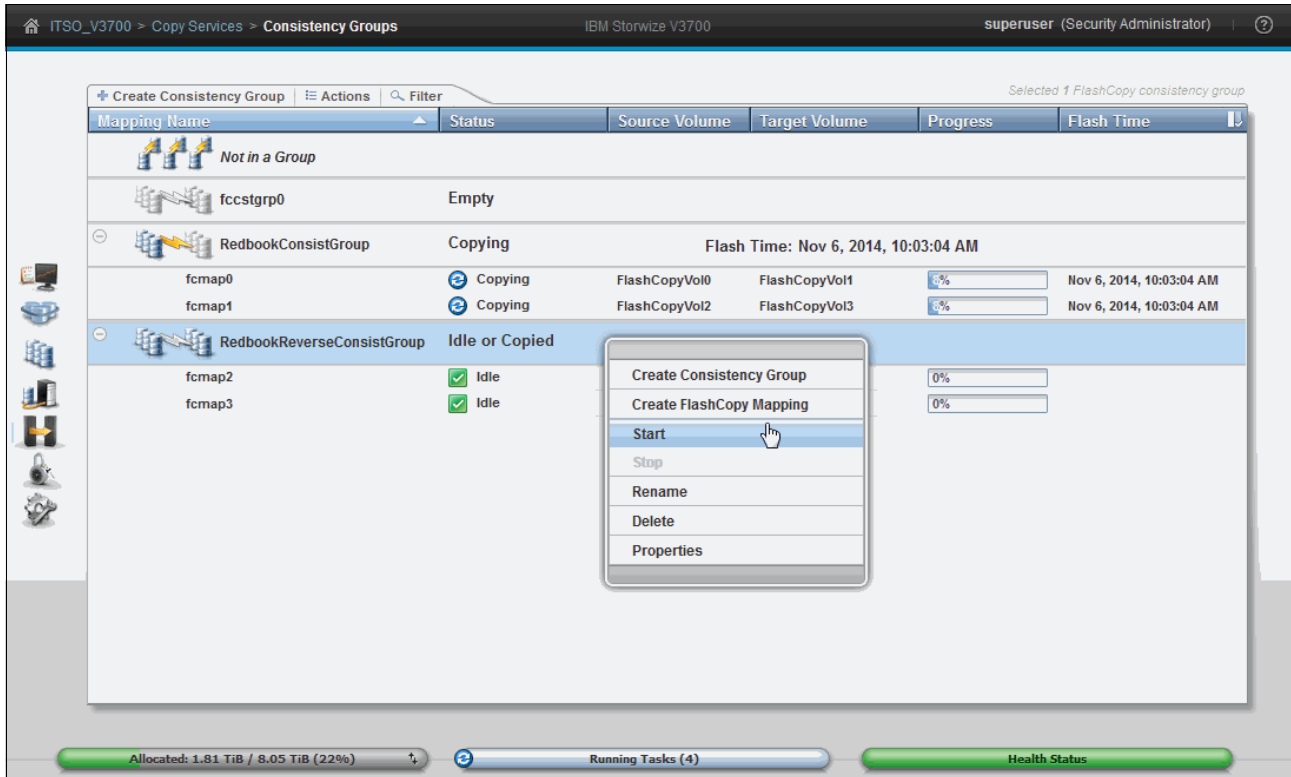


Figure 10-68 Starting Consistency group restore

- This overwrites the source volumes with the original bitmap data that was saved in the snapshots. The command completes, as shown in Figure 10-69.

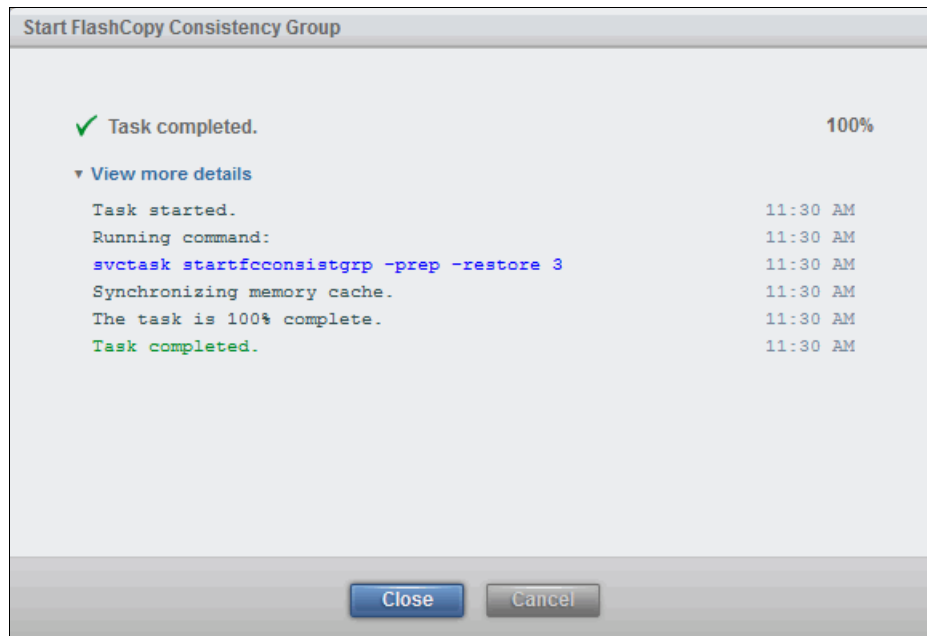


Figure 10-69 Consistency Group restore command

Important: The IBM Storwize V3700 automatically appends the **-restore** option to the command.

7. Click **Close** and the command pane returns to the Consistency Group window. The reverse consistency group is now reported as Copying, as shown in Figure 10-70. The progress reported is the opposite of the original mappings, because the direction is reversed. The restore process is now complete, and the original data from the snapshots is available on the source volumes.

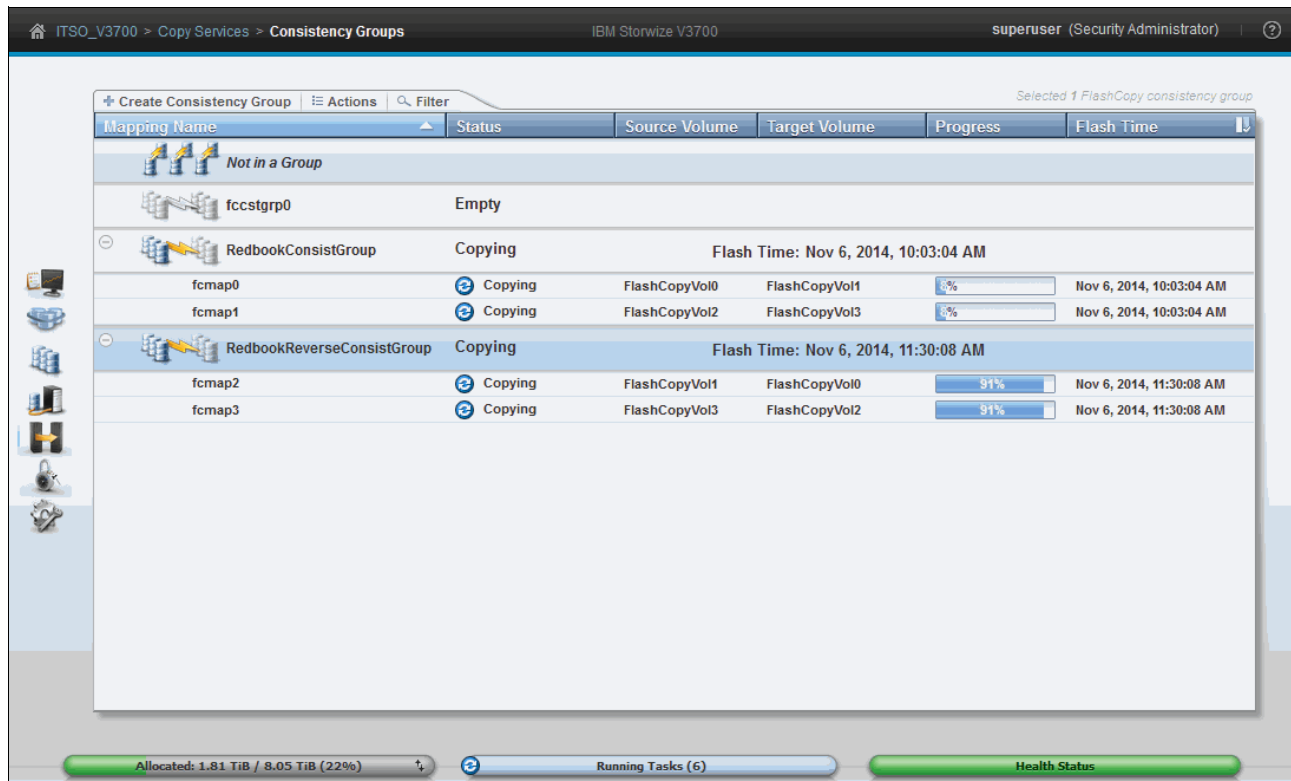


Figure 10-70 Consistency Group restored

10.2 Remote Copy

In this section, we describe how the Remote Copy function works in IBM Storwize V3700. We also provide the implementation steps for Remote Copy configuration and management using the GUI.

Remote Copy consists of three methods for copying: IBM Metro Mirror, IBM Global Mirror, and Global Mirror with Change Volumes. Metro Mirror is designed for metropolitan distances with a synchronous copy requirement. Global Mirror is designed for longer distances without requiring the hosts to wait for the full round-trip delay of the long-distance link through asynchronous methodology. Global Mirror with Change Volumes is an added functionality for Global Mirror that is designed to attain consistency on lower-quality network links.

Metro Mirror and Global Mirror are IBM branded terms for the functions of *synchronous remote copy* and *asynchronous remote copy*. Throughout this book, the term *Remote Copy* is used to refer to both functions where the text applies to each term equally.

10.2.1 Remote Copy concepts

Remote Copy concepts are described in this section.

Partnership

When a partnership is created, you connect two separate IBM Storwize V3700 systems, or an IBM SAN Volume Controller, IBM Storwize V5000, or IBM Storwize V7000, and an IBM Storwize V3700. After the partnership creation is configured on both systems, further communication between the node canisters in each of the storage systems is established and maintained by the storage area network (SAN). All inter-cluster communication goes through the Fibre Channel (FC) network.

The partnership must be defined on both participating Storwize or IBM SAN Volume Controller systems to make the partnership fully functional.

Interconnection: Interconnects between IBM Storwize products were introduced in version 6.3.0. Because IBM Storwize V3700 supports only version 6.4.1 or higher, there is no problem with support for this functionality. However, any other Storwize product must be at a minimum level of 6.3.0 to connect to the IBM Storwize V3700 and the IBM Storwize V3700 must set the replication layer using the `svctask chsystem -layer replication` command. Layer limitations are described next.

Introduction to layers

IBM Storwize V3700 implements the concept of *layers*. Layers determine how the IBM Storwize portfolio interacts with the IBM SAN Volume Controller. Currently, there are two layers: Replication and storage.

The *replication layer* is used when you want to use the IBM Storwize V3700 with one or more IBM SAN Volume Controllers as a Remote Copy partner. The *storage layer* is the default mode of operation for the IBM Storwize V3700, and is used when you want to use the IBM Storwize V3700 to present storage to an IBM SAN Volume Controller as a back-end system.

The layer for the IBM Storwize V3700 can be switched by running `svctask chsystem -layer replication`. Generally, switch the layer while your IBM Storwize V3700 system is not in production. This situation prevents potential disruptions, because layer changes are not I/O-tolerant.

Figure 10-71 shows the effect of layers on IBM SAN Volume Controller and IBM Storwize V3700 partnerships.

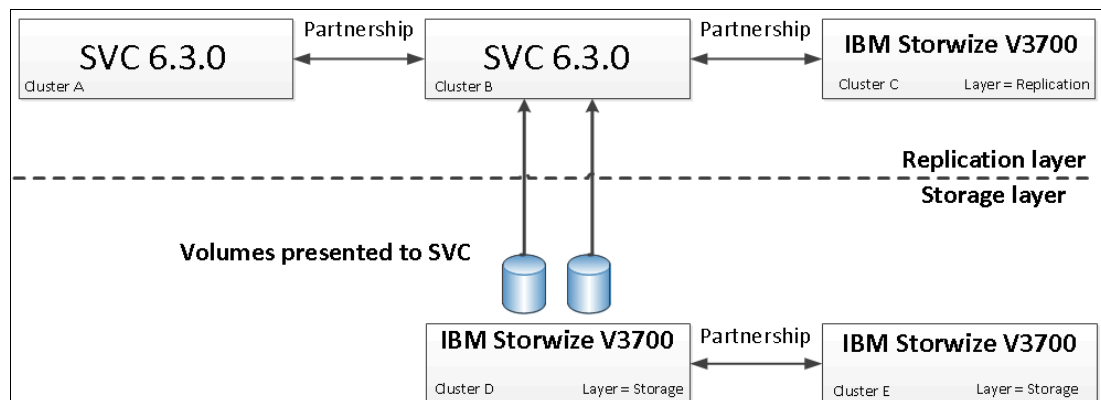


Figure 10-71 IBM Storwize V3700 virtualization layers

The replication layer enables an IBM Storwize V3700 system to be a Remote Copy partner with an IBM SAN Volume Controller. The storage layer enables an IBM Storwize V3700 system to function as back-end storage for an IBM SAN Volume Controller. An IBM Storwize V3700 system cannot be in both layers at the same time.

Limitations on the IBM SAN Volume Controller and IBM Storwize V3700 partnership

IBM SAN Volume Controller and IBM Storwize V3700 systems can be partners in a Remote Copy partnership. However, the following limitations apply:

- ▶ The layer for the Storwize V3700 must be set to replication. The default is storage.
- ▶ If any other SAN Volume Controller or Storwize ports are visible on the SAN (aside from the ones on the cluster where you are making the change), you cannot change the layer. The SAN Volume Controller and IBM Storwize V3700 will need to be zoned together after changing the layer on the IBM Storwize V3700.
- ▶ If any cluster partnership is already defined on the IBM Storwize V3700, you cannot change the layer.

Partnership topologies

A partnership between up to four IBM Storwize V3700 systems is allowed.

The following typical partnership topologies between multiple IBM Storwize V3700 systems are available:

- ▶ Daisy-chain topology, as shown in Figure 10-72.




Figure 10-72 Daisy chain partnership topology for IBM Storwize V3700

- ▶ Triangle topology, as shown in Figure 10-73.

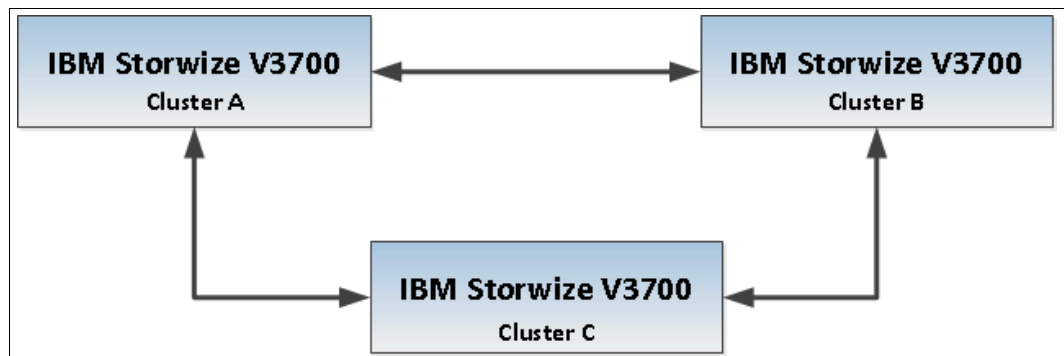


Figure 10-73 Triangle partnership topology for IBM Storwize V3700

- ▶ Star topology, as shown in Figure 10-74.




Figure 10-74 Star topology for IBM Storwize V3700

- ▶ Full-meshed topology, as shown in Figure 10-75.




Figure 10-75 Full meshed IBM Storwize V3700

Partnerships: These partnerships are valid for configurations with IBM SAN Volume Controllers and IBM Storwize V3700 systems if the IBM Storwize V3700 systems are using the replication layer. They are also valid for IBM Storwize V5000 and V7000 products.

Partnership states

A partnership has the following states:

- ▶ Partially Configured

Indicates that only one cluster partner is defined from a local or remote cluster to the displayed cluster and is started. For the displayed cluster to be configured fully, and to complete the partnership, you must define the cluster partnership from the cluster that is displayed to the corresponding local or remote cluster.

- ▶ Fully Configured

Indicates that the partnership is defined on the local and remote clusters, and is started.

- ▶ **Remote Not Present**
Indicates that the remote cluster is not present for the partnership.
- ▶ **Partially Configured (Local Stopped)**
Indicates that the local cluster is only defined to the remote cluster, and the local cluster is stopped.
- ▶ **Fully Configured (Local Stopped)**
Indicates that a partnership is defined on the local and remote clusters, and the remote cluster is present, but the local cluster is stopped.
- ▶ **Fully Configured (Remote Stopped)**
Indicates that a partnership is defined on the local and remote clusters, and the remote cluster is present, but the remote cluster is stopped.
- ▶ **Fully Configured (Local Excluded)**
Indicates that a partnership is defined between a local and remote cluster. However, the local cluster was excluded. This state can occur when the fabric link between the two clusters was compromised by too many fabric errors, or by slow response times of the cluster partnership.
- ▶ **Fully Configured (Remote Excluded)**
Indicates that a partnership is defined between a local and remote cluster. However, the remote cluster was excluded. This state can occur when the fabric link between the two clusters was compromised by too many fabric errors, or by slow response times of the cluster partnership.
- ▶ **Fully Configured (Remote Exceeded)**
Indicates that a partnership is defined between a local and remote cluster, and the remote is available. However, the remote cluster exceeds the number of allowed clusters within a cluster network. A maximum of four clusters can be defined in a network.

If the number of clusters exceeds that limit, the IBM Storwize V3700 system determines the inactive cluster or clusters by sorting all of the clusters by their unique identifier in numerical order. The inactive cluster partner that is not in the top four of the cluster-unique identifiers shows Fully Configured (Remote Exceeded).

Remote Copy relationships

A Remote Copy relationship is a relationship between two individual volumes of the same size. These volumes are called a *master (source) volume* and an *auxiliary (target) volume*. Typically, the master volume contains the production copy of the data, and is the volume that the application normally accesses. The auxiliary volume often contains a backup copy of the data, and is used for disaster recovery (DR).

The master and auxiliary volumes are defined when the relationship is created, and these attributes never change. However, either volume can operate in the primary or secondary role as necessary. The primary volume contains a valid copy of the application data, and receives updates from the host application, which is analogous to a source volume. The secondary volume receives a copy of any updates to the primary volume, because these updates are all transmitted across the mirror link.

Therefore, the secondary volume is analogous to a continuously updated target volume. When a relationship is created, the master volume is assigned the role of primary volume, and the auxiliary volume is assigned the role of secondary volume. The initial copying direction is from master to auxiliary. When the relationship is in a consistent state, you can reverse the copy direction.

The two volumes in a relationship must be the same size. The Remote Copy relationship can be established on the volumes within one IBM Storwize V3700 storage system, which is called an *intra-cluster relationship*. The relationship can also be established in different IBM Storwize V3700 storage systems, or between an IBM Storwize V3700 storage system and an IBM SAN Volume Controller, IBM Storwize V5000, or IBM Storwize V7000, which are called *inter-cluster relationships*.

Important: The use of Remote Copy target volumes as Remote Copy source volumes is not allowed. A FlashCopy target volume can be used as a Remote Copy source volume, and also as a Remote Copy target volume.

Metro Mirror

Metro Mirror is a type of Remote Copy that creates a synchronous copy of data from a master volume to an auxiliary volume. With synchronous copies, host applications write to the master volume, but do not receive confirmation that the write operation completed until the data is written to the auxiliary volume. This action ensures that both volumes have identical data when the copy completes. After the initial copy completes, the Metro Mirror function always maintains a fully synchronized copy of the source data at the target site.

Figure 10-76 shows how a write to the master volume is mirrored to the cache of the auxiliary volume before an acknowledgment of the write is sent back to the host that issued the write. This process ensures that the auxiliary is synchronized in real time if it is needed in a failover situation.

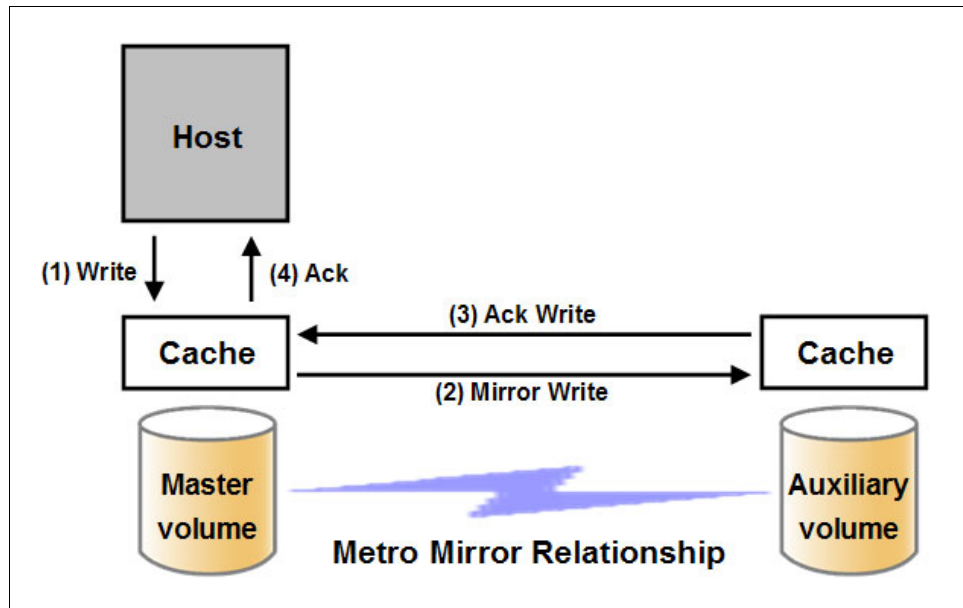


Figure 10-76 Write on volume in a Metro Mirror relationship

The Metro Mirror function supports copy operations between volumes that are separated by distances up to 300 kilometers (km). For disaster recovery purposes, Metro Mirror provides the simplest way to maintain an identical copy on the primary and secondary volumes.

However, as with all synchronous copies over remote distances, there can be a performance effect to host applications. This performance effect is related to the distance between primary and secondary volumes and, depending on application requirements, its use might be limited based on the distance between sites.

Global Mirror

Global Mirror provides an asynchronous copy, which means that the secondary volume is not an exact match of the primary volume at every point. The Global Mirror function provides the same function as Metro Mirror Remote Copy, without requiring the hosts to wait for the full round-trip delay of the long-distance link. However, some delay can be seen on the hosts in congested or overloaded environments. Make sure that you closely monitor and understand your workload.

In asynchronous Remote Copy (which Global Mirror provides), write operations are completed on the primary site, and the write acknowledgment is sent to the host before it is received at the secondary site. An update of this write operation is sent to the secondary site at a later stage, which provides the capability to perform Remote Copy over distances that exceed the limitations of synchronous Remote Copy. The distance of Global Mirror replication is limited primarily by the latency of the wide area network (WAN) link that is provided.

Global Mirror requires a 80 millisecond (ms) round-trip-time for data that is sent to the remote location. The propagation delay is roughly 8.2 microseconds (μs) per mile, or 5 μs per kilometer for FC connections. Each device in the path adds more delay of about 25 μs . Devices that use software (such as some compression devices) add much more time. The time that is added by software-assisted devices is highly variable, and should be measured directly. Be sure to include these times when you are planning your Global Mirror design.

You should also measure application performance based on the expected delays before Global Mirror is fully implemented. The IBM Storwize V3700 storage system provides you with an advanced feature that enables you to test performance implications *before* Global Mirror is deployed and a long-distance link is obtained. This advanced feature is enabled by modifying the IBM Storwize V3700 storage system `gmintradelaysimulation` and `gminterdelaysimulation` parameters.

These parameters can be used to simulate the write delay to the secondary volume. The delay simulation can be enabled separately for each intra-cluster or inter-cluster Global Mirror. You can use this feature to test an application before the full deployment of the Global Mirror feature. Figure 10-77 shows that a write operation to the master volume is acknowledged back to the host before the write operation is mirrored to the cache for the auxiliary volume.

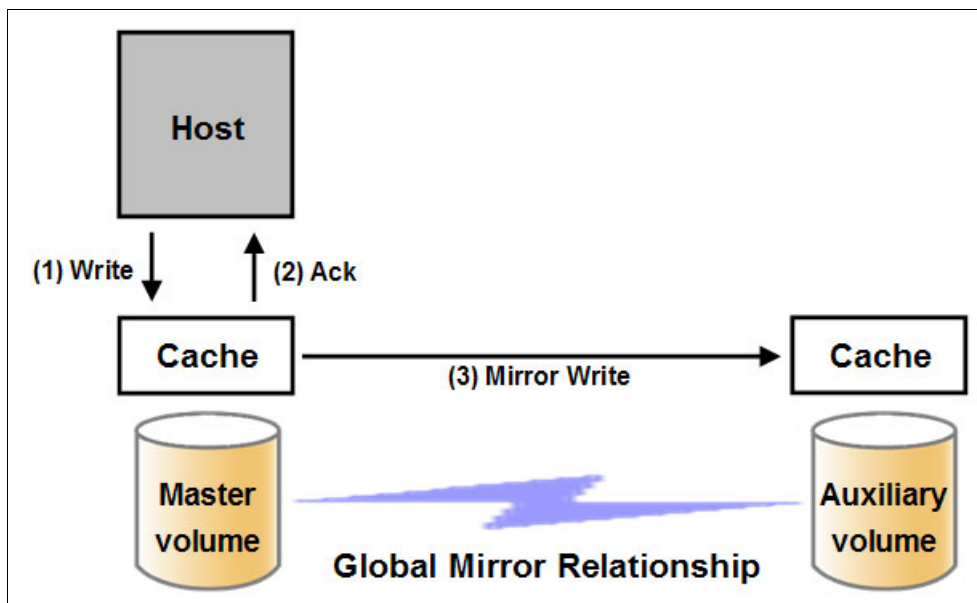


Figure 10-77 Global Mirror write sequence

The Global Mirror algorithms always maintain a consistent image on the auxiliary volume. They achieve this consistent image by identifying sets of I/Os that are active concurrently at the master, assigning an order to those sets, and applying those sets of I/Os in the assigned order at the secondary.

In a failover scenario where the secondary site must become the master source of data (depending on the workload pattern and the bandwidth and distance between local and remote site), certain updates might be missing at the secondary site. Therefore, any applications that use this data must have an external mechanism for recovering the missing updates and reapplying them (for example, a transaction log replay).

10.2.2 Global Mirror with Change Volumes

Global Mirror within the IBM Storwize V3700 is designed to achieve a recovery point objective (RPO) as low as possible so that data is as up-to-date as possible. This capability places some strict requirements on your infrastructure and, in certain situations (with low network link quality, congested hosts, or overloaded hosts), you might be affected by multiple 1920 (congestion) errors.

Congestion errors happen in the following primary situations:

- ▶ Congestion at the source site through the host or network
- ▶ Congestion in the network link or network path
- ▶ Congestion at the target site through the host or network

Global Mirror includes functionality that is designed to address the following conditions that negatively affect some Global Mirror implementations:

- ▶ Estimation of bandwidth requirements tends to be complex.
- ▶ It is often difficult to ensure that the latency and bandwidth requirements can be met.
- ▶ Congested hosts on the source or target site can cause disruption.
- ▶ Congested network links can cause disruption with only intermittent peaks.

To address these issues, *Change Volumes* were added as an option for Global Mirror relationships. Change Volumes use the FlashCopy functionality, but cannot be manipulated as FlashCopy volumes because they are special-purpose only. Change Volumes replicate point-in-time images on a cycling period (the default is 300 seconds). This situation means that your change rate must include only the condition of the data at the point-in-time that the image was taken, rather than all of the updates during the period. This situation can provide significant reductions in replication volume.

Figure 10-78 shows a basic Global Mirror relationship without Change Volumes.

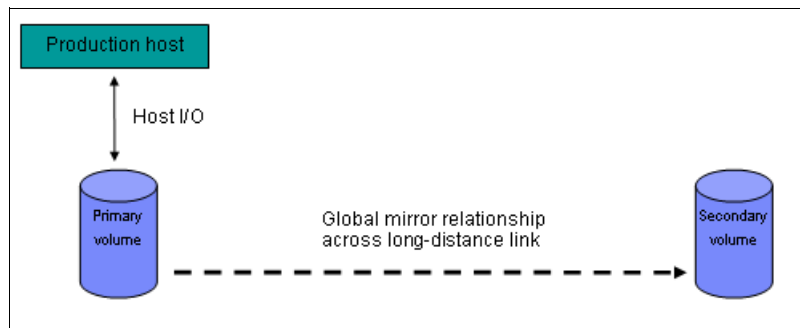


Figure 10-78 Global Mirror without Change Volumes

Figure 10-79 shows a Global Mirror relationship with Change Volumes.

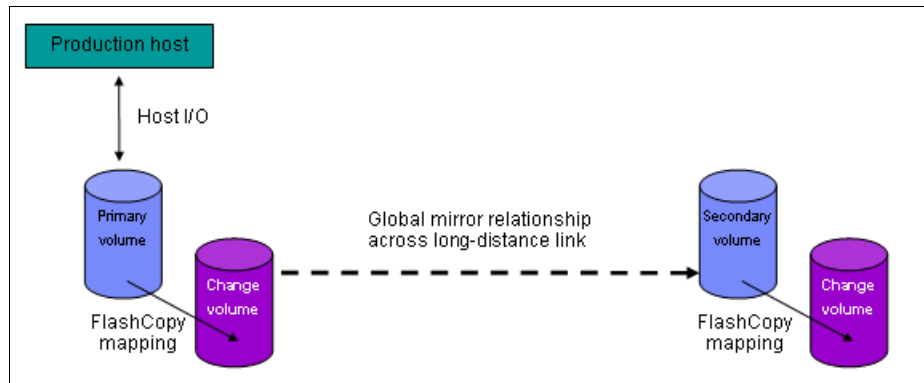


Figure 10-79 Global Mirror with Change Volumes

With Change Volumes, a FlashCopy mapping exists between the primary volume and the primary Change Volume. The mapping is updated during a cycling period (60 seconds - 24 hours). The primary Change Volume is then replicated to the secondary Global Mirror volume at the target site, which is then captured in another change volume on the target site. This situation provides a consistent image at the target site, and protects your data from being inconsistent during resynchronization.

Figure 10-80 shows several I/Os on the source volume, the same number on the target volume, and in the same order. Assuming that this set is the same set of data updated over and over, these updates are wasted network traffic, and the I/O can be completed much more efficiently, as shown in Figure 10-81 on page 525.




Figure 10-80 Global Mirror I/O replication without Change Volumes

In Figure 10-81, the same data is being updated repeatedly, so Change Volumes demonstrate significant I/O transmission savings because you must send only I/O number 16, which was the last I/O before the cycling period.

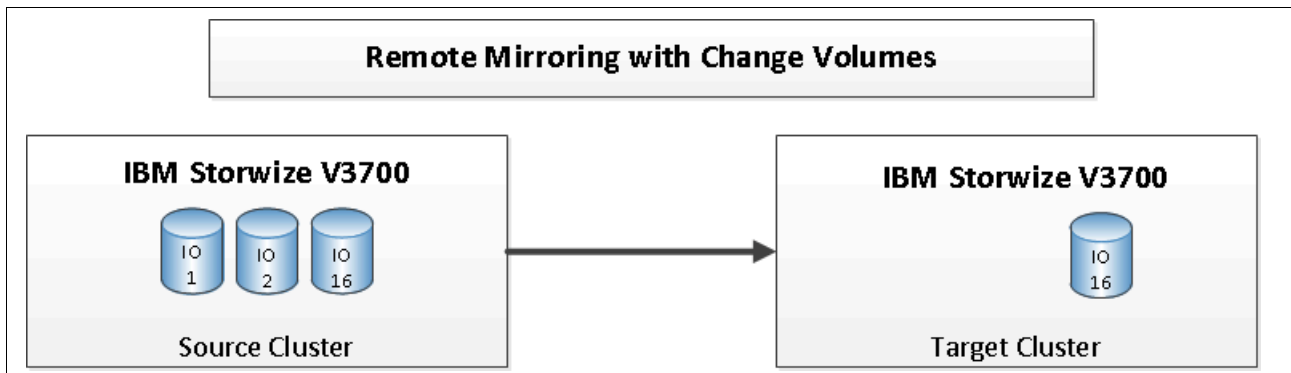


Figure 10-81 Global Mirror I/O replication with Change Volumes

The cycling period can be adjusted by running the `chrcrelationship -cycleperiodseconds <60-86400>` command. If a copy does not complete in the cycle period, the next cycle does not start until the prior cycle completes. It is for this reason that the use of Change Volumes gives you the following possibilities for RPO:

- ▶ If your replication completes in the cycling period, your RPO is twice the cycling period.
- ▶ If your replication does not complete within the cycling period, your RPO is twice the completion time. The next cycling period starts immediately after the prior period is finished.

Careful consideration should be given to balancing your business requirements with the performance of Global Mirror with Change Volumes. Global Mirror with Change Volumes increases the inter-cluster traffic for more frequent cycling periods, so going as short as possible is not always the answer. In most cases, the default should meet your requirements and perform reasonably well.

Important: When Global Mirror volumes with Change Volumes are used, make sure that you remember to select the Change Volume on the auxiliary (target) site. Failure to do so leaves you exposed during a resynchronization operation.

The GUI automatically creates Change Volumes for you. However, it is a limitation of this initial release that they are fully provisioned volumes. To save space, you should create thin-provisioned volumes in advance, and use the Existing Volume option to select your Change Volumes.

Remote Copy consistency groups

A consistency group is a logical entity that groups copy relationships. By grouping the relationships, you can ensure that these relationships are managed in unison, and that the data within the group is in a consistent state. For more information about the necessity of consistency groups, see 10.1.6, “Managing a FlashCopy consistency group” on page 502.

Remote Copy commands can be issued to a Remote Copy consistency group. Therefore, the commands can be issued simultaneously for all Metro Mirror relationships that are defined within that consistency group, or to a single Metro Mirror relationship that is not part of a Metro Mirror consistency group.

Figure 10-82 shows the concept of Remote Copy consistency groups. Because the RC_Relationship1 and RC_Relationship2 are part of the consistency group, they can be handled as one entity, while the stand-alone RC_Relationship3 is handled separately.

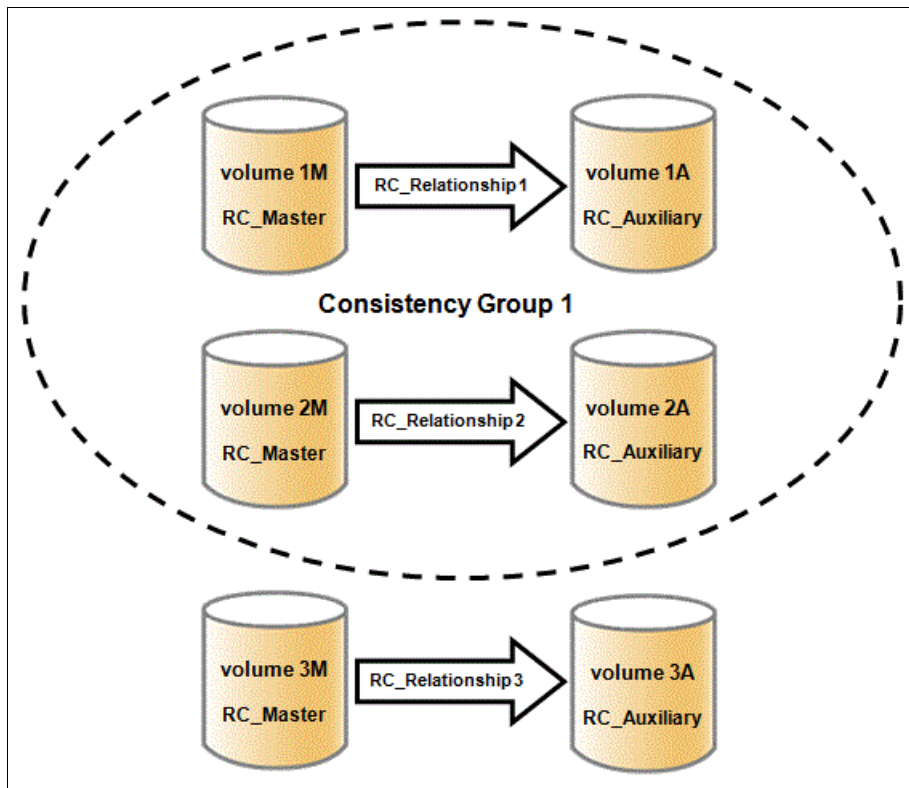


Figure 10-82 Remote Copy consistency group

Remote Copy relationships do not have to belong to a consistency group, but if they do they can only belong to one consistency group. Relationships that are not part of a consistency group are called *stand-alone relationships*.

A consistency group can contain zero or more relationships. All relationships in a consistency group must have matching primary and secondary clusters, which are sometimes referred to as *master clusters* and *auxiliary clusters*. All relationships in a consistency group must also have the same copy direction and state.

Metro Mirror and Global Mirror relationships cannot belong to the same consistency group. A copy type is automatically assigned to a consistency group when the first relationship is added to the consistency group. After the consistency group is assigned a copy type, only relationships of that copy type can be added to the consistency group.

Remote Copy and consistency group states

Stand-alone Remote Copy relationships and consistency groups share a common configuration and state model. All of the relationships in a non-empty consistency group have the same state as the consistency group.

The following states apply to the relationships and the consistency groups, except for the Empty state, which is only for consistency groups:

- ▶ **InconsistentStopped**

The primary volumes are accessible for read and write I/O operations, but the secondary volumes are not accessible for either one. A copy process must be started to make the secondary volumes consistent.
- ▶ **InconsistentCopying**

The primary volumes are accessible for read and write I/O operations, but the secondary volumes are not accessible for either one. This state indicates that a copy process is ongoing from the primary to the secondary volume.
- ▶ **ConsistentStopped**

The secondary volumes contain a consistent image, but it might be outdated compared to the primary volumes. This state can occur when a relationship was in the ConsistentSynchronized state and experienced an error that forced a freeze of the consistency group or the Remote Copy relationship.
- ▶ **ConsistentSynchronized**

The primary volumes are accessible for read and write I/O operations. The secondary volumes are accessible for read-only I/O operations.
- ▶ **Idling**

The primary volumes and the secondary volumes are operating in the primary role. Therefore, the volumes are accessible for write I/O operations.
- ▶ **IdlingDisconnected**

The volumes in this half of the consistency group are all operating in the primary role, and can accept read or write I/O operations.
- ▶ **InconsistentDisconnected**

The volumes in this half of the consistency group are all operating in the secondary role, and cannot accept read or write I/O operations.
- ▶ **ConsistentDisconnected**

The volumes in this half of the consistency group are all operating in the secondary role, and can accept read I/O operations but not write I/O operations.
- ▶ **Empty**

The consistency group does not contain any relationships.

10.2.3 Remote Copy planning

Before you use Remote Copy, you must plan for its usage.

General guidelines for Remote Copy

General guidelines for Remote Copy include the following considerations:

- ▶ Partnerships between up to four IBM Storwize V3700 storage systems, IBM SAN Volume Controller systems, IBM Storwize V7000, or IBM Storwize V5000 is allowed. The partnership must be defined on any partnered IBM Storwize storage systems or IBM SAN Volume Controller systems to make it fully functional.
- ▶ The two volumes in a relationship must be the same size.

- ▶ The Remote Copy relationship can also be established on the volumes within one IBM Storwize V3700 storage system. When the two volumes are in the same cluster, they must be in the same I/O group.
- ▶ You cannot use Remote Copy target volumes as Remote Copy source volumes. However, a FlashCopy target volume can be used as Remote Copy source volume. Other restrictions are outlined in Table 10-5 on page 530.
- ▶ The Metro Mirror function supports copy operations between volumes that are separated by distances up to 300 km.
- ▶ To manage multiple Remote Copy relationships as one entity, relationships can be made part of a Remote Copy consistency group. These relationships ensure data consistency across multiple Remote Copy relationships, and provide ease of management.
- ▶ One Remote Copy relationship can belong only to one consistency group.
- ▶ All relationships in a consistency group must have matching primary and secondary clusters (master clusters and auxiliary clusters). All relationships in a consistency group must also have the same copy direction and state.
- ▶ Metro Mirror and Global Mirror relationships cannot belong to the same consistency group.
- ▶ An IBM Storwize V3700 storage system implements flexible resynchronization support, which enables it to resynchronize volume pairs that experienced write I/Os to both volumes, and to resynchronize only those regions that are known to have changed.
- ▶ Global Mirror with Change Volumes should have Change Volumes defined for the master and auxiliary volumes.

Remote Copy configuration limits

Table 10-4 lists the Remote Copy configuration limits.

Table 10-4 Remote Copy configuration limits

Parameter	Value
Number of Remote Copy consistency groups per cluster	256
Number of Remote Copy relationships per consistency group	No limit beyond the Remote Copy relationships per system.
Number of Remote Copy relationships per I/O Group	2,048 (4,096 per system)
Total Remote Copy volume capacity per I/O Group	1024 terabytes (TB) (This limit is the total capacity for all master and auxiliary volumes in the I/O group.)

SAN planning for Remote Copy

In this section, we describe some guidelines that can be used for SAN planning for Remote Copy.

Zoning recommendation

Node canister ports on each IBM Storwize V3700 must communicate with each other so that the partnership can be created. These ports must be visible to each other on your SAN. Proper switch zoning is critical to facilitating inter-cluster communication.

The following SAN zoning recommendation should be considered:

- ▶ For each node canister, exactly two FC ports should be zoned to exactly two FC ports from each node canister in the partner cluster.
- ▶ If dual-redundant inter-switch links (ISLs) are available, the two ports from each node should be split evenly between the two ISLs. Therefore, exactly one port from each node canister should be zoned across each ISL. For more information about ISLs, see the following website:

<http://www.ibm.com/support/docview.wss?uid=ssg1S1003634&myns=s033&mynp=familyind5329743&mync=E>

- ▶ All local zoning rules should be followed. A properly configured SAN fabric is key to not only local SAN performance, but Remote Copy. For more information about these rules, see the following website:

http://www-01.ibm.com/support/knowledgecenter/STLM5A_7.4.0/com.ibm.storwize.v3700.740.doc/svc_configrulelesssummary_02171530.html?cp=STLM5A%2F0-6-0-0

Important: When a local fabric and a remote fabric are connected for Remote Copy purposes, the ISL “hop” count between a local node and a remote node cannot exceed seven.

Remote Copy link requirements

The following link requirements are valid for Metro Mirror and Global Mirror:

- ▶ Round-trip latency

The total round-trip latency must be less than 80 ms, and less than 40 ms in each direction. Latency simulations should be performed with your applications before any network links are put in place to see whether the applications perform at an acceptable level while they meet the round-trip latency requirement.

- ▶ Bandwidth

The bandwidth must satisfy the following requirements:

- If you are not using Change Volumes, be able to sustain the peak write load for all mirrored volumes and background copy traffic.
- If you are using Change Volumes with Global Mirror, be able to sustain the change rate of Source Change Volumes and background copy traffic.
- Consider other background copy rate (the leading practice is 10% to 20% of maximum peak load) for initial synchronization and resynchronization.
- Remote Copy internal communication at idle with or without Change Volumes is approximately 2.6 megabits per second (Mbps). This amount is the minimum amount.

Redundancy: If the link between two sites is configured with redundancy so that it can tolerate single failures, the link must be sized so that the bandwidth and latency requirements can be met during single failure conditions.

Native Internet Protocol replication

With the advent of version 7.2 code, the IBM Storwize V3700 now supports replication over native Internet Protocol (IP) links using the built-in networking ports of the cluster nodes. This gives greater flexibility in creating Remote Copy links between Storwize and IBM SAN Volume Controller products. FC over IP (FCIP) routers are no longer required. It uses SANslide technology developed by Bridgewater Limited of Christchurch, UK. They specialize in products that can bridge storage protocols and accelerate data transfer over long distances.

Adding this technology at each end of a WAN Transmission Control Protocol/Internet Protocol (TCP/IP) link significantly improves the utilization of the link. It does this by applying patented artificial intelligence (AI) to hide latency normally associated with WANs. Instead, the Storwize technology uses TCP/IP latency to its advantage. In a traditional IP link performance drops off as more data is sent, because each transmission must wait for acknowledgment before the next can be sent.

Rather than wait for the acknowledgment to come back, the Storwize technology sends more sets of packets across other *virtual connections*. The number of virtual connections is controlled by the AI engine. This improves WAN connection use, which results in a data transfer rate approaching full line speed, and is similar to the use of buffer-to-buffer credits in FC.

If packets are lost from any virtual connection, the data is retransmitted, and the remote unit waits for it. Presuming that this is not a frequent problem, overall performance is only marginally affected because of the delay of an extra round trip for the data that is resent. The implementation of this technology can greatly improve the performance of Remote Copy, especially Global Mirror with Change Volumes over long distances.

For more information about Native IP replication and how to use it, see the separate publication:

IBM SAN Volume Controller and Storwize Family Native IP Replication, REDP-5103

Interaction between Remote Copy and FlashCopy

Table 10-5 lists which combinations of FlashCopy and Remote Copy are supported.

Table 10-5 FlashCopy and Remote Copy interaction

Component	Remote Copy primary	Remote Copy secondary
FlashCopy source	Supported	Supported. When the FlashCopy relationship is in the Preparing and Prepared states, the cache at the Remote Copy secondary site operates in write-through mode. This process adds more latency to the already-latent Remote Copy relationship.
FlashCopy target	This combination is supported and has the following restrictions: <ul style="list-style-type: none"> ▶ Running stop -force might cause the Remote Copy relationship to fully resynchronize. ▶ The I/O group must be the same. 	This combination is supported by the restriction that the FlashCopy mapping cannot be copying, stopping, or suspended. Otherwise, the restrictions are the same as at the Remote Copy primary site.

If you are not using Global Mirror with Change Volumes, for disaster recovery purposes, you can use the FlashCopy feature to create a consistent copy of an image before you restart a Global Mirror relationship.

When a consistent relationship is stopped, the relationship enters the `consistent_stopped` state. While in this state, I/O operations at the primary site continue to run. However, updates are not copied to the secondary site. When the relationship is restarted, the synchronization process for new data is started. During this process, the relationship is in the `inconsistent_copying` state.

The secondary volume for the relationship cannot be used until the copy process completes and the relationship returns to the consistent state. When this situation occurs, start a FlashCopy operation for the secondary volume before you restart the relationship. While the relationship is in the Copying state, the FlashCopy feature can provide a consistent copy of the data. If the relationship does not reach the synchronized state, you can use the FlashCopy target volume at the secondary site.

10.3 Troubleshooting Remote Copy

Remote Copy (Global Mirror and Metro Mirror) has the following primary error codes:

- ▶ 1920. This error is a congestion error that means that the source, the link between source and target, or the target cannot keep up with the rate of demand.
- ▶ 1720. This error is a heartbeat or cluster partnership communication error. This error tends to be more serious, because failing communication between your cluster partners involves some extended diagnostic time.

10.3.1 1920 error

A 1920 error (event ID 050010) can have several triggers. The following official probable cause projections are available:

- ▶ Primary cluster or SAN fabric problem (10%)
- ▶ Primary cluster or SAN fabric configuration (10%)
- ▶ Secondary cluster or SAN fabric problem (15%)
- ▶ Secondary cluster or SAN fabric configuration (25%)
- ▶ Inter-cluster link problem (15%)
- ▶ Inter-cluster link configuration (25%)

In practice, the error that is most often overlooked is latency. Global Mirror has a round-trip-time tolerance limit of 80 ms. A message that is sent from your source cluster to your target cluster and the accompanying acknowledgment must have a total time of less than 80 ms (or 40 ms each way).

The primary component of your round-trip time is the physical distance between sites. For every 1,000 km (621.36 miles), there is a 5 ms delay. This delay does not include the time that is added by equipment in the path. Every device adds a varying amount of time depending on the device, but you can expect about 25 μ s for pure hardware devices. For software-based functions (such as compression that is implemented in software), the added delay tends to be much higher (usually in the millisecond-plus range).

Consider this example. Company A has a production site that is 1,900 km from their recovery site. Their network service provider uses five devices to connect the two sites. In addition to those devices, Company A uses a SAN Fibre Channel Router at each site to provide FCIP to encapsulate the Fibre Channel traffic between sites. There are now seven devices, and 1,900 km of distance delay. All of the devices add 200 μ s of delay each way. The distance adds 9.5 ms each way, for a total of 19 ms.

Combined with the device latency, that is 19.4 ms of physical latency at a minimum. This latency is under the 80 ms limit of Global Mirror, but this number is the best-case number. Link quality and bandwidth play a significant role here. Your network provider likely ensures a latency maximum on your network link; be sure to stay below the Global Mirror round-trip time (RTT) limit.

You can easily double or triple the expected physical latency with a lower-quality or lower-bandwidth network link. As a result, you are suddenly within range of exceeding the limit the moment a large flood of I/O happens that exceeds the bandwidth capacity you have in place.

When you get a 1920 error, always check the latency first. The FCIP routing layer can introduce latency if it is not properly configured. If your network provider reports a much lower latency, this report can be an indication of a problem at your FCIP routing layer. Most FCIP routing devices have built-in tools that you can use to check the RTT delay. When you are checking latency, remember that TCP/IP routing devices (including FCIP routers) report RTT by using standard 64-byte ping packets.

Figure 10-83 shows why the effective transit time should be measured only using packets large enough to hold a FC frame. This packet size is 2148 bytes (2112 bytes of payload and 36 bytes of header), and you should allow more capacity to be safe, because different switching vendors have optional features that might increase this size.

Packet Size	Link Size	Serialization Delay (Time Required to Send Data)	Unit
64	256 Kbps	2.0E+03	microseconds
64	1.5 Mbps	3.4E+02	microseconds
64	100 Mbps	5.1E+00	microseconds
64	155 Mbps	3.3E+00	microseconds
64	622 Mbps	8.2E-01	microseconds
64	1 Gbps	5.1E-04	microseconds
64	10 Gbps	5.1E-05	microseconds
1500	256 Kbps	4.7E+04	microseconds
1500	1.5 Mbps	8.0E+03	microseconds
1500	100 Mbps	1.2E+02	microseconds
1500	155 Mbps	7.7E+01	microseconds
1500	622 Mbps	1.9E+01	microseconds
1500	1 Gbps	1.2E+01	microseconds
1500	10 Gbps	1.2E+00	microseconds
2148	256 Kbps	6.7E+04	microseconds
2148	1.5 Mbps	1.1E+04	microseconds
2148	100 Mbps	1.7E+02	microseconds
2148	155 Mbps	1.1E+02	microseconds
2148	622 Mbps	2.8E+01	microseconds
2148	1 Gbps	1.7E+01	microseconds
2148	10 Gbps	1.7E-03	microseconds

Figure 10-83 Effect of packet size (in bytes) versus link size

Before you proceed, take a quick look at the second-largest component of your RTT: *serialization delay*. Serialization delay is the amount of time that is required to move a packet of data of a specific size across a network link of a specific bandwidth. This delay is based on a simple concept that the time that is required to move a specific amount of data decreases as the data transmission rate increases.

In Figure 10-83, there are orders of magnitude of difference between the different link bandwidths. It is easy to see how 1920 errors can arise when your bandwidth is insufficient and why you should never use a TCP/IP ping to measure RTT for FCIP traffic.

Figure 10-83 on page 532 compares the amount of time in microseconds that is required to transmit a packet across network links of varying bandwidth capacity. The following packet sizes are used:

- ▶ 64 bytes: The size of the common ping packet
- ▶ 1500 bytes: The size of the standard TCP/IP packet
- ▶ 2148 bytes: The size of a Fibre Channel frame

Your path maximum transmission unit (MTU) affects the delay that is incurred in getting a packet from one location to another in the following circumstances:

- ▶ When it causes fragmentation
- ▶ When it is too large and causes too many retransmits when a packet is lost

After you verify your latency by using the correct packet size, proceed with normal hardware troubleshooting.

10.3.2 1720 error

The 1720 error (event ID 050020) is the other primary error code of Remote Copy. Because the term *system partnership* implies that all involved virtualization systems are partners, they must communicate with each other. When a partner on either side stops communicating, you see a 1720 error display in your error log. According to official documentation, there are no likely field-replaceable unit breakages or other causes.

In practice, the source of this error is most often a fabric problem, or a problem in the network path between your partners. When you receive this error, if your fabric has more than 64 host bus adapter (HBA) ports zoned, you should check your fabric configuration to see if more than one HBA port for each node per I/O group are zoned together.

The suggested zoning configuration for fabrics is one port of each node per I/O group associated with each host HBA. For those fabrics with 64 or more host ports, this recommendation becomes a rule. You must follow this zoning rule or the configuration is technically unsupported.

Improper zoning leads to SAN congestion, which can inhibit remote link communication intermittently. Checking the zero buffer credit timer through IBM Tivoli Storage Productivity Center and comparing its value against your sample interval might reveal potential SAN congestion. When a zero buffer credit timer is above 2% of the total time of the sample interval, it is likely to cause problems.

Next, always ask your network provider to check the status of the link. If the link is okay, watch for repetition of this error. It is possible in a normal and functional network setup to have occasional 1720 errors, but multiple occurrences indicate a larger problem.

If you receive multiple 1720 errors, recheck your network connection, and then check the IBM Storwize V3700 partnership information to verify the status and settings. Perform diagnostic tests for every piece of equipment in the path between the two systems. It often helps to have a diagram that shows the path of your replication from logical and physical configuration viewpoints.

If your investigation fails to resolve your Remote Copy problems, you should contact your IBM support representative for a complete analysis.

10.4 Managing Remote Copy using the GUI

The IBM Storwize V3700 storage system provides a separate function icon for copy service management. The following windows are available for managing Remote Copy, which are accessed through the Copy Services function icon:

- ▶ Remote Copy
- ▶ Partnerships

As the names imply, these windows are used to manage Remote Copy and the partnership.

10.4.1 Managing cluster partnerships

The Partnership window is used to manage a partnership between clusters. To access the Partnership window, complete the following steps:

1. Click the **Copy Services** function icon.
2. Click **Partnerships**, as shown in Figure 10-84.




Figure 10-84 Partnerships window

Creating a partnership

To create a partnership, complete the following steps:

1. Click **Create Partnership** in the Partnership window, as shown in Figure 10-85.




Figure 10-85 Create a cluster partnership

2. You are prompted to choose between an FC or IP partnership, as shown in Figure 10-86.




Figure 10-86 Creating an FC or IP partnership

3. If there is no partnership candidate, an error window opens, as shown in Figure 10-87.

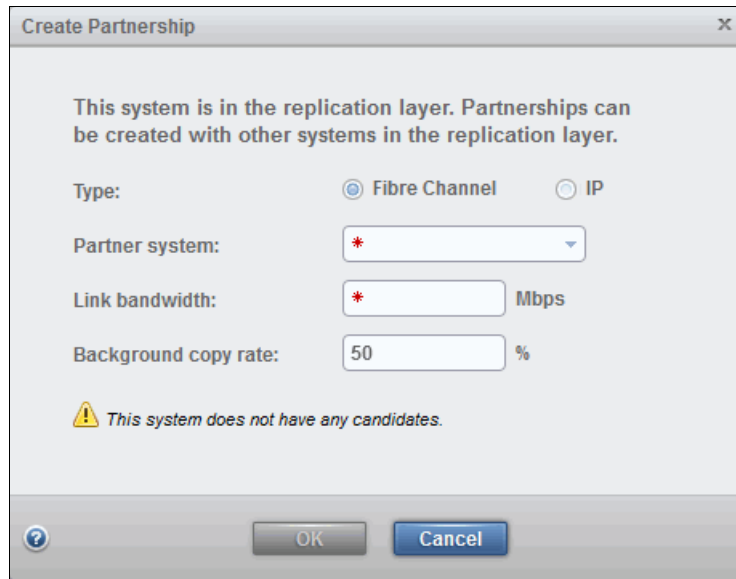


Figure 10-87 No candidates are available to create a partnership

4. Check the zoning and the system status and make sure that the clusters can see each other.

The two systems should not be visible to each other using both FC *and* IP. If you are going to use IP replication, ensure that the two systems are not zoned in any way at the SAN level. A system can have simultaneous partnerships over IP and FC with different systems, but never the same one. For more information, see *IBM SAN Volume Controller and Storwize Family Native IP Replication*, REDP-5103. In this example, we choose to create an FC partnership with an IBM Storwize V5000 system.

5. If the Storwize V3700 can see a candidate system, create your partnership by selecting the appropriate remote storage system from the **Partner system** drop-down, as shown in Figure 10-88, and defining the available bandwidth and the background copy percentage between both systems.

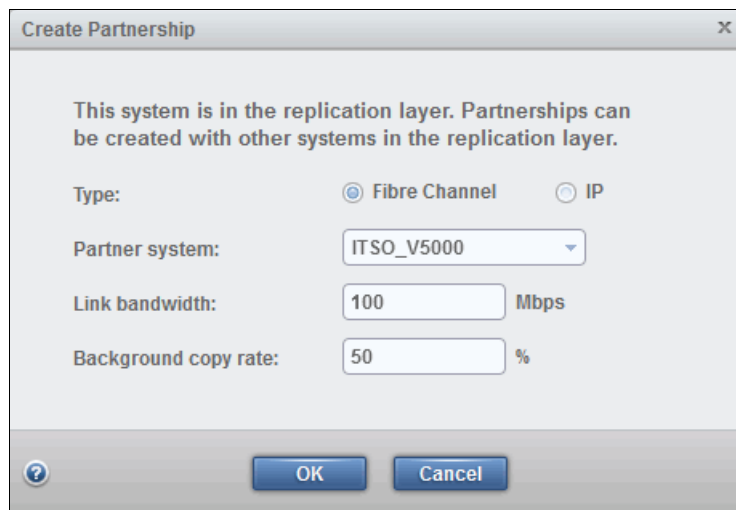


Figure 10-88 Select the remote storage system for the new partnership, and configure partnership properties

6. The bandwidth that you must enter here is used by the background copy process between the clusters in the partnership. To set the background copy bandwidth optimally, make sure that you consider all three resources (primary storage, inter-cluster link bandwidth, and auxiliary storage) to avoid overloading them, which affects the I/O latency.
7. Click **OK** and the partnership task starts.
8. When it finishes, click **Close** and the partnership is complete on this system. It is listed in the Partnership window, but is displayed with an alert and the system health status turns yellow, as shown in Figure 10-89. This is normal, and happens because only one half of the partnership currently exists.

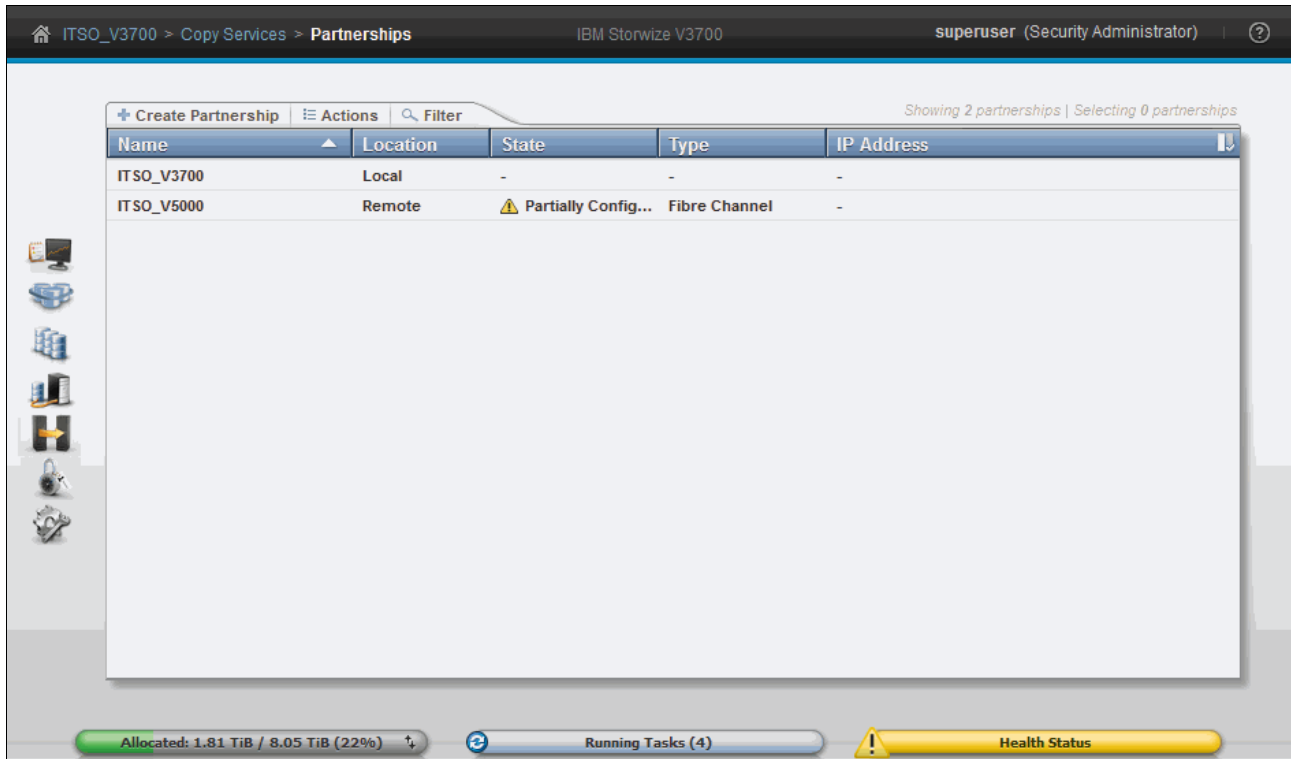


Figure 10-89 Partially configured partnership

9. If you select the partnership, more information is available through the Properties option on the **Actions** menu or by right-clicking, as shown in Figure 10-90.




Figure 10-90 Partnership properties

Important: The partnership is in the Partially Configured: Local state because we have not yet defined it on the other IBM Storwize system. For more information about partnership states, see “Remote Copy and consistency group states” on page 526.

10. From the Properties menu option, the partnership bandwidth and background copy rate can be altered by clicking **Edit**, as shown in Figure 10-91.

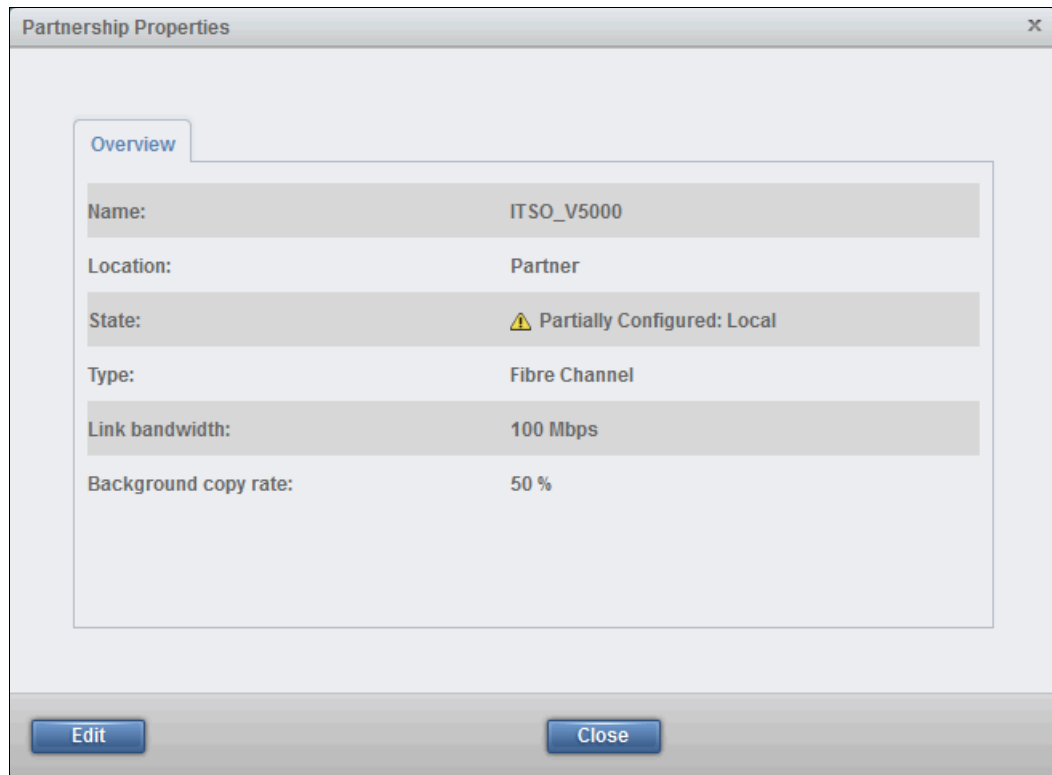


Figure 10-91 Editing partnership bandwidth and background copy rate

11. When you have edited the parameters, click **Save** to save changes or **Cancel** to quit without making changes.

12. Complete the same steps on the second storage system for the partnership to become fully configured. The Remote Copy partnership is now complete between the two IBM Storwize systems, as shown in Figure 10-92, and both systems are ready for Remote Copy relationships to be configured.




Figure 10-92 Fully configured partnership

Stopping and starting a partnership

You can stop the partnership by completing the following steps:

1. Click **Stop** from the **Actions** menu, as shown in Figure 10-93 on page 541. Stopping the partnership disconnects the relationship that uses it.

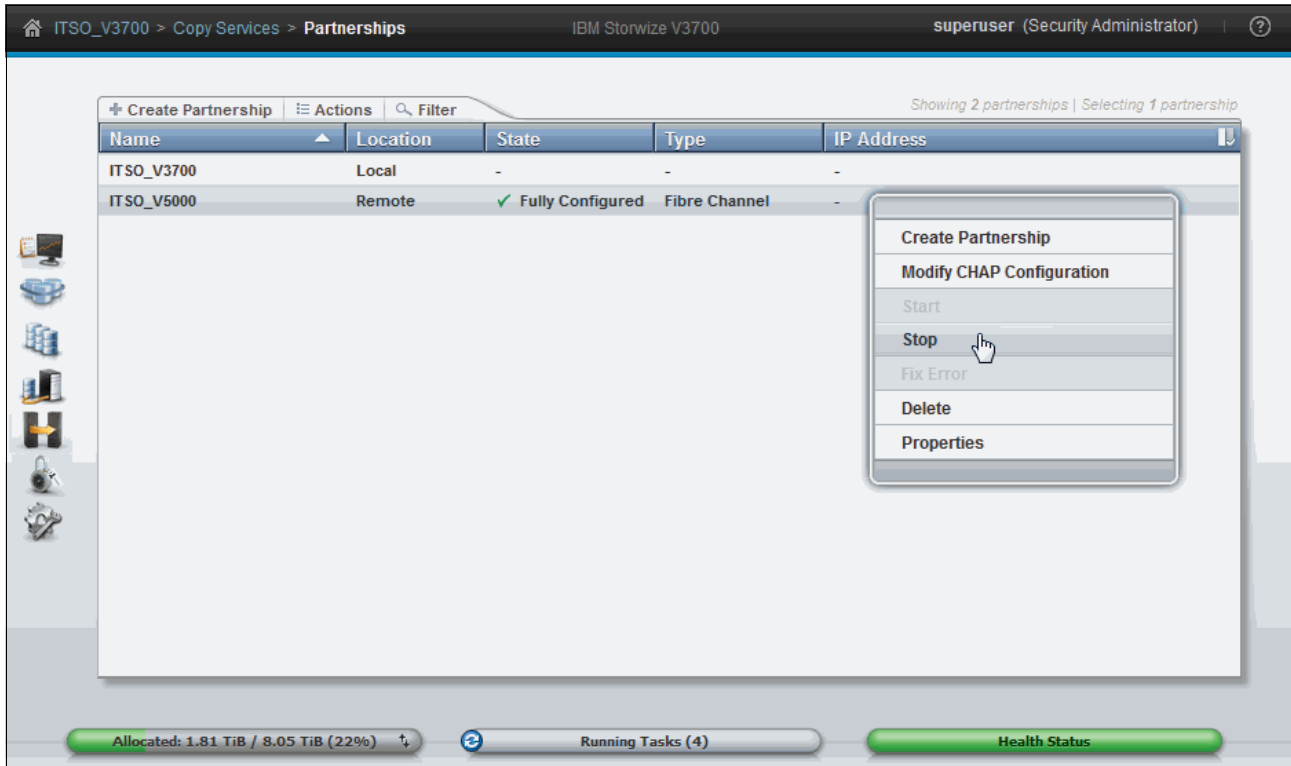


Figure 10-93 Stopping a partnership

2. After you stop the partnership, your partnership is listed as Fully Configured: Stopped, as shown in Figure 10-94.

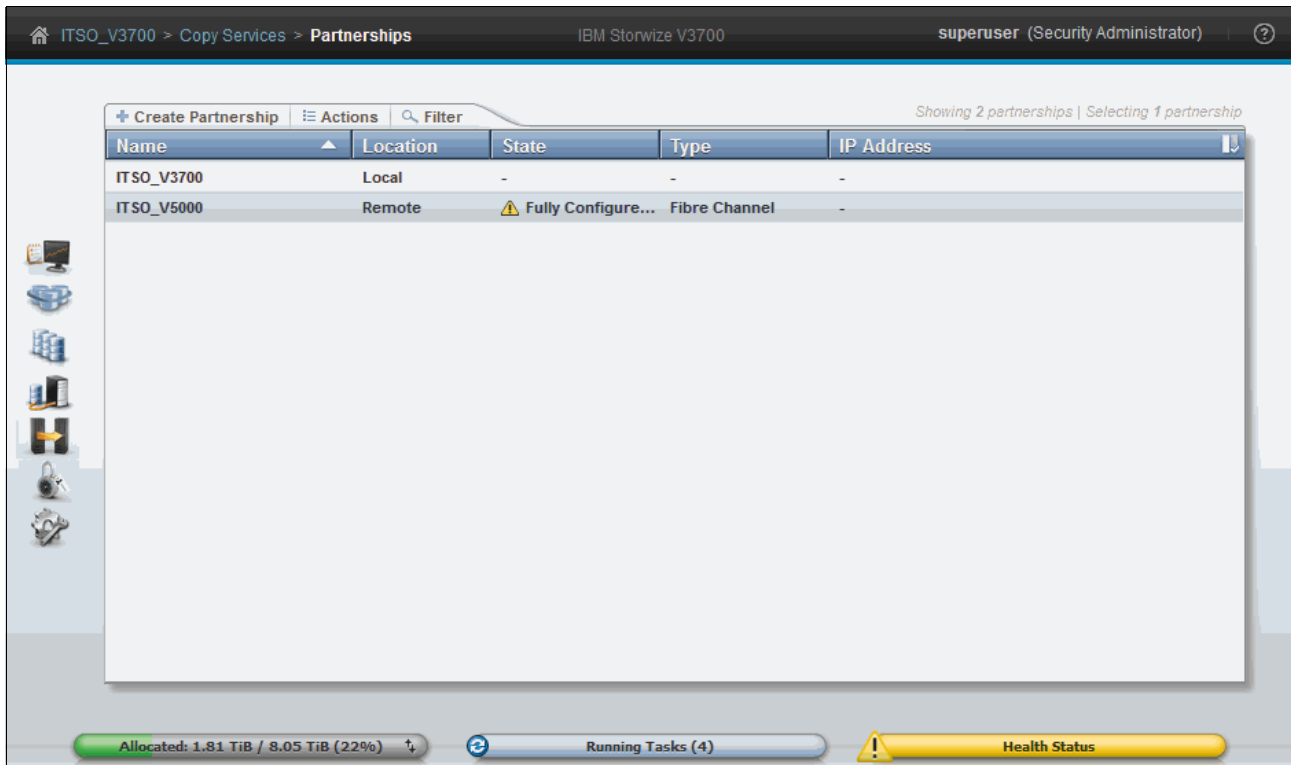


Figure 10-94 Fully configured partnership in Stopped state

- The Health Status turns yellow to indicate that the partnership is down. This is normal. You can restart a stopped partnership by clicking **Start** from the **Actions** menu.

The partnership returns to the fully configured status when restarted, and the health bar returns to green. Initially, when restarting a partnership, the status might change to a Not Present state and the health bar turns red, as shown in Figure 10-95. This is normal, and when the GUI refreshes it will correct itself and go to a Fully Configured state, and the health bar will turn green.

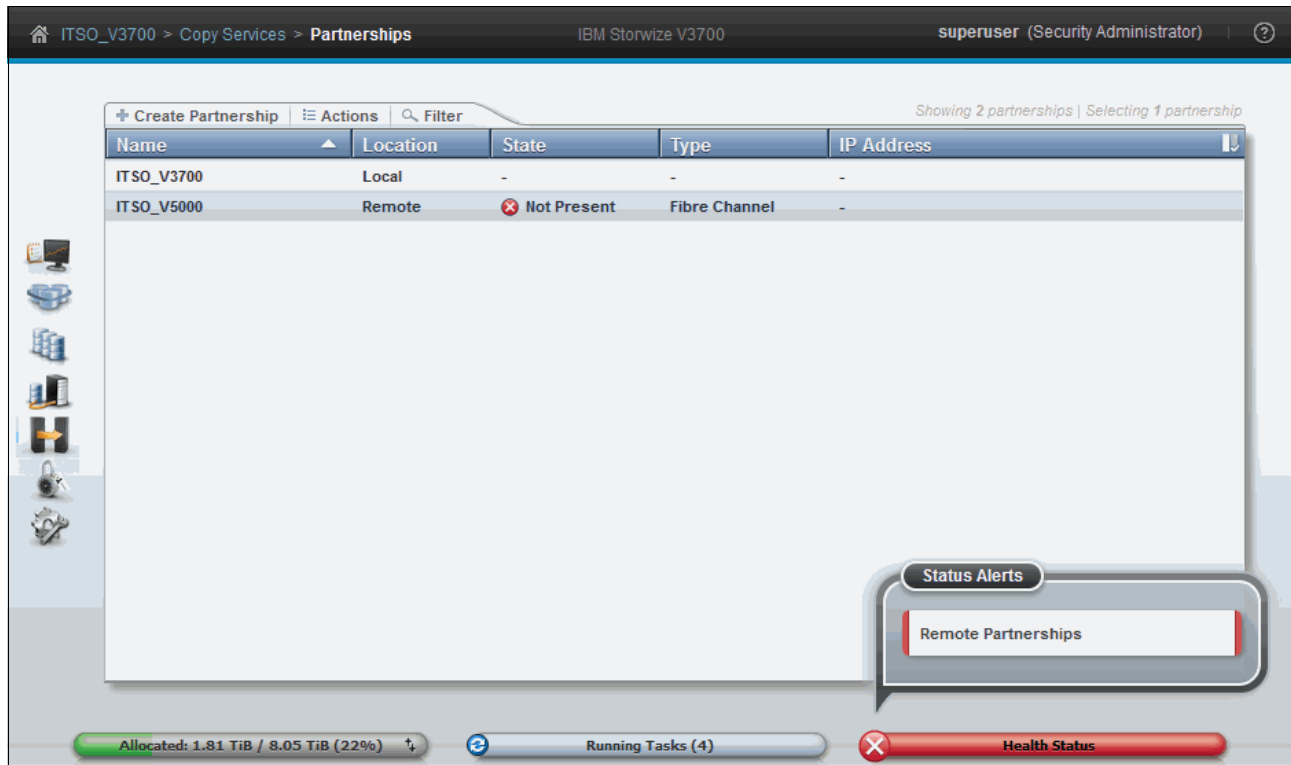


Figure 10-95 Restarting a partnership

Deleting a partnership

You can delete a partnership by selecting **Delete** from the **Actions** menu, as shown in Figure 10-93 on page 541.

10.4.2 Managing stand-alone Remote Copy relationships

A Remote Copy relationship can be defined between two volumes where one is the master (source) and the other one is the auxiliary (target) volume. Use of Remote Copy auxiliary volumes as Remote Copy master volumes is not allowed.

To manage stand-alone Remote Copy relationships, complete the following steps:

- Open the Remote Copy window to manage Remote Copy by clicking the **Copy Services** icon and then clicking **Remote Copy**, as shown in Figure 10-96 on page 543.




Figure 10-96 Open Remote Copy window

The Remote Copy window is where management of Remote Copy relationships and Remote Copy consistency groups is carried out, as shown in Figure 10-97.




Figure 10-97 Remote Copy window

2. The Remote Copy window displays a list of Remote Copy consistency groups. You can also take actions on the Remote Copy relationships and Remote Copy consistency groups. Click **Not in a Group** and all the Remote Copy relationships that are not in any Remote Copy consistency groups are displayed.
3. To customize the blue column heading bar and select different attributes of Remote copy relationships, right-click anywhere in the blue bar.

Creating stand-alone Remote Copy relationships

Important: Before a remote copy relationship can be created, target volumes that are the same size as the source volumes that you want to mirror must be created. For more information about creating volumes, see Chapter 5, “Volume configuration” on page 199.

To create a Remote Copy relationship, complete the following steps:

1. Click **Create Relationship** from the **Actions** menu, as shown in Figure 10-98. A wizard opens and guides you through the Remote Copy relationship creation process.

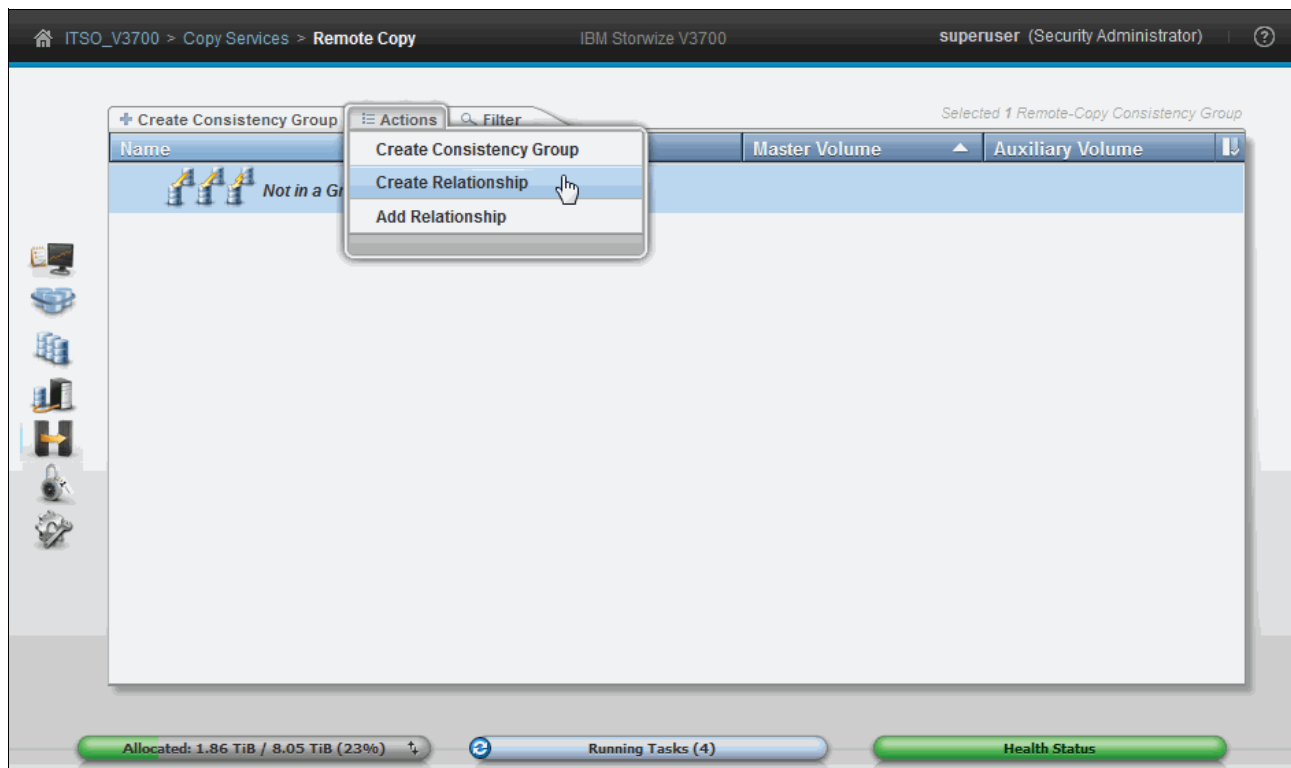


Figure 10-98 Creating a remote copy relationship

- Next, as shown in Figure 10-99, you must set the Remote Copy relationship type. You can select Metro Mirror (synchronous replication) or one of the two forms of Global Mirror (asynchronous replication). Select the appropriate replication type based on your requirements and click **Next**.

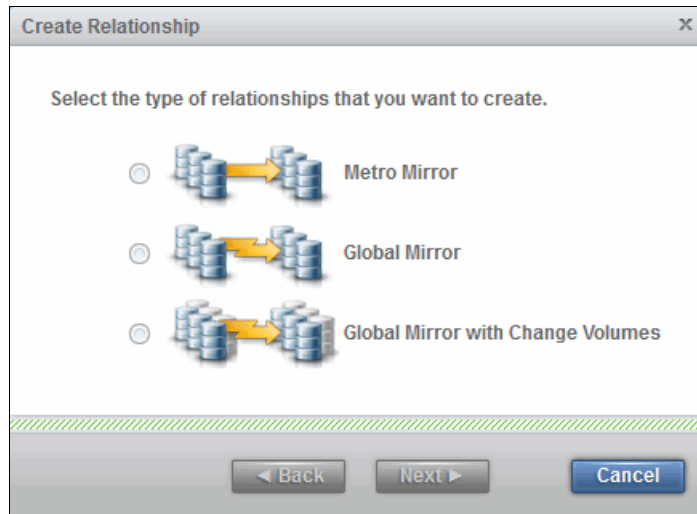


Figure 10-99 Select the appropriate Remote Copy type

- You must select where your auxiliary (target) volumes are:
 - The local system
 - The already-defined second storage system

In our example (as shown in Figure 10-100), we choose another system to build an inter-cluster relationship. Click **Next** to continue.

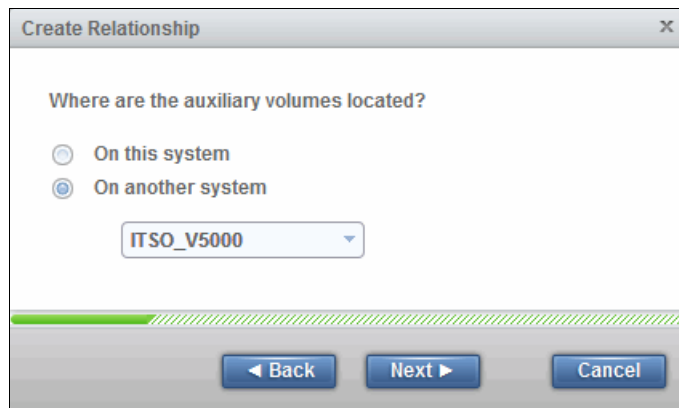


Figure 10-100 Select Remote Copy partner

- Now the Remote Copy master and the auxiliary volume must be specified. Both volumes must have the same size. As shown in Figure 10-101, use the drop-down boxes to select the Master and Auxiliary volumes. The system offers only appropriate auxiliary candidates with the same volume size as the selected master volume. After selecting the volumes, click **Add**.

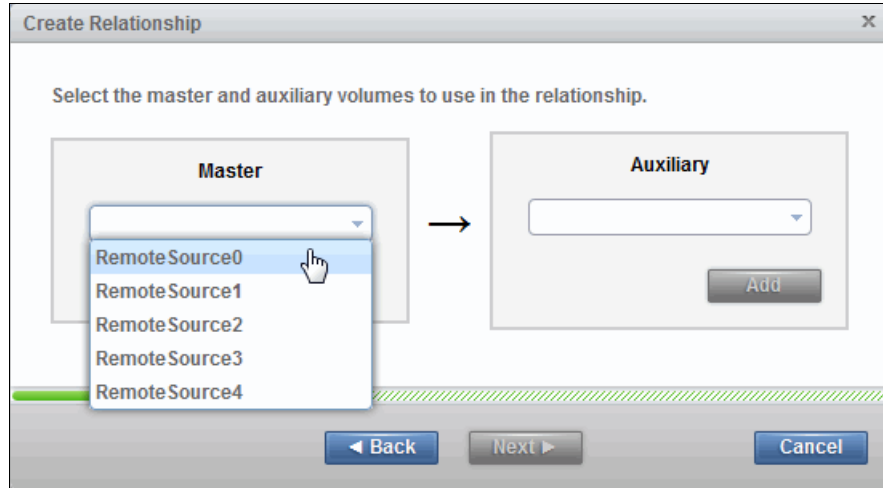


Figure 10-101 Select the master and auxiliary volumes

- You can define multiple and independent relationships by choosing another set of volumes and clicking **Add** again. Incorrectly defined relationships can be removed by clicking the red "X". In our example, we create five independent Remote Copy relationships, as shown in Figure 10-102.

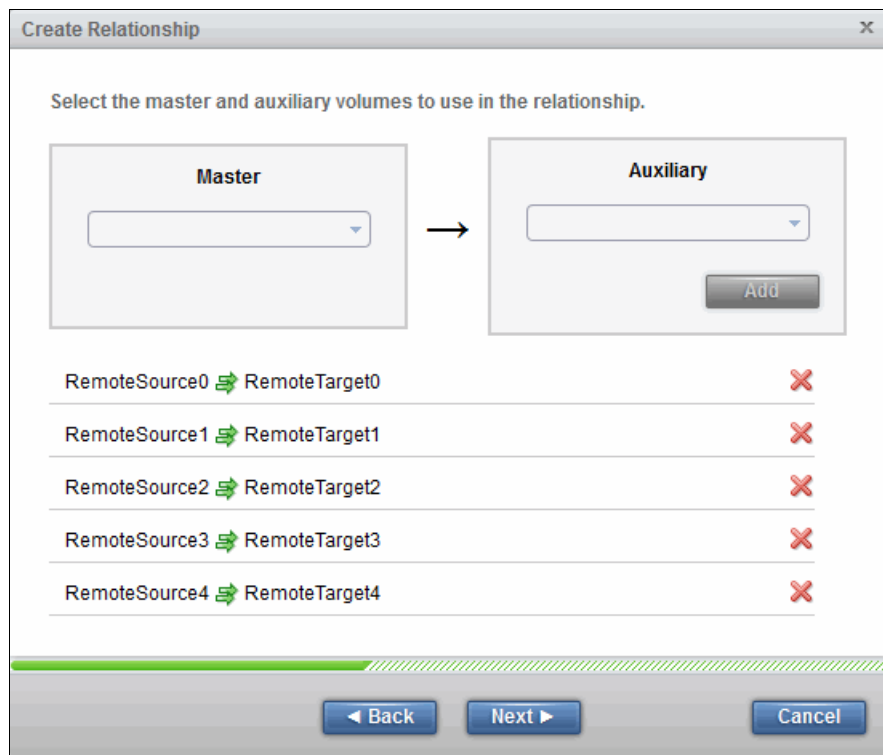


Figure 10-102 Define multiple independent relationships

- The next prompt asks you to select if the volumes in the relationship are already synchronized. In most situations, the data on the master volume and on the auxiliary volume are not identical, so click **No** and then click **Next** to enable an initial copy, as shown in Figure 10-103.




Figure 10-103 Activate initial data copy

- If you select **Yes, the volumes are already synchronized**, a warning message opens, as shown in Figure 10-104. Confirm that the volumes are truly identical, and then click **OK** to continue.

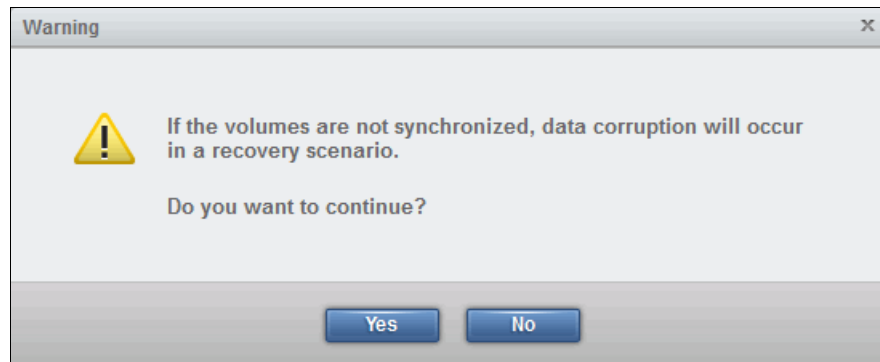


Figure 10-104 Warning message to make sure that the volumes are synchronized

- You can choose to start the initial copying progress now or wait to start it later. In this example, we select **Yes, start copying now** and click **Finish**, as shown in Figure 10-105.

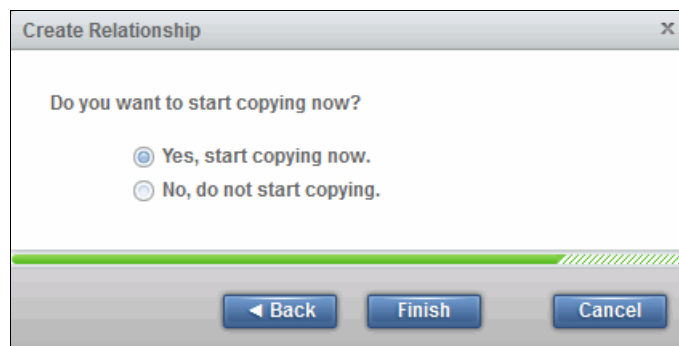


Figure 10-105 Choose if you want to start copying now or later

9. After the Remote Copy relationships creation completes, two independent Remote Copy relationships are defined and displayed in the Not in a Group list, as shown in Figure 10-106.



Figure 10-106 Creating a Remote Copy relationship process completes

10. Optionally, you can monitor the ongoing initial synchronization in the Running Tasks status indicator, as shown in Figure 10-107.

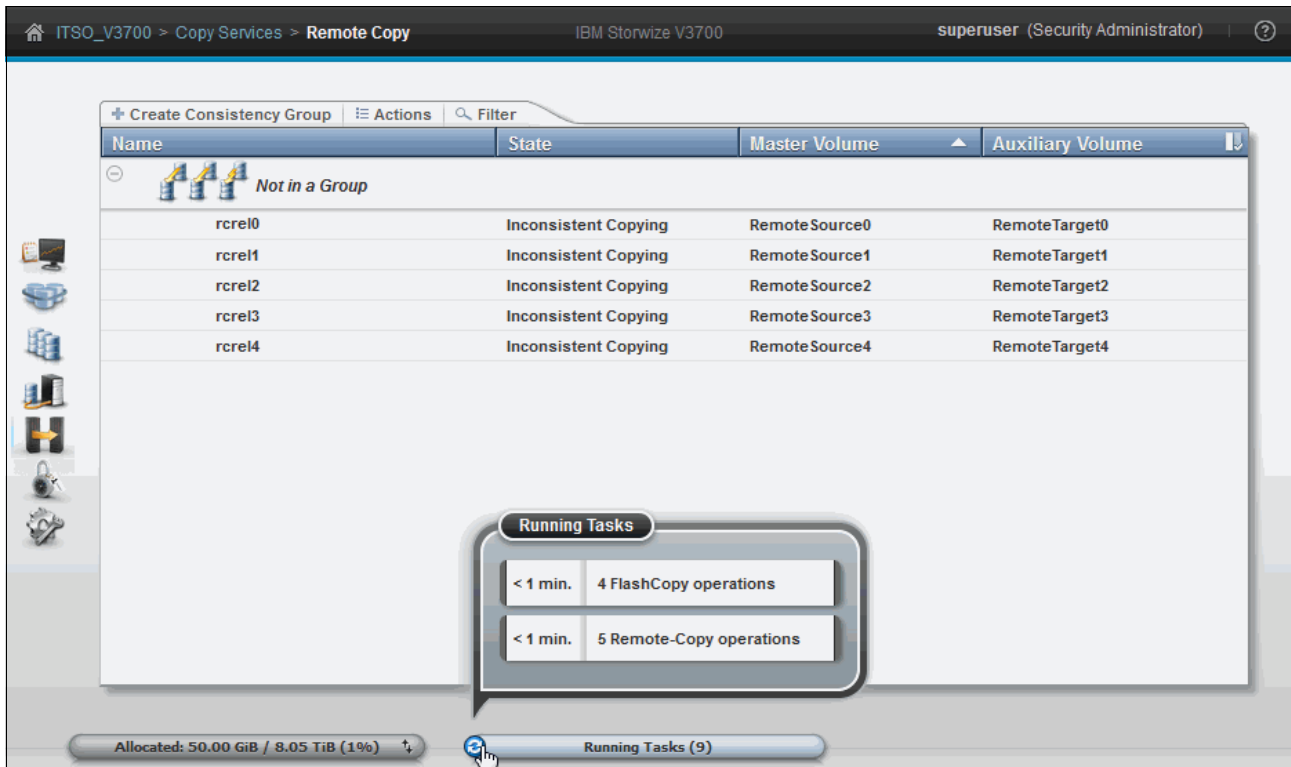


Figure 10-107 Remote copy initialization progress through Running Tasks

11. Highlight one of the operations and click to see the progress, as shown in Figure 10-108.

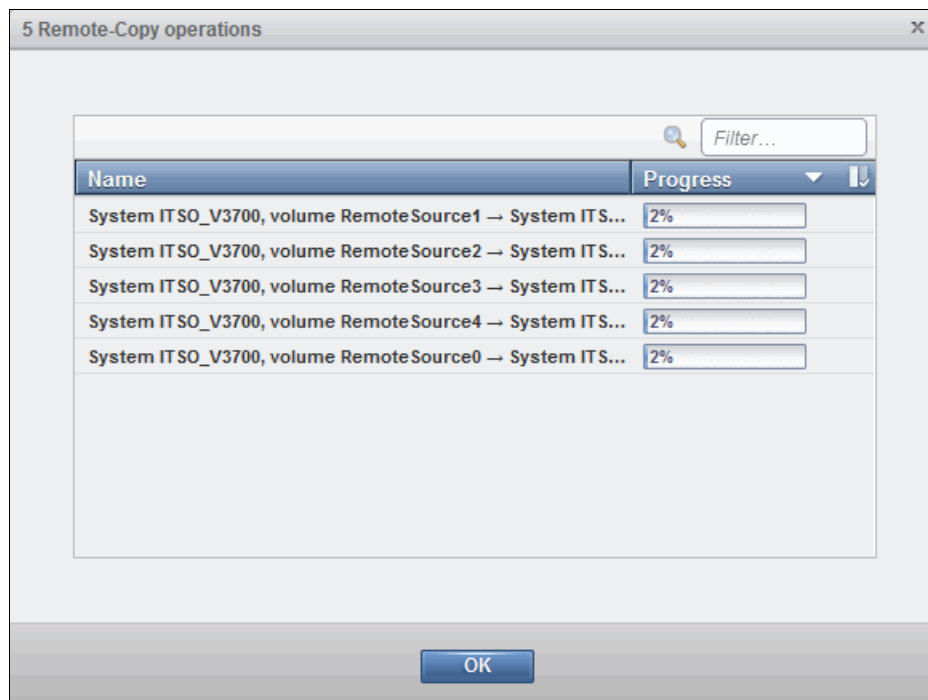


Figure 10-108 Running task expanded to show the percentage complete of each copy operation

Stopping a stand-alone Remote Copy relationship

The Remote Copy relationship can be stopped by completing the following steps:

1. Select the relationship and click **Stop** from the **Actions** menu, as shown in Figure 10-109.

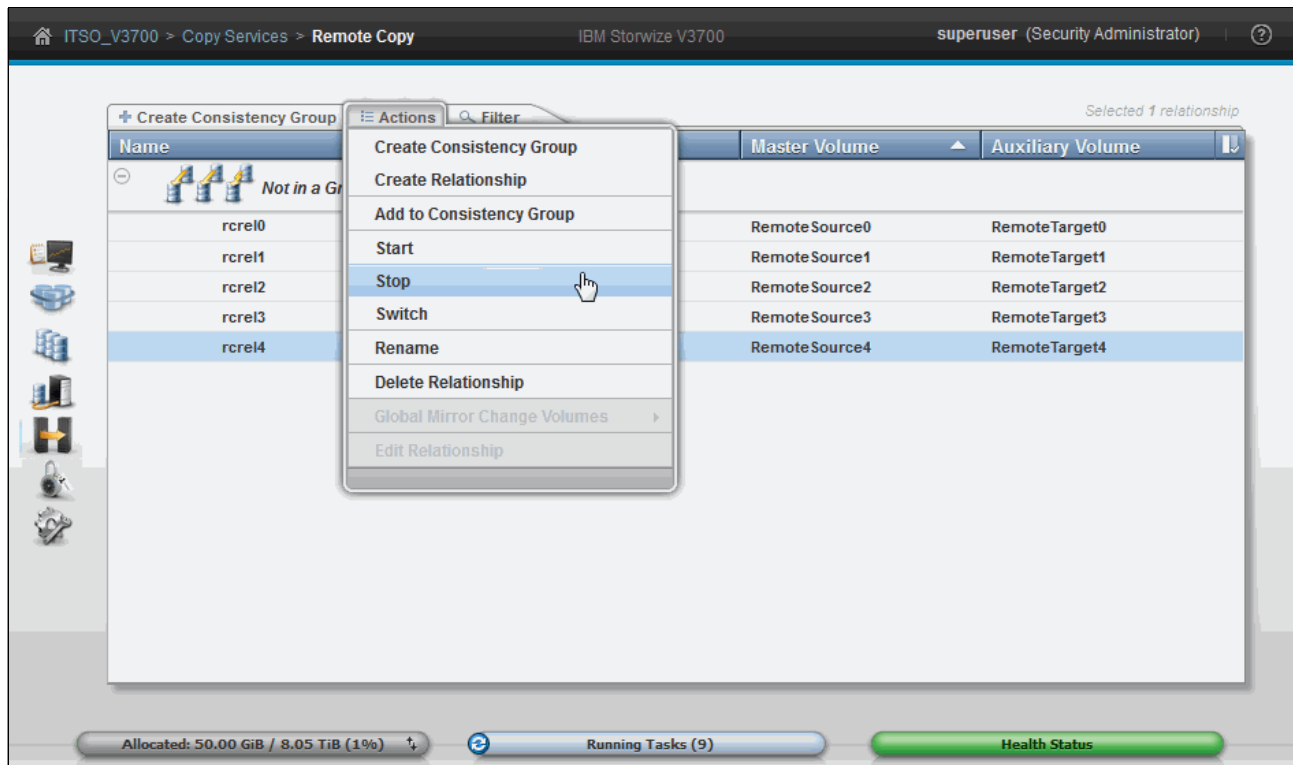


Figure 10-109 Stop Remote Copy relationship

2. A prompt displays. Click to **Allow secondary read/write access**, if required, and then click **Stop Relationship**, as shown in Figure 10-110.

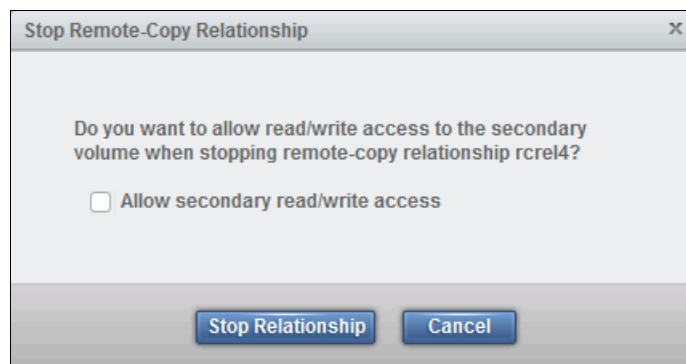


Figure 10-110 Option to allow secondary read/write access

- After the stop completes, the state of the Remote Copy relationship is changed to **Inconsistent Stopped** if it is still in the process of copying, as shown in Figure 10-111. If the relationship was **Consistent Synchronized**, it changes to **Consistent Stopped** if secondary read/write access was not allowed. Otherwise, the state is **Idling**.




Figure 10-111 Remote Copy relationship stop completes

Starting a stand-alone Remote Copy relationship

The Remote Copy relationship can be started by completing the following steps:

1. Select the relationship and click **Start** from the **Actions** menu, or right-click, as shown in Figure 10-112.

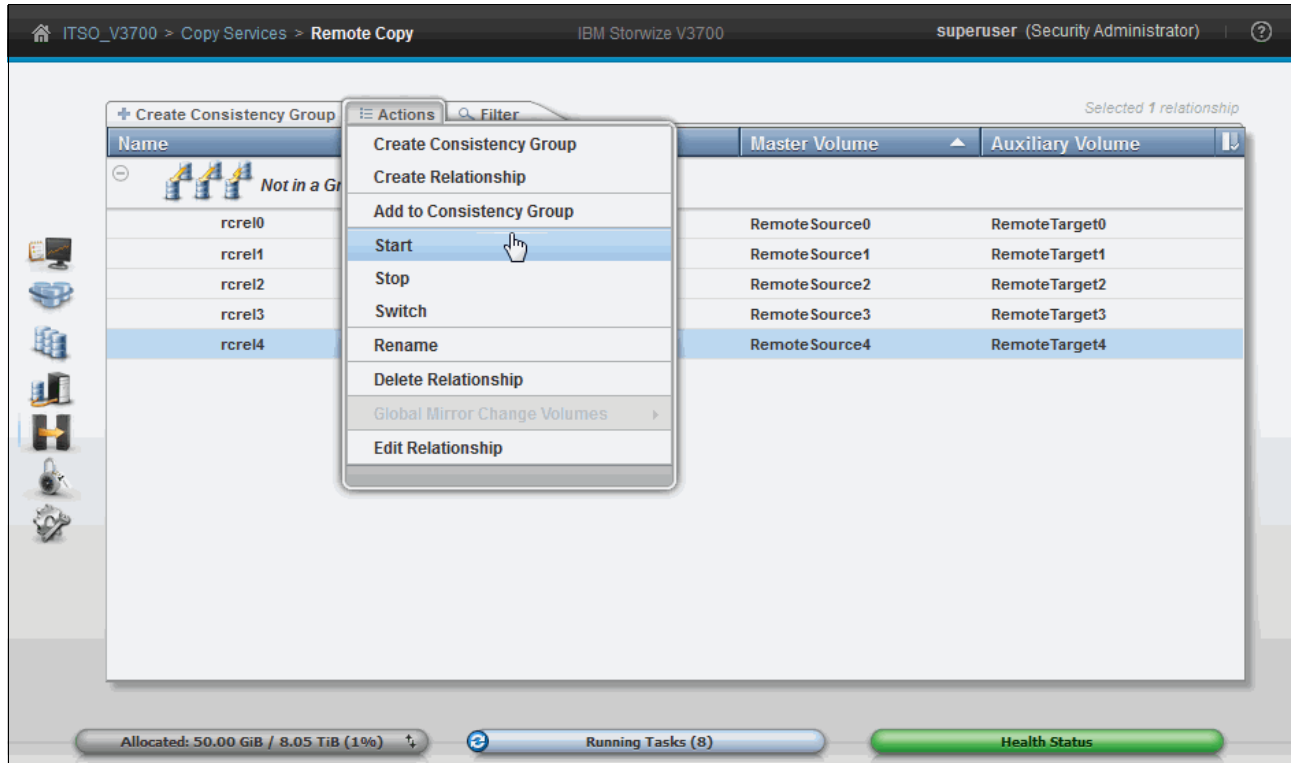


Figure 10-112 Re-start a Remote Copy relationship

- When a Remote Copy relationship is started, the most important item is selecting the copy direction. Both master and auxiliary volumes can be the primary. Make your decision based on your requirements, and click **Start Relationship**. In our example, choose the master volume to be the primary, as shown in Figure 10-113.

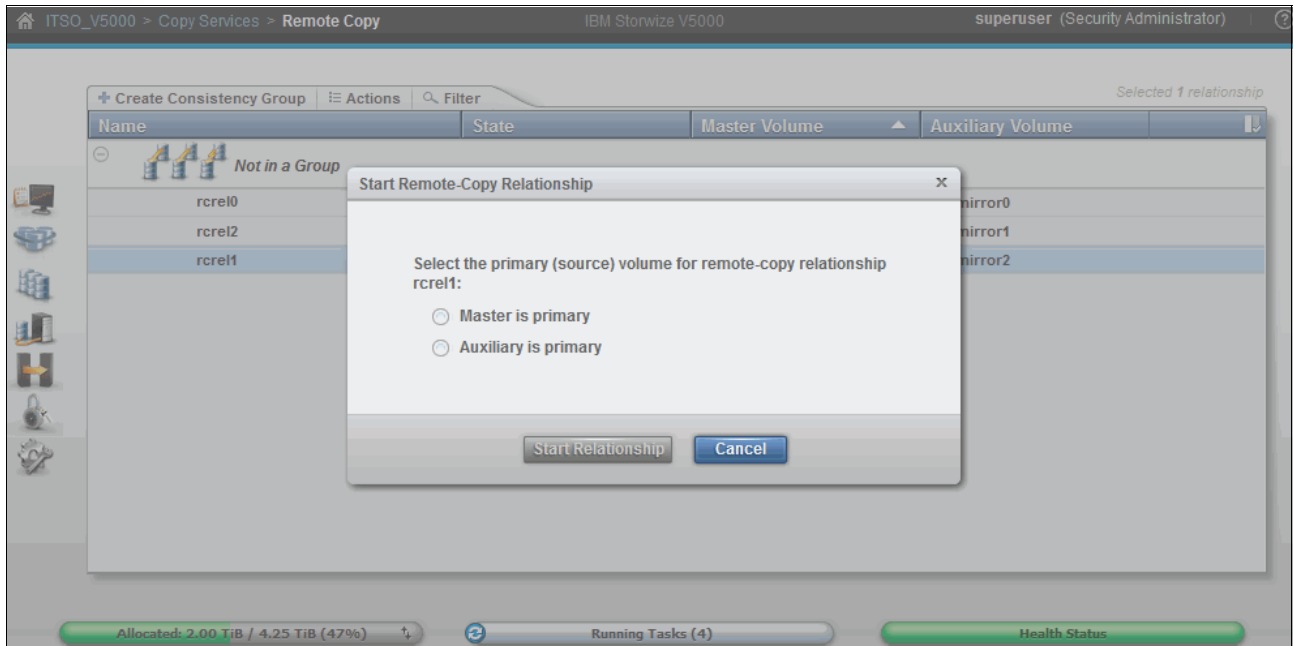


Figure 10-113 Choose the copy direction

Switching the direction of a stand-alone Remote Copy relationship

To switch the copy direction of the Remote Copy relationship, complete the following steps:

1. Select the relationship and click **Switch** from the **Actions** menu, as shown in Figure 10-114.

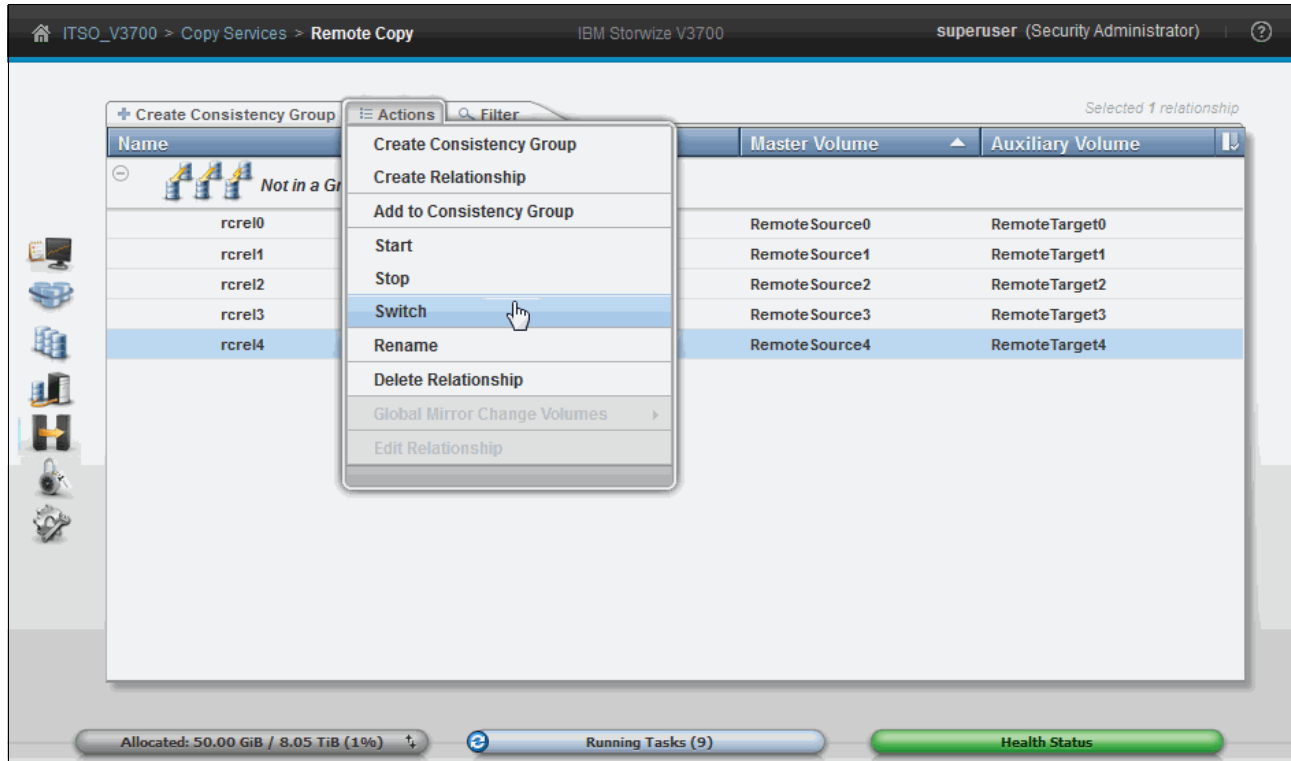


Figure 10-114 Switch Remote Copy relationship

A warning message opens and shows you the consequences of this action, as shown in Figure 10-115 on page 555. If you switch the Remote Copy relationship, the copy direction of the relationship becomes the opposite; that is, the current primary volume becomes the secondary while the current secondary volume becomes the primary.

Write access to the current primary volume is lost, and write access to the current secondary volume is enabled. If it is not a disaster recovery situation, you must stop your host I/O to the current primary volume in advance.

2. Make sure that you are prepared for the consequences. If so, click **Yes** to continue.

Figure 10-115 shows the warning message explaining what will happen.

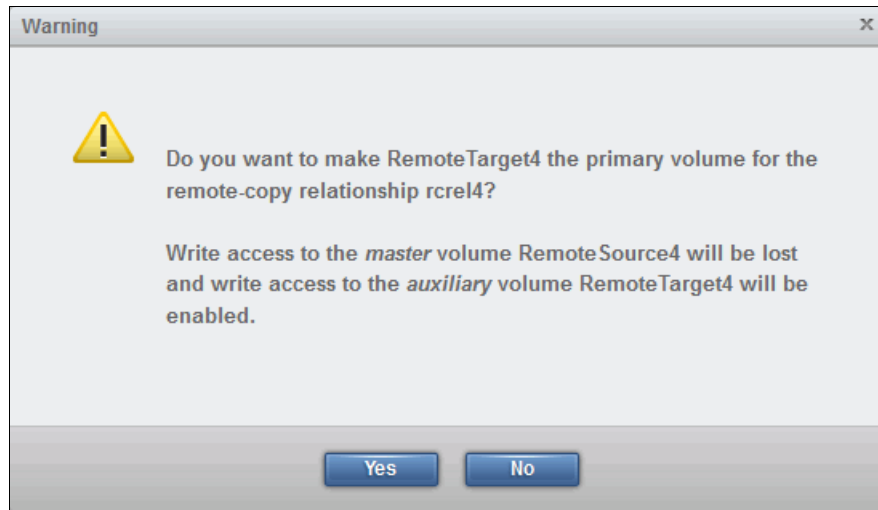


Figure 10-115 Warning message for switching direction of a Remote Copy relationship

3. After the switch completes, your Remote Copy relationship is tagged (as shown in Figure 10-116), and shows you that the primary volume in this relationship was changed to the auxiliary.




Figure 10-116 Switch icon on the state of the relationship

Renaming a stand-alone Remote Copy relationship

The Remote Copy relationship can be renamed by completing the following steps:

1. Select the relationship and then click **Rename** from the **Actions** menu, as shown in Figure 10-117, or highlight the relationships and right-click.

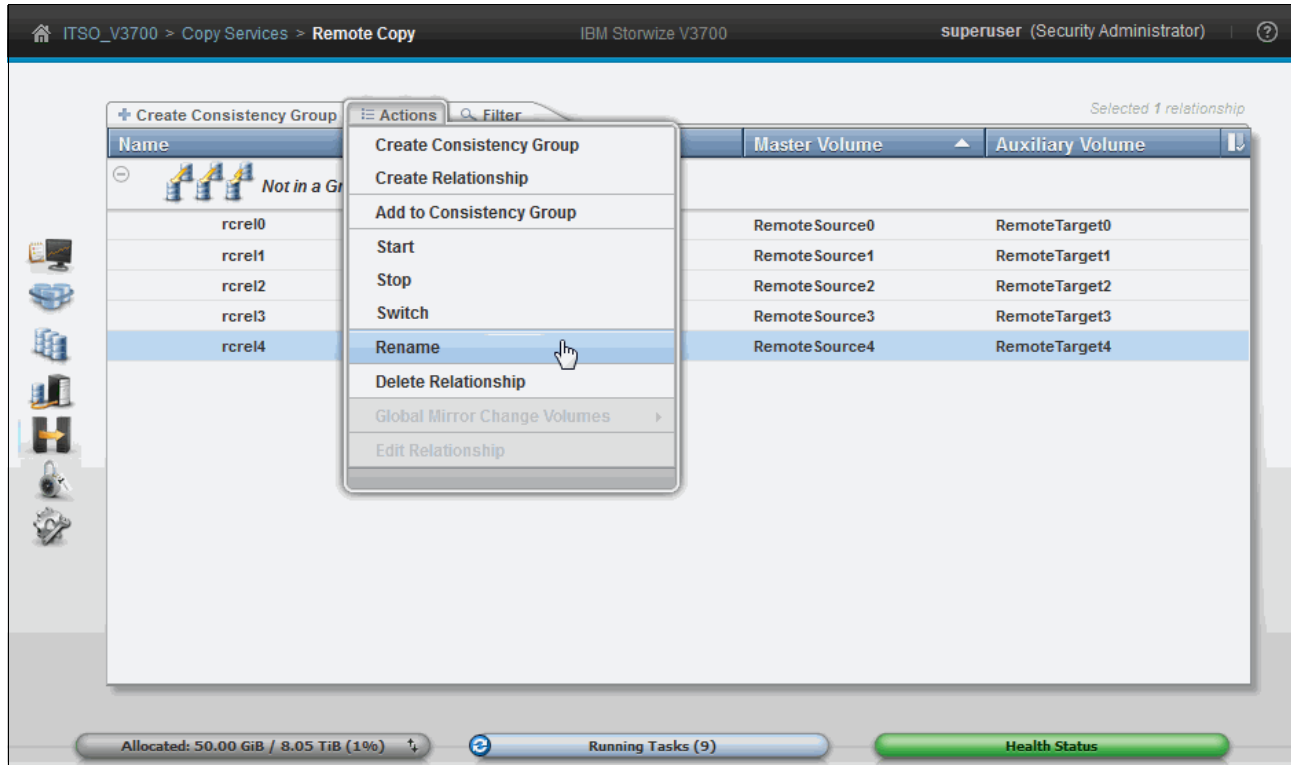


Figure 10-117 Rename the Remote Copy relationship

2. Enter the new name for the Remote Copy relationship and click **Rename**.

Deleting a stand-alone Remote Copy relationship

The Remote Copy relationship can be deleted by completing the following steps:

1. Select the relationship and select **Delete Relationship** from the **Actions** menu, as shown in Figure 10-118 on page 557, or highlight the relationship and right-click.

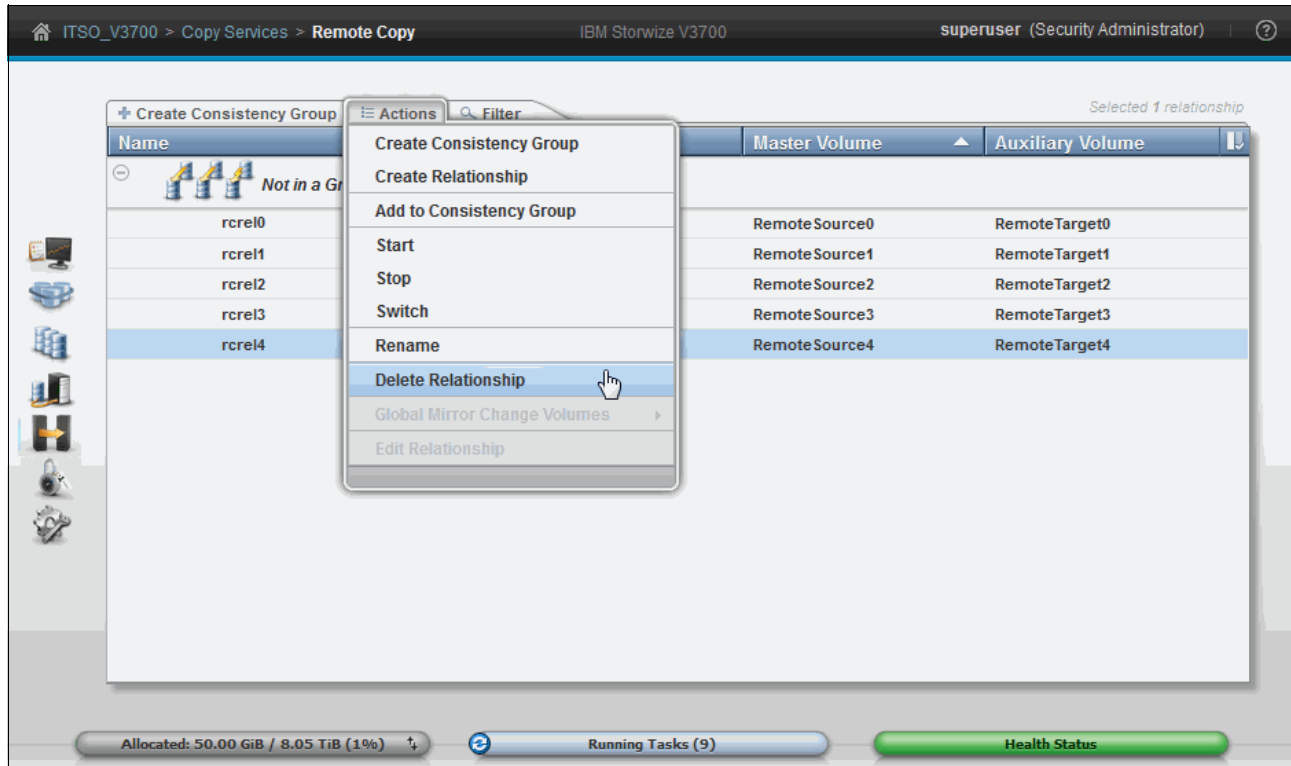


Figure 10-118 Delete a Remote Copy relationship

2. You must confirm this deletion by verifying the number of relationships to be deleted, as shown in Figure 10-119. Click **Delete** to proceed.

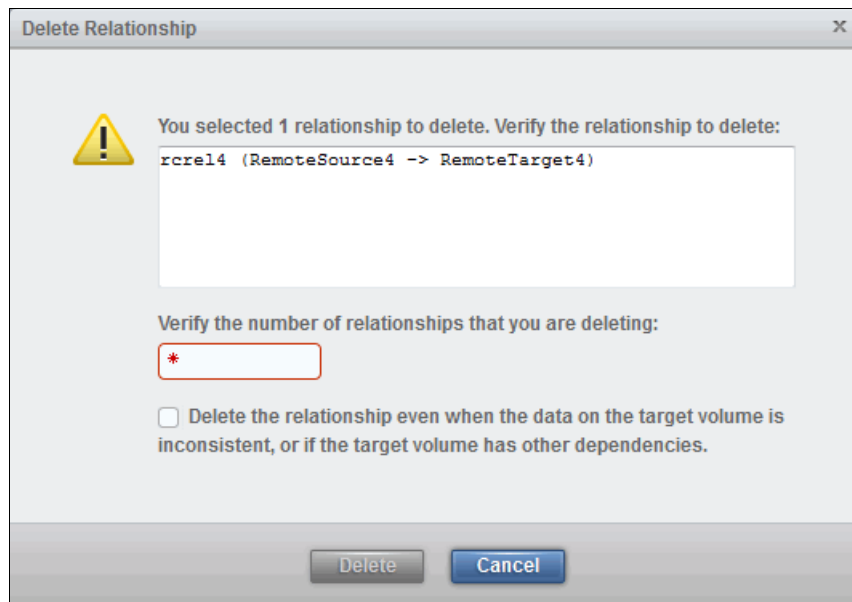


Figure 10-119 Confirm the relationship deletion

10.4.3 Managing a Remote Copy consistency group

A Remote Copy consistency group can be managed from the Remote Copy window as well.

Creating a Remote Copy consistency group

To create a Remote Copy consistency group, complete the following steps:

1. Click **Create Consistency Group**, as shown in Figure 10-120.

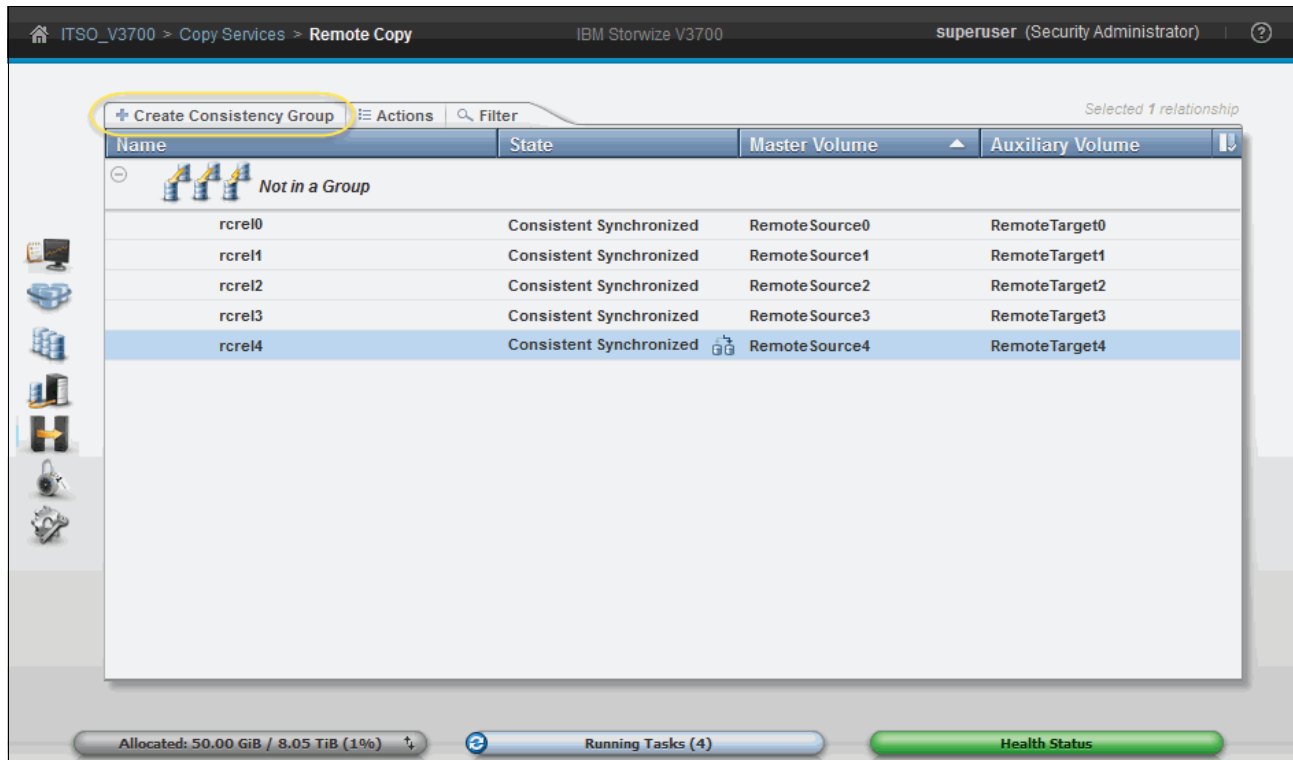


Figure 10-120 Create a Remote Copy consistency group

2. You must enter a name for your new consistency group, as shown in Figure 10-121. In this case, name it RedbookGroup.

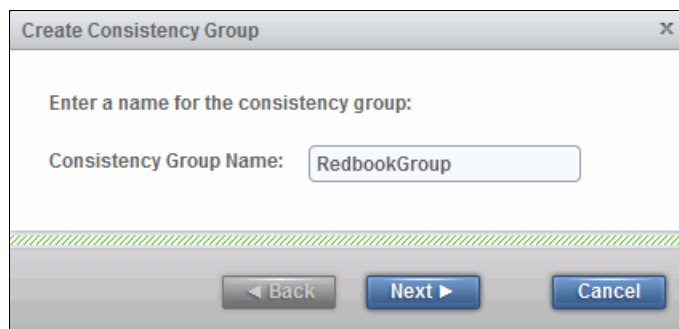


Figure 10-121 Enter a name for the new consistency group

3. You are prompted for the location of auxiliary volumes, as shown in Figure 10-122. In this example, these volumes are on another system. Select the relevant option and, from the drop-down menu, select the correct remote system. In this example, there is only one remote system defined. Click **Next** to continue.

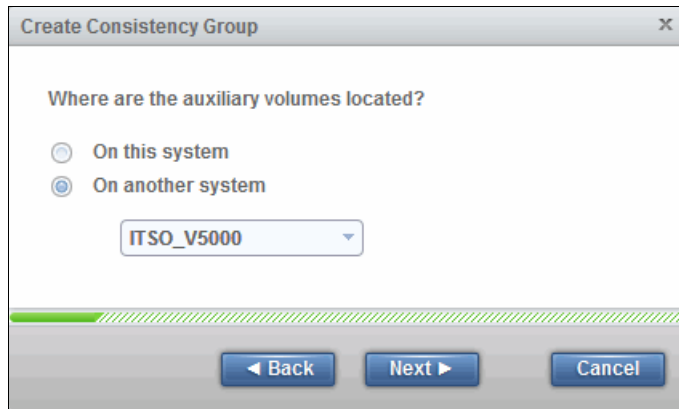


Figure 10-122 Remote Copy consistency group auxiliary volume location window

4. Next you are prompted to create an empty consistency group or add relationships to it, as shown in Figure 10-123.

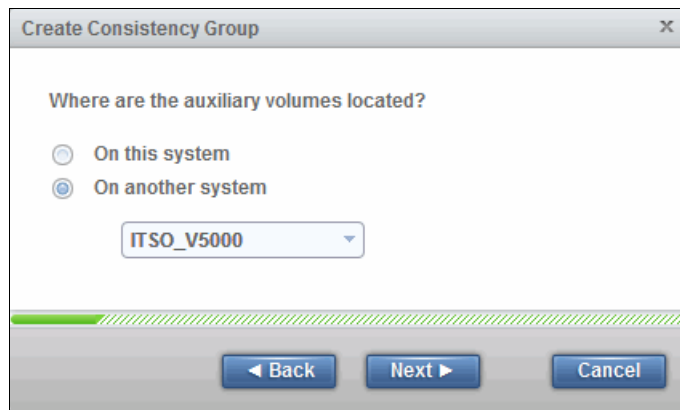


Figure 10-123 Creating an empty consistency group or adding relationships

5. If you select **No** and click **Finish**, the wizard completes and creates an empty Remote Copy Consistency Group. Selecting **Yes** prompts for the type of copy to create, as shown in Figure 10-124.

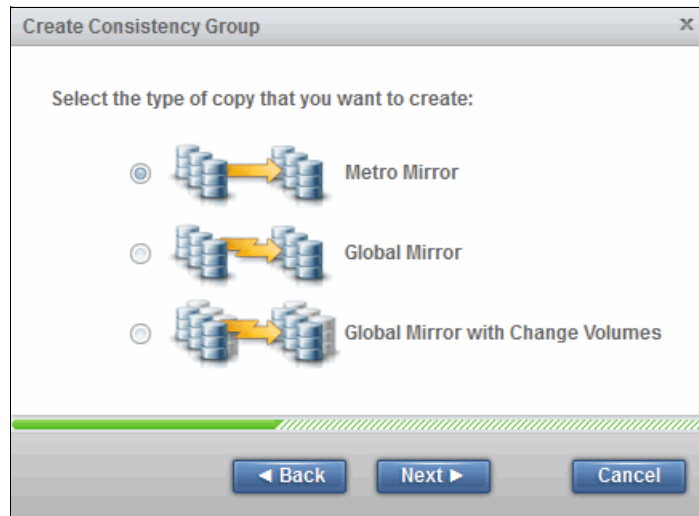


Figure 10-124 Selecting copy type: consistency group wizard

6. Choose the relevant copy type and click **Next**. A consistency group cannot contain both Metro Mirror and Global Mirror volumes. It must contain only one type or the other. In the following window, you can choose existing relationships to add to the new consistency group. Only existing relationships of the type that you chose previously (either Metro Mirror relationships or Global Mirror relationships) are displayed. This step is optional.
7. Use the Ctrl and Shift keys to select multiple relationships to add. If you decide that you do not want to use any of these relationships but you do want to create other relationships, click **Next**.

- However, if you choose relationships and then decide that you do not want any of them, you cannot remove them from the consistency group at this stage. You must stop the wizard and start again, as shown in Figure 10-125.

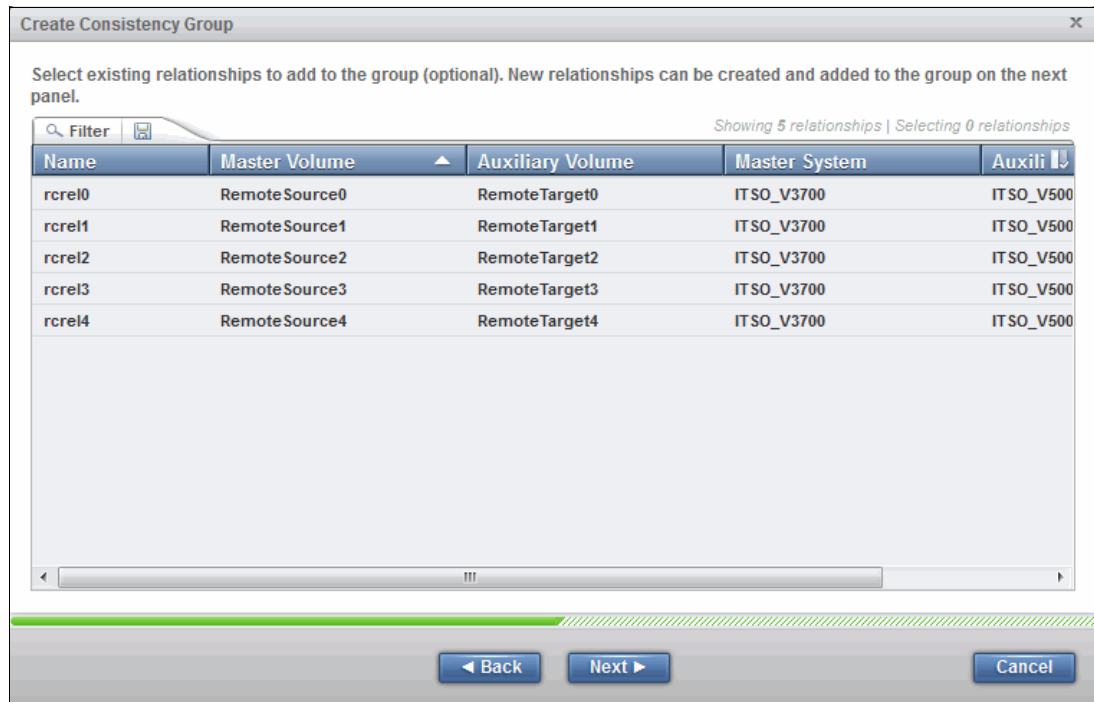


Figure 10-125 Selecting existing relationships

- The next window is optional, and gives the option to create new relationships to add to the consistency group, as shown in Figure 10-126.

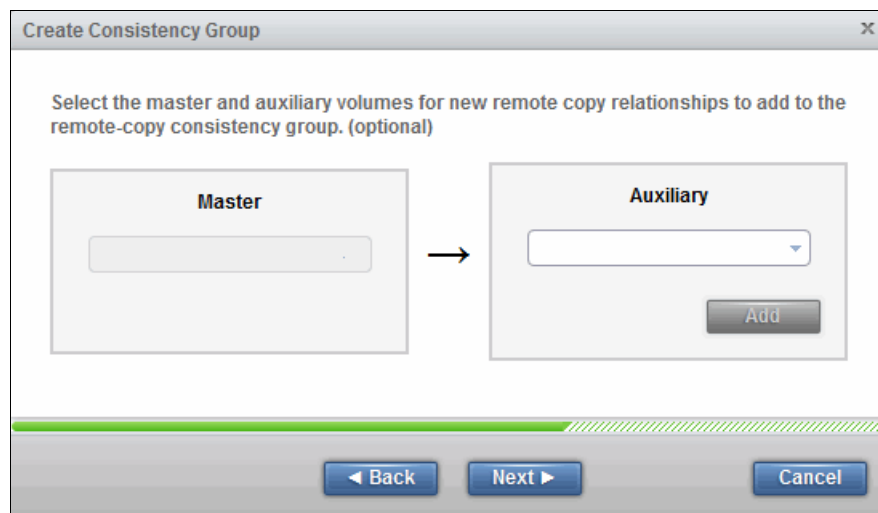


Figure 10-126 Adding new relationships in the consistency group wizard

- Select the relevant Master and Auxiliary volumes for the relationship that you want to create and click **Add**. Multiple relationships can be defined by selecting another Master and Auxiliary volume and clicking **Add** again. You can remove any incorrect choices using the red "X" next to each added relationship.

11. When you finish, click **Next**. The next window prompts for whether the relationships are synchronized, as shown in Figure 10-127.

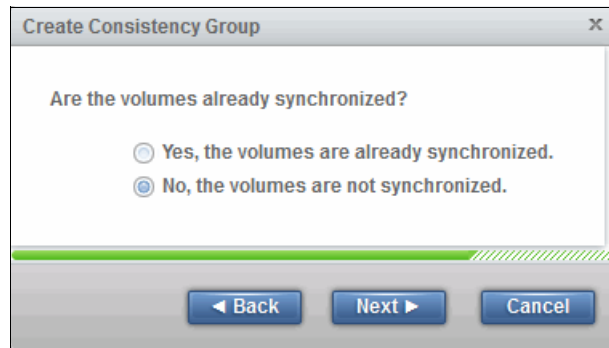


Figure 10-127 Volume synchronization

12. In the next window, you are asked whether you want to start copying the volumes now, as shown in Figure 10-128.

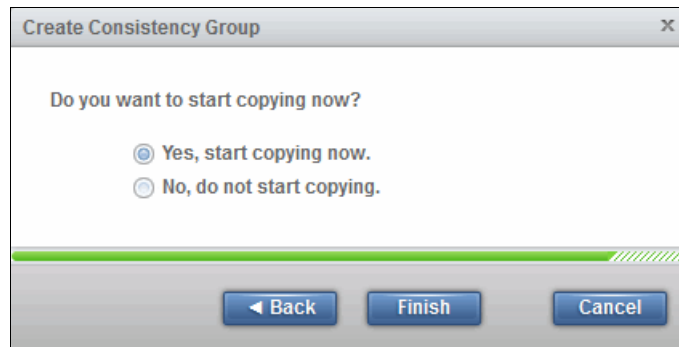


Figure 10-128 Remote Consistency group start copying option

13. After you select this option, click **Finish** to create the Remote Copy Consistency Group. You can encounter an error during the creation task if one of the relationships you requested to be added to the group is not in the correct state, because the group and any relationships must be in the same state to start with. This error is shown in Figure 10-129.

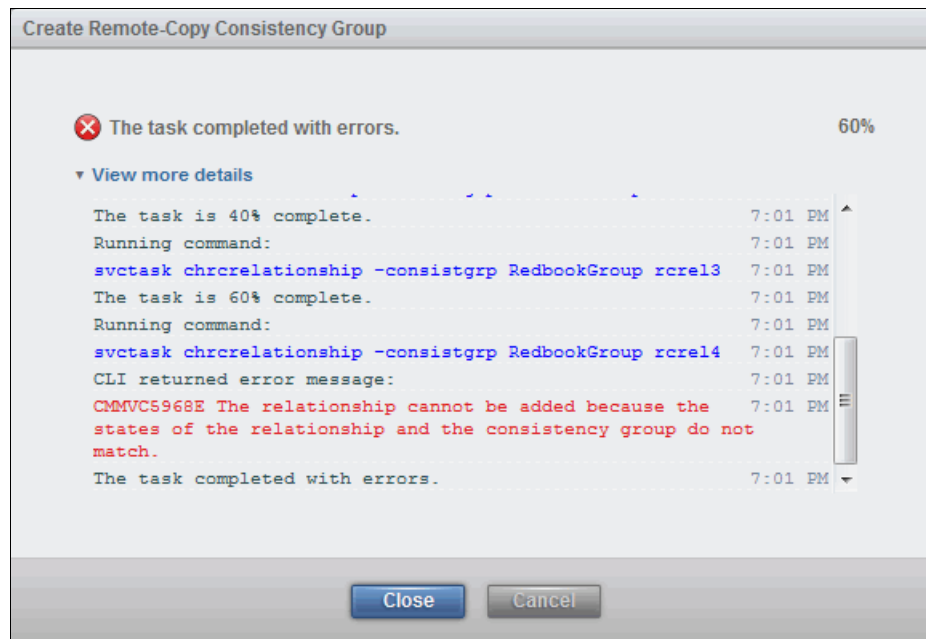


Figure 10-129 Error on creating consistency group when states do not match

- Click **Close** to close the task window, and the new consistency group is now shown in the GUI, as shown in Figure 10-130. Notice that the relationship that we tried to add, **Master Volume (RemoteSource4) → Auxiliary volume (RemoteTarget4)**, is not listed under the new consistency group, but remains under those **Not in a Group**.

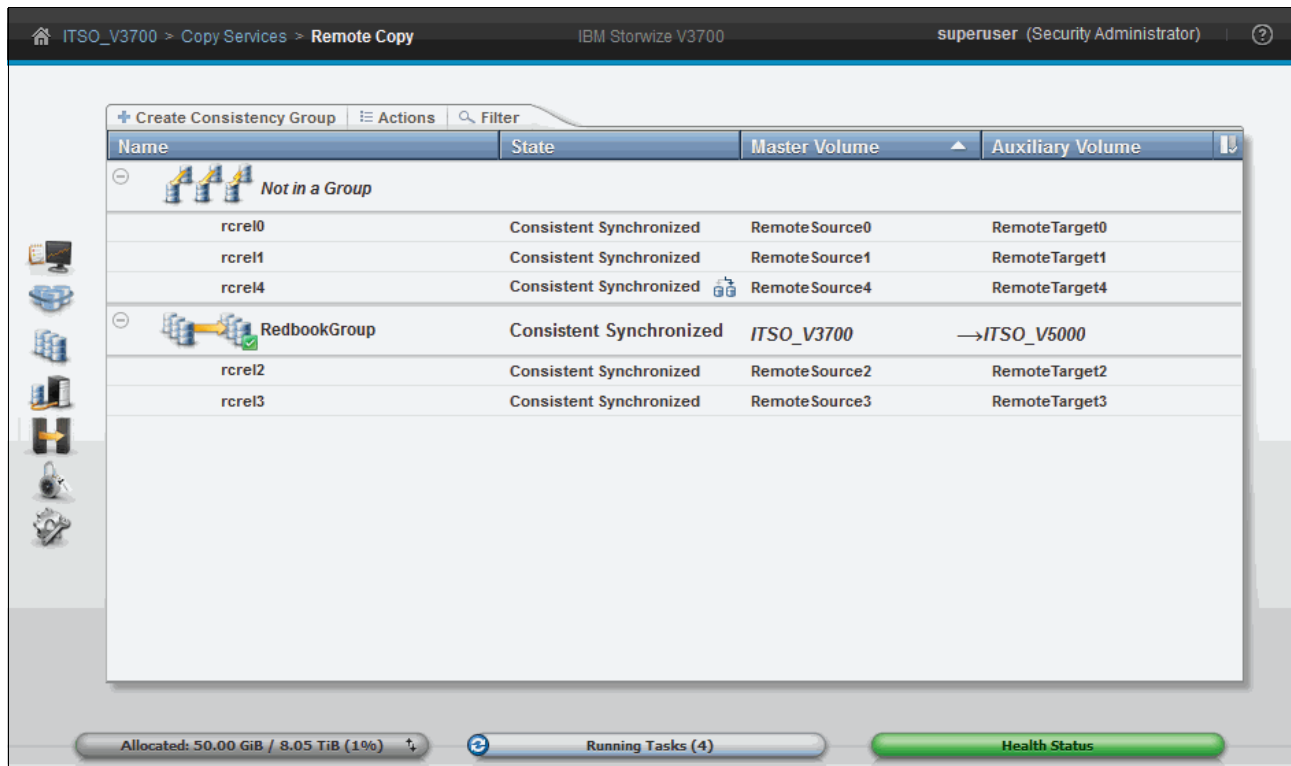


Figure 10-130 New Remote Consistency group created

In this example, we created a consistency group with two relationships. Other Remote Copy relationships can be added to the consistency group later.

Just as for individual Remote Copy relationships, each Remote Copy Consistency Group displays the name and the status of the consistency group beside the Relationship function icon. Also shown is the copy direction (in our case, this is *ITSO_V3700* → *ITSO_V5000*).

- It is easy to change the name of a consistency group by right-clicking the entry, selecting **Rename**, and then entering a new name.
- Alternatively, highlight the consistency group and select **Rename** from the **Actions** menu.

17. Similarly, beneath the Relationship function icon are the Remote Copy relationships in this consistency group. The actions on the Remote Copy relationships can be applied here by using the **Actions** menu or right-clicking the relationships, as shown in Figure 10-131.

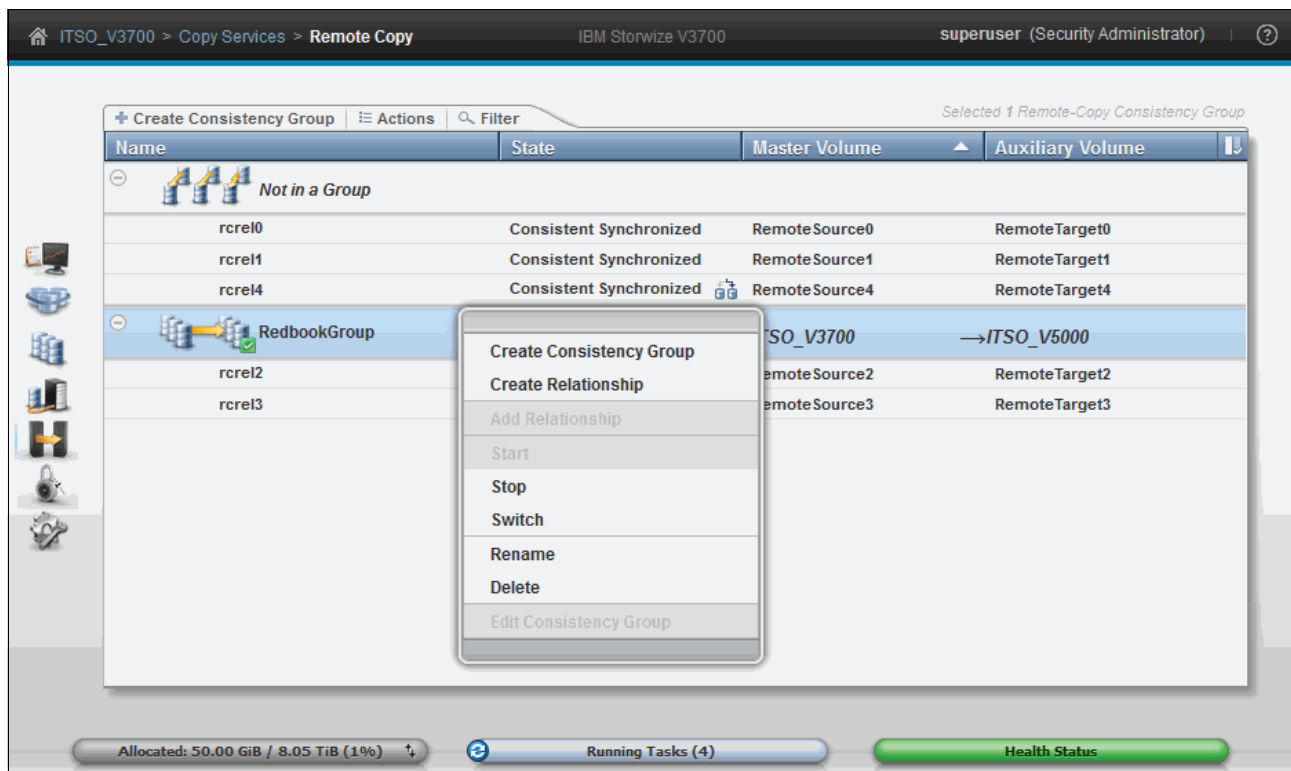


Figure 10-131 Drop-down menu options

18. Notice, however, that unlike those options for individual Remote Copy relationship, those under the control of a consistency group cannot be stopped or started individually. This is shown in Figure 10-132.

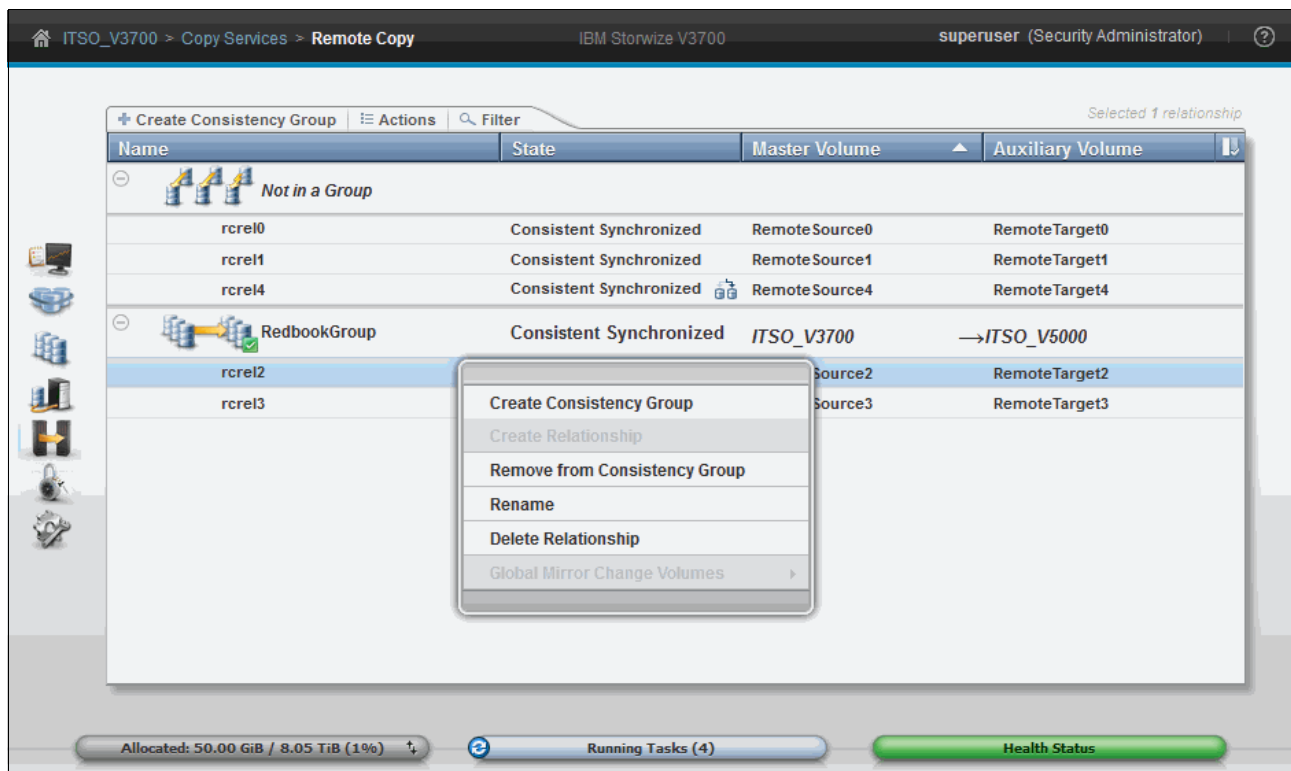


Figure 10-132 Consistency group member options

Adding Remote Copy to a consistency group

The Remote Copy relationships in the Not in a Group list can be added to a consistency group by completing the following steps:

1. Select the volumes and click **Add to Consistency Group** from the **Actions** menu, as shown in Figure 10-133. Now that the relationship that we tried to add earlier in Figure 10-126 on page 561 is in the same state as the consistency group, you can add that to the *ITSOTest* group.

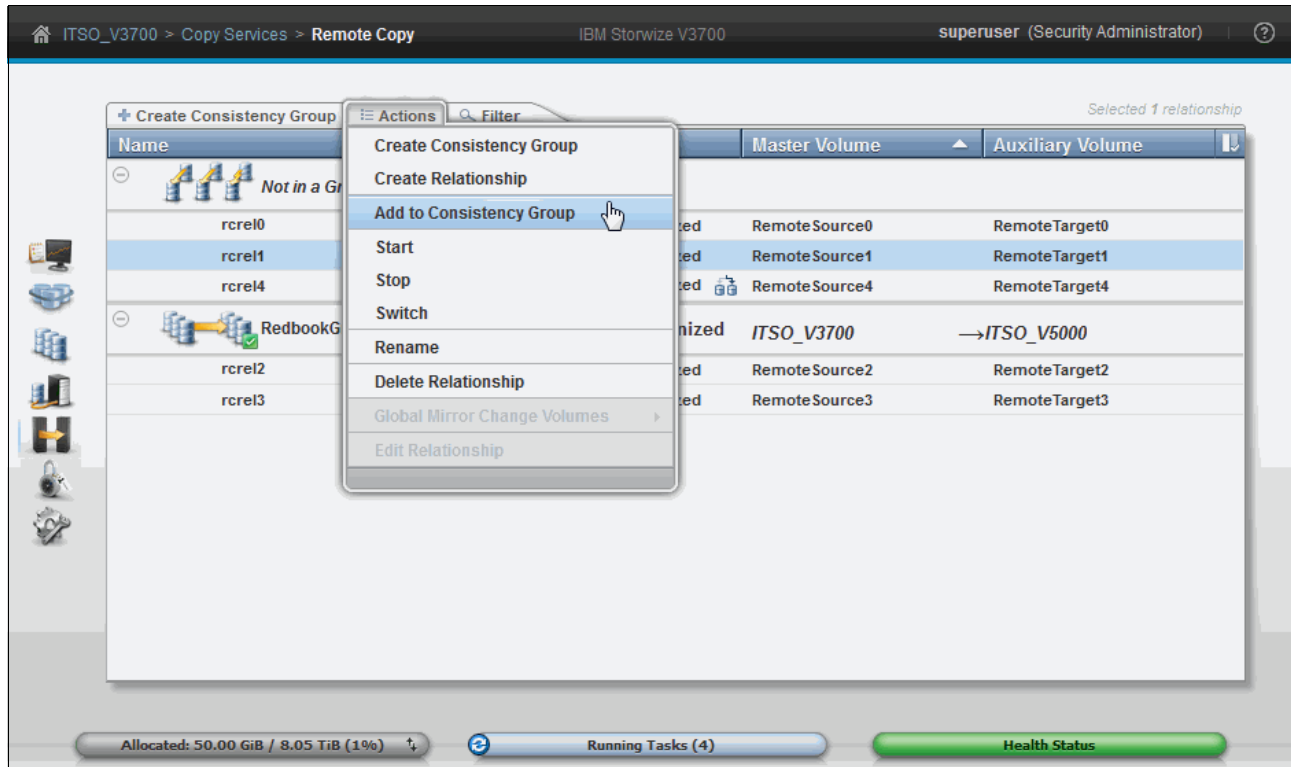


Figure 10-133 Add a Remote Copy relationship to a consistency group

2. You must choose the consistency group to which to add the Remote Copy relationships, select the appropriate consistency group and click **Add to Consistency Group**, as shown in Figure 10-134.

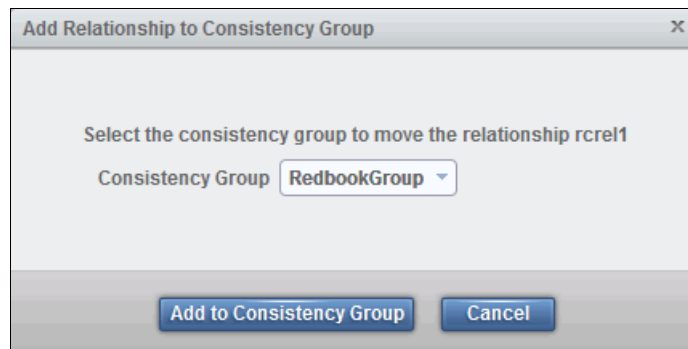


Figure 10-134 Choose the consistency group to add the remote copies

Your Remote Copy relationships are now in the consistency group that you selected. You can only add Remote Copy relationships that are in the same state as the Consistency Group to which you want to add them.

Stopping a consistency group

The Remote Copy relationship can be stopped by completing the following steps:

1. Highlight the consistency group and click **Stop** from the **Actions** menu, as shown in Figure 10-135, or right-click the consistency group.

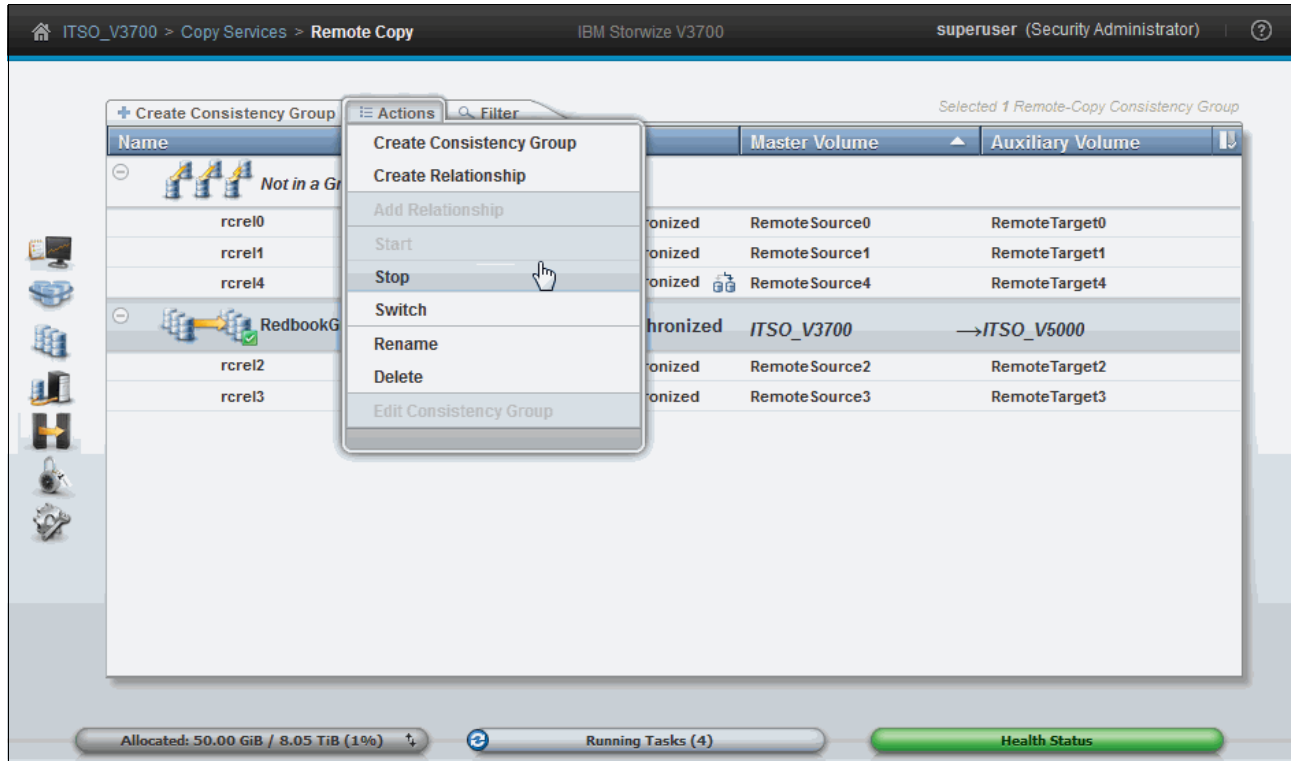


Figure 10-135 Stopping a consistency group

2. You are asked if you want to allow read/write access to secondary volumes, as shown in Figure 10-136. Make your choice and then click **Stop Consistency Group**.

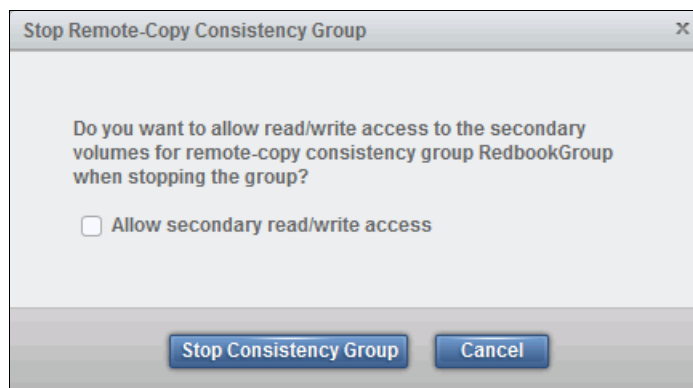


Figure 10-136 Allowing read/write access when stopping a consistency group

- The Consistency Group is now in the **Consistent Stopped** state if you chose not to allow read/write access, or in the **Idling** state if you chose to allow access, as shown in Figure 10-137.




Figure 10-137 A stopped consistency group in the Idling state

Note: The CLI differentiates between stopping consistency groups with or without access using the `-access` flag (for example, `stoprcconsistgrp -access 0`).

Starting a consistency group

The Remote Copy relationship can be started by completing the following steps:

1. Click **Start** in the **Actions** menu, as shown in Figure 10-138, or right-click the consistency group.

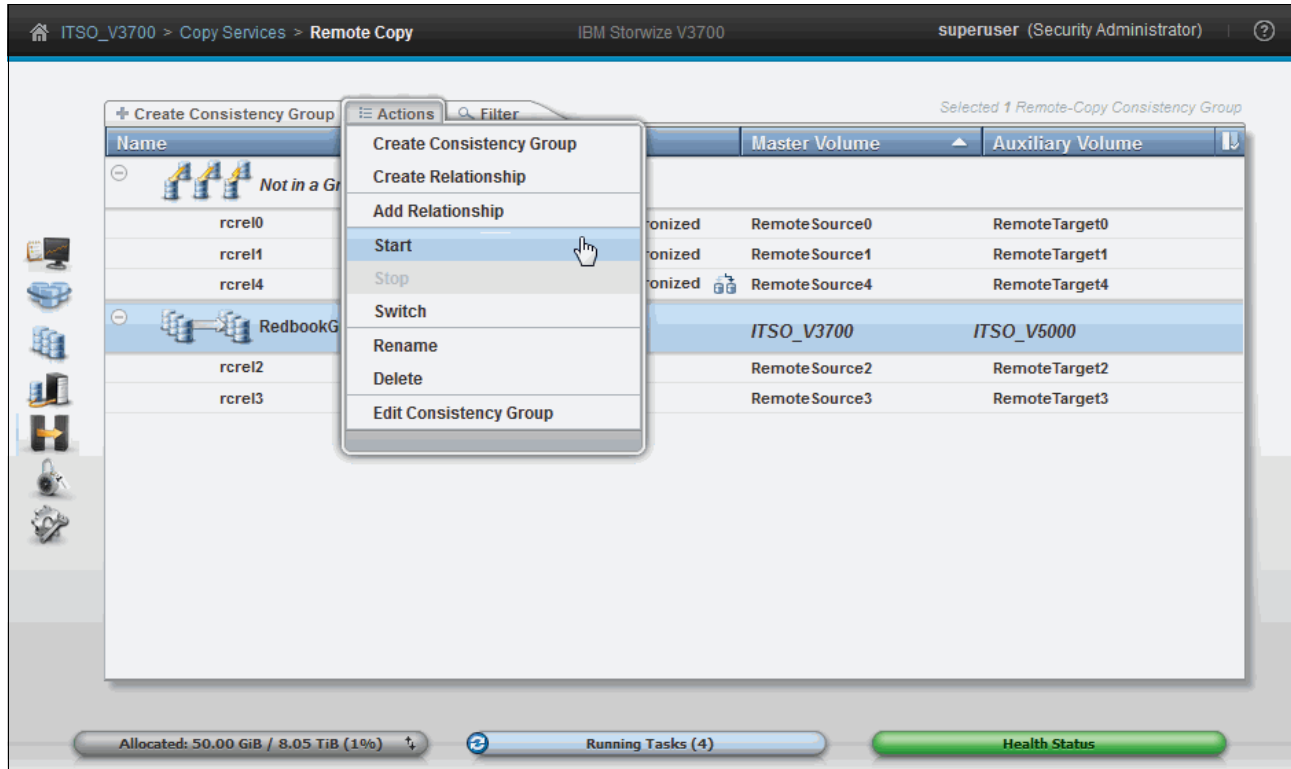


Figure 10-138 Starting the consistency group

2. You are prompted to select whether the Master or Auxiliary volumes are to act as the primary volumes before the consistency group is started as shown in Figure 10-139 on page 571. The consistency group then starts copying data from the primary volumes to the secondary volumes. If the consistency group was stopped without access and is in the **Consistent Stopped** state, you are not prompted to confirm the direction in which to start the group. It starts by default using the original primary volumes from when it was stopped.

Attention: If you are starting a Consistency Group where access was previously allowed to auxiliary volumes when the group was stopped, ensure that you choose the correct volumes to act as the Primary volumes when you restart. Failing to do so might lead to loss of data.

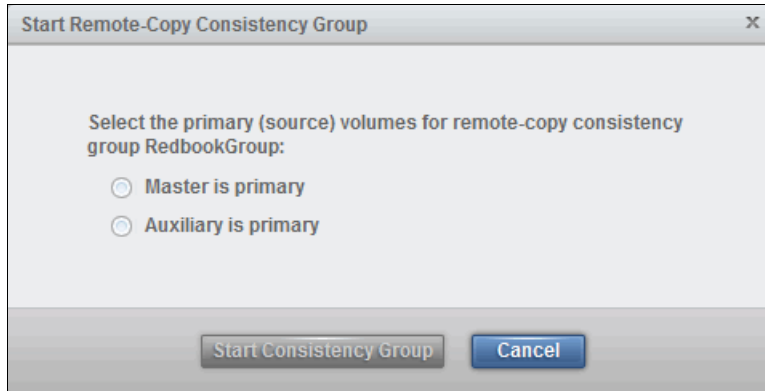


Figure 10-139 Confirming replication direction

Switching a consistency group

As with the switch action on the Remote Copy relationship, you can switch the copy direction of the consistency group:

1. Click **Switch** from the **Actions** menu, as shown in Figure 10-140.

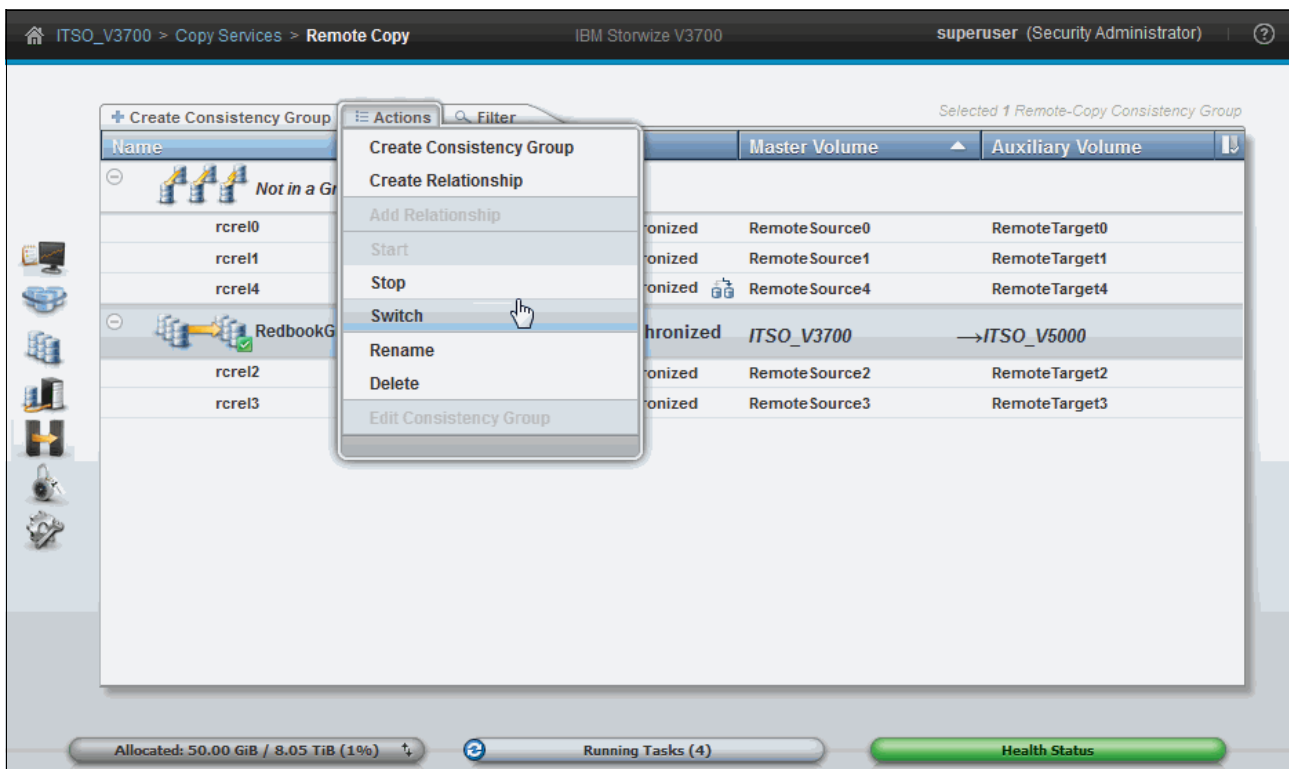


Figure 10-140 Switch the copy direction of a consistency group

2. A warning message opens, as shown in Figure 10-141. After the switch, the primary volumes in the consistency group change. Write access to current master volumes is lost, while write access to the current auxiliary volumes is enabled. This change affects host access, so make sure that these settings are what you need. If so, click **OK** to continue.

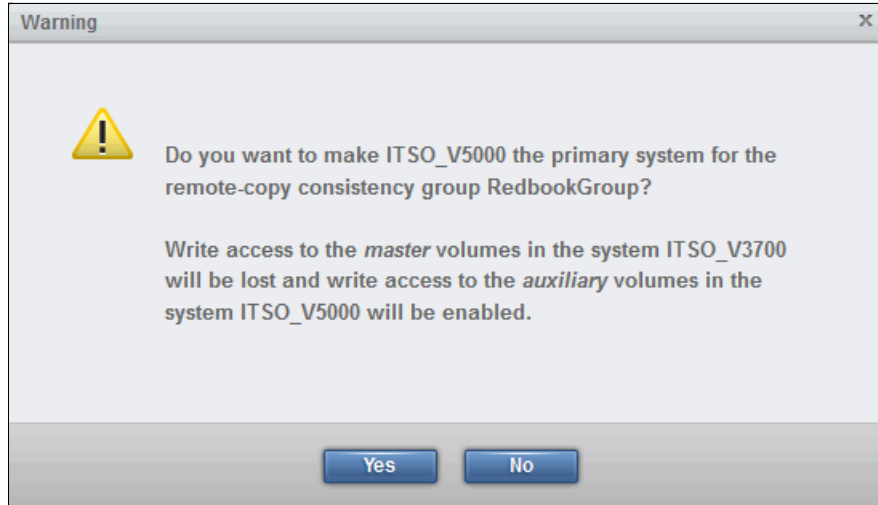


Figure 10-141 Warning message to confirm the switch

3. This option is used for disaster recovery and DR testing. Ensure that host access to the primary volumes is stopped before switching. You can then mount hosts on the auxiliary volumes and continue production from a DR site. As before with individual remote copy relationships, all of the relationships switched as part of the consistency group now show an icon indicating that they have been swapped around. This is shown in Figure 10-142.




Figure 10-142 Consistency group after a switch roles action has taken place

Removing Remote Copy relationships from a consistency group

Remote Copy relationships can be removed from a consistency group by completing the following steps:

1. Select the Remote Copy relationship and click **Remove from Consistency Group** from the **Actions** menu, as shown in Figure 10-143. Alternatively, right-click the relationship to be removed.

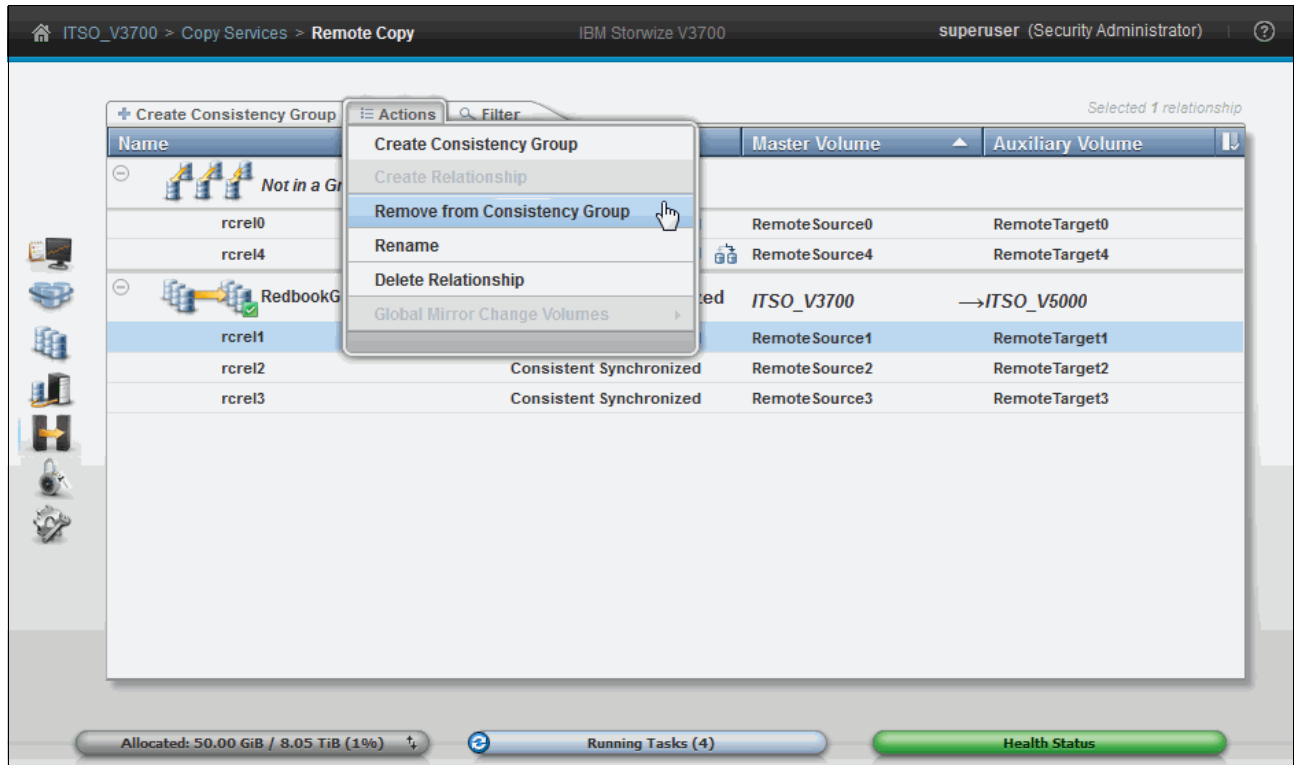


Figure 10-143 Remove Remote Copy relationships from a consistency group

2. A warning displays, as shown in Figure 10-144. Make sure that the Remote Copy relationships that are shown in the field are the ones that you want to remove from the consistency group. Click **Remove** to proceed.

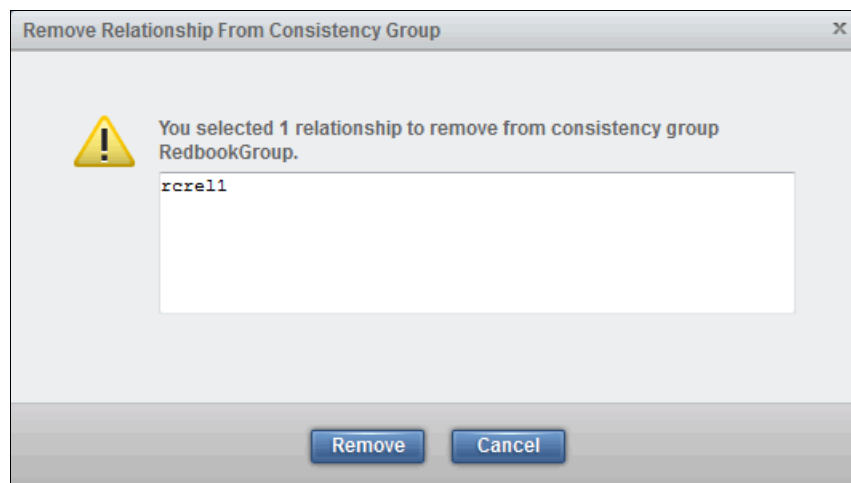


Figure 10-144 Confirm the relationships to remove from the Remote Copy consistency group

After the removal process completes, the Remote Copy relationships are deleted from the consistency group and displayed in the Not in a Group list, as shown in Figure 10-145.

Deleting a consistency group

The consistency group can be deleted by completing the following steps:

1. Click it and select **Delete** from the **Actions** menu, as shown in Figure 10-145. Alternatively, right-click the group to be deleted.

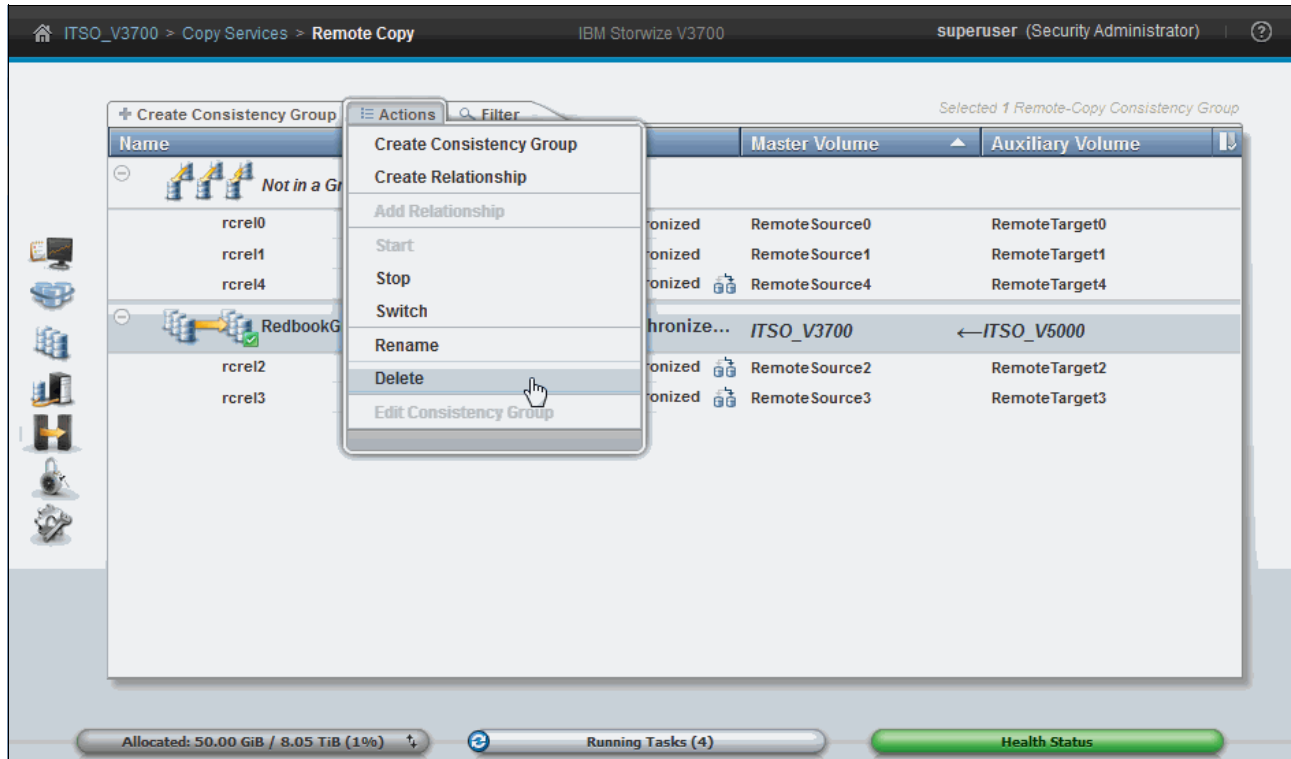


Figure 10-145 Delete a consistency group

2. You must confirm the deletion of the consistency group, as shown in Figure 10-146. Click **Yes** if you are sure that this consistency group should be deleted.

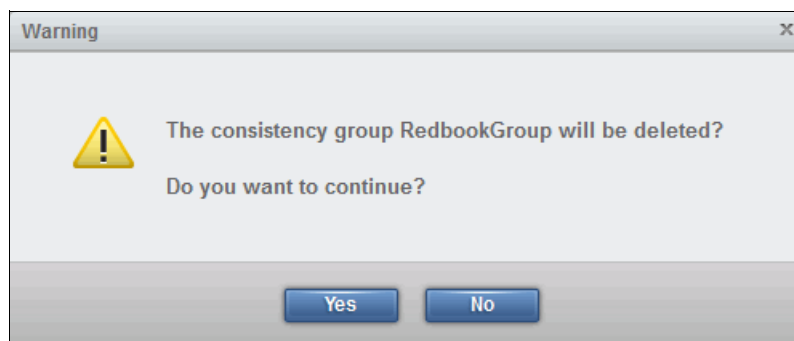


Figure 10-146 Warning to confirm deletion of the consistency group

- The consistency group is deleted. Any relationships that were part of it are returned to the Not in a Group list, as shown in Figure 10-147.




Figure 10-147 Remote Copy relationships returned to Not in a Group



RAS, monitoring, and troubleshooting

This chapter describes the reliability, availability, and serviceability (RAS) features and how to monitor and troubleshoot the IBM Storwize V3700.

This chapter includes the following topics:

- ▶ Reliability, availability, and serviceability
- ▶ IBM Storwize V3700 components
- ▶ Configuration backup procedure
- ▶ Updating software
- ▶ Event log
- ▶ Collecting support information
- ▶ Powering on and powering off the IBM Storwize V3700

11.1 Reliability, availability, and serviceability

This section describes the RAS features of the IBM Storwize V3700, in addition to reviewing monitoring and troubleshooting. RAS features are important concepts in the design of the IBM Storwize V3700. Hardware and software features, design considerations, and operational guidelines all contribute to make the IBM Storwize V3700 reliable.

Fault tolerance and a high level of availability are achieved by the following features:

- ▶ The Redundant Array of Independent Disks (RAID) capabilities of the underlying disk subsystems
- ▶ The software architecture that is used by the IBM Storwize V3700 nodes
- ▶ Auto-restart of nodes that are hung
- ▶ Battery units to provide cache memory protection in the event of a site power failure
- ▶ Host system multipathing and failover support

High levels of serviceability are achieved by providing the following benefits:

- ▶ Cluster error logging
- ▶ Asynchronous error notification
- ▶ Dump capabilities to capture software detected failures
- ▶ Concurrent diagnostic procedures
- ▶ Directed maintenance procedures
- ▶ Concurrent log analysis and memory dump data recovery tools
- ▶ Concurrent maintenance of all IBM Storwize V3700 components
- ▶ Concurrent upgrade of IBM Storwize V3700 software and microcode
- ▶ Concurrent addition or deletion of a node canister in a cluster
- ▶ Software recovery through the Service Assistant (SA) tool
- ▶ Automatic software version correction when a node is replaced
- ▶ Detailed status and error conditions that are displayed in the SA tool
- ▶ Error and event notification through Simple Network Management Protocol (SNMP), syslog, and email
- ▶ Node canister support package gathering using Universal Serial Bus (USB), in case of network connection problems

At the heart of the IBM Storwize V3700 is a redundant pair of *node canisters*. The two canisters share the data transmitting and receiving load between the attached hosts and the disk arrays.

11.2 IBM Storwize V3700 components

This section describes each of the components that make up the IBM Storwize V3700 system. Components are described in terms of location, function, and serviceability.

11.2.1 Enclosure midplane assembly

The enclosure midplane assembly is the unit that contains the node or expansion canisters and the power supply units (PSUs). The enclosure midplane assembly initially is generic and configured as a control enclosure midplane or an expansion enclosure midplane. During the basic system configuration, vital product data (VPD) is written to the enclosure midplane assembly, which decides whether the unit is a control enclosure midplane or an expansion enclosure midplane.

Control enclosure midplane

The control enclosure midplane holds node canisters and the PSUs. The control enclosure midplane assembly has specific VPD, such as worldwide node name (WWNN) 1, WWNN 2, machine type and model, machine part number, and serial number. The control enclosure midplane must be replaced only by a trained service provider. After a generic enclosure midplane assembly is configured as a control enclosure midplane, it is no longer interchangeable with an expansion enclosure midplane.

Expansion enclosure midplane

The expansion enclosure midplane holds expansion canisters and the PSUs. The expansion enclosure midplane assembly also has specific VPD, such as machine type and model, machine part number, and serial number. After a generic enclosure midplane assembly is configured as an expansion enclosure midplane, it is no longer interchangeable with a control enclosure midplane. The expansion enclosure midplane must be replaced only by a trained service provider.

Figure 11-1 shows the back of the Enclosure Midplane Assembly.




Figure 11-1 Rear view of Enclosure Midplane Assembly

For more information about replacing the control or expansion enclosure midplane, see the IBM Storwize V3700 Knowledge Center at this website:

<http://www-01.ibm.com/support/knowledgecenter/STLM5A/welcome>

11.2.2 Node canisters: Ports and light-emitting diode (LED)

There are two node canister slots along the top of the unit. The left slot is canister 1 and the right slot is canister 2.

Figure 11-2 shows the back of a fully equipped node enclosure.




Figure 11-2 Node canister

USB ports

There are two USB connectors side-by-side and they are numbered as 1 on the left and as 2 on the right. There are no indicators associated with the USB ports. Figure 11-3 shows the USB ports.




Figure 11-3 Node Canister USB ports

Ethernet ports

There are two 10/100/1000 megabits per second (Mbps) Ethernet ports side-by-side on the canister, and they are numbered 1 on the left and 2 on the right. Port 1 is required and port 2 is optional. The ports are shown in Figure 11-4.




Figure 11-4 Node canister Ethernet ports

Each port has two LEDs, and their status is shown in Table 11-1.

Table 11-1 Ethernet LEDs status

LED	Color	Meaning
Link state	Green	On: There is an Ethernet link.
Activity	Yellow	Flashing: There is activity on the link.

SAS ports

There are four 6 gigabits per second (Gbps) serial-attached Small Computer System Interface (SAS) ports side-by-side on the canister. They are numbered 1 on the left to 4 on the right. IBM Storwize V3700 uses ports 1, 2, and 3 for host connectivity, and uses port 4 to connect optional expansion enclosures. The ports are shown in Figure 11-5.




Figure 11-5 Node canister SAS ports

The SAS LED status meanings are described in Table 11-2.

Table 11-2 SAS LED Status

State	Meaning
Green	Indicates that at least one of the SAS lanes on this connector is operational. If the light is off when it is connected, there is a problem with the connection.
Amber	If the light is on, one of the following errors occurred: <ul style="list-style-type: none"> ▶ One or more (but not all) of the four lanes are up for this connector (if none of the lanes are up, the activity light is off). ▶ One or more of the up lanes are running at a different speed than the others. ▶ One or more of the up lanes are attached to a different address than the others.

IBM Storwize V3700 uses SFF-8644 mini-SAS HD connector cable to connect enclosures, as shown in Figure 11-6.




Figure 11-6 Mini-SAS HD SFF 8644 connector

Battery status

Each node canister houses a battery, the status of which is displayed on three LEDs on the back of the unit, as shown in Figure 11-7.




Figure 11-7 Node canister battery status

The battery indicator status meanings are described in Table 11-3.

Table 11-3 Battery indicator on Node canister

Color	Name	Definition
Green (left)	Battery Status	<ul style="list-style-type: none"> ▶ Fast flash: Indicates that the battery is charging and has insufficient charge to complete a single memory dump. ▶ Flashing: Indicates that the battery has sufficient charge to complete a single memory dump only. ▶ Solid: Indicates that the battery is fully charged and has sufficient charge to complete two memory dumps.
Amber (mid)	Fault	Indicates a fault with the battery.
Green (right)	Battery in use	Indicates that hardened or critical data is writing to disk.

Canister status

The status of each canister is displayed by three LEDs, as shown in Figure 11-8.




Figure 11-8 System status indicators

The system status LED meanings are described in Table 11-4.

Table 11-4 System status indicator

Color	Name	Definition
Green (left)	System Power	<ul style="list-style-type: none"> ▶ Flashing: The canister is in standby mode (the IBM Storwize V3700 code is not running). ▶ Fast flashing: The cannister is running a self-test. ▶ On: The cannister is powered on and the IBM Storwize V3700 code is running.
Green (mid)	System Status	<ul style="list-style-type: none"> ▶ Off: There is no power to the canister, the canister is in standby mode, power-on self-test (POST) is running on the canister, or the operating system (OS) is loading. ▶ Flashing: The node is in candidate or service state. It cannot perform input/output (I/O). It is safe to remove the node. ▶ Fast flash: A code upgrade is running. ▶ On: The node is part of a cluster.
Amber (right)	Fault	<ul style="list-style-type: none"> ▶ Off: The node is in candidate or active state. This state does not mean that there is no hardware error on the node. Any error that is detected is not severe enough to stop the node from participating in a cluster (or there is no power). ▶ Flashing: Identifies the canister. ▶ On: The node is in service state, or there is an error that is stopping the software from starting.

11.2.3 Node canister replaceable hardware components

The IBM Storwize V3700 node canister contains the following customer-replaceable components:

- ▶ Host interface card (HIC)
- ▶ Memory
- ▶ Battery

Figure 11-9 shows the location of these parts within the node canister.




Figure 11-9 Node canister customer replaceable parts

Host interface card replacement

For more information about the HIC replacement process, see the IBM Storwize V3700 Knowledge Center on the following website:

<http://www-01.ibm.com/support/knowledgecenter/STLM5A/welcome>

At the website, select a specific product version and browse to **Troubleshooting** → **Replacing parts** → **Replacing host interface card**.

The HIC replacement is shown in Figure 11-10.



Figure 11-10 Host Interface card replacement

Memory replacement

For more information about the memory replacement process, see the IBM Storwize V3700 Knowledge Center on the following website:

<http://www-01.ibm.com/support/knowledgecenter/STLM5A/welcome>

At the website, select a specific product version and browse to **Troubleshooting** → **Replacing parts** → **Replacing the node canister memory (2x 4 GB DIMM)**.

Figure 11-11 shows the memory location.




Figure 11-11 Memory replacement

Battery Backup Unit replacement

CAUTION: The battery is a lithium ion battery. To avoid possible explosion, do not incinerate the battery. Exchange the battery only with the part that is approved by IBM.

Because the Battery Backup Unit (BBU) is in the node canister, the BBU replacement leads to a redundancy loss until the replacement is complete. Therefore, it is advised to replace the BBU only when advised to do so. It is also suggested to follow the Directed Maintenance Procedures (DMP).

For more information about how to replace the BBU, see the IBM Storwize V3700 Knowledge Center on the following website:

<http://www-01.ibm.com/support/knowledgecenter/STLM5A/welcome>

At the website, select a specific product version and browse to **Troubleshooting** → **Replacing parts** → **Replacing battery in a node canister**.

Complete the following steps to replace the BBU:

1. Grasp the blue touch points on each end of the battery, as shown in Figure 11-12.




Figure 11-12 BBU replacement: Step 1

2. Lift the battery vertically upwards until the connectors disconnect.

Important: During a BBU change, the battery must be kept parallel to the canister system board while it is removed or replaced, as shown in Figure 11-13. Keep equal force, or pressure, on each end.



Figure 11-13 BBU replacement: Step 2

11.2.4 Expansion canister: Ports and LED

There are two expansion canister slots along the top of the unit, as shown in Figure 11-14.




Figure 11-14 Rear of an expansion enclosure

SAS ports

SAS ports are used to connect the expansion canister to the node canister, or to connect an extra expansion canister to the chain. Figure 11-15 shows the SAS ports that are on the expansion canister.




Figure 11-15 Expansion canister SAS ports

The meaning of the SAS port LEDs is described in Table 11-5.

Table 11-5 SAS LED status meaning

State	Meaning
Green	Indicates that at least one of the SAS lanes on these connectors is operational. If the light is off when connected, there is a problem with the connection.
Amber	If the light is on, one of the following errors occurred: <ul style="list-style-type: none"> ▶ One or more (but not all) of the four lanes are up for this connector (if none of the lanes are up, the activity light is off). ▶ One or more of the up lanes are running at a different speed than the others. ▶ One or more of the up lanes are attached to a different address than the others.

Canister status

Each expansion canister has its status displayed by three LEDs, as shown in Figure 11-16.




Figure 11-16 Enclosure canister status

The LED status is described in Table 11-6.

Table 11-6 Enclosure canister status

Color	Name	Definition
Green (left)	Power	Indicates that the canister is receiving power.
Green (mid)	Status	<ul style="list-style-type: none">▶ If the light is on, the canister is running normally.▶ If the light is flashing, there is an error communicating with the enclosure.
Amber (right)	Fault	If the light is solid, there is an error logged against the canister, or the firmware is not running.

11.2.5 Disk subsystem

The IBM Storwize V3700 disk subsystem is made up of control and expansion enclosures. The system can have only one control enclosure, and cannot be clustered with other control enclosures. As of IBM Storwize software V7.4, each control enclosure can attach to up to nine expansion enclosures.

This section describes the parts of the disk subsystem.

SAS cabling

Expansion enclosures are attached to control enclosures by SAS cables. There is one SAS chain, and up to nine expansion enclosures can be attached to this SAS chain. The node canister uses SAS port 4 for expansion enclosures, and ports 1, 2, and 3 for host connectivity.

Important: When an SAS cable is inserted, ensure that the connector is oriented correctly by confirming that the following conditions are met:

- ▶ The pull tab must be below the connector.
- ▶ Insert the connector gently until it clicks into place. If you feel resistance, the connector is probably oriented the wrong way. Do not force it.
- ▶ When inserted correctly, the connector can be removed only by pulling the tab.

The expansion canister has SAS port 1 for the channel input and SAS port 2 for the output, which connects to another expansion enclosure.

The SAS cabling is shown in Figure 11-17.




Figure 11-17 SAS cabling



A strand starts with an SAS initiator chip inside an IBM Storwize V3700 node canister and progresses through SAS expanders, which connect to the disk drives. Each canister contains an expander. Each drive has two ports, each of which is connected to a different expander and strand. This configuration means that both nodes directly access each drive, and there is no single point of failure.

At system initialization, when devices are added to or removed from strands (and at other times), the IBM Storwize V3700 software performs a discovery process to update the state of the drive and enclosure objects.

Slot numbers in enclosures

The IBM Storwize V3700 is made up of enclosures. There are six types of enclosures, as described in Table 11-7.

Table 11-7 Enclosure slot numbering

Enclosure type	Number of slots
Enclosure 12x 3.5-inch drives: <ul style="list-style-type: none"> ▶ Control enclosure 2072-12C ▶ Expansion enclosure 2077-12E 	Enclosure with 12 slots 
Enclosure 24x 2.5-inch drives: <ul style="list-style-type: none"> ▶ Control enclosure 2072-24C ▶ Expansion enclosure 2072-24E ▶ Direct current (dc)-powered control enclosure 2072-2DC ▶ dc-powered expansion enclosure 2072-2DE 	Enclosure with 24 slots 

Array goal

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

RAID level

An IBM Storwize V3700 disk array supports RAID 0, RAID 1, RAID 5, RAID 6, or RAID 10. Each RAID level is described in Table 11-8.

Table 11-8 RAID levels that are supported by an IBM Storwize V3700 array

RAID level	Where data is striped	Drive count (Min - Max)
0	Arrays have no redundancy, and do not support hot-spare takeover. Data is striped evenly across the drives without parity. Performance is improved at the expense of the redundancy provided by other RAID levels.	1 - 8
1	Provides disk mirroring, which duplicates data between two drives. A RAID 1 array is internally identical to a two-member RAID 10 array.	2
5	Arrays stripe data over the member drives with one parity strip on every stripe. RAID 5 arrays have single redundancy, with higher space efficiency than RAID 10 arrays, but with some performance penalty. RAID 5 arrays can tolerate no more than one member drive failure.	3 - 16

RAID level	Where data is striped	Drive count (Min - Max)
6	Arrays stripe data over the member drives with two parity strips on every stripe. A RAID 6 array can tolerate any two concurrent member drive failures.	5 - 16
10	Arrays stripe data over mirrored pairs of drives. RAID 10 arrays have single redundancy. The mirrored pairs rebuild independently. One member out of every pair can be rebuilding or missing at the same time. RAID 10 combines the features of RAID 0 and RAID 1.	2 - 16

Disk scrubbing

The scrub process runs when arrays do not have any other background processes. The process checks that the drive logical block addresses (LBAs) are readable, and that array parity is synchronized. Arrays are scrubbed independently, and each array is entirely scrubbed every seven days.

Solid-state drives

Solid-state drives (SSDs) are treated no differently to hard disk drives (HDDs) concerning RAID arrays or managed disks (MDisks). The SSDs managed by the IBM Storwize V3700 are combined into an array, usually in RAID 10 or RAID 5 format. It is unlikely that RAID 6 SSD arrays are used, because of the double parity effect, with two SSD drives used for parity only.

Drive replacement procedure

From Storwize software V7.4, it is possible to reseal or replace a failed drive in a Storwize V3700 by removing it from its enclosure and replacing with an appropriate new drive, without requiring the DMP to supervise the service action.

The system automatically performs the drive hardware validation tests. The system then promotes the drive into the configuration if these tests pass, automatically configuring the inserted drive as a spare. The status of the drive following the promotion is recorded in the event log, either as an informational message or an error (if some hardware failure occurs during the system action).

For more information about the drive replacement process, see the IBM Storwize V3700 Information Center on the following website:

<http://www-01.ibm.com/support/knowledgecenter/STLM5A/welcome>

At the website, select a specific product version and browse to **Troubleshooting** → **Replacing parts** → **Replacing a Drive assembly**.

11.2.6 Power supply unit

All enclosures require two PSUs for normal operation. A single PSU can power the entire enclosure for redundancy. Figure 11-18 shows the power supplies.




Figure 11-18 Power supply

The left side PSU is numbered 1, and the right side PSU is numbered 2.

PSU LED indicators

The PSU indicators are the same for the control and expansion units. Figure 11-19 shows the PSU LED indicators.




Figure 11-19 PSU LED Indicators

Table 11-9 shows the colors and meaning of the LEDs.

Table 11-9 PSU LED definitions

Position	Color	Marking	Name	Definition
1	Green	In	AC Status	Main power is delivered.
2	Green	DC	DC Status	DC power is available.
3	Amber	Fault exclamation mark	Fault	Fault on PSU.
4	Blue	OK	Service action that is allowed	N/A

PSU replacement procedure

For more information about the PSU replacement process, see the IBM Storwize V3700 Knowledge Center on the following website:

<http://www-01.ibm.com/support/knowledgecenter/STLM5A/welcome>

At the website, select a specific product version and browse to **Troubleshooting** → **Replacing parts** → **Replacing a Power Supply Unit**.

11.3 Configuration backup procedure

If there is a serious failure that requires the system configuration to be restored, the configuration backup file must be used. The file contains configuration data, such as arrays, pools, and volumes (but no customer data). The backup file is updated by the cluster every day and stored in the /dumps directory.

Even so, it is important to save the configuration backup file after the system configuration is changed. The command-line interface (CLI) is used to start a manual backup.

Regularly saving a configuration backup file on the IBM Storwize V3700 is important, and it must be done manually. Download this file regularly to your management workstation to protect the configuration data (a leading practice is to automate this download procedure by using a script, and to save it daily on a remote system).

11.3.1 Generating a manual configuration backup by using the CLI

To generate a configuration backup using the CLI, complete the following steps:

1. Run the **svcconfig backup** command, as shown in Example 11-1.

Example 11-1 Example for backup CLI command

```
svcconfig backup
```

2. The progress of the command is reported by dots, as shown in Example 11-2.

Example 11-2 Backup CLI command progress and output

```
.....  
.....  
CMMVC6155I SVCCONFIG processing completed successfully
```

3. The **svcconfig backup** command creates three files that provide information about the backup process and cluster configuration. These files are created in the /tmp directory on the configuration node, and can be retrieved using Battery Backup Unit (PSCP), as shown in Example 11-3.

Example 11-3 PuTTY PSCP

Using PuTTY PSCP from the CLI of an Microsoft Windows host:

```
pscp -scp superuser@<Cluster IP addr>:/tmp/svc.config.backup.xml c:\RBtemp
```

The three files that are created by the backup process are described Table 11-10.

Table 11-10 File names that are created by the backup process

File name	Description
svc.config.backup.xml	This file contains the cluster configuration data.
svc.config.backup.sh	This file contains the names of the commands that were issued to create the backup of the cluster.
svc.config.backup.log	This file contains details about the backup, including any error information that might be reported.

11.3.2 Downloading a configuration backup by using the graphical user interface

The IBM Storwize V3700 also automatically saves the configuration, daily, to the /dumps directory.

Note: Although the files in the /dump directory are the same as those generated using the CLI `svconfig backup` command, the user has no control over when they are generated.

To download a configuration backup file using the graphical user interface (GUI), complete the following steps:

1. Browse to **Settings** → **Support**, as shown in Figure 11-20.




Figure 11-20 Configuration backup open support pane

2. Select the **Show full log listing** option to list all of the available log files that are stored on the configuration node, as shown in Figure 11-21.




Figure 11-21 Configuration backup: Show full log listing

3. Search for the files named `svc.config.backup.xml_*`, `svc.config.backup.sh_*`, and `svc.config.backup.log_*`.
4. Select the files, right-click, and select **Download**, as shown in Figure 11-22.

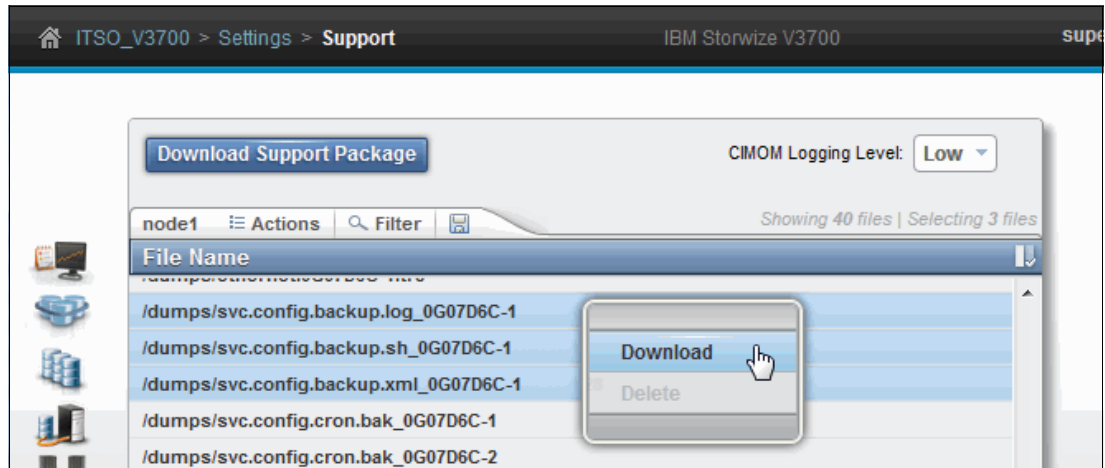


Figure 11-22 Configuration backup: Download

5. Save the configuration backup files to your management workstation, as shown in Figure 11-23.

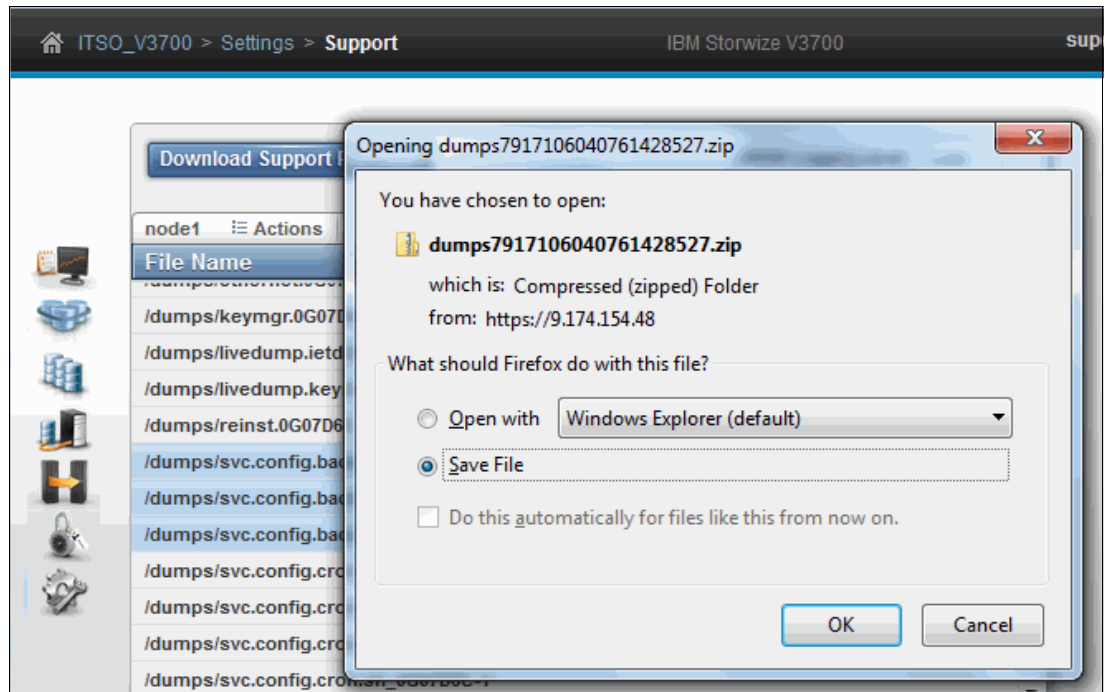


Figure 11-23 Configuration backup: Save File

Even though the configuration backup files are updated automatically on a daily basis, it might be of interest to verify the time stamp of the actual file.

6. To do this, open the `svc.config.backup.xml_xx` file with an editor and search for the string `timestamp=`, which is found near of the top of the file. Figure 11-24 shows the file and the timestamp information.

```

1 <?xml version="1.0" ?>
2
3 <!DOCTYPE xml [
4   <![ELEMENT xml ( object* ) >
5     <![ATTLIST xml label CDATA "" >
6     <![ATTLIST xml timestamp CDATA "" >
7     <![ATTLIST xml version CDATA "" >
8     <![ELEMENT object ( property* ) >
9       <![ATTLIST object type CDATA #REQUIRED >
10      <![ELEMENT property ( #PCDATA ) >
11        <![ATTLIST property name CDATA #REQUIRED >
12        <![ATTLIST property value CDATA "" >
13      ] >
14    ] >
15 <xml
16   label="Configuration Back-up"
17   version="740"
18   file_version="1.206.9.130"
19   timestamp="2014/10/15 18:22:43 BST" >
20
21 <!-- cluster section -->

```

Figure 11-24 Timestamp in backup XML file

11.4 Updating software

The system update process involves updating the entire IBM Storwize V3700 environment. The node canister software and disk firmware are separate updates, and these tasks are described separately.

11.4.1 Node canister software

Allow sufficient time to plan your tasks, review your preparatory update tasks, and complete the update of the IBM Storwize V3700 environment. The update procedures can be divided into the following general update tasks, shown in Table 11-11.

Table 11-11 Software update tasks

Sequence	Upgrade tasks
1	Decide whether you want to update automatically or manually. During an automatic update procedure, the clustered system updates each of the nodes systematically. The automatic method is the preferred procedure for updating software on nodes. However, you can update each node manually.
2	Ensure that Common Information Model (CIM) object manager (CIMOM) clients are working correctly. When necessary, update these clients so that they can support the new version of IBM Storwize V3700 code. Examples might be OS versions and options, such as IBM FlashCopy Manager or VMWare plug-ins.
3	Ensure that multipathing drivers in the environment are fully redundant. If you experience failover issues with multipathing driver support, resolve these issues before you start normal operations.
4	Update other devices in the IBM Storwize V3700 environment. Examples might include updating hosts and switches to the correct levels.
5	Update your IBM Storwize V3700.

Important: The amount of time it takes to perform a node canister update can vary depending on the amount of preparation work that is required, and on the size of the environment. Generally, to actually update the node software, allow 20 - 40 minutes per node canister, and a single 30-minute wait when the update is halfway complete. The 30-minute wait enables the recently updated node canister to come online and be confirmed as operational, and enables time for host multipathing to recover.

Some code levels support upgrades only from specific previous levels. If you upgrade to more than one level later than your current level, you might be required to install an intermediate level. See the following *Concurrent Compatibility and Code Cross Reference for Storwize V3700* support document for further details:

<https://www-304.ibm.com/support/docview.wss?uid=ssg1S1004172>

Important: Ensure that you have no unfixed errors in the log, and that the system date and time are correctly set. Start the fix procedures, and ensure that you fix any outstanding errors before you attempt to concurrently update the code.

Updating software automatically

The software update can be performed concurrently with normal user I/O operations, and each node in the system updates individually. There might be some degradation in the maximum I/O rate that can be sustained by the system while the code is uploaded to a node, during the update, while the node is rebooted, and while the new code is committed. This is because write caching is disabled during the node canister update process.

The updating node is temporarily unavailable, and all I/O operations fail to that node. As a result, the I/O error counts increase, and the failed I/O operations are directed to the partner node of the working pair. Applications do not see any I/O failures. When new nodes are added to the system, the update package is automatically downloaded to the new nodes from the original IBM Storwize V3700 system.

Multipathing requirement

Before an update, ensure that the multipathing drivers are fully redundant, with every path available and online. You might see errors that are related to the paths. These errors go away (failover), and the error count increases during the update. When the paths to the nodes return, the nodes fall back to become a fully redundant system.

GUI node canister software update process

The automatic update process is started in the GUI by starting the Update wizard, as shown in Figure 11-25:

1. Browse to **Settings** → **System** → **Update System** → **Update**.




Figure 11-25 Start Update wizard

2. When updating the node canister software, a Test utility and the node software must first be downloaded from the Internet. This can be downloaded either using the **Download** link in the pane, or manually from the IBM Support site. The Test utility verifies that there are no issues with the current system environment, such as failed components and earlier drive firmware.

3. Select the **Test utility** and **Update package** files by clicking the folder icon in the corresponding input field, and then enter the version of the code level that you are updating to, as shown in Figure 11-26.




Figure 11-26 Select Test utility and software Update package

4. Select whether an automatic or manual code update is required. The **Automatic update** option is the default and suggested choice. Figure 11-27 shows the update mode selection pane.




Figure 11-27 Update mode decision

5. If the **Service Assistant Manual update** option is selected, see “Updating software manually” on page 603.
6. Select **Finish** to start the update process on the nodes.

7. The test utility and node canister software is then uploaded to the system, as shown in Figure 11-28.

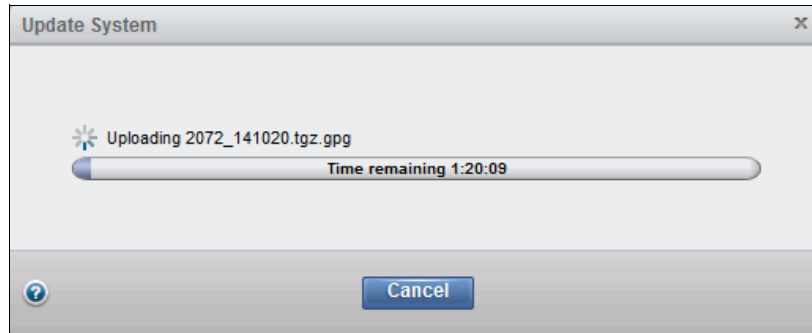


Figure 11-28 Software upload

8. When the software has been uploaded to the IBM Storwize V3700, the test utility is automatically run, as shown in Figure 11-29.

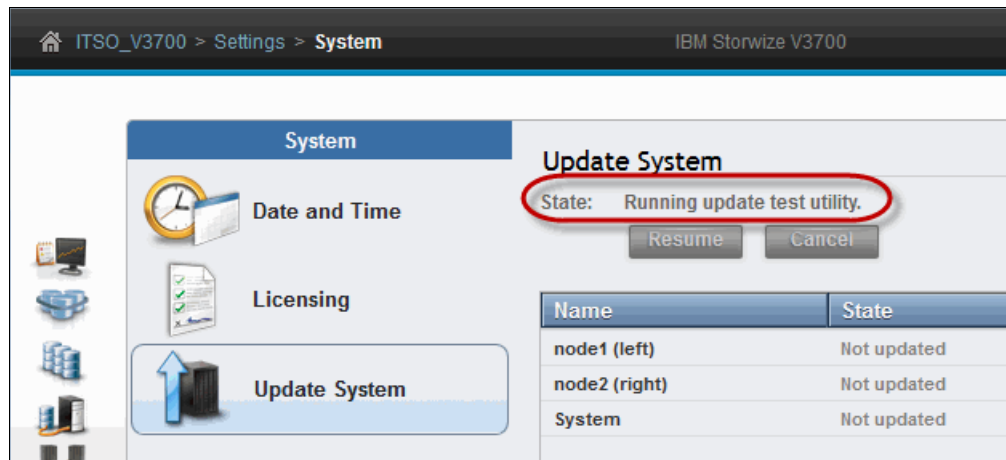


Figure 11-29 Running test utility

9. Messages inform you of any warnings or errors that the test utility finds. Figure 11-30 shows the result of the test utility finding issues with an IBM Storwize V3700 before continuing with the software update.

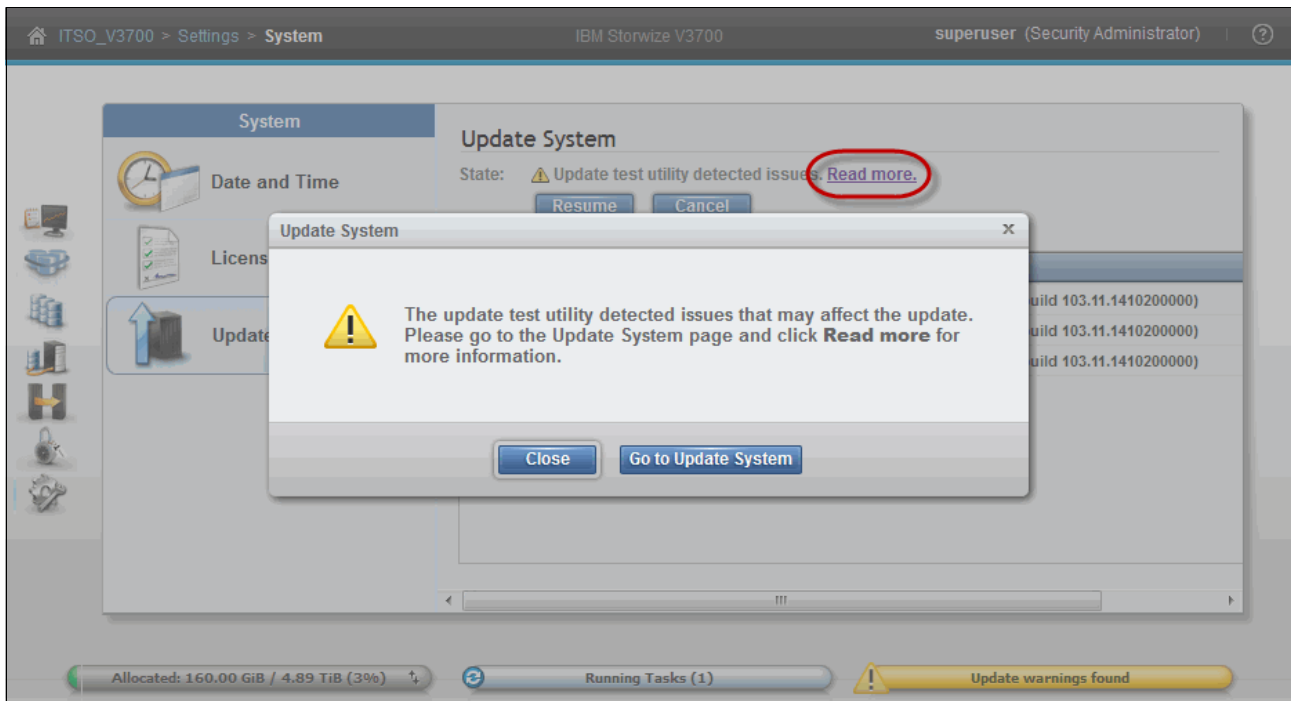


Figure 11-30 Test Utility detected issues

10. Close the warning window and then click the **Read more** link to display the results of the test utility, as shown in Figure 11-31.

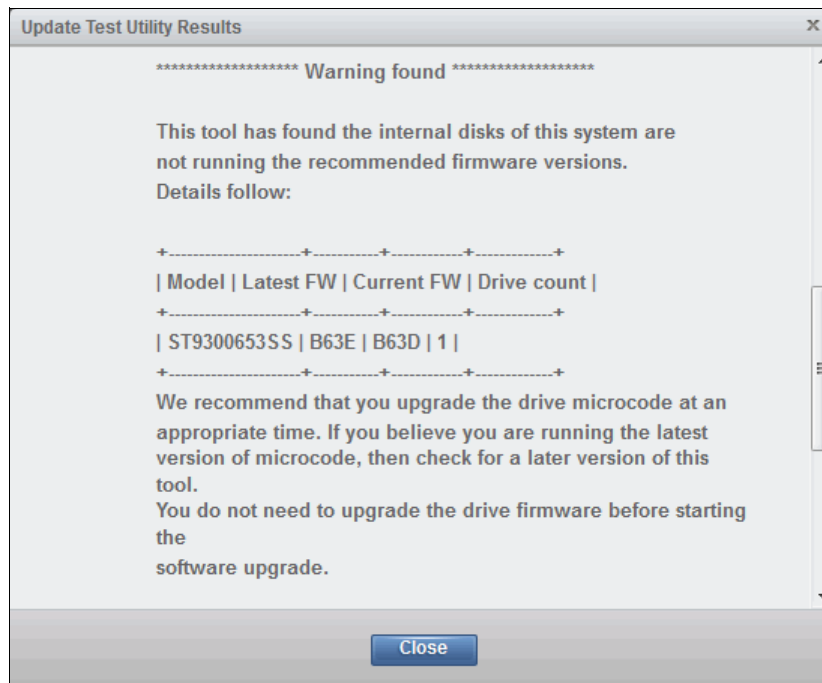


Figure 11-31 Test Utility results

In this example, the test utility has found one warning. This system has a drive with earlier firmware. This issue does not prevent the software update from continuing.

11. Clicking **Close** takes you back to the Update System pane. Click **Resume** to continue with the software update, or click **Cancel** to cancel the software update, enabling the user to correct any issues.

The results of continuing with the software update are shown in Figure 11-32.

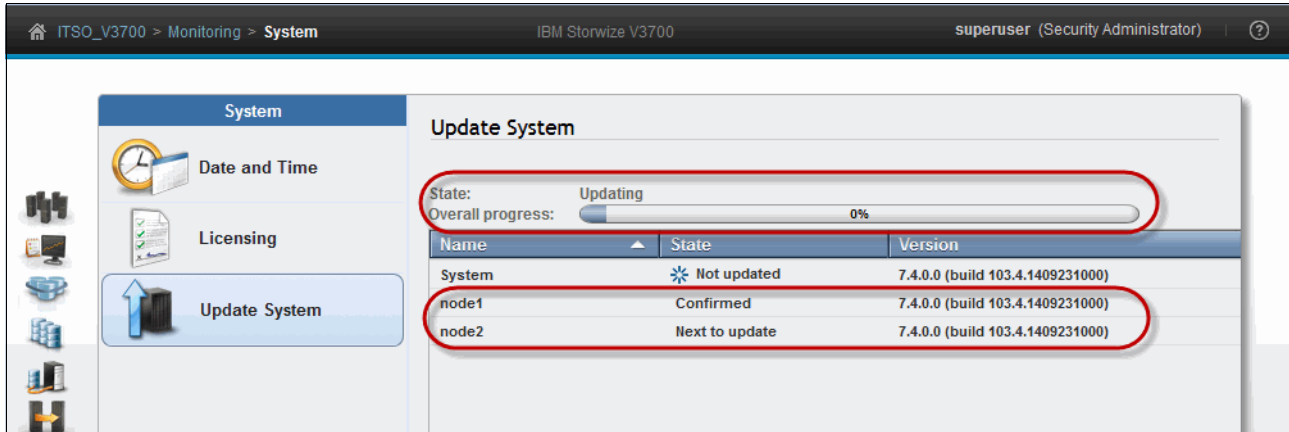


Figure 11-32 Updating system

12. If the canister software is updated to Storwize software V7.4 from V7.3 or lower, an additional confirmation step is included. Click **Confirm Update** to proceed as shown in Figure 11-33.

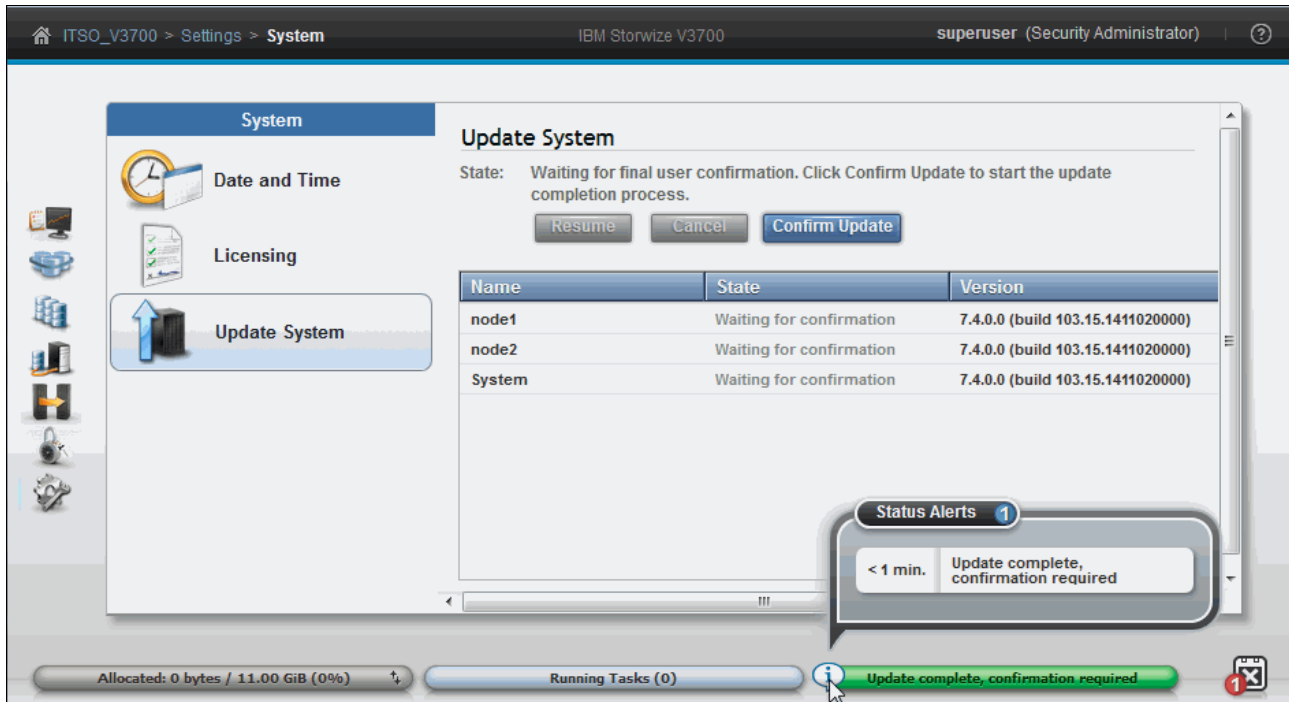


Figure 11-33 Update System: Confirm Update

Nodes are upgraded and rebooted, one at a time, until the upgrade process is complete.

Updating software manually

Important: It is highly advised to upgrade the IBM Storwize V3700 automatically by following the Update wizard. If a manual update is used, make sure that you do not skip any steps.

The steps for a manual update are shown in the Update System help. See Figure 11-34.

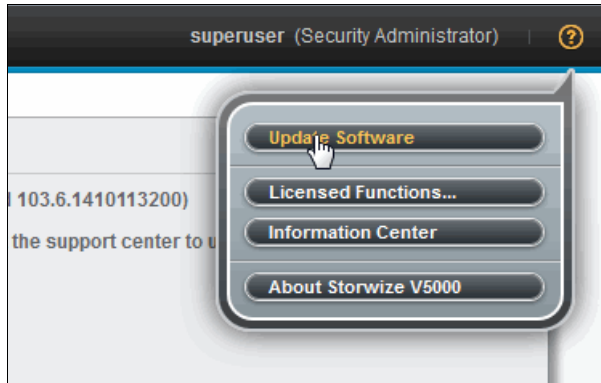


Figure 11-34 Update Software help

Complete the following steps to manually upgrade the software:

1. In the management GUI, click **Settings** → **System** → **Update System**, and run the **Update** wizard.
2. When updating the node canister software, a Test utility and the node software must first be downloaded from the internet. This can be downloaded either using the **Download** link within the pane or manually from the IBM Support site. The Test utility verifies that there are no issues with the current system environment, such as failed components and earlier drive firmware.
3. Select the Test utility and Update package files by clicking the folder icon in the corresponding input field. Enter the version of the code level that you are updating to and click **Update** to continue, as shown in Figure 11-35.

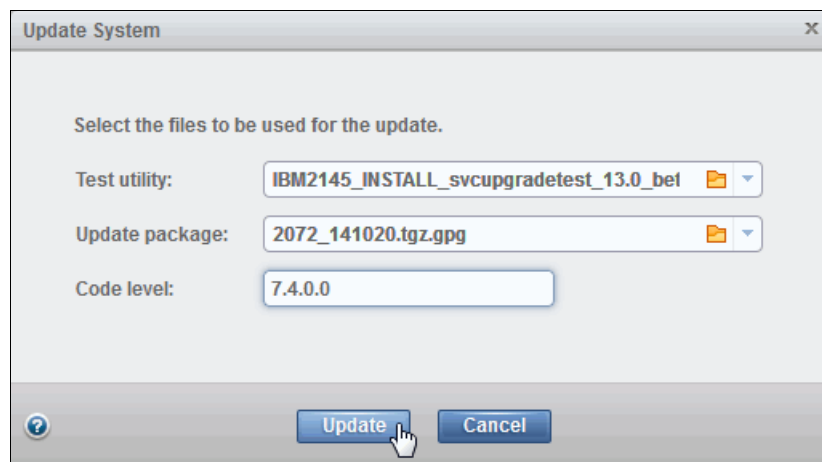


Figure 11-35 Select Test utility and software Update package

4. At the next pane, select **Service Assistant Manual update**, as shown in Figure 11-36.

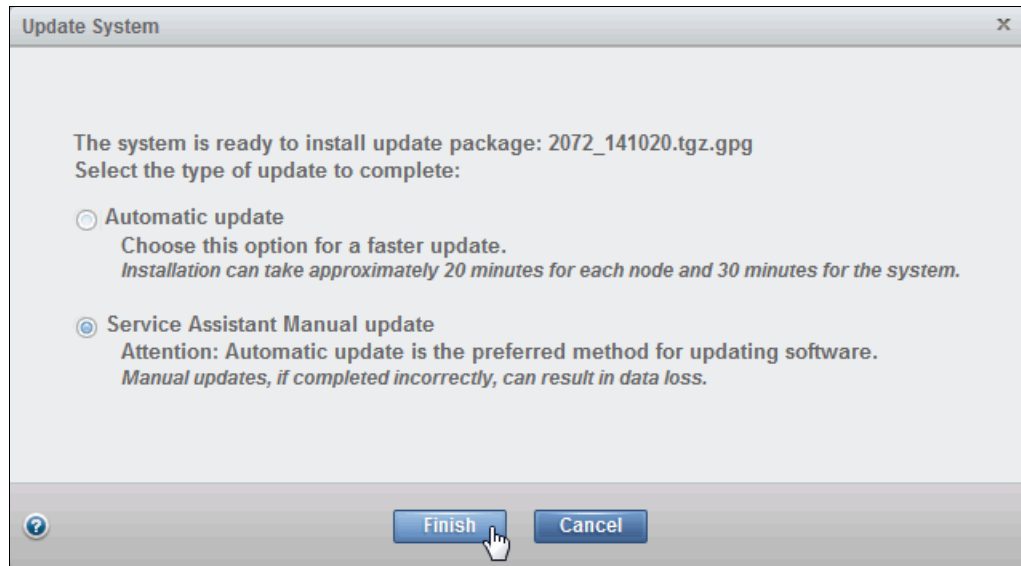


Figure 11-36 Select manual update mode

5. Select **Finish** to start the update process on the nodes.
The test utility and node canister software are then uploaded to the system, as shown in Figure 11-37.




Figure 11-37 Software upload

- After the software has been uploaded, the test utility is automatically run and, if there are no issues, the system is placed in a **Prepared** state, as shown in Figure 11-38.

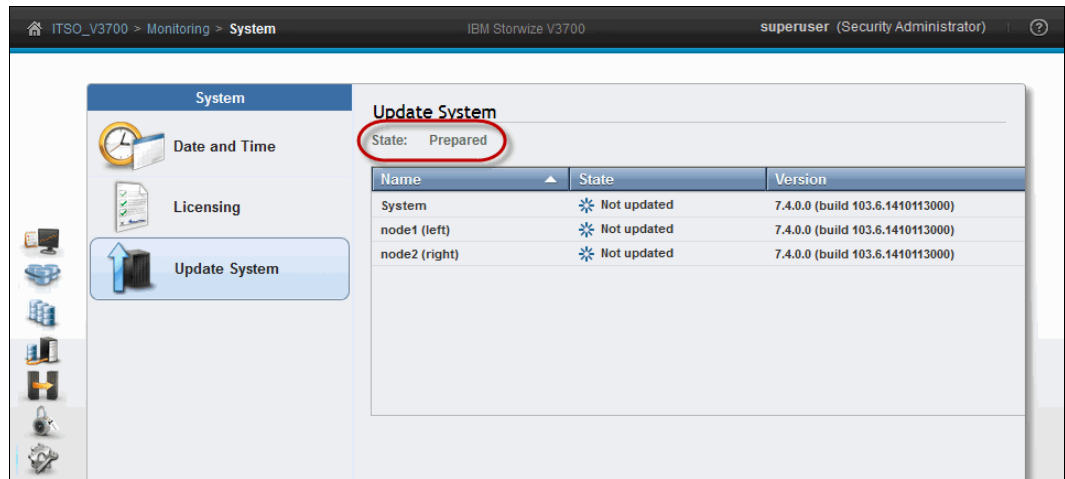


Figure 11-38 Update System status: Prepared state

- Non-configuration nodes are updated first, leaving the configuration node until last. In the management GUI, select **Monitoring** → **System** and hover over the canisters to confirm which is or are the non-configuration node or nodes, as shown in Figure 11-39.

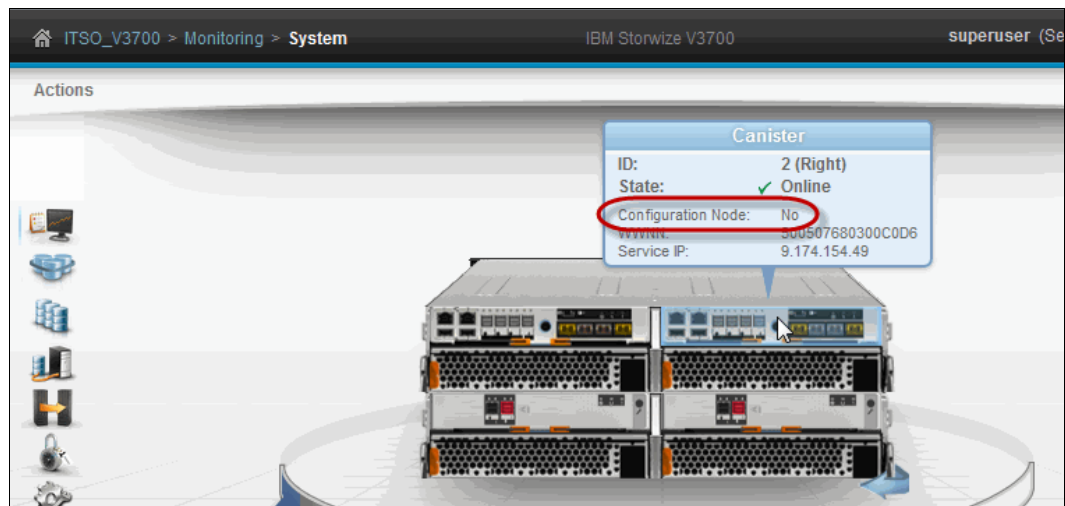


Figure 11-39 Checking the Configuration node status

8. Select the canister that contains the node that you want to update and select **Remove**, as shown in Figure 11-40.




Figure 11-40 Remove the non-config node from the cluster

Important: Make sure that you select the non-configuration nodes first.

9. A warning message displays, as shown in Figure 11-41. Click **Yes** to continue.

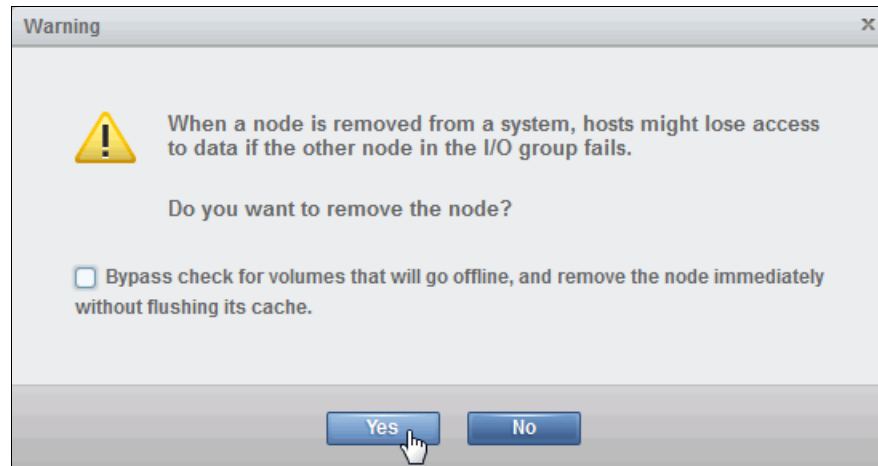


Figure 11-41 Remove node warning message

10. The non-configuration node is removed from the Management GUI Update System pane and shown as Unconfigured when hovering over the node in the **Monitoring** → **System** pane.
11. Open the Service Assistant (SA) tool for the node that you just removed.

12. In the SA tool, the node that is ready for upgrade must be selected. Select the node that shows Node status as service mode, displays an error code of 690, and has no available cluster information, as shown in Figure 11-42.

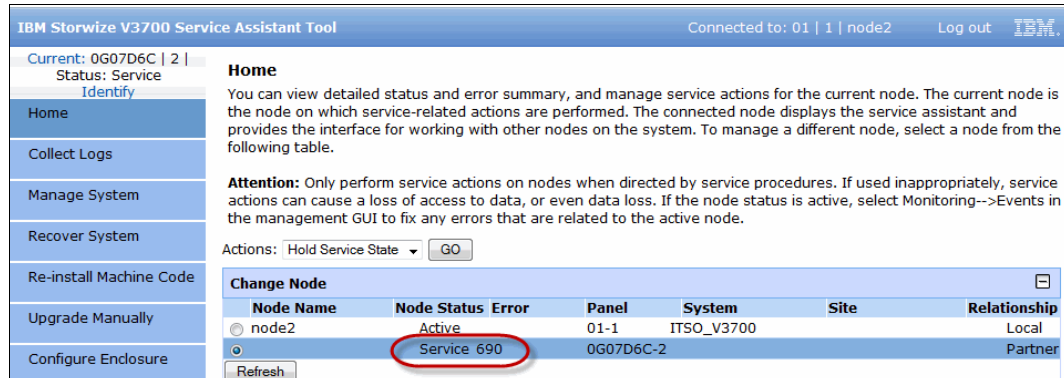


Figure 11-42 Select node in Service GUI for update

13. In the SA tool, select **Upgrade Manually**, and select the required node canister software upgrade file, as shown in Figure 11-43.

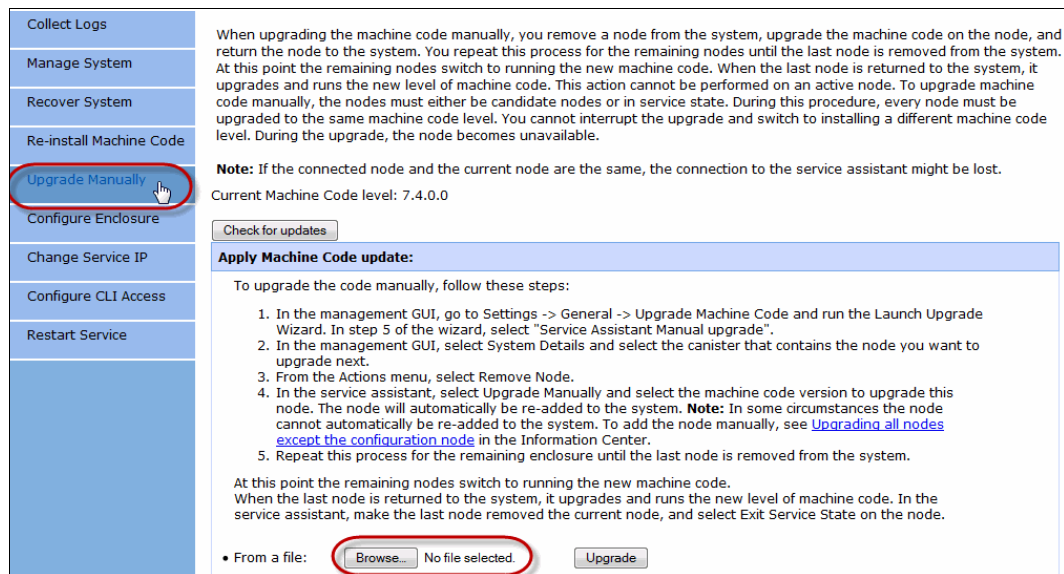


Figure 11-43 Select software file for upgrade

14. Click **Upgrade** to start the upgrade process on the first node.

The node is reintroduced automatically into the system after upgrade. Upgrading and adding the node again can take 20 - 40 minutes, as shown in Figure 11-44.

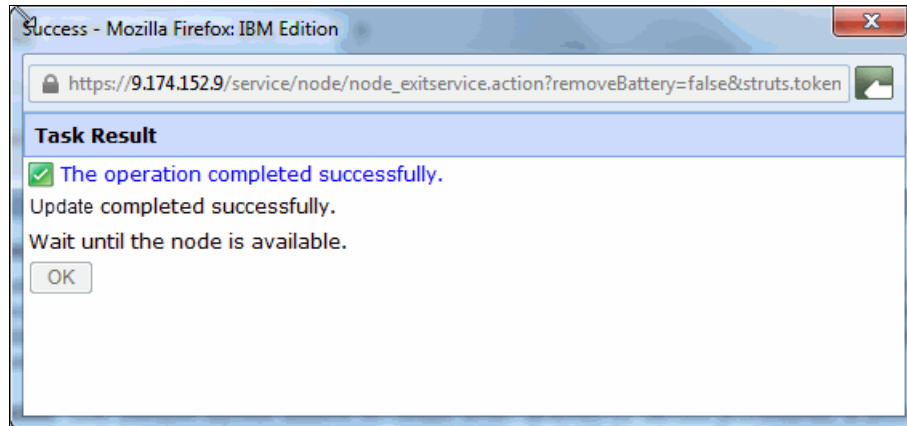


Figure 11-44 Non-configuration node update complete

The management GUI also displays the progress of the update, as shown in Figure 11-45.

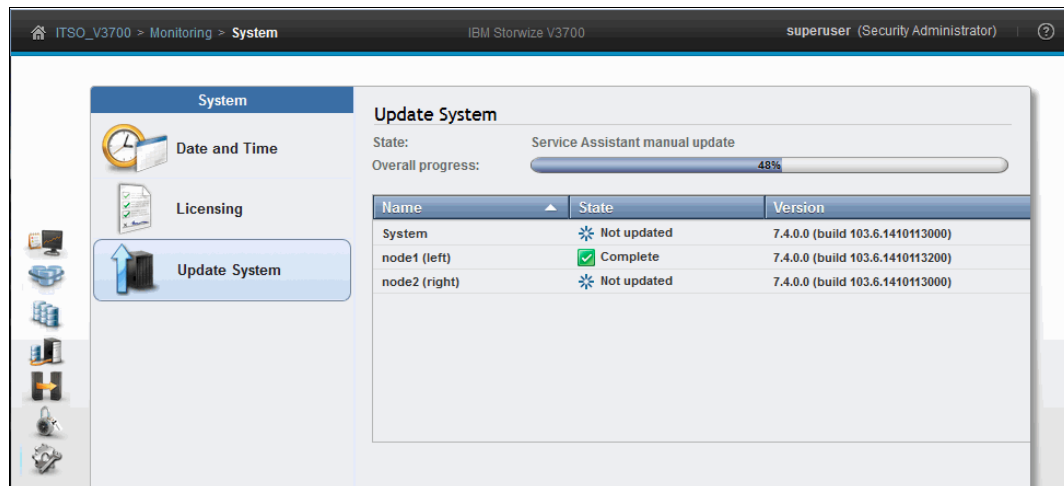


Figure 11-45 Manual node software update progress

15. Repeat steps 1 - 14 for the remaining node (or nodes), leaving the Configuration node until last.

16. After you remove the configuration node from the cluster for the update, a warning displays, as shown in Figure 11-46. Click **Yes** to continue.

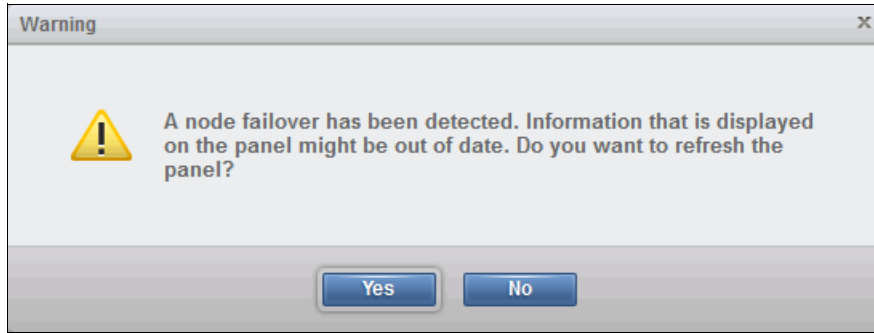


Figure 11-46 Configuration node failover warning

Important: The configuration node remains in Service State when it is readed to the cluster. Therefore, exit Service State manually.

17. To exit from service state, browse to the home pane of the SA tool and open the Action menu. Select **Exit Service State**, as shown in Figure 11-47.

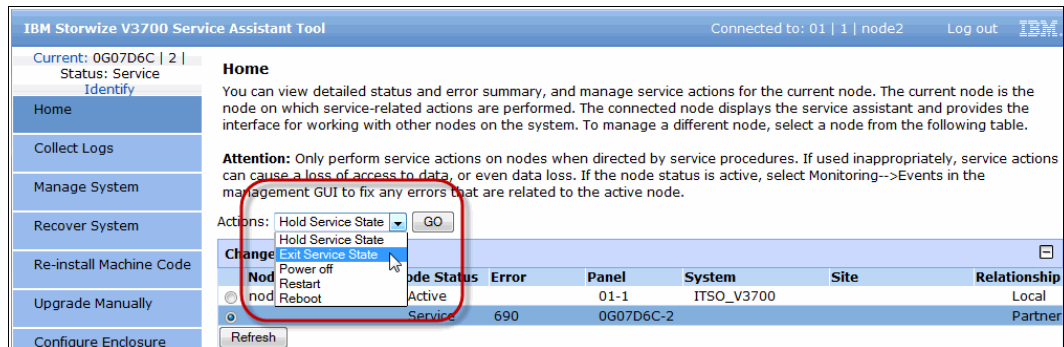


Figure 11-47 Exit service state to add node back in cluster

Both the nodes are now back in the cluster (as shown in Figure 11-48) and the system is running on the new code level.




Figure 11-48 Cluster is active again and running the new code level

18. If the canister software is updated to Storwize software V7.4 from V7.3 or lower, an extra confirmation step is included. Click **Confirm Update** to proceed as shown in Figure 11-49.

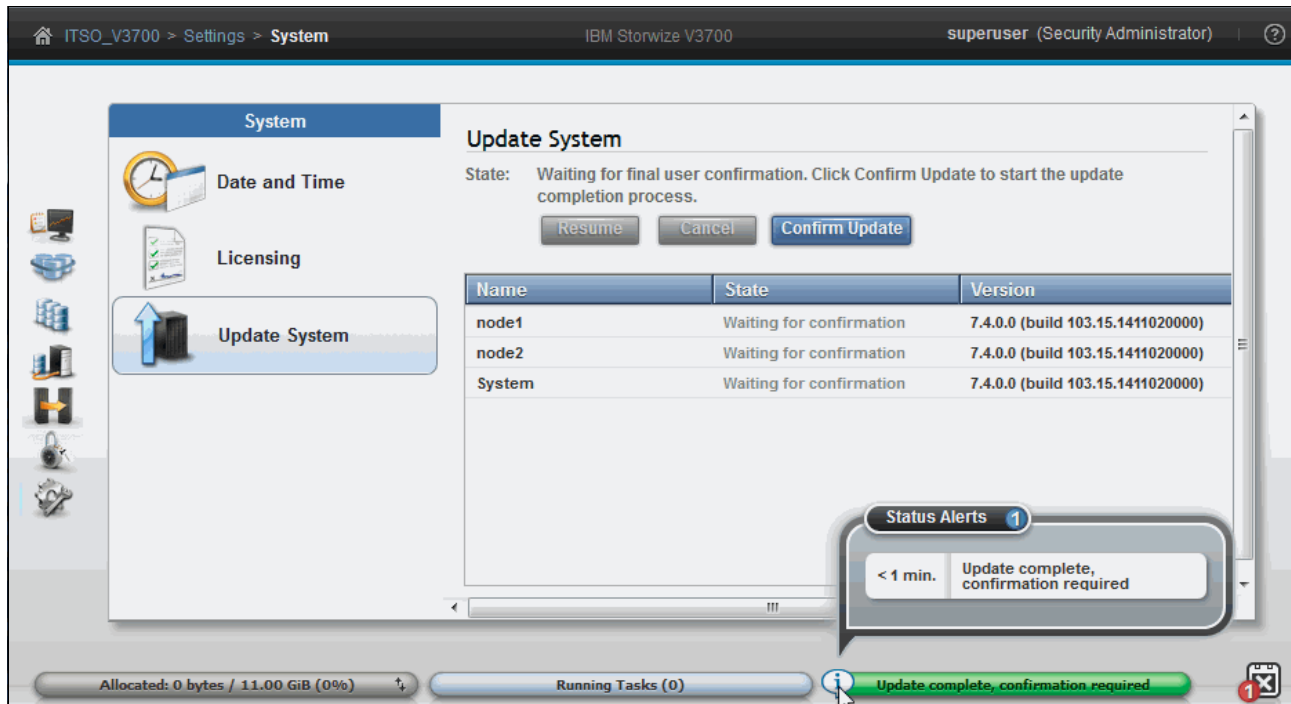


Figure 11-49 Update System: Confirm Update

11.4.2 Upgrading drive firmware

In previous versions of IBM Storwize code, it was only possible to upgrade the drive firmware from the CLI using the `applydrivesoftware` command, after manually uploading the firmware files to the `/home/admin/upgrade` directory of the IBM Storwize V3700. This process is detailed in Appendix A, “Command-line interface setup and SAN Boot” on page 649.

As of Storwize software V7.4, the drive firmware can be upgraded using the GUI. The user can choose to upgrade all drives or select an individual drive.

Download the latest Drive Upgrade package from the IBM Support site:

<http://www.ibm.com/support/entry/portal/support>

Upgrading the firmware on individual drives

To upgrade an individual drive, complete the following steps:

1. Navigate to **Pools** → **Internal Storage**, right-click the drive to be upgraded, and select **Upgrade** from the action menu, as shown in Figure 11-50.



Figure 11-50 Upgrade an individual drive

2. Select the Drive Upgrade package, which was previously downloaded from the IBM Support site, by clicking the folder icon. Click **Upgrade**, as shown in Figure 11-51.



Figure 11-51 Select Drive Upgrade package

- The drive firmware update takes about 2 - 3 minutes per drive. Figure 11-52 shows the completed task.

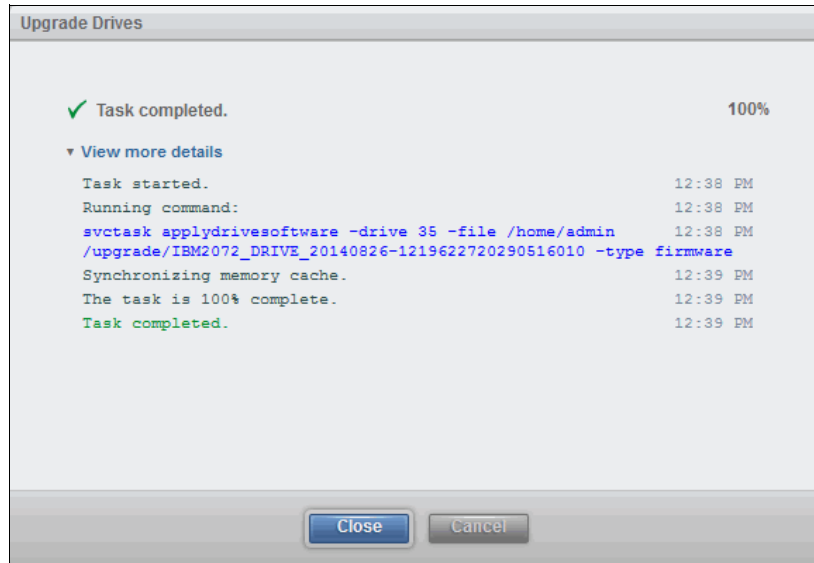


Figure 11-52 Drive firmware update complete

Figure 11-53 shows the **Pools** → **Internal Storage** pane, which displays the result of the drive firmware update (from BD2F to BD2H).

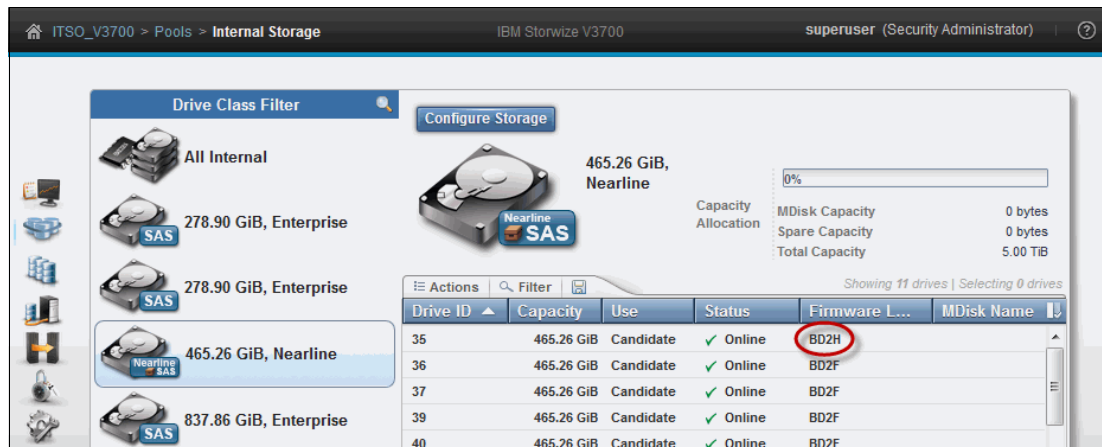


Figure 11-53 Result of drive firmware update

Upgrading the firmware on all drives

There are two ways to upgrade all of the drives in an IBM Storwize V3700 system in the Management GUI:

- ▶ From the **Monitoring** → **System** pane
- ▶ From the **Pools** → **Internal Storage** pane

To upgrade the firmware on all drives, complete the following steps:

1. Upgrade all drives in one of the following ways:
 - Figure 11-54 shows how to upgrade all drives using the Actions menu in the System pane. Select **Update** → **Drives**.

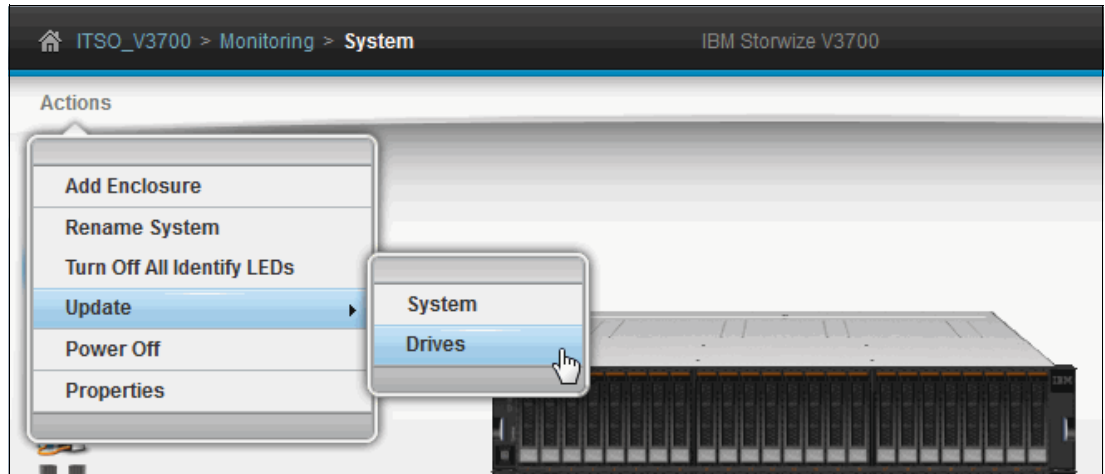


Figure 11-54 Upgrade all drives: System pane

- Figure 11-55 shows how to update all drives using the Action menu in the Internal Storage pane. Select **Upgrade All**.

Note: If any drives are selected, the Action menu displays actions for the selected drives, and the **Upgrade All** option does not display. If a drive is selected, clear it by holding down the Ctrl key and clicking the drive.

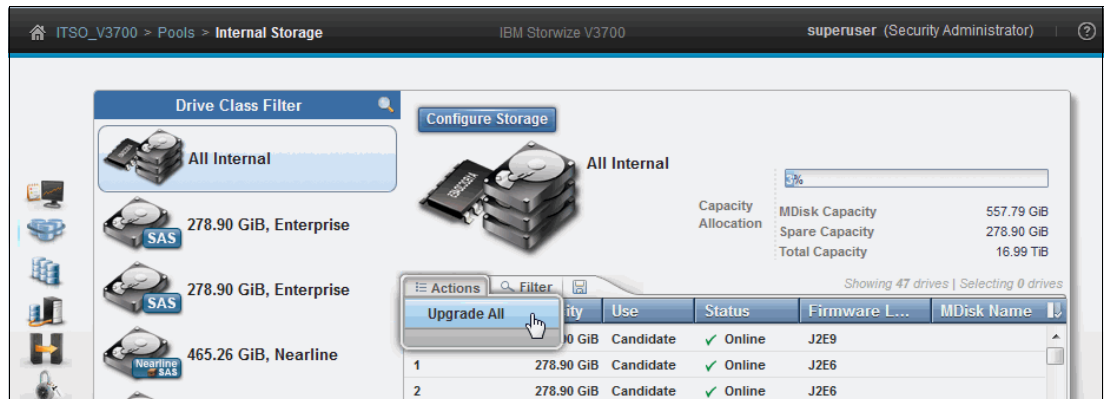


Figure 11-55 Upgrade all drives: Internal Storage pane

2. After initiating the drive upgrade process using either of the previous options, the pane in Figure 11-56 is displayed. Select the Drive Upgrade package, which was previously downloaded from the IBM Support site, by clicking the folder icon, and then click **Upgrade** to continue.

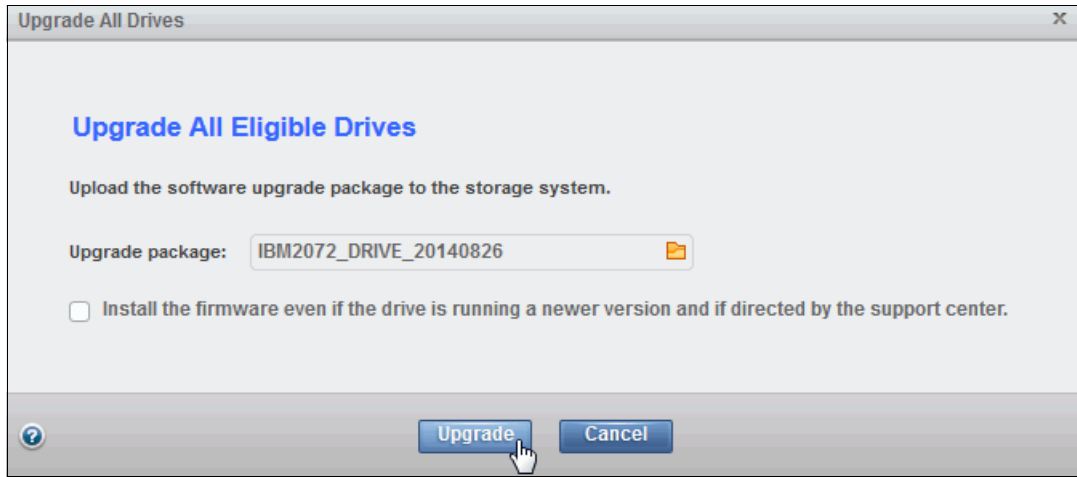


Figure 11-56 Upgrade all drives: select firmware file

All of the drives that require an upgrade are now upgraded.

11.5 Event log

Whenever a significant change in the status of IBM Storwize V3700 is detected, an event is submitted to the event log.

All events are classified as *alerts* or *messages*.

An alert is logged when the event requires action. Some alerts have an associated error code that defines the service action that is required. The service actions are automated through the fix procedures. If the alert does not have an error code, the alert represents an unexpected change in state. This situation must be investigated to see whether it is expected or represents a failure. Investigate an alert and resolve it when it is reported.

A message is logged when a change that is expected is reported; for instance, an IBM FlashCopy operation completes.

Open the event log pane from the GUI by clicking **Monitoring** → **Events** (Figure 11-57).

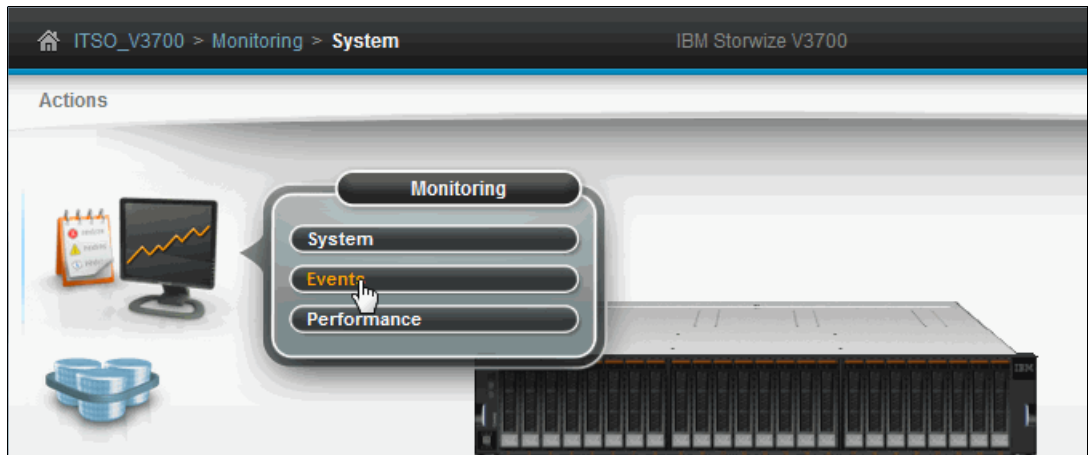


Figure 11-57 Opening the Events pane

Figure 11-58 shows the event log.

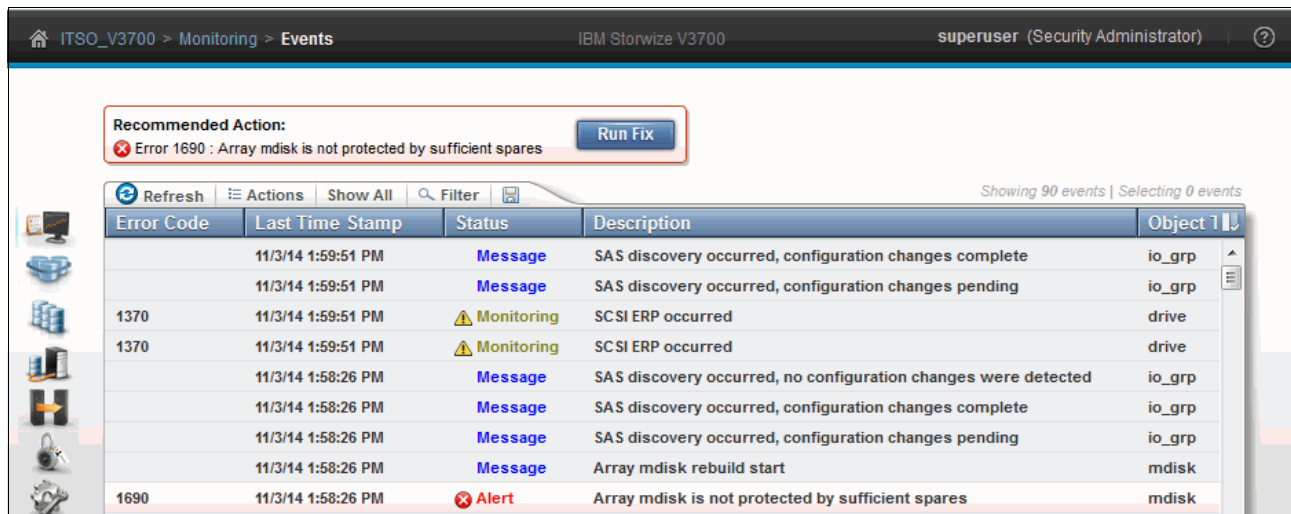


Figure 11-58 The event log view

11.5.1 Managing the event log

The event log features a size limit. After it is full, newer entries replace the older entries, which are not required.

To avoid a repeated event that fills the event log, some records in the event log refer to multiple occurrences of the same event. When event log entries are coalesced in this way, the time stamps of the first occurrence and the last occurrence of the problem are saved in the log entry. A count of the number of times that the error condition occurred is also saved in the log entry. Other data refers to the last occurrence of the event.

Event log pane columns

Right-clicking in any column header opens the option menu in which you can select columns that are shown or hidden. It's also possible to click the Column icon on the far right of the column headers to open the option menu.

Figure 11-59 shows all of the possible columns that can be displayed in the error log view.

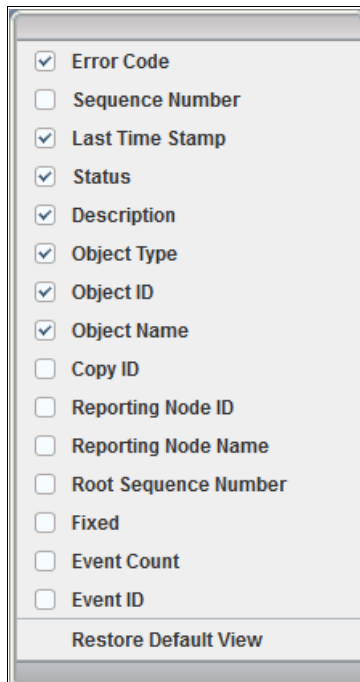


Figure 11-59 Possible event log columns

The following available fields are suggested at a minimum to assist you in diagnosing problems:

► Event ID

This number precisely identifies the reason why the event was logged.

► Error code

This number describes the service action that should be followed to resolve an error condition. Not all events have error codes that are associated with them. Many event IDs can have the same error code because the service action is the same for all of the events.

► Sequence number

A number that identifies the event.

► Event count

The number of events that are coalesced into this event log record.

► Fixed

When an alert is shown for an error condition, it indicates whether the reason for the event was resolved. In many cases, the system automatically marks the events that are fixed when appropriate. There are some events that must be manually marked as fixed. If the event is a message, this field indicates that you read and performed the action. The message must be marked as read.

► Last time

The time when the last instance of this error event was recorded in the log.

► Root sequence number

If set, this number is the sequence number of an event that represents an error that probably caused this event to be reported. Resolve the root event first.

Event log pane options

Figure 11-60 shows the main Event log pane options that should be used to handle system events.

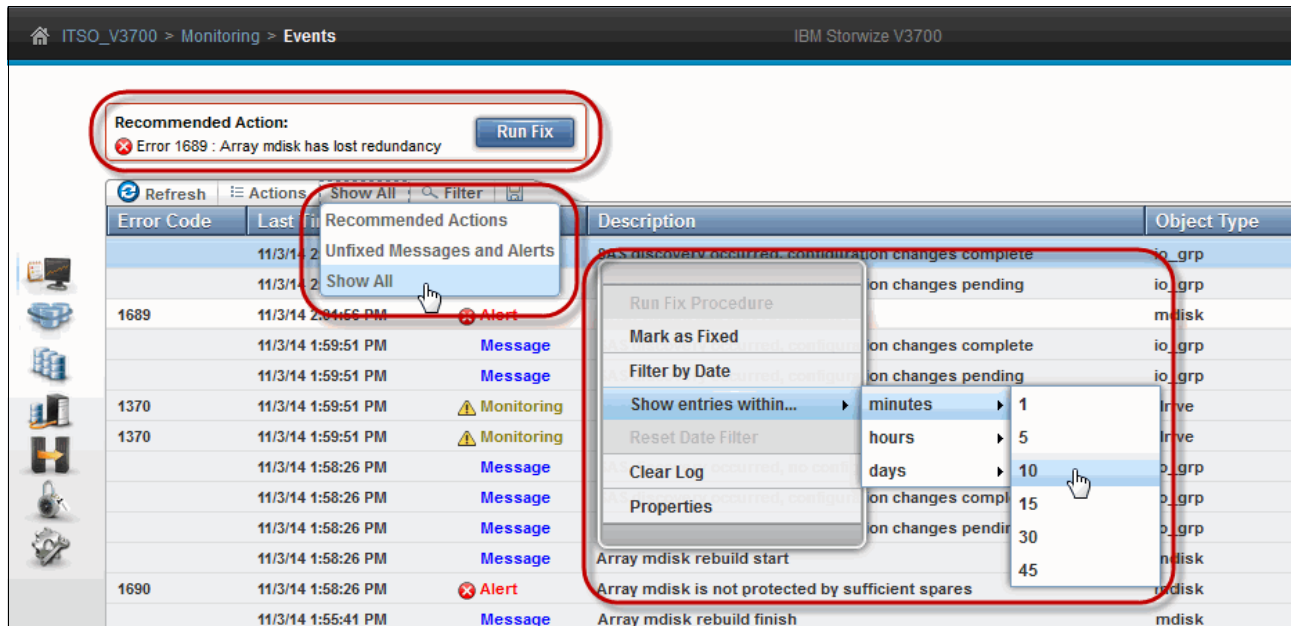


Figure 11-60 Events pane

Event log filter options

The following log filter options are available:

- ▶ Show all
 - This option lists all available events.
- ▶ Unfixed Messages and Alerts
 - This option lists unfixed events. This option is useful to find events that must be handled, but no actions are required or recommended.
- ▶ Recommended Actions (default)
 - Only events with recommended actions (Status Alert) are displayed.

Important: Check for this filter option if no event is listed. There might be events that are not associated to recommended actions.

Figure 11-61 shows an event log with no items found, which does not necessarily mean that the event log is clear. To check whether the log is clear, use the filter option **Show all**.

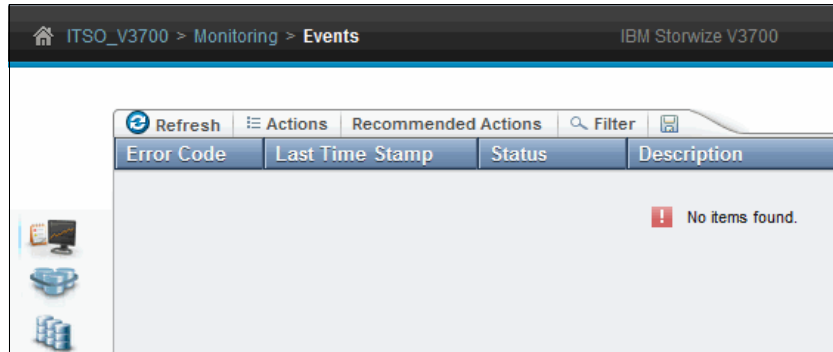


Figure 11-61 No items found in event log

Actions on a single event

Right-clicking a single event gives the following options that might be used for that specific event:

- ▶ **Mark as Fixed**

It is possible to start the Fix Procedure on this specific event, even if it is not the recommended next action.

Some events, such as messages, must be set to Mark as Fixed.

- ▶ **Show entries within... minutes/hours/days**

This option is to limit the error log list to a specific date or a time slot. The following selectable values are available:

- Minutes: 1, 5, 10, 15, 30, and 45
- Hours: 1, 2, 5, and 12
- Days: 1, 4, 7, 15, and 30

- ▶ **Clear Log**

This option clears the complete error log, even if only one event was selected.

Attention: This action cannot be undone, and might prevent the system from being analyzed when severe problems occur.

- ▶ **Properties**

This option provides more sense data for the selected event that is shown in the list.

Recommended Actions

A fix procedure starts a wizard known as a Directed Maintenance Procedure (DMP) that helps to troubleshoot and correct the cause of an error. Some DMPs reconfigure the system based on your responses, ensure that actions are carried out in the correct sequence, and prevent or mitigate loss of data. For this reason, you must always run the fix procedure to fix an error, even if the fix might seem obvious.

To run the fix procedure for the error with the highest priority, go to the Recommended Action pane at the top of the Event page and click **Run This Fix Procedure**. When you fix higher priority events first, the system can often automatically mark lower priority events as fixed.

11.5.2 Alert handling and recommended actions

All events in Alert status require attention. Alerts are listed in priority order, and should be fixed sequentially by using the available fix procedures.

Example: Fibre Channel small form-factor pluggable (SFP) fault

For this example, we emulated an FC SFP failure and to do so, disabled an FC port on the storage area network (SAN) switch.

The following example shows how faults are represented in the error log, how information about the fault can be gathered, and how the recommended action (DMP) can be used to fix the error:

1. Detect an alert.

The Health Status indicator shows a red alert (for more information, see Chapter 3, “Graphical user interface overview” on page 71). The Status Alerts indicator (to the right of the Health status bar) is also visible, and shows two alerts. Click the alert to retrieve the specific information, as shown in Figure 11-62.

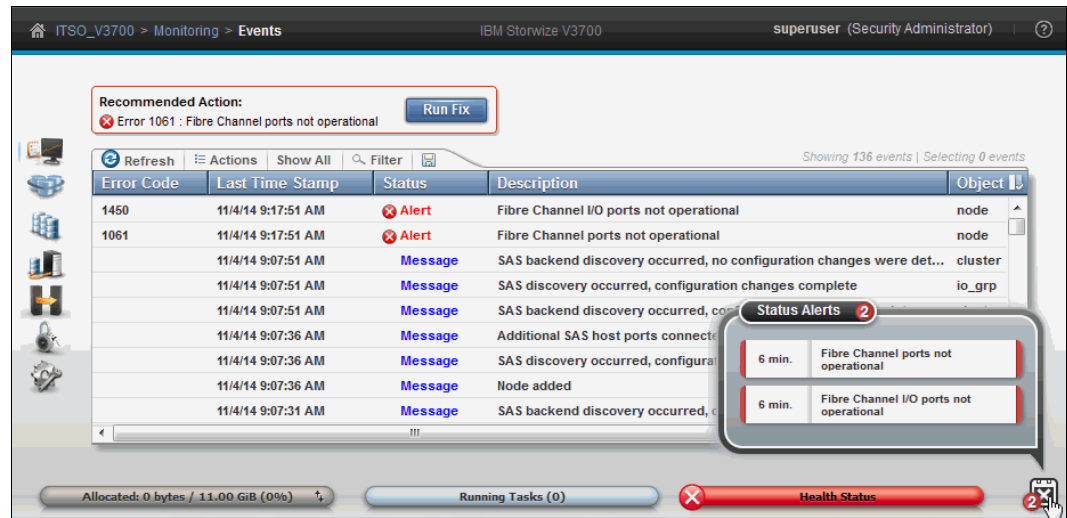


Figure 11-62 Health Status shows degraded system

2. Review the event log for more information.

3. Gather additional information: Alert properties.

More details about the event can be found in the Properties option, as shown in Figure 11-63. This information might be of interest for problem fixing, or for root cause analysis.

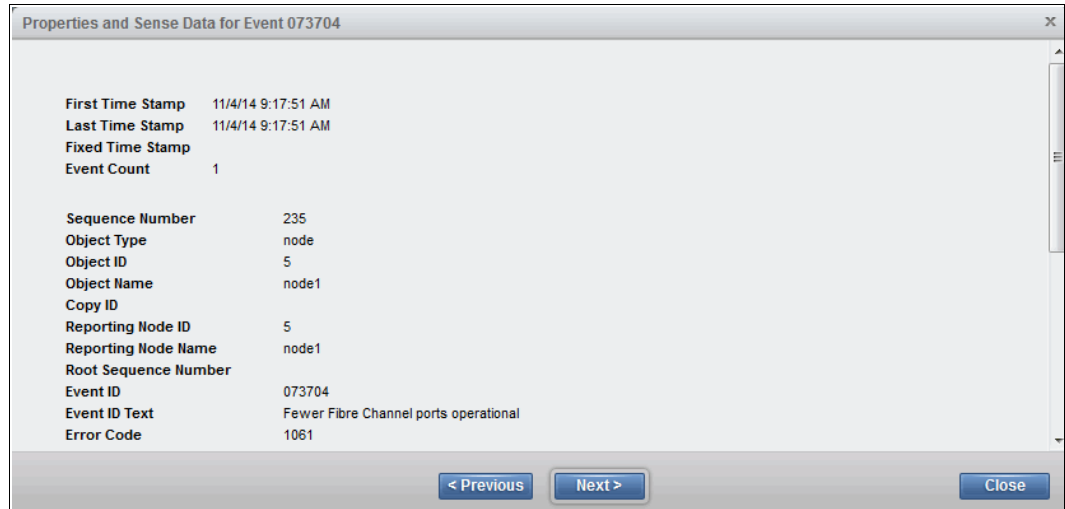


Figure 11-63 Alert properties

4. Run the recommended action (DMP):

- a. Use of the DMP is highly recommended to fix any alerts. It is possible to miss background running tasks when the DMP is bypassed. Not all alerts have DMPs available.
- b. To start the DMP, right-click the alert record or click **Run this fix procedure** at the top of the pane.

- The steps and panes of the DMP are specific to the error. The following figures represent the recommended action (DMP) for the FC SFP failure example, which shows as a **Fibre Channel ports not operational** event.

Figure 11-64 shows step 1 of the DMP for the **Fibre Channel ports not operational** event. In this example, we say that the change is not intentional.

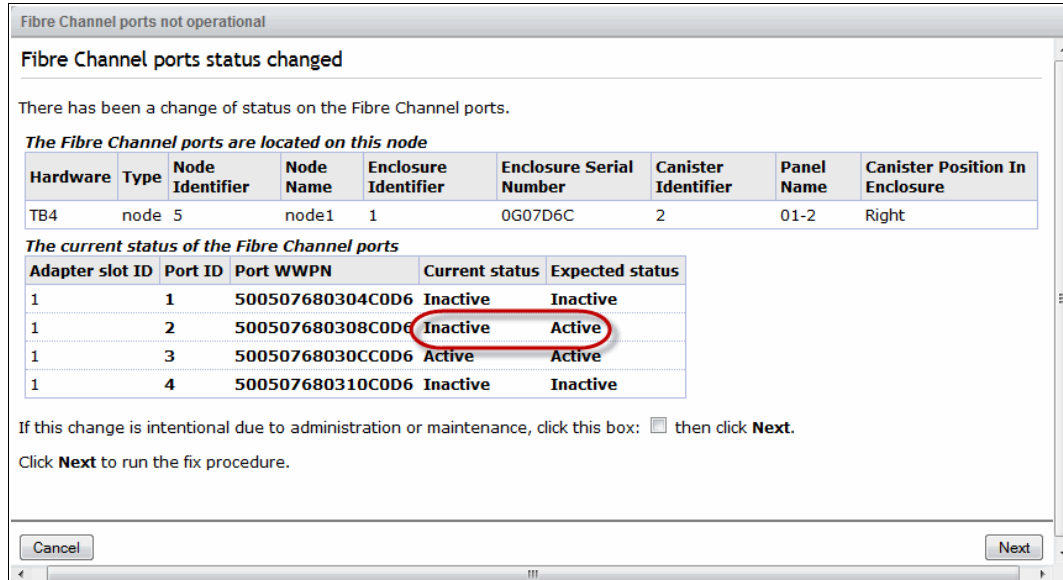


Figure 11-64 Recommended action DMP for Fibre Channel port not operational: Step 1

Figure 11-65 shows step 2 of the DMP for the **Fibre Channel ports not operational** event. The DMP recommends checking the FC cable connected to the inactive port. You will confirm that the FC cables are not defective and move on to the next service action.

Fibre Channel ports not operational

Check the Fibre Channel cabling

For the ports that currently have inactive status and are not expected to be, check the Fibre Channel [cable](#).

- Ensure the correct type of cable is being used.
- If the cable appears damaged, replace it.
- If there are any sharp bends in the cable, re-route or replace it.
- Reseat the cable connector by unplugging the cable for two seconds, and then reconnecting it.

After performing this service action, click **Next** to check the port status.

Select one of these options before clicking **Next**:

Fibre Channel status is incorrect, try next service action
 Fibre Channel status is correct, mark as fixed


The Fibre Channel ports are located on this node

Hardware	Type	Node Identifier	Node Name	Enclosure Identifier	Enclosure Serial Number	Canister Identifier	Panel Name	Canister Position In Enclosure
TB4	node	5	node1	1	0G07D6C	2	01-2	Right

The current status of the Fibre Channel ports

Adapter slot ID	Port ID	Port WWPN	Current status	Expected status
1	1	500507680304C0D6	Inactive	Inactive
1	2	500507680308C0D6	Inactive	Active
1	3	50050768030CC0D6	Active	Active
1	4	500507680310C0D6	Inactive	Inactive

[Refresh port status table](#)



Fibre-channel ports

Cancel Back Next

Figure 11-65 Recommended action DMP for Fibre Channel port not operational: Step 2

Figure 11-66 shows step 3 of the DMP for the **Fibre Channel ports not operational** event, and recommends replacing the SFP connections at each end of the FC cable. After re-enabling the port on the SAN switch to emulate the replacement of the SFP, the link to **Refresh the port status table** was clicked. The port status table now shows that the port is active, and the option to mark the service action as fixed is clicked.

Fibre Channel ports not operational

SFP replacement

For the ports that have a current status of inactive, and are not expected to be, replace the [SFP connectors](#) at both ends of the Fibre Channel cables. After performing the service action, click **Next** to continue.

Select one of these options before clicking **Next**:

Fibre Channel status is incorrect, try next service action
 Fibre Channel status is correct, mark as fixed

The Fibre Channel ports are located on this node

Hardware	Type	Node Identifier	Node Name	Enclosure Identifier	Enclosure Serial Number	Canister Identifier	Panel Name	Canister Position In Enclosure
TB4	node	5	node1	1	0G07D6C	2	01-2	Right

The current status of the Fibre Channel ports

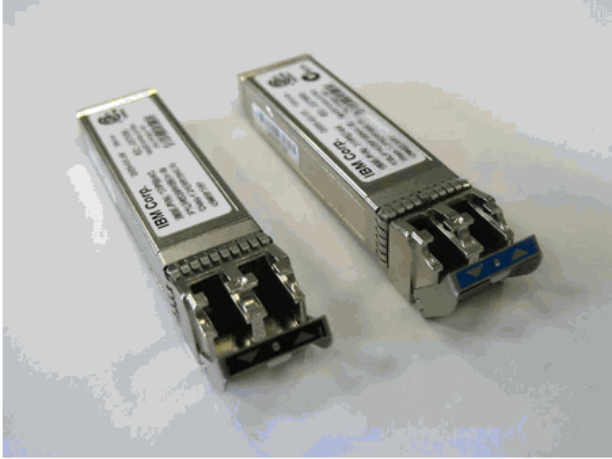
Adapter slot ID	Port ID	Port WWPN	Current status	Expected status
1	1	500507680304C0D6	Inactive	Inactive
1	2	500507680308C0D6	Active	Active
1	3	50050768030CC0D6	Active	Active
1	4	500507680310C0D6	Inactive	Inactive

[Refresh port status table](#)

SFP replacement

Node SFP replacement details

Port number	Part number	Wave length	Connector type	Supported speeds
1	85Y6278	850	LC	2,4,8 Gbps
2	85Y6278	850	LC	2,4,8 Gbps
4	85Y6278	850	LC	2,4,8 Gbps



Short and long wave SFP connectors

Figure 11-66 Recommended action DMP for Fibre Channel port not operational: Step 3

Figure 11-67 shows step 4 of the DMP for the **Fibre Channel ports not operational** event and confirms that the Fibre Channel ports are correctly configured.

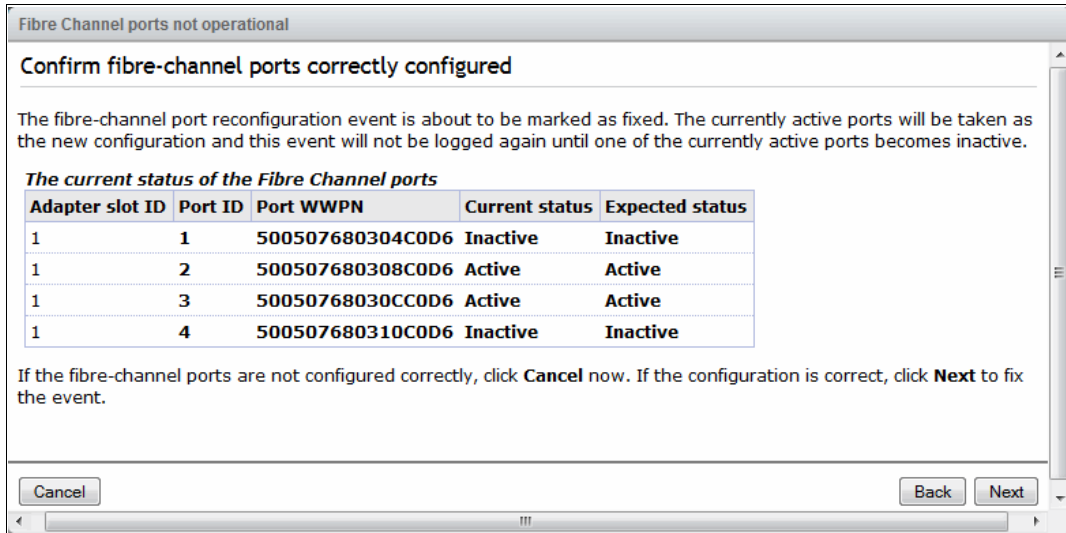


Figure 11-67 Recommended action DMP for Fibre Channel port not operational: Step 4

Figure 11-68 shows step 5 of the DMP for the **Fibre Channel ports not operational** event.

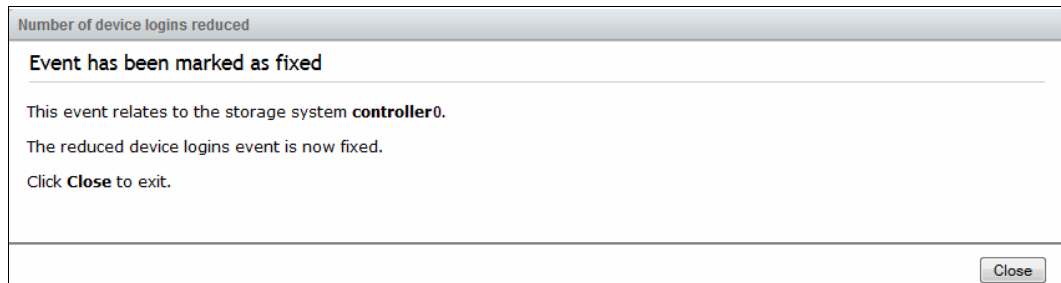


Figure 11-68 Recommended action DMP for Fibre Channel port not operational: Step 5

When all of the steps of the DMP are processed successfully, the recommended action, is complete and the problem should be fixed. Figure 11-69 on page 625 shows the red color of the event status changed to green. The System Health status is green and the Recommended Action box has disappeared, implying that there are no further actions that must be addressed.

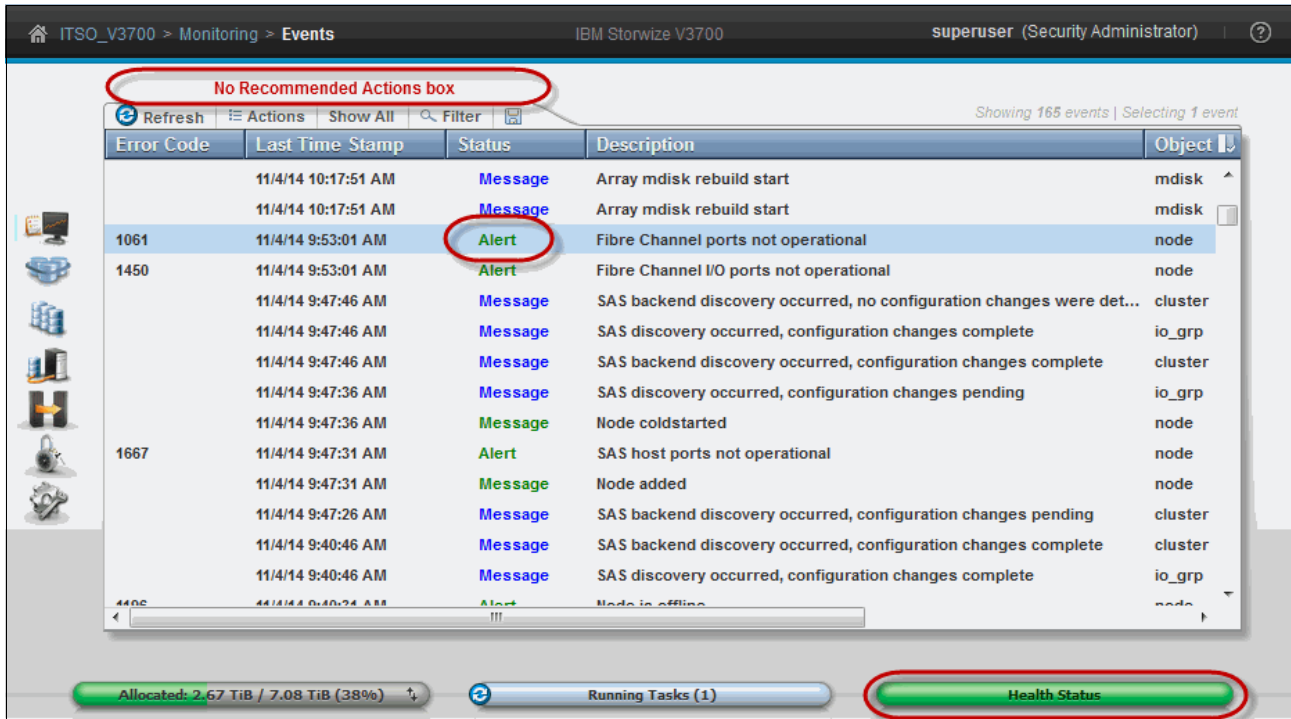


Figure 11-69 Recommended action: Completed

Handling multiple alerts

If there are multiple alerts logged, the IBM Storwize V3700 recommends an action to fix the problem (or problems).

Figure 11-70 shows the event log that displays multiple alerts.

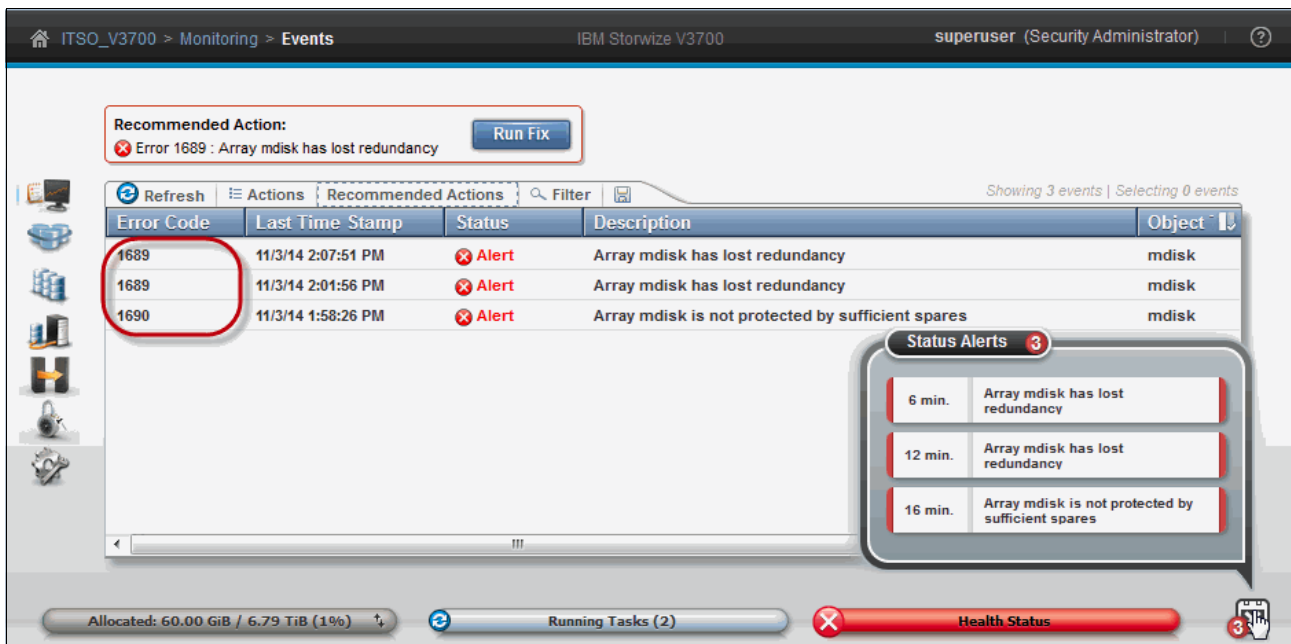


Figure 11-70 Multiple alert events displayed in the event log

The Next Recommended Action function orders the alerts by severity, and displays the events with the highest severity first. If multiple events have the same severity, they are ordered by date, and the oldest event is displayed first.

The following order of severity starts with the most severe condition:

- ▶ Unfixed alerts (sorted by error code, where the lowest error code has the highest severity)
- ▶ Unfixed messages
- ▶ Monitoring events (sorted by error code; the lowest error code has the highest severity)
- ▶ Expired events
- ▶ Fixed alerts and messages

Less severe faults are often fixed with the resolution of the most severe fault.

11.6 Collecting support information

If you have an issue with an IBM Storwize V3700 and call the IBM Support Center, you might be asked to provide support data, as described in the next section.

11.6.1 Support information from the GUI

We show how to obtain support information using the GUI.

Download Support Package wizard

The Download Support Package wizard provides a selection of various package types. IBM Support provides direction on package type selection as required. To download a Support package, complete the following steps:

1. Browse to **Settings** → **Support**, as shown in Figure 11-71.




Figure 11-71 Accessing Support files

2. Click **Download Support Package**, as shown in Figure 11-72.

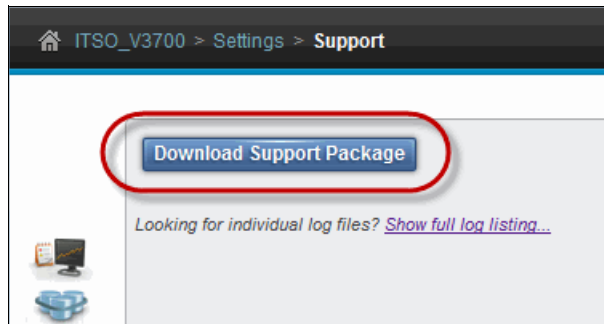


Figure 11-72 Download Support Package

Figure 11-73 shows the Download Support Package pane, from which you can select one of four different types of the support package.

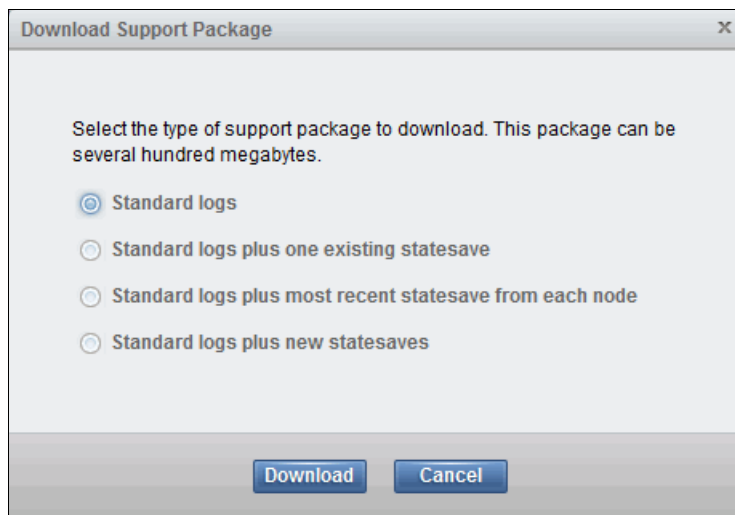


Figure 11-73 Support Package selection

3. Select the type, depending on the event that is being investigated, and click **Download**. IBM Support will notify the user of which package is required.

The following components are included in the support package:

- Standard logs. Contains the most recent logs that were collected from the system. These logs are most commonly used by Support to diagnose and solve problems.
- Standard logs plus one existing statesave. Contains the standard logs from the system and the most recent statesave from any of the nodes in the system. Statesaves are also known as *memory dumps* or *live memory dumps*.
- Standard logs plus most recent statesave from each node. This option is used most often by the support team for problem analysis. They contain the standard logs from system and the most recent statesave from each node in the system.
- Standard logs plus new statesave. This option might be requested by the Support team for problem determination. It generates a new statesave (livedump) for all of the nodes and packages them with the most recent logs.

4. Save the resulting file in a directory for later use, or to upload to IBM Support.

Show full log listing

The Support pane also provides access to the files that are on the node canisters, as shown in Figure 11-74. To show the full log listing, complete the following steps:

1. Click **Show full log listing** to access the node canister files.
2. To save a file to the user's workstation, select a file, right-click the file, and select **Download**.
3. To change to the file listing to show the files on a partner node canister, select the node canister from the Node menu.

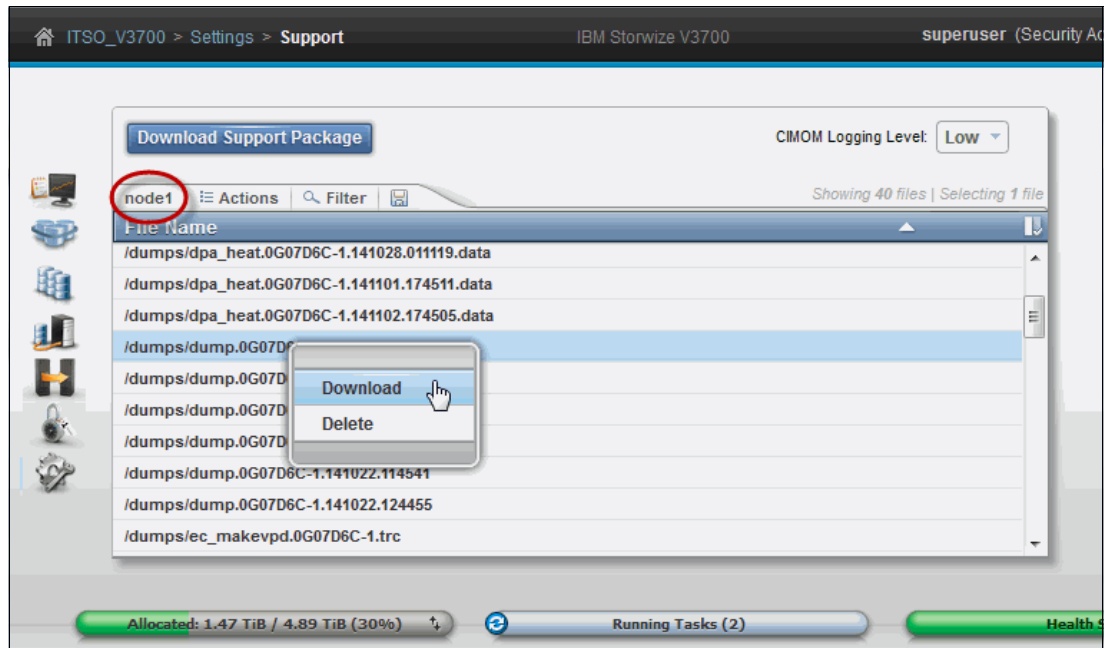


Figure 11-74 Full log listing

11.6.2 Support information from the Service Assistant

The IBM Storwize V3700 management GUI collects information from all of the components in the system. The Service Assistant Tool collects information from all node canisters. The snap file is the information that is collected and packaged in a single file.

To access the gathered information, complete the following steps:

1. If the package is collected by using the SA tool, ensure that the node from which the logs are collected is the current node, as shown in Figure 11-75.

The screenshot shows the IBM Storwize V3700 Service Assistant Tool interface. The top bar indicates the user is connected to '01 | 1 | node1'. The left sidebar contains a navigation menu with 'Collect Logs' highlighted. The main content area shows the 'Home' page with an 'Attention' warning and a 'Change Node' table. In the table, 'node1' is selected with a radio button.

Node Name	Node Status	Error	Panel	System	Site	Relationship
node1	Active		01-1	ITSO_V3700		Local
node2	Active		01-2	ITSO_V3700		Partner

Figure 11-75 Collect logs with the Service Assistance Tool

2. Support information can be downloaded **With latest statesave** or **Without latest statesave**, as shown in Figure 11-76.

The screenshot shows the 'Collect Logs' page in the IBM Storwize V3700 Service Assistant Tool. The left sidebar has 'Collect Logs' selected. The main content area includes a 'Create a support package' section with two radio button options: 'With latest statesave' (selected) and 'Without latest statesave'. Below this is a list of 'Available Files'.

Available Files
<input type="radio"/> 0G07D6C-1.trc
<input type="radio"/> 0G07D6C-1.trc.old
<input type="radio"/> boot.0G07D6C-1.trc
<input type="radio"/> dpa_heat.0G07D6C-1.141027.040122.data
<input type="radio"/> dpa_heat.0G07D6C-1.141028.011119.data
<input type="radio"/> dpa_heat.0G07D6C-1.141101.174511.data
<input type="radio"/> dpa_heat.0G07D6C-1.141102.174505.data

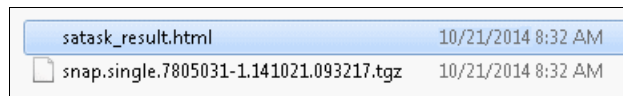
Figure 11-76 Download support file using the Service Assistant Tool

11.6.3 Support information to USB stick

Whenever the management GUI, SA Tool, or a remote connection is unavailable, snaps can be collected from each node using a USB stick.

Complete the following steps to collect snaps by using a USB stick:

1. Create a text file that includes the following command:
`satask snap -dump`
2. Save the file as `satask.txt` in the root directory of the USB stick.
3. Insert the USB stick in the USB port of the node from which the support data should be collected.
4. Wait until no write activity is recognized (this process can take 10 minutes or more).
5. Remove the USB stick and check the results, as shown in Figure 11-77.

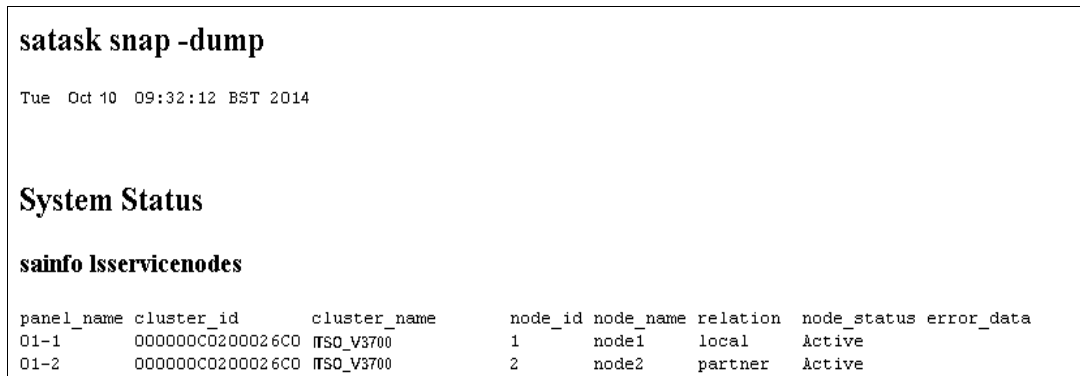


satask_result.html	10/21/2014 8:32 AM
snap.single.7805031-1.141021.093217.tgz	10/21/2014 8:32 AM

Figure 11-77 Single snap result files on USB stick

satask_result file

The `satask_result.html` file is the general response to the command that is issued using the USB stick. If the command did not run successfully, it is noted in this file. Otherwise, any general system information is stored here, as shown in Figure 11-78.



```
satask snap -dump

Tue Oct 10 09:32:12 BST 2014

System Status

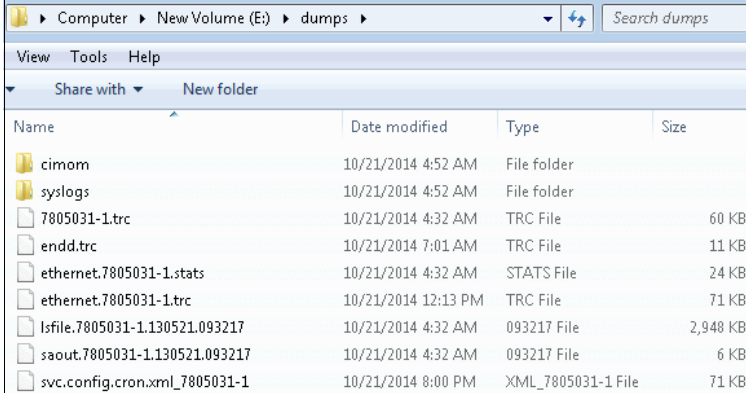
sainfo lsservicenodes

panel_name cluster_id cluster_name node_id node_name relation node_status error_data
01-1 000000C0200026C0 ITSO_V3700 1 node1 local Active
01-2 000000C0200026C0 ITSO_V3700 2 node2 partner Active
```

Figure 11-78 The `satask_result.html` file on a USB stick (header only)

Snap memory dump on USB

A complete statesave of the node in which the USB was inserted is stored in a .tgz file. The name of the file includes the node name and the time stamp. The content of the .tgz file is shown in Figure 11-79.



Name	Date modified	Type	Size
cimom	10/21/2014 4:52 AM	File folder	
syslogs	10/21/2014 4:52 AM	File folder	
7805031-1.trc	10/21/2014 4:32 AM	TRC File	60 KB
endd.trc	10/21/2014 7:01 AM	TRC File	11 KB
ethernet.7805031-1.stats	10/21/2014 4:32 AM	STATS File	24 KB
ethernet.7805031-1.trc	10/21/2014 12:13 PM	TRC File	71 KB
lsfile.7805031-1.130521.093217	10/21/2014 4:32 AM	093217 File	2,948 KB
saout.7805031-1.130521.093217	10/21/2014 4:32 AM	093217 File	6 KB
svc.config.cron.xml_7805031-1	10/21/2014 8:00 PM	XML_7805031-1 File	71 KB

Figure 11-79 Single snap memory dump on USB stick

11.7 Powering on and powering off the IBM Storwize V3700

In the following sections, we describe the process to power on and power off the IBM Storwize V3700 system by using the GUI and the CLI.

11.7.1 Powering off the system

In this section, we show how to power off (shutdown) the IBM Storwize V3700 system by using the GUI and CLI.

Attention: You should *never* power off an IBM Storwize V3700 by powering off the PSUs, removing both PSUs, or removing both power cables from a running system.

Powering down by using the GUI

You can power off a node canister or the entire cluster. When you shut down only one node canister, all of the running tasks remain active, because the remaining node takes over.

Note: If a canister or the enclosure is powered down, a local visit is required to either reseat the canister or power cycle the enclosure.

Complete the following steps to perform a power off:

1. Power off a canister:
 - a. Browse to **Monitoring** → **System**.
 - b. Rotate the enclosure to the rear view.
 - c. Right-click the required canister.
 - d. Select Power Off, as shown in Figure 11-80.

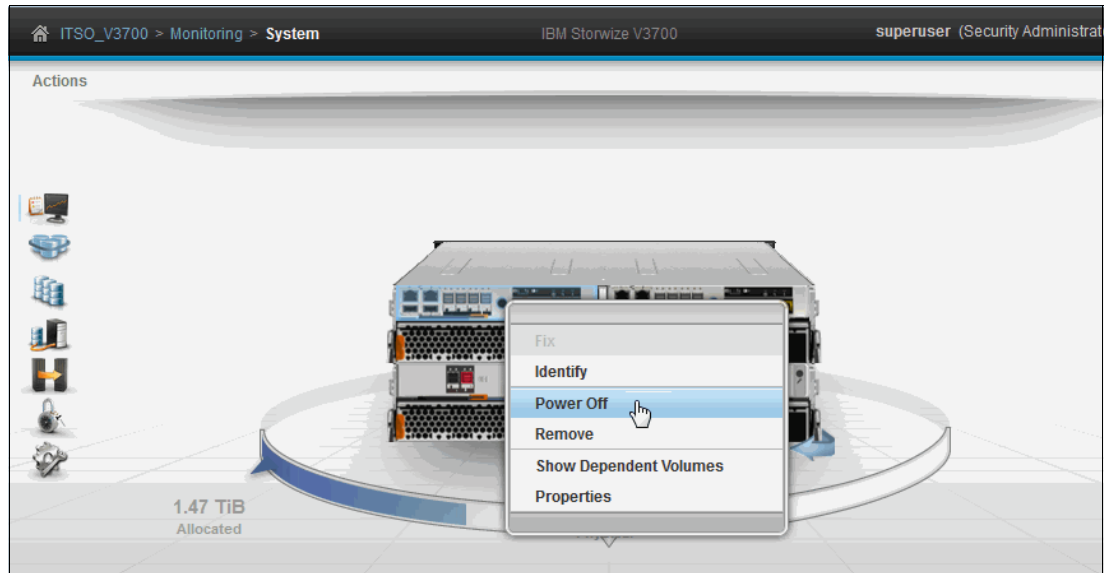


Figure 11-80 Powering off a canister

2. Power down the system:
 - a. Browse to **Monitoring** → **System**.
 - b. Click the Actions menu.
 - c. Select Power Off, as shown in Figure 11-81.

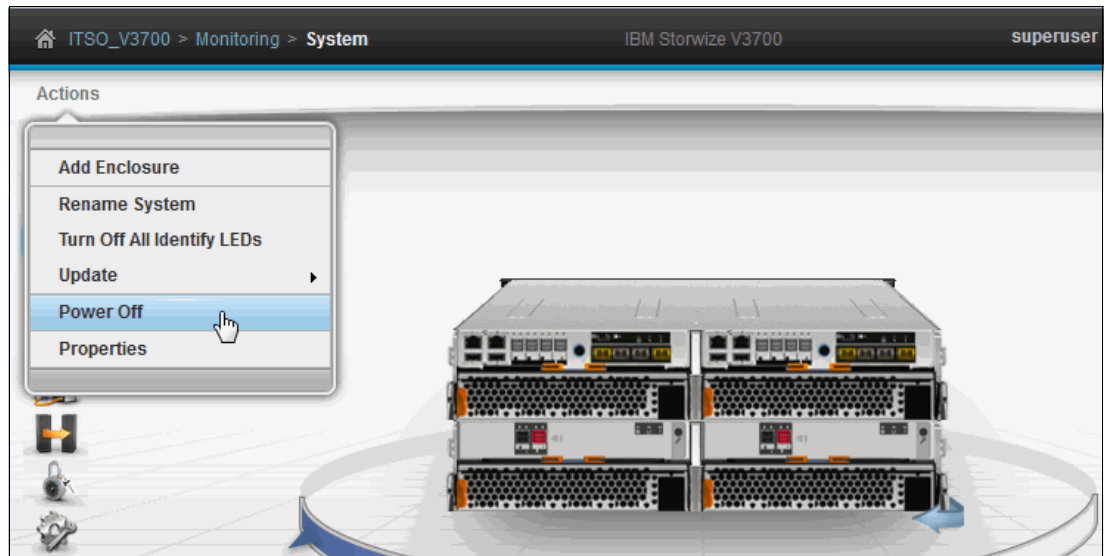


Figure 11-81 Powering off the system

3. A Power Off confirmation window opens, prompting for confirmation to shut down the cluster. Ensure that all FlashCopy, Metro Mirror, Global Mirror, data migration operations, and forced deletions are stopped or allowed to complete before continuing.

4. Enter the provided confirmation code and click **Power Off** to begin the shutdown process, as shown in Figure 11-82.

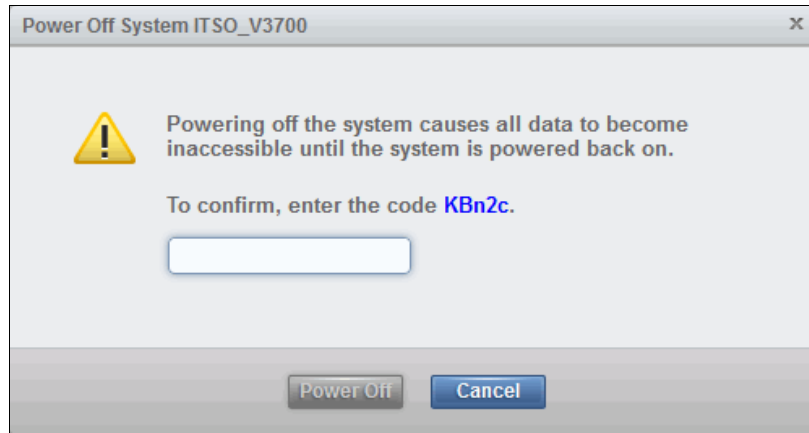


Figure 11-82 Power Off system confirmation

5. Wait for the power LED on both node canisters in the control enclosure to flash at 1 Hz, which indicates that the shutdown operation completed (1 Hz is half as fast as the drive indicator LED).

Tip: When you shut down an IBM Storwize V3700, it does not automatically restart. You must manually restart the system by removing and reapplying power.

Shutting down by using the CLI

The CLI can also be used to shut down an IBM Storwize V3700. The CLI can be accessed on Windows using the PuTTY utility.

Warning: If you are shutting down the entire system, you lose access to all volumes that are provided by this system. Shutting down the system also shuts down all IBM Storwize V3700 nodes. All data is flushed to disk before the power is removed.

Run the **stopssystem** command to shut down a clustered system, as shown in Example 11-4.

Example 11-4 Shut down

```
stopssystem
```

```
Are you sure that you want to continue with the shut down?
```

```
# Type y to shut down the entire clustered system.
```

11.7.2 Powering on

Important: This process assumes that all power is removed from the enclosure. If the control enclosure is shut down but the power is not removed, the power LED on all node canisters flashes at 1 Hz. In this case, remove the power cords from both power supplies and then reinsert them.

Complete the following steps to power on the system:

1. Ensure that any network switches that are connected to the system are powered on.
2. Power on any expansion enclosures by connecting the power cord to both power supplies in the rear of the enclosure, and turning on the power circuit.
3. Power on the control enclosure by connecting the power cords to both power supplies in the rear of the enclosure and turning on the power circuits.

The system starts. The system starts successfully when all node canisters in the control enclosure have their status LED permanently on, which should take no longer than 10 minutes.

4. Start the host applications.

11.8 Tivoli Storage Productivity Center

The IBM Tivoli Storage Productivity Center provides a set of policy-driven automated tools for managing storage capacity, availability, events, performance, and assets in your enterprise environment. Tivoli Storage Productivity Center provides storage management from the host and application to the target storage device. It also provides disk and tape subsystem configuration and management, Performance Management, SAN fabric management and configuration, and usage reporting and monitoring.

In this section, we describe how to use Tivoli Storage Productivity Center to get usage reporting and to monitor performance data. Tivoli Storage Productivity Center can help you to identify, evaluate, control, and predict your enterprise storage management assets. Because it is policy-based, it can detect potential problems and automatically make adjustments that are based on the policies and actions that you define.

For example, it can notify you when your system is running out of disk space, or warn you of an impending storage hardware failure. By alerting you to these and other issues that are related to your stored data, you can prevent unnecessary system and application downtime.

11.8.1 Tivoli Storage Productivity Center benefits

Tivoli Storage Productivity Center includes the following benefits:

- ▶ Simplifies the management of storage infrastructures
- ▶ Manages, configures, and provisions SAN-attached storage
- ▶ Monitors and tracks performance of SAN-attached devices
- ▶ Monitors, manages, and controls (through zones) SAN fabric components
- ▶ Manages the capacity usage and availability of the file systems and databases
- ▶ Offers performance monitoring and reporting
- ▶ Provides report viewing using a web-based GUI

The examples provided in this section are not based on the current version of Tivoli Storage Productivity Center. For detailed guidance on configuring the most current version of Tivoli Storage Productivity Center, see the *IBM Tivoli Storage Productivity Center V5.1 Technical Guide*, SG24-8053, which is another IBM Redbooks publication:

<http://www.redbooks.ibm.com/abstracts/sg248053.html?Open>

11.8.2 Adding IBM Storwize V3700 in Tivoli Storage Productivity Center

After the Tivoli Storage Productivity Center is installed, it is ready to connect to the IBM Storwize V3700 system.

Complete the following steps to connect Tivoli Storage Productivity Center to the IBM Storwize V3700 system:

1. Open your browser and use the following link to start Tivoli Storage Productivity Center, as shown in Figure 11-83:

http://<TPC_system_Hostname>:9550/ITSRM/app/en_US/index.html

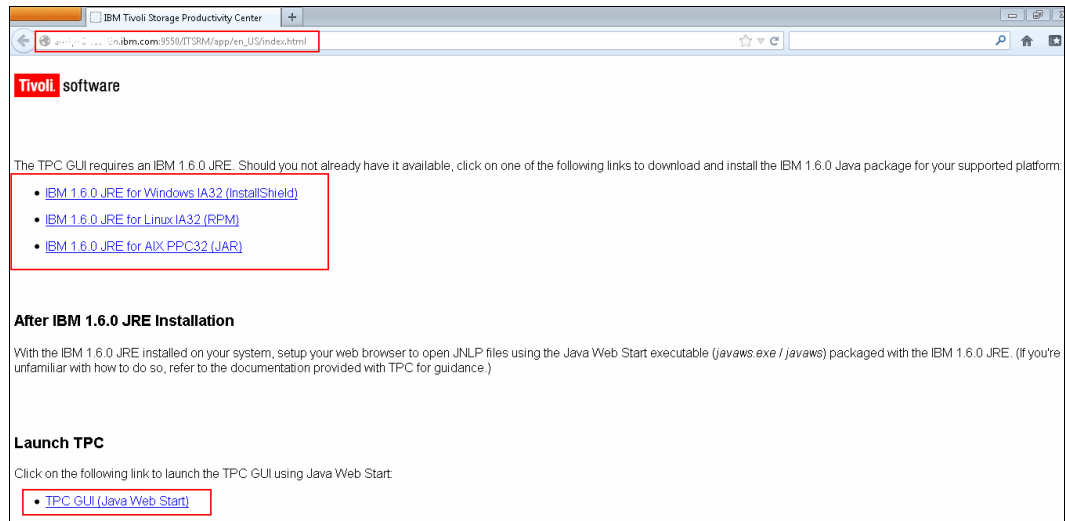


Figure 11-83 Tivoli Storage Productivity Center front page

2. You also can find a link on the web page to download IBM Java, if required. To start Tivoli Storage Productivity Center console, click **TPC GUI (Java Web Start)**.

Tivoli Storage Productivity Center starts an application download, as shown in Figure 11-84. If this is the first time you logged in, it takes time to install the required Java packages to the local system.

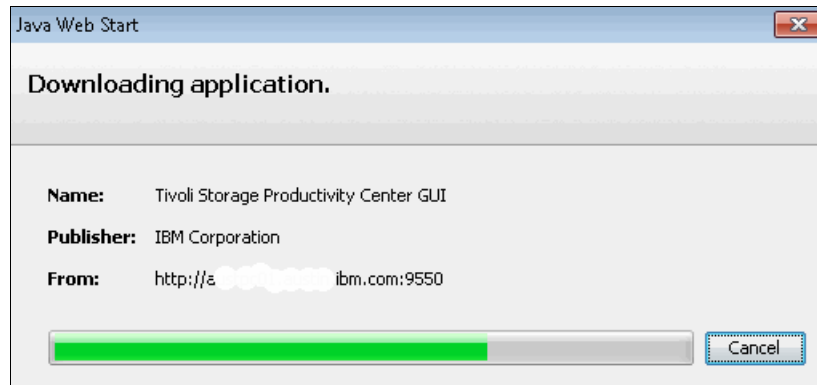


Figure 11-84 Downloading Tivoli Storage Productivity Center application

3. Use your login credentials to access Tivoli Storage Productivity Center, as shown in Figure 11-85.

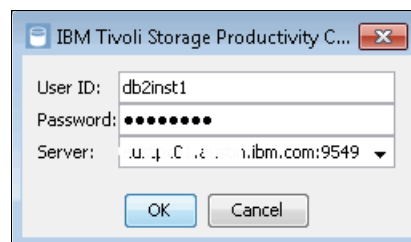


Figure 11-85 Tivoli Storage Productivity Center login access

4. After successfully logging in, you are ready to add storage devices to Tivoli Storage Productivity Center, as shown in Figure 11-86.

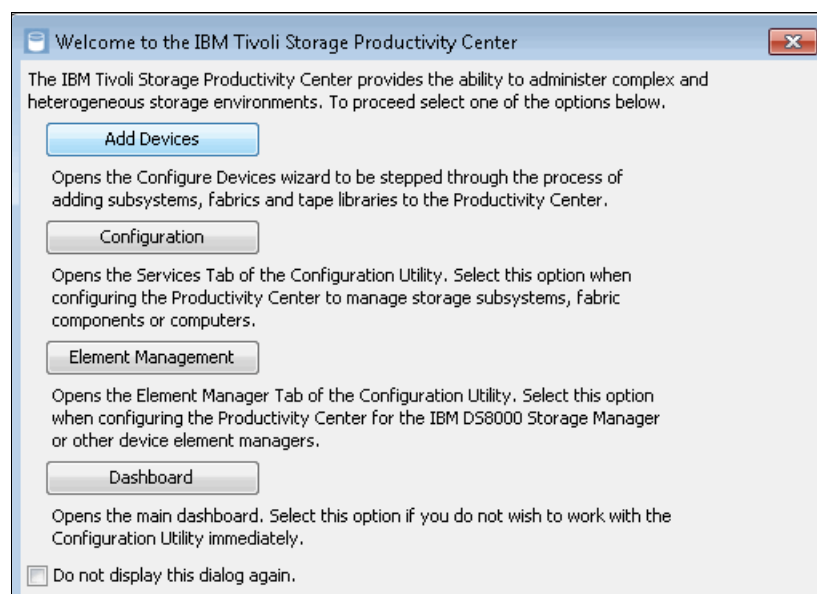


Figure 11-86 Add Devices window

5. Enter the details about your IBM Storwize V3700 into Tivoli Storage Productivity Center, as shown in Figure 11-87.




Figure 11-87 Configure device in Tivoli Storage Productivity Center

6. Continue to follow the wizard after you complete all of the required fields. After the wizard is completed, Tivoli Storage Productivity Center collects information from IBM Storwize V3700. A summary of details is shown at the end of discovery process.

11.9 Using Tivoli Storage Productivity Center to administer and generate reports for an IBM Storwize V3700

This section shows examples of how to use Tivoli Storage Productivity Center to administer, configure, and generate reports for an IBM Storwize V3700 system. A detailed description about Tivoli Storage Productivity Center reporting is beyond the intended scope of this book.

11.9.1 Basic configuration and administration

By using Tivoli Storage Productivity Center, you can administer, monitor, and configure your IBM Storwize V3700 system. However, not all of the options that are normally associated with the IBM Storwize V3700 GUI or CLI are available.

To configure Tivoli Storage Productivity Center, complete the following steps:

1. After successfully adding your IBM Storwize V3700 system, click **Disk Manager** → **Storage Subsystems** to view your configured devices, as shown in Figure 11-88.

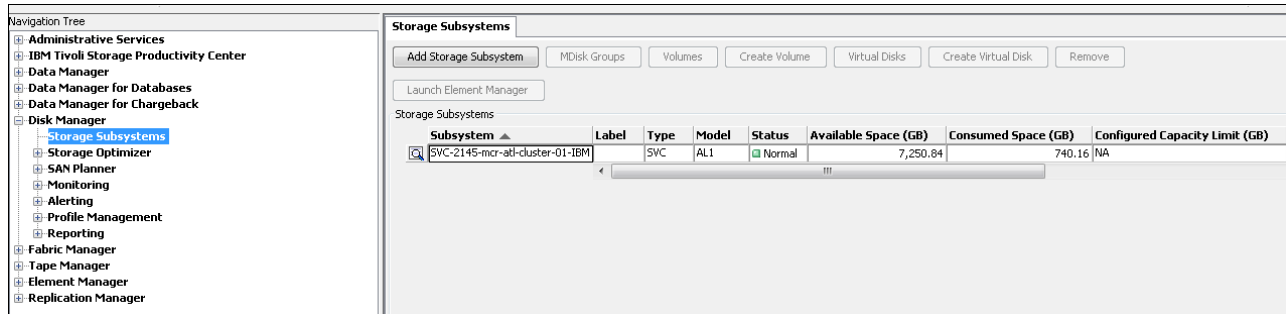


Figure 11-88 Storage Subsystem view

2. When you highlight the IBM Storwize V3700 system, action buttons become available that you can use to view the device configuration or create virtual disks, as shown in Figure 11-89.




Figure 11-89 Action buttons

- a. The MDisk Groups option provides you with a detailed list of the configured MDisk groups, including pool space, available space, configured space, and Easy Tier configuration.
- b. The Virtual Disks option lists all of the configured disks, with the added option to filter them by MDisk Group. The list includes several attributes, such as capacity, volume type, and type.

Important: Tivoli Storage Productivity Center and IBM SAN Volume Controller use the following terms:

- ▶ *MDisk Group* is the equivalent of a Storage Pool on a Storwize device.
- ▶ *Virtual Disk* is the equivalent of a Volume on a Storwize device.

3. If you click **Create Virtual Disk**, the Create Virtual Disk wizard window opens, as shown in Figure 11-90. Use this window to create volumes and specify several options, such as size, name, thin provisioning, and add MDisks to an MDisk group.

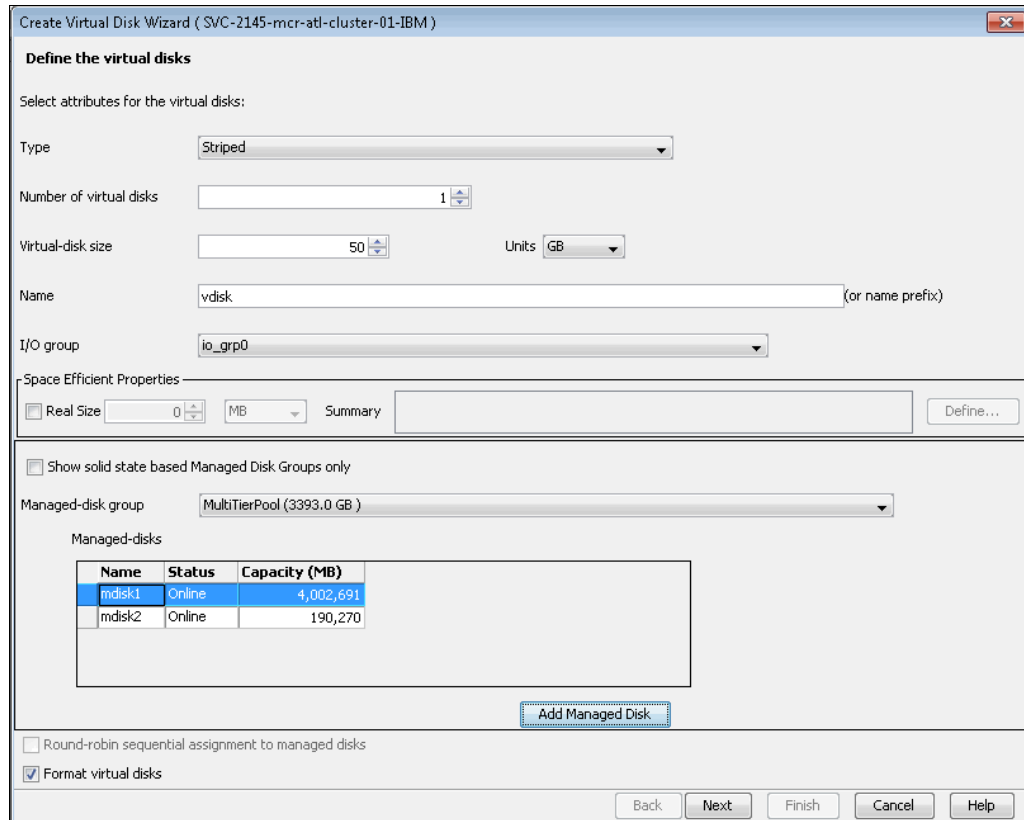


Figure 11-90 Virtual Disk wizard Creation

11.9.2 Generating reports by using Java GUI

In this section, we describe how to generate sample reports by using the GUI. We also create a probe to collect information from IBM Storwize V3700:

1. Navigate to **IBM Tivoli Storage Productivity Center** → **Monitoring** → **Probe**, right-click, and select **Create Probe**, as shown in Figure 11-91.

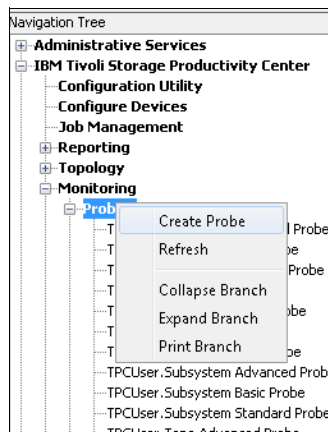


Figure 11-91 Create Probe option

2. Add the IBM Storwize V3700 probe for collecting information, as shown in Figure 11-92.




Figure 11-92 Adding IBM Storwize V3700 in probe

3. After you create a probe, you can click **Create Subsystem Performance Monitor**, as shown in Figure 11-93.

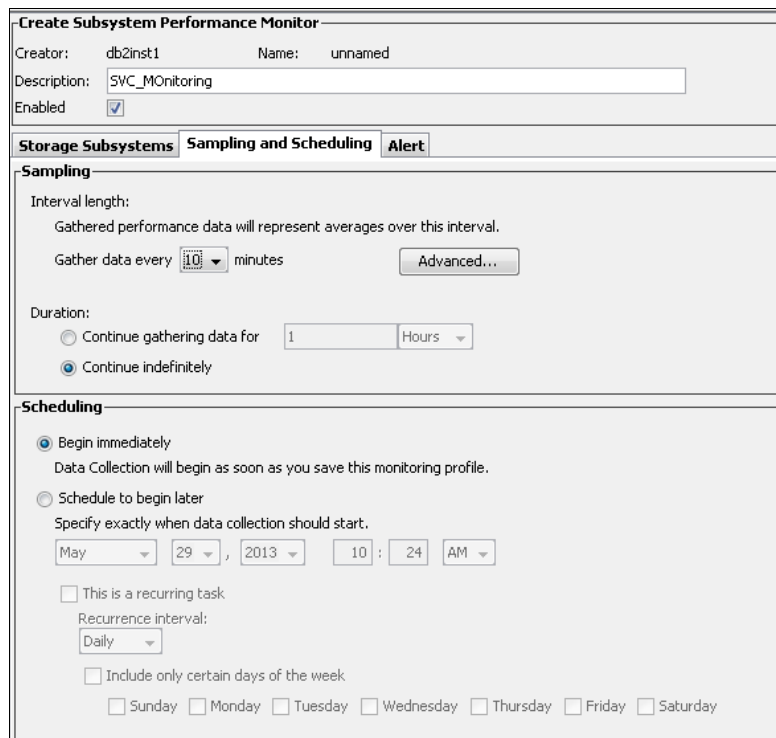


Figure 11-93 Create subsystem performance monitor

- To check the MDisk performance, click **Disk Manager** → **Reporting** → **Storage Subsystem Performance** → **By Managed Disk**. You see many options to include in the wizard to check MDisk performance, as shown in Figure 11-94.

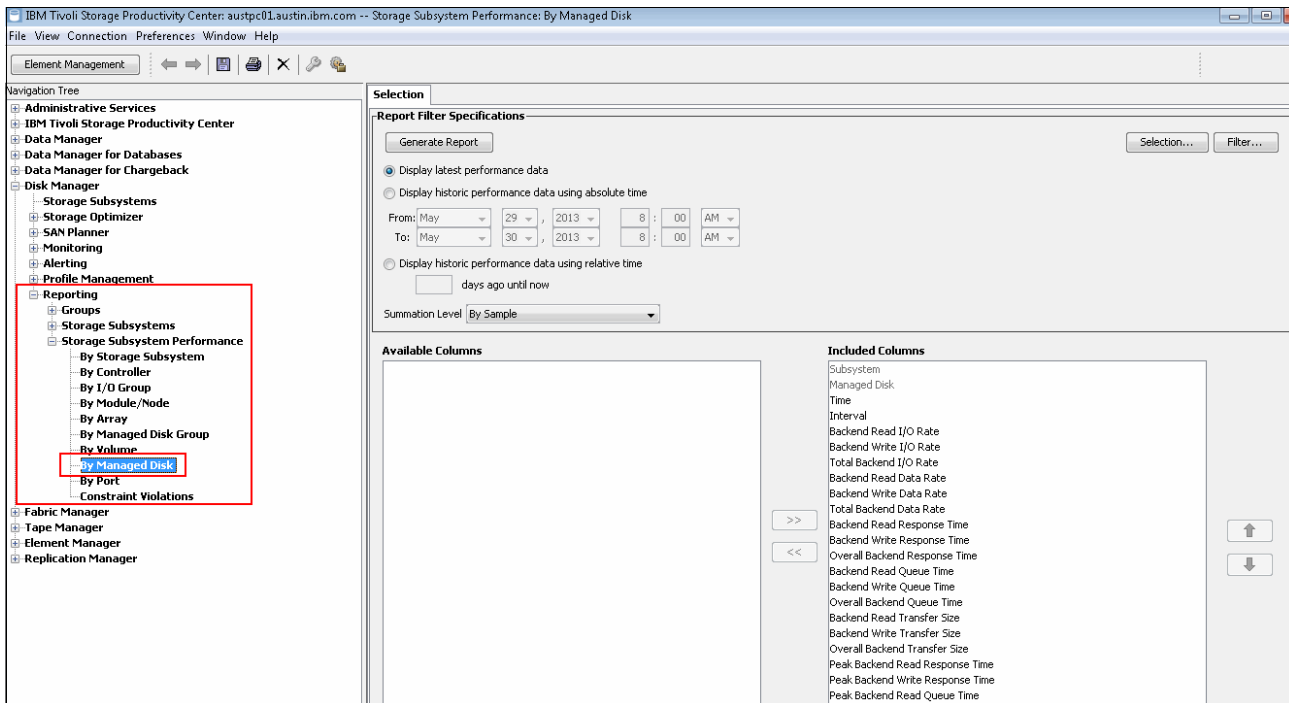


Figure 11-94 Managed disk performance report filter specification

- Click **Generate Report** to see a report, as shown in Figure 11-95.

Subsystem	Managed Disk	Time	Interval	Backend Read I/O Rate	Backend Write I/O Rate	Total Backend I/O Rate
SVC-2145-mcr-atl-cluster-01-IBM	mdisk0	May 30, 2013 7:50:20 AM	600 s	8 ops/s	0 ops/s	8 ops/s
SVC-2145-mcr-atl-cluster-01-IBM	mdisk1	May 30, 2013 7:50:20 AM	600 s	0 ops/s	8.01 ops/s	8.01 ops/s
SVC-2145-mcr-atl-cluster-01-IBM	mdisk2	May 30, 2013 7:50:20 AM	600 s	0 ops/s	0.37 ops/s	0.37 ops/s

Figure 11-95 MDisk performance report

6. Click the upper left icon to see a history chart report of the selected MDisk, as shown in Figure 11-96.

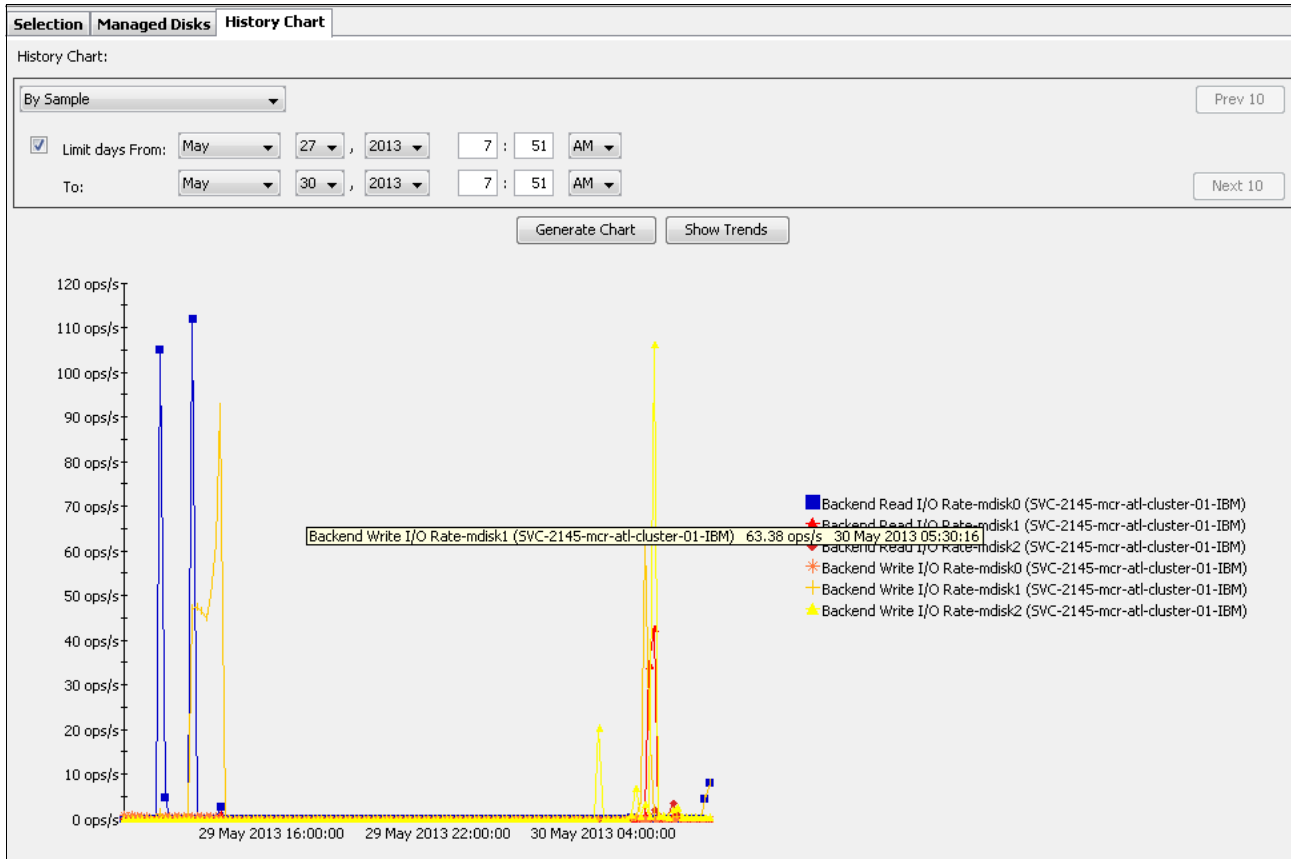


Figure 11-96 MDisk history chart

11.9.3 Generating reports using Tivoli Storage Productivity Center web console

In this section, we describe how to generate reports using the Tivoli Storage Productivity Center web console:

1. To connect to the web page, browse to the following URL:

<https://<tpchostname>.com:9569/srm/>

2. A login pane is displayed (as shown in Figure 11-97), enabling a user to log in using their Tivoli Storage Productivity Center credentials.



Figure 11-97 Tivoli Storage Productivity Center login pane

3. After you log in, the Tivoli Storage Productivity Center web dashboard is displayed, as shown in Figure 11-98. The Tivoli Storage Productivity Center web-based GUI is used to show information about the storage resources in your environment. It contains predefined and custom reports about performance and storage tiering.

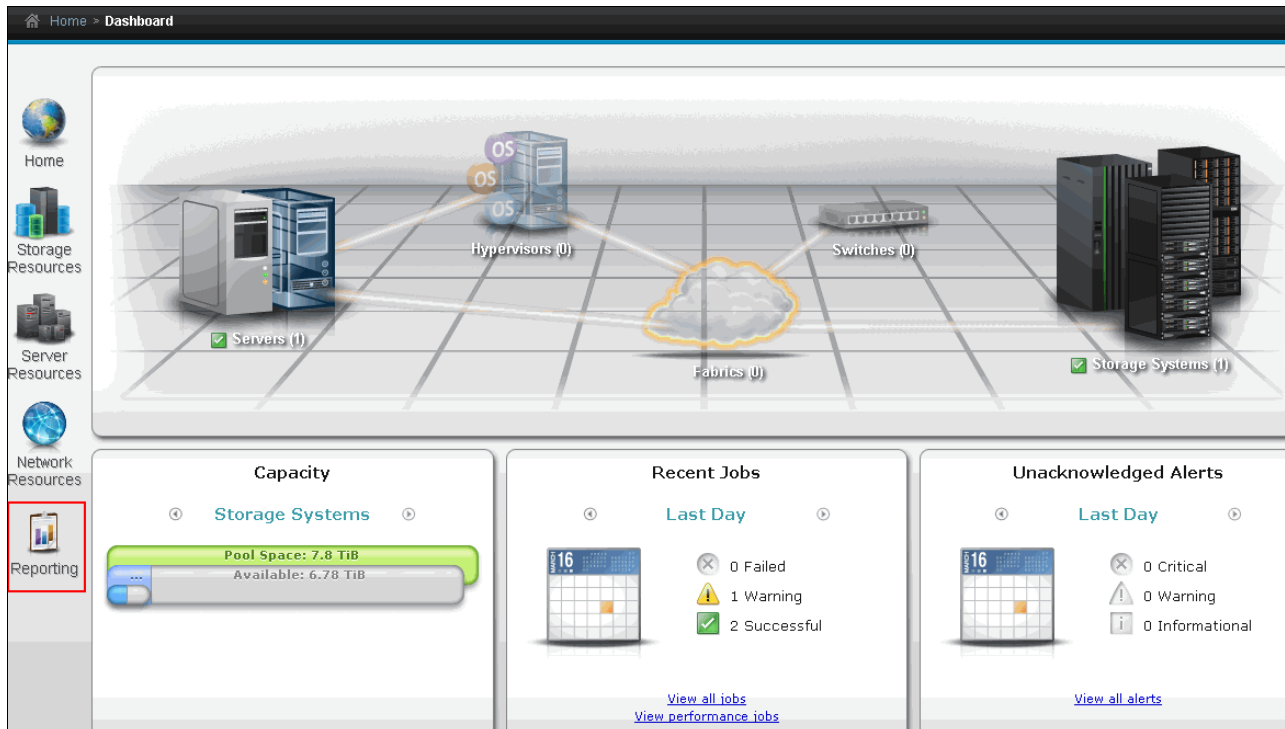


Figure 11-98 Tivoli Storage Productivity Center Dashboard

4. IBM Tivoli Common Reporting can be used to **View predefined reports**, **Create a basic report**, and **Design a report** from the web-based GUI, as shown in Figure 11-99.



Figure 11-99 Tivoli Storage Productivity Center web-based reporting

5. Figure 11-100 shows how to select predefined Storage Tiering reporting.

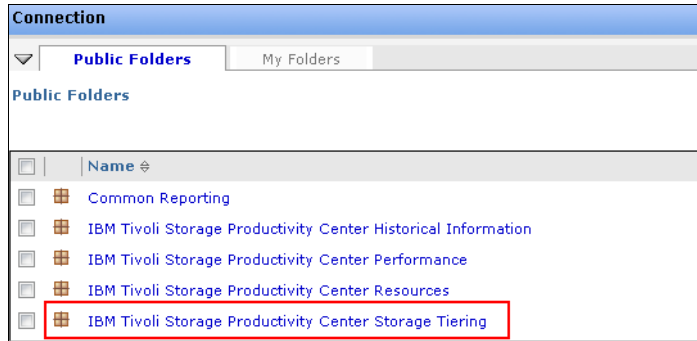


Figure 11-100 Tivoli Storage Productivity Center Storage tiering reporting

6. Figure 11-101 shows the different report options for Storage Tiering.



Figure 11-101 Details reports

7. Figure 11-102 shows the output from VDisk Details report.

MDisk Group		VDisk	VDisk Report Date	VDisk Capacity (GB)	VDisk Thin Provisioning Capacity Status	VDisk Peak Utilization (%)	VDisk Read I/O Rate (ops/s)	VDisk Max Read I/O Rate (ops/s)	VDisk Average Read Response Time (ms/op)	VDisk Max Read Response Time (ms/op)	VDisk Max Read Response Time Threshold of (70 % Exceeded By (hr.:min.))	VDisk Write I/O Rate (ops/s)	VDisk Average Write Response Time (ms/op)	VDisk Max Write Response Time (ms/op)	VDisk Max Write Response Time Threshold of (90 % Exceeded By (hr.:min.))	VDisk Max Read Density Rate (ops/s per GB)	VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)	VDisk Capacity Share of Total VDisk Capacity (%)	VDisk Average Read Cache Hits (%)	VDisk Max Read Cache Hits Actual (%)	VDisk Max Read Cache Hits Time Threshold of (75 % Exceeded By (hr.:min.))	VDisk Max Read I/O Rate Threshold of (80 % Exceeded By (hr.:min.))	VDisk Max Write I/O Rate Threshold of (80 % Exceeded By (hr.:min.))
MultiTierPool		win_2k8-iscsi	May 28, 2013	5.00	OK	0.06	0.00	0.06	8.25	8.36	00:00	0.00	1.34	1.34	00:00	0.01	0.08	0.31	32.65	100.00	00:00	00:00	00:00
		win_2k8-iscsi		Avg 5.00		Avg 0.06	Avg 0.00	Avg 0.06	Avg 8.25	Avg 8.36		Avg 0.00	Avg 1.34		Avg 0.01			Avg 32.65	Avg 100.00				

Figure 11-102 VDisk Details report

8. Figure 11-103 shows the Report Overview in pie-chart format.




Figure 11-103 Reporting Overview

9. Figure 11-104 shows the Easy Tier usage for volumes. To open this report in Tivoli Storage Productivity Center, click **Storage Resources** → **Volumes**.

Status	Pool	Capacit...	Hosts	Allocated Space (GiB)	Physical Allocation (%)	Easy Tier	Tier Capacity SSD (G...	Tier Capacity HDD (G...	Last Data Collection
Online		36.00	None	36.00	Fully Allocated	Enabled/Inactive	0.00	36.00	May 30, 2013 8:57:51 AM CDT
Online		36.00	aus...	36.00	Fully Allocated	Enabled/Inactive	0.00	36.00	May 30, 2013 8:57:51 AM CDT
Online		36.00	aus...	36.00	Fully Allocated	Enabled/Inactive	0.00	36.00	May 30, 2013 8:57:52 AM CDT
Online		36.00	lucia	36.00	Fully Allocated	Enabled/Inactive	0.00	36.00	May 30, 2013 8:57:51 AM CDT
Online		36.00	soaf...	36.00	Fully Allocated	Enabled/Inactive	0.00	36.00	May 30, 2013 8:57:51 AM CDT
Online		36.00	stiaf...	36.00	Fully Allocated	Enabled/Inactive	0.00	36.00	May 30, 2013 8:57:50 AM CDT
Online		46.00	None	46.00	Fully Allocated	Enabled/Inactive	0.00	46.00	May 30, 2013 8:57:51 AM CDT
Online		50.00	rmo...	50.00	Fully Allocated	Enabled/Inactive	0.00	50.00	May 30, 2013 8:57:51 AM CDT
Online		100.00	None	100.00	Fully Allocated	Enabled/Inactive	0.00	100.00	May 30, 2013 8:57:51 AM CDT
Online		500.00	aus...	500.00	Fully Allocated	Enabled/Inactive	0.00	500.00	May 30, 2013 8:57:51 AM CDT
Online		600.00	stiaf...	600.00	Fully Allocated	Enabled/Inactive	0.00	600.00	May 30, 2013 8:57:50 AM CDT
Online		600.00	stiaf...	600.00	Fully Allocated	Enabled/Inactive	0.00	600.00	May 30, 2013 8:57:50 AM CDT
Online		600.00	stiaf...	600.00	Fully Allocated	Enabled/Inactive	0.00	600.00	May 30, 2013 8:57:50 AM CDT

Figure 11-104 Volume Easy Tier usage

10. Figure 11-105 shows a detailed list of Storage Pools.

Name	Physical Allocation (%)	Shortfall (%)	Solid St...	Easy Tier	Tier Capacity SSD (...)	Tier Capacity HDD (G...	Tier Available Space SSD (G...	Last Data Collection
	34%	0%	No	Auto/Inactive	0.00	25,098.75	0.00	May 30, 2013 8:57:26 AM CDT
	34%	0%	No	Auto/Inactive	0.00	25,098.75	0.00	May 30, 2013 8:57:26 AM CDT
	77%	0%	No	Auto/Active	3,339.75	102,710.75	0.00	May 30, 2013 8:57:26 AM CDT
	78%	0%	No	Auto/Active	3,339.75	102,712.25	0.00	May 30, 2013 8:57:26 AM CDT
	97%	0%	No	Auto/Active	3,339.75	102,718.00	0.00	May 30, 2013 8:57:26 AM CDT
	98%	0%	No	Auto/Active	3,339.75	102,712.25	0.00	May 30, 2013 8:57:26 AM CDT
MultiTierPool	99%	100%	No	N/A				May 30, 2013 9:01:01 AM CDT
V5000_Pool_1	0%	12%	No	N/A				May 30, 2013 9:01:01 AM CDT
V5000_Pool_3	1%	0%	No	N/A				May 30, 2013 9:01:01 AM CDT

Figure 11-105 Pool Easy Tier information

11. Figure 11-106 shows Storage Virtualized Pool details in graph format.

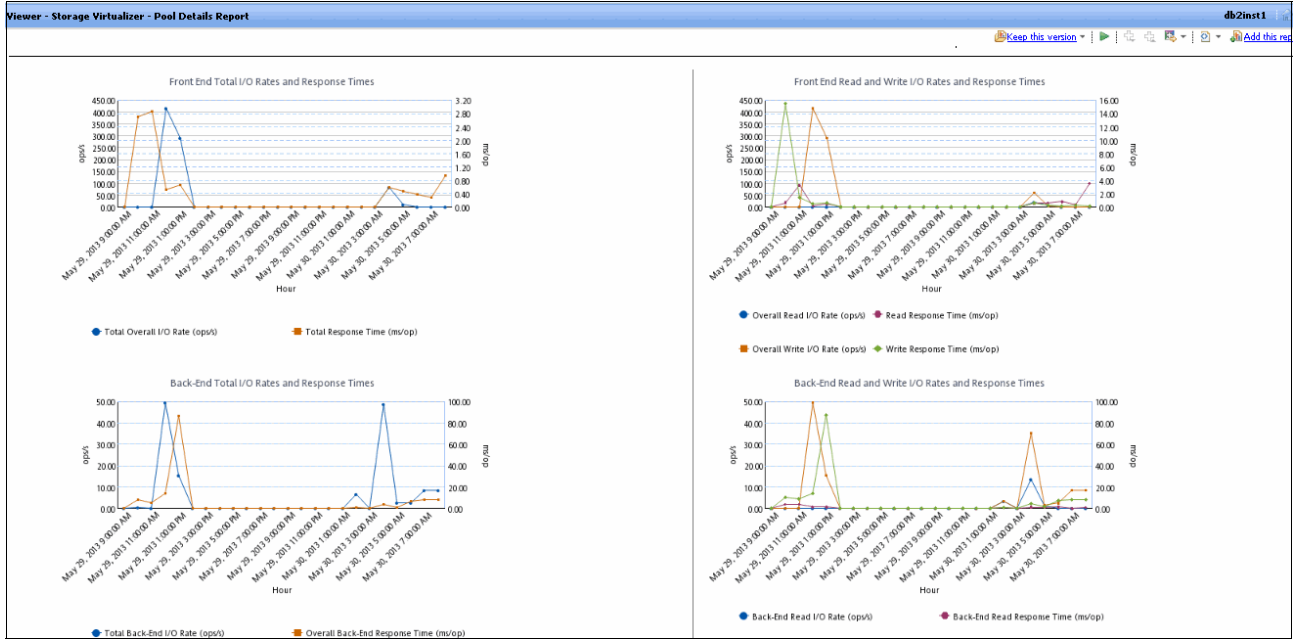


Figure 11-106 Pool details



A

Command-line interface setup and SAN Boot

This appendix describes the setup of the command-line interface (CLI), and provides more information about the storage area network (SAN) Boot function.

This appendix includes the following sections:

- ▶ Command-line interface
- ▶ SAN Boot

Command-line interface

The IBM Storwize V3700 system has a powerful CLI, which offers even more functions than the graphical user interface (GUI). This section is not intended to be a detailed guide to the CLI. That topic is beyond the scope of this book. The basic configuration of the IBM Storwize V3700 CLI and some example commands are covered.

The CLI commands are the same as those in the IBM SAN Volume Controller but, in addition, there are more commands that are available to manage the internal storage of the Storwize V3700. If a task is completed in the GUI, the CLI command is always displayed in the details, as shown throughout this book.

Detailed CLI information is available in the IBM Storwize V3700 Knowledge Center under the Command Line section, which can be found on the following website:

<http://www-01.ibm.com/support/knowledgecenter/STLM5A/welcome>

Implementing the IBM Storwize V7000 V7.2, SG24-7938, also has information about the use of the CLI. The commands in that book also apply to the IBM Storwize V3700 system because it is part of the Storwize family.

Basic setup

In the IBM Storwize V3700 GUI, authentication is done using a user name and a password. The CLI uses a Secure Shell (SSH) to connect from the host to the IBM Storwize V3700 system. A private and public key pair or user name and password is necessary. The following steps are required to enable CLI access with SSH keys:

1. A public key and private key are generated as a pair.
2. A public key is uploaded to the IBM Storwize V3700 system using the GUI.
3. A client SSH tool is configured to authenticate with the private key.
4. A secure connection is established between the client and IBM Storwize V3700 system.

SSH is the communication vehicle that is used between the management workstation and the IBM Storwize V3700 system. The SSH client provides a secure environment from which to connect to a remote machine. It uses the principles of public and private keys for authentication.

SSH keys are generated by the SSH client software. The SSH keys include a public key, which is uploaded and maintained by the clustered system, and a private key, which is kept private on the workstation that is running the SSH client. These keys authorize specific users to access the administration and service functions on the system. Each key pair is associated with a user-defined ID string that can consist of up to 40 characters.

Up to 100 keys can be stored on the system. New IDs and keys can be added, and unwanted IDs and keys can be deleted. To use the CLI, an SSH client must be installed on that system, the SSH key pair must be generated on the client system, and the client's SSH public key must be stored on the IBM Storwize V3700 system.

The SSH client that is used in this book is PuTTY. There is also a PuTTY key generator that can be used to generate the private and public key pair. The PuTTY client can be downloaded at no cost from the following website:

<http://www.chiark.greenend.org.uk>

The following tools should be downloaded:

- ▶ PuTTY SSH client: `putty.exe`
- ▶ PuTTY key generator: `puttygen.exe`

If using a Microsoft Windows OS, it is suggested to download the Windows installer, which includes everything required. This is currently the `putty-0.63-installer.exe` file.

Generating a public and private key pair

To generate a public and private key pair, complete the following steps:

1. Start the PuTTY key generator to generate the public and private key pair, as shown in Figure A-1.




Figure A-1 PuTTY key generator

2. Make sure that the following options are selected:
 - Type of key to generate: **SSH2 RSA**
 - Number of bits in a generated key: 1024

3. Click **Generate** and move the cursor over the blank area to generate the keys, as shown in Figure A-2.

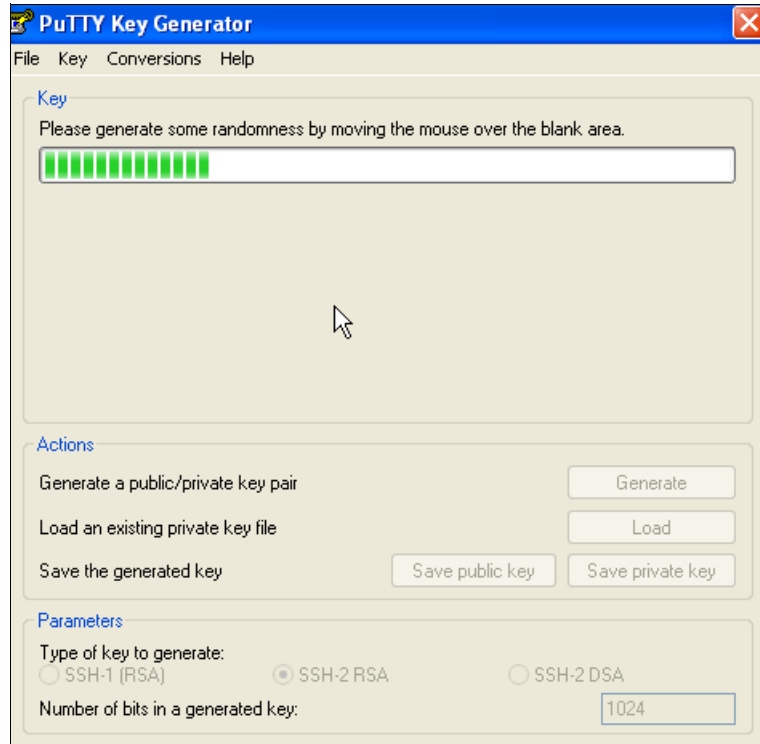


Figure A-2 Generate keys

Generating keys: The blank area that is indicated in this step is the large blank gray rectangle on the GUI inside the section of the GUI labeled Key indicated by the mouse pointer. Continue to move the mouse pointer over the blank area until the progress bar reaches the far right side. This action generates random characters to create a unique key pair.

More information about generating keys can be found in the PuTTY user manual. This is in the Help drop-down menu in the PuTTY GUI.

4. After the keys are generated, save them for later use. Click **Save public key**, as shown in Figure A-3.

Attention: You should always set a key pass phrase before saving the key. Not doing so means that the key will be stored on your workstation unencrypted. Any attacker that gains access to your private key will therefore *gain access to all machines configured to accept it*.

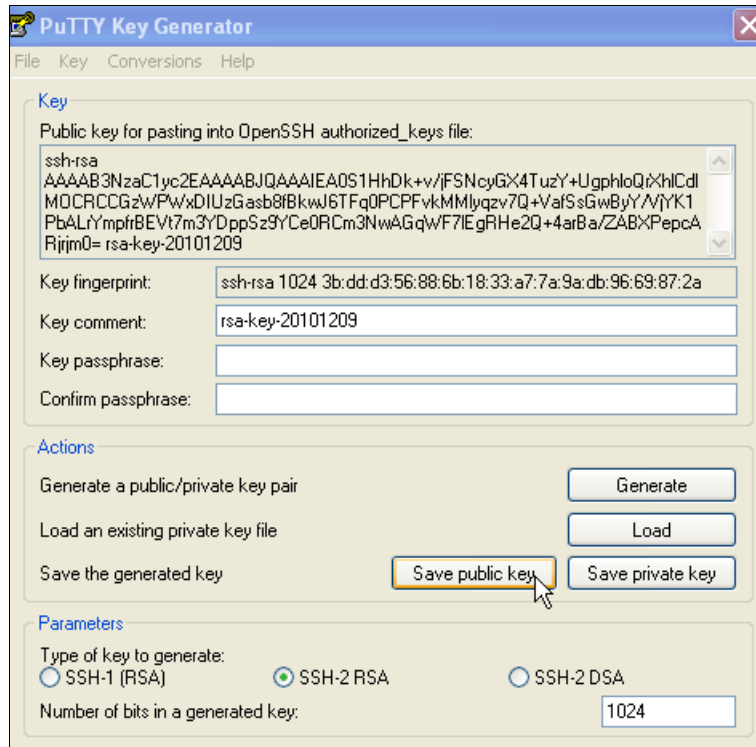


Figure A-3 Save public key

5. You are prompted for a name (for example, pubkey) and a location for the public key (for example, C:\Support Utils\PuTTY). Click **Save**.

Be sure to record the name and location of this SSH public key, because this information must be specified later.

Public key extension: By default, the PuTTY key generator saves the public key with no extension. Use the string pub for naming the public key (for example, pubkey), to differentiate the SSH public key from the SSH private key.

6. Click **Save private key**, as shown in Figure A-4.

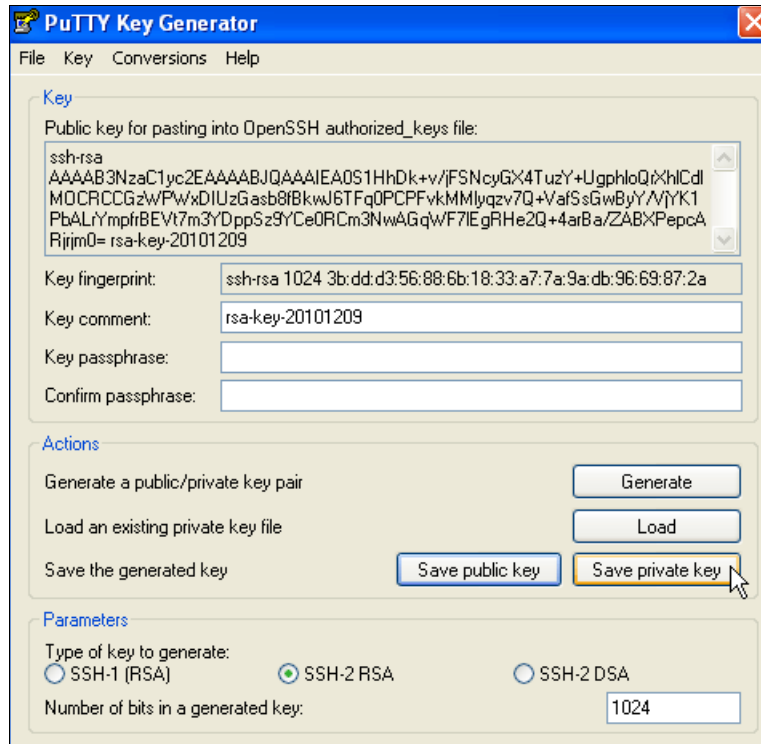


Figure A-4 Save private key

7. You receive a warning message, as shown in Figure A-5. Click **Yes** to save the private key without a passphrase.




Figure A-5 Confirm the security warning

8. When prompted, enter a name (for example, i cat), select a secure place as the location, and click **Save**.

Key generator: The PuTTY key generator saves the private key with a .PPK extension.

9. Close the PuTTY key generator.

For more information about generating the SSH keys, see the PuTTY user manual, found in the Help drop down menu in the PuTTY GUI.

Uploading the SSH public key to the IBM Storwize V3700

After you create your SSH key pair, you must upload your SSH public key onto the IBM Storwize V3700 system. Complete the following steps to upload the key:

1. Open the user section, as shown in Figure A-6.




Figure A-6 Open user section

2. Right-click the user you want to upload the key for, and click **Properties** (Figure A-7).

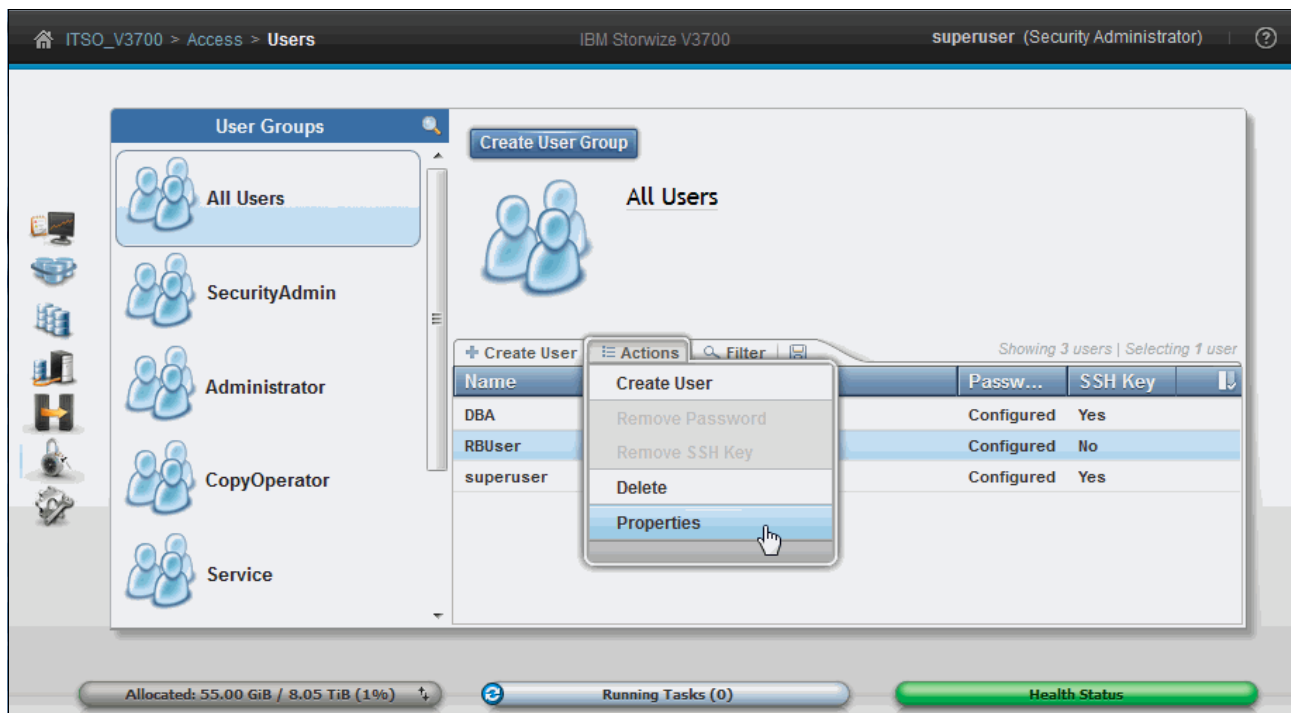


Figure A-7 User properties

3. To upload the public key, click **Browse**, select your public key, and click **OK**, as shown in Figure A-8.



Figure A-8 Select public key

4. Click **OK** and the task to upload the key is started.
5. Click **Close** to return to the GUI.

Configuring the SSH client

Before the CLI can be used, the SSH client must be configured. Complete the following steps to configure the client:

1. Start PuTTY, as shown in Figure A-9.

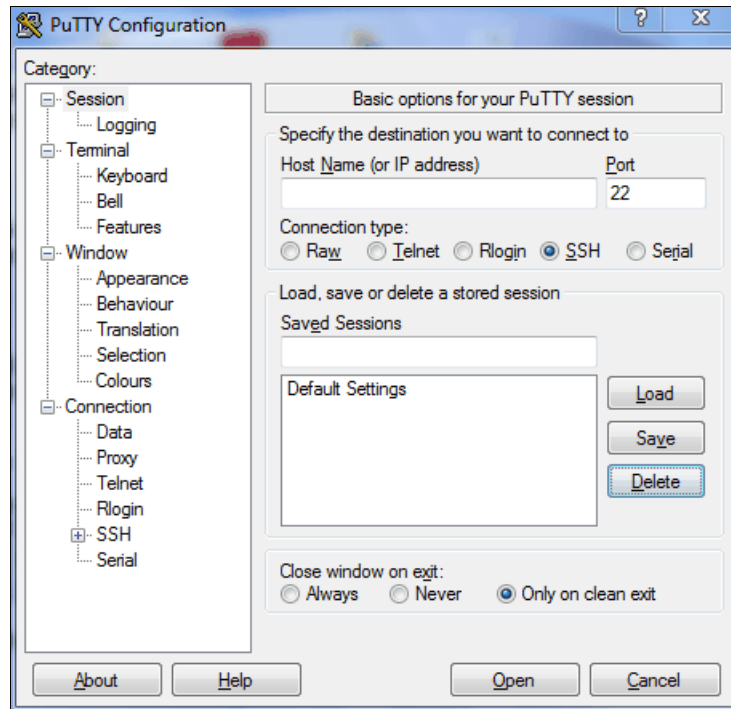


Figure A-9 PuTTY

2. In the right side pane under the **Specify the destination you want to connect to** section, select **SSH**.
3. Under the **Close window on exit** section, select **Only on clean exit**, which ensures that if there are any connection errors, they are displayed in the user's window.

- From the Category pane on the left side of the PuTTY Configuration window, click **Connection** → **SSH** to open the PuTTY SSH Configuration window, as shown in Figure A-10.

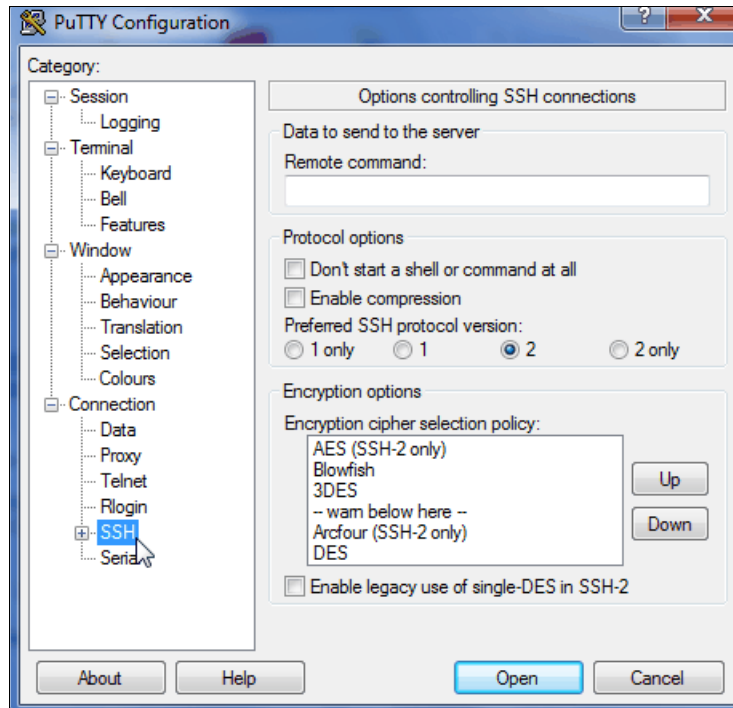


Figure A-10 SSH protocol version 2

- In the right side pane in the **Preferred SSH protocol version** section, select **2**.

- From the Category pane on the left side of the PuTTY Configuration window, click **Connection** → **SSH** → **Auth**.
- As shown in Figure A-11, in the right side pane in the **Private key file for authentication** field under the Authentication Parameters section, browse to or manually enter the fully qualified directory path and file name of the SSH client private key file that was created earlier (for example, C:\Support Utils\PuTTY\icat.PPK).




Figure A-11 SSH authentication

- From the Category pane on the left side of the PuTTY Configuration window, click **Session** to return to the Session view, as shown in Figure A-9 on page 657.

9. In the right side pane, enter the host name or system IP address of the IBM Storwize V3700 clustered system in the Host Name field. Enter a session name in the Saved Sessions field, as shown in Figure A-12.




Figure A-12 Enter session information

10. Click **Save** to save the new session, as shown in Figure A-13.

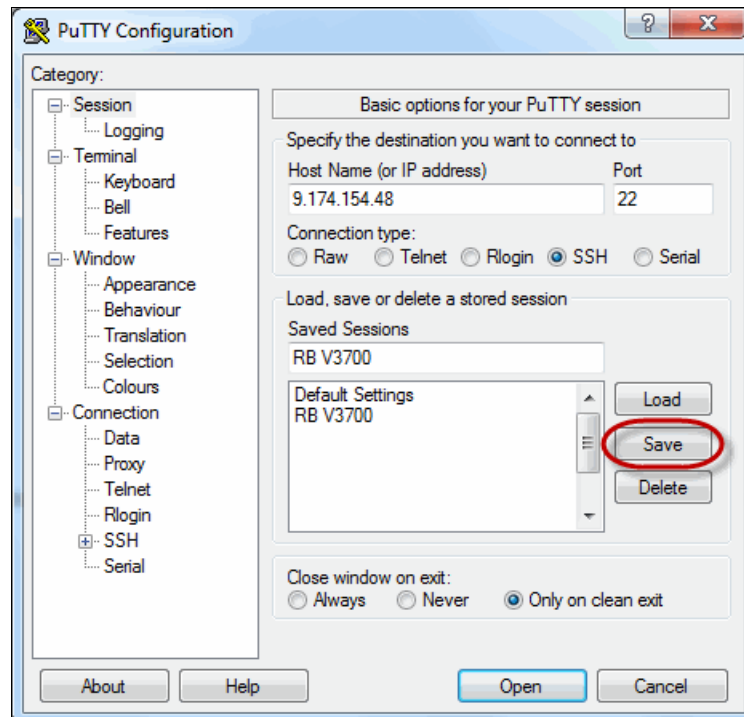


Figure A-13 Save Session

11. Highlight the new session and click **Open** to connect to the IBM Storwize V3700 system.
12. PuTTY now connects to the system and prompts you for a user name. Enter admin as the user name and press Enter (see Example A-1).

Example: A-1 Enter user name

```
login as: superuser
Authenticating with public key "rsa-key-20130521"
Last login: Tue May 21 15:21:55 2013 from 9.174.219.143
IBM_Storwize:ITS0_V3700:superuser>
```

The CLI is now configured for IBM Storwize V3700 administration.

Example commands

A detailed description about all of the available commands is beyond the intended scope of this book. Therefore, we reference some sample commands used elsewhere in this book.

The **svcin** and **svctask** prefixes are no longer needed in IBM Storwize V3700. If you have scripts that use this prefix, they still run without problems, but aren't necessary. If you enter **svcin** or **svctask** and press the Tab key twice, all of the available subcommands are listed. Pressing the Tab key twice also auto-completes commands if the input is valid and unique to the system.

To examine some of the available commands, complete the following steps:

1. Enter **lsvdisk** (as shown in Example A-2) to list all configured volumes on the system. The example shows that six volumes are configured.

Example: A-2 List all volumes

```
IBM_Storwize:ITS0_V3700:superuser>lsvdisk
id name          IO_group_id IO_group_name status mdisk_grp_id mdisk_grp_name
capacity type    FC_id FC_name RC_id RC_name vdisk_UID  opy_count
fast_write_state se_copy_count RC_change compressed_copy_count
0 V3700_Vol1 0          io_grp0      online 0          V3700_Pool
20.00GB striped                6005076300800 empty
1          no          0
1 V3700_Vol2 0          io_grp0      online 0          V3700_Pool
2.00GB striped                6005076300800 empty
1          no          0
2 V3700_Vol3 0          io_grp0      online 0          V3700_Pool
2.00GB striped                6005076300800 empty
1          no          0
3 V3700_Vol4 0          io_grp0      online 0          V3700_Pool
2.00GB striped                6005076300800 empty
1          no          0
4 V3700_Vol5 0          io_grp0      online 0          V3700_Pool
2.00GB striped                6005076300800 empty
1          no          0
5 V3700_Vol6 0          io_grp0      online 0          V3700_Pool
2.00GB striped                6005076300800 empty
1          no          0
```

2. Enter **lshost** to see a list of all configured hosts on the system, as shown in Example A-3.

Example: A-3 List hosts

```
IBM_Storwize:ITSO_V3700:superuser>lshost
id name          port_count iogrp_count status
0 windows2008r2 2           4           online
```

3. To map the volume to the hosts, enter **mkvdiskhostmap**, as shown in Example A-4.

Example: A-4 Map volumes to host

```
IBM_Storwize:ITSO_V3700:superuser>mkvdiskhostmap -host ESXi-1 -scsi 0 -force
ESXi-Redbooks
Virtual Disk to Host map, id [0], successfully created
```

4. To verify the host mapping, enter **lsvdiskhostmap**, as shown in Example A-5.

Example: A-5 List all hosts that are mapped to a volume

```
IBM_Storwize:ITSO_V3700:superuser>lshostvdiskmap ESXi-1
id name  SCSI_id vdisk_id vdisk_name  vdisk_UID
4 ESXi-1 0       2       ESXi-Redbooks 600507680185853FF000000000000011
```

5. In the CLI, there are more options available than in the GUI. All advanced settings can be set (for example, I/O throttling). To enable I/O throttling, change the properties of a volume by using the **changevdisk** command, as shown in Example A-6. To verify the changes, run the **lsvdisk** command.

Example: A-6 Enable advanced properties: I/O throttling

```
IBM_Storwize:ITSO_V3700:superuser>chvdisk -rate 1200 -unit mb ESXi-Redbooks
IBM_Storwize:ITSO_V3700:superuser>
IBM_Storwize:ITSO_V3700:superuser>lsvdisk ESXi-Redbooks

id 2
name ESXi-Redbooks
.
.
vdisk_UID 600507680185853FF000000000000011
virtual_disk_throttling (MB) 1200
preferred_node_id 2
.
.
IBM_Storwize:ITSO_V3700:superuser>
```

Command output: The **lsvdisk** command lists all available properties of a volume and its copies. To make it easier to read, lines in Example A-6 were deleted.

6. If you do not specify the **unit** parameter, the throttling is based on input/outputs (I/Os) rather than throughput, as shown in Example A-7.

Example: A-7 Throttling based on I/O

```
IBM_Storwize:ITS0_V3700:superuser>chvdisk -rate 4000 ESXi-Redbooks
IBM_Storwize:ITS0_V3700:superuser>lsvdisk ESXi-Redbooks
id 2
name ESXi-Redbooks
.
.
vdisk_UID 600507680185853FF000000000000011
throttling 4000
preferred_node_id 2
.
.
IBM_Storwize:ITS0_V3700:superuser>
```

7. To disable I/O throttling, set the I/O rate to 0, as shown in Example A-8.

Example: A-8 Disable I/O Throttling

```
IBM_Storwize:ITS0_V3700:superuser>chvdisk -rate 0 ESXi-Redbooks
IBM_Storwize:ITS0_V3700:superuser>lsvdisk ESXi-Redbooks
id 2
.
.
vdisk_UID 600507680185853FF000000000000011
throttling 0
preferred_node_id 2
IBM_Storwize:ITS0_V3700:superuser>
```

Upgrading Drive Firmware using the CLI

Note: Before upgrading any disk drive firmware, the Storwize system should be checked for any failures, and any that are found should be rectified before continuing.

Disk drives should not be upgraded if their associated arrays are not in a redundant state. Although the update process is designed not to take the drives offline, this cannot be guaranteed.

Running the following command for the drive that you are upgrading informs you of any possible issues:

```
lsdependentvdisks -drive drive_id
```

This returns notification of a possible issue if the drive in question is part of a non-redundant array (for example, a Redundant Array of Independent Disks (RAID) RAID0 array or a RAID5 array with a failed drive).

To perform the disk drive firmware upgrade, you require the following files:

- ▶ Upgrade Test Utility
- ▶ The Drive firmware package

These files can be downloaded from the IBM Support site:

<http://www-947.ibm.com/support/entry/portal/support>

When downloaded, copy the Upgrade Test Utility and Drive Firmware package to your PuTTY folder.

Like the Storwize controller firmware upgrade, the disk drive upgrade requires the use of the Upgrade Test Utility, which shows the drives that need to be upgraded, and checks if there are likely to be any issues.

Copying, installing, and running the Upgrade Test Utility on the Storwize unit

The following steps provide the details to run the upgrade test utility.

Copying the Upgrade Test Utility to the Storwize unit

To copy the Upgrade Test Utility, complete the following steps:

1. From a Windows command prompt, upload the Upgrade Test Utility to the Storwize unit using the PuTTY Secure Copy client (PSCP).
2. Run the following command from the PuTTY folder on your notebook or PC (normally `c:\program files\putty` or `c:\program files (x86)\putty`):

```
pscp -i hursley.ppk IBM2072_INSTALL_upgradetest_12.26  
superuser@9.174.152.17:/home/admin/upgrade
```

In this case, the following components are part of the command:

- `<hursley.ppk>` is the private key file generated when SSH was set up.
- `<IBM2072_INSTALL_upgradetest_12.26>` is the file name of the utility (use the current one).
- `<superuser>` is the user name.
- `<9.174.152.17>` is the management Internet Protocol (IP) address or name of the Storwize unit.

Change these components as appropriate to your environment.

3. It is also possible to upload the Upgrade Test Utility to the Storwize unit using the Storwize user name and password (not using an SSH private key), by using the following command:

```
pscp IBM2072_INSTALL_upgradetest_12.26 superuser@9.174.152.17:/home/admin/upgrade
```

4. You are prompted for the password for whichever user name you specified. This is shown in Example 11-5.

Example 11-5 Uploading test utility

```
C:\Program Files (x86)\PuTTY>pscp IBM2072_INSTALL_upgradetest_12.26 superuser@9  
174.152.17:/home/admin/upgrade  
superuser@9.174.152.17's password:  
IBM2072_INSTALL_upgradete | 59 kB | 59.6 kB/s | ETA: 00:00:00 | 100%
```

```
C:\Program Files (x86)\PuTTY>
```

Installing and running the Upgrade Test Utility

To install and run the Upgrade Test Utility, complete the following steps:

1. From the Storwize CLI, install the Upgrade Test Utility using the following command:

```
svcservicetask applysoftware -file IBM2072_INSTALL_upgradetest_12.26
```

2. Then run the Upgrade Test Utility using the following command:

```
svcupgradetest -f -d
```

The **-f** switch specifies that this is a drive firmware update. The **-d** switch shows firmware details for every disk drive. Omitting the **-d** switch gives a summary.

Example 11-6 shows the output of this command with both switches included.

Example 11-6 Test utility output

```
IBM_Storwize:ITS0_V3700:superuser>svcupgradetest -f -d
svcupgradetest version 12.26
```

Please wait, the test may take several minutes to complete.

```
+***** Warning found *****
```

This tool has found the internal disks of this system are not running the recommended firmware versions. Details follow:

Model	Latest FW	Current FW	Drive Info
ST9146853SS	B63E	B63D	Drive in slot 13 in enclosure 1
			Drive in slot 12 in enclosure 1
			Drive in slot 4 in enclosure 1
			Drive in slot 3 in enclosure 1
MK3001GRRB	SC2E	SC2C	Drive in slot 24 in enclosure 1
			Drive in slot 21 in enclosure 1
			Drive in slot 23 in enclosure 1
			Drive in slot 22 in enclosure 1
			Drive in slot 11 in enclosure 1

We recommend that you upgrade the drive microcode at an appropriate time. If you believe you are running the latest version of microcode, then check for a later version of this tool. You do not need to upgrade the drive firmware before starting the software upgrade.

Results of running svcupgradetest:
=====

The tool has found warnings.
For each warning above, follow the instructions given.

The tool has found 0 errors and 1 warnings
IBM_Storwize:ITS0_V3700:superuser>

The utility lists all drives in the system that require a firmware upgrade (if using the **-d** switch).

Copying the code to the Storwize unit

It is assumed that the disk drive firmware has already been downloaded and placed in the PuTTY folder. To copy the code, complete the following steps:

1. From a Windows command prompt, upload the disk drive firmware to the Storwize unit using the PSCP.
2. Run the following command from the PuTTY folder on your notebook or PC (normally `c:\program files\putty` or `c:\program files (x86)\putty`):

```
pscp -i hursley.ppk IBM2072_DRIVE_20140826  
superuser@9.174.152.17:/home/admin/upgrade
```

In this case, the following components are part of the command:

- `<hursley.ppk>` is the private key file generated when SSH was set up.
- `<IBM2072_DRIVE_20140826>` is the name of the firmware file.
- `<superuser>` is the user name.
- `<9.174.152.17>` is the management IP address or name of the Storwize unit.

Change these components as appropriate to your environment.

3. It is also possible to upload the disk drive firmware to the Storwize unit using the Storwize user name and password (not using a SSH private key), by using the following command:

```
pscp IBM2072_DRIVE_20140826 superuser@9.174.152.17:/home/admin/upgrade
```

You are prompted for a password for the user name in this case.

Applying the disk drive software

Depending on the Storwize code being used, there are several different options for applying the disk drive firmware:

1. If you are using Storwize code V7.1 or older, disk drive firmware can only be manually applied to one drive at a time, using the **applydrivesoftware** command for each individual disk. The output from the test utility shown in Example 11-6 on page 665 gives the drive slot number as the identifier.

However, to run the firmware upgrade on individual drives, the drive ID is required, not the slot ID. To obtain the drive ID from the slot ID use the **lsdrive** output, as shown in Example 11-7 on page 667. The output shown has been abbreviated for the sake of clarity.

Example 11-7 The lsdrive command output

```
IBM_Storwize:ITSO_V3700:superuser>lsdrive
```

id	status	use	tech_type	capacity	mdisk_id	mdisk_name	member_id	enclosure_id	slot_id
0	online	candidate	sas_ssd	185.8GB				1	17
1	online	candidate	sas_ssd	185.8GB				1	18
2	online	candidate	sas_ssd	185.8GB				1	6
3	online	candidate	sas_ssd	185.8GB				1	7
4	online	candidate	sas_ssd	185.8GB				1	16
5	online	candidate	sas_ssd	185.8GB				1	14
6	online	candidate	sas_ssd	185.8GB				1	15
7	online	candidate	sas_ssd	185.8GB				1	10
8	online	candidate	sas_ssd	372.1GB				1	1
9	online	candidate	sas_ssd	185.8GB				1	8
10	online	candidate	sas_ssd	185.8GB				1	9
11	online	candidate	sas_ssd	372.1GB				1	2
12	online	member	sas_hdd	136.2GB	3	mdisk3	3	1	13
13	online	member	sas_hdd	136.2GB	3	mdisk3	2	1	12
14	online	member	sas_hdd	136.2GB	3	mdisk3	1	1	5
15	online	member	sas_hdd	136.2GB	3	mdisk3	0	1	4
16	online	spare	sas_hdd	136.2GB				1	3
17	online	candidate	sas_hdd	278.9GB				1	24
18	online	candidate	sas_hdd	278.9GB				1	21

2. Run the command to apply drive firmware, as shown in Example 11-8.

Example 11-8 Applying drive firmware to a single disk

```
IBM_Storwize:ITSO_V3700:superuser>applydrivesoftware -file  
IBM2072_DRIVE_20140826 -type firmware -drive 12  
IBM_Storwize:ITSO_V3700:superuser>
```

3. With Storwize code V7.2 or later, it is possible to upgrade all drives using the **-a11** switch, as shown in Example 11-9.

Example 11-9 Applying drive firmware to all drives

```
IBM_Storwize:ITSO_V3700:superuser>applydrivesoftware -file  
IBM2072_DRIVE_20140826 -type firmware -a11  
IBM_Storwize:ITSO_V3700:superuser>
```

4. The command takes roughly two minutes per drive to complete. To confirm that all disks have been upgraded, rerun the upgrade test utility or check the internal storage from the GUI.

SAN Boot

IBM Storwize V3700 supports SAN Boot for Windows, VMware, and many other operating systems (OSs). It is also possible to migrate SAN Boot volumes from other storage systems onto the Storwize V3700. Each implementation or migration will be subtly different, depending on the OS, host bus adapter (HBA), and multipath driver to be used. SAN Boot support can also change. This is therefore beyond the scope of this book. For more information about SAN Boot, see the IBM Storwize V3700 Information Center:

<http://www-01.ibm.com/support/knowledgecenter/search/SAN%20Boot?scope=STHGJJ>

For more information about SAN Boot support for different operating systems with IBM Subsystem Device Driver (SDD), see the IBM System Storage *Multipath Subsystem Device Driver User's Guide*, GC52-1309, which is available on the following website:

<http://www.ibm.com/support/docview.wss?uid=ssg1S7000303>

Related publications

The publications listed in this section are considered particularly suitable for a more detailed description of the topics covered in this book.

IBM Redbooks

The following IBM Redbooks publications (or later versions as they become available) provide more information about the topics in this book. Some publications that are referenced in the following list might be available in softcopy only:

- ▶ *Implementing the IBM System Storage SAN Volume Controller V7.2*, SG24-7933
- ▶ *Implementing the IBM Storwize V7000 V7.2*, SG24-7938
- ▶ *IBM System Storage SAN Volume Controller and Storwize V7000 Best Practices and Performance Guidelines*, SG24-7521
- ▶ *Implementing an IBM b-type SAN with 8 Gbps Directors and Switches*, SG24-6116
- ▶ *IBM SAN Volume Controller and Storwize Family Native IP Replication*, REDP-5103

You can search for, view, download, or order these documents and other Redbooks, Redpapers, Web Docs, draft, and other materials, at the following website:

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

IBM Storwize V3700 publications

Storwize V3700 publications are available on the following website:

<https://ibm.biz/BdxyDL>

IBM Storwize V3700 support

Storwize V3700 support is available on the following website:

<https://ibm.biz/BdxyD9>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



Redbooks

Implementing the IBM Storwize V3700

(1.0" spine)

0.875" x 1.498"

460 <-> 788 pages



Implementing the IBM Storwize V3700



Redbooks®

Easily manage and deploy systems with embedded GUI

Experience rapid and flexible provisioning

Protect data with remote mirroring

Organizations of all sizes are faced with the challenge of managing massive volumes of increasingly valuable data. However, storing this data can be costly, and extracting value from the data is becoming more and more difficult.

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

Storwize V3700 delivers efficient, entry-level configurations that are specifically designed to meet the needs of small and midsize businesses. Designed to provide organizations with the ability to consolidate and share data at an affordable price, Storwize V3700 offers advanced software capabilities that are usually found in more expensive systems.

Built on innovative IBM technology, Storwize V3700 addresses the block storage requirements of small and midsize organizations, Storwize V3700 is designed to accommodate the most common storage network technologies. This design enables easy implementation and management.

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

The concepts in this book also relate to the IBM Storwize V3500.

This book was written at a software level of version 7 release 4.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, clients, and IBM Business Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-8107-02

ISBN 0738440361