IBM

Implementing the IBM Storwize V7000 V7.2

- Discover the exciting new version of the IBM virtualization family
- Become familiar with the leading edge intuitive GUI
- See how simple IBM virtualization really is

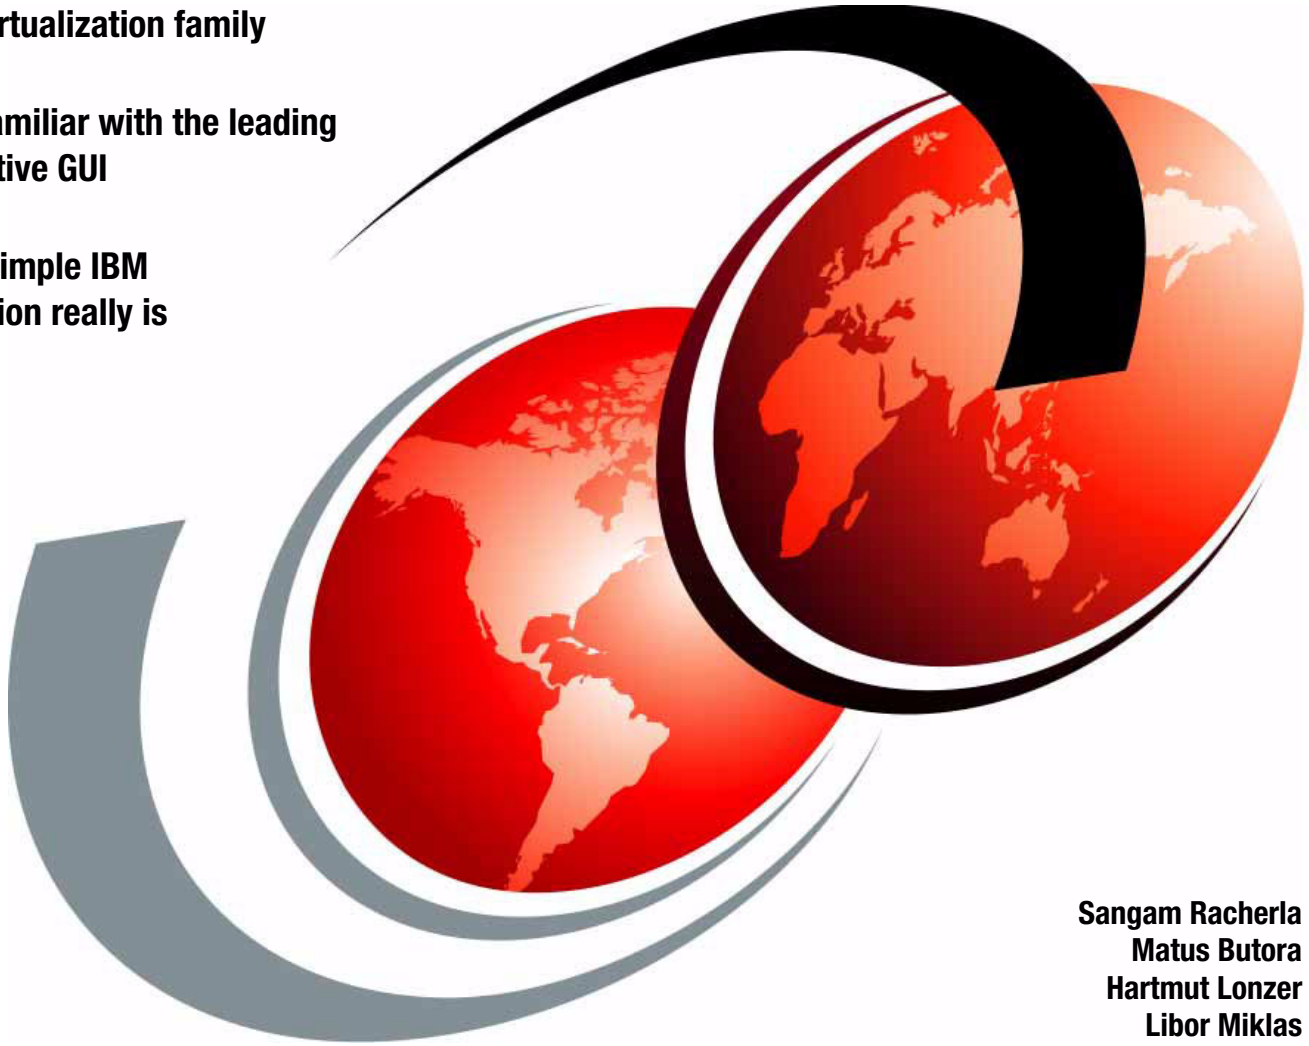

Sangam Racherla
Matus Butora
Hartmut Lonzer
Libor Miklas

Redbooks

ibm.com/redbooks

International Technical Support Organization

Implementing the IBM Storwize V7000 V7.2

July 2014

SG24-7938-02

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

Third Edition (July 2014)

This edition applies to the IBM Storwize V7000 Version 7.2.

Contents

Notices ix
Trademarks x

Preface xi
Authors xi
Now you can become a published author, too! xii
Comments welcome xiii
Stay connected to IBM Redbooks xiii

Summary of changes xv
July 2014, Third Edition xv

Chapter 1. IBM Storwize V7000 system overview 1
1.1 Storage virtualization 2
1.2 IBM Storwize V7000 overview 3
1.3 IBM Storwize V7000 terminology 4
1.3.1 IBM Storwize V7000 models 5
1.3.2 IBM Storwize V7000 attributes 7
1.3.3 IBM Storwize V7000 functions 7
1.3.4 IBM Storwize V7000 licensing 9
1.4 IBM Storwize V7000 hardware 10
1.4.1 Control enclosure 11
1.4.2 Expansion enclosure 13
1.4.3 Disk drive types 14
1.5 IBM Storwize V7000 components 14
1.5.1 Hosts 14
1.5.2 Nodes 15
1.5.3 I/O groups 15
1.5.4 Clustered system 15
1.5.5 RAID 16
1.5.6 Managed disks 17
1.5.7 Quorum disks 17
1.5.8 Storage pools 18
1.5.9 Volumes 20
1.5.10 Thin-provisioned volumes 22
1.5.11 Mirrored volumes 23
1.5.12 Easy Tier 24
1.5.13 iSCSI 24
1.5.14 Real-time Compression 26
1.5.15 IP replication 27
1.6 Advanced copy services 27
1.6.1 Synchronous or asynchronous remote copy 28
1.6.2 FlashCopy 29
1.6.3 Copy Services configuration limits 30
1.7 Management and support tools 31
1.7.1 IBM Assist On-site and remote service 31
1.7.2 Event notifications 31
1.7.3 SNMP traps 31
1.7.4 Syslog messages 32

1.7.5 Call Home email 32
1.8 Useful Storwize V7000 websites 32
1.8.1 IBM Storwize V7000 learning videos on YouTube 33

Chapter 2. Initial configuration 35
2.1 Hardware installation planning 36
2.2 SAN configuration planning 39
2.3 LAN configuration planning 40
2.3.1 Management IP address considerations 41
2.3.2 Service IP address considerations 42
2.4 Host configuration planning 42
2.5 Miscellaneous configuration planning 43
2.6 System management 44
2.6.1 Graphical user interface (GUI) 45
2.6.2 Command-line interface (CLI) 45
2.7 First-time setup of IBM Storwize V7000 46
2.8 Initial configuration 54
2.8.1 Configure Call Home, email alert, and inventory 58

Chapter 3. Graphical user interface overview 65
3.1 Overview of IBM Storwize V7000 management software 66
3.1.1 Access to the storage management software 66
3.1.2 Overview panel layout 67
3.1.3 Navigation 68
3.1.4 Multiple selections 69
3.1.5 Status indicators area 70
3.2 Home menu 72
3.3 Monitoring menu 73
3.3.1 System overview 73
3.3.2 System Details 75
3.3.3 Events 77
3.3.4 Performance 78
3.4 Pools menu 79
3.4.1 Volumes by pool 80
3.4.2 Internal storage 81
3.4.3 External storage 82
3.4.4 MDisks by Pools 84
3.4.5 System Migration 85
3.5 Volumes menu 85
3.5.1 All volumes 86
3.5.2 Volumes by Pool 87
3.5.3 Volumes by Host 87
3.6 Hosts menu 88
3.6.1 Hosts 88
3.6.2 Ports by Host 90
3.6.3 Host Mappings 91
3.6.4 Volumes by Host 91
3.7 Copy Services 92
3.7.1 FlashCopy 92
3.7.2 Consistency Group 93
3.7.3 FlashCopy Mappings 94
3.7.4 Remote Copy 96
3.7.5 Partnerships 97

3.8 Access menu 98
3.8.1 Users 98
3.8.2 Audit Log menu 99
3.9 Settings menu 100
3.9.1 Event notifications 100
3.9.2 Directory services 101
3.9.3 Network menu 102
3.9.4 Support 103
3.9.5 General 104

Chapter 4. Host configuration 105
4.1 Host attachment overview 106
4.2 Preparing the host operating system 107
4.2.1 Windows 2008 R2 and 2012: Preparing for Fibre Channel attachment 107
4.2.2 Windows 2008 R2 and Windows 2012: Preparing for iSCSI attachment 110
4.2.3 VMware ESX: Preparing for Fibre Channel attachment 112
4.2.4 VMware ESX: Preparing for iSCSI attachment 115
4.3 Creating hosts using the GUI 121
4.3.1 Creating Fibre Channel hosts 122
4.3.2 Creating iSCSI hosts 125

Chapter 5. Basic volume configuration 131
5.1 Provisioning storage from IBM Storwize V7000 to hosts 132
5.1.1 Creating generic volumes 133
5.1.2 Creating thin-provisioned volumes 135
5.1.3 Creating mirrored volumes 137
5.1.4 Creating a thin-mirrored volume 140
5.1.5 Creating compressed volumes 141
5.2 Mapping a volume to the host 142
5.2.1 Mapping newly created volumes to the host using the wizard 142
5.3 Discovering volumes on hosts and multipathing 143
5.3.1 Windows 2008 Fibre Channel volume attachment 144
5.3.2 Windows 2008 iSCSI volume attachment 148
5.3.3 VMware ESX Fibre Channel attachment 155
5.3.4 VMware ESX iSCSI attachment 159

Chapter 6. Migration wizard 167
6.1 Preparing for data migration 168
6.2 Data migration using the migration wizard 168

Chapter 7. Storage pools 179
7.1 Working with internal drives 180
7.1.1 Actions on internal drives 183
7.1.2 Configuring internal storage 187
7.2 Working with MDisks 197
7.2.1 Adding MDisks to storage pools 199
7.2.2 Importing MDisks 203
7.2.3 RAID action for MDisks 209
7.2.4 Selecting the tier for MDisks 211
7.2.5 Additional actions on MDisks 213
7.3 Working with storage pools 217

Chapter 8. Advanced host and volume administration 221
8.1 Advanced host administration 222

8.1.1 Modifying Mappings menu 224
8.1.2 Unmapping all volumes from a host 230
8.1.3 Duplicating and Importing Mappings 231
8.1.4 Renaming 236
8.1.5 Deleting a host 237
8.1.6 Host properties 238
8.2 Adding and deleting host ports 241
8.2.1 Adding a Fibre Channel or iSCSI host port 242
8.2.2 Deleting a host port 248
8.3 Host mappings overview 249
8.3.1 Unmapping a volume 250
8.3.2 Properties (Host) 251
8.3.3 Properties (Volume) 251
8.4 Advanced volume administration 252
8.4.1 Advanced volume functions 253
8.4.2 Mapping a volume to a host 255
8.4.3 Unmapping volumes from all hosts 256
8.4.4 Viewing a host mapped to a volume 257
8.4.5 Duplicating a volume 258
8.4.6 Renaming a volume 259
8.4.7 Shrinking a volume 260
8.4.8 Expanding a volume 261
8.4.9 Migrating a volume to another storage pool 261
8.4.10 Exporting to an image mode volume 264
8.4.11 Deleting a volume 266
8.5 Volume properties 267
8.5.1 Overview tab 268
8.5.2 Host Maps tab 271
8.5.3 Member MDisk tab 272
8.5.4 Adding a mirrored volume copy 273
8.5.5 Editing thin-provisioned or compressed volume properties 275
8.6 Advanced volume copy functions 279
8.6.1 Thin-provisioned or compressed volumes 281
8.6.2 Splitting into a new volume 282
8.6.3 Validate Volume Copies 283
8.6.4 Duplicate Volume copy 284
8.6.5 Delete Copy option 285
8.6.6 Migrating volumes using the volume copy features 286
8.7 Volumes by storage pool 286
8.7.1 Changing the Storage Pool function icon 289
8.7.2 Changing the storage pool name 290
8.8 Volumes by Host 292
8.8.1 Renaming a host 294

Chapter 9. External storage virtualization 297
9.1 Planning for external storage virtualization 298
9.1.1 License for external storage virtualization 298
9.1.2 SAN configuration planning 300
9.1.3 External storage configuration planning 301
9.1.4 Guidelines for virtualizing external storage 302
9.2 Working with external storage 302
9.2.1 Adding external storage 303
9.2.2 Managing external storage 303

9.2.3 Removing external storage 307

Chapter 10. Advanced features for storage efficiency 309
10.1 Introduction 310
10.2 Easy Tier 311
10.2.1 Easy Tier concepts 311
10.2.2 Disk tiers 312
10.2.3 Easy Tier process 314
10.2.4 Easy Tier operating modes 315
10.2.5 Implementation considerations 316
10.2.6 More information 317
10.3 Thin provisioning 317
10.3.1 Configuring thin-provisioned volume 318
10.3.2 Performance considerations 320
10.3.3 Limitations of virtual capacity 321
10.4 Real-time Compression 322
10.4.1 Concept of Real-time Compression 323
10.4.2 Configuring compressed volumes 324
10.4.3 Differences from IBM Real-time Compression Appliances 325

Chapter 11. Copy services 327
11.1 FlashCopy 328
11.1.1 Business requirements for FlashCopy 328
11.1.2 FlashCopy functional overview 329
11.1.3 Planning for FlashCopy 337
11.2 Managing FlashCopy using the GUI 338
11.2.1 Quick path to create FlashCopy presets 341
11.2.2 Managing FlashCopy mapping 343
11.2.3 Managing a FlashCopy consistency group 356
11.3 Volume mirroring and migration options 360
11.3.1 Managing volume mirror and migration with the GUI 362
11.4 Remote copy 369
11.4.1 Remote copy concepts 370
11.4.2 Native IP replication technology 373
11.4.3 Remote copy relationships 379
11.4.4 Remote Copy planning 387
11.5 Troubleshooting Remote Copy 389
11.5.1 1920 error 390
11.5.2 1720 error 392
11.6 Managing Remote Copy by using the GUI 393
11.6.1 Managing system partnerships 393
11.6.2 Managing stand-alone Remote Copy relationships 398
11.6.3 Managing a Remote Copy consistency group 407

Chapter 12. IBM Tivoli Storage Productivity Center 415
12.1 Tivoli Storage Productivity Center overview 416
12.1.1 Tivoli Storage Productivity Center V5.1 features and functions 416
12.1.2 Tivoli Storage Productivity Center V5.1 license models 417
12.1.3 IBM Tivoli Storage Productivity Center Select Edition V5.1 417
12.2 Tivoli Storage Productivity Center architecture 418
12.2.1 Data Server 418
12.2.2 Device Server 418
12.2.3 IBM Tivoli Integrated Portal 418
12.2.4 Tivoli Storage Productivity Center for Replication 419

12.2.5 IBM DB2 Database 419
12.2.6 Agents 419
12.2.7 Interfaces 419
12.3 Preparing Windows to install Tivoli Storage Productivity Center 420
12.3.1 Installation overview 420
12.3.2 Product code layout 421
12.3.3 Pre-installation steps for Windows 422
12.4 Installing Tivoli Storage Productivity Center components 434
12.5 Connecting Tivoli Storage Productivity Center to an IBM Storwize V7000 system . . 443
12.6 Administering and reporting an IBM Storwize V7000 system through Tivoli Storage Productivity Center 449
12.6.1 Basic configuration and administration 449
12.6.2 Report generation 451

Chapter 13. RAS, monitoring, and troubleshooting 455
13.1 Reliability, availability, and serviceability (RAS) 456
13.1.1 Node canisters 456
13.1.2 Disk subsystem 459
13.1.3 Power 461
13.2 Configuration backup 463
13.2.1 Configuration backup using the CLI 463
13.2.2 Configuration backup using the GUI 464
13.3 Software upgrade 465
13.4 Critical Fix Notification feature 471
13.5 Monitoring host paths to the IBM Storwize V7000 system 472
13.6 Troubleshooting and fix procedures 473
13.6.1 Managing event log 475
13.6.2 Running a fix procedure 476
13.6.3 Event log details 480
13.7 Monitoring 482
13.7.1 Email notifications and Call Home 482
13.8 Audit log 488
13.9 Collecting support information 490
13.10 Shutting down IBM Storwize V7000 492
13.10.1 Shutting down a node canister 493
13.10.2 Shutting down a system 495
13.10.3 Shutting down and powering on an IBM Storwize V7000 infrastructure 495

Appendix A. CLI setup and SAN Boot 497
Command-line interface (CLI) 498
Basic setup 498
SAN Boot 508
Enabling SAN Boot for Windows 508
Enabling SAN Boot for VMware 508
Windows SAN Boot migration 509

Related publications 511
IBM Redbooks 511
Other publications 511
Online resources 512
Help from IBM 512

Notices

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

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

Trademarks

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

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

AIX®
DB2®
DS8000®
Easy Tier®
FlashCopy®
IBM®
Passport Advantage®
Real-time Compression™
Real-time Compression Appliance™
Redbooks®
Redbooks (logo)  ®
Storwize®
System Storage®
System x®
Tivoli®
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.

Preface

Continuing its commitment to developing and delivering industry-leading storage technologies, IBM® introduces the IBM Storwize® V7000 solution, an innovative new storage offering that delivers essential storage efficiency technologies and exceptional ease of use and performance, all integrated into a compact, modular design that is offered at a competitive, midrange price.

The IBM Storwize V7000 solution incorporates some of the top IBM technologies typically found only in enterprise-class storage systems, raising the standard for storage efficiency in midrange disk systems. This cutting-edge storage system extends the comprehensive storage portfolio from IBM and can help change the way organizations address the ongoing information explosion.

This IBM Redbooks® publication introduces the features and functions of the IBM Storwize V7000 system through several examples. This book is aimed at pre- and post-sales technical support and marketing, storage administrators, and will help you understand the architecture of the Storwize V7000, how to implement it, and take advantage of the industry leading functions and features.

Authors

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, San Jose Center.

Sangam Racherla is an IT Specialist. He holds a degree in Electronics and Communication Engineering and has eleven years of experience in the IT field. His areas of expertise include Microsoft Windows, Linux, IBM AIX®, IBM System x®, and IBM System P servers, and various SAN and storage products.

Matus Butora is an IT Specialist and leader of Storage Support in the IBM ITD Delivery Center in the Czech Republic. He works with enterprise storage environments providing solutions and support for global strategic customers across the globe. Matus has ten years of experience with Open storage hardware and software including IBM DS8000®, IBM Midrange Storage Storwize, NetApp, Tivoli® Storage Management, and Tivoli Storage Productivity Center. Matus is a certified IBM Professional, Brocade, and NetApp certified administrator.

Hartmut Lonzer is a System Partner Technical Sales Specialist in Germany. He works in the German headquarters in Ehningen. As a former Storage Field Technical Sales Support (FTSS) member, his main focus is on storage and System x. He is responsible for educating, supporting, and enabling the Business Partners in his area in technical matters. His experience with SAN Volume Controller and V7000 storage goes back to the beginning of these products. Hartmut has been with IBM in various technical roles for 36 years.

Libor Miklas is an IT Architect working at the IBM Integrated Delivery Center in the Czech Republic. He demonstrates ten years of extensive experience within the IT industry. During the last eight years, his main focus has been on the data protection solutions and on storage management. Libor and his team design, implement, and support midrange and enterprise storage environments for various global and local clients, worldwide. He is an IBM Certified Deployment Professional for the Tivoli Storage Manager family of products and holds a Master's degree in Electrical Engineering and Telecommunications.

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

Megan Gilge
International Technical Support Organization, San Jose Center

Thanks to the authors of the previous editions of this book.

- Authors of the first edition, *Implementing the IBM Storwize V7000*, published in May 2011:

  Brian Cartwright
  Ronda Hruby
  Daniel Koeck
  Xin Liu
  Massimo Rosati
  Jon Tate
  Thomas Vogel
  Bill Wiegand

- Authors of the second edition, *Implementing the IBM Storwize V7000 V6.3*, published in February 2012:

  Alejandro Berardinelli
  Mark Chitti
  Torben Jensen
  Massimo Rosati
  Christian Schroeder
  Jon Tate

Now you can become a published author, too!

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

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

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

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

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

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

- Stay current on recent Redbooks publications with RSS Feeds:

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

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 SG24-7938-02
for Implementing the IBM Storwize V7000 V7.2
as created or updated on July 21, 2014.

July 2014, Third Edition

This revision contains the following new and changed information.

New information

- IP replication

Changed information

- Updates to reflect version 7.2

1

IBM Storwize V7000 system overview

This chapter provides an overview of the IBM Storwize V7000 architecture and includes a brief explanation of storage virtualization.

1.1 Storage virtualization

Storage virtualization, like server virtualization, is one of the foundations of building a flexible and reliable infrastructure solution that allows companies to better align their business and IT needs.

Storage virtualization allows an organization to implement pools of storage across physically separate disk systems (which might be from different vendors). Storage can then be deployed from these pools and can be migrated between pools without any outage of the attached host systems. Storage virtualization provides a single set of tools for advanced functions, such as instant copy and remote mirroring solutions, which means that deploying storage can be performed by using a single tool regardless of the underlying storage hardware.

Figure 1-1 shows a storage virtualization scenario.

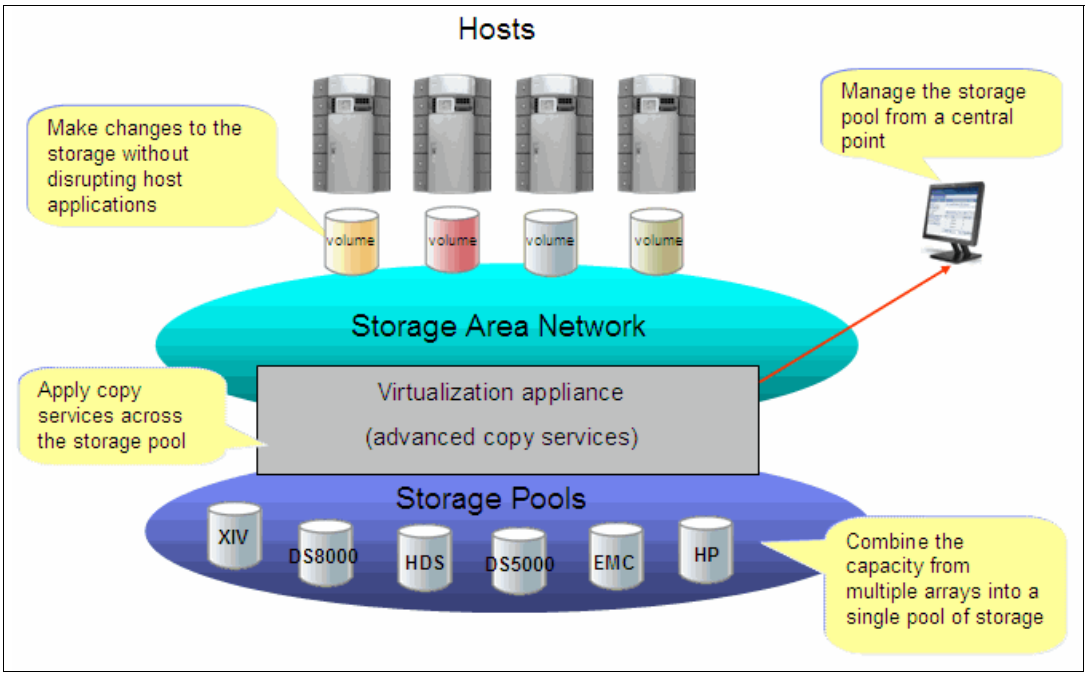

Figure 1-1 Storage virtualization

For a more detailed explanation of storage virtualization, see *Implementing the IBM System Storage SAN Volume Controller V7.2*, SG24-7933.

1.2 IBM Storwize V7000 overview

The IBM Storwize V7000 solution provides a modular storage system that includes the capability to virtualize external SAN-attached storage and its own internal storage. The IBM Storwize V7000 solution is built upon the IBM SAN Volume Controller technology base and uses technology from the IBM System Storage® DS8000 family.

An IBM Storwize V7000 system provides several configuration options that are aimed at simplifying the implementation process. It also provides automated wizards, called *Directed Maintenance Procedures (DMP)*, to assist in resolving any events that might occur. An IBM Storwize V7000 system is a clustered, scalable, and midrange storage system, and an external virtualization device.

Figure 1-2 shows the IBM Storwize V7000 high-level overview.

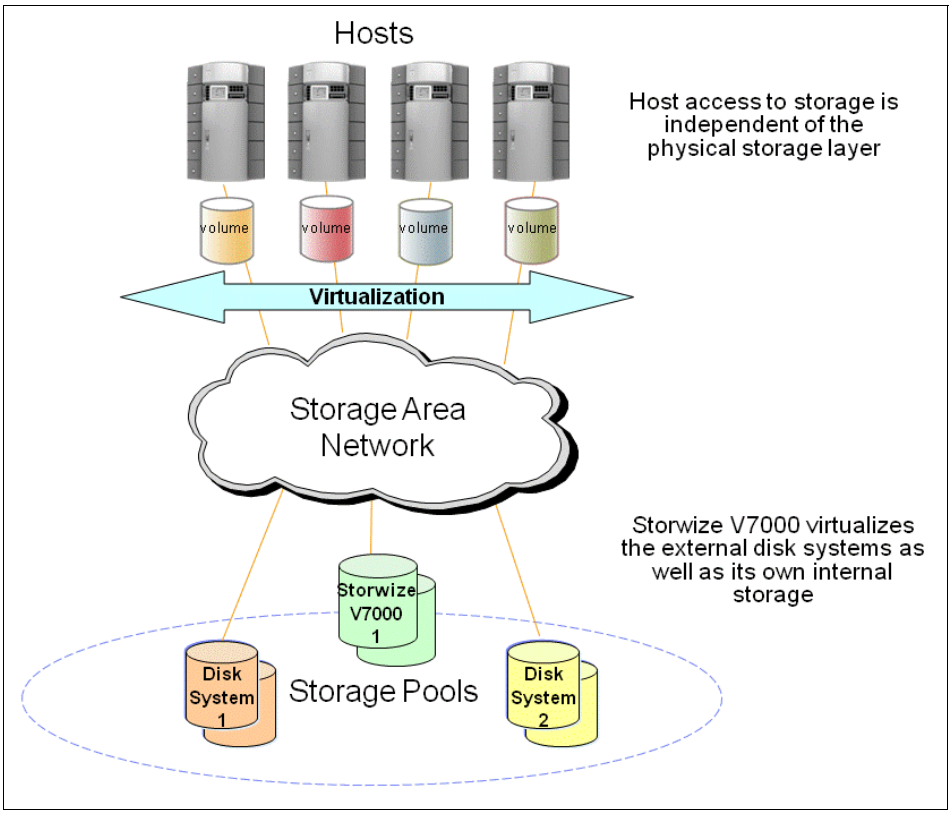

Figure 1-2 IBM Storwize V7000 overview

Included with an IBM Storwize V7000 system is a graphical user interface (GUI) that allows storage to be deployed quickly and efficiently. The GUI runs on the IBM Storwize V7000 system, so there is no need for a separate console. The management GUI contains a series of preestablished configuration options that are called *presets* and that use common 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.

The IBM Storwize V7000 solution provides a choice of up to 480 x 3.5 inch or 960 x 2.5 inch serial-attached SCSI (SAS) drives for the internal storage in a clustered system and uses SAS cables and connectors to attach to the optional expansion enclosures. In a clustered system, the V7000 can provide 1.92 PiB capacity.

When virtualizing external storage arrays, an IBM Storwize V7000 system can provide up to 32 PiB of usable capacity. An IBM Storwize V7000 system supports a range of external disk systems similar to what the SAN Volume Controller supports today.

The IBM Storwize V7000 solution consists of one to four control enclosures and, optionally, up to 36 expansion enclosures (and supports the intermixing of the different expansion enclosures). Within each enclosure are two canisters. Control enclosures contain two node canisters, and expansion enclosures contain two expansion canisters.

1.3 IBM Storwize V7000 terminology

IBM Storwize V7000 system uses the terminology defined in Table 1-1.

Table 1-1 IBM Storwize V7000 terminology

IBM Storwize V7000 term	Definition
Chain	A set of enclosures that are attached to provide redundant access to the drives that are inside the enclosures. Each control enclosure can have one or more chains.
Clone	A copy of a volume on a server at a particular point in time. The contents of the copy can be customized while the contents of the original volume are preserved.
Control enclosure	A hardware unit that includes the chassis, node canisters, drives, and power sources that include batteries.
Event	An occurrence that is significant to a task or system. Events can include completion or failure of an operation, a user action, or the change in the state of a process.
Expansion canister	A hardware unit that includes the serial-attached SCSI (SAS) interface hardware that enables the node hardware to use the drives of the expansion enclosure.
Expansion enclosure	A hardware unit that includes expansion canisters, drives, and power sources that do not include batteries.
External storage	Managed disks (MDisks) that are Small Computer Systems Interface (SCSI) logical units presented by storage systems that are attached to and managed by the clustered system.
Host mapping	The process of controlling which hosts have access to specific volumes within a clustered system.
Internal storage	Array managed disks (MDisks) and drives that are held in enclosures and nodes that are part of the clustered system.
Lane	A single SAS channel.
Managed disk (MDisk)	A component of a storage pool that is managed by a clustered system. An MDisk is either part of a RAID array of internal storage or a Small Computer System Interface (SCSI) logical unit (LU) for external storage. An MDisk is not visible to a host system on the storage area network.
Node canister	A hardware unit that includes the node hardware, fabric and service interfaces, and serial-attached SCSI (SAS) expansion ports.

IBM Storwize V7000 term	Definition
PHY	A single SAS lane. There are four PHYs in each SAS cable.
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. Quorum disks can either be MDisks or internal drives.
Snapshot	An image backup type that consists of a point-in-time view of a volume.
Storage pool	A collection of storage capacity that provides the capacity requirements for a volume.
Strand	The serial-attached SCSI (SAS) connectivity of a set of drives within multiple enclosures. The enclosures can be either control enclosures or expansion enclosures.
Thin provisioning or thin provisioned	The ability to define a storage unit (full system, storage pool, or volume) with a logical capacity size that is larger than the physical capacity assigned to that storage unit.
Volume	A discrete unit of storage on disk, tape, or other data recording medium that supports some form of identifier and parameter list, such as a volume label or input/output control.
Discovery	A process by which two Storwize clusters exchange information about their IP address configuration.
Remote Copy Port Group	The set of local and remote Ethernet ports that can access each other via a long-distance IP link.
Remote Copy Port Group ID	A numeric value that indicates to which group the port belongs.
RC login	A bidirectional full-duplex data path between two Storwize clusters that are remote copy (RC) partners.
Path configuration	The act of setting up RC logins for two partnered Storwize systems.

1.3.1 IBM Storwize V7000 models

The IBM Storwize V7000 platform consists of enclosures and drives. Each enclosure contains two canisters that, although they can be replaced independently, are seen as part of the enclosure.

Additional IBM Storwize V7000 information: For the most up-to-date information about the features, benefits, and specifications of IBM Storwize V7000 models, see the following address:

http://www.ibm.com/systems/storage/disk/storwize_v7000/index.html

The information in this book is valid at the time of writing, but as the IBM Storwize V7000 matures, expect to see new features and enhanced specifications.

The IBM Storwize V7000 models are described in Table 1-2.

Table 1-2 IBM Storwize V7000 models

Model	Cache	FC / iSCSI / SAS ports	Drive slots	Power supply
2076-112 (with two node canisters)	16 GB	8x8 Gb / 4x1 Gb / 4x6 Gb	12 x 3.5-inch	Integrated dual power supplies with battery backup
2076-124 (with two node canisters)	16 GB	8x8 Gb / 4x1 Gb / 4x6 Gb	24 x 2.5-inch	Integrated dual power supplies with battery backup
2076-312 (with two node canisters)	16 GB	8x8 Gb / 4x1 Gb+4x10Gb / 4x6 Gb	12 x 3.5-inch	Integrated dual power supplies with battery backup
2076-324 (with two node canisters)	16 GB	8x8 Gb / 4x1 Gb+4x10Gb / 4x6 Gb	24 x 2.5-inch	Integrated dual power supplies with battery backup
2076-212 (with two expansion canisters)	Not applicable (N/A)	-- / -- / 4x6 Gb	12 x 3.5-inch	Integrated dual power supplies
2076-224 (with two expansion canisters)	N/A	-- / -- / 4x6 Gb	24 x 2.5-inch	Integrated dual power supplies

Figure 1-3 shows the front view of the 2076-112, 212, and 312 enclosures.

Figure 1-3 IBM Storwize V7000 front view for 2076-112, 212, and 312 enclosures

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

Figure 1-4 shows the front view of the 2076-124, 224, and 324 enclosures.

Figure 1-4 IBM Storwize V7000 front view for 2076-124, 224, and 324 enclosures

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

1.3.2 IBM Storwize V7000 attributes

For a complete and updated list of IBM Storwize V7000 configuration limits and restrictions, see the following address:

http://www.ibm.com/support/docview.wss?uid=ssg1S1004511

1.3.3 IBM Storwize V7000 functions

The following functions are available with IBM Storwize V7000:

- Thin provisioning (No license required)

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

- Volume mirroring (No license required)

  Provides a single volume image to the attached host systems while maintaining pointers to two copies of data in separate storage pools. Copies can be on separate disk storage systems that are being virtualized. If one copy failing, IBM Storwize V7000 provides continuous data access by redirecting I/O to the remaining copy. When the copy becomes available, automatic resynchronization occurs.

- FlashCopy (included with the base IBM Storwize V7000 license)

  Provides a volume level point-in-time copy function for any storage being virtualized by IBM Storwize V7000. This function creates copies for backup, parallel processing, testing, and development, and have the copies available almost immediately.

  IBM Storwize V7000 includes the following FlashCopy functions:

  - Full or incremental copy

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

  - Multitarget FlashCopy

    IBM Storwize V7000 supports copying of up to 256 target volumes from a single source volume. Each copy is managed by a unique mapping and, in general, each mapping acts independently and is not affected by other mappings sharing the source volume.

  - Cascaded FlashCopy

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

  - Reverse FlashCopy

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

  - FlashCopy nocopy with thin provisioning

    This function provides a combination of using thin-provisioned volumes and FlashCopy together to help reduce disk space requirements when making copies.

This option has two variations:

- Space-efficient source and target with background copy

  Copies only the allocated space.

- Space-efficient target with no background copy

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

This function may be used with multitarget, cascaded, and incremental FlashCopy.

– Consistency groups

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

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

► Metro Mirror (licensed based on the number of enclosures and includes both Metro and Global Mirror)

Provides a synchronous remote mirroring function up to approximately 300 km between sites. As the host I/O only completes after the data is cached at both locations, performance requirements might limit the practical distance. Metro Mirror provides fully synchronized copies at both sites with zero data loss after the initial copy is completed.

Metro Mirror can operate between multiple IBM Storwize systems.

► Global Mirror (licensed based on capacity being mirrored and includes both Metro and Global Mirror)

Provides long distance asynchronous remote mirroring function up to approximately 8,000 km between sites. With Global Mirror, the host I/O completes locally and the changed data is sent to the remote site later. This function i maintains a consistent recoverable copy of data at the remote site, which lags behind the local site.

Global Mirror can operate between multiple IBM Storwize systems.

► Data migration (no charge for temporary usage)

IBM Storwize V7000 provides a data migration function that can be used to import external storage systems into the IBM Storwize V7000 system. You can do these tasks:

– Move volumes nondisruptively onto a newly installed storage system
– Move volumes to rebalance a changed workload
– Migrate data from other back-end storage to IBM Storwize V7000 managed storage

► IBM System Storage Easy Tier® (no charge)

Provides a mechanism to seamlessly migrate hot spots to the most appropriate tier within the IBM Storwize V7000 solution. This migration could be to internal drives within IBM Storwize V7000 or to external storage systems that are virtualized by IBM Storwize V7000.

This function is shown in Figure 1-5.

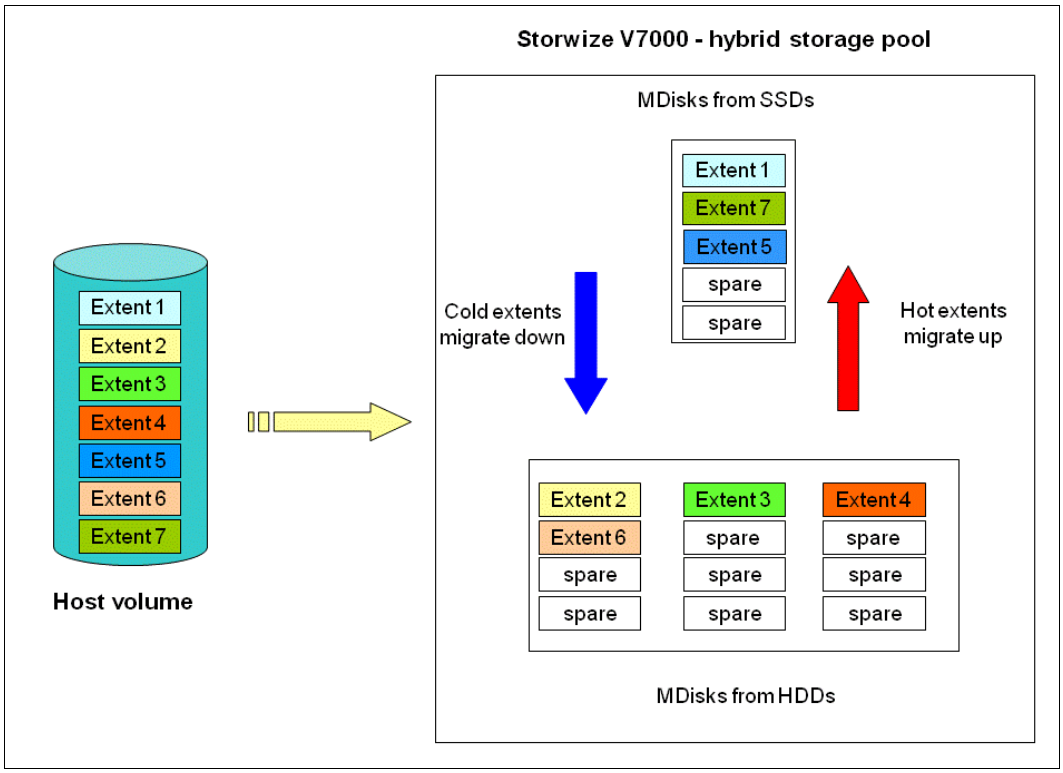

Figure 1-5 Easy Tier overview

1.3.4 IBM Storwize V7000 licensing

IBM Storwize V7000 might require the following licenses:

- Enclosure
- External Virtualization
- Remote Copy (Advanced Copy Services: Metro Mirror and Global Mirror)
- Compression

Table 1-3 summarizes the licenses that might be required.

Table 1-3 Licenses that might be required

License type	Unit	License number	License required?
Enclosure	Base+expansion Physical Enclosure Number	5639-VM7	Yes
External Virtualization	Physical Enclosure Number Of External Storage	5639-EV7	Optional
Remote Copy	Physical Enclosure Number	5639-RM7	Optional
Real-time Compression	Physical Enclosure Number	5639-CP7	Optional
FlashCopy	N/A	N/A	No
Volume Mirroring	N/A	N/A	No
Thin Provisioning	N/A	N/A	No

License type	Unit	License number	License required?
Volume Migration	N/A	N/A	No
Easy Tier	N/A	N/A	No

Migration: If the Storwize V7000 is used as a general migration tool, the appropriate External Virtualization licenses must be ordered. The only exception is if you want to migrate existing data from external storage to IBM Storwize V7000 internal storage; you can temporarily configure your External Storage license for use within 45 days. For a more-than-45-day migration requirement from external storage to IBM Storwize V7000 internal storage, the appropriate External Virtualization license must be ordered.

Real-time Compression: With Version 7.2, the Real-time Compression license is capped after the third license per control enclosure. The capped license covers both internal storage and externally virtualized storage.

1.4 IBM Storwize V7000 hardware

The IBM Storwize V7000 solution is a modular storage system that is built on a common enclosure (control enclosure and expansion enclosure) that is compliant with Storage Bridge Bay (SBB).

SBB is a specification that is created by a non-profit working group that defines a mechanical and electrical interface between a passive backplane drive array and the electronics packages that give the array its “personality.”

Figure 1-6 shows an overview of the hardware components of the IBM Storwize V7000 solution.

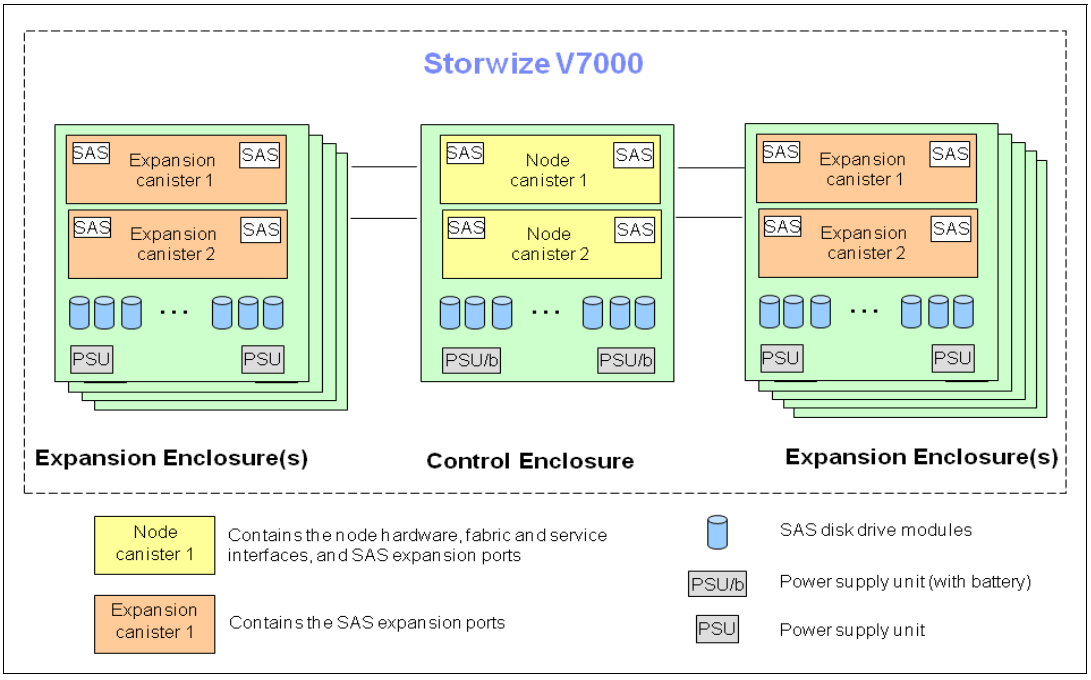

Figure 1-6 IBM Storwize V7000 hardware components

Figure 1-7 shows the controller rear view of models 112 and 124.

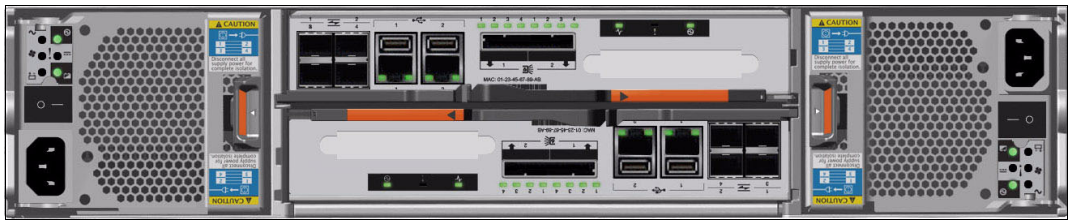

Figure 1-7 IBM Storwize V7000 controller rear view: models 112 and 124

The two power supply slots, on the extreme left and extreme right, each take up the full 2EIA height. The left slot is power supply 1, the right slot is power supply 2. The power supplies are inserted different ways up. Power supply 1 appears the correct way up, and power supply 2 upside down.

There are two *canister* slots, one above the other, in the middle of the chassis. The top slot is canister 1, the bottom slot canister 2. The canisters are inserted different ways up. Canister 1 appears the correct way up, and canister 2 upside down.

The new models, 312 and 324, look similar, but they have two 10 Gb Ethernet ports for use with iSCSI or Fibre Channel over Ethernet (FCoE), as shown in Figure 1-8.

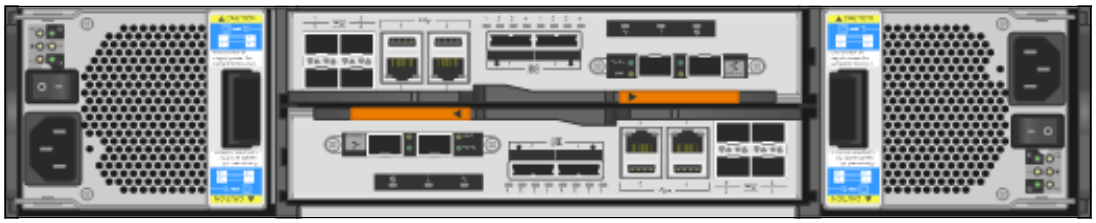

Figure 1-8 IBM Storwize V7000 controller rear view: models 312 and 324

1.4.1 Control enclosure

IBM Storwize V7000 scalability:

- For all systems running Version 6.1.0.7 support is enabled for up to five enclosures.
- For all systems running Version 6.1.0.7 or later, support is enabled for up to 10 enclosures.
- For all systems running Version 6.2.0 or later with two controller enclosures, support is enabled for up to 20 enclosures.
- For all systems running Version 6.4 or later with four controller enclosures, support is enabled for up to 40 enclosures.
- Any software upgrade is concurrent and does not affect any application downtime.

Each IBM Storwize V7000 system has one and optionally up to four control enclosures that contain two *node canisters* each, disk drives, and two power supplies. There are four models of the control enclosure with two models providing 12 3.5-inch disk slots and the other two models providing 24 2.5-inch disk slots. Two models (312 and 324) provide four 10 Gb Ethernet ports for use with iSCSI or FCoE.

Within a control enclosure, each power supply unit (PSU) contains a battery. The battery enables the IBM Storwize V7000 system to perform a memory dump of the cache to internal disks if both power inputs fail.

PSUs: The power supply units for the control enclosure and expansion enclosure are not interchangeable

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

The two nodes provide a highly available fault-tolerant controller so that if one node fails, the surviving node automatically takes over. Nodes are deployed in pairs called *I/O groups*.

With the optional four controllers, you have up to four I/O groups.

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

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

Figure 1-9 shows a single node canister.

Figure 1-9 A single canister for controller models 112 or 124

Four Fibre Channel ports are on the left side of the canister. They are in a block of four in two rows of two connectors. The ports are numbered 1 - 4, from left to right, top to bottom. The ports operate at 2, 4, or 8 Gbps. Use of the ports is optional. There are two green LEDs associated with each port: the speed LED and link activity LED.

Two 10/100/1000 Mbps Ethernet ports are side by side on the canister. They are numbered 1 on the left and 2 on the right. Using port 1 is required; port 2 is optional. Two LEDs are associated with each Ethernet port. The 2076 models 312 or 324 have four extra 10 Gb Ethernet ports (two per canister) for use with iSCSI or FCoE.

Two USB 2.0 ports are side by side on the canister. They are numbered 1 on the left and 2 on the right. Use of the connectors is optional. The only defined usage is with USB flash drives, which is described in Chapter 2, “Initial configuration” on page 35.

Two 6 Gbps SAS ports are side by side on the canister. They are numbered 1 on the left and 2 on the right. These ports are used to connect to the optional expansion enclosures.

1.4.2 Expansion enclosure

The optional expansion enclosure contains two *expansion canisters*, disk drives, and two power supplies. There are two models of the control enclosure: one model provides 12 disk slots and the other provides 24 disk slots.

Figure 1-10 shows the expansion enclosure.

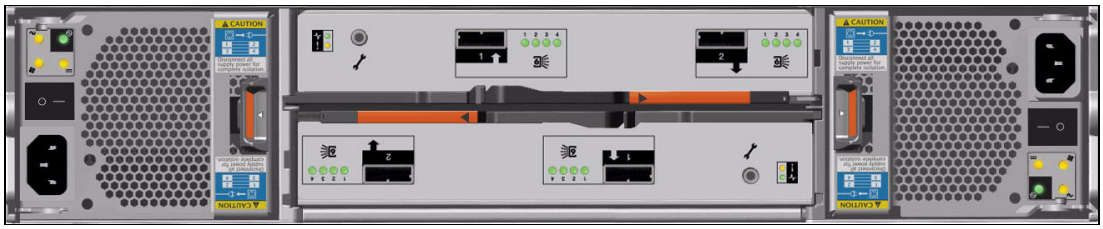

Figure 1-10 Expansion enclosure

The expansion enclosure power supplies are similar to the control enclosure but do not contain the battery. So they are not interchangeable with the controller PSU.There is a single power lead connector on the power supply unit. The PSU has an IEC C14 socket and the mains connection cable has a C13 plug.

Each expansion canister provides two SAS interfaces that are used to connect to the control enclosure and any optional expansion enclosures. The ports are numbered 1 on the left and 2 on the right. SAS port 1 is the IN port and SAS port 2 is the OUT port. There is also a symbol printed above the SAS ports to identify whether it is an IN or an OUT port.

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

Each port connects four PHYs (ports of SAS drives). Each is capable of a 6 Gb transfer rate. An LED is associated with each PHY in each port (eight LEDs in total) to indicate the activity on the PHY. The LEDs are green and are next to the ports; for each port, the LEDs are numbered 1 - 4. If traffic starts, it goes over PHY 1. If the line is saturated, the next PHY starts working. If you see all four PHY LEDs flashing, the back end is fully saturated.

Figure 1-11 shows an expansion canister.

Figure 1-11 A single expansion canister

1.4.3 Disk drive types

IBM Storwize V7000 enclosures currently support SSD, SAS, and Nearline SAS drive types. Each SAS drive has two ports (two PHYs) and I/O can be issued in both paths simultaneously.

Table 1-4 shows the IBM Storwize V7000 disk drive types that are available at the time of writing.

Table 1-4 IBM Storwize V7000 Disk Drive types

Drive type		Speed	Size
2.5-inch form factor	Solid-state drive	N/A	200, 400, and 800 GB
2.5-inch form factor	SAS	10,000 rpm	300, 600, and 900 GB, and 1.2 TB
2.5-inch form factor	SAS	15,000 rpm	146 and 300 GB
2.5-inch form factor	Nearline SAS	7,200 rpm	1 TB
3.5-inch form factor	Nearline SAS	7,200 rpm	2, 3, and 4 TB

1.5 IBM Storwize V7000 components

The IBM Storwize V7000 is a midrange virtualization RAID storage subsystem. It has the following benefits:

- Brings enterprise technology to midrange storage.
- Specialty administrators are not required.
- Client setup and service can be easy.
- The system can grow incrementally as storage capacity and performance needs change.
- Multiple storage tiers are in a single system with nondisruptive migration between them.
- Simple integration can be done into the server environment.

The IBM Storwize V7000 subsystem consists of a set of drive enclosures. Control enclosures contain disk drives and two nodes (an I/O group), which are attached to the SAN fabric. Expansion enclosures contain drives and are attached to control enclosures.

The simplest use of the IBM Storwize V7000 is as a traditional RAID subsystem. The internal drives are configured into RAID arrays and virtual disks created from those arrays.

The IBM Storwize V7000 can also be used to virtualize other storage controllers, as described in Chapter 9, “External storage virtualization” on page 297.

The IBM Storwize V7000 supports regular and solid-state drives and uses IBM System Storage Easy Tier to automatically place volume hot spots on better performing storage.

1.5.1 Hosts

A host system is a server that is connected to IBM Storwize V7000 through a Fibre Channel connection, FCoE, or through an iSCSI connection.

Hosts are defined to IBM Storwize V7000 by identifying their worldwide port names (WWPNs) for Fibre Channel hosts. For iSCSI hosts, they are identified by using their iSCSI names. The iSCSI names can either be iSCSI qualified names (IQNs) or extended unique identifiers (EUIs).

1.5.2 Nodes

IBM Storwize V7000 can have two or up to eight hardware components, called *nodes* or *node canisters*, that provide the virtualization of internal and external volumes, and cache and copy services (remote copy) functions. A clustered system consists of one or four node pairs.

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

1.5.3 I/O groups

Within IBM Storwize V7000, there are one to four pairs of node canisters known as *I/O groups*. The IBM Storwize V7000 supports eight node canisters in the clustered system, which provides four I/O groups.

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

Both nodes of the I/O group act as preferred nodes for their own specific subset of the total number of volumes that the I/O group presents to the host servers (a maximum of 2048 volumes per I/O group). However, both nodes also act as a failover node for its partner node within the I/O group, so a node takes over the I/O workload from its partner node, if required, with no impact to the server's application.

In a Storwize V7000 environment, using active/active architecture, the I/O handling for a volume can be managed by both nodes of the I/O group. Therefore, it is mandatory for servers that are connected through Fibre Channel connectors to use multipath device drivers to be able to handle this capability.

The Storwize V7000 I/O groups are connected to the SAN so that all application servers accessing volumes from the I/O group have access to them. Up to 2048 host server objects can be defined in four I/O groups.

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

1.5.4 Clustered system

A clustered system consists of one to four pairs of nodes. All configuration, monitoring, and service tasks are performed at the system level and the configuration settings are replicated across all node canisters in the clustered system. To facilitate these tasks, one or two management IP addresses are set for the system.

A process is provided 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 Storwize V7000 system configuration information.

System configuration backup: After backing up the system configuration, save the backup data on your hard disk (or at the least outside of the SAN). If you are unable to access the Storwize V7000, you do not have access to the backup data if it is on the SAN.

For the purposes of remote data mirroring, two or more clustered systems (IBM Storwize systems or SAN Volume Controller systems starting from Version 6.3) must form a partnership before creating relationships between mirrored volumes.

Important: IBM Storwize V7000 6.3 introduced the **`layer`** parameter. It can be changed by running **`chsystem`** using only the command-line interface (CLI). The default is the `storage` layer, and you must change it to `replication` if you need to set up a copy services relationship between the IBM Storwize Family and the SAN Volume Controller.

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

The system can be configured using either the IBM Storwize V7000 management software, the command-line interface (CLI), or through an application that uses the IBM Storwize V7000 CIMOM (IBM Tivoli Storage Productivity Center). IBM Systems Director also provides flexible server and storage management capability.

1.5.5 RAID

The Storwize V7000 setup contains several internal drive objects, but these drives cannot be directly added to storage pools. The drives need to be included in a Redundant Array of Independent Disks (RAID) to provide protection against the failure of individual drives.

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

IBM Storwize V7000 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 location and performance of each array. A sequence of drive failures and hot spare takeovers can leave an array unbalanced, that is, with members that do not match these goals. The system automatically rebalances such arrays when the appropriate drives are available.

The following RAID levels are available:

- RAID 0 (striping, no redundancy)

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

- RAID 1 (mirroring between two drives)

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

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

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

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

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

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

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

1.5.6 Managed disks

A managed disk (MDisk) is the unit of storage that IBM Storwize V7000 virtualizes. This unit could be a logical volume on an external storage array presented to the IBM Storwize V7000 or a RAID array consisting of internal drives. The IBM Storwize V7000 can then allocate these MDisks into various storage pools.An MDisk is not visible to a host system on the storage area network, because it is internal or zoned only to the IBM Storwize V7000 system.

An MDisk has four modes:

- Array

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

- Unmanaged

  Unmanaged MDisks are not being used by the system. This situation might occur when an MDisk is first imported into the system, for example.

- Managed

  Managed MDisks are assigned to a storage pool and provide extents so that volumes can use it.

- Image

  Image MDisks are assigned directly to a volume with a one-to-one mapping of extents between the MDisk and the volume. This situation is normally used when importing logical volumes into the clustered system that already have data on them, which ensures that the data is preserved as it is imported into the clustered system.

1.5.7 Quorum disks

A quorum disk is a managed disk (MDisk) that contains a reserved area for use exclusively by the system. In the Storwize V7000, 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 a managed disk (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, then to avoid the possibility of losing all of the quorum disks because of a failure of a single storage system, you should allocate the quorum disk on different storage systems. It is possible to manage the quorum disks by using the CLI.

1.5.8 Storage pools

A storage pool is a collection of MDisks (up to 4096) that are grouped together to provide capacity for volumes. All MDisks in the pool are split into extents with the same size. Volumes are then allocated out of the storage pool and are mapped to a host system.

Storwize V7000 object names: The names must begin with a letter, which cannot be numeric. The name can be a maximum of 63 characters. Valid characters are uppercase (A-Z), lowercase letters (a-z), digits (0 - 9), underscore (_), period (.), hyphen (-), and space. The names must not begin or end with a space.

MDisks can be added to a storage pool at any time to increase the capacity of the storage pool. MDisks can belong in only one storage pool and only MDisks in unmanaged mode can be added to the storage pool. When an MDisk is added to the storage pool, the mode changes from unmanaged to managed and vice versa when you remove it.

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

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

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

Table 1-5 Maximum volume capacity by extent size

Extent size	Maximum volume capacity for normal volumes (GB)
16	2048 (2 TB)
32	4096 (4 TB)
64	8192 (8 TB)
128	16384 (16 TB)
256	32768 (32 TB)
512	65536 (64 TB)
1024	131072 (128 TB)
2048	262144 (256 TB)
4096	528288 (512 TB)
8192	1056576 (1,024 TB)

Table 1-6 compares the maximum volume, MDisk, and system capacity for each extent size.

Table 1-6 Maximum volume, MDIsk, and system capacity for each extent size

Extent size (MB)	Maximum non thin-provisioned volume capacity in GB (and TB)	Maximum thin-provisioned volume capacity in GB	Maximum MDisk capacity in GB (and TB)	Total storage capacity manageable per system[a]
16	2048 (2 TB)	2000	2048 (2 TB)	64 TB
32	4096 (4 TB)	4000	4096 (4 TB)	128 TB
64	8192 (8 TB)	8000	8192 (8 TB)	256 TB
128	16,384 (16 TB)	16,000	16,384 (16 TB)	512 TB
256	32,768 (32 TB	32,000	32,768 (32 TB)	1 PB
512	65,536 (64 TB)	65,000	65,536 (64 TB)	2 PB
1024	131,072 (128 TB)	130,000	131,072 (128 TB)	4 PB
2048	262,144 (256 TB)	260,000	262,144 (256 TB)	8 PB
4096	262,144 (256 TB)	262,144	524,288 (512 TB)	16 PB
8192	262,144 (256 TB)	262,144	1,048,576 (1024 TB)	32 PB

a. The total capacity value assumes that all storage pools in the system use the same extent size.

Use the same extent size for all storage pools in a clustered system, which is a prerequisite for supporting volume migration between two storage pools. If the storage pool extent sizes are not the same, you must use volume mirroring to copy volumes between storage pools, as described in Chapter 7, “Storage pools” on page 179.

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

Default extent size: The GUI of IBM Storwize V7000 has a default extent size value of 256 MB when you define a new storage pool.

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

Single-tiered storage pool

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

- They should have the same hardware characteristics, for example, the same RAID type, RAID array size, disk type, and disk revolutions per minute (rpms).
- The disk subsystems providing the MDisks must have similar characteristics, for example, maximum input/output operations per second (IOPS), response time, cache, and throughput.
- Use MDisks of the same size, and ensure that the MDisks provide the same number of extents. If this configuration is not feasible, you need to 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 containing a mix of generic_hdd AND generic_ssd 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 V7000 Easy Tier function, as described in Chapter 10, “Advanced features for storage efficiency” on page 309.

Figure 1-12 shows these components.

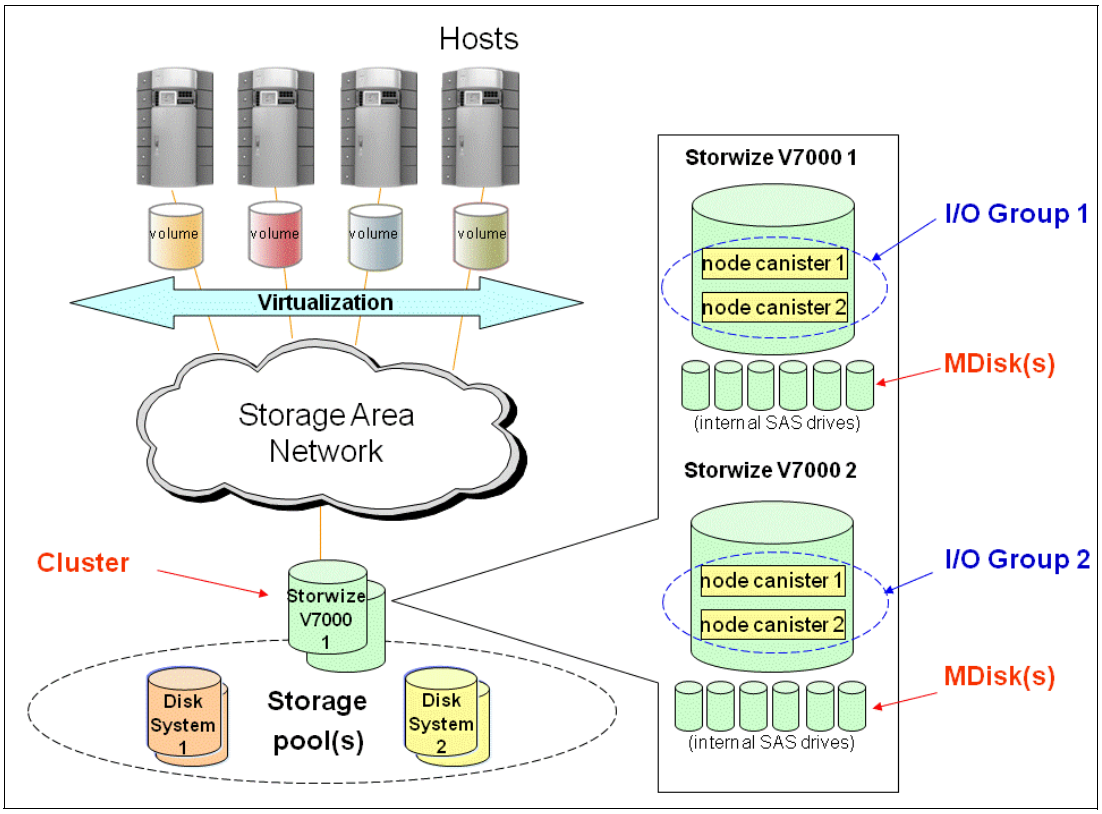

Figure 1-12 IBM Storwize V7000 virtualization components

1.5.9 Volumes

A volume is a logical disk that is presented to a host system by the clustered system. In our virtualized environment, the host system has a volume mapped to it by IBM Storwize V7000. IBM Storwize V7000 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 the extents without impacting the host system.

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

There are three types of volumes:

- Striped

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

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

Figure 1-13 shows how a striped volume is allocated, assuming 10 extents are required,

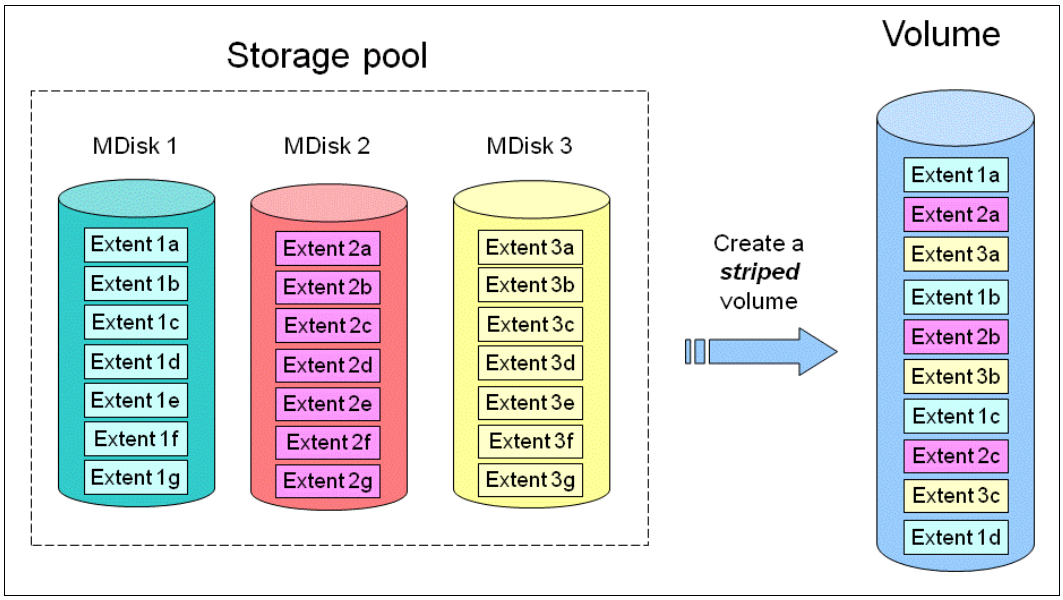

Figure 1-13 Striped volume

- Sequential

  A sequential volume is where the extents are allocated one after the other, from one MDisk to the next MDisk (Figure 1-14).

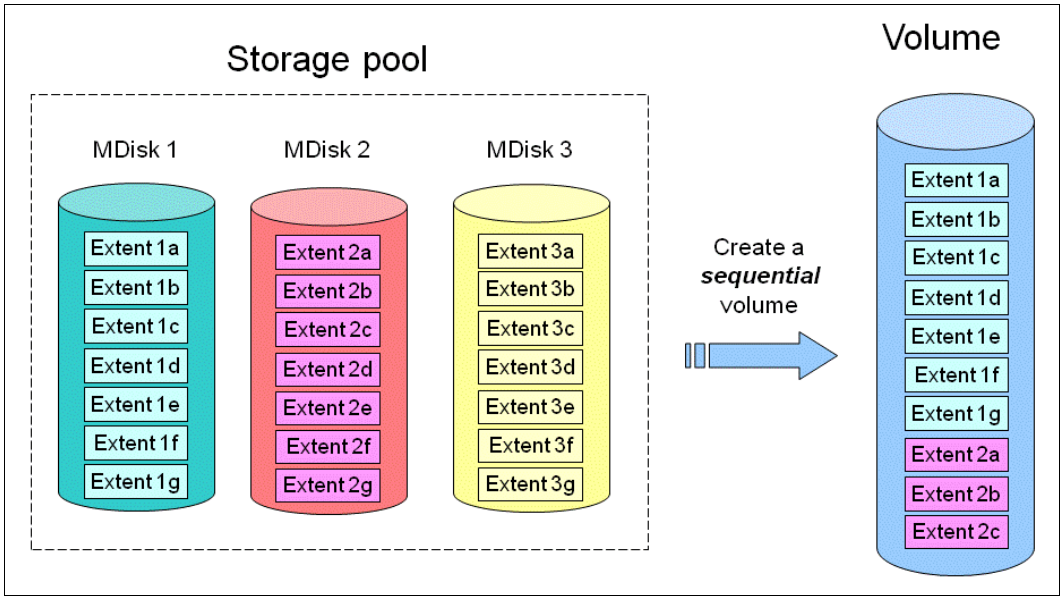

Figure 1-14 Sequential volume

- Image mode

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

When the image mode volume is created, a direct mapping is made between extents that are on the MDisk and the extents that are on the volume. The logical block address (LBA) x on the MDisk is the same as the LBA x on the volume, which ensures that the data on the MDisk is preserved as it is brought into the clustered system (Figure 1-15).

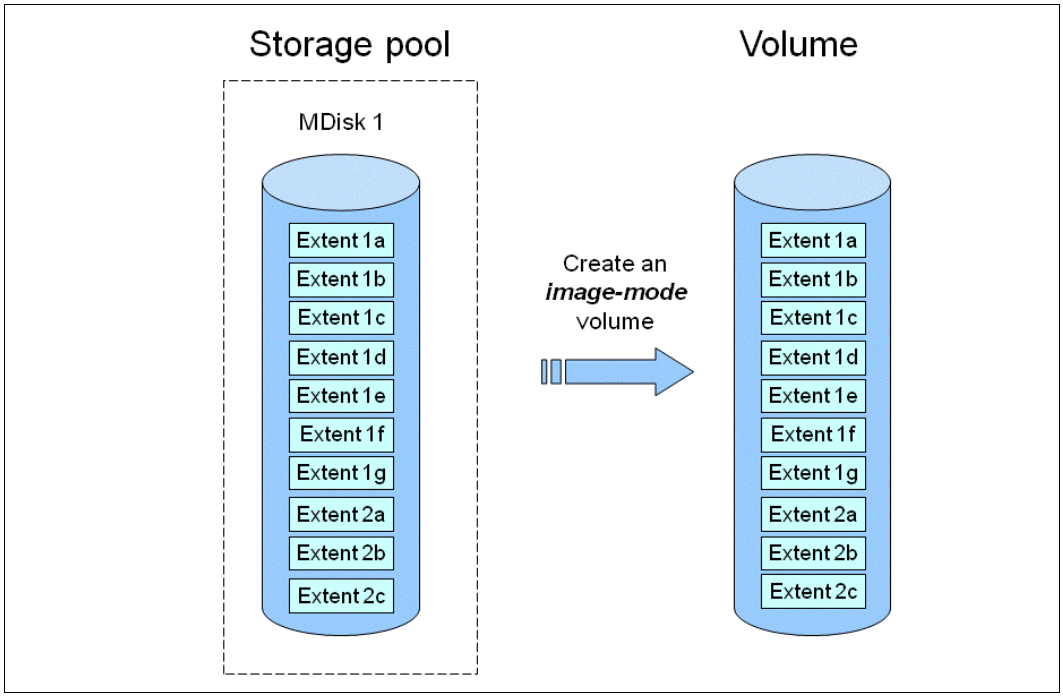

Figure 1-15 Image mode volume

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 add an MDisk containing data to a storage pool, any data on the MDisk is lost. Ensure that you create image mode volumes from MDisks that contain data before adding MDisks to the storage pools.

1.5.10 Thin-provisioned volumes

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

The real capacity determines the quantity of MDisk extents that are allocated for the volume. The virtual capacity is the capacity of the volume reported to IBM Storwize V7000 and to the host servers.

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

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

A thin-provisioned volume can be configured to *autoexpand*, which causes IBM Storwize V7000 to automatically expand the real capacity of a thin-provisioned volume as its real capacity is used. Autoexpand attempts to maintain a fixed amount of unused real capacity on the volume. This amount is known as the *contingency capacity*.

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

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

1.5.11 Mirrored volumes

IBM Storwize V7000 provides a function that is called *volume mirroring*, which enables a volume to have two physical copies. Each volume copy can belong to a different storage pool and can be on different physical storage systems, and provides a high-availability solution.

When a host system issues a write to a mirrored volume, IBM Storwize V7000 writes the data to both copies. When a host system issues a read to a mirrored volume, IBM Storwize V7000 places it into the primary copy. If one of the mirrored volume copies is temporarily unavailable, the IBM Storwize V7000 automatically uses the alternative copy without any outage for the host system. When the mirrored volume copy is repaired, IBM Storwize V7000 resynchronizes the data.

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

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

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

If needed, you can change the time-out value, `latency`, which prioritizes low host latency (default) or `redundancy`, which prioritizes redundancy (longer timeout).

Unmirrored volumes: An unmirrored volume can be migrated from one location to another by adding a second copy to the destination that you want, waiting for the two copies to synchronize, and then removing the original copy. This operation can be stopped at any time.

1.5.12 Easy Tier

IBM Easy Tier is a performance function that automatically migrates or moves extents off a volume to, or from, SSD storage to HDD storage. Easy Tier monitors the host I/O activity and latency on the extent of all volumes with the Easy Tier function turned on in a multitiered storage pool over a 24-hour period. It then creates an extent migration plan based on this activity and then dynamically moves high activity or hot extents to a higher disk tier within the storage pool. It also moves extent activity that has dropped off or cooled from the high tiered MDisk back to lower tiered MDisk.

The Easy Tier function may 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 a solid-state drive. By turning on the Easy Tier function for a single level storage pool and the Easy Tier function for the volumes within that pool, Easy Tier creates a migration report every 24 hours on the number of extents it would move if the pool was a multitiered pool. Easy Tier statistics measurement is enabled.

Using Easy Tier might make it more appropriate to use smaller storage pool extent sizes.

The usage statistics file can be offloaded from IBM Storwize V7000 nodes and then an IBM Storage Tier Advisor Tool (STAT) can be used to create a summary report. You can find the STAT tool at the following website:

`http://www.ibm.com/support/docview.wss?uid=ssg1S4000935`

Contact your IBM representative or IBM Business Partner for more information about the Storage Tier Advisor Tool. Easy Tier is described in more detail in Chapter 10, “Advanced features for storage efficiency” on page 309.

1.5.13 iSCSI

iSCSI is an alternative means of attaching hosts to the IBM Storwize V7000. All communications with back-end storage subsystems, and with other IBM Storwize V7000s, only occur through a Fibre Channel connection.

The iSCSI function is a *software function* that is provided by the IBM Storwize V7000 code, not hardware.

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

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

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

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

The concepts of names and addresses have been carefully separated in iSCSI:

- An *iSCSI name* is a location-independent, permanent identifier for an iSCSI node. An iSCSI node has one iSCSI name, which stays constant for the life of the node. The terms *initiator name* and *target name* also refer to an iSCSI name.
- An *iSCSI address* specifies not only the iSCSI name of an iSCSI node, but also a location of that node. The address consists of a host name or IP address, a TCP port number (for the target), and the iSCSI name of the node. An iSCSI node can have any number of addresses, which can change at any time, particularly if they are assigned by way of Dynamic Host Configuration Protocol (DHCP). An IBM Storwize V7000 node represents an iSCSI node and provides statically allocated IP addresses.

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

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

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

For IBM Storwize V7000, the IQN for its iSCSI target is specified as follows:

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

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

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

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

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

Changing names: Before changing system or node names for an IBM Storwize V7000 clustered system that has servers connected to it using SCSI, be aware that because the system and node name are part of the IQN for the IBM Storwize V7000, you can lose access to your data by changing these names. The IBM Storwize V7000 GUI shows a specific warning, but the CLI does *not*.

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

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

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

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

For further details about configuring iSCSI, see Chapter 4, “Host configuration” on page 105.

1.5.14 Real-time Compression

The IBM Real-time Compression™ solution addresses the challenges listed in the previous section because it was designed for primary storage. Implementing Real-time Compression in Storwize V7000 or SAN Volume Controller offers the following benefits:

- Compression for active primary data: IBM Real-time Compression can be used with active primary data. Therefore, it supports workloads that are not candidates for compression in other solutions. The solution supports online compression of existing data. It allows storage administrators to regain free disk space in an existing storage system without requiring administrators and users to clean up or archive data. This configuration significantly enhances the value of existing storage assets, and the benefits to the business are immediate. The capital expense of upgrading or expanding the storage system is delayed.
- Compression for replicated/mirrored data: Remote volume copies can be compressed in addition to the volumes at the primary storage tier. This process reduces storage requirements in Metro Mirror and Global Mirror destination volumes also.
- No changes to the existing environment are required: IBM Real-time Compression is part of the storage system. It was designed for transparency so that it can be implemented without changes to applications, hosts, networks, fabrics, or external storage systems. The solution is not apparent to hosts, so users and applications continue to work as-is. Compression occurs within the Storwize V7000 or SAN Volume Controller system itself.
- Overall savings in operational expenses: More data is stored in a rack space, so fewer storage expansion enclosures are required to store a data set. This reduced rack space has the following benefits:
  - Reduced power and cooling requirements: More data is stored in a system, therefore requiring less power and cooling per gigabyte or used capacity.
  - Reduced software licensing for additional functions in the system: More data stored per enclosure reduces the overall spending on licensing.
- Disk space savings are immediate: The space reduction occurs when the host writes the data. This process is unlike other compression solutions in which some or all of the reduction is realized only after a post-process compression batch job is run.

1.5.15 IP replication

Starting with Version 7.2, one of the most important functions in the Storwize family is IP replication. IP replication enables the use of lower-cost Ethernet connections for remote mirroring. The capability is available as a chargeable option on all Storwize family systems.

The function is transparent to servers and applications in the same way that traditional FC-based mirroring is. All remote mirroring modes (Metro Mirror, Global Mirror, and Global Mirror with changed volumes) are supported.

Configuration of the system is straightforward: Storwize family systems normally find each other on the network and can be selected from the GUI.

IP replication includes Bridgeworks SANSlide network optimization technology and is available at no additional charge. Remember, remote mirror is a chargeable option but the price does not change with IP replication. Existing remote mirror users have access to the function at no additional charge.

IP connections that are used for replication can have long latency (the time to transmit a signal from one end to the other), which can be caused by distance or by many “hops” between switches and other appliances in the network. Traditional replication solutions transmit data, wait for a response, and then transmit more data, which can result in network utilization as low as 20% (based on IBM measurements). And this scenario gets worse the longer the latency. Bridgeworks SANSlide technology that is integrated with IBM Storwize family requires no separate appliances and so no additional cost and no configuration steps. It uses artificial intelligence (AI) technology to transmit multiple data streams in parallel, adjusting automatically to changing network environments and workloads. SANSlide improves network bandwidth utilization up to 3x so customers can deploy a less costly network infrastructure or take advantage of faster data transfer to speed replication cycles, improve remote data currency, and enjoy faster recovery.

More about IP mirroring is covered in 11.4.2, “Native IP replication technology” on page 373.

1.6 Advanced copy services

IBM Storwize V7000 supports the following copy services:

- Synchronous remote copy
- Asynchronous remote copy
- FlashCopy

Starting with SAN Volume Controller V6.3, copy services functions are implemented within a single IBM Storwize V7000 or between multiple members of the IBM Storwize family. The Copy Services layer sits above and operates independently of the function or characteristics of the underlying disk subsystems used to provide storage resources to an IBM Storwize V7000. Figure 1-16 on page 28 shows an example of copy services with SAN Volume Controllers and IBM Storwize V7000 (they must be at Version 6.3 or later). The **`layer`** parameter of the clustered system properties is introduced in v6.3 and can be changed from `storage` (default) to `replication` if you need to make a new relationship between the IBM Storwize V7000 and SAN Volume Controller systems. You can change it by using only the **`chsystem`** command. In SAN Volume Controller systems, this parameter is fixed to replication and cannot be changed.

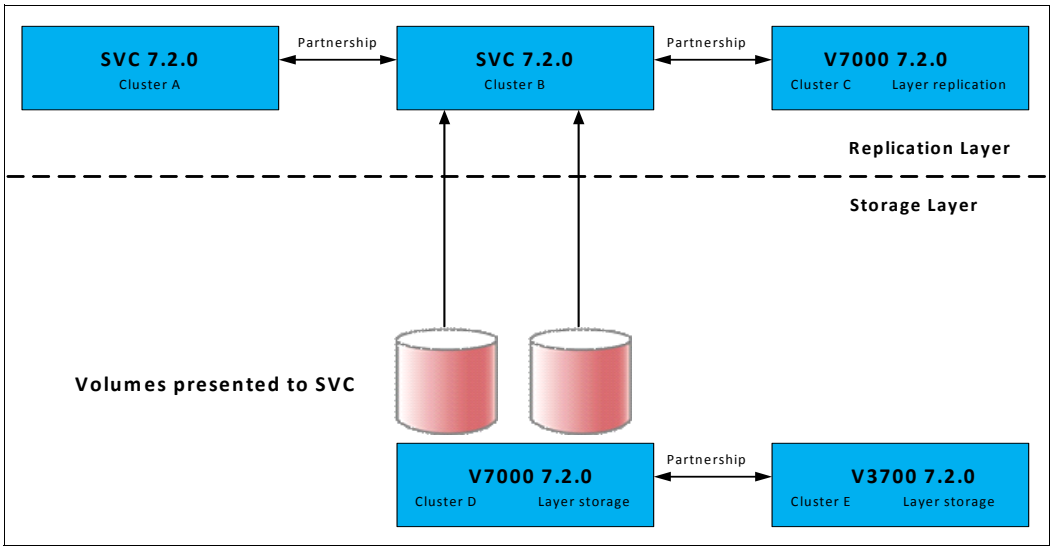

Figure 1-16 Example of SAN Volume Controller and Storwize V7000 copy services relationship

1.6.1 Synchronous or asynchronous remote copy

The general application of remote copy seeks to maintain two copies of data. Often, the two copies are separated by distance, but not always.

The remote copy can be maintained in one of two modes: synchronous or asynchronous.

With the IBM Storwize V7000, Metro Mirror and Global Mirror are the IBM branded terms for the functions that are synchronous remote copy and asynchronous remote copy.

Synchronous remote copy ensures that updates are committed at both the primary and the secondary before the application considers the updates complete; therefore, the secondary is fully up to date if it is needed in a failover. However, the application is fully exposed to the latency and bandwidth limitations of the communication link to the secondary. In a truly remote situation, this extra latency can have a significant adverse effect on application performance.

Special configuration guidelines exist for SAN fabrics that are used for data replication. You must consider the distance and available bandwidth of the intersite links (see Chapter 11, “Copy services” on page 327).

A function of Global Mirror designed for low bandwidth has been introduced in IBM Storwize V7000. It uses change volumes that are associated with the primary and secondary volumes. These volumes are used to record changes to the remote copy volume, the FlashCopy relationship that exists between the secondary volume and the change volume, and between the primary volume and the change volume. This function is called Global Mirror cycling mode. Figure 1-17 on page 29 shows an example of this function where you can see the relationship between volumes and change volumes.

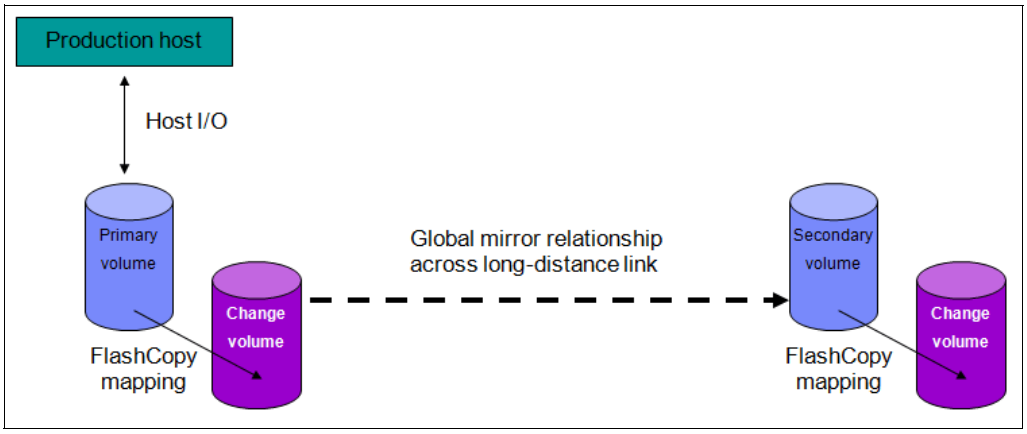

Figure 1-17 Global Mirror with change volumes

In asynchronous remote copy, the application acknowledges that the write is complete before the write is committed at the secondary. Hence, on a failover, certain updates (data) might be missing at the secondary. The application must have an external mechanism for recovering the missing updates, if possible. This mechanism can involve user intervention. Recovery on the secondary site involves starting the application on this recent “backup” and then rolling forward or backward to the most recent commit point.

IBM Support for automation is provided by IBM Tivoli Storage Productivity Center for Replication. The Tivoli documentation can also be accessed online at the IBM Tivoli Storage Productivity Center at the following website:

http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp

1.6.2 FlashCopy

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

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

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

IBM Storwize V7000 also permits multiple target volumes to be FlashCopied from the same source volume. This capability can be used to create images from separate points in time for the source volume, and to create multiple images from a source volume at a common point in time. Source and target volumes can be thin-provisioned volumes.

Reverse FlashCopy enables target volumes to become restore points for the source volume without breaking the FlashCopy relationship and without waiting for the original copy operation to complete. IBM Storwize V7000 supports multiple targets and thus multiple rollback points.

Most clients aim to integrate the FlashCopy feature for point in time copies and quick recovery of their applications and databases. IBM Support is provided by Tivoli Storage FlashCopy Manager, which is described at the following website:

http://www.ibm.com/software/tivoli/products/storage-flashcopy-mgr/

You can read a detailed description about the FlashCopy copy services in Chapter 11, “Copy services” on page 327.

1.6.3 Copy Services configuration limits

Table 1-7 lists the Copy Services configuration limits. For the most up-to-date list of these limits, see the following website:

http://www.ibm.com/support/docview.wss?uid=ssg1S1004511

Table 1-7 Copy Services configuration limits

Properties	Maximum number	Note
Remote copy (Metro Mirror and Global Mirror) relationships per system	4096	This can be any mix of Metro Mirror and Global Mirror relationships.
Remote copy relationships per consistency group	-	No limit is imposed beyond the remote copy relationships per system limit
Remote copy consistency groups per system	256	-
Total Metro Mirror and Global Mirror volume capacity per I/O group	1024 TB	This limit is the total capacity for all master and auxiliary volumes in the I/O group.
Total number of Global Mirror with Change Volumes relationships per system	256	-
FlashCopy mappings per system	4096	-
FlashCopy targets per source	256	-
FlashCopy mappings per consistency group	512	-
FlashCopy consistency groups per system	127	-
Total FlashCopy volume capacity per I/O group	1024 TB	-
Cascaded FlashCopy maps	256	A volume can be the source of up to 256 FlashCopy maps.

1.7 Management and support tools

The IBM Storwize V7000 system can be managed through the native management software that runs in the hardware.

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 either understand how to fix the problem yourself or allow the representative to fix it for you.

To use the IBM Assist On-site tool, the master console must be able to access the Internet. The following website provides further information about this tool:

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

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

You can also use security features to restrict access by the IBM service representative. Your IBM service representative can provide you with more detailed instructions for using the tool.

1.7.2 Event notifications

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

1.7.3 SNMP traps

Simple Network Management Protocol (SNMP) is a standard protocol for managing networks and exchanging messages. IBM Storwize V7000 can send SNMP messages that notify personnel about an event. You can use an SNMP manager to view the SNMP messages that IBM Storwize V7000 sends. You can use the management GUI or the IBM Storwize V7000 command-line interface to configure and modify your SNMP settings.

You can use the Management Information Base (MIB) file for SNMP to configure a network management program to receive SNMP messages that are sent by the IBM Storwize V7000. This file can be used with SNMP messages from all versions of the IBM Storwize V7000 Software.

1.7.4 Syslog messages

The syslog protocol is a standard protocol for forwarding log messages from a sender to a receiver on an IP network. The IP network can be either IPv4 or IPv6. IBM Storwize V7000 can send syslog messages that notify personnel about an event. IBM Storwize V7000 can transmit syslog messages in either expanded or concise format. You can use a syslog manager to view the syslog messages that IBM Storwize V7000 sends. IBM Storwize V7000 uses the User Datagram Protocol (UDP) to transmit the syslog message. You can use the management GUI or the IBM Storwize V7000 command-line interface to configure and modify your syslog settings.

1.7.5 Call Home email

The Call Home feature transmits operational and error-related data to you and IBM through a Simple Mail Transfer Protocol (SMTP) server connection in the form of an event notification email. When configured, this function alerts IBM service personnel about hardware failures and potentially serious configuration or environmental issues. You can use the Call Home function if you have a maintenance contract with IBM or if the Storwize V7000 is within the warranty period.

To send email, you must configure at least one SMTP server. You can specify as many as five additional SMTP servers for backup purposes. The SMTP server must accept the relaying of email from the IBM Storwize V7000 clustered system IP address. You can then use the management GUI or the IBM Storwize V7000 command-line interface to configure the email settings, including contact information and email recipients. Set the reply address to a valid email address. Send a test email to check that all connections and infrastructure are set up correctly. You can disable the Call Home function at any time using the management GUI or the IBM Storwize V7000 command-line interface.

1.8 Useful Storwize V7000 websites

See the following IBM Storwize V7000 web pages for more information:

- Support page:

  http://ibm.co/1nyBrTn
- IBM Storwize V7000 Unified and Storwize V7000 Disk Systems:

  http://www.ibm.com/systems/storage/news/center/storwize_v7000/index.html
- List of supported hardware:

  http://www.ibm.com/support/docview.wss?uid=ssg1S1004450
- Configuration Limits and Restrictions:

  http://www.ibm.com/support/docview.wss?uid=ssg1S1004511
- Documentation:

  http://ibm.co/1nxgxFF
- IBM Knowledge Center:

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

IBM Redbooks publications about IBM Storwize V7000 are available at the following website:

http://www.redbooks.ibm.com/cgi-bin/searchsite.cgi?query=v7000

1.8.1 IBM Storwize V7000 learning videos on YouTube

Several helpful videos that describe the IBM Storwize V7000 system are available on YouTube at the links that are listed in Table 1-8.

Table 1-8 Videos available on YouTube

Description of the video	Web address
IBM Storwize V7000 Storage Virtualization Terminology overview	http://www.youtube.com/watch?v=I2rzt3m2gPO
IBM Storwize V7000 Interface Tour	http://www.youtube.com/watch?v=FPbNRs9HacQ
IBM Storwize V7000 Volume Management	http://www.youtube.com/watch?v=YXeKqH8Sd9o
IBM Storwize V7000 migration	http://www.youtube.com/watch?v=dXxnUN6dk74
IBM Storwize V7000 Introduction to FlashCopy	http://www.youtube.com/watch?v=MXWgGWjBzG4
IBM Storwize V7000 Introduction to Remote Copy Part I: Creating Relationships and Consistency Groups	http://www.youtube.com/watch?v=koFYm--gnEc
IBM Storwize V7000 Introduction to Remote Copy Part II: Recovery and Return	http://www.youtube.com/watch?v=kMOFFS3NYas
IBM Storwize V7000 Rapid Application Storage Solution: Tivoli Storage FlashCopy Manager	http://www.youtube.com/watch?v=1zgEljRVDSc
IBM SAN Volume Controller and Storwize V7000 Performance Panel Sped-up! (HD)	http://www.youtube.com/watch?v=7noC71tLkWs
IBM Real-time Compression and Storwize V7000 and SAN Volume Controller demonstration	http://www.youtube.com/watch?v=rgKj75kn2JO
Storage efficiency with IBM Easy Tier	http://www.youtube.com/watch?v=aggudOKG9go

2

Initial configuration

This chapter describes the IBM Storwize V7000 initial configuration steps and these topics:

- Planning for IBM Storwize V7000 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. Use the following checklist of requirements to plan your installation:

- ☐ Install the hardware as described in *IBM Storwize V7000 Quick Installation Guide,* Form Number GC27-2290. The guide is at the following link:

  http://www.ibm.com/support/knowledgecenter/ST3FR7_7.2.0/com.ibm.storwize.v7000.720.doc/mlt_relatedinfo_224agr.html?lang=en

- ☐ For information about planning the IBM Storwize V7000 environment, go to this address:

  http://pic.dhe.ibm.com/infocenter/storwize/ic/index.jsp?topic=/com.ibm.storwize.v7000.doc/svc_installplan_22qgvs.html

- ☐ An appropriate 19-inch rack with minimum of 2U of space should be available. Space depends on the number of enclosures to install. Each enclosure measures 2U and one control enclosure with up to nine expansion enclosures constitutes an IBM Storwize V7000 system.
- ☐ There should be redundant power outlets in the rack for each of the two power cords included per enclosure. The number of outlets required depends on the number of enclosures to install. The power cords conform to the IEC320 C13 and C14 standards.
- ☐ A minimum of four Fibre Channel (FC) ports attached to the fabric are required, but it is a preferred practice to use eight 2-Gbps, 4-Gbps, or 8-Gbps Fibre Channel ports.

  FC ports: FC ports are required only if FC hosts or remote mirroring are used. You can use the Storwize V7000 with Ethernet-only cabling for iSCSI or FCoE hosts.

- ☐ You must have eight 2-Gbps, 4-Gbps, or 8-Gbps compatible Fibre Channel cable drops.
- ☐ You must have a minimum of two Ethernet ports on the LAN, with four preferred for additional configuration access redundancy or iSCSI host access. If you have a V7000 Model 314 or 324 you can also use the four 10 GB Ethernet ports for iSCSI or FCoE connections.
- ☐ You must have a minimum of two Ethernet cable drops, with four preferred for additional configuration access redundancy or iSCSI host access. Ethernet port one on each node canister must be connected to the LAN with port two being optional.

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

- ☐ Verify that the default IP address configured on Ethernet port 1 on each of the node canisters (192.168.70.121 on node one and 192.168.70.122 on node 2) do not conflict with existing IP addresses on the LAN. The default mask used with these IP addresses is 255.255.255.0 and the default gateway address used is 192.168.70.1.

- ☐ You must have a minimum of three IPv4 or IPv6 IP addresses for system configuration. One is for the clustered system and is what the administrator uses for management, and one for each node canister for service access as needed.

  IP addresses: Use a fourth IP address for backup configuration access. This additional IP address allows a second system IP address to be configured on port 2 of either node canister, which the storage administrator can also use for management of the IBM Storwize V7000 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 V7000. For the models with 10 GB iSCSI ports, you need an additional four addresses.

- ☐ Two 1-meter, 3-meter, or 6-meter SAS cables per expansion enclosure are required. The length of the cables depends on the physical rack location they are installed in relative to the control enclosure or other expansion enclosures.

  Locate the control enclosure so that four enclosures can be located above it and five enclosures below, as shown in Figure 2-1 (which shows the full racking and cabling scheme).

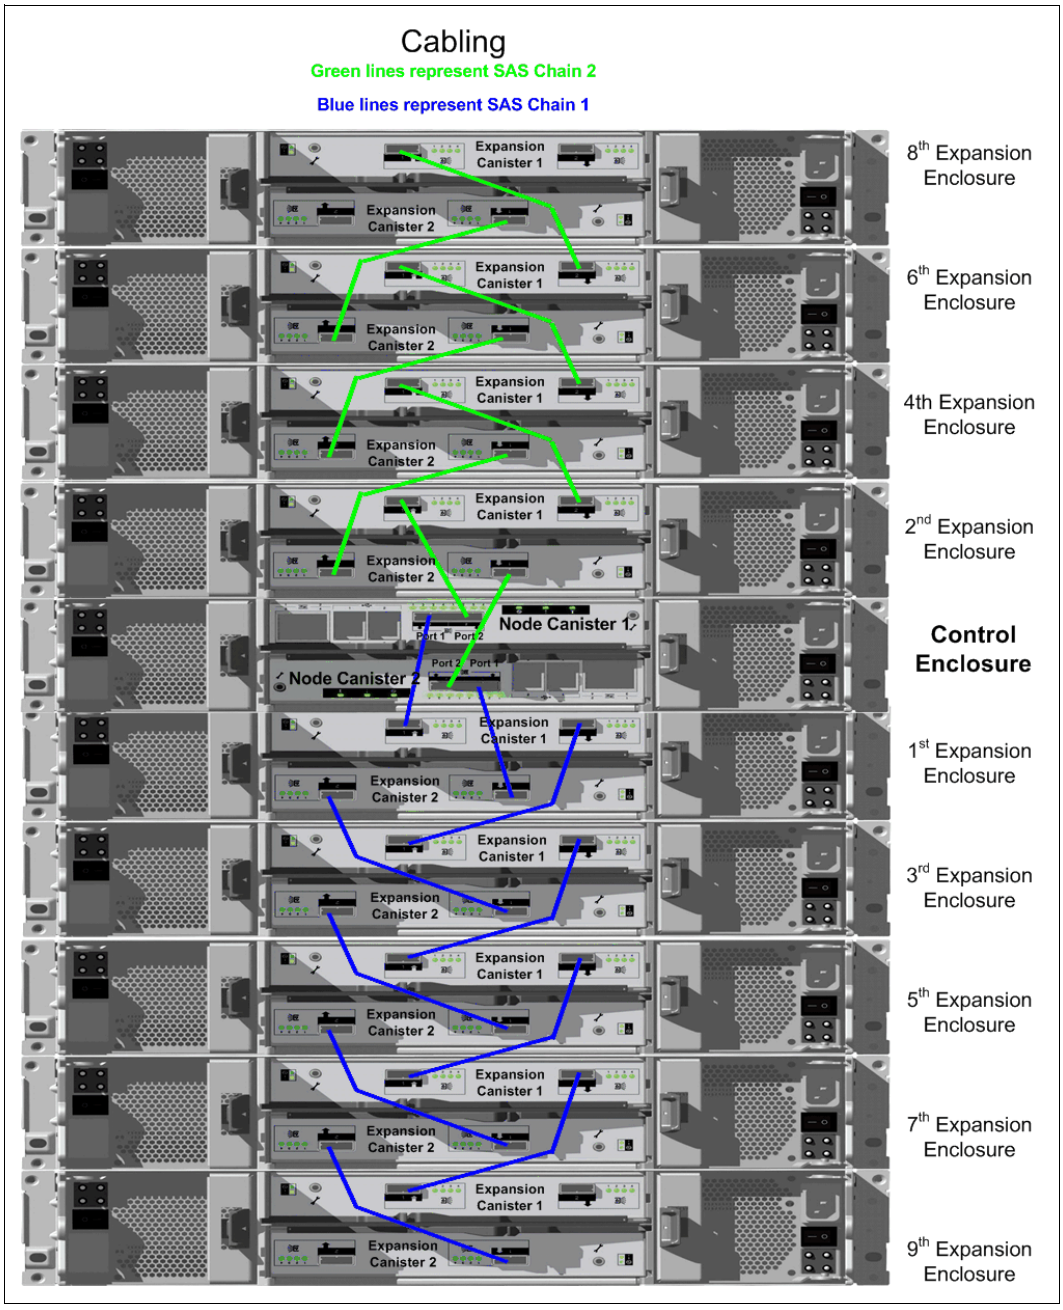

Figure 2-1 Racking and cabling diagram

Disk drives: The disk drives included with the control enclosure, model 2076-124/324 or 2076-112/312, are part of SAS chain number two. Therefore, only four additional expansion enclosures can be connected to this chain. SAS chain number one supports the addition of up to five expansion enclosures. The first expansion enclosure should be connected to SAS chain number one so both chains are used and the full bandwidth of the system is used.

2.2 SAN configuration planning

The recommended SAN configuration is composed of a minimum of two fabrics with all host ports, and any ports on external storage systems to be virtualized by IBM Storwize V7000. The IBM Storwize V7000 ports themselves are evenly split between the two fabrics to provide redundancy in the event one of the fabrics goes offline (either planned or unplanned).

After the IBM Storwize V7000, hosts, and optional external storage systems are connected to the SAN fabrics, zoning needs to be implemented.

In each fabric, create a zone with just the four IBM Storwize V7000 WWPNs, two from each node canister. If there is an external storage system to be virtualized, then in each fabric, create a zone with the four IBM Storwize V7000 WWPNs, two from each node canister, along with up to a maximum of eight WWPNs from the external storage system. Assuming every host has a Fibre Channel connection to each fabric, then in each fabric, create a zone with the host WWPN and one WWPN from each node canister in the IBM Storwize V7000 system.

Maximum ports or WWPNs: IBM Storwize V7000 supports a maximum of 16 ports or WWPNs from a given external storage system that is virtualized.

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

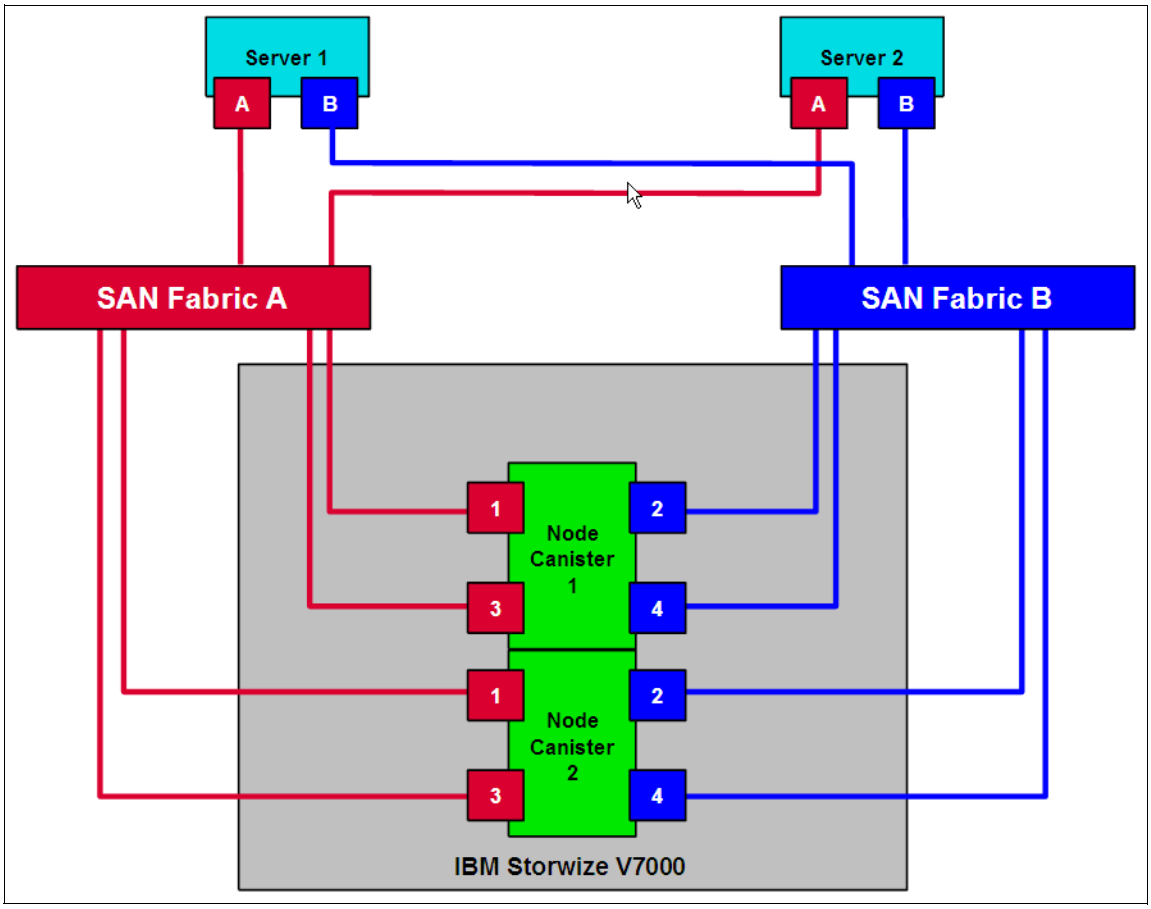

Figure 2-2 SAN cabling and zoning diagram

Create a host and IBM Storwize V7000 zone for each server to which volumes are mapped from the clustered system, as in this example:

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

Verify that the SAN switches or directors to which the IBM Storwize V7000 connects meet the following requirements:

- Switches or directors are at the firmware levels supported by the IBM Storwize V7000.
- IBM Storwize V7000 port login maximums listed in restriction document must not be exceeded.

Connectivity issues: If you have any connectivity issues between IBM Storwize V7000 ports and Brocade SAN Switches or Directors at 8 Gbps, see `http://www-01.ibm.com/support/docview.wss?rs=591&uid=ssg1S1003699` for the correct setting of the fillword port config parameter in the Brocade operating system.

See the following web page for more information:

`http://www.ibm.com/storage/support/Storwize/V7000`

2.3 LAN configuration planning

Two Ethernet ports per node canister are available for connection to the LAN on an IBM Storwize V7000 system. Port 1 on each node canister must be connected to the LAN to provide access for the management of the clustered system and service access. Optionally, port 2 on each node canister can be connected to the LAN to provide redundant management access. In addition, all four ports can also be configured for access to IBM Storwize V7000 volumes by iSCSI hosts.

Configuring IP addresses: No issue exists with configuring multiple IPv4 or IPv6 addresses on a given Ethernet port or using the same Ethernet port for management and iSCSI access. However, you cannot use the same IP address for both management and iSCSI host use.

Figure 2-3 shows a possible IP configuration of Ethernet ports on the IBM Storwize V7000 system.

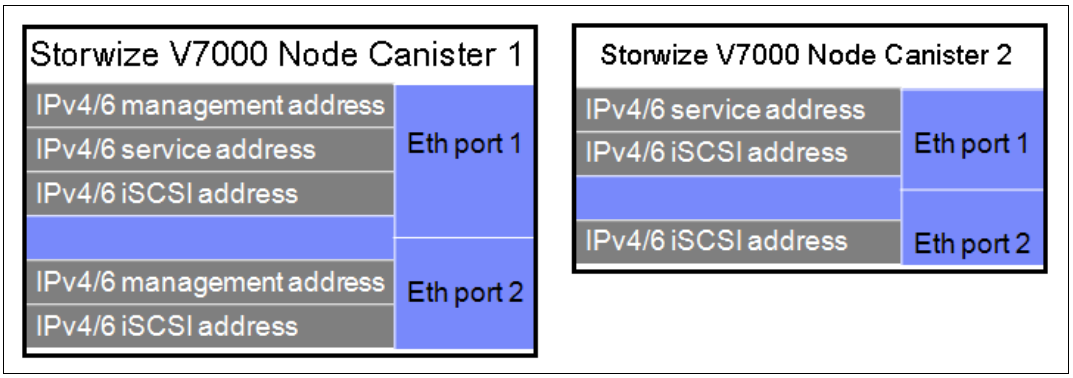

Figure 2-3 Example of IP address configuration options per node canister

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

2.3.1 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 V7000 system.

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

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

Figure 2-4 shows a logical view of the Ethernet ports available for configuration of the one or two management IP addresses. These IP addresses are for the clustered system and therefore associated with only one node, which is then considered the configuration node.

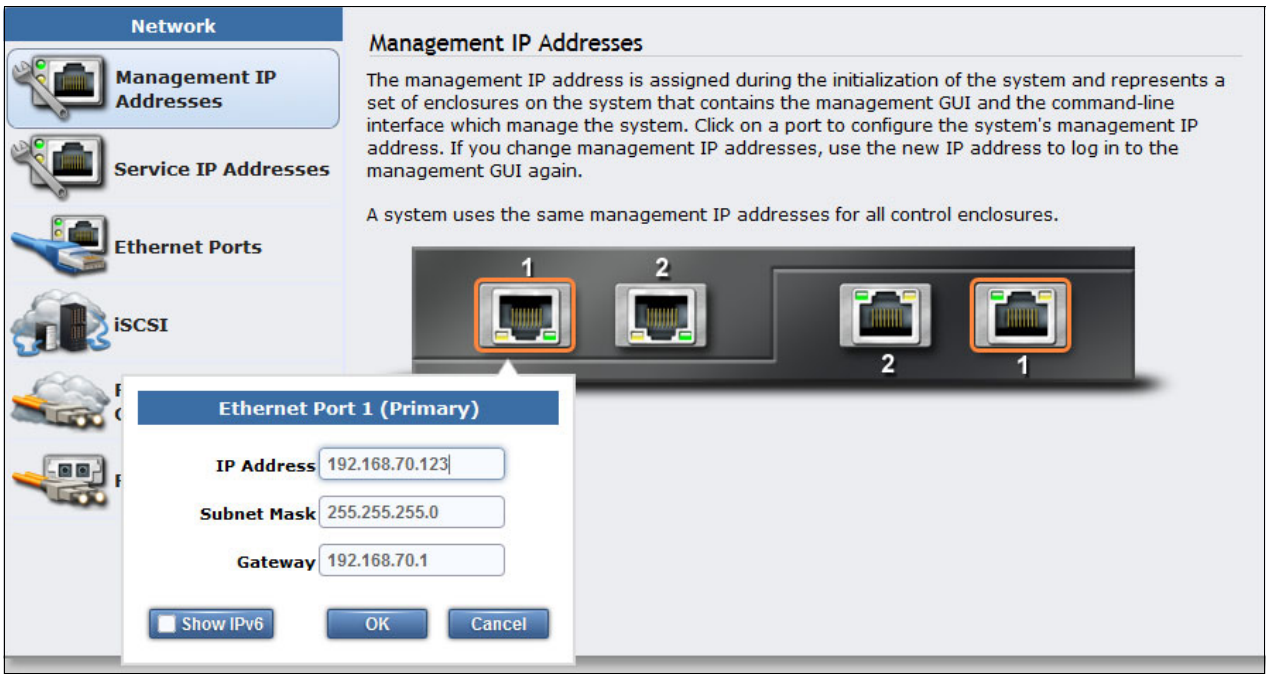

Figure 2-4 Ethernet ports available for configuration

2.3.2 Service IP address considerations

The Ethernet port 1 on each node canister is used for system management, and also used 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.

By default, the service IP address on node canister 1 is 192.168.70.121 and 192.168.70.122 on node canister 2. The default mask is 255.255.255.0 and the default gateway address is 192.168.70.1.

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

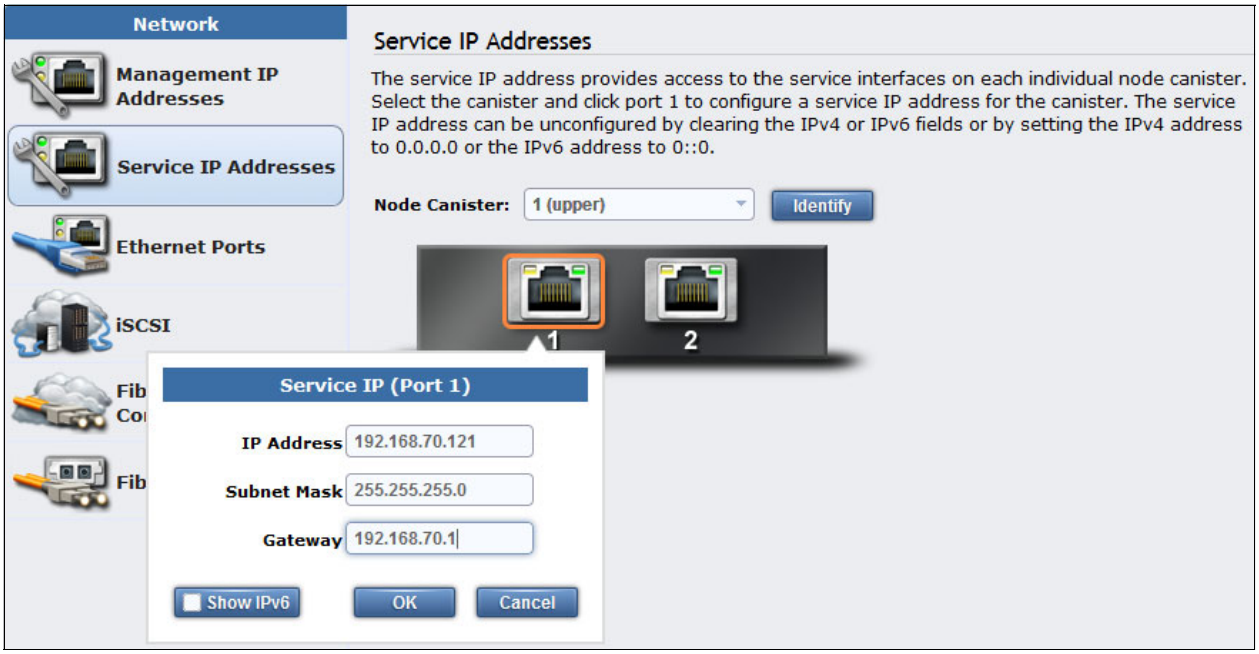

Figure 2-5 Service IP addresses available for configuration

2.4 Host configuration planning

Hosts should have two Fibre Channel connections for redundancy, but IBM Storwize V7000 supports hosts with a single host bus adapter (HBA) port connection. However, if that HBA or its link to the SAN fabric or the fabric itself fails, then the host loses access to its volumes. Even with just a single connection to the SAN, the host has multiple paths to the IBM Storwize V7000 volumes, because that single connection must be zoned with at least one Fibre Channel port per node. Therefore, a multipath driver is required.

SAN Boot with IBM Storwize V7000 is supported. The requirements are listed on the IBM Storwize V7000 support matrix; configuration instructions are provided in the *IBM Storwize V7000 Host Attachment Guide*, which is at the following address:

`http://ibm.co/1pvC5Up`

Verify that the hosts that access volumes from the IBM Storwize V7000 meet the requirements listed at the following website:

`http://www.ibm.com/storage/support/storwize/v7000`

Consider the following items:

- Host operating systems are at the levels supported by the IBM Storwize V7000.
- HBA BIOS, device drivers, and firmware, along with the multipathing drivers, are at the levels supported by IBM Storwize V7000.
- If boot from SAN is required, ensure that it is supported for the operating systems to be deployed.
- If host clustering is required, ensure that it is supported for the operating systems to be deployed.

For more details, see Chapter 4, “Host configuration” on page 105.

2.5 Miscellaneous configuration planning

During the initial setup of the IBM Storwize V7000 system, the installation wizard asks for various information that you should have available during the installation process. If you do not have this information or choose not to configure some of the items at this time, you can configure them later through the GUI.

Gather the information in the following checklist before you do the initial setup. The date and time can be manually entered, but to keep the clock synchronized, use a network time protocol (NTP) service.

- ☐ Current LAN NTP server IP address used for synchronization of devices.
- ☐ Number of external physical enclosures that are virtualized under IBM Storwize V7000.

  If external storage systems are used by IBM Storwize V7000, then a license must be purchased to cover the number of enclosures to be virtualized.
- ☐ Total number of physical enclosures, including both the internal and external virtualized enclosures for the Global Mirror and Metro Mirror features. You must have enough Global Mirror and Metro Mirror licenses for all the enclosures attached to the system, regardless of the amount of Global Mirror and Metro Mirror capacity you intend to use.
- ☐ Number of physical enclosures that will use compression under the IBM Storwize V7000.

The Easy Tier function is included with the IBM Storwize V7000 system and is not a licensed feature. If the system has solid-state drives, and this capability is used to optimize the usage of the SSD drives, then this function is enabled.

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 V7000 system for IBM service (for example, Building 22, first floor).
- ☐ SMTP or email server address to direct alerts to and from the IBM Storwize V7000.

  For the Call Home service to work, the IBM Storwize V7000 system must have access to an SMTP server on the LAN that can forward emails to the default IBM service address callhome0@de.ibm.com.

- ☐ Email address of local administrators needing to be notified of alerts.
- ☐ IP address of SNMP server to which alerts can be directed (for example, operations or the help desk).

After the IBM Storwize V7000 initial configuration, you might want to add extra users who can manage the system. You can create as many users as you need, but currently only five roles generally are configured for users:

- Security Admin

  In this role, the user can do any function on the IBM Storwize V7000.

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

- Administrator

  In this role, the user can do any function on the IBM Storwize V7000 system except create users.

- Copy Operator

  In this role, the user can view anything in the system, but the user can configure and manage only copy functions that include the replication functions of Metro and Global Mirror and the FlashCopy capabilities.

- Monitor

  In this role, the user can view object and system configuration, but cannot configure, manage, or modify any system resource.

- Service

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

2.6 System management

The graphical user interface (GUI) is used to configure, manage, and troubleshoot the IBM Storwize V7000 system. It is used primarily to configure RAID arrays and logical drives, assign logical drives to hosts, replace and rebuild failed disk drives, and expand logical drives.

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

Finally, it offers advanced functions, such as FlashCopy, Volume Mirroring, and Remote Mirroring. A command-line interface (CLI) for the IBM Storwize V7000 system is available too.

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

2.6.1 Graphical user interface (GUI)

A web browser is used for GUI access. You must use a supported web browser to access the management GUI. For a list of supported web browsers, see the *V7.2.x Supported Hardware List, Device Driver, Firmware and Recommended Software Levels for IBM Storwize V7000* document at the following website:

`http://www.ibm.com/support/docview.wss?uid=ssg1S1004450#_MGRGUI`

Open your web browser on a workstation and point it to the IBM Storwize V7000 system management address (this address is defined in 2.7, “First-time setup of IBM Storwize V7000” on page 46).

After the initial configuration described in 2.8, “Initial configuration” on page 54 is completed, the IBM Storwize V7000 Welcome window opens (Figure 2-6).

Figure 2-6 IBM Storwize V7000 Welcome window

2.6.2 Command-line interface (CLI)

The CLI is a flexible tool for system management and uses the SSH protocol. A public/private SSH key pair is optional for SSH access. Appendix A, “CLI setup and SAN Boot” on page 497 describes how to set up SSH access for Windows or UNIX systems. The storage system can be managed by using the CLI (Example 2-1 on page 46).

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

```
IBM_2076:ITSO-Storwize-V7000-2:admin>lsuser
id name      password ssh_key remote usergrp_id usergrp_name
0  superuser yes      no      no     0          SecurityAdmin
1  hartmut   yes      yes     no     0          SecurityAdmin
2  libor     yes      no      no     2          CopyOperator
3  matus     yes      no      no     3          Service
4  support   yes      no      no     3          Administrator

IBM_2076:ITSO-Storwize-V7000-2:admin>
```

Commands: Starting with IBM Storwize V7000 6.3, the **svctask** and **svcinfo** command prefixes are no longer necessary when issuing a command. If you have existing scripts that use those prefixes, they continue to function. You do not need to change the scripts.

Another important change is that you may use the **lssystem** command instead of **svcinfo lscluster**, or you may use **chsytem** instead of **svctask chcluster**. The old commands are maintained for compatibility purposes with any previously written scripts.

Do the initial IBM Storwize V7000 system setup before you use the graphical tools we describe in 2.7, “First-time setup of IBM Storwize V7000” on page 46 (the next section).

2.7 First-time setup of IBM Storwize V7000

IBM Storwize V7000 uses an initial setup that is contained within a USB key. The USB key is delivered with each storage system and has the initialization application, which is the `InitTool.exe` file. A system management IP address, the subnet mask, and the network gateway address are required. The initialization application creates a configuration file on the USB key.

The IBM Storwize V7000 starts the initial setup as soon as you plug in the USB key with the newly created file to the storage system.

USB key: If you are unable to find the official USB key that is supplied with IBM Storwize V7000, you can use any USB key that you have, and then download and copy the package with the `InitTool.exe` application, from the IBM Storwize V7000 Support website at the following address, onto your USB stick:

`http://www.ibm.com/storage/support/Storwize/V7000`

The USB key contains the `InitTool.exe` file (Figure 2-7).

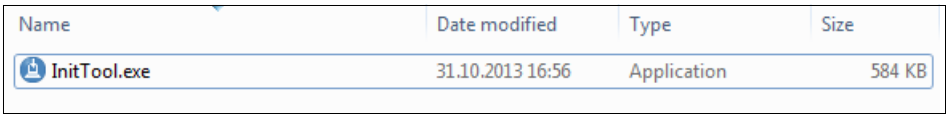

Figure 2-7 InitTool program in the USB key

To complete the initial setup using the USB key, follow these steps:

1. Plug the USB key into an Windows system and start the initialization tool. If the system is configured to autorun USB keys, the initialization tool starts automatically; otherwise, open My Computer and double-click the `InitTool.exe` file. The opening window of the tool is shown in Figure 2-8. After the tool is started, select **Initialize a new system using the USB Key** and click **Next**.

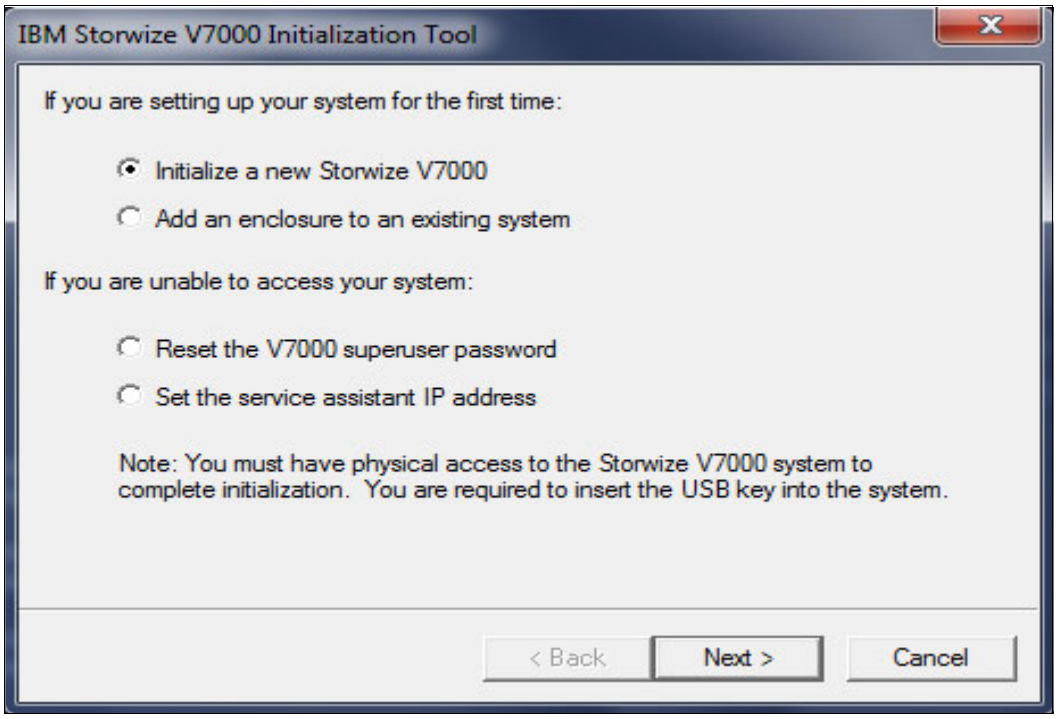

Figure 2-8 Initialization application: Welcome window

2. Enter the IPV4 or IPV6 address, subnet mask, and network gateway address, and then click **Next** (Figure 2-9).

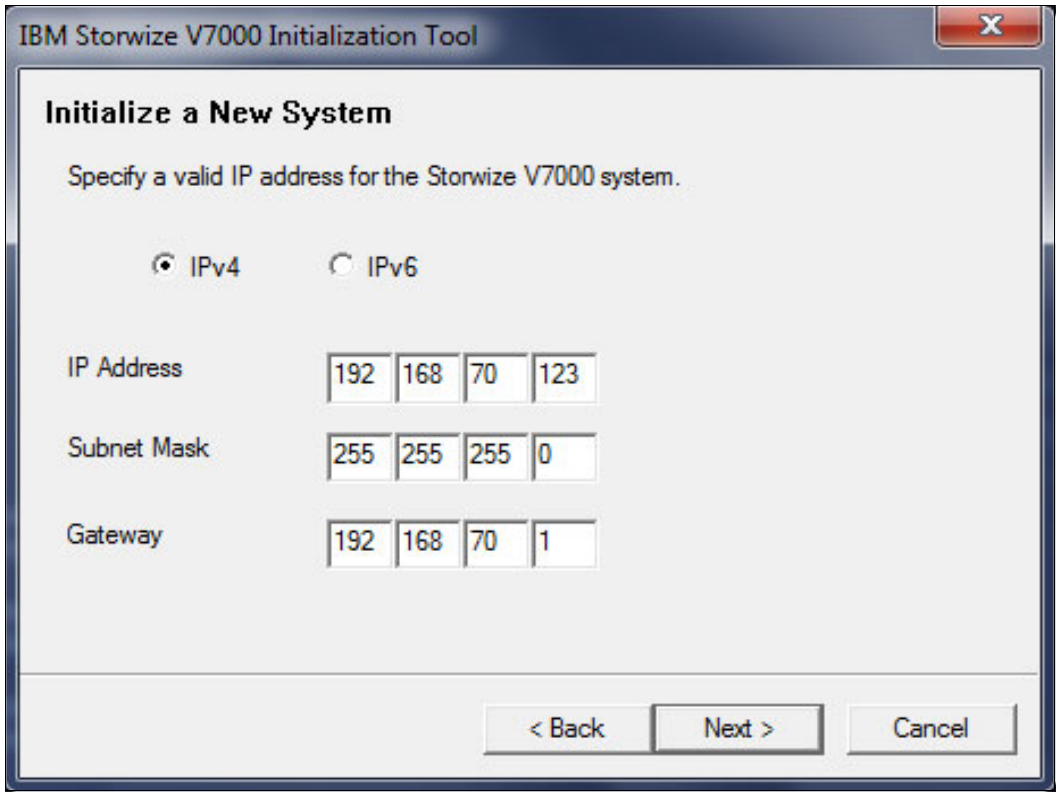

Figure 2-9 Initialization application: Network configuration

3. Click **Finish** (Figure 2-10).

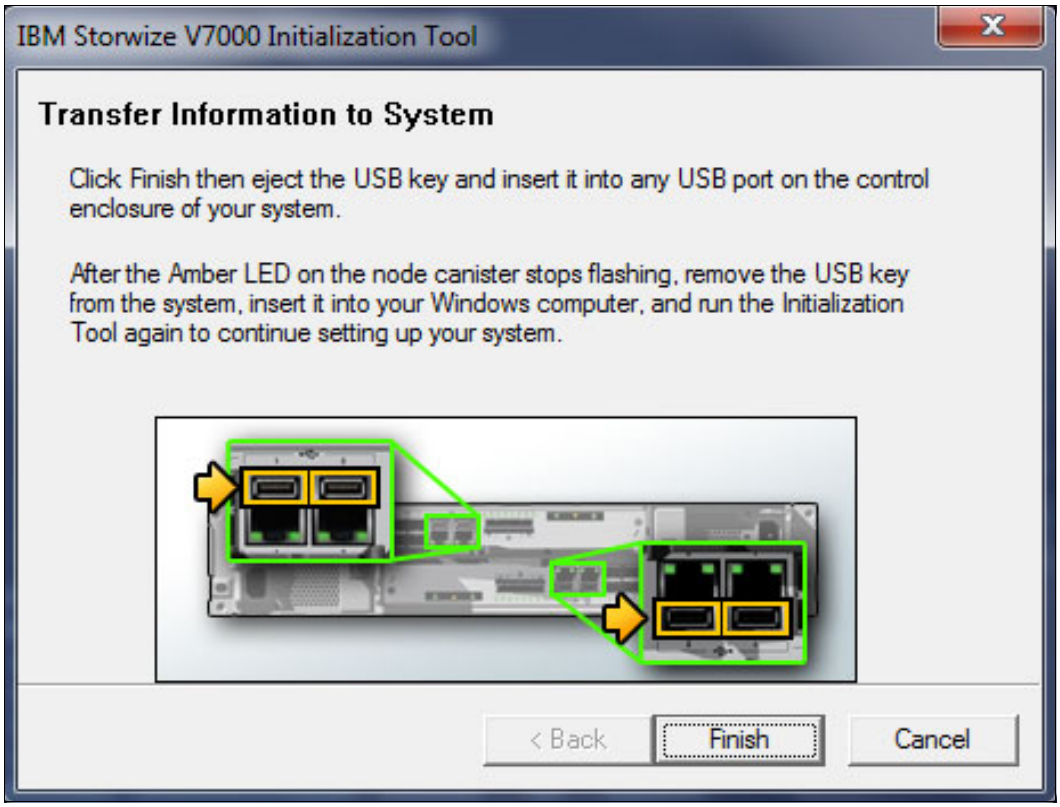

Figure 2-10 Initialization application: Finalizing the initial setup

The application creates a `satask.txt` file on the USB key (Figure 2-11).

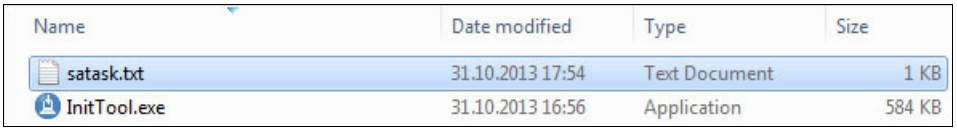

Figure 2-11 The satask.txt file created during the InitTool procedure

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

```
satask mkcluster -clusterip 192.168.70.123 -gw 192.168.70.1 -mask 255.255.255.0
```

4. Unplug the USB key from your Windows system and plug it into one of the IBM Storwize V7000 USB connector slots. The storage system will automatically do these steps:
   a. Detect the key.
   b. Read the `satask.txt` file and run its contents.
   c. Delete the `satask.txt` file.
   d. Write the response to a new file called `satask_result.html` (Figure 2-12).

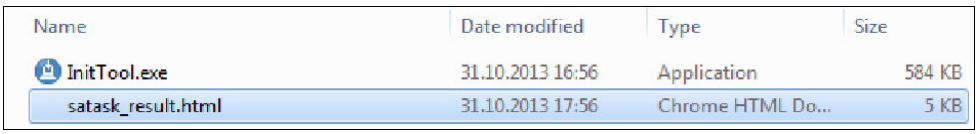

Figure 2-12 The satask_result.html file created during the InitTool procedure

Clustered system creation: While the clustered system is being created, the amber fault LED on the node canister flashes. When the amber fault LED stops flashing, remove the USB key from IBM Storwize V7000 and insert it in your system to check the results.

After this task completes successfully, the initial setup is done. The IBM Storwize V7000 is available for further configuration changes using the newly defined configuration address.

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

Each IBM Storwize V7000 clustered system has one or two system IP addresses. If the configuration node fails, the system IP addresses are transferred to another node in the same clustered system.

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

A sample of the contents of the `satask_result.html` file is shown in Example 2-2.

Example 2-2 satask_result.html file contents

```
Service Command Results
satask mkcluster -clusterip 192.168.70.123 -gw 192.168.70.1 -mask 255.255.255.0

Tue Oct 31 17:56:57 PDT 2013



System Status
sainfo lsservicenodes

panel_name cluster_id       cluster_name          node_id node_name relation
node_status error_data
78G01LB-1 00000200A18007F0 Cluster_192.168.70.123 1       node1     local
Active

sainfo lsservicestatus

panel_name 78G01LB-1
cluster_id 00000200A18007F0
cluster_name ITSO V7000
cluster_status Active
cluster_ip_count 2
cluster_port 1
cluster_ip 192.168.70.123
cluster_gw 192.168.70.1
cluster_mask 255.255.255.0
cluster_ip_6
cluster_gw_6
cluster_prefix_6
cluster_port 2
cluster_ip
cluster_gw
```

```
cluster_mask
cluster_ip_6
cluster_gw_6
cluster_prefix_6
node_id 1
node_name node1
node_status Active
config_node Yes
hardware 100
service_IP_address
service_gateway
service_subnet_mask
service_IP_address_6
service_gateway_6
service_prefix_6
node_sw_version 7.2.0.0
node_sw_build 85.6.1309091200
cluster_sw_build 85.6.1309091200
node_error_count 1
error_code 835
error_data
error_code
error_data
error_code
error_data
error_code
error_data
error_code
error_data
fc_ports 4
port_id 1
port_status Active
port_speed 2Gb
port_WWPN 50050768021005a8
SFP_type Short-wave
port_id 2
port_status Active
port_speed 4Gb
port_WWPN 50050768022005a8
SFP_type Short-wave
port_id 3
port_status Inactive
port_speed N/A
port_WWPN 50050768023005a8
SFP_type Short-wave
port_id 4
port_status Inactive
port_speed N/A
port_WWPN 50050768024005a8
SFP_type Short-wave
ethernet_ports 2
ethernet_port_id 1
port_status Link Online
port_speed 1Gb/s - Full
MAC e4:1f:13:74:09:fb
```

```
ethernet_port_id 2
port_status Not Configured
port_speed
MAC e4:1f:13:74:09:fa
ethernet_port_id 3
port_status
port_speed
MAC
ethernet_port_id 4
port_status
port_speed
MAC
product_mtm 2076-124
product_serial 78G01LB
time_to_charge 0
battery_charging 100
dump_name 78G01LB-1
node_WWNN
disk_WWNN_suffix
panel_WWNN_suffix
UPS_serial_number
UPS_status
enclosure_WWNN_1 50050768020005a8
enclosure_WWNN_2 50050768020005a9
node_part_identity 11S85Y5849YG50CG09827P
node_FRU_part 85Y5899
enclosure_identity 11S85Y5963YG50CGOAN7A0
PSU_count 2
PSU_id 1
PSU_status active
PSU_id 2
PSU_status active
Battery_count 2
Battery_id 1
Battery_status active
Battery_id 2
Battery_status active
node_location_copy 1
node_product_mtm_copy 2076-124
node_product_serial_copy 78G01LB
node_WWNN_1_copy 50050768020005a8
node_WWNN_2_copy 50050768020005a9
latest_cluster_id 200a0a006b2
next_cluster_id 200a0c006b2
console_IP 192.168.70.123:443
has_nas_key no

sainfo lsservicerecommendation

service_action
Follow troubleshooting procedures to fix hardware.

sainfo lshardware

panel_name 78G01LB-1
```

```
node_id 1
node_name node1
node_status Active
hardware 100
actual_different no
actual_valid yes
memory_configured 8
memory_actual 8
memory_valid yes
cpu_count 1
cpu_socket 1
cpu_configured 4 core Intel(R) Xeon(R) CPU C3539 @ 2.13GHz
cpu_actual 4 core Intel(R) Xeon(R) CPU C3539 @ 2.13GHz
cpu_valid yes
cpu_socket
cpu_configured
cpu_actual
cpu_valid
adapter_count 5
adapter_location 0
adapter_configured Four port 8Gb/s FC adapter
adapter_actual Four port 8Gb/s FC adapter
adapter_valid yes
adapter_location 0
adapter_configured 1Gb/s Ethernet adapter
adapter_actual 1Gb/s Ethernet adapter
adapter_valid yes
adapter_location 0
adapter_configured 1Gb/s Ethernet adapter
adapter_actual 1Gb/s Ethernet adapter
adapter_valid yes
adapter_location 0
adapter_configured High Speed SAS adapter
adapter_actual High Speed SAS adapter
adapter_valid yes
adapter_location 0
adapter_configured Midplane bus adapter
adapter_actual Midplane bus adapter
adapter_valid yes
adapter_location
adapter_configured
adapter_actual
adapter_valid
```

2.8 Initial configuration

To complete the initial configuration, these tasks are involved:

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

Complete the following steps:

1. Start the configuration wizard using a web browser on a workstation and point it to the system management IP address defined in 2.7, “First-time setup of IBM Storwize V7000” on page 46. Type in the default superuser password <passw0rd> (where 0 = zero), as shown in Figure 2-13.

Figure 2-13 Setup wizard: Login window

2. Create a new password (change the default password). See Figure 2-14.

Figure 2-14 Default password change

3. System Setup starts with the Welcome window, as shown in Figure 2-15. Click **Next**.

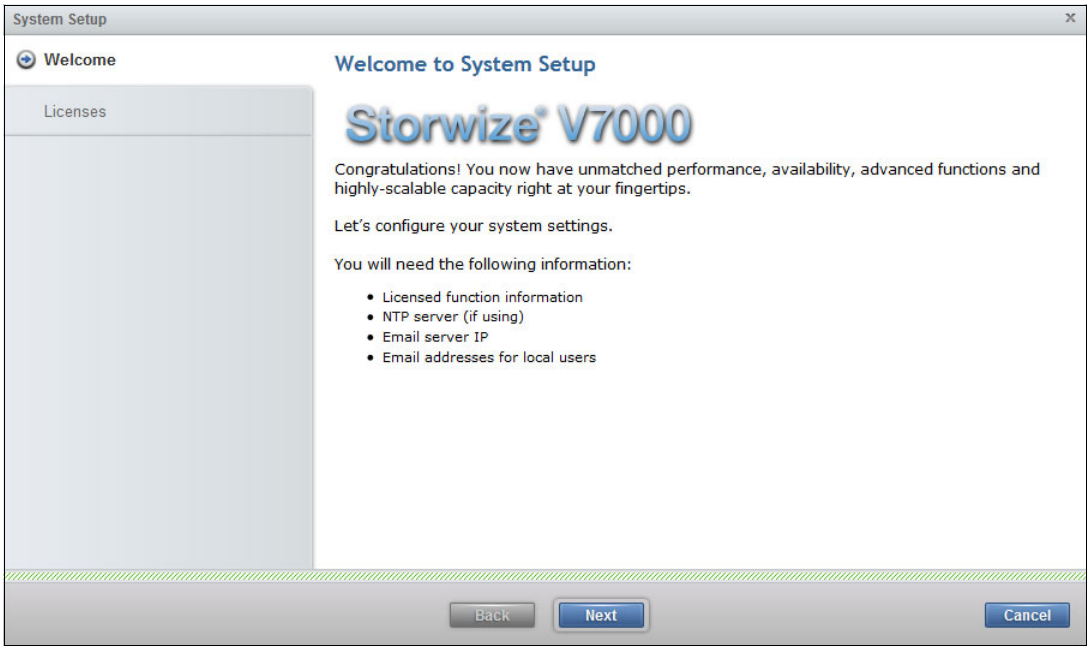

Figure 2-15 Welcome window for system setup

4. Read and accept the license agreement (Figure 2-16).

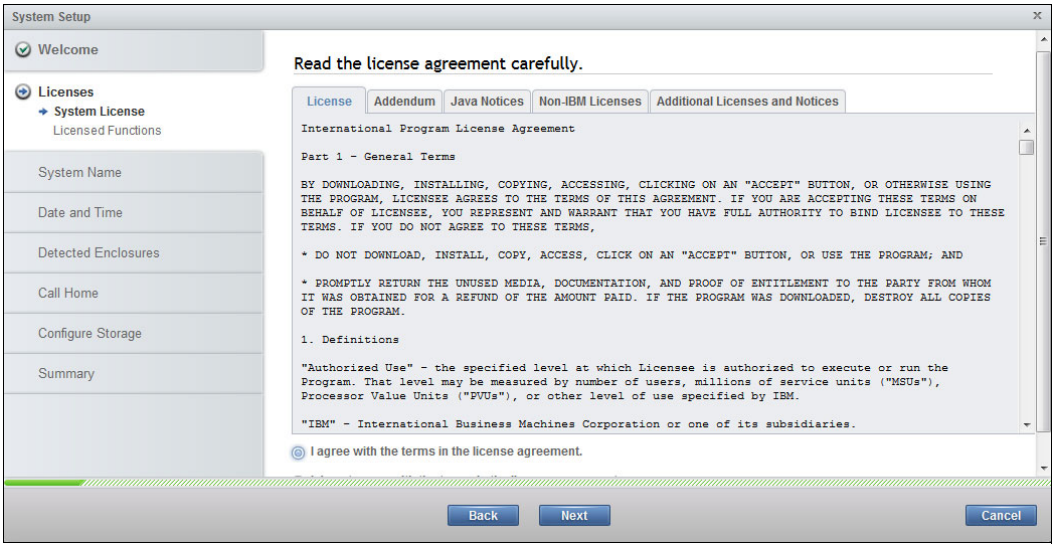

Figure 2-16 Setup wizard: License Agreement

5. Enter the number of purchased licenses (Figure 2-17).

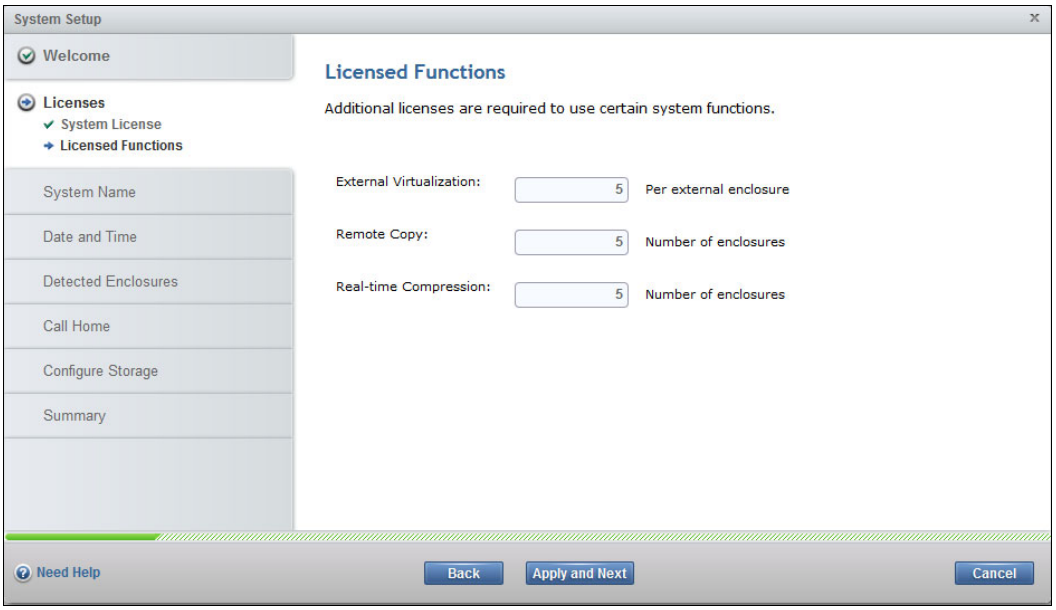

Figure 2-17 Enter the number of purchased licenses

Optional: You can enter advanced licenses for virtualization of external storage devices, Remote Copy limit, or number of Real-time Compression Enclosures as applicable. The virtualization license for all local devices is included in the system and must not be added here.

6. Enter the system name (Figure 2-18).

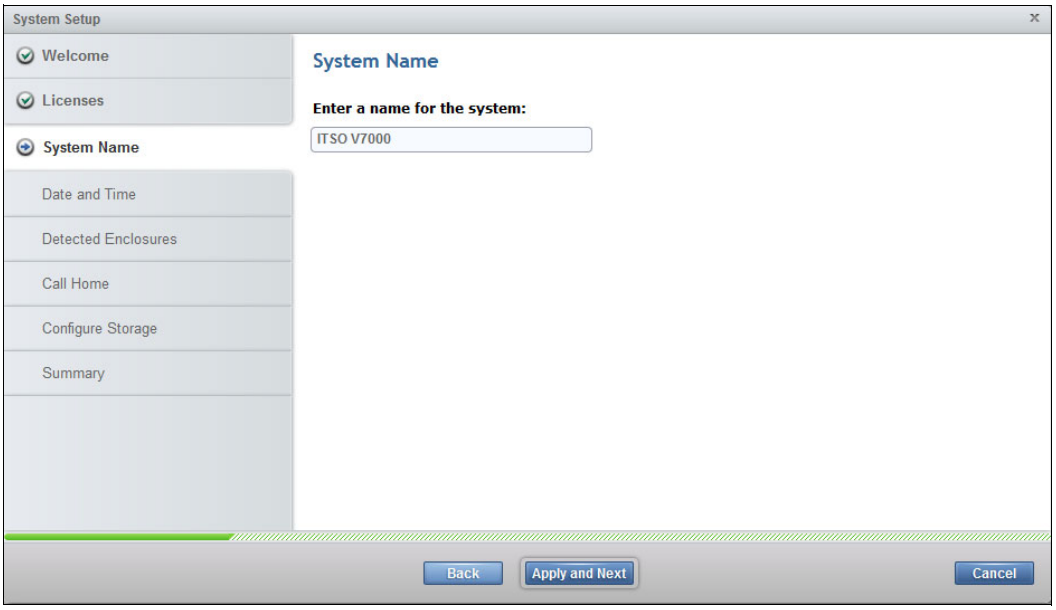

Figure 2-18 System Name

7. Set up current date and time or use an NTP service (Figure 2-19).

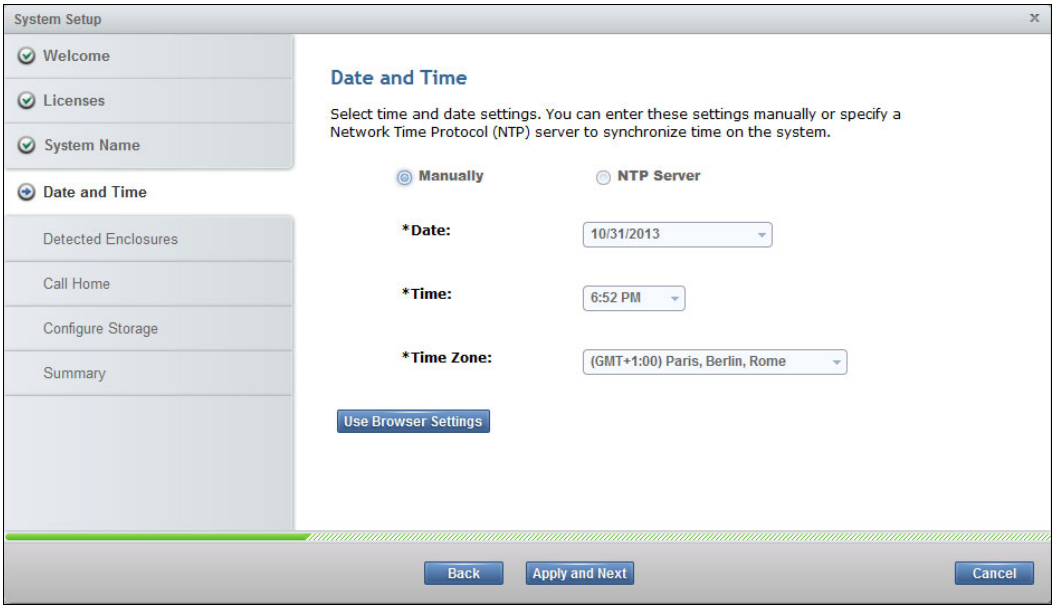

Figure 2-19 Set date and time manually

The detected enclosures are shown. See Figure 2-20.

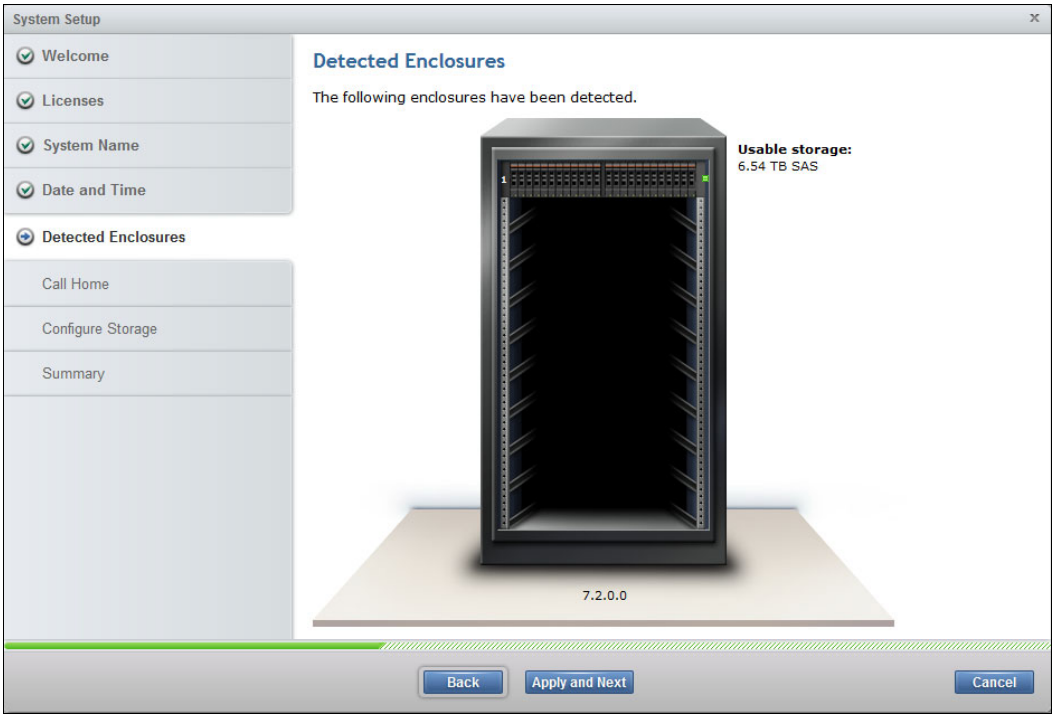

Figure 2-20 Detected enclosures

2.8.1 Configure Call Home, email alert, and inventory

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

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

- `callhome1@de.ibm.com` (USA, Canada, Latin America, and Caribbean Islands)
- `callhome0@de.ibm.com` (All other countries and regions)

To configure Call Home and email alert event notification in IBM Storwize V7000, complete the following steps.

1. Proceed with System Setup. Figure 2-21 shows the start panel for the Call Home configuration wizard. Select **Yes** and click **Next**.

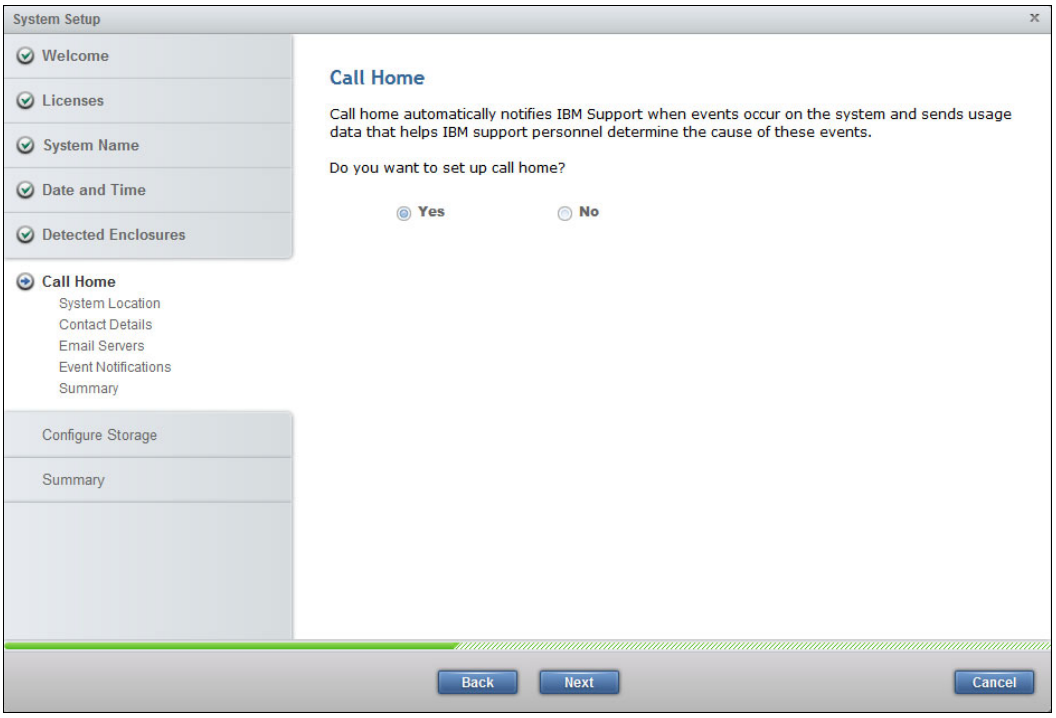

Figure 2-21 Setup wizard option Call Home configuration

2. In the next panel, enter the location of your V7000 (Figure 2-22).

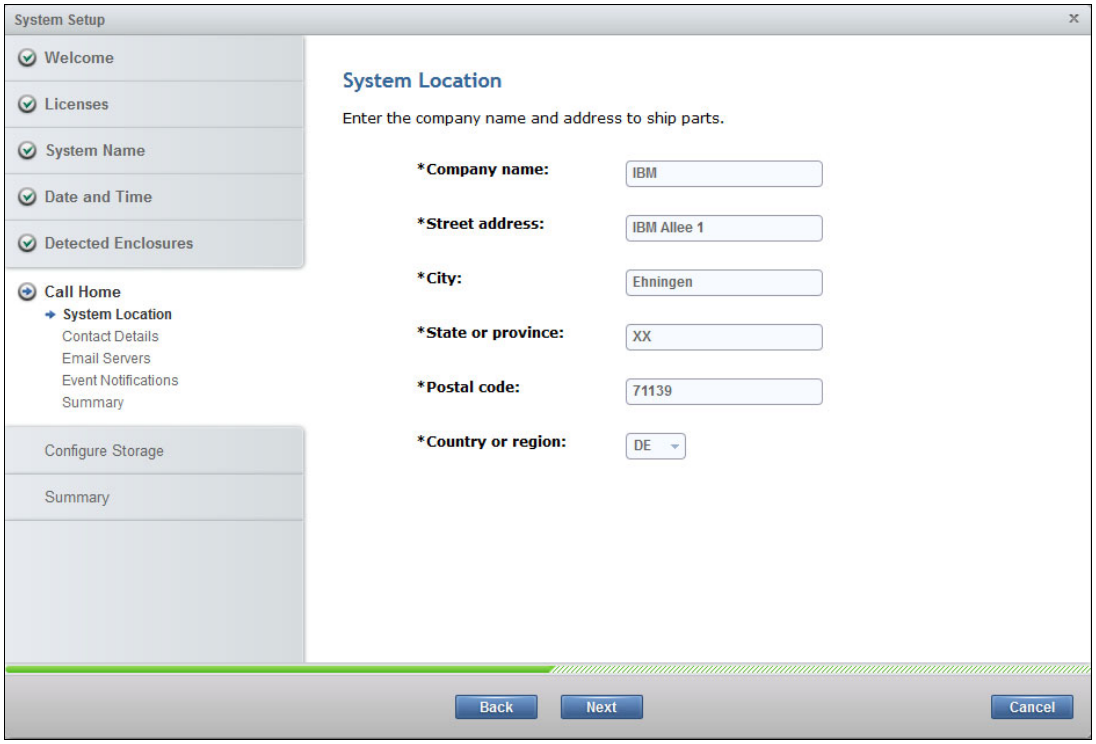

Figure 2-22 System Location of the V7000

3. Specify the contact details (Figure 2-23) for the person who will help IBM Support to resolve problems on the system.

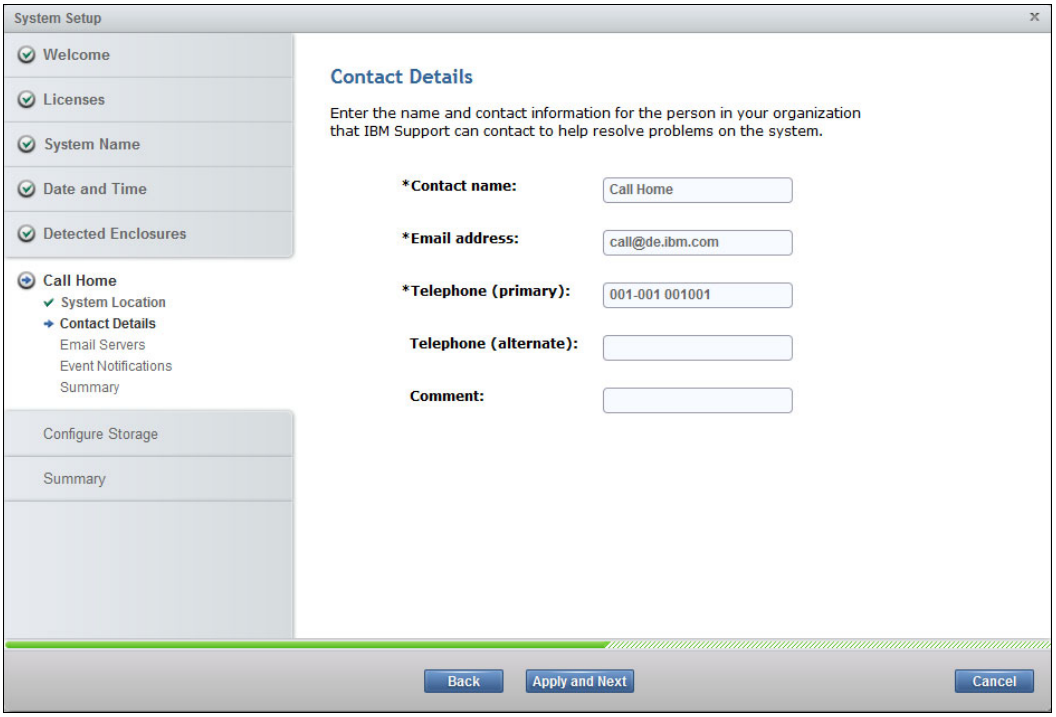

Figure 2-23 Contact Details for the person who will be the contact for the IBM Support

4. Specify the IP address of the email server (Figure 2-24).

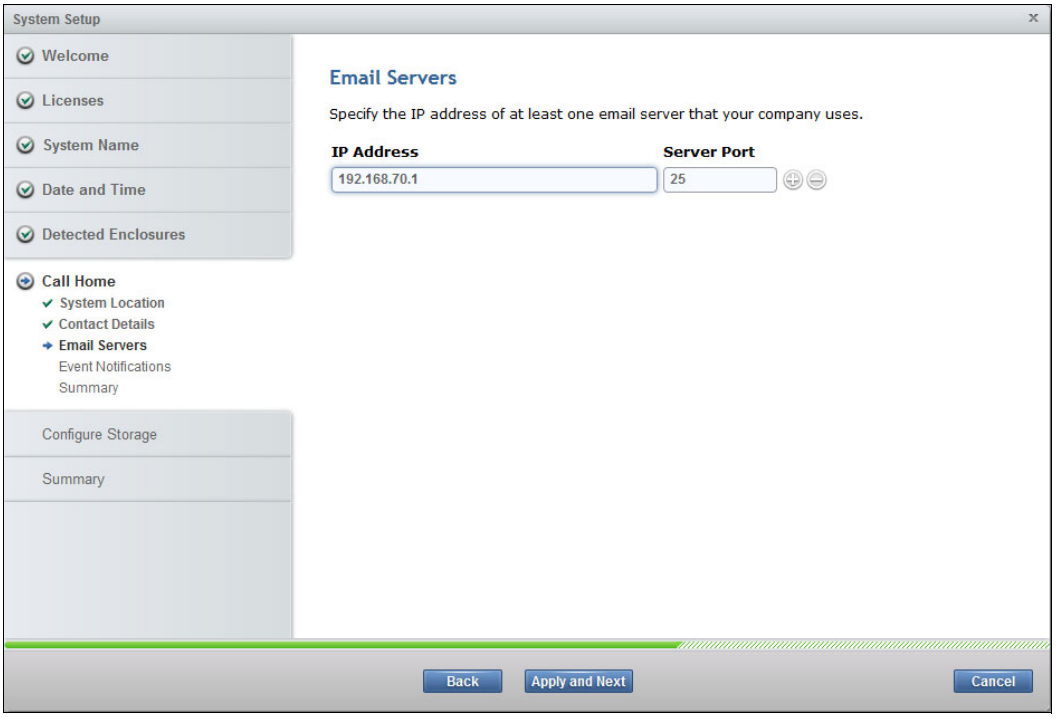

Figure 2-24 IP address for the Email Server

5. Optional: Add another user to receive notifications when events occur (Figure 2-25).

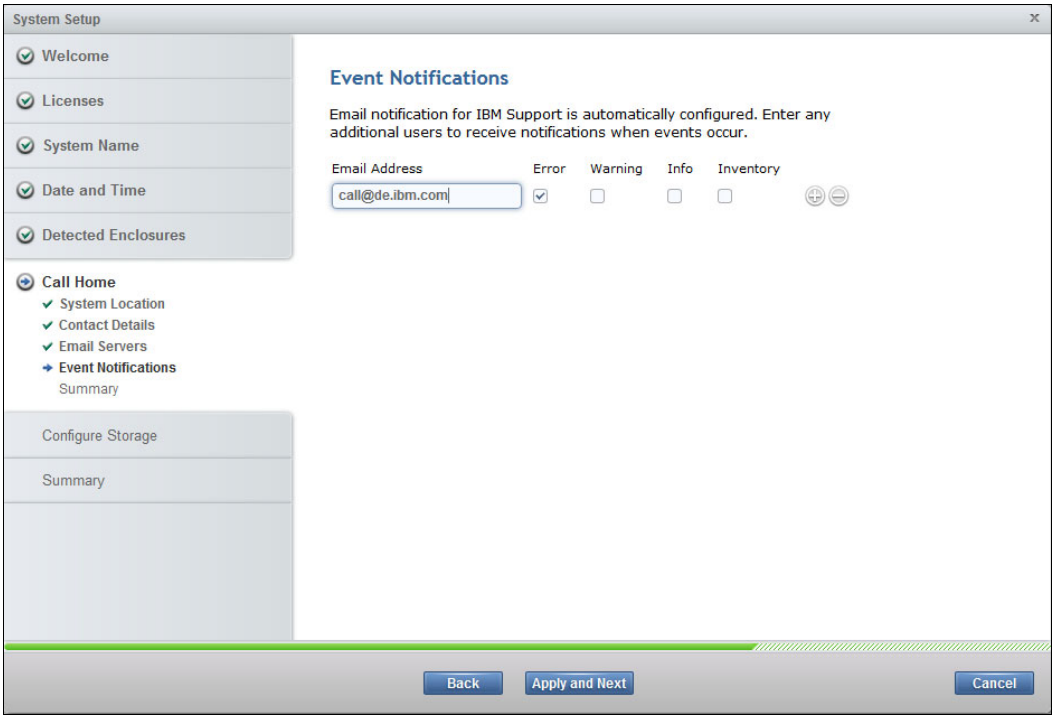

Figure 2-25 Event Notification for additional users

6. Review the summary of changes (Figure 2-26) that are applied in the Call Home panels.

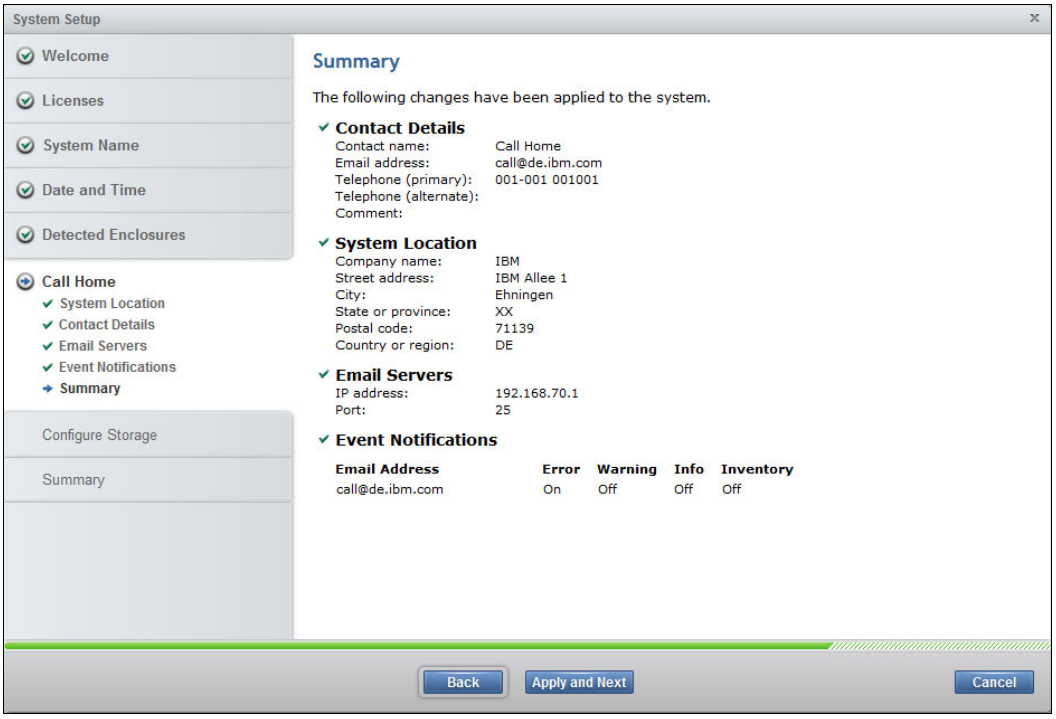

Figure 2-26 Summary panel Call Home

7. Select whether you want storage to be configured automatically now or later (Figure 2-27). For information about how to configure the storage manually, see 7.1.2, “Configuring internal storage” on page 187. To let the system configure the storage, select **Yes**.

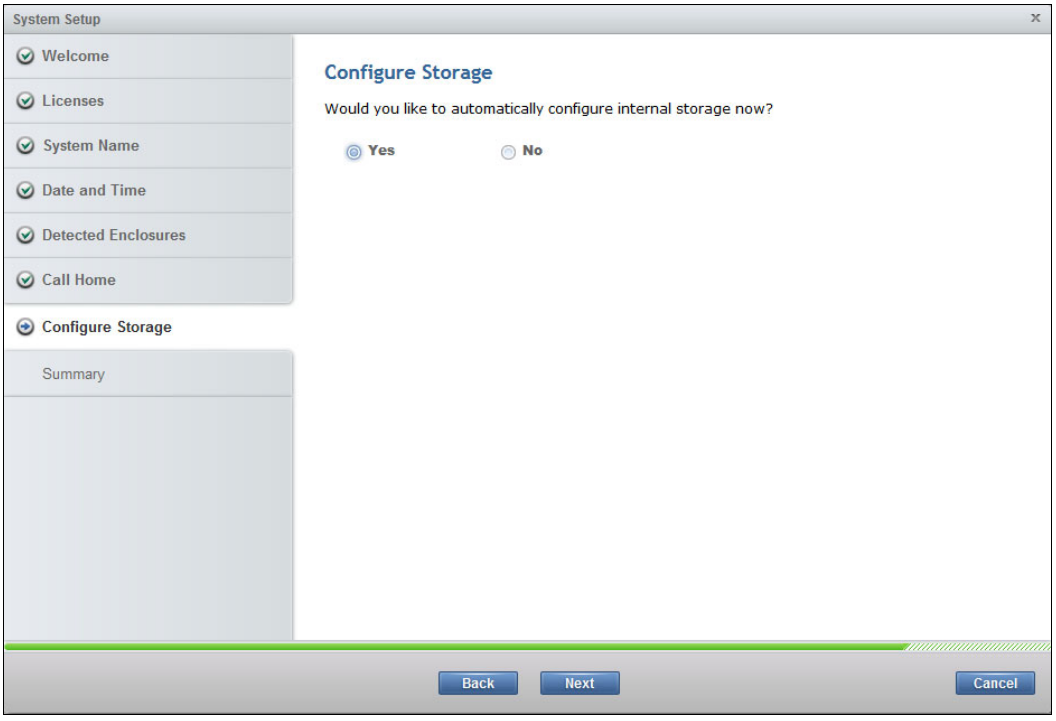

Figure 2-27 Storage Configuration

8. Review the summary of the system configuration (Figure 2-28).

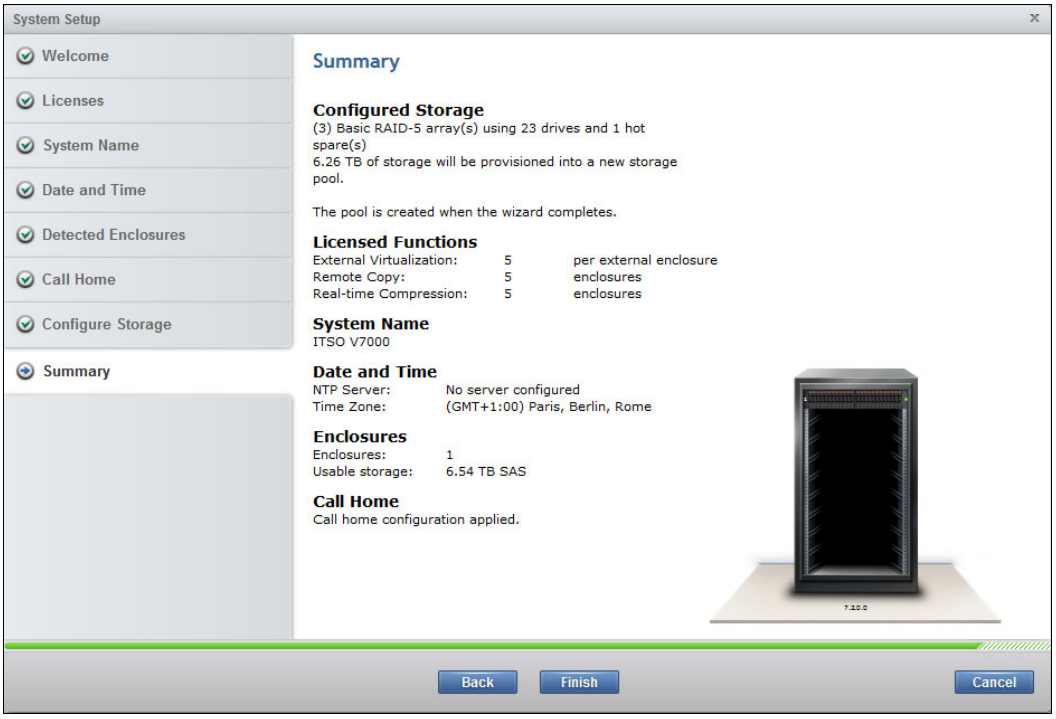

Figure 2-28 System Setup Summary

You have now successfully finalized the initial configuration wizard for the IBM Storwize V7000 system. Figure 2-29 shows the Overview window.

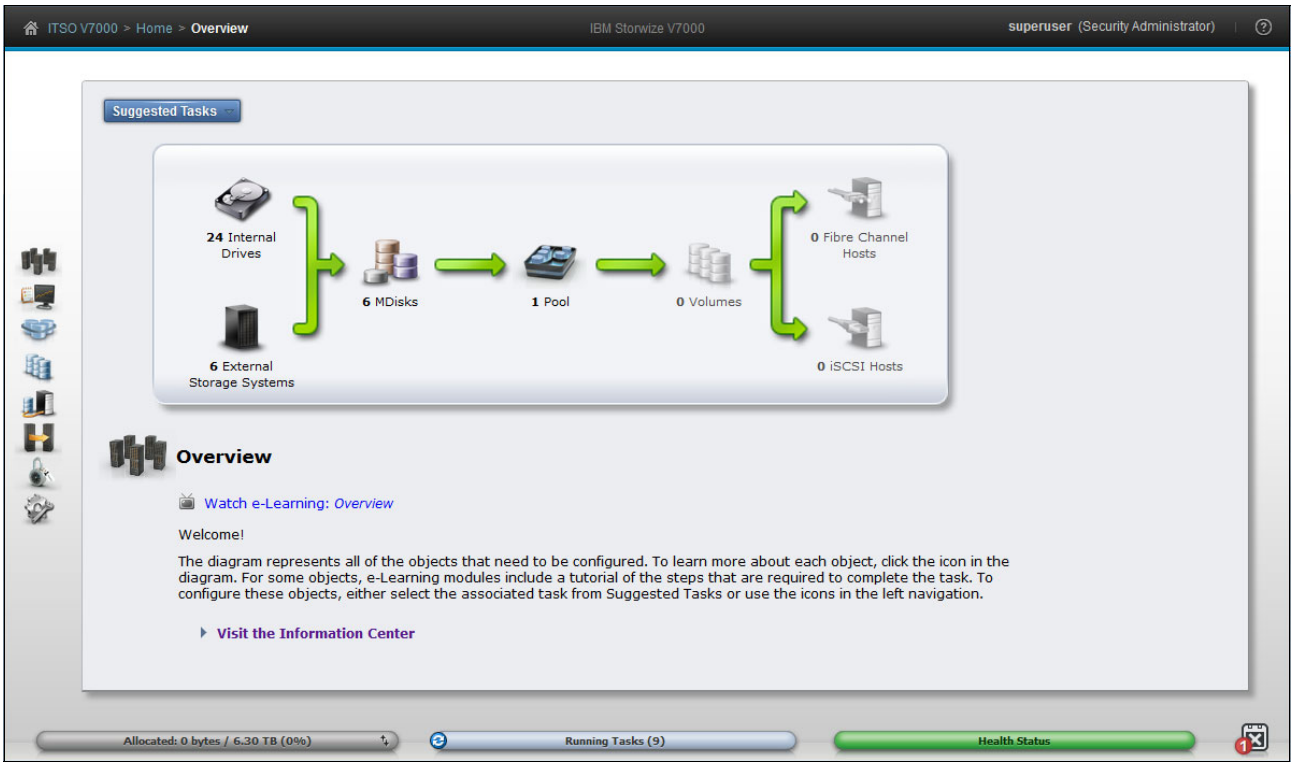

Figure 2-29 IBM Storwize V7000: Getting Started Overview window

3

Graphical user interface overview

This chapter provides an overview of the graphical user interface (GUI) of the IBM Storwize V7000 and shows how to use the navigation tools.

3.1 Overview of IBM Storwize V7000 management software

A graphical user interface can simplify storage management and provide a fast and more efficient management tool. It is loosely based on the IBM System Storage XIV® software and has a similar look and feel.

JavaScript: You must enable JavaScript in your browser. When using Mozilla Firefox, under Advanced JavaScript Settings, click **Disable or replace context menus** and allow cookies.

3.1.1 Access to the storage management software

To log on to the management software, type the IP address that was set during the initial setup process into the address line of your web browser. You can connect from any workstation that can communicate with the system. The login window opens (Figure 3-1).

Figure 3-1 Login window

We suggest that each user who operates the IBM Storwize V7000 has an account that is not shared with someone else. The default user accounts will be disabled for remote access or the passwords will be changed from the default and known only to the system owner, or kept secured for emergency purposes only. Doing this helps to identify the personnel working on the device and to track all the important changes in the systems.

After successful login, the IBM Storwize V7000 Overview panel opens (Figure 3-2).

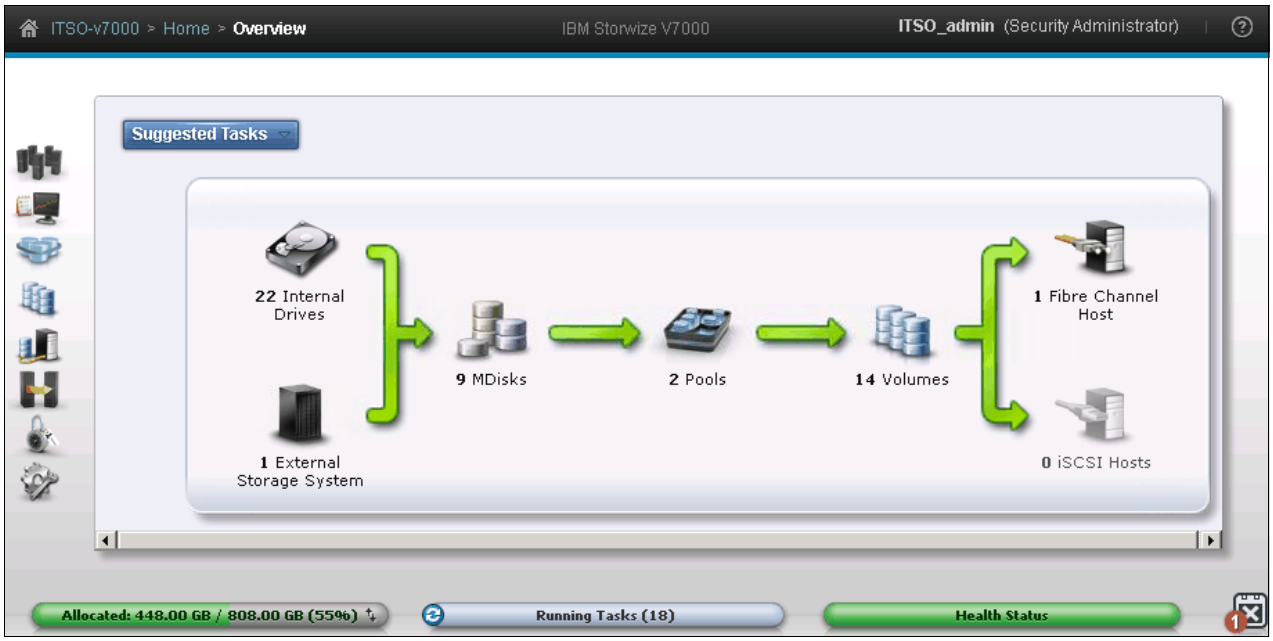

Figure 3-2 Getting started with the IBM Storwize V7000 Overview panel

The IBM Storwize V7000 Overview panel is an important user interface; throughout this chapter we refer to it as the *V7000 Overview panel* or just the *Overview panel*. In the remaining chapters, we do not explain how to access it each time.

3.1.2 Overview panel layout

This Overview panel has three main sections for navigating through the management tool. Figure 3-3 shows the three main areas.

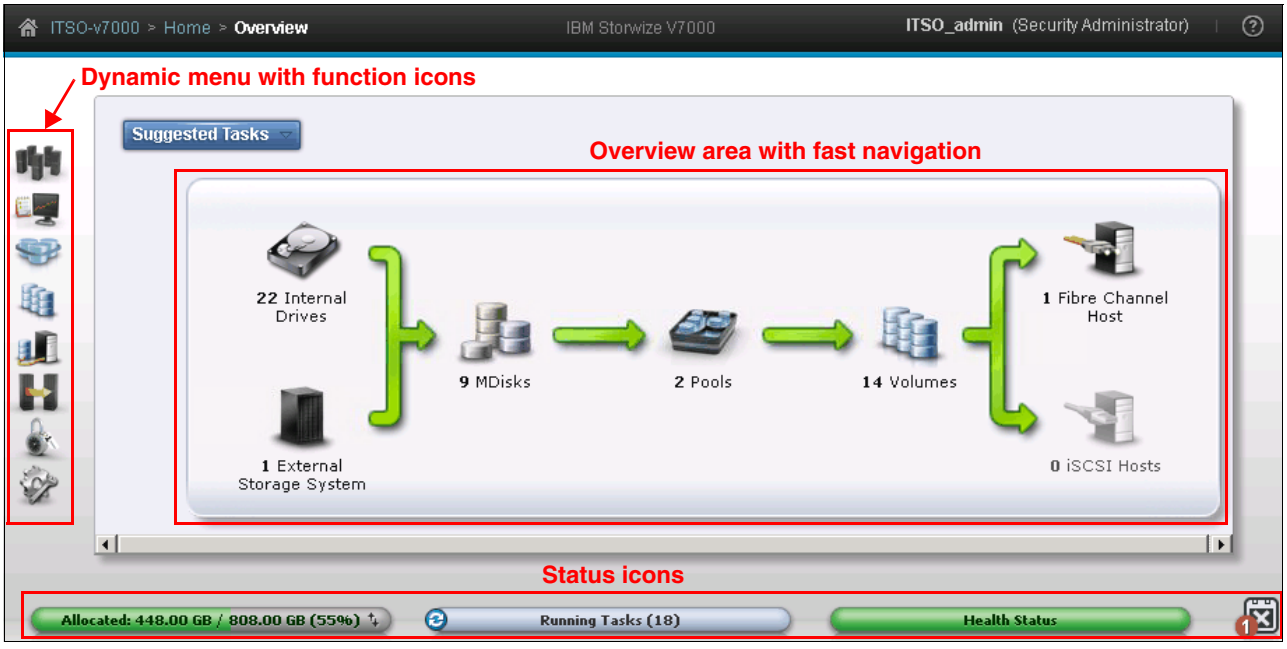

Figure 3-3 Main areas

- The left side of the window shows eight *function icons*. We refer to them collectively as a *dynamic menu*. The dynamic menu includes these function icons:
  - Home menu
  - Monitoring menu
  - Pools menu
  - Volumes menu
  - Hosts menu
  - Copy Services menu
  - Access menu
  - Settings menu
- The middle of the window shows a diagram of the existing configuration. Clicking the function icons in this area provides extended help references, including a link to a short video presentation to explain the topic in more detail. This video is not a navigation tool, but instead is an additional source of information.
- The bottom of the window shows three status indicators. Clicking any of them provides more detailed information about the existing configuration of the IBM Storwize V7000 solution. Click any of these function icons to expand them and minimize them as required or to switch between different information, for example virtual or allocated storage.

3.1.3 Navigation

Navigating in the management tool can be simple. You can hover the cursor over one of the eight function icons to emphasize that icon and display a menu of options. You can move the cursor to an option and click it. Figure 3-4 shows how to access the Support option.

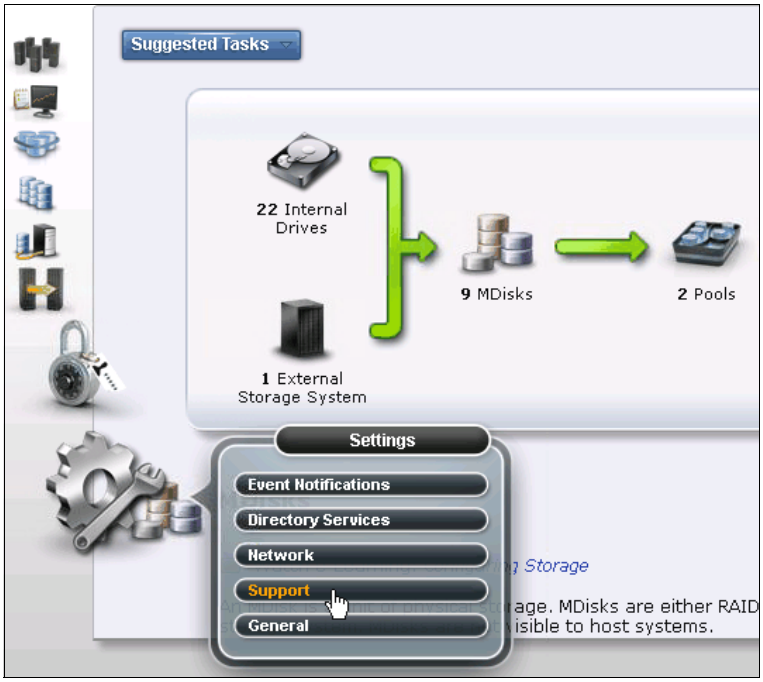

Figure 3-4 Navigation using the menu options

Figure 3-5 shows a list of the IBM Storwize V7000 Software function icons and the associated menu options.

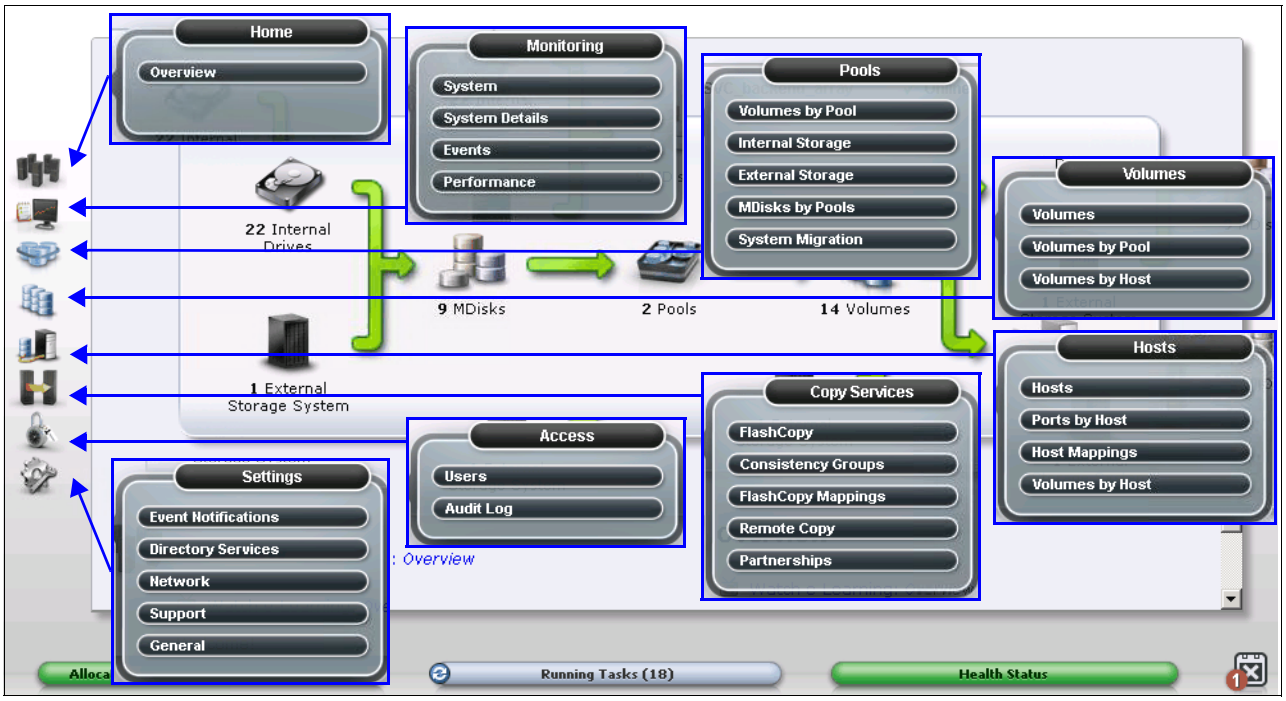

Figure 3-5 BM Storwize V7000 Overview panel and its menu options

3.1.4 Multiple selections

With the improved management tool, you can select multiple items by using a combination of Shift or Ctrl keys. To select multiple items in a display, click the first item, press and hold the Shift key, and click the last item in the list you require. All rows between those two items are selected and highlighted in light blue (Figure 3-6).

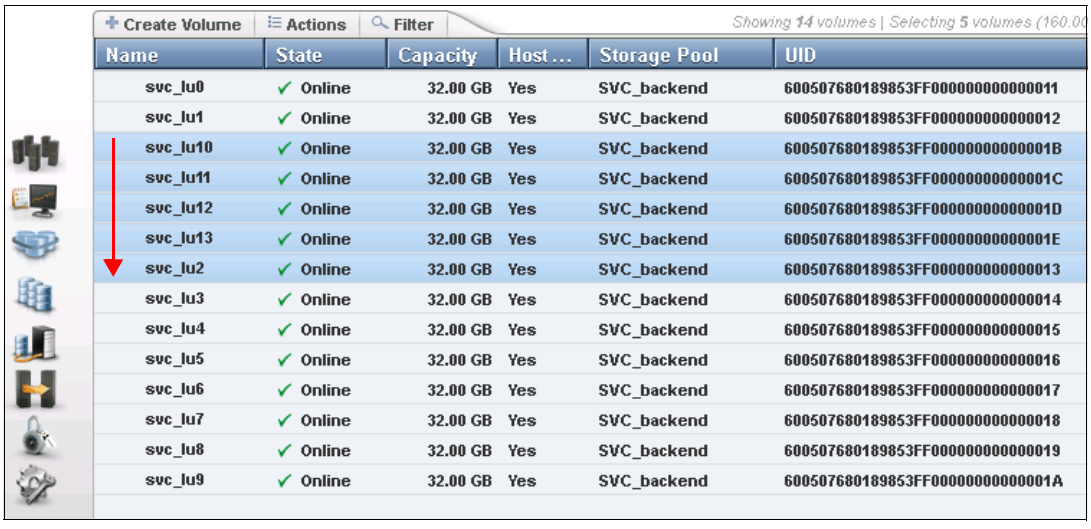

Figure 3-6 Multiple selections using the Shift key

Similarly, if you want to select multiple items that are not in sequential order, click the first item, press and hold the Ctrl key, and click the other items you need.

3.1.5 Status indicators area

The status indicators area at the bottom of the Overview panel (Figure 3-7) show a brief status about the IBM Storwize V7000 storage system. The four indicators are as follows (from left to right):

- Capacity overview
- Running tasks
- Health status
- Status alerts

Figure 3-7 Status Indicators menus

Click each icon or hover the mouse cursor over the icon to see more details about status of the system or task indicated by the icon.

Capacity overview

The left-most indicator shows information about the overall physical capacity (the initial amount of storage that was allocated) and virtual capacity (thin provisioned storage; volume size is dynamically changed as data grows or shrinks, but you still see a fixed capacity). Click the indicator to switch between physical and virtual capacity (Figure 3-8).

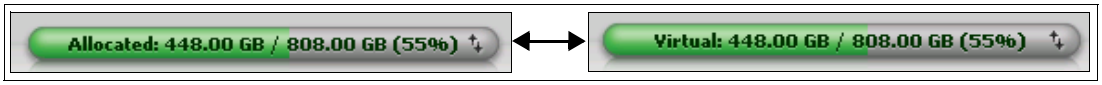

Figure 3-8 Storage capacity overview

Running tasks

The Running tasks indicator provides the information about jobs initiated by the user or system and that typically run for a longer period of time, such array initialization, synchronization, or rebuilt. Hover the mouse cursor over the circular arrows icon on the indicator to open the context menu, which lists all running task (Figure 3-9).

Figure 3-9 Running tasks status

Health status

The Health status indicator warns administrators about critical errors or alerts that need greater attention. In a critical alert, the green color of the indicator becomes red. You can then click the X icon (Figure 3-10) to open the context menu to see the status alerts. If the status is healthy (green, not red), no context menu is available.

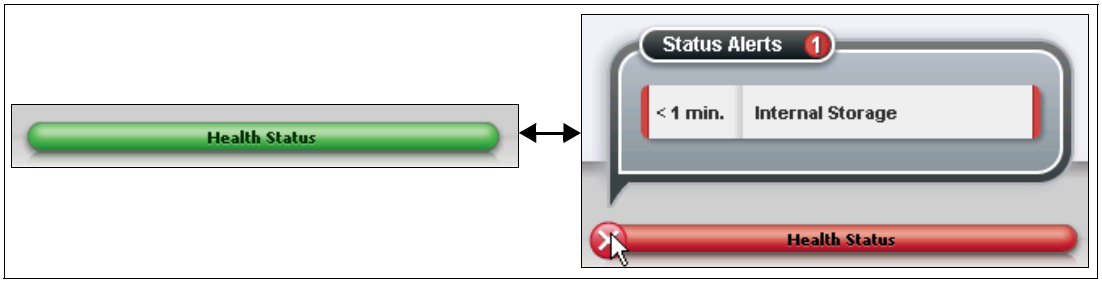

Figure 3-10 Health status of the system

Status alerts

The indicator for status alerts is at the bottom right corner only in case of an important error or warning, which is not yet acknowledged or fixed in the Event log. Hover the mouse cursor over the icon to see the list of important system messages (Figure 3-11). After all alerts are cleared from the system, the Status alerts icon automatically disappears from the view.

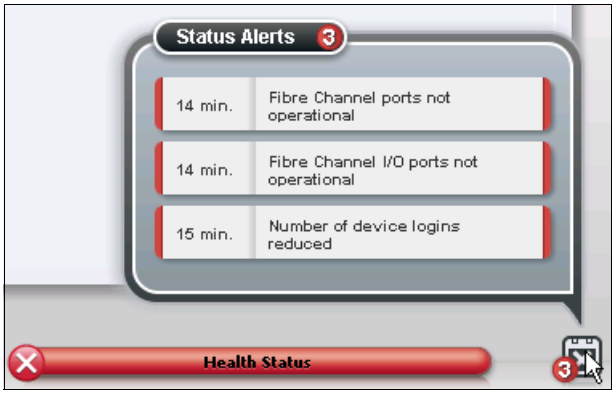

Figure 3-11 IBM Storwize V7000 status alerts

Help

Another useful interface feature is integrated help. You can access help for some fields and objects by hovering the mouse cursor over the question mark icon, next to the field or object (Figure 3-12). Panel-specific help is available by clicking **Need Help** or by using the **Help** link in the upper-right corner of the GUI.

Figure 3-12 Access to panel-specific help

3.2 Home menu

By clicking the Home menu function icon, you can get direct access to the Overview panel (Figure 3-13).

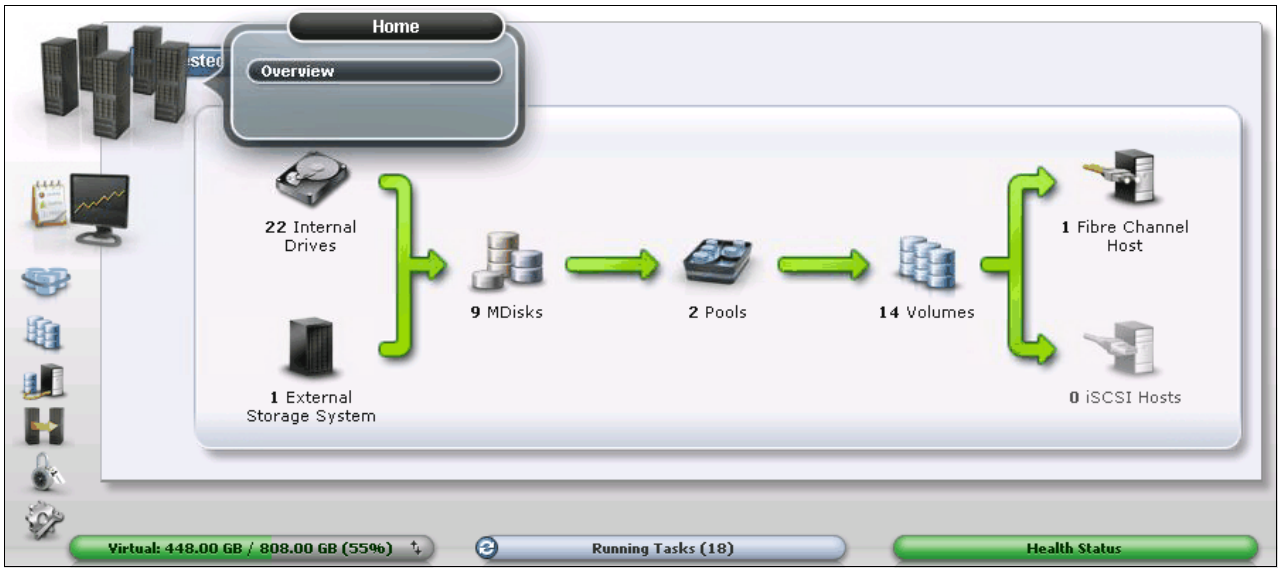

Figure 3-13 Navigate to the Home menu

See 3.1.2, “Overview panel layout” on page 67 to understand the structure of the panel and how to navigate to various system components and manage it more efficient and quickly.

3.3 Monitoring menu

Hover the cursor over the Monitoring function icon to open the Monitoring menu (Figure 3-14).The Monitoring menu offers these navigation directions:

- System: Opens the general overview of your IBM Storwize V7000 system including the depiction of all devices in a rack and the storage capacity. See 3.3.1, “System overview” on page 73.
- System Details: Provides an extended level of information related to the system, including composition of each element in an overall system configuration. In addition to the software version and system unique ID, it shows vital product data of each component. See 3.3.2, “System Details” on page 75.
- Events: Tracks all informational, warning, and error messages that occurred in the system. You can apply various filters to sort them according to your needs or export them to the external comma-separated values (CSV) file. See 3.3.3, “Events” on page 77.
- Performance: Reports the general system statistics with regards to the processor (CPU) utilization, host and internal interfaces, volumes, and MDisks. You can switch between MBps or IOPS. See 3.3.4, “Performance” on page 78.

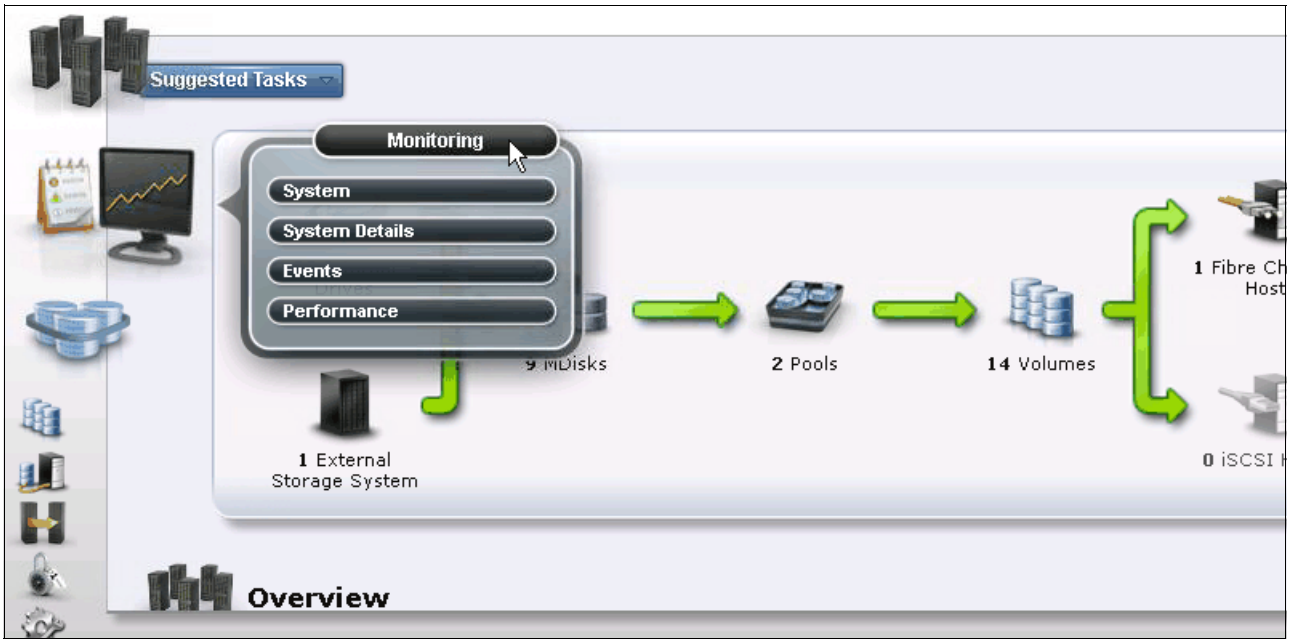

Figure 3-14 Accessing the Monitoring menu

In the following sections, we describe each option on the Monitoring menu. Detailed procedures of how to best use some of these items are described in other chapters (for example, Events in Chapter 13, “RAS, monitoring, and troubleshooting” on page 455).

3.3.1 System overview

The System option on the Monitoring menu provides the general overview about your IBM Storwize V7000 system including the depiction of all devices in a rack and the allocated or physical storage capacity. When thin-provisioned volumes are enabled, the virtual capacity is also shown. See Figure 3-15 on page 74 for details.

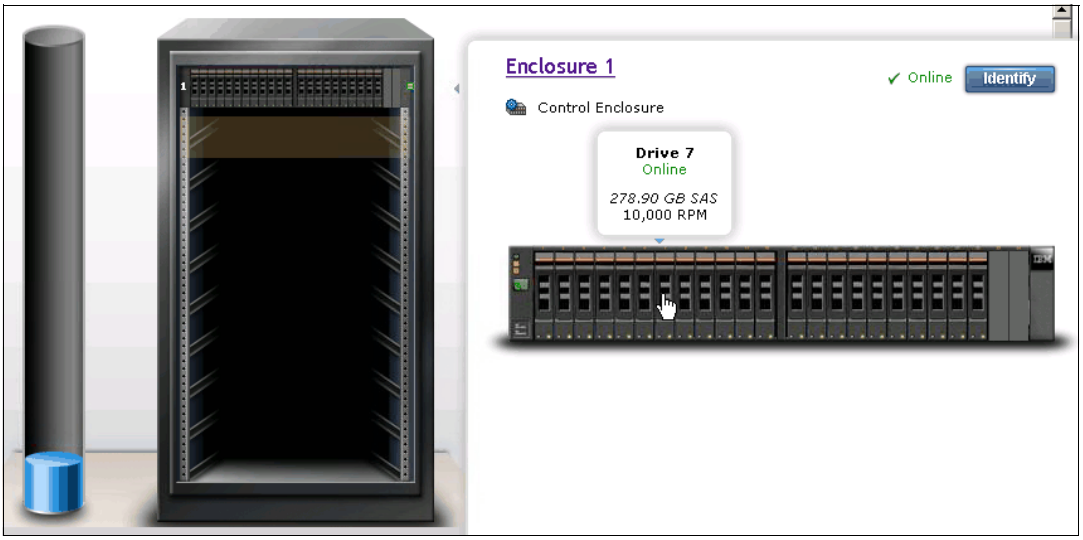

Figure 3-15 System overview with drive details

When you click a specific enclosure, a pop-up window indicates the composition of disk drives in the unit. When you hover the mouse cursor over each disk drive, you see detailed technical parameters such as capacity, interface, rotation speed, and the drive status (online or offline).

In an environment with multiple IBM Storwize V7000 systems, you can easily navigate the onsite personnel or technician to the correct device by enabling the identification LED on the front panel. Click **Identify** in the pop-up window shown in Figure 3-15 and wait for the confirmation of the technician that the device in the data center was properly identified.

After confirmation, click **Turn LED Off** (Figure 3-16).

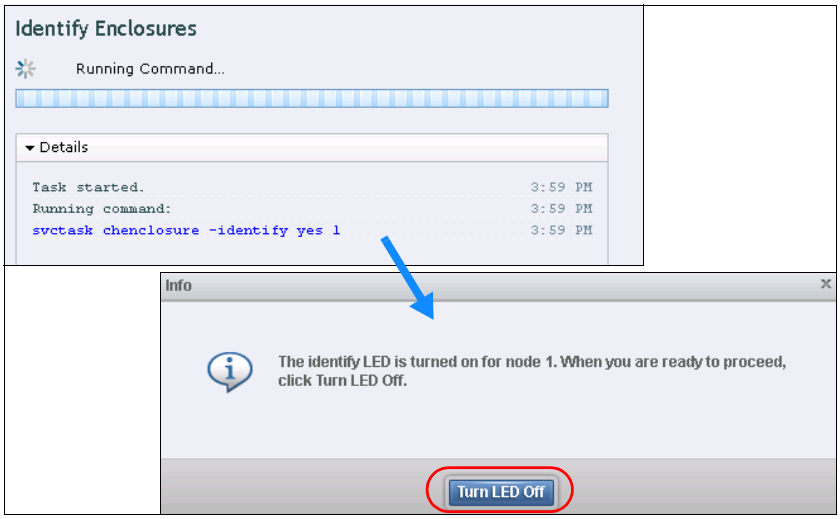

Figure 3-16 Using identification LED

Alternatively, you can use the IBM Storwize V7000 command-line interface (CLI) to get the same results. Type the following sequence of commands:

1. **svctask chenclosure -identify yes 1** (or just **chenclosure -identify yes 1**).
2. **svctask chenclosure -identify no 1** (or just **chenclosure -identify no 1**).

3.3.2 System Details

The System Details option on the Monitoring menu provides the extended level of information related to the system, including the composition of each element in an overall system configuration (Figure 3-17).

Figure 3-17 System details

When you click the reference to the specific enclosure (in our example Enclosure 1), you see vital product data (VPD) about the enclosure (Figure 3-18).

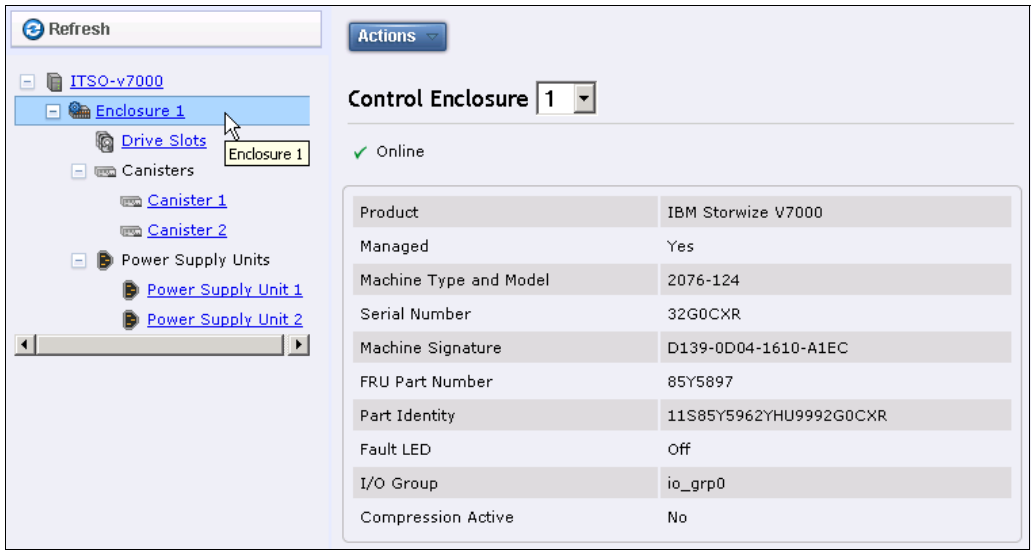

Figure 3-18 Enclosure technical details

You can gather the same level of detail from each canister of an enclosure by clicking its reference link (Figure 3-19).

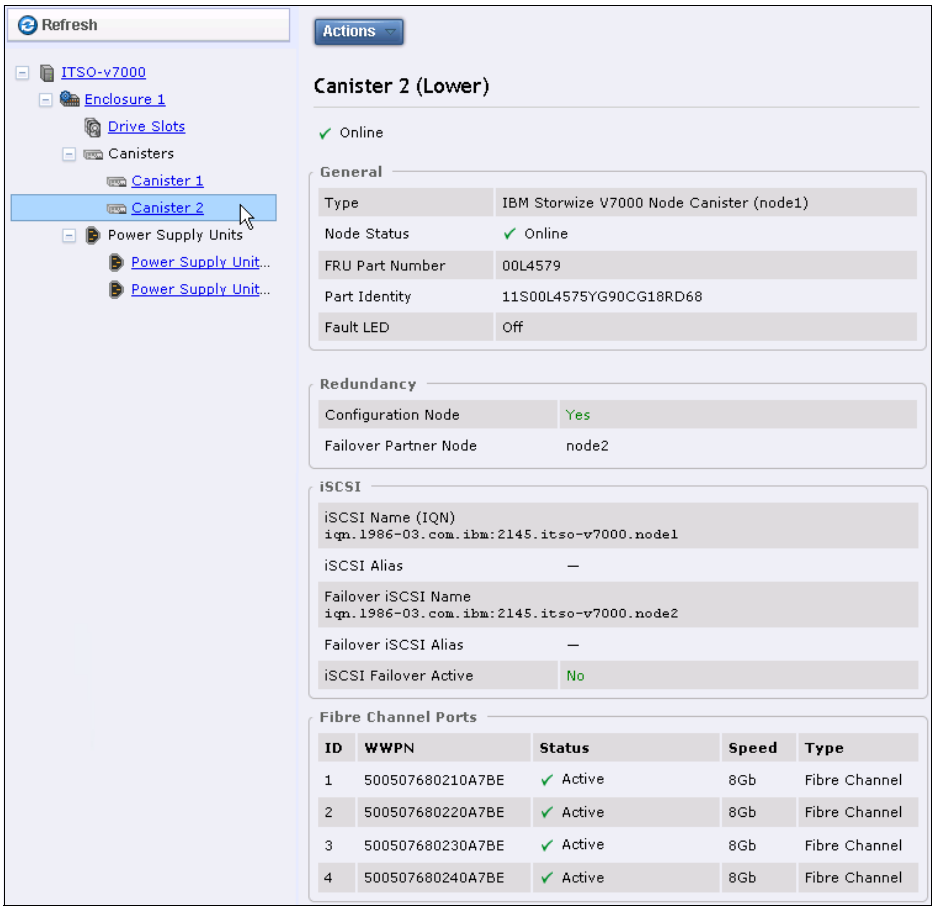

Figure 3-19 Canister details and vital product data

In addition, you can see the status of power supply units with integrated batteries. Click **Power Supply Unit** (Figure 3-20).

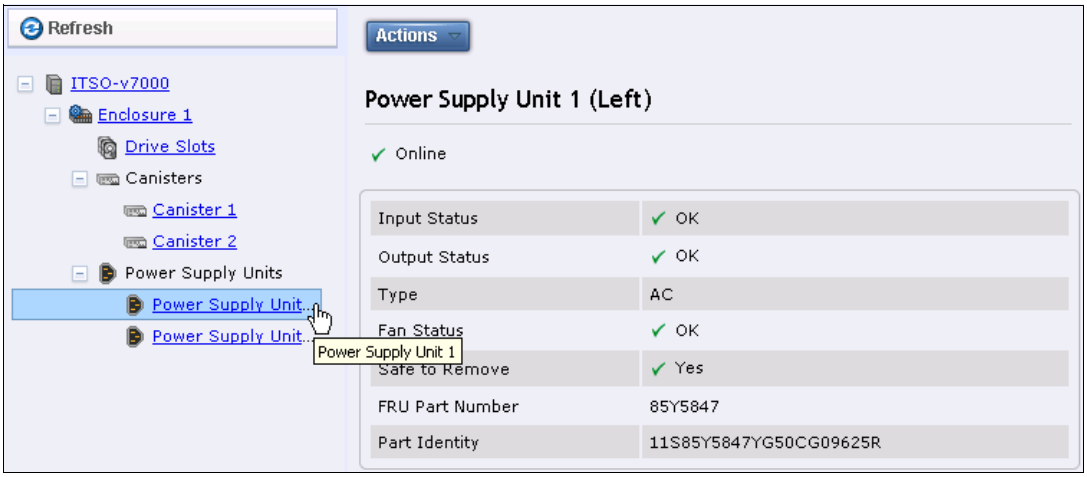

Figure 3-20 Status of power supplies

3.3.3 Events

The Events option, selected from the Monitoring menu, tracks all informational, warning, and error messages that occur in the system. You can apply various filters to sort them or export them to an external comma-separated values (CSV) file. A CSV file can be created from the information that is here. Figure 3-21 shows the display after the Event Log is selected from the menu.

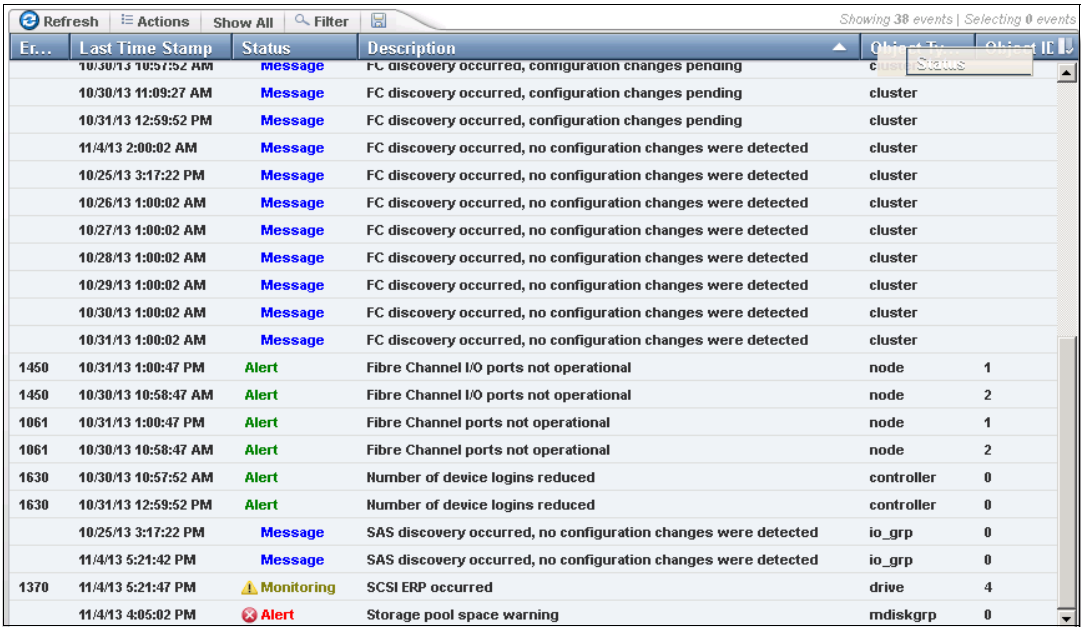

Figure 3-21 Event log menu

The exact procedures for how to work with the Event log and how to run various fix procedures using the Events option are described in Chapter 13, “RAS, monitoring, and troubleshooting” on page 455.

3.3.4 Performance

The Performance pane reports the general system statistics regarding processor (CPU) utilization, host and internal interfaces, volumes, and MDisks. You can switch between MBps or IOPS or even drill down the statistics to the node level. This might be useful when comparing the performance of each node in the system if problems exist after a node failover occurs. See Figure 3-22.

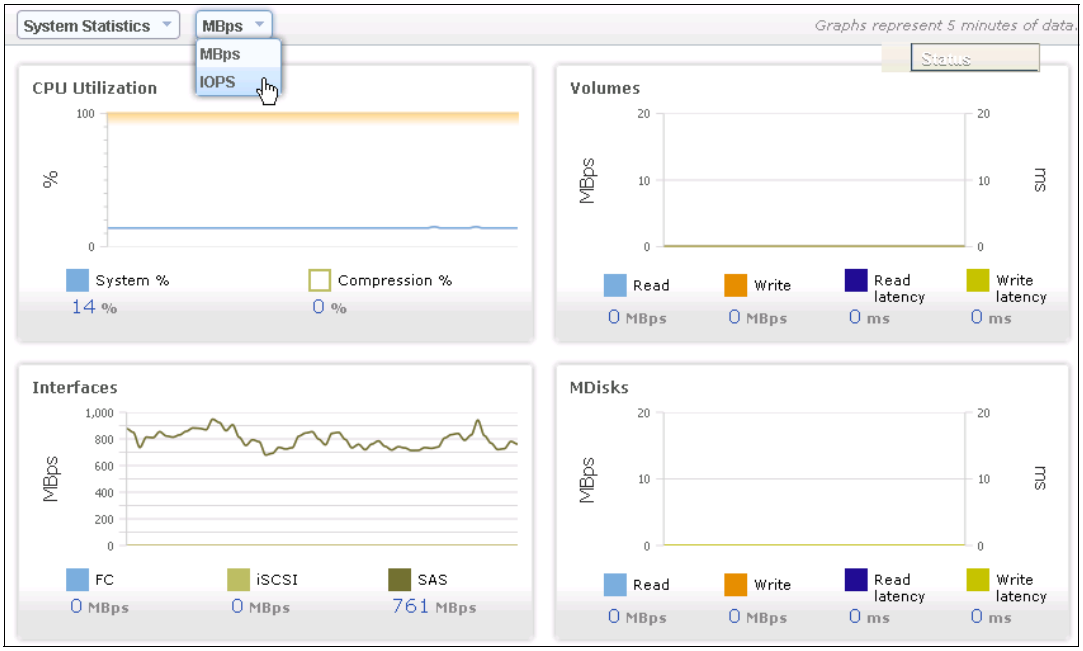

Figure 3-22 Performance statistics of IBM Storwize V7000 system

The performance statistics in the GUI show, by default, the latest five minutes of data. To see details of each sample, click the graph and select the time stamp, as shown in Figure 3-23.

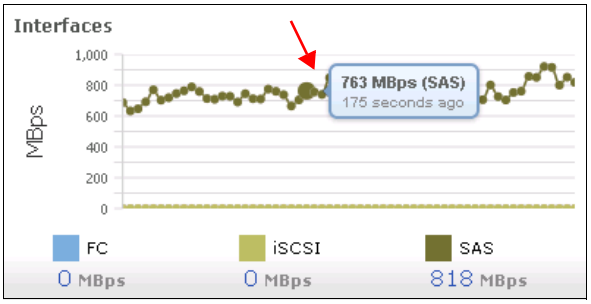

Figure 3-23 Sample details

For in-depth storage monitoring and performance statistics of you IBM Storwize V7000 system with historical data, use the IBM Tivoli Storage Productivity Center for Disk.

3.4 Pools menu

A *pool* or storage pool is a collection of managed disks (MDisks) that jointly contain all of the data for a specified set of volumes. All MDisks in a pool are split into *extents* of the same size. Volumes are created from the extents that are available in the pool. You can add MDisks to a storage pool at any time, either to increase the number of extents that are available for new volume copies or to expand existing volume copies.

Hover the cursor over the Pools function icon to display the Pools menu options (Figure 3-24).

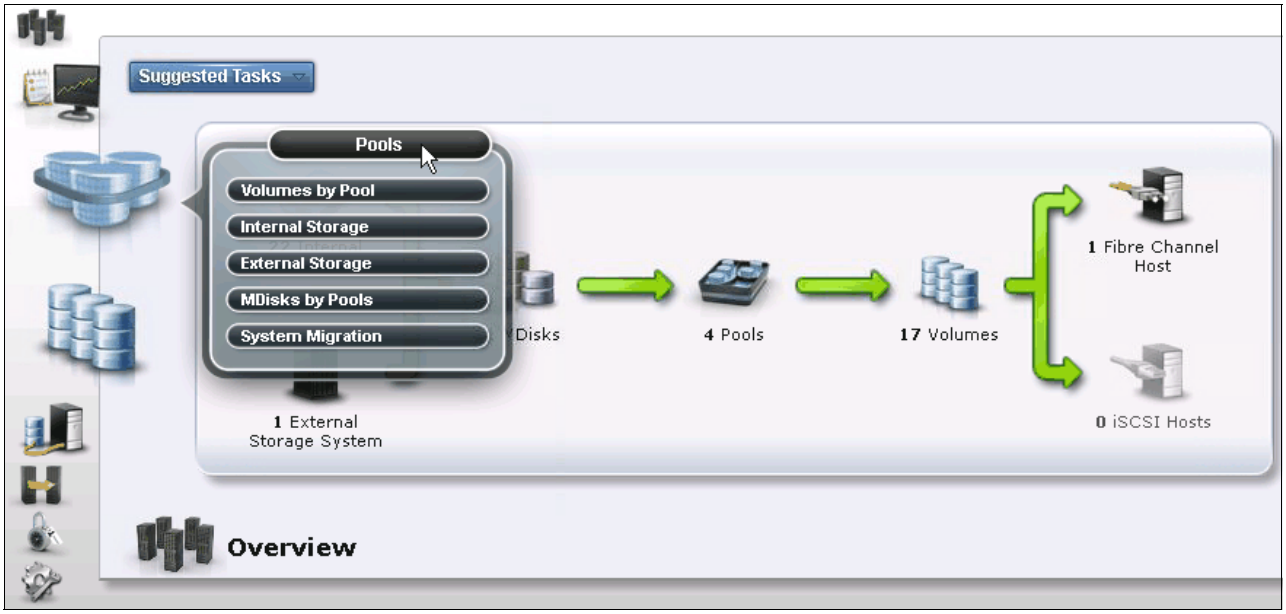

Figure 3-24 Navigate to the Pools menu

The Pools menu has the following options:

- Volumes by Pool: Applies the high-level filter, listing all defined volumes per pool. It also provides a capacity overview, valid only for a specific, selected pool. This view is excellent when planning a migration of a volume to another pool, so you have a common overview of all pools and their associated volumes. Unused volumes are not listed.
- Internal Storage: Provides the overview about all disk drives that are installed in the IBM Storwize V7000 system including its enclosures. You can filter based on disk type and capacity and see also unused volumes not assigned to any pool.
- External Storage: Shows all pools and their volumes created from the systems connected to the IBM Storwize V7000 externally and integrated into the system repository. It does not show any internal pools or volumes. This is also called external virtualization.
- MDisks by Pools: Provides the list of all managed disks either internally or externally connected and associated with one of the defined pools. It also lists all unassigned MDisks separately.
- System Migration: Offers the migration wizard to import data from image-mode MDisks to a specific pool. It is useful when migrating data from an old external storage to the IBM Storwize V7000 non-disruptively to the hosts.

3.4.1 Volumes by pool

This menu option lists all defined volumes, sorted by their pool assignment (Figure 3-25). Unassigned volumes are not visible in this menu. By using this menu you can, for example, create volumes, map or unmap volumes to and from hosts, migrate volumes to another pool, and rename, shrink, or expand volumes.

Figure 3-25 Listing volumes by pool

In addition, you can rename the pool or choose a different icon to represent the pool (Figure 3-26). Choose the appropriate icon.

Figure 3-26 Changing the pool icon

When the pools are defined and the volumes are assigned, the pool can have one of the following four operational status values:

Online The storage pool is online and available. All the MDisks in the storage pool are available.

Degraded path One or more nodes in the clustered system cannot access all the MDisks in the group. A degraded path state is most likely the result of incorrect configuration of either the storage system or the Fibre Channel fabric. However, hardware failures in the storage system, Fibre Channel fabric, or node can also be a contributing factor to this state.

Degraded ports One or more 1220 errors have been logged against the MDisks in the storage pool. The 1220 error indicates that the remote Fibre Channel

port has been excluded from the MDisk. This error might cause reduced performance on the storage system and usually indicates a hardware problem with the storage system. To fix this problem, you must resolve any hardware problems on the storage system and fix the 1220 errors in the event log. To resolve these errors in the log, select **Troubleshooting** → **Recommended Actions** in the management GUI. This action displays a list of unfixed errors that are currently in the event log. For these unfixed errors, select the error name to begin a guided maintenance procedure to resolve them. Errors are listed in descending order with the highest priority error listed first. Resolve highest priority errors first.

Offline The storage pool is offline and unavailable. No nodes in the system can access the MDisks. The most likely cause is that one or more MDisks are offline or excluded.

3.4.2 Internal storage

Click the **Internal Storage** option in the Pools menu to open a window similar to Figure 3-27. From this window, you can allocate RAID arrays of internal disk drives into storage pools. This window also offers the option to display internal drives, based on their capacity and speed.

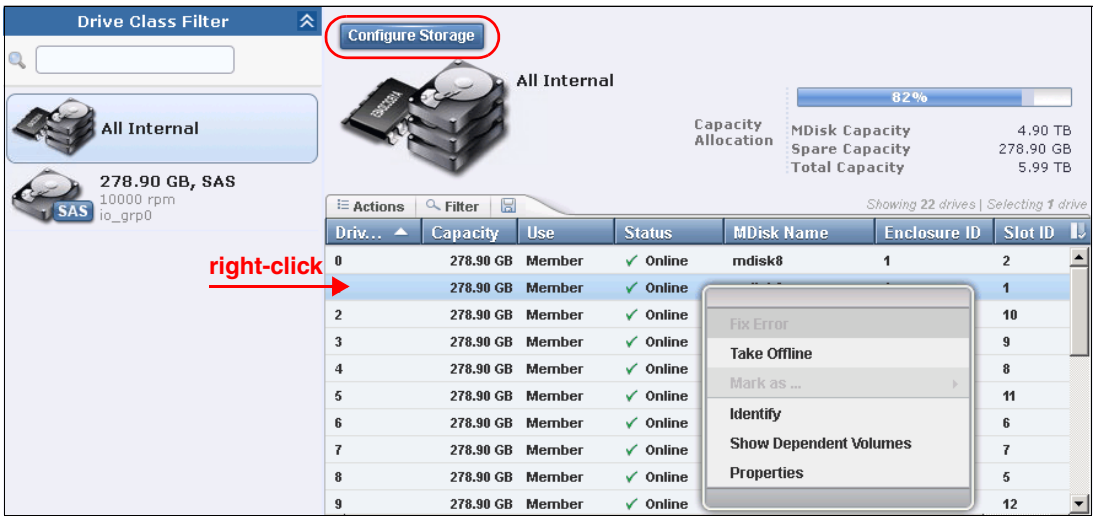

Figure 3-27 Internal storage window

Click **Actions** in the table header or right-click a specific MDisk to take a managed disk offline, show its properties, or mark the MDisk as Unused, Candidate, or Spare.

You can configure a storage array and assign it to the pool by first clicking **Configure Storage** (circled in Figure 3-27). A wizard opens where you choose the recommended array level or select a different one. The IBM Storwize V7000 suggests configuring all available drives, based on recommended values for the RAID level and drive class. However, you can select a different configuration from presets by using the wizard.

Figure 3-28 illustrates how to define the array. To understand the complete procedure for how to configure internal storage, see 7.1.2, “Configuring internal storage” on page 187.

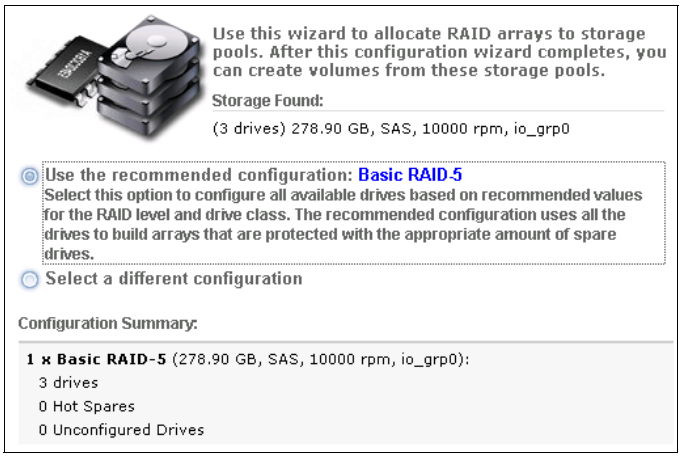

Figure 3-28 Configuring internal storage

For the best performance of your system, always configure and use all disk drives in the IBM Storwize V7000, but with respect to a recommended number of spare drives for the array rebuild process, in case of drive failure.

3.4.3 External storage

Clicking the **External Storage** option opens the window as shown in Figure 3-29. It provides the list of all externally connected (SAN-attached) disk systems to the IBM Storwize V7000.

When the new external storage system is properly zoned to the IBM Storwize V7000, run the Detect MDisks procedure either from the Actions menu in the table header or by right-clicking any of the existing MDisks in the list (Figure 3-29).

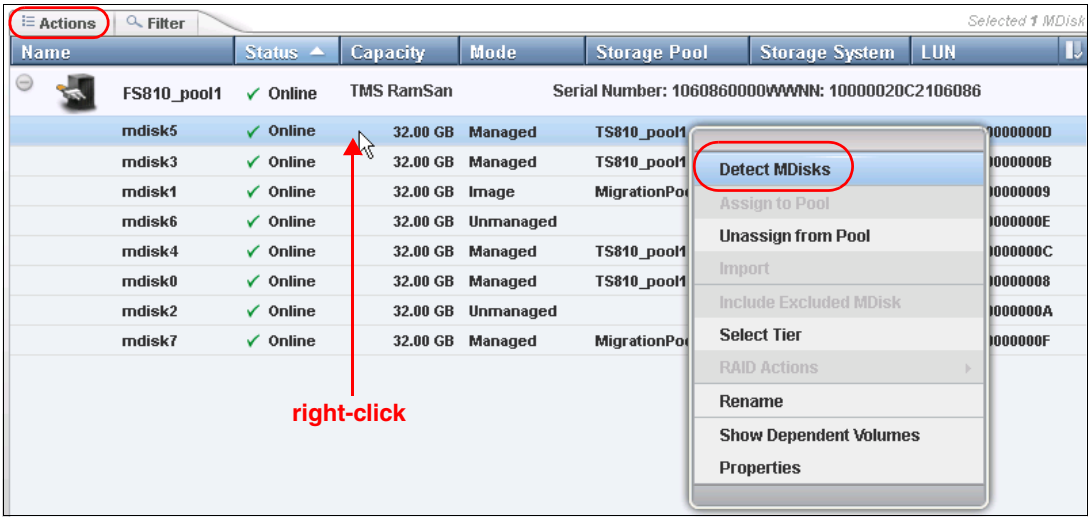

Figure 3-29 Detecting external storage systems

A new storage controller (external storage system) is listed automatically when the SAN zoning configuration is done, but typically without detected disk drives (Figure 3-30).

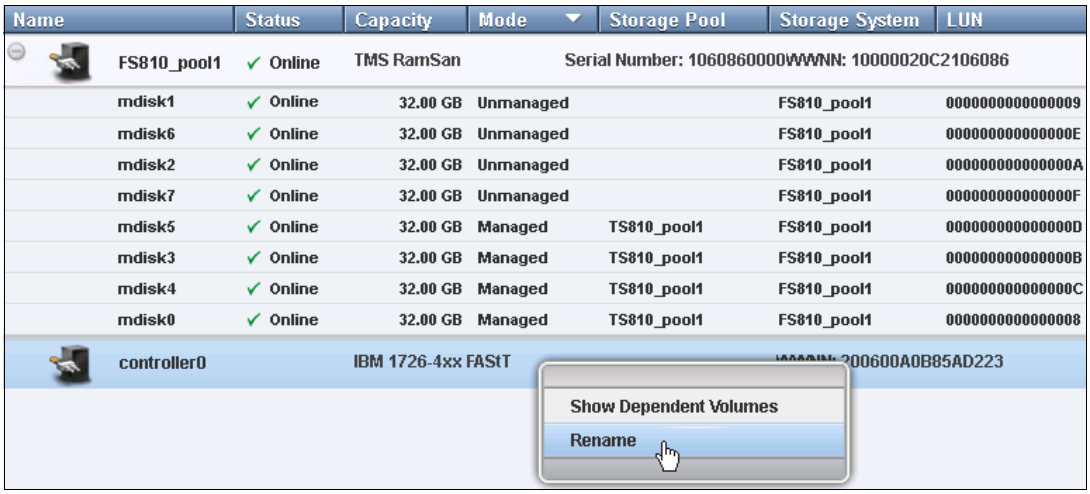

Figure 3-30 Automatically detected new external storage system

By right-clicking a newly detected storage system, you can rename the controller's default name to reflect the real type of the storage device, in our case the IBM System Storage DS3400. We suggest using a simple naming convention, in our case DS3400 (Figure 3-31).

Figure 3-31 Renaming detected system

When the new external storage system has the appropriate name, detect all disks configured on that external storage, in our case IBM System Storage DS3400. Because this IBM entry-level disk storage subsystem is widthdrawn from marketing, it shows a good example of a migration scenario from an old storage system to your new IBM Storwize V7000 system. You can also detect new MDisks from the CLI by using the **svctask detectmdisk** command or just **detectmdisk**. Figure 3-32 shows details about detected managed disks.

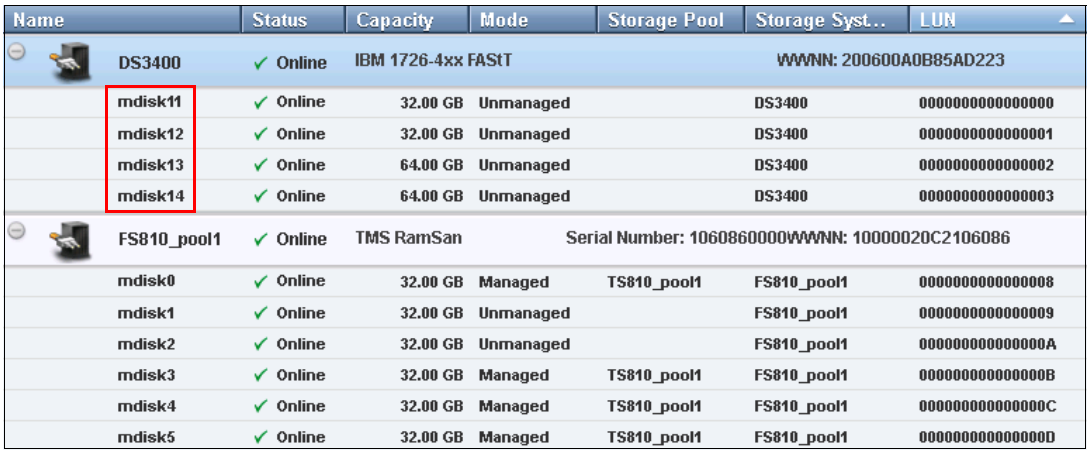

Figure 3-32 Newly discovered managed disks

All newly discovered disks are always interpreted in an *unmanaged* mode. You must assign them to the specific pool to be able to operate them.

Attention: The managed disks (MDisks) are not physical disk drives, but storage arrays configured on external systems.

If you add a managed disk that contains existing data to a managed disk group, you will lose the data that it contains. The *image* mode is the only mode that will preserve this data.

3.4.4 MDisks by Pools

This option on the Pools menu provides the list of all managed disks and arrays of disks either internally or externally connected and associated with one of the defined pools. It also lists all unassigned MDisks separately. An example of the panel is shown in Figure 3-33.

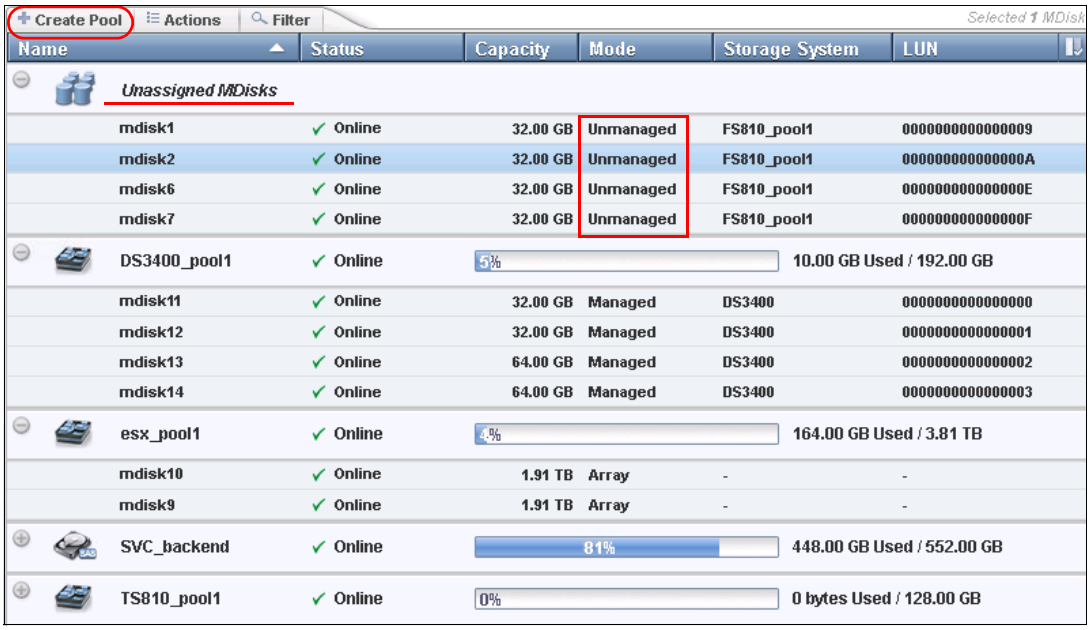

Figure 3-33 List of managed disks sorted within pools

All disks that are not yet assigned to any pool are listed in the Unassigned MDisks section. This section is always at the top of the list, even if you sort the list by pool name (clicking the **Name** header of the table). Right-click a specific disk to open a context menu, where you can assign selected unmanaged disks to the pool.

From the same panel, you can define a new storage pool by using the Create Pool button in the top left corner of the table (highlighted in Figure 3-33). The wizard window opens and you need to specify pool parameters such a Pool Name, Extent Size, and Warning Threshold. You can directly select Unmanaged MDisks that you want to include in the pool, or just skip this task and add MDisks later.

For detailed instructions for how to define pools and assign MDisks, follow the guidance in 7.2.1, “Adding MDisks to storage pools” on page 199.

Note: All sort functions in the header of the table apply to MDisks *within* pools. You cannot sort volumes based on specific criteria *across* all pools.

3.4.5 System Migration

Migrating data from older storage systems to the IBM Storwize V7000 storage system allows applications to benefit from the new features, such as Easy Tier, an intuitive management GUI, and advanced storage replication functions that better support applications.

To migrate existing data, use the IBM Storwize V7000 storage migration wizard to guide you through the procedure. This wizard is available by selecting **Pools** → **System Migration** as shown in Figure 3-34.

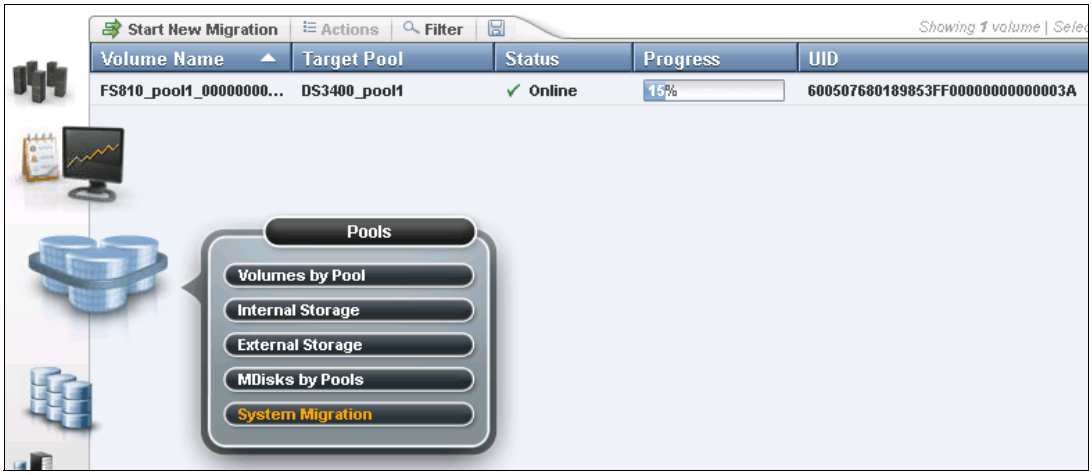

Figure 3-34 System migration

The migration of external volumes into the IBM Storwize V7000 system is one of the key benefits and features of external storage virtualization provided by this product. Therefore, we dedicate the whole chapter to this topic. See Chapter 6, “Migration wizard” on page 167 for detailed steps of the migration process.

3.5 Volumes menu

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 storage pool.

The three types of volumes are striped, sequential, and image. You can also configure volumes with specific advanced features:

Compressed This is a special type of volume where data is compressed as it is written to disk, saving additional space. To use the compression function, you must obtain the IBM Real-time Compression license.

Mirrored By using volume mirroring, a volume can have two physical copies. Each volume copy can belong to a different storage 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.

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

Virtual capacity is the volume storage capacity that is available to a host. Real capacity is the storage capacity that is allocated to a volume copy from a storage pool. In a fully allocated volume, the virtual capacity and real capacity are the same. In a thin-provisioned volume, however, the virtual capacity can be much larger than the real capacity.

Hover the cursor over the Volumes function icon to display the Volumes menu options (Figure 3-35).

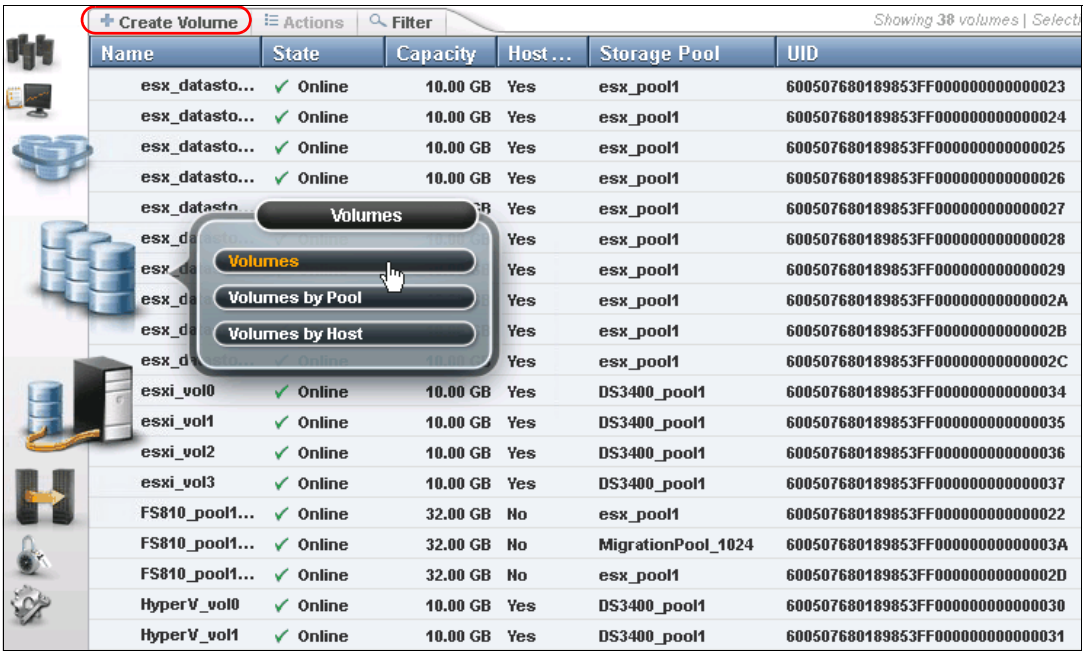

Figure 3-35 Navigate to the Volumes menu

3.5.1 All volumes

Select **Volumes** as shown in Figure 3-35. A list of all defined volumes, alphabetically sorted by the volume name (by default), is displayed. At any time you can change the sort options by clicking a specific header in the table. You can directly configure a new volume by clicking **Create Volume** (highlighted in Figure 3-35). The wizard opens and the list of volume types is displayed (Figure 3-36).

Figure 3-36 Volume types

The description of each type of volume and procedures for how to effectively create these volumes are described in Chapter 5, “Basic volume configuration” on page 131.

In addition to the new volume creation, other direct volume functions are available, such as mapping and unmapping volumes to hosts; renaming, shrinking, or expanding existing volumes; migration to a different pool; or defining a volume copy. All of these tasks are available when you select a specific volume and click **Actions** (Figure 3-37).

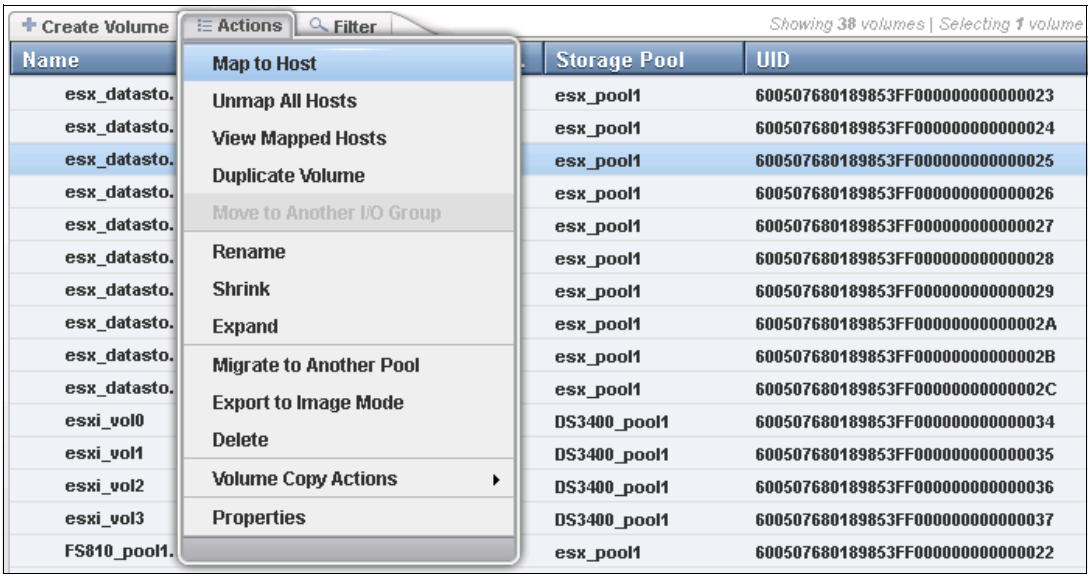

Figure 3-37 Actions menu for volumes

When moving a volume to another I/O group (a different IBM Storwize V7000 system or pair of IBM SAN Volume Controller nodes), be sure that the appropriate host zoning is in place. The target host must have access to both systems: source and target).

3.5.2 Volumes by Pool

This menu is identical to the one described in 3.4.1, “Volumes by pool” on page 80 in a section Pools. See that section for details.

3.5.3 Volumes by Host

Click **Volumes by Host** to open the window shown in Figure 3-38 on page 88. This window shows the volumes that have been mapped to a given host. You can perform the same actions with volumes as in all previous views, either by clicking **Actions** or by using the context menu that opens after you right-click a specific volume. See Figure 3-37 for details.

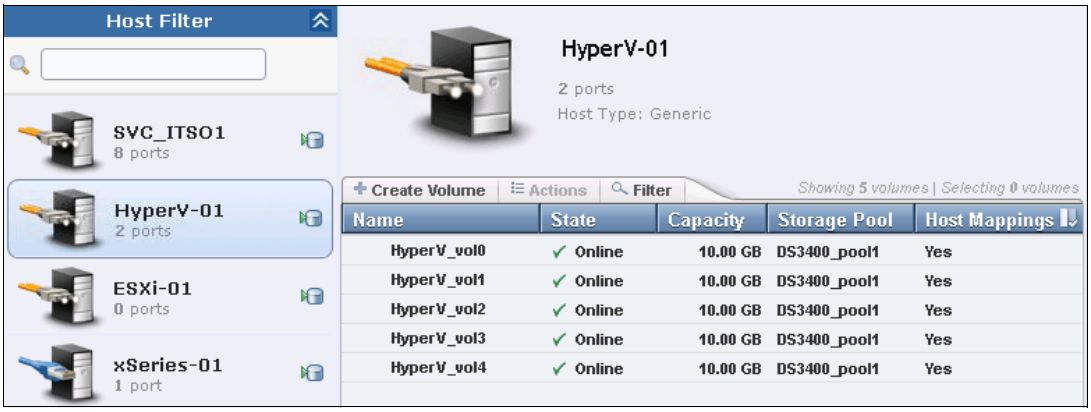

Figure 3-38 Listing volumes by host

In addition, you can rename the host by clicking its name and directly typing the new name into the context line (Figure 3-39).

Be careful: There is no confirmation button and the host is renamed immediately when you click somewhere else or press Enter

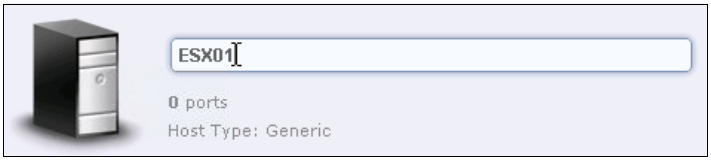

Figure 3-39 Renaming host

3.6 Hosts menu

In a storage area network (SAN) environment, host systems are application servers that access data from the storage controllers that are connected to the SAN. Hosts that are running in several operating environments can connect to the storage using the Storwize V7000. Host connections to the IBM Storwize V7000 are either SCSI over the Fibre Channel SAN, Fibre Channel over Ethernet SAN, or iSCSI over an Ethernet network.

You can use several tools to manage hosts, including the management GUI, the CLI, and specialized utilities for working with host bus adapters. To work with hosts in the management GUI, select **Hosts**. When you hover the cursor over the Host function icon, the Hosts menu opens, which has the following options:

- Hosts
- Ports by Host
- Host Mappings
- Volumes by Host

3.6.1 Hosts

This option provides an overview about all hosts that are connected (zoned) to the system, detected, and configured to be ready for storage allocation. This overview shows the name of the host as defined in the IBM Storwize V7000, the type of the host, its status of accessibility,

the number of ports used for host mapping, and whether host mapping is active. From the same panel, you can create a new host, rename a host, delete a host, or modify a host mapping. The output of the menu selection is shown in Figure 3-40.

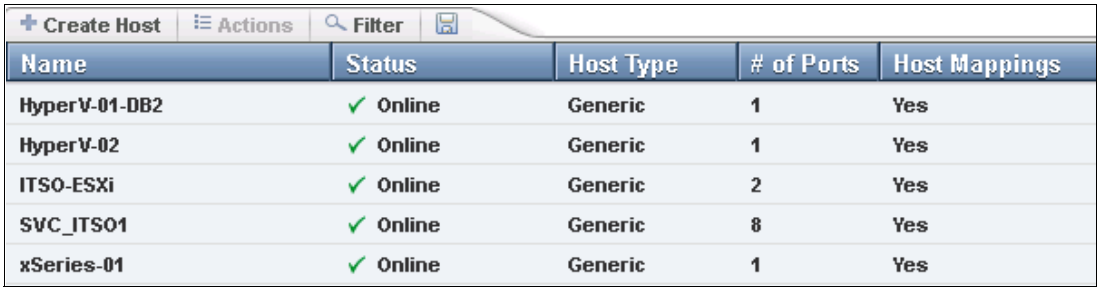

Figure 3-40 Overview of configured hosts

For example, when you click **Create Host** in a panel header, a wizard opens where you define either an iSCSI or Fibre Channel host (Figure 3-41).

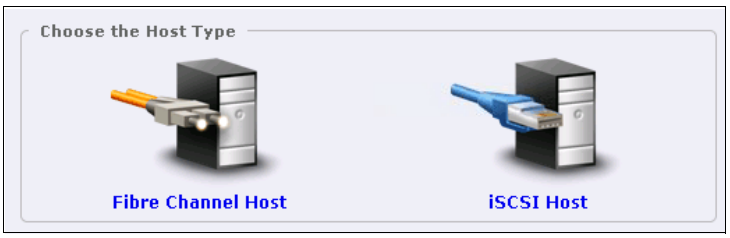

Figure 3-41 Create host wizard

To rename multiple hosts in a single step, mark all hosts you want by using the Control or Shift key, right-click, and then from the opened context menu select **Rename**. The window shown in Figure 3-42 opens.

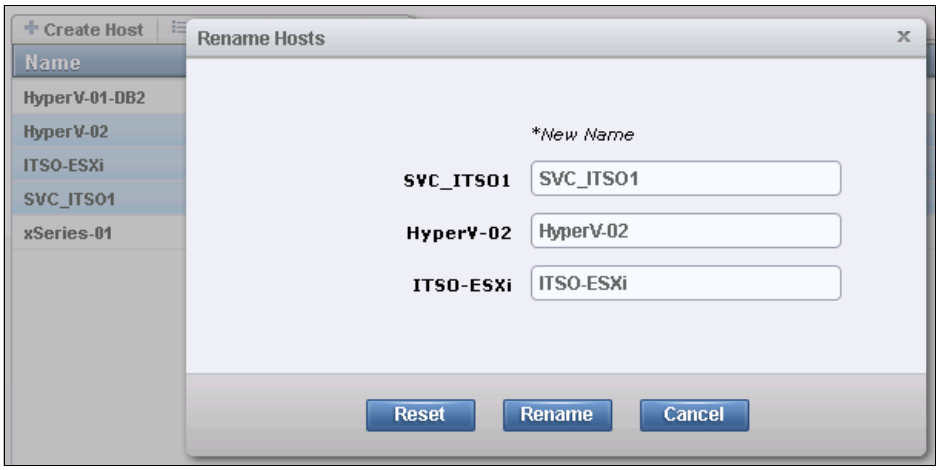

Figure 3-42 Renaming multiple hosts

Some of the actions described here are available from different menus. For example, you can select **Volumes** and its option **Volumes by Hosts**, where you can also rename hosts. This is one of the advantages of the enhanced, redesigned IBM Storwize V7000 management GUI.

3.6.2 Ports by Host

Click **Ports by Hosts** to open the panel shown in Figure 3-43. This panel lists the Fibre Channel and iSCSI ports that are assigned to a particular host.

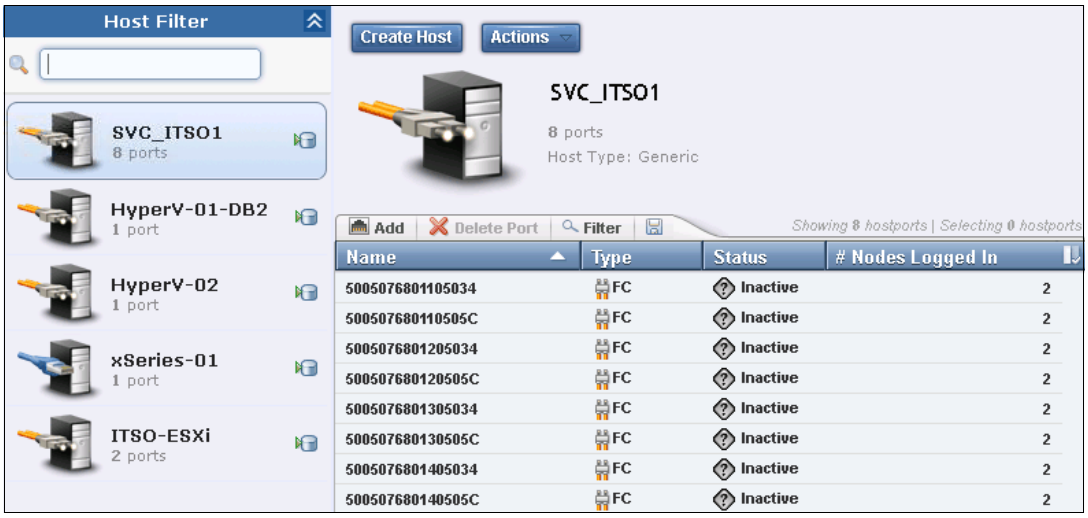

Figure 3-43 Ports by Host window

This overview shows hosts with active, inactive, or degraded ports. You can delete or add a port, or modify its characteristics. Also in this panel, you can create a new host or rename the existing one.

To do any of the tasks shown in Figure 3-44, click Actions and select a menu item.

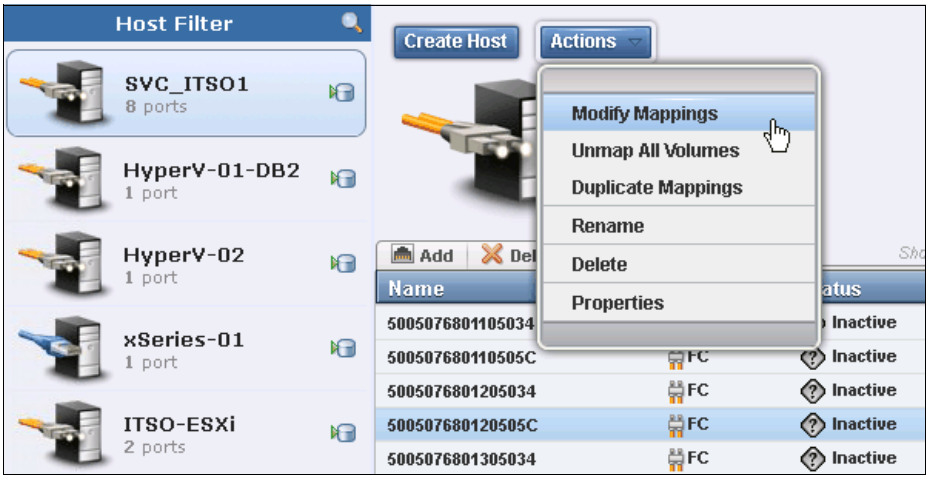

Figure 3-44 Available port actions

To delete multiple ports, select them by using Ctrl or Shift key and click **Delete**.

3.6.3 Host Mappings

Click **Host Mappings** to open the window shown in Figure 3-45. It lists the host name, SCSI identifier, volume name, and volume identifier for all mapped volumes.

Actions | Filter | Showing 33 mappings | Selecting 0 map

Host Name	SCSI ID	Volume Name	Volume Unique Identifier	Caching I/O G..
HyperV-01-DB2	1	HyperV_vol1	600507680189853FF000000000000031	0
HyperV-01-DB2	2	HyperV_vol2	600507680189853FF000000000000032	0
HyperV-01-DB2	3	HyperV_vol3	600507680189853FF000000000000033	0
HyperV-01-DB2	4	HyperV_vol4	600507680189853FF000000000000038	0
HyperV-02	3	HyperV_vol3	600507680189853FF000000000000033	0
HyperV-02	4	HyperV_vol4	600507680189853FF000000000000038	0
HyperV-02	1	HyperV_vol1	600507680189853FF000000000000031	0
HyperV-02	2	HyperV_vol2	600507680189853FF000000000000032	0
ITSO-ESXi	0	esx_datastore_3	600507680189853FF000000000000026	0
SVC_ITSO1	16	esx_datastore_2	600507680189853FF000000000000025	0
SVC_ITSO1	12	svc_lu12	600507680189853FF00000000000001D	0
SVC_ITSO1	1	svc_lu1	600507680189853FF000000000000012	0
SVC_ITSO1	18	esx_datastore_4	600507680189853FF000000000000027	0

Figure 3-45 Host mappings

From this Window you can view the host properties, obtain the list of mapped volumes, or work with port definitions. Right-click the specific host and select **Properties (Host)** from the opened context menu. A window similar to the one in Figure 3-46 opens.

Overview | Mapped Volumes | Port Definitions

Host Name	HyperV-02
Host ID	2
Status	✓ Online
Host Type	Generic
# of FC Ports	1
# of iSCSI Ports	0
# of SAS Ports	0
I/O Group	io_grp0, io_grp1, io_grp2, io_grp3
iSCSI CHAP Secret	

Edit

Figure 3-46 Host properties

You can modify the host name, host type, available I/O group, or iSCSI CHAP Secret by clicking **Edit** (highlighted in Figure 3-46).

3.6.4 Volumes by Host

This option is identical to the option that is available in the dynamic menu Volumes. For a description, see 3.5.3, “Volumes by Host” on page 87.

3.7 Copy Services

The IBM Storwize V7000 Copy Services are part of the IBM Replication Family Services available on IBM SAN Volume Controller, IBM Storwize V7000, and IBM Storwize V5000 products. It consists of the following functions:

- IBM FlashCopy
- Metro Mirror and Global Mirror
- Global Mirror with Changed Volumes
- Volume Mirroring function
- Stretched Cluster (formerly Split Cluster) volume mirroring

In this section, we briefly describe how to navigate in the Copy Services menu.

3.7.1 FlashCopy

FlashCopy is a function you can use to create a point-in-time copy of one of your IBM Storwize V7000 disks. This might be helpful when doing backups or application testing. These copies can be cascaded on one another, read from, written to, and even reversed.

FlashCopy snapshots are able to conserve storage, if needed, by being space-efficient copies (instead of full copies) that record only items that have changed from the originals. Select **FlashCopy** from the dynamic menu to open a panel similar to what is shown in Figure 3-47.

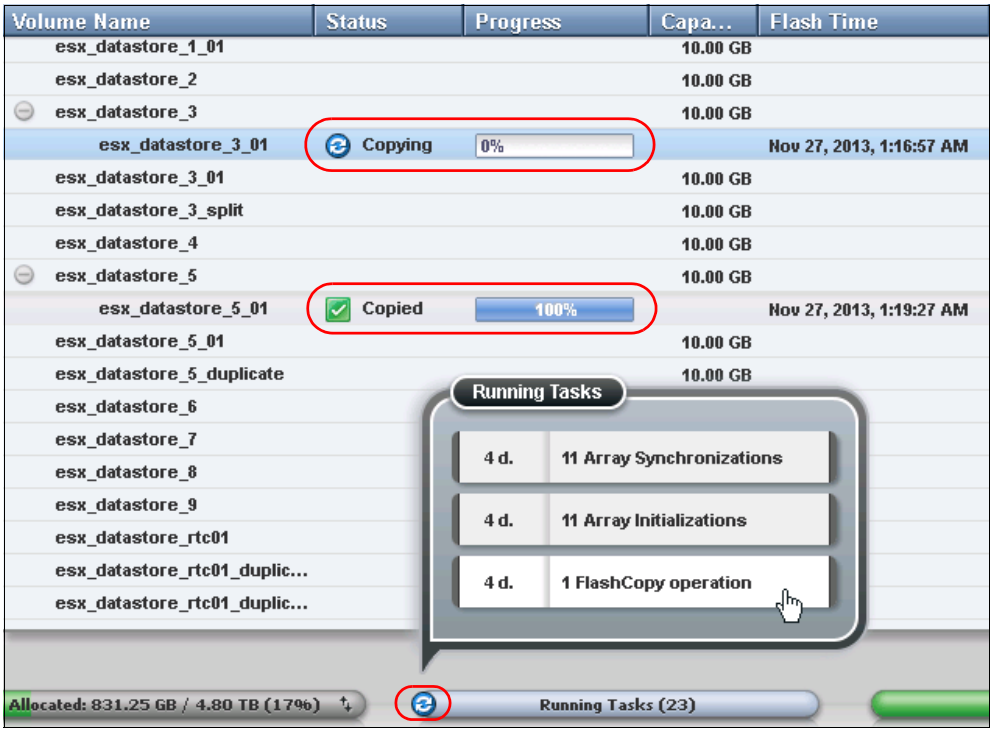

Figure 3-47 FlashCopy operations

If you need to create a FlashCopy of an additional volume, right-click the volume and the list of available functions is listed. From here, you can do tasks such as initiate a new snapshot, and clone or back up a volume.

Clicking the volume name opens the window shown in Figure 3-48. From here, you can click the tabs at the top of the window to display additional information, such as the hosts that the volume or FlashCopy volume is mapped to and its dependant MDisks.

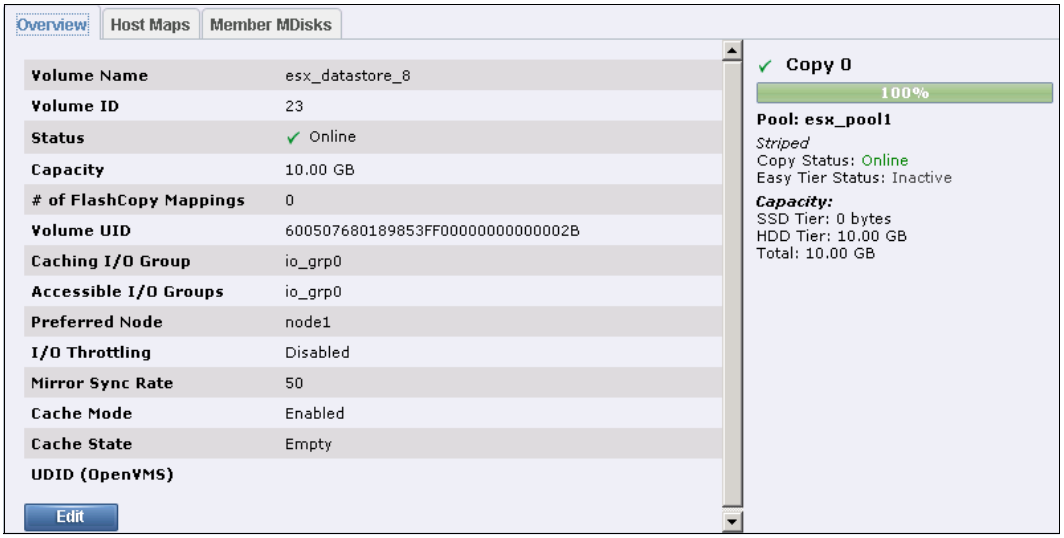

Figure 3-48 FlashCopy volume details

3.7.2 Consistency Group

Click **FlashCopy Consistency Group** to open the window shown in Figure 3-49. Here, FlashCopy relationships can be placed into a consistency group. You can also use start and stop commands against the FlashCopy consistency group from this window by right-clicking the relationship.

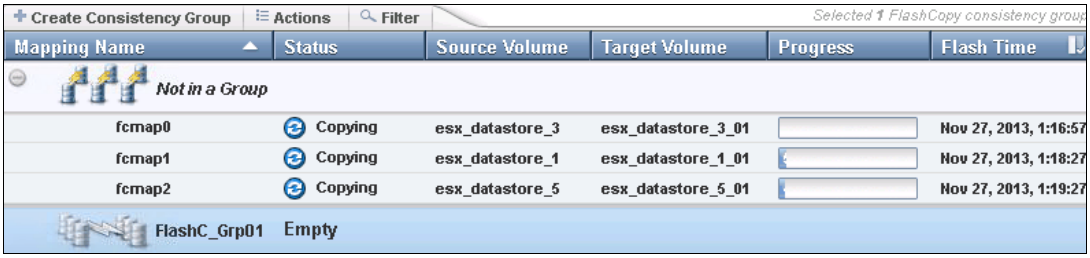

Figure 3-49 FlashCopy Consistency Group window

When any FlashCopy Consistency Group is available, either empty or with existing relations, you can move an existing relation to that group. Right-click a relation and select **Move to Consistency Group** as shown in Figure 3-50 on page 94.

Other actions on the same context menu include remove from consistency group, start (resume) or stop that FlashCopy operation, rename a target volume or FlashCopy mapping, and delete a mapping.

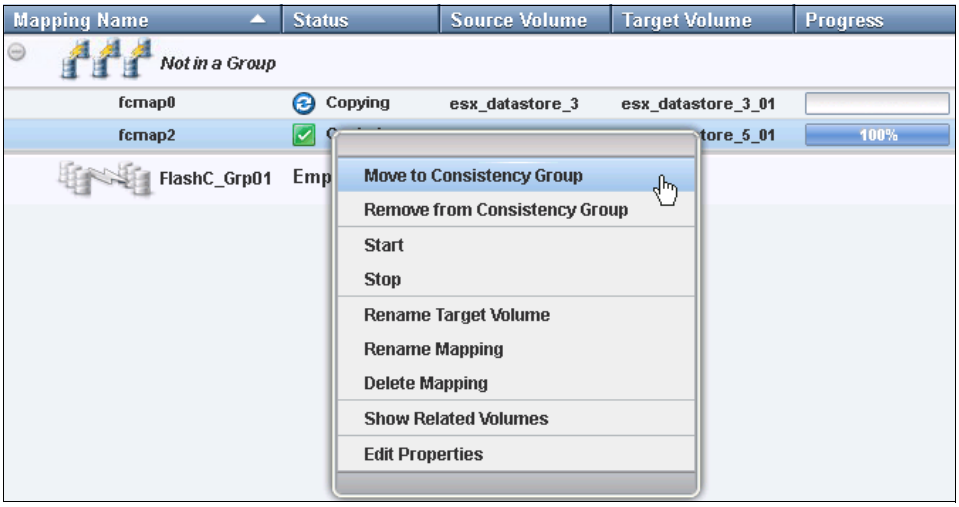

Figure 3-50 Moving relation to the Consistency Group

From the drop-down menu, select the appropriate group (in our case the only one available) and confirm the selection (Figure 3-52).

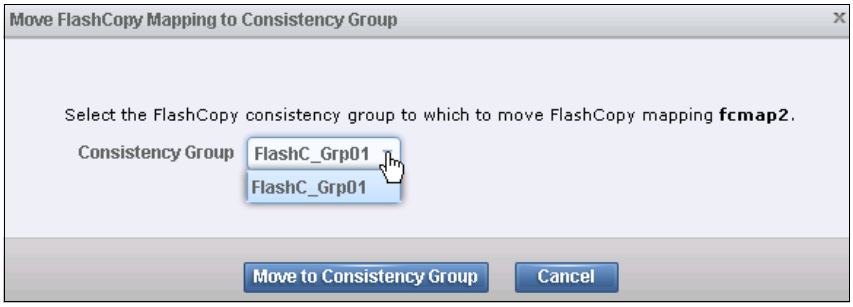

Figure 3-51 Assigning the Consistency Group

The result of the operation is similar to the result shown in Figure 3-52.

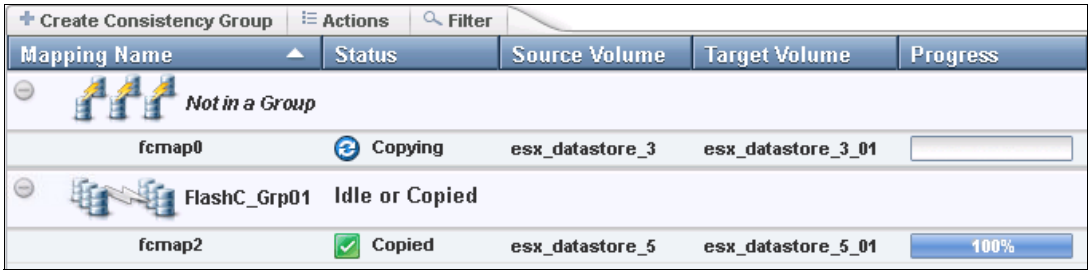

Figure 3-52 Consistency groups

3.7.3 FlashCopy Mappings

To create a new FlashCopy mapping, click Create FlashCopy Mapping (shown in Figure 3-53 on page 95) to start a wizard. This wizard maps a source volume to a target volume to prepare for a subsequent copy. This mapping persists until it is deleted. The mapping specifies the source and destination volumes. The destination must be identical in size to the source or the mapping will fail.

In a single mapping, the source and destination cannot be on the same volume. A mapping is triggered at the point in time when the copy is required. The mapping can optionally be given a name and assigned to a consistency group. These groups of mappings can be triggered at the same time, thus enabling multiple volumes to be copied at the same time. This creates a consistent copy of multiple disks. A consistent copy of multiple disks is required for database products in which the database and log files reside on separate disks.

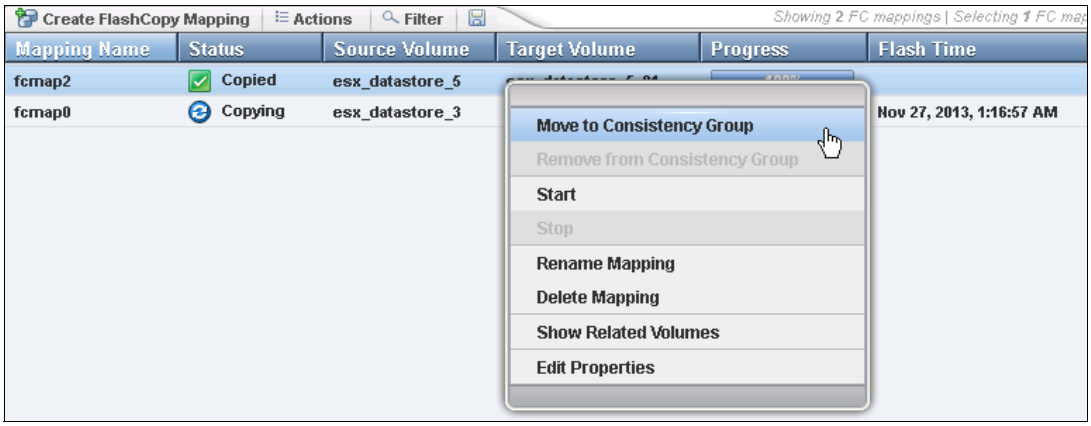

Figure 3-53 FlashCopy mappings

If any consistency group (ID or Name) is not specified, the mapping is assigned to the default *group 0*, which is a special group that cannot be started as a whole. Mappings in this group can be started only on an individual basis.

The example of the wizard for FlashCopy mapping creation is shown in Figure 3-54. Select source and target volumes from the drop-down menu of the wizard.

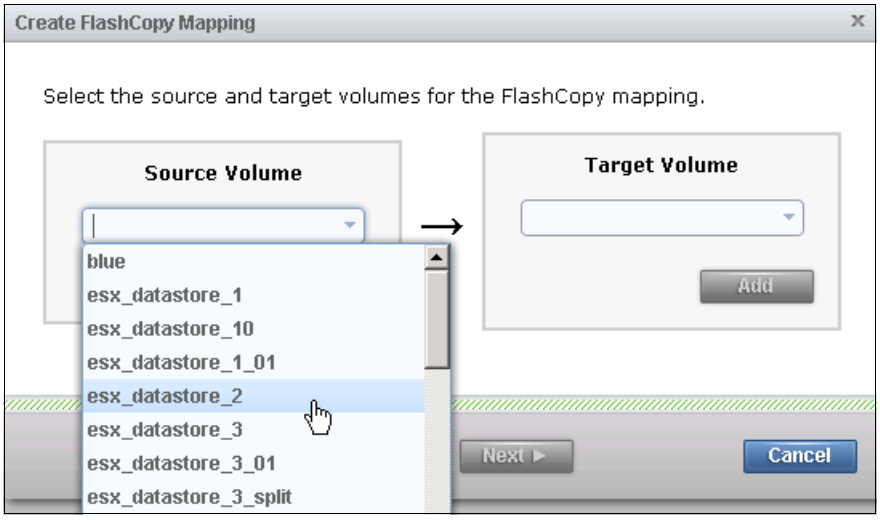

Figure 3-54 Selecting volumes for FlashCopy mappings

You can select the Snapshot (copy-on-write), Clone (replica of the volume without impact on original one), or Backup (data recovery) type of relationship. When selected, you can specify whether you also want to add the mapping to the consistency group.

3.7.4 Remote Copy

Click **Remote Copy** to open the window shown in Figure 3-55. This window shows the existing Remote Copy relationships and you can set up and modify consistency groups. From this window, you can also start and stop relationships, add relationships to a consistency group, and switch the direction of the mirror.

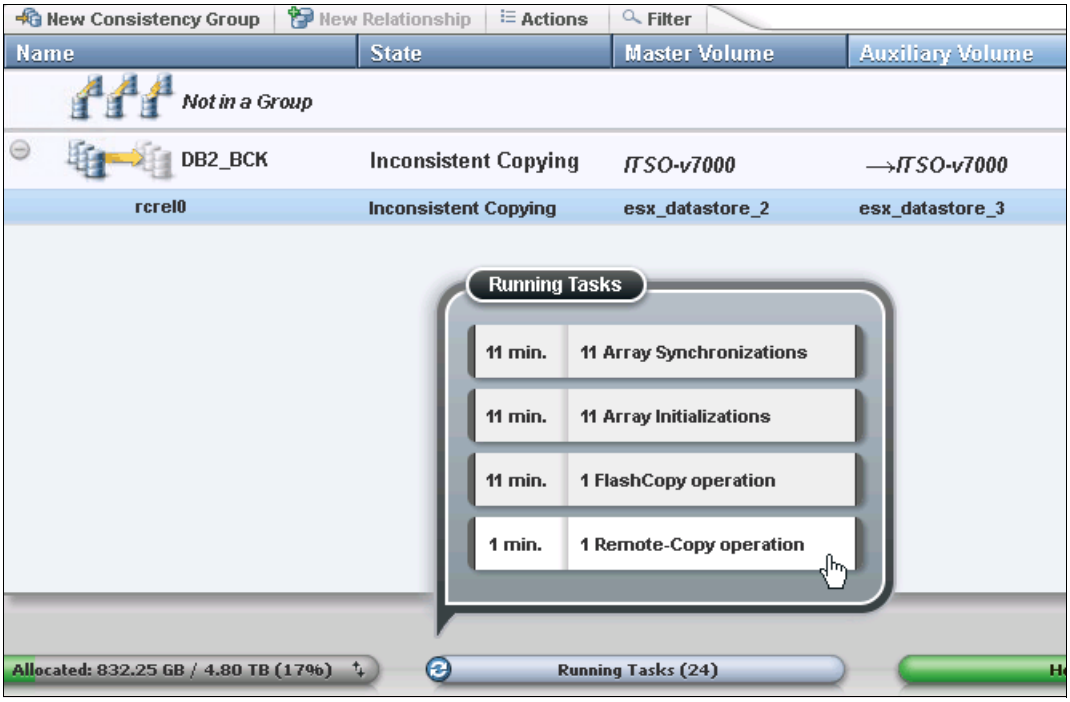

Figure 3-55 Remote Copy window

The menu provides the options to create Metro Mirror, Global Mirror, or Global Mirror with Changed Volumes.

Metro Mirror makes *synchronous* copies, which means that the original write operations are not considered complete until the write operation to the destination disk has been confirmed. The distance between your two sites is usually determined by how much latency your applications can handle.

Global Mirror makes *asynchronous* copies of your disk. This means that the write is considered complete after it is complete at the local disk; it does not wait for the write to be confirmed at the remote cluster as Metro Mirror does. This greatly reduces the latency experienced by your applications if the other cluster is far away. However, it also means that during a failure, the data on the remote copy might not have the most recent changes committed to the local disk.

Global Mirror with Changed Volumes can best be described as “Continuous Remote FlashCopy.” If you use this feature, the IBM Storwize V7000 essentially takes periodic FlashCopies of a disk and writes them to your remote destination. This feature completely isolates the local copy from wide area network (WAN) issues and from sudden spikes in workload that might occur. The drawback is that your remote copy might lag behind the original by a significant amount of data, depending on how you set up the cycle time.

3.7.5 Partnerships

Click **Partnerships** to open the window shown in Figure 3-56. You can use this window to set up a new partnership or delete an existing partnership with another IBM Storwize V7000 system for the purposes of remote mirroring.

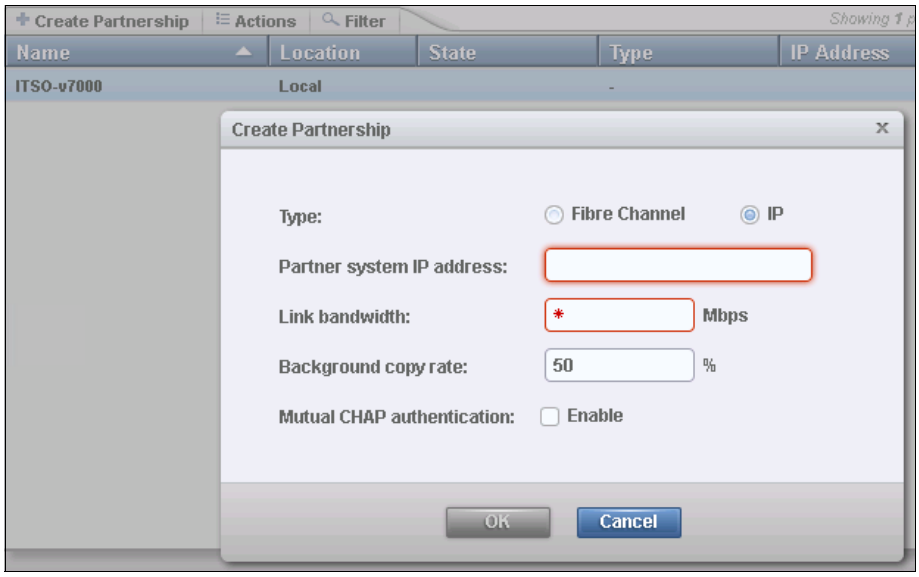

Figure 3-56 Creating the partnership

From this window, you can also set the background copy rate. This rate specifies the bandwidth, in megabytes per second (MBps), that is used by the background copy process between the clusters (Figure 3-57).

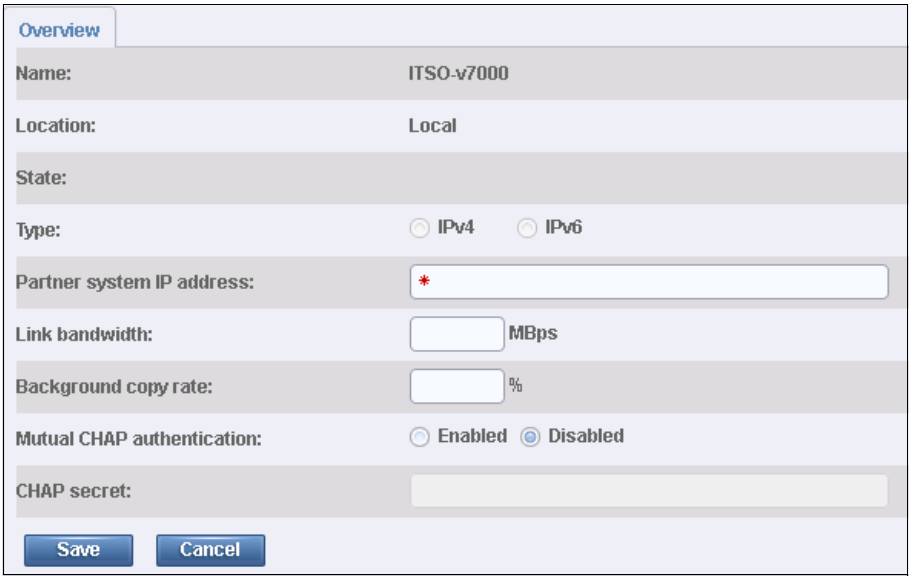

Figure 3-57 Partnership properties

3.8 Access menu

The Access menu has two options: Users (for user management) and Audit Log.

3.8.1 Users

Figure 3-58 shows the Users panel. From here, you can create and delete new users, change and remove passwords, and add and remove SSH keys.

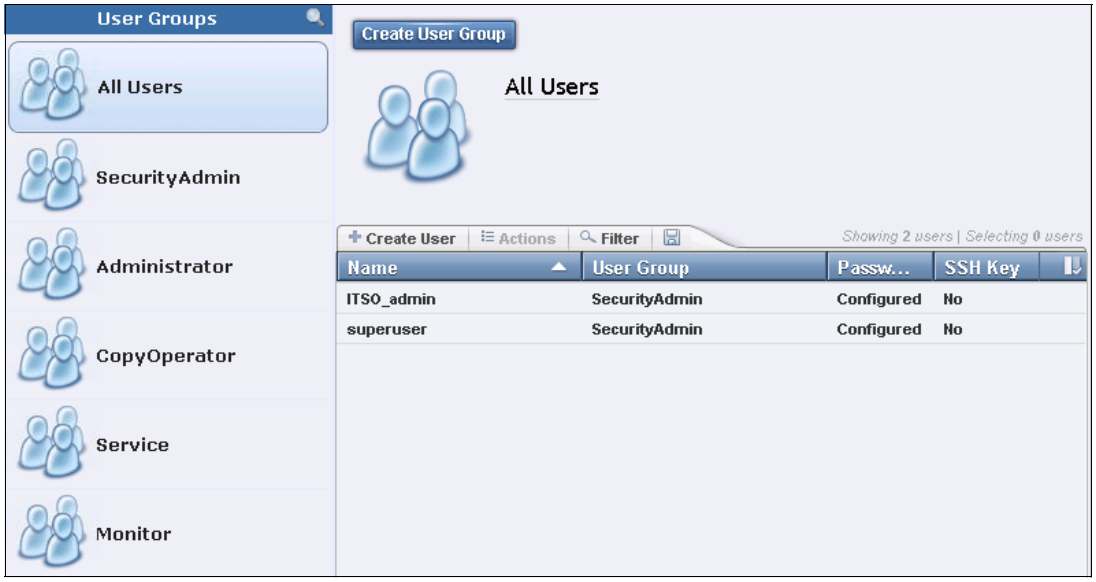

Figure 3-58 Users window

Click **New User** to open the panel shown in Figure 3-59. Use this panel to specify the name and password of the user, and load the SSH key (if the SSH key has been generated). Starting with SAN Volume Controller V6.3, the SSH key is not required for CLI access, and you can choose to use either SSH or a password for CLI authentication.

Figure 3-59 Adding a new user

3.8.2 Audit Log menu

Click **Audit Log** to open the window shown in Figure 3-60. The cluster maintains an audit log of successfully executed commands, indicating which users performed particular actions at certain times.

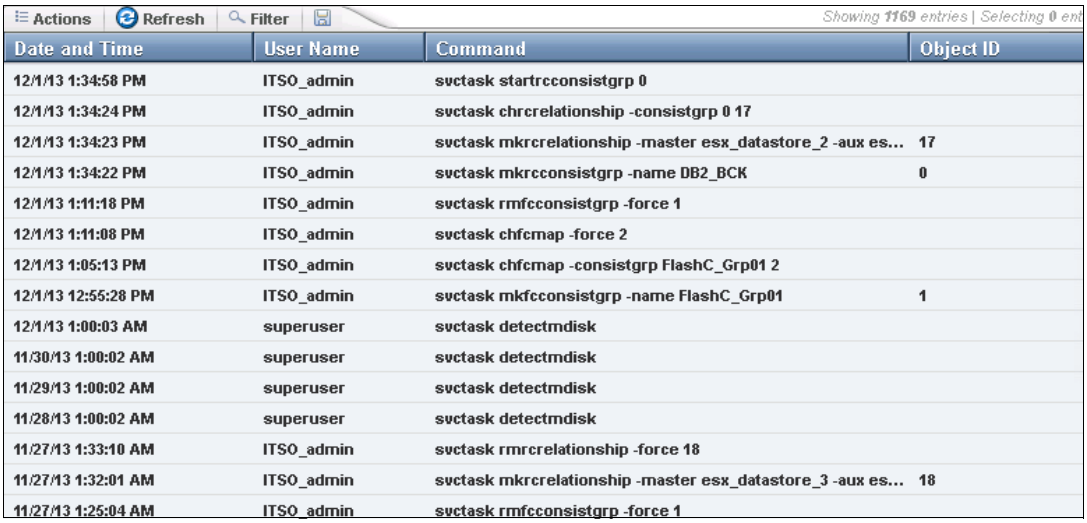

Figure 3-60 Audit Log entries

You can filter audit log records by date or within a specific time frame, defined either by minutes, hours, or days (Figure 3-61).

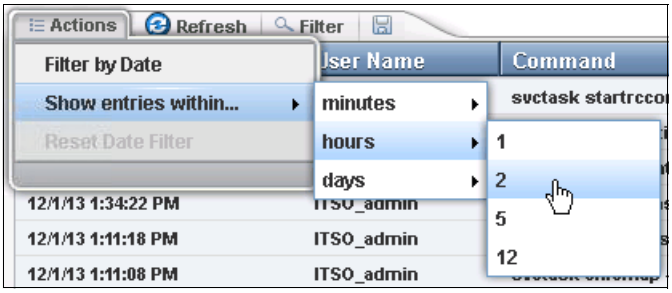

Figure 3-61 Filtering the records

The following commands are not recorded in the audit log:

- All commands that failed
- **dumpconfig**
- **cpdumps**
- **cleardumps**
- **finderr**
- **dumperrlog**
- **dumpinternallog**
- **svcservicetask dumperrlog**
- **svcservicetask finderr**

3.9 Settings menu

The settings menu provides various options to adjust your system parameters according to your needs. You can configure these options:

- Event notifications
- Directory services
- Network
- Support
- General options

These options are described in the following sections. Details for how to use each of them are presented in Chapter 13, “RAS, monitoring, and troubleshooting” on page 455.

3.9.1 Event notifications

It is important to correct any issues reported by your IBM Storwize V7000 system as soon as possible. Configure your system to send automatic notifications when a new event is reported. To avoid monitoring for new events that use the management GUI, select the type of event that you want to be notified about. For example, restrict notifications to just events that require immediate action.

You can use email, SNMP, or Syslog types of notifications. If your system is within warranty, or you have a hardware maintenance agreement, configure your IBM Storwize V7000 system to send email events to IBM directly, if an issue that requires hardware replacement is detected. This mechanism is called *Call Home*. When an event is received, IBM automatically opens a problem report, and, if appropriate, contacts you to verify whether replacement parts are required. The configuration window for email notifications is shown in Figure 3-62.

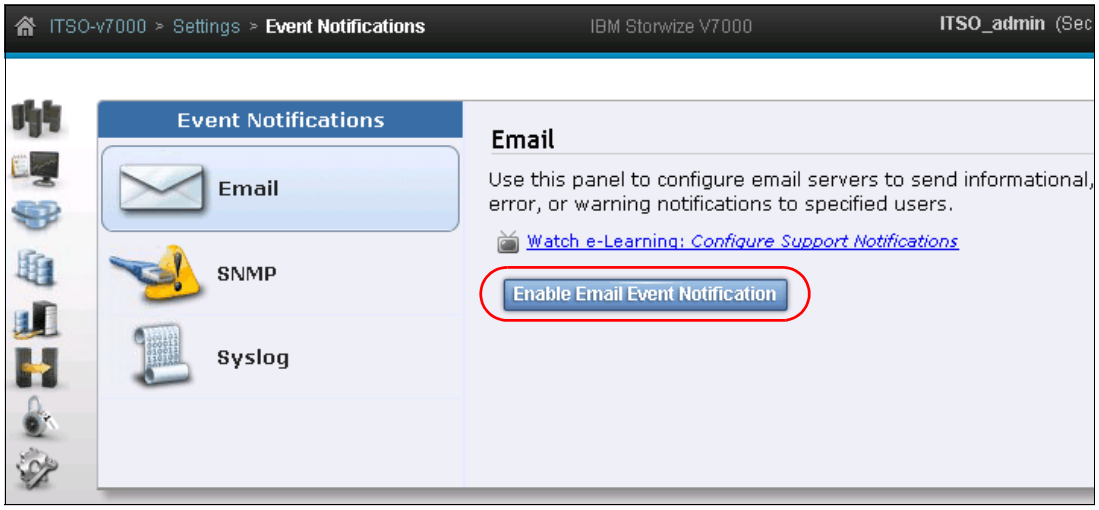

Figure 3-62 Email event notifications

The procedure for how to enable email notifications is described in 13.7.1, “Email notifications and Call Home” on page 482.

3.9.2 Directory services

Directory services allow the user to remotely authenticate to the IBM Storwize V7000 without a need for a local account. This means that when you log on, you authenticate with your domain user ID and password instead of a locally created user ID and password.

The benefits of remote authentication are as follows:

- You do not have to configure every user on every SAN Volume Controller or Storwize V7000. If you have multiple machines, you can more efficiently set up authentication.
- When commands are executed on the SAN Volume Controller or Storwize V7000, the audit log will show the domain user name that issued that command, instead of a local user name, or worse just superuser (that is, determining who mapped a volume, acted as the superuser, and so on, might be difficult).
- You have central control over access. If someone leaves the company, you just remove access at the domain controller, meaning that orphan accounts will not remain on your storage equipment.

The access panel to configure remote authentication is shown in Figure 3-63.

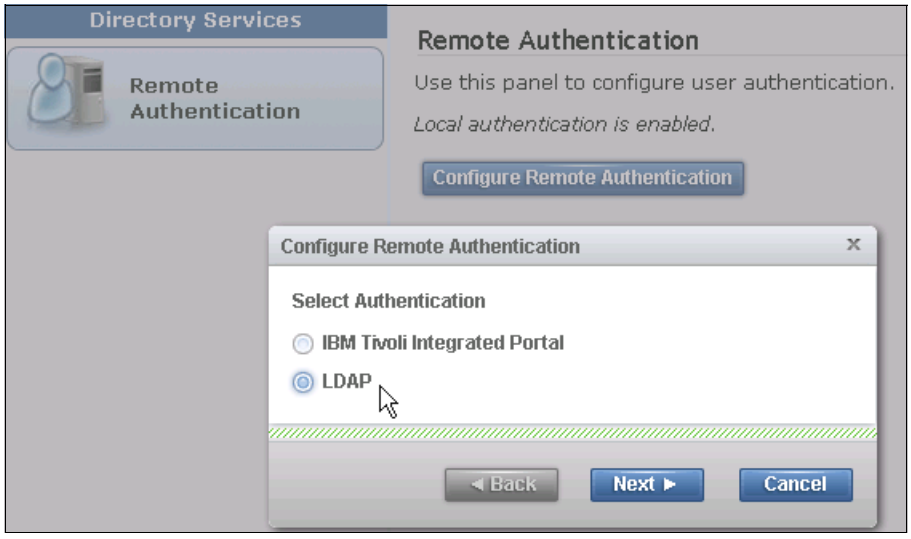

Figure 3-63 Configuring remote authentication

The detailed steps to configure remote logon are described at the following addresses:

- https://www.ibm.com/developerworks/community/blogs/anthonyv/entry/svc_and_storwize_v7000_release_6_3_configuring_ldap1?lang=en
- http://www-01.ibm.com/support/knowledgecenter/ST3FR7_7.2.0/com.ibm.storwize.v7000.720.learning/authentication_launch.html?cp=ST3FR7%2F1-2-10&lang=en

3.9.3 Network menu

Click **Network** to open the window shown in Figure 3-64. From here you can update the network configuration, set up iSCSI definitions, and view information about the Fibre Channel connections.

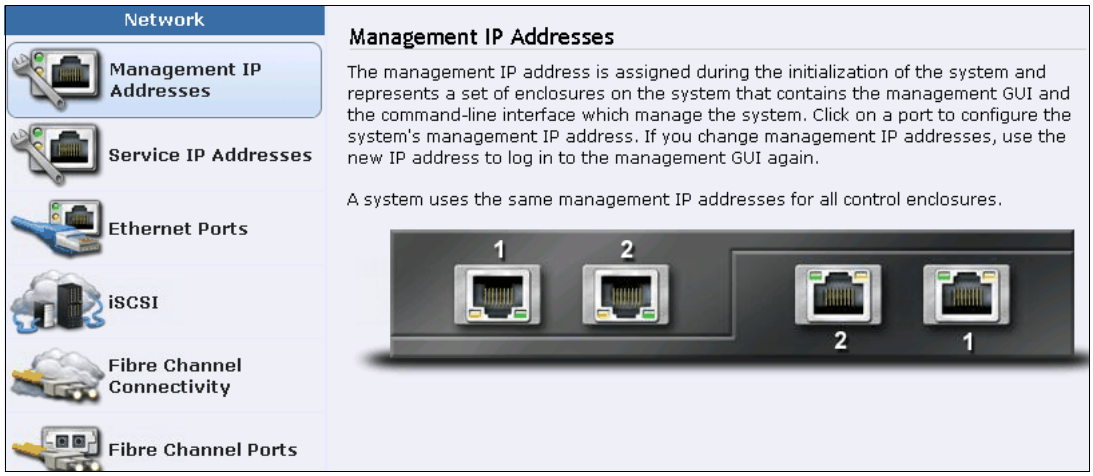

Figure 3-64 Network window

When you click **Fibre Channel Connectivity** (Figure 3-65), useful information is displayed. In this example, we click **Hosts** from the menu and then select to display the details for one specific host, Hurricane, from the list of host systems. Other options that are available from the menu include displaying Fibre Channel details for all devices, for clusters, for nodes, for storage systems, or for hosts.

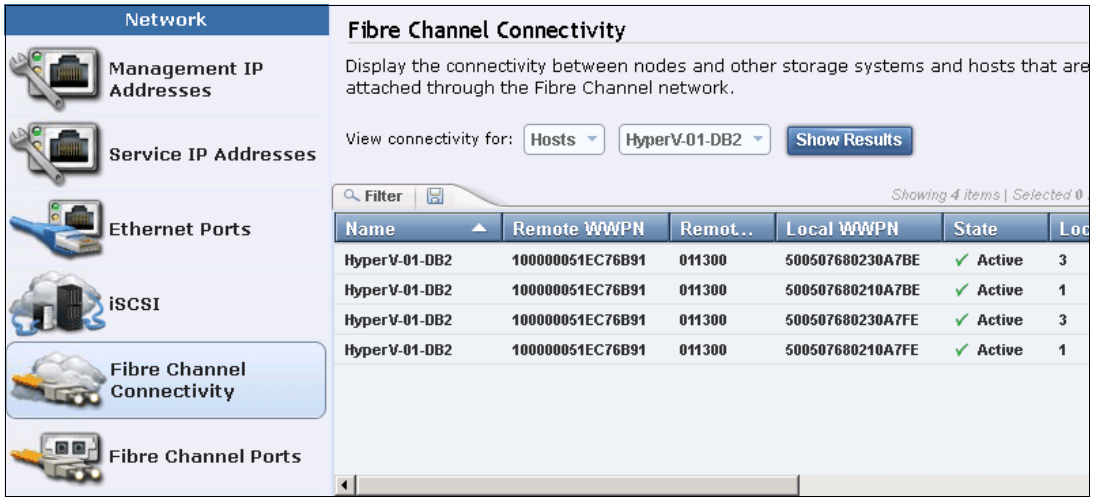

Figure 3-65 Fibre Channel connectivity

3.9.4 Support

The support menu provides access to features that help BM to analyze the critical and uncertain issues with the system, based on downloaded and provided support packages. Figure 3-66 shows how to download support logs.

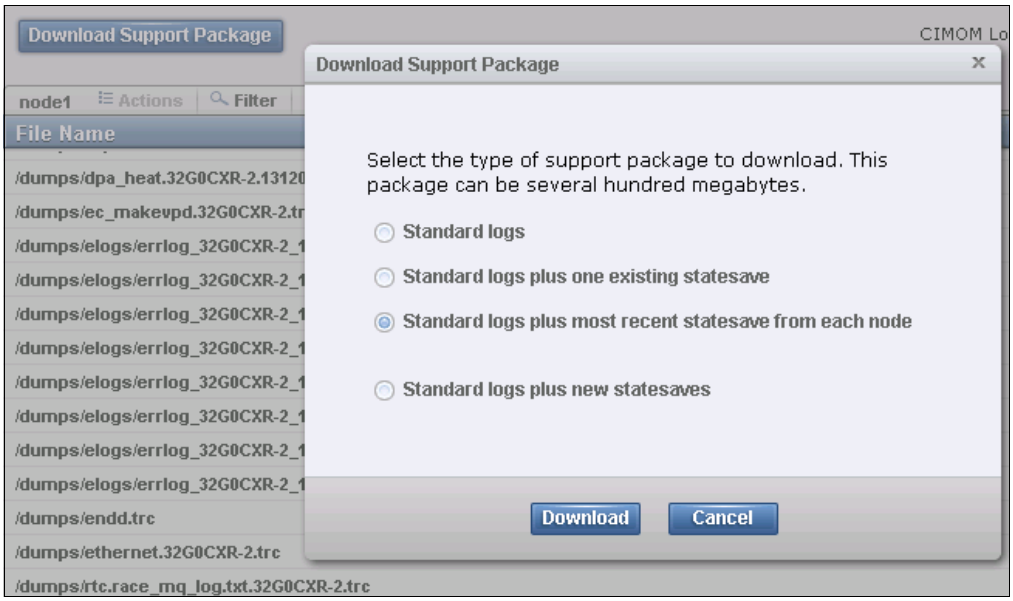

Figure 3-66 Downloading support packages

More details are available in 13.9, “Collecting support information” on page 490.

3.9.5 General

The General window provides four key options:

- Set system date and time
- Manage licenses
- Upgrade software
- Set GUI preferences

The General window opens (Figure 3-67) when you select General from the Settings menu.

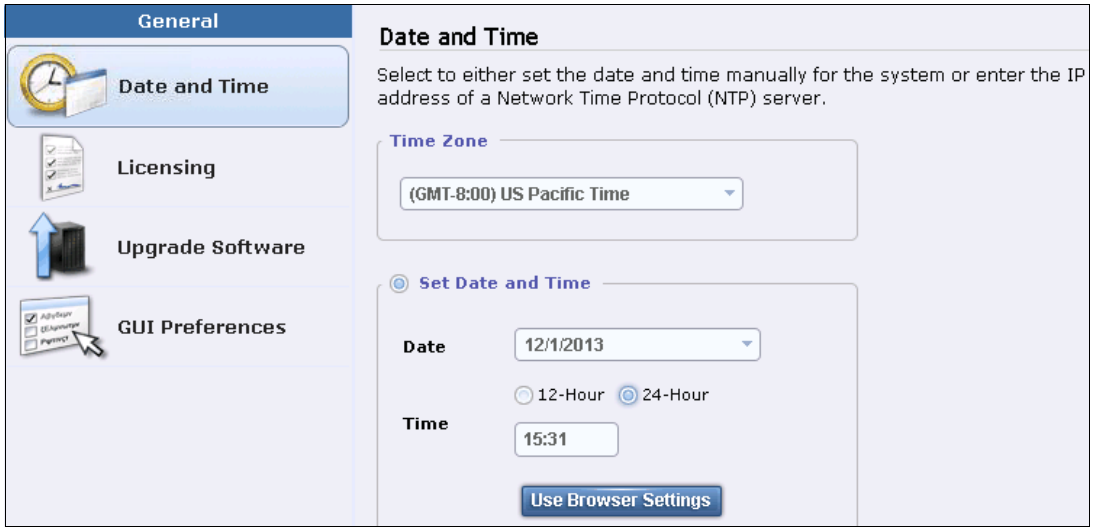

Figure 3-67 General menu

You can also update the license information for specific features as shown in Figure 3-68.

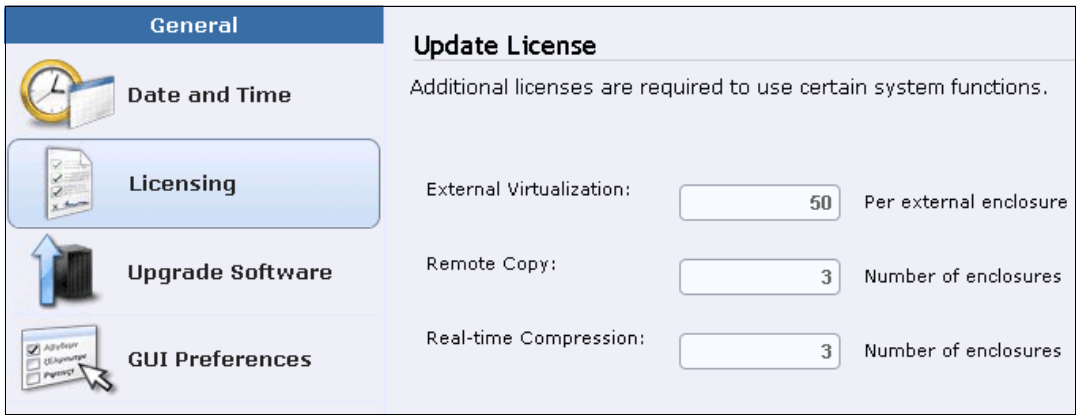

Figure 3-68 Licensing options

To upgrade your IBM Storwize V7000 software, use the procedure that is described in 13.3, “Software upgrade” on page 465.

4

Host configuration

This chapter describes how to use the IBM Storwize V7000 GUI to create hosts, and how to prepare a host to access the volumes that are created. (Volume creation is described in Chapter 5, “Basic volume configuration” on page 131.)

4.1 Host attachment overview

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

This chapter describes the following topics:

- Preparing the host operating system
  - Windows
    - Fibre Channel (FC)
    - iSCSI
  - VMware
    - Fibre Channel
    - iSCSI
- Creating hosts using the Storwize V7000 GUI
  - Creating FC hosts
  - Creating iSCSI hosts

In this chapter, we assume that your hosts are connected to your FC or IP network and you have completed the steps described in Chapter 2, “Initial configuration” on page 35. Follow basic zoning recommendations to ensure that each host has at least two network adapters, that each adapter is on a separate network (or at minimum in a separate zone), and is connected to both canisters. This setup assures four paths for failover and failback purposes.

Before mapping the newly created volumes on the host of your choice, a little preparation goes a long way towards ease of use and reliability. There are several steps required on a host in preparation for mapping new IBM Storwize V7000 volumes to the host. Use the IBM System Storage Interoperation Center (SSIC) to check which code levels are supported to attach your host to your storage. SSIC is a web tool that checks the interoperation of host, storage, switches, and multipathing drivers. It can be found at the following address:

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

The complete support matrix is listed in the *IBM Storwize V7000 Supported Hardware List, Device Driver, Firmware, and Recommended Software Levels V7.2* document, which is available at the following address:

`http://www.ibm.com/support/docview.wss?uid=ssg1S1004450`

This chapter focuses on Windows and VMware. If you must attach any other hosts, for example, IBM AIX, Linux, or even an Apple system, you can find the required information at the following address:

`http://publib.boulder.ibm.com/infocenter/storwize/ic/index.jsp`

4.2 Preparing the host operating system

In this section, we describe how to prepare Windows and VMware hosts for attachment to an IBM Storwize V7000, using either Fibre Channel or iSCSI to connect.

4.2.1 Windows 2008 R2 and 2012: Preparing for Fibre Channel attachment

Complete the following steps to prepare a Windows 2008 R2 or Windows 2012 host to connect to an IBM Storwize V7000 using Fibre Channel:

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

Downloading and installing the supported drivers and firmware

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

The latest supported levels are at the following addresses:

- For Windows 2008 R2:

  `http://www.ibm.com/support/docview.wss?uid=ssg1S1004450#_Win2008`

- For Windows 2012:

  `http://www.ibm.com/support/docview.wss?uid=ssg1S1004450#_Win2012FC`

Configuring Brocade HBAs for Windows

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

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

Configuring QLogic HBAs for Windows

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

After installing the device driver and firmware, you must configure the HBAs. To perform this task, either use the QLogic Sansurfer software or reboot into the HBA BIOS, load the adapter defaults, and set the following values:

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

Configuring Emulex HBAs for Windows

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

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

Setting the Windows timeout value

For Windows hosts, the disk I/O timeout value should be set to 60 seconds. To verify this setting, complete the following steps:

1. Click **Start** → **Run**.
2. In the dialog box, type regedit and press Enter.
3. In the registry editor, locate the `HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Disk\TimeOutValue` key.
4. Confirm that the value for the key is 60 (decimal value), and, if necessary, change the value to 60 (Figure 4-1).

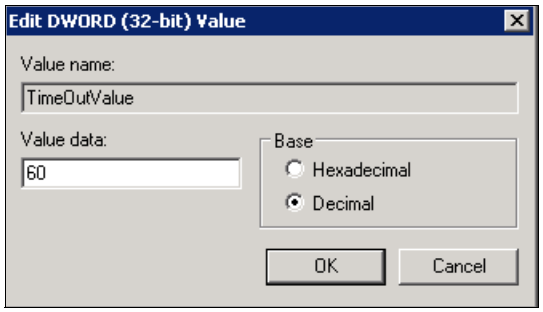

Figure 4-1 Windows timeout value

Installing the multipathing software

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

IBM Subsystem Device Driver DSM (SDDDSM) is the IBM multipath I/O solution that is based on Microsoft MPIO technology. It is a device-specific module that supports IBM storage devices on Windows hosts. The intent of MPIO is to provide better integration of a multipath storage solution with the operating system, and it allows the use of multipath in the SAN infrastructure during the boot process for SAN Boot hosts.

To ensure correct multipathing with IBM Storwize V7000, SDDDSM must be installed on Windows hosts. To install SDDDSM, complete the following steps:

1. Go to the following SDDDSM download matrix to determine the correct level of SDDDSM to install for Windows 2008 R2 or Windows 2012, and then download the package:

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

2. Extract the package to your hard disk drive and run **setup.exe** to install SDDDSM. A command prompt window opens (Figure 4-2). Confirm the installation by entering Y.

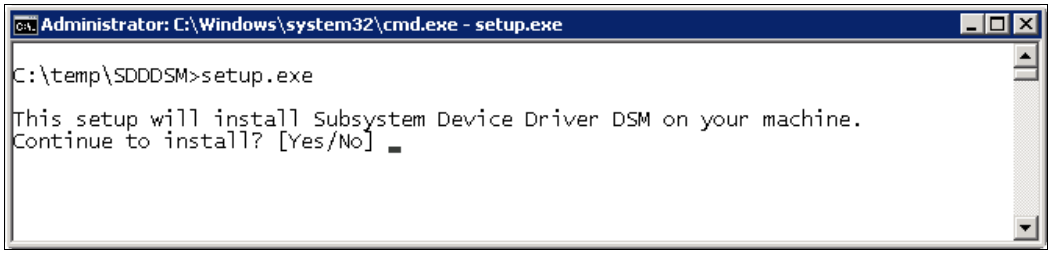

Figure 4-2 SDDDSM setup

3. After the setup completes, you are prompted to restart the system. Confirm this restart by typing yes and pressing Enter (Figure 4-3).

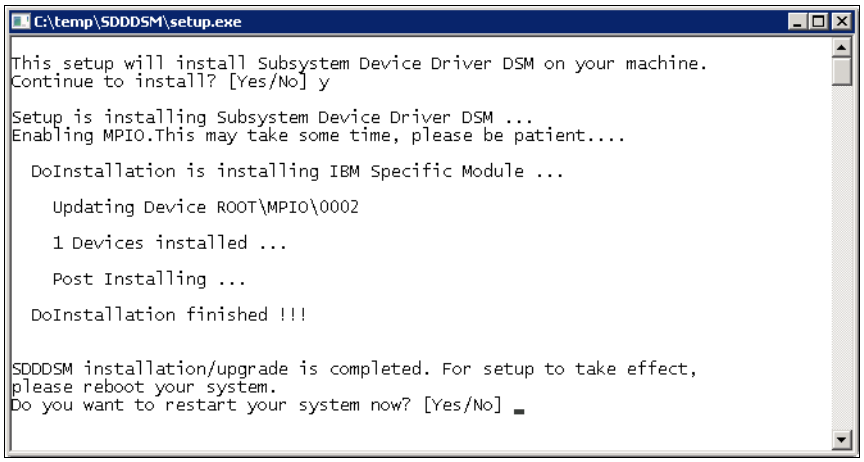

Figure 4-3 Answer yes to restart the host

You have now successfully installed IBM SDDDSM. To check the installed driver version, select **Start** → **All Programs** → **Subsystem Device Driver DSM** → **Subsystem Device Driver DSM**. A command prompt opens; run **datapath query version** to determine the version currently installed (Example 4-1) for this Windows 2008 R2 host.

Example 4-1 datapath query version

```
C:\Program Files\IBM\SDDDSM>datapath.exe query version
IBM SDDDSM Version 2.4.3.4-4
Microsoft MPIO Version 6.1.7601.17514
```

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

Example 4-2 datapath query wwpn

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

If you need more detailed information about SDDDSM, see *Multipath Subsystem Device Driver User's Guide,* GC52-1309.

The Windows host has been prepared to connect to the IBM Storwize V7000 and you know the WWPNs of the host. The next step is to configure a host object for the WWPNs by using the IBM Storwize V7000 GUI. This task is explained in 4.3.1, “Creating Fibre Channel hosts” on page 122.

SAN Boot hosts are beyond the intended scope of this book. For more information about that topic, search for `SAN Boot` at the following address:

`http://www.ibm.com/support/knowledgecenter/ST3FR7_7.2.0/com.ibm.storwize.v7000.720.doc/v7000_ichome_720.html?lang=en`

Windows 2003: The examples focus on Windows 2008 R2 and Windows 2012, but the procedure for Windows 2003 is similar. If you use Windows 2003, do not forget to install Microsoft Hotfix 908980. If you do not install it before performing this procedure, preferred pathing is not available. You can download this hotfix at the following address:

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

4.2.2 Windows 2008 R2 and Windows 2012: Preparing for iSCSI attachment

In Windows 2008 R2 and 2012, the Microsoft iSCSI software initiator is preinstalled. Enter iscsi in the search field of the Windows start menu (Figure 4-4) and click **iSCSI Initiator**.

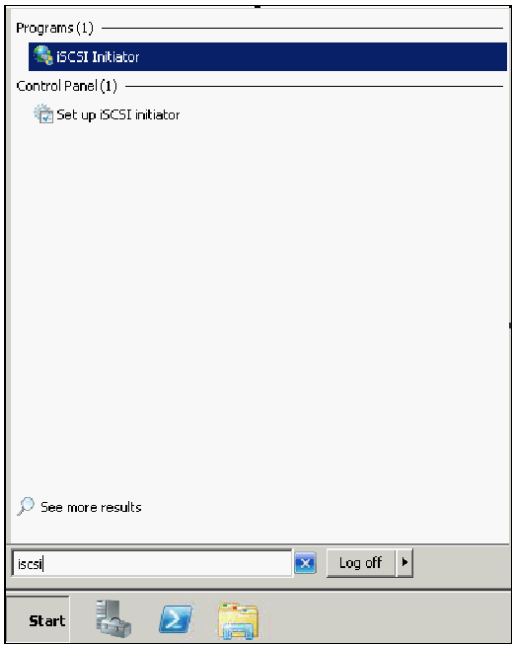

Figure 4-4 Windows iSCSI Initiator

Confirm the automatic startup of the iSCSI Service (Figure 4-5).

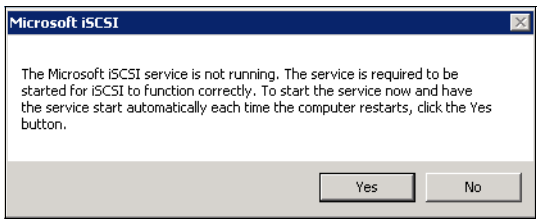

Figure 4-5 Automatic startup of the iSCSI service

The iSCSI Configuration window opens. Select the **Configuration** tab (Figure 4-6). Write down the initiator name of your Windows host, as you need it later.

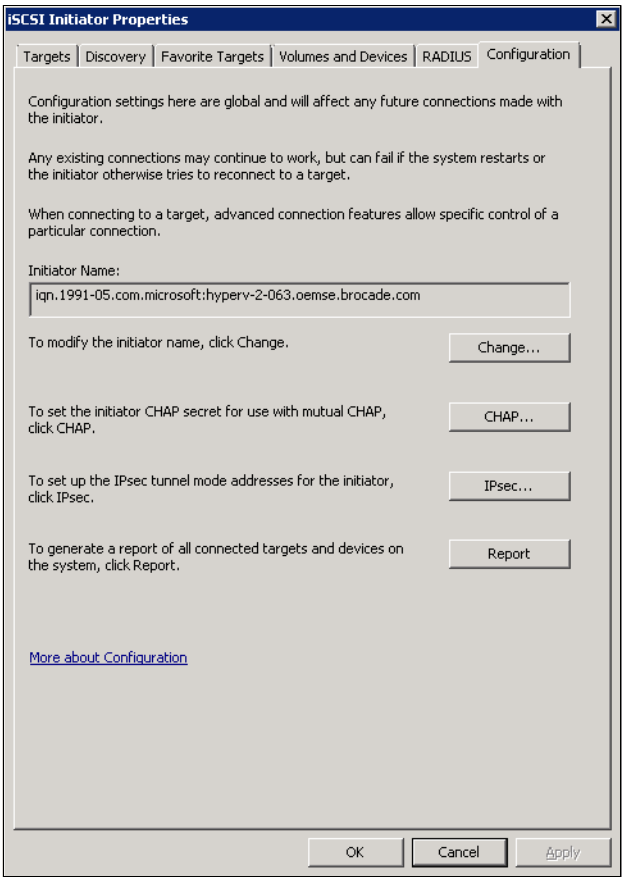

Figure 4-6 iSCSI Initiator Properties window

You can change the initiator name, or enable advanced authentication, but these actions are out of the scope of our basic setup; by default, iSCSI authentication is not enabled. More details are available at the IBM Storwize V7000 Knowledge Center at the following address:

http://publib.boulder.ibm.com/infocenter/storwize/ic/index.jsp?topic=/com.ibm.storwize.v7000.doc/svc_iscsiwindowsauthen_fu67gt.html

Setting the Windows registry keys

Modify the system registry as follows so that your iSCSI operations are more reliable:

1. In the search field of the Windows Start menu, type regedit and click `regedit.exe`.
2. In the registry editor, locate the following key:

   `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<bus ID>\Parameters\LinkDownTime`

   Confirm that the value for the `LinkDownTime` key is 120 (decimal value), and, if necessary, change the value to 120.
3. In the registry editor, locate the following key:

   `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<bus ID>\Parameters\MaxRequestHoldTime`

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

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

   `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<bus ID>\Parameters\ MaxPendingRequests`

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

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

   `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Disk\TimeOutValue`

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

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

These basic steps are to prepare a Windows 2008 R2 or Windows 2012 host for iSCSI attachment. To configure the IBM Storwize V7000 for iSCSI connections, see 4.3.2, “Creating iSCSI hosts” on page 125.

4.2.3 VMware ESX: Preparing for Fibre Channel attachment

Complete the following steps to prepare a VMware ESX host to connect to an IBM Storwize V7000 using Fibre Channel:

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

Downloading and installing the supported firmware

Install the latest firmware levels to your host server. For the HBAs, check the *IBM Storwize V7000 Supported Hardware List, Device Driver, Firmware, and Recommended Software Levels V7.2* document for VMware at the following address:

`http://www.ibm.com/support/docview.wss?uid=ssg1S1004450#_VMWareFC`

Download the latest supported HBA firmware for your configuration and apply it to your system. Some HBAs and especially the new CNA Adapters require an additional driver to be loaded into ESX. Check the VMware Compatibility Guide to see if there are any requirements for your configuration by going to the following address:

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

Configuring Brocade HBAs for VMware ESX

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

Configuring QLogic HBAs for VMware ESX

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

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

Configuring Emulex HBAs for VMware ESX

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

VMware ESX installation

Install your VMware ESX server and load any additional drivers and patches if required. If you are not familiar with the procedure, see the installation guide at this address:

http://www.vmware.com/pdf/vsphere4/r40/vsp_40_esx_vc_installation_guide.pdf

After you have completed your ESX installation, connect to your ESX Server using the vSphere client and navigate to the Configuration tab, click **Storage Adapters**, and scroll down to your FC HBAs (Figure 4-7). Note the WWPNS of the installed adapters for later use.

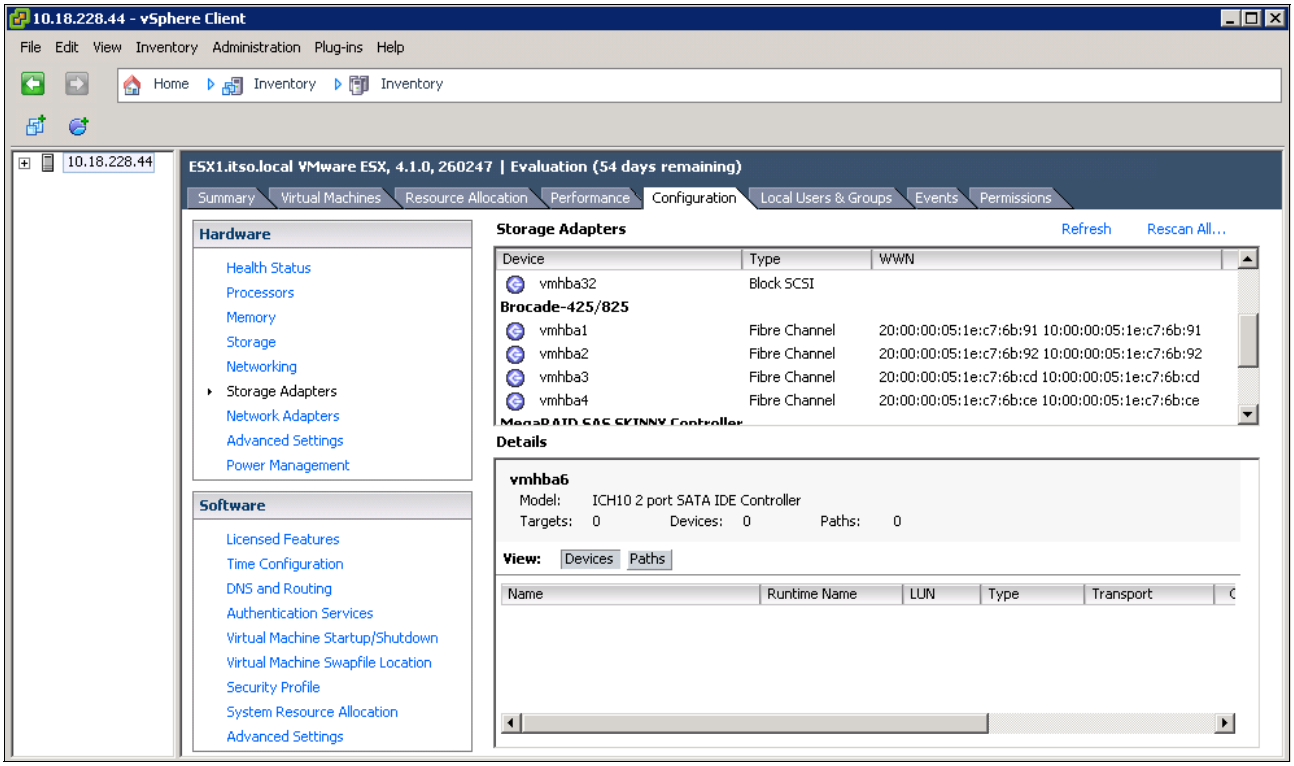

Figure 4-7 Show WWPNs in VMware ESX

VMware ESX multipathing

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

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

The IBM Storwize V7000 is an active/active storage device. Since VMware ESX 4.0 and later, the suggested multipathing policy is *Round Robin*. Round Robin performs static load balancing for I/O. If you do not want to have the I/O balanced over all available paths, the *Fixed* policy is supported also. This policy setting can be selected for every volume. Set this policy after attaching IBM Storwize V7000 LUNs to the ESX host (see 5.3.3, “VMware ESX Fibre Channel attachment” on page 155 for information). If you use an older version of VMware ESX (up to Version 3.5), Fixed is the suggested policy setting.

After all these steps are completed, the ESX host is prepared to connect to the IBM Storwize V7000. Go to 4.3.1, “Creating Fibre Channel hosts” on page 122 to create the ESX FC host in the IBM Storwize V7000 GUI.

4.2.4 VMware ESX: Preparing for iSCSI attachment

This section describes how to enable iSCSI on VMware ESX hosts. We focus on vSphere (ESX 4.0 and later) because the complete iSCSI stack has been rewritten in this level, offers improved performance, and supports useful features, such as jumbo frames and TCP Segmentation Offload. We focus on the basic ESX iSCSI setup; more detailed information is provided in the VMware *iSCSI SAN Configuration Guide*, which is available at the following address:

`http://www.vmware.com/pdf/vsphere4/r40/vsp_40_iscsi_san_cfg.pdf`

Complete the following steps to prepare a VMware ESX host to connect to an IBM Storwize V7000 using iSCSI:

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

The iSCSI initiator is installed by default on your ESX server, and you only have to enable it. To enable it, complete the following steps:

1. Connect to your ESX server using the vSphere Client, navigate to **Configuration**, and select **Networking** (Figure 4-8).

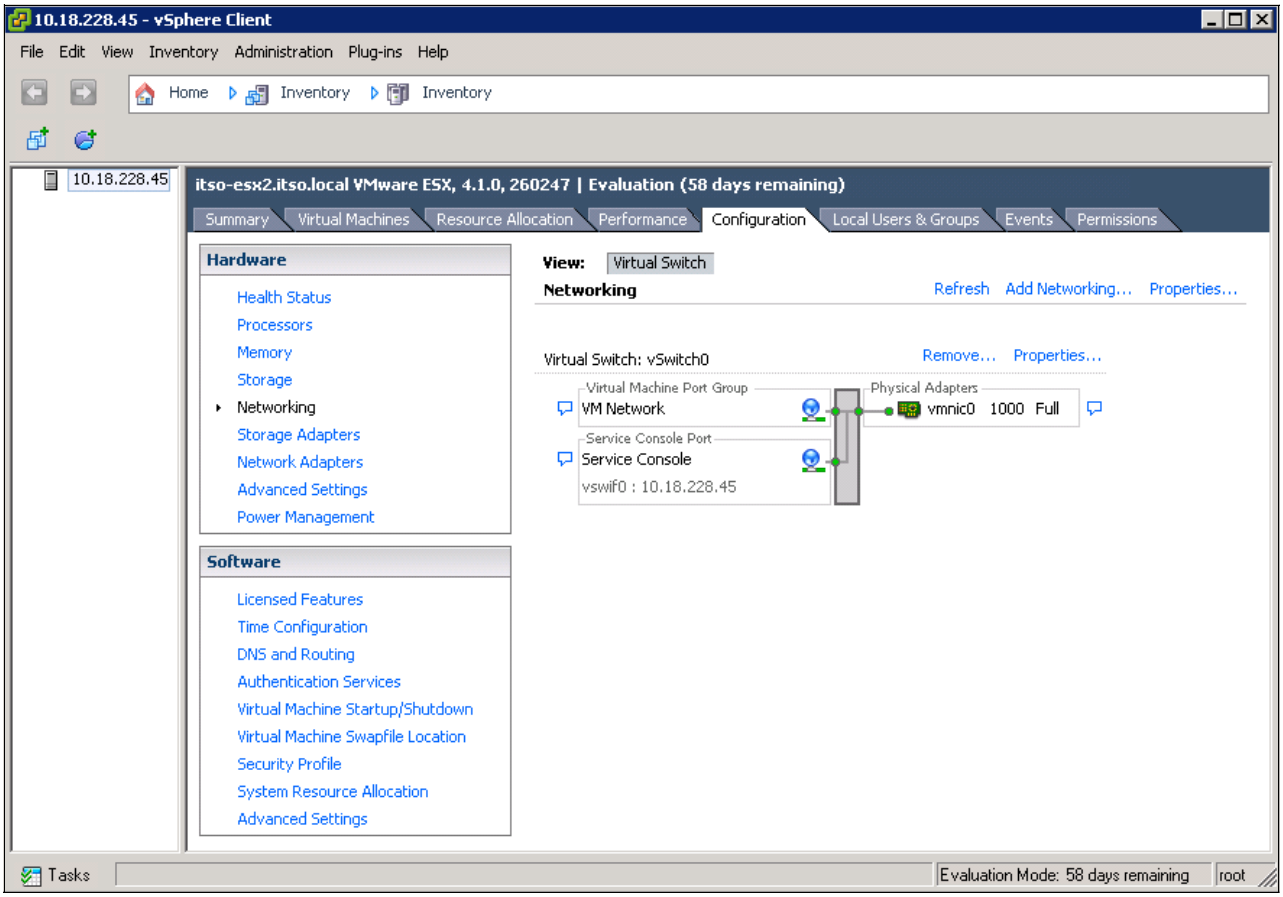

Figure 4-8 Select VMware networking

2. Click **Add Networking** to start the Add Network Wizard (Figure 4-9). Select **VMkernel** and click **Next**.

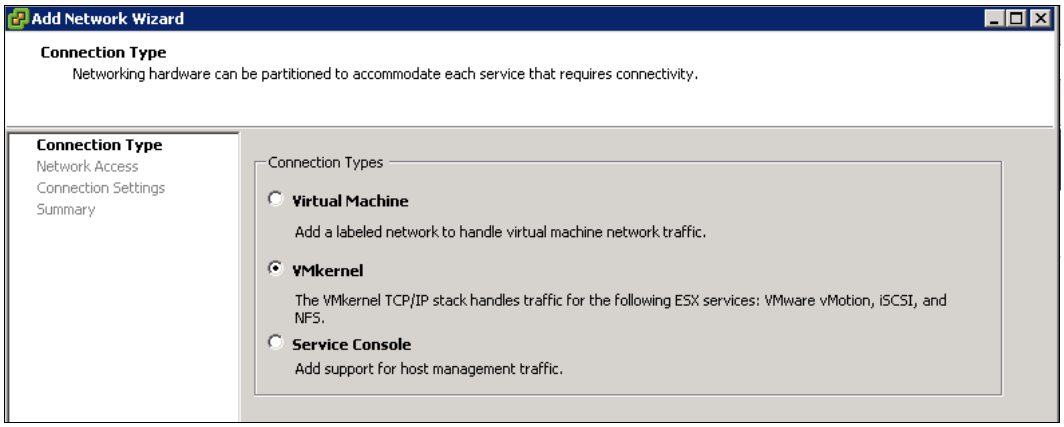

Figure 4-9 VMware: Add Network Wizard

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

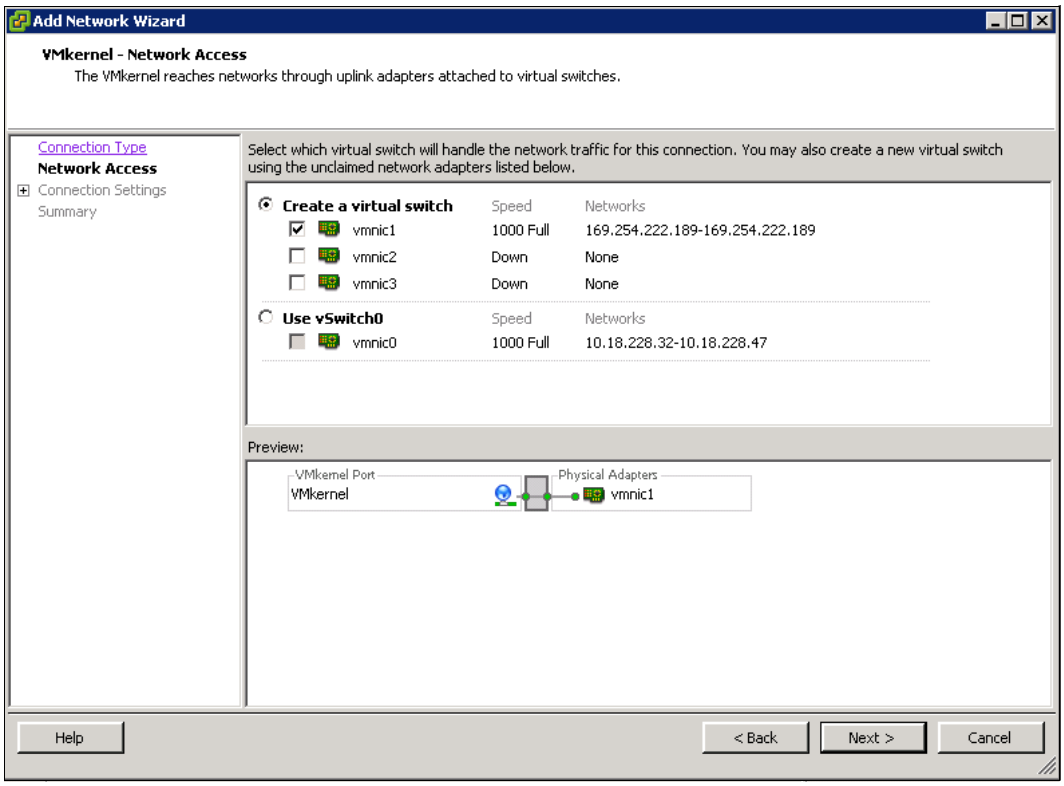

Figure 4-10 VMware: Select an iSCSI interface

4. Enter a meaningful Network Label and click **Next** (Figure 4-11).

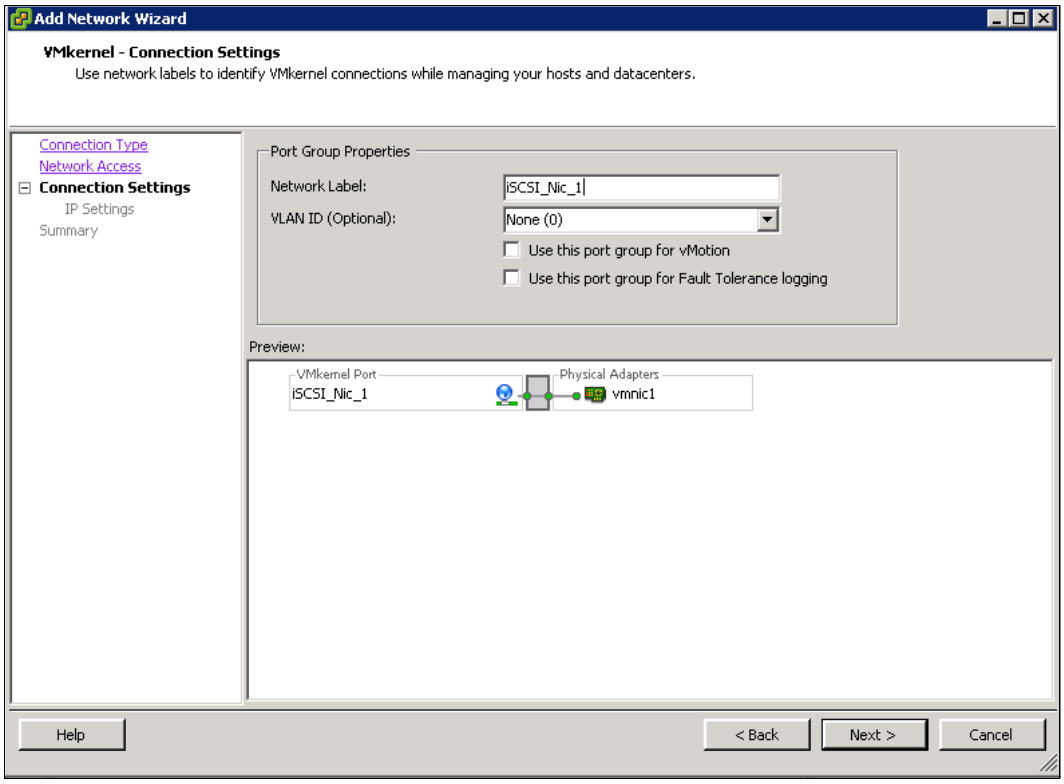

Figure 4-11 VMware: Enter a Network Label

5. Enter an IP address for your iSCSI network. You should use a dedicated network for iSCSI traffic (Figure 4-12).

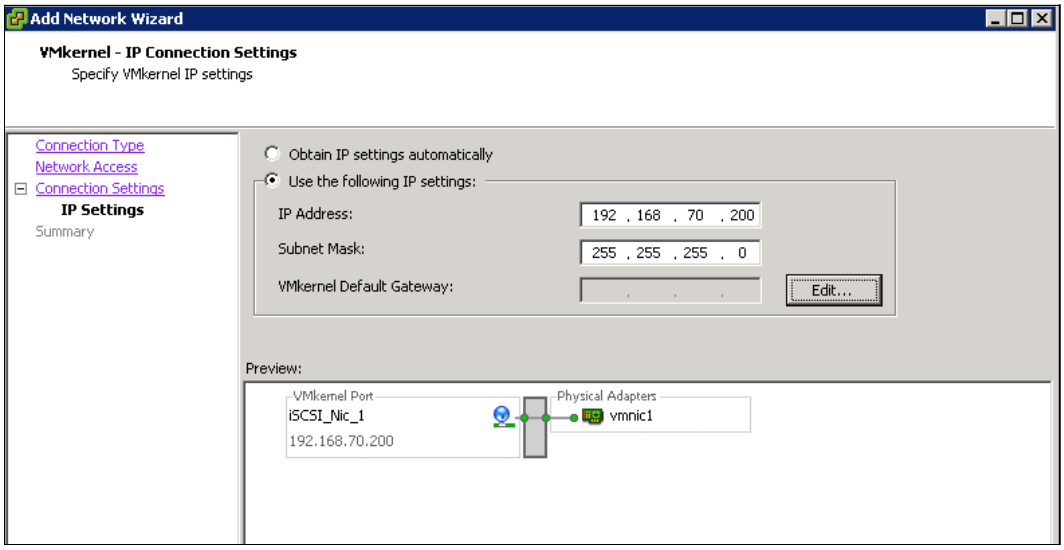

Figure 4-12 VMware: Enter iSCSI Network IP

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

7. Select **Storage Adapters** and scroll to iSCSI Software Adapter (Figure 4-13). Highlight it and click **Properties**.

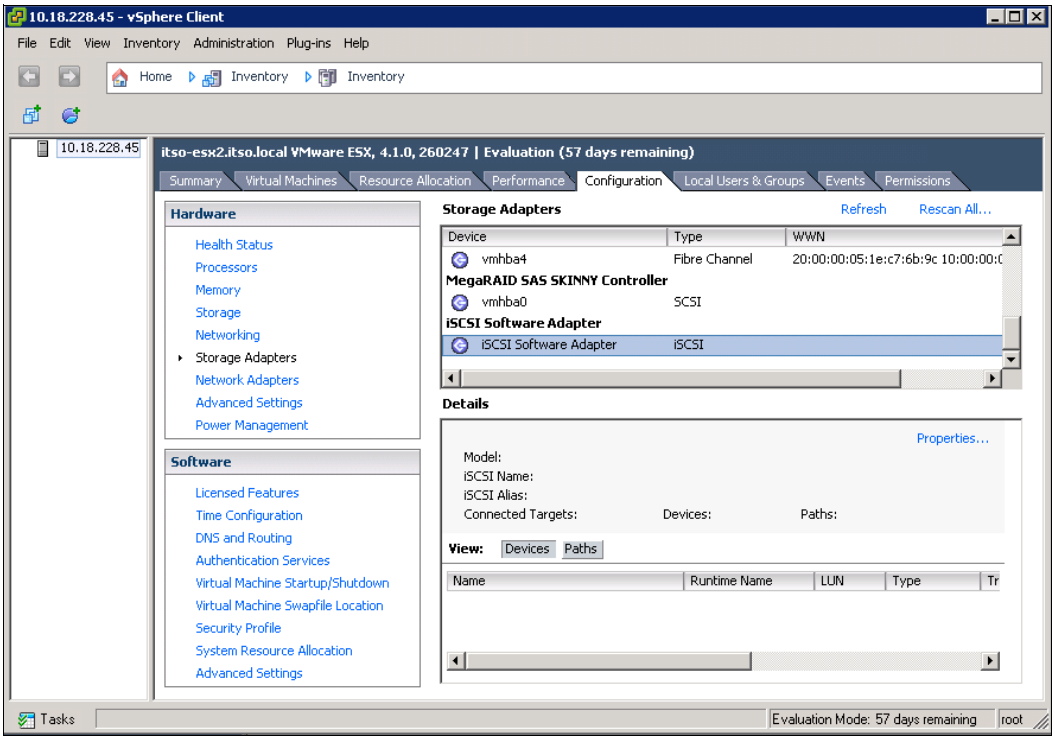

Figure 4-13 VMware: iSCSI Software Adapter

8. The iSCSI Software Adapter Properties window opens. Figure 4-14 shows that the initiator is disabled by default; to change this setting, click **Configure**.

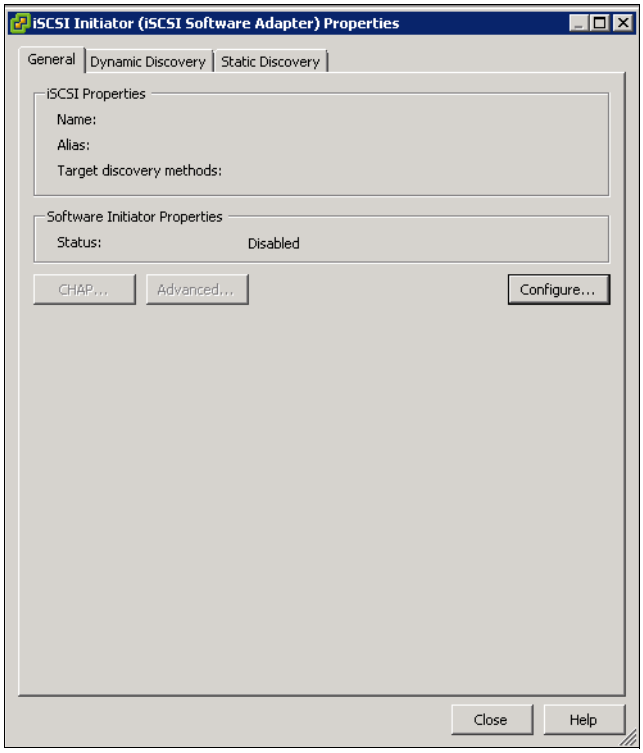

Figure 4-14 VMware: iSCSI Software Adapter Properties

9. Select the **Enabled** check box and click **OK** (Figure 4-15).

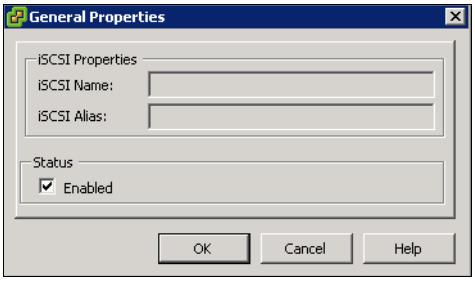

Figure 4-15 VMware: Select the Enabled check box

10. The VMware ESX iSCSI initiator is now successfully enabled (Figure 4-16). Note the initiator name for later use.

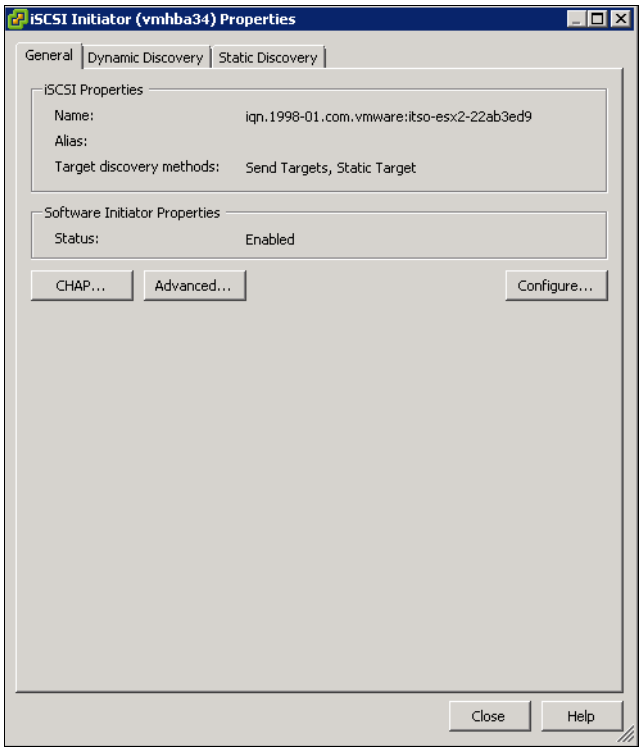

Figure 4-16 VMware iSCSI Initiator enabled

Your VMware ESX host is now prepared to connect to the IBM Storwize V7000. Go to 4.3.2, “Creating iSCSI hosts” on page 125 to create the ESX iSCSI host using the IBM Storwize V7000 GUI.

4.3 Creating hosts using the GUI

This section describes how to create Fibre Channel and iSCSI hosts using the IBM Storwize V7000 GUI. We assume that the hosts are prepared for attachment, as described in 4.2, “Preparing the host operating system” on page 107, and that you know the host WWPNs and their iSCSI initiator names.

To create a host, complete the following steps:

1. Open the host configuration window by clicking **Hosts** (Figure 4-17).

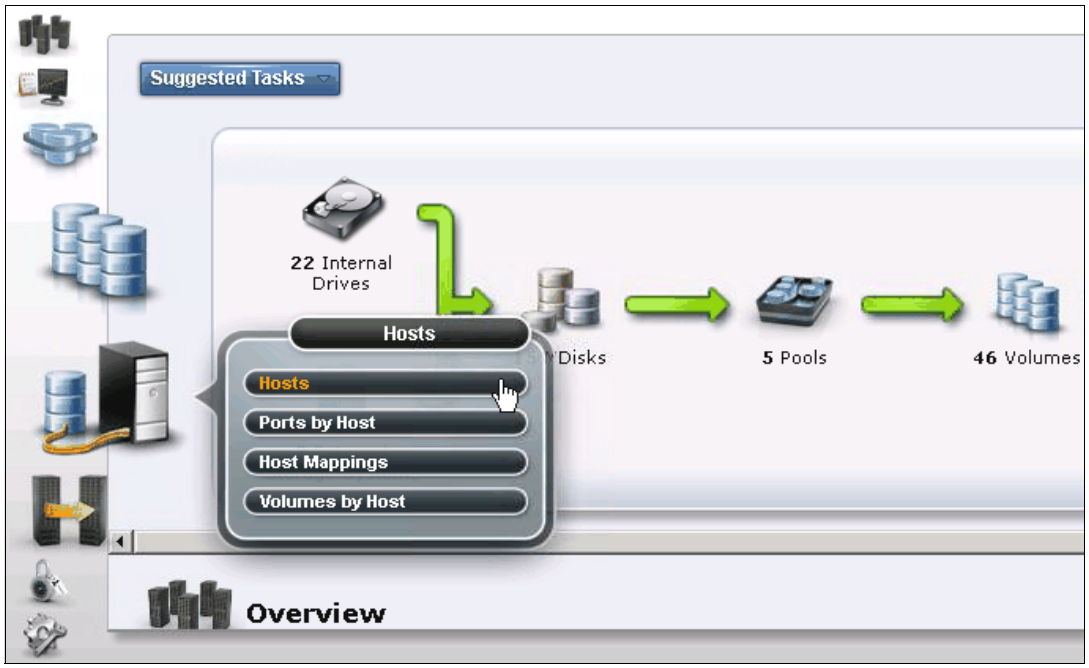

Figure 4-17 Open the host window

2. To create a host, click **New Host** to start the wizard (Figure 4-18).

Figure 4-18 Create a host

If you want to create a Fibre Channel host, continue with 4.3.1, “Creating Fibre Channel hosts” on page 122; to create iSCSI hosts, go to 4.3.2, “Creating iSCSI hosts” on page 125.

4.3.1 Creating Fibre Channel hosts

To create Fibre Channel hosts, complete the following steps:

1. Click **Fibre Channel Host** (Figure 4-18 on page 121). The Fibre Channel configuration wizard opens (Figure 4-19).

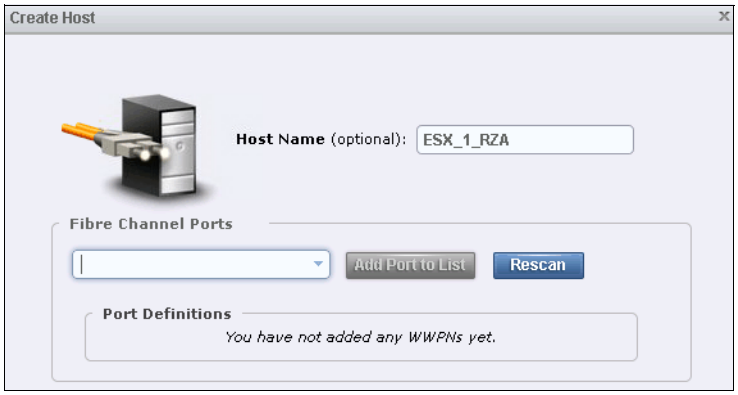

Figure 4-19 Create a Fibre Channel host

2. Enter a host name and click the **Fibre Channel Ports** drop-down menu to get a list of all known WWPNs (Figure 4-20).

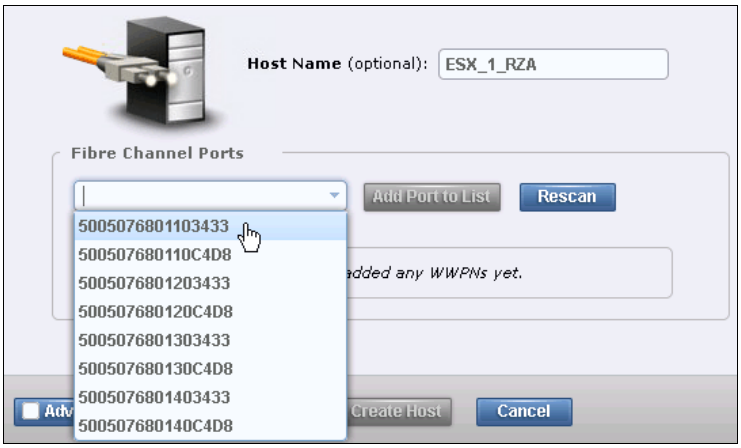

Figure 4-20 Available WWPNs

The IBM Storwize V7000 has the host port WWPNs available if you prepared the hosts, as described in 4.2, “Preparing the host operating system” on page 107. If they do not appear in the list, scan for new disks in your operating system and click **Rescan** in the configuration wizard. If they still do not appear, check your SAN zoning and repeat the scanning.

3. Select the WWPN for your host and click **Add Port to List** (Figure 4-21).

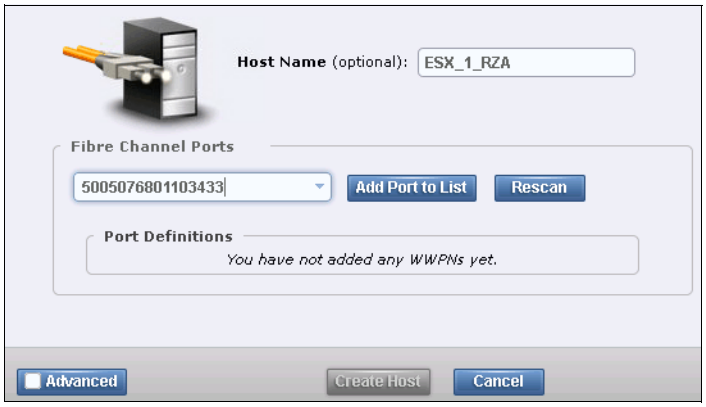

Figure 4-21 Add a port to a list

4. Add all ports that belong to the host (Figure 4-22).

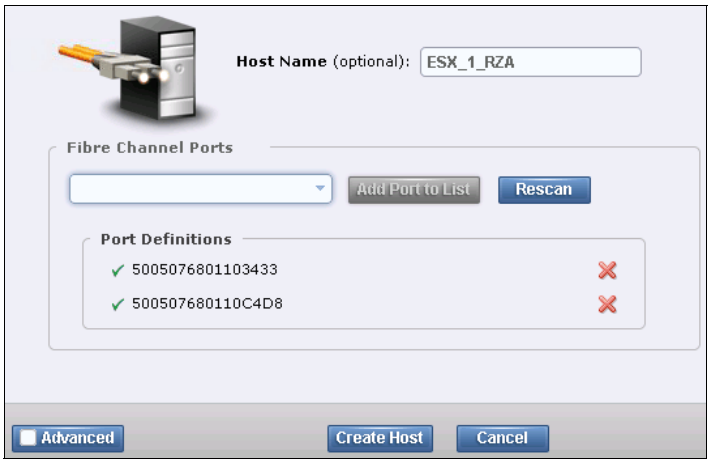

Figure 4-22 Add all WWPNs

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

5. If you are creating an HP/UX or TPGS host, select the **Advanced** check box and more options appear (Figure 4-23). Select your host type.

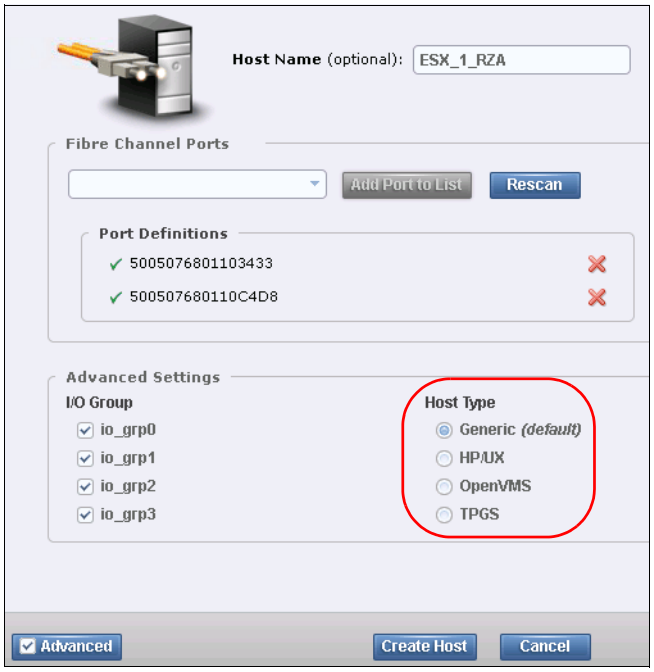

Figure 4-23 Create Host: Advanced Settings

6. Click **Create Host** and the wizard creates the host (Figure 4-24).

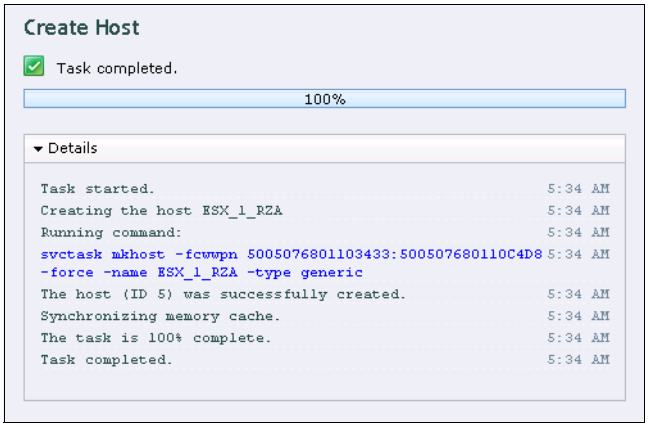

Figure 4-24 Create Host completes

7. Click **Close** to return to the host window. Repeat these steps for all of your Fibre Channel hosts. Figure 4-25 shows the All Hosts window after creating a second host.

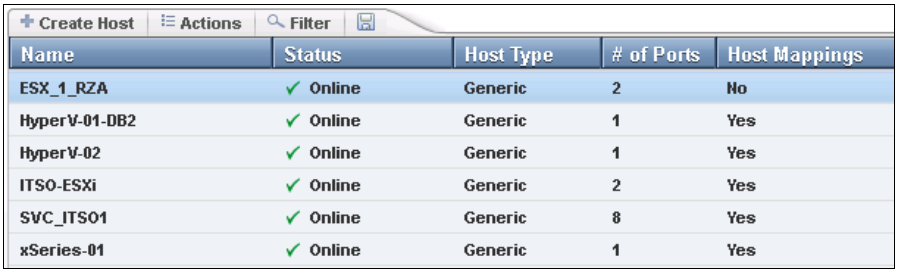

Figure 4-25 All Hosts: After creating a second host

After you complete the creation of Fibre Channel hosts, go to Chapter 5, “Basic volume configuration” on page 131 to create volumes and map them to the created hosts.

4.3.2 Creating iSCSI hosts

To create iSCSI hosts, complete the following steps:

1. Click **iSCSI Host** (Figure 4-18 on page 121) and the iSCSI configuration wizard opens (Figure 4-26).

Figure 4-26 Create an iSCSI host

2. Enter a host name, type the iSCSI initiator name into the iSCSI Ports box, and click **Add Ports to List**. If you want to add several initiator names to one host, repeat this step.

3. If you are connecting an HP/UX or TPGS host, select the **Advanced** check box (Figure 4-27) and select the correct host type.

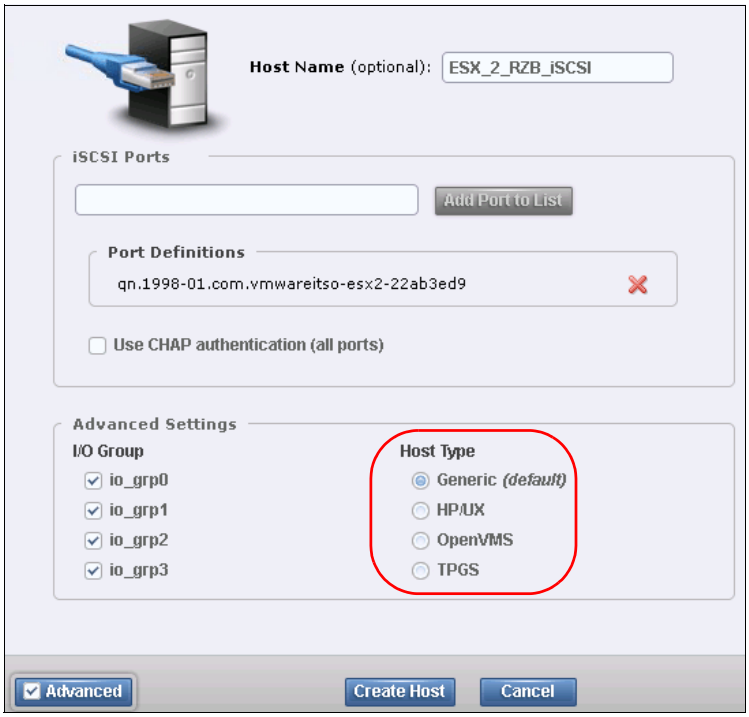

Figure 4-27 Create an iSCSI host: Advanced Settings

4. Click **Create Host** and the wizard completes (Figure 4-28). Click **Close**.

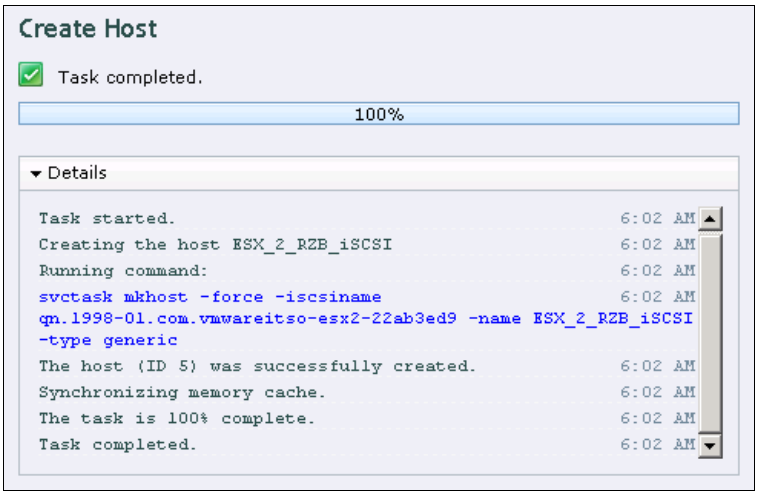

Figure 4-28 Create an iSCSI host: Complete

5. Repeat these steps for every iSCSI host you want to create. Figure 4-29 shows the All Hosts window after creating two Fibre Channel and two iSCSI hosts.

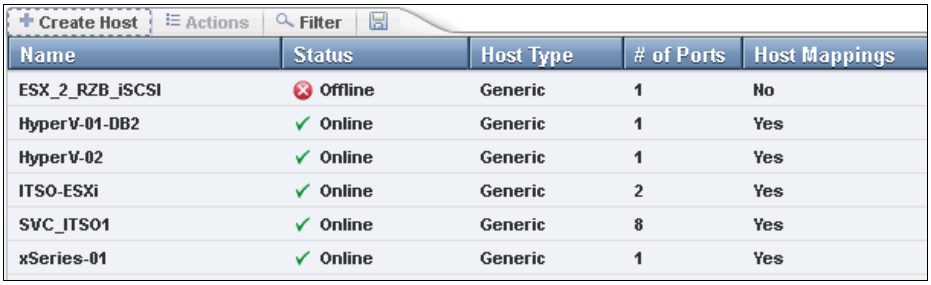

Name	Status	Host Type	# of Ports	Host Mappings
ESX_2_RZB_iSCSI	Offline	Generic	1	No
HyperV-01-DB2	Online	Generic	1	Yes
HyperV-02	Online	Generic	1	Yes
ITSO-ESXi	Online	Generic	2	Yes
SVC_ITSO1	Online	Generic	8	Yes
xSeries-01	Online	Generic	1	Yes

Figure 4-29 All Hosts

The iSCSI hosts are now configured on the IBM Storwize V7000. To provide connectivity, the iSCSI Ethernet ports also must be configured. Complete the following steps to enable iSCSI connectivity:

1. Switch to the Configuration window and select **Network** (Figure 4-30).

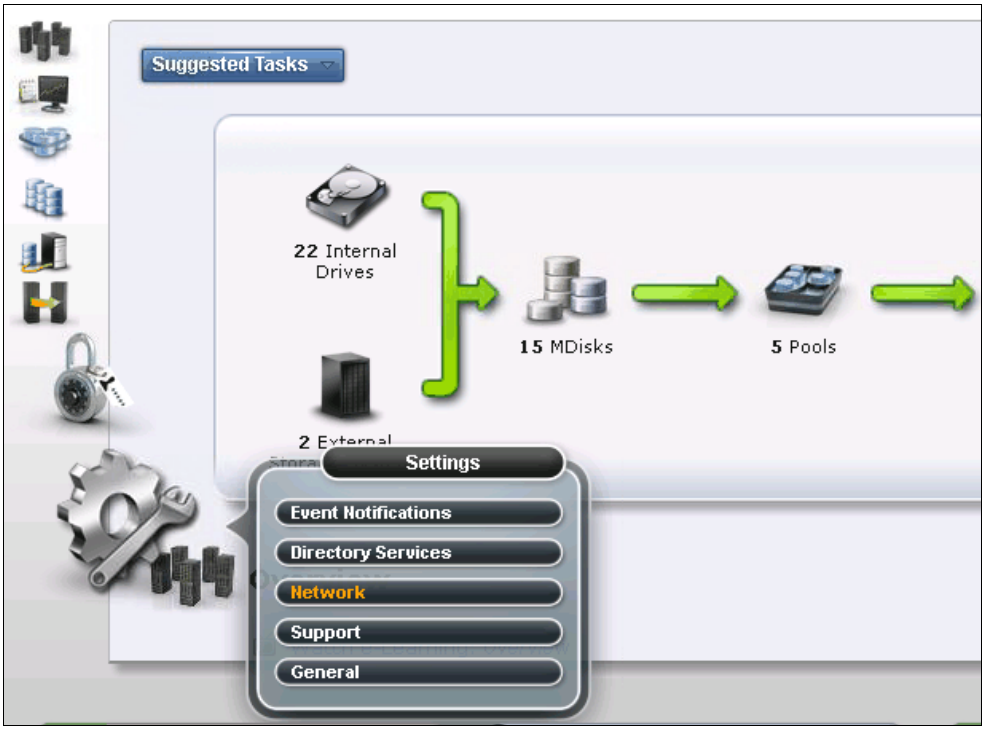

Figure 4-30 Configuration: Network

2. Select **iSCSI** and the iSCSI Configuration window opens (Figure 4-31).

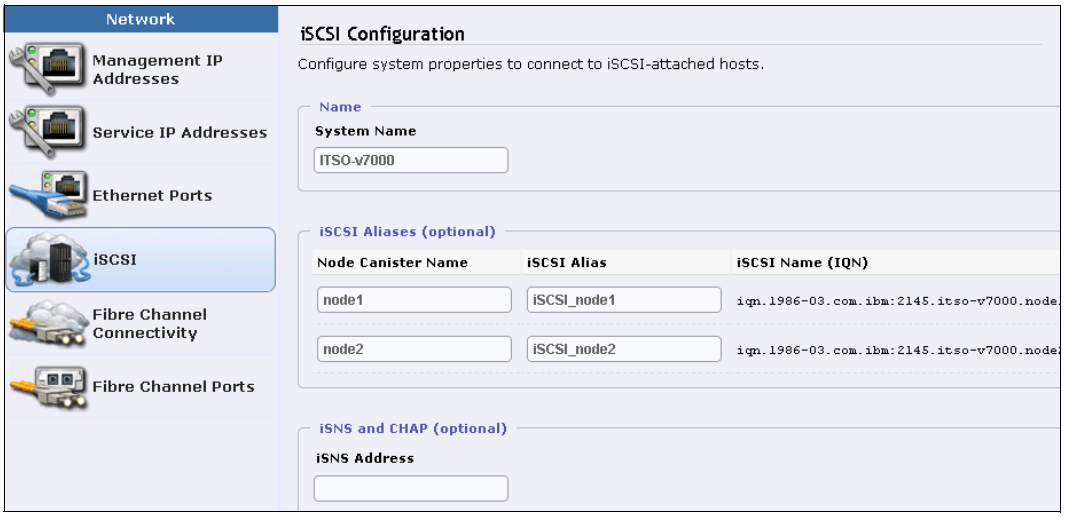

Figure 4-31 iSCSI Configuration window

In the configuration, you have an overview of all the iSCSI settings for the IBM Storwize V7000. You can configure iSCSI Alias, iSNS Addresses, and Chap Authentication Configuration on this window, and the iSCSI IP address, which we also edit in the basic setup.

3. Click **Ethernet Ports** to enter the iSCSI IP address (Figure 4-32). Repeat this step for each port that you want to use for iSCSI traffic.

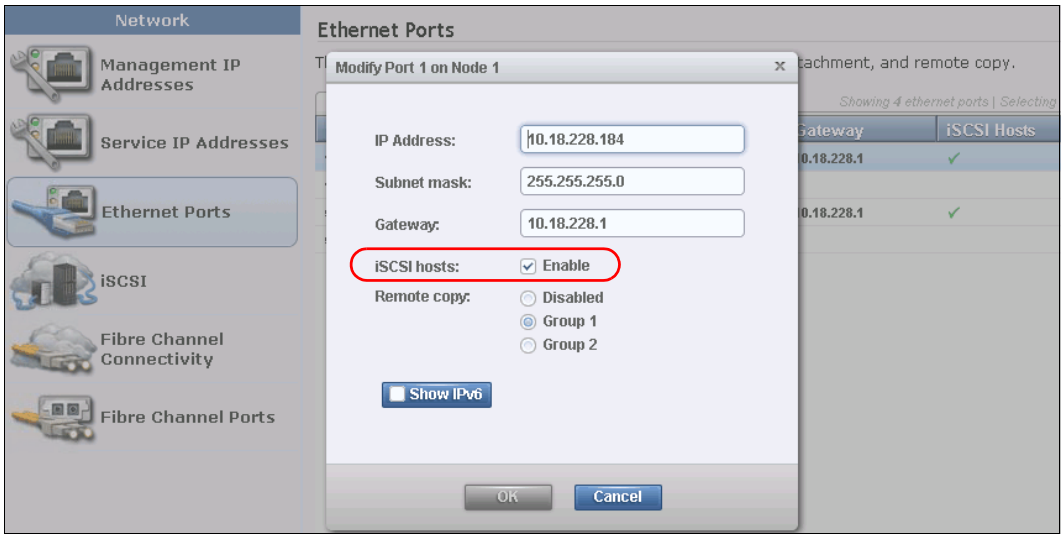

Figure 4-32 Enter an iSCSI IP address

4. After you enter the IP address for each port, click **OK** to enable the configuration.

5. After the changes are successfully applied, click **Close** (Figure 4-33).

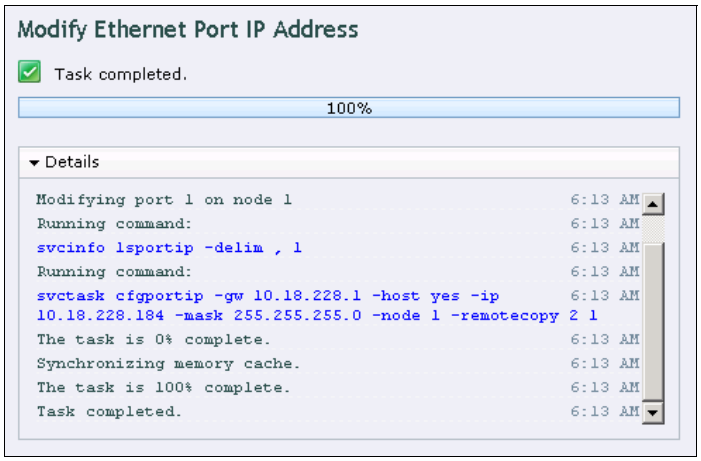

Figure 4-33 iSCSI IP successfully modified

The IBM Storwize V7000 is now configured and ready for iSCSI use. Note the initiator names of your storage canisters (Figure 4-31 on page 128), because you need them later. To create volumes and map them to a host, go to Chapter 5, “Basic volume configuration” on page 131.

5

Basic volume configuration

This chapter describes how to use the IBM Storwize V7000 to create a volume and map a volume to a host. A volume is a logical disk on the IBM Storwize V7000 that is provisioned out of a storage pool and is recognized by a host with an identifier UID field and a parameter list.

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

The second part of this chapter covers how to discover those volumes (5.1, “Provisioning storage from IBM Storwize V7000 to hosts” on page 132). After you finish this chapter, your basic configuration is done and you are able to store data on the IBM Storwize V7000.

Advanced host and volume administration, such as volume migration, creating volume copies, and so on, is described in Chapter 8, “Advanced host and volume administration” on page 221.

5.1 Provisioning storage from IBM Storwize V7000 to hosts

This section describes the setup process and shows how to create volumes and make them accessible from the host. The described steps for completing the basic setup in your environment are as follows:

1. Define new volumes in IBM Storwize V7000.
2. Map volumes to the host.
3. Discover mapped volumes from the host and configure multipathing.

From the dynamic menu (function icons), open the Volumes menu and click the **Volumes** option of the IBM Storwize V7000 GUI (Figure 5-1) to start the process of creating new volumes.

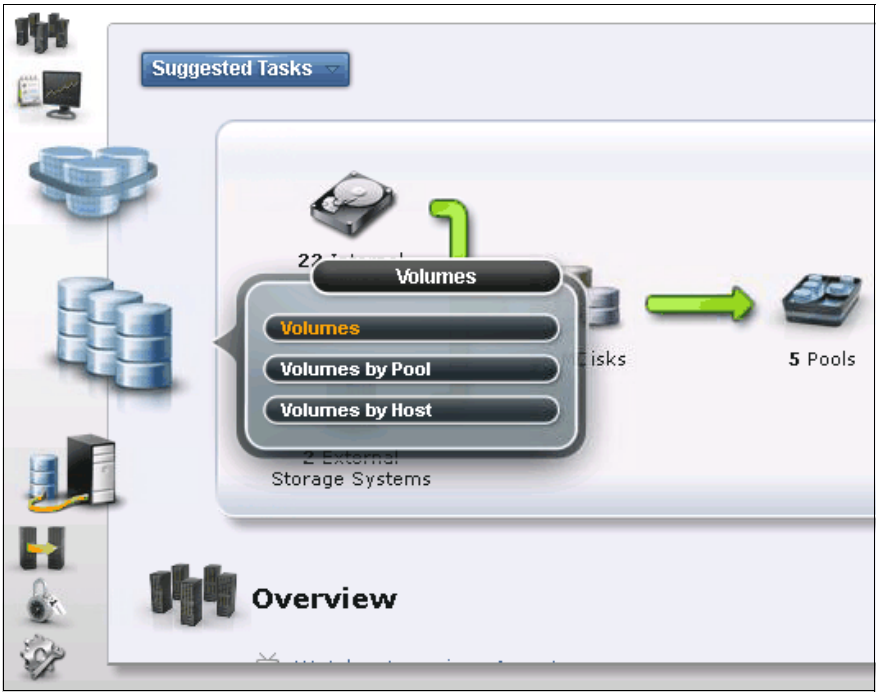

Figure 5-1 Volumes menu

A list of existing volumes that are defined in specific storage pools is displayed. To define a new volume, click **Create Volume** on the tab header (Figure 5-2).

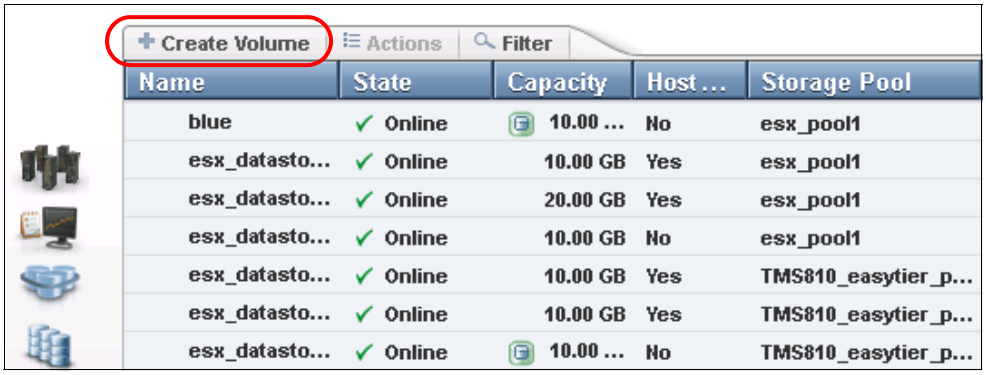

Figure 5-2 New Volume window

The wizard opens (Figure 5-3).

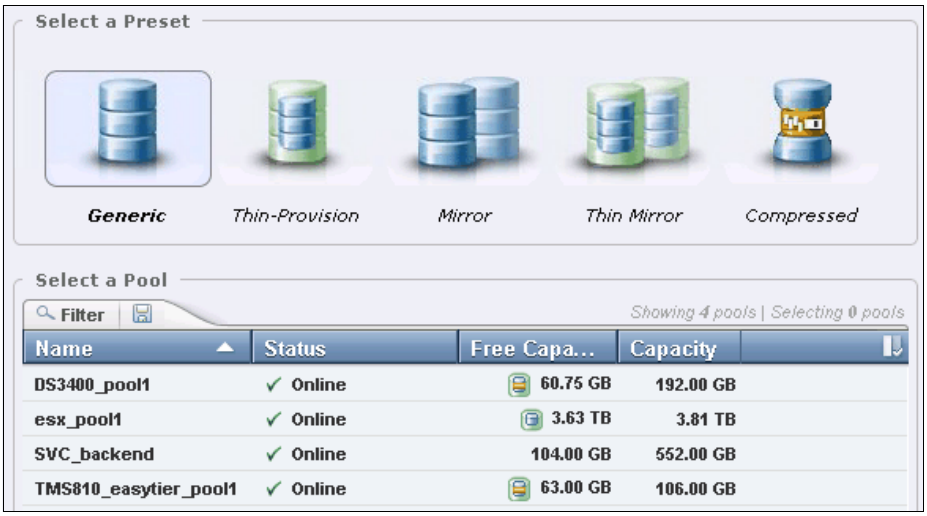

Figure 5-3 Wizard with predefined volumes

By default, all volumes that you create are striped across all available MDisks in one storage pool. The GUI for the IBM Storwize V7000 provides the following preset selections:

- Generic: A striped volume that is fully provisioned, as described in 5.1.1, “Creating generic volumes” on page 133.
- Thin-provisioned: A striped volume that is space-efficient. Click **Advanced** to see more to help you determine how much space is fully allocated initially and how large the volume is able to grow, as described in 5.1.2, “Creating thin-provisioned volumes” on page 135.
- Mirror: The striped volume consists of two striped copies and is synchronized to protect against loss of data if the underlying storage pool of one copy is lost, as described in 5.1.3, “Creating mirrored volumes” on page 137.
- Thin-mirror: Two synchronized copies. Both are thin-provisioned, as described in 5.1.4, “Creating a thin-mirrored volume” on page 140.
- Compressed: The IBM Real-time Compression engine, available in IBM Storwize V7000 v6.4 and later, enables data compression for storage efficiency. More details are in *Real-time Compression in SAN Volume Controller and Storwize V7000*, REDP-4859.

5.1.1 Creating generic volumes

The most commonly used type of volume is the *generic* volume. This type of volume is fully provisioned, with the entire size dedicated to the defined volume. The host and the IBM Storwize V7000 see the fully allocated space without a mirror.

Select a generic volume by clicking the leftmost icon (Figure 5-3) and choose the pool in which the volume should be created. In our example, we selected DS3400_pool1.

Enter the size that you want of the volume and its name.

The result is shown in Figure 5-4. We suggest using an appropriate naming convention of volumes to help you easily identify the associated host or group of hosts. At a minimum, it should contain the name of the pool or at least some tag that identifies the underlying storage subsystem (in our case, the IBM System Storage DS3400 subsystem).

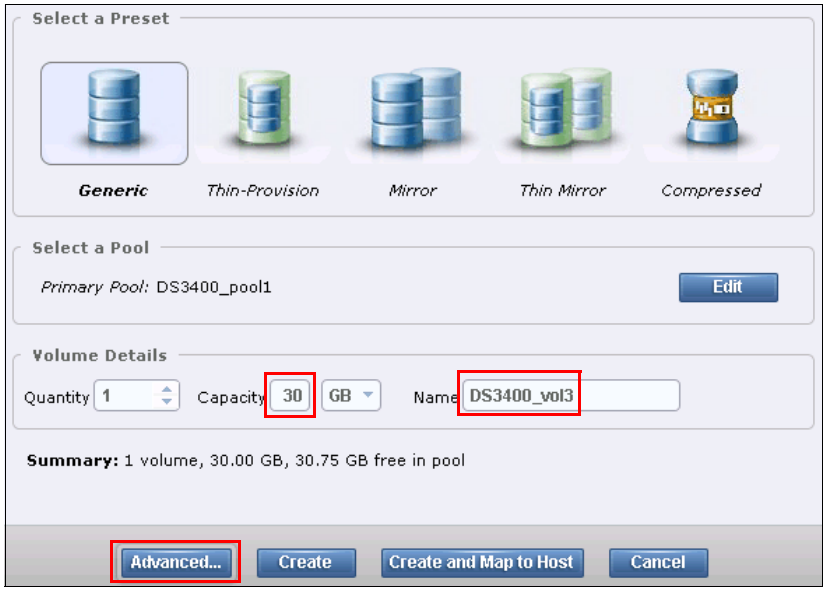

Figure 5-4 Create a generic volume

You can preformat the volume before its use by clicking **Advanced** and selecting **Format Before Use** (Figure 5-5) on the Characteristics tab.

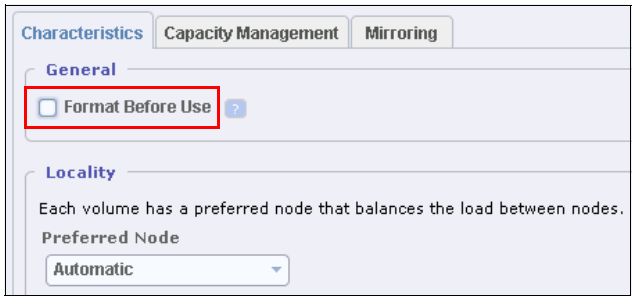

Figure 5-5 Pre-formatting the volume

After you enter a volume name and size, click **Create and Map to Host**.

The new generic volume is created (Figure 5-6). Click **Continue**. A new wizard opens so you can assign a volume to the host as described in 5.2.1, “Mapping newly created volumes to the host using the wizard” on page 142.

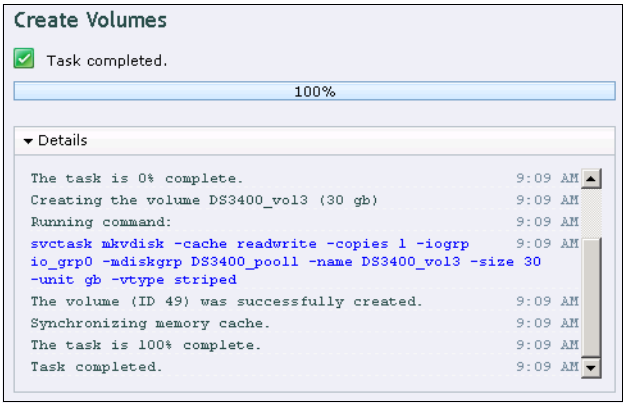

Figure 5-6 Volume created

If you do not want to map the volumes now, click **Create** to complete the task. Volumes can also be mapped later, as described in 5.2, “Mapping a volume to the host” on page 142.

5.1.2 Creating thin-provisioned volumes

Volumes can be configured to be thin-provisioned. With respect to application reads and writes, a thin-provisioned volumes behave as though they were fully allocated. When creating a thin-provisioned volume, you may specify two capacities:

- The real physical capacity allocated to the volume from the storage pool. The real capacity determines the quantity of extents that are initially allocated to the volume.
- Its virtual capacity available to the host. The virtual capacity is the capacity of the volume that is reported to all other components (for example, FlashCopy, cache, and remote copy) and to the host servers.

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

1. From the Volumes selection in the dynamic menu, choose **Create Volume** and select **Thin Provision** (Figure 5-7). Enter the required capacity and volume name.

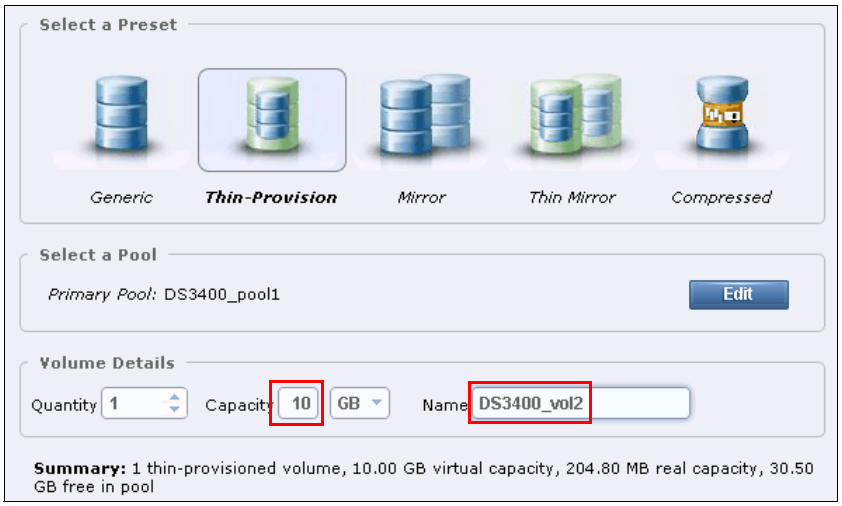

Figure 5-7 Create a thin-provisioned volume

2. Click the **Advanced** button to manage the real and virtual capacity of the volume (Figure 5-8).

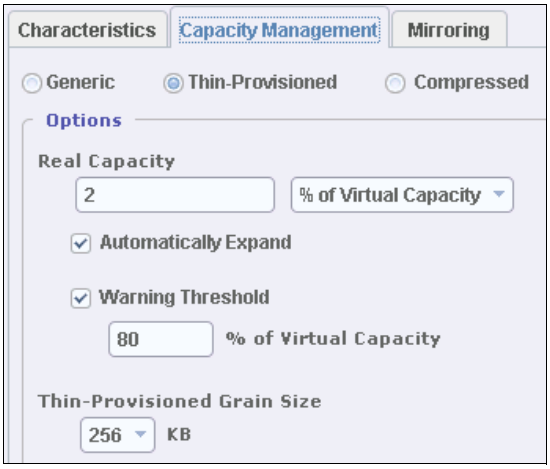

Figure 5-8 Volume capacity management

For the Thin- Provisioned selection on the Capacity Management tab, several advanced options are available:

- Real: Specify the size of the real capacity space used during creation. The default is 2% of virtual capacity.
- Automatically Expand: This option enables the automatic expansion of real capacity, if new capacity must be allocated.
- Warning Threshold: Enter a threshold for receiving capacity alerts.
- Thin-Provisioned Grain Size: Specify the grain size for real capacity.

3. Apply all required changes and click **Create** to define the volume (Figure 5-9).

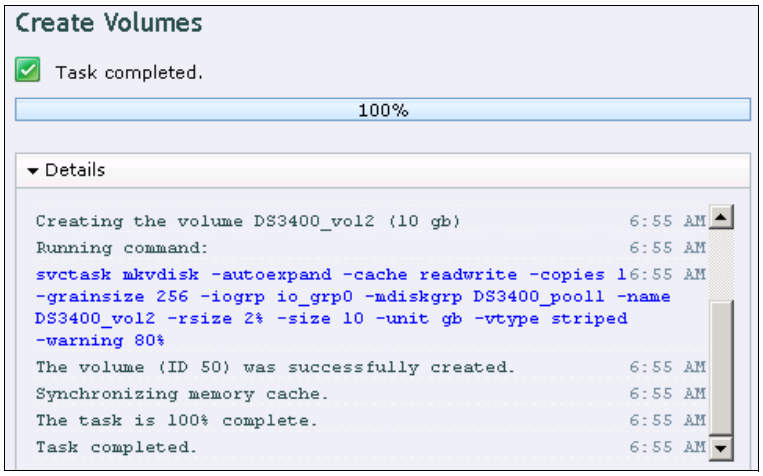

Figure 5-9 Create a thin-provisioned volume

4. Again, you can directly start the wizard for mapping this volume to the host by clicking **Create and Map to Host**.

5.1.3 Creating mirrored volumes

IBM Storwize V7000 offers the capability to mirror volumes, which means a single volume is presented to the host, but two copies exist in the storage back end, usually in different storage pools (all reads are handled by the primary copy). This feature is similar to host based software mirroring such as LVM, but it provides a single point of management for all operating systems, and provides storage high availability to operating systems that do not support software mirroring. With this setup, you can protect against array failures (for example, multiple disk failures), and it offers you more advanced features, as described in Chapter 8, “Advanced host and volume administration” on page 221. It also helps you to greatly improve availability, but it is not a disaster recovery solution, because both copies are accessed by the same node pair and addressable by only a single cluster.

To create a mirrored volume, complete the following steps:

1. Click **Mirror** (Figure 1) and choose the primary and secondary pool by clicking their names. Although the mirrored volume can be created in the same pool, this is not typical. We suggest that you keep mirrored volumes on a separate set of physical disks (separate MDisks).

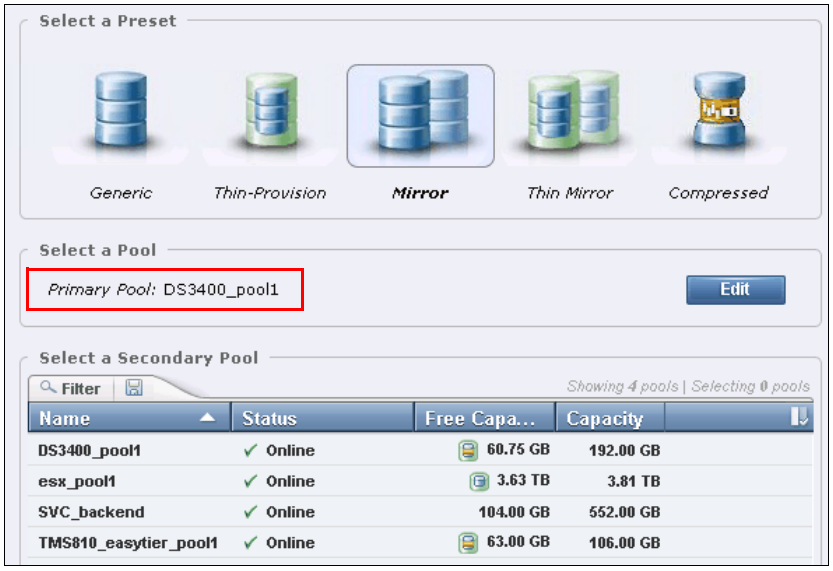

Figure 5-10 Select primary pool

2. Select the secondary pool by clicking it, and then enter a volume name and the required size (Figure 5-11).

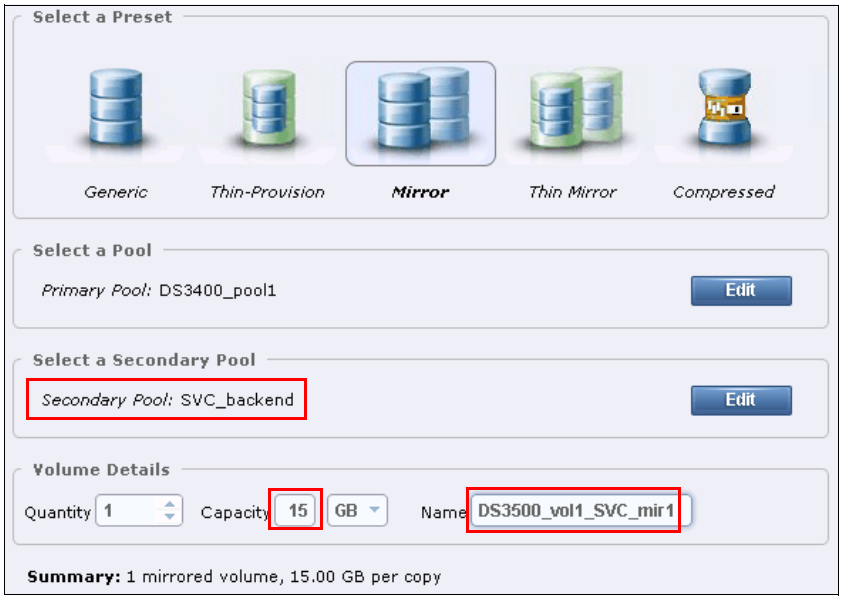

Figure 5-11 Select a secondary pool, volume name and size

The summary shows capacity information about the pool.

3. If you want to select advanced settings, click **Advanced** and click the **Mirroring** tab (Figure 5-12).

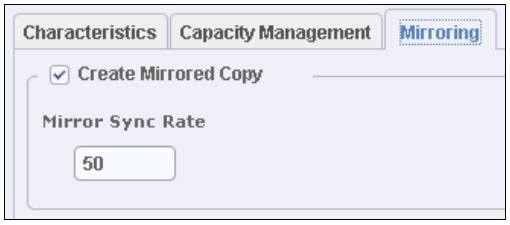

Figure 5-12 Advanced mirroring features

In the advanced mirroring settings, you can specify a synchronization rate. Enter a Mirror Sync Rate in the range of 1 - 100%. This option sets the priority of copy synchronization progress, allowing you to prefer more important volumes for synchronization that is faster than other mirrored volumes. By default, the rate is set to 50% for all volumes. Click **OK** to return to the previous panel.

4. Click **Create and Map to Host**; the mirrored volume is created (Figure 5-13).

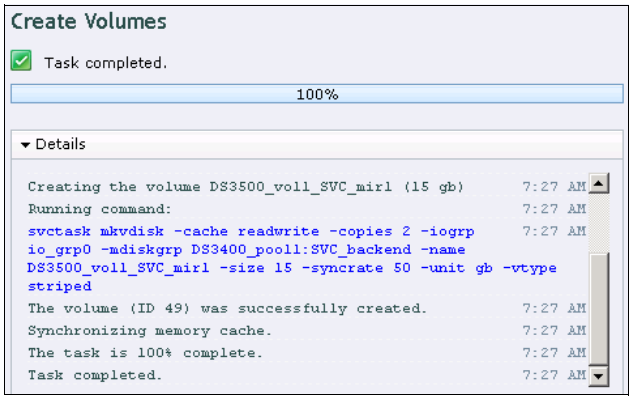

Figure 5-13 Create a mirrored volume

5. Click **Continue.** At this point you can see the progress of a volume synchronization as shown in Figure 5-14. Go to 5.2.1, “Mapping newly created volumes to the host using the wizard” on page 142.

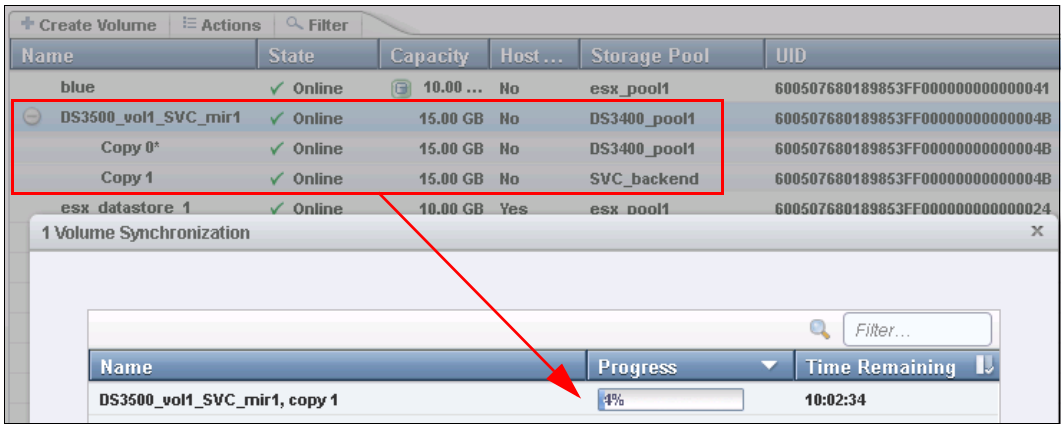

Figure 5-14 Mirror synchronization process

5.1.4 Creating a thin-mirrored volume

With a thin-mirror volume, you can allocate the required space on demand (described in 5.1.2, “Creating thin-provisioned volumes” on page 135) and have several copies of a volume (described in 5.1.3, “Creating mirrored volumes” on page 137) available.

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

1. Select **Thin Mirror** (Figure 5-15) and the primary pool.

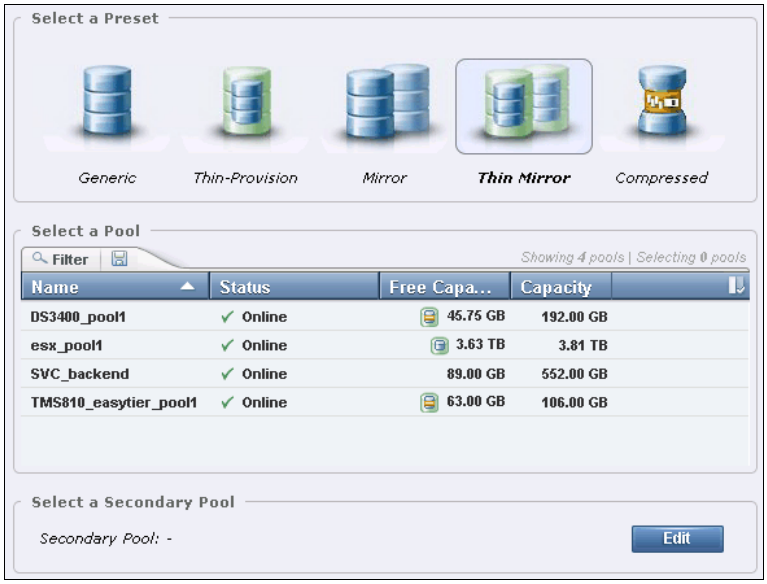

Figure 5-15 Thin Mirror

2. Select the secondary pool and make any changes that you want to the volume attributes (Figure 5-16).

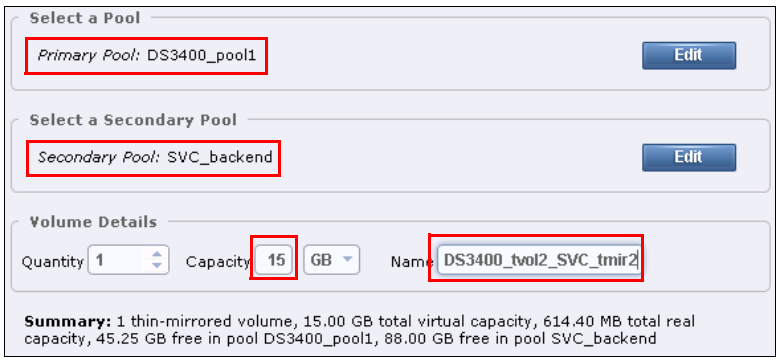

Figure 5-16 Select pools, capacity and name

3. The summary shows you the capacity information and the allocated space. You can click **Advanced** and customize the thin-provision settings or the mirror synchronization rate. After you create and map to host, the confirmation window opens (Figure 5-17).

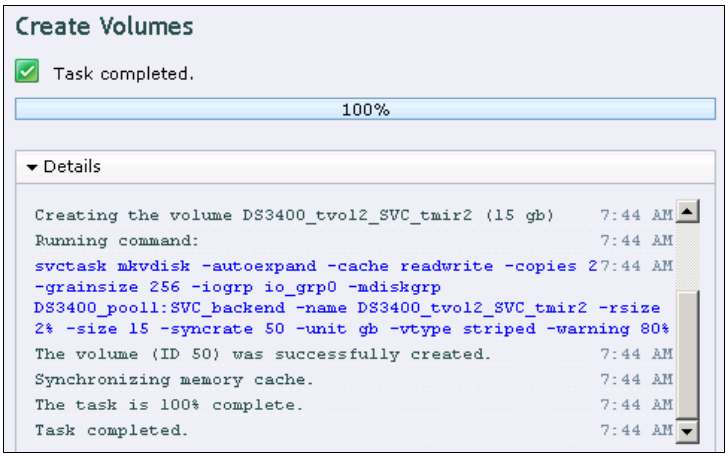

Figure 5-17 Confirmation window

4. The initial synchronization of thin-mirrored volumes is fast when a small amount of real and virtual capacity is used. The amount of time that is required for synchronization might grow rapidly over time (Figure 5-18).

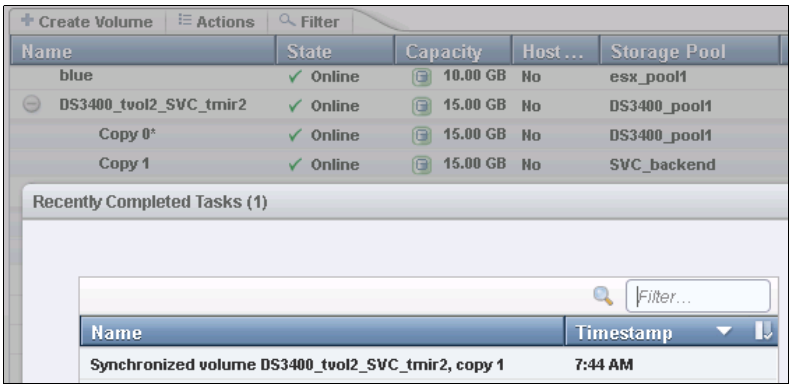

Figure 5-18 Synchronization progress

5.1.5 Creating compressed volumes

IBM Real-time Compression is an advanced software function for storage efficiency, enabled by an IBM SAN Volume Controller family of products, including the IBM Storwize V7000 storage subsystem.

The configuration is similar to generic volumes and is transparent to users. From the Volumes selection in the dynamic menu, click **Create Volume** and click **Compressed** (Figure 5-19 on page 142). Choose a storage pool and enter the required capacity and volume name.

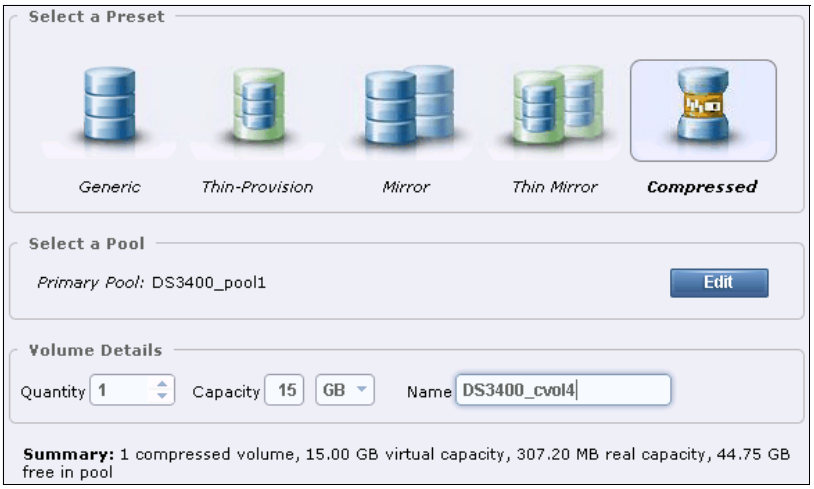

Figure 5-19 Definition of compressed volume

The summary at the bottom of the wizard shows information about allocated (virtual) capacity and the real capacity that data consumes on this volume. In our example, we defined a 15 GB volume, but the real capacity is only 307 MB (because there is no data from the host).

5.2 Mapping a volume to the host

You can map the newly created volume to the host at creation time or map it later. If you did not click the **Create and Map to Host** button when you created the volume, follow the steps in 5.2.1, “Mapping newly created volumes to the host using the wizard” on page 142.

5.2.1 Mapping newly created volumes to the host using the wizard

We continue to map the volume that was created in 5.1, “Provisioning storage from IBM Storwize V7000 to hosts” on page 132. We assume that you followed that procedure and clicked **Continue** as, for example, shown in Figure 5-4 on page 134.

To map the volumes, complete the following steps:

1. Select a host to which the new volume should be attached (Figure 5-20).

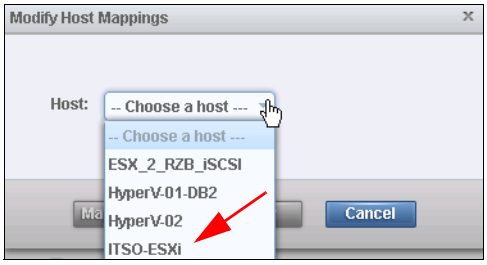

Figure 5-20 Choose a host

2. The Modify Mappings window opens, and your host and the newly created volume are already selected. Click **OK** to map the volume to the host (Figure 5-21).

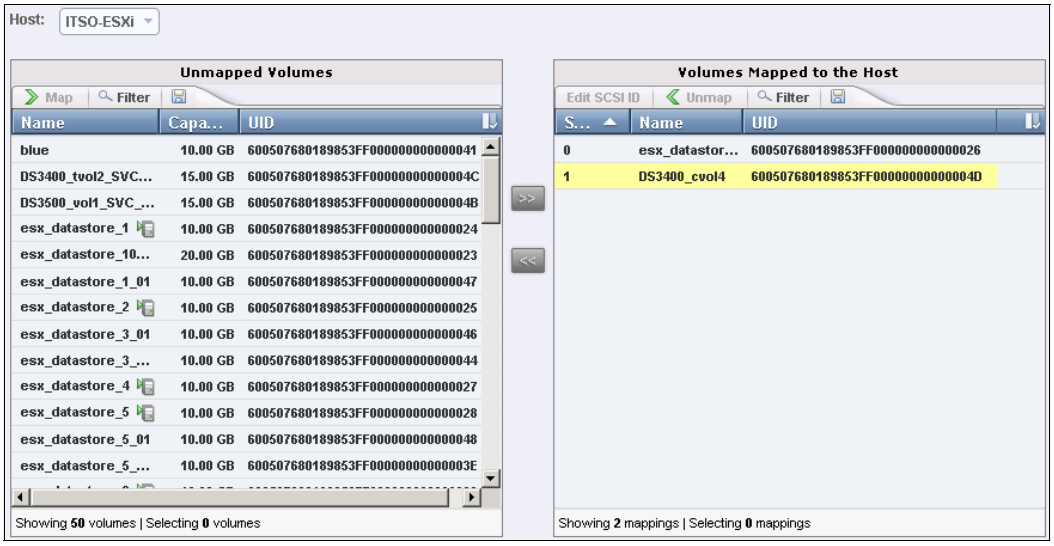

Figure 5-21 Modify mappings

3. After the task completes, click **Apply** and **Close**. The wizard returns to the Volumes window. By double-clicking the volume, you can see the Host Maps (Figure 5-22).

Figure 5-22 Host maps

The host is now able to access the volumes and store data on them. See 5.3, “Discovering volumes on hosts and multipathing” on page 143 for information about discovering the volumes on the host and making additional host settings, if required.

You can also create multiple volumes in preparation for discovering them later, and customize mappings. Advanced host configuration is described in 8.1.1, “Modifying Mappings menu” on page 224.

5.3 Discovering volumes on hosts and multipathing

This section explains how to discover the volumes that were created and mapped in 5.1, “Provisioning storage from IBM Storwize V7000 to hosts” on page 132 and 5.2, “Mapping a volume to the host” on page 142, and how to configure additional multipath settings, if required. We assume that you have completed all previous steps in this book so that the hosts and the IBM Storwize V7000 are prepared:

- Prepare your operating systems for attachment (Chapter 4, “Host configuration” on page 105).
- Create hosts using the GUI (4.3, “Creating hosts using the GUI” on page 121).
- Perform basic volume configuration and host mapping (5.1, “Provisioning storage from IBM Storwize V7000 to hosts” on page 132).

Our examples illustrate how to discover Fibre Channel and iSCSI volumes on Microsoft Windows 2008 and VMware ESX 4.x hosts.

From the dynamic menu of the IBM Storwize V7000 GUI click the **Hosts** icon to open the Hosts menu, and click **Hosts** option (Figure 5-23).

Figure 5-23 Navigate to hosts menu

An overview of the all configured and mapped hosts is displayed (Figure 5-24).

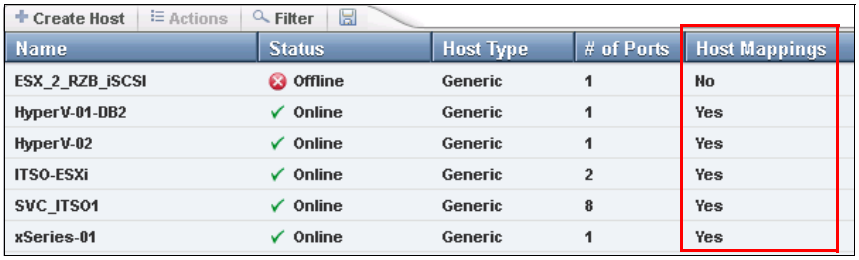

Figure 5-24 The existing hosts

5.3.1 Windows 2008 Fibre Channel volume attachment

To complete Fibre Channel volume attachment in Windows 2008, use the following steps:

1. Right-click your Windows 2008 Fibre Channel host in the Hosts view (Figure 5-25) and select **Properties**.

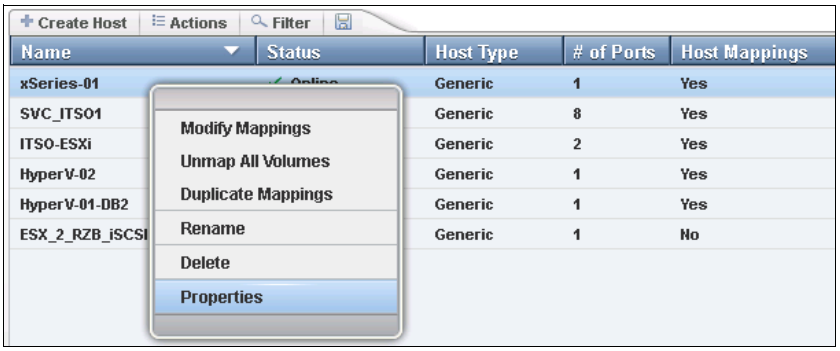

Figure 5-25 Host properties

Navigate to the **Mapped Volumes** tab (Figure 5-26).

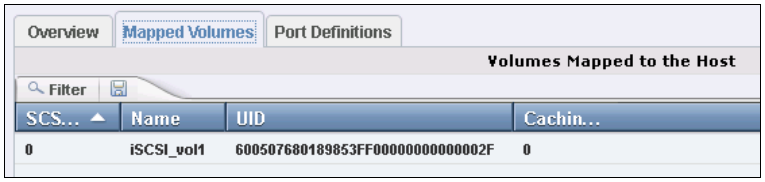

Figure 5-26 Mapped volumes to a host

The host details show you which volumes are currently mapped to the host, and you also see the volume UID and the SCSI ID. In our example, one volume with SCSI ID 0 is mapped to the host.

2. Log on to your Microsoft host and click **Start** → **All Programs** → **Subsystem Device Driver DSM** → **Subsystem Device Driver DSM**. A command-line interface opens. Type the **`datapath query device`** command and press Enter to display IBM Storwize V7000 disks that are connected to this host (Example 5-1).

Example 5-1 datapath query device

```
C:\Program Files\IBM\SDDDSM>datapath query device

Total Devices : 1
DEV#:   0  DEVICE NAME: Disk14 Part0  TYPE: 2145       POLICY: OPTIMIZED
SERIAL: 600507680189853FF00000000000002F
===========================================================================
Path#            Adapter/Hard Disk        State  Mode       Select      Errors
    0    Scsi Port5 Bus0/Disk14 Part0     OPEN   NORMAL          0           0
    1    Scsi Port5 Bus0/Disk14 Part0     OPEN   NORMAL         22           0
    2    Scsi Port6 Bus0/Disk14 Part0     OPEN   NORMAL         19           0
    3    Scsi Port6 Bus0/Disk14 Part0     OPEN   NORMAL          0           0
```

The output provides information about mapped volumes. In our example, one disk is connected, `Disk14`, and four paths to the disk are available (State indicates `OPEN`).

3. Open the Windows Disk Management window (Figure 5-27 on page 146) by clicking **Start** → **Run**, and then type `diskmgmt.msc`, and click **OK**.

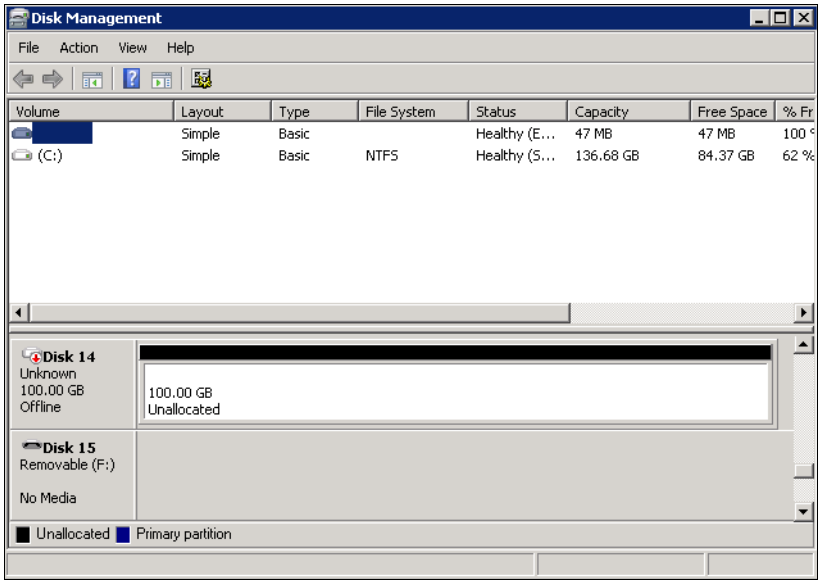

Figure 5-27 Windows Disk Management

Windows device discovery: Usually, Windows discovers new devices, such as disks, by itself (Plug&Play function). If you completed all the steps but do not see any disks, click **Actions** → **Rescan Disk** in Disk Management to discover potential volumes.

4. Right-click the disk in the left pane and select **Online** (Figure 5-28).

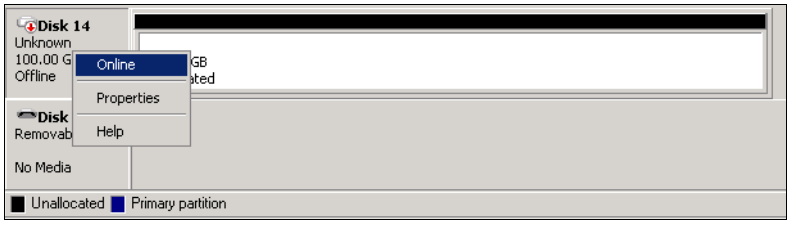

Figure 5-28 Place a disk online

5. Right-click the disk again, select **Initialize Disk** (Figure 5-29), and click **OK**.

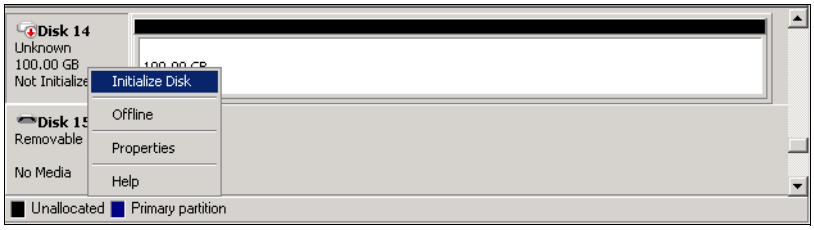

Figure 5-29 Initialize Disk menu

6. Right-click in the right pane and select **New Simple Volume** (Figure 5-30).

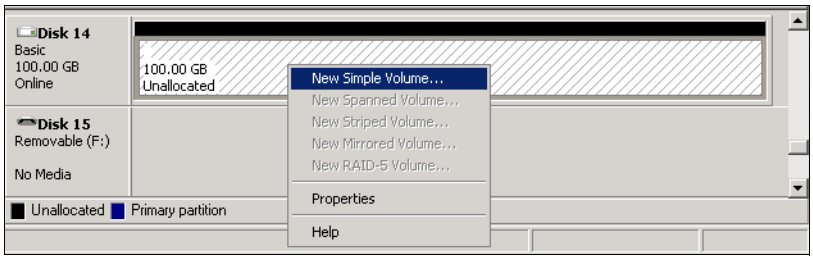

Figure 5-30 New Simple Volume

7. Follow the wizard and the volume is ready to use from your Windows host (Figure 5-31).

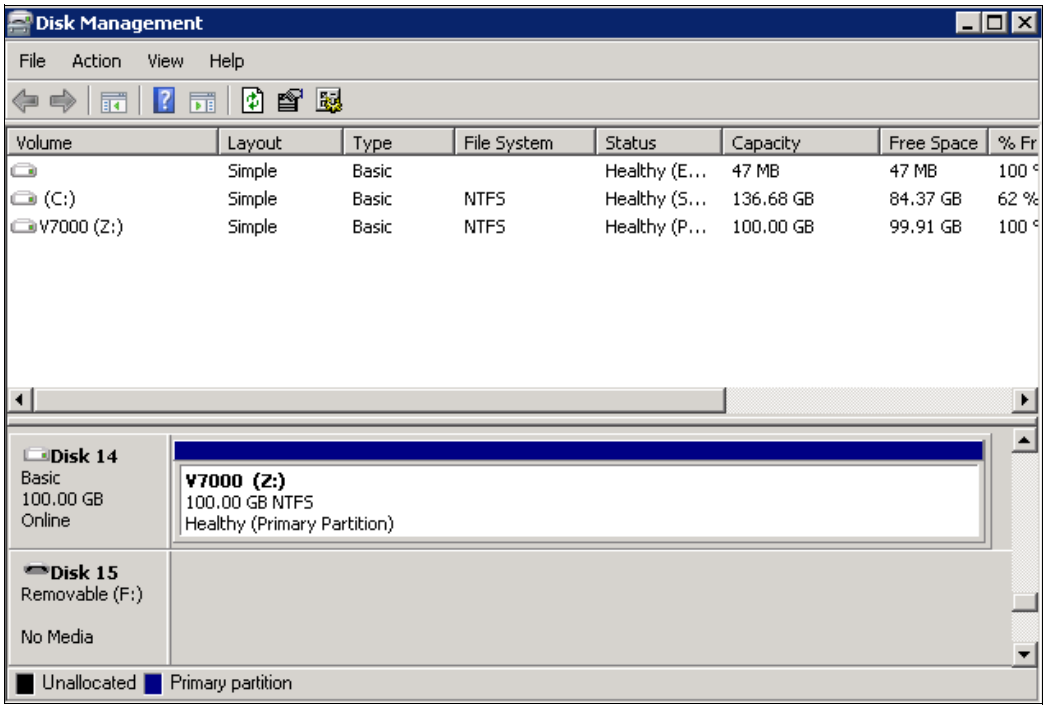

Figure 5-31 Volume is ready to use

The basic setup is now complete, the IBM Storwize V7000 is configured, and the host is prepared to access the volumes over several paths and is able to store data on the storage subsystem.

5.3.2 Windows 2008 iSCSI volume attachment

To perform iSCSI volume attachment in Windows 2008, complete the following steps:

1. Right-click your Windows 2008 iSCSI host in the **Hosts** view (Figure 1), click **Properties**, and click the **Port Definitions** tab (Figure 5-32).

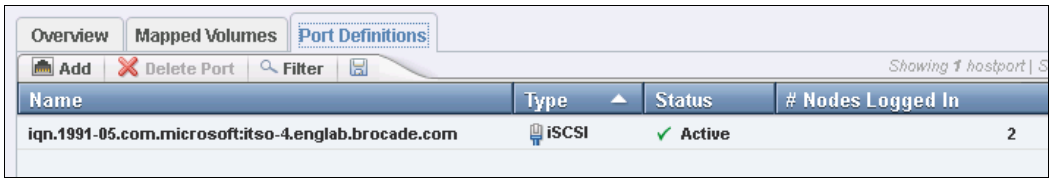

Figure 5-32 iSCSI host address

The host details show you which volumes are currently mapped to the host, and you also see the volume UID and the SCSI ID. In our example, one volume with SCSI ID 0 is mapped to the host.

2. Log on to your Windows 2008 host and click **Start** → **Administrative Tools** → **iSCSI Initiator** to open the iSCSI Configuration tab (Figure 5-33).

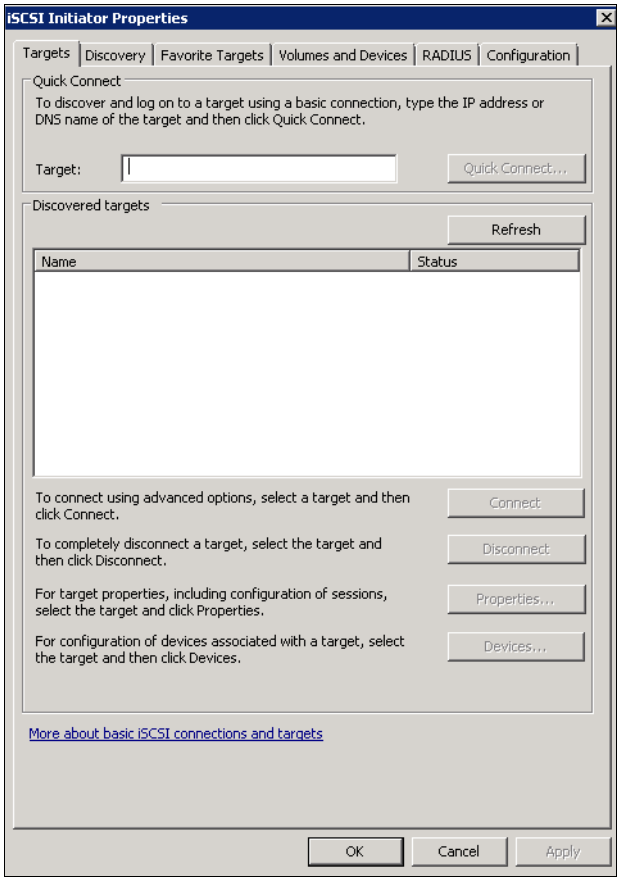

Figure 5-33 Windows iSCSI Configuration tab

3. Enter the IP address of one of the IBM Storwize V7000 iSCSI ports and click **Quick Connect** (Figure 5-34).

> **iSCSI IP addresses:** The iSCSI IP addresses are different from the cluster and canister IP addresses, and they are configured in 4.3.2, “Creating iSCSI hosts” on page 125.

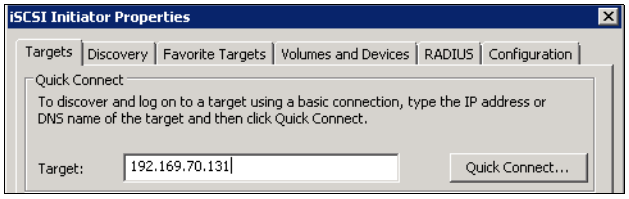

Figure 5-34 iSCSI Quick Connect

The IBM Storwize V7000 initiator is discovered and connected (Figure 5-35).

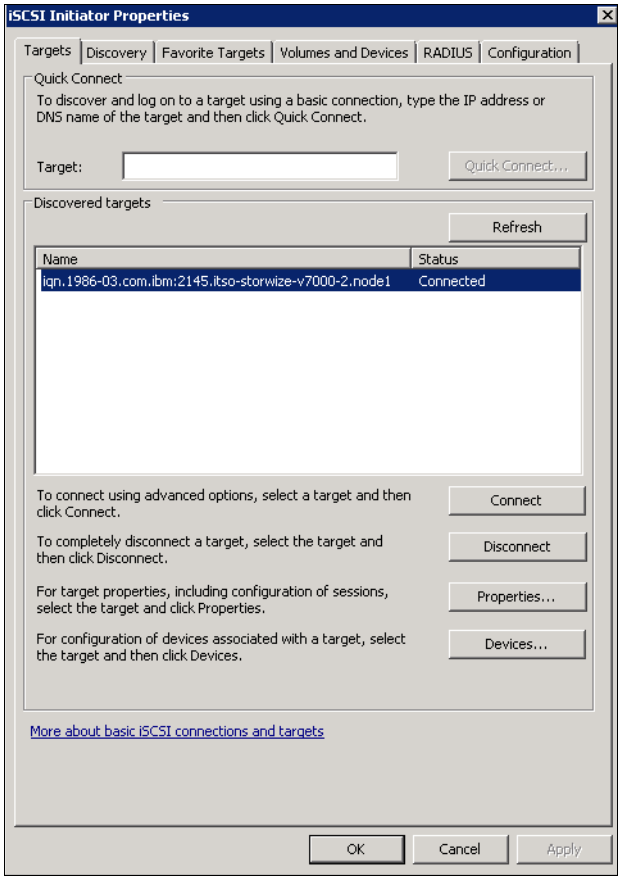

Figure 5-35 iSCSI Initiator target is connected

Now you have completed the steps to connect the storage disk to your iSCSI host, but you are using only a single path at the moment. To enable multipathing for iSCSI targets, more actions are required. Complete the following steps:

1. Click **Start** → **Run** and type `cmd` to open a command prompt. Run the following command and press Enter (Example 5-2):

   ```
   ServerManagerCMD.exe -install Multipath-IO
   ```

   Example 5-2 Installing MPIO

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

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

   Success: Installation succeeded.
   ```

2. Click **Start** → **Administrative Tools** → **MPIO**, click the **Discover Multi-Paths** tab, and select the **Add support for iSCSI devices** check box (Figure 5-36).

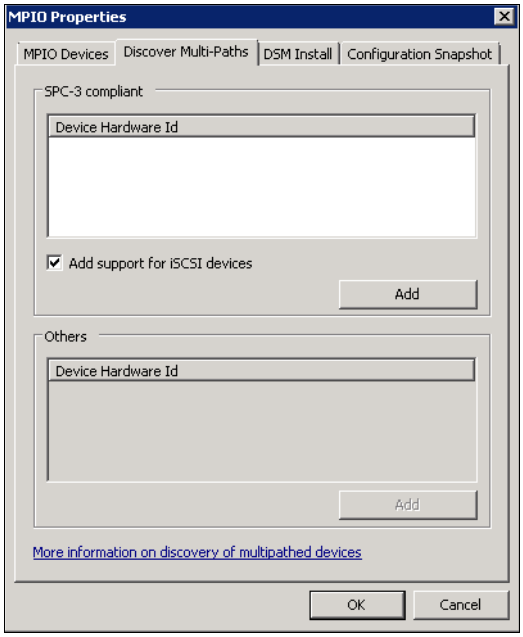

Figure 5-36 Enable iSCSI MPIO

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

4. After reboot, select **Start** → **Administrative Tools** → **iSCSI Initiator** to open the iSCSI Initiator Properties window (Configuration tab). Click the **Discovery** tab (Figure 5-37).

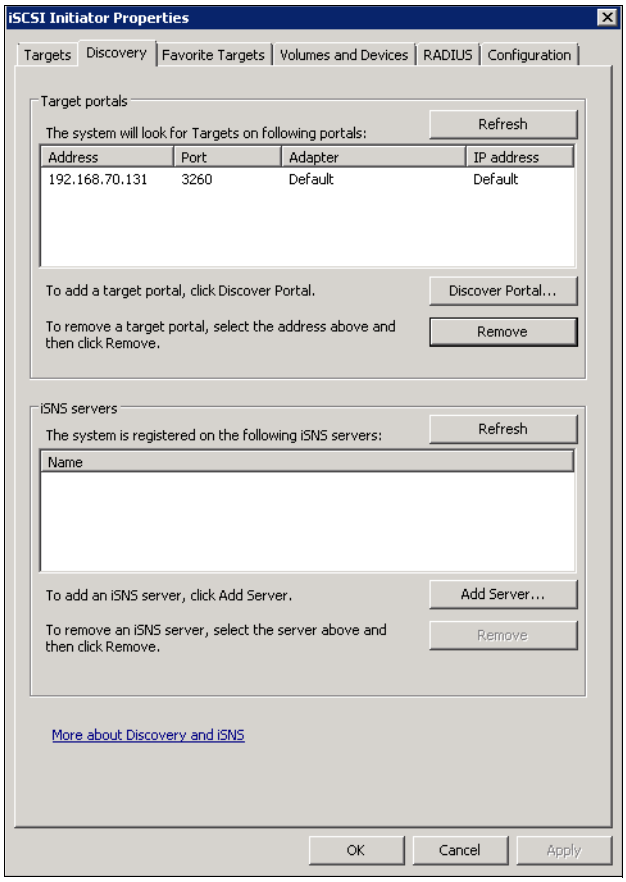

Figure 5-37 iSCSI Properties Discovery tab

5. Click **Discover Portal**, enter the IP address of another IBM Storwize V7000 iSCSI port (Figure 5-38), and click **OK**.

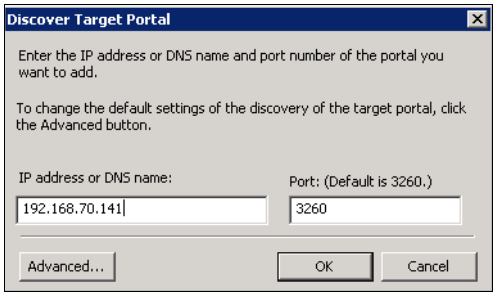

Figure 5-38 Discover Target Portal window

6. Return to the Targets tab (Figure 5-39); the new connection is listed there as Inactive.

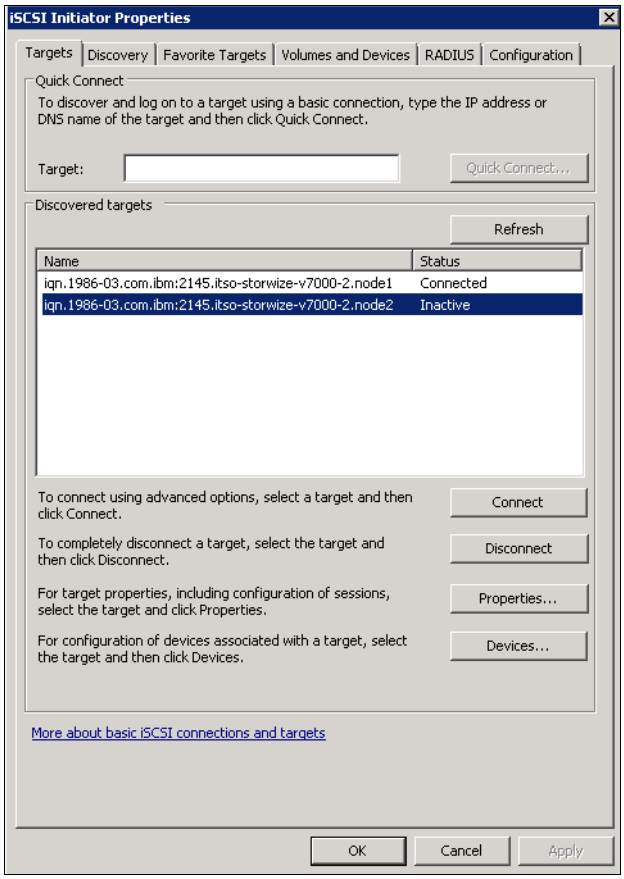

Figure 5-39 Inactive target ports

7. Highlight the inactive port and click **Connect**. The Connect to Target window opens (Figure 5-40).

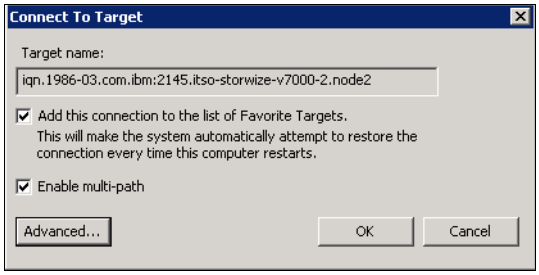

Figure 5-40 Connect to a target

8. Select the **Enable Multipath** check box and click **OK**. The status of the second port now indicates Connected (Figure 5-41).

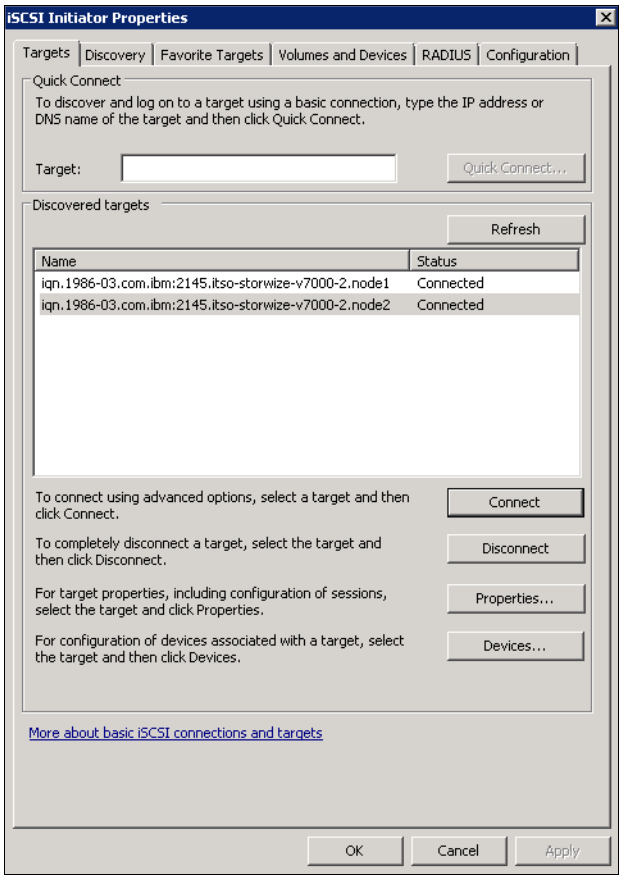

Figure 5-41 Second target port connected

Repeat this step for each IBM Storwize V7000 port you want to use for iSCSI traffic. You may have up to four port paths to the system.

9. Click **Devices** → **MPIO to** make sure that the multipath policy for Windows 2008 is set to the default, which is Round Robin with Subset (Figure 9), and click **OK** to close this view.

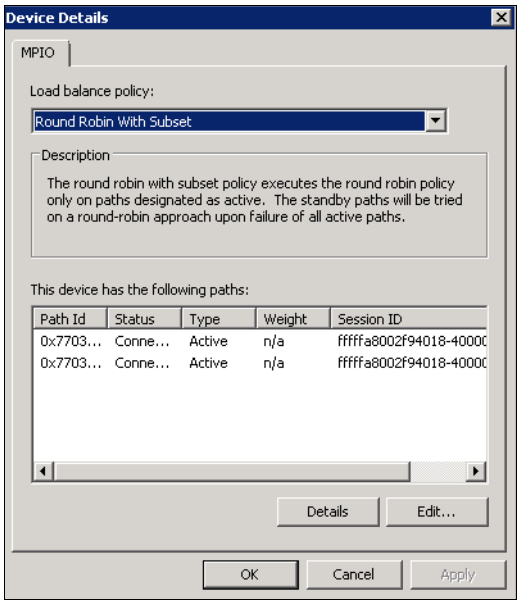

Figure 5-42 Round Robin with Subset

10. Open the Windows Disk Management window (Figure 5-43) by clicking **Start** → **Run**, and then type `diskmgmt.msc` and click **OK**.

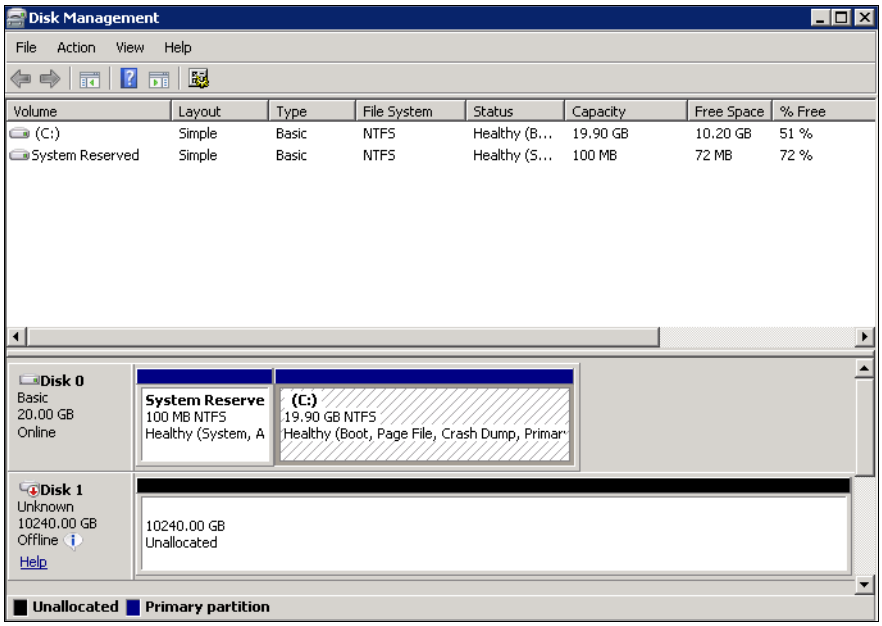

Figure 5-43 Windows Disk Management

11. Set the disk online, initialize it, create a file system on it, and then it is ready to use. The detailed steps of this process are the same as described in 5.3.1, “Windows 2008 Fibre Channel volume attachment” on page 144.

Now the storage disk is ready for use (Figure 5-44). In our example, we mapped a 10 TB disk, which is thin-provisioned, on the IBM Storwize V7000, to a Windows 2008 host using iSCSI.

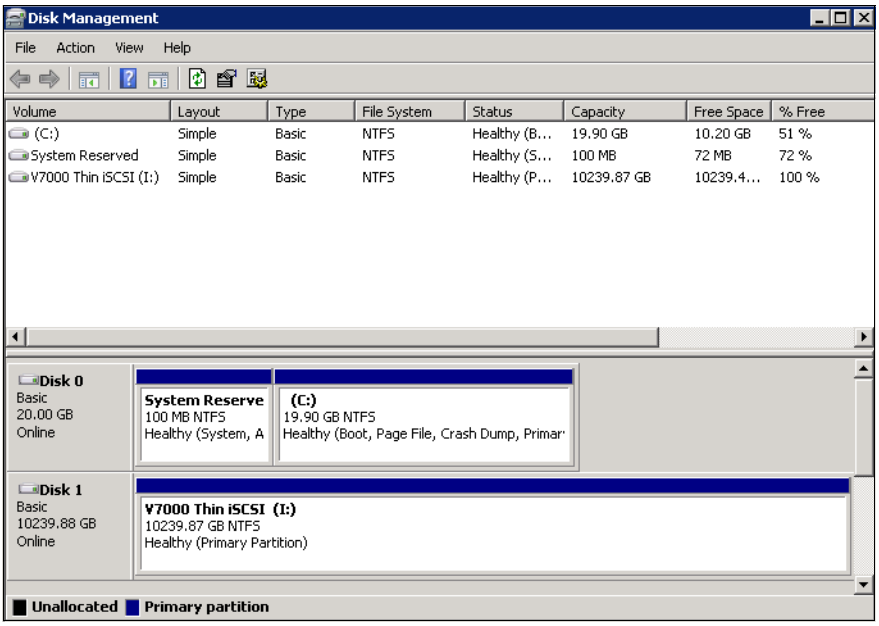

Figure 5-44 Windows Disk Management: Disk is ready to use

5.3.3 VMware ESX Fibre Channel attachment

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

1. Right-click your VMware ESX Fibre Channel host in the Hosts view, select **Properties**, and then click the **Mapped Volumes** tab (Figure 5-45).

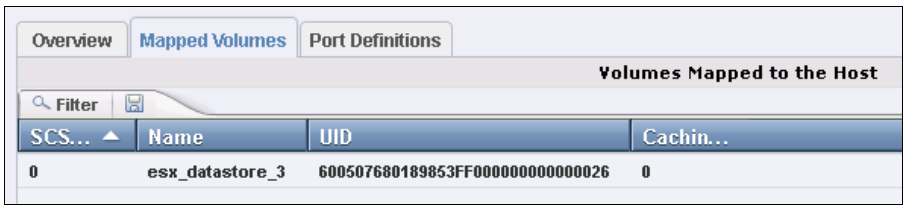

Figure 5-45 Mapped volume to ESX FC host

The Host Details window shows that one volume is connected to the ESX FC host using SCSI ID 1. The UID of the volume is also displayed.

2. Connect to your VMware ESX Server using the vSphere client, navigate to the **Configuration** tab, and select **Storage Adapters** (Figure 5-46).

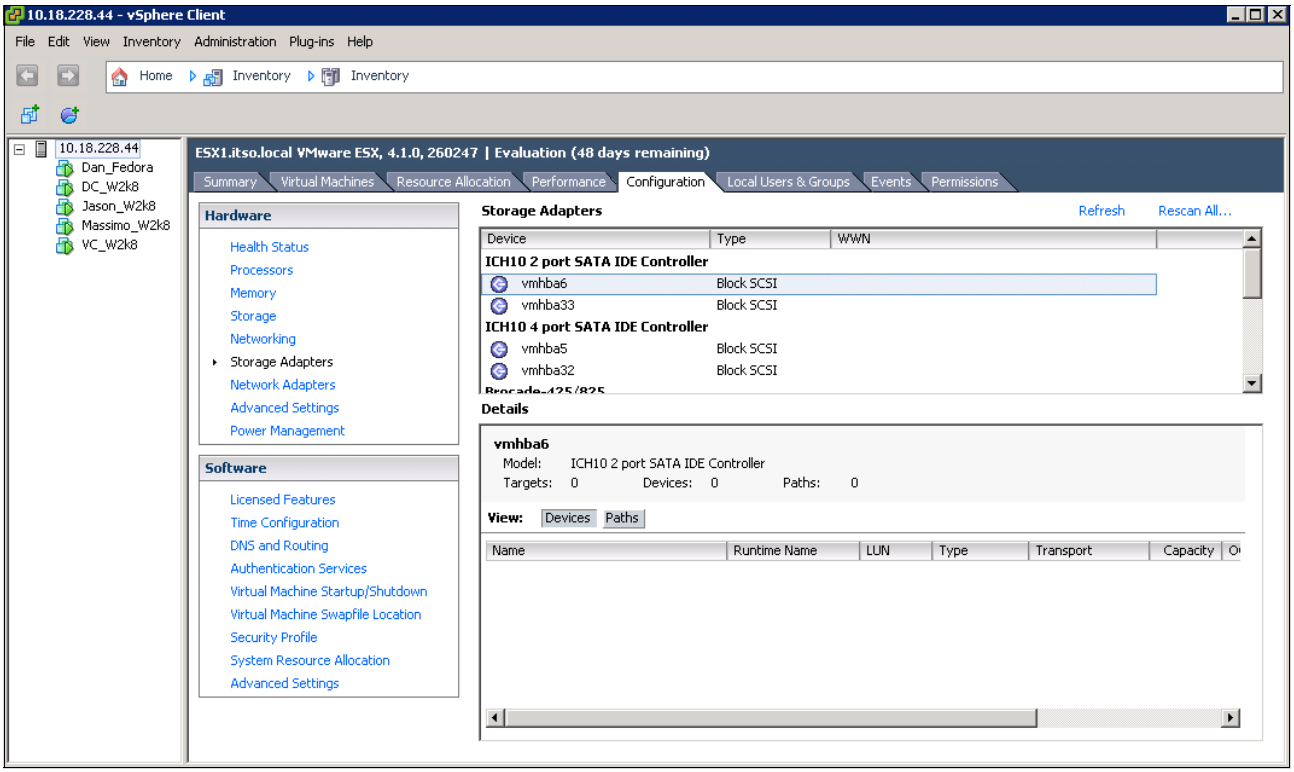

Figure 5-46 vSphere Client: Storage adapters

3. Select **Rescan All** and click **OK** (Figure 5-47) to scan for new storage devices.

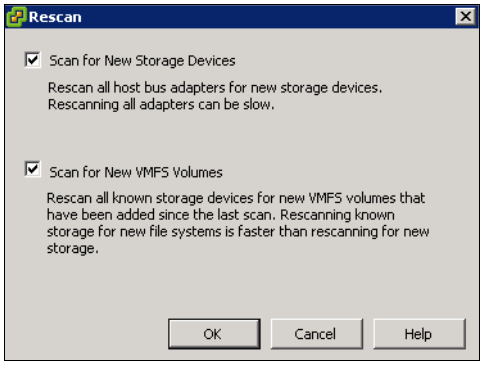

Figure 5-47 Rescan devices

4. Select **Storage** and click **Add Storage** (Figure 5-48).

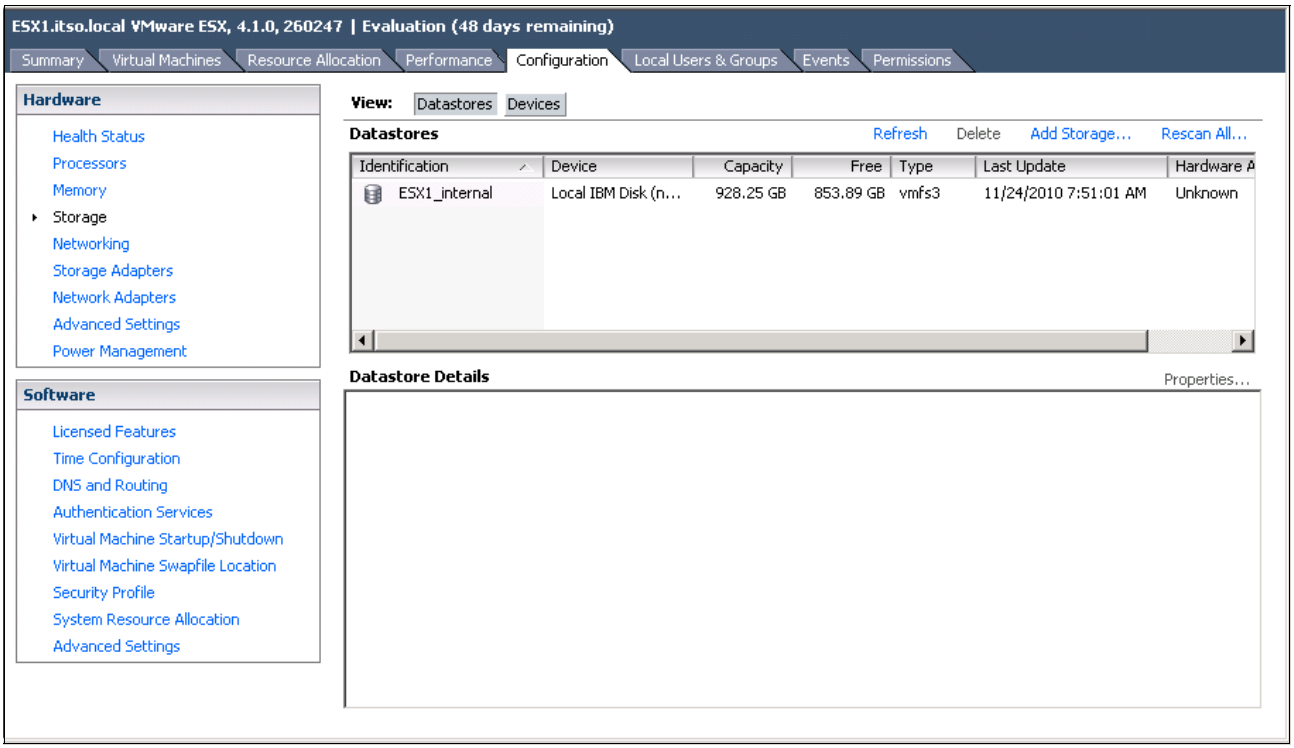

Figure 5-48 vSphere Client: Storage

5. The Add Storage wizard opens. Click **Select Disk/LUN** and click **Next**. The IBM Storwize V7000 disk is displayed (Figure 5-49). Select it and click **Next**.

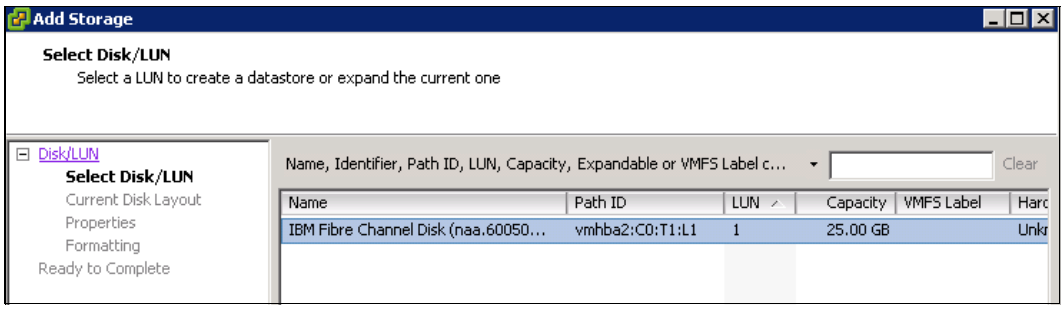

Figure 5-49 Select Disk/LUN menu

6. Follow the wizard to complete the attachment of the disk. After you click **Finish**, the wizard closes and you return to the storage view.

Figure 5-50 shows that the new volume is added to the configuration.

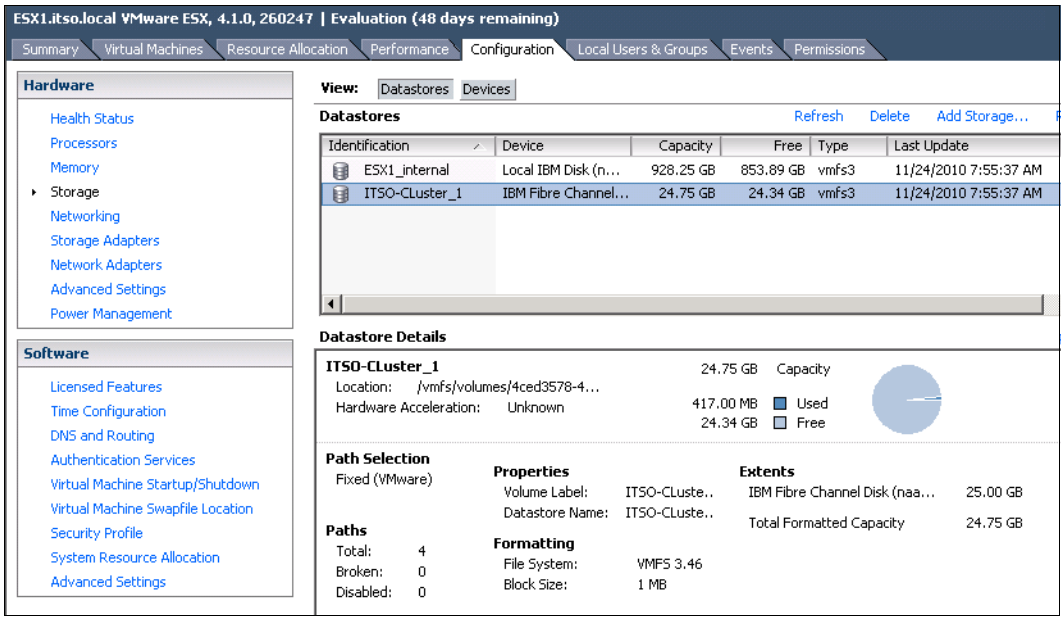

Figure 5-50 Add Storage task complete

7. Highlight the new data store and click **Properties** to see the details of it (Figure 5-51).

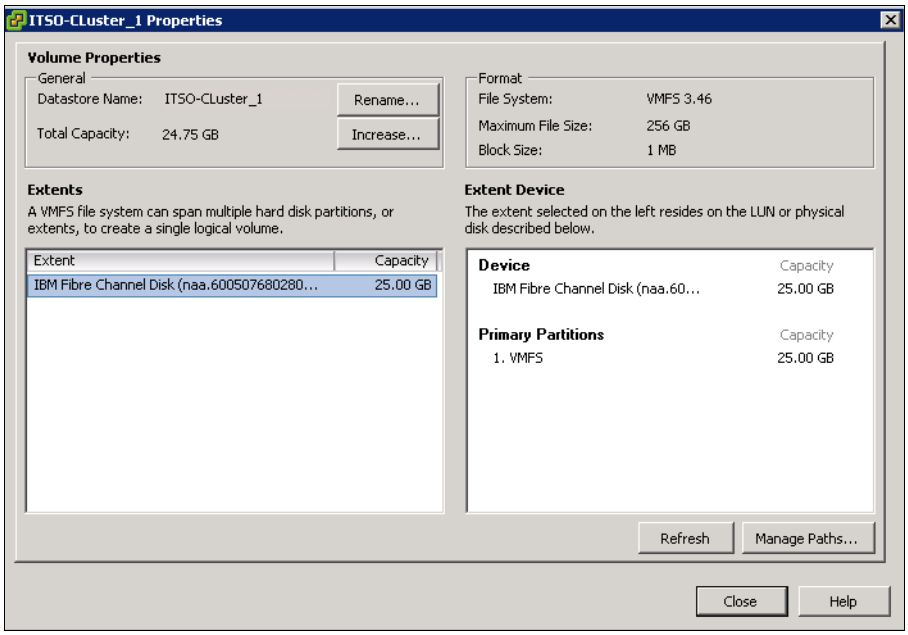

Figure 5-51 Data store properties

8. Click **Manage Paths** to customize the multipath settings. Select **Round Robin** (Figure 5-52) and click **Change**.

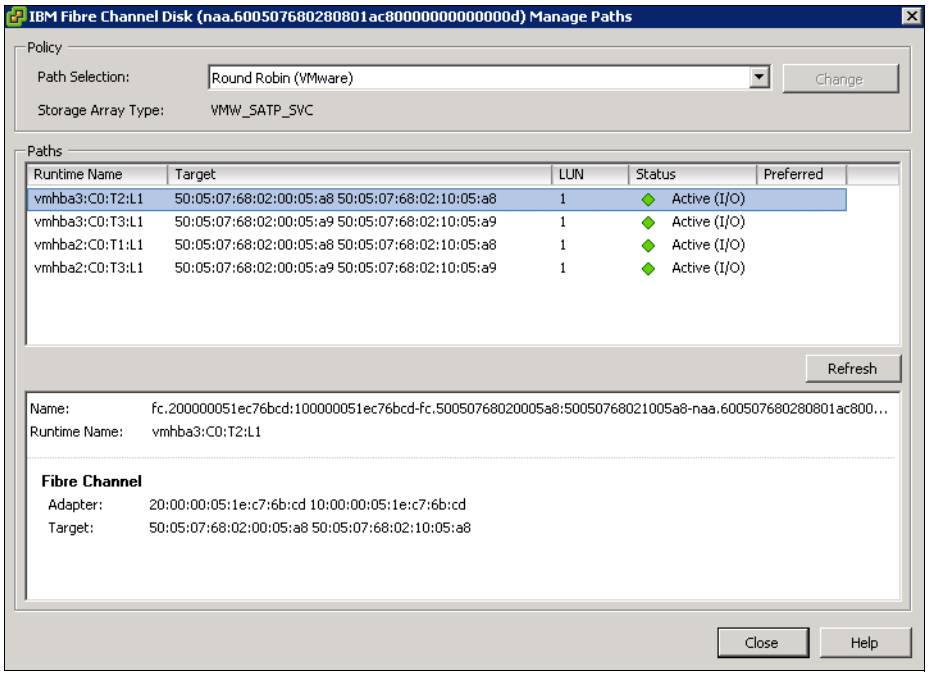

Figure 5-52 Select a data store multipath setting

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

5.3.4 VMware ESX iSCSI attachment

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

1. Right-click your VMware ESX Fibre iSCSI host in the Hosts view and select **Properties**. Click the **Mapped Volumes** tab (Figure 5-53).

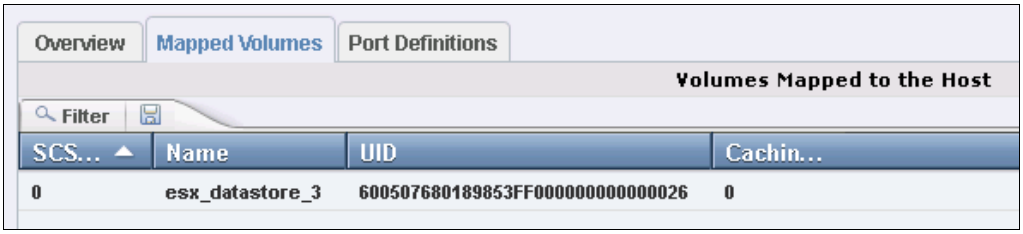

Figure 5-53 iSCSI ESX host properties

The Host Details window shows that one volume is connected to the ESX iSCSI host using SCSI ID 0. The UID of the volume is also displayed.

2. Connect to your VMware ESX Server using the vSphere Client, click the **Configuration** tab (Figure 5-54), and select **Storage Adapters**.

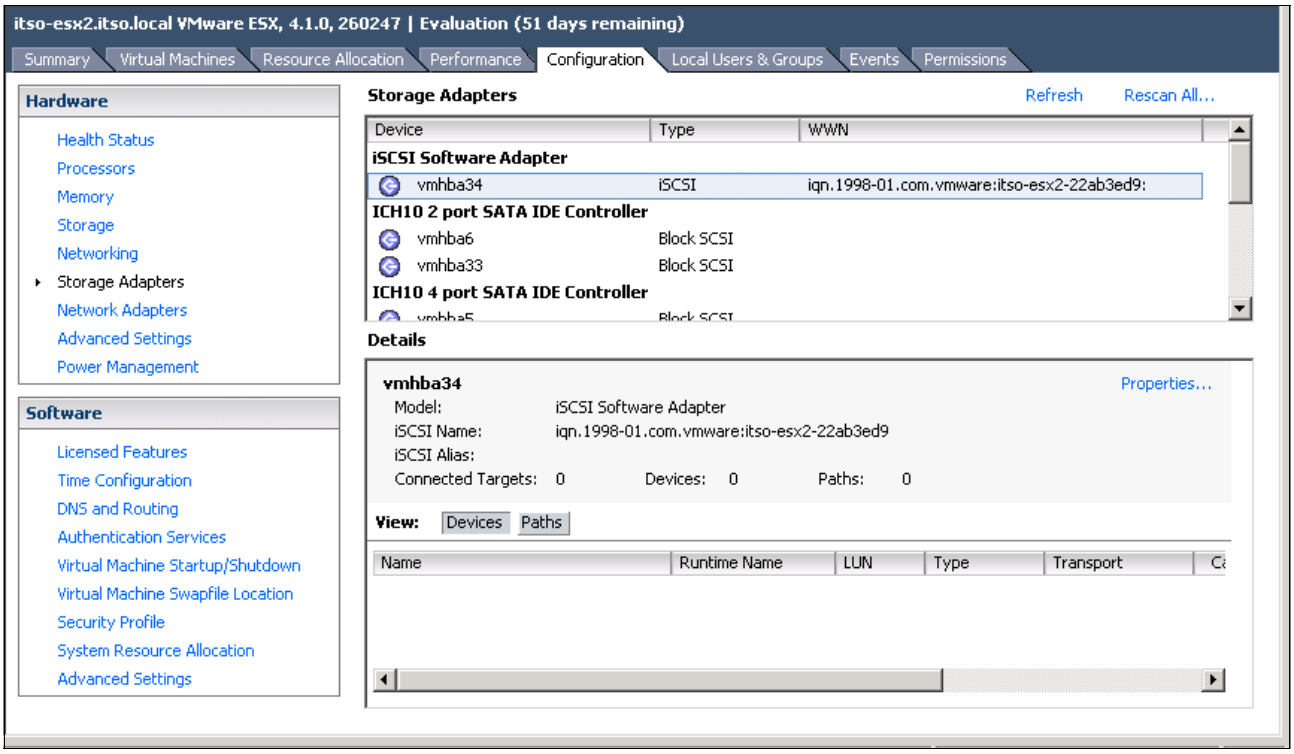

Figure 5-54 vSphere Client: Storage

3. Select **iSCSI Software Initiator** and click **Properties**. The iSCSI initiator properties window opens. Select the **Dynamic Discovery** tab (Figure 5-55) and click **Add**.

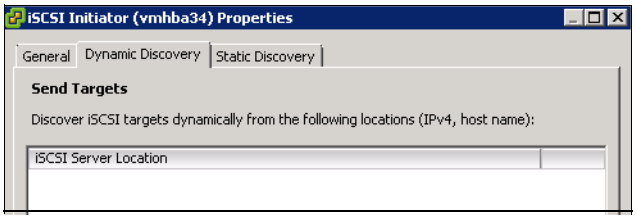

Figure 5-55 iSCSI Initiator properties

4. To add a target, enter the target IP address (Figure 5-56 on page 161). The target IP address is the IP address of a node in the I/O group from which you are mapping the iSCSI volume. Keep the IP port number at the default value of 3260, and click **OK**. The connection between the initiator and target is established.

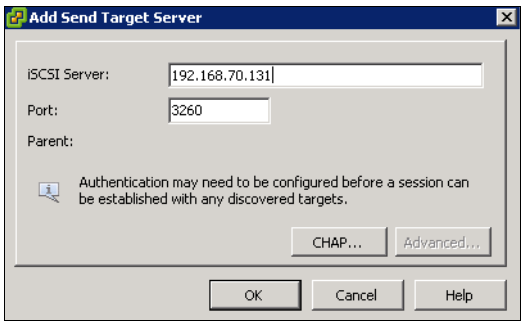

Figure 5-56 Enter a target IP address

Repeat this step for each IBM Storwize V7000 iSCSI port that you want to use for iSCSI connections.

iSCSI IP addresses: The iSCSI IP addresses are different from the cluster and canister IP addresses; they have been configured in 4.3.2, “Creating iSCSI hosts” on page 125.

5. After you have added all the ports required, close the iSCSI Initiator properties by clicking **Close** (Figure 5-55 on page 160).

   You are prompted to rescan for new storage devices. Confirm the scan by clicking **Yes** (Figure 5-57).

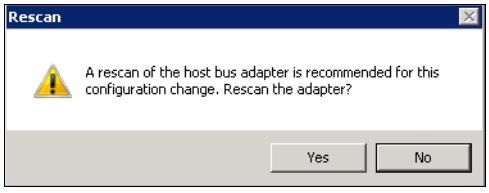

Figure 5-57 Confirm the rescan

6. Go to the storage view shown in Figure 5-58 and click **Add Storage**.

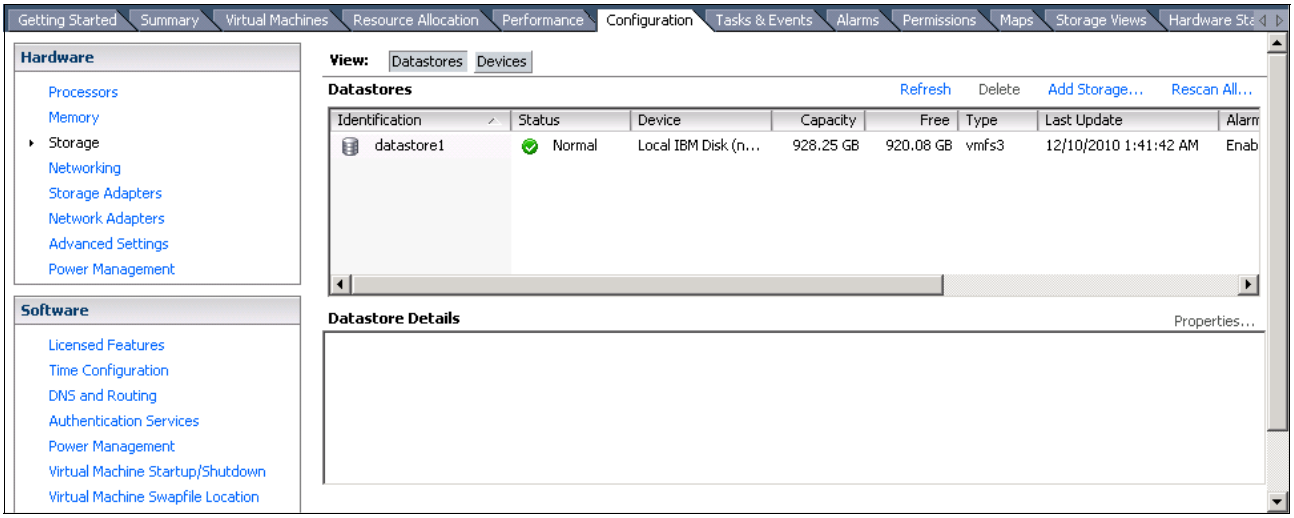

Figure 5-58 Click the Add Storage menu

7. The Add Storage wizard opens (Figure 5-59). Select **Disk/LUN** and click **Next**.

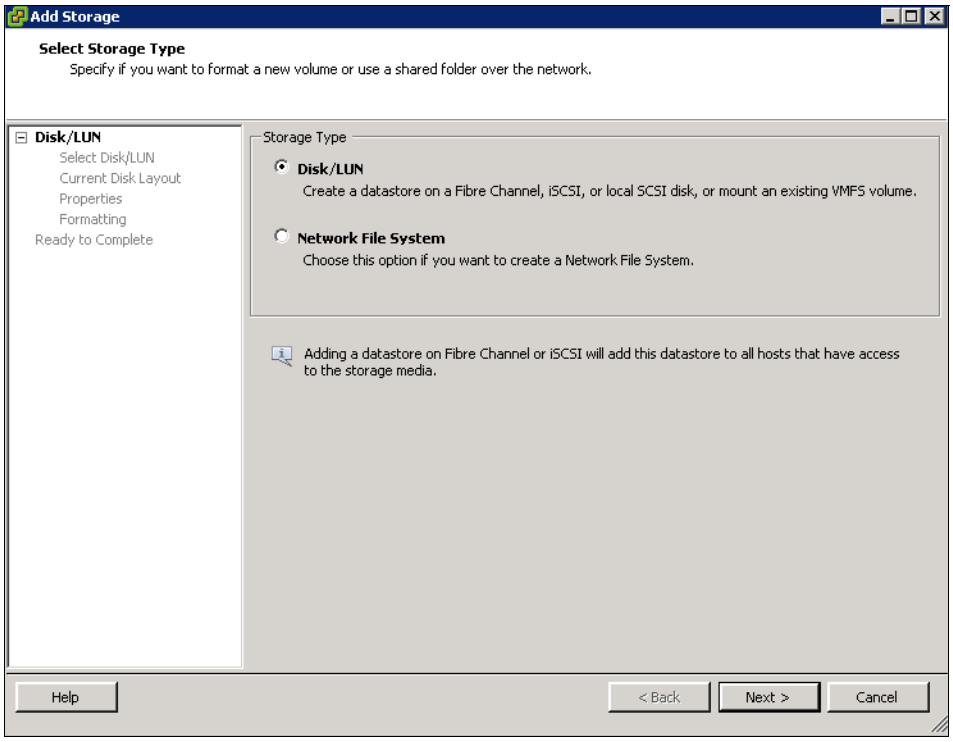

Figure 5-59 Select Disk/LUN menu

8. The new iSCSI LUN displays. Select it and click **Next** (Figure 5-60).

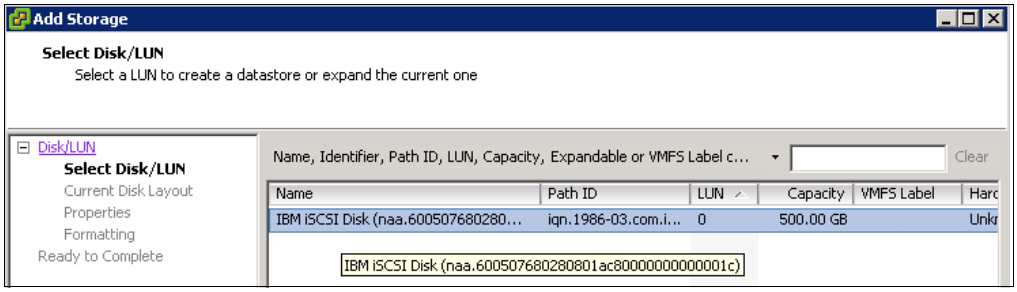

Figure 5-60 Select iSCSI LUN menu

9. Review the disk layout and click **Next** (Figure 5-61).

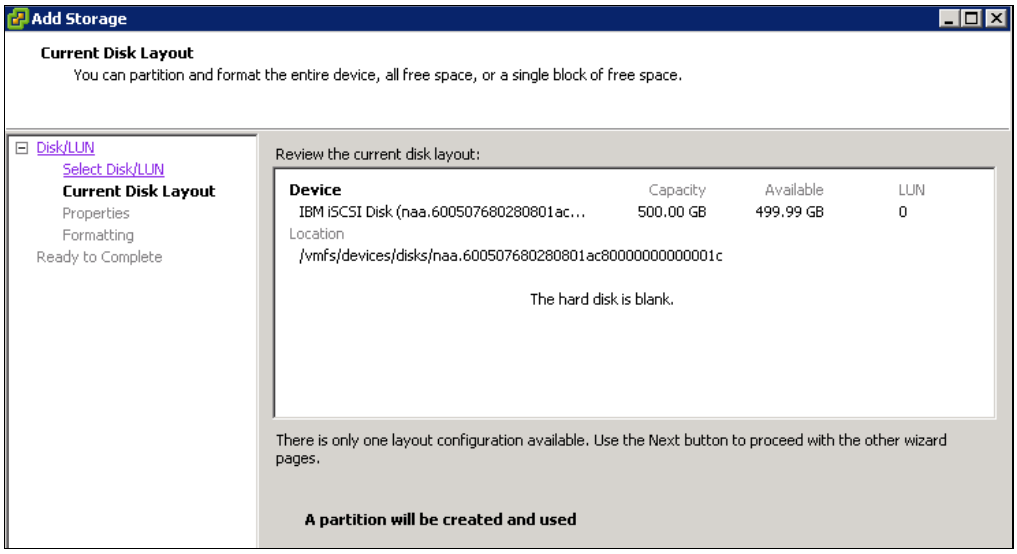

Figure 5-61 Current Disk Layout

10. Enter a name for the data store and click **Next** (Figure 5-62).

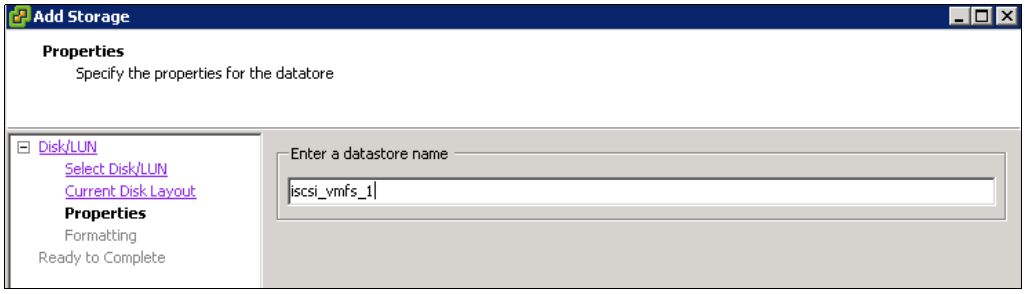

Figure 5-62 Enter a data store name

11. Select the maximum file size and click **Next** (Figure 5-63).

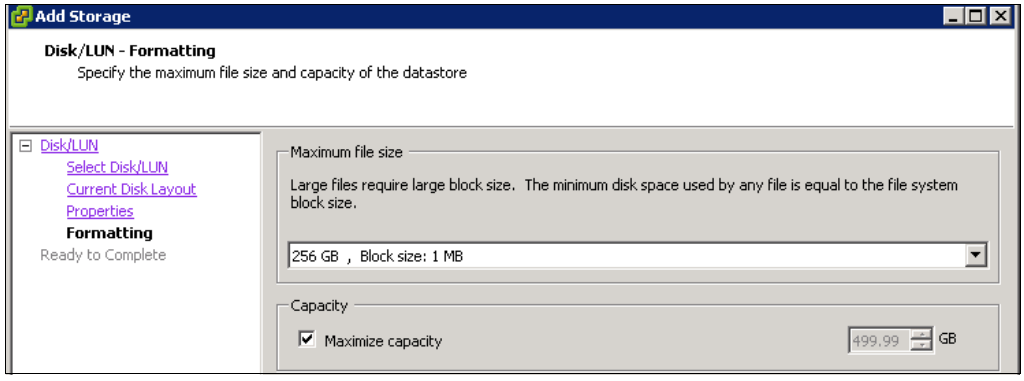

Figure 5-63 Maximum file size

12. Review your selections and click **Finish** (Figure 5-64).

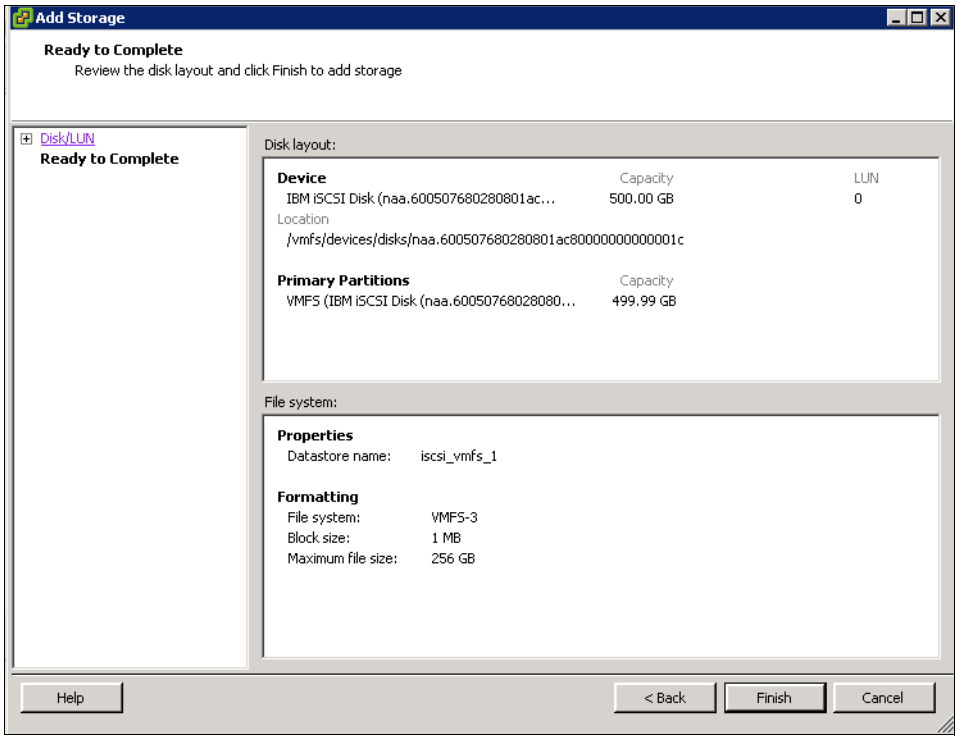

Figure 5-64 Finish the wizard

The new iSCSI LUN is now in the process of being added; this task can be several minutes. After the tasks complete, the new data store is listed in the storage view (Figure 5-65).

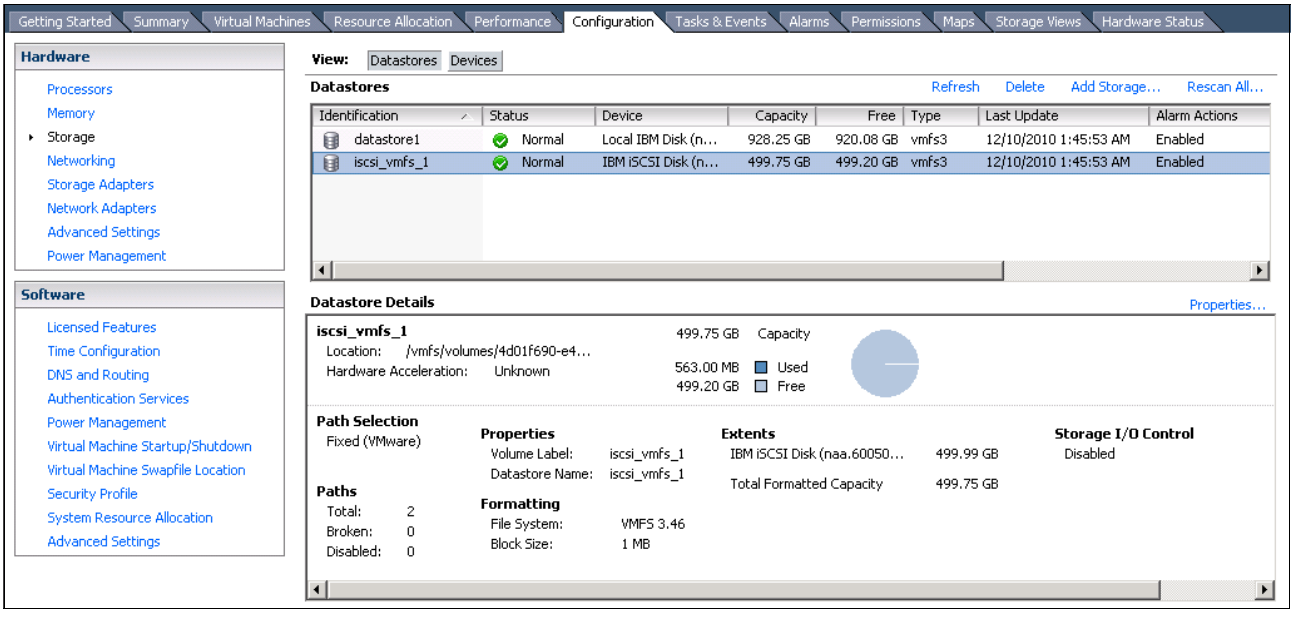

Figure 5-65 New data store available

13. Highlight the new data store and click **Properties** to open and review the data store settings (Figure 5-66).

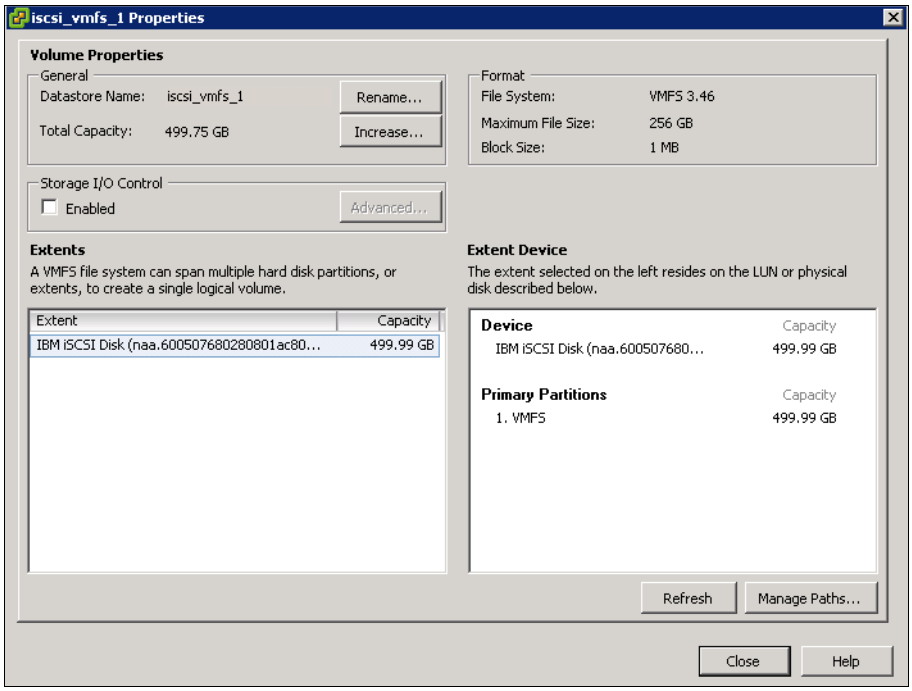

Figure 5-66 Data store properties

14. Click **Manage Paths**, select **Round Robin** as the multipath policy (Figure 5-67), and then click **Change**.

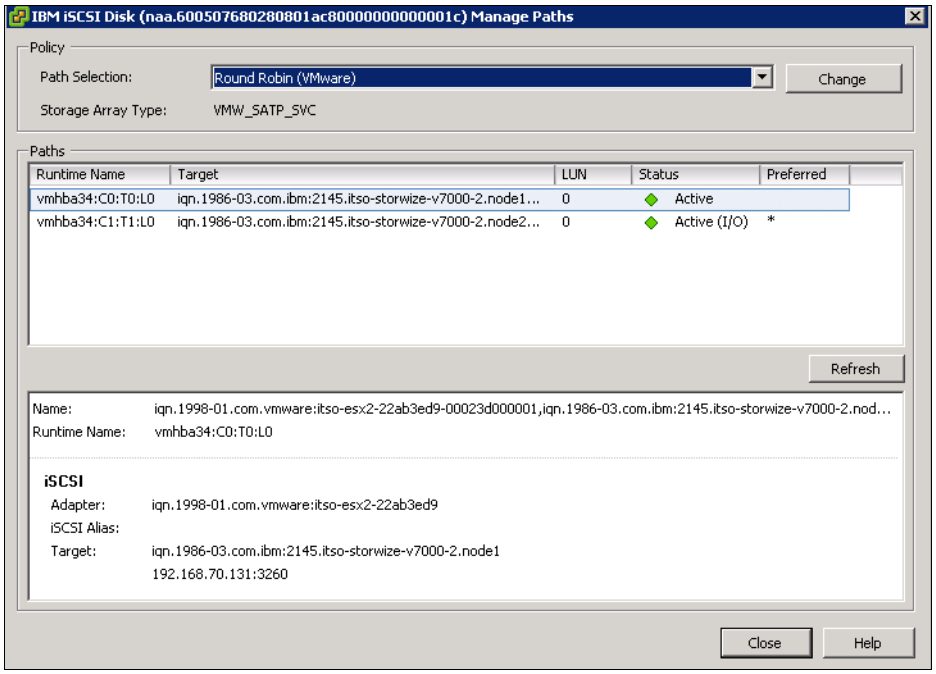

Figure 5-67 Change the multipath policy

15. Click **Close** twice to return to the storage view, and now the storage disk is available and ready to use for your VMware ESX server using an iSCSI attachment.

6

Migration wizard

In this chapter, we describe how to migrate existing data from an older storage system to the IBM Storwize V7000 storage system. Because of this migration applications can benefit from the new features, such as Easy Tier, Thin Provisioning, Real-time Compression, an intuitive management GUI, and advanced storage replication functions that better support applications.

IBM Storwize V7000 provides a storage migration wizard to guide you through the entire procedure.

6.1 Preparing for data migration

In this section, we show you how to prepare external storage systems and IBM Storwize V7000 for data migration. Data migration is always an important step when setting up new storage systems in an existing storage infrastructure.

Before attaching any external storage systems to the IBM Storwize V7000, see the IBM Storwize V7000 support matrix at the following address:

http://www.ibm.com/support/docview.wss?uid=ssg1S1004450

When migrating data from an external storage system to IBM Storwize V7000 (external storage system is removed when completed), you can temporarily configure the external virtualization license setting. Configuring the external license setting prevents messages from being sent that indicate that you are in violation of the license agreement. When the migration is complete, the external virtualization license must be reset to its original limit.

External storage: When external storage is virtualized by IBM Storwize V7000, a per-enclosure external virtualization license is required. Contact your IBM account team or IBM Business Partner for further assistance, if required.

You can temporarily set the license without any charge only during the migration process.

To prepare the data migration, external storage systems must be configured so they are under IBM Storwize V7000 control. Complete the following steps:

1. Stop host I/O to the external storage LUNs that need to be migrated.
2. Remove zones between the hosts and the storage system from which you are migrating.
3. Update your host device drivers, including your multipath driver, and configure them for attachment to the IBM Storwize V7000 system.
4. Create a storage system zone between the storage system being migrated and IBM Storwize V7000 system, and the host zones for the host attachment.
5. Unmap the LUNs in the external storage system to the host and map them to the IBM Storwize V7000 system.
6. Verify that the IBM Storwize V7000 has discovered the LUNs as unmanaged MDisks.

6.2 Data migration using the migration wizard

In this section, we describe how to migrate existing data on external storage systems using the IBM Storwize V7000 storage migration wizard.

We use the wizard that was designed specifically for this scenario to guide you through the process.

Two ways are available to access the menu options for starting a migration:

- From the dynamic menu select **Home** → **Overview**, and then click **Migrate Storage** from the **Suggested Tasks** drop-down menu (Figure 6-1).

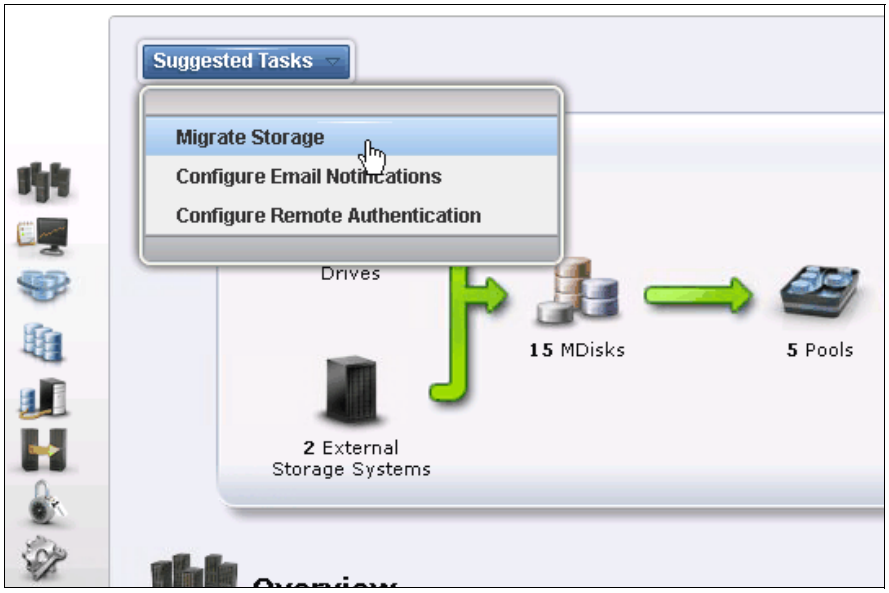

Figure 6-1 Getting Started window with Migrate Storage option displayed

- Click the **Pools** icon and click **System Migration** (Figure 6-2).

Figure 6-2 Pools function icon with Migration option

Whichever method you use, the storage migration window opens (Figure 6-3). Click **Start New Migration** to start the storage migration wizard.

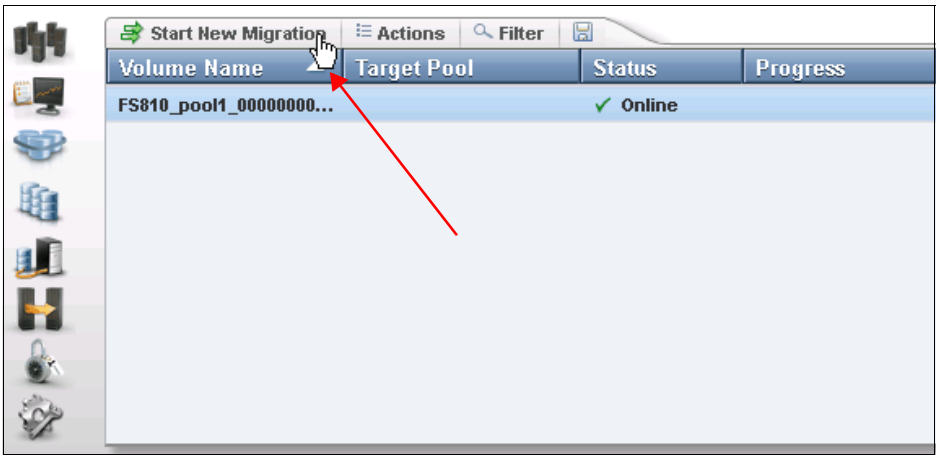

Figure 6-3 Storage migration window

With the IBM Storwize V7000 storage migration wizard, you can more easily migrate your existing data. Complete the following steps:

1. When the wizard starts, confirm the restrictions and prerequisites by selecting the check boxes, and click **Next** (Figure 6-4).

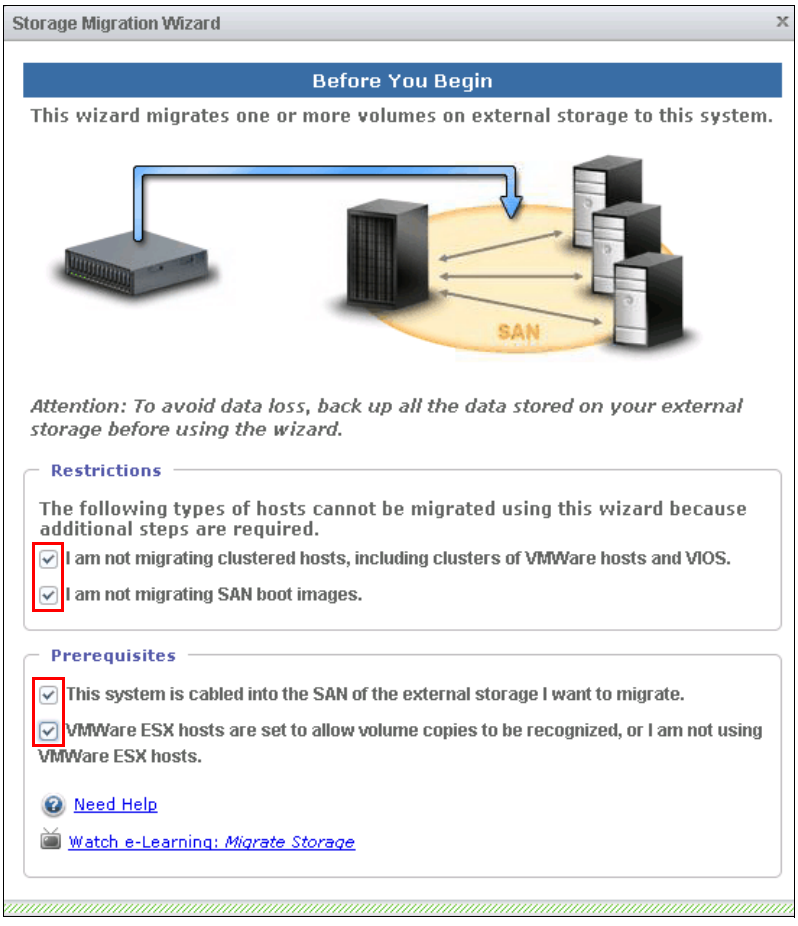

Figure 6-4 Before You Begin window

Avoiding data loss: To avoid any potential data loss, back up all the data stored on your external storage before using the wizard.

The restrictions and prerequisites for using the storage migration wizard are as follows:

– Restrictions:

- Do not use the storage migration wizard to migrate cluster hosts, including clusters of VMware hosts and VIOS.
- Do not use the storage migration wizard to migrate SAN Boot images.

If you have either of these two environments, you must migrate them outside of this wizard. You can find more information about this topic at the following address:

`http://publib.boulder.ibm.com/infocenter/storwize/ic/index.jsp`

The VMware ESX Storage vMotion feature might be an alternative to migrating VMware clusters. For more information about this topic, see the following address:

`http://www.vmware.com/products/vsphere/features/vmotion.html`

– Prerequisites:

- Make sure that the IBM Storwize V7000 Fibre Channel ports have been connected to the SAN fabrics to which the external disk controller and hosts you want to migrate from are connected.
- If you have VMware ESX hosts in the migration, make sure the VMware ESX hosts are set to allow volume copies to be recognized.

Click **Next** to go to the next step.

2. In the second step of the storage migration wizard, prepare the environment for migration, and click **Next** (Figure 6-5).

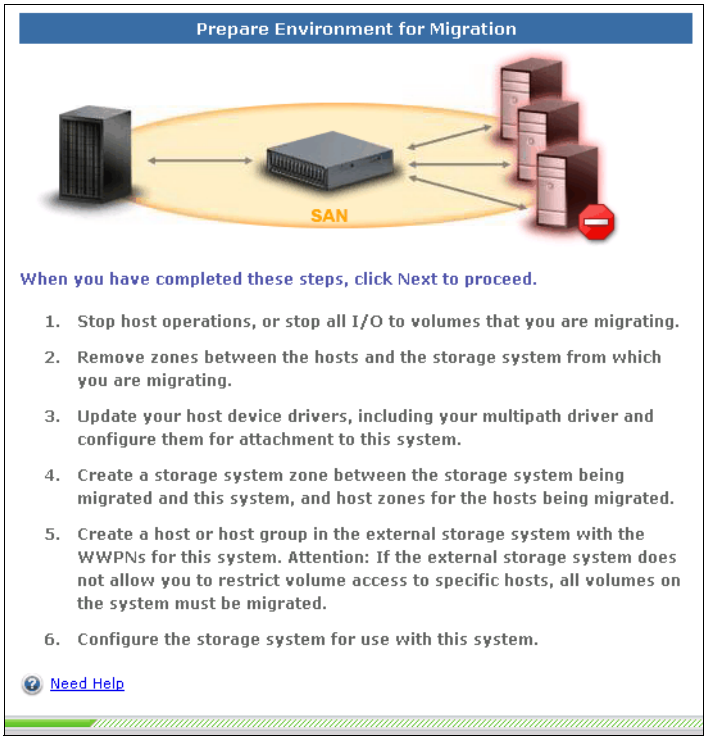

Figure 6-5 Prepare the environment for migration

3. Figure 6-6 shows the third step (panel) of the storage migration wizard. Do everything that is listed as part of step 3. and complete. Be sure to record the information mentioned in this step, because you need it for later steps. Click **Next**.

SCSI ID: You might need to record the SCSI ID of the volume to mapped to the host. Some operating systems do not support changing the SCSI ID during the migration.

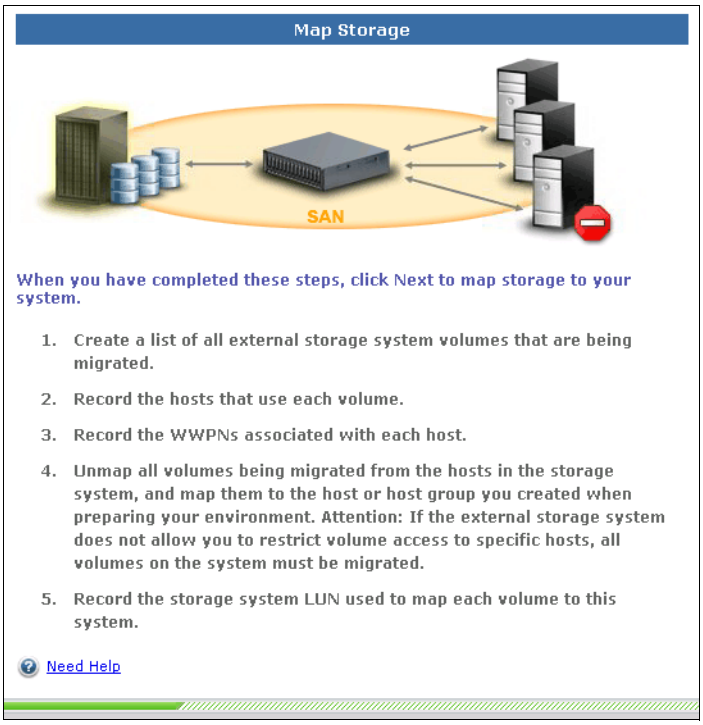

Figure 6-6 Map Storage

After you click **Next,** the IBM Storwize V7000 starts to discover external devices (if you correctly zoned the external storage systems with the IBM Storwize V7000 and mapped the LUNs; see Figure 6-7).

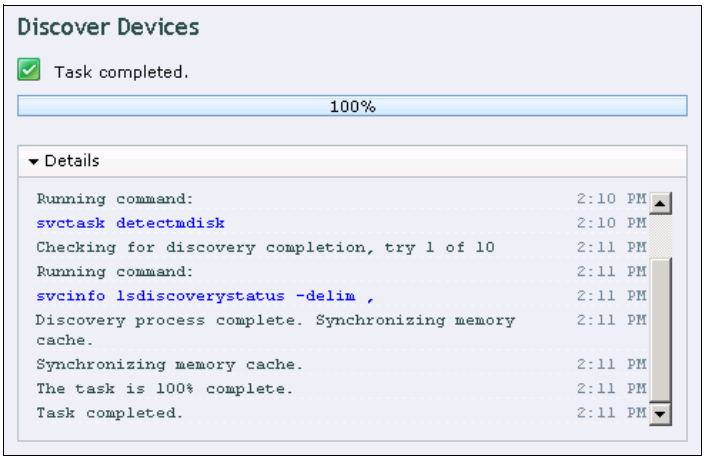

Figure 6-7 Discovering MDisks

4. When discovery completes, the IBM Storwize V7000 shows the MDisks found (Figure 6-8). Select the MDisks you want to migrate and click **Next**.

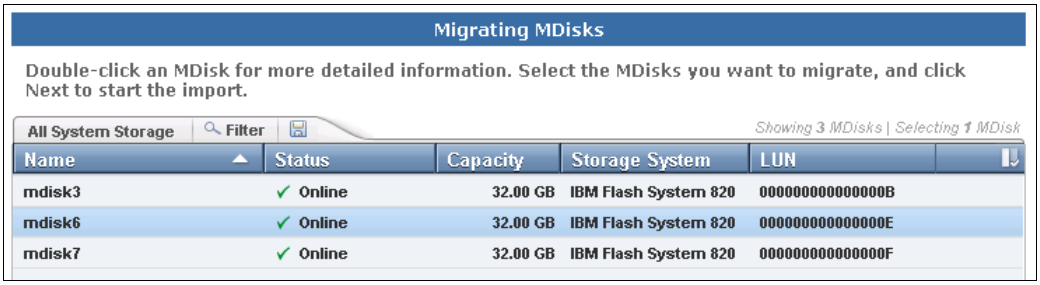

Figure 6-8 Discovered MDisks

If the MDisks that are the subject of the migration are present in the list, select them and click **Next**. The IBM Storwize V7000 starts to import the MDisks that you selected. If the MDisks that need migration are not in the list, check your zone configuration and LUN mapping, and reissue the **Detect MDisks** to trigger the IBM Storwize V7000 discovery procedure again.

You can select one or more MDisks as required, and detailed information about the MDisk is visible by double-clicking it.

Figure 6-8 shows that three LUNs are discovered as MDisks that are candidates for migration. In your particular situation, you might need to reference the information you recorded earlier to identify these MDisks. In our example, MDisk6 has been planned for migration to the new Storwize V7000.

When you click **Next** in this step, the IBM Storwize V7000 completes the import of the MDisk with the host's data and a storage pool is created (Figure 6-9). The MDisks are added to the pool and image mode of the volumes (with the same size as the MDisks) is enabled. These MDisks are ready for mapping back to the original hosts.

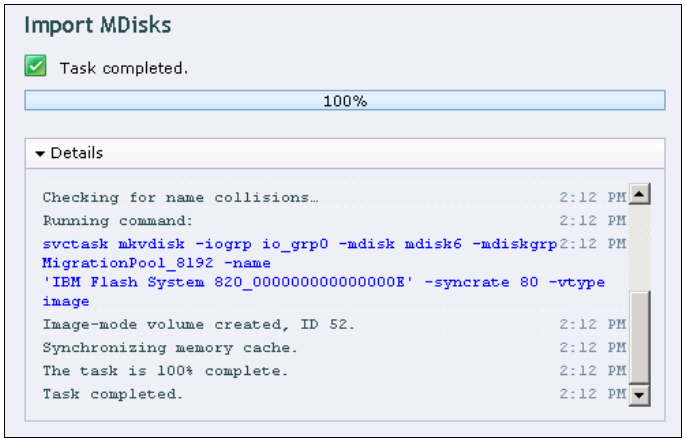

Figure 6-9 Importing discovered MDisks

5. Configure the host that needs to access the data after the migration, or create new hosts as needed, and click **Next** (Figure 6-10). In our case, we plan to assign the migrated volume to the host SVC_ITSO1.

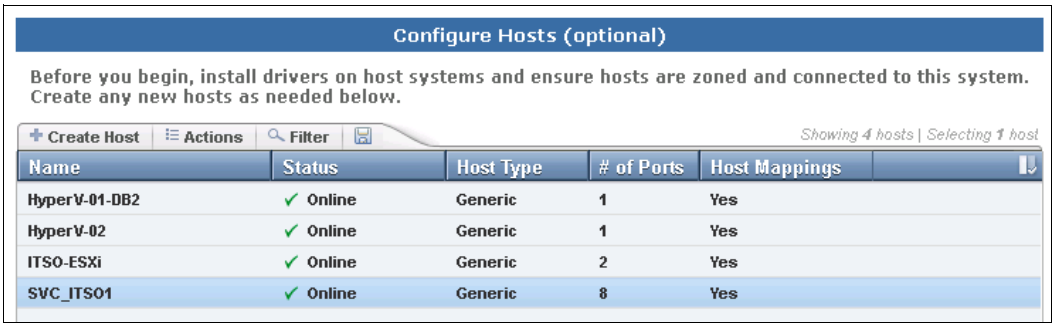

Figure 6-10 Configure Hosts window

Before you configure any hosts, make sure that the appropriate drivers have been installed on the host and that the host zones have been zoned correctly.

If the host that needs to access the data on the volume after the migration is complete is in the list, click **Next**.

If the host has not been created on the IBM Storwize V7000 storage system, click **New Host** to create it as required.

6. Map the newly migrated volume to the host, and when the mapping is complete, click **Next** (Figure 6-11).

Figure 6-11 Map Volumes to Hosts window

In step 6 of the migration wizard, the volumes from the imported MDisks that need to be migrated have been listed. The names of the volumes are assigned automatically by the IBM Storwize V7000 storage system. You can change the names to any words that would be meaningful to you by selecting the volume and clicking **Rename** in the **Actions** drop-down menu.

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

To map the volumes to the hosts, select the volumes and click **Map to Host**. A window opens with a drop-down list of the hosts. At this time, the IBM Storwize V7000 lets you choose which host to which you need to map the volumes. Choose the correct host and click **Next**.

Volume mapping: As a preferred practice, map the volume to the host with the same SCSI ID before the migration, which you should have recorded in step 3.

Figure 6-12 shows the menu for choosing the host to which you map the volumes.

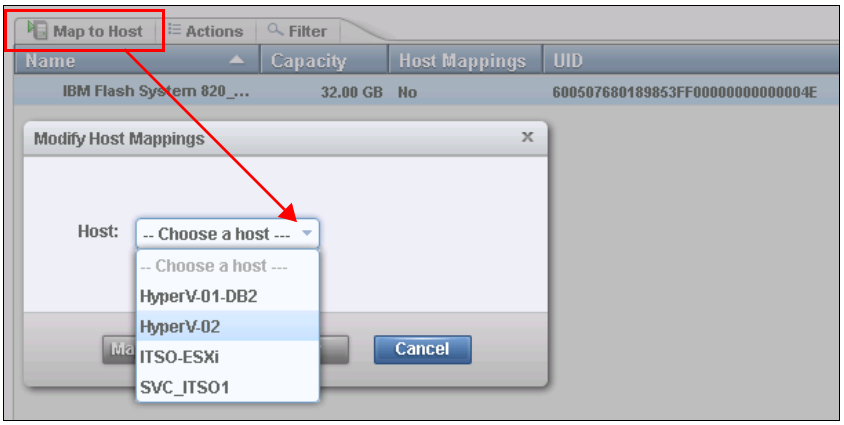

Figure 6-12 Choose the host to which you map the volumes

After choosing the host, the Modify Mappings window opens. On the right, you can find your newly mapped volumes highlighted in yellow. You can change the SCSI ID of the new mappings. Click **OK** to complete the mapping.

Figure 6-13 shows the Modify Mappings window in the migration wizard.

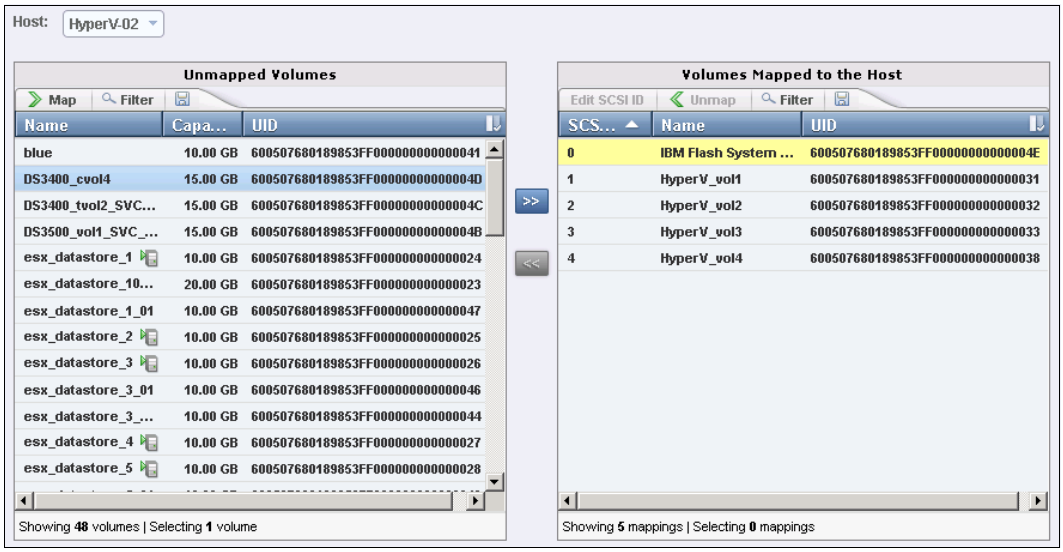

Figure 6-13 Modify Mappings window in the migration wizard

When the mapping completes, you see that the Host Mappings column of the volumes changed from No to Yes. A scan can be performed to discover the new devices on the host for verification. Click **Next** to go to the next step of the storage migration wizard.

7. Select the destination storage pool for data migration and click **Next** (Figure 6-14).

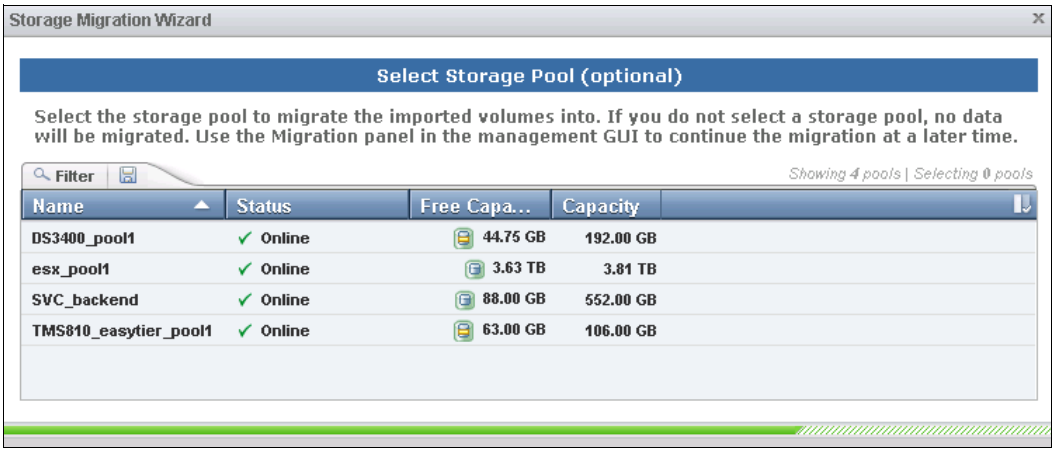

Figure 6-14 Select a storage pool

The destination storage pool of the data migration might be an external storage pool or an internal storage pool. Make sure that there is enough space in the storage pools.

When you click **Next**, the migration begins (Figure 6-15). The migration runs in the background and results in a copy of the data being placed on the MDisks in the storage pool selected. The process uses the volume mirroring function included with the IBM Storwize V7000, and when complete, the volumes have pointers to both the new copy on the storage pool selected and on the original external storage system.

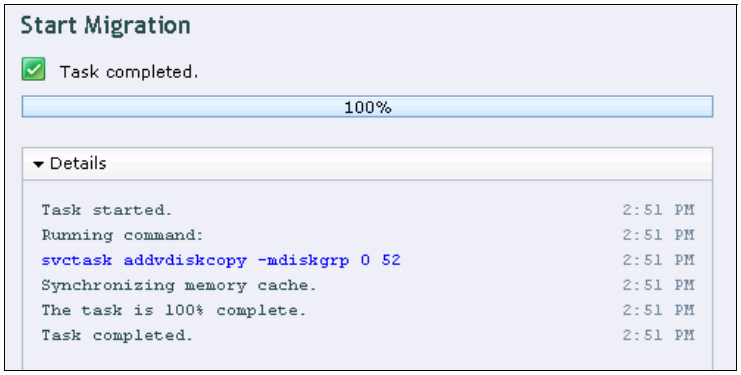

Figure 6-15 Starting migration

8. Click **Finish** to end the storage migration wizard (Figure 6-16).

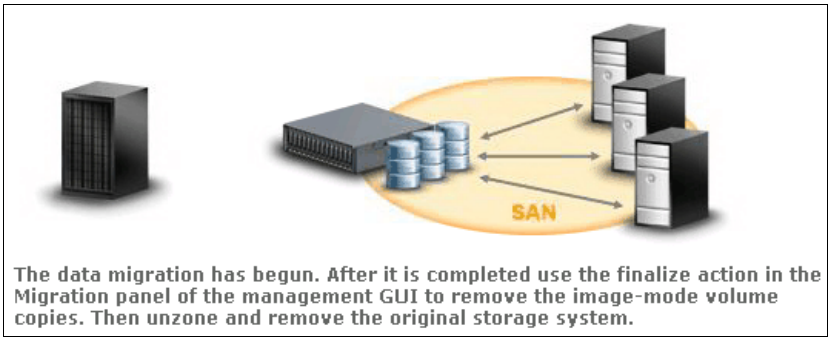

Figure 6-16 End of the storage migration wizard

The end of the storage migration wizard is not the end of the data migration. The data migration has begun, and after clicking **Finish** in step 8, you can find the migration progress in the migration window. You can also find the target storage pool to which your volumes are being migrated, along with the status of the volumes.

Figure 6-17 shows the data migration progress in the migration window.

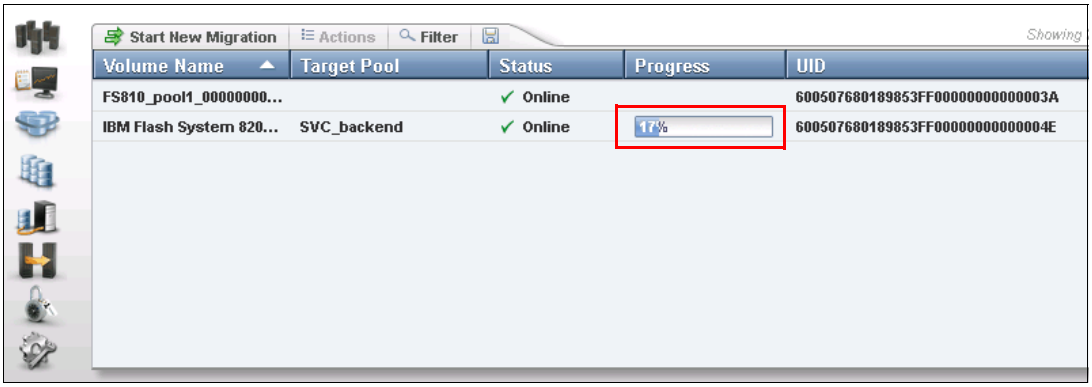

Figure 6-17 Migration progress in the migration window

When the migration progress reaches 100%, select the volumes and click **Finalize** in the **Actions** drop-down menu in the migration window (Figure 6-18). The image mode copy of the volumes on the original external storage system is deleted and the associated MDisks from the storage pool are removed; the status of those MDisks is unmanaged.

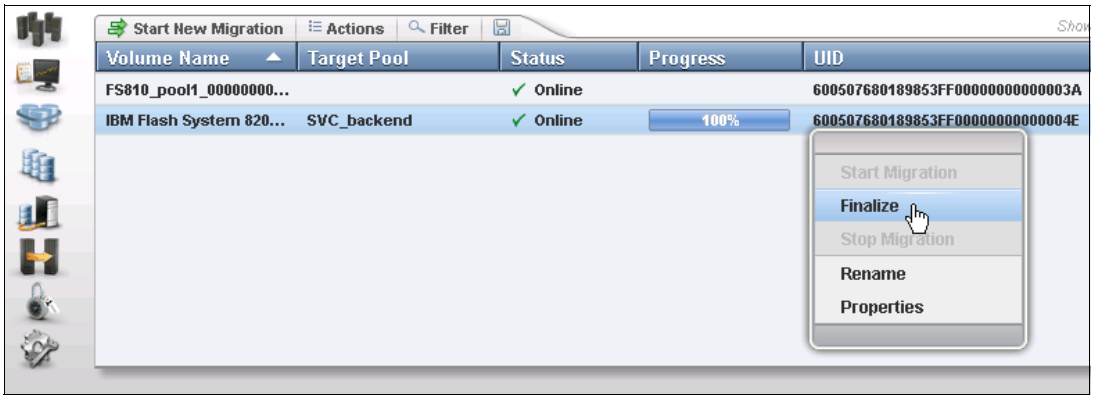

Figure 6-18 Finalize the migration

The IBM Storwize V7000 requests that you confirm the migration of the volumes. Verify the volume name and the number of migrations and click **OK**. Figure 6-19 shows the Finalize Volume Migration window.

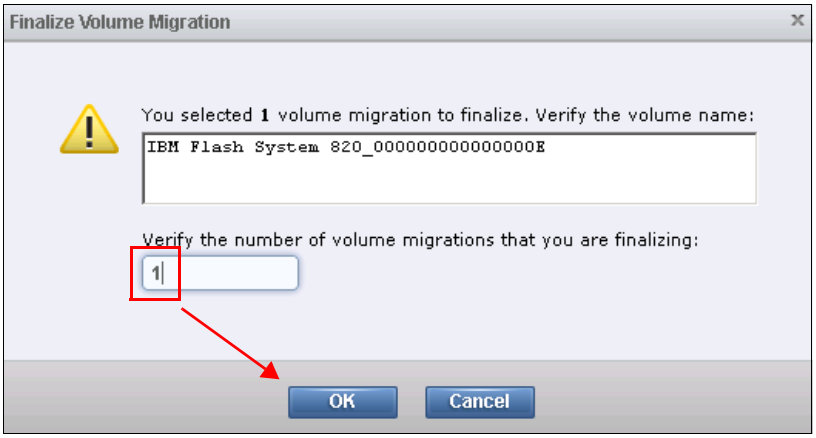

Figure 6-19 Verify the migration finalization

When the finalization completes, the data migration to the IBM Storwize V7000 is done. You can un-zone and remove the older storage system from the IBM Storwize V7000.

For more information about the advanced migration function, see Chapter 7, “Storage pools” on page 179 and Chapter 9, “External storage virtualization” on page 297.

7

Storage pools

This chapter explains how IBM Storwize V7000 manages physical storage resources. All storage resources under IBM Storwize V7000 control are managed using *storage pools*. Storage pools aggregate disk capacity on a storage area network (SAN) and provide the containers in which volumes can be created. Storage pools make it easy to dynamically allocate resources, maximize productivity, and reduce costs. Advanced internal storage, MDisks, and storage pool management are covered in this chapter; external storage is covered in Chapter 9, “External storage virtualization” on page 297.

7.1 Working with internal drives

This section describes how to configure the internal storage disk drives using different RAID levels and different optimization strategies.

In this section, we start with the environment shown in Figure 7-1. The majority of the internal drives are unused. Currently, the existing MDisks come from external storage, and sample storage pools, volumes, and hosts have been created for use.

You can learn how to manage MDisks in 7.2, “Working with MDisks” on page 197, how to manage storage pools in 7.3, “Working with storage pools” on page 217, how to work with external storage in Chapter 9, “External storage virtualization” on page 297, how to create volumes in Chapter 5, “Basic volume configuration” on page 131, and how to create hosts in Chapter 4, “Host configuration” on page 105.

The IBM Storwize V7000 storage system provides an individual Internal Storage panel for managing all internal drives. To access the Internal Storage panel, open the IBM Storwize V7000 Overview panel, click the **Internal Drives** function icon, and then click **Pools**. Figure 7-1 shows how to do this.

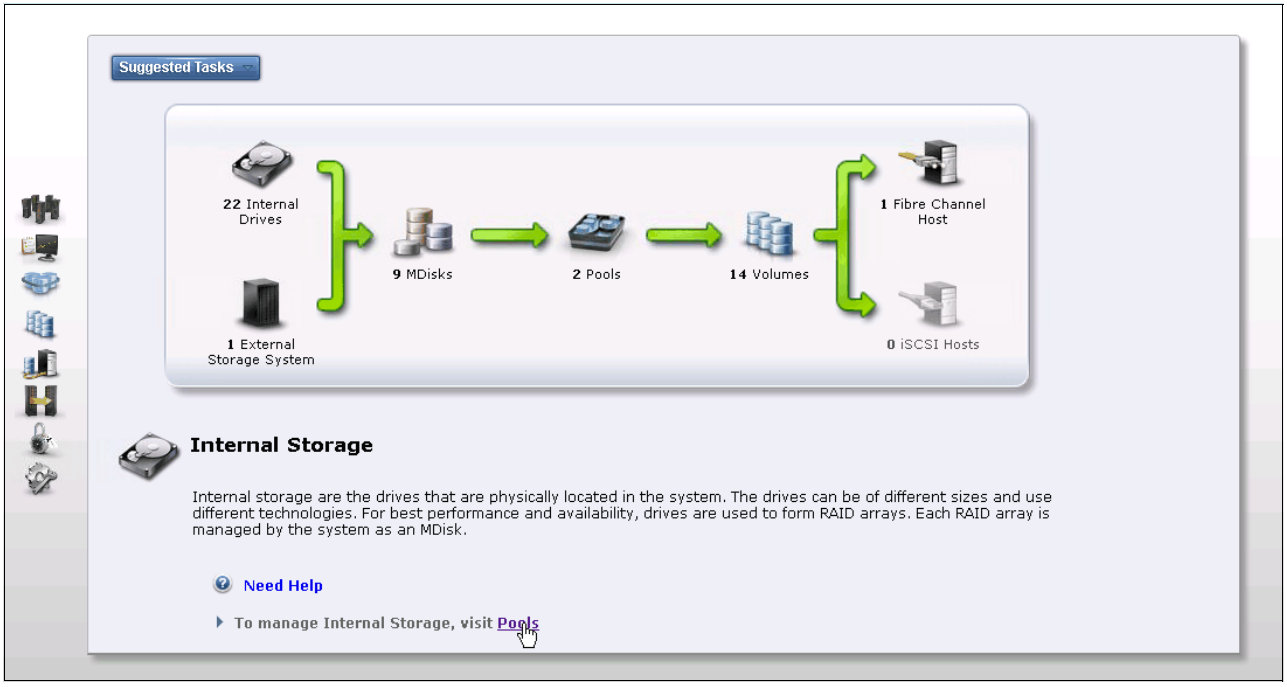

Figure 7-1 Access the Internal Storage panel from the IBM Storwize V7000 Overview

Another way to access the Internal Storage panel is to hover the mouse cursor over the **Pools** selection in the dynamic menu and select **Internal Storage**. (Figure 7-2).

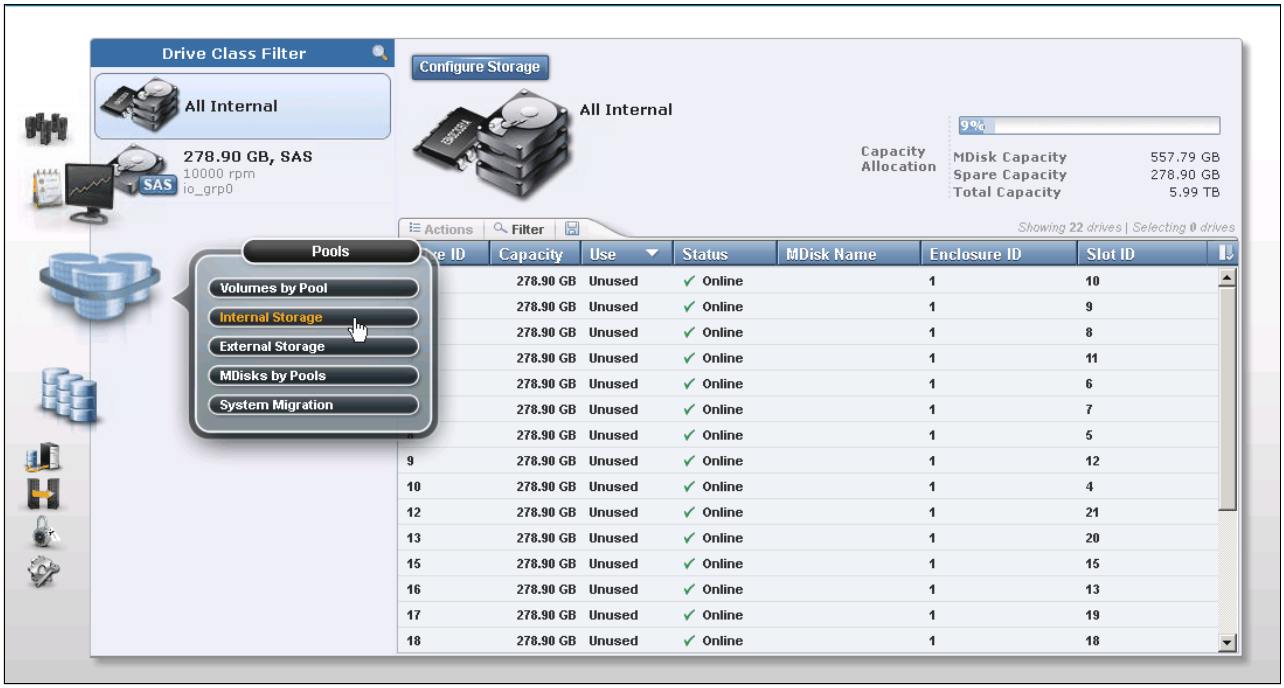

Figure 7-2 Access the Internal Storage panel using the Pools selection

The Internal Storage panel (Figure 7-3) gives an overview of the internal drives installed in the IBM Storwize V7000 storage system. Selecting **All Internal** in the Drive Class Filter shows all drives installed in the managed system, including attached expansion enclosures. Alternatively, you can filter the drives by their type or class, for example, you can choose to show only SAS, SATA or SSD drives.

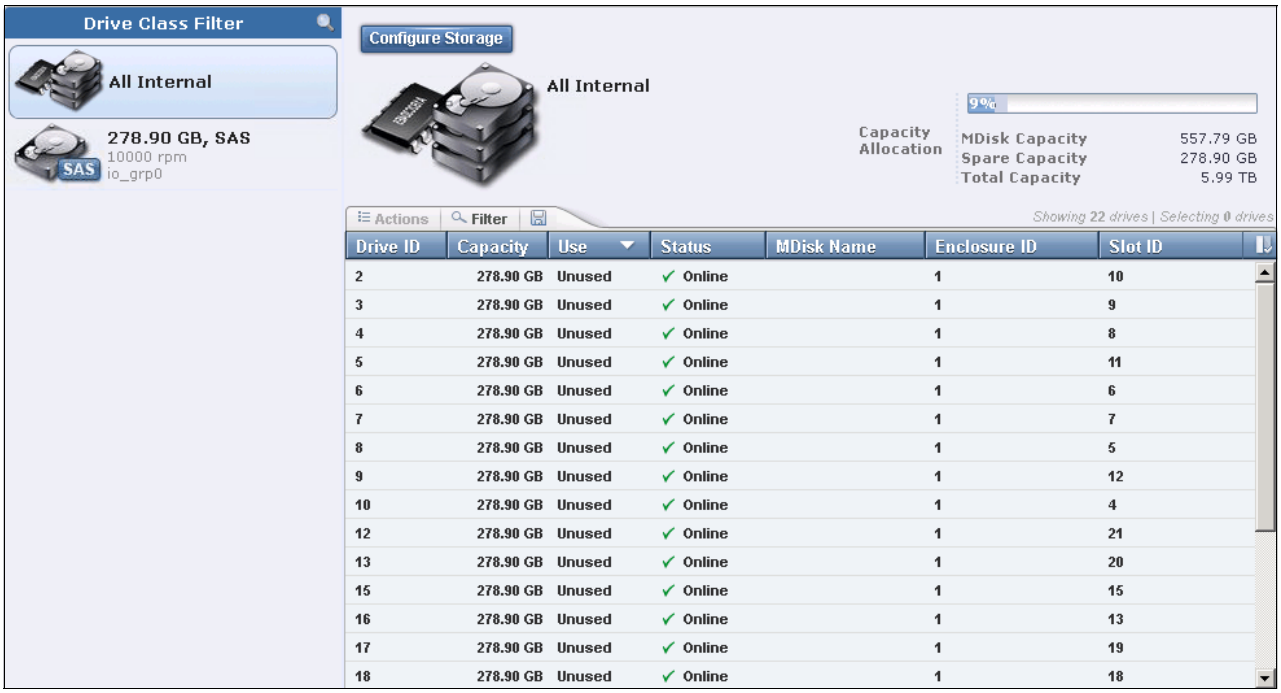

Figure 7-3 Internal Storage panel

The right side of the Internal Storage panel lists the internal disk drives of the selected type. By default, the following information is listed with it:

- Logical drive ID
- Drive's capacity
- Current type of use (unused, candidate, member, spare, or failed)
- Status (online, offline, and degraded)
- MDisk's name that the drive is a member of
- Enclosure ID that it is installed in
- Physical Drive Slot ID of the enclosure that it is installed in

More details, for example, the drive's RPM speed or its MDisk member ID, are available by clicking the blue header bar of the table (Figure 7-4). Also, if you click the table headers, the table content can be ascending or descending, sorted by the value represented in that column.

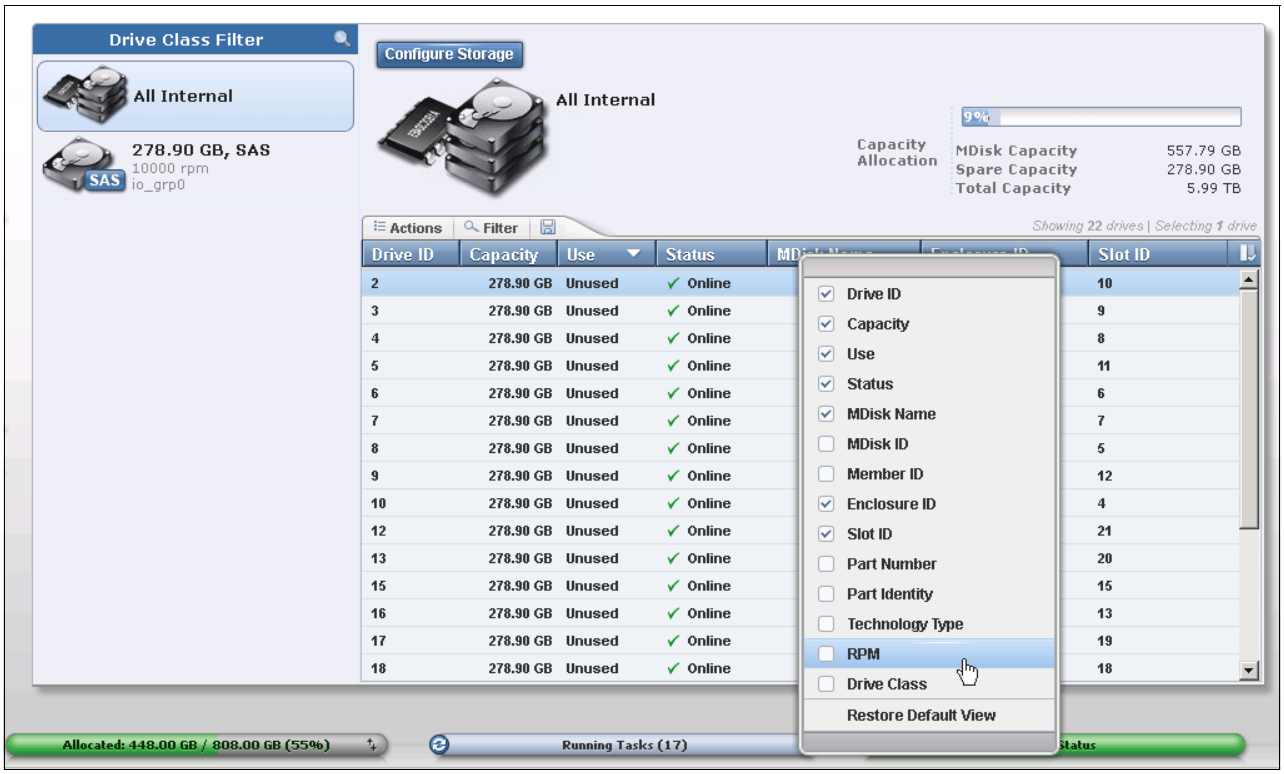

Figure 7-4 Internal Storage panel details selection

In addition, you can find the current internal storage capacity allocation indicator at the upper right. The Total Capacity shows the overall capacity of the internal storage that is installed in this IBM Storwize V7000 storage system. The MDisk Capacity shows the internal storage capacity that is assigned to the MDisks. The Spare Capacity shows total capacity of all drives that were designated as spare drives.

The percentage bar indicates the percentage of the total capacity that is allocated to the MDisk capacity, with the grayed area showing the capacity of the spare disks (Figure 7-5).

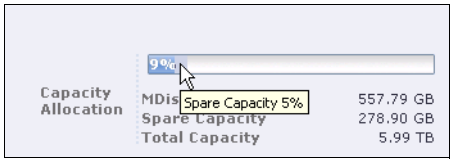

Figure 7-5 The percentage bar for Internal Storage Capacity

7.1.1 Actions on internal drives

There are a few actions can be taken on internal drives when you select the drive and right-click it, or click the **Actions** drop-down menu (Figure 7-6).

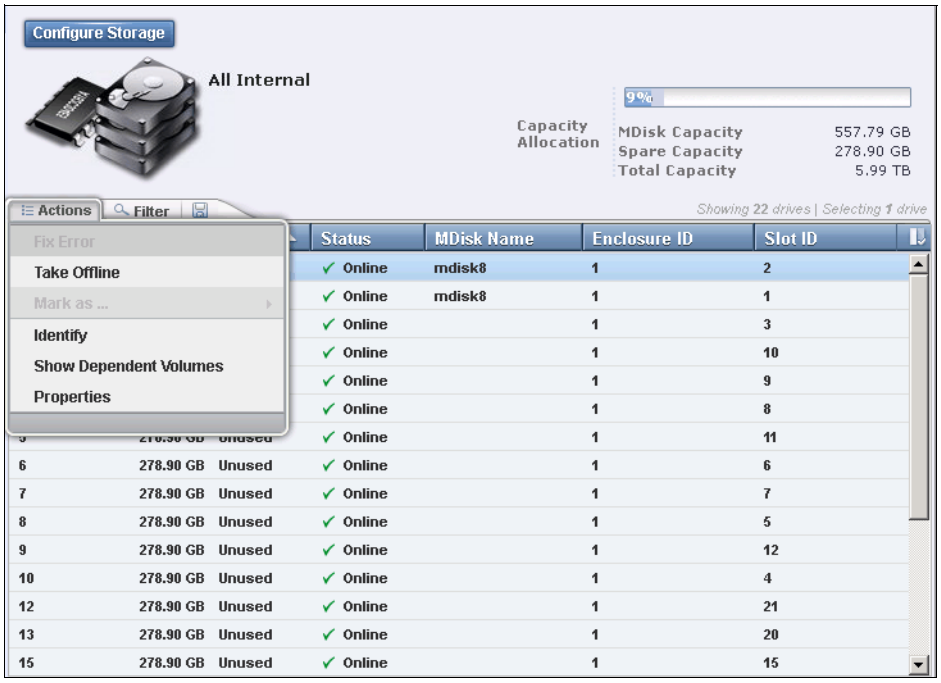

Figure 7-6 Take actions on internal drives

Options on the Actions menu are as follows:

- Fix Error: If the physical drive is in an error condition, the Fix Error action starts the Directed Maintenance Procedure (DMP) for a defective drive. For more information about this topic, see Chapter 13, “RAS, monitoring, and troubleshooting” on page 455.
- Take Offline: When problems occur on the drives, the internal drives can be taken offline by selecting this option in the Actions drop-down menu. A confirmation window opens (Figure 7-7 on page 184). The IBM Storwize V7000 storage system prevents the drive from being taken offline if the result might produce data loss. A drive should be taken offline only if a spare drive is available.

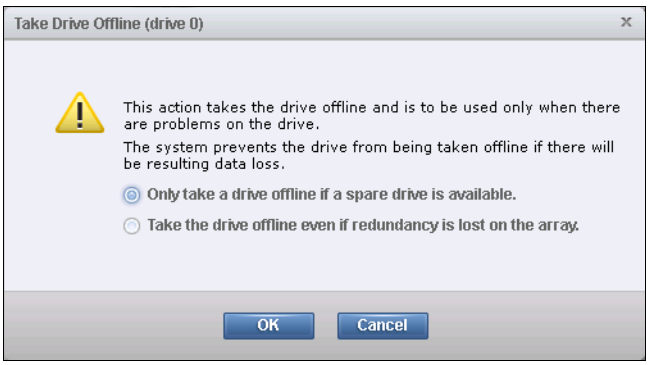

Figure 7-7 Take Drive Offline warning window

Internal drives: Choosing the option to take internal drives offline, even ifredundancy is lost on the array, can lead to potential data loss.

- Mark as: The internal drives in the IBM Storwize V7000 storage system can be assigned to several usage roles. These roles are as follows:
  - Unused: The drive is not a member of an MDisk (array).
  - Candidate: The drive is available for use in an array.
  - Spare: The drive can be used as a hot spare if required.
  - Failed: The drive was either intentionally taken offline or failed because of an error.

  Select **Mark as** from the Actions drop-down menu, and select the role you want to assign to the drive (Figure 7-8).

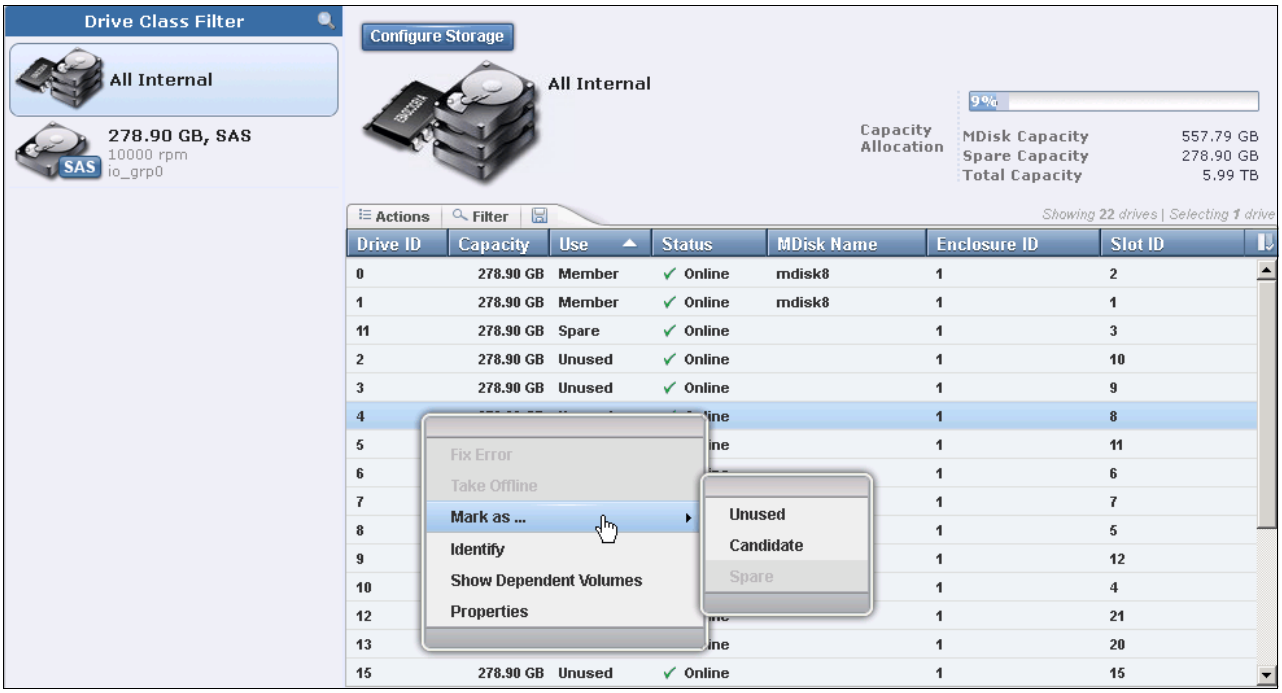

Figure 7-8 Change a drive's usage

- Identify: Use the Identify action to turn on the LED light so you can easily identify a drive that needs to be replaced, or that you want to troubleshoot.

  Figure 7-9 shows the information that is displayed when you click the **Identify** action.

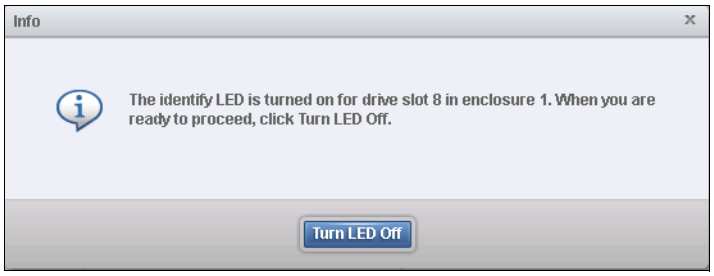

Figure 7-9 Drive identification LED feature

  Click **Turn LED Off** when you are done.

- Show Dependent Volumes: Shows you the volumes that are dependent on the drives (Figure 7-10).

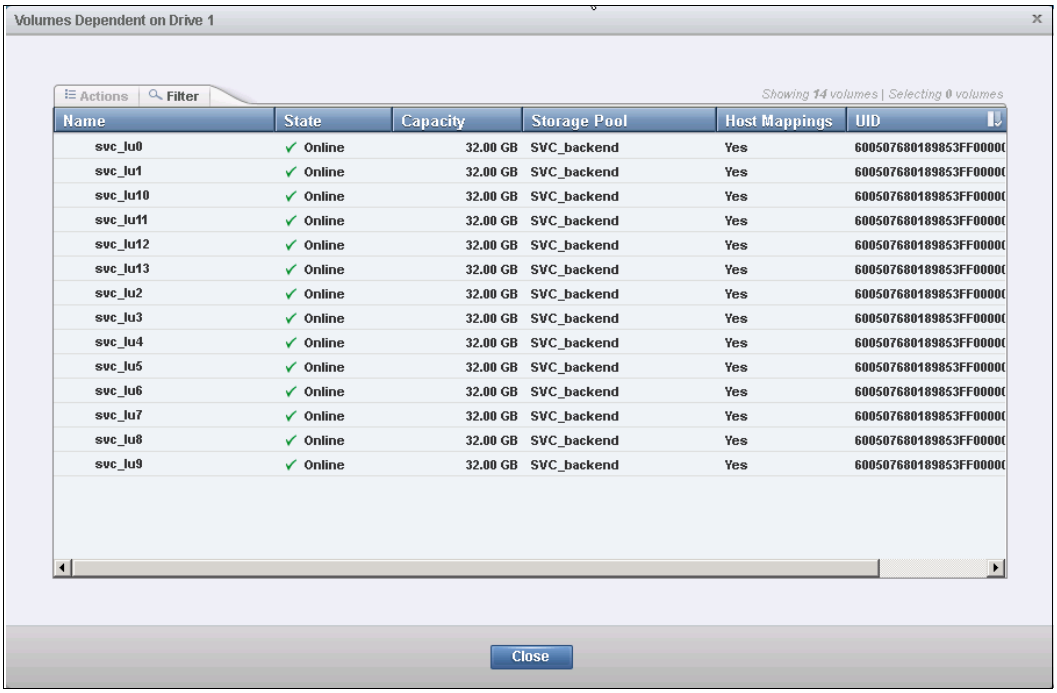

Figure 7-10 Drive's dependent volumes

- Properties: Clicking **Properties** in the Actions drop-down menu, or double-clicking the drive, provides more information about the drives (Figure 7-11).

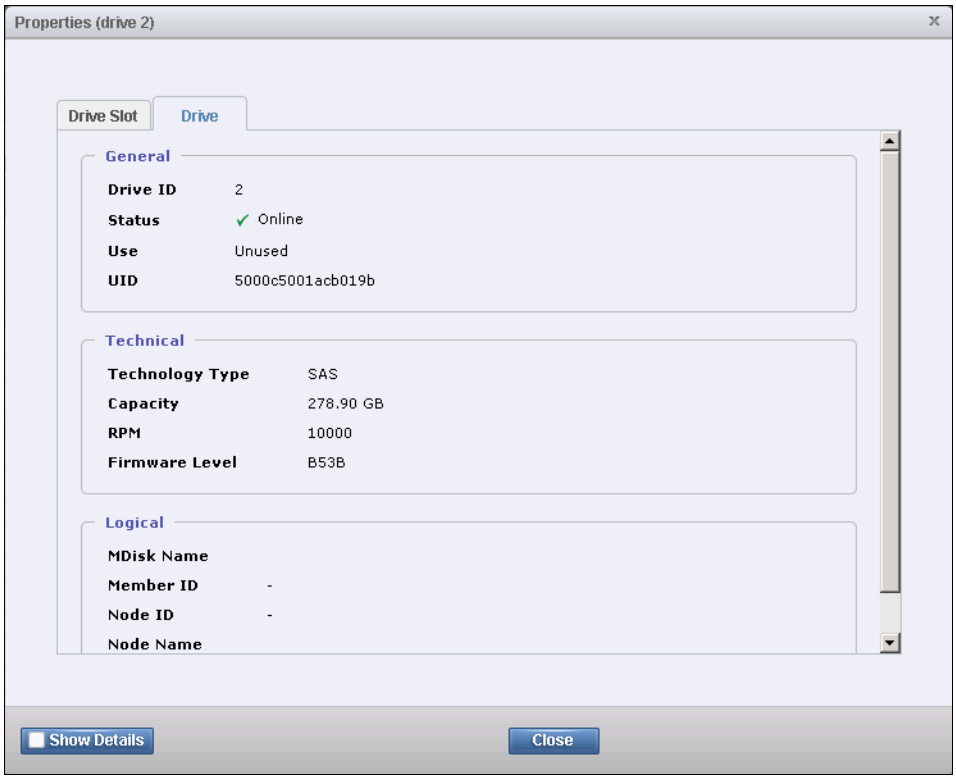

Figure 7-11 Properties tab with default format

If you select the **Show Details** check box, you can discover more detailed information, including vendor ID, FRU Part Number, and ID (Figure 7-12).

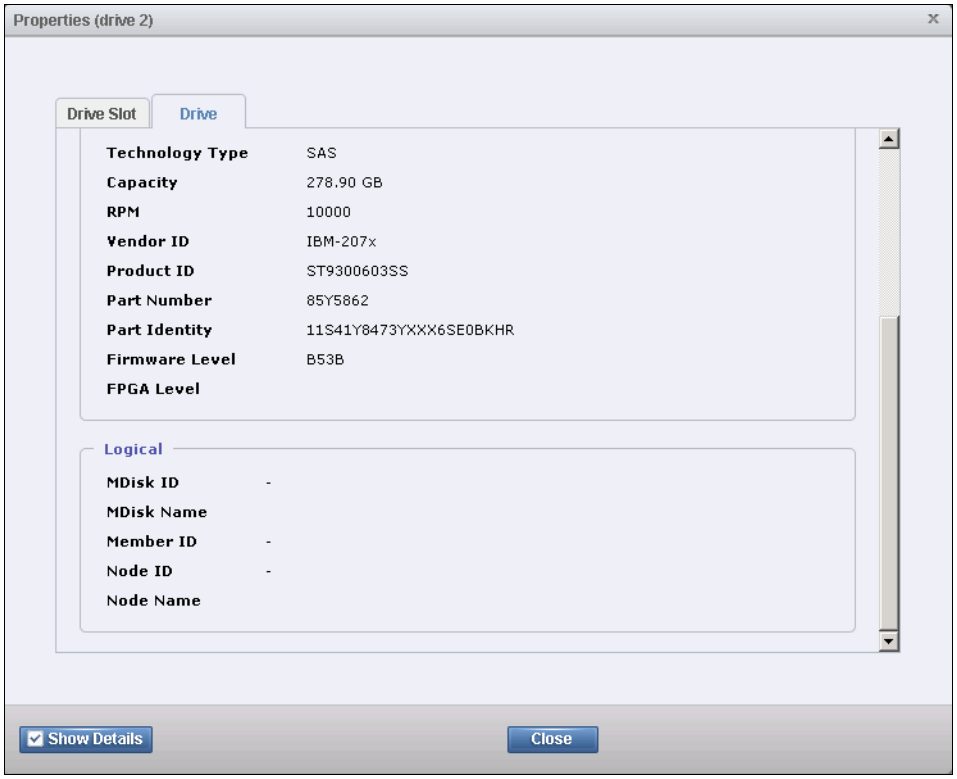

Figure 7-12 Properties tab with details

7.1.2 Configuring internal storage

To configure internal storage for use with hosts, click move the mouse cursor over the **Pools** selection and click **Internal Storage**, and then click the **Configure Storage** button (Figure 7-4 on page 182).

A configuration wizard opens and guides you through the process of configuring internal storage (Figure 7-13).

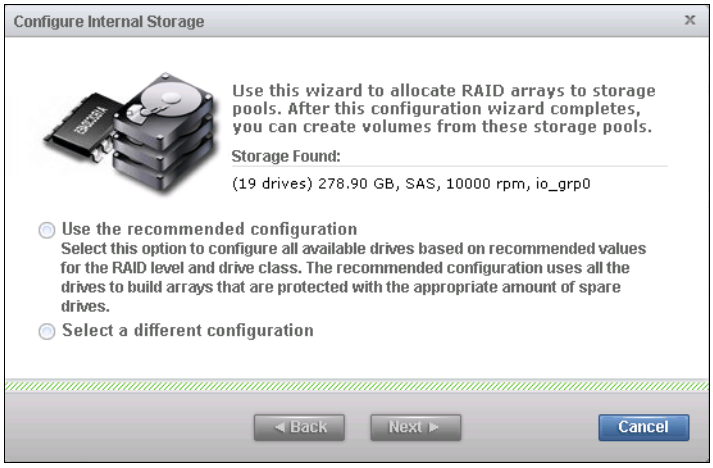

Figure 7-13 Wizard to Configure Internal Storage

The wizard shows all internal drives with a status of candidate available for configuration. If there are internal drives with a status of unused, a window opens, offering to include them in the RAID configuration.

The storage configuration wizard simplifies the initial disk drive setup and offers two options:

- Use the recommended configuration

  This option guides you through the wizard described in “Using the recommended configuration” on page 190.
- Select a different configuration

  This option uses the wizard described in “Selecting a different configuration” on page 192.

Before going through the storage configuration wizard, we first describe IBM Storwize V7000 RAID configuration presets.

RAID configuration presets

RAID configuration presets are used to configure internal drives based on recommended values for the RAID level and drive class. Each preset has a specific goal for the number of drives per array, the number of spare drives to maintain redundancy, and whether the drives in the array are balanced across enclosure chains, thus protecting the array from enclosure failures. For details about enclosure chains, see 1.3, “IBM Storwize V7000 terminology” on page 4.

Table 7-1 describes the presets that are used for solid-state drives (SSDs) for the IBM Storwize V7000 storage system.

Table 7-1 SSD RAID presets

Preset	Purpose	RAID level	Drives per array goal	Spare drive goal
SSD RAID 5	Protects against a single drive failure. Data and one stripe of parity are striped across all array members.	5	8	1
SSD RAID 6	Protects against two drive failures. Data and two stripes of parity are striped across all array members.	6	12	1
SSD RAID 10	Protects against at least one drive failure. All data is mirrored on two array members.	10	8	1
SSD RAID 0	Provides no protection against drive failures.	0	8	0
SSD Easy Tier	Mirrors data to protect against drive failure. The mirrored pairs are spread between storage pools to be used for the Easy Tier function.	10	2	1

SSD RAID instances: In all SSD RAID instances, drives in the array are balanced across enclosure chains if possible.

Table 7-2 describes the RAID presets that are used for hard disk drives for the IBM Storwize V7000 storage system.

Table 7-2 HDD RAID presets

Preset	Purpose	RAID level	Drives per array goal	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	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	1	All drives in the array are from the same chain wherever possible.
Basic RAID 10	Protects against at least one drive failure. All data is mirrored on two array members.	10	8	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	1	Exactly half of the drives are from each chain.
RAID 0	Provides no protection against drive failures.	0	8	0	All drives in the array are from the same chain wherever possible.

Using the recommended configuration

As shown in Figure 7-14, when you click **Use the recommended configuration**, the wizard offers a recommended storage configuration at the bottom of the window.

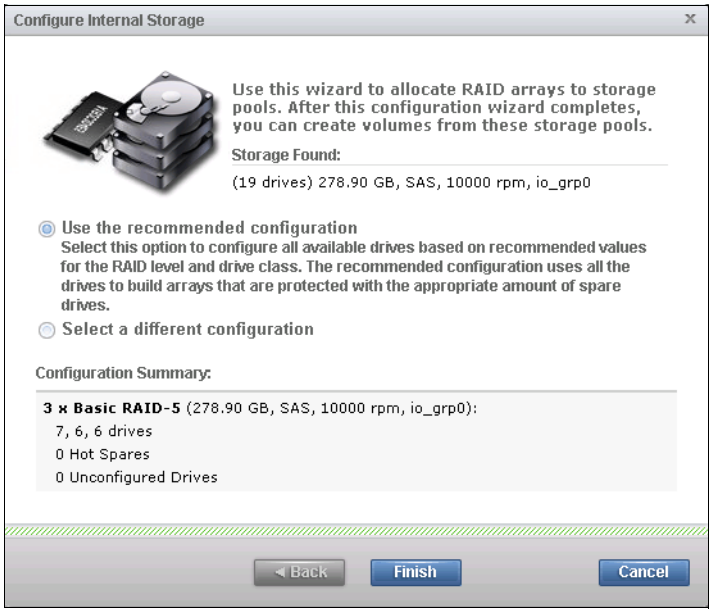

Figure 7-14 Recommended configuration

The recommended RAID presets for different drive classes are as follows:

- SSD Easy Tier preset for solid-state drives
- Basic RAID 5 for SAS drives
- Basic RAID 6 for Nearline SAS drives

Using the recommended configuration, spare drives are also automatically created to meet the spare goals according to the preset chosen; one spare drive is created out of every 24 disk drives of the same drive class on a single chain.

For example, if you have 20 x 450 GB 10 K SAS drives on one chain, one drive in these 20 drives is dedicated as a spare drive. If you have 20 x 450 GB 10 K SAS drives on both chains, which means that there are 10 drives in each chain, then one spare drive on each chain is created. So, if you have 40 x 450 GB 10k SAS drives on both chains, then two spare drives on each chain are created and you have a total of 36 drives that can be array members for the RAID setup.

Spare drives in the IBM Storwize V7000 are *global spares*, which means that any spare drive having at least the same capacity as the drive that needs to be replaced can be used in any array. Thus, an SSD array with no SSD spare available will use an HDD spare instead.

In our example of using the recommended configuration, three arrays that use the Basic RAID 5 preset are proposed with no hot spare, because the hot spare is already assigned (see Figure 7-8 on page 184). If the proposed configuration meets your requirements, click **Finish**, and the system automatically creates the array MDisks with a size according to the chosen RAID level.

Storage pools are also automatically created to contain the MDisks with similar performance characteristics, including the consideration of RAID level, number of member drives, drive class, and so on.

After an array is created, the Array MDisk members are synchronized with each other through a background initialization process. To monitor the progress of the initialization process, click the icon at the left of the **Running Tasks** status bar and then click the displayed task (Figure 7-15).

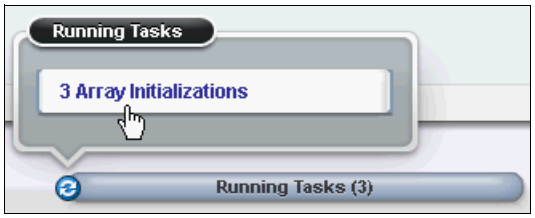

Figure 7-15 Running Tasks status bar

The progress window opens (Figure 7-16). The array is available for I/O during this process. The initialization does not affect the availability because of possible member drive failures.

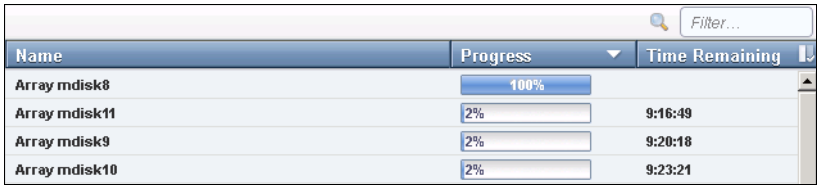

Figure 7-16 Array initialization

The capacity allocation indicator shows that allocation capacity has reached 95% after the configuration is set (Figure 7-17).

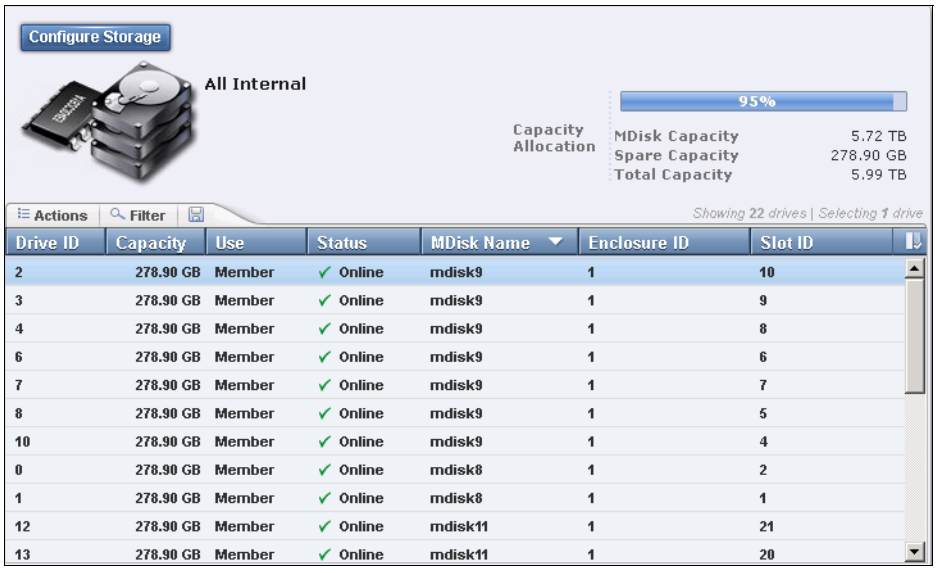

Figure 7-17 Capacity allocation after applying the recommended configuration

If the proposed configuration does not meet your requirements, click **Select a different configuration** on the initial window of the Configure Storage wizard and continue with the more flexible setup shown in “Selecting a different configuration” on page 192.

Selecting a different configuration

The option **Select a different configuration** offers a more flexible way for the configuration of the internal storage as compared to the **Use the recommended configuration** preset in terms of drive selection, RAID level, and storage pool to be used.

Complete the following steps:

1. Click **Select a different configuration** and select the **Drive Class** you want to configure (Figure 7-18).

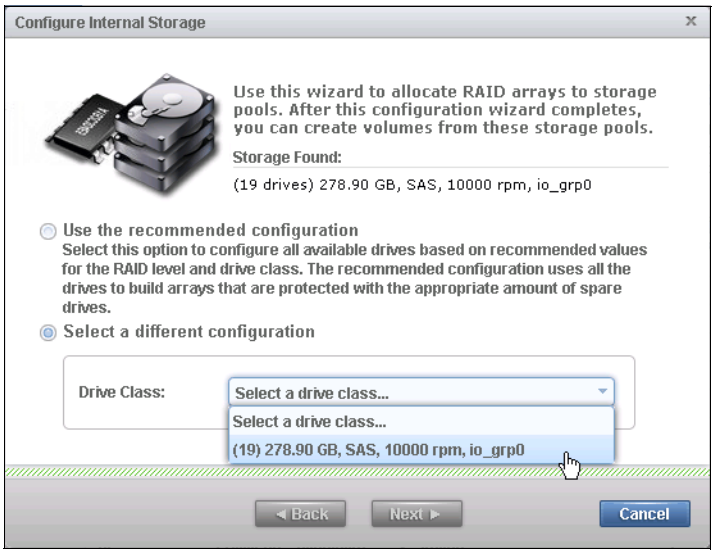

Figure 7-18 Select a drive class

2. Select an appropriate RAID preset (Figure 7-19).

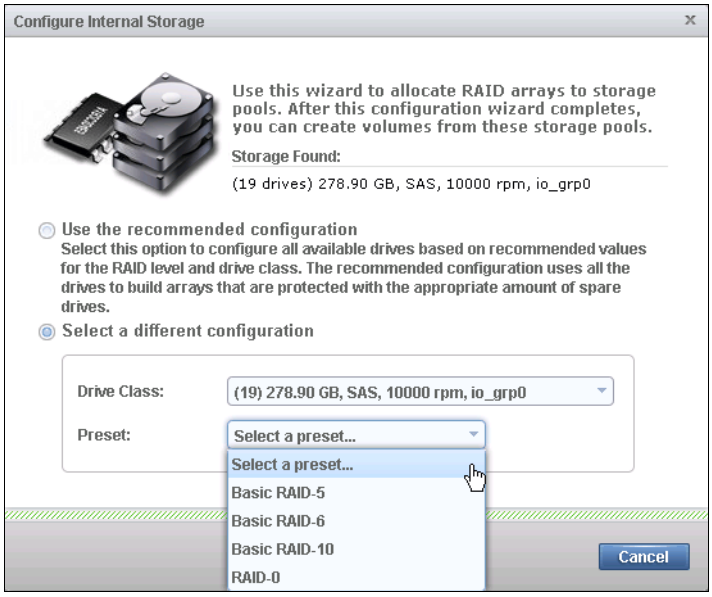

Figure 7-19 Select a RAID preset

3. Select the **Number of drives to provision** in the configuration. By default, all drives with the Candidate status are included, and drives with the Unused status are made available for use in the array also. Also, the wizard offers to automatically configure spares to match

the spare goal, as explained in “RAID configuration presets” on page 188. Disabling this option enables you to define spare drives manually by marking drives with the status Candidate as Spare. Then, decide which of the following two optimization options to use:

– Optimize for Performance (see “Optimize for Performance” on page 193)
– Optimize for Capacity (see “Optimize for Capacity” on page 195)

Optimize for Performance

The goal of this algorithm is to create arrays with identical capacity and performance characteristics to achieve the best possible performance for the volumes to be created. In a performance optimized setup, the IBM Storwize V7000 provisions eight physical disk drives in a single array MDisk, except for the following situations:

- RAID 6 uses 12 disk drives.
- SSD Easy Tier uses two disk drives.

Therefore, creating an Optimized for Performance configuration is possible only if a sufficient number of drives are 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 then in another array.

Figure 7-20 shows an example of the performance optimized setup.

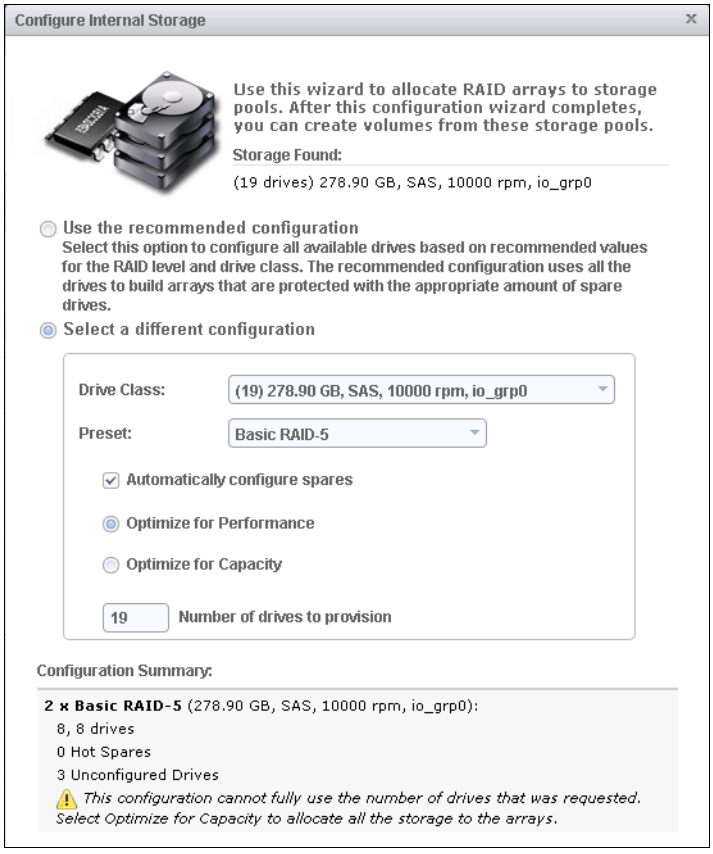

Figure 7-20 Example of a performance optimized setup

Complete the following steps:

1. Choose the storage pool you want to assign the capacity to (Figure 7-21). Select either an existing storage pool that does not contain MDisks, or a pool that contains MDisks with similar performance characteristics, which is listed automatically.

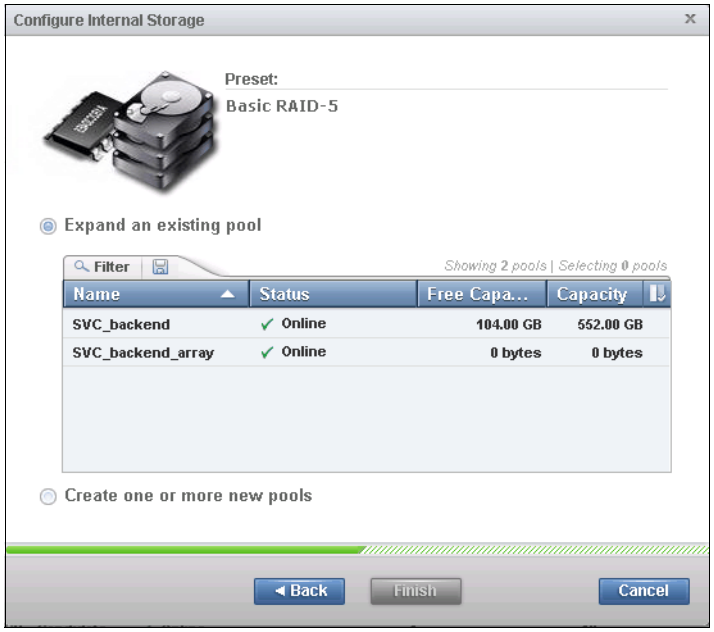

Figure 7-21 Assign capacity to an existing storage pool

2. Alternatively, create a new storage pool (Figure 7-22).

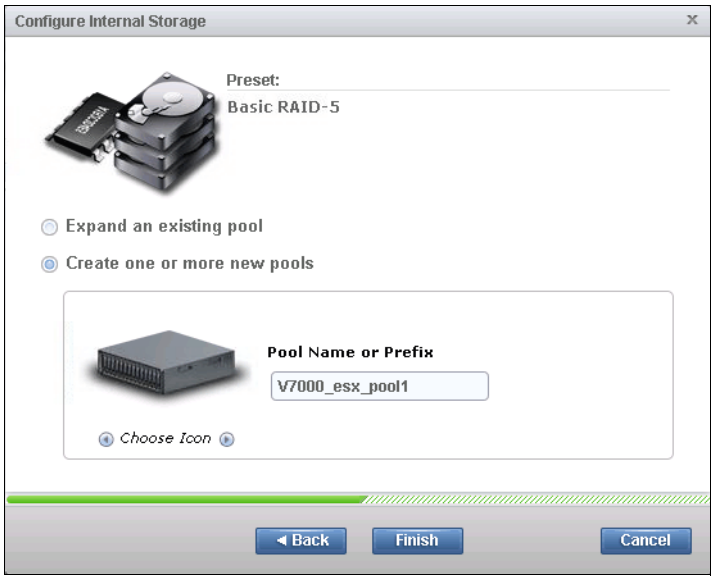

Figure 7-22 Create storage pool

3. Click **Finish** to finalize the wizard. After the wizard completes, the configuration changes are reflected at the upper right of the Internal Storage panel (Figure 7-23).

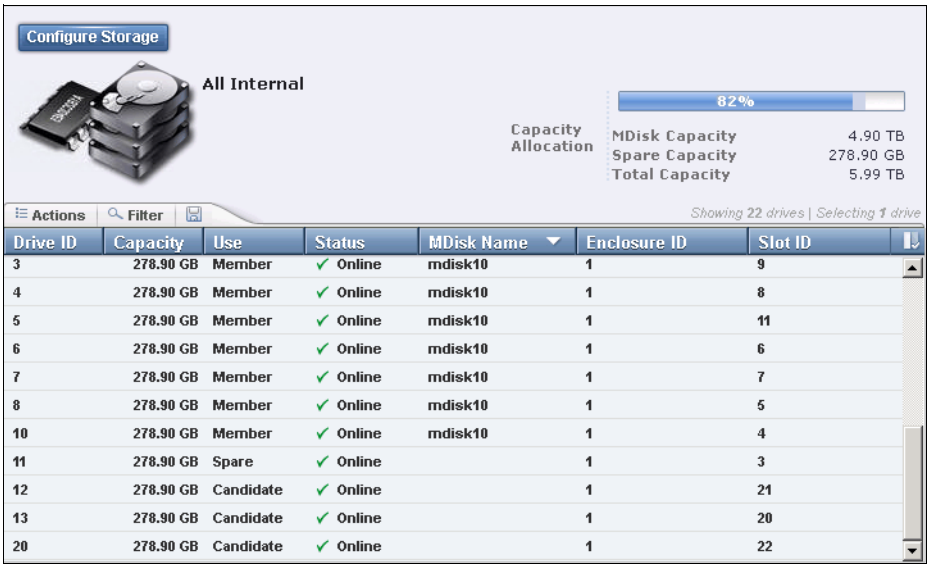

Drive ID	Capacity	Use	Status	MDisk Name	Enclosure ID	Slot ID
3	278.90 GB	Member	Online	mdisk10	1	9
4	278.90 GB	Member	Online	mdisk10	1	8
5	278.90 GB	Member	Online	mdisk10	1	11
6	278.90 GB	Member	Online	mdisk10	1	6
7	278.90 GB	Member	Online	mdisk10	1	7
8	278.90 GB	Member	Online	mdisk10	1	5
10	278.90 GB	Member	Online	mdisk10	1	4
11	278.90 GB	Spare	Online		1	3
12	278.90 GB	Candidate	Online		1	21
13	278.90 GB	Candidate	Online		1	20
20	278.90 GB	Candidate	Online		1	22

Figure 7-23 Configuration complete with performance optimized setup

With the performance optimized setup, two array MDisks are created, each containing eight drives. Three drives remain unconfigured because they did not meet the goal of eight drives for a Basic RAID 5 array to create another *performance optimized* MDisk. The progress of the array initialization process is displayed in the Running Tasks status indicator, and the capacity allocation indicator has been updated to 82% also.

Optimize for Capacity

The goal of this algorithm is to create a setup with the maximum usable capacity, depending on the selected RAID level.

While creating arrays, the IBM Storwize V7000 system attempts to fulfill the *width goal* for each array before creating another one. This setup is an example with 15 *unused* drives **Select a different configuration** was chosen with the following options:

- Preset: Basic RAID-5.
- Automatically configure spares was checked.
- Optimize for capacity.

The results are as follows:

- No spare drive is defined because it is already assigned on the chain (Figure 7-8 on page 184).
- Two basic RAID 5 arrays are created. One of the arrays contains eight drives that are width-goal matched. One array contains the remaining seven drives that are not fully width-goal matched.
- Four unused drives remain.

The *width goals* for the array levels are shown in Table 7-3.

Table 7-3 Array width goals

RAID level	Array width goal
RAID 5, 10, or 0	8 disks
RAID 6	12 disks
SSD Easy Tier	2 disks

If configuring one fully populated IBM Storwize V7000 SFF disk enclosure with 24 disk drives in the Capacity Optimized setup, all available Candidate disk drives are used, no “unconfigured drives” remain, unlike in the Performance Optimized setup. Figure 7-24 shows an example of the capacity optimized setup.

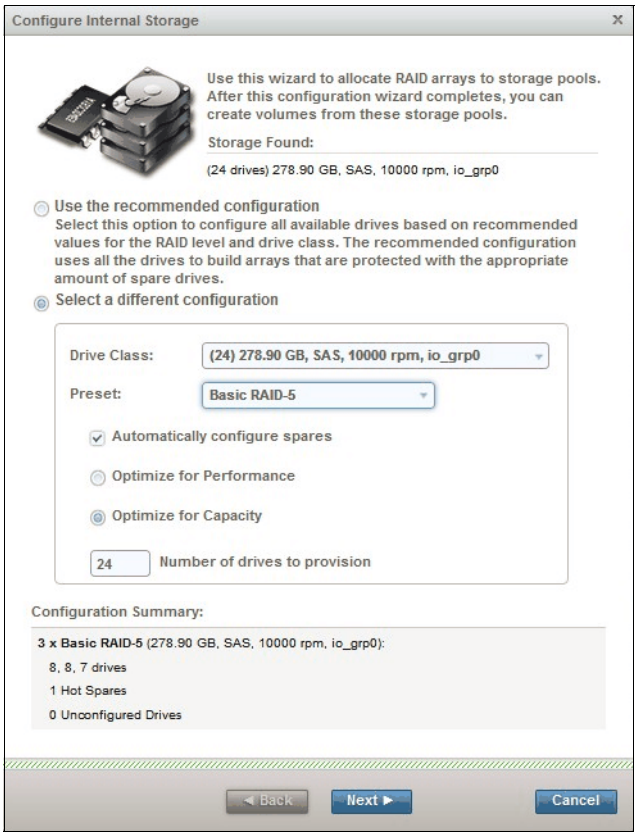

Figure 7-24 Capacity optimized setup

Provisioned drives: You can choose the number of drives to be provisioned. The wizard proposes that you create arrays from all available candidate drives. If you want to configure only a subset of drives now, you can change the number in the **Drives to provision** field to the value that you want. If you want to manually select the drives to provision, assign Candidate status to only the drives that you want. Alternatively, you can use the CLI for full control over the drive configuration, which is beyond the intended scope of this book.

7.2 Working with MDisks

After the configuration is completed for the internal storage, you can find the MDisks that are created on the internal arrays in the MDisks by Pools panel.

To access the MDisks by Pools panel from the IBM Storwize V7000 Overview panel, click the **MDisks** function icon. In the extended help information panel, click **Pools**.

Figure 7-25 shows how to access the Pools panel.

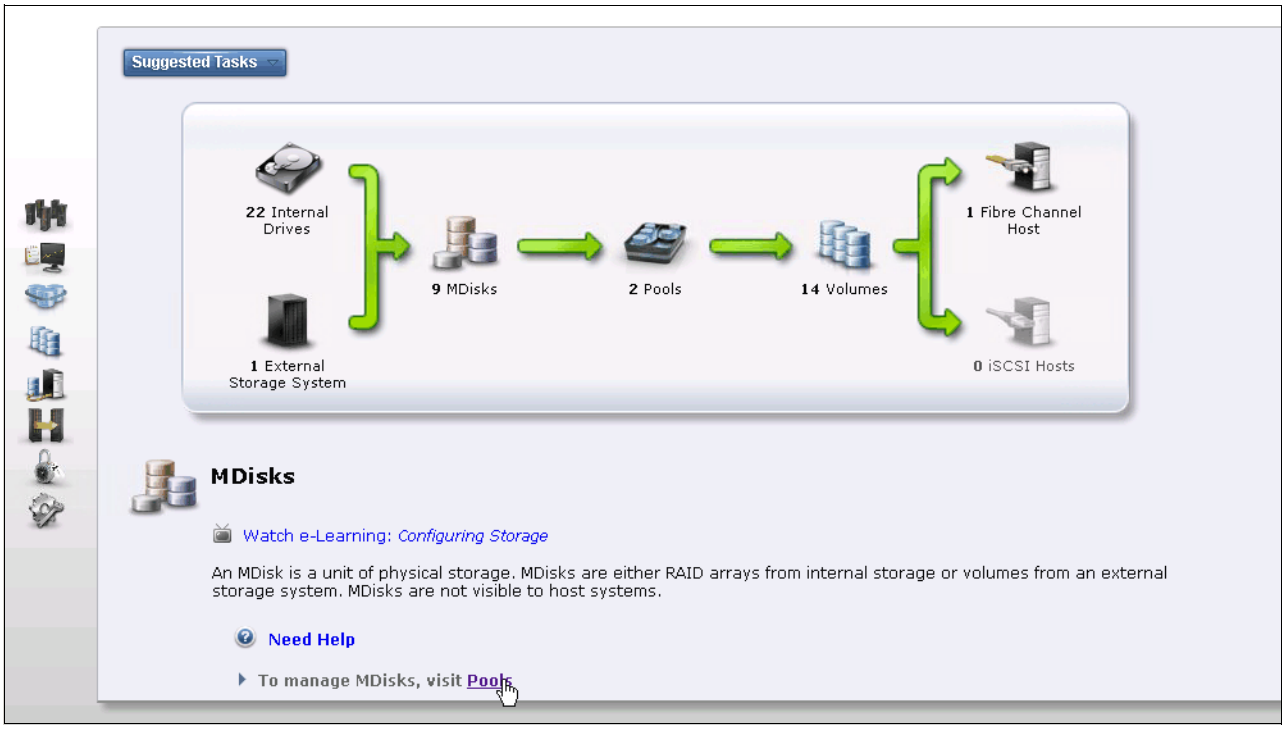

Figure 7-25 Access the Pools management panel from the Overview panel

Another way to access the MDisks by Pools panel is by using the Pools selection in the dynamic menu on the left of the IBM Storwize V7000 GUI (Figure 7-26).

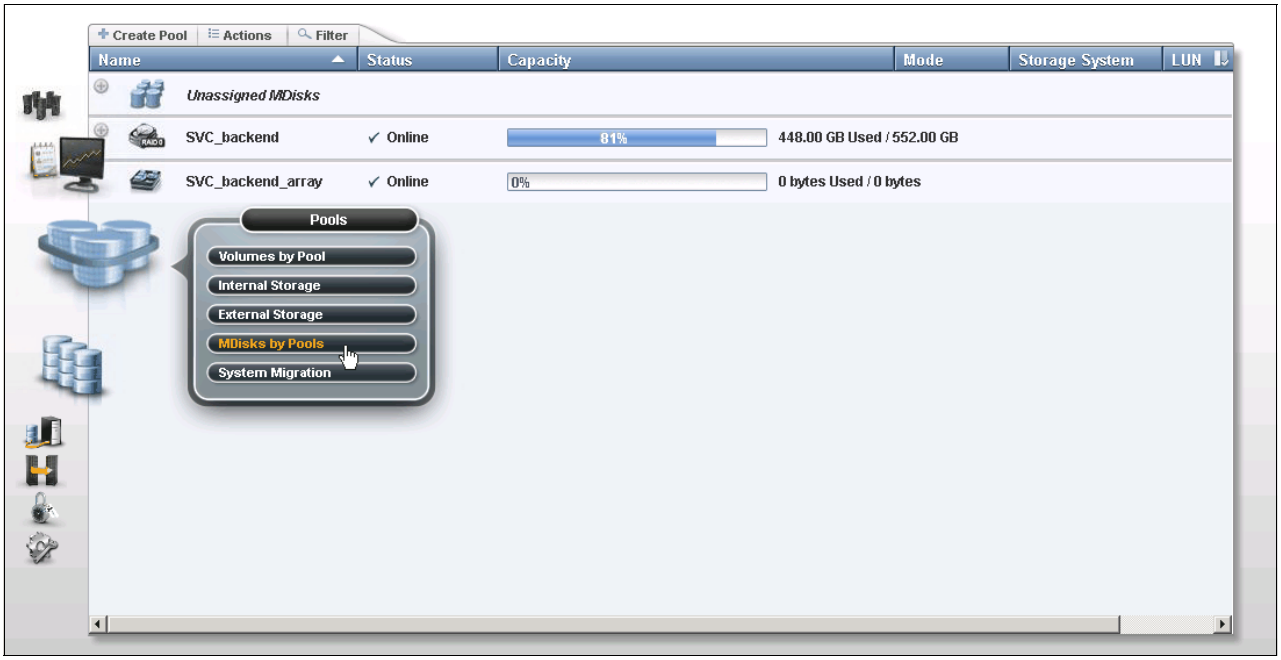

Figure 7-26 Access the MDisks by Pools panel using the Pools selection

The MDisks by Pools panel (Figure 7-27) lets you manage all MDisks made of both internal and external storage.

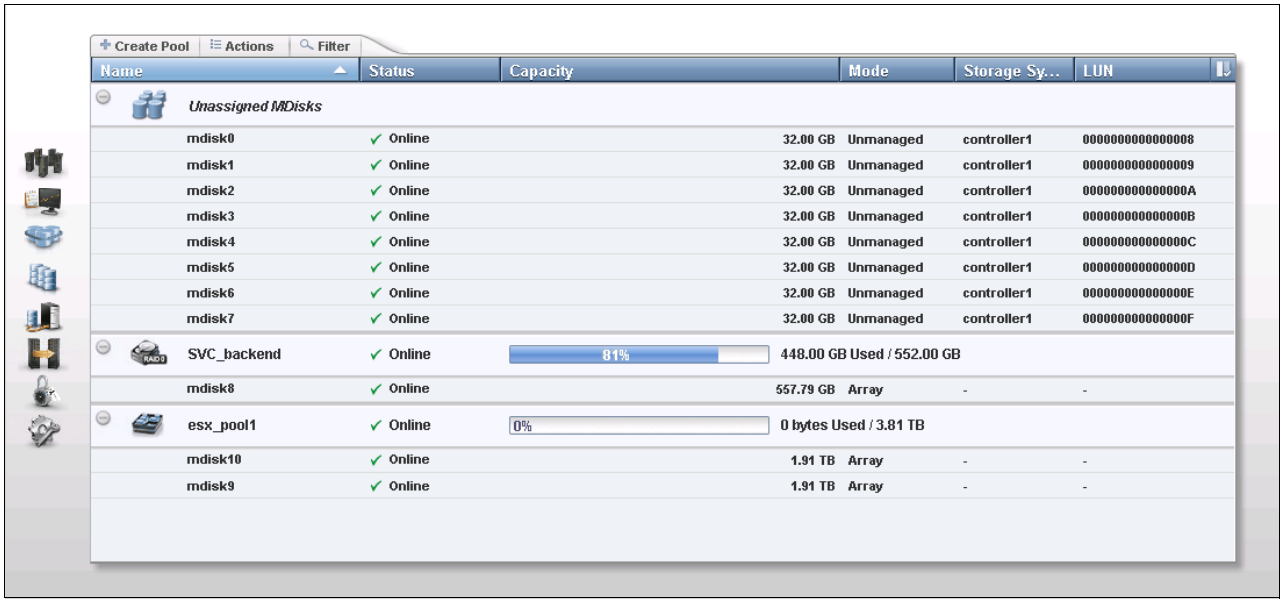

Figure 7-27 MDisks by Pools panel

The panel provides the following information:

- MDisk name
- Numerical ID
- Status
- Capacity
- Mode
- Name of the storage pool it belongs to
- Name of the back-end storage system for MDisk on external storage
- MDisk's LUN ID from external storage systems
- Assigned storage tier

You can find more information about how to attach external storage to an IBM Storwize V7000 storage system in Chapter 9, "External storage virtualization" on page 297.

IBM Storwize V7000 has four modes for MDisks:

- Array: Array mode MDisks are constructed from internal drives using the RAID functionality. Array MDisks are always associated with storage pools (also known as MDisk Groups).
- Unmanaged: The MDisk is not a member of any storage pools, which means it is not being used by the IBM Storwize V7000 storage system. LUNs presented by external storage systems to IBM Storwize V7000 are discovered as unmanaged MDisks.
- Managed: The MDisk is assigned to a storage pool and provides extents to be used by volumes.
- Image: The MDisk is assigned directly to a volume with a one-to-one mapping of extents between the MDisk and the volume.

7.2.1 Adding MDisks to storage pools

By adding *unmanaged MDisks* to a pool, their status changes to *managed MDisks*. Managed MDisks can belong to only one pool. Unmanaged MDisks can either be added to a newly created pool or to an existing pool to expand its capacity. Pools are commonly used to group MDisks from the same storage subsystem.

A new pool can be created in the MDisks by Pools panel by clicking the **Create Pool** icon. Optionally assign a name to the pool and choose an icon (Figure 7-28).

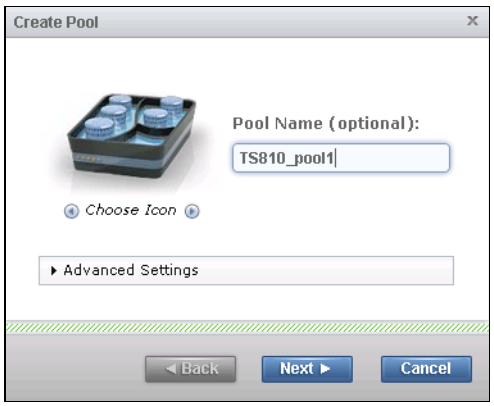

Figure 7-28 Create Pool: part 1 of 2

In the Create Pool window (Figure 7-29), you can include unmanaged MDisks in the new pool. Several filter options are available so you can limit the selection, for example, by storage subsystem, capacity, and so on. To select multiple MDisks press the Ctrl key while using the mouse to click the MDisks that you want from the list.

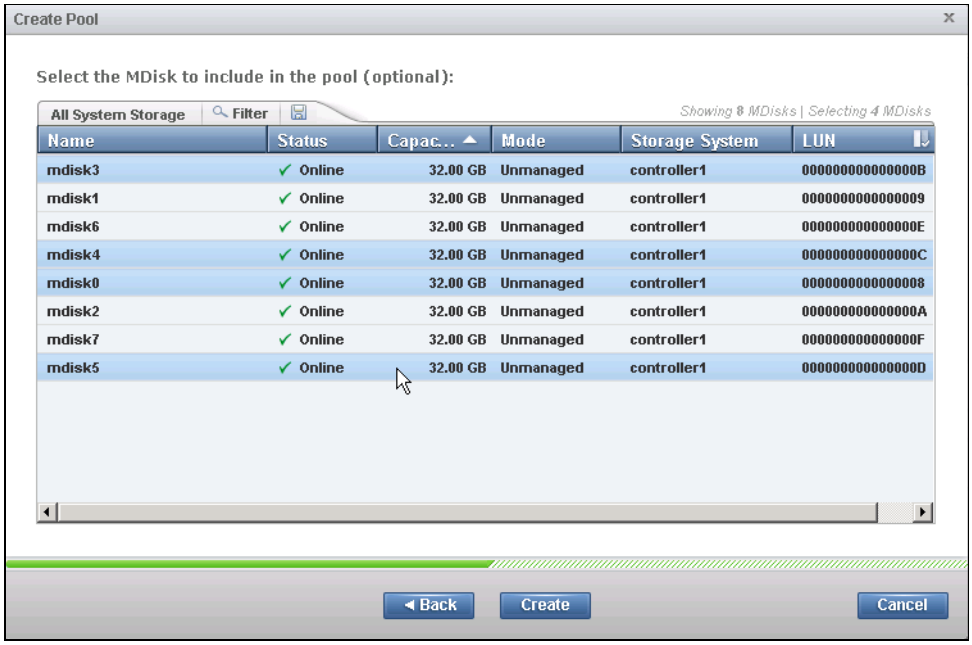

Figure 7-29 Create Pools: part 2 of 2

To assign an unmanaged MDisk to an existing pool, select the MDisk, click **Actions** → **Assign to Pool** (Figure 7-30).

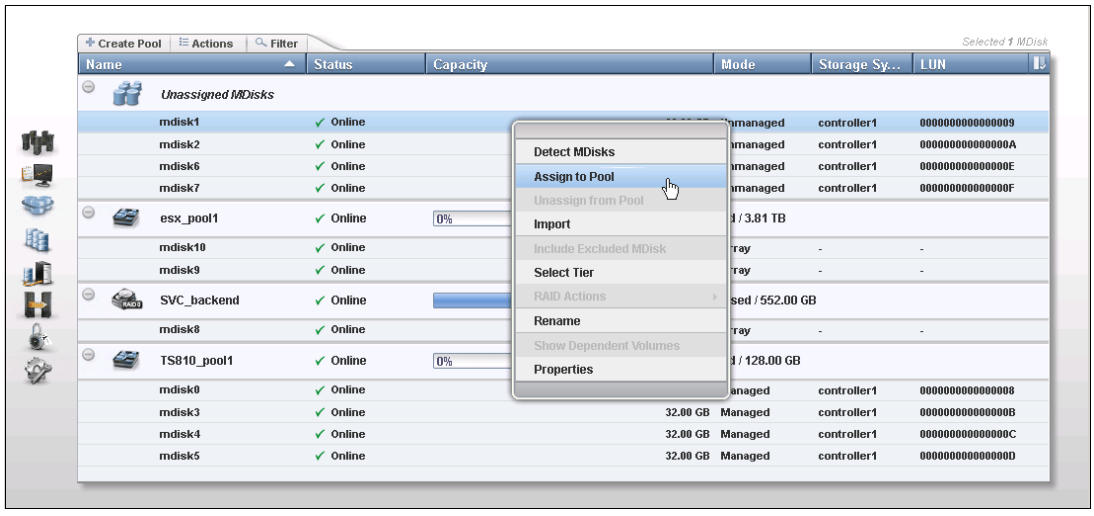

Figure 7-30 Assign an unmanaged MDisk to a storage pool

Existing data: If existing data is on the unmanaged MDisks that you need to preserve, do *not* select **Assign to Pool** on this LUN because this action deletes the data. Use **Import** instead, which is described in 7.2.2, “Importing MDisks” on page 203.

Choose the storage pool to which you want to add the MDisk and click **Add to Pool** (Figure 7-31).

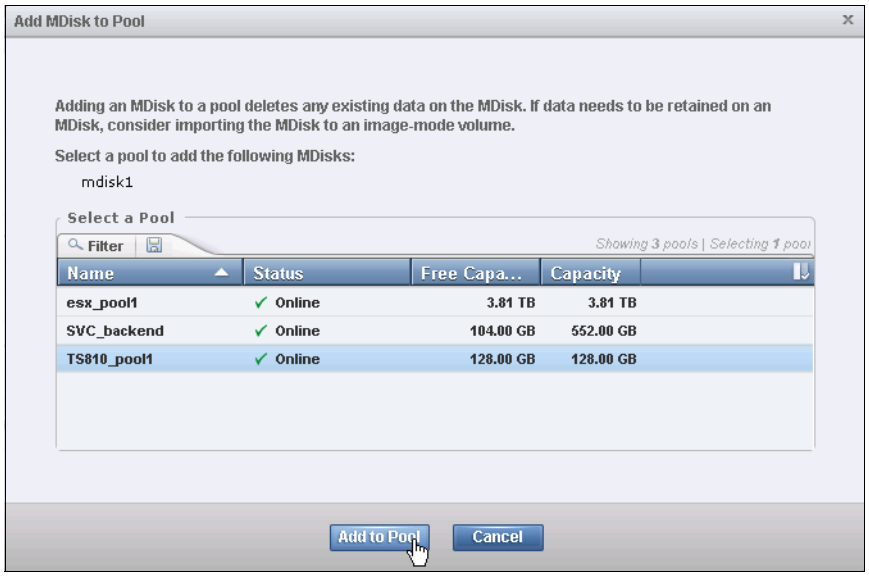

Figure 7-31 Add MDisk to Pool

After the IBM Storwize V7000 system completes this action, the MDisk shows up in the pool that it was added to (Figure 7-32).

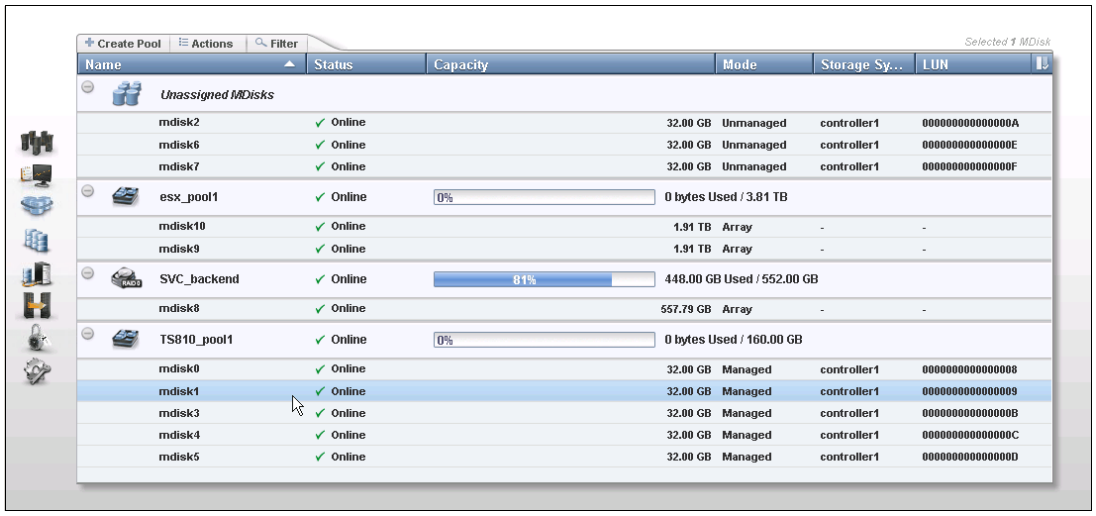

Figure 7-32 MDisk added to pool

In some cases, you might want to remove MDisks from storage pools to reorganize your storage allocation. You can remove MDisks from storage pools by selecting the MDisks and clicking **Unassign from Pool** from the **Actions** drop-down menu (Figure 7-33).

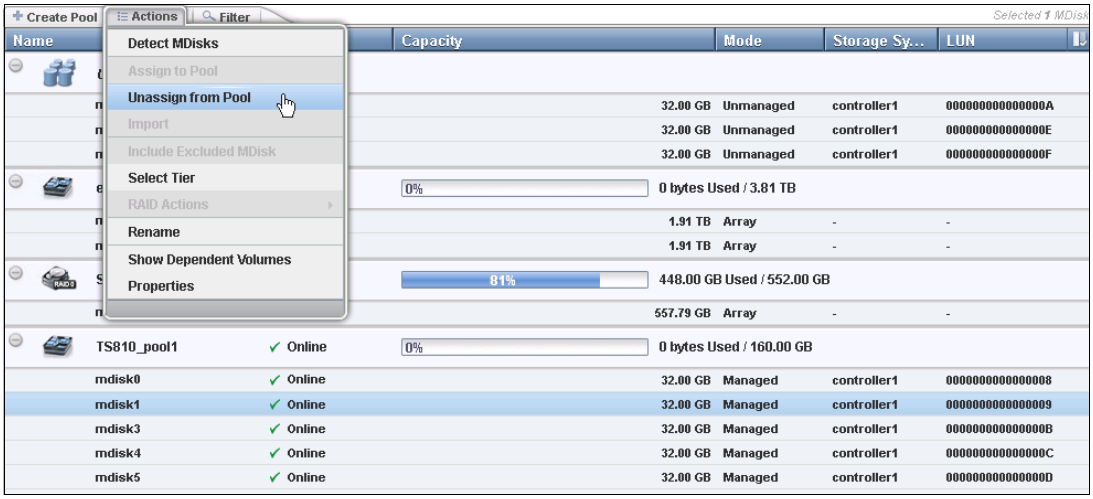

Figure 7-33 Unassign an MDisk from the storage pool

Confirm the number of MDisks you want to remove (Figure 7-34). If you have data on the MDisks, and you still need to remove the MDisks from the pool, select the following check box: **Remove the MDisk from the storage pool even if it has data on it. The system migrates the data to other MDisks in the pool**.

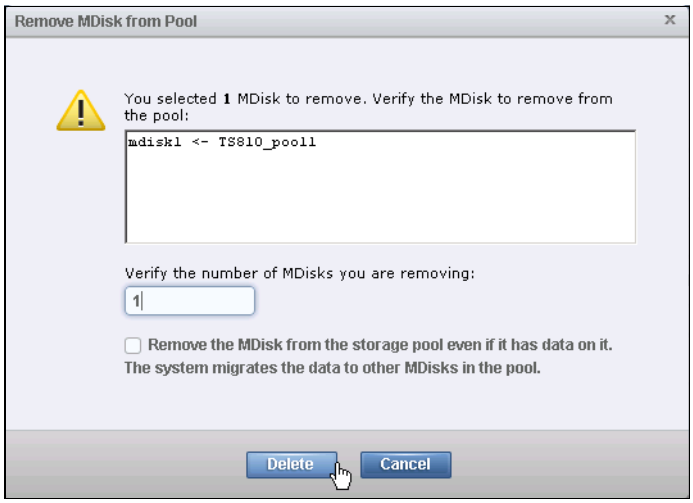

Figure 7-34 Confirm the removal of MDisk from the pool

Available capacity: The data on the MDisk that is being unassigned from the storage pool will be migrated to available space on the remaining MDisks. Make sure that you have enough available capacity remaining in the storage pool for the data on the removed MDisk to be migrated.

After you click **Delete**, data migration from the MDisks to be removed starts. You can find the migration progress in the Running Tasks status indicator (Figure 7-35).

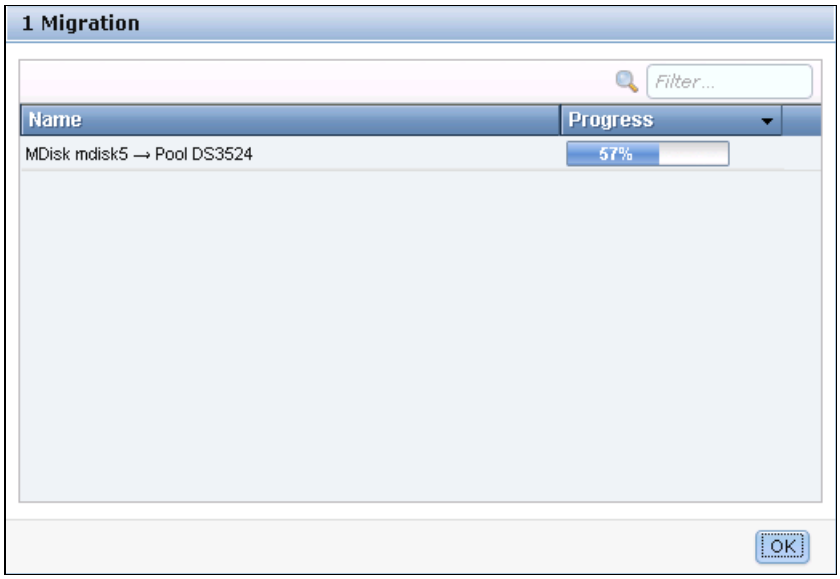

Figure 7-35 Data migration progress when removing MDisks from the pool

7.2.2 Importing MDisks

LUNs that are hosted on external storage systems can be migrated into IBM Storwize V7000 internal storage. Also, hosts that were previously directly attached to those external storage systems can continue to use their storage that is now presented by the IBM Storwize V7000 instead.

To achieve this configuration, the existing external LUNs must be imported as an *image-mode volume* using the Import option. This action is possible for unmanaged MDisks only. Those disks must not have been added to a pool, as described in 7.2.1, “Adding MDisks to storage pools” on page 199.

If the Import option is used and no existing storage pool is chosen, a temporary *migration pool* is created to hold the new image-mode volume. This image-mode volume has a direct block-for-block translation from the imported MDisk to the volume presented from now on by the IBM Storwize V7000 and existing data is being preserved.

Figure 7-36 shows an example of how to import an unmanaged MDisk. Select the unmanaged MDisk and click **Import** from the **Actions** drop-down menu.

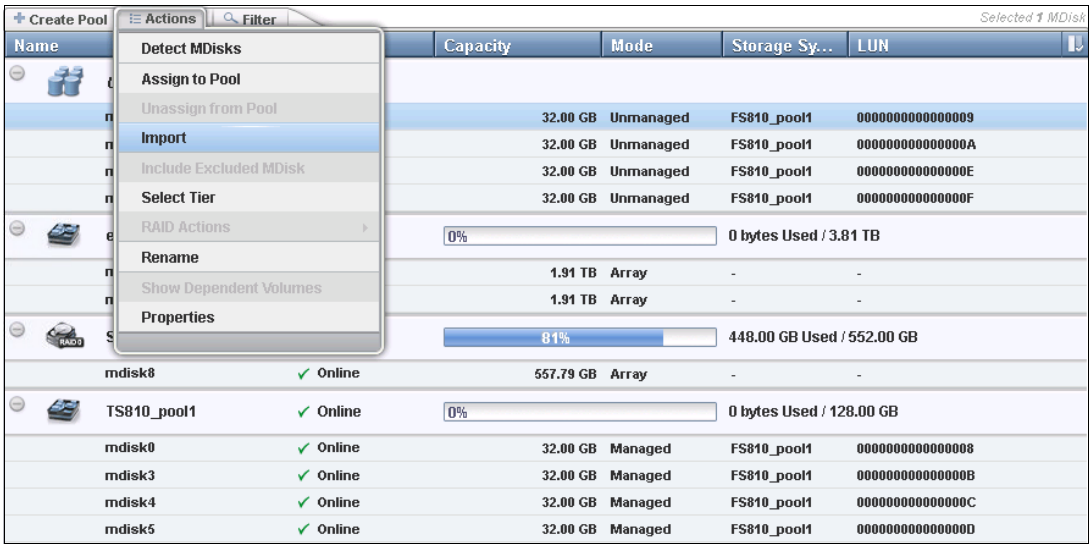

Figure 7-36 Import MDisk as image-mode volume

Figure 7-37 shows that the import wizard is activated; the wizard guides you through the import process:

1. Select whether to enable or disable caching on the image-mode volume, and then click **Next**. Clear the **Enable Caching** check box if you use copy services on the external storage system hosting the LUN. A preferred practice is to use the copy services of IBM Storwize V7000 for volumes that are virtualized. More information about virtualizing external storage is in Chapter 9, “External storage virtualization” on page 297.

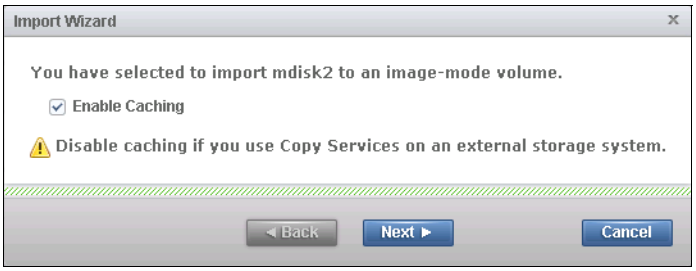

Figure 7-37 Import Wizard, enable caching on image-mode volume

2. Choose to import the volume into the existing storage pool or use a temporary pool for migration (Figure 7-38 on page 205). If you choose to migrate to a temporary pool, select the extent size characteristics of the storage pool. Select the extent size and click **Finish**.

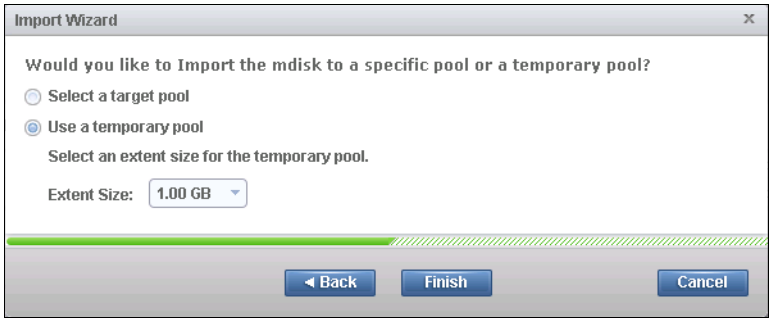

Figure 7-38 Import Wizard, select target storage pool

3. You can check the Task completion window on the action performed. Note that an image-mode volume from exactly one MDisk was created and placed into a new storage pool. See Example 7-1.

Example 7-1 Task window showing CLI command on Migrate MDisk action

```
svctask mkvdisk -iogrp io_grp0 -mdisk mdisk1 -mdiskgrp MigrationPool_1024 -name
FS810_pool1_0000000000000009 -syncrate 80 -vtype image
```

4. This volume is now presented through IBM Storwize V7000 pool and can be mapped to the original host in the *pass through* mode. The data is still physically present on the physical disk of the original storage controller system and no automatic migration process is running at this time. If needed, the image-mode volume can be migrated manually into an internal pool (array) by selecting either **Migration to Another Pool** or **Volume Copy Actions** on the Volumes panel of IBM Storwize V7000 GUI. More information about these volume actions can be found in Chapter 8, “Advanced host and volume administration” on page 221.

   Figure 7-39 shows the Volume panel with the imported image-mode MDisk.

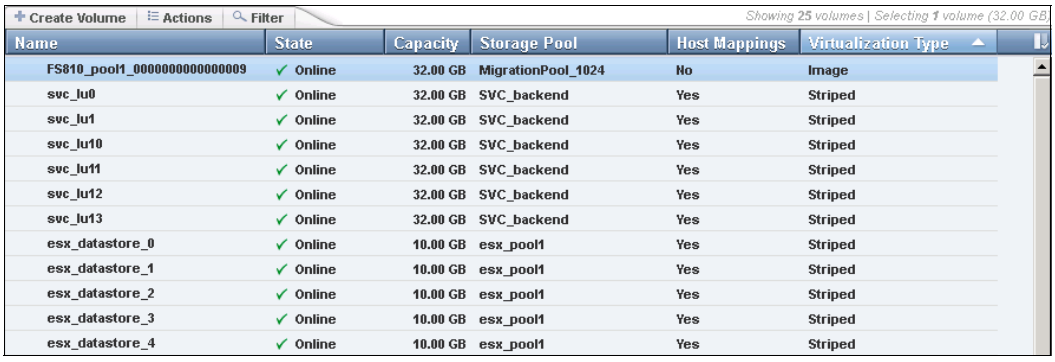

Name	State	Capacity	Storage Pool	Host Mappings	Virtualization Type
FS810_pool1_0000000000000009	Online	32.00 GB	MigrationPool_1024	No	Image
svc_lu0	Online	32.00 GB	SVC_backend	Yes	Striped
svc_lu1	Online	32.00 GB	SVC_backend	Yes	Striped
svc_lu10	Online	32.00 GB	SVC_backend	Yes	Striped
svc_lu11	Online	32.00 GB	SVC_backend	Yes	Striped
svc_lu12	Online	32.00 GB	SVC_backend	Yes	Striped
svc_lu13	Online	32.00 GB	SVC_backend	Yes	Striped
esx_datastore_0	Online	10.00 GB	esx_pool1	Yes	Striped
esx_datastore_1	Online	10.00 GB	esx_pool1	Yes	Striped
esx_datastore_2	Online	10.00 GB	esx_pool1	Yes	Striped
esx_datastore_3	Online	10.00 GB	esx_pool1	Yes	Striped
esx_datastore_4	Online	10.00 GB	esx_pool1	Yes	Striped

Figure 7-39 Imported image-mode MDisk presented as an image-mode volume

5. If you want to migrate the data from the external volume (that is being imported) to existing V7000 internal MDisks (arrays), choose an existing destination storage pool in step 2 on page 204; only pools with sufficient free extent capacity are listed (Figure 7-40). The data migration begins automatically after the MDisk is imported successfully.

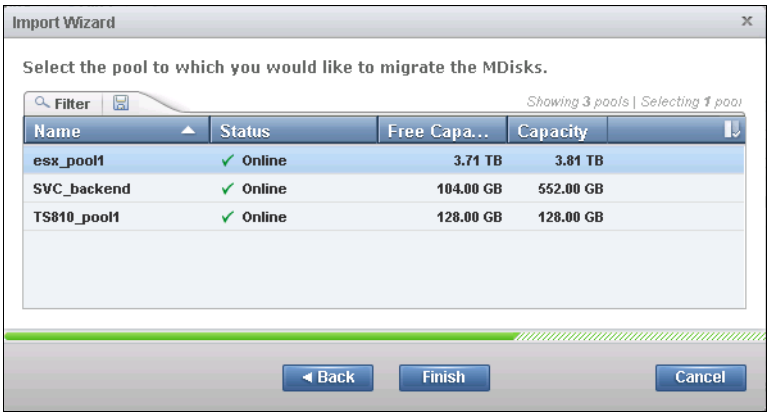

Figure 7-40 Import Wizard: Select existing storage pool

You can check the migration progress in the Running Tasks status indicator (Figure 7-41) or by selecting **Pools** → **System Migration**.

Figure 7-41 Migration progress in the status indicator of Running Tasks

Note: Notice the difference between importing an MDisk into a new temporary storage pool and importing an MDisk into an existing storage pool (array). In the first case we only *hide* the original storage controller LUN behind IBM Storwize V7000 system, not moving any of its data blocks; in the latter case, we directly migrate the data from the original storage controller LUN to IBM Storwize V7000 internal storage pool (array).

6. After the migration completes, you can find the volume in the chosen destination pool (Figure 7-42).

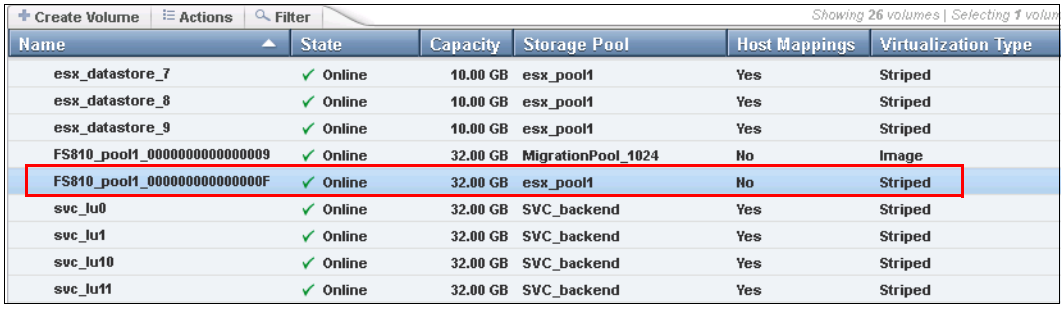

Figure 7-42 Imported Image MDisk migrated to destination pool as a striped volume

All data has been migrated off the source MDisk and striped to MDisks in the destination storage pool. The source MDisk status is changed to managed and is associated with an automatically created migration pool. It can be further used as a regular MDisk to host volumes (Figure 7-43).

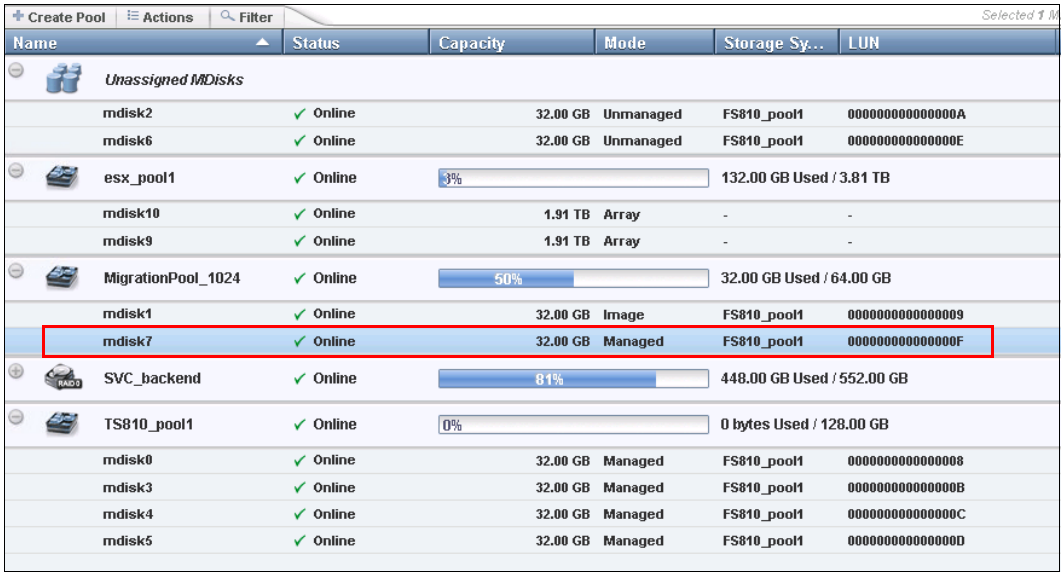

Figure 7-43 MDisk mode changed to managed

Alternatively, the migration into another pool can be done by selecting **Pools** → **System Migration**. Migration is described in more detail in Chapter 6, “Migration wizard” on page 167.

This feature is normally used as a vehicle to migrate data from existing external LUNs into storage pools, either located internally or externally, on the IBM Storwize V7000. You should *not* use image-mode volumes as a *long-term solution* for reasons of performance and reliability.

To migrate an image-mode volume into a regular storage pool, select **Pools** → **System Migration** from the dynamic menu, select the volume to be migrated, and click **Actions** → **Start Migration**. Choose a target storage pool into which to migrate the data and click **Add to Pool** (Figure 7-44). Only regular pools, and no migration pools, are offered as a destination.

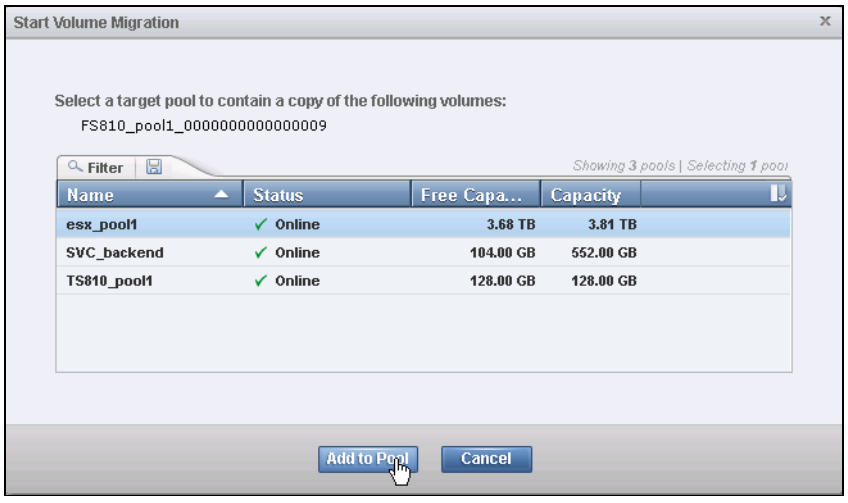

Figure 7-44 Migrate an Image-Mode Volume into a regular storage pool

The migration internally uses the volume copy function, which creates a second copy of the existing volume in the chosen target pool. For more information about the volume copy function, see 8.6, “Advanced volume copy functions” on page 279. If the migration process is interrupted for any reason, select the volume being migrated and click **Actions** → **Stop Migration** (Figure 7-45). The process of adding a secondary volume copy is interrupted.

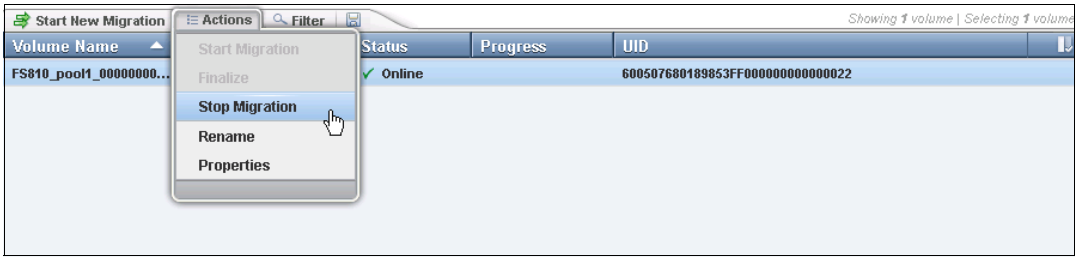

Figure 7-45 Stop Migration

When the migration progress reaches 100%, click **Actions** → **Finalize** to complete the migration process (Figure 7-46).

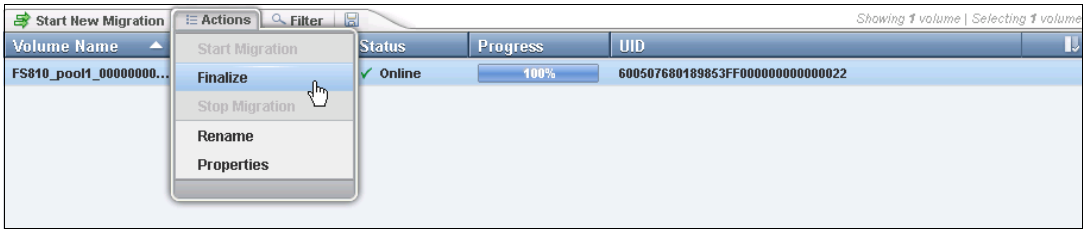

Figure 7-46 Finalize volume migration

The original volume copy on the image-mode MDisk is deleted and the newly created copy is kept. The source image-mode MDisk is removed from the Migration Pool and changes its mode back to unmanaged.

7.2.3 RAID action for MDisks

Internal drives in the IBM Storwize V7000 are managed as Array mode MDisks, on which several RAID actions can be performed. Choose the appropriate Array MDisk by selecting **Pools** → **MDisks by Pools** from the dynamic menu, and then click **Actions** → **RAID Actions** (Figure 7-47).

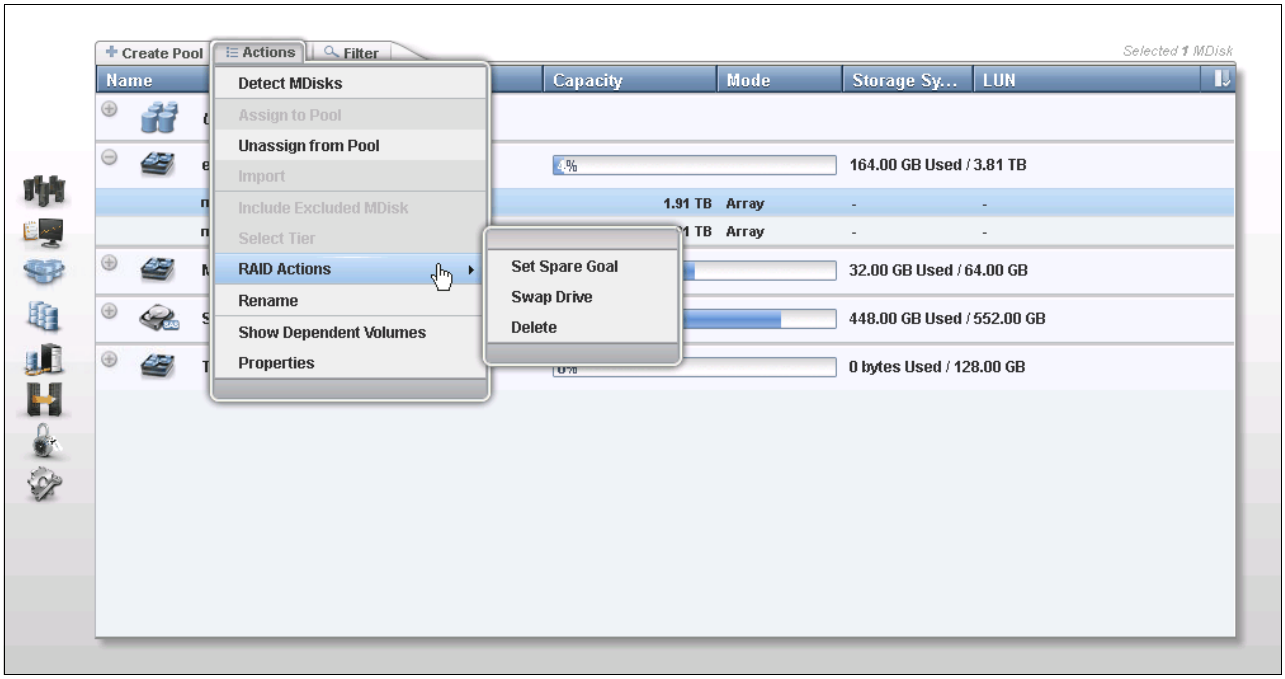

Figure 7-47 RAID actions on Array MDisks

You can choose the following RAID actions:

- Set Spare Goal: You can set the number of spare drives that is required to protect the array from drive failures. If the number of drives 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 either replacing defunct drives, which occupy a spare drive, or by adding additional drives as spare. Remember, during the internal drive configuration, spare drives are automatically assigned according to the chosen RAID preset’s spare goals, as described in 7.1.2, “Configuring internal storage” on page 187.

- Swap Drive: This 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 has failed, or that is expected to fail soon, for example, as indicated by an error message in the event log. Select an MDisk that contains the drive to be replaced and click **RAID Actions** → **Swap Drive**. In the Swap Drive window (Figure 7-48), select the member drive to be replaced and click **Next**.

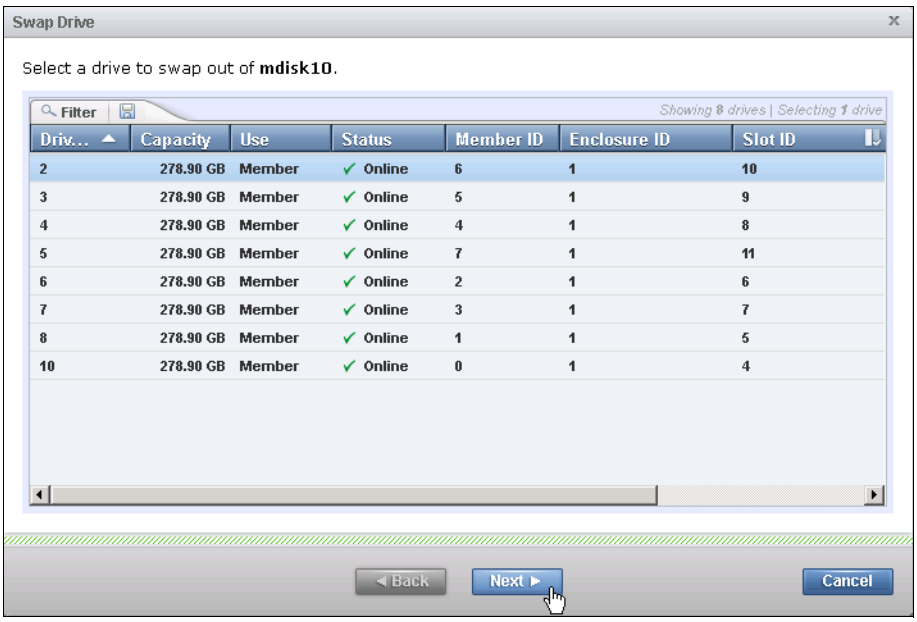

Figure 7-48 Select a drive to swap

In shown in Figure 7-49, select an available Candidate or Spare drive to swap into the Array MDisk and click **Finish**.

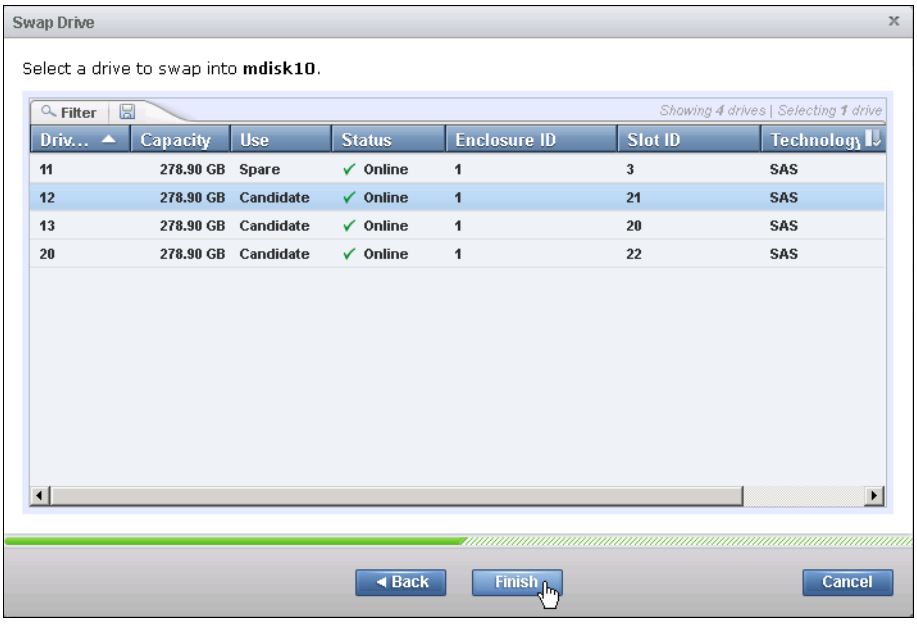

Figure 7-49 Select a drive to swap into MDisk

The exchange of the drives starts and runs in the background, and the volumes on the affected MDisk remain accessible.

- Delete: An Array MDisk can be deleted from the MDisks by Pools panel by clicking **RAID Actions** → **Delete**. You must confirm the deletion by entering the correct number of array MDisks to be deleted (Figure 7-50). If data is on the MDisk, force the deletion by selecting the check box (**Delete the RAID array MDisk even if it has data on it. The system migrates the data to other MDisks in the pool**).

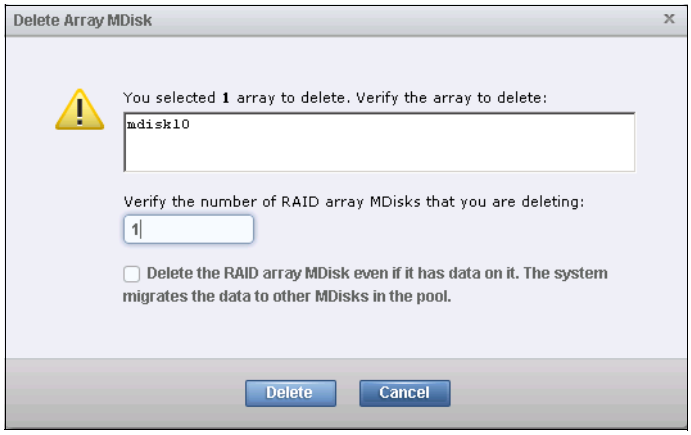

Figure 7-50 Delete Array MDisks window

Available capacity: Make sure that you have enough available capacity left in the storage pool for the data on the MDisks to be removed.

Data on MDisks is migrated to other MDisks in the pool, assuming enough space is available on the remaining MDisks in the pool. After deleting an MDisk from a pool, its member drives return to candidate mode.

7.2.4 Selecting the tier for MDisks

The IBM Storwize V7000 Easy Tier feature is described in Chapter 10, “Advanced features for storage efficiency” on page 309. In this section, we show how to adjust the tier settings.

Two tiers exist:

- Generic SSD tier is for storage made of solid-state drives, which is the faster-performing storage.
- Generic HDD tier is for everything else.

Internal drives have their tier assigned automatically by the IBM Storwize V7000. MDisks on external storage systems are assigned the *generic HDD tier* by default. This setting can be changed manually by the user. To assign a specific tier to an MDisk, select **Pools** → **MDisks by Pool** from the dynamic menu and select **Actions** → **Select Tier** (Figure 7-51).

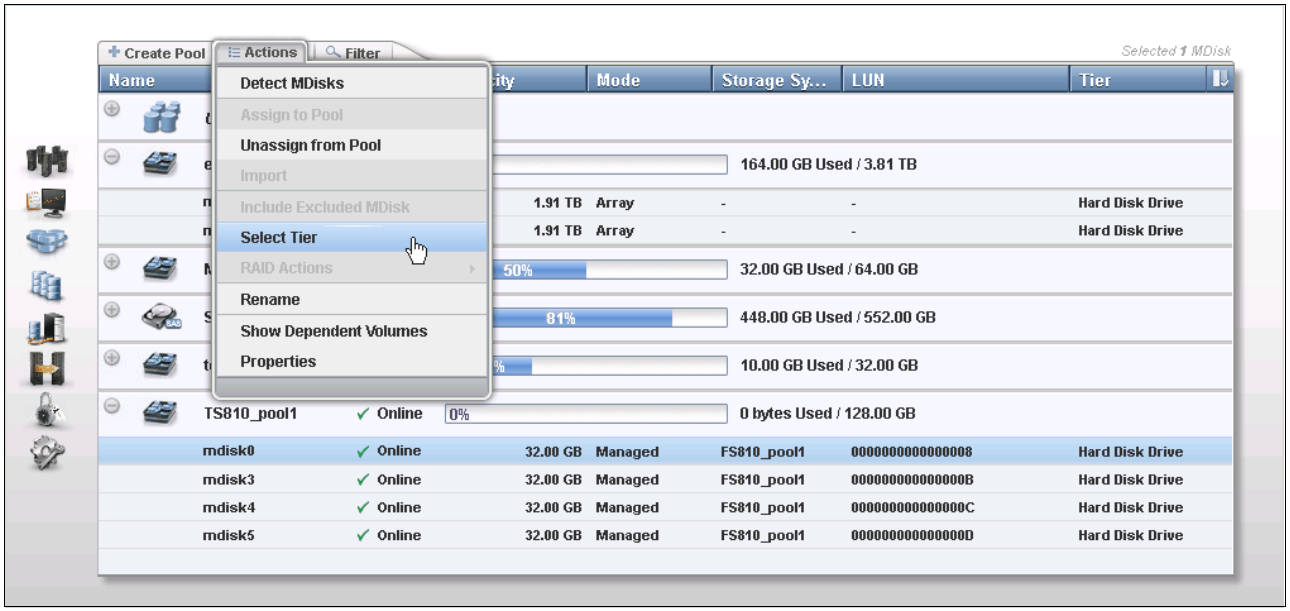

Figure 7-51 Select Tier for an MDisk

We assign the **Solid-State Drive** tier to mdisk0 (Figure 7-52). This MDisk is a LUN made of real SSDs in an external storage IBM Flash System. The tier that was assigned by default is Hard-Disk Drive.

Figure 7-52 Assign desired tier to an MDisk

After the action completes successfully, the MDisk can be found in the SSD tier (Figure 7-53).

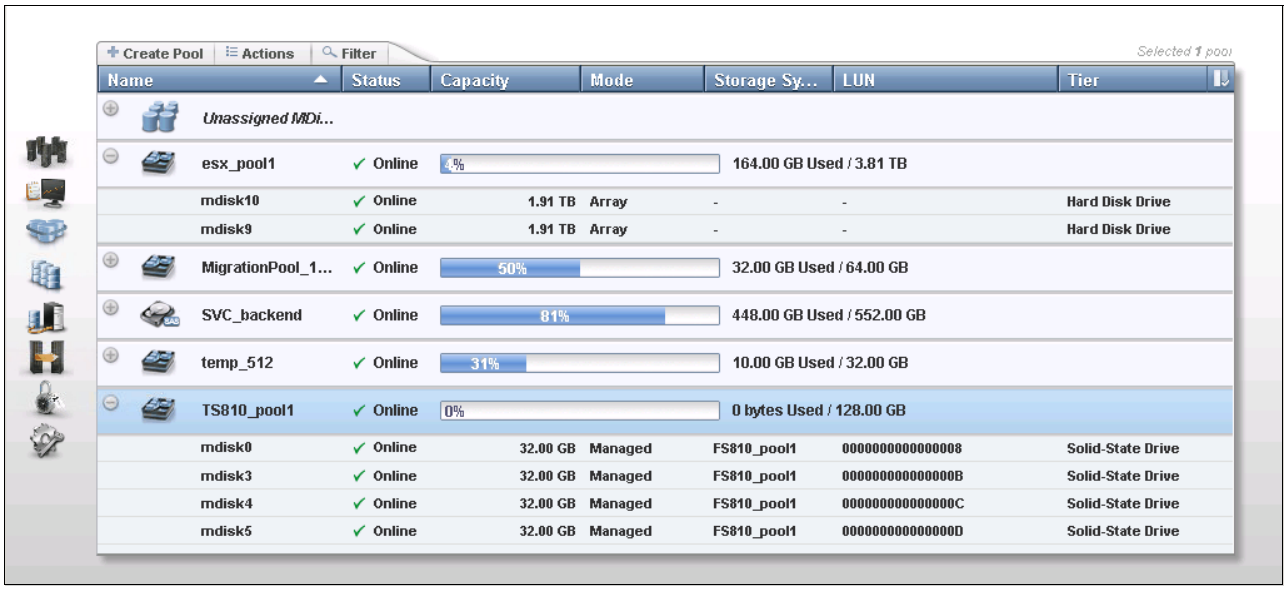

Figure 7-53 The tier has been assigned to the MDisk

7.2.5 Additional actions on MDisks

Several additional actions can be performed on MDisks:

- Detect MDisks: Selecting **Actions** → **Detect MDisks** is useful if you have external storage controllers in your environment (for more information, see Chapter 9, “External storage virtualization” on page 297). The Detect MDisks action initiates a rescan of the Fibre Channel network. It discovers any new MDisks that are mapped to the IBM Storwize V7000 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 V7000 configuration to reflect any changes.

  When external storage controllers are added to the IBM Storwize V7000 environment, the IBM Storwize V7000 automatically discovers the controllers and the LUNs that are presented by those controllers and lists them as unmanaged MDisks. However, if you attached new storage and the IBM Storwize V7000 does not detect it, you might need to use the Detect MDisks action before the system automatically detects the new LUNs. If the configuration of the external controllers is modified after, the IBM Storwize V7000 might be unaware of these configuration changes. Use the Detect MDisks action to rescan the Fibre Channel network and update the list of unmanaged MDisks.

  MDisks detection: The Detect MDisks action is asynchronous. Although the task appears to be finished, it still might be running in the background.

- Include Excluded MDisk: An MDisk might be excluded from the IBM Storwize V7000 because of multiple I/O failures. These failures might be caused, for example, by link errors. After a fabric-related problem is corrected, the excluded disk can be added back into the IBM Storwize V7000 by selecting the MDisk and clicking **Actions** → **Include Excluded MDisk**.

- Rename: MDisks can be renamed by selecting the MDisk and clicking **Actions** → **Rename**. Input the new name of your MDisk and click **Rename** (Figure 7-54).

Figure 7-54 Rename MDisk

- Show Dependent Volumes: The volumes that depend on a particular MDisk can be displayed by selecting the MDisk and clicking **Actions** → **Show Dependent Volumes**. The volumes are listed with general information (Figure 7-55).

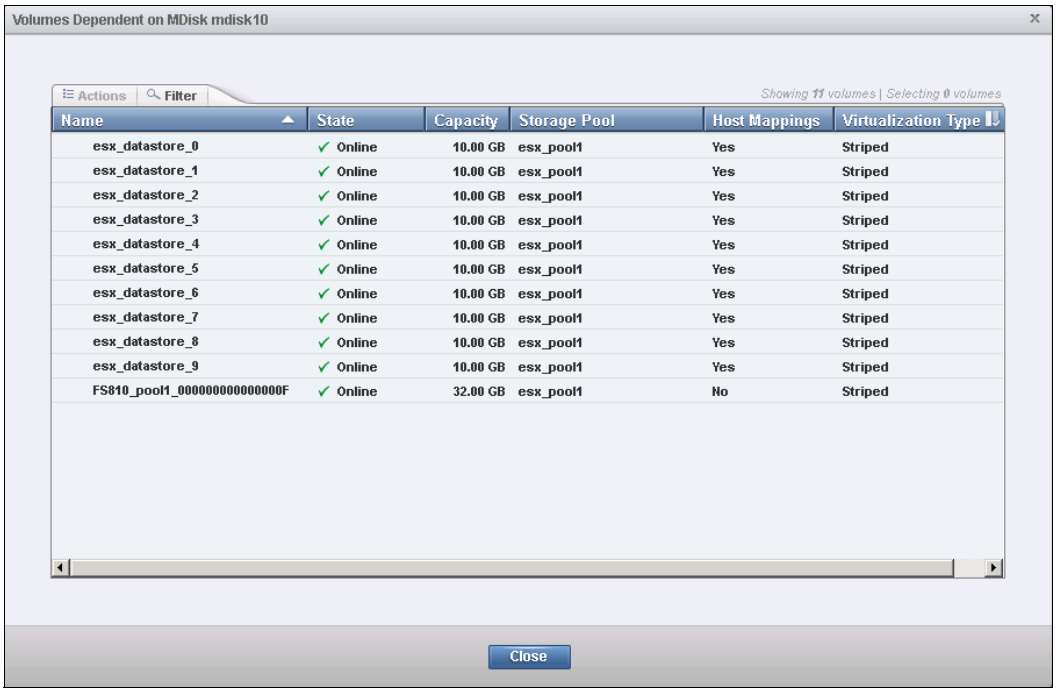

Figure 7-55 Show Dependent Volumes

From the Volumes panel of the IBM Storwize V7000 GUI, several actions can be done to volumes by selecting the volume and choosing an action from the **Actions** menu (Figure 7-56). For more information about volume actions, see Chapter 5, “Basic volume configuration” on page 131.

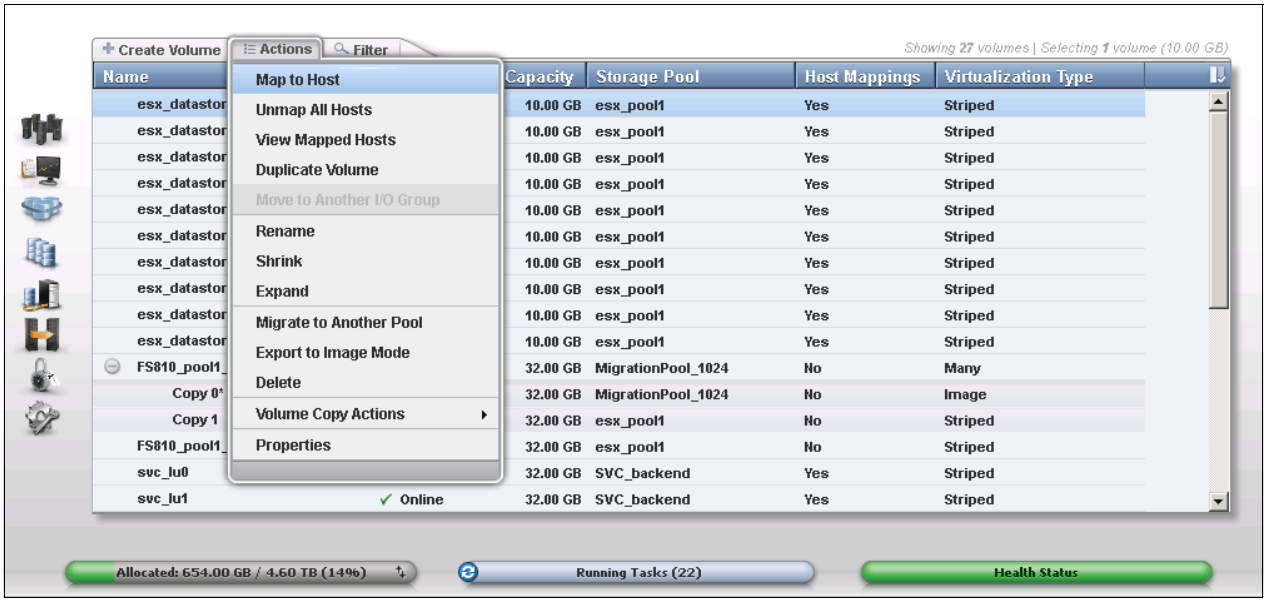

Figure 7-56 Volume actions

To view information about MDisks that are related to the specific volume, select volume **Properties**, and on the **Member MDisks** tab right-click an MDisk and select **Properties**. A window opens.

This information window has three tabs:

– Overview: Lists information about the MDisk. To show more details, click **Show Details** (Figure 7-57).

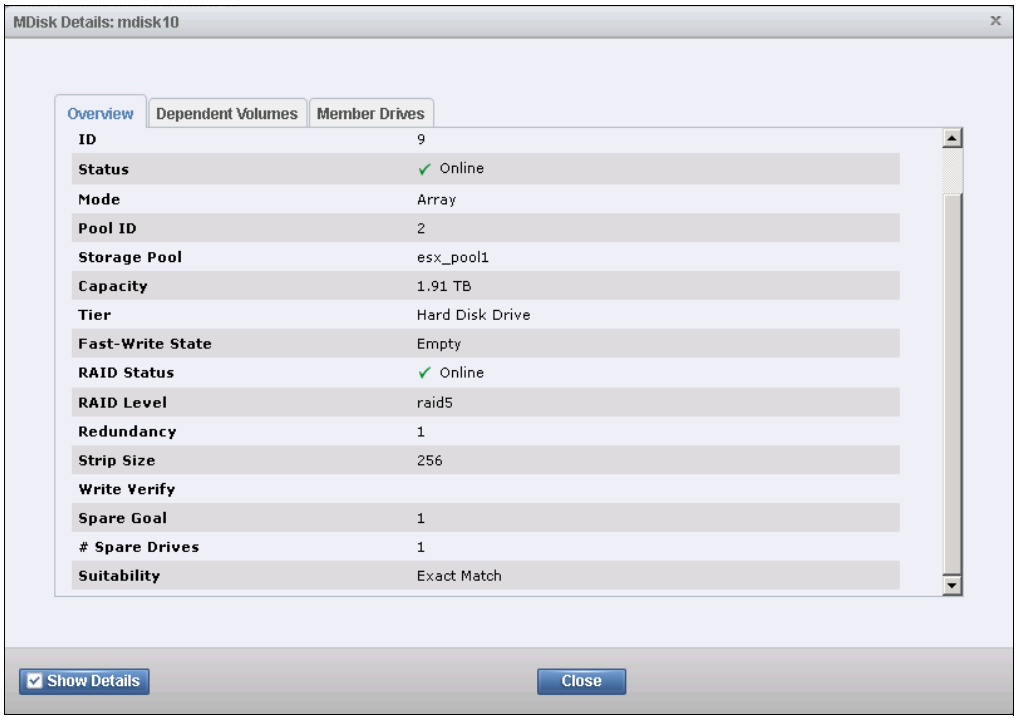

Figure 7-57 MDisk overview with detailed information

– Dependent Volumes: Lists all volumes that use extents on this MDisk (Figure 7-58).

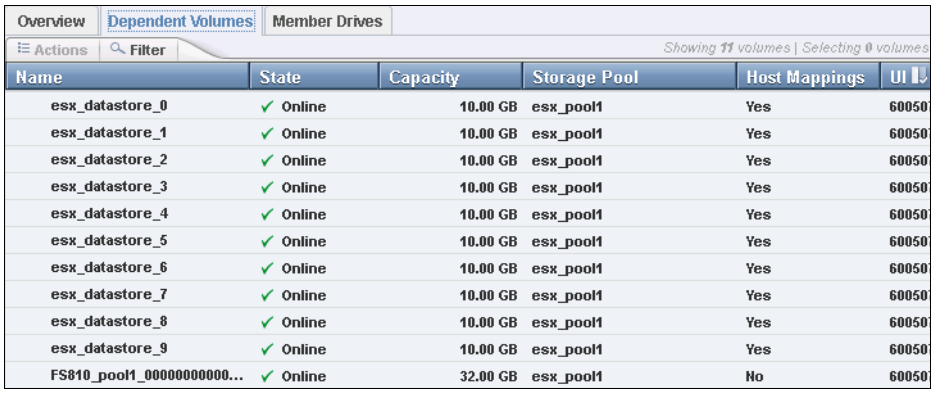

Figure 7-58 Volumes containing extents on this MDisk

– Member Drives: Is used only with array MDisks and lists all the member drives of this MDisk (Figure 7-59). Also, all actions described in 7.1.1, “Actions on internal drives” on page 183 can be done on the drives listed here.

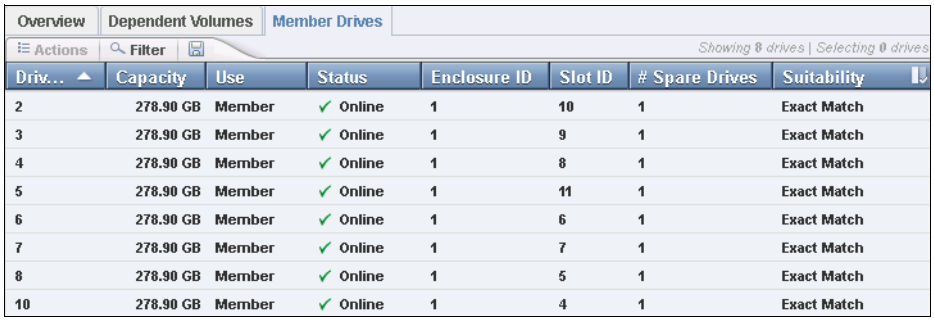

Driv...	Capacity	Use	Status	Enclosure ID	Slot ID	# Spare Drives	Suitability
2	278.90 GB	Member	Online	1	10	1	Exact Match
3	278.90 GB	Member	Online	1	9	1	Exact Match
4	278.90 GB	Member	Online	1	8	1	Exact Match
5	278.90 GB	Member	Online	1	11	1	Exact Match
6	278.90 GB	Member	Online	1	6	1	Exact Match
7	278.90 GB	Member	Online	1	7	1	Exact Match
8	278.90 GB	Member	Online	1	5	1	Exact Match
10	278.90 GB	Member	Online	1	4	1	Exact Match

Figure 7-59 Member Drives tab

7.3 Working with storage pools

Storage pools aggregate a disk capacity into a single container for MDisks and provision the storage space for volumes that are allocated to storage hosts. They are the same as storage pools in an IBM SAN Volume Controller when they are using external storage for their capacity, or similar to RAID arrays on a traditional storage subsystem when they are built by using internal drives. IBM Storwize V7000 organizes storage in storage pools to ease storage management, making it more efficient. Storage pools and MDisks are managed in the MDisks by Pools panel. You can access the MDisks by Pools panel from the IBM Storwize V7000 **Overview** panel clicking the **MDisks** icon. Extended help information for storage pools is displayed, and after you click **Visit Pools**, the MDisks by Pools panel opens.

Figure 7-60 shows how to access the Pools panel from the V7000 Overview panel.

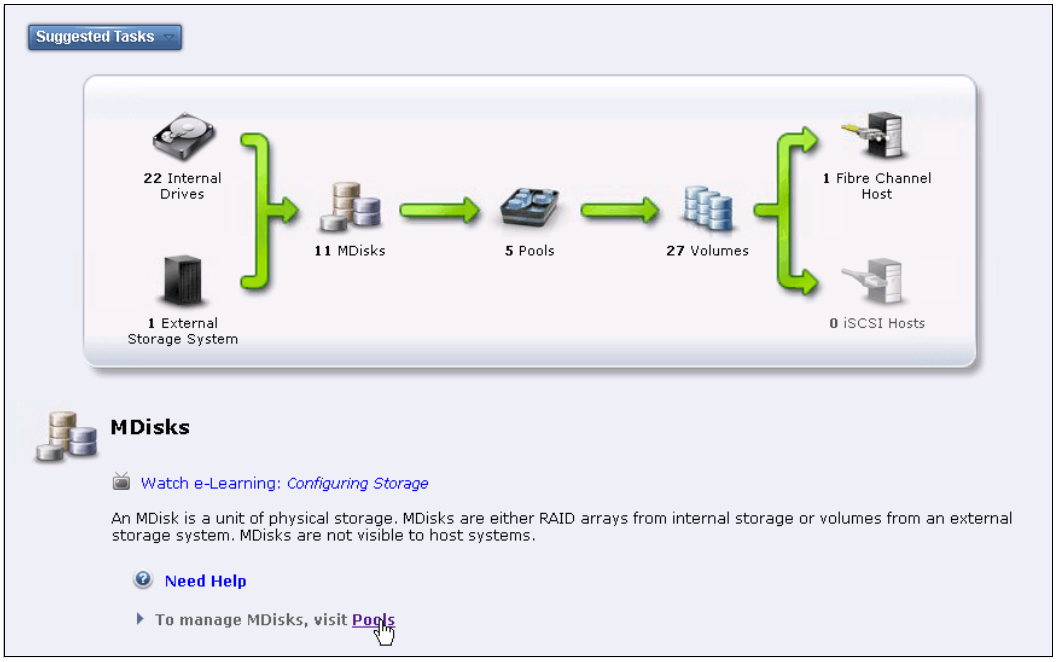

Figure 7-60 Access Pool management panel from the Overview panel

The other way to access the Pools panel is to move the mouse cursor over the **Pools** selection in the dynamic menu and click **MDisks by Pools** (Figure 7-61).

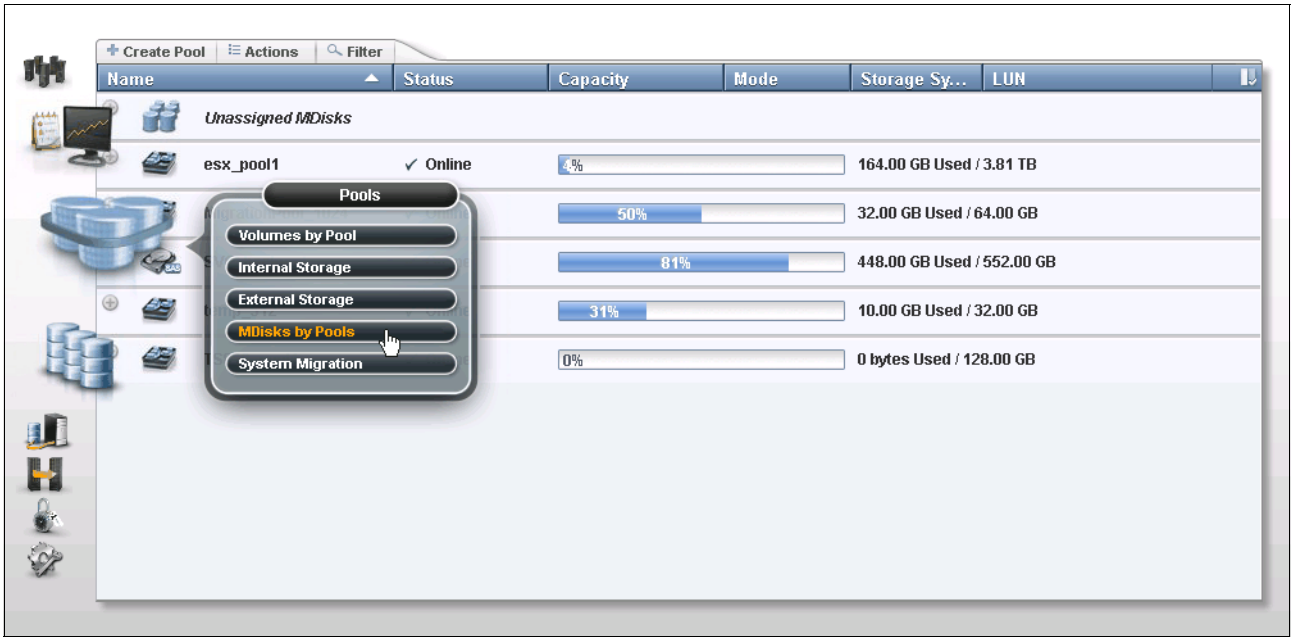

Figure 7-61 Access MDisks by Pools panel from dynamic menu

The MDisk by Pools panel (Figure 7-62) allows you to manage storage pools, both internal and external. All existing storage pools are displayed row-by-row. The first row is the *Unassigned MDisks* item, containing all unmanaged MDisks, if any exist.

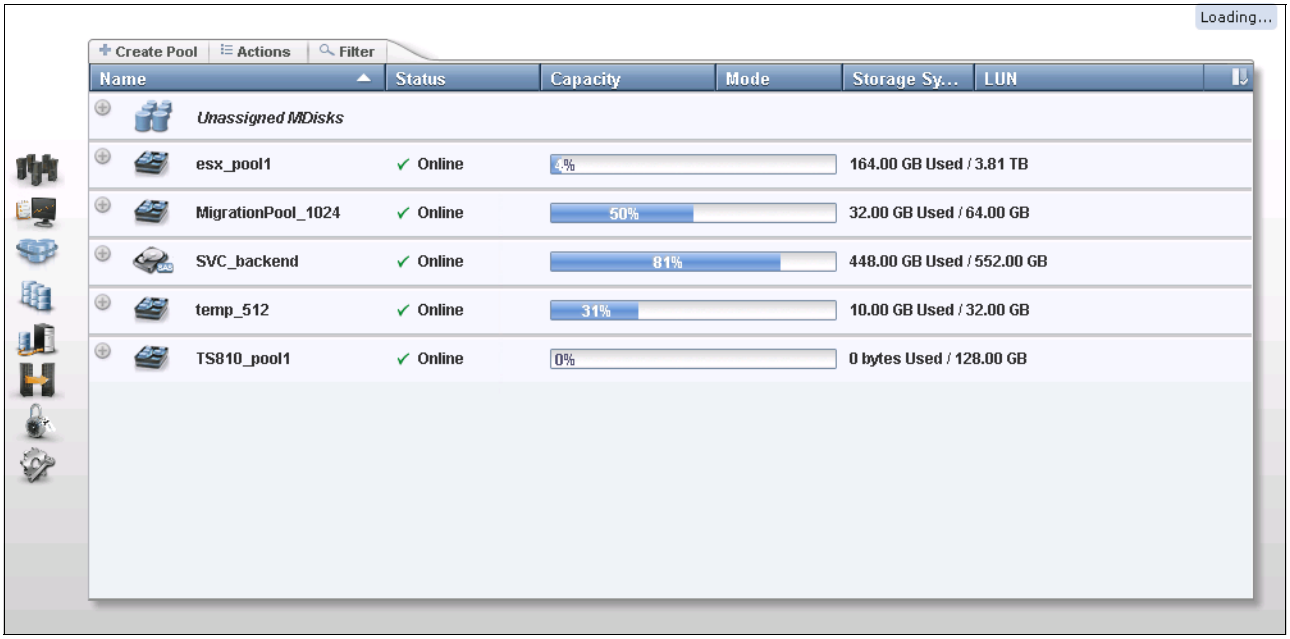

Figure 7-62 MDisks by Pools panel

To make the view more clear and structured, filters can be used to show only those storage pools that match conditions that are defined in the filter field in the upper part of the panel (Figure 7-63).

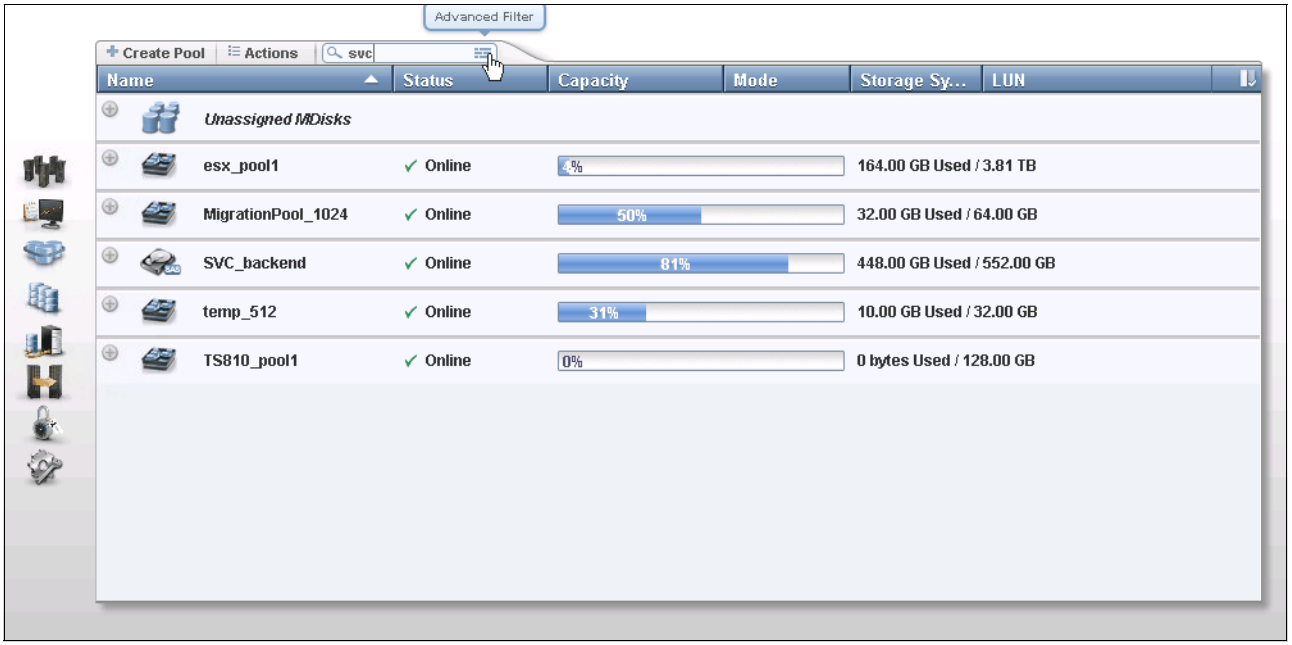

Figure 7-63 Filter the pools view

A few actions can be performed on storage pools using the **Actions** menu. A pool can be *renamed*, its *icon* can be changed, and it can be deleted from here also.

When you expand a pool's entry by clicking the plus (+) sign to the left of the pool's icon, you get access to the MDisks that are associated with this pool. You can perform all actions on them, as described in 7.2, "Working with MDisks" on page 197.

In its default view, each defined storage pool is displayed along with its assigned icon and name, status, graphical indicator showing the pool's capacity allocated to volumes, mode of the containing MDisks, storage system, and MDisk LUN ID.

To remove a storage pool that is not needed, click **Actions** → **Delete Pool** (Figure 7-64).

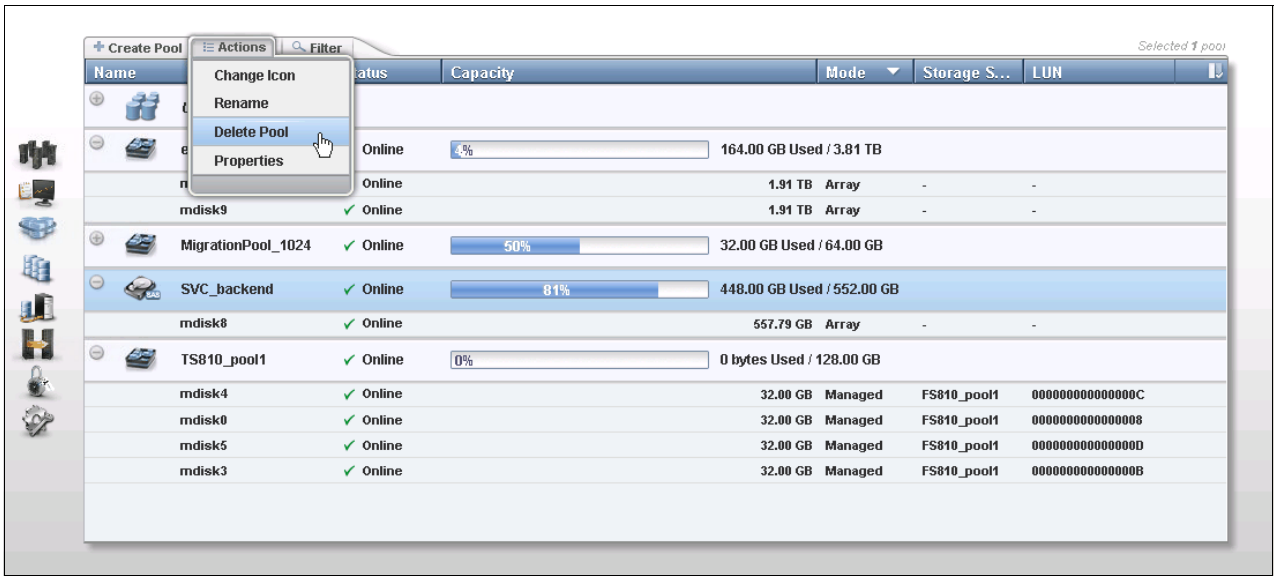

Figure 7-64 Delete Pool action

A confirmation window opens (Figure 7-65). If there are volumes with their only copy on MDisks in the pool to be deleted, they are listed in the confirmation window. If safely deleting the volumes and MDisks in this pool is possible, even though they still might be mapped to a host, select the **Delete all volumes, host mappings and MDisks that are associated with this pool** check box and click **Delete**.

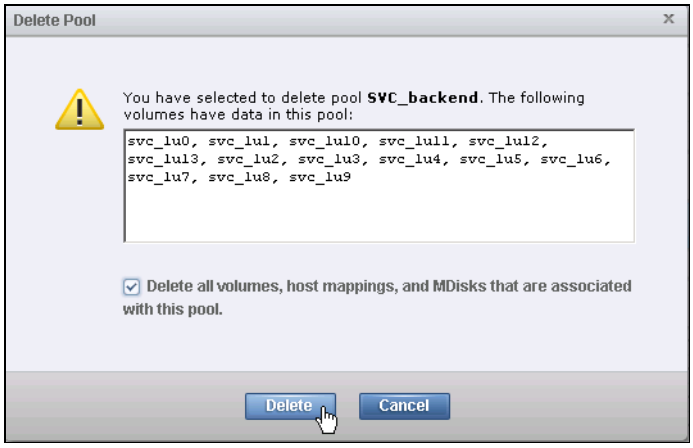

Figure 7-65 Confirm the deletion of the pool

Attention: After you delete the pool, all the data in the pool is lost except for the image-mode MDisks; their volume definition is deleted, yet the data on the imported MDisk remains untouched.

After you delete the pool, all the associated volumes and their host mappings are removed. All the managed or image-mode MDisks in the pool return to a status of unmanaged after the pool is deleted. If the pool is deleted, all the array mode MDisks in the pool are removed and all the member drives return to candidate status.

8

Advanced host and volume administration

The IBM Storwize V7000 offers many functions for volume and hosts creation and configuration. In Chapter 4, “Host configuration” on page 105 and in Chapter 5, “Basic volume configuration” on page 131, the basic host and volumes features of IBM Storwize V7000 are covered. Those chapters show how to create hosts and volumes, and how to map volumes to a host. This chapter describes other advanced host and volume administration tasks apart from Easy Tier, which is described in Chapter 10, “Advanced features for storage efficiency” on page 309.

This chapter contains the following topics:

- Advanced host administration
- Advanced volume administration

8.1 Advanced host administration

This section covers host administration, including topics such as host modification, host mappings, and deleting hosts. Basic host creation using Fibre Channel and iSCSI connectivity is described in 4.3.1, “Creating Fibre Channel hosts” on page 122 and 4.3.2, “Creating iSCSI hosts” on page 125.

We assume that you created several hosts in your IBM Storwize V7000 GUI and that some volumes are already mapped to them. We describe the three functions that are covered in the Hosts section of the IBM Storwize V7000 GUI (Figure 8-1):

- Hosts (8.1.1, “Modifying Mappings menu” on page 224).
- Ports by Host (8.2, “Adding and deleting host ports” on page 241).
- Host Mappings (8.3, “Host mappings overview” on page 249).

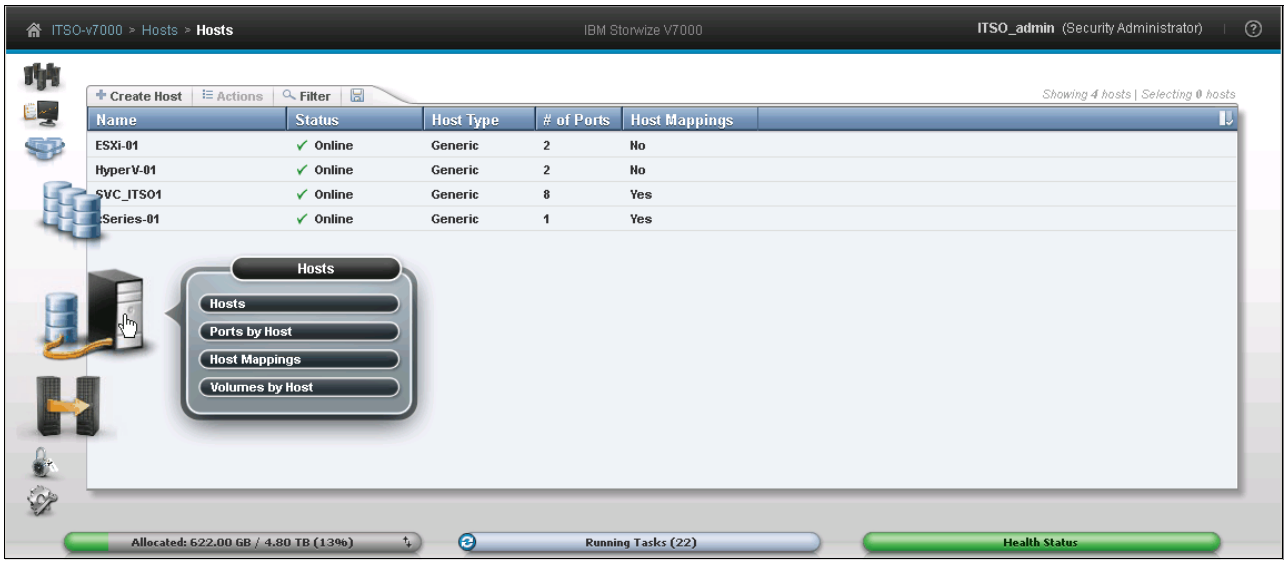

Figure 8-1 IBM Storwize V7000 Hosts menu

If you hover the mouse cursor over the **Hosts** selection in the IBM Storwize V7000 dynamic menu, the Hosts panel opens (Figure 8-2).

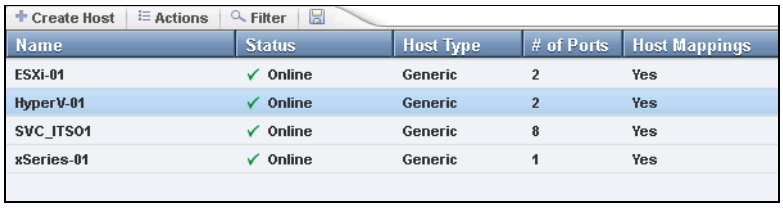

Figure 8-2 Hosts panel

As our example shows, four hosts are created and volumes are already mapped to all of them; we use these hosts to show the modification possibilities.

If you highlight a host, you can either click **Actions** (Figure 8-3) or right-click the host (Figure 8-4) to see all the available tasks.

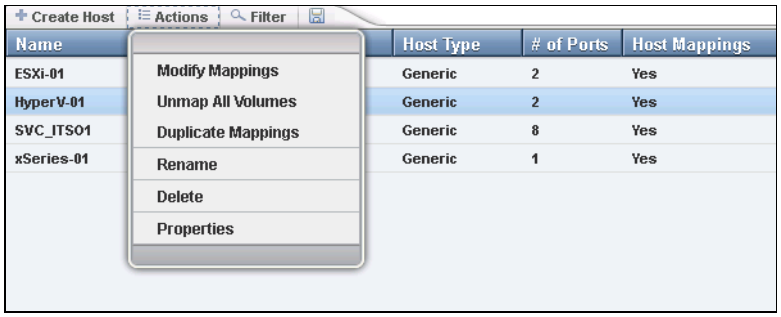

Figure 8-3 Host Actions

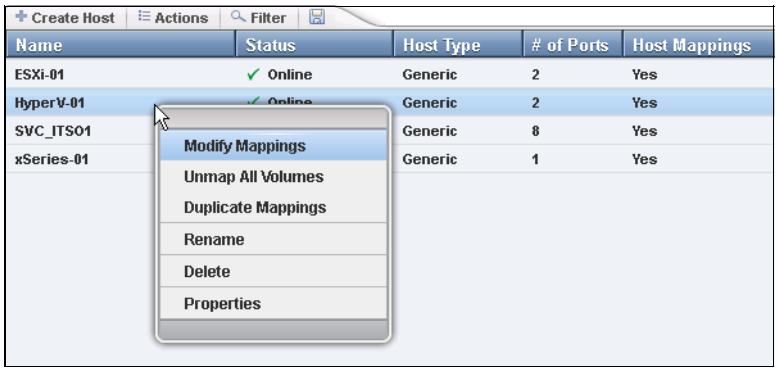

Figure 8-4 Host menu

8.1.1 Modifying Mappings menu

On the Hosts panel, select a host and click **Actions** → **Modify Mappings** (Figure 8-3 on page 223). The window shown in Figure 8-5 opens. At the upper left, you see that the highlighted host is selected. The two list boxes show all available unmapped and mapped volumes. The left list includes the volumes that are ready for mapping to this host; the right list includes the volumes already mapped. In our example, four volumes with SCSI IDs 0, 1, 2, and 3 are mapped to the host HyperV-01, and 32 more volumes are available (unmapped).

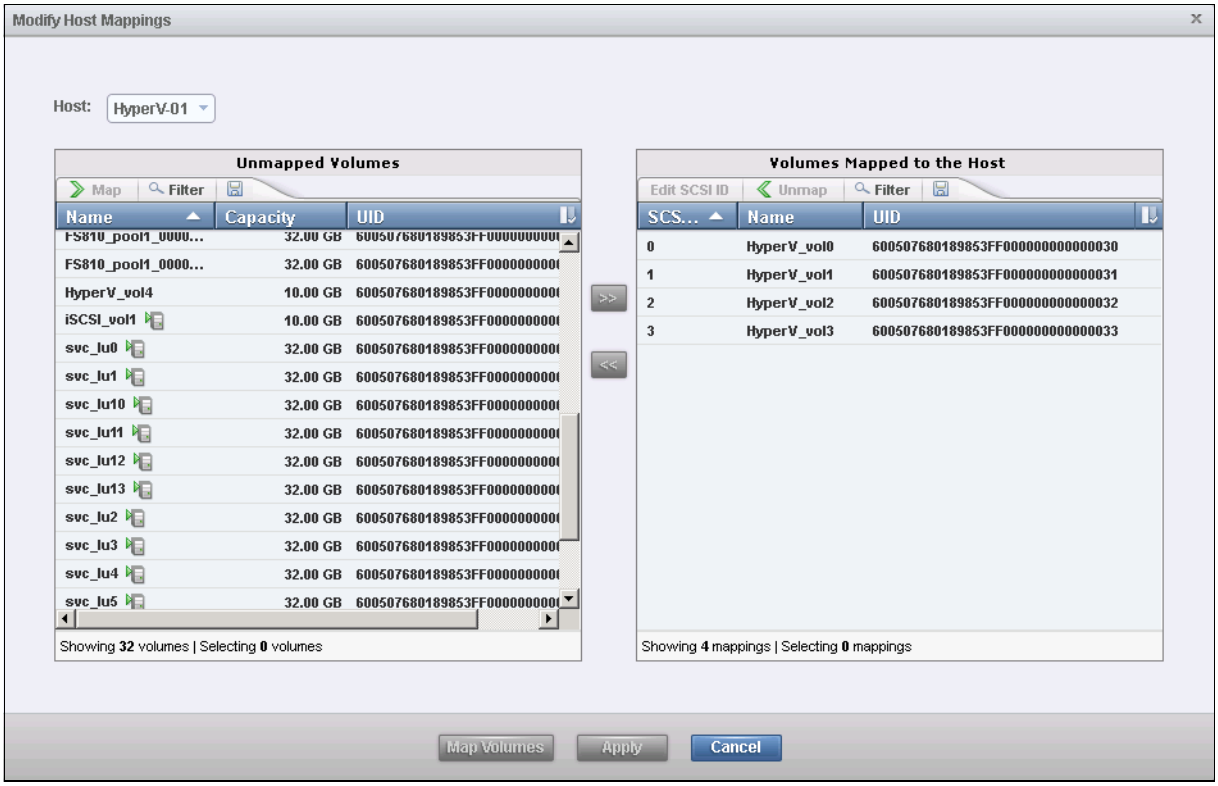

Figure 8-5 Modify Host Mappings window

To map a volume, select it in the left pane, and then click the right arrow (>>) to move the volume to the list on the right to map another volume to a host (Figure 8-6).

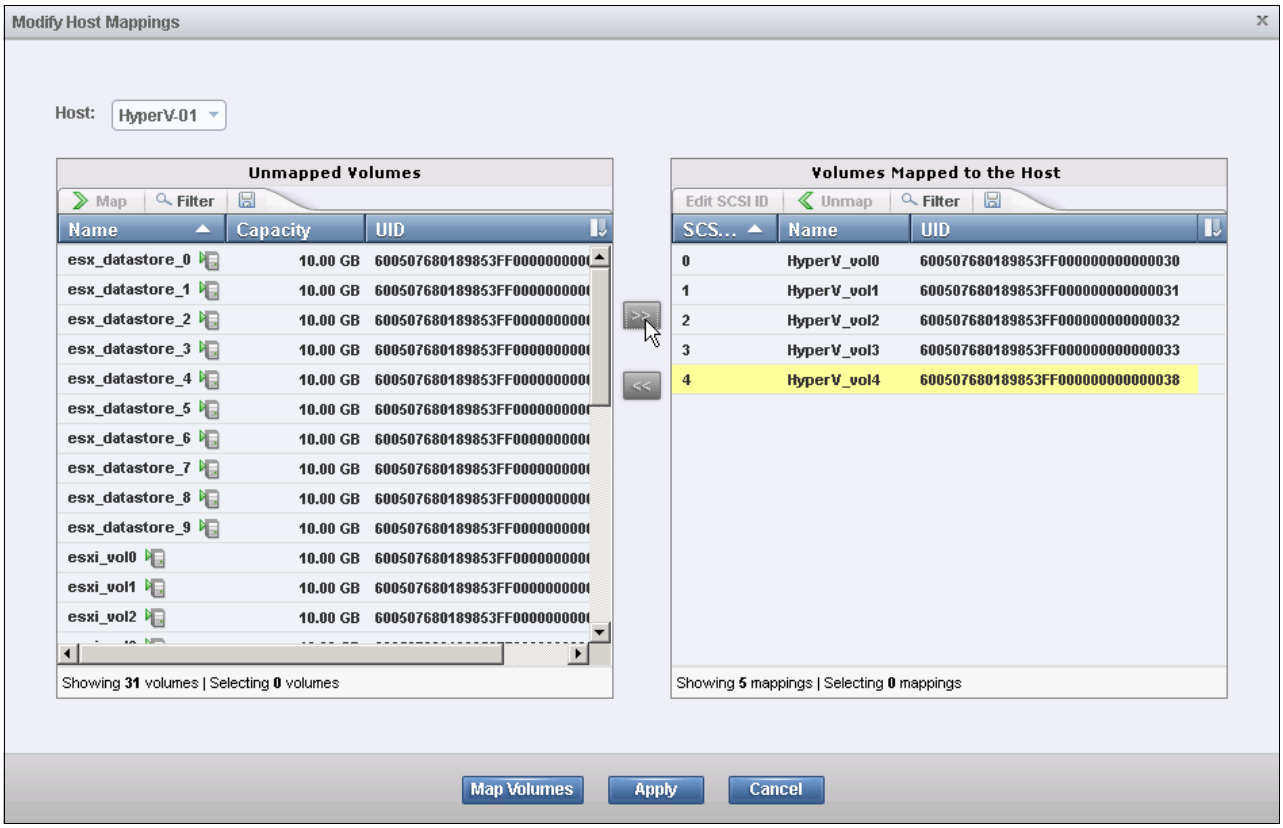

Figure 8-6 Modify Host Mappings: map volume

The changes are marked in yellow and now the Map Volumes and Apply buttons are enabled.

If you click **Map Volumes**, the changes are applied (Figure 8-7) and the window closes. If you click **Apply**, the changes are submitted to the system (Figure 8-7), but the window remains open for further changes (Figure 8-8).

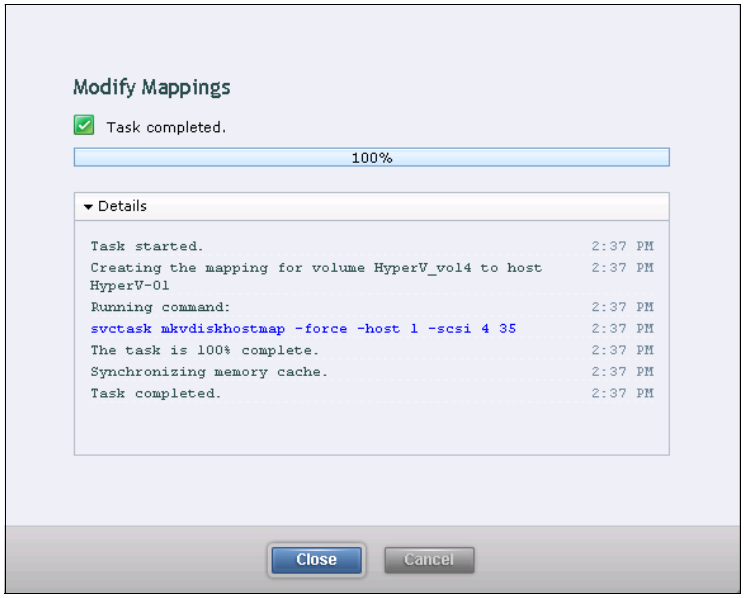

Figure 8-7 Modify Mappings task completion window

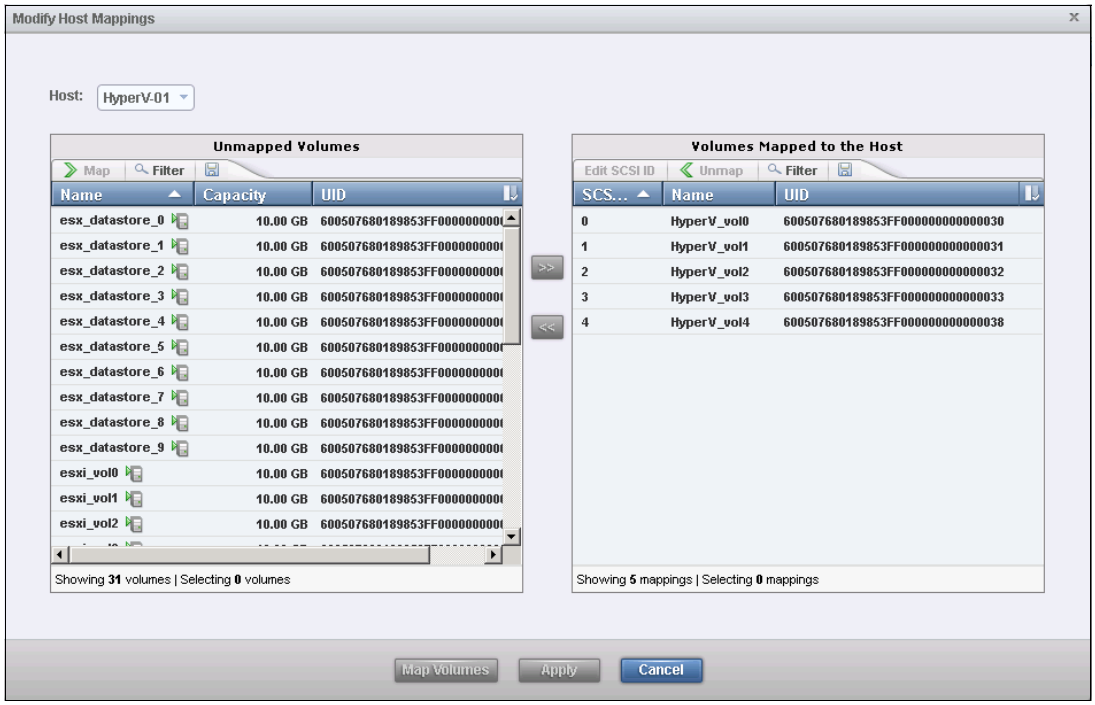

Figure 8-8 Modify Mappings: applied changes

You can now select another host in the Host drop-down menu (Figure 8-9) to modify the host settings for it or continue working with the one that is selected (Figure 8-8 on page 226).

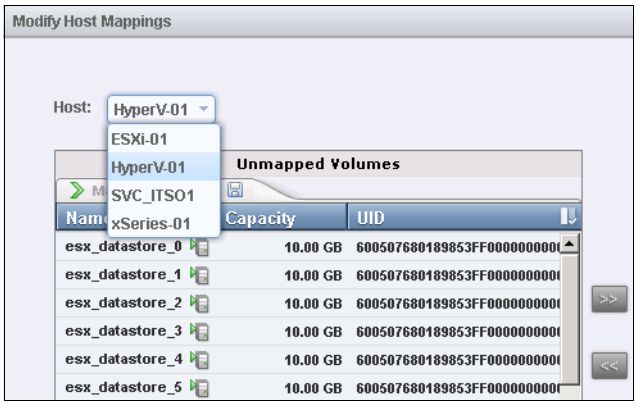

Figure 8-9 Modify another hosts mappings

Select the volume to be modified, again, and click the right arrow to move the volume to the list in the right pane. The changes are shown in yellow there. If you right-click the highlighted volume, you are able to change the SCSI ID, which is used for the host mapping (Figure 8-10). Select **Edit SCSI ID**

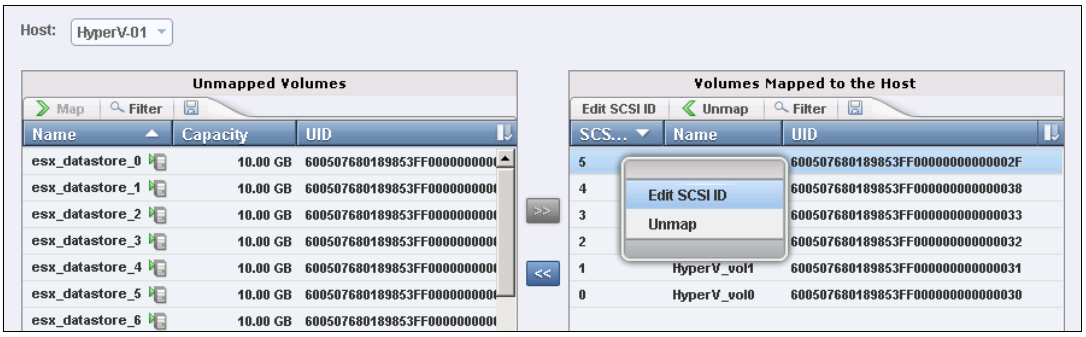

Figure 8-10 Edit the SCSI ID

Enter a SCSI ID and click **OK** to change the ID (Figure 8-11).

Figure 8-11 Enter a new SCSI ID

The changes are shown in the Modify Host Mappings window (Figure 8-12).

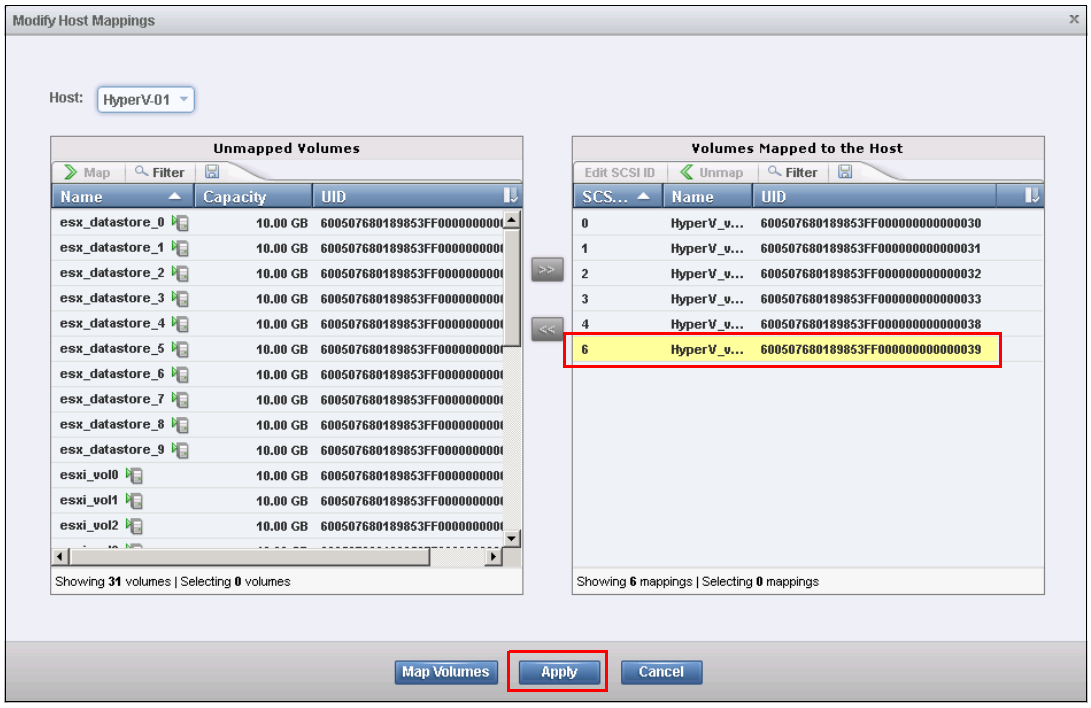

Figure 8-12 Modify Host Mappings

Click **Apply** to submit the changes. The resulting output is shown in Figure 8-13.

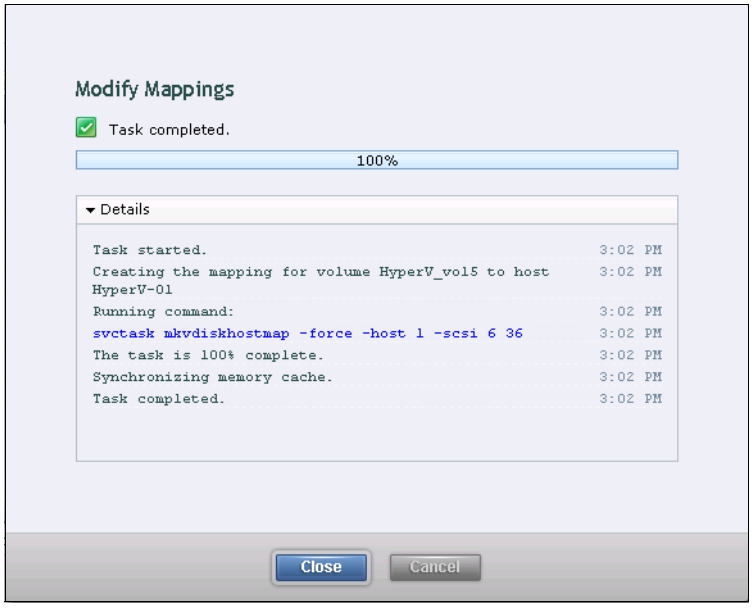

Figure 8-13 Modify Mappings task completion window

Note: You can change the volume SCSI ID only before it is mapped to a host. Changing it after is not possible unless you unmap it again.

If you want to remove a host mapping, the steps are similar, except that you select a volume in the right pane and click the left arrow (**<<**) button to remove the mapping (Figure 8-14).

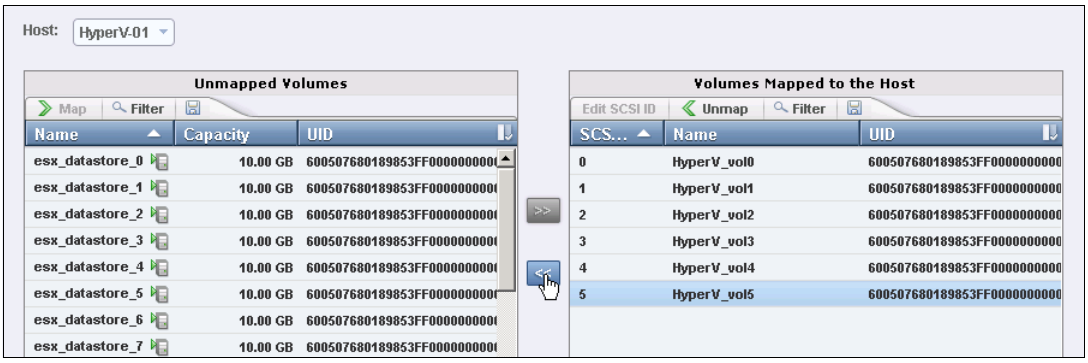

Figure 8-14 Modify Host Mappings: remove mapping

Figure 8-15 shows that the selected volume has been moved to the left pane for unmapping.

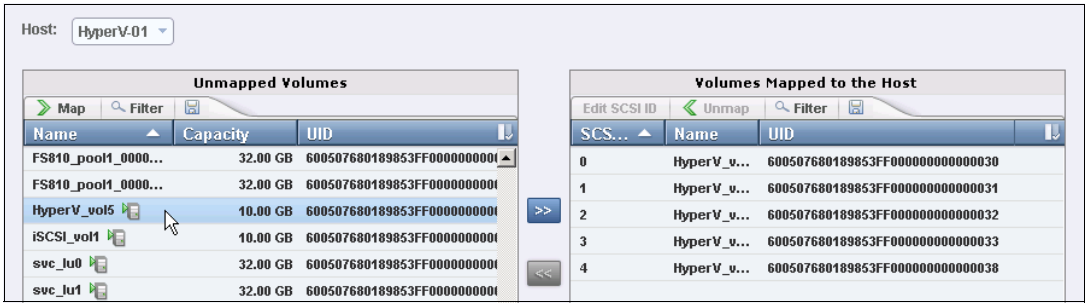

Figure 8-15 Modify Host Mappings: mapping removed

Click **Apply** to submit the changes to the system (Figure 8-16).

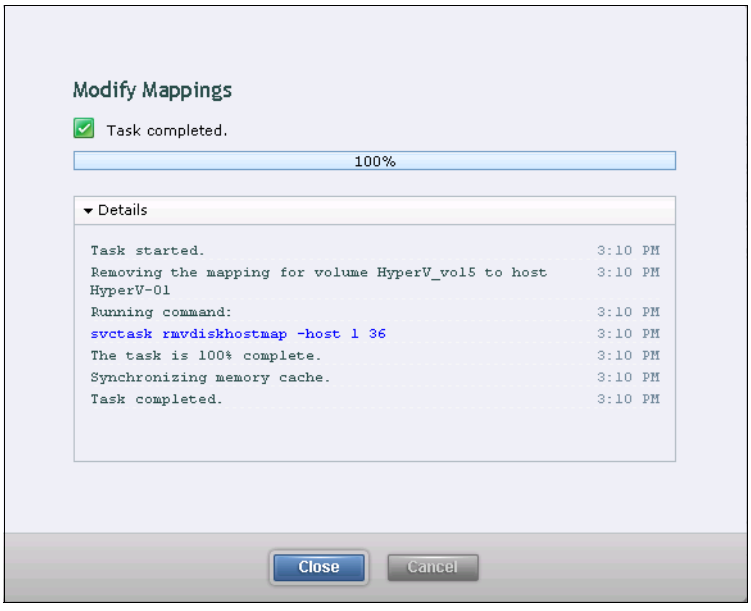

Figure 8-16 Modify Host Mappings: removal complete

After you are done with all host mapping modifications, click **Close** to return to the Modify Mappings window (Figure 8-5 on page 224).

8.1.2 Unmapping all volumes from a host

A host is able to access only those volumes on your IBM Storwize V7000 system that are mapped to it. If you want to remove access to all volumes for one host, regardless of how many volumes are mapped to it, you can do this task in one step. From the Hosts panel, select the host, click **Actions** → **Unmap All Volumes** to remove all access that this host has to its volumes (Figure 8-17).

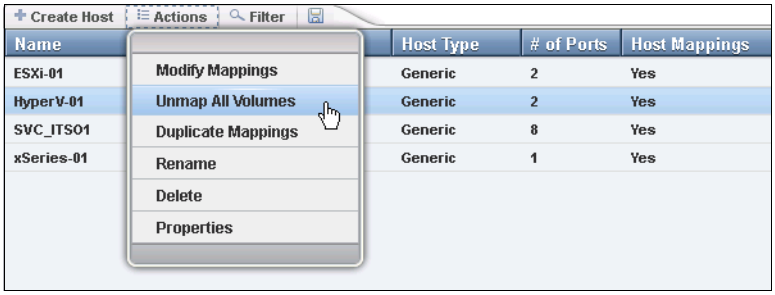

Figure 8-17 Unmap All Volumes action

You are prompted about the number of mappings you want to remove. Enter the number and click **Unmap** (Figure 8-18). In our example, we remove two mappings.

Unmapping: If you click **Unmap**, all access for this host to volumes that are controlled by IBM Storwize V7000 system is removed. Make sure that you run the required procedures on your host operating system before removing the volume mappings from your host object on IBM Storwize V7000.

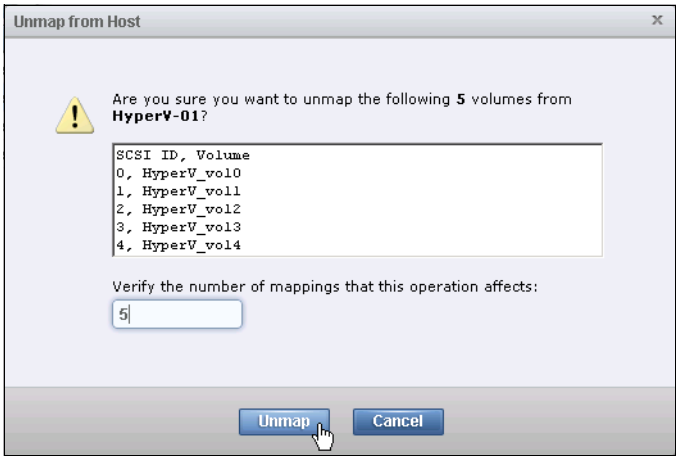

Figure 8-18 Enter the number of mappings to be removed

The changes are applied to the system (Figure 8-19). Click **Close** after you review the output.

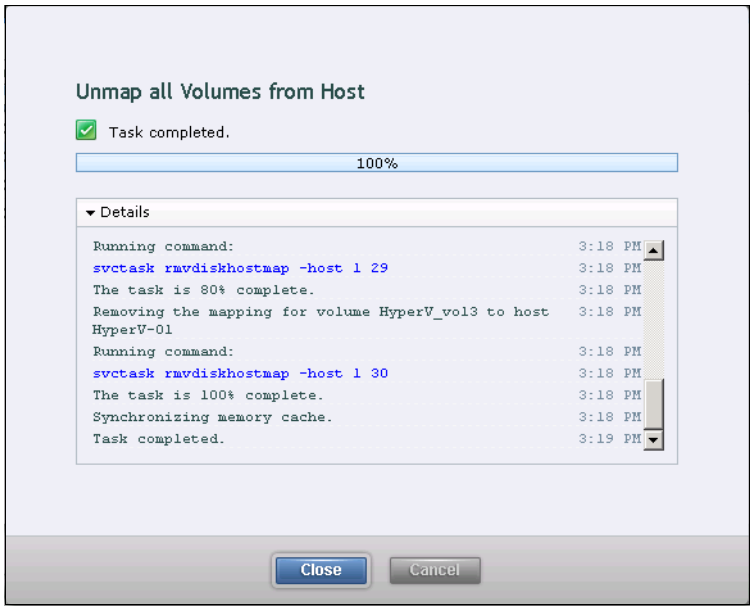

Figure 8-19 Unmap all Volumes from Host task completion window

Figure 8-20 shows that the selected host does not have host mappings anymore.

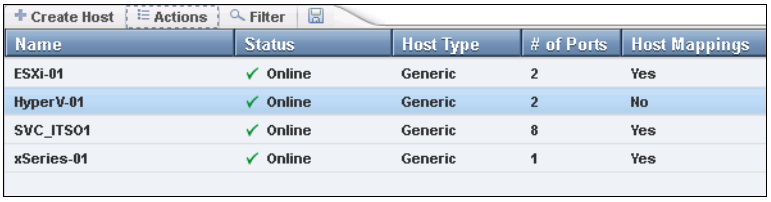

Figure 8-20 All mappings for host HyperV have been removed

8.1.3 Duplicating and Importing Mappings

Volumes that are assigned to a host can be mapped to another host object. You can do this for example when you add a new node to the host cluster and want to ensure that the new host node has access to same set of volumes as the source host.

Verify the mappings on the existing source host object. From the Hosts Actions menu (Figure 8-3 on page 223), right-click the host, select **Properties** and click the **Mapped Volumes** tab (Figure 8-21).

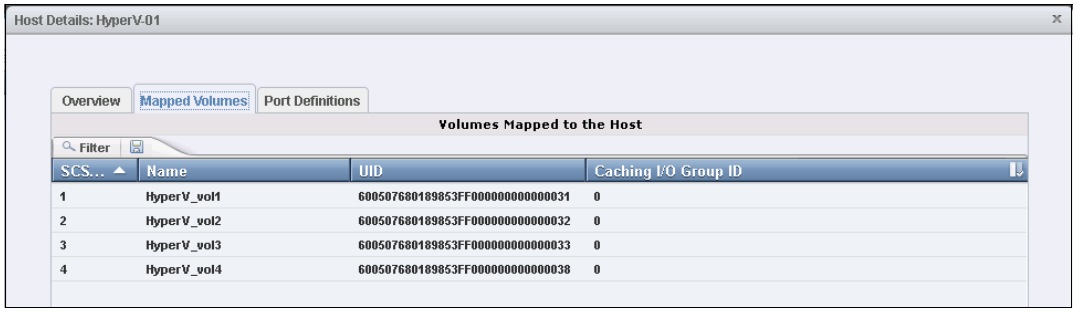

Figure 8-21 Host Details: Mapped Volumes

Select a host whose mappings you want to duplicate, and then click **Actions** → **Duplicate Mappings** (Figure 8-22).

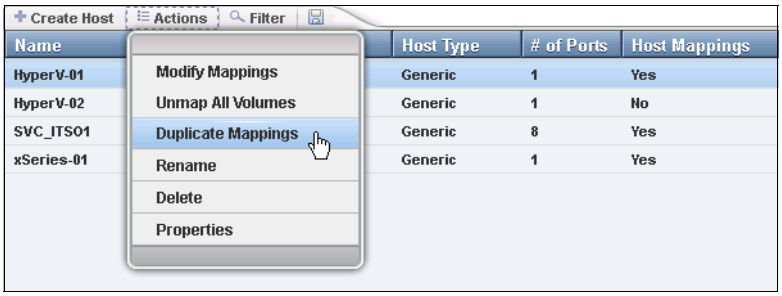

Figure 8-22 Duplicate Mappings action on source host object

The Duplicate Mappings window opens. Select a target host object to which you want to add all the existing source host mappings and click **Duplicate** (Figure on page 233).

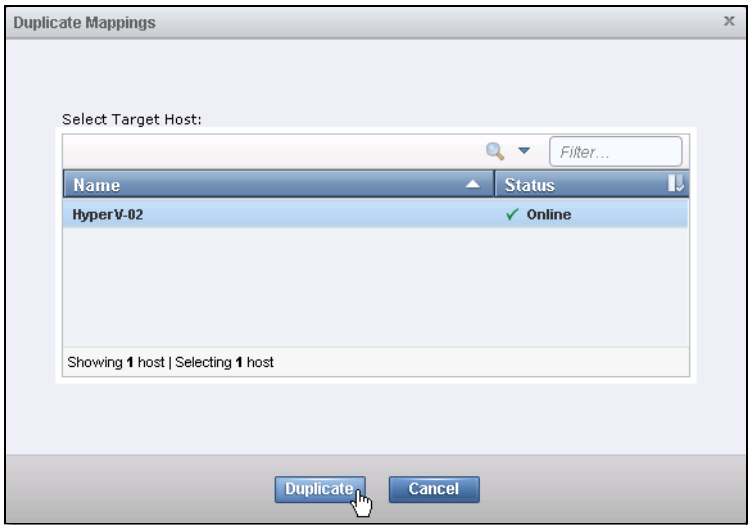

Figure 8-23 Duplicate mappings window

After the task completion is displayed (Figure 8-24), verify the new mappings on the target host object. From the Hosts Actions menu (Figure 8-3 on page 223), right-click the target host, select **Properties,** and click the **Mapped Volumes** tab (Figure 8-25).

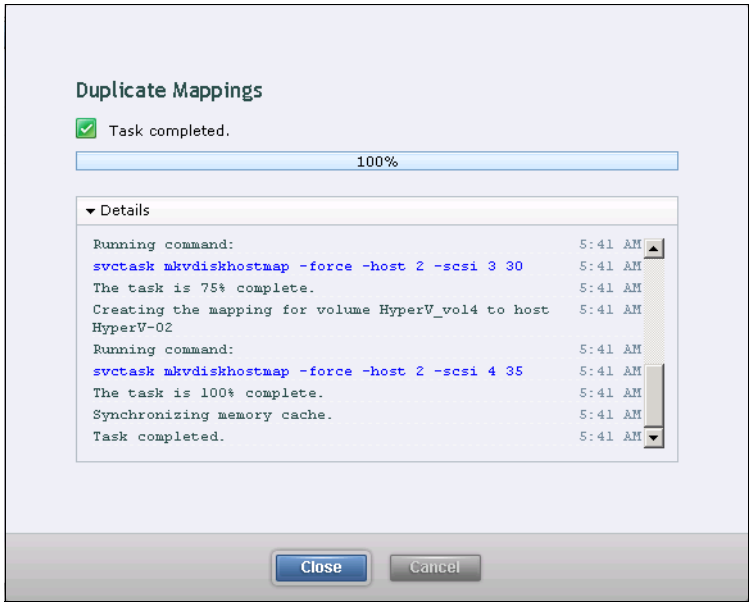

Figure 8-24 Duplicate Mappings task completion window

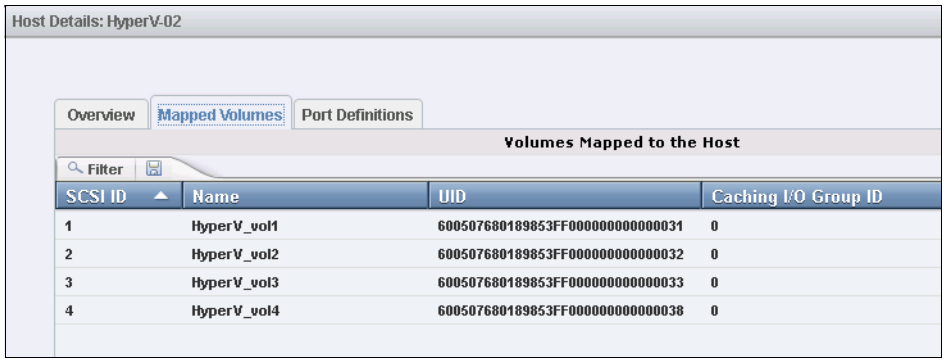

Figure 8-25 Host Details, new mappings on target host

You can perform the same action from the Actions menu of the target host object. You can import existing source host mappings provided the target host has no existing mappings defined.

Verify that no mappings are on the target host object. From the Hosts Actions menu (Figure 8-3 on page 223), right-click the host, select **Properties**, and then click the **Mapped Volumes** tab (Figure 8-26).

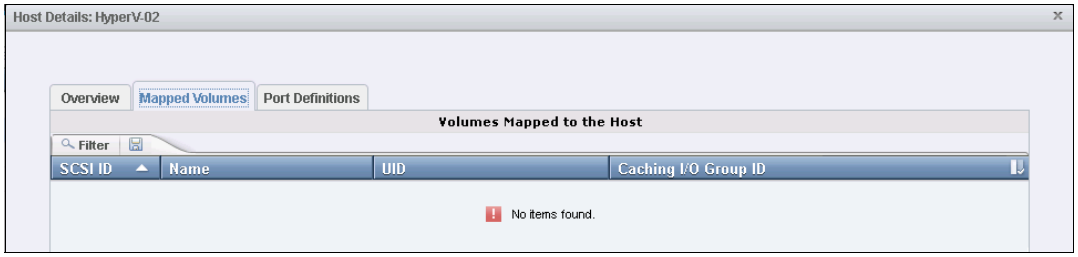

Figure 8-26 Import Mappings, verify there are no mappings on target host object

From the Hosts Actions menu, select a host to which you want to import existing mappings, and then click **Actions** → **Import Mappings** (Figure 8-27).

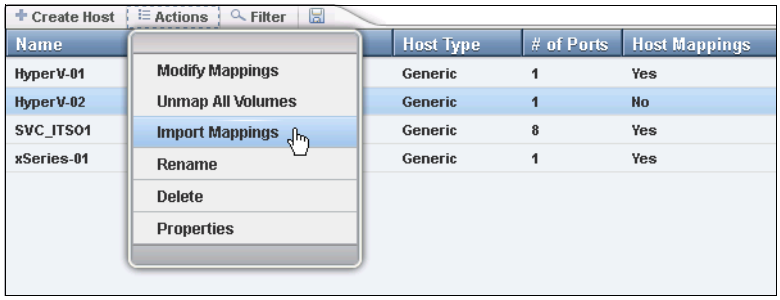

Figure 8-27 Import Mappings action on target host object

The Import Mappings window opens; select the source host object from the drop-down list and click **Import** (Figure 8-28).

Figure 8-28 Import Mappings window

After the task completion window is displayed (Figure 8-29), verify the new mappings on the target host object. From the Hosts Actions menu (Figure 8-3 on page 223), right-click the target host, select **Properties** and click the **Mapped Volumes** tab (Figure 8-30).

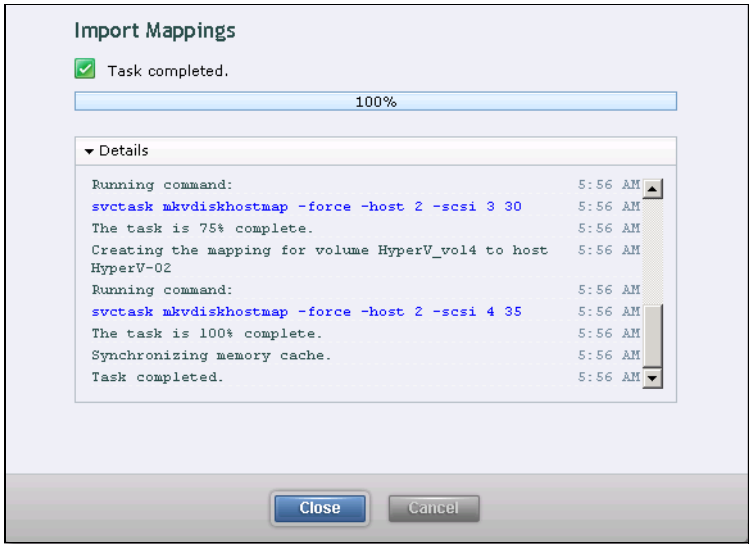

Figure 8-29 Import Mappings task completion window

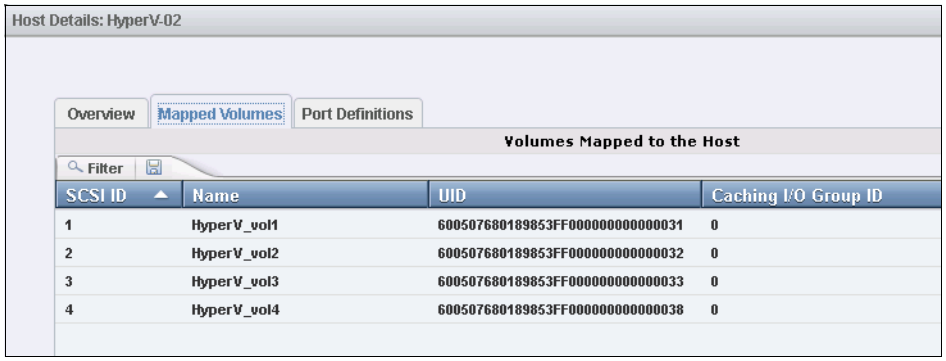

Figure 8-30 Host Details, new mappings on target host

8.1.4 Renaming

To rename a host, select it, and then right-click and select **Rename** (Figure 8-31).

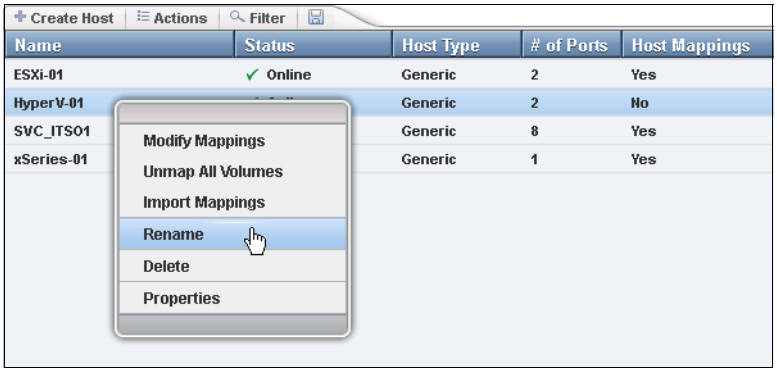

Figure 8-31 Rename a host

Enter a new name and click **Rename** (Figure 8-32). If you click **Reset**, the changes are reset to the original host name.

Figure 8-32 Rename Host window

After the changes are applied to the system, click **Close** (Figure 8-33).

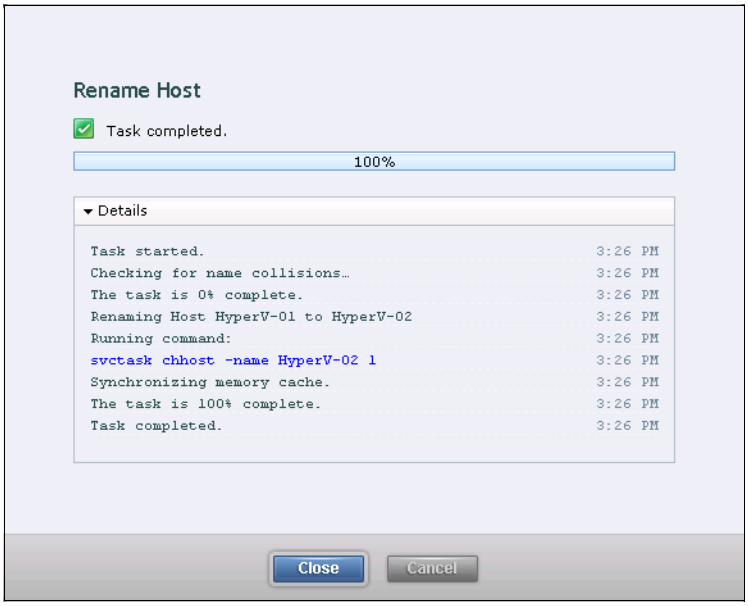

Figure 8-33 Rename Host task completion window

8.1.5 Deleting a host

To delete a host, from the IBM Storwize V7000 Hosts panel, select the host and click **Actions** → **Delete** (Figure 8-34).

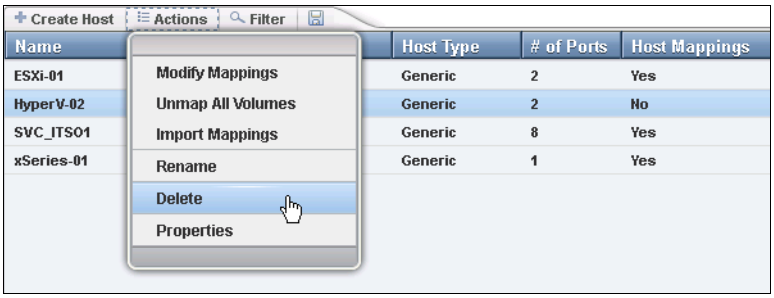

Figure 8-34 Delete a host

Confirm the number of hosts you want to delete and click **Delete** (Figure 8-35).

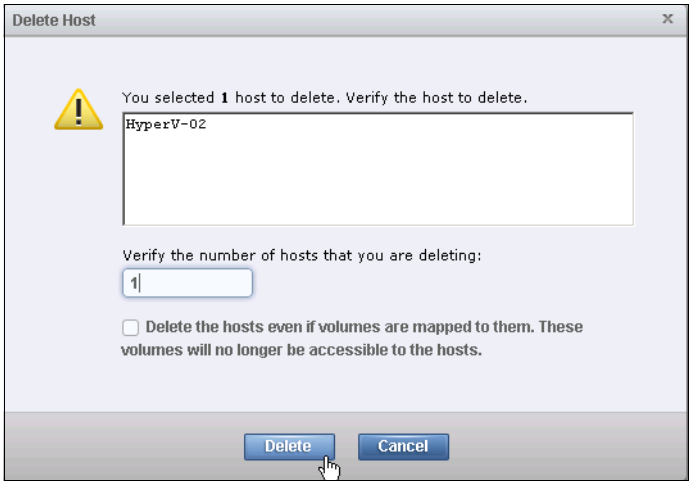

Figure 8-35 Confirm the deletion of the host

If you want to delete a host that has volumes mapped, you must force the deletion by selecting the check box in the lower part of the window. If you select this check box, the host is deleted, and it no longer has access to this system.

After the task is completed, click **Close** (Figure 8-36) to return to the mappings window.

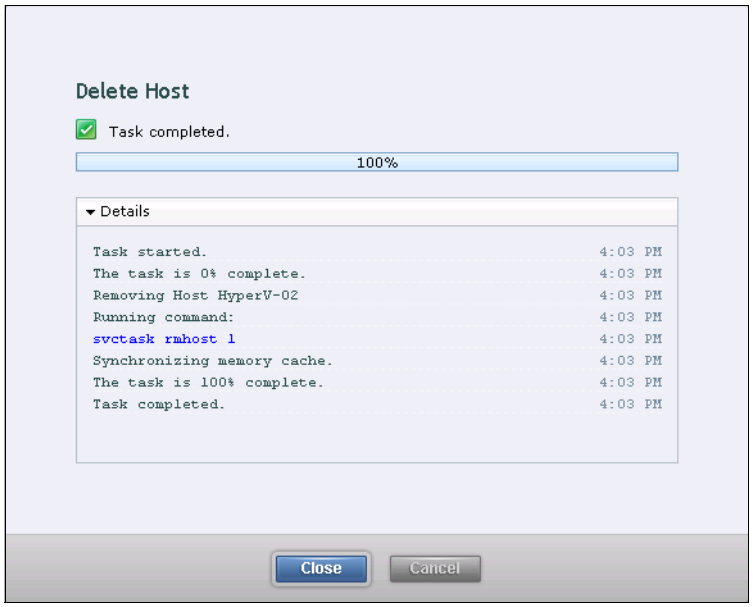

Figure 8-36 Delete host task completion window

8.1.6 Host properties

From the IBM Storwize V7000 Hosts panel, select the host and click **Actions** → **Properties** (Figure 8-37).

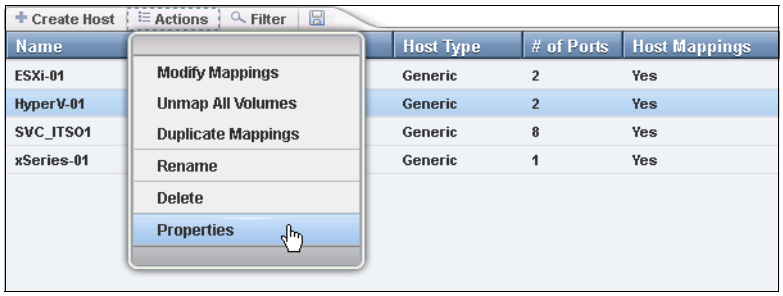

Figure 8-37 Host properties

The Host Details window opens (Figure 8-38).

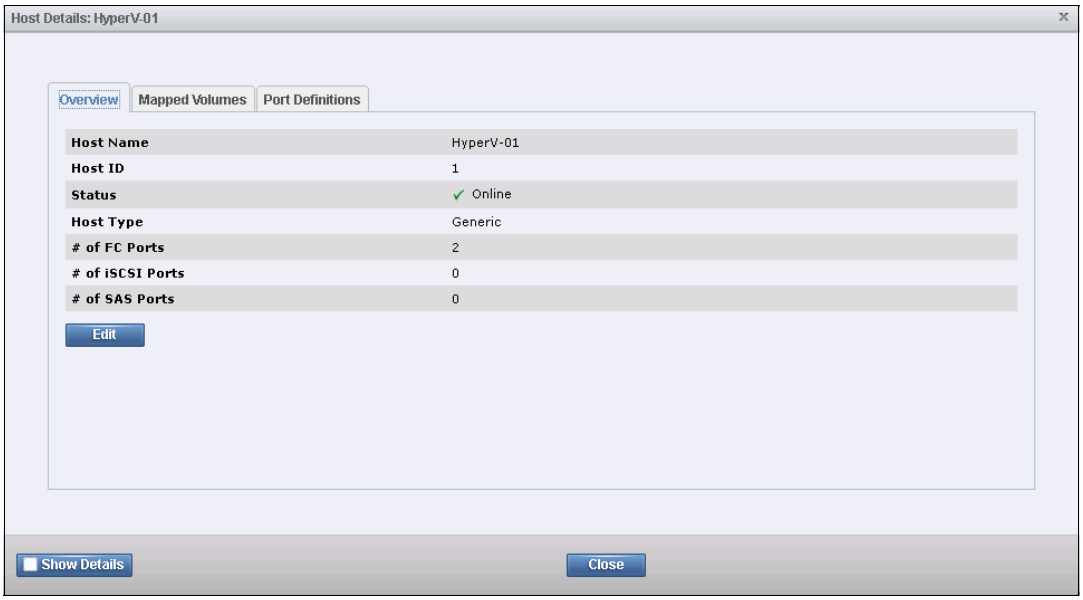

Figure 8-38 Host properties overview

The Host Details window shows an overview of your host properties. It has three tabs: Overview, Mapped Volumes, and Port Definitions.

The Overview tab is shown in Figure 8-38. Select the **Show Details** check box to see more information about the host (Figure 8-39).

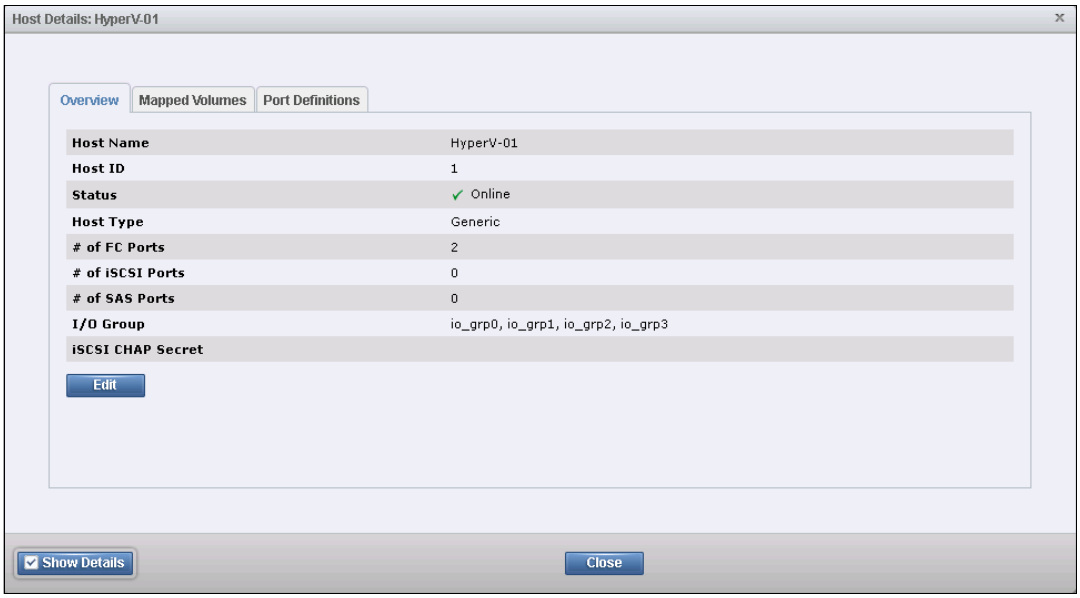

Figure 8-39 Host Properties: Show details

Click **Edit** if you want to change some host properties (Figure 8-40).

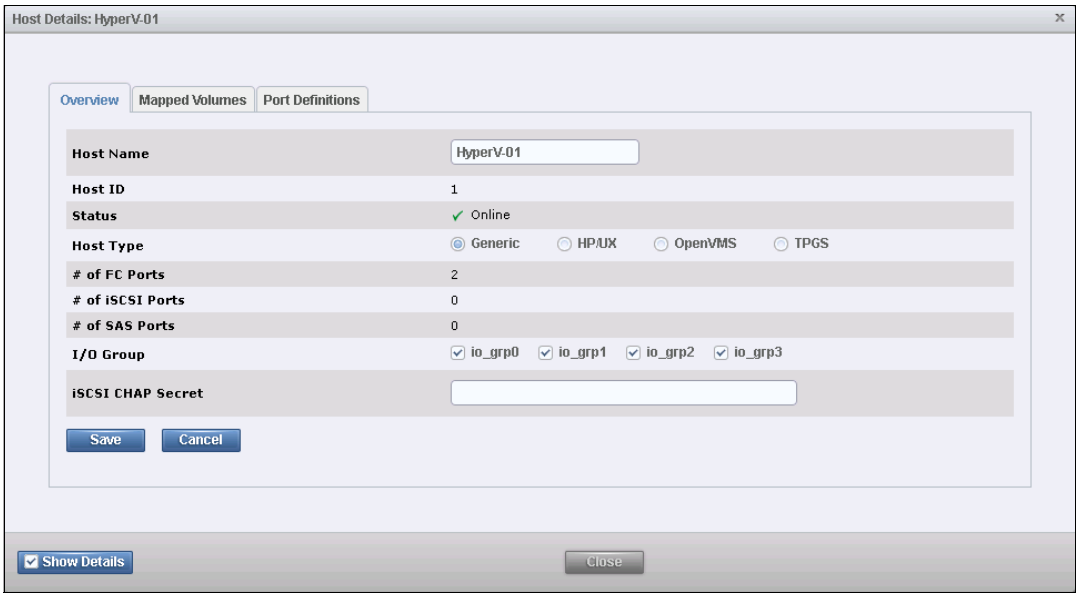

Figure 8-40 Edit host properties

In this window, you can modify the following items:

- Host Name: Change the host name.
- Host Type: Change this setting if you are going to attach HP/UX, OpenVMS, or TPGS hosts.
- iSCSI CHAP Secret: Enter or change the iSCSI CHAP secret for this host.

When you finish making changes (if required), click **Save** to apply them (Figure 8-40). The editing window closes.

The Mapped Volumes tab shows an overview of which volumes are currently mapped with which SCSI ID and UID to this host (Figure 8-41). The Show Details check box does not show any additional information.

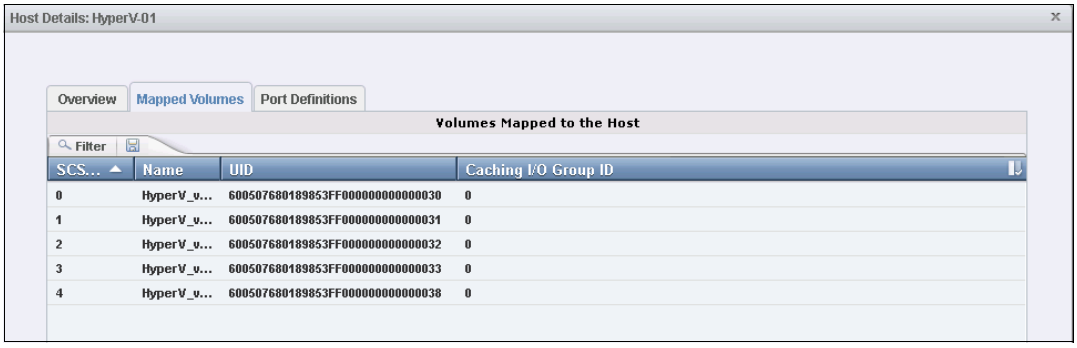

Figure 8-41 Mapped volumes tab

The Port Definitions tab shows the configured host ports of a host and gives you status information about them (Figure 8-42).

Figure 8-42 Port definitions

This window offers you the option to start Add and Delete Port actions, as described in 8.2, “Adding and deleting host ports” on page 241.

Click **Close** to close the Host Details window.

8.2 Adding and deleting host ports

To configure host ports, from the IBM Storwize V7000 select **Ports by Host** to open the associated panel (Figure 8-43).

Figure 8-43 Ports by Host

The left pane lists all the hosts; the function icons indicate whether the host is Fibre Channel (*orange* cable) or iSCSI (*blue* cable). The properties of the highlighted host are shown in the right pane. If you click **New Host**, the wizard starts; it is described in 4.3, “Creating hosts using the GUI” on page 121. If you click **Actions** (Figure 8-44), the tasks described in 8.1.1, “Modifying Mappings menu” on page 224 may be started from this location.

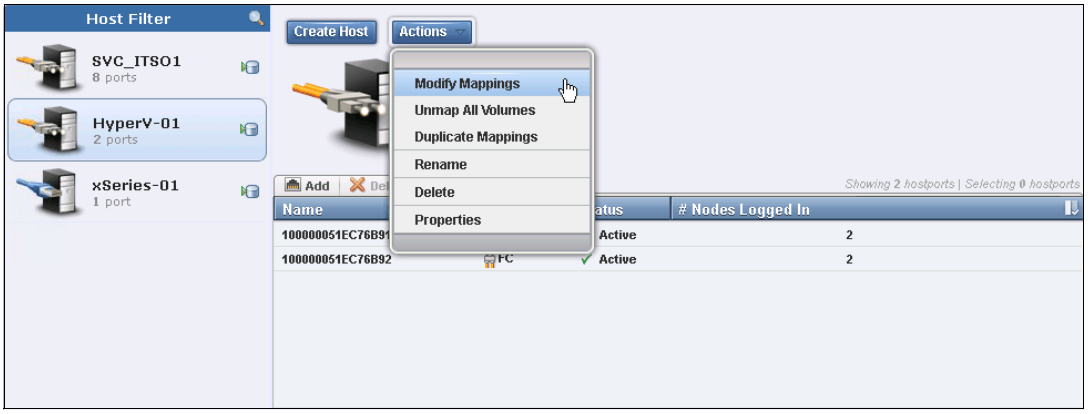

Figure 8-44 Ports by Host actions

8.2.1 Adding a Fibre Channel or iSCSI host port

To add a host port, highlight the host, click **Add** (Figure 8-45), and select Fibre Channel Port (see “Adding a Fibre Channel port”) or an iSCSI Port (see “Adding an iSCSI host port”).

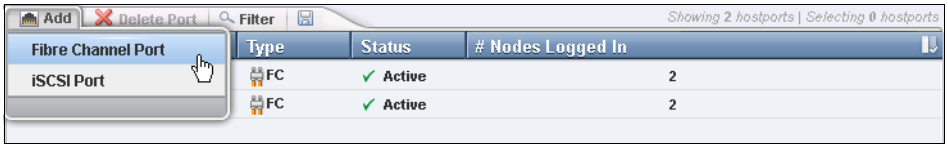

Figure 8-45 Add host ports

Adding a Fibre Channel port

Click **Fibre Channel Port** (Figure 8-45). The Add Fibre Channel Ports window opens (Figure 8-46).

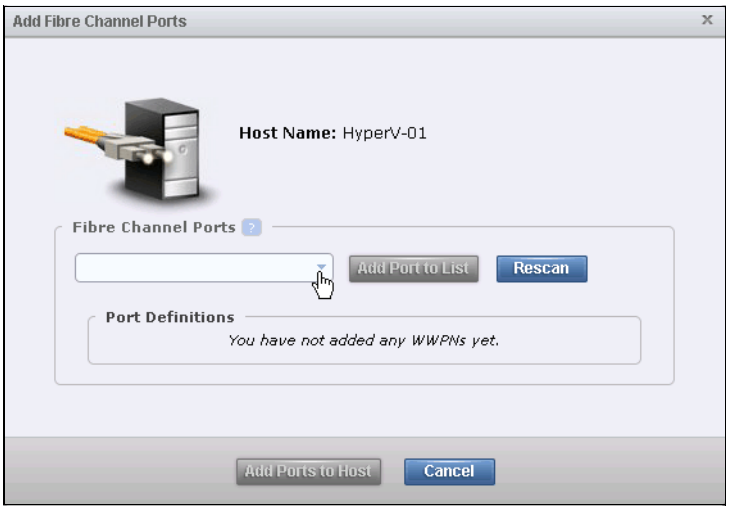

Figure 8-46 Add Fibre Channel Ports window

If you click the drop-down menu, a list of all known Fibre Channel host ports (Figure 8-47) is displayed. If the WWPN of your host is not available in the menu, check your SAN zoning and rescan the SAN from the host. Then, click **Rescan**; the new port is now available in the menu.

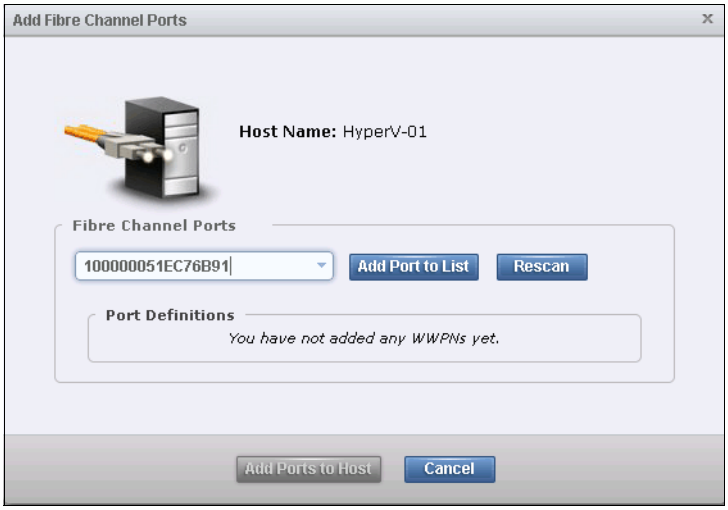

Figure 8-47 Add Fibre Channel Ports: Known WWPNs

Select the WWPN you want to add and click **Add Port to List** (Figure 8-48).

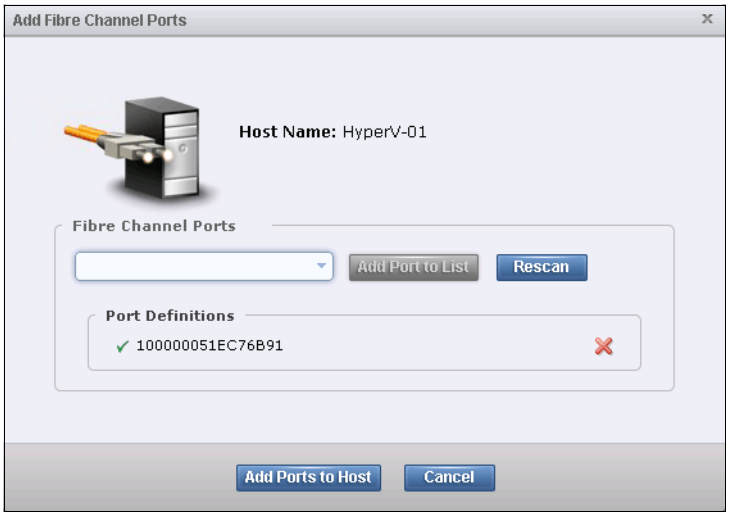

Figure 8-48 Add a port to list

You may repeat this step to add more ports to a host.

If you want to add an offline port (if the WWPN of your host is not available in the drop-down menu), manually enter the WWPN of the port into the Fibre Channel Ports field (Figure 8-49) and click **Add Port to List**.

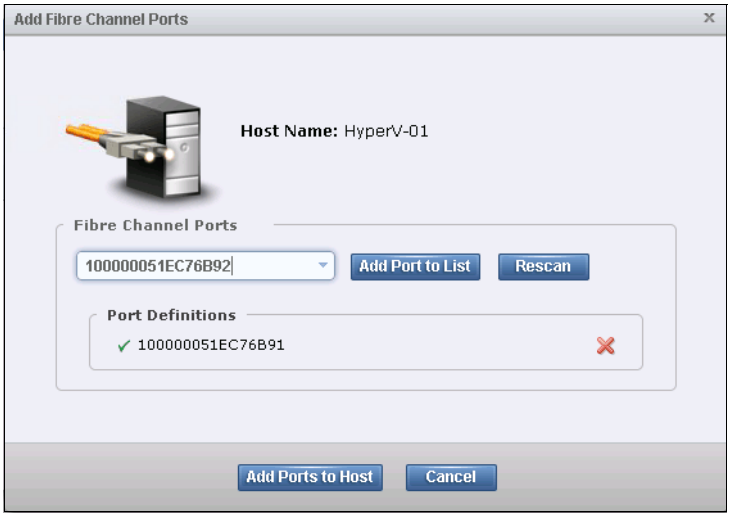

Figure 8-49 Add an offline port

The port is unverified (Figure 8-50) because it is not logged on to the IBM Storwize V7000. The first time it logs on, its state is automatically changed to online and the mapping is applied to this port.

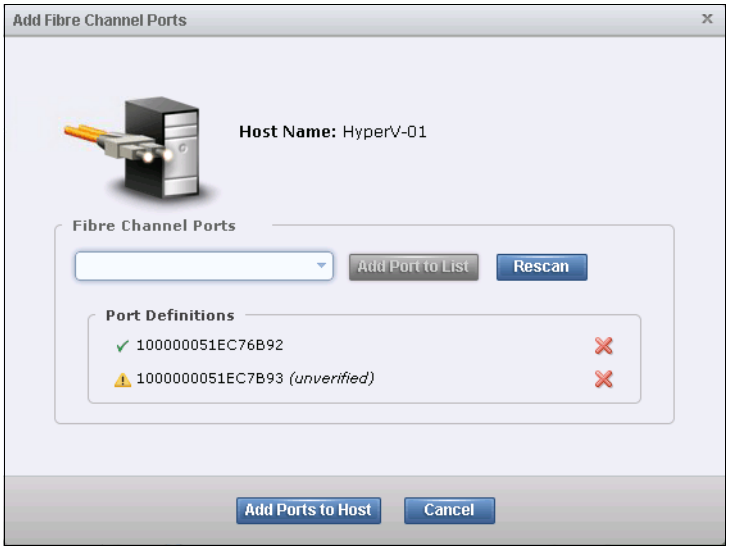

Figure 8-50 Unverified port

To remove a port from the list, click the red X next to the port (Figure 8-51). In this example, we delete the manually added FC port so only the detected port remains.

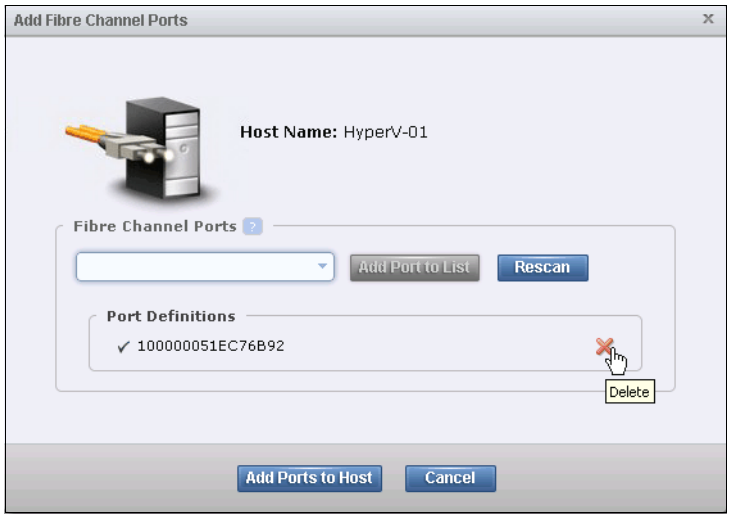

Figure 8-51 Remove a port from a list

Click **Add Ports to Host** and the changes are applied (Figure 8-52).

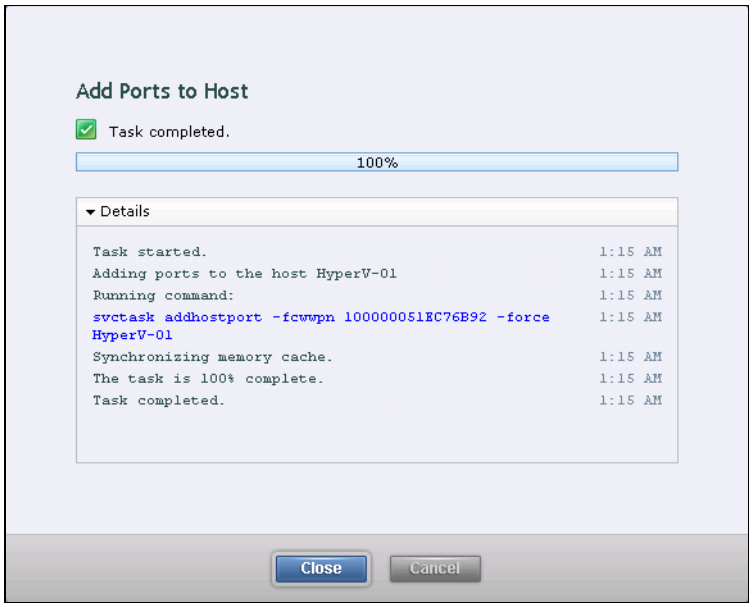

Figure 8-52 Add Ports to Host task completion window

Click **Close** to return to the Ports to Host window.

Adding an iSCSI host port

To add an iSCSI host port, click **iSCSI Port** (Figure 8-45 on page 242). The Add iSCSI Port window opens (Figure 8-53).

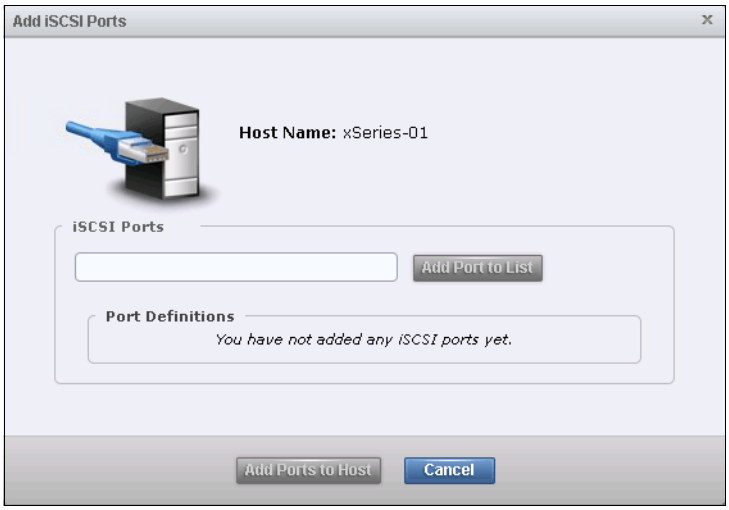

Figure 8-53 Add iSCSI host ports

Enter the initiator name of your host (Figure 8-54) and click **Add Port to List**.

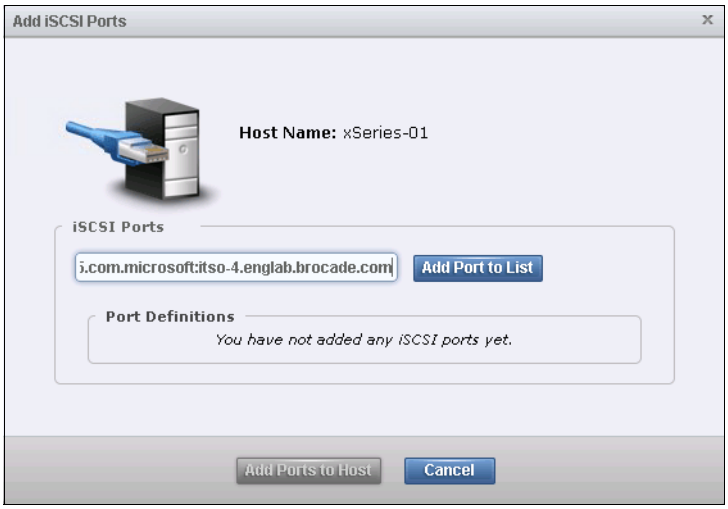

Figure 8-54 Enter the initiator name

Click **Add Ports to Host** (Figure 8-55).

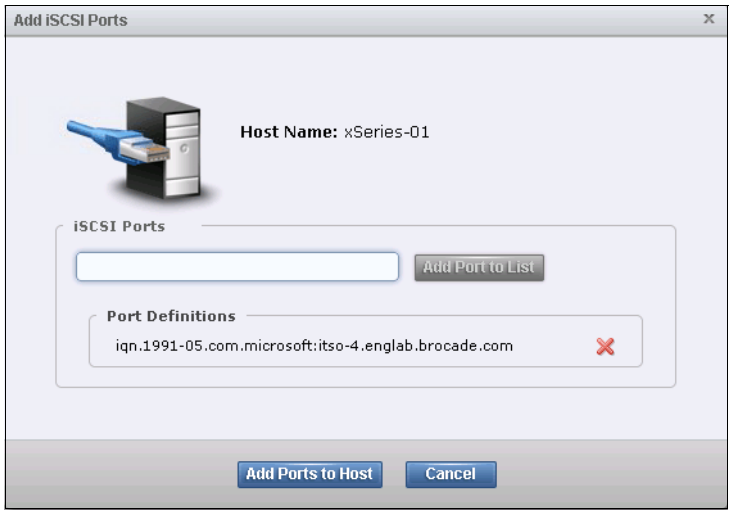

Figure 8-55 Add ports to the ports definitions

The tasks are completed and changes to the system are applied (Figure 8-56). Click **Close** to return to the Ports by Host window.

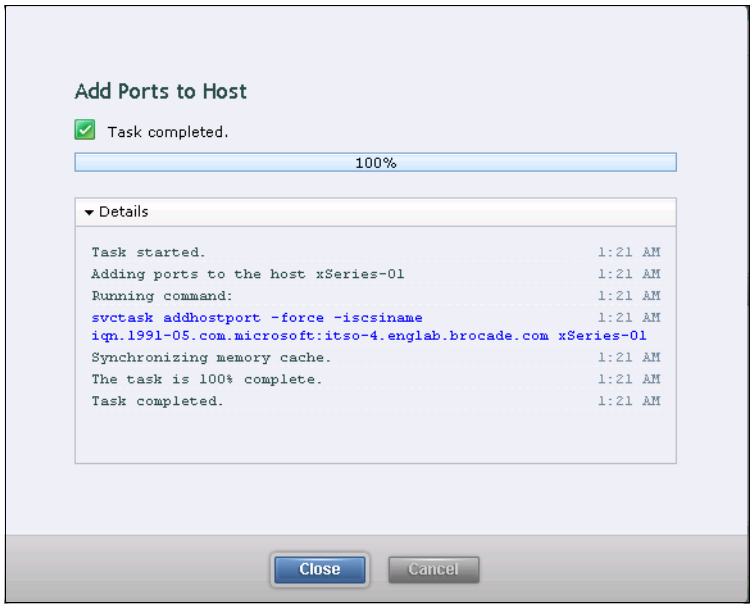

Figure 8-56 Add Ports to Host task completion window

8.2.2 Deleting a host port

To delete a host port, highlight it and click **Delete Port** (Figure 8-57).

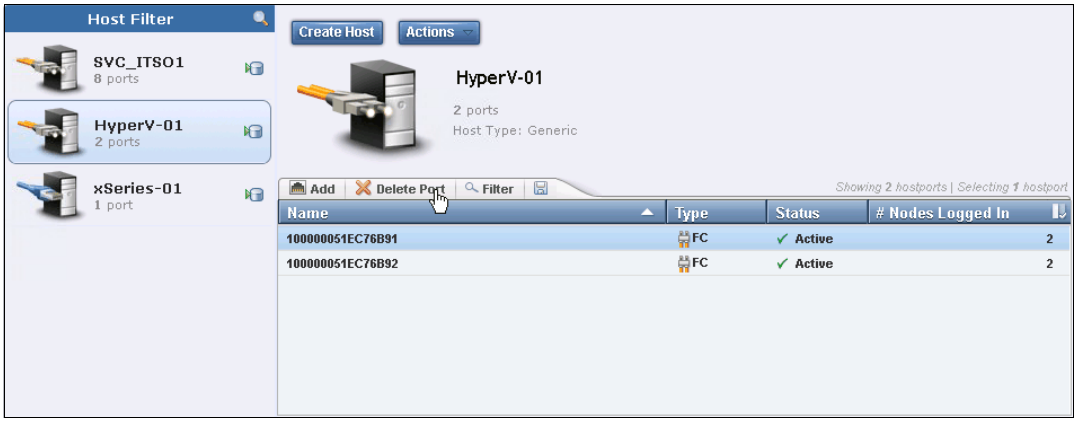

Figure 8-57 Delete host port

You may also press the Ctrl key while you select several host ports to delete (Figure 8-58).

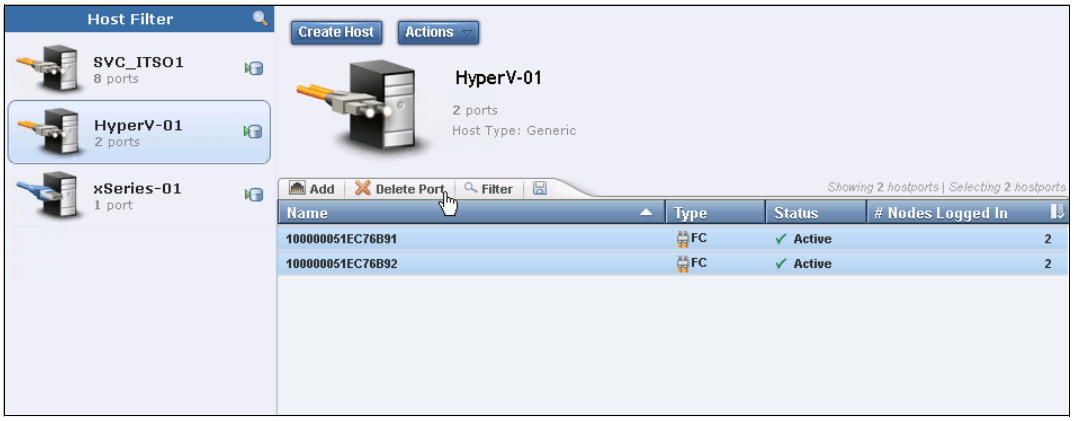

Figure 8-58 Delete several host ports

Click **Delete** and enter the number of host ports you want to remove (Figure 8-59).

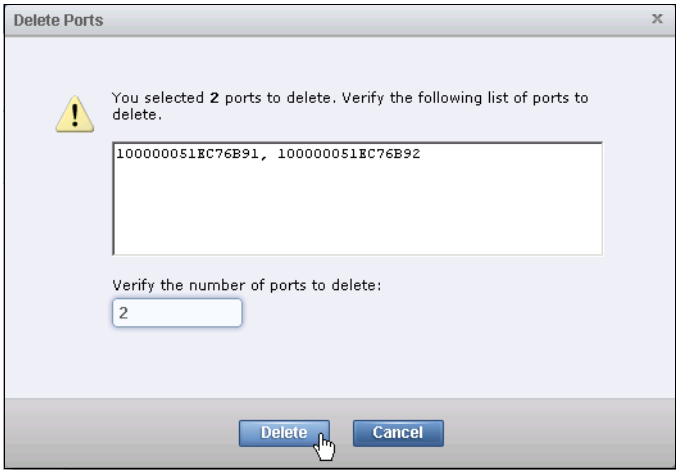

Figure 8-59 Enter the number of host ports to delete

Click **Delete** to apply the changes to the system (Figure 8-60).

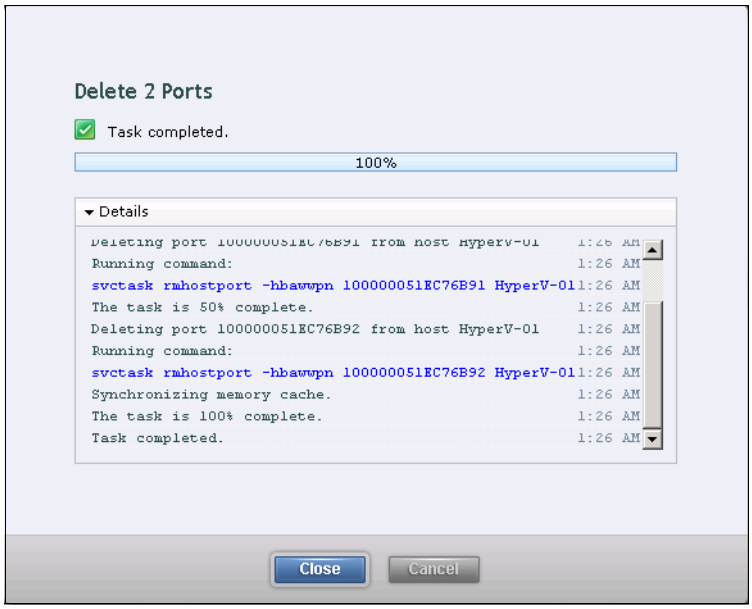

Figure 8-60 Delete Host Ports task completion window

Click **Close** to return to the Host by Ports window.

8.3 Host mappings overview

Select **Host Mappings** (Figure 8-1 on page 222) to open the host mappings panel (Figure 8-61).

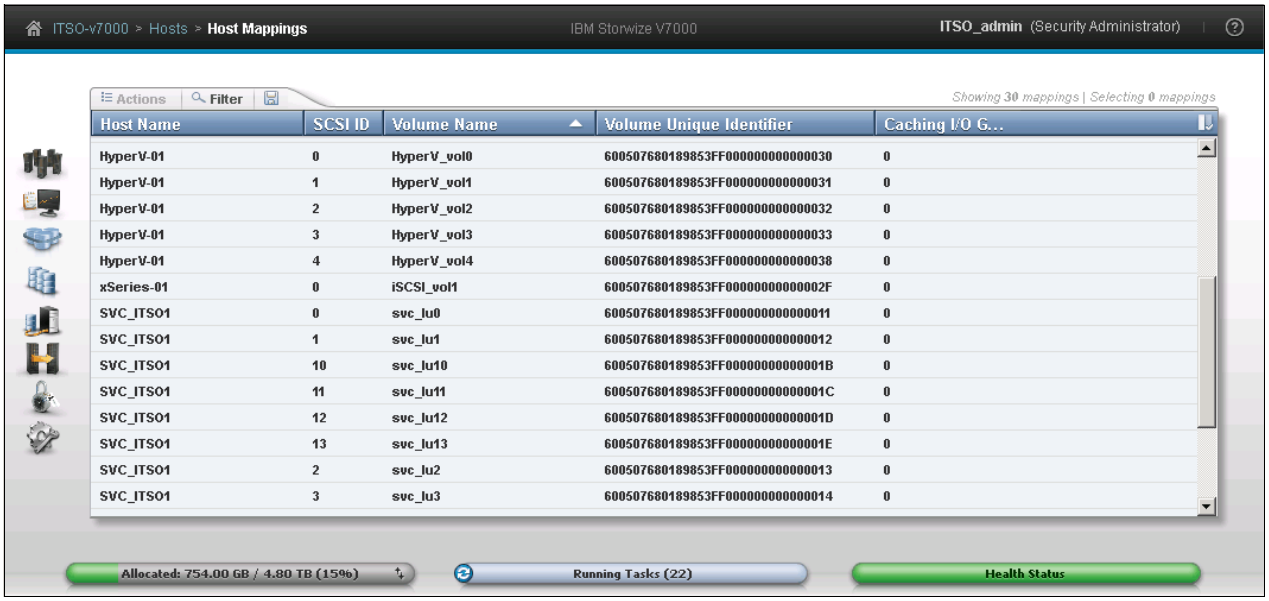

Figure 8-61 Host mappings

This panel lists all hosts and volumes. Our example shows that the host *HyperV-01* has five mapped volumes, and their associated SCSI ID, Volume Name, and Volume Unique Identifier (UID). If you have more than one caching I/O group, you also see which volume is handled by which I/O group.

If you select one line and click **Actions** (Figure 8-62), the following tasks are available:

- Unmapping a volume
- Properties (Host)
- Properties (Volume)

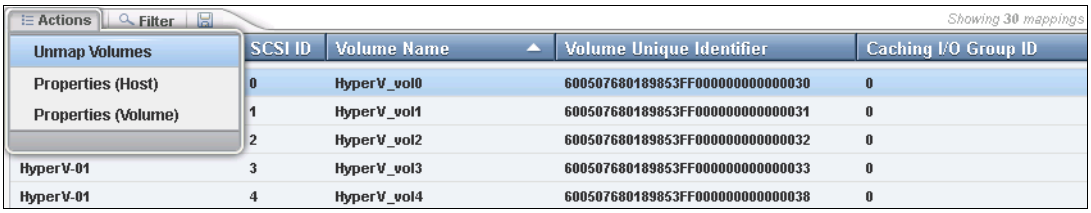

Figure 8-62 Host Mappings Actions drop-down menu

8.3.1 Unmapping a volume

Select one or more lines, click **Unmap Volumes**, enter the number of entries to remove (Figure 8-63), and then click **Unmap**.

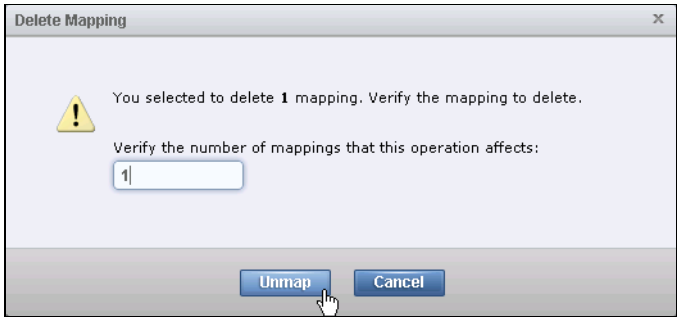

Figure 8-63 Unmap selected volumes

This action removes the mappings for all selected entries (Figure 8-64).

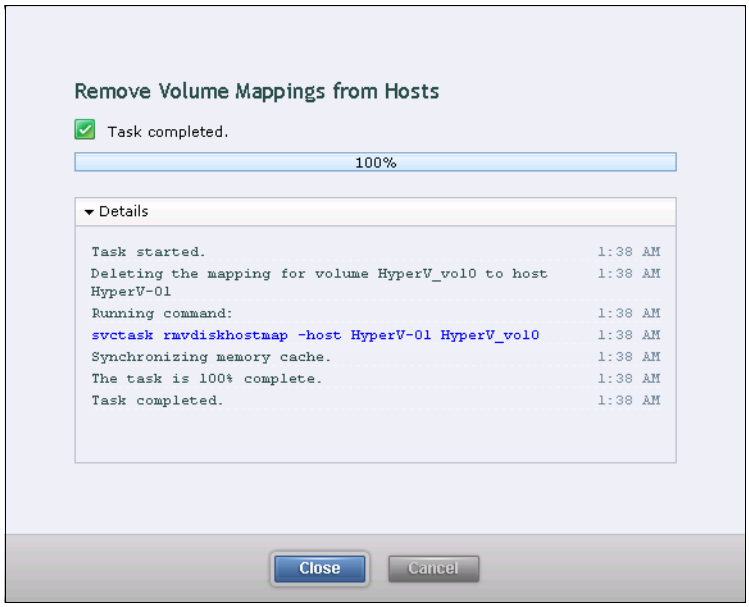

Figure 8-64 Remove Volume Mappings from Hosts task completion window

8.3.2 Properties (Host)

Selecting an entry and clicking **Properties (Host)**, as shown in Figure 8-62 on page 250, opens the host properties window. The contents of this window are described in 8.1.6, “Host properties” on page 238.

8.3.3 Properties (Volume)

Selecting an entry and clicking **Properties (Volume)**, as shown in Figure 8-62 on page 250, opens the volume properties view. The contents of this window are described in 5.1, “Provisioning storage from IBM Storwize V7000 to hosts” on page 132.

8.4 Advanced volume administration

This section covers volume administration, such as volume modification and the migration or creation of new volume copies. Basic volume creation is covered in 5.1, “Provisioning storage from IBM Storwize V7000 to hosts” on page 132. In this section, we assume that you already have created some volumes in your IBM Storwize V7000 and that your are familiar with generic, thin-provisioned, and mirrored volumes.

Figure 8-65 shows that three volumes options are available to administer advanced features:

- Volumes (8.4.1, “Advanced volume functions” on page 253 and 8.6, “Advanced volume copy functions” on page 279)
- Volumes by Pool (8.7, “Volumes by storage pool” on page 286)
- Volumes by Host (8.8, “Volumes by Host” on page 292)

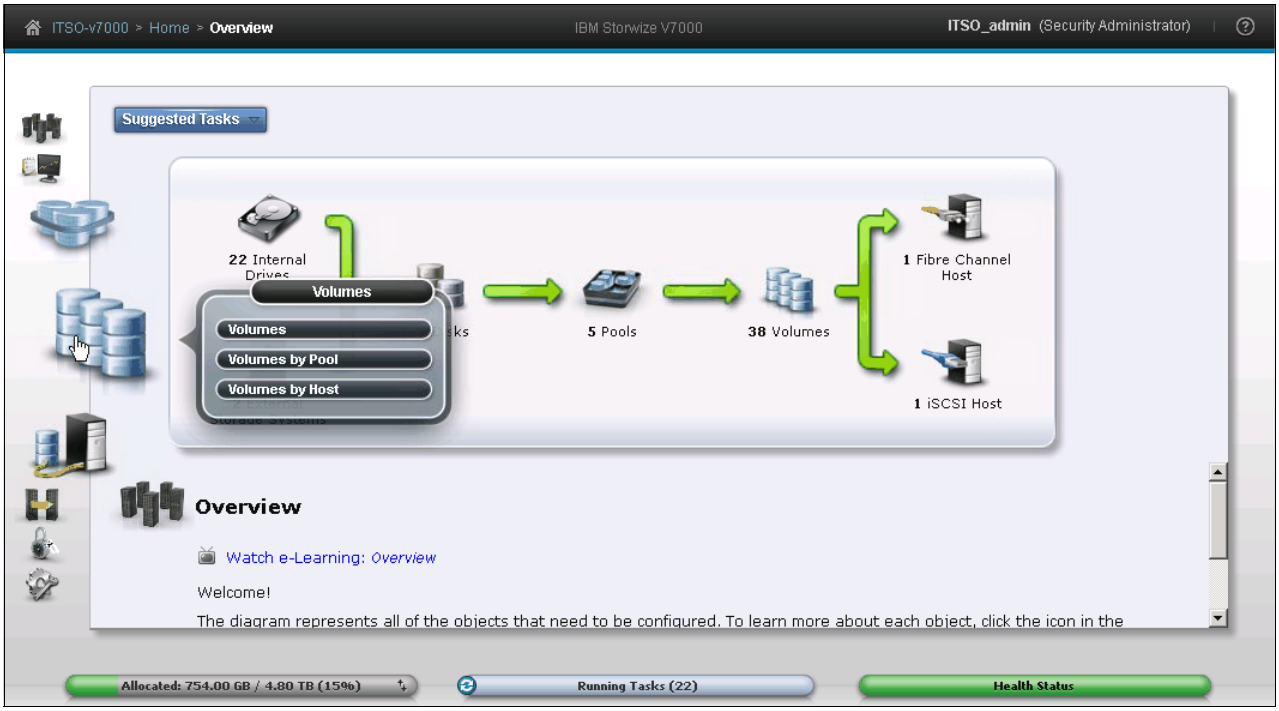

Figure 8-65 Volumes menu

8.4.1 Advanced volume functions

Click **Volumes** (Figure 8-65 on page 252); the Volumes panel opens (Figure 8-66).

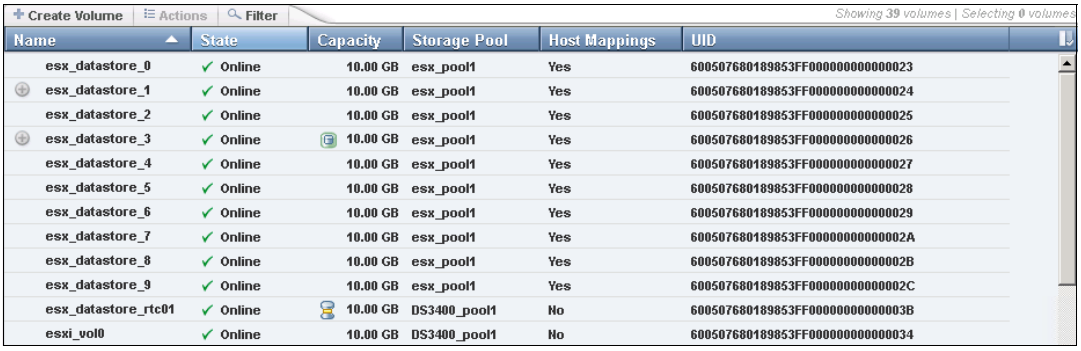

Figure 8-66 Volumes panel

This window lists all configured volumes on the system and provides the following information:

- Name: Shows the name of the volume. A plus sign (⊕) next to the name means that several copies of this volume exist. Click it to expand the view and list the copies (Figure 8-67).
- State: Gives you information about the volume state, which can be online, offline, or degraded.
- Capacity: The capacity that is presented to the host is listed here. A *blue* volume icon listed next to the capacity means that this volume is thin-provisioned, and that the listed capacity is the virtual capacity, which might be less than the real capacity on the system. Similarly, the *compressed* volume icon indicates the data on the volume are compressed by Real-time Compression feature.
- Storage Pool: Shows in which Storage Pool the volume is stored. If you have several volume copies, it shows you the pool of the primary copy.
- Host Mappings: Indicates whether the volume is mapped to at least one host or not.
- UID: This is the volume unique identifier.

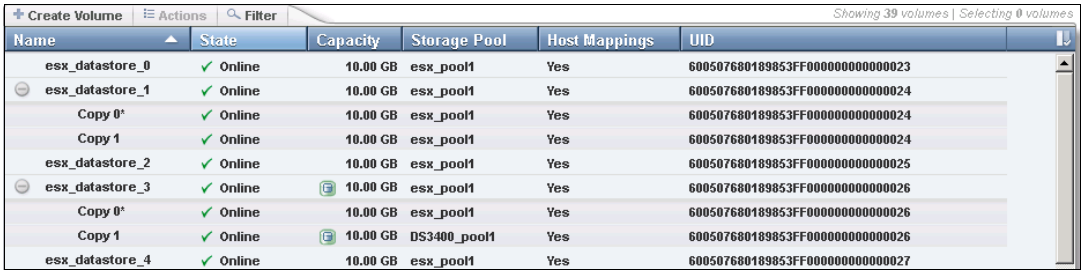

Figure 8-67 Volume copies

To create a volume, click **New Volume** and complete the steps described in 5.1, “Provisioning storage from IBM Storwize V7000 to hosts” on page 132.

Select a volume and click **Actions** to see the available actions for a volume (Figure 8-68).

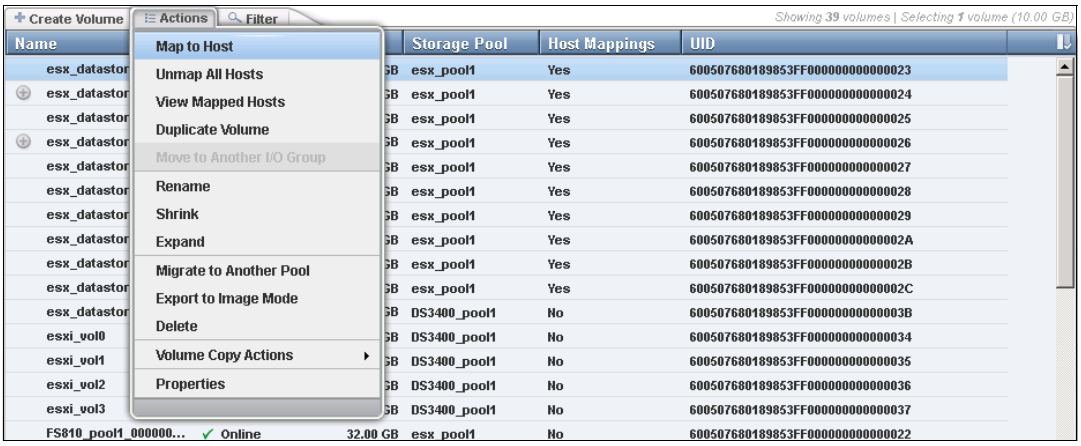

Figure 8-68 Volume Actions menu

You may also right-click a volume and select the actions in the menu (Figure 8-69).

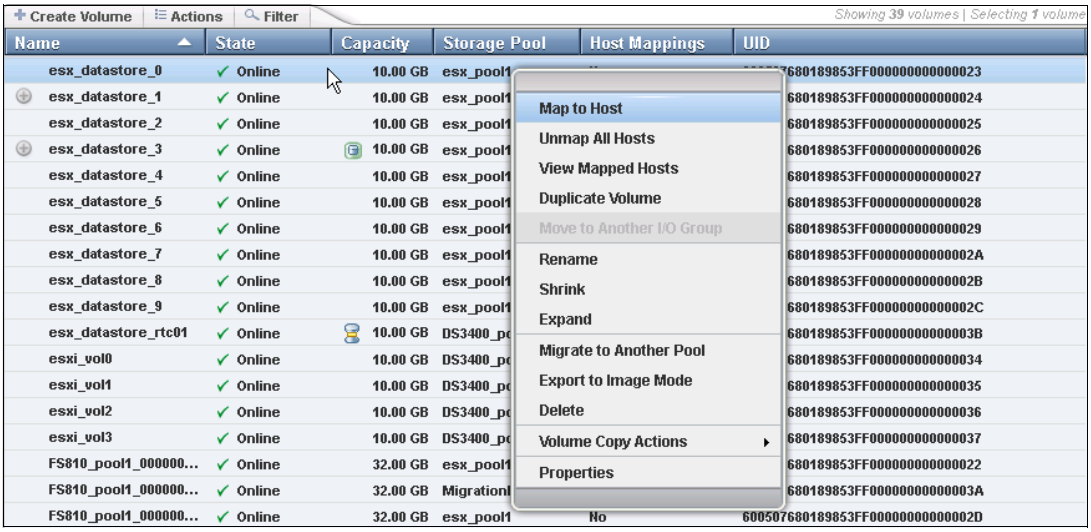

Figure 8-69 More volume actions

The following volume options are available:

- Map to Host (“Mapping a volume to a host” on page 255)
- Unmap All Hosts (“Unmapping volumes from all hosts” on page 256)
- View Mapped Host (“Viewing a host mapped to a volume” on page 257)
- Duplicate Volume (“Duplicating a volume” on page 258)
- Move to Another I/O Group (only applicable to multiple I/O group systems
- Rename (“Renaming a volume” on page 259)
- Shrink (“Shrinking a volume” on page 260)
- Expand (“Expanding a volume” on page 261)
- Migrate to Another Pool (“Migrating a volume to another storage pool” on page 261)
- Export to image Mode (“Exporting to an image mode volume” on page 264)
- Delete (“Deleting a volume” on page 266)
- Properties (“Volume properties” on page 267)

Depending on which volume preset you view, other actions might available for Volume Copy:

- Add Mirror Copy: Available only for generic volumes (see 8.5.4, “Adding a mirrored volume copy” on page 273)
- Thin Provisioned: Available only for thin-provisioned volumes (see 8.5.5, “Editing thin-provisioned or compressed volume properties” on page 275)
- Compressed: Available only for compressed volumes

8.4.2 Mapping a volume to a host

To map a volume to a host, select **Map to Host** from the menu shown in Figure 8-68 on page 254. Select the host to which you want to map the volume (Figure 8-70).

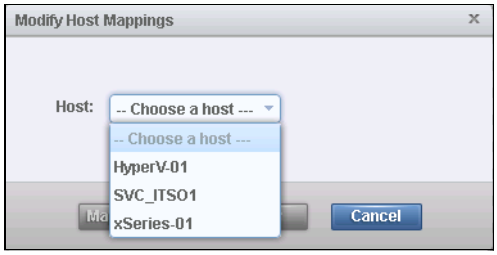

Figure 8-70 Modify Host Mappings, select a host

The Modify Mappings window opens. The Host field indicates your selected host; the yellow volume in the list at the right is the selected volume that will be mapped (Figure 8-71). Click **Map Volumes** to apply the changes to the system.

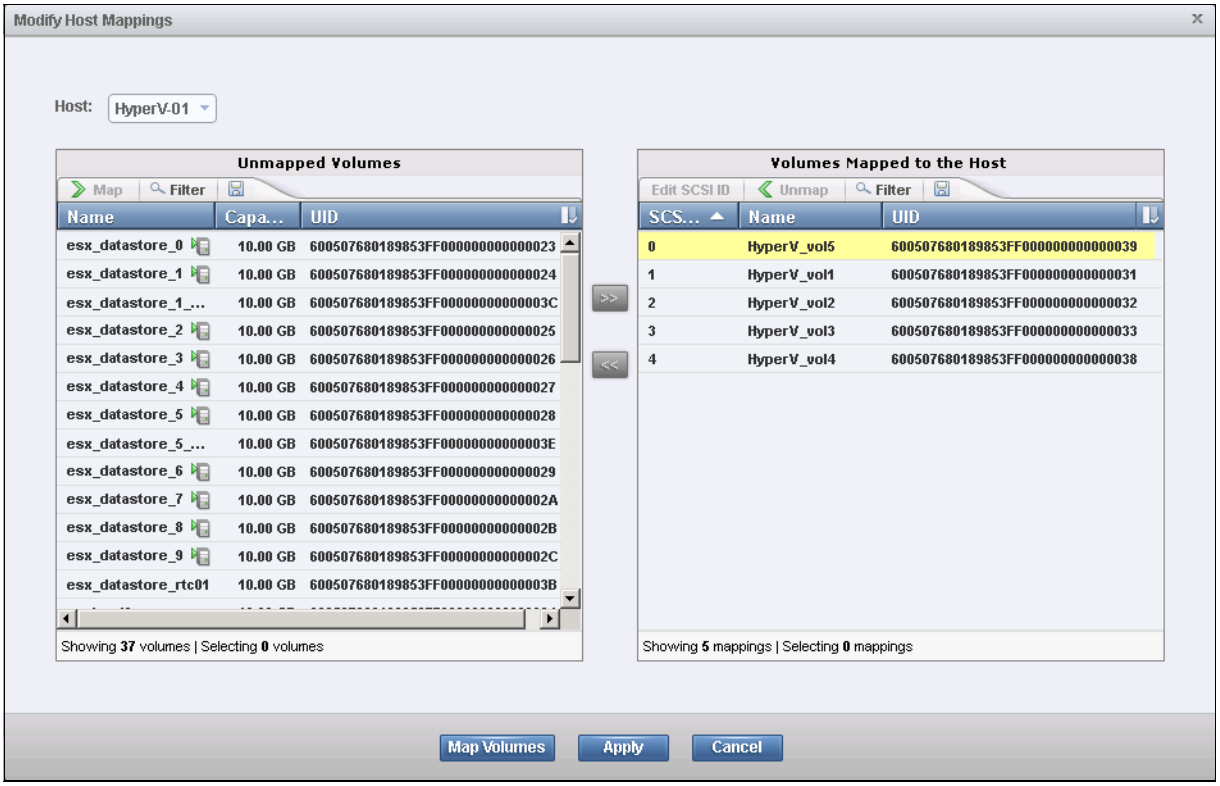

Figure 8-71 Modify Host Mappings window

After the changes are completed, click **Close** to return to the Volumes panel (Figure 8-72).

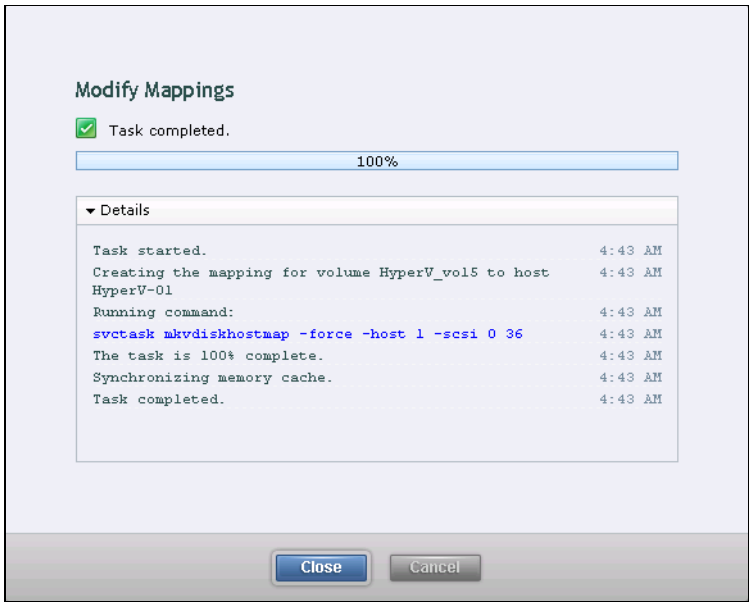

Figure 8-72 Modify Mappings task completion window

Modify Mappings window: The Modify Mappings window is described in 8.1.1, “Modifying Mappings menu” on page 224.

8.4.3 Unmapping volumes from all hosts

To remove all host mappings from a volume, click **Unmap All Host** (Figure 8-68 on page 254). This action removes all host mappings from this volume, which means that no hosts are able to access this volume anymore. Enter the number of mappings that are affected and click **Unmap** to perform this action (Figure 8-73).

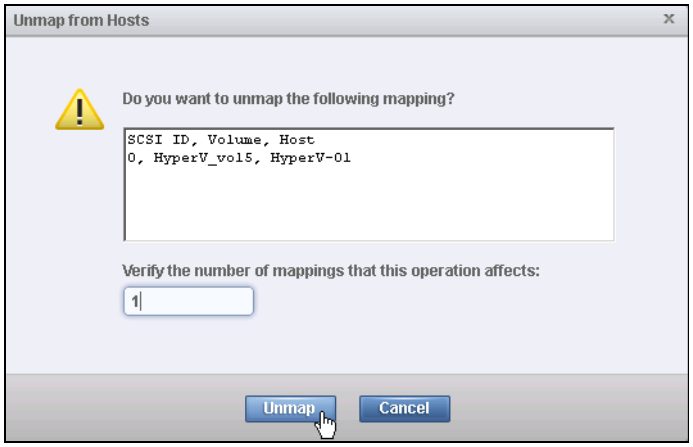

Figure 8-73 Unmap Volumes: Enter number of mappings to be removed

After the task completes, click **Close** (Figure 8-74) to return to the Volumes panel.

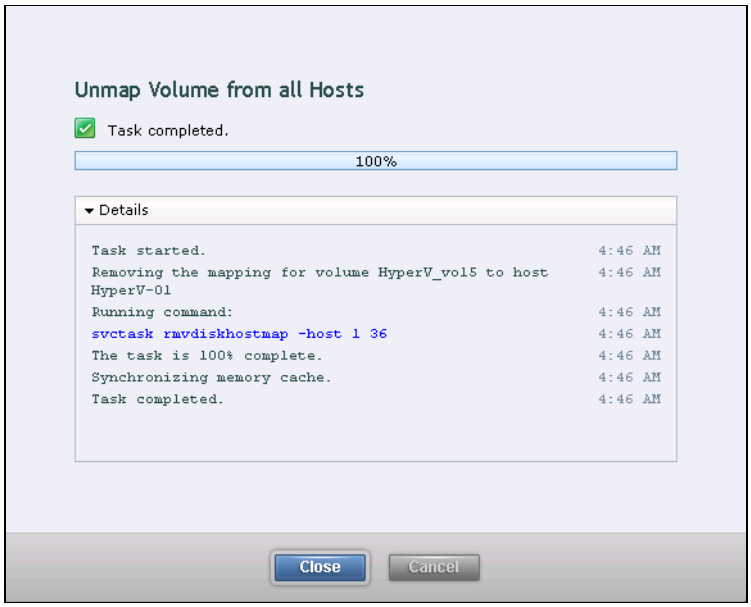

Figure 8-74 Unmap Volume from all Hosts task completion window

8.4.4 Viewing a host mapped to a volume

If you want to know which host mappings are currently configured, highlight a volume and click **View Mapped Host** (Figure 8-68 on page 254). This action opens the Host Maps tab of the Volume Details window (Figure 8-75). This example shows one existing host mapping for the *esx_datastore_2* volume.

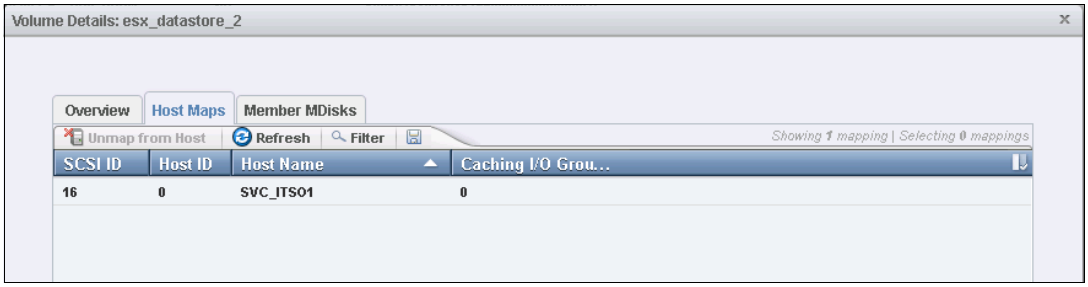

Figure 8-75 Volume details

To remove a mapping, select the host and click **Unmap from Host**, which removes the access for the selected host (after you confirm it). If several hosts are mapped to this volume (for example, in a cluster), only the highlighted host is removed.

8.4.5 Duplicating a volume

You can create a new volume using the same presets and parameters as an existing volume. These parameters are as follows:

- Volume preset (generic, thin-provision, compressed)
- Volume size
- Storage pool
- Access and Caching I/O group
- Caching mode
- Easy Tier status
- Virtualization type

Note: Duplicating a volume creates a new volume with the same preset and volume parameters as the source volume. Duplicating a volume does not duplicate volume data. Duplicating mirrored and image-mode volumes is not supported.

To duplicate a volume, select **Duplicate Volume** from the Actions menu (Figure 8-68 on page 254). The Duplicate Volume window opens (Figure 8-76). By default, a sequence integer is appended to the name of the volume you are duplicating. You can change this name as shown in this figure. Click **Duplicate**.

Figure 8-76 Duplicate Volume

Close the task window after the volume duplicate task is finished (Figure 8-77).

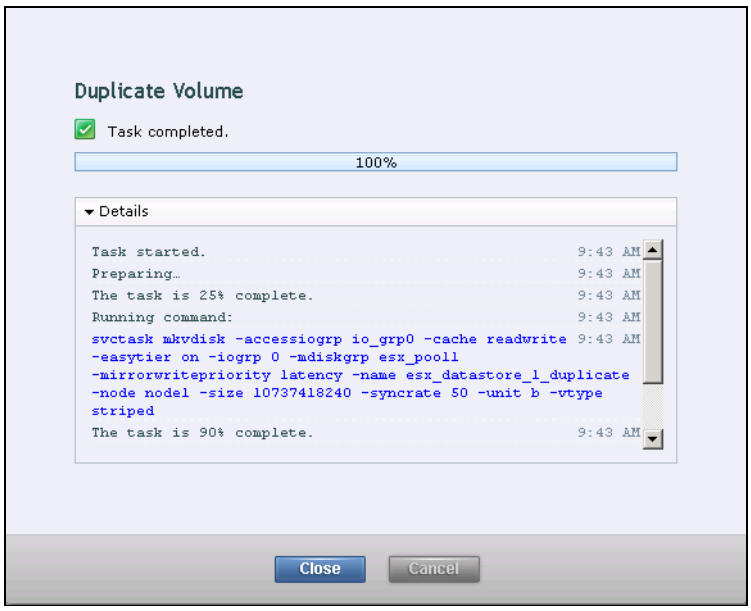

Figure 8-77 Duplicate volume task completion window

8.4.6 Renaming a volume

To rename a volume, select **Rename** (Figure 8-68 on page 254). A window opens.

Enter the new name (Figure 8-78). If you click **Reset**, the name field is always reset to the currently active name of the volume. Click **Rename** to apply the changes.

Figure 8-78 Rename Volume window

Click **Close** when you are done (Figure 8-79).

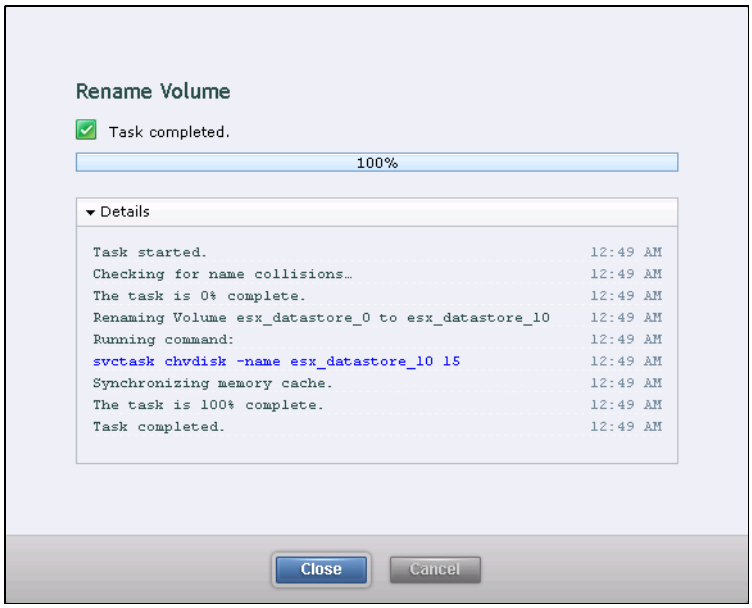

Figure 8-79 Rename Volume task completion window

8.4.7 Shrinking a volume

IBM Storwize V7000 offers you the ability to shrink volumes. However, you should only use this feature if your host operating system supports this feature. Before shrinking a volume, complete the preparation required in your host operating system to shrink a volume on the storage system. After you have prepared your OS, click **Shrink** (Figure 8-68 on page 254). You can either enter the new size, or enter how much the volume should shrink. If you enter a value, the other line updates itself (Figure 8-80). Click **Shrink** to start the process.

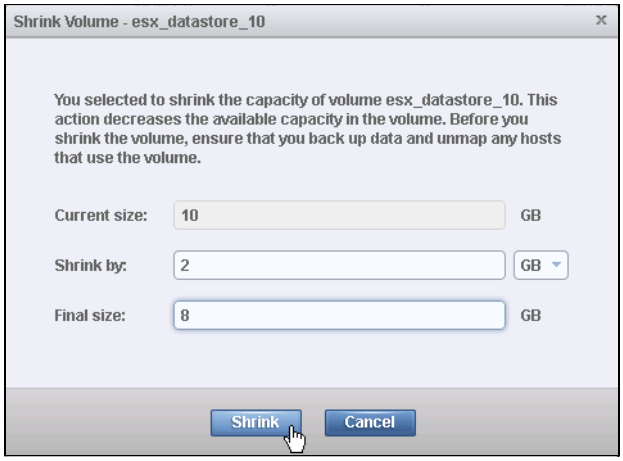

Figure 8-80 Shrink Volume window

Click **Close** (Figure 8-81) to return to the Volumes window.

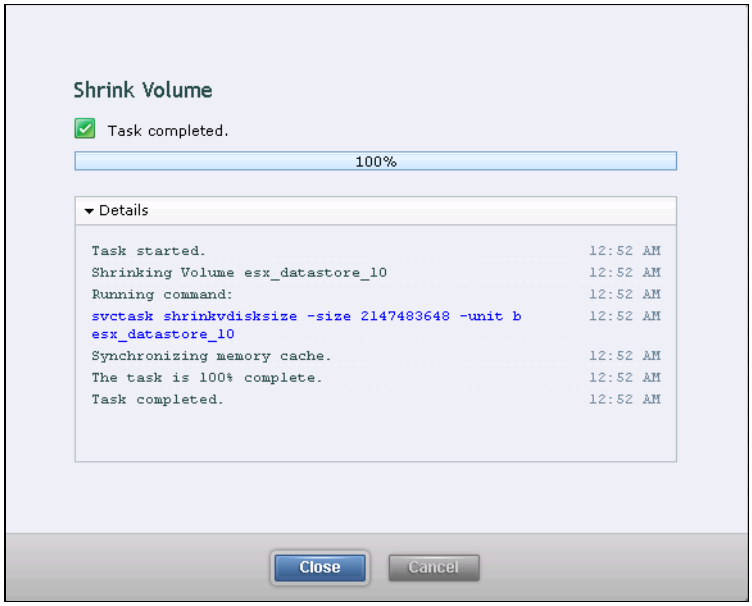

Figure 8-81 Shrink Volume task completion window

Run the required procedures on your host to complete the shrinking process.

8.4.8 Expanding a volume

To expand a volume, click **Expand** (Figure 8-68 on page 254); the Expand Volume window opens. Before you continue, determine whether your operating system supports online volume expansion. Enter the new volume size and click **Expand** (Figure 8-82).

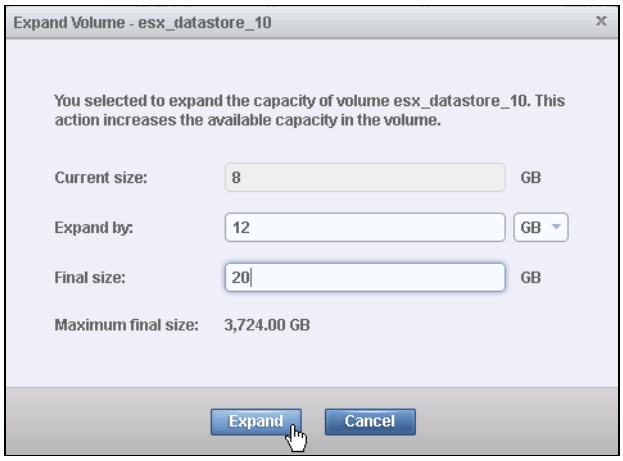

Figure 8-82 Expand Volume window

After the tasks are completed, click **Close** (Figure 8-83) to return to the Volumes panel.

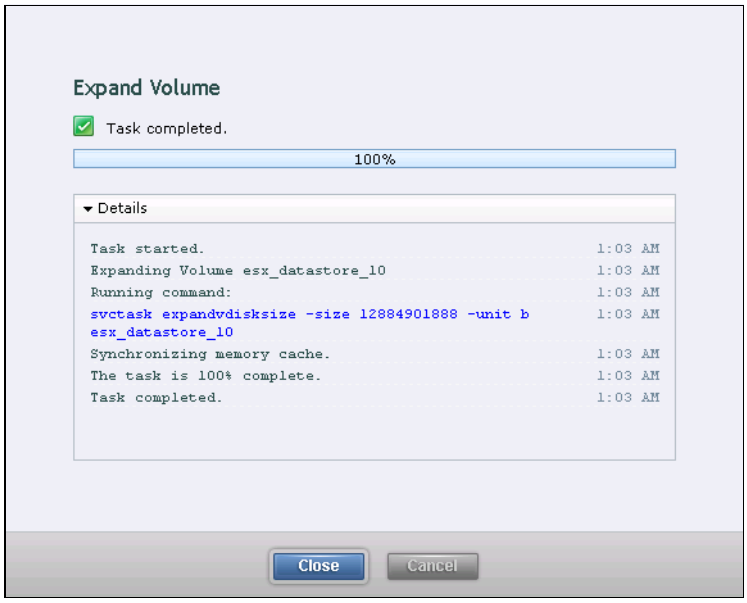

Figure 8-83 Expand Volume task completion window

Run the required procedures in your operating system to use the newly available space.

8.4.9 Migrating a volume to another storage pool

IBM Storwize V7000 provides online volume migration while applications are running. Storage pools are managed disk groups, as described in Chapter 7, “Storage pools” on page 179. With volume migration, you can move the data between these storage pools, regardless of whether the pool is an internal pool, or a pool on another external storage system. This migration is done without the server and application knowing that it even

occurred. The migration process itself is a low priority process that does not affect the performance of the IBM Storwize V7000. However, it moves one extent after another to the new storage pool, so the performance of the volume is affected by the performance of the new storage pool after the migration process.

To migrate a volume to another storage pool, click **Migrate to Another Pool** (Figure 8-68 on page 254). The Migrate Volume Copy window opens. If your volume consists of more than one copy, select the copy (from the menu shown in Figure 8-84) that you want to migrate to another storage pool. If the selected volume consists of one copy, this selection menu is not available.

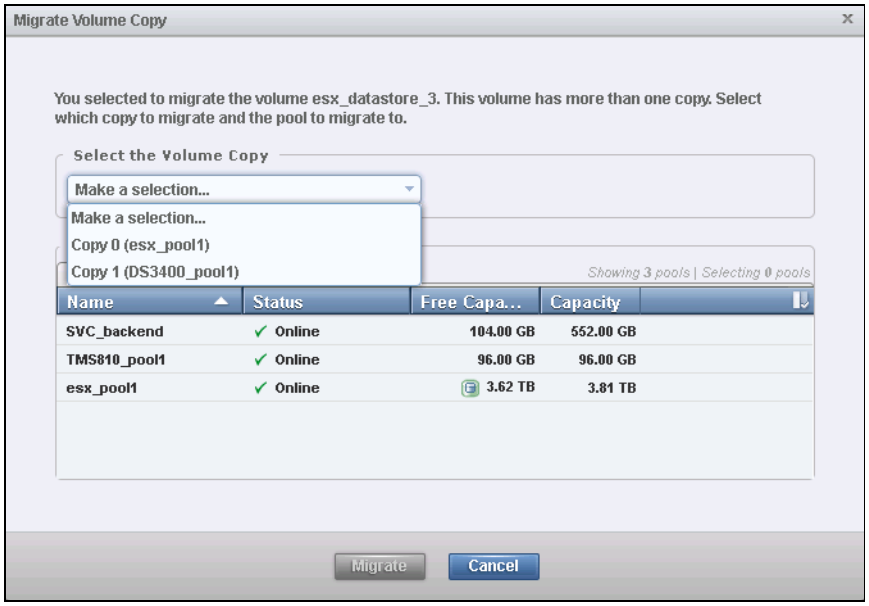

Figure 8-84 Migrate Volume Copy window: Select copy

Select the new target storage pool and click **Migrate** (Figure 8-85).

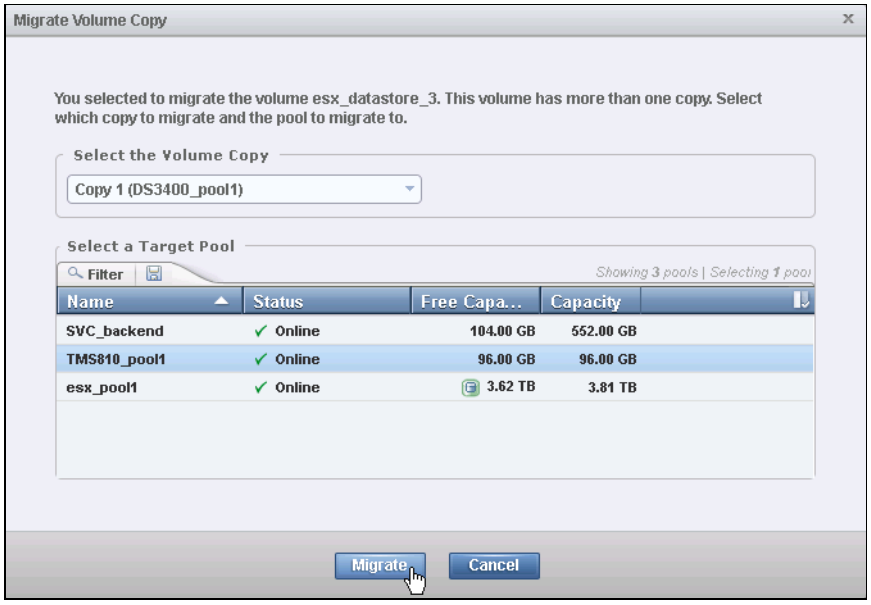

Figure 8-85 Migrate Volume Copy, Select target pool

The volume copy migration starts (Figure 8-86). Click **Close** to return to the Volumes panel.

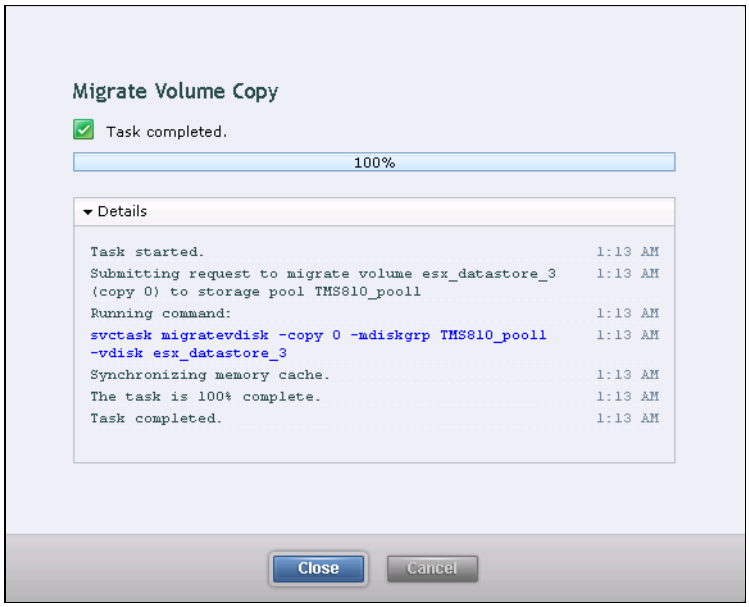

Figure 8-86 Migrate Volume Copy started

Depending on the size of the volume, the migration process takes some time, but you can monitor the status of the migration in the running tasks bar at the bottom of the GUI (Figure 8-87).

Figure 8-87 Migration progress

After the migration is completed, the volume is shown in the new storage pool. Figure 8-88 shows that it was moved from the mdiskgrp0 pool to the mdiskgrp1 pool.

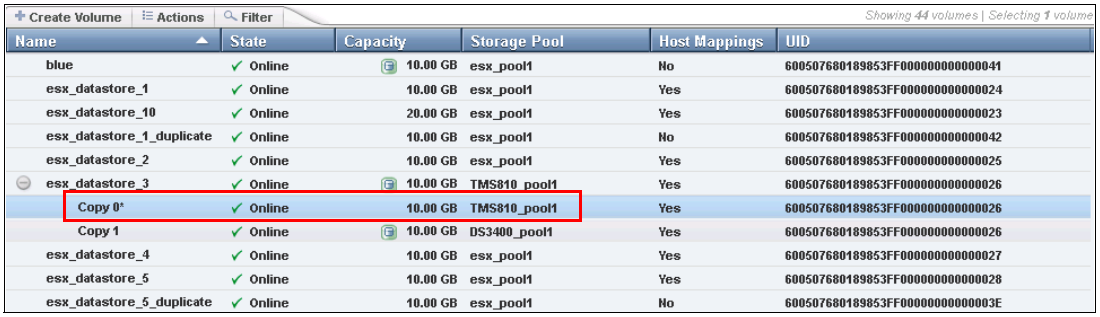

Figure 8-88 Migration complete

The volume copy has now been migrated without any host or application downtime to the new storage pool. It is also possible to migrate both volume copies to other pools online.

Another way to migrate volumes to another pool is by performing the migration using the volume copies, as described in 8.6.6, “Migrating volumes using the volume copy features” on page 286.

Note: Migrating a volume between storage pools with different extent sizes is not supported. If you need to migrate a volume to a storage pool with a different extent size, use volume copy features instead.

8.4.10 Exporting to an image mode volume

Image mode provides a direct block-for-block translation from the MDisk to the volume with no virtualization. An image mode MDisk is associated with exactly one volume. This feature can be used to export a volume to a non-virtualized disk and to remove the volume from storage virtualization. To export a volume to an image mode volume, click **Export to Image Mode** (Figure 8-68 on page 254), and the window opens (Figure 8-89). Select the MDisk to which you want to export the volume and click **Next**.

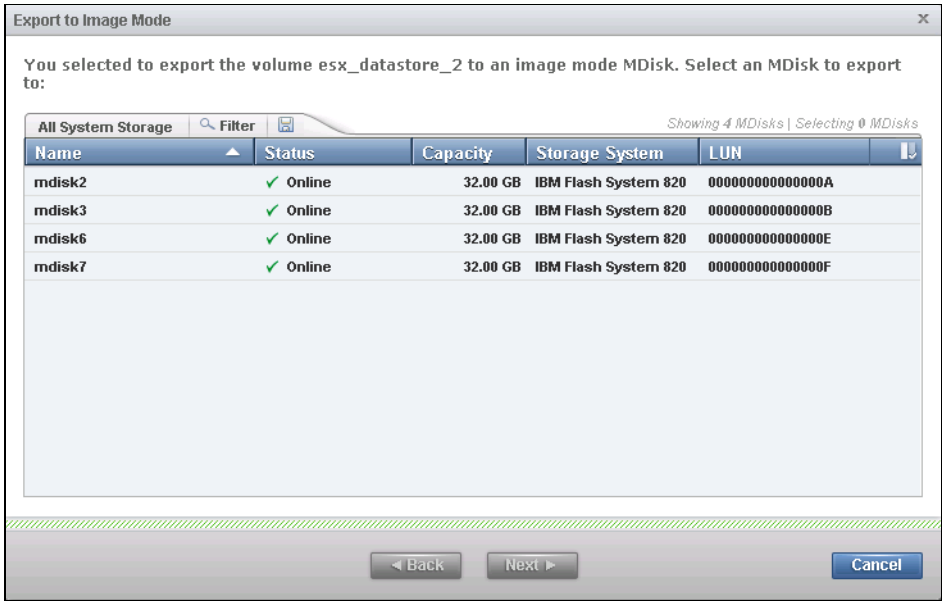

Figure 8-89 Export to Image Mode

Select a storage pool for the new image mode volume and click **Finish** (Figure 8-90).

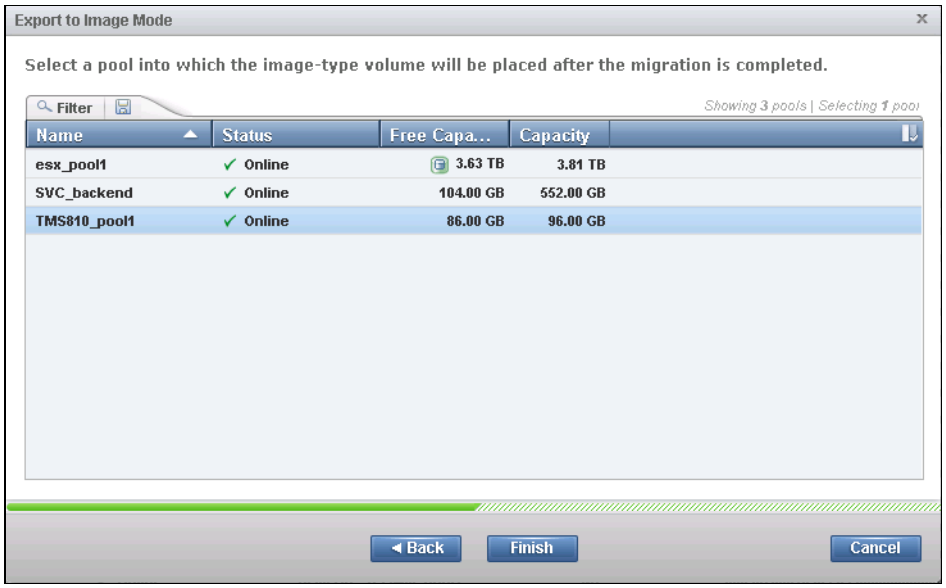

Figure 8-90 Export to Image Mode

The migration starts (Figure 8-91). Click **Close** to return to the Volumes panel.

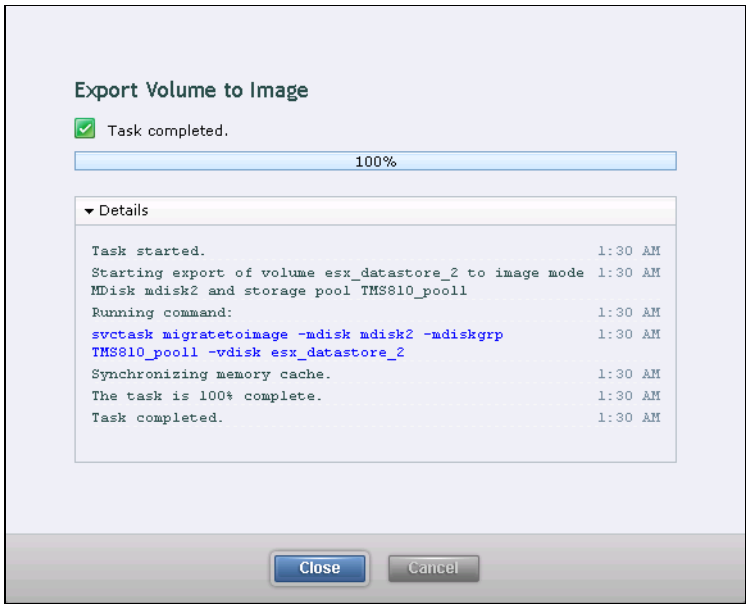

Figure 8-91 Export Volume To Image task completion window

8.4.11 Deleting a volume

To delete a volume, select **Delete** (Figure 8-68 on page 254). Enter the number of volumes you want to delete and select the check box if you want to force the deletion (Figure 8-92). You must force the deletion if the volume has host mappings or is used in FlashCopy mappings or Remote Copy relationships.

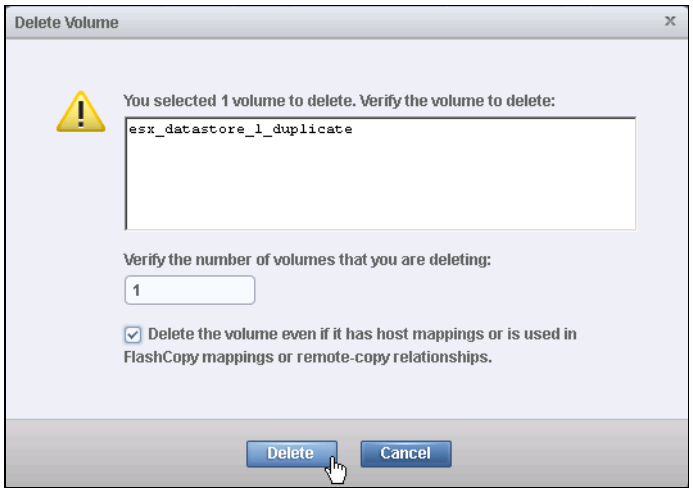

Figure 8-92 Delete Volume

Click **Delete** and the volume is removed from the system (Figure 8-93).

Attention: This action removes all copies from your storage system and the data on the volume is lost. Before you perform this step, be sure that you do not need the volume anymore.

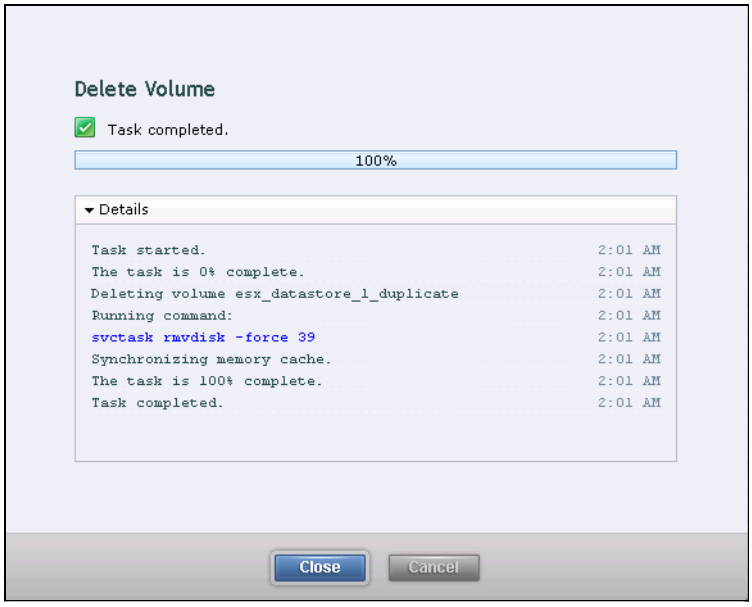

Figure 8-93 Delete Volume task completion window

Click **Close** to return to the Volumes window.

8.5 Volume properties

To open the advanced view of a volume, select **Properties** (Figure 8-68 on page 254); the Volume Details window opens (Figure 8-94). In this window, three tabs are available:

- Overview tab
- Host Maps tab
- Member MDisk tab

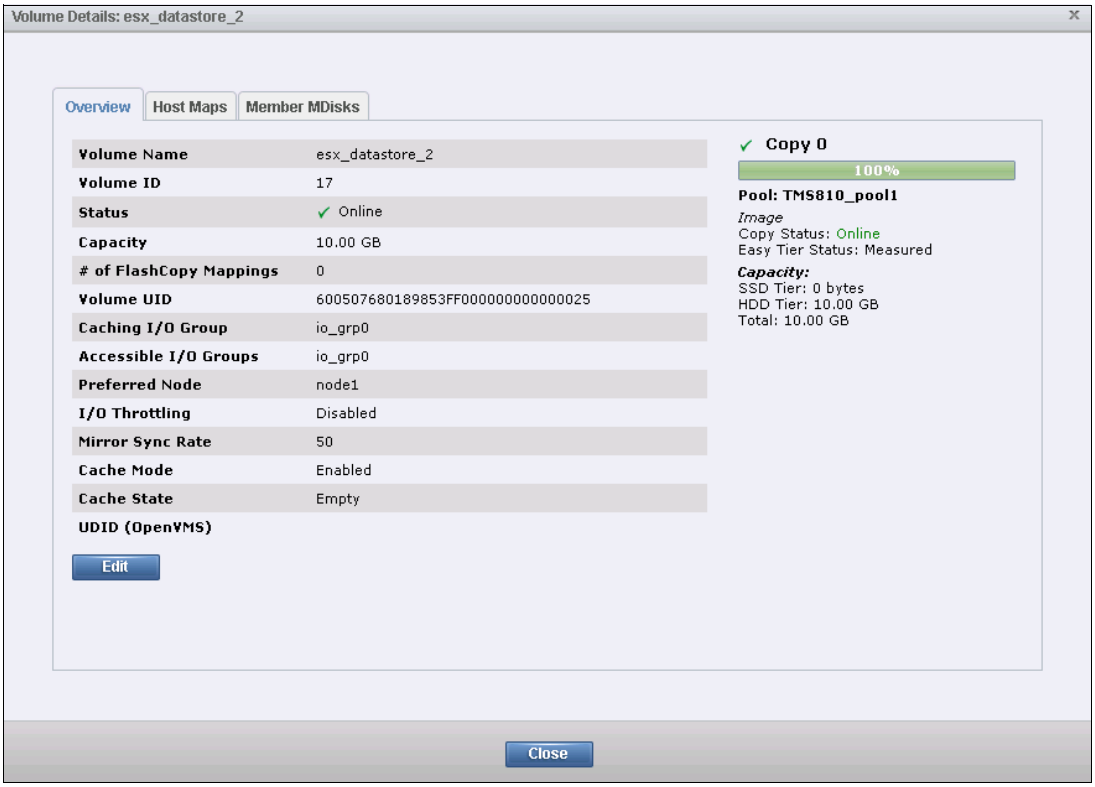

Figure 8-94 Volume properties

8.5.1 Overview tab

The Overview tab (Figure 8-94 on page 267) shows a basic overview of the volume properties. In the left part of the window, you find common volume properties, and in the right part, you see information about the volume copies (Figure 8-95).

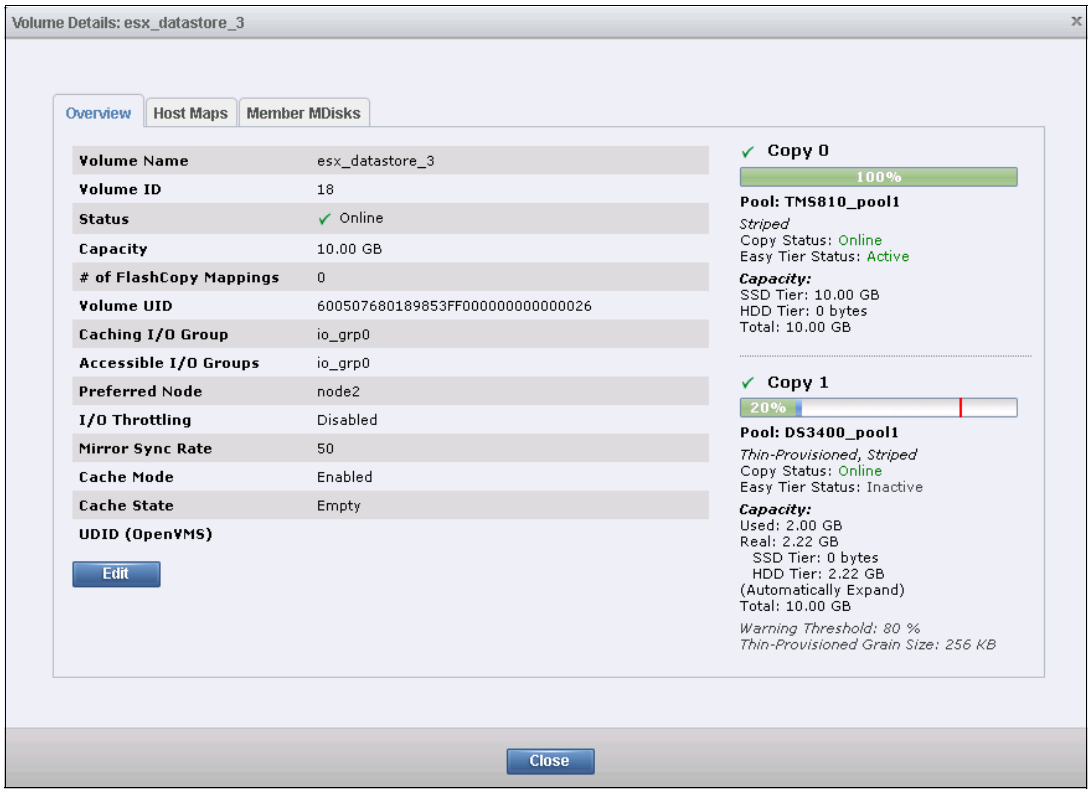

Figure 8-95 Volume Details: Overview tab

The following details are available:

- Volume Properties:
  - Volume Name: Shows the name of the volume.
  - Volume ID: Shows the ID of the volume. Every volume has a system-wide unique ID.
  - Status: Gives status information about the volume which can be online, offline, or degraded.
  - Capacity: Shows the capacity of the volume. If the volume is thin-provisioned, this number is the virtual capacity; the real capacity is displayed for each copy.
  - # of FlashCopy Mappings: Number of existing FlashCopy relationships. A detailed description about FlashCopy is covered in Chapter 11, “Copy services” on page 327.
  - Volume UID: The volume unique identifier.
  - Caching and Accessible I/O Group: Specifies the I/O group to which the volume belongs.
  - Preferred Node: Specifies the ID of the preferred node for the volume.
  - I/O Throttling: You may set a maximum rate at which the volume processes I/O requests. The limit can be set in IOPS or MBps. This is an advanced feature and you may enable it only through the CLI, as described in Appendix A, “CLI setup and SAN Boot” on page 497.

- Mirror Sync Rate: After creation, or if a volume copy is offline, the mirror sync rate weights the synchronization process. Volumes with a high sync rate (100%) complete the synchronization faster than volumes with a lower priority. By default, the rate is set to 50% for all volumes.
- Cache Mode: Shows whether the cache is enabled or disabled for this volume.
- Cache State: Gives you feedback if open I/O requests are inside the cache that have not been destaged to the disks.
- UDID (OpenVMS): The unit device identifiers are used by OpenVMS host to access the volume.

- Copy Properties:
  - Storage Pool: Gives you information about in which pool the copy rests, what copy type it is (generic, thin-provisioned, or compressed), and gives you status information.
  - Capacity: Shows the allocated (used) and the virtual (real) capacity, and the warning threshold and the set grain size.

To edit any of these settings, click **Edit**; the window changes to the modify mode (Figure 8-96).

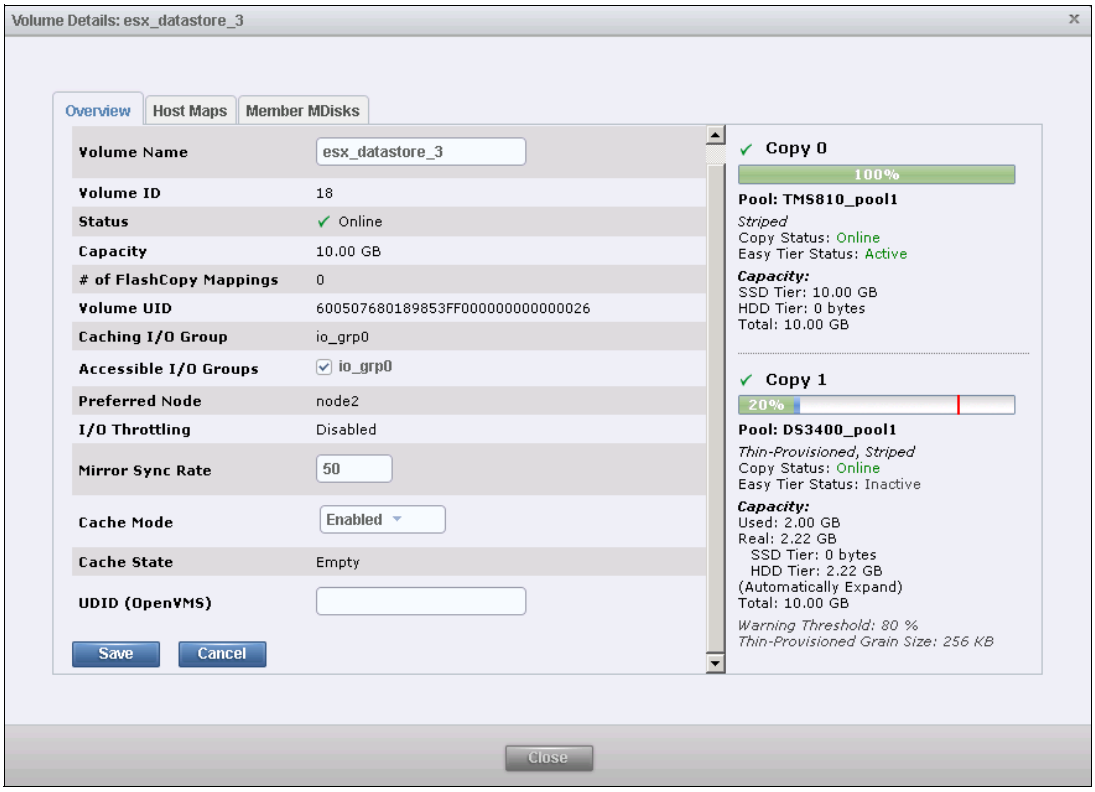

Figure 8-96 Edit volume details

In the Volume Details window, you may change the following properties:

- Volume Name
- Accessible I/O Group
- Mirror Sync Rate
- Cache Mode
- UDID

Make any required changes and click **Save** (Figure 8-97). For this example, we change the Mirror Sync Rate.

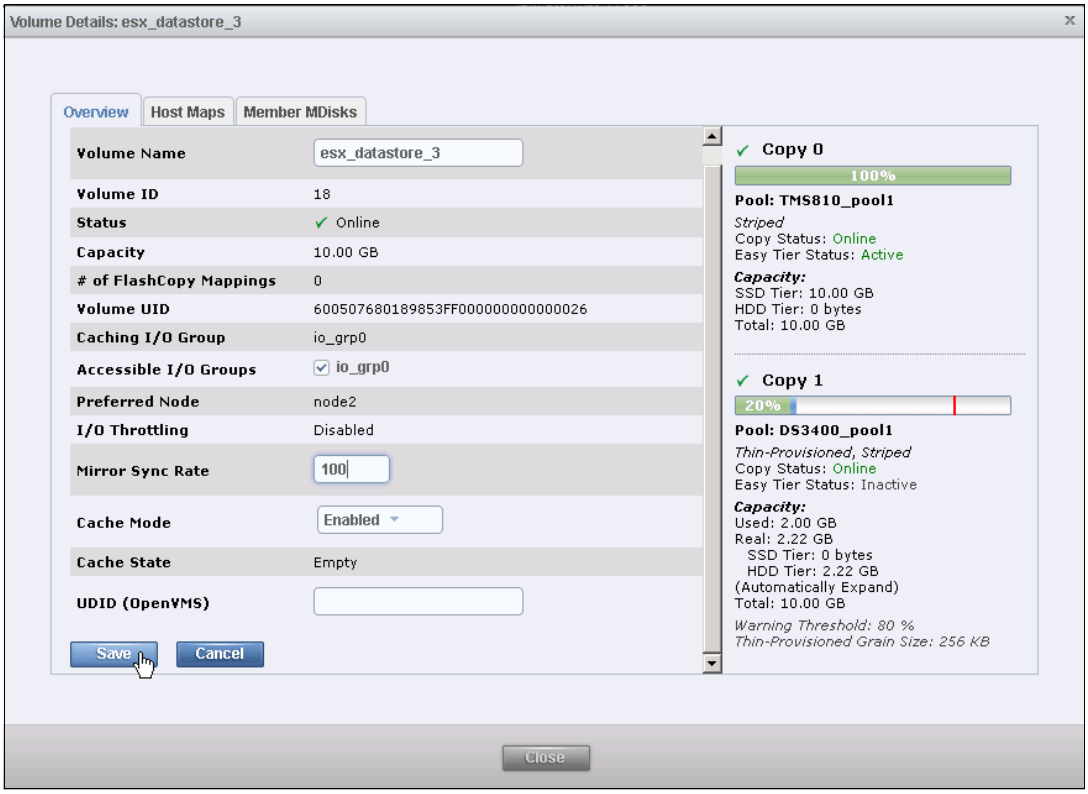

Figure 8-97 Change properties

The changes are applied to the system (Figure 8-98).

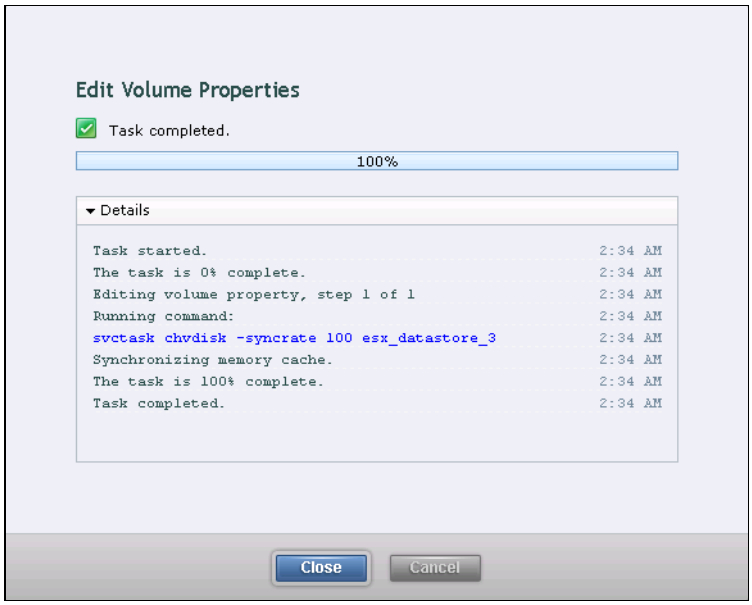

Figure 8-98 Edit Volume Properties task completion window

Mirror Sync Rate: Setting the Mirror Sync Rate to 0% disables synchronization.

8.5.2 Host Maps tab

The second tab of the volume properties is Host Maps (Figure 8-99). All hosts that are currently mapped to this volume are listed in this view.

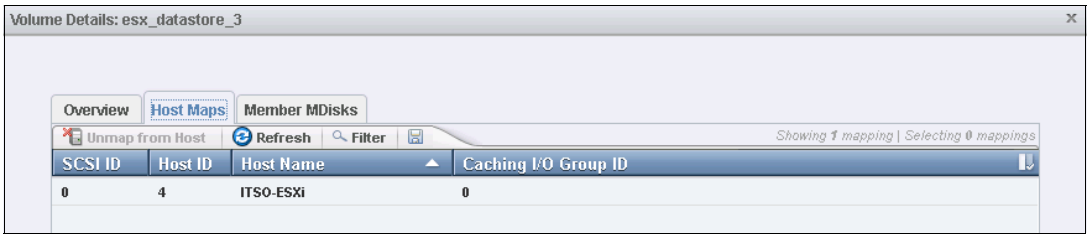

Figure 8-99 Volume Details: Host Maps tab

If you want to unmap a host from this volume, select it and click **Unmap from Host**. Confirm the number of mappings to remove and click **Unmap** (Figure 8-100).

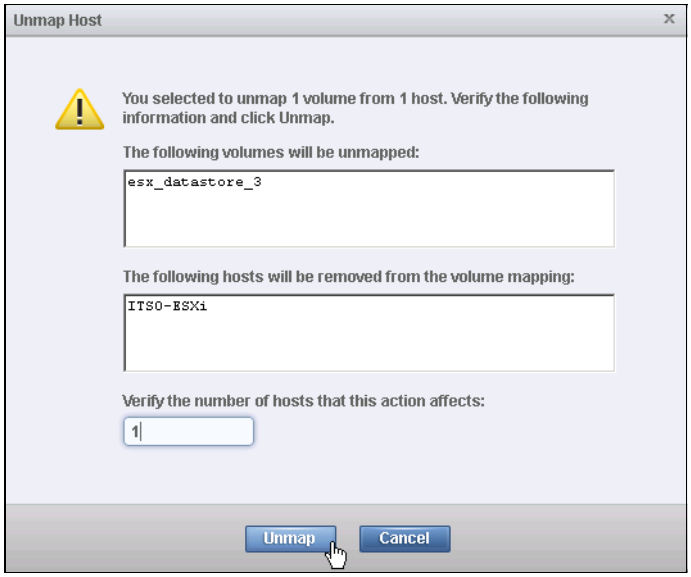

Figure 8-100 Unmap from host

The changes are applied to the system (Figure 8-101). The selected host no longer has access to this volume. Click **Close** to return to the Volumes panel.

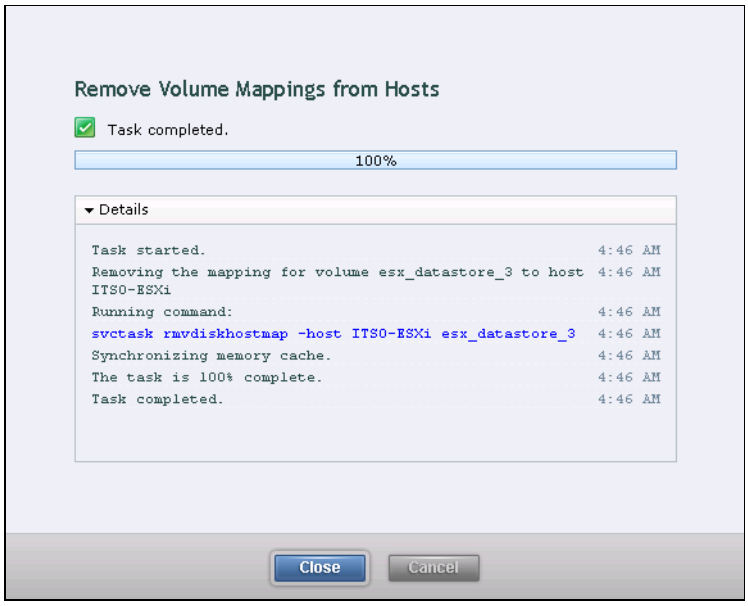

Figure 8-101 Remove Volume Mappings task completion window

8.5.3 Member MDisk tab

The third tab is Member MDisk, which lists all MDisks on which the volume is located. Select a copy and the associated MDisks are then listed in the window (Figure 8-102).

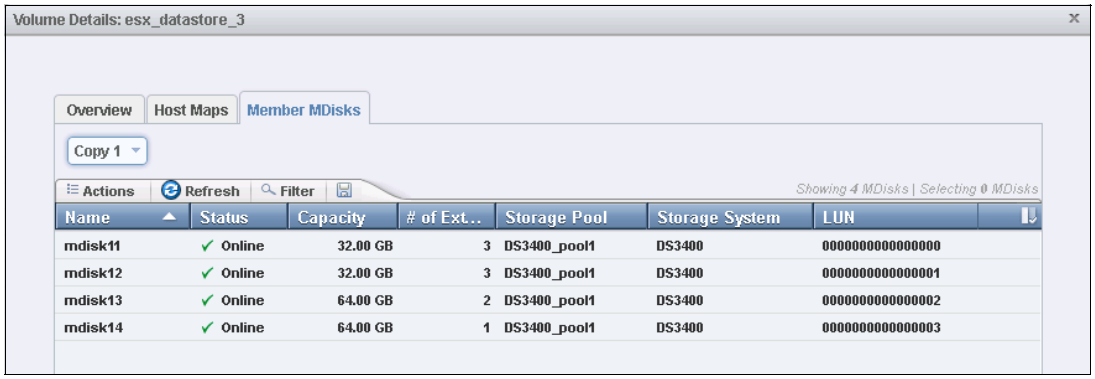

Figure 8-102 Volume Details: Member MDisks

Select an MDisk and click **Actions** to view the available tasks (Figure 8-103). The tasks are described in Chapter 7, “Storage pools” on page 179.

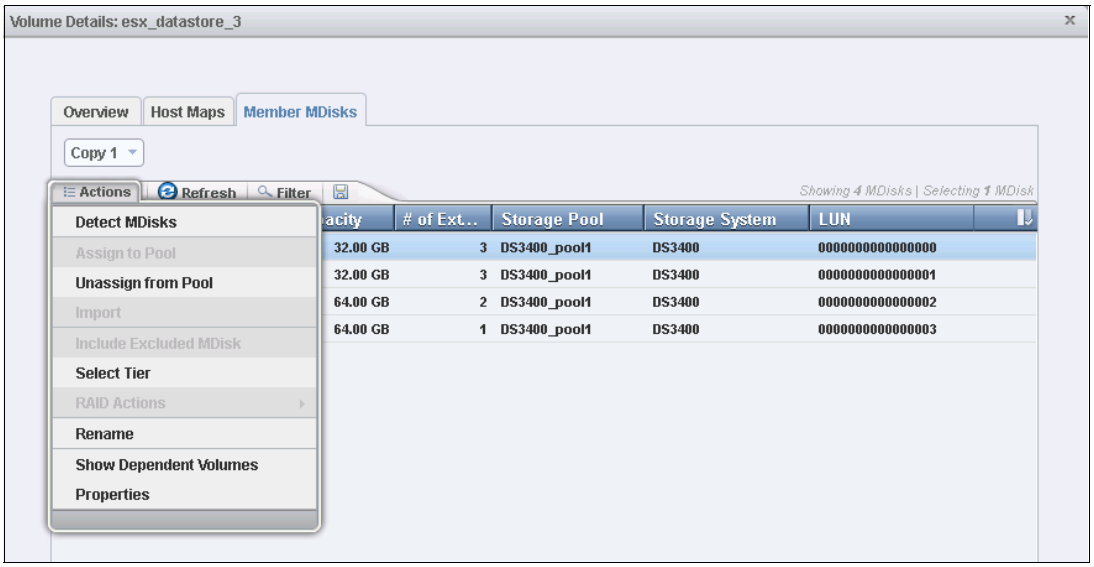

Figure 8-103 Member MDisk actions

Click **Close** to return to the Volumes panel.

8.5.4 Adding a mirrored volume copy

If you have an existing volume that consists of only one copy, you can add a mirrored volume copy to it. This action creates an extra copy of your volume online. The second copy can be generic, thin-provisioned, or compressed, and it can be created in any storage pool. To add a second copy, click **Volume Copy Actions** → **Add Mirrored Copy** (Figure 8-104).

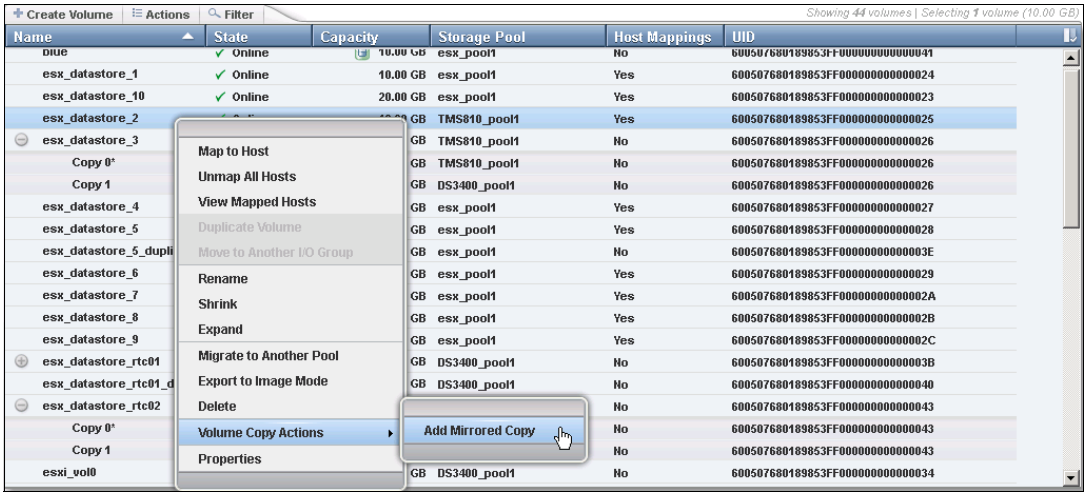

Figure 8-104 Add Mirrored Copy

Select the storage pool in which to create the new copy; if you want the new copy to be thin-provisioned or compressed, select the appropriate preset (Figure 8-105). Click **Add Copy**.

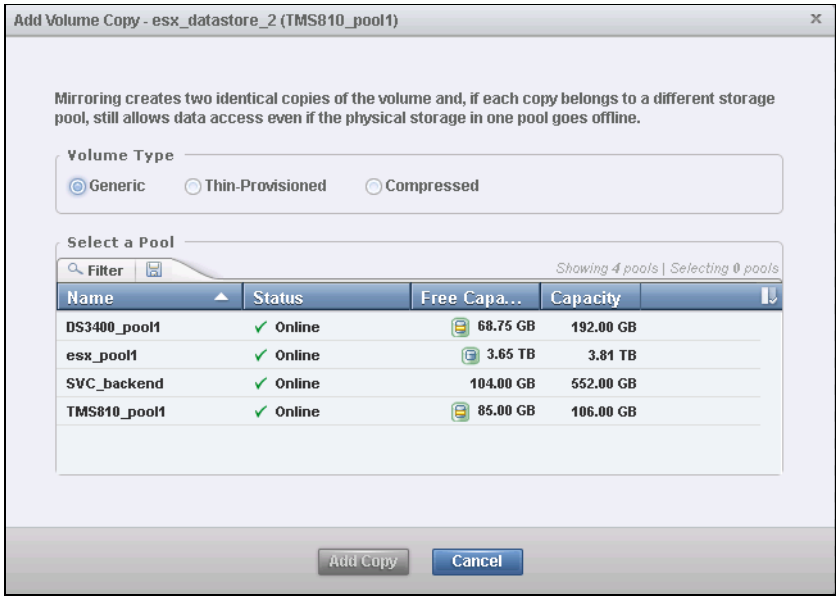

Figure 8-105 Add Volume Copy window

The new copy is created (Figure 8-106) and the data is synchronized as a background task.

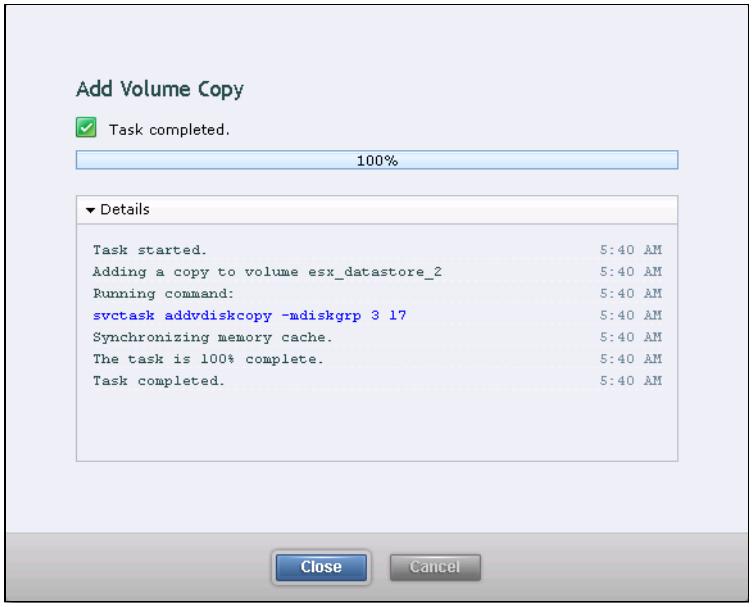

Figure 8-106 Add Volume Copy task completion window

Figure 8-107 shows that the volume named *esx_datastore_2* now has two volume copies, and one synchronization is task running.

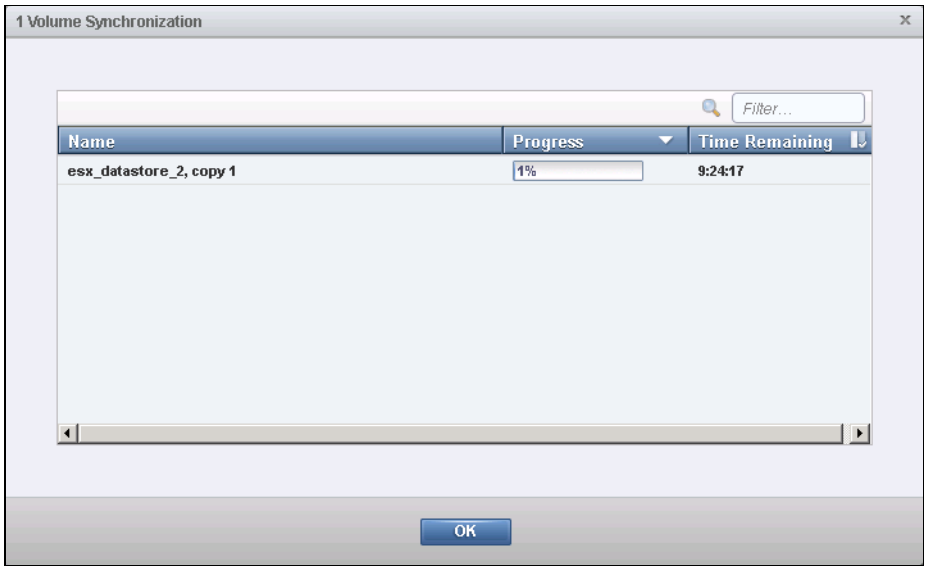

Figure 8-107 Volume copy synchronization

8.5.5 Editing thin-provisioned or compressed volume properties

From a host's perspective, the virtual capacity expansion and shrinkage of a volume affects the host access. The real capacity expansion or shrinkage of a volume, which is described in this section, is transparent to the hosts.

Note: In the following sections, we demonstrate real capacity operations; we use a thin-provisioned volume as an example. However, the same actions apply to the compressed volume preset.

modifying the volume size that is presented to a host is covered in 8.4.7, "Shrinking a volume" on page 260 and 8.4.8, "Expanding a volume" on page 261. However, if you have a thin-provisioned or compressed volume, you can also edit the allocated (real) size and the warning thresholds. To accomplish this task, select the thin-provisioned or compressed volume, click **Actions** → **Volume Copy Actions** → **Thin Provisioned (or Compressed)** (Figure 8-69 on page 254). Three more options are available, as shown in Figure 8-108 on page 276:

- Shrink: See "Shrinking thin-provisioned or compressed volume space" on page 276.
- Expand: See "Expanding thin-provisioned or compressed volume space" on page 277.
- Edit Properties: See "Editing thin-provisioned or compressed volume properties" on page 278.

These changes are only to the internal storage, so you do not have to make any changes on your host.

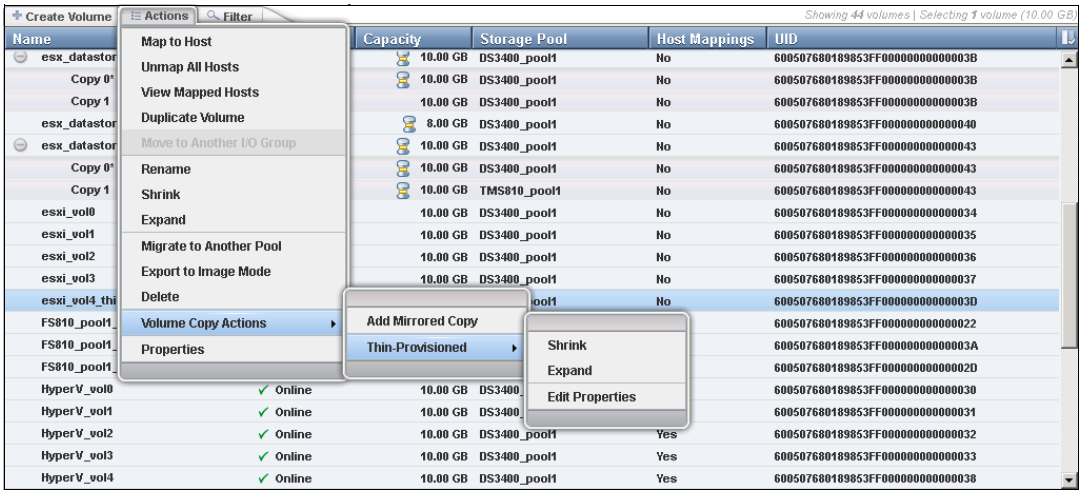

Figure 8-108 Thin-provisioned actions

Shrinking thin-provisioned or compressed volume space

Select **Shrink** (Figure 8-108) to reduce the allocated space of a thin-provisioned or compressed volume. Enter either how much the volume should shrink or the new final size (Figure 8-109), and click **Shrink**.

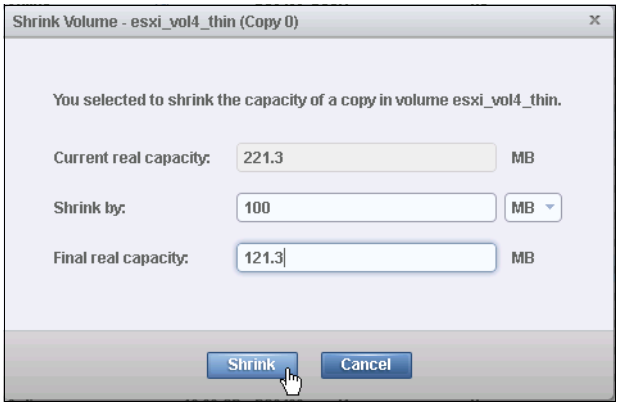

Figure 8-109 Shrink Volume real capacity window

After the task completes, click **Close** (Figure 8-110).

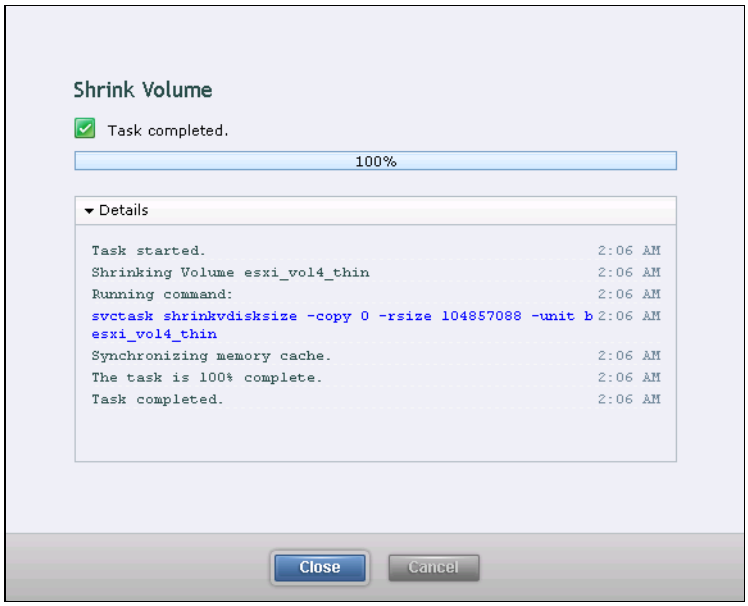

Figure 8-110 Shrink Volume task completion window

The allocated space of the thin-provisioned volume is now reduced.

Deallocating extents: You can only deallocate extents that do not have stored data on them. If the space is allocated, because there is data on them, you are not able to shrink the allocated space.

Expanding thin-provisioned or compressed volume space

To expand the allocated space of a thin-provisioned or compressed volume, select **Expand** (Figure 8-108 on page 276). Enter either how much space should be allocated or enter the new final size and click **Expand** (Figure 8-111).

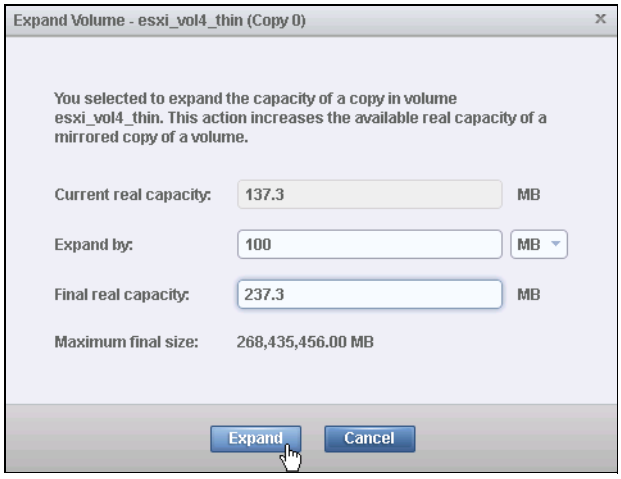

Figure 8-111 Expand Volume real capacity window

The new space is now allocated. Click **Close** (Figure 8-112).

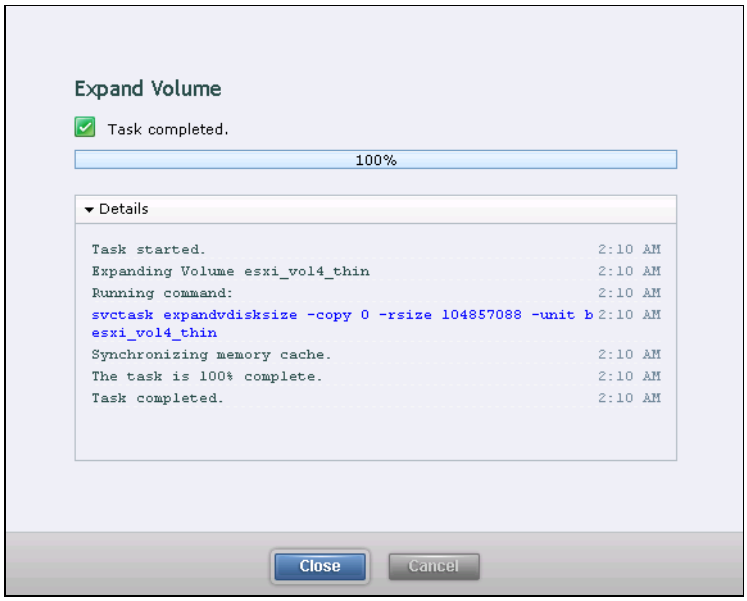

Figure 8-112 Expand Volume task completion window

Editing thin-provisioned or compressed volume properties

To edit thin-provisioned or compressed volume properties, select **Edit Properties** (Figure 8-108 on page 276). Edit the settings if required (Figure 8-113) and click **OK** to apply the changes.

Figure 8-113 Editing thin-provisioned properties

After the task completes, click **Close** (Figure 8-114) to return to the Volumes panel.

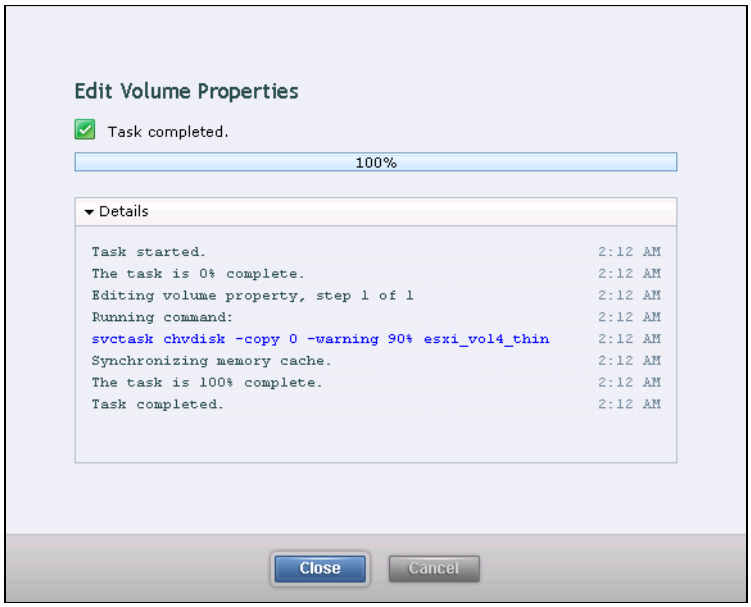

Figure 8-114 Edit Volume Properties task completion window

8.6 Advanced volume copy functions

In 8.4.1, “Advanced volume functions” on page 253, we describe all the available actions at a volume level and how to create a second volume copy. In this section, we focus on volumes consisting of two volume copies. If you expand the volume and highlight a copy, more volume copy actions become available (Figure 8-115):

- Make Primary (when selecting secondary copy)
- Thin-provisioned or compressed volumes
- Split into New Volume
- Validate Volume Copies
- Duplicate Volume Copy
- Delete Copy option

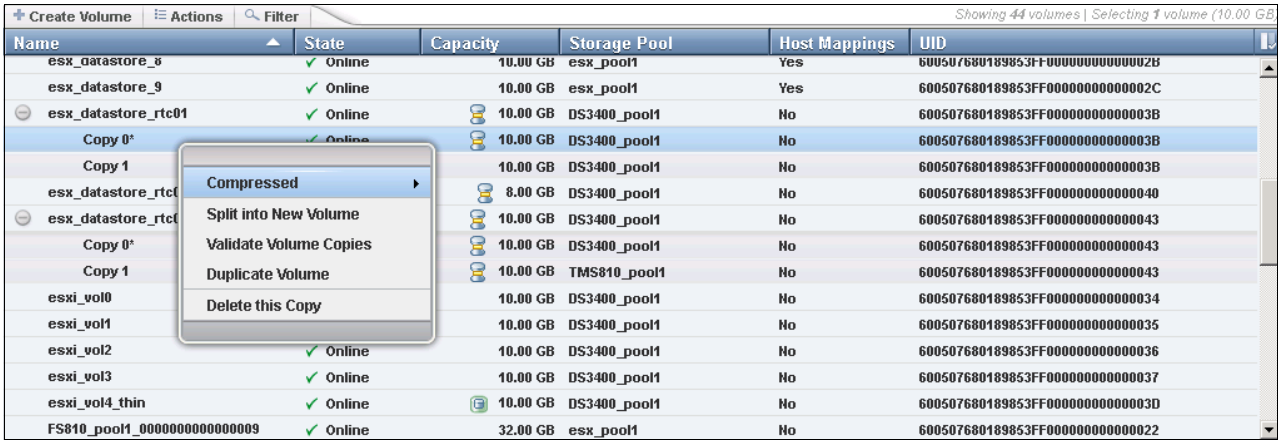

Figure 8-115 Advanced volume copy functions

If you look at the volume copies shown in Figure 8-115 on page 279, you notice that one of the copies has an asterisk (*) next to its name (Figure 8-116).

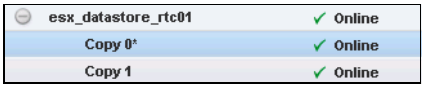

Figure 8-116 Primary copy indicated by an asterisk (*)

Each volume has a primary and a secondary copy, and the star indicates the primary copy. The two copies are always synchronized, which means that all writes are destaged to both copies, but all reads are always done from the primary copy. By default, the primary and secondary copies always switch between Copy 0 and Copy 1 during creation to balance the reads across your storage pools. However, always be sure that the I/Os to the primary copies are customized to the performance of all your storage pools, so therefore you can change the roles of your copies.

To accomplish this task, right-click the secondary copy and select **Make Primary** (Figure 8-117). Usually, a preferred practice is to place the volume copies on storage pools with similar performance because the write performance is constrained if one copy is on a lower performance pool. Writes must complete to both copies before the volume can provide acknowledgment to the host that the write completed successfully.

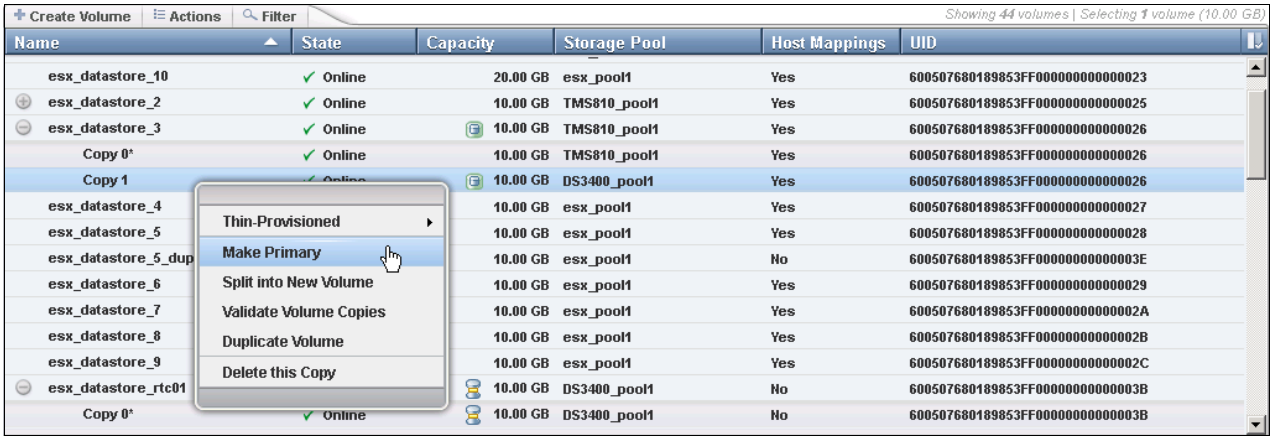

Figure 8-117 Make Primary menu

If you need only high-read performance, another possibility is to place the primary copy in an SSD pool, and the secondary copy in a normal disk pool. This action maximizes the read performance of the volume and makes sure that you have a synchronized second copy in your less expensive disk pool. It is also possible to migrate the copies online between storage pools; in 8.4.9, “Migrating a volume to another storage pool” on page 261, you can select which copy you want to migrate.

Click **Make Primary** and the role of the copy is changed to online. Click **Close** when the task completes (Figure 8-118).

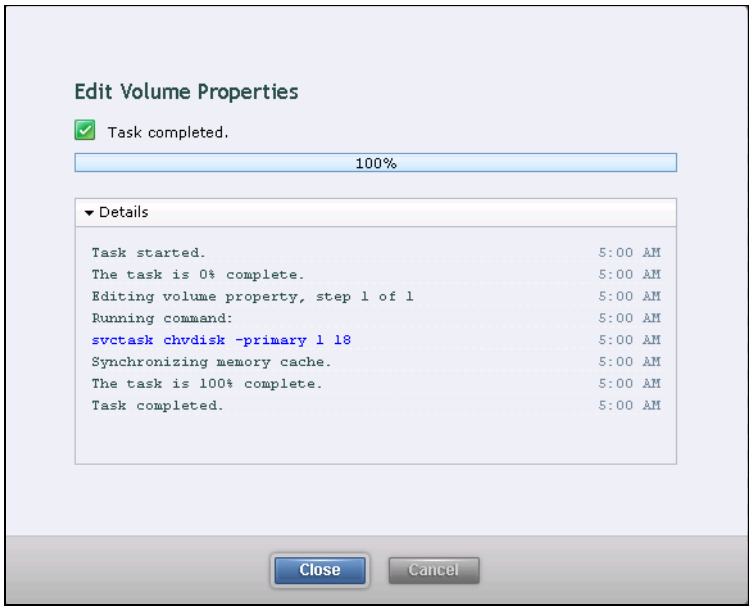

Figure 8-118 Edit Volume Properties task completion window

The volume copy feature is also a powerful option for migrating volumes, as described in 8.6.6, “Migrating volumes using the volume copy features” on page 286.

8.6.1 Thin-provisioned or compressed volumes

These functions (Figure 8-119) are the same as described in the following sections. You can specify the same settings for each volume copy.

- “Shrinking thin-provisioned or compressed volume space” on page 276
- “Expanding thin-provisioned or compressed volume space” on page 277
- “Editing thin-provisioned or compressed volume properties” on page 278

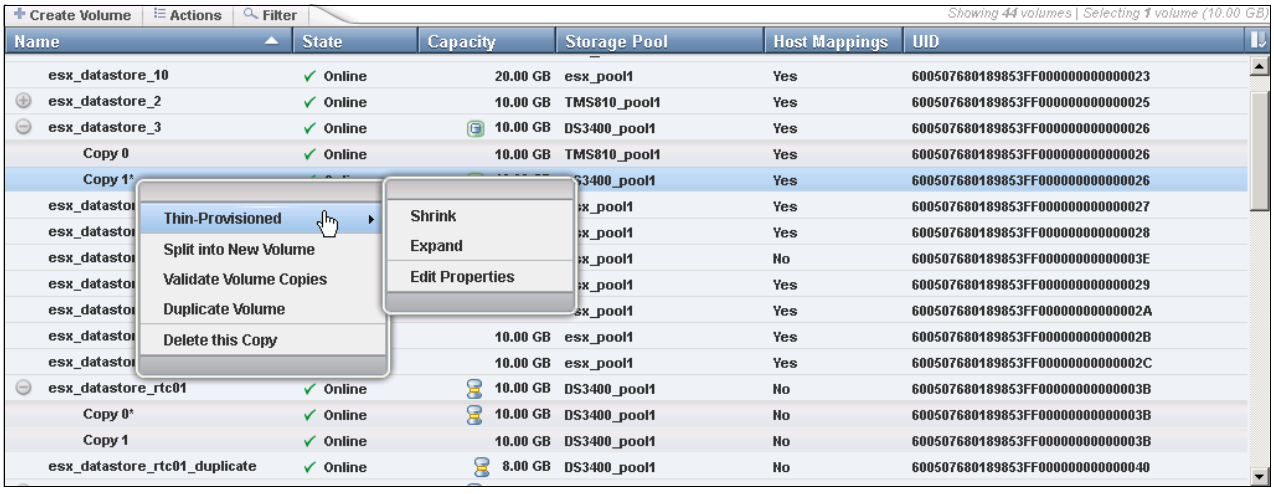

Figure 8-119 Thin-provisioned volume copy tasks

8.6.2 Splitting into a new volume

If your two volume copies are synchronized, you can split one of the copies to a new volume and map this new volume to another host. From a storage point of view, this procedure can be performed online, which means you could split one copy from the volume, and create a copy from the remaining one without any host impact. However, if you want to use the split copy for testing or backup purposes, you must make sure that the data inside the volume is consistent. Therefore, you must flush the data to storage to make the copies consistent. See your operating system documentation for details about flushing the data. The easiest way is to shut down the hosts before splitting a copy.

To split a copy, click **Split into New Volume** (Figure 8-115 on page 279). If you perform this action on the primary copy, the remaining secondary copy automatically becomes the primary for the source volume. Enter a name for the new volume and click **Split Volume Copy** (Figure 8-120).

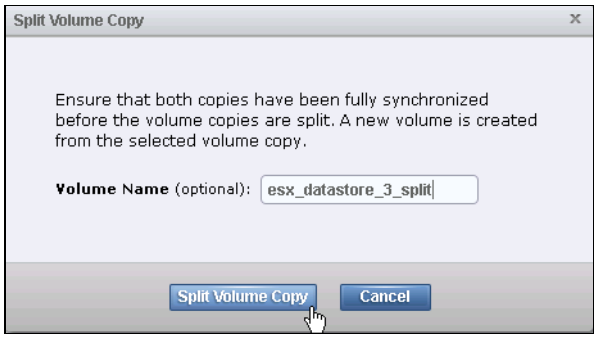

Figure 8-120 Split Volume Copy window

After the task completes, click **Close** to return to the Volumes panel, where the copy appears as a new volume that can be mapped to a host (Figure 8-121).

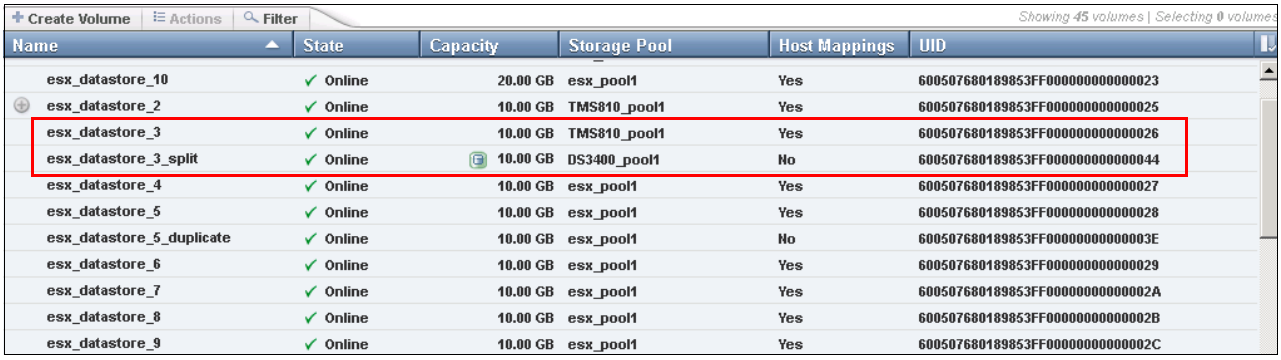

Figure 8-121 All Volumes: New volume from split copy

Note: After splitting the volume copies, the new volume presents itself with a new UID. Compare the UIDs in Figure 8-119 on page 281 and Figure 8-121.

8.6.3 Validate Volume Copies

You must validate the volume copies in certain cases. An example is when one of the storage back-end arrays that is hosting one of the copies becomes unavailable or lost because of the planned or unplanned outage, storage array maintenance, or any other situation that causes the two copies to become out of sync.

To validate the copies of a mirrored volume, complete the following steps:

1. Select **Validate Volume Copies** (Figure 8-115 on page 279). The Validate Volume Copies window opens (Figure 8-122).

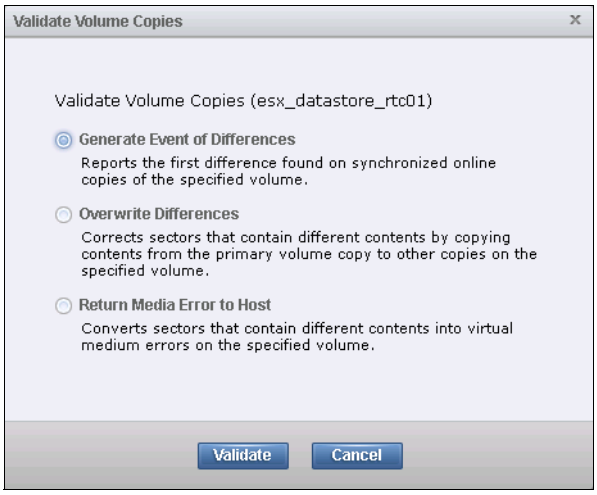

Figure 8-122 Validate Volume Copies window

The following options are available:

- Generate Event of Differences: Use this option only if you want to verify 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 being compared. Upon completion, the command process logs an event, which indicates the number of differences that were corrected. Use this option if you are sure that either the primary volume copy data is correct or that your host applications can handle incorrect data.
- Return Media Error to Host: Use this option to convert sectors on all volumes copies that contain different contents into virtual medium errors. Upon completion, the command logs an event, which indicates the number of differences that were found, the number that were converted into medium errors, and the number that were not converted. Use this option if you are unsure what the correct data is, and you do not want an incorrect version of the data to be used.

2. Select which action to perform and click **Validate** to start the task. The volume is now being checked. Click **Close** (Figure 8-123).

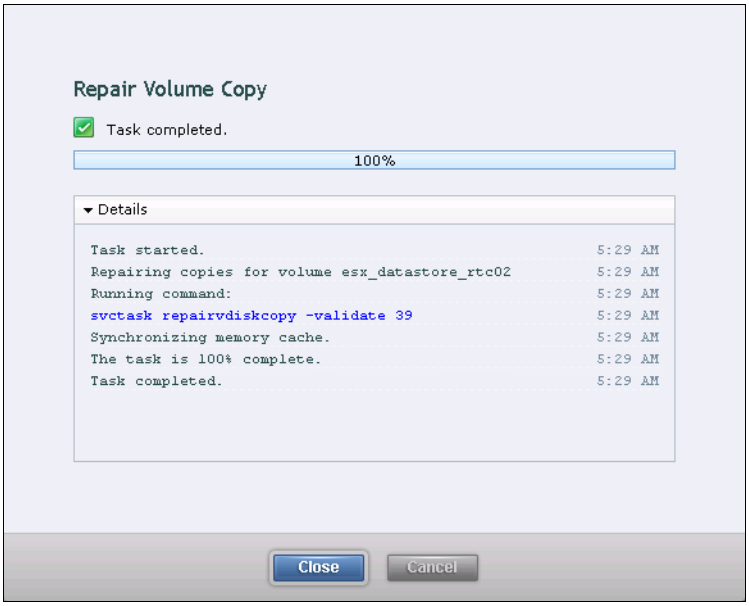

Figure 8-123 Repair Volume Copy task completion window

The validation process takes some time depending on the volume size. You can check the status in the Running Tasks window (Figure 8-124).

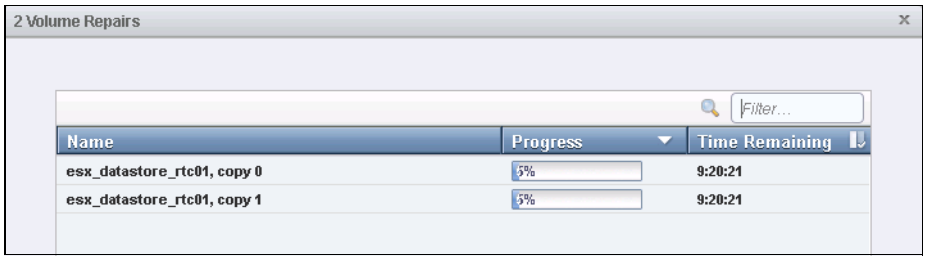

Figure 8-124 Repair Volume Copies: Running tasks

8.6.4 Duplicate Volume copy

You can create a new volume by using the same presets and parameters as an existing volume copy. The parameters are as follows:

- Volume preset (generic, thin-provision, compressed)
- Volume size
- Storage pool
- Access and Caching I/O group
- Caching mode
- Easy Tier status
- Virtualization type

See 8.4.5, “Duplicating a volume” on page 258 for more information.

Note: Duplicating a volume creates a new volume with the same preset and volume parameters as the source volume. Duplicating a volume does not duplicate volume data. Duplicating mirrored and image-mode volumes is not supported. Duplicating a mirrored copy of a volume is supported.

8.6.5 Delete Copy option

To delete a volume copy, select **Delete this Copy** (Figure 8-115 on page 279). The copy is deleted, but the volume remains online by using the remaining copy. Confirm the deletion process by clicking **Yes** (Figure 8-125).

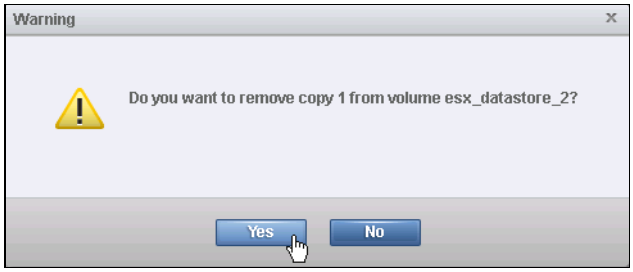

Figure 8-125 Deleting a copy

The copy is deleted. Click **Close** (Figure 8-126) to return to the Volumes panel.

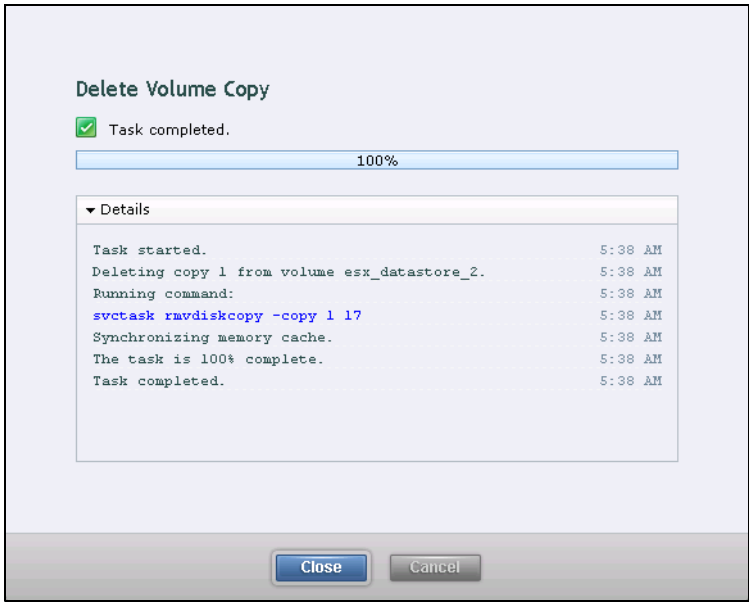

Figure 8-126 Delete Volume Copy task completion window

8.6.6 Migrating volumes using the volume copy features

The previous sections show that creating, synchronizing, splitting, and deleting volume copies are possible. 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 described in 8.4.9, “Migrating a volume to another storage pool” on page 261. If you use this feature, one extent after another is migrated to the new storage pool. However, using volume copies provides another possibility to migrate volumes. To migrate a volume, complete the following steps:

1. Create a second copy of your volume in the target storage pool (see 8.5.4, “Adding a mirrored volume copy” on page 273).
2. Wait until the copies are synchronized.
3. Change the role of the copies and make the new copy of the primary copy (see 8.6, “Advanced volume copy functions” on page 279).
4. Split or delete the old copy from the volume (see 8.6.2, “Splitting into a new volume” on page 282 or 8.6.5, “Delete Copy option” on page 285).

This migration process requires more user interaction, but it offers some benefits, for example, if you migrate 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; all reads are still performed in the tier 1 pool to the primary copy. After the synchronization, all writes are destaged to both pools, but the reads are still only done from the primary copy. Now you can switch the role of the copies online (step 3), and test the performance of the new pool. If you are done testing your lower performance pool, you can split or delete the old copy in tier 1, or switch back to tier 1 in seconds, because the tier 2 pool did not meet your performance requirements.

8.7 Volumes by storage pool

To see an overview of which volumes are in which storage pool, hover the mouse cursor over the Volumes selection in the IBM Storwize V7000 dynamic menu and click **Volumes by Pool** (Figure 8-65 on page 252). The Volumes by Pool panel opens (Figure 8-127).

Figure 8-127 Volumes by Pool panel

The left pane lists all of your existing storage pools. You can filter them by using the *Pool Filter*. More details about storage pools are in Chapter 7, “Storage pools” on page 179.

In the upper right, you see information about the pool that you selected in the pool filter, and the following information is also there:

- Pool Icon: Storage Pools have different characteristics. It is possible to change the pool function icon to identify the pool type (see 8.7.1, “Changing the Storage Pool function icon” on page 289).
- Pool Name: This name is the name of the storage pool given during creation. It is possible to change this name from this window (see 8.7.2, “Changing the storage pool name” on page 290).
- Pool Details: Gives you status information about the pool, such as the number of MDisks and volume copies, and the Easy Tier status.
- Volume allocation: Provides details about the available, allocated, compressed and virtual space in this pool.

The lower right lists all volumes that have at least one copy in this storage pool and provides the following information about them:

- Name: Shows the name of the volume.
- State: Provides state information about the volume.
- Capacity: Shows the capacity that is presented to the host of the volume. A *green* volume sign next to it means that this volume is thin-provisioned and the capacity shown is the virtual size. Similarly, a *zipped* icon next to the volume means the volume is compressed using the Real-time Compression feature.
- Host Mappings: Shows whether (at least) one host mapping exists.
- UID: Shows the Volume unique identifier.

You may also create a volume in the selected pool. Click **Create Volume**. The wizard starts (described in 5.1, “Provisioning storage from IBM Storwize V7000 to hosts” on page 132).

If you select a volume and click **Actions**, or right-click the volume (Figure 8-128), the same options from the Volumes panel (described in 8.4.1, “Advanced volume functions” on page 253) are listed. Detailed instructions about each task are covered in that section.

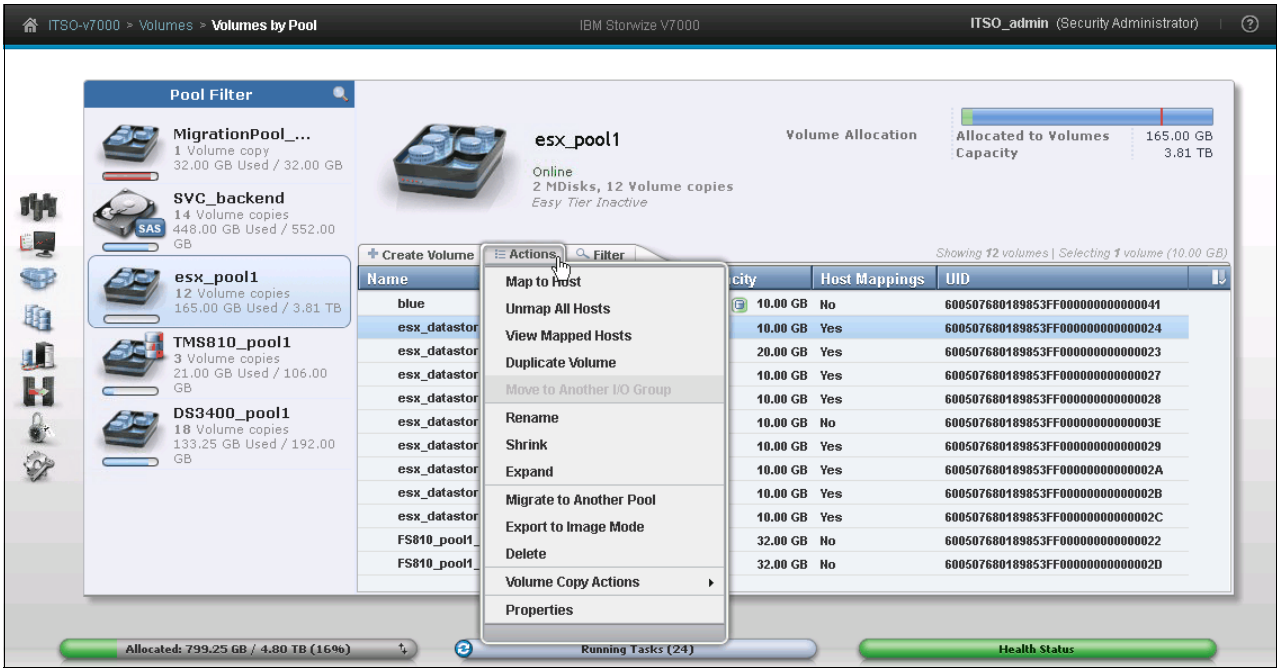

Figure 8-128 Volumes by Pool actions

If you select a volume copy and click **Actions**, or right-click the volume (Figure 8-129), the same options from the Volumes panel appear. Detailed instructions about each task are covered 8.6, “Advanced volume copy functions” on page 279.

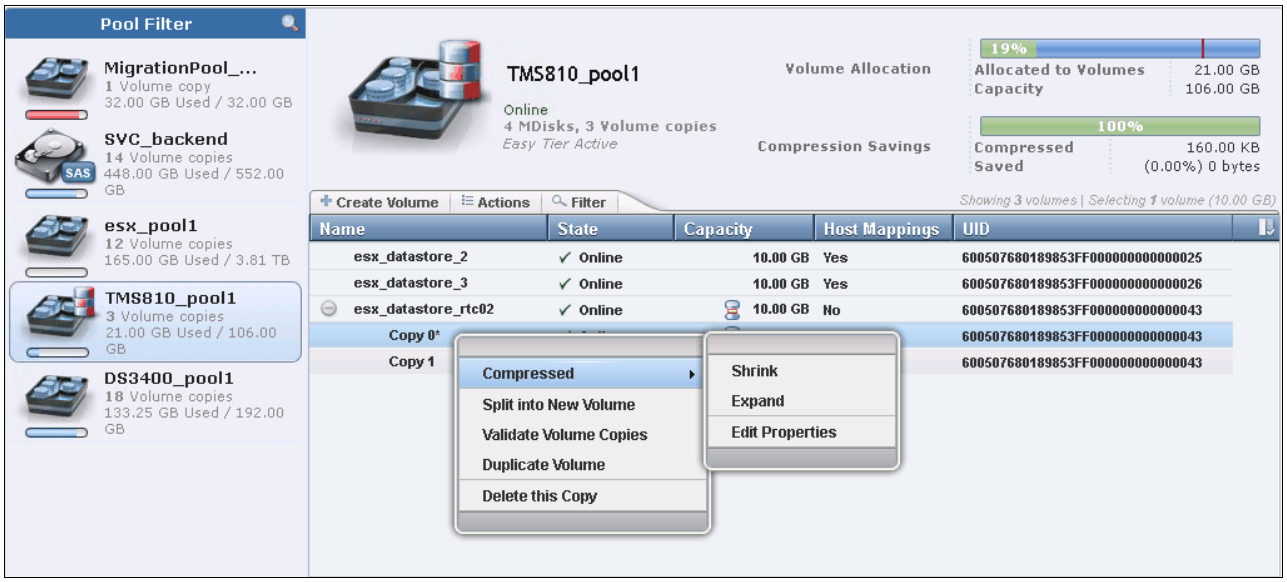

Figure 8-129 Volume copy by pool actions

8.7.1 Changing the Storage Pool function icon

Usually, storage pools have different characteristics; you can change the Storage Pool function icon to reflect these differences. Click the pool function icon; the Change Icon window opens (Figure 8-130).

Figure 8-130 Change Icon window

Use the left and right arrows to select a new function icon (Figure 8-131). Several options are available. Click **OK**.

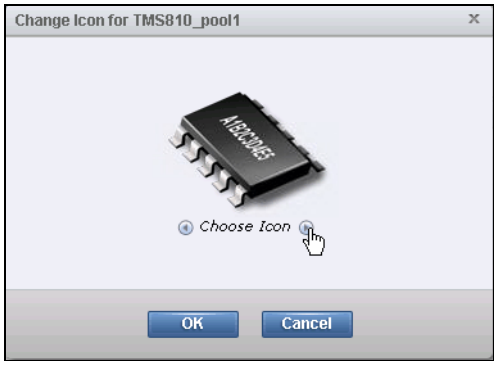

Figure 8-131 Select a new function icon

The changes are applied to the system (Figure 8-132). Click **Close**.

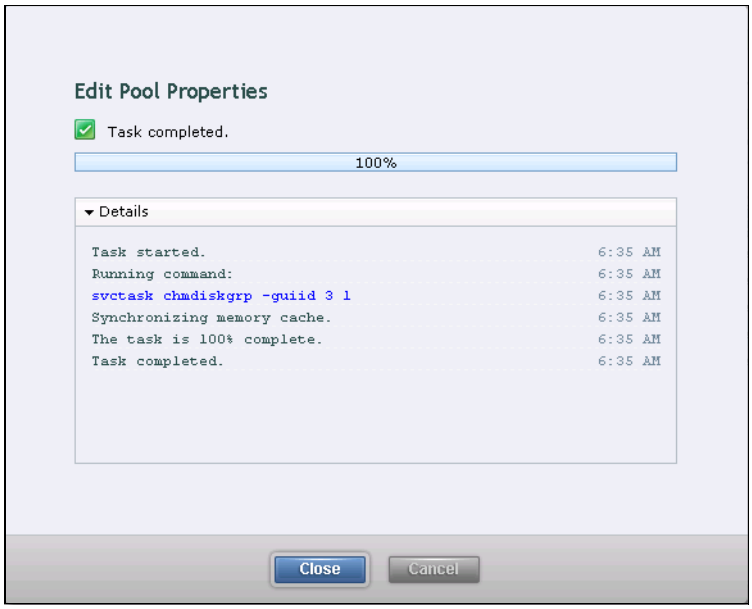

Figure 8-132 Edit Pool Properties task completion window

The function icon has been changed to so that identifying the pool is easier (Figure 8-133).

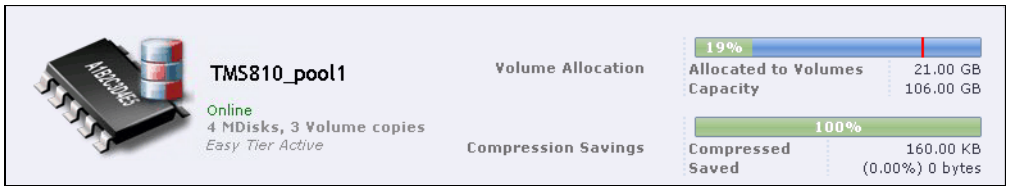

Figure 8-133 Changed pool icon

8.7.2 Changing the storage pool name

To change the given name of a storage pool, click the name; you can then edit it (Figure 8-134).

Figure 8-134 Change pool name

Enter a new name (Figure 8-135) and press Enter.

Figure 8-135 Type new pool name

The changes are applied to the system (Figure 8-136).

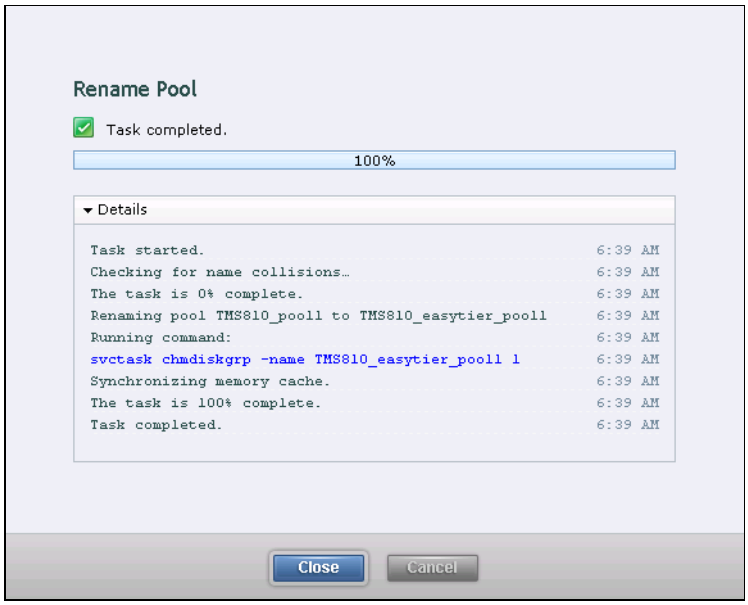

Figure 8-136 Rename pool task completion window

The name for the storage pool is now changed (Figure 8-137).

Figure 8-137 Change pool name: new name

8.8 Volumes by Host

To see an overview of which volumes a host is able to access, hover the mouse cursor over Volumes in the IBM Storwize V7000 dynamic menu and click **Volumes by Host** (Figure 8-65 on page 252). The Volumes by Host panel opens (Figure 8-138).

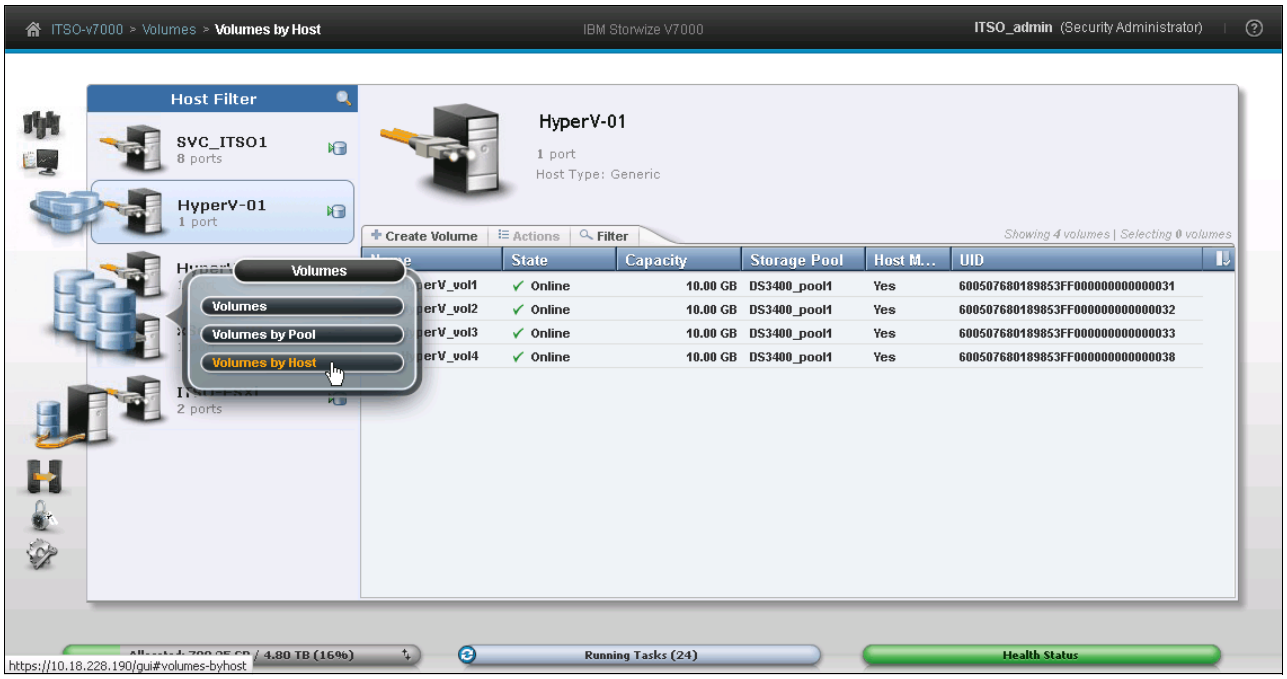

Figure 8-138 Volumes by Host panel

The left pane lists all host objects that are defined on the IBM Storwize V7000 system. You can use *Host Filter* to search for a specific host. When you select a host, its properties are listed in the right pane of the view. The hosts with the *orange* cable represent Fibre Channel hosts; the *blue* cable represents iSCSI hosts. The upper right side shows the host function icon, the host name, the number of host ports, and the host type. The volumes that are mapped to this host are listed in the pane below it.

If you want to create a volume for this host, click **Create Volume**. The wizard starts (this is the same wizard described in 5.1, “Provisioning storage from IBM Storwize V7000 to hosts” on page 132).

Select a volume and click **Actions** or right-click the volume (Figure 8-139); the same options from the Volumes panel appear. Details about each task are covered in 8.4.1, “Advanced volume functions” on page 253.

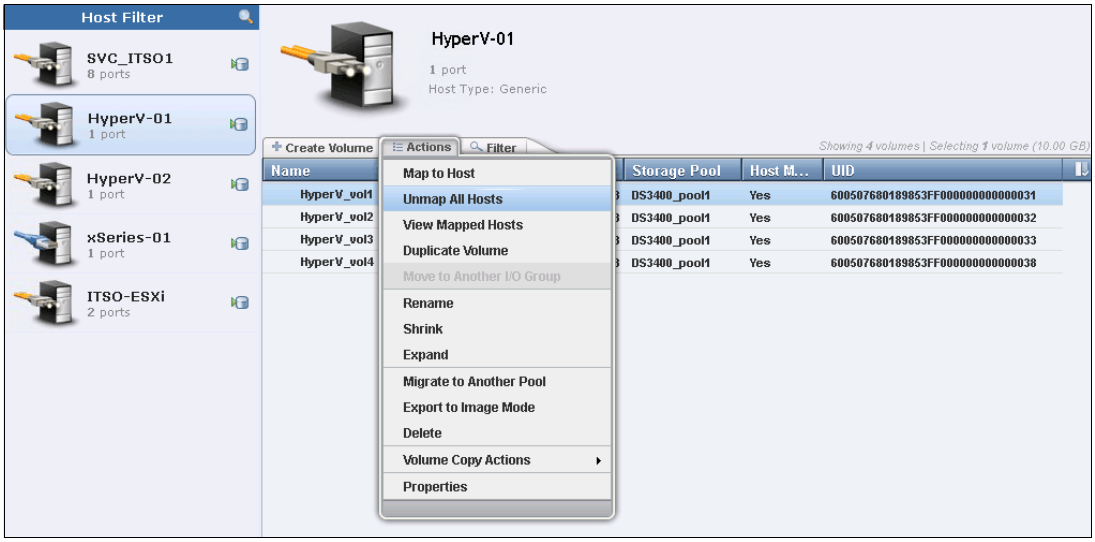

Figure 8-139 Volume actions

If the volume owns more than one copy, you can select a volume copy and click **Actions** or right-click the copy (Figure 8-140), and the same options from the Volumes panel appear. Details about each task are covered in 8.6, “Advanced volume copy functions” on page 279.

Figure 8-140 Volume copy actions

8.8.1 Renaming a host

To rename a host in the Volumes by Host panel, click it and then edit the name (Figure 8-141).

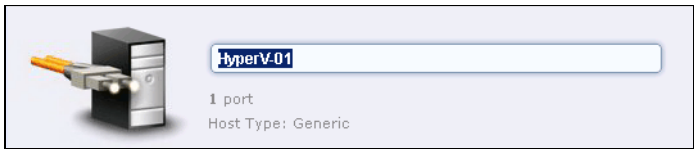

Figure 8-141 Rename a host

Enter a new name (Figure 8-142) and press Enter.

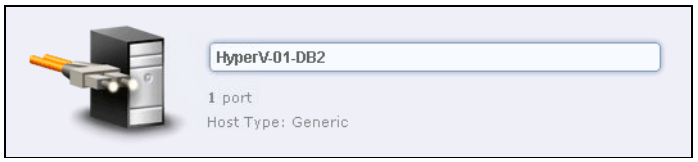

Figure 8-142 Enter a new name

The changes to the system (Figure 8-143).

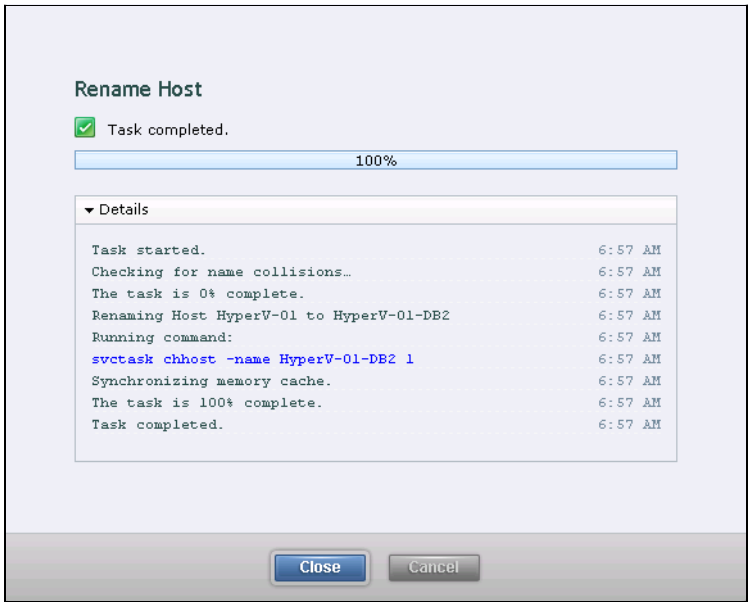

Figure 8-143 Rename Host task completion window

Click **Close** to return to the Volumes by Host panel (Figure 8-144).

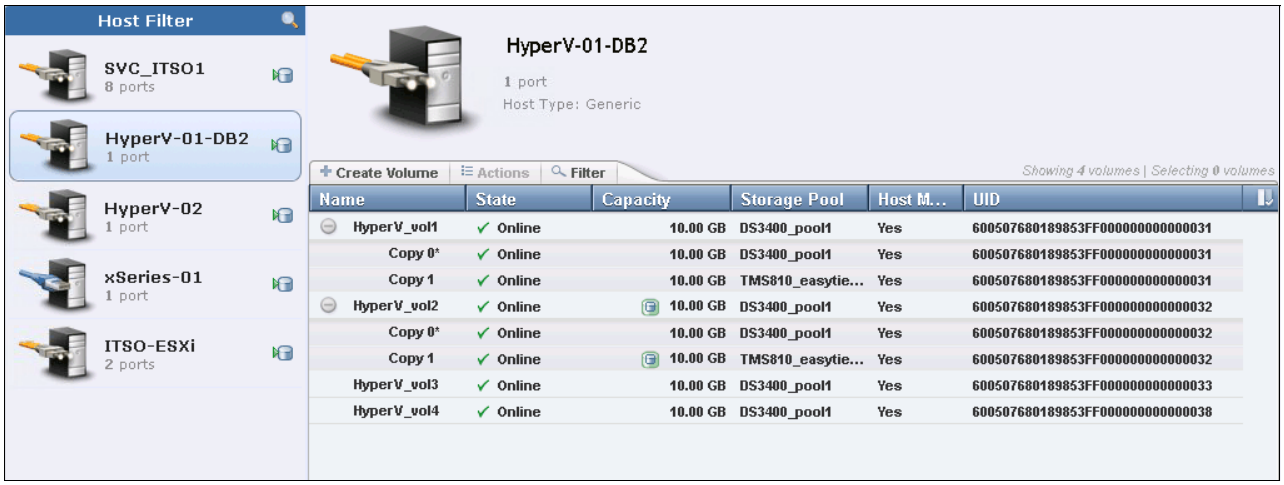

Figure 8-144 Volumes by Host

9

External storage virtualization

In this chapter, we describe how to incorporate external storage systems into the virtualized world of the IBM Storwize V7000. A key feature of IBM Storwize V7000 is its ability to consolidate disk controllers from various vendors into pools of storage. In this way, the storage administrator can, from a single user interface, manage and provision storage to applications, and use a common set of advanced functions across all the storage systems under the control of the IBM Storwize V7000.

9.1 Planning for external storage virtualization

In this section, we describe how to plan for virtualizing external storage with IBM Storwize V7000. Virtualizing the storage infrastructure with IBM Storwize V7000 makes your storage environment more flexible, cost-effective, and easy to manage. The combination of IBM Storwize V7000 and an external storage system allows more storage capacity benefits from the powerful software function within the IBM Storwize V7000.

The external storage systems that are incorporated into the IBM Storwize V7000 environment can new or existing systems. The data on existing storage systems can be easily migrated to the IBM Storwize V7000 managed environment, as described in Chapter 6, “Migration wizard” on page 167 and Chapter 7, “Storage pools” on page 179.

The following web page has the latest supported hardware list, device drivers, firmware, and recommended software levels for IBM Storwize V7000:

http://www.ibm.com/support/docview.wss?uid=ssg1S1004450#_Controllers

9.1.1 License for external storage virtualization

From a licensing standpoint, when external storage systems are virtualized by IBM Storwize V7000, a per-enclosure External Virtualization license is required. You can find more licensing information in 1.3.4, “IBM Storwize V7000 licensing” on page 9, and contact your IBM account team or IBM Business Partner for further assistance, if required.

Migration: If the IBM Storwize V7000 is used as a general migration tool, then the appropriate External Virtualization licenses must be ordered. The only exception is if you want to migrate existing data from external storage systems to IBM Storwize V7000 internal storage, as you can temporarily configure your External Storage license within 45 days. For a more than 45 day migration requirement from external storage to IBM Storwize V7000 internal storage, an appropriate External Virtualization license must be ordered.

You can configure the IBM Storwize V7000 licenses by going to the Settings section in the IBM Storwize V7000 dynamic menu and then clicking **General** (Figure 9-1).

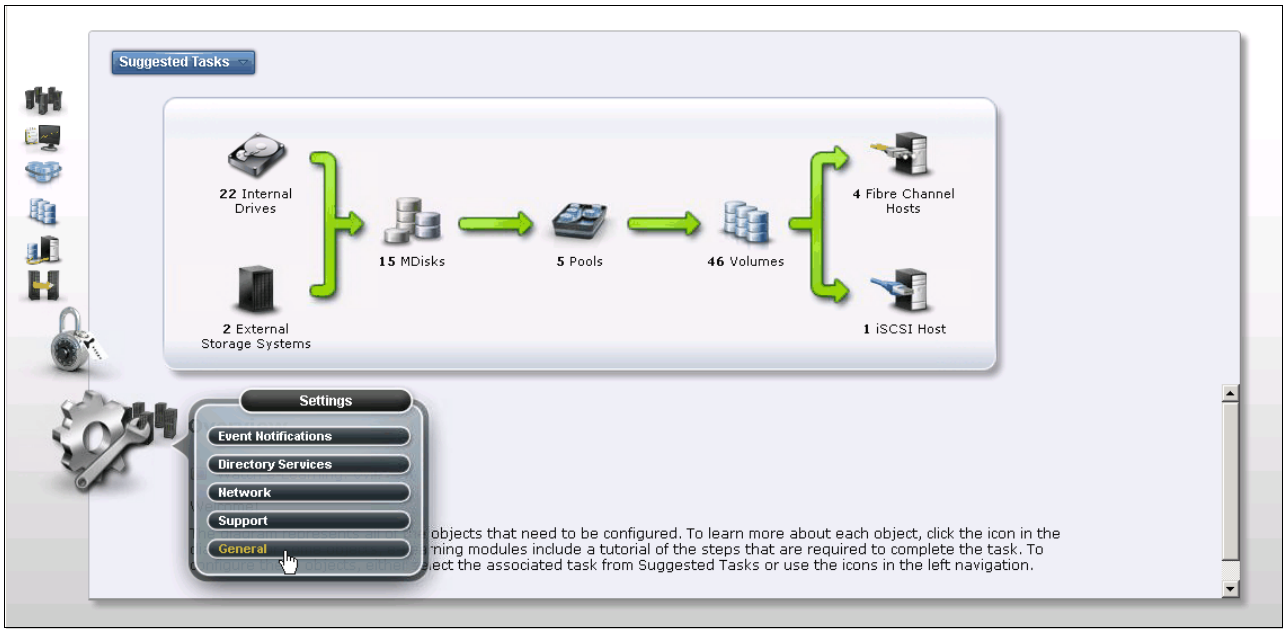

Figure 9-1 Settings, General panel

In the General pane on the left, click **Licensing**. The Update License view opens in the right pane (Figure 9-2). This Update License view shows three license options that you can set. Set these options to the limit you obtain from IBM:

- External Virtualization
- Remote Copy
- Real-time Compression

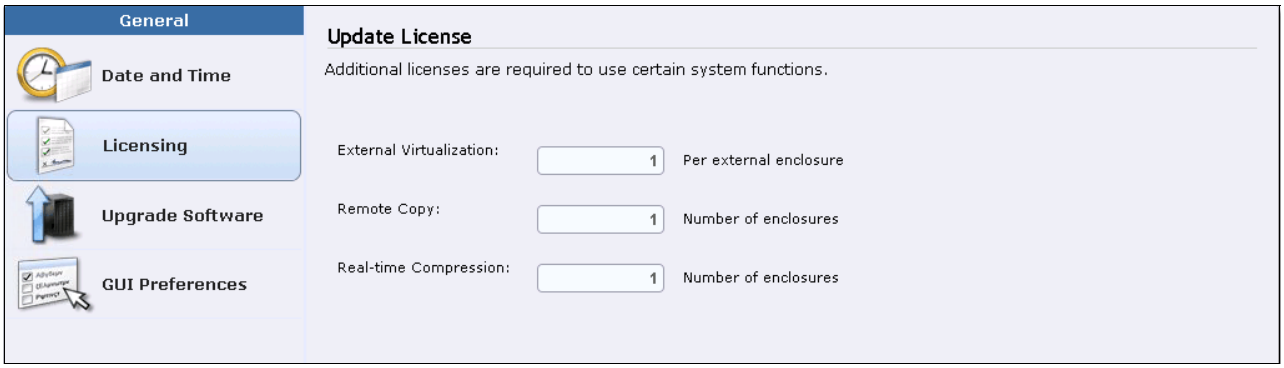

Figure 9-2 Update License window

For assistance with licensing questions or to purchase an External Virtualization, Remote Copy, or Real-time Compression license, contact your IBM account team or IBM Business Partner.

9.1.2 SAN configuration planning

External storage controllers virtualized by IBM Storwize V7000 must be connected through SAN switches. A direct connection between the IBM Storwize V7000 and storage controllers or hosts ports is not supported.

Make sure that the switches or directors are at the firmware levels supported by the IBM Storwize V7000 and that the IBM Storwize V7000 port login maximums listed in the restriction document will not be exceeded. The configuration restrictions can be found by navigating to the Support home page at the following address:

http://www.ibm.com/support

The recommended SAN configuration is composed of a minimum of two fabrics. The ports on external storage systems that will be virtualized by the IBM Storwize V7000 and the IBM Storwize V7000 ports are evenly split between the two fabrics to provide redundancy if one of the fabrics goes offline.

After the IBM Storwize V7000 and external storage systems are connected to the SAN fabrics, SAN zoning must be implemented. In each fabric, create a zone with the four IBM Storwize V7000 port WWPNs, two from each node canister, along with up to a maximum of eight port WWPNs from each external storage system.

Ports: IBM Storwize V7000 supports a maximum of 16 ports or WWPNs from a given external storage system that will be virtualized.

Figure 9-3 is an example of how to cable the devices to the SAN. Refer to this example as we describe the zoning.

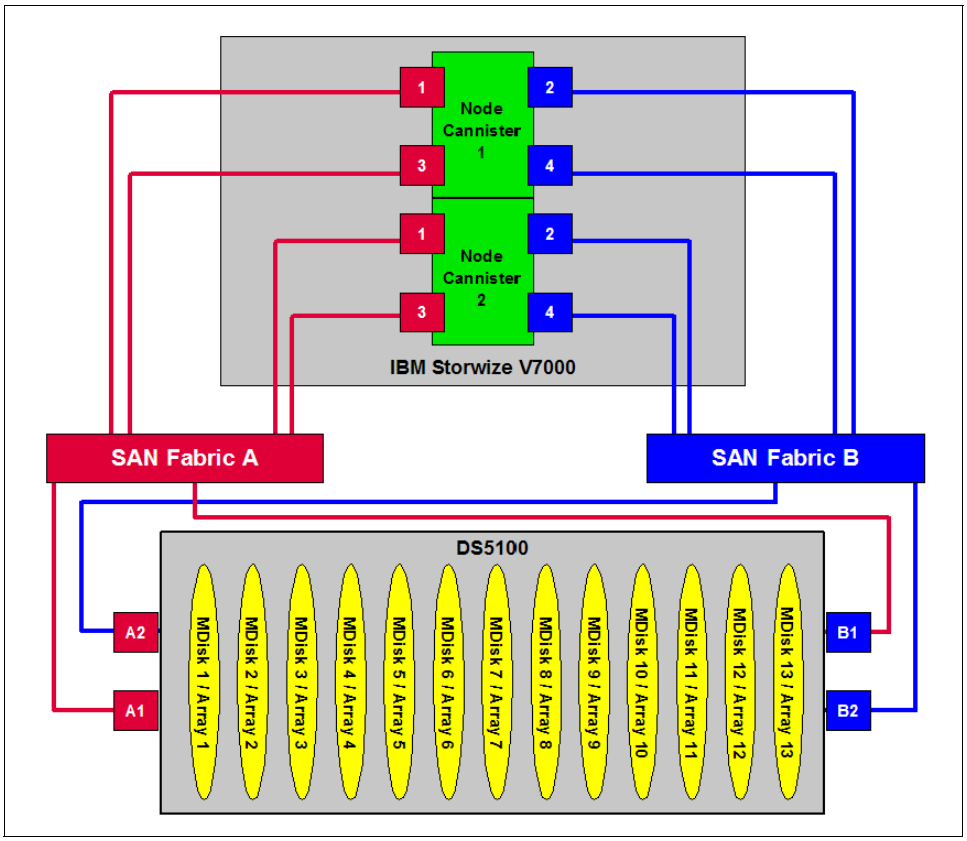

Figure 9-3 SAN cabling and zoning example diagram

Create an IBM Storwize V7000 and external storage zone for each storage system to be virtualized as in the following example:

- Zone DS5100 controller ports A1 and B1 with all node ports 1 and 3 in the RED fabric
- Zone DS5100 controller ports A2 and B2 with all node ports 2 and 4 in the BLUE fabric

9.1.3 External storage configuration planning

External storage systems provide redundancy through various RAID levels, which prevents a single physical disk failure from causing an MDisk, storage pool, or associated host volume, from going offline. To minimize the risk of data loss, only virtualize storage systems where LUNs are configured using a RAID level other than RAID 0 (for example RAID 1, RAID 10, RAID 0+1, RAID 5, or RAID 6).

Verify that the storage controllers to be virtualized by IBM Storwize V7000 meet the requirements, and that you found the configuration restrictions by navigating at the Support Portal website:

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

Make sure that the firmware or microcode levels of the storage controllers to be virtualized are supported by IBM Storwize V7000.

IBM Storwize V7000 must have exclusive access to the LUNs from the external storage system mapped to it. LUN access cannot be shared between IBM Storwize V7000s or

between an IBM Storwize V7000 and other storage virtualization platforms or between an IBM Storwize V7000 and hosts. However, different LUNs could be mapped from one external storage system to an IBM Storwize V7000 and other hosts in the SAN through different storage ports.

Make sure to configure the storage subsystem LUN masking settings to map all LUNs to all the WWPNs in the IBM Storwize V7000 storage system.

Be sure to go to the IBM Storwize V7000 page and review the *Configuring and servicing external storage system* topic before you prepare the external storage systems for discovery by the IBM Storwize V7000 system. This website is at the following location:

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

You may also go to this link:

`http://www-01.ibm.com/support/knowledgecenter/ST3FR7_7.2.0/com.ibm.storwize.v7000.720.doc/svc_configdiskcontrollersovr_22n9uf.html?lang=en`

9.1.4 Guidelines for virtualizing external storage

When virtualizing external storage with the IBM Storwize V7000, follow these guidelines:

- Avoid splitting arrays into multiple LUNs at the external storage system level. When possible, create a single LUN per array for mapping to the IBM Storwize V7000.
- Except for Easy Tier, do not mix MDisks that vary in performance or reliability in the same storage pool. Always put similarly sized MDisks into one storage pool. For more information about Easy Tier, see Chapter 10, “Advanced features for storage efficiency” on page 309.
- Do not leave volumes in image mode. Only use image mode to import or export existing data into or out of the IBM Storwize V7000. Migrate such data from image mode MDisks to other storage pools to benefit from storage virtualization.
- Using the copy services in Storwize V7000 gives you a unified method to manage data integrity across heterogeneous storage systems.
- The Easy Tier function is included with the IBM Storwize V7000 system, and the external storage system could benefit from this powerful storage tiering function to remove hot spots and improve overall performance.

9.2 Working with external storage

In this section, we describe how to manage external storage using an IBM Storwize V7000.

The basic concepts of managing external storage system are the same as internal storage. IBM Storwize V7000 discovers LUNs from the external storage system as one or more MDisks. These MDisks are ultimately added to a storage pool in which volumes are created and mapped to hosts as needed.

9.2.1 Adding external storage

To add new external storage systems to the IBM Storwize V7000 virtualized environment, complete the following steps:

1. Zone a minimum of two and a maximum of 16 Fibre Channel ports from the external storage system with all eight Fibre Channel ports on the IBM Storwize V7000 system. As a preferred practice, have two fabrics for redundancy in the SAN. Then in each fabric, zone two ports from each node canister in the IBM Storwize V7000 system with half the ports from the external system. As the IBM Storwize V7000 is virtualizing your storage, hosts should be zoned with the V7000 controller port WWPNs.
2. Using the storage partitioning or LUN masking feature of the external storage system, create a group that includes all eight IBM Storwize V7000 WWPNs.
3. Create equal size arrays on the external system using any RAID level but zero.
4. Create a single LUN per RAID array.
5. Map the LUNs to all eight Fibre Channel ports on the IBM Storwize V7000 system by assigning them to the group created in step 2.
6. Verify that IBM Storwize V7000 discovered the LUNs as unmanaged MDisks. If they do not show up automatically, click the **Detect MDisk** option from the MDisk window of the GUI, as described in 7.2.5, “Additional actions on MDisks” on page 213. You should see the MDisks mapped to the IBM Storwize V7000 under the respective Storage system.
7. Select the storage tier for the MDisks, as described in 7.2.4, “Selecting the tier for MDisks” on page 211.
8. Create a storage pool, as described in 7.3, “Working with storage pools” on page 217.
9. Add the MDisks to the pool, as described in 7.2.1, “Adding MDisks to storage pools” on page 199.
10. Create volumes and map them to hosts as needed, as described in Chapter 4, “Host configuration” on page 105 and Chapter 5, “Basic volume configuration” on page 131.

If the external storage systems are not the new systems, that is, existing data on the LUNs must be kept after virtualization, complete the steps in Chapter 6, “Migration wizard” on page 167 to prepare the environment. Then, you can migrate the existing data (with or without using wizard) to IBM Storwize V7000 internal storage or another external storage system.

To manually import MDisks and migrate the data to other storage pools, see 7.2.2, “Importing MDisks” on page 203. Whether you migrate the data with the wizard or not, you can select your destination storage pools to be internal storage pools or external storage pools.

9.2.2 Managing external storage

The IBM Storwize V7000 provides an individual External Storage panel for managing external storage systems.

You can access the External Storage pane by opening the IBM Storwize V7000 Home pane and clicking the External Storage System function icon. Extended help information for external storage is displayed. Click **Pools** and the External Storage pane opens.

Figure 9-4 shows how to access the External Storage pane from the IBM Storwize V7000 Home pane.

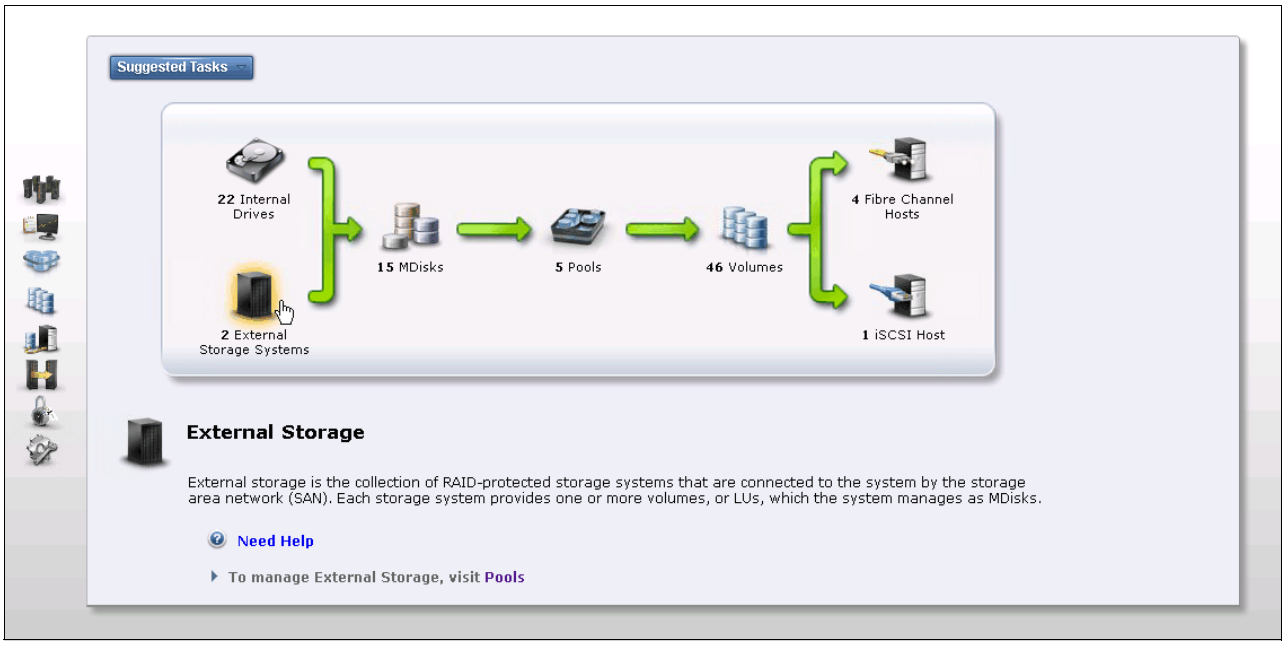

Figure 9-4 Access the External Storage pane from the IBM Storwize V7000 Home pane

The other method to access the External Storage pane is to use the Pools selection from the IBM Storwize V7000 dynamic menu that is shown in the left pane of the GUI (Figure 9-5).

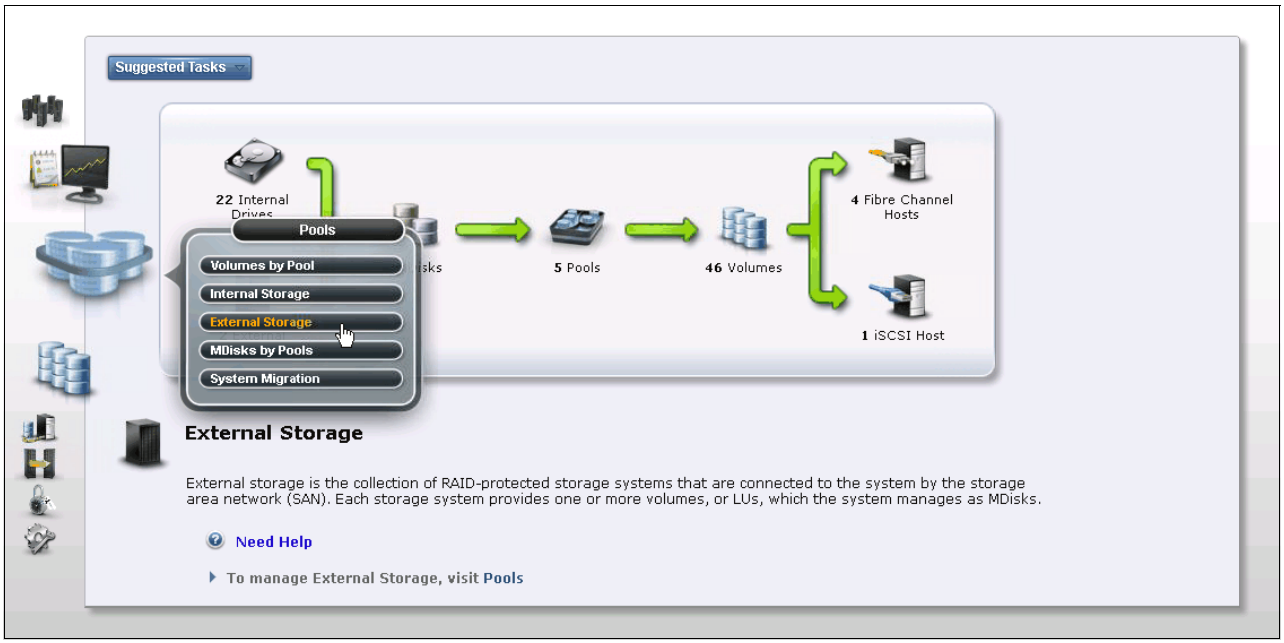

Figure 9-5 Access the External Storage pane from the IBM Storwize V7000 dynamic menu

The External Storage window (Figure 9-6) shows an overview of all your external storage systems. To see a list of managed disks (MDisks) presented by each storage system to your IBM Storwize V7000, expand each external storage system by clicking the plus sign (+).

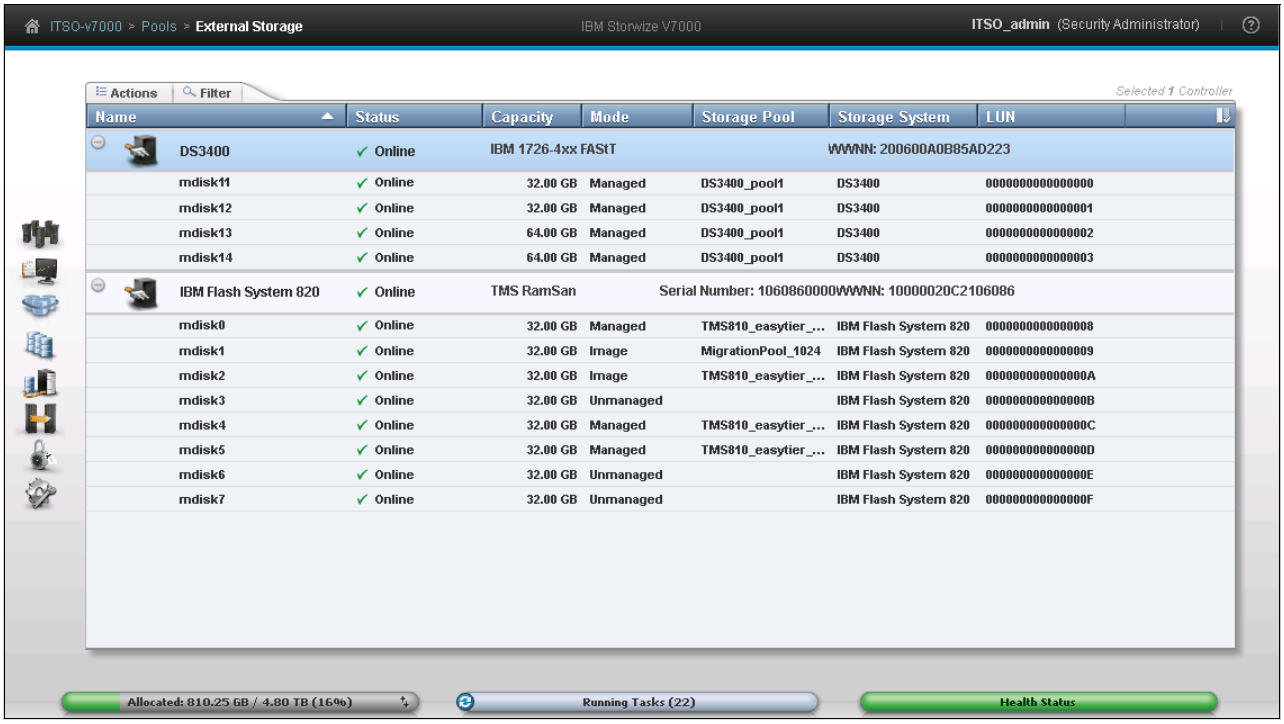

Figure 9-6 External Storage window

You can change the name of an external storage system by selecting the storage system and clicking **Actions** or right-clicking the storage system. The status of the storage system, its type, and WWPN are also listed in the table.

If you select the storage system and click **Actions**, you can do these tasks (Figure 9-7):

- Show Dependent Volumes
- Rename the storage system

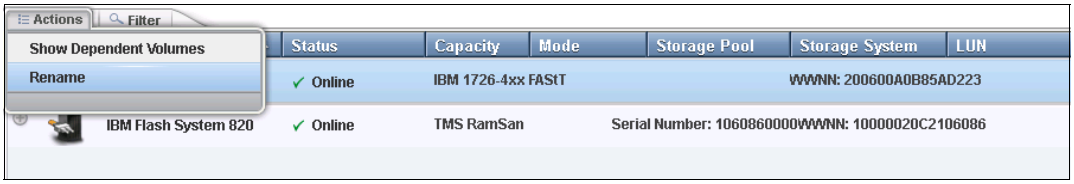

Figure 9-7 External Storage panel, the Actions drop-down menu

Click **Show Dependent Volumes** to view all the volumes located in this external storage system (Figure 9-8).

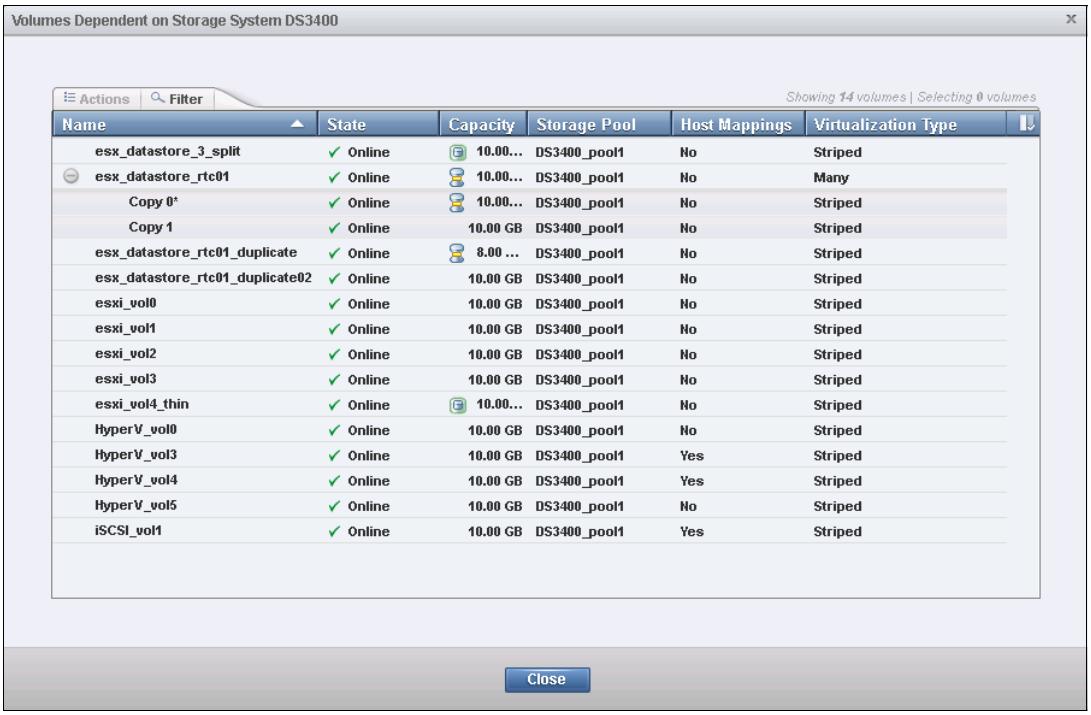

Figure 9-8 Volumes dependent on external storage

If you click **Actions** in this window (Figure 9-9), you can do volume tasks, such as Map to Host, Shrink, Expand, Migrate to Another Pool, Volume Copy Actions, and other actions that are described in Chapter 8, “Advanced host and volume administration” on page 221.

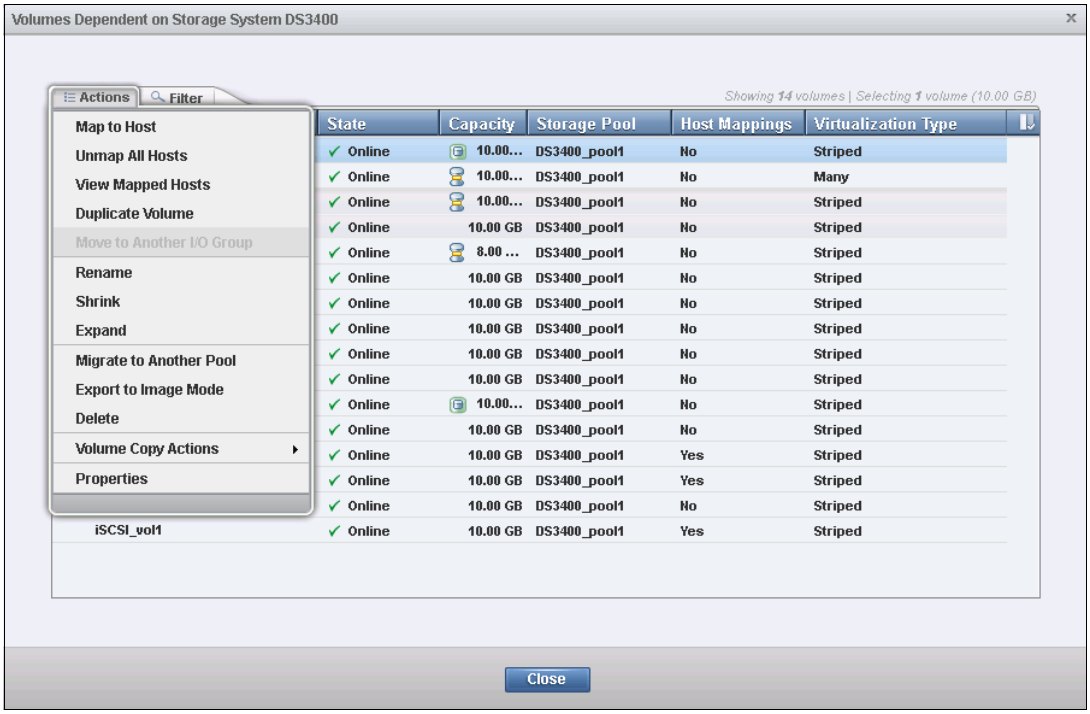

Figure 9-9 Actions you can take with dependent volumes

One of the features of the IBM Storwize V7000 storage system is that it can be used as a data migration tool. In the IBM Storwize V7000 virtualization environment, you can migrate your application data nondisruptively from one internal or external storage system to another storage system, which makes storage management much simpler and reduces risk.

Volume copy is another key feature that you can benefit from by using IBM Storwize V7000 virtualization. Two copies could be applied to your data to enhance the availability for a critical application. A volume copy could be also used to generate test data or data migration.

Chapter 8, “Advanced host and volume administration” on page 221 provides more information about the volume actions of the IBM Storwize V7000 storage system.

Returning to the External Storage window, you can perform various MDisk tasks. In the list, you can find the name of an MDisk, its capacity and mode, the storage pool, and the storage system it belongs to in the list. The actions on MDisks can also be made through the Actions menu (Figure 9-10), including Detect MDisks, Add to Pool, Import, and so on. This menu is the same as the one in the MDisks window, which is described in 7.2, “Working with MDisks” on page 197.

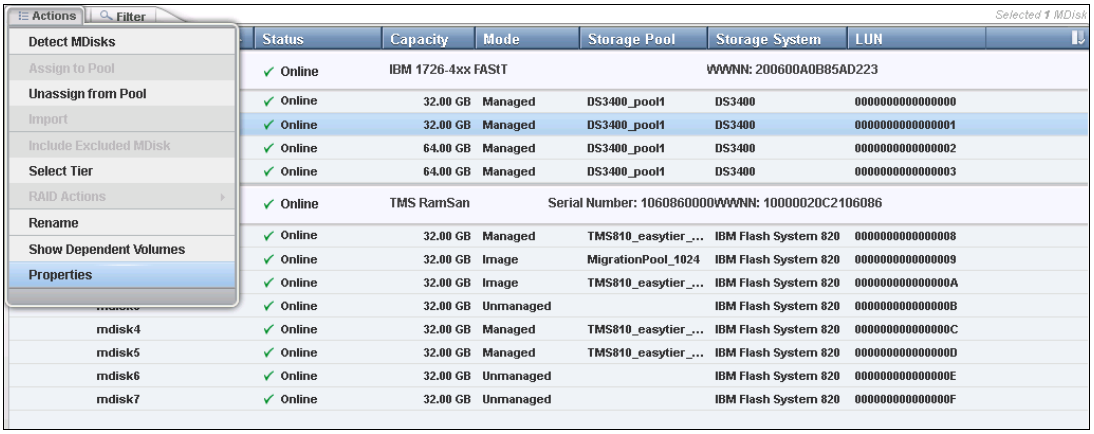

Figure 9-10 MDisk menu in the External Storage panel

9.2.3 Removing external storage

If you want to remove the external storage systems from the IBM Storwize V7000 virtualized environment, you have several options:

- If you want to remove the external storage systems and discard the data on it, complete the following steps:
  a. Stop any host I/O on the volumes.
  b. Remove the volumes from the host file system, logical volume, or volume group, and remove the volumes from the host device inventory.
  c. Remove the host mappings of volumes and the volumes themselves on IBM Storwize V7000, as described in 8.1.1, “Modifying Mappings menu” on page 224.
  d. Remove the storage pools to which the external storage systems belong, as described in 7.3, “Working with storage pools” on page 217, or you can keep the storage pool and remove the MDisks of the external storage from the storage pools, as described 7.2.1, “Adding MDisks to storage pools” on page 199.
  e. Unzone and disconnect the external storage systems from the IBM Storwize V7000.
  f. Click **Detect MDisks** to make IBM Storwize V7000 discover the removal of the external storage systems, as described in 7.2.5, “Additional actions on MDisks” on page 213.

- If you want to remove the external storage systems and keep the volumes and their data on the IBM Storwize V7000, complete the following steps:
  a. Migrate volumes and their data to the other storage pools that are on IBM Storwize V7000 internal storage or other external storage systems, as described in 5.3, “Discovering volumes on hosts and multipathing” on page 143.
  b. Remove the storage pools to which the external storage systems belong, or you can keep the storage pools and remove the MDisks of the external storage from the storage pools, as described in Chapter 7, “Storage pools” on page 179.
  c. Unzone and disconnect the external storage systems from the IBM Storwize V7000.
  d. Click **Detect MDisks** to make IBM Storwize V7000 discover the removal of the external storage systems, as described in 7.2.5, “Additional actions on MDisks” on page 213.
- If you want to remove the external storage systems from IBM Storwize V7000 control and keep the volumes and their data on external storage systems, complete the following steps:
  a. Migrate volumes and their data to the other storage pools that are on IBM Storwize V7000 internal storage or other external storage systems, as described in Chapter 5, “Basic volume configuration” on page 131.
  b. Remove the storage pools to which the external storage systems belong, or you can keep the storage pools and remove the MDisks of the external storage from the storage pools, as described in Chapter 7, “Storage pools” on page 179.
  c. Export volumes to image mode with the MDisks on the external storage systems. The restrictions and prerequisites for migration can be found in Chapter 6, “Migration wizard” on page 167. You also need to record pre-migration information, for example, the original SCSI IDs the volumes used to be mapped to hosts. Some operating systems do not support changing the SCSI ID during the migration. More information about migration is at the IBM Storwize V7000 welcome page:

     http://pic.dhe.ibm.com/infocenter/storwize/ic/index.jsp
  d. Unzone and disconnect the external storage systems from the IBM Storwize V7000.
  a. Click **Detect MDisks** to make IBM Storwize V7000 discover the removal of the external storage systems, as described Chapter 7, “Storage pools” on page 179.

10

Advanced features for storage efficiency

In this chapter, we introduce the basic concepts of dynamic data relocation and storage optimization features. The IBM SAN Volume Controller family of products (including IBM Storwize V7000) offers software functions for storage efficiency:

- Easy Tier
- Thin Provisioning
- Real-time Compression

We provide only basic technical overview and benefits of each feature. For more details about planning and configuration, see the following resources:

- Easy Tier
  - *Implementing IBM Easy Tier with IBM Real-time Compression*, TIPS1072
  - *IBM System Storage SAN Volume Controller and Storwize V7000 Best Practices and Performance Guidelines*, SG24-7521
  - *IBM DS8000 Easy Tier*, REDP-4667 (the concept is similar to SAN Volume Controller Easy Tier)
- Thin Provisioning
  - *Thin Provisioning in an IBM SAN or IP SAN Enterprise Environment*, REDP-4265
  - *DS8000 Thin Provisioning*, REDP-4554 (the concept is similar to SAN Volume Controller thin provisioning)
- Real-time Compression
  - *Real-time Compression in SAN Volume Controller and Storwize V7000*, REDP-4859
  - *Implementing IBM Real-time Compression in SAN Volume Controller and IBM Storwize V7000*, TIPS1083
  - *Implementing IBM Easy Tier with IBM Real-time Compression*, TIPS1072

10.1 Introduction

In modern and complex application environments, the increasing and often unpredictable demands for storage capacity and performance lead to issues of planning and optimization of storage resources.

Consider the following typical storage management issues:

- Usually when a storage system is implemented, only a portion of the configurable physical capacity is deployed. When the storage system runs out of the installed capacity and more capacity is requested, a hardware upgrade is implemented to add new physical resources to the storage system. This new physical capacity hardly can be configured to keep an even spread of the overall storage resources. Typically, the new capacity is allocated to fulfill only new storage requests. The existing storage allocations do not benefit from the new physical resources. Similarly, the new storage requests do not benefit from the existing resources; only new resources are used.
- In a complex production environment, optimizing storage allocation for performance is not always possible. The unpredictable rate of storage growth and the fluctuations in throughput requirements, which are I/O per second (IOPS), often lead to inadequate performance. Furthermore, the tendency to use even larger volumes to simplify storage management works against the granularity of storage allocation, and a cost-efficient storage tiering becomes difficult to achieve. With the introduction of high performing, but expensive, technologies such as solid-state drives (SSD), this challenge becomes even more important.
- The move to larger physical disk drive capacities means that previous access densities that were achieved with low-capacity drives can no longer be sustained.
- Any business has applications that are more critical than others, and there is a need for specific application optimization. Therefore, there is a need for the ability to relocate specific application data to faster storage media.
- Although more servers are purchased with local SSDs attached for better application response time, the data distribution across these direct-attached SSDs and external storage arrays must be carefully addressed. An integrated and automated approach is crucial to achieve performance improvement without compromise to data consistency, especially in a disaster recovery situation.

All of these issues deal with data placement and relocation capabilities or data volume reduction. Most of these challenges can be managed by having spare resources available and by moving data, the use of data mobility tools or operating systems features (such as host level mirroring), to optimize storage configurations. However, all of these corrective actions are expensive in terms of hardware resources, labor, and service availability. Relocating data among the physical storage resources dynamically or effectively reducing the amount of data, that is, transparently to the attached host systems, is becoming increasingly important.

10.2 Easy Tier

In today's storage market, solid-state drives (SSDs) are emerging as an attractive alternative to hard disk drives (HDDs). Because of their low response times, high throughput, and IOPS-energy-efficient characteristics, SSDs have the potential to allow your storage infrastructure to achieve significant savings in operational costs. However, the current acquisition cost per GB for SSDs is much higher than for HDDs. SSD performance depends a lot on workload characteristics, so SSDs need to be used with HDDs. It is critical to choose the right mix of drives and the right data placement to achieve optimal performance at low cost. Maximum value can be derived by placing "hot" data with high I/O density and low response time requirements on SSDs; target HDDs for "cooler" data that is accessed more sequentially and at lower rates.

Easy Tier automates the placement of data among different storage tiers, and can be enabled for internal and external storage also. This IBM Storwize V7000 no charge feature boosts your storage infrastructure performance to achieve optimal performance through a software, server, and storage solution.

10.2.1 Easy Tier concepts

Easy Tier is a no-charge feature of IBM Storwize V7000 that brings the enterprise storage functions (originally available on IBM DS8000 and IBM XIV enterprise class storage systems) to the midrange segment. It enables automated subvolume data placement throughout different storage tiers to intelligently align the system with current workload requirements and to optimize the usage of SSDs. This functionality includes the ability to automatically and nondisruptively relocate data (at the extent level) from one tier to another tier in either direction to achieve the best available storage performance for your workload in your environment. Easy Tier reduces the I/O latency for hot spots, but it does not replace storage cache. Both Easy Tier and storage cache solve a similar access latency workload problem, but these two methods weigh differently in the algorithmic construction, based on "locality of reference," recency, and frequency. Because Easy Tier monitors I/O performance from the device end (after cache), it can pick up the performance issues that cache cannot solve and complement the overall storage system performance.

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

The SSDs are treated no differently by the SAN Volume Controller than hard disk drives (HDDs) regarding RAID arrays or MDisks.

The individual SSDs in the storage, managed by the SAN Volume Controller, are combined into an array, usually in RAID 10 or RAID 5 format. It is unlikely that RAID6 SSD arrays will be used because of the double parity overhead, with two logical SSDs used for parity only. A LUN is created on the array and is then presented to the SAN Volume Controller as a normal managed disk (MDisk).

As is the case for HDDs, the SSD RAID array format helps to protect against *individual* SSD failures. Depending on your requirements, you can achieve more high availability protection above the RAID level by using volume mirroring.

Easy Tier is a performance optimization function, as it automatically migrates, or moves, extents belonging to a volume between different storage tiers (Figure 10-1). Because this migration works at the extent level, it is often referred to as *sub-LUN migration*.

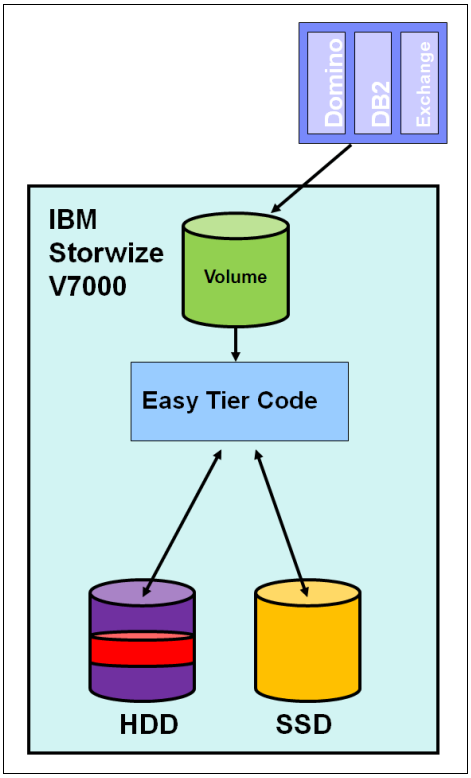

Figure 10-1 Easy Tier

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

10.2.2 Disk tiers

The MDisks (LUNs) presented to the SAN Volume Controller cluster are likely to have different performance attributes because of the type of disk or RAID array on which they reside. The MDisks can be on 15 K RPM Fibre Channel or SAS disk, Near-line SAS or SATA, or even SSDs.

Thus, a storage tier attribute is assigned to each MDisk. The default is generic_hdd. With SAN Volume Controller V6.1, a new disk tier attribute is available for SSDs and is known as generic_ssd.

Consider that the SAN Volume Controller does not automatically detect SSD MDisks. Instead, all external MDisks are initially put into the generic_hdd tier by default. Then, the administrator must manually change the SSD tier to generic_ssd by using the command-line interface (CLI) or graphical user interface (GUI).

Single tier storage pools

Figure 10-2 shows a scenario in which a single storage pool is populated with MDisks that are presented by an external storage controller. In this solution, the striped or mirrored volume can be measured by Easy Tier, but no action to optimize the performance occurs.

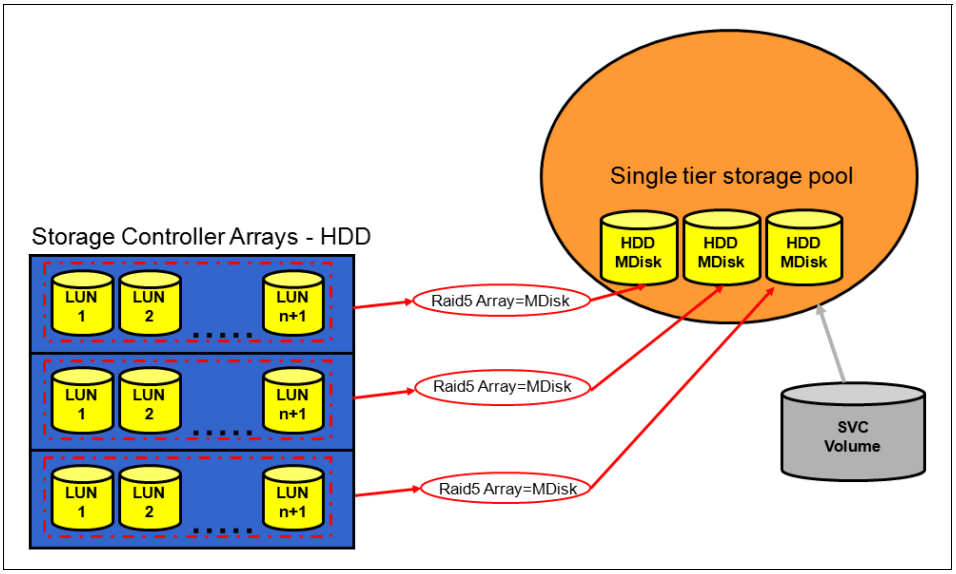

Figure 10-2 Single tier storage pool with striped volume

MDisks that are used in a single-tier storage pool should have the same hardware characteristics, for example, the same RAID type, RAID array size, disk type, and disk revolutions per minute (RPM) and controller performance characteristics.

Multitier storage pools

A multitier storage pool has a mix of MDisks with more than one type of disk tier attribute, for example, a storage pool that contains a mix of generic_hdd and generic_ssd MDisks. Figure 10-3 on page 314 shows a scenario in which a storage pool is populated with two different MDisk types: one belonging to an SSD array and one belonging to an HDD array. Although this example shows RAID 5 arrays, other RAID types can be used.

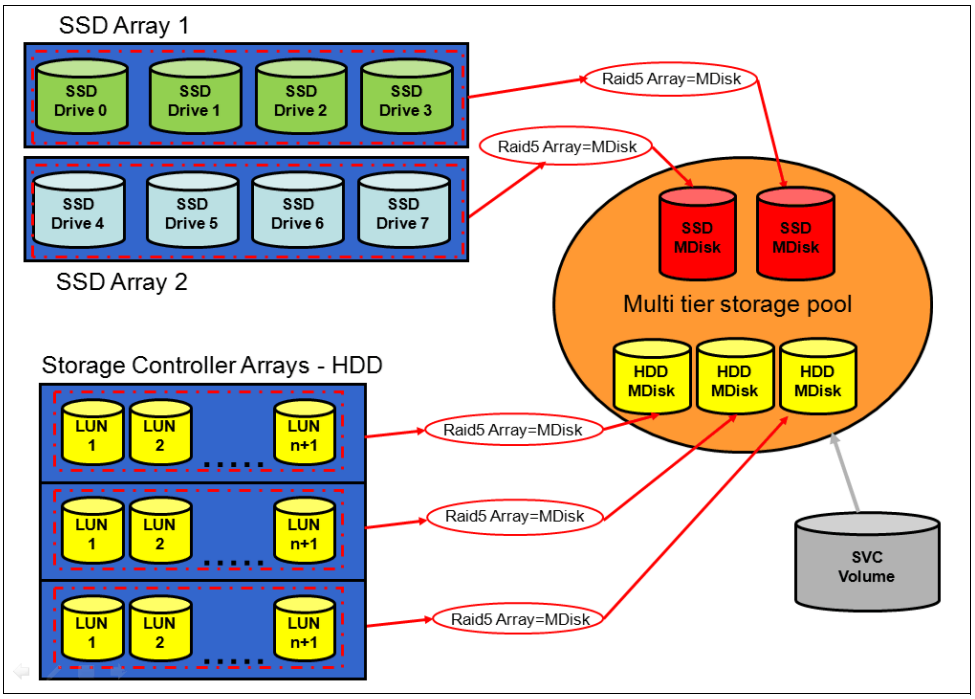

Figure 10-3 Multitier storage pool with striped volume

Adding SSD to the pool means that additional space is also now available for new volumes or volume expansion.

Attention: Image mode and sequential volumes are *not* candidates for Easy Tier automatic data placement.

10.2.3 Easy Tier process

The Easy Tier function has four main processes:

- I/O Monitoring

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

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

  This process is efficient and adds negligible processing overhead to the SAN Volume Controller nodes.

- Data Placement Advisor

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

  This process also identifies extents that need to be migrated back to a lower (HDD) tier.

- Data Migration Planner

  By using the extents that were previously identified, the Data Migration Planner step builds the extent migration plan for the storage pool.

- Data Migrator

  This process involves 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 so that a maximum of up to 30 MBps is migrated, which equates to around 3 TB a day that are migrated between disk tiers.

When it relocates volume extents, Easy Tier performs these actions:

- It attempts to migrate the most active volume extents up to SSD first.

  To ensure that a free extent is available, you might need to first migrate a less frequently accessed extent back to the HDD.

- A previous migration plan and any queued extents that are not yet relocated are abandoned.

10.2.4 Easy Tier operating modes

Easy Tier has three main operating modes:

- Off mode
- Evaluation or measurement only mode
- Automatic Data Placement or extent migration mode

Easy Tier off mode

With Easy Tier is in off mode, no statistics are recorded and no extent migration occurs.

Evaluation or measurement only mode

Easy Tier Evaluation or measurement only mode collects usage statistics for each extent in a single tier storage pool where the Easy Tier value is set to *on* for both the volume and the pool. This collection is typically done for a single-tier pool that contains only HDDs so that the benefits of adding SSDs to the pool can be evaluated before any major hardware acquisition.

A `dpa_heat.nodeid.yymmdd.hhmmss.data` statistics summary file is created in the `/dumps` directory of the SAN Volume Controller nodes. This file can be offloaded from the SAN Volume Controller nodes with PSCP **`-load`** or by using the GUI as shown in the "Measuring by using the Storage Advisor Tool" topic in *IBM System Storage SAN Volume Controller and Storwize V7000 Best Practices and Performance Guidelines*, SG24-7521. A web browser is used to view the report that is created by the tool.

Automatic Data Placement or extent migration mode

In Automatic Data Placement or extent migration operating mode, the **`-easytier on`** or **`auto`** storage pool parameter must be set, and the volumes in the pool must have **`-easytier on`**. The storage pool must also contain MDisks with different disk tiers, thus being a multitiered storage pool.

Dynamic data movement is transparent to the host server and application users of the data, other than providing improved performance. Extents are automatically migrated as explained in the implementation rules in *IBM System Storage SAN Volume Controller and Storwize V7000 Best Practices and Performance Guidelines*, SG24-7521.

The statistic summary file is also created in this mode. This file can be offloaded for input to the advisor tool. The tool produces a report on the extents that are moved to SSD and a prediction of performance improvement that can be gained if more SSD arrays are available.

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

10.2.5 Implementation considerations

No Easy Tier license is required for the SAN Volume Controller. Easy Tier comes as part of the V6.1 code. For Easy Tier to migrate extents, disk storage must be available that has different tiers, for example a mix of SSD and HDD.

Implementation rules

Consider the following implementation and operation rules when you use the IBM System Storage Easy Tier function on the SAN Volume Controller:

- Easy Tier automatic data placement is not supported on 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.
- Automatic data placement and extent I/O activity monitors are supported on each copy of a mirrored volume. Easy Tier works with each copy independently of the other copy.

  Volume mirroring consideration: Volume mirroring can have different workload characteristics on each copy of the data because reads are normally directed to the primary copy and writes occur to both. Thus, the number of extents that Easy Tier migrates to the SSD tier might be different for each copy.

- If possible, the SAN Volume Controller creates new volumes or volume expansions by using extents from MDisks from the HDD tier. However, it uses extents from MDisks from the SSD tier if necessary.

When a volume is migrated out of a storage pool that is managed with Easy Tier, Easy Tier automatic data placement mode is no longer active on that volume. Automatic data placement is also turned off while a volume is being 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.

Limitations

When you use IBM System Storage Easy Tier on the SAN Volume Controller, Easy Tier has the following limitations:

- Removing an MDisk by using the **`-force`** parameter

  When an MDisk is deleted from a storage pool with the **`-force`** parameter, extents in use are migrated to MDisks in the *same* tier as the MDisk that is being removed, if possible. If insufficient extents exist in that tier, extents from the other tier are used.

- Migrating extents

  When Easy Tier automatic data placement is enabled for a volume, you cannot use the **`svctask migrateexts`** CLI command on that volume.

- Migrating a volume to another storage pool

  When the SAN Volume Controller migrates a volume to a new storage pool, Easy Tier automatic data placement between the two tiers is temporarily suspended. After the volume is migrated to its new storage pool, Easy Tier automatic data placement between the generic SSD tier and the generic HDD tier resumes for the moved volume, if appropriate.

  When the SAN Volume Controller migrates a volume from one storage pool to another, it attempts to migrate each extent to an extent in the new storage pool from the same tier as the original extent. In several cases, such as where a target tier is unavailable, the other tier is used. For example, the generic SSD tier might be unavailable in the new storage pool.

- Migrating a volume to image mode.

  Easy Tier automatic data placement does not support image mode. When a volume with Easy Tier automatic data placement mode active is migrated to image mode, Easy Tier automatic data placement mode is no longer active on that volume.

- Image mode and sequential volumes cannot be candidates for automatic data placement, however, Easy Tier supports evaluation mode for image mode volumes.

Preferred practices:

- Always set the storage pool **`-easytier`** value to `on` rather than to the default value `auto`. With this setting, you can more easily turn on evaluation mode for existing single tier pools, and no further changes are needed when you move to multitier pools. For more information about the mix of pool and volume settings, see the Easy Tier activation topic in *IBM System Storage SAN Volume Controller and Storwize V7000 Best Practices and Performance Guidelines*, SG24-7521.
- With Easy Tier, using smaller storage pool extent sizes can become more appropriate.

10.2.6 More information

Detailed planning and configuration considerations, preferred practices, and monitoring and measurement tools description are available in *IBM System Storage SAN Volume Controller and Storwize V7000 Best Practices and Performance Guidelines*, SG24-7521 and *Implementing IBM Easy Tier with IBM Real-time Compression*, TIPS1072.

10.3 Thin provisioning

Thin provisioning, in a shared storage environment, is a method for optimizing utilization of available storage. It relies on allocation of blocks of data-on-demand, versus the traditional method of allocating all the blocks in the beginning. This way eliminates almost all white space, which helps avoid the poor utilization rates, often as low as 10%, that occur in the traditional storage allocation method where large pools of storage capacity are allocated to individual servers but remain unused (not written to).

Thin provisioning presents more storage space to the hosts or servers that are connected to the storage system than is actually available on the storage system. The IBM Storwize V7000 has this capability for both, Fibre Channel and iSCSI provisioned volumes.

An example of thin provisioning is when a storage system contains 5000 GB of usable storage capacity, but the storage administrator has mapped volumes of 500 GB each to 15 hosts. In this example, the storage administrator makes 7500 GB of storage space visible to the hosts even though the storage system has only 5000 GB of usable space (Figure 10-4). If all 15 hosts immediately use all 500 GB provisioned to them, there would be a problem. The storage administrator has to monitor the system and add storage as needed.

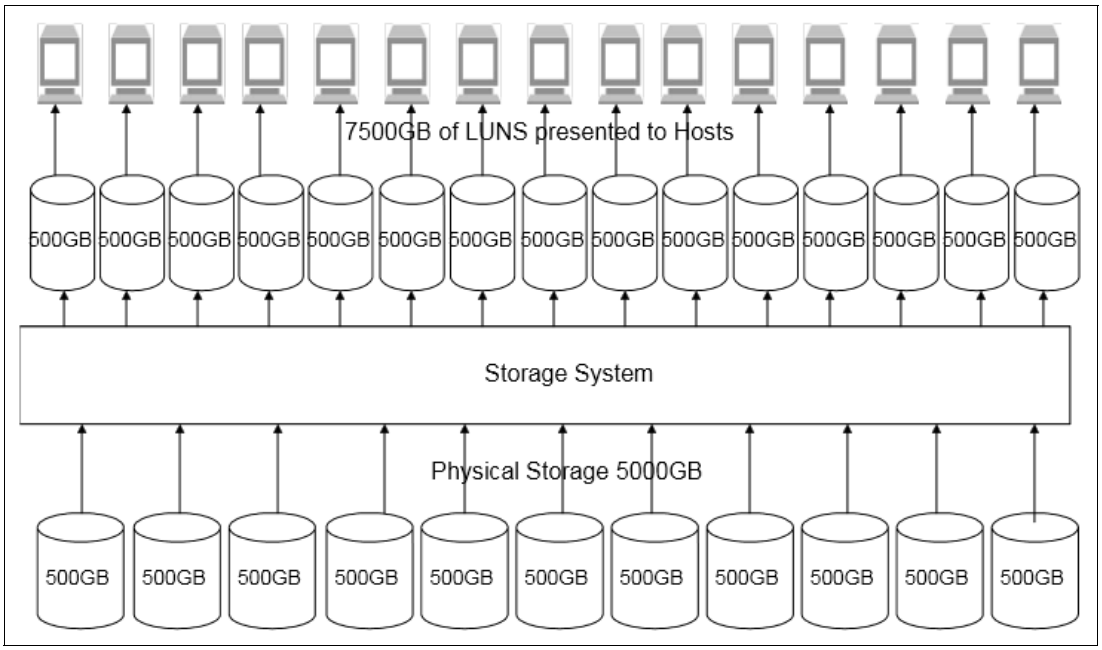

Figure 10-4 Concept of thin provisioning

You can imagine thin provisioning as the same process as when airlines sell more tickets on a flight than available physical seats, assuming that some passengers will not appear at checkin. They do not assign actual seats at the time of sale, avoiding each client having a claim on a specific seat number. The same concept applies to thin provisioning (airline) IBM Storwize V7000 (plane) and its volumes (seats). The storage administrator (airline ticketing system) has to closely monitor the allocation process and set proper thresholds.

10.3.1 Configuring thin-provisioned volume

Volumes can be configured as *thin-provisioned* or *fully allocated*. Thin-provisioned volumes are created with real and virtual capacities. You can still create volumes by using a striped, sequential, or image mode virtualization policy, just as you can with any other volume.

Real capacity defines how much disk space is allocated to a volume. *Virtual capacity* is the capacity of the volume that is reported to other IBM System Storage SAN Volume Controller components (such as FlashCopy or remote copy) and to the hosts.

A directory maps the virtual address space to the real address space. The directory and the user data share the real capacity.

Thin-provisioned volumes are available in two operating modes: autoexpand and non-autoexpand. You can switch the mode at any time. If you select the autoexpand feature, the SAN Volume Controller automatically adds a fixed amount of additional real capacity to the thin volume as required. Therefore, the autoexpand feature attempts to maintain a fixed amount of unused real capacity for the volume. This amount is known as the *contingency*

capacity. 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 without the autoexpand feature, and thus has a zero contingency capacity, goes offline as soon as the real capacity is used and needs to expand.

Warning threshold: Enable the warning threshold (by using email or an SNMP trap) when working with thin-provisioned volumes, on the volume, and on the storage pool side, especially when you do not use the autoexpand mode. Otherwise, the thin volume goes offline if it runs out of space.

Autoexpand mode does not cause 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.

A thin-provisioned volume can be converted nondisruptively to a fully allocated volume, or vice versa, by using the volume mirroring function. For example, you can add a thin-provisioned copy to a fully allocated primary volume and then remove the fully allocated copy from the volume after they are synchronized.

The fully allocated to thin-provisioned migration procedure uses a zero-detection algorithm so that grains that contain all zeros do not cause any real capacity to be used.

Tip: Consider using thin-provisioned volumes as targets in FlashCopy relationships.

Space allocation

When a thin-provisioned volume is initially created, a small amount of the real capacity is used for initial metadata. Write I/Os to the grains of the thin volume (that were not previously written to) cause grains of the real capacity to be used to store metadata and user data. Write I/Os to the grains (that were previously written to) update the grain where data was previously written.

Grain definition: The grain is defined when the volume is created and can be 32 KB, 64 KB, 128 KB, or 256 KB.

Smaller granularities can save more space, but they have larger directories. When you use thin provisioning with FlashCopy, specify the same grain size for both the thin-provisioned volume and FlashCopy.

To create thin-provisioned volume, choose **Create Volume** from Volumes in the dynamic menu and select **Thin Provision** (Figure 10-5). Enter the capacity and volume name.

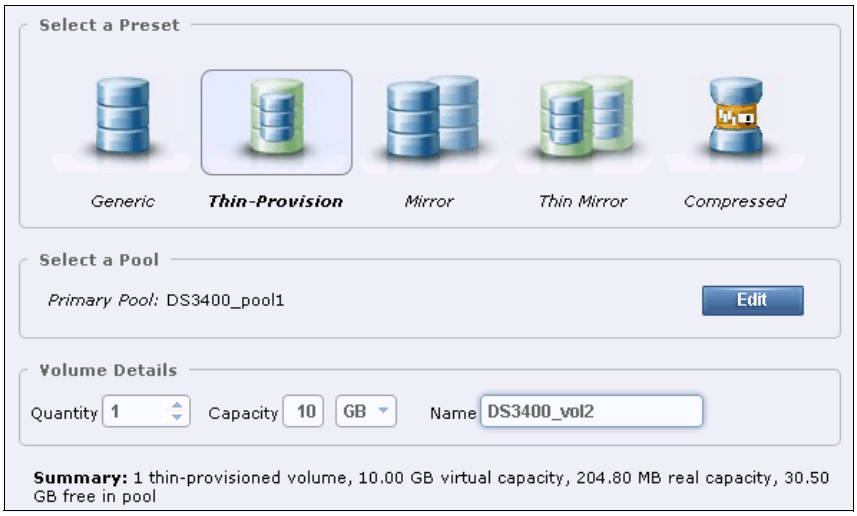

Figure 10-5 Thin-provisioned volume creation

In an Advanced menu of this wizard, you can set virtual and real capacity, warning thresholds, and grain size (Figure 10-6).

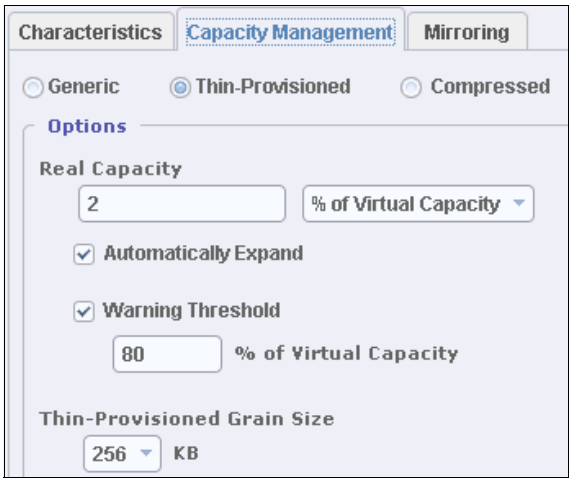

Figure 10-6 Advanced options

For the whole procedure of configuring a thin-provisioned volume, follow the guidance in 5.1.2, “Creating thin-provisioned volumes” on page 135.

10.3.2 Performance considerations

Thin-provisioned volumes require more I/Os because of the directory accesses:

- For truly random workloads, a thin-provisioned volume requires approximately one directory I/O for every user I/O so that performance is 50% of a normal volume.
- The directory is two-way write-back cache (similar to the SAN Volume Controller fast-write cache) so that certain applications perform better.
- Thin-provisioned volumes require more CPU processing so that the performance per I/O group is lower.

Use the striping policy to spread thin-provisioned volumes across many storage pools.

Important: Do not use thin-provisioned volumes where high I/O performance is required.

Thin-provisioned volumes save capacity only if the host server does not write to whole volumes. Whether the thin-provisioned volume works well partly depends on how the file system allocated the space:

- Some file systems (for example, New Technology File System (NTFS)) write to the whole volume before they overwrite the deleted files. Other file systems reuse space in preference to allocating new space.
- File system problems can be moderated by tools, such as “defrag,” or by managing storage by using host Logical Volume Managers (LVMs).

The thin-provisioned volume also depends on how applications use the file system. For example, some applications delete log files only when the file system is nearly full.

There is no recommendation for thin-provisioned volumes. As explained previously, the performance of thin-provisioned volumes depends on what is used in the particular environment. For the absolute best performance, use fully allocated volumes instead of a thin provisioned volume.

10.3.3 Limitations of virtual capacity

Several factors (extent and grain size) limit the virtual capacity of thin-provisioned volumes beyond the factors that limit the capacity of regular volumes.

Table 10-1 shows the maximum thin-provisioned volume virtual capacities for an extent size.

Table 10-1 Maximum thin volume virtual capacities for an extent size

Extent size in MB	Maximum volume real capacity in GB	Maximum thin virtual capacity in GB
16	2,048	2,000
32	4,096	4,000
64	8,192	8,000
128	16,384	16,000
256	32,768	32,000
512	65,536	65,000
1024	131,072	130,000
2048	262,144	260,000
4096	524,288	520,000
8192	1,048,576	1,040,000

Table 10-2 show the maximum thin-provisioned volume virtual capacities for a grain size.

Table 10-2 Maximum thin volume virtual capacities for a grain size

Grain size in KB	Maximum thin virtual capacity in GB
32	260,000
64	520,000
128	1,040,000
256	2,080,000

For more information and detailed performance considerations when you configure thin provisioning, study the *IBM System Storage SAN Volume Controller and Storwize V7000 Best Practices and Performance Guidelines*, SG24-7521.

10.4 Real-time Compression

The IBM Real-time Compression software embedded in IBM System Storage SAN Volume Controller and IBM Storwize V7000 solution addresses the requirements of primary storage data reduction, including performance. It does so by using a purpose-built technology called Real-time Compression (RACE engine). It offers the following benefits:

- Compression for active primary data: IBM Real-time Compression can be used with active primary data. Therefore, it supports workloads that are not candidates for compression in other solutions. The solution supports online compression of existing data. It allows storage administrators to regain free disk space in an existing storage system without requiring administrators and users to clean up or archive data. This configuration significantly enhances the value of existing storage assets, and the benefits to the business are immediate. The capital expense of upgrading or expanding the storage system is delayed.
- Compression for replicated and mirrored data: Remote volume copies can be compressed in addition to the volumes at the primary storage tier. This process reduces storage requirements in Metro Mirror and Global Mirror destination volumes also.
- No changes to the existing environment are required: IBM Real-time Compression is part of the storage system. It was designed for transparency so that it can be implemented without changes to applications, hosts, networks, fabrics, or external storage systems. The solution is not apparent to hosts, so users and applications continue to work as-is. Compression occurs within the Storwize V7000 or SAN Volume Controller system.
- Overall savings in operational expenses: More data is stored in a rack space, so fewer storage expansion enclosures are required to store a data set. This reduced rack space has the following benefits:
  - Reduced power and cooling requirements: More data is stored in a system, therefore requiring less power and cooling per gigabyte or used capacity.
  - Reduced software licensing for additional functions in the system: More data stored per enclosure reduces the overall spending on licensing.

Tip: Implementing compression in SAN Volume Controller provides the same benefits to externally virtualized storage systems.

- Disk space savings are immediate: The space reduction occurs when the host writes the data. This process is unlike other compression solutions in which some or all of the reduction is realized only after a post-process compression batch job is run.

10.4.1 Concept of Real-time Compression

Random Access Compression Engine (RACE) technology is based on over 40 patents that are not primarily about compression. Rather they define how to make industry standard Lempel-Ziv (LZ) compression of primary storage operate in real-time and allow random access. The primary intellectual property behind this is the RACE engine. At a high level, the IBM RACE component compresses data that is written into the storage system dynamically. This compression occurs transparently, so Fibre Channel and iSCSI connected hosts are not aware of the compression. RACE is an in-line compression technology, meaning that each host write is compressed as it passes through the SAN Volume Controller software to the disks. This has a clear benefit over other compression technologies that are based on post-processing. These technologies do not provide immediate capacity savings, and therefore are not a good fit for primary storage workloads such as databases and active data-set applications.

When a host sends a write request, it is acknowledged by the write cache of the system, and then staged to the storage pool. As part of its staging, it passes through the compression engine, and is then stored in compressed format onto the storage pool. Writes are therefore acknowledged immediately after they are received by the write cache, with compression occurring as part of the staging to internal or external physical storage.

Capacity is saved when the data is written by the host because the host writes are smaller when written to the storage pool.

IBM Real-time Compression is a self-tuning solution, similar to the Storwize V7000 and SAN Volume Controller system itself. It is adapted to the workload that runs on the system at any particular moment.

Random Access Compression Engine

The IBM patented Random Access Compression Engine turns over the traditional approach to compression. It uses variable-size chunks for the input, and fixed-size chunks for the output. This method enables an efficient and consistent method to index the compressed data because it is stored in fixed-size containers.

RACE technology is implemented into the SAN Volume Controller and Storwize V7000 thin provisioning layer, and is an organic part of the stack. The SAN Volume Controller and Storwize V7000 software stack is shown in Figure 10-7 on page 324. Compression is transparently integrated with existing system management design. All of the SAN Volume Controller and Storwize V7000 advanced features are supported on compressed volumes. You can create, delete, migrate, map (assign), and unmap (unassign) a compressed volume as though it were a fully allocated volume. This compression method provides nondisruptive conversion between compressed and uncompressed volumes. This conversion provides a uniform user experience and eliminates the need for special procedures when dealing with compressed volumes.

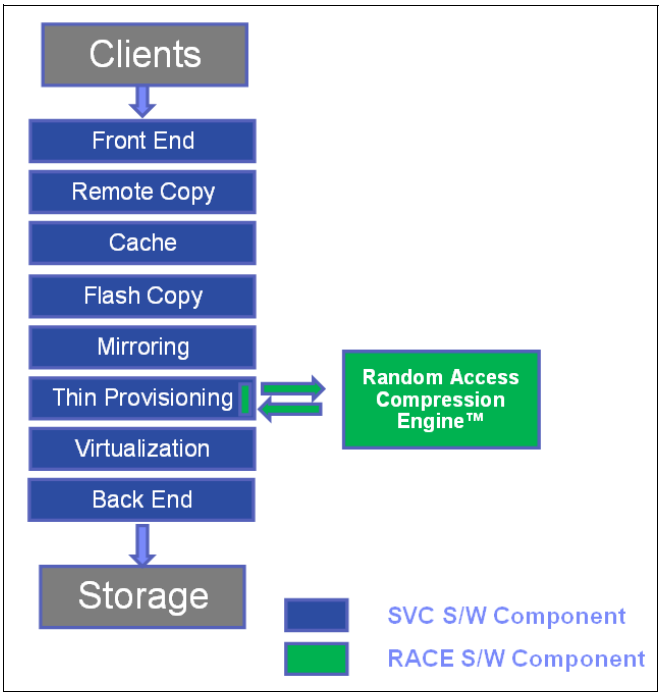

Figure 10-7 RACE integration in IBM Storwize V7000

When a host sends a *write request* to Storwize V7000 or SAN Volume Controller, it reaches the cache layer. The host is immediately sent an acknowledgement of its I/Os. When the cache layer destages to the RACE, the I/Os are sent to the thin-provisioning layer. They are then sent to RACE and, if necessary, the original host write or writes. The metadata that holds the index of the compressed volume is updated if needed, and is compressed also.

When a host sends a *read request* to Storwize V7000 or SAN Volume Controller, it reaches the cache layer as follows:

- If there is a cache hit, the SAN Volume Controller cache replies to the host with the requested data.
- If there is a cache miss, the SAN Volume Controller cache sends an internal read request to the thin-provisioning layer and then to RACE.

10.4.2 Configuring compressed volumes

To use compression on the SAN Volume Controller and Storwize V7000, licensing is required. With the SAN Volume Controller, Real-time Compression is licensed by capacity, per terabyte of virtual data. In the Storwize V7000, Real-time Compression is licensed per enclosure.

The configuration is similar to generic volumes and is transparent to users. From the Volumes in a dynamic menu select **Create Volume** and click **Compressed** (Figure 10-8 on page 325). Select a storage pool and enter the required capacity and volume name.

The summary in the bottom of the wizard provides the information about allocated (virtual) capacity and the real capacity that data consume on this volume. In our example we have defined a 15 GB volume but the real capacity is only 307 MB (because there is no data from the host yet).

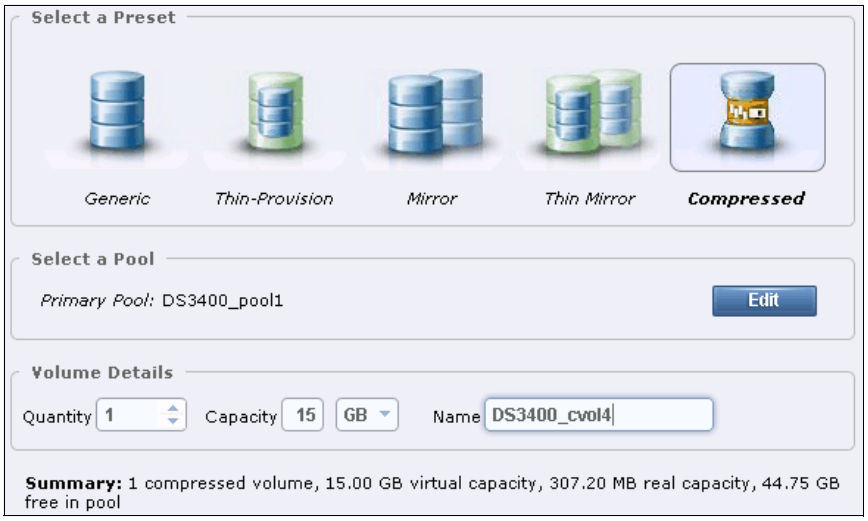

Figure 10-8 Configuring compressed volume

When compressed volume is configured, you can directly map it to the host or do so later. For more details about definition of compressed volumes, see 5.1.5, “Creating compressed volumes” on page 141.

10.4.3 Differences from IBM Real-time Compression Appliances

Although the underlying technology of the IBM Real-time Compression Appliance™ and the built-in Real-time Compression in SAN Volume Controller and Storwize V7000 is the same, there are some notable differences:

- Appliance-based versus internal: IBM Real-time Compression Appliance is implemented as an add-on to an NAS storage system, whereas SAN Volume Controller and Storwize V7000 compression is integral to the system.
- File-based versus block-based: Real-time Compression Appliance compresses data written to files, whereas SAN Volume Controller and Storwize V7000 compression is for data written to block devices. However, this distinction makes no practical difference to the user. Block devices typically contain file systems, so the SAN Volume Controller and Storwize V7000 compression is applicable to files also.
- Exports versus volumes: Real-time Compression Appliance configuration of which data to compress is based on exports and shares, whereas in SAN Volume Controller and Storwize V7000 the configuration is per volume.
- Supported external storage systems: Real-time Compression Appliance support matrix is focused on the major NAS storage systems. SAN Volume Controller and Storwize V7000 support the major SAN storage systems.

For more information about Real-time Compression and its deployment in IBM Storwize V7000, see *Real-time Compression in SAN Volume Controller and Storwize V7000*, REDP-4859.

11

Copy services

In this chapter, we discuss the Copy Services functions that are available for the IBM Storwize product family. These functions are also available for IBM System Storage SAN Volume Controller.

We describe the copy services functions provided by the IBM Storwize V7000 storage system, including FlashCopy and Remote Copy. Copy services functions are useful for making data copies for backup, application test, recovery, and so on. The IBM Storwize V7000 system can help you more easily apply these functions to your environment through its intuitive GUI.

11.1 FlashCopy

With the FlashCopy function of the IBM Storwize V7000 storage system, you 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 duplication of data of on your source volume. Volumes may remain online and active while you create consistent copies of the data sets. Because the copy is performed at the block level, it operates below the host operating system and cache and is therefore 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, the source volume is frozen briefly to initialize the FlashCopy bitmap and then I/O is allowed to resume. Although several FlashCopy options require the data to be copied from the source to the target in the background, which can take time to complete, the resulting data on the target volume copy appears to have completed immediately. This task is accomplished through the use of a bitmap (or bit array) that tracks changes to the data after the FlashCopy is initiated, and an indirection layer, which allows data to be read from the source volume transparently.

11.1.1 Business requirements for FlashCopy

When deciding whether FlashCopy addresses your needs, adopt a combined business and technical view of the problems you need to solve. Determine what your needs are from a business perspective, and then determine if FlashCopy fulfills the technical needs of those business requirements.

With an immediately available copy of the data, FlashCopy can be used in various business scenarios:

- Rapidly creating consistent backups of dynamically changing data

  FlashCopy can be used to create backups through periodic execution of FlashCopy; the FlashCopy target volumes can be used to rapidly restore individual files or the entire volume through Reverse FlashCopy (using the **-restore** option).

  The target volumes created by FlashCopy may be also used for backup to tape, by attaching them to another server, which to a great extent offloads the production server. After the copy to tape completes, the target volumes can be discarded if required or kept as a rapid restore copy of the data.

- Rapidly creating consistent copies of production data to facilitate data movement or migration between hosts

  FlashCopy can be used to facilitate the movement or migration of data between hosts while minimizing downtime for applications. FlashCopy allows application data to be copied from source volumes to new target volumes while applications remain online. After the volumes are fully copied and synchronized, the application can be stopped and then immediately started on the new server accessing the new FlashCopy target volumes. This mode of migration is swifter than other migration methods that are available through the IBM Storwize V7000 because the size or the speed of the migration is not as limited.

- Rapidly creating copies of production data sets for application development and testing

  Under normal circumstances, to perform application development and testing, data must be restored from traditional backup media, such as tape. Depending the amount of data and the technology in use, this process could easily take a day or more. With FlashCopy, a copy can be created and online for use in just a few minutes. The time varies based on the application and the data set size.

- Rapidly creating copies of production data sets for auditing purposes and data mining

  Auditing or data mining normally require the usage of the production applications. This situation can cause high loads for databases track inventories or similar data. With FlashCopy, you can create copies for your reporting and data mining activities. This situation reduces the load on your production systems, increasing their performance.

- Rapidly creating copies of production data sets for quality assurance

  Quality assurance is an interesting case for FlashCopy. Because traditional methods involve so much time and labor, the refresh cycle is typically extended. This reduction in time allows much more frequent refreshes of the quality assurance database.

11.1.2 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 using consistency groups. FlashCopy consistency groups are further described in “FlashCopy consistency groups” on page 334.

The minimum granularity that IBM Storwize V7000 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 V7000 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 specified) you must run either **prestartfcmap** or **prestartfcconsistgrp**, which puts the SAN Volume Controller Cache into write-through mode, flushing the I/O currently bound for your volume. After FlashCopy is started, an effective copy of a source volume to a target volume has been 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, both 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 11-1 shows redirection of the host I/O toward the source volume and target volume.

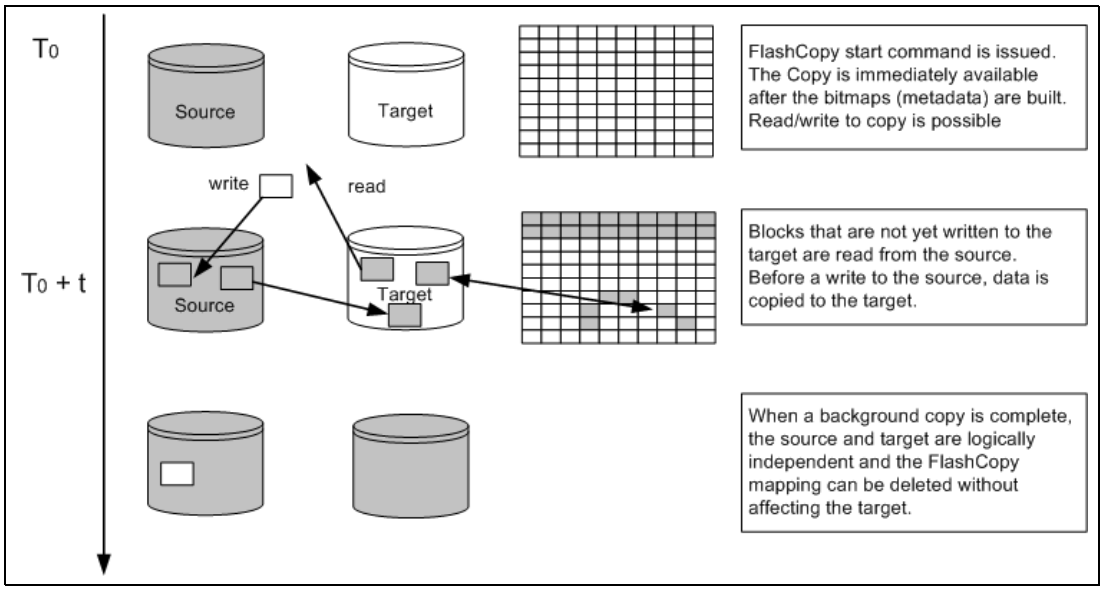

Figure 11-1 Redirection of host I/O

When data is copied between volumes, it is copied in units of address space known as *grains*. Grains are units of data grouped together to optimize the use of the bitmap that tracks changes to the data between the source and target volume. You have the option of using 64 KB or 256 KB grain sizes; 256 KB is the default. The FlashCopy bitmap contains 1 bit for each grain and is used to track whether the source grain has been copied to the target. The 64 KB grain size consumes bitmap space at a rate of four times the 256 KB default size.

The FlashCopy bitmap dictates read and write behavior for both the source and target volumes as follows:

- Read I/O request to source: Reads are performed from the source volume, which is the same as for non-FlashCopy volumes.
- Write I/O request to source: Writes to the source cause the grains to be copied to the target if it has not already been copied, and then the write is performed to the source.
- Read I/O request to target: Reads are performed from the target if the grains have already been copied; otherwise, the read is performed from the source.
- Write I/O request to target: Writes to the target cause the grain to be copied from the source to the target unless the entire grain is being written, and then the write completes to the target.

FlashCopy mappings

A FlashCopy mapping defines the relationship between a source volume and a target volume. FlashCopy mappings can be either stand-alone mappings or a member of a consistency group, as described in “FlashCopy consistency groups” on page 334.

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 only copy data that has been modified since the initial FlashCopy mapping. This action reduces the amount of time spent 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 referring to the source. A dependency relationship exists between a particular target and all newer targets that share a source until all data has been copied to this target and all older targets.

Cascaded FlashCopy mappings

The cascaded FlashCopy function allows a FlashCopy target volume to be the source volume of another FlashCopy mapping. Up to 256 mappings can exist in a cascade. If cascaded mappings and multiple target mappings are used, a tree of up to 256 mappings can be created.

Cascaded mappings differ from multiple target FlashCopy mappings in depth. Cascaded mappings have an association in the manner of A > B > C; 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 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 enabled background copy is that the target volume becomes a real independent clone of the FlashCopy mapping source volume after the copy is complete. When the background copy is disabled, the target volume only remains a valid copy of the source data while the FlashCopy mapping remains in place. Copying only the changed data saves your storage capacity (assuming it is thin provisioned and **`-rsize`** has been correctly set up.)

The relationship of the background copy rate value to the amount of data copied per second is shown in Table 11-1.

Table 11-1 Background copy rate

Value	Data copied per second	Grains per second (256 KB grain)	Grains per second (64 KB grain)
1 - 10	128 KB	0.5	2
11 - 20	256 KB	1	4
21 - 30	512 KB	2	8
31 - 40	1 MB	4	16
41 - 50	2 MB	8	32
51 - 60	4 MB	16	64
61 - 70	8 MB	32	128
71 - 80	16 MB	64	256
81 - 90	32 MB	128	512
91 - 100	64 MB	256	1024

Data copy rate: The data copy rate remains the same regardless of the FlashCopy grain size. The difference is the number of grains copied per second. The grain size can be either 64 KB or 256 KB. The smaller size consumes more bitmap space and thus limits the total amount of FlashCopy space possible, but might be more efficient regarding the amount of data moved, depending on your environment.

Cleaning rate

The cleaning rate provides a method for FlashCopy copies with dependant mappings (either multiple target or cascaded) to be able to complete their background copies before their source goes offline or is deleted after a stop has been issued.

When you create or modify a FlashCopy mapping, you can specify a cleaning rate for the FlashCopy mapping that is independent of the background copy rate. The cleaning rate is also defined as a value of 0 - 100, which has the same relationship to data copied per second with the backup copy rate (Table 11-1).

The cleaning rates 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 maybe 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

At any point in time, a mapping is in one of the following states:

- Idle or Copied

  The source and target volumes act as independent volumes even if a mapping exists between the two. Read and write caching is enabled for both source and target volumes.

  If the mapping is incremental and the background copy is complete, the mapping records only the differences between the source and target volumes. If the connection to both nodes in the IBM Storwize V7000 storage system that the mapping is assigned to is lost, the source and target volumes go offline.

- Copying

  The copy is in progress. Read and write caching is enabled on 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 SCSI front end as a hardware error. If the mapping is incremental and a previous mapping has completed, the mapping only records the differences between the source and target volumes. If the connection to both nodes in the IBM Storwize V7000 storage system that the mapping is assigned to is lost, the source and target volumes go offline.

- Preparing

  The target volume is online, but not accessible. The target volume cannot perform read or write caching. Read and write caching is failed by the SCSI front end as a hardware error. Any changed write data for the source volume is flushed from the cache. Any read or write data for the target volume is discarded from the cache. If the mapping is incremental and a previous mapping has completed, the mapping records only the differences between the source and target volumes. If the connection to both nodes in the IBM Storwize V7000 storage system that the mapping is assigned to is lost, the source and target volumes go offline.

- Stopped

  The mapping is stopped because either 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. If the connection to both nodes in the IBM Storwize V7000 storage system that the mapping is assigned to is lost, the source and target volumes go offline.

- Stopping

  The mapping is in the process of copying data to another mapping. If the background copy process is complete, the target volume is online while the stopping copy process completes. If the background copy process is 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 started, but it did not complete. Access to the metadata is lost, which causes both 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 has not been flushed and has been 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 having related data that spans them. For these volumes, consistency groups maintain the integrity of the FlashCopy by ensuring that *dependent writes*, which are described in more detail in “Dependent writes” on page 334, are executed in the application’s intended sequence.

When consistency groups are used, the FlashCopy commands are issued to the FlashCopy consistency group, which performs the operation on all FlashCopy mappings contained within the consistency group.

Figure 11-2 shows a consistency group consisting of two FlashCopy mappings.

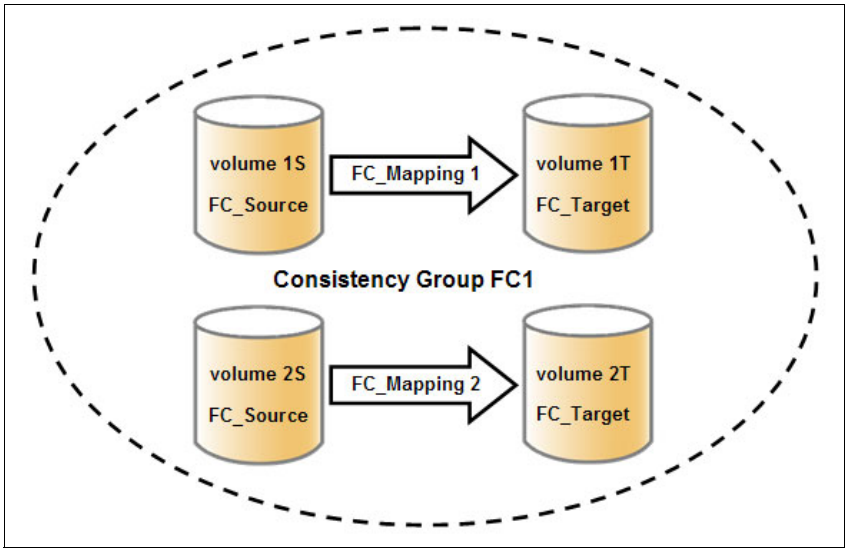

Figure 11-2 FlashCopy consistency group

FlashCopy mapping management: After an individual FlashCopy mapping has been 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 illustrate why using consistency groups is crucial when a data set spans multiple volumes, consider the following typical sequence of writes for a database update transaction:

1. A write is executed to update the database log, indicating that a database update is about to be performed.
2. A second write is executed to complete the actual update to the database.
3. A third write is executed to update the database log, indicating that the database update has completed successfully.

The database ensures the correct ordering of these writes by waiting for each step to complete before starting the next step. However, if the database log (updates 1 and 3) and the database itself (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 had 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 had 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, performing a FlashCopy operation on multiple volumes as an atomic operation using consistency groups is necessary.

A FlashCopy consistency group can contain up to 512 FlashCopy mappings. The more mappings you have, the more time is required 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 starting the FlashCopy for the consistency group, all FlashCopy mappings in the consistency group are started at the same time, resulting in a point-in-time copy that is consistent across all FlashCopy mappings that are contained in the consistency group.

A consistency group aggregates FlashCopy mappings, not volumes. Thus, where a source volume has multiple FlashCopy mappings, they can be in the same or separate consistency groups. If a particular volume is the source volume for multiple FlashCopy mappings, you might want to create separate consistency groups to separate each mapping of the same source volume. Regardless of whether the source volume with multiple target volumes is in the same consistency group or in separate consistency groups, the resulting FlashCopy produces multiple identical copies of the source data.

The consistency group can be specified when the mapping is created. You can also add the FlashCopy mapping to a consistency group or change the consistency group of a FlashCopy mapping later. Do not place stand-alone mappings into a consistency group, because they become controlled as part of that consistency group.

FlashCopy consistency group states

At any point in time, a FlashCopy consistency group is in one of the following states:

- Idle or Copied

  All FlashCopy Mappings in this consistency group are in the Idle or Copied state.

- Preparing

  At least one FlashCopy mapping in this consistency group is in the Preparing state.

- Prepared

  The consistency group is ready to start. While in this state, the target volumes of all FlashCopy mappings in this consistency group are 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 either you issued a command or an I/O error occurred.

- Suspended

  At least one FlashCopy mapping in the consistency group is in the Suspended state.
- Empty

  The consistency group does not have any FlashCopy mappings.

Reverse FlashCopy

Reverse FlashCopy enables FlashCopy targets to become restore points for the source without breaking the FlashCopy relationship and without waiting for the original copy operation to complete. It supports multiple targets and multiple rollback points.

A key advantage of Reverse FlashCopy is that it does not delete the original target, thus allowing processes using 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 starting the reverse copy operation. This copy restores the original source data, which can be useful for diagnostic purposes.

Figure 11-3 shows an example of the reverse FlashCopy scenario.

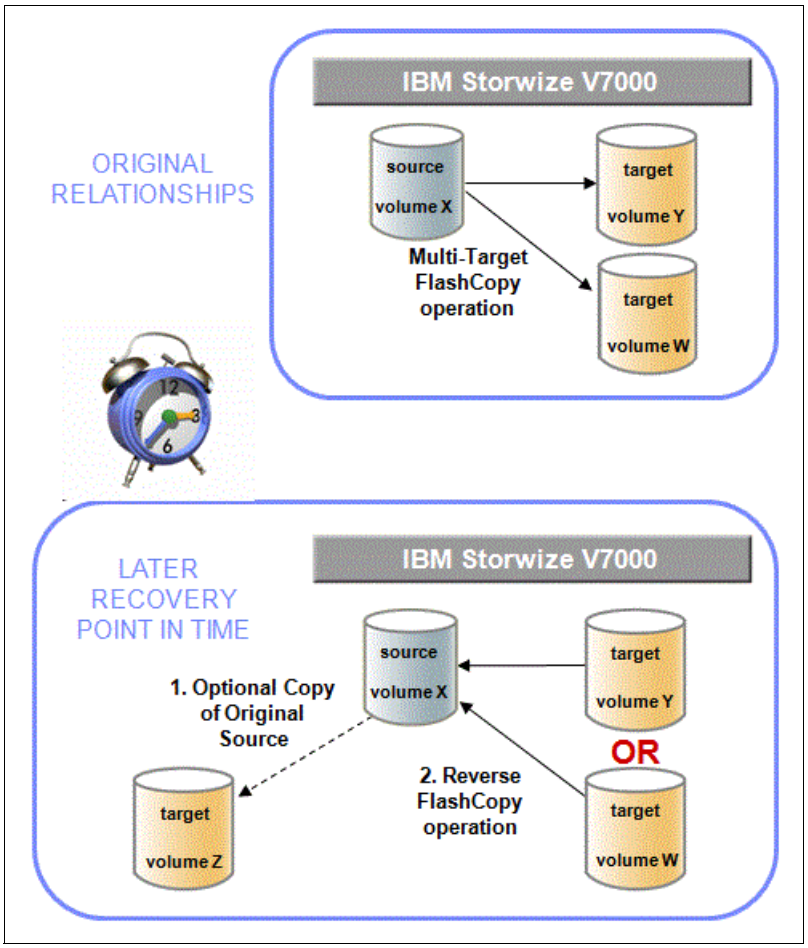

Figure 11-3 Reverse FlashCopy scenario

To restore from an FlashCopy backup, complete the following steps:

1. Optional: Create a target volume (volume Z) and run FlashCopy on the production volume (volume X) to copy data on to the new target for later problem analysis.
2. Create a FlashCopy map with the backup to be restored (volume Y) or (volume W) as the source volume and volume X as the target volume, if this map does not exist.
3. Start the FlashCopy map (volume Y to volume X) with the **-restore** option to copy the backup data onto the production disk.

> **The -restore option:** In the GUI, the **-restore** option is applied automatically when you start the FlashCopy mapping from volume Y to volume X; in the CLI, you need to add the **-restore** option to the command manually. More information about using the CLI is in Appendix A, "CLI setup and SAN Boot" on page 497.

The production disk is instantly available with the backup data.

Regardless of whether the initial FlashCopy map (volume X to volume Y) is incremental, the Reverse FlashCopy operation only copies 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.

11.1.3 Planning for FlashCopy

Before performing a FlashCopy, read the planning considerations 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 V7000 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 them 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 11-2 lists several maximum configurations for FlashCopy.

Table 11-2 FlashCopy maximum configurations

FlashCopy property	Maximum
FlashCopy targets per source	256
FlashCopy mappings per system	4,096
FlashCopy consistency groups per system	127
FlashCopy mappings per consistency group	512

FlashCopy presets

IBM Storwize V7000 storage system provides three FlashCopy presets, listed in Table 11-3, to simplify the more common FlashCopy operations.

Table 11-3 FlashCopy presets

Preset	Purpose
Snapshot	Creates a point-in-time view of the production data. The snapshot is not intended to be an independent copy, but is used to maintain a view of the production data at the time the snapshot is created. This preset automatically creates a thin-provisioned target volume with 0% of the capacity allocated at the time of creation. The preset uses a FlashCopy mapping with 0% background copy so that only data written to the source or target is copied to the target volume.
Clone	Creates an exact replica of the volume, which can be changed without affecting the original volume. After the copy operation completes, the mapping that was created by the preset is automatically deleted. This preset automatically creates a volume with the same properties as the source volume and creates a FlashCopy mapping with a background copy rate of 50. The FlashCopy mapping is configured to automatically delete itself when the FlashCopy mapping reaches 100% completion.
Backup	Creates a point-in-time replica of the production data. After the copy completes, the backup view can be refreshed from the production data, with minimal copying of data from the production volume to the backup volume. This preset automatically creates a volume with the same properties as the source volume. The preset creates an incremental FlashCopy mapping with a background copy rate of 50.

Presets: All of the presets can be adjusted by using the Advanced Settings expandable section in the GUI.

11.2 Managing FlashCopy using the GUI

IBM Storwize V7000 storage system provides a separate menu to access copy service management. Three selections are used for managing FlashCopy under the Copy Services panel:

- FlashCopy
- Consistency Groups
- FlashCopy Mappings

Figure 11-4 on page 339 shows the Copy Services dynamic menu.

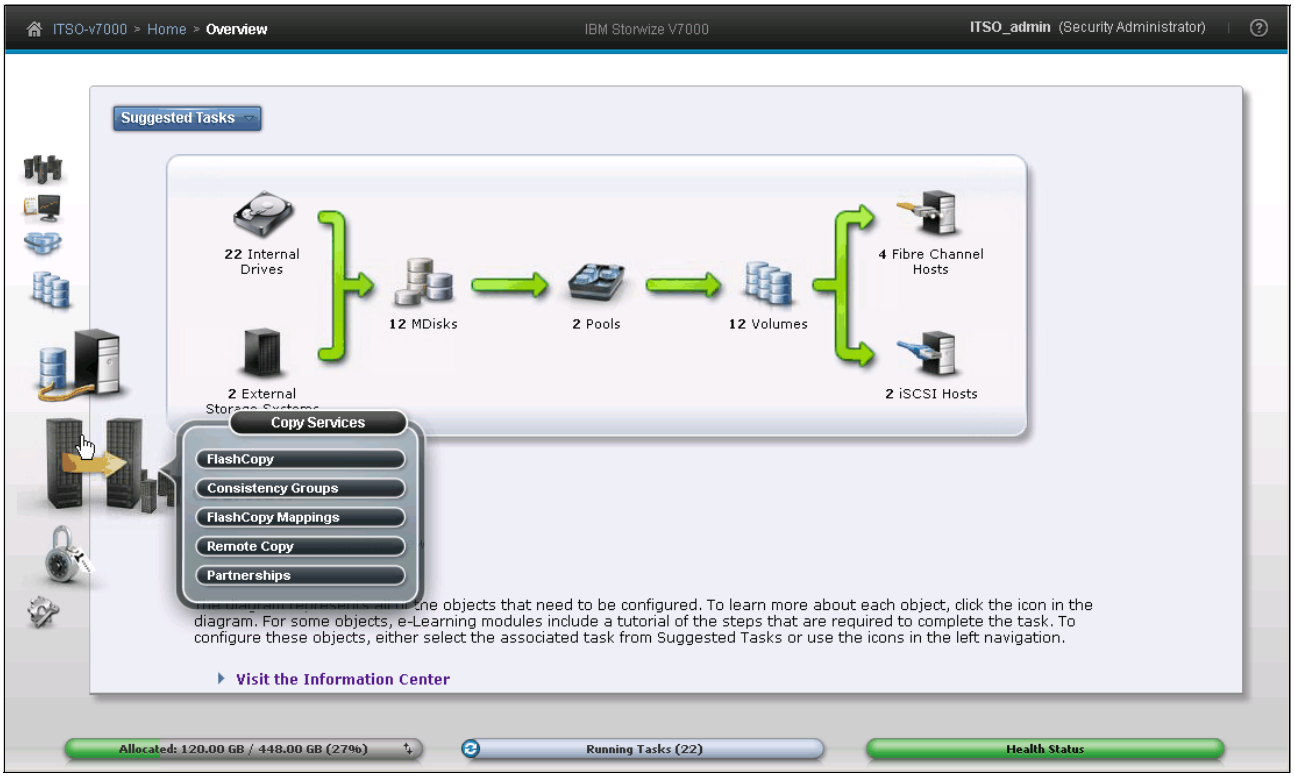

Figure 11-4 Copy Services menu

Most of the actions to manage the FlashCopy mapping can be done both in the FlashCopy panel or the FlashCopy Mappings panel. However, the quick path to create FlashCopy presets can be found only in the FlashCopy panel.

Click **FlashCopy** in the Copy Services dynamic menu, and the FlashCopy panel opens (Figure 11-5). In the FlashCopy panel, the FlashCopy mappings are sorted by volumes.

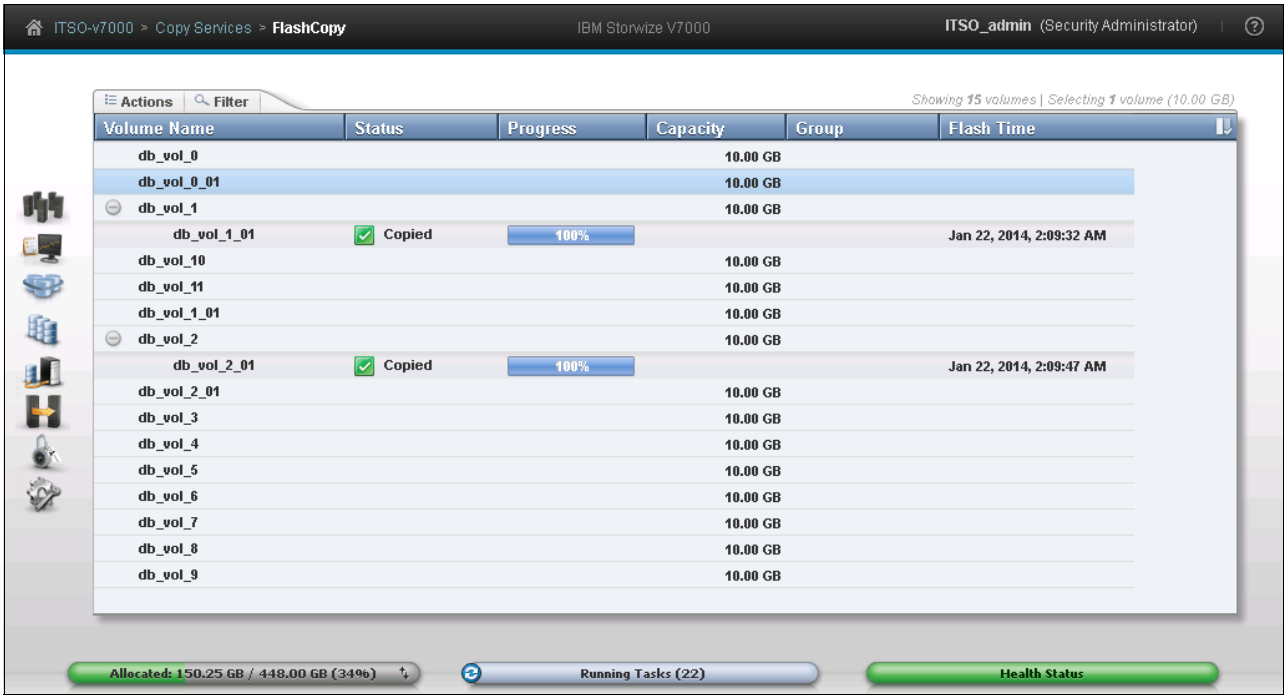

Figure 11-5 FlashCopy panel

Click **FlashCopy Mappings** in the Copy Services menu, and the FlashCopy Mappings panel opens (Figure 11-6). In the FlashCopy Mappings panel, the FlashCopy mappings are listed on separate rows.

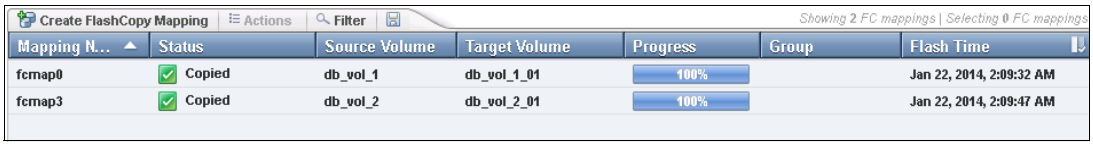

Figure 11-6 FlashCopy Mappings panel

The Consistency Groups panel can be used to manage the consistency groups for FlashCopy mappings. Click **Consistency Groups** in the Copy Services menu, and the Consistency Groups panel opens (Figure 11-7).

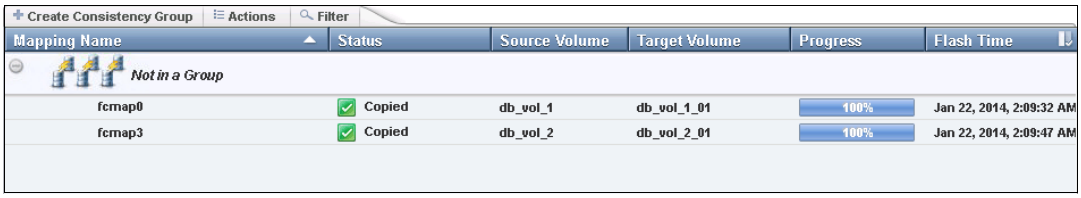

Figure 11-7 Consistency Groups panel

11.2.1 Quick path to create FlashCopy presets

Use the presets in the FlashCopy window to more easily create FlashCopy.

Creating a snapshot

In the FlashCopy window, right-click a volume and select **Create Snapshot** from the pop-up menu (Figure 11-8).

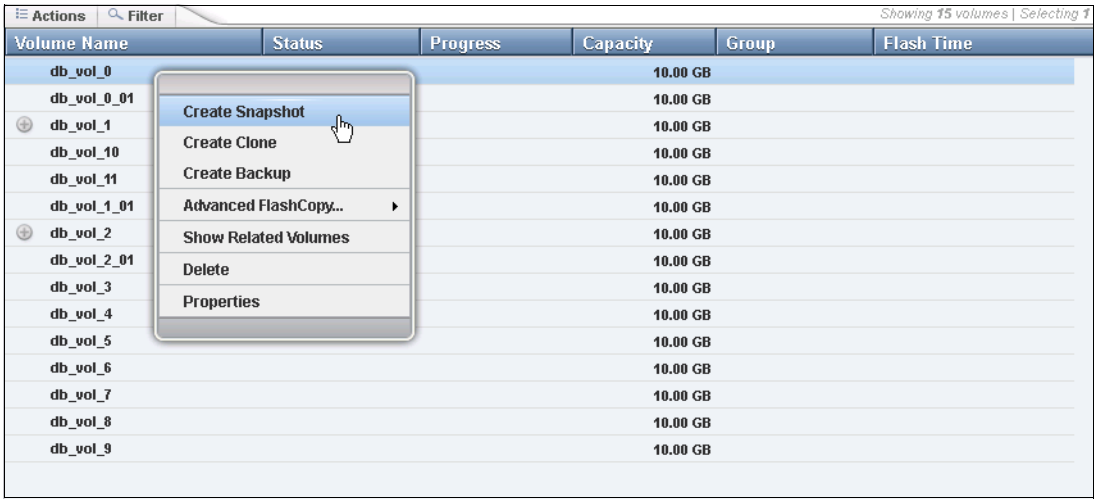

Figure 11-8 Create a snapshot from the preset

You now have a snapshot volume for the volume you choose.

Creating a clone

In the FlashCopy panel, right-click a volume and select **Create Clone** from the pop-up menu (Figure 11-9).

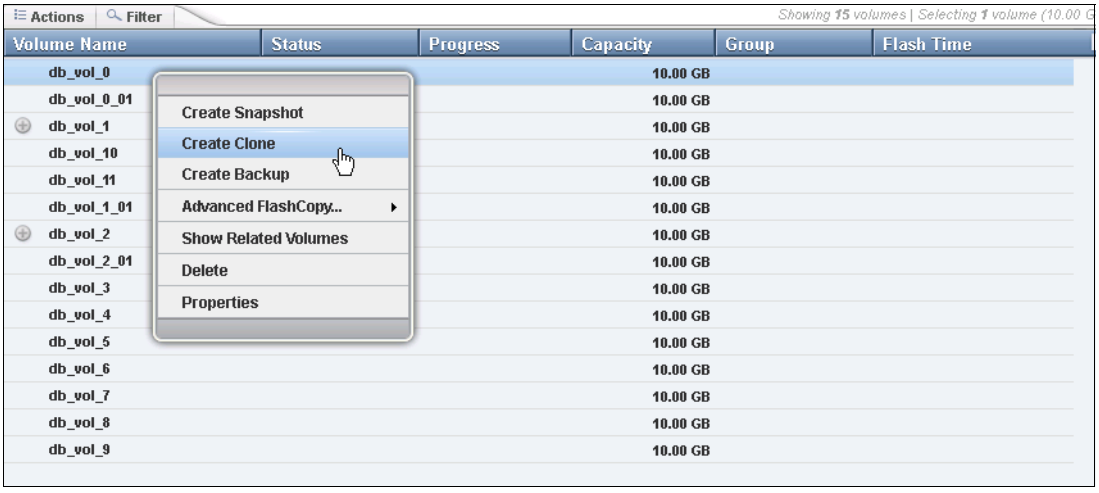

Figure 11-9 Create a clone from the preset

You now have a clone volume for the volume you choose.

Creating a backup

In the FlashCopy panel, right-click a volume and select **Create Backup** from the pop-up menu (Figure 11-10).

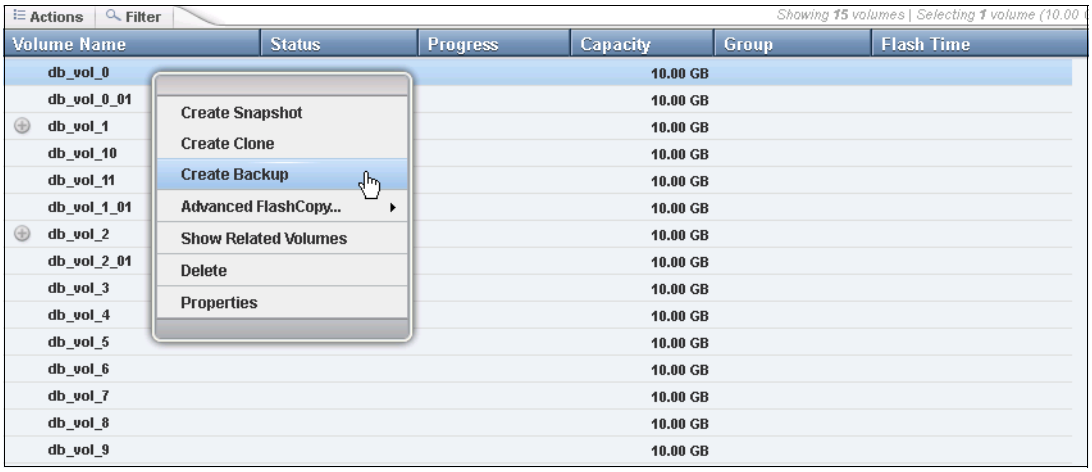

Figure 11-10 Create a backup from the preset

You now have a backup volume for the volume you choose.

Now, in the FlashCopy panel, you find three FlashCopy target volumes under the source volume (Figure 11-11). The progress bars next to the target volumes indicate the copy progress as a percentage. The copy progress remains 0% for snapshot; there is no change happening at the moment because only the changed data is being copied. At the same time, the copy progresses for clone and backup keep increasing.

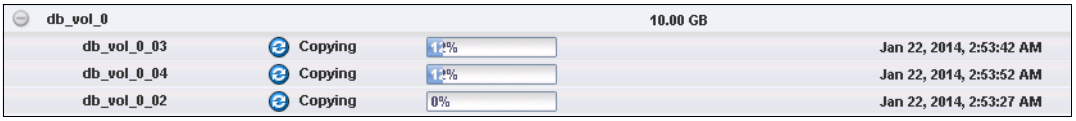

Figure 11-11 FlashCopy in progress viewed for the FlashCopy panel

The copy progress is also located under the Running Tasks status indicator (Figure 11-12). This view is slightly different than the Remote Copy Running Tasks status indicator.

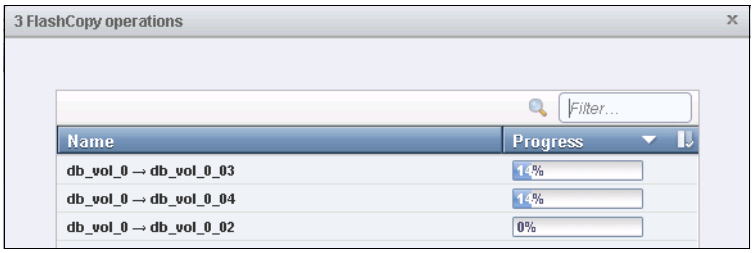

Figure 11-12 The FlashCopy progress viewed from the Running Tasks status indicator

After the copy progress is complete, the FlashCopy with the clone preset is automatically removed (Figure 11-13).

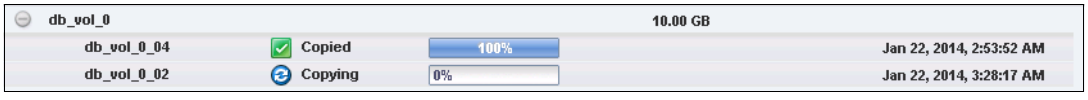

Figure 11-13 FlashCopy progresses complete

11.2.2 Managing FlashCopy mapping

The FlashCopy presets cover the most used FlashCopy configurations for general situations. However, customized FlashCopy mappings are still necessary in some more complex scenarios.

Creating customized FlashCopy mappings

You can create FlashCopy mappings through the FlashCopy panel. Right-click the volume you want to be the source volume for the FlashCopy mapping and select **Advanced FlashCopy** → ***(either Create New or Use Existing option)*** from the pop-up menu (Figure 11-14).

You have two options, which are described next:

- Create New Target Volumes
- Use Existing Target Volumes

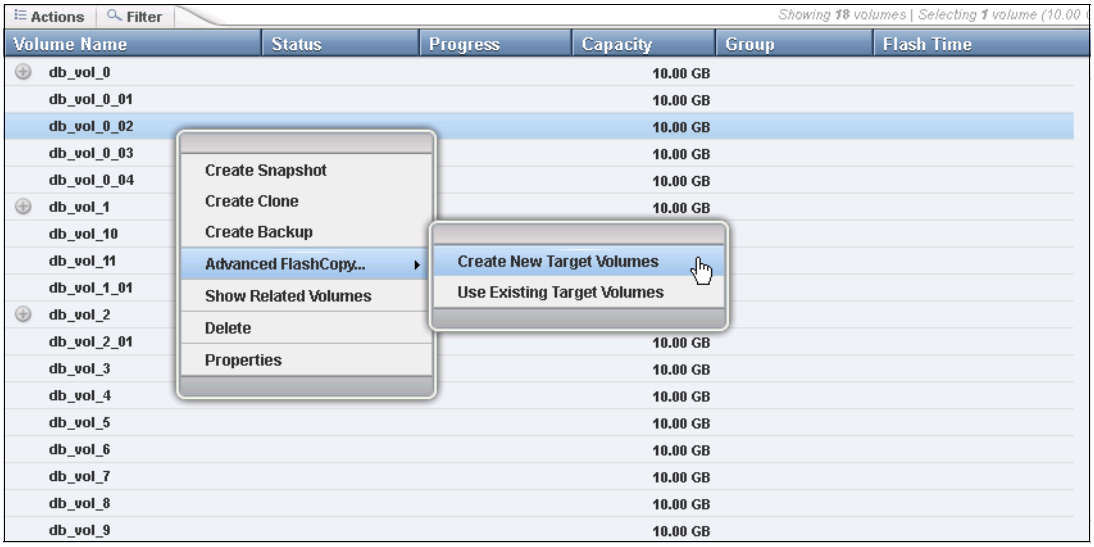

Figure 11-14 Create advanced FlashCopy

Create New Target Volumes

This is the option to use if you have not created the target volume yet.

1. Click **Create new target volumes**,
2. Choose the preset, but no matter which preset you choose, you can modify the setting of the FlashCopy mapping. Choose one preset that has the most similar configuration to the one required, and click **Advanced Settings** to make any appropriate adjustments on the properties (Figure 11-15 on page 344). Figure 11-16 on page 344 shows the default setting.

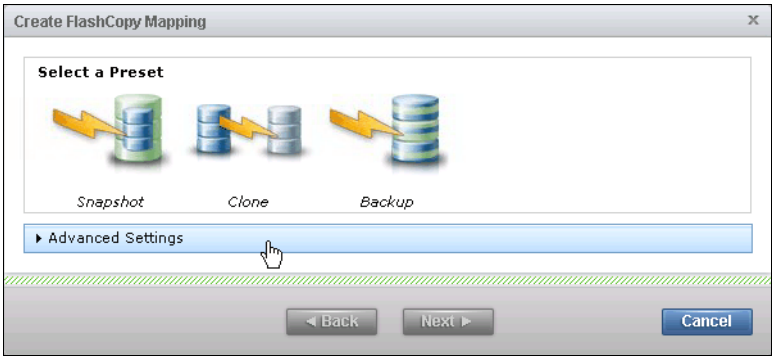

Figure 11-15 Choose a preset most similar to your requirement

Figure 11-16 Default setting for the snapshot preset

For example, if the snapshot preset was selected, the default settings can be found when you click **Advanced Settings**, which are as follows:

- Background Copy: 0
- Incremental: No
- Auto Delete after completion: No
- Cleaning Rate: 0

The same is true with clone preset, and the default settings of clone preset can be found when you click **Advanced Settings**, after you select clone preset (Figure 11-17 on page 345), which are as follows:

- Background Copy: 50
- Incremental: No
- Auto Delete after completion: Yes
- Cleaning Rate: 50

Figure 11-17 Default settings for the clone preset

The same is true with backup preset. The default settings of backup preset are found when you click **Advanced Settings**, after you select the backup preset (Figure 11-18), which are as follows:

- Background Copy: 50
- Incremental: Yes
- Auto Delete after completion: No
- Cleaning Rate: 50

Figure 11-18 Default settings for the backup preset

Change the settings of the FlashCopy mapping according to your requirements and click **Next**.

3. You can add your FlashCopy mapping to a consistency group (Figure 11-19 on page 346). If the consistency group is not ready, the FlashCopy mapping can be added to the consistency group after. Click **Next** to continue.

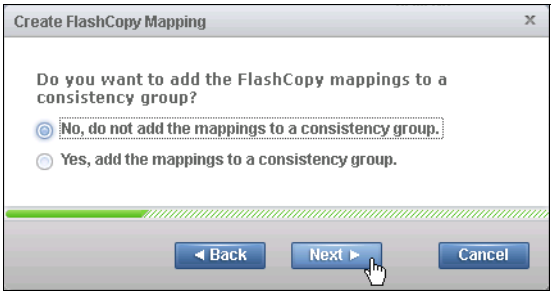

Figure 11-19 Add FlashCopy mapping to a consistency group

4. You are prompted to select how you want to manage volume capacity. The default choice is **Create a generic volume**. If you select a thin-provisioned or compressed volume, more options are available (Figure 11-20).

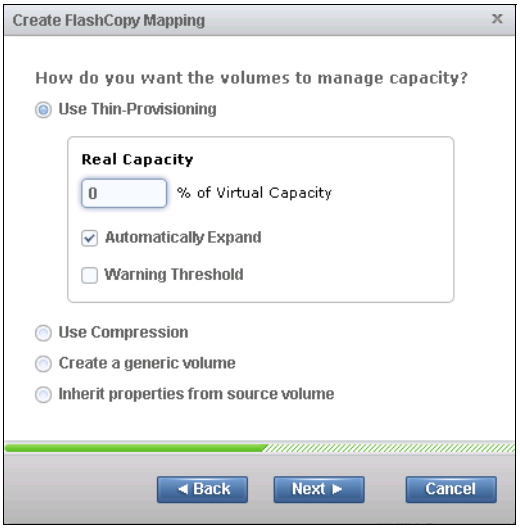

Figure 11-20 Create a thin provisioned target volume

If you chose the clone or backup preset at the beginning of this wizard, the default choice is **Create a generic volume** (Figure 11-21).

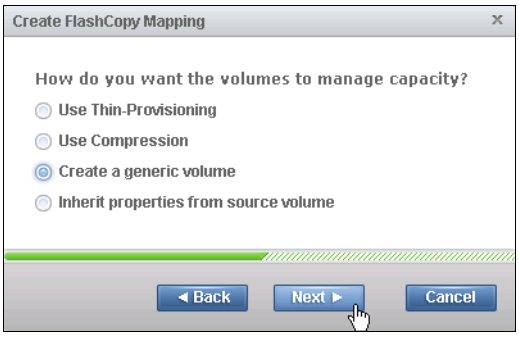

Figure 11-21 The default choice when creating a FlashCopy mapping

5. Select the generic preset and click **Next**.
6. You can choose from which storage pool you want to create your volume. As shown in Figure 11-22 on page 347, you can select the same storage pool that is used by the source volume. Click **Finish**.

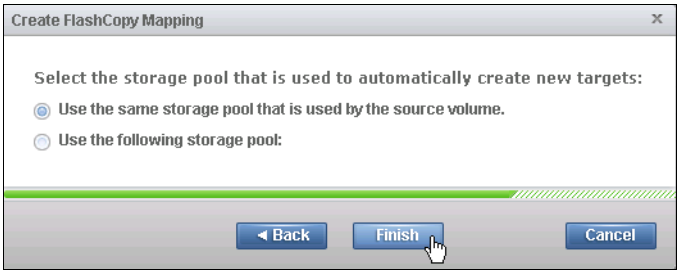

Figure 11-22 Choose use the same storage pool with the source volume

7. You can also specify another storage pool for your new volume (Figure 11-23).

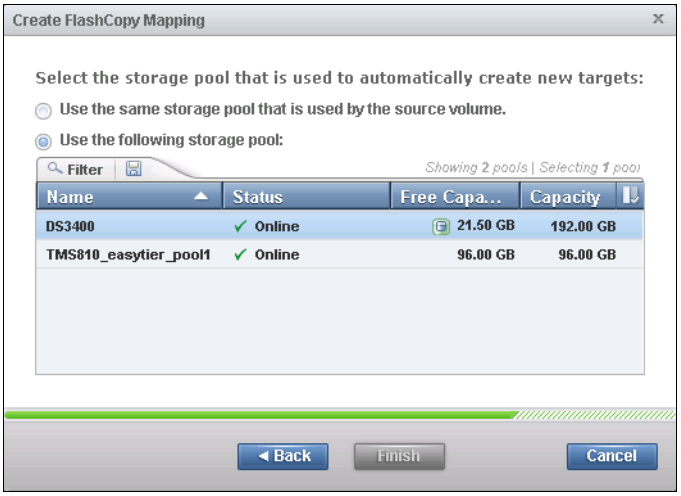

Figure 11-23 Choose another storage pool to create the volume

8. Click **Finish** after you make your decision. The FlashCopy mapping is created on your volume with a new target (Figure 11-24). The status of the newly created FlashCopy mapping is Idle. It can be started, as described in “Starting a FlashCopy mapping” on page 350.

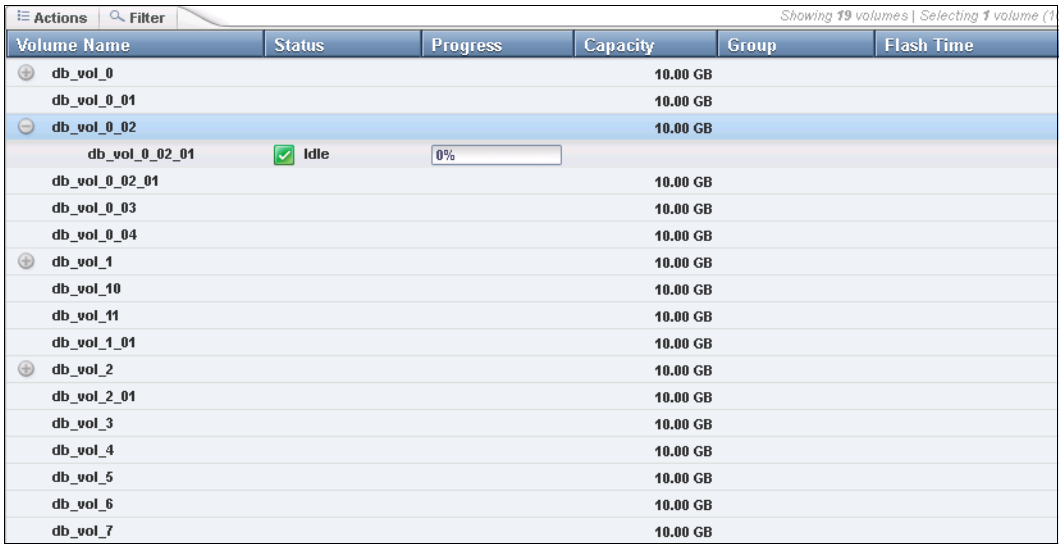

Figure 11-24 New FlashCopy mapping has been created with a new target

Use Existing Target Volumes

This is the option to use if you already have candidate target volumes.

1. In the Advanced FlashCopy menu, select **Use existing target volumes** (Figure 11-25).

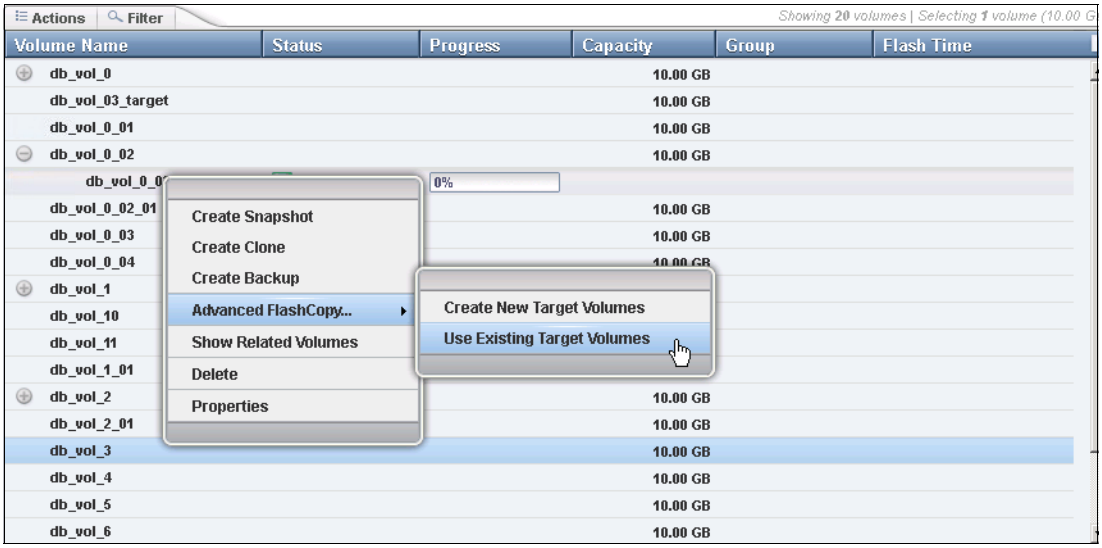

Figure 11-25 Create FlashCopy mapping using existing target volume

2. You must choose the target volume for the source volume you selected. Select the target volume from the drop-down menu in the right pane and click **Add** (Figure 11-26).

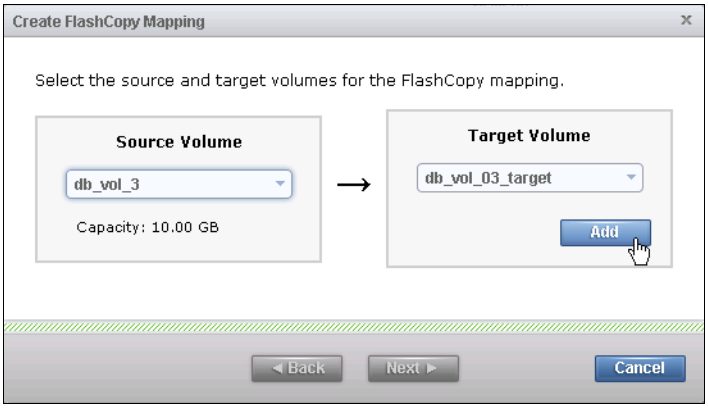

Figure 11-26 Select the target volume

3. The FlashCopy mapping is listed (Figure 11-27 on page 349). Click the red **X** if the FlashCopy mapping is not the one you want to create. If the FlashCopy mapping is what you want, click **Next** to continue.

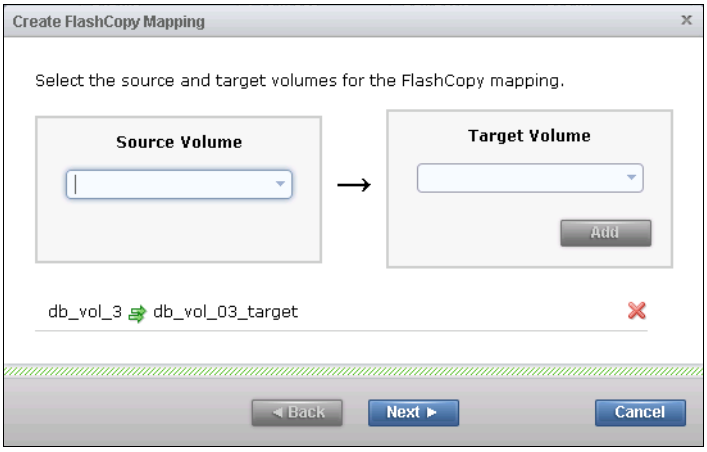

Figure 11-27 Add FlashCopy mapping

4. Select the preset and adjust the settings (Figure 11-28). Make sure that the settings meet your requirements, and click **Next**.

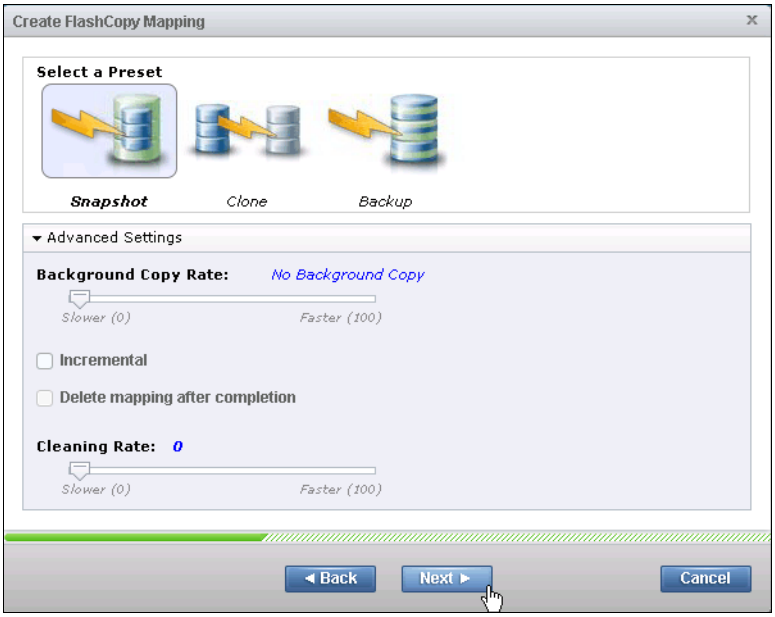

Figure 11-28 Select a preset and make your adjustments

5. Now you can add the FlashCopy mapping to a consistency group (Figure 11-29). Click **Finish** and the FlashCopy mapping is created with the status of Idle.

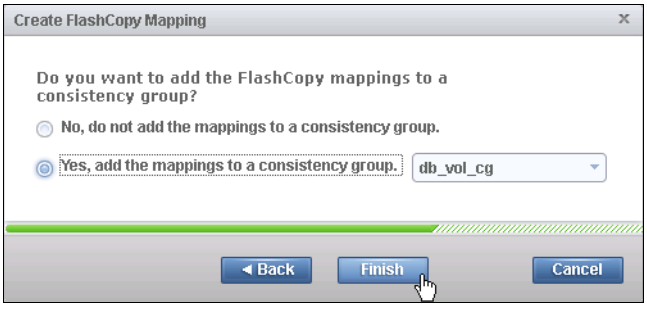

Figure 11-29 Select a consistency group to add the FlashCopy mapping to

You can also create the FlashCopy mappings in the FlashCopy Mappings panel by clicking **Create FlashCopy Mapping** at the upper left, as shown in Figure 11-30.

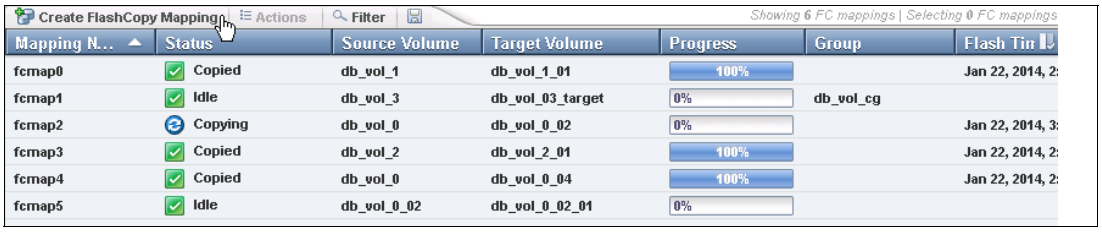

Figure 11-30 Create a FlashCopy mapping in the FlashCopy Mappings panel

A wizard opens to guide you through the creation of a FlashCopy mapping, and the steps are the same as for creating an advanced FlashCopy mapping using existing target volumes in the FlashCopy panel.

Starting a FlashCopy mapping

Most FlashCopy mapping actions can be done from both the FlashCopy panel or the FlashCopy Mappings panel. Although the steps for these actions are the same in both panels, we show only the steps in the FlashCopy panel.

To start the mapping, right-click the FlashCopy target volume in the FlashCopy panel and select **Start** from the pop-up menu (Figure 11-31). The status of the FlashCopy mapping changes from Idle to Copying.

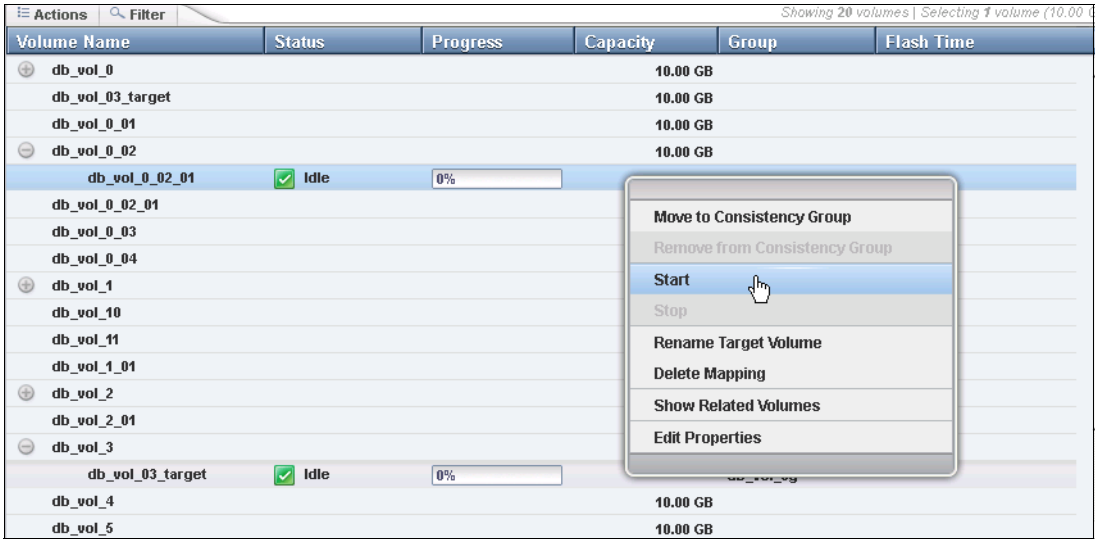

Figure 11-31 Start FlashCopy mapping

Stopping a FlashCopy mapping

To stop the FlashCopy mapping, right-click the FlashCopy target volume in the FlashCopy panel and select **Stop** from the pop-up menu (Figure 11-32). After the stopping process completes, the status of the FlashCopy mapping is changed to Stopped.

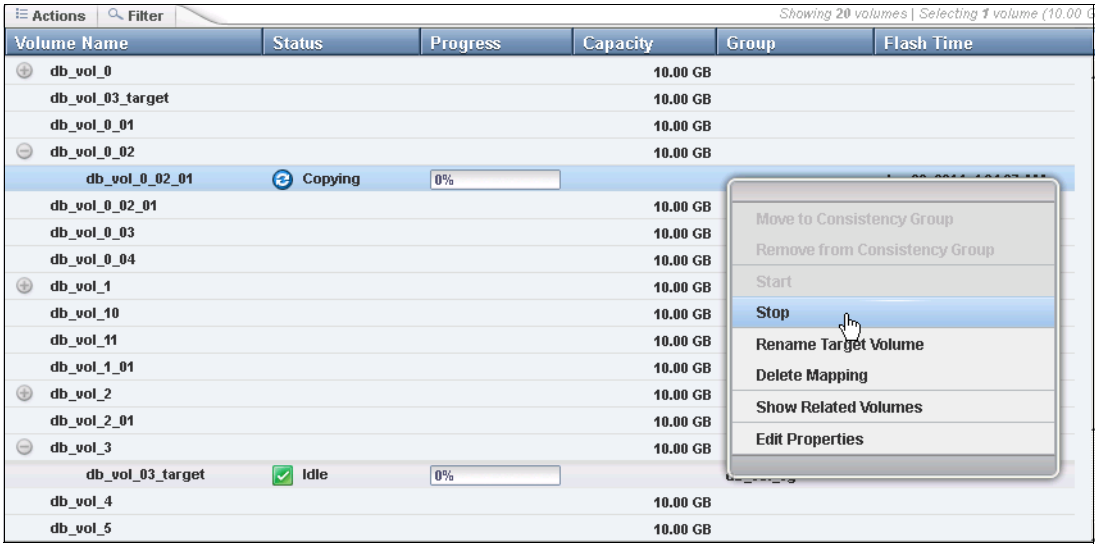

Figure 11-32 Stopping a FlashCopy mapping

Renaming the target volume

If the FlashCopy target volumes were created automatically by an IBM Storwize V7000 storage system, the name of the target volumes is the source volume name plus a suffix that contains numbers. The name of the target volumes can be changed to be more meaningful in your environment.

To change the name of the target volume, right-click the FlashCopy target volume in the FlashCopy panel and select **Rename Target Volume** from the pop-up menu (Figure 11-33).

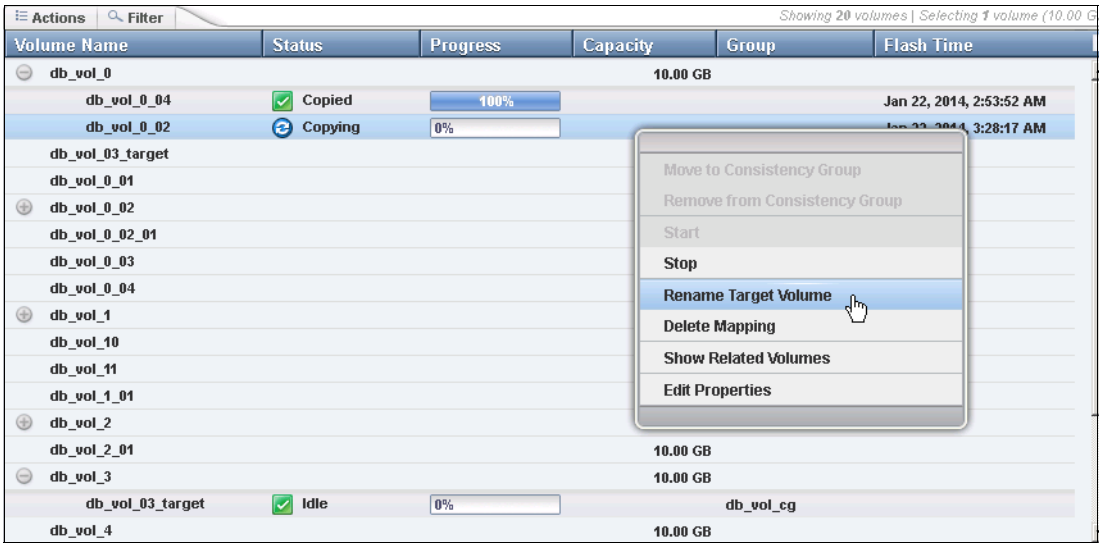

Figure 11-33 Rename a target volume

Input your new name for the target volume (Figure 11-34). Click **Rename** to finish.

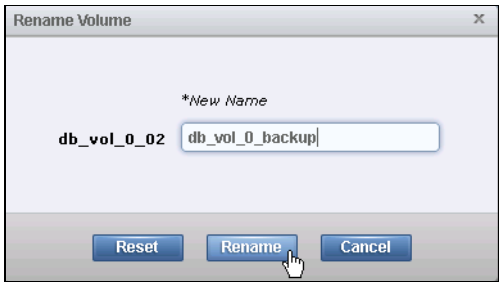

Figure 11-34 Rename a target volume

Renaming a FlashCopy mapping

The FlashCopy mappings are created with names that begin with *fcmap*. Change the name of FlashCopy mappings if you want it to be more meaningful to you.

To change the name of a FlashCopy mapping, right-click the FlashCopy mapping in the FlashCopy Mappings panel and select **Rename Mapping** from the pop-up menu (Figure 11-35).

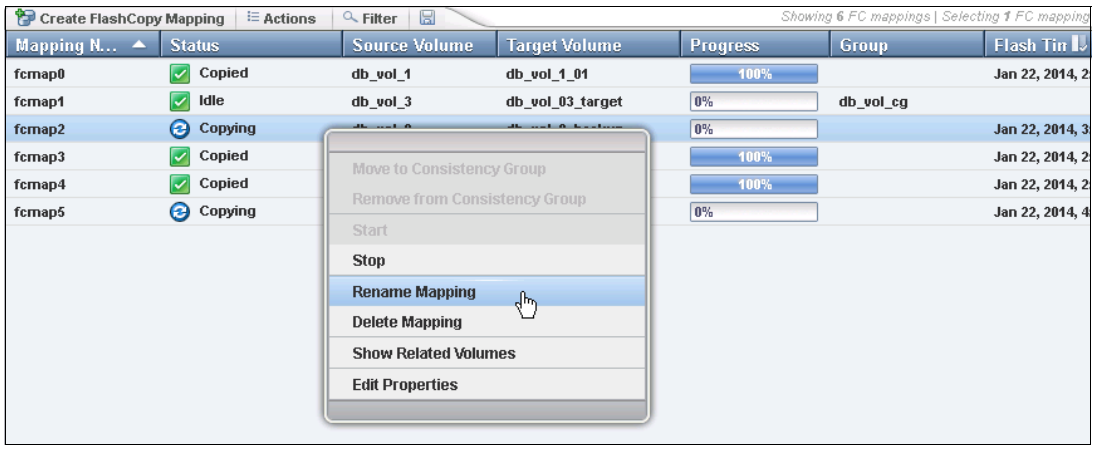

Figure 11-35 Rename a FlashCopy mapping

Enter a new name for the FlashCopy mapping (Figure 11-36). Click **Rename** to finish.

Figure 11-36 Enter a new name for the FlashCopy mapping

Deleting a FlashCopy mapping

To delete FlashCopy mapping, stop it first, then right-click the FlashCopy target volume in the FlashCopy panel and select **Delete Mapping** from the pop-up menu (Figure 11-37).

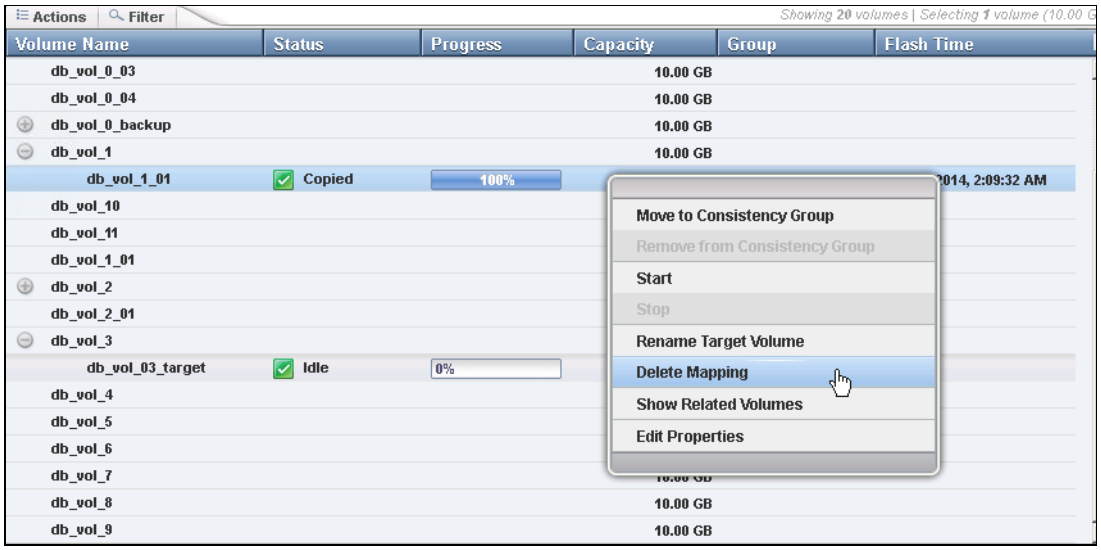

Figure 11-37 Select Delete Mapping

FlashCopy Mapping state: If the FlashCopy mapping is in the Copying state, it must be stopped before being deleted.

Confirm your action to delete FlashCopy mappings in the window that opens (Figure 11-38). Verify the number of FlashCopy mappings that you need to delete, and if you want to delete the FlashCopy mappings while the data on the target volume is inconsistent with the source volume, select the check box. Click **Delete**. Your FlashCopy mapping is removed.

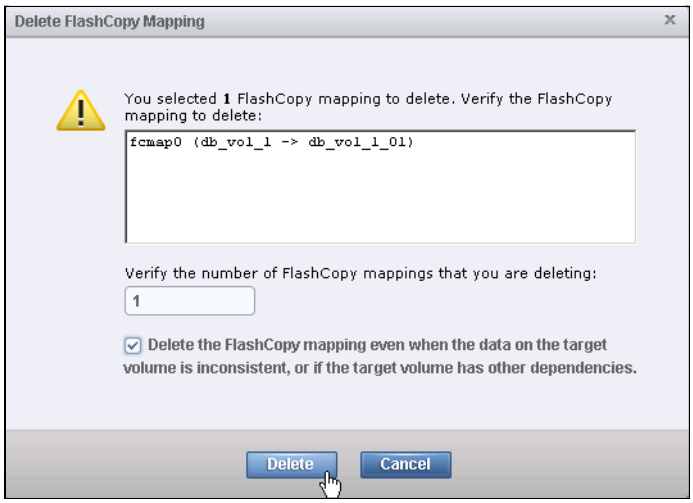

Figure 11-38 Confirm the deletion of FlashCopy mappings

Target volume not deleted: Deleting the FlashCopy mapping does not delete the target volume. If you need to reclaim the storage space occupied by the target volume, you need to delete the target volume manually.

Showing related volumes

To show the FlashCopy volume dependency, right-click the FlashCopy source volume in the FlashCopy panel and select **Show Related Volumes** from the pop-up menu (Figure 11-39).

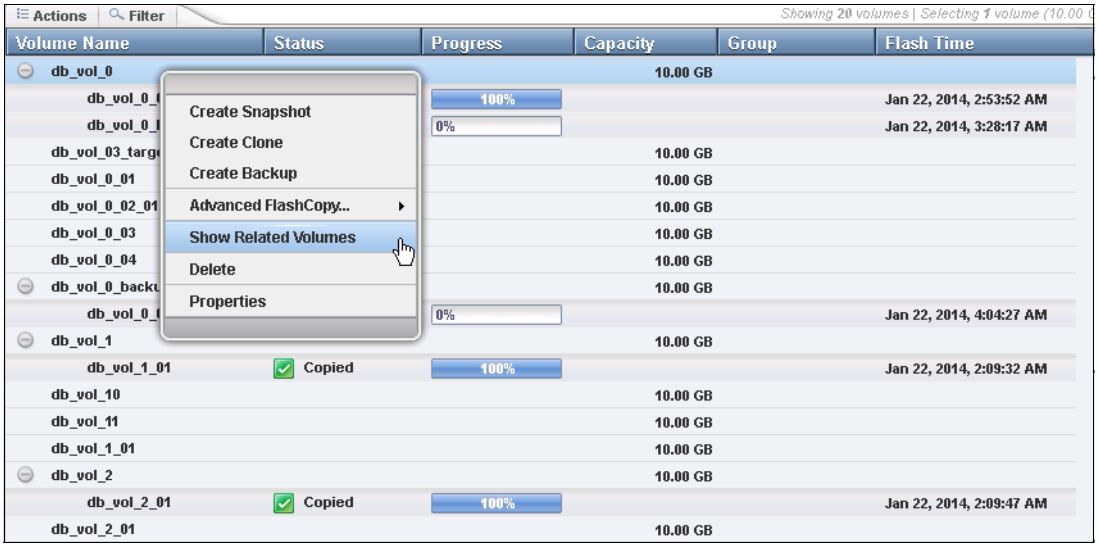

Figure 11-39 Show Related Volumes menu

The FlashCopy mapping dependency tree opens (Figure 11-40).

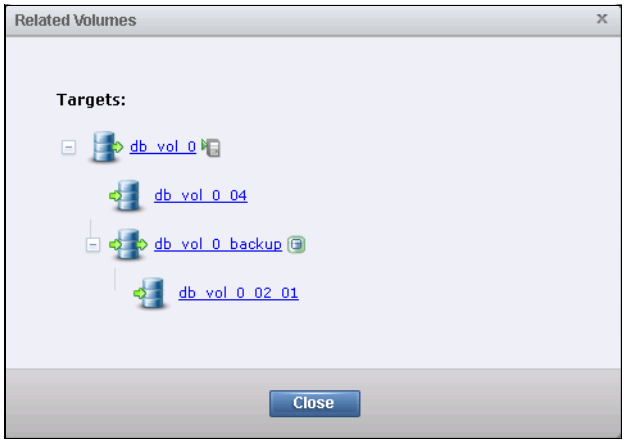

Figure 11-40 FlashCopy related volumes

Editing properties

To change the background copy rate and cleaning rate after the FlashCopy mapping is created, right-click the FlashCopy target volume in the FlashCopy panel and select **Edit Properties** from the pop-up menu (Figure 11-41).

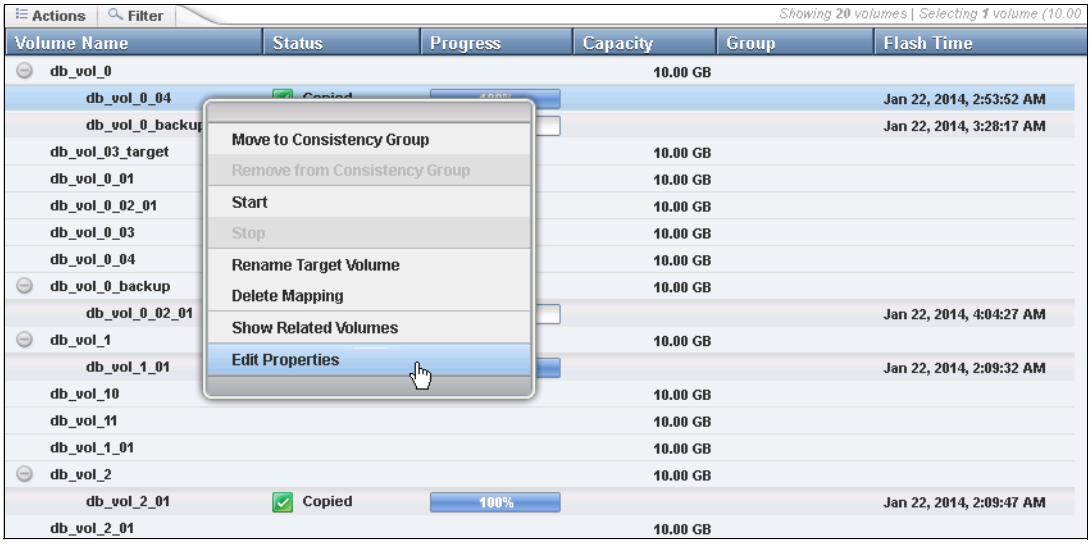

Figure 11-41 Edit Properties menu

You can then modify the value of the background copy rate and cleaning rate by moving the pointers on the bars (Figure 11-42). Click **Save** to save changes.

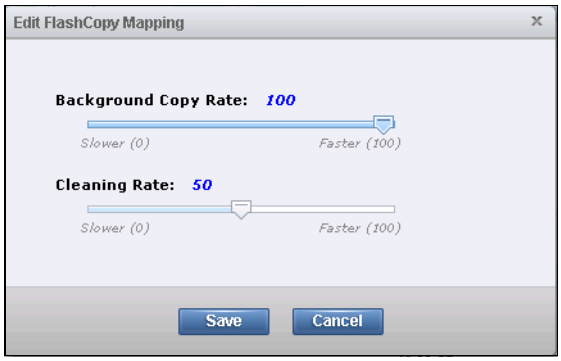

Figure 11-42 Change the copy rate

11.2.3 Managing a FlashCopy consistency group

To manage FlashCopy consistency groups, click the **Consistency Groups** selection on the Copy Services dynamic menu (Figure 11-43).

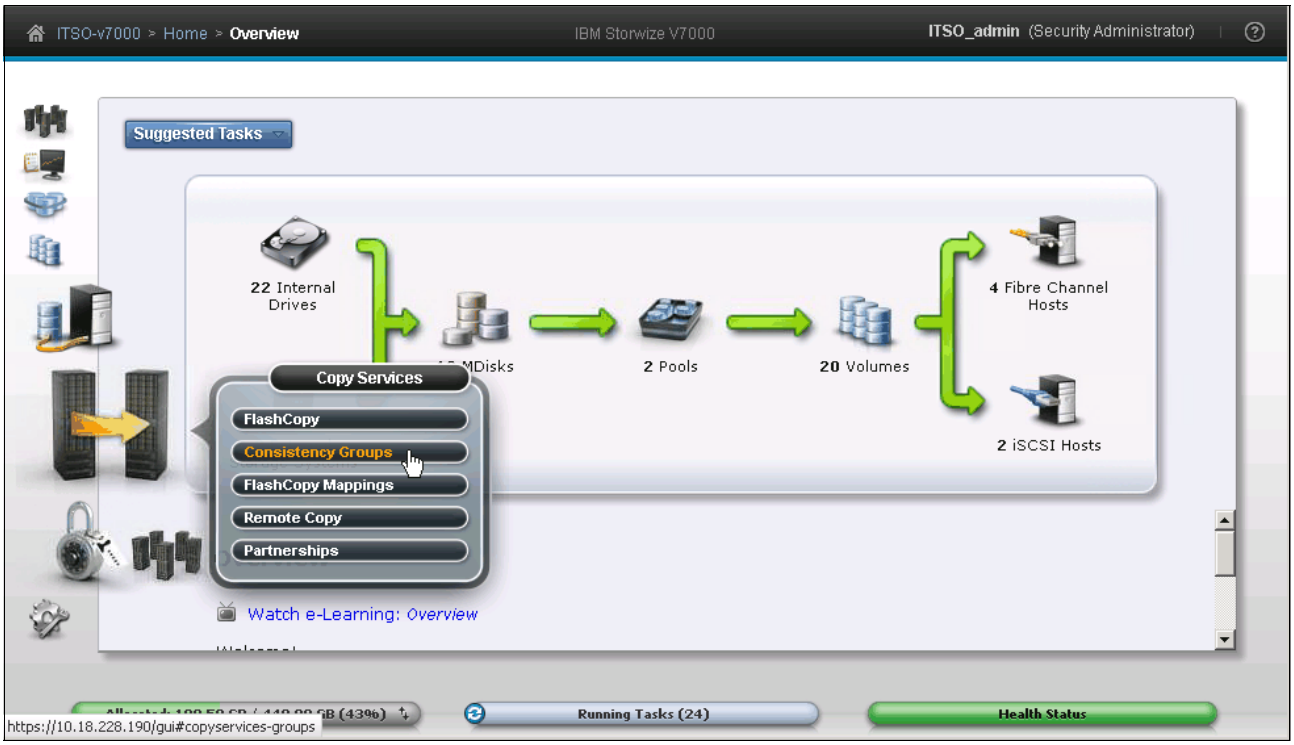

Figure 11-43 Access to the Consistency Groups panel

The Consistency Groups panel (Figure 11-44) is where you can manage both consistency groups and FlashCopy mappings.

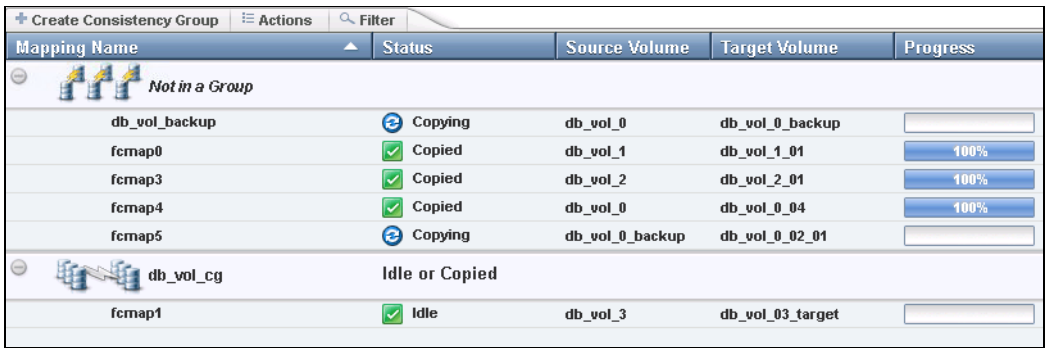

Figure 11-44 Consistency Groups panel

In the left pane of the Consistency Groups panel, you can list the consistency groups you need. Click **Not in a Group**, and then expand your selection by clicking the plus icon next to **Not in a Group**. All the FlashCopy mappings that are not in any consistency groups are displayed in the lower pane.

In the lower pane of the Consistency Groups panel, you can discover the properties of a consistency group and the FlashCopy mappings in it. You can also take actions on any consistency groups and FlashCopy mappings within the Consistency Groups panel, as

allowed by their state. All the actions allowed for the FlashCopy mapping are described in 11.2.2, “Managing FlashCopy mapping” on page 343.

Creating a FlashCopy consistency group

To create a FlashCopy consistency group, click **Create Consistency Group** at the top of the Consistency Groups panel (Figure 11-44 on page 356).

When you are prompted, enter the name of the new consistency group (Figure 11-45). Following your naming conventions, enter the name of the new consistency group in the box and click **Create**.

Figure 11-45 Enter the name for the consistency group

After the creation process completes, you find a new consistency group in the left pane of the Consistency Groups panel. Select the new consistency group, and you see more detailed information about this consistency group in the right pane (Figure 11-46).

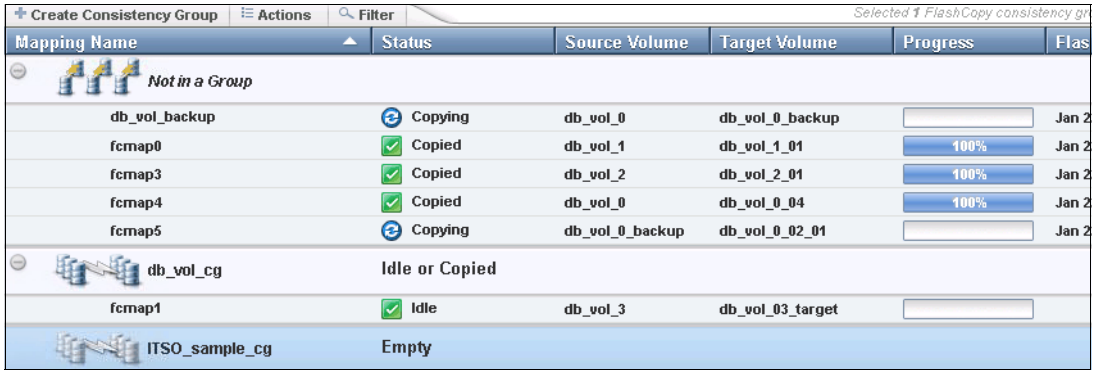

Figure 11-46 New consistency group

You can rename the consistency group by clicking the name of the consistency group in the main pane and enter a name for it (following your naming convention). Next to the name of the consistency group, the state shows that it is now an empty consistency group with no FlashCopy mapping in it.

Adding FlashCopy mappings to a consistency group

Click **Not in a Group** to list all the FlashCopy mappings with no consistency group. To add FlashCopy mappings to a consistency group, select them and then select **Move to Consistency Group** from the pop-up menu (Figure 11-47 on page 358). To select a sequential range of mappings, select a mapping, press the Shift key, and click the last item in the range. To non-sequentially select mappings, press the Ctrl key while you click each individual mapping.

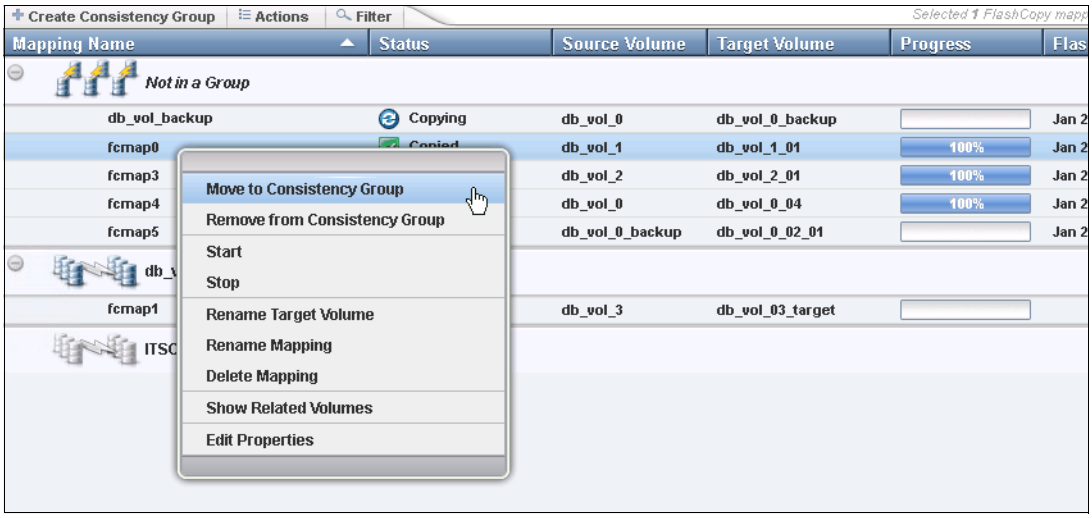

Figure 11-47 Select the FlashCopy mappings to add to a consistency group

When prompted (Figure 11-48), specify the consistency group into which you want to move the FlashCopy mapping. Click **Move to Consistency Group** to continue.

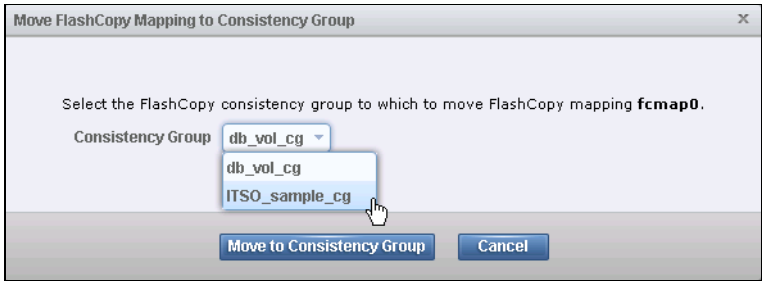

Figure 11-48 Select consistency group

After the action completes, the FlashCopy mappings you selected are removed from the Not In a Group list of the consistency group you chose.

Starting a consistency group

To start a consistency group, click **Start** from the pop-up menu (Figure 11-49).

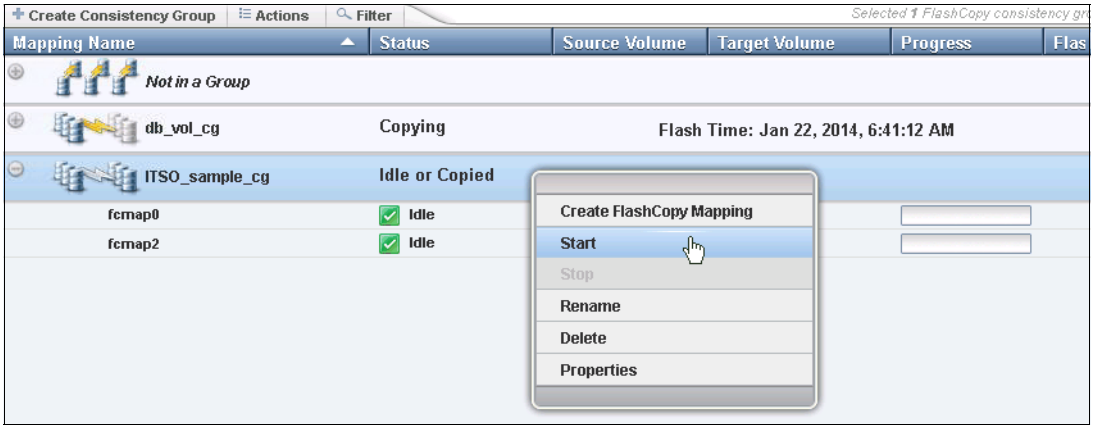

Figure 11-49 Start a consistency group

After you start the consistency group, all the FlashCopy mappings in the consistency group are started. The state of the FlashCopy changes to Copying (Figure 11-50).

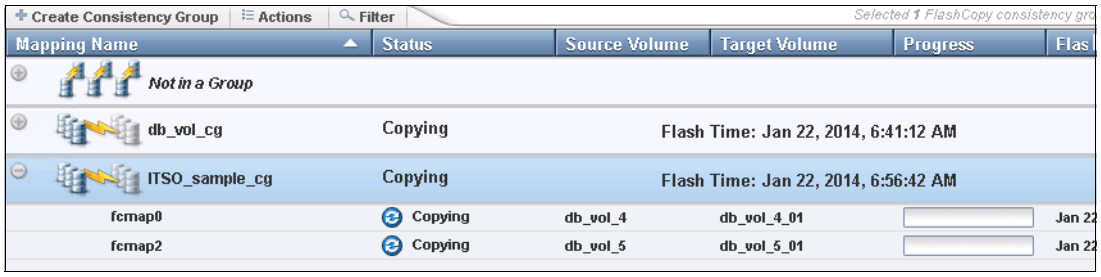

Figure 11-50 Consistency group start completes

Stopping a consistency group

To stop the consistency group, right-click a consistency group and then select **Stop** from the pop-up menu (Figure 11-51).

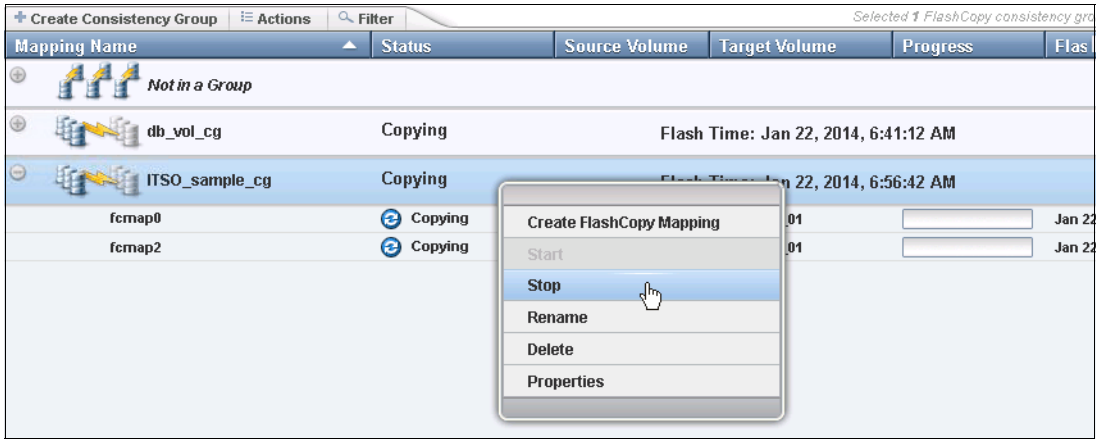

Figure 11-51 Stop a consistency group

After the stop process completes, the FlashCopy mappings in the consistency group are in the Stopped state, and a red circle with an “X” in it overlays the function icon of this consistency group to indicate an alert (Figure 11-52).

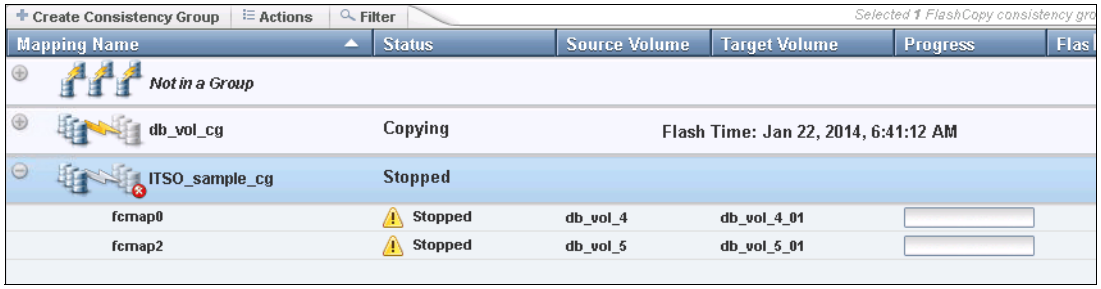

Figure 11-52 Consistency group stop completes

Previously copied relationships that were added to a consistency group that was later stopped before all members of the consistency group completed synchronization do not go out of the Copied state.

Removing FlashCopy mappings from a consistency group

To remove the FlashCopy mappings from a consistency group, right-click the FlashCopy mappings and select **Remove from Consistency Group** from the pop-up menu (Figure 11-53).

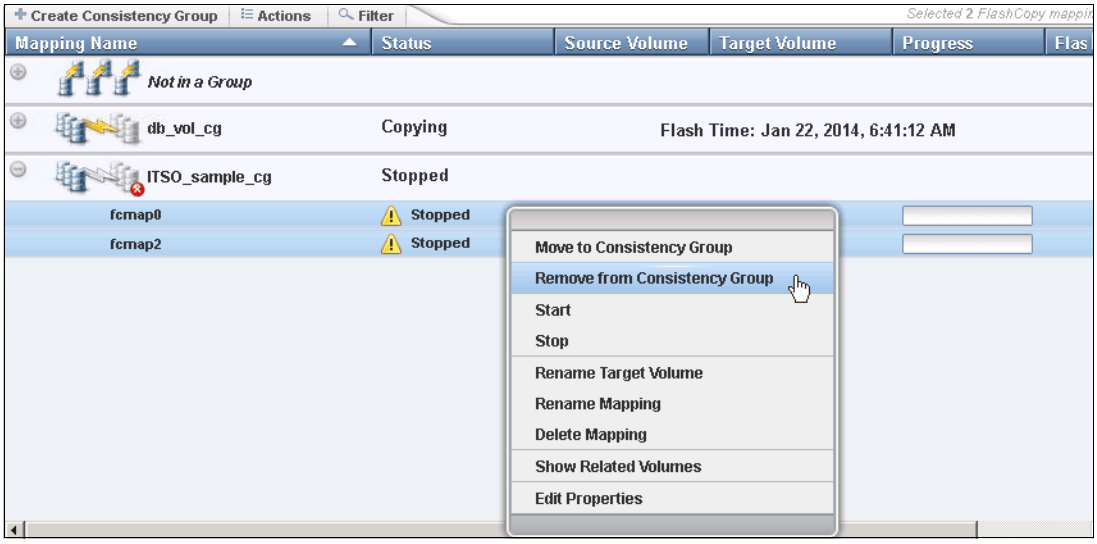

Figure 11-53 Remove from consistency group

The FlashCopy mappings are returned to the Not in a Group list after being removed from the consistency group.

Deleting a consistency group

To delete a consistency group, right-click a consistency group and then select **Delete** from the pop-up menu (Figure 11-54).

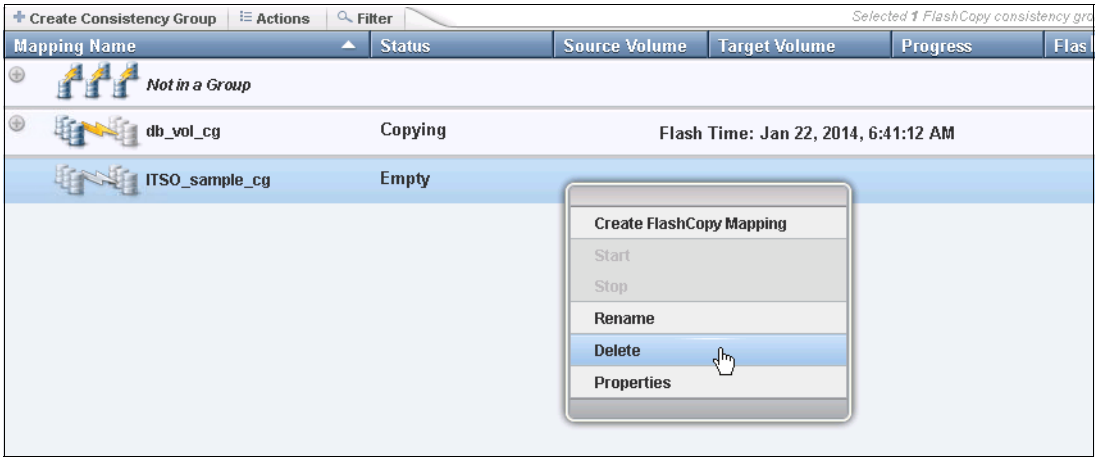

Figure 11-54 Delete a consistency group

11.3 Volume mirroring and migration options

Volume mirroring is a simple function that is like RAID 1 that allows a volume to remain online even when the storage pool backing becomes inaccessible. This function protects the volume from storage infrastructure failures by providing seamless mirroring between storage pools.

This function is provided by a specific volume mirroring function in the I/O stack, and cannot be manipulated like a FlashCopy or other types of copy volumes. This feature does, however, provide migration functionality, which can be obtained by splitting the mirrored copy from the source or using the **migrate to** command. This feature does not control back-end storage mirror or replication.

With this feature, host I/O completes when both copies are written. Before Version 6.3, this feature would take a copy offline when it had an I/O timeout, and then resynchronize with the online copy after it recovered. Since Version 6.3, this feature has been enhanced with a tunable latency tolerance. This feature provides an option to give preference to losing the redundancy between the two copies. This tunable timeout value is either Latency or Redundancy.

The Latency tuning option (set by running **svctask chvdisk -mirrowritepriority latency**) is the default and was the behavior found in releases before Version 6.3, and prioritizes host I/O latency. This setting produces a preference to host I/O over availability.

However, you might need to give preference to Redundancy when availability is more important than I/O response time. To accomplish this task, run **svctask chvdisk -mirrowritepriority redundancy**.

Regardless of which option you choose, volume mirroring can provide extra protection for your environment.

Several options are available regarding migration:

- Export to Image mode: This option allows you to move storage from *managed mode* to *image mode*, which is helpful if you are using the IBM Storwize V7000 as a migration device. For example, vendor A's product cannot communicate with vendor B's product, but you need to migrate existing data from vendor A to vendor B. By using the Export to Image mode, you can migrate data using copy services functions and then return control to the native array, while maintaining access to the hosts.
- Import to Image mode: This option allows you to import an existing storage MDisk or LUN with existing data, from an external storage system without putting metadata on it, so the existing data remains intact. After they are imported, all copy services functions maybe used to migrate the storage to other locations, while the data remains accessible to your hosts.
- Volume migration using Volume Mirroring then Split into New Volume: This option allows you to use the RAID 1 functionality to create two copies of data that initially have a set relationship (one primary and one secondary) but then break the relationship (both primary and no relationship) and make them independent copies of data. You can use this mode to migrate data between storage pools and devices. You might use this option if you want to move volumes to multiple different storage pools. You can only mirror one volume at a time.
- Volume migration using Migrate to Another Pool: This option allows any volume to be migrated between storage pools without interruption to host access. This option is effectively a quicker version of clicking **Volume Mirroring** → **Split into New Volume**. You might use this option to move volumes in a single step or if you do not already have a volume mirror copy.

Migration methods: Although the migration methods listed do not distribute access, you must provide a brief outage to install the host drivers for your IBM Storwize V7000. See *IBM System Storage SAN Volume Controller Host Attachment User's Guide*, SC26-7905 for more details. Make sure to consult the revision of the document that applies for your Storwize V7000.

11.3.1 Managing volume mirror and migration with the GUI

To make a volume mirror using the GUI, click **Actions** → **Add Mirrored Copy** from the **Volumes** window (Figure 11-55).

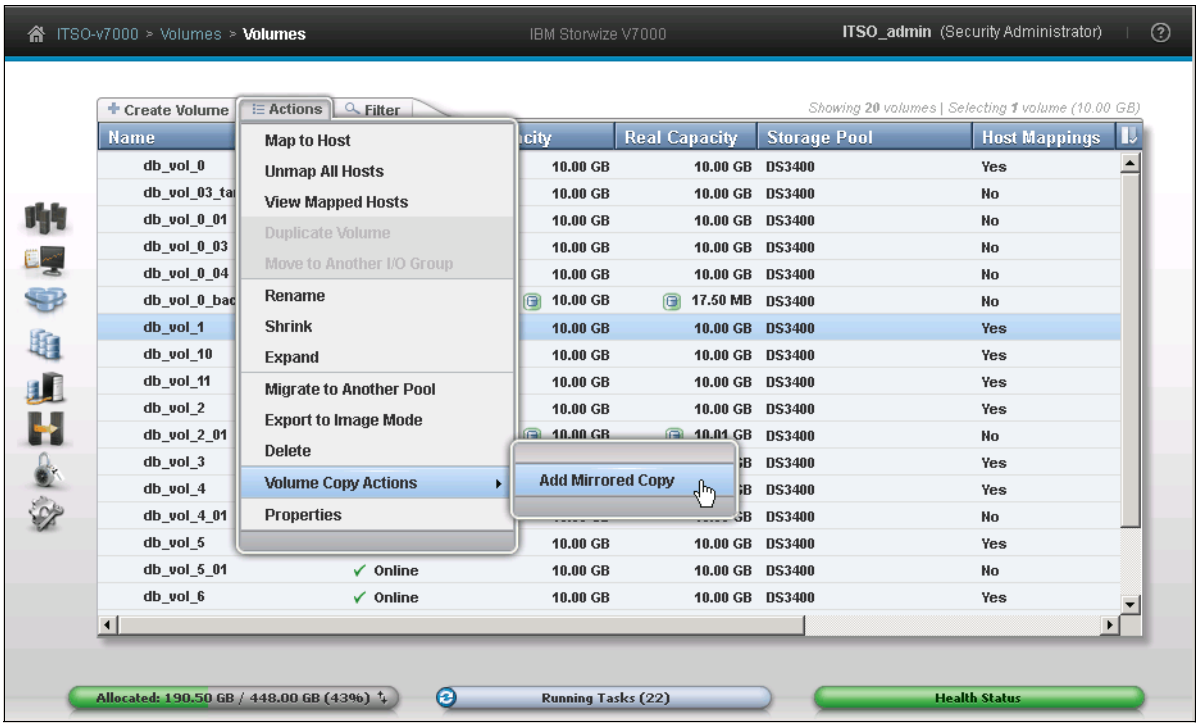

Figure 11-55 Add Mirrored Copy

Next, you can specify the type of volume mirror (generic, thin-provisioned, or compressed) and select the storage pool to use for the copy (Figure 11-56). Be sure to select a storage pool with sufficient space and similar performance characteristics, and then click **Add Copy**.

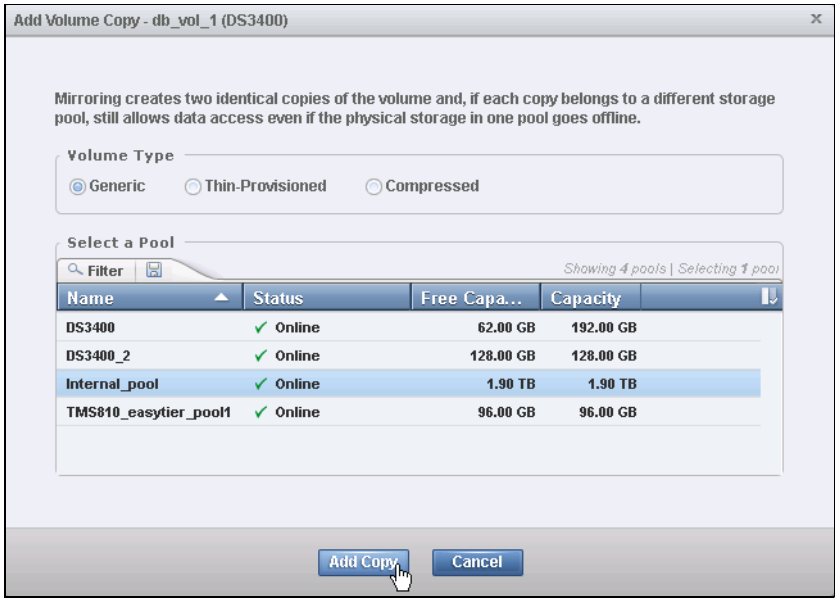

Figure 11-56 Confirm Volume Mirror type and storage pool to use for the mirror

After you create your mirror, you can view the distribution of extents (Figure 11-57) or you can view the mirroring progress percentage by using Running Tasks (Figure 11-58).

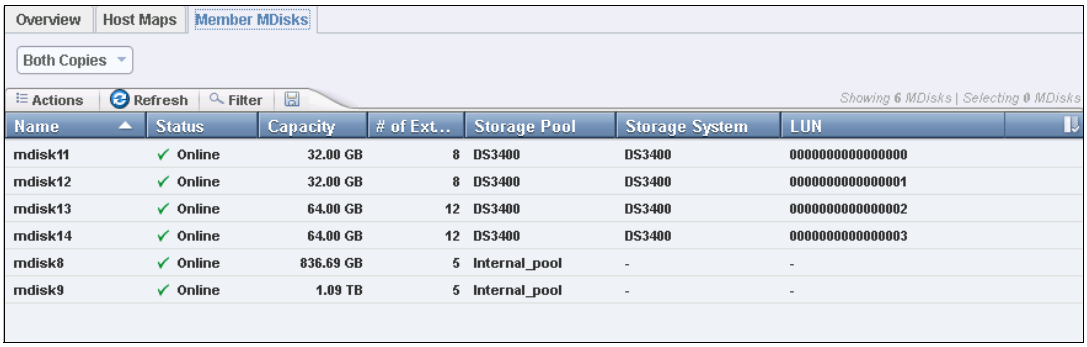

Figure 11-57 The distribution of extents for primary and mirror copy of a volume

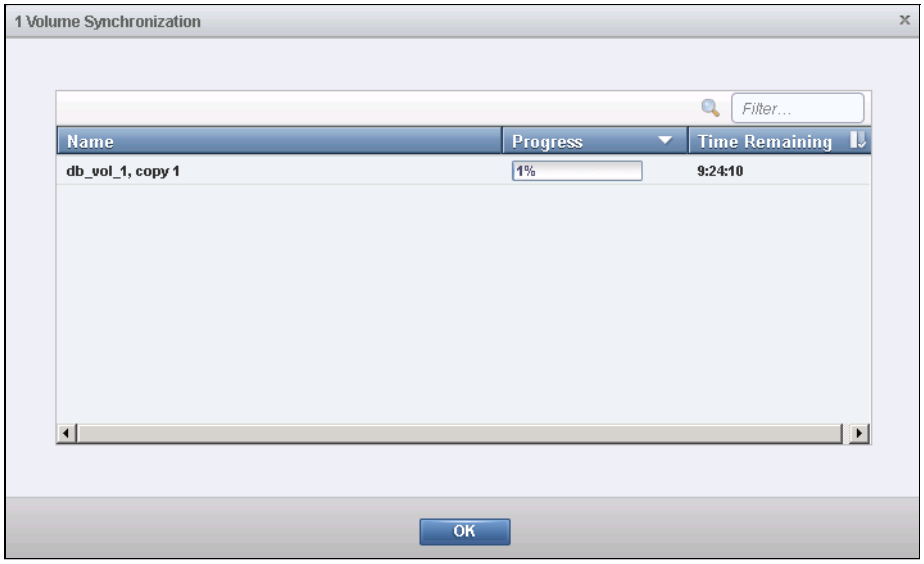

Figure 11-58 Progress of a mirror copy creation as viewed by using Running Tasks

Extent distribution: Extent distribution for the mirror copy is automatically balanced as well as possible within the storage pool selected.

After the copy completes, you have the option of splitting either copy of the mirror into a new stand-alone volume (Figure 11-59).

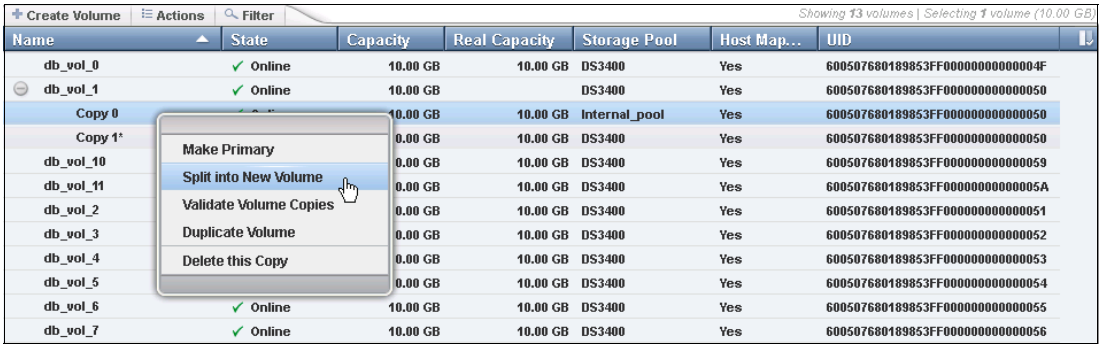

Figure 11-59 Selection of Split into New Volume

If you clicked **Split into New Volume** on either Copy0 or Copy1, you can specify a new volume name and confirm the split (Figure 11-60).

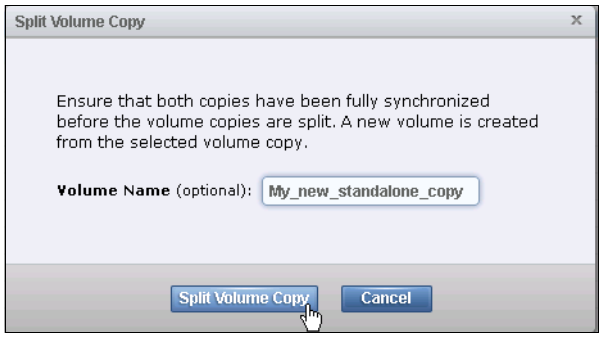

Figure 11-60 Confirmation of volume mirror split

After providing a new volume name (optional) and confirming the split, the results are available, as in Figure 11-61.

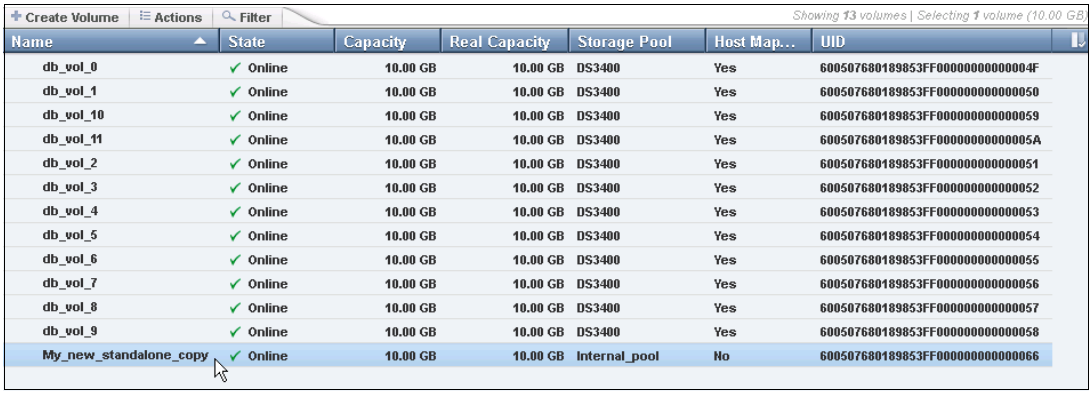

Name	State	Capacity	Real Capacity	Storage Pool	Host Map...	UID
db_vol_0	Online	10.00 GB	10.00 GB	DS3400	Yes	600507680189853FF00000000000004F
db_vol_1	Online	10.00 GB	10.00 GB	DS3400	Yes	600507680189853FF00000000000050
db_vol_10	Online	10.00 GB	10.00 GB	DS3400	Yes	600507680189853FF00000000000059
db_vol_11	Online	10.00 GB	10.00 GB	DS3400	Yes	600507680189853FF0000000000005A
db_vol_2	Online	10.00 GB	10.00 GB	DS3400	Yes	600507680189853FF00000000000051
db_vol_3	Online	10.00 GB	10.00 GB	DS3400	Yes	600507680189853FF00000000000052
db_vol_4	Online	10.00 GB	10.00 GB	DS3400	Yes	600507680189853FF00000000000053
db_vol_5	Online	10.00 GB	10.00 GB	DS3400	Yes	600507680189853FF00000000000054
db_vol_6	Online	10.00 GB	10.00 GB	DS3400	Yes	600507680189853FF00000000000055
db_vol_7	Online	10.00 GB	10.00 GB	DS3400	Yes	600507680189853FF00000000000056
db_vol_8	Online	10.00 GB	10.00 GB	DS3400	Yes	600507680189853FF00000000000057
db_vol_9	Online	10.00 GB	10.00 GB	DS3400	Yes	600507680189853FF00000000000058
My_new_standalone_copy	Online	10.00 GB	10.00 GB	Internal_pool	No	600507680189853FF00000000000066

Figure 11-61 Results of volume mirror split

Splitting volume copies: When you split a volume copy, the view returns to the pool in which it was created, not where the primary copy existed.

If you want to migrate your volumes to another storage pool in one step instead of two, you can use the **Migrate to Another Pool** option (Figure 11-62).

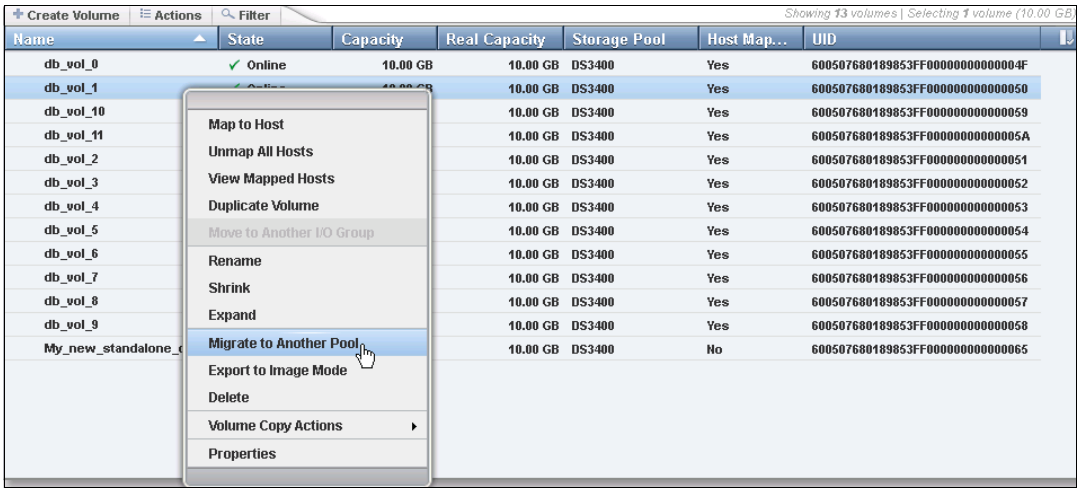

Figure 11-62 Using the Migrate to Another Pool option

Migration limitation: You cannot migrate more than one volume at a time. For this reason, Copy Services functions are more expedient, if available.

If the volume has only one copy, a storage pool selection window opens. If it has two, the same window with a slight variation that allows you to choose the copy to migrate opens (Figure 11-63).

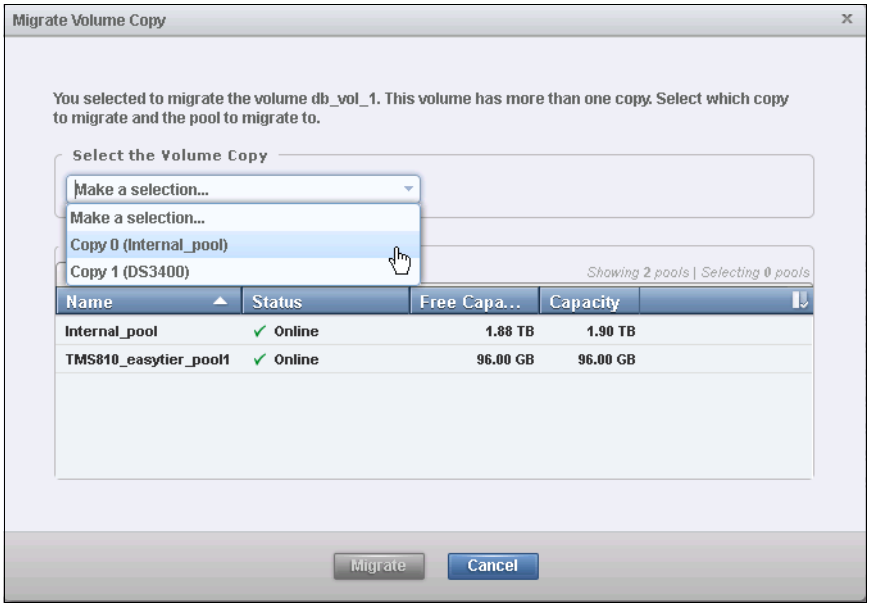

Figure 11-63 Selecting the destination storage pool of a mirrored volume

The selection you are presented with in Figure 11-63 denotes the current pool of each volume copy, so you can better determine which storage pool to use.

Finally, we explore the image mode import and image mode export. Both of these methods allow you to use all copy services functions on storage that contains pre-existing data.

To import pre-existing storage, click **Pools** → **MDisks by Pool** → **Unassigned MDisks**. Right-click the MDisk you want to import and then select **Import** (Figure 11-64).

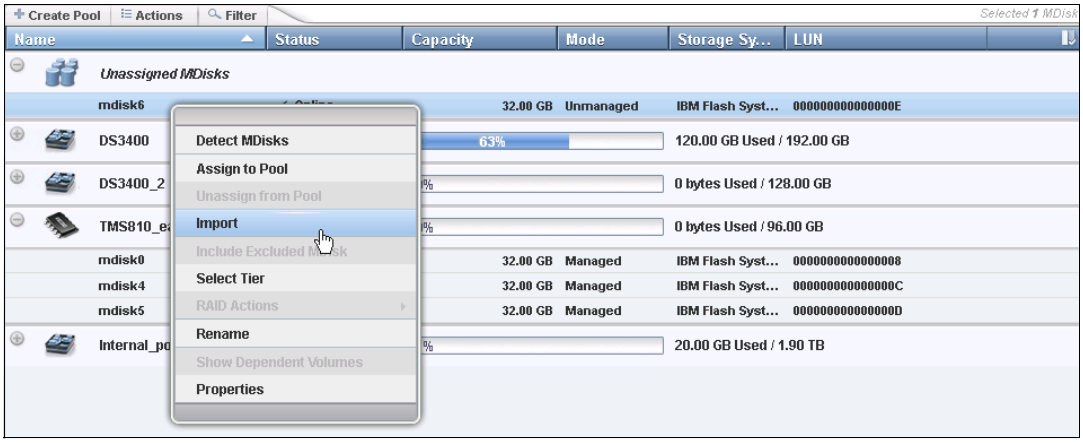

Figure 11-64 Import in image mode option

In the next window (Figure 11-65), enable or disable volume caching. Make your selection and click **Next**.

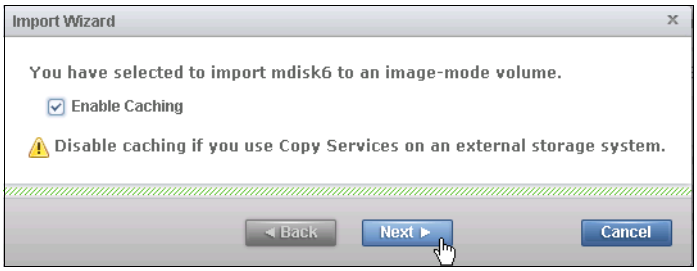

Figure 11-65 Enable or disable caching

In the next window, you can select to import the image-mode volume into existing or temporary storage pool (Figure 11-66). If you choose temporary pool, select the extent size. Click **Finish**.

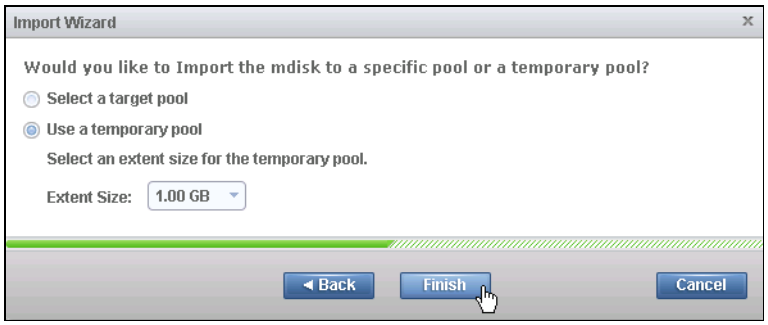

Figure 11-66 Select storage pool

The volume is imported into a new temporary pool as image-mode volume (Figure 11-67).

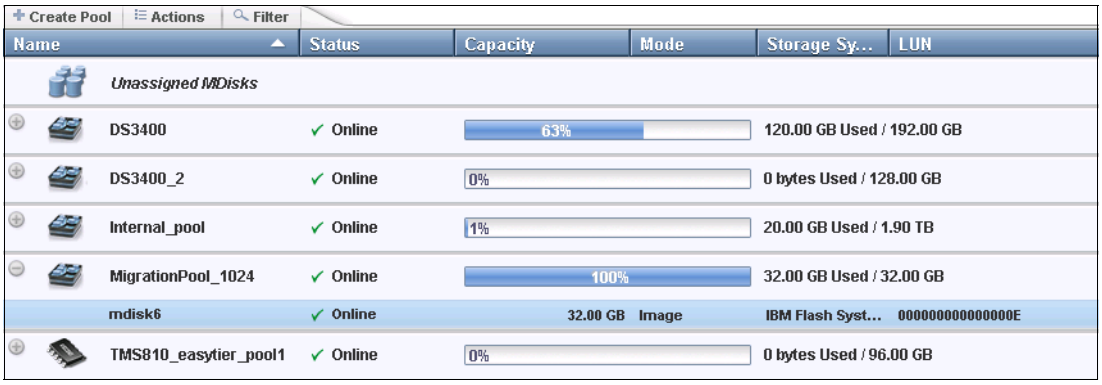

Figure 11-67 Temporary migration pool for image mode volume

To export a volume, it must not be in an image mode. Right-click on the volume, and select **Export to Image Mode** (Figure 11-68).

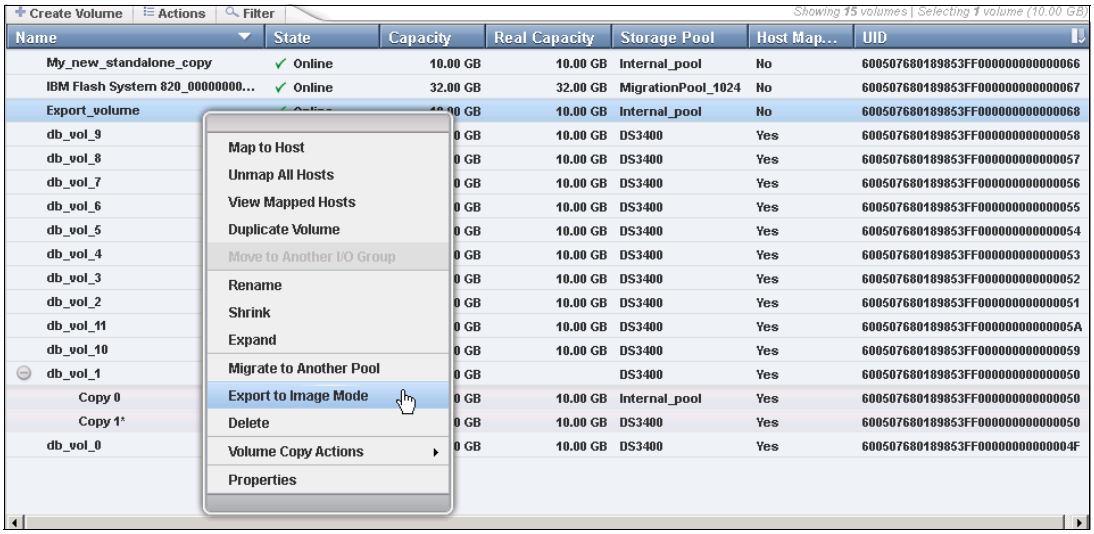

Figure 11-68 Export to image mode option

You can export only one volume or copy at a time, and you must select an MDisk for the volume when you export it (Figure 11-69).

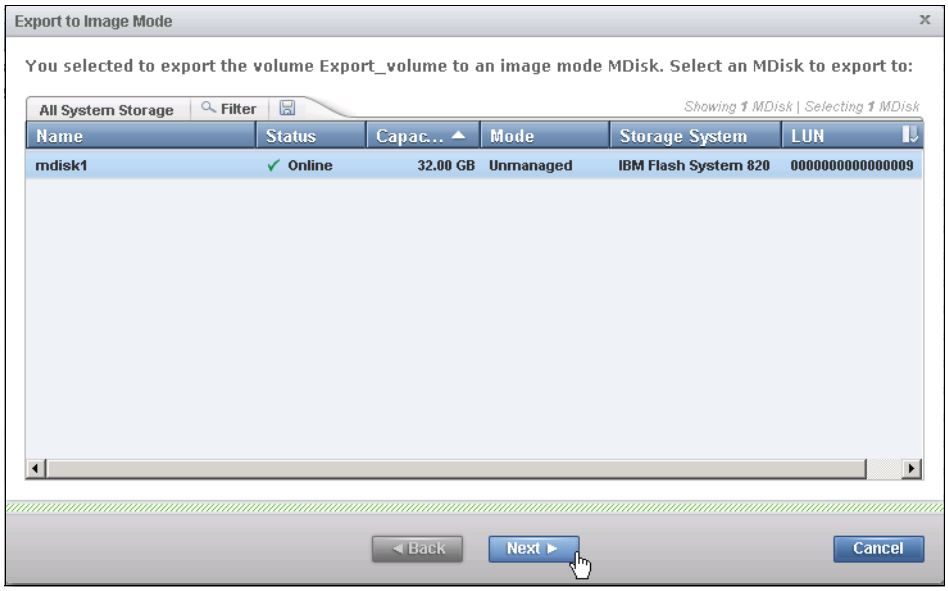

Figure 11-69 Select MDisk to export managed mode volume to image mode volume

When you click **Next**, on the window you need to select a storage pool into which the volume will be exported. Select the pool and click **Finish** (Figure 11-70).

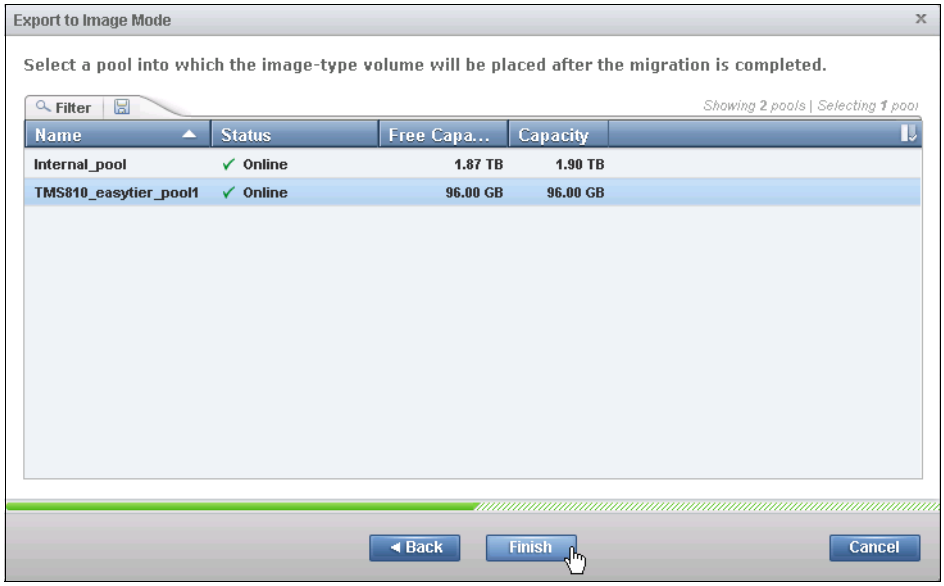

Figure 11-70 Select storage pool

The volume is then migrated to a selected MDisk and placed into target storage pool (Figure 11-71). You can view the export progress in the running task window (Figure 11-72.)

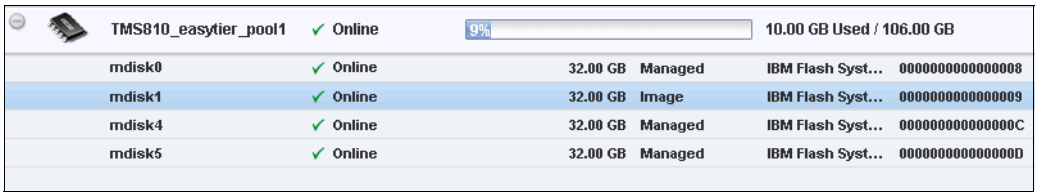

Figure 11-71 Storage pool with image mode MDisk

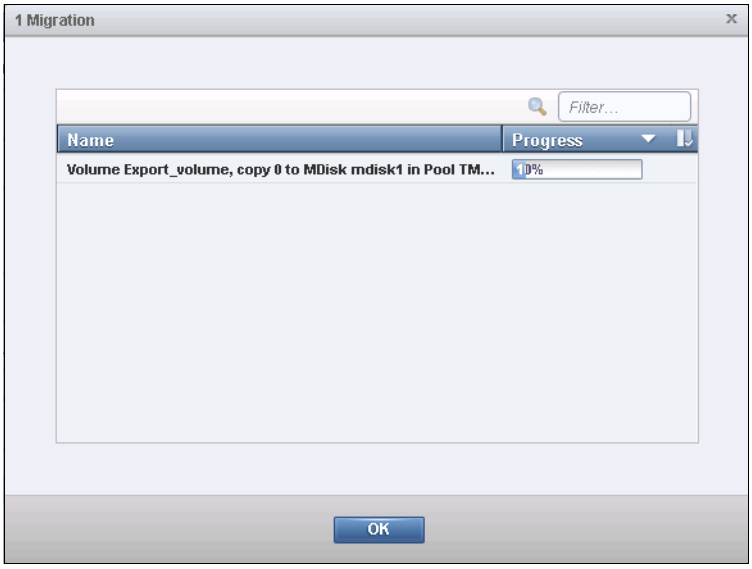

Figure 11-72 Volume migration progress

11.4 Remote copy

In this section, we describe how the Remote Copy function works in IBM Storwize V7000. We also provide the implementation steps, using the GUI, for Remote Copy configuration and management.

Remote copy consists of three methods for copying:

- Metro Mirror: This is for metropolitan distances with a synchronous copy requirement.
- Global Mirror: This is 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: This is an added piece of functionality for Global Mirror to attain consistency on higher latency network links.

Metro Mirror and Global Mirror are IBM branded terms for the functions Synchronous Remote Copy and Asynchronous Remote Copy. Throughout this book, the term *remote copy* is used to refer to both functions where the text applies to each term equally.

11.4.1 Remote copy concepts

Be sure you understand the remote copy concepts.

Partnership

When creating a partnership, we connect two separate IBM Storwize V7000 systems or an IBM SAN Volume Controller and an IBM Storwize V7000, which are separated by distance, but this distance is not required. After the partnership creation is configured on both systems, further communication between the node canisters in each of the storage systems is established and maintained by the SAN or IP network. All inter-cluster communication goes through the Fibre Channel or IP network. Partnership must be defined on both IBM Storwize V7000 systems or on the IBM Storwize V7000 and then IBM SAN Volume Controller systems, to make the partnership fully functional.

Interconnection: Interconnects between IBM Storwize V7000 and IBM SAN Volume Controller are introduced in Version 6.3. Both devices must be at that level, and the IBM Storwize V7000 must set the replication layer using the **`svctask chsystem -layer replication`** limitations described in the next section, “Introduction to layers”.

Introduction to layers

Version 6.3 introduces a new concept called *layers* for the IBM Storwize V7000. Layers determine how the V7000 interacts with the IBM SAN Volume Controller. Currently two layers are available:

- Replication: This layer is for when you want to use the IBM Storwize V7000 with one or more SAN Volume Controllers as a Remote Copy partner.
- Storage: This layer is the default mode of operation for the IBM Storwize V7000, and is for when you want to use the IBM Storwize V7000 to present storage to a SAN Volume Controller as an external storage system.

The layer for the IBM Storwize V7000 can be switched by running **`svctask chsystem -layer replication`**. Although this can be done while I/O is in progress, some restrictions exist regarding when the command can be used. See “Limitations on the SAN Volume Controller and Storwize V7000 partnership” on page 371.

Figure 11-73 shows the effect of layers on SAN Volume Controller and V7000 partnerships.

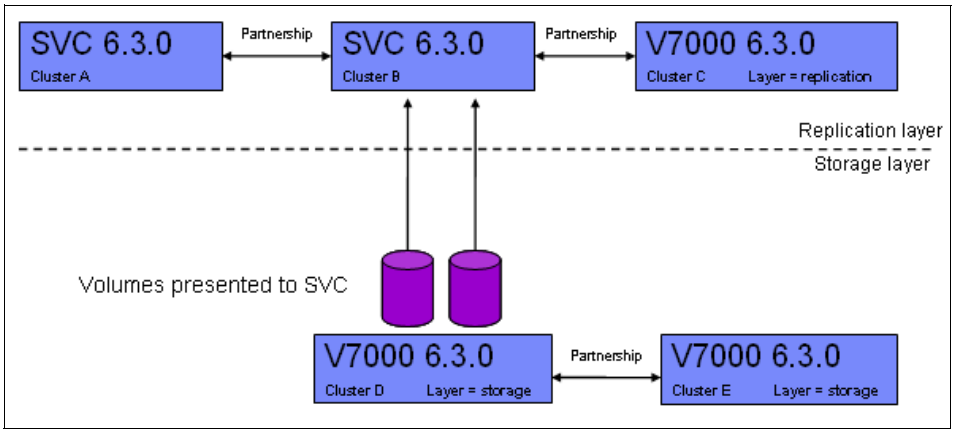

Figure 11-73 The relationship of layers and SAN Volume Controller and Storwize V7000 partnerships

The replication layer allows an IBM Storwize V7000 system to be a Remote Copy partner with a SAN Volume Controller; the storage layer allows an IBM Storwize V7000 system to function as back-end storage for a SAN Volume Controller. An IBM Storwize V7000 system cannot be in both layers at once.

Note: For creating a remote copy partnership between two or more Storwize V7000 systems, a system layer must be the same on all systems; it can be both, or either the storage or replication layer.

Limitations on the SAN Volume Controller and Storwize V7000 partnership

SAN Volume Controller and IBM Storwize V7000 systems can be partners in a Remote Copy partnership, although consider the following limitations:

- The layer for the V7000 must be set to replication. The default is storage.
- If any other SAN Volume Controller or IBM Storwize V7000 ports are visible on the SAN, aside from the ones on the system where you are making changes, you cannot change the layer.
- If any host object is defined to a SAN Volume Controller or IBM Storwize V7000 system, you cannot change the layer.
- If any MDisks from an IBM Storwize V7000 other than the one you are making the layer change on are visible, you cannot change the layer.
- If any system partnership is defined, you cannot change the layer.

Partnership topologies

A partnership between up to four IBM Storwize V7000 systems are allowed.

Typical partnership topologies between multiple IBM Storwize V7000s are as follows:

- Daisy-chain topology (Figure 11-74)

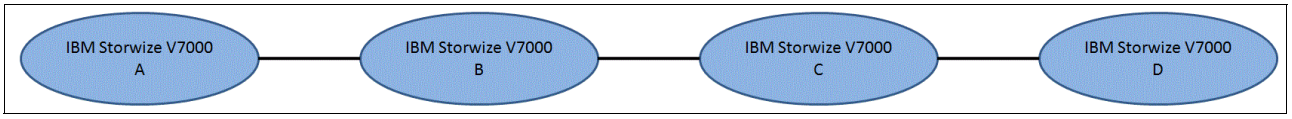

Figure 11-74 Daisy-chain topology for IBM Storwize V7000 Triangle topology

- Triangle topology (Figure 11-75)

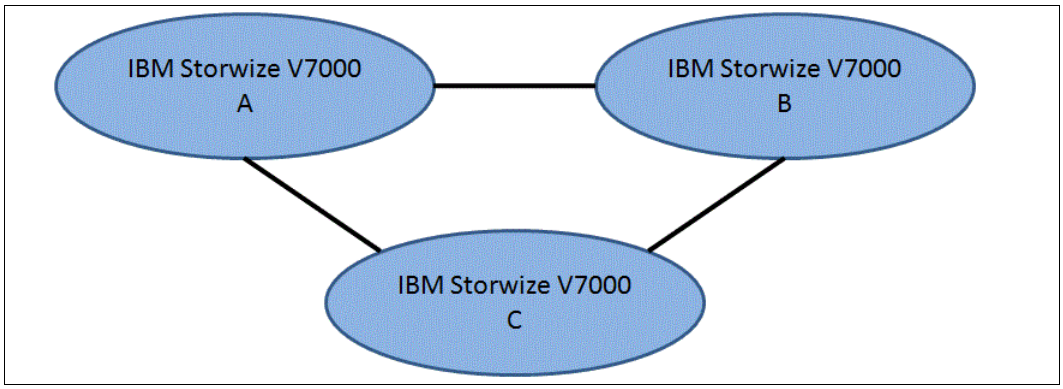

Figure 11-75 Triangle topology for IBM Storwize V7000

- Star topology (Figure 11-76)

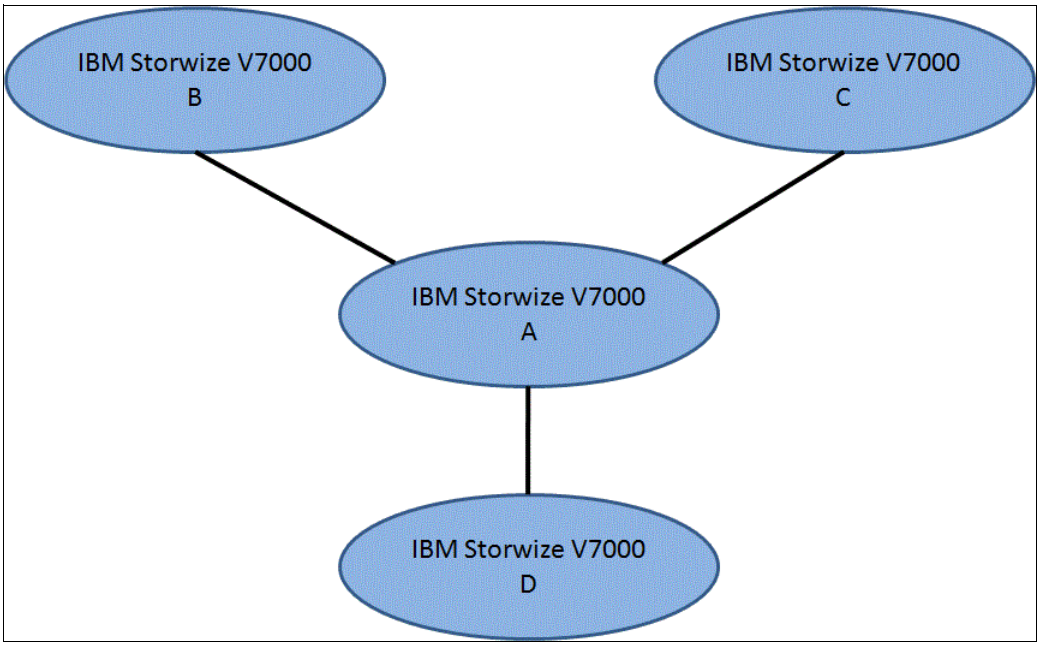

Figure 11-76 Star topology for IBM Storwize V7000

- Full-meshed topology (Figure 11-77)

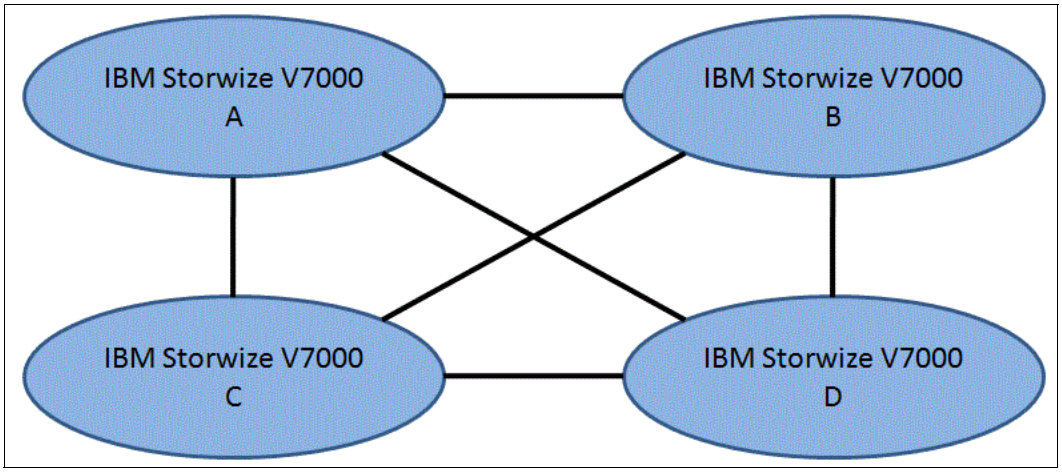

Figure 11-77 Full-meshed IBM Storwize V7000

Partnerships: These partnerships are valid for configurations with SAN Volume Controllers and IBM Storwize V7000 systems, while the IBM Storwize V7000 systems are using the replication layer.

Partnership states

A partnership has the following states:

- Partially Configured

  Indicates that only one system partner is defined from a local or remote system to the displayed system and is started. For the displayed system to be configured fully and to complete the partnership, you must define the system partnership from the system that is displayed to the corresponding local or remote system.

- Fully Configured

  Indicates that the partnership is defined on the local and remote systems and is started.

- Remote Not Present

  Indicates that the remote system is not present for the partnership.

- Partially Configured (Local Stopped)

  Indicates that the local system is only defined to the remote system and the local system is stopped.

- Fully Configured (Local Stopped)

  Indicates that a partnership is defined on both the local and remote systems and the remote system is present, but the local system is stopped.

- Fully Configured (Remote Stopped)

  Indicates that a partnership is defined on both the local and remote systems and the remote system is present, but the remote system is stopped.

- Fully Configured (Local Excluded)

  Indicates that a partnership is defined between a local and remote system; however, the local system has been excluded. Usually this state occurs when the fabric link between the two systems has been compromised by too many fabric errors or slow response times of the system partnership.

- Fully Configured (Remote Excluded)

  Indicates that a partnership is defined between a local and remote system; however, the remote system has been excluded. Usually this state occurs when the fabric link between the two systems has been compromised by too many fabric errors or slow response times of the system partnership.

- Fully Configured (Remote Exceeded)

  Indicates that a partnership is defined between a local and remote system and the remote is available; however, the remote system exceeds the number of allowed systems within a system network. The maximum of four systems can be defined in a network. If the number of systems exceeds that limit, the IBM Storwize V7000 system determines the inactive system or systems by sorting all the systems by their unique identifier in numerical order. The inactive system partner that is not in the top four of the system unique identifiers shows Fully Configured (Remote Exceeded).

11.4.2 Native IP replication technology

With the 7.2 code version, remote mirroring over IP communication is now supported on IBM SAN Volume Controller and Storwize Family systems using Ethernet communication links. IP replication uses innovative *Bridgeworks SANSlide* technology to optimize network bandwidth and utilization. This new function enables the use of lower speed and lower cost networking infrastructure for data replication. Bridgeworks SANSlide technology integrated into IBM SAN Volume Controller and Storwize Family Software uses artificial intelligence to help optimize

network bandwidth utilization and adapt to changing workload and network conditions. This technology can improve remote mirroring network bandwidth utilization up to three times, so that clients can deploy a less costly network infrastructure or speed remote replication cycles to enhance disaster recovery effectiveness.

Note: For creating remote partnerships between IBM SAN Volume Controller and Storwize Family systems consider these rules:

- The SAN Volume Controller is always in the *replication* layer.
- By default, Storwize V7000 is in *storage* layer.
- A system can only form partnerships with systems in the same layer.
- A SAN Volume Controller can virtualize a Storwize V7000 only if the Storwize V7000 is in the Storage layer.
- With version 6.4, a Storwize V7000 in the replication layer can virtualize a Storwize V7000 in the storage layer.

In a typical Ethernet network data flow, the data transfer slows over time. This is because of the latency caused by waiting for acknowledgement of each set of packets sent, because the next packet set cannot be sent until the previous packet is acknowledged (Figure 11-78).

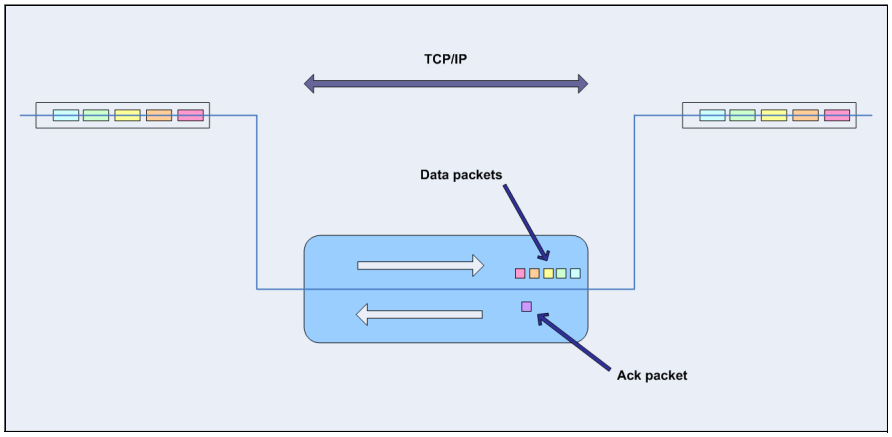

Figure 11-78 Typical Ethernet network data flow

Using the Bridgeworks SANSlide technology, this typical behavior can be dramatically eliminated with enhanced parallelism of the data flow by using multiple virtual connections (VC) that share the same IP links and addresses. The Artificial Intelligence engine can dynamically adjust the number of VCs, the receive window size, and the packet size as appropriate to maintain optimum performance. While waiting for one VC's ACK, it sends more packets across other VCs. If packets are lost from any VC, data is automatically retransmitted (Figure 11-79 on page 375).

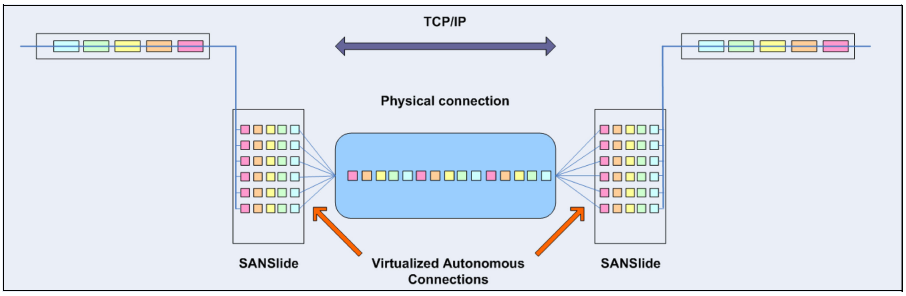

Figure 11-79 Optimized network data flow using Bridgeworks SANSlide technology

For more information about Bridgeworks SANSlide technology, see *IBM Storwize V7000 and SANSlide Implementation*, REDP-5023.

With native IP partnership, the following Copy Services features are supported.

Metro Mirror

Referred to as *synchronous replication*, Metro Mirror provides a consistent copy of a source virtual disk on a target virtual disk. Data is synchronously written to the target virtual disk after it is written to the source virtual disk, so that the copy is continuously updated.

Global Mirror and Global Mirror Change Volumes

Referred to as *asynchronous replication*, Global Mirror provides a consistent copy of a source virtual disk on a target virtual disk. Data is asynchronously written to the target virtual disk, so that the copy is continuously updated, but the copy might not contain the last few updates in the event that a disaster recovery operation is performed. An added extension to Global Mirror is Global Mirror with Change Volumes. Global Mirror with Change Volumes is the preferred method for use with native IP replication.

IP partnership limitations

Consider the following prerequisites and assumptions before an IP partnership between two SAN Volume Controller or Storwize V7000 systems can be established:

- SAN Volume Controller and Storwize systems are successfully installed with the latest IBM SAN Volume Controller 7.2 code levels.
- SAN Volume Controller and Storwize systems have the necessary licenses that allow remote copy partnerships to be configured between two systems. No separate license is required to enable IP partnership.
- The storage SAN configurations are properly done and the infrastructure to support SAN Volume Controller and Storwize systems in remote copy partnerships over IP links is properly in place.
- The two systems should be able to ping each other and do the discovery.
- Maximum partnerships between the local and remote systems, including both IP and Fibre Channel (FC) partnerships, is limited to the current maximum supported, which is three partnerships (four systems total).
- In 7.2.0 only a single partnership over IP is supported.
- A system can have simultaneous partnerships over FC and IP but with separate systems. FC zones between two systems must be removed before configuring an IP partnership.
- IP partnerships are supported on both 10 Gbps links and 1 Gbps links. However, the intermix of both on a single link is not currently supported.

- The maximum supported round-trip time is 80 milliseconds (ms) for 1 Gbps links.
- The maximum supported round-trip time is 10 ms for 10 Gbps links.
- The minimum supported link bandwidth is 10 Mbps.
- The inter-cluster heartbeat traffic consumes 1 Mbps per link
- Only nodes from two I/O groups can have ports that are configured for an IP partnership.
- Migrations of remote copy relationships directly from Fibre Channel based partnerships to IP partnerships are not supported.
- IP partnerships between the two systems can be either over IPv4 or IPv6 only and not both.
- VLAN tagging of the IP addresses configured for remote copy is not supported.
- SAN Volume Controller and Storwize systems allow two system IPs (management IPs) to be configured. If planning to use the second system IP for IP partnerships, both the system IPs should be in the same subnet.
- An added layer of security is provided by means of Challenge Handshake Authentication Protocol (CHAP) authentication
- TCP ports 3260 and 3265 are used for IP partnership communications and therefore these ports will need to be opened in firewalls between the systems.
- Maximum throughput is currently restricted based on use of 1 Gbps or 10 Gbps ports:
  - One 1 Gbps port can transfer up to 110 MBps
  - Two 1 Gbps ports can transfer up to 220 MBps
  - One 10 Gbps port can transfer up to 190 MBps
  - One 10 Gbps ports can transfer up to 280 MBps

Bandwidth: The Bandwidth setting definition, when creating the IP partnerships, has changed: Previously, the bandwidth setting defaults to 50 MB and it was the maximum transfer rate from primary to secondary site for initial sync/resync of volumes.

The *Link bandwidth* setting is now configured using Mbits not MBytes and you set this to a value that the communication link can actually sustain or what is actually allocated for replication. The *Background copy rate* setting is now a percentage of the Link bandwidth and it determines the bandwidth available for initial synchronizations and resynchronizations or for Global Mirror with Change Volumes.

IP partnership terminology

The IP partnership terminology and abbreviations are listed in Table 11-4.

Table 11-4 IP partnership terminology

Term	Definition
Remote copy group or Remote copy port group	This is a number that groups a set of IP addresses that are connected to the same physical link. Therefore, only IP addresses that are part of the same remote copy group can form remote copy connections with the partner system. 0 – Ports that are not configured for remote copy 1 – Ports that belong to remote copy port group 1 2 – Ports that belong to remote copy port group 2 **Note**: Each IP address can be shared for iSCSI host attach and remote copy function. Therefore, appropriate settings must be applied to each IP address.
IP partnership	Two SAN Volume Controller or Storwize systems are partnered to perform remote copy over native IP links.
FC partnership	Two SAN Volume Controller or Storwize systems are partnered to perform remote copy over native Fibre Channel links.
Failover	Failure of a node within an I/O group fails and causes virtual disk access through the surviving node. The IP addresses failover to the surviving node in the I/O group. When the configuration node of the system fails, management IPs also fail over to an alternate node.
Failback	When the failed node rejoins the system, all failed over IP addresses are failed back from the surviving node to the rejoined node and virtual disk access is restored through this node.
linkbandwidthmbits	Aggregate bandwidth of all physical links between two sites in Mbps.
IP partnership or partnership over native IP links	These terms describe the IP partnership feature.

States of IP partnership

Table 11-5 describes the partnership states in an IP partnership.

Table 11-5 States of IP partnership

State	Systems connected	Support for active remote copy I/O	Description
Partially_Configured_Local	No	No	This state indicates that the initial discovery is completed.
Fully_Configured	Yes	Yes	Discovery has successfully completed between two systems and the two systems can establish remote copy relationships.
Fully_Configured_Stopped	Yes	Yes	The partnership is stopped on the system.
Fully_Configured_Remote_Stopped	Yes	No	The partnership is stopped on the remote system.
Not_Present	Yes	No	The two systems cannot communicate with each other. This state is also seen when data paths between the two systems have not been established.

State	Systems connected	Support for active remote copy I/O	Description
Fully_Configured_Exceeded	Yes	No	Too many systems are in the network and the partnership from the local system to the remote system has been disabled.
Fully_Configured_Excluded	No	No	The connection is excluded because of too many problems, or either system is unable to support the I/O workload for the Metro Mirror and Global Mirror relationships.

To establish two systems in IP partnerships, complete these steps, as the administrator:

1. Optional: Configure the CHAP secret on both systems. This is not a mandatory step and users can choose to not configure the CHAP secret.
2. If required, configure the system IP addresses on both the local and remote system, so that they can discover each other over the network.
3. Configure the SAN Volume Controller ports on each node in both the systems using the **`svctask cfgportip`** command:
   a. Configure the IP addresses for remote copy data.
   b. Add the IP addresses in the respective remote copy port group.
   c. Define whether the host access on these ports over iSCSI will be allowed.
4. Establish partnership with the remote system, from the local system, where the partnership state then transitions to the `Partially_Configured_Local` state.
5. Establish partnership from the remote system with the local system and if successful the partnership state then transitions to the Fully_Configured state, thus implying that the partnerships over IP network are successfully established. The partnership state momentarily remains in the Not_Present state before transitioning to `fully_configured`.
6. Create Metro Mirror, Global Mirror, and Global Mirror with Change Volume relationships.

Partnership consideration: When creating the partnership, there is no master/auxiliary status defined or implied. The partnership is equal and the concepts of master/auxiliary and primary/secondary apply only to volume relationships, not to system partnerships.

Remote Copy groups

This section discusses remote copy groups (or remote copy port groups) and the different ways in which the links between the two remote systems can be configured. The two SAN Volume Controller or Storwize systems can be connected to each other over one or at most two links. To address the requirement to let the system know about the physical links between two sites, the concept of remote copy port groups is introduced.

SAN Volume Controller or Storwize IP addresses that are connected to the same physical link are designated with identical remote copy port groups. The SAN Volume Controller or Storwize system supports three remote copy groups: 0, 1, and 2. The SAN Volume Controller or Storwize system IP addresses are, by default, in remote copy port group 0. Ports in port group 0 are not considered for creating Remote Copy data paths between two systems. For partnerships to be established over IP links directly, IP ports must be configured in either remote copy group 1 if there is a single inter-site link or both 1 and 2 if there are two inter-site links.

You can assign an IPv4 address and an IPv6 address to each Ethernet port on the SAN Volume Controller or Storwize platforms. Each of these IP addresses can be shared between iSCSI host attach and IP partnership. The user must configure the required IP address (IPv4 or IPv6) on an Ethernet port with a remote copy port group. The administrator might want to use IPv6 addresses for remote copy operations and use IPv4 addresses on that same port for iSCSI host attach. This also implies that for two systems to establish IP partnership both must have IPv6 addresses configured.

Administrators can choose to dedicate an Ethernet port for IP partnership only. In that case, host access must be explicitly disabled for that IP address and any other IP address that is configured on that Ethernet port.

Note: To establish an IP partnership, each SAN Volume Controller node or Storwize node canister must have only a single remote copy port group configured (either 1 or 2). The remaining IP addresses must be in remote copy port group 0.

11.4.3 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 typically contains a backup copy of the data and is used for disaster recovery.

The master and auxiliary volumes are defined when the relationship is created, and these attributes never change. However, either volume can operate in the primary or secondary role as necessary. The primary volume contains a valid copy of the application data and receives updates from the host application, which is analogous to a source volume. The secondary volume receives a copy of any updates to the primary volume, because these updates are all transmitted across the mirror link. Therefore, the secondary volume is analogous to a continuously updated target volume. When a relationship is created, the master volume is assigned the role of primary volume and the auxiliary volume is assigned the role of secondary volume. The initial copying direction is from master to auxiliary. When the relationship is in a consistent state, you can reverse the copy direction.

The two volumes in a relationship must be the same size. The Remote Copy relationship can be established on the volumes within one IBM Storwize V7000 storage system, which is called an *intra*-cluster relationship. The relationship can also be established in different IBM Storwize V7000 storage systems or between an IBM Storwize V7000 storage system and an IBM SAN Volume Controller, which are called *inter*-cluster relationships.

Usage of Remote Copy target volumes as Remote Copy source volumes is not allowed. A FlashCopy target volume can be used as Remote Copy source volume and also as a Remote Copy target volume.

FlashCopy: Using target volumes as remote copy sources and target volumes for FlashCopy requires Version 6.2.x or later.

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 has 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 maintains a fully synchronized copy of the source data at the target site at all times.

Figure 11-80 shows a write operation to the master volume is mirrored to the cache of the auxiliary volume before an acknowledgement of the write is sent back to the host that issued the write. This process ensures that the auxiliary is synchronized in real time, if it is needed in a failover situation.

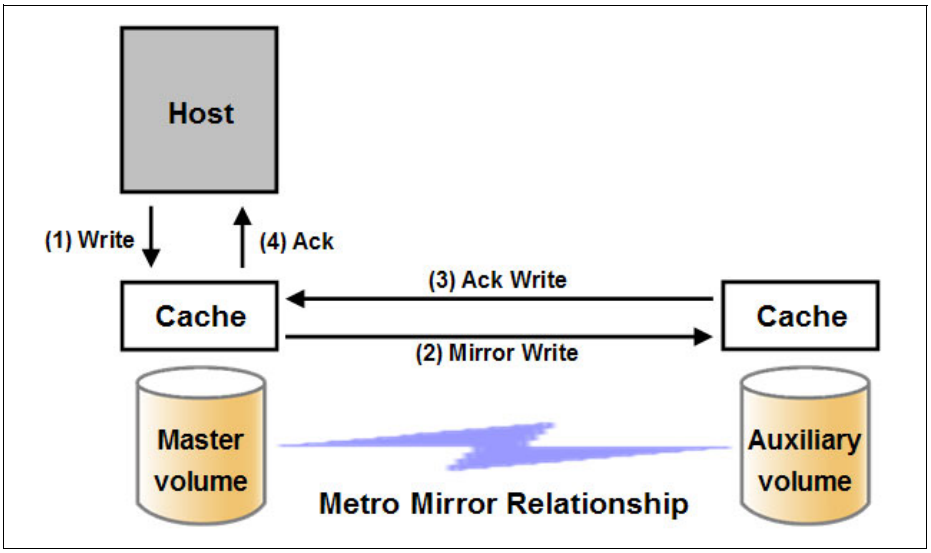

Figure 11-80 Write on volume in a Metro Mirror relationship

The Metro Mirror function supports copy operations between volumes that are separated by distances up to 300 km. For disaster recovery purposes, Metro Mirror provides the simplest way to maintain an identical copy on both the primary and secondary volumes. However, as with all synchronous copies over remote distances, there can be a performance impact to host applications. This performance impact is related to the distance between primary and secondary volumes and, depending on application requirements, its use might be limited based on the distance between sites.

Global Mirror

Global Mirror provides an asynchronous copy, which means that the secondary volume is not an exact match of the primary volume at every point in time. 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 acknowledgement is sent to the host before it is received at the secondary site. An update of this write operation is sent to the secondary site at a later stage, which provides the capability to perform Remote Copy over distances exceeding the limitations of synchronous Remote Copy.

The distance of Global Mirror replication is limited primarily by the latency of the WAN Link provided. Global Mirror has a requirement of 80 ms round-trip-time for data sent to the remote location. The propagation delay is roughly 8.2 µs per mile or 5 µs per kilometer for Fibre Channel connections. Each device in the path adds addition delay of about 25 µs. Devices that use software (such as some compression devices) add much more time. The time added

by software-assisted devices is highly variable and should be measured directly. Be sure to include these times when 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 V7000 storage system provides you with an advanced feature of Global Mirror that permits you to test performance implications before deploying Global Mirror and obtaining a long-distance link. This advanced feature is enabled by modifying the IBM Storwize V7000 storage system parameters **`gmintradelaysimulation`** and **`gminterdelaysimulation`**. These two 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. You can find more information about how to enable the CLI feature in Appendix A, “CLI setup and SAN Boot” on page 497.

Figure 11-81 shows that a write operation to the master volume is acknowledged back to the host that issued the write before the write operation is mirrored to the cache for the auxiliary volume.

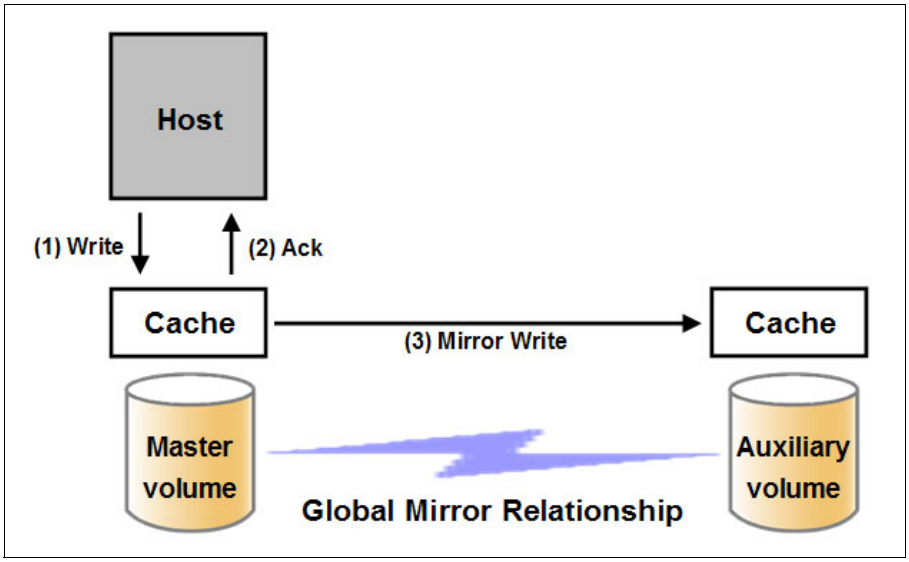

Figure 11-81 Global Mirror write sequence

The Global Mirror algorithms maintain a consistent image on the auxiliary volume at all times. 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.

The multiple I/Os within a single set are applied concurrently. The process that marshals the sequential sets of I/Os operates at the secondary system, and it is therefore not subject to the latency of the long-distance link. These two elements of the protocol ensure that the throughput of the total system can be grown by increasing system size, while maintaining consistency across a growing data set.

In the V7.2 code, these algorithms have been enhanced to optimize Global Mirror behavior and latency even further. As stated before, Global Mirror write I/O from production to a secondary SAN Volume Controller system requires serialization and sequence-tagging before being sent across network to remote site (to maintain a write-order consistent copy of data). Sequence-tagged Global Mirror writes on the secondary system are processed without parallelism and management of write I/O sequencing imposes additional latency on write I/Os

in code version before V7.2. As a result, high-bandwidth Global Mirror throughput environments can experience performance impacts on primary system during high I/O peak periods. The SAN Volume Controller code V7.2 now allows more parallelism in processing and managing Global Mirror writes on secondary system using these methods:

- Nodes on the secondary system store replication writes in new redundant non-volatile cache.
- Cache content details are shared between nodes.
- Cache content details are batched to make node-to-node latency less of an issue.
- Nodes intelligently apply these batches in parallel as soon as possible.
- Nodes internally manage and optimize Global Mirror secondary write I/O processing.

Note: Although the V7.2 enhancements of Global Mirror require no changes in administration and management, before you upgrade to code V7.2, stop all Global Mirror relationships. The proper checks are provided in the latest **`svcupgradetest`** utility.

In a failover scenario, where the secondary site needs to 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.

Global Mirror with Change Volumes

Global Mirror within the IBM Storwize V7000 can achieve a recovery point objective (RPO) as low as possible, so that data is as up-to-date as possible. This capability places strict requirements on your infrastructure and in certain situations; with low network link quality or congested or overloaded hosts, you might be impacted by multiple 1920 (congestion) errors.

Congestion errors happen in three primary situations:

- At the source site through the host or network
- In the network link or network path
- At the target site through the host or network

Global Mirror functionality can address several conditions that negatively affect some Global Mirror implementations:

- Estimation of bandwidth requirements tends to be complex.
- Meeting the latency and bandwidth requirements is often difficult to guarantee.
- Congested hosts on either the source or target site can cause disruption.
- Congested network links can cause disruption with only intermittent peaks.

To address these issues, *Change Volumes* have been 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, instead of all the updates during the period. This situation can provide significant reductions in replication volume.

Figure 11-82 shows a basic Global Mirror relationship without Change Volumes.

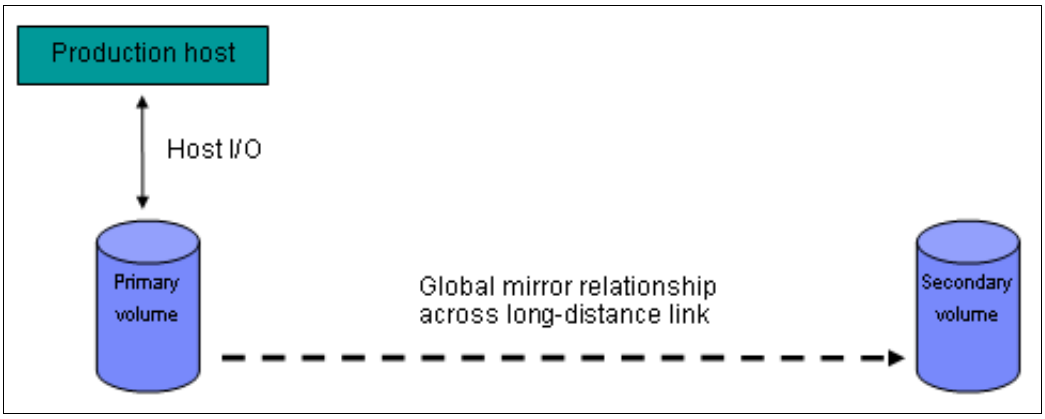

Figure 11-82 Global Mirror without Change Volumes

Figure 11-83 shows a relationship with the Change Volumes.

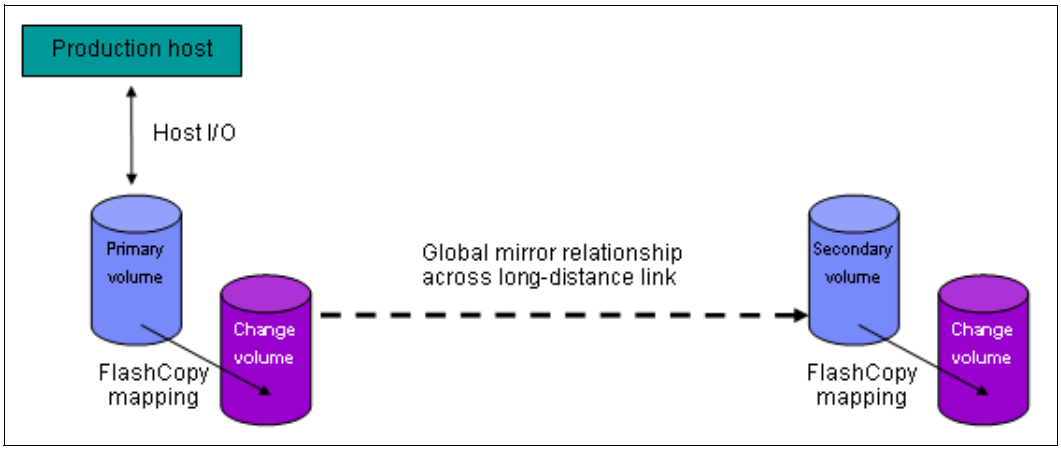

Figure 11-83 Global Mirror with Change Volumes

With Change Volumes, a FlashCopy mapping exists between the primary volume and the primary Change Volume. The mapping is updated during a cycling period (every 60 seconds to one day.) The primary Change Volume is then replicated to the secondary Global Mirror volume at the target site, which is then captured in another change volume on the target site. This situation provides an always consistent image at the target site and protects your data from being inconsistent during resynchronization.

Look more closely at how Change Volumes might reduce replication traffic.

Figure 11-84 shows several I/Os on the source volume and the same number on the target volume, and in the same order. Assuming that this set is the same set of data being updated repeatedly, then these updates are wasted network traffic and the I/O can be completed much more efficiently (Figure 11-85).

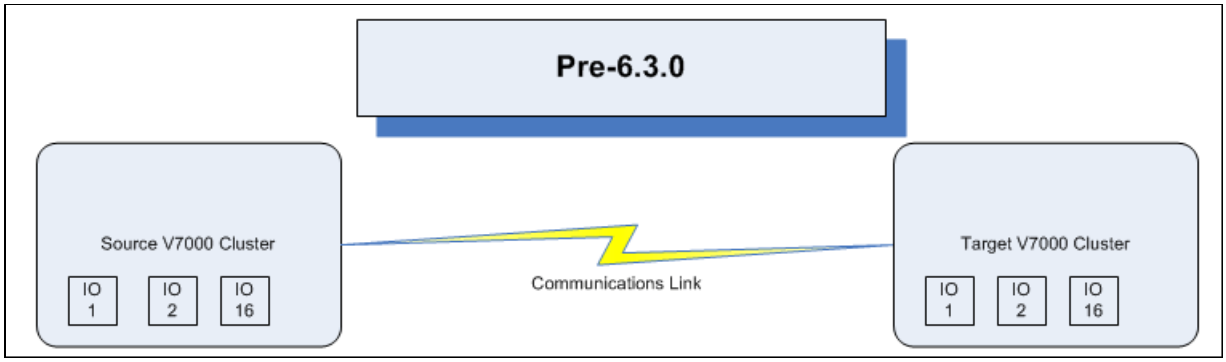

Figure 11-84 Global Mirror I/O replication without Change Volumes

In Figure 11-85, the same data is being updated repeatedly, so Change Volumes demonstrate significant I/O transmission savings, because you only need to send I/O number 16, which was the last I/O before the cycling period.

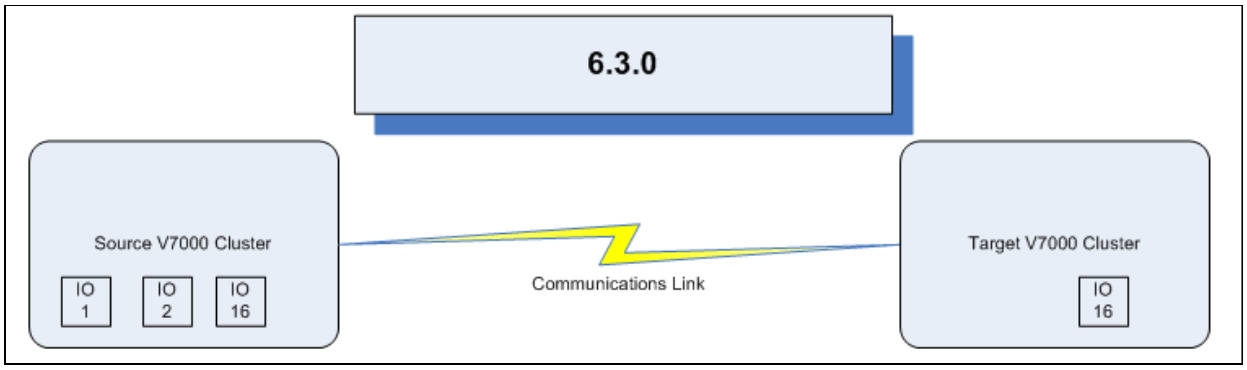

Figure 11-85 Global Mirror I/O replication with Change Volumes

The cycling period can be adjusted by running **`chrcrelationship -cycleperiodseconds <60-86400>`**. If a copy does not complete in the cycle period, the next cycle does not start until the prior one has completed. For this reason, using Change Volumes gives you two 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 one is finished.

Carefully consider the balancing of 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 using Global Mirror volumes with Change Volumes, 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.

Remote Copy consistency groups

A consistency group is a logical entity that groups copy relationships. By grouping the relationships, you can ensure that these relationships are managed in unison and the data within the group is in a consistent state. More information about the necessity of consistency groups is in “FlashCopy consistency groups” on page 334.

Remote Copy commands can be issued to a Remote Copy consistency group, and therefore simultaneously for all Metro Mirror relationships defined within that consistency group, or to a single Metro Mirror relationship that is not part of a Metro Mirror consistency group.

Figure 11-86 shows the concept of Remote Copy consistency groups. Because the RC_Relationships 1 and 2 are part of the consistency group, they can be handled as one entity, while the stand-alone RC_Relationship 3 is handled separately.

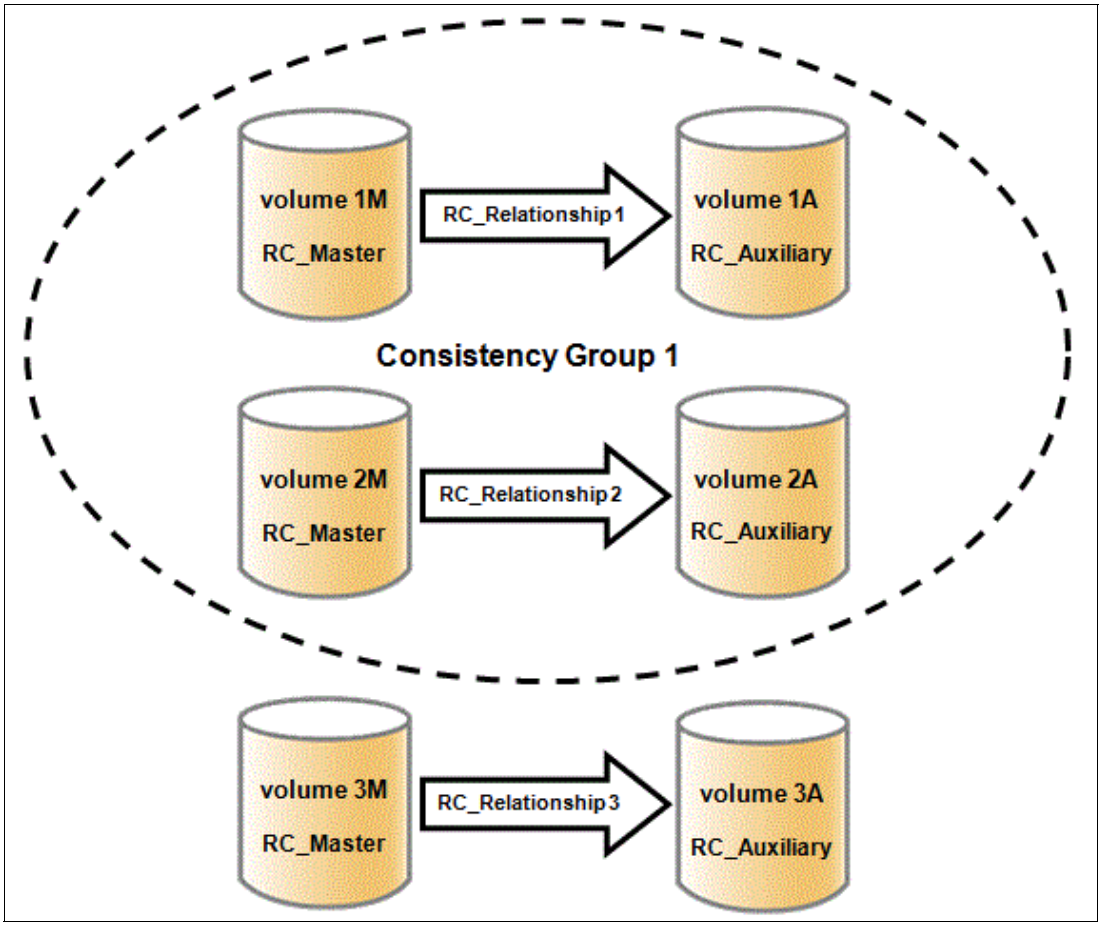

Figure 11-86 Remote Copy consistency group

Remote Copy relationships can belong to only one consistency group, but they do not have to belong to a consistency group. Relationships that are not part of a consistency group are called *stand-alone relationships*. A consistency group can contain zero or more relationships. All relationships in a consistency group must have matching primary and secondary systems, which are sometimes referred to as *master* and *auxiliary systems*. 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.

These states apply to both 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 out-of-date with respect to the primary volumes. This state can occur when a relationship was in the ConsistentSynchronized state and experiences an error that forces a freeze of the consistency group or the Remote Copy relationship.

- ConsistentSynchronized

  The primary volumes are accessible for read and write I/O operations. The secondary volumes are accessible for read-only I/O operations.

- Idling

  Both the primary volumes and the secondary volumes are operating in the primary role. Consequently, 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.

11.4.4 Remote Copy planning

Before you use Remote Copy, you need to plan for its usage.

General guidelines for Remote Copy

Consider the following general guidelines for Remote Copy:

- Partnerships between up to four IBM Storwize V7000 storage systems or IBM SAN Volume Controller systems is allowed. The partnership must be defined on any partnered IBM Storwize V7000 storage systems or IBM SAN Volume Controller systems to make it fully functional.
- The two volumes in a relationship must be the same size.
- The Remote Copy relationship can be established on the volumes within one IBM Storwize V7000 storage system or in different IBM Storwize V7000 storage systems. When the two volumes are in the same system, they must be in the same I/O group.
- You can use Remote Copy target volumes as Remote Copy source volumes. A FlashCopy target volume can be used as Remote Copy source volume, as of Version 6.2.0. There are additional restrictions outlined in Table 11-7 on page 389.
- The Metro Mirror function supports copy operations between volumes that are separated by distances up to 300 km.
- One Remote Copy relationship can only belong to one consistency group.
- All relationships in a consistency group must have matching primary and secondary systems, which are sometimes referred to as master and auxiliary systems. All relationships in a consistency group must also have the same copy direction and state.
- Metro Mirror and Global Mirror relationships cannot belong to the same consistency group.
- To manage multiple Remote Copy relationships as one entity, relationships can be made part of a Remote Copy consistency group, which ensures data consistency across multiple Remote Copy relationships and provides ease of management.
- An IBM Storwize V7000 storage system implements flexible resynchronization support, enabling it to resynchronize volume pairs that have experienced write I/Os to both disks 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 11-6 lists the Remote Copy configuration limits.

Table 11-6 Remote Copy configuration limits

Parameter	Value
Number of Remote Copy consistency groups per system	256
Number of Remote Copy relationships per consistency group	8,192
Number of Remote Copy relationships per I/O Group	2,048
Total Remote Copy volume capacity per I/O Group	1024 TB (This limit is the total capacity for all master and auxiliary volumes in the I/O group.)

SAN planning for Remote Copy

Here are guidelines for planning for a SAN for Remote Copy.

Zoning recommendation

Node canister ports on each IBM Storwize V7000 must be able to communicate with each other so that the partnership creation can be performed. These ports must be visible to each other on your SAN. Proper switch zoning is critical to facilitating inter-cluster communication.

The SAN zoning recommendations are as follows:

- For each node canister, exactly two Fibre Channel ports should be zoned to exactly two Fibre Channel ports from each node canister in the partner system.
- If dual-redundant inter-switch links (ISLs) are available, then the two ports from each node should be split evenly between the two ISLs, that is, exactly one port from each node canister should be zoned across each ISL. More details are available at the following address:

  `http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003634&myns=s033&mynp=familyind5329743&mync=E`
- In addition, all local zoning rules should be followed. A properly configured SAN fabric is key to local SAN performance and Remote Copy. You can learn more about these rules at the following address:

  `http://pic.dhe.ibm.com/infocenter/storwize/ic/index.jsp?topic=%2Fcom.ibm.storwize.v7000.doc%2Fsvc_configrulessummary_02171530.html`

Fabrics: When a local fabric and a remote fabric are connected together for Remote Copy purposes, the ISL hop count between a local node and a remote node cannot exceed seven.

Remote Copy link requirements

These 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 putting any network links in place to see if the applications perform at an acceptable level while meeting the round-trip latency requirement.
- Bandwidth

  The bandwidth must satisfy the following requirements:

  – If you are not using Change Volumes: Be able to sustain peak write load for all mirrored volumes and background copy traffic.

  – If you are using Change Volumes with Global Mirror: Be able to sustain change rate of Source Change Volumes and background copy traffic.

  – Additional background copy rate (the preferred practice is 10 - 20% of maximum peak load) for initial synchronization and resynchronization.

  – Remote Copy internal communication at idle with or without Change Volumes is approximately 2.6 Mbps. This amount is the minimum amount.

Redundancy: If the link between two sites is configured with redundancy so that it can tolerate single failures, the link must be sized so that the bandwidth and latency requirement can be met during single failure conditions.

Interaction between Remote Copy and FlashCopy

Table 11-7 lists supported combinations of FlashCopy and Remote Copy.

Table 11-7 FlashCopy and Remote Copy interaction

Component	Remote Copy primary site	Remote Copy secondary site
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 destination	This combination is supported. It has several restrictions: ► Issuing **`stop -force`** might cause the Remote Copy relationship to fully resynchronize. ► Code level must be V6.2.x or later. ► I/O group must be the same.	This combination is supported with 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, you can, for disaster recovery purposes, 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.

11.5 Troubleshooting Remote Copy

Remote Copy (Global Mirror and Metro Mirror) has two primary error codes:

- 1920 error: This is a congestion error and means that either the source, the link between source and target, or the target were not able to keep up with the rate of demand.
- 1720 error: This is a heartbeat or system partnership communication error. It tends to be more serious because failing communication between your system partners involves some extended diagnostic time.

11.5.1 1920 error

A 1920 error (event ID 050010) can have several triggers. Official probable cause projections are as follows:

- Primary 2145 system or SAN fabric problem (10%)
- Primary 2145 system or SAN fabric configuration (10%)
- Secondary 2145 system or SAN fabric problem (15%)
- Secondary 2145 system 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 SAN Volume Controller system to your target SAN Volume Controller System and the accompanying acknowledgement must have a total time of 80 ms or 40 ms each way (for Version 4.1.1.x and later).

Round-trip time (RTT): For Version 4.1.0.x and earlier, this limit was 68 ms or 34 ms one way for Fibre Channel extenders, and for SAN routers it was 10 ms one way or 20 ms round trip. Make sure to use the correct values for the correct versions!

The primary component of your round-trip time is the physical distance between sites. For every 1000 km (621.36 miles), there is a 5 ms delay. This delay does not include the time added by equipment in the path. Every device adds a varying amount of time, depending on the device, but expect about 25 µs for pure hardware devices. For software-based functions (such as compression implemented in software), the delay added tends to be much higher (usually in the millisecond plus range).

Consider an example. Company A has a production site that is 1900 km distant from their recovery site. Their network service provider uses a total of five devices to connect the two sites. In addition to those devices, Company A employs a SAN Fibre Channel Router at each site to provide Fibre Channel over IP (FCIP) to encapsulate the Fibre Channel traffic between sites. There are now seven devices, and 1900 km of distance delay. All the devices add 200 µs of delay *each way.* The distance adds 9.5 ms each way, for a total of 19 ms. Combined with the device latency, that is 19.4 ms of *physical* latency at a minimum. This latency is under the 80 ms limit of Global Mirror, but this number is the best case number. Link quality and bandwidth play a significant role here. Your network provider likely guarantees a latency maximum on your network link; be sure to stay below the Global Mirror 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. Keep in mind that the FCIP routing layer can introduce latency if it is not properly configured. If your network provider reports a much lower latency, this report might be an indication of a problem at your FCIP Routing layer. Most FCIP Routing devices have built-in tools so you can check the RTT. When checking latency, remember that TCP/IP routing devices (including FCIP routers) report RTT using standard 64-byte ping packets.

Figure 11-87 shows why the effective transit time should be measured by using only packets large enough to hold a Fibre Channel frame. This packet size is 2148 bytes (2112 bytes of payload and 36 bytes of header) and you should allow some additional capacity to be safe, because different switching vendors have optional features that might increase this size. After you verify the latency using the correct packet size, continue with normal hardware troubleshooting.

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 11-87 The effect of packet size (in bytes) versus the link size

Before you proceed, look at the second largest component of your round-trip-time: serialization delay. *Serialization delay* is the amount of time required to move a packet of data of a specific size across a network link of a given bandwidth. This delay is based on a simple concept: the time required to move a specific amount of data decreases as the data transmission rate increases.

Figure 11-87 shows orders of magnitude of difference between the various link bandwidths. You can easily see how 1920 errors can arise when bandwidth is insufficient and why you should never use a TCP/IP **`ping`** to measure RTT for FCIP traffic.

Figure 11-87 compares the amount of time in microseconds required to transmit a packet across network links of varying bandwidth capacity. Three 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

Remember, your path MTU affects the delay that is incurred in getting a packet from one location to another, when it causes fragmentation or is too large and causes too many retransmits when a packet is lost.

11.5.2 1720 error

A 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 be able to communicate with each other. When a partner on either side stops communicating, you see a 1720 error 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 of the network path between your partners. When you receive this error and if your fabric has more than 64 HBA ports that are zoned, check your fabric configuration for zoning of more than one HBA port for each node per I/O group. One port for each node per I/O group, associated with the host, is the recommended zoning configuration for fabrics. For those fabrics with 64 or more host ports, this recommendation becomes a rule. You must follow this zoning rule or the configuration is technically unsupported.

Improper zoning leads to SAN congestion, which can inhibit remote link communication intermittently. Checking the zero buffer credit timer through IBM Tivoli Storage Productivity Center and comparing its value against your sample interval might reveal potential SAN congestion. Anytime 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 point to a larger problem.

If you receive multiple 1720 errors, recheck your network connection and then check the IBM Storwize V7000 partnership information to verify their status and settings. Perform diagnostic tests for every piece of equipment in the path between the two systems. Having a diagram that shows the path of your replication from both logical and physical configuration viewpoints is helpful.

If your investigation fails to resolve your Remote Copy problems, contact your IBM support representative for a complete analysis.

11.6 Managing Remote Copy by using the GUI

The IBM Storwize V7000 storage system provides a separate menu icon for copy service management. Two panels are available for managing Remote Copy, which are accessed through the Copy Services menu icon; these two panels are used to manage Remote Copy and the partnership:

- Remote Copy
- Partnerships

11.6.1 Managing system partnerships

The Partnership panel is used to manage a partnership between systems. To access the Partnership panel, click the **Copy Services** icon and click **Partnerships** (Figure 11-88).

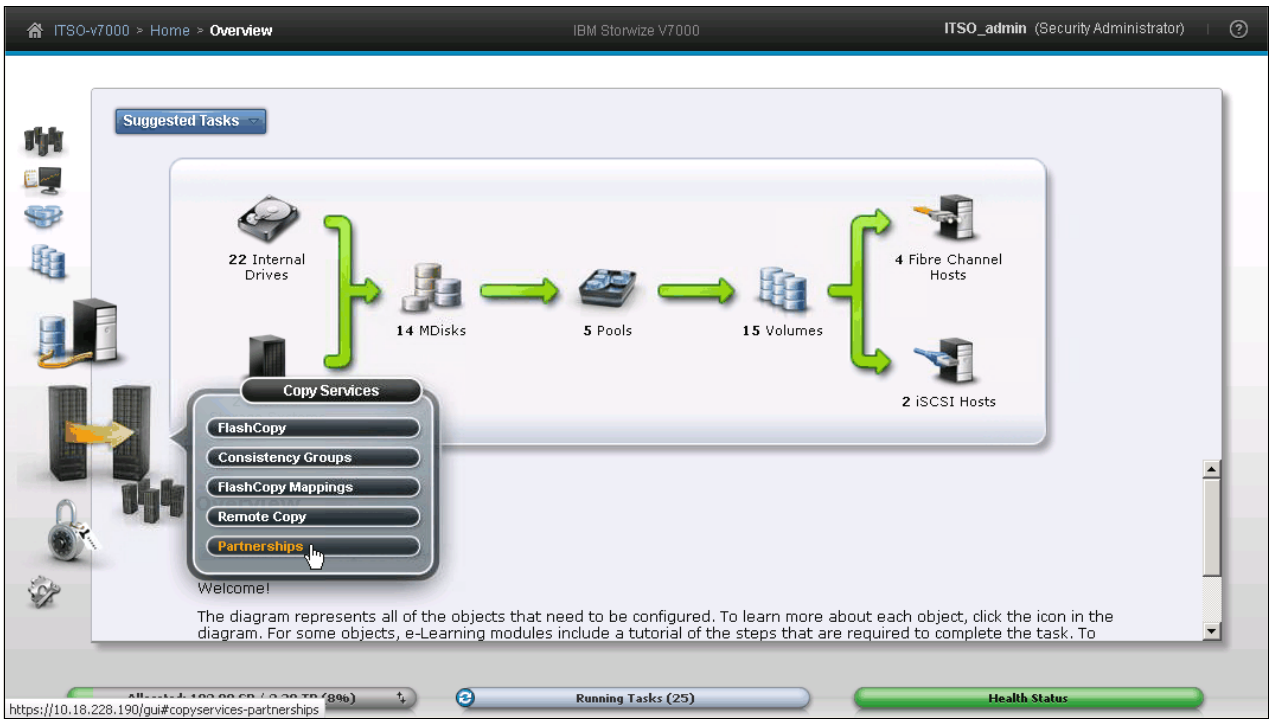

Figure 11-88 Partnership panel

Creating a partnership

No partnership is defined in our example (Figure 11-89), so you must create a partnership between the two partner systems. Click **Create Partnership** on the Partnership panel.

Figure 11-89 Create a system partnership

Depending on which type of partnership you want to create, select either Fibre Channel or IP (Figure 11-90).

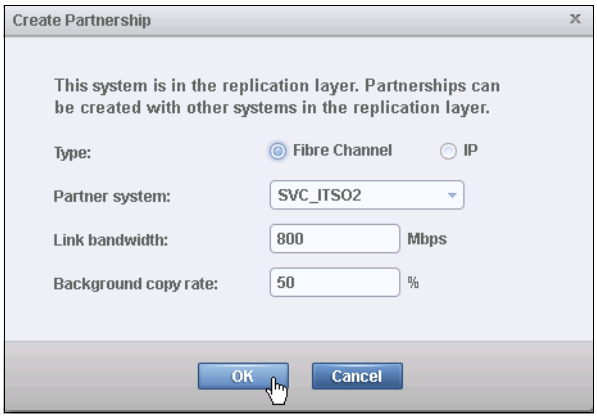

Figure 11-90 Create a Fibre Channel partnership

For Fibre Channel partnership, provide the following information:

- Partner system: Select the partner system for partnership from the drop-down list (available partner candidates are listed automatically after proper SAN zoning is in place)
- Link bandwidth
- Background copy rate (optional)

If there is no partnership candidate, a window shows an alert (Figure 11-91).

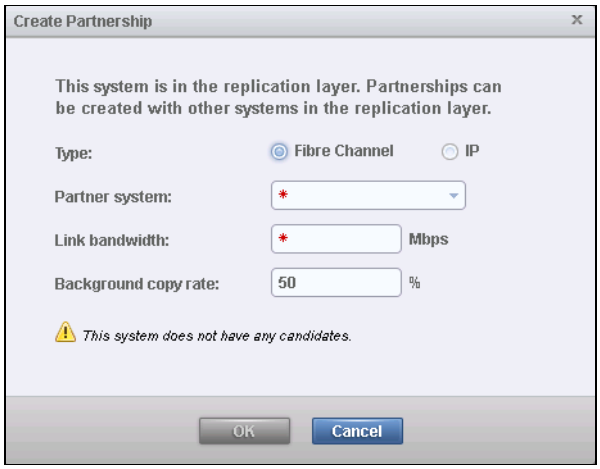

Figure 11-91 No candidates are available to create a partnership

For native IP partnership, provide the following information (Figure 11-92 on page 395):

- Partner system IP address
- Link bandwidth
- Background copy rate (optional)
- Mutual CHAP authentication (optional)
- CHAP secret (optional)

Figure 11-92 Create an IP partnership

The link bandwidth determines the threshold that is used by the background copy process between the systems in the partnership. Set this value so that it is less than or equal to the bandwidth that can be sustained by the communication link between the systems. The link must be able to sustain any host requests and the rate of background copy.

The background copy rate determines the priority that is given to the copy process. A faster rate increases the priority of the process, which might affect performance of other operations.

Click **OK** and the partnership definition is complete on the first partner system. You can find the partnership listed in the Partnership panel in the partially configured state (Figure 11-93).

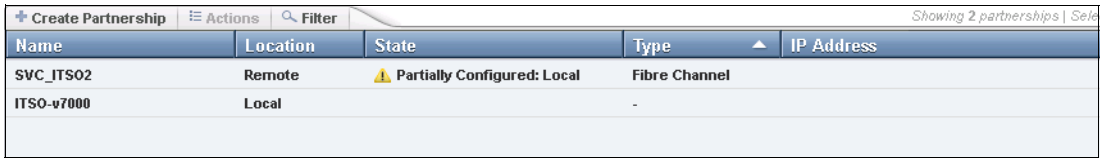

Figure 11-93 Partially configured partnership

Complete the same steps on the second storage system that becomes a fully configured partner (Figure 11-94).

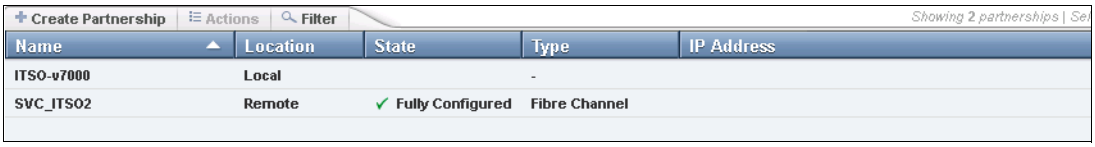

Figure 11-94 Fully configured partnership

The Remote Copy partnership is now implemented between two systems and both systems are ready for further configuration of Remote Copy relationships (Figure 11-95).

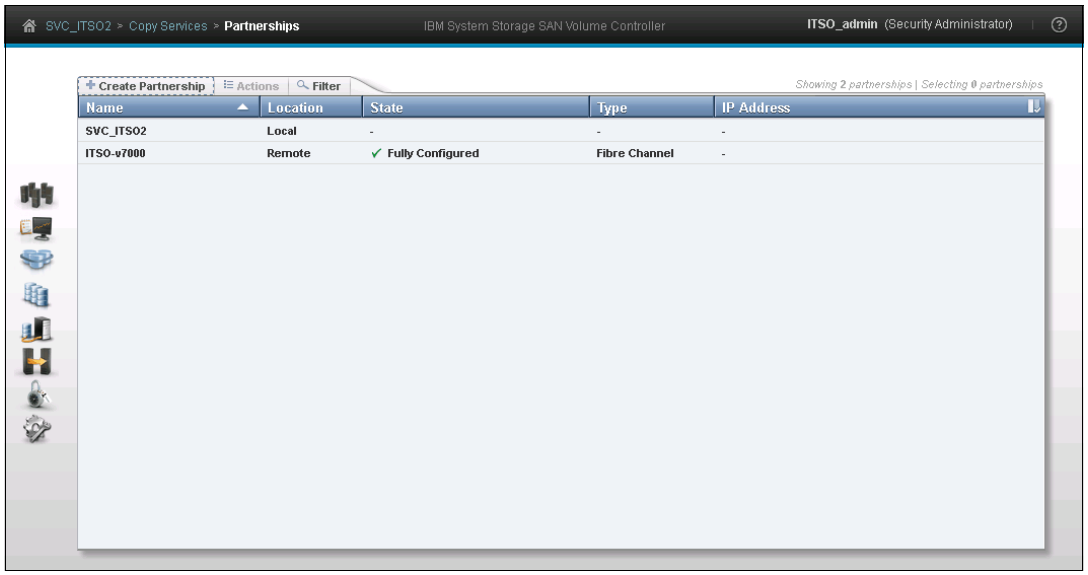

Figure 11-95 Fully configured partnership

You can also edit the partnership properties in the Partnerships panel by clicking **Actions** → **Properties**. Click **Save** to confirm the modifications (Figure 11-96).

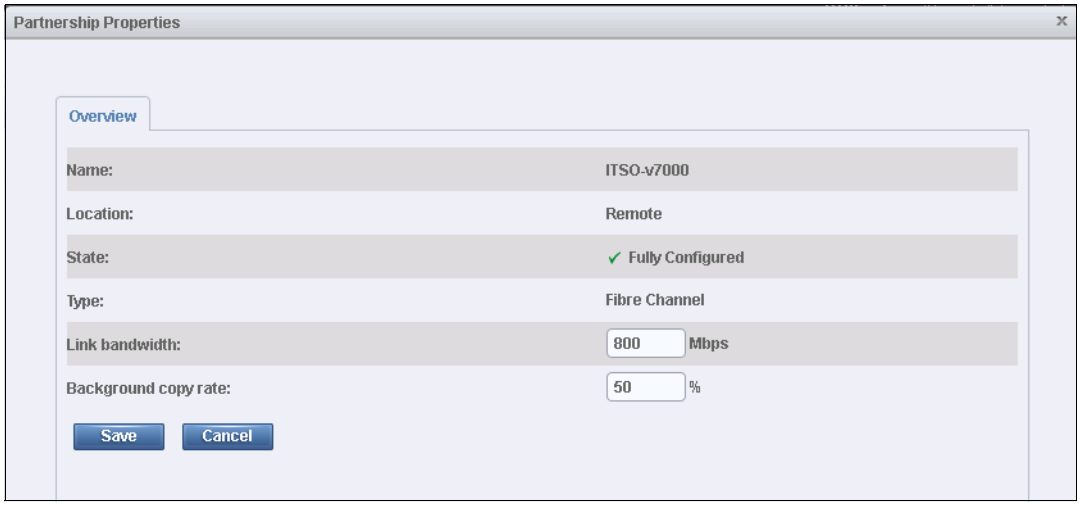

Figure 11-96 Edit partnership properties

Stopping and starting a partnership

To stop the partnership select **Stop** from the pop-up menu (Figure 11-97). If you stop the partnership, the relationships using this partnership are disconnected.

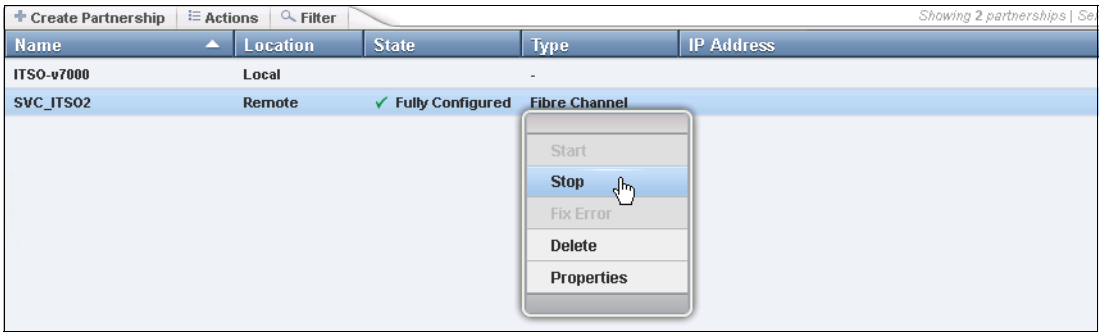

Figure 11-97 Stop the partnership

After you stop the partnership, it is listed as Fully Configured: Stopped (Figure 11-98).

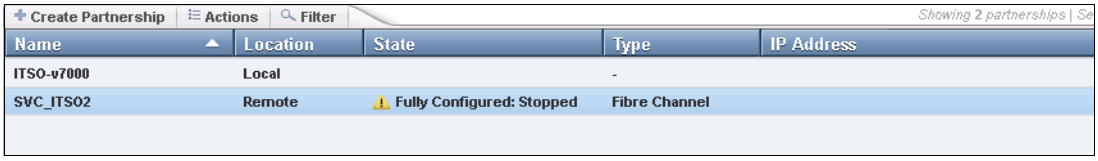

Figure 11-98 Fully configured partnership in stopped state

To restart a stopped partnership, select it and click **Actions** → **Start** (Figure 11-99).

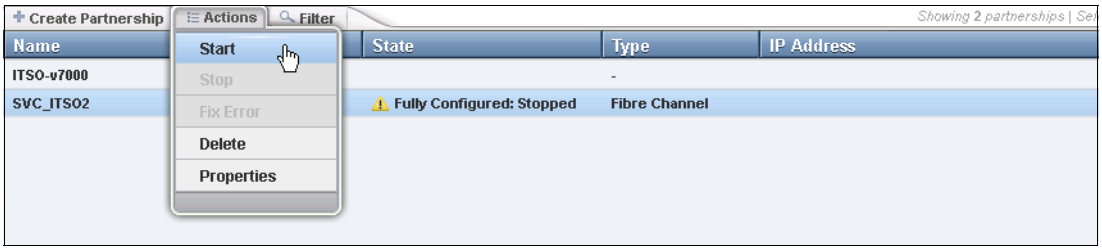

Figure 11-99 Start partnership

The partnership returns to the fully configured status when it has been restarted.

Deleting a partnership

To delete a partnership, right-click it and select **Delete** from the pop-up menu (Figure 11-100).

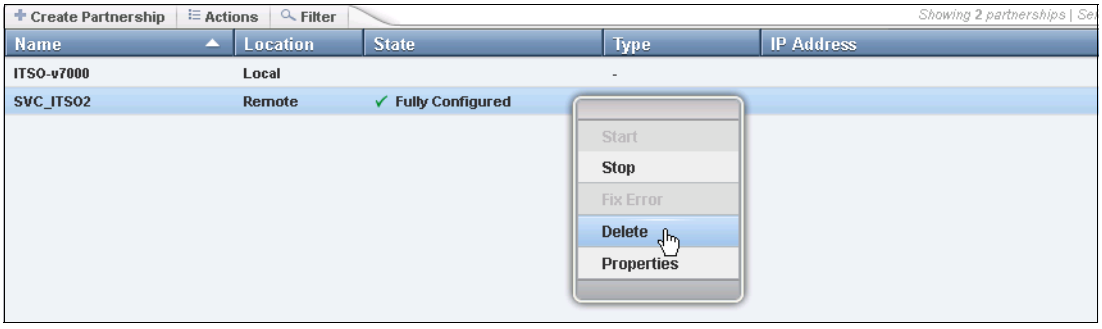

Figure 11-100 Delete a partnership

11.6.2 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. Usage of Remote Copy auxiliary volumes as Remote Copy master volumes is not allowed.

Open the Remote Copy panel to manage Remote Copy by clicking the **Copy Services** icon and clicking **Remote Copy** (Figure 11-101).

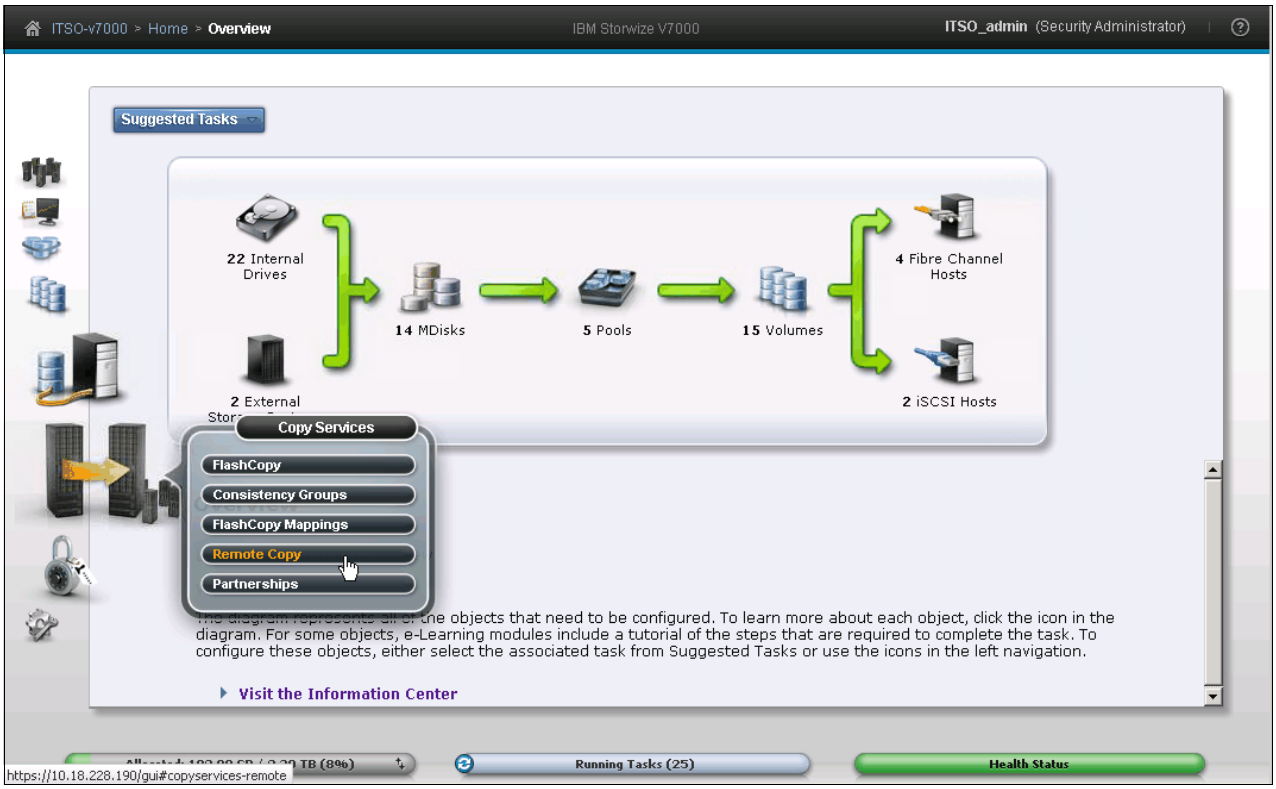

Figure 11-101 Open Remote Copy panel

The Remote Copy panel (Figure 11-102) is where you can manage Remote Copy relationships and Remote Copy consistency groups.

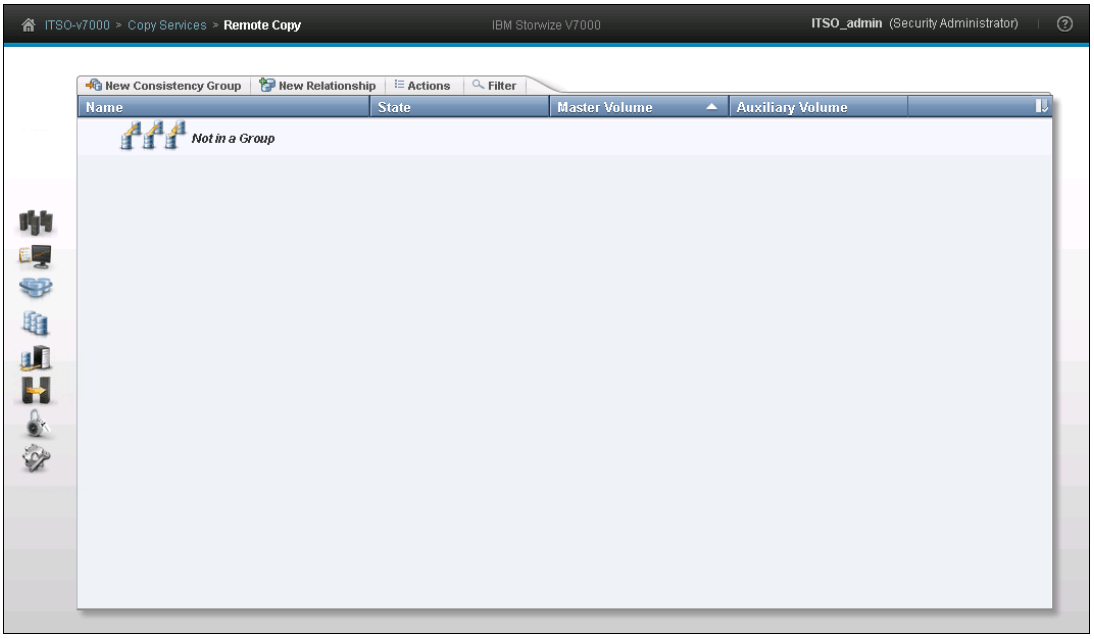

Figure 11-102 Remote Copy panel

The Remote Copy panel has a consistency group filter to list Remote Copy consistency groups that meet your requirements. You can also take actions on the Remote Copy relationships and Remote Copy consistency group. Click **Not in a Group** to display all Remote Copy relationships that are not in any Remote Copy consistency groups.

Creating stand-alone Remote Copy relationships

To create a Remote Copy relationship, click **New Relationship** at the top of the Remote Copy panel (Figure 11-102). A wizard opens and guides you through the Remote Copy relationship creation process.

As shown in Figure 11-103, you must set the Remote Copy relationship type first. Based on your requirements, you can select Metro Mirror (synchronous replication) or Global Mirror (asynchronous replication). Select the appropriate replication type and click **Next**.

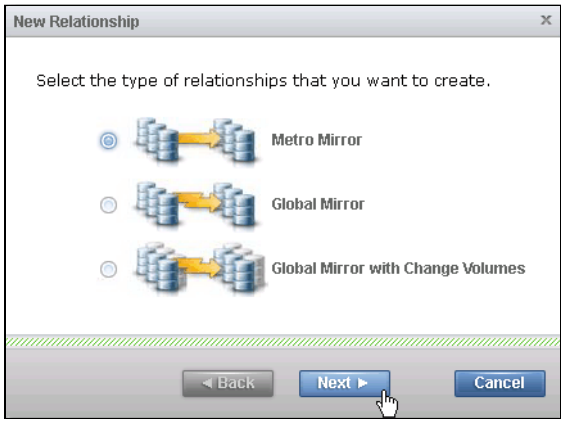

Figure 11-103 Select the appropriate Remote Copy type

Select your Remote Copy auxiliary (target) storage system, the local system, or the already-defined second storage system as the Remote Copy partner. For this example (Figure 11-104), choose another system to build an inter-cluster relationship. Click **Next**.

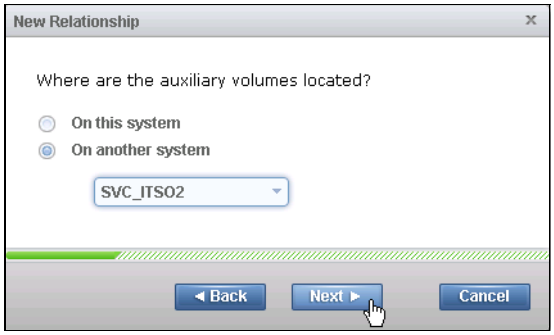

Figure 11-104 Select Remote Copy partner

The Remote Copy master and auxiliary volume need to be specified. Both volumes must have the same size. As shown in Figure 11-105, the system offers only appropriate auxiliary candidates with the same volume size as the selected master volume. After you select the volumes based on your requirement, click **Add**.

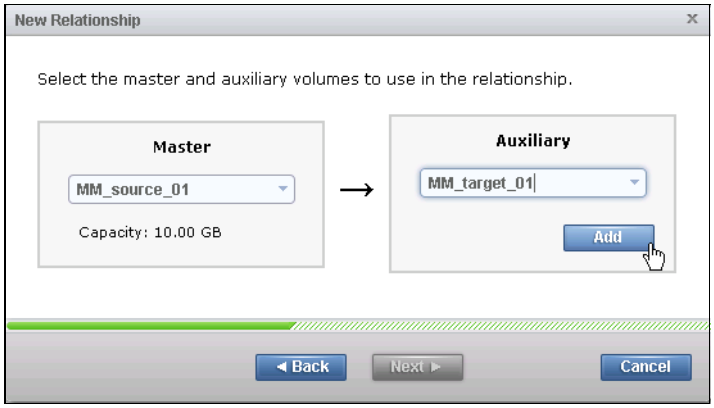

Figure 11-105 Select the master and auxiliary volume

To define multiple and independent relationships click **Add**. To remove a relationship, click the red **X**. In this example, create two independent Remote Copy relationships (Figure 11-106).

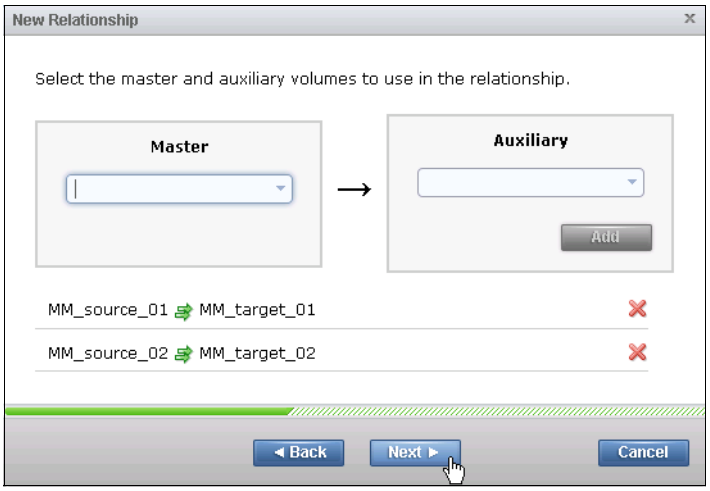

Figure 11-106 Define multiple independent relationships

A window opens and asks if the volumes in the relationship are already synchronized (Figure 11-107). In most situations, the data on the master volume and on the auxiliary volume are not identical, so select **No** and click **Next** to enable an initial copy.

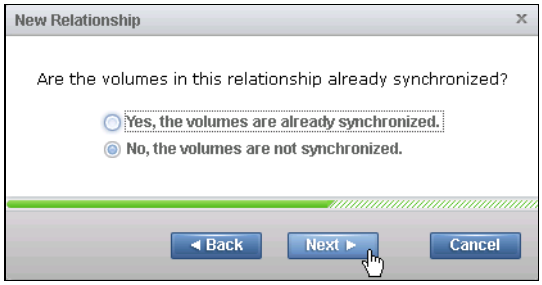

Figure 11-107 Activate initial data copy

If you select **Yes, the volumes are already synchronized**, a warning message opens (Figure 11-108). Be sure that the volumes are identical, and then click **Yes**.

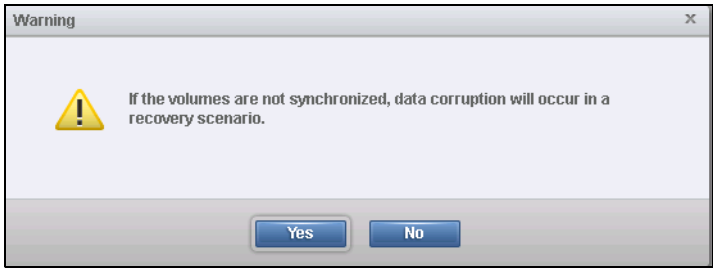

Figure 11-108 Warning message to make sure that the volumes are synchronized

You can choose to start the initial copying progress now, or wait to start it at a later time. In this example, select **Yes, start copying now** and click **Finish** (Figure 11-109).

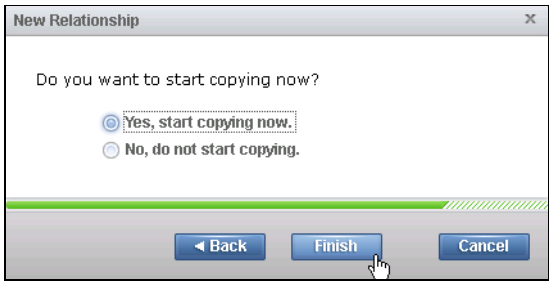

Figure 11-109 Choose if you want to start copying now or later

After the Remote Copy relationships creation completes, two independent Remote Copy relationships are defined and displayed in the Not in a Group list (Figure 11-110).

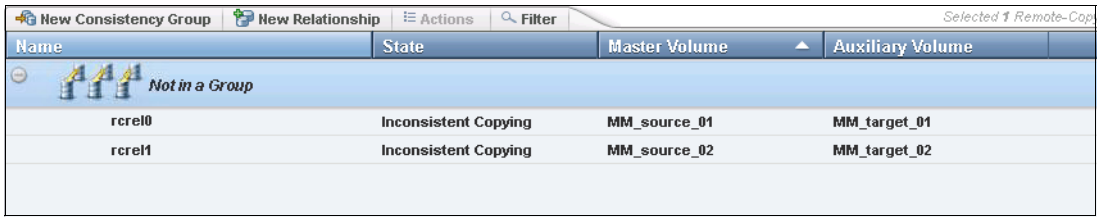

Figure 11-110 Remote Copy relationship creation completes

Optional: You can monitor the ongoing initial synchronization in the Running Tasks status indicator (Figure 11-111).

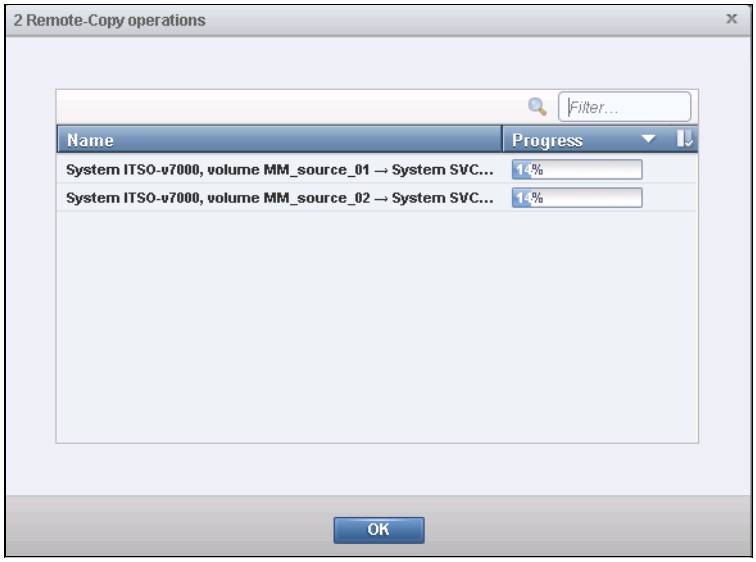

Figure 11-111 Remote copy initialization progress through Running Tasks

Stopping a stand-alone Remote Copy relationship

To stop the Remote Copy relationship, right-click the relationship and select **Stop** from the pop-up menu (Figure 11-112).

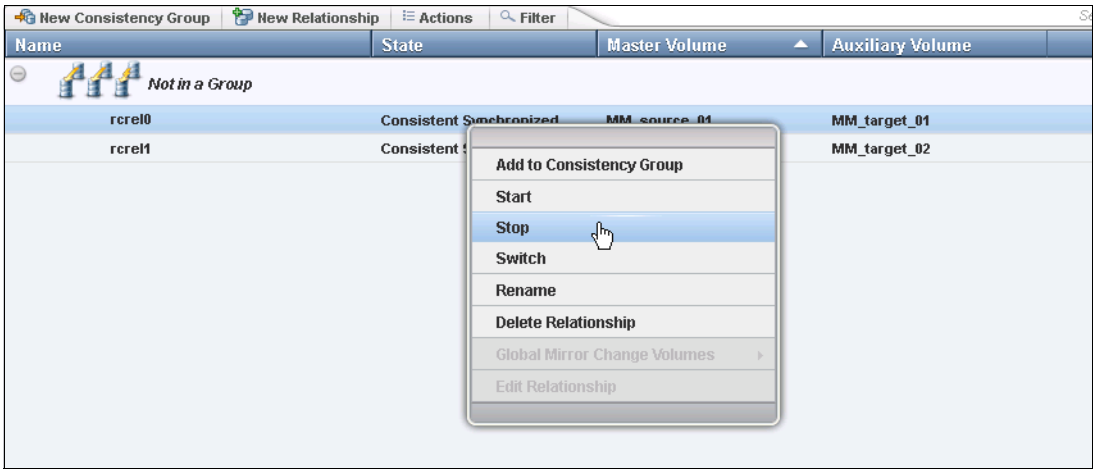

Figure 11-112 Stop Remote Copy relationship

Allow secondary read/write access, if required, and click **Stop Relationship** (Figure 11-113).

Figure 11-113 Option to allow secondary read/write access

After the stop completes, the state of the Remote Copy relationship is changed from Consistent Synchronized to Idling (Figure 11-114). Read and write access to both volumes is now allowed.

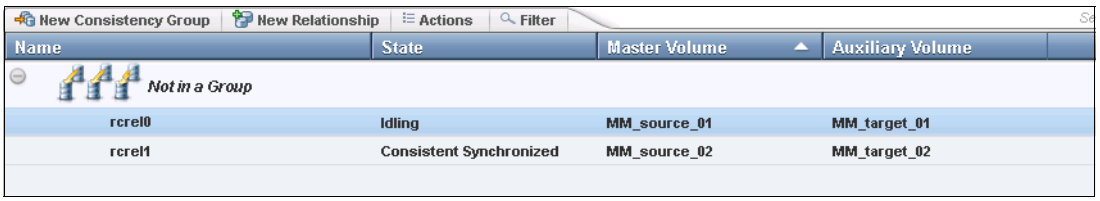

Figure 11-114 Remote Copy relationship stop completes

Starting a stand-alone Remote Copy relationship

To start the Remote Copy relationship, right-click the relationship and select **Start** from the pop-up menu (Figure 11-115).

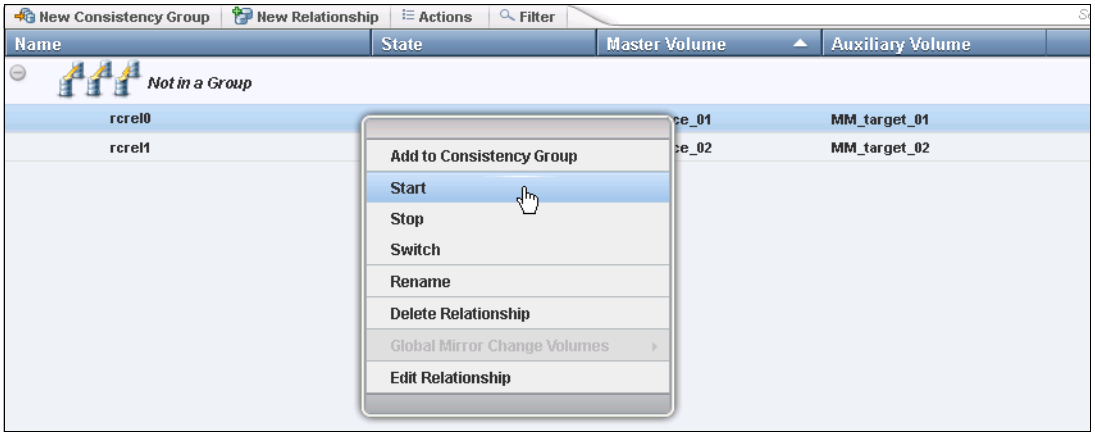

Figure 11-115 Start a Remote Copy relationship

When you start a Remote Copy relationship, the most important item is selecting the copy direction. Both master and auxiliary volumes can be the primary. Base your decision on your requirements and click **Start Relationship**. In this example, choose the master volume to be the primary (Figure 11-116).

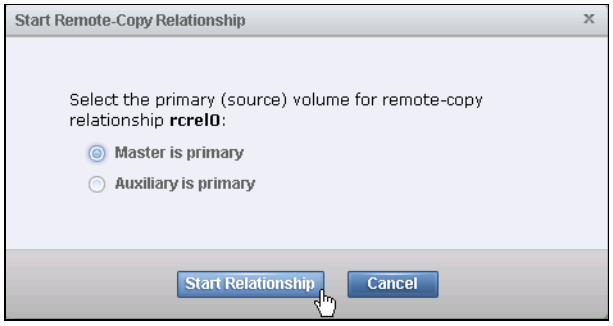

Figure 11-116 Choose the copy direction

Switching the direction of a stand-alone Remote Copy relationship

To switch the copy direction of the Remote Copy relationship, right-click the relationship and select **Switch** from the pop-up menu (Figure 11-117).

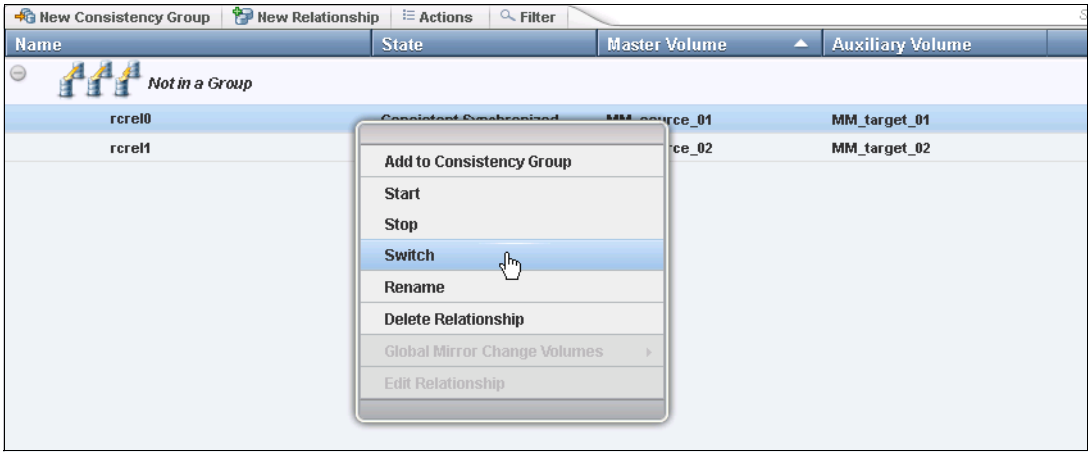

Figure 11-117 Switch Remote Copy relationship

A warning message opens, showing the consequences of this action (Figure 11-118). 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, and 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, stop your host I/O to the current primary volume in advance. Make sure that you are prepared for the consequences, and, if so, click **Yes** to continue.

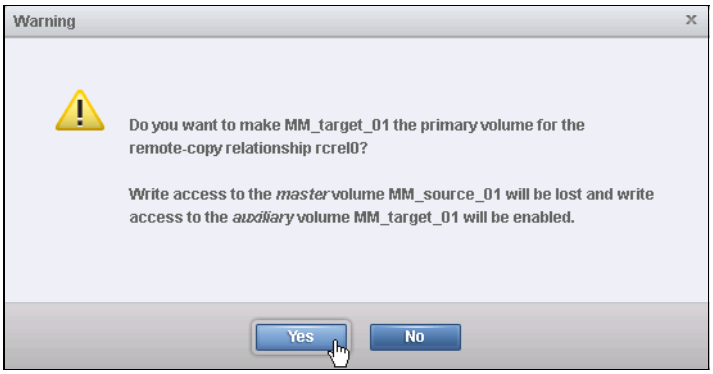

Figure 11-118 Warning message for switching direction of a Remote Copy relationship

After the switch completes, your Remote Copy relationship is tagged (Figure 11-119), and shows that the primary volume in this relationship has been changed.

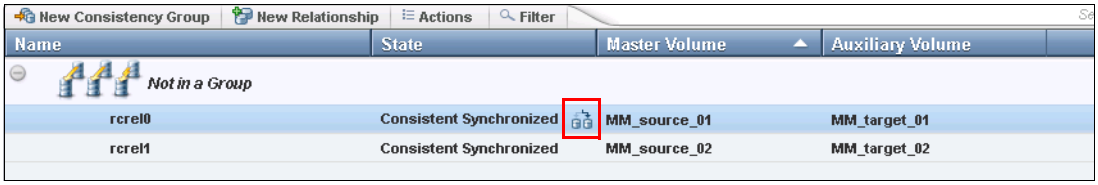

Figure 11-119 The switch icon on the state of the relationship

Renaming a stand-alone Remote Copy relationship

To rename the Remote Copy relationship, right-click the relationship and select **Rename** from the pop-up menu (Figure 11-120).

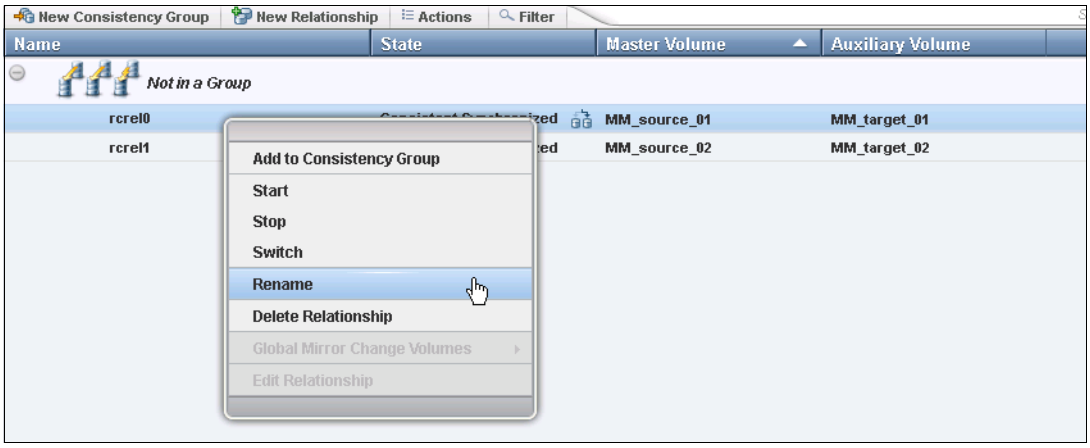

Figure 11-120 Rename the Remote Copy relationship

Enter the new name for the Remote Copy relationship and click **Rename** (Figure 11-121).

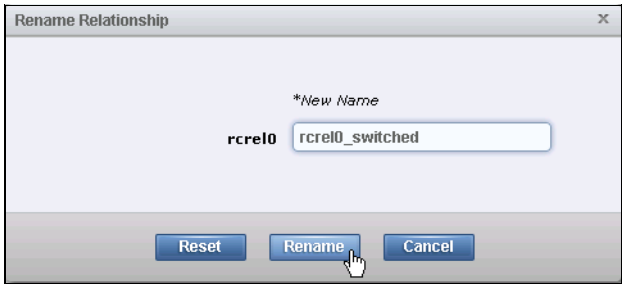

Figure 11-121 Enter a new name for the Remote Copy relationship

Deleting a stand-alone Remote Copy relationship

To delete the Remote Copy relationship, right-click the relationship and select **Delete Relationship** from the pop-up menu (Figure 11-122).

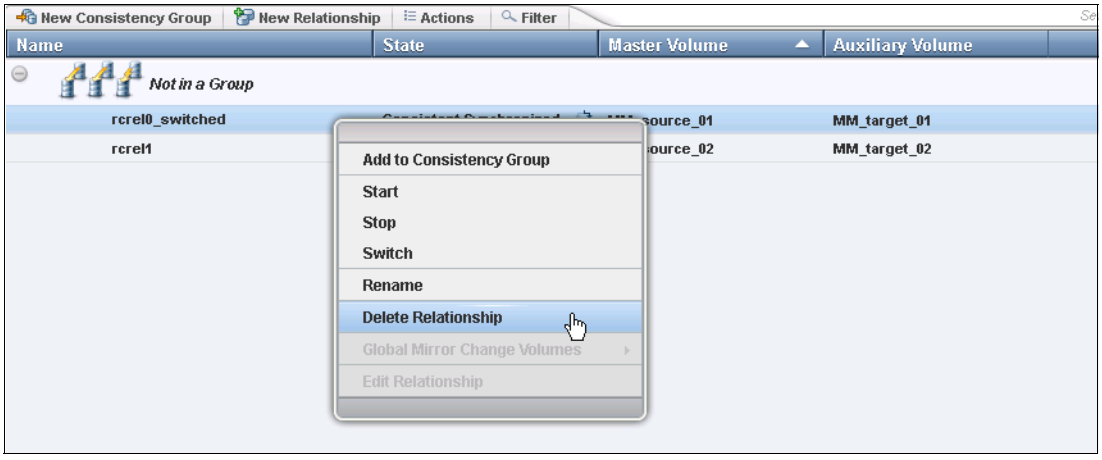

Figure 11-122 Delete a Remote Copy relationship

Confirm this deletion by verifying the number of relationships to be deleted (Figure 11-123). Click **Delete** to proceed.

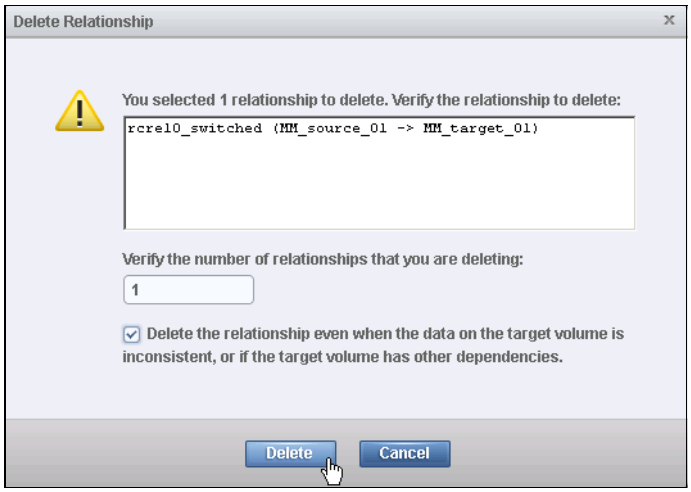

Figure 11-123 Confirm the relationship deletion

11.6.3 Managing a Remote Copy consistency group

A Remote Copy consistency group can also be managed from the Remote Copy panel.

Creating a Remote Copy consistency group

To create a Remote Copy consistency group, click **New Consistency Group** (Figure 11-124).

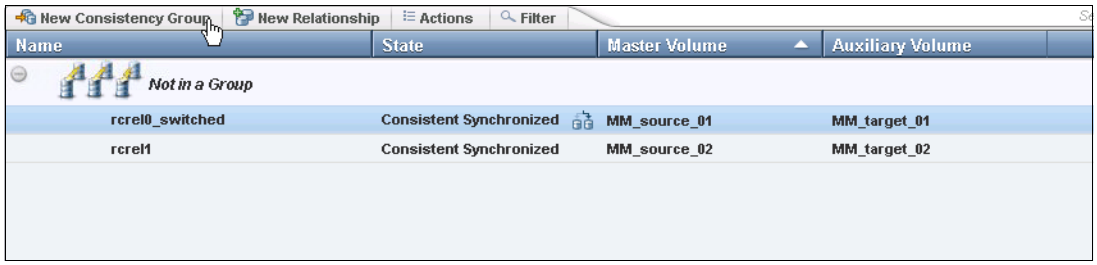

Figure 11-124 Create a consistency group

Enter a name for your new consistency group (Figure 11-125).

Figure 11-125 Enter a name for the new consistency group

In the next window (Figure 11-126), select the system on which the auxiliary volumes reside and click **Next**.

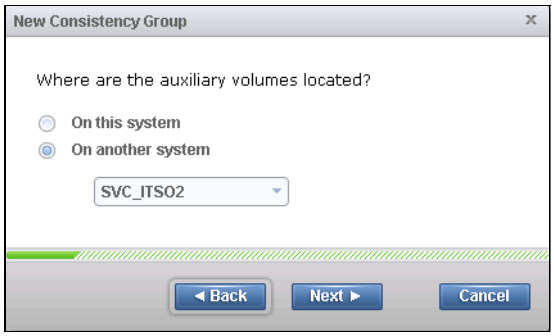

Figure 11-126 Select the system with auxiliary volumes

You can choose to create an empty consistency group or, if you want, add Remote Copy relationships into the consistency group now (Figure 11-127). If you select **Yes, add relationships to this group**, you can select existing relationships or create new ones to add to the consistency group. In our example, create an empty consistency group, and add Remote Copy relationships to the consistency group afterward. Click **Finish** to proceed.

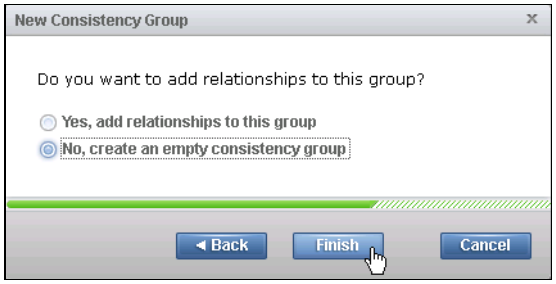

Figure 11-127 Create an empty consistency group

After the creation process completes, a new empty consistency group is listed on the Remote Copy panel (Figure 11-128).

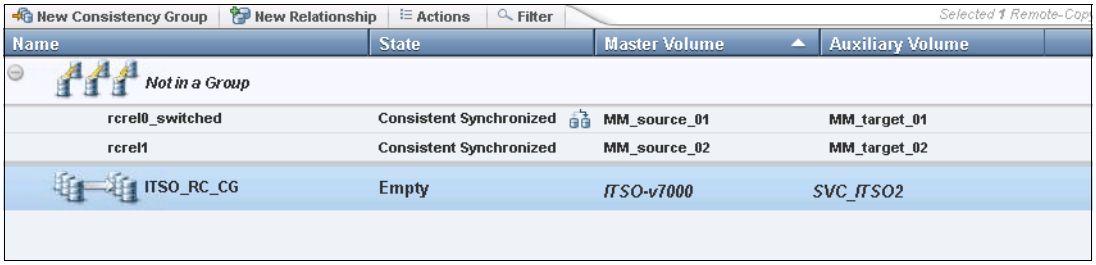

Figure 11-128 New Remote Copy consistency group

You can find the name and the status of the consistency group next to the Relationship icon. At the upper part of the Remote Copy panel, you can take actions on the Remote Copy consistency group. The actions on the Remote Copy relationships can be applied here by using the **Actions** drop-down menu.

Adding a Remote Copy relationship to a consistency group

To add the Remote Copy relationships in the *Not in a Group* list to a consistency group, right-click the volumes and select **Add to Consistency Group** from the pop-up menu (Figure 11-129).

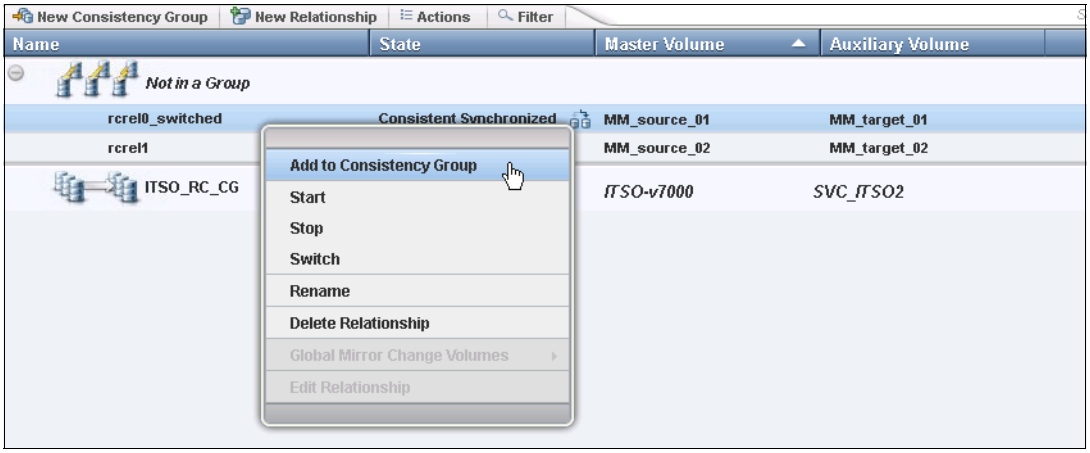

Figure 11-129 Add Remote Copy relationships to a consistency group

You must choose the consistency group to which to add the Remote Copy relationships. Based on your requirements, select the appropriate consistency group and click **Add to Consistency Group** (Figure 11-130).

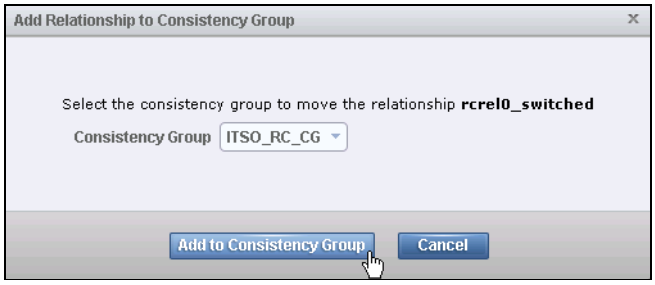

Figure 11-130 Choose the consistency group to add the remote copies to

Your Remote Copy relationships are now in the consistency group you selected.

Starting a consistency group

To start the Remote Copy relationship, right-click the consistency group and select **Start** from the pop-up menu (Figure 11-131).

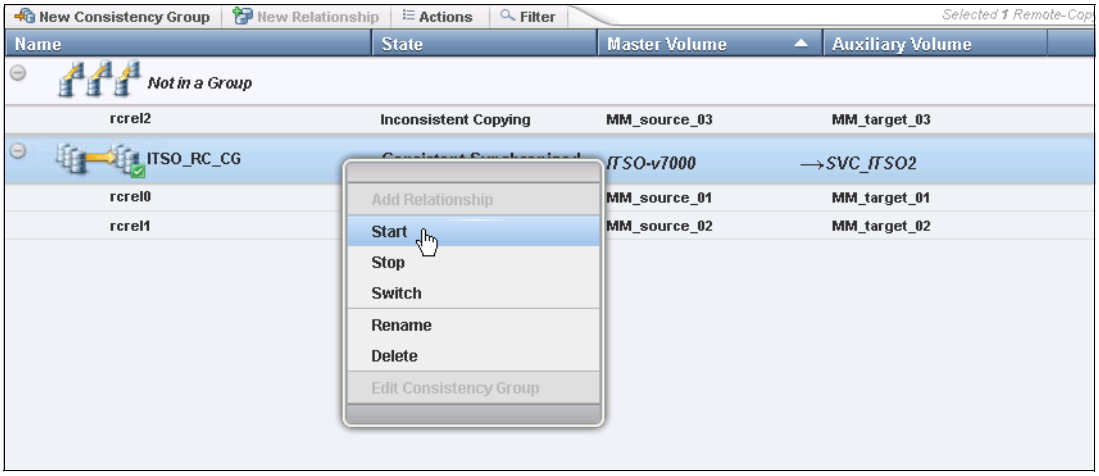

Figure 11-131 Start the consistency group

Stopping a consistency group

To stop the Remote Copy relationship, right-click it and select **Stop** from the pop-up menu (Figure 11-132).

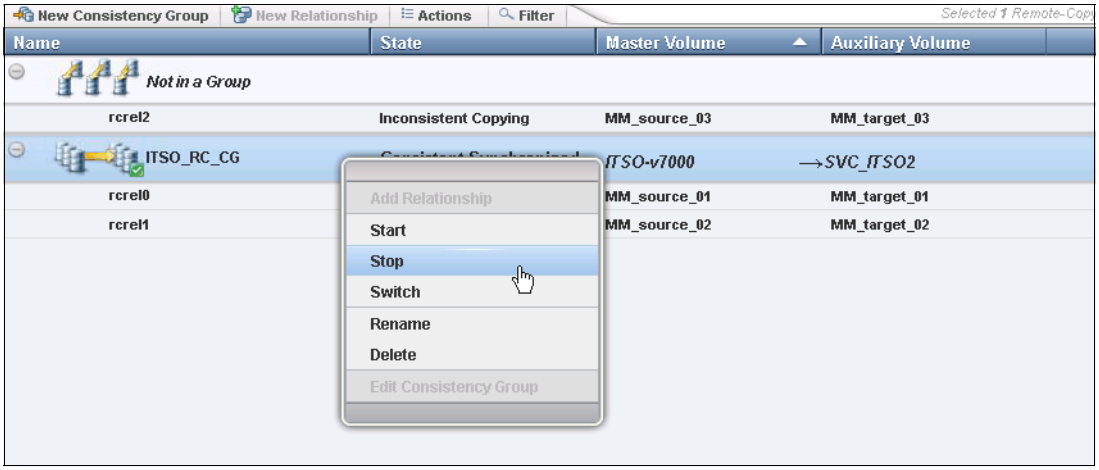

Figure 11-132 Stop the consistency group

To allow read/write access to secondary volumes, select the check box (Figure 11-133) and click **Stop Consistency Group**.

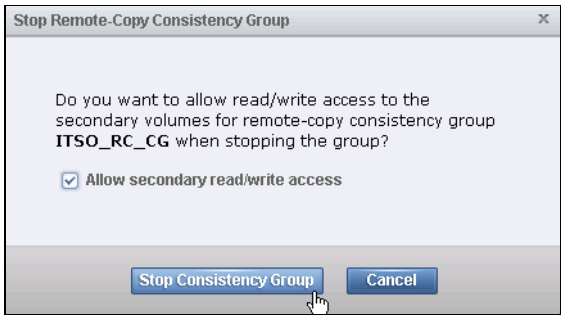

Figure 11-133 Confirm consistency group stop and decide to allow secondary read/write access

Switching a consistency group

As with the switch action on the Remote Copy relationship, you can switch the copy direction of the consistency group also. To switch the copy direction of the consistency group, select **Switch** from the pop-up menu (Figure 11-134).

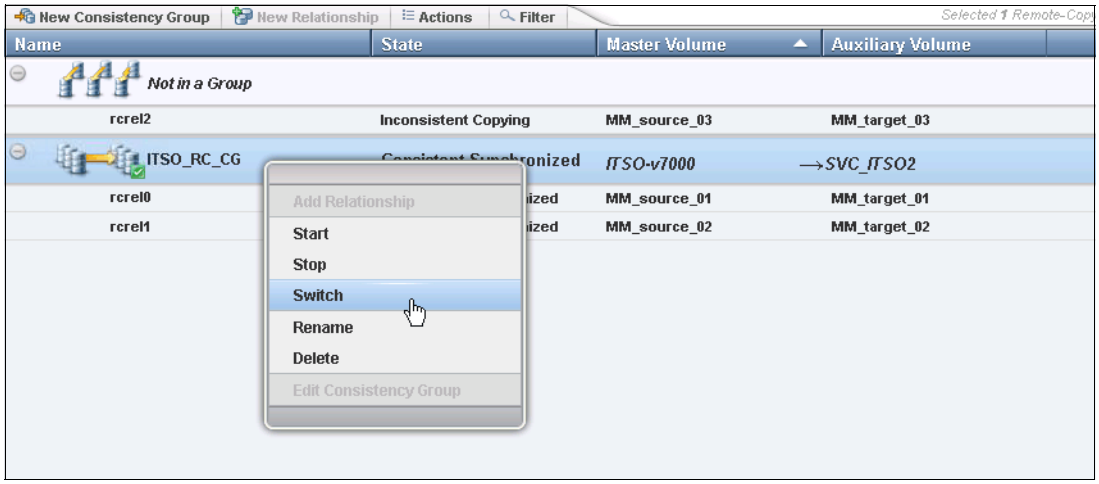

Figure 11-134 Switch the copy direction of a consistency group

A warning message is issued (Figure 11-135 on page 412). After the switch, the primary system in the consistency group is changed. Write access to current master volumes is lost; write access to the current auxiliary volumes is enabled. Make sure that these settings are what you need, and, if so, click **Yes** to continue.

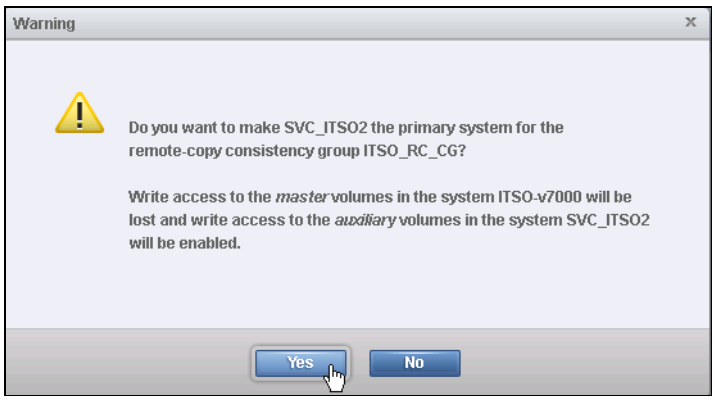

Figure 11-135 Warning message to confirm the switch

Removing Remote Copy relationships from a consistency group

To remove the Remote Copy relationships from the consistency group, select the Remote Copy relationships, and select **Remove from Consistency Group** from the pop-up menu (Figure 11-136).

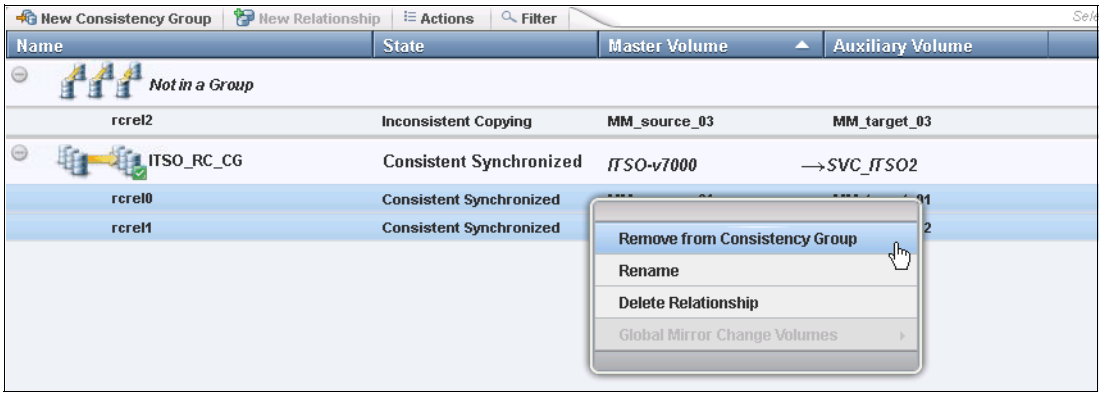

Figure 11-136 Remove Remote Copy relationships from a consistency group

You are prompted to confirm the Remote Copy relationships you want to remove from the consistency group (Figure 11-137). Make sure the Remote Copy relationships listed in the box are those that you want to remove from the consistency group, and click **Remove**.

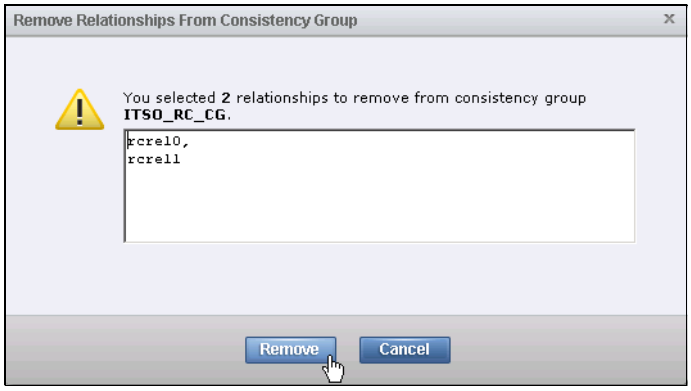

Figure 11-137 Confirm the relationships to remove from the Remote Copy consistency group

After the removal process completes, the Remote Copy relationships are deleted from the consistency group and displayed in the *Not in a Group* list.

Deleting a consistency group

To delete the consistency group, right-click it and select **Delete** from the pop-up menu (Figure 11-138).

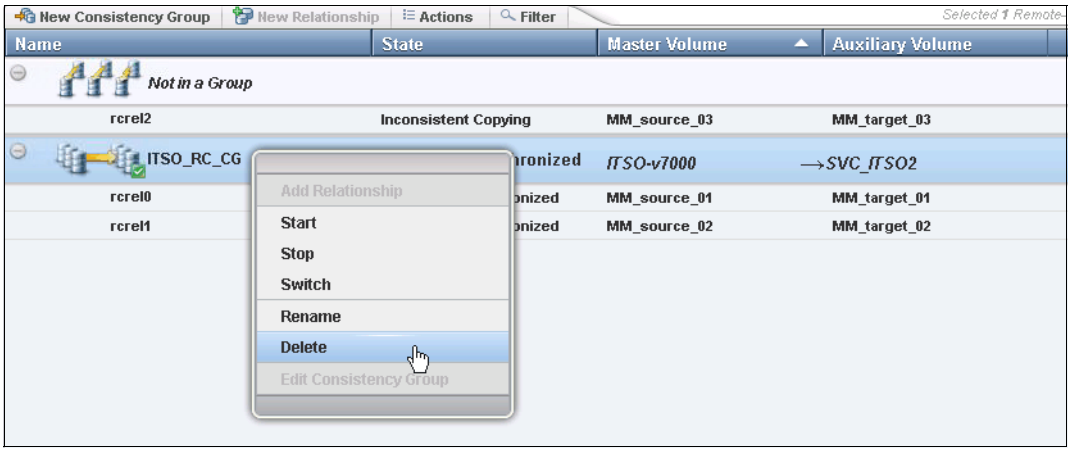

Figure 11-138 Delete a consistency group

Confirm the deletion of the consistency group (Figure 11-139). Click **Yes** if you are sure that this consistency group should be deleted.

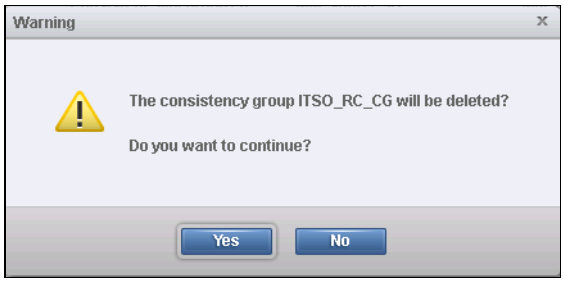

Figure 11-139 Warning to confirm deletion of the consistency group

The consistency group is deleted. After the deletion process completes, the Remote Copy relationships are removed from the consistency group and shown in the *Not in a Group* list.

12

IBM Tivoli Storage Productivity Center

IBM Tivoli Storage Productivity Center offers storage infrastructure management that helps optimize storage management by centralizing, simplifying, automating, and optimizing storage tasks associated with storage systems, data disaster recovery, storage networks, and capacity management. This chapter provides an overview of the main functions in Tivoli Storage Productivity Center. We also discuss the new packaging and licensing in the V5.1 release and demonstrate the installation and initial application configuration.

Tivoli Storage Productivity Center offers these advantages:

- Helps centralize the management of your storage infrastructure from a single interface using role-based administration and single sign-on.
- Provides a single management application with modular integrated components that are easy to install and provide common services for simple and consistent configuration and consistent operations across host, fabric, and storage systems.
- Manages performance and connectivity from the host file system to the physical disk, including in-depth performance monitoring and analysis on SAN fabric performance.
- Manages the capacity utilization and availability of storage systems, file systems, and databases.
- Monitors, manages, and controls (zone) SAN fabric components.
- Automates capacity provisioning of file systems.
- Monitors and tracks the performance of SAN-attached native API or SMI-S compliant storage devices.
- Manages advanced replication services (Global Mirror, Metro Mirror, and FlashCopy).

Tivoli Storage Productivity Center can help you manage the capacity utilization of storage systems, file systems, and databases. It can help automate file system capacity provisioning, and perform device configuration and management of multiple devices from a single user interface. It can tune and proactively manage the performance of storage devices on a storage area network (SAN) and manage, monitor, and control your SAN fabric.

12.1 Tivoli Storage Productivity Center overview

In this book, we give a basic overview of Tivoli Storage Productivity Center, and show how to perform a typical setup (in 12.4, “Installing Tivoli Storage Productivity Center components” on page 434) and connect it to the IBM Storwize V7000 (in 12.5, “Connecting Tivoli Storage Productivity Center to an IBM Storwize V7000 system” on page 443). Remember that the functions depend on the type of Tivoli Storage Productivity Center license you have and are not included with IBM Storwize V7000.

Contact your IBM Business Partner or IBM Representative to obtain the correct license for your requirements.

12.1.1 Tivoli Storage Productivity Center V5.1 features and functions

Tivoli Storage Productivity Center V5.1 includes these features and functions:

- A next-generation, web-based user interface offers ease of use. The user interface provides a common “look and feel” that is based on the current user interfaces for IBM XIV Storage System, IBM Storwize V7000, and IBM System Storage SAN Volume Controller. It enables quick access to key storage assets, status, and performance information.
- Reporting based on IBM Cognos helps create and integrate custom reports on capacity, performance, and utilization.
- Replication features include support for external failover by external products, allowing users to define Warning and Severe thresholds for recovery point objective (RPO) alerts, and providing the ability to export replication history to a table for analysis.
- An integrated Install Anywhere wizard helps simplify installation and post-installation configuration to help accelerate administrator time-to-value. The integrated installer is a major improvement to the install process. Now, one installer can do single-server installation, multiple-server installations, and check for prerequisites and tell you where to get them if you do not have them.
- Offers better user management and integration with external user repositories, like Microsoft Active Directory.
- Enhanced management for virtual environments provides enhanced reporting for virtual servers (VMware).
- Tiered Storage Optimization provides integration with the existing storage optimizer and storage tiering reporting. Tiered Storage Optimization is policy-driven information lifecycle management (ILM), leveraging virtualization technology to provide recommendations for storage re-location. It provides recommendations for workload migration based on user defined policy based on file system level data, performance and capacity utilization. This ensures only highest performing workloads are allocated to most expensive storage.

12.1.2 Tivoli Storage Productivity Center V5.1 license models

IBM Tivoli Storage Productivity Center V5.1 products include two license models:

- IBM Tivoli Storage Productivity Center V5.1

  Tivoli Storage Productivity Center V5.1 provides device management capabilities such as automated system discovery, provisioning, data replication, configuration, and performance monitoring for storage systems and storage networks.

- IBM Tivoli Storage Productivity Center Select Edition V5.1

  Tivoli Storage Productivity Center Select Edition V5.1 offers the same features as Tivoli Storage Productivity Center V5.1 but at attractive entry-level pricing for operations with smaller capacities. It is licensed per storage device, such as disk controllers and their respective expansion units.

12.1.3 IBM Tivoli Storage Productivity Center Select Edition V5.1

Tivoli Storage Productivity Center now enables enterprise-wide management of IBM Storwize V7000 systems. In addition to device-level management software that is packaged with IBM Storwize V7000 systems, Tivoli Storage Productivity Center offers incremental benefits that are included in its latest version. Tivoli Storage Productivity Center Select Edition is ideally suited for use with IBM Storwize V7000 systems.

The benefits of Tivoli Storage Productivity Center Select Edition are as follows:

- Simplified deployment and visibility: Tivoli Storage Productivity Center Select Edition supports IBM Storwize V7000 systems during discovery as a new device type with different types of managed disks. The quick discovery and configuration capabilities of Tivoli Storage Productivity Center enable you to attach the storage device with ease and helps you configure efficiently, that is, plan for replication while provisioning the device. Tivoli Storage Productivity Center Select Edition also enables launch-in-context and single sign-on, significantly reducing the burden on storage administrators.
- The Tivoli Storage Productivity Center topology viewer offers a collective view that includes IBM Storwize V7000 systems, helps differentiate between external array-based disks and local disks, and shows tiering information. Tivoli Storage Productivity Center Select Edition also extends thin provisioning support for IBM Storwize V7000 systems, enabling increased utilization and lowered costs.
- Performance management: An IBM Storwize V7000 system offers unmatched performance and availability among midrange disk systems. Adding Tivoli Storage Productivity Center Select Edition enhances performance monitoring by capturing metrics such as input and output (I/O) and data rates, and cache utilization from a single console. Tivoli Storage Productivity Center helps establish threshold levels based on business priorities and alerts when these levels are breached, while proactively managing performance and service levels by tracking historical information.
- Replication management: Tivoli Storage Productivity Center for Replication gives IBM Storwize V7000 systems superior disaster recovery (DR) management, provides central control of the replication environment, and helps establish FlashCopy and Metro Mirror and Global Mirror relationships.
- Comprehensive reporting: Tivoli Storage Productivity Center Select Edition offers detailed metrics that include performance data for storage subsystems, controllers, cache, I/O, arrays, disk groups, and ports. These performance statistics can be stored in database tables for later use, so that storage administrators can track and measure service levels.

- Easy Tier: Tivoli Storage Productivity Center Select Edition also provides information about Easy Tier to monitor automatic migration of data assets to high-performance disk tiers (such as solid-state drives (SSDs)).
- Advanced management: Tivoli Storage Productivity Center offers performance metrics-based recommendations for provisioning, including SAN planning (and DR planning with Tivoli Storage Productivity Center for Replication). The *Disk Magic* model of Tivoli Storage Productivity Center helps identify "hot spots," improving storage optimization for IBM Storwize V7000.

12.2 Tivoli Storage Productivity Center architecture

The IBM Tivoli Storage Productivity Center consists of the following key components:

- Data Server
- Device Server
- IBM Tivoli Integrated Portal
- Tivoli Storage Productivity Center for Replication
- IBM DB2 Database
- Agents
- Interfaces

12.2.1 Data Server

This component is the control point for product scheduling functions, configuration, event information, reporting, and graphical user interface (GUI) support. It coordinates communication with and data collection from agents that scan file systems and databases to gather storage demographics and populate the database with results. Automated actions can be defined to perform file system extension, data deletion, backup or archiving, or event reporting when defined thresholds are encountered. The Data Server is the primary contact point for GUI user interface functions. It also includes functions that schedule data collection and discovery for the Device Server.

12.2.2 Device Server

This component discovers, gathers information from, analyzes performance of, and controls storage subsystems and SAN fabrics. It coordinates communication with and data collection from agents that scan SAN fabrics and storage devices.

12.2.3 IBM Tivoli Integrated Portal

Tivoli Storage Productivity Center V4 is integrated with Tivoli Integrated Portal. This integration provides functions such as single sign-on and the use of IBM Tivoli Common Reporting:

- Single sign-on: Enables you to access Tivoli Storage Productivity Center and then Tivoli Storage Productivity Center for Replication using a single user ID and password.
- Tivoli Common Reporting: Tivoli Common Reporting is a component provided by Tivoli Integrated Portal. It is one possible option to implement customized reporting solutions using SQL database access, providing output in HTML, PDF, or Microsoft Excel.

12.2.4 Tivoli Storage Productivity Center for Replication

Starting with Tivoli Storage Productivity Center V4.1, the Tivoli Storage Productivity Center for Replication product is integrated into Tivoli Storage Productivity Center. Currently, the integration is limited to basic functions, such as providing launch-in-context links in the Tivoli Storage Productivity Center GUI, and cross-checks when a volume is deleted with Tivoli Storage Productivity Center and mapping of user roles.

12.2.5 IBM DB2 Database

A single database instance serves as the repository for all Tivoli Storage Productivity Center components. This repository is where all of your storage information and usage statistics are stored. All agent and user interface access to the central repository is done through a series of calls and requests made to the server. All database access is done using the server component to maximize performance and to eliminate the need to install database connectivity software on your agent and UI systems.

12.2.6 Agents

Outside of the server, there are several interfaces that are used to gather information about the environment. The most important sources of information are the Tivoli Storage Productivity Center agents (Storage Resource Agent, Data Agent, and Fabric Agent), native APIs, and SMI-S enabled storage devices that use a CIMOM agent (either embedded or as a proxy agent). Storage Resource Agent, native APIs, CIM agents, and Out of Band Fabric Agents gather host, application, storage system, and SAN fabric information and send that information to the Data Server or Device Server.

12.2.7 Interfaces

As Tivoli Storage Productivity Center gathers information from your storage (servers, subsystems, and switches) across your enterprise, it accumulates a repository of knowledge about your storage assets and how they are used. You can use the reports provided in the user interface view and analyze that repository of information from various perspectives to gain insight into the use of storage across your enterprise. The user interfaces (UI) enable users to request information and then generate and display reports based on that information. Certain user interfaces can also be used for configuration of Tivoli Storage Productivity Center or storage provisioning for supported devices.

The following interfaces are available for Tivoli Storage Productivity Center:

- Tivoli Storage Productivity Center GUI: This GUI is the central point of Tivoli Storage Productivity Center administration. Here you have the choice of configuring Tivoli Storage Productivity Center after installation, defining jobs to gather information, initiating provisioning functions, viewing reports, and working with the advanced analytics functions.
- Java Web Start GUI: When you use Java Web Start, the regular Tivoli Storage Productivity Center GUI is downloaded to your workstation and started automatically, so you do not have to install the GUI separately. The main reason for using Java Web Start is that it can be integrated into other products (for example, Tivoli Integrated Portal). By using launch-in-context from those products, you are guided directly to the Select window. The launch-in-context URLs can also be assembled manually and be used as bookmarks.

- TPCTOOL: This interface is a command-line interface (CLI) based program that interacts with the Tivoli Storage Productivity Center Device Server. It is used to extract performance data from the Tivoli Storage Productivity Center repository database to create graphs and charts with multiple metrics, with various unit types and for multiple entities (for example, subsystems, volumes, controller, and arrays) using charting software. Commands are entered as lines of text (that is, sequences of types of characters) and output can be received as text. Furthermore, the tool provides queries, management, and reporting capabilities, but you cannot initiate discoveries, probes, and performance collection from the tool.
- Database access: Starting with Tivoli Storage Productivity Center V4, the Tivoli Storage Productivity Center database provides views that provide access to the data stored in the repository, which allows you to create customized reports. The views and the required functions are grouped together into a database schema called TPCREPORT. For this function, you need to have sufficient knowledge about SQL. To access the views, IBM DB2® supports various interfaces, for example, JDBC and ODBC.

12.3 Preparing Windows to install Tivoli Storage Productivity Center

Tivoli Storage Productivity Center can be installed on one of your host servers. In this section, we cover the Tivoli Storage Productivity Center V4.2 installation wizard for Windows. The installation wizard covers two installation paths, Typical and Custom. We guide you through the installation of the Typical path. The installation in this chapter is not related to any of the different licenses that are available. All editions use the same code base and as such each window is the same for each license.

Note: Because of lab limitations, we demonstrate the Tivoli Storage Productivity Center use case on Tivoli Storage Productivity Center V4.2. To install and configure the Tivoli Storage Productivity Center version 5.1 see the *IBM Tivoli Storage Productivity Center V5.1 Technical Guide*, SG24-8053.

12.3.1 Installation overview

To get Tivoli Storage Productivity Center V4.2 to work, the prerequisite components must be installed before starting the installation wizard. Complete the following steps:

1. Check that the system meets the prerequisites (see “Verifying the system requirements for Windows” on page 422).
2. Install the prerequisite components (see 12.3.3, “Pre-installation steps for Windows” on page 422).
3. Install the Tivoli Storage Productivity Center components (see 12.4, “Installing Tivoli Storage Productivity Center components” on page 434).
4. Connect Tivoli Storage Productivity Center to the IBM Storwize V7000 system (see 12.5, “Connecting Tivoli Storage Productivity Center to an IBM Storwize V7000 system” on page 443).

With the Typical installation you can install all the components of the Tivoli Storage Productivity Center on the local server in one step, but you still can decide which components to install:

- Server:Data Server, Device Server Replication Manager, and Tivoli Integrated Portal
- Clients: Tivoli Storage Productivity Center GUI
- Storage Resource Agent

The drawback of using the Typical installation is that everything besides the selections in the earlier list are set to the defaults. At about 75% of the way into the installation, the installer starts the Tivoli Storage Productivity Center for Replication installation wizard to give you the options to change some installation parameters. You basically must step through it, and click **Finish** to start its installation procedure. After this procedure is done, click **Finish** to return to the Tivoli Storage Productivity Center installer to complete the last few steps of the installation.

12.3.2 Product code layout

In this section, we describe the contents of the product media at the time of the writing of this book. The IBM Passport Advantage® and web media content consists of a disk image and a SRA compressed file. The disk image is broken up into four parts:

- Disk1 part 1 contains these Tivoli Productivity Center components:
  - Database Schema
  - Data Server
  - Device Server
  - GUI
  - CLI
  - Storage Resource Agent (SRA)
- Disk1 part 2 contains these Tivoli Productivity Center components:
  - Tivoli Integrated Portal
  - Tivoli Storage Productivity Center for Replication
- Disk1 part 3 contains IBM Tivoli Integrated Portal Fix Pack.

Disk notes: Part 1, part 2, and part 3 are required for every Tivoli Storage Productivity Center installation and need to be downloaded and extracted to a single directory.

- Disk1 part 4 contains an optional component named IBM Tivoli Storage Productivity Center Monitoring Agent for IBM Tivoli Monitoring.

Directory name: On *Windows*, ensure that the directory name where the installation images are on has *no* spaces or special characters, or the Tivoli Storage Productivity Center installation fails. For example, a failure occurs if you happened to have a directory name such as the following example:

```
C:\tpc 42 standard edition\disk1
```

The SRA compressed file contains Tivoli Storage Productivity Center Storage Resource Agents (SRAs). A Tivoli Storage Productivity Center Storage Resource Agent contains the local agent installation components:

- Storage Resource Agent
- Installation scripts for the Virtual I/O Server

The content of this disk is as follows:

- The `readme` directory
- The `sra` directory
- The `version.txt` file

In addition to the images mentioned, the following images are available:

- Tivoli Storage Productivity Center Storage globalization
- IBM Tivoli Storage Productivity Center for Replication Two Site Business Continuity License, which is available for Windows, Linux, and AIX
- IBM Tivoli Storage Productivity Center for Replication Three Site Business Continuity License, which is available for Windows, Linux, and AIX

12.3.3 Pre-installation steps for Windows

Certain prerequisite components need to be installed before proceeding with the Tivoli Storage Productivity Center storage installation. They are IBM DB2 UDB Enterprise Server Edition 9.1 Fix Pack 2 or later, 9.5 Fix Pack 3a or later, or 9.7 without any Fix Pack.

Fix Pack and Version 9.7: Do not use Version 9.7 with any fix pack, even if it is a later version, as the Tivoli Storage Productivity Center installation will fail.

Verifying the system requirements for Windows

To install Tivoli Storage Productivity Center, the following system requirements must be met:

- One Intel Quad Core Xeon or greater
- 8 GB of memory
- 6 GB of free disk space and 500 MB in the temp directory.

The latest operating system support level matrix is available at the following address:

`http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp?topic=%2Fcom.ibm.tpc_V42.doc%2FfqzO_t_upgrading_all.html`

Verifying the primary domain name systems

Before you start the installation, verify whether a primary domain name system (DNS) suffix is set. This action might require a computer restart.

To verify the primary DNS name, complete the following steps:

1. Right-click **My Computer** on your desktop.
2. Click **Properties**.
3. Click the **Computer Name** tab. On the window that opens, click **Change**.
4. Enter the host name in the Computer name field. Click **More** to continue.
5. In the next window, verify that Primary DNS suffix field displays the correct domain name. Click **OK**.
6. If you made any changes, you must restart your computer for the changes to take effect.

Activating the NetBIOS settings

If NetBIOS is not enabled on Microsoft Windows 2003, then a GUID is not generated. Verify and activate the NetBIOS settings. On your Tivoli Storage Productivity Center Server, click **Start** → **Control Panel** → **Network Connections**. Click **Local Area Connections**. From the Local Area Connection Properties window, double-click **Internet Protocol (TCP/IP)**. In the next window, Internet Protocol (TCP/IP) Properties, click **Advanced**. On the WINS tab, select **Enable NetBIOS over TCP/IP** and click **OK** (Figure 12-1).

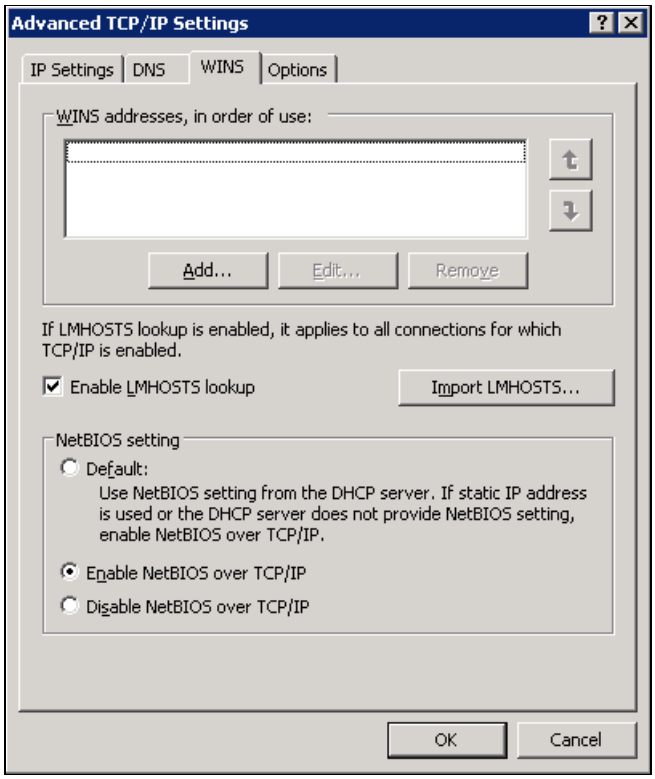

Figure 12-1 Advanced TCP/IP properties

Installing the Tivoli Storage Productivity Center prerequisites

In this section, we show how to install the Tivoli Storage Productivity Center prerequisites on Windows. We perform a typical installation of DB2 9.7 without a fix pack.

Before beginning the installation, it is important that you log on to your system as a local administrator with Administrator authority.

DB2 installation

To begin the installation of DB2, complete the following steps:

1. Insert the IBM DB2 Installer CD into the CD drive.

   If Windows autorun is enabled, the installation program starts automatically. If it does not start automatically, open Windows Explorer, go to the DB2 Installation image path, and double-click the `setup.exe` file.

> **User ID:** Only the user ID that has installed the DB2 product has the privilege to issue the **`db2start`** and **`db2stop`** commands.

The Welcome window opens (Figure 12-2). Click **Install a Product** to proceed with the installation.

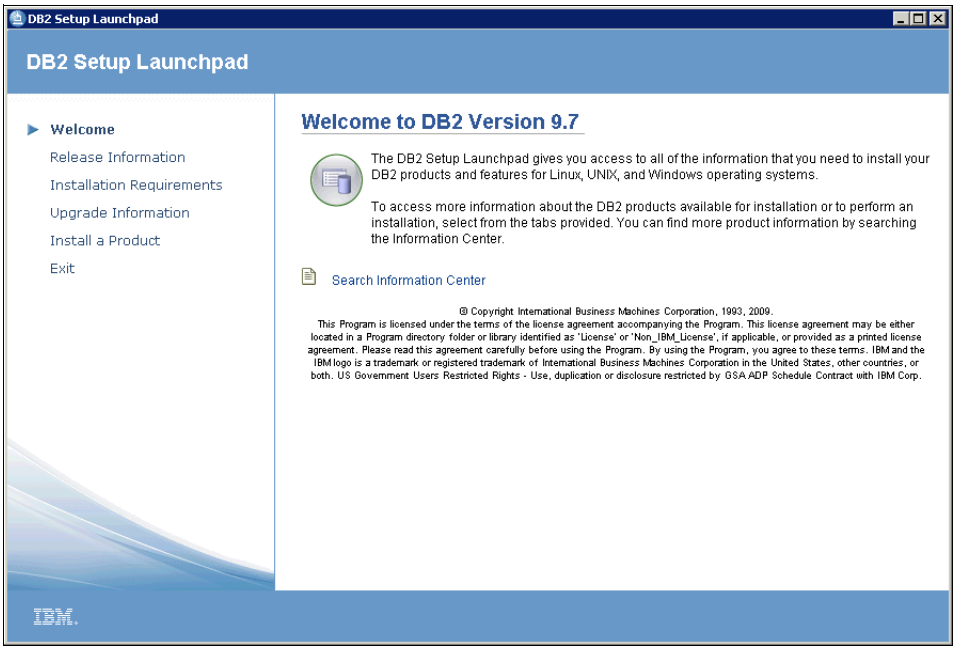

Figure 12-2 DB2 Setup Welcome window

2. The next window allows you to select the DB2 product to be installed. Select **DB2 Enterprise Server Edition Version 9.7** by clicking **Install New** to proceed (Figure 12-3).

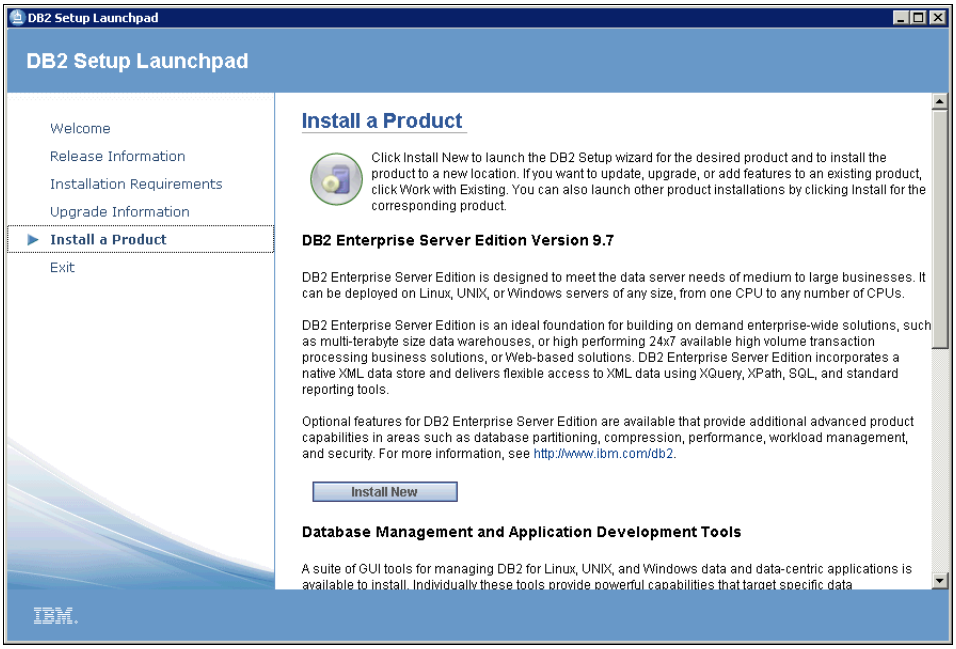

Figure 12-3 Select product

3. The DB2 Setup wizard window opens (Figure 12-4). Click **Next** to proceed.

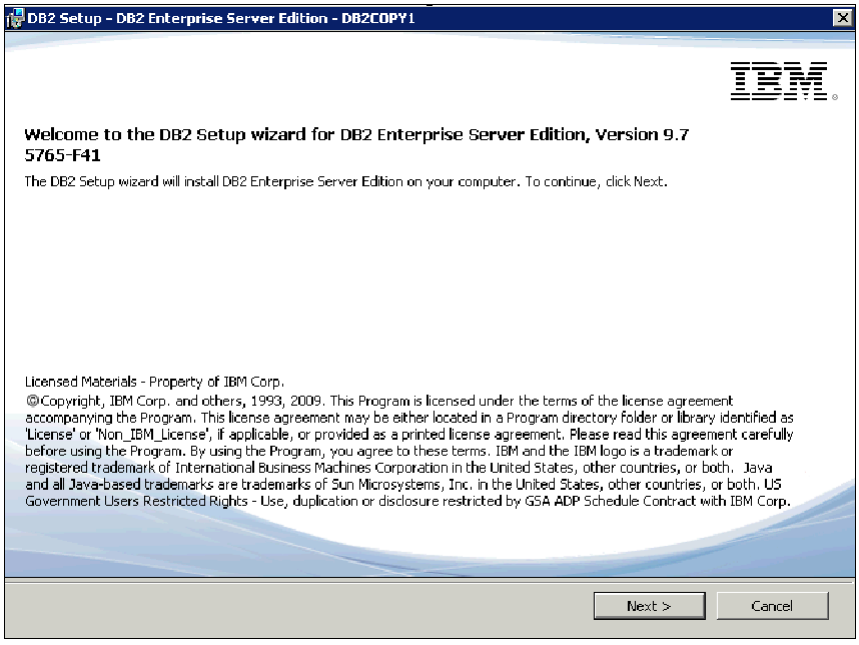

Figure 12-4 Setup wizard

4. Read the license agreement; click **I accept the terms in the license agreement** (Figure 12-5).

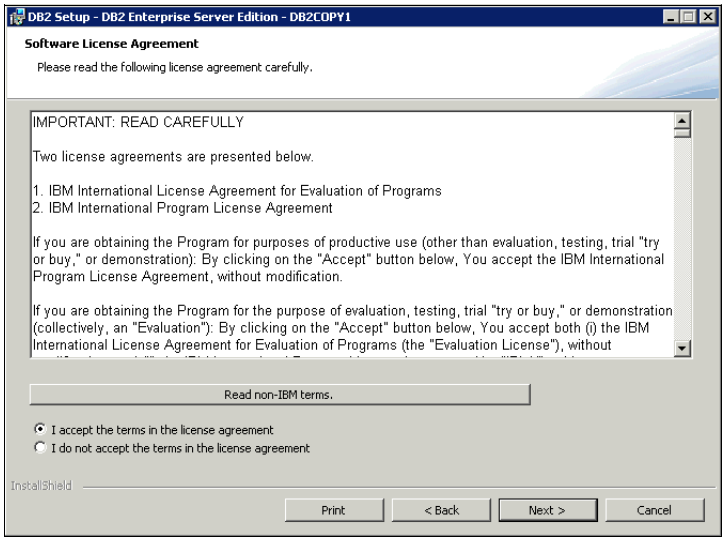

Figure 12-5 License agreement

5. Accept the default installation type, **Typical**, and click **Next** to continue (Figure 12-6).

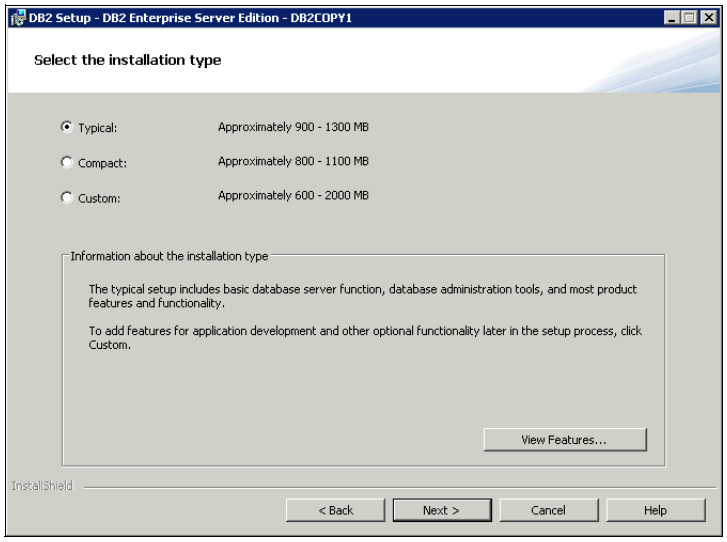

Figure 12-6 Typical installation

6. Select **Install DB2 Enterprise Server Edition on this computer** (Figure 12-7). Click **Next** to continue.

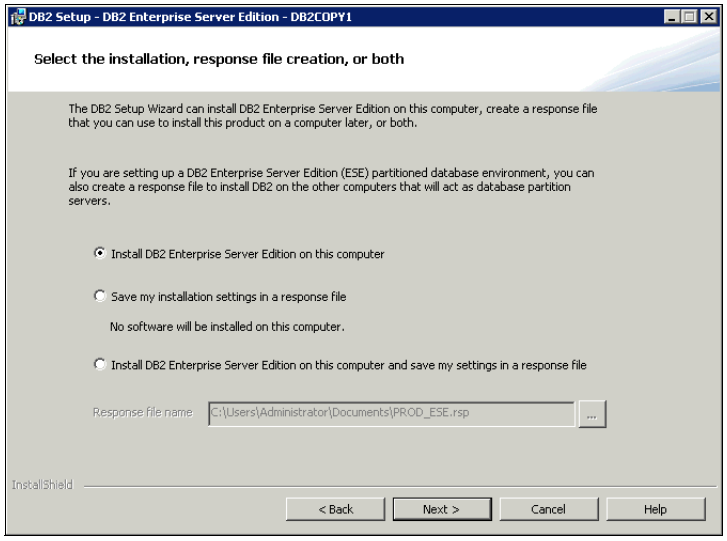

Figure 12-7 Installation action

7. Figure 12-8 shows the default values for the drive and directory to be used as the installation folder. You can change these values or accept the defaults, and then click **Next** to continue. In our installation, we accept to install on the `C:\` drive.

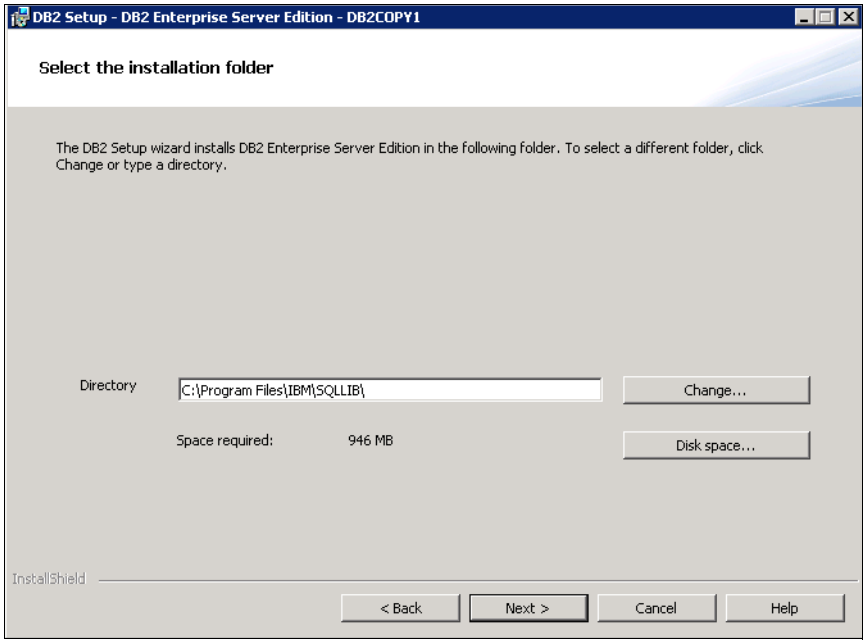

Figure 12-8 Windows installation folder

8. The next window requires user information for the DB2 Administration Server; it can be a Windows domain user. If it is a local user, select **None - use local user account** for the Domain field.

   The user name field has a default user name already present. You can change it or leave the default and type the password of the DB2 user account that you want to create (Figure 12-9). Leave **Use the same user name and password for the remaining DB2 services** cleared and click **Next** to continue.

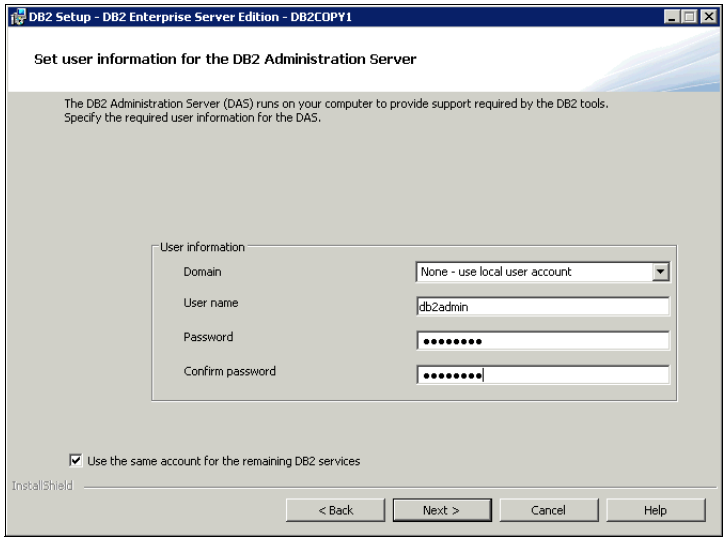

Figure 12-9 User information

DB2 creates a user with the following administrative rights:

- Act as a part of an operating system.
- Create a token object.
- Increase quotas.
- Replace a process-level token.
- Log on as a service.

9. Select **Create the default DB2 Instance** and click **Next** (Figure 12-10).

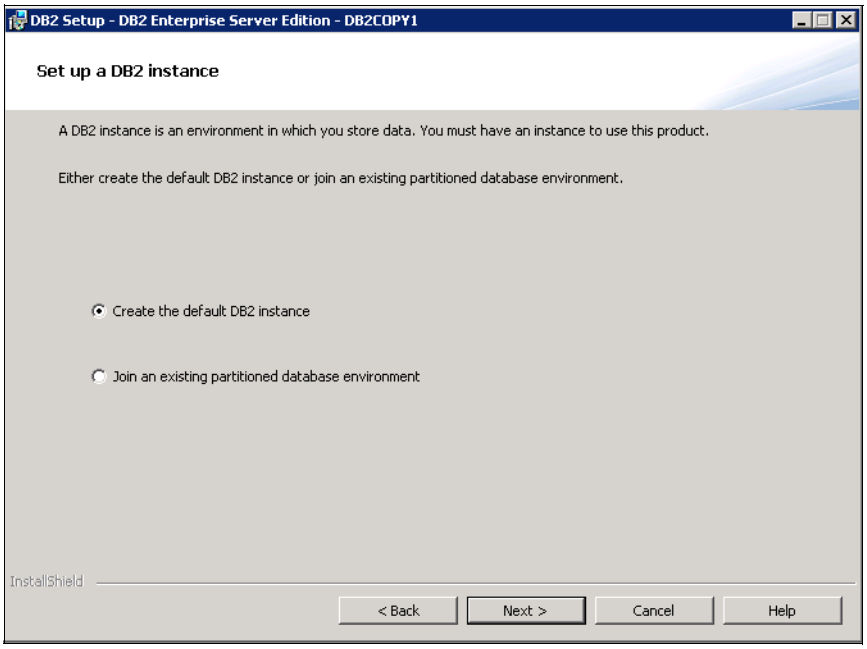

Figure 12-10 Create default DB2 instance

10. Select **Single Partition Instance** and click **Next** (Figure 12-11).

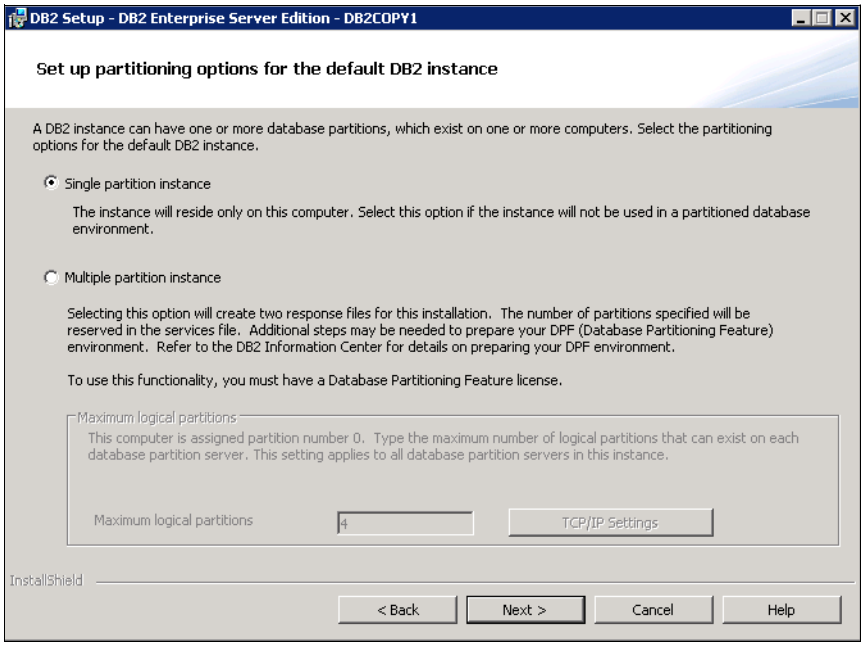

Figure 12-11 Single partition instance

11. Accept the default DB2 Instance and click **Next** to continue (Figure 12-12).

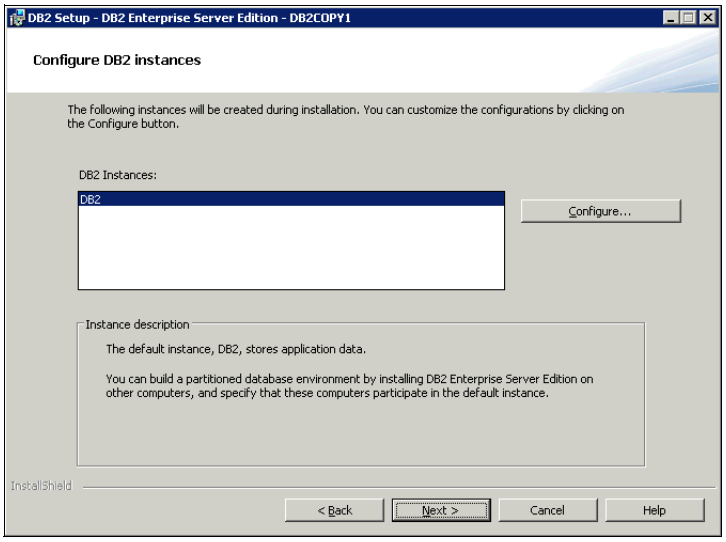

Figure 12-12 Configure DB2 instances

12. Specify options to prepare the DB2 tools catalog (Figure 12-13). Verify that **Prepare the DB2 tools catalog on this computer** is clear. Click **Next** to continue.

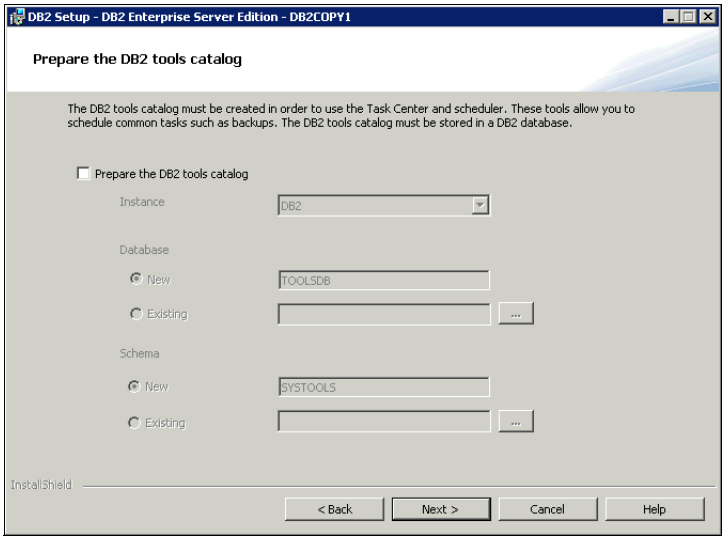

Figure 12-13 Prepare the DB2 tools catalog

13. Set the DB2 server to send notifications when the database needs attention (Figure 12-14). Ensure that **Set up your DB2 server to send notification** is clear and click **Next** to continue.

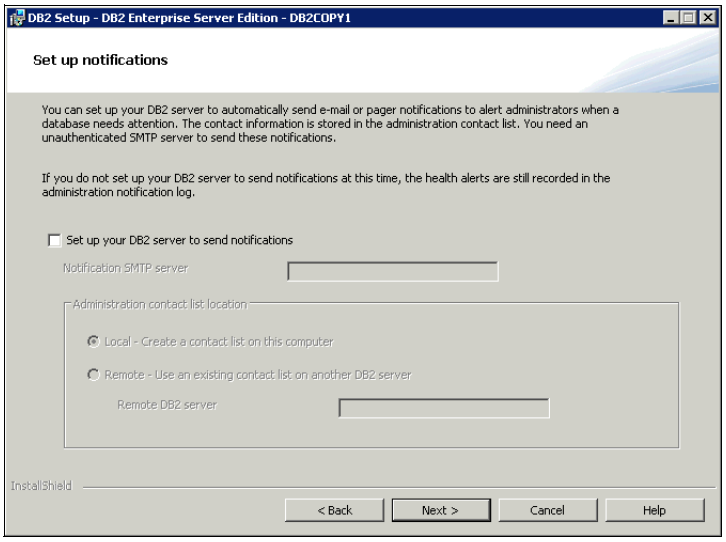

Figure 12-14 Health Monitor

14. Accept the defaults for the DB2 administrators group and DB2 users group in the Enable operating system security for DB2 objects window (Figure 12-15) and click **Next** to proceed.

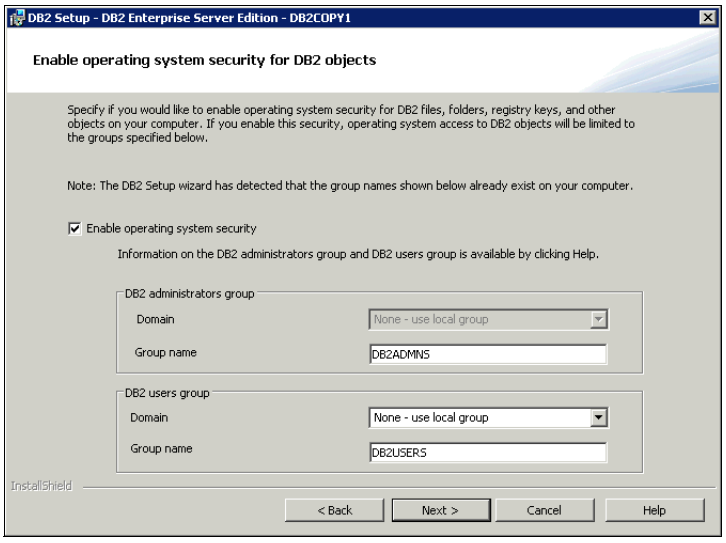

Figure 12-15 Enable operating system security for DB2 objects

15. Figure 12-16 shows the summary window of what will be installed, based on your input. Review the settings and click **Finish** to continue.

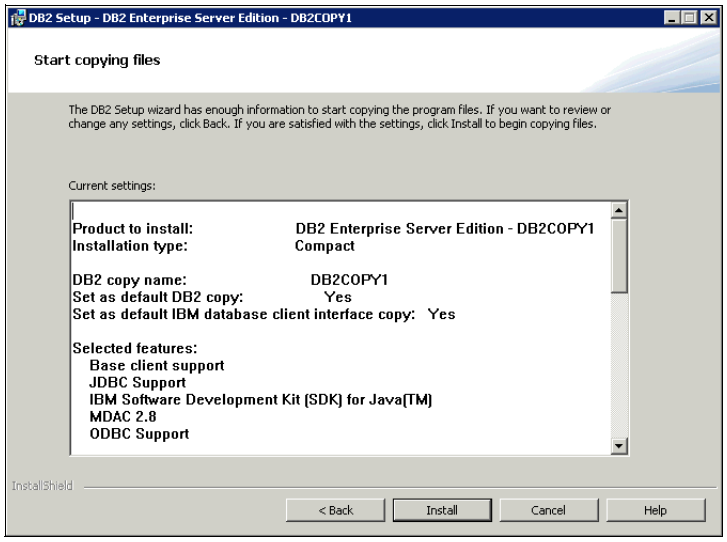

Figure 12-16 Summary window

As DB2 installation proceeds, you see a progress window similar to the one shown in Figure 12-17.

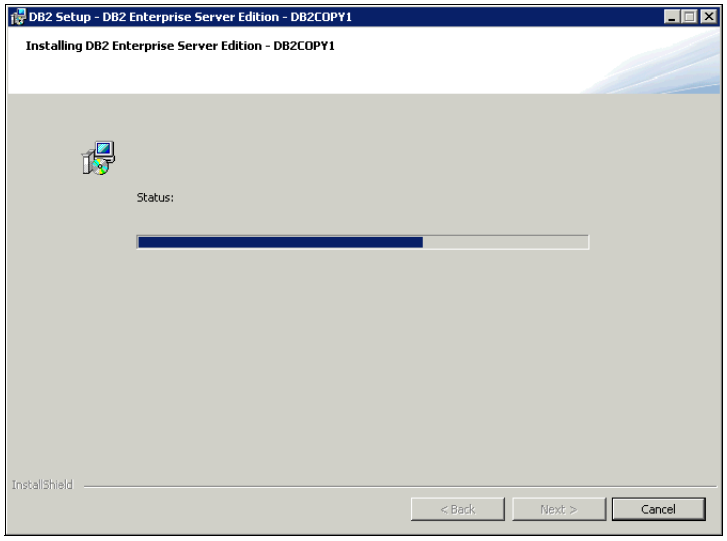

Figure 12-17 DB2 Enterprise Server Edition installation progress

16. When the setup completes, click **Finish** (Figure 12-18).

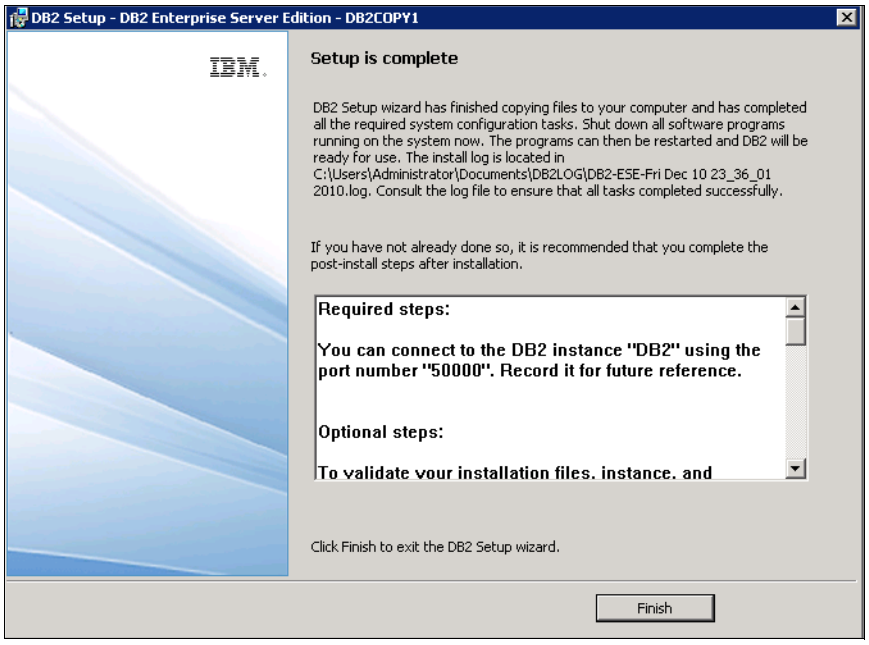

Figure 12-18 Setup is complete

17. Close the DB2 Setup Launchpad (Figure 12-19) to complete the installation.

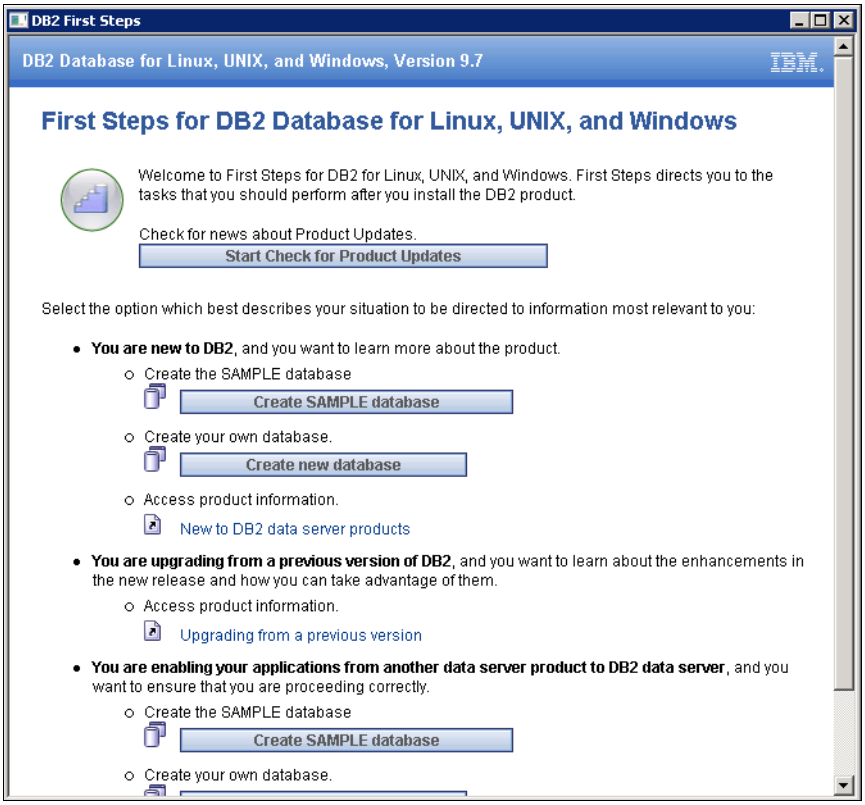

Figure 12-19 DB2 first steps

Verifying the DB2 installation

Complete the following steps to verify the DB2 installation:

1. Open a DB2 Command window by clicking **Start** → **IBM DB2** → **DB2COPY1 (Default)** → **Command Line Tools** → **Command Window**.
2. Create the SAMPLE database by running **db2sampl** (Figure 12-20).

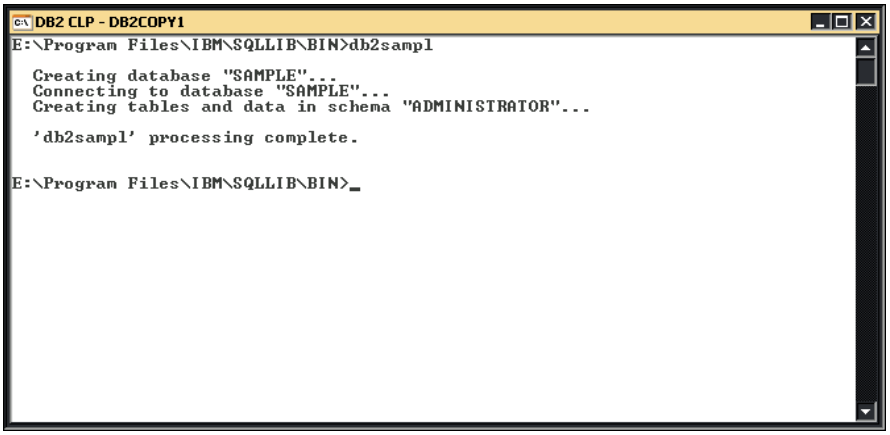

Figure 12-20 Create the SAMPLE database

3. Enter the following DB2 commands, then connect to the SAMPLE database, issue a simple SQL query, and reset the database connection.

```
db2 connect to sample
db2 "select * from staff where dept = 20"
db2 connect reset
```

The result of these commands is shown in Figure 12-21.

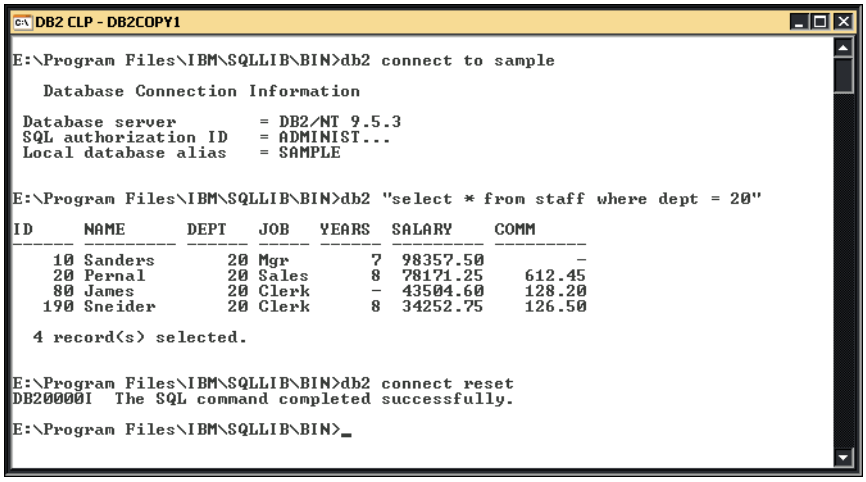

Figure 12-21 DB2 commands results

4. If you installed DB2 Enterprise Server 9.7, restart the system now.

Rebooting: If you do not reboot your server or at least the DB2 service after installing DB2 level 9.7, the Tivoli Storage Productivity Center installation fails. For more details, see the following address:

http://www-01.ibm.com/support/docview.wss?uid=swg21452614

12.4 Installing Tivoli Storage Productivity Center components

Make sure that all the prerequisites described in 12.3.3, “Pre-installation steps for Windows” on page 422 have been installed. If they have, you can continue to install the Tivoli Storage Productivity Center components, keeping in mind that with Tivoli Storage Productivity Center V4.2, both Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication are installed.

Before starting the installation, verify that a supported version of DB2 Enterprise Server Edition has been installed and it has been started.

Important: Log on to your system as a local administrator with database authority.

Complete the following steps:

1. If Windows autorun is enabled, the installation program starts automatically. If it does not, open Windows Explorer and go to the Tivoli Storage Productivity Center CD–ROM drive or directory. Double-click `setup.exe`.
2. Choose your language and click **OK** (Figure 12-22).

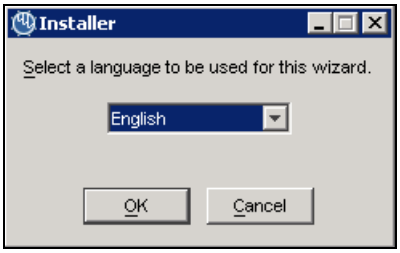

Figure 12-22 Language selection window

3. Read the terms of the license agreement, select **I accept**, and click **Next** (Figure 12-23).

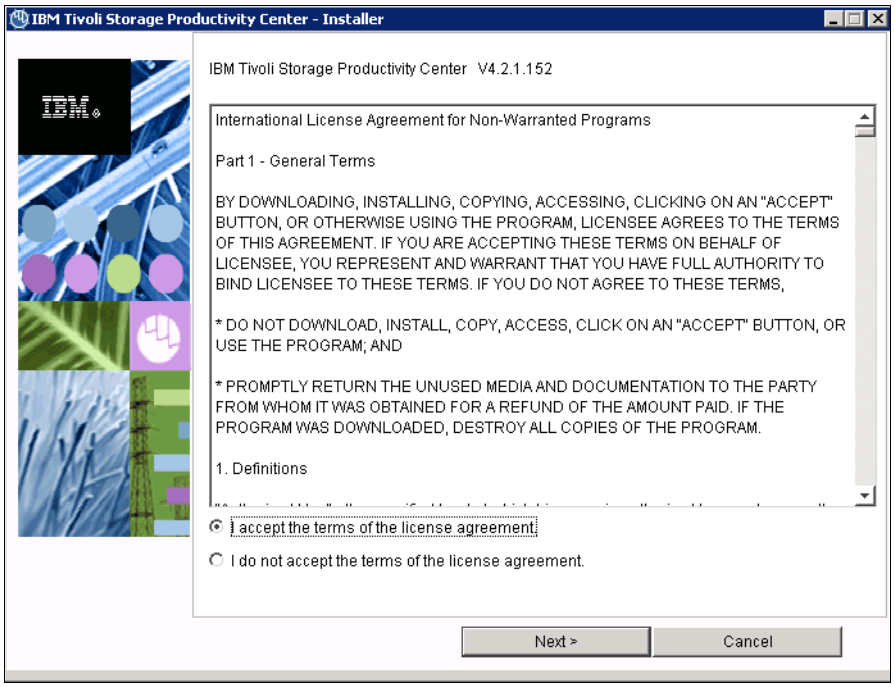

Figure 12-23 Accept the license agreement

4. Select **Typical installation** and make sure that you have selected all the check boxes (Figure 12-24).

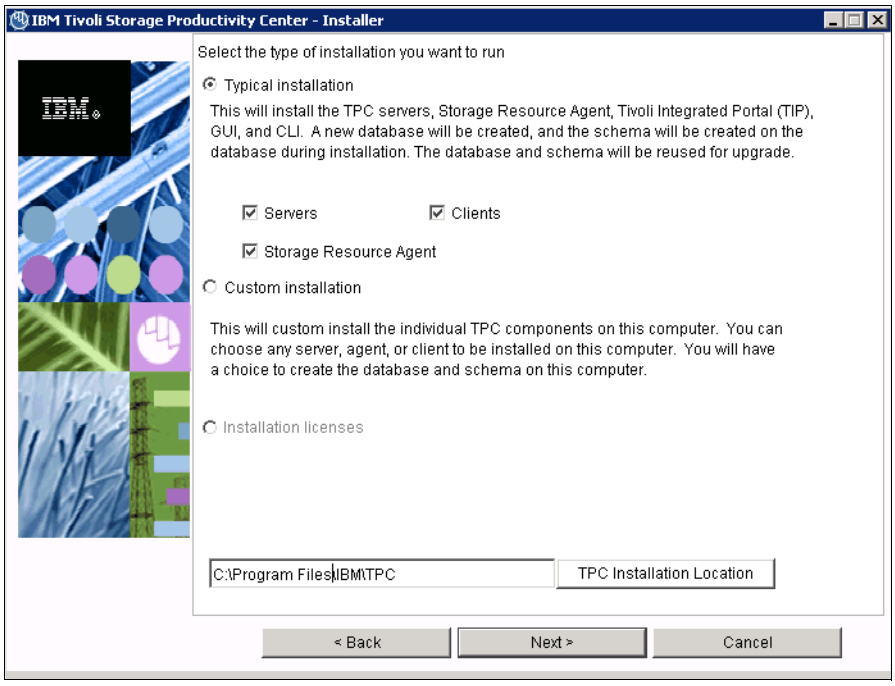

Figure 12-24 Typical Installation

5. Enter the user ID and the password you selected during DB2 installation (Figure 12-25).

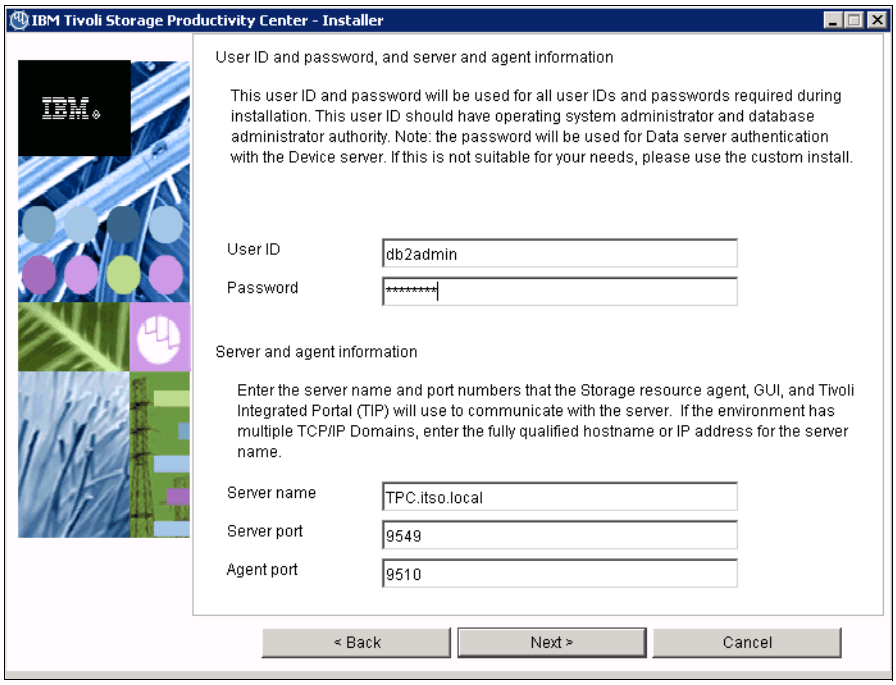

Figure 12-25 Enter a DB2 user

6. Specify the location to install Tivoli Integrated Portal (Figure 12-26). If you install Tivoli Storage Productivity Center on 64-bit Windows, the default path includes `Program Files (x86)`. Remove x86 or an error message is issued and you are not able to proceed.

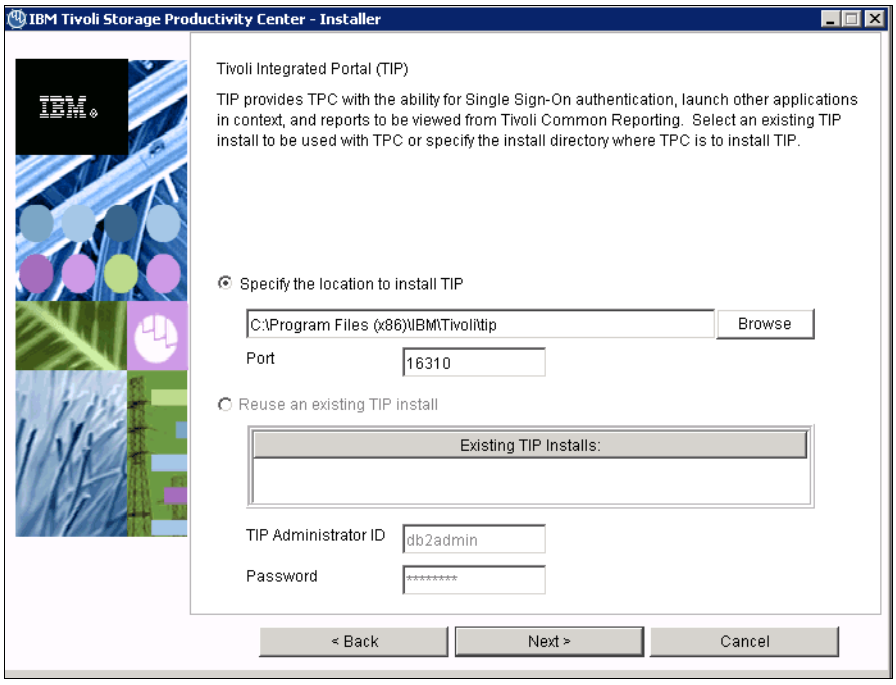

Figure 12-26 Select IBM Tivoli Integrated Portal location

7. Select your preferred authentication method and click **Next** (Figure 12-27).

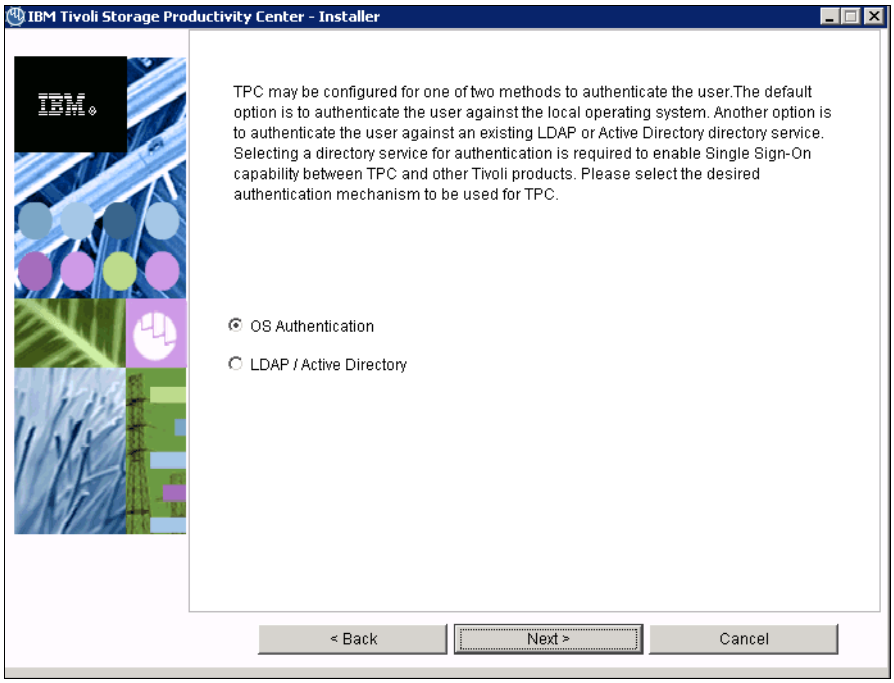

Figure 12-27 Select authentication method

8. Review the summary report and click **Install** to start the installation task (Figure 12-28).

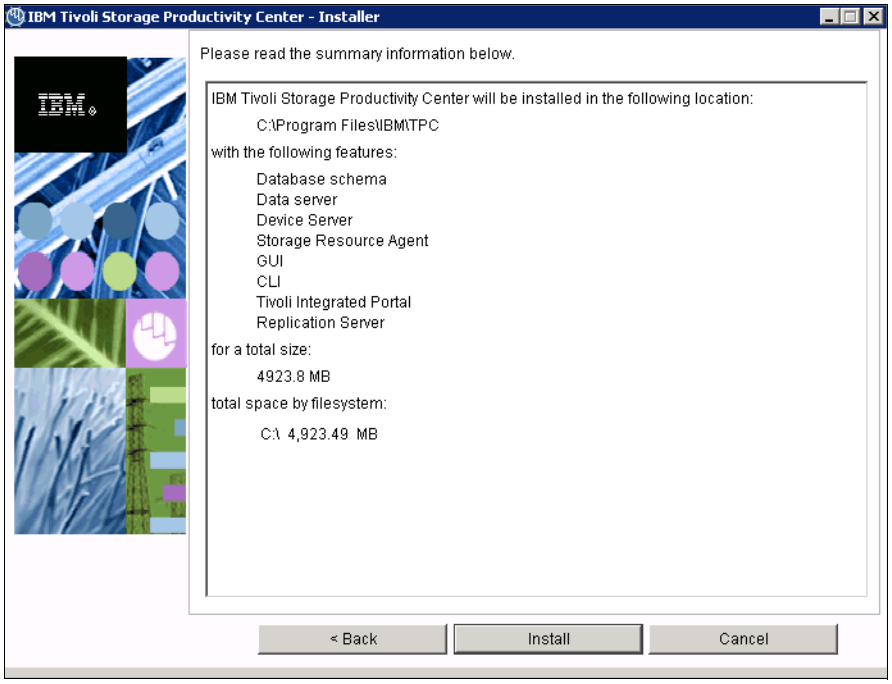

Figure 12-28 Summary information

9. The installation task now starts and takes some time. During the installation, you are prompted to select the Tivoli Storage Productivity Center for Replication settings (Figure 12-29). Click **Next**.

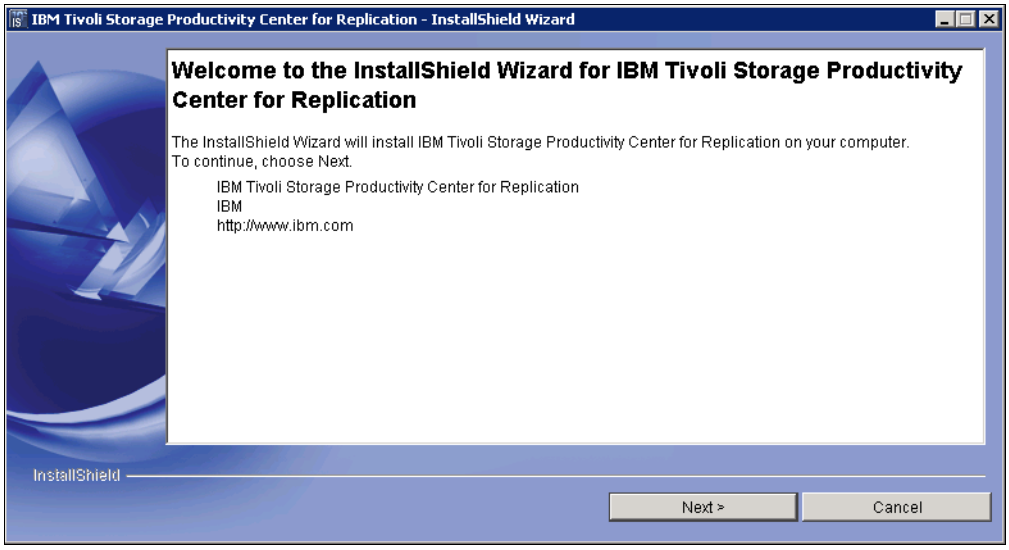

Figure 12-29 Tivoli Storage Productivity Center for Replication wizard

10. To perform a prerequisites check, click **Next** (Figure 12-30).

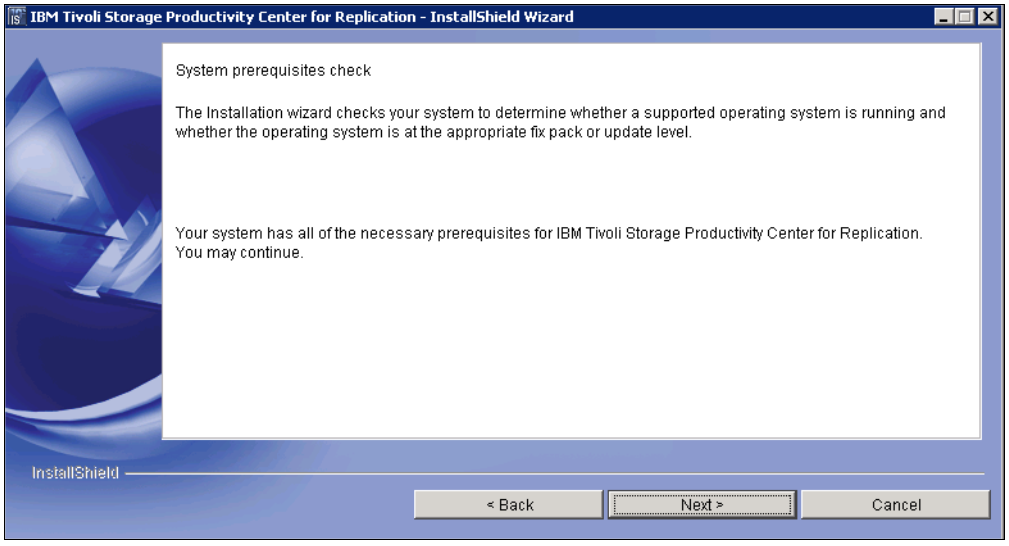

Figure 12-30 System prerequisites check

11. Select **I accept the terms of the license agreement** and click **Next** (Figure 12-31).

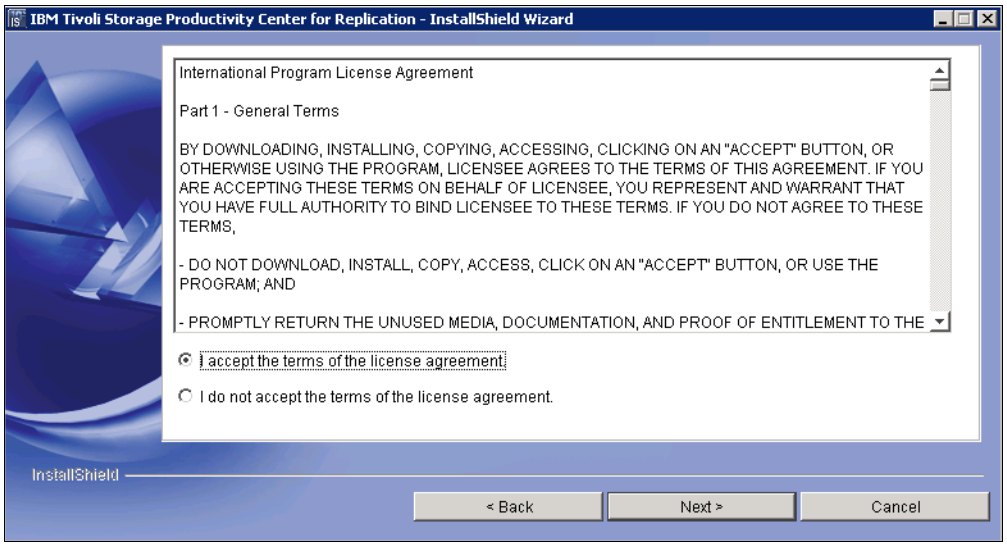

Figure 12-31 Accept the license agreement

12. Select an installation path and click **Next** (Figure 12-32).

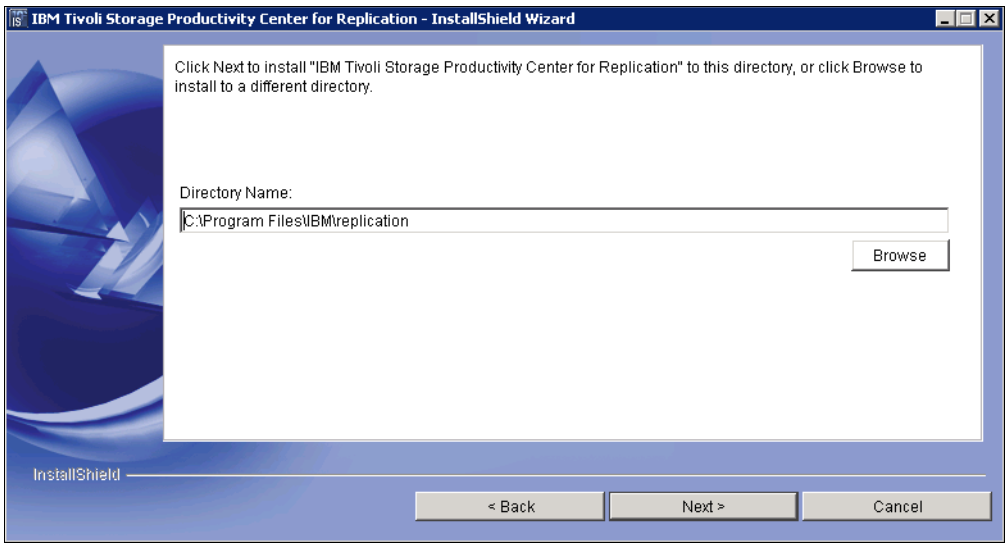

Figure 12-32 Select the installation path

13. Keep the default ports and click **Next** (Figure 12-33).

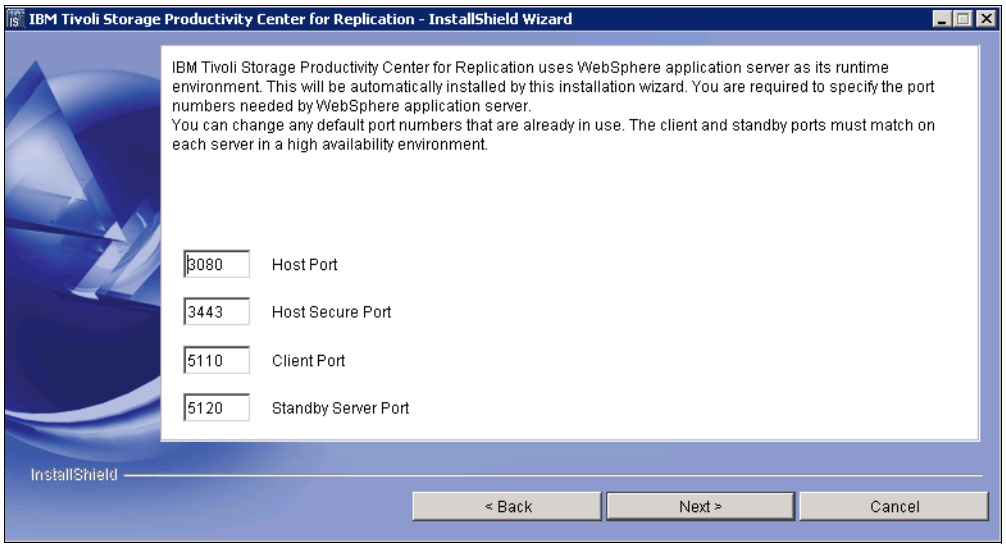

Figure 12-33 Select ports

14. Review the installation summary and click **Install** (Figure 12-34).

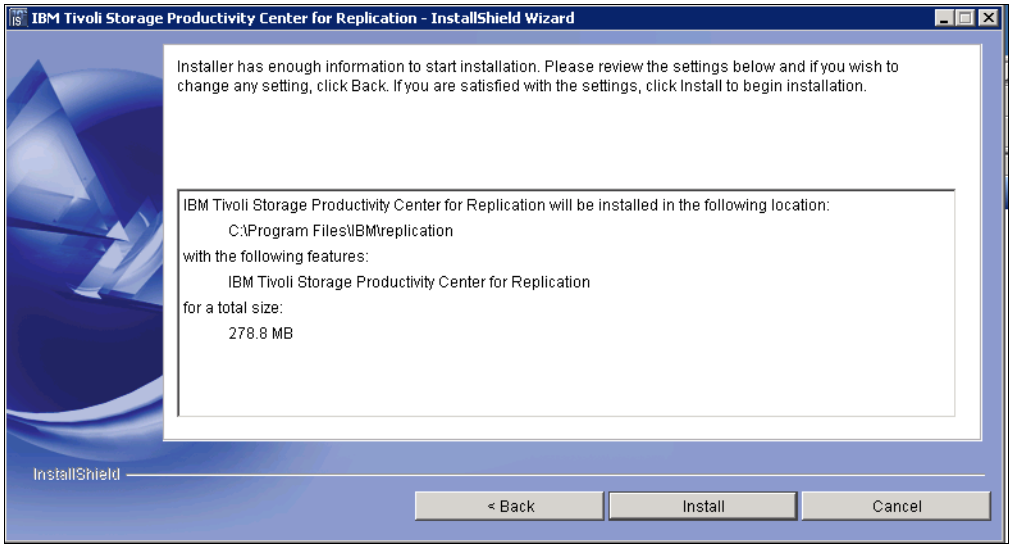

Figure 12-34 Tivoli Storage Productivity Center for Replication installation summary

The installation progress for Tivoli Storage Productivity Center for Replication is shown in Figure 12-35.

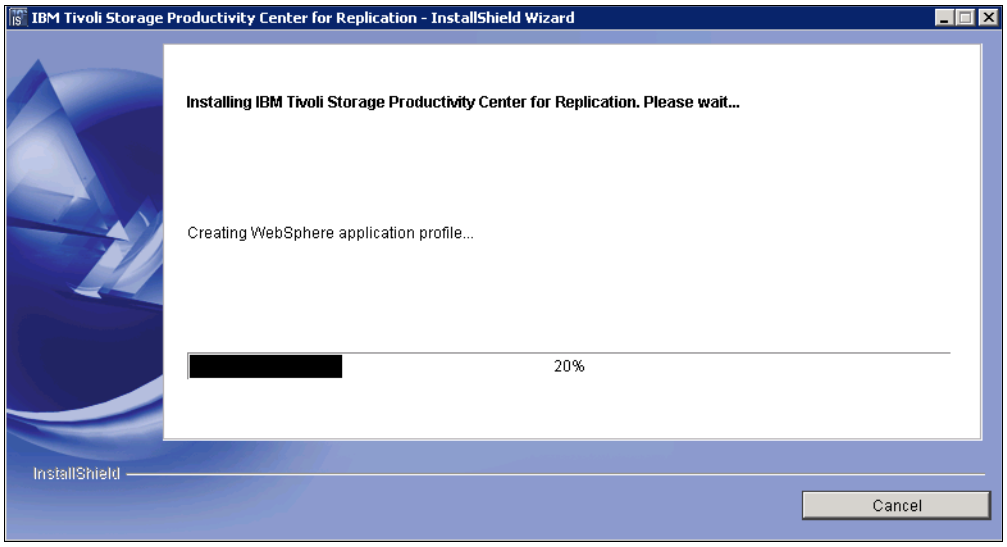

Figure 12-35 Installation progress

15. After the installation completes, click **Finish** (Figure 12-36).

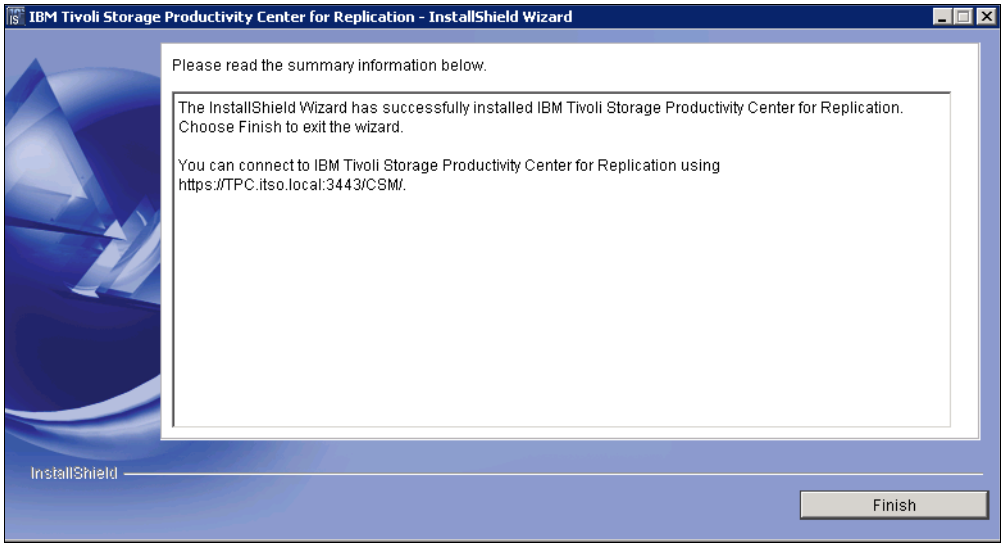

Figure 12-36 Tivoli Storage Productivity Center for Replication installation complete

The view returns to the Tivoli Storage Productivity Center installation (Figure 12-37).

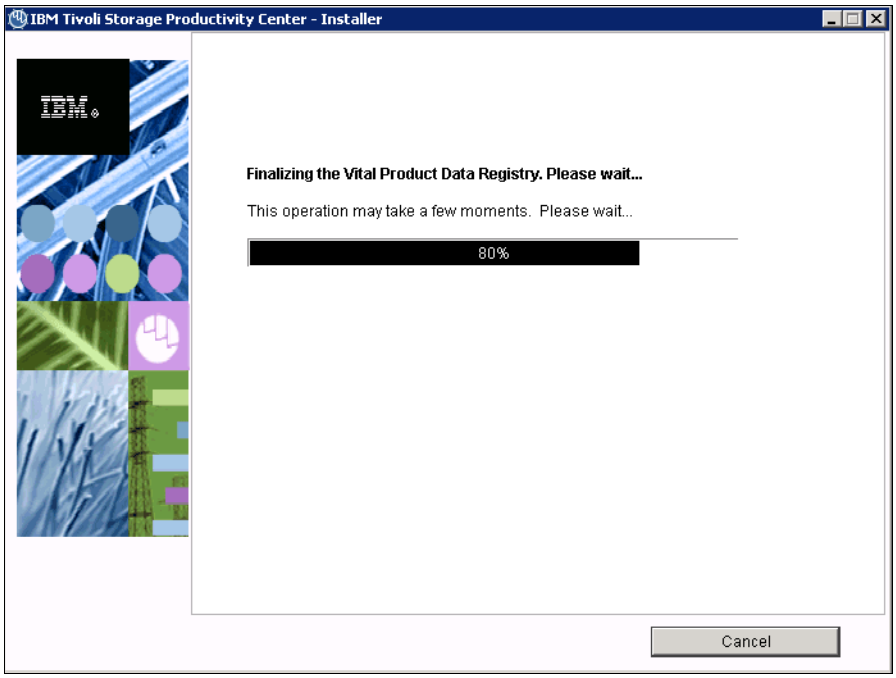

Figure 12-37 Tivoli Storage Productivity Center installation progress

16. After several minutes, the Tivoli Storage Productivity Center installation process completes. Click **Finish** to close the wizard (Figure 12-38).

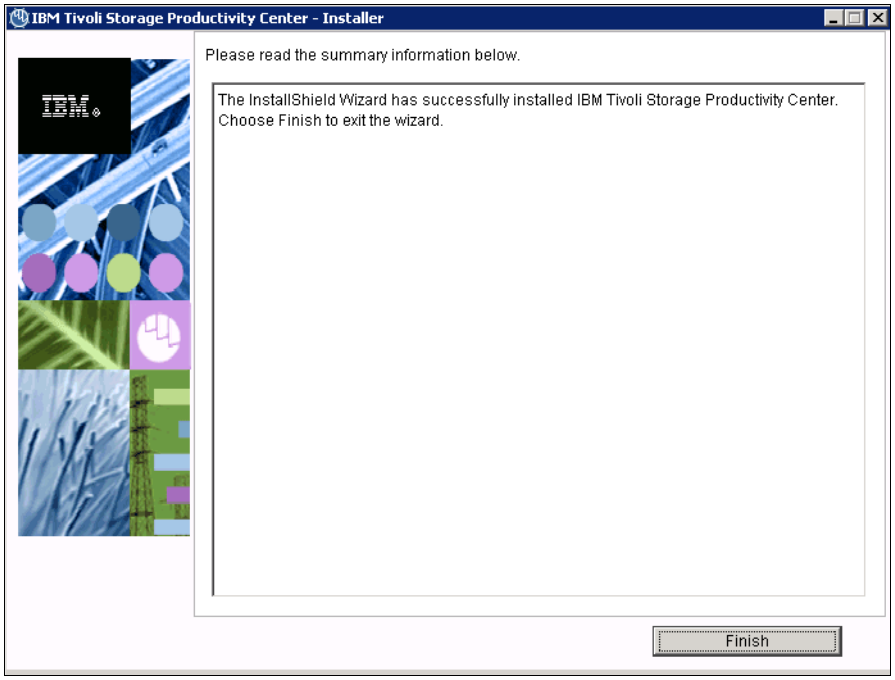

Figure 12-38 Tivoli Storage Productivity Center Typical installation wizard complete

12.5 Connecting Tivoli Storage Productivity Center to an IBM Storwize V7000 system

After the Typical installation, Tivoli Storage Productivity Center is ready to connect to the IBM Storwize V7000 system. Complete the following steps:

1. Start the Tivoli Storage Productivity Center GUI by clicking **Start** → **All Programs** → **IBM Tivoli Storage Productivity Center** → **> Productivity Center**. You are prompted to authenticate (Figure 12-39). Click **OK**.

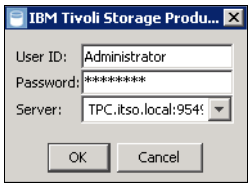

Figure 12-39 Authentication

2. The welcome window opens (Figure 12-40). Click **Add Devices** to connect to a system.

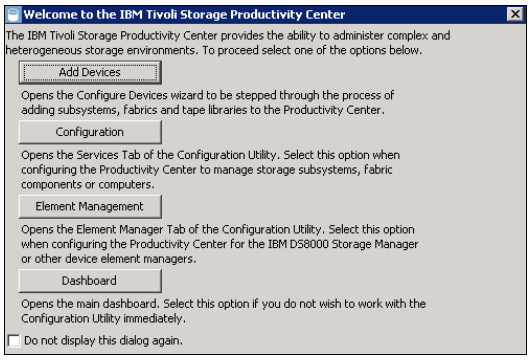

Figure 12-40 Welcome window

3. Select **Storage Subsystem** and click **Next** (Figure 12-41).

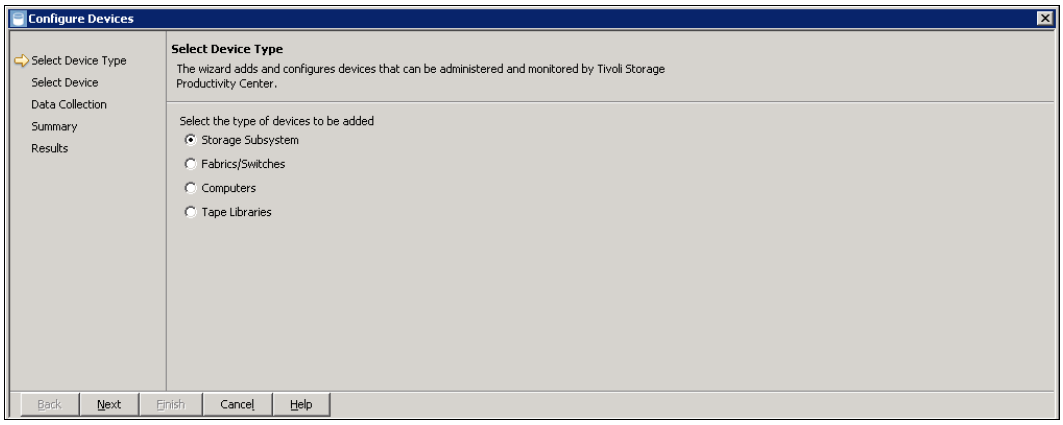

Figure 12-41 Select Device Type

4. Select **Add and configure new storage subsystems** and click **Next** (Figure 12-42).

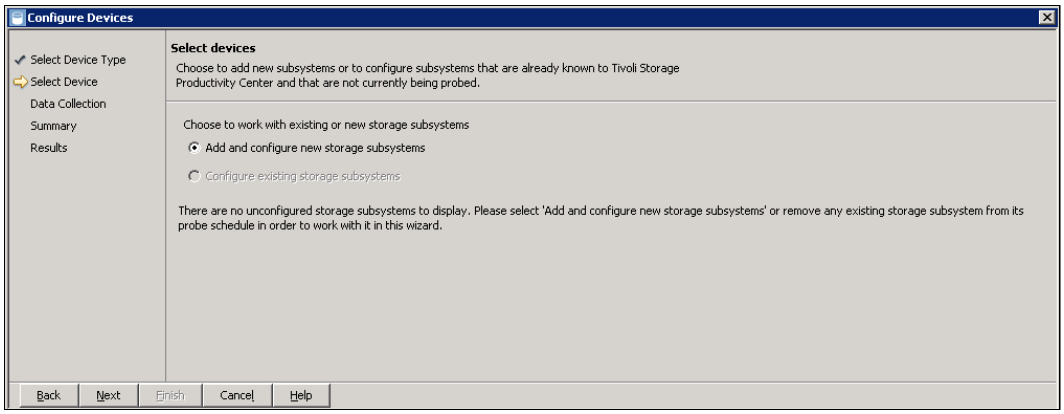

Figure 12-42 Add and configure new storage subsystems

5. Select or enter the following values to connect to the IBM Storwize V7000 system (shown in Figure 12-43):
   - Device Type: IBM SAN Volume Controller / IBM Storwize V7000.
   - Software Version: 5+
   - IP Address: Specify the IBM Storwize V7000 system IP address.
   - Select Key: Upload new key
   - Admin Username: superuser
   - Admin Password: Enter the superuser password.
   - Username: Select a user to connect.
   - Private SSH Key: Specify the location of a private key. Because you are authenticated with the user name and password, this key does not really have to be associated with the selected user.

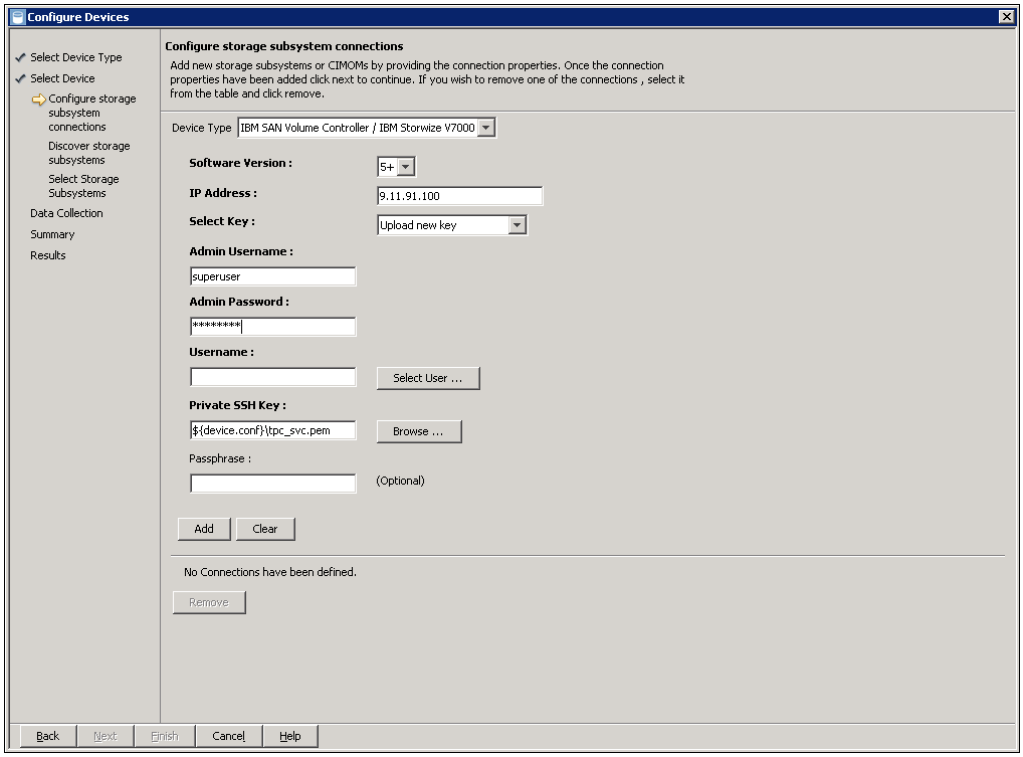

Figure 12-43 Configure storage subsystem connections

6. Click **Select User** to retrieve the list of users (Figure 12-44). For this example, choose the **monitor** user. Click **OK**.

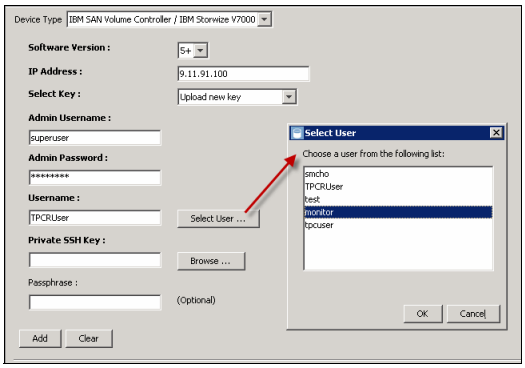

Figure 12-44 Select User window

7. Click **Add** to connect to your IBM Storwize V7000 system (Figure 12-45).

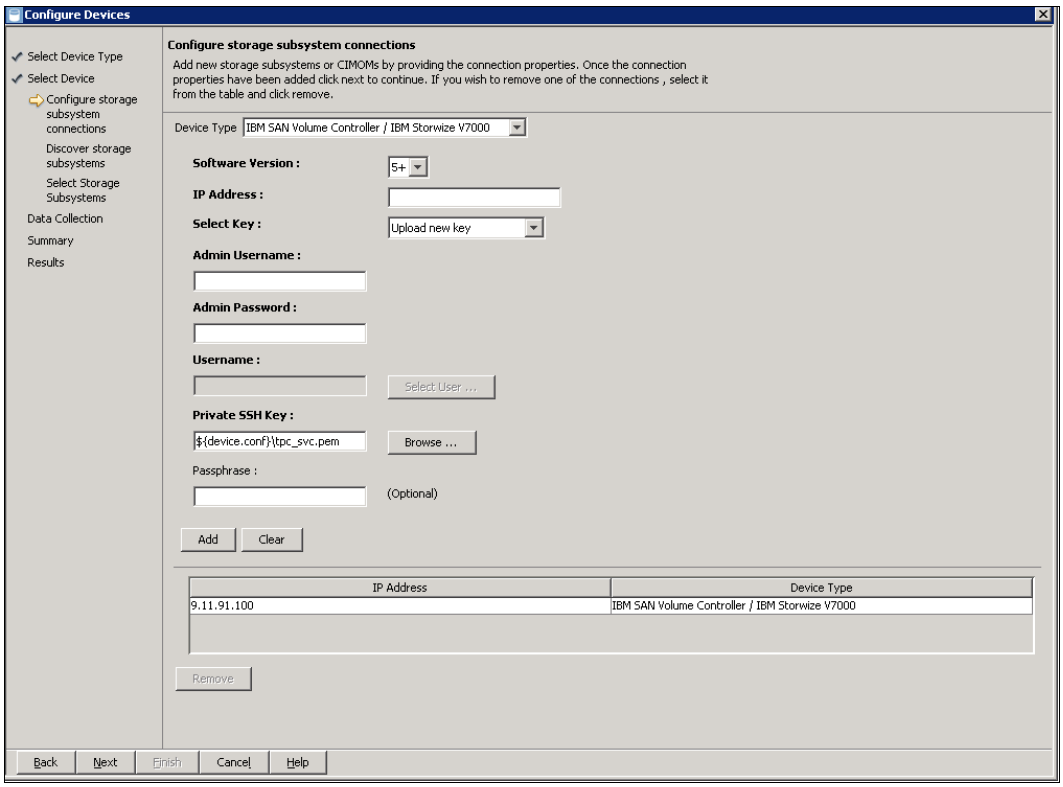

Figure 12-45 Configure storage subsystem connections: System discovered

8. Repeat step 2 on page 443 through step 7 on page 445 to add another system, or click **Next** to complete the discovery of the new storage subsystems (Figure 12-46).

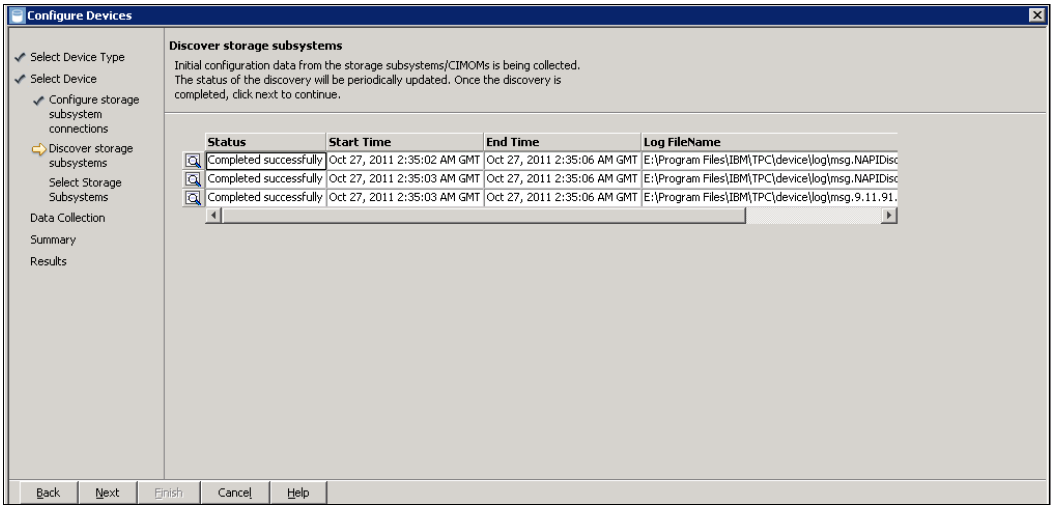

Figure 12-46 Discover subsystems

9. After the discovery, select the new storage subsystem that you want to add to the configuration (Figure 12-47).

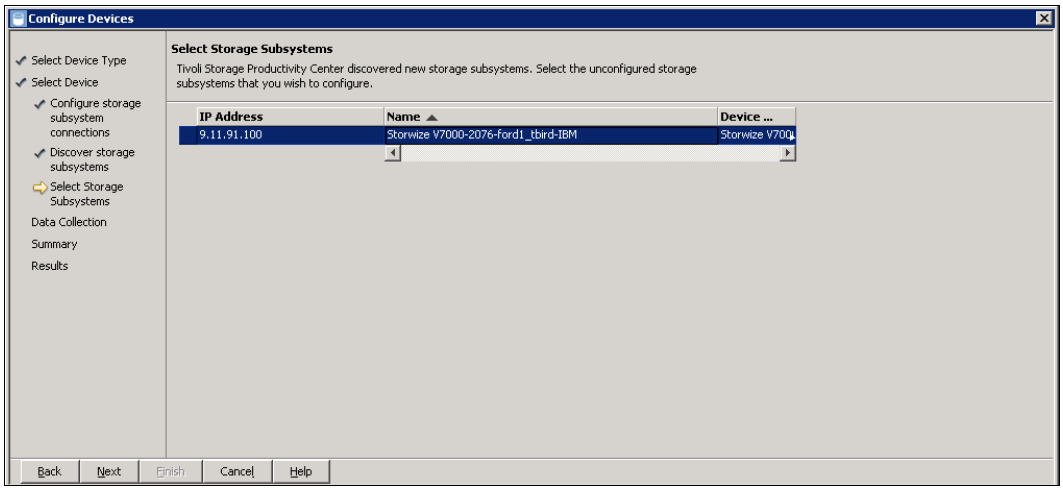

Figure 12-47 Select storage subsystem to connect

10. Specify the data collection settings. To add the new system to the default group, select **Monitoring Group** and **Subsystem Standard Group** (Figure 12-48).

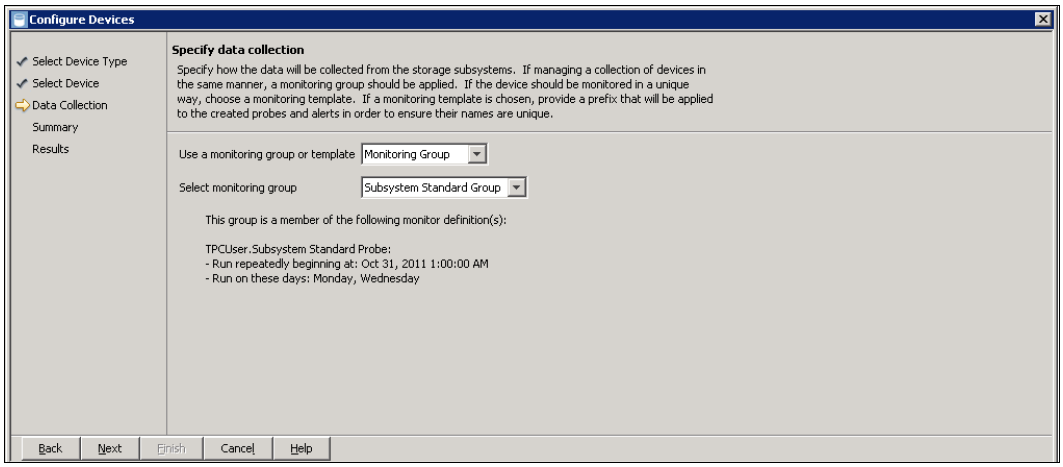

Figure 12-48 Specify data collection settings

11. Review your selections and click **Next** (Figure 12-49) to add the new device.

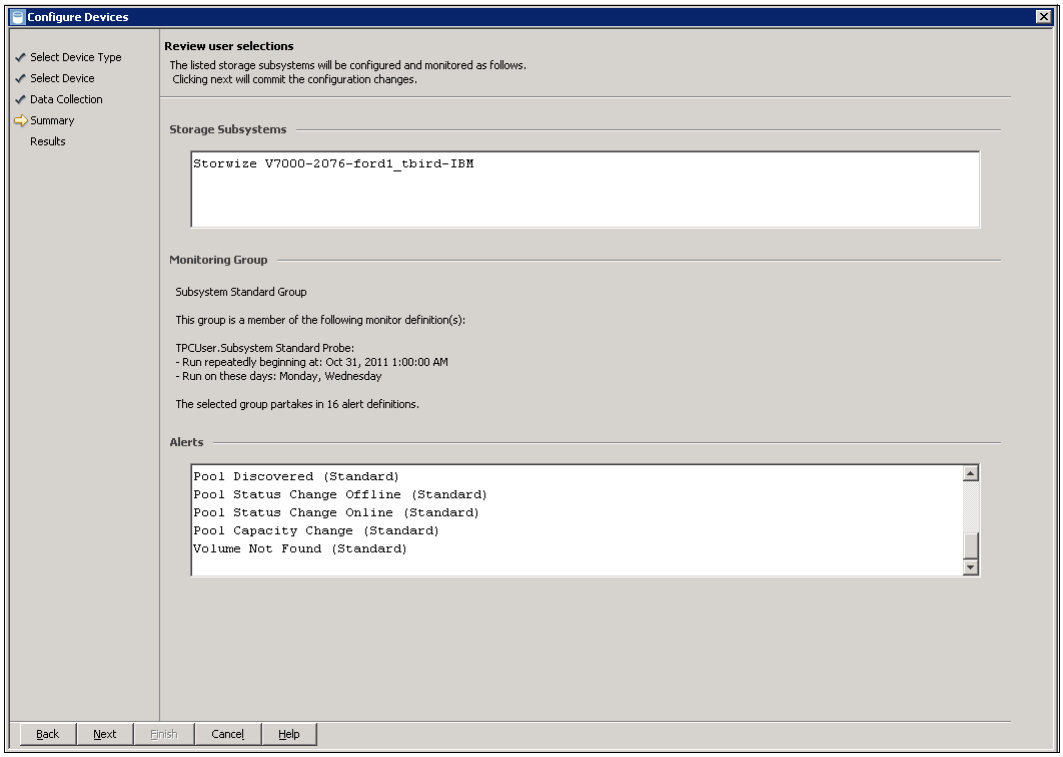

Figure 12-49 Review user selections

12. The IBM Storwize V7000 system has now been added to Tivoli Storage Productivity Center. Click **Finish** to close the wizard (Figure 12-50).

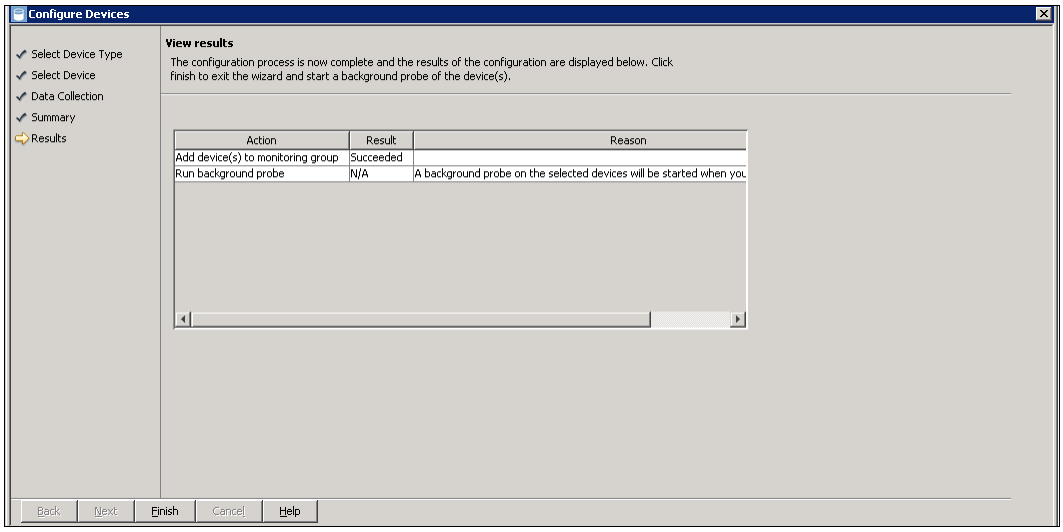

Figure 12-50 View results

13. You are prompted whether you would like to view the job history, but you do not need to view it at this time. Click **Close** (Figure 12-51).

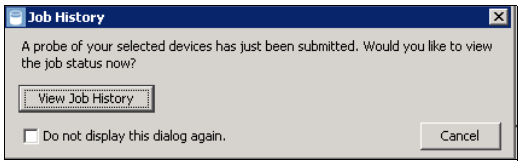

Figure 12-51 Job History

The IBM Storwize V7000 system has now been added successfully and can be administered by Tivoli Storage Productivity Center. The normal IBM Storwize V7000 GUI and CLI are still available, and can be used to manage the system also.

12.6 Administering and reporting an IBM Storwize V7000 system through Tivoli Storage Productivity Center

This section shows examples of how to use Tivoli Storage Productivity Center to administer, configure, and generate reports for your IBM Storwize V7000 system. A detailed description about Tivoli Storage Productivity Center reporting is beyond the intended scope of this book.

12.6.1 Basic configuration and administration

Tivoli Storage Productivity Center allows you to administer and configure your IBM Storwize V7000 system, but not all of the options normally associated with the IBM Storwize V7000 GUI or CLI are available.

After successfully adding your IBM Storwize V7000 system, click **Disk Manager** → **Storage Subsystems** to view your configured devices (Figure 12-52).

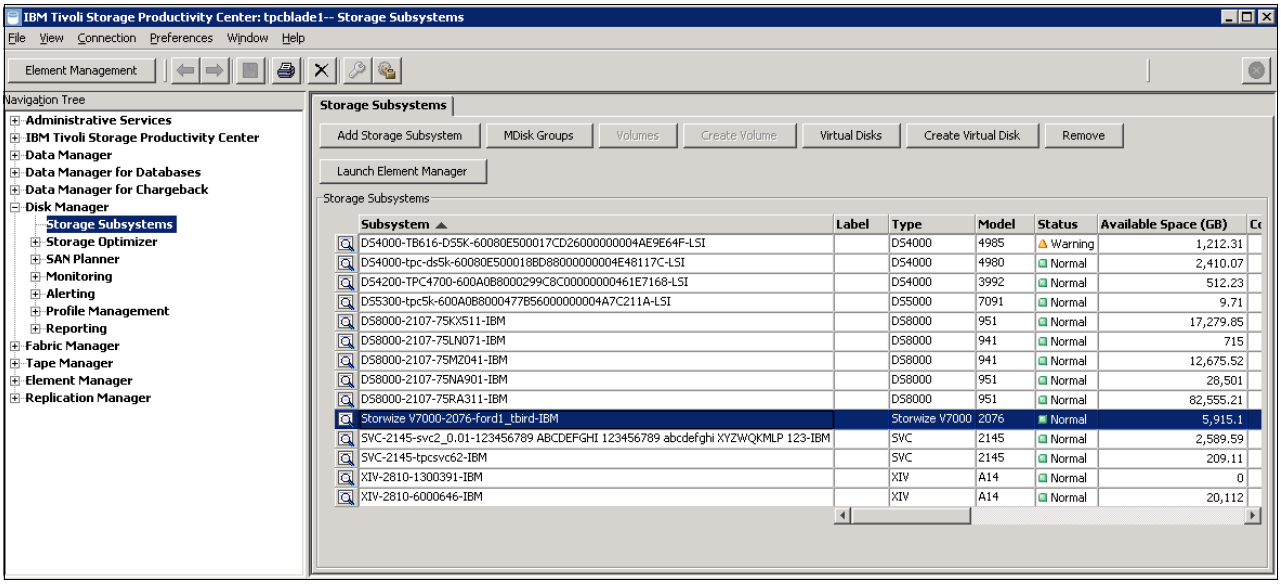

Figure 12-52 Storage Subsystems view

When you highlight the IBM Storwize V7000 system, action buttons become available so you can view the device configuration or create virtual disks (Figure 12-53).

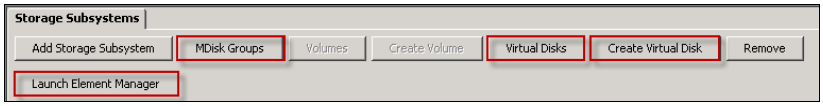

Figure 12-53 Action buttons

- MDisk Groups provides a detailed list of the configured MDisk groups including, pool space, available space, configured space, and Easy Tier configuration.
- Virtual Disks lists all the configured ones with the added option to filter them by MDisk Group. The list includes several attributes, such as capacity, volume type, and type.
- Create Virtual Disk opens a wizard to create volumes and add MDisks to a group.
- Launch Element Manager opens the IBM Storwize V7000 GUI from your default browser.

If you click **Create Virtual Disk**, the Create Virtual Disk wizard window opens (Figure 12-54). Use this window to create volumes, specifying several options, such as size, name, thin provisioning, and even add MDisks to an MDisk Group.

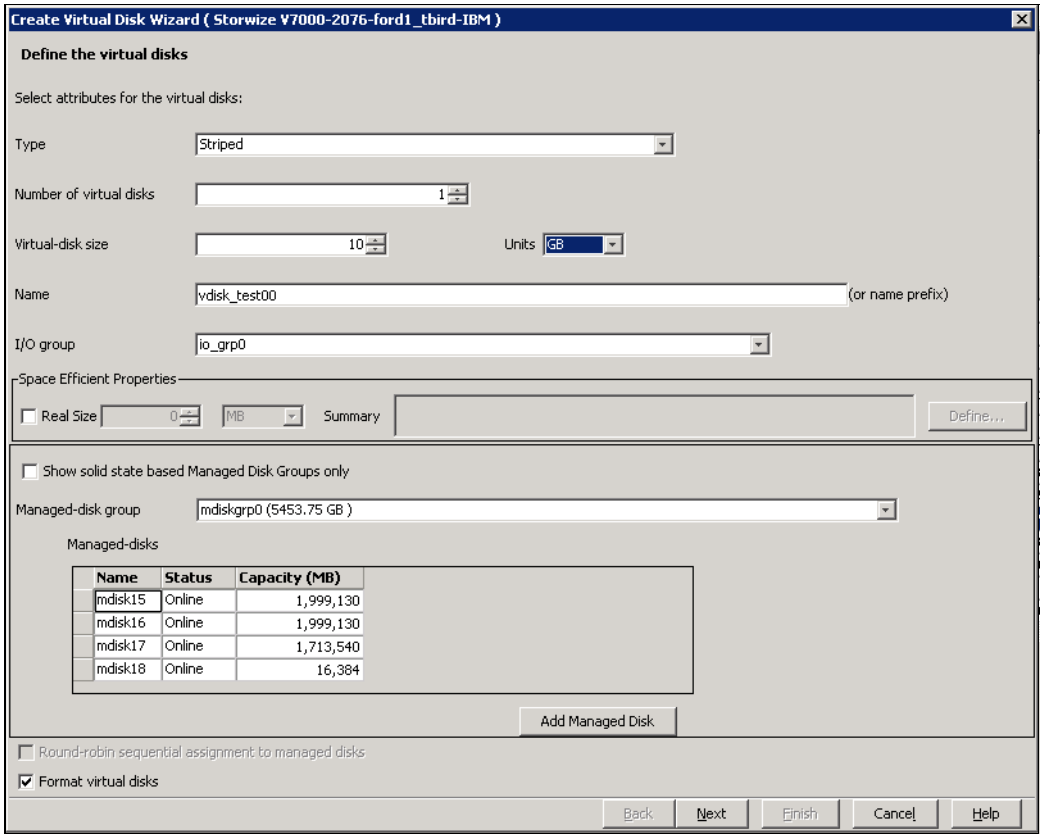

Figure 12-54 Virtual Disk creation wizard

12.6.2 Report generation

In this section, we show some basic report generation. Click **Disk Manager** → **Reporting** → **Storage Subsystems** to view the available options to select various reports (Figure 12-55).

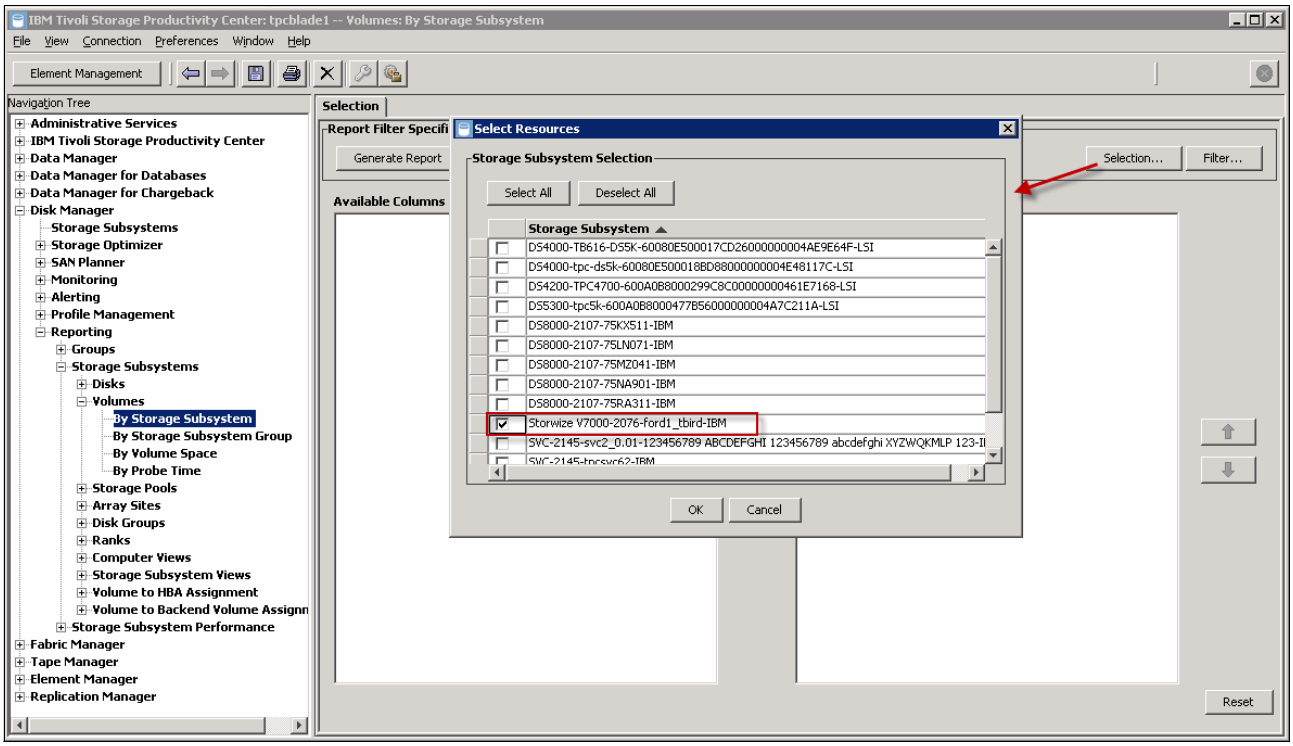

Figure 12-55 V7000 volumes report generation

If you highlight a report in the left pane, the associated view appears in the right pane. In Figure 12-55, we selected to report on Volumes filtered by Storage Subsystem. To report only on IBM Storwize V7000 volumes, click **Selection** and select the check box next to the IBM Storwize V7000 system.

Figure 12-56 shows the generated report; Figure 12-57 shows the configured storage pools report.

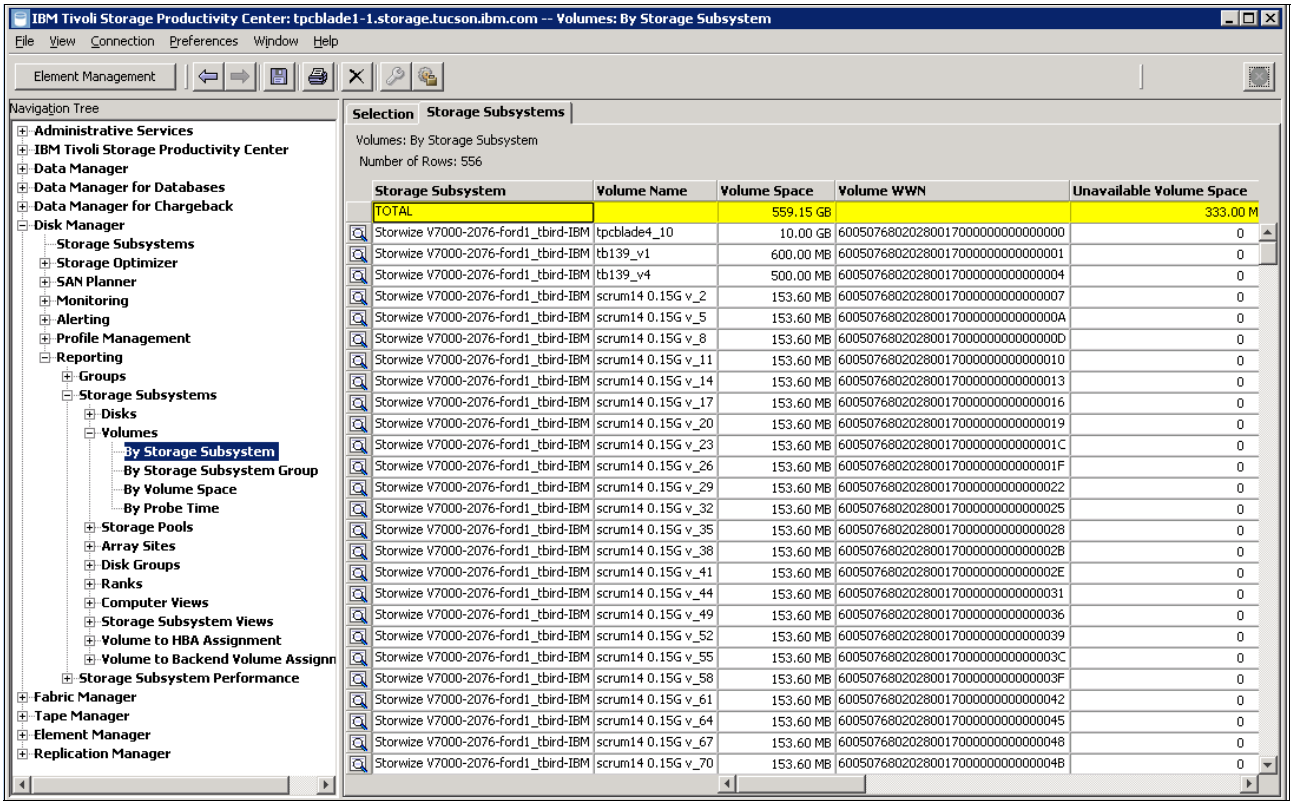

Figure 12-56 Volumes by storage subsystem

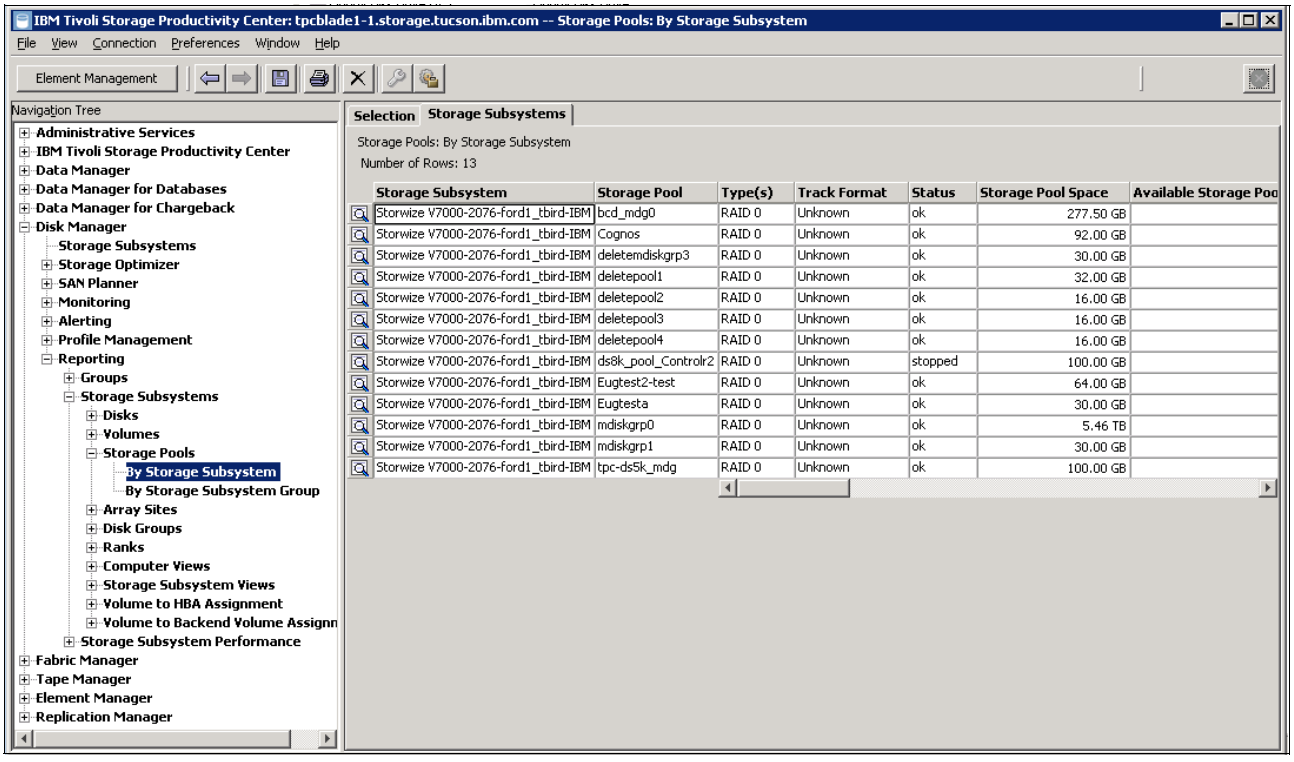

Figure 12-57 Storage pools by storage subsystem

In addition to being able to monitor configuration data, there are many more options, such as performance statistics, health monitoring, path usage, and so on.

To monitor performance, create a subsystem performance monitor by clicking **Disk Manager** → **Monitoring**, right-clicking **Subsystem Performance Monitors**, and click **Create Subsystem Performance Monitor**.

Figure 12-58 shows you the statistics for a single disk pool.

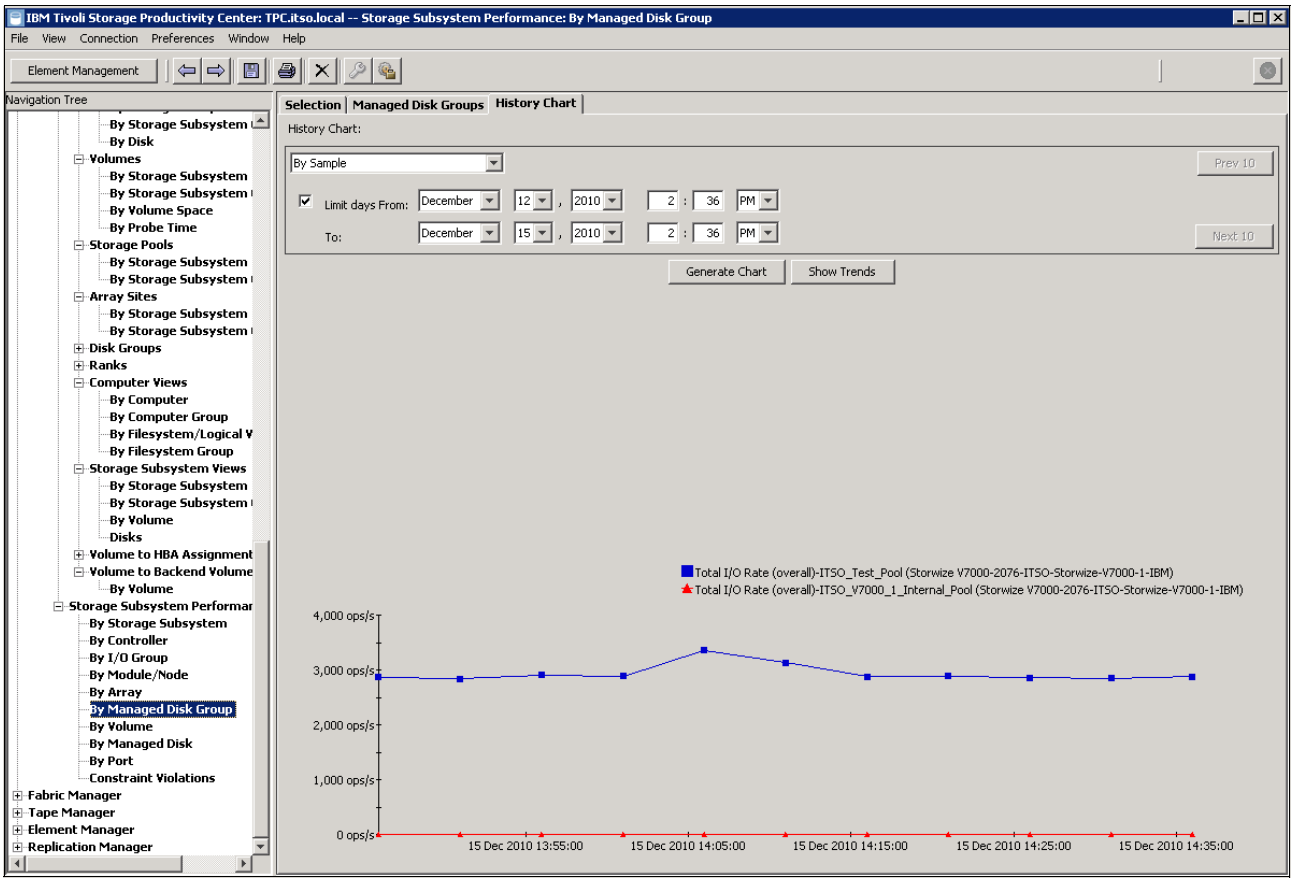

Figure 12-58 Storage subsystem performance by managed disk group (storage pool): Total I/O rate

13

RAS, monitoring, and troubleshooting

There are many ways to monitor and troubleshoot the IBM Storwize V7000 system. This chapter describes the following topics:

- Reliability, availability, and serviceability (RAS)
- Hardware and LED descriptions
- Monitoring from a host
- Monitoring from the IBM Storwize V7000 system
- Backup procedure
- Software upgrade procedure
- Troubleshooting
- Recommended actions
- Event log navigation
- Audit log navigation
- Support
- Shutting down the IBM Storwize V7000 system
- Shutting down the infrastructure using the IBM Storwize V7000 system

13.1 Reliability, availability, and serviceability (RAS)

Reliability, availability, and serviceability are important concepts in the design of the IBM Storwize V7000 system. Hardware features, software features, design considerations, and operational guidelines all contribute to make the IBM Storwize V7000 system reliable.

Fault tolerance and high levels of availability are achieved by these methods:

- The RAID capabilities of the underlying disk subsystems
- IBM Storwize V7000 nodes clustering using a *Compass* architecture
- Auto-restart of hung nodes
- UPS units to provide memory protection in the event of a site power failure
- Host system failover capabilities

High levels of serviceability are available through these methods:

- 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 V7000 components
- Concurrent upgrade of IBM Storwize V7000 Software and microcode
- Concurrent addition or deletion of node canisters in a clustered system
- Software recovery through a service panel push button
- Automatic software version correction when replacing a node
- Detailed status and error conditions displayed on the service panel
- Error and event notification through SNMP, syslog, and email

At the heart of the IBM Storwize V7000 system is a pair of *node canisters*. These two canisters share the data transmitting and receiving load between the attached hosts and the disk arrays. This section examines the RAS features of the IBM Storwize V7000 system, monitoring, and troubleshooting.

13.1.1 Node canisters

The two node canisters are in the control enclosure and they work as a clustered system. Figure 13-1 shows the ports and indicator lights of a node canister. The second canister is placed below the first one in an upside down position.

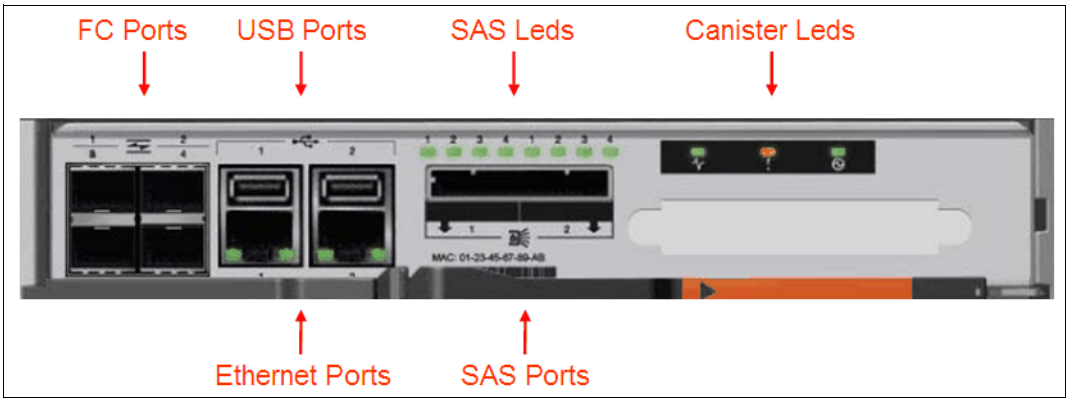

Figure 13-1 Ports and indicators of a node canister (112 or 124 controller models)

Fibre Channel

Four Fibre Channel (FC) ports are on the left side of the canister. They are in two rows of two ports each. The ports are numbered 1 - 4 from left to right, top to bottom. The ports operate at 2, 4, or 8 Gbps. Each port has two green LEDs associated with it. These LEDs are not shown in the figure, but are located between the two rows of ports and are triangular, pointing toward the port to which they refer. The left LED indicates speed; the right LED indicates link status. Table 13-1 describes the status of the indicators.

Table 13-1 Fibre Channel link LEDs statuses

Speed LED (left)	Link LED (right)	Link status
Off	Off	Inactive
Off	On or flashing	Idle or busy at 2 Gbps
Flashing	On or flashing	Idle or busy at 4 Gbps
On	On or flashing	Idle or busy at 8 Gbps

Left and right: As the bottom canister is placed upside down, do not misinterpret the status LEDs on it. Left becomes right and right becomes left.

USB

Two active USB connectors are side by side; they are numbered as 1 on the left and 2 on the right. No indicators are associated with the USB ports.

Ethernet and LED status

Two 10/100/1000 Mbps Ethernet ports are side by side on the canister; they are numbered as 1 on the left and 2 on the right. Each port has two LEDs and their status values are shown in Table 13-2.

Table 13-2 Ethernet LEDs statuses

LED	Color	Meaning
Link state	Green	It is on when there is an Ethernet link.
Activity	Orange	It is flashing when there is activity on the link.

Serial Attached SCSI ports

Two 6 Gbps Serial Attached SCSI (SAS) ports are side by side on the canister with indicator LEDs on top of them. They are numbered 1 on the left and 2 on the right. Each port connects four *PHYceivers* (a physical transceiver that operates at the physical layer); each PHY is associated with an LED. These LEDs are green and are directly above the ports. For each port, the LEDs are numbered 1 - 4. The LED indicates activity on the link; the status values are listed in Table 13-3.

Table 13-3 SAS LED statuses

Status	Meaning
On	Link is connected.
Flashing	Link is connected and being used for I/O.
Off	Link is not connected.

Node canister status LEDs

There are three LEDs in a row at the upper right of the canister that indicate the status and identification for the node (Table 13-4).

Table 13-4 Node canister LED statuses

Position	Color	Name	State	Meaning
Left	Green	Cluster status	On	The node is in the active or starting state. It might not be safe to remove the canister. If the fault LED is off, then the node is an active member of a cluster. If the fault LED is also on, there is a problem establishing a cluster, for example because of a lack of quorum.
			Flashing	The node is in a candidate or service state. It is safe to remove the canister.
			Off	There is no power to the canister or it is in a state where IBM Storwize V7000 code is not running, for example, standby mode or power-on self-test (POST).
Middle	Amber	Fault	On	The node is supposed to be in a cluster but is unable to reach an active state.
			Flashing	The canister is being identified. A fault might exist.
			Off	The node is in a candidate or active state. This status does *not* mean that a hardware error does not exist on the node.
Right	Green	Power	On	The canister is powered on and processors are running.
			Fast flashing	The canister is running a POST.
			Slow flashing	AC is available, but the canister is in the standby mode that follows a request by the software to power off.
			Off	AC is not available.

Expansion canister status LEDs

As Figure 13-2 shows, two 6 Gbps SAS ports are side by side on the canister of every enclosure. They are numbered 1 on the left and 2 on the right. Each port connects four PHYs; each PHY is associated with an LED. These LEDs are green and are next to the ports.

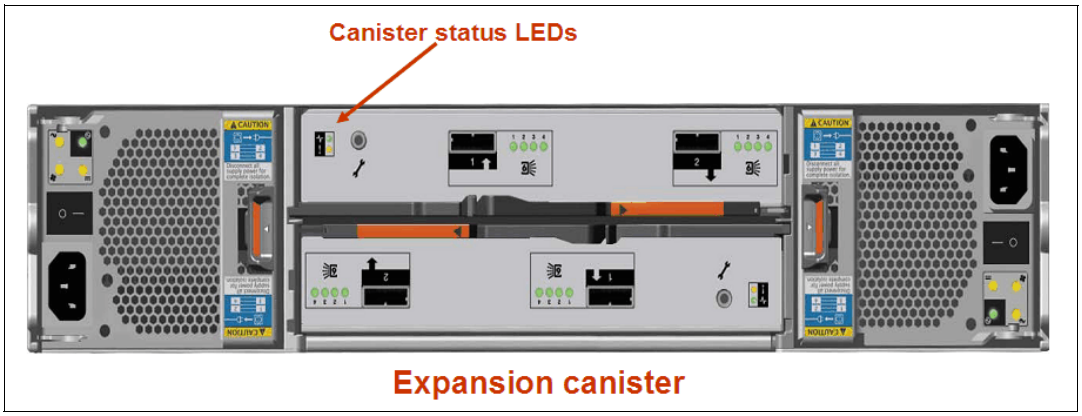

Figure 13-2 Expansion canister status LEDs

Table 13-5 shows the LED state values of the expansion canister.

Table 13-5 Expansion canister LEDs statuses

Position	Color	Name	State	Meaning
Top	Green	Status	On	The canister is active.
			Flashing	The canister has a vpd error.
			Off	The canister is not active.
Bottom	Amber	Fault	On	The canister hardware is faulty.
			Flashing	The canister is being identified.
			Off	No fault canister is being identified.

13.1.2 Disk subsystem

The IBM Storwize V7000 system consists of disk enclosures. Two types of enclosures are available:

- *2U12* that supports 12 3.5-inch drives
- *2U24* that supports 24 2.5-inch drives

The drives fit into the front of the enclosure, and the rear of the enclosures are identical and have slots for two canisters and two power supplies. Enclosures are used as either control enclosures or expansion enclosures. They are differentiated by the type of canister and power supply they contain.

An array is a type of MDisk made up of disk drives that are in the enclosures. 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 on the number of members in the array. An IBM Storwize V7000 system 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 location and performance of each array member. A sequence of drive failures and hot spare takeovers can leave an array unbalanced, that is, with members that do not match these goals. The system automatically rebalances such arrays when appropriate drives are available.

An IBM Storwize V7000 system supports the RAID levels listed in Table 13-6.

Table 13-6 RAID levels that are supported by IBM Storwize V7000 system

RAID level	Where data is striped	Minimum to maximum members
0	Data is striped on one or more drives.	1 - 8
1	Data is mirrored between two drives.	2
5	Data is striped across several drives with one parity.	3 - 16
6	Data is striped across several drives with two parities.	5 - 16
10	Data is striped across pairs of mirrored drives.	2 - 16

Disk scrubbing

The scrub process runs when arrays do not have any other background processes. The process checks that the drive LBAs are readable and array parity is in synchronization. Arrays are scrubbed independently and each array is entirely scrubbed every seven days.

Solid-state drives

The IBM Storwize V7000 system treats solid-state drives (SSDs) the same as spinning hard disk drives (HDDs), with respect to RAID arrays or MDisks. The individual SSD drives in the storage managed by the IBM Storwize V7000 are combined into an array, usually in RAID 10 or RAID 5 format. It is unlikely that RAID 6 SSD arrays are used because of the double parity impact, with two SSD logical drives used for parity only.

A LUN is created on the array, which is then presented to the IBM Storwize V7000 as a normal managed disk (MDisk). As is the case for HDDs, the SSD RAID array format helps protect against individual SSD failures. Depending on your requirements, additional high availability protection, above the RAID level, can be achieved by using volume mirroring.

SAS cabling

Expansion enclosures are attached to control enclosures using SAS cables (Figure 13-3). The SAS network is made up of strands and chains.

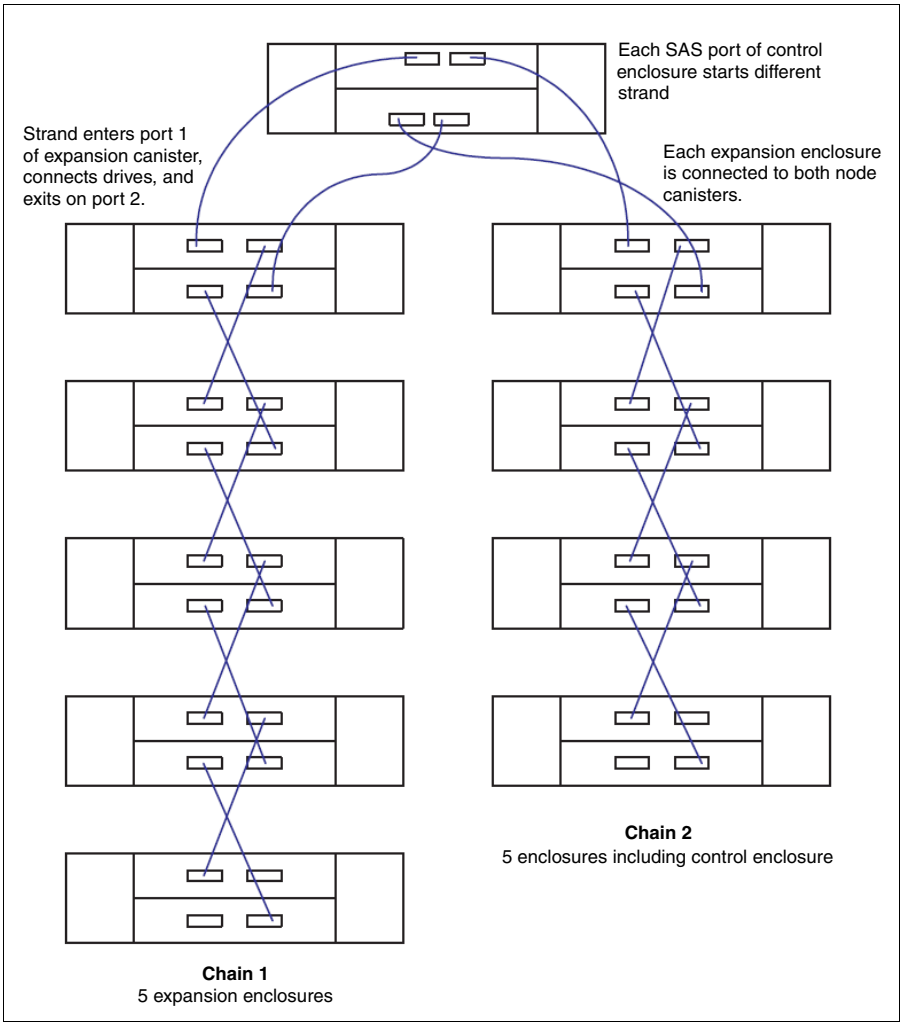

Figure 13-3 SAS cabling

A strand starts with an SAS initiator chip inside an IBM Storwize V7000 node canister and progresses through SAS expanders, which connect disk drives; each canister contains an expander. Figure 13-3 on page 460 shows how the SAS connectivity works inside the node and expansion canisters. Each drive has two ports, each connected to a different expander and strand. This configuration assures that both nodes in the I/O group have direct access to each drive and there is no single point of failure.

A chain consists of a set of enclosures, correctly interconnected (Figure 13-3 on page 460). Chain 1 of an I/O group is connected to SAS port 1 of both node canisters; chain 2 is connected to SAS port 2. This configuration means that chain 2 includes the expander and drives of the control enclosure.

At system initialization, when devices are added to or removed from strands, and at other times, the IBM Storwize V7000 Software performs a discovery process to update the state of the drive and enclosure objects.

13.1.3 Power

All enclosures require two power supply units (PSUs) for normal operation. A single PSU can power the entire enclosure for redundancy.

Control enclosure PSUs contain batteries and expansion enclosure PSUs do not. The additional battery function requires two additional LEDs, which is the main difference between the PSUs when viewed from outside.

There is a power switch on the power supply. The switch must be on for the PSU to be operational. If the power switch is turned off, then the PSU stops providing power to the system. For control enclosure PSUs, the integrated battery continues to be able to supply power to the node canisters.

Figure 13-4 shows the two PSUs present in the controller and expansion enclosure. The controller PSU has two more LEDs than the expansion enclosure because of the battery status indication.

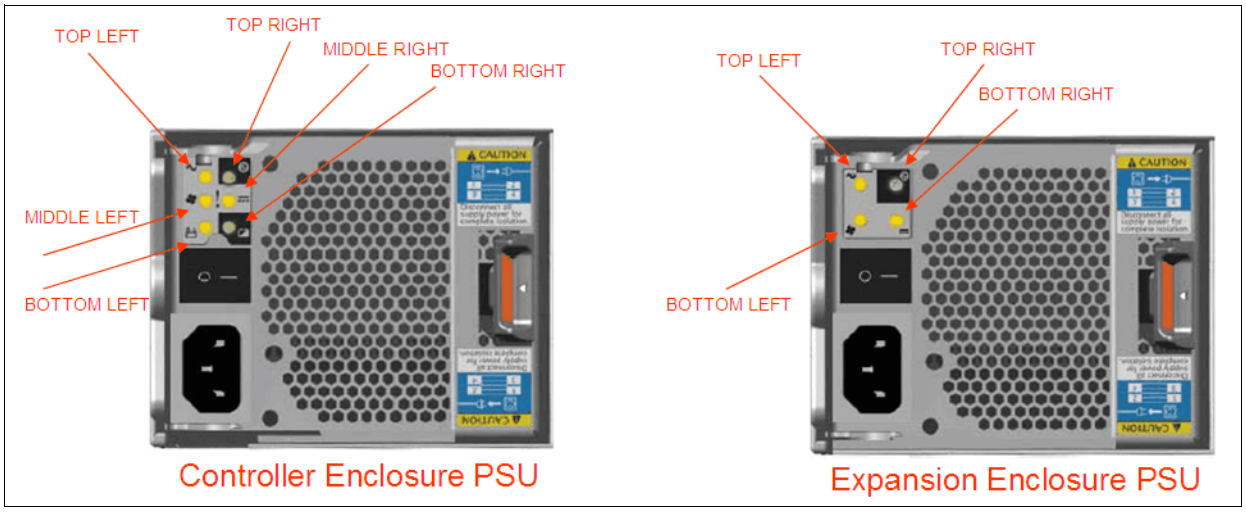

Figure 13-4 Controller and expansion enclosure LED status indicator

Table 13-7 describes the meaning of the LEDs in both enclosures.

Table 13-7 LED statuses of power in the controller and expansion enclosure

Position	Color	Meaning
Upper left	Amber	AC failure
Upper right	Green	PSU OK
Middle left (Controller Enclosure) Lower left (Expansion Enclosure)	Amber	Fan failure
Middle right (Controller Enclosure) Lower left (Expansion Enclosure)	Amber	AC failure
Lower left (Controller Enclosure only)	Amber	Battery fault
Lower right (Controller Enclosure only)	Green	Battery good

Table 13-8 describes the LEDs for the PSU on the controller and expansion enclosure.

Table 13-8 PSU LEDs statuses

PSU OK	AC failure	Fan failure	DC failure	Meaning
On	On	On	On	A communication failure exists between the PSU and enclosure midplane.
On	Off	Off	Off	The PSU is on and working correctly.
Flashing	Any	Any	Any	Neither canister is working; it might be mis-seated, unplugged, or broken. (If the PSU OK is flashing, then the AC, fan, or DC light can be either on or off.)
Off	On	On	On	A communication failure and PSU problem exist.
Off	Flashing	Flashing	Flashing	A PSU firmware download is in progress.
Off	On	Off	On	There is no AC input.
Off	Off	On	Off	A fan has failed.
Off	Off	Off	On	The PSU is not seated correctly.
Off	Off	Off	Off	There is no power to the enclosure.

Table 13-9 describes the LEDs for the battery on the controller enclosure.

Table 13-9 Status of battery LED

Battery is good	Battery is at fault	Meaning
On	Off	The battery is OK and fully charged.
Flashing	Off	The battery is OK but not fully charged. It is either charging or doing a maintenance discharge.
Off	On	The battery is broken.
Off	Flashing	The battery is broken but can be serviced.

13.2 Configuration backup

You can download and save the configuration backup file using the IBM Storwize V7000 GUI or command-line interface (CLI). On an ad-hoc basis, we suggest manually doing this procedure because it is able to save the file directly to your workstation. The command-line option requires login to the system and downloading the dumped file using specific SCP protocol. The command-line option is a good practice for an automated backup of the configuration.

Important: Save configuration files of IBM Storwize V7000 regularly. The best approach is to do this daily and automate this task. Always perform the additional backup before any critical maintenance task such an upgrade of the microcode, software version, and so on.

The backup file is updated by the cluster every day. Saving it after any changes to your system configuration is also important. It contains configuration data of arrays, pools, volumes, and so on. The backup never contains any client data.

To successfully perform the configuration backup, follow the prerequisites and requirements:

- All nodes must be online.
- No independent operations that change the configuration can be running in parallel.
- No object name can begin with an underscore.
- All objects should have non-default names.

Although objects should have non-default names at the time that the backup is taken, this prerequisite is not mandatory. The backup command reports an error when the default name is discovered, but the configuration is saved. However, the default object names for controllers, I/O groups, and MDisks do not restore correctly if the ID of the object differs from what is recorded in the current configuration data file. All other objects with default names are renamed during the restore process.

13.2.1 Configuration backup using the CLI

You can use the CLI to trigger configuration backup. The **svcconfig backup** command generates a new backup file. Triggering a parallel backup using the GUI is not possible, but you can save it by using only the GUI.

Example 13-1 shows output of the **svcconfig backup** command.

Example 13-1 Saving configuration using CLI

```
IBM_Storwize:ITSO-v7000:ITSO_admin>svcconfig backup
.............
CMMVC6130W Cluster ITSO-v7000-2 with inter-cluster partnership fully_configured
will not be restored
..................................................
..................................................
..................................................
CMMVC6155I SVCCONFIG processing completed successfully
IBM_Storwize:ITSO-v7000:ITSO_admin>
```

The **svcconfig backup** command generates three files that provide information about the backup process and cluster configuration. These files are dumped into the /tmp directory on the configuration node.

Table 13-10 describes the three files that are created by the backup process.

Table 13-10 File names created by the backup process

File name	Description
`svc.config.backup.xml`	This file contains your cluster configuration data.
`svc.config.backup.sh`	This file contains the names of the commands that were issued to create the backup of the cluster.
`svc.config.backup.log`	This file contains details about the backup, including any error information that might have been reported.

More information about the CLI of the IBM Storwize V7000 system is in Appendix A, “CLI setup and SAN Boot” on page 497.

13.2.2 Configuration backup using the GUI

To download a backup of the configuration using the GUI, complete the following steps:

1. Navigate to the **Settings** icon and click **Support** (Figure 13-5).

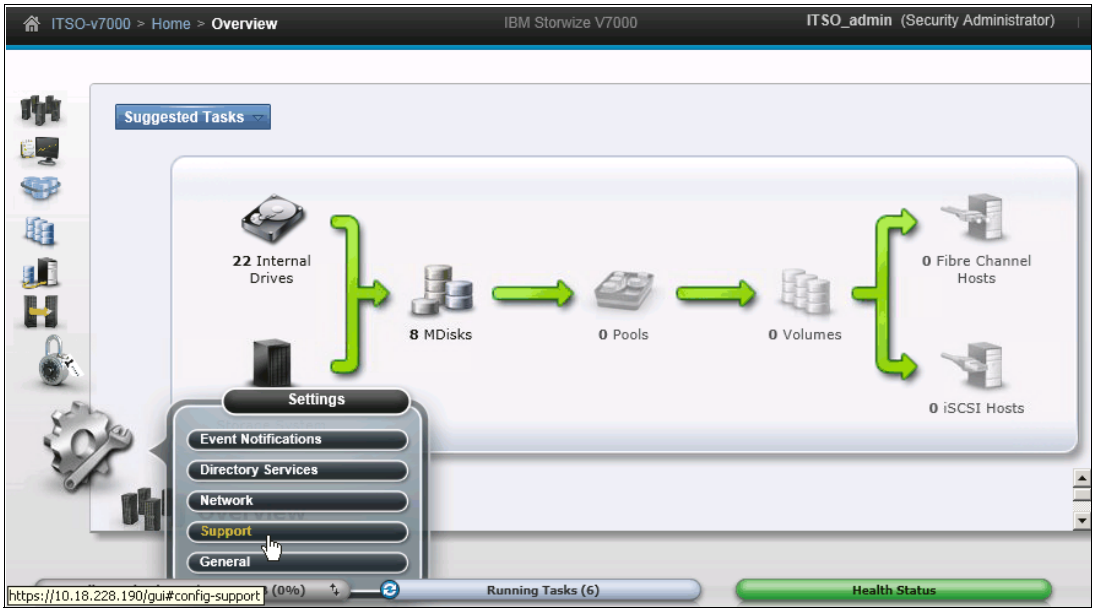

Figure 13-5 Support option

2. The window shown in Figure 13-6 opens. Click **Show full log listing** to show all log files.

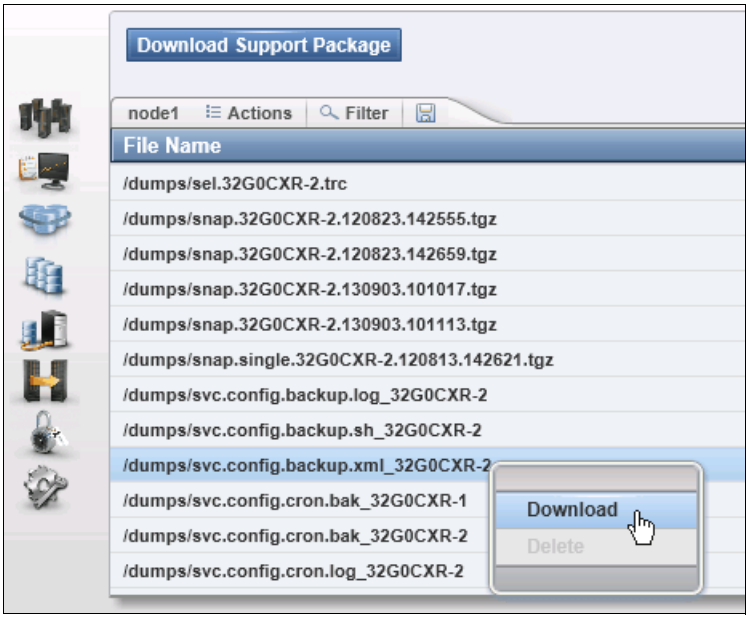

Figure 13-6 Show full log files window

3. Search for and right-click the following file name, and then select **Download** to transfer the file to your workstation:

   `/dumps/svc.config.backup.xml_*`

13.3 Software upgrade

Concurrent software update of all components is supported through the standard Ethernet management interfaces.

To upgrade the IBM Storwize V7000 software, complete the following steps:

1. Open a supported web browser and navigate to your cluster IP address. A login window opens (Figure 13-7).

Figure 13-7 IBM Storwize V7000 GUI login window

2. Log in as your superuser. The IBM Storwize V7000 management home window opens. Move the mouse cursor over **Settings** → **General** (Figure 13-8) and click **General**.

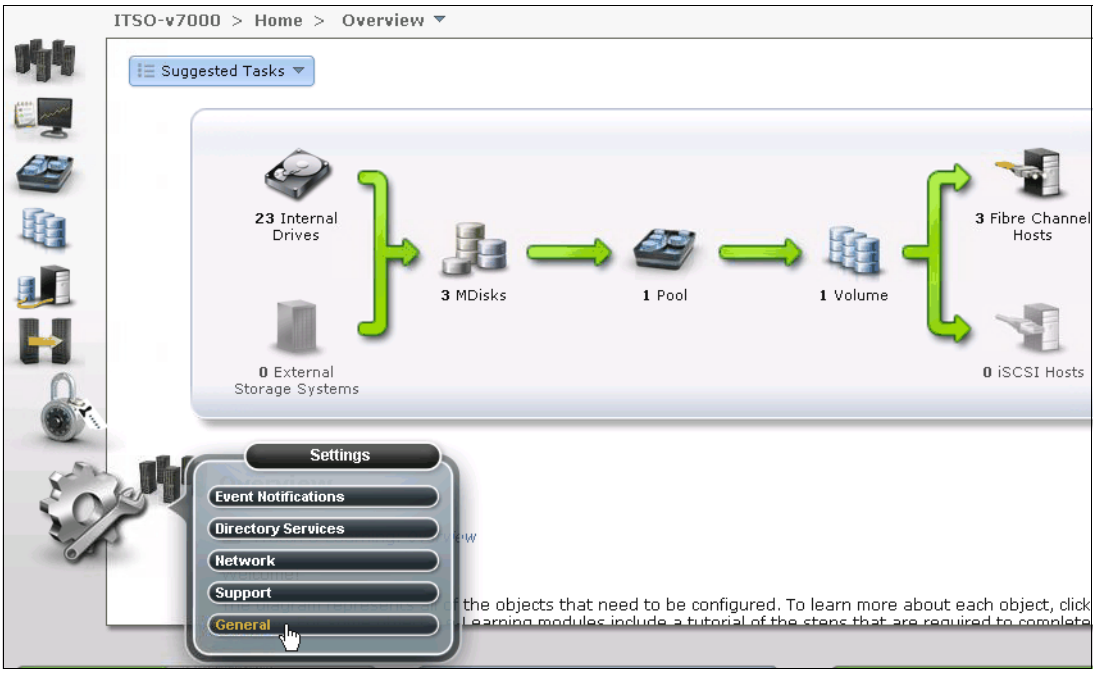

Figure 13-8 Settings menu

3. In the General menu, choose **Upgrade Software**. The Upgrade Software pane opens (Figure 13-9).

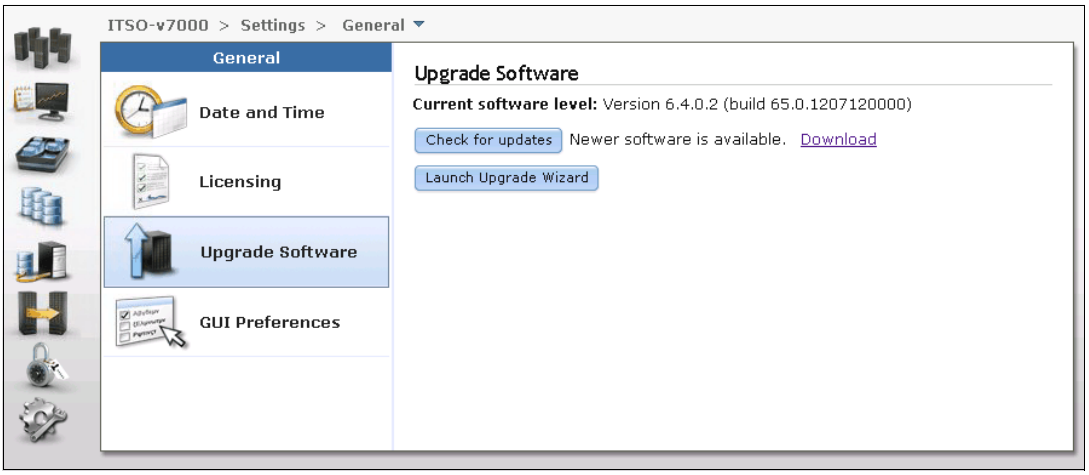

Figure 13-9 Upgrade Software menu

This window offers two options:

– Check for updates: This function checks with the IBM website about whether a version of IBM Storwize V7000 software is newer than the version you have installed. You must have an Internet connection to use this function.

My Notifications: Use the My Notifications tool to receive notifications of new and updated support information to better maintain your system environment, especially in an environment where a direct Internet connection is not possible. Go to the following address (an IBM account is required) and add your IBM Storwize V7000 system to the notifications list to be advised of support information and to download the latest code to your workstation for later upload:

http://www.ibm.com/software/support/einfo.html

– Launch Upgrade wizard: This function starts the software upgrade process.

4. Click **Launch Upgrade Wizard** to start the upgrade process. You are redirected to the window shown in Figure 13-10.

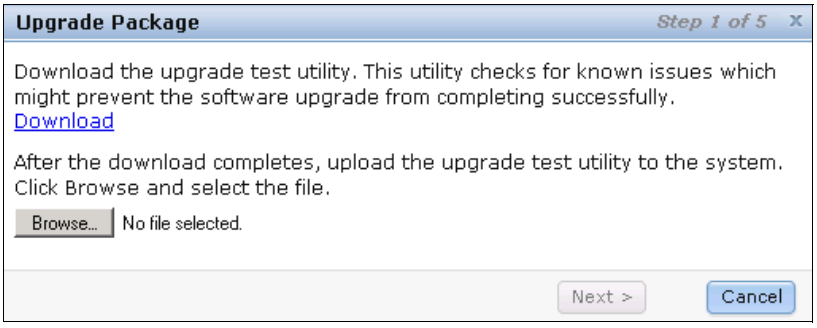

Figure 13-10 Upgrade Package window

From this window, you can download the upgrade test utility. If you downloaded it previously, you can browse to the location where it is saved (Figure 13-11).

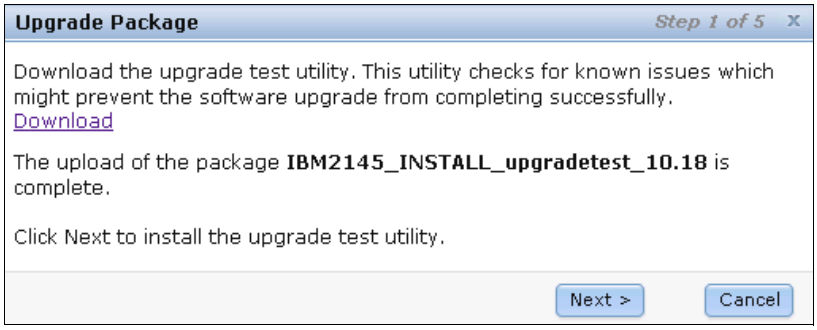

Figure 13-11 Upload Test Utility

Naming: The upgrade test utility is common for all products from the IBM SAN Volume Controller family. Do not be confused by the `IBM2145_` naming convention, which refers to SAN Volume Controller products (while IBM Storwize V7000 is marked as `IBM2076_`).

5. Click **Next**; the upgrade test utility is applied and the next window opens (Figure 13-12).

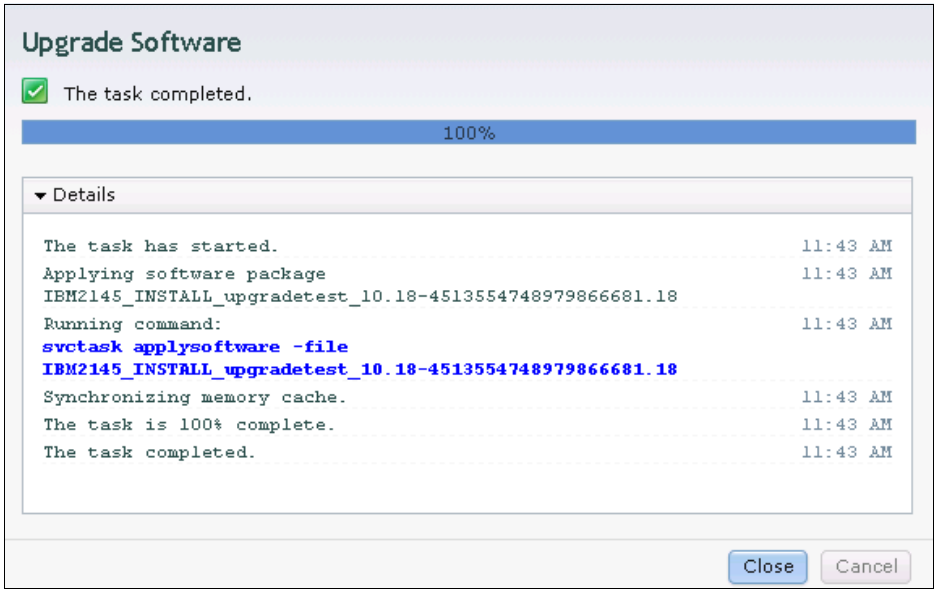

Figure 13-12 Upgrade Test Utility applied

6. Click **Close**. The window shown in Figure 13-13 opens, where you can run the Upgrade Test Utility.

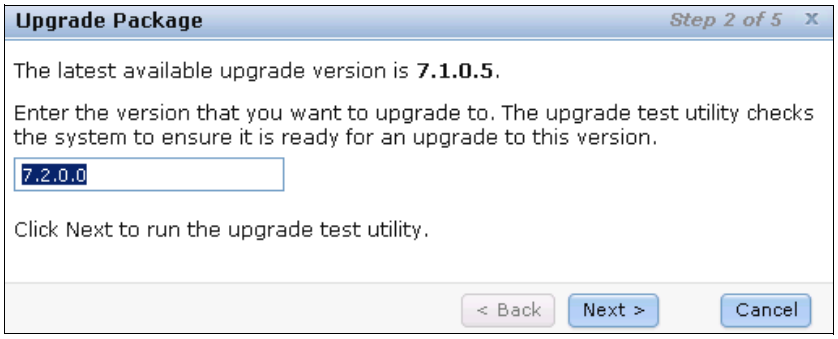

Figure 13-13 Run Upgrade Test Utility

7. Click **Next**. You are redirected to the window shown in Figure 13-14. At this time, the Upgrade Test Utility runs and you are able to see the suggested actions, if any, in the window shown in Figure 13-14.

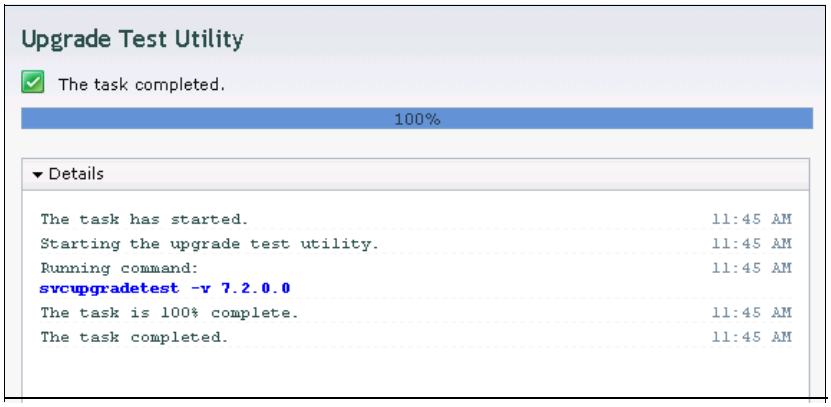

Figure 13-14 Results of the Test Utility

8. Click **Next** to start the software upload procedure, and you are redirected to the window shown in Figure 13-15.

Figure 13-15 Downloading software code

9. Download the IBM Storwize V7000 software upgrade package, or browse and upload the software upgrade package from the location where you saved it. Verify that file is successfully loaded into the system (Figure 13-16).

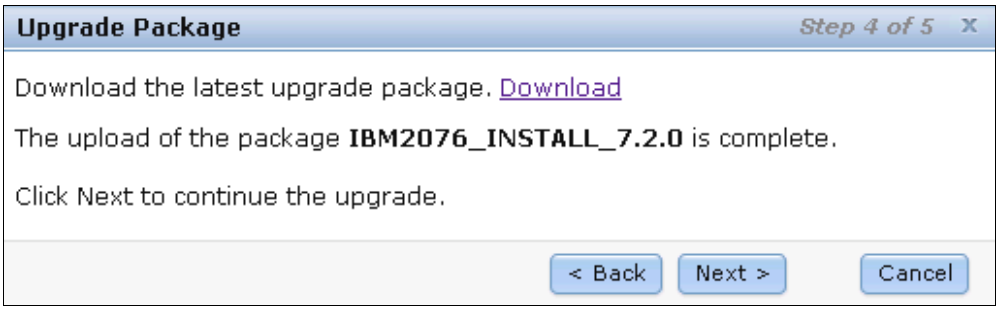

Figure 13-16 System ready to upgrade

When prompted, choose **Automatic upgrade** rather than **Service Assistant Manual upgrade**. Manual upgrade is eligible for the cases when the action is suggested and monitored by IBM support personnel (Figure 13-17).

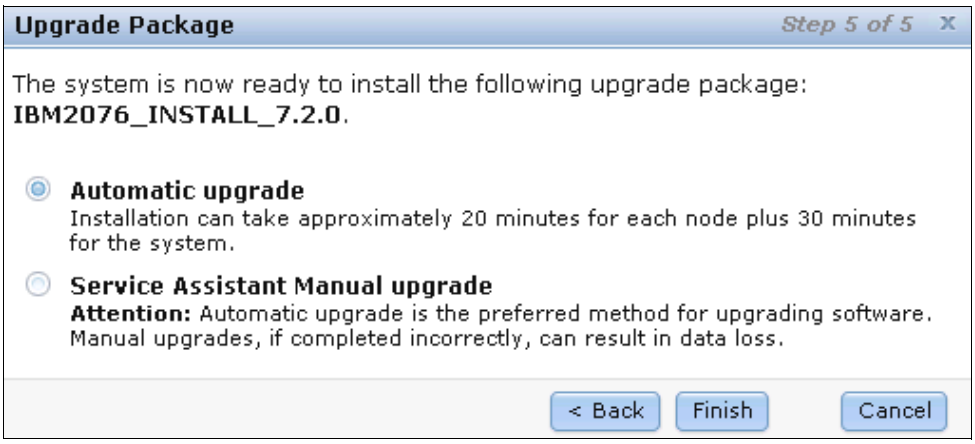

Figure 13-17 The automatic upgrade selection

10. Click **Finish**. The software upgrade starts. You are redirected to the window illustrated in Figure 13-18. Ignore the warning message that is related to the detection of cluster node if it appears.

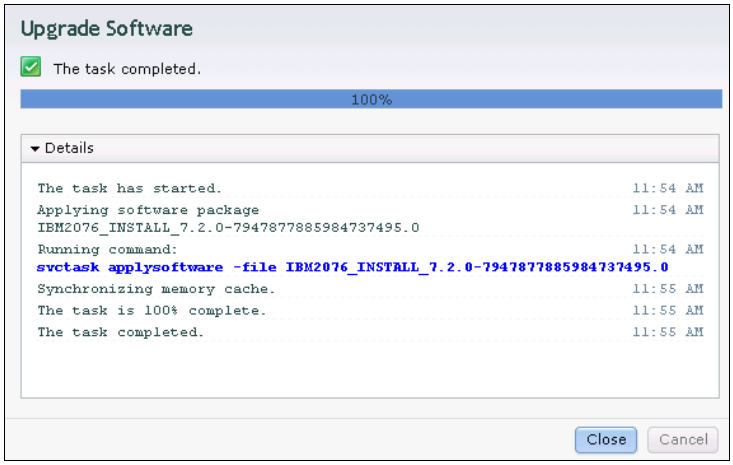

Figure 13-18 Upgrade the software code

11. Click **OK**. You have completed your task to upgrade the SAN Volume Controller software. You receive messages that inform you that first one node, then the other, has been upgraded. When both nodes have been rebooted, you have completed your SAN Volume Controller software upgrade. The section of messages is illustrated Figure 13-19 on page 471.

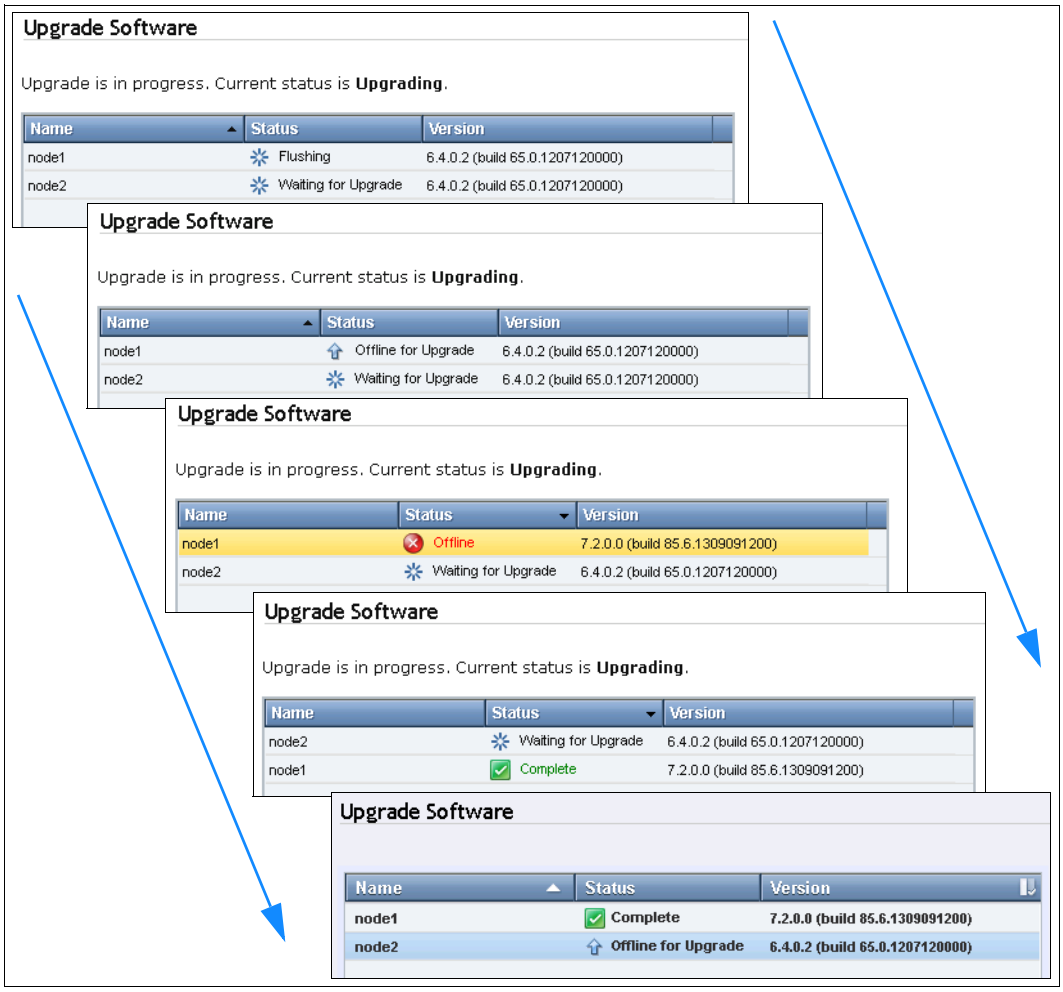

Figure 13-19 Information messages during upgrade process

13.4 Critical Fix Notification feature

Starting with the IBM Storwize V7000 software V6.3.0.0, the new *Critical Fix Notification* function enables IBM to warn IBM Storwize V7000 and SAN Volume Controller users when a critical issue exists in the level of code that they are using. The system notifies users when they log on to the GUI using an web browser connected to the Internet.

Consider the following information about this function:

- It warns users only about critical fixes and does not warn them that they are running a previous version of the software.
- It works only if the browser also has access to the Internet (the IBM Storwize V7000 and SAN Volume Controller systems themselves do not need to be connected to the Internet).
- The function cannot be disabled, and each time it displays a warning, it must be acknowledged (with the option to not warn the user again for that issue).

The decision about what is a “critical” fix is subjective and requires judgment, which is exercised by the development team. As a result, clients might still encounter bugs in code that were not deemed critical. They should continue to review information about new code levels to determine if they are supposed to upgrade even without a critical fix notification.

13.5 Monitoring host paths to the IBM Storwize V7000 system

In this section, we show how to monitor paths to the IBM Storwize V7000 system. On host systems using the IBM multipath module, such as Windows and other operating systems, a powerful CLI tool is available (see the basic Windows configuration in Chapter 4, “Host configuration” on page 105 for information about monitoring IBM Storwize V7000 connectivity from the host). Example 13-2 shows the output of **datapath query device** command with four paths that should be available, but two of the paths have been closed, indicating a problem.

Example 13-2 pcmpath query device showing a problem with one canister

```
C:\Program Files\IBM\SDDDSM>datapath query device
Total Devices : 2

DEV#:   0  DEVICE NAME: Disk2 Part0  TYPE: 2145       POLICY: OPTIMIZED
SERIAL: 600507680280801AC80000000000001F
============================================================================
Path#            Adapter/Hard Disk        State  Mode       Select     Errors
    0     Scsi Port5 Bus0/Disk2 Part0     OPEN   NORMAL         78          1
    1     Scsi Port5 Bus0/Disk2 Part0     OPEN   NORMAL          0          0
    2     Scsi Port6 Bus0/Disk2 Part0    CLOSE   NORMAL          0          0
    3     Scsi Port6 Bus0/Disk2 Part0    CLOSE   NORMAL         25          1

DEV#:   1  DEVICE NAME: Disk3 Part0  TYPE: 2145       POLICY: OPTIMIZED
SERIAL: 600507680280801AC800000000000020
============================================================================
Path#            Adapter/Hard Disk        State  Mode       Select     Errors
    0     Scsi Port5 Bus0/Disk3 Part0     OPEN   NORMAL          1          1
    1     Scsi Port5 Bus0/Disk3 Part0     OPEN   NORMAL         78          1
    2     Scsi Port6 Bus0/Disk3 Part0    CLOSE   NORMAL         28          1
    3     Scsi Port6 Bus0/Disk3 Part0    CLOSE   NORMAL          1          1
```

The **datapath query adapter** command (Example 13-3) shows all IBM Storwize V7000 paths that are available to the host. It shows that only Adapter 0 is available, and that the state of Adapter 1 is FAILED.

Example 13-3 datapath query adapter command

```
C:\Program Files\IBM\SDDDSM>datapath query adapter
Active Adapters :2

Adpt#        Name        State    Mode       Select     Errors  Paths   Active
    0  Scsi Port5 Bus0   NORMAL   ACTIVE         36          0      2        2
    1  Scsi Port6 Bus0   FAILED   ACTIVE         17          1      2        0
```

After the problem is fixed, scan for new disks on your host, and verify that all paths are available again (Example 13-4).

Example 13-4 datapath query device command

```
C:\Program Files\IBM\SDDDSM>datapath query device
Total Devices : 2

DEV#:   0  DEVICE NAME: Disk2 Part0  TYPE: 2145       POLICY: OPTIMIZED
SERIAL: 600507680280801AC80000000000001F
============================================================================
Path#            Adapter/Hard Disk        State  Mode       Select     Errors
    0     Scsi Port5 Bus0/Disk2 Part0     OPEN   NORMAL         95          1
    1     Scsi Port5 Bus0/Disk2 Part0     OPEN   NORMAL          0          0
    2     Scsi Port6 Bus0/Disk2 Part0     OPEN   NORMAL          0          0
    3     Scsi Port6 Bus0/Disk2 Part0     OPEN   NORMAL         28          1

DEV#:   1  DEVICE NAME: Disk3 Part0  TYPE: 2145       POLICY: OPTIMIZED
SERIAL: 600507680280801AC800000000000020
============================================================================
Path#            Adapter/Hard Disk        State  Mode       Select     Errors
    0     Scsi Port5 Bus0/Disk3 Part0     OPEN   NORMAL          1          1
    1     Scsi Port5 Bus0/Disk3 Part0     OPEN   NORMAL         95          1
    2     Scsi Port6 Bus0/Disk3 Part0     OPEN   NORMAL         31          1
    3     Scsi Port6 Bus0/Disk3 Part0     OPEN   NORMAL          1          1
```

You can also use the **`datapath query adapter`** command and check that the FAILED path is back online, where the State is indicated as `NORMAL` (Example 13-5).

Example 13-5 datapath query adapter to check if a failed path is online

```
C:\Program Files\IBM\SDDDSM>datapath query adapter
Active Adapters :2

Adpt#         Name        State    Mode       Select     Errors  Paths   Active
    0  Scsi Port5 Bus0   NORMAL   ACTIVE         52          0      2        2
    1  Scsi Port6 Bus0   NORMAL   ACTIVE         21          1      2        2
```

13.6 Troubleshooting and fix procedures

The management GUI of IBM Storwize V7000 is a browser-based GUI for configuring and managing all aspects of your system. It provides extensive facilities to help troubleshoot and correct problems. We explain how to effectively use its features to avoid service disruption of your IBM Storwize V7000.

Figure 13-20 on page 474 shows the menu to start the Monitoring wizard for *System* information, gathering *System Details*, viewing *Events*, or seeing real-time *Performance* statistics.

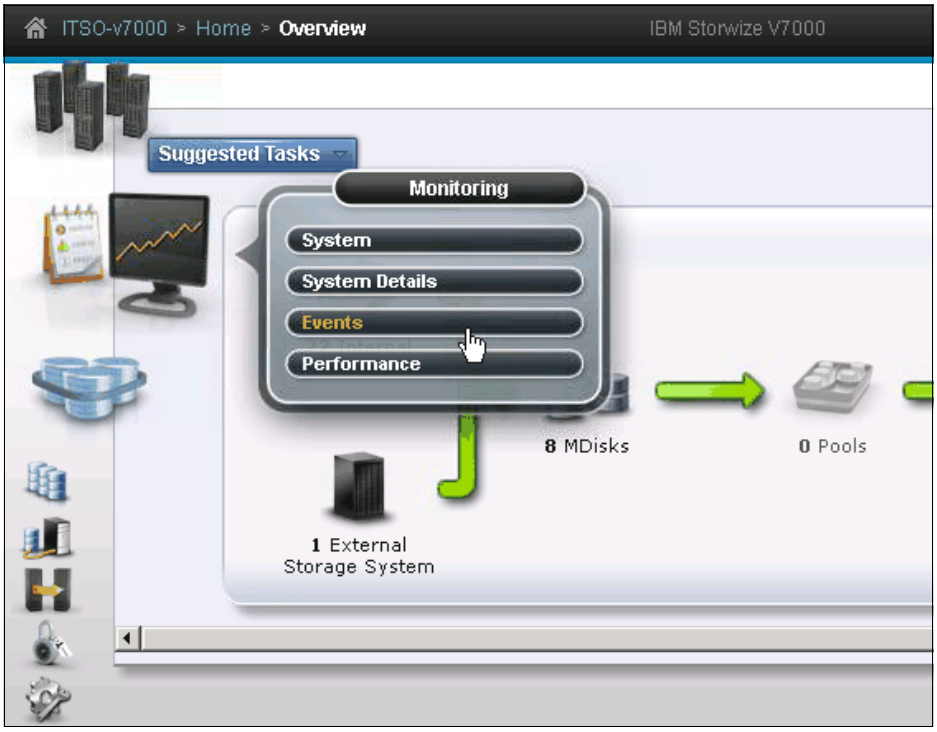

Figure 13-20 Monitoring options

Use the management GUI to manage and service your system. Select **Monitoring** → **Events** to access problems that must be fixed and maintenance procedures that step you through the process of correcting problems. Information in the Events panel can be filtered in three ways:

- Recommended Actions

  Shows only the alerts that require attention. Alerts are listed in priority order and should be fixed sequentially by using the available fix procedures. For each problem that is selected, you can do these tasks:

  - Run a fix procedure
  - View the properties

- Unfixed Messages and Alerts

  Displays only the alerts and messages that are not fixed. For each entry that is selected, you can do these tasks:

  - Run a fix procedure
  - Mark an event as fixed
  - Filter the entries to show them by specific minutes, hours, or dates
  - Reset the date filter
  - View the properties

- Show All

  Displays all event types whether they are fixed or unfixed. For each entry that is selected, you can do these tasks:

  - Run a fix procedure
  - Mark an event as fixed
  - Filter the entries to show them by specific minutes, hours, or dates
  - Reset the date filter
  - View the properties

Some events require a certain number of occurrences in 25 hours before they are displayed as unfixed. If they do not reach this threshold in 25 hours, they are flagged as *expired*. Monitoring events are below the coalesce threshold and are usually transient.

Important: The management GUI is the primary tool that is used to *operate* and *service* your system; the real-time *monitoring* should be established by SNMP traps, email notifications, or syslog messaging on an automatic manner.

13.6.1 Managing event log

Regularly check the status of the system using the management GUI: If you suspect a problem, first use the management GUI to diagnose and resolve the problem.

Use the views that are available in the management GUI to verify the status of the system, the hardware devices, the physical storage, and the available volumes.

1. Select **Monitoring** → **Events** to see all problems that exist on the system (Figure 13-21).

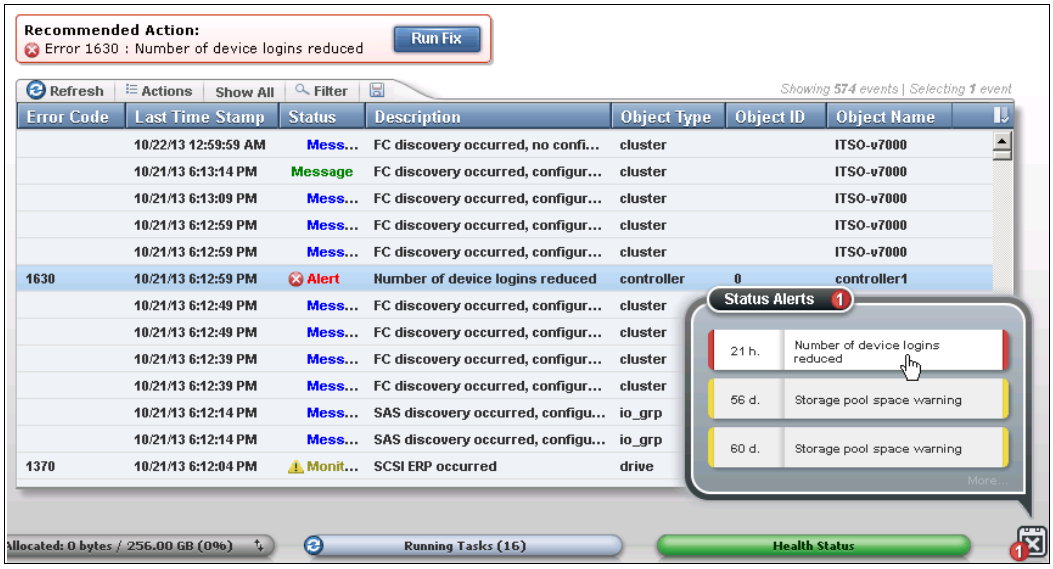

Figure 13-21 Messages in the event log

2. Select **Show All** → **Recommended Actions** to display the most important events to be resolved (Figure 13-22). The Recommended Actions tab shows the highest priority maintenance procedure that must be run. Use the troubleshooting wizard so that IBM Storwize V7000 system can determine the proper order of maintenance procedures.

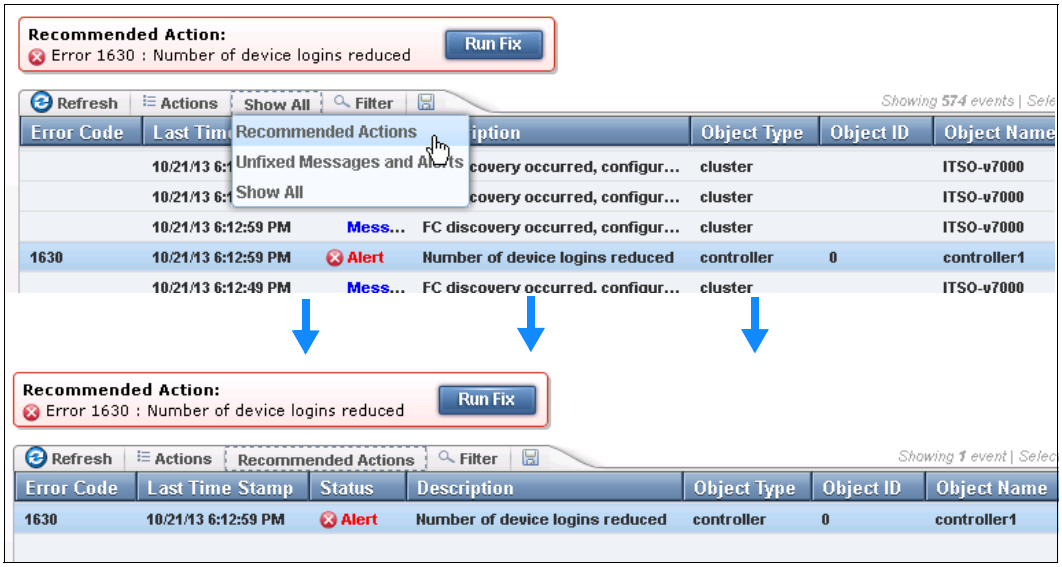

Figure 13-22 Recommended Actions

In this example, the *number of device logins reduced* is listed (service error code 1630). Review the physical FC cabling to determine the issue and then click **Run Fix**. At any time and from any GUI panel, you can directly navigate to this menu by using the Status Alerts icon at the right bottom corner of the GUI (Figure 13-23).

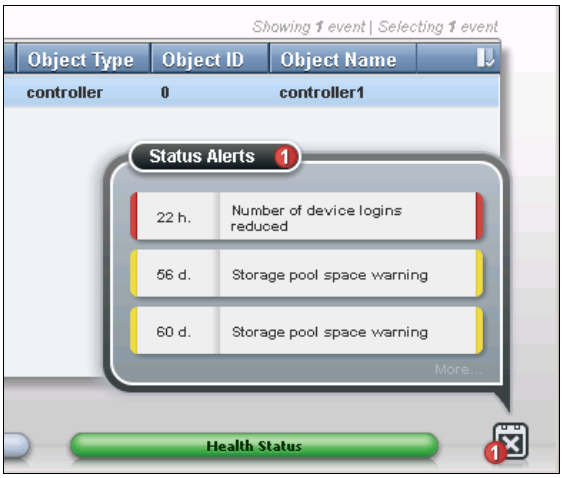

Figure 13-23 Status alerts

13.6.2 Running a fix procedure

If there is a service error code for the alert, you can run a fix procedure that assists you in resolving the problem. These fix procedures analyze the system and provide more information about the problem. They suggest actions to take and step you through the actions that automatically manage the system where necessary. Finally, they check that the problem is resolved.

If an error is reported, always use the fix procedures within the management GUI to resolve the problem. Always use the fix procedures for both software configuration problems and hardware failures. The fix procedures analyze the system to ensure that the required changes do not cause volumes to be inaccessible to the hosts. The fix procedures automatically perform configuration changes that are required to return the system to its optimum state.

Preferred practice: Before running a fix procedure, take a backup of the system configuration, as suggested in 13.2, “Configuration backup” on page 463. Also a good practice is to also perform the additional backup after the procedure.

The fix procedure displays information that is relevant to the problem and provides various options to correct the problem. Where possible, the fix procedure runs the commands that are required to reconfigure the system. Always use the recommended action for an alert because these actions ensure that all required steps are performed. Use the recommended actions even in cases where the service action seems obvious, such as a drive showing a fault. In this case, the drive must be replaced and reconfiguration must be performed. The fix procedure performs the reconfiguration for you.

The fix procedure also checks that another existing problem does not result in a fix procedure that causes volume data to be lost. For example, if a power supply unit in a node enclosure must be replaced, the fix procedure checks and warns you if the integrated battery in the other power supply unit is not sufficiently charged to protect the system.

Hint: If possible, fix the alerts in the order shown to resolve the most serious issues first. Often, other alerts are corrected automatically because they were the result of a more serious issue.

The following example demonstrates how to clear the error that is related to the malfunctioning FC connectivity between control canisters of the IBM Storwize V7000:

1. From the dynamic menu (the icons on the left) select **Monitoring** → **Events**, and then focus on the errors with the highest priority first. List only the recommended actions by selecting the filters in the **Actions** menu (Figure 13-24). Click **Run Fix**.

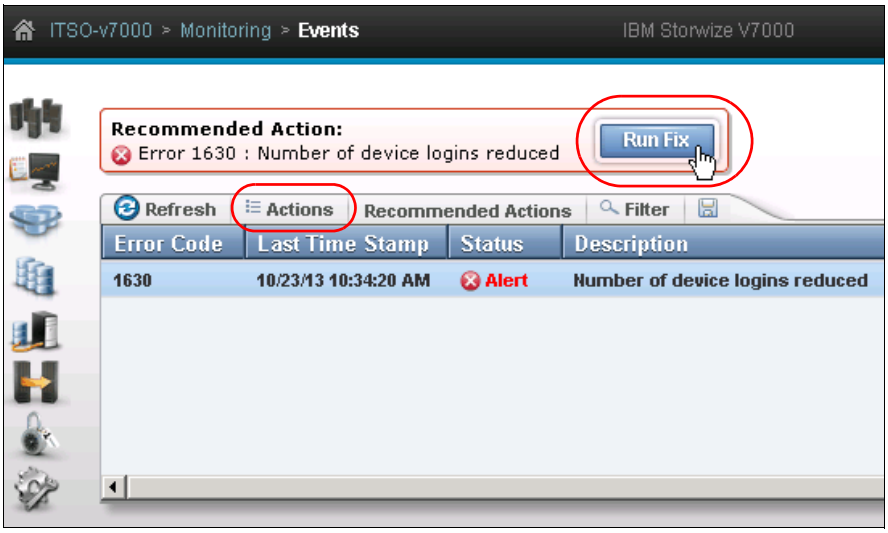

Figure 13-24 Initiate Run Fix Procedure from management GUI

2. The pop-up window asks whether the issue was caused by a planned change or maintenance task or whether it appeared in an uncontrolled manner (Figure 13-25).

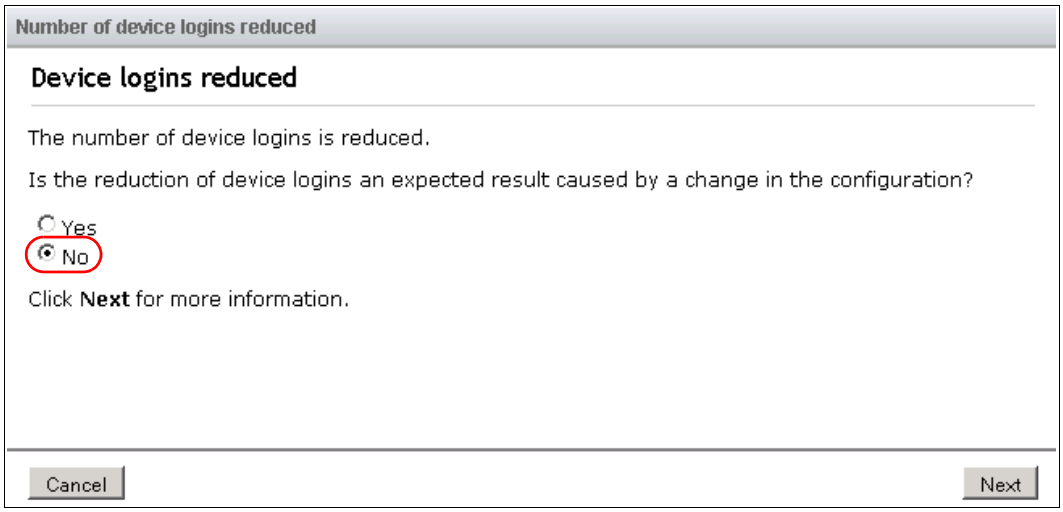

Figure 13-25 Determination of planned action

If you answer **Yes**, the fix procedure finishes, assuming that all changes in the system are done on purpose and no other action is necessary. However, our example simulates a broken FC cable and follow the complete fix procedure. Select **No** and click **Next**.

3. In the next window (Figure 13-26), IBM Storwize V7000 lists suggested actions and which components must be checked to fix and close the error. When you are sure that all possible technical requirements are met (in our case we replaced a broken FC cable), click **Next**.

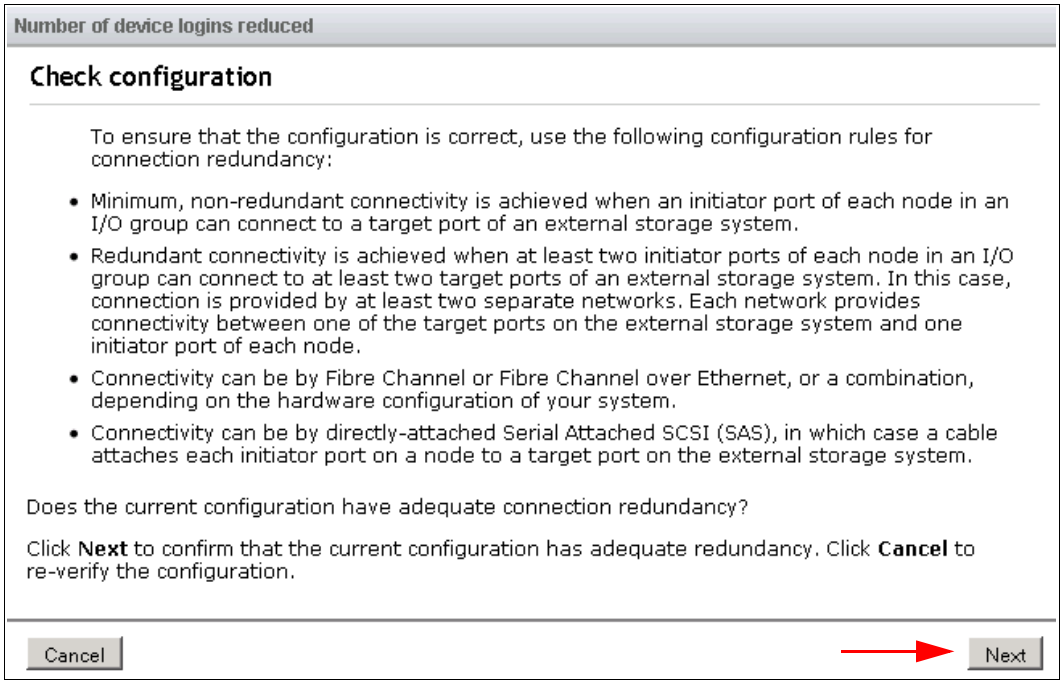

Figure 13-26 Verification steps to eliminate single point of failure

The discovery of managed disks starts (Figure 13-27).

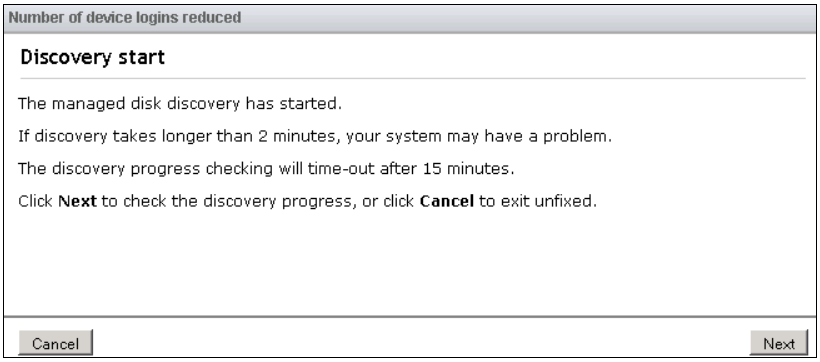

Figure 13-27 Starting the discovery of managed disks

If no other important issue exists, it finishes maximally within 2 minutes, depending on the number of enclosures and installed disk drives (Figure 13-28).

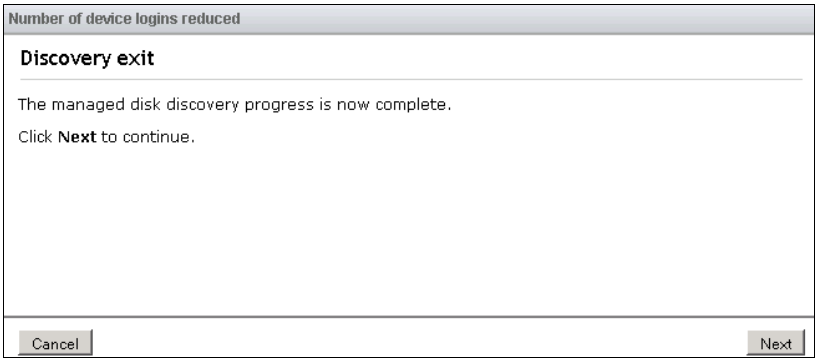

Figure 13-28 Discovery complete

4. An event has been marked as fixed and you can safely finish the fix procedure. Click **Close** and the event is removed from the list of events (Figure 13-29).

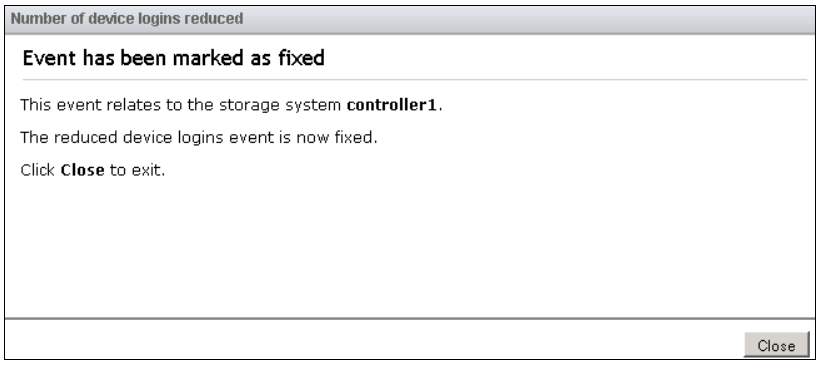

Figure 13-29 Correctly finished fix procedure

Resolve alerts in a timely manner

Perform the recommended actions as quickly as possible after the problem is reported. Your system is designed to be resilient to most single hardware failures. However, if it operates for any period of time with a hardware failure, the possibility increases that a second hardware failure can result in some volume data that is unavailable. If several unfixed alerts exist, fixing any one alert might become more difficult because of the effects of the others.

13.6.3 Event log details

Multiple views of the events and recommended actions are available. The GUI works like a typical Windows context menu, so the event log grid is manipulated through the row that contains the column headings (Figure 13-30). When you right-click a table heading, a menu for the column choices opens.

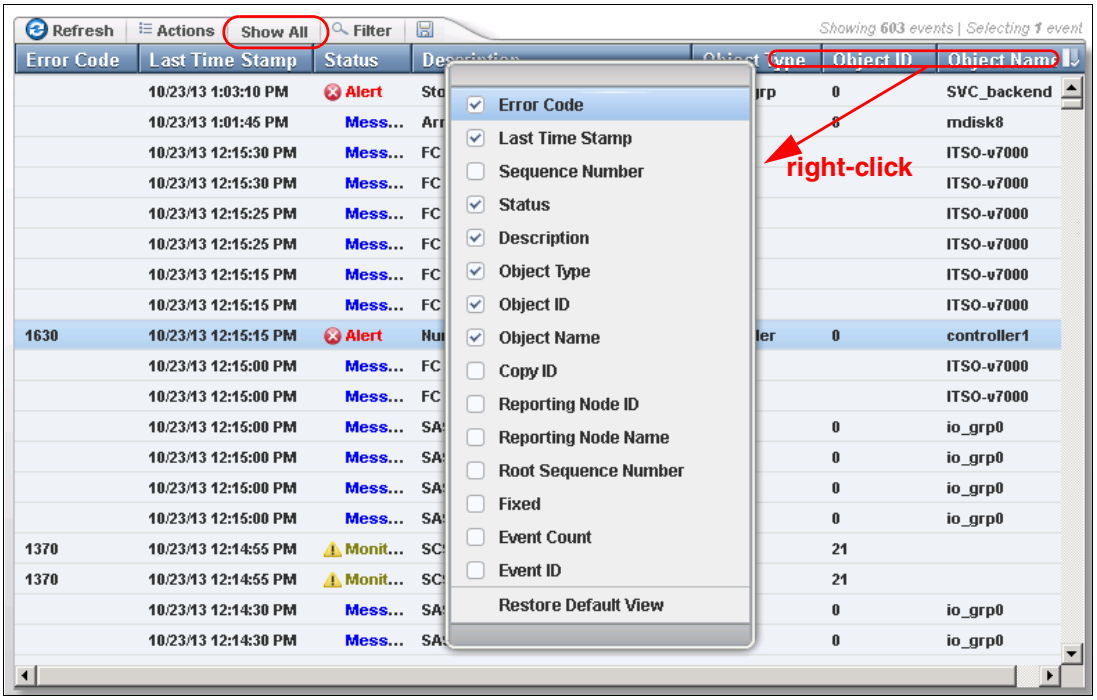

Figure 13-30 Grid options of the event log

Select or remove columns as needed. You can then also extend or shrink the width of the column to fit your screen resolution and size. This way is manipulated majority of grids in the management GUI of IBM Storwize V7000, not only the events panel.

Every field of the event log is available as a column in the event log grid. Several fields are useful when you work with IBM Support. The preferred method in this case is to use the **Show All** filter, with events sorted by time stamp. All fields have the sequence number, event count, and the fixed state. Using Reset Grid Preferences sets the grid back to the factory defaults.

You might want see more details about each critical event. Some details are not shown in the main grid. To access properties and sense data of specific event, right-click the specific event (anywhere in its row) and choose **Properties** from the context menu.

The properties window opens (Figure 13-31) with all the relevant sense data such as first and last time of an event occurrence, worldwide port name (WWPN) and worldwide node name (WWNN), enabled or disabled automatic fix, and more.

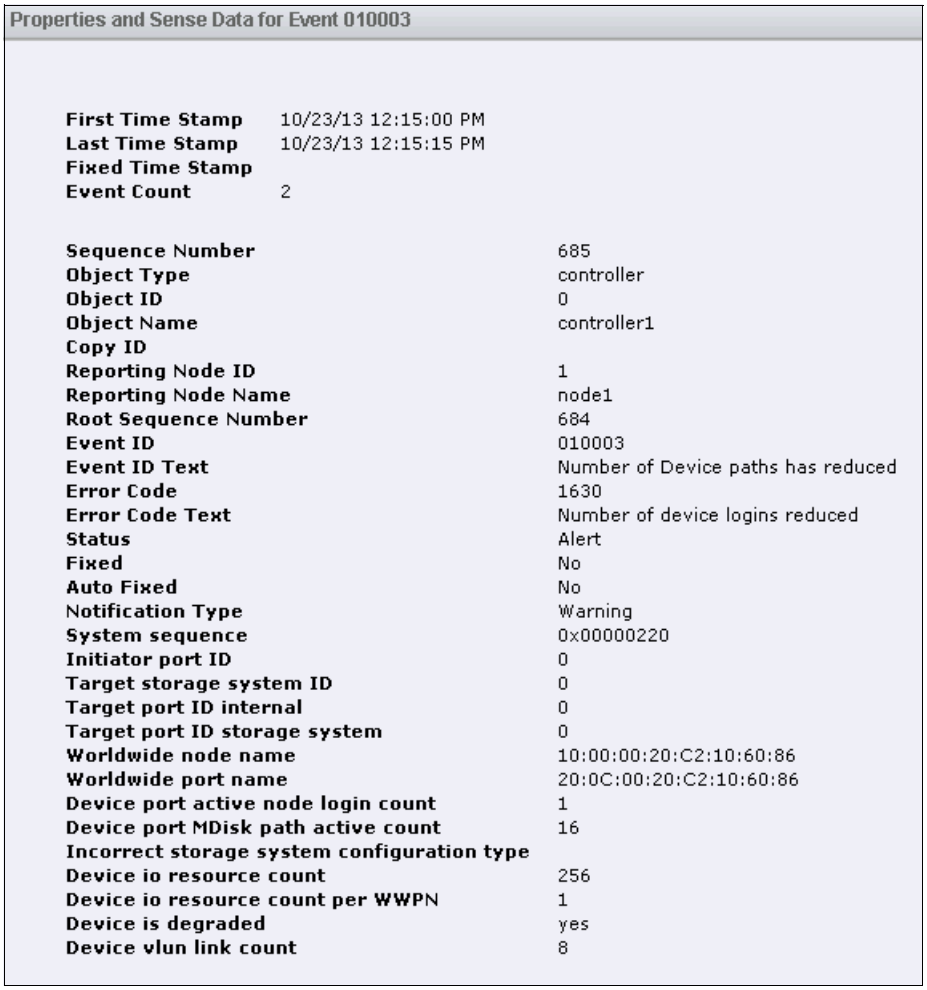

Figure 13-31 Event sense data and properties

For more details about troubleshooting options, see the IBM Storwize V7000 *Troubleshooting, Recovery, and Maintenance Guide*, GC27-2291, which is available at the following location:

http://pic.dhe.ibm.com/infocenter/storwize/ic/topic/com.ibm.storwize.v7000.doc/tbrd_bkmap_trblshoot.pdf

13.7 Monitoring

An important steps is to correct any issues that are reported by your IBM Storwize V7000 system as soon as possible. Configure your system to send automatic notifications when a new event is reported. To avoid monitoring for new events that use the management GUI, select the type of event for which you want to be notified. For example, restrict notifications to just events that require immediate action. Several event notification mechanisms exist:

Email An event notification can be sent to one or more email addresses. This mechanism notifies individuals of problems. Individuals can receive notifications wherever they have email access, including mobile devices.

SNMP A Simple Network Management Protocol (SNMP) traps report can be sent to a data-center management system, such as IBM Systems Director, that consolidates SNMP reports from multiple systems. With this mechanism, you can monitor your data center from a single workstation.

Syslog A syslog report can be sent to a data center management system that consolidates syslog reports from multiple systems. With this option, you can monitor your data center from a single location.

If your system is within warranty, or you have a hardware maintenance agreement, configure your IBM Storwize V7000 system to send email events directly to IBM, if an issue that requires hardware replacement is detected. This mechanism is known as *Call Home*. When this event is received, IBM automatically opens a problem report, and if appropriate, contacts you to verify whether replacement parts are required.

Important: If you set up Call Home to IBM, ensure that the contact details that you configure are correct and kept up to date because personnel can change.

13.7.1 Email notifications and Call Home

The Call Home function of the IBM Storwize V7000 utilizes the email notification being sent to the specific IBM support center, therefore the configuration is similar as in case of sending emails to the specific person or system owner. The following procedure summarizes how to configure email notifications and emphasizes what is specific to Call Home.

1. Prepare your contact information that you want to use for the email notification and verify the accuracy of the data. From the dynamic menu, select **Setting** → **Event Notifications** (Figure 13-32 on page 483).

Figure 13-32 Configuration of event notifications

2. Select **Email** and then click **Enable Email Event Notification** (Figure 13-33). You can also access the IBM eLearning movie for more technical details.

   http://www.ibm.com/support/knowledgecenter/ST3FR7_7.2.0/com.ibm.storwize.v7000.720.learning/events_launch.html?lang=en

   For the correct functionality of email notifications, ask your network administrator if SMTP is enabled on the management network and is not, for example, blocked by firewalls. Be sure to test the accessibility to the SMTP server using the **`telnet`** command (port 25 for a non-secured connection, port 465 for SSL-encrypted communication) using any server in the same network segment.

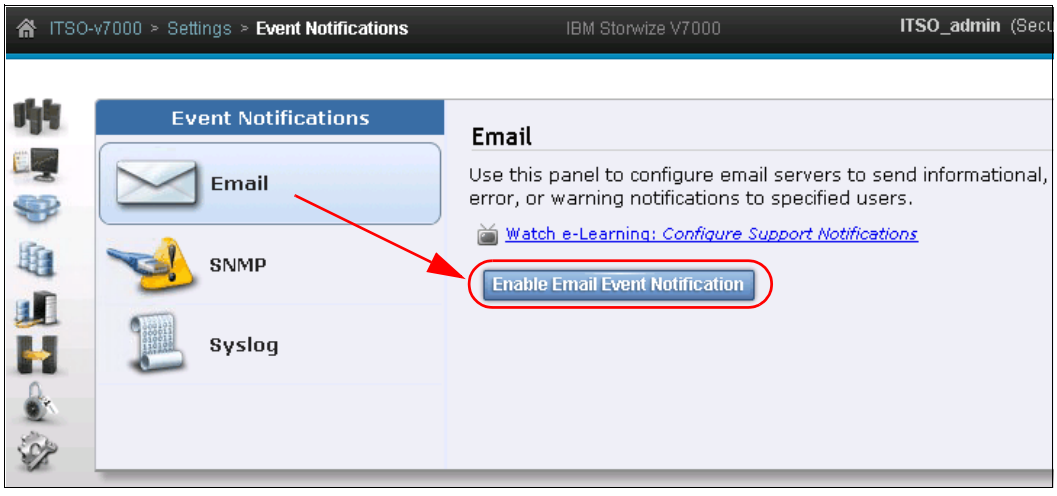

Figure 13-33 Configuration of email notifications

3. Provide the information about the location of the system (Figure 13-34) and contact information of the IBM Storwize V7000 owner (Figure 13-35) in order to be reachable by IBM support. *Always* keep this information current.

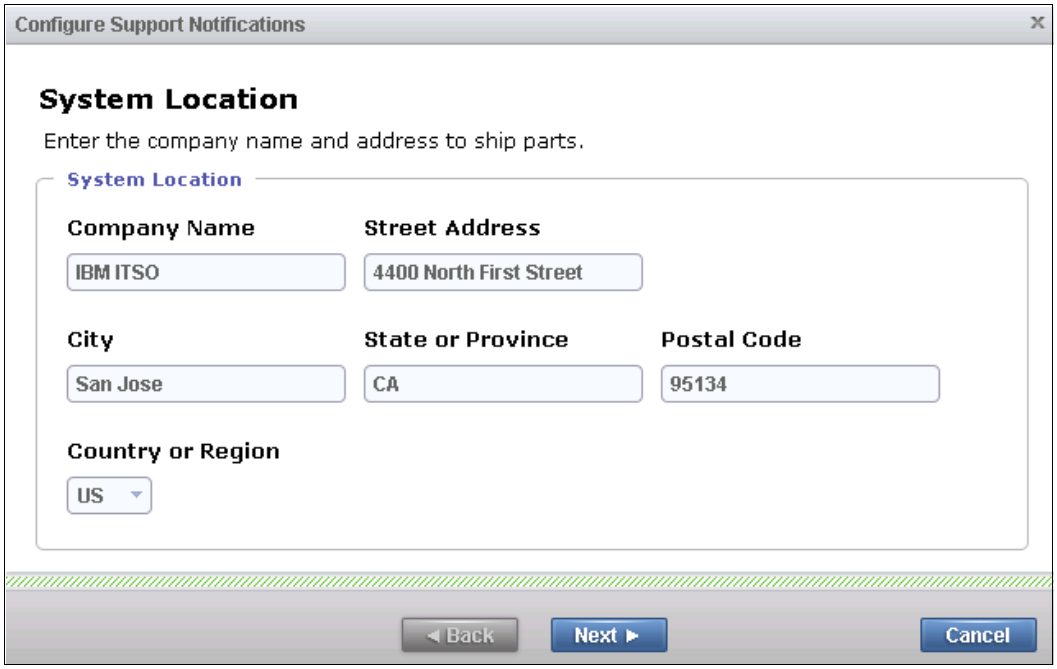

Figure 13-34 Location of the device

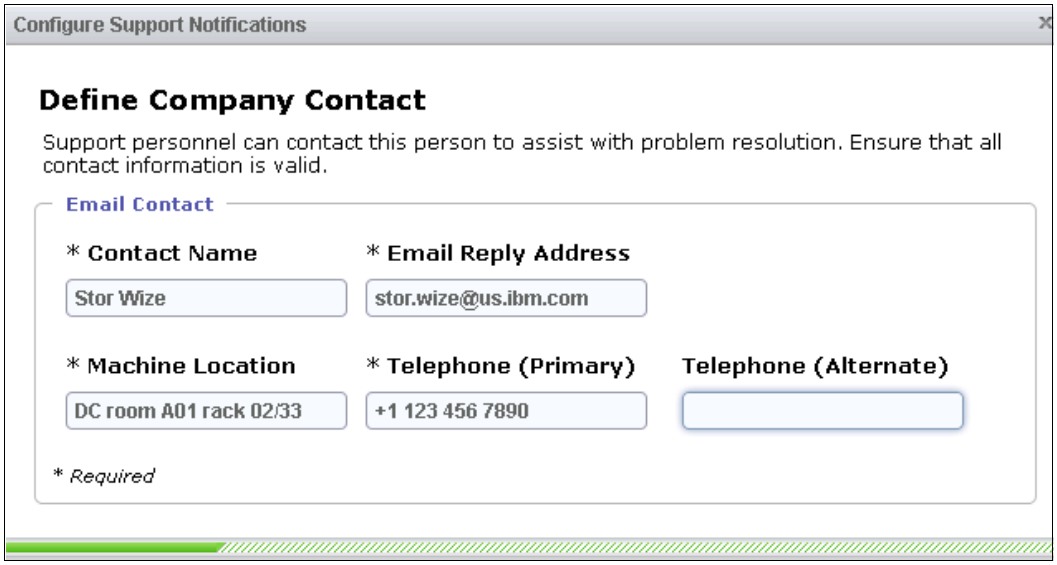

Figure 13-35 Contact information

4. Configure the Simple Mail Transfer Protocol (SMTP) server according to the instruction in Figure 13-36.

   You can also set up the **Inventory Reporting** function as and extension of the Home Call feature. Rather than reporting a problem, an email is sent to IBM that describes your system hardware and critical configuration information. Object names and other information, such as IP addresses, are not sent. The inventory email is sent on a regular basis. Based on the information that is received, IBM can inform you if the hardware or software that you are using requires an upgrade because of a known issue.

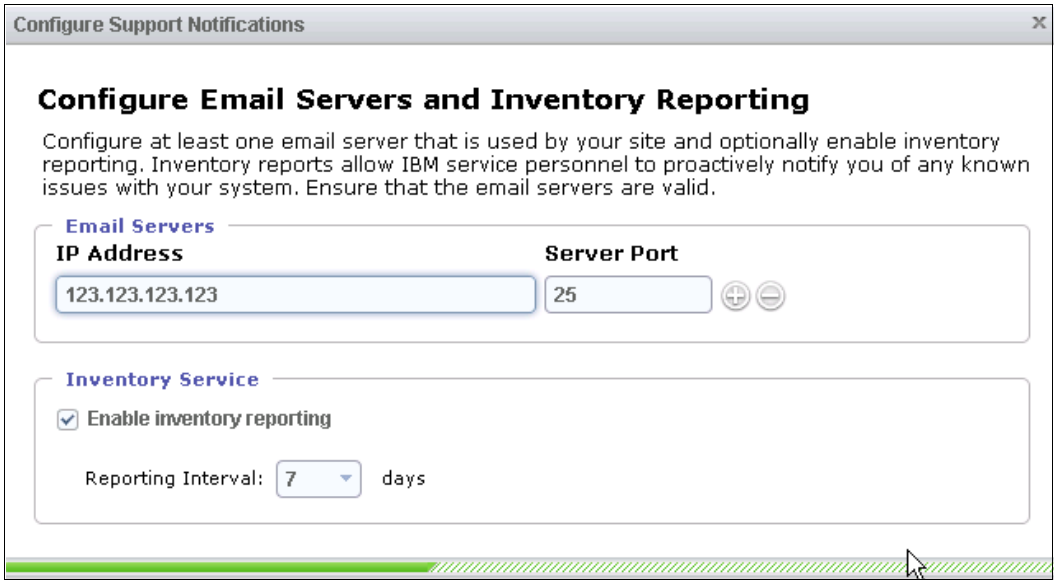

Figure 13-36 Configure email servers and inventory reporting

   The inventory reports can be submitted to the IBM support in the interval from 1 to 15 days. We suggest to keep the weekly default frequency.

5. In the next step, configure and confirm email addresses of the IBM support (drop-down menu **Support**) or local user (menu **Local**). See Figure 13-37 on page 486 for details. In this wizard you can set only one email address. Other recipients can be added afterwards.

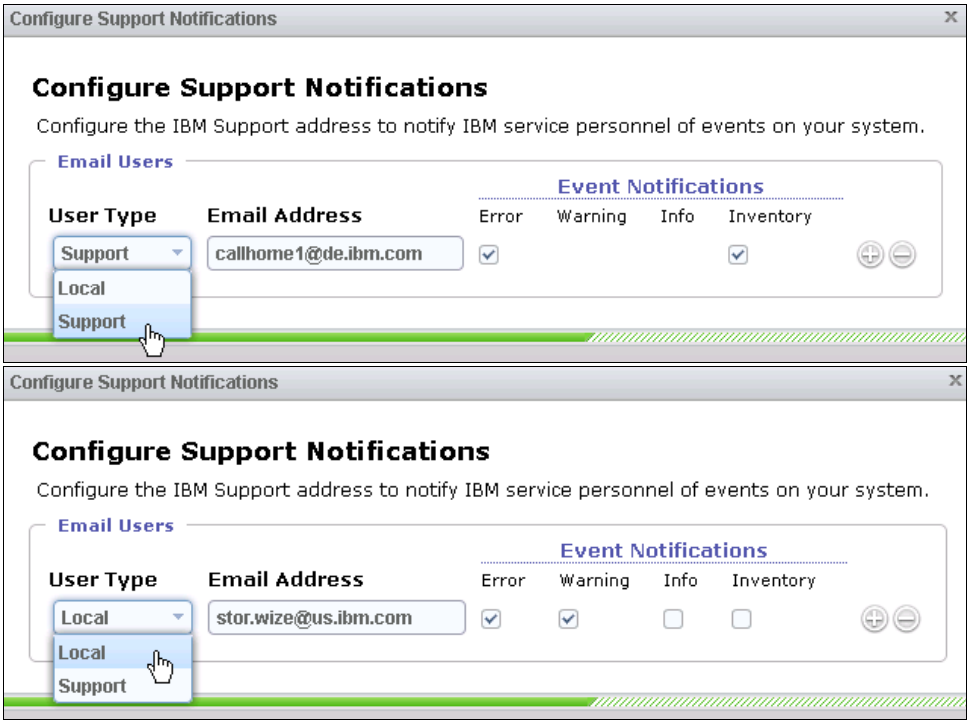

Figure 13-37 Setting email recipients

6. Complete the configuration wizard and test the email function. To do so, you have to enter **Edit** mode as illustrated in Figure 13-38. In the same window, you can define the additional email recipient, either from IBM support or local users.

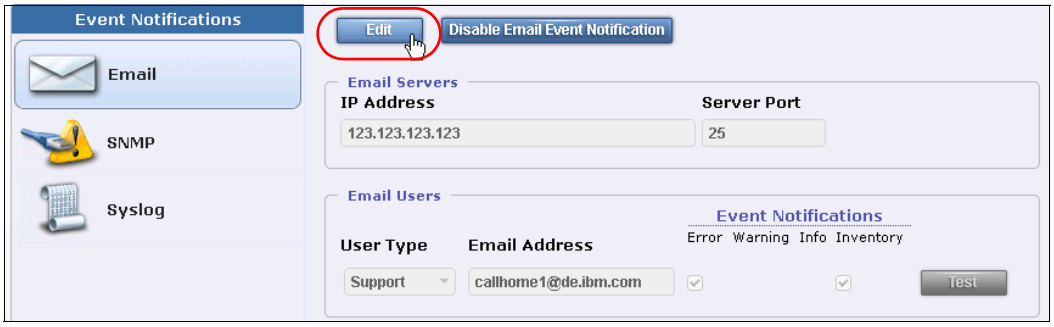

Figure 13-38 Entering edit mode

We strongly suggest to keep sending inventory enabled to at least IBM support. However, it might be beneficial to do the same for local users. The email output can serve as a basis for the client's inventory and asset management to keep track of all hardware devices installed in the environment.

7. In Edit mode, you are allowed to change any of the previously configured settings. When you are finished editing these parameters, have added more recipients, or just tested the connection, you can save configuration to make the changes take effect (Figure 13-39 on page 487).

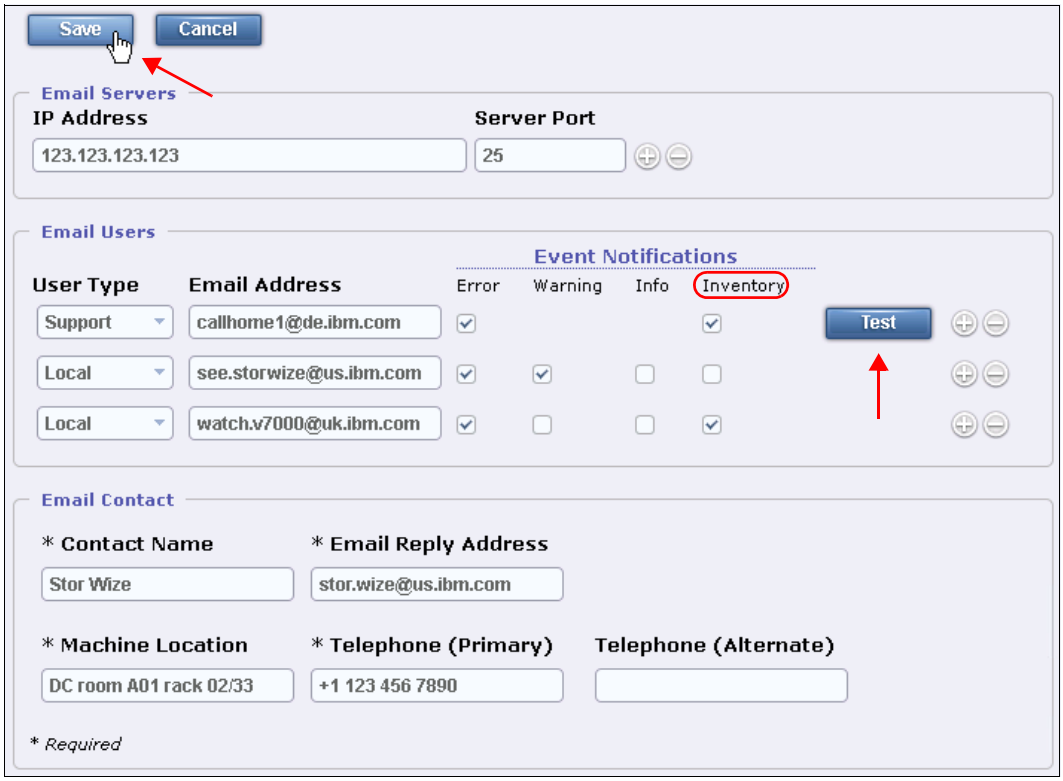

Figure 13-39 Saving modified configuration

Disabling and enabling notifications

At any time you can temporarily or permanently disable email notifications as shown in Figure 13-40. This is good practice when running maintenance tasks on your IBM Storwize V7000, such as code upgrade or replacement of malfunctioning parts. After the maintenance operation, do not forget to re-enable the email notification function. The same results can be achieved with the CLI **`svctask stopmail`** and **`svctask startmail`** commands.

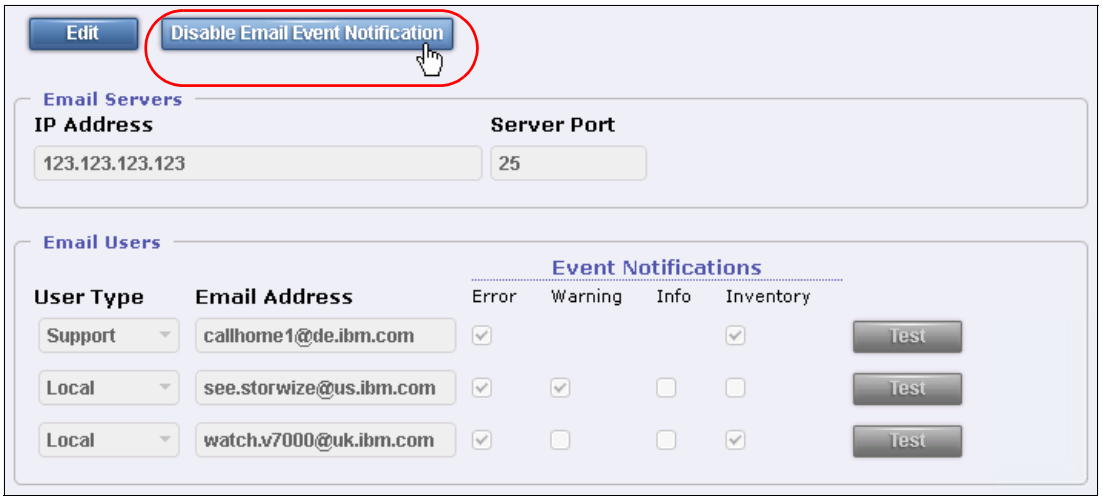

Figure 13-40 Disabling or enabling email notifications

13.8 Audit log

The audit log is useful when analyzing past configuration events, especially when trying to determine, for example, how a volume ended up being shared by two hosts, or why the volume was overwritten. The audit log is also included in the *svc_snap* support data to aid in problem determination.

An audit log tracks action commands that are issued through a Secure Shell (SSH) session or through the management GUI. It provides the following entries:

- Identity of the user who issued the action command
- Name of the actionable command
- Time stamp of when the actionable command was issued on the configuration node
- Parameters that were issued with the actionable command

The following items are not documented in the audit log:

- Commands that fail are not logged
- A result code of 0 (success) or 1 (success in progress) is not logged
- Result object ID of node type (for the addnode command) is not logged
- Views are not logged

Several specific service commands are not included in the audit log:

- `dumpconfig`
- `cpdumps`
- `cleardumps`
- `finderr`
- `dumperrlog`
- `dumpintervallog`
- `svcservicetak dumperrlog`
- `svcservicetask finderr`

Figure 13-41 shows the access to the Audit Log. Click **Audit Log** in the dynamic menu to see which configuration CLI commands have been run on the IBM Storwize V7000 system.

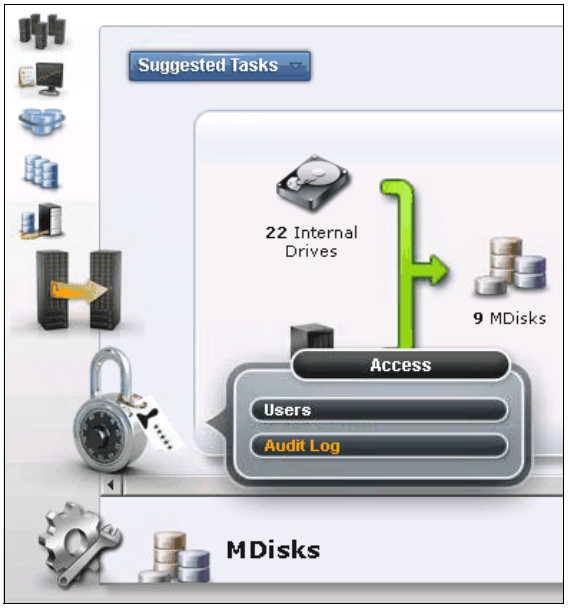

Figure 13-41 Audit Log from Access Management window

Figure 13-42 shows an example of the audit log after creating volumes and mapping them to hosts, with a command highlighted. The Running Tasks button is available at the bottom of the window in the status panel; if you click that button, the progress of the currently running tasks is displayed.

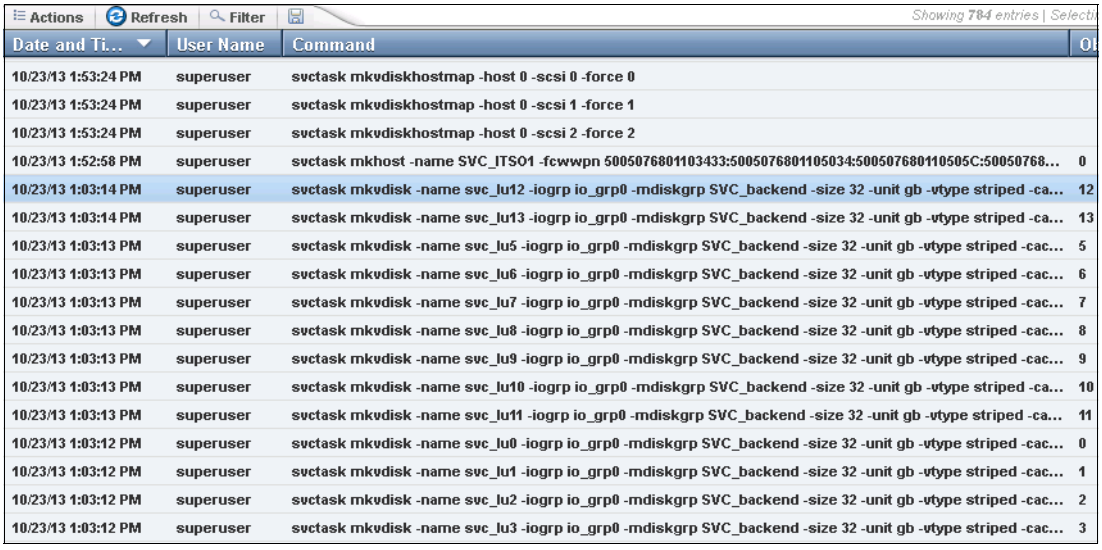

Figure 13-42 Audit log

Changing the view of the Audit Log grid by is also possible by right-clicking column headings (Figure 13-43). The grid layout and sorting is completely under the user's control, so you can view everything in the audit log, sort different columns, or reset the default grid preferences.

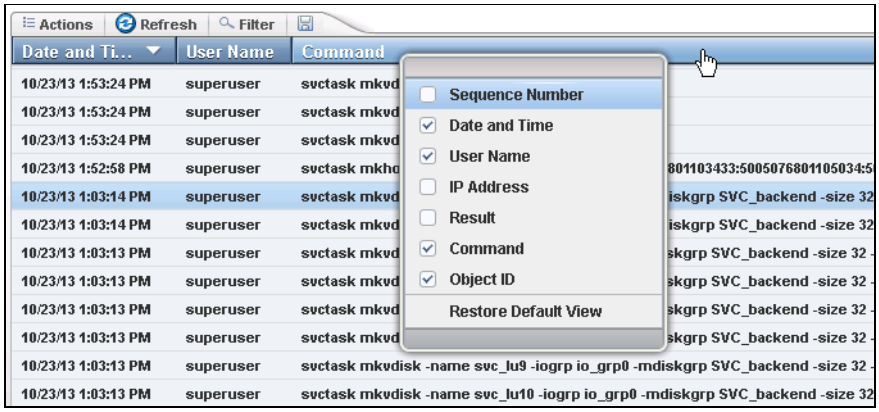

Figure 13-43 Right-click audit log column headings

13.9 Collecting support information

Occasionally, if you have a problem and call the IBM Support Center, they might ask you to provide support data. You can find this data under the Support tab of the Troubleshooting navigation window.

Click **Settings** and then the **Support** tab to begin the procedure of collecting support data (Figure 13-44). Assuming that the node restarts, use the menu shown in Figure 13-45 to collect the default logs plus all the existing statesaves to capture the maximum data for support.

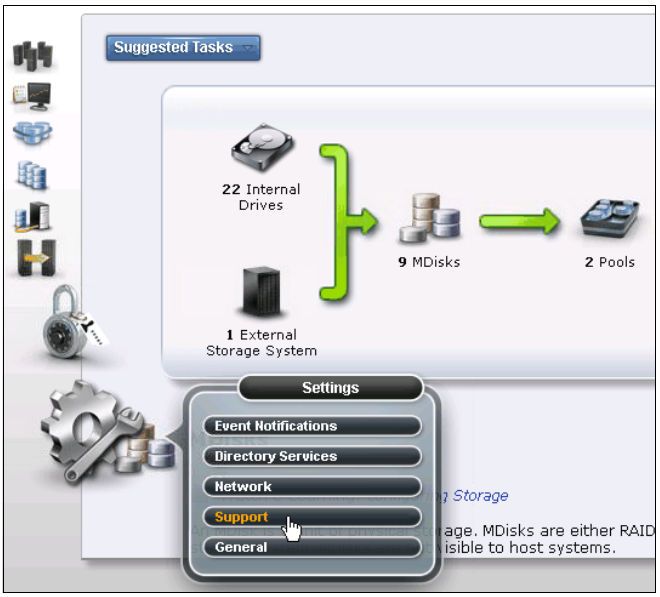

Figure 13-44 Support option Data Collection

Click **Download Support Package** (Figure 13-45). To list all individual log files, click the **Show All Individual Log Files** menu.

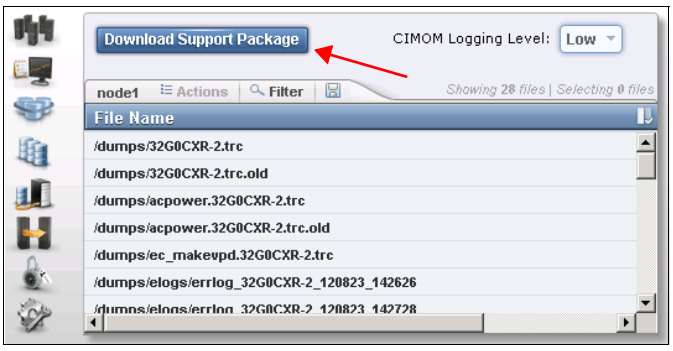

Figure 13-45 Download Support Package window

The window for collecting various versions of svc_snap opens (Figure 13-46); the version you download depends on the event that is investigated. For example, if you notice that a node was restarted in the event log, capture the snap with the latest existing statesave.

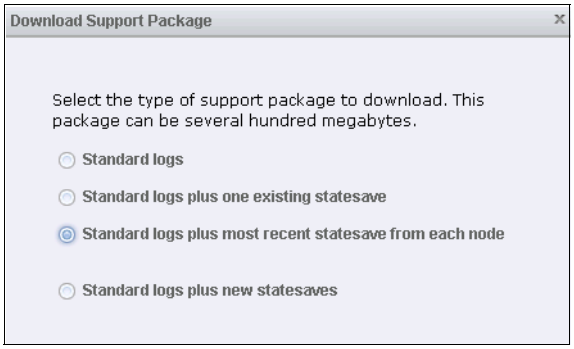

Figure 13-46 Download support package choices

The procedure to create the snap on the IBM Storwize V7000 system, including the latest statesave from each node canister, starts. This process might take a few minutes (Figure 13-47).

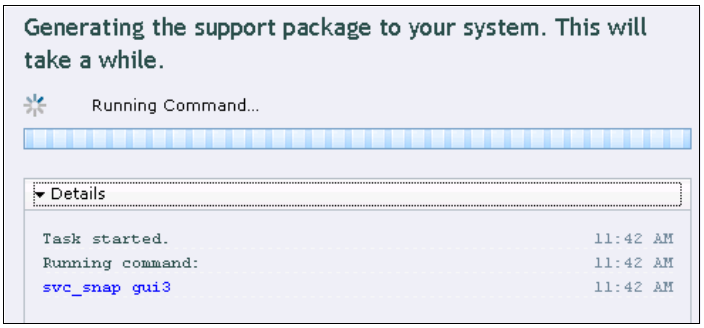

Figure 13-47 Task detail window

A window opens that gives you the choice to save the file on your local Windows system (Figure 13-48).

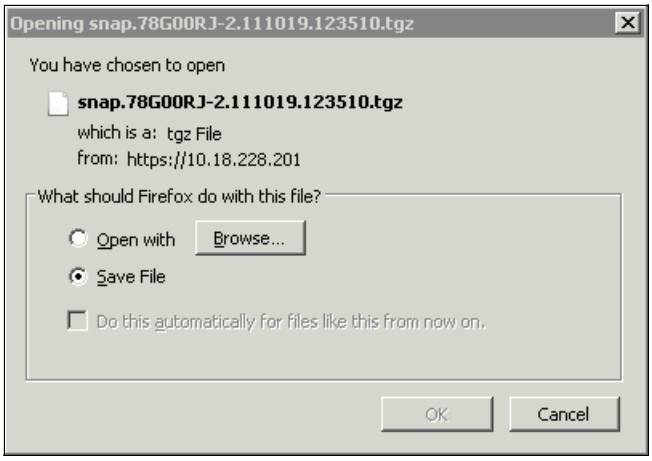

Figure 13-48 Save file to the local Windows system

Save the resulting snap file in a directory (Figure 13-49).

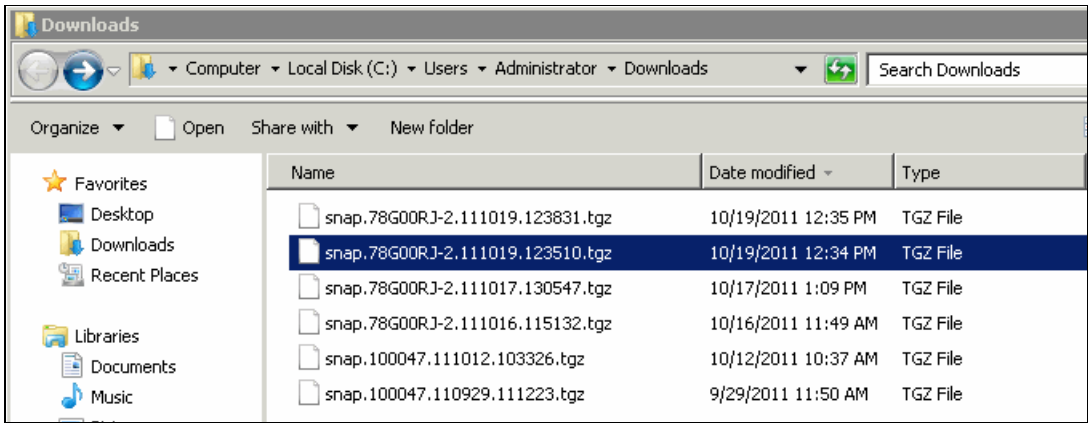

Figure 13-49 Save the resulting snap file to a local directory

Before you open a call with IBM support, be prepared to upload the resulting snap file to the IBM Support portal at the following address:

`http://www.ecurep.ibm.com/app/upload`

You are ready to call the IBM Support Line or use the IBM Support Portal to open a call. If you use the latter option, go to the following address:

`http://www.ibm.com/support/entry/portal/Open_service_request?brandind=Hardware`

13.10 Shutting down IBM Storwize V7000

You can safely shut down an IBM Storwize V7000 system using both the GUI or the CLI.

Important: Never shut down your IBM Storwize V7000 system by powering off the PSUs, removing both PSUs, or removing both power cables from a running system. It can lead to inconsistency or loss of the data staged in the cache.

Before shutting down the IBM Storwize V7000, stop all hosts that have allocated volumes from the device. This step can be skipped for hosts that have volumes that are also provisioned with mirroring (host-based mirror) from different storage device. However, doing so incurs errors that are related to lost storage paths/disks in the host error log.

You can shut down only one node canister or the entire cluster. When you shut down only one node canister, all activities remain active; when you shut down the entire cluster, you need to power on locally to start the system.

13.10.1 Shutting down a node canister

To shut down a single node canister using the GUI, complete the following steps:

1. Hover the cursor over the **Monitoring** function icon and click **System Details** (Figure 13-50).

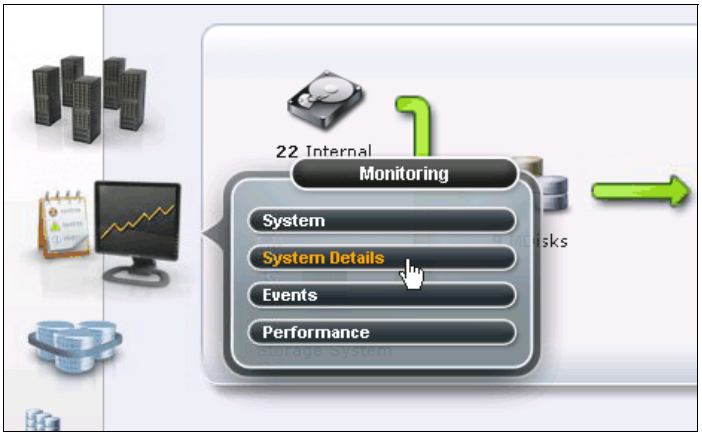

Figure 13-50 System Device option of the Monitoring function icon

2. Expand the **Canisters** section on the left and select the canister that you want to stop. Click **Actions** → **Shut Down Node Canister** (Figure 13-51).

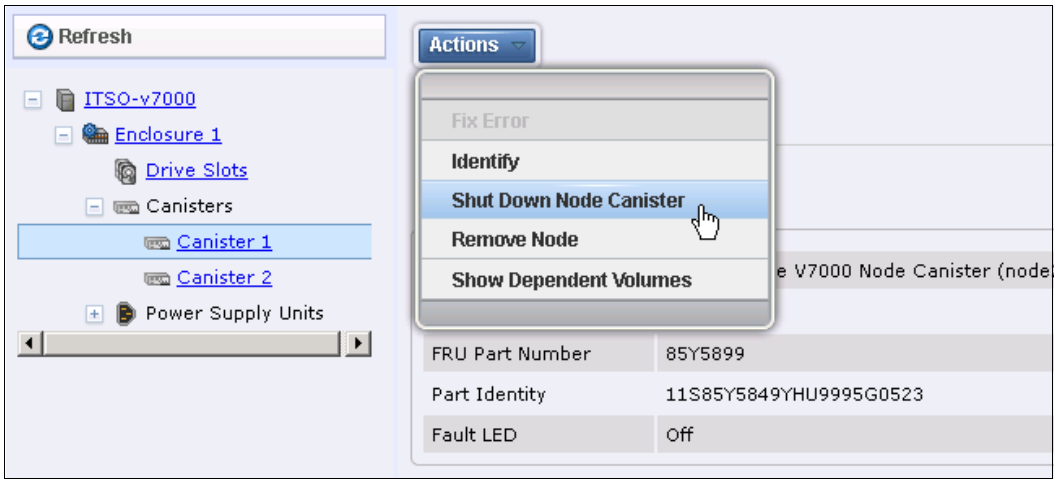

Figure 13-51 Shut Down Node Canister option

3. The confirmation window opens (Figure 13-52). Confirm whether you want to shut down the node. Type YES and click **OK**. If the node is active as a Configuration node, the control is moved automatically to the second node canister. Your session to the GUI will probably interrupt. Re-establish it again from the browser after takeover happens.

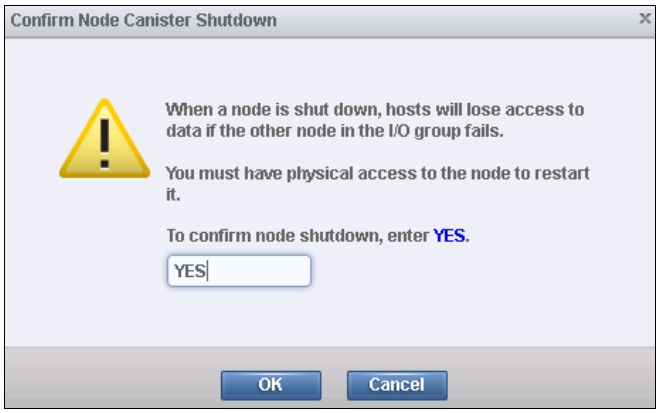

Figure 13-52 Confirm Shutdown window

Shut down is complete (Figure 13-53).

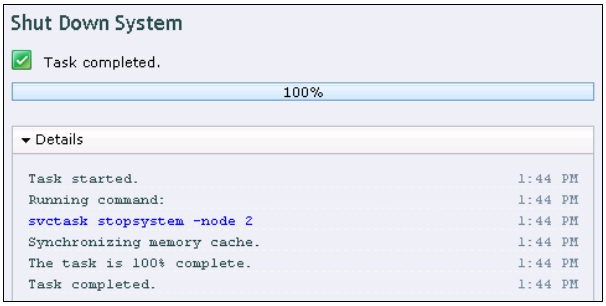

Figure 13-53 Shutdown complete

To shut down a node canister from the CLI, run the **svctask stopsystem -node 2** command.

13.10.2 Shutting down a system

The procedure to shut down a system is similar to shutting down a node canister. Instead of selecting a specific canister in the menu, select the whole Storwize V7000 system, and then from the **Actions** context menu, select **Shut Down System** (Figure 13-54).

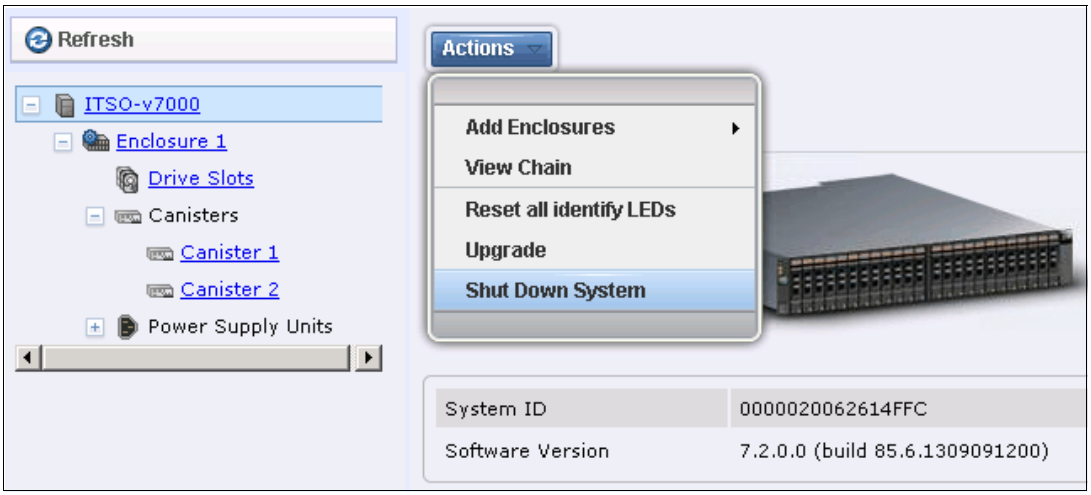

Figure 13-54 Shutdown of the IBM Storwize V7000

Confirm the validity of your decision to shut down the IBM Storwize V7000 clustered systems by typing `YES` in the pop-up window (Figure 13-55).

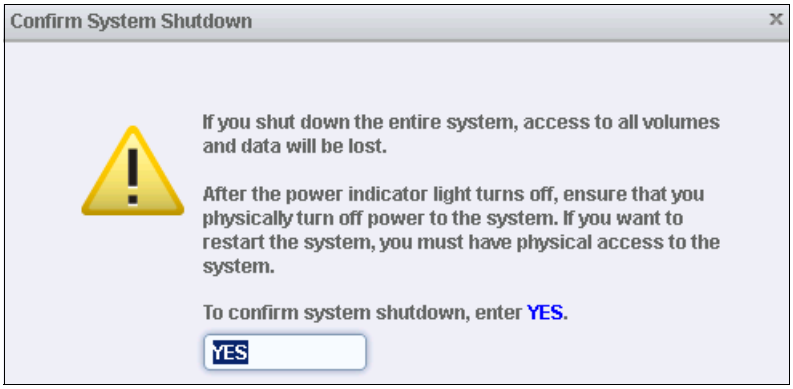

Figure 13-55 Shutdown confirmation

The whole system shutdown is typically planned in case of site maintenance (power outage, building construction, and so on), because all components of IBM Storwize V7000 are redundant and replaceable while the system is running. To start the device again after shutdown, you must have physical access to the system and then turn on the switches on the power supplies.

13.10.3 Shutting down and powering on an IBM Storwize V7000 infrastructure

When you shut down or power on the entire infrastructure (storage, servers, and applications), you must follow a particular sequence for both the shutdown and the power-on actions. Here is an example sequence of a shutdown, and then a power-on, of an infrastructure that includes an IBM Storwize V7000 system.

Figure 13-56 shows the location of power switches on the controller and expansion enclosures.

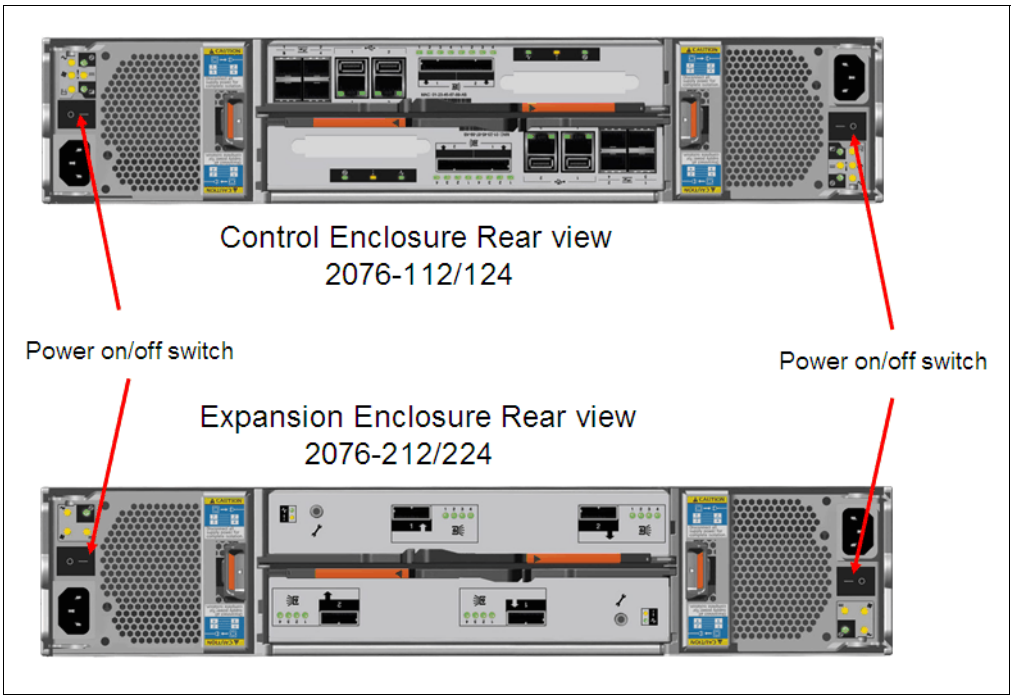

Figure 13-56 Location of power switches

Shutting down

To shut down the infrastructure, complete the following steps:

1. Shut down your servers and all applications.
2. Shut down your IBM Storwize V7000 system:
   a. Shut down the cluster using either the GUI or CLI.
   b. Power off both switches of the controller enclosure.
   c. Power off both switches of all the expansion enclosures.
3. Shut down your SAN switches.

Powering on

To power on your infrastructure, complete the following steps:

1. Power on your SAN switches and wait until the boot completes.
2. Power on your storage systems and wait until the systems are up, and then:
   a. Power on both switches of all the expansion enclosures.
   b. Power on both switches of the controller enclosure.
3. Power on your servers and start your applications.

A

CLI setup and SAN Boot

This appendix describes the setup of the CLI and provides extra information about the SAN Boot function.

Command-line interface (CLI)

The IBM Storwize V7000 system has a powerful CLI, which offers even more functions than the GUI. This section is not intended to be a detailed guide to the CLI, as that topic is beyond the scope of this book. The basic configuration of the IBM Storwize V7000 CLI and some example commands are covered. However, the CLI commands are the same as in the SAN Volume Controller, and in addition, there are more commands that are available to manage internal storage. If a task completes in the GUI, the CLI command is always displayed in the details, as shown throughout this book.

Detailed CLI information is available at the IBM Storwize V7000 web page under the command-line section, which is at the following address:

http://www.ibm.com/support/knowledgecenter/ST3FR7_7.2.0/com.ibm.storwize.v7000.720.doc/svc_clicommandscontainer_229g0r.html?lang=en

Implementing the IBM System Storage SAN Volume Controller V7.2, SG24-7933 provides more detailed information about using the CLI. The commands in that book also apply to the IBM Storwize V7000 system.

Basic setup

In the IBM Storwize V7000 GUI, authentication is done by using a user name and password. The CLI uses a Secure Shell to connect from the host to the IBM Storwize V7000 system. As of IBM Storwize V7000 V6.3, either a private and a public key pair or user name and password is necessary. The following steps are required to enable CLI access with SSH keys:

- A public key and a private key are generated together as a pair.
- A public key is uploaded to the IBM Storwize V7000 system through the GUI.
- A client Secure Shell (SSH) tool must be configured to authenticate with the private key.
- A secure connection can be established between the client and IBM Storwize V7000.

Secure Shell is the communication vehicle between the management workstation and the IBM Storwize V7000 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 V7000 systems.

The SSH client used in this book is PuTTY. Also, a PuTTY key generator can be used to generate the private and public key pair. The PuTTY client can be downloaded from the following address at no cost:

http://www.chiark.greenend.org.uk

Download the following tools:

- PuTTY SSH client: `putty.exe`
- PuTTY key generator: `puttygen.exe`

Generating a public and private key pair

To generate a public and private key pair, complete the following steps:

1. Start the PuTTY key generator to generate the public and private key pair (Figure A-1).

Figure A-1 PuTTY key generator

Make sure that the following options are selected:

- SSH2 RSA
- Number of bits in a generated key: 1024

2. Click **Generate** and move the cursor over the blank area to generate keys (Figure A-2).

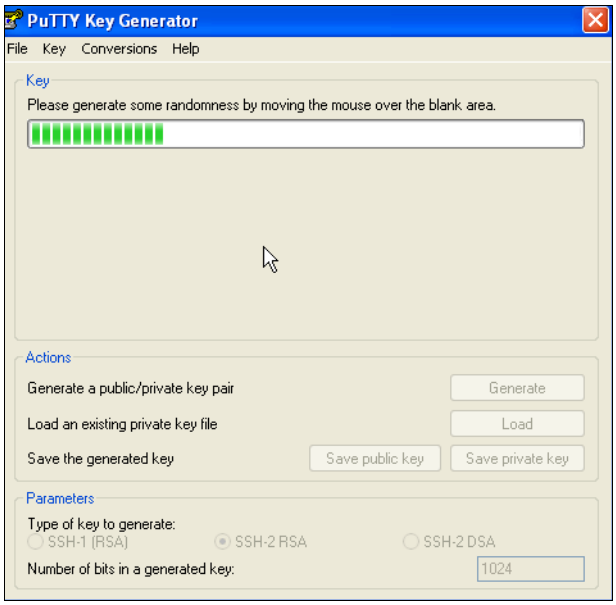

Figure A-2 Generate keys

To generate keys: The blank area that is indicated by the message is the large blank rectangle on the GUI inside the section of the GUI labeled Key. Continue to move the mouse pointer over the blank area until the progress bar reaches the far right. This action generates random characters to create a unique key pair.

3. After the keys are generated, save them for later use. Click **Save public key** (Figure A-3).

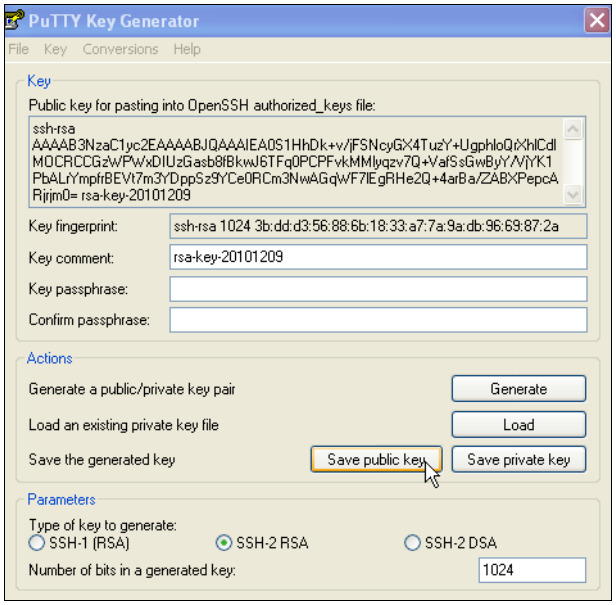

Figure A-3 Save public key

4. You are prompted for a name (for example, pubkey) and a location for the public key (for example, `C:\Support Utils\PuTTY`). Click **Save**.

   Ensure that you record the name and location, because the name and location of this SSH public key must be specified later.

Public key extension: By default, the PuTTY key generator saves the public key with no extension. Use the string “pub” for naming the public key, for example, “pubkey”, to easily differentiate the SSH public key from the SSH private key.

5. Click **Save private key** (Figure A-4).

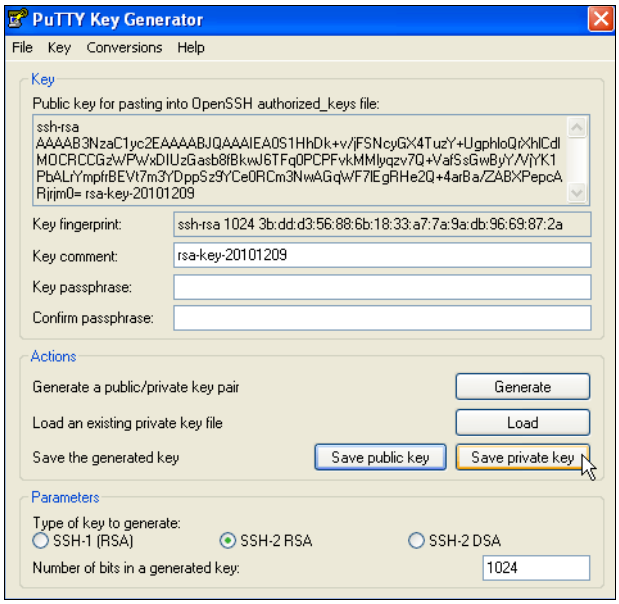

Figure A-4 Save private key

6. You are prompted with a warning message (Figure A-5). Click **Yes** to save the private key without a passphrase.

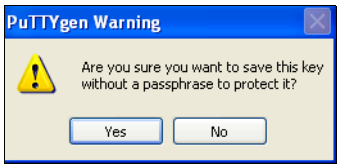

Figure A-5 Confirm the security warning

7. When prompted, enter a name (for example, “icat”), select a secure place as the location, and click **Save**.

> **Key generator:** The PuTTY key generator saves the private key with the PPK extension.

8. Close the PuTTY key generator.

Uploading the SSH public key to the IBM Storwize V7000

After you create your SSH key pair, upload your SSH public key onto the SAN Volume Controller system. Complete the following steps:

1. Open the user section (Figure A-6).

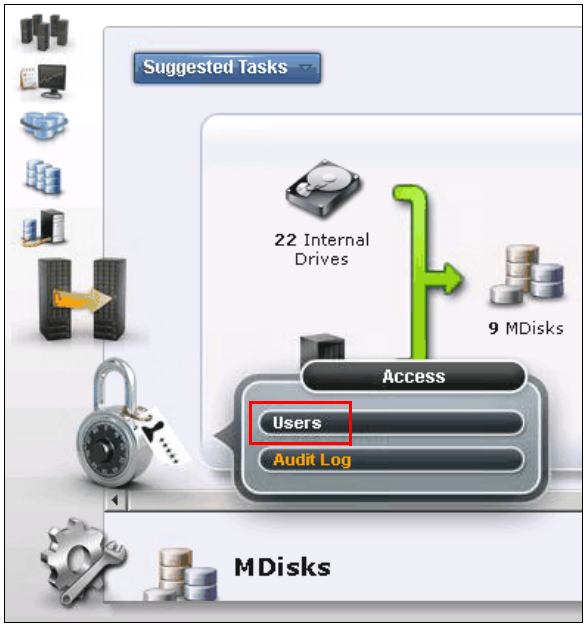

Figure A-6 Open user section

2. Right-click the user name for which you want to upload the key and click **Properties** (Figure A-7).

Figure A-7 Superuser properties

3. To upload the public key, click **Browse**, select your public key, and click **OK** (Figure A-8).

Figure A-8 Select public key

4. Click **OK** and the key is uploaded. Click **Close** to close the wizard and return to the GUI.

Configuring the SSH client

Before the CLI can be used, the SSH client must be configured as follows:

1. Start PuTTY. The PuTTY Configuration window opens (Figure A-9).

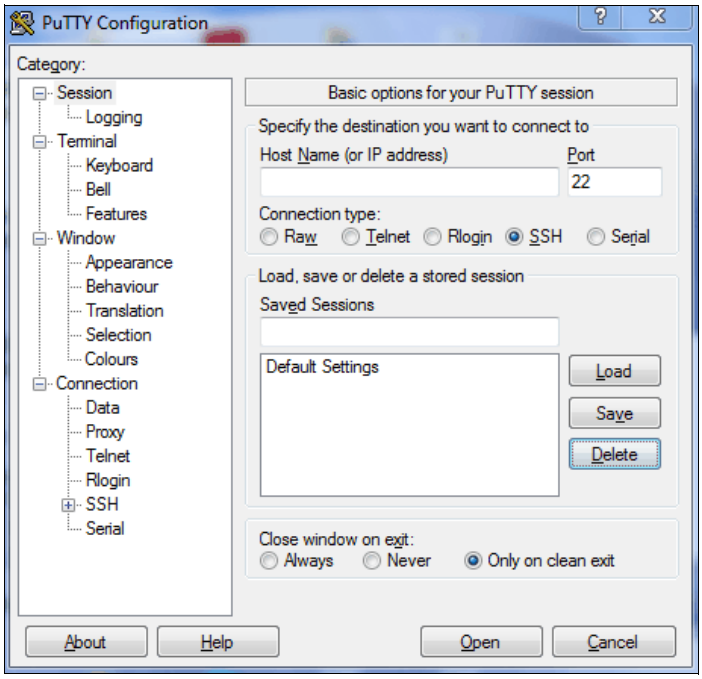

Figure A-9 PuTTY

In the right pane, select **SSH** as the connection type. Under the “Close window on exit” section, select **Only on clean exit**, which ensures that if any connection errors occur, they are displayed on the user’s window.

2. In the Category pane, on the left side of the PuTTY Configuration window (Figure A-10), click **Connection** → **SSH** to open the PuTTY SSH Configuration window.

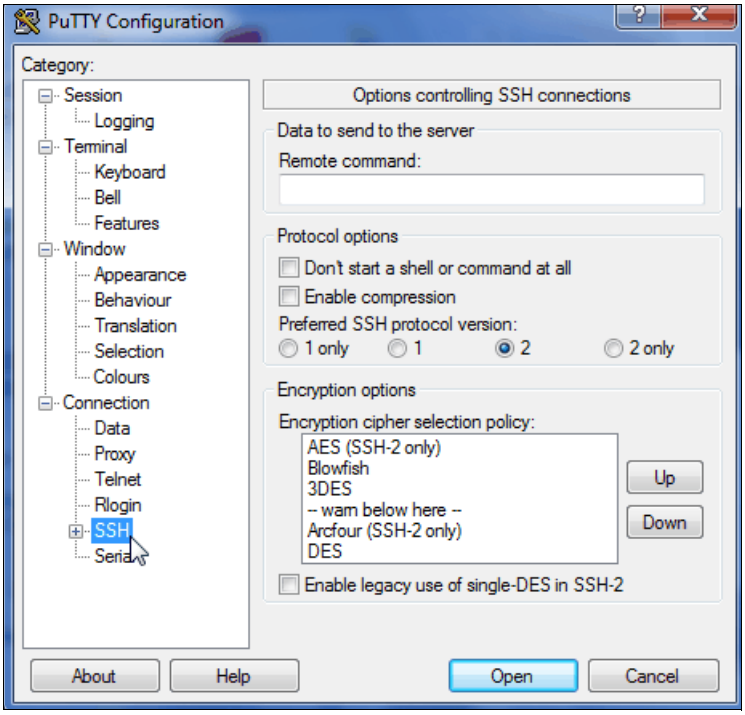

Figure A-10 SSH protocol version 2

3. Under the “Preferred SSH protocol version” section, select **2**.
4. In the Category pane on the left, click **Connection** → **SSH** → **Auth**. More options are displayed for controlling SSH authentication.

In the **Private key file for authentication** field (Figure A-11), either browse to or type the fully qualified directory path and file name of the SSH client private key file, which was created previously (for example, `C:\Support Utils\PuTTY\icat.PPK`).

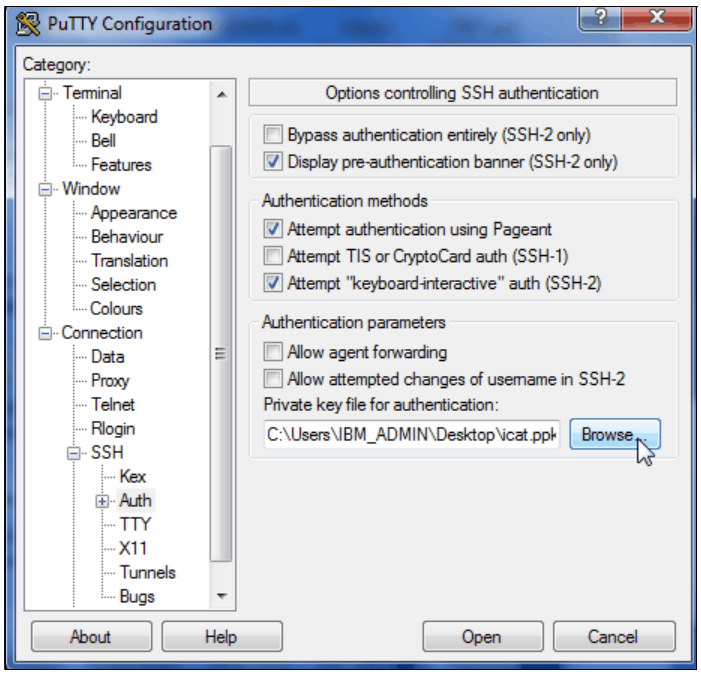

Figure A-11 SSH authentication

5. In the Category pane, click **Session** to return to the Basic options for your PuTTY session view (Figure A-12 on page 506).

6. Enter the following information in these fields (Figure A-12) in the right pane:
   - Host Name: Specify the host name or system IP address of the IBM Storwize V7000 clustered system.
   - Saved Sessions, and enter a session name.

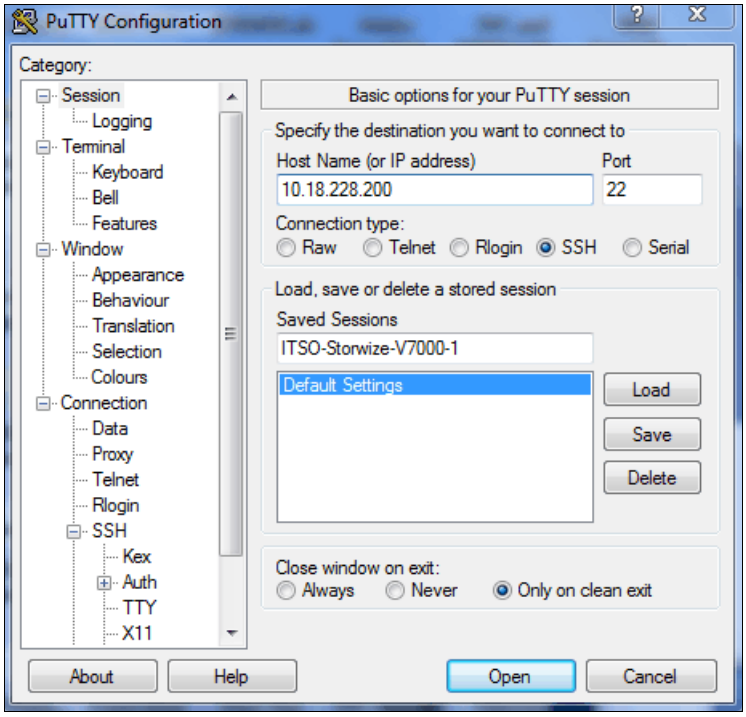

Figure A-12 Enter session information

7. Click **Save** to save the new session (Figure A-13).

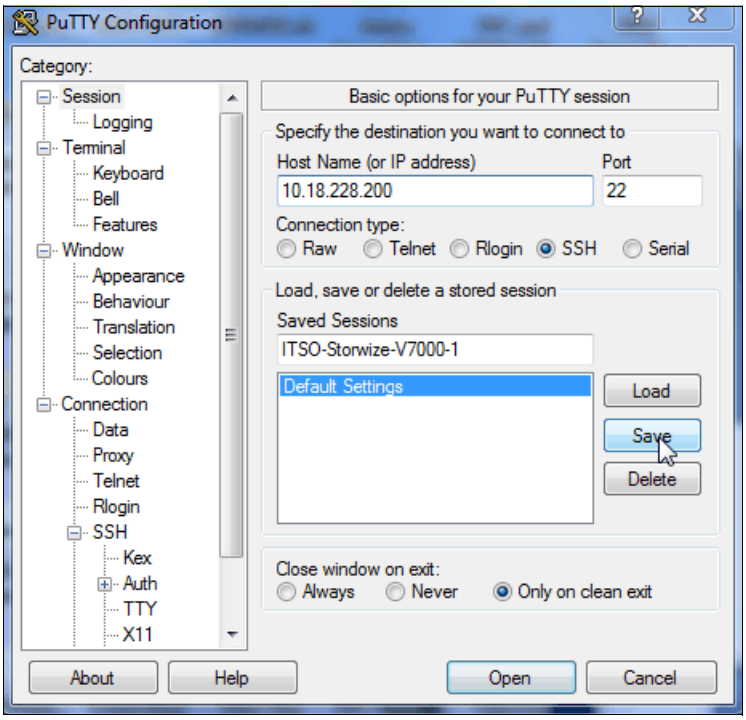

Figure A-13 Save the new session

8. Select the new session and click **Open** to connect to the IBM Storwize V7000 system. A PuTTY Security Alert opens; confirm it by clicking **Yes** (Figure A-14).

Figure A-14 Confirm Security Alert

9. PuTTY now connects to the system and prompts you for a user name to log in as. Enter `admin` as the user name (Example A-1) and press Enter.

Example: A-1 Enter user name

```
login as: admin
Authenticating with public key "rsa-key-20101006"
IBM_2076:ITSO-Storwize-V7000-1:admin>
```

You have now completed the tasks to configure the CLI for IBM Storwize V7000 administration.

SAN Boot

IBM Storwize V7000 supports SAN Boot for Windows, VMware, and many other operating systems. SAN Boot support can change, so regularly check the IBM Storwize V7000 interoperability matrix at this address:

http://www.ibm.com/systems/storage/disk/storwize_v7000/interop.html

The IBM Knowledge Center for Storwize V7000 has much information about SAN Boot in combination with various operating systems. For more information, go to this address:

http://publib.boulder.ibm.com/infocenter/storwize/ic/index.jsp

More information about SAN Boot is also in the *IBM Multipath Subsystem Device Driver User's Guide*, which is available at the following address:

http://www.ibm.com/support/docview.wss?rs=503&context=HW26L&uid=ssg1S7000303

Enabling SAN Boot for Windows

Complete the following procedure if you want to install Windows host using SAN Boot:

1. Configure the IBM Storwize V7000 system so that only the boot volume is mapped to the host.
2. Configure the Fibre Channel SAN so that the host only sees one IBM Storwize V7000 system node port. Multiple paths during installation are not supported.
3. Configure and enable the HBA BIOS.
4. Install the operating system using the normal procedure, selecting the volume as the partition on which to install.

   HBAs: You might need to load an additional HBA device driver during installation, depending on your Windows version and the HBA type.

5. Install SDDDSM after the installation has completed.
6. Modify your SAN zoning to allow multiple paths.
7. Check your host to see if all paths are available.
8. Set redundant boot devices in the HBA BIOS to allow the host to boot when its original path has failed.

Enabling SAN Boot for VMware

Complete the following steps if you want to install a VMware ESXhost using SAN Boot:

1. Configure the IBM Storwize V7000 system so that only the boot volume is mapped to the host.
2. Configure the Fibre Channel SAN so that the host only sees one IBM Storwize V7000 system node port. Multiple paths during installation are not supported.
3. Configure and enable the HBA BIOS.

4. Install the operating system using the normal procedure, selecting the volume as the partition on which to install.

   HBAs: You might need to load an additional HBA device driver during installation, depending on your ESX level and the HBA type.

5. Modify your SAN zoning to allow multiple paths.
6. Check your host if all paths are available and modify the multipath policy if required.

Windows SAN Boot migration

If you have a host that runs a Windows 2000 Server, Windows Server 2003, or Windows Server 2008 operating system, and have existing SAN Boot images that are controlled by storage controllers, you can migrate these images to image-mode volumes that are controlled by the IBM Storwize V7000 system.

SAN Boot procedures: For SAN Boot procedures for other operating systems, check the IBM Knowledge Center for Storwize V7000:

`http://publib.boulder.ibm.com/infocenter/storwize/ic/index.jsp`

Complete the following steps to migrate your existing SAN Boot images:

1. If the existing SAN Boot images are controlled by an IBM storage controller that uses the IBM Subsystem Device Driver (SDD) as the multipathing driver, you must use SDD V1.6 or later. Run the SDD **`datapath set bootdiskmigrate 2076`** command to prepare the host for image migration. See the Multipath Subsystem Device Driver (SDD) documentation for more information.
2. Shut down the host.
3. Complete the following configuration changes on the storage controller:
   a. Write down the SCSI LUN ID each volume is using (for example, boot LUN SCSI ID 0, Swap LUN SCSI ID 1, Database Lun SCSID 2, and so on).
   b. Remove all the image-to-host mappings from the storage controller.
   c. Map the existing SAN Boot image and any other disks to the IBM Storwize V7000 system.
4. Change the zoning so that the host is able to see the IBM Storwize V7000 I/O group for the target image mode volume.
5. Complete the following configuration changes on the IBM Storwize V7000 system:
   a. Create an image mode volume for the managed disk (MDisk) that contains the SAN Boot image. Use the MDisk unique identifier to specify the correct MDisk.
   b. Create a host object and assign the host HBA ports.
   c. Map the image mode volume to the host using the same SCSI ID as before. For example, you might map the boot disk to the host with SCSI LUN ID 0.
   d. Map the swap disk to the host, if required. For example, you might map the swap disk to the host with SCSI LUN ID 1.

6. Change the boot address of the host by completing the following steps:
   a. Restart the host and open the HBA BIOS utility of the host during the booting process.
   b. Set the BIOS settings on the host to find the boot image at the worldwide port name (WWPN) of the node that is zoned to the HBA port.
7. If SDD V1.6 or later is installed and you ran **bootdiskmigrate** in step 1 on page 509, reboot your host, update SDDDSM to the latest level, and go to step 14. If SDD V1.6 is not installed, go to step 8.
8. Modify the SAN Zoning so that the host only sees one path to the IBM Storwize V7000.
9. Boot the host in single-path mode.
10. Uninstall any multipathing driver that is not supported for IBM Storwize V7000 system hosts that run the applicable Windows Server operating system.
11. Install SDDDSM.
12. Restart the host in single-path mode and ensure that SDDDSM was properly installed.
13. Modify the SAN Zoning to enable multipathing.
14. Rescan drives on your host and check that all paths are available.
15. Reboot your host and enter the HBA BIOS.
16. Configure the HBA settings on the host. Ensure that all HBA ports are boot-enabled and can see both nodes in the I/O group that contains the SAN Boot image. Configure the HBA ports for redundant paths.
17. Exit the BIOS utility and finish starting the host.
18. Map any additional volumes to the host as required.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- *Get More Out of Your SAN with IBM Tivoli Storage Manager*, SG24-6687
- *IBM Storwize V7000 and SANSlide Implementation*, REDP-5023
- *IBM System Storage Business Continuity: Part 1 Planning Guide*, SG24-6547
- *IBM System Storage Business Continuity: Part 2 Solutions Guide*, SG24-6548
- *IBM System Storage SAN Volume Controller and Storwize V7000 Replication Family Services*, SG24-7574
- *IBM System Storage SAN Volume Controller and Storwize V7000 Best Practices and Performance Guidelines*, SG24-7521
- *IBM Tivoli Storage Area Network Manager: A Practical Introduction*, SG24-6848
- *IBM Tivoli Storage Productivity Center V5.1 Technical Guide*, SG24-8053
- *Implementing an IBM b-type SAN with 8 Gbps Directors and Switches*, SG24-6116
- *Implementing IBM Easy Tier with IBM Real-time Compression*, TIPS1072
- *Implementing IBM Real-time Compression in SAN Volume Controller and IBM Storwize V7000*, TIPS1083
- *Implementing the IBM System Storage SAN Volume Controller V7.2*, SG24-7933
- *Introduction to Storage Area Networks and System Networking*, SG24-5470
- *Using the SVC for Business Continuity*, SG24-7371

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Other publications

These publications are also relevant as further information sources:

- *Command Line Interface User's Guide,* SG26-7903
- *IBM System Storage Master Console: Installation and User's Guide*, GC30-4090
- *IBM System Storage Open Software Family SAN Volume Controller: CIM Agent Developers Reference*, SC26-7545
- *IBM System Storage Open Software Family SAN Volume Controller: Command-Line Interface User's Guide*, SC26-7544

- *IBM System Storage Open Software Family SAN Volume Controller: Configuration Guide*, SC26-7543
- *IBM System Storage Open Software Family SAN Volume Controller: Host Attachment Guide*, SC26-7563
- *IBM System Storage Open Software Family SAN Volume Controller: Installation Guide*, SC26-7541
- *IBM System Storage Open Software Family SAN Volume Controller: Planning Guide*, GA22-1052
- *IBM System Storage Open Software Family SAN Volume Controller: Service Guide*, SC26-7542

Online resources

These websites are also relevant as further information sources:

- Open source site for SSH for Windows and Mac:

  http://www.openssh.com/windows.html

- SAN Volume Controller supported platform:

  http://www.ibm.com/support/entry/portal/product/system_storage/storage_software/storage_virtualization/san_volume_controller_%282145%29?productContext=-1948454624

- Subsystem Device Driver download site:

  http://www.ibm.com/support/entry/portal/product/system_storage/disk_systems/mid-range_disk_systems/ibm_storwize_v7000_%282076%29?productContext=-1546771614

- SAN Volume Controller online documentation:

  http://www.ibm.com/support/knowledgecenter/STPVGU/com.ibm.storage.svc.console.640.doc/mlt_relatedinfo_224agr.html

- Sysinternals home page:

  http://www.sysinternals.com

- VMware information:

  http://www.vmware.com/products/vsphere/features/vmotion.html

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

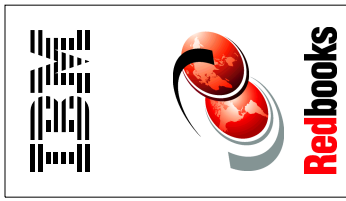

Implementing the IBM Storwize V7000 V7.2

(1.0” spine)
0.875”<->1.498”
460 <-> 788 pages

Implementing the IBM Storwize V7000 V7.2

Discover the exciting new version of the IBM virtualization family

Become familiar with the leading edge intuitive GUI

See how simple IBM virtualization really is

Continuing its commitment to developing and delivering industry-leading storage technologies, IBM introduces the IBM Storwize V7000 solution, an innovative new storage offering that delivers essential storage efficiency technologies and exceptional ease of use and performance, all integrated into a compact, modular design that is offered at a competitive, midrange price.

The IBM Storwize V7000 solution incorporates some of the top IBM technologies typically found only in enterprise-class storage systems, raising the standard for storage efficiency in midrange disk systems. This cutting-edge storage system extends the comprehensive storage portfolio from IBM and can help change the way organizations address the ongoing information explosion.

This IBM Redbooks publication introduces the features and functions of the IBM Storwize V7000 system through several examples. This book is aimed at pre- and post-sales technical support and marketing, storage administrators, and will help you understand the architecture of the Storwize V7000, how to implement it, and take advantage of the industry leading functions and features.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-7938-02 ISBN 0738439614