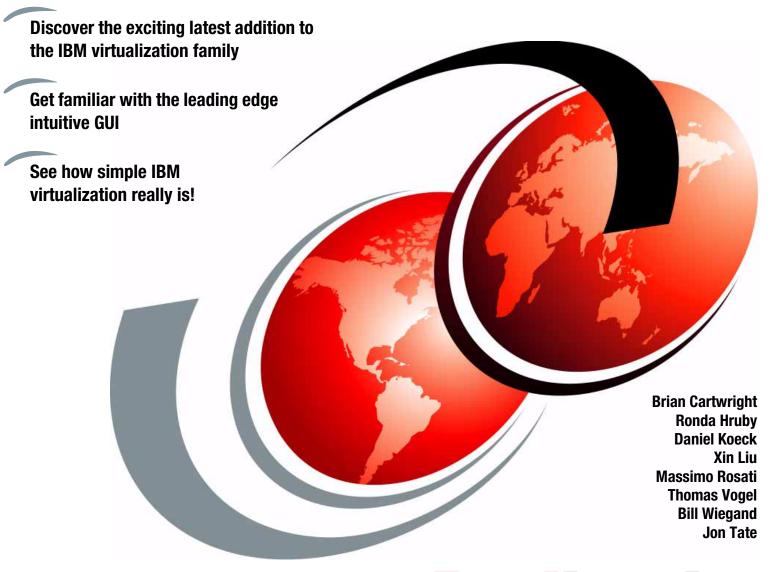


SG24-7938-00

Implementing the IBM Storwize V7000



Redbooks



International Technical Support Organization

Implementing the IBM Storwize V7000

December 2010

Note: Before using this information and the product it supports, read the information in "Notices" on page xiii.

First Edition (December 2010)

This edition applies to the IBM Storwize V7000 6.1.0.

This document created or updated on March 3, 2011.

© Copyright International Business Machines Corporation 2010. All rights reserved.

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

Contents

| Trademarks | |
|---|---------|
| Preface | |
| The team who wrote this book | XV |
| Now you can become a published author, too! | . xviii |
| Comments welcome | xix |
| Stay connected to IBM Redbooks | xix |
| Chapter 1. Overview of the IBM Storwize V7000 | -1 |
| 1.1 Storage Virtualization | |
| 1.2 IBM Storwize V7000 Overview | |
| 1.3 IBM Storwize V7000 terminology | |
| 1.3.1 IBM Storwize V7000 models | |
| 1.3.2 IBM Storwize V7000 attributes | |
| 1.3.3 IBM Storwize V7000 Functions | |
| 1.3.4 IBM Storwize V7000 Licensing | |
| 1.4 IBM Storwize V7000 hardware | |
| 1.4.1 Control enclosure | |
| 1.4.2 Expansion enclosure | |
| 1.4.3 Disk Drive Types | |
| 1.5 IBM Storwize V7000 Components | |
| 1.5.1 Hosts | |
| 1.5.2 Nodes | |
| 1.5.3 I/O Group | |
| 1.5.4 Cluster | |
| 1.5.5 RAID | |
| 1.5.6 Managed Disks | |
| 1.5.7 Quorum disks | |
| 1.5.8 Storage pools | |
| 1.5.9 Volumes | |
| 1.5.10 Thin provisioned volumes | |
| 1.5.11 Mirrored Volumes | |
| 1.5.12 Easy Tier | |
| 1.5.13 iSCSI | |
| 1.6 Advanced Copy Services | |
| 1.6.1 Synchronous/Asynchronous remote copy. | |
| 1.6.2 FlashCopy | |
| 1.6.3 Copy Services configuration limits | |
| 1.7 Management and Support Tools | |
| 1.8 Useful Storwize V7000 Web Links | |
| | |
| 1.8.1 IBM Storwize V7000 learning videos on youtube | 32 |
| Chapter 2. Initial configuration | 35 |
| 2.1 Hardware Installation Planning | 36 |
| 2.2 SAN Configuration Planning | 38 |
| 2.3 LAN Configuration Planning | 41 |
| 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) | |
| 2.6.2 Command line interface | |
| 2.7 First time setup | |
| 2.8 Initial configuration | |
| 2.8.1 Configure Call Home, E-Mail alert and Inventory | |
| 2.0.1 Comigure Gail Floride, E Mail alort and inventory | J-T |
| Chapter 3. Graphical User Interface overview | 63 |
| 3.1 Overview of IBM Storwize V7000 management software | |
| 3.2 Home menu | |
| 3.3 Troubleshooting menu | |
| 3.4 Physical Storage menu | |
| 3.5 Volumes menu | |
| | |
| 3.6 Hosts menu | |
| 3.7 Copy Services menu | |
| 3.8 User Management menu | |
| 3.9 Configuration menu | 20 |
| | |
| Chapter 4. Host Configuration | |
| 4.1 Preparing the Host Operating System | |
| 4.1.1 Windows 2008 - Preparing for Fibre Channel Attachment | |
| 4.1.2 Windows 2008 - preparing for iSCSI Attachment | 32 |
| 4.1.3 VMware ESX - Preparing for Fibre Channel Attachment | 35 |
| 4.1.4 VMware ESX - Preparing for iSCSI Attachment | 39 |
| 4.2 Creating Hosts using the GUI | |
| 4.2.1 Creating FC Hosts | |
| 4.2.2 Creating iSCSI Hosts | |
| | |
| Chapter 5. Basic Volume Configuration | 57 |
| 5.1 Provisioning storage from the IBM Storwize V7000 and making it available for the host | |
| 158 | |
| 5.1.1 Creating a Generic Volume | 60 |
| 5.1.2 Creating a Thin-provisioned Volume | |
| 5.1.3 Creating a Mirrored Volume | |
| 5.1.4 Creating a Thin-mirror Volume | |
| 5.1.4 Cleaning a Thirl-Hillor Volume | |
| | |
| 5.2.1 Mapping newly created Volumes to the Host using the wizard | |
| 5.3 Discover the volumes from the host and specify multipath settings | |
| 5.3.1 Windows 2008 Fibre Channel volume attachment | |
| 5.3.2 Windows 2008 iSCSI Attachment | |
| 5.3.3 VMware ESX Fibre Channel Attachment | 83 |
| 5.3.4 VMware ESX iSCSI Attachment | 89 |
| | |
| | 99 |
| 6.1 Preparation for Data Migration | 00 |
| 6.2 Migrate the Data Using the Migration Wizard | 00 |
| | |
| Chapter 7. Storage Pools | |
| 7.1 Work with Internal Drives | 20 |
| 7.1.1 Actions on Internal Drives | 22 |
| 7.1.2 Configuring Internal Storage | 25 |
| 7.2 Work with MDisks | |

| 7.2.1 Add MDisks to Storage Pools | 242 |
|--|-----|
| 7.2.2 Import MDisks | |
| 7.2.3 RAID Action for MDisks | 258 |
| 7.2.4 Select the tier for MDisks | |
| 7.2.5 Additional Actions on MDisks | 263 |
| 7.3 Work with Storage Pools | 266 |
| | |
| Chapter 8. Advanced Host and Volume Administration | |
| 8.1 Advanced Host Administration | |
| 8.1.1 Modify Mappings | |
| 8.1.2 Unmap all volumes from a host | |
| 8.1.3 Rename | |
| 8.1.4 Delete a Host | |
| 8.1.5 Host Properties | |
| 8.2 Add and delete host ports | |
| 8.2.1 Add Host Port | |
| 8.2.2 Add Fibre Channel Port | |
| 8.2.3 Add iSCSI Host Port | |
| 8.2.4 Delete Host Port | |
| 8.3 Host mappings overview | |
| 8.3.1 Unmap Volume | |
| 8.3.2 Properties Host | |
| 8.3.3 Properties Volume | |
| 8.4 Advanced Volume Administration | |
| 8.4.1 Advanced Volume Functions | |
| 8.4.2 Map Volume to Host | |
| 8.4.3 Unmap Volumes from all Hosts | |
| 8.4.4 View mapped Host to a Volume | |
| 8.4.5 Rename a Volume | 306 |
| 8.4.6 Shrink a Volume | 307 |
| 8.4.7 Expand a Volume | 308 |
| 8.4.8 Migrate volume to another Storage Pool | |
| 8.4.9 Export to Image Mode Volume | |
| 8.4.10 Delete Volume | 313 |
| 8.5 Volume Properties | 314 |
| 8.5.1 Overview | |
| 8.5.2 Host Maps | 317 |
| 8.5.3 Member MDisk | |
| 8.5.4 Add a Mirrored Volume Copy | 320 |
| 8.5.5 Edit Thin-provision Volume Properties | 321 |
| 8.6 Advanced Volume Copy Functions | 324 |
| 8.6.1 Thin-provisioned | 326 |
| 8.6.2 Split into New Volume | 326 |
| 8.6.3 Validate Volume Copies: | 327 |
| 8.6.4 Delete this Copy | 329 |
| 8.6.5 Migrate Volumes using the Volume Copy features | 329 |
| 8.7 Volumes by Storage Pool | |
| 8.7.1 Changing the Storage Pool Icon | 332 |
| 8.7.2 Changing the Storage Pool Name | 333 |
| 8.8 Volumes by Host | 334 |
| 8.8.1 Rename a Host | 336 |
| | |
| Chapter 9. External Storage Virtualization | 339 |

| 9.1 Planning for External Storage Virtualization | |
|---|--|
| 9.1.1 License for External Storage Virtualization | |
| 9.1.2 SAN Configuration Planning | |
| 9.1.3 External Storage Configuration Planning | |
| 9.1.4 Guidelines for Virtualizing External Storage | |
| 9.2 Work with External Storage | |
| 9.2.1 Add External Storage | |
| 9.2.2 Manage External Storage | 345 |
| 9.2.3 Remove External Storage | 351 |
| Observator Francisco | 050 |
| Chapter 10. Easy Tier | |
| 10.1 Easy Tier Overview | |
| 10.2 Easy Tier for IBM Storwize V7000 | |
| 10.2.1 Disk Tiers | |
| 10.2.2 Tiered Storage Pools | |
| 10.2.3 The Easy Tier Process | |
| 10.2.4 Easy Tier Operating Modes | |
| 10.2.5 Easy Tier Rules | |
| 10.3 Easy Tier Configuration using the GUI | 361 |
| 10.3.1 Creating Multi Tiered Pools - Enable Easy Tier | 361 |
| 10.3.2 Downloading Easy Tier I/O measurements | 369 |
| 10.4 Easy Tier Configuration using the CLI | |
| 10.4.1 Enable Easy Tier Evaluation Mode | |
| 10.4.2 Enable or Disable Easy Tier on Single Volumes | |
| 10.5 IBM Storage Tier Advisor Tool | |
| 10.5.1 Creating graphical reports | |
| 10.5.2 STAT Reports | |
| 10.0.2 Civil Reporter | 0,, |
| | |
| Chapter 11. Copy Services | |
| Chapter 11. Copy Services | |
| • | 382 |
| 11.1 FlashCopy | 382 382 |
| 11.1 FlashCopy | 382 382 390 |
| 11.1 FlashCopy | 382 382 390 391 |
| 11.1 FlashCopy | 382 382 390 391 398 |
| 11.1 FlashCopy | 382 382 390 391 398 423 |
| 11.1 FlashCopy 11.1.1 FlashCopy Concepts. 11.1.2 Planning for FlashCopy. 11.1.3 Managing FlashCopy Using the GUI. 11.1.4 Manage FlashCopy Mapping 11.2 Remote Copy 11.2.1 Concepts Inside Remote Copy. | 382 382 390 391 398 423 423 |
| 11.1 FlashCopy 11.1.1 FlashCopy Concepts. 11.1.2 Planning for FlashCopy. 11.1.3 Managing FlashCopy Using the GUI. 11.1.4 Manage FlashCopy Mapping 11.2 Remote Copy 11.2.1 Concepts Inside Remote Copy 11.2.2 Planning for Remote Copy | 382 382 390 391 398 423 423 430 |
| 11.1 FlashCopy 11.1.1 FlashCopy Concepts. 11.1.2 Planning for FlashCopy. 11.1.3 Managing FlashCopy Using the GUI. 11.1.4 Manage FlashCopy Mapping 11.2 Remote Copy 11.2.1 Concepts Inside Remote Copy. | 382 382 390 391 398 423 423 430 |
| 11.1 FlashCopy 11.1.1 FlashCopy Concepts. 11.1.2 Planning for FlashCopy. 11.1.3 Managing FlashCopy Using the GUI. 11.1.4 Manage FlashCopy Mapping 11.2 Remote Copy 11.2.1 Concepts Inside Remote Copy 11.2.2 Planning for Remote Copy 11.2.3 Manage Remote Copy using GUI. | 382 382 390 391 398 423 423 430 432 |
| 11.1 FlashCopy 11.1.1 FlashCopy Concepts. 11.1.2 Planning for FlashCopy. 11.1.3 Managing FlashCopy Using the GUI. 11.1.4 Manage FlashCopy Mapping 11.2 Remote Copy 11.2.1 Concepts Inside Remote Copy 11.2.2 Planning for Remote Copy | 382 382 390 391 398 423 423 430 432 |
| 11.1 FlashCopy 11.1.1 FlashCopy Concepts. 11.1.2 Planning for FlashCopy. 11.1.3 Managing FlashCopy Using the GUI. 11.1.4 Manage FlashCopy Mapping 11.2 Remote Copy 11.2.1 Concepts Inside Remote Copy. 11.2.2 Planning for Remote Copy. 11.2.3 Manage Remote Copy using GUI. Chapter 12. Tivoli Storage Productivity Center 12.1 TPC Overview. | 382 382 390 391 398 423 423 430 432 459 460 |
| 11.1 FlashCopy 11.1.1 FlashCopy Concepts. 11.1.2 Planning for FlashCopy. 11.1.3 Managing FlashCopy Using the GUI. 11.1.4 Manage FlashCopy Mapping. 11.2 Remote Copy 11.2.1 Concepts Inside Remote Copy. 11.2.2 Planning for Remote Copy. 11.2.3 Manage Remote Copy using GUI. Chapter 12. Tivoli Storage Productivity Center. 12.1 TPC Overview. 12.1.1 TPC for Disk Midrange Edition. | 382 382 390 391 398 423 423 430 432 459 460 460 |
| 11.1 FlashCopy 11.1.1 FlashCopy Concepts. 11.1.2 Planning for FlashCopy. 11.1.3 Managing FlashCopy Using the GUI. 11.1.4 Manage FlashCopy Mapping. 11.2 Remote Copy 11.2.1 Concepts Inside Remote Copy 11.2.2 Planning for Remote Copy 11.2.3 Manage Remote Copy using GUI. Chapter 12. Tivoli Storage Productivity Center 12.1 TPC Overview 12.1.1 TPC for Disk Midrange Edition 12.2 TPC Architecture. | 382 382 390 391 398 423 423 430 432 459 460 460 |
| 11.1 FlashCopy 11.1.1 FlashCopy Concepts. 11.1.2 Planning for FlashCopy. 11.1.3 Managing FlashCopy Using the GUI. 11.1.4 Manage FlashCopy Mapping 11.2 Remote Copy 11.2.1 Concepts Inside Remote Copy 11.2.2 Planning for Remote Copy 11.2.3 Manage Remote Copy using GUI. Chapter 12. Tivoli Storage Productivity Center 12.1 TPC Overview. 12.1.1 TPC for Disk Midrange Edition 12.2 TPC Architecture. 12.2.1 Data Server. | 382 382 390 391 398 423 423 430 432 459 460 461 461 |
| 11.1 FlashCopy 11.1.1 FlashCopy Concepts 11.1.2 Planning for FlashCopy 11.1.3 Managing FlashCopy Using the GUI. 11.1.4 Manage FlashCopy Mapping 11.2 Remote Copy 11.2.1 Concepts Inside Remote Copy 11.2.2 Planning for Remote Copy 11.2.3 Manage Remote Copy using GUI. Chapter 12. Tivoli Storage Productivity Center 12.1 TPC Overview 12.1.1 TPC for Disk Midrange Edition 12.2 TPC Architecture 12.2.1 Data Server 12.2.2 Device Server | 382 382 390 391 398 423 423 430 432 459 460 461 461 |
| 11.1 FlashCopy 11.1.1 FlashCopy Concepts. 11.1.2 Planning for FlashCopy. 11.1.3 Managing FlashCopy Using the GUI. 11.1.4 Manage FlashCopy Mapping. 11.2 Remote Copy. 11.2.1 Concepts Inside Remote Copy. 11.2.2 Planning for Remote Copy. 11.2.3 Manage Remote Copy using GUI. Chapter 12. Tivoli Storage Productivity Center. 12.1 TPC Overview. 12.1.1 TPC for Disk Midrange Edition. 12.2 TPC Architecture. 12.2.1 Data Server. 12.2.2 Device Server. 12.2.3 Tivoli Integrated Portal. | 382 382 390 391 398 423 423 430 432 459 460 461 461 461 |
| 11.1 FlashCopy 11.1.1 FlashCopy Concepts. 11.1.2 Planning for FlashCopy. 11.1.3 Managing FlashCopy Using the GUI. 11.1.4 Manage FlashCopy Mapping 11.2 Remote Copy 11.2.1 Concepts Inside Remote Copy. 11.2.2 Planning for Remote Copy. 11.2.3 Manage Remote Copy using GUI. Chapter 12. Tivoli Storage Productivity Center 12.1 TPC Overview. 12.1.1 TPC for Disk Midrange Edition 12.2 TPC Architecture. 12.2.1 Data Server. 12.2.2 Device Server 12.2.3 Tivoli Integrated Portal 12.2.4 Tivoli Storage Productivity Center for Replication. | 382 382 390 391 398 423 423 430 459 460 461 461 461 461 |
| 11.1 FlashCopy 11.1.1 FlashCopy Concepts. 11.1.2 Planning for FlashCopy. 11.1.3 Managing FlashCopy Using the GUI. 11.1.4 Manage FlashCopy Mapping 11.2 Remote Copy 11.2.1 Concepts Inside Remote Copy. 11.2.2 Planning for Remote Copy. 11.2.3 Manage Remote Copy using GUI. Chapter 12. Tivoli Storage Productivity Center 12.1 TPC Overview. 12.1.1 TPC for Disk Midrange Edition 12.2 TPC Architecture. 12.2.1 Data Server. 12.2.2 Device Server 12.2.3 Tivoli Integrated Portal 12.2.4 Tivoli Storage Productivity Center for Replication. 12.2.5 DB2 Database. | 382 382 390 391 398 423 423 430 459 460 461 461 461 461 461 |
| 11.1 FlashCopy 11.1.1 FlashCopy Concepts. 11.1.2 Planning for FlashCopy. 11.1.3 Managing FlashCopy Using the GUI. 11.1.4 Manage FlashCopy Mapping 11.2 Remote Copy. 11.2.1 Concepts Inside Remote Copy. 11.2.2 Planning for Remote Copy. 11.2.3 Manage Remote Copy using GUI. Chapter 12. Tivoli Storage Productivity Center. 12.1 TPC Overview. 12.1.1 TPC for Disk Midrange Edition. 12.2 TPC Architecture. 12.2.1 Data Server. 12.2.2 Device Server. 12.2.3 Tivoli Integrated Portal. 12.2.4 Tivoli Storage Productivity Center for Replication. 12.2.5 DB2 Database. 12.2.6 Agents. | 382 382 390 391 398 423 423 430 432 459 460 461 461 461 461 462 462 |
| 11.1 FlashCopy 11.1.1 FlashCopy Concepts 11.1.2 Planning for FlashCopy 11.1.3 Managing FlashCopy Using the GUI. 11.1.4 Manage FlashCopy Mapping 11.2 Remote Copy 11.2.1 Concepts Inside Remote Copy 11.2.2 Planning for Remote Copy 11.2.3 Manage Remote Copy using GUI. Chapter 12. Tivoli Storage Productivity Center 12.1 TPC Overview 12.1.1 TPC for Disk Midrange Edition 12.2 TPC Architecture. 12.2.1 Data Server. 12.2.2 Device Server 12.2.3 Tivoli Integrated Portal 12.2.4 Tivoli Storage Productivity Center for Replication. 12.2.5 DB2 Database. 12.2.6 Agents. 12.2.7 Interfaces | 382 382 390 391 398 423 423 430 432 459 460 461 461 461 461 462 462 462 |
| 11.1 FlashCopy 11.1.1 FlashCopy Concepts 11.1.2 Planning for FlashCopy 11.1.3 Managing FlashCopy Using the GUI. 11.1.4 Manage FlashCopy Mapping 11.2 Remote Copy 11.2.1 Concepts Inside Remote Copy 11.2.2 Planning for Remote Copy 11.2.3 Manage Remote Copy using GUI. Chapter 12. Tivoli Storage Productivity Center 12.1 TPC Overview 12.1.1 TPC for Disk Midrange Edition 12.2 TPC Architecture 12.2.1 Data Server 12.2.2 Device Server 12.2.3 Tivoli Integrated Portal 12.2.4 Tivoli Storage Productivity Center for Replication 12.2.5 DB2 Database. 12.2.6 Agents 12.2.7 Interfaces 12.3 Preparing Windows to Install TPC | 382 382 390 391 398 423 423 430 459 460 461 461 461 461 462 462 462 463 |
| 11.1 FlashCopy 11.1.1 FlashCopy Concepts 11.1.2 Planning for FlashCopy 11.1.3 Managing FlashCopy Using the GUI. 11.1.4 Manage FlashCopy Mapping 11.2 Remote Copy 11.2.1 Concepts Inside Remote Copy 11.2.2 Planning for Remote Copy 11.2.3 Manage Remote Copy using GUI. Chapter 12. Tivoli Storage Productivity Center 12.1 TPC Overview 12.1.1 TPC for Disk Midrange Edition 12.2 TPC Architecture 12.2.1 Data Server 12.2.2 Device Server 12.2.3 Tivoli Integrated Portal 12.2.4 Tivoli Storage Productivity Center for Replication 12.2.5 DB2 Database. 12.2.6 Agents 12.2.7 Interfaces 12.3 Preparing Windows to Install TPC 12.3.1 Installation overview | 382 382 390 391 398 423 423 430 432 460 461 461 461 461 461 462 462 462 463 463 |
| 11.1 FlashCopy 11.1.1 FlashCopy Concepts 11.1.2 Planning for FlashCopy 11.1.3 Managing FlashCopy Using the GUI. 11.1.4 Manage FlashCopy Mapping 11.2 Remote Copy 11.2.1 Concepts Inside Remote Copy 11.2.2 Planning for Remote Copy 11.2.3 Manage Remote Copy using GUI. Chapter 12. Tivoli Storage Productivity Center 12.1 TPC Overview 12.1.1 TPC for Disk Midrange Edition 12.2 TPC Architecture 12.2.1 Data Server 12.2.2 Device Server 12.2.3 Tivoli Integrated Portal 12.2.4 Tivoli Storage Productivity Center for Replication 12.2.5 DB2 Database. 12.2.6 Agents 12.2.7 Interfaces 12.3 Preparing Windows to Install TPC | 382 382 390 391 398 423 423 430 432 459 460 461 461 461 461 462 462 463 463 463 |

| 12.4 Installing Tivoli Storage Productivity Center components | 477 |
|--|-------------|
| 12.5 Connect TPC to IBM Storwize V7000 | 486 |
| 12.6 TPC Example Views | 493 |
| Chapter 13. RAS, Monitoring and Troubleshooting | 497 |
| 13.1 Reliability, Availability, and Serviceability on the IBM Storwize V7000 | |
| 13.1.1 Node Canister | |
| 13.1.2 Disk subsystem | |
| 13.1.3 Power | |
| 13.2 Configuration backup procedure | |
| 13.3 Software upgrade | |
| 13.4 Monitoring host paths to the IBM Storwize V7000 | |
| 13.5 Troubleshooting overview | |
| 13.6 Manipulating the Event Log and Recommended Actions Grid | |
| 13.7 Recommended Actions - details | 520 |
| 13.8 Audit Log | 533 |
| 13.9 Collecting Support Information | 535 |
| 13.10 Shutting down the IBM Storwize V7000 | 538 |
| 13.10.1 Shutting down using the GUI | |
| 13.10.2 Shutting down using the Command Line Interface | 540 |
| 13.11 Shutting down and powering on an IBM Storwize V7000 infrastructure | 540 |
| Appendix A. CLI setup and SAN Boot | 543 |
| Command Line Interface | |
| Basic Setup | 544 |
| Example Commands | 552 |
| SAN Boot | |
| Enable SAN Boot for Windows | 557 |
| Enable SAN Boot for VMware | |
| Windows SAN Boot Migration | 558 |
| Related publications | 561 |
| IBM Redbooks publications | |
| Other publications | |
| Online resources | 562 |
| Help from IBM | 563 |
| In along | - 0- |

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 give 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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites 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.

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® IBM® Solid®

DB2® Passport Advantage® System Storage®

DS6000™ POWER® Tivoli® DS8000® Redbooks® XIV®

FlashCopy® Redbooks (logo) @®

The following terms are trademarks of other companies:

Snapshot, and the NetApp logo are trademarks or registered trademarks of NetApp, Inc. in the U.S. and other countries.

Java, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. 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.

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.

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

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

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® is introducing the IBM Storwize V7000, 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's offered at a competitive, midrange price.

IBM Storwize V7000 incorporates some of IBM's top 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 book will introduce the features and functions of the Storwize V7000 by example.

The team who wrote this book

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

Brian Cartwright is a Consulting I/T Specialist based in Brisbane, Australia. He has 23 years experience in the storage market covering both open systems and mainframe storage systems. His areas of expertise include Storage Area Networks, Disk and Tape subsystems, Storage virtualization and IBM storage software. Brian holds a degree in Computing Studies from Canberra University and is both an IBM Certified I/T Specialist and an IBM Certified Storage Specialist.

Ronda Hruby is a SAN Volume Controller Level 2 Support Engineer at the Almaden Research Center in San Jose, California. She also supports multipathing software, and virtual tape products. Before joining the IBM Storage Software PFE organization in 2002, she worked in hardware and microcode development for more than 20 years and is a SNIA certified professional.

Daniel Koeck is a Systems Engineer in IBM Systems and Technology Group in Austria. He has a graduate degree in applied computer science and IT security. His daily business is mainly focused on virtualized environments and his areas of expertise include storage, server, desktop and I/O virtualization, as well as high-availability solutions, datacenter migrations and software engineering. Daniel is author of several technical publications, including 3 previous redbooks and enjoys sharing his knowledge as a speaker on national and international conferences.

Xin Liu is an Advisory IT Specialist in IBM China with a focus on IBM system hardware and software products he has undertaken various roles in the technical support area. Starting as a field support engineer in IBM Global Technology Services China, now he is a member of the Advanced Technical Skills organization in the IBM Greater China Group specializing in storage virtualization. Xin holds a Master's degree in Electronic Engineering from Tsinghua University, China.

Massimo Rosati is a Certified ITS Senior Storage and SAN Software Specialist in IBM Italy. He has 25 years of experience in the delivery of Professional Services and SW Support. His areas of expertise include storage hardware, Storage Area Network, storage virtualization,

Disaster Recovery and Business Continuity solutions. He has written other IBM Redbooks® on storage virtualization products.

Thomas Vogel is an IBM Certified IT Specialist working for the European Advanced Technical Support organization. He joined IBM in 1997 and since 2000 he has worked at the European Storage Competence Center (ESCC) in Mainz, Germany. He has worked with IBM midrange and high end products in the Open Systems environment and he has been involved in new product introductions for FAStT, DS6000™, DS8000®. Since 2007 his focus is on the SAN Volume Controller (SVC) and covers European ATS support, and he is involved in product updates, ESP and Beta programs. Thomas holds a degree in Electrical Engineering from the Technical University Ilmenau, Germany.

Bill Wiegand is a Senior IT Specialist for IBM System Storage® and is part of the Global Advanced Technical Skills organization. Bill has been with IBM 32 years and in the information technology field for 35 in various roles from hardware service and software technical support to team lead for a UNIX® and DBA group in a Fortune 500 company. For the last 10 years he has been focused on storage and is the America's storage virtualization expert for the SAN Volume Controller (SVC). In this role he has worldwide responsibilities for educating the technical sales teams and business partners, helping to design and sell solutions around SVC, assisting with customer performance and problem analysis as well as acting as the technical liaison between the field and the SVC development team.

The project manager was:

Jon Tate is a Project Manager for IBM System Storage System Networking and Virtualization solutions at the International Technical Support Organization, San Jose Center. Before joining the ITSO in 1999, he worked in the IBM Technical Support Center, providing level 2 and 3 support for IBM storage products. Jon has 25 years of experience in storage software and management, services, and support, and is both an IBM Certified IT Specialist and an IBM SAN Certified Specialist. He is also the UK Chair of the Storage Networking Industry Association.

Figure 1 shows the authors outside IBM Hursley.



Figure 1 L-R: Brian, Bill, Ronda, Thomas, and Jon

Figure 2 shows the San Jose team of authors.

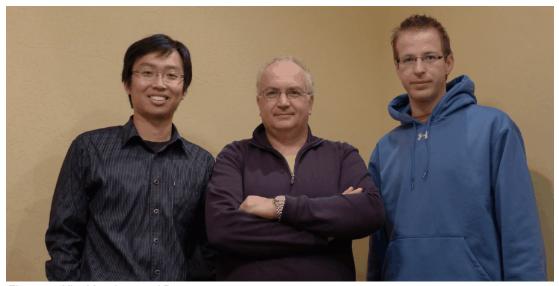


Figure 2 Xin, Massimo, and Dan

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

Tayfun Arli Chris Canto Peter Eccles John Fairhurst Carlos Fuente **Howard Greaves**

Geoff Lane

Alex Howell

Andrew Martin

Cameron McAllister

Paul Merrison

Steve Randle

Lucy Harris

Bill Scales

Dave Sinclair

Matt Smith

Steve White

Barry Whyte

Muhammad Zuba

IBM Hursley

Duane Bolland

Evelyn Wick

Larry Wiedenhoft

IBM Beaverton

Terry Niemeyer

IBM Austin

Mary Lovelace

IBM San Jose

Larry Chiu

Paul Muench

IBM Almaden

Sharon Wang

IBM Chicago

Chris Saul

IBM San Jose

Torsten Rothenwaldt

IBM Germany

A special thanks also goes to Brocade who hosted the second team of authors in their headquarters building in San Jose, California:

Mansi Botadra

Yong Choi

Silviano Gaona

Brian Steffler

Marcus Thordal

Steven Tong

Brocade Communications Systems

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



1

Overview of the IBM Storwize V7000

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 now one of the foundations for a building a flexible and reliable infrastructure solution which allows companies to better align their business and IT needs.

Storage Virtualization allows an organization to implement pools of storage across physically separate disk system (which may even be from different vendors). Storage can then be deployed from these pools and can be migrated between pools without any outage to the attached host systems. It provides a single set of tools for advanced functions such as instant copy and remote mirroring solutions. This means that deploying storage can be performed using a single tool regardless of the underlying storage hardware. This is illustrated below in Figure 1-1.

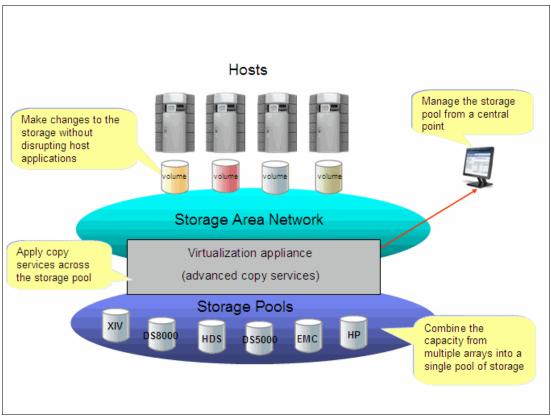


Figure 1-1 Storage virtualization

It is not our intent to provide a detailed description of virtualization in this paper, for a more detailed explanation of Storage Virtualization refer to *Implementing the IBM System Storage SAN Volume Controller V6.1*, 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 as well as its own internal storage. The IBM Storwize V7000 solution is built upon the IBM SAN Volume Controller (SVC) technology base and utilizes technology from the IBM DS8000 family.

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

In Figure 1-2 we show the IBM Storwize V7000 high-level overview.

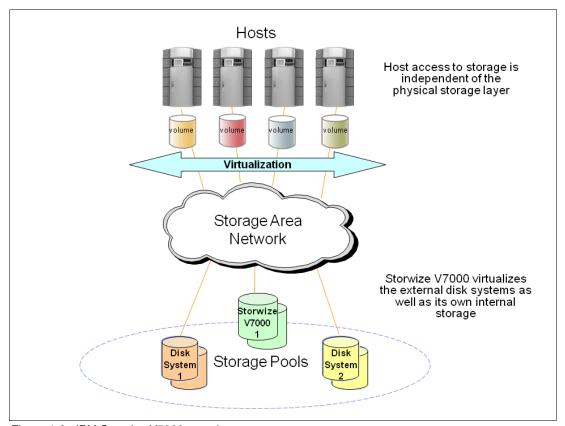


Figure 1-2 IBM Storwize V7000 overview

Included with IBM Storwize V7000 is a simple and easy to use Graphical User Interface which is designed to allow storage to 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 called presets that use commonly used settings to quickly configure objects on the system. Presets are available for creating volumes and FlashCopy® mappings and for setting up RAID configuration.

The IBM Storwize V7000 solution provides a choice of up to 120 x 3.5 inch or 240 x 2.5 inch Serial Attached SCSI (SAS) drives for the internal storage and uses SAS cables and connectors to attach to the optional expansion enclosures.

Note: Be aware that at the time of writing this book there is a temporary restriction, that limit the maximum enclosures to 5 (1 controller enclosure and 4 expansion enclosures), this limitation will be removed in the near future.

When virtualizing external storage arrays IBM Storwize V7000 can provide up to 32PB of usable capacity. IBM Storwize V7000 supports a range of external disk systems similar to what the SVC supports today.

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

1.3 IBM Storwize V7000 terminology

IBM Storwize V7000 has introduced some new terminology which is 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 of significant to a task or system. Events can include completion or failure of an operation, a user action, or the change in 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 cluster. | |
| host mapping | The process of controlling which hosts have access to specific volumes within a cluster. | |
| internal storage | Array managed disks (MDisks) and drives that are held in enclosures and nodes that are part of the cluster. | |
| lane | A single SAS channel | |
| managed disk (MDisk) | A component of a storage pool that is managed by a cluster. 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. | |
| PHY | A single SAS lane. There are 4 PHY 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. | |

| IBM Storwize V7000 term | Definition | |
|---------------------------------------|---|--|
| 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, 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 | |

1.3.1 IBM Storwize V7000 models

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

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® |
|---|-------|-----------------------|-------------|---|
| 2076-112 (with two node canisters) | 16GB | 8x8Gb/4x1Gb/4x 6Gb | 12x3,5" | Integrated dual power supplies with battery backup |
| 2076-124 (with two node canisters) | 16GB | 8x8Gb/4x1Gb/4x 6Gb | 24x2,5" | Integrated dual power supplies with battery backup |
| 2076-212 (with two expansion canisters) | N/A | //4x6Gb | 12x3,5" | Integrated dual power supplies |
| 2076-224 (with two expansion canisters) | N/A | //4x6Gb | 24x2,5" | Integrated dual power supplies |

Shown in Figure 1-3 is the front view of the 2076-112 and 2076-212 enclosures



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

The drives are positioned in 4 columns of 3 horizontal mounted drive assemblies. The drive slots are numbered 1 to 12, starting top left and going left to right, top to bottom.

Shown in Figure 1-4 is the front view of the 2076-124 and 224 enclosures.



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

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

1.3.2 IBM Storwize V7000 attributes

IBM Storwize V7000 provides an integrated 2EIA storage controller with the following attributes:

- Dual controller with 16GB cache and battery backup
- DS8000 technology for RAID functionality
- ► Support for RAID 0, 1, 5, 6, 10
- ► Maximum member drives per RAID 0 = 8 drives
- Maximum member drives per RAID 1 = 2 drives
- Maximum member drives per RAID 5/6/10 = 16 drives
- Maximum RAID arrays 128
- Maximum of 240 internal drives
- Maximum of 9 expansion enclosures
- 8 x 8 Gbps Fibre Channel ports
- Maximum 16 Storage Fibre Channel ports per external storage
- Maximum 1024 external storage WWNNs
- Maximum 1024 external storage WWPNs
- ► Maximum of 2048 host volumes
- Maximum of 512 volumes per host
- Maximum capacity of 256TB per volume
- Maximum virtualization addressability for back-end storage of 32PB
- Maximum of 256 Fibre Channel attached hosts
- Maximum of 256 iSCSI attached hosts
- Maximum of 256 combined (iSCSI and Fibre Channel hosts)
- Maximum of 4096 MDisks
- Maximum of 2TB capacity per MDisk

For a complete and updated list of IBM Storwize V7000 configuration limits and restrictions check the Web site at:

 $\label{lem:http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003702\&myns=s028\&mynp=familyind5402112\&mync=E$

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 never used. Thin provisioned volumes provide the ability to allocate real physical disk capacity only when data is actually 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 completely separate disk storage systems that are being virtualized. In the case of one copy failing IBM Storwize V7000 will provide 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 volume level point-in-time copy function for any storage being virtualized by IBM Storwize V7000. This is designed to create copies for backup, parallel processing, test and development and have the copies available almost immediately.

IBM Storwize V7000 includes the following FlashCopy functions:

Full/Incremental Copy

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

Multi-target FlashCopy

IBM Storwize V7000 supports copying 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 same 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 for data from an earlier point-in-time copy to be restored with minimal disruption to the host.

- FlashCopy Nocopy with Thin Provisioning

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

- Space-efficient source and target with background copy
 - This copies only the allocated space
- Space-efficient target with no background copy

This copies only the space used for changes between the source and target and is generally referred to as "snapshots".

May be used with multi-target, cascaded and incremental FlashCopy

- Consistency Groups

Consistency groups address the issue where application data resides across multiple volumes. By placing the FlashCopy relationships into a Consistency Group commands can be issued against all of the volumes residing in the group. This enables a consistent point-in-time copy of all of the data even though it may reside on physical separate volume.

FlashCopy mappings can be members of a consistency group, or they can be operated in a stand-alone manner, 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 synchronous remote mirroring function up to approximately 300km between sites. As the host I/O only completes once the data is cached at both locations performance requirements may limit the practical distance. Metro Mirror is designed to provide fully synchronized copies at both sites with zero data loss once the initial copy is completed.

Metro Mirror can operate between multiple IBM Storwize V7000 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 8,000km between sites. With Global Mirror the host I/O completes locally and the changed data is sent to the remote site later. This is designed to maintain a consistent recoverable copy of data at the remote site which lags behind the local site.

Global Mirror can operate between multiple IBM Storwize V7000 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.

It allows you to:

- Move volumes non-disruptively 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

► Easy Tier (no charge)

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

This is shown in Figure 1-5 on page 10.

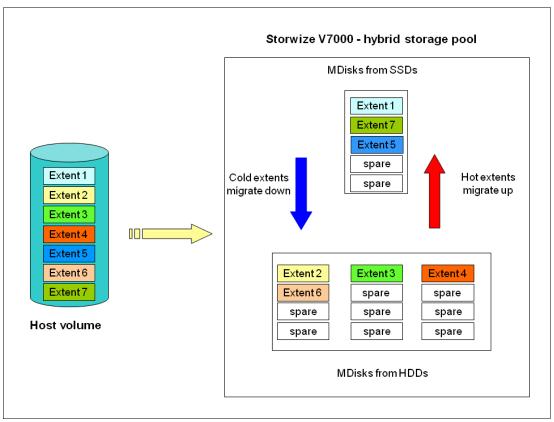


Figure 1-5 Easy Tier overview

1.3.4 IBM Storwize V7000 Licensing

IBM Storwize V7000 may require the following licenses:

- ► Enclosure
- ► External Virtualization
- ► Remote Copy (Advanced Copy Services: Metro/Global Mirror)

In Table 1-3 there is a summary of all licenses that may be required.

Table 1-3 Licenses that may be required

| License type | Unit | License number | License required? |
|----------------------------|---|----------------|-------------------|
| ENCLOSURE | BASE+ESPANSION PHYSICAL ENCLOSURE NUMBER | 5639-VM1 | YES |
| EXTERNAL VIRTUALIZATION | PHYSICAL ENCLOSURE NUMBER OF EXTERNAL STORAGE | 5639-EV1 | OPTIONAL |
| REMOTE COPY | PHYSICAL ENCLOSURE NUMBER | 5639-RM1 | OPTIONAL |

| License type | Unit | License number | License required? |
|---------------------|------|----------------|-------------------|
| FLASHCOPY | N/A | N/A | NO |
| VOLUME MIRRORING | N/A | N/A | NO |
| THIN PROVISIONING | N/A | N/A | NO |
| VOLUME MIGRATION | N/A | N/A | NO |
| EASY TIER | N/A | N/A | NO |

Note: If the Storwize V7000 is to be 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 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, then the appropriate External Virtualization license must be ordered too.

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 Storage Bridge Bay (SBB) compliant.

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

Shown in Figure 1-6 on page 12 is an overview of the hardware components of the IBM Storwize V7000 solution.

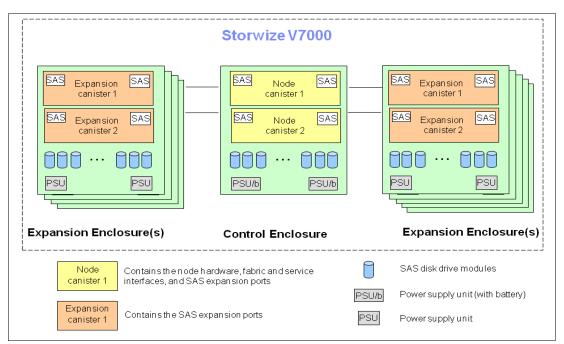


Figure 1-6 IBM Storwize V7000 hardware components

The controller is shown in Figure 1-7.



Figure 1-7 IBM Storwize V7000 controller rear

There are two power supply slots, on the extreme left and extreme right, each taking up the full 2EIA height. The left hand slot is power supply 1, the right hand 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.

1.4.1 Control enclosure

Each IBM Storwize V7000 system will have one control enclosure which contains two *node canisters*, disk drives and two power supplies. There are two models of the control enclosure with one model providing 12 3,5" disk slots and the other providing 24 2,5" disk slots.

Within a control enclosure, each power supply unit (PSU) contains a battery. The battery is designed to enable the IBM Storwize V7000 system to perform a dump of the cache to internal disks in the event of both power inputs failing.

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 if one node was to fail the surviving node would automatically take over. Nodes are deployed in pairs called I/O groups.

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

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

Shown in Figure 1-8 is a single node canister.



Figure 1-8 A single canister

There are four Fibre Channel ports 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 to 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.

There are two 10/100/1000 Mbps Ethernet ports side by side on the canister. They are numbered 1 on the left and 2 on the right. Use of port 1 is required while use of port 2 is optional. There are two LEDs associated with each Ethernet port.

There are two USB 2.0 ports 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 use is with a USB memory sticks and is described in Chapter 2, "Initial configuration" on page 35.

There are two 6 Gbps SAS ports 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 with one model providing 12 disk slots and the other providing 24 disk slots.

Figure 1-9 shows the expansion enclosure.



Figure 1-9 Expansion enclosure

The expansion enclosure power supplies are similar to the control enclosure but do not contain the battery. 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 2 SAS interfaces which 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 will be used to attach to additional expansion enclosures.

Each port connects four PHYs (ports of SAS drive). There is an LED associated with each PHY in each port (8 LEDs in total). The LEDs are green and are next to the ports, and for each port they are numbered 1 through 4. The LED indicates activity on the PHY.

Shown in Figure 1-10 is an expansion canister.



Figure 1-10 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 2 ports (2 PHY) and I/O can be issued down both paths simultaneously.

Shown in Table 1-4 are the IBM Storwize V7000 Disk Drive types that are available at the time of writing.

Drive Type Speed Size 2.5" form Solid® State n/a 300GB factor Disk enterprise-grade multilevel cell (E-MLC) 300, 450 and 2.5" form SAS 10,000 RPM factor 600GB 3.5" form Nearline SAS 7,200 RPM 2TB factor

Table 1-4 IBM Storwize V7000 Disk Drive types

1.5 IBM Storwize V7000 Components

The IBM Storwize V7000 is an entry/midrange virtualization RAID storage subsystem. It includes the following benefit:

- ▶ Brings enterprise technology to entry/midrange storage.
- Specialist administrators are not required.
- ► Easy client setup and service.
- Ability to grow the system incrementally as storage capacity and performance needs change.
- Multiple storage tiers in a single system with nondisruptive migration between them.
- ► Simple integration 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 discussed in Chapter 9, "External Storage Virtualization" on page 339.

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 and/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 has two hardware components called nodes or node canisters that provide the virtualization of internal and external volumes, cache and copy services (remote copy) functions. A cluster consists of one node pair.

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

1.5.3 I/O Group

Within IBM Storwize V7000 a pair of node canisters is known as an I/O group. The IBM Storwize V7000 supports two node canisters in the cluster which provides one I/O group.

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

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

So 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 FC to use multipath drivers to be able to handle these capability.

The Storwize V7000 I/O Group is connected to the SAN so that all application servers accessing volumes from the I/O Group have access to them. Up to 256 host server objects can be defined in the I/O Group.

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

1.5.4 Cluster

A cluster consists of a pair of nodes. All configuration, monitoring, and service tasks are performed at the cluster level and the configuration settings are replicated across both node canisters in the cluster. To facilitate these tasks one or two management ip addresses are set for the cluster.

There is a process provided to backup the cluster configuration data onto disk so that the cluster can be restored in the event of a disaster. This method does not backup application data, only Storwize V7000 cluster configuration information.

Note: After backup of the cluster configuration, remember to save the backup data on your hard disk (or at the very least outside of the SAN), because in the event that you are unable to access the Storwize V7000 you will not have access to the backup data if it is on the SAN.

For the purposes of remote data mirroring, two or more clusters (IBM Storwize V7000 systems) must form a partnership prior to creating relationships between mirrored volumes.

Note: At this time remote copy services are supported only between Storwize V7000 systems. Any remote copy relationship between Storwize V7000 and SVC is not supported.

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

The cluster can be configured using either the IBM Storwize V7000 management software, the Command Line Interface (CLI) or via an application that uses the IBM Storwize V7000 CIMOM (that is to say TPC). IBM Systems Director also provides flexible server and storage management capability.

1.5.5 **RAID**

The Storwize V7000 setup will contain a number of 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 different degrees of redundancy and performance, and have different restrictions on 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 desired location and performance of each array. A sequence of drive failures and hot spare takeovers can leave an array unbalanced, that is with members that do not match these goals. The system automatically rebalances such arrays when appropriate drives are available.

Available RAID levels are:

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

RAID 0 arrays stripe data across the drives. The system supports RAID 0 arrays with just one member which is similar to traditional JBOD attach. RAID 0 arrays have no redundancy so do

not support hot spare takeover or immediate exchange. RAID 0 array can be formed by1 to 8 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 2 drives only.

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

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

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

For a detailed description about RAID levels see the Web site at :

http://en.wikipedia.org/wiki/RAID

1.5.6 Managed Disks

A managed disk (MDisk) refers to the unit of storage that IBM Storwize V7000 virtualizes. This 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 as it is internal or only zoned to the IBM Storwize V7000 system.

An MDisk has four modes:

Array

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

Unmanaged

This is when an MDisk is not being used by the cluster. This might occur when an MDisk is first imported into the cluster for example.

Managed

This is when an MDisk is assigned to a storage pool and provides extents that volumes can use it.

► Image

This is when an MDisk is assigned directly to a volume with a one-to-one mapping of extents between the MDisk and the volume. This is normally used when importing logical volumes into the cluster that already have data on them. This ensures the data is preserved as it is imported into the cluster.

1.5.7 Quorum disks

A quorum disk is a managed disk (MDisk) that contains a reserved area for use exclusively by the cluster. In the Storwize V7000 internal drives can be considered as quorom candidates. The cluster uses quorum disks to break a tie when exactly half the nodes in the cluster remain after a SAN failure.

The cluster automatically forms the quorum disk by taking a small amount of space from a managed disk (MDisk). It will allocate 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 quorum disks with a failure of a single storage system it is recommended to allocate the quorum disk on different storage systems. Quorum disk management is possible using the CLI.

1.5.8 Storage pools

A storage pool is a collection of MDisks (up to 128) that are grouped together in order 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 will change from unmanaged to managed and vice-versa when you remove it.

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

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

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

| Extent Size | Maximum volume capacity for normal volumes (GB) | | | |
|-------------|---|--|--|--|
| 16 | 2048 (2TB) | | | |
| 32 | 4096 (4TB) | | | |
| 64 | 8192 (8TB) | | | |
| 128 | 16384 (16TB) | | | |
| 256 | 32768 (32TB) | | | |

Table 1-5 Maximum volume capacity by extent size

| Extent Size | Maximum volume capacity for normal volumes (GB) | | | |
|-------------|---|--|--|--|
| 512 | 65536 (64TB) | | | |
| 1024 | 131072 (128TB) | | | |
| 2048 | 262144 (256TB) | | | |
| 4096 | 528288 (512TB) | | | |
| 8192 | 1056576 (1,024TB) | | | |

The effect of extent size on the maximum Cluster capacity is shown in Table 1-6.

Table 1-6 Extent size and affect on cluster capacity

| Extent Size | Maximum storage capacity of cluster | | | | |
|-------------|-------------------------------------|--|--|--|--|
| 16MB | 64TB | | | | |
| 32MB | 128TB | | | | |
| 64MB | 256TB | | | | |
| 128MB | 512TB | | | | |
| 256MB | 1PB | | | | |
| 512MB | 2PB | | | | |
| 1024MB | 4PB | | | | |
| 2048MB | 8PB | | | | |
| 4096MB | 16PB | | | | |
| 8192MB | 32PB | | | | |

We recommend that you use the same extent size for all storage pools in a cluster, which is a prerequisite for supporting volume migration between two storage pool. 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 219.

For most clusters a capacity of 1PB is sufficient. We therefore recommend that you use a value of 256MB.

Note: The GUI of IBM Storwize V7000 has a default value of 256 Extent Size when you define a new storage pool.

A storage pool can have a threshold warning set that will automatically issue 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 inducing performance and other problems:

► They 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.
- ▶ We recommend that you use MDisks of the same size and therefore MDisks that provide the same number of extents. If that is not feasible you will need to check the distribution of the volumes' extents in that storage pool.

Multi Tiered Storage Pool

A multiple tiered storage pool will have 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 multi-tiered storage pool will therefore contain MDisks with different characteristics as opposed to the single tier storage pool. However, we recommend that each tier has MDisks of the same size and MDisks that provide the same number of extents.

A multi-tiered storage pool is used to enable automatic migration of extents between disk tiers using the IBM Storwize V7000 Easy Tier function, Chapter 10, "Easy Tier" on page 353.

Figure 1-11 shows these components.

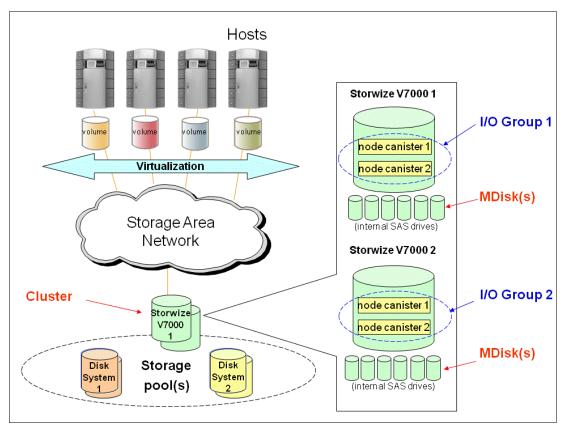


Figure 1-11 IBM Storwize V7000 virtualization components

1.5.9 Volumes

A volume is a logical disk that is presented to a host system by the cluster. In our virtualized environment the host system has a volume mapped to it by IBM Storwize V7000. IBM Storwize V7000 translates this volume into a number of extents which are allocated across

MDisks. The advantage with storage virtualization is that the host is "decoupled" from the underlying storage so the virtualization appliance can move the extents around without impacting the host system.

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

There are three types of volumes:

Striped

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

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

This is illustrated in Figure 1-12 below which shows how an striped volume is allocated assuming 10 extents were required,

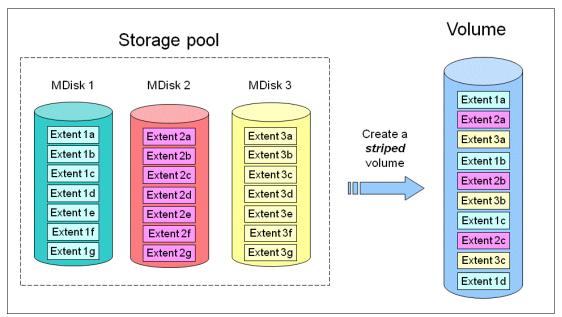


Figure 1-12 Striped volume

Seguential

A sequential volume is where the extents are allocated one after the other from one MDisk then progresses to the next MDisk.

This is shown in Figure 1-13 on page 23.

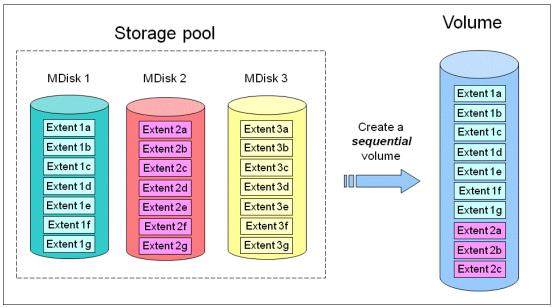


Figure 1-13 Sequential volume

Image Mode

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

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

This is illustrated in Figure 1-14 on page 23.

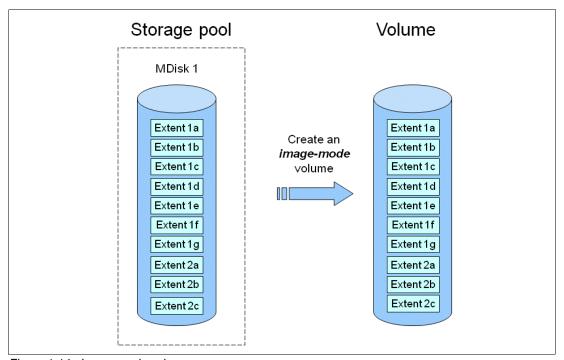


Figure 1-14 Image mode volume

Some virtualization functions are not available for image mode volumes so it is often useful to migrate the volume into a new storage pool. Once 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 will behave with respect to application reads and writes as though they were fully allocated. When a volume is created, the user will specify two capacities: the real capacity of the volume and its virtual capacity.

The real capacity will determine the quantity of MDisk extents that will be allocated for the volume. The virtual capacity will be the capacity of the volume reported to IBM Storwize V7000 and to the host servers.

The real capacity will be 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 over-allocated volumes, or it can be used in conjunction with FlashCopy. Thin Provisioned volumes can be used in conjunction with the mirrored volume feature as well.

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 will go offline as soon as it needs to expand whereas a volume with a non-zero contingency capacity will stay online until it has been used up.

Autoexpand will 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 will be recalculated.

To support the autoexpansion of thin provisioned volumes the storage pools from which they are allocated have a configurable warning capacity. When the used free capacity of the group exceeds the warning capacity, a warning is logged.

For example, if a warning of 80% has been specified, the warning will be logged when 20% of the free capacity remains.

Note: Thin-provisioned volumes require additional I/O operation to read and write metadata to the internal disks or to the back-end storage which will also add additional load to the IBM Storwize V7000 nodes. We therefore do not recommend the use of these volumes for high performance applications, or any workload with a high write I/O component.

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 called *volume mirroring* which enables a volume to have two physical copies. Each volume copy can belong to a different storage pool and can reside on different physical storage systems, and it provides a high-availability solution.

When a host system issues a write to a mirrored volume IBM Storwize V7000 will write the data to both copies. When a host system issues a read to a mirrored volume IBM Storwize V7000 will place it to the primary copy. If one of the mirrored volume copies is temporarily unavailable, the IBM Storwize V7000 will automatically use the alternate copy without any outage for the host system. When the mirrored volume copy is repaired IBM Storwize V7000 will resynchronize 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 completely 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.

Note: An unmirrored volume can be migrated from one location to another by simply adding a second copy to the desired destination, wait for the two copies to synchronize, and then remove the original copy. This operation can be stopped at any time.

1.5.12 Easy Tier

Easy Tier is a performance function that will automatically migrate or move 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 multi-tier storage pool over a 24 hours period. It then creates an extent migration plan based on this activity and will then dynamically move high activity or hot extents to a higher disk tier within the storage pool. It will also move extents whose activity has dropped off or cooled from the high tier 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 get an understanding of the potential benefit of Easy Tier in your environment before installing solid-state disk. By turning the Easy Tier function ON for a single level storage pool and also turning Easy Tier Function ON for the volumes within that pool, Easy Tier will create a migration report every 24 hours on the number of extents it would move if the pool was a multi-tiered pool. Easy Tier statistics measurement is enabled.

The use of Easy Tier may make it more appropriate to use smaller storage pool extent sizes.

The usage statistics file can be off-loaded from IBM Storwize V7000 nodes and then an IBM Storage Advisor Tool can be used to create a summary report.

Contact your IBM representative or IBM Business Partner for more information on the Storage Advisor Tool.

We describe Easy Tier in more detail in Chapter 10, "Easy Tier" on page 353.

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 V7000, only occur via FC.

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 thereby leverages 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 executes 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 TCP/IP 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) these elements:

- ► The string "iqn".
- ► A date code specifying the year and month in which the organization registered the domain or sub-domain 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:

ign.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:

ign.1991-05.com.microsoft:<computer name>

The IQNs can be abbreviated used 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 have 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*.

Be careful: Before changing cluster or node names for an IBM Storwize V7000 cluster that has servers connected to it by way of SCSI, be aware that because the cluster and node name are part of the IBM Storwize V7000's IQN, you can lose access to your data by changing these names. The IBM Storwize V7000-GUI will display a specific warning, 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 will be 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 in configuring iSCSI refer to Chapter 4, "Host Configuration" on page 127.

1.6 Advanced Copy Services

The IBM Storwize V7000 supports the following copy services:

- Synchronous remote copy
- Asynchronous remote copy
- FlashCopy

Copy services functions are implemented within a single IBM Storwize V7000 or between multiple IBM Storwize V7000s. 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.

1.6.1 Synchronous/Asynchronous remote copy

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

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. It is necessary to consider the distance and available bandwidth of the intersite links. Refer to Chapter 11, "Copy Services" on page 381.

In asynchronous remote copy the application is provided acknowledgement that the write is complete prior to the write being 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 bringing up 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 information center:

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 it existed at a single point in time. Although the copy operation takes time, the resulting data at the target appears as though the copy was made instantaneously.

FlashCopy is sometimes described as an instance of a Time-Zero (T0) copy or a Point in Time (PiT) copy technology.

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

IBM Storwize 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, as well as create multiple images from a Source volume at a common point in time. Source and/or 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 having to wait 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:

http://www-01.ibm.com/software/tivoli/products/storage-flashcopy-mgr/

You can read a detailed description of FlashCopy copy services in Chapter 11, "Copy Services" on page 381.

1.6.3 Copy Services configuration limits

In Table 1-7 we describe the Copy Services configuration limits. For the most up to date list of these limits refer to:

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

Table 1-7 Copy Services configuration limits

| Properties | Maximum number | Note |
|--|----------------|---|
| Remote Copy (Metro Mirror and Global Mirror) relationships per cluster | 2040 | 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 cluster limit |
| Remote Copy consistency groups per cluster | 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. |
| FlashCopy mappings per cluster | 4096 | |
| FlashCopy targets per source | 256 | |
| Cascaded Incremental FlashCopy maps | 4 | A volume can be the source of up to 4 incremental FlashCopy maps. If this number of maps is exceeded then the FlashCopy behaviour for that cascade becomes non-incremental. |
| FlashCopy mappings per consistency group | 512 | |
| FlashCopy consistency groups per cluster | 127 | |
| Total FlashCopy volume capacity per I/O group | 1024 TB | |

1.7 Management and Support Tools

The IBM Storwize V7000 system can be managed through the IBM Systems Storage Productivity Center (SSPC) or by using the native management software which runs in the hardware itself. The IBM System Storage Productivity Center (SSPC) is an integrated hardware and software solution that provides a single point of entry for managing IBM Storwize V7000 clusters, IBM System Storage DS8000 systems, and other components of your data storage infrastructure.

SSPC simplifies storage management in the following ways:

- Centralizing the management of storage network resources with IBM storage management software
- Providing greater synergy between storage management software and IBM storage devices
- Reducing the number of servers that are required to manage your software infrastructure
- Providing simple migration from basic device management to storage management applications that
- ► Provide higher-level functions

Assist On-site and remote service

The IBM Assist On-site tool is a remote desktop-sharing solution that is offered through the IBM Web site. 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 the 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 SSPC or master console must be able to access the Internet. The following Web site provides further information about this tool:

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

When you access the Web site, 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 onto your SSPC or master console to connect you and your IBM service representative to the remote service session. The IBM Assist On-site contains several layers of security to protect your applications and your computers.

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

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

Event notifications

IBM Storwize V7000 can use Simple Network Management Protocol (SNMP) traps, syslog messages, and Call Home email to notify you and the IBM Support Center when significant events are detected. Any combination of these notification methods can be used simultaneously.

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.

SNMP traps

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 IBM Storwize V7000 software.

Syslog messages

The syslog protocol is a standard protocol for forwarding log messages from a sender to a receiver on an IP network. The 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.

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 functionality 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 cluster 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 Web Links

The IBM Storwize V7000 Support page is at this Web site:

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

The IBM Storwize V7000 Home page is at this Web site:

http://www-03.ibm.com/systems/storage/news/center/storwize v7000/index.html

The IBM Storwize V7000 Supported hardware list is at this Web site:

http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003703

The IBM Storwize V7000 Configuration Limit and Restrictions are at this Web site:

 $\label{lem:http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003702\&myns=s028\&mynp=familyind5402112\&mync=E$

The IBM Storwize V7000 Online documentation is at this Web site:

http://www-947.ibm.com/support/entry/portal/Documentation/Hardware/System_Storage/Disk_systems/Mid-range_disk_systems/IBM_Storwize_V7000_%282076%29

The IBM Storwize V7000 Online infocenter is at this Web site:

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

You can see the IBM Redbooks publications about IBM Storwize V7000 at this Web site:

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

1.8.1 IBM Storwize V7000 learning videos on youtube

Some interesting videos are available on youtube that describe the IBM Storwize V7000 GUI and are available at the web sites shown in Table 1-8:

Table 1-8 Videos available on youtube

| Description of the video | Link WEB | | | |
|---|--|--|--|--|
| IBM Storwize V7000 Storage Virtualization Terminology Overview | http://www.youtube.com/watch?v=I2rzt3m 2gP0 | | | |
| IBM Storwize V7000 Interface tour | http://www.youtube.com/watch?v=FPbNRs9 HacQ | | | |
| IBM Storwize V7000 Volume management | http://www.youtube.com/watch?v=YXeKqH8 Sd9o | | | |
| IBM Storwize V7000 Migration | http://www.youtube.com/watch?v=dXxnUN6 dk74 | | | |
| IBM Storwize V7000 Introduction to Flashcopy | http://www.youtube.com/watch?v=MXWgGWj BzG4 | | | |
| IBM Storwize V7000 Introduction to Remote Copy 1 | http://www.youtube.com/watch?v=koFYm gnEc | | | |

| Description of the video | Link WEB | | |
|--|--|--|--|
| IBM Storwize V7000 Introduction to Remote Copy 2 | http://www.youtube.com/watch?v=kMOFFS3 NYas | | |



2

Initial configuration

This chapter provides a description of the IBM Storwize V7000 initial configuration steps.

In this chapter, we discuss these topics:

- ► Planning for IBM Storwize V7000 installation
- First time setup
- Initial configuration steps
- ► Call Home, Email event alert and inventory setting

2.1 Hardware Installation Planning

| oper planning prior to the actual physical installation of the hardware is required. Below is a ecklist of requirements to use in preparation. |
|--|
| Install the hardware as described in the <i>IBM Storwize V7000 Quick Installation Guide Version 6.1.0</i> , GC27-2290-00 |
| A physical location for the equipment with adequate cooling available as recommended in the IBM Storwize V7000 Planning Guide. |
| An appropriate 19 inch rack with 2-20U of space available depending 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. |
| Redundant power outlets in the rack for each of the two power cords included per enclosure. From 2-20 outlets are required depending on the number of enclosures to install. The power cords conform to the IEC320 C13/C14 standards. |
| A minimum of 4 Fibre Channel ports attached to the fabric are required, but it is recommended to use eight 2/4/8 Gbps Fibre Channel ports. |
| Note: FC ports are required only if FC hosts or remote mirroring will be used. You can use the Storwize V7000 with Ethernet-only cabling for iSCSI hosts. |
| Eight 2/4/8Gbps compatible Fibre Channel cable drops. |
| A minimum of two Ethernet ports on the LAN with four preferred for additional configuration access redundancy and/or iSCSI host access. |
| A minimum of two Ethernet cable drops with four preferred for additional configuration access redundancy and/or iSCSI host access. Ethernet port one on each node canister must be connected to the LAN with port two being optional. |
| Note: 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. |
| A minimum of three IPv4 or IPv6 IP addresses for cluster configuration. One is for the cluster and is what the administrator will use for management and one for each node canister for service access as needed. |
| Note: A fourth IP address is recommended for backup configuration access. This will allow a second cluster IP address to be configured on port two of either node canister that 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 will access volumes from the IBM Storwize V7000. |
| Two one, three or six 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 and/or other expansion enclosures. The recommendation is to locate the control enclosure such that four enclosures can be located above it and five enclosures below it as shown in Figure 2-1.

Note: The disk drives included with the control enclosure, model 2076-224 or 2076-212, 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 utilized.

Figure 2-1 on page 38 shows the recommended racking and cabling scheme.

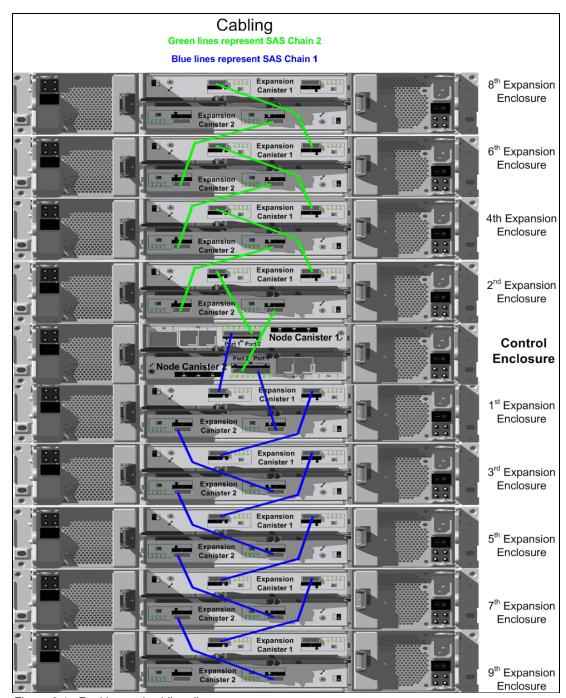


Figure 2-1 Racking and cabling diagram

2.2 SAN Configuration Planning

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

Once IBM Storwize V7000, hosts and optional external storage systems are connected to the SAN fabrics, zoning will need 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.

Note: IBM Storwize V7000 supports a maximum of sixteen ports or WWPNs from a given external storage system that will be virtualized.

Figure 2-2 on page 39 is an example of how to cable devices to the SAN. Reference this example as we discuss the recommended zoning below.

Create a Host/IBM Storwize V7000 zone for each server that volumes will be mapped to from the cluster. For example:

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

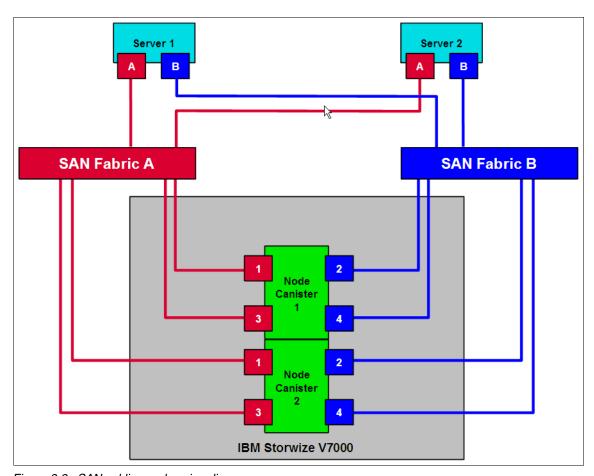


Figure 2-2 SAN cabling and zoning diagram

Verify that the SAN switches and/or directors that the IBM Storwize V7000 will connect to meet the following requirements as noted at:

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

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

Note: If you have any connectivity issues between IBM Storwize V7000 ports and Brocade SAN Switches or Directors at 8 Gbps, refer to the Web site at:

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

for the correct setting of the **fillword port config** parameter in the Brocade operating system.

2.3 LAN Configuration Planning

There are two Ethernet ports per node canister 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 management of the cluster as well as 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.

Note: There is no issue with configuring multiple IPv4 and/or IPv6 addresses on a given Ethernet port nor using the same Ethernet port for management and iSCSI access. However, you can not use the same IP address for both management and iSCSI host use.

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

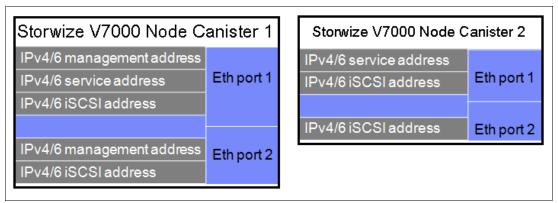


Figure 2-3 Example of IP address configuration options per node canister

Note: The IP management address shown on node canister 1 is an example of the configuration node, in case of failover these addresses will transfer to node canister 2 and this node canister will be the new configuration node. Be aware that 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

With Ethernet port 1 from each node canister required to be connected to the LAN, a single management IP address for the cluster will be configured as part of the initial setup of the IBM Storwize V7000 system.

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

For additional cluster management redundancy the recommendation is to connect Ethernet port 2 on each of the node canisters to the LAN. This 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 cluster itself and therefore only associated with one node which is then considered the configuration node.

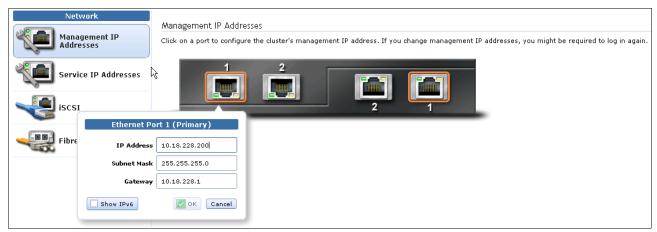


Figure 2-4 Ethernet ports available for configuration

2.3.2 Service IP Address Considerations

Ethernet port 1 on each node canister is used for cluster management as discussed above and also used for service access when required. In normal operations the service IP addresses will not be needed. However, in the event of a node canister problem it may be necessary for service personnel to log in to the node to perform service actions.

By default the service IP address on node canister 1 is 192.168.70.121 and on node canister 2 it is 192.168.70.122. 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 does support hosts with a single HBA port connection. However, if that HBA or its link to the

SAN fabric or the fabric itself fails then the host will lose access to its volumes. Even with just a single connection to the SAN the host will have 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 and requirements are listed on the IBM Storwize V7000 support matrix and configuration instructions are provided in the IBM Storwize V7000 Host Attachment Guide.

Verify that the hosts which will access volumes from the IBM Storwize V7000 meet the following requirements as found at the URL below.

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

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

For further details see Chapter 4, "Host Configuration" on page 127.

2.5 Miscellaneous Configuration Planning

During initial setup of the IBM Storwize V7000 system the installation wizard will ask for various information that you will want to have available during the 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 via the GUI. Below is a checklist of information to gather before performing the initial setup.

The date and time can be manually entered but to keep the clock synchronized the recommendation is to use the network time protocol (NTP) service.

Document the current LAN NTP server IP address used for synchronization of devices.

If external storage systems will be used by IBM Storwize V7000 then a license must be purchased to cover the number of enclosures to be virtualized.

Document the number of external physical enclosures that will be virtualized under the IBM Storwize V7000.

Document the total number of physical enclosures, both internal and external virtualized enclosures for the GM and MM feature. You must have enough GM and MM license entitlement for all enclosures attached to the system, regardless of the amount of GM and MM capacity you intend to use..

The Easy Tier function is included with the IBM Storwize V7000 system and is not a purchased feature. if the system has solid state drives, SSDs, and this capability is to be used to optimize the utilization of the SSD drives then this function will be enabled.

For alerts to be sent to storage administrators and to set up call home to IBM for service and support you will need the information below:

☐ Name of primary storage administrator for IBM to contact if necessary.

| | E-mail address of the above storage administrator for IBM to contact if necessary. |
|-----|--|
| | Phone number of the above storage administrator for IBM to contact if necessary. |
| | Physical location of the IBM Storwize V7000 system for IBM service. (E.g. Building 22, 1st floor) |
| | SMTP or E-mail server address to direct alerts to and from the IBM Storwize V7000. |
| | For call home service to work, the IBM Storwize V7000 system must have access to a SMTP server on the LAN that can forward E-mails to the default IBM service address 'callhome0@de.ibm.com'. |
| | E-mail address of local administrators needing to be notified of alerts. |
| | IP address of SNMP server to direct alerts to if desired. (E.g. Operations or help desk) |
| cai | ter the IBM Storwize V7000 initial configuration you may want to add additional users who n manage the cluster. You can create as many users as you need but currently there are ly three roles generally configured for users, administrator, copyoperator and monitor. |

The administrator role allows the user to perform any function on the IBM Storwize V7000 system except create users.

Note: Create users is allowed by the superuser role only and should be limited to as few users as possible.

The copyoperator role allows the user to view anything in the system but the user can only configure and manage copy functions which include the replication functions of Metro and Global Mirror and the FlashCopy capabilities.

The monitor role allows the user to view anything in the system but they can not create, modify or change anything in the system nor perform any actions that change the state of the cluster.

The only other role available is the service role which would generally be used if you create a user ID for the IBM service representative. This role allows IBM service personnel to view anything on the system just as the monitor role provides, plus perform service related commands such as adding a node back to the cluster after being serviced.

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 the logical drives.

It allows for troubleshooting and management tasks, such as checking the status of the storage server components, updating the firmware, and managing the storage server. Finally, it offers advanced functions, such as FlashCopy, Volume Mirroring, and Remote Mirroring. A Command Line Interface (CLI) for the IBM Storwize V7000 system is available as well.

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

2.6.1 Graphical user interface (GUI)

A Web browser is used for GUI access and you should open your favorite Web browser on a workstation and point to the IBM Storwize V7000 system management address, and we show how this is defined in 2.7, "First time setup" on page 46.

After the initial configuration as described in 2.8, "Initial configuration" on page 52 you get the IBM Storwize V7000 system welcome screen as shown in Figure 2-6 on page 45

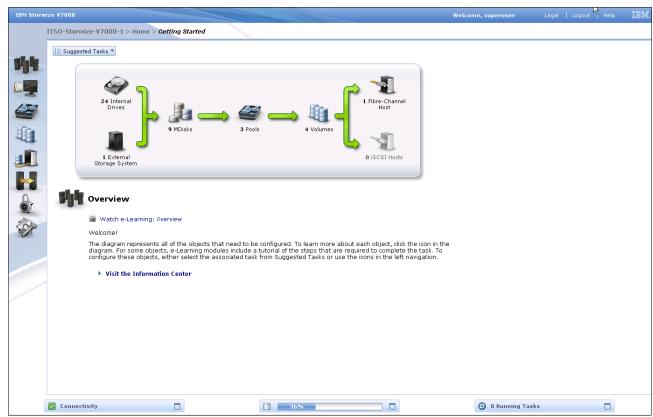


Figure 2-6 IBM Storwize V7000 Welcome screen

2.6.2 Command line interface

The Command Line Interface (CLI) is a very flexible tool for system management and uses the SSH protocol. A public / private SSH key pair is required for SSH access. The Appendix A, "CLI setup and SAN Boot" on page 543 describes how to set up SSH access for MS Windows or Unix systems. As a result the storage system can be managed using the CLI as shown in Example 2-1 on page 45.

Example 2-1 System management using the Command Line Interface (CLI)

| ΙBI | IBM_2076:ITSO-Storwize-V7000-1:admin>svcinfo lsuser | | | | | |
|---------------------------------------|---|----------|---------|--------|------------|---------------|
| id | name | password | ssh_key | remote | usergrp_id | usergrp_name |
| 0 | superuser | yes | yes | no | 0 | SecurityAdmin |
| 1 | MASSIMO | yes | no | no | 1 | Administrator |
| 2 | JASON | yes | no | no | 1 | Administrator |
| 3 | DAN | yes | no | no | 1 | Administrator |
| IBM_2076:ITSO-Storwize-V7000-1:admin> | | | | | | |

The initial IBM Storwize V7000 system setup should be done using the graphical tools we describe starting in 2.7, "First time setup" on page 46.

2.7 First time setup

This chapter describes how to perform a first time IBM Storwize V7000 system setup.

The IBM Storwize V7000 provides an easy to use initial setup contained within a USB key. The USB key is delivered with each storage system and contains the initialization application called "InitTool.exe". A system management ip address, the subnet mask and the network gateway address are required. The initialization application creates a configuration file on the USB key.

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

Note: In case you are unable to find the official USB key supplied with the IBM Storwize V7000 you can use any USB key you have, and download and copy the initTool.exe application from the IBM Storwize V7000 Support site at http://www.ibm.com/storage/support/Storwize/V7000

The USB key contains the initTool.exe file as shown in Figure 2-7 on page 46.



Figure 2-7 InitTool program in the USB key

These are the steps to follow to perform the initial setup using the USB key.

1. Initially plug in the USB key to an MS Windows system and start the initialization tool, if the system is configured to autorun for USB keys, the initialization tool starts automatically, otherwise, open the USB key from My Computer and double-click on the InitTool.exe file. The tool is shown in Figure 2-8 on page 47, after the tool is started select "Initialize a new system using the USB Key" and click next as shown in Figure 2-8 on page 47.

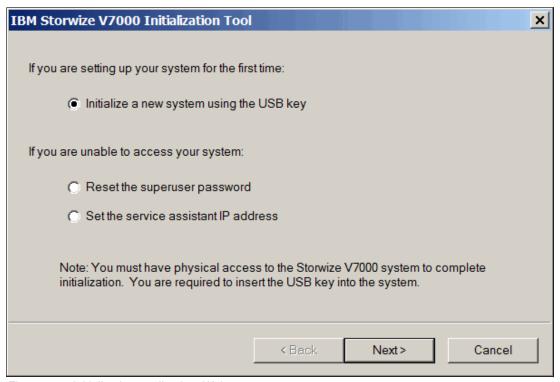


Figure 2-8 Initialization application: Welcome screen

2. Type the IPV4 or IPV6 address, subnet mask and network gateway address, then click next as shown in Figure 2-9.

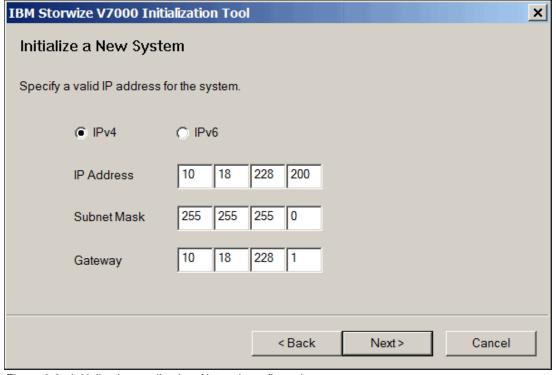


Figure 2-9 Initialization application: Network configuration

3. Click Next to finish the initial application as shown in Figure 2-10 on page 48.

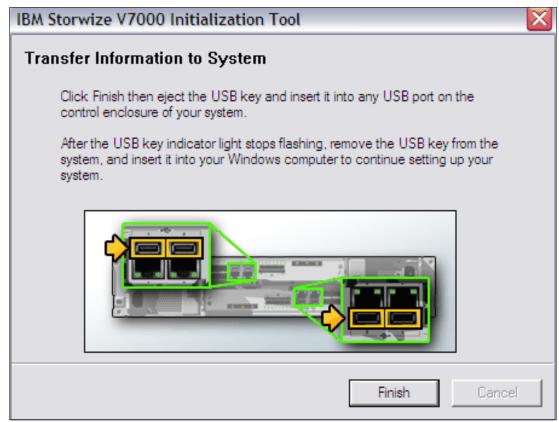


Figure 2-10 Initialization application: finalize the initial setup

After these steps the application creates a new file called satask.txt on the USB key as shown in Figure 2-11 on page 48.

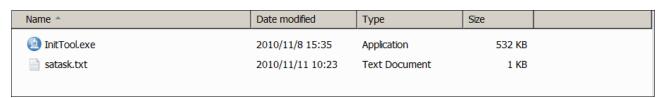


Figure 2-11 File satask.txt created during InitTool procedure

The content of this file will be similar to the following command:

satask mkcluster -clusterip 10.18.228.200 -gw 10.18.228.200 -mask 255.255.255.0

- 4. Unplug the USB key from your Windows system and plug it into the IBM Storwize V7000 USB connector slot and the storage system will automatically:
 - a. Detect the Key.
 - b. Read the satask.txt file and execute the content.
 - c. Delete the satask.txt file.
 - d. Write the response to a new file called satask_result.html as shown in Figure 2-12 on page 49.

Note: While the cluster 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.

| Name ^ | Date modified | Туре | Size | |
|--------------------|------------------|---------------|--------|--|
| InitTool.exe | 2010/11/8 15:35 | Application | 532 KB | |
| satask_result.html | 2010/11/11 18:24 | HTML Document | 4 KB | |
| | | | | |

Figure 2-12 File satask_result.html created during InitTool procedure

After this has completed successfully the initial setup is successfully 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 Cluster has one or two cluster ip addresses. If the configuration node fails, the cluster IP addresses are transferred to another node in the same cluster.

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 content of the satask_result.html file is shown in Example 2-2.

Example 2-2 satask_result.html file content.

```
Service Command Results
satask mkcluster -clusterip 10.18.228.200 -gw 10.18.228.1 -mask 255.255.255.0
Thu Nov 11 10:34:51 PST 2010
System Status
sainfo lsservicenodes
panel name cluster id
                            cluster_name
                                                  node_id node_name relation
node status error data
32GOCXR-1 0000020060E14FFC Cluster_10.18.228.200 1
                                                           32GOCXR-1 local
Starting
32G0CXR-2
                                                                     partner
Candidate
sainfo lsservicestatus
panel name 32G0CXR-1
cluster id 0000020060e14ffc
cluster name Cluster 10.18.228.200
cluster_status Active
cluster_ip_count 2
```

```
cluster port 1
cluster_ip 10.18.228.200
cluster gw 10.18.228.1
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 32G0CXR-1
node status Starting
config node No
hardware 100
service IP address 192.168.70.121
service gateway 192.168.70.1
service subnet mask 255.255.255.0
service IP address 6
service_gateway_6
service prefix 6
node sw version 6.1.0.0
node sw build 49.0.1010130001
cluster sw build 49.0.1010130001
node_error_count 0
error code
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 Inactive
port speed N/A
port WWPN 500507680110a7fe
SFP type Short-wave
port id 2
port status Inactive
port speed N/A
port WWPN 500507680120a7fe
SFP type Short-wave
port_id 3
port_status Inactive
port speed N/A
port WWPN 500507680130a7fe
SFP_type Short-wave
```

```
port id 4
port_status Inactive
port speed N/A
port WWPN 500507680140a7fe
SFP type Short-wave
ethernet ports 2
ethernet port id 1
port status Link Online
port speed 1Gb/s - Full
MAC e4:1f:13:74:0a:81
ethernet port id 2
port status Not Configured
port speed
MAC e4:1f:13:74:0a:80
product mtm 2076-124
product serial 32GOCXR
time to charge 0
battery charging 100
dump name 32GOCXR-1
node WWNN
disk WWNN suffix
panel WWNN suffix
UPS serial number
UPS status active
enclosure WWNN 1 500507680100a7fe
enclosure WWNN 2 500507680100a7be
node part identity 11S85Y5849YHU9995G051A
node FRU part 85Y5899
enclosure identity 11S85Y5962YHU9992GOCXR
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 32GOCXR
node WWNN 1 copy 500507680100a7fe
node WWNN 2 copy 500507680100a7be
latest cluster id 20060e14ffc
next_cluster_id 20061014ffc
sainfo lsservicerecommendation
service action
No service action required, use console to manage node.
```

2.8 Initial configuration

This section describes how to perform the initial configuration including:

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

Follow these steps:

 Start the configuration wizard using a Web browser on a workstation and point to the system management IP address defined in Chapter 2.7, "First time setup" on page 46.
 Type in the default superuser password passw0rd> (0=zero) as shown in Figure 2-13 on page 52.



Figure 2-13 Setup wizard: login screen

2. Read and accept the license agreement as shown in Figure 2-14 on page 53.

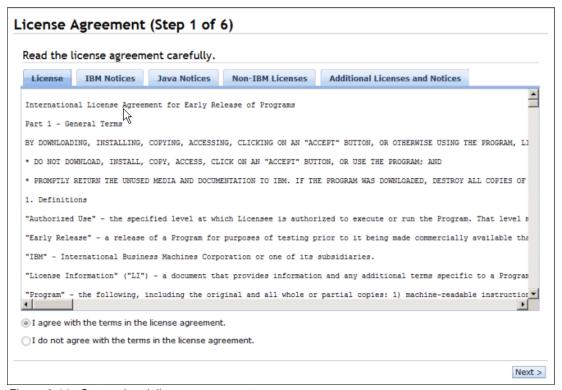


Figure 2-14 Setup wizard: license agreement

3. Set up the system name. current date, and time as shown in Figure 2-15.

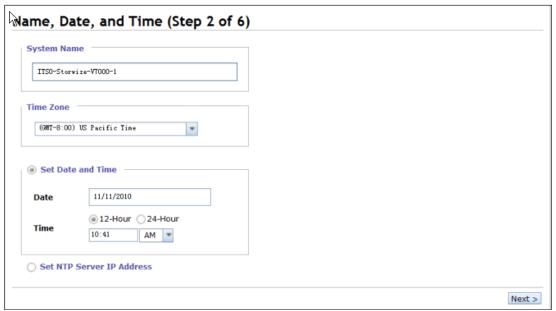


Figure 2-15 Setup wizard: set name, date and time

4. Optionally you can type in advanced licenses for virtualization of external storage devices and a remote copy limit as applicable. The virtualization license for all local devices is already included in the system and must not be added here.

Figure 2-16 shows an example of how to add it for external storage and remote copy.

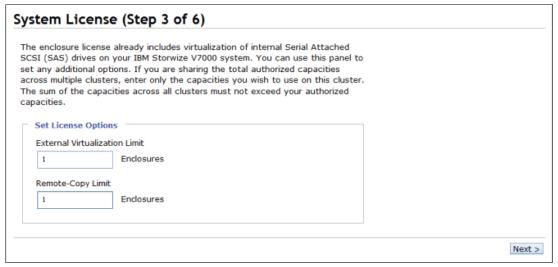


Figure 2-16 Setup wizard: license setting

2.8.1 Configure Call Home, E-Mail alert and Inventory

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

Note: When configuring email call home to IBM Support, use one of the following email addresses depending on country or region of installation:

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

To configure the call home and E-Mail alert event notification in the IBM Storwize V7000 perform the following steps.

Clicking **Configure Email Event Notification** as shown in Figure 2-17 on page 54 starts an Email configuration wizard.



Figure 2-17 Setup wizard option E-mail and event configuration:

Define the sending Email account as shown in Figure 2-18 on page 55.

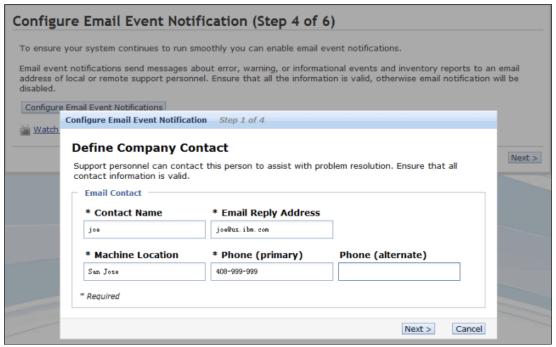


Figure 2-18 E-mail and event configuration: set contact information

Configure the Email server as shown in Figure 2-19 on page 55.

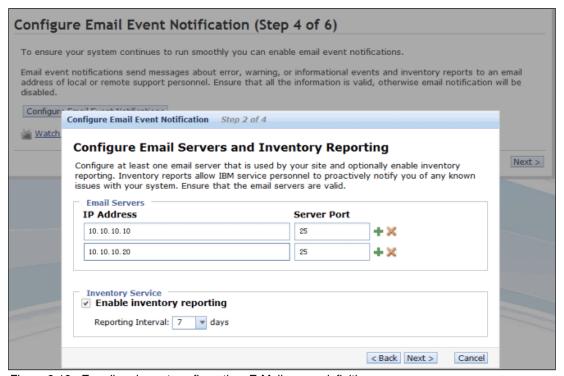


Figure 2-19 E-mail and event configuration: E-Mail server definition

Type in the receiving Email address for local users and for support and select the appropriate message type for each user as shown in Figure 2-20 on page 56.

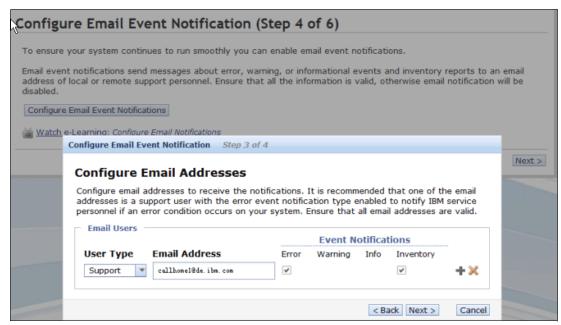


Figure 2-20 Email and Event configuration: set E-Mail address and event level for call home

Click on **Next** verify the Email configuration and click **Finish** as shown in Figure 2-21 on page 56.

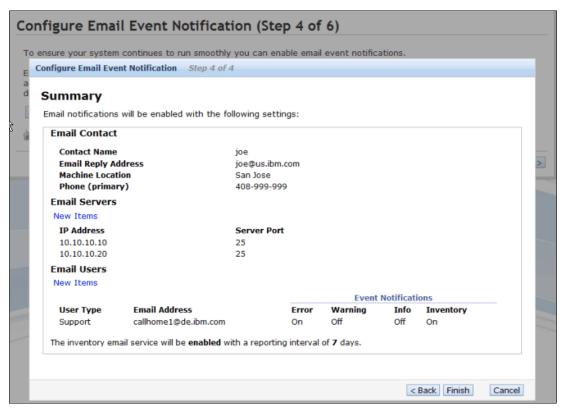


Figure 2-21 Email and event configuration: verify Email addresses and finish configuration

The Email Event Notification Setup is successfully done; the Setup wizard shows the new Email configuration.

You can modify the setting, discard or continue with the configuration by clicking on **Next** as shown in Figure 2-22 on page 57.

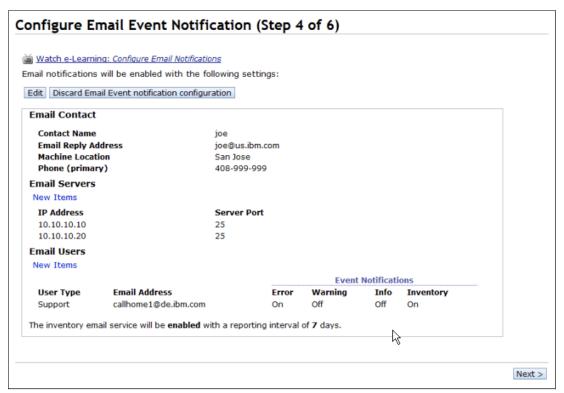


Figure 2-22 Setup wizard finalize E-Mail setting

The configuration Wizard continues with the hardware configuration.

Verify the hardware as shown in Figure 2-23 on page 58.

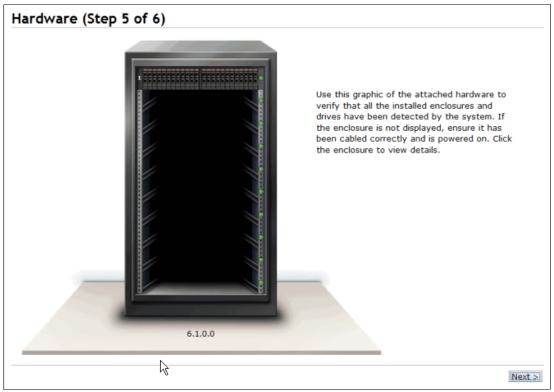


Figure 2-23 Setup wizard: verify the detected hardware

After you click on **Next** an **Add Enclosure** window appears as shown in Figure 2-24 on page 58.

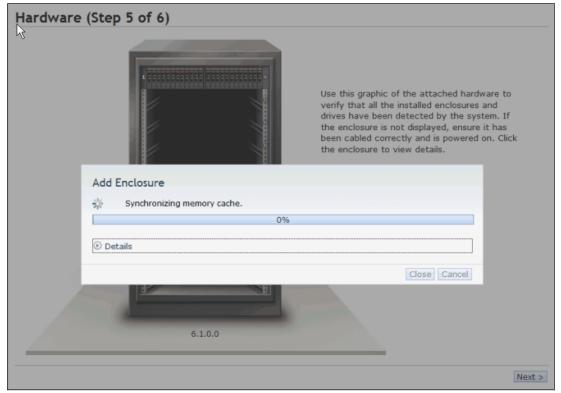


Figure 2-24 Setup wizard: synchronizing memory cache

Hardware (Step 5 of 6)

Use this graphic of the attached hardware to verify that all the installed enclosures and drives have been detected by the system. If the enclosure is not displayed, ensure it has been cabled correctly and is powered on. Click the enclosure to view details.

Add Enclosure

The task completed.

100%

Close Cancel

Next ➤

Wait until The task completed screen appears as shown in Figure 2-25 on page 59.

Figure 2-25 Setup wizard: task completed

Finally the configuration wizard detects the internal disk drives and offer you a storage configuration for all internal disk drives.

- ► If the proposed configuration is acceptable to you, select the check box and click **Finish** to finalize the initial configuration as shown in Figure 2-26 on page 60.
- ▶ If you prefer a customized configuration leave the check box blank and click **Finish** and continue with the storage configuration using Advanced Storage Configuration described in Chapter 8, "Advanced Host and Volume Administration" on page 273 and Chapter 11, "Copy Services" on page 381.
- ▶ Optionally you can start the migration wizard and let the system guide you.

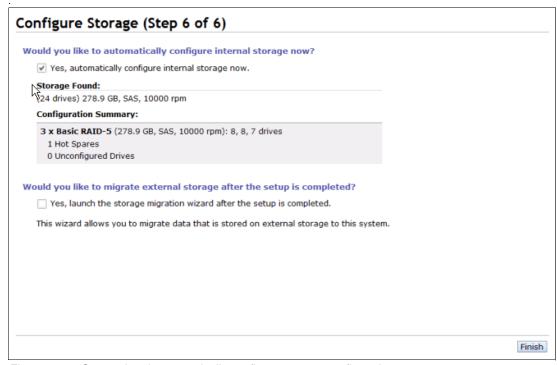


Figure 2-26 Setup wizard: automatically configure storage configuration

After you click on **Finish** a "Synchronizing memory cache window appears as shown in Figure 2-27 on page 60.

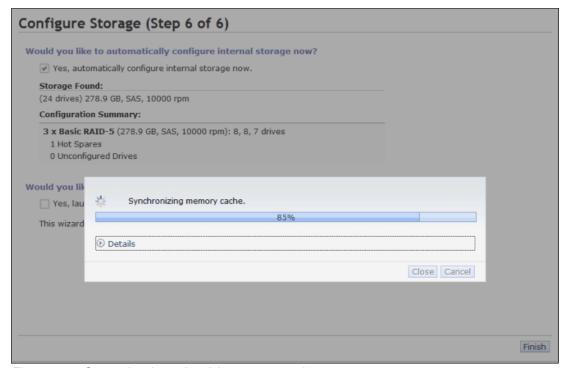


Figure 2-27 Setup wizard: synchronizing memory cache

Wait until The task completed window appears as shown in Figure 2-28 on page 61.

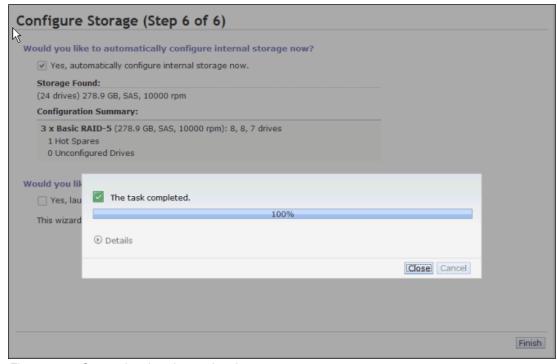


Figure 2-28 Setup wizard: task completed

You have now successfully finalized the initial configuration wizard for IBM Storwize V7000 system. Figure 2-29 on page 61 shows the **Getting Started** page and is ready for use.

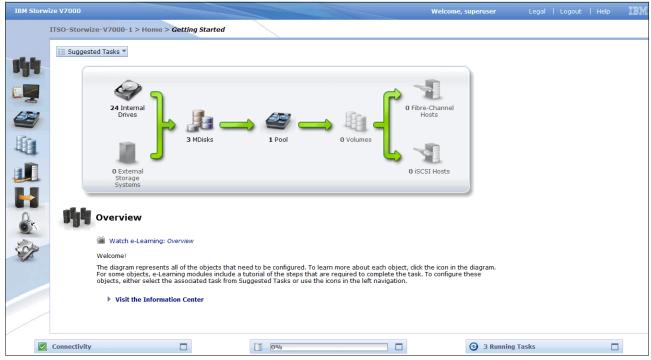


Figure 2-29 IBM Storwize V7000: Getting started page

| 7938 IBM Storwize V7000 Initial Configuration Massimo.fm Draft Document for Review March 3, 2011 2:46 pm | |
|--|--|
| | |
| | |
| | |
| | |
| | |



3

Graphical User Interface overview

This chapter provides an overview of the graphical user interface for IBM Storwize V7000 and covers how to use the navigation tools.

3.1 Overview of IBM Storwize V7000 management software

The new graphical user interface is designed to 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.

Note: You may need to enable JavaScript in your browser. Additionally, if using Firefox under Advanced JavaScript Settings you will need to Disable or replace context menus, and allow cookies.

Access to the storage management software

To login to the management software point your web browser at the IP Address that was used during the initial setup process. The login screen will appear as shown in Figure 3-1.

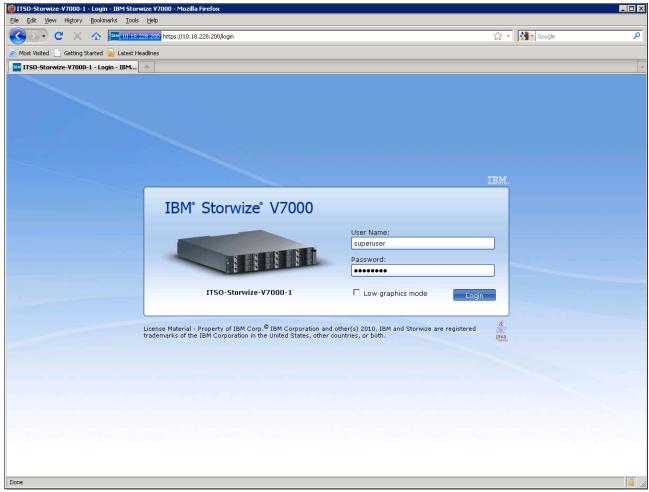


Figure 3-1 Login Screen

Once logged in successfully the Getting Started screen is displayed. This is shown in Figure 3-2 on page 65.

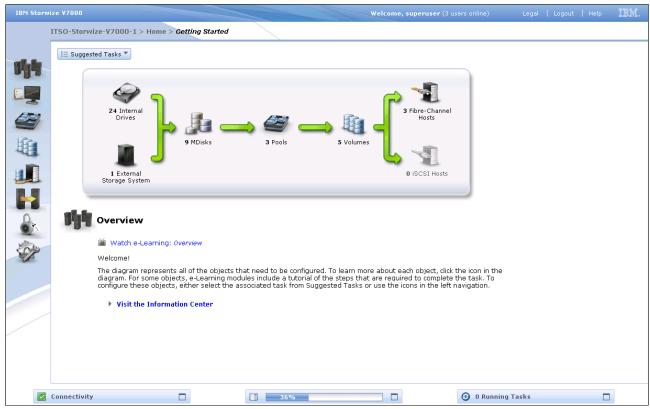


Figure 3-2 Home menu showing the Getting Started panel

Screen layout

This screen has three main sections for navigating through the management tool. On the far left hand side of the screen are eight Function Icons. The eight icons represent:

- ► Home menu
- ► Troubleshooting menu
- Physical Storage menu
- ▶ Volumes menu
- ► Hosts menu
- Copy Services menu
- User Management menu
- Configuration menu

In the middle of the screen is a diagram illustrating the existing configuration. Clicking on the icons in this area will provide extended help references including a link to a short video presentation to explain the topic in more detail. This is not a navigation tool but rather an extended help screen which includes configuration information.

At the bottom of the screen are three status indicators. Clicking on these will provide more detailed information about the existing configuration of the IBM Storwize V7000 solution. Simply click on these icons to expand them and minimize them as required.

The diagram in Figure 3-3 on page 66 shows the main screen areas.

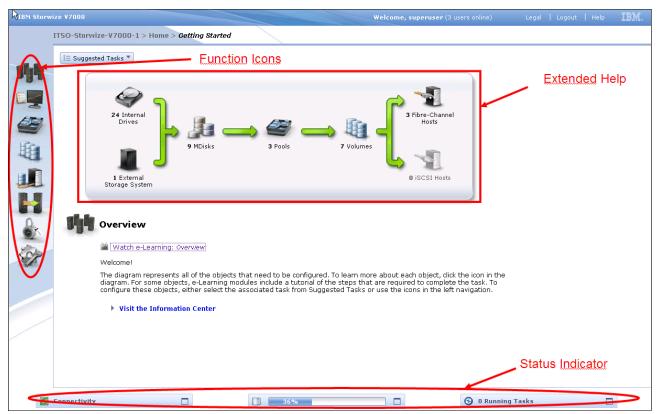


Figure 3-3 Main screen areas

Navigation

Navigating around the management tool is very simple.

You can hover the cursor over one of the eight Function Icons on the left hand side of the screen which will highlight the Function Icon and will then display a list of options. You can then move the cursor to the desired option and select it. This method is illustrated below in Figure 3-4 on page 67 below.

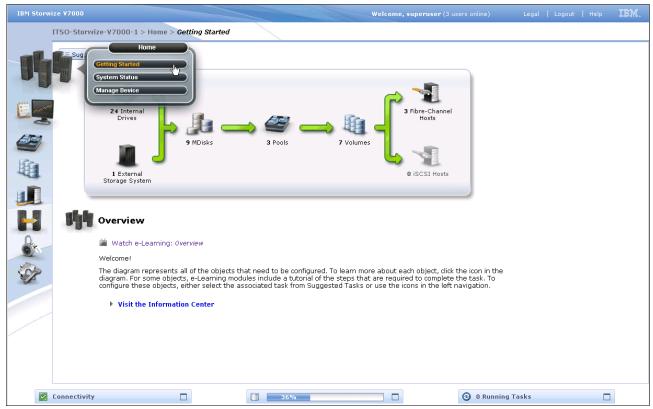


Figure 3-4 Navigation using the menu options

An alternative method is to click on the desired Function Icon. This will navigate you directly to that screen. This is illustrated below in Figure 3-5 on page 68.

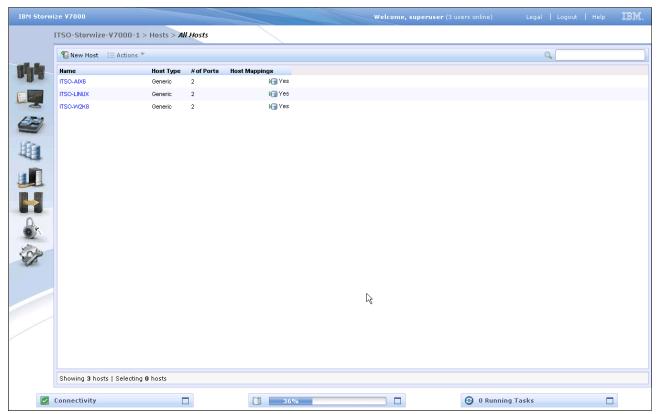


Figure 3-5 Navigation using the menu tabs

Clicking on the top of the screen in **All Hosts** you can change the view as shown in Figure 3-6 on page 69. This apply to any other menu options.

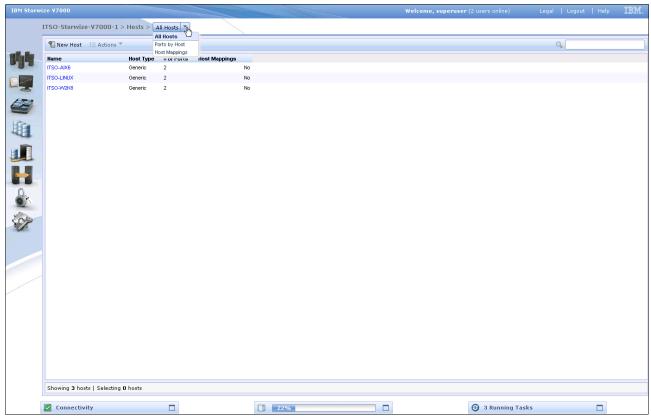


Figure 3-6 Navigation using change view

Shown in Figure 3-7 on page 70 is a list of the IBM Storwize V7000 software Function Icons and the associated menu options.

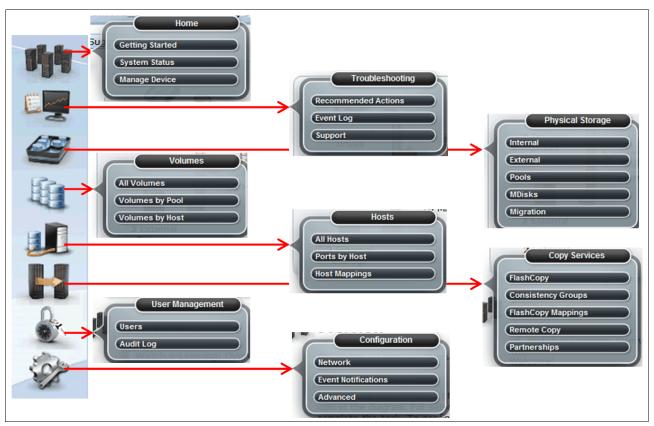


Figure 3-7 BM Storwize V7000 software - menu options

Multiple selections

The new management tool also provides the ability to select multiple items by using a combination of the SHIFT or CTRL keys. To select multiple items in a display click on the first item and then hold down the SHIFT key and click in the last item in the list you require. This will cause all the items in between to be selected as shown in Figure 3-8 on page 71.

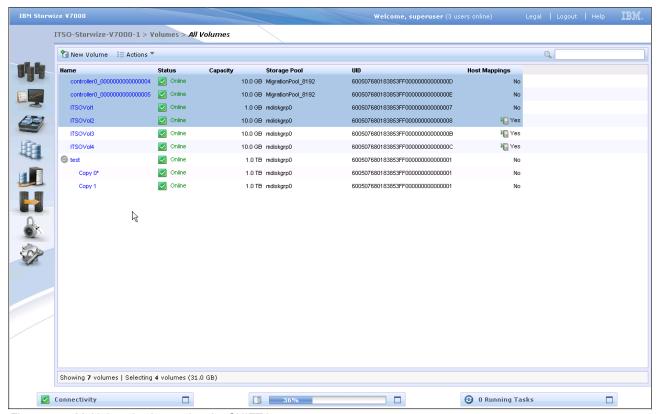


Figure 3-8 Multiple selections using the SHIFT key

If you wish to use multiple select on items that are not in sequential order you can click on the first item and then hold down the CTRL key and click on the other items you require as shown in Figure 3-9 on page 72.

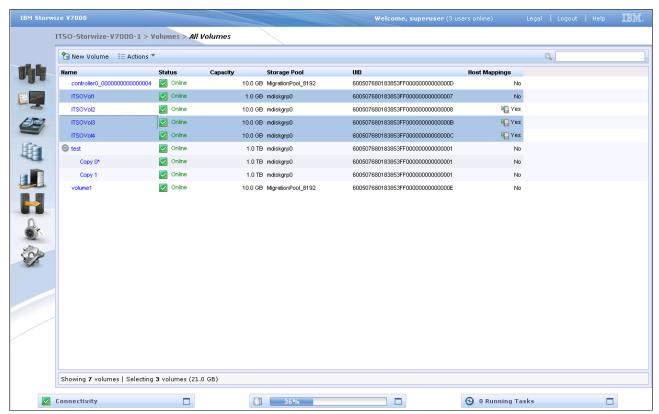


Figure 3-9 Multiple selections using the CTRL key

Status Indicators

Another useful tool is the Status Indicator menus which appear at the bottom of the screen as shown in Figure 3-10 on page 73. These can be maximized and minimized and provide links to connectivity options, storage allocation display and a display of long running tasks.

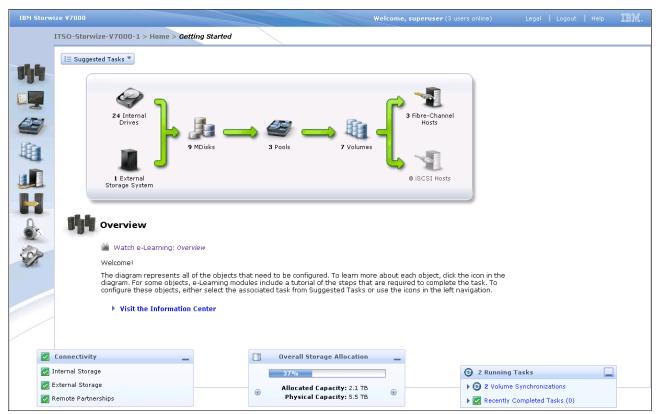


Figure 3-10 Status Indicators

You can see the details of the running task by clicking on the most right status indicator as shown in Figure 3-11 on page 74.

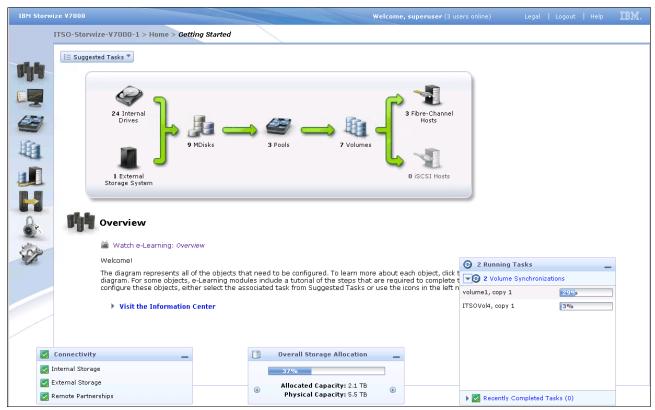


Figure 3-11 Showing running task details

3.2 Home menu

The IBM Storwize V7000 management software provides a very efficient and quick mechanism of navigating between the various different functions. Clicking on one of the eight Function Icons on the left hand side of the screen will cause a menu option to appear which allows you to navigate directly to the selected option.

This is shown in Figure 3-12 on page 75.

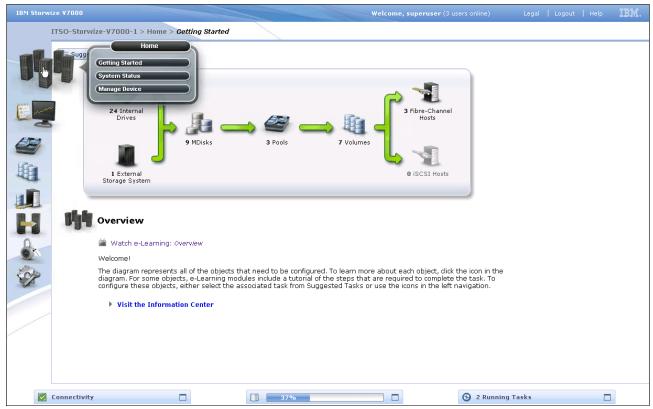


Figure 3-12 Navigating to the Home menu

System Status menu

The System Status menu is a useful screen for identifying configuration information about the IBM Storwize V7000 solution. This screen displays a bar on the left which lists the physical and usable capacity of the system when the cursor hovers over it as shown in Figure 3-13 on page 76. Also displayed is a picture of the IBM Storwize V7000 solution showing the Control enclosure and any Expansion enclosures.

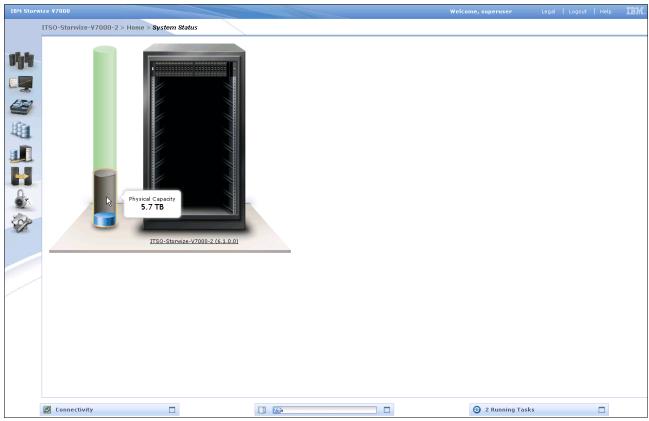


Figure 3-13 System Status screen showing Physical Capacity

By moving the cursor down this bar the Used Capacity is displayed as shown in Figure 3-14 on page 76.

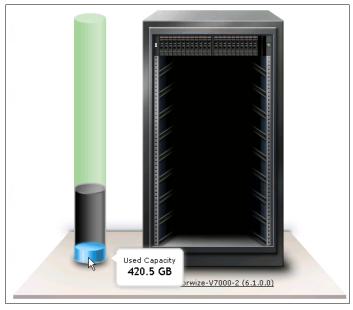


Figure 3-14 System Status screen showing Usable Capacity

By moving the cursor up this bar the Virtual Capacity is displayed as shown in Figure 3-15 on page 77. The virtual capacity is the capacity reported by IBM Storwize V7000 to the host when you are using thin-provisioning volumes.

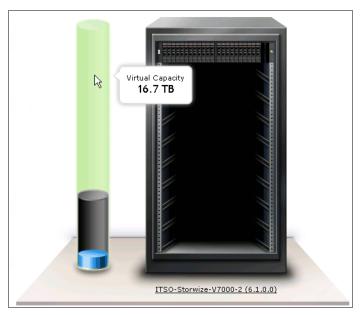


Figure 3-15 System Status screen showing Virtual Capacity

The System Status menu can also display the status of the various IBM Storwize V7000 components. In this example one of the enclosures is reporting a drive problem. By clicking on the enclosure and then hovering the cursor over the affected drive, a status report is displayed for the drive in question. This is shown in Figure .



Figure 3-16 System Status screen showing an offline component

More detailed configuration information is displayed as shown in Figure 3-17 by clicking on the cluster name which is displayed at the bottom of the IBM Storwize V7000 picture.

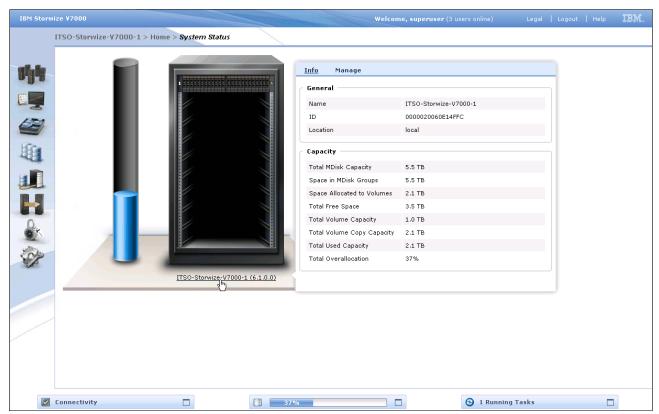


Figure 3-17 System Status screen showing cluster information

3.3 Troubleshooting menu

The following section shows the **Troubleshooting** menu and its options. By hovering the cursor over the Troubleshooting icon the popup menu will appear as shown in Figure 3-18 on page 80.

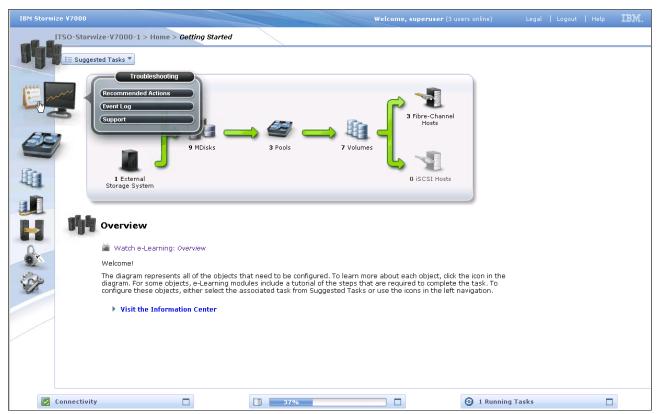


Figure 3-18 Navigating to the Troubleshooting menu

Recommended Actions

Selecting the **Recommended Actions** option will display the screen shown in Figure 3-19 on page 81. This screen lists any events that have occurred that may impact the IBM Storwize V7000 solution. To fix an event you can click on run fix procedure this will run a Directed Maintenance Procedure (DMP) that will walk you through the process to fix that particular event.

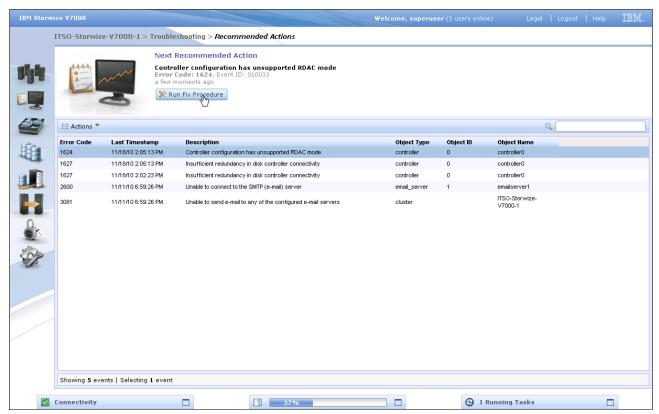


Figure 3-19 Recommended Actions option

Another way to fix an event is right click on the particular event and click on the **Run Fix Procedure**. To go into detail on a specific event click on **Properties** as shown in Figure 3-20 on page 82,

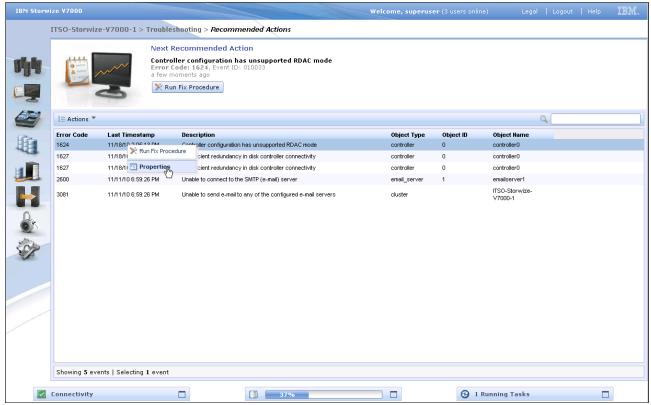


Figure 3-20 Right click on the single Event

Clicking on the **Properties** option you will get the screen shown in Figure 3-21 on page 83.

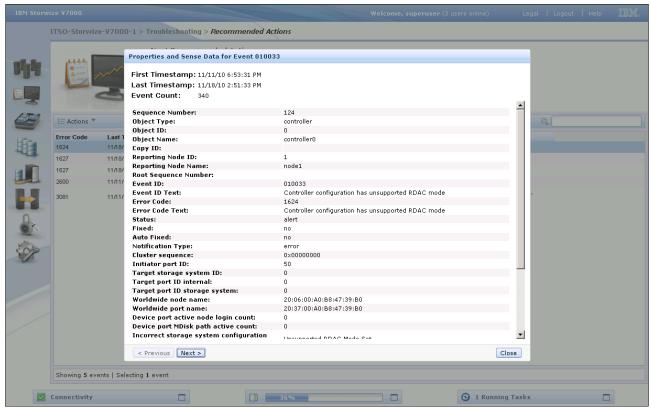


Figure 3-21 Properties of the Event

Event Log

Clicking on the **Event Log** option will display the screen shown in Figure 3-22 on page 84. From this screen you can use the action buttons to either mark an event as fixed or clear the log.

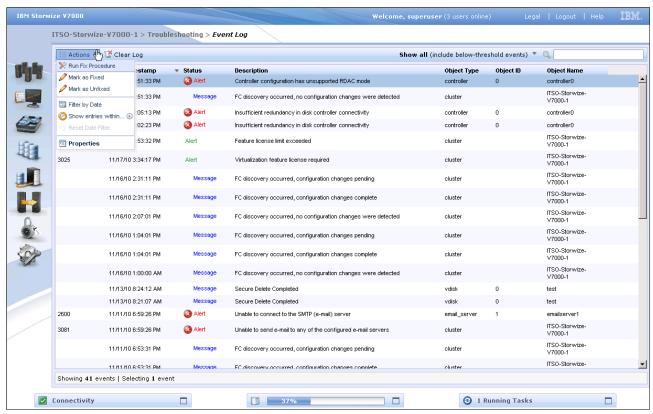


Figure 3-22 Event Log option from the action button

Alternatively, you can select the particular event you are interested in and right click on it which will present the options as shown in Figure 3-23 on page 85.

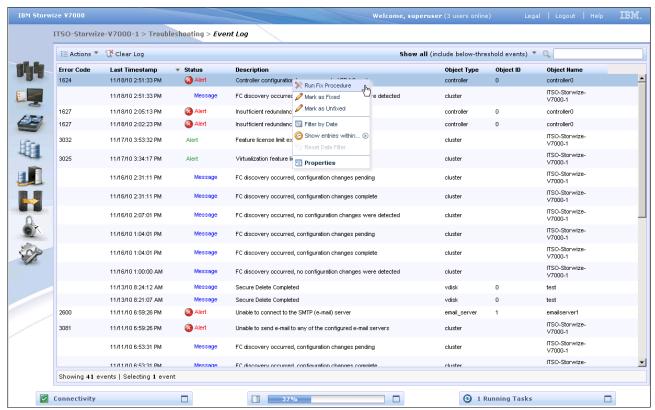


Figure 3-23 Event Log option with right click

Support

Clicking on the **Support** option will display the screen shown in Figure 3-24 on page 86. From this screen click on **Show full log** listing to display all log files.

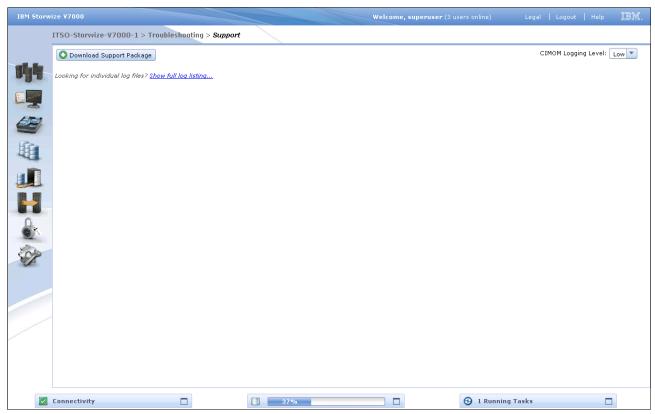


Figure 3-24 Support option

You can download the various log files or delete them selecting a single item as shown in Figure 3-25 on page 87 and click on either the **Download** or **Delete** options in the Actions button.

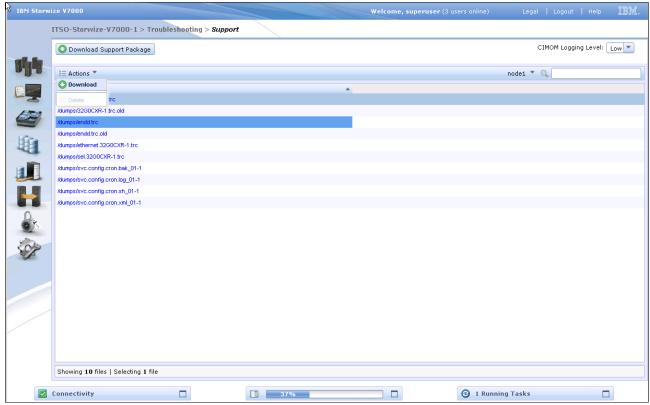


Figure 3-25 Download/Delete options of Actions button

Note: When the delete option is not highlighted, the file can not be deleted because it is a file used by the system.

Towards the upper right side of the screen there is Node option to display node canister 1 or 2 log files as shown in Figure 3-26 on page 88.

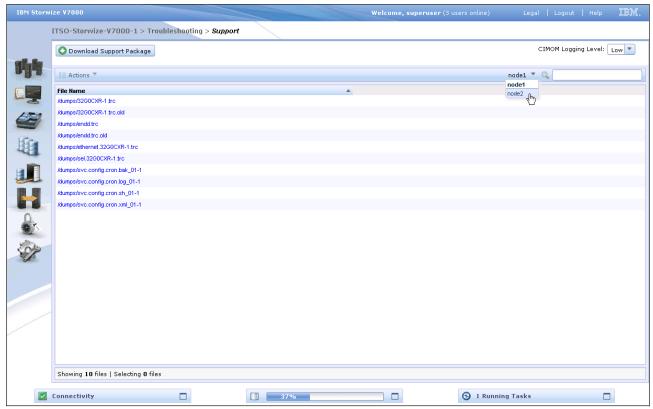


Figure 3-26 Changing log listing of the Nodes Canister

Selecting the **Download Support Package** option will display the screen shown in Figure 3-27 on page 89.

This provides a number of different options to collect logs and state save information from the cluster.

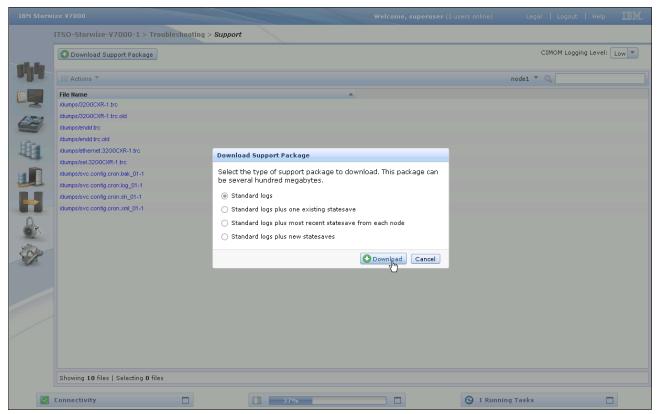


Figure 3-27 Downloading support package

Selecting **Download** generates the support package as shown in Figure 3-28 on page 90.

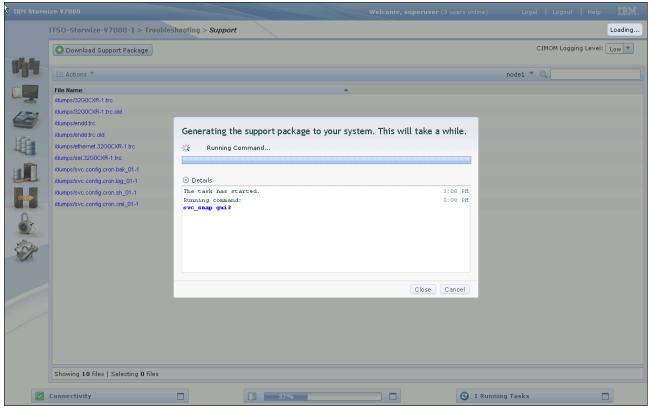


Figure 3-28 Generating support package

Click on **Save File** and then **OK** to save a copy of the package as shown in Figure 3-29 on page 90.

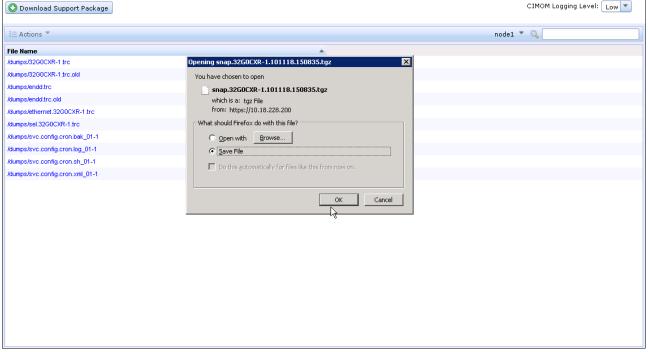


Figure 3-29 Save file

3.4 Physical Storage menu

By moving the cursor onto the Physical Storage Function Icon a subset of menu options will appear as shown in Figure 3-30 on page 91.



Figure 3-30 Navigating to the Physical storage menu

Internal

Selecting the **Internal** option will display the screen shown in Figure 3-31 on page 92. From this screen you can configure the internal disk drives into storage pools. The panel also provides the option to display the internal drives based on their capacity and speed.

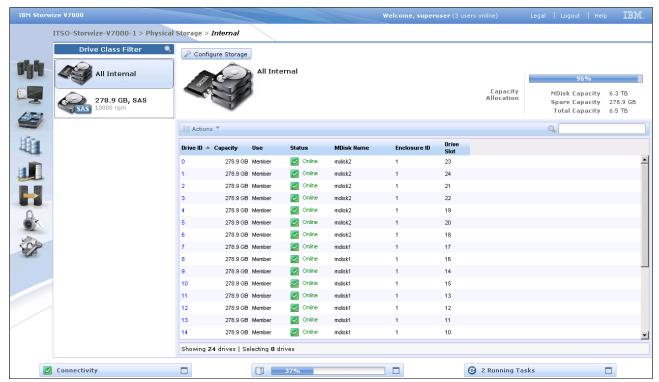


Figure 3-31 Internal storage option

External

Selecting the **External** option will display the screen in Figure 3-32 on page 93. This will display any External Disk Systems that IBM Storwize V7000 is virtualizing. From this screen MDisks can be added to existing pools, imported as image mode volumes or renamed. By highlighting a MDisk you can also display any dependent volumes.

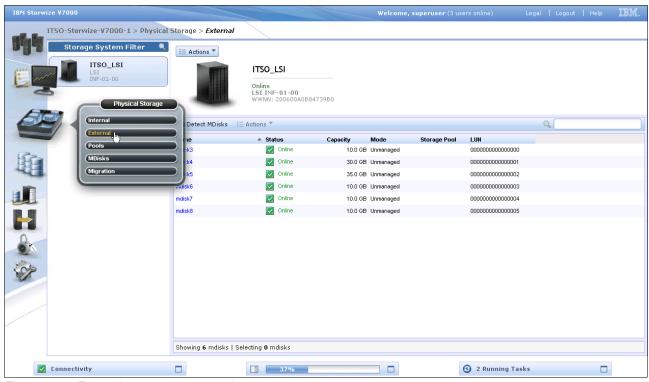


Figure 3-32 External storage systems option

Storage Pools

The screen shown in Figure 3-33 on page 93 displays the **Storage Pools**. From here you can create or delete Storage pools.

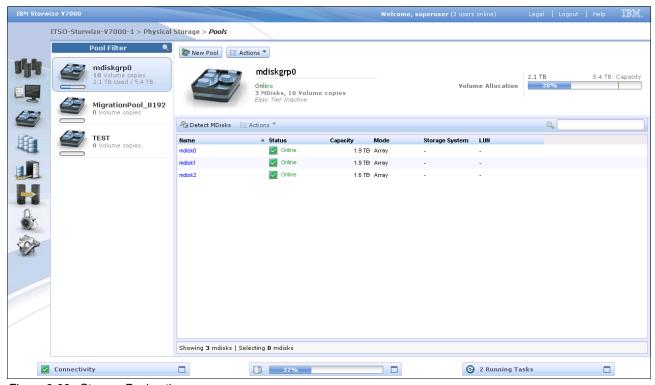


Figure 3-33 Storage Pool option

MDisks

The screen shown in Figure 3-34 on page 94 displays the **MDisks** that are available to the IBM Storwize V7000 system. The MDisks displayed will show whether they are *managed* in which case the storage pool will be displayed or whether they are *unmanaged* in which case they can be added to a new pool.

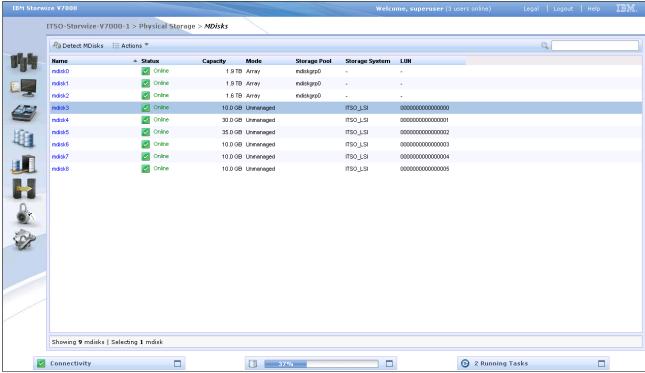


Figure 3-34 MDisk option

By clicking on the display bar as indicated in Figure 3-35 on page 95 you can choose to change the fields that are displayed. Just select the items you wish to be displayed.

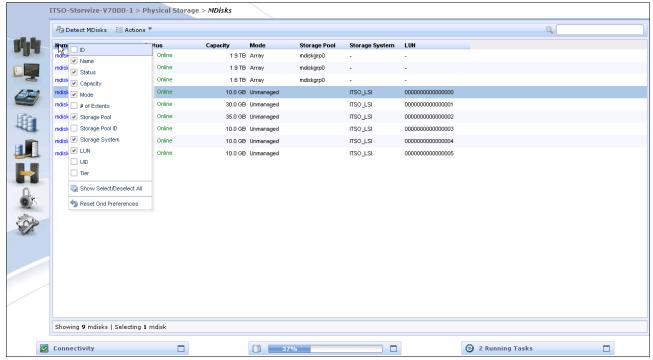


Figure 3-35 Displaying additional fields

From this screen you can choose to either use the options from the **Action** button or you could choose to highlight the particular MDisk(s) that you require and click on the right mouse button as shown in Figure 3-36 on page 95.

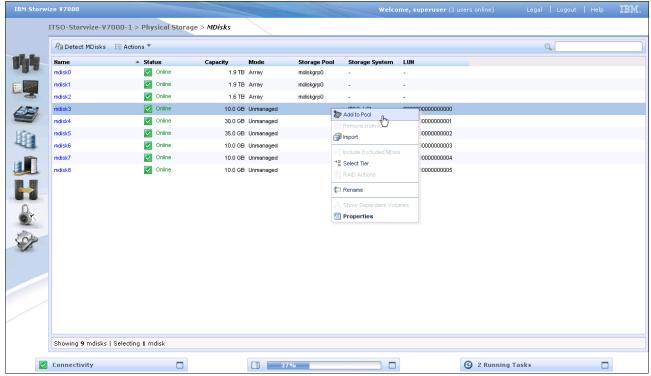


Figure 3-36 Commands for a single MDisk from the MDisks option

3.5 Volumes menu

The following section shows the **Volumes** menu and its options. By hovering the cursor over the **Volumes** icon the popup menu will appear as shown in Figure 3-37 on page 96.

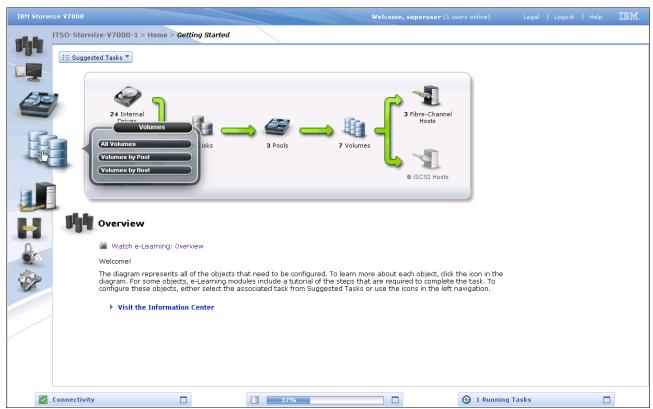


Figure 3-37 Navigating to the Volumes menu

All Volumes

Clicking on the **All Volumes** option will display the screen as shown in Figure 3-38 on page 97. From here you can perform tasks on the volumes such as shrink, enlarge, map to a host or migrate a volume.

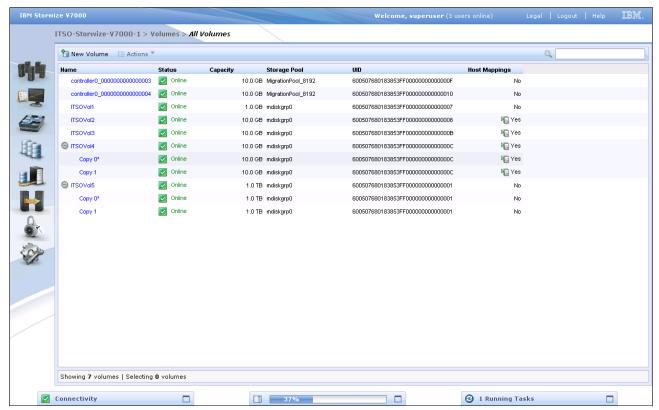


Figure 3-38 All Volumes option

From this menu you can perform a variety of operations on the volumes. You can use the **Action** button or you can right click on the Volume name which will display a list of operations that can be performed against the volume as shown in Figure 3-39 on page 98.

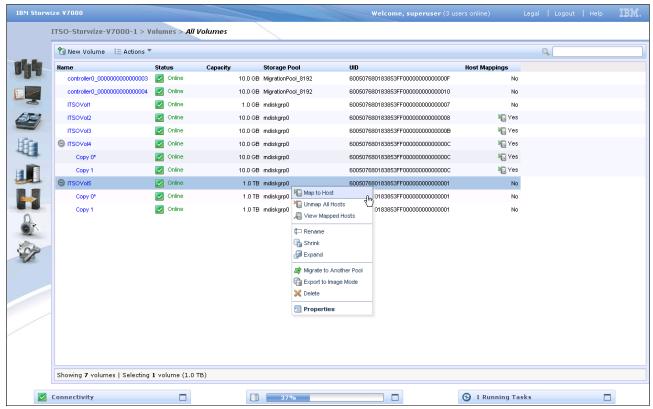


Figure 3-39 Volume operations

Volumes by Pool

Selecting the **Volumes by Pool** option will display the screen shown in Figure 3-40 on page 99.

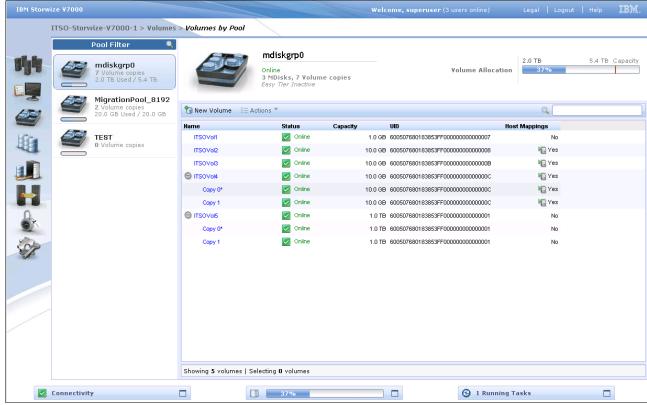


Figure 3-40 Volumes by Pool option

Similar to the previous screen you can either use the Action button or you can click on the Pool to display a list of valid commands as shown in Figure 3-41 on page 100.

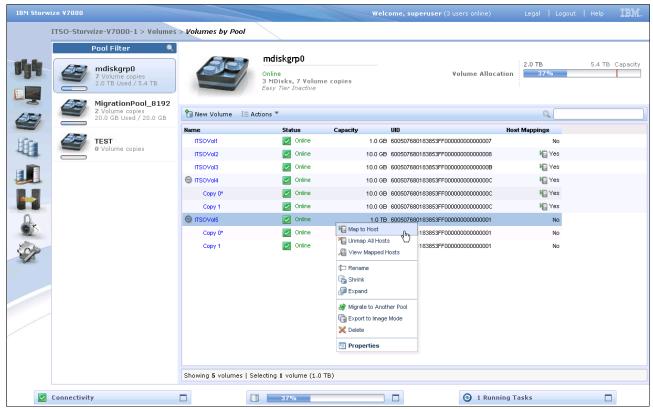


Figure 3-41 Commands for a single volume from the Volume by Pool option

Volumes by Host

Selecting the **Volumes by Host** option will display the screen shown in Figure 3-42 on page 101. This will display the volumes that have been mapped to a given host.

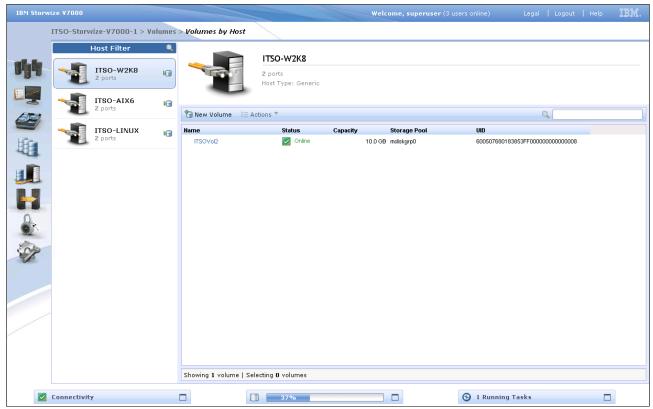


Figure 3-42 Volumes by Host option

Use the Action button or you can click on the Pool to display a list of valid commands as shown in Figure 3-43 on page 102.

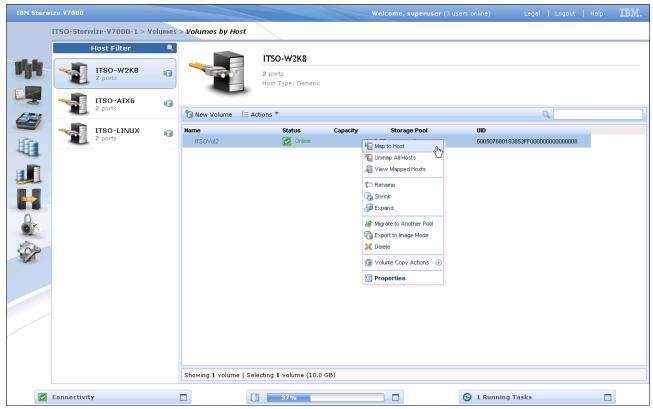


Figure 3-43 Commands for a single volume from the Volume by Host option

3.6 Hosts menu

The following section shows the **Hosts** menu and its options. By hovering the cursor over the **Host** icon the popup menu will appear as shown in Figure 3-44 on page 103.

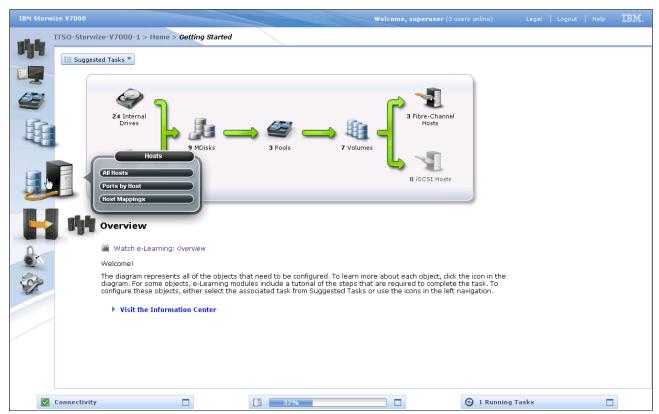


Figure 3-44 Navigating to the Hosts menu

All Hosts

Selecting the **All Hosts** option will display the screen shown in Figure 3-45 on page 104. From here you can modify host mappings, unmap hosts, rename hosts and create new hosts.

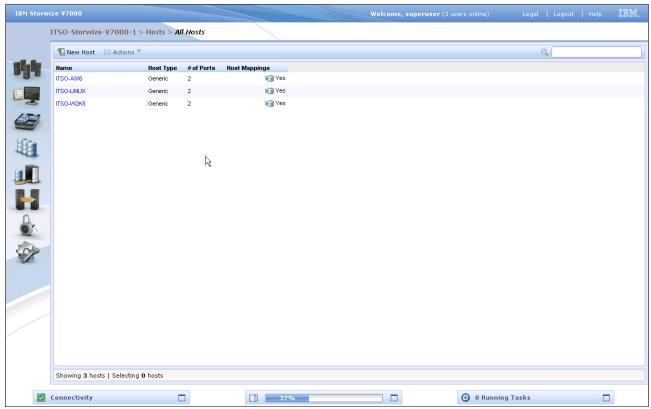


Figure 3-45 All Hosts option

As with a number of other screens you can use the command buttons or you can select a host and right click on it to access the commands as shown in Figure 3-46 on page 105.

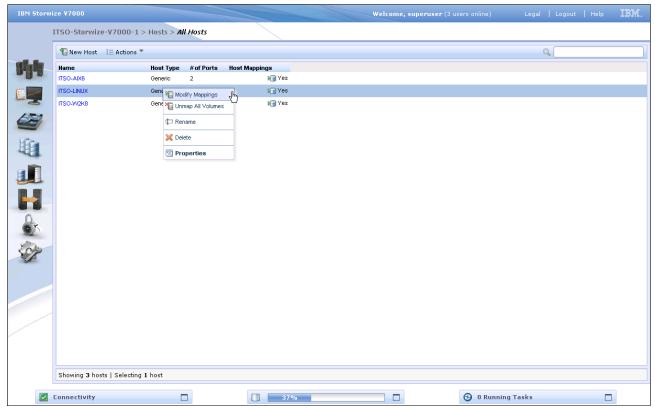


Figure 3-46 Commands for a single Host from the All Hosts option

Ports by Host

Selecting the **Ports by Hosts** option will display the screen shown in Figure 3-47 on page 106. This screen will display the Fibre Channel and iSCSI ports that are assigned to a particular host.

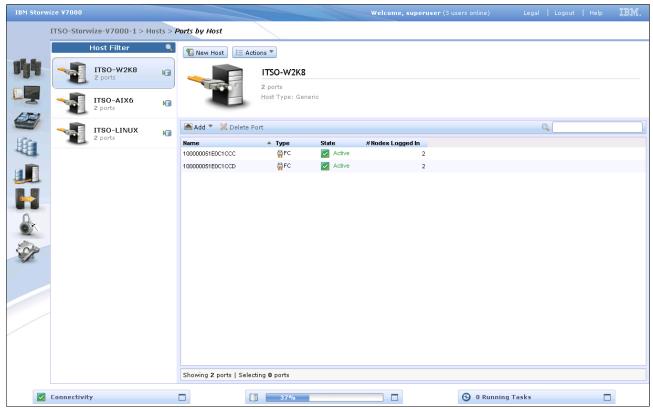


Figure 3-47 Ports by Host option

From this screen by clicking on the **Actions** button you can modify the mappings, unmap volumes, rename hosts and delete ports as shown in Figure 3-48 on page 107.

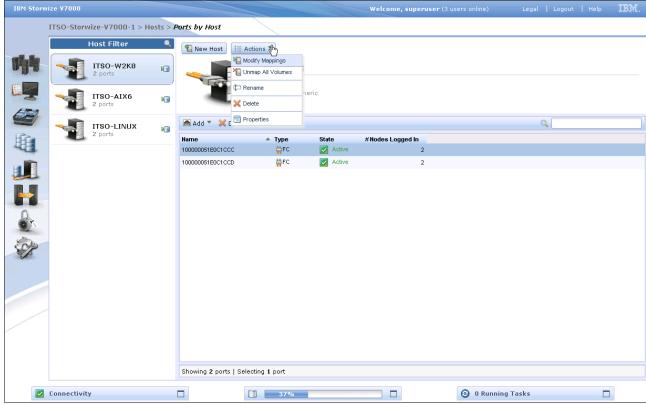


Figure 3-48 Commands for a single Host from the Ports by Host option

Host Mappings

Clicking on the **Host Mappings** option will display the screen shown in Figure 3-49 on page 108. This screen displays the host id, SCSI identifier and the Volume identifier for all the mapped volumes.

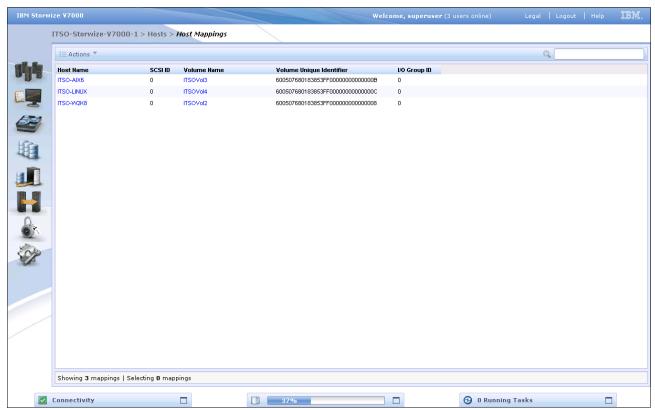


Figure 3-49 Host Mapping option

As with a number of other screens you can use the command buttons as shown in Figure 3-50 on page 109 or you can select a host and right click on it to access the commands.

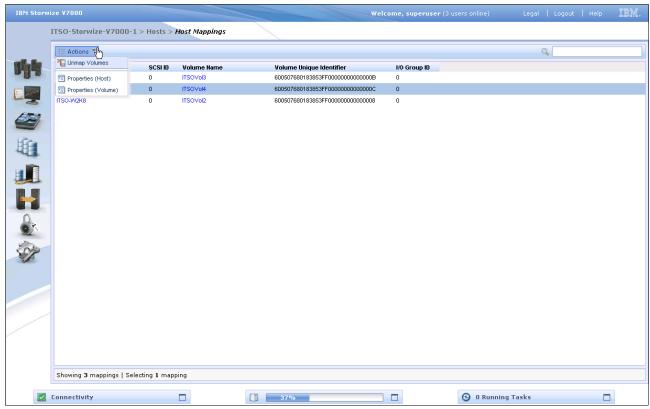


Figure 3-50 Commands for a single Host from the Host Mapping option

3.7 Copy Services menu

The following section shows the **Copy Services** menu and its options. By hovering the cursor over the Copy Services icon the popup menu will appear as shown in Figure 3-51 on page 110.

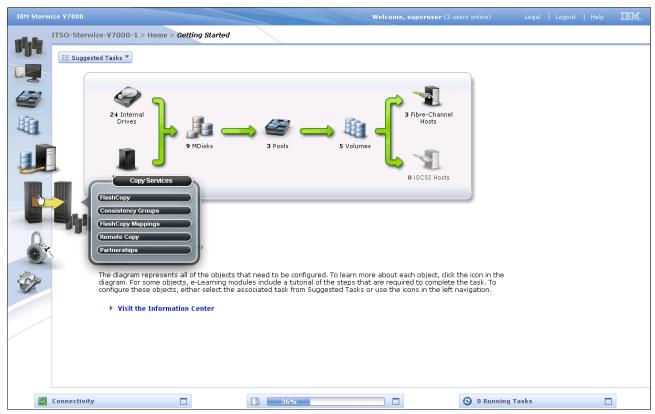


Figure 3-51 Navigating to the Copy Services menu

FlashCopy

Clicking on the **FlashCopy** option will display the screen shown in Figure 3-52 on page 111. This screen shows the volumes that are available and by right clicking on a volume a list of operations is displayed. From here we can perform tasks such as initiate a new snapshot, clone or backup just by clicking on the volume name.

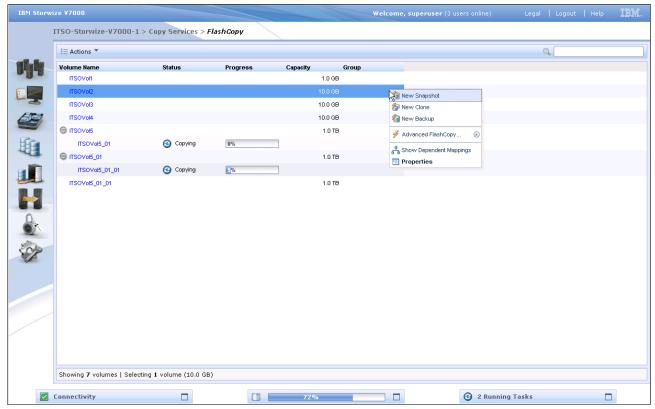


Figure 3-52 FlashCopy option

Clicking on the volume name will display the screen shown in Figure 3-53 on page 112. From here you can click on the tabs at the top of the screen to display additional information such as the hosts that the volume or FlashCopy volume is mapped to and its dependant MDisks.

By clicking on the **Action** button you can also initiate a mirrored copy of the volume or migrate it to a different storage pool.

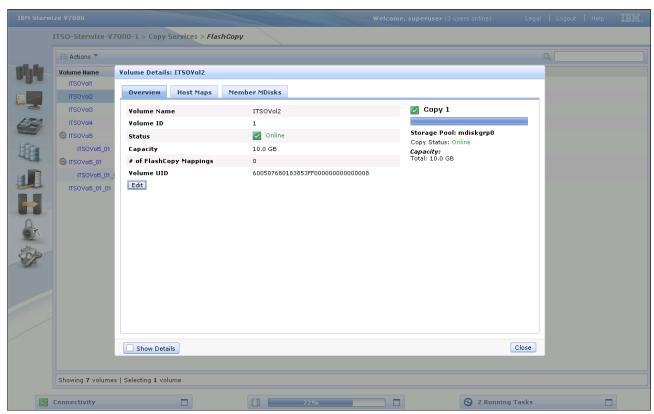


Figure 3-53 FlashCopy volume details

FlashCopy Consistency Group

Selecting the **FlashCopy Consistency Group** option will display the screen shown in Figure 3-54 on page 113. This screens enables FlashCopy relationships to be placed into a Consistency Group. You can also use start and stop commands to the FlashCopy Consistency Group from this screen by right clicking on the relationship.

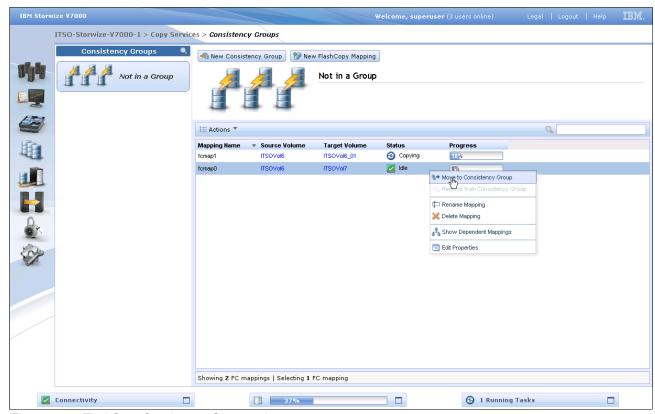


Figure 3-54 FlashCopy Consistency Group option

FlashCopy Mappings

Clicking on the **FlashCopy Mapping** option will display the screen shown in Figure 3-55 on page 114. From this screen we can start, stop, delete and rename the FlashCopy mappings. There is also an option to move the relationship into a Consistency Group.

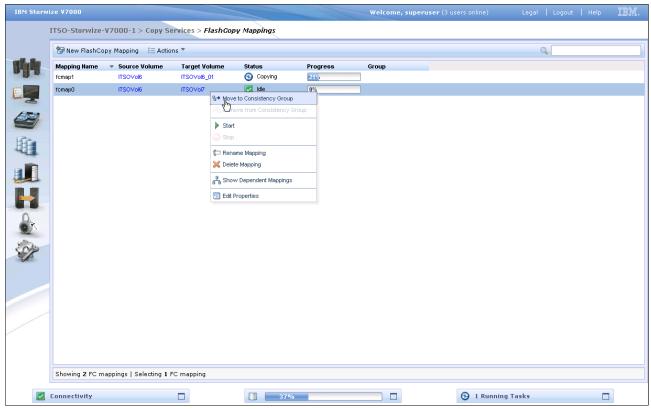


Figure 3-55 FlashCopy Mapping options

Remote Copy

Selecting the **Remote Copy** option will display the screen shown in Figure 3-56 on page 115. This screen displays the existing Remote Copy relationships and allows us to setup and modify consistency groups. From this screen we can also start and stop relationships, add relationships to a consistency group and switch the direction of the mirror.

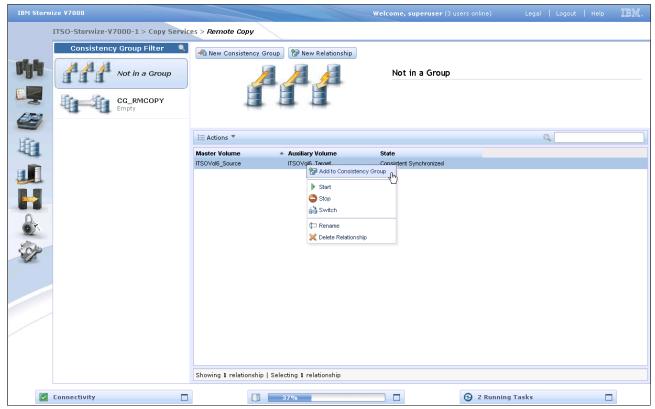


Figure 3-56 Remote Copy option

Partnerships

Selecting the **Partnerships** option will display the screen shown in Figure 3-57 on page 116. This allows us to setup a new partnership or delete an existing partnership with another IBM Storwize V7000 system for the purposes of remote mirroring.

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

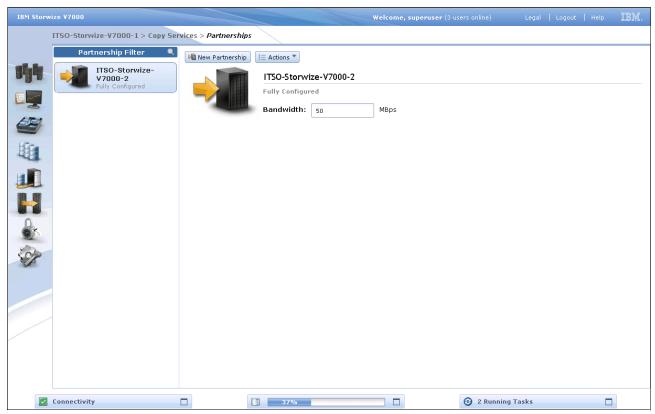


Figure 3-57 Partnerships option

3.8 User Management menu

The following section shows the **User Management** menu and its options. By hovering the cursor over the User Management icon the popup menu will appear as shown in Figure 3-58 on page 117.

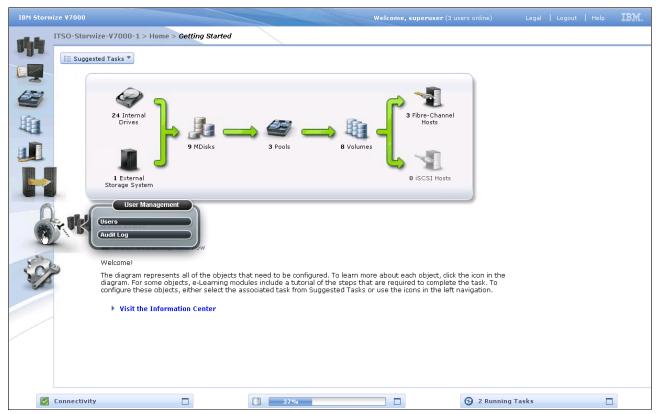


Figure 3-58 User Management menu

Users

Figure 3-59 on page 118 shows the **Users** option screen. This screen enables you to create and delete new users, change and remove passwords, add and remove SSH keys.

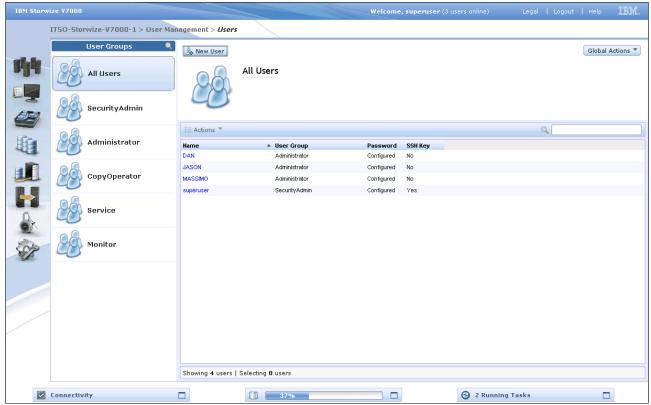


Figure 3-59 All Users option

Clicking on the **New User** button will display a window as shown in Figure 3-60 on page 119. From here you can input the name of the user, the password and load the ssh key.

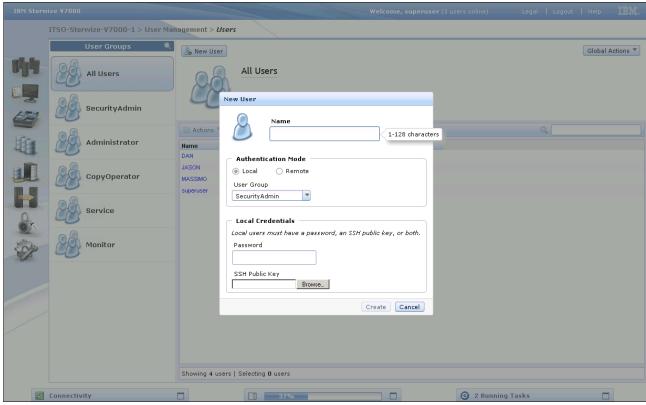


Figure 3-60 New user option

Audit Log

Selecting the **Audit Log** option will display the screen shown in Figure 3-61 on page 120. The cluster maintains an audit log of successfully executed commands, indicating which users performed particular actions at certain times.

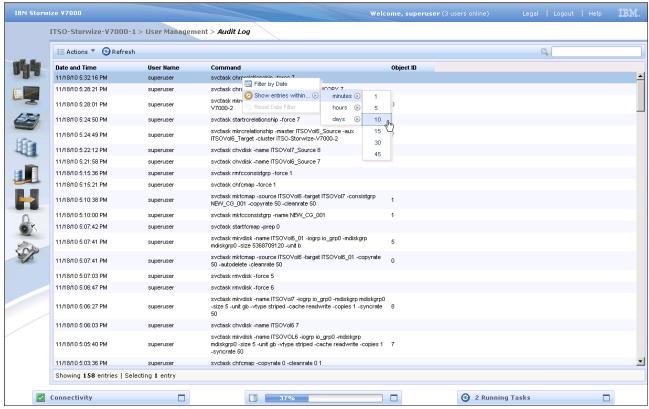


Figure 3-61 Audit Log option

3.9 Configuration menu

The following section shows the **Configuration** menu and its options. By hovering the cursor over the Configuration icon the popup menu will appear as shown in Figure 3-62 on page 121.

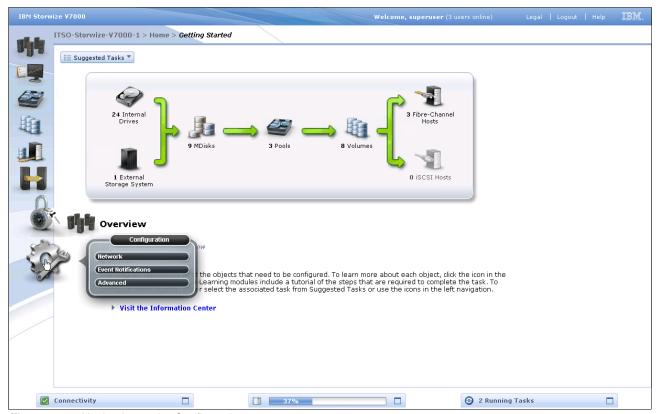


Figure 3-62 Navigating to the Configuration menu

Network

Clicking on the **Network** option will display the screen shown in Figure 3-63 on page 122. From here we can update the network configuration, setup iSCSI definitions and display information on the Fibre Channel connections.

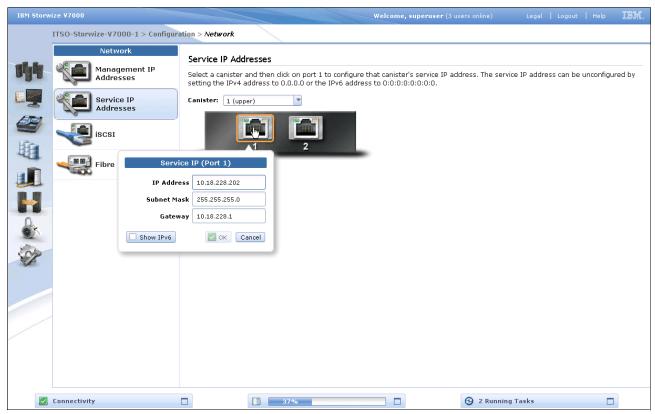


Figure 3-63 Network option

By selecting the **Fibre Channel** option as shown in Figure 3-64 on page 123 some very useful information is presented. In this example we have selected the **Hosts** option from the drop down box and then selected to display the details for one specific host - Hurricane - from the list of host systems. Other options available from the drop down box include displaying Fibre Channel details: for all devices, for clusters, for nodes, for storage systems or for hosts.

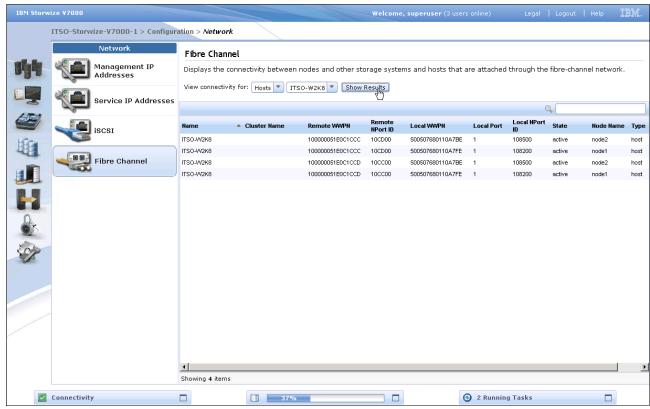


Figure 3-64 Fibre Channel details

Event notification

Shown in Figure 3-65 on page 124 is the screen displayed when selecting the **Event notification** option. From this screen you can configure the e-mail alerts, (included the Call Home function), SNMP monitoring and define syslog servers and the message types.

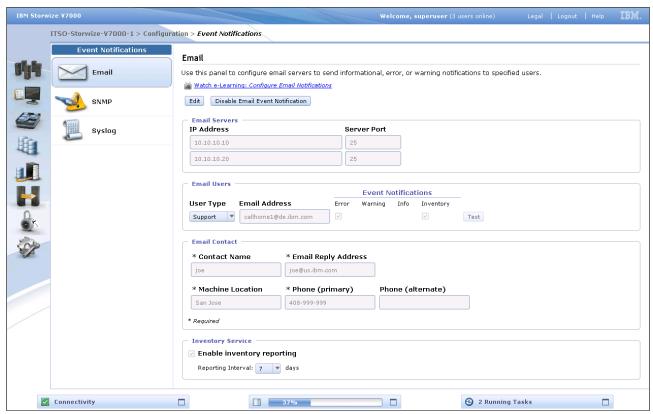


Figure 3-65 Event notification option

Advanced

By selecting the **Advanced** option as shown in Figure 3-66 on page 125 the following screen is displayed. This screen provides options to set the date and time, update the software licensing levels, upgrade the firmware and set GUI preferences.

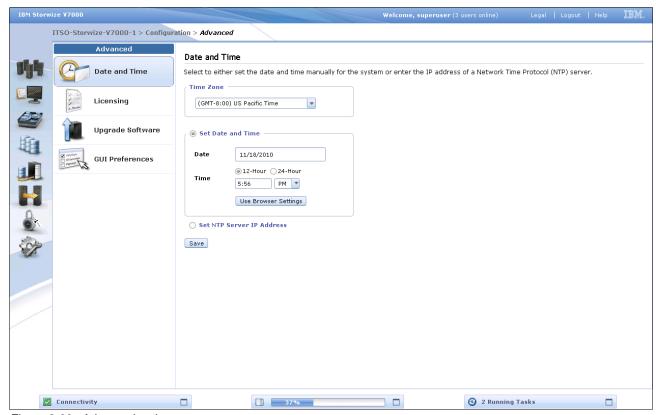


Figure 3-66 Advanced option



4

Host Configuration

This chapter describes how to use the IBM Storwize V7000 GUI to create hosts, and how to prepare a host in order to access the volumes that will be created later in the volumes chapter (Chapter 5, "Basic Volume Configuration" on page 157).

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

This chapter contains the following:

- Preparing the host operating system
 - Windows
 - Fibre Channel
 - 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, and each adapter is on a separate network (or at minimum in a separate zone), and is connected to both canisters. This assures four paths for failover and failback purposes.

Prior to mapping the newly created volumes on the host of your choice, a little preparation will go 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. First use the System Storage Interoperation Center (SSIC), to check which code levels are supported to attach your host to your storage. SSIC is a web tool to check the interoperation of host, storage, switches, and multipathing drivers.

http://www-03.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 V6.1 document, which is available at:

http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003703

This chapter will focus on Windows and VMware. If you do have to attach any other hosts, for example AIX®, Linux®, or even Apple, then you will find the required information in the IBM Storwize V7000 Information Center available at:

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

4.1 Preparing the Host Operating System

This part describes how to prepare Windows and VMware hosts for attachment to an IBM Storwize V7000, using either Fibre Channel or iSCSI to connect.

4.1.1 Windows 2008 - Preparing for Fibre Channel Attachment

Perform the following steps to prepare a Windows 2008 host to connect to an IBM Storwize V7000 using Fibre Channel.

- Make sure that the latest OS service pack and hotfixes are applied to your Microsoft® server.
- ► Use the latest firmware and driver levels on your host system.
- Install HBA or HBAs on the Windows server using the latest bios and driver.
- Connect the FC Host Adapter ports to the switches.
- Configure the switches (zoning).
- Configure the HBA for hosts running Windows.
- Set windows time-out value.
- Install the multipath module.

Download and install the supported drivers and firmware

Install a supported HBA driver for your configuration. Use windows device manager or vendor tools like 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 available at:

http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003703#_Win2008

Configuring Brocade HBAs for Windows

This section applies to Windows hosts that have Brocade HBAs installed and after installing the device driver and firmware you must configure the HBAs. To do that, 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 do that, 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 point to point only
- LUNs Per Target: 0
- ► Port Down Retry Count: 15

Configuring Emulex HBAs for Windows

This section applies to Windows hosts that have Emulex HBAs installed, after installing the device driver and firmware you must configure the HBAs. To do that, either use the Emulex HBAnyware Software or reboot into the HBA Bios load the defaults, and set topology to 1 (10F_Port Fabric).

Set Windows time-out value

For Windows hosts the disk I/O time-out value should be set to 60 seconds. To verify this perform the following steps:

- 1. Click Start, and select Run.
- 2. In the dialog box, type regedit and press Enter.
- In the registry browsing tool, locate the HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Disk\TimeOutValue key.
- 4. Confirm that the value for the key is 60 (decimal value), and, if necessary, change the value to 60, as shown in Figure 4-1

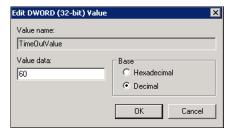


Figure 4-1 Windows Time-out Value

Installing the multipath module

Microsoft Multi Path Input Output (MPIO) solutions are designed to work in conjunction 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 specifically designed to support IBM storage devices on Windows hosts. The intention of MPIO is to get a better integration of 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 proper multipathing with IBM Storwize V7000, SDDDSM has to be installed on Windows hosts.

1. Check the SDDDSM download matrix to determine the correct level of SDDDSM to install for Windows 2008 and download the package at:

http://www-01.ibm.com/support/docview.wss?rs=540&uid=ssg1S7001350#WindowsSDDDSM

2. Extract the package and start the setup application as shown in Figure 4-2 on page 130.

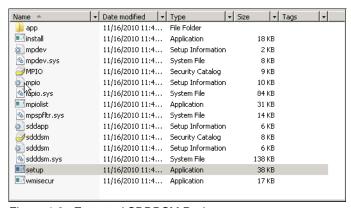


Figure 4-2 Extracted SDDDSM Package

3. If security warnings are enabled on your host, you will be prompted to click **Run** as shown in Figure 4-3



Figure 4-3 Confirm Security Warning

4. The setup CLI appears, type **yes** to install the SDDDSM and press **Enter** as shown in Figure 5.

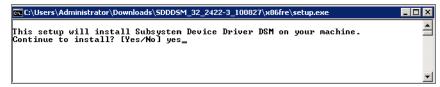


Figure 4-4 Confirm SDDDSM Installation

5. After the setup completes you will be asked to restart the system. Confirm this by typing **yes** and press **Enter** as shown in Figure 4-5 on page 131.

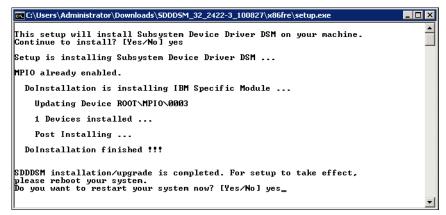


Figure 4-5 Answer "yes" to restart the host

After the reboot you have now successfully installed the IBM SDDDSM. You can check the installed driver version if you click Start \rightarrow All Programs \rightarrow Subsystem Device Driver DSM \rightarrow Subsystem Device Driver DSM. A command prompt will open and the command **datapath query version** may be used to determine the version currently installed as in Example 4-1 for this Windows 2008 host.

Example 4-1 datapath query version

```
C:\Program Files\IBM\SDDDSM>datapath query version
IBM SDDDSM Version 2.4.2.2-3
Microsoft MPIO Version 6.0.6002.18005
```

This tool can also be used to determine the wwpns of the host. Type datapath query wwpn as shown in Example 4-2 and note down the wwpns of your host, as you will need them later.

Example 4-2 Datapath query wwpn

```
C:\Program Files\IBM\SDDDSM>datapath query wwpn
Adapter Name PortWWN
Scsi Port3: 100000051EAF08A0
Scsi Port4: 100000051EAF08A1
Scsi Port5: 100000051EOC1CCC
Scsi Port6: 100000051EOC1CCD

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

If you need more detailed information about SDDDSM you will find it in the *Multipath Subsystem Device Driver User's Guide*, GC52-1309-02.

http://www-01.ibm.com/support/docview.wss?rs=540&context=ST52G7&uid=ssg1S7000303

Now the Windows host has been prepared to connect to the IBM Storwize V7000 and we know the wwpns of the host. The next step is to configure a host object for the wwpn's identified using the IBM Storwize V7000 GUI. This is explained in "Creating FC Hosts" on page 147.

Note that SAN boot hosts are beyond the intended scope of this paper and for more information follow the steps in the Information Center available from the IBM Support Portal.

Author Comment: If we have time we should cover SAN Boot

Note: This book focuses on Windows 2008, however the procedure for Windows 2003 is very similar. If you use Windows 2003 do not forget to install Microsoft Hotfix 908980. If you do not install it before operation, preferred pathing is not available:

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

4.1.2 Windows 2008 - preparing for iSCSI Attachment

In Windows 2008 the Microsoft iSCSI software initiator is pre-installed. Select **ISCSI Initiator** in the Administrative Tools menu as shown in Figure 4-6 on page 133.

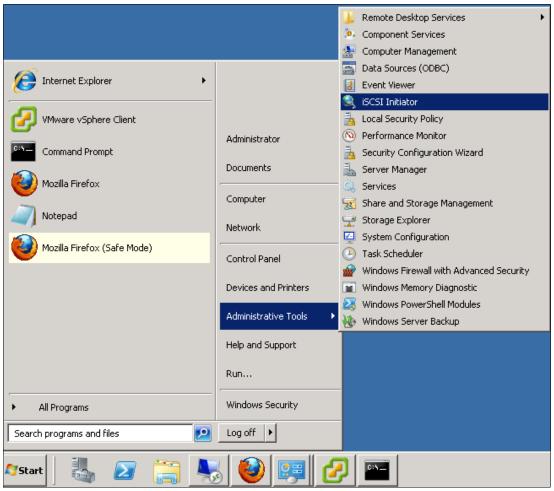


Figure 4-6 Windows iSCSI Initiator

Confirm the automatic startup of the iSCSI Service as shown in Figure 4-7.

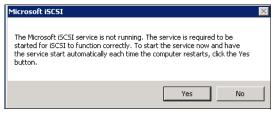


Figure 4-7 Auto start iSCSI Service

The iSCSI Configuration window will appear, select the configuration tab as shown in Figure 4-8 on page 134 and remember the initiator name of your Windows host, you will need it later.

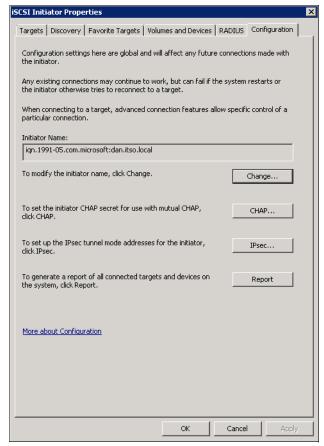


Figure 4-8 iSCSI Initiator Properties

You can change the initiator name, or enable advanced authentication, but this is out of the scope of our basic setup. More detailed information is available in the IBM Storwize V7000 Information Center:

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

Set the Windows registry keys

It is recommended to make the following changes to the system registry, this will make your iSCSI operations more reliable.

- 1. Click **Start**, and select **Run**.
- 2. In the dialog box, type regedit and press Enter.
- In the registry browsing tool, locate the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\
bus id>\Parameters\LinkDownTime key.
- 4. Confirm that the value for the LinkDownTime key is 120 (decimal value), and, if necessary, change the value to 120 as shown in Figure 4-9 on page 135.

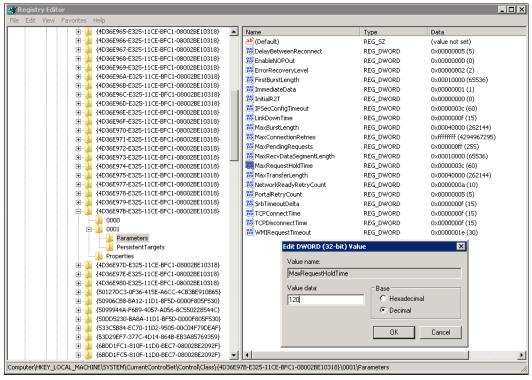


Figure 4-9 Edit windows registry for iSCSI hosts

- In the registry browsing tool, locate the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\
bus id>\Parameters\MaxRequestHoldTime key.
- Confirm that the value for the MaxRequestHoldTime key is 120 (decimal value), and, if necessary, change the value to 120.
- 7. In the registry browsing tool, locate the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\
bus id>\Parameters\ MaxPendingRequests key.
- 8. Confirm that the value for the MaxPendingRequests key is 2048 (decimal value), and, if necessary, change the value to 2048.
- In the registry browsing tool, locate the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Disk\TimeOutValue key.
- 10. Confirm that the value for the TimeOutValue key is 60 (decimal value), and, if necessary, change the value to 60.
- 11. Reboot your host for these changes to take effect.

This are the basic steps to prepare a Windows 2008 iSCSI host, next go to "Creating iSCSI Hosts" on page 150 to configure the IBM Storwize V7000 for iSCSI connections:

4.1.3 VMware ESX - Preparing for Fibre Channel Attachment

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

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

Download and install 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 V6.1" list for VMware available at:

http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003703#_VMware

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 if there are any requirements for your configuration:

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 in the host system and make sure that the Adapter Bios is disabled, unless you are using boot from SAN.

Configuring Qlogic HBAs for VMware ESX

This section applies to ESX hosts that have Qlogic HBAs installed and after installing the firmware you must configure the HBAs. To do that, either use the Qlogic Sansurfer Software or reboot into 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.: 8Link Down Timeout: 10Command Timeout: 20

Extended event logging: Disabled (only enable it for debugging)

RIO Operation Mode: 0Interrupt 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 in the host system and make sure that the Adapter Bios is disabled, unless you are using boot from SAN.

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 you will find a detailed installation guide at:

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, select **Storage Adapters** and scroll down to your FC HBAs as shown in Figure 4-10 on page 138. Remember the wwpns of the installed adapters for later use.

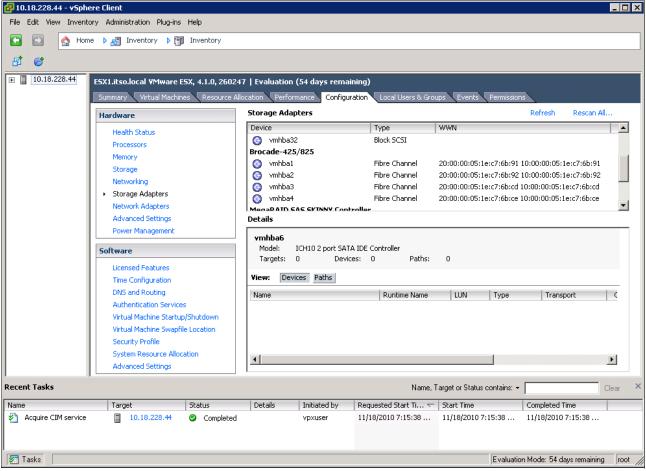


Figure 4-10 Show wwpns in VMware ESX

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 higher the recommended multipathing policy is **Round Robin**. This will perform 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 as well. This policy setting can be selected for every volume. But this will be done afterwards once we have attached IBM Storwize V7000 LUNs to the ESX host (5.3.3, "VMware ESX Fibre Channel Attachment" on page 183). If you use an older version of VMware ESX up to 3.5 **Fixed** is the recommended policy setting.

After all these steps are completed the ESX host is prepared to connect to the IBM Storwize V7000, go to 4.2.1, "Creating FC Hosts" on page 147 to create the ESX FC host in the IBM Storwize V7000 GUI.

4.1.4 VMware ESX - Preparing for iSCSI Attachment

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

http://www.vmware.com/pdf/vsphere4/r40/vsp 40 iscsi san cfg.pdf

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

- ► Make sure that the latest firmware levels are applied on your host system.
- ► Install VMware ESX and load additional drivers if required
- ► Connect the ESX server to your network, it is recommended to use separate network interfaces for iSCSI traffic.
- ► 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.

1. Connect to your ESX server using the vSphere Client and navigate to **Configuration** and select **Networking** as shown in Figure 4-11.

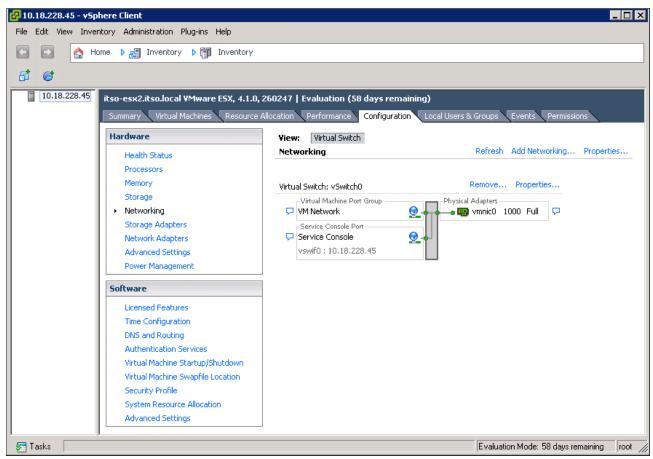


Figure 4-11 Select VMware Networking

2. Click **Add Networking** to start the Add Network Wizard as shown in Figure 4-12 on page 140. Select VMkernel and click **Next**.

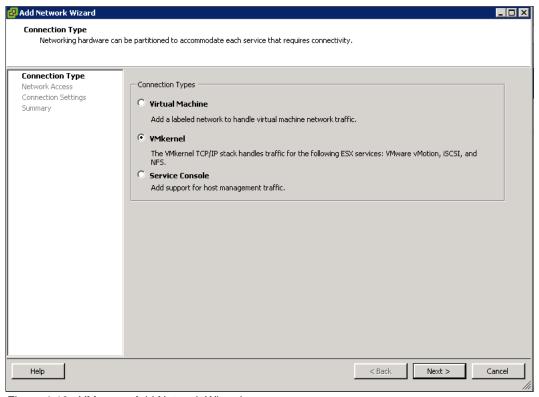


Figure 4-12 VMware - Add Network Wizard

3. Select one or more network interface(s) you want to use for iSCSI traffic and click next as shown in Figure 4-13 on page 141.

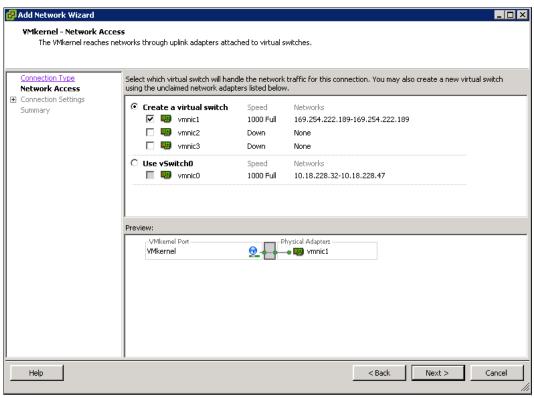


Figure 4-13 VMware - Select iSCSI interface

4. Enter a meaningful Network Label and click Next as shown in Figure 4-14.

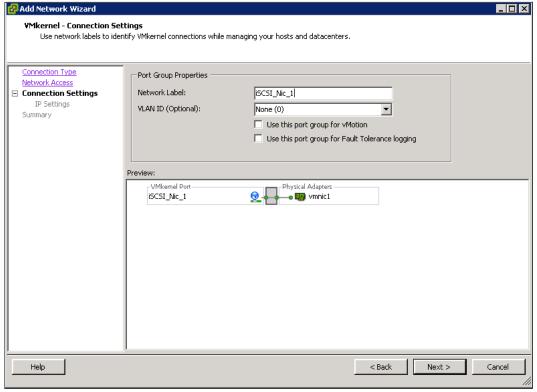


Figure 4-14 VMware - Enter Network Label

5. Enter an IP address for your iSCSI network, it is strongly recommended to use a dedicated network for iSCSI traffic, as shown in Figure 4-15.

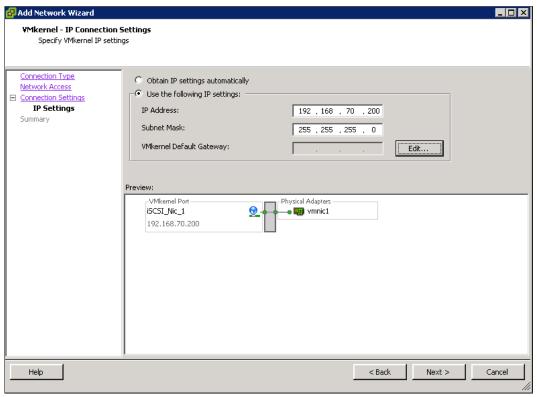


Figure 4-15 VMware - Enter iSCSI Network IP

- 6. Click **Finish** to complete the setup.
- 7. Select **Storage Adapters** and scroll down to iSCSI Software Adapter as shown in Figure 4-16 on page 143. Highlight it and click **Properties**.

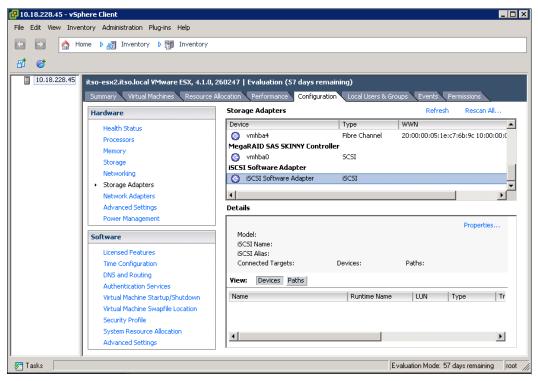


Figure 4-16 VMware - iSCSI Software Adapter

8. The iSCSI Software Adapter Properties window appears. As you can see in Figure 4-17 the initiator is disabled by default, to change this click **Configure.**

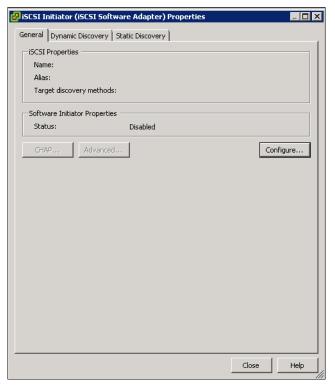


Figure 4-17 VMware - iSCSI Software Adapter Properties

9. Mark the **Enabled** checkbox and click **OK** as shown in Figure 4-18.

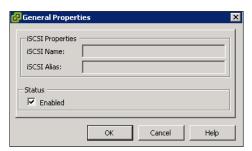


Figure 4-18 VMware - Mark Enabled Checkbox

10. The VMware ESX iSCSI initiator is now successfully enabled as shown in Figure 4-19. Remember your initiator name for later use.

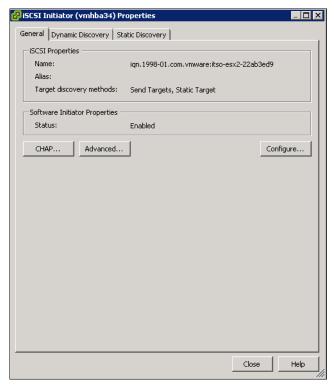


Figure 4-19 VMware iSCSI Initiator enabled

Your VMware ESX host is now prepared to connect to the IBM Storwize V7000. Go to "Creating iSCSI Hosts" on page 150 to create the an ESX iSCSI host in the IBM Storwize V7000 GUI.

4.2 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 "Preparing the Host Operating System" on page 128 and you are aware of the host wwpns or iSCSI initiator names.

Open the host configuration as shown in Figure 4-20 by selecting All Hosts.

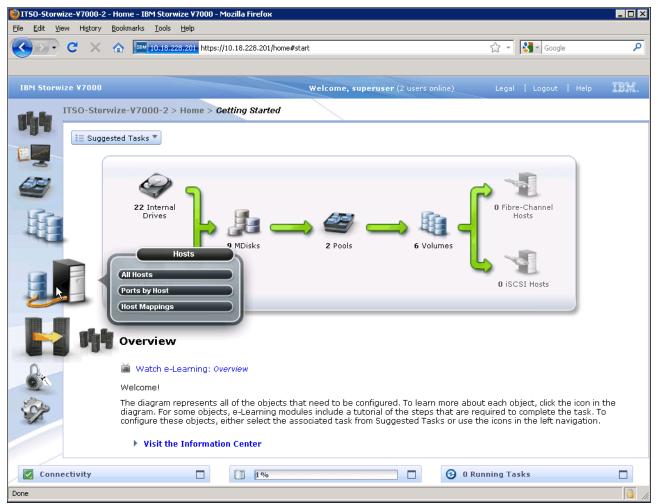


Figure 4-20 Opening the Host Section

The **All Host** Section will appear as shown in Figure 4-21 on page 146.

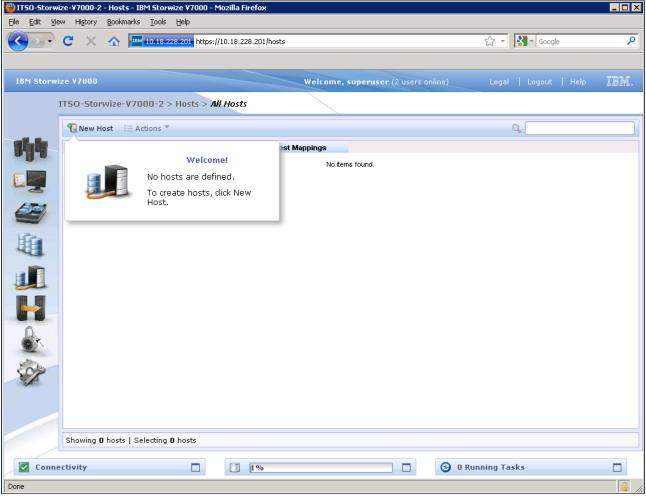


Figure 4-21 All Hosts

To create a new host click **New Host** to start the wizard as shown in Figure 4-22.



Figure 4-22 Create Host

If you want to create a Fibre Channel host continue with 4.2.1, "Creating FC Hosts", for iSCSI Hosts go to "Creating iSCSI Hosts" on page 150.

4.2.1 Creating FC Hosts

1. Click **Fibre Channel Host** as shown in Figure 4-22 on page 146, and the Fibre Channel configuration wizard will appear as shown in Figure 4-23.



Figure 4-23 Create FC Host

2. Enter a Host Name and click the Fibre Channel Ports box to get a list of all known wwpns as shown in Figure 4-24.

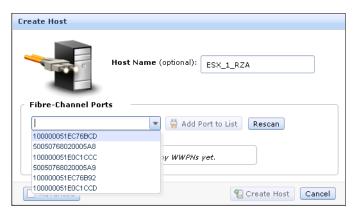


Figure 4-24 Available wwpns

The IBM Storwize V7000 will have the host port wwpn's available if you prepared the host(s) and know your wwpns as described in "Preparing the Host Operating System" on page 128. 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** as shown in Figure 4-25 on page 148.

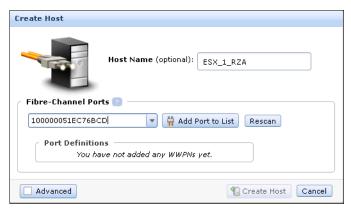


Figure 4-25 Add Port to List

4. Add all ports that belong to the host as shown in Figure 4-26.

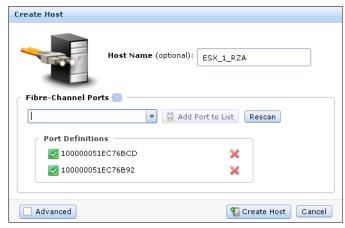


Figure 4-26 Add all wwpns

Note: If you want to create hosts that are offline, or not connected at the moment, it is also possible to enter the wwpns manually. Just type them into the Fibre Channel Ports Box and add them to the list as well.

5. If you are creating an HP/UX or TPGS host, check the **Advanced** checkbox and more options as shown in Figure 4-27 on page 149 will appear. Select your host type.

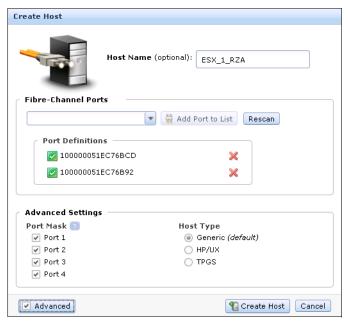


Figure 4-27 Create Host - Advanced Settings

6. Click Create Host and the wizard creates the host as shown in Figure 4-28.

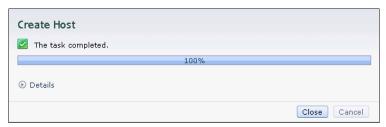


Figure 4-28 Create Host Completes

7. Click Close to return to the host section as shown in Figure 4-29.

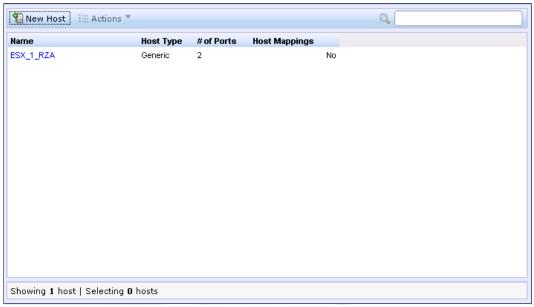


Figure 4-29 All Hosts

8. Repeat these steps for all of your Fibre Channel hosts. Figure 4-30 shows the **All Hosts** section after creating a second host.

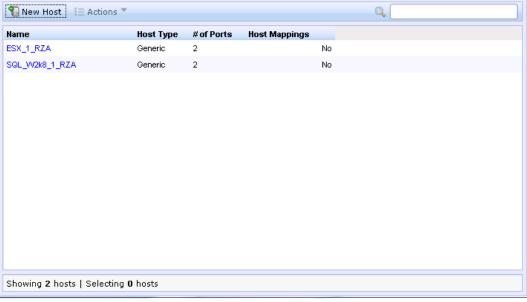


Figure 4-30 All Hosts - after creation of second host

Once you are completed creating Fibre Channel hosts go to Chapter 5, "Basic Volume Configuration" on page 157 to create Volumes and map them to the created hosts.

4.2.2 Creating iSCSI Hosts

1. Click **iSCSI Host** as shown in Figure 4-22 on page 146, and the iSCSI configuration wizard will appear as shown in Figure 4-31 on page 151.

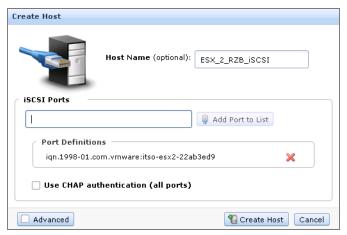


Figure 4-31 Create iSCSI Host

2. Enter a host name, type the iSCSI initiator name into the iSCSI Ports box, and click **Add Ports to List** as shown in Figure 4-32. If you want to add several initiator names to one host repeat this step.

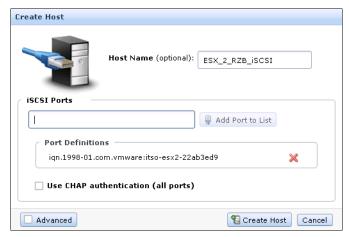


Figure 4-32 Create iSCSI Host - Enter name and iSCSI Ports

3. If you are connecting an HP/UX or TPGS host mark the **Advanced** checkbox as shown in Figure 4-33 on page 152 and select the correct host type.

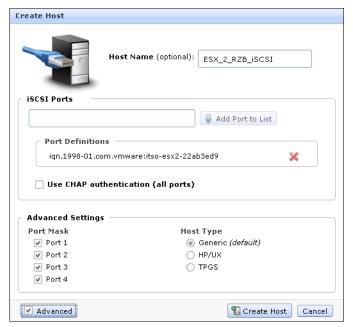


Figure 4-33 Create iSCSI Host - Advanced Settings

4. Click Create Host and the wizard will complete as shown in Figure 4-34. Click Close.

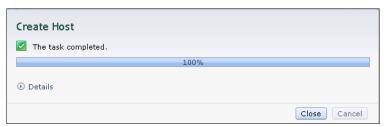


Figure 4-34 Create iSCSI Host - Complete

5. Repeat these steps for every iSCSI host you want to create, Figure 4-35 shows the All Hosts section after creating two Fibre Channel and two iSCSI hosts.

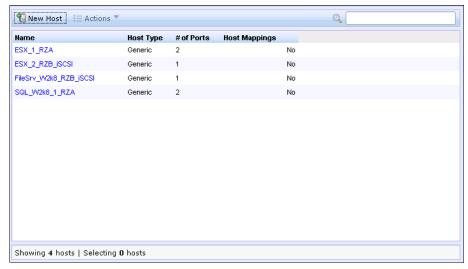


Figure 4-35 All Hosts

Now the iSCSI hosts are configured on the IBM Storwize V7000. But to provide connectivity, the iSCSI ethernet ports also have to be configured. Perform the following steps to enable iSCSI connectivity:

1. Switch to the configuration section and select **Network** as shown in Figure 4-36.

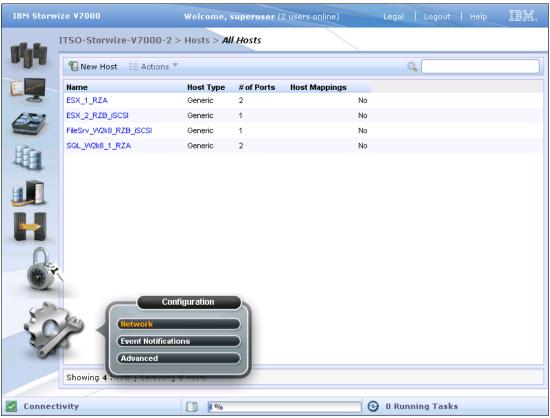


Figure 4-36 Configuration -Network

2. Select **ISCSI** and the iSCSI Configuration section will appear as shown in Figure 4-37 on page 154.

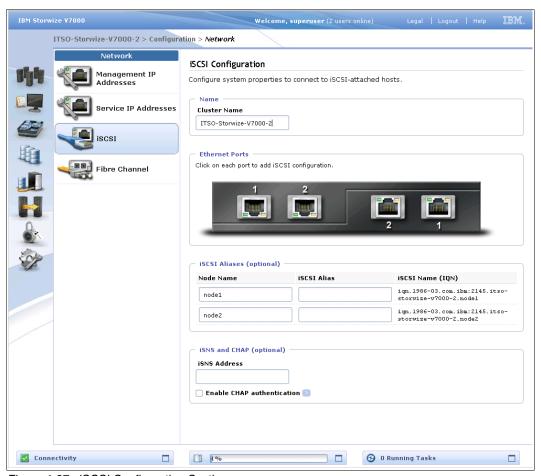


Figure 4-37 iSCSI Configuration Section

In the configuration you do have an overview about all iSCSI settings for the IBM Storwize V7000. You can add iSCSI Alias, iSNS Addresses and Chap authentication configuration in this section, as well as the iSCSI IP Address which we will also edit in the basic setup now.

3. Click on the Ethernet Ports to enter the iSCSI IP Address as shown in Figure 4-38 on page 155. Repeat this step for each port you want to use for iSCSI traffic.



Figure 4-38 Enter iSCSI IP Address

4. After you have entered the IP Address for each port click **Apply Changes** to enable the configuration as shown in Figure 4-39.

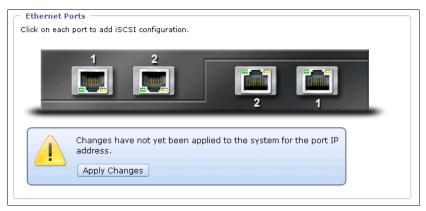


Figure 4-39 Apply ISCSI IP Changes

5. After the changes are successfully applied click **Close** as shown in Figure 4-40.

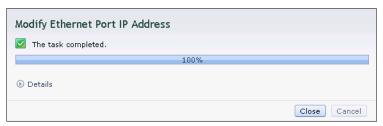


Figure 4-40 iSCSI IP successfully modified

The IBM Storwize V7000 is now configured and ready for iSCSI use. Remember the initiator names of your storage canisters as shown in Figure 4-37 on page 154 as you will need them

afterwards. Go to Chapter 5., "Basic Volume Configuration" to create volumes and map them to a host.



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 will describe how to create volumes and map them to the defined hosts.

The second part covers how to discover those volumes (5.1, "Provisioning storage from the IBM Storwize V7000 and making it available for the host"). Once you have finished with this chapter your basic configuration is done and you are able to store data on the IBM Storwize V7000.

Advanced host and volume administration, like volume migration, creating volume copies... and much more, is covered in Chapter 8, "Advanced Host and Volume Administration" on page 273.

5.1 Provisioning storage from the IBM Storwize V7000 and making it available for the host

This part follows the setup process of the book and shows how to create volumes and make them accessible from the host. The following steps are required to complete the basic setup of your environment:

- ► Create new volumes
- ► Map volumes to the host (5.2, "Map a Volume to the Host" on page 168)
- ▶ Discover the volumes from the host and specify multipath settings (5.3, "Discover the volumes from the host and specify multipath settings" on page 170)

To start, open the All Volumes section of the IBM Storwize V7000 GUI as shown in Figure 5-1 on page 158 to start the process to create new volumes.

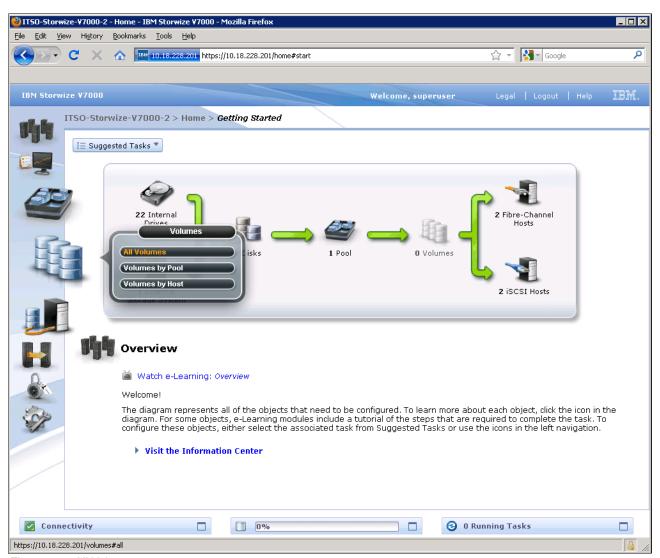


Figure 5-1 All Volumes

The "All Volumes" section will appear as shown in Figure 5-2 on page 159.



Figure 5-2 All Volumes section

At the moment we have not created any volumes, therefore the Welcome message is displayed. Follow the recommendation, click **New Volume** and a new window will appear as shown in Figure 5-3.

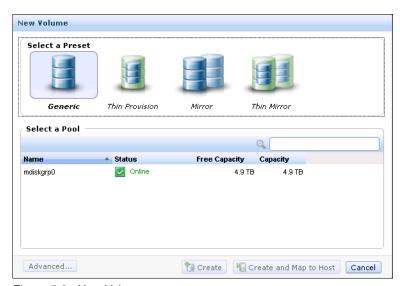


Figure 5-3 New Volume

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 for the user:

► **Generic**: a striped volume that is fully provisioned as shown in "Creating a Generic Volume" on page 160.

- ► Thin-provision: a striped volume that is space efficient. There are choices in the Advanced button to determine how much space is fully allocated initially and how large the volume is able to grow as shown in "Creating a Thin-provisioned Volume" on page 161.
- ▶ 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 shown in "Creating a Mirrored Volume" on page 163.
- ► **Thin-mirror**: two synchronized copies, both are thin-provisioned as shown in "Creating a Thin-mirror Volume" on page 166.

Select which volume you want to create and go to the section as listed above.

5.1.1 Creating a Generic Volume

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

We choose a generic volume as shown in Figure 5-3 on page 159, and afterwards we select the pool in which the volume should be created. Select the pool by clicking on it. In our example we have done that by clicking on mdiskgrp0. The result is shown in Figure 5-4.

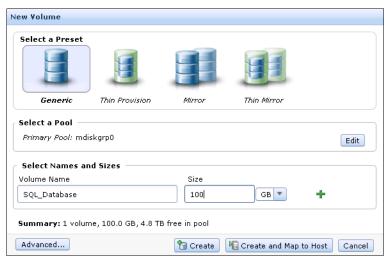


Figure 5-4 Creating a Generic Volume

Enter a Volume Name and a size, click **Create and Map to Host.** The new Generic Volume will be created as shown in Figure 5-5 on page 161. Click **Continue** and go to "Mapping newly created Volumes to the Host using the wizard" on page 168.

If you do not want to map the volumes now, just click **Create** to complete the task. Volumes can also be mapped later as described in 5.2, "Map a Volume to the Host" on page 168.

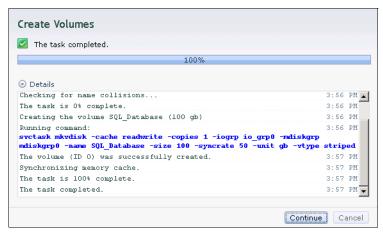


Figure 5-5 Volume Created

5.1.2 Creating a Thin-provisioned Volume

Volumes can be configured to be "Thin-provisioned". A thin-provisioned volume will behave with respect to application reads and writes as though they were fully allocated. When creating a thin-provisioned volume it is possible to specify two capacities: the real physical capacity allocated to the volume from the storage pool and its virtual capacity available to the host. So the real capacity will determine the quantity of extents that will be initially allocated to the volume. The virtual capacity will be the capacity of the volume reported to all other components (for example, FlashCopy, Cache, and Remote Copy) and to the host servers.

Note: Thin-provisioned volumes require additional I/O operations to read and write metadata to the back-end storage which also generates additional load to the system. Therefore it is recommended to not use thin-provisioned volumes for applications with very high I/O write workloads.

To create a thin-provisioned volume select **Thin Provision** as shown in Figure 5-6.



Figure 5-6 Create Thin Provisioned Volume

Select the pool in which the thin-provisioned volume should be created by clicking on it and enter the volume name and size as shown in Figure 5-7 on page 162.



Figure 5-7 Enter name and size

Note that under the name field there is a Summary showing that you are about to make a thin-provisioned volume, how much the virtual space will be, the space that will be allocated (real size) and the free capacity in the pool. By default the real capacity is 2% of the virtual capacity, you can change this setting by clicking **Advanced**. On the "Thin Provision" Tab as shown in Figure 5-8 are several advanced options available.



Figure 5-8 Advanced Settings

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

Make advanced settings if required and click **OK** to return to Figure 5-7 and click **Create and Map to Host** and the creation task will complete as shown in Figure 5-9 on page 163.

If you do not want to map the volumes now, just click **Create** to complete the task. Volumes can also be mapped later as described in 5.2, "Map a Volume to the Host" on page 168

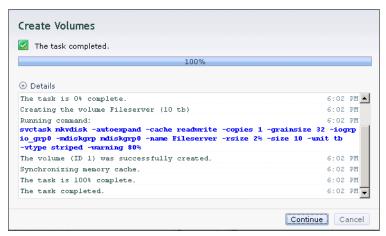


Figure 5-9 Create Thin-provisioned volume

Click **Continue** and go to "Mapping newly created Volumes to the Host using the wizard" on page 168.

5.1.3 Creating a Mirrored Volume

IBM Storwize V7000 offers the possibility to mirror volumes. That means a single volume is presented to the host, but two copies exist in the storage backend, 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. This lets you protect against array failures (for example multiple disk failures) and offers you more advanced features as described in 8.6, "Advanced Volume Copy Functions" on page 324. It also allows you to greatly improve the availability but it is not a disaster recovery solution due to both copies being accessed by the same node pair and only addressable by a single cluster.

To create a mirrored volume select **Mirror** as shown in Figure 5-10 on page 164.

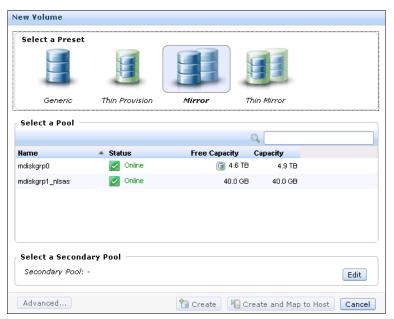


Figure 5-10 Creating a Mirrored Volume

Select the primary pool by clicking on it and the view will change to second pool as shown in Figure 5-11.

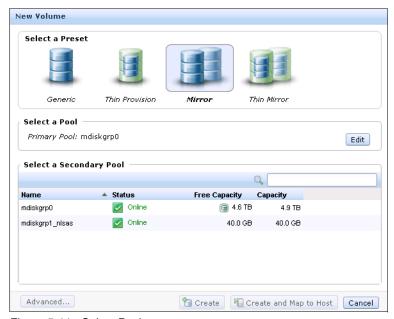


Figure 5-11 Select Pool

Select the secondary pool by clicking on it and enter a volume name and the required size as shown in Figure 5-12 on page 165.

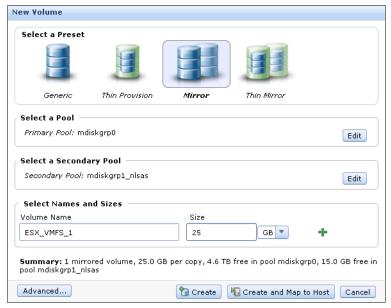


Figure 5-12 Select Pool, Name and Size

The summary shows you capacity information about the pool, and there is also the possibility to click **Advanced** and select the mirror tab as shown in Figure 5-13.

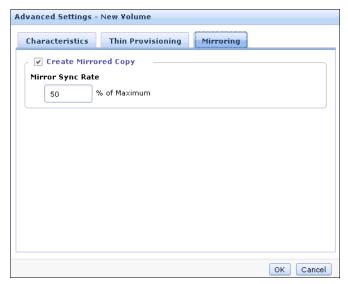


Figure 5-13 Advanced Mirror Features

In the advanced mirroring settings you are able to specify a synchronisation rate. Enter a **Mirror Sync Rate** between 1 and 100%. With this option you can set the importance of the copy synchronisation progress. This enables you to prefer more important volumes to synchronize faster than other mirrored volumes. By default the rate is set to 50% for all volumes. Click **OK** to return to Figure 5-12 on page 165.

Click **Create and Map to Host** and the mirrored volume will be created as shown in Figure 5-14 on page 166.



Figure 5-14 Creating a Mirrored Volume

Click **Continue** and go to "Mapping newly created Volumes to the Host using the wizard" on page 168

5.1.4 Creating a Thin-mirror Volume

A Thin-mirror volume combines only allocating the required space on demand (described in "Creating a Thin-provisioned Volume" on page 161) and having several copies of a volume (described in "Creating a Mirrored Volume" on page 163) available. To create a thin-mirror volume select **Thin Mirror** as shown in Figure 5-15.

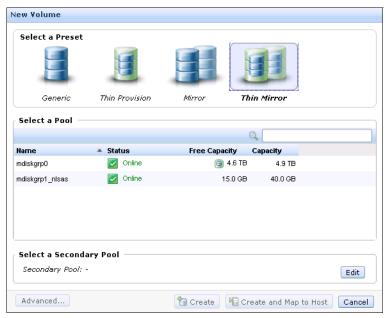


Figure 5-15 Thin Mirror

Select the primary pool by clicking on it and the view will change to the second pool as shown in Figure 5-16 on page 167.

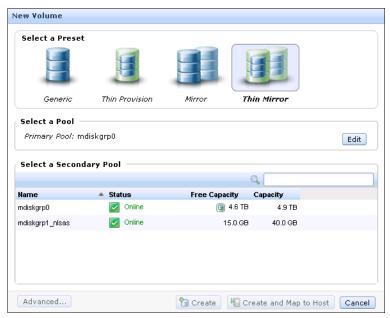


Figure 5-16 Select Pools

Select the pool for the secondary copy and enter a name and a size for the new volume as shown in Figure 5-17.

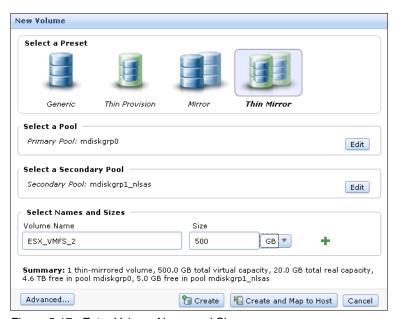


Figure 5-17 Enter Volume Name and Size

The summary shows you the capacity information, and the allocated space. Now you can click **Advanced** and customize the thin provision settings as described and shown in Figure 5-8, or the mirror synchronization rate as described and shown in Figure 5-13. If you have opened the advanced settings click **OK** to return to Figure 5-17.

Click **Create and Map to Host** and the mirrored volume will be created as shown in Figure 5-18 on page 168.

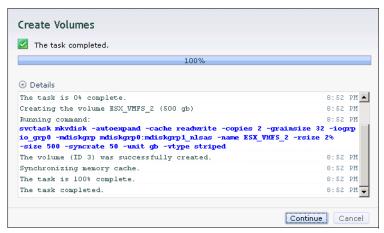


Figure 5-18 Create and Map to Host

Click **Continue** and go to "Mapping newly created Volumes to the Host using the wizard" on page 168

5.2 Map a Volume to the Host

The first part of this section ("Mapping newly created Volumes to the Host using the wizard") describes how to map a volume to a host if you have clicked the **Create and Map to Host** Button. The second part (5.2, "Map a Volume to the Host" on page 168) explains the manual host mapping process to create customized mappings.

5.2.1 Mapping newly created Volumes to the Host using the wizard

At this point, we continue to map the volume we have created in 5.2, "Map a Volume to the Host" on page 168. We assume that you have followed the procedure and clicked the **Continue** Button as, for example, shown in Figure 5-4 on page 160.

As the first step of the mapping process, you have to select a host to which the new volume should be attached to as shown in Figure 5-19.

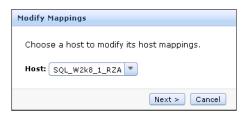


Figure 5-19 Choose Host

Now the wizards opens the Modify Mappings section, and has already preselected your host and the newly created volume. Just click **OK** and the volume will be mapped to the host as shown in Figure 5-20 on page 169.

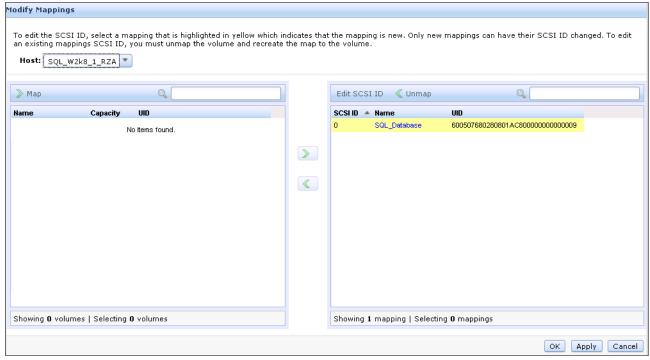


Figure 5-20 Modify Mappings

After the task is completed click as shown in Figure 5-21 and the wizard will return to the All Volumes section.

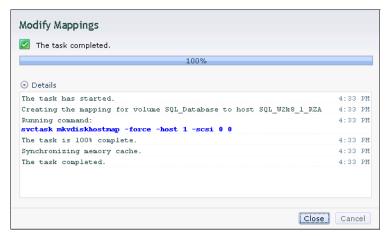


Figure 5-21 Modify Mappings complete

Now the newly created volume is displayed and we also see that it is already mapped to a host as shown in Figure 5-22 on page 170.

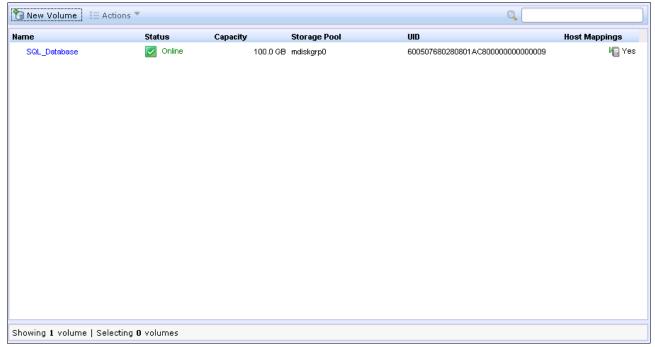


Figure 5-22 New Volume mapped to host

The host is now able to access the volumes and store data on it, go to 5.3, "Discover the volumes from the host and specify multipath settings" on page 170 to discover the volumes on the host and make some additional host settings if required.

Or create multiple volumes in preparation for discovering them later. Mappings can be customized as well. Advanced host configuration is covered in 8.1.1, "Modify Mappings" on page 276.

5.3 Discover the volumes from the host and specify multipath settings

This section shows how to discover the volumes created and mapped in the previous section and set additional multipath settings if required.

We assume that you have completed all steps described previously in the book so that the hosts and the IBM Storwize V7000 are prepared:

- Prepare your operating systems for attachment (Chapter 4, "Host Configuration" on page 127).
- Create Hosts using the GUI (4.2, "Creating Hosts using the GUI").
- Basic Volume configuration and host mapping (5.1, "Provisioning storage from the IBM Storwize V7000 and making it available for the host").

This section shows how to discover Fibre Channel and iSCSI Volumes from Windows 2008 and VMware ESX 4.x hosts.

In the IBM Storwize V7000 GUI select **Hosts** and **All Hosts** as shown in Figure 5-23 on page 171.

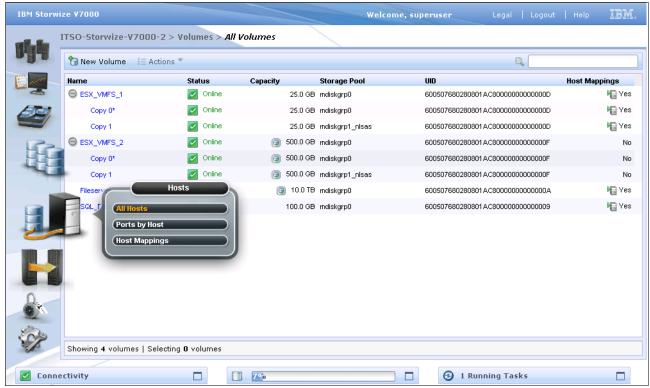


Figure 5-23 Open All Hosts

The view gives you an overview about the currently configured and mapped hosts as shown in Figure 5-24.

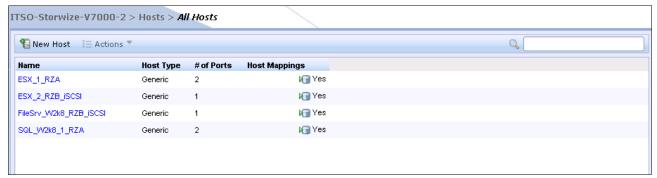


Figure 5-24 All Hosts

5.3.1 Windows 2008 Fibre Channel volume attachment

Right click your Windows 2008 Fibre Channel host in the All Hosts view (Figure 5-25 on page 172), select **Properties** and navigate to the **Mapped Volumes** tab as shown in Figure 5-26 on page 172.

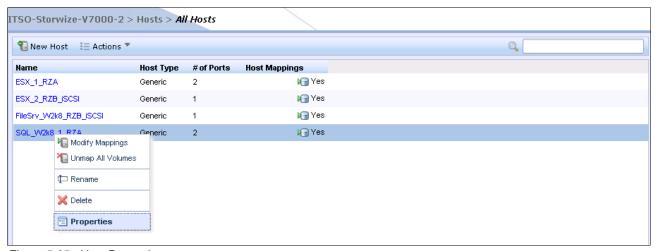


Figure 5-25 Host Properties

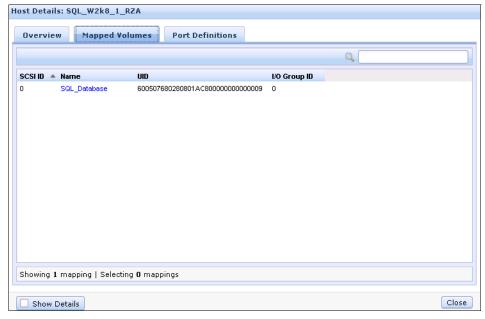


Figure 5-26 Mapped Volumes to a Host

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.

Log on to your Microsoft host and click "Start → All Programs → Subsystem Device Driver DSM → Subsystem Device Driver DSM". A command line box will appear. Enter the datapath query device command and press enter to see if there are IBM Storwize V7000 disks connected to this host as shown in Example 5-1.

Example 5-1 datapath query device

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

Total Devices: 1

DEV#: 0 DEVICE NAME: Disk14 Part0 TYPE: 2145 POLICY: OPTIMIZED

SERIAL: 600507680280801AC800000000000009

| ====== | :====================================== | | | | |
|--------|---|-------|--------|--------|--------|
| Path# | Adapter/Hard Disk | State | Mode | Select | Errors |
| 0 | Scsi Port5 BusO/Disk14 PartO | OPEN | NORMAL | 0 | 0 |
| 1 | Scsi Port5 BusO/Disk14 PartO | OPEN | NORMAL | 22 | 0 |
| 2 | Scsi Port6 BusO/Disk14 Part0 | OPEN | NORMAL | 19 | 0 |
| 3 | Scsi Port6 BusO/Disk14 Part0 | OPEN | NORMAL | 0 | 0 |

C:\Program Files\IBM\SDDDSM>

The output provides information about the connected volume(s), in the example shown there is one disk connected which is **Disk 14** for the Windows host, and 4 paths to the disk are available (State = Open). Open Windows Disk Management (Figure 5-27) by clicking **Start** → **Run**, type **diskmgmt.msc** and click **OK**.

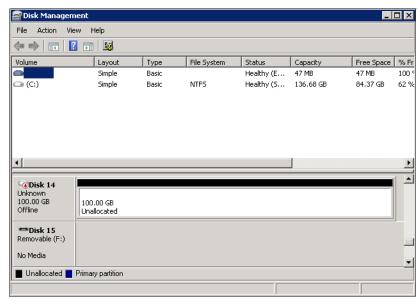


Figure 5-27 Windows disk management

Note: Usually Windows discovers new devices such as disks by itself,. If you have completed all steps and do not see any disks select **Actions** \rightarrow **Rescan Disk** in Disk Management to discover the new volumes.

Right click the disk in the left section and place it online as shown in Figure 5-28 on page 173.



Figure 5-28 Place disk online

Right click it again to initialize the disk as shown in Figure 5-29 and to confirm the task click **OK**.



Figure 5-29 Initialize Disk

Now right click the right section of the disk and select **Create simple Volume** as shown in Figure 5-30.



Figure 5-30 New Simple Volume

Follow the wizard and the volume will be ready to use from your Windows host as shown in Figure 5-31 on page 174.

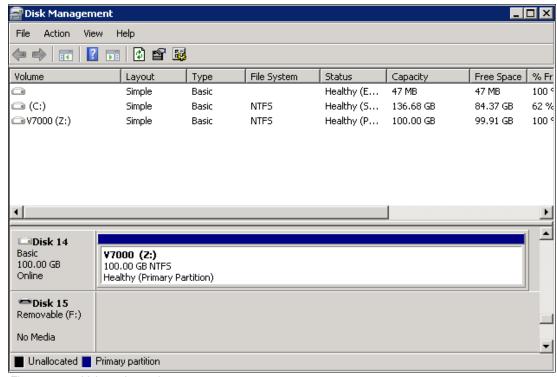


Figure 5-31 Volume is ready to use

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 Attachment

Right click your Windows 2008 iSCSI host in the all hosts view (Figure 5-32) select **Properties** and navigate to the **Mapped Volumes** tab as shown in Figure 5-33.

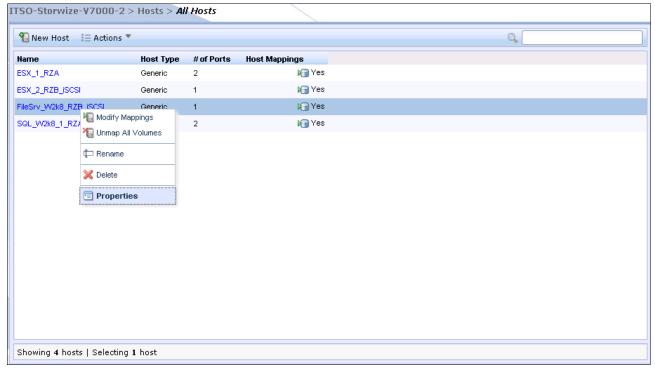


Figure 5-32 All Hosts view

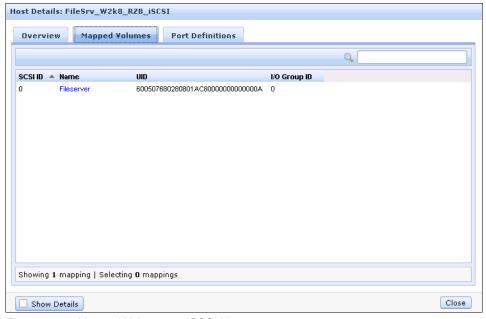


Figure 5-33 Mapped Volumes to iSCSI Host

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.

Log on to your Windows 2008 host and click **Start** \rightarrow **Administrative Tools** \rightarrow **iSCSI Initiator** to open the iSCSI configuration tab as shown in Figure 5-34 on page 176.

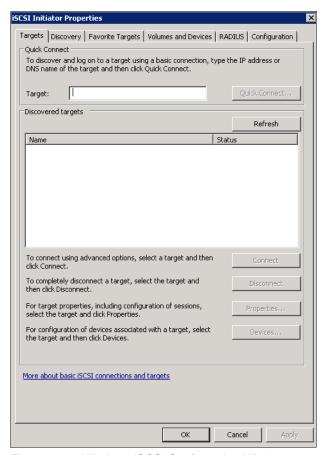


Figure 5-34 Windows iSCSI Configuration Windows

Enter the IP address of one of the IBM Storwize V7000 iSCSI ports and click **Quickconnect** as shown in Figure 5-35.

Note: The iSCSI IP addresses are different to the cluster and canister IP addresses, and they have been configured in 4.2.2, "Creating iSCSI Hosts" on page 150.



Figure 5-35 iSCSI Quickconnect

The IBM Storwize V7000 initiator will be discovered and connected as shown in Figure 5-36 on page 177.



Figure 5-36 iSCSI Initiator target is connected

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

Click **Start** → **Run** and type **cmd** to open a command prompt. Enter **ServerManagerCMD.exe**-install **Multipath-I0** and press enter as shown in Example 5-2.

Example 5-2 Installing MPIO

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

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

Success: Installation succeeded.
```

Click Start \rightarrow Administrative Tools \rightarrow MPIO and open the Discover Multi-Paths Tab and enable the Add support for iSCSI devices checkbox as shown in Figure 5-37 on page 178.

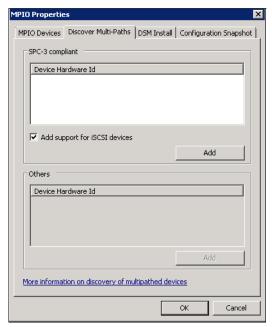


Figure 5-37 Enable iSCSI MPIO

Click **Add** and confirm the action message to reboot your host.

After the reboot log on again, click **Start** --> **Administrative Tools** --> **iSCSI Initiator** to open the iSCSI configuration tab and navigate to the **Discovery** Tab as shown in Figure 5-38 on page 179.

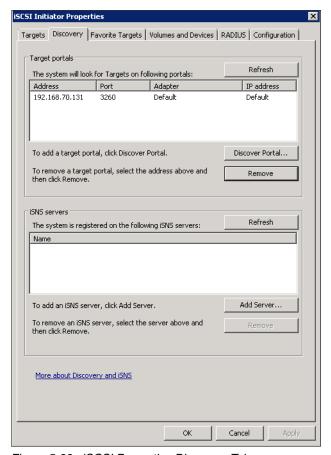


Figure 5-38 iSCSI Properties Discovery Tab

Click **Discover Portal...** and enter the IP Address of another IBM Storwize V7000 iSCSI port as shown in Figure 5-39 and click **OK**.

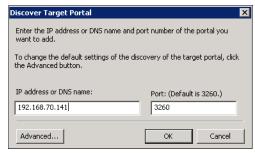


Figure 5-39 Discover Target Portal

Return to the **Targets** Tab as shown in Figure 5-40 on page 180 and you will find the new connection there listed as inactive.

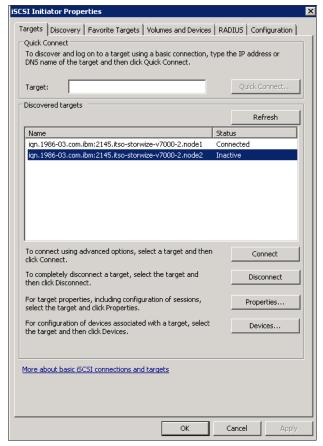


Figure 5-40 Inactive Target Ports

Highlight the inactive port and click **Connect**, the Connect to Target window will appear as shown in Figure 5-41.

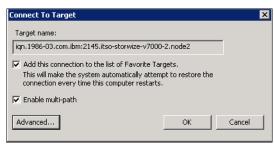


Figure 5-41 Connect to target

Make sure that you check the **Enable Multipath** checkbox and click **OK**, the second port will now appear as connected as shown in Figure 5-42 on page 181.

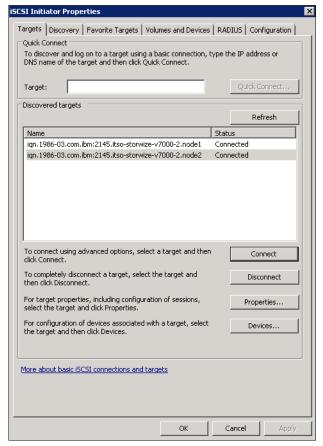


Figure 5-42 Second Target Port Connected

Repeat this step for each IBM Storwize V7000 port you want to use for iSCSI traffic, it is possible to have up to 4 port paths to the system.

Click **Devices** → **MPIO** to make sure that the multipath policy for Windows 2008 is set to default, **RoundRobin with Subset** as shown in (Figure on page 181), and click **OK** to close this view.

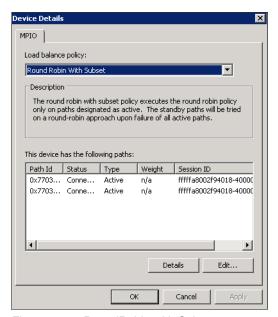


Figure 5-43 RoundRobin with Subset

Open the Windows Disk Management (Figure 5-44) by clicking $\mathbf{Start} \to \mathbf{Run}$ type $\mathbf{diskmgmt.msc}$ and click \mathbf{OK} .

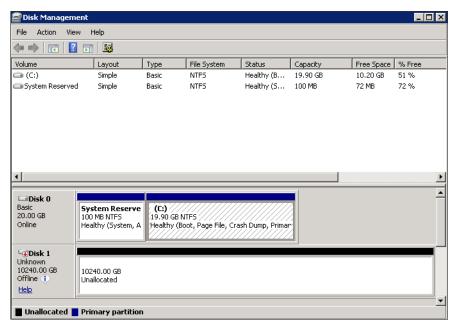


Figure 5-44 Windows Disk Management

Place the disk online, initialize it, create a filesystem 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 171.

Now the storage disk is ready for use as shown in Figure 5-45. In our example we have mapped a 10 TB disk, which is thin-provisioned on the IBM Storwize V7000 to a Windows 2008 host using iSCSI.

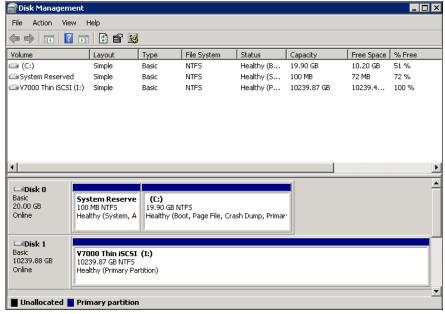


Figure 5-45 Windows Disk Management, Disk is Ready to Use

5.3.3 VMware ESX Fibre Channel Attachment

Right click your VMware ESX Fibre Channel host in the All Hosts view (Figure 5-46) select **Properties** and navigate to the **Mapped Volumes** tab as shown in Figure 5-47.

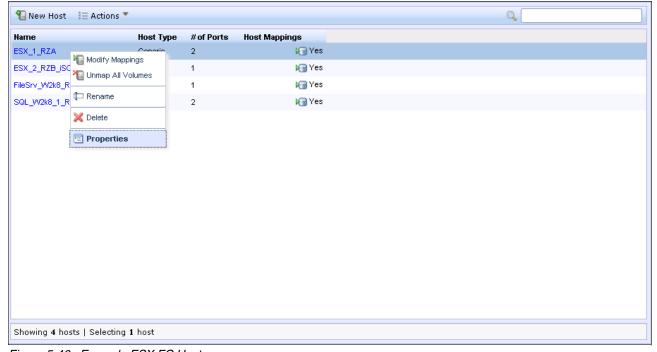


Figure 5-46 Example ESX FC Host

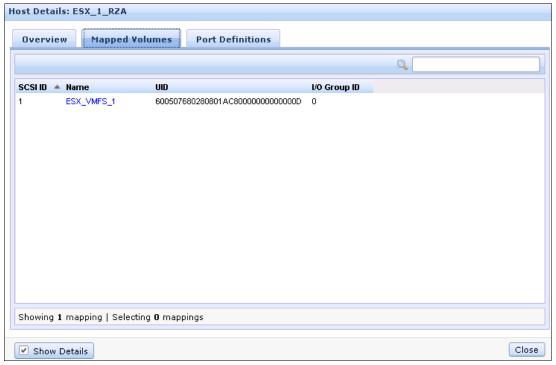


Figure 5-47 Mapped Volumes to ESX FC Host

In the host details section, you will see that at the moment there is one volume connected to the ESX FC host using SCSI ID 1. The UID of the volume is also displayed.

Connect to your VMware ESX Server using the vSphere client, navigate to the **Configuration** tab and select **Storage Adapters** as shown in Figure 5-48.

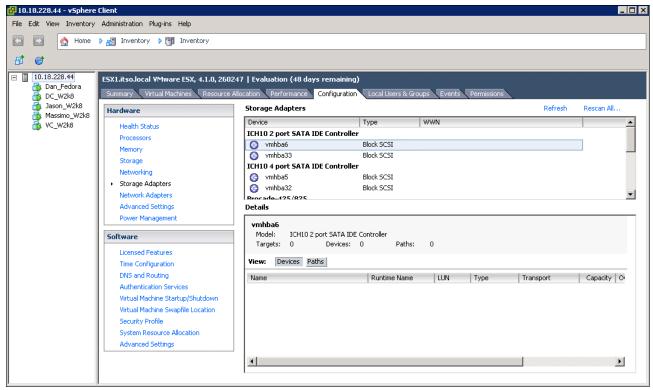


Figure 5-48 vSphere Client - Storage Adapters

Click **Rescan All...** and click **OK** (Figure 5-49) to scan for new storage devices.



Figure 5-49 Rescan

Select Storage and click Add Storage as shown in Figure 5-50 on page 186.

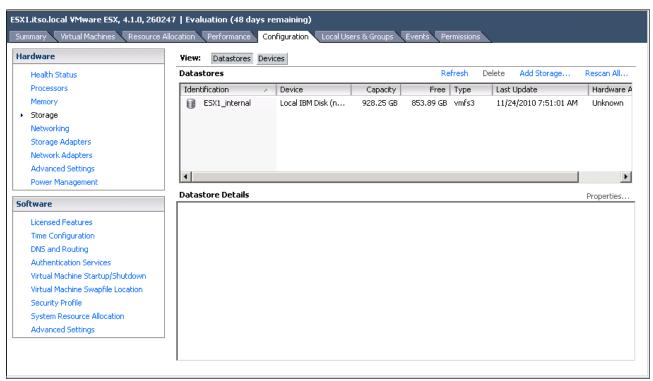


Figure 5-50 vSphere Client - Storage

The Add Storage wizard will appear, select **Select Disk/LUN** click **Next**. Now the IBM Storwize V7000 disk appears as shown in Figure 5-51. Highlight it and click **Next**.

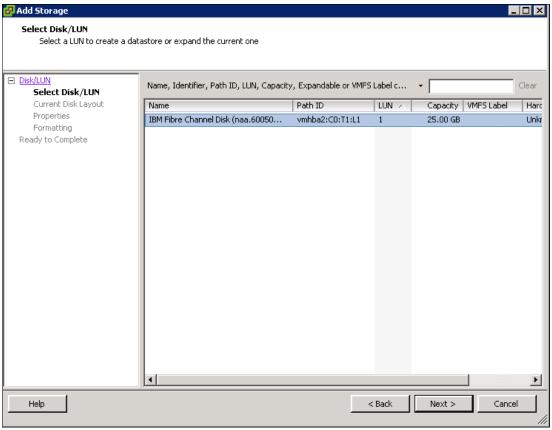


Figure 5-51 Select Disk/LUN

Follow the wizard to complete the attachment of the disk, after you select **Finish** the wizard closes and you return to the storage view. In Figure 5-52 on page 188 you see that the new volume has been added to the configuration.

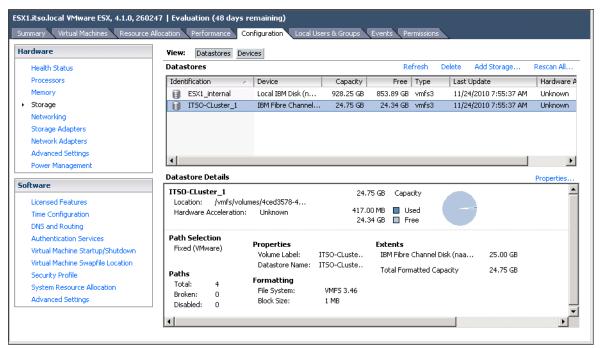


Figure 5-52 Add Storage task complete

Highlight the new Datastore and click **Properties** to see the details of it as shown in Figure 5-53.

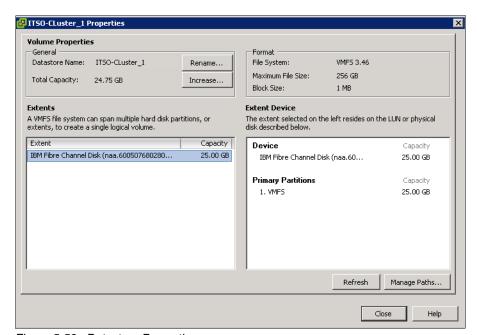


Figure 5-53 Datastore Properties

Click **Manage Paths** to customize the multipath settings. Select **RoundRobin** as shown in Figure 5-54 on page 189 and click **Change**.

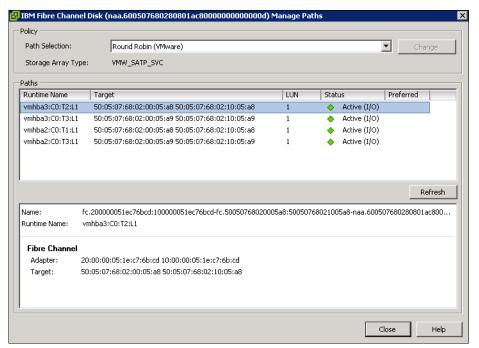


Figure 5-54 Select datastore multipath setting

Now 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

Right click your VMware ESX Fibre iSCSI host in the all hosts view (Figure 5-55) select **Properties** and navigate to the **Mapped Volumes** tab as shown in Figure 5-56.

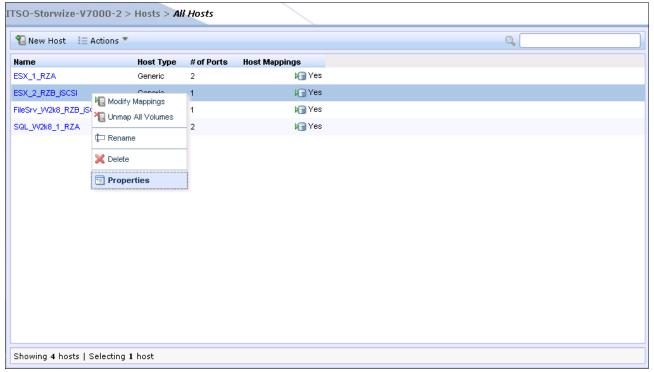


Figure 5-55 Select iSCSI ESX Host Properties

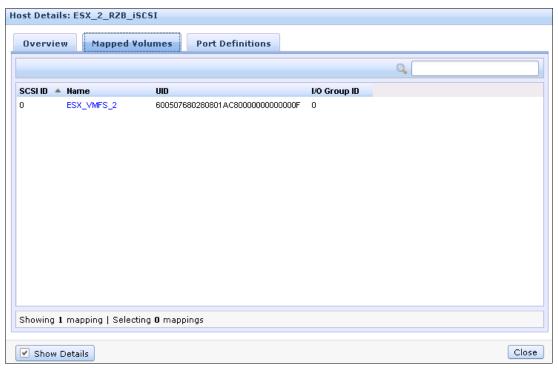


Figure 5-56 iSCSI ESX Host Properties

In the host details section, you will see that at the moment there is one volume connected to the ESX iSCSI host using SCSI ID 0. The UID of the volume is also displayed.

Connect to your VMware ESX Server using the vSphere client, navigate to the **Configuration** tab and select **Storage Adapters** as shown in Figure 5-57.

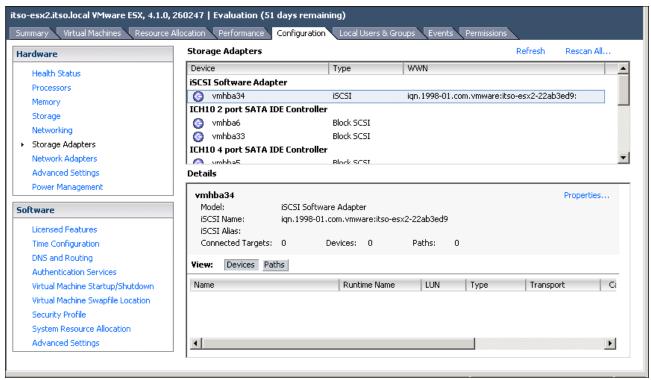


Figure 5-57 vSphere Client - Storage

Highlight the **iSCSI Software Initiator** and click **Properties**. The iSCSI initiator properties will appear, select the **Dynamic Discovery** Tab (Figure 5-58 on page 192) and click **Add**.

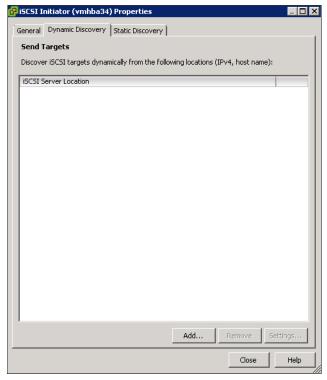


Figure 5-58 iSCSI Initiator Properties

To add a target enter the target IP address as shown in Figure 5-59. The target IP address is the IP address of a node in the I/O group from which you are mapping the iSCSI volume. Leave the IP port number at the default value of 3260, and click OK. The connection between the initiator and target is established after clicking OK

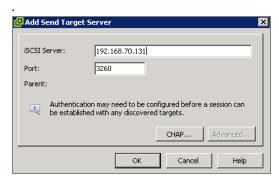


Figure 5-59 Enter Target IP Address

Repeat this step for each IBM Storwize V7000 iSCSI Port you want to use for iSCSI connections.

Note: The iSCSI IP addresses are different to the cluster and canister IP addresses, they have been configured in 4.2.2, "Creating iSCSI Hosts" on page 150.

After you have added all the ports required close the iSCSI Initiator Properties by clicking **Close** (Figure 5-58).

You will be prompted to rescan for new storage devices, confirm the scan by clicking **Yes** as shown in Figure 5-60 on page 193.



Figure 5-60 Confirm Rescan

Go to the storage view as shown in Figure 5-61 and click Add Storage.

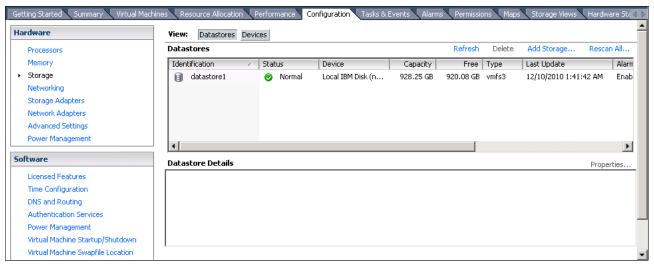


Figure 5-61 Add Storage

The "Add Storage" wizard will appear (Figure 5-62 on page 194). Select **Disk/Lun** and click **Next**.

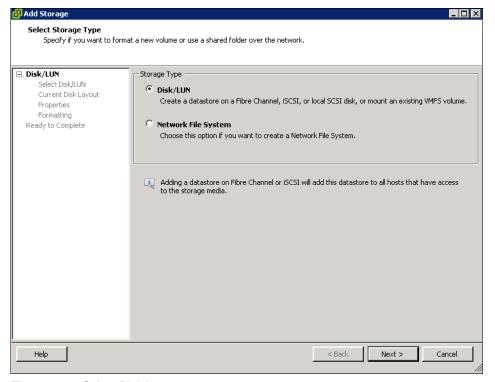


Figure 5-62 Select Disk Lun

The new iSCSI LUN will appear, highlight it and click **Next** as shown in Figure 5-63.

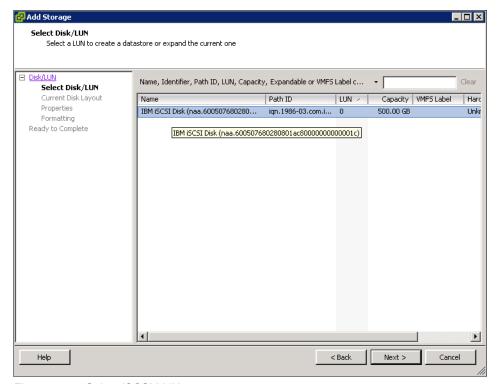


Figure 5-63 Select iSCSI LUN

Review the disk layout and click **Next** as shown in Figure 5-64 on page 195.

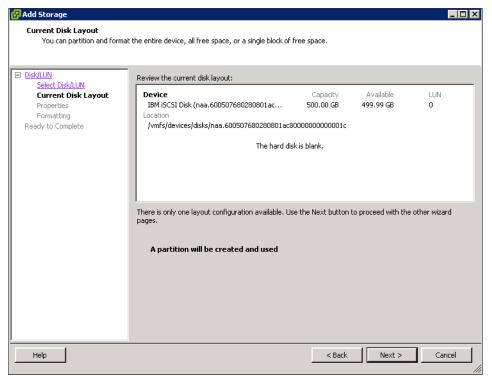


Figure 5-64 Current Disk Layout

Enter a name for the datastore and click **Next** as shown in Figure 5-65.

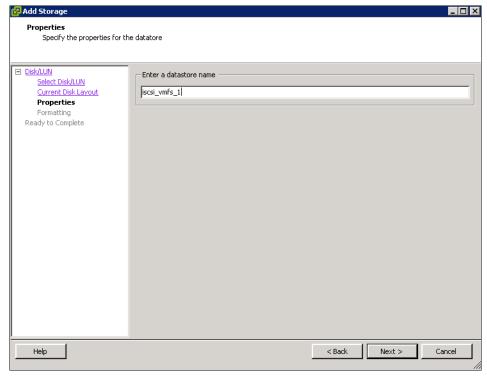


Figure 5-65 Enter a datastore name

Select the maximum file size and click **Next** as shown in Figure 5-66.

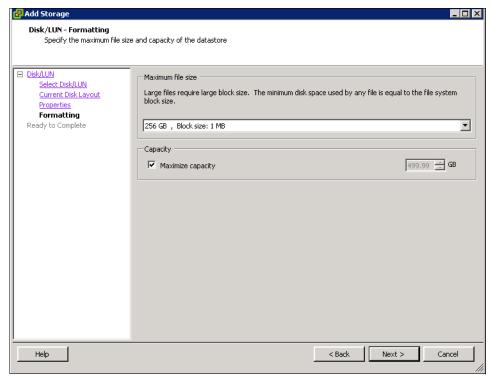


Figure 5-66 Maximum file size

Review your selections and click **Finish** as shown in Figure 5-67.

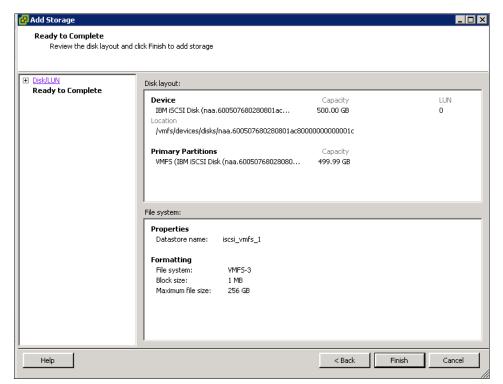


Figure 5-67 Finish Wizard

The new iSCSI LUN is now in the process of being added and this can take a few minutes. Once the tasks complete, the new datastore appears in the storage view as shown in Figure 5-68 on page 197.

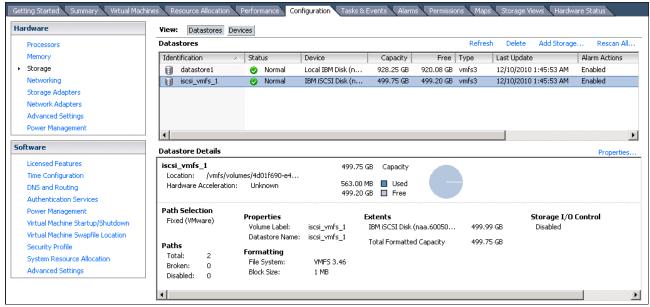


Figure 5-68 New Datastore available

Highlight the new datastore and click **Properties** to open and review the datastore settings as shown in Figure 5-69.

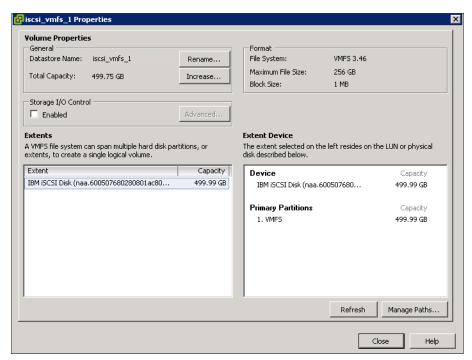


Figure 5-69 Datastore Properties

Click **Manage Paths**, select **Round Robin** as the multipath policy as shown in Figure 5-70 on page 198 and click **Change**.

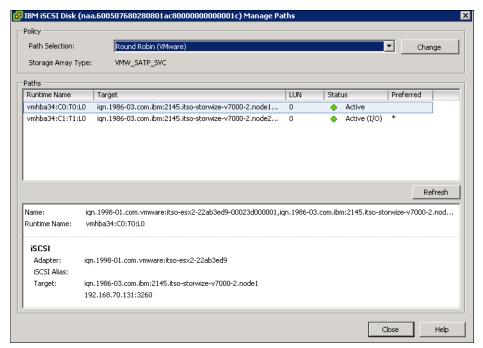


Figure 5-70 Change Multipath Policy

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 on legacy storage systems into the IBM Storwize V7000 storage system. Migrating data from legacy storage systems to the IBM Storwize V7000 storage system will let applications benefit from the new features such as Easy Tier, an intuitive management GUI, and advanced storage replication functions to better support applications.

To migrate existing data the IBM Storwize V7000 provides a storage migration wizard to guide you through the entire procedure.

6.1 Preparation for Data Migration

In this section we will show how to prepare external storage systems and IBM Storwize V7000 for data migration. Data migration is always an important step when involving new storage systems into existing storage infrastructure.

Before attaching any external storage systems to IBM Storwize V7000, check the IBM Storwize V7000 support matrix at the website below:

http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003703

When migrating data from an external storage system to IBM Storwize V7000, where the external storage system will be removed from IBM Storwize V7000 control when complete, IBM allows you to temporarily configure the external virtualization license setting. Configuring the external license setting will prevent messages from being sent indicating you are in violation of the license agreement. When the migration is complete the external virtualization license must be reset to its original limit.

Note: When external storage is to be virtualized by IBM Storwize V7000 then a per-enclosure external virtualization license is required. See your IBM account team or IBM Authorized Business Partner for further assistance if required.

To prepare the data migration, external storage systems need to be configured to be under IBM Storwize V7000 control. These are the steps to follow:

- 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.
- Update your host device drivers, including your multipath driver and configure them for attachment to IBM Storwize V7000 system.
- 4. Create a storage system zone between the storage system being migrated and IBM Storwize V7000 system, and 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 Migrate the Data Using the Migration Wizard

In this section we will discuss how to migrate existing data on external storage systems using the IBM Storwize V7000 storage migration wizard.

We will use the IBM Storwize V7000 wizard that has been designed specifically for this scenario to guide you through the process and discuss the steps along the way.

There are two ways to access the menu options for starting a migration. One starts from the the initial **Getting Started** panel where you can select **Migrate Storage** from the **Suggested Tasks** drop down, as shown in Figure 6-1.

Figure 6-1 displays the **Getting Started** panel and the **Migrate Storage** option in the **Suggested Tasks** drop down.

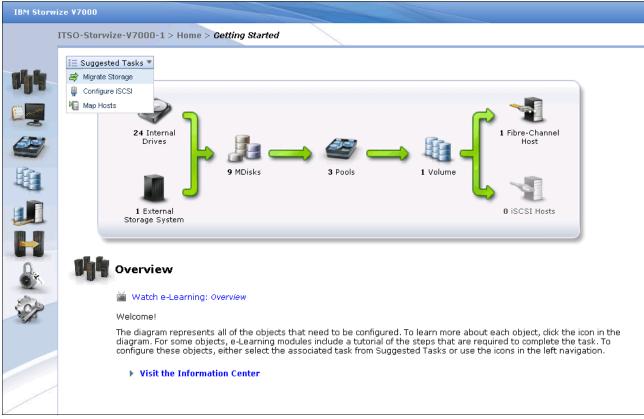


Figure 6-1 Getting Started panel with Migrate Storage option displayed

The other way is to navigate to the **Migration** option via the **Physical Storage** function icon, the third one on the left side of the panel. Figure 6-2 on page 202 shows the **Migration** option via the **Physical Storage** function icon.



Figure 6-2 Physical Storage icon with Migration option

Whichever method is chosen, the storage migration panel will appear. Click on **Start New Migration** to start the storage migration wizard.

Figure 6-3 shows the storage migration panel.

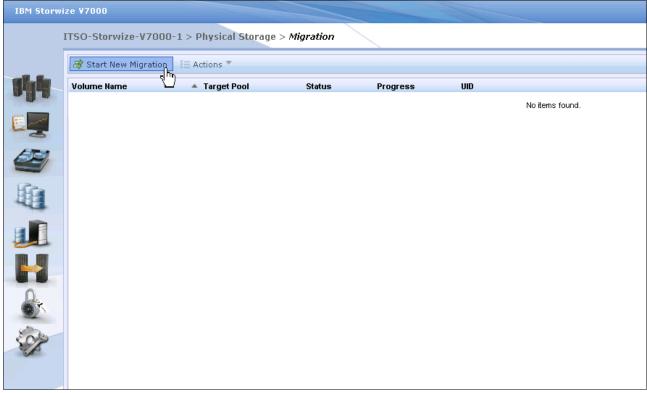


Figure 6-3 Storage migration panel

Now using the IBM Storwize V7000 storage migration wizard, you can easily migrate your existing data. Follow the steps below:

1. Follow Step 1 of storage migration wizard, check the restrictions and prerequisites, and Click **Next**.

Figure 6-4 on page 204 shows Step 1 in the storage migration wizard.

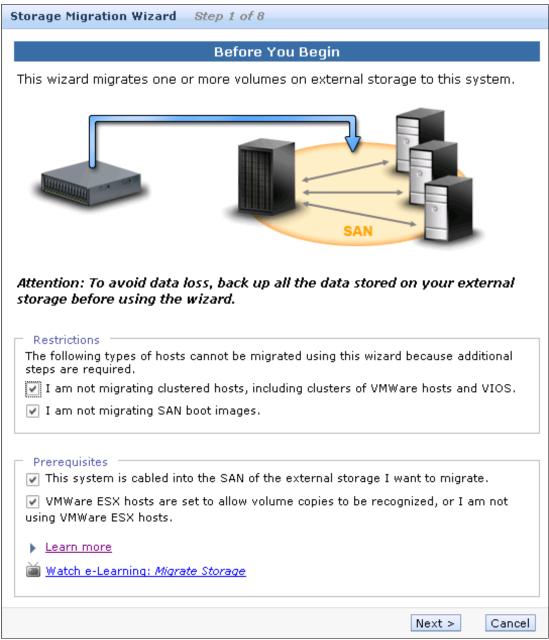


Figure 6-4 Before you begin

Note: To avoid any potential data loss, back up all the data stored on your external storage before using the wizard.

Step 1 shows there are some restrictions and prerequisites in using the storage migration wizard.

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 will need to migrate them outside of this wizard. You can find more information on IBM Storwize V7000 information center:

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

The VMware ESX Storage vMotion feature might be an alternative to migrate VMware clusters, consult VMware for more information.

Prerequisites:

- Make sure the IBM Storwize V7000 Fibre Channel ports have been connected to the SAN fabric(s) that the external disk controller and hosts you wish to migrate from are connected to.
- 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, after you have understood the restrictions and satisfied any prerequisites.

2. Follow Step 2 of the storage migration wizard, complete the environment preparation for migration, and Click **Next**.

Figure 6-5 on page 206 shows Step 2 of the storage migration wizard.

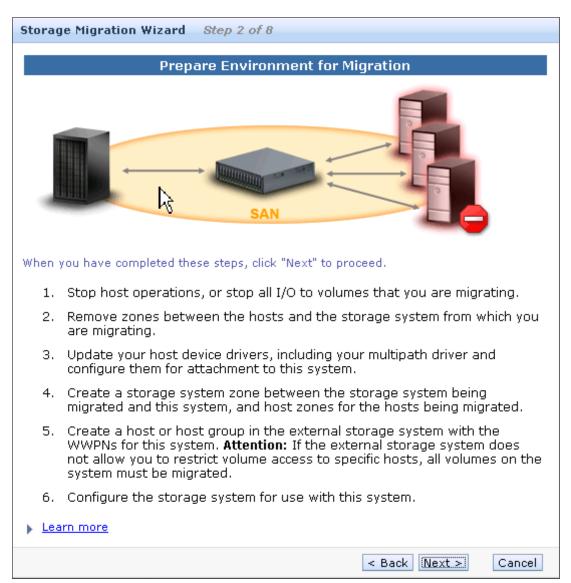


Figure 6-5 Prepare environment for migration

3. Follow Step 3 of the storage migration wizard, complete the mapping of external storage LUNS. Make sure you record the information mentioned in this step, as it makes it much easier for your following steps. Then click **Next**.

Note: You may need to record the SCSI ID that the volume is mapped to the host with. Some operating systems do not support changing the SCSI ID during the migration.

Figure 6-6 shows Step 3 of the storage migration wizard.

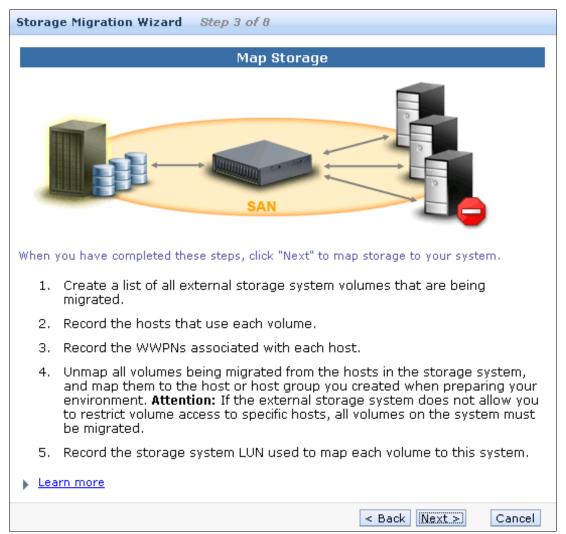


Figure 6-6 Map storage

After you click **Next**, the IBM Storwize V7000 will start to discover external devices (if you have correctly zoned the external storage systems with the IBM Storwize V7000 and mapped the LUNs). When the discovery complete, IBM Storwize V7000 will show the MDisks found.

4. Choose the MDisks you want to migrate, and Click Next.

Figure 6-7 on page 208 shows step 4 of the storage migration wizard.

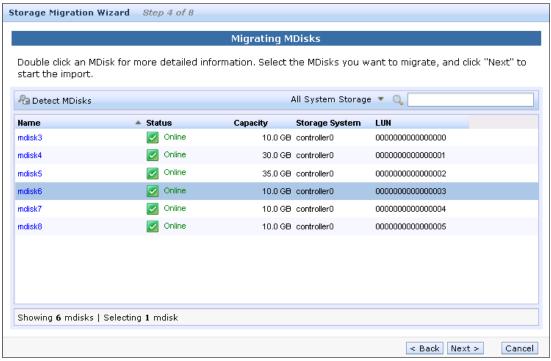


Figure 6-7 Migrating MDisks

If the MDisks that need migrating are in the list, select them and click **Next**, then IBM Storwize V7000 will start to import the MDisks that you have chosen. If the MDisks that need migrating are not in the list, you may need to check your zone configuration and LUN mapping, and click **Detect MDisks** to let IBM Storwize V7000 run the discovery procedure again.

You can select one or more MDisks as required, and detailed information of the MDisk can be shown by double clicking on it.

In Figure 6-7 there are six LUNs discovered as MDisks that are candidates for migration. In your particular situation, you may need to reference the information you recorded earlier to identify these MDisks. In our example the MDisks have been selected to go forward to the next step.

When you click **Next** in this step, IBM Storwize V7000 will complete importing the MDisks with the host's data and a storage pool has been created. The MDisks are added to the pool and Image Mode volumes (with the same size as the MDisks) are created which are ready for mapping back to the original hosts.

Configure the host which needs to access the data after the migration, and/or create new hosts as needed and Click Next.

Figure 6-8 on page 209 shows Step 5 of the storage migration wizard.

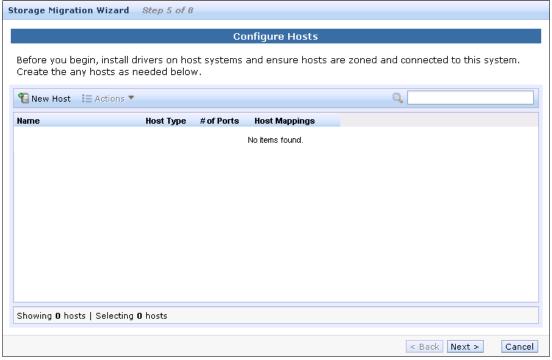


Figure 6-8 Configure Hosts

Before you configure any hosts, make sure appropriate drivers have been installed on the host and host zones have been zoned correctly.

If the host which 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.

Figure 6-9 on page 210 shows the Create Host window.

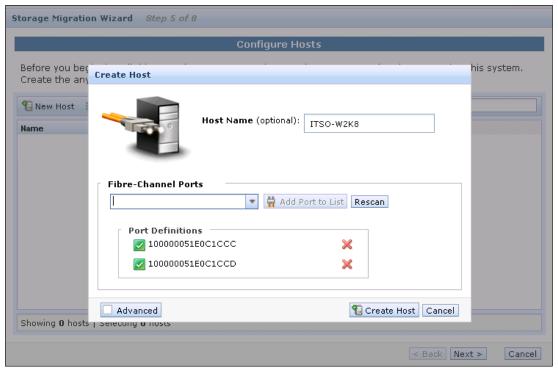


Figure 6-9 Create host window

In the **Create Host** window, give the host a name and from the **Fibre-Channel Ports** drop-down list select the WWPNs for this host that you recorded earlier. If WWPNs of this host do not appear in the drop down list. Click the **Rescan** button and try again. If they still do not show up in the list you can manually enter them.

When ready to proceed click on **Create Host**.

After the host has been created, you can find it in the host list in Step 5. Then click **Next** to continue with the migration wizard.

Figure 6-10 shows Step 5 of the storage migration wizard with the new host.

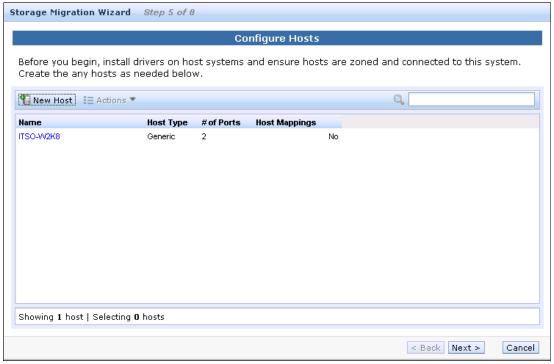


Figure 6-10 Step 5 with newly created host ITSO_W2K8.

6. Map the newly migrated volume to the host, and when mapping is complete, click Next.

Figure 6-11 shows Step 6 of the storage migration wizard.

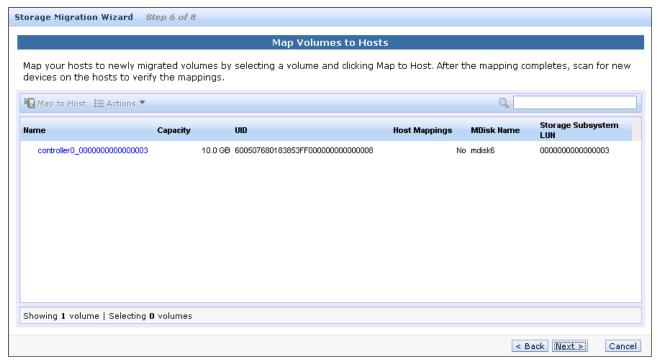


Figure 6-11 Map volumes to hosts

In Step 6 of the migration wizard, the volumes from the imported MDisks which need migrating have been listed. The names of the volumes have been 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 click on **Rename** in the **Actions** drop-down menu.

Note: The names must begin with a letter, which cannot be numeric. 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. And 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 will pop up with a drop-down list of the hosts. At this time, the IBM Storwize V7000 will let you choose which host you need to map the volumes to. Choose the correct host and click **Next**.

Note: Best practice recommends to map the volume to 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 map the volumes to.

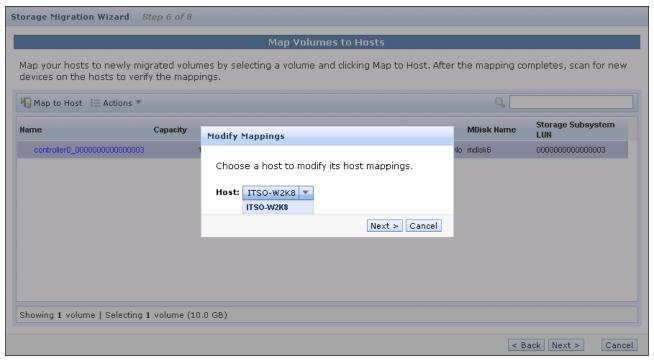


Figure 6-12 Choose the host to map the volumes to

After choosing the host, you will enter the **Modify Mappings** panel of the IBM Storwize V7000. 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 on page 213 shows the **Modify Mappings** panel in the migration wizard.

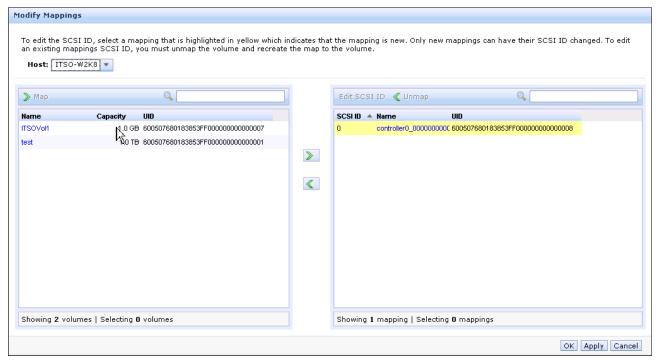


Figure 6-13 Modify Mappings panel in migration wizard

When the mapping completes, in Step 6 of the storage migration wizard, you will find the Host Mappings column of the volumes change from No to Yes. A scan could 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.

Figure 6-14 shows Step 6 of the storage migration wizard with host mappings modified.

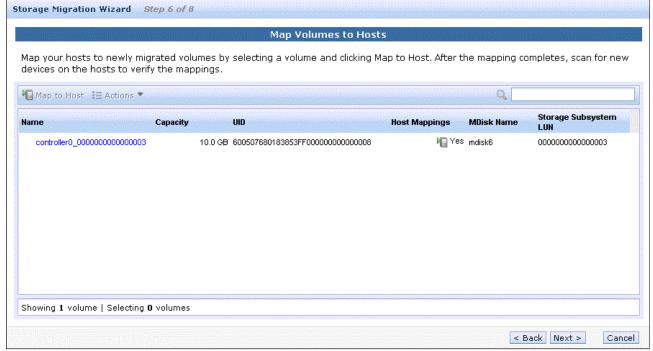


Figure 6-14 Completion of the host mappings

7. Select the destination storage pool for data migration, and click Next.

Figure 6-15 shows Step 7 of the storage migration wizard.

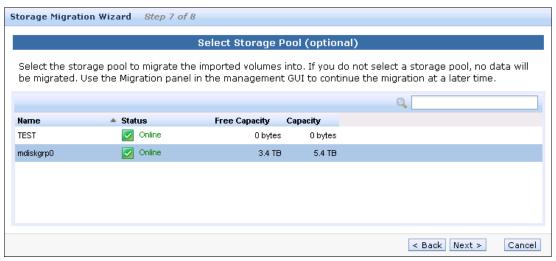


Figure 6-15 Select storage pool

The destination storage pool of data migration could be an external storage pool or an internal storage pool. Make sure there is enough space in the storage pools.

After you click **Next**, the migration begins. The migration will run 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 will have pointers to both the new copy on the storage pool selected as well as on the original external storage system.

8. Click **Finish** to end the storage migration wizard in Step 8.

Figure 6-16 shows Step 8 of the storage migration wizard.

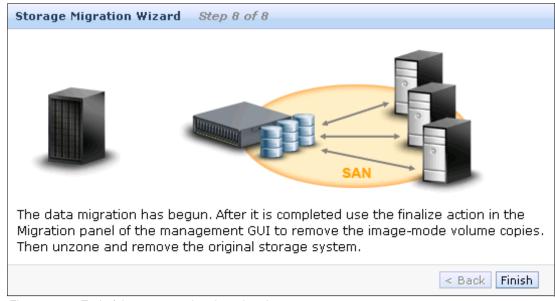


Figure 6-16 End of the storage migration wizard

The end of the storage migration wizard is not the end of the data migration. The data migration has just begun, and after clicking **Finish** in Step 8, you can find the migration progress in the migration panel. You can also find the target storage pool which your volumes are being migrated to, along with the status of the volumes.

Figure 6-17 on page 215 shows the data migration progress in the migration panel.

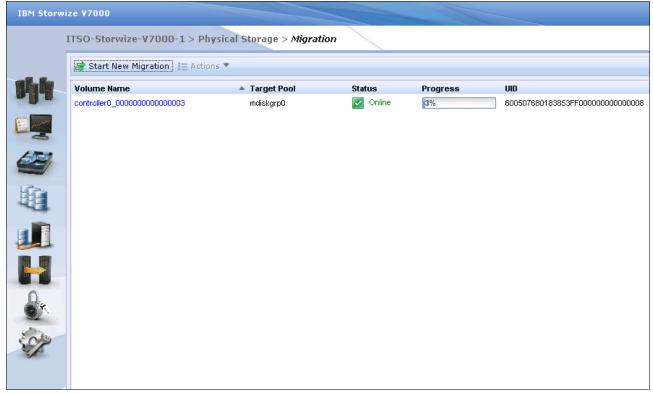


Figure 6-17 Migration progress in the migration panel

When the migration progress reaches 100%, select the volumes and click **Finalize** in the **Actions** drop-down menu in the migration panel, as shown in Figure 6-18. The Image Mode copy of the volumes on the original external storage system will be deleted and the associated MDisks from the storage pool will be removed, and the status of those MDisks will be unmanaged.

Figure 6-18 shows how to finalize the data migration in IBM Storwize V7000.

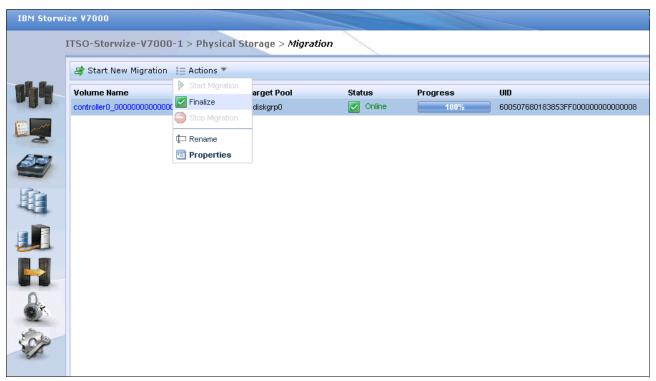


Figure 6-18 Finalize the migration

After the you select **Finalize** in the **Action** list, the IBM Storwize V7000 will need your confirmation on finalizing the migration for the volumes. Verify the volume name and the number of migrations you are finalizing, and if you are satisfied click **OK**.

Figure 6-19 shows the Finalize Volume Migration window.

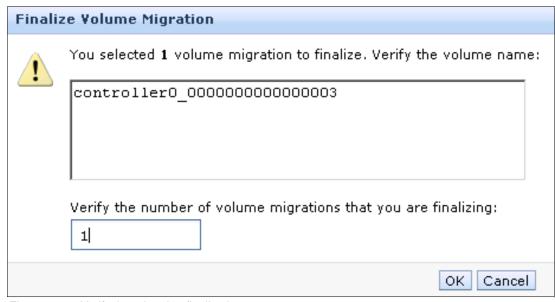


Figure 6-19 Verify the migration finalization

When the finalization completes, the data migration to the IBM Storwize V7000 is done, you can un-zone and remove the legacy storage system from the IBM Storwize V7000.

For more information on advanced migration function, refer to Chapter 7, "Storage Pools" on page 219 and Chapter 9, "External Storage Virtualization" on page 339.



7

Storage Pools

This chapter explains how IBM Storwize V7000 manages physical storage resources. All storage resources under IBM Storwize V7000 control are managed with the concept of storage pools. Storage pools make the storage easy to dynamically allocate resources, maximize its productivity and reduce the cost. Advanced internal storage, MDisk and storage pool management will be covered in this chapter, while external storage will be covered in Chapter 9, "External Storage Virtualization" on page 339.

7.1 Work with Internal Drives

This section discusses how to configure the internal storage disk drives using different RAID levels and different optimization strategies.

In this section, we are starting with the environment as shown in Figure 7-1. At this moment in time all the Internal drives remain unconfigured, and we will cover the steps to configure internal storage later. Currently the existing MDisks come from external storage, and example storage pools, volumes and hosts have been created for use.

You can learn how to manage MDisks in 7.2, "Work with MDisks" on page 239; how to manage storage pools in 7.3, "Work with Storage Pools" on page 266; how to work with external storage in << refer to external storage chapter>>; how to create volumes in << refer to basic volume configuration chapter>>; and how to create hosts in << refer to basic host configuration chapter>>.

The IBM Storwize V7000 storage system provides an individual **Internal** panel for managing all internal drives.

You can access the **Internal** panel through the **Getting Started** panel, and click the **Internal Drives** icon. Then extended help information for internal drives will show up below, click **Physical Storage** and you will be taken to the **Internal** panel.

Figure 7-1 shows how to access the **Internal** panel from the **Getting Started** panel.

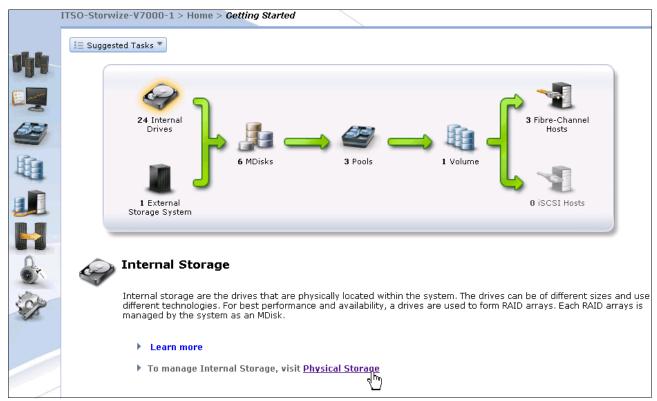


Figure 7-1 Access the Internal panel from Getting Started panel

The other way to access the **Internal** panel will be from the **Physical Storage** functional icons on the left hand side.

Figure 7-2 on page 221shows how to access the **Internal** panel from the **Physical Storage** functional icons on the left hand side.

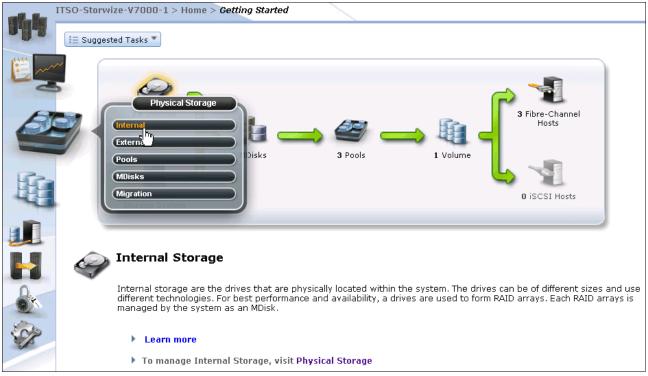


Figure 7-2 Access Internal panel from Physical Storage function icon

The **Internal** panel, as shown in Figure 7-3, gives an overview of all your internal drives. On the left of the panel there is a catalog of the internal drives. In this catalog you can find out how many different types of internal disks are in this IBM Storwize V7000 storage system. Select any type on the left, and the internal disks of this type will show up on the right. Select **All Internal**, and all of your internal disks will list on the right.

Figure 7-3 shows the Internal panel.

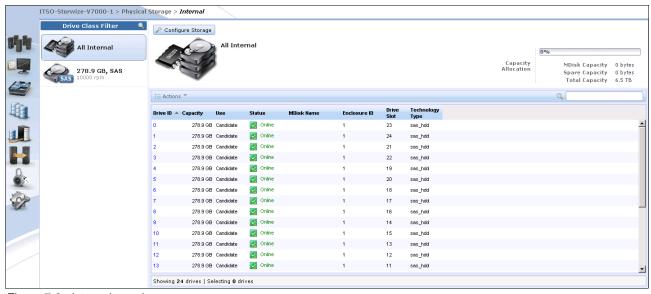


Figure 7-3 Internal panel

On the right hand side of the **Internal** panel, the internal disks of the type you selected will be listed along with their general information, including the drive ID, capacity, the role of drive, status, MDisk name, enclosure ID, drive slot, drive type and so on.

In addition, you can find the current internal storage capacity allocation indicator on the top right. The Total Capacity shows the internal storage capacity you have overall in this IBM Storwize V7000 storage system; The MDisk Capacity shows the internal storage capacity have been assigned to MDisks; The Spare Capacity shows the internal storage capacity used for hot spare. Now the percentage bar indicates 0% capacity allocated, because no internal storage have been configured in this example.

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 or click the **Actions** drop down, as shown in Figure 7-4.

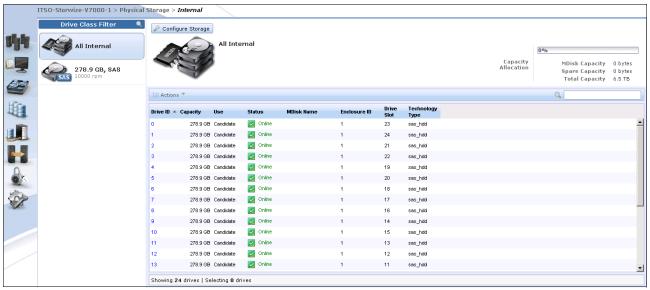


Figure 7-4 Take actions on internal drives

- Fix Error action will start the fix procedure, << refer to RAS chapter>>.
- 2. The internal drives could be taken offline when there are problems on the drives by clicking on **Take Offline** in the **Actions** drop down list. A confirmation window appears as shown in Figure 7-5 on page 223. The IBM Storwize V7000 storage system prevents the drive from being taken offline if there may be data loss as a result. It is recommended to only take the drive offline if a spare drive is available.

Note: Choosing the option to take internal drives offline even if redundancy is lost on the array could lead to potential data loss.

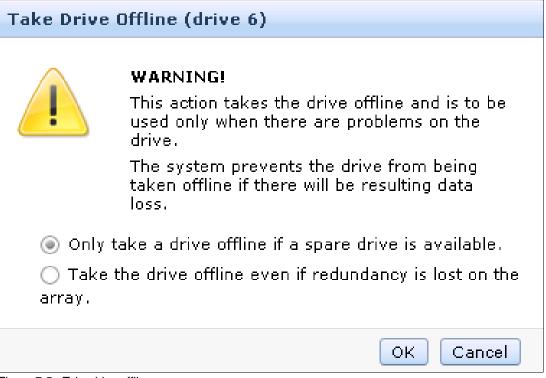


Figure 7-5 Take drive offline

- 3. The internal drives in the IBM Storwize V7000 storage system could be assigned to several roles and are designated as unused, candidate or spare. The meaning of these roles are:
 - Unused: the drive is not in use and will not be used as a spare.
 - Candidate: the drive is available for use in an array.
 - Spare: the drive can be used as a hot spare if required.

Select the **Mark as...** in the **Actions** drop down, and select the role you want the drive to be assigned as shown in Figure 7-6 on page 224.

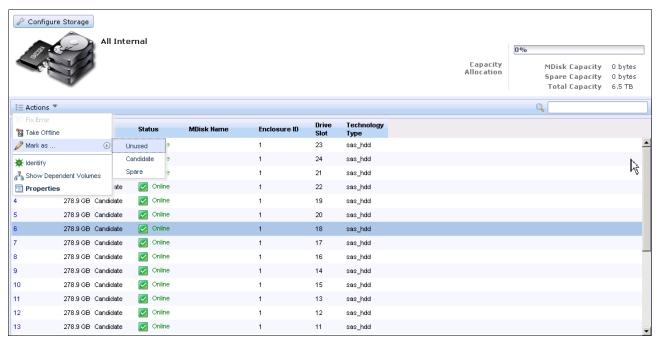


Figure 7-6 Assign role to internal drive

4. 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 7-7 shows the information when you click the **Identify** action.

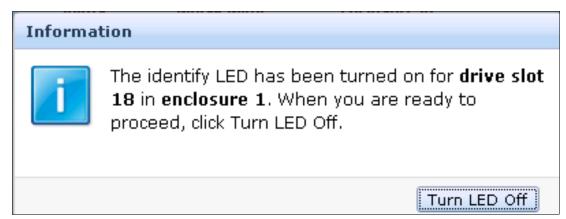


Figure 7-7 Drive identification information

Click Turn LED Off, when you are done.

- Show Dependent Volumes will show you the volumes which are dependent on the drives.
- 6. The **Properties** action in the **Actions** drop down list will show you more information about the drives, as shown in Figure 7-8 on page 225. With the **Show Details** checkbox selected, you can find much more detailed information, including vendor ID, FRU Part Number and ID, as shown in Figure 7-9 on page 225.

Figure 7-8 on page 225 shows the **Properties** tab with default format.

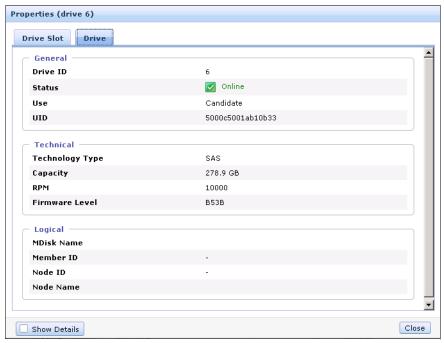


Figure 7-8 Properties tab with default format

Figure 7-9 shows the **Properties** tab with detail format.

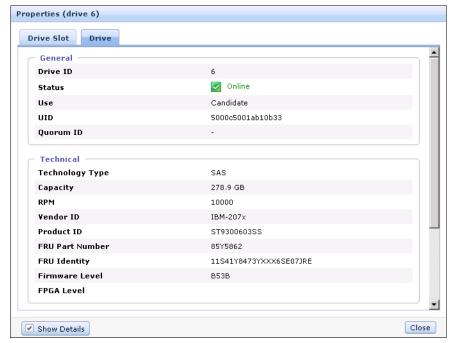


Figure 7-9 Properties tab with detail format

7.1.2 Configuring Internal Storage

To configure internal storage for use, click the **Configure Storage** button on the top left on the right part of the **Internal** panel, as shown in Figure 7-3 on page 221. A configuration wizard

will pop up to guide you the process of configuring internal storage, as shown in Figure 7-3 on page 221.

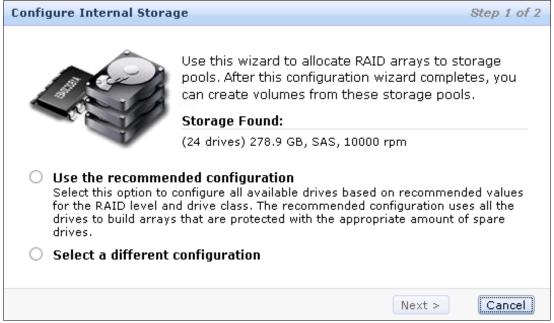


Figure 7-10 Wizard for configure internal storage

The wizard will include all the candidate internal drives in the configuration. If you have internal drives with the unused role, a pop-up window will appear to ask you if you want to include the unused drives in the configuration.

Usage of the storage configuration wizard simplifies the initial disk drive setup and offers two options:

- Use the recommended configuration
- ► Select a different configuration

Selecting the **Use the recommended configuration** will guide us through the wizard discussed in "Usage of the Recommended Configuration" on page 228. **Select a different configuration** uses the wizard discussed in "Select a Different Configuration" on page 231.

Before going through the storage configuration wizard, the concept of the IBM Storwize V7000 RAID configuration presets will be introduced first with the basic concept of RAID itself described in << refer to the overview chapter>>.

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.

Table 7-1 on page 227 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 strip of parity are striped across all array members. | 5 | 8 | 1 |
| SSD RAID 6 | Protects against two drive failures. Data and two strips 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 Easy Tier function. | 10 | 2 | 1 |

Note: 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 strip 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 strips 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. |

| Preset | Purpose | RAID level | Drives per array goal | Spare goal | Chain balance |
|--------|--|---------------|--------------------------------|---------------|--|
| RAID 0 | Provides no protection against drive failures. | 0 | 8 | 0 | All drives in the array are from the same chain wherever possible. |

Usage of the Recommended Configuration

As shown in Figure 7-11 when you select Use the recommended configuration, the wizard offers a recommended storage configuration at the bottom of the window.



Figure 7-11 Recommended configuration

Here are the recommended RAID presets for different drive classes:

- 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 as needed to meet the spare goals of the presets. Under automatic creation, one spare drive will be created out of every 24 disk drives with the same drive class on a single chain.

For example, if you have 20 x 450GB 10K SAS drives on one chain, one drive in these 20 drives will be marked as a spare drive; if you have 20 x 450GB 10K SAS drives on both

chains, which means 10 drives in each chain, then one spare drive on each chain will be created. So if you have 40 x 450GB 10K SAS drives on both chains, then two spare drives on each chain will be created and you have a total of 36 drives that can be the members for the RAID setup.

Spare drives in the IBM Storwize V7000 are global spares. This means that any spare which is at least as big as the drive which is being replaced can be used in an array. Thus an SSD array with no SSD spare available would use an HDD spare instead.

In our example, after the recommended configuration, three arrays using the basic RAID5 presets are offered with one host spare. If the proposed configuration meets your requirements click **Finish** and the system automatically creates array MDisks with the size equivalent to RAIDs, as shown in Figure 7-13.

Storage pools are also automatically created to contain the MDisks with similar performance characteristics, including the consideration of RAID type, number of member drives, drive class and so on.

When an array is created, the array members are synchronized with each other by a background initialization process. You can monitor the progress of the initialization process by clicking the status indicator of **Running Tasks**, as shown in Figure 7-12. The array is available for I/O during this process, and initialization has no impact on availability due to member drive failures.

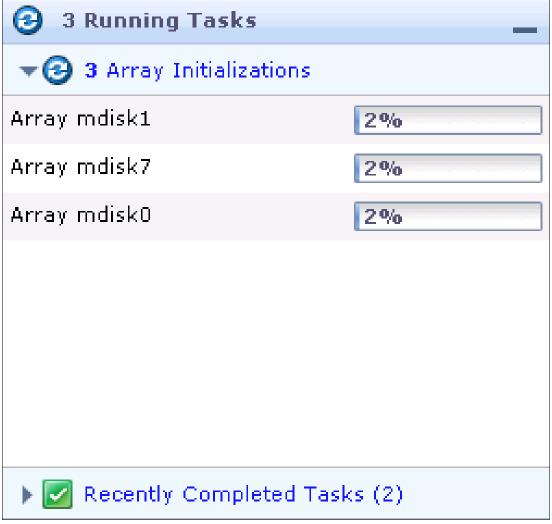


Figure 7-12 Array initialization

The capacity allocation indicator shows the allocation capacity has reached 96% after configuration as shown in Figure 7-13 on page 231.

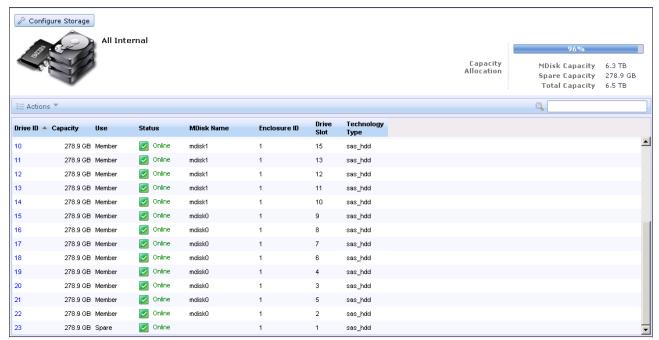


Figure 7-13 Configuration complete with recommended configuration

If the recommended configuration is not what you need, choose **Select a different configuration** and continue with the more flexible setup as shown in "Select a Different Configuration" on page 231.

Select a Different Configuration

The "different configuration" provides more flexibility during the setup process than the "recommended configuration" preset. Multiple options for drive, RAID level, and pool selection are available. Click **Select a different configuration** and choose the drive class you want to configure as shown in Figure 7-14 on page 232.

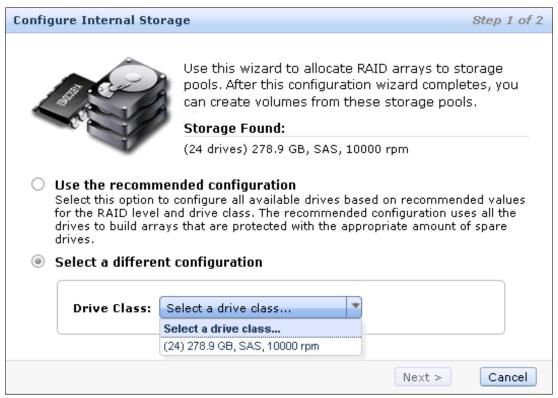


Figure 7-14 Select a drive class

Click **Next** and select an appropriate RAID preset as shown in Figure 7-15 on page 233.

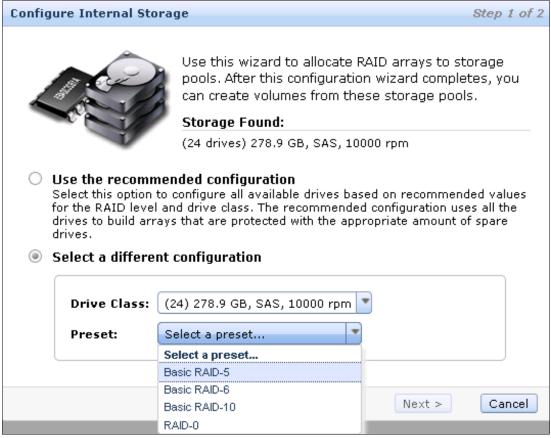


Figure 7-15 Select RAID Preset

Choose RAID5 for example, and click **Next**. Then you can select the number of drives to provision in the configuration and decide if you want the wizard to automatically configure hot spare drives for you. Although the usage of this policy is optional, spare drives can be configured manually as well.

Furthermore here, you have two options as to how to configure your storage:

1. Performance optimized setup:

The goal of this algorithm is to create pools of identical arrays.

In a performance optimized setup the IBM Storwize V7000 uses always 8 physical disk drives in a single array, except in the following situations:

- RAID 6 uses 12 disk drives
- ► SSD Easy Tier uses 2 disk drives

As a consequence all arrays with similar physical disks provide the same performance. The remaining disks can be used in a different array. Figure 7-16 on page 234 shows an example of the performance optimized setup.

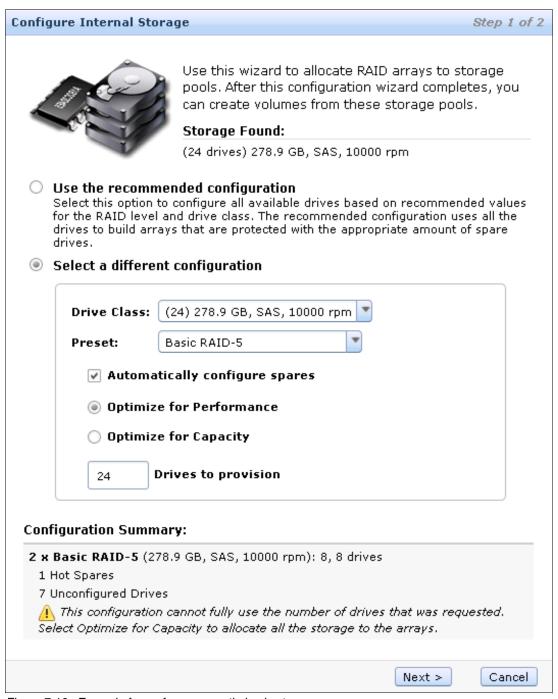


Figure 7-16 Example for performance optimized setup

Note: If the performance optimized configuration cannot fully use the number of drives you choose to provision, you need to configure the remaining unconfigured drives separately.

Capacity optimized setup

The goal of this algorithm is to create a setup with the maximum usable capacity, depending on the selected RAID level.

The IBM Storwize V7000 system tries to achieve the width goal for each RAID array before it creates a new one. 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, 0 | 8 disks |
| RAID 6 | 12 disks |
| SSD Easy Tier | 2 disks |

All disks are used in the capacity optimized setup. There are no "unconfigured devices" as in a performance optimized setup.

Figure 7-17 shows an example of the capacity optimized setup.

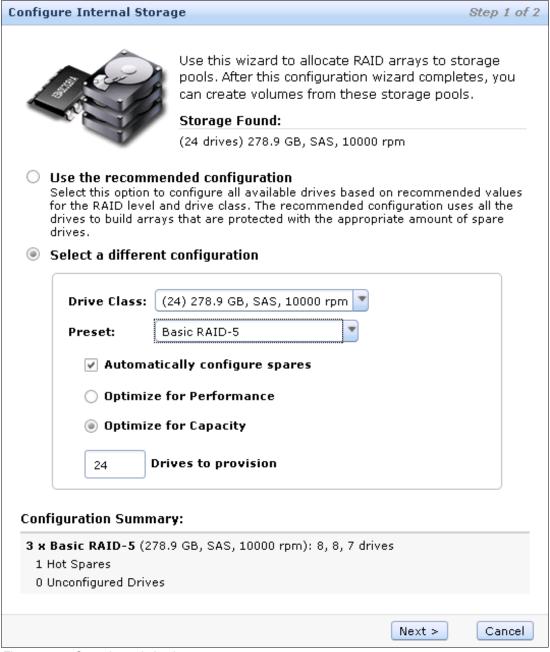


Figure 7-17 Capacity optimized setup

After you have decided on the setup to be applied to your configuration, you can proceed to the next step.

Note: You can select a different number of drives to provision. The wizard will propose that it creates arrays from all candidate drives. But if you want to configure only a subset of drives at this moment in time, then you can change the number in the Drives to provision box to a lower number, and the wizard will apply the policy accordingly.

In this example, we choose the performance optimized setup as shown in Figure 7-16 on page 234 and click **Next**.

In the next step you need to choose the storage pool to assign the new capacity to. You can select an existing storage pool with no MDisk(s) in it or similar performance characteristics which has been listed automatically by the IBM Storwize V7000 as shown in Figure 7-18 on page 237, or create a new storage pool as shown in Figure 7-19 on page 238.



Figure 7-18 Assign new capacity to an existing storage pool

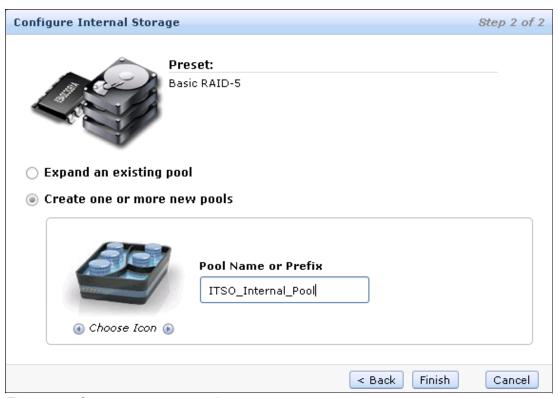


Figure 7-19 Create a new storage pool

Click **Finish** to finalize the wizard. After the wizard completes, our configuration can be found on the right side of the **Internal** panel, as shown in Figure 7-20 on page 239.

With the performance optimized setup, two MDisks with the size equivalent to arrays have been created. The storage pool has also been created if you chose to create a new storage pool to contain the new internal capacity. Seven unconfigured drives can also be found in the drive list for further actions. You will find that the array initialization process is going on in the status indicator of **Running Tasks**, and the capacity allocation indicator has been updated to 67% as well.

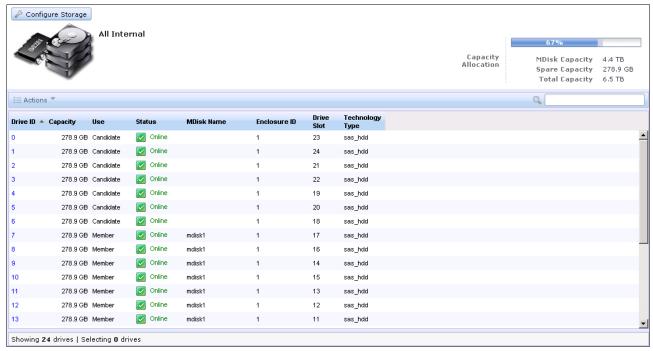


Figure 7-20 Configuration completes with performance optimized setup

7.2 Work with MDisks

After the configuration completes for the internal storage, you can find the MDisks created on the internal arrays in the **MDisks** panel.

You can access the **MDisks** panel through the **Getting Started** panel, and click the **MDisks** icon. The extended help information for MDisks will show up below, click on **Physical Storage** and you will be taken to the **MDisks** panel.

Figure 7-21 on page 240 shows how to access the **MDisks** panel from the **Getting Started** panel.

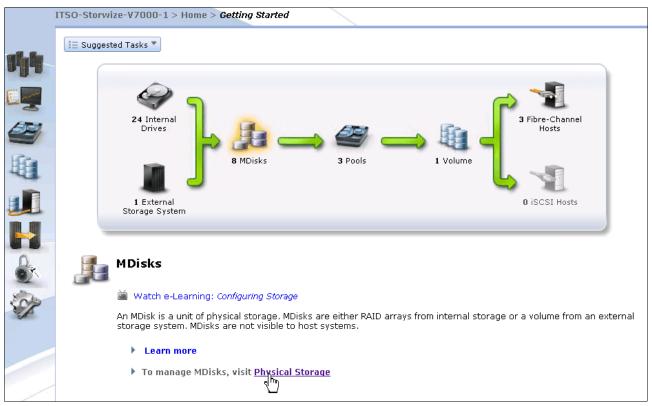


Figure 7-21 Access the MDisks panel from Getting Started panel

Another way to access the **MDisks** panel will be from the **Physical Storage** functional icons on the left hand side.

Figure 7-22 on page 241 shows how to access the **MDisks** panel from the **Physical Storage** functional icons on the left hand side.



Figure 7-22 Access MDisks panel from Physical Storage function icon

The **MDisks** panel, as shown in Figure 7-23 on page 242, provides you with easy access to manage all your MDisks. All MDisks are listed in the MDisks panel including internal and external storage, with general information including name, status, capacity, mode, storage pool it belongs, storage system it comes from for external LUNs as well as its LUN ID, and tier assigned.

You can find more information on how to attach external storage to IBM Storwize V7000 storage system in << refer to external storage>>.

In IBM Storwize V7000 there are four modes for MDisks:

Array:

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

▶ Unmanaged:

The MDisk is not a member of any storage pools, which means it has not been actually used by the IBM Storwize V7000 storage system. The LUNs presented by external storage to IBM Storwize V7000 will be discovered as unmanaged MDisks.

Managed:

The MDisk is assigned to a storage pool and provides extents that volumes can use.

► Image:

The MDisk is assigned directly to a volume with a one-to-one mapping of extents between the MDisk and the volume.

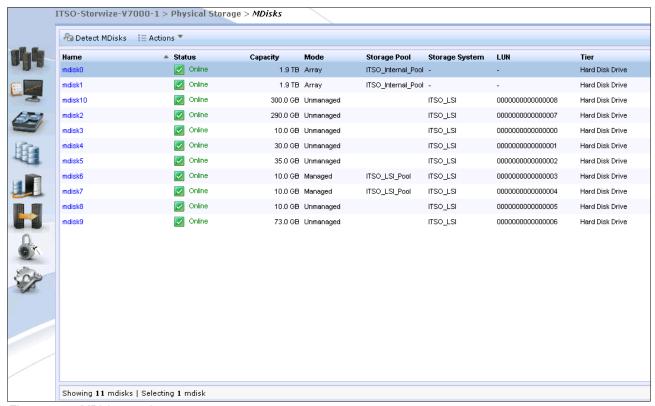


Figure 7-23 MDisks panel

7.2.1 Add MDisks to Storage Pools

The unmanaged MDisks could be changed to managed MDisks by selecting the **Add to Pool** action. You can select an unmanaged MDisk, choose **Add to Pool** in the right click menu or in the **Actions** drop down list, as shown in Figure 7-24.

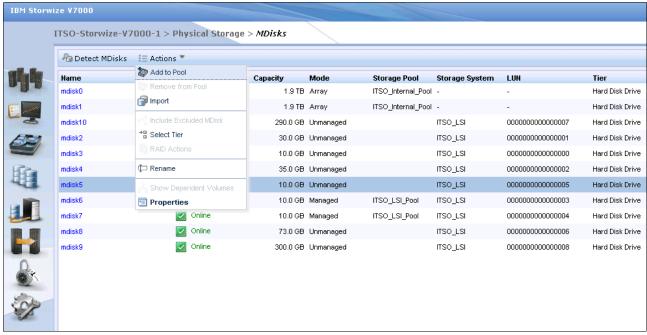


Figure 7-24 Add unmanaged MDisk to storage pool

Note: If there is existing data on the unmanaged MDisks you need to preserve, *do not* select **Add to pool** on this LUN as this will destroy the data. Use **Import** instead, which is described in 7.2.2, "Import MDisks" on page 246.

In the next step, choose the storage pool you want to add the MDisk to, and click **Add to Pool** on the bottom as shown in Figure 7-25.



Figure 7-25 Choose the storage pool to add the MDisk to

After the IBM Storwize V7000 system completes the action, you will see that the MDisk is in the pool you selected as shown in Figure 7-26.

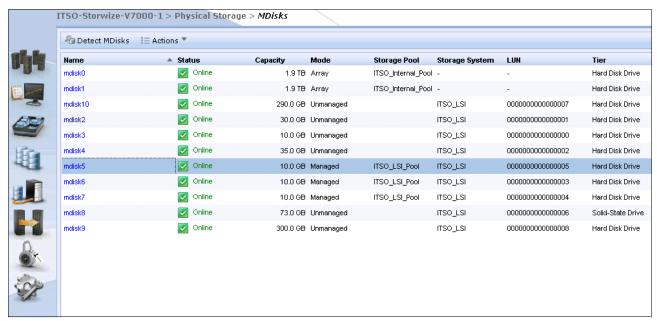


Figure 7-26 MDisk in the pool

In some cases, you may need to remove MDisks from storage pools to reorganize your storage allocation. You can remove MDisks from storage pools by selecting the MDisks and choosing **Remove from Pool** from the right click menu or the **Actions** drop down list, as shown Figure 7-27.

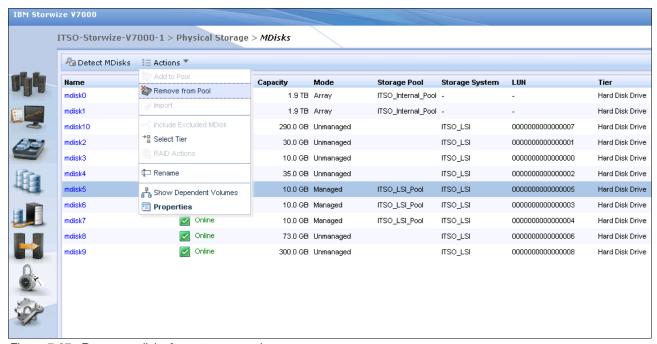


Figure 7-27 Remove mdisks from storage pool

The next step to remove MDisks from storage pool is shown in Figure 7-28 on page 245. You need to confirm the number of MDisks you need to remove. If you have data on the MDisks, and you still need to remove the MDisks from the pool, check the **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** box.

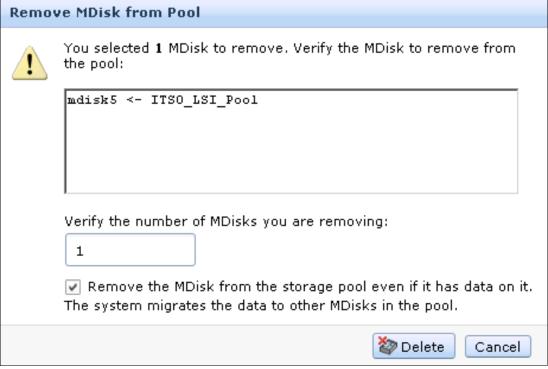


Figure 7-28 Confirm to remove MDisks from the pool

Note: Make sure you have enough available capacity left in the storage pool for the data on the MDisks to be removed

After you click **Delete**, data migration off the MDisks to be removed starts. You can find the migration progress in the status indicator of Running Tasks, as shown in Figure 7-29 on page 246.

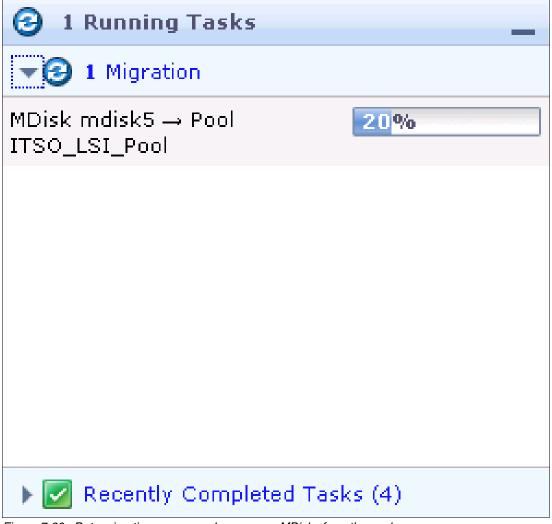


Figure 7-29 Data migration progress when remove MDisks from the pool

7.2.2 Import MDisks

You can also bring the unmanaged MDisks into production by importing them if there is existing data on the MDisks. Using the **Import** option and the IBM Storwize V7000 will create a default pool to put it in, and will create an Image mode volume the same size as the MDisk and preserve the data. After import, the MDisk will be in image mode. Furthermore, it will optionally allow you to select a storage pool to migrate the existing data to.

Image mode is the only mode that preserves the data on the MDisks. As mentioned in 7.2.1, "Add MDisks to Storage Pools" on page 242, if you add an MDisk that contains existing data to a storage pool, you lose the data that it contains.

Select an unmanaged MDisk, choose **Import** in the right click menu or in the **Actions** drop down list, as shown in Figure 7-30 on page 247.

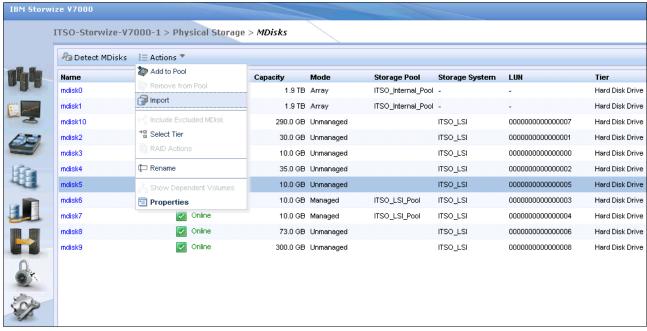


Figure 7-30 Import MDisks

As shown in Figure 7-31 on page 248, the import wizard will be activated to guide the import process.

In the first step of the import wizard, there are two checkboxes which apply to two special scenarios:

Tick the first checkbox only if the volume that you are importing was originally a thin-provisioned volume that was exported to an image-mode volume on the cluster. You may find more information on export volumes in Chapter 8, "Advanced Host and Volume Administration" on page 273.

The second checkbox is ticked by default, which means the IBM Storwize V7000 will enable caching for image mode volumes. Although you can remove the tick on the second checkbox to disable caching if you are using copy services on the external storage to maintain the data consistency to the application server between the local and remote site on the external disk controller level. It is recommended to use the copy services of IBM Storwize V7000 for disaster recovery in the virtualized environment under it. You can find more information about virtualizing external storage in Chapter 9, "External Storage Virtualization" on page 339

Make your choices and click Next.

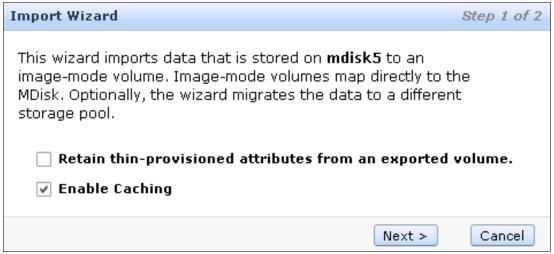


Figure 7-31 First step in Import Wizard

In the next step you can choose the destination storage pool if you want to migrate the data on image mode MDisks. If you select the destination pool here and click **Finish**, as shown in Figure 7-32, the migration will begin after the import of the MDisks.

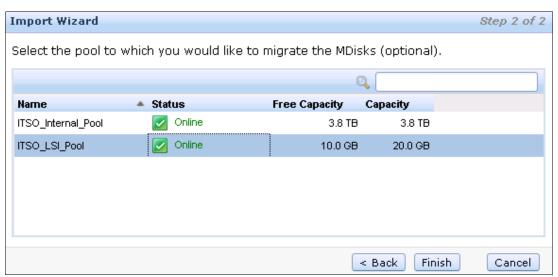


Figure 7-32 Select migration destination pool

You can find the migration progress in the status indicator of Running Tasks, as shown in Figure 7-33 on page 249.

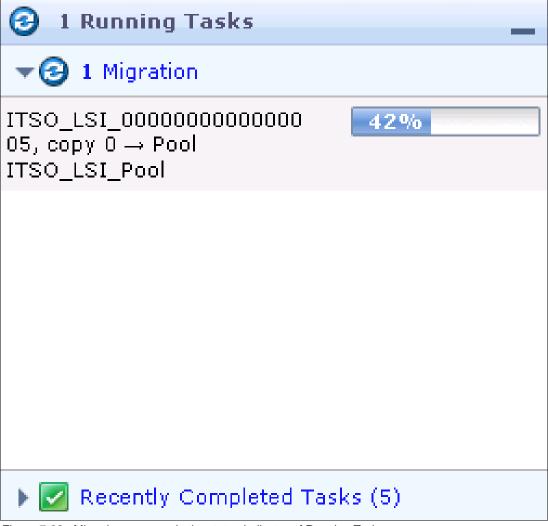


Figure 7-33 Migration progress in the status indicator of Running Tasks

The migration progress can also be found in the **Migration** panel as shown in Figure 7-34 on page 250.

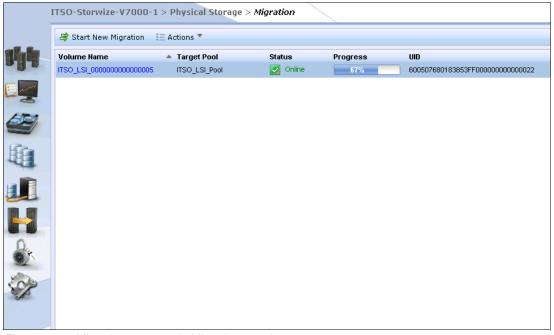


Figure 7-34 Migration progress in Migration panel

After the migration you will find the volume in the destination pool, as shown in Figure 7-35.

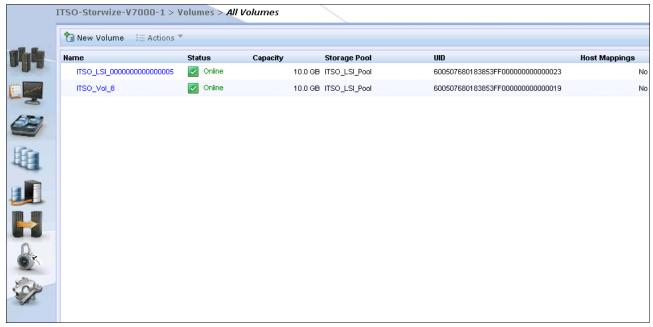


Figure 7-35 Volume migrated to destination pool

All the data has now been migrated from the source MDisk to the target storage pool. With no data on it, the source MDisk has been changed to managed mode, and could be used to serve other volumes, as shown in Figure 7-36 on page 251.

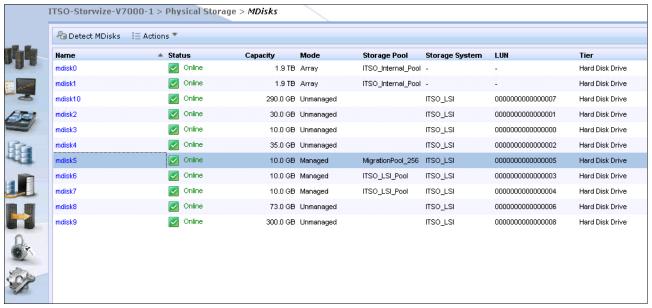


Figure 7-36 MDisk mode change to managed

On the other hand, you can leave nothing selected in the second step in import wizard and then click **Finish**, as shown in Figure 7-37.

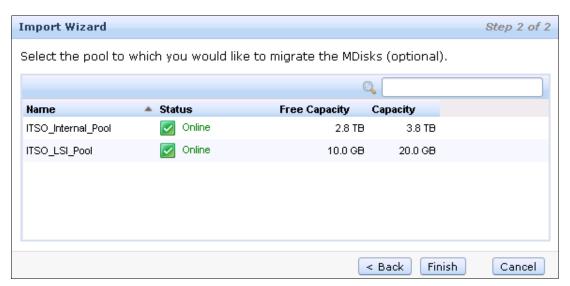


Figure 7-37 Nothing selected in second step in Import Wizard

A warning message will appear after you click **Finish** without select anything, as shown in Figure 7-38 on page 252. Click **OK** to proceed.

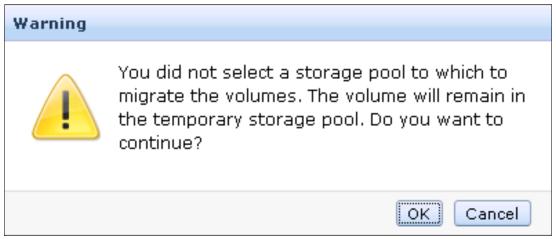


Figure 7-38 Warning message to import without migration

After import, you can see that the MDisk mode has changed to image mode as shown in Figure 7-39.

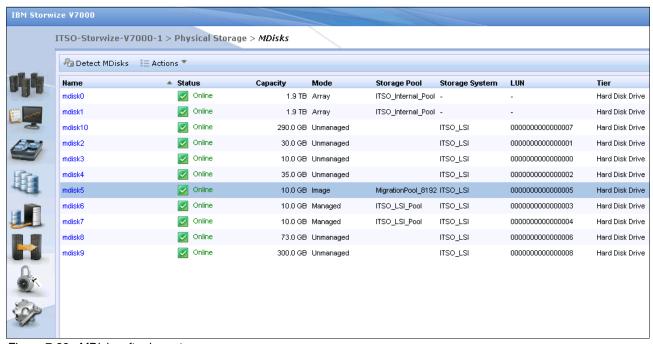


Figure 7-39 MDisks after import

Without migration, the related image mode volume will remain in the temporary storage pool, which is automatically created by IBM Storwize V7000, as shown in Figure 7-40 on page 253.

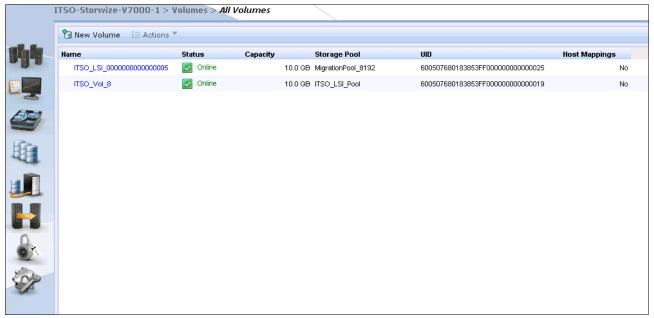


Figure 7-40 Volume after import

You can manually start your migration from the image mode volume through the volume actions **Migration to Another Pool** or **Volume Copy Actions**. More information can be found about volume actions in Chapter 5, "Basic Volume Configuration" on page 157.

You can also migrate the data on imported MDisks to other storage pools in the **Migration** panel. How to access the Migration panel could be found in Chapter 6, "Migration Wizard" on page 199.

If you have already imported MDisks from external storage systems without data migration, you will find it in the migration panel as shown in Figure 7-41 on page 253.

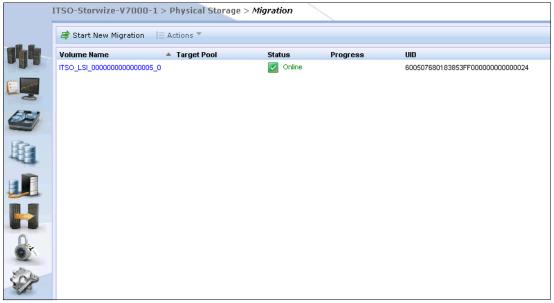


Figure 7-41 Imported MDisk in Migration panel

The candidate volumes for migration in the list, are directly mapped to the image mode MDisks you have imported. Select the volume you want to migrate and click 'Start Migration' in the 'Actions' drop down menu, as shown in Figure 7-42.

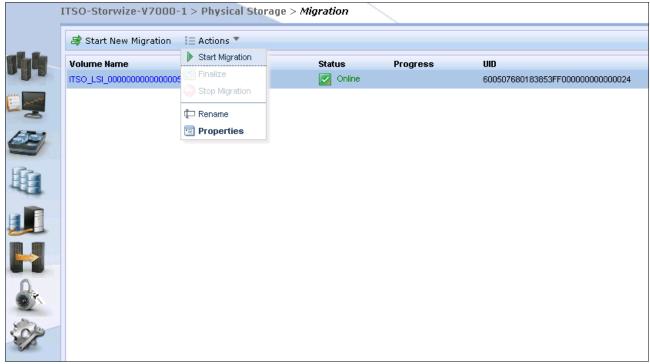


Figure 7-42 Start migration

After you start the migration, IBM Storwize V7000 will let you choose the destination storage pool. Choose the target storage pool you want to migrate the data from the volumes to and click **Add to Pool**, as shown in Figure 7-43 on page 254.

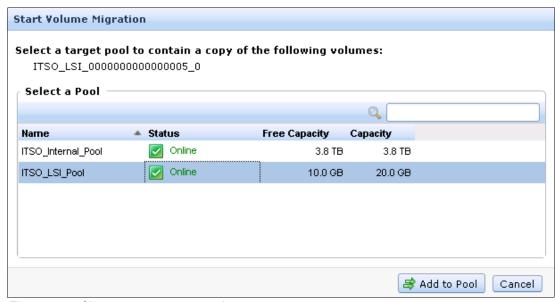


Figure 7-43 Choose target storage pool

Now the data migration has begun, and you can find the migration progress in the migration panel, as shown in Figure 7-44.

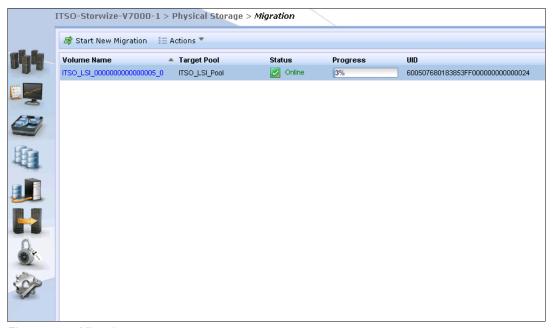


Figure 7-44 Migration starts

Since the migration starts in the **Migration** panel using the volume copy function, you can also find the progress of volume copy synchronization in the status indicator of **Running Tasks**, as shown in Figure 7-45 on page 256. You can find more explanation on the volume copy function in 8.6, "Advanced Volume Copy Functions" on page 324.

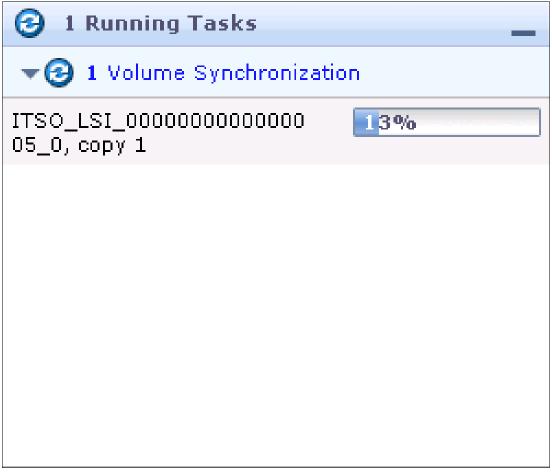


Figure 7-45 Volume synchronization progress for migration

As described in Chapter 6, "Migration Wizard" on page 199, when the migration progress reaches 100%, click 'Finalize' to complete the migration process, which will delete the copy on the image mode MDisk and keep the other one. The source MDisk which was in image mode originally will be removed from its storage pool and change its mode to unmanaged as shown in Figure 7-46 on page 256.

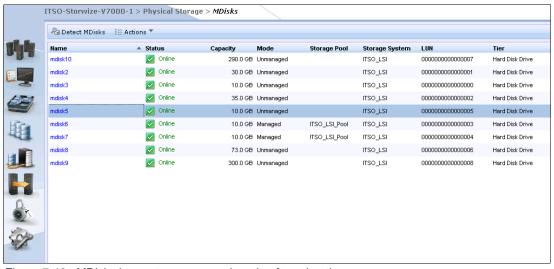


Figure 7-46 MDisk change to unmanaged mode after migration

In the **Migration** panel, you can also stop the migration progress after it starts by clicking **Stop Migration** in the **Actions** drop down list as shown in Figure 7-47. After the migration is forced to stop, the second copy of the data will be removed. Start the migration again if needed.

Figure 7-47 shows how to stop the data migration progress.

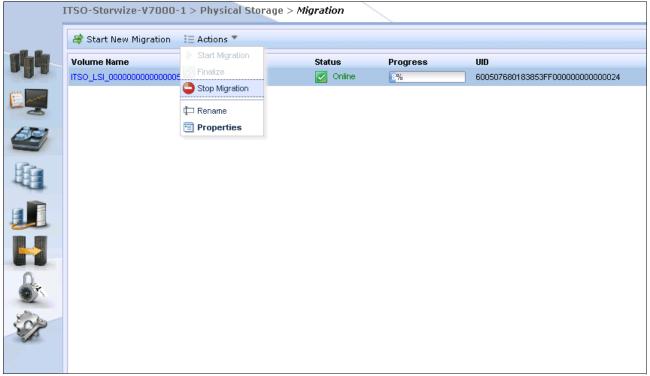


Figure 7-47 Stop migration

The IBM Storwize V7000 will need your verification and confirmation on the migrations you choose to stop. Verify the volume names in the migrations and the number of migrations you want to stop, and click 'OK' to confirm your decision.

Figure 7-48 shows the verification window for stopping migrations.

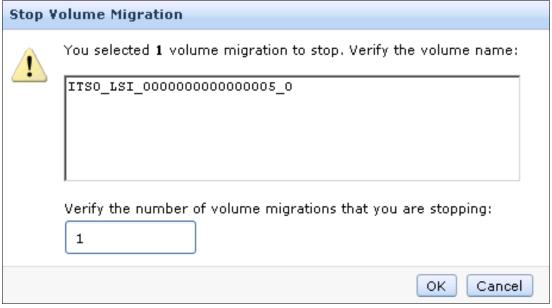


Figure 7-48 Verify the migrations you are stopping

7.2.3 RAID Action for MDisks

With the array mode MDisks, you can perform several RAID actions by selecting the MDisk and choosing **RAID Actions** in the right click menu or in the **Actions** drop down list, as shown in Figure 7-49.

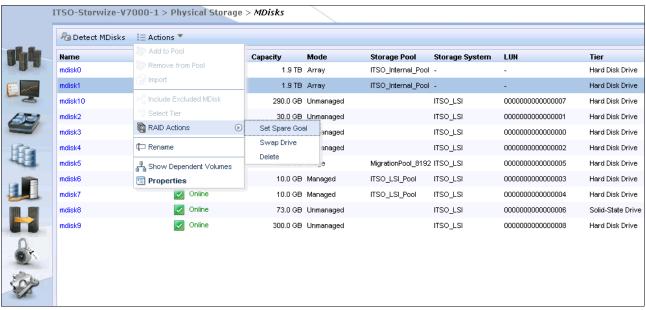


Figure 7-49 RAID Actions

You can set the spare goal of the MDisk by selecting **Set Spare Goal** in the **RAID Actions** menu. Then input your new spare goal for the MDisk as shown in Figure 7-50 on page 259.



Figure 7-50 Set Spare Goal

After you click **Save**, the IBM Storwize V7000 will remember your goal for spare drives. If your current spare drives do not meet your goal for the MDisk, you will receive a warning in the event log that says 'Array mdisk is not protected by sufficient spares'. Manually mark the drives as spares as needed as described in 7.1.1, "Actions on Internal Drives".

Swap Drive action could be used to replace drives in the array. Select an MDisk and choose **Swap Drive** in the **RAID Actions** menu. You need to select a drive to swap out of the MDisk first as shown in Figure 7-51, and click **Next**.

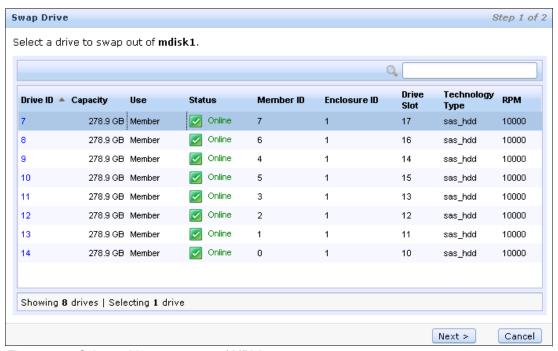


Figure 7-51 Select a drive to swap out of MDisk

Then you need to select a drive to swap into the MDisk. The IBM Storwize V7000 will list the drives with candidate and spare mode for selecting to swap in as shown in Figure 7-52 on page 260.

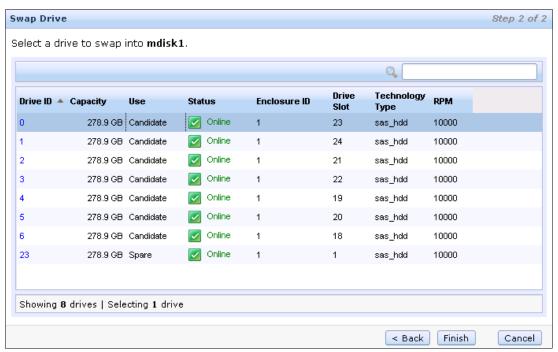


Figure 7-52 Select a drive to swap into MDisk

After you click Finish, IBM Storwize V7000 will start exchanging the array members in background. When it is done, you will find the members of the array have been changed. The swap-out drive will keep the original mode of the one swapped in.

If the array MDisks have reached the end of their lifecycle, they can be deleted by **Select the MDisks** and choose **Delete** in the **RAID Actions** menu. Then IBM Storwize V7000 will ask you to confirm your delete action as shown in Figure 7-53 on page 261. You need to confirm the number of array MDisks you want to delete. If you have data on the MDisks and you still need to delete them select **Delete the RAID array MDisk even if it has data on it. The system migrates the data to other MDisks in the pool.**

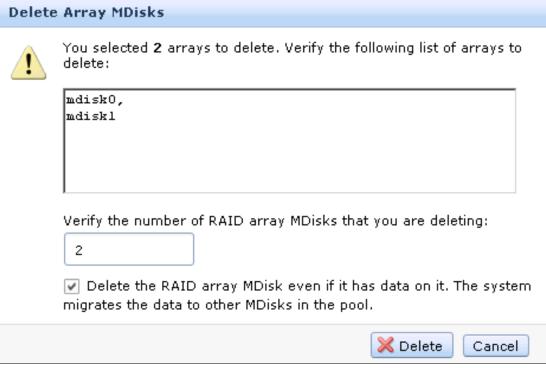


Figure 7-53 Delete Array MDisks

Note: Make sure you have enough available capacity left in the storage pool for the data on the MDisks to be removed.

After you click **Delete**, all the member drives of the MDisks would return to candidate mode, and the data on the MDisks would be migrated to other MDisks in the pool.

7.2.4 Select the tier for MDisks

IBM Storwize V7000 defines two tiers:

- Generic SSD tier for storage made of solid-state drives, which is the faster-performing storage.
- ► Generic HDD tier for everything else.

While the Internal drives will have their tier chosen automatically by IBM Storwize V7000, the tier of external MDisks will be assigned to the generic HDD tier by default and can be manually changed by the user.

To assign a specific tier to one MDisk, select it and choose **Select Tier** in the right click menu or in the **Actions** drop down list as shown in Figure 7-54 on page 262.

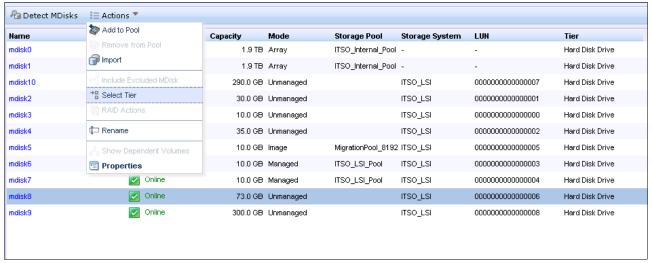


Figure 7-54 Select Tier for MDisk

As the MDisk we select in this example has been assigned to the HDD tier by default, we select the desired tier for this MDisk to be SSD, and **click OK**, as shown in Figure 7-55.



Figure 7-55 Select desired tier to be assigned to MDisk

After the action has been completed successfully, the MDisk can be found in the tier SSD as shown in Figure 7-56.



Figure 7-56 Desired tier has been assigned to MDisk

7.2.5 Additional Actions on MDisks

The **Detect MDisk** button on the top left corner of the **MDisks** panel, is useful if you have external storage controllers in your environment, refer to Chapter 9, "External Storage Virtualization" on page 339 for more information. The **Detect MDisk** action manually rescans the Fibre Channel network, and discovers any new MDisks that have been 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, as described in Chapter 9, "External Storage Virtualization" on page 339, generally the IBM Storwize V7000 will automatically discover the controllers and the LUNs that are presented by the controllers are displayed as unmanaged MDisks. However, if you have attached new storage and the IBM Storwize V7000 has not detected it, you might need to run this command before the cluster detects the new LUNs. And if the configuration of the external controllers are modified afterwards, the IBM Storwize V7000 might be unaware of these configuration changes. Use this action to rescan the Fibre Channel network and update the list of unmanaged MDisks.

Note: Detect MDisk is asynchronous and returns a prompt while it continues to run in the background.

MDisks could be excluded from the IBM Storwize V7000 because of multiple I/O failures. These failures might be caused by link errors. Once a fabric-related problem has been fixed, the excluded disk can be added back into the IBM Storwize V7000 by selecting the MDisks and choosing **Include Excluded MDisk** in the right click menu or in the **Actions** drop down list.

The MDisk can be renamed easily by selecting the MDisk and choosing **Rename** in the right click menu, or in the **Actions** drop down list. Input the new name of your MDisk and click **Rename**, as shown in Figure 7-57.



Figure 7-57 Rename MDisk

The volumes dependent on an MDisk can be shown by selecting the MDisk and choosing **Show Dependent Volumes** in the right click menu, or in the **Actions** drop down list. The volumes will be listed with general information as shown in Figure 7-58 on page 264.

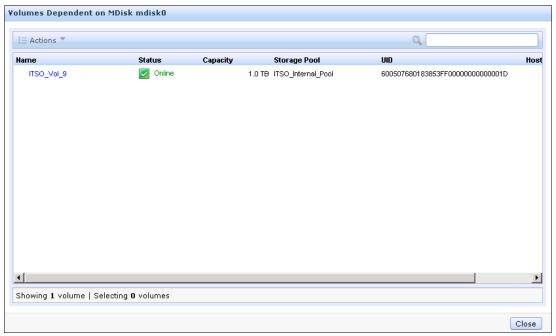


Figure 7-58 Show dependent volumes

Several volume actions could be taken on the volumes by selecting the volume and choosing the action needed in the right click menu or in the **Actions** drop down list as shown in Figure 7-59. More information could be found on volume actions in Chapter 5, "Basic Volume Configuration" on page 157.

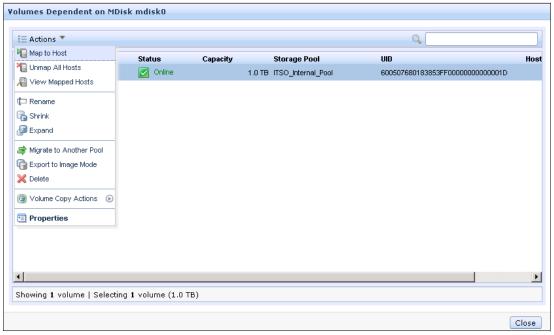


Figure 7-59 Available actions on volumes

The **Properties** action on the MDisk will show the information you need to identify it. On the **MDisks** panel, select the MDisk and choose **Properties** in the right click menu or in the **Actions** drop down list, and you will get a pop-up window that displays the information as shown in Figure 7-60 on page 265.

There are three tabs in this information window. The **Overview** tab contains information about the MDisk itself.

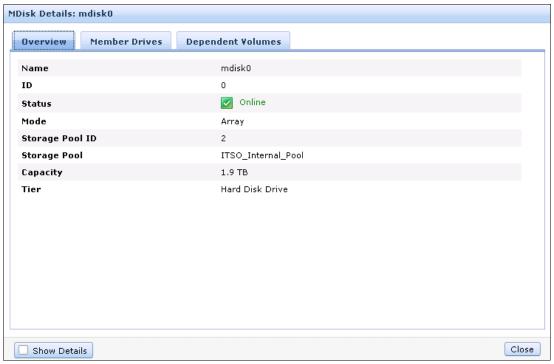


Figure 7-60 MDisk overview with general information

With the default format in the **Overview** tab, there is only general information displayed. Choose Show Details to get more detailed information about the MDisk, as shown in Figure 7-61.



Figure 7-61 MDisk overview with detailed information

The **Member Drives** tab is only for array MDisks. In the **Member Drives** tab you will find all the member drives of this MDisk displayed as shown in Figure 7-62. You can also take required actions here as described in 7.1.1, "Actions on Internal Drives".

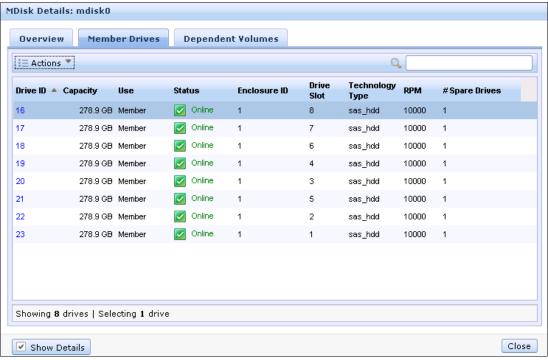


Figure 7-62 Member Drives tab

The last tab is the **Dependent Volumes** tab and the content is the same as described in the previous **Show Dependent Volumes** action.

7.3 Work with Storage Pools

Storage Pools act as a container for MDisks and provision the capacity to volumes. IBM Storwize V7000 organizes storage with storage pools to make it able to be managed easily and efficiently. An independent panel is there for storage pool management, called **Pools**.

You can access the **Pools** panel through **Getting Started** panel, and click the **Pools** icon. Then, extended help information for storage pools will show up below and if you click **Physical Storage** you will go to the **Pools** panel.

Figure 7-63 on page 267 shows how to access the **Pools** panel from the **Getting Started** panel.

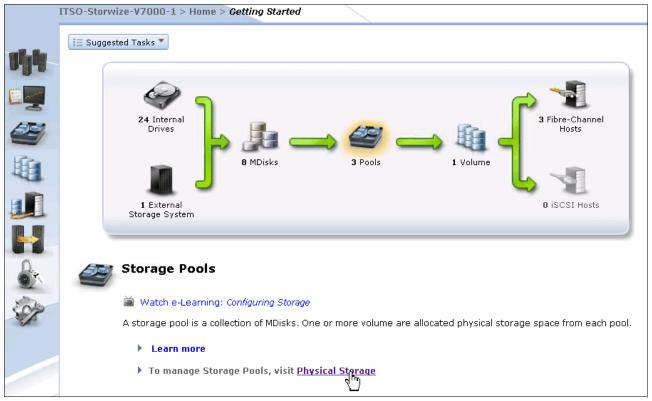


Figure 7-63 Access the Pools panel from Getting Started panel

The other way to access the **Pools** panel will be from the **Physical Storage** functional icons in the left.

Figure 7-64 on page 268 shows how to access the **Pools** panel from the **Physical Storage** functional icons in the left.



Figure 7-64 Access Pools panel from Physical Storage function icon

The **Pools** panel, as shown Figure 7-65 on page 268, is where you can manage all your storage pools. On the left side of the **Pools** panel you can find your storage pools listed. With the help of the filter you can list only the storage pools you need. Choose one storage pool from the left of the **Pools** panel and you will find more information about the pool on the right.

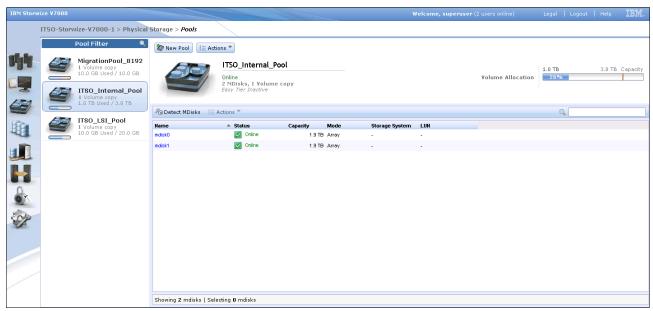


Figure 7-65 Pools panel

The name of the storage pool and its status are displayed beside the storage pool icon. The name of the storage pool could be changed by clicking on the name and entering the new name.

Below the name of the storage pool you will find the status of the storage pool and the summary of MDisks and volumes associated with the pool. The Easy Tier status can also be found here, more information about Easy Tier can be found in Chapter 10, "Easy Tier" on page 353

On the top right hand side of the **Pools** panel you can find a volume allocation indicator which shows the capacity allocated to volumes versus the total capacity. With the percentage calculated in the indicator it is easy to know the utilization of your storage pool which makes management much easier.

The MDisks in the storage pool are listed at the below section on the right with general information, as shown in Figure 7-65 on page 268. The actions that can be taken on the MDisks will not be duplicated here as they are the same as described in 7.2, "Work with MDisks" on page 239.

To create a new storage pool, you can click the **New Pool** button on the top left, then a simple wizard will guide you the process, as shown in Figure 7-66 on page 269. Input a name for your new storage pool, or the name will be assigned automatically by IBM Storwize V7000. Do not forget to choose a icon for your storage pool, it is useful and easy to manage to let the icon of the storage pool represent the feature of MDisks in the pool. Click **Next** when you have made your decision.



Figure 7-66 Step 1 to create a storage pool

In the next step the MDisks that can be added to this storage pool have been listed as shown in Figure 7-67. It is possible to have no MDisks selected which means you are not ready to put any MDisks into the storage pool yet, and an empty storage pool will be created. Click **Finish** to proceed.

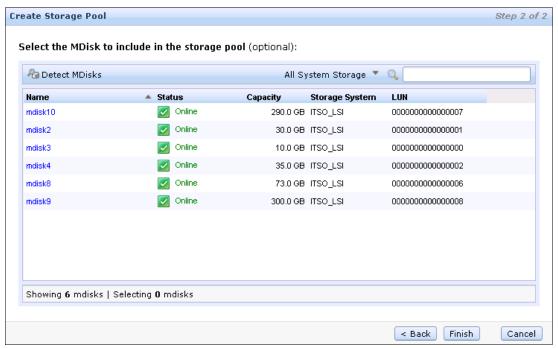


Figure 7-67 Step 2 to create a storage pool

If you choose not to put any MDisks into the pool a warning message will appear to confirm that you wish to create an empty pool as shown in Figure 7-68 on page 270. Click **OK** to continue.



Figure 7-68 Confirm to create an empty pool

After the creation you will find your new pool listed on the left hand side of the **Pools** panel. If the new pool is empty, you can add MDisks to the pool afterwards in the **MDisks** panel as described in 7.2.1, "Add MDisks to Storage Pools".

To remove one storage pool you can select **Delete Pool** from the **Actions** drop down list as shown in Figure 7-69.

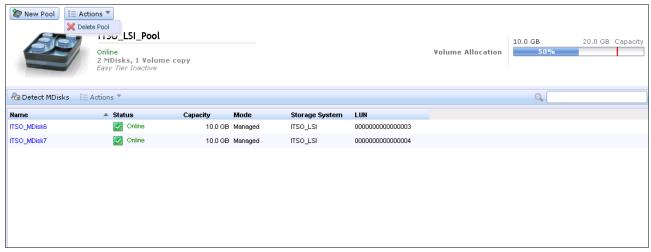


Figure 7-69 Delete Pool

A confirmation window will appear as shown in Figure 7-70 on page 271. The volumes listed in the box are the ones with the only copy in this pool. These volumes will be deleted as well when you delete a storage pool. Check all the volumes listed and make sure they are all required to be deleted. If you have volumes or MDisks in this pool you can proceed only if you have ticked the checkbox **Delete all volumes**, **host mappings**, **and MDisks that are associated with this pool**. Make your selection and confirm you wish to continue by clicking **Next**.

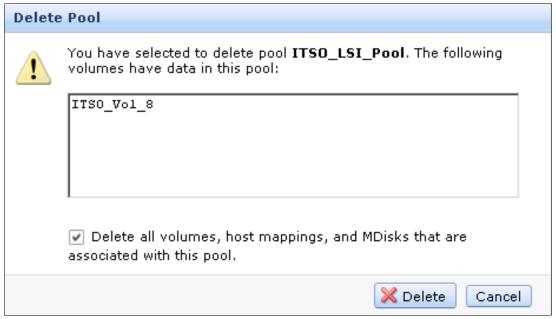


Figure 7-70 Confirm to delete pool

Note: Take care as after you delete the pool, all the data in the pool will be lost except for image mode MDisks.

After you delete the pool all the associated volumes and their host mappings will be removed. All the managed or image mode MDisks in the pool will return to a status of unmanaged after

the pool is deleted. And if the pool is deleted, all the array mode MDisks in the pool will be removed and all the member drives will return to candidate.



8

Advanced Host and Volume Administration

Chapter 4, "Host Configuration" on page 127 and Chapter 5, "Basic Volume Configuration" on page 157 cover the basic host and volumes features of IBM Storwize V7000. They show how to create hosts and volumes, and how to map them to a host. However the IBM Storwize V7000 offers a lot more functions, and this chapter describes all the features that the IBM Storwize V7000 offers you.

This chapter covers all the host and volumes actions, except the basic configuration (covered previously) and EasyTier, which is described in Chapter 10, "Easy Tier" on page 353.

This chapter contains the following:

- ► Advanced Host Administration (8.1, "Advanced Host Administration" on page 274)
- Advanced Volume Administration (8.4, "Advanced Volume Administration" on page 300)

8.1 Advanced Host Administration

This section covers host administration such as host modification, host mappings and deleting hosts. The basic host creation and mapping it to a host using Fibre Channel and iSCSI is shown in 4.2.1, "Creating FC Hosts" on page 147 and 4.2.2, "Creating iSCSI Hosts" on page 150.

We assume that you have already created some hosts in your Storwize V7000 GUI and that some volumes are already mapped to them. In this section we will cover the three functions that are covered in the host section of the IBM Storwize V7000 GUI as shown in Figure 8-1.

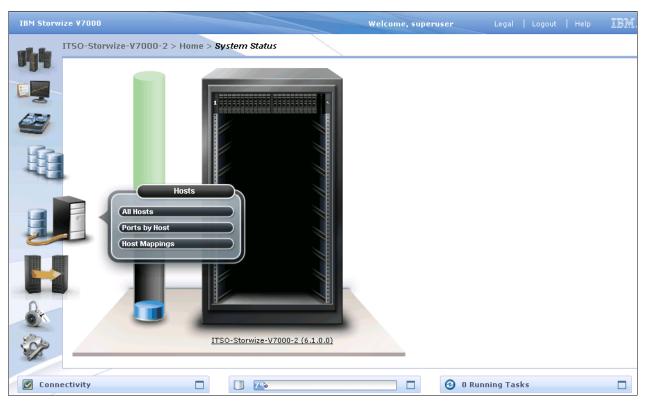


Figure 8-1 IBM Storwize V7000 Host Section

- ► All Hosts (8.1.1, "Modify Mappings" on page 276).
- ▶ Ports by Host (8.2, "Add and delete host ports" on page 289).
- ► Host Mappings (8.3, "Host mappings overview" on page 297).

Go to the **All Host** section as shown in Figure 8-1, and it will appear as shown in Figure 8-2 on page 275.

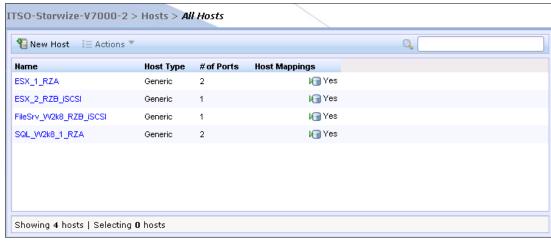


Figure 8-2 All Hosts

As you can see in our example, four hosts are created and volumes are already mapped to them, we will use these hosts to show the modification possibilities. If you highlight a host you can either click **Action** (Figure 8-3) or right click it (Figure 8-4) to see all available tasks.

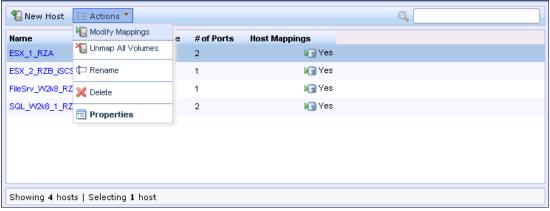


Figure 8-3 Host Action Panel

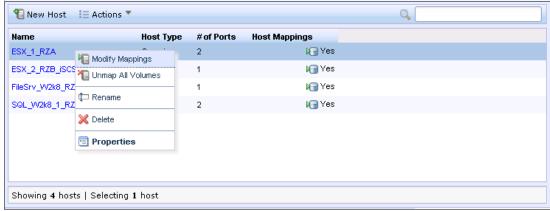


Figure 8-4 Host Right Click Panel

8.1.1 Modify Mappings

Highlight a host and select **Modify Mappings** as shown in Figure 8-3 to open the associated view as shown in (Figure 8-5).

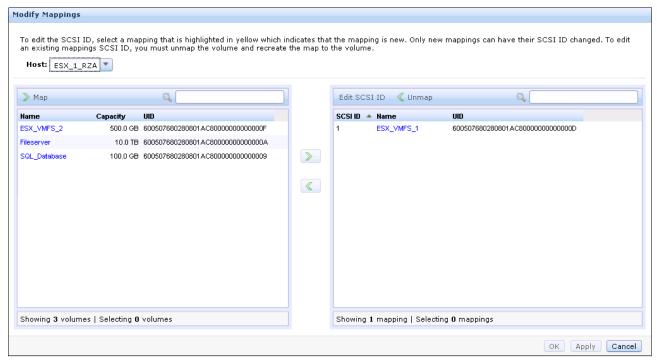


Figure 8-5 Modify Mappings

In the upper left you will see that the highlighted host is selected. The two boxes below show all available volumes. The left box shows the volumes that are ready for mapping to this host, and the right box includes the volumes already mapped. In our example one volume with SCSI ID 1 is mapped to the host, and there are three more available. Highlight a volume in the left box, and select the upper arrow (pointing to the right) to map another volume to a host as shown in (Figure 8-6 on page 277).

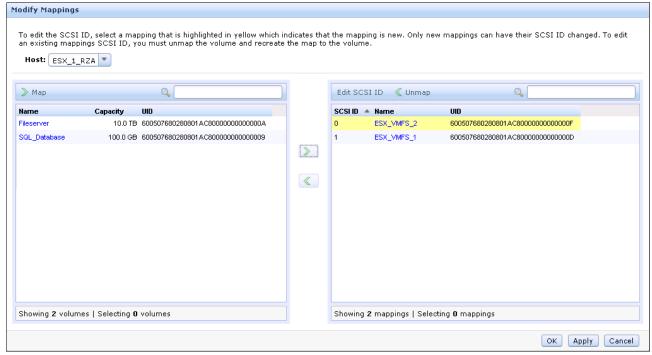


Figure 8-6 Modify Mappings - Add Volume

The changes are marked in yellow and now the **OK** and **Apply** button are enabled. If you click **OK** the changes will be applied (Figure 8-7) and the view will be closed, if you click **Apply** the changes will be submitted to the system (Figure 8-7), but the view will remain open for further changes as shown in (Figure 8-8).

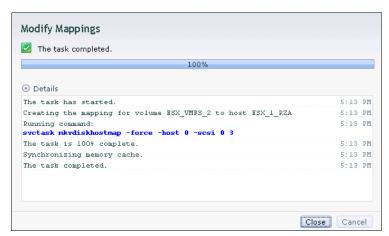


Figure 8-7 Modify Mappings - Changes Applied

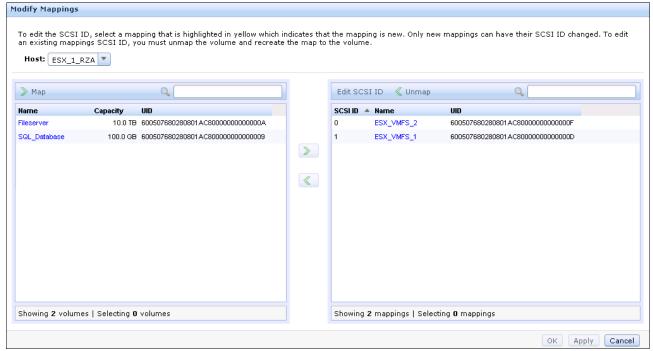


Figure 8-8 Modify Mappings - Applied Changes

You can now select another host in the hosts drop down box as shown in Figure 8-9 to modify the host settings for it.

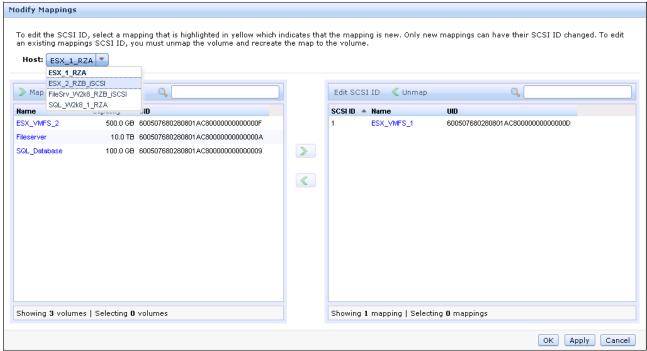


Figure 8-9 Modify another host

Highlight the volume to be modified again and click the right arrow button to move it to the right box. The changes will again be shown in yellow there. If you right click the yellow

volume, you are able to change the SCSI ID, which will be used for the host mapping, as shown in Figure 8-10 on page 279.

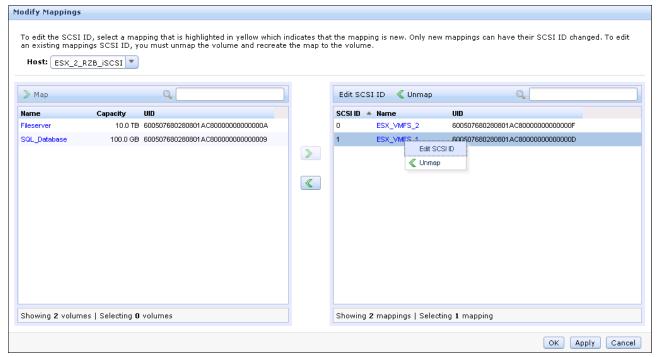


Figure 8-10 Edit SCSI ID

Click Edit SCSI ID and click OK to change the ID as shown in Figure 8-11.



Figure 8-11 Enter new SCSI ID

The changes are displayed in the modify mappings view as shown in Figure 8-12, click **Apply** to submit the changes as shown in Figure 8-13.

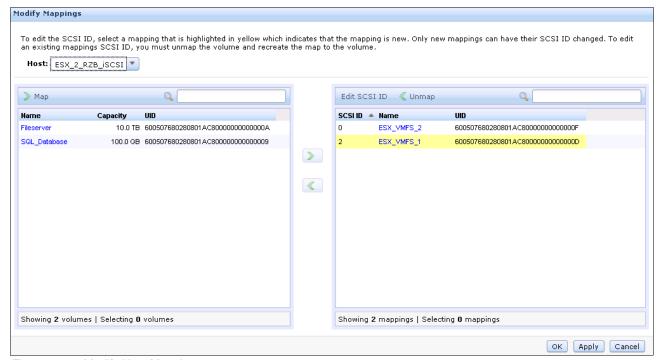


Figure 8-12 Modify Host Mappings

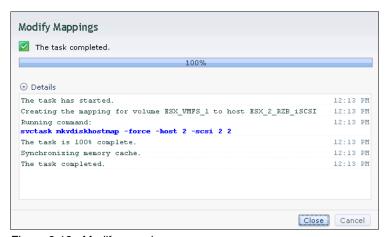


Figure 8-13 Modify mappings

If you would like to remove a host mapping the required steps are pretty much the same, except that you select a volume in the right box and click the left arrow button to remove the mapping again as shown in Figure 8-14 on page 281.

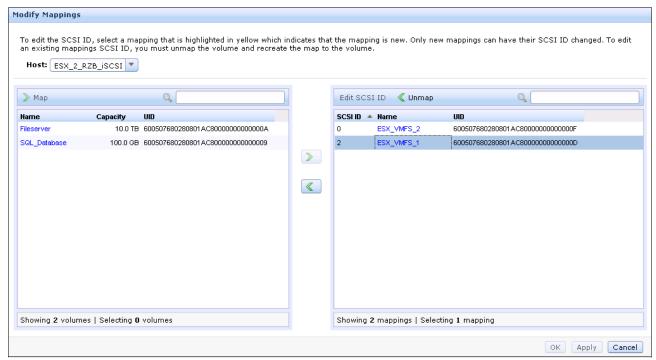


Figure 8-14 Modify Mappings - Remove

Figure 8-15 shows that the selected volume has been moved to the left box, to unmap it.

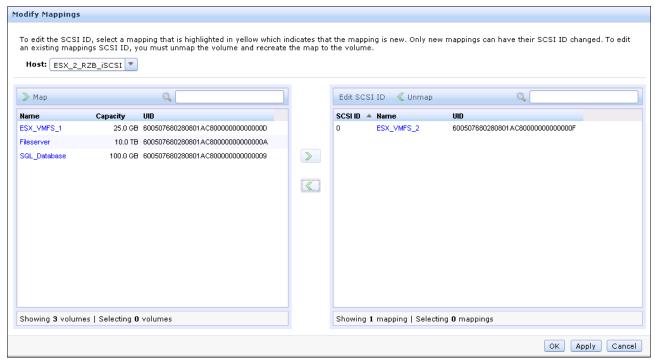


Figure 8-15 Modify mappings - Remove

Click **OK** or **Apply** to submit the changes to the system (Figure 8-16 on page 282).

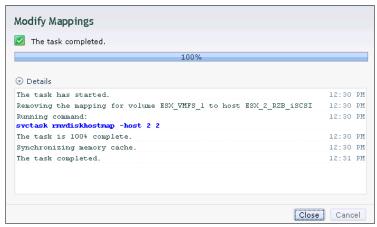


Figure 8-16 Modify Mappings - Removal Complete

Once you are done with all host mapping modifications click **OK** to return to the modify mappings view (Figure 8-5).

8.1.2 Unmap all volumes from a host

A host is only able to access volumes on your IBM Storwize V7000 which are mapped to it. If you want to remove complete access for one host, regardless of how many volumes are mapped to it, you can do this in one step. Highlight the host and click **Unmap All Volumes** to remove all access this host has to its volumes as shown in Figure 8-17.

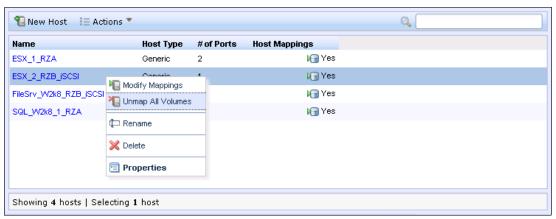


Figure 8-17 Unmap All Volumes

You will be prompted as to the number of mappings you want to remove, enter the number and click **Unmap** as shown in Figure 8-18.

Note: If you click **Unmap** all access for this host to volumes controlled by IBM Storwize V7000 system will be removed, make sure that you run the required procedures in your host operating system before removing the volumes from your host.

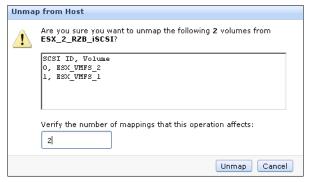


Figure 8-18 Enter number of mappings to be removed

The changes will be applied to the system as shown in Figure 8-19.

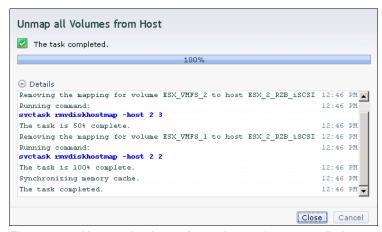


Figure 8-19 Unmap all volumes from a host - changes applied

Figure 8-20 shows that the selected host does not have any host mappings anymore.

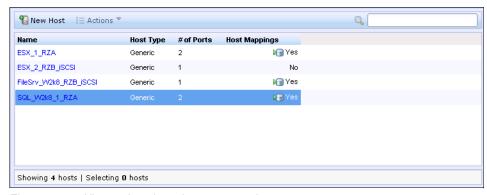


Figure 8-20 All mappings have been removed

8.1.3 Rename

To rename a host highlight it and click **Rename** as shown in Figure 8-21 on page 284.

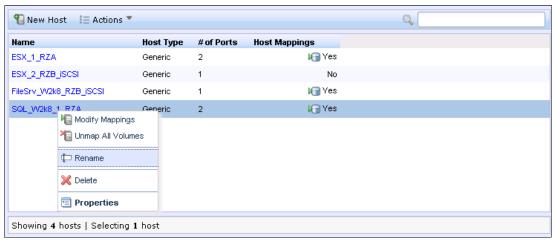


Figure 8-21 Rename Host

Enter a new name and click **Rename** as shown in Figure 8-22. If you click **Reset** the changes will be reset to the primary host name.



Figure 8-22 Rename Host

After the changes have been applied to the system click Close as shown in Figure 8-23.

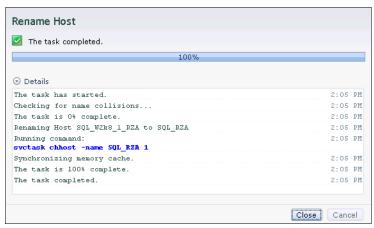


Figure 8-23 Rename Host - Applied Changes

8.1.4 Delete a Host

To delete a host highlight it and click **Delete** as shown in Figure 8-24 on page 285.

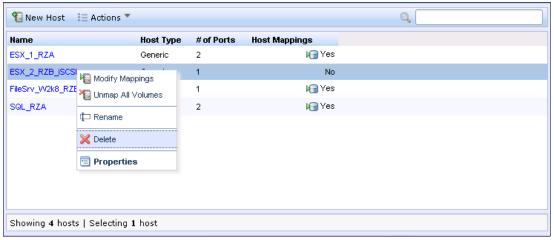


Figure 8-24 Delete a Host

Enter the number of hosts you want to delete and click **Delete** as shown in Figure 8-25.



Figure 8-25 Confirm Delete Host

If you want to delete a host that has volumes mapped, you have to force it by selecting the checkbox in the lower part of the window. If you enable this checkbox the host will be deleted, and it will no longer have access to this system.

After the task has completed click **Close** as shown in Figure 8-26 to return to the mappings view.

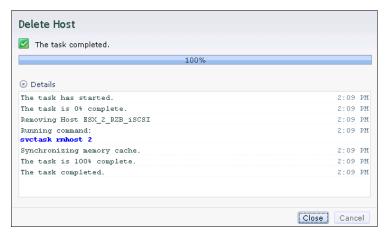


Figure 8-26 Delete Host Complete

8.1.5 Host Properties

Highlight a host and select **Properties** (Figure 8-27) to open the Host Details view as shown in Figure 8-28.

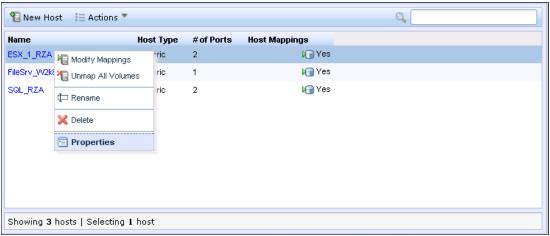


Figure 8-27 Host Properties

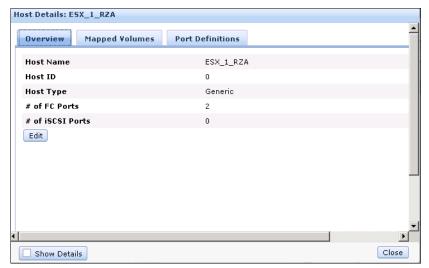


Figure 8-28 Host Properties Overview

The Host Details view gives you a overview about your host properties. There are three tabs available: Overview, Mapped Volumes, and Port Definitions.

The overview tab is shown in Figure 8-28, select the **Show Details checkbox** to get more information of the host as shown in Figure 8-29.



Figure 8-29 Host Properties - Show Details

Click the **Edit** button and you will be able to change the host properties as shown in Figure 8-30.

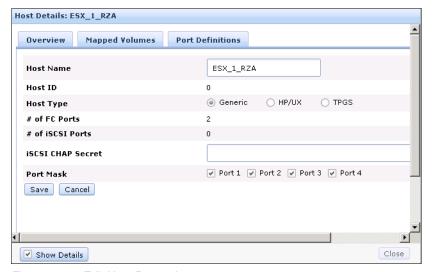


Figure 8-30 Edit Host Properties

It is possible to make the following changes:

- ► Host Name: Change the host name
- ► Host Type: Change this setting if you are going to attach HP/UX or TPGS Hosts
- ► iSCSI CHAP Secret: Enter or change the iSCSI chap secret for this host
- Port Mask: To change the port mask, disable ports for this host by unchecking the checkbox.

Note: You can use port masks to control the node target ports that a host can access, this can be useful to limit the number of logins with mapped volumes visible to a host multipathing driver, instead of changing the SAN zoning. The best practice recommendation is to limit each host to 4 paths.

Make the changes, if required, and click **Save** to apply them (Figure 8-30) to close the editing view.

The **Mapped Volume** tab gives you an overview as to which volumes are mapped with which SCSI ID and UID to this host as shown in Figure 8-31.



Figure 8-31 Mapped Volumes

The **Port Definitions** tab shows you the configured host ports of a host and gives you status information as shown in Figure 8-32.

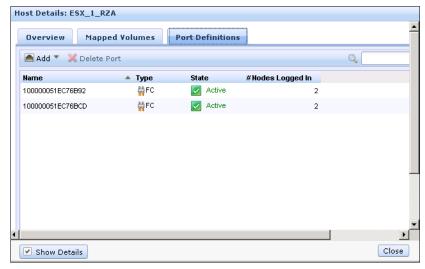


Figure 8-32 Port Definitions

This view offers you the option to start **Add** and **Delete Port** actions, and they are detailed as described in 8.2, "Add and delete host ports" on page 289.

Click Close to close the Host Details section.

8.2 Add and delete host ports

To configure host ports click **Ports by Host** as shown in Figure 8-1 to open the associated view as shown in Figure 8-33.

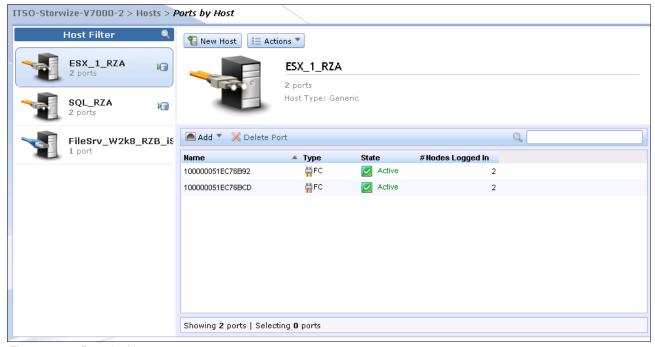


Figure 8-33 Ports by Host

In the left side of the view you will find all hosts listed, and the icons give you a overview if it is a Fibre Channel (orange cable) or iSCSI (blue cable) host. The properties of the highlighted host are shown in the right view. If you select the **New Host** button, the wizard as described in 4.2, "Creating Hosts using the GUI" on page 144 will start. If you expand the **Action** panel as shown in Figure 8-34 on page 290 the same tasks as described in 8.1.1, "Modify Mappings" on page 276 are able to be started from this location.

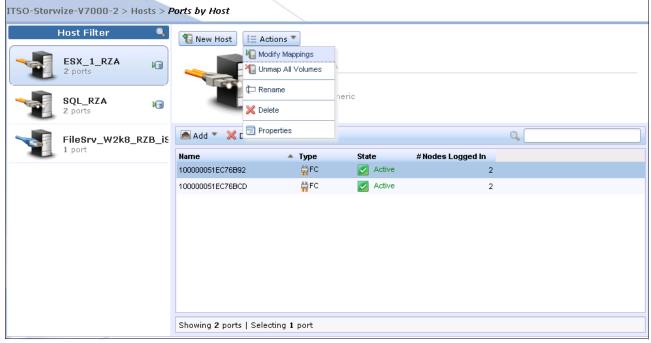


Figure 8-34 Ports by Host Actions

8.2.1 Add Host Port

To add a host port highlight the host, and click **Add** as shown in Figure 8-33 and select if you want to add a Fibre Channel or an iSCSI port as shown in Figure 8-35.



Figure 8-35 Add Host Ports

8.2.2 Add Fibre Channel Port

Click **Add Fibre-Channel Port** and the **Add Fibre-Channel Ports** view will appear as shown in Figure 8-36 on page 291.



Figure 8-36 Add Fibre-Channel Ports

If you click the drop down bar, you will get a list of all known Fibre-Channel host ports as shown in Figure 8-37.



Figure 8-37 Add Fibre Channel Ports - Known WWPNs

Select the WWPN you want to add and click **Add Port to List** as shown in Figure 8-38.



Figure 8-38 Add Port to List

It is possible to repeat this step to add more ports to a host. If the WWPN of your host is not available in the drop down list, check your SAN zoning, and rescan the SAN from the host. Afterwards click **Rescan** and the new port should now be available in the drop down list.

If you want to add an offline port simply type the WWPN of the port into the Fibre Channel Ports box as shown in Figure 8-39 and click **Add Port to List**.

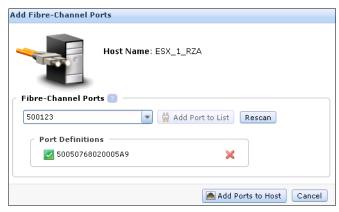


Figure 8-39 Add Offline Port

The Port will appear as **unverified** as shown in Figure 8-40 because it is not logged in to the IBM Storwize V7000. The first time it logs on, it will change is state to online automatically and the mapping will be applied to this port.



Figure 8-40 Unverified Port

To remove one of the ports from the list simply click the red x next to it as shown in Figure 8-41.

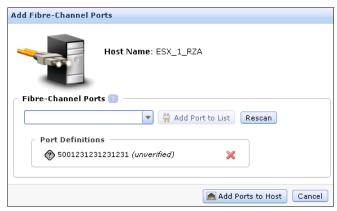


Figure 8-41 Remove Port from List

Click Add Ports to Host and the changes will be applied as shown in Figure 8-42.

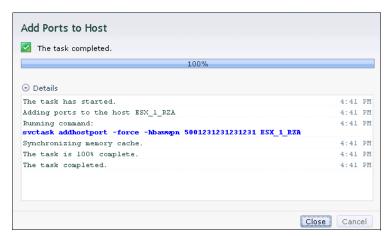


Figure 8-42 Add Ports to Host

Click Close to return to the Ports to Host view.

8.2.3 Add iSCSI Host Port

To add an iSCSI host port select **iSCSI Port** as shown in Figure 8-35 and the Add iSCSI Port view will appear as shown in Figure 8-43.



Figure 8-43 Add iSCSI Host Ports

Enter the initiator name of your host (Figure 8-44) and click **Add Port to List** as shown in Figure 8-45.



Figure 8-44 Enter Initiator Name



Figure 8-45 Add Ports to List

Click **Add Ports to Host** to completed the tasks and apply the changes to the system as shown in Figure 8-46 on page 295.

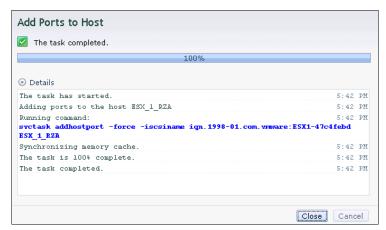


Figure 8-46 Add Host Ports- Apply Changes

Click Close to return to the Ports by Host view.

8.2.4 Delete Host Port

To delete a host port highlight it and click **Delete** as shown in Figure 8-47.

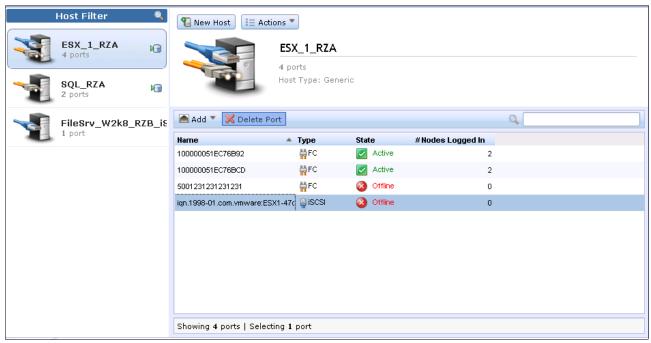


Figure 8-47 Delete Host Ports

If you hold the **Ctrl** Key on your keyboard, you can also select several host ports to delete as shown in Figure 8-48 on page 296.

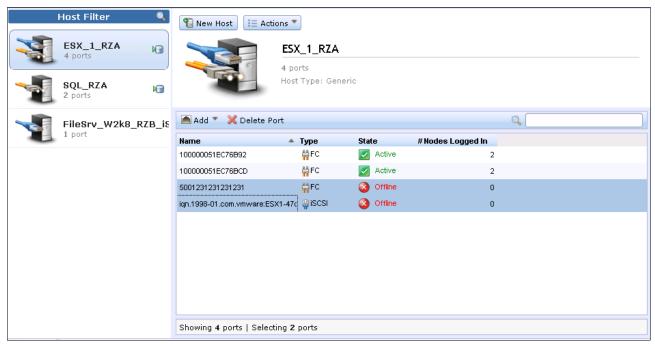


Figure 8-48 Delete several Host Ports

Click **Delete** and enter the number of host ports you want to remove as shown in Figure 8-49.

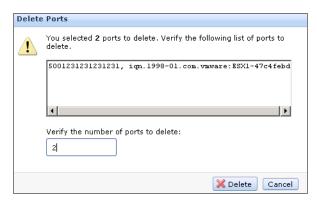


Figure 8-49 Enter # of Host Ports to delete

Click **Delete** and the changes will be applied to the system as shown in Figure 8-50 on page 297.

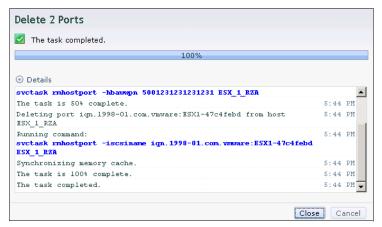


Figure 8-50 Delete Host Ports

Click Close to return to the Host by Ports view.

8.3 Host mappings overview

Select **Host Mappings** as shown in Figure 8-1 to open the host mappings overview shown in Figure 8-51.

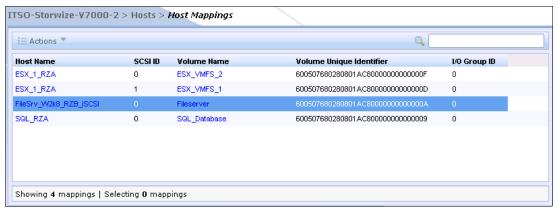


Figure 8-51 Host Mappings

This will generate a list of all host and volumes. In our example you can see the host "ESX_RZA" has two volumes mapped, and the associated SCSI ID, Volume Name, and the Volume Unique Identifier (UID). If you have more than one I/O Group, you also see which volume is handled by which I/O group.

If you highlight one line and click **Actions** as shown in Figure 8-52 on page 298 you will have the following tasks available:

- ► Unmap Volume ("Unmap Volume" on page 298)
- Properties Host ("Properties Host" on page 299)
- ► Properties Volume ("Properties Volume" on page 299)

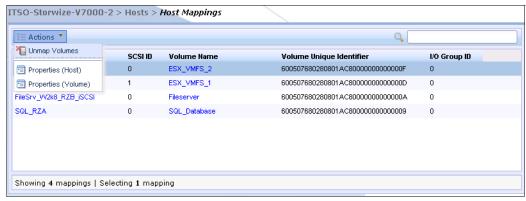


Figure 8-52 Host Mappings Action

8.3.1 Unmap Volume

Highlight one or more lines and select **Unmap Volume**, enter the number of entries to remove (Figure 8-53) and click **Unmap**, this will remove the mapping(s) for all selected entries as shown in Figure 8-54.

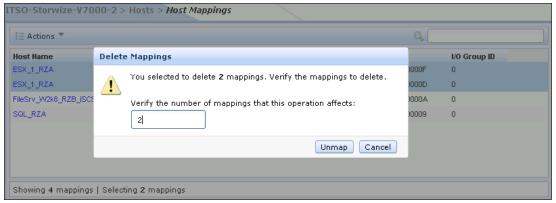


Figure 8-53 Unmap selected Volumes

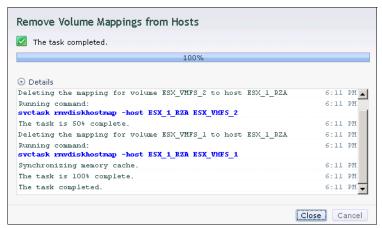


Figure 8-54 Unmap Complete

8.3.2 Properties Host

Selecting an entry and clicking on **Properties Host** as shown in Figure 8-52 on page 298 will open the host properties view. The possibilities of this view are described in "Host Properties" on page 286.

8.3.3 Properties Volume

Selecting an entry and clicking on **Properties Volume** as shown in Figure 8-52 on page 298 will open the volume properties view. The possibilities of this view are described in 5.1, "Provisioning storage from the IBM Storwize V7000 and making it available for the host" on page 158.

8.4 Advanced Volume Administration

This section covers volume administration such as volume modification, migration or creating new volume copies. Basic volume creation is covered in 5.1, "Provisioning storage from the IBM Storwize V7000 and making it available for the host" on page 158. In this chapter 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 8-55 shows that there are three volumes sections available to administer advanced features:

- ► **All Volumes** (8.4.1, "Advanced Volume Functions" on page 300 and 8.6, "Advanced Volume Copy Functions" on page 324)
- ▶ Volumes per Pool (8.7, "Volumes by Storage Pool" on page 330)
- ▶ Volumes per Host (8.8, "Volumes by Host" on page 334)

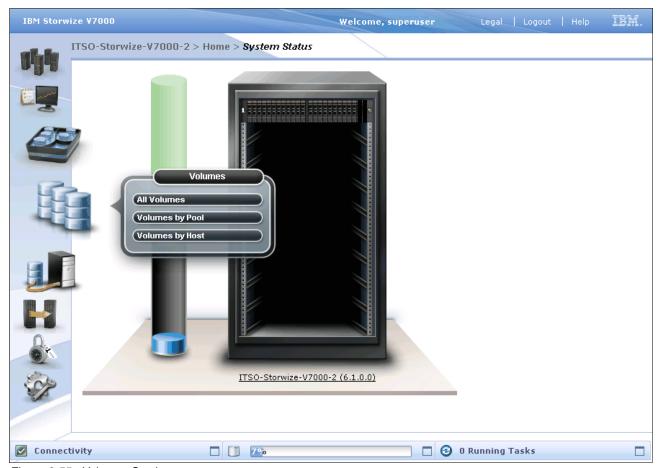


Figure 8-55 Volumes Section

8.4.1 Advanced Volume Functions

Select **All Volumes** as shown in Figure 8-55 and the associated view will appear as shown in Figure 8-56 on page 301.

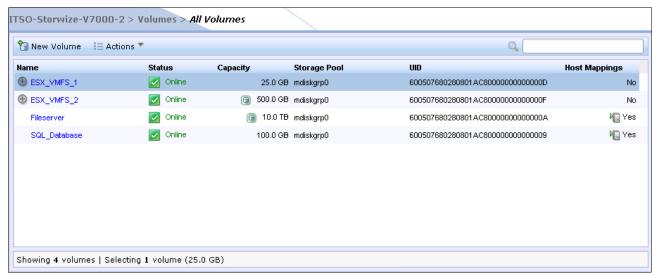


Figure 8-56 All Volumes

This view lists all configured volumes on the system and provides the following information to you:

- ▶ Name: Displays the name of the volume, if there is a + next to the name this means that there are several copies of this volume. Click it to expand the view and list the copies as shown in Figure 8-57.
- ▶ Status: Gives you status information about the volume
- ► Capacity: The capacity which is presented to the host is listed here. If there is a little blue volume listed to the capacity, this means that this volume is a Thin-provisioned volume, and that the listed capacity is the virtual capacity, which may be less than the real capacity on the system.
- Storage Pool: Shows in which Storage Pool the volume is stored. If you have several volume copies it shows you the pool of the primary copy.
- ► UID: Is the volume Unique IDentifier
- ► Host Mappings: Provides information if the volume is mapped to at least one host or not.

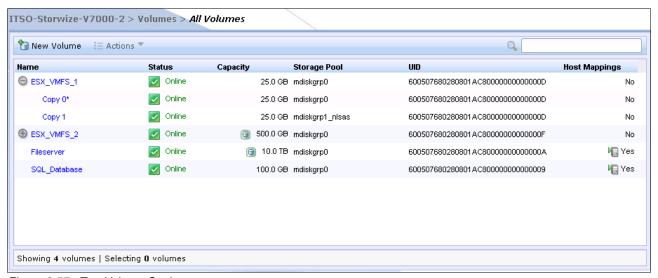


Figure 8-57 Two Volume Copies

To create a new volume click **New Volume** and follow the steps as described in 5.1, "Provisioning storage from the IBM Storwize V7000 and making it available for the host" on page 158.

Highlight a volume and click **Actions** to see the available actions for a volume as shown in Figure 8-58.

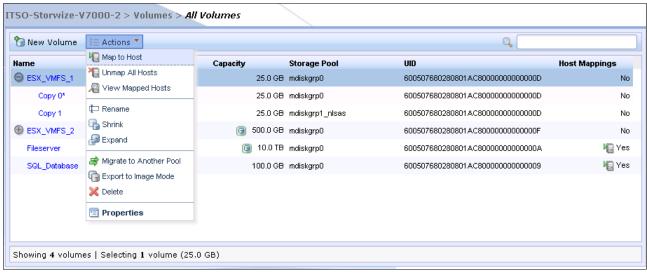


Figure 8-58 Volume Actions

It is also possible to right click a volume and select the actions there as shown in Figure 8-59 on page 302.

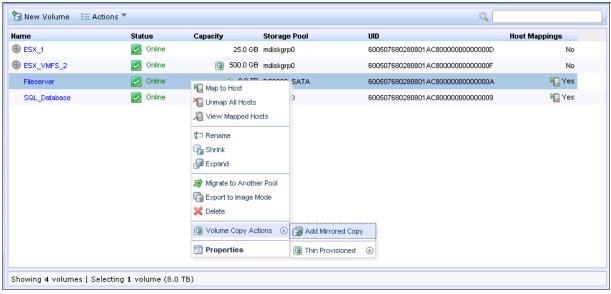


Figure 8-59 More Volume Actions

The following volume options are available:

- Map to Host ("Map Volume to Host" on page 303)
- Unmap All Hosts ("Unmap Volumes from all Hosts" on page 304)
- ► View Mapped Host ("View mapped Host to a Volume" on page 305)

- ► Rename ("Rename a Volume" on page 306)
- ► Shrink ("Shrink a Volume" on page 307)
- ► Expand ("Expand a Volume" on page 308)
- ▶ Migration to Another Pool ("Migrate volume to another Storage Pool" on page 309)
- ► Export to image Mode ("Export to Image Mode Volume" on page 311)
- ▶ **Delete** ("Delete Volume" on page 313)
- ► **Properties** ("Volume Properties" on page 314)

Depending on which kind of volume you view, there may be two more available Actions as also shown in Figure 8-59.

- ► Volume Copy Actions:
 - Add Mirror Copy: only available for generic volumes ("Add a Mirrored Volume Copy" on page 320)
 - Thin Provisioned: only available for Thin-provisioned volumes ("Edit Thin-provision Volume Properties" on page 321)

8.4.2 Map Volume to Host

If you want to map a volume to a host select **Map to Host** as shown in. Figure 8-58. Select the host to which you want to map the volume and click **Next** as shown in Figure 8-60.

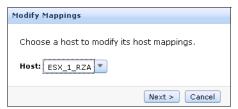


Figure 8-60 Select Host

The modify mappings view will appear, in the left upper corner, you will see your selected host, and the yellow volume is the selected volume that will be mapped as shown in Figure 8-61 on page 304. Click **OK** to apply the changes to the system.

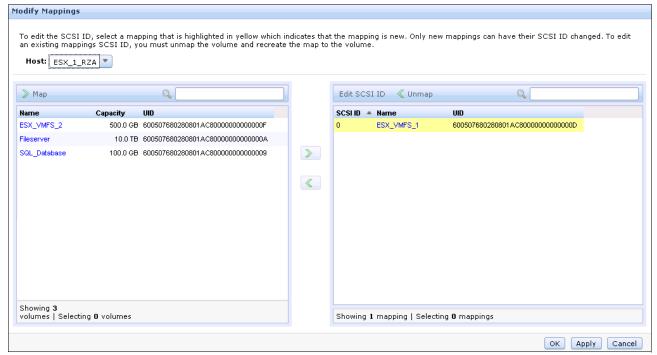


Figure 8-61 Modify Mappings

After the changes have completed click **Close** to return to the All Volumes view as shown in Figure 8-62.

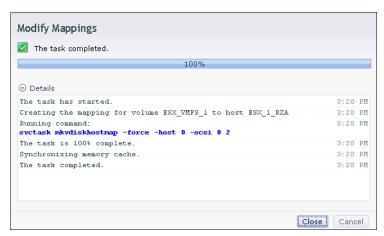


Figure 8-62 Modify Mappings

Note: The Modify Mappings view is described detailed in Chapter 8.1.1, "Modify Mappings" on page 276.

8.4.3 Unmap Volumes from all Hosts

If you want to remove all host mappings from a volume select **Unmap All Host** as shown in Figure 8-58. This will remove all host mappings from this volume, which means that no host will have access to this volume anymore. Enter the number of mappings that are affected and click **Unmap** to perform this action as shown in Figure 8-63 on page 305.

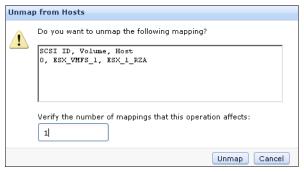


Figure 8-63 Unmap Volumes - Enter number of mappings to be removed

After the task has completed click **Close** as shown in Figure 8-64 to return to the All Volumes view.



Figure 8-64 Unmap Volume from all Host complete

8.4.4 View mapped Host to a Volume

If you want to know which host mappings are currently configured highlight a volume and select **View Mapped Host** as shown in Figure 8-58. This will open the **Host Maps** tab of the Volume Details views shown in Figure 8-65 on page 306. In the example you will see there is one existing host mapping for the volume SQL_Database.

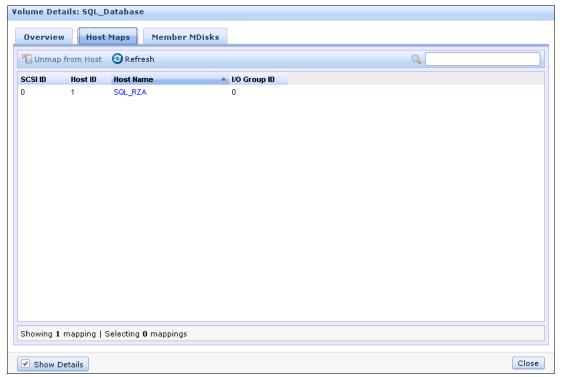


Figure 8-65 Volume Details

If you want to remove a mapping highlight the host and click **Unmap from Host** this will remove the access for the selected host after you have confirmed it. If there are several hosts mapped to this volume (for example in a cluster) only the highlighted host will be removed.

8.4.5 Rename a Volume

To rename a volume select **Rename** as shown in Figure 8-58 on page 302. A new window will appear, enter the new name there as shown in Figure 8-66.



Figure 8-66 Rename Volume

If you click **Reset** the name field will always be reset to the currently active name of the volume. Click **Rename** to apply the changes and click **Close** when it is done as shown in Figure 8-67 on page 307.

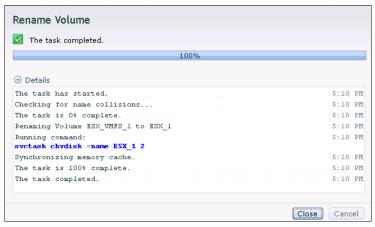


Figure 8-67 Rename Volume Complete

8.4.6 Shrink a Volume

IBM Storwize V7000 offers you the possibility to shrink volumes. However, you should only use this feature if your host operating systems supports this as well. Before shrinking a volume perform the preparation required in your host operating system to shrink a volume on the storage. After you have prepared it select **Shrink** as shown in Figure 8-58 on page 302. In the opened view you can either enter the new size, or enter how much the volume should shrink. If you enter a value the other line will update itself as shown in Figure 8-68.

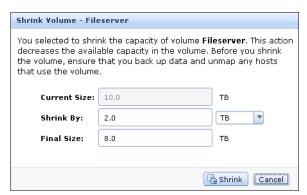


Figure 8-68 Shrink Volume

Click **Shrink** to start the process and click **Close** as shown in Figure 8-69 on page 308 to return to the All Volumes view.

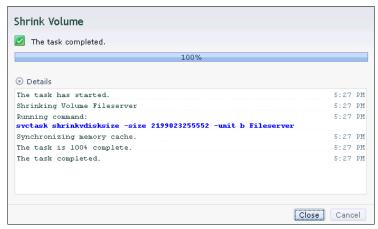


Figure 8-69 Shrink Volume Complete

Run the required procedures on your host to complete the shrink process.

8.4.7 Expand a Volume

If you want to expand a volume click **Expand** as shown in Figure 8-58 on page 302 and the expand view will appear. Before you continue check if your operating system supports online volume expansion. Enter the new volume size and click **Expand** as shown in Figure 8-70.

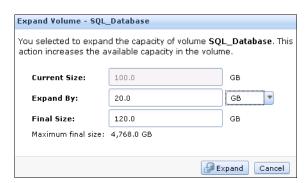


Figure 8-70 Expand Volume

After the tasks have completed click **Close** as shown in Figure 8-71 to return to the All Volumes View.

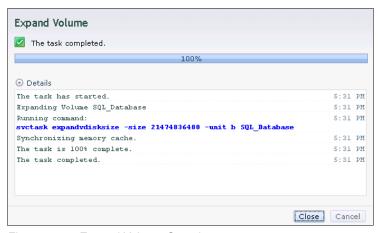


Figure 8-71 Expand Volume Complete

Run the required procedures in your operating system to use the newly available space.

8.4.8 Migrate 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 219. Volume migration allows you to move the data between these storage pools, regardless if it is an internal pool, or a pool residing on another external storage system. This migration is done without the servers and applications knowledge 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 will become more like the performance of the new storage pool during the migration process.

To migrate a volume to another storage pool select **Migrate to Another Pool** as shown in Figure 8-58 on page 302. The Migrate Volume Copy will appear. If your volume consists of more than one copy you will be asked which copy you want to migrate to another storage pool as shown in Figure 8-72. If the selected volume consists of one copy this selection box will not appear.

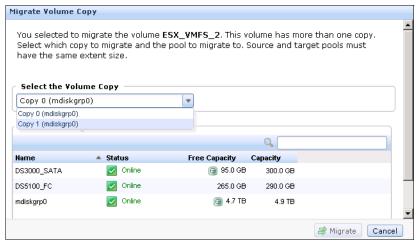


Figure 8-72 Migrate Volume - Select Copy

Migrate Volume Copy You selected to migrate the volume ESX_VMFS_2. This volume has more than one copy. Select which copy to migrate and the pool to migrate to. Source and target pools must have the same extent size. Select the Volume Copy Copy 1 (mdiskgrp0) Select a Pool 📝 Q Status Free Capacity Capacity 95.0 GB 300.0 GB Online DS3000_SATA DS5100_FC Online 265.0 GB 290.0 GB Online ⊕ 4.7 TB 4.9 TB mdiskgrp0

Select the new target storage pool and click Migrate as shown in Figure 8-73.

Figure 8-73 Migrate Volume Copy - Select Target Pool

The volume copy migration will start as shown in Figure 8-74. Click **Close** to return to the all volumes view.

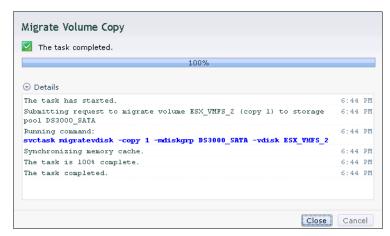


Figure 8-74 Migrate volume Copy Started

Depending on the size of the volume the migration process will take some time, but you can monitor the status of the migration in the running tasks bar as shown in Figure 8-75.

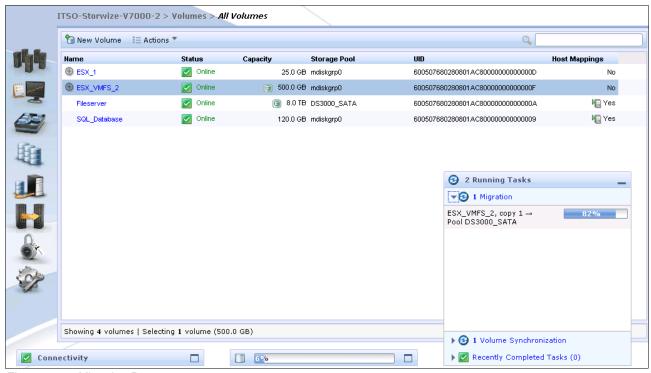


Figure 8-75 Migration Process

After the migration has completed volume is shown in the new storage pool. In Figure 8-76 on page 311 you can see that it was moved from pool mdiskgrp0 to pool DS3000_SATA.

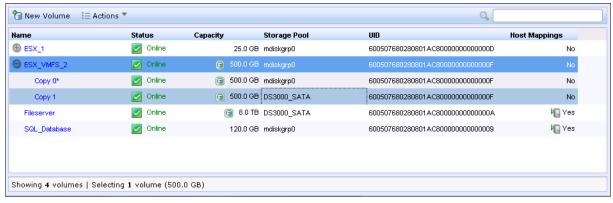


Figure 8-76 Migration Complete

The volume copy has now been migrated without any downtime to the new storage pool. It would also be possible to migrate both volume copies to other pools online.

Another way to migrate volumes to another pool is doing the migration with using the volume copies as described in "Migrate Volumes using the Volume Copy features" on page 329.

8.4.9 Export to 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 select **Export to Image Mode** as shown in Figure 8-58 on page 302 and the view will appear as shown in Figure 8-77. Select the MDisk to which you want to export the volume and click **Next.**

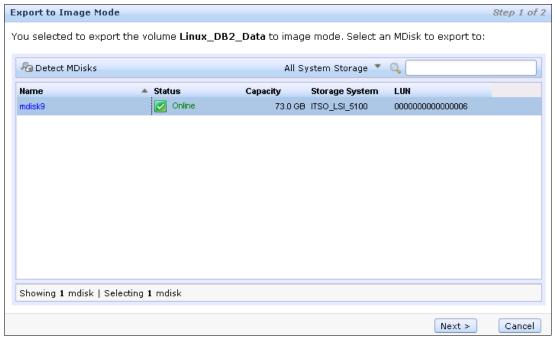


Figure 8-77 Export to Image Mode

Select a storage pool for the new image mode volume and click **Finish** as shown in Figure 8-78 on page 312.

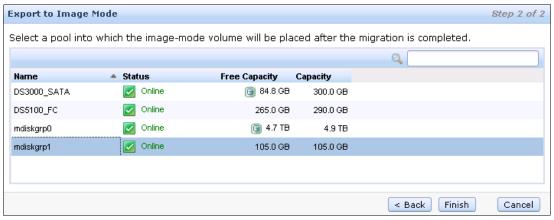


Figure 8-78

The migration will start as shown in Figure 8-79. Click Close to return to the All Volumes view.

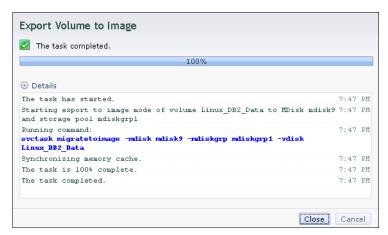


Figure 8-79 Export Volume to Image

8.4.10 Delete Volume

To delete a volume select **Delete** as shown in Figure 8-58. Enter the number of volumes you want to delete and mark the checkbox if you want to force the deletion as shown in Figure 8-80. You have to force it if the volume has host mappings or is used in FlashCopy mappings or remote-copy relationships.

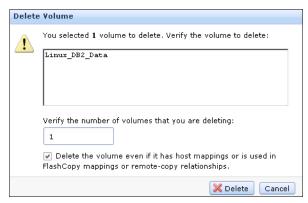


Figure 8-80 Delete Volume

Click **Delete** and the volume will be removed from the system as shown in Figure 8-81.

Note: This will remove all copies from your storage system and the data on the volume will be lost. Before you perform this step make sure you do not need it anymore.

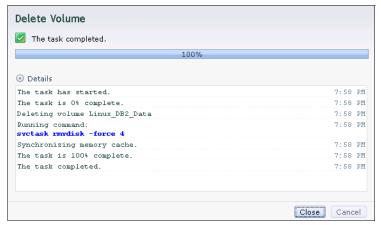


Figure 8-81 Delete Volume Complete

Click Close to return to the All Volumes view.

8.5 Volume Properties

To open an advanced view of a volume select **Properties** as shown in Figure 8-58 on page 302 and the volume properties will be displayed as shown in Figure 8-82 on page 314. In this view three tabs are available:

- Overview ("Overview" on page 315)
- ► Host Maps ("Host Maps" on page 317)
- Member MDisk ("Member MDisk" on page 319)

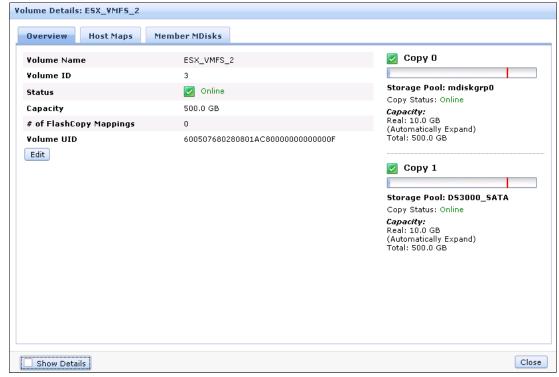


Figure 8-82 Volume Properties

8.5.1 Overview

The overview tab as shown in Figure 8-82 gives you a basic overview about the volume properties. In the left part of the view you will find common volume properties, and in the right part you are shown information about the volume copies. To get a more detailed view check the **Show Details** checkbox in the bottom left corner as shown in Figure 8-83.

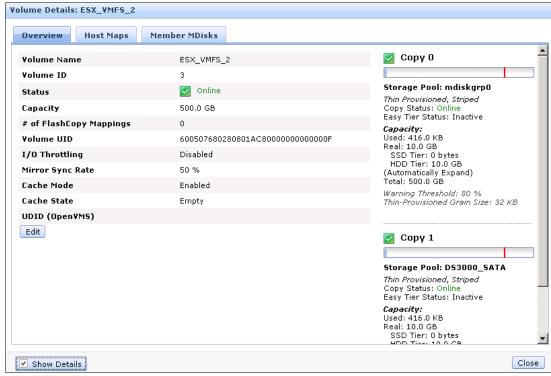


Figure 8-83 Show Details

The following details are available:

▶ Volume Properties:

- Volume Name: Displays 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.
- 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 relations. A detailed description about FlashCopy is covered in Chapter 11, "Copy Services" on page 381
- Volume UID: Is the Volume Unique IDentifier
- I/O Throttling: It is possible to set a maximum rate at which the volume process I/O requests. The limit can be set in I/Os to MB/s. This is an advanced feature and it is only possible to enable it via the CLI as described in Appendix A, "CLI setup and SAN Boot" on page 543.
- Mirror Sync Rate: After creation, or if a volume copy was offline, the mirror sync rate
 weights the synchronisation process. Volumes with a high sync rate (100%) will
 complete the synchronisation faster than volumes with a lower priority. Per default the
 rate is set to 50% for all volumes.

- Cache Mode: Displays if 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 now.
- UDID: The unit device identifiers is used by OpenVMS host to access the volume.

► Copy Properties:

- Storage Pool: Gives you information in which pool the copy resists, what kind of copy it is (Generic, Thin Provisioned), and gives you status information.
- Capacity: Displays the allocated (used) and the virtual (Real) capacity, as well as the warning threshold and the set grain size.

If you want to edit any of these settings click the **Edit** button and the view will change in the modify mode as shown in Figure 8-84.

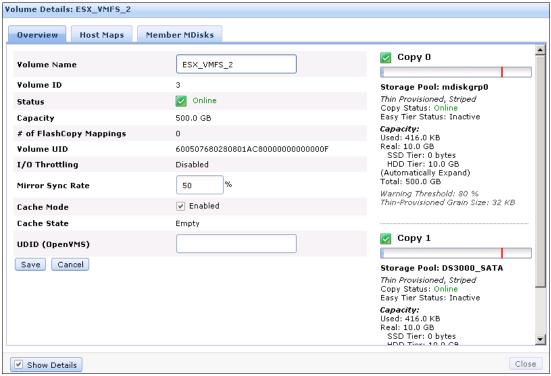


Figure 8-84 Edit Volume Details

Inside the volume details view it is possible to change the following properties:

- Volume Name:
- Mirror Sync Rate
- Cache Mode
- ▶ UDID

Make the changes if required and click **Save** as shown in Figure 8-85 on page 317.

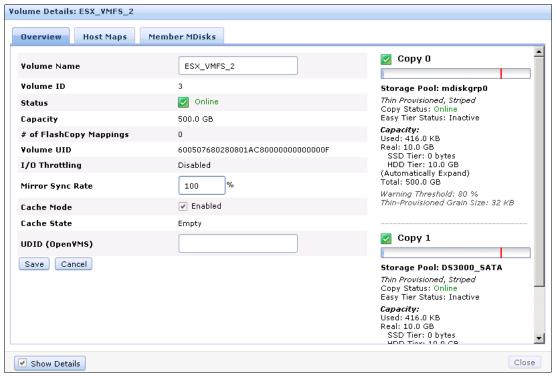


Figure 8-85 Change Properties

The changes will be applied to the system as shown in Figure 8-86.

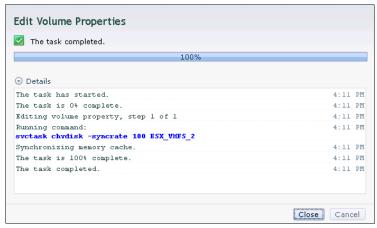


Figure 8-86 Apply Changes

Note: Setting the Mirror Sync Rate to 0% would disable the synchronization.

8.5.2 Host Maps

The second tab of the volume properties is **Host Map** as shown in Figure 8-87 on page 318. All hosts that are currently mapped to this volume are listed in this view.

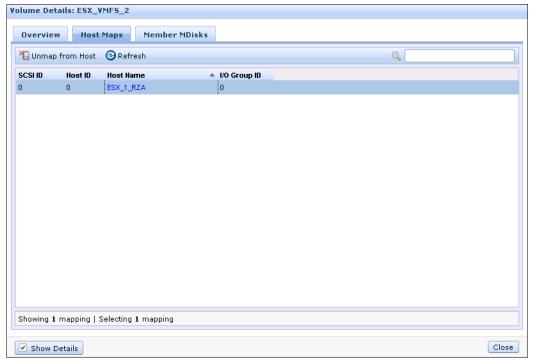


Figure 8-87 Host Maps

If you want to unmap a Host from this volume highlight it and select **Unmap from Host.** Confirm the number of mappings to remove and click **Unmap** as shown in Figure 8-88.

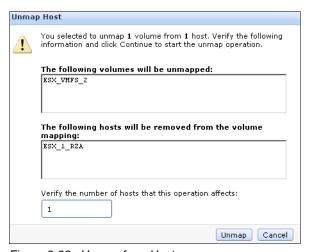


Figure 8-88 Unmap from Host

The changes will be applied to the system as shown in Figure 8-89 on page 319. The selected host will no longer have access to this volume. Click **Close** to return to the Host Maps view.

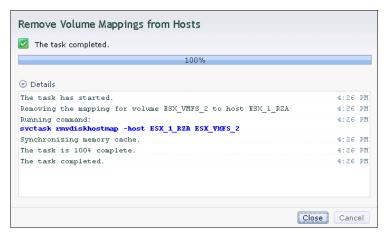


Figure 8-89 Unmap Host Complete

8.5.3 Member MDisk

The third tab is Member MDisk, which lists all MDisks on which the volume resides. Select a copy and the associated MDisks will appear in the view as shown in Figure 8-90.

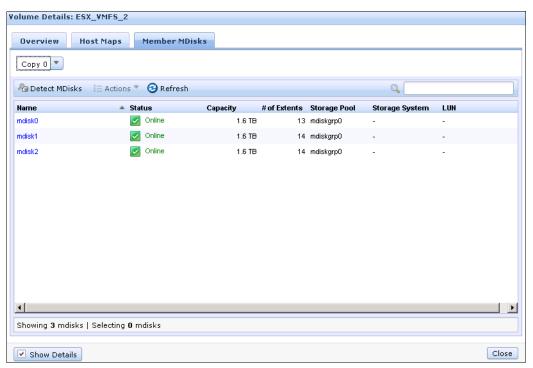


Figure 8-90 Member MDisk

Highlight a MDisk and click **Actions** to get a view of available tasks as shown in Figure 8-91 on page 320. The tasks are described in Chapter 7, "Storage Pools" on page 219.

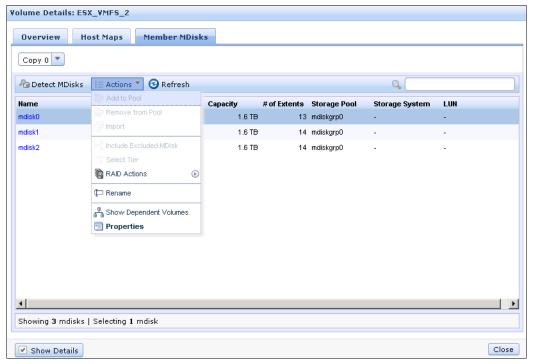


Figure 8-91 Member MDisk Actions

Click Close to return to the All Volumes view.

8.5.4 Add 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 will create an additional copy of your volume online. The second copy can either be generic or Thin-provisioned. To add a second copy select **Volume Copy Actions** \rightarrow **Add Mirrored Copy** as shown in Figure 8-59 on page 302. Select the storage pool in which the new copy should be migrated and if the new copy should be Thin-provisioned mark the checkbox as shown in Figure 8-92. Click **Add Copy**.

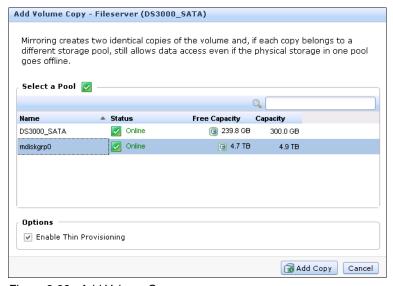


Figure 8-92 Add Volume Copy

The new copy will be created as shown in Figure 8-93.

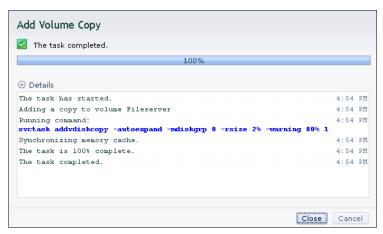


Figure 8-93 Add Volume Copy Complete

The new copy now has been created and the data is synchronized as a background task. Figure 8-94 shows you that the volume named "Fileserver" now has two volume copies, and there is one synchronisation task running.

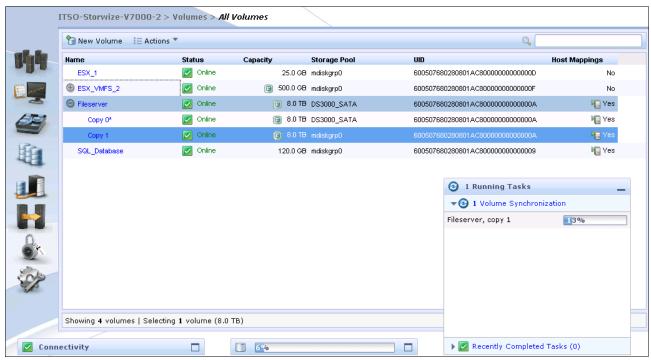


Figure 8-94 Volume Copy Synchronization

8.5.5 Edit Thin-provision Volume Properties

In "Shrink a Volume" on page 307 and in "Expand a Volume" on page 308 is already covered how to modify the volume size that is presented to a host. But if you have a Thin-provisioned volume, you can also edit the allocated size and the warning thresholds. To do this select **Volume Copy Actions --> Thin Provisioned** as shown in Figure 8-59 on page 302 and there are three more options available as shown in Figure 8-95 on page 322:

Shrink: ("Shrink Thin-provisioned space" on page 322)

- ► Expand: ("Expand Thin Provisioned Space" on page 323)
- Edit Properties: ("Edit Thin-provisioned Properties" on page 323)

This changes are only to the internal storage, so you do not have to make any changes on your host.



Figure 8-95 Thin Provisioned Options

Shrink Thin-provisioned space

Select **Shrink** as shown in Figure 8-95 on page 322 to reduce the allocated space of a thin provisioned volume. Enter either how much the volume should shrink, or the new final size as shown in Figure 8-96 and click **Shrink**.

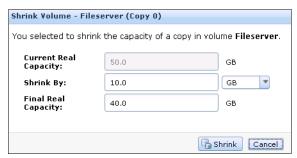


Figure 8-96 Shrink

After the task has completed click Close as shown in Figure 8-97.

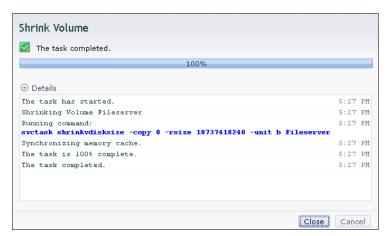


Figure 8-97 Shrink Thin-provisioned allocation

The allocated space of the Thin-provisioned volume has now been reduced.

Note: You can only unallocate extends that do not have stored data on it, if the space is allocated, because there are data on them you will not be able to shrink the allocated space.

Expand Thin Provisioned Space

To expand the allocated space of a Thin-provisioned volume select **Expand** as shown in Figure 8-95 on page 322. Enter either how much space should be allocated, or enter the new final size and click **Expand** as shown in Figure 8-98.

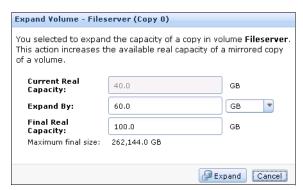


Figure 8-98 Expand Thin-provisioned Space

The new space is now allocated, click Close as shown in Figure 8-99.

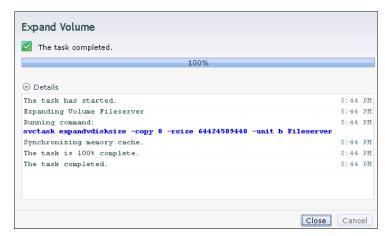


Figure 8-99 Expand Thin-provisioned Space Complete

Edit Thin-provisioned Properties

To edit Thin-provisioned properties select **Edit Properties** as shown in Figure 8-95 on page 322. Edit the settings if required as shown in Figure 8-100 and click **OK** to apply the changes.

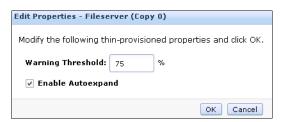


Figure 8-100 Edit Thin-provisioned Properties

After the task has completed click **Close** as shown in Figure 8-101 on page 324 to return to the All Volumes view.

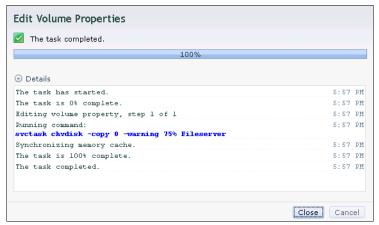


Figure 8-101 Edit Thin-provisioned Properties Complete

8.6 Advanced Volume Copy Functions

In 8.4.1, "Advanced Volume Functions" we discussed all the available actions at a volume level, and how to create a second volume copy. In this part we focus on volumes consisting of two volume copies, if you expand the volume and highlight a copy, there are some more volume copy actions are available as shown in Figure 8-102:

- ► Thin-provisioned: ("Thin-provisioned" on page 326)
- ► Split into New Volume: ("Thin-provisioned" on page 326)
- ► Validate Volume Copies: ("Validate Volume Copies:" on page 327)
- ▶ **Delete this Copy:** ("Delete this Copy" on page 329)

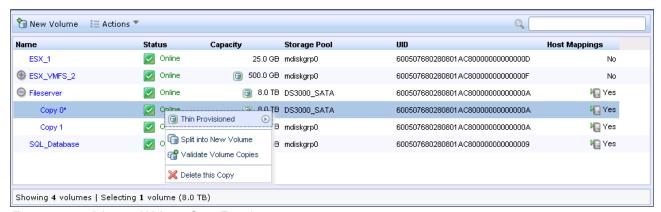


Figure 8-102 Advanced Volume Copy Functions

If you look at the volume copies as shown in Figure 8-102, you will notice that one of the copies, next to the Name, has a star displayed (Figure 8-103).



Figure 8-103 Primary Copy

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 copy always switches between Copy 0 and Copy 1 during creation to balance the reads across your storage pools. However you should make sure that the I/Os to the primary copies are customized to the performance of all your storage pools, and therefore you can change the roles of your copies.

To do this highlight the secondary copy and right click it or select **Actions** and you will see that there is also a task "Make Primary" as shown in Figure 8-104 on page 325. Usually it is a best practice to place the volume copies on storage pools with similar performance, because the write performance will be constrained if one copy is on a lower performance pool. This is because writes must complete to both copies before the volume can provide acknowledgment to the host that the write completed successfully. If you need high read performance only, another possibility can be to place the primary copy in an SSD pool, and the secondary copy to a normal disk pool. This will maximise the read performance of the volume, and make sure that you have a synchronized second copy in your less expensive disk pool. Of course, it is also possible to migrate the copies online between storage pools, remember, in "Migrate volume to another Storage Pool" on page 309 you can select which copy you want to migrate.

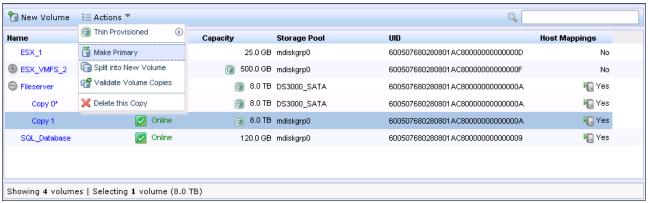


Figure 8-104 Make Primary

Select **Make Primary** and the role of the copy will be changed online. Click **Close** when the task has completed as shown in Figure 8-105.

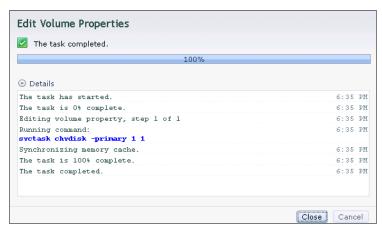


Figure 8-105 Make Primary Complete

The volume copy feature can also be a powerful option to migrate volumes as described later in "Migrate Volumes using the Volume Copy features" on page 329.

8.6.1 Thin-provisioned

This section includes the same functions as described in "Shrink Thin-provisioned space" on page 322 to "Edit Thin-provisioned Properties" on page 323 as shown in Figure 8-106. You can specify exactly the same settings for each volume copy.

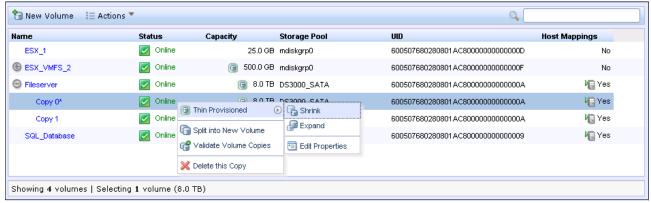


Figure 8-106 Thin-provisioned Volume Copy Tasks

8.6.2 Split into 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 new copy from the remaining one without any host impact. However if you want to use the split copy for testing or backup purposes, you have to make sure that the data inside the volume is consistent. Therefore you have to flush the data to the storage to get them consistent. Refer to your operating system documentation on how to flush the data. The easiest way is to shut the hosts down before splitting a copy.

To split a copy select **Split into New Volume** as shown in Figure 8-102. If you perform this action on the primary copy, the remaining secondary copy will automatically become the primary for the source volume. Enter a name for the new volume and click **Split Volume Copy** as shown in Figure 8-107.



Figure 8-107 Split Volume Copy

After the task has completed click **Close** to return to the All Volumes view where the copy appears now as a new volume that can be mapped to a host as shown in Figure 8-108 on page 327.

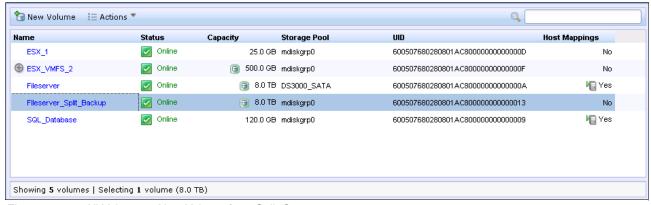


Figure 8-108 All Volumes - New Volume from Split Copy

8.6.3 Validate Volume Copies:

To validate the copies of a mirrored volume, perform the following steps:

Select **Validate Volume Copies** as shown in Figure 8-102 on page 324. The Validate Volume Copies window opens as shown in Figure 8-109 and there are the following options available:

- ► Generate Event of differences: Use this option if you only 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. This 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.

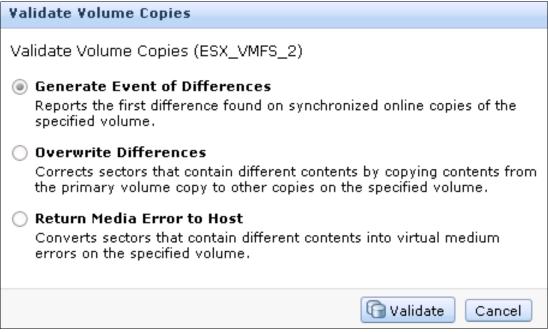


Figure 8-109 Validate Volume Copies

Select which action to perform and click **Validate** to start the task. The volume is now checked click **Close** (Figure 8-110).

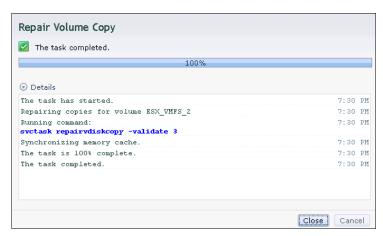


Figure 8-110 Validate Volume Copies - Started

The validation process takes some time depending on the volume size, and you can check the status in the running tasks view as shown in Figure 8-111 on page 329.

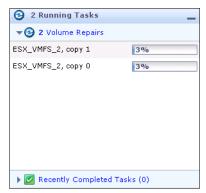


Figure 8-111 Validate Volume Copies - Running Tasks

8.6.4 Delete this Copy

Select **Delete** as shown in Figure 8-102 to delete a volume copy. The copy will be deleted, but the volume will remain online using the remaining copy. Confirm the delete process by clicking **OK** as shown in Figure 8-112.



Figure 8-112 Delete Copy

The copy will be deleted, click Close (Figure 8-113) to return to the All Volumes view.

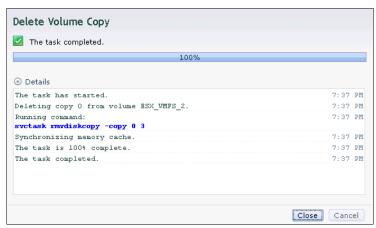


Figure 8-113 Delete Copy Complete

8.6.5 Migrate Volumes using the Volume Copy features

In the previous sections we have shown that it is possible to create, synchronize, split and delete volume copies. A combination of these tasks can be used to migrate volumes to other storage pools. The easiest way to migrate volume copies is to use the migration feature as shown in "Migrate volume to another Storage Pool" on page 309. If you use this one extent

after another would be migrated to the new storage pool. However using volume copies gives you another possibility to migrate volumes. To migrate a volume perform the following steps:

- 1. Create a second copy of your volume in the target storage pool ("Add a Mirrored Volume Copy" on page 320).
- 2. Wait until the copies are synchronized.
- 3. Change the role of the copies and make the new copy primary (8.6, "Advanced Volume Copy Functions" on page 324).
- 4. Split or delete the old copy from the volume ("Split into New Volume" on page 326 or "Delete this Copy" on page 329).

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 Tier2 storage pool. In Step 1 you will create the copy on the Tier 2 pool, while all reads will be still be 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 with the testing against your lower performance pool, you can decide to either split or delete the old copy in Tier 1, or to easily switch back to Tier1 in seconds, because the Tier2 pool did not meet your requirements.

8.7 Volumes by Storage Pool

To get an overview as to which volumes reside on which storage pool click **Volumes by Pool** as shown in Figure 8-55 and the Volumes by Pool view will appear as shown in Figure 8-114 on page 330.

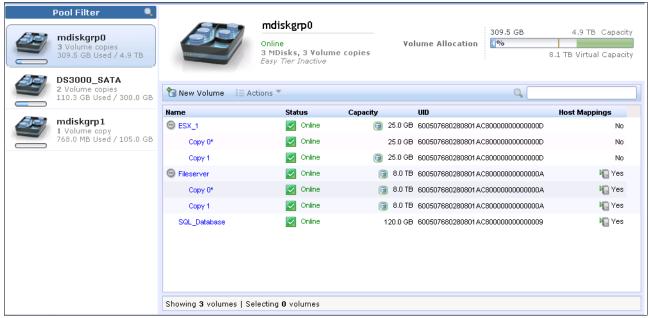


Figure 8-114 Volumes by Pool

The left part of the view is named "Pool Filter", and all of your existing storage pools are displayed there. More detailed information about storage pools is covered in Chapter 7, "Storage Pools" on page 219.

To the upper right you will see information about the pool which you selected in the pool filter, and the following information is covered there:

- ▶ **Pool Icon:** Storage Pools have different characteristics, it is possible to change the pool icon to identify the pool type ("Changing the Storage Pool Icon" on page 332).
- ▶ **Pool Name:** This is the name of the storage pool given during creation, it is possible to change this name from this view ("Changing the Storage Pool Name" on page 333).
- ► **Pool Details:** Gives you status information about the pool, like the number of MDisks and Volume copies, as well as the EasyTier Status.
- ▶ **Volume allocation:** Provides details about the available, allocated and virtual space in this pool.

The lower right lists all volumes that have at least one copy in this storage pool and provides you the following information about them:

- ▶ Name: Shows the name of the volume.
- ► Status: Gives status feedback about the volume.
- ► Capacity: Displays the capacity which is presented to the host of the volume, if it has the green volume sign next to it, this means this volume is Thin-provisioned and it is the virtual size.
- UID: Shows the Volume Unique IDentifier
- Host Mappings: Displays if at least one host mapping exists or not.

It is also possible to create a new volume in the selected pool. Click **Create Volume** and the same wizard as described in 5.1, "Provisioning storage from the IBM Storwize V7000 and making it available for the host" on page 158 will appear.

If you highlight a volume and select **Actions** or right click as shown in Figure 8-115 the same options as in the All Volumes section (8.4.1, "Advanced Volume Functions" on page 300) will appear. Detailed instructions about each task are covered there.

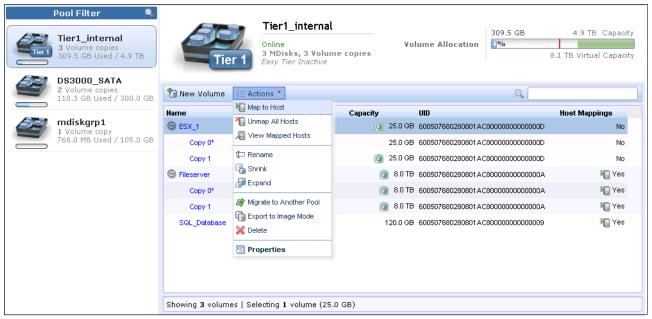


Figure 8-115 Volumes by Pool Actions

If you highlight a volume copy and select **Actions** or right click it as shown in Figure 8-116, the same options as in the All Volumes section (8.6, "Advanced Volume Copy Functions" on page 324) will appear. Detailed instructions about each task are covered there.

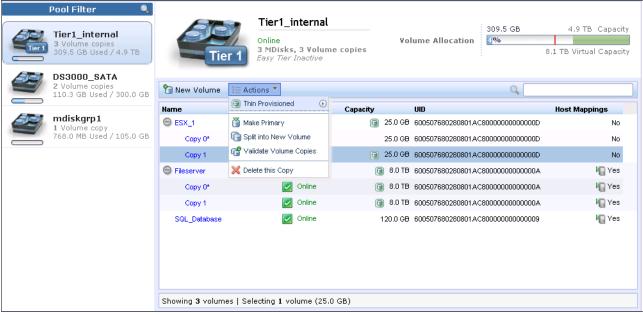


Figure 8-116 Volume Copy by Pool Actions

8.7.1 Changing the Storage Pool Icon

Usually Storage Pools have different characteristics and to reflect this in the view it is possible to change the pool icon. Click on the pool icon and the change icon view will appear (Figure 8-117 on page 332).



Figure 8-117 Change Icon

Use the left and right arrows to select a new icon as shown in Figure 8-118. There are several options available.



Figure 8-118 Select new Icon

Click **OK** and the changes will be applied to the system (Figure 8-119). Click **Close**.

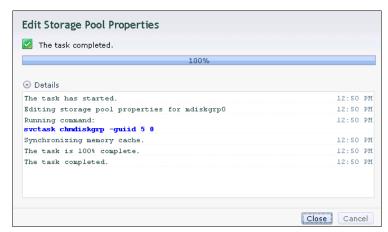


Figure 8-119 Change Icon - Complete

The icon has been changed to make it easier to identify the pool as shown in Figure 8-120.



Figure 8-120 Changed Icon

8.7.2 Changing the Storage Pool Name

To change the given name of a storage pool simply click on the name and you will be able to edit it as shown in Figure 8-121.



Figure 8-121 Change Name

Type a new name as shown in Figure 8-122.



Figure 8-122 Type New Name

Press Enter and the changes will be applied to the system as shown in Figure 8-123.

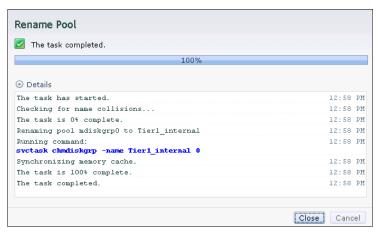


Figure 8-123 Rename Storage Pool - Complete

The name for the storage pool has now been changed as shown in Figure 8-124.



Figure 8-124 Change Name - Complete

8.8 Volumes by Host

To get an overview as to which volumes a host is able to access select **Volumes by Host** as shown in Figure 8-55 and the Volumes by Host view will appear as shown in Figure 8-125.

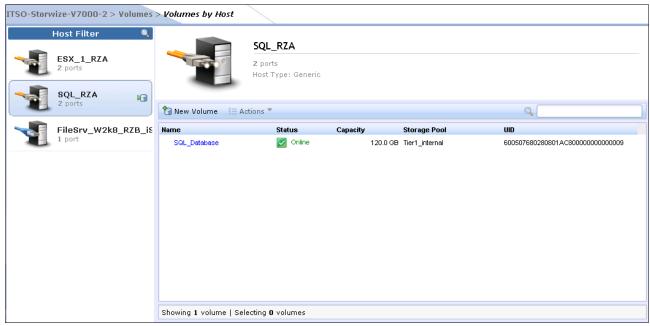


Figure 8-125 Volumes by Host

In the left side of the view is the "Host Filter" and if you select a host there its properties will appear in the right side of the view. The hosts with the orange cable represent Fibre Channel hosts, and the blue cable represents iSCSI hosts. In the upper right side you will see the host icon, the host name, the number of host ports and the host type. Below all the volumes that are mapped to this host are listed.

If you want to create a new volume for this host select **New Volume** and the same wizard as already described in 5.1, "Provisioning storage from the IBM Storwize V7000 and making it available for the host" on page 158 will appear.

Highlight a volume and select **Actions** or right click as shown in Figure 8-126 on page 335, and the same options as in the All Volumes section (8.4.1, "Advanced Volume Functions" on page 300) will appear. Detailed instructions about each task are covered there.

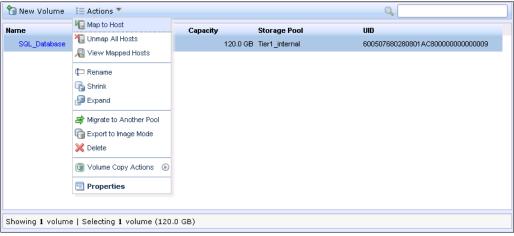


Figure 8-126 Volume Actions

If the volume owns more than one copy you can also highlight a volume copy and select **Actions** or right click it as shown in Figure 8-127, and the same options as in the All Volumes section (8.6, "Advanced Volume Copy Functions" on page 324) will appear. Detailed instructions about each task are covered there

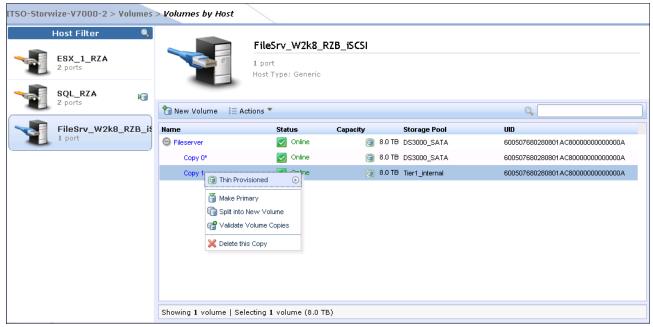


Figure 8-127 Volume Copy Actions

8.8.1 Rename a Host

To rename a host in the Volumes by host view simply click on it and you will be able to edit the name as shown in Figure 8-128.



Figure 8-128 Rename a Host

Enter a new name as shown in Figure 8-129.

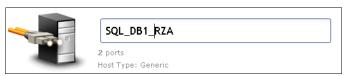


Figure 8-129 Enter new Name

Press **Enter** to apply the changes to the system as shown in Figure 8-130.

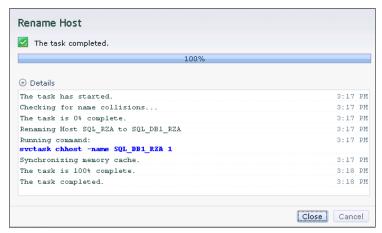


Figure 8-130 Rename Host Complete

Click **Close** to return to the Volumes by Host view as shown in Figure 8-131.

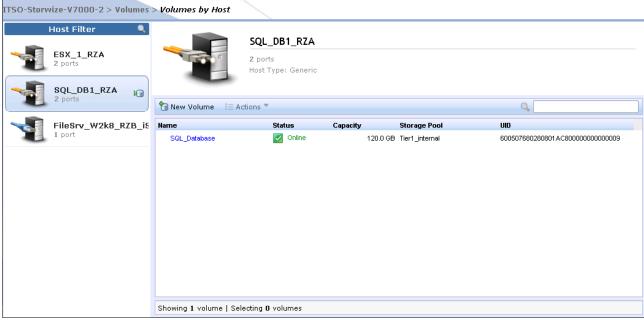


Figure 8-131 Volumes by Host



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 utilize a common set of advanced functions across all the storage systems under the IBM Storwize V7000's control.

9.1 Planning for External Storage Virtualization

In this section we will describe how to plan for virtualizing external storage with IBM Storwize V7000. Virtualizing the storage infrastructure with IBM Storwize V7000 will make your storage environment more flexible, cost-effective and easy to manage. The combination of IBM Storwize V7000 and external storage system will allow more storage capacity benefit from the powerful software function within the IBM Storwize V7000.

These external storage systems that will be incorporated into the IBM Storwize V7000 environment could be brand new systems or existing systems. The data on the existing storage systems could be easily migrated to the IBM Storwize V7000 managed environment, as described in Chapter 6, "Migration Wizard" on page 199 and Chapter 7, "Storage Pools" on page 219

9.1.1 License for External Storage Virtualization

From a licensing standpoint, when external storage systems are to be 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 10, and refer to your IBM account team or IBM Business Partner for further assistance if required.

Note: If the Storwize V7000 is to be used as a general migration tool, then 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 more-than-45-day migration requirement from external storage to IBM Storwize V7000 internal storage, then the appropriate External Virtualization license must be ordered too.

You can configure the IBM Storwize V7000 license by selecting the **Configuration** button, which is the bottom icon of the navigation panel on the left hand side, and choose the **Advanced** option.

Figure 9-1 on page 341 shows the **Advanced** option in **Configuration** button.

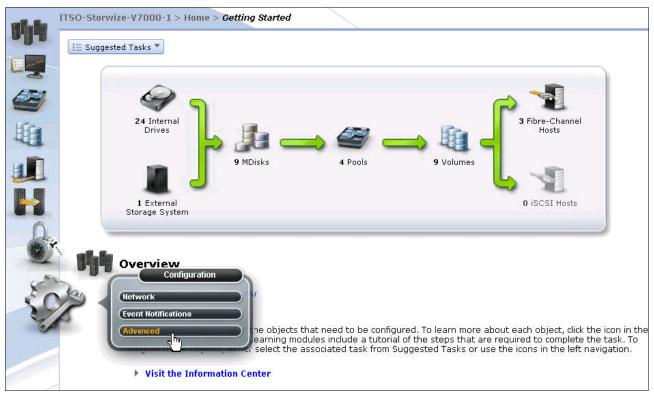


Figure 9-1 Advanced option in Configuration button

In the **Advanced** panel, select **Licensing** and you will find the **Update License** panel on the right, as shown in Figure 9-2 on page 342.

In the **Update License** panel there are two license options you can set, **External Virtualization Limit** and **Remote-Copy Limit**. Set these two license options to the limit you obtained from IBM.



Figure 9-2 Update License Panel

For assistance with licensing questions or to purchase an external virtualization or remote copy license, contact your IBM account team or IBM Authorized Business Partner.

9.1.2 SAN Configuration Planning

External Storage controllers to be virtualized by IBM Storwize V7000 must be connected via SAN switches. Direct connection between the IBM Storwize V7000 and storage controllers is not supported.

Verify that the SAN switches and/or directors that the IBM Storwize V7000 will connect to meet the following requirements as noted at:

http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003703# Switches

Make sure switches and/or directors are at firmware levels supported by the IBM Storwize V7000 and IBM Storwize V7000 port login maximums listed in restriction document will not be exceeded. The configuration restrictions can be found at:

 $\label{lem:http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003702\&myns=s028\&mynp=familyind5402112\&mync=E$

The recommended SAN configuration comprises a minimum of two fabrics with the ports on external storage systems to be virtualized by IBM Storwize V7000 and the IBM Storwize V7000 ports themselves evenly split between the two fabrics to provide redundancy in case one of the fabrics goes offline whether planned or unplanned.

Once IBM Storwize V7000 and external storage systems are connected to the SAN fabrics, zoning will need to be implemented. 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 each external storage system.

Note: IBM Storwize V7000 supports a maximum of sixteen ports or WWPNs from a given external storage system that will be virtualized.

Figure 9-3 on page 343 is an example of how to cable devices to the SAN. Reference this example as we discuss the recommended zoning below.

Create an IBM Storwize V7000/external storage zone for each storage system to be virtualized. For 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.

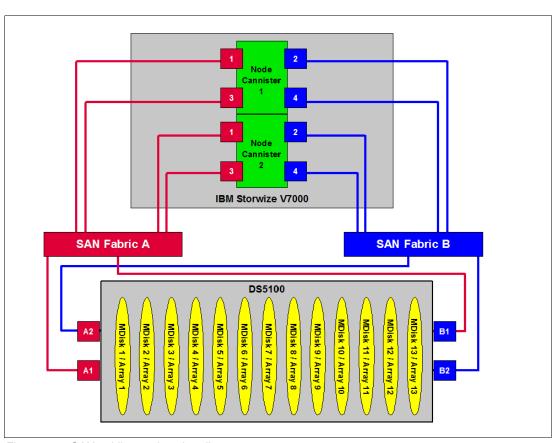


Figure 9-3 SAN cabling and zoning diagram

9.1.3 External Storage Configuration Planning

External storage systems provide redundancy via 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 then 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 following requirements found at the URL below:

http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003703

Make sure 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 external storage system mapped to it. LUNs cannot be shared between IBM Storwize V7000s or between IBM Storwize V7000 and other storage virtualization platforms or between IBM Storwize V7000 and hosts. However, different LUNs could be mapped from one external storage system to 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.

Review the IBM Storwize V7000 Information Center under the section titled **Configuring and servicing external storage system** to prepare the external storage systems for discovery by the IBM Storwize V7000 system. This is the link for IBM Storwize V7000 Information Center:

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

9.1.4 Guidelines for Virtualizing External Storage

When virtualizing external storage with the IBM Storwize V7000, there are a few guidelines to follow:

- ► Avoid splitting arrays into multiple LUNs at the external storage system level. When possible, create a single LUN per array for mapping to 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 on Easy Tier, you can find it in Chapter 10, "Easy Tier" on page 353.
- ▶ 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 service in Storwize V7000 will give a unified method to manage the 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 Work with External Storage

In this section we will discuss how to manage external storage with IBM Storwize V7000.

The basic concept 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 will ultimately be added to a storage pool in which volumes will be created and mapped to hosts as needed.

Therefore, in this section the steps for managing MDisks/storage pools, which have been mentioned in the Chapter 7, "Storage Pools" on page 219, will not be described again.

9.2.1 Add External Storage

To add new external storage systems to the IBM Storwize V7000 virtualized environment, these are the basic steps:

- ► Zone a minimum of two and a maximum of sixteen Fibre Channel ports from the external storage system with all eight Fibre Channel ports on the IBM Storwize V7000 system. It is recommended to 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.
- ► Using the storage partitioning or LUN masking feature of the external storage system create a group that includes all eight IBM Storwize V7000 WWPNs.
- ► Create equal size arrays on the external system using any RAID level but zero.
- Create a single LUN per RAID array.
- ► Map the LUNs to all eight Fibre Channel ports on the IBM Storwize V7000 system by assigning them to the group created above.
- Verify IBM Storwize V7000 discovered the LUNs as unmanaged MDisks.
 - If they don't show up automatically then select the **Detect MDisk** option from the MDisk panel of the GUI, as described in 7.2.5, "Additional Actions on MDisks" on page 263
- ► Select storage tier for the MDisks, as described in 7.2.4, "Select the tier for MDisks" on page 261.
- ► Create a new storage pool, as described in 7.3, "Work with Storage Pools" on page 266
- ► Add the MDisks to the pool, as described in 7.2.1, "Add MDisks to Storage Pools" on page 242.
- ► Create volumes and map them to hosts as needed, as described in Chapter 4, "Host Configuration" on page 127 and Chapter 5, "Basic Volume Configuration" on page 157.

If the external storage systems are not brand new ones, which means there is existing data on the LUNs that need to be kept after virtualization, follow the steps inChapter 6, "Migration Wizard" on page 199 to prepare the environment. Then you can migrate the existing data with or without the wizard to IBM Storwize V7000 internal storage or some other external storage system.

In 7.2.2, "Import MDisks" on page 246, it shows how to manually import MDisks and migrate the data to other storage pools. 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 Manage External Storage

The IBM Storwize V7000 provides an individual **External** panel for managing external storage systems.

You can access the **External** panel through the **Getting Started** panel, and click the **External Storage System** icon. Extended help information for external storage will show up click on **Physical Storage** and you will go to the **External** panel.

Figure 9-4 shows how to access the **External** panel from the **Getting Started** panel.



Figure 9-4 Access External panel from Getting Started panel

The other way to access the **External** panel is rom the **Physical Storage** functional icons on the left hand side.

Figure 9-5 on page 347 shows how to access the **External** panel from the **Physical Storage** functional icons on the left hand side.

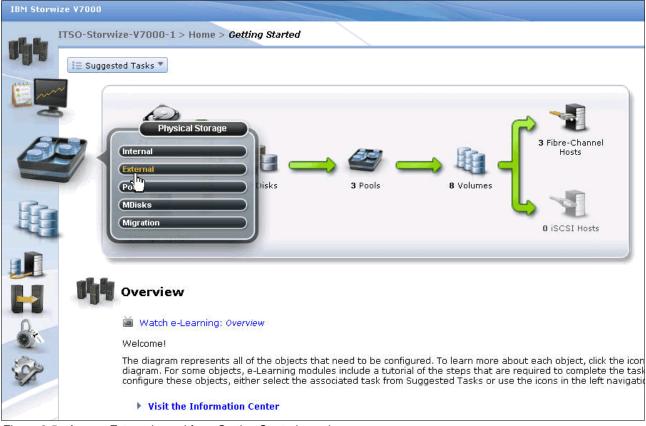


Figure 9-5 Access External panel from Getting Started panel

The **External** panel, as shown in Figure 9-6, gives you an overview of all your external storage systems. On the left hand side of the panel, there is a list of the external storage systems. With the help of the filter you can display only the external storage systems you need to take action on to be listed. If you click and highlight the external storage system, you will find detailed information shown on the right hand side including all the MDisks provided by it.

On the right of the panel you can change the name of external storage system by clicking on the name beside the picture of the external storage box. The status of the external storage system and its WWNN can also be found under the name.

Figure 9-6 on page 348 shows the External panel.

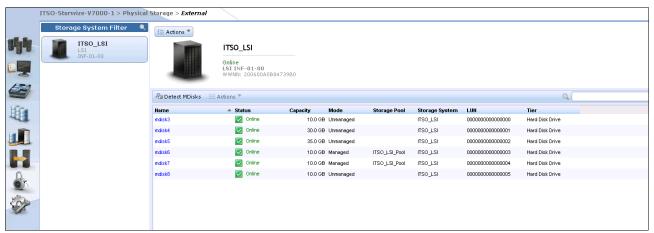


Figure 9-6 External panel

From the **Actions** drop down list on the top of the name of external storage on the right part of the **External** panel, you can find **Show Dependent Volumes** option as shown in Figure 9-7.

Figure 9-7 shows the Show Dependent Volumes option in Actions drop down list.

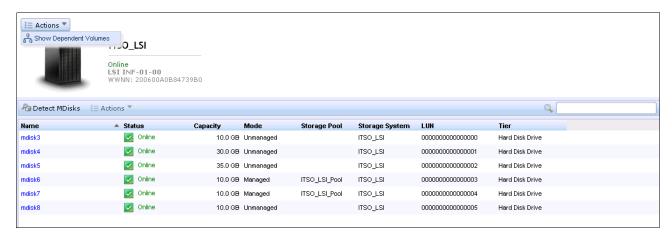


Figure 9-7 Show Dependent Volumes option in Actions drop down list

Clicking the **Show Dependent Volumes** option will show you the volumes in this external storage system.

Figure 9-8 on page 349 shows the volumes dependent on external storage.

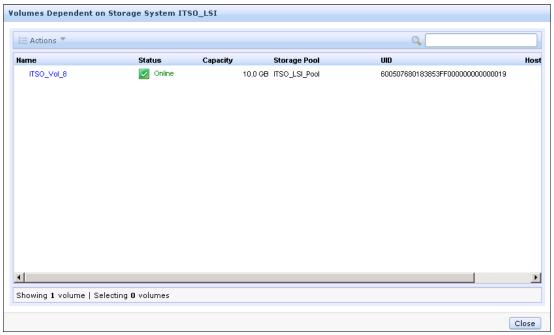


Figure 9-8 Volumes dependent on external storage

In the display of dependent volumes on external storage you can take volume actions, including **Map to Host**, **Shrink**, **Expand**, **Migrate to Another Pool**, **Volume Copy Actions** and so on as shown in Figure 9-9 on page 350.

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 non-disruptively from one internal or external storage system to another, which will make storage management much simpler with less risk.

Volume copy is another key feature that you can benefit from by introducing IBM Storwize V7000 virtualization. Two copies could be applied to your data to enhance the availability for a critical application. Volume copy could be also used for generating test data or data migration.

In Chapter 8, "Advanced Host and Volume Administration" on page 273, you can find more information on the volume actions of IBM Storwize V7000 storage system.

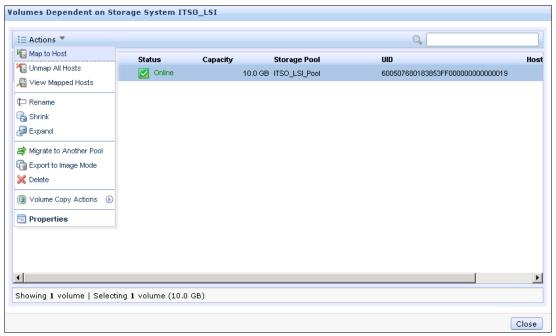


Figure 9-9 Actions you can take with volumes

Returning to the **External** panel you can find an MDisk menu on the right, including an MDisk list which shows the MDisks provided by this external storage system. You can find the name of an MDisk, its capacity, the storage pool and the storage system it belongs to in the list. The actions on MDisks can also be taken through the menu, including **Detect MDisks**, **Add to Pool**, **Import**, and so on. This menu is the same as in the **MDisks** panel which has already been covered in 7.2, "Work with MDisks" on page 239.

Figure 9-10 shows the MDisk menu in the External panel.

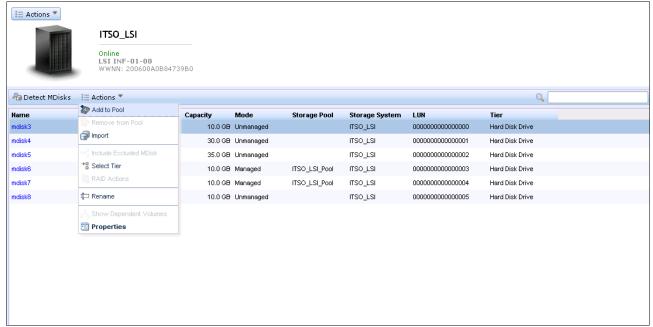


Figure 9-10 MDisk menu in the External panel

9.2.3 Remove External Storage

To remove the external storage systems from the IBM Storwize V7000 virtualized environment, the steps are dependent on your requirement:

- 1. If you want to remove the external storage systems and discard the data on it:
 - a. Stop any host I/O on the volumes.
 - b. Remove the volumes from host file system/logical volume/volume group, remove the volumes from host device inventory.
 - c. Remove the host mapping of volumes and the volumes themselves on IBM Storwize V7000. 8.1.1, "Modify Mappings" on page 276.
 - d. Remove the storage pools which the external storage systems belong to, as described in 7.3, "Work with Storage Pools" on page 266 or you can keep the storage pool and remove the MDisks of the external storage from the storage pools, as described 7.2.1, "Add MDisks to Storage Pools" on page 242
 - e. Unzone and disconnect the external storage systems wit.h IBM Storwize V7000.
 - f. Run **Detect MDisks** to make IBM Storwize V7000 discover the removal of the external storage systems, refer to 7.2.5, "Additional Actions on MDisks" on page 263.
- 2. if you want to remove the external storage systems and keep the volumes and their data on the IBM Storwize V7000:
 - a. Migrate volumes and their data to the other storage pools which are on IBM Storwize V7000 internal storage or other external storage systems, refer to 5.3, "Discover the volumes from the host and specify multipath settings" on page 170.
 - b. Remove the storage pools which the external storage systems belong to, 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 219.
 - c. Unzone and disconnect the external storage systems with IBM Storwize V7000.
 - d. Run **Detect MDisks** to make IBM Storwize V7000 discover the removal of the external storage systems. 7.2.5, "Additional Actions on MDisks" on page 263.
- 3. 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:
 - a. Migrate volumes and their data to the other storage pools which are on IBM Storwize V7000 internal storage or other external storage systems, described inChapter 5, "Basic Volume Configuration" on page 157.
 - b. Remove the storage pools which the external storage systems belong to, 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 219.
 - c. Export volumes to image mode with the MDisks on the external storage systems. For the restrictions and prerequisites for migration, you will find them in Chapter 6, "Migration Wizard" on page 199. You also need to record pre-migration information for use, for example, the original SCSI IDs the volumes used to be mapped to hosts. Some operating systems do not support changing SCSI ID during the migration. More information about migration you may find in the IBM Storwize V7000 Information Center at:

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

- d. Unzone and disconnect the external storage systems with IBM Storwize V7000.
- a. Run **Detect MDisks** to make IBM Storwize V7000 discover the removal of the external storage systems as described in Chapter 7, "Storage Pools" on page 219.



10

Easy Tier

In today's storage market Solid State Drives (SSDs) are emerging as an attractive alternative to hard disk drives (HDDs). Because of very 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 is currently much higher then HDDs. SSD performance depends a lot on workload characteristics, therefore SSDs need to be used in conjunction with HDDs and it is critical to choose the right mix of drives and the right data placement in order to achieve optimal performance at low cost. Maximum value can be derived by placing "hot" data with high IO density and low response time requirements on SSDs, while targeting HDDs for "cooler" data that is accessed more sequentially and at lower rates.

Easy Tier automates the placement of data amongst different storage tiers, and can be enabled for internal and external storage as well. This IBM Storwize V7000 no charge feature will boost your storage infrastructure performance to achieve optimal performance through a software, server and storage solution.

This chapter describes the function provided by the EasyTier disk performance optimization feature of the IBM Storwize V7000, and it also covers how to activate the Easy Tier process for both evaluation purposes and for automatic extent migration.

10.1 Easy Tier Overview

Easy Tier is a no charge feature with IBM Storwize V7000 which brings enterprise storage functions to the midrange segment. It enables automated sub-volume data placement throughout different storage tiers to intelligently align the system with current workload requirements, to optimize the utilization of SSDs drives. This includes the ability to automatically and non-disruptively relocate data (at the extent level) from one tier to the other 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 solves a similar access latency workload problem, but these two methods weigh differently in the algorithmic construction based on "locality of reference", recency and frequency. Since 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 to monitor I/O statistics on single extents and move them manually to an appropriate storage tier and react to workload changes.

Easy Tier is a performance optimization function that overcomes this issue as it will automatically migrate, or move, extents belonging to a volume between different storage tiers as shown in Figure 10-1 on page 355. As this migration works at the extent level, in other publications this technology is often referred to as "Sub-LUN Migration".

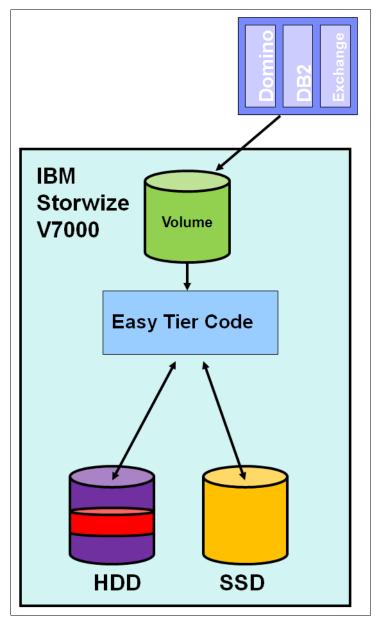


Figure 10-1 Easy Tier

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 storage pool, as well as moving extents whose activity has dropped off, or cooled, from higher disk tier MDisks back to a lower tier MDisk.

To enable this migration between MDisks with different tier levels, the target storage pool has to consist of different characteristic MDisks. These pools are named "Multi Tiered Storage Pools". IBM Storwize V7000 Easy Tier is optimized to boost the performance of storage pools containing HDDs and SSDs.

To identify the potential benefits of Easy Tier in your environment before actually installing higher MDisk tiers, such as SSDs, it is possible to enable the Easy Tier monitoring volumes in single tiered storage pools. Even though the Easy Tier extent migration is not possible within a single tier pool, the Easy Tier statistical measurement function is. Enabling Easy Tier on a

single tiered storage pool will start the monitoring process and log the activity of the volume extents. In this case, Easy Tier will create a migration plan file which can then be used to display a report on the number of extents that would be appropriate for migration to higher level MDisk tiers, such as SSDs. The IBM Storage Tier Advisor Tool is a no additional cost tool, that helps you in analyzing this data. If you do not have an IBM Storwize V7000 now, use Disk Magic to get a better idea about the required number of SSDs appropriate for your workload. In the worst case where you do not have any workload performance data, a good starting point can be to add about 5% of nett capacity as SSD to your configuration. But this ratio is heuristics based and will change according to different applications or different disk tier performance in each configuration. For database transactions a ratio of fast SAS or FC drives to SSD is about 6:1 to achieve the optimal performance but, again, this is just theory and depends a lot on the workload.

Easy Tier is available for IBM Storwize V7000 internal volumes, and for volumes residing on external virtualized storage subsystems as well, as the SSDs can be either internal and/or external drives. However, from the fabric point of view, it is recommended to use SSDs inside the IBM Storwize V7000, even if the lower tiered disk pool resides on an external storage because this will reduce the traffic traversing the SAN environment.

10.2 Easy Tier for IBM Storwize V7000

This section covers a description of the terms and and example implementation of Easy Tier on the IBM Storwize V7000. After reading this section you will be able to understand the Easy Tier concept as it relates to the IBM Storwize V7000.

10.2.1 Disk Tiers

It is likely that IBM Storwize V7000 internal disks, and also external disks, will have different performance attributes. As described in Chapter 7, "Storage Pools" on page 219 without Easy Tier it is a best practise to place drives with the same attributes (RPM, Size, Type) in the same storage pool, and not to intermix different drives with different attributes. This is also valid for external MDisks that are grouped together into storage pools. All internal HDDs and external MDisks are initially put into the generic_hdd tier by default. Internal SSD will be identified as a high performance tier MDisk by IBM Storwize V7000, and all external SSD MDisks must be changed to the high performance tier as shown in 7.2.4, "Select the tier for MDisks" on page 261.

10.2.2 Tiered Storage Pools

With IBM Storwize V7000 we have to differentiate between two types of storage pools:

- Single Tiered Storage Pools
- ► Multi Tiered Storage Pools

As shown in Figure 10-2 on page 357, Single Tiered Storage Pools include one type for disk tier attribute. Each disk should have the same size and performance characteristics. Multi Tiered Storage Pools are populated with two different disk tier attributes, which means high performance tier SSD disks, and generic HDD disks. A volume migration as described in "Migrate volume to another Storage Pool" on page 309 and is when the complete volume is migrated from one Storage Pool to another. An Easy Tier Data Migration will only move extents inside the storage pool to different performance attributes.

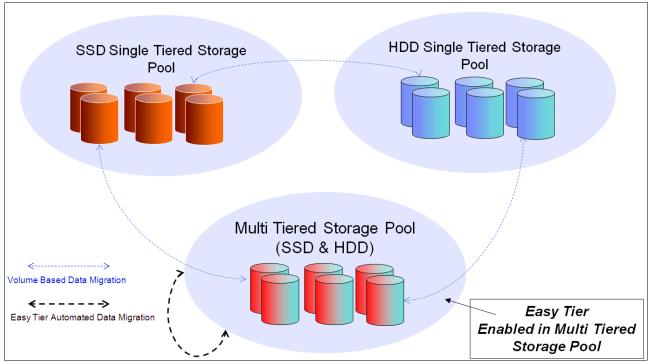


Figure 10-2 Tiered Storage Pools

10.2.3 The Easy Tier Process

The Easy Tier feature consists of four main processes. These processes make sure that the extent allocation in Multi Tiered Storage Pools is optimized for the best performance as monitored on your workload in the last 24 hours. Be aware that the 24 hours is not a fixed point in time each day, it is an ongoing window which always moves.

The four main processes are:

- 1. I/O Monitoring: IOM ("I/O Monitoring IOM" on page 358)
- 2. Data Placement Advisor: DPA ("Data Placement Advisor DPA" on page 358)
- 3. Data Migration Planer: DMP ("Data Migration Planer DMP" on page 358)
- **4. Data Migrator:** DM ("Data Migrator DM" on page 359)

Figure 10-3 on page 358 shows the Easy Tier processes.

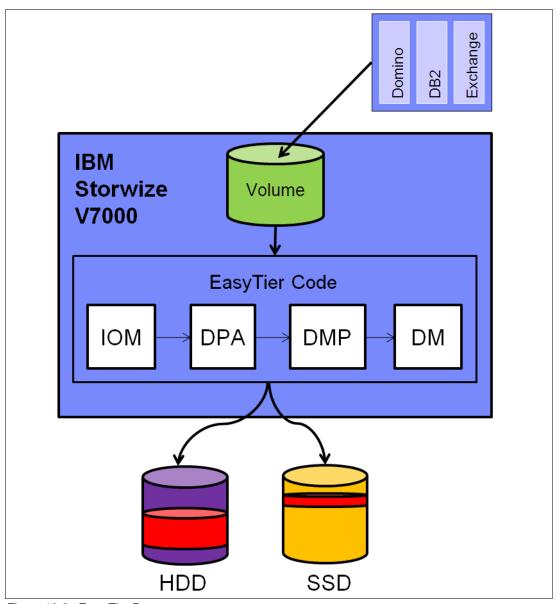


Figure 10-3 Easy Tier Processes

I/O Monitoring - IOM

This process operates continuously on all Easy Tier enabled volumes and monitors and logs the I/O activity. It generates the usage statistics for each extent and derives averages for a rolling 24 hour period of activity. This process is very efficient and adds only negligible processing overhead to the IBM Storwize V7000 canisters.

Data Placement Advisor - DPA

The DPA uses the workload statistics and makes a cost benefit decision on which extents should be candidates for migration to a higher or lower MDisk tier. It outputs migration recommendations to the DMP.

Data Migration Planer - DMP

The DMP performs analysis to deliver a recommended data migration plan to the DM based on the physical storage characteristics

Data Migrator - DM

DM confirms and schedules data migration activity based on the data migration plan using IBM Storwize V7000 migration function to seamlessly relocate extents up to, or down from, the high disk tier as shown in Figure 10-4.

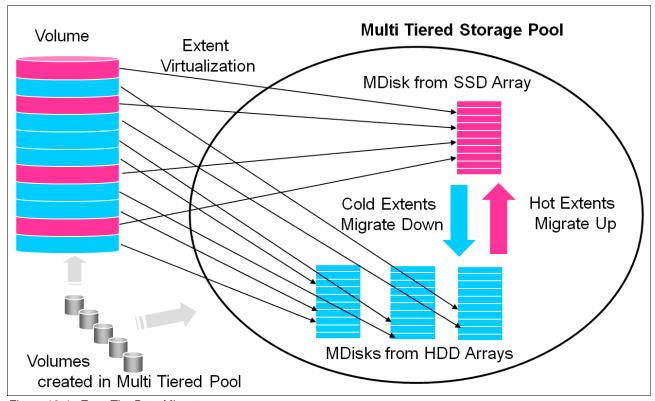


Figure 10-4 Easy Tier Data Migrator

The extent migration is capped so that a maximum of up to 30 MBps is migrated. This equates around 3TB a day that will be at maximum be migrated between disk tiers.

10.2.4 Easy Tier Operating Modes

IBM Storwize V7000 offers three operating modes for Easy Tier:

- ► Easy Tier OFF: ("Easy Tier OFF" on page 359)
- ► Evaluation Mode: ("Evaluation Mode" on page 359)
- ► Auto Data Placement Mode: ("Auto Data Placement Mode" on page 360)

Easy Tier - OFF

Easy Tier can be turned off, no statistics are recorded or extents will be moved.

Evaluation Mode

If you turn on Easy Tier in a Single Tiered Storage pool, it runs in evaluation mode. This means it measures the I/O activity for all extents. A statistic summary file is created and can be off-loaded from the IBM Storwize V7000 as described later in 10.4.1, "Enable Easy Tier Evaluation Mode" on page 372. This file can be analyzed with the IBM Storage Tier Advisory Tool as shown in 10.5, "IBM Storage Tier Advisor Tool" on page 377. This will give you an

understanding about the benefits for your workload if you were to add SSDs to your pool, prior to any hardware acquisition.

Auto Data Placement Mode

This operating mode is enabled by default if you create a Multi Tiered Storage Pool. On all volumes inside the Multi Tiered Storage Pool Easy Tier is also set to ON. The extents will be migrated dynamically be the Easy Tier processes to achieve the best performance. The movement is transparent to the host server and applications, and it will only provide increased performance.

If you do want to disable Auto Data Placement Mode for single volumes inside a Multi Tiered Storage Pool it is possible to turn the mode off at a volume level. This will exclude it from Auto Data Placement Mode and will measure the I/O statistics only.

The statistic summary file can be off-loaded for input to the advisor tool. The tool will produce a report on the extents moved to SSD and a prediction of performance improvement that could be gained if more SSD was available.

10.2.5 Easy Tier Rules

Some operating rules exist when using IBM System Storage Easy Tier on the IBM Storwize V7000.

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.

Note: Volume mirroring can have different workload characteristics on each copy of the data as reads are normally directed to the primary copy and writes occur to both. Thus the amount of extents that Easy Tier will migrate to SSD tier will probably be different for each copy.

- ► Easy Tier is supported to work with all striped volumes, which includes:
 - Generic Volumes
 - Thin-provisioned Volumes
 - Mirrored Volumes
 - Thin-Mirror Volumes
 - Global and Metro Mirror source and target
- ► 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.
- ► If possible, IBM Storwize V7000 creates new volumes or volume expansions using extents from MDisks from the HDD tier, but will use extents from MDisks from the SSD tier if no HDD space is available.
- ▶ When a volume is migrated out of a storage pool that is managed with Easy Tier, automatic data placement mode is no longer active on that volume. Automatic data placement is also turned off while a volume is 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.
- ► SSD performance is dependent on block sizes, and small blocks perform much better than larger ones. As Easy Tier is optimized to work with SSD it decides if an extent is hot by

measuring I/O smaller than 64KB, but it will migrate the entire extent to the appropriate disk tier.

- As extents are migrated, the use of smaller extents makes Easy Tier more efficient.
- ► The first migration of hot data to SSD will start about 1 hour after enabling Auto Data Placement Mode, but it will take up to 24 hours to achieve optimal performance.
- ▶ In the current IBM Storwize Easy Tier V7000 implementation it will take about two days until hot spots are considered to be moved from SSDs, this prevents a case where hot spots are being moved from SSDs if the workload changes over a weekend.
- ▶ If you run an unusual workload over a longer time period, Auto Data Placement can be turned off and on online, to avoid data movement.

Depending on which storage pool and which Easy Tier configuration is set, a volume copy can have the Easy Tier states as shown in Table 10-1:

| Table 10-1 Easy Hel States | | | |
|----------------------------|--------------------------------------|------------------------------------|--------------------------------|
| Storage Pool | Single or Multi Tier Storage Pool | VolumeCopy Easy Tier Setting | Easy Tier Status |
| Off | Single | Off | None |
| Off | Single | On | None |
| Off | Multi | Off | None |
| Off | Multi | On | None |
| Auto | Single | Off | None |
| Auto | Single | On | None |
| Auto | Multi | Off | Evaluation |
| Auto | Multi | On | Evaluation + AutoDataPlacement |
| On | Single | Off | Evaluation |
| On | Single | On | Evaluation |
| On | Multi | Off | Evaluation |
| On | Multi | On | Evaluation + AutoDataPlacement |

Table 10-1 Easy Tier States

10.3 Easy Tier Configuration using the GUI

This section describes how to activate Easy Tier using the IBM Storwize V7000 GUI.

10.3.1 Creating Multi Tiered Pools - Enable Easy Tier

In this section we describe how to create Multi Tiered Storage pools using the GUI. As soon as a storage pool changes from single tier to multi tier, Easy Tier is enabled by default for the pool and on all volume copies inside this pool. Figure 10-5 shows that in our example two single tiered storage pools exist and Easy Tier is inactive. Now SSD drives have to be added to the SAS storage pool to enable Easy tier.



Figure 10-5 Single Tiered Pool

Figure 10-6 on page 362 shows that two internal SSD drives are available and they are in "Candidate" Status. Internal SSD are added to the high performance tier automatically by IBM Storwize V7000.

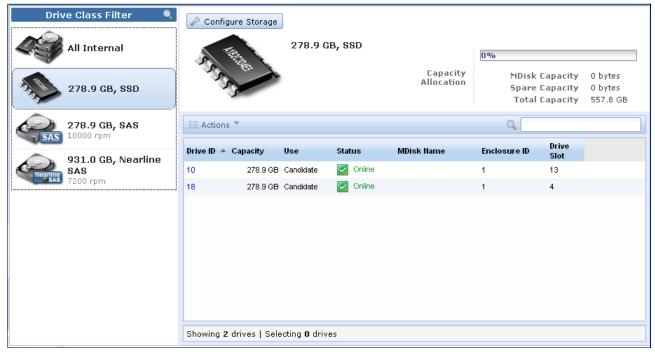


Figure 10-6 Internal SSD

Click **Configure Storage** (Figure 10-6 on page 362) and the Storage Configuration wizard will appear as shown in Figure 10-7.

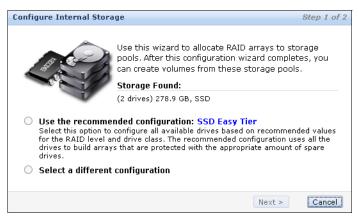


Figure 10-7 Configure Internal Storage

The wizard recommends using the SSD drives to enable Easy Tier. If you select "Use recommended configuration" it will select the recommended RAID level and hot spare coverage for your system automatically. If you choose "Select a different configuration" as shown in Figure 10-8 on page 363 you can select the preset.



Figure 10-8 Select a Preset

Choose a custom RAID level, or you can also select the SSD Easy Tier preset there to review and modify the recommended configuration. In our example we select SSD Easy Tier and the view will be expanded as shown in Figure 10-9.

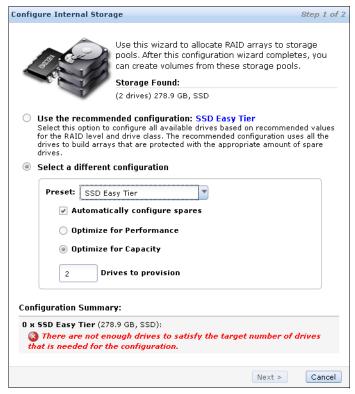


Figure 10-9 Select SSD Easy Tier

In the IBM Storwize V7000 used for this example, there are only two SSDs installed. The SSD Easy Tier preset tries to choose a RAID level and protect it with hot spare drives. As this is not possible with two drives only we get an Error message. To resolve it, unmark the "Automatically configure spares" option to create a RAID without spare drives as shown in Figure 10-10 on page 365 and click **Next**.

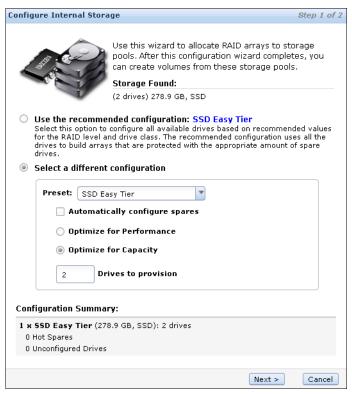


Figure 10-10 Review Preset Configuration

To create a Multi Tiered Storage pool, the SSD drives have to be added to an existing generic HDD pool. Select **Expand an existing Pool** as shown in Figure 10-11 and select the pool you want to change to a multi tiered storage pool. In the example the "SAS" storage pool is selected, click **Finish**.

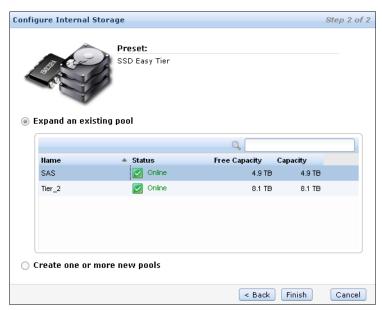


Figure 10-11 Expand an existing Pool

Now the array is configured on the SSD drives, and they are added to the selected storage pool. Click **Close** after the task has completed as shown in Figure 10-12 on page 366.

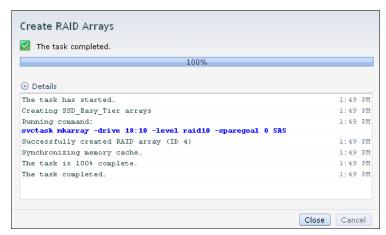


Figure 10-12 Complete Task

Figure 10-13 shows that the internal SSD drives usage has now changed to "Member" and that the wizard created an MDisk name of "mdisk4".

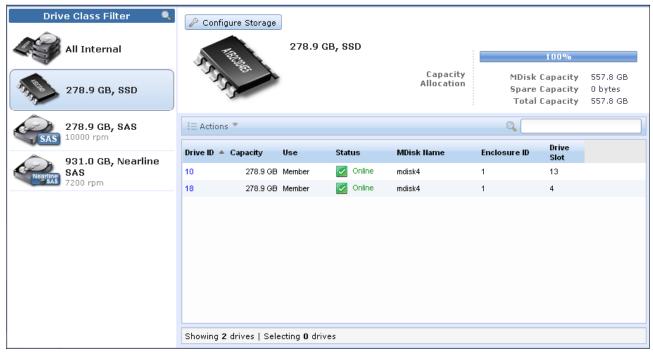


Figure 10-13 SSD Use has changed

In Figure 10-14 on page 367 you can see that the new MDisk is now part of the storage pool "SAS" and that the status of the Easy Tier has changed to active. In this pool, Auto Placement Mode is started and the Easy Tier processes start to work.

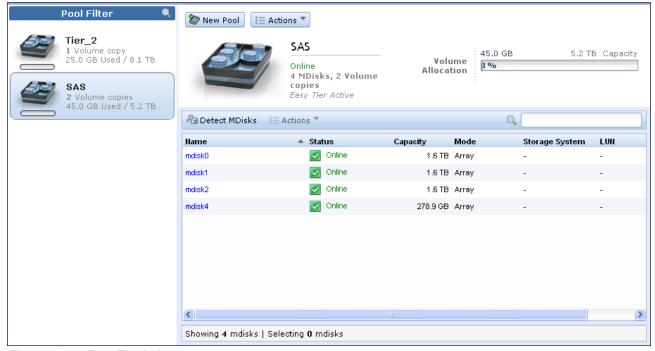


Figure 10-14 Easy Tier Active

Now the storage pool has been successfully changed to a Multi Tiered Storage Pool and Easy Tier has been activated by default. To reflect this change better we rename the storage pool ("Changing the Storage Pool Name" on page 333) and change the icon ("Changing the Storage Pool Icon" on page 332) as shown in Figure 10-15.



Figure 10-15 Multi Tier Storage Pool

By default Easy Tier is now active in this pool and for all volumes residing in this storage pool. Figure 10-16 on page 368 shows that in the example two volumes reside on the Multi Tier

Storage Pool. If you open the properties of a volume (Right Click or Actions \rightarrow Properties) you can also see that Easy Tier is enabled on the volume by default as shown in Figure 10-17.

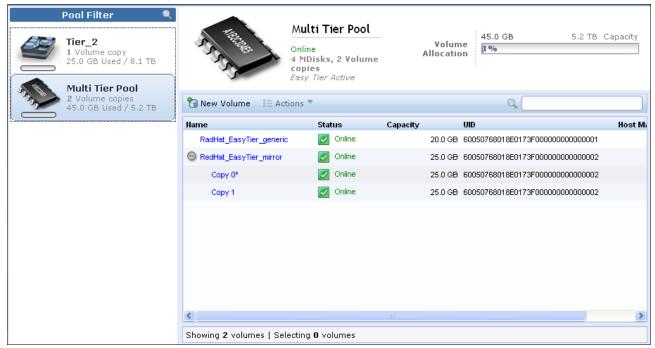


Figure 10-16 Volumes by Pool

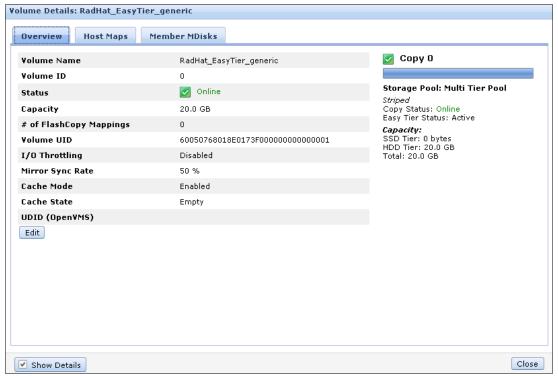


Figure 10-17 Easy Tier enabled volume

If a volume owns more than one copy, Easy Tier can be enabled and disabled on each copy separately as shown in Figure 10-18.

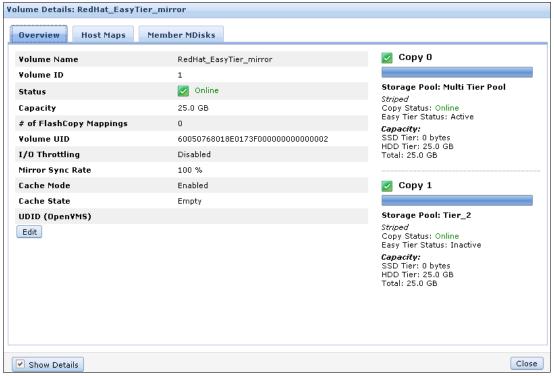


Figure 10-18 Easy Tier by Copy

If you would like to enable Easy Tier on the second copy, you have to also change the storage pool of the second copy to a Multi Tier Storage Pool by repeating these steps.

If external SSD is used it is required to select the tier manually as described in 7.2.4, "Select the tier for MDisks" on page 261 and then add the external SSD MDisk to a storage pool as covered in 7.2.1, "Add MDisks to Storage Pools" on page 242. This will also change the storage pools to Multi Tier Storage Pools and enable Easy Tier on the pool and the volumes.

10.3.2 Downloading Easy Tier I/O measurements

As described before, Easy Tier has now been enabled and Auto Data Placement Mode is active now. Extents will be automatically migrated to, or from, high performance disk tiers, and the statistic summary collection is now active. The statistics log file can be downloaded to analyze how many extents have been migrated, and to monitor if it would make sense to add more SSD to the Multi Tiered Pool.

To download the statistics file open the support view as shown in Figure 10-19 on page 370.



Figure 10-19 Open Support Section

Click **Show full log** listing as shown in Figure 10-20 to display all available logs.



Figure 10-20 Show full log listing

All available logs are displayed as shown in Figure 10-21 on page 371. The Easy Tier log files are always named "dpa_heat.canister_name_date.time.data".

Note: Depending on your workload it can take up to 24 hours until a new Easy Tier log file is created.

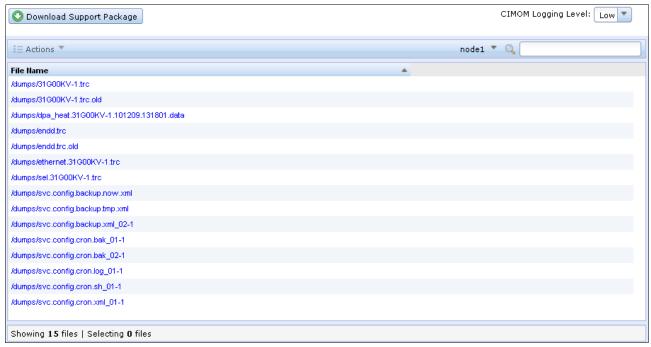


Figure 10-21 Show full log listing

To make the search easier type **heat** into the search field and press Enter as shown in Figure 10-22.

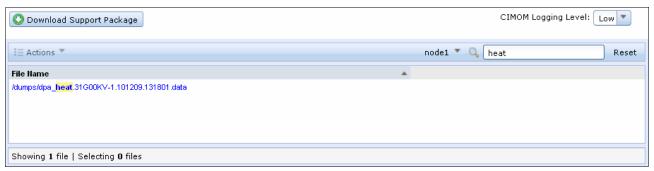


Figure 10-22 Filter Heat Files

If you run Easy Tier for a longer time period, it will generate a heat file at least every 24 hours. The time and date of the file creation is included in the filename. The heat log file always includes the measured I/O activity of the last 24 hours. Select the file for Easy Tier measurement for the most representative time period, right click it and select **Download** as shown in Figure 10-23.



Figure 10-23 Download Heat File

Depending on your browser settings the file will be downloaded to your default location, or you will be asked to save it to your computer. This file can be analyzed as described in 10.5, "IBM Storage Tier Advisor Tool" on page 377.

10.4 Easy Tier Configuration using the CLI

Section 10.3, "Easy Tier Configuration using the GUI" showed how to enable IBM Storwize V7000 Easy Tier using the GUI. Easy Tier can also be configured using the CLI, and for the advanced and more confident user this will offer several more options for Easy Tier Configuration.

Before using the CLI you have to configure CLI access as described in Appendix A, "CLI setup and SAN Boot" on page 543.

Note: In most examples shown in this section, many lines have been deleted in the command output or responses. The aim being to concentrate on the Easy Tier related information only.

10.4.1 Enable Easy Tier Evaluation Mode

If you want to enable Easy Tier in evaluation mode you have to enable Easy Tier on a Single Tiered Storage Pool. Connect to your IBM Storwize V7000 using the CLI and type **svcinfo 1 smdiskgrp** as shown in Example 10-1. This will give an overview about all configured storage pools and the Easy Tier Status of the pool. In the example the are two storage pools listed. Tier_2 with Easy Tier inactive, and the Multi Tier Pool with Easy Tier enabled.

Example 10-1 List Storage Pools

To get a more detailed view about the Single Tier storage pool type svcinfo lsmdiskgrp "storage pool name" as shown in Example 10-2.

Example 10-2 Storage Pools details - Easy Tier OFF

IBM 2076:ras-stand-7-tb1:admin>svcinfo lsmdiskgrp Tier 2

```
id 0
name Tier 2
status online
mdisk count 1
vdisk_count 1
easy tier auto
easy tier status inactive
tier generic ssd
tier mdisk count 0
tier capacity 0.00MB
tier free capacity 0.00MB
tier generic hdd
tier mdisk count 1
tier capacity 8.14TB
tier free capacity 8.11TB
IBM 2076:ras-stand-7-tb1:admin>
```

To enable Easy Tier on the Single Tier Storage Pool type svctask chmdiskgrp -easytier on "Storage Pool name" as shown in Example 10-3.

Example 10-3 Enable Easy Tier on Single Tier Storage Pool

```
IBM_2076:ras-stand-7-tb1:admin>svctask chmdiskgrp -easytier on Tier_2
IBM_2076:ras-stand-7-tb1:admin>
```

Check the status of the storage pool again by repeating the svcinfo 1smdiskgrp "storage pool name" command as shown in Example 10-4.

Example 10-4 Storage Pools details - Easy Tier ON

```
IBM 2076:ras-stand-7-tb1:admin>svcinfo lsmdiskgrp Tier 2
id 0
name Tier 2
status online
mdisk count 1
vdisk_count 1
easy tier on
easy tier status active
tier generic ssd
tier mdisk count 0
tier capacity 0.00MB
tier free capacity 0.00MB
tier generic_hdd
tier mdisk count 1
tier capacity 8.14TB
tier free capacity 8.11TB
IBM 2076:ras-stand-7-tb1:admin>
```

Also repeat the **svcinfo 1smdiskgrp** command as shown in Example 10-5 and you will see that Easy Tier is turned on the storage pool now, but that Auto Data Placement Mode is not active on the Multi Tier Storage Pool.

Example 10-5 List Storage Pools

Type svcinfo 1svdisk to get a list of all volumes as shown in Example 10-6.

Example 10-6 List All Volumes

To get a more detailed view about a volume, use svcinfo 1svdisk "volume name" as shown in Example 10-7 on page 374. The example output shows two copies of a volume. Copy 0 is in a Multi Tier Storage Pool and Auto Data Placement is active, Copy 1 is in the Single Tier Storage Pool and Easy Tier Evaluation mode is active, this is indicated by the line "easy tier status measured".

Example 10-7 volume details

```
IBM 2076:ras-stand-7-tb1:admin>svcinfo lsvdisk RedHat EasyTier mirror
id 1
name RedHat EasyTier mirror
mdisk grp id many
mdisk_grp_name many
type many
copy count 2
copy id 0
status online
sync yes
primary yes
mdisk grp id 3
mdisk grp name Multi Tier Pool
type striped
easy tier on
easy tier status active
tier generic ssd
tier capacity 512.00MB
tier generic hdd
tier_capacity 24.50GB
copy id 1
status online
sync yes
primary no
mdisk grp id 0
mdisk grp name Tier 2
```

type striped
easy_tier on
easy_tier_status measured
tier generic_ssd
tier_capacity 0.00MB
tier generic_hdd
tier_capacity 25.00GB
IBM 2076:ras-stand-7-tb1:admin>

These changes are also reflected in the GUI as shown in Figure 10-24 on page 375.

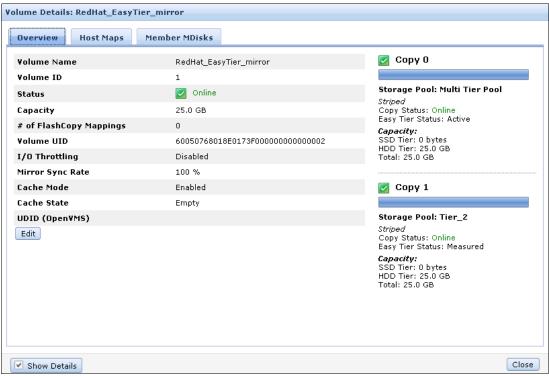


Figure 10-24 Easy Tier Status by volume

Easy Tier Evaluation Mode is now active on the Single Tier Storage Pools. To download the I/O statistics and analyze it go to 10.3.2, "Downloading Easy Tier I/O measurements".

10.4.2 Enable or Disable Easy Tier on Single Volumes

If you enable Easy Tier on a storage pool, by default all volume copies inside the pools Easy Tier also have it enabled. This applies both to Multi Tier and Single Tier Storage Pools. But it is also possible to turn Easy Tier on and off for single volume copies.

To disable Easy Tier on single volumes type svctask chvdisk -easytier off "volume name" as shown in Example 10-8.

Example 10-8 Disable Easy Tier on Single Volume Copy

IBM_2076:ras-stand-7-tb1:admin>svctask chvdisk -easytier off
RedHat EasyTier mirror

IBM 2076:ras-stand-7-tb1:admin>

This will disable Easy Tier on all copies of this volume. Example 10-9 shows that the Easy Tier status of the copies has changed, even if Easy Tier is still enabled on the storage pool.

Example 10-9 Easy Tier Disabled

```
IBM 2076:ras-stand-7-tb1:admin>svcinfo lsvdisk RedHat EasyTier mirror
id 1
name RedHat_EasyTier_mirror
mdisk_grp_id many
mdisk grp name many
type many
copy_count 2
copy_id 0
status online
sync yes
primary yes
mdisk_grp_id 3
mdisk_grp_name Multi Tier Pool
type striped
easy tier off
easy_tier_status measured
tier generic ssd
tier_capacity 512.00MB
tier generic hdd
tier_capacity 24.50GB
copy_id 1
status online
sync yes
primary no
mdisk grp id 0
mdisk_grp_name Tier_2
type striped
easy_tier off
easy_tier_status measured
tier generic_ssd
tier capacity 0.00MB
tier generic_hdd
tier capacity 25.00GB
IBM_2076:ras-stand-7-tb1:admin>
```

To enable Easy Tier on a volume copy type svctask chvdisk -easytier on "volume name" as shown in Example 10-10 in and the Easy Tier Status will change back to enabled as shown in Example 10-7 on page 374.

Example 10-10 Easy Tier enabled

IBM_2076:ras-stand-7-tb1:admin>svctask chvdisk -easytier on RedHat_EasyTier_mirror
IBM 2076:ras-stand-7-tb1:admin>

10.5 IBM Storage Tier Advisor Tool

The IBM Storage Tier Advisor Tool (STAT) is a command line tool that runs on Windows based systems. If you run Easy Tier in Evaluation Mode it will analyze your extents and estimate how much benefit you will derive if you implement Easy Tier Auto Data Placement with various amounts of SSD. If Auto Data Placement Mode is already active, the analysis will also include an overview of migrated heat data. and recommendations as to whether you could derive any benefit by adding more SSD drives. The output provides a graphical representation of the performance data collected by Easy Tier over a 24 hour operational cycle.

10.5.1 Creating graphical reports

The STAT takes input from the dpa_heat log file and produces an HTML file that contains the report. Download the heat_log file as described in section 10.3.2, "Downloading Easy Tier I/O measurements" on page 369 and save it to the hard drive of a Windows system.

You can read more about the tool, and download it, at the IBM website:

http://www-01.ibm.com/support/docview.wss?uid=ssg1S4000935

Click **Start** \rightarrow **Run** type **cmd** and click **OK** to open a command prompt.

Enter the following command to generate the report (Example 10-11):

STAT.exe -o c:\directory_where_you_want_the output_to_go c:\location_of_dpa_heat_data_file

If you do not specify -o c:\directory_where_you_want_the output_to_go the output will go to the directory where the STAT.exe file is located.

Example 10-11 Generate HTML File

C:\EasyTier>STAT.exe -o C:\EasyTier C:\StorwizeV7000_Logs\dpa_heat.31G00KV-1.101 209.131801.data

CMUA00019I The STAT.exe command has completed.

C:\EasyTier>

Browse the directory where you directed the output file, and there will be a file named **index.html**. Open it with your browser to view the report.

10.5.2 STAT Reports

If you open the report (ie. the index.html file) of an input file from an IBM Storwize V7000 system which was in Easy Tier Evaluation mode, a view will open which will give you an estimate about the benefit if you were to add SSD drives as shown in Figure 10-25 on page 378.

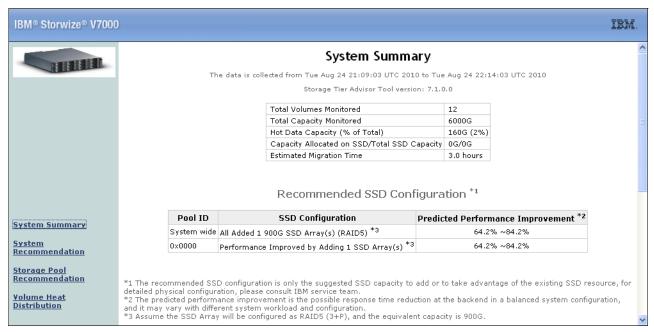


Figure 10-25 STAT Report - System Summary

The System Summary already provides the most important numbers. In Figure 10-25 on page 378 it is shown that 12 volumes have been mirrored with a total capacity of 6000GB. The results of the analysis of the hot extents is that about 160GB (which means about 2%) should be migrated to the high performance disk tier. It also recommends that one SSD Array should be added as a high performance tier consisting of 4 SSD drives (RAID 5 3+P). This predicted performance improvement in response time reduction at the backend in a balanced system would be between 64% and 84%.

Click Volume Heat Distribution to change to a more detailed view as shown in Figure 10-26.

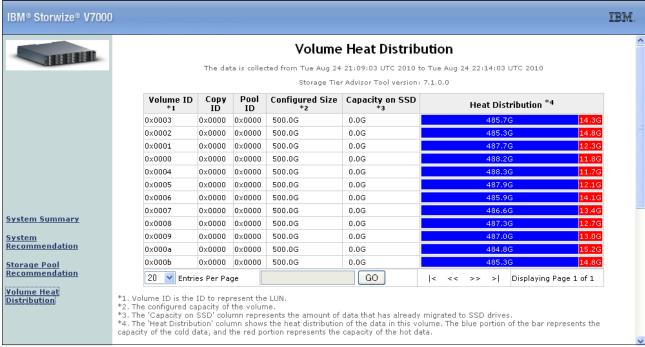


Figure 10-26 Volume Heat Distribution

The table in this view gives you a more detailed view as to how the hot extents are distributed across your system. It contains the following information:

- Volume ID: Is the unique ID of each Volume on IBM Storwize V7000.
- Copy ID: If a volume owns more than one copy the data is measured for each copy.
- Pool ID: Is the unique ID of each Pool configured on IBM Storwize V7000
- Configured Size: Is the configured size of each volume that is represented to the host.
- Capacity on SSD: Capacity of the volumes on high performance disk tier (even in evaluation mode volumes can reside on high performance disk tiers if they have been moved there before).
- ► **Heat Distribution**: Shows the heat distribution of the data in this volume. The blue portion of the bar represents the capacity of the cold data, and the red portion represents the capacity of the hot data. The red hot data are candidates to be moved to high performance disk tier.





Copy Services

In this chapter we describe the copy services functions provided by IBM Storwize V7000 storage system, including FlashCopy and Remote Copy. Copy services functions are useful to make data copies for backup, application test, recovery and so on. The IBM Storwize V7000 makes it easy to apply these functions to your environment using the intuitive GUI.

11.1 FlashCopy

In this section we describe how the FlashCopy function works in IBM Storwize V7000 storage system and the implementation steps using GUI are provided for FlashCopy configuration and management.

FlashCopy, also known as point-in-time copy, can be used to help creating immediately available consistent copies of data sets while they remain online and actively in use.

The FlashCopy is performed at the block level on IBM Storwize V7000 storage system, which operates below the host operating system and is therefore transparent to the host. However, to ensure the integrity of the copy that is made, it is necessary to flush the host cache for any outstanding reads or writes prior to performing the FlashCopy operation. Failing to do so will produce what's referred to as a crash consistent copy, meaning the resulting copy will require the same type of recovery procedure (such as log replay and filesystem checks) as would be required following a host crash.

Some operating systems and applications provide facilities to stop I/O operations and ensure all data is flushed from host cache. If these facilities are available they can be used to prepare and start a FlashCopy operation. When this type of facility is not available the host cache must be flushed manually by quiescing the application and unmounting the filesystem or drives.

With the immediately available copy of data, FlashCopy could be used in various business scenarios, which include:

- Creating consistent and fast backups of application data.
 - The backup data copies can be taken by FlashCopy periodically and the target volumes can be used to preform a rapid restore of individual files or the entire volume.
 - The data copy created by FlashCopy may be also used for backup to tape, by attaching them to another server, which to a great extent off-loaded the production server. Once the copy to tape has been completed the target volumes can be discarded if required.
- Creating data copies for moving or migrating data.
 - FlashCopy can be used to facilitate the movement or migration of data between hosts while minimizing downtime for applications. FlashCopy will allow application data to be copied from source volumes to new target volumes while applications remain online. Once the volumes are fully copied and synchronized the application can be brought down and then immediately brought back up on the new server accessing the new FlashCopy target volumes.
- Creating copies of production datasets for development/testing, auditing purposes or data mining.
 - It is usually required to have the copies of actual production data for development/testing, auditing purposes or data mining. FlashCopy makes this easy to accomplish without putting the production data at risk or requiring downtime to create a constant copy.

11.1.1 FlashCopy Concepts

FlashCopy occurs between a source volume and a target volume. The source and target volumes must be the same size. 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.

The source and target volumes must belong to the same IBM Storwize V7000 storage system, however they do not have to be in the storage pool.

FlashCopy works by defining a FlashCopy mapping that consists of one source volume together with one target volume. Multiple FlashCopy mappings can be defined, and point-in-time consistency can be maintained across multiple individual mappings using consistency groups, which would be described in "FlashCopy Consistency Groups" on page 386.

When 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 (T0).

Immediately following the FlashCopy operation both the source and target volumes are available for use. The FlashCopy operation creates a bitmap which is referenced and maintained to direct I/O requests within the source and target relationship. This bitmap is updated to reflect the active block locations as data is copied in the background from the source to target and updates are made to the source.

When data is copied between volumes it is copied in units of address space known as grains. The FlashCopy bitmap contains one bit for each grain and is used to keep track of whether the source grain has been copied to the target.

Figure 11-1 illustrates the redirection of the host I/O toward the source volume and the target volume.

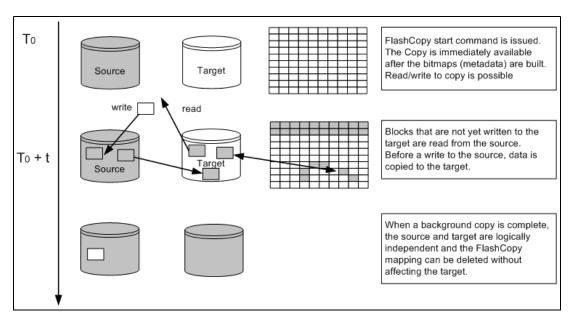


Figure 11-1 Redirection of host I/O

The FlashCopy bitmap dictates read and write behavior for both the source and target volumes, as below:

► Read I/O request to source:

Reads are performed from the source volume, this is the same as for non-FlashCopy volumes.

► Write I/O request to source:

Writes to the source will cause the grains to be copied to the target if it has not already been copied, then the write will be performed to the source.

► Read I/O request to target:

Reads are performed from the target if the grains has already been copied, otherwise the read is performed from the source.

Write I/O request to target:

Writes to the target will cause the grain to be copied from the source to the target unless the entire grain is being written, then the write will complete 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 or a member of a consistency group, as described in "FlashCopy Consistency Groups" on page 386.

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 reduces the amount of time that it takes to re-create an independent FlashCopy image. You can define a FlashCopy mapping as incremental only when you create the FlashCopy mapping.

Multiple target FlashCopy Mappings

You can copy up to 256 target volumes from a single source volume. Each relationship between a source and target volume is managed by a unique mapping such that a single volume can be the source volume in 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 have 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 will refer to new targets first before referring to the source. A dependency relationship exists between a particular target and all newer targets that share the same 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.

Background Copy

The background copy rate is a property of a FlashCopy mapping defined as a value between 0 and 100. The background copy rate can be defined and dynamically changed for individual FlashCopy mappings. A value of 0 disables background copy.

With FlashCopy background copy, the source volume data will be copied to the corresponding target volume in the FlashCopy mapping. If the background copy rate is set to 0, which means to disable the FlashCopy background copy, only data that changed on the source volume will be copied to the target volume. The benefit of using a FlashCopy mapping with background

copy enabled is that the target volume becomes a real independent clone of the FlashCopy mapping source volume once 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. Although on the other hand, copying only the changed data saves your storage capacity (assuming it is thin provisioned and -rsize has been correctly setup.

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 Rate value to data copied per second

| Value | Data copied per second |
|----------|------------------------|
| 1 - 10 | 128 KB |
| 11 - 20 | 256 KB |
| 21 - 30 | 512 KB |
| 31 - 40 | 1 MB |
| 41 - 50 | 2 MB |
| 51 - 60 | 4 MB |
| 61 - 70 | 8 MB |
| 71 - 80 | 16 MB |
| 81 - 90 | 32 MB |
| 91 - 100 | 64 MB |

Cleaning Rate

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 between 0 and 100, which has the same relationship to data copied per second with the backup copy rate, as shown in Table 11-1.

The cleaning rates controls the rate at which the cleaning process operates. The cleaning process copies data from the target volume of a mapping to the target volumes of other mappings that are dependent on this data. The cleaning process must complete before the FlashCopy mapping can go to the stopped state.

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 the source and the target volumes.

If the mapping is incremental and the background copy is complete, 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 will be offline.

Copying

The copy is in progress. Read and write caching is enabled on the source and the target volumes.

Prepared

The mapping is ready to start. The target volume is online, but is not accessible. The target volume cannot perform read or write caching. Read and write caching is failed by the SCSI front end as a hardware error. If the mapping is incremental and a previous mapping 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 would be described in more detail in , "Dependent Writes" on page 387, are executed in the application's intended sequence.

When consistency groups are used the FlashCopy commands are issued to the FlashCopy consistency group which preforms the operation on all FlashCopy mappings contained within the consistency group.

Figure 11-2 on page 387 illustrates a consistency group consisting of two FlashCopy mappings.

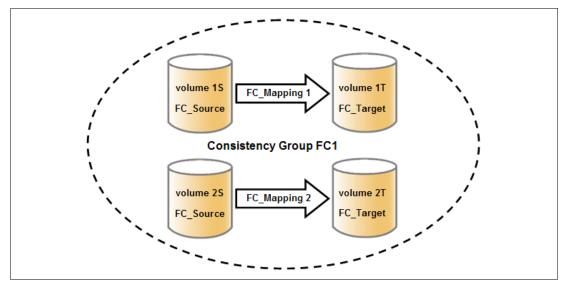


Figure 11-2 FlashCopy consistency group

Note: After an individual FlashCopy mapping has been added to a consistency group it can only be managed as part of the group, operations such as start and stop are no longer allowed on the individual mapping.

Dependent Writes

To illustrate why it is crucial to use consistency groups when a data set spans multiple volumes, consider the following typical sequence of writes for a database update transaction:

- 1. A write is executed to update the database log indicating that a database update is about to be performed.
- 2. A second write is executed to perform 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 prior to the FlashCopy of the database log. This 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, since 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, it is necessary to perform a FlashCopy operation on multiple volumes as an atomic operation with the use of consistency groups.

A FlashCopy consistency group can contain up to 512 FlashCopy mappings. FlashCopy commands can then be issued to the FlashCopy consistency group and thereby simultaneously for all of the FlashCopy mappings that are defined in the consistency group. For example, when start the FlashCopy for the consistency group, all of the FlashCopy mappings in the consistency group are started at the same time, resulting in a point-in-time copy that is consistent across all of the FlashCopy mappings that are contained in the consistency group.

It is important to note that 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. If the source volume with multiple target volumes is in the same consistency group, the resulting FlashCopy will produce 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 having to wait for the original copy operation to complete. It supports multiple targets and thus multiple rollback points.

A key advantage of reverse FlashCopy function is that it does not destroy 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 to be made prior to starting the reverse copy operation. This provides the ability to restore back to the original source data which can be useful for diagnostic purposes.

Figure 11-3 shows an example of the reverse FlashCopy scenario.

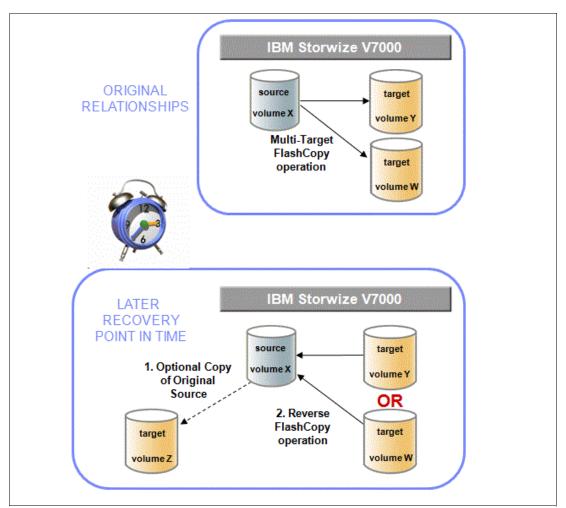


Figure 11-3 Reverse FlashCopy scenario

The following outlines the steps required to restore from an FlashCopy backup

- 1. (Optional) Create a new target volume (volume Z) and FlashCopy the production volume (volume X) onto the new target for later problem analysis.
- 2. Create a new 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 already exist.
- 3. Start the FlashCopy map (volume Y -> volume X) with the -restore option to copy the backup data onto the production disk.

Note: In GUI, the -restore option will be applied automatically when you start the FlashCopy mapping from volume Y to volume X; in the CLI, you need to add **-restore** option to the command manually. You will find more information on using the CLI in Appendix A, "CLI setup and SAN Boot" on page 543.

4. The production disk is instantly available with the backup data.

Note that regardless of whether the initial FlashCopy map (volume X -> 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.

11.1.2 Planning for FlashCopy

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.
- ► FlashCopy source and target volumes can be thin-provisioned.
- ► The source and target volumes must be the same size. The size of the source and target volumes cannot be altered (increased or decreased) while a FlashCopy mapping is defined.
- ► FlashCopy operations perform in direct proportion to the performance of the source and target disks. If you have a fast source disk and slow target disk, the performance of the source disk is reduced because it has to wait for the write operation to occur at the target before it can write to the source.

Maximum Configurations for FlashCopy

Table 11-2 shows some of the FlashCopy maximum configurations.

Table 11-2 FlashCopy maximum configurations

| FlashCopy property | Maximum |
|--|---------|
| FlashCopy targets per source | 256 |
| FlashCopy mappings per cluster | 4,096 |
| FlashCopy consistency groups per cluster | 127 |
| FlashCopy mappings per consistency group | 512 |

FlashCopy Presets

The IBM Storwize V7000 storage system provides three FlashCopy presets, named Snapshot, Clone and Backup, to simplify the more common FlashCopy operations, as shown in Table 11-3 on page 391.

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. |

11.1.3 Managing FlashCopy Using the GUI

IBM Storwize V7000 storage system provides a separate function icon to access copy service management. There are three panels for managing FlashCopy under the **Copy Services** function icon, including:

- ► FlashCopy
- ► Consistency Groups
- FlashCopy Mappings

Figure 11-4 on page 392 shows the **Copy Services** function icon.

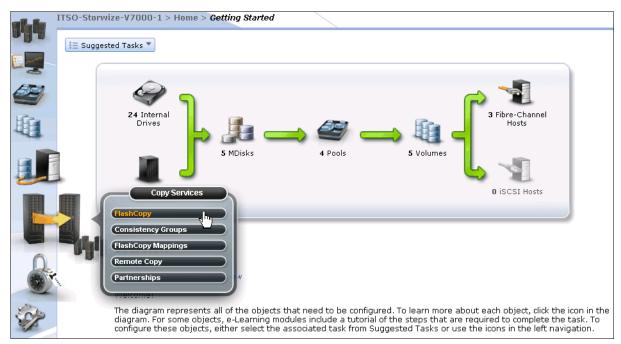


Figure 11-4 Copy Service function icon

Most of actions to manage the FlashCopy mapping could be done both in the **FlashCopy** panel or the **FlashCopy Mappings** panels. Although the quick path to create FlashCopy presets could only be found in **FlashCopy** panel.

Select **FlashCopy** in the **Copy Services** function icon, and **FlashCopy** panel will be displayed as shown in Figure 11-5. In the **FlashCopy** panel, the FlashCopy mappings are organized by volumes.

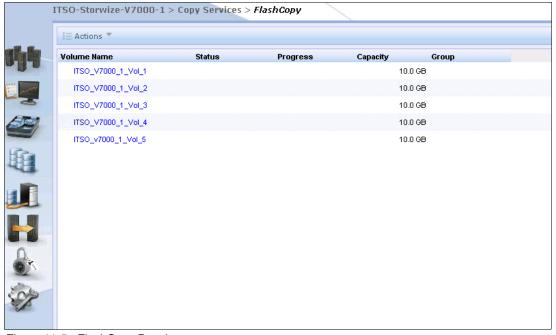


Figure 11-5 FlashCopy Panel

Select FlashCopy Mappings in the Copy Services function icon, and FlashCopy Mappings panel will be displayed as shown in Figure 11-6. In the FlashCopy Mappings panel, the FlashCopy mappings are listed one by one.

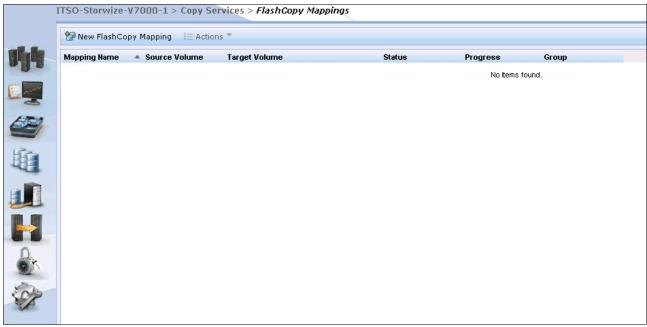


Figure 11-6 FlashCopy Mappings Panel

The consistency groups panel could be used to manage the consistency groups for FlashCopy mappings. Select **consistency groups** in the **Copy Services** function icon, and **consistency groups** panel will be displayed as shown in Figure 11-7.



Figure 11-7 consistency groups panel

Quick Path to Create FlashCopy Presets

It is very easy to create FlashCopy using the presets in the **FlashCopy** panel.

Create Snapshot with one click

In the **FlashCopy** panel, choose a volume and select **New Snapshot** in the right click menu or in the **Actions** drop down list, as shown in Figure 11-8.

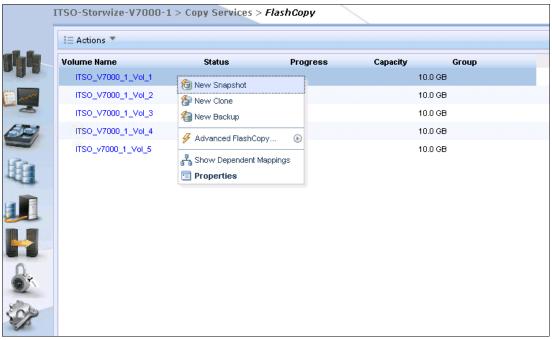


Figure 11-8 Create Snapshot with one click

After only one click, you will have a snapshot volume for the volume you choose.

Create Clone with one click

In the **FlashCopy** panel, choose a volume and select **New Clone** in the right click menu or in the **Actions** drop down list, as shown in Figure 11-9 on page 394.

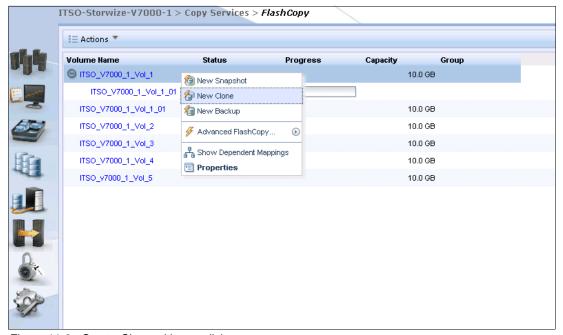


Figure 11-9 Create Clone with one click

After only one click, you will have a clone volume for the volume you choose.

Create Backup with one click

In the **FlashCopy** panel, choose a volume and select **New Backup** in the right click menu or in the **Actions** drop down list, as shown in Figure 11-10.

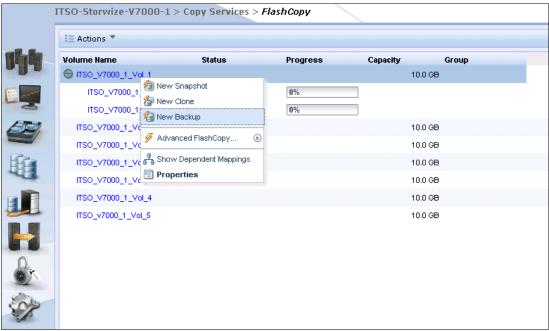


Figure 11-10 Create Backup with one click

After one click, you will have a backup volume for the volume you choose.

Now, in the **FlashCopy** panel, you will find three FlashCopy target volumes under the source volume, as shown in Figure 11-11. The progress bars behind the target volumes indicate the copy progress in percentage. The copy progress remains 0% for snapshot while there is no change happening now since only the changed data would be copied. At the same time, the copy progress for clone and backup keep increasing.

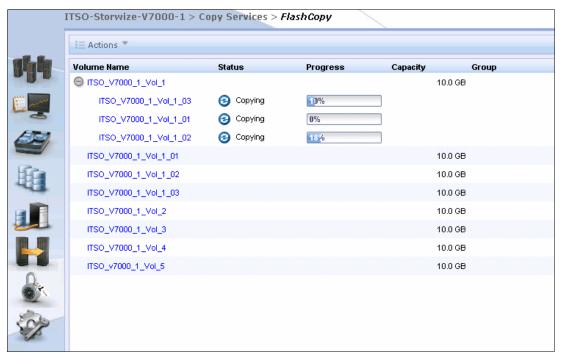


Figure 11-11 FlashCopy in progress in FlashCopy panel

The copy progress could be also found in the status indicator of **Running Tasks**, as shown in Figure 11-12 on page 397.

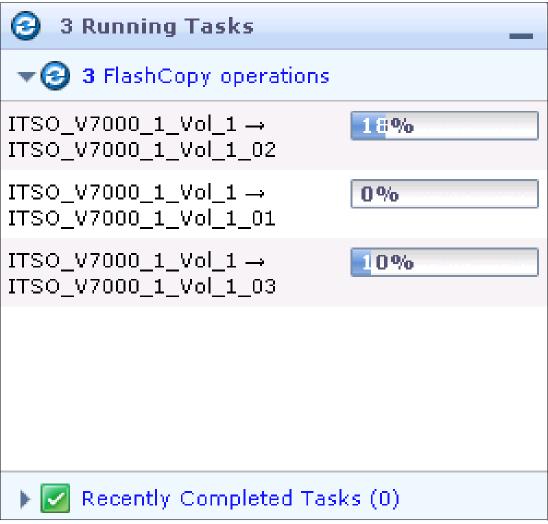


Figure 11-12 FlashCopy progress in Running Tasks

After the copy progresses complete, you will find the FlashCopy with the clone preset has been deleted automatically, as shown in Figure 11-13 on page 398.

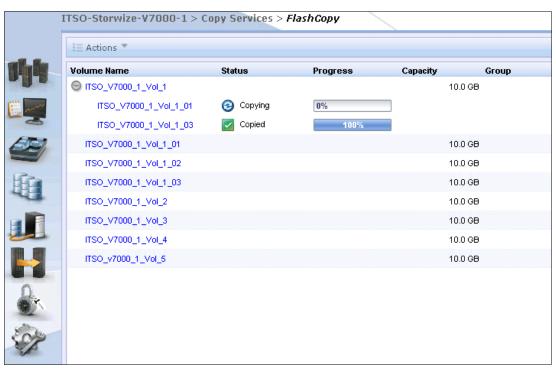


Figure 11-13 Copy progresses complete

11.1.4 Manage FlashCopy Mapping

The FlashCopy presets cover mostly used FlashCopy configuration in general situations. Although, customized FlashCopy mappings are still necessary in some complicated scenarios.

Create FlashCopy Mapping

You can create FlashCopy mappings in the **FlashCopy** panel, select the volume you want to be the source volume in FlashCopy mapping, and choose **Advanced FlashCopy...** option in the right click menu or the **Actions** drop down list, as shown in Figure 11-14 on page 399.

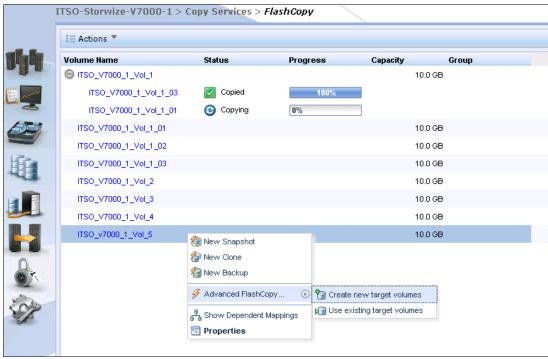


Figure 11-14 Create advanced FlashCopy

Then you have two options:

Select Create new target volumes, if you have not created the target volume yet.

After you select **Create new target volumes**, the wizard will let you choose the preset, as shown in Figure 11-15, but no matter which preset you choose, you can modify the setting of the FlashCopy mapping. Therefore choose one preset which has the most similar configuration to the one required, and click **Advanced Settings** to make any appropriate adjustments on the properties, as shown in Figure 11-16 on page 400.

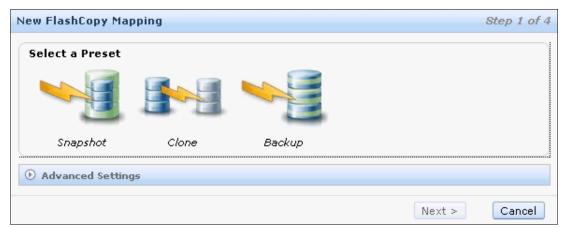


Figure 11-15 Choose a preset most similar to your requirement

For example, if the snapshot preset has been chosen, the default settings can be found when you click **Advanced Settings**, which are:

- ► Background Copy: 0
- Incremental: No

- Auto Delete after completion: No
- ► Cleaning Rate: 0



Figure 11-16 Default setting for snapshot

It is the same with clone preset, and the default settings of clone preset can be found when you click **Advanced Settings**, after you select clone preset, as shown in Figure 11-17 on page 401, which including:

- ► Background Copy: 50
- ► Incremental: No
- ► Auto Delete after completion: Yes
- ► Cleaning Rate: 50

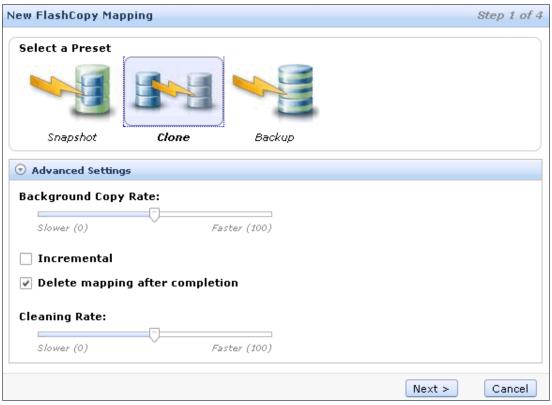


Figure 11-17 Default setting for clone

It is the same with backup preset and the default settings of backup preset could be found when you click **Advanced Settings**, after you select backup preset, as shown in Figure 11-18 on page 402, which includes:

► Background Copy: 50

► Incremental: Yes

► Auto Delete after completion: No

► Cleaning Rate: 50



Figure 11-18 Default setting for backup

Change settings of the FlashCopy mapping according to your requirement, and click **Next**.

In the next step, you could add your FlashCopy mapping to a consistency group, as shown in Figure 11-19. If the consistency group is not ready now, the FlashCopy mapping could be added to consistency group afterwards. Click **Next** to continue.

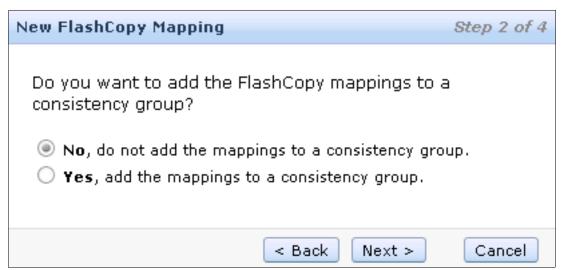


Figure 11-19 Add FlashCopy mapping to consistency group

Then you have the options to choose which storage pool you want to create your volume from. As shown in Figure 11-20 on page 403, you can select the same storage pool that is used by the source volume.

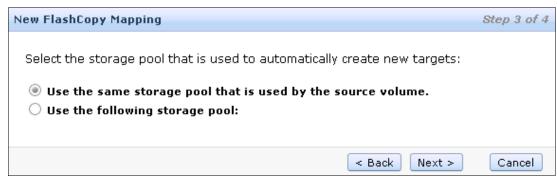


Figure 11-20 Choose to use the same storage pool with the source volume

And you can also specify some other storage pool for your new volume, as shown in Figure 11-21. Click **Next** to continue.

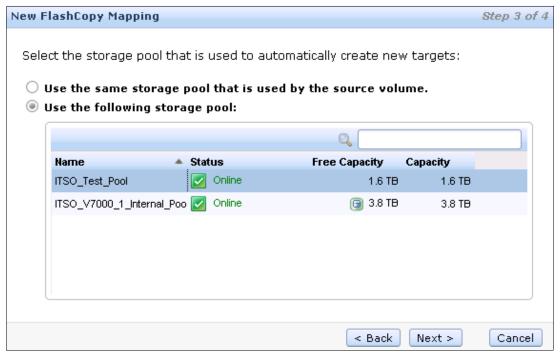


Figure 11-21 Choose other storage pool to create the new volume

In the next step, you will be asked if you want to create a thin provisioning volume. If you have chosen the snapshot preset in the beginning of this wizard, **Yes** would be your default choice. You can specify the detail setting of your thin provisioning volume, as shown in Figure 11-22 on page 404.

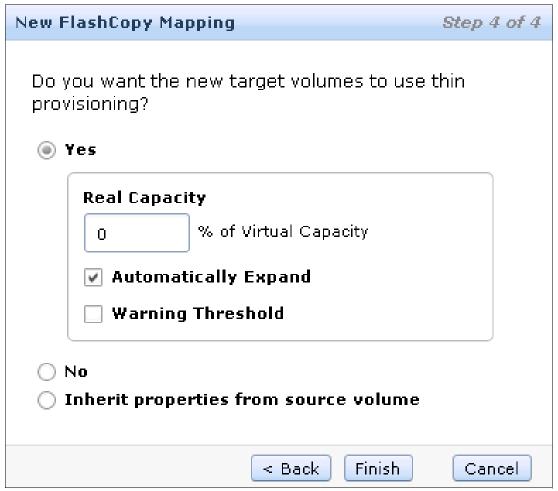


Figure 11-22 Create a thin provisioning target volume

Otherwise, if you have chosen the clone or backup preset at the beginning of this wizard, **Inherit properties from source volume** would be your default choice, as shown in Figure 11-23 on page 405.

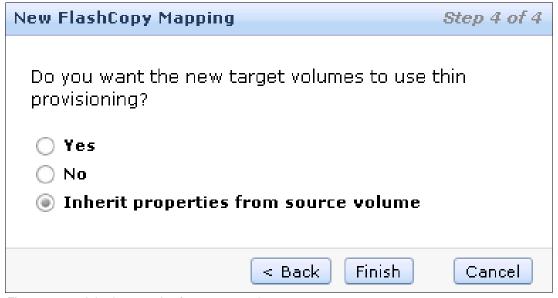


Figure 11-23 Inherit properties from source volume

Click Finish when you make your decision, and the FlashCopy mapping will be created on your volume with a new target, as shown in Figure 11-24. The status of the newly created FlashCopy mapping will be **Idle**, it could be started, as described in, "Start FlashCopy Mapping" on page 408.



Figure 11-24 New FlashCopy mapping has been created with new target

2. In **Advanced FlashCopy...**, if you already have candidate target volumes, select **Use existing target volumes**, as shown in Figure 11-25 on page 406.

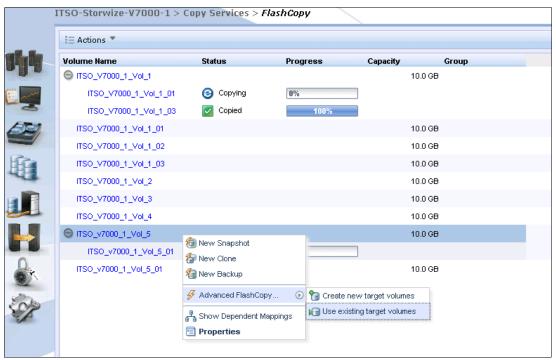


Figure 11-25 Create FlashCopy mapping using existing target volumes

Then you need to choose the target volume for the source volume you selected. Select the target volume in the drop down list on the right of the window, click **Add**, as shown in Figure 11-26.

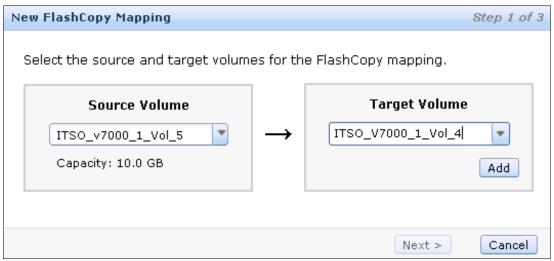


Figure 11-26 Select the target volume

After you click Add, the FlashCopy mapping will be listed as shown in Figure 11-27 on page 407. Click the red cross if the FlashCopy mapping is not the one you want to create. If the FlashCopy mapping is what you want, click **Next** to continue.

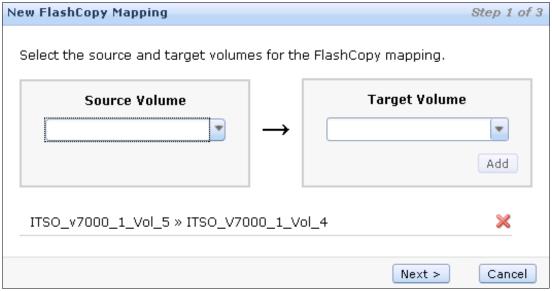


Figure 11-27 Add FlashCopy mapping

In the next step, you can select the preset and make your adjustment to the settings, as shown in Example 11-28. Make sure the settings meet your requirements, and click **Next**.



Figure 11-28 Select preset and make your adjustment to the settings

Now you can add the FlashCopy mapping to a consistency group, as shown in Figure 11-29 on page 408. Click **Finish** and the FlashCopy mapping will be created with the status of **Idle**.

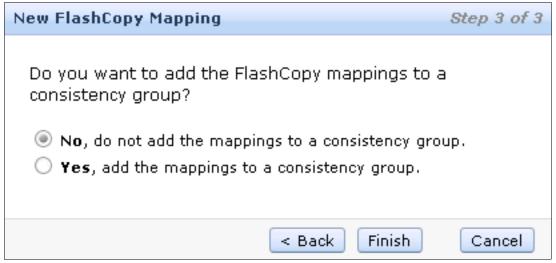


Figure 11-29 Select consistency group to add the FlashCopy mapping to

You can also create the FlashCopy mappings in the **FlashCopy Mapping** panel by clicking New FlashCopy Mapping on the top left, as shown in Figure 11-30.



Figure 11-30 Create new FlashCopy mapping in FlashCopy panel

A wizard will pop up to guide you to create a FlashCopy mapping, and the steps are the same as creating an advanced FlashCopy mapping using existing target volumes in the **FlashCopy** panel, which has been just described above.

Start FlashCopy Mapping

Most of the actions on FlashCopy mapping can be taken both in the **FlashCopy** panel or the **FlashCopy Mapping** panel. For the actions available on both panels, to avoid duplication, in

the following sections, we will show the steps based on the **FlashCopy** panel, although the steps would be the same if you were to use the **FlashCopy Mapping** panel.

You can start by selecting the FlashCopy target volume in the **FlashCopy** panel and choosing the **Start** option in the right click menu or the **Actions** drop down list, as shown in Figure 11-31. Then the status of the FlashCopy mapping will be changed to **Copying** from **Idle**.

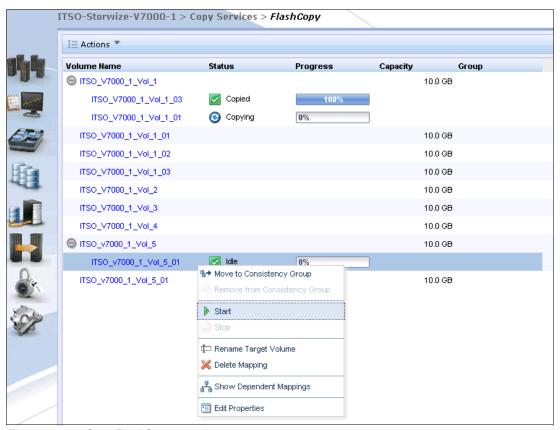


Figure 11-31 Start FlashCopy mapping

Stop FlashCopy Mapping

The FlashCopy mapping could be stopped by selecting the FlashCopy target volume in the **FlashCopy** panel and choosing the **Stop** option in the right click menu or the **Actions** drop down list, as shown in Figure 11-32 on page 410. After the stopping process completes, the status of the FlashCopy mapping will be changed to **Stopped**.

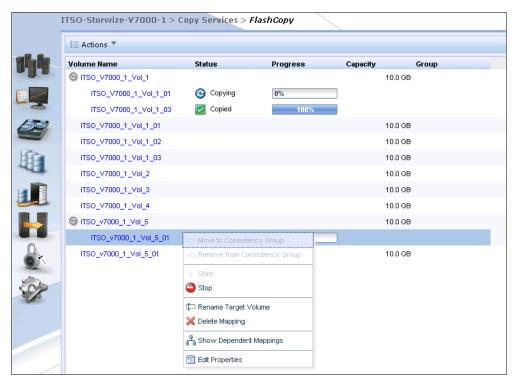


Figure 11-32 Stop the FlashCopy mapping

Rename Target Volume

If the FlashCopy target volumes were created automatically by IBM Storwize V7000 storage system, the name of the target volumes would be the source volume name plus a suffix containing numbers. The name of target volumes could be changed to be more meaningful to your environment.

To change the name of the target volume, selecting the FlashCopy target volume in the **FlashCopy** panel and choose the **Rename Target Volume** option in the right click menu or the **Actions** drop down list, as shown in Figure 11-33 on page 411.

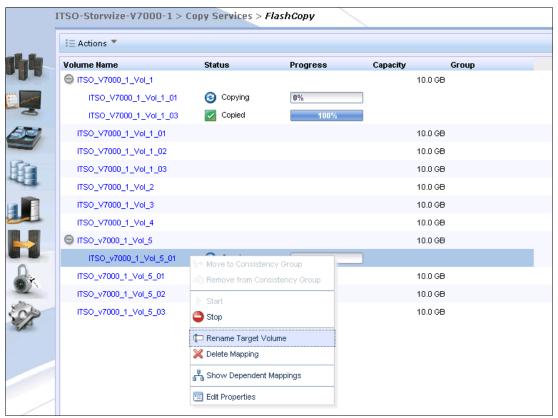


Figure 11-33 Rename Target Volume

Input your new name for the target volume, as shown in Figure 11-34. Click Rename to finish.



Figure 11-34 Input new name for the target volume

Rename FlashCopy Mapping

The FlashCopy mappings will be created with names beginning with **fcmap**. The name of FlashCopy mappings can be changed to be more meaningful to you.

To change the name of a FlashCopy mapping, selecting the FlashCopy mapping in the **FlashCopy Mappings** panel and choose the **Rename Mapping** option in the right click menu or the **Actions** drop down list as shown in Figure 11-35 on page 412.



Figure 11-35 Rename FlashCopy mapping

Then you need to input your new name for the FlashCopy mapping, as shown in Figure 11-36. Click **Rename** to finish.



Figure 11-36 Input new name for FlashCopy mapping

Delete FlashCopy Mapping

The FlashCopy mapping can be deleted by selecting the FlashCopy target volume in the **FlashCopy** panel and choosing the **Delete Mapping** option in the right click menu or the **Actions** drop down list, as shown in Figure 11-37 on page 413.

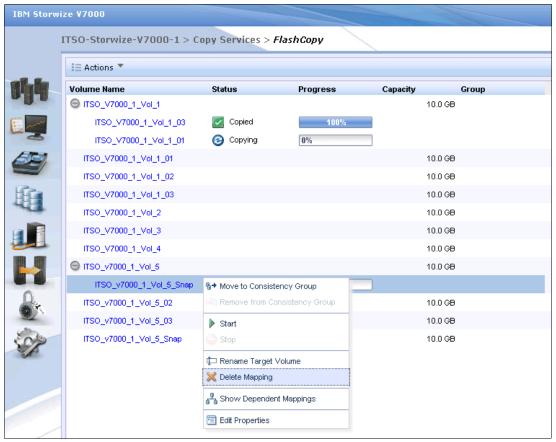


Figure 11-37 Delete Mapping

Note: If the FlashCopy mapping is in the **Copying** state, it must be stopped before being deleted.

Then you need to confirm your action to delete FlashCopy mappings in the pop-up window, as shown in Figure 11-38 on page 414. Verify the number of FlashCopy mappings you need to delete, and if you are sure that you want to delete the FlashCopy mappings when the data on the target volume is inconsistent with the source volume, tick the checkbox below. Click **Delete** and your FlashCopy mapping will be removed.

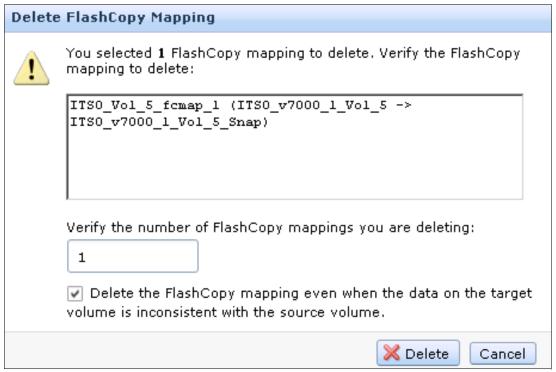


Figure 11-38 Confirm to delete FlashCopy mappings

Note: Deleting the FlashCopy mapping will not delete the target volume. If it is needed to reclaim the storage space occupied by the target volume, you need to delete the target volume manually.

Show Dependent Mapping

You can show the FlashCopy mapping dependency by selecting the FlashCopy source volume in the **FlashCopy** panel and choosing the **Show Dependent Mappings** option in the right click menu or the **Actions** drop down list, as shown in Figure 11-39 on page 415.

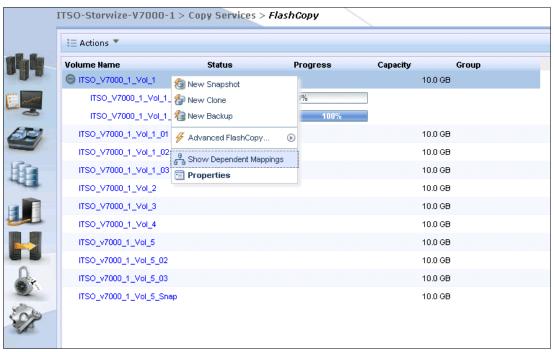


Figure 11-39 Show Dependent Mappings

The FlashCopy mapping dependency tree will be displayed, as shown in Figure 11-40.



Figure 11-40 FlashCopy mapping dependency

Edit Property

The background copy rate and cleaning rate can be changed after the FlashCopy mapping has been created, by selecting the FlashCopy target volume in the **FlashCopy** panel and choosing **Edit Property** option in the right click menu or the **Actions** drop down list, as shown in Figure 11-41 on page 416.

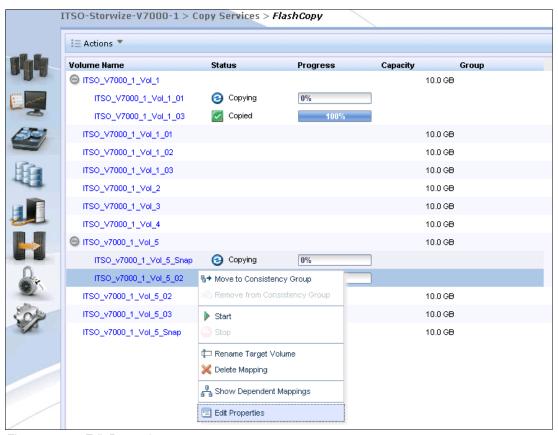


Figure 11-41 Edit Properties

You can then modify the value of background copy rate and cleaning rate by moving the pointer on the bar, as shown in Figure 11-42. Click **Save** to save changes.

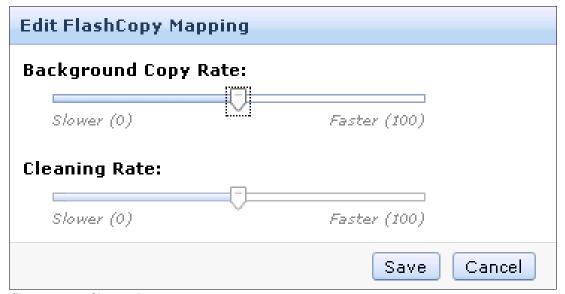


Figure 11-42 Change the copy rate

Manage FlashCopy Consistency Group

FlashCopy consistency groups can be managed in the **Consistency Groups** panel under **Copy Services** functional icon, as shown in Figure 11-43.

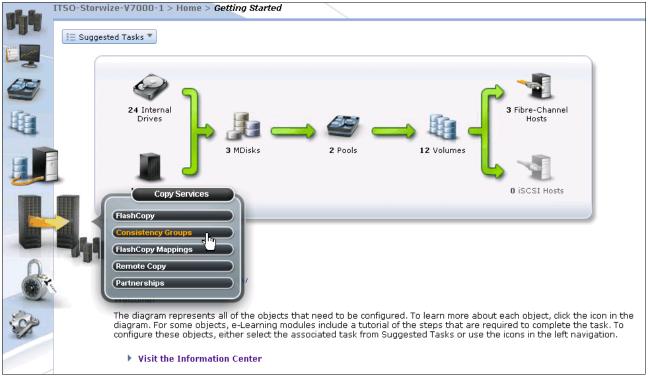


Figure 11-43 Access to the Consistency Groups panel

The **Consistency Groups** panel, as shown in Figure 11-44, is where you can manage both consistency groups and FlashCopy mappings.

On the left of the **Consistency Groups** panel, you can list out the consistency groups you need. Select **Not in a Group**, and all the FlashCopy mappings which are not in any consistency groups will be displayed on the right.

On the right of the **Consistency Groups** panel you can find out the properties of a consistency group and the FlashCopy mappings in it. You can also take actions on the consistency group and the FlashCopy mappings in it on the right. All the actions allowed on the FlashCopy mapping have been introduced in "Manage FlashCopy Mapping" on page 398, and will not be described again in this section.

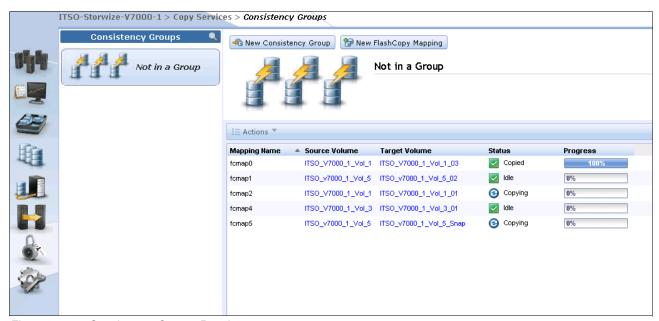


Figure 11-44 Consistency Groups Panel

Create FlashCopy Consistency Group

To create a new FlashCopy consistency group, click New Consistency Group on the top of the **Consistency Groups** panel, as shown in Figure 11-44.

Then you would be asked to input the name of the new consistency group, as shown in Figure 11-45. Following the naming conventions, type the name of the new consistency group in the box and click **Create**.



Figure 11-45 Input the name for new consistency group

After the creation process has been completed, you can find a new consistency group on the left of the Consistency Groups panel. Select the new consistency group, and you will see more detail information of this consistency group on the right, as shown in Figure 11-46 on page 419.

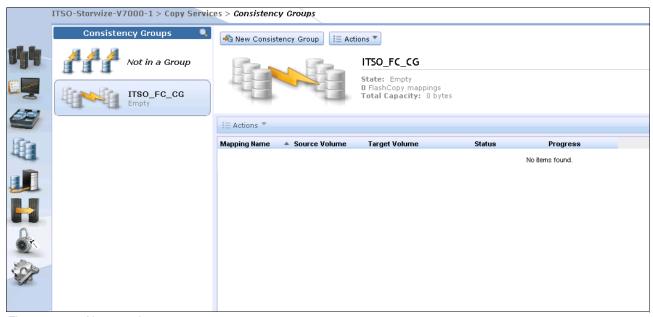


Figure 11-46 New consistency group

You simply rename the consistency group by clicking the name of the consistency group on the right and type a name for it following the naming convention. Under the name of the consistency group, the state shows that it is now an empty consistency group with no FlashCopy mapping in it.

Add FlashCopy Mappings to Consistency Group

Select **Not in a Group** to list all FlashCopy mappings with no consistency group. Then you can add FlashCopy mappings to a consistency group by selecting them and clicking **Move to Consistency Group** option in the right click menu or the **Actions** drop down list, as shown in Figure 11-47.



Figure 11-47 Select the FlashCopy mappings to add to a consistency group

Then you will be asked to specify which consistency group you want to move the FlashCopy mapping into as shown in Figure 11-48. Click **Move to Consistency Group** to continue.



Figure 11-48 Select consistency group

After the action completes you will find the FlashCopy mappings you selected have been removed from the **Not In a Group** list to the consistency group you chosen.

Start Consistency Group

To start a consistency group, select **Start** from the **Actions** drop down list on the top, as shown in Figure 11-49.

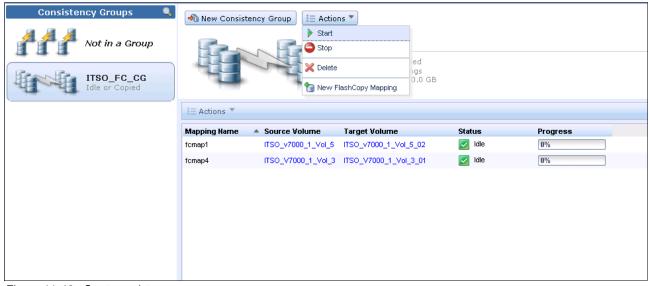


Figure 11-49 Start consistency group

After you start the consistency group, all the FlashCopy mappings in the consistency group will be started. The start time of the FlashCopy will also be recorded on the top right of the panel, as shown in Figure 11-50 on page 421.

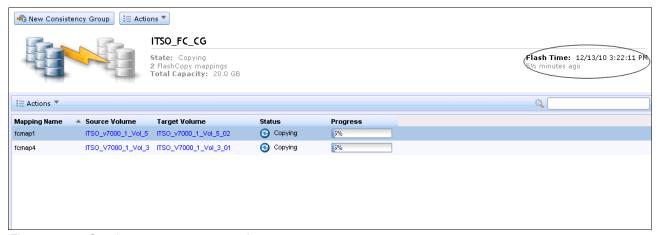


Figure 11-50 Consistency group start completes

Stop Consistency Group

The consistency group can be stopped by selecting **Stop** from the **Actions** drop down list on the top, as shown in Figure 11-51.

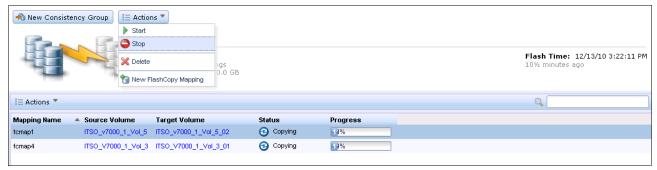


Figure 11-51 Stop consistency group

After the stop process completes, the FlashCopy mappings in the consistency group would be in Stopped state, and a red cross would appear on the icon of this consistency group to show an alert, as shown in Figure 11-52 on page 421.

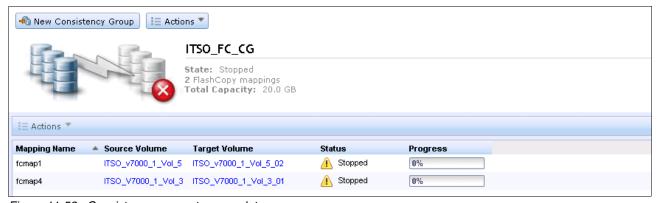


Figure 11-52 Consistency group stop completes

Remove FlashCopy Mappings from Consistency Group

The FlashCopy mappings can be removed from the consistency group by selecting the FlashCopy mappings and click **Remove from Consistency Group** in the right click menu or the **Actions** drop down list of the FlashCopy mapping, as shown in Figure 11-53.



Figure 11-53 Remove from consistency group

The FlashCopy mappings will be returned to the **Not in a Group list** after the removal from the consistency group.

Delete Consistency Group

The consistency group can be deleted by selecting **Delete** from the **Actions** drop down list on the top, as shown in Figure 11-54 on page 422.



Figure 11-54 Delete the consistency group

11.2 Remote Copy

In this section we describe how the Remote Copy function works in IBM Storwize V7000 and the implementation steps using the GUI are provided for Remote Copy configuration and management.

Remote Copy consists of two methods for copying: Metro Mirror and Global Mirror. Metro Mirror is designed for metropolitan distances in conjunction with a synchronous copy requirement, while Global Mirror is designed for longer distances without requiring the hosts to wait for the full round-trip delay of the long distance link.

Metro Mirror and Global Mirror are IBM branded terms for the functions Synchronous Remote Copy and Asynchronous Remote Copy, respectively. Throughout this paper, the term "Remote Copy" is used to refer to both functions where the text applies to each term equally.

11.2.1 Concepts Inside Remote Copy

Partnership

When creating a partnership, we connect two IBM Storwize V7000 that are separated by distance. After the partnership creation has been configured on both systems, further communication between the node canisters in each of the storage systems is established and maintained by the SAN network. All inter-cluster communication goes through the Fibre Channel network. Partnership must be defined on both IBM Storwize V7000s to make it fully functional.

Partnership topology

Building up partnerships between up to four IBM Storwize V7000 is allowed.

Typical partnership topologies between multiple IBM Storwize V7000s are described below:

1. Daisy-chain topology, as shown in Figure 11-55:

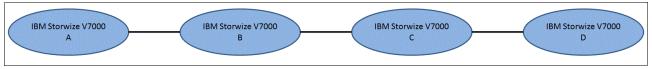


Figure 11-55 Daisy-chain topology for IBM Storwize V7000

2. Triangle topology, as shown in Figure 11-56 on page 424:

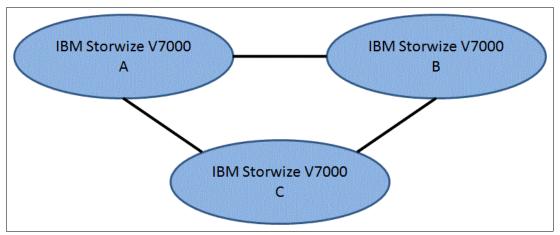


Figure 11-56 triangle topology for IBM Storwize V7000

3. Star topology, as shown in Figure 11-57:

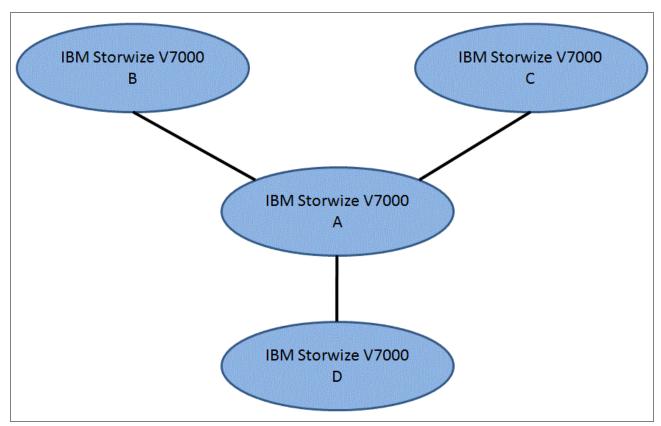


Figure 11-57 Star topology for IBM Storwize V7000

4. Full-meshed topology, as shown in Figure 11-58 on page 425:

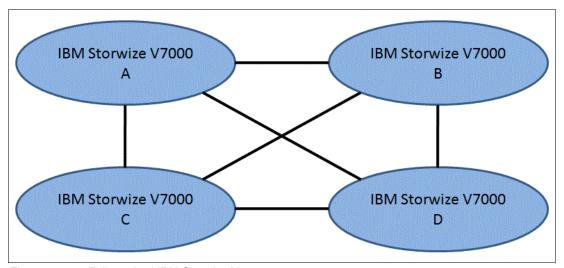


Figure 11-58 Full-meshed IBM Storwize V7000

Partnership States:

Partially Configured

Indicates that only one cluster partner is defined from a local or remote cluster to the displayed cluster and is started. For the displayed cluster to be configured fully and to complete the partnership, you must define the cluster partnership from the cluster that is displayed to the corresponding local or remote cluster.

Fully Configured

Indicates that the partnership is defined on the local and remote clusters and is started.

Remote Not Present

Indicates that the remote cluster is not present to the partnership.

Partially Configured (Local Stopped)

Indicates that the local cluster is only defined to remote cluster and the local cluster is stopped.

Fully Configured (Local Stopped)

Indicates that a partnership is defined on both the local and remote clusters and the remote cluster is present, but the local cluster is stopped.

Fully Configured (Remote Stopped)

Indicates that a partnership is defined on both the local and remote clusters and the remote cluster is present, but the remote cluster is stopped.

Fully Configured (Local Excluded)

Indicates that a partnership is defined between a local and remote cluster; however, the local cluster has been excluded. Usually this state occurs when the fabric link between the two clusters has been compromised by too many fabric errors or slow response times of the cluster partnership.

Fully Configured (Remote Excluded)

Indicates that a partnership is defined between a local and remote cluster; however, the remote cluster has been excluded. Usually this state occurs when the fabric link between the

two clusters has been compromised by too many fabric errors or slow response times of the cluster partnership.

Fully Configured (Remote Exceeded)

Indicates that a partnership is defined between a local and remote cluster and the remote is available; however, the remote cluster exceeds the number of allowed clusters within a cluster network. The maximum of four clusters can be defined in a network. If the number of clusters exceeds that limit, IBM Storwize V7000 determines the inactive cluster or clusters by sorting all the clusters by their unique identifier in numerical order. The inactive cluster partner which is not in the top four of the cluster unique identifiers displays Fully Configured (Remote Exceeded).

Remote Copy Relationship

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, 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. Therefore, 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 is supported to be established on the volumes within one IBM Storwize V7000 storage system, which called intra-cluster relationship, or in different IBM Storwize V7000 storage systems, which called inter-cluster relationship.

Usage of Remote Copy target volumes as Remote Copy source volumes is not allowed. A FlashCopy target volume cannot be used as Remote Copy source volume.

Metro Mirror

Metro Mirror is a type of remote copy that creates a synchronous copy of data from a master volume to an auxiliary volume. With synchronous copies, host applications write to the master volume but do not receive confirmation that the write operation has completed until the data is written to the auxiliary volume. This ensures that both the 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 11-59 on page 427 illustrates how a write 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, in case it is needed in a failover situation.

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, like

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.

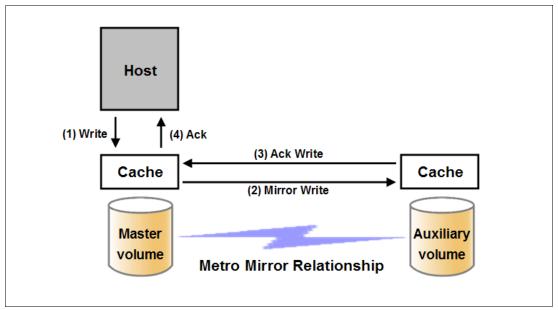


Figure 11-59 Write on volume in Metro Mirror relationship

Global Mirror

The 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.

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.

Figure 11-60 on page 428 shows that a write operation to the master volume is acknowledged back to the host issuing the write before the write operation is mirrored to the cache for the auxiliary volume.

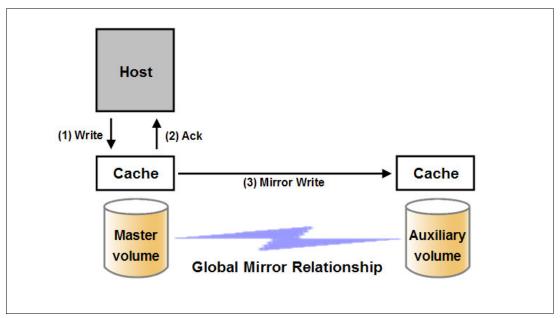


Figure 11-60 Global Mirror write sequence

The Global Mirror algorithms maintain a consistent image on the auxiliary 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.

In a failover scenario, where the secondary site needs to become the master source of data, depending on the workload pattern as well as the bandwidth and distance between local and remote site, certain updates might be missing at the secondary site. Therefore, any applications that will use this data must have an external mechanism for recovering the missing updates and reapplying them, for example, such as a transaction log replay.

IBM Storwize V7000 storage system provide you the advanced feature on Global Mirror to permit you to test performance implications before deploying Global Mirror and obtaining a long distance link. This advanced feature could be enabled by modifying the IBM Storwize V7000 storage system parameters, *gmintradelaysimulation* and *gminterdelaysimulation*, by using the CLI. These two parameters could 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 on how to enable the CLI feature in Appendix A, "CLI setup and SAN Boot" on page 543.

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. You may more information in "FlashCopy Consistency Groups" on page 386 regarding the necessity of consistency groups.

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 11-61 on page 429 illustrates the concept of Remote Copy consistency groups.

Because the RC_Relationship 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.

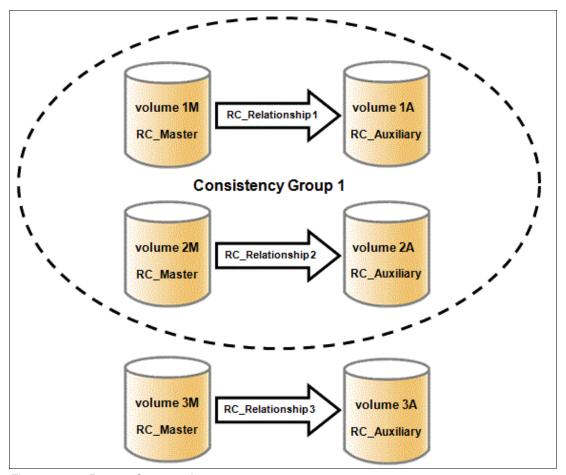


Figure 11-61 Remote Copy consistency group

Remote Copy relationships can only belong to one consistency group; however, 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 clusters, which are sometimes referred to as master and auxiliary clusters. All relationships in a consistency group must also have the same copy direction and state.

Metro Mirror and Global Mirror relationships cannot belong to the same consistency group. A copy type is automatically assigned to a consistency group when the first relationship is added to the consistency group. After the consistency group is assigned a copy type, only relationships of that copy type can be added to the consistency group.

Remote Copy and consistency group States

Stand-alone Remote Copy relationships and consistency groups share a common configuration and state model. All of the relationships in a non-empty consistency group have the same state as the consistency group.

These states apply to both the relationships and the consistency groups, except **Empty** state 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. A copy process must be started to make the secondary volumes consistent.

InconsistentCopying

The primary volume are accessible for read and write I/O operations, but the secondary volume are not accessible for either. This state indicates 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.

11.2.2 Planning for Remote Copy

General Guidelines for Remote Copy

- Building up partnerships between up to four IBM Storwize V7000 storage systems is allowed.Partnership must be defined on both IBM Storwize V7000 storage systems to make it fully functional.
- ► The two volumes in a relationship must be the same size.
- ► The Remote Copy relationship is supported to 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 cluster, they must be in the same I/O group.
- ► Usage of Remote Copy target volumes as Remote Copy source volumes is not allowed. A FlashCopy target volume cannot be used as Remote Copy source volume.

- ► The Metro Mirror function supports copy operations between volumes that are separated by distances up to 300 km.
- One Remote Copy relationships can only belong to one consistency group.
- ▶ All relationships in a consistency group must have matching primary and secondary clusters, which are sometimes referred to as master and auxiliary clusters. All relationships in a consistency group must also have the same copy direction and state.
- ► Metro Mirror and Global Mirror relationships cannot belong to the same consistency group.
- ► 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.
- ▶ 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.

Remote Copy Configuration Limits

Table 11-4 lists the Remote Copy configuration limits.

Table 11-4 Metro Mirror configuration limits

| Parameter | Value |
|--|--|
| Number of Remote Copy consistency groups per cluster | 256 |
| Number of Remote Copy relationships per cluster | 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

Zoning Recommendation

Node canister ports on each IBM Storwize V7000 must be able to communicate with each other in order for the partnership creation to be performed. Switch zoning is critical to facilitating intercluster communication.

The SAN zoning recommendation is 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 cluster.
- ▶ If dual-redundant ISLs are available, then the two ports from each node should be split evenly between the two ISLs, i.e. exactly one port from each node canister should be zoned across each ISL.

More details are available in the following link:

http://www-01.ibm.com/support/docview.wss?uid=ssg1S1003634&myns=s033&mynp=familyind5329743&mync=E

Note: When a local and a remote fabric are connected together for Remote Copy purposes, the inter-switch link (ISL) hop count between a local node and a remote node cannot exceed seven.

Remote Copy Link Requirements

The link requirements are valid for Metro Mirror and Global mirror. Basically there are two link requirements:

1. Round trip latency

Total round trip latency must be <80ms, and in each direction <40ms

2. Bandwidth

The bandwidth must satisfy the following requirement:

- a. Peak write load for all mirrored volumes.
- b. Additional background copy rate (best practice is 10% to 20% of maximum peak load) for initial synchronization and resynchronization.
- c. Remote Copy internal communication with approximately 2.6Mbps

Note: 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 could be met during single failure conditions.

Interaction between Remote Copy and FlashCopy

Table 11-5 lists which combinations of FlashCopy and Remote Copy are supported.

| Table 11-5 | FlashCopy and | I remote copy interaction |
|------------|---------------|---------------------------|
|------------|---------------|---------------------------|

| Component | Remote Copy primary site | Remote Copy secondary site |
|--------------------------|--------------------------|----------------------------|
| FlashCopy Source | Supported | Supported |
| FlashCopy Destination | Not supported | Not supported |

For disaster recovery purposes, you can use the FlashCopy feature to create a consistent copy of an image before you restart a Global Mirror relationship.

When a consistent relationship is stopped, the relationship enters the consistent_stopped state. While in this state, I/O operations at the primary site continue to run. However, updates are not copied to the secondary site. When the relationship is restarted, the synchronization process for new data is started. During this process, the relationship is in the inconsistent_copying state. The secondary volume for the relationship cannot be used until the copy process completes and the relationship returns to the consistent state. When this 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.

11.2.3 Manage Remote Copy using GUI

IBM Storwize V7000 storage system provide a separate function icon for copy service management. There are two panels for managing Remote Copy under the **Copy Services** function icon, including:

- Remote Copy
- Partnerships

As the name implies, these two panels are used to manage Remote Copy and the Partnership respectively.

Manage cluster partnerships

Partnership panel will be used to manage partnership between clusters. To access the Partnership panel, choose **Copy Services** functional icon and click **Partnerships** as shown in Figure 11-62.

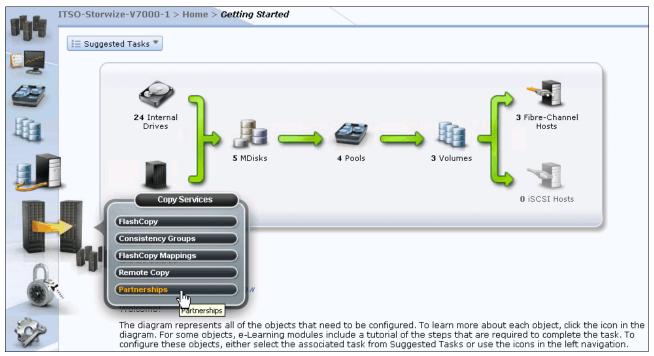


Figure 11-62 Partnership panel

Create New Partnership

No partnership is defined in our example as shown in Figure 11-63 on page 434, so we must create a new partnership between the IBM Storwize V7000 systems. Click **New Partnership** on the right of the **Partnership** panel.



Figure 11-63 Create new partnership between IBM Storwize V7000 storage systems

If there is no partnership candidate, a pop-up error will be displayed, as shown in Figure 11-64.

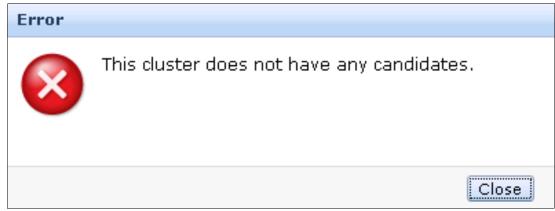


Figure 11-64 No Candidates available to create a partnership

Check the zoning and the system status, make sure the clusters can 'see' each other. Then you can create your partnership by selecting the appropriate remote storage system as shown in Figure 11-65 on page 435 and defining the available bandwidth between both systems.



Figure 11-65 Select the remote IBM Storwize V7000 storage system for a partnership

The bandwidth you need to input here is used by the background copy process between the clusters in the partnership. In order to set the background copy bandwidth optimally, make sure that you consider all three resources (the primary storage, the intercluster link bandwidth, and the secondary storage) to avoid overloading them so as to affect the foreground I/O latency.

Click **Create** and the partnership definition is complete on the first IBM Storwize V7000 system, as shown in Figure 11-66. You can find the partnership listed on the left of the **Partnership** panel, and with the selection of the partnership, more information for this partnership will be displayed on the right.

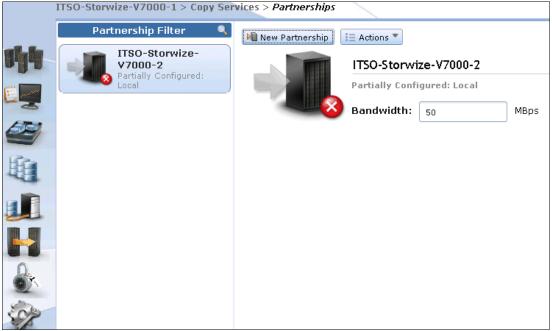


Figure 11-66 Partially configured relationship

Perform the same steps on the second storage system which will become a fully configured partner, as shown in Figure 11-67.

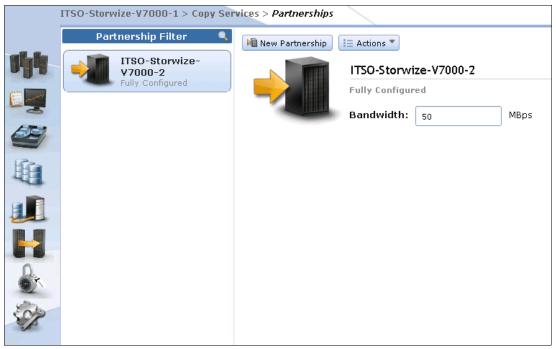


Figure 11-67 Fully Configured partnership

The Remote Copy partnership is now implemented between two IBM Storwize V7000 systems and both systems are ready for further configuration of Remote Copy relationships.

You can also change the bandwidth setting for the partnership in the **Partnerships** panel, as shown in Figure 11-68. Select **Apply Changes** to confirm your modification.

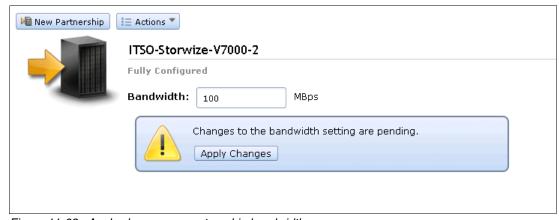


Figure 11-68 Apply changes on partnership bandwidth

Stop and Start Partnership

You can stop the partnership, by select **Stop Partnership** in the **Actions** drop down list, as shown in Figure 11-69. Although if you stop the partnership, the relationship using this partnership will be disconnected.

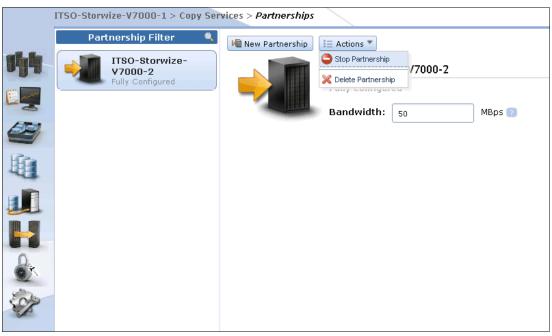


Figure 11-69 Stop Partnership

After you have stopped the partnership, your partnership will be marked as Fully Configured: Stopped, as shown in Figure 11-70.



Figure 11-70

With a stopped partnership, you can restart it as needed by selecting **Start Partnership** in the **Actions** drop down list, as shown in Figure 11-71 on page 438.

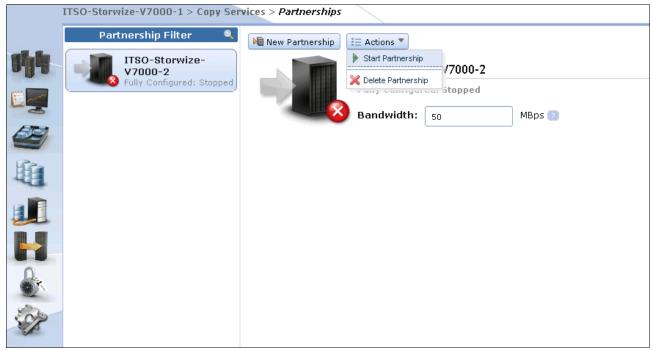


Figure 11-71 Start Partnership

The partnership will return to fully configured status when it has been restarted.

Delete Partnership

You can delete a partnership by selecting **Delete Partnership** in the **Actions** drop down list, as shown in Figure 11-72.

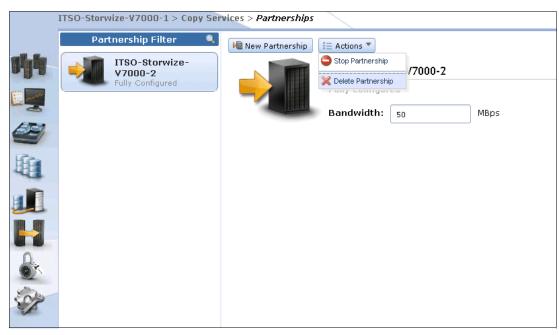


Figure 11-72 Delete Partnership

Manage Stand-alone Remote Copy Relationships

A Remote Copy relationship can be defined between two volumes, 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.

Access the Remote Copy panel to manage remote copy, as shown in Figure 11-73.

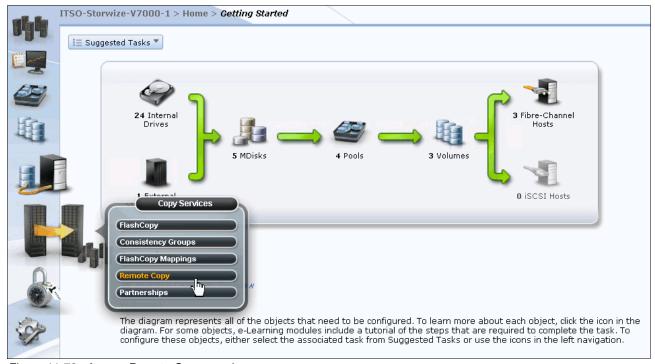


Figure 11-73 Access Remote Copy panel

The **Remote Copy** panel, as shown in Figure 11-74 on page 440, is where you can manage Remote Copy relationship and Remote Copy consistency groups.

On the left of the **Remote Copy** panel, there is a consistency group filter to list Remote Copy consistency groups that meet your requirement. On the right of the **Remote Copy** panel, you can find out the properties of a consistency group and the Remote Copy relationships in it. You can also take actions on Remote Copy relationships and Remote Copy consistency group on the right. Select **Not in a Group**, all the Remote Copy relationships which are not in any Remote Copy consistency groups will be displayed on the right.

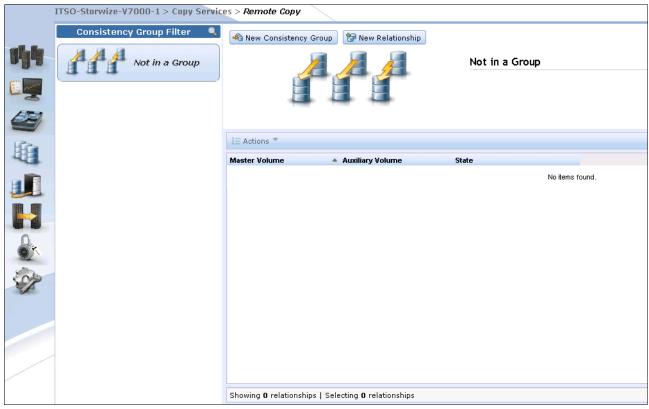


Figure 11-74 Remote Copy panel

Create Stand-alone Remote Copy Relationships

To create a new Remote Copy relationship, click on **New Relationship** on the right of the **Remote Copy** panel, as shown in Figure 11-74. Then a wizard will guide you through the Remote Copy relationship creation process.

As shown in Figure 11-75 on page 441, you need to set the Remote Copy relationship type first. Based on your requirement, you can select Metro Mirror (synchronous replication) or Global Mirror (asynchronous replication). Select the appropriate replication type and click **Next**.



Figure 11-75 Select the appropriate Remote Copy type

In the next step you will need to select your Remote Copy auxiliary (target) storage system, the local system or the already defined second storage system as the Remote Copy partner. In our example, shown in Figure 11-76, we choose another system to build an inter-cluster relationship. Click **Next** to continue.

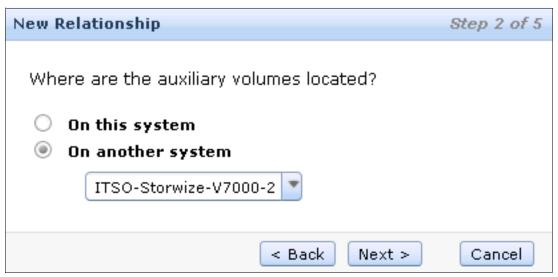


Figure 11-76 Select Remote Copy partner

Then, the Remote Copy master and auxiliary volume needed to be specified. Both volumes must have the same size. As shown in Figure 11-77 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**.

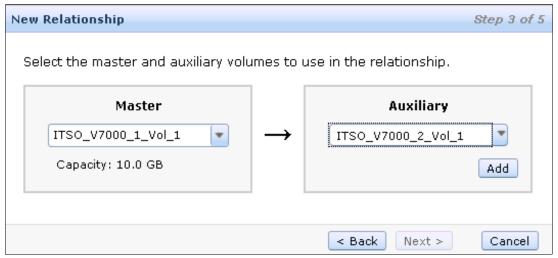


Figure 11-77 Select Master and Auxiliary volumes

You can define multiple, independent relationships using the **Add** button and you can remove a relationship by clicking on the red cross. In our example, we create two independent Remote Copy relationships, as shown in Figure 11-78.

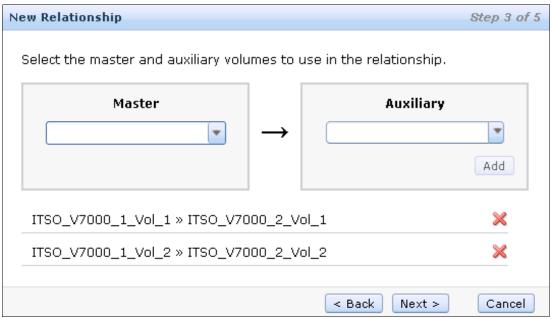


Figure 11-78 Define multiple, independent relationships

In the next step you are asked if the volumes in the relationship are already synchronized. In most situations the data on the master volume and on the auxiliary volume are not identical. So select **No** and click **Next** to enable an initial copy as shown in Figure 11-79 on page 443.

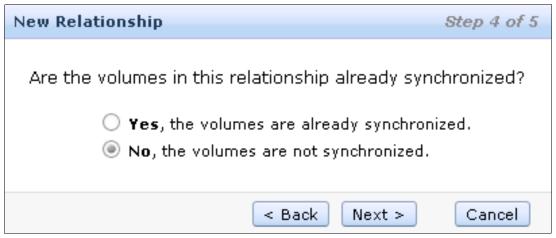


Figure 11-79 Activate an initial data copy

If you select **Yes, the volumes are already synchronized** in this step, A warning message will pop up as shown in Figure 11-80. Make sure the volumes are truly identical, then click **OK** to continue.

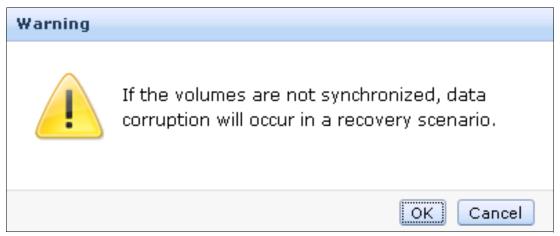


Figure 11-80 Warning massage to make sure the volumes are synchronized

In the last step you can choose to start the initial copying progress now, or wait to start at a later time. In our example, we select **Yes, start copying now** and click **Finish**, as shown in Figure 11-81 on page 444.



Figure 11-81 Choose if you want to start copying now or later

After the Remote Copy relationships creation completes, two independent Remote Copy relationships will be defined and displayed in **Not in a Group list**, as shown in Figure 11-82.

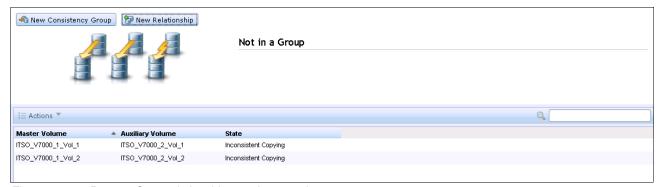


Figure 11-82 Remote Copy relationship creation completes

Optionally you can monitor the ongoing initial synchronisation in the status indicator of **Running Tasks**, as shown in Figure 11-83 on page 445.

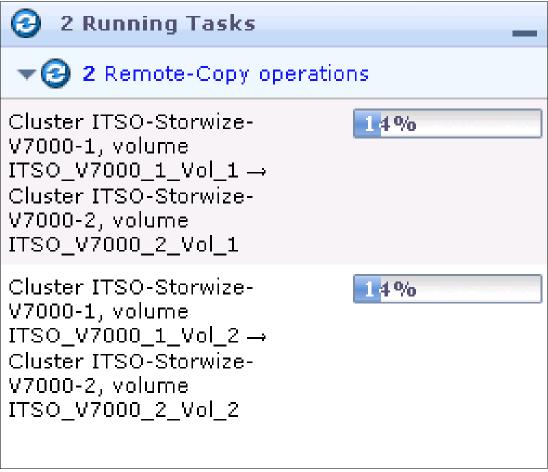


Figure 11-83 Remote Copy initialization progress

Stop Stand-alone Remote Copy Relationship

The Remote Copy relationship can be stopped by selecting the relationship and choosing the **Stop** in the right click menu or in the **Actions** drop down list, as shown in Figure 11-84 on page 445.

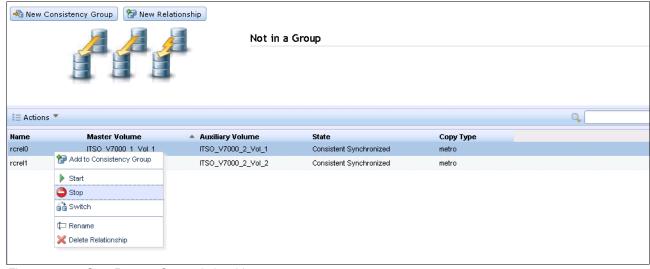


Figure 11-84 Stop Remote Copy relationship

In the next step, allow secondary read/write access if required, and click **Stop Relationship**, as shown in Figure 11-85.

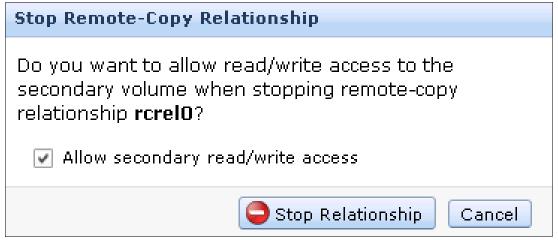


Figure 11-85 Options to allow the secondary rw access

After the stop completes, the state of the Remote Copy relationship will be changed to **Idling** from **Consistent Synchronized**, as shown in Figure 11-86 on page 446. Now read/write access to both volumes is allowed.

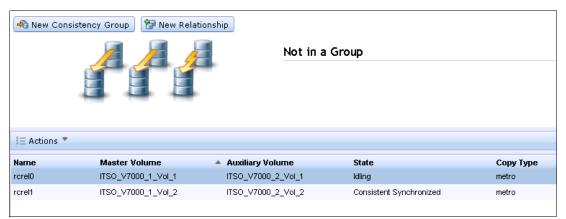


Figure 11-86 Remote Copy relationship stop completes

Start Stand-alone Remote Copy Relationship

The Remote Copy relationship can be started by selecting the relationship and choosing **Stop** in the right click menu or in the **Actions** drop down list, as shown in Figure 11-87.



Figure 11-87 Start Remote Copy relationship

To start a Remote Copy relationship, the most important item is selecting the copy direction. Both master and auxiliary volumes could be the primary. Make your decision based on the requirement in your situation and click **Start Relationship**. In our example, we choose the master volume to be the primary, as shown in Figure 11-88.

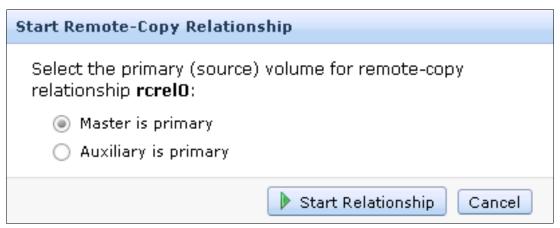


Figure 11-88 Choose copy direction

Switch Stand-alone Remote Copy Relationship

The copy direction of the Remote Copy relationship could be switched by selecting the relationship and choosing **Switch** in the right click menu or in the **Actions** drop down list, as shown in Figure 11-89.



Figure 11-89 Switch Remote Copy relationship

A warning message will pop up and show you the consequences, as shown in Figure 11-90 on page 448. If you switch the Remote Copy relationship, the copy direction of the relationship will become the opposite: the current primary volume will become the secondary, while the current secondary volume will become the primary. What is more, write access to the current primary volume will be lost and write access to the current secondary volume will be enabled. If it is not in a disaster recovery situation, you need to stop your host I/O to the current primary volume in advance. Make sure you are prepared for the consequences, and if so click **OK** to continue.

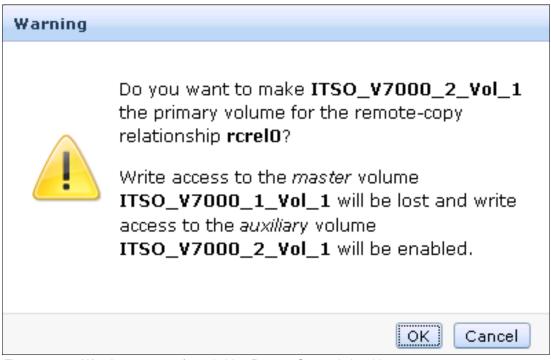


Figure 11-90 Warning message for switching Remote Copy relationship

After the switch completes your Remote Copy relationship will be tagged, as shown in Figure 11-91, to show you the primary volume in this relationship has been changed.

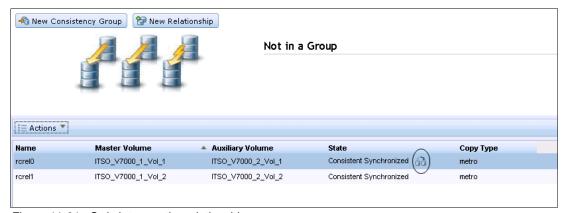


Figure 11-91 Switch tag on the relationship

Rename Stand-alone Remote Copy Relationship

The Remote Copy relationship can be renamed by selecting the relationship and choosing **Rename** in the right click menu or in the **Actions** drop down list, as shown in Figure 11-92.



Figure 11-92 Rename the Remote Copy relationship

Then input the new name for the Remote Copy relationship and click **Rename**, as shown in Figure 11-93.



Figure 11-93 Input a new name for the Remote Copy relationship

Delete Stand-alone Remote Copy Relationship

The Remote Copy relationship can be deleted by selecting the relationship and choosing **Delete Relationship** in the right click menu or in the **Actions** drop down list, as shown in Figure 11-94 on page 450.



Figure 11-94 Delete Remote Copy relationship

Confirmation is needed by verifying the number of relationships to be deleted, as shown in Figure 11-95. Click **Delete** to proceed.

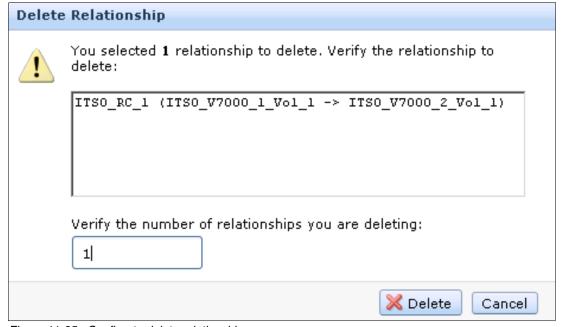


Figure 11-95 Confirm to delete relationships

Manage Remote Copy Consistency Group

Remote Copy consistency group can be managed from the **Remote Copy** panel as well.

Create Remote Copy Consistency Group

To create a new Remote Copy consistency group, click the **New Consistency Group** button on the top of the **Remote Copy** panel, as shown in Figure 11-96.

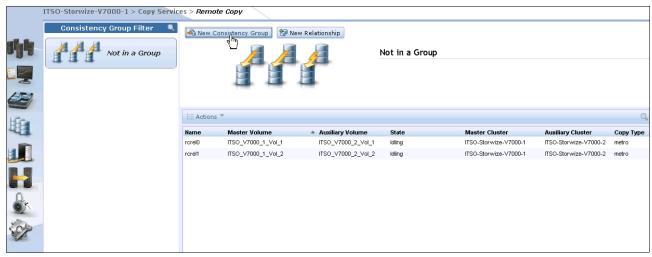


Figure 11-96 Create new consistency group

Then you need to input a name for your new consistency group as shown in Figure 11-97.



Figure 11-97 Enter a name for the new consistency group

In the next step, as shown in Figure 11-98, You can choose to create an empty consistency group or if you want to add Remote Copy relationships into the consistency group now. 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, we create a empty consistency group now, and add Remote Copy relationships to the consistency group afterwards. Click **Finish** to process.



Figure 11-98 Create a empty consistency group

After the creation process completes, a new empty consistency group appears on the left part of the **Remote Copy** panel. Click the new consistency group, and you can find more information on the right, as shown in Figure 11-99.

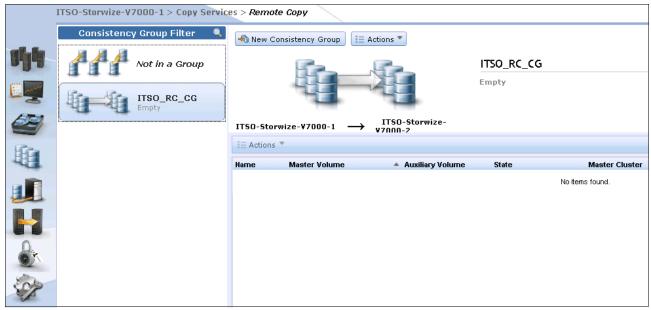


Figure 11-99 New Remote Copy consistency group created

You can find the name and the status of the consistency group beside the relationship icon. It is easy to change the name of consistency group by clicking the name and entering a new one. On the top of the right part in the **Remote Copy** panel, you can take actions on the Remote Copy consistency group. And below the relationship icon, you can find all the Remote Copy relationships in this consistency group. The actions on the Remote Copy relationships can be applied here as well as with the right click menu or the **Actions** drop down list.

Add Remote Copy to Consistency Group

The Remote Copy relationships in **Not in a Group** list could be added to a consistency group by selecting the volumes and choosing **Add to Consistency Group** in the right click menu or from the **Actions** drop down list, as shown in Figure 11-100 on page 453

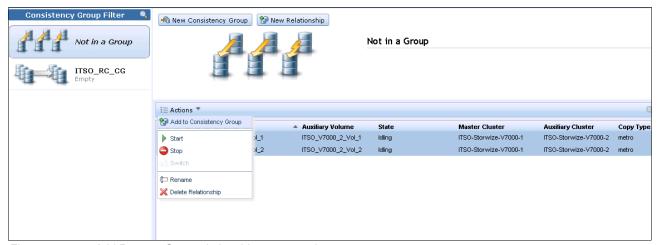


Figure 11-100 Add Remote Copy relationships to a consistency group

In the next step, you are asked to choose the consistency group to add the Remote Copy relationships to. Based on your requirements, select the appropriate consistency group and click **Add to Consistency Group**, as shown in Figure 11-101.



Figure 11-101 Choose the right consistency group to add the Remote Copy relationships to

Then you will find your Remote Copy relationships in the consistency group you selected.

Start Consistency Group

The Remote Copy relationship could be started by selecting **Start** in the **Actions** drop down list on the top of the **Remote Copy** panel, as shown in Figure 11-102 on page 454.

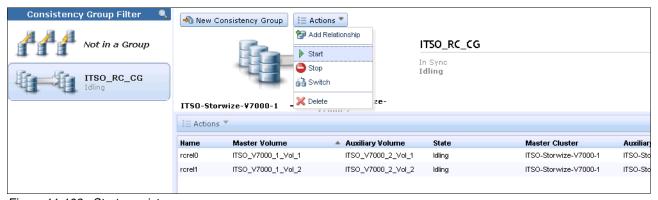


Figure 11-102 Start consistency group

Next, you can select the copy direction of the consistency group as required as shown in Figure 11-103, and we choose **Master is primary** in our example and click Start **Consistency Group**. Then the consistency group will start copying data from the primary to the secondary.

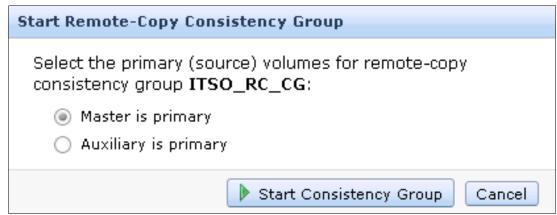


Figure 11-103 Choose the primary to start the consistency group

Stop Consistency Group

The Remote Copy relationship can be stopped by selecting **Stop** in the **Actions** drop down list on the top of the **Remote Copy** panel, as shown in Figure 11-104.



Figure 11-104 Stop consistency group

In the next step, you can allow read/write access to secondary volumes by ticking the checkbox, as shown in Figure 11-105, and click Stop Consistency Group to proceed.

Stop Remote-Copy Consistency Group Do you want to allow read/write access to the secondary volumes for remote-copy consistency group ITSO_RC_CG when stopping the group? ✓ Allow secondary read/write access Stop Consistency Group Cancel

Figure 11-105 Allow secondary read/write access

Switch Consistency Group

As with the switch action on the Remote Copy relationship, you can switch the copy direction of the consistency group as well. To switch the copy direction of the consistency group, select **Switch** in the **Actions** drop down list on the top of the **Remote Copy** panel, as shown in Figure 11-106.



Figure 11-106 Switch consistency group

A warning message will pop up to you to show the consequence, as shown in Figure 11-107 on page 456. After the switch, the primary cluster in the consistency group will be changed. And write access to current master volumes will be lost, while write access to the current auxiliary volumes will be enabled. Make sure that is what you need, and if so click **OK** to continue.

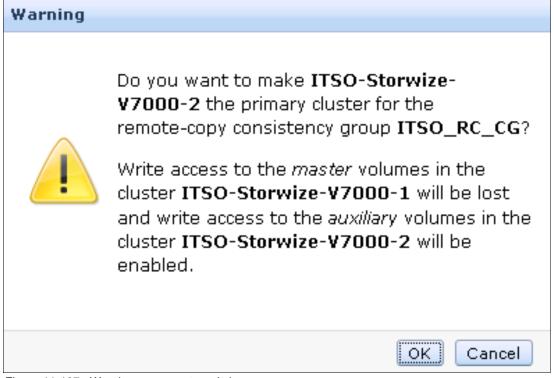


Figure 11-107 Warning message to switch

Remove Remote Copy relationships from Consistency Group

The Remote Copy relationships could be removed from the consistency group by selecting the Remote Copy relationships and choosing **Remove from Consistency Group** in the right click menu or in the **Actions** drop down list, as shown in Figure 11-108.

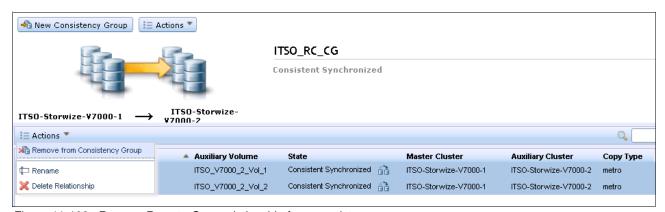


Figure 11-108 Remove Remote Copy relationship from consistency group

You will be asked to confirm the Remote Copy relationships you want to delete from the consistency group, as shown in Figure 11-109 on page 457. Make sure the Remote Copy relationships shown in the box are the ones you need to remove from the consistency group, and click **Remove** to proceed.



Figure 11-109 Confirm the Remote Copy relationships to be removed from consistency group

After the removal process completes, the Remote Copy relationships will be deleted from the consistency group and displayed in the **Not in a Group** list.

Delete Consistency Group

The consistency group can be deleted by selecting **Delete** in the **Actions** drop down list on the top of the **Remote Copy** penal, as shown in Figure 11-110.

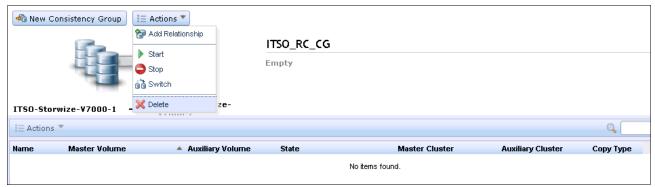


Figure 11-110 Delete consistency group

Next, you will need to confirm your delete of the consistency group, as shown in Figure 11-111. Click **OK** if you are sure that this consistency group should be deleted.

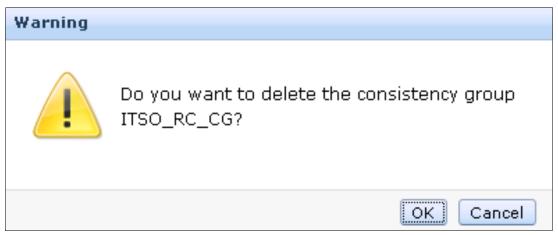


Figure 11-111 Confirm to delete consistency group

The consistency group has been deleted.



Chapter 12.

Tivoli Storage Productivity Center

The IBM Tivoli Storage Productivity Center (TPC) suite of storage infrastructure management tools can help customers improve time to value, as well as reduce the complexity of managing their storage environments by centralizing, simplifying and optimizing storage tasks associated with storage systems, storage networks, replication services and capacity management.

The IBM Tivoli Storage Productivity Center is designed to:

- ► Help centralize the management of your storage infrastructure from a single interface using role-based administration and single sign-on.
- Provide a single management application with modular integrated components that are easy to install and provide common services for simple/consistent configuration and consistent operations across host, fabric and storage systems
- Manage performance and connectivity from the host file system to the physical disk, including in-depth performance monitoring and analysis on SAN fabric performance.
- Manage the capacity utilization and availability of storage systems, file systems and databases
- Monitor, manage and control (zone) SAN fabric components
- Automate capacity provisioning of file systems
- Monitor and track the performance of SAN-attached SMI-S compliant storage devices
- Manage advanced replication services (Global Mirror, Metro Mirror and Flash Copy)

IBM 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, perform device configuration and management of multiple devices from a single user interface. And it can tune and proactively manage the performance of storage devices on the Storage Area Network (SAN) and manage, monitor and control your SAN fabric. Now there is also a Plug in for IBM System Director available named "IBM System Director Storage Productivity Center". It is designed to be used as an embedded version of TPC 4.2.1 available for IBM System Director 6.2.1.

12.1 TPC Overview

In this book we will cover a basic overview of TPC, and show how to perform a typical setup (12.4, "Installing Tivoli Storage Productivity Center components") and connect it to the IBM Storwize V7000 (12.5, "Connect TPC to IBM Storwize V7000"). Remember that the functions are dependent on the type of TPC license you have and is 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 TPC for Disk Midrange Edition

Tivoli Storage Productivity Center (TPC) now enables enterprise-wide management of IBM Storwize V7000. In addition to device-level management software that comes packaged with IBM Storwize V7000, TPC offers incremental benefits that are included its latest version 4.2.1. TPC for Disk Midrange Edition (MRE), is ideally suited and recommended for use with Storwize V7000.

- ▶ Simplified deployment and visibility: TPC for DISK MRE supports Storwize V7000 during discovery as a new device type with different types of managed disks. TPC's quick discovery and configuration capabilities enables you to attach the storage device with ease and helps you to configure efficiently i.e. plan for replication while provisioning the device. Disk MRE also enables Launch in Context and Single Sign-on, significantly reducing the burden on storage administrators.
- ► TPC's topology viewer offers a collective view that includes Storwize V7000, helps differentiate between external array-based disk and local disk, and displays tiering information. TPC for Disk MRE also extends thin provisioning support for Storwize V7000, enabling increased utilization and lowered costs.
- ▶ Performance management: IBM Storwize V7000 offers unmatched performance and availability among midrange disk systems. Adding TPC for Disk MRE enhances performance monitoring by capturing metrics such as input and output (I/O) and data rates, and cache utilization from a single console. TPC helps establish threshold levels based on business priorities and alerts when these levels are breached. This helps administrators to avoid the 'knee curves', while proactively managing performance and service levels by tracking historical information.
- ► Replication management: Tivoli Storage Productivity Center for Replication enables IBM Storwize V7000 with superior disaster recovery (DR) management providing central control of the replication environment, and helps establish Flash Copy, Metro and Global Mirror relationships.
- ► Comprehensive reporting: TPC for Disk MRE offers detailed metrics that include performance data for storage subsystem, controller, cache, I/O, array, disk group and port. These performance statistics can be stored in database tables for later use, so that storage administrators can track and measure service levels.
- ► Easy Tier: TPC for Disk MRE also provides information on Easy Tier to monitor automatic migration of data assets to high-performance disk tiers (SSD).
- ► Advanced management: TPC Standard Edition offers performance metrics-based recommendations for provisioning, including SAN planning (and DR planning with TPC for Replication). TPC's new 'Disk Magic' model helps identify 'hot spots', improving storage optimization for IBM Storwize V7000.

12.2 TPC Architecture

The IBM Tivoli Storage Productivity Center consists of the following key components:

- Data Server: ("Data Server" on page 461)
- ► Device Server: ("Device Server" on page 461)
- ► TIP: ("Tivoli Integrated Portal" on page 461)
- ► TPC for Replication: ("Tivoli Storage Productivity Center for Replication" on page 461)
- ► DB2®: ("DB2 Database" on page 462)
- ► Agents: ("Agents" on page 462)
- ► Interfaces: ("Interfaces" on page 462)

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 Tivoli Integrated Portal

IBM Tivoli Storage Productivity Center V4 is integrated with IBM Tivoli Integrated Portal (TIP). This integration provides functionalities like single sign-on and the use of 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 (TCR) is a component provided by TIP. 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 TPC V4.1, the IBM Tivoli Storage Productivity Center for Replication product is starting to get integrated into TPC. Currently the integration is limited to basic functions such as providing Launch in Context links in the TPC GUI, as well as crosschecks when a volume is deleted with TPC and mapping of user roles.

12.2.5 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 machines.

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 TPC agents (Storage resource agent, Data agent and Fabric agent) as well as SMI-S enabled storage devices that use a CIMOM agent (either embedded or as a proxy agent). Storage Resource agent, 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 TPC 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) enables users to request information and then generate and display reports based on that information. Certain user interfaces can also be used for configuration of TPC or storage provisioning for supported devices.

The following interfaces are available for TPC:

- ► TPC GUI: This is the central point of TPC administration. Here you have the choice of configuring TPC after installation, define jobs to gather information, initiate provisioning functions, view reports, and work with the advanced analytics functions.
- ▶ Java[™] Web Start GUI: When you use Java Web Start, the regular TPC GUI will be downloaded to your workstation and started automatically, so you do not have to install the GUI separately. The main reason for using the Java Web Start is that it can be integrated into other products (for example, TIP). By using Launch in Context from those products, you will be guided directly to the select panel. The Launch in Context URLs can also be assembled manually and be used as bookmarks.
- ▶ **TPCTOOL:** This is a command line (CLI) based program which interacts with the TPC Device Server. Most frequently it is used to extract performance data from the TPC repository database in order to create graphs and charts with multiple metrics, with various unit types and for multiple entities (for example, Subsystems, Volumes, Controller, 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 TPC V4, the TPC 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, you need to have sufficient knowledge about SQL. To access the views, DB2 supports various interfaces, for example, JDBC and ODBC.

12.3 Preparing Windows to Install TPC

TPC can be installed on one of your host servers, or you can also purchase the IBM System Storage Productivity Center (SSPC). It is a hardware appliance that consolidates IBM storage administration and configuration utilities into a single console. The System Storage Productivity Center is comprised of:

- Physical hardware (an IBM System x3550 M2 Type 7946 server running the Microsoft Windows Server 2008 32-bit Enterprise Edition operating system)
- Management software (including the IBM Tivoli Storage Productivity Center and other IBM Software)

If you have SSPC the software is pre-installed already.

In this section, we cover the Windows TPC installation wizard. The installation wizard covers two installation paths, "Typical" and "Custom". We will 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 all the panels look the same.

12.3.1 Installation overview

In order to get Tivoli Storage Productivity Center V4.2 to work, the prerequisite components have to be installed prior to invoking the installation wizard. You need to follow the following steps:

- ► Check that the system meets the prerequisites ("Verifying system requirements for Windows" on page 465).
- ▶ Install the prerequisite components (12.3.3, "Pre-installation steps for Windows" on page 464).
- ► Install Tivoli Storage Productivity Center components (12.4, "Installing Tivoli Storage Productivity Center components" on page 477).
- Connect TPC to IBM Storwize V7000 (12.5, "Connect TPC to IBM Storwize V7000" on page 486).

The Typical installation allows you to 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 TIP
- ► Clients: PC GUI
- Storage Resource Agent

The drawback of using the typical installation is that everything besides the above selection will be set to default. At about 75% of the installation the installer will launch the Tivoli Storage Productivity Center for Replication installation wizard to give you the options to also change some installation parameters. You basically have to step through it, and press Finish to start its installation procedure. Once this is done, you have to 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 writing. The Passport Advantage® and Web media content consists of a disk image and a SRA zip file. The disk image is broken up into 4 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:
 - IBM Tivoli Integrated Portal
 - IBM Tivoli Storage Productivity Center for Replication
- ► Disk1 part 3 contains:
 - IBM Tivoli Integrated Portal Fixpack

Note: Part 1, part2 and part 3 are required for every TPC installation and need to be downloaded and extracted to a single directory.

- ► Disk1 part 4 contains an optional component:
 - IBM Tivoli Storage Productivity Center Monitoring Agent for IBM Tivoli Monitoring

Note: On **Windows**, ensure that the directory name where the installation images reside has *no* spaces or special characters. This will cause the Tivoli Storage Productivity Center installation to fail. For example, a failure will occur if you happened to have a directory name such as:

C:\tpc 42 standard edition\disk1

The SRA zip file contains Tivoli Storage Productivity Center Storage Resource Agents (SRAs). 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:

- Directory: readme
- Directory: sra
- ► File: version.txt

In addition to the images mentioned above there are these images available:

- Tivoli Storage Productivity Center Storage National Language Support
- ► 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, or 9.5 Fix Pack 3a or later, or 9.7 without any Fix Pack

Fixpack and 9.7: Do not use 9.7 with any fixpack, even if it is a later version, as the TPC installation will fail.

Verifying system requirements for Windows

To install TPC the following system requirements have to be fulfilled:

- 1 x Intel® Quad Core Xeon or greater
- ▶ 8 GB of RAM
- ▶ 6 GB of free disk space and 500 MB in the temp directory.
- ► The latest operating system support level matrix is available at:

http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp?topic=%2Fcom.ibm.tpc V42.doc%2Ffqz0 t upgrading all.html

Verifying primary domain name systems

Before you start the installation, we recommend that you verify if a primary domain name system (DNS) suffix is set. This can require a computer restart.

To verify the primary DNS name, follow these steps:

- 1. Right-click **My Computer** on your desktop.
- 2. Click Properties.
- 3. Click the Computer Name tab. On the panel that is displayed, click Change.
- 4. Enter the host name in the Computer name field. Click More to continue.
- In the next panel, 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 NetBIOS settings

If NetBIOS is not enabled on Microsoft Windows 2003, then GUID is not generated. You must verify and activate NetBIOS settings. On your Tivoli Storage Productivity Center Server, go to Start → Control Panel → Network Connections. Select your Local Area Connections. From the Local Area Connection Properties panel, double-click Internet Protocol (TCP/IP). The next panel is the Internet Protocol (TCP/IP) Properties. Click Advanced. On the WINS tab, select Enable NetBIOS over TCP/IP and click OK as shown in Figure 12-1.

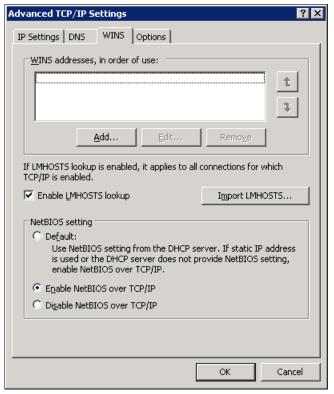


Figure 12-1 Advanced TCP/IP properties

Installing TPC 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 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, follow these steps:

1. Insert the IBM DB2 Installer CD into the CD-ROM drive.

If Windows autorun is enabled, the installation program ought to start automatically. If it does not, open Windows Explorer and go to the DB2 Installation image path and double-click **setup.exe**.

Note: Only the user ID that has installed the DB2 product has the privilege to issues the **db2start** and **db2stop** commands.

You will see the Welcome panel, as shown in Figure 12-2. Select **Install a Product** to proceed with the installation.

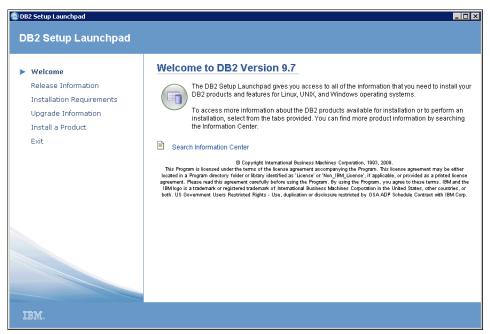


Figure 12-2 DB2 Setup Welcome panel

The next panel allows you to select the DB2 product to be installed. Select the DB2
 Enterprise Server Edition Version 9.7 by clicking Install New to proceed as shown in Figure 12-3.

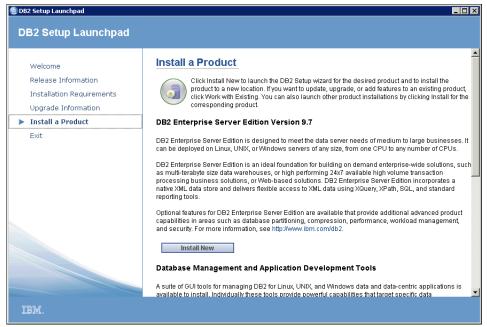


Figure 12-3 Select product

3. The DB2 Setup wizard panel is displayed, as shown in Figure 12-4. Click Next to proceed.

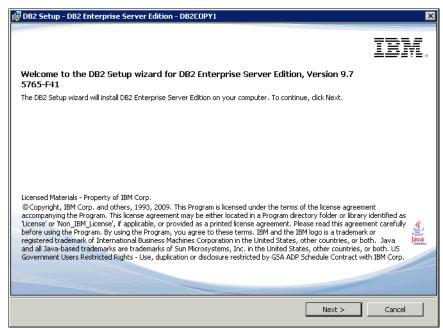


Figure 12-4 Setup wizard

 The next panel displays the license agreement; click I accept the terms in the license agreement (Figure 12-5).

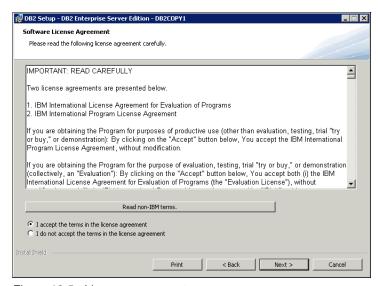


Figure 12-5 License agreement

5. To select the installation type, accept the default of **Typical** and click **Next** to continue (see Figure 12-6).

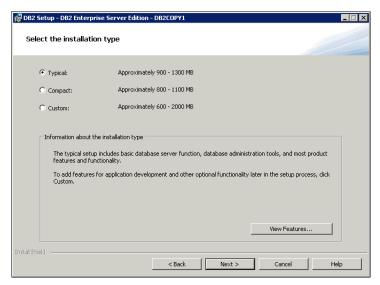


Figure 12-6 Typical installation

Select Install DB2 Enterprise Server Edition on this computer (see Figure 12-7). Click Next to continue.

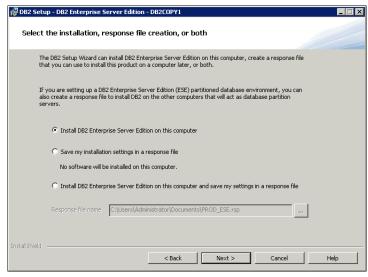


Figure 12-7 Installation action

7. The panel shown in Figure 12-8 shows the default values for the drive and directory to be used as the installation folder. You can change these or accept the defaults, then click **Next** to continue. In our installation, we accept to install on the C: drive.

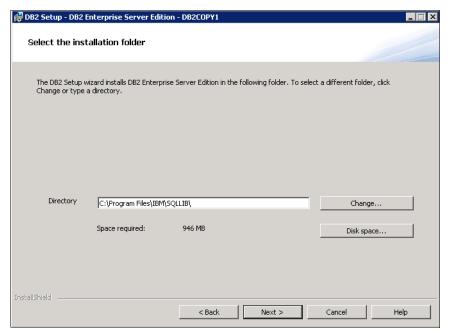


Figure 12-8 Windows Installation Folder

The next panel 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 is pre-filled with a default user name. You can change it or leave the default and type the password of the DB2 user account that you want to create (see Figure 12-9). Leave the check-box **Use the same user name and password for the remaining DB2 services** checked and click **Next** to continue.

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.



Figure 12-9 User Information

9. In the Configure DB2 instances panel, select **Create the default DB2 Instance** and click **Next** as shown in Figure 12-10.

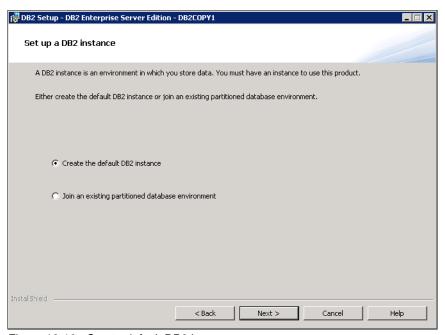


Figure 12-10 Create default DB2 instance

10. Select **Single Partition Instance** and click **Next** as shown in Figure 12-11 on page 472.

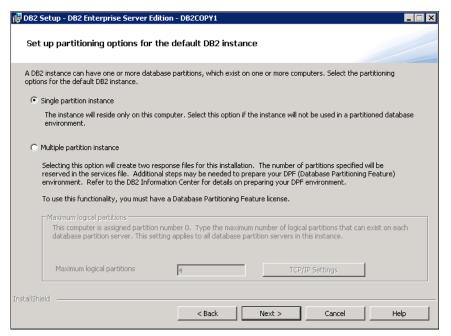


Figure 12-11 Single Partition Instance

11. Accept the default DB2 Instance and click **Next** to continue (see Figure 12-12).

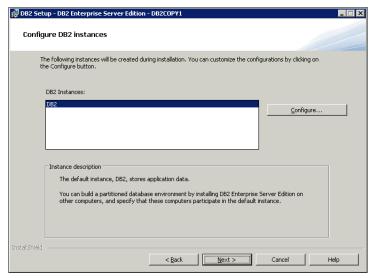


Figure 12-12 Configure DB2 instances

12. The next panel allows you to specify options to prepare the DB2 tools catalog, as shown in Figure 12-13. Verify that **Prepare the DB2 tools catalog on this computer** is not selected. Click **Next** to continue.

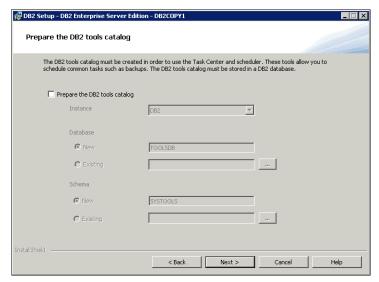


Figure 12-13 Prepare db2 tools catalog

13. The next panel, shown in Figure 12-14, allows you to set the DB2 server to send notifications when the database needs attention. Ensure that the check-box **Set up your DB2 server to send notification** is unchecked and then click **Next** to continue.

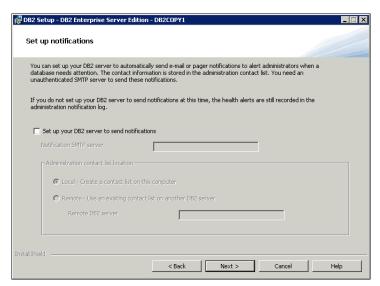


Figure 12-14 Health Monitor

14. Accept the defaults for the DB2 administrators group and DB2 users group in the Enable operating system security for DB2 objects panel shown in Figure 12-15 and click **Next** to proceed.

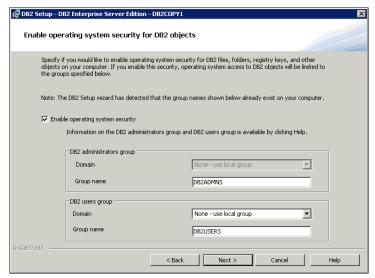


Figure 12-15 Enable operating system security for DB2 objects

15. Figure 12-16 shows the summary panel about what is going to be installed, based on your input. Review the settings and click **Finish** to continue.

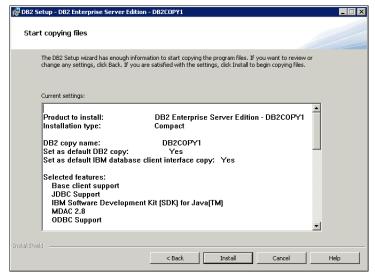


Figure 12-16 Summary panel

The DB2 installation proceeds and you see a progress panel similar to the one shown in Figure 12-17.

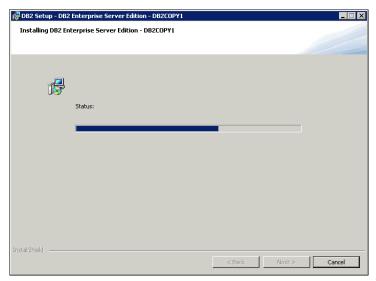


Figure 12-17 DB2 Enterprise Server Edition installation progress

16. When the setup completes, click **Finish**, as shown in Figure 12-18.

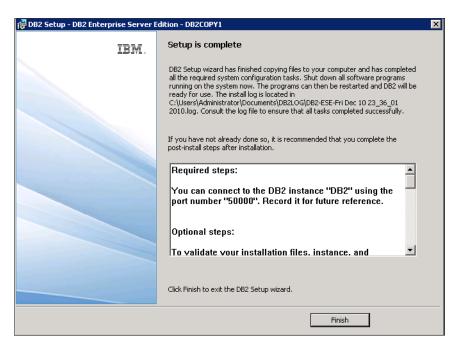


Figure 12-18 Setup is Complete

III DB2 First Steps DB2 Database for Linux, UNIX, and Windows, Version 9.7 First Steps for DB2 Database for Linux, UNIX, and Windows Welcome to First Steps for DB2 for Linux, UNIX, and Windows, First Steps directs you to the tasks that you should perform after you install the DB2 product. Check for news about Product Updates. Start Check for Product Updates Select the option which best describes your situation to be directed to information most relevant to you: . You are new to DB2, and you want to learn more about the product. o Create the SAMPLE database Create SAMPLE database o Create your own database. o [Create new database O Access product information. New to DB2 data server products You are upgrading from a previous version of DB2, and you want to learn about the enhancements in the new release and how you can take advantage of them. o Access product information. Upgrading from a previous version • You are enabling your applications from another data server product to DB2 data server, and you want to ensure that you are proceeding correctly. o Create the SAMPLE database

17. Close the DB2 Setup Launchpad (Figure 12-19) to complete the installation.

Figure 12-19 DB2 First Steps

Verifying the DB2 installation

Follow these steps to verify the DB2 installation:

o Create your own database.

Create SAMPLE database

- Launch a DB2 Command window: Start → IBM DB2 → DB2COPY1 (Default) → Command Line Tools → Command Window.
- 2. Create the SAMPLE database, entering the db2samp1 command as shown in Figure 12-20.



Figure 12-20 Create the SAMPLE database

3. Enter the following DB2 commands. 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.

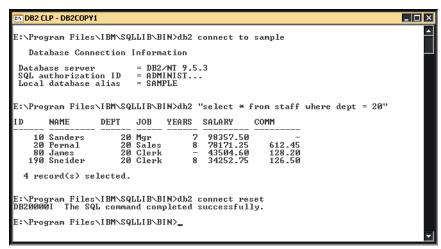


Figure 12-21 DB2 commands results

4. If you did install DB2 Enterprise Server 9.7 restart the system now.

Note: If you do not reboot your server or at least the DB2 service after installing DB2 level 9.7 the TPC installation will fail. For more details refer to:

http://www-01.ibm.com/support/docview.wss?uid=swg21452614

12.4 Installing Tivoli Storage Productivity Center components

Make sure that all prerequisites as described in 12.3.3, "Pre-installation steps for Windows" on page 464 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.

Follow these steps:

- If Windows autorun is enabled, the installation program ought to start 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** (see Figure 12-22).



Figure 12-22 Language selection panel

3. The License Agreement panel is displayed. Read the terms and select I accept the terms of the license agreement. Then click Next to continue as shown in Figure 12-23.

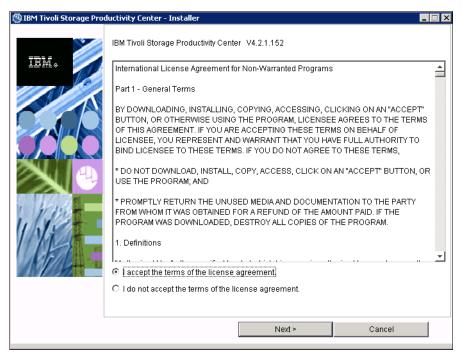


Figure 12-23 Accept License Agreement

4. Select **Typical installation** and make sure you have marked all checkboxes as shown in Figure 12-24.

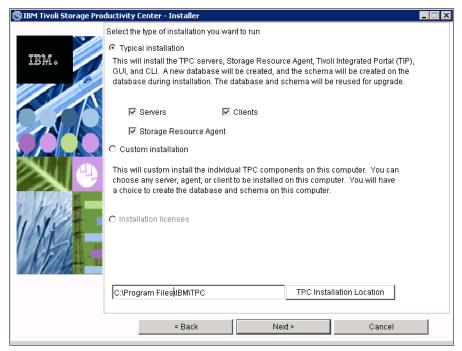


Figure 12-24 Typical Installation

5. Enter the User ID and the Password you have selected during DB2 installation as shown in Figure 12-25.

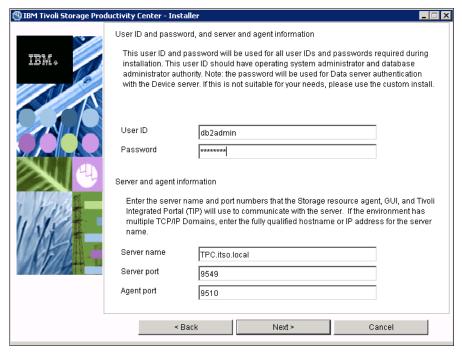


Figure 12-25 Enter DB2 User

6. Specify the location to install TIP as shown in Figure 12-26. Note if you install TPC on 64 bit windows the default path will include "Program Files (x86)". Remove x86 otherwise you will get an error message and you will not be able to proceed.

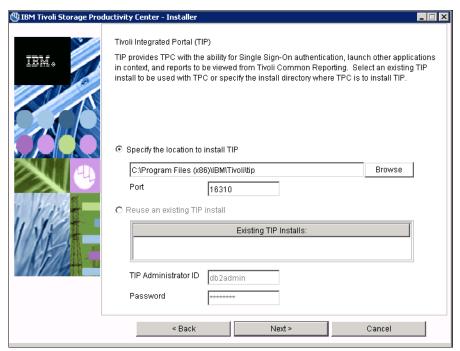


Figure 12-26 Select TIP location

7. Select your preferred authentication method and click **Next** as shown in Figure 12-27.

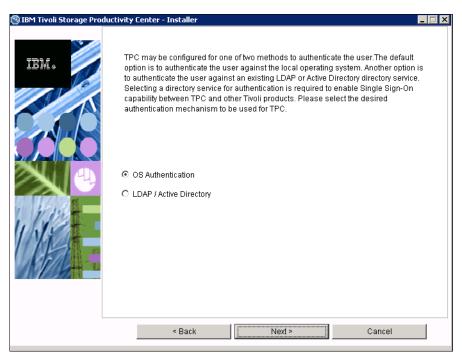


Figure 12-27 Select Authentication Method

8. Review the summary report and click **Install** to start the installation task as shown in Figure 12-28 on page 480.

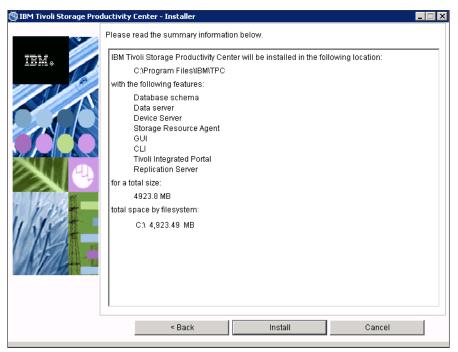


Figure 12-28 Summary Information

The installation task will now start and take some time. During the installation you will be prompted to select the TPC for Replication settings as shown in Figure 12-29. Click Next.



Figure 12-29 TPC for Replication Wizard

10. To perform a prerequisites check click Next (Figure 12-30 on page 481).

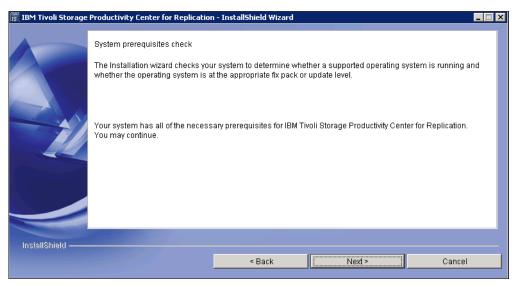


Figure 12-30 System Prerequisites Check

11. Select I accept the terms of the license agreement and click Next (Figure 12-31).



Figure 12-31 Accept License Agreement

12. Select an installation path and click **Next** as shown in Figure 12-32 on page 482.

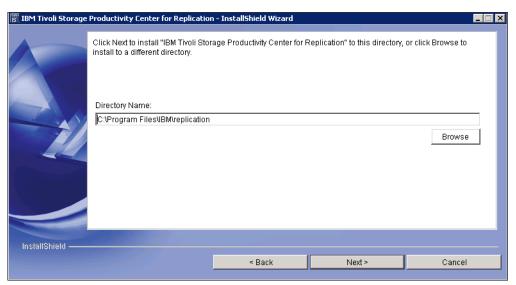


Figure 12-32 Select Installation Path

13. Keep the default ports and click **Next** as shown in Figure 12-33.

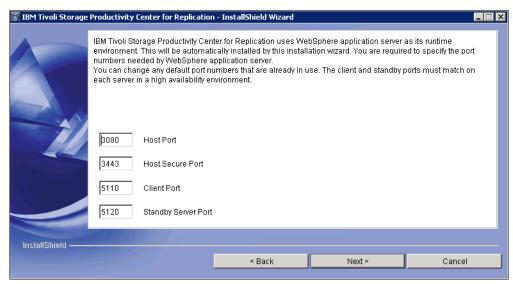


Figure 12-33 Select Ports

14. Review the installation summary and click Install as shown in Figure 12-34 on page 483.

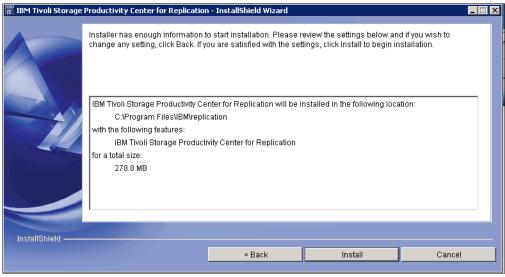


Figure 12-34 TPC for Replication Install Summary

15.TPC for Replication will be installed as shown in Figure 12-35.

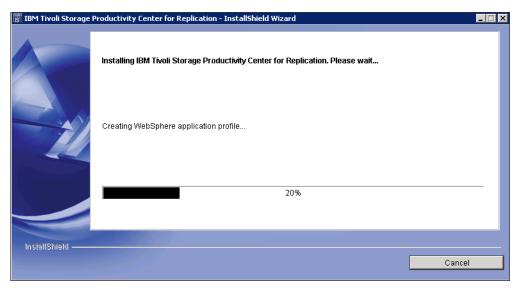


Figure 12-35 Installation Progress

16. After the installation has completed click **Finish** as shown in Figure 12-36 on page 484.

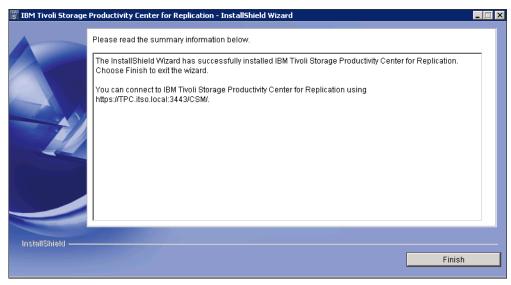


Figure 12-36 TPC for Replication Installation Complete

17. The view will jump back to the TPC Installation as shown in Figure 12-37.

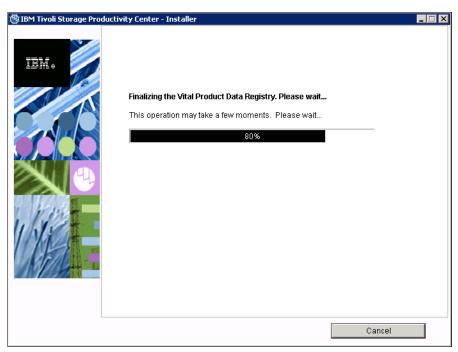


Figure 12-37 TPC Installation Progress

18. After some minutes the TPC installation process completes, click **Finish** to close the wizard (Figure 12-38 on page 485).

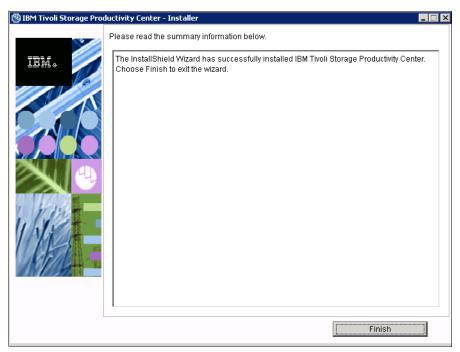


Figure 12-38 TPC Typical Installation Wizard Complete

TPC installation using the typical path is complete now.

12.5 Connect TPC to IBM Storwize V7000

After the typical installation, TPC is ready to connect to IBM Storwize V7000 now. Perform the following steps:

Start the TPC GUI by selecting Start --> All Programs --> IBM Tivoli Storage
 Productivity Center --> Productivity Center. You will be prompted to authenticate as shown in Figure 12-39. Click OK.



Figure 12-39 Enter Credits

2. The welcome panel will appear (Figure 12-40 on page 486). Click **Add Devices** to connect to a new system.

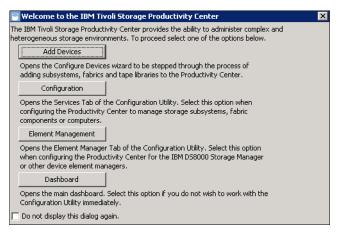


Figure 12-40 Welcome Panel

3. Select **Storage Subsystem** and click **Next** as shown in Figure 12-41.

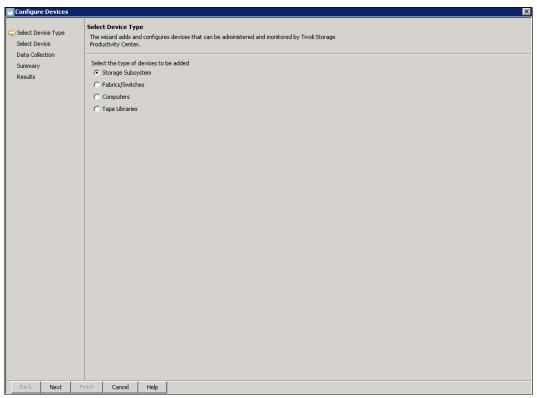


Figure 12-41 Select Device Type

4. Select **Add and configure new storage subsystems** and click **Next** (Figure 12-42 on page 487).

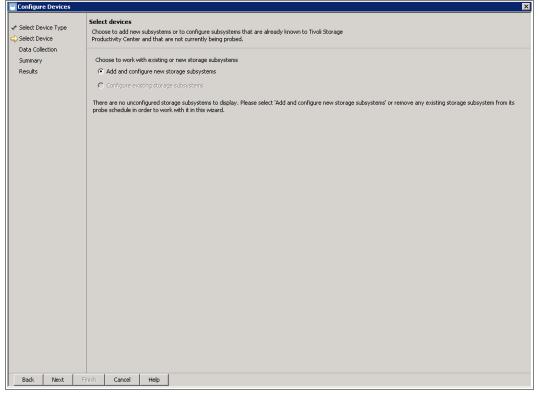


Figure 12-42 Add and configure new storage subsystems

- 5. Select or enter the following settings to connect to IBM Storwize V7000 (Figure 12-43 on page 488).
 - a. Device Type: IBM SAN Volume Controller / IBM Storwize V7000
 - b. Software Version: 5+
 - c. IP Address: Enter Storwize V7000 Cluster IP
 - d. Select Key: Upload new key
 - e. Admin Username: superuser
 - f. Admin Password: enter the superuser password (default = passw0rd)
 - g. Username: Select a User to connect
 - h. Private SSH Key: Specify the location of the private ssh key (if you don't have a key now create one as described in Appendix A, "CLI setup and SAN Boot" on page 543.

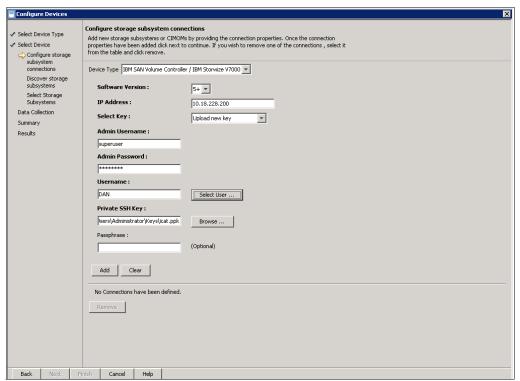


Figure 12-43 Configure storage subsystem connections

6. Click **Add** to connect to your IBM Storwize V7000 System as shown in Figure 12-44 on page 489.

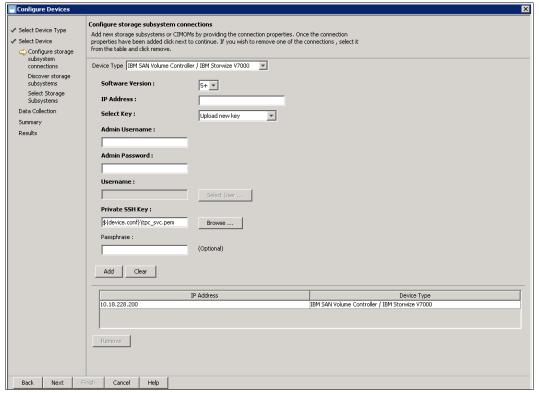


Figure 12-44 Configure storage subsystem connections - system discovered

7. Repeat these steps to add another system or click **Next** to complete the discovery of the new storage subsystems (Figure 12-45).

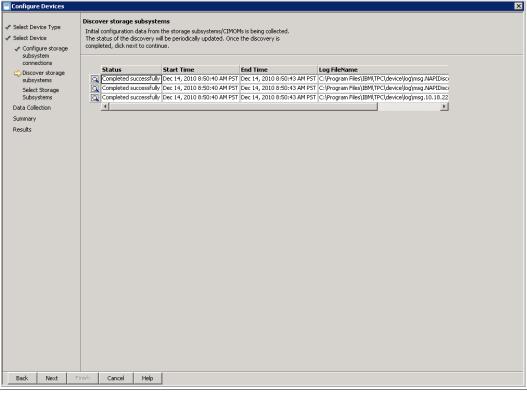


Figure 12-45 Discover Subsystems

8. After the discovery select the new storage subsystem that you wish to add to the configuration as shown in Figure 12-46 on page 490.

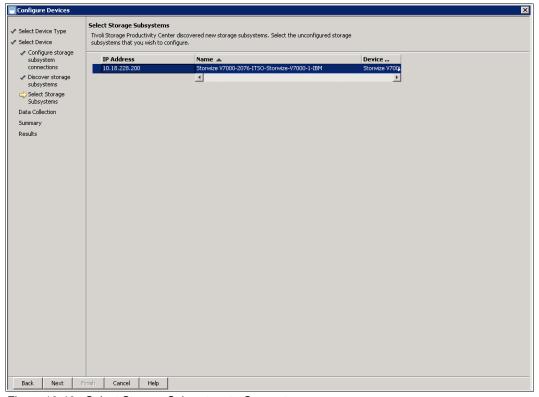


Figure 12-46 Select Storage Subsystem to Connect

 Specify the data collection settings and to add the new system to the default group select Monitoring Group and Subsystem Standard Group as shown in Figure 12-47 on page 491.

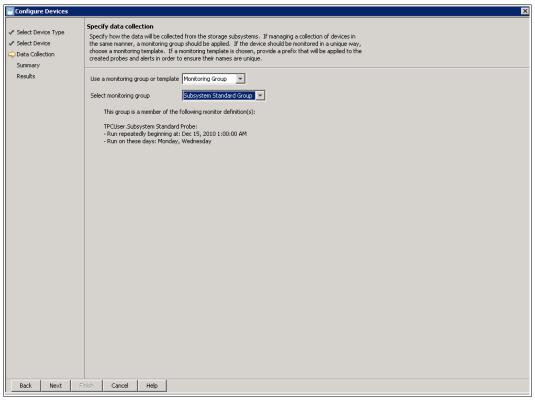


Figure 12-47 Specify Data Collection

10. Review your selections and click **Next** as shown in Figure 12-48 on page 492 to add the new device.

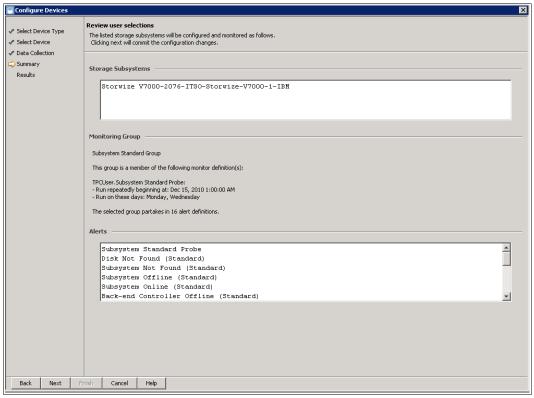


Figure 12-48 Review User Selections

11. The IBM Storwize V7000 has now been added to TPC, click **Finish** to close the wizard (Figure 12-49 on page 493).

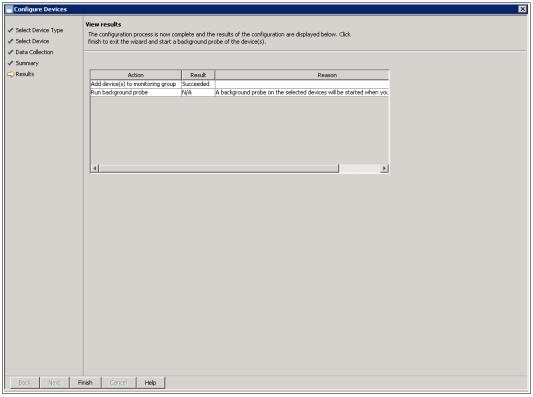


Figure 12-49 View Results

12. You will be prompted if you would like the job history, but we will not view it now, click **Close** (Figure 12-50).



Figure 12-50 Job History

The IBM Storwize V7000 has now been added successfully to be administered by Tivoli Storage Productivity Center. Of course the normal IBM Storwize V7000 GUI and CLI is still available, and can be used to manage the system as well.

12.6 TPC Example Views

This section shows example views to give a first impression of TPC with the Storwize V7000. The left tree gives you many selection possibilities to select different reports as shown in Figure 12-51 on page 494. If you highlight a report in the left panel the associated view appears on the right. This example shows a report about the configured volumes on the IBM Storwize V7000, and Figure 12-52 on page 495 shows you details about the configured MDisks.

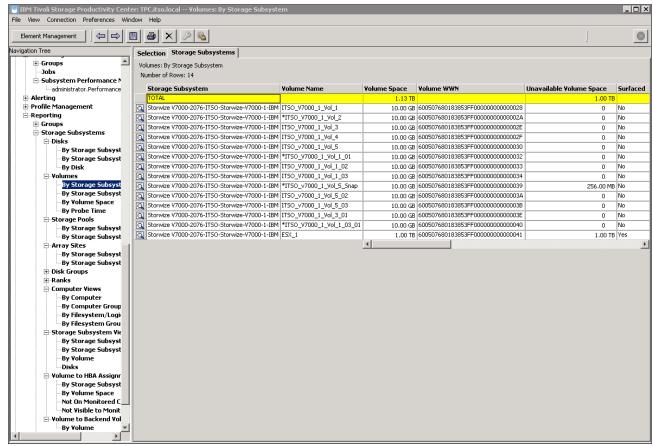


Figure 12-51 Volumes by Storage Subsystem

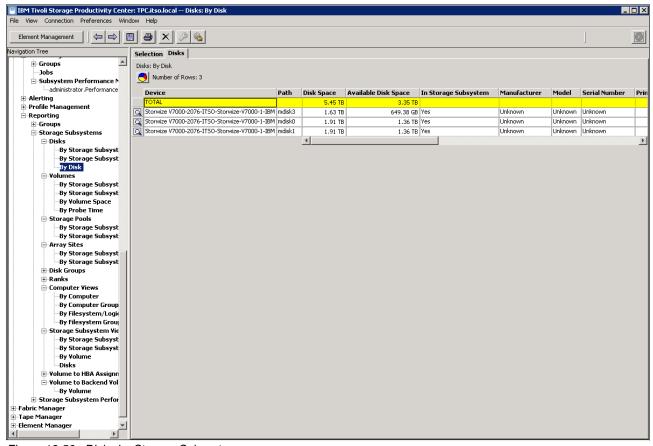


Figure 12-52 Disks by Storage Subsystem

But it is not only possible to monitor configuration data, there are many more possibilities such as performance statistics, health monitoring, path usage, and so on. Figure 12-53 on page 496 shows you statistics for a single disk pool.

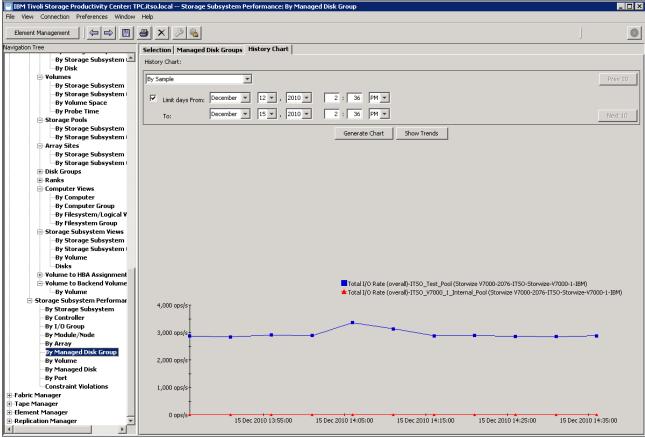


Figure 12-53 Storage Subsystem performance by Managed Disk Group (Storage Pool): Total I/O Rate

A detailed description about TPC is beyond the intended scope of this book.



13

RAS, Monitoring and Troubleshooting

There are many ways to monitor and troubleshoot the IBM Storwize V7000. This chapter contains the following topics:

- ► RAS
- ► Hardware and LEDs description
- Monitoring from a Host
- Monitoring from the IBM Storwize V7000
- ► Backup procedure
- Software upgrade procedure
- ► Troubleshooting
- ► Recommended Actions
- Event Log Navigation
- Audit Log Navigation
- Support
- ► Shutting down the IBM Storwize V7000
- ► Shutting down the infrastructure with IBM Storwize V7000

13.1 Reliability, Availability, and Serviceability on the IBM Storwize V7000

Reliability, availability, and serviceability (RAS) are important concepts in the design of the IBM Storwize V7000. Hardware features, software features, design considerations, and operational guidelines all contribute to make the IBM Storwize V7000 extremely reliable.

Fault tolerance and high levels of availability are achieved by:

- The RAID capabilities of the underlying disk subsystems
- ► IBM Storwize V7000 nodes clustering using "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 achieved by providing:

- Cluster error logging
- ► Asynchronous error notification
- Dump capabilities to capture software detected failures
- ► Concurrent diagnostics
- Directed maintenance procedures
- Concurrent log analysis and 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 IBM Storwize V7000 nodes canister in a cluster.
- ► Software recovery via 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 e-mail

At the heart of the IBM Storwize V7000 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 looks at the RAS features of the IBM Storwize V7000, monitoring and troubleshooting.

13.1.1 Node Canister

The two node canisters are located in the control enclosure and they work as a cluster.

Figure 13-1shows the ports and indicator lights of a node canister.

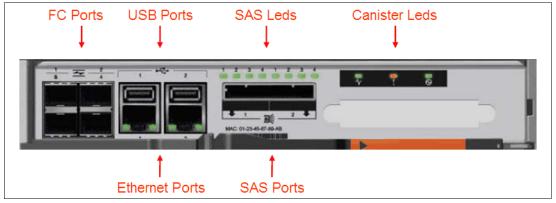


Figure 13-1 Ports and indicators of a node canister

Fibre Channel

There are four FC ports on the left hand side of the canister. They are in two rows of two connectors. The ports are numbered 1 to 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 are not shown in the figure but are located between the two rows of ports and are triangular, pointing towards the port to which they refer. On the left is the Speed LED and on the right the Link LED. In Table 13-1 we explain the status of the indicators:

Table 13-1 Fiber Channel link LEDs status

| Speed LED | Link LED | Link Status |
|-----------|----------------|-----------------------|
| OFF | OFF | INACTIVE |
| OFF | ON OR BLINKING | IDLE OR BUSY AT 2Gbps |
| BLINKING | ON OR BLINKING | IDLE OR BUSY AT 4Gbps |
| ON | ON OR BLINKING | IDLE OR BUSY AT 8Gbps |

USB

There are two USB connectors side by side, and they are numbered as 1 on the left and as number 2 on the right. There are no indicators associated with the USB ports.

Ethernet and LED status

There are two 10/100/1000 Mbps Ethernet ports side by side on the canister and they are numbered 1 on the left and 2 on the right. Each port has two LEDs and the status is shown in Table 13-2.

Table 13-2 Ethernet LEDs status

| LED | Color | Meaning |
|------------|--------|--|
| Link state | Green | It is ON when there is an ethernet link. |
| Activity | Yellow | It is flashing when there is activity on the link. |

SAS and LED status

There are two 6 Gbps SAS ports side by side on the canister. 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 directly above the ports. For each port they are numbered 1 through 4. The LED indicates activity on the link and the status is shown in Table 13-3.

Table 13-3 SAS LED status

| State | Meaning |
|----------|--|
| ON | Link is connected |
| FLASHING | Link is connected and being used for I/O |
| OFF | Link is not connected |

Canister and LED status

There are three LEDs in a row towards the top right of the canister that indicate the status and identification for the node as shown in Table 13-4 on page 500.

Table 13-4 Node Canister LEDs status

| Position | Color | Name | State | Meaning |
|----------|--------------|-------|------------------|---|
| Left | Left Green C | | ON | The node is in the active or starting state. It may 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 then there is a problem establishing a cluster for example due to lack of quorum. |
| | | | FLASHING | The node is in 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 system POST. |
| Middle | Amber | Fault | ON | The node is supposed to be in a cluster but is unable to reach active state. |
| | | | FLASHING | The canister is being identified. There may be a fault |
| | | | OFF | The node is in candidate or active state. This does not mean there is not a hardware error on the node. |
| Right | Green | Power | ON | Canister is powered on and the CPUs are running. |
| | | | FAST FLASHING | Canister is running Power On Self Test (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 and LED status

As shown in Figure 13-2 on page 501 there are two 6 Gbps SAS ports side by side on the canister. They are numbered 1 on the left and 2 on the right. Each port connects four PHYs; each PHY is associated with a LED. These LEDs are green and are next to the ports.

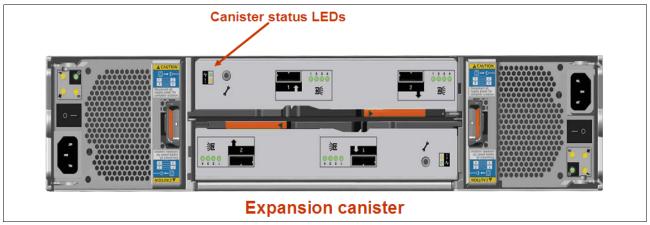


Figure 13-2 Canister status LEDs

In Table 13-5 on page 501 we describe the LED status of the expansion canister.

Table 13-5 Expansion Canister LEDs status

| Position | Color | Name | State | Meaning |
|----------|-------|--------|----------|---|
| Тор | Green | Status | ON | Canister active |
| | | | FLASHING | Canister has a vpd error |
| | | | OFF | Canister not active |
| Bottom | Amber | Fault | ON | Canister hardware is Faulty |
| | | | FLASHING | Canister is being identified. |
| | | | OFF | No Fault Canister is not being identified |

13.1.2 Disk subsystem

The IBM Storwize V7000 system is made up of enclosures. There are two kind of enclosures: a 2U12 that takes 12 3.5" drives; and a 2U24 that takes 24 2.5" drives. The drives fit into the front of the enclosure, 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. 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 desired 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.

IBM Storwize V7000 supports the RAID levels shown in Table 13-6.

Table 13-6 RAID levels supported by IBM Storwize V7000

| RAID level | where data is striped | Minimum-Maximum member |
|------------|---|------------------------|
| 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 on 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 process. The process is to check that 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 Disks

The SSD drives are treated no differently by IBM Storwize V7000 than HDDs with respect to RAID arrays or MDisks. The individual SSD drives in the storage managed by the IBM Storwize V7000 will be combined into an array, usually in RAID10 or RAID5 format. It is unlikely that RAID6 SSD arrays will be used due to the double parity overhead, with two SSD logical drives used for parity only.

A LUN will be 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 will help 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 as shown in Figure 13-3 on page 503 The SAS network is made up of strands and chains.

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 503 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 means 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 connected as shown in Figure 13-3 on page 503. 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 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.

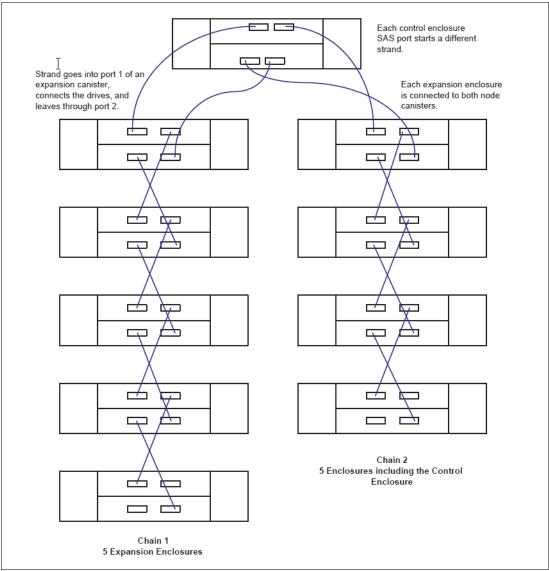


Figure 13-3 SAS Cabling

13.1.3 Power

All enclosures require two 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 on page 504 shows the two PSUs present in the controller and expansion enclosure, the controller PSU have 2 more LEDS than the expansion enclosure due to the battery status indication. Table 13-7 on page 504 shows the meaning of the LEDs in both enclosure:

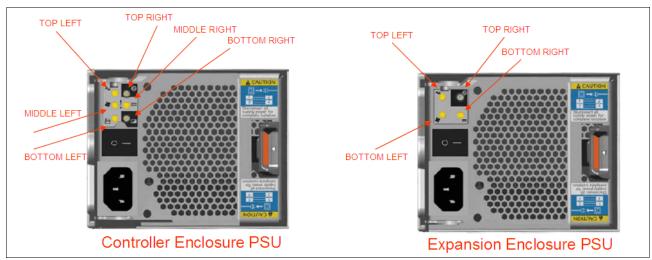


Figure 13-4 Controller and Expansion Enclosure LED status indicator

Table 13-7 LED status of Power in the Controller and Expansion Enclosure

| Position | Color | Meaning |
|--|-------|---------------|
| Top left | Amber | Ac Failure |
| Top right | Green | PSU Ok |
| Middle left (Controller Enclosure) Bottom left (Expansion Enclosure) | Amber | Fan Failure |
| Middle right (Controller Enclosure) Bottom left (Expansion Enclosure) | Amber | Ac Failure |
| Bottom left (Controller Enclosure only) | Amber | Battery Fault |
| Bottom right (Controller Enclosure only) | Green | Battery Good |

Table 13-8 shows the meaning of the LEDs for the PSU on the controller and expansion enclosure.

Table 13-8 PSU LEDs status

| PSU Ok | Ac Failure | Fan Failure | Dc Failure | Meaning |
|--------|------------|-------------|------------|---|
| ON | ON | ON | ON | Communication failure between PSU and enclosure midplane. |

| PSU Ok | Ac Failure | Fan Failure | Dc Failure | Meaning |
|----------|------------|-------------|------------|--|
| ON | OFF | OFF | OFF | PSU is on and working correctly |
| FLASHING | ANY | ANY | ANY | Neither canister is working: mis-seated, unplugged, broken |
| OFF | ON | ON | ON | Communication failure and PSU problem |
| OFF | FLASHING | FLASHING | FLASHING | PSU firmware download in progress |
| OFF | ON | OFF | ON | No AC input |
| OFF | OFF | ON | OFF | Fan failed |
| OFF | OFF | OFF | ON | PSU not seated correctly |
| OFF | OFF | OFF | OFF | No power to enclosure |

Table 13-9 shows the meaning of the LEDs for the battery on the controller enclosure

Table 13-9 Battery LEDs status

| Battery Good | Battery Fault | Meaning |
|-----------------|------------------|---|
| ON | OFF | Battery is OK and fully charged |
| FLASHING | OFF | Battery is OK but not fully charged. Either charging or doing maintenance discharge |
| OFF | ON | Battery is broken |
| OFF | FLASHING | Battery is broken but can be serviced |

13.2 Configuration backup procedure

You can download a Configuration Backup file using the IBM Storwize V7000 GUI. This is a manual procedure and is not automatic.

Regularly saving a configuration backup file of the IBM Storwize V7000 is important, we recommend that you download this file regularly to your management workstation to protect the data (best practice is automate with script and save it every day).

This file must be used if there is a serious failure that requires you to restore your system configuration.

The backup file is updated by the cluster everyday, It is important to save it after any changes to your system configuration. It contains configuration data such as arrays, pools, volumes and so on (no customer applications data)

You can use the GUI or the CLI to generate it. You must use the CLI command svcconfig backup to produce a new backup file, it is not possible at this time generate a new backup file with the GUI but you can only save it.

Generating a configuration backup using the CLI

Further information on the Command Line Interface in the IBM Storwize V7000 is in Appendix A, "CLI setup and SAN Boot" on page 543.

You must have the following prerequisites in place:

- ► All nodes must be online.
- ▶ No object name can begin with an underscore.
- All objects must have non-default names, that is, names that are not assigned by the IBM Storwize V7000.

Although we recommend that objects have non-default names at the time that the backup is taken, this prerequisite is not mandatory. Objects with default names are renamed when they are restored.

To generate a Configuration backup from the Command Line Interface enter this command:

svcconfig backup

Example 13-1 shows an example of the svcconfig backup command.

Example 13-1 svcconfig backup command

The **svcconfig backup** command creates three files that provide information about the backup process and cluster configuration. These files are created in the /tmp directory of the configuration node.

Table 13-10 describes the three files that are created by the backup process:

Table 13-10

| 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. |

Downloading a configuration backup using the Web GUI

To download a backup of the configuration using the GUI follows these steps:

1. Open the **Troubleshooting** menu and choose the **Support** option as shown in Figure 13-5.

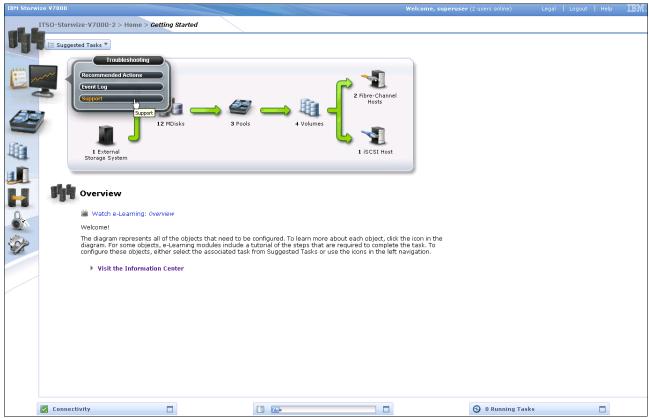


Figure 13-5 Support option

2. Clicking on the **Support** option will display the screen shown in Figure 13-6 on page 508. From this screen, click on **Show full log listing...** to display all log files.

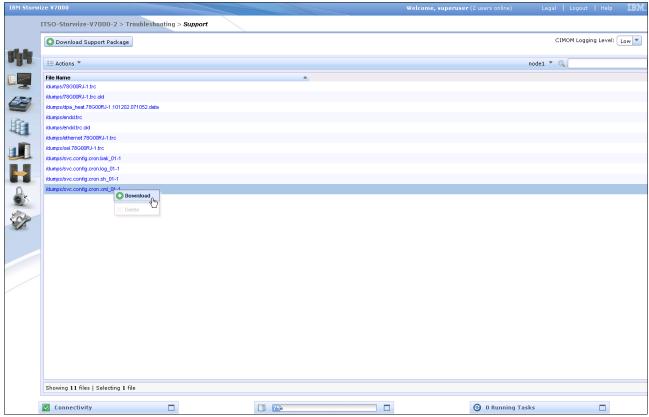


Figure 13-6 Show full log files screen

3. Look for a file named "/dumps/svc.config.backup.xml_*", and selecting it you can right click on it to **Download** it to your workstation.

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, perform the following steps:

1. With a supported Web browser, replace <your-cluster-ip-address> with your cluster IP address here:

http://<your-cluster-ip-address>

You will be taken to the IBM Storwize V7000 GUI login screen as shown in Figure 13-7.



Figure 13-7 IBM Storwize V7000 GUI Login screen

Login with your superuser/password and you will get into the IBM Storwize V7000
management home page. From there, go to Configuration → Advanced menu as shown in
Figure 13-8 on page 509 and click on Advanced.

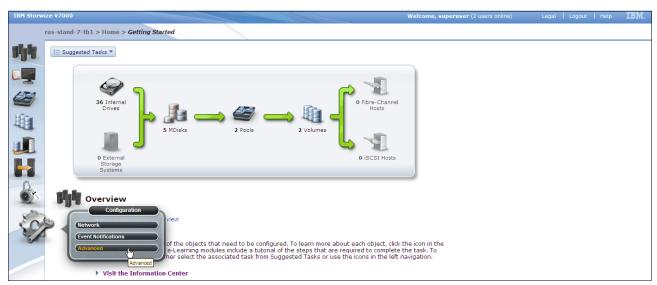


Figure 13-8 Configuration menu

3. In the Advanced menu click on **Upgrade Software** and you will get to the screen as shown in Figure 13-9.



Figure 13-9 Upgrade Software menu

From the screen shown in Figure 13-9, you can click on the following buttons:

► Check for updates — this will check directly on the IBM Web site if there is a newer IBM Storwize V7000 software version rather than the version you have installed in your IBM Storwize V7000 and you will need an internet connection to be able to do this.

Note: We recommend that you 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 https://www.ibm.com/support/mynotifications (an IBM ID is required). Add the Storwize V7000 to the notifications list to be advised of support information and to download the latest code to your workstation for later upload.

- ► Launch the Upgrade Wizard this will launch the software upgrade process.
- 4. Click on **Launch Upgrade Wizard** to start the upgrade process and you will be redirected to the screen shown in Figure 13-10 on page 510.

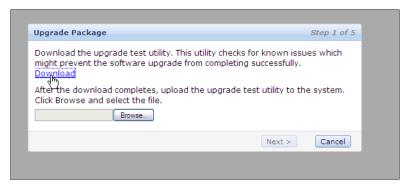


Figure 13-10 upgrade package screen

From the screen shown in Figure 13-10 you can download the Upgrade Test Utility, or if you downloaded it previously you can browse to the location where you saved it as shown in Figure 13-11.

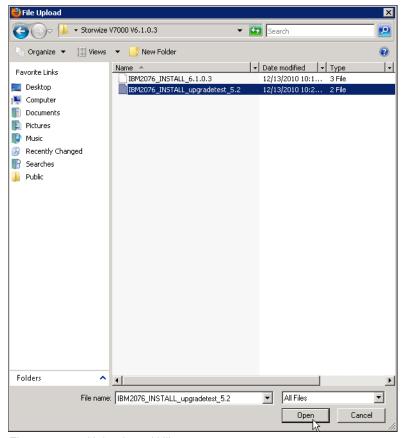


Figure 13-11 Upload test Utility

5. When the Upgrade Test Utility has been uploaded, you will get the screen as shown in Figure 13-12 on page 511.

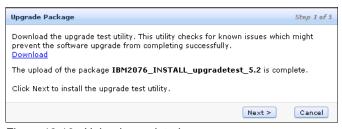


Figure 13-12 Upload completed

6. Click **Next** in Figure 13-12 on page 511 and the Upgrade Test Utility will be applied and you will be redirected to the screen as shown in Figure 13-13.

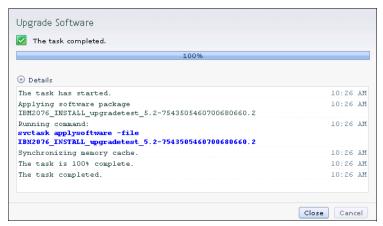


Figure 13-13 Upgrade test utility applied

Click Close on Figure 13-13 and you will be get the screen shown in Figure 13-14, where you can run the Upgrade Test Utility.



Figure 13-14 Run upgrade test utility

8. Click **Next** and you will be redirected to the screen shown in Figure 13-14. At this time the Upgrade Test Utility will run and you will be able to see the suggested actions, if any, or the screen shown in Figure 13-15.

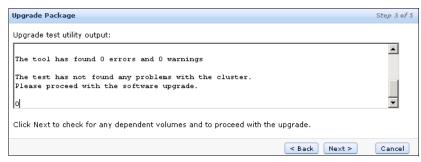


Figure 13-15 Results test utility

9. Click **Next** on Figure 13-15 to start the software upload procedure, and you will be redirected to the screen shown in Figure 13-16.



Figure 13-16 Downloading Software code

From the screen shown in Figure 13-16 you can download the SVC software upgrade package, or you can browse and upload the software upgrade package from the location where you saved it as shown in Figure 13-17.

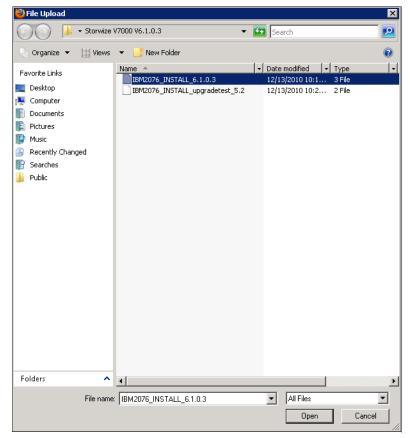


Figure 13-17 Downloaded code

Click **Open** in Figure 13-17 and you will be redirected to the screens shown in Figure 13-18 and Figure 13-19.



Figure 13-18 Uploading software code

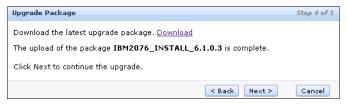


Figure 13-19 Upload complete

10. Click **Next** and you will be redirected to the screen shown in Figure 13-20.



Figure 13-20 System ready to upgrade

11. Click **Finish** on Figure 13-20 and the software upgrade will start and you will be redirected to the screen shown in Figure 13-21. Clicking Close as in Figure 13-21 will give you the warning message as shown in Figure 13-22.

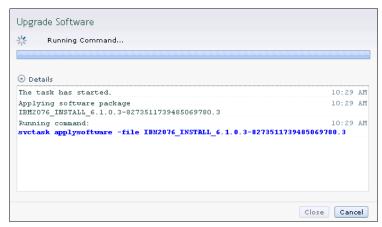


Figure 13-21 Upgrading software code



Figure 13-22 Warning message

- 12. Click OK on Figure 13-22 and you have completed your task to upgrade the SVC software
- 13. You will get messages that let you know that first one node, then the other, has been upgraded. When both nodes have been rebooted, you have completed your SVC software upgrade.

13.4 Monitoring host paths to the IBM Storwize V7000

In this section we show how to monitor paths to the IBM Storwize V7000. On host systems using the IBM multipath module such as Windows, and other operating systems a powerful command line tool is available (refer to basic Windows configuration in Chapter 4, "Host Configuration" on page 127 to monitor IBM Storwize V7000 connectivity from the host. Example 13-2 shows the output of the **datapath query device** command and it shows that four paths should be available, but two of the paths have been closed which may indicate a problem.

Example 13-2 pcmpath query device showing a problem with one canister

C:\Program Files\IBM\SDDDSM>datapath guery device

Total Devices: 2

DEV#: 0 DEVICE NAME: Disk2 Part0 TYPE: 2145 POLICY: OPTIMIZED

SERIAL: 600507680280801AC8000000000001F

| 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 BusO/Disk3 PartO | OPEN | NORMAL | 1 | 1 | |
| 1 | Scsi Port5 BusO/Disk3 PartO | OPEN | NORMAL | 78 | 1 | |
| 2 | Scsi Port6 BusO/Disk3 Part0 | CLOSE | NORMAL | 28 | 1 | |
| 3 | Scsi Port6 BusO/Disk3 Part0 | CLOSE | NORMAL | 1 | 1 | |

C:\Program

Files\IBM\SDDDSM>

The datapath query adapter command, as shown in Example 13-3, shows all IBM Storwize V7000 paths that are available to the host. We can see that only Adapter 0 is available, and that Adapter 1 state 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 |

C:\Program

Files\IBM\SDDDSM>

Once the problem is fixed, scan for new disks on your host, and check if all paths are available again as shown in 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: 600507680280801AC8000000000001F

| Path# | Adapter/Hard Disk | State Mode | Select | Errors |
|-------|-----------------------------|-------------|--------|--------|
| 0 | Scsi Port5 BusO/Disk2 PartO | OPEN NORMAL | 95 | 1 |
| 1 | Scsi Port5 BusO/Disk2 PartO | 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

| === | ==== | | ===== | ======== | ====== | | ====== | ======== | ====== |
|-----|------|-------------------|-------|------------|--------|------|--------|----------|--------|
| Pat | h# | Adapter/Hard Disk | | | State | Mode | Select | Errors | |
| | 0 | Scsi | Port5 | BusO/Disk3 | Part0 | OPEN | NORMAL | 1 | 1 |
| | 1 | Scsi | Port5 | BusO/Disk3 | Part0 | OPEN | NORMAL | 95 | 1 |
| | 2 | Scsi | Port6 | BusO/Disk3 | Part0 | OPEN | NORMAL | 31 | 1 |
| | 3 | Scsi | Port6 | Bus0/Disk3 | Part0 | OPEN | NORMAL | 1 | 1 |

C:\Program Files\IBM\SDDDSM>

You can also use again the **datapath query adapter** and check that the FAILED path is back ONLINE (State=NORMAL) as shown in Example 13-5 on page 516.

Example 13-5 datapath query adapter to check 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 |

C:\Program Files\IBM\SDDDSM>

13.5 Troubleshooting overview

Figure 13-23 shows the navigation panel to invoke the Troubleshooting wizard for viewing event logs, running maintenance procedures, and collecting support information if you require IBM assistance to determine the appropriate repair procedures.

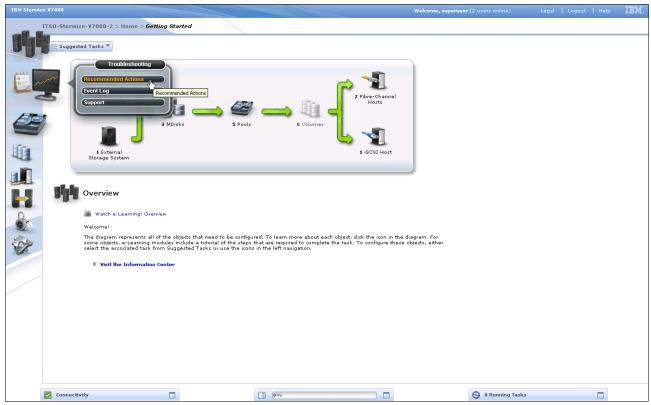


Figure 13-23 Troubleshooting

Opening the Troubleshooting panels brings up a navigation panel (along the top) with the main functions:

- ► Recommended Actions
- Event Log
- ► Support

The Recommended Actions tab will indicate the highest priority maintenance procedure that needs to be executed. Use the troubleshooting wizard first to allow the IBM Storwize V7000 to determine the proper order of maintenance procedures. We clicked on the Recommended Actions and the highest priority event that needs to be fixed is shown in Figure 13-24 on page 518.

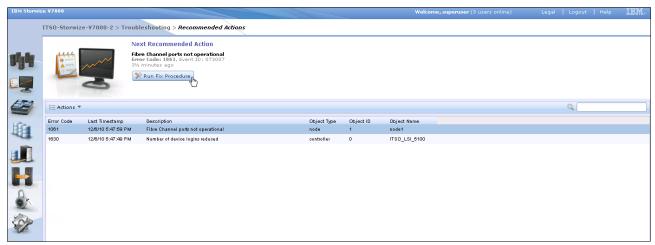


Figure 13-24 Recommended Actions

In our example one Fibre Channel Port was not operational. The next step in this example is to review the physical FC cabling to determine the issue, then click the button **Run Fix Procedure**.

Best practice is to review the event logs and recommended actions periodically to ensure there are no unexpected events, and to configure callhome so that notification of serious events is done immediately.

The Event Log tab includes events that are important to know about including both errors and informational events. An example of the event log is shown in Figure 13-25.

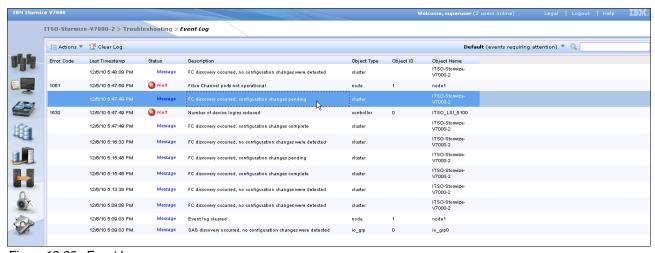


Figure 13-25 Event Log

Another choice from the Event Log is the Support option as shown in Figure 13-26 on page 519. This selection is useful for providing data necessary for IBM Support to determine what the current status is of the IBM Storwize V7000. This function provides several versions of the svc_snap command embedded within the GUI choices.

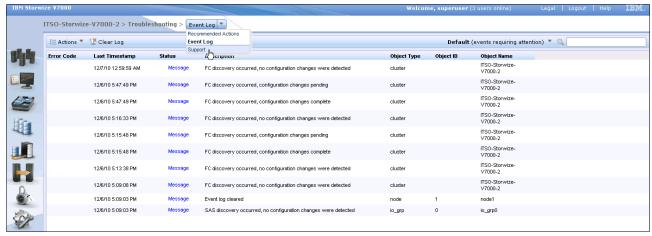


Figure 13-26 Support option from event log

You can also click on **Troubleshooting** and then **Support** as shown in Figure 13-27 and the choices are to either download a support package to your local system, or to display the available logs.

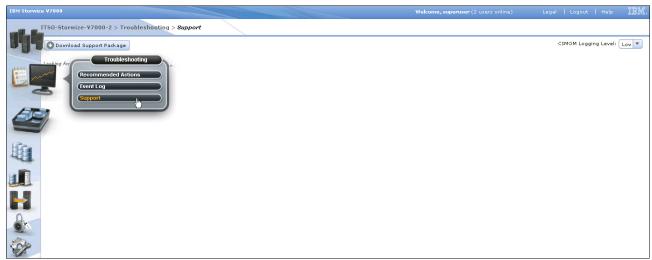


Figure 13-27 Support Tab of Troubleshooting

13.6 Manipulating the Event Log and Recommended Actions Grid

There are multiple views possible of the events and recommended actions available. The GUI is designed to work like any Windows application menu, and therefore the event log grid is manipulated via the row containing the column headings as shown in Figure 13-28 on page 520.

The row is highlighted and then a right-mouse click will pop-up column choices. Check or uncheck the column preferences as needed. This manipulability of the menu columns is also available on the Audit Log Grid as demonstrated briefly in 13.7, "Recommended Actions - details" on page 520.

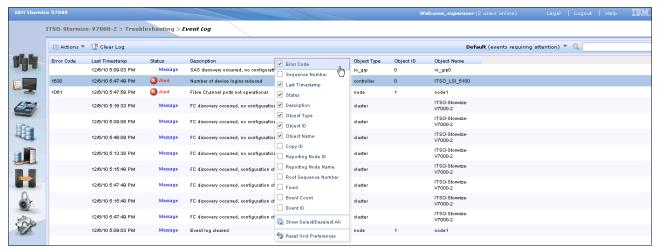


Figure 13-28 Right-mouse click on Event log grid column headings

Every field of the event log is available as a column in the Event Log grid. There are several fields that are useful to add when working with IBM Support and using the **Show ALL** filter, with events sorted by timestamp. These fields are the sequence number, event count, and the fixed state.

Using the selection Reset Grid Preferences will set the grid back to the defaults.

It is also possible to adjust the width of the columns to match your preference as shown in Figure 13-29 by holding the mouse over the column marker, clicking and holding, and then moving the column width to match your preference. This can also be reset by the Reset Grid Preferences.

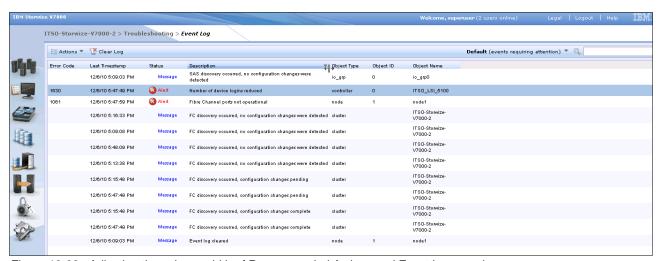


Figure 13-29 Adjusting the column width of Recommended Actions and Event Log panels

13.7 Recommended Actions - details

The IBM Storwize V7000 will determine the most important action that is currently presenting itself in the event log, and highlight this at the top of the Recommended Actions panel. We recommend to check this panel periodically to monitor for any unexpected informational events or errors.

As an example we show how a Fibre Channel cable that is loose presents itself and the IBM Storwize V7000 reported this problem as shown in Figure 13-30 as the highest priority event at the time.



Figure 13-30 Recommended Action: Fibre Channel Ports Not Operational

Next we clicked the Run Fix Procedure button, and started the dialogue in Figure 13-31.

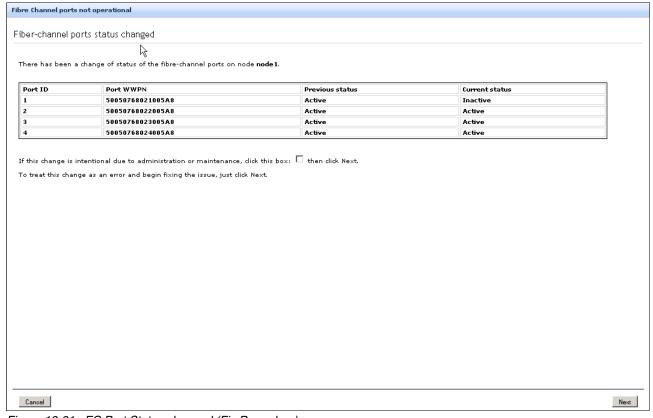


Figure 13-31 FC Port Status changed (Fix Procedure)

Unfortunately the Fibre Channel link is inactive so we select the option shown in Figure 13-32 on page 522 and then click **Next**.

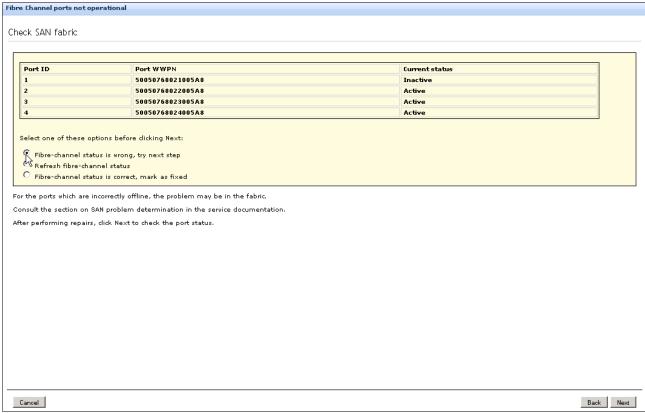


Figure 13-32 FC port status inactive

As suggested by the DMP Recommended Action we changed the FC cable and then we clicked on **Next** as shown in Figure 13-33 on page 523.

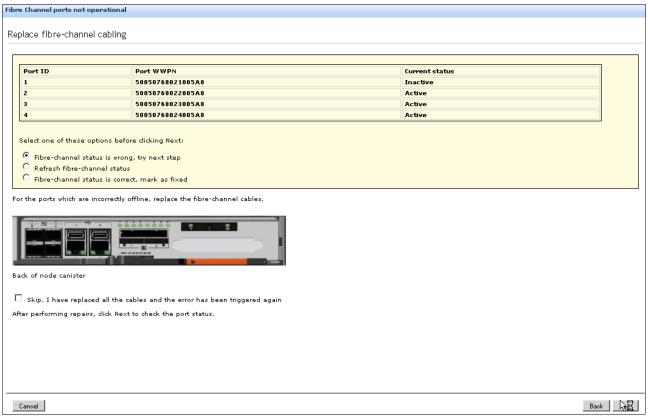


Figure 13-33 Recommended Action DMP Process

Then we select **Refresh fibre-channel status** to verify if the status changes as shown in Figure 13-34 on page 524.

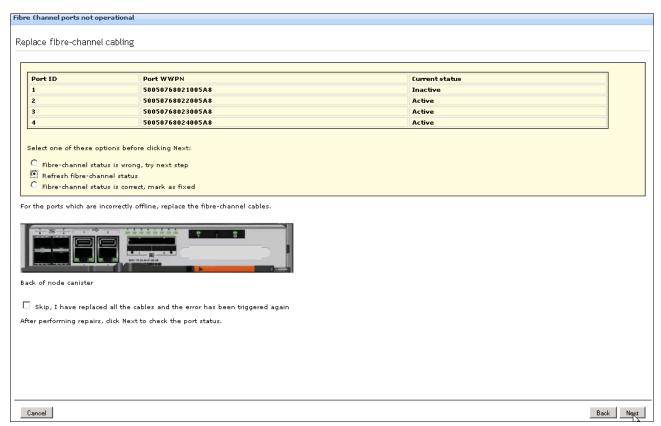


Figure 13-34 Recommended Action DMP Process

After replacing the cable the status is changed from **Inactive** to **Active** as shown in Figure 13-35 on page 525.

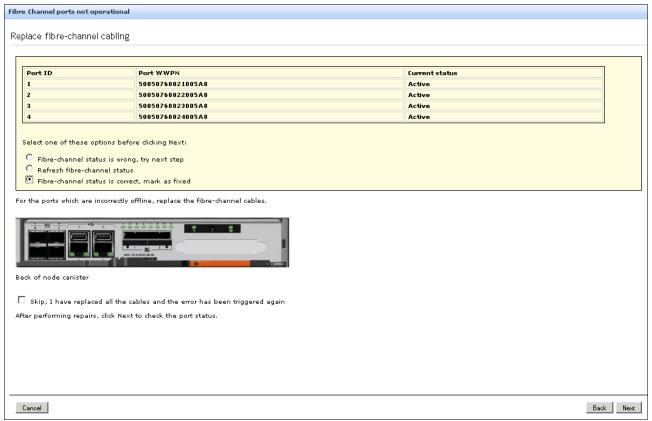


Figure 13-35 Recommended Action DMP Process

Clicking on the **Next** button the error is about to be marked as fixed as shown in Figure 13-36 on page 526.

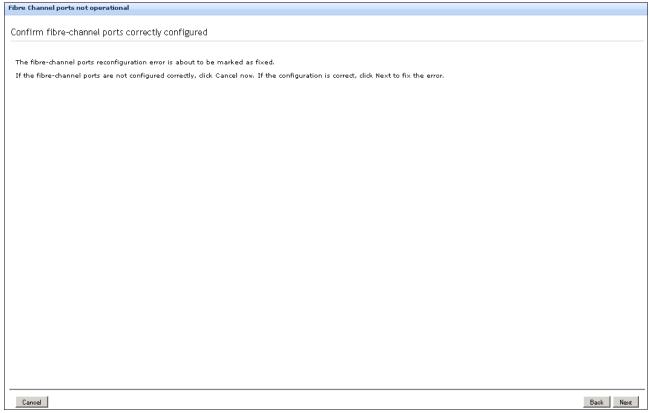


Figure 13-36 Recommended Action DMP Process

Finally, you get the last panel of the DMP procedure and the error has been marked as fixed. Click on the **Close** button to end the procedure as shown in Figure 13-37 on page 527.

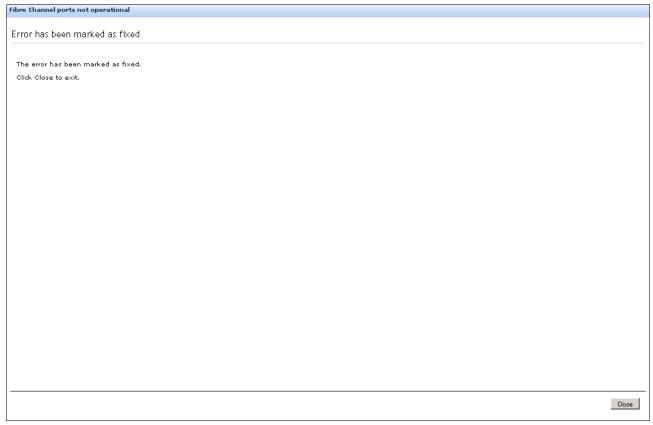


Figure 13-37 Recommended Action DMP Error marked as fixed

Returning to the main Recommended Actions grid, it is also possible to mark this error fixed or view the event details. This is shown in Figure 13-38.

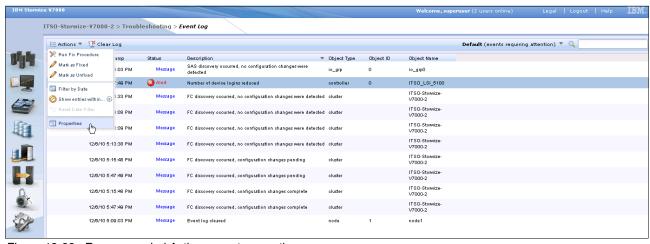


Figure 13-38 Recommended Actions event properties

Clicking on the **Properties** button drop-down one can view all the properties of this event log, such as which node canister, which port, which wwpn, and which event code as displayed in Figure 13-39 on page 528.

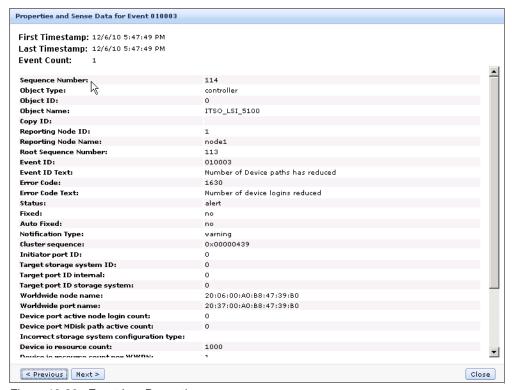


Figure 13-39 Event Log Properties

Note that the event log properties menu is accessible from both the recommended actions and the event log navigation tabs. Also, the Next button will skip to the next event log properties as displayed in Figure 13-40 on page 528.

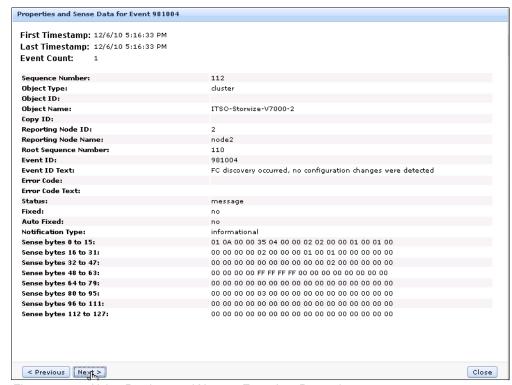


Figure 13-40 Using Previous and Next on Event Log Properties

At this point we reset the loose Fibre Channel cable. Once we return to the GUI, we use the Mark As Fixed drop-down button to indicate the problem is fixed, or we use the Run Fix Procedure as shown in Figure 13-41.

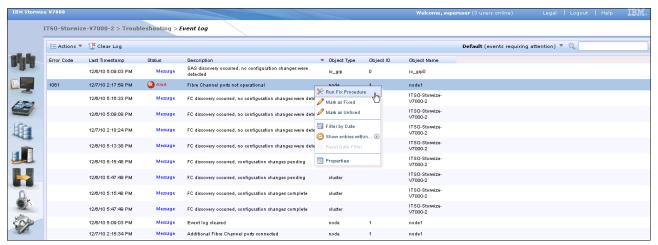


Figure 13-41 Run Fix Procedure

Since we clicked the Run Fix Procedure, the directed maintenance procedure (DMP) starts showing the current status of the ports as shown in Figure 13-42 on page 529.

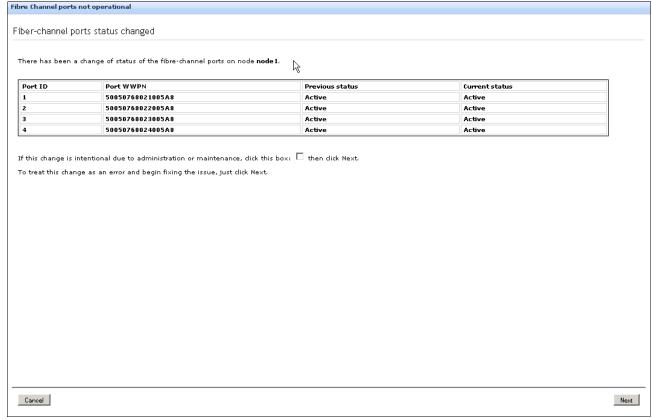


Figure 13-42 Run Fix Procedure, Mark FC Port fixed, details

Figure 13-43 shows the next screen in the procedure, giving you another look at the current status of the ports, and since we find the status is correct (all active), we click on the *Fibre channel status is correct, mark as fixed* button.

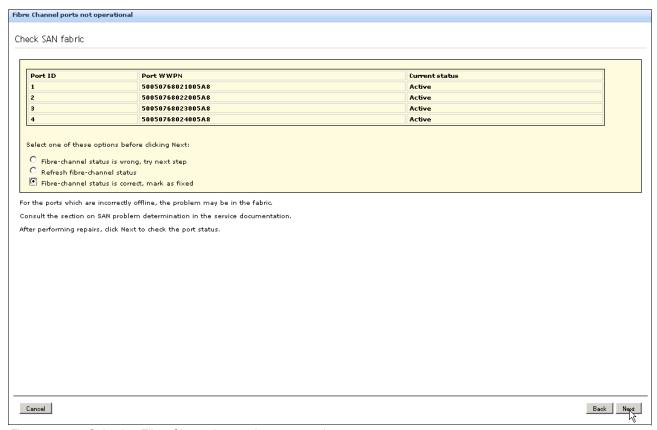


Figure 13-43 Selecting Fibre Channel status is correct option

As shown in Figure 13-44 on page 531 there is a confirmation panel before the error is marked as fixed where you can cancel the DMP procedure, and we click on **Next** button.

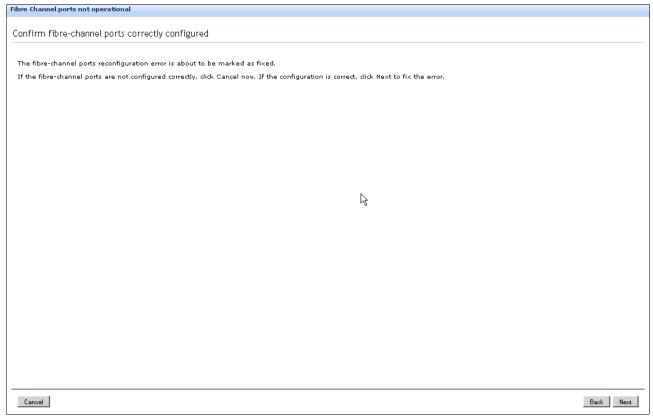


Figure 13-44 Error is about to be marked as fixed

At this point the error is marked as fixed as shown by Figure 13-45 on page 532, and the error will no longer be visible on the Recommended Actions grid.

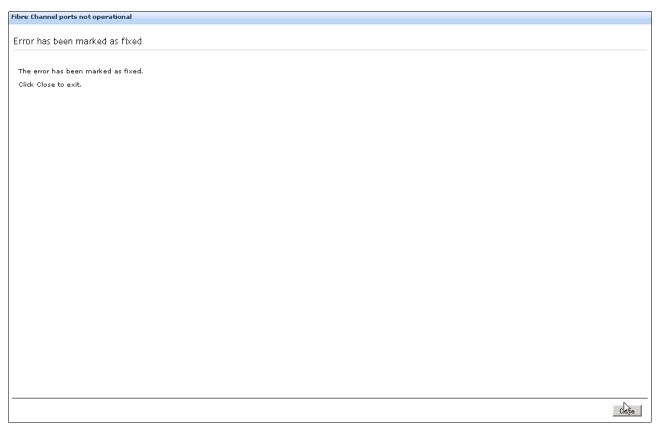


Figure 13-45 Confirmation panel the error is fixed.

It is possible to confirm the error is fixed by selecting the **Event Log** tab, using the **Show All** filter, and including the 'fixed' column as shown in Figure 13-46 on page 533.

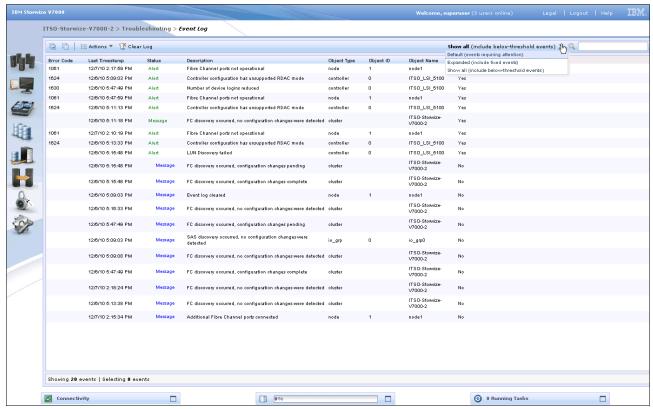


Figure 13-46 Using the Event Log to review a fixed error

13.8 Audit Log

As shown in Figure 13-47 on page 534 along the navigation panel at the bottom of the User Management navigation menu is the Audit Log tab. By clicking on this button we are able to view which configuration CLI commands have been executed so far on the IBM Storwize V7000 (information and service commands are not included).

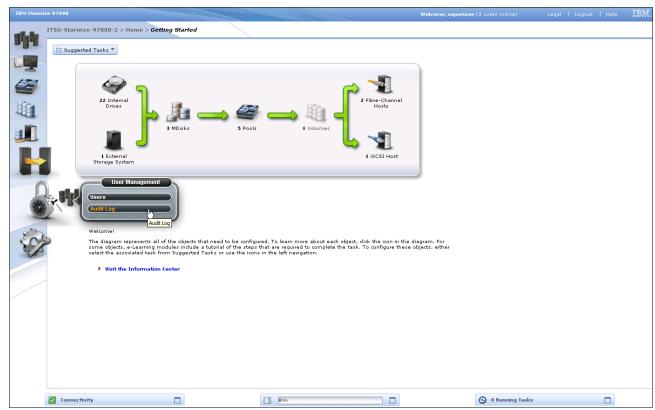


Figure 13-47 Audit log from user management pane

An example of the Audit Log viewed after creating volumes and mapping them to hosts is shown in Figure 13-48 with a command highlighted.

Notice also that the running tasks button is available at the bottom of the screen in the status pod, and if clicked shows the progress of currently executing tasks.

The Audit Log is especially useful in determining past configuration events when trying to determine how perhaps a volume ended up shared by two hosts, or may have been overwritten. The audit log is also included in the **svc_snap** support data in order to aid in problem determination.

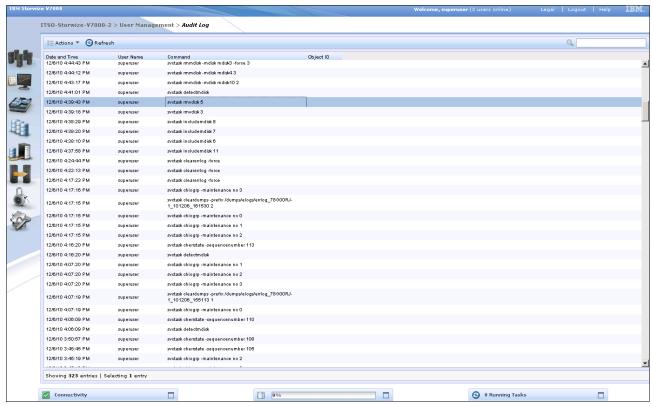


Figure 13-48 Audit Log

It is also possible to change the view of the Audit Log grid just like we were able to manipulate the event logs by right-clicking on the column headings as demonstrated in Figure 13-49.

The grid layout and sorting is completely under the users control so that one can show everything in the audit log, sort on different columns, or reset back to the default grid preferences.



Figure 13-49 Right click on audit log column headings

13.9 Collecting Support Information

Occasionally, if you get a problem and call the IBM Support Center, they may ask you to provide support data. This is available in the Support Tab of the Troubleshooting navigation panel.

Click on the **Support Tab** to begin the procedure of collecting support data as shown in Figure 13-50.

Launching the drop-down as shown in Figure 13-52, and assuming that the node restarted, we collect the default logs plus all the existing states are in order to capture the maximum data for support.

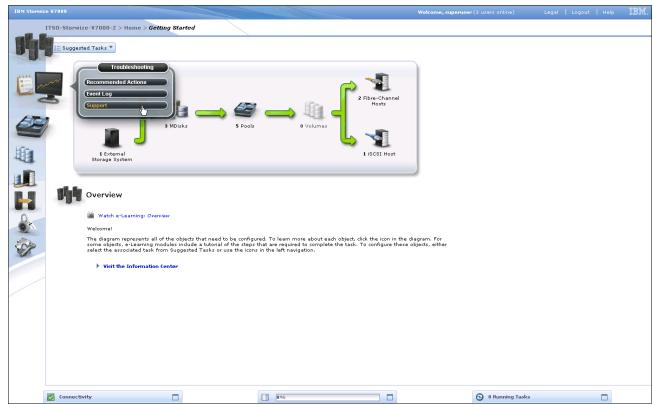


Figure 13-50 Support option

Then click on the Download Support Package button as shown in Figure 13-51.



Figure 13-51 Download Support Package panel

As shown in Figure 13-52 on page 537 this will launch the menus for collecting different versions of <code>svc_snap</code>, depending on the event that is to be investigated. For example, if you notice that a node was restarted in the event log, we recommend that you capture the snap with the latest existing states aves.



Figure 13-52 Download support package choices

The next panel creates the snap on the IBM Storwize V7000, including the latest states ave from each node canister. This process may take a few minutes as shown in Figure 13-53.

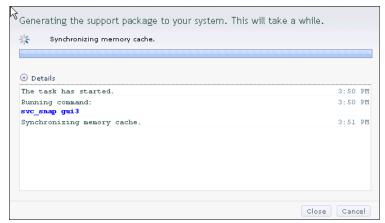


Figure 13-53 task detail panel

Figure 13-54 shows that the GUI gives you a choice to save the file on your local Windows system.

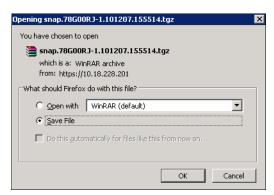


Figure 13-54 Saving file to local Windows system

As shown in Figure 13-55, save the resulting snap file in a directory.

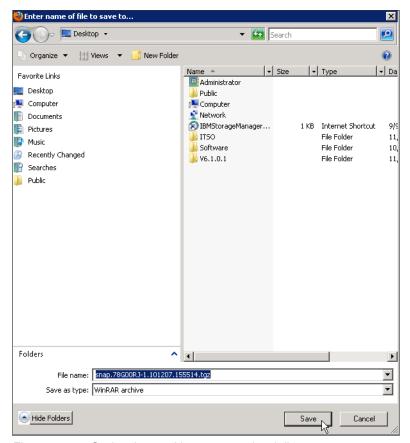


Figure 13-55 Saving the resulting snap to a local directory

As you will want to upload the resulting snap to the IBM Support portal once you open a call with IBM support, navigate to:

http://www.ecurep.ibm.com/app/upload

At this point you are ready to call the IBM Support Line or use the IBM Support Portal to open a call:

http://www-947.ibm.com/support/entry/portal/Open_service_request?brandind=Hardware

13.10 Shutting down the IBM Storwize V7000

We show how to shutdown the Storwize V7000 using both the GUI and the CLI.

Note: You should never shut down your IBM Storwize V7000 by powering off the PSUs, removing both PSUs, or removing both power cables from a running system.

13.10.1 Shutting down using the GUI

You can shutdown only one node canister or the entire cluster. In the case you shutdown only one node canister all the activities remain active, when you shutdown the entire cluster, you need to power on locally to restart the system.

This is the sequence of commands using the GUI:

1. Hover the cursor on the **HOME** function icon and click on **Manage Device** as shown in Figure 13-56 on page 539.

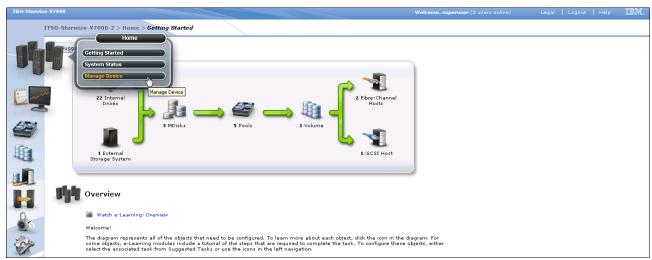


Figure 13-56 Manage Device option of the HOME function icon

2. From the **Action** button you can click on the **Shutdown Cluster** option as shown in Figure 13-57 on page 539.



Figure 13-57 Shutdown cluster option

3. The Confirm IBM Storwize V7000 Shutdown IBM Storwize V7000 window (Figure 13-58 on page 540) opens. You will get a message asking you to confirm whether you want to shut down the cluster.

Ensure that you have stopped all FlashCopy mappings, Remote Copy relationships, data migration operations, and forced deletions before continuing. Click Yes to begin the shutdown process.

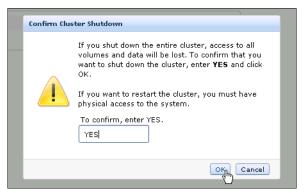


Figure 13-58 Confirm Shutdown panel

Tip: When you shut down the IBM Storwize V7000, it will not automatically start. You must manually start the IBM Storwize V7000.

13.10.2 Shutting down using the Command Line Interface

You can also shutdown the IBM Storwize V7000 via the CLI using the putty program. You must use this command:

SVCTASK STOP CLUSTER

13.11 Shutting down and powering on an IBM Storwize V7000 infrastructure

When you shutdown, or power on, the entire infrastructure (storage, servers, and applications) you must follow a particular sequence for both shutdown and then power on. Here we describe an example sequence of shutdown, and then power on, of an infrastructure that includes the Storwize V7000. Figure 13-59 on page 541 shows the location of power switches on Controller and Expansion enclosures:

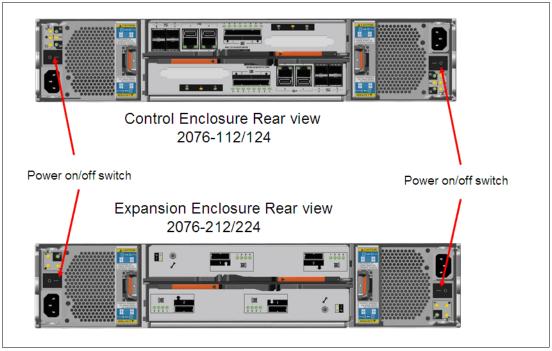


Figure 13-59 Location of power switches

Shutting down

- 1. Shutdown your servers and all applications
- 2. Shutdown your IBM Storwize V7000:
 - a. Shutdown the cluster via the GUI or Command Line Interface
 - b. Power off both switches of the Controller enclosure
 - c. Power off both switches of all Expansion enclosures
- 3. Shutdown your SAN switches

Powering on

- 1. Power on you SAN switches and wait until the boot has completed
- 2. Power on your storage systems and wait until the systems are up, then:
 - a. Power on both switches of all Expansion enclosures
 - b. Power on both switches of the Controller enclosure
- 3. Power on your servers and start your Applications





CLI setup and SAN Boot

This appendix describes setup of the CLI and provides additional information about SAN boot.

Command Line Interface

The IBM Storwize V7000 offers a very powerful CLI, which offers even more functions than the GUI. This is not intended to be a detailed CLI guide as this is not the scope of the 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 SVC, and in addition there are some more available to manage internal storage. If a task completes in the GUI, the CLI command is always displayed in the details as shown throughout the book. This should help to make it easy to use the CLI.

Detailed CLI information is available in the IBM Storwize V7000 Infocenter Command Line Section available at:

http://publib.boulder.ibm.com/infocenter/storwize/ic/index.jsp?topic=/com.ibm.storwize.v7000.doc/svc clicommandscontainer 229g0r.html

In *Implementing the IBM System Storage SAN Volume Controller V6.1*, SG24-7933, is also a lot of information about using the CLI and the commands in that book also apply to the Storwize V7000.

Basic Setup

In the Storwize V7000 GUI, authentication is done with a username and a password. The CLI uses a secure shell to connect from the host to the IBM Storwize V7000. Therefore a private and a public key pair has to be used. The following steps are required to enable CLI access:

- A public key and a private key are generated together as a pair
- A public key is uploaded to the IBM Storwize V7000 via the GUI.
- A client ssh tool has to 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 SVC cluster.

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 cluster, and a private key that is kept private to the workstation that is running the SSH client. These keys authorize specific users to access the administration and service functions on the cluster. 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 cluster. 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 SVC cluster(s).

The recommended ssh client is PuTTY, there is also a PuTTY key generator which could be used to generate the private and public key pair. The free download is available at:

http://www.chiark.greenend.org.uk

Download the following tools:

PuTTY ssh client: putty.exe

► PuTTY key generator: puttygen.exe

Generate public and private key pair

Start the PuTTY key generator to generate the public and private key pair as shown in Figure A-1 on page 545.



Figure A-1 PuTTY Key Generator

Make sure the following options are selected:

- SSH2 RSA
- Number of bits in a generated key: 1024

Click **Generate** and move the cursor on the blank area in order to generate the keys as shown in Figure A-2.

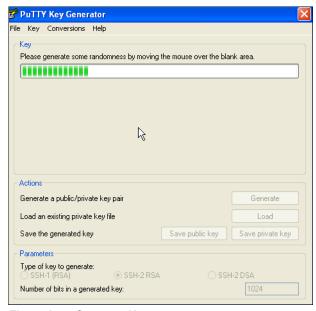


Figure A-2 Generate Keys

To generate keys: The blank area 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.

After the keys are generated save them for later use. Click **Save public key** as shown in Figure A-3.



Figure A-3 Save Public Key

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.

Note: By default the PuTTY Key Generator saves the public key with no extension. We recommend that you use the string "pub" in naming the public key, for example, "pubkey", to easily differentiate the SSH public key from the SSH private key.

Click **Save private key** as shown in Figure A-4 on page 547.

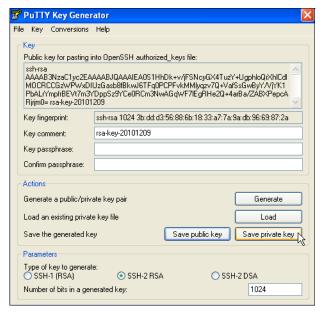


Figure A-4 Save Private Key

You are prompted with a warning message as shown in Figure A-5. Click **Yes** to save the private key without a passphrase.



Figure A-5 Confirm the Security Warning

When prompted enter a name (for example "icat"), select a secure place as location and click **Save**.

Note: The PuTTY Key Generator saves the private key with the PPK extension.

Close the PuTTY key generator.

Upload the SSH public key to the IBM Storwize V7000

After you have created your SSH key pair you need to upload your SSH public key onto the SVC cluster. Open the user section as shown in Figure A-6 on page 548.

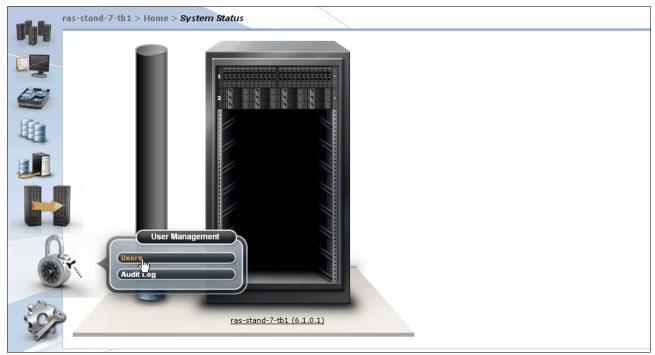


Figure A-6 Open User Section

Right click the user for which you want to upload the key and click **Properties** as shown in Figure A-7.



Figure A-7 Superuser Properties

To upload the public key click **Browse**, select your public key and click **OK** as shown in Figure A-8 on page 549.



Figure A-8 Select Public Key

Click **OK** and the key will be uploaded as shown in Figure A-9.

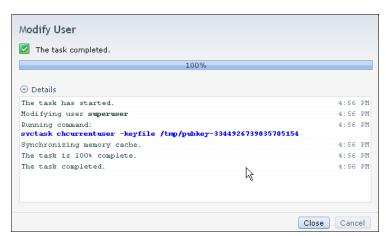


Figure A-9 Public Key Upload Complete

Click Close to return to the GUI.

Configure the SSH Client

Before the CLI can be used the SSH client must be configured. To do this start PuTTY as shown in Figure A-10 on page 550.

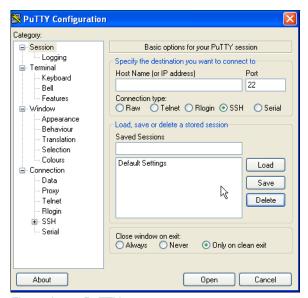


Figure A-10 PuTTY

In the right pane under the "Specify the destination you want to connect to" section, select SSH. Under the "Close window on exit" section, select Only on clean exit, which ensures that if there are any connection errors, they will be displayed on the user's window.

From the Category pane on the left side of the PuTTY Configuration window, click **Connection --> SSH** to display the PuTTY SSH Configuration window, as shown in Figure A-11.

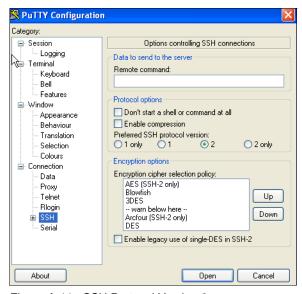


Figure A-11 SSH Protocol Version 2

In the right pane, in the "Preferred SSH protocol version" section, select 2.

From the Category pane on the left side of the PuTTY Configuration window, select **Connection --> SSH --> Auth**. As shown in Figure A-12 on page 551, in the right pane, in the "Private key file for authentication:" field under the Authentication Parameters section, either browse to or type the fully qualified directory path and file name of the SSH client private key file created earlier (for example, C:\Support Utils\PuTTY\icat.PPK).



Figure A-12 SSH Authentication

From the Category pane on the left side of the PuTTY Configuration window, click **Session** to return to the Session view as shown in Figure A-10 on page 550.

In the right pane, enter the hostname or cluster IP address of the IBM Storwize V7000 cluster in the Host Name field, and enter a session name in the Saved Sessions field as shown in Figure A-13.

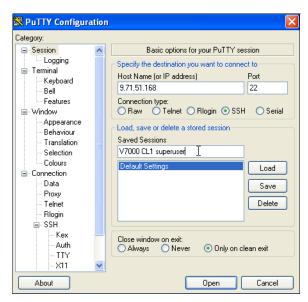


Figure A-13 Enter Session Information

Click Save to save the new session as shown in Figure A-14 on page 552.

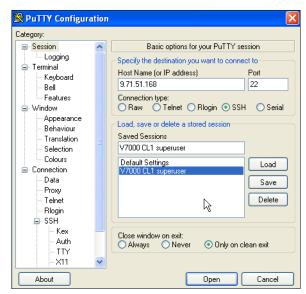


Figure A-14 Save Session

Highlight the new session and click **Open** to connect to the IBM Storwize V7000. A PuTTY Security Alert will appear, confirm it by clicking **Yes** as shown in Figure A-15



Figure A-15 Confirm Security Alert

PuTTY will now connect to the cluster and prompt you for a username. Enter **admin** as username and press **Enter** as shown in Example A-1.

Example: A-1 Enter User Name

login as: admin
Authenticating with public key "rsa-key-20101209"
IBM 2076:ras-stand-7-tb1:admin>

You have now completed the tasks that are required to configure the CLI for IBM Storwize V7000 administration.

Example Commands

A detailed description about all the available commands is beyond the intended scope of this book. In this section are sample commands listed that we have referenced in the book.

All commands that have the intention to generate any output start with **svcinfo**. All commands which will make system changes start with **svctask**. If you type **svcinfo** or **svctask** and press the "Tab-Key" twice all the available sub-commands are listed. Pressing

the "Tab Key" twice will also auto complete commands if the input is valid and unique to the system.

Enter **svcinfo 1svdisk** as shown in Example A-2 to list all configured volumes on the system. The example shows that three volumes are configured.

Example: A-2 List All Volumes

```
IBM 2076:ras-stand-7-tb1:admin>svcinfo lsvdisk
id name
                      IO group id IO group name status mdisk grp id
mdisk grp name capacity type FC id FC name RC id RC name vdisk UID
fc map count copy count fast write state se copy count
O RadHat EasyTier generic O
                               io grp0
                                           online 3
                                                           Multi
Tier Pool 20.00GB striped
0
                                              empty
1 RedHat EasyTier mirror 0
                             io grp0
                                        online many
                                                       many
                               25.00GB many
                    0
2 ESX_vmfs_1
                             io grp0
                                       online many
                                                       many
200.00GB many
                               empty
                      0
```

Enter **svcinfo 1shost** to get a list of all configured hosts on the system as shown in Example A-3.

Example: A-3 List Hosts

```
IBM_2076:ras-stand-7-tb1:admin>svcinfo lshost
id name port_count iogrp_count
0 ESX_1_RZA 2 4
1 ESX_2_RZB 2 4
```

To map the volume to the hosts use the **svctask mkvdiskhostmap** command as shown in Example A-4.

Example: A-4 Map volumes to host

```
IBM_2076:ras-stand-7-tb1:admin>svctask mkvdiskhostmap -host ESX_1_RZA -scsi 0 -force ESX_vmfs_1
Virtual Disk to Host map, id [0], successfully created
IBM_2076:ras-stand-7-tb1:admin>svctask mkvdiskhostmap -host ESX_2_RZB -scsi 0 -force ESX_vmfs_1
Virtual Disk to Host map, id [0], successfully created
IBM_2076:ras-stand-7-tb1:admin>
```

To verify the host mapping use the **svcinfo lsvdiskhostmap** command as shown in Example A-5.

Example: A-5 List all hosts mapped to a volume

In the CLI there are more options available than in the GUI. All advanced settings can be set, for example I/O throttling. To enable I/O throttling the properties of a volume can be changed using the **svctask changevdisk** command as shown in Example A-6. To verify the changes the **svcinfo lsvdisk** command is used.

Note: The **svcinfo 1svdisk** command lists all available properties of a volume and its copies, however to make it easier to read lines in the example output have been deleted.

Example: A-6 Enable Advanced Properties - for example I/O Throttling

If you do not specify the unit parameter the throttling will be based on I/Os instead of throughput as shown in Example A-7.

Example: A-7 Throttling based on I/O

To disable I/O Throttling set the I/O rate to 0 as shown in Example A-8.

Example: A-8 Disable I/O Throttling

```
IBM_2076:ras-stand-7-tb1:admin>svctask chvdisk -rate 0 ESX_vmfs_1
IBM_2076:ras-stand-7-tb1:admin>svcinfo lsvdisk ESX_vmfs_1
id 2
.
.vdisk_UID 60050768018E0173F0000000000000
throttling 0
preferred_node_id 2
.
```

IBM 2076:ras-stand-7-tb1:admin>

FlashCopy Example Commands

IBM storwize V7000 offers powerful FlashCopy functions as described in Chapter 11, "Copy Services" on page 381. FlashCopy creation is a process which is often scripted. The commands to create a FlashCopy are easy to find out if you create a Flash Copy in the GUI, and the commands for your system are listed if the creation completes.

Example A-9 shows you the required steps to prepare a Reverse FlashCopy, and shows you the FlashCopy command using the "Reverse" option. As you can see at the end of Example A-9 FCMAP_rev_1 shows a restoring value of **yes** while the FlashCopy mapping is copying. After it has finished copying, the restoring value field will change to **no**.

Example: A-9 Prepare Reverse Flash Copy

```
IBM 2076:ITSO-Storwize-V7000-1:admin>svcinfo lsvdisk
id name IO_group_id IO_group_name status mdisk_grp_id
mdisk grp name capacity type FC id FC name RC id RC name
vdisk UID fc map count copy count fast write state
se copy count
4 Volume FC S O io grpO online 2
STGPool DS4700 1 1.00GB striped
60050768018281BEE00000000000000 0 1 empty 0
5 Volume FC T S1 0 io grp0 online 2
STGPool DS4700 1 1.00GB striped
60050768018281BEE00000000000000 0 1 empty 0
6 Volume FC T1 0 io grp0 online 2
STGPool DS4700 1 1.00GB striped
60050768018281BEE00000000000000 0 1 empty 0
IBM 2076:ITSO-Storwize-V7000-1:admin>svctask mkfcmap -source Volume FC S -target
Volume FC T S1 -name FCMAP 1 -copyrate 50
FlashCopy Mapping, id [0], successfully created
IBM 2076:ITSO-Storwize-V7000-1:admin>svctask mkfcmap -source Volume FC T S1
-target
Volume FC S -name FCMAP rev 1 -copyrate 50
FlashCopy Mapping, id [1], successfully created
IBM 2076:ITSO-Storwize-V7000-1:admin>svctask mkfcmap -source Volume FC T S1
-target
Volume FC T1 -name FCMAP 2 -copyrate 50
FlashCopy Mapping, id [2], successfully created
IBM 2076:ITSO-Storwize-V7000-1:admin>svcinfo lsfcmap
id name source vdisk id source vdisk name target vdisk id target vdisk name
group id group name status progress copy rate clean progress incremental
partner_FC_id partner_FC_name restoring
O FCMAP 1 4 Volume FC S 5 Volume FC T S1
idle or copied 0 50 100 off 1 FCMAP rev 1 no
1 FCMAP rev 1 5 Volume FC T S1 4 Volume FC S
idle or copied 0 50 100 off
O FCMAP 1 no 2 FCMAP 2 5 Volume FC T S1 6 Volume FC T1
idle or copied 0 50 100 off no
```

```
IBM 2076:ITSO-Storwize-V7000-1:admin>svctask startfcmap -prep FCMAP 1
IBM 2076:ITSO-Storwize-V7000-1:admin>svcinfo lsfcmap
id name source vdisk id source vdisk name target vdisk id target vdisk name
group id group name status progress copy rate clean progress incremental
partner FC id partner FC name restoring
O FCMAP_1 4 Volume_FC_S 5 Volume_FC_T_S1
copying 0 50 100 off 1 FCMAP rev 1 no
1 FCMAP rev 1 5 Volume FC T S1 4 Volume FC S
idle or copied 0 50 100 off 0 FCMAP 1 no
2 FCMAP_2 5 Volume_FC_T_S1 6 Volume_FC_T1
idle or copied 0 50 100 off no
IBM 2076:ITSO-Storwize-V7000-1:admin>svctask startfcmap -prep FCMAP 2
IBM 2076:ITSO-Storwize-V7000-1:admin>svcinfo lsfcmap
id name source vdisk id source vdisk name target vdisk id target vdisk name
group id group name status progress copy rate clean progress incremental
partner FC id partner FC name restoring
O FCMAP_1 4 Volume_FC_S 5 Volume_FC_T_S1 copying 8 50 91 off 1 FCMAP_rev_1 no
1 FCMAP rev 1 5 Volume FC T S1 4 Volume FC S idle or copied 0 50 100 off 0 FCMAP 1
2 FCMAP 2 5 Volume FC T S1 6 Volume FC T1 copying 0 50 100 off no
IBM 2076:ITSO-Storwize-V7000-1:admin>svctask startfcmap -prep FCMAP rev 1
CMMVC6298E The command failed because a target VDisk has dependent FlashCopy
mappings.
IBM 2076:ITSO-Storwize-V7000-1:admin>svctask startfcmap -prep -restore FCMAP rev 1
IBM 2076:ITSO-Storwize-V7000-1:admin>admin>svcinfo lsfcmap
id name source_vdisk_id source_vdisk_name target_vdisk_id target_vdisk_name
group id group name status progress copy rate clean progress incremental
partner FC id partner FC name restoring
O FCMAP 1 4 Volume FC S 5 Volume FC T S1
copying 20 50 86 off 1 FCMAP rev 1 no
1 FCMAP rev 1 5 Volume FC T S1 4 Volume FC S
copying 86 50 16 off 0 FCMAP 1 yes
2 FCMAP 2 5 Volume FC T S1 6 Volume FC T1
copying 12 50 100 off no
```

Modify Global Mirror Delay for Testing

As described in 11.2, "Remote Copy" on page 423 a delay can be simulated on Global Mirror Links. This is an advanced feature only available in the CLI. Example A-10 shows you how to add a delay to the Global Mirror connection.

Example: A-10 Add Global Mirror Delay

```
total_allocated_extent_capacity 2.10TB

IBM_2076:ITSO-Storwize-V7000-1:admin>svctask chcluster -gmintradelaysimulation 20
IBM_2076:ITSO-Storwize-V7000-1:admin>svcinfo lscluster 0000020060E14FFC
id 0000020060E14FFC
name ITSO-Storwize-V7000-1
gm_inter_cluster_delay_simulation 0
gm_intra_cluster_delay_simulation 20
total_allocated_extent_capacity 2.10TB

IBM_2076:ITSO-Storwize-V7000-1:admin>svctask chcluster -gminterdelaysimulation 20
IBM_2076:ITSO-Storwize-V7000-1:admin>svcinfo lscluster 0000020060E14FFC
id 0000020060E14FFC
name ITSO-Storwize-V7000-1
.
gm_inter_cluster_delay_simulation 20
gm_intra_cluster_delay_simulation 20
.
total_allocated_extent_capacity 2.10TB
```

SAN Boot

IBM Storwize V7000 supports SAN Boot for Windows, VMware and many other operating systems. SAN Boot support can change from time to time, so we recommend regularly checking the IBM Storwize V7000 interoperability matrix:

http://www-03.ibm.com/systems/storage/disk/storwize v7000/interop.html

Also in the IBM Storwize V7000 Information Center a lot of information about SAN Boot in combination with different operating systems is provided:

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

Additional information about SAN Boot is also covered in the IBM Multipath Subsystem Device Driver User's Guide available at:

http://www-01.ibm.com/support/docview.wss?rs=503&context=HW26L&uid=ssg1S7000303

Enable SAN Boot for Windows

Perform the following procedure if you want to install a Windows host using SAN boot:

- 1. Configure the Storwize V7000 so that only the boot volume is mapped to the host.
- 2. Configure the Fibre Channel SAN so that the host only sees one 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.

Note: It might be required to load an additional HBA device driver during installation, depending on your Windows version and the HBA type.

- 5. Don't forget to 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.

Enable SAN Boot for VMware

Perform the following procedure if you want to install a VMware ESXhost using SAN boot:

- 1. Configure the Storwize V7000 so that only the boot volume is mapped to the host.
- 2. Configure the Fibre Channel SAN so that the host only sees one 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.

Note: It might be required 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 Windows 2000 Server, Windows Server 2003, or Windows Server 2008 operating system, and existing SAN boot images that are controlled by storage controllers, you can migrate these images to image-mode volumes that are controlled by the Storwize V7000 system.

Note: For SAN Boot Procedures for other operating systems check the IBM Storwize V7000 Information Center available at:

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

Perform the following steps to migrate your existing SAN boot images:

- If the existing SAN boot images are controlled by an IBM storage controller that uses SDD as the multipathing driver, you must use SDD v1.6 or higher. Run the SDD command datapath set bootdiskmigrate 2076 to prepare the host for image migration. See the Multipath Subsystem Device Driver (SDD) documentation for more information.
- 2. Shut down the host.
- 3. Perform 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,....)
 - 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 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. Perform the following configuration changes on the 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 performing 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 v 1.6 or higher is installed and you ran the **bootdiskmigrate** command in step 1, reboot your host, update SDDDSM to the latest level and go to step 14. If SDD v.1.6 is not installed go to the next step.
- 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 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 booting 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 publications

Note that several of the documents referenced here might be available in softcopy only.

- ▶ Implementing the IBM System Storage SAN Volume Controller V6.1, SG24-7933
- ► Get More Out of Your SAN with IBM Tivoli Storage Manager, SG24-6687
- ▶ IBM Tivoli Storage Area Network Manager: A Practical Introduction, SG24-6848
- IBM System Storage: Implementing an IBM SAN, SG24-6116
- Introduction to Storage Area Networks, SG24-5470
- ► SAN Volume Controller V4.3.0 Advanced Copy Services, SG24-7574
- SAN Volume Controller: Best Practices and Performance Guidelines, SG24-7521
- Using the SVC for Business Continuity, SG24-7371
- ▶ IBM System Storage Business Continuity: Part 1 Planning Guide, SG24-6547
- IBM System Storage Business Continuity: Part 2 Solutions Guide, SG24-6548

Other publications

These publications are also relevant as further information sources:

- ► IBM System Storage Open Software Family SAN Volume Controller: Planning Guide, GA22-1052
- ► IBM System Storage Master Console: Installation and User's Guide, GC30-4090
- Subsystem Device Driver User's Guide for the IBM TotalStorage Enterprise Storage Server and the IBM System Storage SAN Volume Controller, SC26-7540
- ► IBM System Storage Open Software Family SAN Volume Controller: Installation Guide, SC26-7541
- ► IBM System Storage Open Software Family SAN Volume Controller: Service Guide, SC26-7542
- ► IBM System Storage Open Software Family SAN Volume Controller: Configuration Guide, SC26-7543
- ► 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: CIM Agent Developers Reference, SC26-7545
- IBM TotalStorage Multipath Subsystem Device Driver User's Guide, SC30-4096
- ► IBM System Storage Open Software Family SAN Volume Controller: Host Attachment Guide, SC26-7563

- ► IBM System Storage SAN Volume Controller Model 2145-CF8 Hardware Installation Guide, GC52-1356
- ► IBM System Storage SAN Volume Controller Model 2145-8A4 Hardware Installation Guide, GC27-2219
- ► IBM System Storage SAN Volume Controller Model 2145-8G4 Hardware Installation Guide, GC27-2220
- ► IBM System Storage SAN Volume Controller Models 2145-8F2 and 2145-8F4 Hardware Installation Guide, GC27-2221
- ► IBM System Storage SAN Volume Controller V5.1.0 Host Attachment Guide, SG26-7905-05
- ► Command Line Interface User's Guide, SG26-7903-05

Online resources

These Web sites are also relevant as further information sources:

► IBM TotalStorage home page:

```
http://www.storage.ibm.com
```

► SAN Volume Controller supported platform:

http://www-1.ibm.com/servers/storage/support/software/sanvc/index.html

► Download site for Windows Secure Shell (SSH) freeware:

http://www.chiark.greenend.org.uk/~sgtatham/putty

▶ IBM site to download SSH for AIX:

http://oss.software.ibm.com/developerworks/projects/openssh

▶ Open source site for SSH for Windows and Mac:

http://www.openssh.com/windows.html

Cygwin Linux-like environment for Windows:

http://www.cygwin.com

► IBM Tivoli Storage Area Network Manager site:

http://www-306.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageAreaNetworkManager.html

► Microsoft Knowledge Base Article 131658:

http://support.microsoft.com/support/kb/articles/Q131/6/58.asp

Microsoft Knowledge Base Article 149927:

http://support.microsoft.com/support/kb/articles/Q149/9/27.asp

Sysinternals home page:

http://www.sysinternals.com

Subsystem Device Driver download site:

http://www-1.ibm.com/servers/storage/support/software/sdd/index.html

► IBM TotalStorage Virtualization home page:

http://www-1.ibm.com/servers/storage/software/virtualization/index.html

SVC support page:

 $\label{lem:http://www-947.ibm.com/systems/support/supportsite.wss/selectproduct?taskind=4\& brandind=5000033\&familyind=5329743\&typeind=0\&modelind=0\&osind=0\&psid=sr\&continue.x=1$

► SVC online documentation:

http://publib.boulder.ibm.com/infocenter/svcic/v3r1m0/index.jsp

► IBM Redbooks publications about SVC:

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

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

161-162, 165, 167, 222, 230, 234-235, 237-238, 241, Α 245, 261, 266, 269, 301, 315–316, 331, 340, 350, 356, active-active architecture 16 378-379, 385, 391, 431, 459, 553, 555 adding 24-25, 44, 377, 390 Cascaded FlashCopy 8, 384 administrative rights CD layout 463 DB2 user 470 **CDB 26** Advanced Copy Services 10, 27 chain 5, 37, 227-228, 423, 502 Advanced Settings 149, 152, 399-401 **CHAP 288** alert 20, 26, 35, 54, 421 chat session 30 algorithm 18, 233-234 CIMOM 17, 462 alias 27 CLI 17, 19, 27, 44-45, 130, 315, 372, 390, 428, 462, 464, aliases 27 493, 506, 533, 538, 540, 543-544, 549, 552, 554, 556 analysis xvi, 358, 377-378, 389, 459, 498 clone 5, 110, 385, 395, 397, 400-401, 404 architecture 1, 16, 26, 498 Cluster 16, 20, 49, 488, 498, 500, 506, 539 Array 17–18, 235, 241, 259, 261, 378 cluster 5, 16-21, 23, 27, 29, 31, 36, 39, 41-42, 44, 49, Array managed disks 5 78-79, 88, 119, 163, 176, 192, 204, 247, 263, 306, 390, asynchronous remote 9, 28, 427 423, 425-426, 428, 430-431, 433, 441, 455, 498, 500, Asynchronous remote copy 27 506, 508, 538-539, 541, 544, 547, 551-552 asynchronous remote copy 28 cluster partnership 425, 506 asynchronous replication 440 clusters 17, 20, 30, 115, 122, 204-205, 425-426, 429, audit log 119, 534-535 431, 433-435 Authentication 480, 550-551 Colliding writes 431 authentication 27, 134, 154, 479, 544, 550 Command Descriptor Block 26 Autoexpand 24 command-line interface 31 autoexpand 24 commands 9, 26, 44, 99, 101, 104, 108, 112, 119, 386, autoexpansion 24 388, 428, 466, 476–477, 506, 533, 539, 544, 552, 555 automatically discover 263 concepts 26, 498 automation 28 configurable warning capacity 24 autorun enabled 466 configuration 3, 8, 13, 16-17, 28-29, 31, 35-36, 38, auxiliary 29, 426-429, 431, 439, 441-442, 447, 455 41-47, 49, 52, 54-57, 59-61, 65, 75, 78, 121, 128, 133, auxiliary VDisk 427 136, 145, 147, 150, 153–155, 157, 170, 176, 178, 187, 208, 220, 225–226, 228–231, 233–234, 236, 238–239, 263, 273, 342, 356, 361, 363, 382, 398-399, 423, 429, 431, 436, 459-463, 490, 495, 505-507, 515, 533-534, background copy 8, 115, 384-386, 391, 415-416, 432, 544, 558-559 435 configuration limits 29, 431 background copy bandwidth 435 configuration node 16-17, 41-42, 49, 506 background copy rate 384-385, 391 connected xix, 15-16, 27, 36-37, 39, 41, 49, 127, 148, backup 5-8, 16-17, 28, 31, 36, 41, 110, 326, 381-382, 172-173, 176-177, 180, 184, 190, 205, 342, 431, 500, 385, 387, 389-391, 395, 401-402, 404, 426, 461, 505-508 connectivity 6, 40, 72, 153, 459, 462, 502, 515 bandwidth 28, 37, 115, 428, 432, 434-436 consistency 9, 29, 114, 247, 383, 386-388, 390, 417, battery 6-7, 12, 14, 503-505 419, 421-422, 428-431, 439, 452-453, 458 bitmap 383, 387 Consistency Group 9, 112-113, 388, 417-422, 428, boot 43, 130, 132, 136-137, 204, 541, 543, 557-559 450-457 consistency group 9, 29, 114, 384, 386-388, 390, 393, C 402, 407-408, 417-422, 428-431, 439, 450-458 cache 7, 12, 16, 21, 58, 60, 316, 354, 382, 386, 426-427, Consistency Groups 9, 386, 391, 417-418, 428 consistency groups 29, 386-388, 428-429, 431 caching 247, 385-386 consistent 9, 326, 382, 387-388, 426, 428, 430, 432, 459 Call Home email 30-31 consistent data set 382 Call Home feature 31 ConsistentDisconnected 430 canisters 4-6, 12-13, 16, 36, 41, 49, 127, 155, 358, 423, ConsistentStopped 430 498, 501-503 ConsistentSynchronized 430 container 266 capacity 4, 6-9, 15, 19-20, 24, 29, 43, 75, 77, 91,

contingency 24 drive fault 17 contingency capacity 24 drive faults 17 control enclosure 4-5, 12-14, 36-37, 498, 502-503 drives 3, 5-7, 9, 12-13, 15, 17-19, 37, 43-44, 59, 91, cooling 36 220-224, 226-229, 233-234, 236, 238, 241, 259, 261, Copied 385, 388 266, 272, 353-354, 356, 361-366, 377-378, 382, copy bandwidth 435 501-502, 559 copy operation 28, 389, 391 **DSMs** 129 dump 12, 498 copy process 115, 386, 430, 432, 435 copy rate 115, 384-385, 391, 415-416, 432 Copy Services 10, 27, 29, 65, 109–110, 391–393, 417, Ε 432-433 Easy Tier 9, 15, 21, 25-26, 43, 199, 227-228, 233, 235, Copy services functions 27, 381 269, 344, 353–361, 363–364, 366–367, 369–377, 460 copy services layer 27 empty 269-270, 419, 429, 451-452, 553, 555 Copying 385, 388, 409, 413 Enclosure 10, 58, 504 copying state 432 enclosures 3-7, 9, 13-15, 36-37, 43, 75, 77, 501-503, create a FlashCopy 555 540-541 creating 3, 17, 148, 150, 152, 157, 161, 300, 382, 390, entire VDisk 382 408, 423, 534 error 31, 327, 386, 388, 430, 434, 479, 498, 500-501, creation 19, 150, 162-163, 228, 270, 274, 300, 315, 325, 506, 525-527, 530-533 331, 371, 391, 418, 423, 431, 440, 444, 452, 555 error notification 498 Ethernet 13, 26, 36, 41-42, 49, 154, 499, 508 D Ethernet port 13, 36, 41, 49 data 5-6, 8-9, 11, 16-18, 23-28, 30-31, 157, 160, 170, EUIs 15 174, 199-200, 203-204, 208-209, 214-216, 222, 227, event 5, 12, 16-17, 30-31, 35, 38, 42, 54-56, 80-81, 243-248, 250, 253-254, 257, 260-261, 271, 309, 313, 83-84, 137, 259, 327, 461, 498, 516-521, 527-528, 321-322, 326-327, 340, 343-345, 349, 351, 353-354, 535-536 356, 358-361, 370, 377, 379, 381-391, 395, 413, 426, event log 518-520, 527-528, 536 428, 431-432, 442-443, 454, 460-462, 490, 495, 498, Event notifications 30 502, 505-506, 518, 534-536, 539 events 3, 30, 80, 518-520, 534 Data Migration 9, 200, 356-357 Excluded 263, 425 data migration 9, 200, 214-215, 245, 253-254, 257, 349, expand xix, 24, 44, 65, 290, 301, 308, 323-324 359 expansion canister 5, 14, 501 expansion canisters 4-6, 13 data mining 382 expansion enclosure 5, 13-14, 36-37, 503-504 database attention required 473 expansion enclosures 4, 6-7, 14, 36-37, 501 date code 26 extended unique identifiers 15 DB2 Administration Server 470 Extent 19-20 DB2 install 466 extent 19-20, 22, 25, 309, 329, 353-355, 357-361, 382 DB2 installation extent level 354 verify installation 476 extent size 19 DB2 tools catalog 473 extent sizes 20, 25 DB2 user account 470 extents 18-19, 21-25, 161, 241, 354-356, 358-361, db2sampl command 476 369, 377-379 decoupled 22 external storage 4-5, 7, 9, 11, 18, 38-39, 43, 53, 200, default location 372 204, 206–207, 214–215, 219–220, 241, 247, 253, 263, delete 86-87, 93, 106, 113, 115, 117, 256, 260, 271, 309, 339-340, 342-345, 347-351, 353, 356 284-285, 289, 295-296, 313, 329-330, 386, 391, External Virtualization 10-11, 340-341 413-414, 438, 450, 456-458 deleting 25, 274 F dependent writes 386-387 fabric 5, 13, 15, 26, 36, 39, 43, 205, 263, 342-343, 345, destaged 316, 325, 330 356, 425, 431, 459, 462 detected 30, 54, 58, 263, 498 failover 16, 28, 41, 127, 428 disconnected 436

fabric 5, 13, 15, 26, 36, 39, 43, 205, 263, 342–343, 345, 356, 425, 431, 459, 462 failover 16, 28, 41, 127, 428 failover situation 426 failure 5, 17–19, 227, 343, 432, 464, 498, 502, 504–505 FAStT xvi fault tolerant 13 Fibre Channel attached 7 Fibre Channel Port Login 27 Fibre Channel ports 7, 13, 36, 205, 345, 431

discovering 170

DNS suffix 465

disk controller 205, 247

distance 9, 28, 423, 427-428

domain name system 465

documentation 28, 32, 326, 558

Disk Drive Types 14

file system 351, 459, 461 FlashCopy 3, 8-9, 24, 27-29, 44, 110-113, 161, 313, 315, 381–399, 405, 408–415, 417–420, 422, 426, 430, 432, 539, 555-556 FlashCopy mapping 9, 384-385, 387-388, 390-391, 398-399, 402, 405-415, 417, 419-422, 555 FlashCopy mapping states Idling/Copied 386 FlashCopy mappings 9, 29, 113, 383-384, 386, 388, 390 FlashCopy Nocopy 8 FlashCopy properties 390 flexibility 231 foreground I/O latency 435 FRU 224 Full Feature Phase 27 Full/Incremental Copy 8 fully allocated 8, 24-25, 160-161

G

generator 544–545, 547 Global Mirror 9–10, 28–29, 44, 423, 427–429, 431, 440, 459–460, 556 Global Mirror relationship 432 gminterdelaysimulation 428, 557 gmintradelaysimulation 428, 557 goals 17, 228, 235, 502 grain 162, 316, 383–384 granularity 382 GUI 3, 20, 27, 31–32, 43–45, 124, 127, 132, 138, 144, 158–159, 170, 199, 274, 345, 361, 372, 375, 381–382, 390–391, 423, 432, 461–464, 486, 493, 505–509, 518–519, 529, 537–539, 541, 544, 546, 549, 554–555 GUID 465

Н

hardware failures 31 HBA 42-43, 128-129, 136, 557-559 HBA ports 559 help xv, xviii-xix, 8, 44, 65, 220, 239, 266, 268, 345, 347, 382, 459, 502, 544 high availability 163, 502 host 2, 5, 7–9, 15–16, 18–19, 21–22, 24–26, 36, 38–39, 41, 43, 77, 96, 100, 103–105, 107–108, 122, 127–133, 135-139, 144-152, 156-158, 160-161, 163, 168-178, 182-184, 189-190, 200, 206, 208-213, 220, 229, 271, 273-276, 278, 280, 282-292, 294-297, 299, 301, 303-308, 313, 316, 318, 321-322, 325-327, 331, 334-336, 343, 351, 360, 379, 382-383, 426-427, 448, 459, 462-463, 465, 515-516, 544, 553, 557-559 Host Type 288 hot spare 17, 222-223, 233, 363-364, 501-502

Ī

I/O Group 16, 297, 431 I/O group 13, 15–16, 29, 192, 297, 430–431, 502, 559 I/O statistics 354, 360, 375 IBM Assist On-site tool 30 IBM Storwize V7000 xv, 1, 3-12, 14-18, 20-21, 24-28, 30-33, 35-36, 38-46, 48-49, 54, 61, 63-65, 69, 74-75, 77-78, 80, 92, 94, 115, 127-128, 130, 132, 134-136, 138–139, 144, 147, 153–155, 157–160, 163, 170, 172, 174, 176, 179, 181–182, 186, 192, 199–200, 203, 205, 207-209, 212, 214-216, 233, 254, 257, 273-274, 282, 292, 300, 307, 309, 339-340, 342-345, 349, 351, 353-356, 358-362, 364, 372, 377, 379, 381-383, 385-386, 390-391, 410, 423-426, 428, 430-436, 460, 463, 486, 488, 492-493, 497-498, 500-503, 505-506, 508-510, 515, 517-518, 520-521, 533, 537-541, 544, 547, 551-552, 557-559 IBM Support Center 30, 535 IBM Systems Storage Productivity Center 30 identical data 426 IdlingDisconnected 430 Image 18, 241, 246, 311, 313 Image Mode 23, 208, 215, 311-312 Image mode 23, 246 image mode 23-24, 92, 246-248, 252-254, 256, 271, 311-312, 344, 351, 360, 559 including xv, 31, 52, 65, 200, 204, 222, 224, 229, 241, 347, 349–350, 381, 391, 400, 432, 459–460, 463, 506, 518, 532, 537 inconsistent 413 InconsistentCopying 430 InconsistentDisconnected InconsistentStopped 430 information xv-xvi, 13, 16, 25, 28, 30-31, 43, 55, 65, 75, 78-79, 88, 111, 121-122, 128, 131-132, 134, 139, 165, 167, 173, 205-206, 208, 217, 220, 222, 224, 239, 241, 247, 253, 263–266, 268–269, 287, 289, 301, 315–316, 330-331, 340, 344-345, 347, 349, 351, 372, 379, 390, 418, 428, 435, 452, 460–462, 470, 506, 510, 516, 533, 543-544, 557-558 initiator 26-27, 132-134, 139, 143-144, 151, 155, 176, 191-192, 294, 502 initiator name 26, 144 install 36, 128, 130, 132, 136, 138, 177, 459, 462-463, 465-466, 470, 477, 479, 557-558 integrity 386-387 intercluster link bandwidth 435 internal drive objects 17 internal drives 5, 7, 15, 18, 91, 220-222, 226 internal storage 3, 5, 11, 214, 219-220, 222, 225-226, 239, 322, 340, 344-345, 351, 544 IP address 17, 26, 31, 36, 41-44, 49, 52, 142, 176, 192, 508, 551 IP addresses 17, 26, 36, 42, 49, 176, 192 IP network 26, 31 IQN 26-27 IQNs 15, 27 iSCSI 6-7, 15, 26-27, 36, 41, 105, 121, 127-128, 132-135, 139-144, 146, 150-155, 170, 175-179, 181-182, 189-192, 194, 197-198, 274, 288, 290-291, 294, 335 iSCSI Address 26 iSCSI Name 26 iSCSI node 26 iSCSI Qualified Name 26

ISL hop count 431 iSNS 154

J
JBOD 17

Κ

key 45–46, 48–49, 70–72, 118, 129, 134–135, 339, 349, 389, 461, 488, 544–550, 552

L

lane 5 latency 25, 28, 354-355, 432, 435 LBA 23, 327 learning videos 32 LED 13-14, 49, 224, 499-501, 504 LEDs 13-14, 497, 499-501, 503-505 license 8, 11, 43, 52-54, 200, 340-342, 460, 468, 478, licenses 10-11, 53, 340, 463 Licensing 10, 341 limits 8, 29, 431 link errors 263 Linux 128, 464 listing 85, 88, 370-371, 507 local cluster 425 log messages 31 logged 24, 64, 292 logical block address 23, 327 logical disk 21, 157 Login Phase 27 login phase 27 logins 288 LU 5 LUNs 129, 138, 200, 207-208, 241, 263, 343-345

М

maintenance contract 31 Managed 5, 18, 241, 496 managed 5, 8-9, 16, 19, 24, 30, 41, 45, 94, 219, 242, 250-251, 266, 271, 340, 360, 384, 387, 417, 428, 450, 460, 502 managed disk 5, 18-19, 309, 502, 559 management xvi, 3, 5, 16-17, 19, 28, 30-32, 36, 41-42, 44-46, 49, 52, 64-66, 70, 74, 163, 173, 199, 219, 266, 269, 349, 382, 391, 423, 431–432, 459–460, 462, 505, 508-509, 534, 544 Management Information Base 31 managing 30-31, 44, 220, 344-345, 391, 432, 459-460 mapped to this host 335 mapping 5, 8-9, 18, 23, 127, 168, 170, 192, 206, 208, 211-213, 241, 274, 276, 279-280, 282, 292, 298, 305-306, 331, 344, 351, 383-388, 390-392, 461, 534, 553, 555 Master xv, 442, 454 master 29-30, 426-429, 431, 439, 441-442, 447, 455 master console 30

241-244, 246, 250, 252-253, 256, 258-266, 311, 314, 319-320, 343-345, 347, 350-351, 355-356, 358, 366, 369, 501-502, 559 MDisks 5, 7, 18–22, 24, 92, 94–95, 111, 159, 200, 207-208, 211, 214-215, 220, 222, 229, 238-248, 253-254, 258, 260-261, 263-264, 266, 269-271, 319, 331, 344-345, 350-351, 355-356, 360, 493, 502 metadata 24, 161, 386 Metro Mirror 9, 28-29, 360, 423, 426-429, 431-432, 440, 459 Metro Mirror consistency group 428 Metro Mirror relationship 428 MIB 31 microcode xv, 344, 498 Microsoft Multi Path Input Output 129 migrate 9, 11, 23–25, 96, 111, 199–200, 203–205, 207, 246, 248, 253-254, 309, 311, 325-326, 329-330, 340, 345, 349, 354, 360, 558 migrating 25, 200, 208, 211, 382 migration 9, 11, 15, 20–21, 25, 30, 59, 157, 199–200, 202-217, 245-246, 248-250, 252-257, 300, 309-312, 329-330, 340, 349, 351, 353-356, 358-361, 382, 460, 539, 558 migration process 256, 309-310 migration progress 215, 245, 248, 254, 256-257 mirrored 9, 17-18, 24-25, 111, 163, 165, 167, 227, 320, 327, 360, 378, 426-427, 432, 502 mirrored pair 18 mirrored volume 25 mirroring 2, 9, 17, 20, 25, 36, 115, 163, 165, 214, 360, 502 MLC 15 modes 18, 28, 138, 241, 359 modifying 428 monitoring 16, 123, 355, 360, 459-460, 495, 498 MPIO 129-131, 177-178, 181 multipath I/O 130 multipath storage solution 130 multiple paths 43, 558 multiple source 28 multiple target 28, 384 Multi-target FlashCopy 8 multi-tiered storage pool 21

N

naming authority string 26
Nearline SAS 14–15, 228
NETBIOS 465
Node 4, 41, 87, 431, 498
node 5–6, 12–13, 16–17, 26–27, 36, 39, 41–42, 44, 49, 51, 87, 163, 192, 288, 343, 345, 423, 431, 498–500, 502–503, 506, 514, 527, 536–538, 557–559
node canister 5, 13, 17, 36, 41–42, 431
nodes 5, 13, 15–16, 19, 24–26, 49, 122, 385–386, 498, 502, 506, 514, 559
non-mirrored volume 25
non-zero contingency 24
NTP 43

MDisk 5, 7, 18–19, 22–25, 92, 95, 208, 219, 222, 237,

| 0 | Q |
|---|---|
| object names 19 | Q-parity 18 |
| operations 8, 21, 28, 42, 97–98, 110, 134, 161, 382, | quorom candidates 19 |
| 386–387, 390, 426–427, 430–432, 459, 539 | quorum disk 5, 19 |
| ordering 387 | quorum disks 19 |
| overview 1, 3, 10–11, 63, 154, 171, 221, 226, 265, | quotum disks 15 |
| | |
| 287–288, 290, 297, 315, 330, 334, 347, 372, 377, 460, | R |
| 463, 516 | RAID 3, 5, 7, 15, 17–18, 20, 22, 44, 220, 226–229, |
| overwritten 383, 534 | 231–235, 241, 258–260, 343, 345, 363–364, 378, 498, |
| | 501–502 |
| P | RAID 0 17, 227, 343 |
| parity strips 18 | RAID 1 7, 17–18, 343 |
| partner node 16 | RAID 10 17–18, 227, 343 |
| partnership 17, 115, 423, 425–426, 431, 433–438, 506 | RAID 5 7, 17–18, 227, 235, 343 |
| passphrase 547 | RAID 6 17–18, 227, 235, 343 |
| peak 432 | RAID array 7, 18, 502 |
| per cluster 29, 390, 431 | RAID arrays 15 |
| • | RAS 222, 497–498 |
| performance xv–xvi, 9, 15, 17, 20, 24–25, 28, 139, 229, | , |
| 233–239, 309, 325, 330, 344, 353–357, 360–362, 369, | real capacity 24, 161–162, 301, 315 |
| 377–379, 390, 427–428, 459–462, 495–496, 501–502 | rebalances 17, 263, 502 |
| performance improvement 360, 378 | receiver 31 |
| performance requirements 139 | Redbooks Web site |
| performance statistics 460 | Contact us xix |
| PHY 5, 14, 499–500 | redundancy 17–19, 36, 38, 41–42, 222, 226, 342–343, |
| PHYs 14, 499–500 | 345, 432, 501, 503 |
| physical location 27, 36 | redundant 5, 41, 431, 558–559 |
| PiT 28 | Redundant Array of Independent Disks 17 |
| PLOGI 27 | relationship 17, 23, 28, 112–113, 383–385, 389, |
| Point in Time 28 | 426–430, 432, 435–436, 439–442, 445–450, 452 |
| point in time 5, 28–29, 357, 385–386, 388, 427 | reliability 127, 344 |
| point-in-time copy 8–9, 382, 388 | Remote 10, 29, 32–33, 44, 114, 161, 341, 423, 425–427, |
| port 13–15, 26–27, 36, 39–42, 49, 147, 154–155, | 429–432, 436, 439–441, 444–447, 449–450, 457, 539 |
| 179–181, 192, 263, 288, 291–292, 294–295, 342, 431, | remote xix, 2, 9, 16–17, 27–28, 30, 36, 45, 53, 115, 247, |
| 460, 499–500, 502, 522, 527, 557–559 | 313, 342, 425–428, 431–432, 434–435, 439, 544 |
| Port Mask 288 | remote copy 28, 381, 423, 426, 428, 430-433, 439, 442, |
| ports 5–7, 13–14, 36, 38–42, 49, 105–106, 128, 136, | 450–455 |
| 148, 153, 176, 192, 205, 263, 288–289, 291–293, | removing 282, 538 |
| 295–296, 335, 342–345, 431, 482, 498–500, 502, 529, | restore points 28, 389 |
| 559 | restrict access 30 |
| ports of SAS drive 14 | Reverse FlashCopy 8, 28, 389, 555 |
| power supply 12, 14, 501, 503 | RFC3720 26 |
| power supply unit 12 | |
| P-parity 18 | 0 |
| Prepare 170, 473 | S |
| Prepared 386, 388 | SAN fabric 459, 462 |
| Preparing 127–128, 135, 139, 386, 388 | SAN planning 460 |
| presets 3, 226–229, 390, 392–393, 398 | SAN Volume Controller xv-xvi, 3, 488 |
| primary 25, 28, 43, 163–164, 166, 284, 301, 325–327, | SAN zoning 147, 288, 292, 431, 558 |
| 330, 360, 374, 376, 426–432, 435, 447–448, 454–455, | SAS 3, 5–6, 14–15, 36–37, 228, 356, 361, 365–366, |
| 461, 465 | 499–500, 502 |
| priority 309, 315, 517, 521 | SAS connector 14 |
| private key 544–547, 550 | SAS interfaces 14 |
| provisioning 6, 24, 77, 403–404, 459–460, 462 | SAS ports 13, 499 |
| PSU 12, 14, 503–504 | SBB 11 |
| public key 544, 546–548 | scalable 3 |
| PuTTY 544–547, 549–552 | scripts 464 |
| PuTTY Key Generator 546 | SCSI 3, 5-6, 26-27, 107, 172, 176, 184, 190, 206, 212, |
| Tutti Ney Generator 340 | 276, 279, 288, 297, 351, 386, 558–559 |
| | SDD 558-559 |

SDDDSM 130-131, 172-173, 515-516, 558-559

| secondary 28, 164, 167, 325–326, 426–432, 435, 446, 448, 454–455 | 360–361, 365–367, 369, 372–373, 375–376, 383, 390, 402–403 |
|--|--|
| secondary site 28, 427–428, 432 | storage pools 8, 17–18, 20, 24–25, 91, 163, 214, |
| security xv, 30, 130, 139, 474 | 219–220, 227, 241, 244, 253, 266, 268, 309, 325, |
| security features 30 | 329–330, 344–345, 351, 355–356, 361, 369, 372 |
| sender 31 | Storwize V7000 xv, 1, 3–12, 14–21, 24–28, 30–33, |
| separate console 3 | 35–36, 38–46, 48–49, 54, 61, 63–65, 69–70, 74–75, |
| | |
| Sequential 22 | 77–78, 80, 92, 94, 115, 127–128, 130, 132, 134–136, |
| sequential 22, 25, 71, 360 service tasks 16 | 138–139, 144, 147, 153–155, 157–160, 163, 170, 172, 174, 176, 179, 181–182, 186, 192, 199–200, 203, 205, |
| | 207–208, 212, 214, 233, 273–274, 282, 292, 300, 307, |
| showing 65, 75–79, 162, 515, 529 | |
| shrinking 307 | 309, 339–340, 342–345, 349, 353–356, 358–362, 364, |
| shut down 538–540 Simple Network Management Protocol 30 | 372, 377, 379, 423, 426, 430, 433, 435–436, 460, 463, |
| Simple Network Management Protocol 30 single sign-on 459, 461 | 486, 488, 492–493, 497–498, 500–503, 505–506, 508–510, 515, 517–518, 520–521, 533, 537–541, 544, |
| single sign-on 435, 461 single tiered 20, 355, 361 | 547, 551–552, 557–559 |
| single hered 20, 333, 361 site 8–9, 18, 28, 30, 32, 40, 46, 247, 426–428, 432, 498, | |
| | strand 6, 502 stripe 17–18 |
| 510 oloto 7 12 12 501 | Striped 22 |
| slots 7, 12–13, 501 | • |
| SMTP server 31, 44 | striping 17 Subsystem Device Driver DSM 120 121 172 |
| snapshot 5, 110, 391, 394–395, 399–400, 403 | Subsystem Device Driver DSM 130–131, 172 |
| SNIA XV | summary report 25, 480 |
| SNMP 30–31, 44, 123, 498 SNMP manager 31 | superuser 44–45, 52, 488, 509 |
| SNMP messages 31 | surviving node 13 |
| | Suspended 386, 388 SVC xvi, 3–4, 17, 513–514, 544, 547 |
| SNMP settings 31 | svcinfo 45, 372–374, 376, 552–557 |
| software upgrade 510, 513–514 Solid State Drives 353 | svctask 373, 375–376, 552–557 |
| solution 2–4, 9, 11, 25, 30, 65, 75, 80, 129–130, 163, 353 | svctask mkfcmap 555 |
| sort 535 | • |
| sorting 426, 535 | synchronized 9, 43, 160, 229, 321, 325–326, 330, 382, 426, 432, 442–443 |
| source 8–9, 28–29, 250, 256, 326, 360, 382–386, | synchronizing 58, 60 |
| 388–391, 395, 398, 402–406, 410, 413–414, 426, 428, | syslog manager 31 |
| 430, 439, 555 | Syslog messages 31 |
| source volume 8, 28, 383–384, 386, 388–389, 391, 426 | syslog messages 30–31 |
| Split 324, 326–327, 330 | syslog protocol 31 |
| splitting 25, 326, 344 | System Storage Productivity Center 30, 463 |
| SSD 14, 25, 43, 227–229, 233, 235, 261–262, 325, 353, | systems xv, 2, 4–5, 8–9, 17, 19, 25, 30, 32, 38–39, 43, |
| 356, 360–366, 369, 377–379, 460, 502 | 45, 93, 122, 127–128, 138, 163, 170, 199–200, 206–207, |
| SSH 45, 117, 488, 544, 546–547, 549–550 | 253, 307, 339–340, 342–345, 347, 351, 377, 382, 423, |
| SSH client 544, 549–550 | 426, 430, 433–434, 436, 459–461, 465, 515, 541, |
| SSH key 544, 547 | 557–558 |
| SSH keys 544 | 337 330 |
| SSPC 30, 463 | _ |
| stack 139 | T |
| state 5, 13, 15, 25, 43–44, 88, 226, 228, 261, 292, | T0 28, 383 |
| 385–386, 388, 413, 419, 421, 425–426, 428–432, 446, | target 8-9, 26-28, 177, 180, 192, 215, 250, 254, 288, |
| 499–500, 503, 515, 520 | 310, 330, 355, 360, 382–391, 395, 399, 404–406, |
| states 361, 384–385, 388, 429 | 408–415, 426, 430, 432, 439, 441, 555–556, 559 |
| statistics 25, 354, 358–360, 369, 375, 460, 462, 495 | target name 26 |
| stop 112–114, 257, 382, 386–387, 421, 436, 446, 448 | target volume 28, 382-384, 386, 389, 391, 406, 410, |
| Stopped 386, 388, 409, 421, 425, 437 | 414, 426 |
| stopping 257–258, 384, 386, 409 | thin provisioned 6, 24–25, 322, 385 |
| Storage Bridge Bay 11 | thin provisioned volume 24 |
| storage cache 354 | Thin provisioned volumes 8, 24 |
| storage capacity 19–20, 222 | Thin Provisioning 8 |
| storage pool 5–6, 18–22, 24–25, 94, 111, 157, 159–161, | thin-provisioned 28, 160–162, 182, 247, 390–391 |
| 208, 214–215, 219, 237–238, 241, 243–246, 248, 250, | thrashing 16 |
| 252, 254, 256, 261, 266, 268–271, 309–312, 320, | threshold warning 20 |
| 330–331, 333–334, 343–345, 350–351, 355–356, | tie 19 |
| , , | time xviii, 4-5, 8-9, 14, 17-19, 21, 25-26, 28-29, 31, 35, |

43, 46, 52-53, 124, 128-129, 132, 212, 220, 236, 292, 310, 327–328, 353, 357, 361, 370–371, 378, 382–386, 388, 391, 395, 420, 426–427, 443, 459, 463, 480, 506, 512, 521, 557 Time-Zero 28 Tivoli Integrated Portal 461, 464 Tivoli Storage Productivity Center 28, 459-466, 477, 486, 493 Tivoli Storage Productivity Center for Replication 460-461, 463-464, 477 **TotalStorage Productivity Center** component install 477 TPC 17, 459-466, 477, 479-480, 483-486, 492-493, 496 traffic 139-140, 142, 154, 181, 356 troubleshoot 30, 44, 224, 497 typical install 469

U

UDP 31 unbalanced 17, 502 Unmanaged 18, 241 unmanaged 19, 94, 200, 215, 241–243, 246, 256, 263, 271, 345 unmanaged MDisk 246 update 121, 124, 128, 263, 307, 387, 427, 503, 508, 559 updates xvi, 28, 263, 383, 387, 426, 428, 432, 510 upgrade 124, 497–498, 508, 510, 512–514 usage statistics 25 USB 13, 46, 48–49, 499 used capacity 20, 24 used free capacity 24 User Datagram Protocol 31

V

VDisk 556 viewing 516 virtual capacity 24, 161-162, 301 virtual disks 15 virtualization xv-xvi, 1-3, 7, 15-16, 21-22, 24, 53, 200, 311-312, 342, 344-345, 349 volume 5-9, 13, 15-16, 18-25, 28-29, 96-97, 100, 102, 110-112, 138, 157, 159-169, 171-174, 176, 184, 187, 190, 192, 205–206, 209, 211–212, 214, 216, 220, 241, 246-247, 250, 252-255, 257, 264, 269, 276, 278-281, 297, 299-332, 335-336, 343, 349, 351, 354-356, 360-361, 368-369, 374-376, 379, 382-391, 394-395, 398-399, 402-406, 409-415, 426-428, 430-432, 439, 441-442, 447-448, 461, 502, 534, 553-554, 557-559 Volume Mirroring 8, 44 volume mirroring 20, 25

W

warning capacity 24 warning threshold 316 warranty period 31 Windows 27, 45–46, 48, 127–135, 170–171, 173–176, 181–182, 377, 463–466, 470, 477, 515, 519, 537, 557–559 working with 520 WWNNs 7 WWPNs 7, 15, 39, 210, 291, 343–345

Ζ

zero contingency 24 zone 39, 127, 200, 208, 216, 342–343, 345, 459 zoning recommendation 431

250 by 526 which equals a spine width of .4752". In this case, you would use the .5" spine. Now select the Spine width for the book and hide the others: Special>Conditional To determine the spine width of a book, you divide the paper PPI into the number of pages in the book. An example is a 250 page book using Plainfield opaque 50# smooth which has a PPI of 526. Divided

Conditional Text Settings (ONLY!) to the book files. Text>Show/Hide>SpineSize(->Hide:)>Set. Move the changed Conditional text settings to all files in your book by opening the book file with the spine.fm still open and File>Import>Formats the

Draft Document for Review March 3, 2011 2:46 pm

7938spine.fm



Redbooks

Implementing the IBM Storwize **V7000**

1.5"<-> 1.998" (1.5" spine)

789 <->1051 pages





Implementing the IBM Storwize V7000

460 <-> 788 pages 0.875"<->1.498" (1.0" spine)

Implementing the IBM Storwize V7000

IBM

Redbooks







Implementing the IBM Storwize V7000

90<->249 pages 0.17"<->0.473 (0.2"spine)

53<->89 pages 0.1"<->0.169 (0.1"spine)

Text>Show/Hide>SpineSize(->Hide:)>Set. Move the changed Conditional text settings to all files in your book by opening the book file with the spine.fm still open and File>Import>Formats the 250 by 526 which equals a spine width of .4752". In this case, you would use the .5" spine. Now select the Spine width for the book and hide the others: Special>Conditional To determine the spine width of a book, you divide the paper PPI into the number of pages in the book. An example is a 250 page book using Plainfield opaque 50# smooth which has a PPI of 526. Divided

Draft Document for Review March 3, 2011 2:46 pm

Conditional Text Settings (ONLY!) to the book files.

7938spine.fm 574



IBM

V7000 Implementing the IBM Storwize

(2.5" spine) 2.5"<->nnn.n" 1315<-> nnnn pages

Implementing the IBM Storwize

IBM

Redbooks

Implementing the IBM Storwize V7000



Discover the exciting latest addition to the IBM virtualization family

Get 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® is introducing the IBM Storwize V7000, 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's offered at a competitive, midrange price.

IBM Storwize V7000 incorporates some of IBM's top 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 book will introduce the features and functions of the Storwize V7000 by example.

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

SG24-7938-00

ISBN