IBM

IBM System Storage DS3000: Introduction and Implementation Guide

- Sample configurations with step by step instructions
- Configuration and administration with Storage Manager
- DS3200, DS3300, DS3400, and EXP3000



Sangam Racherla
Corne Lottering
Mahendran Ramasamy
Vaclav Sindelar

Redbooks

ibm.com/redbooks



International Technical Support Organization

IBM System Storage DS3000: Introduction and Implementation Guide

January 2009

SG24-7065-01

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

Second Edition (January 2009)

This edition applies to the Version 10.35 of DS3000 Storage Manager, and to the IBM System Storage DS3200, DS3300, and DS3400.

© Copyright International Business Machines Corporation 2007, 2009. All rights reserved.
Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices xiii
Trademarks xiv

Preface xv
The team that wrote this book xv
Become a published author xviii
Comments welcome xviii

Summary of changes xix
January 2009, Second Edition xix

Part 1. Introduction 1

Chapter 1. Disk attachment technology 3
1.1 Fibre Channel disk attachment 4
1.2 Serial Attached SCSI (SAS) disk attachment 5
1.3 iSCSI disk attachment 8
1.3.1 iSCSI initiators and targets 9
1.3.2 iSCSI discovery 10
1.3.3 iSCSI security considerations 11

Chapter 2. What is new 13
2.1 DS3000 Storage Manager enhancements 14
2.1.1 New features 14

Chapter 3. DS3000 introduction 17
3.1 IBM TotalStorage Portfolio 18
3.2 DS3000 products 18
3.2.1 IBM System Storage DS3200 20
3.2.2 IBM System Storage DS3300 21
3.2.3 IBM System Storage DS3400 22
3.2.4 IBM System Storage EXP3000 22
3.2.5 Product comparison of DS3200, DS3300, and DS3400 24
3.3 Premium Features 25
3.4 DS3000 Storage Manager 26
3.4.1 DS3000 Storage Manager components 27
3.4.2 Management methods 30

Part 2. Getting started 33

Chapter 4. Cabling 35
4.1 Enclosure ID settings 36
4.2 SAS cabling 36
4.3 EXP3000 attachment 37
4.4 Ethernet management 40
4.5 Host attachment 40
4.6 Power cabling 45
4.7 DS3000 Front and Power LEDs 46
4.8 Powering on and off 47

© Copyright IBM Corp. 2007, 2009. All rights reserved.

Chapter 5. Basic hardware configuration 49
5.1 Storage subsystem concepts 51
5.1.1 RAID levels 51
5.1.2 Arrays and logical drives 55
5.1.3 Hot spares 57
5.1.4 Storage partitioning 58
5.2 DS3000 Storage Manager installation 59
5.3 Add a DS3000 subsystem in Storage Manager 68
5.4 Initial Setup Tasks wizard 71
5.4.1 Locating the storage subsystem 72
5.4.2 Renaming the storage subsystem 72
5.4.3 Setting a storage subsystem password 73
5.4.4 (Optional) Viewing and enabling Premium Features 74
5.4.5 (Optional) Changing the network configuration 76

Chapter 6. Installing IBM System Storage DS3000 Storage Manager 81
6.1 Installing DS3000 Storage Manager on Microsoft Windows 2003 and 2008 82
6.1.1 Installation preparation 82
6.1.2 Installing or upgrading the Storage Manager Client on Microsoft Windows 2003 or 2008 82
6.2 Installing Storage Manager on Linux 90
6.2.1 Preparing for install 90
6.2.2 Installing or upgrading Storage Manager using the GUI 91
6.2.3 Installing Storage Manager using a text console 96

Chapter 7. Host configuration 101
7.1 HBA configuration and driver installation 102
7.2 Multipath support 104
7.2.1 Microsoft MPIO 104
7.2.2 Linux RDAC 105
7.2.3 AIX MPIO 106
7.2.4 IBM System i with VIOS 107

Part 3. Administration 109

Chapter 8. Administration - Enterprise 111
8.1 Enterprise Management window overview 112
8.1.1 Task Assistant 112
8.1.2 Enterprise Management window 113
8.2 Functions in the Enterprise window 115
8.2.1 Subsystem context menu 115
8.2.2 The Enterprise Management window menu bar 124
8.2.3 The Quick Access buttons 125

Chapter 9. Administration - Configure 131
9.1 Configure host access 132
9.1.1 What a host is 132
9.1.2 Preparations in the host server 132
9.2 Configure hosts 140
9.2.1 Configure host access (automatic) 141
9.2.2 Configure host access (manual) 145
9.2.3 Configure host group 150
9.2.4 Create host-to-logical drive mappings 151
9.3 Configure storage 155

9.3.1 Configuring storage - automatic 156
9.4 Advanced functions - FlashCopy 183
9.4.1 FlashCopy: How it works. 183
9.4.2 Estimating FlashCopy repository logical drive capacity 185
9.4.3 Checking the status of the FlashCopy Premium Feature 187
9.4.4 Creating a FlashCopy using the wizard 187
9.4.5 Using the FlashCopy logical drive. 195
9.5 Advanced functions - VolumeCopy 196
9.5.1 Introduction to VolumeCopy 196
9.5.2 Copying data for greater access 198
9.5.3 Creating and managing VolumeCopy copies 198
9.5.4 Understanding VolumeCopy 198
9.5.5 VolumeCopy and performance considerations 200
9.5.6 VolumeCopy: Step-by-step 201
9.5.7 Create a VolumeCopy pair 201

Chapter 10. Administration - Summary 207
10.1 Status box 210
10.1.1 Storage subsystem status. 210
10.1.2 Operations in Progress 210
10.1.3 Alert status 211
10.1.4 Connection lost 211
10.2 Hardware components 212
10.2.1 Storage Subsystem Profile 212
10.2.2 Drives 213
10.2.3 Hot spare drives 214
10.3 Capacity 215
10.4 Hosts and mappings 216
10.4.1 Configured hosts. 216
10.4.2 Host-to-Logical Drive Mappings 217
10.4.3 Storage partitions 217
10.5 Arrays and logical drives 218
10.6 Information Center 219

Chapter 11. Administration - Modify 221
11.1 Modify hosts 223
11.1.1 Edit Host Topology 223
11.1.2 Edit host-to-logical drive mappings 228
11.2 Modify storage. 231
11.2.1 Rename Arrays and Logical Drives. 231
11.2.2 Delete Arrays and Logical Drives 232
11.2.3 Add Free Capacity (Drives). 234
11.2.4 Change Array RAID Level. 237
11.2.5 Replace drives 238
11.2.6 Change Logical Drive Ownership/Preferred Path 241
11.2.7 Modify FlashCopy Logical Drives 243
11.2.8 Manage Logical Drive Copies 250

Chapter 12. Administration - Tools 257
12.1 Storage Subsystem tools 259
12.1.1 Rename Storage Subsystem 259
12.1.2 Set or Change Password 260
12.1.3 Locate 262
12.1.4 View/Enable Premium Features 265

12.1.5 Change Enclosure ID Numbers 272
12.2 Logical drive tools 272
12.2.1 Change Media Scan Settings 272
12.3 Controllers tools 278
12.3.1 Synchronize Controller Clocks 278
12.3.2 Configure Ethernet Management Ports. 278
12.3.3 Change Battery Settings 281
12.4 Accessibility tool 281
12.4.1 Inherit System Settings 281

Chapter 13. Administration - iSCSI 283
13.1 Authentication 285
13.1.1 Change Target Authentication 285
13.1.2 Enter Mutual Authentication Permissions 286
13.2 Identification and networking. 288
13.2.1 Change Target Identification. 289
13.2.2 Change Target Discovery 290
13.2.3 Configure iSCSI Host Ports. 291
13.3 Sessions and statistics 295
13.3.1 View/End iSCSI Sessions 295
13.3.2 View iSCSI Statistics. 297

Chapter 14. Administration - Support 299
14.1 Gather Support Information. 301
14.1.1 Save Support Information 302
14.1.2 Change Automatic Collection Setting 303
14.2 View Storage Subsystem Profile. 304
14.3 Download firmware 307
14.3.1 Upgrade by DS3000 Storage Manager. 308
14.3.2 Upgrade using the DS3000 Controller Firmware Upgrade Tool 324
14.4 Manage Controllers. 330
14.4.1 Place Controller Online or Offline 331
14.4.2 Redistribute Logical Drives 333
14.5 Advanced Support Tasks 334
14.5.1 View Event Log 335
14.5.2 Import/Export Array. 340
14.5.3 Change Remote Login Setting 346
14.6 Recover from Failure. 347
14.7 View Online Help. 349
14.8 About IBM System Storage 3000 350

Chapter 15. Remote Support Manager 353
15.1 DS3000 Remote Support Manager for Storage 354
15.1.1 Hardware and software requirements. 355
15.1.2 How RSM for Storage works. 356
15.1.3 Notification e-mail and events filtering 357
15.1.4 Remote access methods 358
15.1.5 RSM management interface 359
15.1.6 RSM security considerations. 361
15.2 Installing and setting up RSM 362
15.2.1 Installing the host OS 362
15.2.2 Installing RSM. 362
15.2.3 Setting up RSM. 362
15.2.4 Configuring SNMP traps in Storage Manager. 375

15.2.5 Remote Access connectivity 376
15.2.6 Activating RSM 382
15.2.7 Managing alerts 383

Part 4. Sample configurations 389

Chapter 16. SAS configuration 1 - Windows 2003 host 391
16.1 Equipment required 392
16.2 Installing SAS HBAs 392
16.2.1 SAS HBA driver installation 393
16.3 Installing DS3000 Storage Manager host software 395
16.4 Host access configuration 396

Chapter 17. SAS configuration 2 - Linux host 405
17.1 Equipment required 406
17.2 Installing SAS HBAs 406
17.2.1 SAS HBA driver installation 407
17.3 Installing DS3000 Storage Manager software 407
17.4 Installing RDAC for Linux 408
17.4.1 Building the RDAC driver 409
17.4.2 Installing the RDAC driver 409
17.4.3 Verification steps 410
17.4.4 Configuring disk space in Linux 411

Chapter 18. SAS configuration 3 - Windows 2008 on an IBM BladeCenter HS21 server 413
18.1 Equipment required 414
18.2 IBM BladeCenter setup 414
18.2.1 Installing Windows Server 2008 414
18.2.2 HS21 SAS Expansion Cards 415
18.2.3 Recording the SAS Expansion Card WWPN 415
18.2.4 HS21 SAS Expansion Card device driver 418
18.2.5 SAS Connectivity Modules 418
18.2.6 SAS Connectivity Module firmware update 419
18.2.7 Configuring the SAS connectivity module 421
18.2.8 SAS Connectivity Module zoning 422
18.3 Installing DS3000 Storage Manager host software 425
18.4 Configure the disk space in Windows Server 2008 426

Chapter 19. SAS configuration 4 - RHEL V5.2 Linux on an IBM BladeCenter HS21 server 429
19.1 Equipment required 430
19.2 IBM BladeCenter setup 431
19.2.1 Installing Red Hat Enterprise Linux 5 Update 2 431
19.2.2 HS21 SAS Expansion Cards 431
19.2.3 Recording the SAS Expansion Card WWPN 431
19.2.4 HS21 SAS Expansion Card driver installation 434
19.2.5 SAS Connectivity Modules 434
19.2.6 SAS Connectivity Module firmware update 435
19.2.7 Configuring the SAS connectivity module 437
19.2.8 SAS Connectivity Module zoning 438
19.3 Installing DS3000 Storage Manager software 441
19.4 Installing RDAC for Linux 442
19.4.1 Building the RDAC driver 442

19.4.2 Installing the RDAC driver 443
19.4.3 Verification steps 443
19.4.4 Configuring disk space in Linux 444

Chapter 20. SAS configuration 5 - Boot from SAN with Windows 2008 on an IBM BladeCenter HS21 server 445
20.1 Equipment required 446
20.2 IBM BladeCenter setup 447
20.2.1 HS21 SAS Expansion Card 447
20.2.2 SAS Connectivity Modules 447
20.2.3 SAS Connectivity Module firmware update 448
20.3 Boot from SAN setup 451
20.3.1 SAS Connectivity Module configuration 451
20.3.2 SAS Connectivity Module zoning 451
20.3.3 SAS controller and SAS Expansion Card configuration 454
20.3.4 DS3200 boot volume setup 457
20.3.5 DS3200 host mapping 463
20.3.6 Installing Windows Server 2008 470
20.3.7 Installing DS3000 Storage Manager host software 472
20.3.8 Enabling a failover path for a Windows 2008 boot volume 473
20.3.9 Failover testing 478

Chapter 21. iSCSI configuration 1 - Windows 2003 and 2008 with a QLogic QMC4052 HBA 481
21.1 Network configuration 483
21.2 Installing Windows Server 2003 484
21.3 Installing Windows Server 2008 484
21.4 Installing the iSCSI HBA driver 485
21.4.1 Preparation 485
21.4.2 Installing the driver 486
21.5 Installing SANsurfer 489
21.6 Configuring the iSCSI HBA network settings with SANsurfer 495
21.6.1 Connecting with SANsurfer to the host where the HBA is installed 495
21.6.2 HBA data collection 496
21.6.3 Flash iSCSI HBA to latest level 499
21.6.4 Configure network settings 503
21.6.5 Test network settings 508
21.7 Setting up the DS3300 logical drives and host mapping 509
21.8 Configuring iSCSI targets in the host 510
21.9 Installing RDAC driver for DS3000 514
21.10 Test configuration 515
21.11 Enhancing iSCSI connection security 515
21.11.1 Setting up initiator authentication 515
21.11.2 Setting up target authentication 523

Chapter 22. iSCSI configuration 2 - Software Initiator on Red Hat Linux V5.2 529
22.1 Network configuration 530
22.2 Installing Red Hat Enterprise Linux 5 Update 2 531
22.3 Installing and configuring the Open iSCSI software initiator 532
22.4 Installing the RDAC failover driver 533
22.4.1 Preparing the host to install the RDAC driver 534
22.4.2 Installing the RDAC driver 534
22.5 Setting up the DS3300 logical drives and host mapping 536
22.5.1 Finishing the configuration 537

22.6 Changing the configuration at run time 541
22.7 Enhancing connection security 542

Chapter 23. iSCSI configuration 3 - Windows 2003 cluster. 545
23.1 Network configuration 546
23.2 Operating system installation 547
23.3 Preparing the cluster installation 547
23.4 Installing iSCSI software initiator. 548
23.5 Collecting the iSCSI qualified name of each cluster node 550
23.6 Installing IBM System Storage DS3000 Storage Manager 551
23.7 Configuring the IBM System Storage DS3300 554
23.8 Configuring the Microsoft iSCSI software initiator. 555
23.8.1 Configuring target discovery using iSNS server 556
23.8.2 Configuring target discovery using sendTargets. 558
23.8.3 Configuring login to discovered targets. 562
23.9 Preparing shared drives on a cluster node 568
23.10 Installing the cluster service on the first cluster node 569
23.11 Installing cluster service on the remaining nodes 575
23.12 Configuring the cluster network usage 580
23.13 Testing the cluster configuration 584

Chapter 24. FC configuration 1 - Emulex HBA boot blade server from SAN. 585
24.1 Cabling and zoning 586
24.2 Configuring the DS3400 591
24.3 Enabling boot from SAN with Emulex adapters 592
24.3.1 Set port switch speed 592
24.3.2 Enable adapter BIOS 593
24.3.3 Select boot device. 594
24.4 Installing Windows 2003 with the Emulex HBA. 597
24.4.1 Creating the Emulex HBA driver diskette 597
24.4.2 Installing Windows 2003 598
24.4.3 Installing multipath driver 600
24.5 Installing Emulex HBAnyware 602
24.6 Determining what to do if the boot path fails 608
24.6.1 Boot from SAN with one HBA 608
24.6.2 Multi-port failover Boot from SAN 609

Chapter 25. FC configuration 2 - Blade server Microsoft Cluster. 611
25.1 Preparing the environment 612
25.1.1 Cabling and zoning 612
25.1.2 Creating logical drives and host mapping with an SMcli script 612
25.2 Installing Microsoft Windows and device drivers. 613
25.2.1 Installing HBA drivers 614
25.2.2 Installing the multipath driver 617
25.2.3 Windows Disk Manager 620
25.3 Creating a Microsoft Cluster 626
25.3.1 Installing the second node 626
25.3.2 Editing the storage subsystem configuration for clustering by SMcli 627
25.3.3 Enabling the Cluster Service. 629
25.3.4 Joining the second node to the cluster 634

Chapter 26. FC configuration 3 - Linux SAN boot from a DS3400 with an IBM System x server 639
26.1 What is required 640

26.2 First FC HBA port configuration 641
26.3 SAN switch zoning for the first I/O path 644
26.4 Creating a logical drive 646
26.4.1 Host server definition 646
26.4.2 Creating a logical drive 650
26.4.3 Logical drive ownership 656
26.5 Configuring logical drive as boot device 658
26.6 Installing Linux 659
26.6.1 Installing the kernel source tree (RHEL4) 660
26.6.2 Updating the QLA2400 driver (RHEL4). 660
26.7 Installing Linux RDAC in RHEL5. 663
26.7.1 Building the RDAC driver 663
26.7.2 Installing the RDAC driver. 663
26.7.3 Verification 664
26.8 Enabling the second I/O path 665
26.8.1 Adding a second FC HBA port to the host server definition 665
26.8.2 Enabling a second FC HBA port for SAN boot 667
26.9 Verifying the installation 669
26.9.1 RDAC commands 670

Chapter 27. FC configuration 4 - QLogic HBA boot for IBM System x from SAN . . . 671
27.1 What is required 672
27.2 First FC HBA adapter configuration 673
27.3 Configuring the system BIOS 678
27.4 SAN switch zoning for the first I/O path 681
27.5 Creating a logical drive 682
27.5.1 Host server definition 682
27.5.2 Creating a logical drive 686
27.5.3 Logical drive ownership 692
27.6 Configuring a logical drive as a boot device 695
27.7 Installing Windows Server 2008 696
27.7.1 Installing DS3000 Storage Manager host software. 697
27.8 Enabling the second I/O path 698
27.8.1 Zoning configuration for second path 698
27.8.2 Adding a second FC HBA port to the DS3400 host server definition 698
27.8.3 Second HBA adapter configuration. 701
27.8.4 Verifying Multi-path in Windows 705
27.8.5 Failover testing 707

Chapter 28. FC configuration 5 - Boot from SAN with VMware ESX Server 711
28.1 What is required 712
28.1.1 Installing ESX Server on a DS3400 server - configuration procedure overview 712
28.1.2 Installing ESX Server SAN boot on a DS3400 - step-by-step procedure 713

Part 5. Appendices 723

Appendix A. FlashCopy example - Windows 2003 and 2008 725
Windows 2003 and 2008 - basic/regular disks. 726
Process overview 727
Detailed instructions for Windows 2003 and 2008 basic disks. 728

Appendix B. IBM Support Web site 733
Sample navigation procedure 734
My notifications 740

Appendix C. CLI 747
Running the CLI 748
Script Editor 748
General SMcli syntax 750
Adding a storage subsystem to the Storage Manager configuration 753
Showing defined subsystems in the Storage Manager configuration 754
Configuring alerts 755
Defining the mail server and e-mail address to send out the e-mail alerts 755
Defining email alert recipients 756
Deleting e-mail alert recipients 757
SNMP alert recipients 758
Issuing commands to the storage subsystem 758
Sample command: Save configuration script file 760

Related publications 763
IBM Redbooks publications 763
Other publications 763
Online resources 763
How to get IBM Redbooks publications 764
Help from IBM 764

Index 765

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.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

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:

AFS®
AIX 5L™
AIX®
Alerts®
BladeCenter®
DS4000™
DS6000™
DS8000™
FlashCopy®
i5/OS®
IBM®
POWER™
POWER5™
POWER6™
Redbooks®
Redbooks (logo) ®
System i®
System p®
System Storage™
System Storage DS®
System x™
TotalStorage®

The following terms are trademarks of other companies:

Novell, SUSE, the Novell logo, and the N logo are registered trademarks of Novell, Inc. in the United States and other countries.

QLogic, SANsurfer, and the QLogic logo are registered trademarks of QLogic Corporation. SANblade is a registered trademark in the United States.

VMware, the VMware "boxes" logo and design are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions.

Java, RSM, Solaris, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Active Directory, Microsoft, MS, Windows Server, Windows Vista, 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, other countries, or both.

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

This IBM® Redbooks® publication introduces the IBM System Storage™ DS3000, providing an overview of its design and specifications, and describing in detail how to set up, configure, and administer it. Since the DS3000 has different disk attachment options, we describe these different methods, including SAS and Fibre Channel. This edition covers updates and additional functions available with the DS3000 Storage Manager Version 10.35 (firmware level 7.35).

You learn how to install the DS3000 Storage Manager, and how to use its GUI and command-line options. We cover quick setup of the DS3000: creating arrays and logical drives and making the storage available to a hosts. You will also find advanced management tasks, including setting up and managing FlashCopy® and Volume Copy. Finally, we provide several configuration scenarios. The scenarios describe different methods for attaching various DS3000 models under Windows® and Linux®, including clustered and boot-from-SAN environments. These examples include step by step instructions using both the GUI as well as command-line scripts.

This book is intended for customers, IBM Business Partners, and IBM technical professionals who want to learn more about the capabilities and advanced functions of the DS3000 Series of storage servers with Storage Manager Software. It also targets those who have a DS3000 storage system and need detailed advice on how to configure and manage it.

Note: The file names and the screen captures shown in this entire book are only for basic illustration of steps and instructions to be followed. The actual version numbers displayed may vary depending on the version of the latest Storage Manager and Firmware that are available for download on the IBM Support Web site.

The team that 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.

© Copyright IBM Corp. 2007, 2009. All rights reserved.



Figure 1 The team: Corne, Mahendran, Sangam, and Vaclav

Sangam Racherla is an IT specialist and Project Leader working at the International Technical Support Organization, San Jose Center. He holds a degree in Electronics and Communication Engineering and has 8 years of experience in the IT field. He has been with the International Technical Support Organization for the past 5 years and has extensive experience installing and supporting the ITSO lab equipment for various Redbooks projects. His areas of expertise include Microsoft® Windows, Linux, AIX®, IBM System x™ and IBM System p® servers and various SAN and Storage products.

Corne Lottering is a Systems Storage Field Technical Sales Specialist (FTSS) supporting the System Storage sales team in IBM South Africa, Systems and Technology Group. He has been with IBM for 8 years and has experience in a wide variety of storage technologies, including the DS4000™, DS8000™, IBM SAN switches, N-Series, tape, and storage software. His daily support activities include pre-sales and post-sales support, including developing and presenting technical solutions and proposals. Since joining the System Storage team, he has been responsible for various implementation and support projects for customers across Africa.

Mahendran Ramasamy is a Project Manager supporting SAN Storage, Data Backup & Recovery and File System Management for India Software Labs, SWG Group, IBM India. He is a Brocade Certified Engineer and he has been with IBM for 4 years. He has experience in a wide variety of storage technologies, including the DS4000, DS6000™, DS8000, IBM SAN switches, ITSM, Tape Libraries, AFS®, GSA, storage management software, and HP Storage. His daily support activities include operations management and Project Management for India Software Lab's new initiatives. Since joining the ISL IS/IT Operations Team, he has been responsible for various implementation and support projects for internal customers across India Software Lab Locations.

Vaclav Sindelar is a Field Technical Support Specialist (FTSS) for IBM System Storage at the IBM Czech Republic headquarters in Prague. His daily support activities include pre-sales support for IBM Storage products. He has 5 years of FTSS Storage experience with a focus on IBM disk arrays and SAN. He has been with IBM since 2001 and worked as storage specialist before he came to IBM. He holds a Master's degree in computer science from the Technical University of Brno in the Czech Republic.

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

Deanna Polm, Jon Tate, Alex Osuna, Bertrand Dufrasne, and Emma Jacobs,
International Technical Support Organization, San Jose Center

The authors would like to express their thanks to the following people, whose expertise and support were integral to the writing of this book.

Doris Konieczny
Bruce Allworth
John Fasano
Gene Cullum
Danh T Le
Eric W. Konieczny
Michael Quillen
Harold Pike
Shawn Andrews
Chuck Grimm
Bob D Dalton
Paul Goetz
Dariusz Ferenc
John Sexton
Florian Graf
Nikolaj B Kjeldsen
Bill Willson
David Watts
Richard Mancini
Margaret Ticknor
Don Brennan
Jon Etkins
Hernando Bedoya
Khalid M Ansari
IBM

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

The authors of the first edition, *IBM System Storage DS3000: Introduction and Implementation Guide*, SG24-7065-00, published in October 2007, were:

Charlotte Brooks
Jure Arzensek
Hendrik Ernst
Daniel Koeck

Become a published author

Join us for a two- to six-week residency program! Help write a book dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you will develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Learn 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 in one of the following ways:

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

  ibm.com/redbooks

- Send your comments in an e-mail 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

Summary of changes

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

Summary of Changes
for SG24-7065-01
for IBM System Storage DS3000: Introduction and Implementation Guide
as created or updated on January 4, 2010.

January 2009, Second Edition

This revision reflects the addition, deletion, or modification of new and changed information described below.

New information

- New DS3000 Storage Manager V10.35 features (with firmware Version 7.35)
- Firmware upgrade validation and upgrade utility
- RAID 6
- Increase number of drives in RAID 0 or 1 volume group
- Greater than 2 TB logical drives
- IPv6 on the management ports
- Support bundle and diagnostic data capture
- Dynamic RAID migration
- Automatic support data collection
- Configurable failed drive replacement

Changed information

- Increased number of partitions on storage subsystem.
- Software Feature Pack and the EXP3000 Expansion License have been incorporated into the base models and are no longer needed as prerequisites to the other advanced function options.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

Part 1

Introduction

In this part of the book, we introduce the IBM System Storage DS3000.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

1

Disk attachment technology

In this chapter, we describe basic disk attachment methods in the context of the IBM System Storage DS3000. We discuss the following technologies:

- Fibre Channel (FC)
- Serial-Attached SCSI (SAS)
- Internet SCSI (iSCSI)

Fibre Channel has traditionally been used to attach storage subsystems in midrange and large scale environments. However, as the DS3000 products are geared towards Small and Medium Business (SMB) and departmental environments, SAS and iSCSI attachment technologies are supported as well.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

1.1 Fibre Channel disk attachment

Fibre Channel (FC) is a high-speed disk attachment technology, designed to connect a large number of storage devices to a number of host servers across a Storage Area Network (SAN). Fibre Channel Protocol (FCP) transfers SCSI commands and data across physical FC links.

FC supports a much higher number of devices and much longer cable lengths than SCSI. It has become the preferred disk attachment technology in midrange and large scale datacenter solutions.

At the time of writing, DS3000 Storage maximum FC throughput is 4 Gbps. In fact, 10 Gbps links can be used today, but only for SAN switch interconnection.

Host servers contain one or more FC Host Bus Adapters (HBA). The HBAs provide connectivity to the storage devices using FC cabling and SAN Switch.

For more information about Fibre Channel and SANs, see *Introduction to Storage Area Networks*, SG24-5470.

FC topologies

FC supports several connectivity topologies:

- Point-to-point

  This is the simplest topology and provides a direct link between an FC HBA inside a host server and a storage device.

- Arbitrated loop

  This topology can be used to interconnect several FC devices. A typical example would be to attach a certain number of host servers to an FC storage subsystem. A loop can consist of up to 126 devices.

  Devices on the loop use one-way ring communication. In any given moment, only two devices on the loop can communicate. This means the devices share bandwidth, so the arbitrated loop topology is not suitable for high performance requirements.

  Arbitrated loops were commonly implemented with the use of an FC hub. Even though this is physically a star topology, logically it will be a loop. Alternatively, devices can be connected in a daisy chain manner.

  Arbitrated loops are rarely seen these days, as switched fabrics have become the norm.

- Switched fabric

  The most commonly used topology in a typical SAN today is switched fabric. SAN switches are used to provide FC connectivity between the host servers and storage devices. Switched fabrics can become very complex in large scenarios, connecting hundreds of host servers to a very large number of storage subsystems.

  SAN switches provide optimized traffic flow and increased performance by allowing concurrent data transfers between many connected hosts and storage devices. Switched fabrics can provide dedicated bandwidth, as opposed to arbitrated loop technology, where the bandwidth is shared among all the devices in the loop.

FC protocol layers

The FC protocol is split into five layers, named FC0 through FC4. Let us look briefly at them:

- FC0 is the physical layer, which describes cabling, connectors, signalling, and so on. This layer defines the physical media implementation.
- FC1 is the data link layer. This layer contains the 8b/10b encoding and decoding of signals for transmission across the physical media.
- FC2 is the network layer and defines the main FC protocols. This layer defines how the frames are transferred.
- FC3 is the common services layer. This layer provides services such as multi-casting and striping.
- FC4 is the application protocol mapping layer. In storage connectivity applications, FCP protocol is used to encapsulate SCSI data into FC frames.

FC cable types

FC implementations can utilize either single-mode or multi-mode FC cables.

The name *multi-mode fiber* indicates that multiple modes, or rays of light, can travel through the cable core simultaneously. The multi-mode fiber cable uses a larger diameter core, which makes it easier to couple than the single-mode fibre cable. With a throughput of 8 Gbps, the length of the cable can be up to 150 m.

Single-mode fibre transfers a single ray of light. The core diameter is much smaller than the core of multi-mode cable. Therefore, coupling is much more demanding and tolerances for single-mode connectors and splices are very low. However, single-mode fiber cables can be much longer. Cable length can exceed 50 km.

Multi-mode cabling is much more common, as it is easier to work with and meets the requirements of most customer scenarios. However, in situations where very long cable lengths are needed, single-mode cabling will be required.

FC World Wide Names (WWN)

FC devices are presented with a unique identifier called *World Wide Name* (WWN). The WWNs are somewhat similar to the MAC addresses in Ethernet terms. For example, each FC HBA has its own WWN, which is hard-coded (or burned-in) during manufacturing. The HBA will be uniquely identified by the storage subsystem using its WWN.

1.2 Serial Attached SCSI (SAS) disk attachment

SAS technology evolved from parallel SCSI. While traditional SCSI is a parallel interface, SAS uses the serial communication protocol. However, SAS does use the SCSI command set to communicate with the attached devices. When parallel SCSI reached the transfer rate of 320 MBps, it became apparent that further transfer rate increases would be difficult to achieve. Even though the Ultra-640 parallel SCSI technology is available, it is not very practical due to very limited cable lengths. Therefore, most manufacturers moved from Ultra-320 SCSI to SAS, rather than adopting the Ultra-640 SCSI.

At the time of writing, typical SAS throughput is 3 Gbps full duplex. However, 6 Gbps capable SAS devices are already appearing and will become more widely used in the near future. SAS has the capability to reach 12 Gbps if the host can drive it at that speed. When the first 3 Gbps connection is full, the next 3 Gbps connection is used, and so on, up to four connections.

A SAS domain is a set of SAS devices connected together with SAS cabling. Each SAS device has a unique identifier, a World Wide Name (WWN). The role of this unique identifier might look similar to a SCSI ID, but there are some differences:

- On a SCSI bus, there can be up to 16 SCSI IDs, while a SAS domain can have up to 16256 devices.
- SCSI IDs are only unique within their SCSI bus. SAS WWNs are globally unique identifiers.

The SAS WWN is also referred to as the *Phy SAS address*.

SAS expanders

Basically, SAS uses point-to-point serial links. Point-to-point topology essentially dictates that only two devices can be connected; however, with the use of SAS expanders, the number of devices in a SAS domain can greatly increase. There are two types of expanders:

- Fan-out expanders

  A SAS domain can have only one fan-out expander. Devices attached to such a fan-out expander can be initiators, targets and edge expanders. The fan-out expander can attach to a maximum of 128 SAS devices.

- Edge expanders

  Edge expanders typically connect to the fan-out expander and to a mixture of SAS initiators and targets. Up to 128 devices are allowed on a single edge expander.

The maximum number of devices in a SAS domain is 128 edge expanders multiplied by 128 devices minus 128 connections to the fan-out expander. This gives a total of up to 16256 SAS devices.

In the current DS3000 implementation, up to 48 drives can be configured in a single DS3000 using three EXP3000 expansion units.

SAS protocol layers

The SAS protocol consists of four layers:

- The physical (or *phy*) layer

  This layer represents the hardware components, such as transceivers, which send and receive electrical signals on the wire.

- The link layer

  The link layer manages connections across phy interfaces.

- The port layer

  The port layer passes the SAS frames to the link layer. It also selects the most appropriate physical layer for data transmission (when multiple layers are available).

- The transport layer

  This layer encapsulates information into SAS frames and passes the frames to the port layer. It also disassembles the frames received from other SAS devices and relays the information to the driver/application.

Benefits of SAS

SAS technology provides several benefits over SCSI:

- SAS uses point-to-point technology, while SCSI uses multidrop cables. Because SAS devices have dedicated connections to the initiators, there is no bus conjunction; this provides more efficient throughput. On a SCSI bus, the bandwidth is shared between the devices, so as the number of devices on the bus increases, the throughput efficiency decreases.
- SAS has a lower signalling impact than SCSI, which again improves efficiency.
- SCSI configurations are sensitive to improper termination. SAS does not have such problems.
- SAS supports much longer cable lengths and higher maximum device attachment than SCSI.

SAS environments also support attachment of SATA disk drives. In fact, SAS could be viewed as a technology to transfer SCSI commands and data across a SATA-like physical interface. SAS was designed to be compatible with SATA on the physical level, but supporting higher transfer rates and longer cables.

SAS wide ports

Each SAS port includes four full duplex links or *lanes* within a single connector, as shown in Figure 1-1, with each lane running a speed of 3 Gbps. A single lane is used as the path to the drives; the second, third, and fourth lanes are used as overflow when concurrent I/Os overload the channel. For example, suppose the first link is transmitting data at 3 gigabits per second. If another block of data then needs to be written to disk, while the link one is still busy, then link two will manage the overflow of data that cannot be transmitted by link one. If link one finishes its transmission of data, then the next block of data will be transmitted on link one again; otherwise, another link will be used. In this way, for heavy I/O workloads, it is possible that all links are being used at certain times, providing a simultaneous data speed of 12 Gbps.

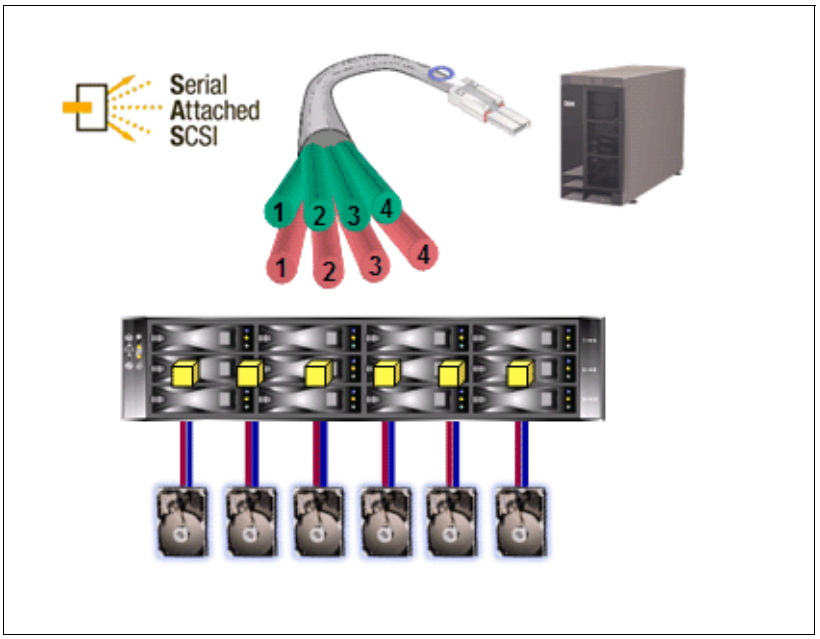


Figure 1-1 SAS wide ports

SAS drive technology

Figure 1-2 shows how SCSI and SAS drives are attached to the controllers. The older SCSI technology used a loop architecture, as shown on the left. If a controller has to communicate with the fourth drive in the loop, the communication must pass through drives one, two, and three. This is slow, and causes problems if any individual drive fails. The point-to-point topology used in SAS configurations, as shown on the right, means that there is a direct path to each drive from each controller, so communication can take place directly, with no effects caused by an individual drive failure.

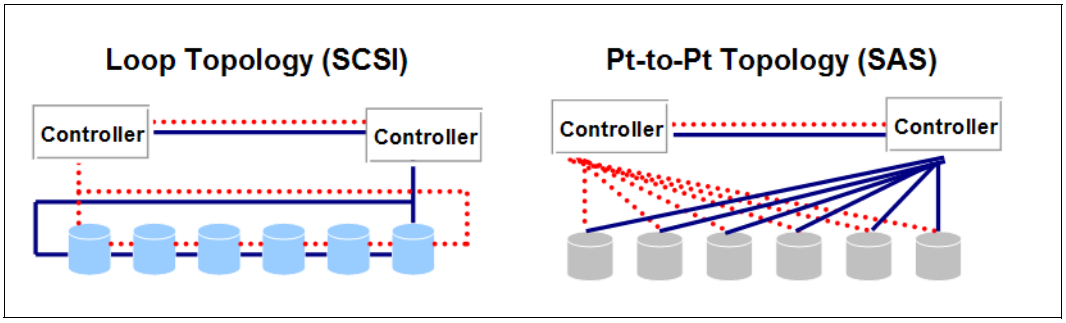


Figure 1-2 SCSI and SAS disk technology compared

1.3 iSCSI disk attachment

iSCSI stands for *Internet SCSI*. The basic idea of iSCSI is to use the Ethernet infrastructure on the physical layer and TCP/IP communication for transferring encapsulated SCSI commands and data. Using iSCSI, *IP SANs* can be implemented. Because an Ethernet infrastructure is readily available at low cost, iSCSI is a good choice for the SMB and also departmental SANs of larger companies.

IP SANs are a cheaper alternative to FC SANs; however the lower cost of iSCSI also implies lower performance and scalability. Encapsulation has an impact and the transfer rate is lower. A typical Ethernet network operates at 1 Gbps, while an FC SAN can run up to 8 Gbps. However, there are ways to address iSCSI performance:

- While the host servers can use almost any Ethernet network interface card for iSCSI traffic, this does mean that the CPUs in the host server have to run the iSCSI stack (to perform encapsulation of SCSI commands and data). This causes CPU and memory overhead, which can impact performance.

  For increased performance, it is better to use dedicated iSCSI HBAs to process the TCP/IP stack. This technology is known as *TCP Offload Engine (TOE)*. TOE technology relieves the CPUs in the host server from having to process the SCSI encapsulation, which can lead to better performance.

- Ethernet transfer rate is growing. 10 Gbps Ethernet is coming. The standards for 10 GbE have been defined; however, at the time of writing, this is a fairly new technology, which has yet to gain wider commercial acceptance. Initial offerings will also carry significant cost, which is expected to drop as the user base increases.

1.3.1 iSCSI initiators and targets

iSCSI uses the concept of initiators and targets, as shown in Figure 1-3.

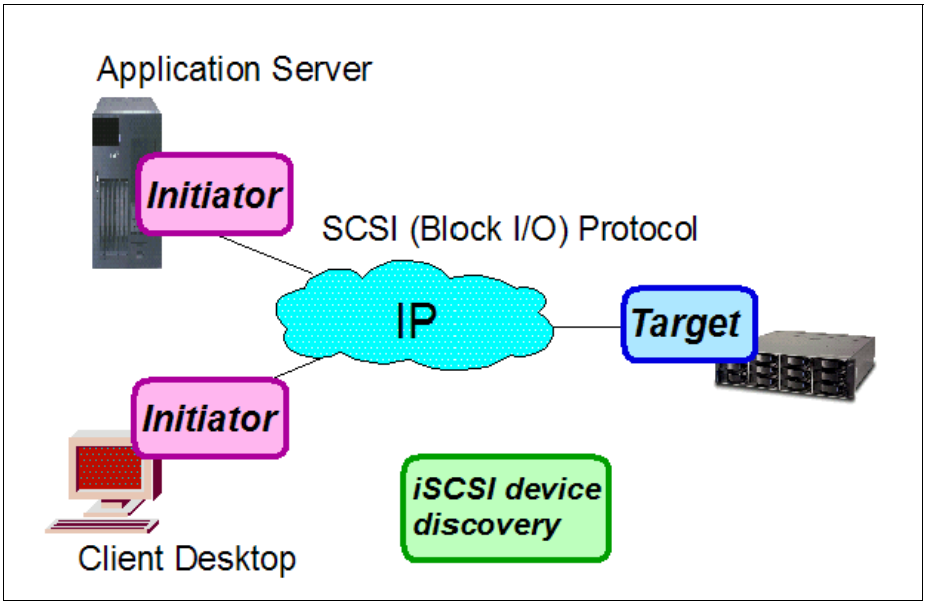


Figure 1-3 iSCSI components

The iSCSI initiator can be either an iSCSI HBA inside a host server, or you can define a software iSCSI initiator by using an iSCSI stack and an Ethernet network adapter. An example of an iSCSI stack is the Microsoft iSCSI Software Initiator, which runs on Windows Server® 2003 and Windows Server 2008. At the time of writing, the current version is V2.07 and is available for download from the Microsoft Web site. For Windows Sever 2008, the iSCSI initiator is included in the box.

Note: Even though the Microsoft iSCSI Software Initiator is available on different operating systems, including Windows XP, Windows 2000, and Windows Vista®, DS3300 only supports Windows Server 2003 and Windows Server 2008.

Refer to the IBM System Storage DS3000 series Interoperability Matrix for the complete list of the supported operating systems. This matrix can be found at:

`http://www-03.ibm.com/systems/resources/systems_storage_disk_ds3000_pdf_interop.pdf`

For the IBM AIX operating system, refer to the “iSCSI software initiator and software target” topic at `http://publib.boulder.ibm.com/infocenter/systems/index.jsp`.

iSCSI HBAs are physically very similar to Ethernet network cards. However, they typically contain a ROM-based BIOS utility that can be used to configure the iSCSI settings and HBA operational parameters. In this respect, they resemble SCSI and FC HBAs. Inside the operating system, the iSCSI HBAs are classified as storage adapters.

The iSCSI target is typically the storage device or subsystem, for example, DS3300. Other types of peripheral devices, like tape drives and medium changers, can act as iSCSI targets as well.

iSCSI naming

The iSCSI initiators and targets on a SAN are known by their respective iSCSI names, which must be unique. The iSCSI name is used as part of an ISCSI address, as part of all sessions established between initiators and targets. Next, we describe two types of iSCSI names:

- IQN
- EUI

IQN

A commonly used format of iSCSI names is the *iSCSI Qualified Name* (IQN). The format of an IQN is:

```
iqn.yyyy-mm.{reversed domain name}
```

For example, an iSCSI HBA inside a host server named Rhine in the domain rivers.local would be assigned the following IQN:

```
iqn.2007-06.local.rivers.rhine
```

The DS3300 uses IQN names.

EUI

An alternative type of iSCSI name is the Enterprise Unique Identifier (EUI). The format of an EUI is *eui* plus 16 hex digits. For example:

```
eui.acdc15882005bdef
```

iSCSI addressing

The iSCSI address has the following format:

```
<IP address>[:<port>]/<iSCSI name>
```

The IP address can be either IPv4, IPv6, or the fully qualified domain name. The <port> is optional; it specifies the TCP port that the target is listening for connections on. If it is not used, the well-known iSCSI port (3260) is assumed. The <iSCSI name> is the IQN or EUI name of the device. It is optional.

The iSCSI address specifies a single path to an iSCSI target. The iSCSI address is primarily used during discovery.

1.3.2 iSCSI discovery

iSCSI discovery allows an initiator to find the target(s) to which it has access. This requires a minimum of user configuration. Several methods of discovery may be used:

A list of targets at the initiator

An administrator can statically define the iSCSI targets to the host system initiator. This process allows the administrator to specify the iSCSI target node name and IP address:port to the host system initiator or its host bus adapter (HBA). iSCSI HBAs should support an administrator defining this information. This type of discovery is useful in small installations and is known as *static discovery*.

Queries to known iSCSI servers

An iSCSI initiator can probe its environment and, when a possible iSCSI target is found, start a *discovery session* with the target by issuing a **`SendTargets`** command. The target can reply to a **`SendTargets`** command by returning a list of all iSCSI target nodes it knows about.

Queries to an Internet Storage Name Server (iSNS)

The Internet Storage Name Server permits iSCSI targets to register with a central point. The administrator can set up discovery domains so that when a host iSCSI initiator queries the central control point for the locations of iSCSI storage controllers, only the authorized controllers are reported. The iSNS server can be located by one of the following techniques:

- iSCSI initiators multicasting to the iSNS server
- Setting the iSNS server IP address in the DHCP server
- Setting the iSNS server IP address in the iSCSI initiator or target
- Setting the iSNS server IP address in the SLP server (see “Service Location Protocol (SLP)”)

Service Location Protocol (SLP)

The Service Location Protocol can be used to locate iSCSI target devices. SLP operates with three agents:

- User agent (UA): Works on the client (iSCSI initiator) to help establish contact with a service (iSCSI target). It does this by retrieving information from service agents (SA) or directory agents (DA).
- Service agent (SA): Runs on the iSCSI target device to advertise the service and its capabilities.
- Directory agent (DA): Collects service advertisements from the iSCSI targets.

1.3.3 iSCSI security considerations

FC disk attachment uses a separate FC SAN, not accessible to Ethernet network users. iSCSI, on the other hand, is a SAN technology that uses the Ethernet network, which is a lot more vulnerable to intrusion. Therefore, iSCSI security is very important.

iSCSI connection authentication

One way to provide secure data transmission is by means of authentication. Although authentication is optional, its use should be considered mandatory in all security-sensitive environments. When enabled, the iSCSI target will authenticate the initiator. Optionally, the initiator can authenticate the target as well. Each connection within a session has to be authenticated. Several authentication methods can be used:

- Challenge Handshake Authentication Protocol (CHAP)
- Secure Remote Password (SRP)
- Kerberos V5 (KRB5)
- Simple Public-Key generic security service API Mechanism (SPKM1)
- Simple Public-Key generic security service API Mechanism (SPKM2)

In our sample configurations, we used CHAP.

IP Security (IPSec)

The authentication methods listed above prevent unauthorized initiators from accessing the targets, but there is no protection on the packet level. For increased security, the actual packets should be protected as well. As iSCSI relies on TCP/IP communication, IP Security (IPSec) can be used to achieve this.

IPSec authenticates and encrypts each packet in the IP data stream. There are two IPSec modes:

- Transport mode

  With transport mode, only the payload in each packet is encrypted. The IP header is left unencrypted, so the routing works just the same as without IPSec.

- Tunnel mode

  With tunnel mode, the entire packet is encrypted, including the IP header. This means that the whole encrypted packet must be encapsulated in a new IP packet, so that routing will function properly.

IPsec is commonly used to set up Virtual Private Networks (VPN).

2

What is new

This chapter is a brief overview of additions and enhancements available with Storage Manager V10.35 (controller code Version 07.35).

This chapter is intended for readers already familiar with most of the DS3000 concepts and features who just need a quick overview of the latest changes.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

2.1 DS3000 Storage Manager enhancements

IBM System Storage DS3000 Storage Manager Version 10.35 (controller code Version 07.35) includes all the functions already available in previous Storage Manager releases, and also offers several significant new features.

Note: Because of the design changes, upgrading the firmware to Version 07.35 is a non-concurrent upgrade and it must be done offline.

The controller code (firmware) Version 7.35 is based on a new layered code design that provides better scalability and flexibility for future code enhancements. As part of the new design, the configuration database (DACstore region) has also been restructured. This new structure provides the foundation for supporting some of the new features, such as greater than 2 TB logical drives, RAID 6, and an increased number of logical partitions.

2.1.1 New features

Note: The DS3000 Software Feature Pack and the EXP3000 Expansion License have been incorporated into the base models and are no longer needed as prerequisites to the other advanced function options. The required controller firmware that enables this function must be at firmware level V06.70.23.00 or later, and it is available on the DS3000 product support Web site at:

`http://www-03.ibm.com/systems/storage/disk/entry/index.html`

You can also refer to the IBM System Storage DS3000 series Interoperability Matrix at:

`http://www-03.ibm.com/systems/resources/systems_storage_disk_ds3000_pdf_interop.pdf`

IBM System Storage DS3000 Storage Manager Version 10.35 introduces the new features discussed in this section.

RAID 6 support

RAID 6 technology is designed to provide improved data protection against multiple disk drive failure. RAID 6 uses a double parity check implementation (designated as p+q). Dynamic RAID-level migration is supported, but may require capacity expansion prior to RAID migration.

Support for more than 30 drives for RAID 0 and 1

A logical drive can be created using all the drives in the storage system. Although not a recommendation, you could now create a logical drive using a possible maximum of 48 drives for a DS3000.

Support for greater than 2 TB logical drives

You can now create greater than 2 TB (terabyte) logical drives. If you combine this feature with the ability to also now support more than 30 physical drives (assuming RAID 0 or RAID 1), and if you use 450 GB SAS drives in a DS3000 with its maximum of 48 disk drives, you can create a 21 TB logical drive. (Keep in mind, however, that the OS making use of the logical drive can impose other limitations.)

Increased maximum number of partitions

The maximum number of supported partitions for DS3000 prior to 07.35 firmware was 16; this number was increased to 32. DS3000 also has a new option for eight partitions now.

Support bundle collected upon critical events

The event monitor on the client system, upon occurrence of a critical event, saves diagnostic data (support bundle) to the local hard drive of the client system in the same area already used for other recovery information. This information is stored and is not overwritten for 72 hours. E-mail notification can be set up for all critical events.

Use Diagnostic Data Capture for additional failure events

The goal of this feature is to capture enough information about the controller state at the time of an unusual event and store that diagnostic data for later retrieval.

Host support

Solaris™ 10 on SPARC host support was introduced with SM 10.35 and Firmware 7.35. This supports extends to I/O only and the storage subsystem has to be managed from a different operating system.

Host support is also updated for the existing operating systems listed below:

- AIX 5L™ V5.3.N and AIX V6.1.B for DS3400
- Red Hat V4.7 and Red Hat V5.2 on Intel® and Power systems for all DS3000 products
- SLES V9.4 and SLES V10.2 on Intel and Power systems for all the DS3000 products
- VMware® V3.0.2 for DS3400, VMware V3.5.1 / 3i for DS3200, DS3300, and DS3400
- NetWare V6.5 SP8 for DS3200 and DS3400

Additional support was also announced for Virtual IO Server (VIOS) V1.5.2 with DS3400 and SVC V4.2 and V4.3 with DS3400.

Note: Refer to the IBM System Storage DS3000 series Interoperability Matrix for the latest information at:

http://www-03.ibm.com/systems/resources/systems_storage_disk_ds3000_pdf_interop.pdf

Hardware enhancements and support

There are now the following hardware enhancements and support:

- All DS3000 products now support 450 GB SAS drives.
- IBM BladeCenter® 8 Gb FC switches and 8 Gb Daughter Cards are now supported
- DS3200 can now be attached to BladeCenter SAS Switch with the following enhancements.
  - Multi-Initiator support with BladeCenter SAS switch supporting 14 blades within one BladeCenter.
  - 48 SAS and SATA drives are now supported.
  - Support now extends to both booting from SAN and data for DS3200.
  - All blade types (with exception of workstation blade HC10) are now supported.
  - Red Hat V5.2, SLES V10.2, Windows Server 2003 R3, Windows Server 2008, and VMware V3.5.x are now supported.

- AIX (Versions 5L V5.3N and V6.1B) is now supported on Power Blades.

- IBM System p servers with Linux is now supported with DS3300.

Controller Ethernet Management

IPv6 on the management ports is now supported and has to be manually enabled.

Increased number of Global Hot Spares (GHS)

The new controller code allows the creation of unlimited Global Hot Spares. You can use either the Automatic option or Manual option to set the GHS.

Migration during a drive tray relocation

Storage Manager V10.35 (with controller firmware V7.35) incorporates new export and import options to safely move arrays between different DS3000 systems without losing data. This allows you to transfer configurations between enclosures (including SAS to FC, FC to iSCSI, and so on) without having to worry about enabling and disabling NVSRAM options.

This is very helpful when you have to upgrade or replace a DS3000 system with a new model or faster disks, but want to preserve the expansions and their data. The export/import options check that all the conditions to support the disk migration are met before placing the disks offline and allowing removal of the disks. Now, instead of using the option to place an array offline, just select the Export Array option from the source machine. Select the **Import Array** option in the destination machine to accept the exported disks with their data.

Under the Advanced Support Tab, there is an option to flag the array to be exported.

Drive Migration warning

This feature allows the user to manually intervene and permits safer migration scenarios. (When migrating configured logical drives to an existing storage subsystem, there is a warning given if the number of the logical drives being added will overrun the maximum number of drives allowed on this particular subsystem.) This feature thus prevents users from importing more logical drives than a particular subsystem can support.

Configurable failed drive replacement

You can now designate a hot spare to be a permanent member of a logical drive, thus eliminating the copyback portion once the failed drive has been replaced. The replacement drive in that case is in an unassigned state.

RAID redundant data verification prior to read

This feature (Pre-Read Redundancy Check) is supported on all RAID levels (except RAID 0) and enables verification of redundant RAID data consistency before returning the requested read data back to the host. This must be enabled on a logical drive basis. The function is automatically disabled when a logical drive becomes degraded. If a mismatch is detected, the system returns a media error - unrecoverable read error (03/11/00). The system does not attempt to correct or recalculate parity automatically.

Windows RDAC no longer supported

RDAC for Windows is not supported with this new version of Storage Manager (V10.35). For multipathing and failover support, you have to install the SM failover driver (also referred to as MPIO DSM).

3

DS3000 introduction

In this chapter, we introduce the IBM System Storage DS3000 product line and functionality. The product line consists of three storage subsystems: DS3200, DS3300, and DS3400. We describe the EXP3000 SAS disk drive enclosure as well.

We also explain the Premium Features philosophy and how the DS3000 Storage Manager utility works.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

3.1 IBM TotalStorage Portfolio

IBM has brought together into one family, known as the DS family, a broad range of disk systems to help small to large size enterprises select the right solutions for their needs. The DS family combines the high-performance IBM System Storage DS6000 and DS8000 Series of enterprise servers that they inherit from the ESS, with the DS4000 series of mid-range systems, and other line of entry systems (DS3000).

The DS3000 series is a good fit for the entry to mid-range SAN and direct attach market space. DS3200 and DS3400 systems provide a smooth transition to the DS4000 series systems, offering robust functionality, exceptional reliability, and simple storage management. The overall positioning of the DS3000 series within the IBM System Storage DS® family is shown in Figure 3-1.

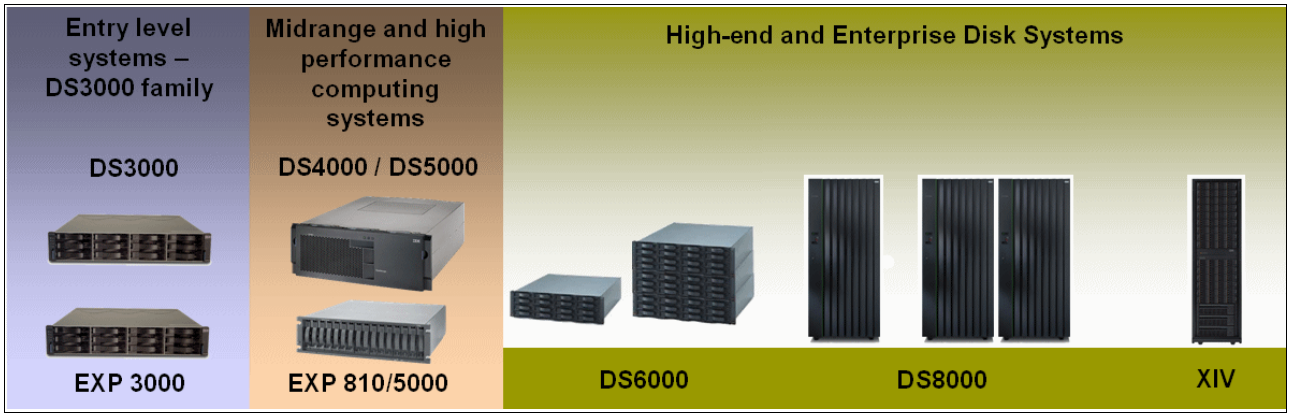


Figure 3-1 IBM TotalStorage® DS family

3.2 DS3000 products

The IBM System Storage DS3000 is a family of storage subsystems supporting up to two redundant RAID controllers and one of the following host server connectivity options:

- SAS on DS3200
- iSCSI on DS3300
- FC on DS3400

The basic differences between the DS3000 families from the hard disk point of view is shown in Table 3-1.

Table 3-1 DS3000 families HDD support

Feature	DS3200	DS3300	DS3400
Drive support	SAS/SATA	SAS/SATA	SAS/SATA
Host ports	SAS	iSCSI	FC
Maximum drives	48	48	48
Drive capacity (max)	450 GB SAS / 1 TB SATA	450 GB SAS / 1 TB SATA	450 GB SAS / 1 TB SATA

Note: In DS3000 families, you can add three expansion units and each unit can have a maximum of 12 drives. Thus, you will get a maximum of 48 drives in total.

Each DS3000 is a 2U rack mountable enclosure, containing one or two RAID controllers, two power supply units, and up to 12 SAS disk drives. The disk drives are installed at the front, as shown in Figure 3-2.

For larger disk configurations, you can attach up to three EXP3000 SAS expansion enclosures for a maximum of 48 disk drives. When using 450 GB SAS disk drives, the maximum storage capacity exceeds 21 TB.



Figure 3-2 DS3000 front view

RAID controllers and power supply units are installed in the rear, as shown in Figure 3-3.



Figure 3-3 DS3200 rear view

RAID controllers support RAID levels 0, 1, 3, 5, 6, and 10. Each controller has 512 MB of cache (upgradeable to 1 GB) and a cache battery. The battery can maintain the data in cache for up to three days in the event of a power outage.

In dual controller configurations, the left controller is A and the right is B, when seen from the rear view of the subsystem.

Dual controller configurations offer redundant access to disk storage. In case of controller or I/O path failure, the other controller will continue to provide access to disk drives.

DS3000 RAID controllers have connectors for the following ports:

- Host server attachment ports (SAS, iSCSI, or FC)
- Drive side SAS ports
- Ethernet management port
- Serial management port

As an example, Figure 3-4 shows the DS3200 RAID controller. You can identify all the ports mentioned above.

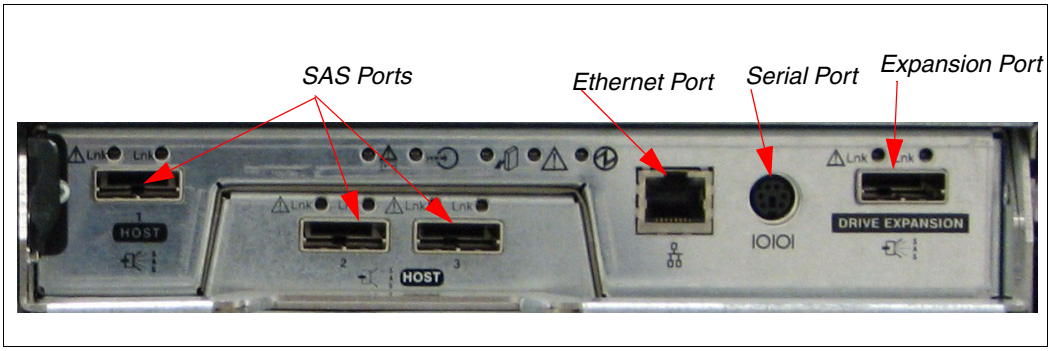


Figure 3-4 DS3200 RAID controller

Note: The serial management port is not meant for general use; it is reserved for advanced troubleshooting activities by trained IBM system service representatives.

Each power supply unit contains two fans. Power and cooling are both redundant; in case of a power supply or fan failure, the DS3000 will remain operational and you can replace the failed component without downtime. However, to ensure sufficient cooling airflow, do not remove the failed component until you are ready to replace it with a new one. If you operate the DS3000 subsystem with certain components removed, the opening will disturb the airflow and decrease cooling efficiency.

We describe the specific features of DS3200, DS3300, and DS3400 in the following sections.

3.2.1 IBM System Storage DS3200

IBM System Storage DS3200 supports SAS attachment of host servers.

You can directly attach up to three host servers to a DS3200 RAID controller. However, the standard version of the DS3200 RAID controller contains only one SAS host port, indicated by the arrow in Figure 3-5.

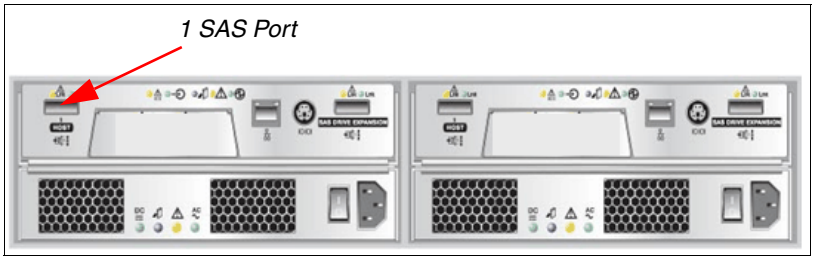


Figure 3-5 DS3200 RAID controller with one SAS host port

The optional DS3200 SAS 2-Port Daughter Card provides two additional host ports. The arrows in Figure 3-6 show all three ports.

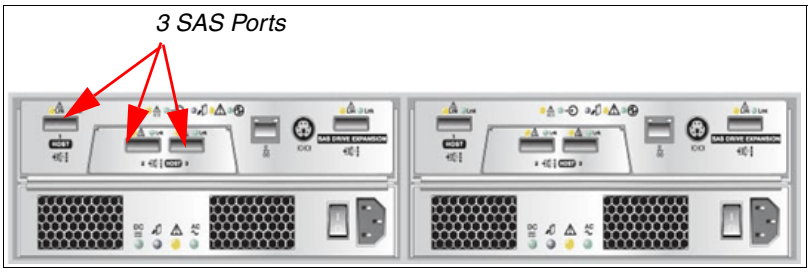


Figure 3-6 DS3200 RAID controller with three SAS host ports

The DS3200 single controller version does not provide redundant host attachment, therefore we recommend using the dual controller setup whenever possible.

3.2.2 IBM System Storage DS3300

The DS3300 offers iSCSI host servers connectivity for IBM System x and IBM BladeCenter servers. Each controller has two 1 Gbps iSCSI host ports (shown in Figure 3-7) that can be used for host server attachment. You would normally use Ethernet switches between the host servers and DS3300 controller iSCSI ports in order to establish Ethernet-based SAN. Up to 32 hosts and 16 maximum host types can be attached.

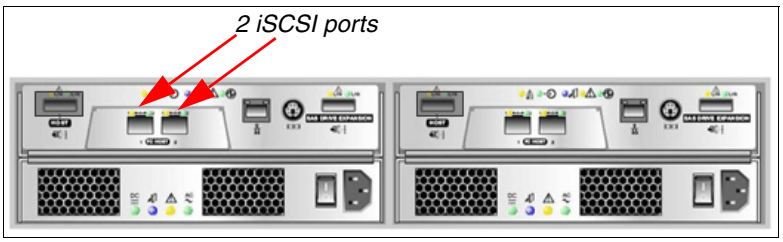


Figure 3-7 DS3300 RAID controller

3.2.3 IBM System Storage DS3400

The DS3400 can have one or two RAID controllers. Each controller has two 4 Gbps capable FC ports for host server attachment (as shown in Figure 3-8). You need to install an SFP module in each FC connector. The FC ports can operate at 2 or 4 Gbps, depending on the SFP module installed.

You can attach up to two host servers directly, but a SAN switch infrastructure is required for more than two servers.

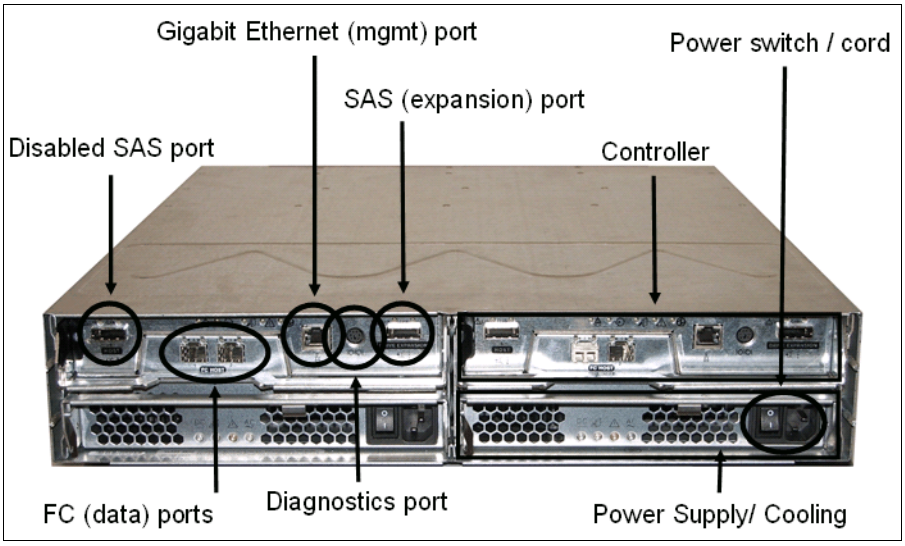


Figure 3-8 DS3400 RAID controller

As with the DS3200 and DS3300, we recommend using the dual RAID controller configuration. This will provide redundant disk storage access to the host servers.

3.2.4 IBM System Storage EXP3000

The EXP3000 is a 2U rack-mountable SAS disk drive enclosure. It supports up to 12 SAS 3.5-inch disk drives. If you use 450 GB disk drives, the maximum storage capacity in a single EXP3000 will be 5.4 TB. You can attach up to three EXP3000 enclosures to a DS3000 storage subsystem, and this gives you a maximum storage capacity of 16.2 TB.

The EXP3000 enclosures connect to the drive side SAS port of the DS3000 through a 3 Gbps SAS interface.

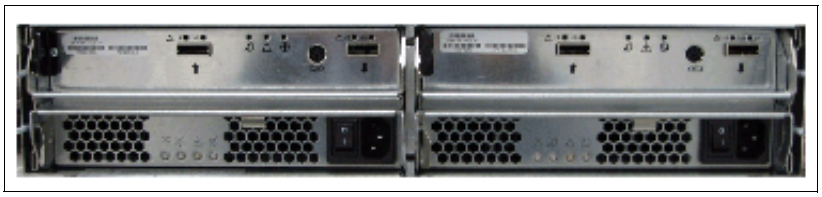


Figure 3-9 EXP3000 rear view

The main components of an EXP3000 are one or two Environmental Services Modules (ESM), two power supply modules, and up to 12 SAS disk drives. ESMs and power supply modules are in the rear of the enclosure (as shown in Figure 3-9), while the SAS drives are installed in the front.

Let us now describe these components:

- Environmental Services Module (ESM)

  An ESM contains two SAS ports labeled *In* and *Out*. Use the SAS In port to attach the EXP3000 to the DS3000 subsystem or to the SAS Out port of another EXP3000. The SAS Out port should be connected to a SAS In port of another EXP3000.

  We show an example of the drive-side SAS cabling in Figure 3-10.

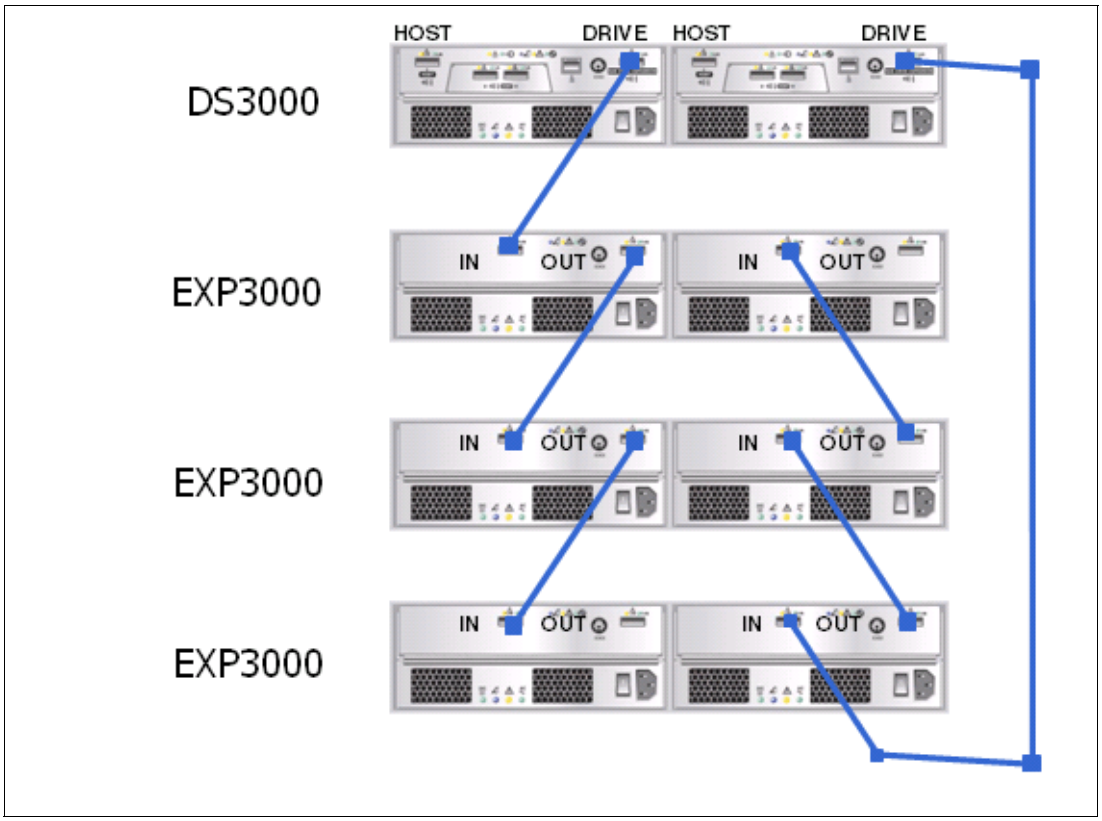


Figure 3-10 SAS connections between the ESMs and the DS3000

  The EXP3000 comes with one ESM in the standard configuration. Another ESM can optionally be installed. Both ESMs must be installed in order to connect the EXP3000 to a dual controller DS3000 storage subsystem.

- Power supply module

  Each EXP3000 is equipped with two hot-swappable and redundant power supply modules. Each power supply module contains two fans for cooling redundancy. If a power supply or a fan fails, you can replace it while the EXP3000 remains fully operational. But remember to not operate the EXP3000 with a power supply unit removed for long periods of time; this will affect cooling efficiency. You should replace the failed power supply module within 10 minutes of removal.

- SAS disk drives

  The EXP3000 can house up to 12 hot-swappable SAS disk drives, installed at the front of the enclosure. Drives of 10000 and 15000 rpm and capacities of up to 450 GB are supported.

3.2.5 Product comparison of DS3200, DS3300, and DS3400

Table 3-2 will help you understand the different features available in the DS family.

Table 3-2 Comparison of DS3200, DS330, and DS3400

DS3200	DS3300	DS3400
Up to three 3 Gbps SAS host ports per controller (host IO interface).	Two 1 Gbps iSCSI host connections per controller (host IO interface).	Two 4 Gbps FC connectors per controller (host IO interface).
Intel X-Scale 80331 (667 MHz) processor.	Intel X-Scale 80331 (667 MHz) processor.	Intel X-Scale 80331 (667 MHz) processor.
SAS back-end interfaces using the LSI1068 SAS IOC.	SAS back-end interfaces using the LSI1068 SAS IOC.	SAS back-end interfaces using the LSI1068 SAS IOC.
512 MB or 1 GB DDR II (per controller).	512 MB or 1 GB DDR II (per controller.	512 MB or 1 GB DDR II (per controller).
Clustering.	Clustering.	Clustering.
Up to 255 volumes.	Up to 255 volumes.	Up to 255 volumes.
Supports Storage Manager and SMcli.	Supports Storage Manager and SMcli.	Supports Storage Manager and SMcli.
Major Event Log (MEL) for error logging.	Major Event Log (MEL) for error logging.	Major Event Log (MEL) for error logging.
HotAdd and Dynamic logical drive Expansion.	HotAdd and Dynamic logical drive Expansion.	HotAdd and Dynamic logical drive Expansion.
RAID Level Migration using SMcli.	RAID Level Migration using SMcli.	RAID Level Migration using SMcli.
Physical disk state indications.	Physical disk state indications.	Physical disk state indications.
Media Scan with or without consistency check.	Media Scan with or without consistency check.	Media Scan with or without consistency check.
Immediate available feature (IAF).	Immediate available feature (IAF).	Immediate available feature (IAF).
Up to three additional EXP3000 enclosures for a total of 48 physical disks.	Up to three additional EXP3000 enclosures for a total of 48 physical disks.	Up to three additional EXP3000 enclosures for a total of 48 physical disks.
Cache Battery Backup - 3 Cell (18650) Li-ion Battery Pack.	Cache Battery Backup - 3 Cell (18650) Li-ion Battery Pack.	Cache Battery Backup - 3 Cell (18650) Li-ion Battery Pack.
RAID controller firmware: 07.35.xx.xx. Note that when updating firmware, the Storage Manager software may also require an update. Consult the compatibility matrix for information. If an update is required, always update the Storage Manager software first.	RAID controller firmware: 07.35.XX.XX.	RAID controller firmware: 07.35.XX.XX.

3.3 Premium Features

Standard configurations of DS3000 storage subsystems come with the following Premium Features.

Note: The DS3200 Software Feature Pack, DS3300 Software Feature Pack, DS3400 Software Feature Pack, and the EXP3000 Expansion License have been incorporated into the base models and are no longer needed as prerequisites to the other advanced function options. The required controller firmware that enables this function must be at firmware level V06.70.23.00 or later and is available on the DS3000 product support Web site at:

http://www-03.ibm.com/systems/storage/disk/entry/index.html

You can also refer to the IBM System Storage DS3000 series Interoperability Matrix, found at the following link:

http://www-03.ibm.com/systems/resources/systems_storage_disk_ds3000_pdf_interop.pdf

DS3000 Partition Expansion License

As part of the standard configuration, four storage partitions are enabled on the DS3200, DS3300, and DS3400 with Windows and Linux on Intel host attach license (this can be increased to 8, 16, or 32). The maximum number of storage partitions is 32 on all DS3000 products. Use the DS3000 Partition Expansion License to enable all 32 storage partitions.

DS3000 FlashCopy Expansion License

This feature enables FlashCopy. FlashCopy is a point-in-time copy of a *source logical drive*. The FlashCopy logical drive becomes available almost instantaneously.

FlashCopy requires the use of a defined *FlashCopy repository*, which will contain the original content of the data that has since been altered. FlashCopy logical drives are often used as a source for a backup operation. They can also be used to simply and quickly roll back to an original data state, thus providing a restore point. However, if the source logical drive is lost, the point-in-time FlashCopy will be lost as well. For more information about FlashCopy, see 9.4, “Advanced functions - FlashCopy” on page 183.

As part of the standard configuration, two FlashCopies are enabled on every DS3000 storage subsystem and this Premium Feature enables up to 64 FlashCopies.

DS3000 VolumeCopy License

VolumeCopy is a way to provide a complete point-in-time copy of a source logical drive. As opposed to FlashCopy (where only the original values of changed data are copied to the repository), the whole source logical drive is copied to target. You can use this functionality for data replication, relocation, backup, or to restore snapshot data to the original logical drive. The time required to establish a copy will depend on the size of the source data and the operation priority settings. While establishing the copy, the source logical drive will be in read-only state.

Once all the data is copied to the target, the target will remain available if the source logical drive is lost. For more information about VolumeCopy, see 9.5, “Advanced functions - VolumeCopy” on page 196.

The VolumeCopy Premium Feature allows for up to 128 VolumeCopies. Be aware that FlashCopy is a prerequisite for VolumeCopy.

DS3000 FlashCopy VolumeCopy License

As stated above, the FlashCopy Premium Feature must be enabled before you can enable and use VolumeCopy. For this reason, IBM provides the FlashCopy VolumeCopy license; this is actually a bundle of both Premium Features.

DS3000 EXP3000 Expansion License

The Expansion License is enabled by default for one to three EXP3000 expansion units.

3.4 DS3000 Storage Manager

The DS3000 Storage Manager is used to manage the DS3200, DS3300, and DS3400 storage subsystems. Along with your Storage box, you might get DS Storage Manager Software and Host Kit CDs; otherwise, you can download it from the IBM support Web site. Using DS3000 Storage Manager software, you can perform tasks such as creating arrays and logical drives, assigning logical drives to the host servers, setting up FlashCopy and VolumeCopy, capturing logs for troubleshooting, and so on. When discussing the DS3000 Storage Manager, it is important to differentiate between the following two terms:

- Host server

  This is a server attached to the DS3000 storage subsystem through the I/O path (SAS, iSCSI, or Fibre Channel). The host server has access to certain logical drives defined on the DS3000.

- Management station

  The management station is a system used for DS3000 management. The DS3000 Storage Manager GUI runs on the management station. You need to establish a management connection between the management station and the DS3000 storage subsystem. This can be done in two ways:

  – Out-of-band

    When using out-of-band management, the management station is connected to the Ethernet management port in each DS3000 RAID controller. All management communication flows across the TCP/IP connection between the management station and the DS3000. We also call this method *direct-attached management*. The management station in this case only requires an Ethernet connection to the DS3000.

  – In-band

    This method utilizes the I/O path between a host server and the DS3000. In this case, the management station does not have direct TCP/IP connection to the DS3000, but rather communicates with the DS3000 through an HBA, which acts as a gateway to the DS3000 storage subsystem, that is, communication between the management station and the host server is across the FC/SAS I/O path.

    We also call this method *host-attached management*.

We discuss both management methods in more detail in 3.4.2, “Management methods” on page 30.

3.4.1 DS3000 Storage Manager components

The DS3000 Storage Manager package consists of the following components:

- Storage Manager Client (SMclient/SMcli)
- Storage Manager Agent (SMagent)
- Storage Manager Utility (SMutil)
- Storage Manager multipath support
- Java™ access bridge (for Windows only)

In the following sections, we provide more details on these components.

Storage Manager Client (SMclient)

This is the actual graphical user interface (GUI) that you use to manage the DS3000 subsystems. It has two distinct parts:

- Enterprise Management Window

  This window opens when the DS3000 Storage Manager is launched. It lists the storage subsystems it knows about. You can add new storage subsystems and perform various tasks on the enterprise level. We show a sample Enterprise Management Window in Figure 3-11.

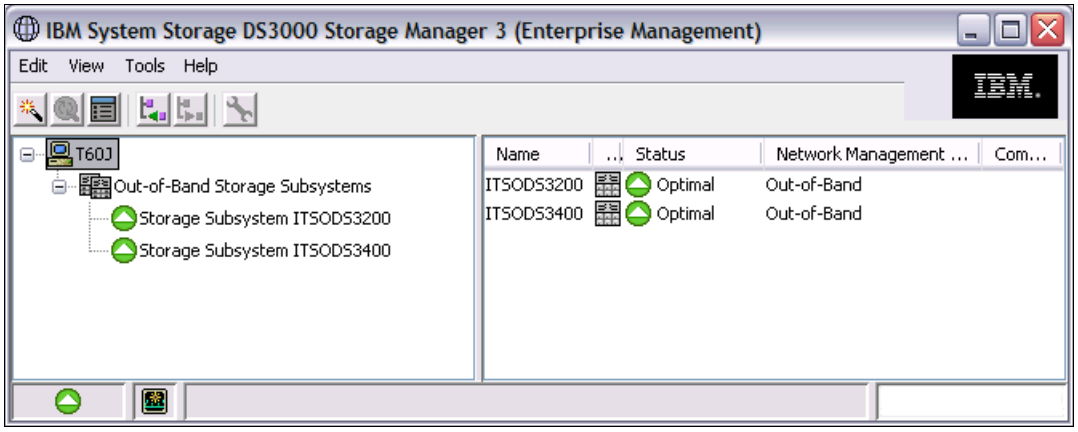


Figure 3-11 Enterprise Management Window

- Subsystem Management Window

  This allows you to manage a particular DS3000 Storage Subsystem. Management tasks such as creating arrays, logical drives, storage partitions, FlashCopy, and VolumeCopy are all performed from within the Subsystem Management Window. See Figure 3-12.

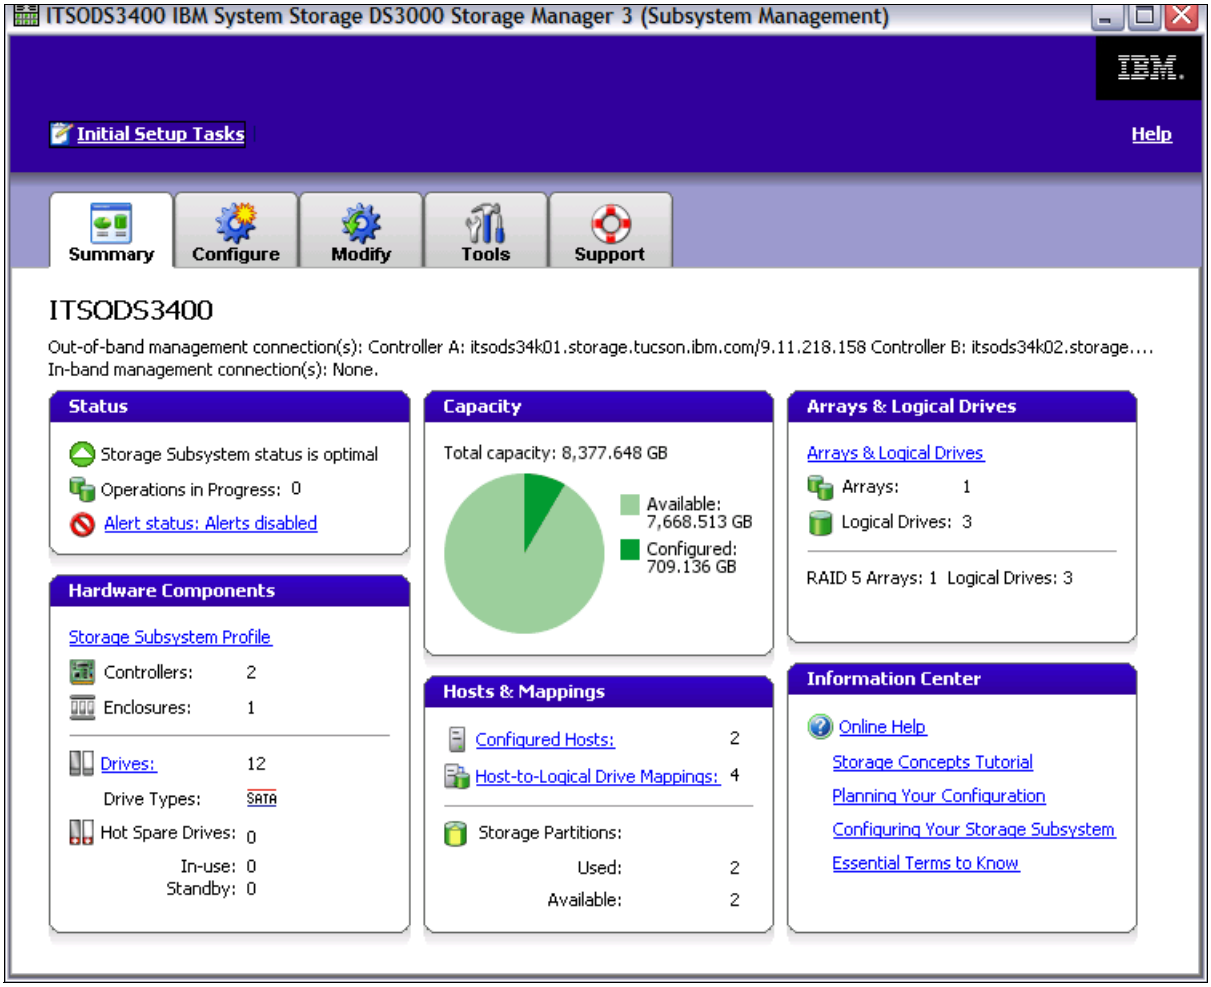


Figure 3-12 Subsystem Management window

You install SMclient on the management station. This is usually not one of the host servers, but rather a workstation that belongs to a system administrator. However, it is also possible to install SMclient on a host server. Figure 3-12 shows a sample Subsystem Management Window.

SMclient is available for Windows and Linux operating systems.

SMcli command-line interface

Besides providing a GUI for the management tasks, SMclient also includes a component called SMcli. SMcli provides a powerful command-line interface (CLI). All the tasks available in the Storage Manager GUI can also be run using the CLI. In addition, there are certain tasks that are only available in the CLI, but not in the GUI. There are two ways to use the CLI:

- The SMcli executable, which runs the CLI commands from the operating system command prompt.
- The Script Editor in the DS3000 Storage Manager, launched from the Enterprise Management Window.

You can either run a single command at a time or execute pre-written scripts. See Appendix C, “CLI” on page 747 for more information about the CLI and Script Editor.

Storage Manager Agent (SMagent)

SMagent is an optional component that is only required for in-band management. SMagent is installed in the host server and allows the SMclient to communicate with DS3000 across the I/O path. At the time of writing, only the FC I/O path is supported for management communication, so only the DS3400 can be managed in-band.

Storage Manager Utility (SMutil)

The Storage Manager Utility package contains the following components:

- Hot-add utility

  You can use this utility to dynamically add newly created logical drives to the operating system running in the host without having to reboot the host server. This utility is available for Windows only.

- SMdevices utility

  This utility can be used to associate logical drives to device names in the operating system. It is installed with both the Windows and Linux package.

- SMrepassist utility

  SMrepassist is a Windows utility used with FlashCopy and VolumeCopy. The utility allows you to flush cached data prior to creating a FlashCopy/VolumeCopy image. In other operating systems, you need to unmount the file system. SMrepassist does not exist in the Linux package.

SMutil is a required package in the host server.

Multipath support

We recommend installing two HBAs in the host servers and using the dual controller DS3000 subsystems. This provides higher availability through I/O path and controller redundancy. However, dual path configurations will only work correctly if you install the appropriate multipath driver for your particular operating system in the host server.

DS3000 Storage Manager for Windows includes multipath support, based on the Microsoft MPIO framework. This component is actually called *RDAC* in the DS3000 Storage Manager installation utility, but be careful not to confuse this with the RDAC as we know it in the DS4000 context. The DS4000 RDAC is a proprietary multipath solution, while the DS3000 “RDAC” conforms to the Microsoft MPIO specification.

DS3000 Storage Manager for Linux does not include any multipath support. RDAC for Linux is available as a separate open source package called MPP.

Java Access Bridge

This is included in the Windows Storage Manager package only. Java Access Bridge for Microsoft Windows makes it possible for Windows based assistive technology to access and interact with the application.

3.4.2 Management methods

In this section, we discuss the two management methods (in-band and out-of-band) in more detail, to help you select which is more appropriate for your environment, because each method has some associated advantages as well as disadvantages. Both methods offer identical functionality: You can perform any management task with either of these methods.

In-band management

The in-band management method uses the I/O path between the host server and the DS3000 to transfer management commands and information.

This method does not use the management Ethernet ports on DS3000 RAID controllers and does not require a management TCP/IP network. However, it does require a special *access logical drive* to manage the DS3000 controllers. This means that you cannot configure the maximum number of logical drives, because one of them is reserved for the access logical drive. But this is usually not a problem, because virtually all customers will find the maximum number of logical drives more than sufficient.

An example of in-band management is shown in Figure 3-13. Two host servers are attached to the DS3400 subsystem with FC cables. They both run SMagent code. The management workstation runs SMclient code. SMclient communicates with SMagent through Ethernet, and SMagent communicates with the DS3000 across the FC I/O path.

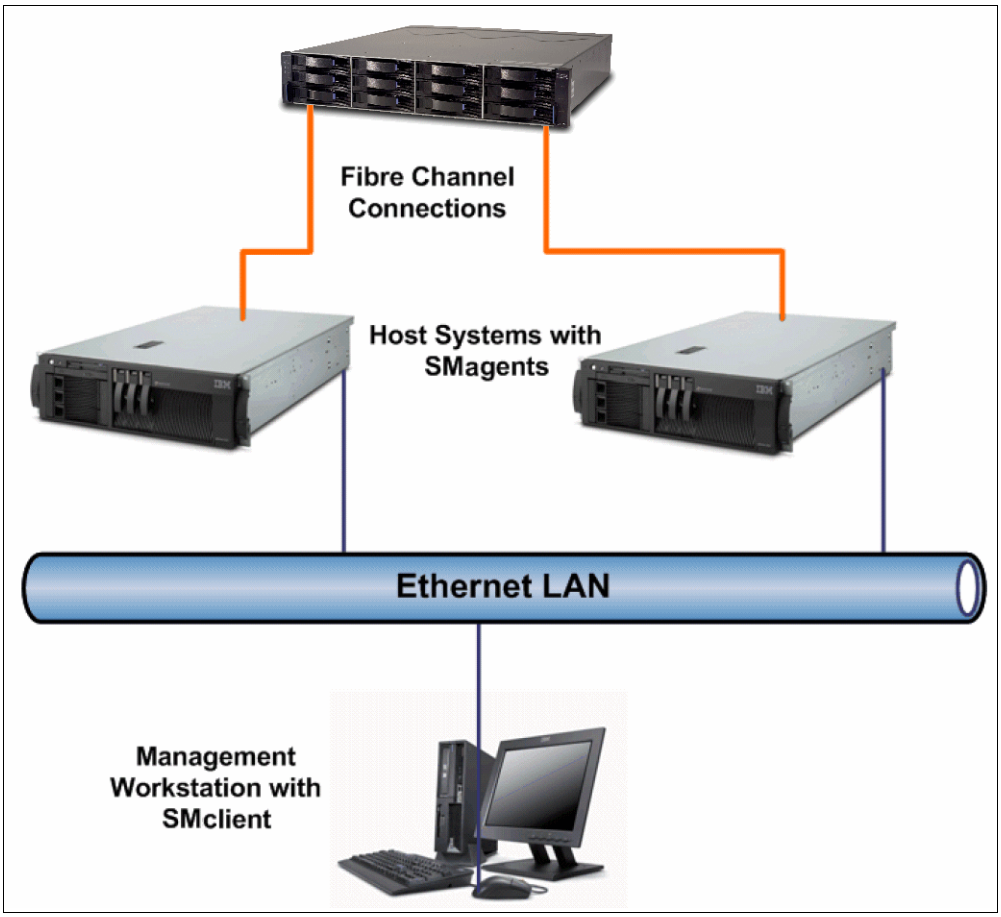


Figure 3-13 In-band management

Access logical drive

The access logical drive exists in each storage partition by default (no manual configuration is required). This is not actually a real logical drive, although it is presented to the operating system as a drive on LUN 31. In order to allow for in-band management communication across the I/O path, we need a target device for SCSI commands. The SCSI commands to this target are used as a vehicle for Storage Management communication along the I/O path. The access logical drive is that target device. The access logical drive is also sometimes referred to as the *Universal Transport Mechanism (UTM)* device or the Universal XPort device.

Out-of-band management

Out-of-band management requires that the management IP addresses are configured on both controllers and that the controllers' management ports are connected to the management network. This should be a separate LAN or a VLAN, as we do not recommend using the production LAN or VLAN for the management network traffic.

A separate management workstation is another requirement; typically, the system administrator uses their own workstation for this purpose. Figure 3-14 shows the management workstation and the DS3000 subsystem, connected on the Ethernet management network.

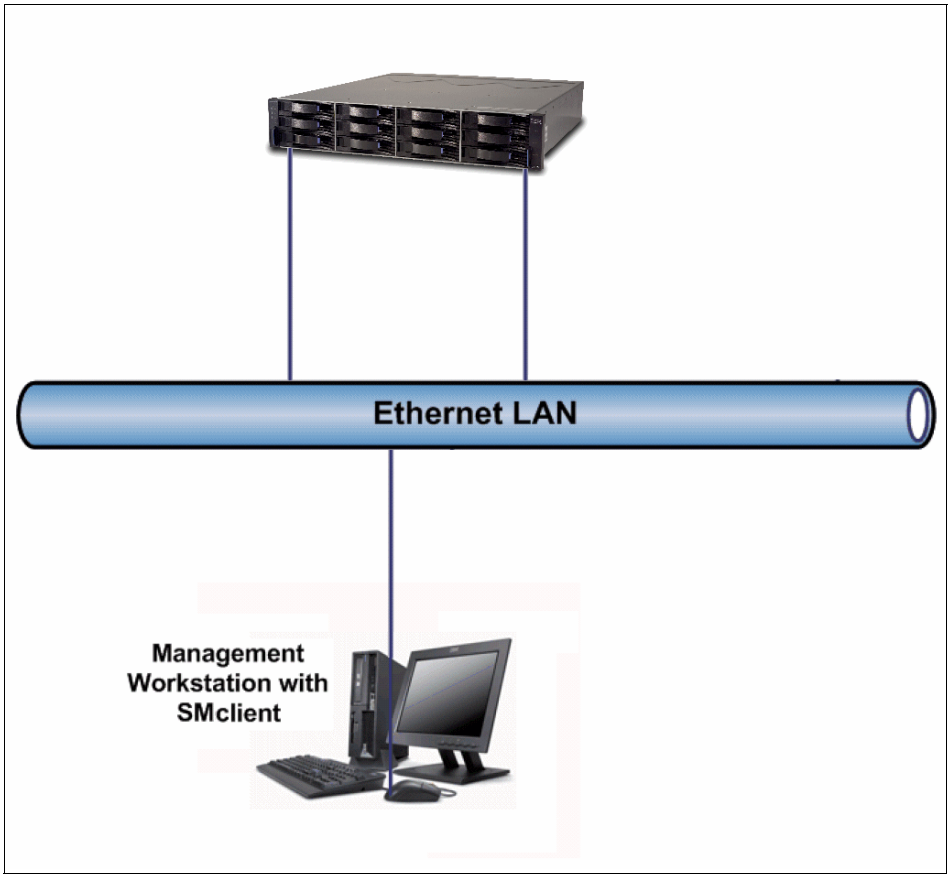


Figure 3-14 Out-of-band management

Out-of-band management offers the following advantages:

- There is no need for an access logical drive (unlike for in-band management). Therefore, you can use the maximum number of logical drives supported by the host servers' operating system.
- If the I/O paths fail, you can still access the DS3000 storage subsystem out-of-band, check the status, and capture logs for effective troubleshooting. Access to both controllers is required for almost all in-band management functions.
- If in-band management cannot be used (for example, when the SMagent is not available for the host server operating system), you can effectively use out-of-band management.

We recommend setting up and using both methods, if at all possible. This will introduce some redundancy to your management setup and provide management access to the DS3000 subsystem even in the case of I/O path or Ethernet management network failure.

Part 2

Getting started

In this part of the book, we show how to get started using the IBM System Storage DS3000.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

4

Cabling

In this chapter, we explain how to properly cable the DS3000 subsystems. We discuss the following cabling areas:

- SAS cabling between the DS3000 and EXP3000 enclosures
- Ethernet management cabling
- Host server attachment (SAS, iSCSI, and Fibre Channel)
- Power cabling

We also describe the correct power-on procedures for the IBM System Storage DS3000.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

4.1 Enclosure ID settings

The DS3000 storage subsystem can be attached to up to three EXP3000 enclosures. The controller automatically sets the enclosure ID number. You can change the setting through the storage management software, if necessary. For proper operation, you have to verify or set a unique enclosure ID for the DS3000 and for each of the EXP3000s. This is very important, because the enclosure ID, along with the drive's position within the enclosure, is used to determine the hardware address of each hard disk drive. If two enclosures are configured with the same enclosure ID, the disk drives will have conflicting hardware addresses.

If the DS3000 is used standalone and is not connected to any EXP3000s, then the enclosure ID setting does not really matter. But if any EXP3000s will ever be attached, make sure you verify the ID settings on each enclosure.

The factory default enclosure ID of the DS3000 controller is 85. We recommend that you leave this as it is. If you do decide to change it, we recommend using values of 80 or higher, since some storage subsystems may not support values lower than 80.

In a dual controller DS3000 subsystem, both controllers should have the same enclosure ID. Similarly, on an EXP3000 with two ESMs, both ESMs must have the same enclosure ID.

You can view and change the enclosure IDs with the DS3000 Storage Manager or using the CLI. Valid enclosure ID values are 0 to 99.

See 12.1.5, "Change Enclosure ID Numbers" on page 272 for information about how to change the enclosure ID. On the CLI, use the following command:

```
set enclosure [ serial-number ] id = enclosure-id
```

4.2 SAS cabling

All DS3000 storage subsystems, regardless of the type of host attachment, use SAS disk drives/SATA Disk Drives and SAS cabling to attach expansion enclosures. Therefore, let us have a look at the SAS cables and connectors used.

The SAS ports on the DS3000 controller and EXP3000 ESM all support mini-SAS 4x multilane connectors. SAS cables with mini-SAS connectors that will fit in these ports are required, as shown in Figure 4-1 on page 37. IBM provides SAS cables in two cable lengths: 1 and 3 meters.

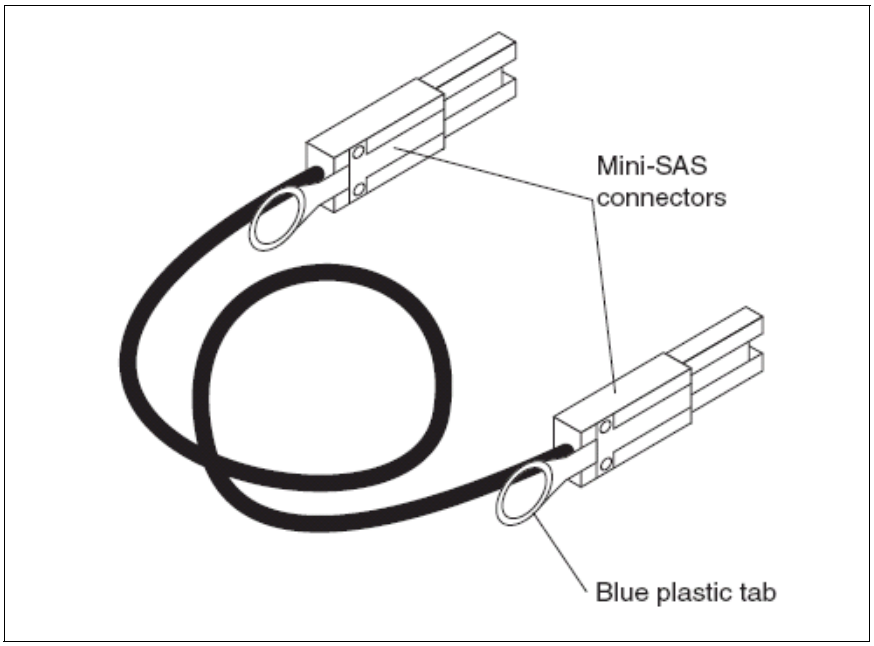


Figure 4-1 SAS cable

To avoid damage to the SAS cables, consider the following precautions:

- When you route the cable along a folding cable-management arm, leave enough slack in the cable.
- Route the cable away from places where it can be damaged by other devices in the rack cabinet.
- Do not put excess weight on the cable at the connection point. Make sure that the cable is well supported.

To connect a mini-SAS cable, insert the mini-SAS connector into a mini-SAS port. Make sure that it locks into place.

To remove a mini-SAS cable, complete the following steps:

1. Put one finger into the hole on the blue plastic tab on the mini-SAS connector and gently pull on the tab to release the locking mechanism.
2. As you pull on the tab, pull out the connector to remove it from the port.

4.3 EXP3000 attachment

Here are some general rules to follow when attaching EXP3000 enclosures to a DS3000:

- Connect the drive-side SAS port on the DS3000 controller to the SAS In port on the EXP3000.
- Connect the SAS Out Port on the EXP3000 to the SAS In port on the next EXP3000.

Single controller DS3000

Figure 4-2 shows a single controller DS3000 subsystem and three EXP3000 enclosures, with a single ESM installed in each EXP3000. Following the rules above, connect the SAS ports on the DS3000 and the EXP3000s as shown.

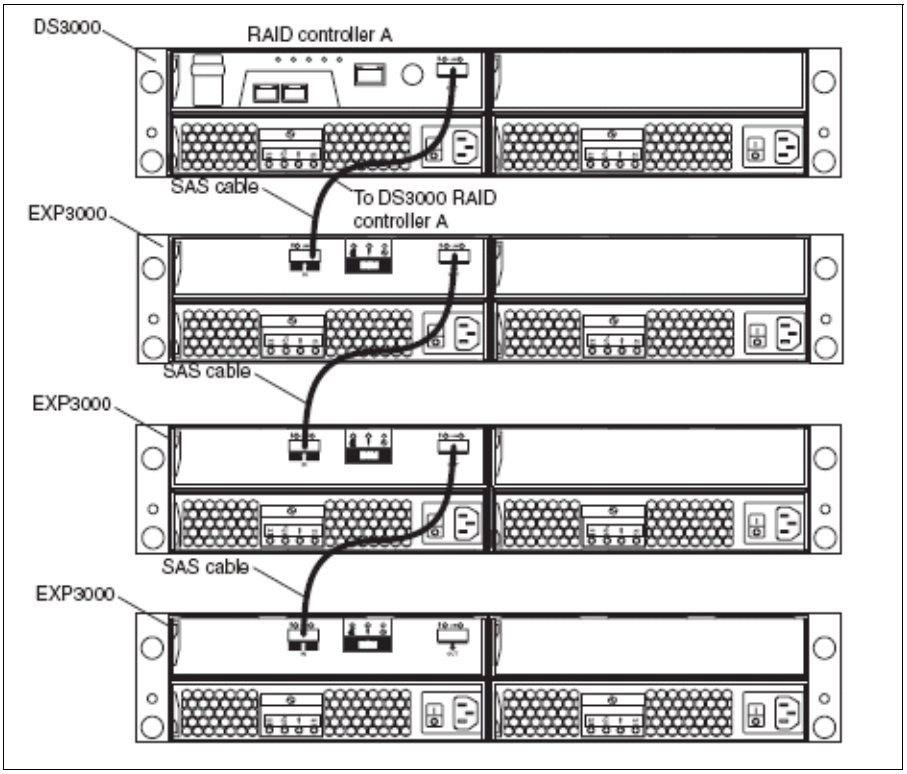


Figure 4-2 EXP3000 attachment - single controller and single ESM configuration

This is a fully functional configuration, but there is no redundancy on the drive side. If there is an EXP3000 failure or a broken SAS cable, some EXP3000s could lose connection unnecessarily.

Dual controller DS3000

We therefore recommend that you use a dual controller DS3000 and two ESMs in each EXP3000. Figure 4-3 on page 39 shows the SAS cable connections in this scenario.

Use top-to-bottom order on the left-side controller and EXP3000 ESMs. The right-side controller and EXP3000 ESMs should be connected in bottom-to-top order. This will eliminate a single point of failure on the drive side. If any EXP3000 fails completely, the remaining EXP3000s will still be accessible through one or the other DS3000 controller.

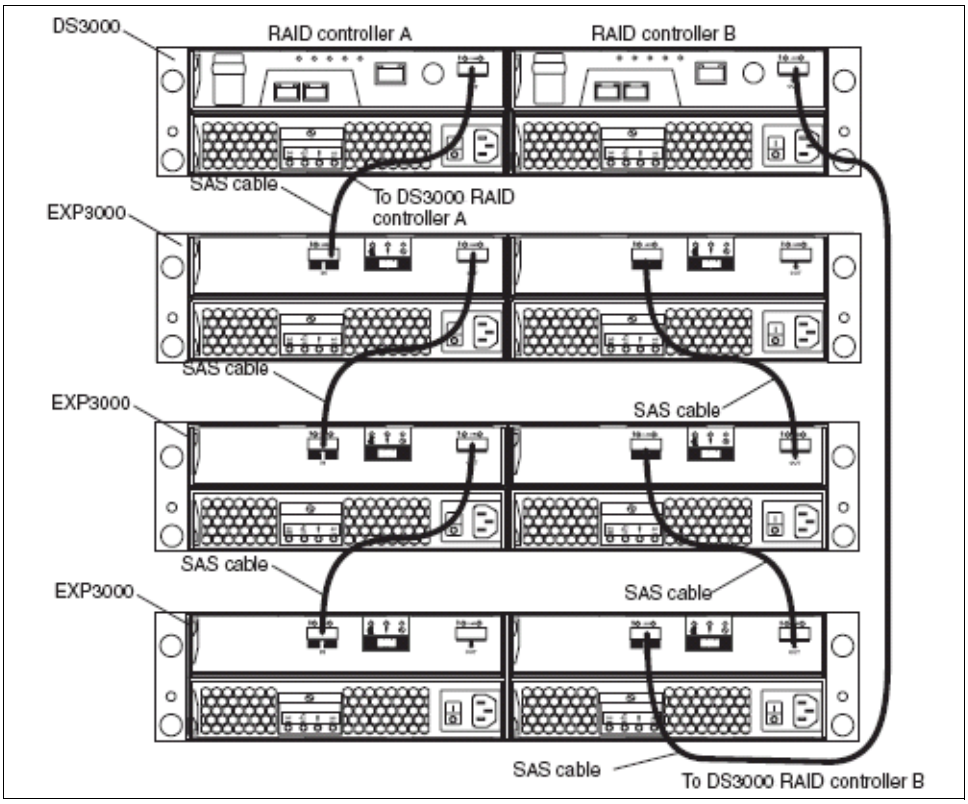


Figure 4-3 EXP3000 attachment

If the SAS cabling on the drive side is incorrect, it will be detected by the DS3000 controller firmware. The DS3000 Storage Manager will alert you about the condition by logging a critical Major Event Log event. In addition, Recovery Guru will point to the miswire condition and advise you of the steps to correct the problem. See 14.6, “Recover from Failure” on page 347 for more information about failure recovery.

Keep in mind that a miswire condition will only be reported if the wrong SAS connections will result in a non-working configuration. It is also possible to attach the EXP3000 enclosures in a technically correct manner, which might not be optimal. Such configurations will not produce a miswire event. Therefore, you should carefully follow the recommendations here for best results.

Here are a few examples of miswired configurations:

- EXP3000 enclosures are connected to SAS host ports instead of to SAS drive ports on the DS3000 controller.
- The SAS In port is connected to the SAS In port on another EXP3000.
- The SAS Out port is connected to the SAS Out port on another EXP3000.
- SAS ports are interconnected within a single EXP3000.

An example of a non-optimal, but technically correct configuration would be to connect both left and right side ESMs in top-to-bottom order. While this configuration does work and does not result in a miswire condition, it does not provide redundancy. For example, if an entire EXP3000 fails, the EXP3000 enclosures beneath it will all lose access to the DS3000 subsystem.

4.4 Ethernet management

Each DS3000 RAID controller contains an Ethernet management port, which you can use for out-of-band management. If you have a dual controller DS3000 subsystem, make sure the management workstation can access the management port on each controller. If only one controller is assessable by the management machine, the DS3000 Storage Manager will not be able to manage the enclosure.

Do not use your public LAN for DS3000 out-of-band management. Instead, we strongly recommend setting up a dedicated LAN or VLAN just for management purpose. This will provide increased security of your DS3000 subsystem. If the DS3000 RAID controllers are on a public LAN, a knowledgeable user could install the DS3000 Storage Manager on a workstation, or use the CLI to run potentially destructive tasks. For an additional layer of security, we also recommend that you enable password protection on the DS3000 subsystem (refer to 12.1.2, “Set or Change Password” on page 260).

4.5 Host attachment

In this section, we describe the host side connections between the DS3000 storage subsystem and host servers. As the DS3000 product range supports different host attachment technologies, we now take a look at each of these possibilities.

DS3200 — SAS

Host server SAS HBAs are used to connect to the DS3200 SAS ports. The DS3200 RAID controller can have up to three SAS host ports, so you can connect up to three host servers, as shown in Figure 4-4.

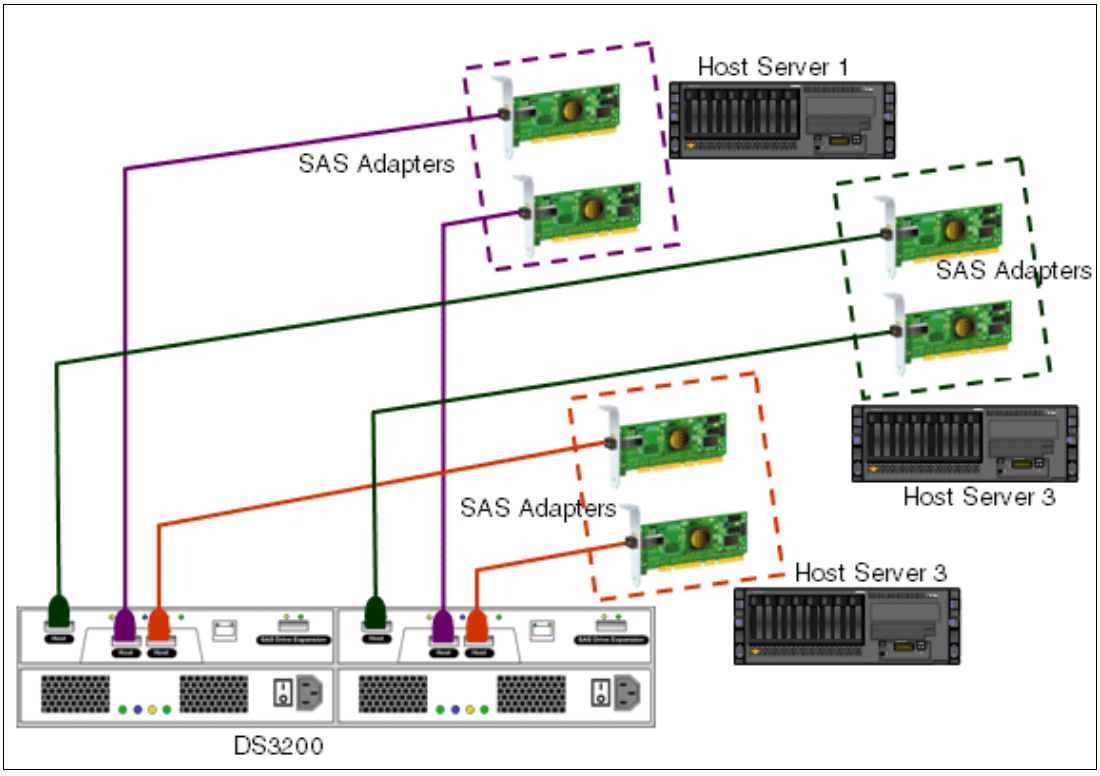


Figure 4-4 Dual controller DS3200 with three host servers

In this scenario, each host server has two HBAs and each HBA is connected to a different controller. This provides I/O path and controller redundancy.

If the host servers only have one HBA, it is possible to attach up to six hosts to a dual controller DS3200 subsystem. However, such a configuration is not recommended, as there is no redundancy for the I/O path or controller. If an HBA, SAS cable, or RAID controller fails, certain hosts will inevitably lose access to logical drives. With two HBAs per host for redundancy, then three individual servers can be attached to a dual controller DS3200.

SAS switches are coming onto the market, which, when available, will allow connection of a higher number of host servers.

DS3300 — iSCSI

The DS3300 supports iSCSI SAN, which utilizes the standard Ethernet infrastructure, using regular Ethernet cables on the host side. The simplest case is a direct connection to a single host server. It is more typical to attach multiple host servers through an Ethernet switch. See Figure 4-5 for an example.

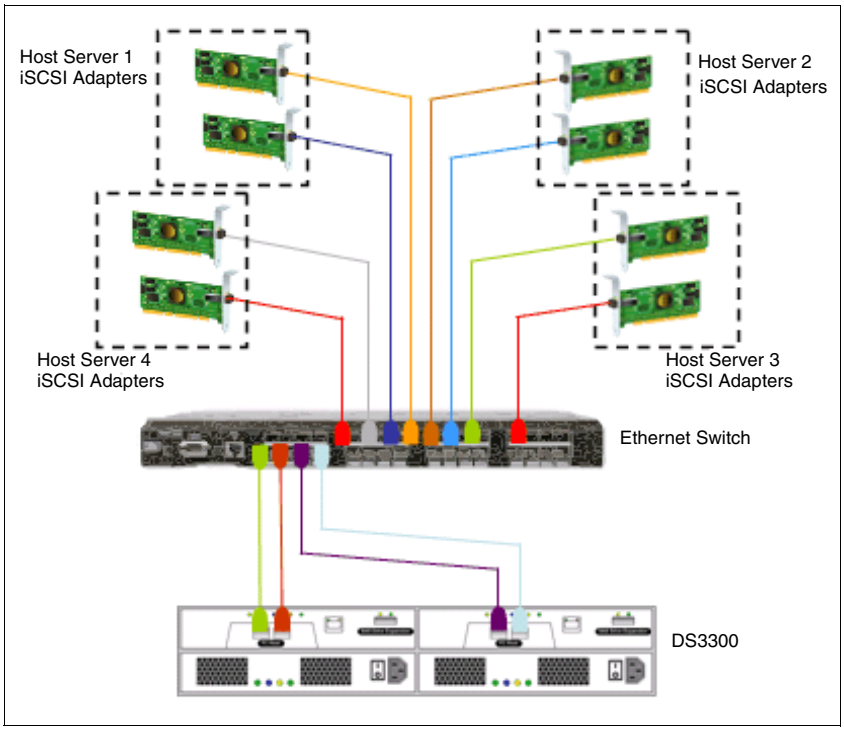


Figure 4-5 Four host servers attached to a DS3300 subsystem

DS3400 — Fibre Channel

The DS3400 RAID controller has two FC host ports; you can therefore connect up to two host servers directly. For a larger number of host servers, you have to use an FC SAN infrastructure.

You need to install an SFP module in each FC host port you intend to use. The SFP modules convert electrical signals from the FC port into optical signals suitable for transmission over FC media and vice versa. Plug one end of the FC cable into the SFP module and the other end into the host server HBA or a SAN switch port.

The SFP modules are keyed in order to prevent incorrect insertion. Do not use force when inserting them, as this could damage the SFP module and the FC port. When inserted correctly, the locking mechanism will lock the module in place.

When removing an SFP module, make sure that you correctly release the locking mechanism. SFP modules can have either a plastic or wire tab, which is used to unlock the latch and allow for easy removal. Before trying to remove the SFP module, make sure that you disconnect the FC cable.

Even though the SFP modules can be either long wave or short wave, the DS3400 only supports short wave SFP modules. Therefore, make sure that you use only supported SFPs.

Use an LC-LC FC cable to connect the SFP module to a SAN switch port or to an FC HBA. Figure 4-6 shows such a cable.

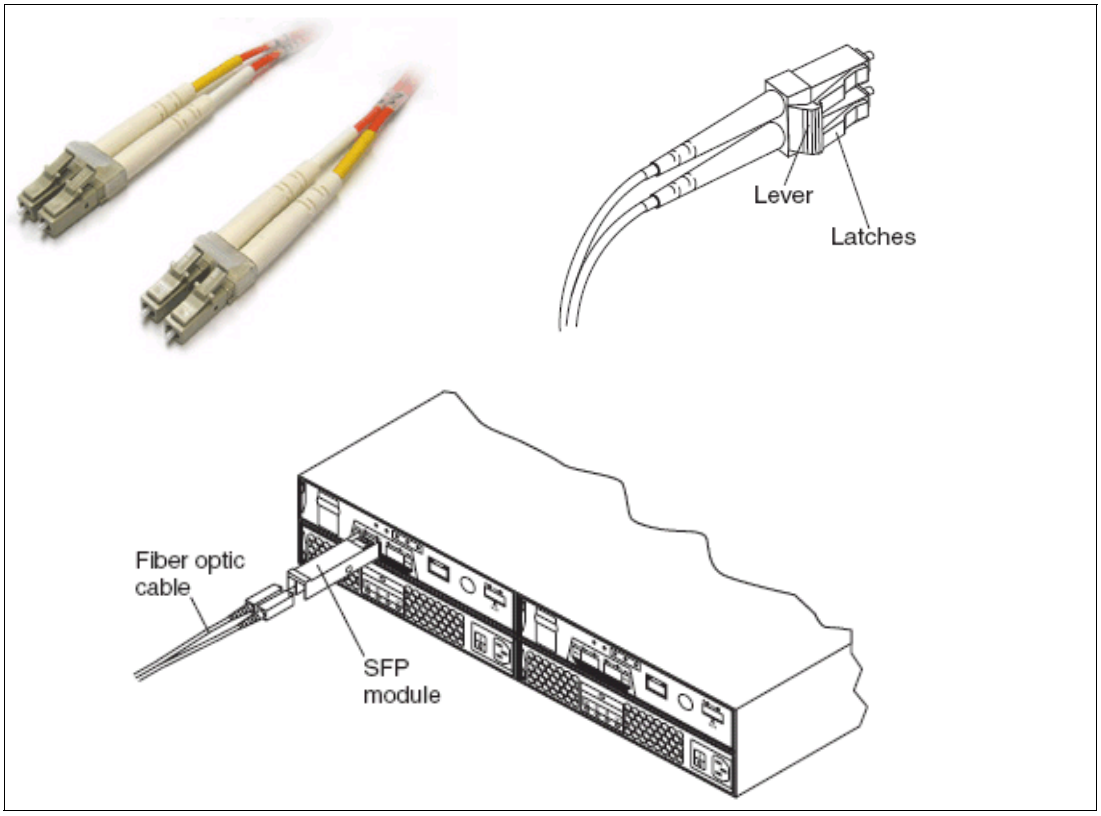


Figure 4-6 LC-LC FC cable

Direct-attached DS3400

Figure 4-7 shows a simple direct-attached configuration. Two host servers are connected to a dual-controller DS3400. Each server uses two FC HBAs, so this is a fully redundant setup. If an HBA, FC cable, or RAID controller fails, the host servers will still have access to logical drives. This type of setup is suitable for a two-node Microsoft Cluster Server configuration.

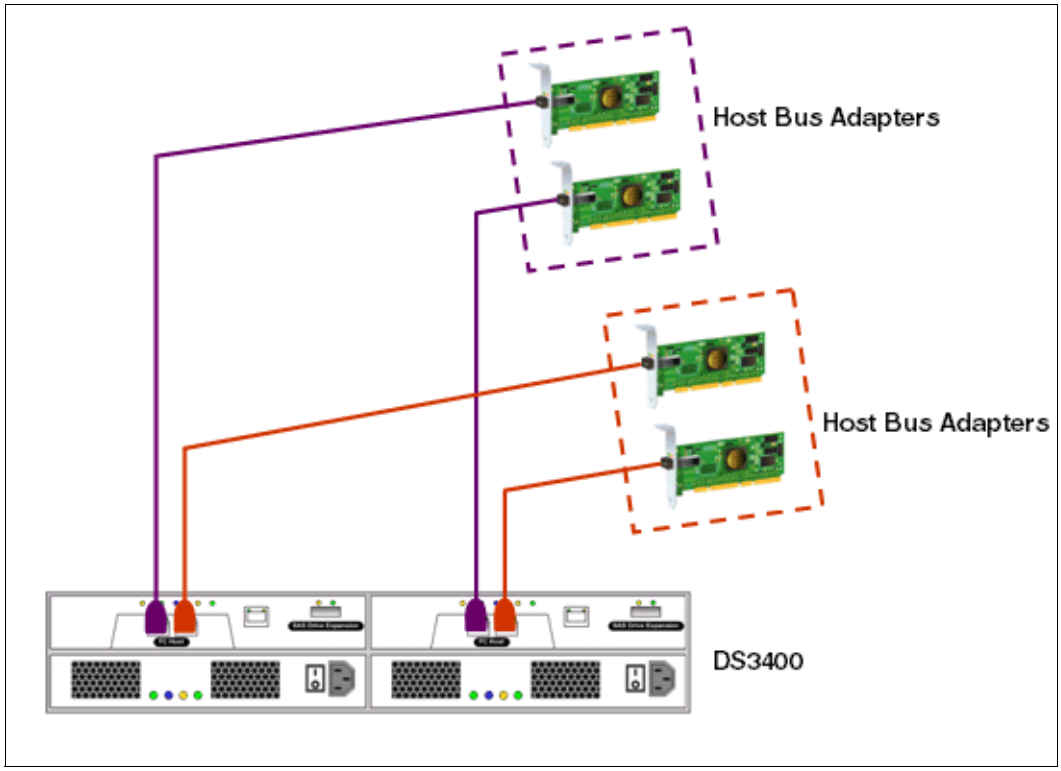


Figure 4-7 Two host servers attached to dual controller DS3400

Switch-attached DS3400

As the DS3400 subsystem can support many more than two host servers, let us consider a larger configuration. For more than two servers, a SAN switch is required. Figure 4-8 displays a sample configuration with four host servers attached to a dual controller DS3400. Each host has two FC HBAs for redundant I/O path support.

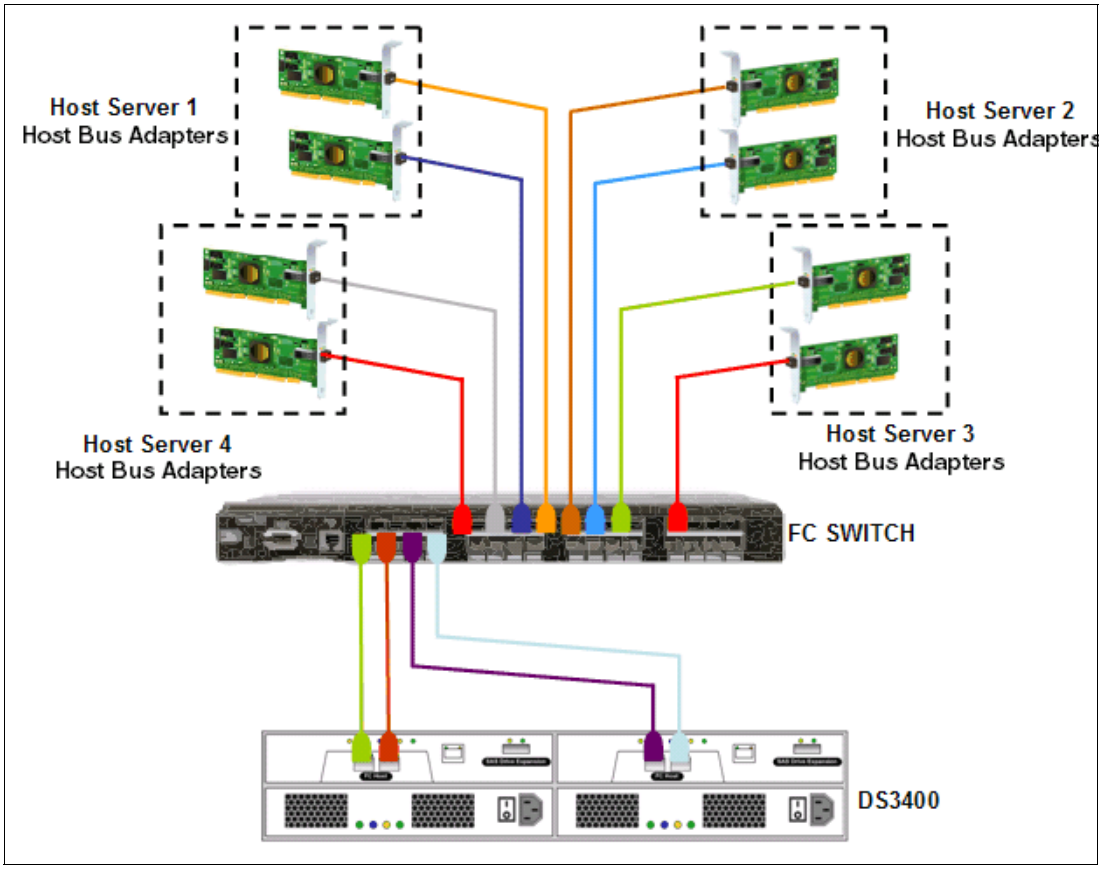


Figure 4-8 Four host servers attached to a dual controller DS3400

The SAN switch must be zoned so that each HBA in a particular server is connected to a different DS3400 RAID controller. For maximum availability, you can add another SAN switch, which eliminates a single point of failure on the entire I/O path. We show this in Figure 4-9.

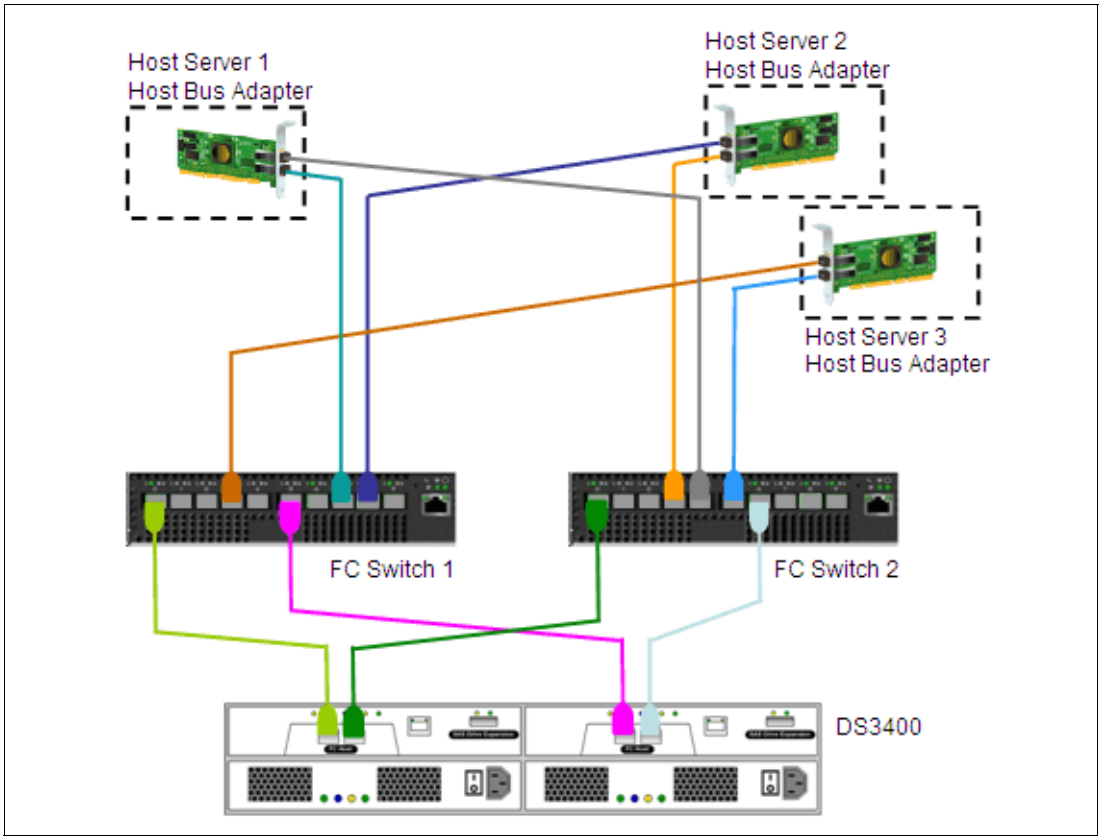


Figure 4-9 Redundant SAN switches

4.6 Power cabling

Connecting the power cables to DS3000 subsystems and EXP3000 enclosures is a simple task. Each of the two power supplies has a power cord connector. You simply have to connect each power supply to a power source. For increased protection in case of power loss, consider connecting each power supply to a different power source.

Attention: Make sure the power switch on each power supply is in the *Off* position before attaching the power cable.

4.7 DS3000 Front and Power LEDs

This section describes the primary LEDs on the DS3000 Front and power supplies. These LEDs are shown in Figure 4-10.

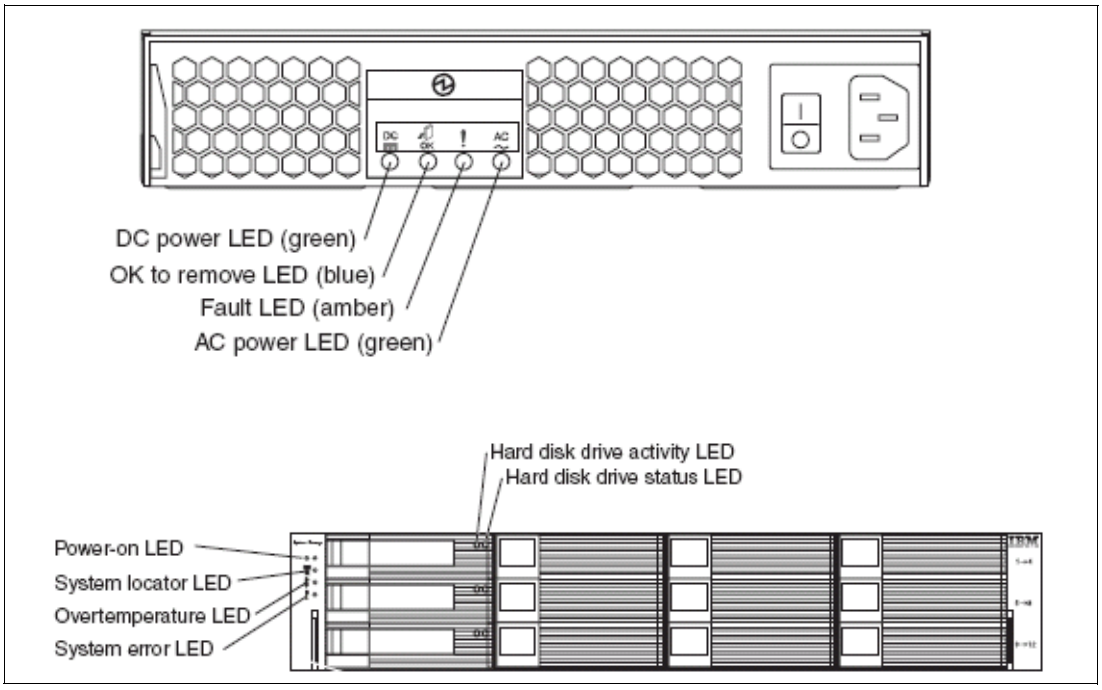


Figure 4-10 DS3000 Front and Power Supply LEDs

Where:

- Power Supply LEDs
  - DC power LED (green)

    When this green LED is lit, it indicates that the DS3000 is turned on and is supplying both 5-volt and 12-volt DC power to the DS3300.
  - OK to remove LED (blue)

    When this blue LED is lit, it indicates that it is safe to remove the power-supply unit.
  - Fault LED (amber)

    When this amber LED is lit, it indicates that a power supply or fan has failed or that a redundant power supply is not turned on.
  - AC power LED (green)

    When this green LED is lit, it indicates that the DS3300 is receiving AC power.
- Front Panel LEDs
  - Power-on LED (green)

    When this green LED is lit, it indicates that the power supply is turned on and is supplying both 5-volt and 12-volt DC power to the DS3300.
  - System locator LED (blue)

    This blue LED can be lit by the DS3000 Storage Manager software to aid in visually locating the DS3300.

– Over temperature LED (amber)

  When this amber LED is lit, it indicates that the DS3300 is in an over temperature condition.

– System error LED (amber)

  When this amber LED is lit, it indicates that the unit has a fault, such as in a power supply, controller, or hard disk drive.

– Hard disk drive activity LED (green)

  Each hard disk drive has an activity LED. When this green LED is flashing, it indicates drive activity.

– Hard disk drive status LED (amber)

  Each hard disk drive has a status LED. When this amber LED is lit continuously, it indicates a drive failure. When it is flashing, it indicates that a drive identify or rebuild activity is in progress.

Note: For controller LED and other hardware information, refer to the appropriate System Storage DS3000 Storage Subsystem: Installation, User's, and Maintenance Guide manual, which can be found at the following URL.

http://www-304.ibm.com/systems/support/

4.8 Powering on and off

A typical DS3000 configuration includes the following components:

- Host servers
- DS3000 storage subsystem
- EXP3000 enclosures
- (Optional) SAN infrastructure

Powering on

It is important that you follow the correct power-on sequence, or the system might not work properly. Power the components on in this order:

1. SAN switches should be powered on first. In a typical SAN environment, the SAN switches are kept up and running all the time and are typically only be powered down in case of emergency.
2. EXP3000 enclosures are next. Turn on the switches on both power supplies on each EXP3000 enclosure. Allow enough time for the enclosures to fully power up, which can take several minutes.
3. After all EXP3000 enclosures are up and running, power on the DS3000. Be aware that large configurations can take up to ten minutes to fully power up.
4. Finally, power on all the host servers.

Important: We do not recommend powering on a DS3000 subsystem without any disk drives installed. If you do so, a new Feature Enable Identifier (FEI) will be generated, which could cause Premium Features to go out of compliance.

Should this happen, power on with the drives reinserted to restore the original FEI and make the Premium Features compliant again.

Powering off

In normal circumstances, the DS3000 subsystem would rarely be powered off; it is designed to run continuously. However, in certain events (scheduled maintenance, controller replacement on a single controller subsystem, emergency due to unforeseen circumstances, and so on), you will have to power the subsystem off. Use the same steps as above, but in the reverse order. If you cannot shut down and power off the host servers, then make sure the DS3000 logical drives are unmounted at the operating system level.

Important: Unless it is an emergency, the DS3000 subsystem should never be turned off if any amber LEDs are lit. You should use the Recovery Guru and other troubleshooting tools to correct the failure first.

5

Basic hardware configuration

In this chapter, we first describe the basic concepts you will need to understand in order to effectively configure and use the IBM System Storage DS3000. We then explain the procedures required to perform the initial setup of the DS3000:

- Initial management tasks (discovering and naming the subsystem)
- Enabling the software feature pack and Premium Features
- Creating arrays and logical drives
- Defining host servers and host groups
- Assigning host servers and logical drives into storage partitions

For clarity's sake, we use a real-life example throughout this chapter. We configure a DS3400 by performing all the actions listed above. Figure 5-1 on page 50 shows our sample hardware setup.

We use an IBM System x server with two FC HBAs, which is attached to a dual controller IBM TotalStorage DS3400 through a SAN switch. The IBM TotalStorage DS3400 is connected to an IBM TotalStorage EXP3000 enclosure.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

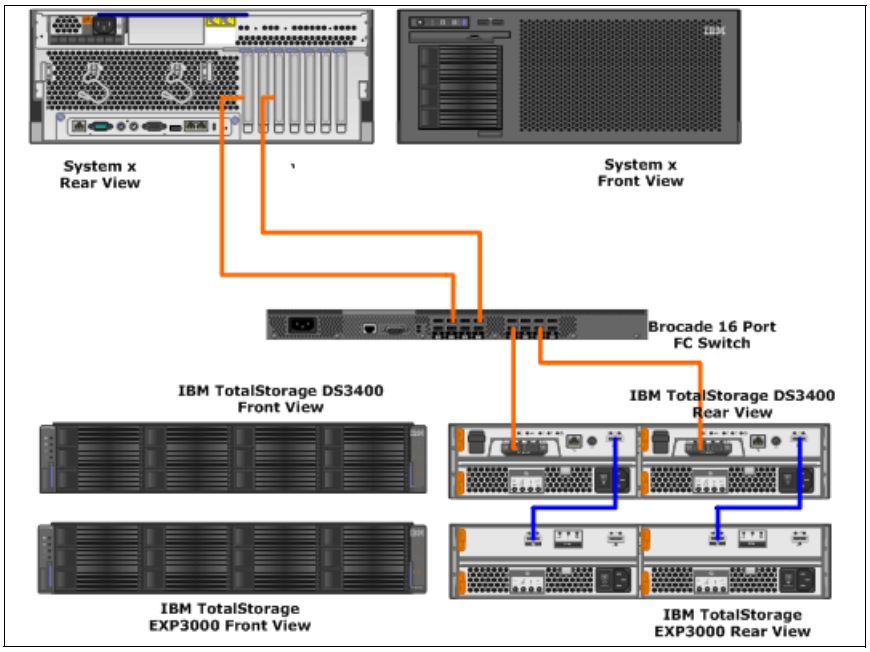


Figure 5-1 Sample hardware setup

The *Initial Setup Tasks* wizard is designed for a quick and easy DS3000 installation. This wizard contains the essential steps required to get your DS3000 up and running.

The tasks available through the wizard allow you to:

1. Locate the storage subsystem.
2. Rename the storage subsystem.
3. Set a storage subsystem password.
4. Configure host access.
5. Configure the storage subsystem (automatic or manual creation of arrays and logical drives).
6. (Optional) View and enable Premium Features.
7. (Optional) Change the network configuration.

We describe the Initial Setup Tasks wizard in 5.4, “Initial Setup Tasks wizard” on page 71.

Note: In this chapter, we focus mainly on the functions provided by the wizard. This wizard tries to automate functions where possible. However, experienced users will want to know how to perform these same functions manually. Therefore, you can consider this chapter to be a “Quick Start”; we provide more detailed information about these functions in Part 3, “Administration” on page 109.

5.1 Storage subsystem concepts

Let us discuss some basic concepts and terms applicable to the IBM System Storage DS3000 storage subsystem. The DS3000 is an external enclosure that provides shared storage to the attached host servers. We can configure a large number of physical disk drives into *arrays* and *logical drives*, and then specify which host servers can access certain logical drives. A logical drive on the DS3000 appears as a locally attached physical disk drive to the host server. When you configure an array, you have to specify a RAID level for this array. The RAID level selection will depend on your particular performance and availability requirements.

5.1.1 RAID levels

Here we discuss the various RAID levels.

What RAID is

The basic structure of RAID is the array. An array is a collection of drives that is configured, formatted, and managed in a particular way. The number of drives in the array, and the way that data is split between them, is what determines the RAID level, the capacity of the array, and its overall performance and data protection characteristics. Deciding what types of arrays to set up, and how to configure them, is the first thing you do when setting up a RAID implementation.

DS3000 subsystems support RAID levels 0, 1, 3, 5, 6, and 10. We describe each of these levels in the following sections. Each different RAID level provides a different balance between performance, data availability, and usable storage capacity.

RAID 0

RAID 0 is also known as *data striping*. In this level, the data is striped sequentially across all participating physical drives, as shown in Figure 5-2. This RAID level can provide very fast data access, but there is no redundancy. If a physical drive fails, the data in the array will be lost. You should therefore not use RAID 0 for storing any critical data. A good choice for RAID 0 would be data that might require very fast access and can be easily restored. Data access performance increases with the number of disk drives used in the array.

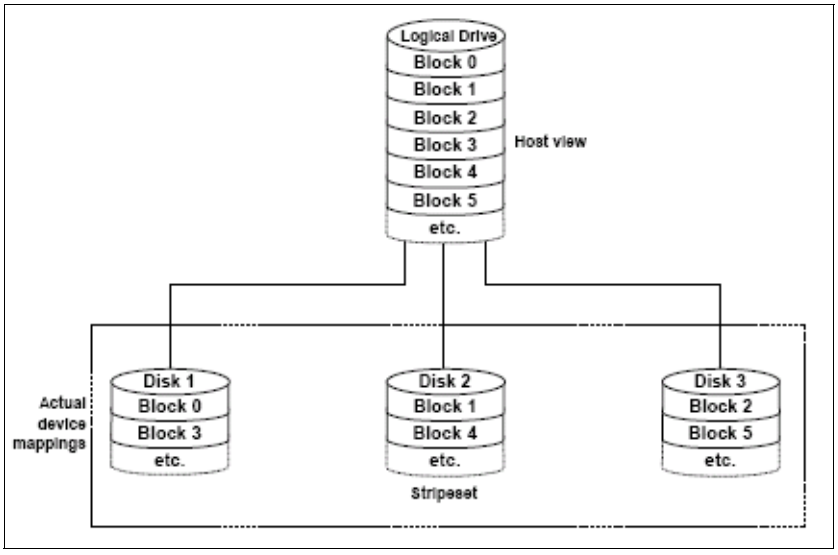


Figure 5-2 RAID 0

RAID 1

RAID 1 is also known as *disk mirroring*, that is, a mirrored set without parity. It uses exactly two disk drives to mirror the data between them (Figure 5-3). All read and write operations are performed on both physical disks. This RAID level provides good redundancy; in the case of a disk failure, all the data is still available and accessible. However, this comes at a cost: We can only use 50% of the total storage capacity.

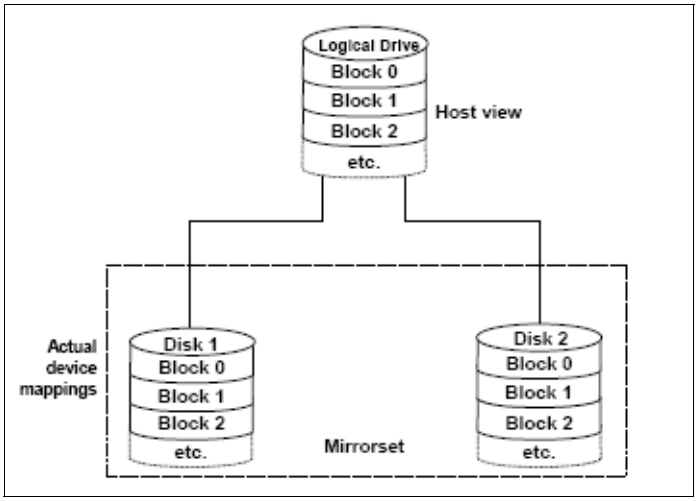


Figure 5-3 RAID 1

RAID 3

RAID 3 is a striped set with dedicated parity/bit interleaved parity. This RAID level is an attempt to overcome the low usable capacity limitation of RAID 1 configurations. The idea is to have a larger number of disk drives (three or more) and use one dedicated disk drive to store parity information. If any disk drive in the array fails, parity information can be used to calculate the missing information on the failed drive. Usable capacity is more efficient than in RAID 1, and as the number of drives increases, so does the storage space efficiency. The capacity of one drive is sacrificed for redundancy, so if you have five drives in your RAID 3 array, the storage space efficiency will be 80%. With ten drives, the efficiency will be 90%.

However, with RAID 3, the dedicated parity drive is a performance bottleneck. Each write operation requires the parity to be re-computed and updated, and this means that the parity drive is accessed every time a block is written onto the array. Because of this, RAID 3 is rarely used.

RAID 5

RAID 5 is a striped set with distributed parity. The idea of RAID 5 is similar to that of RAID 3, but with an important improvement: RAID 5 does not use a dedicated parity drive. Instead, parity blocks are evenly distributed across all physical disk drives, as shown in Figure 5-4. This means that RAID 5 offers better write performance than RAID 3, which improves greatly when using a larger number of disk drives.

RAID 5 is very popular, due to its good balance between performance, availability, and usable capacity.

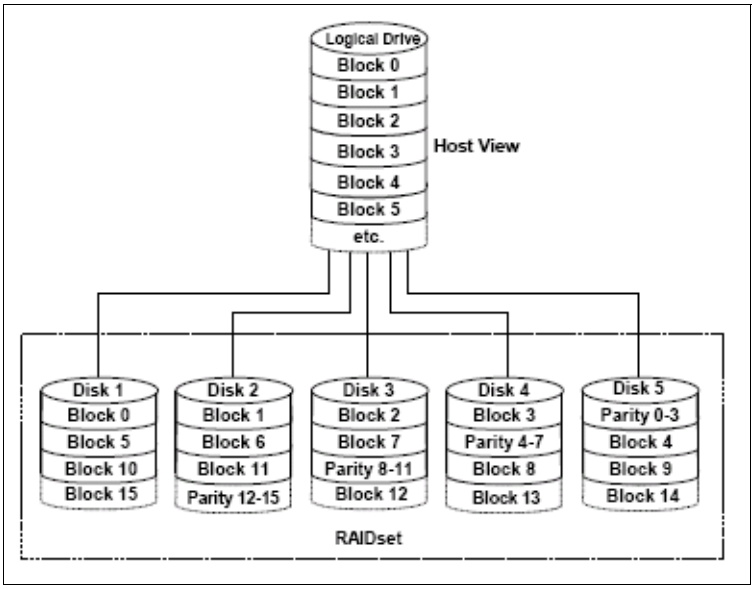


Figure 5-4 RAID 5

RAID 6

RAID 6 is a striped set with dual distributed parity, as shown in Figure 5-5 on page 54. This RAID level has high availability with additional fault tolerance. RAID 6 employs n+2 drives, which can survive the failure of any two drives.

RAID 6 stripes blocks of data and parity across an array of drives and it calculates two sets of information for each block of data (p+q). For the purposes of RAID 6 p+q, they can be used to generate up to two missing values from a set of data.The key to this method is the q, which is a codeword based upon Reed-Solomon error correction. As such, q is more like a CRC than parity. Based upon principles of set theory and linear algebra, Reed-Solomon codes are well-known codes that are also maximum distance separable.

The calculation of q is complex. In the case of the DS3000, this calculation is made by the hardware and thus performs better than the software-based implementation found in other storage systems.

By storing two sets of distributed parities, RAID 6 is designed to tolerate two simultaneous disk failures. This is a good implementation for environments using SATA disks.

Due to the impact of more parity calculations, in terms of writing data, RAID 6 is slower than RAID 5, but may be faster in random reads thanks to the spreading of data over one more disks.

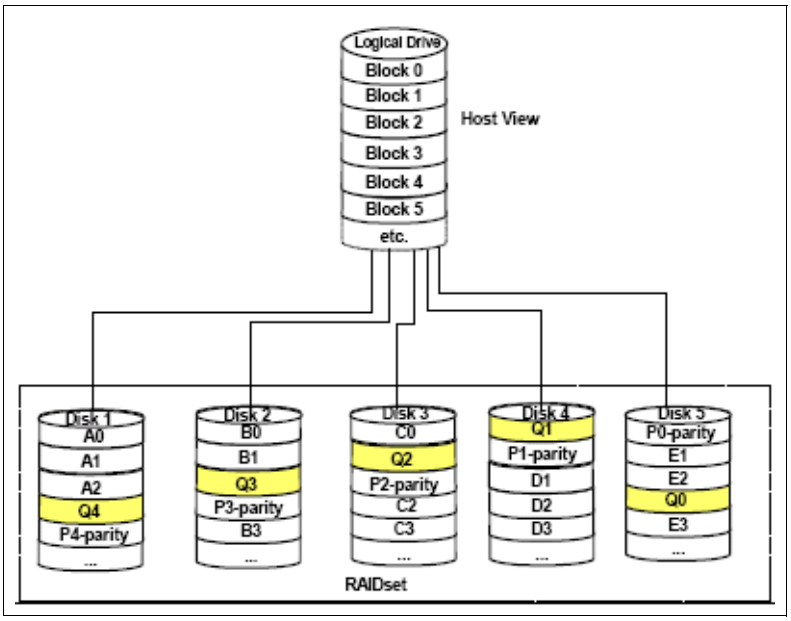


Figure 5-5 RAID 6

The q value allows the system to calculate up to two missing values (if you lose two disks) and can be used to recover the missing data.

It is good for multi-user environments, such as database or file system storage, where typical I/O size is small, and there is a high proportion of read activity, and in situations where additional fault tolerance is required.

RAID 10

RAID 10 is also known as $RAID\ 1 + 0$. As stated above, RAID 1 can only be configured on exactly two disk drives. RAID 10 is a way to expand mirroring onto four or more disk drives. It is actually a combination of mirroring and RAID 0 (striping), as shown in Figure 5-6 on page 55. RAID 10 provides excellent data availability. If a disk drive fails, all the data can still be accessed. Performance is also excellent, as the data is striped in a way similar to RAID 0. The only drawback is that the usable capacity is only 50% of total storage capacity.

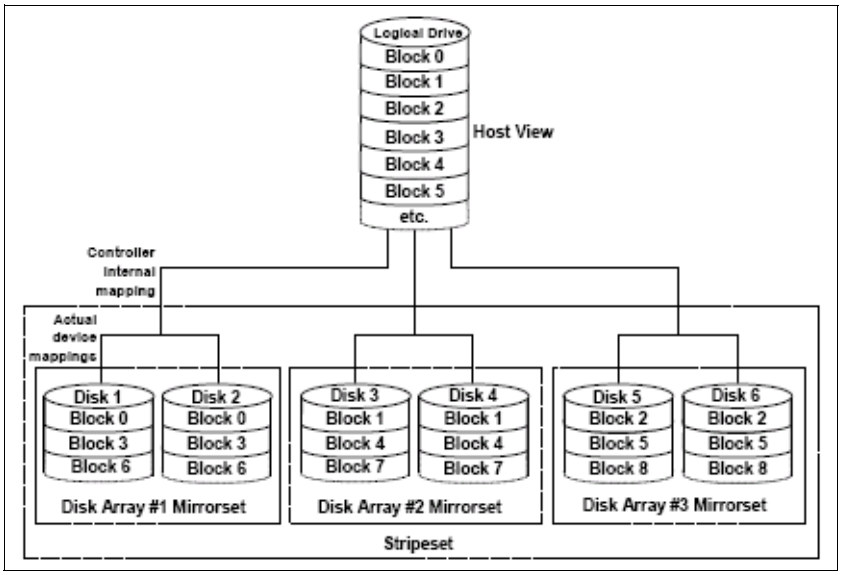


Figure 5-6 RAID 10

RAID levels summary

Based on the respective level, RAID offers the following performance results:

- RAID 0 offers high performance, but does not provide any data redundancy.
- RAID 1 offers high performance for write-intensive applications.
- RAID 3 is good for large data transfers in applications, such as multimedia or medical imaging, that write and read large sequential chunks of data.
- RAID 5 is good for multi-user environments, such as database or file system storage, where the typical I/O size is small, and there is a high proportion of read activity.
- RAID 6 is good for multi-user environments, such as database or file system storage, where the typical I/O size is small and there is a high proportion of read activity, and in situations where additional fault tolerance is required.
- RAID 10 offers higher performance than RAID 1, and more reliability than RAID 5.

5.1.2 Arrays and logical drives

Before the host servers can start using the disk drives on the DS3000, you have to create at least one array and logical drive. An *array* is simply a set of physical disk drives. Each array has to be carved into one or more logical drives. Finally, you have to assign the logical drives to the host servers. Host servers see the logical drives as physical disk drives.

The RAID level is defined per array. This means that all logical drives inside an array will use the same RAID level.

Figure 5-7 illustrates the concept of arrays and logical drives. As indicated already, DS3400 is connected to an System x server. Now let us configure it. The diagram shows nine physical disk drives, labeled *Drive 1* through *Drive 9*. Disk drives 1 to 4 are grouped into *Array 1* and set to operate as RAID level 10. Drives 5 to 9 operate as RAID 5 and are members of *Array 2*.

Array 1 is carved into two logical drives (1 and 2) and there is some free space left. Array 2 is split into logical drives 3, 4, 5 and 6, which take up an entire array space.

Logical drives 1 through 6 are now ready to be assigned to the host servers.

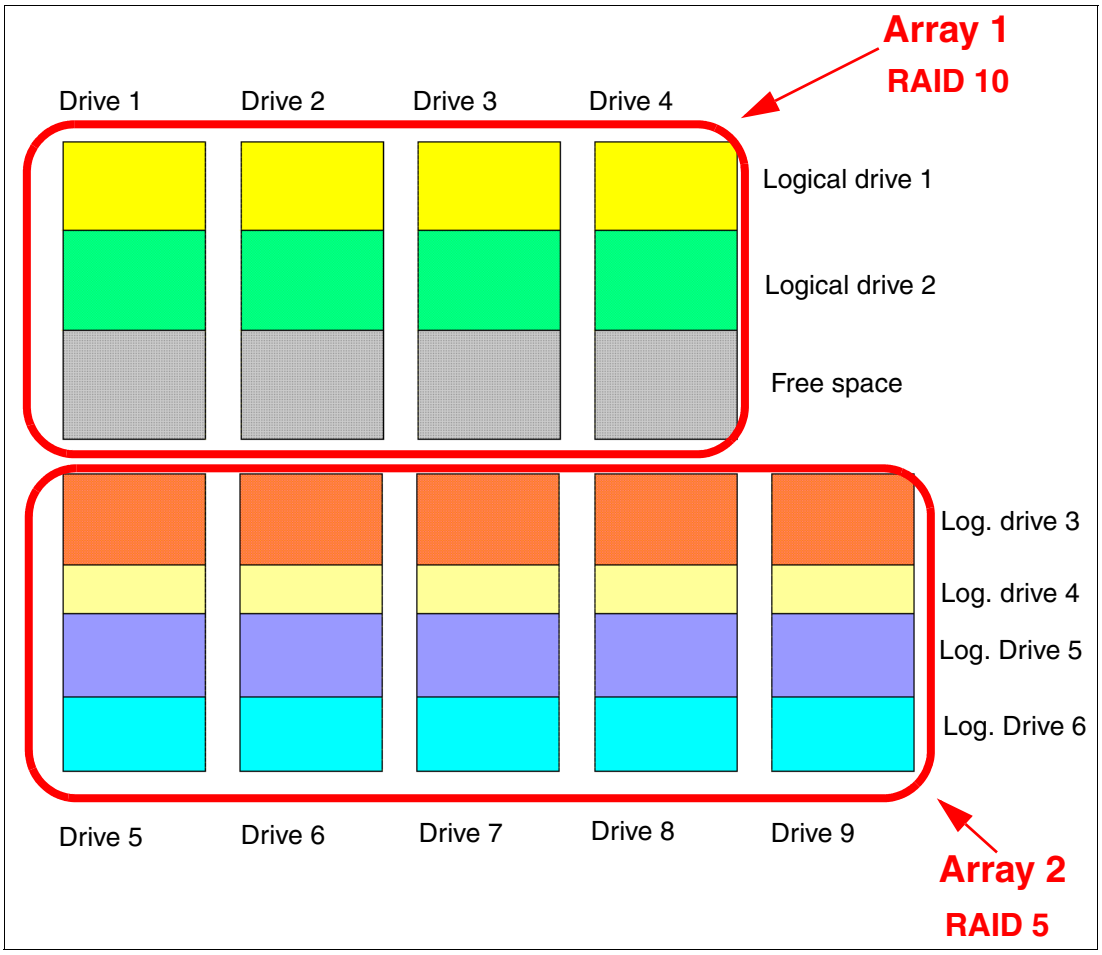


Figure 5-7 Physical drives, arrays, and logical drives

Normally, each logical drive should be assigned to exactly one host server. The only exception is clustered host servers, which all need to access the same set of logical drives.

Segment size

Segment size is the maximum amount of contiguous data that is placed on one drive before striping to the next drive.The segment size can be 8, 16, 32, 64, 128, 256, or 512 KB. In order to select the appropriate segment size for a logical drive, you should know the data access pattern. For example, large sequential data access will benefit from a large segment size. On the other hand, random workload in small blocks will perform better with smaller segment sizes. Incorrect segment size selection will impact performance, so you should plan and carefully select the value suitable for your environment. Although the segment size can be changed later, it is best to think about the right selection before creating the logical drives.

Tips: The possible segment sizes available are 8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, and 512 KB. You should consider the following:

- Storage Manager sets a default block size of 64 KB.
- For database applications, block sizes between 32–128 KB have shown to be more effective.
- In a large file environment, such as on media streaming or CAD, 128 KB and bigger is recommended.
- For a Web server or file and print server, the range should be between 16–64 KB.

Enclosure loss protection

When each physical drive belonging to an array is in a different enclosure, the array is protected against enclosure loss. If an EXP3000 enclosure fails, only one drive in the array will be lost, so the array will remain operational. The maximum number of EXP3000s supported behind the DS3000 subsystem is three. This means that only arrays with up to four disk drives can be protected against enclosure loss.

If you create larger arrays, enclosure loss protection will not be possible, since some enclosures will need more than just one drive from the array. In addition, only arrays with RAID levels 1, 3, 5, and 10 can be protected against enclosure loss. A RAID 0 array has no redundancy, so enclosure loss protection is not possible.

Figure 5-8 shows an example of an array protected against enclosure loss. The array consists of four physical drives, marked *Drive1, Drive2, Drive3,* and *Drive4*. As you can see, each drive is in a different enclosure.



Figure 5-8 Enclosure loss protection

5.1.3 Hot spares

Even though disk drives are becoming increasingly more reliable, disk failures can occur. To protect your data in case of a drive failure, you should primarily use a redundant RAID level, for example, RAID 1, 5, or 10. This way, your data will remain intact should a hard drive fail. However, the array will then operate in a degraded state until the failed disk drive is replaced and data reconstruction completes on the new disk drive.

To ensure as high availability as possible, we strongly recommend that you use hot spare drives. Under normal circumstances, the hot spare drives do not participate in any array. When one of the disk drives fails, the host spare will automatically be used as a replacement drive. Data reconstruction onto the hot spare will take place immediately. When reconstruction finishes, the array will again be in a fully operational status and the hot spare will become a member of the array.

Depending on your configuration, it might be wise to use more than just one hot spare. In small DS3000 configurations with a relatively low number of disk drives, one hot spare might suffice. But for larger configurations, we recommend that you define several hot spares.

When defining hot spare drives, you need to consider the drive size. For example, if you use a mixture of different disk drive sizes in your arrays, the hot spare drive size must be large enough so that it can effectively replace any of the drives.

The following methods are available to allocate hot spare drives in the storage subsystem:

- Automatic assignment: The storage subsystem automatically calculates the number of hot spare drives needed and allocates accordingly. This can be used on unconfigured storage subsystems.
- Explicit assignment: The hot spare drives are manually selected and assigned.

Use the DS3000 Storage Manager to configure the above options.

5.1.4 Storage partitioning

The DS3000 subsystem can provide storage access to many host servers. Access to the RAID controllers is shared among the attached host servers. However, the actual logical drives access cannot be shared. Each logical drive has to be assigned to one server only, except when you have a group of clustered servers; in this case, the whole group will have access to the same set of logical drives.

If two or more independent, non-clustered servers try to access the same logical drive, this setup will not work properly, which will lead to application errors, data corruption, and operating system crashes, because there is no locking mechanism in place to arbitrate for the logical drive access.

When the servers are clustered, it is a different story. Clustering middleware provides a locking mechanism so that all the host servers can access the logical drives in consistent manner and without conflicts. For example, Microsoft Cluster Server uses a simple mechanism; at any given time, a particular logical drive access is allowed to one host server only. Clustering software prevents access to that logical drive from other host servers.

To ensure proper access to the logical drives, we have to create assignments, or mappings, of host servers and groups to the logical drives. We call this *storage partitioning*. A storage partition is a set of logical drives that can be accessed by the same host or a host group.

Let us take a look at an example: In Figure 5-7 on page 56, we have nine physical disk drives and we divide them in two arrays. Array 1 uses RAID 10, consists of four disk drives, and is divided into two logical drives. Array 2 contains five disk drives, is configured as RAID 5, and contains four logical drives. We therefore have six logical drives available for host server use.

We want to map the logical drives to four host servers:

- Host 1 is a stand-alone server and needs access to logical drive 1.
- Host 2 is also a stand-alone server. It requires access to logical drives 3 and 6.

- Host 3 and Host 4 operate in a cluster. They need access to logical drives 2, 4, and 5.

The host-to-logical-drive mappings have to be done as shown in Figure 5-9. As you can see, we need to create three storage partitions.

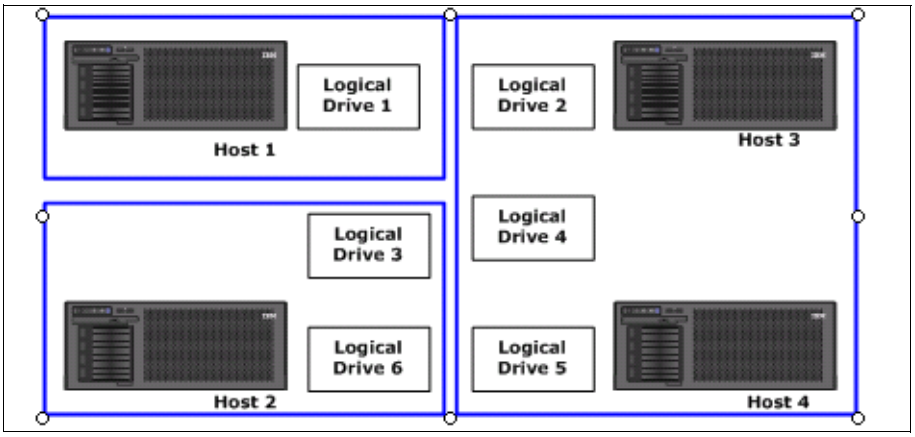


Figure 5-9 Storage partitions

The standard configuration is four storage partitions enabled on all DS3000 products, and a maximum of 32 partitions can be enabled using different combinations of the available DS3000 Partition Expansion License.

5.2 DS3000 Storage Manager installation

The DS3000 Storage Manager package contains several components, explained in 3.4.1, “DS3000 Storage Manager components” on page 27. The installation process consists of installing the appropriate set of Storage Manager components on the target computer.

We will install Storage Manager on two systems: a host server and a management station. The following components will be installed in the host server:

- SMclient
- SMagent
- SMutil
- Multipath support

When the installation completes, the host server will have all the components necessary for accessing the logical drives on the DS3000 and for in-band management. Refer to in-band management and out-of-band management in 3.4.2, “Management methods” on page 30.

Note: Before installing your DS3000 hardware and software, be sure to meet all the requirements for your DS3000 model, HBAs, firmware version, SAN, and OS level as specified in the IBM Support Web site.

Go to the Web page for the System Storage Interoperation Center to get the latest DS3000 compatibility information:

`http://www-01.ibm.com/systems/support/storage/config/ssic/index.jsp`

The management station only needs one component: SMclient.

The DS3000 Storage Manager can be downloaded from the IBM Systems support Web site at:

http://www-304.ibm.com/jct01004c/systems/support

Storage Manager packages are available for:

- Windows Server 2003/2008 (32-bit and 64-bit versions)
- Linux 2.6 kernel
- IBM AIX V6.1

In our example, we will install Windows versions on both the host server and the management station. More detailed installation instructions can be found in Chapter 6, “Installing IBM System Storage DS3000 Storage Manager” on page 81.

1. Launch the installation executable file, which will be in SMIA-WS32-xx.xx.xx.xx for 32-bit Windows. This is a self extracting archive (see Figure 5-10).

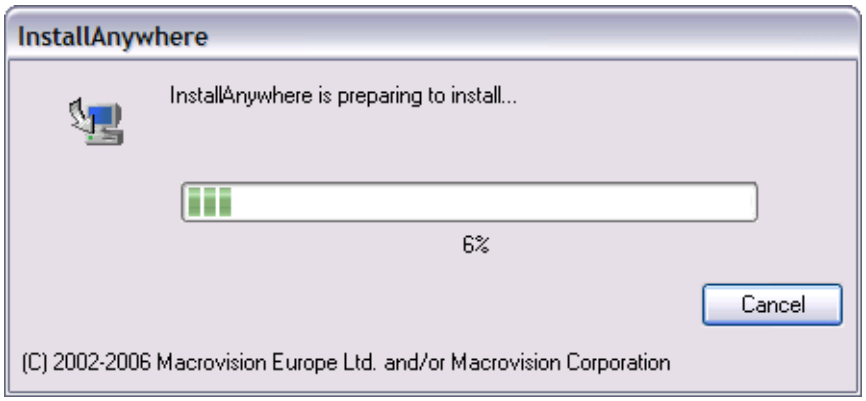


Figure 5-10 SM Manager Self Extract

2. Select the language, as shown in Figure 5-11 on page 61, and click **OK**.



Figure 5-11 Language selection

3. Next you will see the Introduction and Copyright Statement windows. Read the information, and then click **Next** on each window.

4. The License Agreement displays (Figure 5-12). Review the agreement, then accept it. If you do not accept the License Agreement, the installation process will not continue.

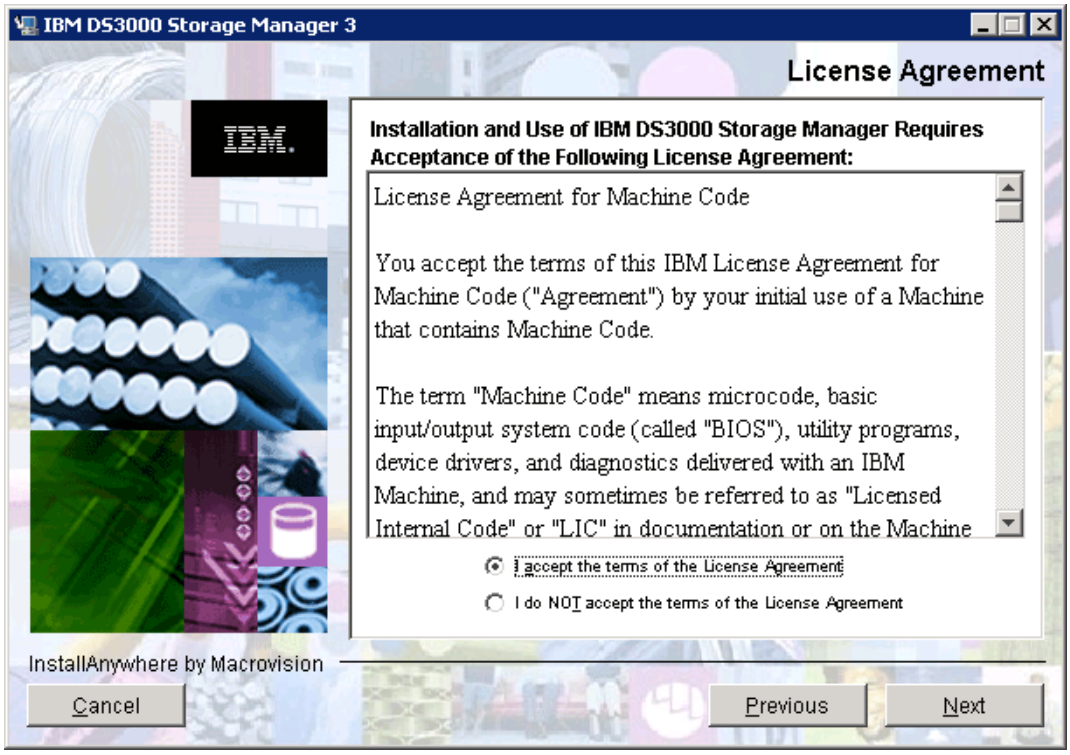


Figure 5-12 License Agreement

5. If you do not already have Storage Manager installed, you are prompted for the installation directory, as shown in Figure 5-13 on page 63. The default is C:\Program Files\DS3000.

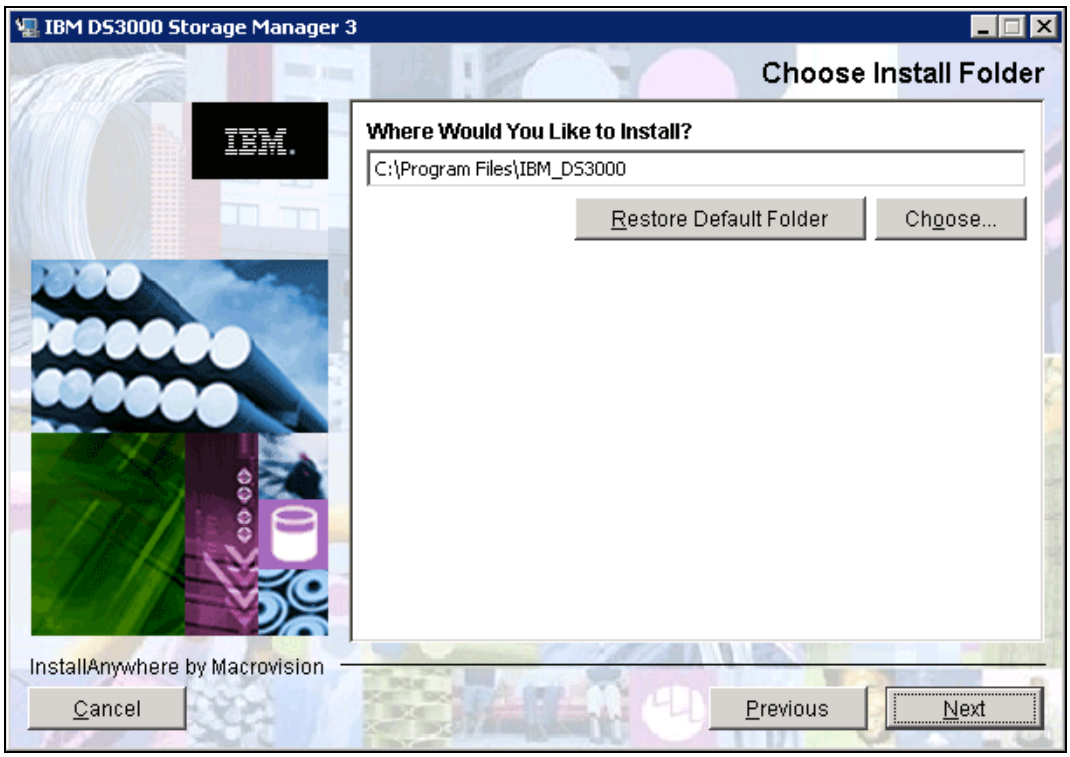


Figure 5-13 Select installation directory

6. Figure 5-14 shows the Installation Type selection. Select one of the installation types, depending on whether this is a host server or a management station.

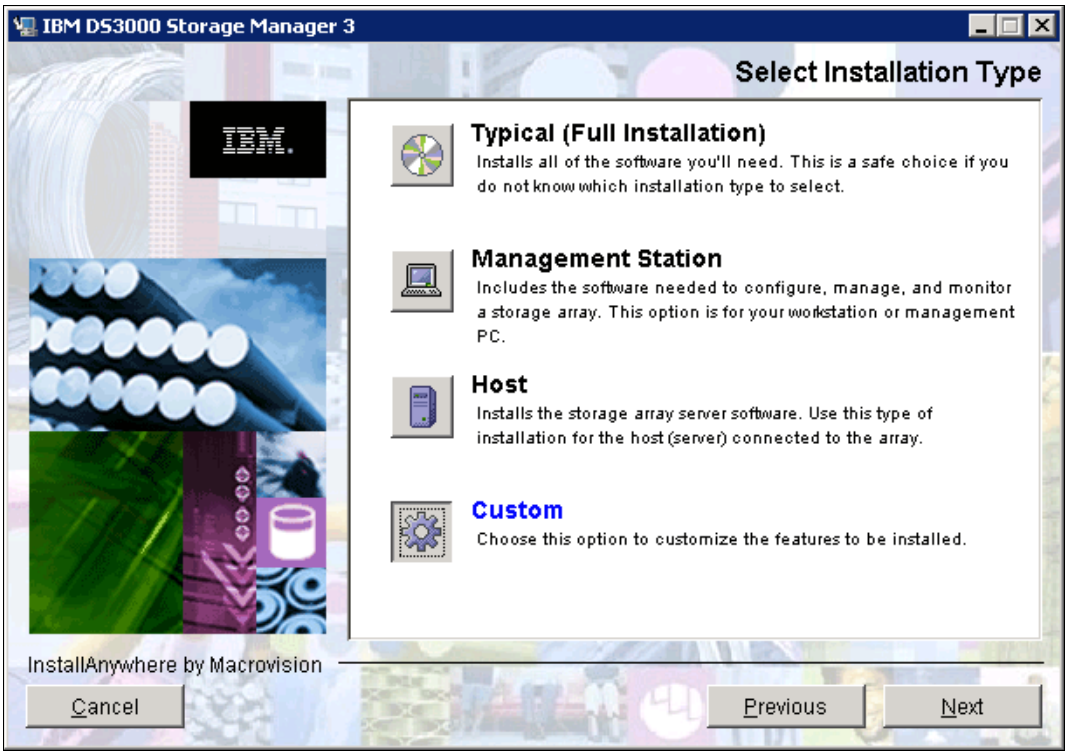


Figure 5-14 Select installation type

The first two selections (*Management Station* or *Host*) will install a predefined set of components. There are two additional options:

– Typical (Full Installation)

  Select this installation type if you want to install all Storage Manager components.

– Custom

  Use this selection if you want to specify your own set of components to install. We show this in Figure 5-15.



Figure 5-15 Custom installation - select components

Besides the components that have already been explained (SMclient, SMutil, SMagent, and multipath support), you also see *Java Access Bridge* on the list. This is an optional component, which provides support for assistive technology (for example, screen readers), so that visually impaired users can use Storage Manager.

7. If an older version of the DS3000 Storage Manager is already installed, the installation process displays a warning, as shown in Figure 5-16 on page 65. The message advises you to back up the configuration files emwback.bin and emwdata.bin before continuing the installation.



Figure 5-16 Overwrite Warning

8. In Figure 5-17, you are asked whether to automatically start the DS3000 Storage Manager monitor. The monitor should only run on one management station, therefore if you are installing your first management station, select **Automatically Start Monitor (Recommended)**. If you are installing an additional management station, then you do not need to select this.

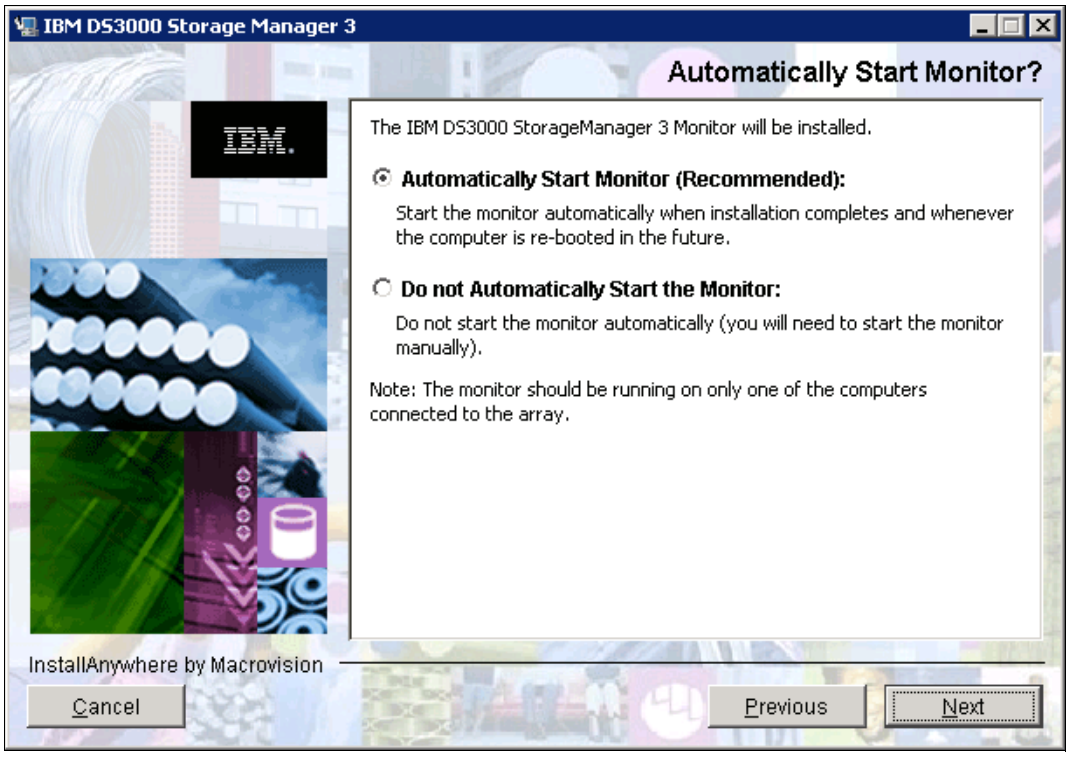


Figure 5-17 Automatically Start Monitor

9. The Pre-Installation Summary window displays, with all the options selected so far. Review the information and click **Install** to initiate the actual installation process. You can monitor the installation progress bar, as shown in Figure 5-18.



Figure 5-18 Installing DS3000 Storage Manager

10. In addition, another installation status messaging window will appear during the installation, as shown in Figure 5-19.

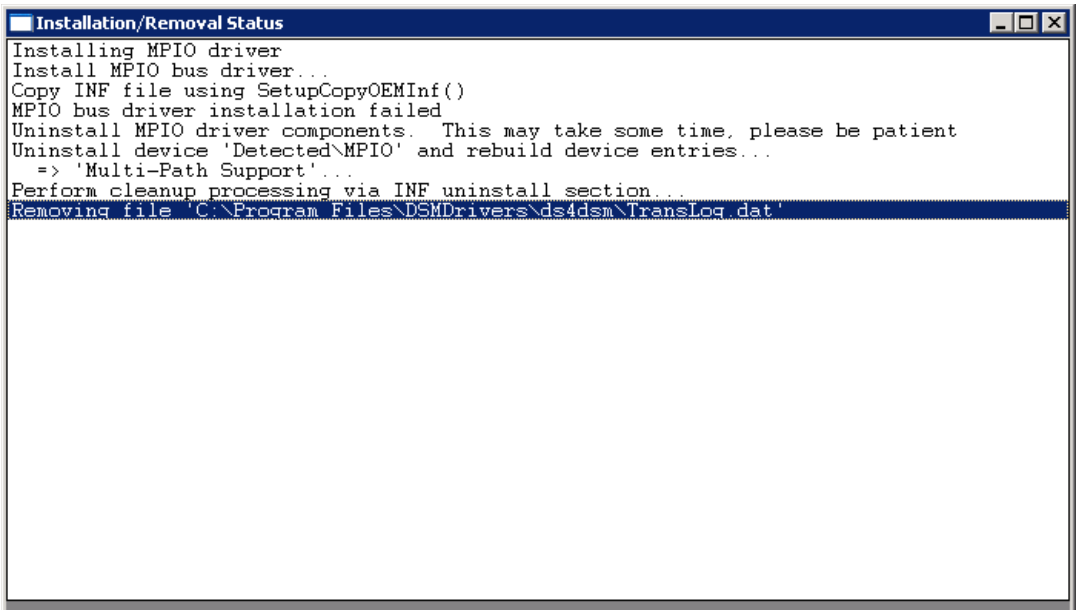


Figure 5-19 Installation Status window

11. When the installation finishes, the Install Complete window appears. Click **Done** to exit the installation utility, as shown in Figure 5-20.

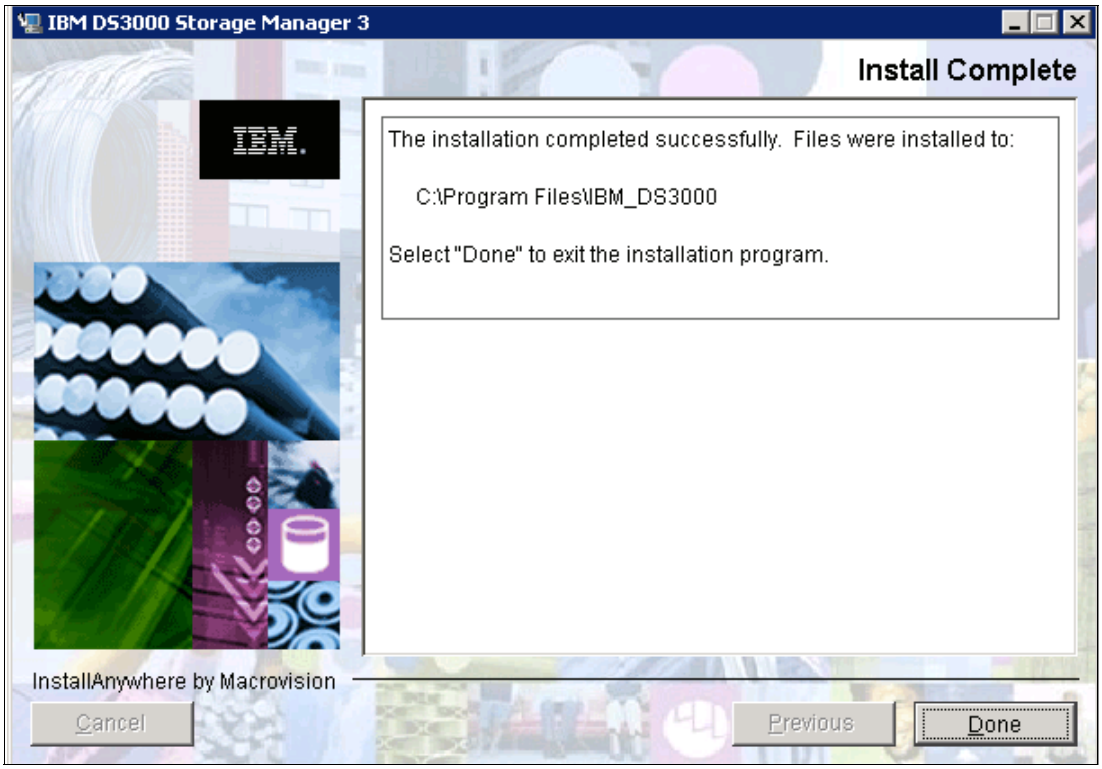


Figure 5-20 Installation Complete

12. Host server components installation requires that the host server to be restarted.

The DS3000 Storage Manager is installed and ready for use. To start Storage Manager, select **Start** → **All Programs** → **Storage Manager 3 Client** → **Storage Manager 3 Client**.

5.3 Add a DS3000 subsystem in Storage Manager

Right after the installation, the Storage Manager will be unconfigured, since no DS3000 subsystems have been defined in the Enterprise Management Window. Therefore, you will be prompted to add a storage subsystem as follows:

1. The Select Addition Method window appears, as shown in Figure 5-21.

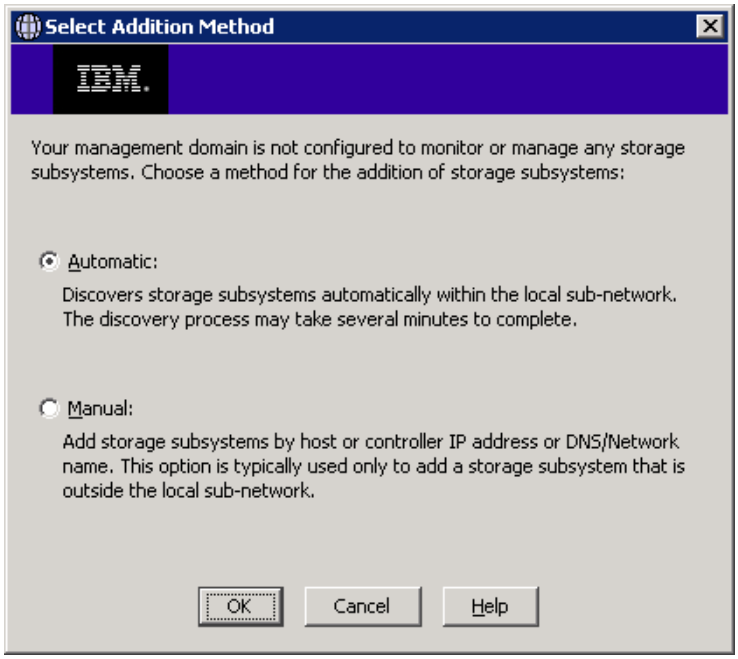


Figure 5-21 Select Addition Method

Automatic subsystem discovery will only work if the management station and the storage subsystem are on the same subnet. In our sample case, they are on different subnets, so we will have to proceed with the *manual* addition method (by clicking the **Manual** radio button).

2. Click **OK** to display the Add Storage Subsystem window (see Figure 5-22 on page 69).

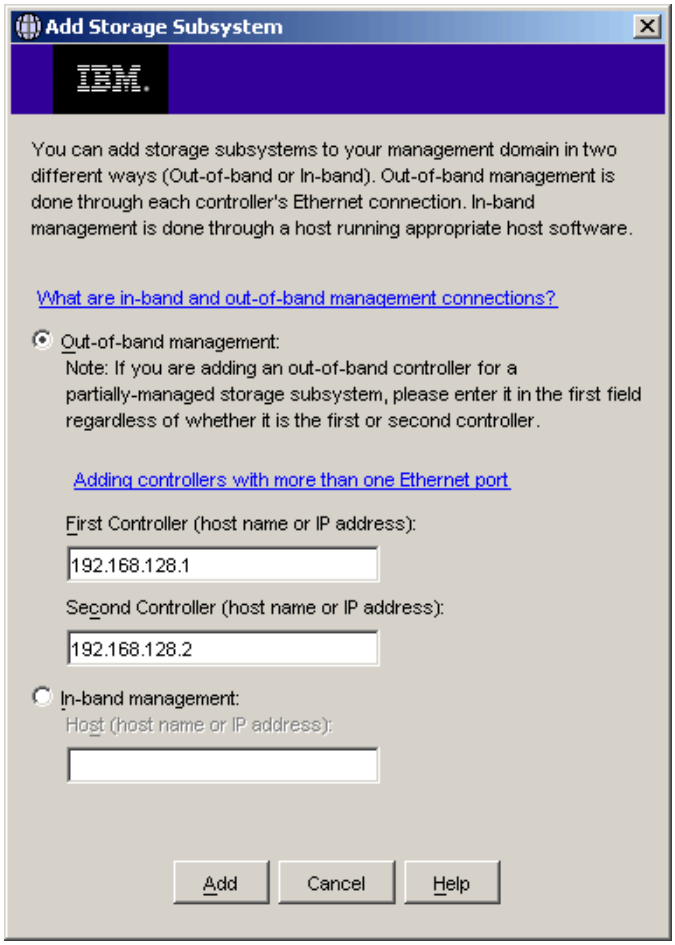


Figure 5-22 Add Storage Subsystem

This is where you specify the management type. In our case, we select **Out-of-band management**. As this is the initial installation, both controllers are configured with the default management IP addresses. Enter the two default IP addresses, supplied with the factory condition.

- 192.168.128.101 for the first controller.
- 192.168.128.102 for the second controller.

Note: Before using the default management IP addresses, the controllers actually try to obtain addresses through DHCP/BOOTP first. If this fails, the default IP addresses are used.

Each DS3000 controller has a label showing its management port MAC address; you can use this to determine which IP address was assigned if the DHCP/BOOTP is successful.

3. Click **Add**. The storage subsystem will be added to the Enterprise Management window and you will receive confirmation, as shown in Figure 5-23. The dialog box offers an opportunity to add another storage subsystem. As this is the only DS3000 we want to add at this point, we click **No**.

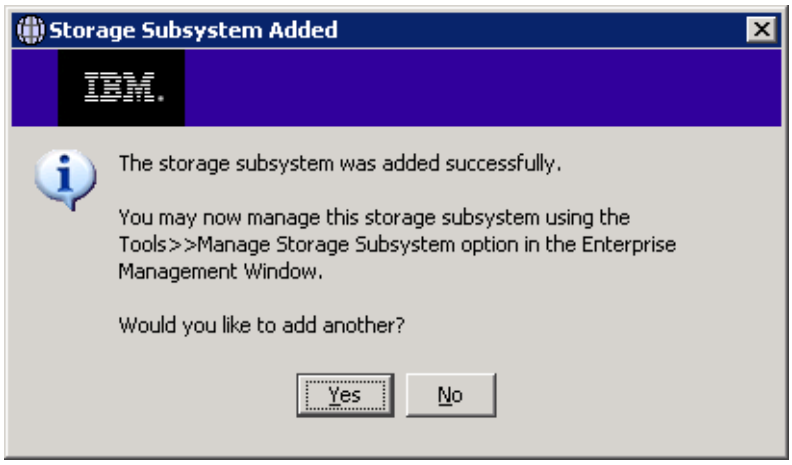


Figure 5-23 Storage Subsystem Added

4. The Enterprise Management window now contains our DS3400, accessible through the out-of-band management connection (see Figure 5-24).

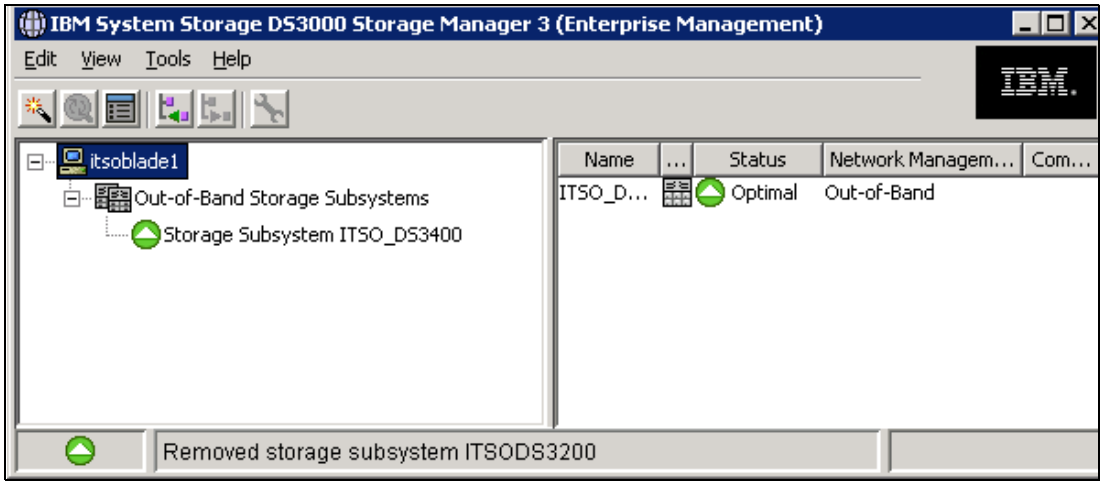


Figure 5-24 DS3400 in Enterprise Management window

5. Double-click the DS3400 to launch the Subsystem Management window. Besides the Subsystem Management window, another window opens up. This is the Initial Setup Tasks window, shown in Figure 5-25 on page 71. We will discuss each contained task.

5.4 Initial Setup Tasks wizard

We use the Initial Setup Tasks wizard to perform most of the basic configuration steps. Figure 5-25 displays the tasks available through this wizard.

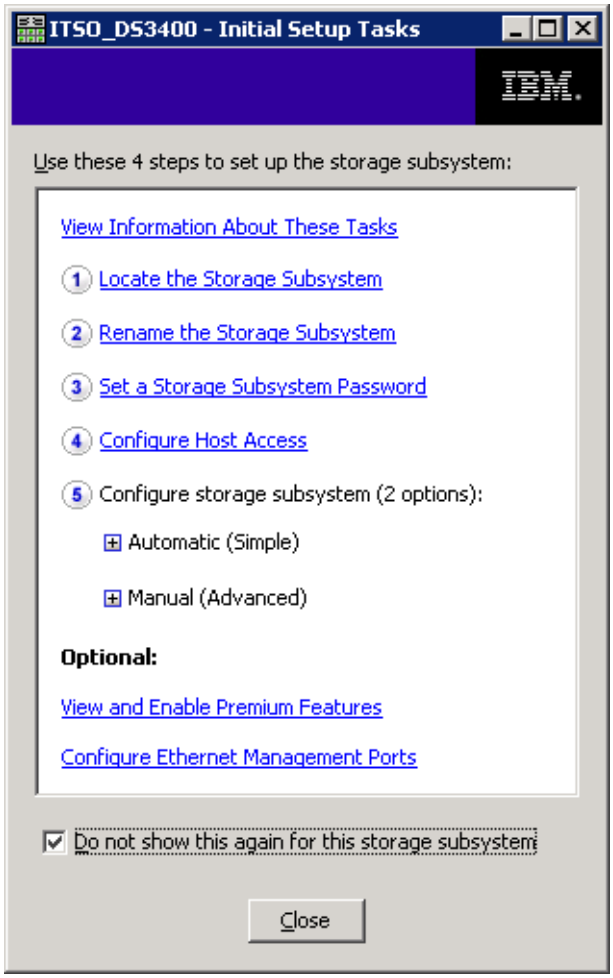


Figure 5-25 Initial Setup Tasks

The wizard is basically a set of shortcuts to the listed tasks.

5.4.1 Locating the storage subsystem

When you select this task, a window in the subsystem management window similar to Figure 5-26 appears. Use this option to locate a particular DS3000 within a rack full of other equipment (host servers, storage subsystems, LAN and SAN switches, and so on). You can locate a DS3000 subsystem, a particular controller, or an EXP3000 enclosure.

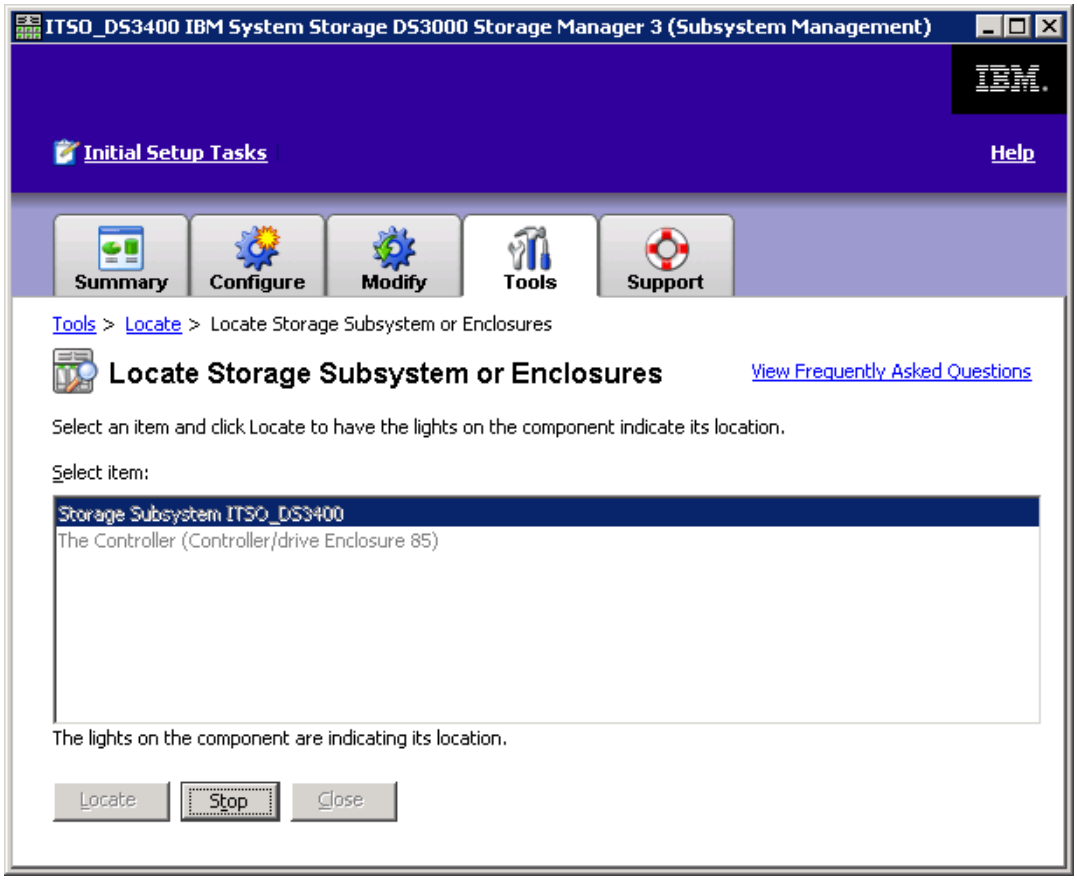


Figure 5-26 Locate Storage Subsystem window

Click **Locate**, and the LED on the selected component will indicate the component's location.

5.4.2 Renaming the storage subsystem

Out-of-the-box DS3000 will not have a meaningful name in DS3000 Storage Manager. Especially if you have multiple DS3000s, you will want to assign distinguishing names to appear in the Enterprise Management window. We show the Rename Storage Subsystem window in Figure 5-27 on page 73. Enter a new name, and click **OK** to confirm it.

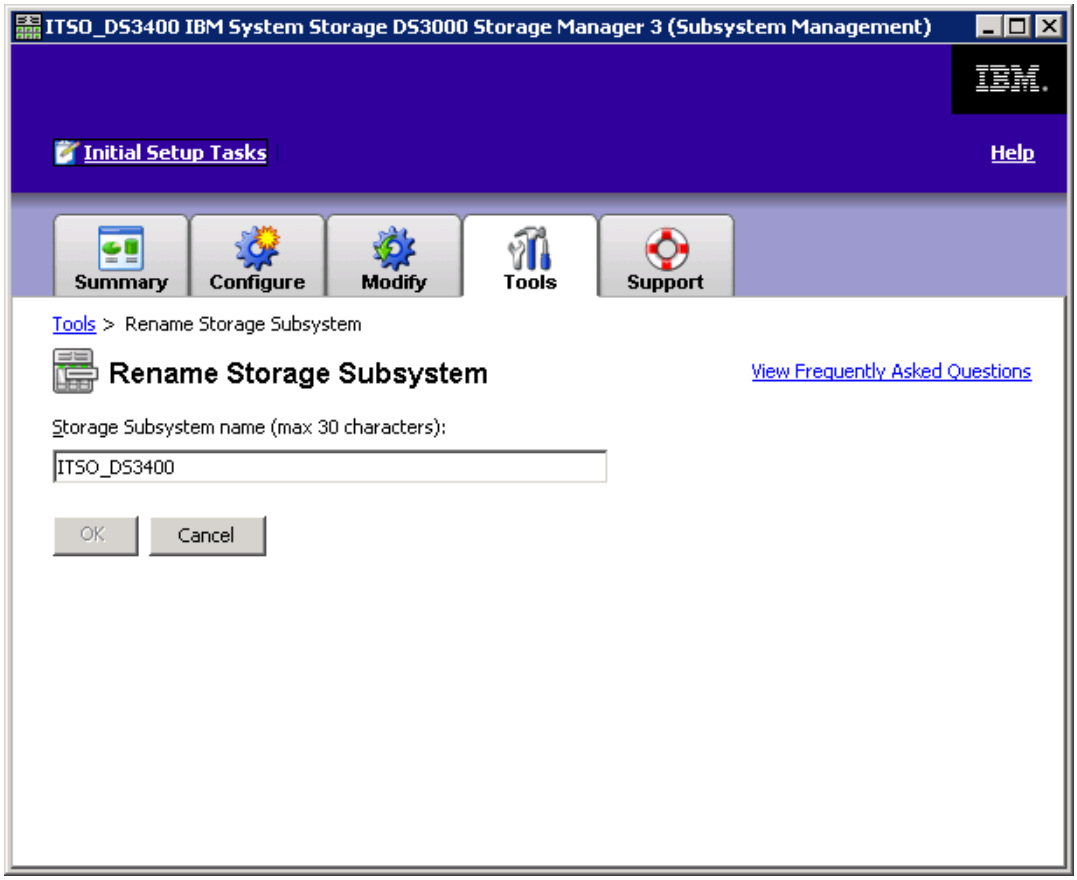


Figure 5-27 Rename Storage Subsystem window

Note: Ensure that you make the right choice for the storage subsystem name at this time. Once the host servers begin accessing logical drives, you should not change the subsystem name, because some operating systems use the subsystem name for identification.

5.4.3 Setting a storage subsystem password

We recommend implementing security and access control for DS3000 storage subsystem management, because DS3000 Storage Manager is freely available on the IBM Support Web site. Without security in place, anyone with access to the management network could download and install the Storage Manager and look for storage subsystems. An unauthorized user could then connect to a certain subsystem and be able to perform destructive tasks, such as deleting logical drives and arrays.

To prevent this, you can set a password for the storage subsystem. If the password is enabled, then it will be required for any potentially destructive tasks. Set the password as shown in Figure 5-28.

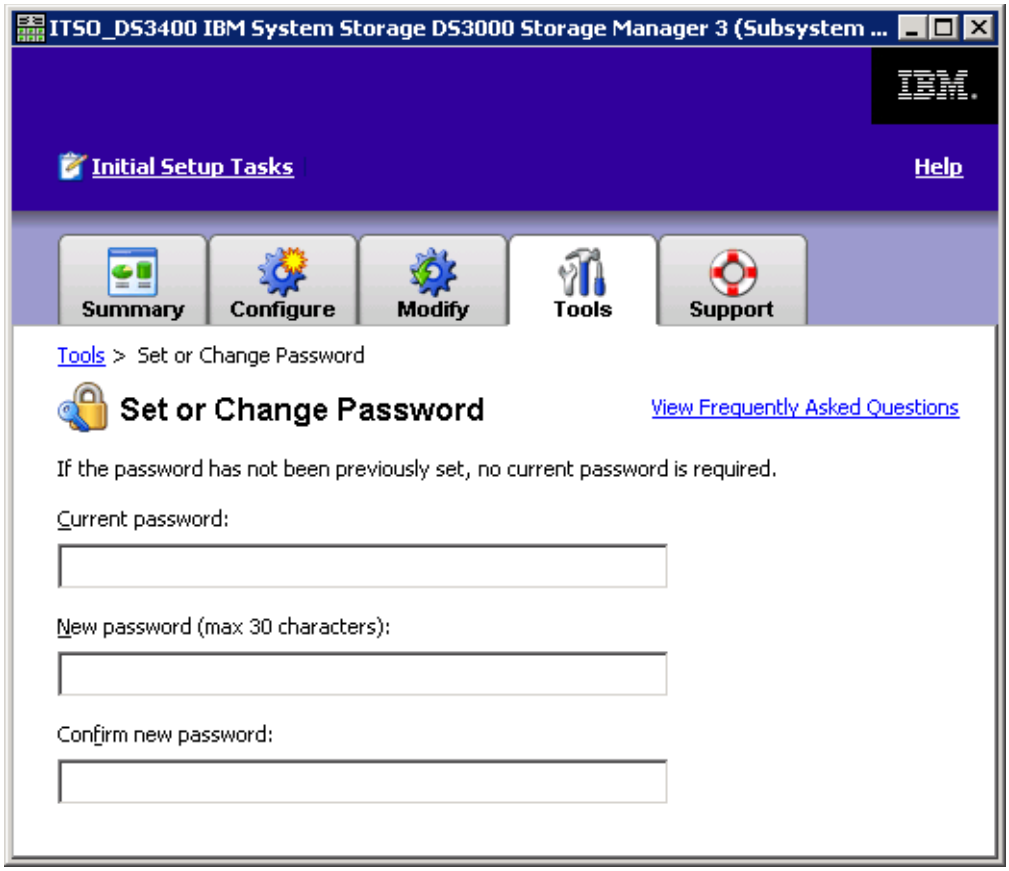


Figure 5-28 Set or Change Password window

5.4.4 (Optional) Viewing and enabling Premium Features

Click the **Tools** tab and click **View/Enable Premium Features**, which launches the View/Enable Premium Features task, shown in Figure 5-29 on page 75.

In our example, some Premium Features have already been enabled.

Note: The DS3200 Software Feature Pack, DS3300 Software Feature Pack, DS3400 Software Feature Pack, and the EXP3000 Expansion License have been incorporated into the base models and are no longer needed as prerequisites to the other advanced function options. The required controller firmware that enables this function must be at firmware level V06.70.23.00 or later and is available on the DS3000 product support Web site:

`http://www-03.ibm.com/systems/storage/disk/entry/index.html`

Also, you can refer the IBM System Storage DS3000 series Interoperability Matrix at the following link:

`http://www-03.ibm.com/systems/resources/systems_storage_disk_ds3000_pdf_interop.pdf`

For more information about Premium Features, refer to 3.3, “Premium Features” on page 25 and 12.1.4, “View/Enable Premium Features” on page 265 for information about applying and managing the Premium Features.

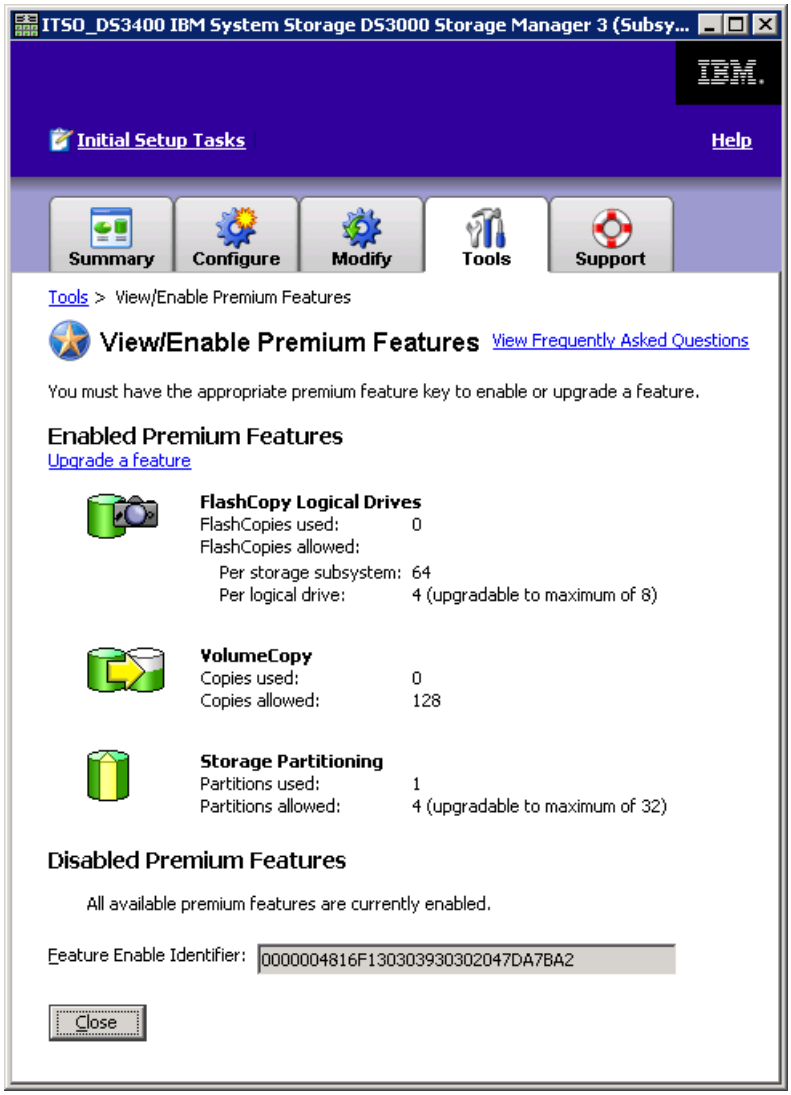


Figure 5-29 View/Enable Premium Features window

5.4.5 (Optional) Changing the network configuration

At this point, most of the basic configuration tasks on the DS3000 subsystem are done. However, the management IP addresses are still set to default values. As these addresses are more than likely not compatible with the typical customer IP assignments, it is now time to change the addresses to proper values. We use the Change Network Configuration task to do this.

Begin by clicking the highlighted task, as shown in Figure 5-30.

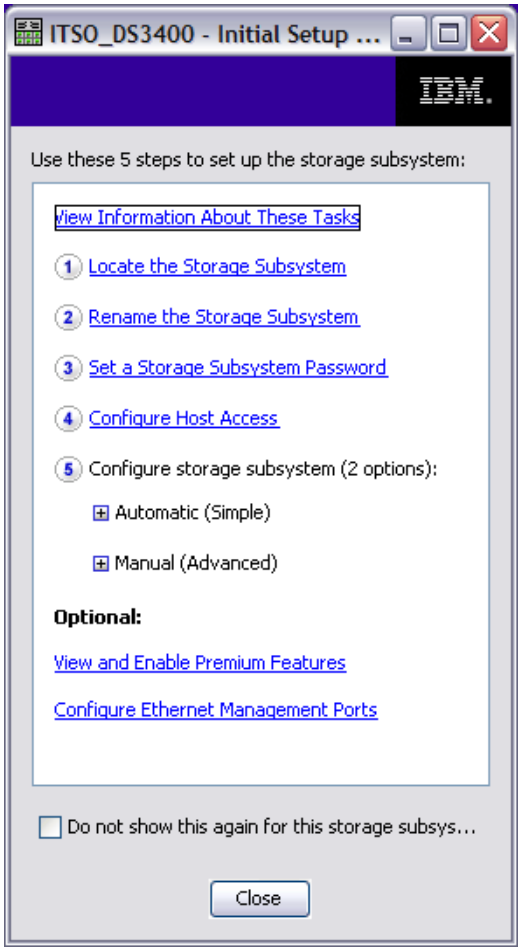


Figure 5-30 Initial Setup Tasks window

This action brings up the window shown in Figure 5-31 on page 77. You can select the controller (A or B) that you want to configure. The MAC address of the Ethernet management port is displayed, and you can set the IP address, subnet mask, and gateway to the values appropriate for your specific environment. Also, IPv6 settings and the port speed can be selected.

Make sure to select and configure both controllers if you have two of them.

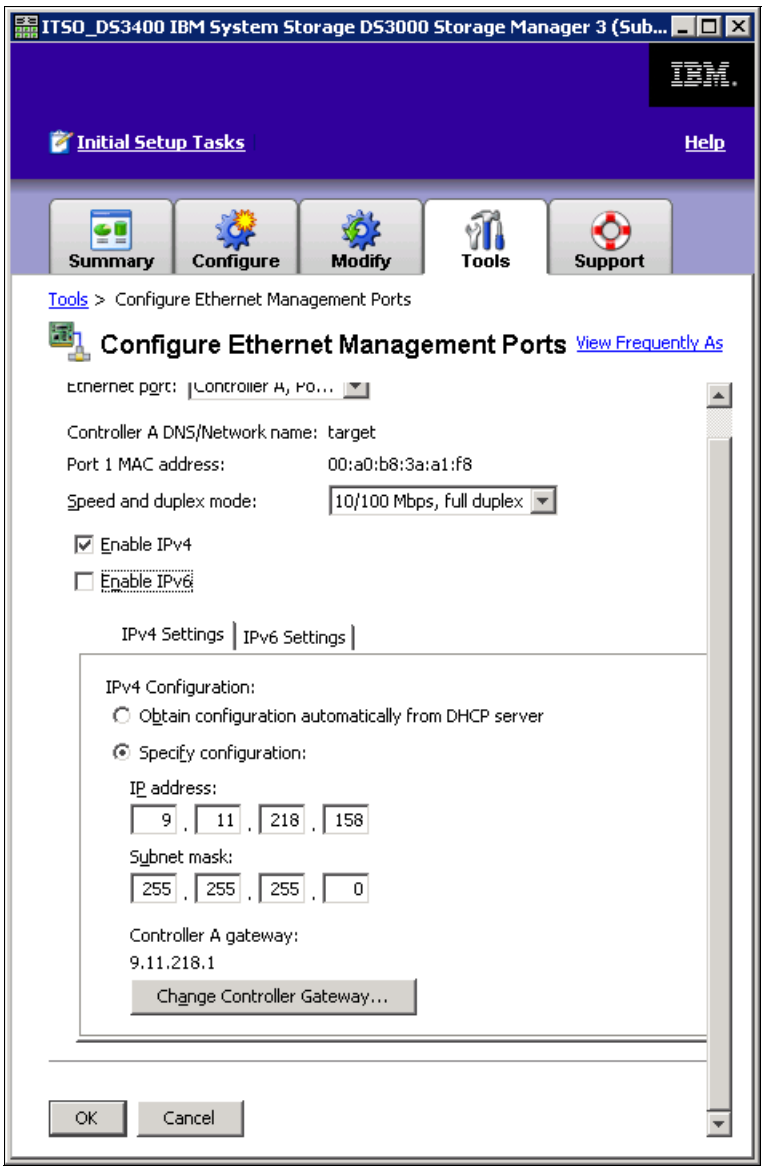


Figure 5-31 Change network configuration

Because controllers A and B have had their IP addresses changed, you have to add the storage subsystem to the Enterprise Management window again (and remove the storage subsystem that was added with default IP addresses).

To remove a subsystem, right-click it and select **Remove** from the context menu. This is shown in Figure 5-32.

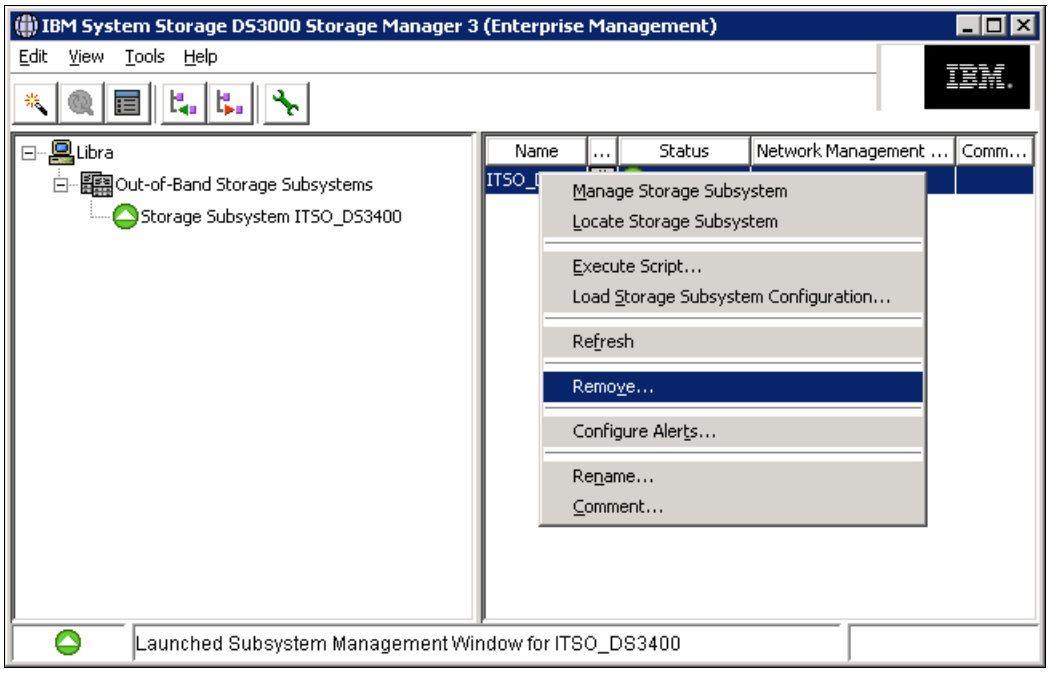


Figure 5-32 Remove the subsystem

Now let us add the subsystem with these IP addresses, as shown in Figure 5-33.

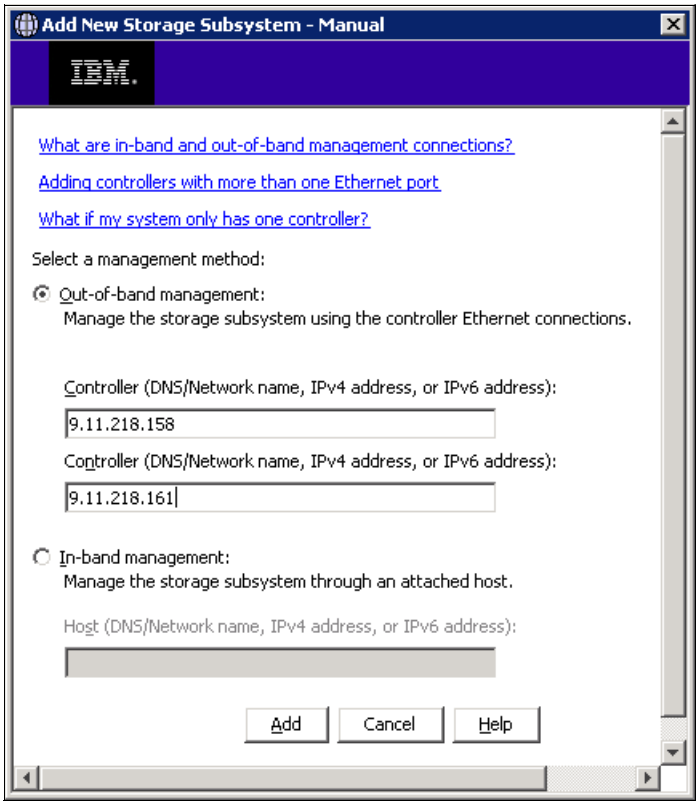


Figure 5-33 Add Storage Subsystem with reconfigured IP addresses

The DS3000 will be added to the Enterprise Management window using the new IP addresses.

This completes the basic configuration of DS3000 storage subsystem.

6

Installing IBM System Storage DS3000 Storage Manager

In this chapter, we describe how to install the IBM System Storage DS3000 Storage Manager on Microsoft Windows and Linux. This chapter is organized in the following sections:

- 6.1, “Installing DS3000 Storage Manager on Microsoft Windows 2003 and 2008” on page 82
- 6.2, “Installing Storage Manager on Linux” on page 90
  - 6.2.2, “Installing or upgrading Storage Manager using the GUI” on page 91
  - 6.2.3, “Installing Storage Manager using a text console” on page 96

The sections here describe installing DS3000 Storage Manager on a system that is used only for management purposes. Depending on the usage of the system, additional Storage Manager components have to be installed. For more details about Storage Manager components and their use, see 3.4.1, “DS3000 Storage Manager components” on page 27.

Note: The file names shown in this chapter are only for a basic illustration of the steps to be followed, and they will vary depending on the latest version of the Storage Manager available for download.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

6.1 Installing DS3000 Storage Manager on Microsoft Windows 2003 and 2008

This section describes how to install DS3000 Storage Manager on a Windows 2008 system.

Note: the installation on a Windows 2003 and 2008 server is same. See the installation procedure and configuration procedure in 5.2, “DS3000 Storage Manager installation” on page 59.

Depending on the usage of the system, different components need to be installed, as described in 3.4.1, “DS3000 Storage Manager components” on page 27.

6.1.1 Installation preparation

Follow these steps:

1. Download the Storage Manager installation package from the IBM Support Web site. See Appendix B, “IBM Support Web site” on page 733 for information about how to get downloads from the IBM Support site. You will see a 32-bit (x86) and 64-bit (x86-64) version; make sure to download the package corresponding to the Windows version of the host or management station. For more details about how to download DS3000 fixes, code, and documentation, see Appendix B, “IBM Support Web site” on page 733.
2. Unpack the ZIP archive into a local directory. The ZIP archive contains:
   – A README file with information about the package and the installation
   – A change history of the package
   – The installer file in the format SMIA-WS<32|64>-*xx.xx.xx.xx*.exe, where the Xs represent the version information
   – A subdirectory with MIB files
3. Decide on the installation type and proceed with the installation.

The above steps are applicable for Manager and Client options.

6.1.2 Installing or upgrading the Storage Manager Client on Microsoft Windows 2003 or 2008

Follow these steps:

1. Run the installer file; our version is SMIA-WS32-03.35.35.11.exe. The files unpack, as shown in Figure 6-1 on page 83.

Note: The file name shown above is only for illustration purposes and it may vary depending on the latest version of the Storage Manager available for download.

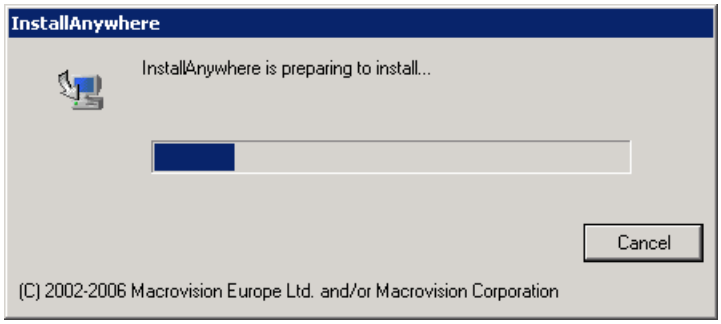


Figure 6-1 Preparation

2. Select the language to display the license agreement, and click **OK** (see Figure 6-2).

Note: The installation program will always use English; the selected locale only affects the license agreement.



Figure 6-2 Select language

3. Read the introduction and click **Next**, as shown in Figure 6-3.

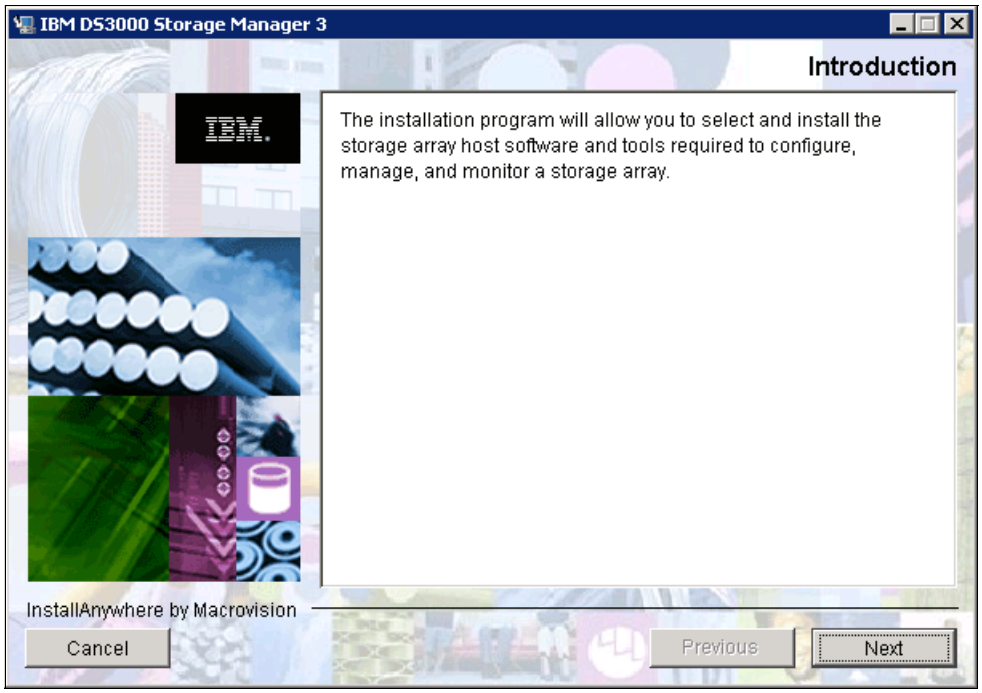


Figure 6-3 Introduction

4. Read the copyright statement and click **Next (**Figure 6-4).



Figure 6-4 Copyright Statement

5. Read the license agreement and click **Next** (Figure 6-5 on page 85).

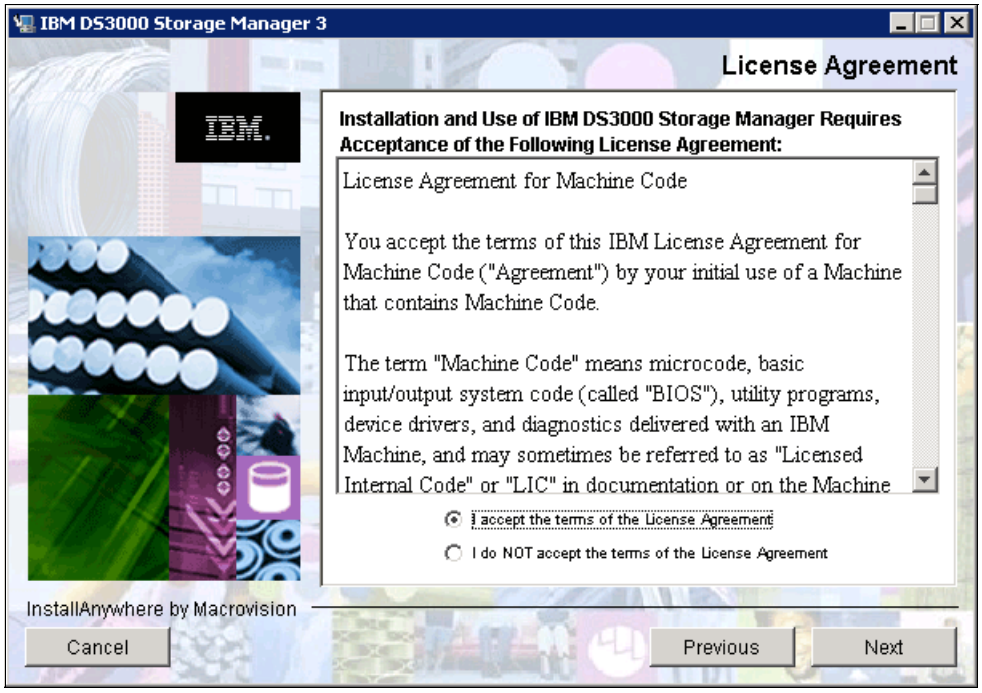


Figure 6-5 License Agreement

6. Specify the directory where you want to install, as shown in Figure 6-6. The default is C:\Program Files\IBM_DS3000.

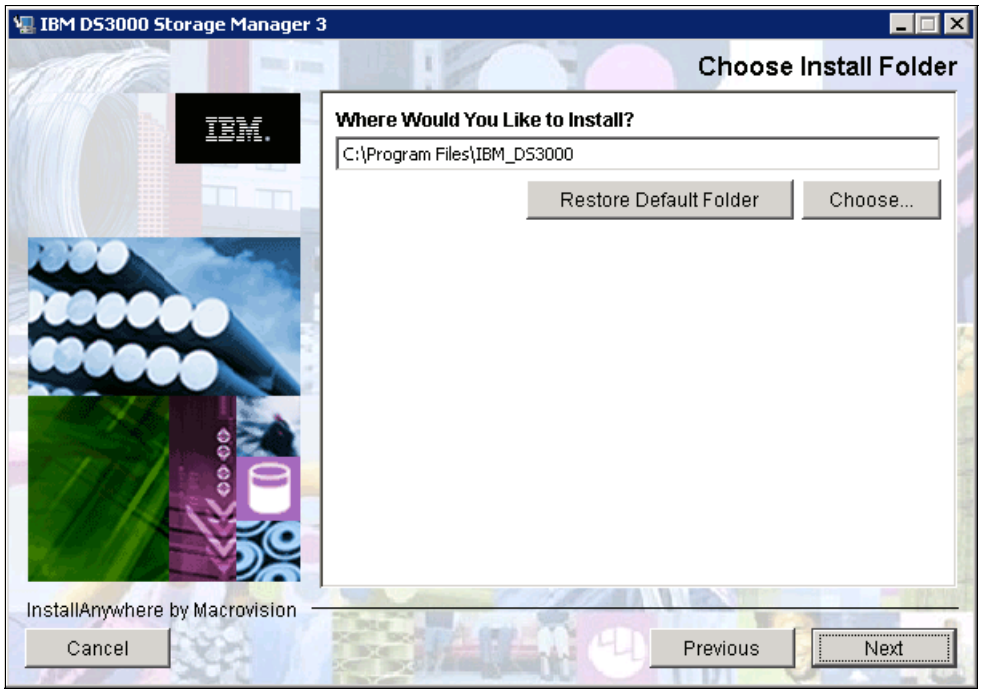


Figure 6-6 Select installation directory

7. For Installation Type (Figure 6-7), select **Custom** and click **Next**. We recommend the Custom installation because you can see what components of Storage Manager will be installed in each Install Set (Typical, Management Station, or Host) and you can also modify them.

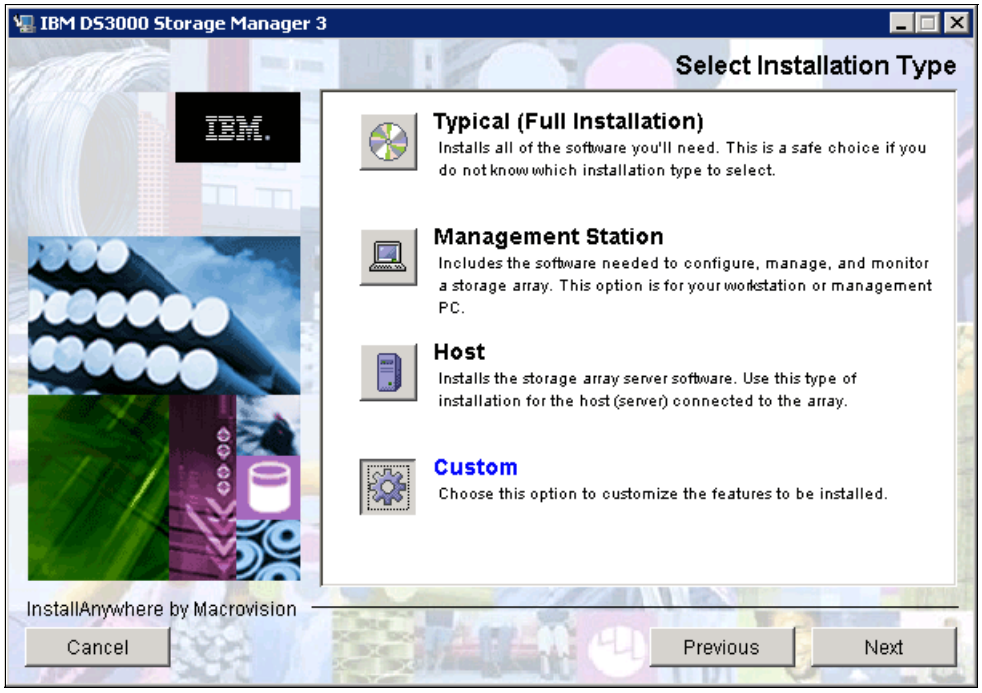


Figure 6-7 Select Installation Type

8. In Figure 6-8 on page 87, select **Install Set** from the drop-down menu. The same options shown in Figure 6-7 appear. When you select an Install Set, you can see what components will be installed, and modify them if you wish. See 3.4.1, “DS3000 Storage Manager components” on page 27 for a description of each component. Our example shows the client selection because the system being installed is for management purposes and the other components are not required for it.



Figure 6-8 Select product features to install

9. If a version of Storage Manager is already installed, you will see the warning shown in Figure 6-9. You can select to overwrite the existing installation or to cancel at this point. Click **OK** to continue. If you do not have Storage Manager already installed, you do not see this warning.

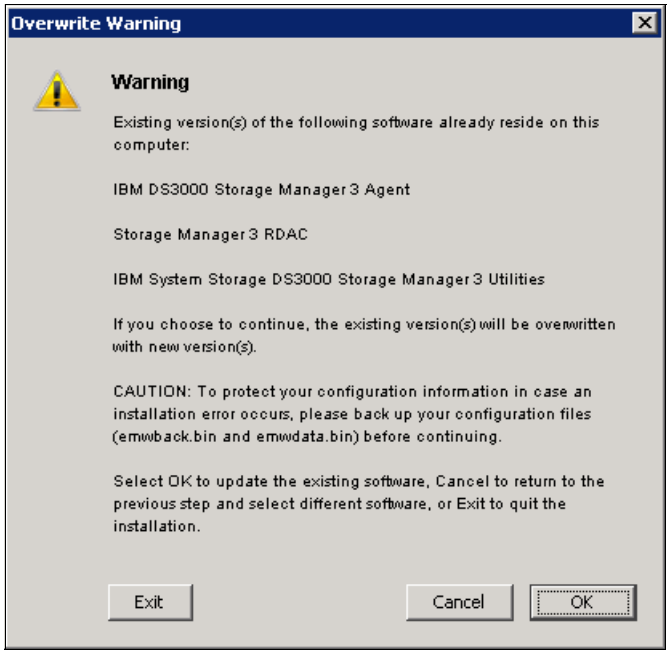


Figure 6-9 Overwrite Warning

10. In Figure 6-10, you can select whether to automatically start the Storage Manager Monitor.

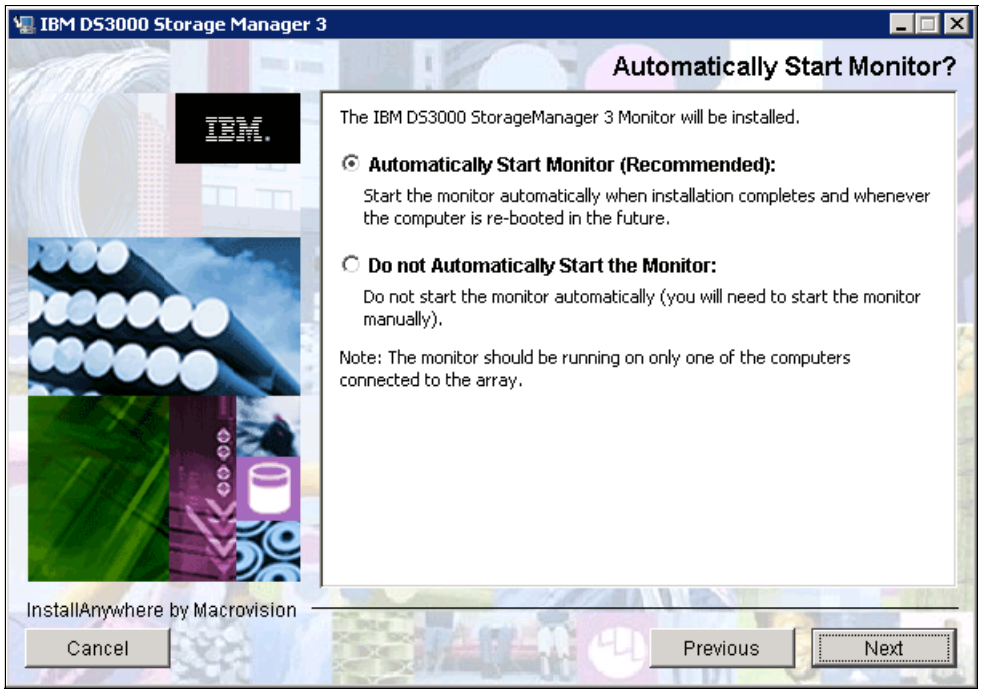


Figure 6-10 Automatically Start Monitor

11. Verify the pre-installation summary and click **Install**. See Figure 6-11.

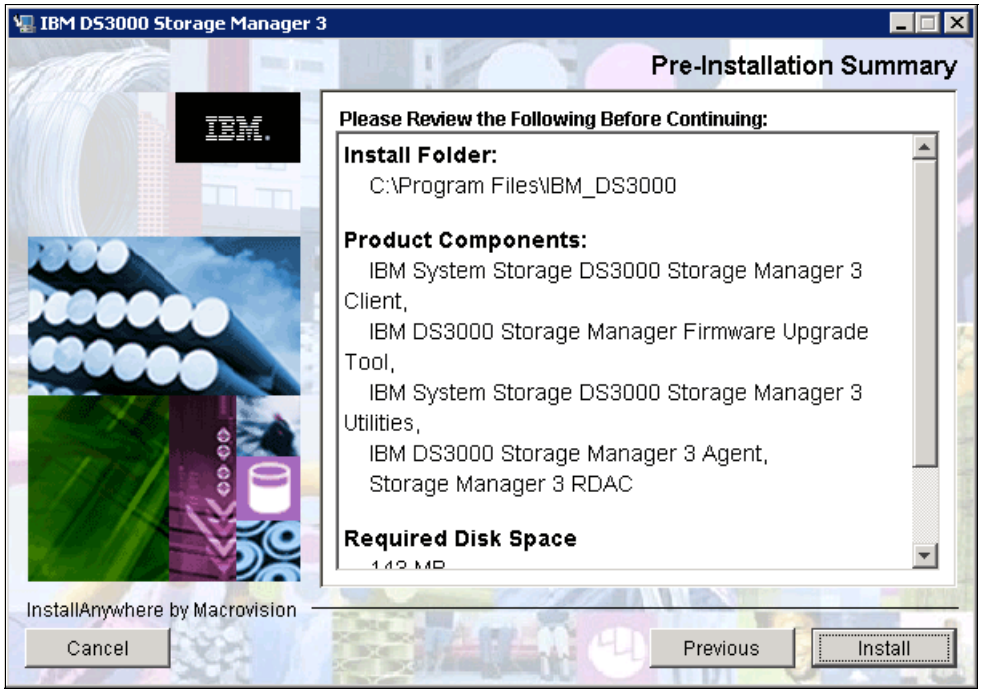


Figure 6-11 Pre-installation Summary

12. The installation status window will appear during the installation, as shown in Figure 6-12.

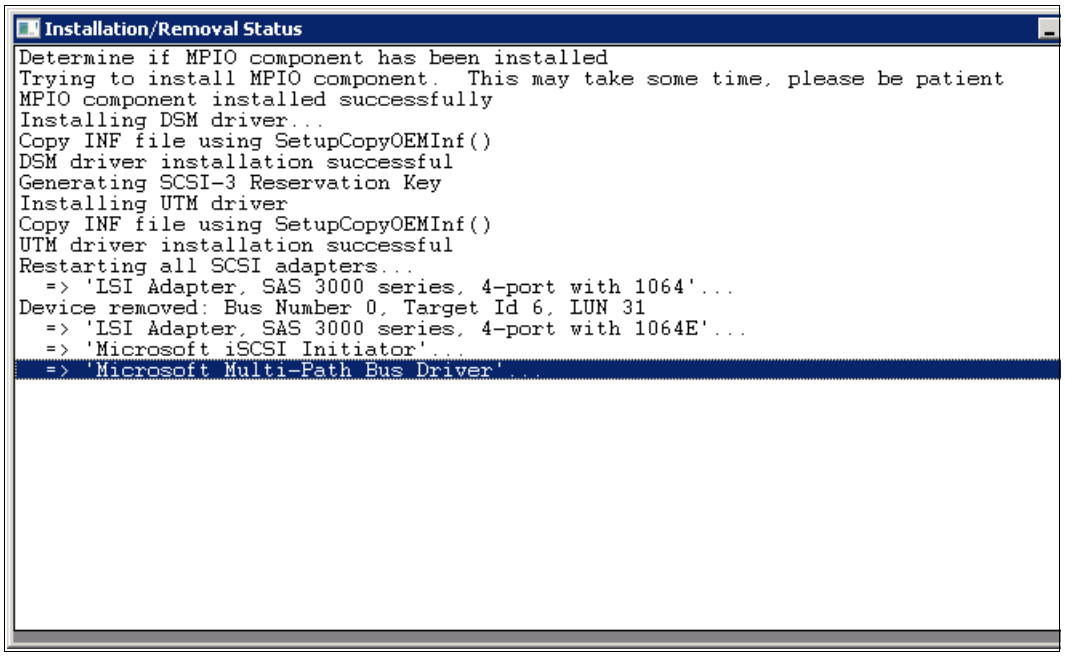


Figure 6-12 Installation status window

After the installation has finished, verify that there were no errors and click **Done**. See Figure 6-13.

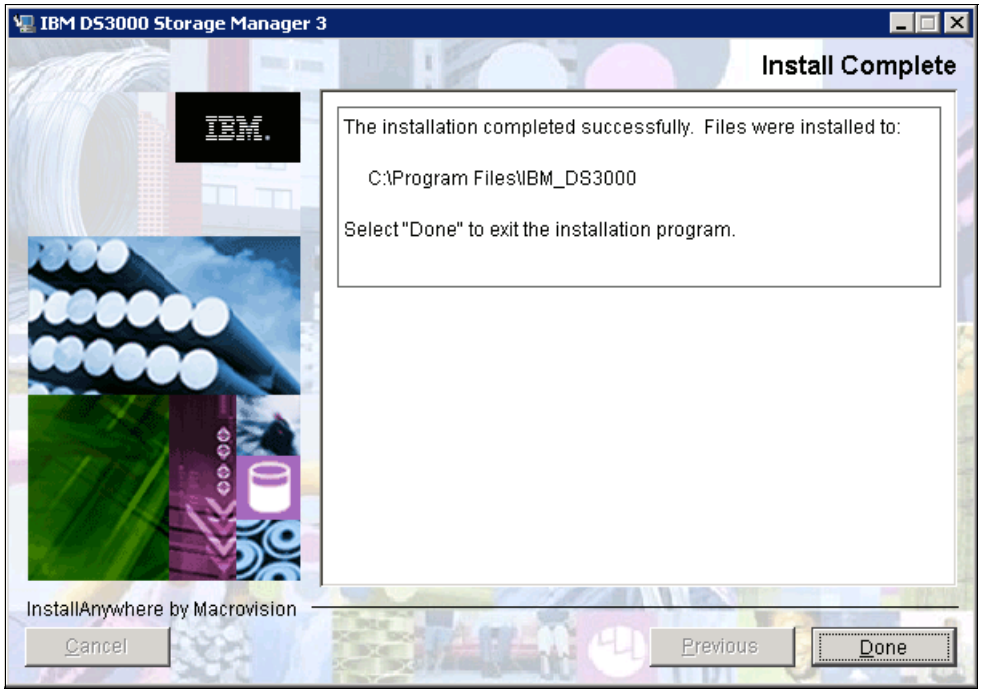


Figure 6-13 Install Complete

6.2 Installing Storage Manager on Linux

Installing Storage Manager on Linux is similar to Windows, with one major difference: The failover driver known as RDAC is not included in the Storage Manager Linux installation package; it must be installed separately from the Storage Manager package.

Storage Manager on Linux can be installed using a GUI or a text-based interface. Both installation variants are described here. The first steps are common to all methods:

- 6.2.1, “Preparing for install” on page 90
- 6.2.2, “Installing or upgrading Storage Manager using the GUI” on page 91
- 6.2.3, “Installing Storage Manager using a text console” on page 96

6.2.1 Preparing for install

To prepare for the installation, do the following:

1. Download the latest version of the DS3000 Storage Manager Client to the management station or host connected to the DS3000. See Appendix B, “IBM Support Web site” on page 733 for information about getting downloads from the IBM Support site.
2. Extract the archive to a local directory, as shown in Example 6-1.

Example 6-1 Extract package archive

```
# tar xvfz ibm_sw_ds3k_0335A5xx_linux2.6_32-64.tgz
SMIA-Linux-03.35.A5.xx.bin
ibm_sw_ds3k_0335A5xx_linux2.6_32-64.txt
ibm_sw_ds3k_0335A5xx_linux2.6_32-64.chg
SM2MIB/README_First.txt
SM2MIB/SM2.MIB
#
```

3. The following files are unpacked:
   - The installation package and executable
   - The readme file that describes the package
   - The change history of the package
   - A subdirectory SM2MIB that contains the MIB files

   The directory contents are shown in Example 6-2.

Example 6-2 Directory contents

```
# ls -R
.:
SM2MIB
SMIA-Linux-03.35.A5.xx.bin
ibm_sw_ds3k_0335A5xx_linux2.6_32-64.txt
ibm_sw_ds3k_0335A5xx_linux2.6_32-64.chg

./SM2MIB:
README_First.txt  SM2.MIB
#
```

4. The file with a.bin extension is the installer that contains also the installation files. By default, the installer requires a graphical console to launch, but can be also installed on a text console or left unattended.

Note: The command to install the installable is **$SMIA-LINUX-*xx.xx-xx.xx*.bin -i *options***.

Execute this file with the option -? to display the help on the available options to install Storage Manager.

The options are:

- awt: A GUI installation that allows you to select the language of the license agreement. Note that only the license agreement is displayed in the selected language; the installation itself always uses English.
- swing: See awt. It uses a different user interface.
- console: Use a text based console for the installation.
- silent: Perform an unattended installation.

5. Choose which method you want to use, and launch the install file, specifying your option with the -i parameter. For example:

```
$SMIA-LINUX-03.35-A5.11.bin -i console
```

6.2.2 Installing or upgrading Storage Manager using the GUI

Do the following steps:

1. Log on to the Linux server with administrative rights, that is, usually the root user, on a graphical console.
2. Start the installer that was extracted in step 2 on page 90 in a terminal window.

3. Select the language to display the license agreement and click **OK** (Figure 6-14).

Note: The installation will always be performed in English.



Figure 6-14 Installation

4. Read the introduction and click **Next** (Figure 6-15).



Figure 6-15 Introduction

5. Read the copyright statement and click **Next** (Figure 6-16).



Figure 6-16 Copyright Statement

6. Read and accept the license agreement. Click **Next** (Figure 6-17).

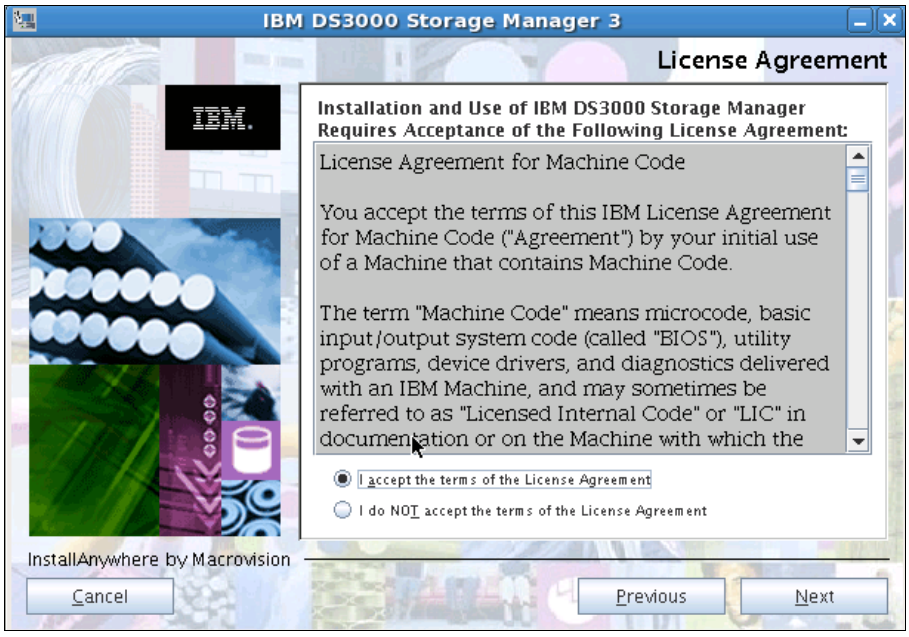


Figure 6-17 License Agreement

7. Select installation type **Custom** and click **Next** (see Figure 6-18). We recommend the Custom installation because you can see what components of Storage Manager will be installed in each Install Set (Typical, Management Station, or Host) and you can also modify them.



Figure 6-18 Installation type

8. In Figure 6-19, select the Install Set to be chosen (**Management Station**, in our case). You will then see which Storage Manager components will be included. Click **Next**.



Figure 6-19 Select Product Features

9. If a version of Storage Manager is already installed, you will see the warning in Figure 6-20. You can select to overwrite the existing installation or to cancel at this point. If you do not have Storage Manager already installed, you do not see this warning. Click **OK** to remove the old version and proceed with the installation.

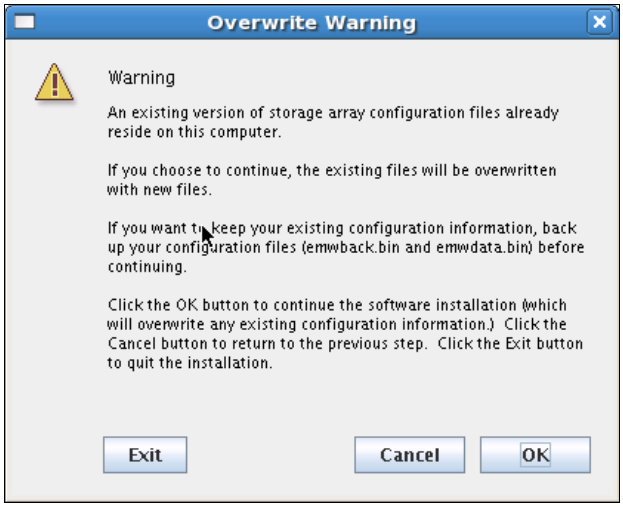


Figure 6-20 Existing Storage Manager detected

10. Verify the installation options and click **Next**. See Figure 6-21.



Figure 6-21 Pre-installation Summary

11. After a successful installation, you will see messages similar to Figure 6-22. If there are no errors, click **Done** to exit.

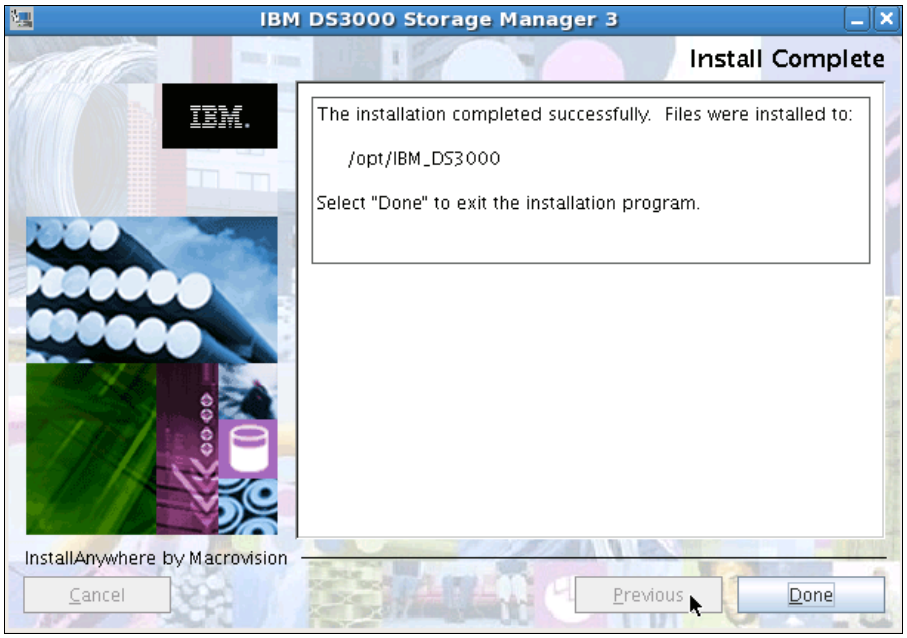


Figure 6-22 Installation Complete

12. Optionally, verify the installation on the Desktop by finding the icon shown in Figure 6-23.



Figure 6-23 SM Manager on the Linux desktop

6.2.3 Installing Storage Manager using a text console

If your Linux workstation does not have a graphical console, you can install IBM DS3000 Storage Manager on a text console. After installing the Storage Manager Agent, the system can be managed from any host that has the client installed and a network connection to this server.

Do these steps:

1. Start the installer with the option -i console. Select the locale (the default is English), and press Enter. See Example 6-3 on page 97.

Example 6-3 Text based installation

```
# sh SMIA-LINUX-03.35-A5.11.bin -i console
Preparing to install...
Extracting the JRE from the installer archive...
Unpacking the JRE...
Extracting the installation resources from the installer archive...
Configuring the installer for this system's environment...
awk: cmd. line:6: warning: escape sequence `\.' treated as plain `.'

Launching installer...

Preparing CONSOLE Mode Installation...

===============================================================================
Choose Locale...
----------------

    1- Deutsch
  ->2- English
    3- Español
    4- Français
    5- Italiano
    6- Português  (Brasil)

CHOOSE LOCALE BY NUMBER: 2
```

2. Read the introduction and press Enter to proceed.
3. Read the copyright statement, and select Y to accept it.
4. In Example 6-4, choose the type of installation to perform. A full installation installs the Storage Manager Client, Utilities, and Agent. The option Management Station will install the Storage Manager Client and Utilities and the option Host will install only the Storage Manager Agent. We select 4 for Custom and press Enter.

Example 6-4 Installation type

```
Select Installation Type
------------------------

Choose the Install Set to be installed by this installer.

  ->1- Typical (Full Installation)
    2- Management Station
    3- Host

    4- Customize...

ENTER THE NUMBER FOR THE INSTALL SET, OR PRESS <ENTER> TO ACCEPT THE DEFAULT
   :
```

5. Because we selected Custom, we can select or deselect individual components. By default, all components will be installed. Enter the number of any component that is not needed for you particular installation. In Example 6-5, we select 3 and press Enter. This deselects the Storage Manager Agent.

Example 6-5 Choose Product Features

```
Choose Product Features
-----------------------

ENTER A COMMA_SEPARATED LIST OF NUMBERS REPRESENTING THE FEATURES YOU WOULD
LIKE TO SELECT, OR DESELECT. TO VIEW A FEATURE'S DESCRIPTION, ENTER
'?<NUMBER>'.  PRESS <RETURN> WHEN YOU ARE DONE:

    1- [X] IBM System Storage DS3000 Storage Manager 3 Client
    2- [X] IBM System Storage DS3000 Storage Manager 3 Utilities
    3- [X] IBM System Storage DS3000 Storage Manager 3 Agent

Choose the Features to be installed by this installer.
   :
```

6. Confirm the pre-installation summary and press Enter. See Example 6-6.

Example 6-6 Storage Manager - Pre-installation summary

```
Pre-Installation Summary
------------------------

Review the Following Before Continuing:

Install Folder:
    /opt/IBM_DS3000

Product Components:
    IBM System Storage DS3000 Storage Manager 3 Client,
    IBM System Storage DS3000 Storage Manager 3 Utilities

Required Disk Space
    105 MB

Available Disk Space
    2,661 MB



PRESS <ENTER> TO CONTINUE:
```

7. The installation is now performed into the directory /opt/IBM_DS3000; you cannot change this directory. You will see a message similar to Example 6-7 on page 99 when the installation process is complete.

Example 6-7 Installation complete

```
Installation Complete
---------------------

Congratulations. IBM_DS3000 has been successfully installed to:

   /opt/IBM_DS3000

PRESS <ENTER> TO EXIT THE INSTALLER:
#
```

7

Host configuration

We have already described how to configure logical drives and assign them to host servers in Chapter 5, “Basic hardware configuration” on page 49. Now we need to configure the host servers so that they will be able to access their respective logical drives. Basically, we need to install HBAs, configure them, and add their drivers in the operating system. Finally, we also have to install multipath support when using redundant paths between our host server and the DS3000.

In this chapter, we are going to learn about the following topics:

- 7.1, “HBA configuration and driver installation” on page 102
- 7.2, “Multipath support” on page 104
  - 7.2.1, “Microsoft MPIO” on page 104
  - 7.2.2, “Linux RDAC” on page 105
  - 7.2.3, “AIX MPIO” on page 106
  - 7.2.4, “IBM System i with VIOS” on page 107

© Copyright IBM Corp. 2007, 2009. All rights reserved.

7.1 HBA configuration and driver installation

Each host server needs at least one HBA so that a physical connection to the DS3000 can be established. When connecting to a dual-controller DS3000, the host server would typically have two HBAs installed. The type of HBA depends on the DS3000 subsystem:

- For DS3200 connections, use SAS HBAs.
- DS3300 will require iSCSI HBAs (or iSCSI software initiators) to be installed in the host servers.
- To attach to a DS3400, use FC HBAs.

We provide many details of configuring HBAs in Part 4, “Sample configurations” on page 389.

Each of these HBA types typically has a ROM-based configuration utility; use this utility to configure the basic HBA settings. A sample window from the QLogic® Fast!UTIL utility is shown in Figure 7-1.

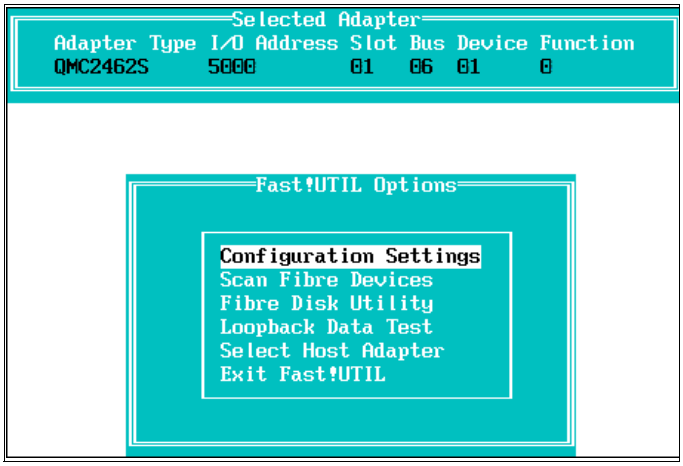


Figure 7-1 QLogic Fast!UTIL

Alternatively, you can also use an operating system based HBA configuration utility, for example, SANSurfer for QLogic based HBAs (see Figure 7-2 on page 103 and Figure 7-3 on page 103) or Emulex HBAnyware.

These utilities can be used to determine the FC HBA WWN or set the iSCSI HBA IQN. The HBA identifier (WWN or IQN) is very important; you will need it when assigning logical drives access to the host servers. The host server definition in DS3000 Storage Manager is based on these HBA identifiers. HBA identifiers are the sole criteria used by the DS3000 subsystem to allow or prevent host access to logical drives.

Other functions of these utilities include setting connection parameters, scanning for attached devices, running diagnostic tests, and so on.

The utilities also allow you to enable or disable the HBA BIOS. The BIOS normally does not need to be enabled; we actually recommend disabling it, unless you are booting from the DS3000 subsystem. The HBA BIOS must be enabled for this to work.

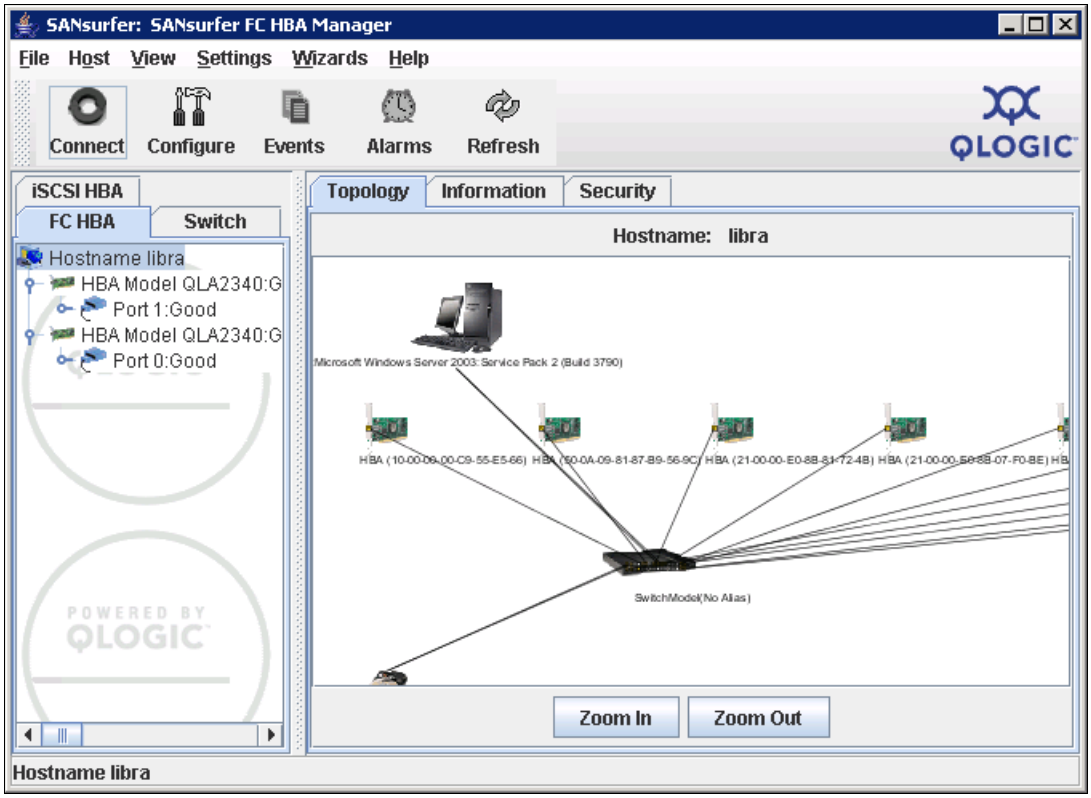


Figure 7-2 QLogic SANsurfer® - Host View

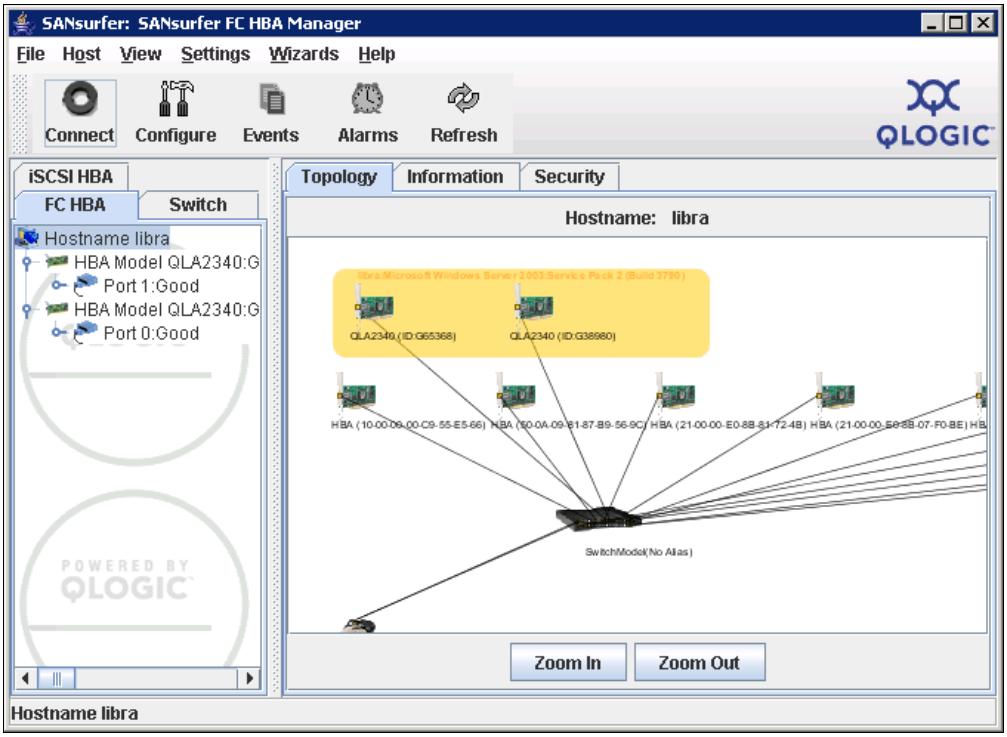


Figure 7-3 SANsurfer HBA View

Next, you will have to install the correct driver for your HBA in the operating system. You can find the latest drivers on IBM Systems support Web site at:

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

We provide some examples of HBA driver installation in this book, including:

- Linux driver for LSI Logic SAS HBA (See 17.2.1, “SAS HBA driver installation” on page 407.)
- Linux driver for QLogic 2462 FC HBA (See 26.6.2, “Updating the QLA2400 driver (RHEL4)” on page 660.)
- Windows driver for LSI Logic SAS FC HBA (See 16.2.1, “SAS HBA driver installation” on page 393.)
- QLogic iSCSI HBA (See 21.4, “Installing the iSCSI HBA driver” on page 485.)
- Windows driver for QLogic FC HBA (See 25.2.1, “Installing HBA drivers” on page 614.)

Download the appropriate driver for your particular HBA type and carefully follow the instructions in the readme file.

7.2 Multipath support

If you use two HBAs for attachment to a dual-controller subsystem, you will have to install multipath support to allow the failover procedures to function. If a controller or path fails, the host server can continue to access the logical drives through the alternate path.

If you are using the iSCSI software initiator under Windows or Linux with access to more than one iSCSI port, you will also have to install RDAC/MPP/MPIO.

On Windows Server 2003 and Windows 2008 hosts, we use Microsoft MPIO for multipath support (see Figure 7-4 on page 105). On Linux, the open-source RDAC solution is required. For AIX, MPIO is the preferred method for attachment. In the following sections, we take a closer look at these three implementations. Additional details and examples can be found in Part 4, “Sample configurations” on page 389.

7.2.1 Microsoft MPIO

Microsoft MPIO is a framework that provides integration of various storage subsystem vendors’ multipath solutions into the Windows architecture. MPIO usually interacts with a Device-Specific Module (DSM) that should be provided by the storage subsystem vendor. The idea is to offer a standardized way for storage vendors to integrate their hardware solutions with Windows, rather than each vendor having to develop their own proprietary multipath implementation.

One of many examples of the configuration of MPIO in Microsoft Windows is described in 24.4.3, “Installing multipath driver” on page 600.

Microsoft MPIO consists of the following three drivers:

- mpdev.sys
- mpio.sys
- mpfltr.sys

As stated before, we also need the DSM, which supports DS3000. The DS3000 DSM is installed as a part of the DS3000 Storage Manager host server package.

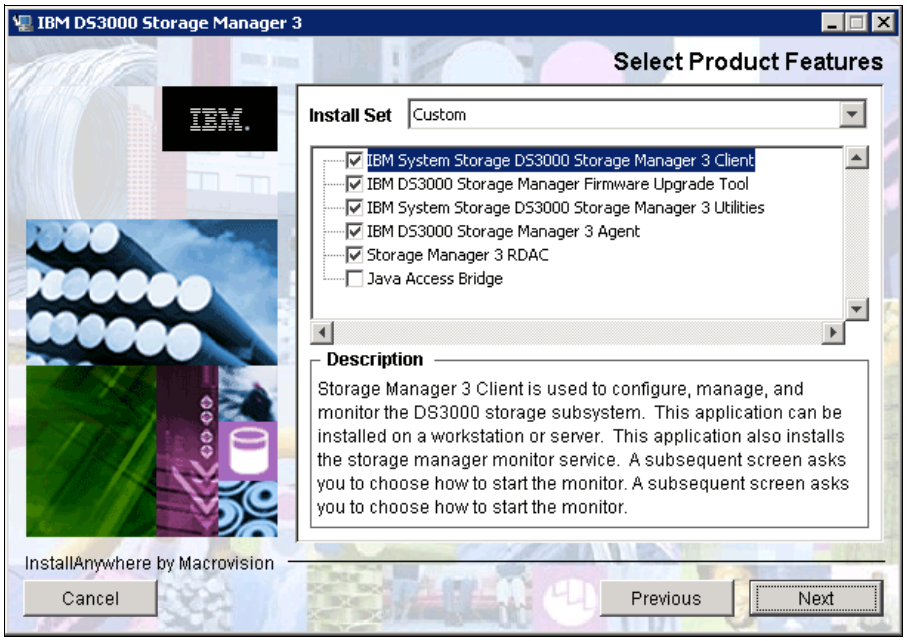


Figure 7-4 Storage Manager RDAC (consists of Microsoft MPIO and DS3000 DSM)

As you can see in Figure 7-4, the multipath support component is referred to as Storage Manager 3 RDAC in the installation wizard. If you are already familiar with the IBM System Storage DS4000, then you know that the DS4000 RDAC is a proprietary multipath solution specific to DS4000. It does not use Microsoft MPIO.

The DS3000 Storage Manager RDAC, on the other hand, is a very different implementation altogether. It consists of Microsoft MPIO and the DS3000 DSM.

The DS3000 DSM consists of the following two drivers:

- ds4dsm.sys
- ds4utm.sys

The driver names suggest that this DSM solution supports DS4000 storage subsystems as well. This is indeed true; DS3000 DSM actually does support both the DS3000 and DS4000 storage subsystems. But remember, this is not the same as the proprietary DS4000 RDAC package.

Now let us summarize:

- DS4000 subsystems can work either with the proprietary DS4000 RDAC or with the Microsoft MPIO compatible DSM.
- DS3000 subsystems should use the Microsoft MPIO compatible DSM only. Incidentally, the DS3000 DSM is called RDAC, but has nothing in common with the DS4000 RDAC.

7.2.2 Linux RDAC

Linux RDAC is not included with DS3000 Storage Manager for Linux; it is a separate package, also known as the MPP driver. It consists of two parts: *mppUpper* and *mppVhba* (virtual HBA).

An example of a configuration with RDAC in a Linux environment is shown in 26.7, “Installing Linux RDAC in RHEL5” on page 663.

You can download the RDAC package from the following Web site:

http://www.lsi.com/rdac/ds3000.html

Alternatively, you can start by visiting the IBM Systems support Web site; however, you will eventually be redirected to the URL listed above.

There are two different Linux RDAC packages: one for kernel Version 2.4 and the other for Version 2.6. (The format would be IBM DS3000 Linux 2.x RDAC v*xx.xx.xx.xx*.)

The packages each contain a compressed tar file of the source code for the RDAC driver. Use the package to build the RDAC kernel driver by carefully following instructions in the readme file.

As the multipath and failover functionality will be provided by RDAC, it is important that you use the non-failover driver with your HBA. Linux RDAC cannot coexist with failover HBA drivers.

You should also make sure each HBA in the host server only sees one DS3000 RAID controller. If any individual HBA has access to both controllers, then the RDAC driver will not work properly. To prevent this, make sure you implement correct zoning on the SAN switch (in the case of FC HBAs and DS3400).

Another important consideration is the order of LUNs assigned through host-to-logical-drive mapping. The LUNs assigned to a Linux host must be a contiguous set of numbers, because, by default, the Linux kernel does not detect so-called *sparse LUNs*, that is, it does not scan any LUNs after a skipped number. Therefore, no LUNs after a skipped number would be available to the host server.

Never assign the access logical drive to LUN 0, as this can prevent access to other mapped logical drives. The access logical drive should be assigned to LUN 31.

You need to make sure that the HBA driver is installed and the DS3000 subsystem attached correctly *before* you install the Linux RDAC driver.

The Linux RDAC installation steps for SAS attached DS3200 are also provided in 17.4, “Installing RDAC for Linux” on page 408.

Note for DS3300: If using the iSCSI Software Initiator, MPP with failover support is not supported if there are HBAs of any kind (LSI/QLogic/Emulex) also installed in the host. At the time of writing, when installing the MPP for Software Initiator, no HBAs must be available. See Chapter 22, “iSCSI configuration 2 - Software Initiator on Red Hat Linux V5.2” on page 529 for more information.

7.2.3 AIX MPIO

Since the beginning of 2008, the DS3400 Storage Subsystem is also supported in an IBM System p environment. It can be used with AIX 5L V5.2 and V5.3, AIX V6.1, and also with VIOS V1.5.2. IBM Power Blades JSxx with Linux on Power also support DS3200 and DS3300. In all the cases, the attachment requires that you buy the AIX/VIOS Host Attach license (P/N 13N1923) for DS3000.

More information regarding the attachment of DS3000 to AIX and VIOS can be found on the IBM Support Web site:

http://www.ibm.com/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-5073850&brandind=5000028

Also, verify all the readme files supplied with the DS3000 Storage Manager software and DS3000 controller firmware if you use AIX host attachment.

The MPIO drivers are already included in standard AIX installations, but verify that devices.common.IBM.mpio is installed.

The following operating system software versions are supported:

- AIX 5L V5.2: The operating system must be at Version 5200-10-04-0750 or later.
- AIX 5L V5.3: The operating system must be at Version 5300-07-03-0811 or later.
- AIX V6.1: The operating system must be at Version 6100-00-03-0808 or later.

After you do appropriate SAN zoning and after you assign DS3000 logical drives to the HBAs in the host server, as described in 9.2, “Configure hosts” on page 140, you can discover new devices in AIX by running the following command:

```
# cfgmgr -v
```

To verify the devices were added successfully, run:

```
# lsdev -Cc disk | grep DS3K
```

Another very useful command is:

```
# mpio_get_config -Av
```

This command shows detailed information regarding DS3000, data paths, and LUNs in the AIX operating system.

7.2.4 IBM System i with VIOS

As mentioned in 7.2.3, “AIX MPIO” on page 106, DS3400 is supported by VIOS V1.5.2. It enables you to use a DS3400 Storage Subsystem to save data from IBM System i® LPARs configured on POWER™ servers.

System i servers and the new operating system IBM i (formerly known as i5/OS®) can use DS3400 as a device connected to the IBM Virtual I/O Server (VIOS).

VIOS is an AIX-based software appliance that has been specifically developed for virtualizing I/O resources on IBM POWER5™ and POWER6™ systems. The Virtual I/O Sever is installed in its own logical partition (LPAR) and owns the physical I/O resources like Ethernet, SCSI, and FC HBA adapters. That devices are virtualized for other client LPARs to share them remotely using the built-in Hypervisor services on IBM POWER platform. Other LPARs can run AIX, Linux on Power, or IBM i operating systems.

As it is beyond the intended scope of this book to describe the configuration of the IBM i operating system for the DS3400, we recommend you refer to *IBM i and Midrange External Storage*, SG24-7668.

In this book, you can find examples and all the setup steps for using the IBM i operating system with DS3400.

Part 3

Administration

In this part of the book, we describe the administration of a IBM System Storage DS3000 server using the DS3000 Storage Manager.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

8

Administration - Enterprise

In this chapter, we describe the IBM System Storage DS3000 Storage Manager Enterprise Management window, including its functions and how to use it.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

8.1 Enterprise Management window overview

When you launch Storage Manager, you get the Enterprise Management window (shown in Figure 8-2 on page 113), and also, by default, the Task Assistant.

8.1.1 Task Assistant

Initially, the Enterprise Management Window Task Assistant also displays, as shown in Figure 8-1. The Task Assistant gives you a quick way to access common tasks.

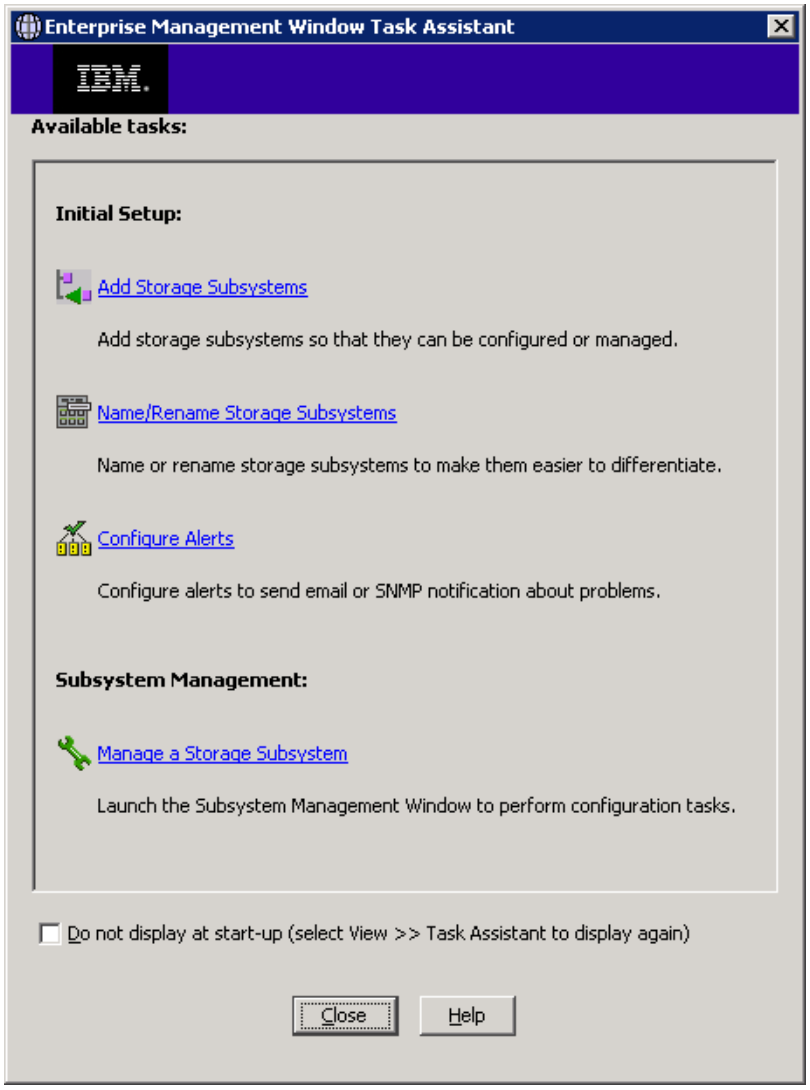


Figure 8-1 Task Assistant

You can access these common tasks from this window, or you can use the Storage Manager itself. Here is a list of the functions provided in Task Assistant, with information about where the use of these functions is described in this book:

- Initial Setup:
  - Add Storage Subsystems: See 5.3, “Add a DS3000 subsystem in Storage Manager” on page 68 or “Add Storage Subsystem to view” on page 127.

- Name/Rename Storage Subsystems: See 12.1.1, “Rename Storage Subsystem” on page 259.
- Configure Alerts®: See “Configure Alerts” on page 118.

- Subsystem Management:
  - Manage a Storage Subsystem: Starts the main Storage Manager application, as described in the following chapters.

If you do not want to see the Task Assistant each time you start Storage Manager, check **Don’t show the task assistant at start-up again**. Click **Close** to exit from Task Assistant. To start Task Assistant subsequently, you can:

- Select **View** → **Task Assistant** (see Figure 8-19 on page 125).
- Right-click your management host in the left pane and select **Task Assistant**.
- Click the Task Assistant Quick Access button near the top of the window (see “Display Task Assistant” on page 126).

8.1.2 Enterprise Management window

The Enterprise Management window (Figure 8-2) is the entry point to manage each DS3000 storage subsystem. We described how to add storage subsystems in 5.3, “Add a DS3000 subsystem in Storage Manager” on page 68. Once they have been added, they will appear every time you start Storage Manager.

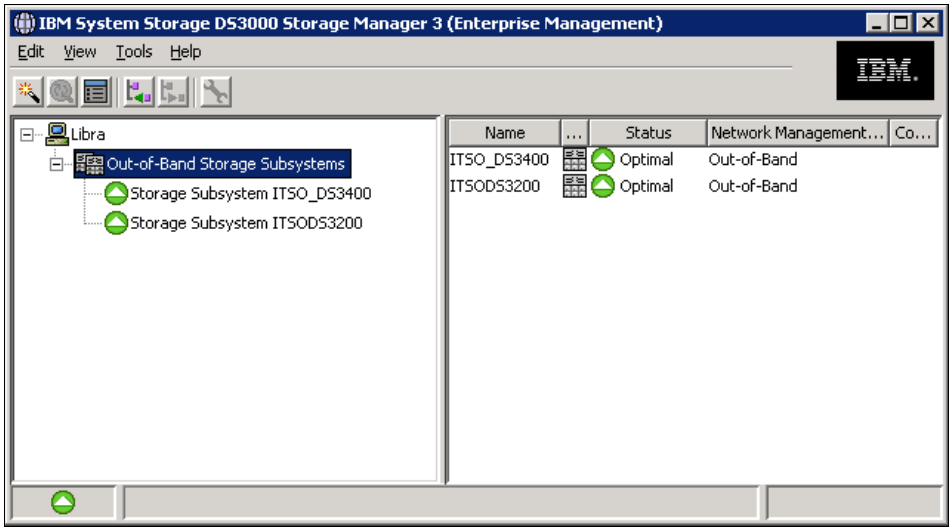


Figure 8-2 Enterprise Management window

The Enterprise Management window displays a list of all DS3000 storage subsystems that the client can access either directly or through the host agents. If you can access a certain storage server in both ways, and possibly through several host agents, you see it listed not just once, but many times in the Enterprise Management window.

Note: Although a single storage server could appear listed several times in the left pane when it is accessed by various host agents or directly attached, it only appears once in the right pane.

In the left window pane, you see your management station, and your managed subsystems. The subsystems are divided into two groups, In-Band Storage Subsystems and Out-of-Band Storage Subsystems. You can also see the status of your subsystems in this pane. If your subsystem appears green, the status is optimal. If you highlight the subsystem in the left part of the window, you also see a short summary about this system in the right window pane. If the subsystem appears red in this view, then it needs some attention. What you do in this case is described in Chapter 14, “Administration - Support” on page 299.

Figure 8-3 shows the various status icons for a storage subsystem. Table 8-1 explains the meaning of each icon.

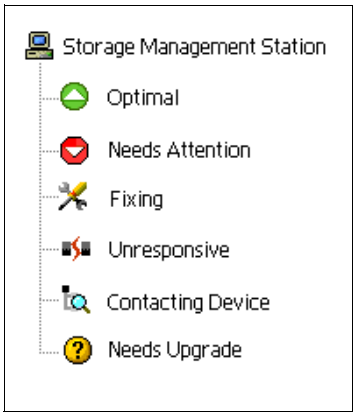


Figure 8-3 Status icons

Table 8-1 Status icons and their meanings

Status	Indicates
Optimal	Every component in the managed storage subsystem is in the desired working condition.
Needs Attention	There is a problem with the managed storage subsystem that requires intervention to correct it.
Fixing	A Needs Attention condition has been corrected and the managed storage subsystem is currently transitioning to an Optimal state.
Unresponsive	The storage management station cannot communicate with the storage subsystem, or one or both controllers in the storage subsystem.
Contacting Storage Subsystem	The Enterprise Management window has started an the storage management software is establishing contact with the storage subsystem.
Needs Upgrade	The storage subsystem is running a level of firmware that is no longer supported by the storage management software.

8.2 Functions in the Enterprise window

Next, we describe the various functions that are available from the Enterprise window.

8.2.1 Subsystem context menu

Right-click one of the subsystems to see the context menu with the tasks shown in Figure 8-4.

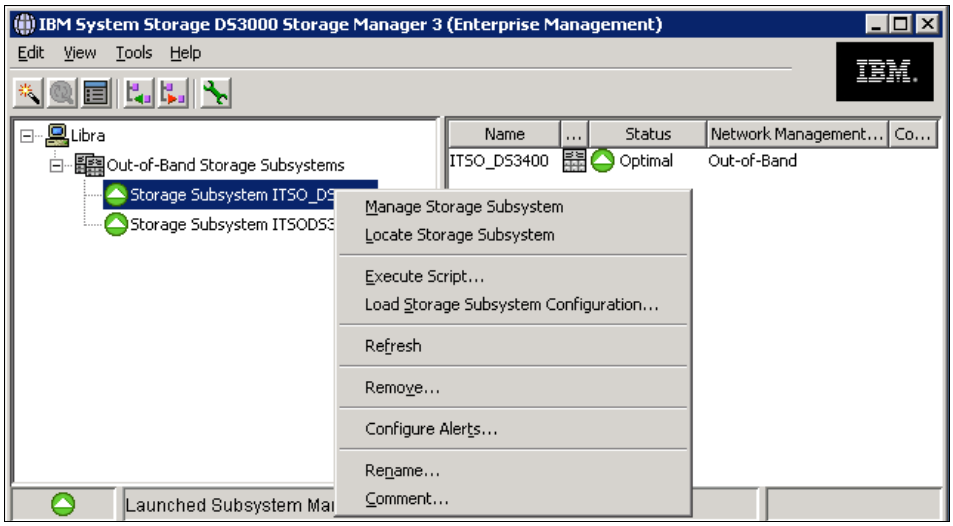


Figure 8-4 Context menu

Manage Storage Subsystem

This opens the Subsystem Management window. You can also double-click the system to manage it.

Locate Storage Subsystem

If you have several storage subsystems installed, it can be hard to locate the right one. If you select this option, the indicator light on the selected subsystem starts flashing (Figure 8-5).

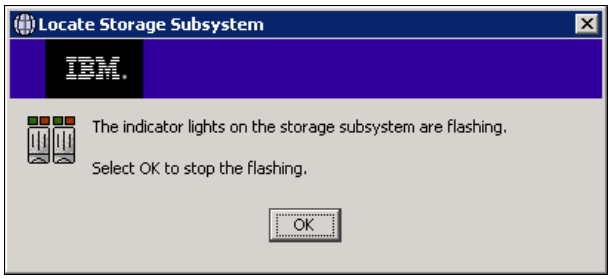


Figure 8-5 Indicator lights are flashing

Execute Script

Here you can open the DS3000 Script Editor. This Editor can be a very effective and useful tool to manage your subsystem. For more information about how to use it, see “Script Editor” on page 748.

Load Storage Subsystem Configuration option

In “Sample command: Save configuration script file” on page 760, we describe how to save a configuration script file based on the current environment. This is done on the CLI using the **`save StorageSubsystem configuration`** command. With the Load Storage Subsystem Configuration option in Storage Manager, you can load a saved file so that the configuration is applied to another storage system

This option will append the uploaded configuration to your existing one unless you do either of the following operations:

- Reset the storage subsystem’s configuration before loading the new one.
- Edit the saved script so that it contains the **`reset configuration`** command.

To upload a configuration file, do the following steps:

1. Select this option, and the load configuration notes will appear. Read them carefully, and then click **Yes**.
2. The Script Editor will open and you will be prompted to select the configuration file. Select your configuration file and click **OK**.
3. Now the file gets loaded into the Script Editor and the Confirm Load Configuration window appears:
   a. If you want to append a configuration to the existing one, click **Execute**, and the script will append the uploaded configuration to the existing one.
   b. If you want to restore your subsystem’s configuration before the new one is uploaded, select **Edit** and look in the Script Editor for the following lines and uncomment them, as shown in Figure 8-6. After uncommenting these two lines, select **Tools** → **Verify** → **Execute** from the Menu bar to execute the script and upload the new configuration.

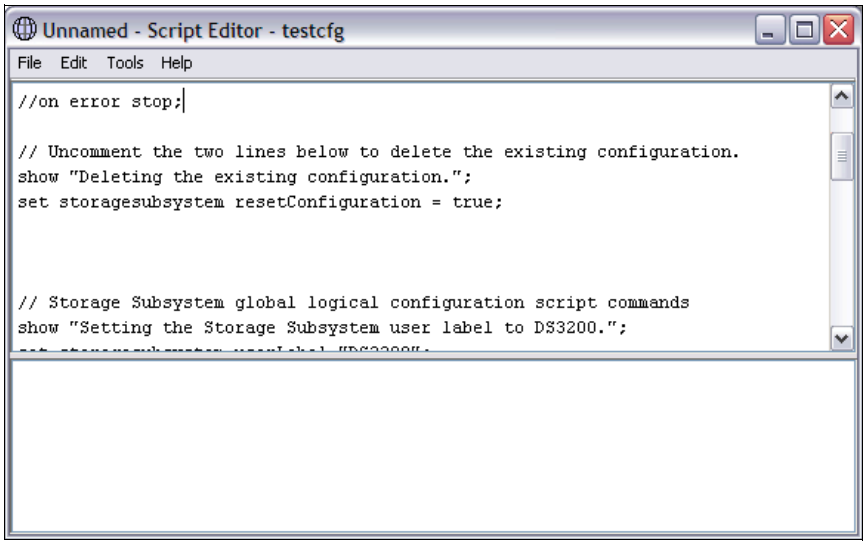


Figure 8-6 Uncomment lines to delete the existing configuration before uploading a new one

Note: We recommend that you use the load configuration file option for new DS3000s that you want to configure similarly to an existing DS3000. For example, if you have a number of storage subsystems with similar configurations, you can configure one, save the configuration file, modify it if necessary (for example, to use a different storage subsystem, or define different hosts), then use it to quickly configure the remaining storage subsystems.

You could also use this option to re-create replacement DS3000s if your original hardware was destroyed in a disaster..

Existing configurations: The script will not overwrite an existing configuration, that is, the script will fail if it attempts to build an array using drives that are already part of an existing array.

However, if you have lost an array that had data stored on it, do not reapply the configuration using this script; it will initialize the drives and delete the data. Instead, you should contact IBM technical support to see if your data can be recovered.

Refresh

Select this option to refresh the view of the Enterprise Management window (Figure 8-7). You might use this if a system's status has changed, for example, after an error recovery.

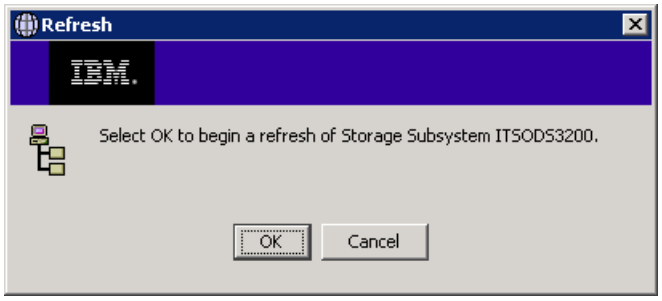


Figure 8-7 Refresh View

Remove

Here you can remove the selected storage subsystem from your view. This will not affect the views of any other storage management stations. To remove a storage subsystem, select this option and click **OK** (Figure 8-8).

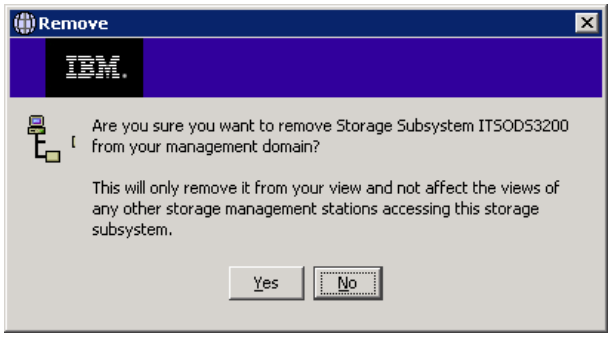


Figure 8-8 Remove Subsystem

Configure Alerts

The IBM System Storage DS3000 Manager includes the Event Monitor Service, which enables the host running this monitor to send out alerts through e-mail (SMTP) or traps (SNMP). The Event Monitor can be used to alert you of critical events for any of the DS3000 storage subsystems in your environment. For high availability, the event monitor service should run on two separate host systems that are available 24 hours a day. Both servers should be capable of out-of-band and in-band management. This ensures proper alerting, even if one server is down or a connection type has failed.

Note: The DS3000 does not send the e-mail or SNMP trap itself. The management station (running the event monitor service) sends the notification on behalf of the storage server. If the management station is down, no alerting will be done.

What systems are monitored

Depending on how you set up alerts, different storage systems are monitored by the Event Monitor.

- Alerts can be configured as follows:
  - To configure alerts for all storage systems listed in the Enterprise Management window, right-click your local management system in the Enterprise Management window (at the top of the tree) and choose **Configure Alerts**, as shown in Figure 8-9.

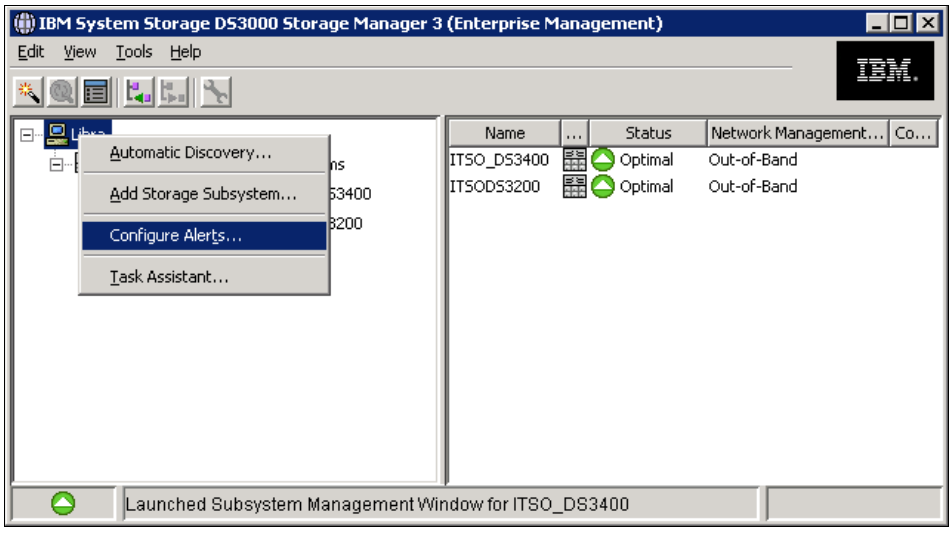


Figure 8-9 Configure Alerts for all systems

  - If you can see the same storage subsystem through different paths, directly attached and through different hosts running the host agent, you will receive multiple alerts. Right-click a specific storage system to define the alerting specifically for this storage server, as shown in Figure 8-10 on page 119.

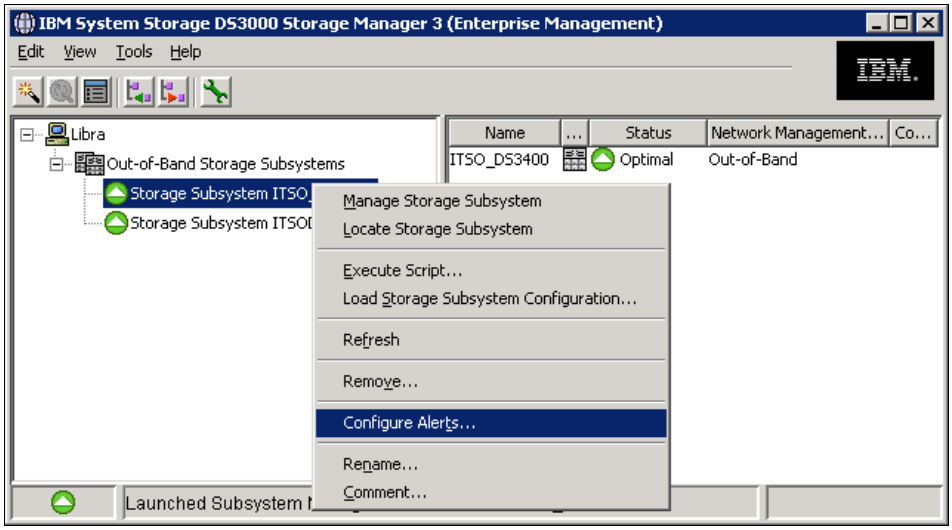


Figure 8-10 Configure Alerts for particular subsystem

– Select **Configure Alerts** for one or more systems and the alert window appears.

E-mail alerts

Here we discuss how to work with e-mail alerts:

1. To send e-mail alerts, you have to first define an SMTP server. Enter the name or IP address of your mail server and the sender e-mail address, as shown in Figure 2.

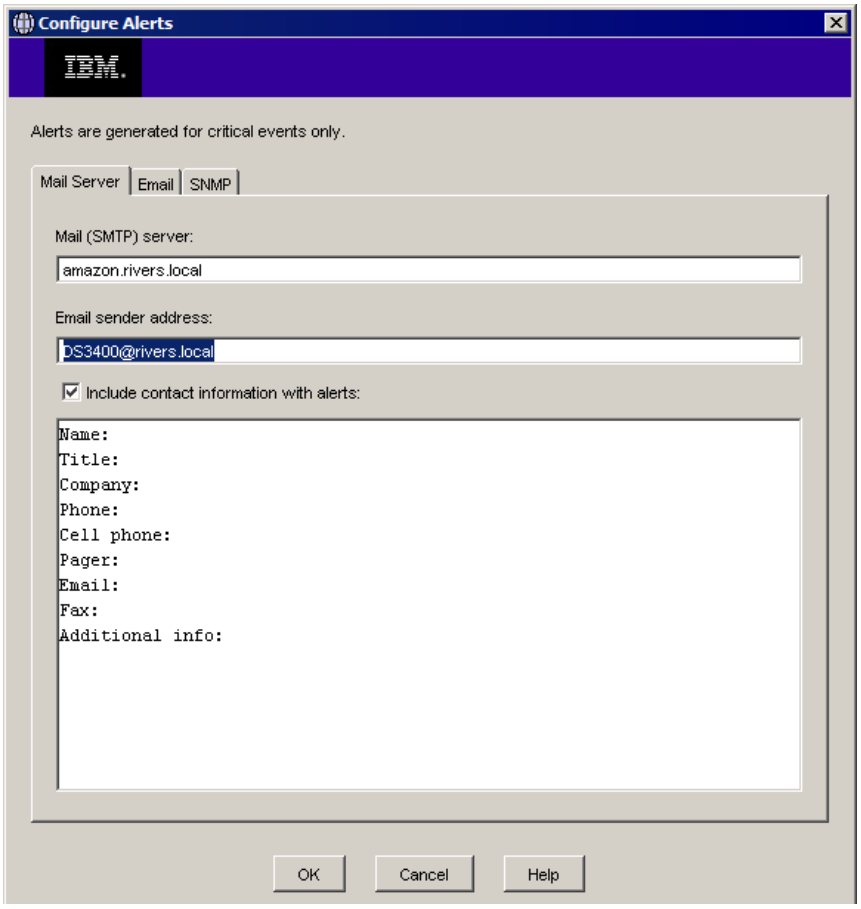


Figure 8-11 Configure Alerts - define mailserver

2. Select the **Email** tab to configure the e-mail address to which the alerts are sent. Enter the e-mail addresses and click **Add.**

3. Now your window should look similar to Figure 8-12. In the drop-down menu next to the Address, you can select which information should be sent and how often the data should be delivered. If you highlight an Address, you can also send a test e-mail to validate that your configuration is working.

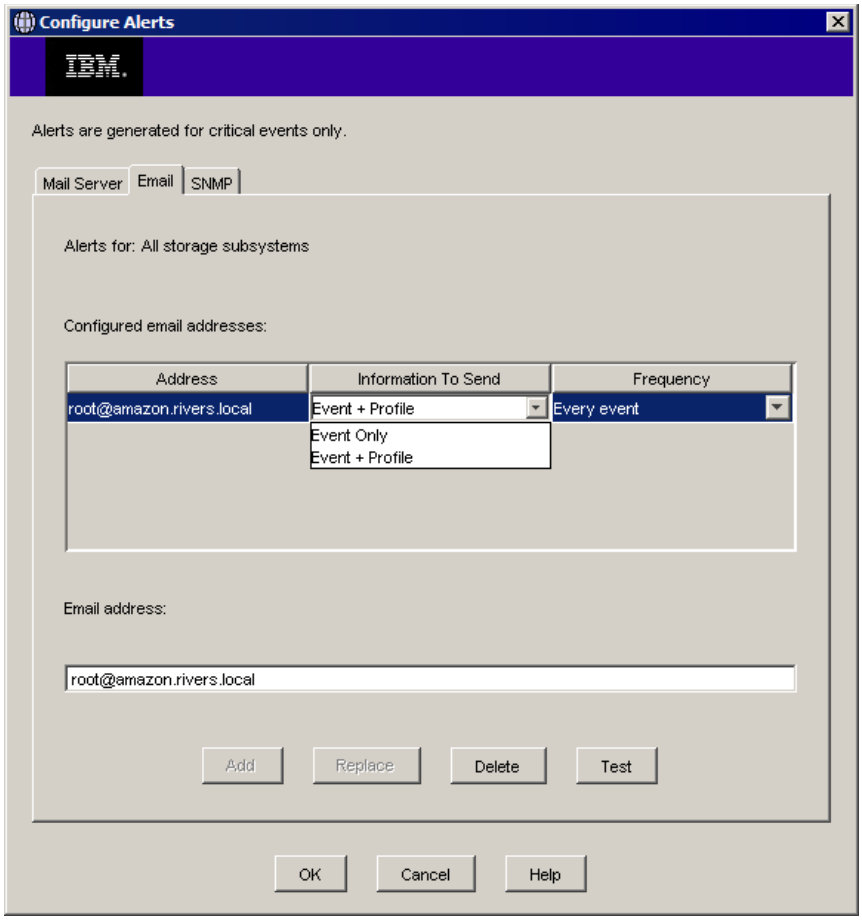


Figure 8-12 Configure Alerts - Defining e-mail address

SNMP alerts

Similarly, select the SNMP console for receiving and handling the traps sent by the service. Enter your community name and the destination dress of your SNMP Server and click **Add**, as shown in Figure 8-13.

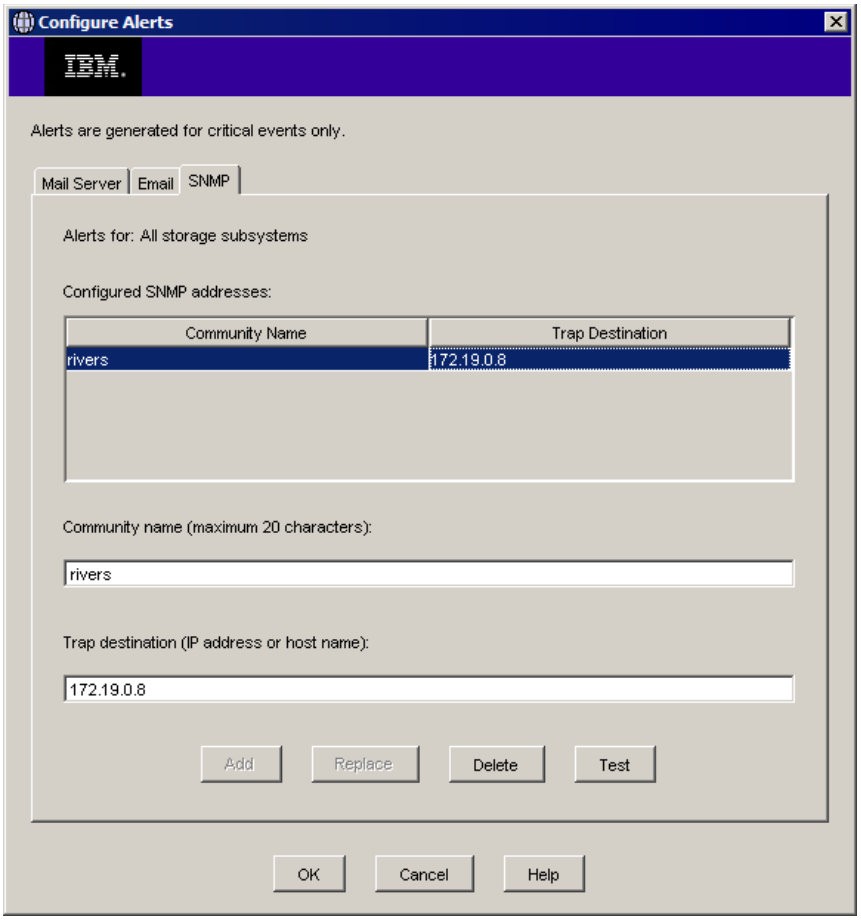


Figure 8-13 Configure Alerts - SNMP

Note: An SNMP manager is required to receive and translate your SNMP traps, for example, IBM Director. For more information about IBM Director and how to configure it, see *Implementing IBM Director 5.20*, SG24-6188.

Rename

If you have multiple DS3000s installed, we recommend giving each one a user-specified name for easier identification. To rename a DS3000, choose this option, enter a new name for the subsystem, and click **OK** (Figure 8-14 on page 123).

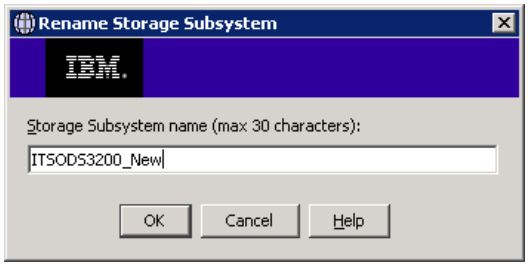


Figure 8-14 Rename Storage Subsystem

Note: We recommend that you set the name only during the initial setup of the box and do not change it later. Some operating systems use the storage subsystem name for identification.

Comment

Here you can add a comment to a disk subsystem in the Enterprise Management window. Choose this option, enter your comment, and click **OK**, as shown in Figure 8-15.

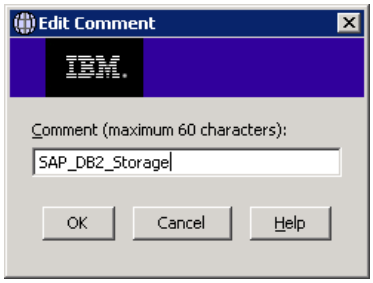


Figure 8-15 Edit comment

Your comment will appear in the right pane of the Enterprise Management window, as shown in Figure 8-16. Your comment could be some kind of description of the disk subsystem.

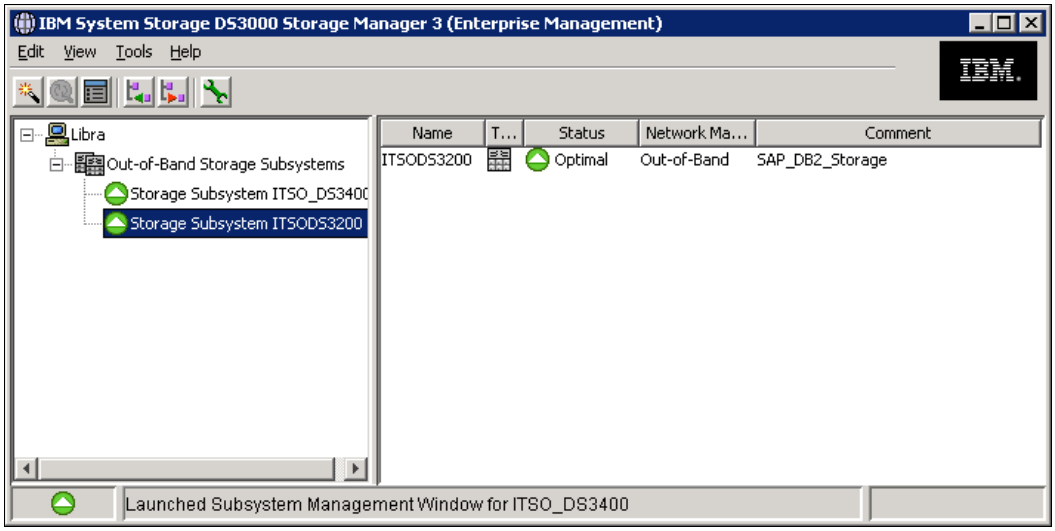


Figure 8-16 Add comment to Storage Subsystem

To delete or edit the comment, just re-select the **Comment** option and modify it as needed. Click **OK** when done.

8.2.2 The Enterprise Management window menu bar

You will see various options in the window menu bar at the top of the Management window.

Edit and Tools submenu

All the tasks described in 8.2.1, “Subsystem context menu” on page 115 can also be accessed from the task menu bar. To do this, highlight a subsystem and open the appropriate menu. The Edit (Figure 8-17) and Tools (Figure 8-18) submenus contain all of the tasks. Performing these tasks is done exactly as described in the previous section.

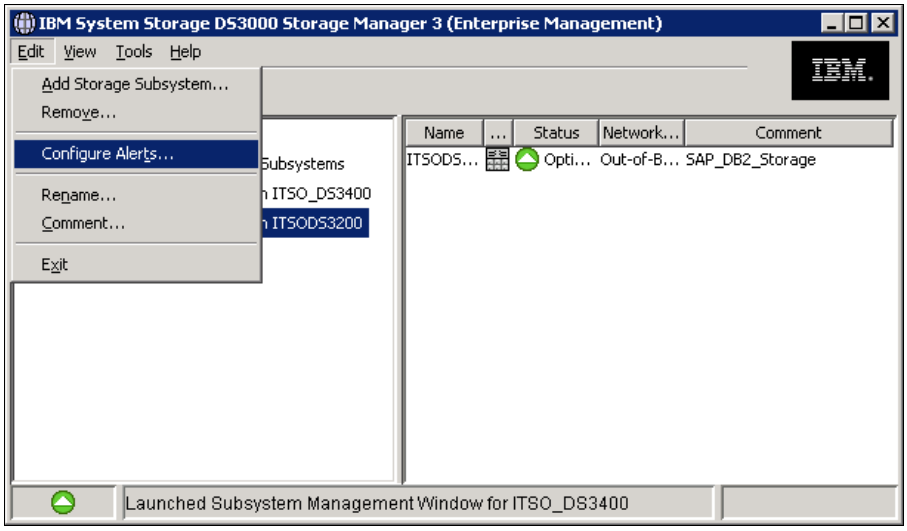


Figure 8-17 Edit menu

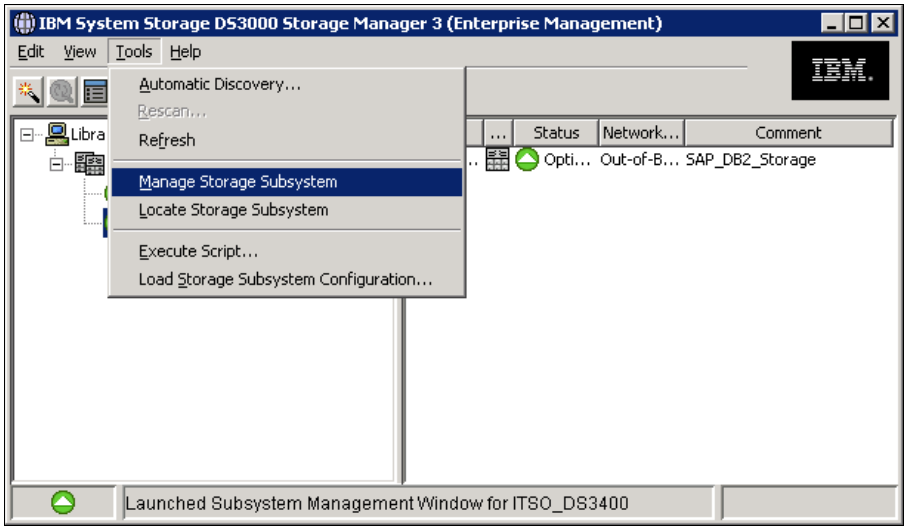


Figure 8-18 Tools menu

The View submenu

Use this menu to start the Enterprise Management window Task Assistant (Figure 8-19) or to customize the view of your enterprise window by selecting one of the sort options. The Task Assistant is described in 8.1.1, “Task Assistant” on page 112.

Sort the subsystems in the Enterprise Management window

To sort your subsystems, choose one of the following fields from the menu as a sort key:

- By Name
- By Status
- By Management Type
- By Comment
- Partially Managed Storage Subsystems

Note: Partially Managed Storage Subsystems are storage servers where only one IP Address (Controller) is connected to the management station. You should ensure that you always add both Ethernet connections. Some maintenance tasks require a connection to both controllers.

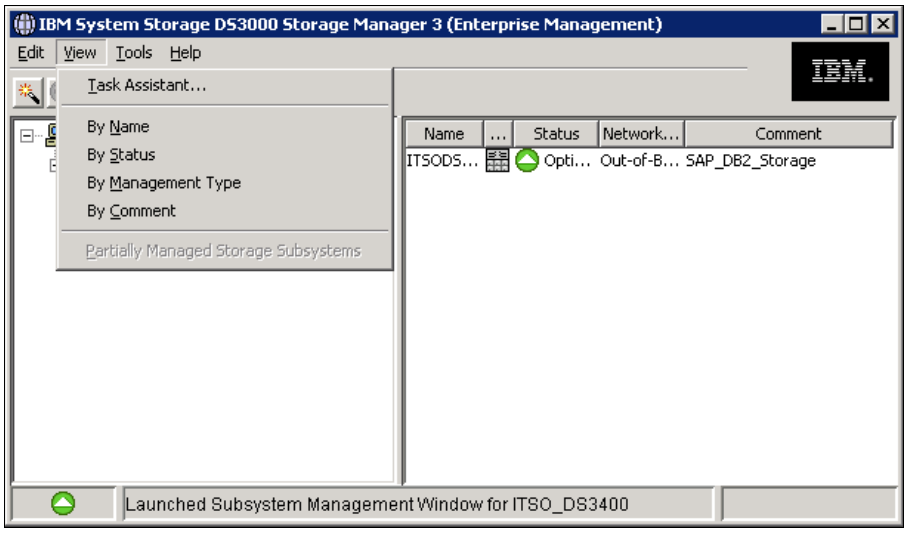


Figure 8-19 View submenu

8.2.3 The Quick Access buttons

Below the menu bar you see the Quick Access buttons, which are used to directly activate some functions of the Enterprise Management window. Click them to activate the associated task.

Automatically discover new storage subsystems



Figure 8-20 Automatically discover new storage systems button

This button starts an automatic discovery of new storage subsystems (Figure 8-21) by sending in-band and out-of-band broadcasts. If it finds directly attached storage systems or hosts running the DS3000 Storage Manager agent (with an attached Storage subsystem), it adds these storage subsystems to the Enterprise Management window.

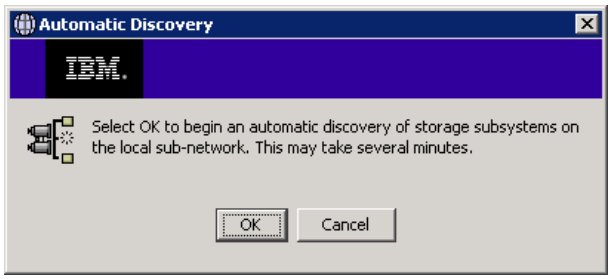


Figure 8-21 Automatic Discovery

Note: If your subsystem is not discovered automatically by the wizard (for example, because it is in another broadcast domain), you can add it manually, as described in “Add Storage Subsystem to view” on page 127.

Rescan selected host for new storage subsystems



Figure 8-22 Rescan Host button

If you are using a host running an agent to manage your storage subsystems, you can use this button to only rescan the selected host. To do this task, highlight the host, click this button, and the Rescan window appears (Figure 8-23). Select **OK** and the host will scan for newly attached storage devices.

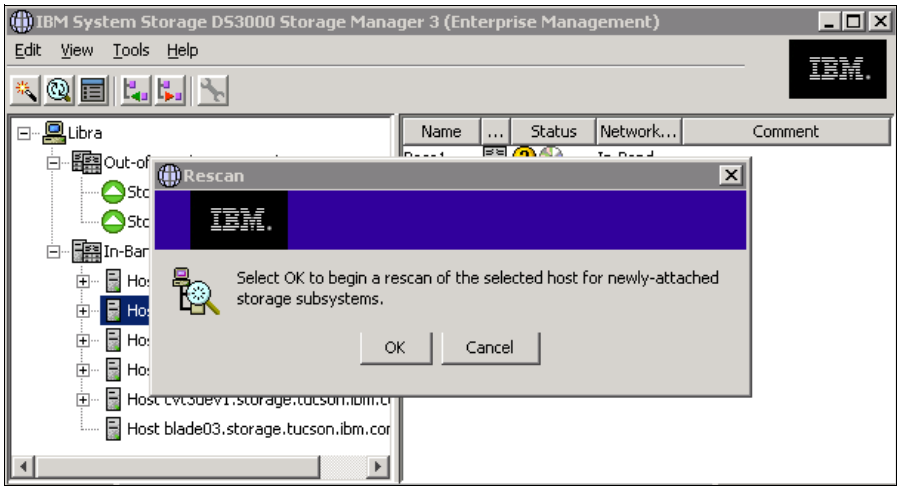


Figure 8-23 Rescan Single Host

Display Task Assistant



Figure 8-24 Display Task Assistant

This will start the Task Assistant wizard described in 8.1.1, “Task Assistant” on page 112.

Add Storage Subsystem to view



Figure 8-25 Add Storage Subsystem to view

Click this button to display the Add Storage Subsystem window. Here you can add systems manually to your Enterprise window view. For out-of-band management, you have to provide the IP address or host names of the controllers (Figure 8-26). If the storage subsystem has dual controllers, you must provide the IP address of both controllers.

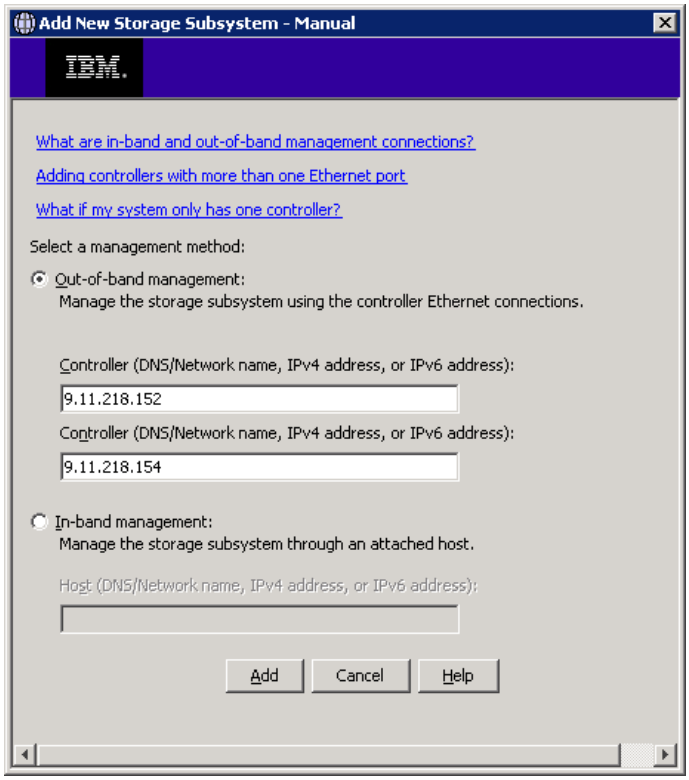


Figure 8-26 Manually add out-of-band system

For in-band management, specify the IP address or host name of the host attached to the storage server (Figure 8-27).

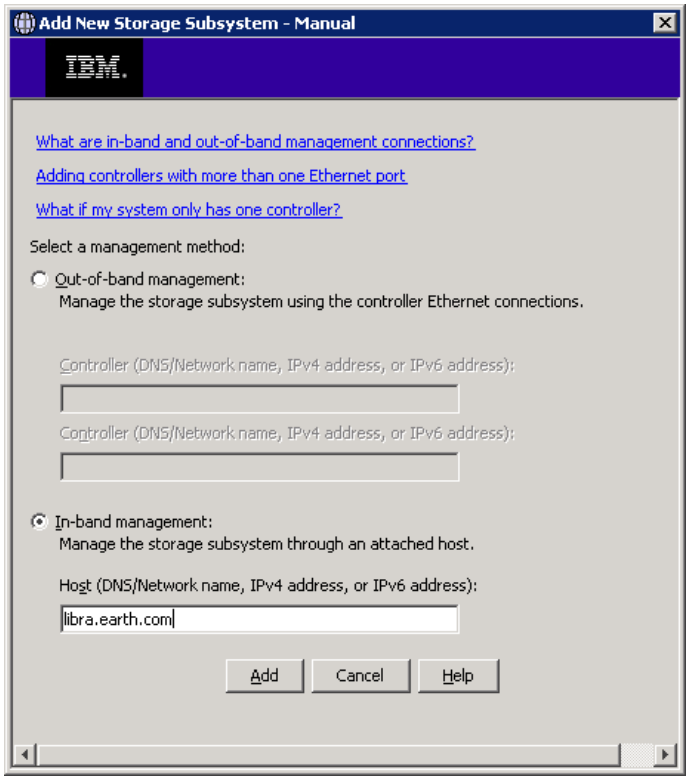


Figure 8-27 Manually add in-band system

Remove Storage Subsystem from View



Figure 8-28 Remove Storage Subsystem from View

To remove a storage subsystem from your view, highlight it, click this button, and click **OK** to confirm (Figure 8-29 on page 129). This subsystem is only removed from your view; the views of other management stations are not affected.

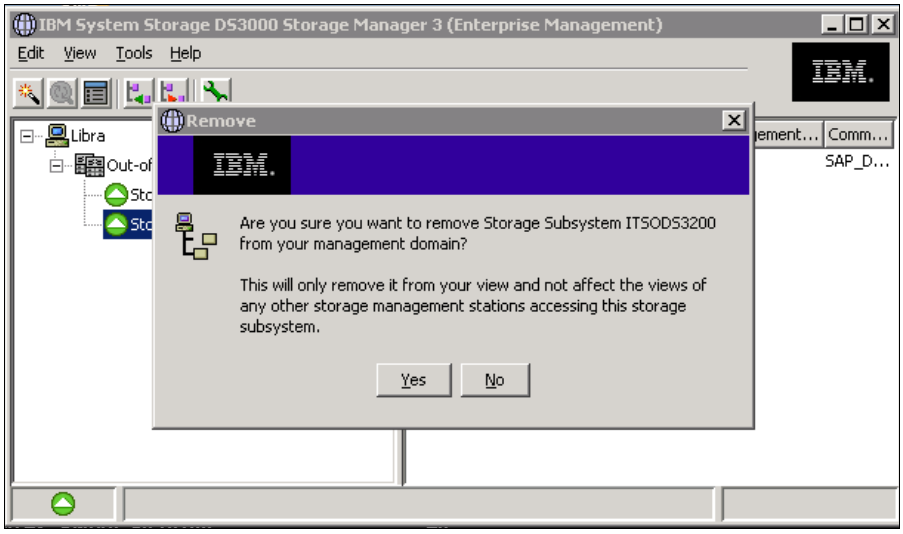


Figure 8-29 Remove Subsystem from view

Launch Management window



Figure 8-30 Launch Management window

Highlight a storage subsystem and click this button to launch the Subsystem Management window for the selected system, as described in the window for the selected system..

9

Administration - Configure

In this chapter, we describe the Configure tab of the IBM System Storage DS3000 Storage Manager. In this tab, you can configure a DS3000 storage subsystem, including functions for creating hosts, arrays/logical drives, and mapping them together, so that the host can access the storage subsystem.

You can use the Subsystem Management Configure tab to:

- Hosts
  - Configure host access (automatic).
  - Configure host access (manual).
  - Create host groups.
  - Create host-to-logical drive mappings.
- Storage
  - Automatic configuration.
  - Configure hot spares.
  - Create arrays and logical drives.
  - Create flashcopy logical drives.
  - Create logical drive copies.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

9.1 Configure host access

Host access is how a host (server) can access storage on the DS3000.

9.1.1 What a host is

First, we take a closer look at configuring hosts. Hosts in a DS3000 context can have a number of meanings:

- Host Port

  A host port is a single Host Bus Adapter Port in a server that is attached to the storage subsystem. The port can be a port on a SAS, iSCSI, or FC Host Bus Adapter (HBA), depending on the technology used. Each port uses a unique identifier known as the *World Wide Name* (WWN). If you are using dual port adapters, they have two WWNs, one for each port. iSCSI software initiators also generate a WWN for identification.

- Hosts

  A host is a single system that can be mapped to a logical drive. The DS3000 only communicates through the use of a WWN. The DS3000 is not by default aware of which HBAs are in the same server or of servers that have a certain relationship, such as a cluster. This is something you have to configure on your storage server. You can create hosts and assign to them to one or several WWNs of the host ports.

- Host groups

  A host group is a collection of hosts that are allowed to access certain logical drives, for example, a cluster of two systems.

To summarize, every HBA port in a server has its own unique WWN. This WWN can be assigned to create single hosts. If several hosts have to access the same logical drive, you can use host groups to achieve that.

9.1.2 Preparations in the host server

Before you configure your storage subsystem, you should do some preparation in the host server and your environment to make the configuration simpler:

- Install the host HBAs in the host server (see your server's documentation for any special instructions). Some servers have recommended slots to use for installing an HBA. For redundancy, we always recommend using two HBAs per host. Update the HBA firmware and the host server firmware to the current levels and shut down your host server.

  Note: IBM server firmware can be found on the IBM Support Web page at:

  `http://www-304.ibm.com/jct01004c/systems/support/`

- Attach all host cabling to the storage subsystem.
- If you are not using a direct attached topology, configure your switches to ensure the right zoning. We show some sample zoning configurations in Chapter 24, "FC configuration 1 - Emulex HBA boot blade server from SAN" on page 585 and Chapter 26, "FC configuration 3 - Linux SAN boot from a DS3400 with an IBM System x server" on page 639. For the DS3400, you can use QLogic or Emulex HBAs. Power on your server and enter the HBAs' BIOS setup. This is done by entering a particular key sequence while the BIOS is loading. If you do not see the BIOS loading during the bootup sequence, make sure that ROM Control Execution is Enabled on your host server.

Tip: If you are using an IBM System x server, to enable ROM Control Execution, enter the System BIOS (F1) and select **Advanced Setup** → **PCI Bus Control** → **PCI Control Execution**.

FC QLogic HBAs

Do the following steps:

1. To enter the QLogic HBA BIOS (called Fast!UTIL), press Ctrl + Q when prompted. (Figure 9-1).

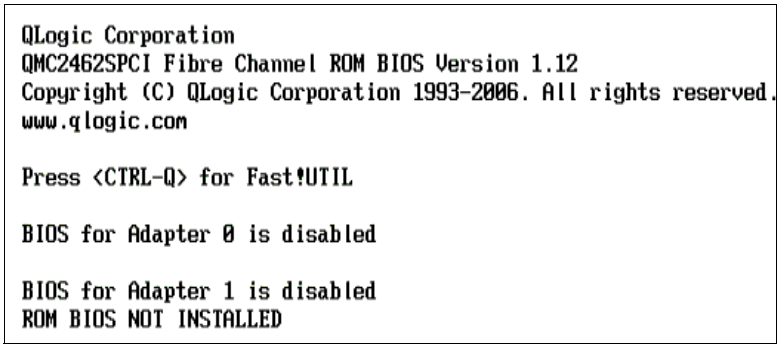


Figure 9-1 Enter QLogic BIOS

2. A screen similar to Figure 9-2 displays, listing the adapters found. Choose one and press Enter to load the ports menu.

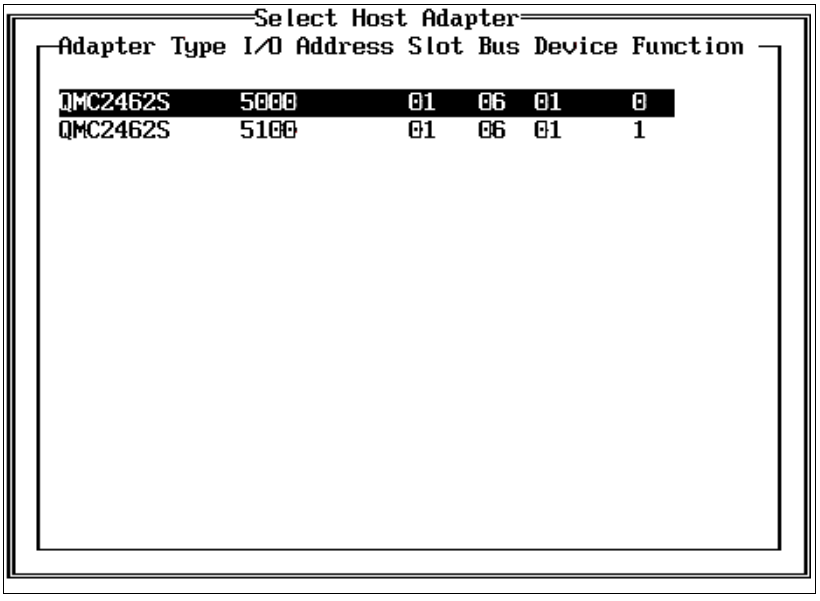


Figure 9-2 QLogic - Select adapter

3. Figure 9-3 shows the ports menu that is displayed when you select an adapter.

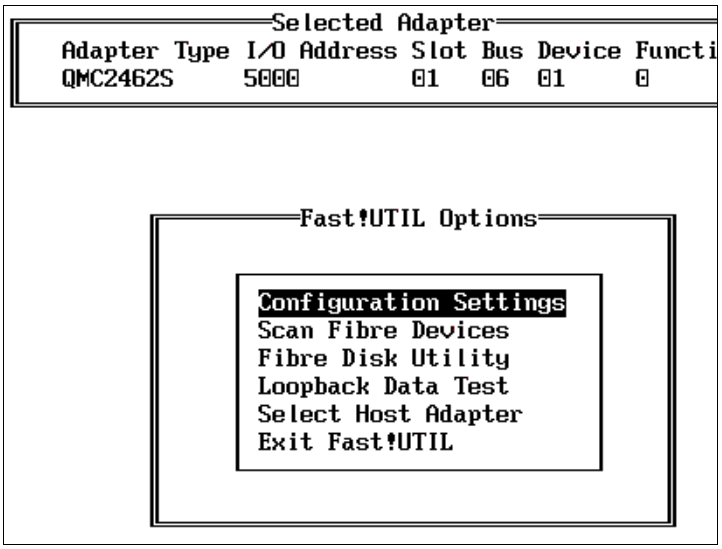


Figure 9-3 QLogic - Adapter menu

This ports menu is for configuring the HBA's port. In this book we will not cover all possible FC HBA settings, but only the most important ones for a basic setup to connect to your DS3000 storage subsystem. If you are familiar with FC environments, you will have your own procedures to implement the DS3000 storage servers in your environment.

4. You need to discover the WWN of your FC port. The WWN is the unique identifier, which will be seen by the DS3000. Select Configuration Settings and in the next submenu Adapter Settings (Figure 9-4).

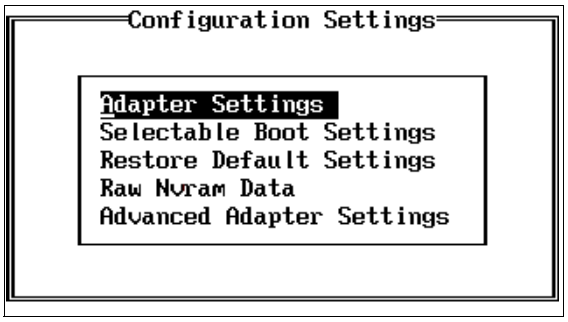


Figure 9-4 QLogic - Configuration submenu

5. The Adapter Settings are displayed (Figure 9-5 on page 135). The entry Adapter Port Name shows the WWN of this port. Record this WWN with the associated HBA (port) number and host; you will need it for later configuration steps. In our example, the WWN port number is 210000E0B157800.

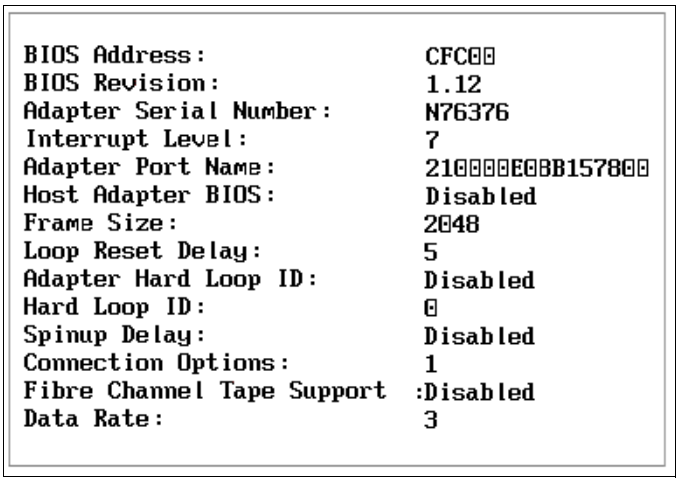


Figure 9-5 QLogic - WWN

6. If your server has more than one port (more than one HBA or a dual port HBA), press ESC until you return to the QLogic Adapter menu (shown in Figure 9-3 on page 134). Highlight the Select Host Adapter line and press Enter to return to the Select Host Adapter screen (shown in Figure 9-2 on page 133). Select the other port and repeat these steps to discover the WWN of the second port.
7. After you have recorded all the WWNs, return to the main menu of one port and select Scan Fibre Devices (Figure 9-6). The port will run a scan to discover all attached devices.

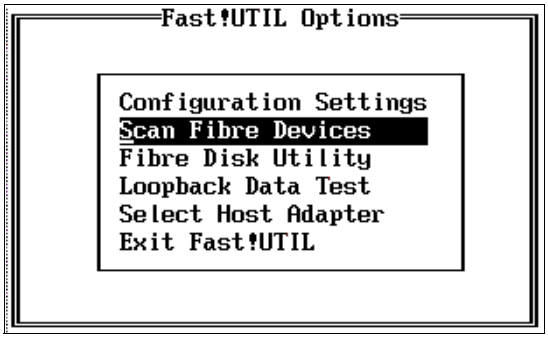


Figure 9-6 QLogic - Scan Fibre Devices

8. Depending on your cabling and zoning, you should see at least one port of your attached DS3400, as shown in Figure 9-7.

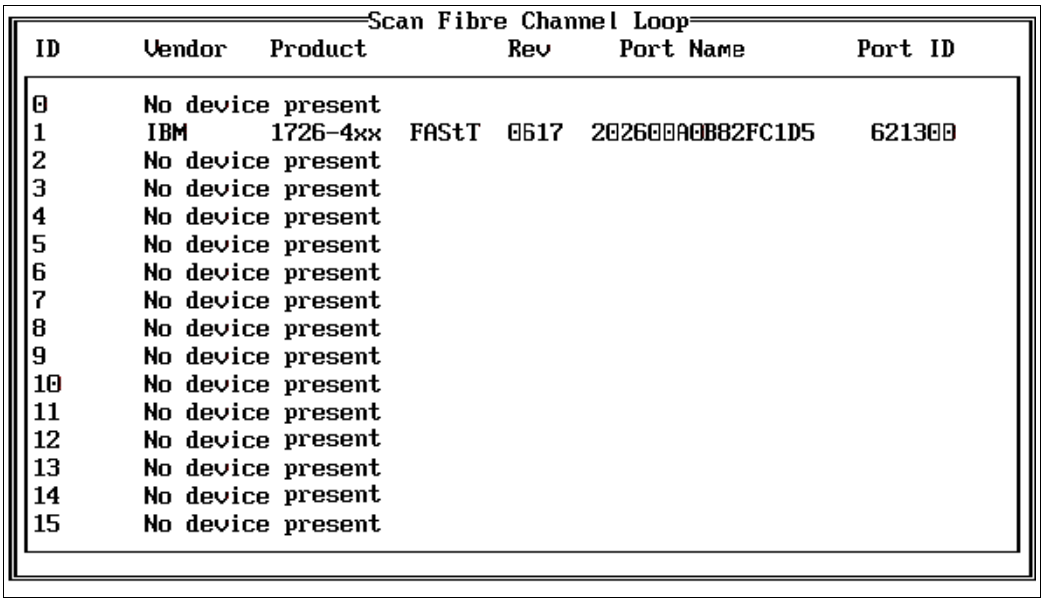


Figure 9-7 QLogic - One DS3400 Port discovered

Repeat these steps on each HBA port that is installed in the system. If you see the correct DS3400 (depending on your configuration), you have finished the basic hardware setup and have the QLogic Adapters installed.

FC Emulex HBA

To enter the Emulex HBA BIOS, press Ctrl +E or Alt + E when prompted (Figure 9-8).

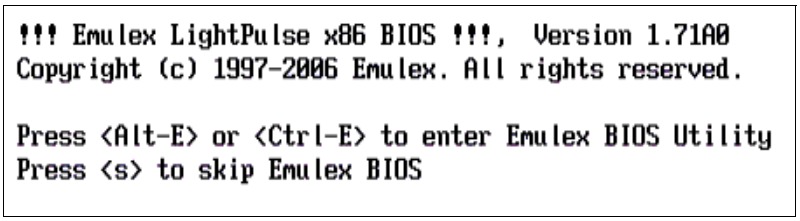


Figure 9-8 Emulex - Enter HBA BIOS

Figure 9-9 shows the initial screen of the BIOS Utility.

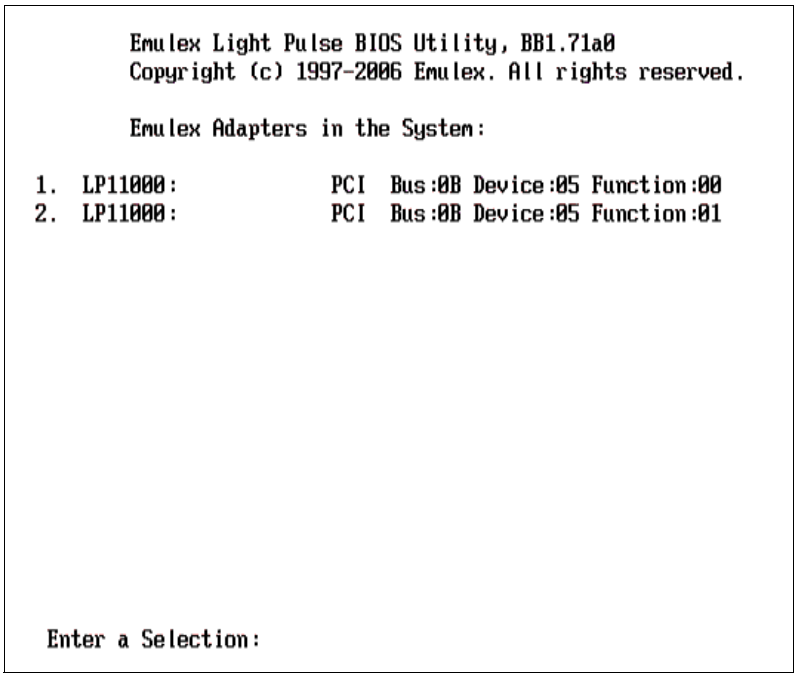


Figure 9-9 Emulex - BIOS

Initially, you will see that all the HBA ports have been detected. Select one of the port numbers and press Enter to open the port menu (Figure 9-10). This menu displays the port's WWN; our example shows 10000000 C93542C2. Repeat this for the second port of the adapter and record the WWNs.

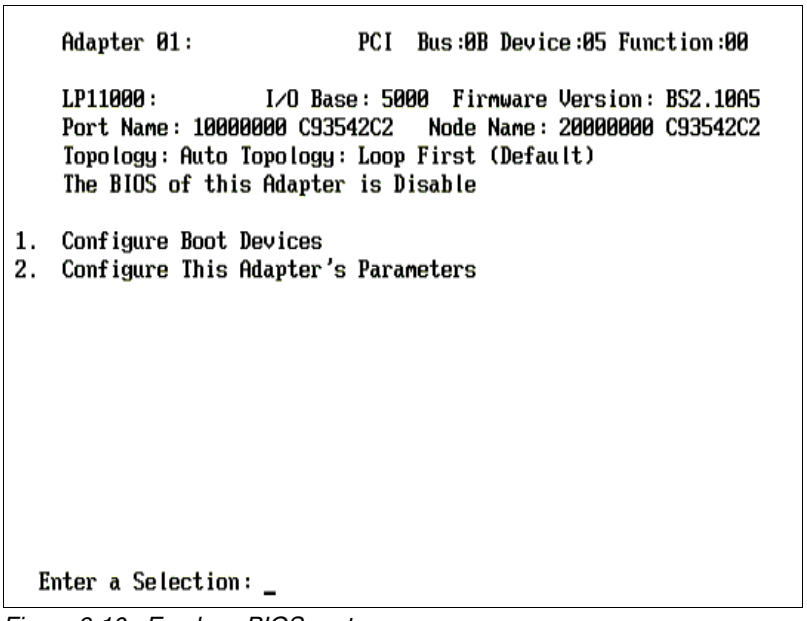


Figure 9-10 Emulex - BIOS port menu

LSI SAS HBA

Do the following steps:

1. To enter the LSI SAS HBA BIOS, press Ctrl + C when the BIOS sequence displays, as shown in Figure 9-11.

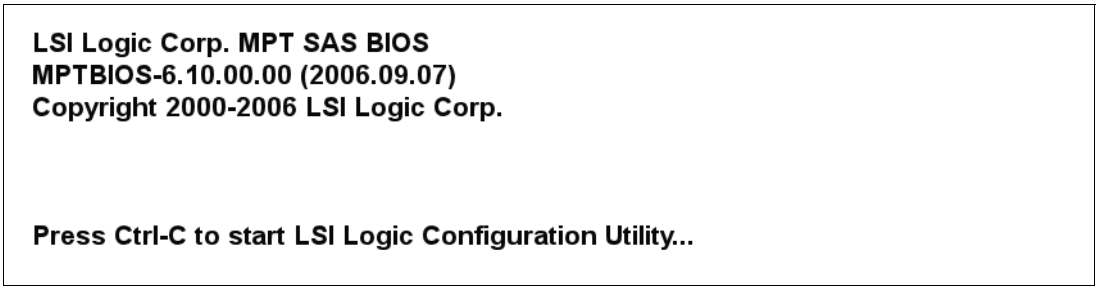


Figure 9-11 Enter LSI HBA BIOS

2. The utility starts, as shown in Figure 9-12. Highlight the adapter and press Enter to open the adapter's BIOS.

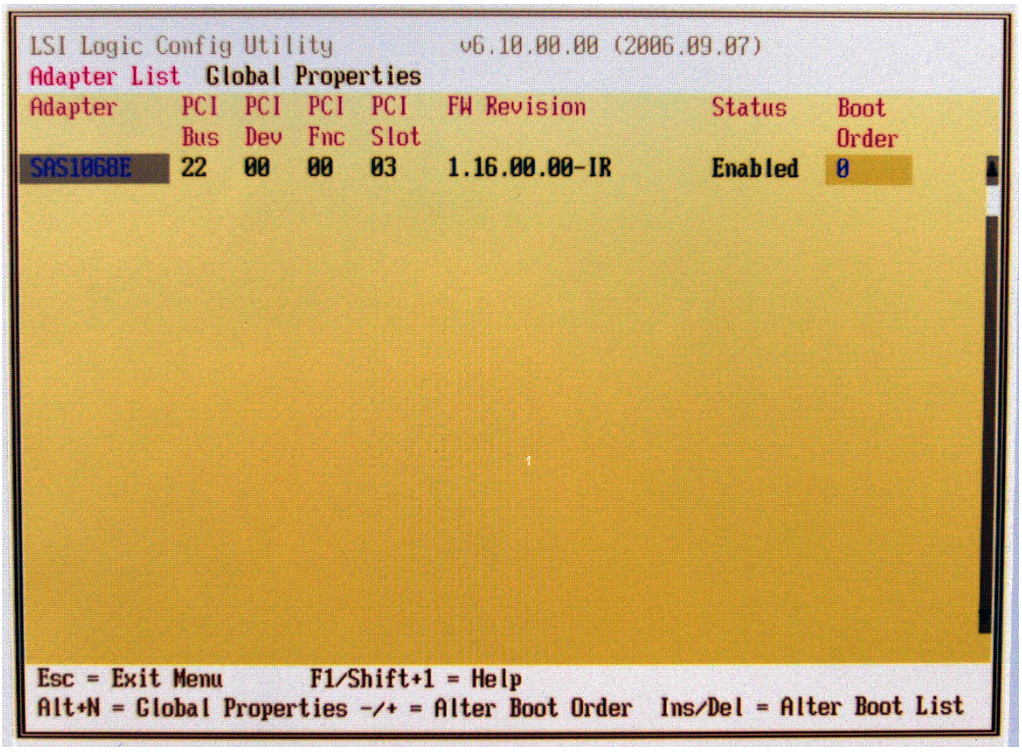


Figure 9-12 Select SAS adapter

3. The adapter BIOS is now displayed, which shows you the SAS Address (Figure 9-13). This is the identifier of the SAS HBA, in our example, this is 500605B0:000701F0.

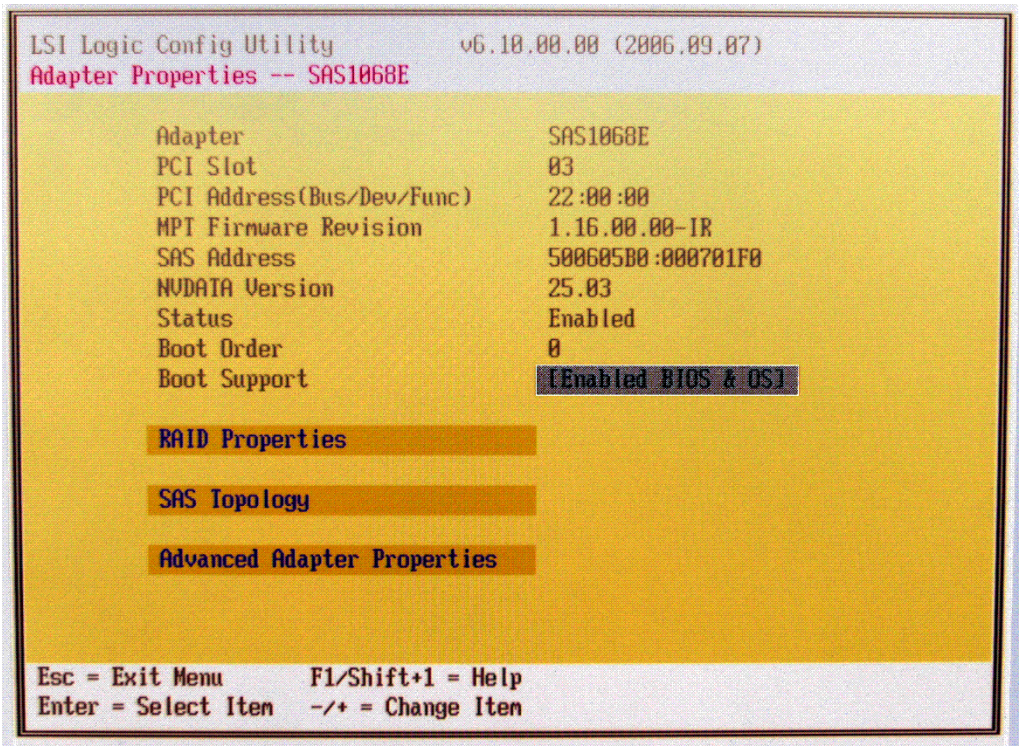


Figure 9-13 SAS Address

4. Record the SAS address and host. Repeat for any other SAS HBAs. Press Esc to exit the SAS HBA BIOS.

9.2 Configure hosts

The recorded WWN of the host server HBAs will be used here to set up and configure the DS3000.

In the first part of the *configure* tab (Figure 9-14 on page 141), you can:

- Configure host access (automatic).
- Configure host access (manual).
- Create host groups.
- Create host-to-logical drive mappings.

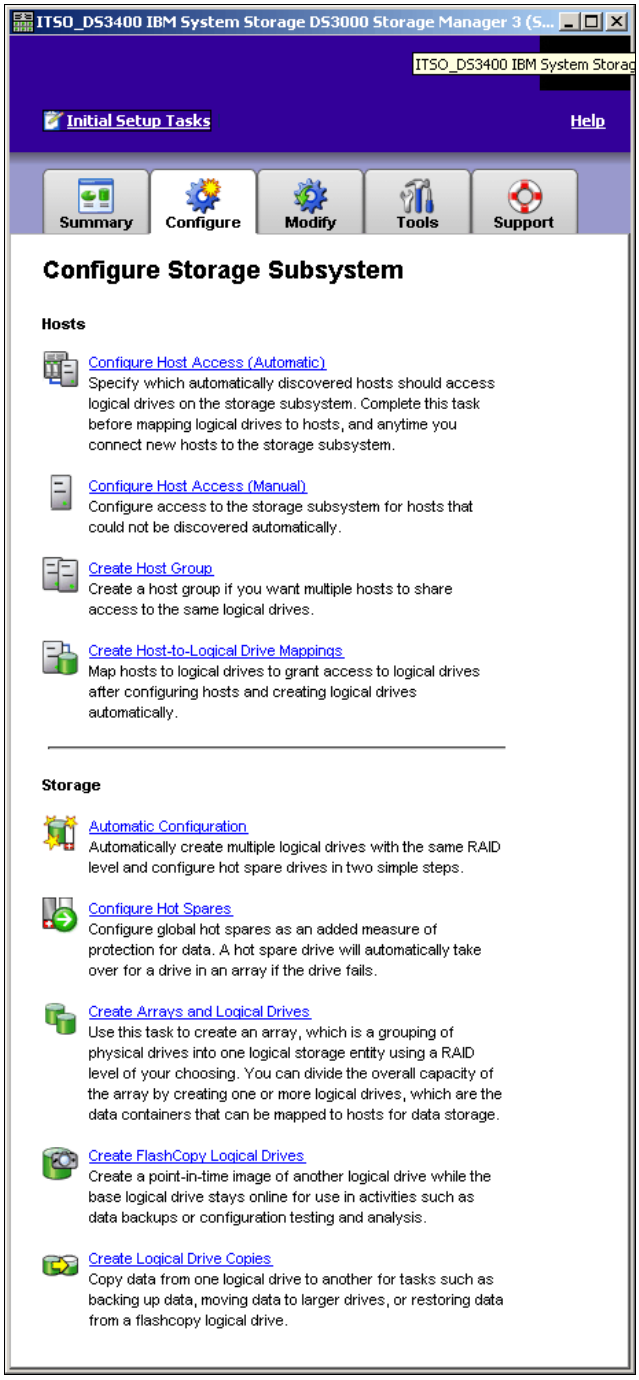


Figure 9-14 The Configure tab

9.2.1 Configure host access (automatic)

The Initial Setup Tasks wizard provides a shortcut to automatic host access configuration. For completion's sake, we also describe manual host access configuration in this section. The goal of either of these two methods is to define a host server that will later be used for logical drives assignment.

Automatic host access configuration is straightforward: Move hosts discovered in the Available Hosts window to the Selected Hosts window (see Figure 9-15). “What hosts are discovered” on page 144 describes the conditions under which hosts will be discovered in the Available Hosts window.

In our example, the host server named Libra was automatically discovered and listed in the Available Hosts window.

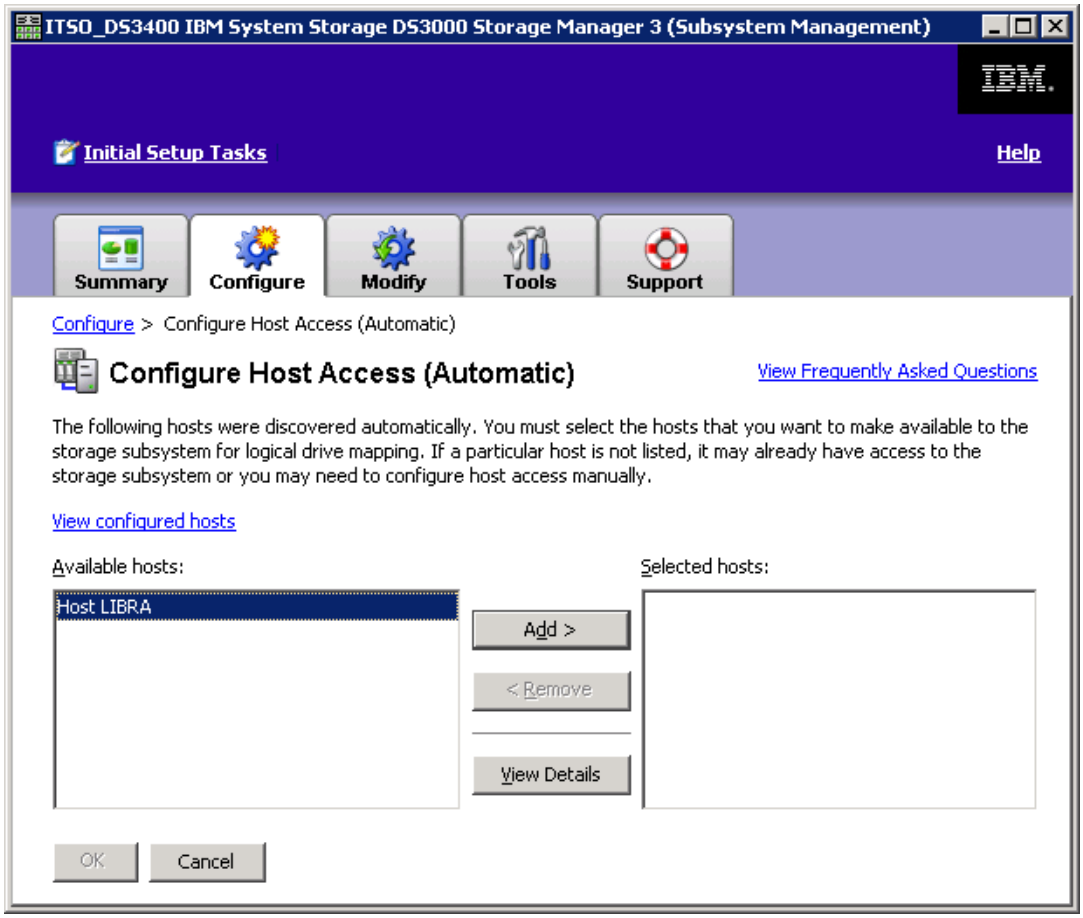


Figure 9-15 Configure Host Access (Automatic) - available hosts

Highlight the host, and click **Add** to move the host server to the Selected Hosts window, as shown in Figure 9-16 on page 143.

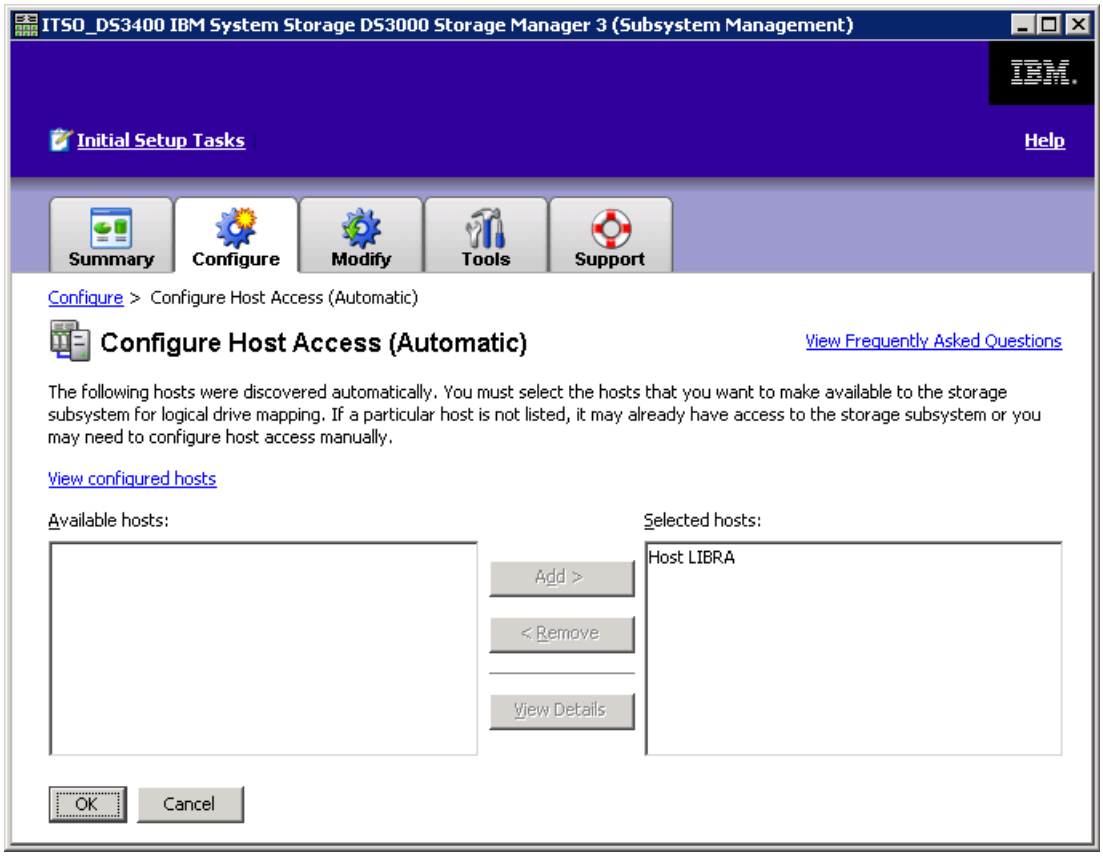


Figure 9-16 Configure Host Access (Automatic) - selected hosts

Click **OK**, which will display the Configuration Completed window. Your host server is now available for mapping to drives.

It is a good idea to check that all host ports have been correctly detected. You can do this by using the Edit Topology task under the Modify tab. In Figure 9-17, the two HBA host ports have been correctly detected in our sample scenario.

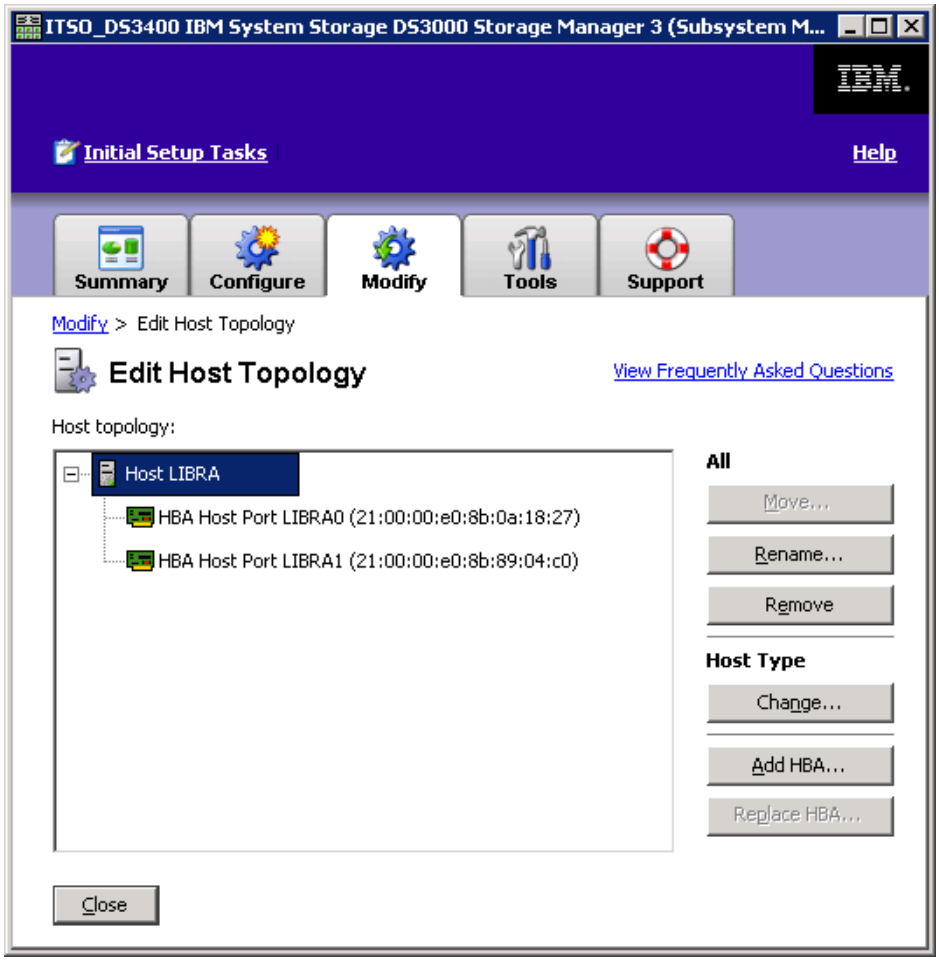


Figure 9-17 Verifying the host and host ports

What hosts are discovered

For a host server to be automatically discovered:

- The DS3000 Storage Manager host software components must be installed on the server.
- The SMagent service must be up and running.
- The host server can communicate with the access logical drive.

On a Windows host server, you can easily verify that SMagent is running in the Application Event Log. Check the last two SMagent events. The last event will state whether the SMagent service is running. The next to last event will display the access logical drives that are detected by the SMagent (see Figure 9-18 on page 145).

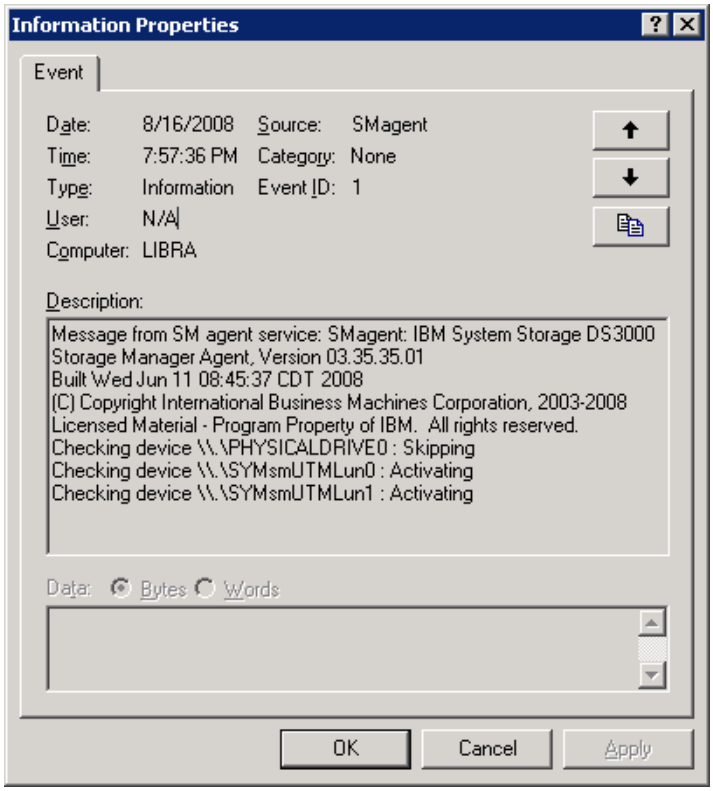


Figure 9-18 Application event log - event properties

9.2.2 Configure host access (manual)

As described in 9.1.1, “What a host is” on page 132, the DS3000 communicates only through WWNs.

The storage subsystem by default does not relate ports and HBAs to the servers where they are installed. You have to configure these relationships in this section of Storage Manager.

Specify the following parameters in order to define a host:

- Host server name.
- Host server type.
- Identify the HBAs (for example, in case of FC HBAs, we need to specify their WWNs).

Using the window shown in Figure 9-19, you can set the host server name and type.

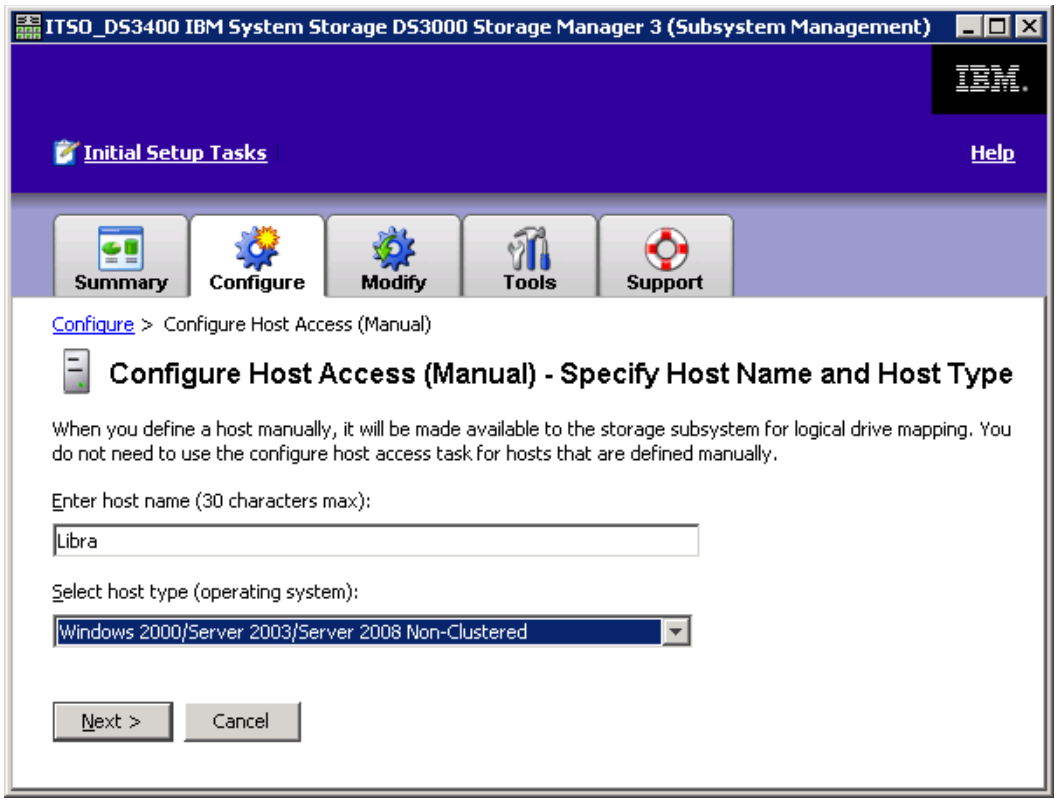


Figure 9-19 Specify Host Name and Type

Our sample server is a System x running Windows Server 2003 and it is not a part of any cluster. The host type you select must reflect this.

Click **Next**; this brings you to a window similar to Figure 9-20 on page 147. Here, you specify the HBA host ports. In our example, our server uses a dual-port 4 Gbps FC HBA. Therefore, we need to define two host ports and specify the correct WWN for each port. In this case, SM has detected the WWN number automatically; otherwise, click **New** to add WWN numbers manually.

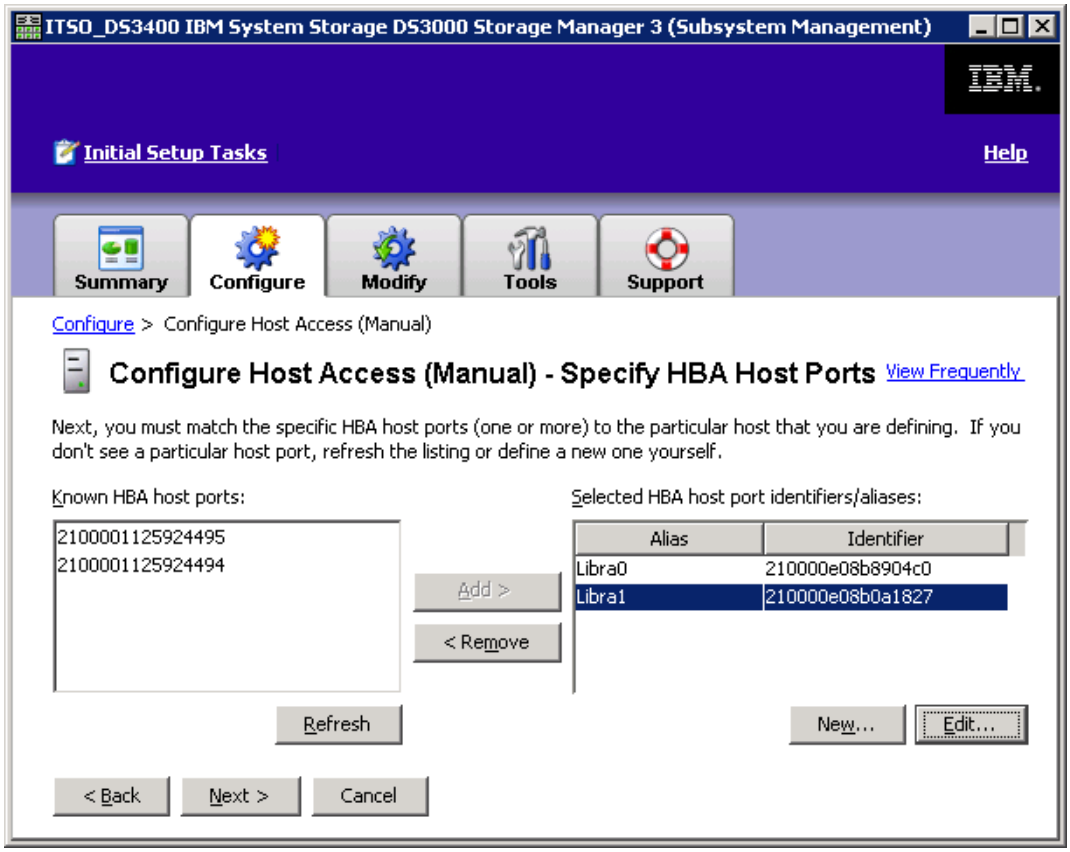


Figure 9-20 Specify HBA Host Ports

Clicking **Edit** will open the window shown in Figure 9-21. You can now edit the HBA host port identifier and an alias to allow it to be identified. As we are using an FC HBA, this will be the WWN of FC HBA port 0.

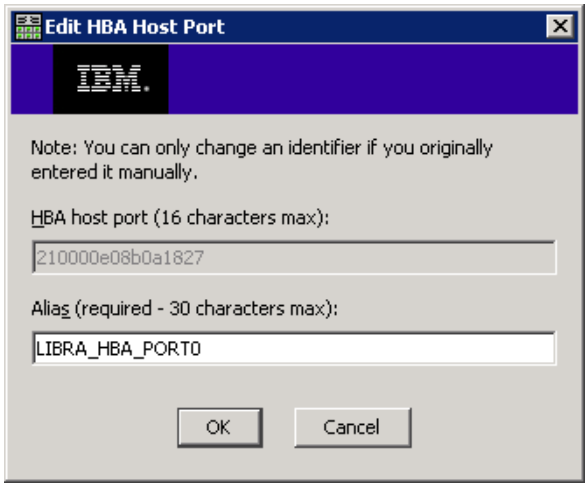


Figure 9-21 Edit the new HBA host port

Repeat the same steps for the second HBA host port. When this is complete, both host port identifiers should be visible in the window, as shown in Figure 9-22.

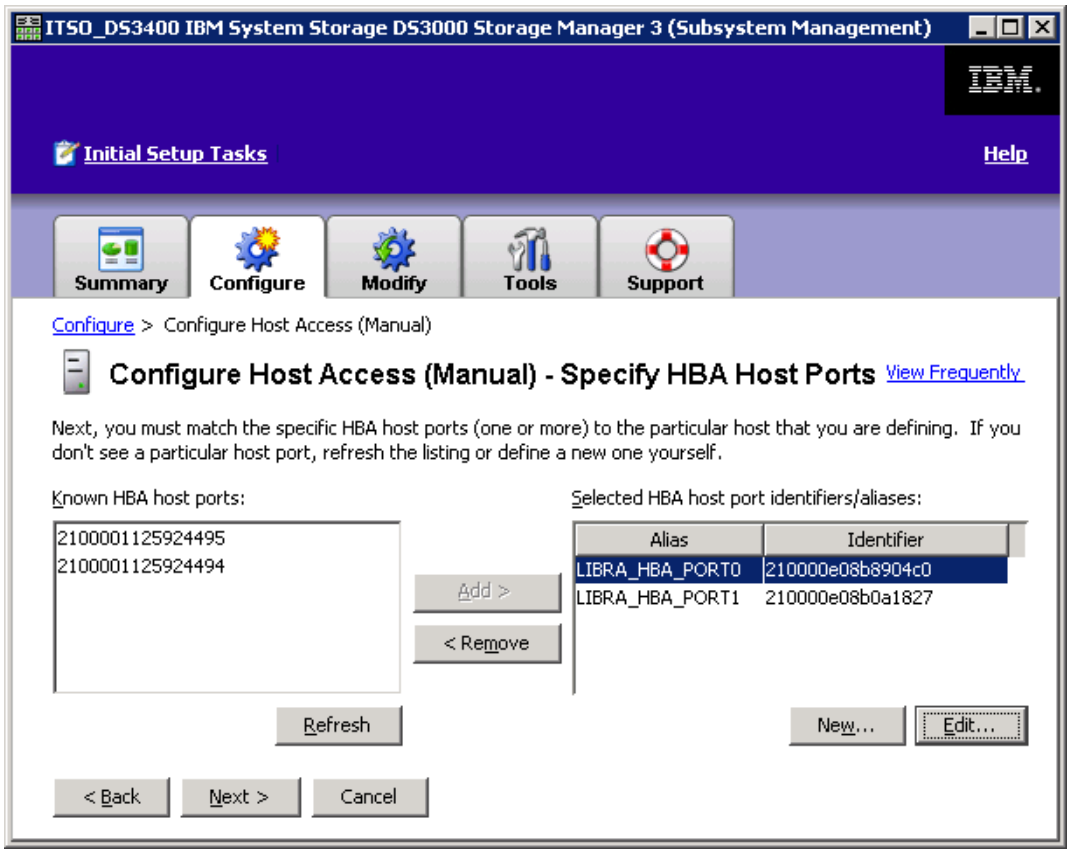


Figure 9-22 Both host port identifiers have been defined

Host group

The next step is to define a *host group*. This is only required if the host is part of a cluster, because a group is basically a container for the definitions of clustered host servers.

In our sample case, the host server is a stand-alone machine, so we will not create a group. See the Specify Host Group window in Figure 9-23. We select **No**, and click **Next**.

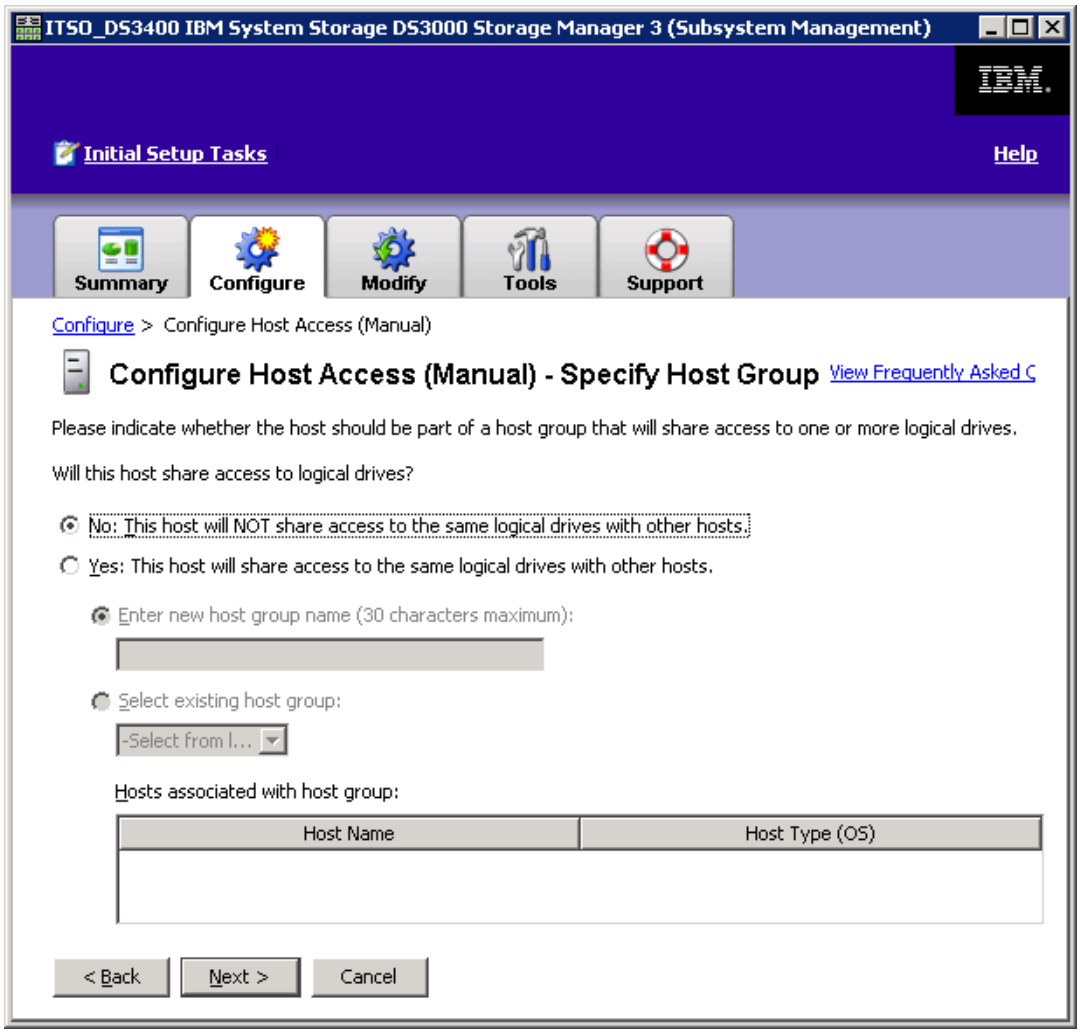


Figure 9-23 Specify Host Group (optional)

Completion

Finally, the host definition summary window appears, as shown in Figure 9-24. You can correct the settings if required; otherwise, click **Finish** to complete the process.

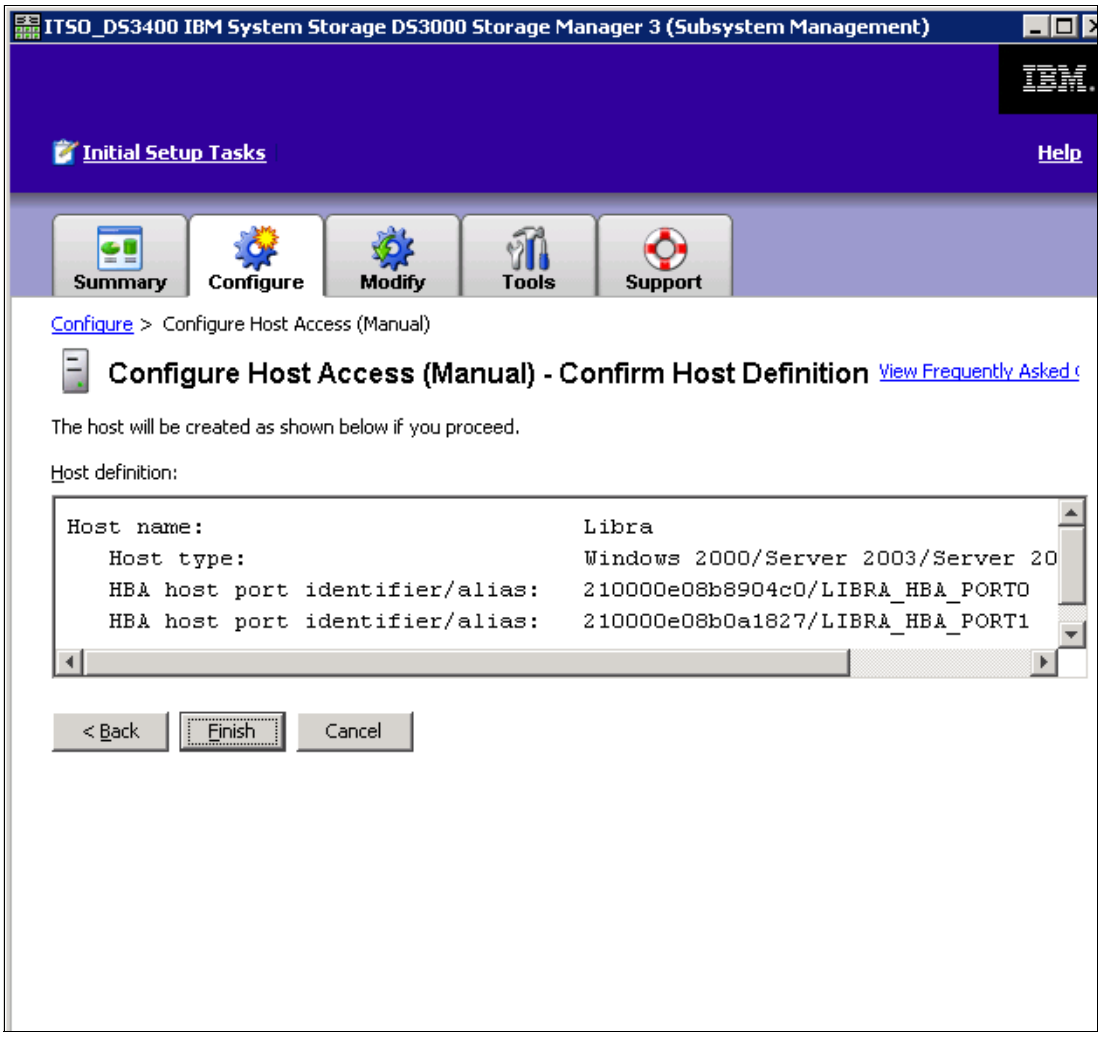


Figure 9-24 Confirm Host Definition

The host server is now available for logical drive mapping.

9.2.3 Configure host group

One of the options on the Configure tab is Configure Host Group. You have to define a host group if you want multiple hosts to share access to the same logical drive. We show how to do this when creating the host as a separate task here:

1. In the window shown in Figure 9-14 on page 141, click **Create Host Group** to display the Create Host Group window. Enter a unique name for your new host group, select the hosts that you want to add to the group, and click **Add** (Figure 9-25 on page 151).

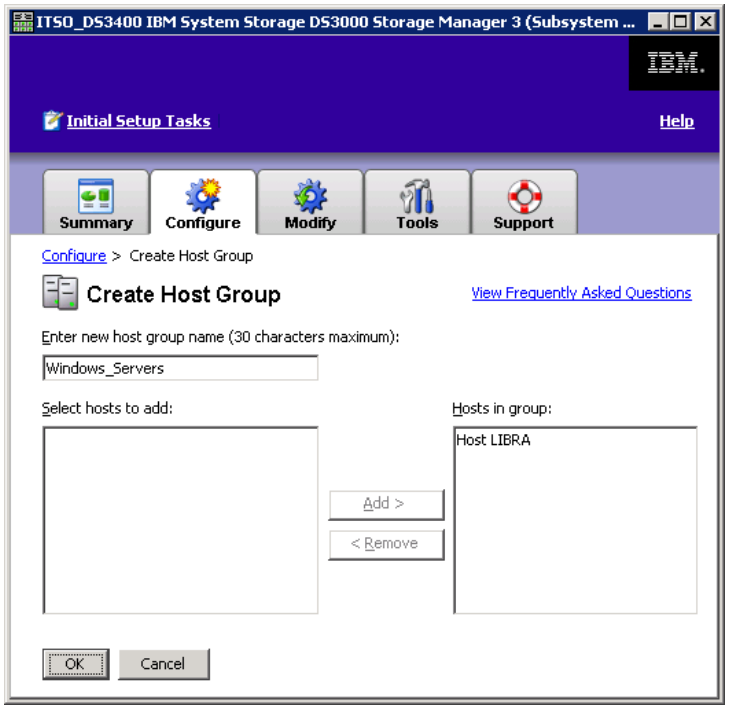


Figure 9-25 Create Host Group - enter name and add hosts

2. Click **OK** to confirm your settings and the new host group will be created, as shown in Figure 9-26. Click **OK** to return to the main Configure tab.

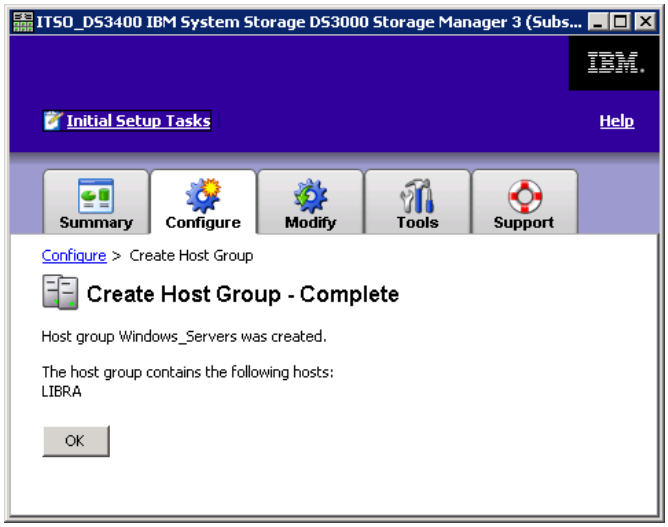


Figure 9-26 Create Host Group - Complete

9.2.4 Create host-to-logical drive mappings

This is the third part of the Configure tab (Figure 9-14 on page 141). We take a closer look at the host to logical drive mappings. First., we need to discuss storage partitioning.

What storage partitioning is

Storage partitioning adds a high level of flexibility to the DS3000 storage server. It enables you to connect to the same storage server multiple and heterogeneous host systems, either in stand-alone or clustered mode. The term storage partitioning is somewhat misleading, as it actually represents a host or a group of hosts and the logical disks they access.

Without storage partitioning, the logical drives configured on a DS3000 storage server can only be accessed by a single host system or by a single cluster. In most environments, this is undesirable, as we want to share the same storage system among multiple servers.

Storage partitioning, on the other hand, allows you to create sets, containing the hosts with their HBAs and the logical drives. We call these sets *storage partitions*. Now the host systems can only access their assigned logical drives, just as though these logical drives were locally attached to them. Storage partitioning adapts the SAN idea of globally accessible storage to the local-storage-minded operating systems.

Storage partitioning lets you map and mask LUNs (that is why it is also referred to as LUN masking). This means that once you have assigned a LUN to a host, it is hidden from all other hosts connected to the same storage server, so that access to that LUN is exclusively reserved for that host.

It is a good practice to do your storage partitioning prior to connecting multiple hosts. Operating systems like to write their signatures to any device they can access.

Heterogeneous host support means that the host systems can run different operating systems. But be aware that all the host systems within a particular storage partition must run the same operating system, as all host systems within a particular storage partition have unlimited access to all logical drives in this partition. Therefore, file systems on these logical drives must be compatible with all host systems. To ensure this, the most secure way is to define your storage partitions to include only hosts running the same operating system, even though some operating systems might be able to mount foreign file systems.

A storage partition contains several components:

- Hosts groups
- Hosts
- Host ports
- Logical drive mappings

The storage partition is a combination of all these components. It ensures proper access to the different logical drives even if there are several hosts or clusters connected. Storage partitioning associates a logical drive to a host or host group, so that every logical drive has its own mapping (logical drive mappings)

Every unassigned logical drive is mapped to the undefined mappings group. This means no host (or host port, to be precise) can access these logical drives until they are mapped.

A DS3000 supports up to 32 storage partitions (see Table 10-1 on page 217). Each time a logical drive is mapped to a new host or host group, a new storage partition is created automatically. If you map more logical drives to the same host or host group, this does not count as a new storage partition. For example, a cluster with two nodes with redundant I/O paths would be configured as one host group with two hosts. Each host would have two host ports for redundancy. Several logical drives would be mapped to this host group. All these components represent one storage partition. If you attach another unclustered host system to the same storage system and map a logical drive to that host, you create another storage partition. If you then define a new logical drive and map it to either the cluster or the single host, you are still using two storage partitions.

Heterogeneous hosts

When implementing a DS3000 storage server, you can use a mixture of different operating systems and clustered and non-clustered variants of the same operating systems. However, all logical drives in a single storage partition must be configured for the same operating system. Also, all hosts in that same storage partition must run the same defined operating system. Before connecting the systems, always check the latest interoperability information at:

http://www.ibm.com/servers/storage/support/disk/

Create host-to-logical drive mappings

Follow the below steps to create host-to-logical drive mappings.

1. From the Configure tab (Figure 9-14 on page 141), select **Create Host-to-Logical Drive Mappings**.
2. The wizard starts (Figure 9-27). Select the host or the host group that you want to assign to a logical drive. If you have not created all the required hosts and host groups yet, do this before continuing this wizard. You will see the current number of partitions used, and the maximum number of partitions allowed.

   The partitions are managed automatically for you by Storage Manager (you do not need to create or delete them), and they will be created and the number used will be updated whenever a new partition is required by a particular host/LUN mapping.

   Select a host group (as in Figure 9-27) or a single host and click **Next.**

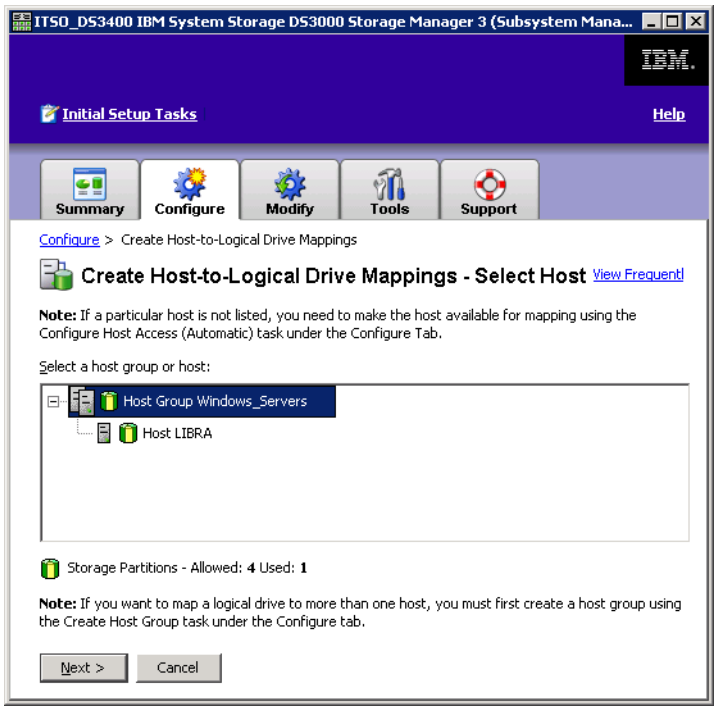


Figure 9-27 Storage Partitioning - Select Host

Note: Remember, if several hosts need to access one logical drive, create a host group and select the host group here.

3. In Figure 9-28, decide which logical drive to assign to the selected host or host group.

   This view displays all logical drives without a host mapping. Select the logical drive you want to assign, or check **Select all logical drives** to map all remaining logical drives to one host. Leave the LUN number as the default, or if you have special requirements, select a LUN number and click **Finish**.

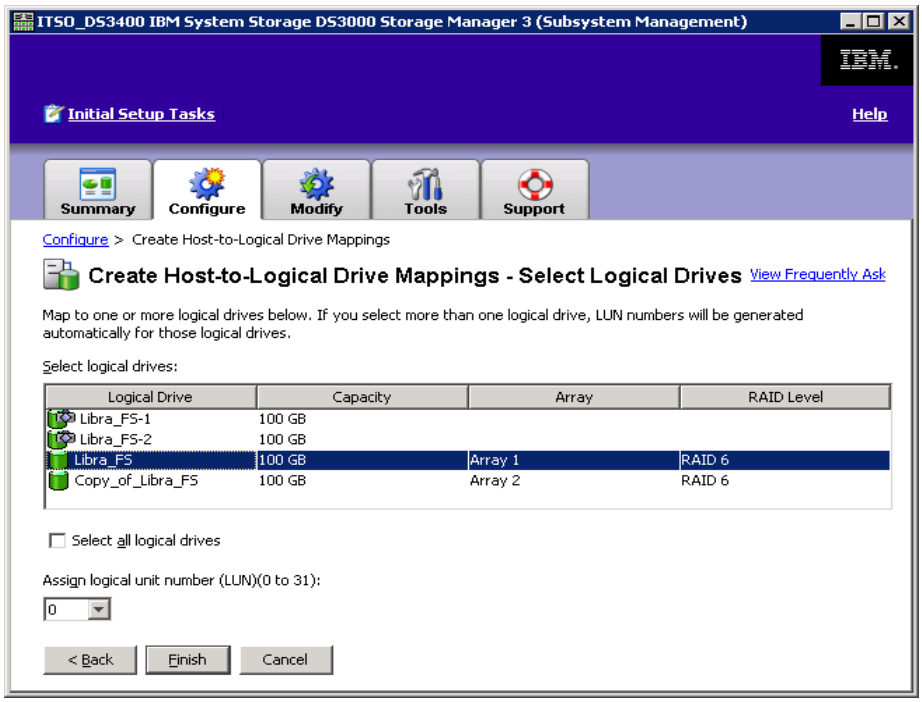


Figure 9-28 Storage Partitioning - Select Logical Drives

4. The wizard will perform the storage partitioning. You will see a completion message, as shown in Figure 9-29, when partitioning is complete. Click **OK to** exit. You will be prompted to create another host-to-logical drive mapping or to return to the main Configure tab.

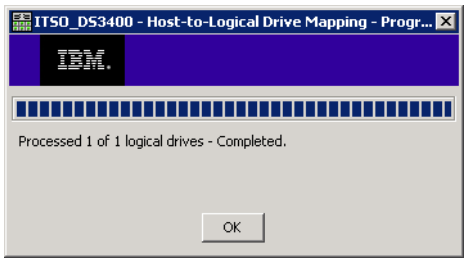


Figure 9-29 Storage Partitioning - complete

Note: When you run the logical drives wizard, it prompts you at the end, asking if you want to map the generated drive now. If you do map the drive at this point, the storage partitioning is also created then. However, if you prefer to do it at a later point of time, or if you had not yet configured all the hosts and host groups, you can add storage partitions as shown in this section.

9.3 Configure storage

In the second part of the Configure tab (Figure 9-14 on page 141), you will find the following options:

- Automatic Configuration
- Configure Hot Spares
- Create Logical Drives
- Create FlashCopy Logical Drives
- Create Logical Drive Copies

The first three tasks are considered basic functions, and Create FlashCopy Logical Drives and Create Logical Drive Copies are considered advanced functions.

These tasks set up the arrays and logical drives on your storage subsystem. Arrays are sets of physical drives that the controller logically groups together to provide one or more logical drives to an application or an cluster. A logical drive is a logical structure you create on a storage system or data storage.

Creating arrays is one of the most basic steps and is required before you can start using the physical disk space, that is, you divide your disk drives into arrays and create one or more logical drives inside each array. The first three wizards allow you to set up this basic configuration on your storage subsystem. More information about arrays and their RAID Levels can be found in 5.1, “Storage subsystem concepts” on page 51.

9.3.1 Configuring storage - automatic

Automatic configuration is a quick and easy way to create arrays and logical drives and begin using the storage subsystem in the shortest time possible. You only have to specify the RAID level; everything else is done automatically. When the arrays and logical drives are created, you will have to map the host servers or groups to the new logical drives. Figure 9-30 shows the shortcuts to the automatic storage subsystem configuration tasks.

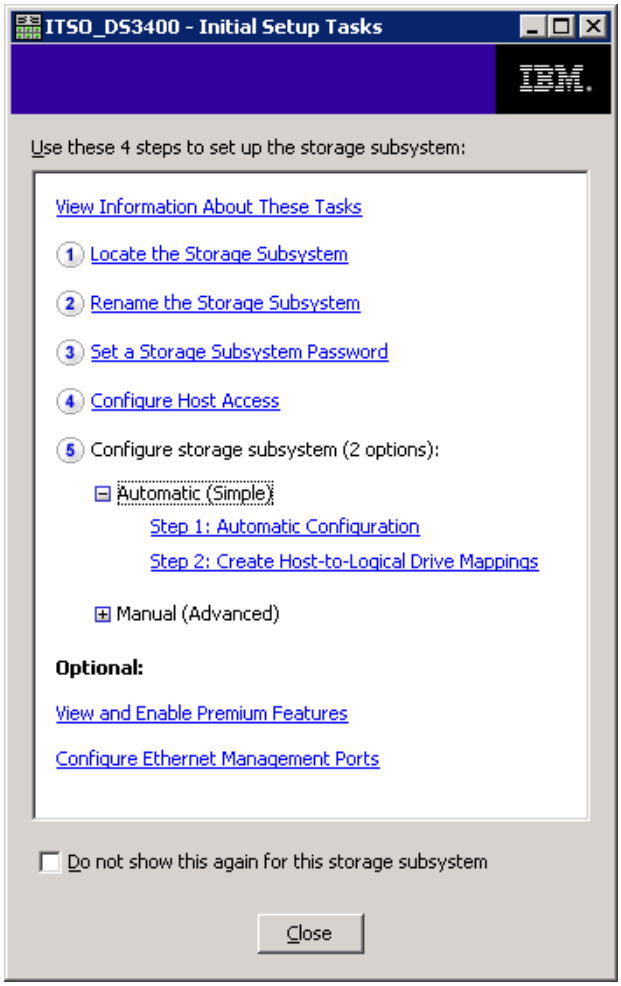


Figure 9-30 Configure storage subsystem - automatic

We will discuss option 5 from Figure 9-30; for more information about the rest of the tasks in this window, refer to 5.4, “Initial Setup Tasks wizard” on page 71. The tasks listed under option 5 are:

Configure Storage Subsystem (two options):

- Automatic (Simple)
  - Step 1: Automatic Configuration
  - Step 2: Create Host-to-Logical Drive Mappings
- Manual (Advanced)
  - Step 1: Configure Hot Spare Drives
  - Step 2: Create Logical Drives

Automatic Configuration

As stated, this is a one-click way to set up the arrays and logical drives. The only parameter you need to specify is the RAID level; Storage Manager will then group any unconfigured physical disk drives into arrays and carve the arrays into logical drives. Figure 9-31 shows an example of an automatic configuration.

In the drop-down list, select the RAID level, from 0, 1, 3, 5, or 6. To select RAID 10, select 1; if there are at least four drives in the array, then RAID 10 will be automatically used.

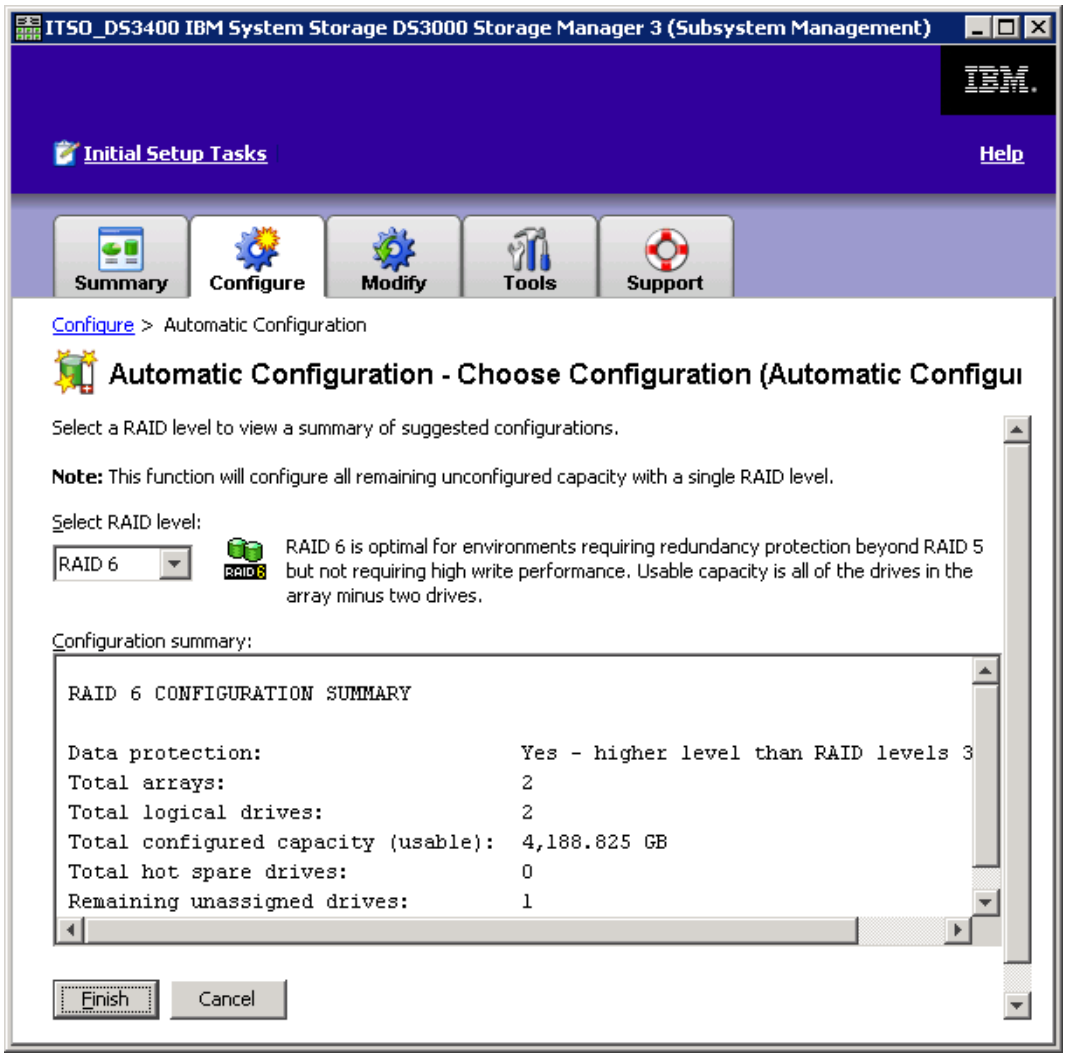


Figure 9-31 Automatic Configuration - Choose Configuration

We selected RAID 6. The Configuration summary shows that two arrays and two logical drives will be created. Click **Finish** to proceed. This configuration starts, as shown in Figure 9-32.

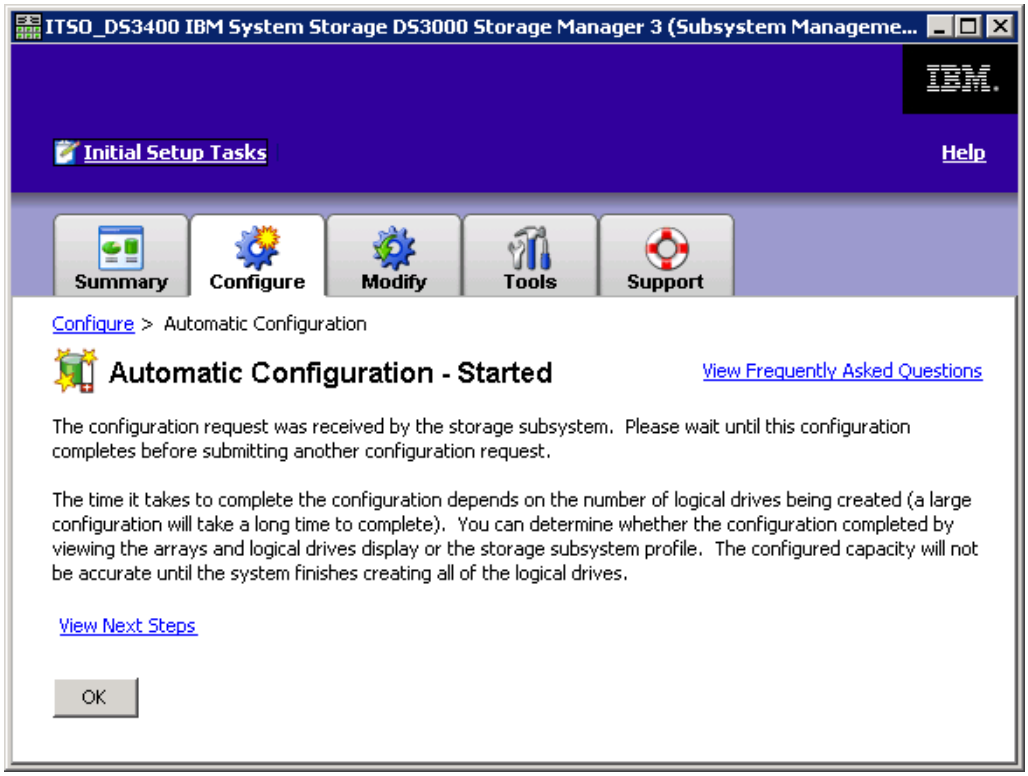


Figure 9-32 Automatic Configuration - Started

We can now verify what was created. A convenient way to do this is to first view the Operation Progress status, as shown in Figure 9-33.

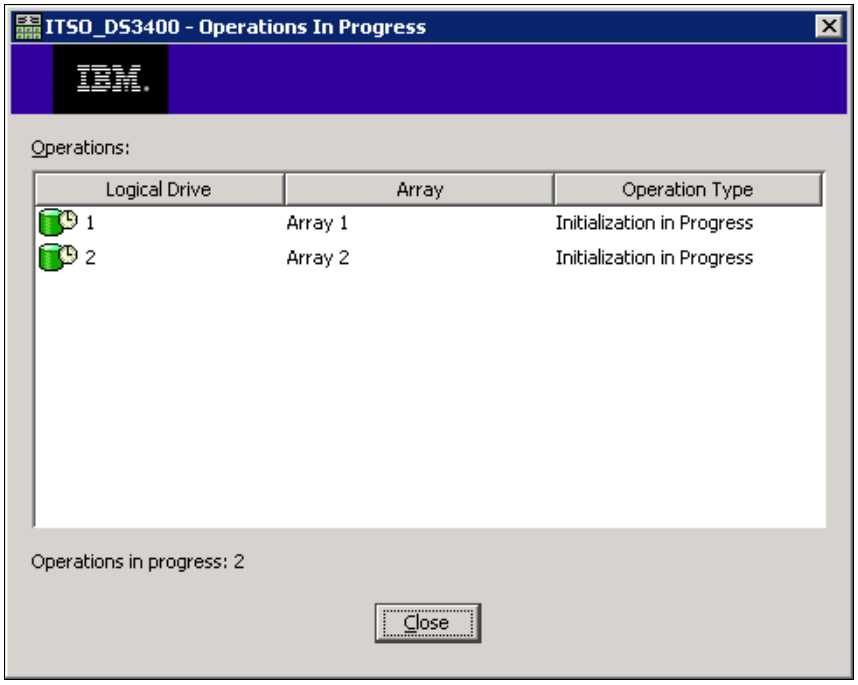


Figure 9-33 Operation status progress view

Then you can click **Array and logical drives** in the Summary tab to view the Array Status. As shown in Figure 9-34, two RAID 6 arrays (numbered 1 and 2) were created. Each array contains one logical drive, called simply numbers 1 and 2.

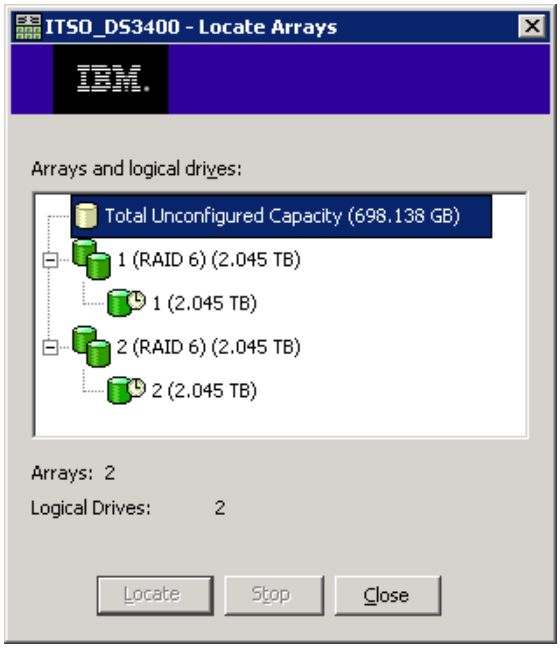


Figure 9-34 Two new arrays and logical drives

While simple to use, automatic configuration gives you only limited control over how the arrays and logical drives will actually be set up. You can only select the RAID level; all other parameters (such as the number of arrays, number and position of disk drives in the arrays, sizes of logical drives, and logical drive I/O settings) are chosen automatically and are outside of your control. The resulting configuration might not suit your requirements. For this reason, we recommend manually configuring the storage subsystem instead. Detailed instructions for manual storage subsystem configuration of arrays and logical drives are given in “Configuring the storage subsystem (manual)” on page 163.

Create host-to-logical-drive mappings

This task allows us to map a host server or a group to the logical drives we just created.

1. Select the host server (or a group) that you want to map to the drives, as shown in Figure 9-35.



Figure 9-35 Create Host-to-Logical Drive Mappings - Select Host

In our case, we select the host server **Libra**, and click **Next**.

2. Select one or more logical drives to map to the host or group. We selected both the logical drives just created, called 3 and 4 (see Figure 9-36).

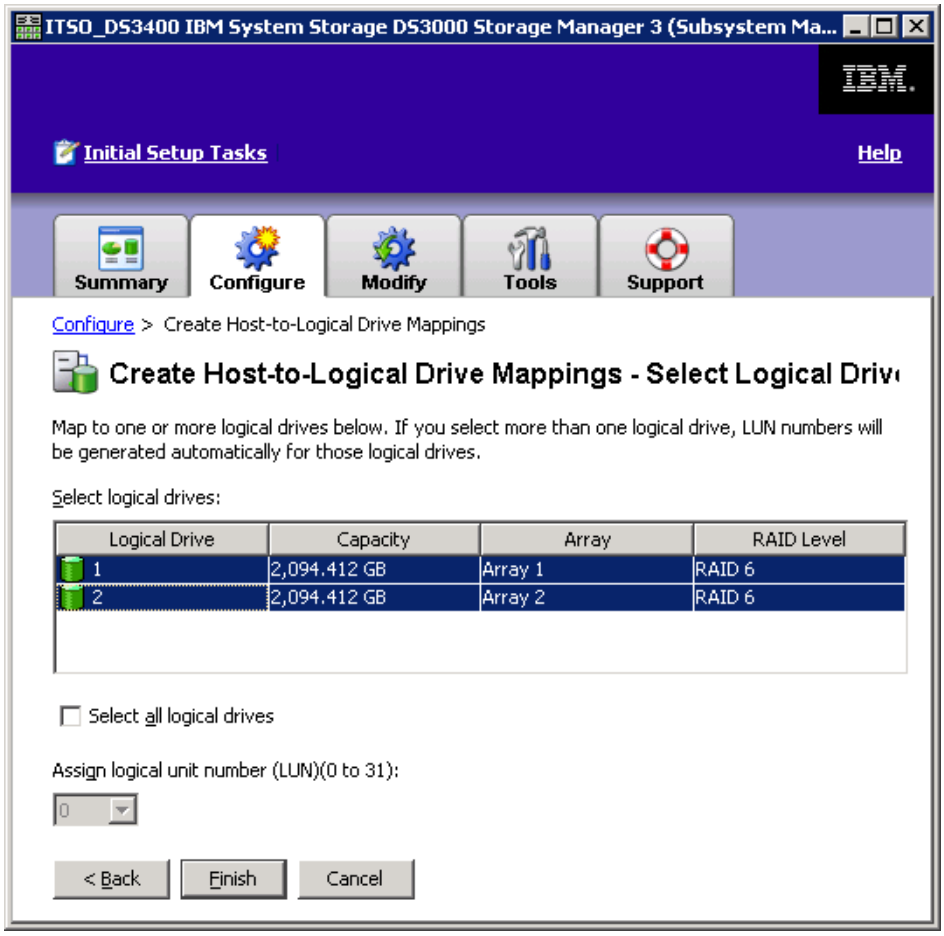


Figure 9-36 Create Host-to-Logical Drive Mappings - Select Logical Drives

3. Click **Finish**, and you will see a progress bar while the mappings are created. When the task completes, you will see the confirmation window, as shown in Figure 9-37.

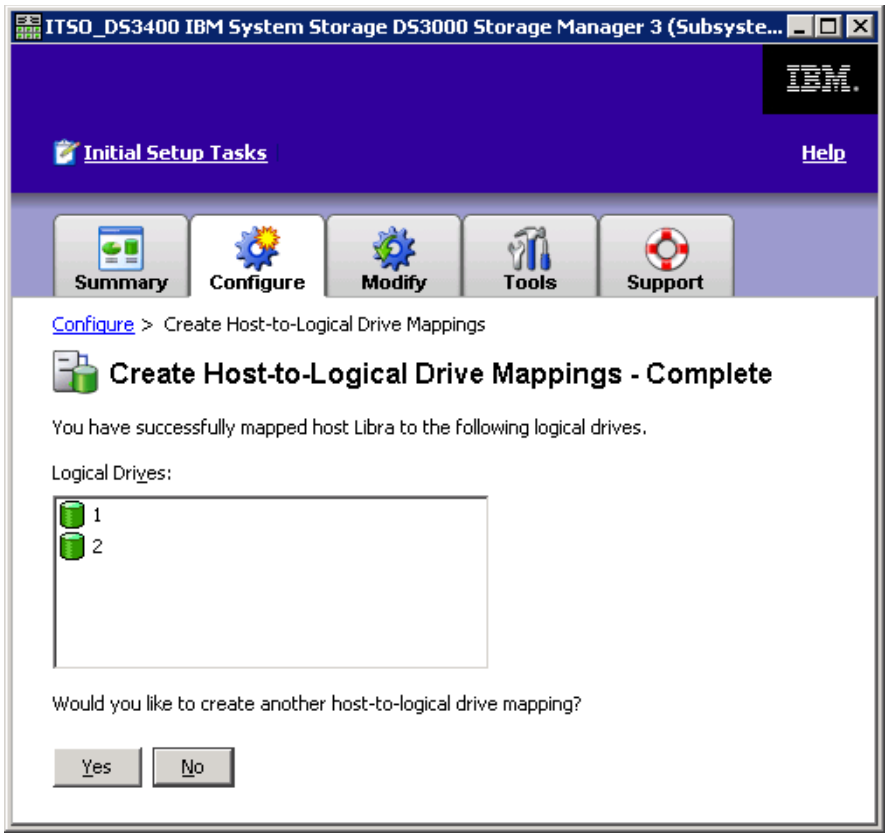


Figure 9-37 Create Host-to-Logical Drive Mappings - Complete

The two new logical drives should now be accessible from the host server Libra.

Configuring the storage subsystem (manual)

Figure 9-38 shows the Configure Storage Subsystem - manual task expanded and highlighted.

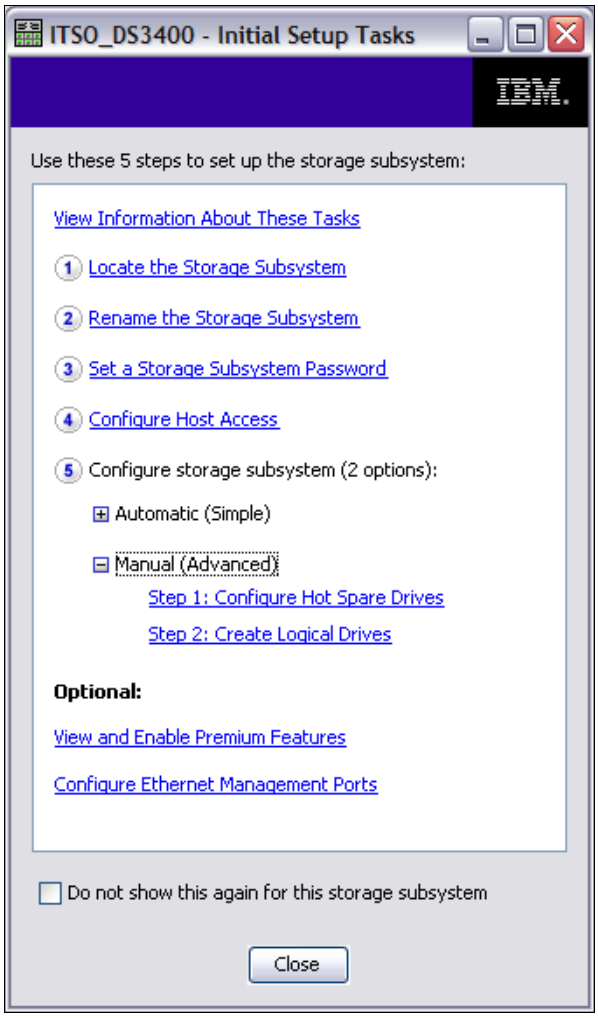


Figure 9-38 Configure Storage Subsystem - manual

The Step 1: Configure Hot Spare Drives option has two additional options, as shown in Figure 9-39.

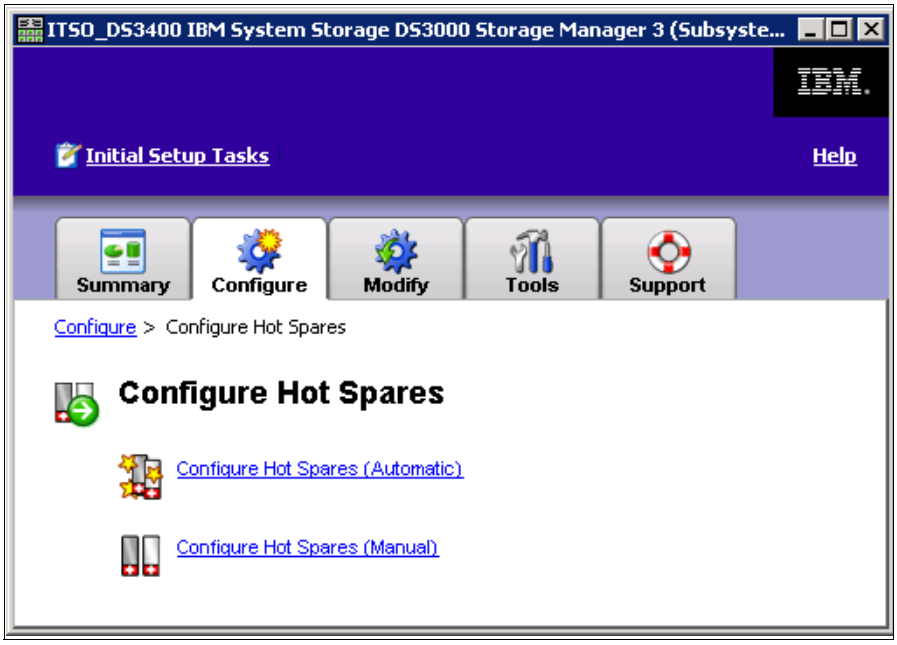


Figure 9-39 Configure Hot Spares - options

Configure Hot Spare Drives - Automatic

There are two ways you can do hot spare assignment. By choosing the automatic method, hot spare can be assigned automatically based on the enclosure connection. Storage Manager assigns the appropriate number of hot spare drives to provide optimal coverage for all physical disks. Figure 9-40 on page 165 shows the automatic hot spares configuration window.

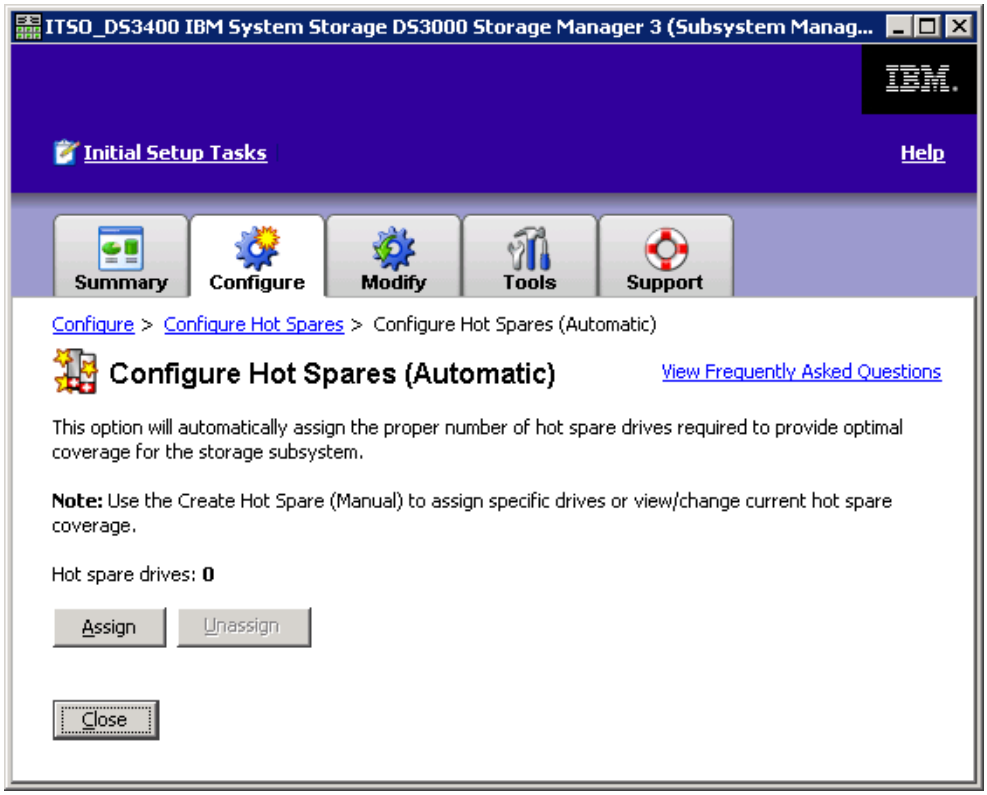


Figure 9-40 Configure Hot Spares (Automatic)

Click **Assign** and the hot spare configuration will take place. When it finishes, the completion window appears, as shown in Figure 9-41.

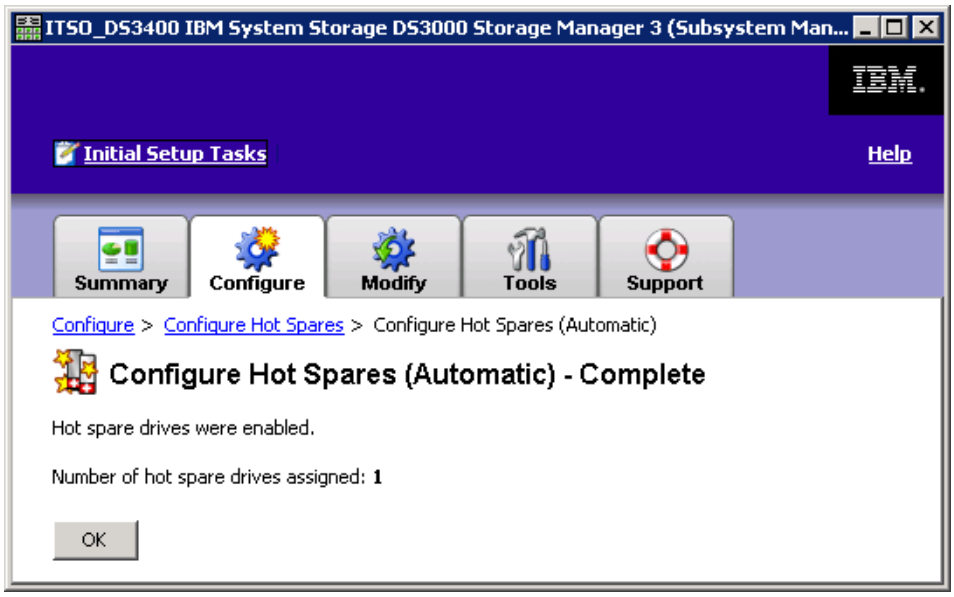


Figure 9-41 Host spares configuration complete

Configure Hot Spare Drives - Manual

Click the manual option to assign hot spare manually, as shown in Figure 9-39 on page 164.

Array information and number of disk information will be displayed in the next window. Click the **Assign** button, as shown in Figure 9-42.

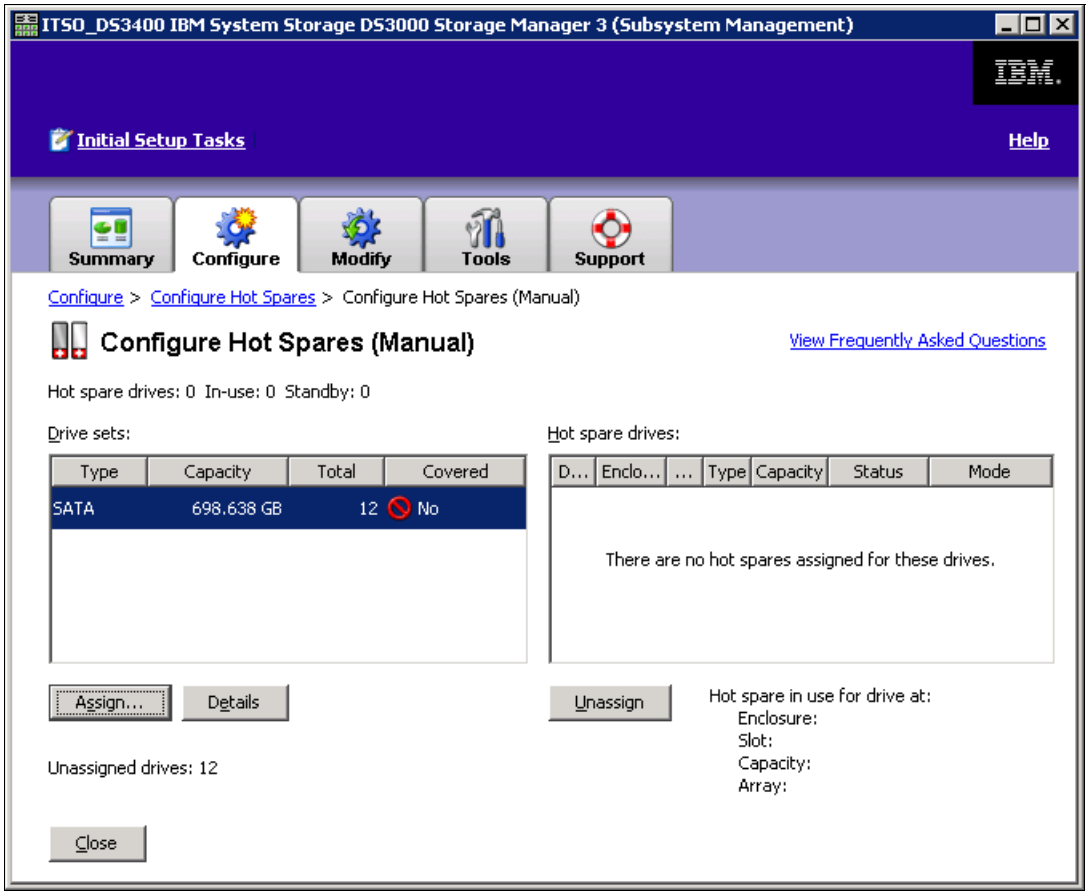


Figure 9-42 Configure Hot Spare (Manual)

The hard disk information and location will be displayed accordingly. Select the hard disk to assign the hot spare disk, as shown in Figure 9-43.

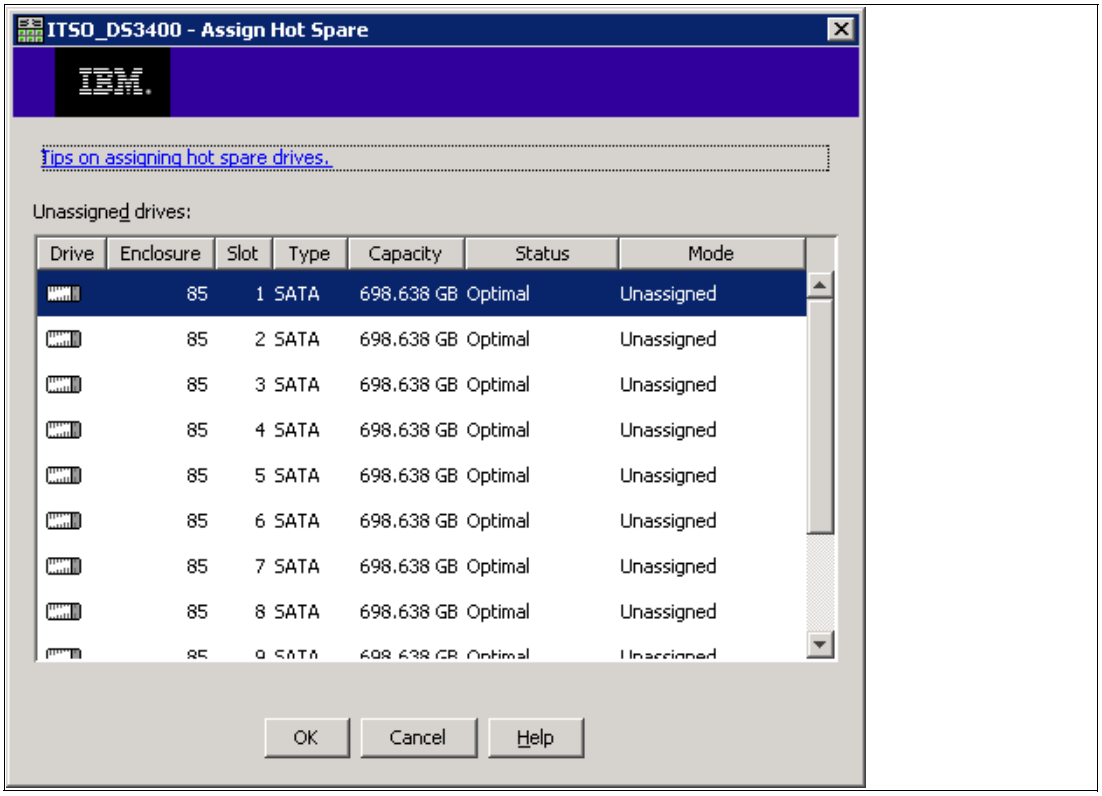


Figure 9-43 Hard disk selection and location (manual)

Select any hard disk to assign. In this case, hard disk 12 has been selected as a hot spare (Figure 9-44). Also, you can select more than one hot spare if your application insists.

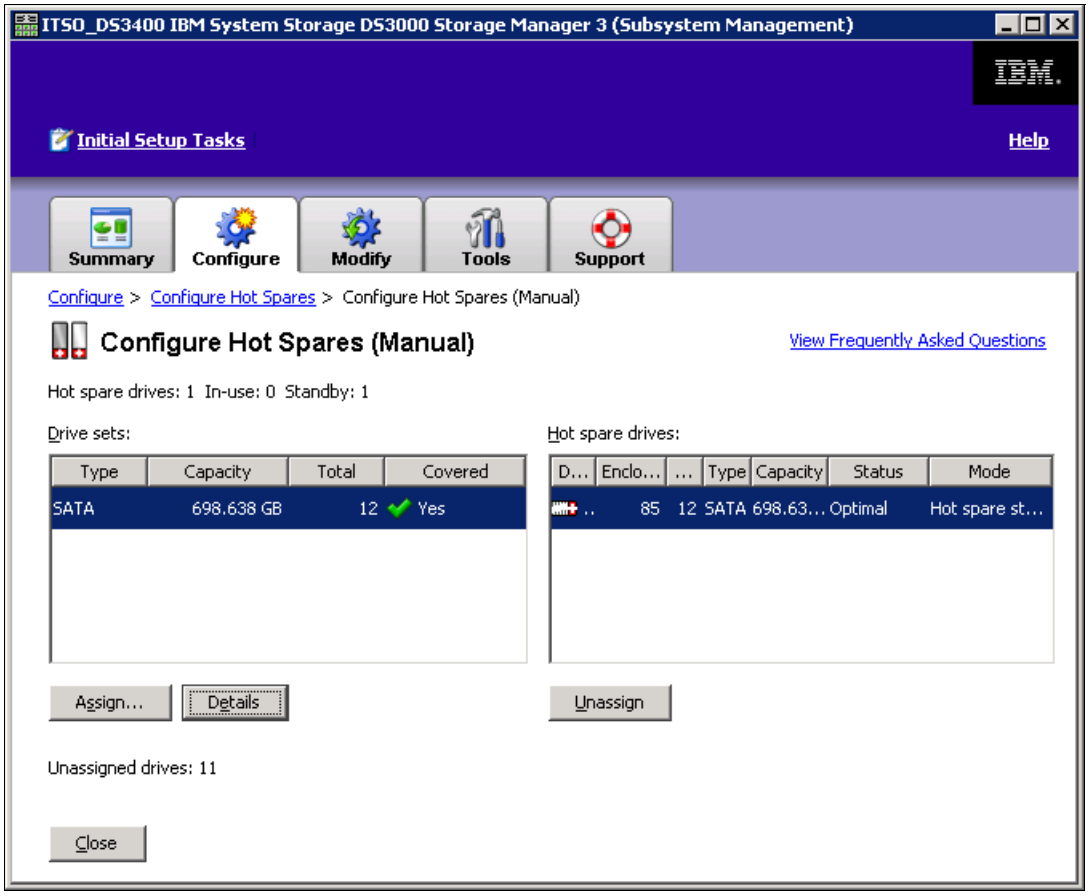


Figure 9-44 Assigned hot spare (manual)

Create logical drives

This action is more than just a logical drive creation; it also includes array creation (if needed) and host-to-logical-drive mapping. Before you create logical drives, you should understand the differences between the possible RAID levels, as described in 5.1, “Storage subsystem concepts” on page 51.

A logical drive is the entity that is available to the host attached to the storage subsystem. You can make arrays out of the unconfigured capacity of a storage subsystem, and you can make logical drives in the free capacity of an array. Unconfigured capacity is the disk drives that are not already assigned to an array. Free capacity is space in an array that has not been assigned to a logical drive.

Figure 9-45 shows a configured storage subsystem with nine disk drives. There is one array of six disk drives, which has been configured into three logical drives. One disk drive has been configured as a hot spare and two disk drives remain unconfigured.

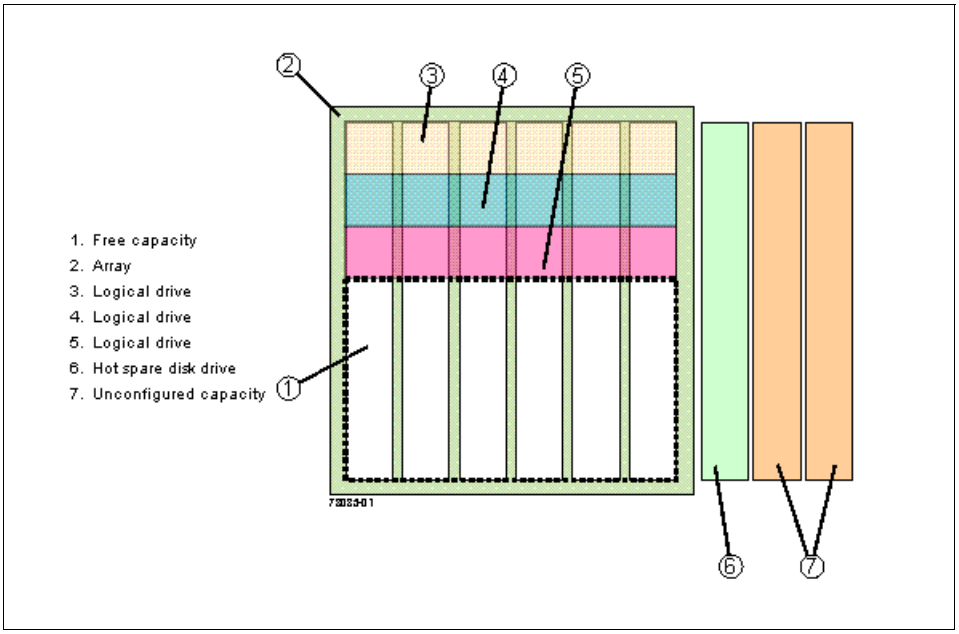


Figure 9-45 Array and logical drive configuration

With the Create Logical Drives wizard (Figure 9-14 on page 141), you can:

- Create arrays and logical drives from unconfigured capacity.
- Use free capacity in existing arrays to create logical drives.

To start, click **Create Logical Drives**. You will see two options (Figure 9-46).

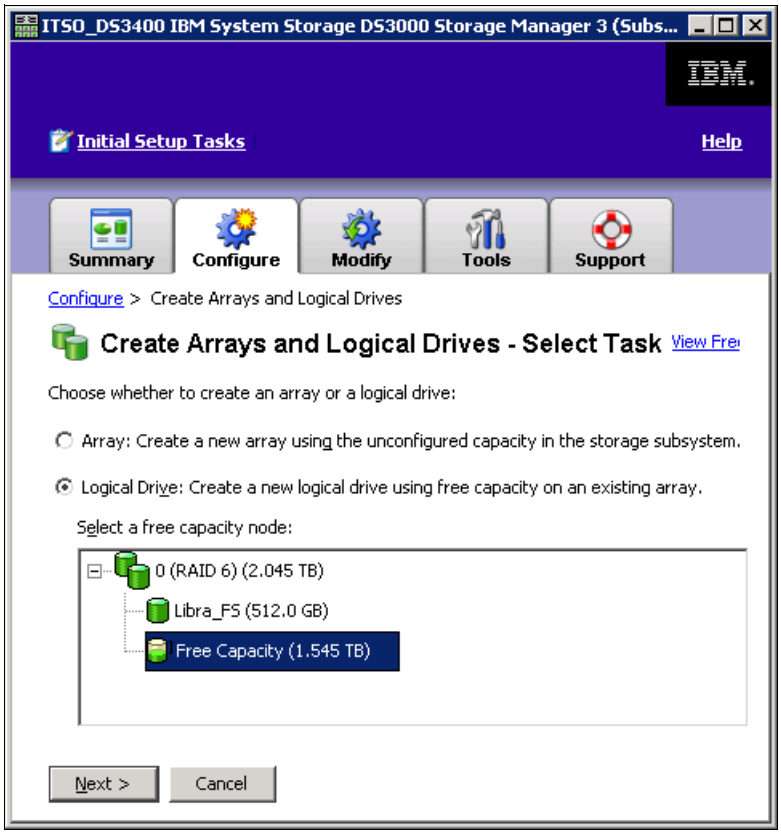


Figure 9-46 Select capacity type

Note: If you are configuring a new system from scratch, you will not see this selection because you have not created any arrays.

Unconfigured capacity

To create a new array from unconfigured capacity, select the first option from Figure 9-47 on page 171, and click **Next**.The options for Free Capacity will be grayed out as there are no arrays currently defined.

Do the following steps:

1. First, we have to select **Array: Create a new array using the unconfigured capacity in the storage subsystem**. We show this in Figure 9-47.

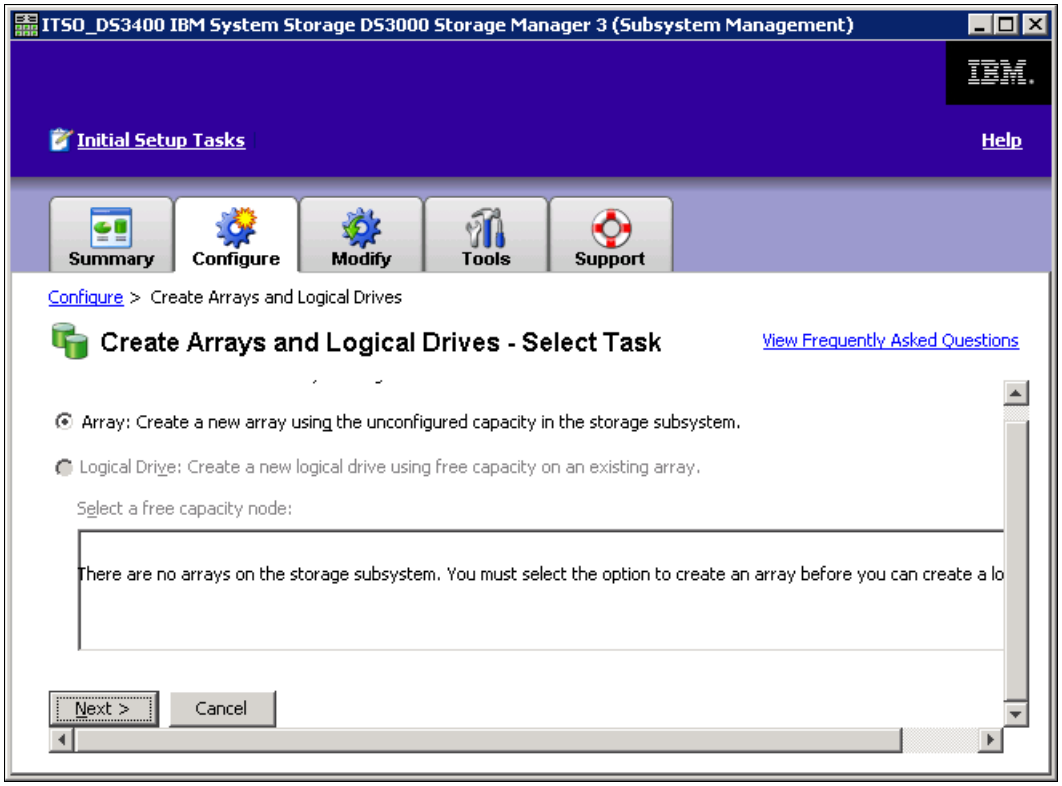


Figure 9-47 Create Logical Drives - select capacity type

2. Next we define the array. This requires two parameters: RAID level and a set of physical disk drives to be used in the array. As you can see in Figure 9-48, the array can be defined automatically (recommended) or manually.

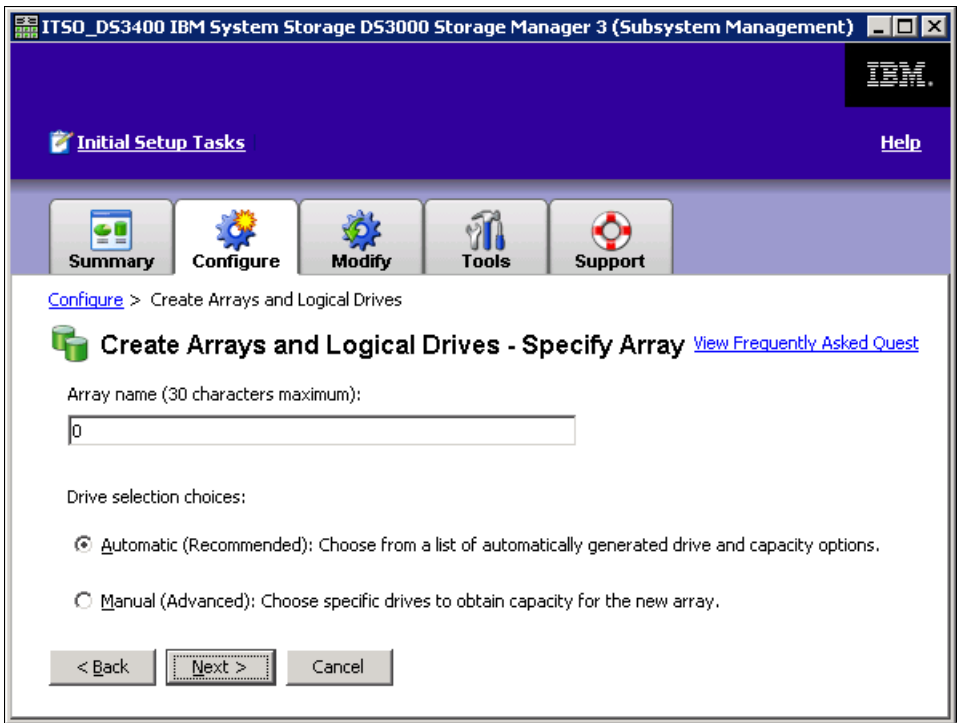


Figure 9-48 Create Logical Drives - Drive selection choices

To define an array automatically, continue with step 3.

To define an array manually, go to step 4 on page 173.

3. To define the array automatically, first select the RAID level. Depending on your selection, you will then be presented with multiple possible array configurations of different numbers of physical drives that may or may not offer enclosure loss protection. Enclosure loss protection means that all array members are in different DS3000 and EXP3000 enclosures.

   You can choose between RAID 0, 1, 3, 5 or 6; if you pick RAID 1 and select four or more drives, then RAID 10 (striping and mirroring) will be automatically used on the array. If you select RAID 6, then it creates RAID 6, as shown in Figure 9-49 on page 173.

   Select the array capacity that best suits your needs and click **Next** to continue. This will create the array with selected parameters.

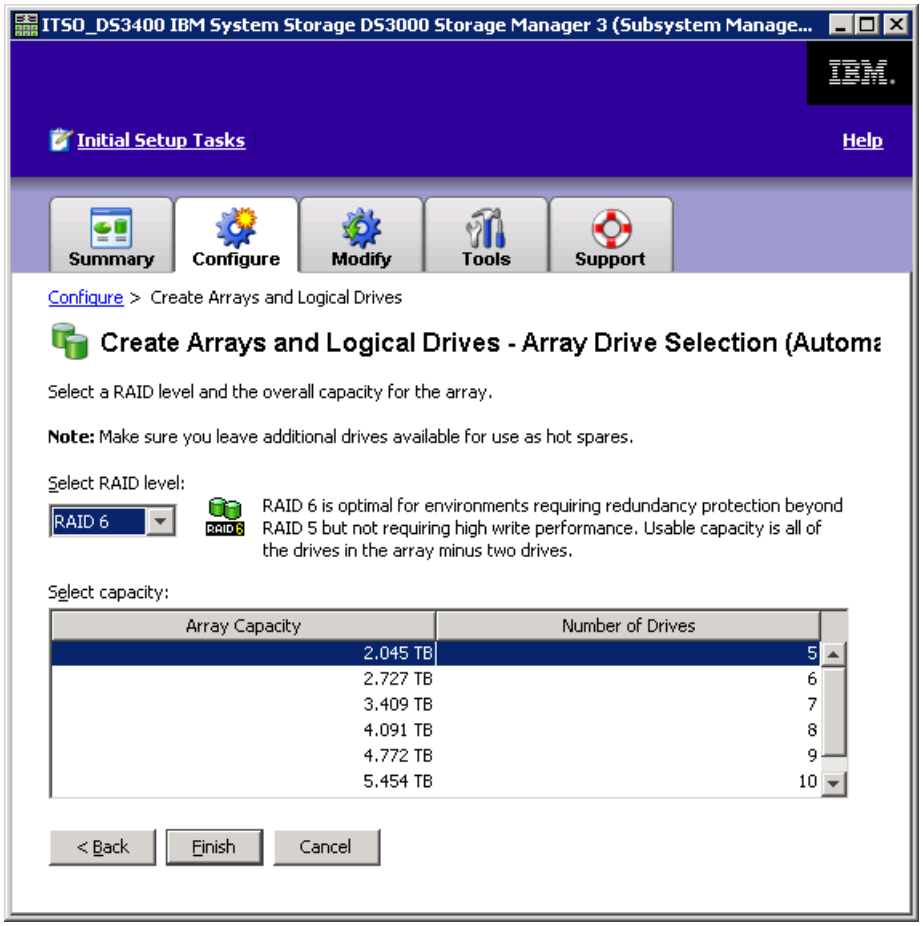


Figure 9-49 Create Logical Drives - specify capacity (automatic)

In our case, we selected a RAID 6 array across two physical drives. The drives are in different enclosures, so we do have enclosure loss protection.

Go to step 5 on page 176 to continue with logical drive definition.

4. This is the manual array definition step. You need to specify the RAID level and then select the physical drives that will become array members. The number of drives will depend on the RAID level:
   - For RAID 0, you can select any number of drives.
   - RAID 1 requires an even number of drives. If you select two disks, this will indeed be a RAID 1 array. If you select four or more, this will actually become a RAID 10 array.
   - RAID 3 and 5 require at least three disk drives.
   - RAID 6 requires a minimum of five disk drives (max 30 drives)

We show the manual drive selection in Figure 9-50.

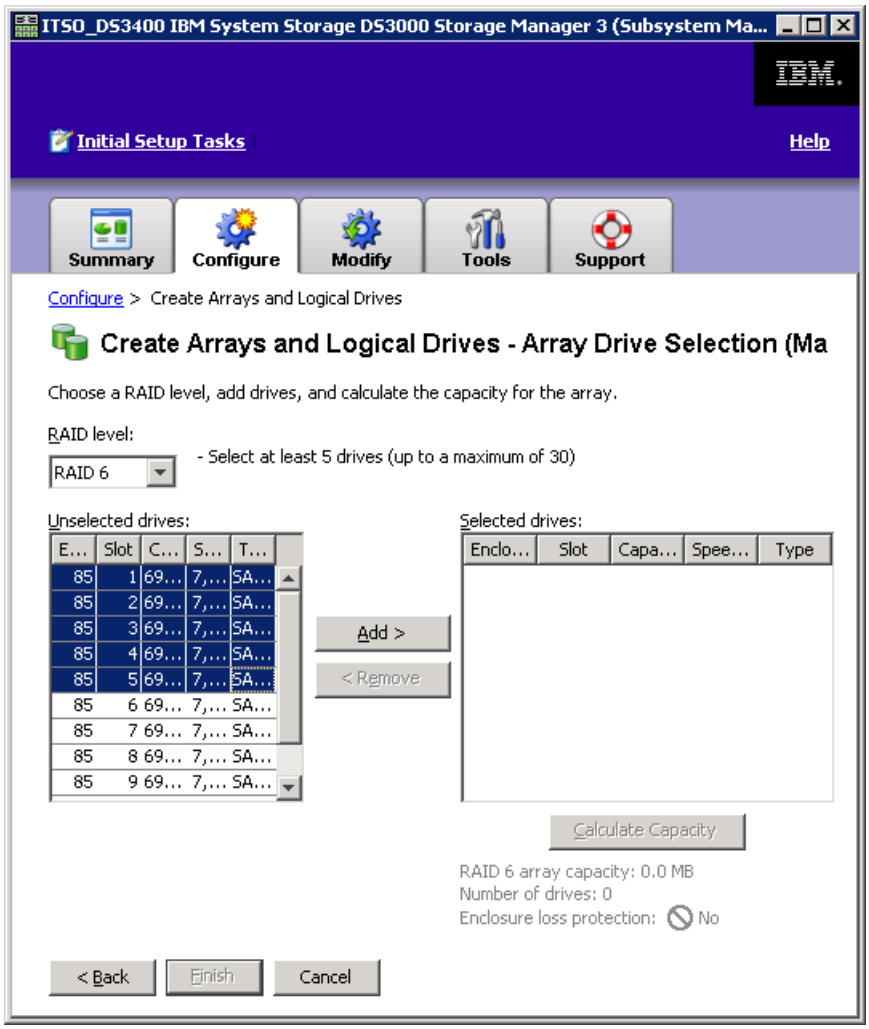


Figure 9-50 Create Logical Drives - manual drive selection

After selecting the RAID level, highlight the drives you wish to add to your array, then click **Add**. The chosen drives will appear in the Selected drives window. Now click **Calculate Capacity** and you will see the array summary (capacity, number of drives, and enclosure loss support), as shown in Figure 9-51.

Click **Next** to create the array with selected parameters, then proceed to step 5 on page 176.

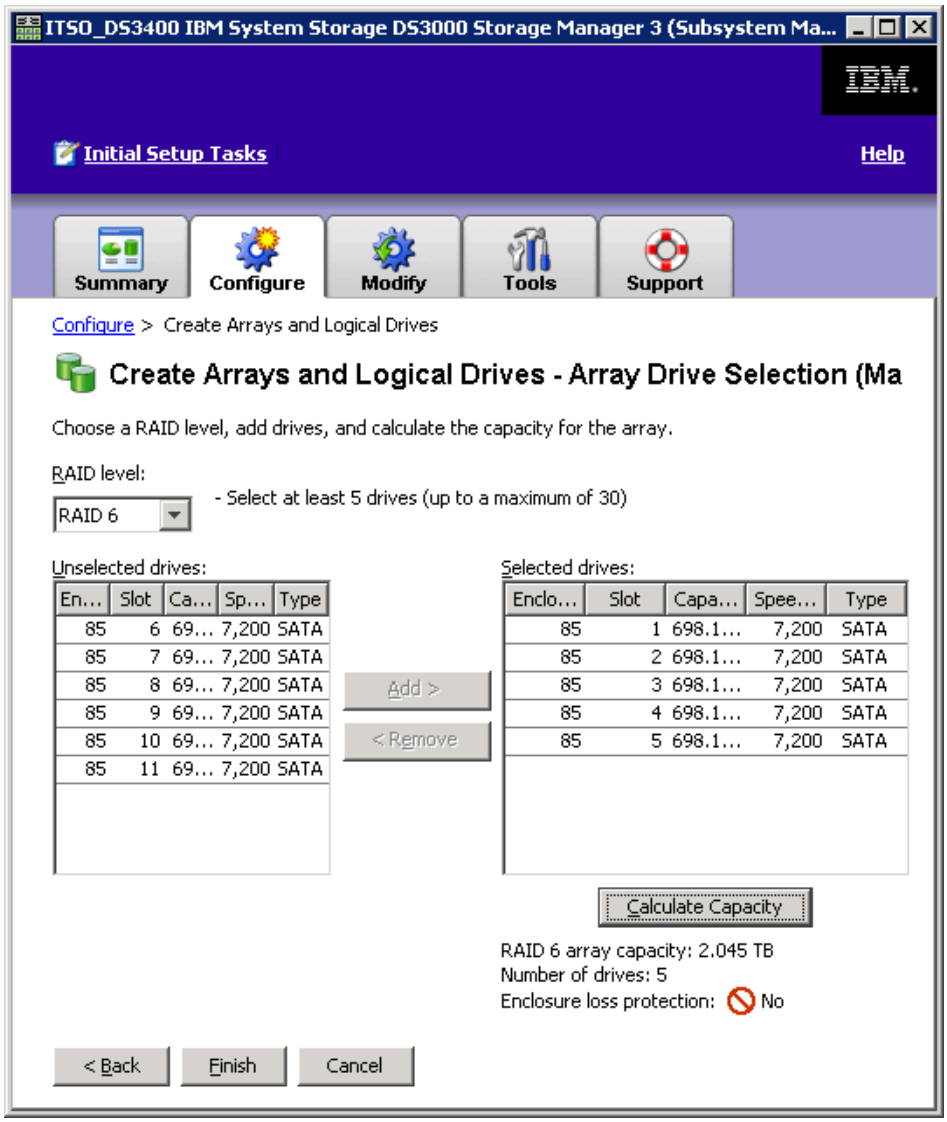


Figure 9-51 Manual Drive Selection - Calculate Capacity

5. After the array is defined, we then create a logical drive, as shown in Figure 9-52 on page 177.

   Logical drive definition requires several parameters:

   – Logical drive size

     You can either create one large logical drive across the entire array, or you can carve the array into several smaller logical drives. In our example, we created a logical drive of 512 GB. The total array capacity is 2.05 TB; the remaining free capacity can be used for additional logical drives.

   – Logical drive name

     Each logical drive must have a unique name. The name will be used in all management tasks to refer to a particular logical drive. For clarity, the name should reflect the role of the logical drive.

   – Logical drive I/O characteristics

     The logical drive I/O access pattern can be very different, depending on the host server role. For example, file servers use the disk storage in a different way than database servers. For best performance, it is important to match the logical drive parameters, such as segment size and cache pre-fetch, to the actual I/O access pattern. Select one of the predefined options (file server, database server, or multimedia), as shown in Figure 9-52 on page 177.

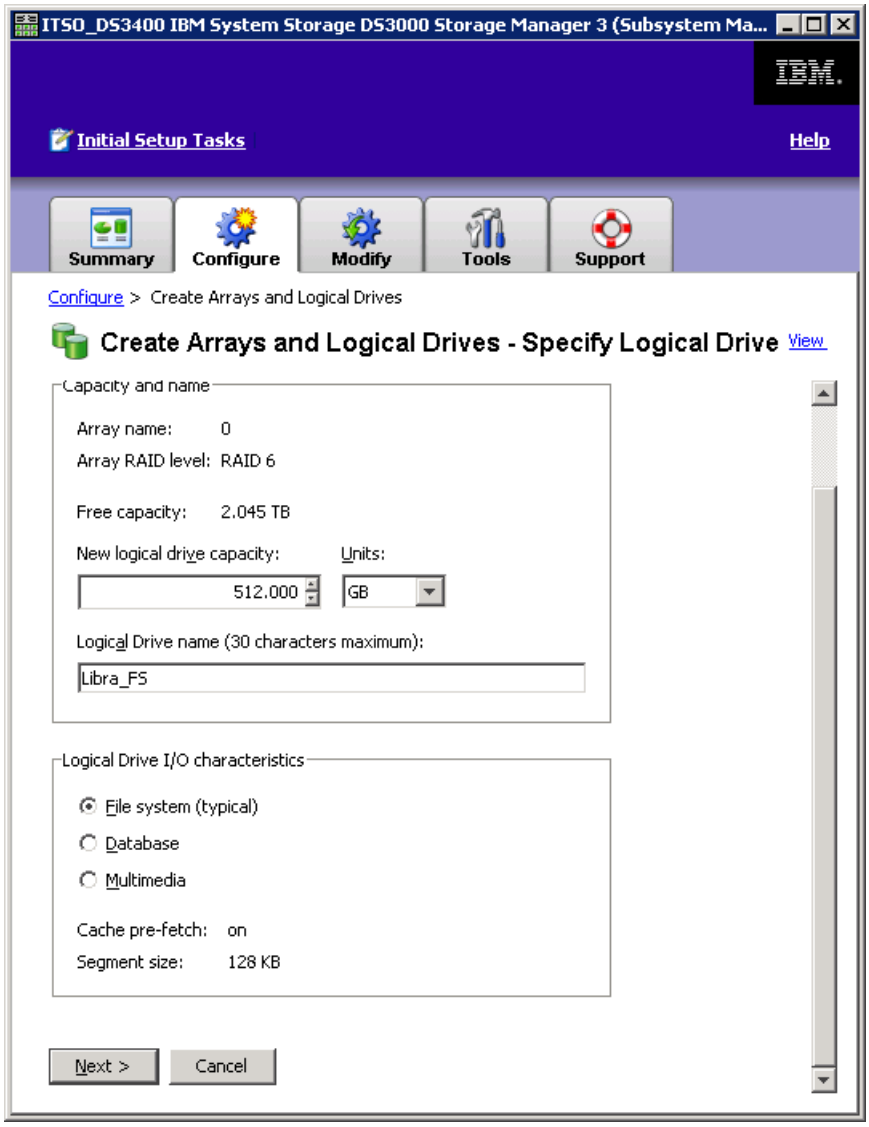


Figure 9-52 Create Logical Drives - Specify Logical Drive

In our case, we named our logical drive Libra_FS, set capacity to 512 GB, and configured it for File system use.

6. The final step is to configure the host-to-logical-drive mapping. The newly created logical drive must be assigned to a host server (or to a group of clustered servers). You can decide to either Map now or Map later.

   We will do the mapping right away, so we need to provide two parameters:

   – A host or a host group that will have access to this logical drive.

   – Logical Unit Number (LUN) of our logical drive in the host.

As you can see in Figure 9-53, we decided to map the logical drive to host server Libra and assigned it as LUN 0.

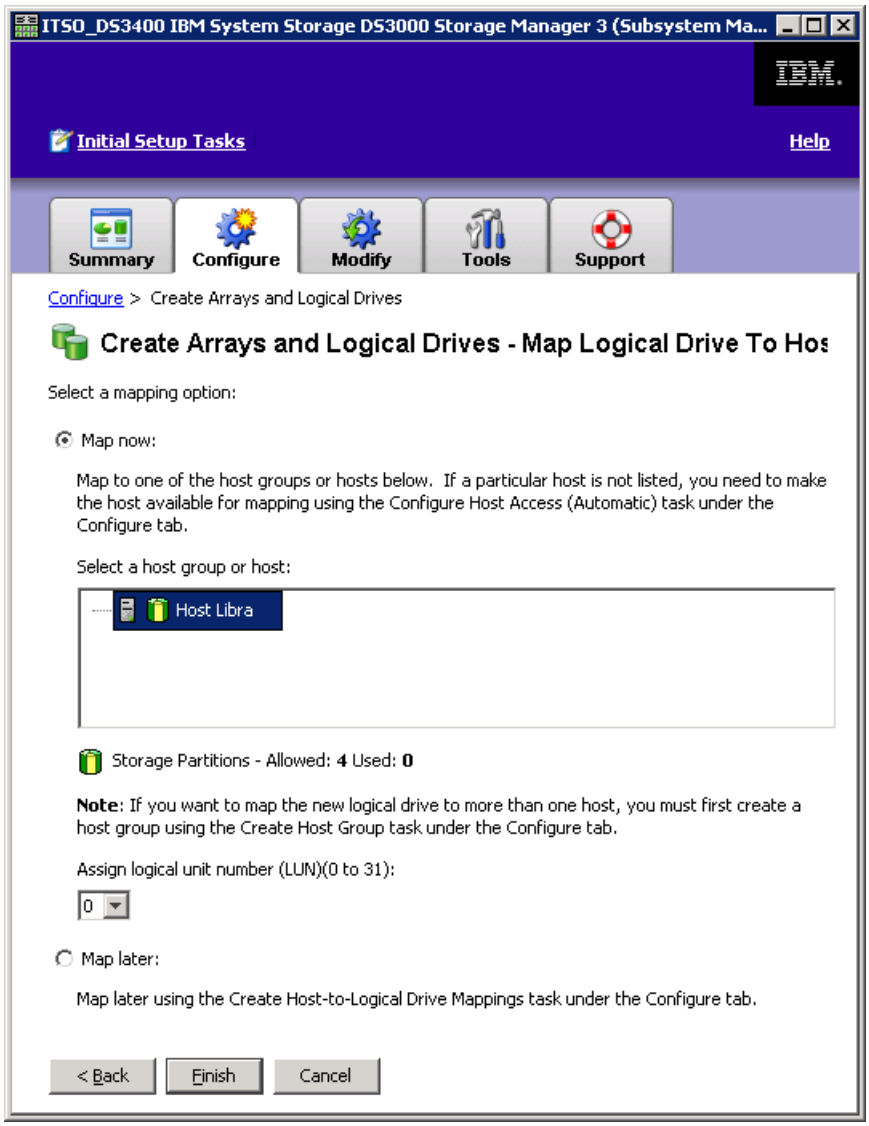


Figure 9-53 Create Logical Drives - Map Logical Drive To Host

Click **Finish** to complete the manual logical drive creation. This brings you to the final confirmation window (see Figure 9-54). The new logical drive is now ready for use.



Figure 9-54 Create Logical Drive - Complete

Free capacity

In this section, we show you how to use free capacity. If free capacity is available in an existing array, you must remember the RAID for the existing array. You cannot change the RAID for the logical drive. If you need a different RAID, then you must configure a new array, if you have space, or select a different array.

Select **Logical Drive: Create a new logical drive using the free capacity on an existing array** and highlight the free capacity of the array on which you want to create a new logical drive. Click **Next**, as shown in Figure 9-55.

.

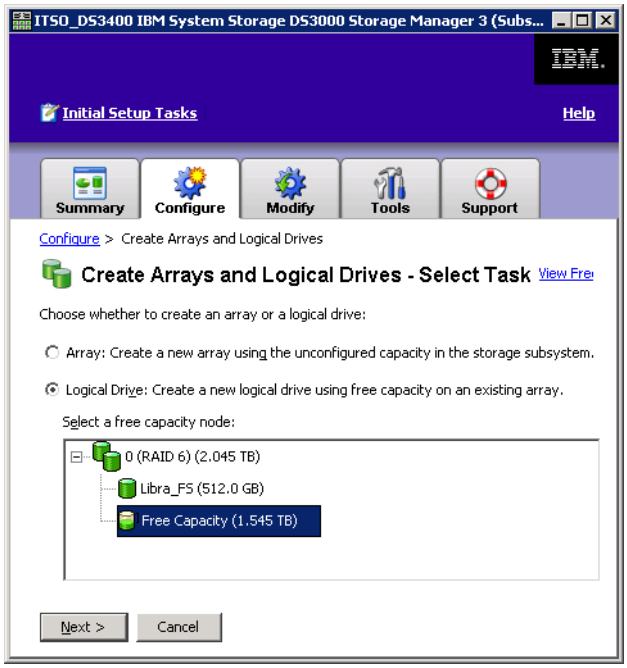


Figure 9-55 Create logical drive in free capacity

Specify the size of the logical disk and the logical drive I/O characteristics and then click **Next**, as shown in Figure 9-56.

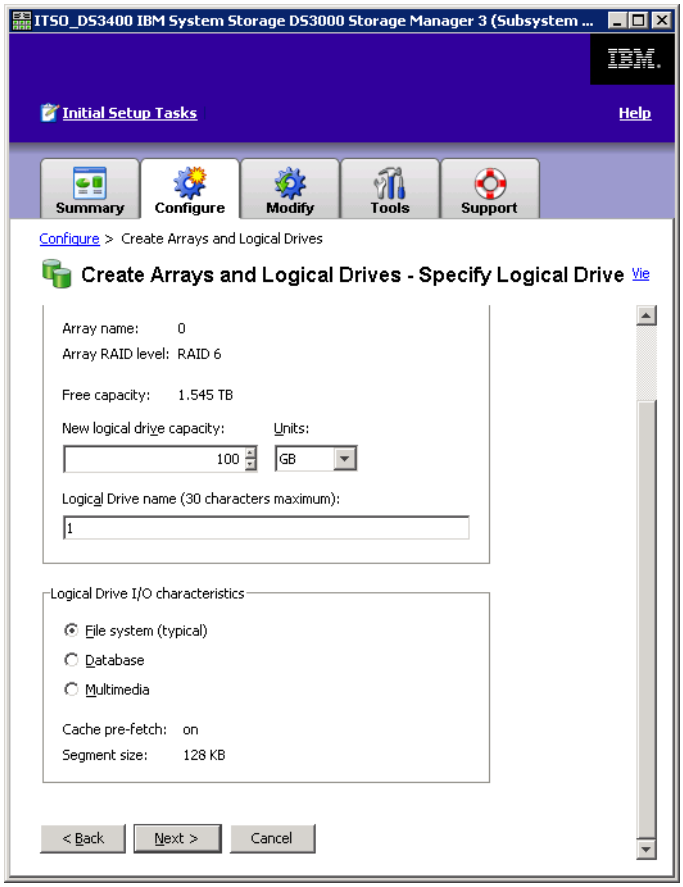


Figure 9-56 Creating Logical Drive - size

Once you have set the size and I/O characteristics, you will map the host. You can do it now or later. In this example, we have mapped it now, as shown in Figure 9-57.

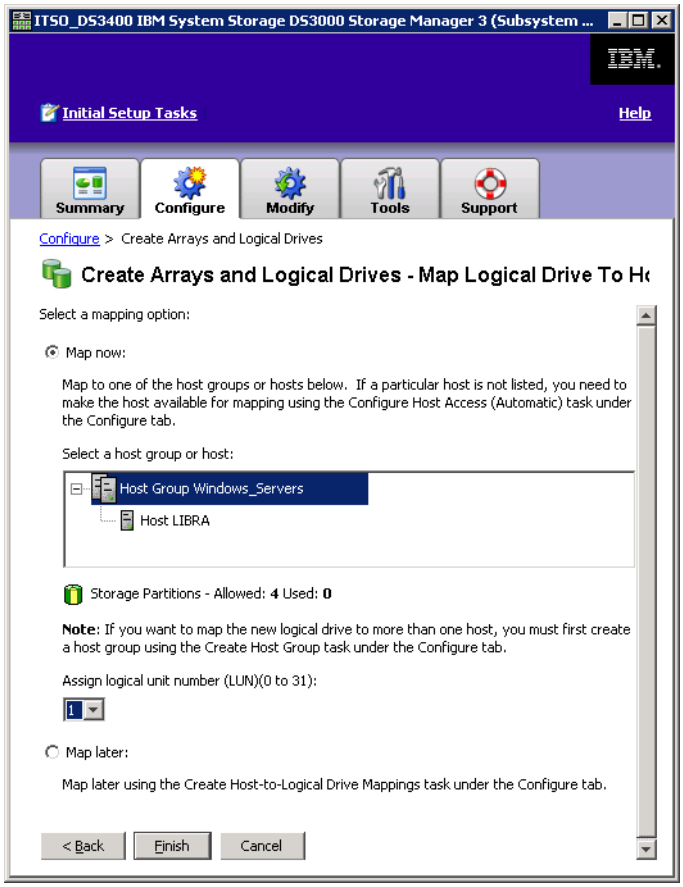


Figure 9-57 Logical drive creation - mapping view

Click **Next** to complete the task, as shown in Figure 9-58 on page 183



Figure 9-58 Logical drive creation - completed

Note: Unless you have unique requirements, we recommend using the tasks under Automatic (Simple) under 5: Configure Storage Subsystem, as shown in Figure 9-30 on page 156. This usually ensures the most optimal balance between capacity, performance, and redundancy. A manual configuration will typically not have the most optimal settings.

9.4 Advanced functions - FlashCopy

A FlashCopy logical drive is a point-in-time image of a logical drive. It is the logical equivalent of a complete physical copy, but it is created much more quickly than a physical copy. It also requires less disk space. On the other hand, it is not a real physical copy, because it does not copy all the data. Consequently, if the source logical drive is damaged, the FlashCopy logical drive cannot be used for recovery, but it works if the source data is corrupted.

9.4.1 FlashCopy: How it works

In the DS3000 Storage Manager, the logical drive that will be the source of the FlashCopy is called the *base logical drive.* Typically, you create a FlashCopy so that an application (for example, a backup application) can access the FlashCopy and read the data while the base logical drive remains online to users and applications. The FlashCopy drive, in this case, is temporary; it is usually disabled rather than deleted when the backup completes.

You can also create multiple FlashCopies of a base logical drive and use the copies in write mode to perform testing and analysis. Before you upgrade your database management system, for example, you can use FlashCopy logical drives to test different configurations. Then you can use the performance data provided by the Storage Management software to help decide how to configure the live database system.

Tip: For analysis, data mining, and testing without any degradation of the production logical drive performance, you can use FlashCopy together with VolumeCopy.

When you initiate a FlashCopy, the controller suspends I/O to the base logical drive for only a few seconds. Meanwhile, it creates a new logical drive called the *FlashCopy repository logical drive*, where it stores FlashCopy metadata and copy-on-write data (Figure 9-59). It will build a metadata database, which contains only pointers. When the controller finishes creating the FlashCopy repository logical drive, I/O write requests to the base logical drive can resume.

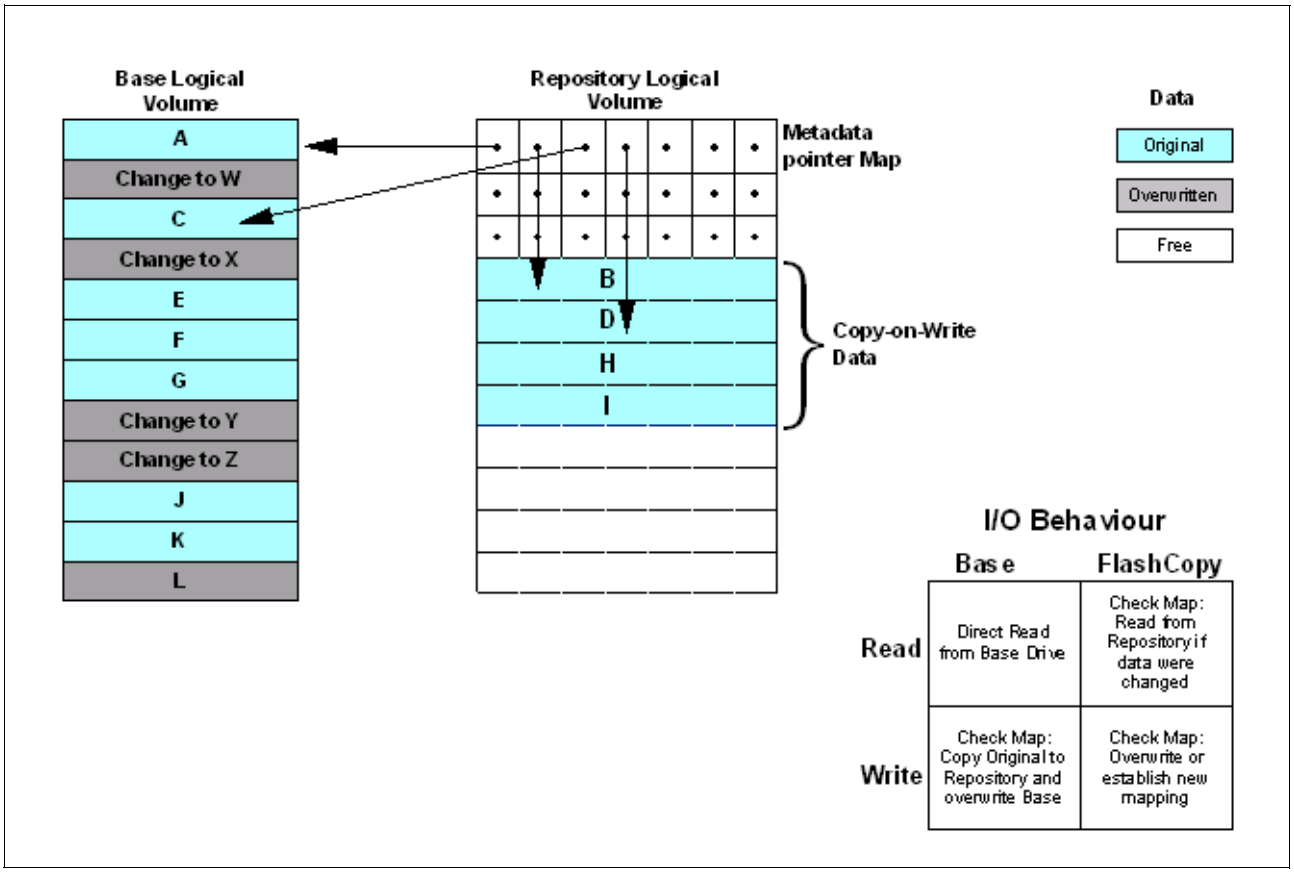


Figure 9-59 Flash Copy Read and write schema

If a data block on the base logical drive is modified, a copy-on-write occurs first, which copies the contents of blocks that are to be modified into the FlashCopy repository logical drive, for safekeeping. Subsequently, the corresponding pointer in the metadata database changes. Since the FlashCopy repository logical drive stores copies of the original data in those data blocks, further changes to those data blocks write directly to the base logical drive without another copy-on-write. And, since the only data blocks that are physically stored in the FlashCopy repository logical drive are those that have changed since the time of the FlashCopy, the FlashCopy technology uses less disk space than a full physical copy.

When you create a FlashCopy logical drive, you specify where to create the FlashCopy repository logical drive, its capacity, warning threshold, and other parameters. You can disable the FlashCopy when you are finished with it, for example, after a backup completes. The next time you re-create the FlashCopy, it reuses the existing FlashCopy repository logical drive. Deleting a FlashCopy logical drive will also delete the associated FlashCopy repository logical drive.

9.4.2 Estimating FlashCopy repository logical drive capacity

The *FlashCopy repository logical drive* is created to store FlashCopy metadata (data about the FlashCopy) and any copy-on-write data needed during the life of the FlashCopy logical drive. By default, the FlashCopy repository logical drive's capacity is 20% of the base logical drive's capacity. In general, this capacity should be sufficient. However, you can use the following information to help determine the appropriate capacity of the FlashCopy repository logical drive:

- A FlashCopy repository logical drive must be at least 8 MB.
- Set a larger percentage if a large percentage of data blocks will change on the base logical drive during the life of the FlashCopy logical drive.
- Set a larger percentage if you need to keep the FlashCopy logical drive for a longer period.

Calculating expected overhead

Normally, the default setting will be sufficient. You will see a warning when the FlashCopy repository logical drive reaches a certain percentage full. You can increase its capacity by selecting **Logical Drive** → **Increase Capacity** in Storage Manager.

Use the following formula to calculate the amount of management overhead required to store FlashCopy data on the FlashCopy repository logical drive. This formula should be used merely as a guide, and FlashCopy repository logical drive capacity should be re-estimated periodically.

Note: To use this formula, you need to convert from bytes to kilobytes, and then to megabytes.

The formula to calculate the amount of management overhead required is:

192 KB + (*X*/2000)

Here, *X* is the capacity of the base logical drive in bytes.

Example

For a 5 GB base logical drive, where 30% of the data blocks are expected to change on the base logical drive, the estimated FlashCopy repository logical drive capacity can be calculated as follows:

1. Convert the base logical drive's capacity to bytes:

   [5 x 1024(K) x 1024(M) x 1024(G)]

   When converted, 5 GB equals 5,368,709,120 bytes.
2. Divide the base logical drive's capacity (in bytes) by 2000.

   When divided, the result is 2,684,354.56 bytes.
3. Convert the result from step 2 (in bytes) to kilobytes (KB):

   [2,684,354.56 / 1024(K)]

   When converted, the result is 2621.44 KB.
4. Add 192 KB to the results from step 3:

   192 KB + 2621.44 KB = 2813.44 KB
5. Convert the result from step 4 to megabytes (MB).

   When converted, the amount of management overhead required is calculated to be 2.75 MB (or 0.002686 GB).

6. In this example, 30% of the data blocks on the base logical drive are expected to change. To accurately calculate the FlashCopy repository logical drive capacity, sufficient space needs to be allowed for the copy-on-write data as well as the management overhead (calculated in step 5).

   To calculate the copy-on-write space required, calculate the percentage of the base logical drive expected change:

   30 percent of 5 GB = 1.5 GB

   The final estimated FlashCopy repository logical drive capacity for this example is:

   1.5 GB + 0.002686 GB = 1.502686 GB

7. In the Create FlashCopy logical drive wizard, you will use this calculated value in the Repository Capacity window to specify the percentage (%) full box of base logical drive to set the estimated FlashCopy repository logical drive capacity (see Figure 9-69 on page 194).

   Note: The percentage (%) full box sets the FlashCopy repository logical drive capacity as a percentage of the base logical drive. Using the percentage (%) full box, increase or decrease the percentage until the FlashCopy Repository logical drive Capacity value matches the estimated capacity calculated in step 6. (Some rounding up might be required.)

FlashCopy failure policy

The FlashCopy repository logical drive's failure policy determines what happens if the FlashCopy repository logical drive becomes full (that is, all of its capacity has been used). The choices are:

- Fail FlashCopy logical drive

  When the FlashCopy repository logical drive becomes full, the associated FlashCopy logical drive becomes invalid. This is the default policy. The base logical drive remains online and user accessible and writable.

- Fail writes to base logical drive

  When the FlashCopy repository logical drive becomes full, writes to the base logical drive fail. The FlashCopy logical drive is preserved and remains valid because no new copy-on-write data is generated for the FlashCopy repository logical drive.

Attention: Do not ignore the FlashCopy repository logical drive `threshold exceeded` notification.

If a FlashCopy logical drive or FlashCopy repository logical drive is displayed as a missing logical drive, this means that the storage system could no longer access these drives. Missing logical drives, in most cases, are recoverable.

Note: Using a defragmentation utility to defragment a base logical drive with an associated FlashCopy repository logical drive will cause a copy-on-write of every data block in the base logical drive. This can cause the FlashCopy repository logical drive to become full before the defragmentation operation is completed.

To prevent this from happening, during a base logical drive defragmentation, ensure that the FlashCopy repository logical drives' capacity is set to 105% of the size of the base logical drive before starting the defragmentation utility. This is the minimum size needed to support a copy-on-write of every data block in the base logical drive.

9.4.3 Checking the status of the FlashCopy Premium Feature

Use the following procedure to check the Premium Features enabled on the DS3000.

From the DS3000 Subsystem Management window, in the Tools tab, select **View / Enable Premium Features**. The List Premium Features view opens and lists the enabled Premium Features (Figure 9-60). Verify that FlashCopy Logical Drives: is Enabled. If the FlashCopy feature is not enabled, enable it as described in 5.4.4, “(Optional) Viewing and enabling Premium Features” on page 74. Click **Configure** again to return to the configuration options.



Figure 9-60 View/Enable Premium Features

9.4.4 Creating a FlashCopy using the wizard

Now we will see how to set up a FlashCopy. You can use either the GUI or command-line interface (the **`create FlashCopyLogicaldrive`** command). The CLI can be scripted to support automatic operations.

This procedure is done by using the GUI. We summarize the CLI commands that achieve the same task in Appendix A, “FlashCopy example - Windows 2003 and 2008” on page 725.

To create a FlashCopy drive, you need to complete these main steps:

1. Stop/suspend host application I/Os.
2. Launch the Create FlashCopy Logical Drive Wizard.
3. Restart the application.

Step 1: Stop/suspend host application I/Os

This step is necessary to quiesce I/Os to the base drive.

1. Stop (or suspend) the host application that is accessing the base logical drive.
2. Unmount the base logical drive from the operating system if possible, so you can take a valid consistent copy.

> **Note:** It might not be convenient to stop database applications; however, typically you can put the database into backup or another quiesced mode while the FlashCopy is created. Refer to your application documentation for details.
>
> If you are using a database application, you should also back up application recovery files, such as rollback and redo logs, at the same time you create the FlashCopy. This is to ensure that the FlashCopy can be restored consistently if necessary.

3. In Windows, run a tool named *SMrepassist* (replication assistance) in the Storage Manager directory. This flushes all the memory-resident data for the file system indicated by [file system-identifier] to ensure that the storage subsystem hardware creates an accurate FlashCopy logical drive or logical drive copy, and that signatures and file system partitions are recognized.

   The command entered to run this utility is:

   ```
   SMrepassist -f [filesystem-identifier]
   ```

   For example:

   ```
   SMrepassist -f e:
   ```

   Where e: is the logical drive on the DS3000.

> **Important:** Windows 2003 operating system specific procedures can be found in Appendix A, “FlashCopy example - Windows 2003 and 2008” on page 725. Linux procedures are similar, but there are specific methods for mounting the FlashCopy disk depending on which variant of Linux is used. Refer to your Linux documentation for details.

Step 2: Launch the Create FlashCopy Logical Drive wizard

1. On the Configure tab, select the **Create Logical Drive Copies** option to start the wizard (Figure 9-14 on page 141). Review the information shown in the additional instructions window (Figure 9-61) and close it.

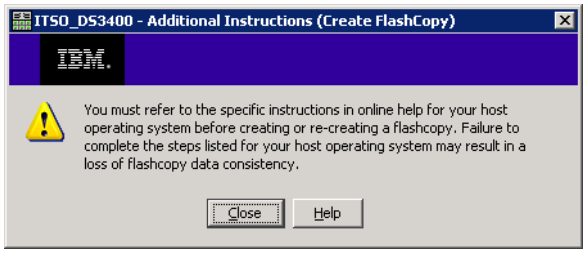


Figure 9-61 FlashCopy - Additional Instructions

2. The drive selection window appears (Figure 9-62 on page 189). Select the base logical drive from which to create the FlashCopy. Select a drive and click **Next**.

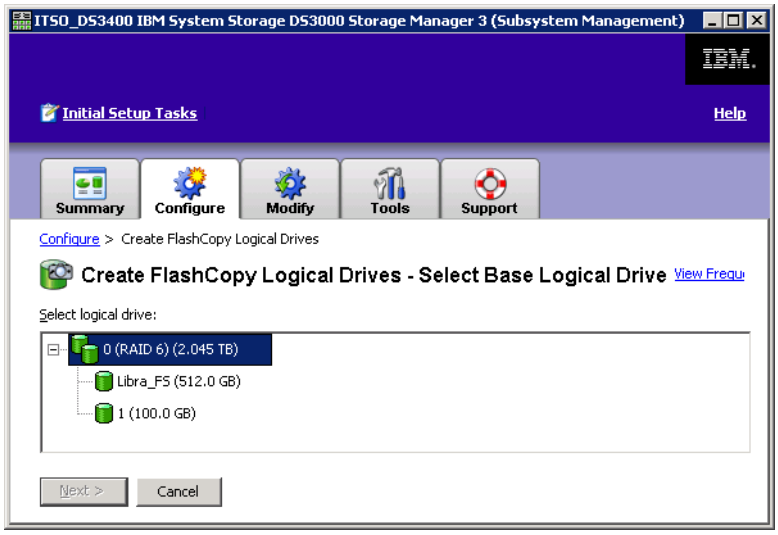


Figure 9-62 FlashCopy - Select Base Logical Drive

3. Select the path (Figure 9-63). You can select either a simple or advanced path:
   - Simple path (recommended): Create a FlashCopy logical drive using the recommended default values.
   - Advanced path: Customize logical drive attributes, policies, and notifications.

   If no free capacity exists or the available free capacity is unusable, a warning message is issued.

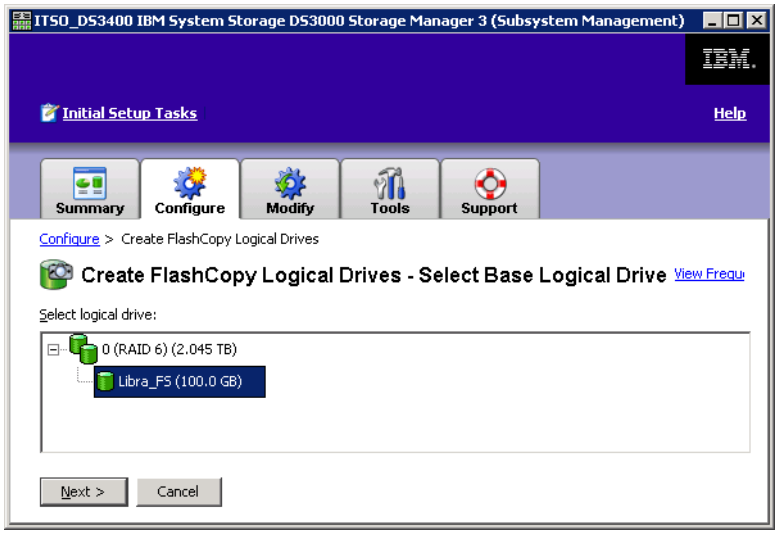


Figure 9-63 FlashCopy - select path

Create logical drives — simple path

First, we describe the simple path using the default values. The advanced path is covered in “Create logical drives — advanced” on page 192.

1. Select **Simple** and click **Next**.

2. The Specify Logical Drive Names window appears. Choose names for the FlashCopy logical drive and its associated repository logical drive (Figure 9-64).

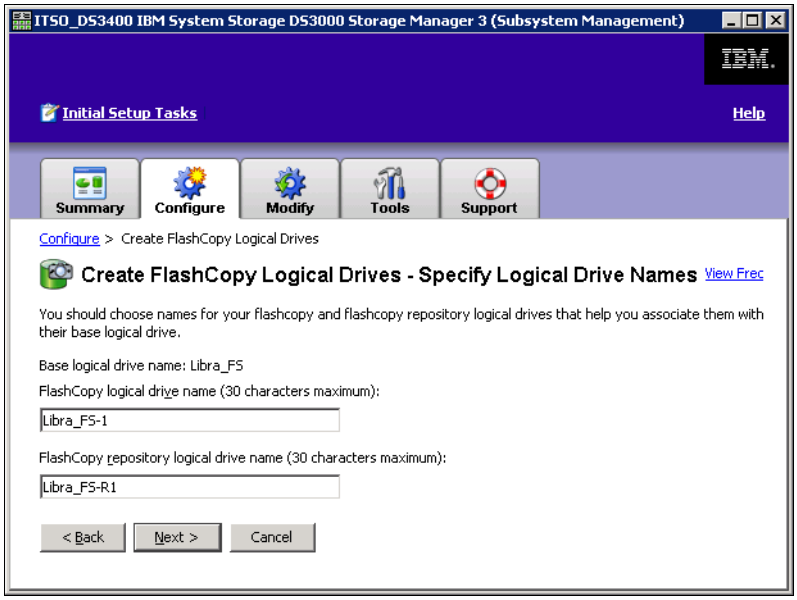


Figure 9-64 FlashCopy - Specify Drive Names

The default naming convention for the first FlashCopy is to use the base logical drive name and add a suffix of "-1" for the logical drive and "-R1" for the repository drive. The second FlashCopy uses 2 instead of 1. This is repeated up to the fourth logical drive.

For example, if you are creating the first FlashCopy logical drive for a base logical drive called Data 1, then the default FlashCopy logical drive name is Data 1-1, and the associated FlashCopy repository logical drive default name is Data 1-R1. The default name of the next FlashCopy logical drive you create based on Data 1 is Data 1-2, with the corresponding FlashCopy repository logical drive named Data 1-R2 by default.

Enter the names you want, or accept the defaults, and click **Next**.

Tips: Consider the following recommendations:

- Regardless of whether you use the software-supplied sequence number that (by default) populates the "FlashCopy logical drive name or FlashCopy repository logical drive name" field, the next default name for a FlashCopy or FlashCopy repository logical drive still uses the sequence number determined by the software. For example, if you override the name of the first FlashCopy of base logical drive Data 1 *DataVolMay28*, the default name for the next FlashCopy of Accounting is still *Data 1-2*.
- The next available sequence number is based on the number of existing FlashCopies of a base logical drive. If you delete a FlashCopy logical drive, its sequence number becomes available again.
- You must choose unique names for the FlashCopy and FlashCopy repository logical drives. Otherwise, an error message is displayed.

3. The Map FlashCopy Logical Drive to Hosts view displays (Figure 9-65). Here you assign logical drive-to-LUN mappings between the FlashCopy logical drive and the host that will access the FlashCopy logical drive, using the Mappings View of the Subsystem Management window. In some cases, depending on the host operating system and if any logical drive manager software in use, mapping the same host to both a base logical drive and its associated FlashCopy logical drive might result in conflicts.

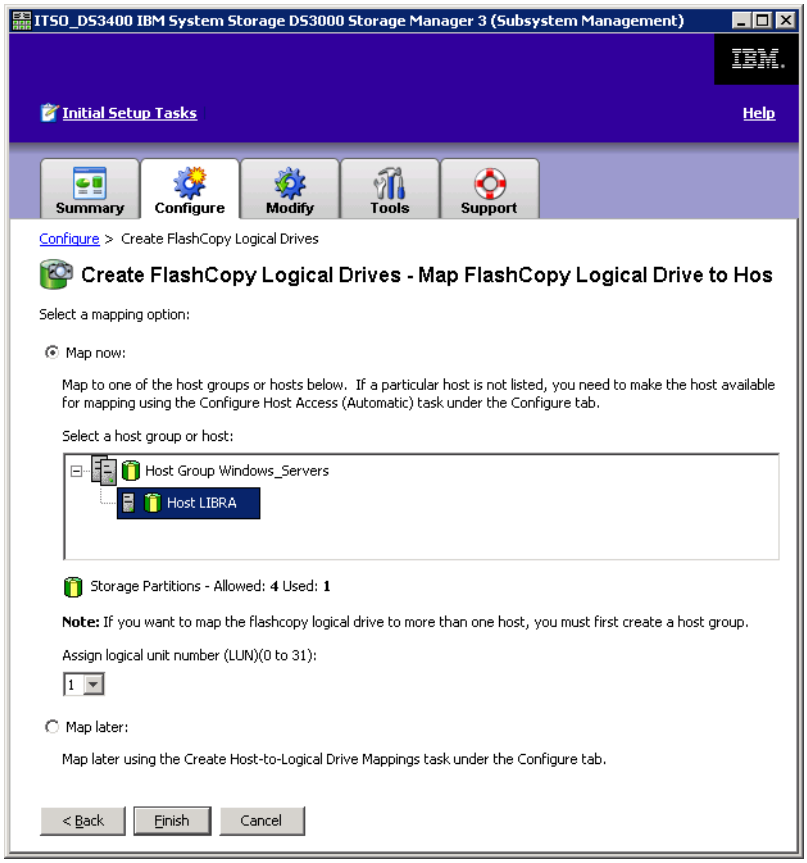


Figure 9-65 FlashCopy - Map logical drive to host

To map the FlashCopy logical drive to a host, select the host and click **Finish**. If you want to do all your mappings later, select **Map later**.

4. Click **Finish** to save the configuration. You will be prompted to create another FlashCopy Logical drive (Figure 9-66). If you plan to create multiple FlashCopies, we advise you to create and map all of the FlashCopy logical drives before registering them with hosts, so you can register all logical drives together.

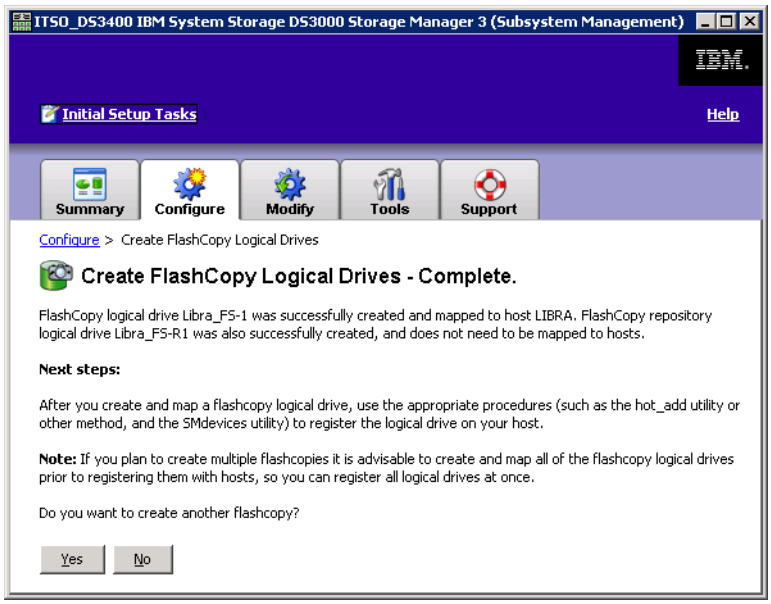


Figure 9-66 FlashCopy - Complete

In the example just shown in Figure 9-65 on page 191, the FlashCopy logical disk is made available to the host Libra. This enables the host Libra to access and mount the FlashCopied disk as part of its own file system.

It is possible to map the FlashCopy logical disk to the same server that owns the base logical disk. However, note that the two logical disks, immediately after creating the FlashCopy, appear exactly the same (a block-by-block copy). Many operating systems do not tolerate seeing an exact duplicate volume. In Windows, for example, when a FlashCopy source logical drive is a dynamic disk, the FlashCopy source and target logical drive will get the same DiskID in Logical Disk Manager (LDM), which will cause problems. When the FlashCopy source logical drive is a basic disk, FlashCopy source and target logical drives will get the same disk signature. We recommend that you map the FlashCopy destination logical drive to a different host.

For more information about Windows disk best practices, see the Microsoft Web site at:

`http://support.microsoft.com/kb/816307/en-us`

Important: If you use a FlashCopy drive regularly (for example, for backup purposes), use the Disable FlashCopy and Re-create FlashCopy options to reuse the FlashCopy repository logical drive, as described in “Create logical drives — advanced” on page 192. Using these options preserves the existing mappings to the FlashCopy logical drive.

Continue to “Step 3: Restart host application” on page 195.

Create logical drives — advanced

For more control over the FlashCopy logical drive, select **Advanced** from the Select Path view (Figure 9-67 on page 193) and click **Next**.

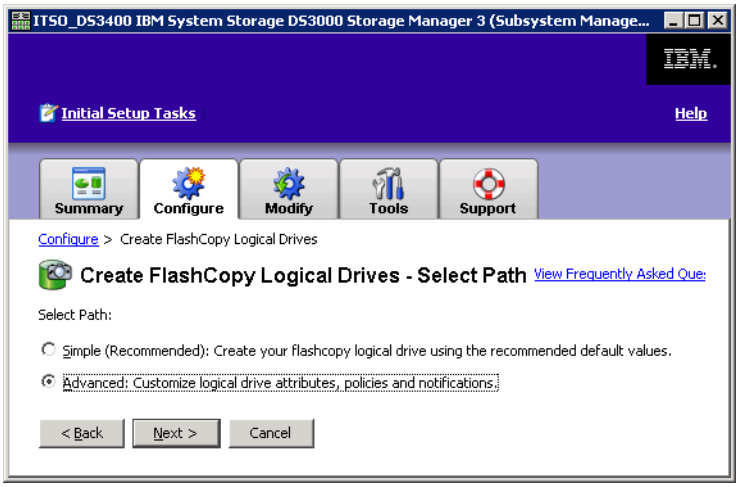


Figure 9-67 FlashCopy - advanced path

Do the following steps:

1. The Specify Logical Drive Names window appears, just as it appeared in Figure 9-64 on page 190, with the same naming conventions. If required, edit the names and click **Next**.
2. The Select Capacity Type view appears (Figure 9-68). You can configure your FlashCopy repository logical drive to reside in any array. However, to make it easier to move disk drives, put the FlashCopy repository logical drive in the same array as the base logical drive. When you do this, if disk drives associated with the array are moved to another storage subsystem, all logical drives associated with the FlashCopy logical drive also reside in that group. Also, no more disk drives need to be moved.

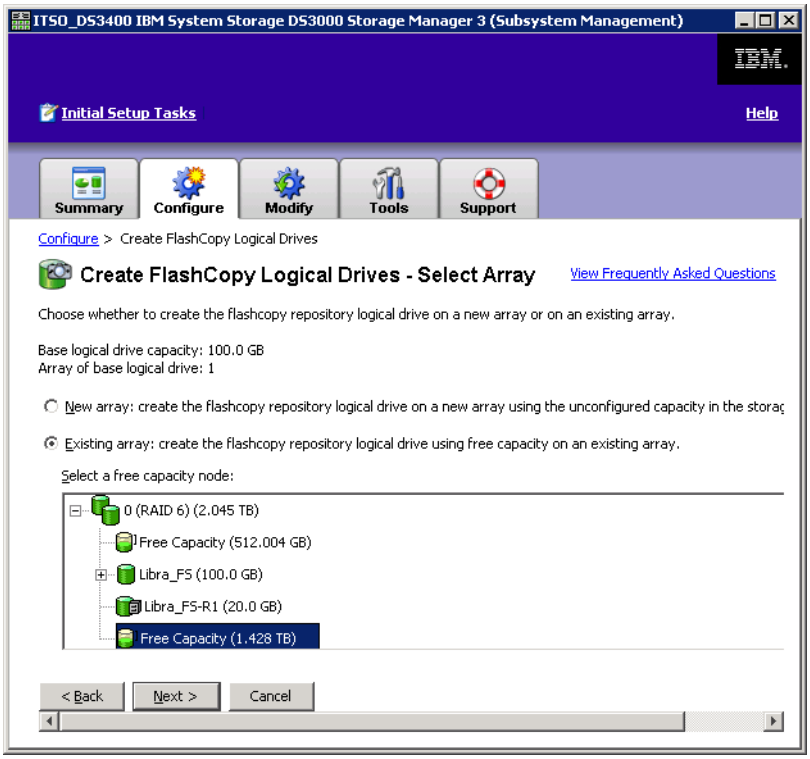


Figure 9-68 FlashCopy - select capacity type

If you decide to create the FlashCopy Logical drive in a new array, select **Unconfigured capacity** option and click **Next** to create an array. The process to create a new array is discussed in “Unconfigured capacity” on page 170. Alternatively, as shown in Figure 9-68 on page 193, select **Free Capacity** and click **Next**.

3. Both ways will bring you to the Specify Repository Logical Drive Capacity view (Figure 9-69). Here you specify how much of the array’s free capacity you want to allocate to the FlashCopy repository logical drive, as a percentage of the base logical drive’s capacity. Calculating this capacity is described in 9.4.2, “Estimating FlashCopy repository logical drive capacity” on page 185. Enter a value, depending on your requirements, and click **Next**.

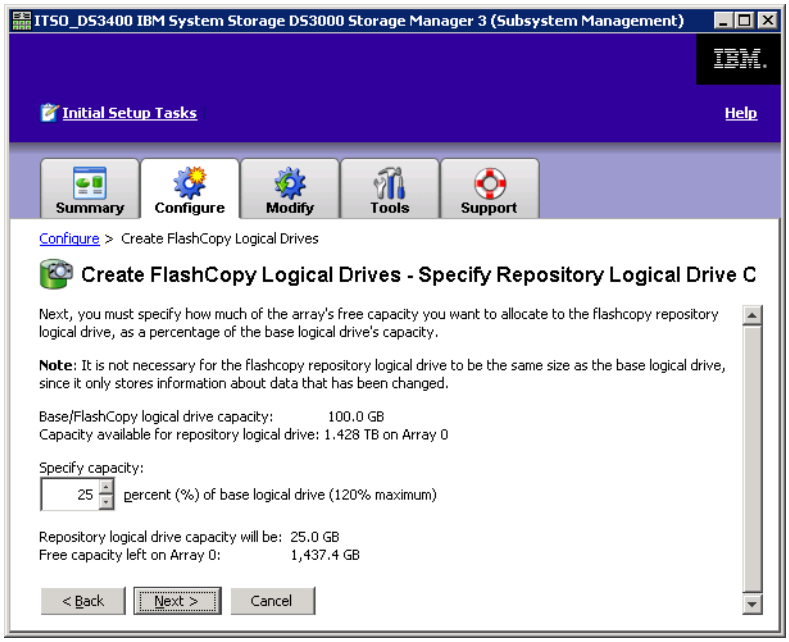


Figure 9-69 FlashCopy -Specify Repository

4. In the Specify Repository Policies window (Figure 9-70 on page 195), select which repository full policies you want to use. We described these options in “FlashCopy failure policy” on page 186. Enter the notification threshold and your choice for what to do if the repository logical drive fills and click **Next**.

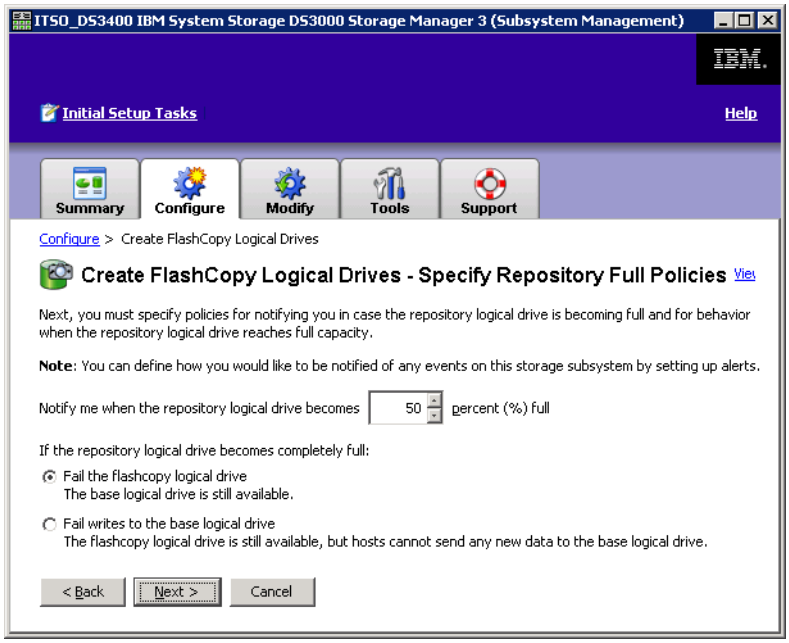


Figure 9-70 FlashCopy - Specify Repository Full Policies

5. You are prompted for the host mappings, as shown in Figure 9-65 on page 191. Follow the instructions given in step 3 on page 191 to complete your FlashCopy logical drive configuration. The steps are identical to the automatic path.

You have completed the basic setup of your FlashCopy logical drives. More FlashCopy functions are covered in 11.2.7, “Modify FlashCopy Logical Drives” on page 243.

Step 3: Restart host application

The FlashCopy is now created. You can restart the host application that was active on the base logical disk (or take your database out of backup/quiesce mode).

9.4.5 Using the FlashCopy logical drive

You will use *specific operating system and host utilities* to mount and use the mapped FlashCopy drive. The basic procedure is:

1. Run the host-based hot_add (or an operating-system specific utility) utility to register the FlashCopy logical drive with the host operating system.
2. Run the host-based SMdevices utility to associate the mapping between the physical device name and the logical drive name.
3. Mount the logical drive to the host and use the FlashCopied volume as required, for example, to take an application backup.

For information about specific procedures for Windows 2003, see Appendix A, “FlashCopy example - Windows 2003 and 2008” on page 725.

9.5 Advanced functions - VolumeCopy

The VolumeCopy Premium Feature copies data from one logical drive (source) to another logical drive (target) in a single storage subsystem. The target logical drive is an exact copy or *clone* of the source logical drive. VolumeCopy can be used to copy data from arrays that use smaller capacity drives to arrays that use larger capacity drives, to back up data, or to restore FlashCopy logical drive data to the base logical drive. The VolumeCopy Premium Feature includes a Create Copy Wizard, to assist in creating a logical drive copy, and a Copy Manager, to monitor logical drive copies after they have been created. See 11.2.8, “Manage Logical Drive Copies” on page 250 for information about Copy Manager.

9.5.1 Introduction to VolumeCopy

The VolumeCopy Premium Feature must be enabled by purchasing a feature key. FlashCopy must be installed as well. VolumeCopy is only available as a bundle that includes a FlashCopy license.

VolumeCopy is a full point-in-time replication (Figure 9-71). It allows for analysis, mining, and testing of data without any degradation of the production logical drive performance. It also improve backup and restore operations, making them faster and eliminating I/O contention on the primary (source) logical drive.

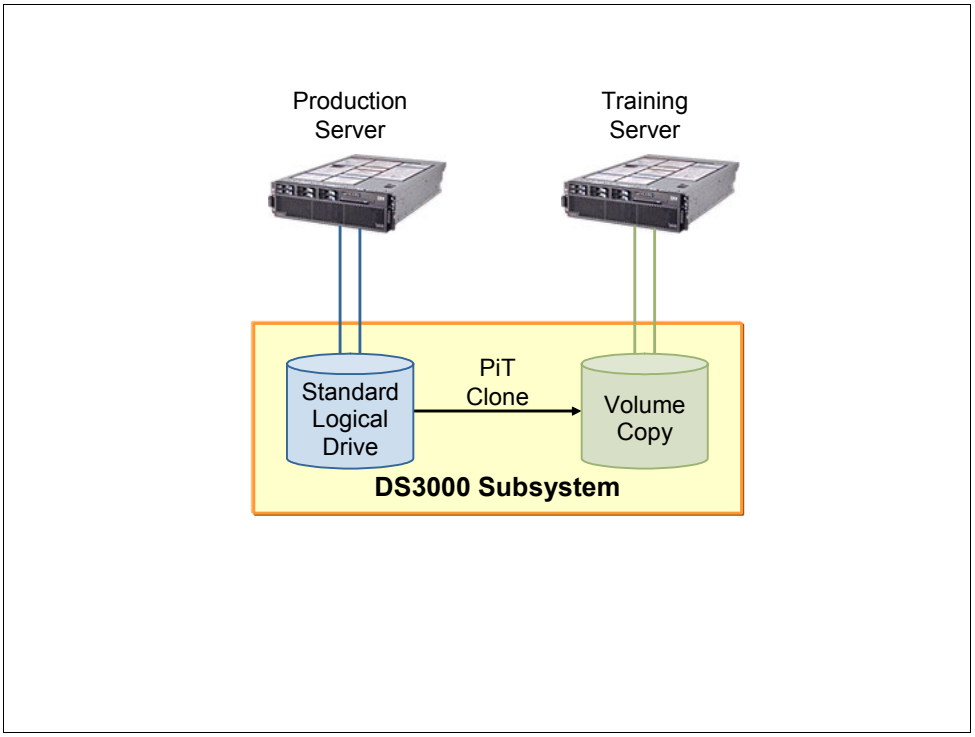


Figure 9-71 VolumeCopy

To create a VolumeCopy request, specify two compatible volumes (logical drives). One logical drive is designated as the source, and the other is the target.

Copying data is a background operation managed by the controller firmware, which reads the source logical drive and writes the data to the target logical drive. If the storage system controller experiences a reset, the copy request is restored and the copy process resumes from the last known progress boundary.

After submitting a copy request, the source logical drive is only available for read I/O activity while a logical drive copy has a status of In Progress, Pending, or Failed. Write requests are allowed after the logical drive copy is completed. Read and write requests to the target logical drive will not take place while the logical drive copy has a status of In Progress, Pending, or Failed.

These restrictions are necessary to ensure the integrity of the point-in-time copy. If the logical drive being copied is large, this can result in an extended period of time when a production application cannot make updates or changes to the data.

Important: Because all write requests to the source logical drive are rejected when the VolumeCopy is taking place, it is essential to minimize the time it takes to complete the copy operation. This is why VolumeCopy must always be used in conjunction with FlashCopy. In other words, first make a FlashCopy of the source logical drive, then perform a VolumeCopy of the FlashCopy (see Figure 9-72).

As shown in Figure 9-72, FlashCopy, which allows a point-in-time copy to be made while maintaining read/write access, enables a complete copy to be created without interrupting the I/O activity of the production logical drive.

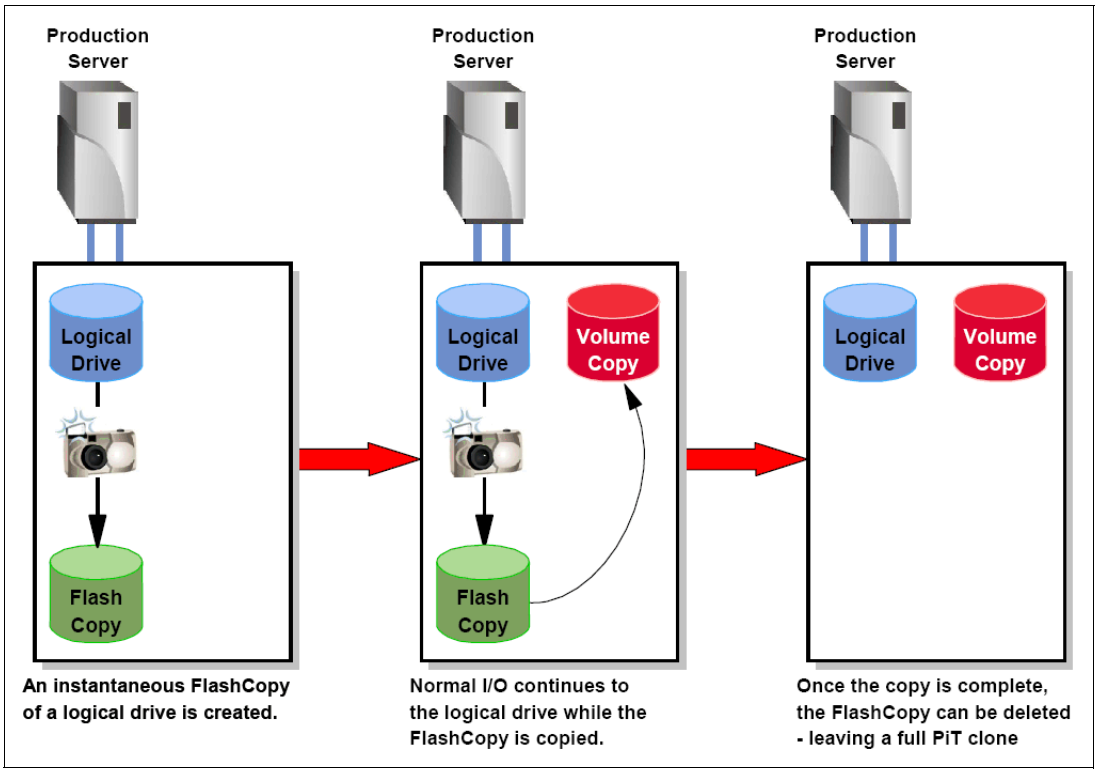


Figure 9-72 VolumeCopy integrated with FlashCopy

9.5.2 Copying data for greater access

In the following sections, we describe some of the advantages of VolumeCopy.

Migration

As storage requirements for a logical drive change, the VolumeCopy function can be used to copy data to a logical drive in an array that utilizes larger capacity disk drives within the same storage system. This provides an opportunity to move data to larger drives (for example, 73 GB to 300 GB).

Backing up data

The VolumeCopy function can create a backup of a logical drive by copying data from one logical drive to another logical drive in the same storage system. The target logical drive can be used as a backup of the source logical drive, for system testing, or to do a backup to another device, such as a tape drive.

Restoring FlashCopy logical drive data to the base logical drive

If you need to restore data to the base logical drive from its associated FlashCopy logical drive, the VolumeCopy function can be used to copy the data from the FlashCopy logical drive to the base logical drive. You can create a VolumeCopy of the data on the FlashCopy logical drive, then copy the data to the base logical drive.

9.5.3 Creating and managing VolumeCopy copies

A VolumeCopy can be made and managed using the IBM DS3000 Storage Manager or with command-line interface (CLI) commands and scripts.

9.5.4 Understanding VolumeCopy

A VolumeCopy operation will invalidate (fail) all FlashCopy logical drives associated with the target logical drive, if any exist. If you select a base logical drive of a FlashCopy logical drive, you must disable all FlashCopy logical drives associated with the base logical drive before you can select it as a target logical drive; otherwise, the base logical drive cannot be used as a target logical drive.

Important: A VolumeCopy overwrites data on the target logical drive and automatically makes the target logical drive *read-only* to hosts.

If there are eight logical drive copies with a status of In Progress, any subsequent VolumeCopy will have a status of Pending until one of the eight logical drive copies completes.

VolumeCopy and modification operations

If a modification operation is running on a source logical drive or target logical drive, and the VolumeCopy has a status of In Progress, Pending, or Failed, then the VolumeCopy will not take place. If a modification operation is running on a source logical drive or target logical drive after a VolumeCopy has been created, then the modification operation must complete before the VolumeCopy can start. If a VolumeCopy has a status of In Progress, then any modification operation will not take place.

Modification operations include:

- Defragmenting an array
- Copying back to a drive that is part of an array
- Initialization of a logical drive
- Dynamic Segment Sizing (DSS) change for a logical drive
- Dynamic Reconstruction Rate (DRR) of a drive that is part of an array
- Dynamic RAID Level Migration (DRM) change for an array
- Dynamic Capacity Expansion (DCE), to increase an array's capacity using unconfigured capacity (in the form of unused drives)
- Dynamic Logical Drive Expansion (DVE), to increase a logical drive capacity using free capacity available on the array of the standard or FlashCopy repository logical drive

Failed VolumeCopy

A VolumeCopy can fail because of a read error from the source logical drive, a write error to the target logical drive, or because of a failure on the storage system that affects the source logical drive or target logical drive. A critical event is logged in the event log when the VolumeCopy fails, and a Needs-Attention icon is displayed in the Enterprise Management Window. While a VolumeCopy has this status, the host has read-only access to the source logical drive, and read and write requests to the target logical drive will not take place until the failure is corrected by using the Recovery Guru.

Preferred controller ownership

During a VolumeCopy, the same controller must own both the source logical drive and target logical drive. If both logical drives do not have the same preferred controller when the VolumeCopy starts, the ownership of the target logical drive is automatically transferred to the preferred controller of the source logical drive. When the VolumeCopy is completed or is stopped, ownership of the target logical drive is restored to its preferred controller. If ownership of the source logical drive is changed during the VolumeCopy, ownership of the target logical drive is also changed.

Failed controller

Controller ownership must be manually changed to the alternate controller to allow the VolumeCopy to complete under the following conditions: If a VolumeCopy has a status of In Progress, and its preferred controller fails, the ownership transfer does not occur automatically in the failover.

Restrictions

The following restrictions apply to the source logical drive, target logical drive, and storage system:

- The source logical drive is available for read I/O activity only while a VolumeCopy has a status of In Progress or Pending. Write requests are allowed after the VolumeCopy is completed.

  Tip: In practice, VolumeCopy must only be used with FlashCopy, in order to shorten the source logical driver read-only period of time.

- A logical drive can be used as a target logical drive in only one VolumeCopy at a time.
- The maximum allowable number of logical drive copies per storage system is dependent upon the number of target logical drives available on your storage system.

- A storage system can have up to eight VolumeCopies running at any given time.
- The target logical drive capacity must be equal to or greater than the source logical drive capacity.
- A source logical drive can be a standard logical drive, a FlashCopy logical drive, or a FlashCopy base logical drive.
- A target logical drive can be a standard logical drive or a base logical drive of a disabled or failed FlashCopy logical drive.

Important: If you choose a base logical drive of a FlashCopy logical drive as your target logical drive, you must disable all FlashCopy logical drives associated with the base logical drive before you can select it as a target logical drive; otherwise, the base logical drive cannot be used as a target logical drive.

Logical drives that have the following statuses cannot be used as a source logical drive or target logical drive:

- A logical drive that is reserved by the host cannot be selected as a source or target logical drive.
- A logical drive that is in a modification operation.
- A logical drive that is the source logical drive or target logical drive in another VolumeCopy with a status of Failed, In Progress, or Pending.
- A logical drive with a status of Failed.
- A logical drive with a status of Degraded.

9.5.5 VolumeCopy and performance considerations

During a VolumeCopy, data is read from the source logical drive and written to the target logical drive in the same storage system. Because the VolumeCopy diverts controller processes resources from I/O activity, it can affect I/O activity on the storage system. The copy priority defines how much processing time is allocated for a VolumeCopy versus I/O activity.

VolumeCopy priority rates

Several factors contribute to system performance, including I/O activity, logical drive RAID level, logical drive configuration (number of drives in the array or cache parameters), and logical drive type (FlashCopy logical drives can take more time to copy than standard logical drives).

You can select the copy priority while you are creating a new VolumeCopy, or change it later using the Copy Manager (described in 11.2.8, “Manage Logical Drive Copies” on page 250). The following copy priority rates are available: lowest, low, medium, high, and highest.

Note: The lowest priority rate supports I/O activity, but the VolumeCopy will take longer. The highest priority rate supports the VolumeCopy, but I/O activity might be affected. Remember, however, that VolumeCopy should always be used in conjunction with FlashCopy.

9.5.6 VolumeCopy: Step-by-step

In this section, we show how to set up VolumeCopy.

Check the status of the VolumeCopy Premium Feature

You have to check that the VolumeCopy Premium Feature is licensed.

From the DS3000 Subsystem Management window, click the **Tools** tab and select **View/Enable Premium Features**. The List Premium Features view opens and lists the enabled Premium Features (Figure 9-73). Verify that FlashCopy Logical Drives: is Enabled. If the FlashCopy feature is not enabled, see 5.4.4, "(Optional) Viewing and enabling Premium Features" on page 74 for instructions about how to enable it. Click **Configure** again to return to the configuration options.

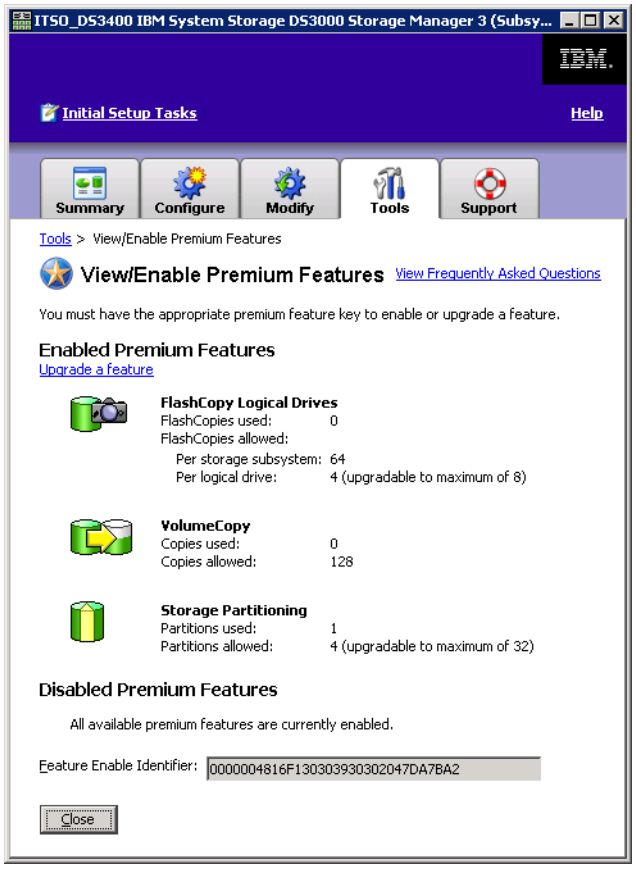


Figure 9-73 VolumeCopy - status of the VolumeCopy Premium Feature

Note: If you receive a `Premium Features - Out of Compliance` error message during a management session, use the Recovery Guru (see 14.6, "Recover from Failure" on page 347) to resolve the problem.

9.5.7 Create a VolumeCopy pair

This section describes in detail how to create VolumeCopies using the Create Copy Wizard from the Subsystem Management console. The wizard will guide you through the VolumeCopy creation process.

Create Copy Wizard

The Create Copy Wizard helps you select a source logical drive from a list of available logical drives, select a target logical drive from a list of available logical drives, and set the copy priority for the VolumeCopy. When you have completed the Wizard dialogs, the VolumeCopy starts and data is read from the source logical drive and written to the target logical drive.

Important: Remember that we recommend first making a FlashCopy of the logical drive, and then using the FlashCopy as the source of the VolumeCopy.

Launch the Create Copy Wizard:

1. Stop all I/O activity to the source logical drive and target logical drive.
2. Unmount any file systems on the source logical drive and target logical drive.
3. In the DS3000 Storage Manager Configure tab, select **Create Logical Drive Copies** (Figure 9-14 on page 141).
4. Figure 9-74 appears, where you start the wizard. Here you choose a source logical drive.

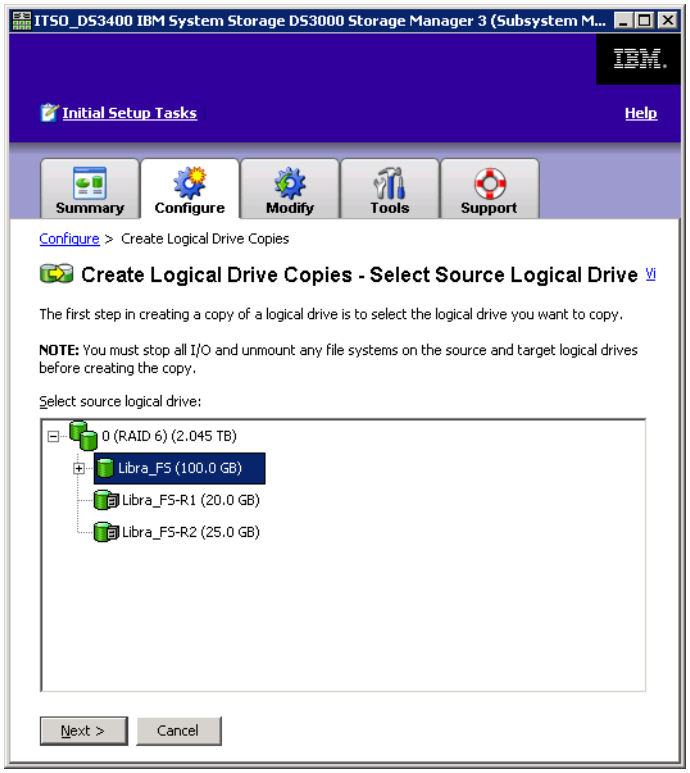


Figure 9-74 VolumeCopy - Select Source Logical Drive

The source logical drive can be used for read requests only while a VolumeCopy has a status of In Progress, Pending, or Failed. Writes are allowed only after the VolumeCopy is completed. Highlight the source logical drive you want to copy and click the **Next** button.

Note: The following logical drives are not valid source logical drives and you will not be able to click **Next** if you select them:

- FlashCopy repository logical drive
- Failed logical drive
- Missing logical drive
- A logical drive currently in a modification operation
- A logical drive holding a legacy or persistent reservation
- A logical drive that is a source logical drive or target logical drive in another VolumeCopy with a status of In Progress, Failed, or Pending

5. Select a target logical drive. You have two choices (Figure 9-75):
   - Use an existing logical drive with the same or a larger capacity than the target drive.
   - Create a new logical drive in the array (if enough free capacity exists)

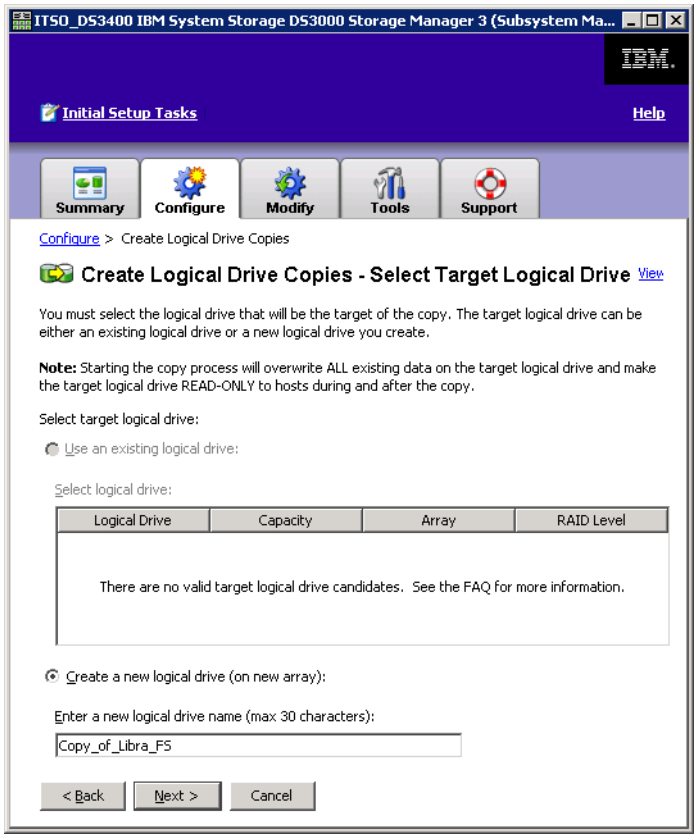


Figure 9-75 VolumeCopy - Select Target Logical Drive

Select one of the two possibilities. If you create a new logical drive, you also have to specify a name for it. In this case, it is selected automatically.

6. Set the priority of the VolumeCopy (Figure 9-76) using the sliding scale. The priority represents a balance between I/O activity and the VolumeCopy operation. The lowest priority rate supports I/O activity, but the VolumeCopy will take longer. The highest priority rate supports the VolumeCopy, but I/O activity might be affected.

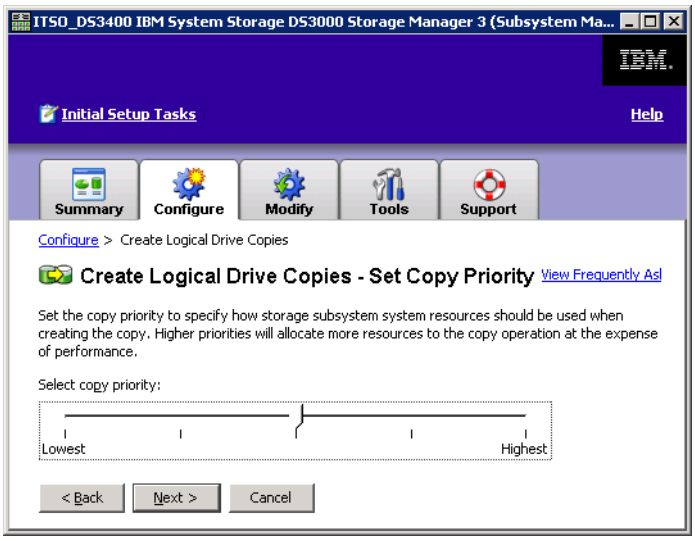


Figure 9-76 VolumeCopy - Set Copy Priority

Position the scale and click **Next**.

7. Figure 9-77 on page 205 is the confirmation window. It provides a summary of the current VolumeCopy configuration. Review the source, target, and priority. Remember that a VolumeCopy will overwrite data on the target logical drive and automatically makes the target logical drive read-only to hosts.

   If your configuration is correct and you want to continue, type yes in the box and click **Finish** (Figure 9-77 on page 205).



Figure 9-77 VolumeCopy - confirm configuration

8. The VolumeCopy process starts. You can either create another VolumeCopy or return to the Configuration tab (Figure 9-78).

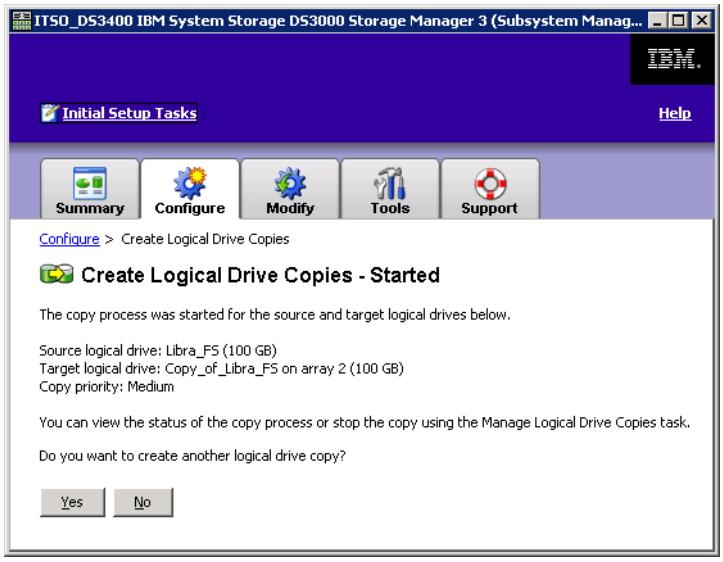


Figure 9-78 VolumeCopy - Started

After the VolumeCopy has started, you can use Copy Manager (described in 11.2.8, “Manage Logical Drive Copies” on page 250) to manage your VolumeCopies. There you must also disable the Read-Only attribute for the target logical drive when the VolumeCopy has completed.

Tip: After you have selected or created a target logical drive, give it a unique name so that it is easily recognizable in the Array View. For example, if the source logical drive name is Accounting, you could call the target logical drive Accounting_Copy. That way you can quickly identify the source logical drives and target logical drives available on the storage system. Renaming logical drives is described in 11.2.1, “Rename Arrays and Logical Drives” on page 231.

10

Administration - Summary

In this chapter, we give you an overview of the Summary tab of the DS3000 Storage Manager.

We explain:

- How to start the initial setup wizard again
- How to monitor the status of your storage subsystem
- How to get an overview of your installed hardware components and the current configuration
- How to get, view, and save a profile

Chapter 8, “Administration - Enterprise” on page 111 describes how to start managing a system from the Enterprise Window.

If you select to manage a specific storage subsystem (by double-clicking it), first the IBM DS3000 Storage Manager Subsystem Management Window will open, along with the Initial Setup Task pop-up window (Figure 10-1 on page 208).

With this task, you can perform a quick basic installation, as described in 5.4, “Initial Setup Tasks wizard” on page 71.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

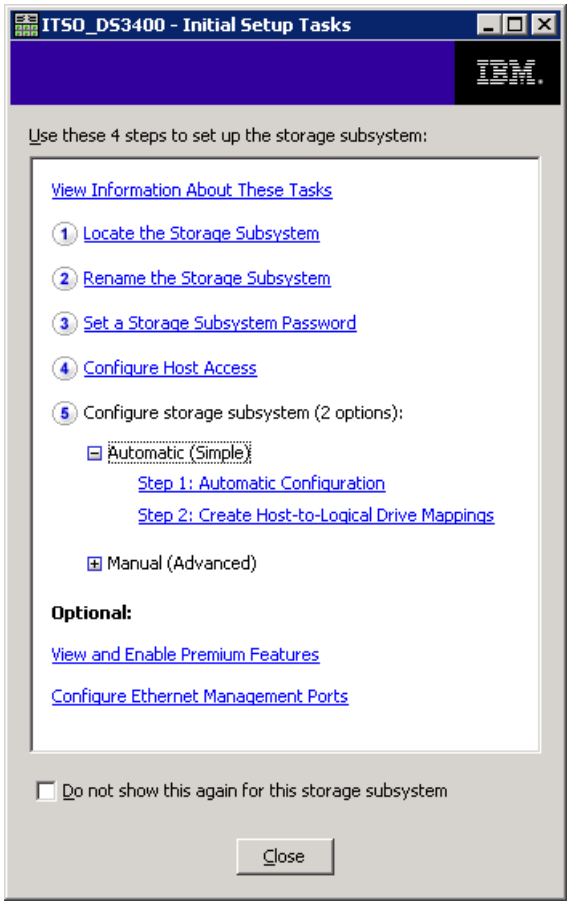


Figure 10-1 Initial Setup Tasks

In this chapter, we describe the configuration and management tasks in the graphical user interface (GUI).

Close the Initial Setup Tasks window and you will see the main Summary tab. Figure 10-2 on page 209 shows an unconfigured system that is running in an optimal status. The top of the tab shows the name of the DS3000 system, with the link next to it, to re-start the Initial Setup Tasks if you need them again. The name of the Subsystem can be edited in the Enterprise Management Window, as described in “Rename” on page 122.

The Summary tab GUI is divided into six groups, which we describe in the following sections:

- 10.1, “Status box” on page 210
- 10.2, “Hardware components” on page 212
- 10.3, “Capacity” on page 215
- 10.4, “Hosts and mappings” on page 216
- 10.5, “Arrays and logical drives” on page 218
- 10.6, “Information Center” on page 219

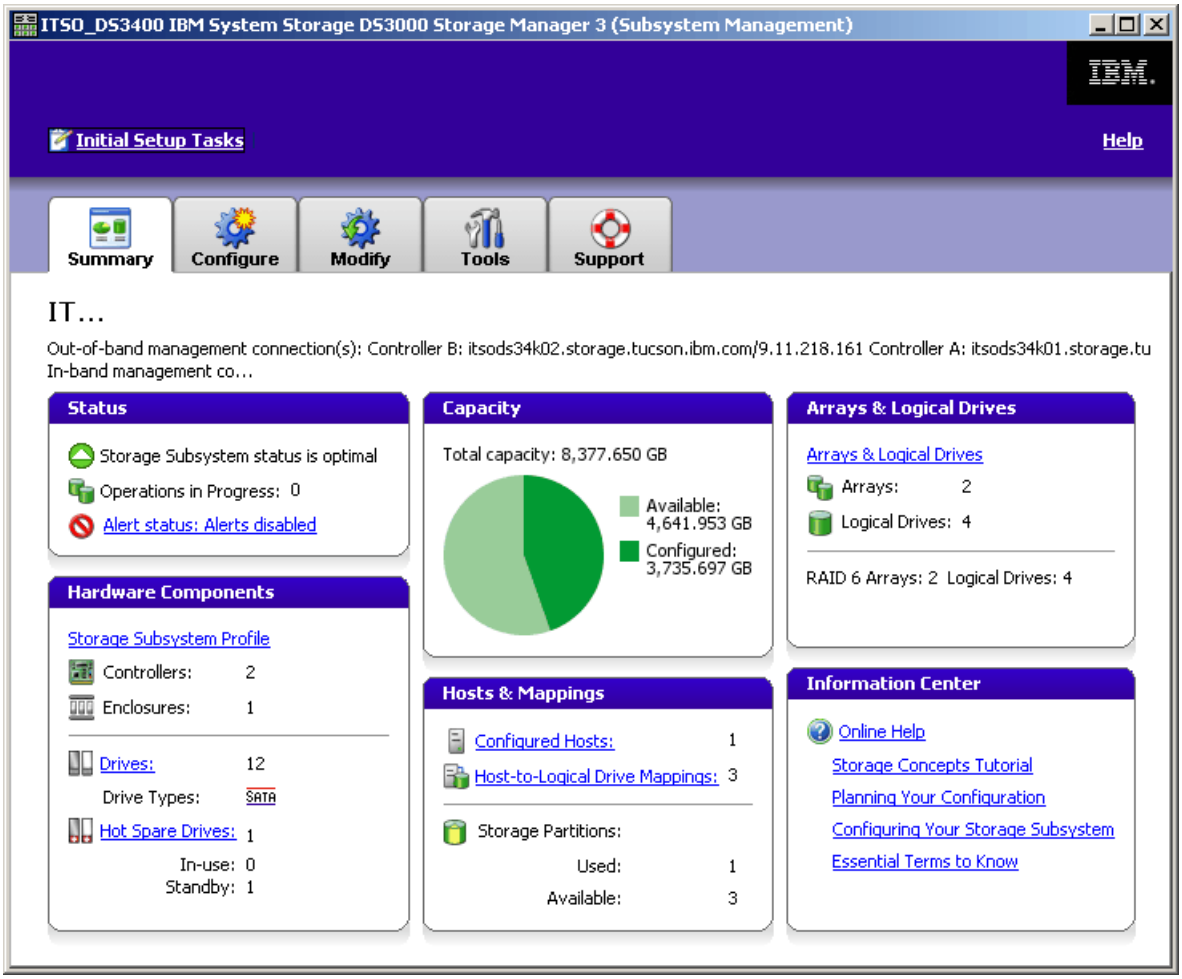


Figure 10-2 IBM System Storage DS3000 Storage Manager Summary tab

10.1 Status box

The first box is the Status Box, which displays the status of your storage subsystem (Figure 10-3).

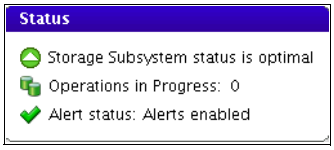


Figure 10-3 Status box

10.1.1 Storage subsystem status

The first line shows the current subsystem status. If the status is green (Figure 10-3), the subsystem is running in an optimal status. If there is any problem with the subsystem, the status will change to red (Figure 10-4). Click the link **Storage Subsystem Needs Attention to** open the Recovery Guru. How to troubleshoot and recover from problems is covered in Chapter 14, “Administration - Support” on page 299.

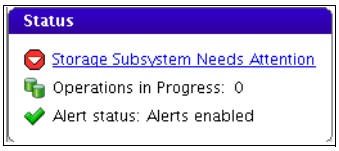


Figure 10-4 Status box - Storage Subsystem Needs Attention

10.1.2 Operations in Progress

The second line shows operations that are currently in progress (Figure 10-5). If there is at least one operation in progress, you can click the link to display a window with more detailed information (Figure 10-6 on page 211).

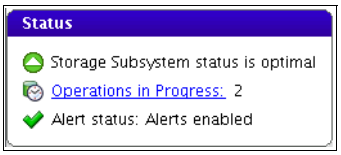


Figure 10-5 Status Box - operations in progress / alerts enabled

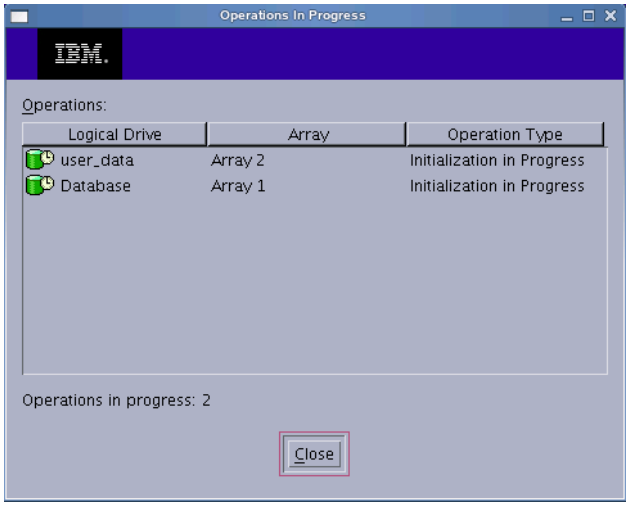


Figure 10-6 Detailed progress information

10.1.3 Alert status

The third line of the box shows if there is an alert configuration enabled. If alerts are configured, this line looks like Figure 10-5 on page 210. If there are currently no alerts configured, it looks like Figure 10-7. “Configure Alerts” on page 118 describes how to enable alerts.

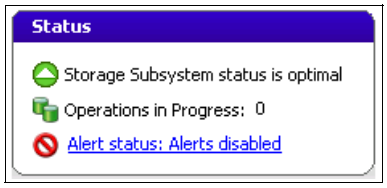


Figure 10-7 Status Box - alerts disabled

10.1.4 Connection lost

If the management connection from the management station to the storage server is lost, the icons and text change, and an additional line appears in the box to notify you of this change (Figure 10-8). If only the management connection is lost, the attached hosts can still access data on the disks, but you will not be able to manage your subsystem or to receive alerts. The subsystem summary will display the last known state of the subsystem, but you will not be able to perform any administrative operations.

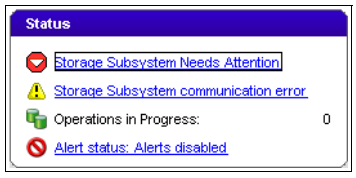


Figure 10-8 Status box - connection lost

10.2 Hardware components

Below the status box is an overview of the installed hardware components (Figure 10-9), that is, the physically installed components in this storage subsystem. The upper part of the box shows the number of installed controllers and the number of connected enclosures. The lower part of the box shows the number of installed hard drives, the drive type, and an overview of the hot spare drives.

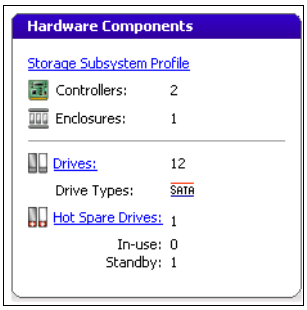


Figure 10-9 Hardware Components

The box also includes three links to more detailed information, described in the following sections.

10.2.1 Storage Subsystem Profile

Configuring a storage server can be a complex task: It is essential to document the configuration and every configuration change. This documentation can be saved in a file known as a *subsystem profile*. This profile stores information about the controllers, attached drives and enclosures, microcode levels, arrays, logical drives, and storage partitioning.

To generate the subsystem profile, click the **Storage Subsystem Profile** link in the Hardware Components box (Figure 10-9). The information will be gathered from the various components and the profile window appears (Figure 10-10). Use the **Save As** button to save the profile locally, to help document a change history of your storage subsystem. It will be saved as a text file. You can look at the profile later, as described in 14.2, “View Storage Subsystem Profile” on page 304.

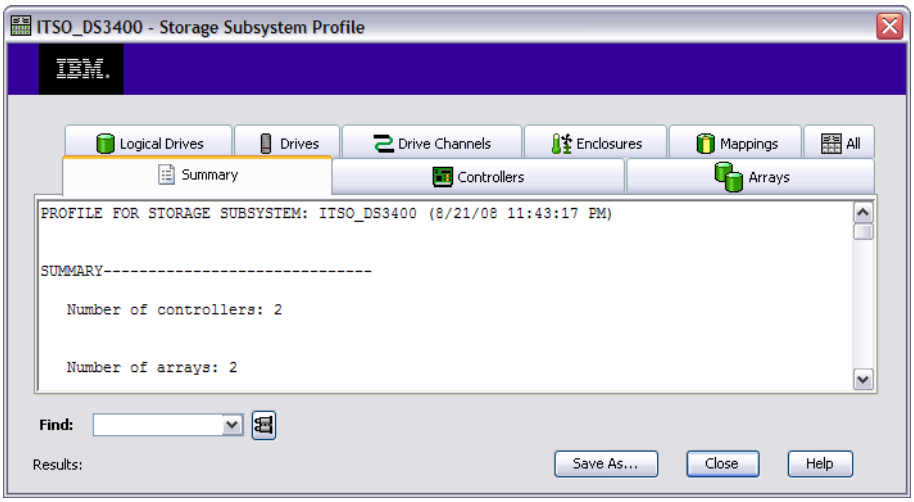


Figure 10-10 Storage Subsystem Profile

Note: We recommend you that you save a new version of the profile and store it securely whenever the configuration changes. Even in the case of a complete configuration loss, you can restore the arrays and logical drives configuration as well as the mappings for the storage partitioning. The profile should be stored locally, and also be included in any offsite disaster recovery documentation.

10.2.2 Drives

This section displays the overall installed hard drives in this box, and the drive type. If you click the **Drives** link (Figure 10-11), the Locate Drives Window (Figure 10-12) appears.

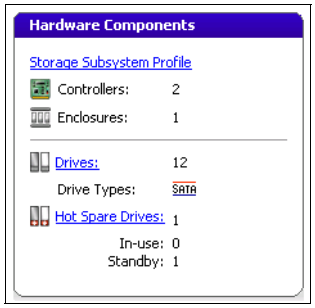


Figure 10-11 Hardware Components - Drives link

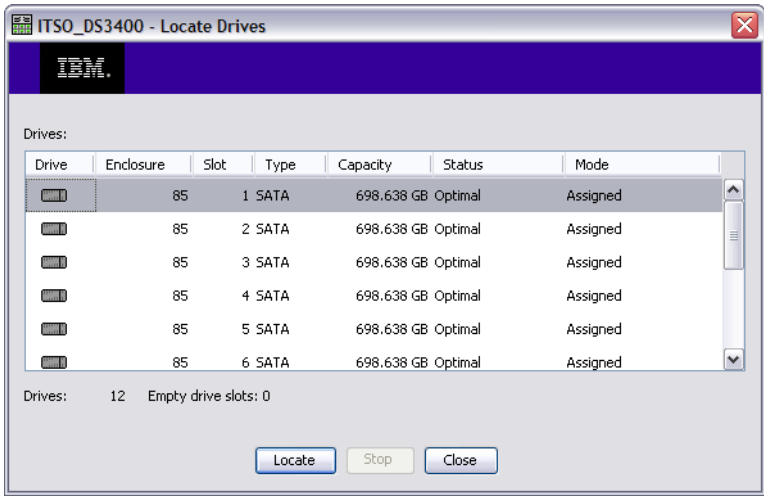


Figure 10-12 Hardware Components - Locate Drives

The Locate Drives windows lists the slots and enclosure where each drive is installed, and also the type, capacity, status, and mode of each drive. The drive mode can be:

- Assigned: The drive is assigned to an array.
- Unassigned: The drive is not a member of an array.
- Hot Spare Drive: The drive is a hot spare drive.

For more information about assigning drives to an array, see “Create logical drives” on page 168.

To locate one or more drives physically in your disk subsystem, highlight them and click **Locate** (Figure 10-13). The yellow light on the drive will start flashing, so that you can locate them.

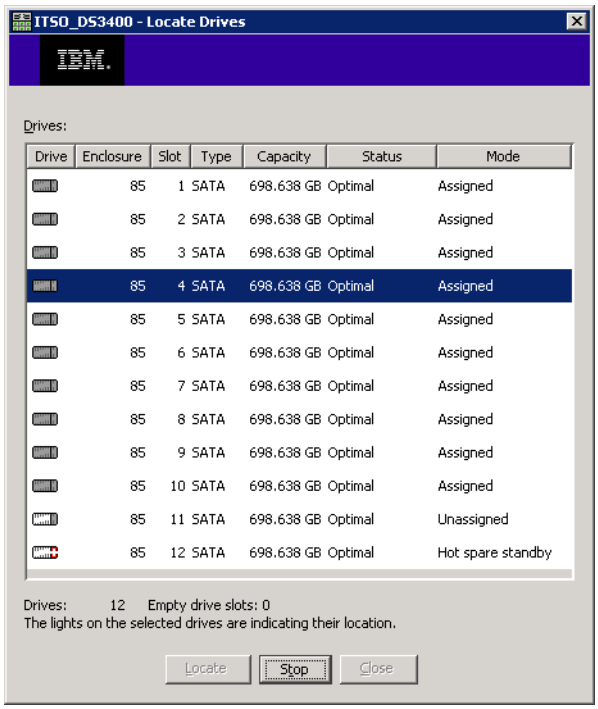


Figure 10-13 Hardware Components - Locate Drives flashing

10.2.3 Hot spare drives

This box displays information about hot spare drives (Figure 10-14).

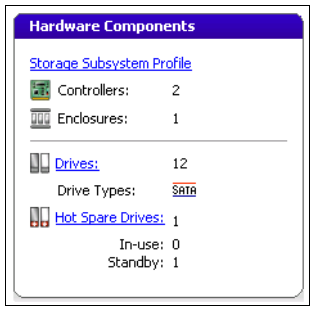


Figure 10-14 Hardware Components - hot spare overview

Hot spare disk drives provide additional protection to be used in the event of a disk drive fault in a RAID array (RAID 1, 3, or 5). A hot spare drive is like having a replacement drive in advance. This window displays only an overview of how many hot spare drives are assigned, and how many of them are in use.

To get a more detailed view, click the **Hot Spare Drive** link to open the Locate Drives window (Figure 10-15). This window covers the same view and functions as the previous Locate Drives window (10.2.2, “Drives” on page 213) for hot spare drives. More information about assigning hot spare drives is given in “Configure Hot Spare Drives - Automatic” on page 164.

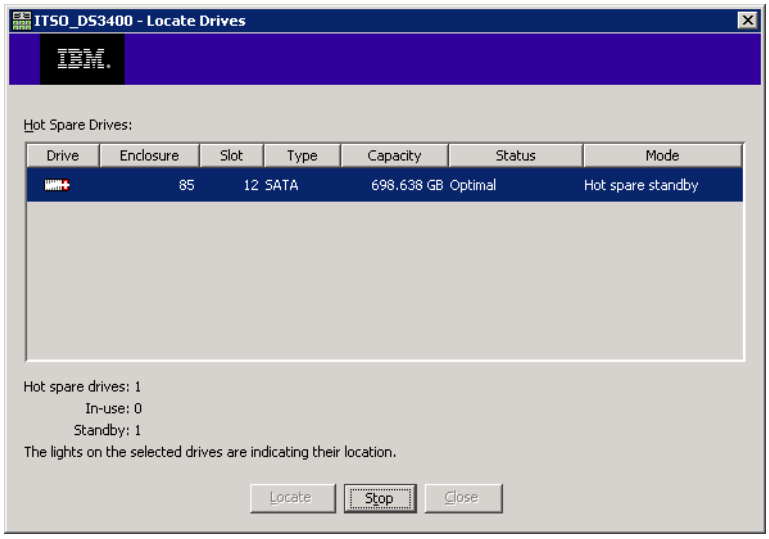


Figure 10-15 Hardware components - locate hot spare

10.3 Capacity

The Capacity window is located at the top of the middle column (Figure 10-2 on page 209). It displays the overall capacity of your system, that is, the total capacity of the installed drives. Figure 10-2 on page 209 shows an unconfigured subsystem where all space is available. After you create logical drives, this view gets updated to show the current total capacity (Figure 10-16).

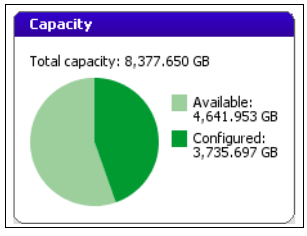


Figure 10-16 Capacity - Configured and unconfigured total capacity

10.4 Hosts and mappings

Below the Capacity box is an overview of your hosts and mappings (Figure 10-17).

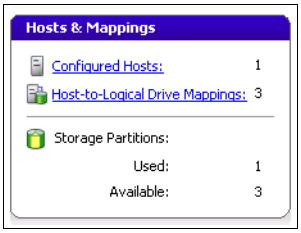


Figure 10-17 Hosts & Mappings

Hosts are single systems that can be mapped to a drive. This means a host is a system that accesses the logical drive through the attached technology (SAS, Fiber, or iSCSi). If multiple hosts access a group of logical drives, they all would have the equal access to these drives. Since most operating systems do not usually allow multiple hosts to access the same logical drives, you must create storage partitions. You do this by mapping specific logical drives to the host ports of the host systems. Configuring host access is described in 9.2.4, “Create host-to-logical drive mappings” on page 151.

10.4.1 Configured hosts

This line shows you how many hosts are currently configured on this storage subsystem (Figure 10-17). Click the link to display the configured host names (Figure 10-18). This view is only an overview of the configured hosts. The DS3000 server only communicates through the use of Word Wide Names (WWN). The storage subsystem is not aware of which Host Bus Adapters are in the same server or cluster unless you configure it. More information about host configuration is covered in 9.2, “Configure hosts” on page 140.

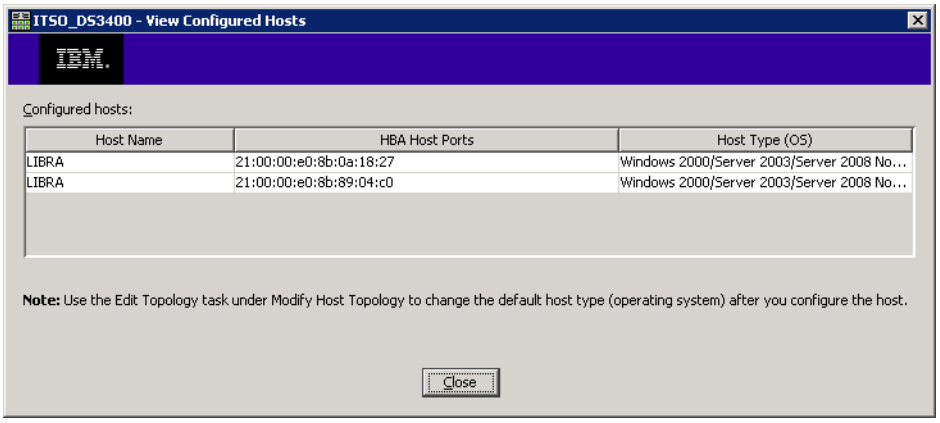


Figure 10-18 Hosts & Mappings - configured host list

10.4.2 Host-to-Logical Drive Mappings

Host-to-Logical Drive Mappings is the second link in this box. Next to the link you see the number of hosts that are currently mapped to a logical drive. This means the number of hosts that are allowed to see a logical drive. Opening the link displays a more detailed view (Figure 10-19).

In this view, the Logical Drive Name is displayed next to the mapped host (Accessible By column), so it is easy to discover which host is allowed to access which logical drive. The view also includes the LUN Number, the Logical Drive Capacity, and the Type. More information about drive mappings is provided in 9.2.4, “Create host-to-logical drive mappings” on page 151.

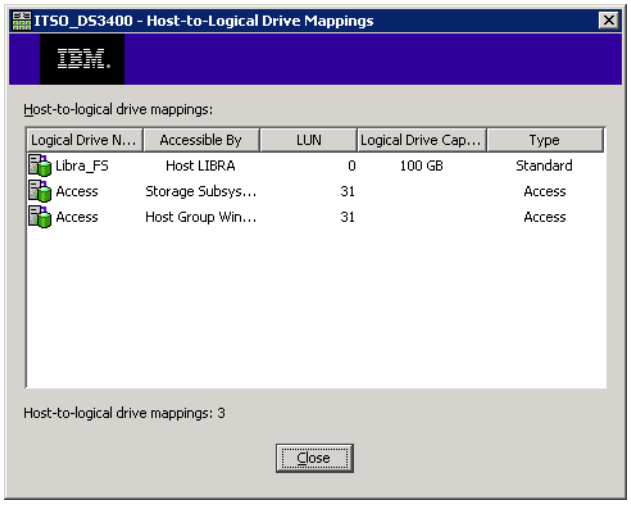


Figure 10-19 Hosts & Mappings - Host-to-Logical Drive Mappings

10.4.3 Storage partitions

This part of the box displays information about the number of storage partitions that are used and available on this storage subsystem. Storage partitions are enabled by an additional feature pack, so you can either order enough storage partitions for your planned configuration at the initial setup, or you can also get them later as an upgrade. Table 10-1 shows how many storage partitions are delivered as standard with each model DS3000 and which packs are available as an upgrade. More information about Storage Partitions is provided in “Create host-to-logical drive mappings” on page 153. The Storage Partition Box on the Summary tab shows you how many of them are activated on your subsystem and how many are in use.

Table 10-1 Storage partitions overview

Model	Standard storage partitions	Upgrade to
DS3200	4	8 or 16 or 32
DS3300	4	8 or 16 or 32
DS3400	4	8 or 16 or 32

10.5 Arrays and logical drives

At the top of the right column is the Arrays & Logical Drives Box (Figure 10-2 on page 209). This box shows how many arrays and logical drives are configured on your system (Figure 10-20). An array is a set of physical drives that the controller groups together to provide one or more logical drives to the host or host group.

While creating an array, you can also choose the array's RAID; for more details about configuring arrays and logical drives, see 9.3, "Configure storage" on page 155. Depending on your configuration, this affects the system's fault tolerance and performance. The bottom part of the box shows which levels are configured on your storage subsystems, how many arrays of this RAID level exist, and how many logical drives are in each.

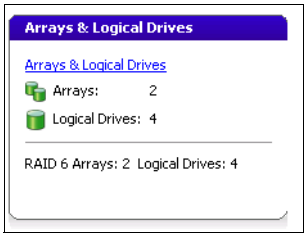


Figure 10-20 Arrays & Logical Drives

For a closer look at your arrays, click the **Arrays & Logical Drives** link. The Locate Arrays window (Figure 10-21) will appear, showing the total unconfigured capacity (if there is any left), and the individual arrays. Expand the arrays to see which logical drives they include. In our Example, Array 1 has two logical drives, and all other arrays have one logical drive. On the two Logical drives on Array 1, you can see a small clock. This means that there is a task in progress, for example, initializing arrays or editing segment size.

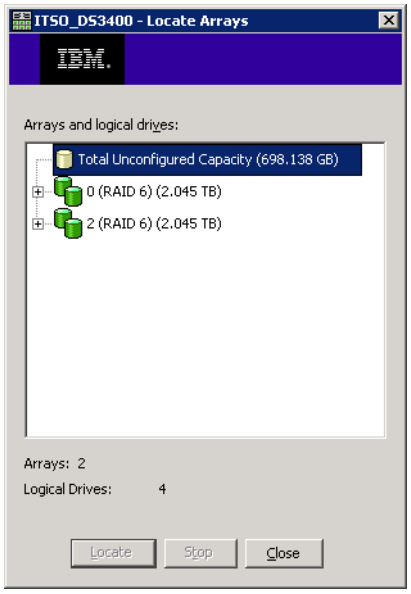


Figure 10-21 Arrays & Logical Drives - Locate Arrays

If you highlight one of the arrays, you can then click **Locate**. If you do that, the physical drives that belong to the array will start flashing (Figure 10-22). This helps identify the physical locations of the grouped drives.

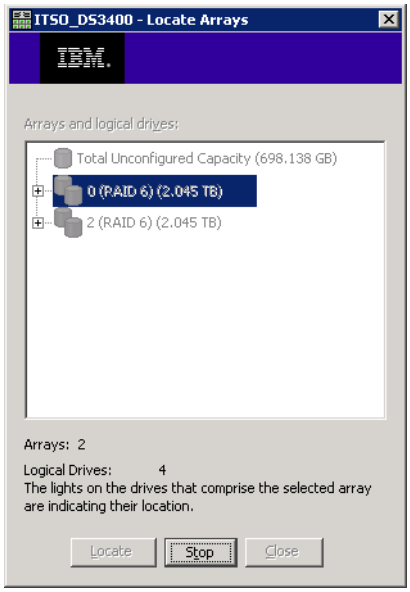


Figure 10-22 Arrays & Logical Drives - drives flashing

10.6 Information Center

At the right bottom edge is the Information Center (Figure 10-23). It includes several links to help you understand and set up a storage subsystem, as well as answering any questions you might have.

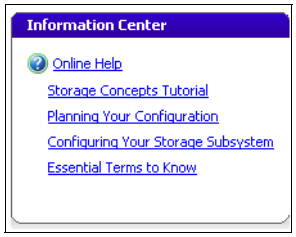


Figure 10-23 Information Center

- Online Help

  This link opens the Online Help at the start window (Figure 10-24 on page 220). Click one of the other links in the Information Center window to navigate in the online help directly to the related content:

  - Storage Concepts Tutorial
  - Planning your Configuration
  - Configuring Your Storage Subsystem
  - Essential Terms to Know

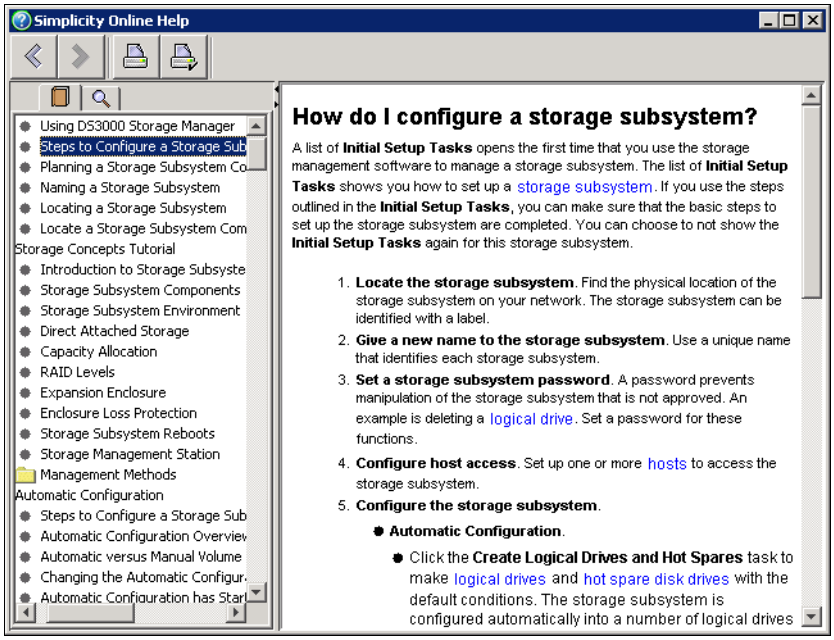


Figure 10-24 Information Center - Online Help

11

Administration - Modify

We describe how to initially configure the DS3000 Storage Subsystem in Chapter 9, “Administration - Configure” on page 131 by using the Configure tab. The Modify tab (Figure 11-1 on page 222) contains options to change your existing storage subsystem configuration.

Like the Configure tab, the Modify tab also contains two parts:

- Hosts
  - Viewing and editing the host servers configuration
  - Editing the host-to-logical volume mappings
- Storage
  - Renaming or deleting logical drives and arrays
  - Adding free capacity
  - Changing RAID types
  - Drive replacing
  - Changing controller ownership
  - Modifying FlashCopy logical drives
  - Managing logical drive copies

We will now look at each of these in turn.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

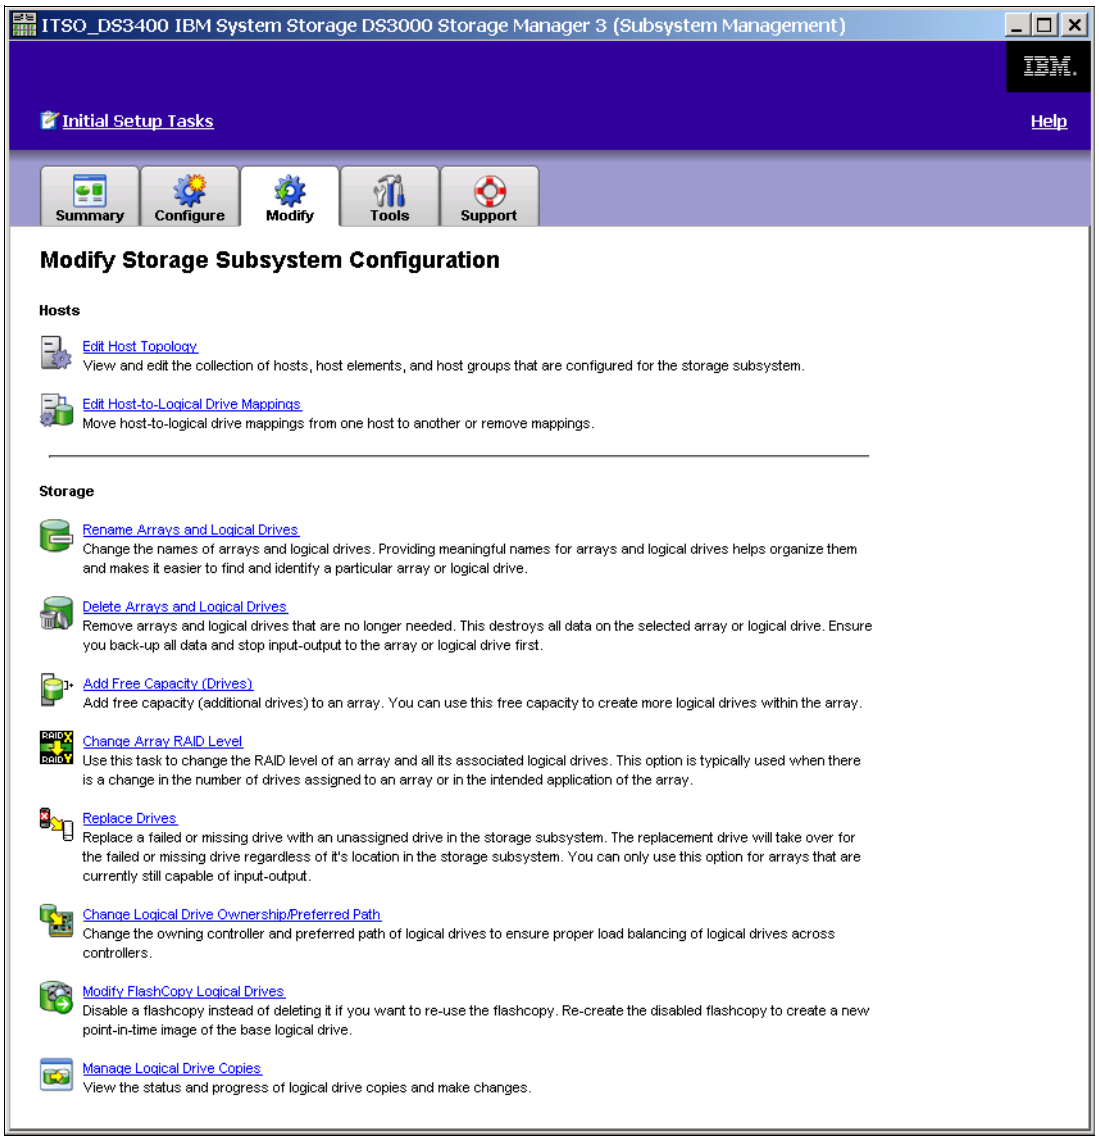


Figure 11-1 The Modify tab

11.1 Modify hosts

Here we discuss the modification of hosts.

11.1.1 Edit Host Topology

Click the **Edit Host Topology** link in Hosts section of the Modify tab to access functions to change your host configuration, including hosts, host groups, and host ports.

Figure 11-2 shows the main window.

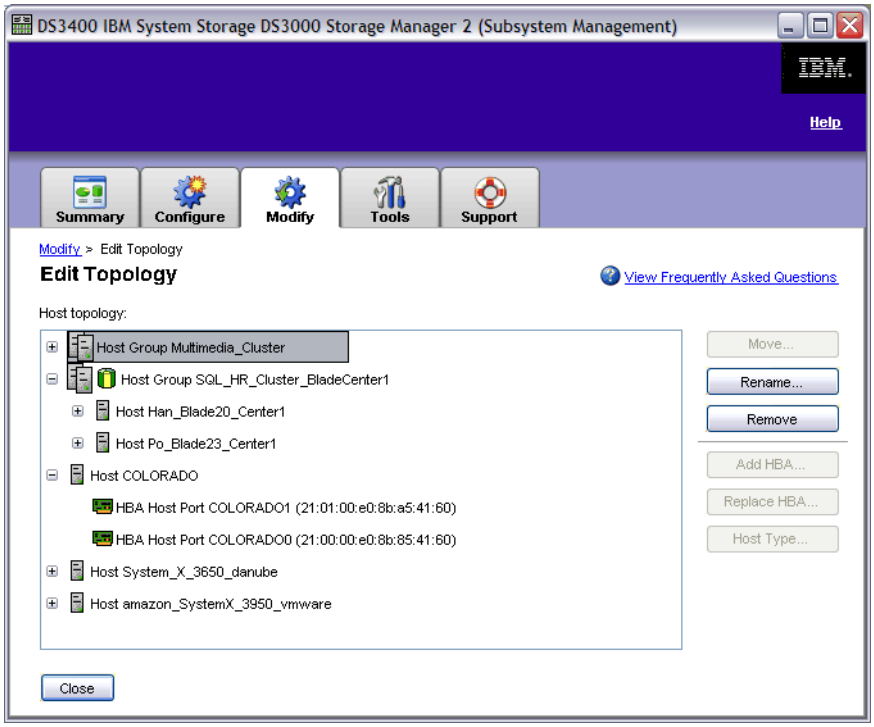


Figure 11-2 Modify - Edit Topology

The “Host topology” item displays the current host configuration, including all host groups, hosts, and host ports (HBAs) with their associated World Wide Port Name (WWPN). On the right side of the window there are buttons for the possible configuration changes. The list of valid (highlighted) buttons depends on your selection from the “Host topology” view:

- Move
- Rename
- Remove
- Host Type Change
- Add HBA
- Replace HBA

Move

The Move button is active when you select either a host (Figure 11-3) or an HBA.

Move host

If you select a host (either a single host or a member of a host group), you can:

- Move the host to an existing host group.
- Move the host from one host group to the other.
- Move the host out of the host group.

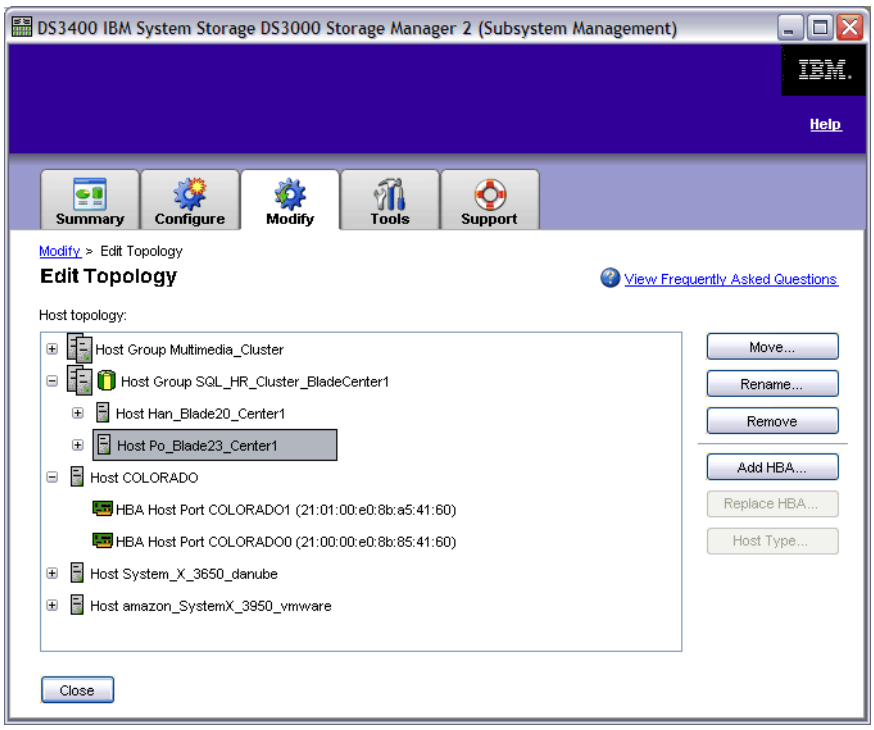


Figure 11-3 Edit Topology - select host

Highlight the host you want to change and click **Move**. The Move Host window appears (Figure 11-4). From the drop-down menu, either select a different host group to which you want to move the host or select **Move out of host group** and click **OK**.

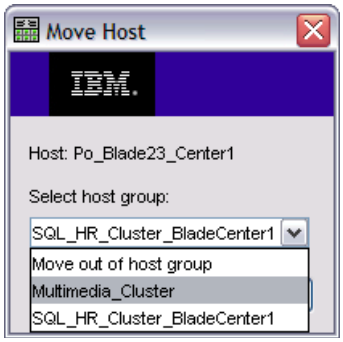


Figure 11-4 Edit Topology - Move Host

If you are moving the host to a different group, you will see a confirmation window. Click **Yes** to move the host.

The configuration change will be submitted to the storage subsystem and the host topology view will be updated.

Move HBA

You can also move an HBA from one host to the other. Select the HBA (Figure 11-5) and click **Move**.

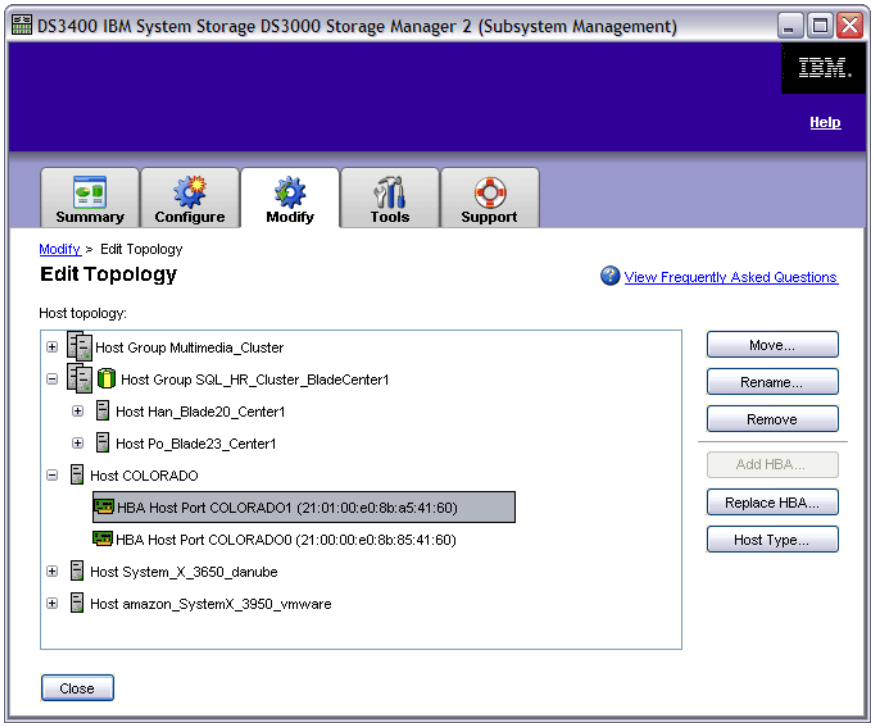


Figure 11-5 Edit Topology - select HBA

The Move HBA Host Port window appears (Figure 11-6).

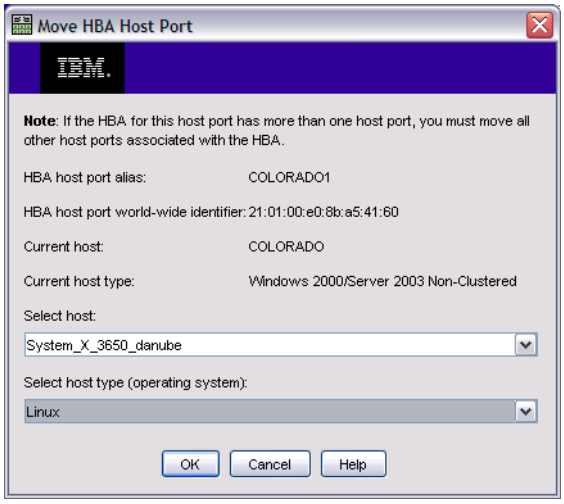


Figure 11-6 Edit Topology - Move HBA

The window shows a summary of the selected HBA properties; review them to be sure you have selected the right HBA. You also see two drop-down menus:

- In the Select Host drop-down menu, choose the host where you want to move the HBA.
- In the Select Host type drop-down menu, change the operating system if the new host runs a different one.

Click **OK** to start the configuration change. The new configuration will be submitted to the storage server and the host topology view will be updated.

Rename

The Rename button is always active, no matter what the selection. You can rename all listed hots, host groups, and HBA Ports. Highlight an item in the topology and click **Rename**. The Rename window (Figure 11-7) will appear. Enter a new unique name and click **OK**.

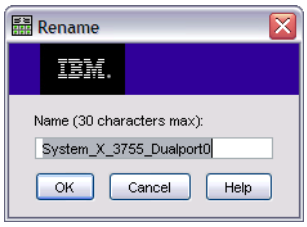


Figure 11-7 Edit Topology - rename host

Remove

The Remove button is also always active. Highlight an item and click **Remove**. Depending on what you remove, there will be a warning, as shown in Figure 11-8. Review the information, and click **Yes** to confirm. The selected item will be deleted.

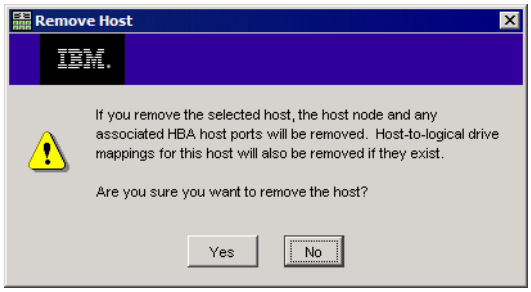


Figure 11-8 Edit Topology - Remove Host

Note: If you remove an item, all of its associated sub-items will be deleted too. This means that if you delete a host, the HBA Ports will also be deleted. Deleting a host group will delete the group and all the hosts the group includes. Deleting a host also deletes the host-to-logical drive mappings for this item.

After deleting a host, that host will lose all its previously defined access to the logical drives. The WWPNs will appear again in the “Known HBA host ports” view of the configure Host Access Wizard (Figure 9-20 on page 147). To recreate this host and give it access to logical drives, see 9.1, “Configure host access” on page 132.

Add HBA

Here you can add HBA Ports to an existing host. Highlight a host and click **Add HBA**. The Add HBA Host Port window (Figure 11-9) appears.



Figure 11-9 Edit Topology - Add HBA Host Port

If the port WWPN is already visible to the storage server, you can select this WWPN from the drop-down list. If it is not listed, enter the port WWPN manually in this field. To rescan for new HBAs, click **Refresh**.

Select the host operating system from the “Host Type” drop-down menu, and also enter a unique name for the new host port in the “HBA host port alias” field. Click **Add** and the new host port will be added to your host.

Replace HBA

Remember that storage partition assignments are based on WWPN of ports; the DS3000 only communicates through the WWPNs (see 9.1.1, “What a host is” on page 132). If you have to replace an HBA, the adapter WWPN will change. Since the new WWPN does not appear in any of the storage partition assignments after replacement, this host port will have no access to any logical drives. To re-enable the new WWPN’s access to the assigned logical drive, you have to tell the DS3000 about the new HBA.

After you have replaced the HBA and connected the cable, highlight the old HBA host port in the Host topology view. Click the **Replace HBA** button, and the “Replace HBA Host Port” window appears (Figure 11-10 on page 228). If the storage subsystem already sees the new HBA, you will be able to select the WWPN in the “New HBA host port world wide identifier” drop-down list; otherwise, type the new WWPN manually in this field. If you want to edit the HBA host port alias, you can also do this here. Click **OK** to save the settings to the storage subsystem. Now the host will be able to access its logical drives again over the new host port.

Note: If you want to add the HBA’s WWPN to the host topology before the DS3000 recognizes the HBA on SAN, for example, if proper SAN zoning is not yet done, make sure to type the correct WWPN of the HBA. By setting the World Wide Node Name (WWNN) of the HBA, you will not be able to enable Host-to-LUN mapping correctly. The WWPN and the WWNN differs in one digit only in many cases, but it is important.

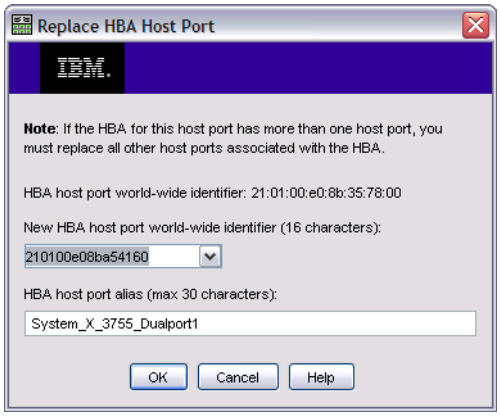


Figure 11-10 Edit Topology - Replace HBA

Host Type Change

Every host has an associated host type, indicating the operating system installed. The Host Type Change button allows you to change this host type. Select the host in the Host topology view and click **Change**. The new window shows the host name and the current host type; you can select the new host type in a drop-down menu. Select the appropriate choice and click **OK** to make the change (Figure 11-11).

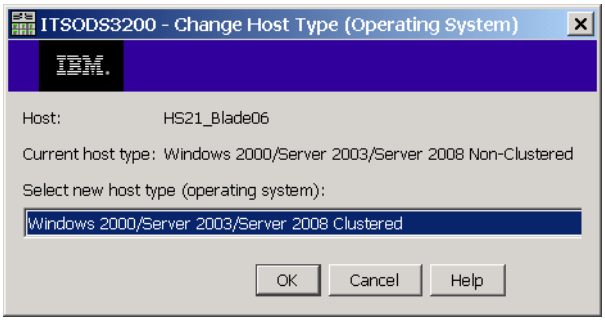


Figure 11-11 Edit Topology - Change Host Type

Note: In the previous version of DS3000 firmware (6.xx), you had to change the host type on every single host port (HBA) of the changed host. In redundant configurations (recommended), you have at least two host ports in one system, so you had to change the host type of every host port that was installed in this host server. Starting with Version 7.35 of the DS3000 firmware, you change the host type of the host server only, as shown in Figure 11-11 on page 228.

11.1.2 Edit host-to-logical drive mappings

In 9.2.4, “Create host-to-logical drive mappings” on page 151, we describe the principles of storage partitioning and host-to-logical drive mapping. In this part of the Modify tab (Figure 11-1 on page 222), you can change previous defined mappings.

Some reasons for changing the mapping could be:

- Changing the logical unit number (LUN)
- Changing the host that is mapped to a logical drive
- Reconfiguration of your storage subsystem

When you change a logical drive mapping, the changes apply to both the hosts and the host groups.

Note: Before you move a LUN from a storage partition, you should stop any I/O to this logical drive and you should unmount it in the OS of the host server.

Host-to-logical drive mapping requires that the Storage Partitions feature be enabled and, by default, four storage partitions are enabled on all the DS3000 systems. When you need more than four, you can enable up to 32 partitions using the Storage Partitions Premium Feature (see 12.1.4, “View/Enable Premium Features” on page 265). The number of licensed storage partitions means how many host servers can be exclusively mapped.

Change mapping

The procedure for moving a LUN from one storage partition to other is as follows:

1. Click **Edit Host-to-Logical Drives Mapping** (Figure 11-1 on page 222) to open the associated view (Figure 11-12).

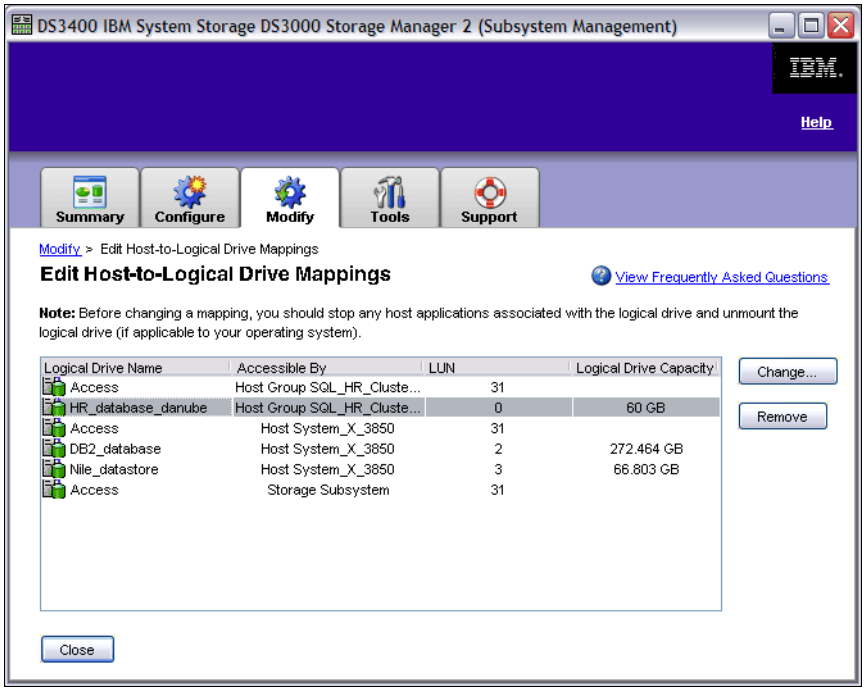


Figure 11-12 Edit Host-to-Logical Drive Mappings

2. You can see the logical drive name, the host group or the host that can access it, the LUN number, and the capacity. Select the entry you want to change and click **Change**.

The “Change Mapping” window appears (Figure 11-13).

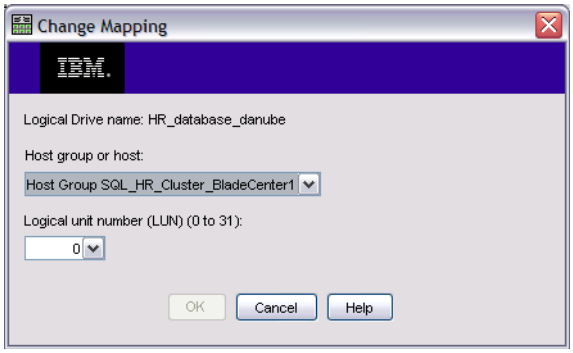


Figure 11-13 Change Mapping - change LUN

3. Select a new host or host group from the first drop-down menu and select the required LUN number from the second drop-down menu. Click **OK**, and a confirmation window appears (Figure 11-14).

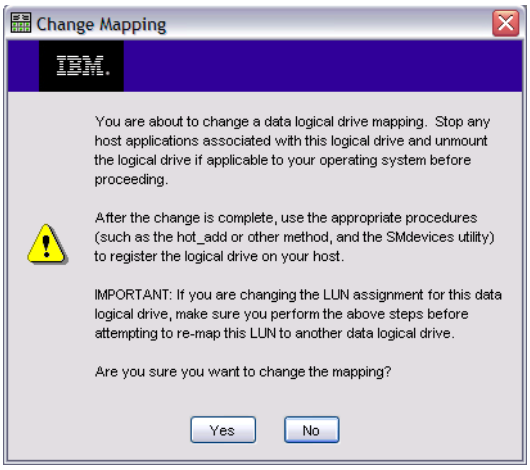


Figure 11-14 Change Mapping

4. Read it and click **Yes** to confirm.

Remove mapping

The procedure for removing a LUN from the storage partition is as follows:

1. Click **Edit Host-to-Logical Drives Mapping** (Figure 11-1 on page 222) to open the associated view (Figure 11-12 on page 229).
2. You can see the logical drive name, the host group or the host that can access it, the LUN number, and the capacity. Select the entry you want to change and click **Remove**.
3. The Remove Mapping window appears (Figure 11-15 on page 231).

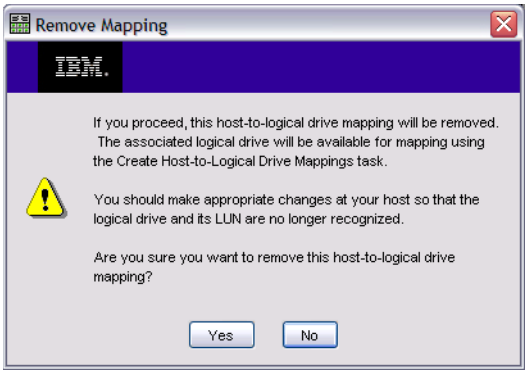


Figure 11-15 Remove Mapping

4. Read the information and click **Yes** to confirm.
5. The mapping is now removed from your configuration. It means that no host server sees this logical drive.

11.2 Modify storage

The second section shown in Figure 11-1 on page 222 is *Storage*. It contains the following functions to modify, delete, or manage all kinds of logical drives and arrays:

- Rename Arrays and Logical Drives
- Delete Array and Logical Drives
- Add Free Capacity (Drives)
- Change Array RAID Level
- Replace Drives
- Change Logical Drive Ownership/Preferred Path
- Modify FlashCopy Logical Drives
- Manage Logical Drive Copies

The DS3000 firmware V7.35 introduces some new features:

- Rename of array
- Array without any logical drives defined
- RAID level migration
- Drive replacement

These features are not available in previous 6.xx versions of the DS3000 firmware.

11.2.1 Rename Arrays and Logical Drives

Throughout this book, we have highly recommended using meaningful names for storage subsystem components, in order to make identification and management tasks easier.

To rename, click **Rename Array and Logical Drives** to start the wizard (Figure 11-16 on page 232). Expand the arrays and select the array or the logical drive you want to rename. Type the new name in the “Name” field and click **Rename**. The selected array or logical drive will be renamed immediately. To rename another logical drive, select another drive; otherwise, click **Close**.

Names can have maximum of 30 characters, can include letters, numbers, and the special characters: underscore (_), minus (-), and pound/number sign (#). No other special characters are permitted.

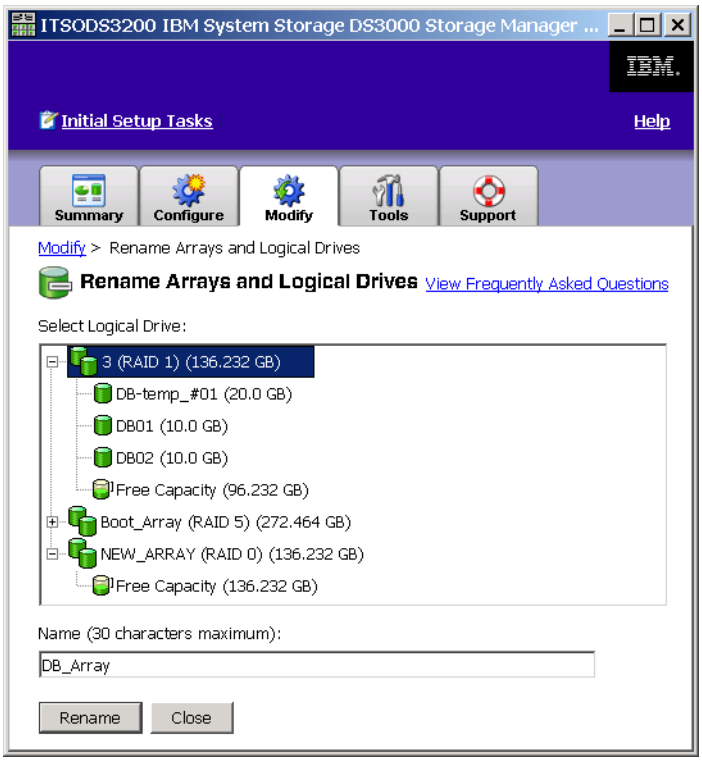


Figure 11-16 Rename Array and Logical Drives

11.2.2 Delete Arrays and Logical Drives

Click this option to remove arrays or logical drives from your system. The Delete Arrays and Logical Drives view appears (Figure 11-17 on page 233). Select the logical drive or array you want to delete and click **Delete** to remove it from the DS3000 subsystem.

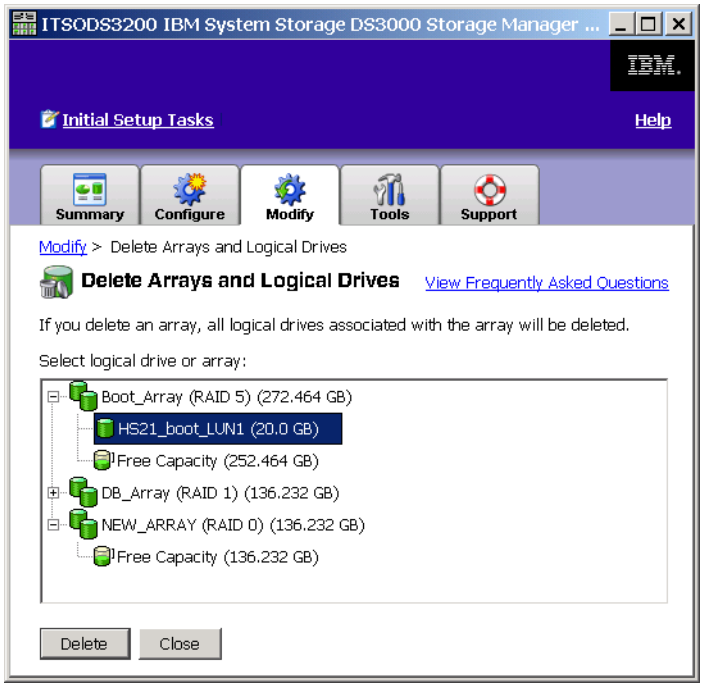


Figure 11-17 Delete Arrays and Logical Drives

Important: Before deleting a logical drive, observe the following considerations:

- Deleting a Logical Drive causes loss of all data on the logical drive. Back up the data and stop all I/O before performing this operation, if necessary.
- If you delete an array, all logical drives associated with the array will be also deleted. If you delete the last remaining logical drive in an array, the array will not be deleted. It is a new feature coming with the DS3000 firmware V7.35.
- If a file system is mounted on the logical drive, unmount it before attempting to delete that drive.
- Deleting a base logical drive automatically deletes any associated FlashCopy logical drives and FlashCopy repository logical drives.
- Deleting a FlashCopy logical drive automatically deletes the associated FlashCopy Repository logical drive.

After you click **Delete**, a warning window appears (Figure 11-18). To continue, type yes in the prompted box and click **OK**.

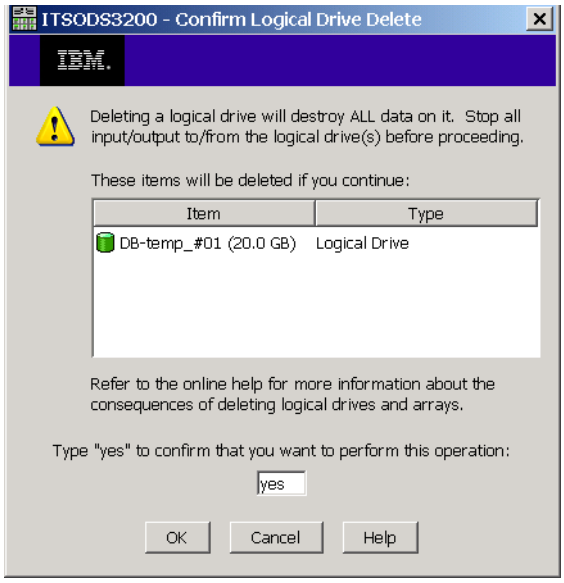


Figure 11-18 Confirm Logical Drive Delete

This task will take a few seconds, as shown in the progress window (Figure 11-19). Click **OK** to return to the Delete Logical Drives view. To delete another logical drive or array, repeat these steps.

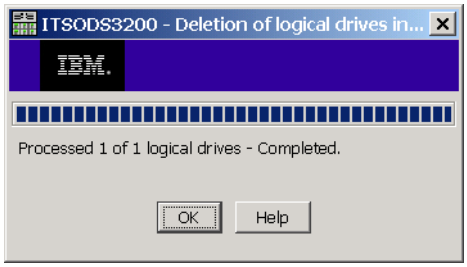


Figure 11-19 Logical drive deleted

11.2.3 Add Free Capacity (Drives)

With this option, you can expand an array by adding additional disk drives to an existing array. Use this option if you have installed new drives in a DS3000, or if there is spare disk capacity that is not yet assigned to any array.

1. Click the link to start the Add Free Capacity (Drives) wizard. If there is no free capacity in the DS3000, you will see the error message shown in Figure 11-20 on page 235. You must first install additional drives.

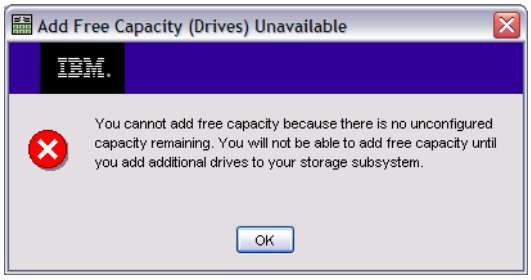


Figure 11-20 No free capacity available

2. If there is free capacity available, the Select Array view will appear (Figure 11-21). Select the array to expand and click **Next**.

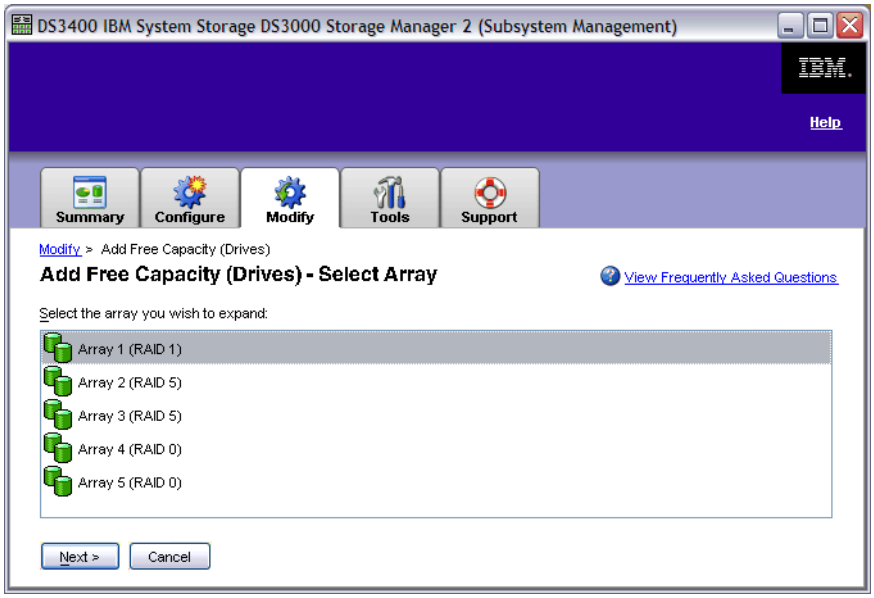


Figure 11-21 Add Free Capacity (Drives) - Select Array

3. The **Select Capacity** view will appear (Figure 11-22), showing a summary of the array.

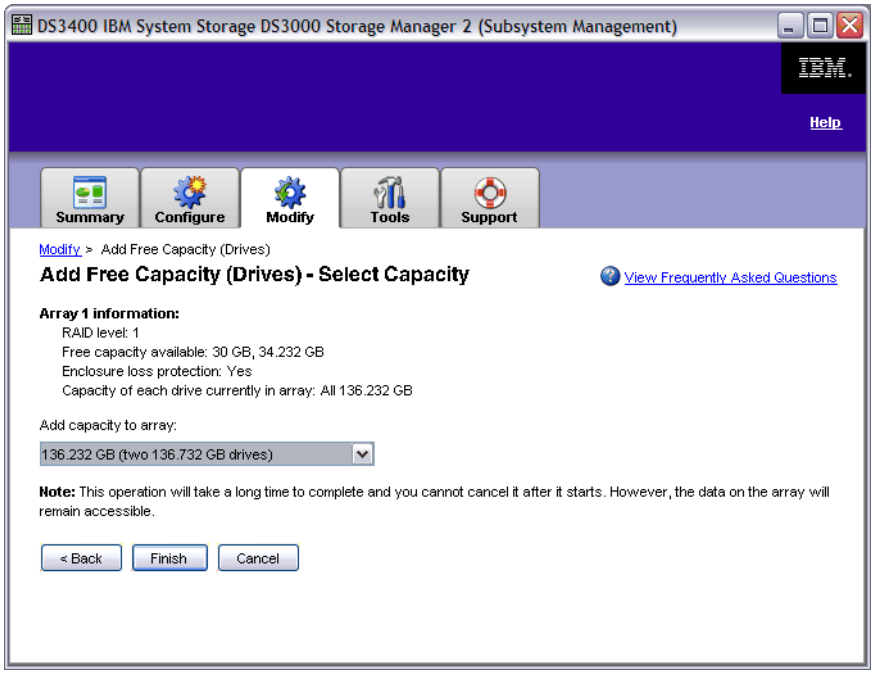


Figure 11-22 Select capacity

4. From the drop-down menu, select how many drives you want to add to this array. You must add two drives for a RAID 1 or RAID 10 array. For other types of RAID, you can select one or two drives in one step. If you need to add more than two disk drives, you have to wait while the first two drives are added successfully (it can take hours) and then you can add more drives. Click **Finish** to expand the array (Figure 11-23). Your operation will only take few minutes to complete, but the background procedure of adding disk drive(s) can take hours to complete and you cannot cancel it after it starts. However, existing data on the array remains accessible. Click **OK** to return to the Modify tab.

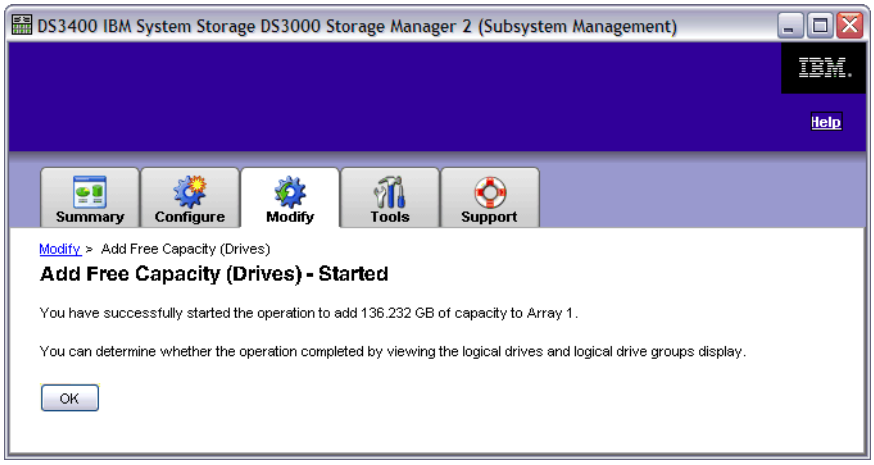


Figure 11-23 Add Free Capacity (Drives) - Started

11.2.4 Change Array RAID Level

This new option in the V7.35 firmware enables you to change the RAID level of given array directly from the DS3000 Storage Manager. In the past, it was only possible by running SMcli commands. Do these steps to change the array's RAID level:

1. Open the appropriate window (Figure 11-24) by clicking the **Change Array RAID Level** link in the Modify tab (Figure 11-1 on page 222).

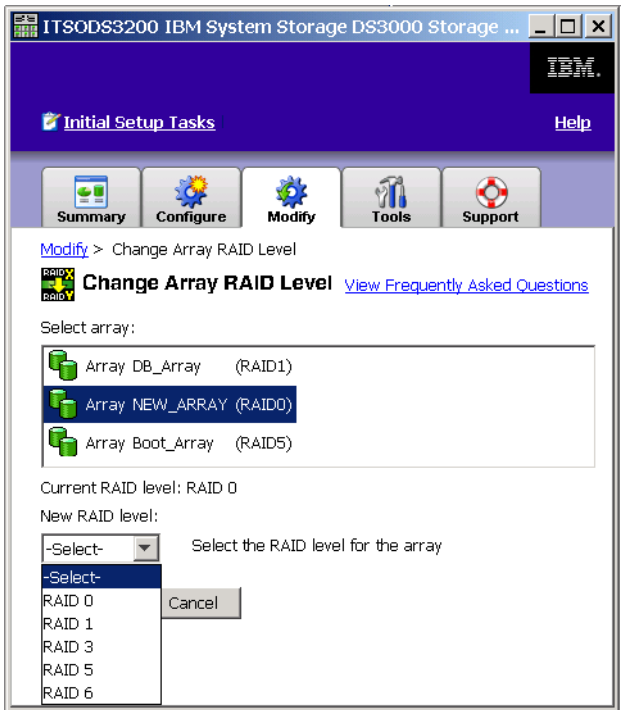


Figure 11-24 Change RAID type

2. Select an array for which the RAID level has to be changed.
3. Select target RAID level for the highlighted array in the New RAID level drop-down menu.
4. Click the **Change** button to start the RAID change process.
5. The confirmation window (Figure 11-25) appears. Read the important message and click **Yes** to continue.

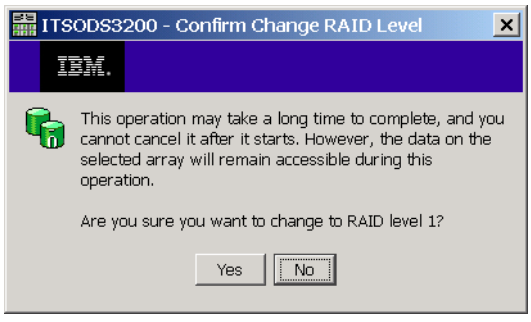


Figure 11-25 RAID level migration - Confirm

6. The final window (Figure 11-26) summarizes your RAID level change. Click **OK** to return to the Modify tab.

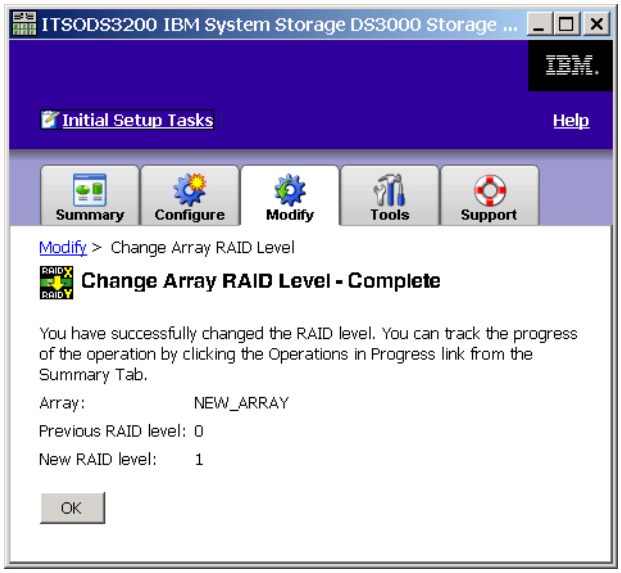


Figure 11-26 RAID level change confirmation

11.2.5 Replace drives

With the new DS3000 firmware V7.35, you can replace a failed or missing disk drive that is used by an array. The array is in “Degraded” status at this time (see the example of the Recovery Guru window in Figure 11-27). You can assign a working - unassigned disk drive, which is already inserted in the DS3000 Storage Subsystem, as a replacement drive.

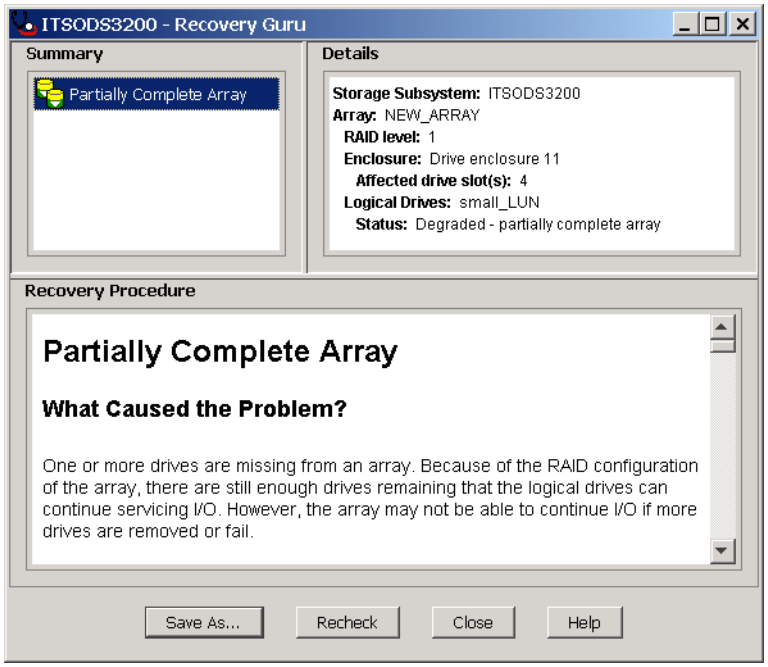


Figure 11-27 Degraded array - missing drive example

1. Click **Replace Drives** in the Modify tab (Figure 11-1 on page 222) if some disk drive is missing or failed in your DS3000. The list of eligible replacement disk drives appears in the Replace Drives window (Figure 11-28).

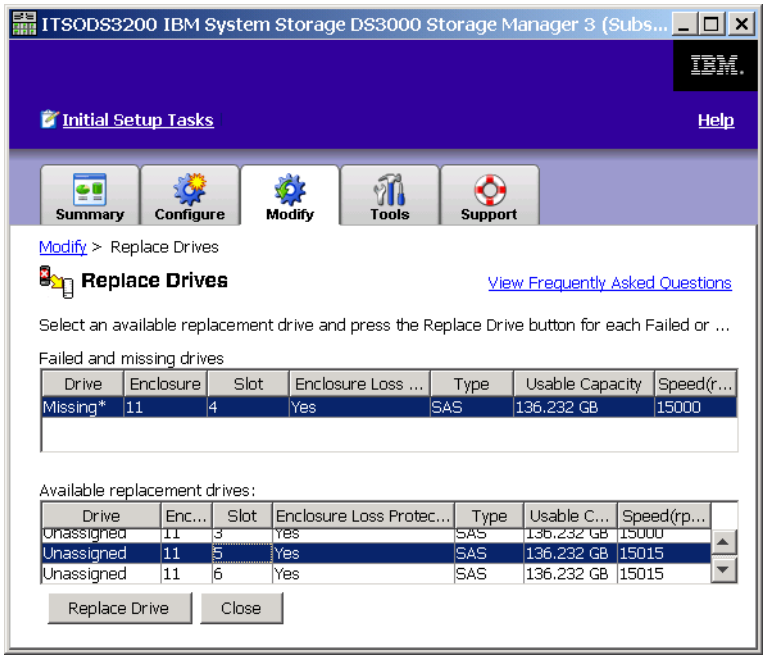


Figure 11-28 Replace Drives

2. Select the missing or failed drive from the first list (if more than one drive is missing or failed). Select the replacement drive from the second list of available drives. The available replacement disk drives are equal to or greater in size, must have the same or higher RPM speed, and must have the disk drive type as the failed disk drive or the missing disk drive. We also recommend selecting the drive that has a Yes entry in the column “Enclosure Loss Protection”, if such disk exists in the list.

   You also have a choice to select the HotSpare drive as a replacement drive (see Figure 11-30 on page 240). If a HotSpare disk drive is defined and is in use by the array with a missing or failed disk, you have to wait until the DS3000 controller reconstructs all the logical drives in that array. If you try it earlier, the error window (Figure 11-29) appears.

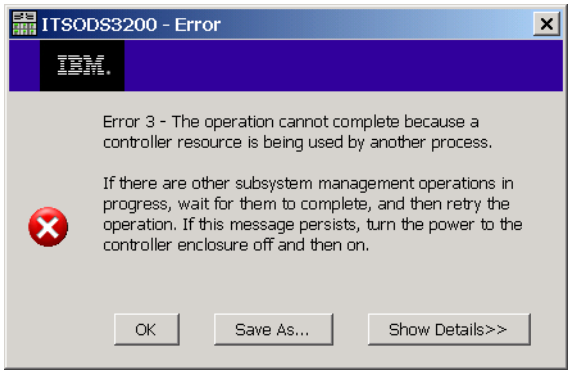


Figure 11-29 Drive replacement if array is reconstructed to HotSpare

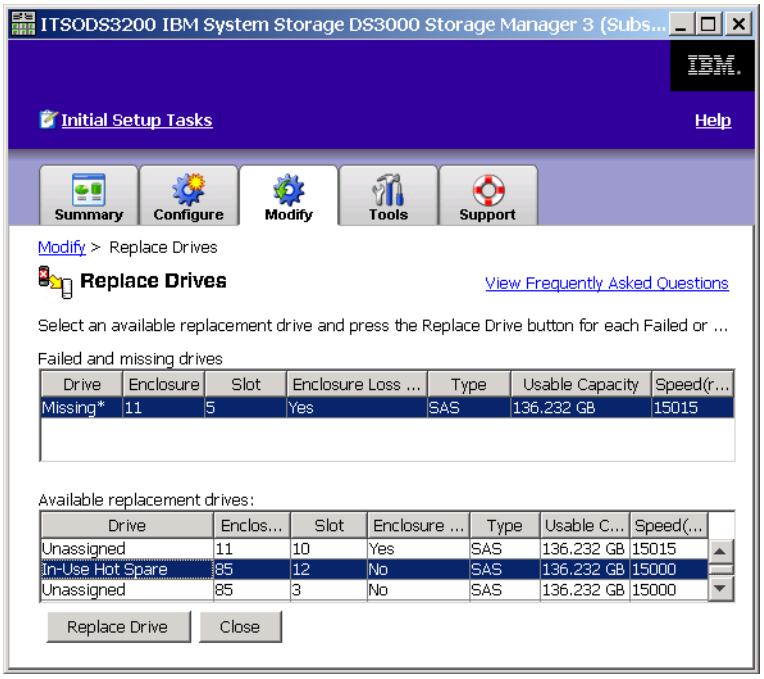


Figure 11-30 HotSpare as replacement drive

3. Click the **Replace Drive** button to start the operation. The confirmation window (Figure 11-31) appears.

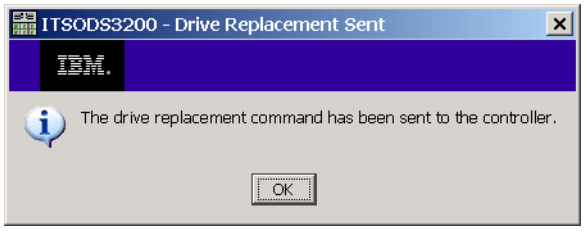


Figure 11-31 Drive Replacement Sent

Click **OK** and a new window (Figure 11-28 on page 239) shows the remaining missed or failed drives to replace, or no drives for replacement are displayed. Continue with the replacement or click **Close** to return to the Modify tab.

Warning: Because of possible loss of data access, use the Replace Drives option to logically change disk drives *only*. Physical removal and replacement of a failed disk drive may damage the array and the data on the drive.

If a HotSpare disk drive is used as replacement drive, do not forget to define another disk drive as HotSpare. Returning a missing disk or repaired disk does not set the HotSpare setting automatically.

11.2.6 Change Logical Drive Ownership/Preferred Path

If you are using a dual controller version of the DS3000, each logical drive has a preferred controller of ownership. This controller normally handles all I/O requests for this particular logical drive. In other words, each logical drive is owned by one and only one controller at any point in time. The alternate controller only takes over and handles the I/O requests if there is a failure in the I/O path, for example, a defective HBA or switch.

When defining logical drives, the system normally alternates ownership between the two controllers (workload is not taken into account). Therefore, you could end up with a configuration where there is an imbalance, so that one controller is handling much more I/O than the other. To balance the workload between the controllers, you can change the preferred ownership of a logical drive to the other controller. To change the preferred ownership of a logical drive, the controller path itself must be active; you cannot change the ownership of a drive if it is temporarily being handled by its alternate controller.

Important: Be sure that the operating system using the logical drive uses a multipath I/O driver; otherwise, it loses access to the logical drive.

Balancing traffic is unfortunately not always a trivial task. For example, if an application requires large disk space to be located and accessed in one chunk, it becomes harder to balance traffic by spreading the smaller volumes among controllers.

In addition, the load across controllers and logical drives is constantly changing. The logical drives and data accessed at any given time depend on which applications and users are active during that time period, which is why it is important to monitor the system.

The preferred controller ownership of a logical drive or array is the controller of an active-active pair that is designated to own these logical drives. The preferred controller owner is the controller that currently owns the logical drive or array.

If the preferred controller is undergoing a firmware download, ownership of the logical drives is automatically shifted to the other controller, and that controller becomes the current owner of the logical drives. If the preferred controller need to be replaced, you should disable the controller first by putting it into off-line mode. This will intentionally cause a failover of LUNs to the other controller and allow the preferred controller to be removed and replaced. This is considered a routine ownership change and is reported with an informational entry in the event log.

There can also be a forced failover from the preferred controller to the other controller because of I/O path errors. This is reported with a critical entry in the event log, and will be reported by the Enterprise Management software to e-mail and SNMP alert destinations.

During a VolumeCopy, the same controller must own both the source logical drive and target logical drive. If both logical drives do not have the same preferred controller when the VolumeCopy starts, the ownership of the target logical drive is automatically transferred to the preferred controller of the source logical drive. When the VolumeCopy is completed or is stopped, ownership of the target logical drive is restored to its preferred controller. If ownership of the source logical drive is changed during the VolumeCopy, ownership of the target logical drive is also changed.

Change Logical Drive Ownership/Preferred Path

Follow these steps:

1. Click the link to display the associated view (Figure 11-32). Select the logical drive to display the current owner and preferred path. The other controller is selected under Select new controller and preferred path. To change the logical drive to the other controller, click **Change**.

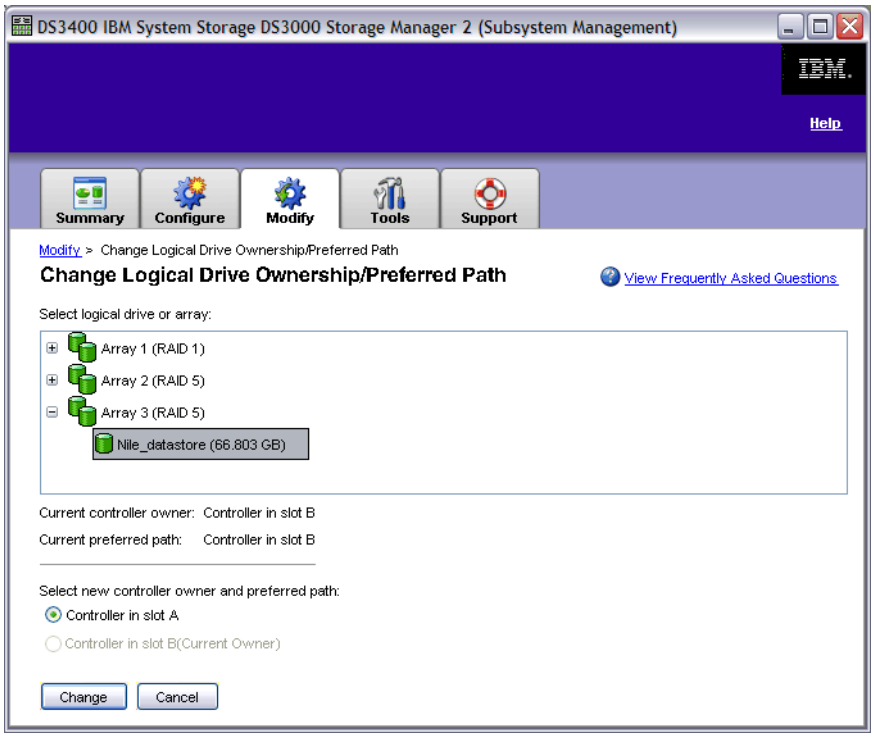


Figure 11-32 Current owner

2. Review the message in Figure 11-33 and click **Yes**.

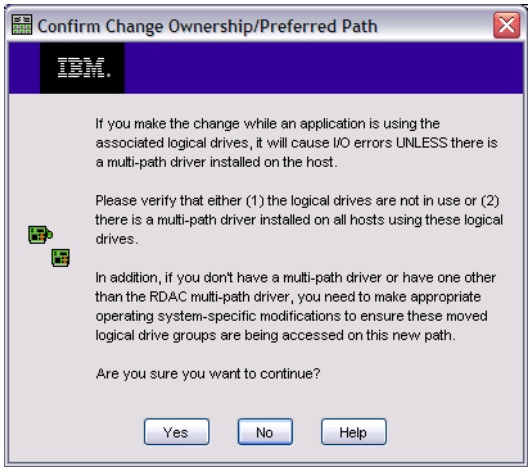


Figure 11-33 Confirm Change Ownership/Preferred Path

3. The command to change the logical drive ownership and preferred path is sent to the storage subsystem, as shown in Figure 11-34 on page 243. Click **OK**.

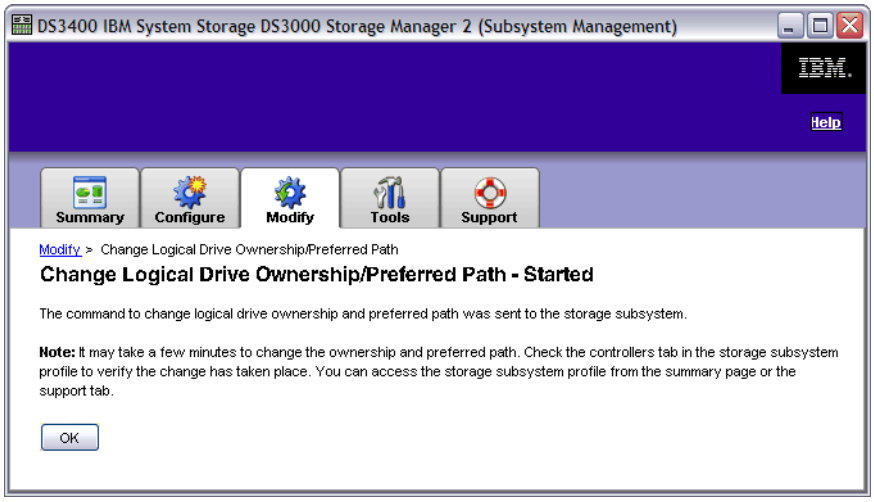


Figure 11-34 Change ownership - started

11.2.7 Modify FlashCopy Logical Drives

In 9.4, “Advanced functions - FlashCopy” on page 183, we describe how to create FlashCopy drives. In the Modify FlashCopy Logical Drives link, you can manage and modify FlashCopy drives with the following options, as shown in Figure 11-35:

- Disable FlashCopy Logical Drive
- Re-create FlashCopy Logical Drives
- Expand FlashCopy Repository

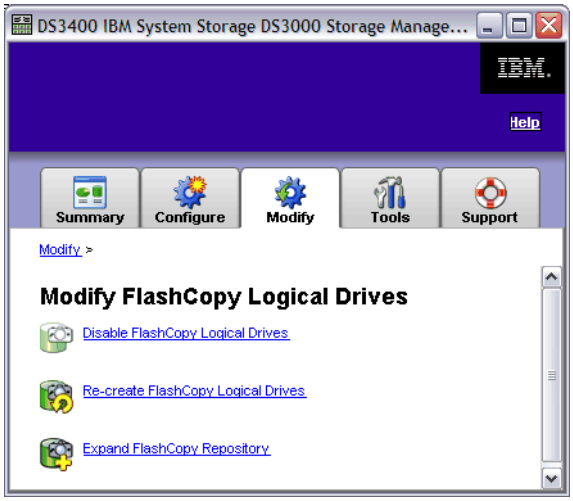


Figure 11-35 Modify FlashCopy Logical Drives

Disable FlashCopy Logical Drives

When you no longer need a FlashCopy logical drive, you can disable it temporarily. Enabled FlashCopy logical drives impact storage system performance by the copy-on-write activity to the associated FlashCopy Repository logical drive. Disabling a FlashCopy logical drive stops the copy-on-write activity.

Disabling the FlashCopy logical drive rather than deleting it will retain the FlashCopy drive and its associated repository. Then, when you need to create a different snapshot data of the same base logical drive, just use the re-create option to reuse the disabled FlashCopy drive. This takes less time than creating a new one.

Note: When you disable a FlashCopy logical drive:

- You cannot use that FlashCopy logical drive again until you use the re-create option on that logical drive. So, the host server that has this FlashCopy assigned as a LUN will lose access to data on it. Be sure you unmount any file system on the FlashCopy drive before you disable it.
- Only that FlashCopy logical drive is disabled. All other FlashCopy logical drives remain functional.

If you do not intend to ever re-create a FlashCopy, you can delete that FlashCopy logical drive (see 11.2.2, “Delete Arrays and Logical Drives” on page 232) instead of disabling it.

To disable a FlashCopy logical drive, follow these steps:

1. Click **Disable FlashCopy Logical Drives** to open the associated window (Figure 11-36).

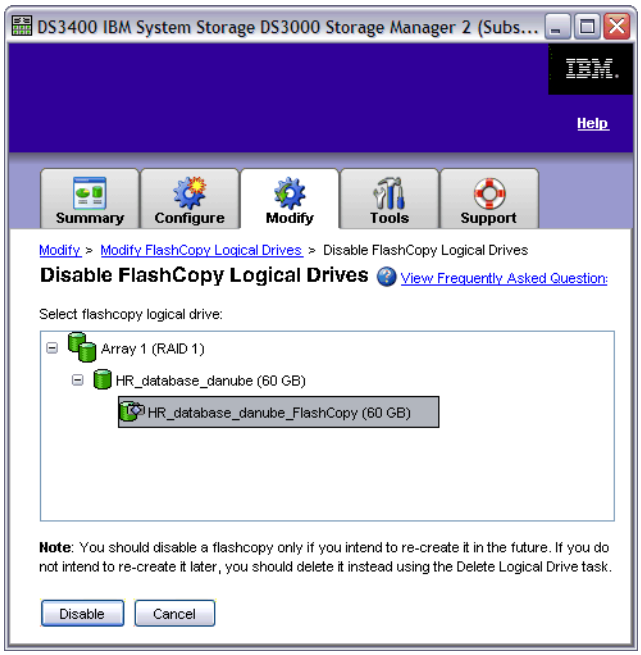


Figure 11-36 Modify - Disable FlashCopy Logical Drives

2. Select the FlashCopy Logical Drive you want to disable and click **Disable**.
3. Review the warning (Figure 11-37 on page 245), confirm it by typing yes, and click **OK**.

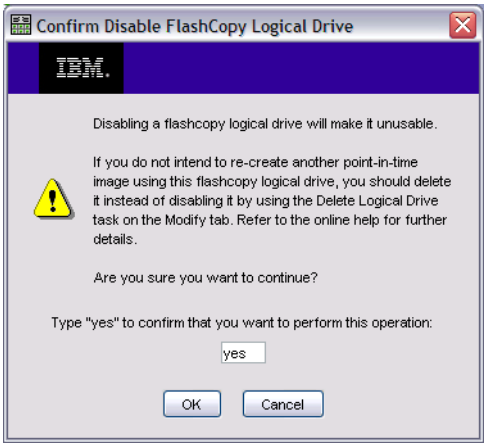


Figure 11-37 Modify - confirm disabling FlashCopy drive

4. When the operation is complete, you will see a message similar to Figure 11-38. Click **OK**.



Figure 11-38 Modify - disable FlashCopy logical drive successful

Re-create a FlashCopy logical drive

Re-creating a FlashCopy logical drive takes less time than creating a new one. If you have a FlashCopy logical drive that you no longer need, instead of deleting it, you can reuse it (and its associated FlashCopy Repository logical drive) to create a different FlashCopy logical drive of the same base logical drive.

When you re-create a FlashCopy logical drive, here are some considerations:

- The FlashCopy logical drive must be either in an optimal or a disabled state.
- All previous copy-on-write data on the FlashCopy Repository logical drive will be deleted.
- The FlashCopy logical drive and FlashCopy Repository logical drive parameters remain the same as the previously disabled FlashCopy logical drive and its associated FlashCopy Repository logical drive. After the FlashCopy logical drive is re-created, you can change parameters on the FlashCopy Repository logical drive through the appropriate menu options.
- The original names for the FlashCopy logical drive and FlashCopy Repository logical drives are retained. You can change these names after the re-create option completes.

- When using this option, the previously configured FlashCopy name, parameters, and FlashCopy repository logical drive are used.

To recreate a FlashCopy drive, follow these steps:

1. Click **Re-create FlashCopy Logical Drives** (Figure 11-35 on page 243) to display the associated view (Figure 11-39).

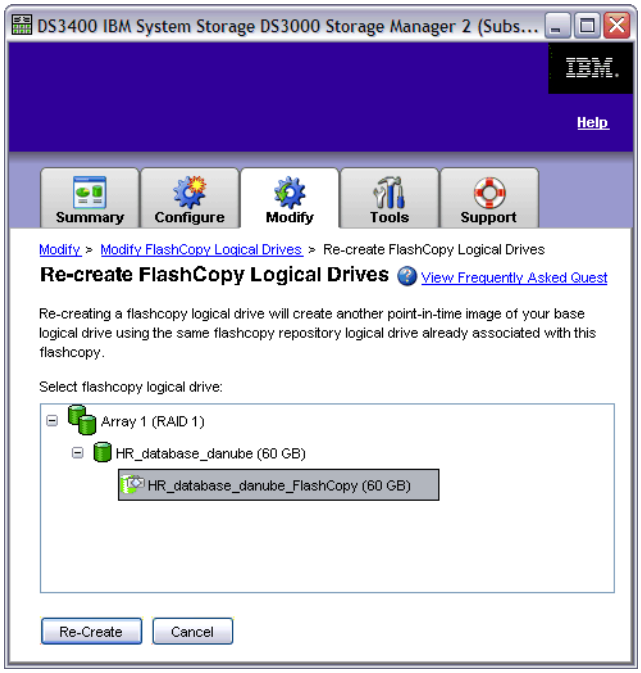


Figure 11-39 Modify - Re-create FlashCopy Logical Drives

2. Select the FlashCopy logical drive you want to re-create and click **Re-Create**.
3. Review the warning (Figure 11-40), confirm it by typing **yes, and** click **OK**.

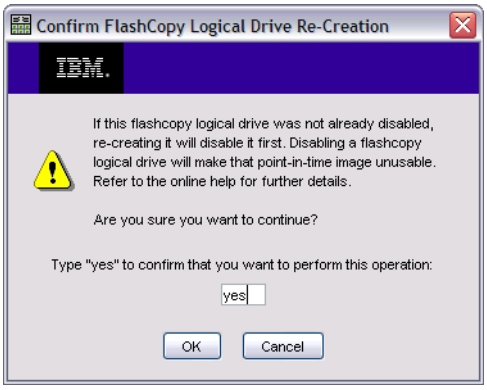


Figure 11-40 Modify - Confirm FlashCopy Logical Drive Re-Creation

4. When the operation is complete, you will see a message similar to Figure 11-41. Click **OK**.

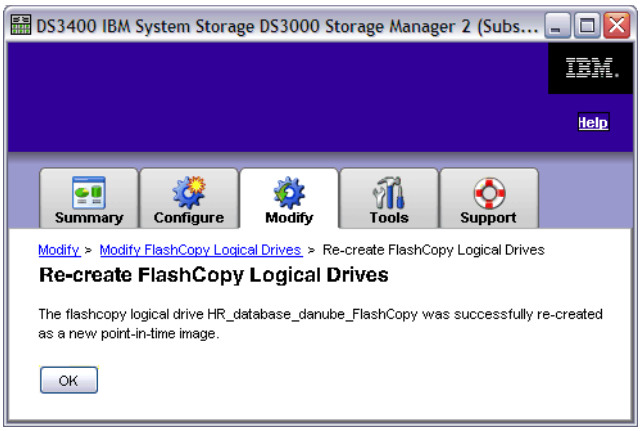


Figure 11-41 Modify - FlashCopy logical drive re-creation successful

Increase the size of a FlashCopy Repository logical drive

Use the Expand FlashCopy Repository option to increase the storage capacity of an existing FlashCopy Repository logical drive. Typically, this option is used when a warning is received that the FlashCopy Repository logical drive is in danger of becoming full.

You can increase the storage capacity by:

- Using free capacity available on the array of the FlashCopy Repository logical drive.
- Adding unconfigured capacity (in the form of unused disk drives) to the array where the FlashCopy Repository logical drive is placed. Use this option when less then needed or no free capacity exists on the array.

Note: A maximum of two drives can be added to the array at the same time to increase FlashCopy Repository logical drive capacity. For RAID 1 or 10, the minimum is two drives. See details about adding disk(s) to the array in 11.2.3, “Add Free Capacity (Drives)” on page 234.

A FlashCopy Repository logical drive storage capacity cannot be increased if:

- The repository logical drive does not have Optimal status.
- Any logical drive in the array is in any state of modification.
- The controller that owns this repository logical drive is in the process of adding capacity to another logical drive (each controller can add capacity to only one logical drive in the array at a time).
- One or more HotSpare drives are in use in the array.

To expand a FlashCopy drive repository, follow these steps:

1. Click **Expand FlashCopy Repository** (Figure 11-35 on page 243) to display the associated view (Figure 11-42).

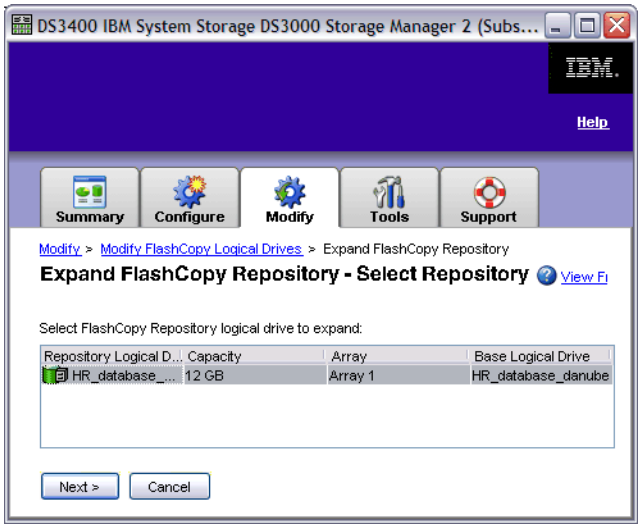


Figure 11-42 Modify - Expand FlashCopy Repository

2. Select the FlashCopy Repository you want to expand and click **Next**.

   The Increase Capacity view appears (Figure 11-43), showing the size of the current Base Logical Drive and of the FlashCopy Repository.

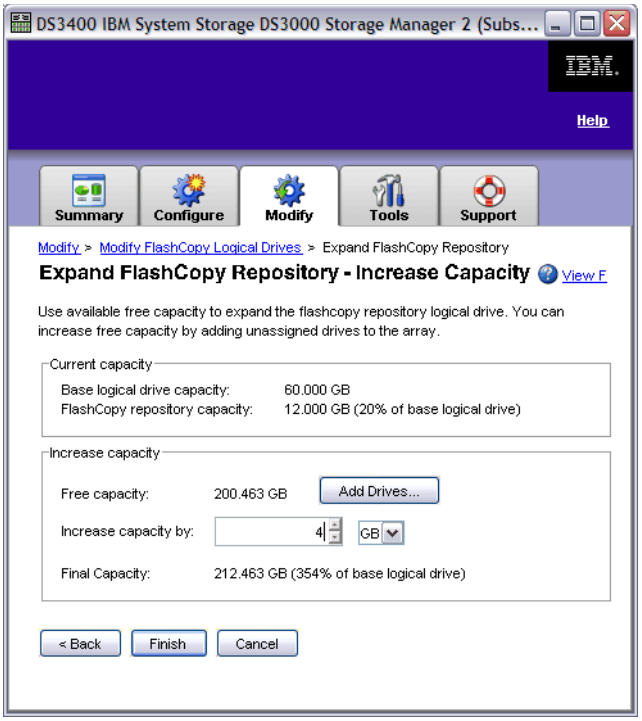


Figure 11-43 Modify - FlashCopy Repository - Increase Capacity

3. If there is free capacity available in this array left, just enter the amount to increase capacity by in the Increase Capacity field and click **Finish**.
4. If there is not enough free capacity left, you can expand it here by adding new drives. To do this, click **Add Drives** to open the “Add Drives” window (Figure 11-44). Select from the drop-down menu how many drives you want to add and click **Add**.

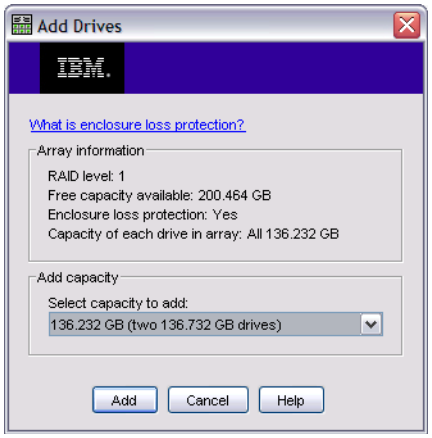


Figure 11-44 Modify - FlashCopy - Add Drives

5. The new free capacity will be immediately available in the “Increase Capacity” part of window (Figure 11-45).

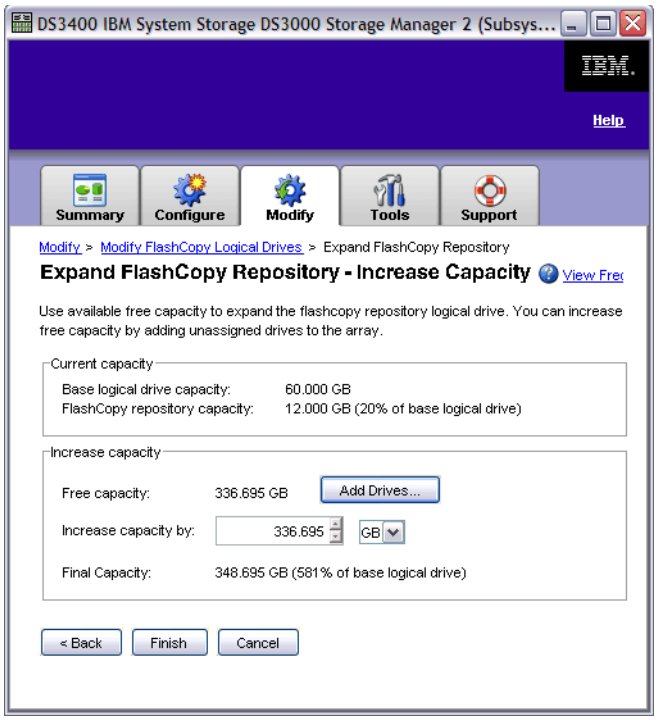


Figure 11-45 Modify - drives added - more free capacity available

6. Type the amount of capacity you want to add to the repository drive, click **Finish**, review the warning window (Figure 11-46), and click **Yes** to continue.

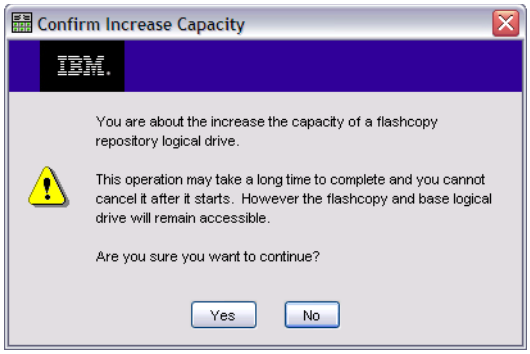


Figure 11-46 Modify - increase FlashCopy Repository drive size confirmation

Note: After you have confirmed the capacity expansion process, the capacity will be added and the FlashCopy Repository drive will be expanded.

This can take some time (a very long time if you add disk drives) and you will not be able to stop this process once you have started it. However, increased capacity of FlashCopy Repository can be used immediately.

11.2.8 Manage Logical Drive Copies

The Copy Manager window (Figure 11-47 on page 251) is used to monitor the progress of a VolumeCopy feature and perform the following tasks for all logical drive copies on the storage subsystem:

- Re-Copy a logical drive.
- Stop VolumeCopy.
- Set the Read/Write permissions for a target logical drive.
- Change VolumeCopy priority.
- Remove VolumeCopy pairs.

Information is displayed in one line for each of the logical drive copies on the storage subsystem, including the source logical drive and the target logical drive (they form a VolumeCopy pair), the status of the VolumeCopy, a time stamp for any completed copies, and the VolumeCopy priority assigned to the VolumeCopy pair. If a logical drive is read-only to hosts, a lock icon is displayed in appropriate column. You can view the progress of a VolumeCopy in the Status column.

The progress of the VolumeCopy is displayed in the status bar. There you see the completed percentage. If you highlight one VolumeCopy pair, and the status is in progress, the estimated remaining time is displayed below the status box.

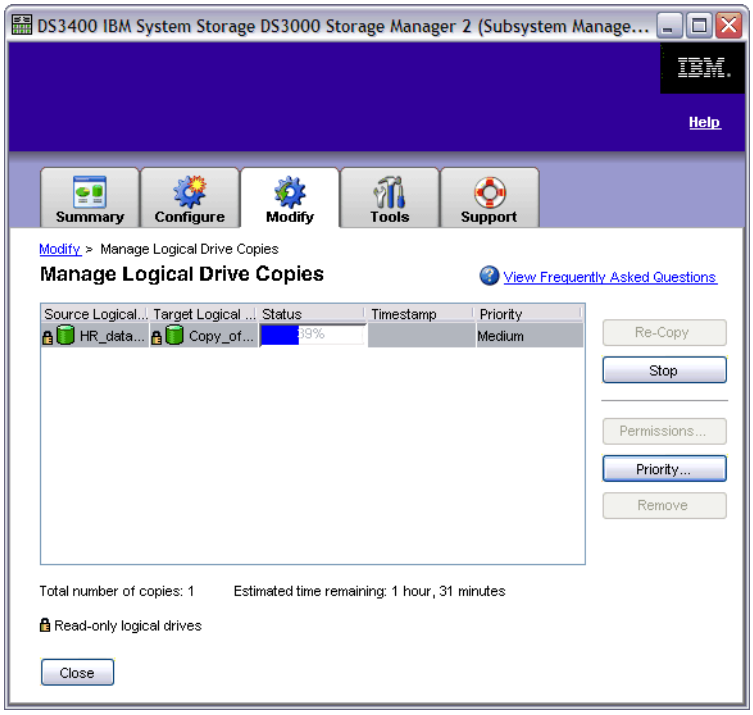


Figure 11-47 Manage Logical Drive Copies

Re-copying a logical drive

Use the Re-Copy option in Copy Manager to create a new (fresh) VolumeCopy for an already selected source and target logical drive (VolumeCopy pair). You can use this option when the status of a VolumeCopy pair is “Completed”, “Stopped”, or “Failed”. The VolumeCopy starts over from the beginning.

Warning: The Re-Copy option:

- Overwrites all existing data on the target logical drive.
- Makes the target logical drive read-only to the host until you set the Read/Write attribute in the Copy Manager when the copy process ends.
- Makes the source logical drive read-only until the copy process ends. Be sure the host that is mapped to the source logical drive does not need to write to the source drive! We recommend unmounting the source drive from OS.

Important: Also, consider the following information:

- To use the Re-Copy option, you can select only one VolumeCopy pair in Copy Manager.
- A logical drive currently in a modification operation cannot be used as a source logical drive or target logical drive.
- A logical drive that has a status of Degraded cannot be used as a target logical drive.
- A logical drive that has a status of Failed cannot be used as a source logical drive or target logical drive.

To re-copy a VolumeCopy, do the following steps:

1. Select a VolumeCopy pair that you want to re-copy and click **Re-Copy**, as shown in Figure 11-47 on page 251.
2. The Re-Copy window opens (Figure 11-48). Review the information, select a copy priority, and type yes to confirm that you want to perform this operation. Click **OK** to start.

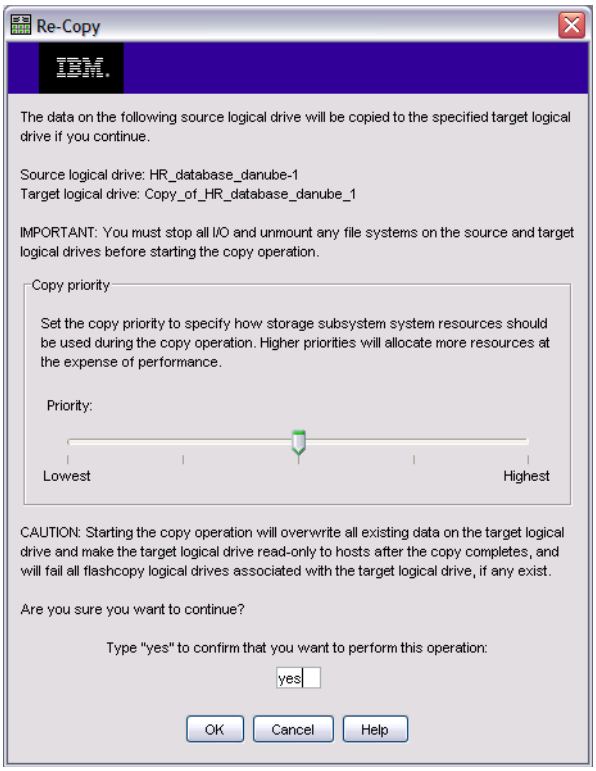


Figure 11-48 Logical Drive Copies - Re-Copy

3. The Re-Copy process will start immediately. You see the progress in the Copy Manager window.

Stopping a VolumeCopy

This option in the Copy Manager is used to manually stop a copy process in progress or if the VolumeCopy pair has a status or “Pending” or “Failed”. The status changes to “Stopped”.

Note: Using this option on a VolumeCopy pair with a status of “Failed” also clears the Needs-Attention condition on the affected DS3000 Storage Subsystem.

To stop a VolumeCopy pair:

1. Select a VolumeCopy pair and click **Stop**.
2. The Stop Copy window appears (Figure 11-49).

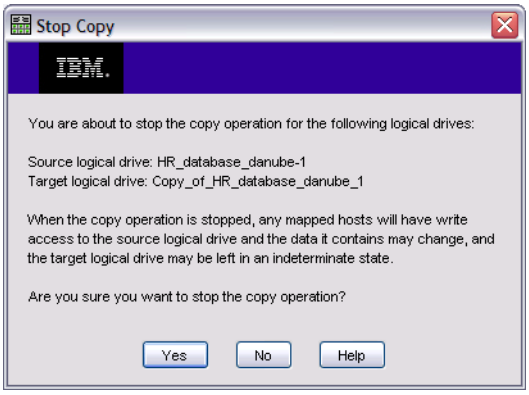


Figure 11-49 Stop VolumeCopy

3. Read the information and click **Yes.** The VolumeCopy process stops immediately.

Permissions

Read and write requests to the target logical drive are not possible while the VolumeCopy pair has a status of “Pending”, copying is in progress, or if the VolumeCopy operation fails before completing. Note that the source logical drive is read-only at this time.

After the VolumeCopy completes (Figure 11-50) or is stopped manually, the target logical drive automatically becomes read-only to hosts and a lock icon is displayed in the “Target Logical Drive” column. Note that the source drive starts to be read/write accessible again.

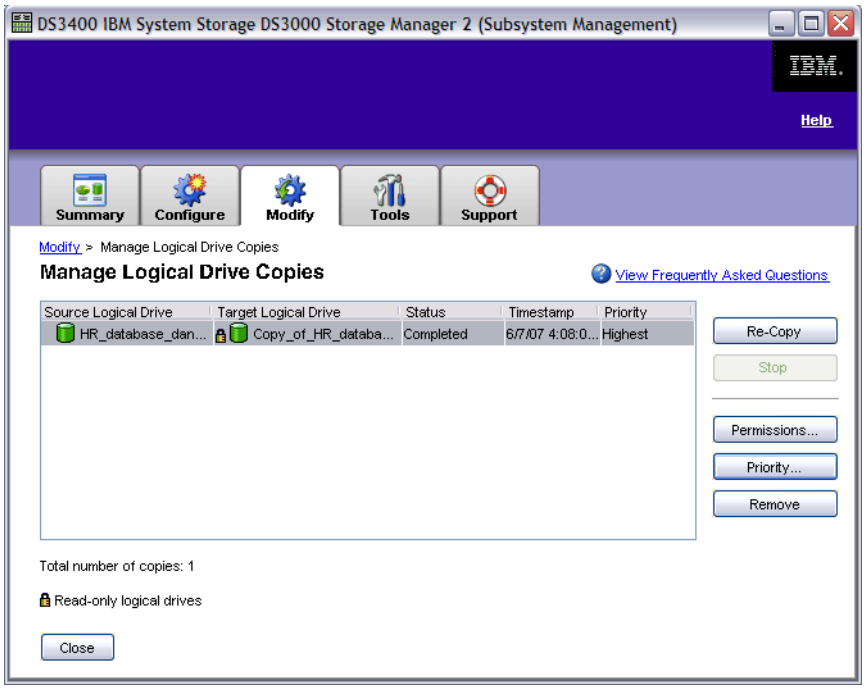


Figure 11-50 VolumeCopy - completed

You may need to keep the read-only attribute enabled in order to preserve the data on the target logical drive. For example, if you can keep the target logical drive for backup purposes, if you are copying data from one array to a larger array for greater accessibility, or if you are using the data on the target logical drive to copy back to the base logical drive of a disabled or failed FlashCopy logical drive.

If you decide not to preserve the data on the target logical drive after the VolumeCopy is completed, use Copy Manager to disable the Read-Only attribute for the target logical drive.

To disable the Read-Only attribute:

1. Select a VolumeCopy pair in the window, as shown in Figure 11-47 on page 251, and click the **Permissions** button to display the Set Permissions Window (Figure 11-51).

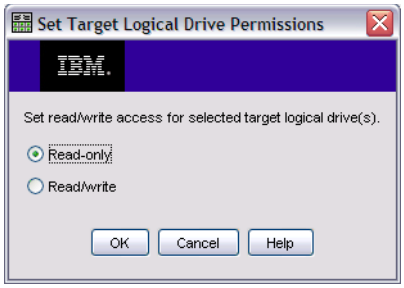


Figure 11-51 Set Permissions

2. Select either Read-only or Read/Write and click **OK**.
3. The progress window shows you when the change is completed (Figure 11-52).

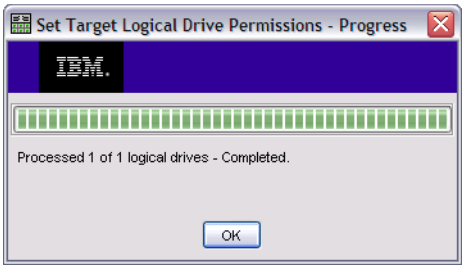


Figure 11-52 Change Permissions Completed

You can set the volume back to Read-only using the same procedure.

Change VolumeCopy priority

Use this option in the Copy Manager to select the rate at which a VolumeCopy completes for a selected VolumeCopy pair. You can change the copy priority for a copy pair before the VolumeCopy begins, while the VolumeCopy process is in progress, or after the VolumeCopy has completed (for a Re-Copy).

To change the priority, select the VolumeCopy pair and click **Priority**. The “Change Priority” window appears (Figure 11-53).

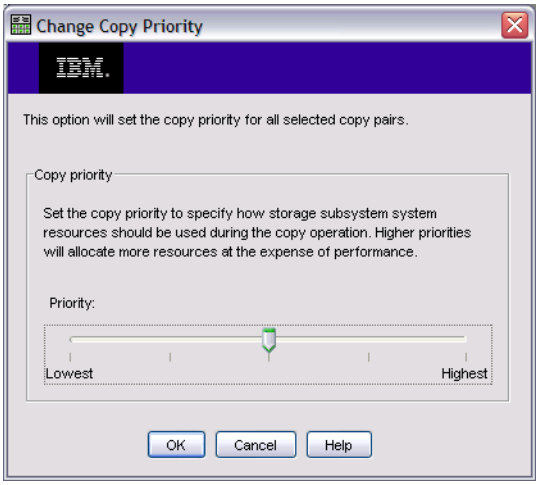


Figure 11-53 Volume Copy - Change copy Priority

Review the information, set a new priority, and click **OK.** You will be informed when the priority change is done (Figure 11-54). Click **OK** to accept it.

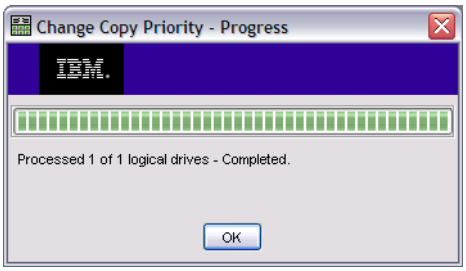


Figure 11-54 VolumeCopy - priority change complete

Remove VolumeCopy pairs

Use this option to remove one or more logical drive copies from Copy Manager. After the VolumeCopy is removed from Copy Manager, the target logical drive can be selected as a source logical drive or target logical drive for a new VolumeCopy.

If you remove a VolumeCopy pair, the source logical drive and target logical drive are no longer displayed in Copy Manager.

Important: Consider the following information:

- This option does not delete the data on the source logical drive or target logical drive.
- If the VolumeCopy process is in progress, it must be stopped before you can remove the VolumeCopy pair from the Copy Manager.

The procedure for removing VolumeCopy pairs is as follows:

1. Select the copy pair(s) and click **Remove**. You can use the Ctrl or Shift key for multiple selections. The Remove Copy Pairs confirmation window appears (Figure 11-55). Read the information and click **Yes**.

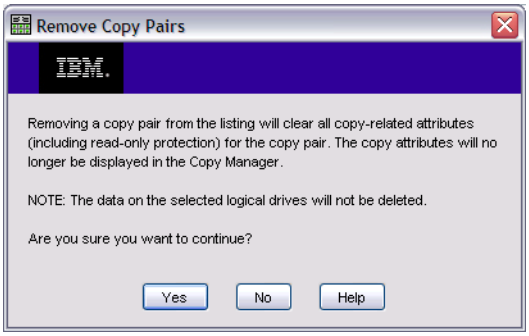


Figure 11-55 Remove Copy Pairs

2. The progress window appears (Figure 11-56). When the process is complete, click **OK**.

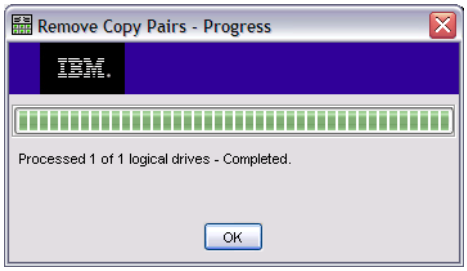


Figure 11-56 Remove Copy Pairs - completed

12

Administration - Tools

In this chapter, we describe the Tools tab of the IBM System Storage DS3000 Storage Manager.

We cover the following topics in four sections:

- Storage Subsystem
  - Naming of storage subsystem
  - Security of storage subsystem
  - Locating components
  - Managing Premium Features
  - Managing enclosure IDs
- Logical drives
  - Media Scan settings
  - Pre-Read Redundancy Check settings
- Controllers
  - Controller clocks synchronization
  - Ethernet management ports configuration
  - Battery settings
- Accessibility
  - Inheriting of system settings

© Copyright IBM Corp. 2007, 2009. All rights reserved.

To access the Storage Subsystem Management Tools, click the **Tools** tab in DS3000 Storage Manager (Figure 12-1).



Figure 12-1 The Storage Subsystem Management Tools Tab

12.1 Storage Subsystem tools

The Storage Subsystem section of the Tools tab includes the following wizards:

- Rename Storage Subsystem
- Set or Change Password
- Locate
- View/Enable Premium Features
- Change Enclosure ID Numbers

12.1.1 Rename Storage Subsystem

If you have several storage subsystems installed in your environment, you should give each subsystem a meaningful name, so that you can distinguish between the systems. When you first add the storage subsystem to your Storage Manager, it is identified as “Unnamed”. You can assign a name to it during the initial setup (see 5.4.2, “Renaming the storage subsystem” on page 72), or later using this option.

If you are not sure which name is assigned to a specific storage subsystem, see 5.4.1, “Locating the storage subsystem” on page 72 or “Locate Storage Subsystem or Enclosures” on page 263 for instructions.

Rules of naming

Use these rules when assigning names to your disk subsystems:

- Names can contain a maximum of 30 characters.
- Names can include letters, numbers, and the special characters underscore (_), minus (-), and pound/number sign (#). No other special characters are permitted.
- Use a clear meaningful name that is easy to understand and to remember. Avoid arbitrary names or names that lose their meaning in the future.

Rename subsystem

Follow this procedure to rename the subsystem:

1. Click **Rename Storage Subsystem** (Figure 12-1 on page 258) to open the associated view (Figure 12-2).

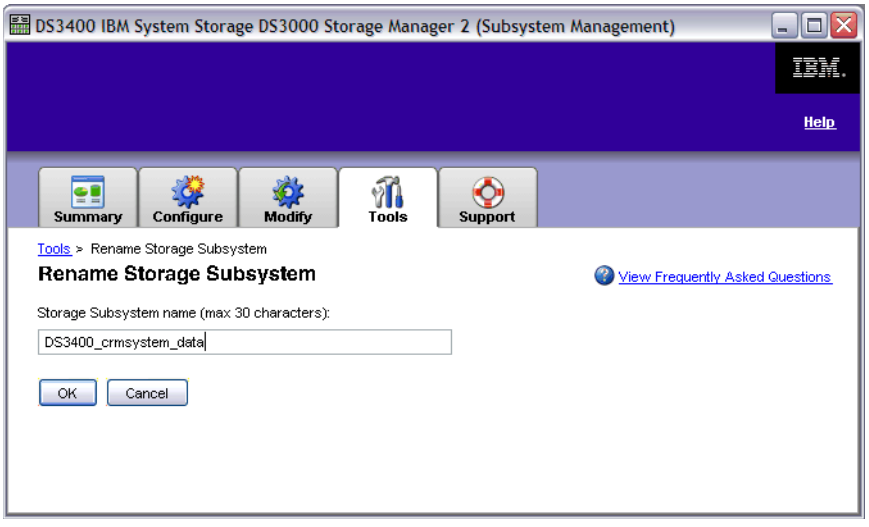


Figure 12-2 Rename Storage Subsystem

2. Type the new storage subsystem name into the name field and click **OK**.
3. Read the advice in the window (Figure 12-3), and click **Yes** to confirm. Renaming the storage subsystem will also influence, for example, some customized SMcli scripts, where the subsystem is referenced by its name.

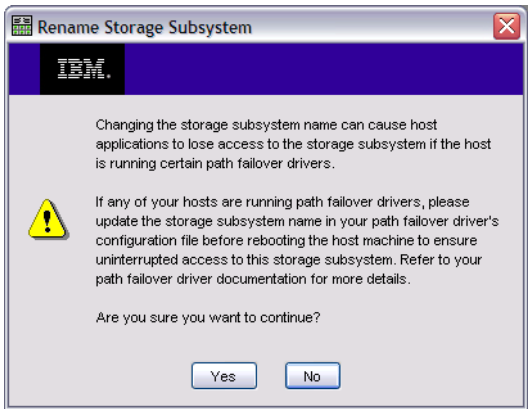


Figure 12-3 Rename Storage Subsystem Information

4. When the renaming process is complete, you will see a confirmation similar to Figure 12-4. Click **OK** to return to the Tools tab.

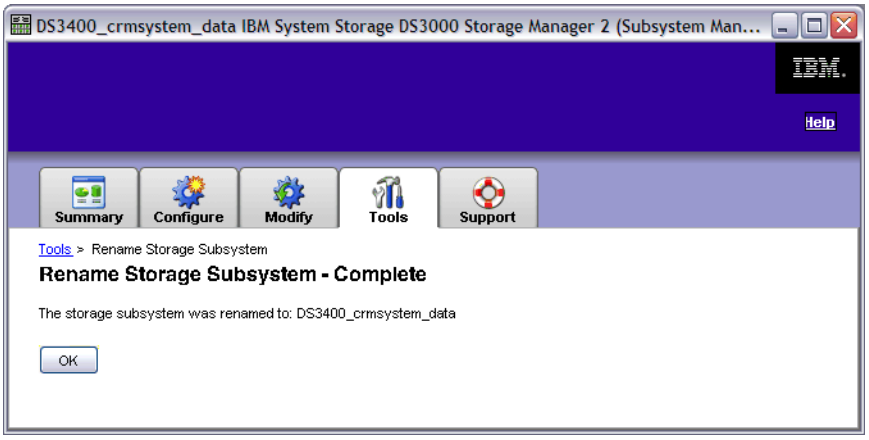


Figure 12-4 Rename Storage Subsystem - Complete

12.1.2 Set or Change Password

You can configure each storage subsystem with a password to protect it from unauthorized access. If password protection is enabled, you will be prompted for the password when changing the configuration, such as when you create or delete a logical drive. The password is not required again if you do more configuration changes in the same session. View operations do not require a password to be entered. To disable password protection, follow the process described here. Enter the current password and leave the New password and Confirm password text boxes blank.

Note: If you forget your password, contact IBM Technical Support. Password reset is possible with direct DS3000 Storage Subsystem HW unit attachment only.

Remember that:

- Passwords are case sensitive.
- The maximum length is 30 characters.
- Spaces included in a password are saved.
- For security reasons, you can attempt to enter a password ten times before the DS3000 storage subsystem enters a “lockout” state. Before you can try to enter a password again, you must wait ten minutes for the “lockout” state to be reset.

Set password

Follow this procedure:

1. Click **Set or Change Password** (Figure 12-1 on page 258) to open the view (Figure 12-5).

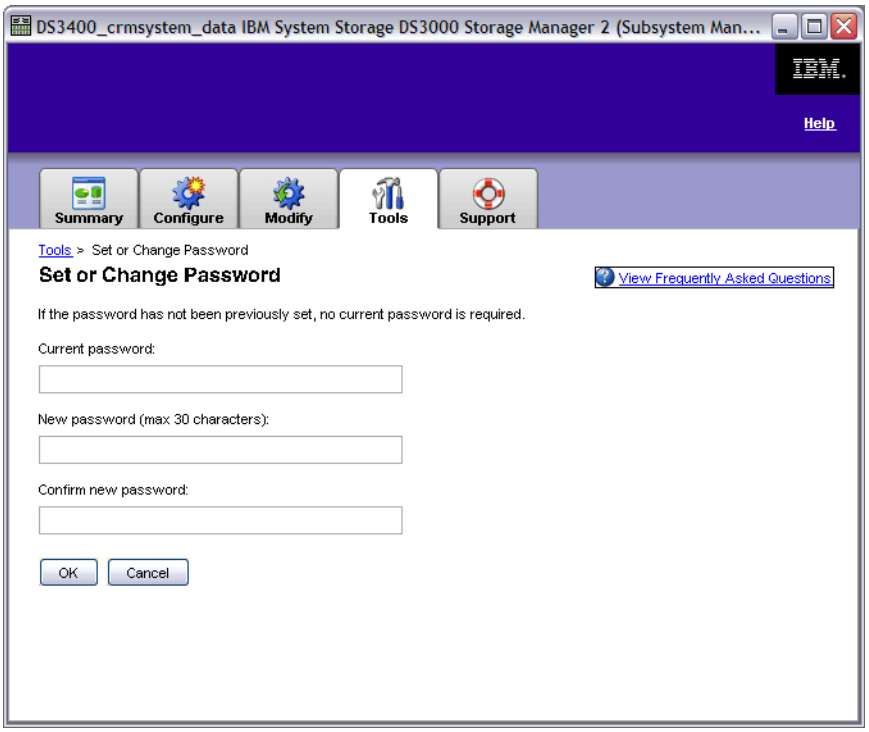


Figure 12-5 Set or Change Password

2. Enter the current password in the first text box; if there is no current password, leave it blank.
3. Enter the new password in the appropriate field.
4. Re-enter the new password in the Confirm new password field to confirm it and click **OK**.

5. The DS3000 password is now set (Figure 12-6). Click **OK** to return to the Tools tab.

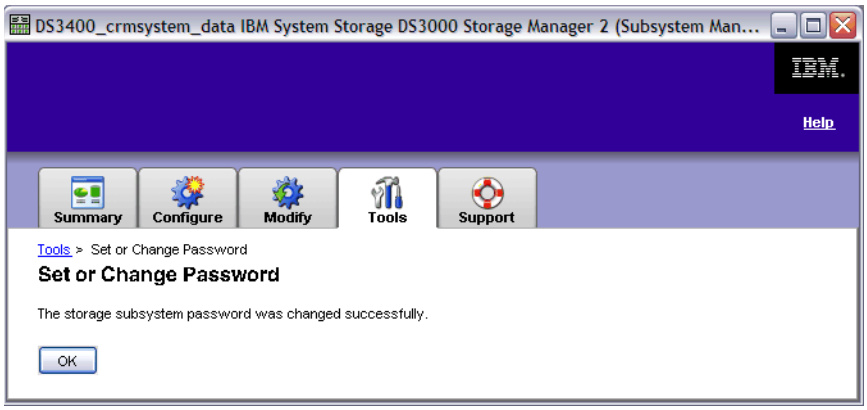


Figure 12-6 Set or Change Password successful

Verify password change

To verify that your password change is working, close all Storage Manager windows, including the Enterprise Management Window. Re-start Storage Manager, and try, for example, to rename a logical drive. Now you will be prompted for the password (Figure 12-7) before you can run this task.



Figure 12-7 Enter Password

12.1.3 Locate

This section turns on the indicator LED lights for various components in the storage subsystem to identify them physically in the rack or in your server room.

If you have several DS3000 storage subsystems, the Locate function is very helpful in:

- Locating a storage subsystem or enclosure that has failed
- Labeling the storage subsystem or the enclosures with the name used in the storage management software
- Locating and labeling the disk drives of the array that you plan to migrate

The Locate function works with these storage subsystem components:

- Disk Drive(s)
- Array
- EXP Enclosure
- Controller
- Storage Subsystem

Click **Locate** (Figure 12-1 on page 258) to open the Locate window (Figure 12-8 on page 263).

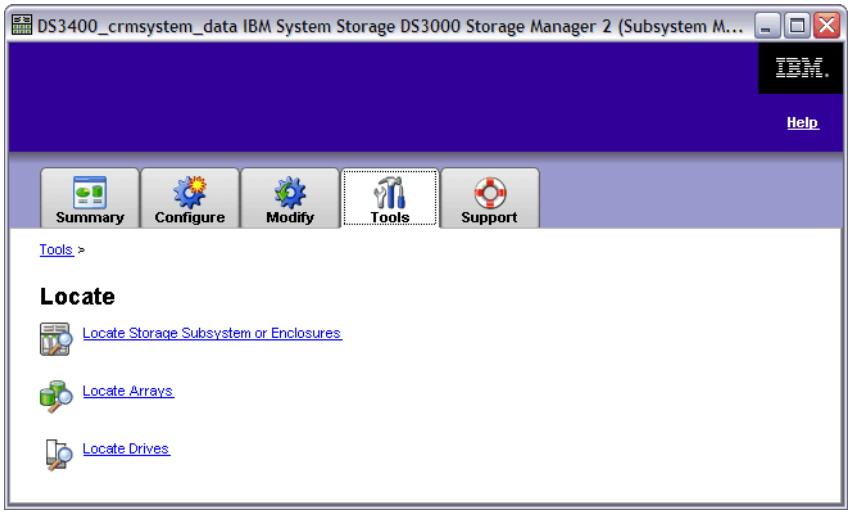


Figure 12-8 Locate

Locate Storage Subsystem or Enclosures

Click **Locate Storage Subsystem or Enclosures** and a window similar to Figure 12-9 appears. Select the subsystem or enclosure you want to locate and click **Locate**.

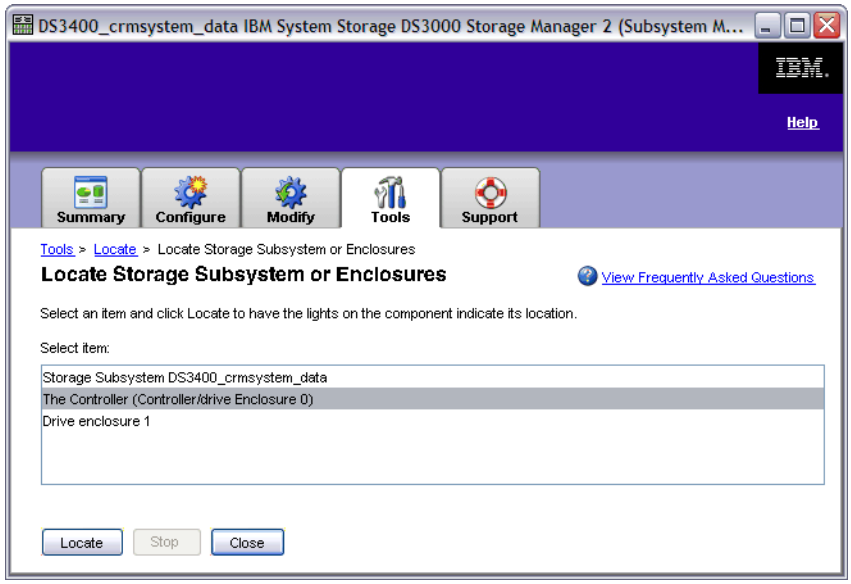


Figure 12-9 Locate Storage Subsystem - start

The LED lights on the selected component will start flashing. After you have located it, click **Stop** (Figure 12-10) to disable the flashing and click **Close** to return to the Locate view (Figure 12-8 on page 263).

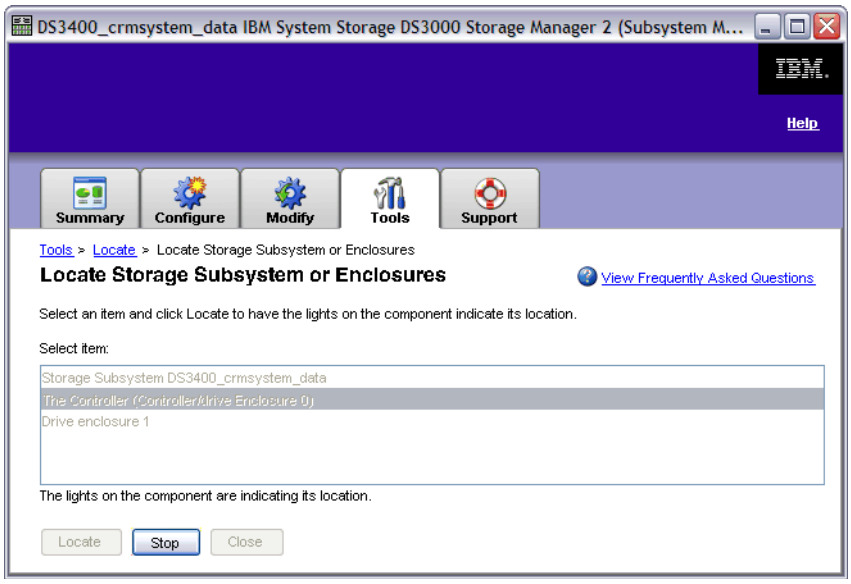


Figure 12-10 Locate Storage Subsystem or Enclosures - stop

Locate Arrays

Click **Locate Arrays** to open the **Locate Arrays view** (Figure 12-11). Select the array you want to locate and click **Locate**. The lights on the disk drives associated with this array will start flashing. Click **Stop** to disable the flashing and **Close** to return (Figure 12-8 on page 263).

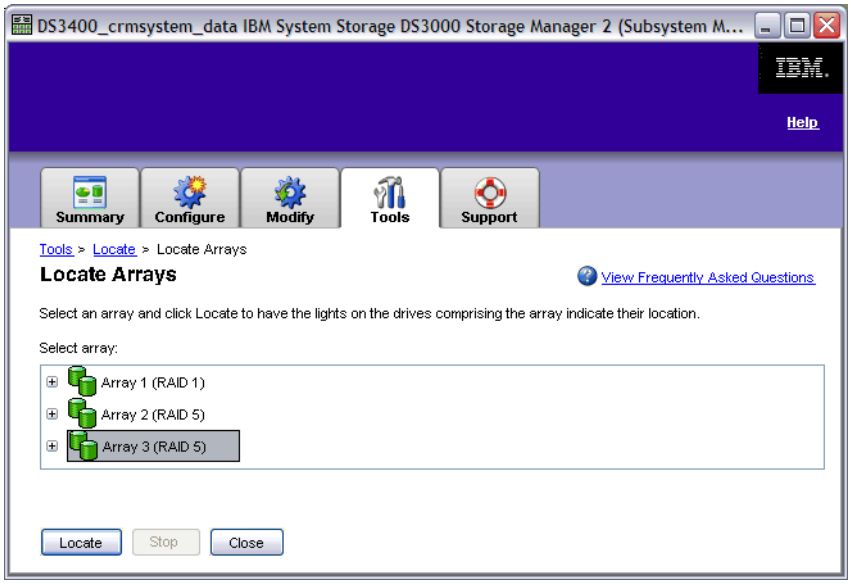


Figure 12-11 Locate Arrays

Locate Drives

Click **Locate Drives** to display the window shown in Figure 12-12. Select the drive(s) you want to locate and click **Locate**. The drive lights will start flashing. Click **Stop** and **Close** to disable the flashing and return to the Locate view (Figure 12-8 on page 263).

Tip: To highlight several drives in one step, use the Ctrl or Shift key while selecting/deselecting them.

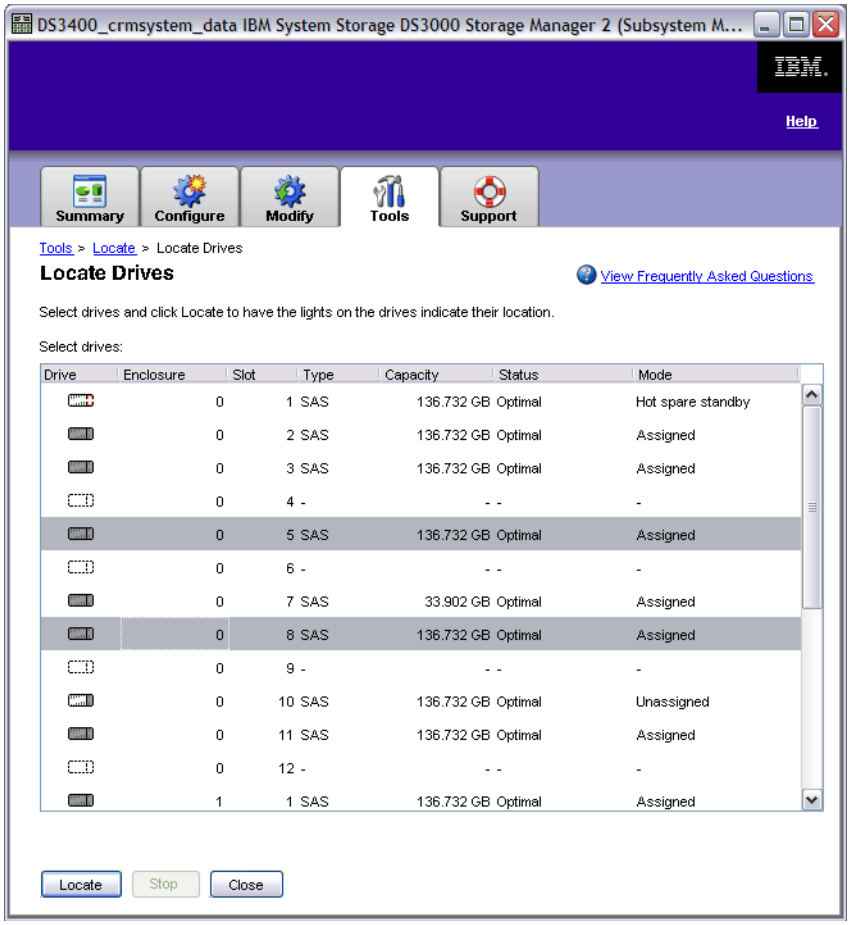


Figure 12-12 Locate Drives

12.1.4 View/Enable Premium Features

As discussed in 3.3, “Premium Features” on page 25, you can license your DS3000 storage subsystems with different optional Premium Features. In this part of the Tools tab, you can view or modify the status of Premium Features and also enable them if it is not already done.

Premium Features - obtaining the license key file

You need a special license *.key file to activate each Premium Feature. This file must be downloaded from the IBM feature activation Web site (Figure 12-13 on page 266), where each feature is also activated based on the serial number of your DS3000:

http://www-912.ibm.com/PremiumFeatures/

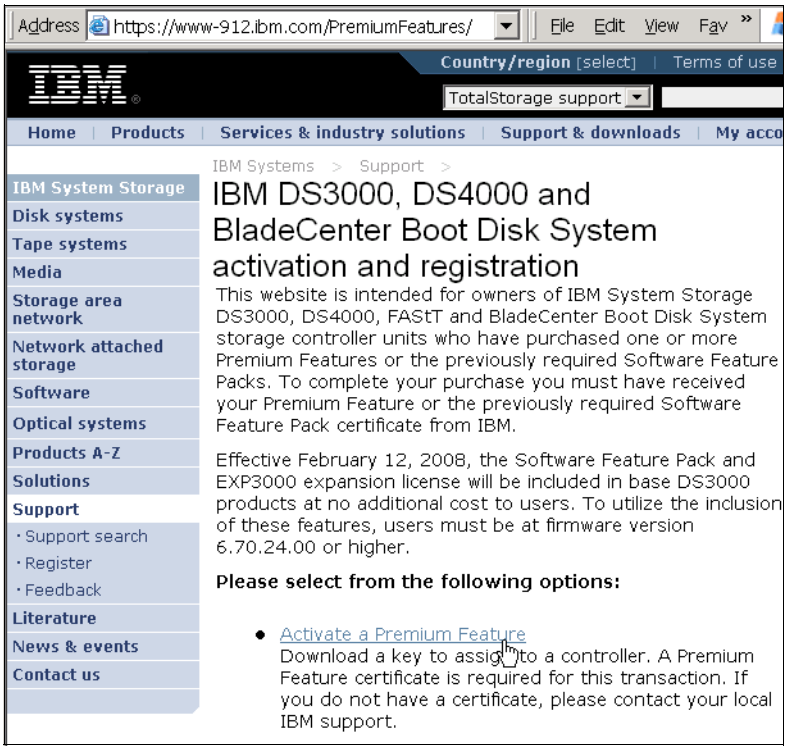


Figure 12-13 Premium Feature activation site

If your DS3000 Storage Subsystem is delivered with optional (additional charge) Premium Feature (more than two FlashCopies, VolumeCopy, or more than four Storage Partitions), you can find instructions and the actual procedure to perform to obtain the license key file in the Premium Feature certificate, which is included in the license envelope and is part of your DS3000 Premium Feature delivery. You can also find a unique access code for this IBM Web site, where the*.key file must be generated and downloaded.

Warning: Store the document from the license envelope in safe place if you want to postpone your Premium Feature key file download. If you lose the access code for the IBM activation Web site, it will be difficult to get it again from IBM in a short time frame. You will need a proof of purchase and need to contact an authorized IBM Support representative.

Note: If you lose the downloaded *.key file, you can download it again, without the access code from the license envelope, at the same IBM Web site.

Premium Feature - Enable Software Feature Pack

Note: Starting February 2008, the DS3000 Software Feature Pack has been incorporated into the base models and is no longer needed as prerequisite for the other advanced function options. The required controller firmware that enables this function must be at firmware level V06.70.23.00 or later.

If you are using an older system running an older firmware level, the Software Feature Pack has to be enabled before you can use any Premium Feature functions like FlashCopy, VolumeCopy, or Storage Partitioning.

If there are no Feature Packs enabled on your DS3000, the Enable Feature Pack link is displayed in the Tools tab (Figure 12-14). Enable the Software Feature Pack first, as it is a prerequisite for all other Premium Features.



Figure 12-14 The Tools tab - Enable Feature Pack

Warning: Installing the Software Feature Pack requires a reboot of the storage subsystem.

To enable the Software Feature Pack, do the following steps:

1. Click **Enable Feature Pack** to display the Select Feature Pack Key file window, as shown in Figure 12-15. Navigate to the key file you received from IBM and click **OK** to install it.

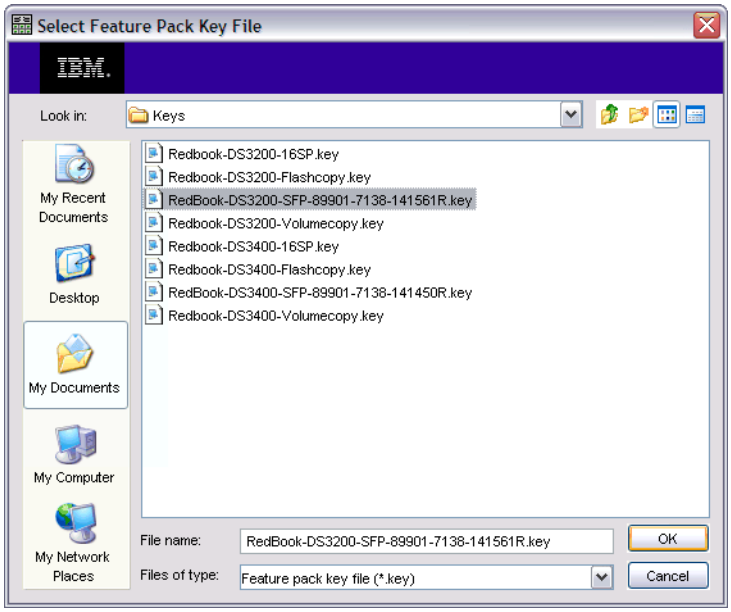


Figure 12-15 Select Feature Pack Key File

2. In the Confirm Enable Feature Pack window, read the information and click **Yes** to confirm it (Figure 12-16).

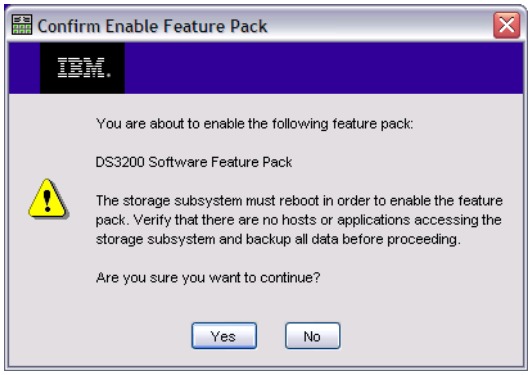


Figure 12-16 Confirm Enable Feature Pack

3. The key file is committed to the system. You will see a confirmation message similar to Figure 12-17. Click **Close** to go back to the Tools tab. The DS3000 system will reboot.

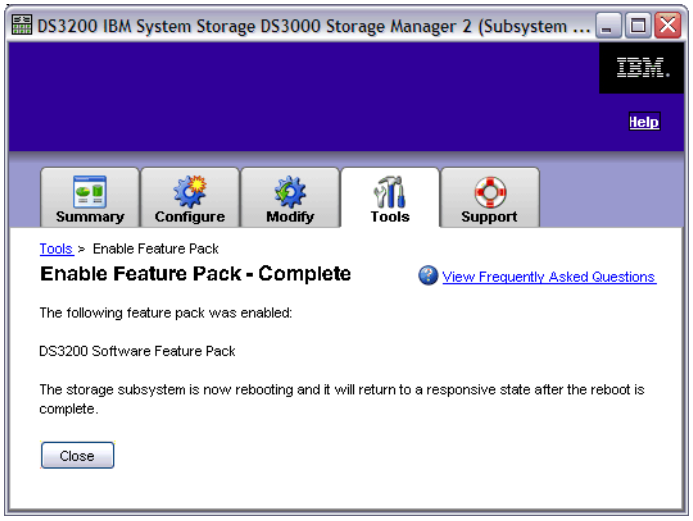


Figure 12-17 Enable Feature Pack - Complete

If you have already installed a Feature Pack, the link Enable Feature Pack in the Tools tab (Figure 12-1 on page 258 or Figure 12-14 on page 267) changes to View/Enable Premium Feature.

Premium Feature - View/Enable Premium Feature

Click **View/Enable Premium Feature** in the Tools tab to display the window shown in Figure 12-18.

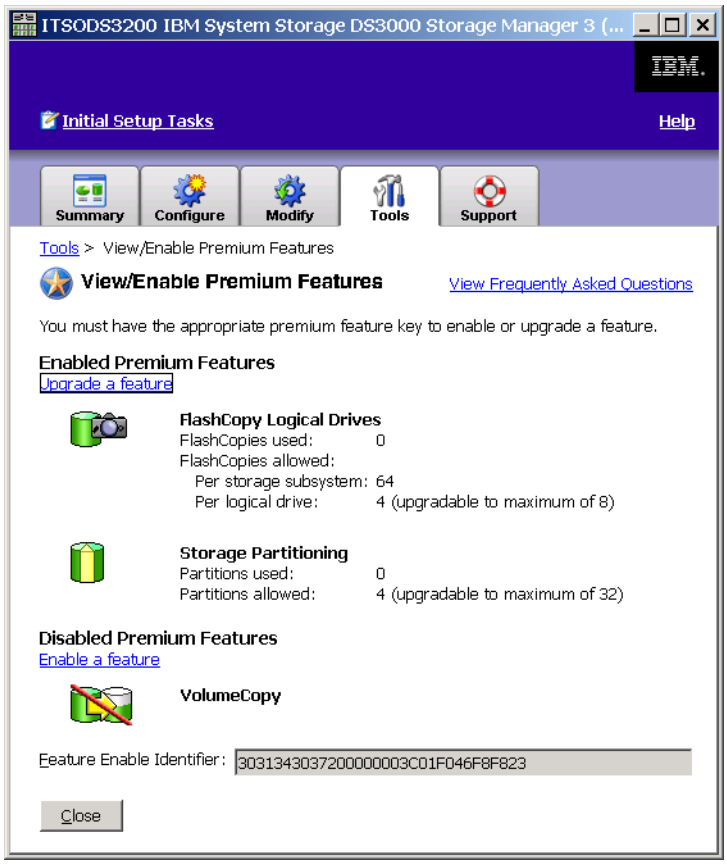


Figure 12-18 View/Enable Premium Features

This window displays the installed Premium Features licenses and the numbers of objects that are in use by them. To enable additional Premium Features, click **Enable a feature**. To modify the licensed feature, click **Upgrade a feature**. Enter the license key file, as shown in Figure 12-19 on page 271, confirm the key, and repeat the step for each additional Premium Feature you want to enable.

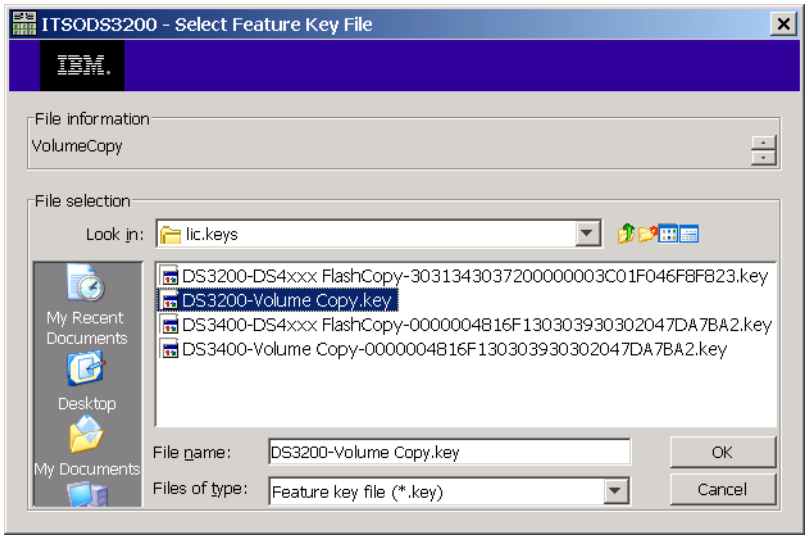


Figure 12-19 Select Feature Key File

Figure 12-20 shows a DS3000 with all basic Premium Features enabled. If you want more FlashCopies or more Storage Partitions, another Premium Feature must be ordered and activated. Click **Close** to return to the Tools tab.

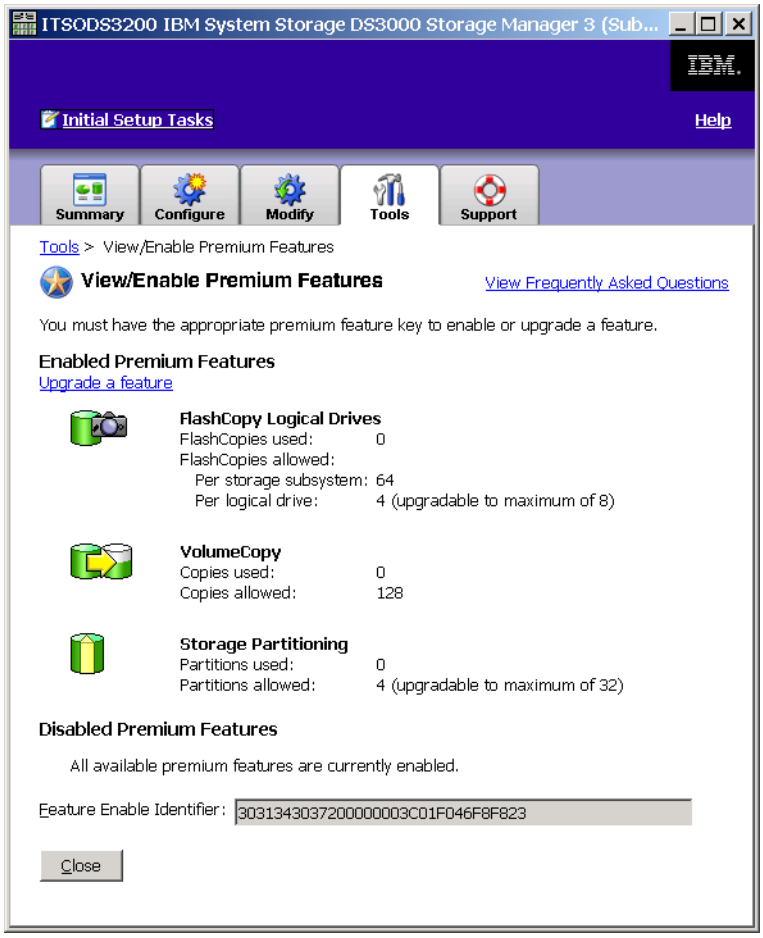


Figure 12-20 View/Enable Premium Features - all features enabled

12.1.5 Change Enclosure ID Numbers

Each DS3000 Enclosure (including the Drive and Controller Enclosures) has to have a unique Enclosure ID, as described in 4.1, “Enclosure ID settings” on page 36. In this section, you can set these IDs.

To do this, click **Change Enclosure ID Numbers**. A window similar to Figure 12-21 appears.

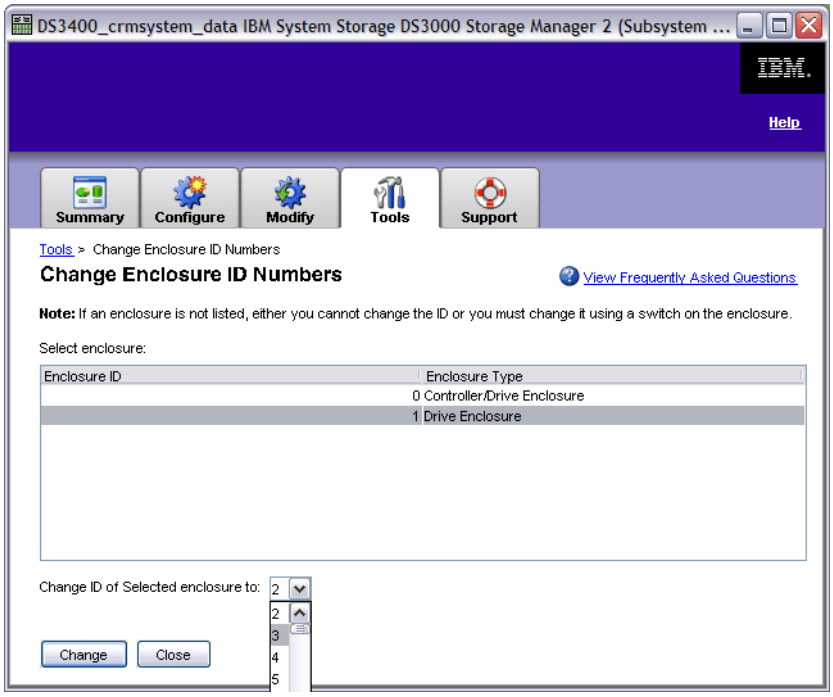


Figure 12-21 Change Enclosure ID Numbers

Highlight an enclosure and select the ID number you want to assign it from the Change ID of Selected enclosure to drop-down menu, and click **Change**.

12.2 Logical drive tools

Here we discuss logical drive tools.

12.2.1 Change Media Scan Settings

Media scan is a background process that checks the physical disks for defects by reading the raw data from the disk and writing it back. This detects possible problems caused by bad sectors of the physical disks before they disrupt normal data reads or writes. This process is sometimes known as *data scrubbing*.

A media scan runs continuously in the background, using spare cycles of controllers to complete its work. The default is to run a media scan every 30 days, which is the maximum time a media scan will have to complete the task. During the scan process, the DS3000 calculates how much longer the scan process will take to complete, and adjusts the priority of the scan to ensure that the scan completes within the time setting allocated. Once the media scan has completed, it will start over and reset its time for completion to the current setting. This media scan setting can be reduced; however, if the setting is too low, priority will be given

to the media scan over host activity to ensure that the scan completes in the allocated time. This scan can impact performance, but improves data integrity.

A media scan should be enabled for the entire storage subsystem. This system wide enabling specifies the duration over which the media scan will run. The logical drive enabling specifies whether or not to do a redundancy check as well as media scan.

A media scan can be considered a surface scan of the hard drives, while a redundancy check scans the blocks of a RAID 3, RAID 5, or RAID 6 logical drive and compares it against the calculated redundancy data. In the case of a RAID 1 or RAID10 logical drive, the redundancy scan compares blocks between copies on mirrored drives.

The media scan function is designed to minimize its impact on other system I/O; in most cases, no effect on I/O should be observed with a 30 day setting unless the processor is utilized at more than 95%. The length of time that it will take to scan the LUNs depends on the capacity of all the LUNs on the system and the utilization of the controller. Table 12-1 shows the possible errors and describes some of the actions that the DS3000 will take as a result of media scan and redundancy check operations. The media scan reports any errors that are found to the event log.

Table 12-1 Media scan errors

Reported error	Description	Result
Unrecovered media error	The data could not be read on its first attempt, or on any subsequent retries.	With redundancy check: Data is reconstructed and scanned again. Without redundancy check: No error correction.
Recovered media error	The drive could not read the requested data on its first attempt, but succeeded on a subsequent attempt.	Data is written to drive and verified.
Redundancy mismatches	Redundancy errors are found.	The first 10 redundancy mismatches found on a logical drive are reported. Operating system data checking operations should be executed.
Unfixable error	The data could not be read, and parity or redundancy information could not be used to regenerate it.	An error is reported.

A media scan is a long running process with the lowest priority, so you can only run a media scan on logical drives that meet the following conditions:

- Optimal status
- No modification operation in progress

Change Media Scan Settings

Click **Change Media Scan Settings** (Figure 12-1 on page 258) to open the associated view (Figure 12-22).

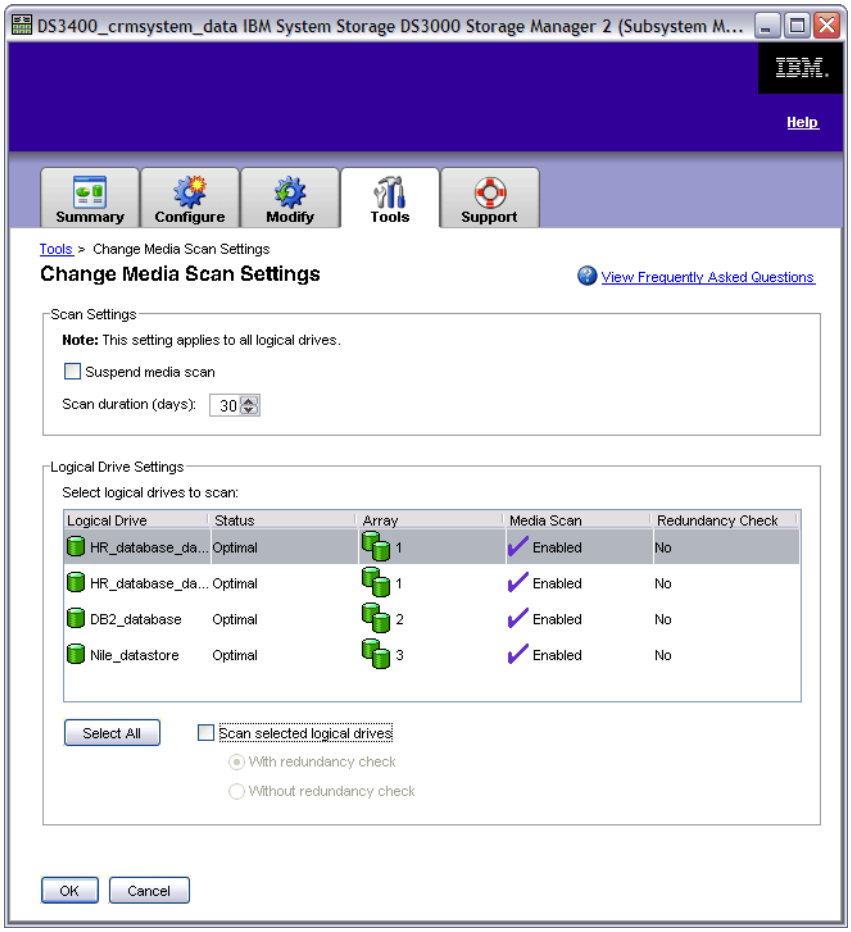


Figure 12-22 Media Scan Settings - standard settings

The Change Media Scan Settings window has two sections. In the first, you can suspend media scan or set the scan duration, and in the second, you can enable or disable scanning for selected logical drives. You can decide if it will be done with or without the redundancy check.

Suspend Media Scan

To disable the media scan, select the **Suspend Media Scan** check box and click **OK**. All media scan functions will be disabled for all logical drives.

Edit scan duration days

To edit the scan duration, make sure **Suspend Media Scan** is unchecked and select the new value in the **Scan duration** box. Click **OK** to save the changes.

Logical drive settings

If the media scan is active, you can enable a redundancy check for individual logical drives:

1. Highlight the logical drive(s) you want to set (Figure 12-22).
2. To enable the media scan, select the **Scan selected logical drives** check box, choose with or without redundancy check, and click **OK**.

3. The changes will be saved on the storage subsystem, and you are informed (Figure 12-23) when it is completed. Click **OK** to return to the Tools tab.

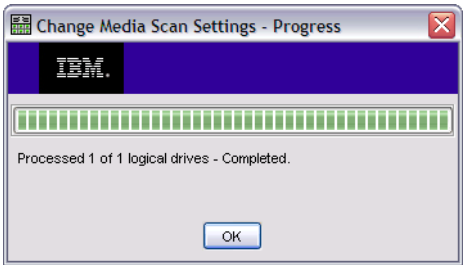


Figure 12-23 Change Media Scan Settings completed

Note: To edit the scan settings for multiple drives, select them with the Ctrl or Shift key. To edit all drives, click **Select All**.

Figure 12-24 shows a Change Media Scan Settings window with customized scan settings. Compare it with Figure 12-22 on page 274: In the Select logical drives to scan box, you can see the current media scan configuration for every logical drive.

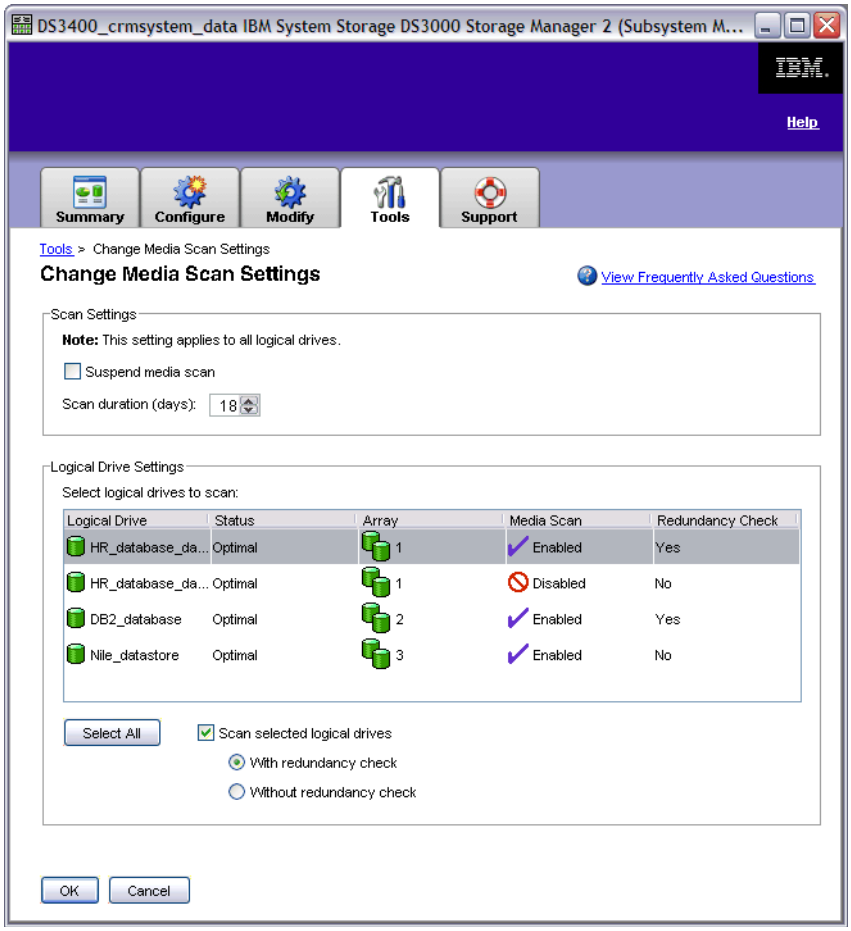


Figure 12-24 Customized media scan settings

Change Pre-Read Redundancy Check Settings

Do the following steps:

1. Click **Change Pre-Read Redundancy Check Settings** in the Tools tab (Figure 12-1 on page 258) to open the associated window (Figure 12-25).

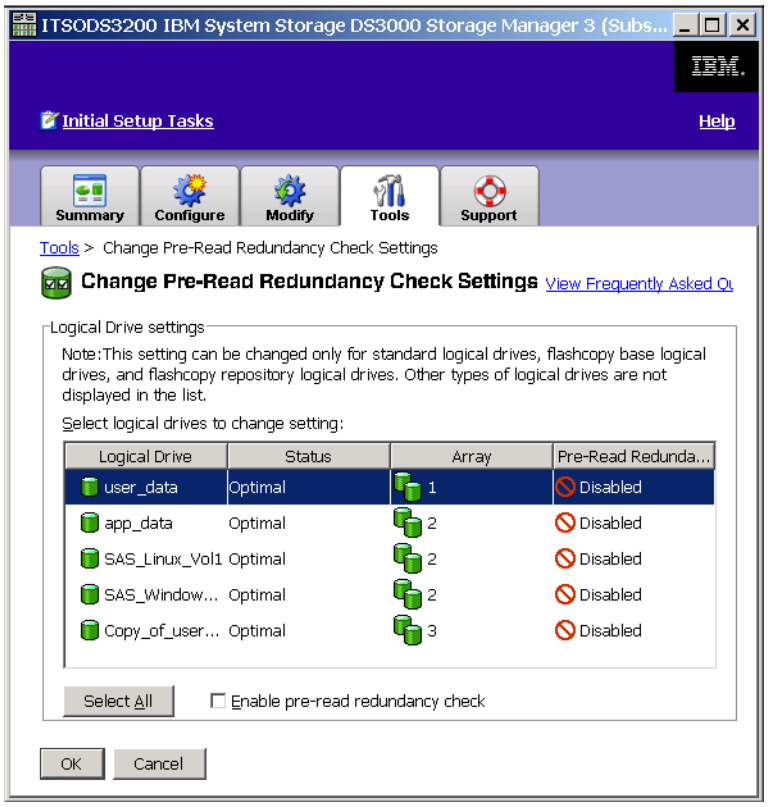


Figure 12-25 Change Pre-Read Redundancy Check Settings

The Pre-Read Redundancy Check option is implemented initially in DS3000 Storage Subsystems with firmware V7.35.

You can use this option to define the capability of the new DS3000 Storage Subsystem firmware to pre-read logical drive redundancy information and determine if the data of a given logical drive is consistent (redundant). This check is done on each read operation on the selected logical drive.

A logical drive that has this feature enabled returns read errors if the data is determined to be inconsistent by the controller firmware. It happens, for example, if data from one disk of a RAID 10 array is readable and mirrored data from a second disk is not readable.

Remember that:

- You cannot enable this feature for logical drives that do not contain redundancy information.
- You cannot enable this check for high-level logical drives (FlashCopy logical drives).
- You can enable pre-read redundancy check on each logical drive separately.
- You can enable this check for normal logical drives, and also on FlashCopy base logical drives, VolumeCopy source and target drives, and FlashCopy repository drives.

If a logical drive that is configured with a pre-read check is migrated to a RAID 0 type that does not maintain redundancy information, the metadata of the logical drive continues to reflect this setting. The read operations from that logical drive ignores pre-read redundancy check. If the logical drive is subsequently migrated back to a RAID type that supports redundancy, the feature will become active again.

A similar function is provided if the background media scan is set with a redundancy check. The Pre-Read Redundancy Check is similar to the Media Scan setting with redundancy check enabled (see 12.2.1, “Change Media Scan Settings” on page 272). The main difference is that the pre-read check is done online on each read operation and Media Scan checks for redundancy offline at a predefined time interval in days.

Note: Enabling this feature can influence the overall performance of read operations of the selected logical drives. This feature should be used for data verification mainly in environments where data consistency is a key requirement.

2. Select the logical drives that are required for a pre-read redundancy check, as shown in Figure 12-26. You can use the Ctrl or Shift key to select some of the list, or click **Select All** to select all the LUNs. Select the **Enable pre-redundancy check** check box and click **OK** to activate this feature.

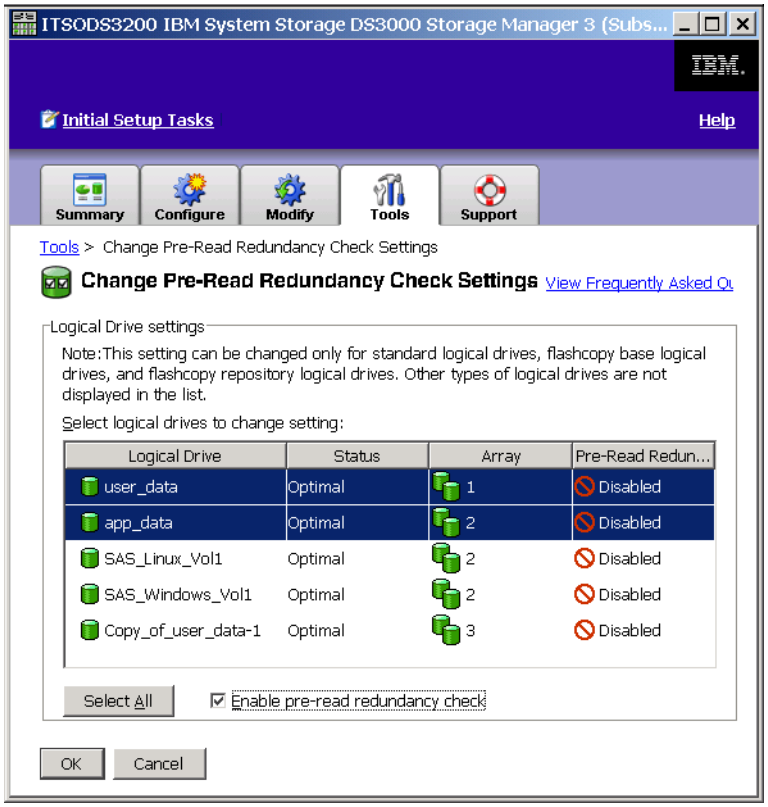


Figure 12-26 Select LUNs for pre-read redundancy check

3. Confirm your selection by clicking **Yes** in next window, wait for the process to complete, and click **OK**. When done, you will see a list of LUNs with changed settings.
4. You can see the new overall setting if you click **Change Pre-Read Redundancy Check Settings** in the Tools tab (Figure 12-1 on page 258).

12.3 Controllers tools

The third section of the Tools tab (Figure 12-1 on page 258) has three functions for controllers' settings:

- Synchronize Controller Clocks
- Configure Ethernet Management Ports
- Change Battery Settings

12.3.1 Synchronize Controller Clocks

Here you can synchronize the internal clocks of the storage subsystem controllers of the storage management station. The storage subsystem stores its own event log, so we recommend having all clocks set to the same time. If more than one storage management station is connected to a controller, the controller synchronizes to the storage management station that is giving the command. You should synchronize the controller clocks if the event time stamps written by the controller no longer match the event time stamps written to the host or storage management station.

Do the following steps:

1. Click **Synchronize Controller Clocks** (Figure 12-1 on page 258). Figure 12-27 displays the time of Controllers A and B and the Management Station.

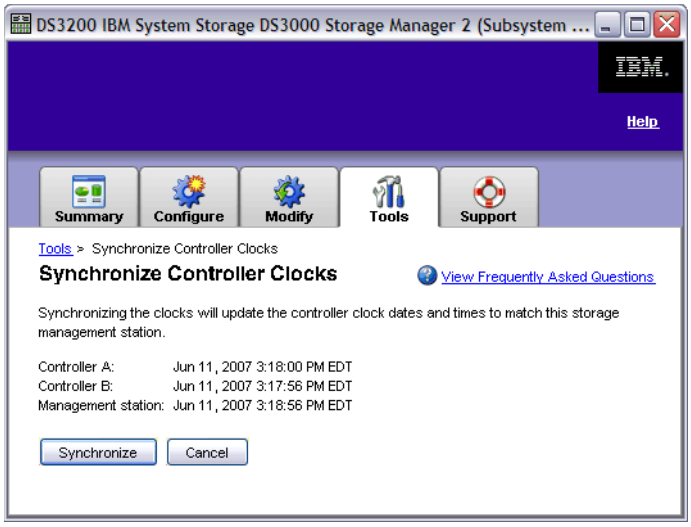


Figure 12-27 Synchronize Controller Clocks

2. Click **Synchronize** to update the time.

12.3.2 Configure Ethernet Management Ports

In 3.4, "DS3000 Storage Manager" on page 26, we describe the differences between in-band and out-of-band management. For out-of-band management, you have to assign an IP address to each controller. This can be done in the initial setup wizard (5.4.5, "(Optional) Changing the network configuration" on page 76) or you can also set and change your storage subsystem management IP addresses here. Setting up the IP addresses of the iSCSI ports is described in 13.2, "Identification and networking" on page 288.

Do the following steps:

1. Click **Configure Ethernet Management Ports** in the Tools tab in Figure 12-1 on page 258. Figure 12-28 displays the current network configuration of the controller. In the first drop-down menu, you can switch between the controllers A and B. Select the controller you want to view or change.

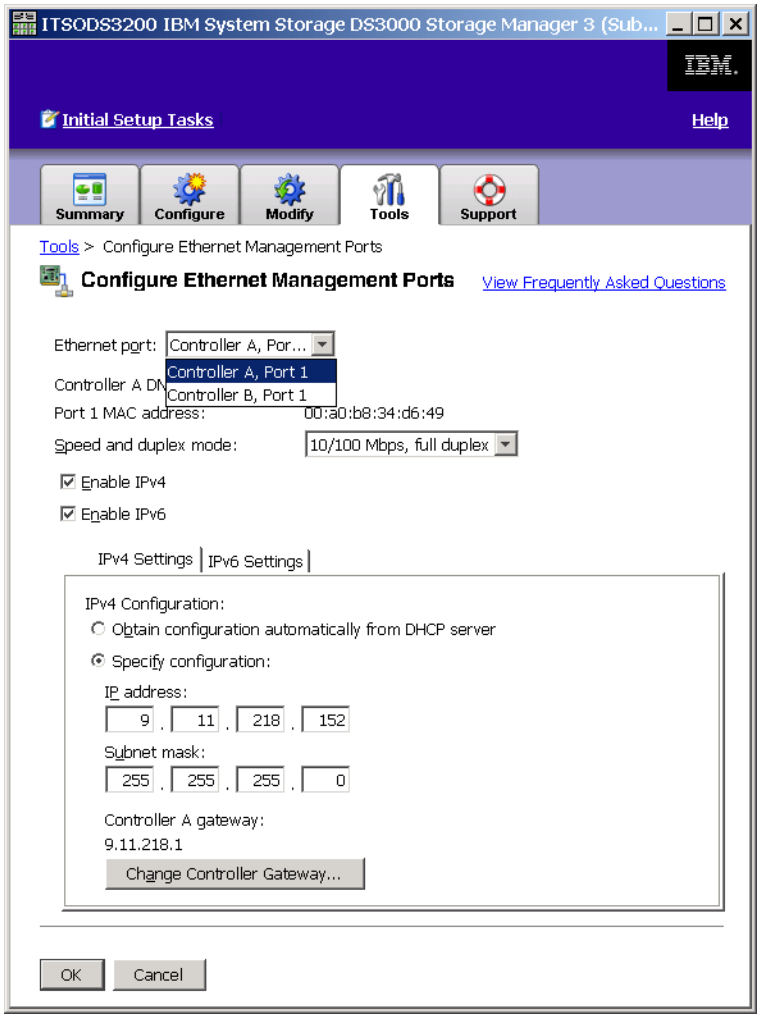


Figure 12-28 Change network configuration

2. The current IP network configuration is displayed. You can enable/disable and configure the IPv6 setting (new in firmware V7.35) if you plan to use it. We highly recommend the use of static IP addresses, as shown in Figure 12-28.

   You can also obtain IP addresses from a DHCP server and assign an IP address automatically, if you select the appropriate check box. If you do this, we recommend that you assign DHCP reservations so that the leases are maintained consistently across restarts of the DS3300 storage subsystem. If a controller's IP address changes to unknown values, you will have to automatically discover it again in the Enterprise view, as discussed in "Automatically discover new storage subsystems" on page 125. This task will be more difficult if your management station is in a different LAN subnet than your storage subsystem.

3. You can manually set the Speed and duplex mode of the Ethernet interface port in the second drop-down menu. "Auto-negotiate" is the default value.

4. There is also a button for Change Controller Gateway (Change Controller Router IP-Address for IPv6) in this window. The gateway (router) IP address should be set to the same address for both controllers.
5. When you finish the network changes, click **OK** to submit them. You can change the settings for controller A and B in one step.
6. A warning will appear (Figure 12-29) that, after changing the network configuration, your storage subsystem may not be discovered in the Enterprise view. Read it and click **Yes** to confirm.

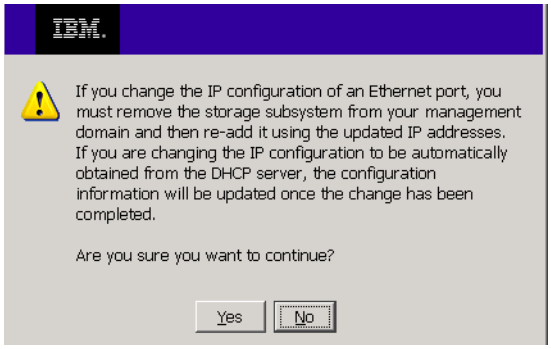


Figure 12-29 Confirm change network configuration

7. Now the changes are submitted to the controllers. Storage Manager displays these changes when the changes are completed and reminds you again to re-add your subsystem to the enterprise view (Figure 12-30). We explain how to do this in “Add Storage Subsystem to view” on page 127.

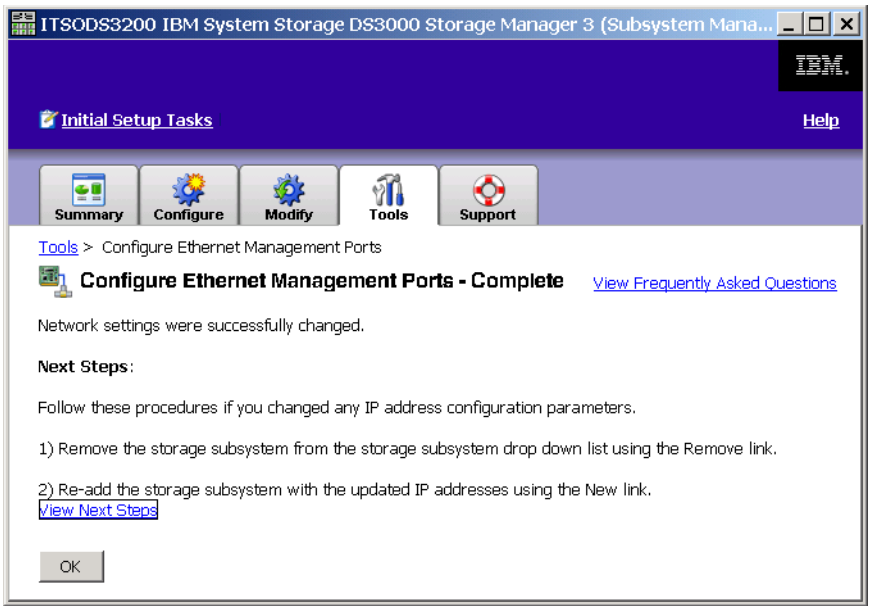


Figure 12-30 Change network configuration - complete

Note: You can lose the management connection to your DS3000, and then your management GUI stops responding if you change the IP addresses of both controllers. Go to the **Enterprise** menu and remove and add the Storage Subsystem again with the new IP addresses.

12.3.3 Change Battery Settings

Each controller has a battery to protect the cached data in case of a power loss. Open this view to monitor the age of your batteries (Figure 12-31). If the batteries are nearly at end of life, contact IBM Support to replace the batteries. After replacing the batteries, you should reset the age for the new batteries. To reset the counter, just click the **Reset Age** button for each battery you selected in the drop-down menu.

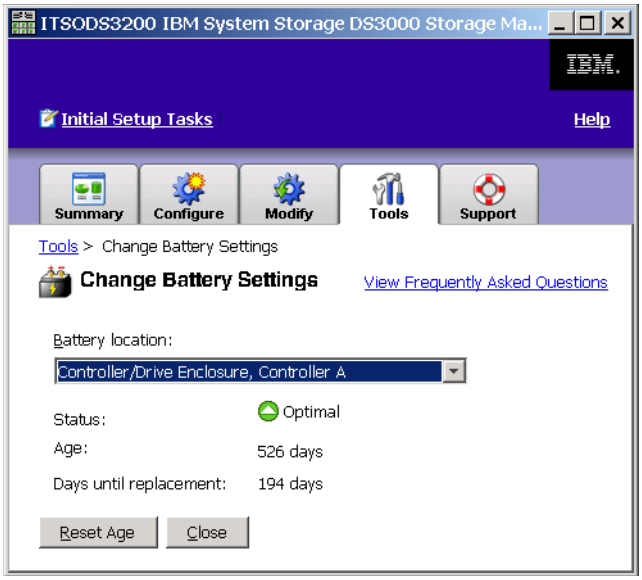


Figure 12-31 Reset Battery Age

12.4 Accessibility tool

The final section of the Tools tab has options for configuring the appearance of Storage Manager for easier viewing.

12.4.1 Inherit System Settings

This setting allows the Storage Manager software to use the color and font settings defined in the operating system. It should be useful especially if you want increase or decrease the font size in the DS3000 Storage Manger windows.

Do the following steps:

1. Click **Inherit System Settings** (see Figure 12-1 on page 258).

2. In Figure 12-32, select the **Inherit system settings for color and font** check box if you want to inherit the settings from the operating system and click **OK**.

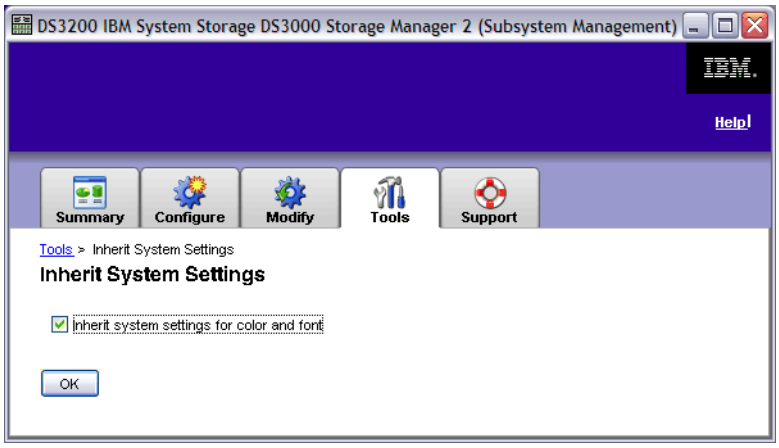


Figure 12-32 Inherit System Settings

13

Administration - iSCSI

In this chapter, we describe the Manage iSCSI Settings tab of the DS3000 Storage Manager. These functions are specific for an IBM System Storage DS3300 and, for this reason, are only available when a DS3300 is managed.

We cover the following topics:

- Authentication
  - Change Target Authentication
  - Enter Mutual Authentication Permissions
- Identification and networking
  - Change Target Identification
  - Change Target Discovery
  - Configure iSCSI Host Ports
- Session and statistics
  - View / End iSCSI Sessions
  - View iSCSI Statistics

© Copyright IBM Corp. 2007, 2009. All rights reserved.

Figure 13-1 shows the iSCSI tab in Storage Manager.

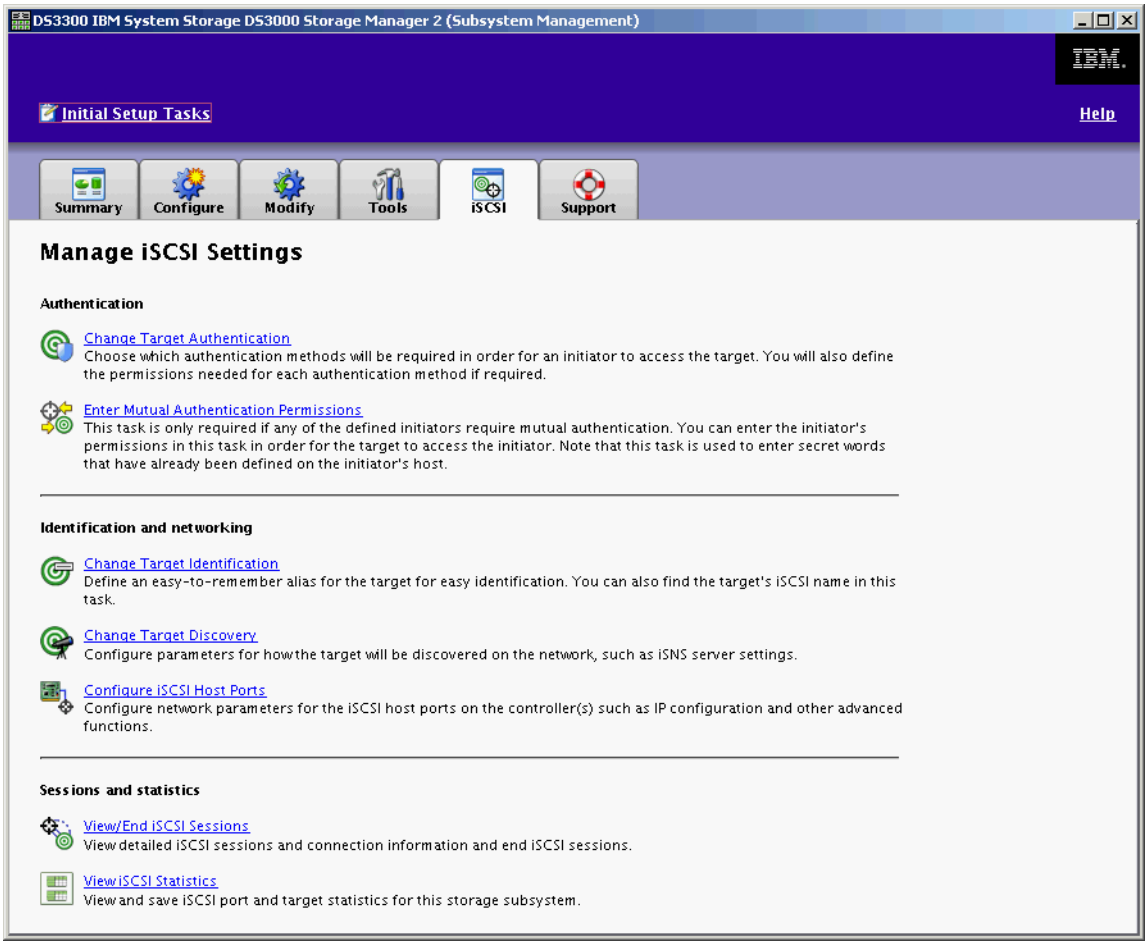


Figure 13-1 DS3000 Storage Manager - iSCSI

13.1 Authentication

Settings that affect the security of iSCSI connections can be done here. We discuss iSCSI security in 1.3.3, “iSCSI security considerations” on page 11.

13.1.1 Change Target Authentication

Target authentication is used to allow only configured HBAs to access the storage subsystem. When target authentication is enabled, the target needs to authenticate itself against the initiator that attempts to access the storage subsystem. When using a QLogic iSCSI HBA, target authentication is called *bidirectional authentication* and can only be configured on the HBA when the initiator itself authenticates against the storage subsystem. If this is required, you have to also configure an initiator CHAP secret on the DS3300 by following the steps in 13.1.2, “Enter Mutual Authentication Permissions” on page 286.

To configure target authentication, follow these steps:

1. On the iSCSI tab (Figure 13-1 on page 284), click **Change Target Authentication**.
2. A new dialog opens (Figure 13-2 on page 286) that provides two options:
   - None
   - CHAP

   This setting affects the connection between an iSCSI initiator and a DS3300 iSCSI port. Option None allows any iSCSI initiator to establish an iSCSI connection to this target. When option CHAP is selected, an initiator is required to provide a CHAP password to get a session established. CHAP needs to be enabled if mutual authentication is required by an iSCSI initiator. Both options (None and CHAP) can be enabled together; in this case, initiators with and without a target secret can access the storage subsystem.

   Note: Only one CHAP target secret can be defined. All initiators using target authentication must use the same secret.

   From a security perspective, we recommend that you enable CHAP. However, since the configuration of CHAP adds some complexity, we suggest that you set up and test all connections with no CHAP, and later implement the security.

Figure 13-2 shows both target authentication options checked.

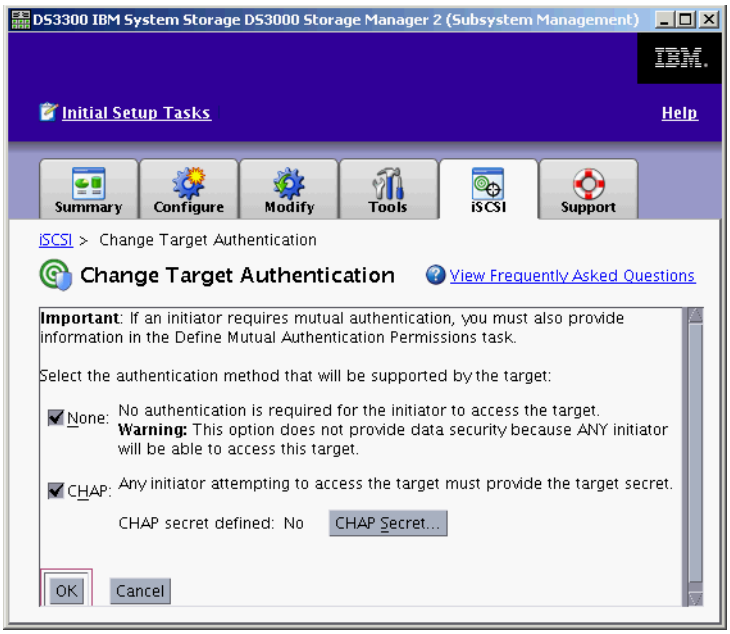


Figure 13-2 DS3000 Storage Manager - iSCSI target authentication

3. If CHAP is enabled, you have to define the CHAP target secret. Click **CHAP Secret...**.
4. Figure 13-3 shows the Enter Target CHAP secret dialog. Enter a CHAP secret that is at least 12 characters long. Repeat the text to confirm.

   Alternatively, click **Generate Random Secret** to generate a 16 character long secret. Click **OK**.

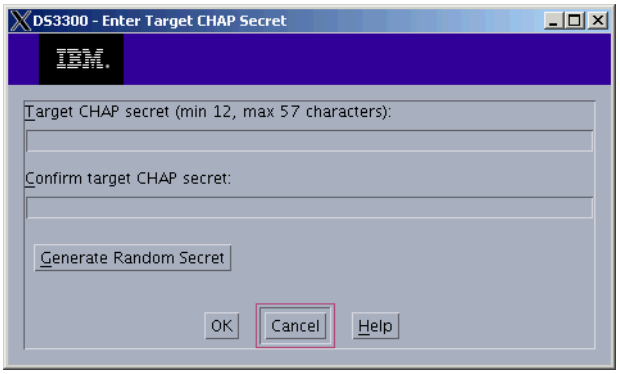


Figure 13-3 DS3000 Storage Manager - Target CHAP secret

5. You will return to the Change Target Authentication window (Figure 13-2). Click **OK** to save the modifications.

13.1.2 Enter Mutual Authentication Permissions

With Mutual Authentication, a CHAP secret is configured that a specific initiator must use it to get access to the storage subsystem. That is also called initiator authentication. Each initiator can have its own CHAP secret. Initiator authentication works without target authentication for QLogic iSCSI HBAs, but not vice versa.

Do the following steps:

1. On the iSCSI tab (see Figure 13-1 on page 284), click **Enter Mutual Authentication Permissions**.
2. Select a host port that has already been defined, as described in 9.2, “Configure hosts” on page 140. In these iSCSI menus, the host ports are referred to as Initiator Labels. Figure 13-4 shows a selected host port, Hudson-Port0. There is no CHAP secret defined for this initiator, as shown in the last column, Permission Set?. Click **CHAP Secret...**.

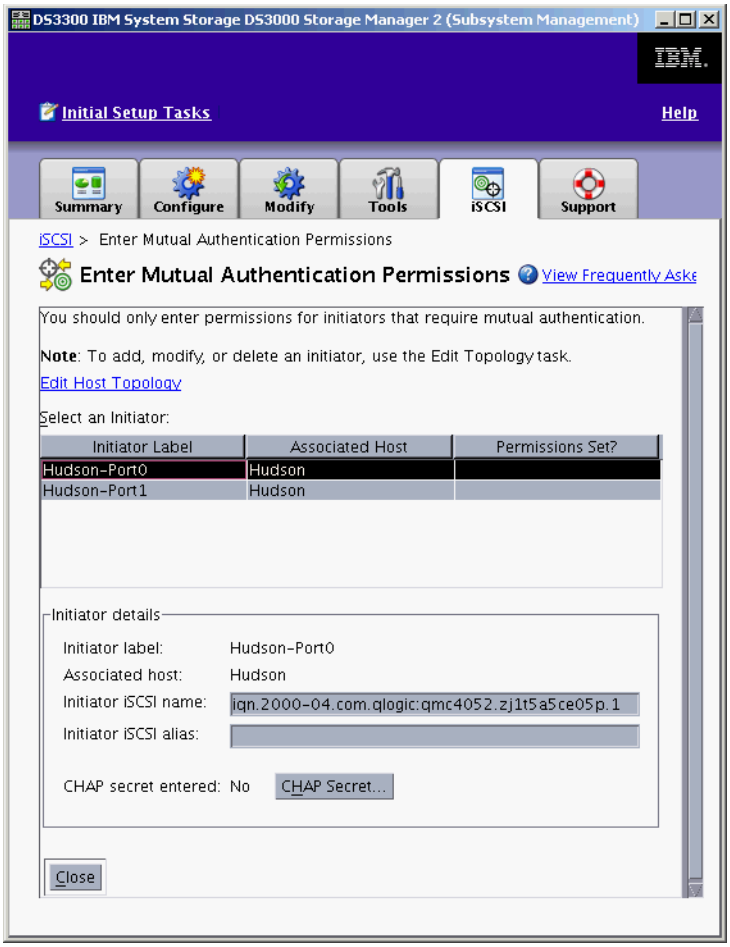


Figure 13-4 Storage Manager - Enter Mutual Authentication Permissions

3. Enter an initiator CHAP secret with at least 12 characters twice, as shown in Figure 13-5. Click **OK**.

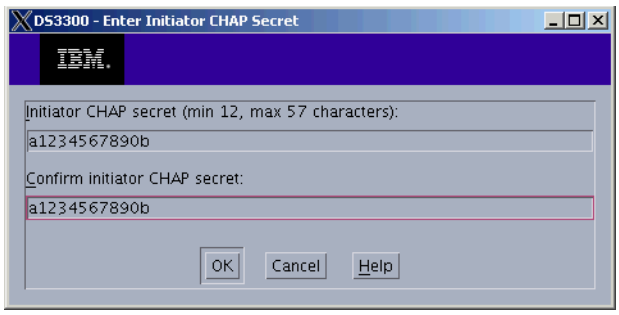


Figure 13-5 Storage Manager - Initiator CHAP Secret

4. In the confirmation window, Figure 13-6, click **OK**.

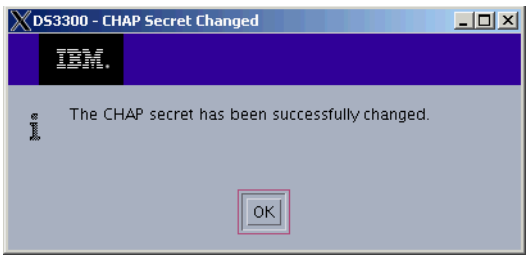


Figure 13-6 Storage Manager - Initiator CHAP Secret Changed

5. The Permission Set? column now indicates that the port has a CHAP secret defined, as shown in Figure 13-7. Define CHAP secrets for all remaining host ports using the same procedure. When this is complete, click **OK** to exit this task.

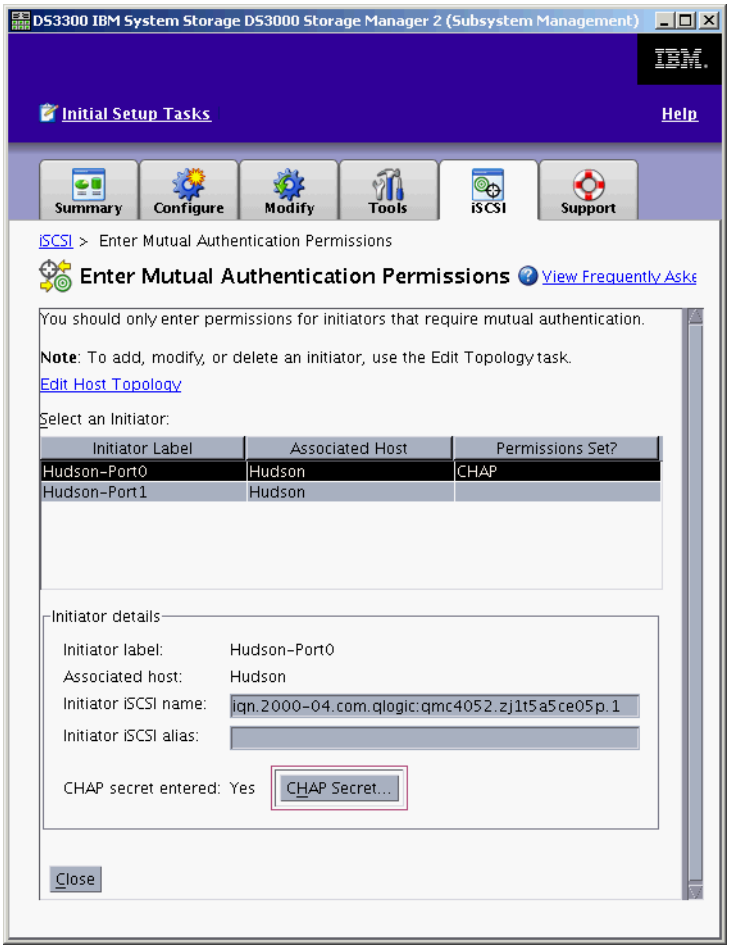


Figure 13-7 Storage Manager - Defined mutual authentication settings

13.2 Identification and networking

Identification and networking allows you to configure the networking settings for the two iSCSI ports on each controller, the target discovery method that can be used by initiators, and an alias for the subsystem.

13.2.1 Change Target Identification

Change Target Identification allows you to set up an alias name for the iSCSI storage subsystem that can be seen in the iSCSI session information. To set up an iSCSI target alias:

1. Click **Change Target Identification**.
2. Enter an alias name (DS3300 in our example in Figure 13-8) and click **OK**.

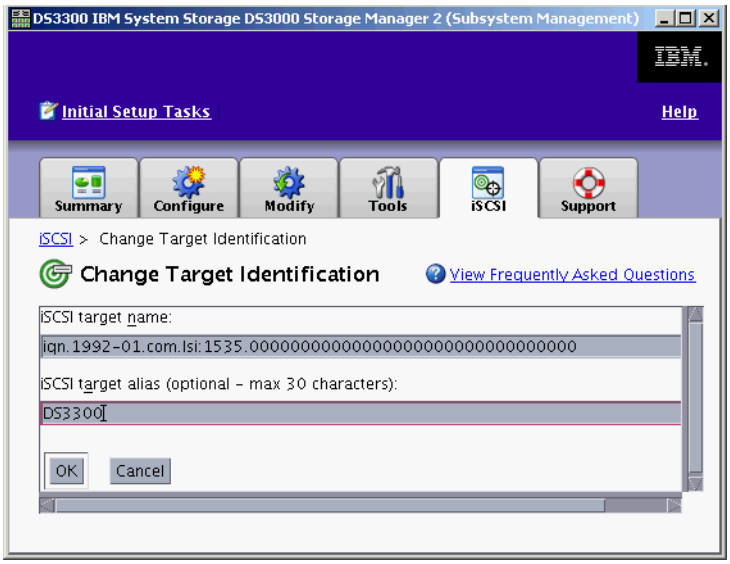


Figure 13-8 Storage Manager - Change Target Identification

3. You will see a confirmation message, as shown in Figure 13-9. Click **OK** to return to the iSCSI tab.

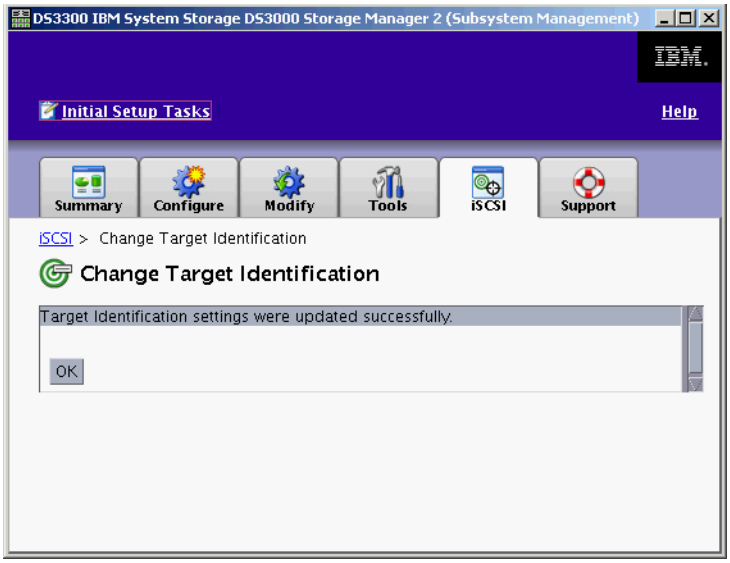


Figure 13-9 Storage Manager - Target identification changed

13.2.2 Change Target Discovery

We discussed techniques for target discovery in 1.3.2, “iSCSI discovery” on page 10. The Change Target Discovery dialog can be used to specify if an Internet Storage Name Service (iSNS) will be used to register this target for easier location by iSCSI initiators.

By default, the Send Targets and Static methods are supported from the initiator to discover iSCSI targets. An iSNS server can also be specified, using either its IPv4 or IPv6 address or by DHCP. If using iSNS, your DHCP should be configured with option 43 (Vendor Specific Info) to inform the DHCP client about the IP address of the iSNS server. See your DHCP server’s documentation for more information about how to configure this option.

The change target discovery, do these steps:

1. To modify the discovery settings, click **Change Target Discovery**.
2. Select **Use iSNS server**, as shown in Figure 13-10.

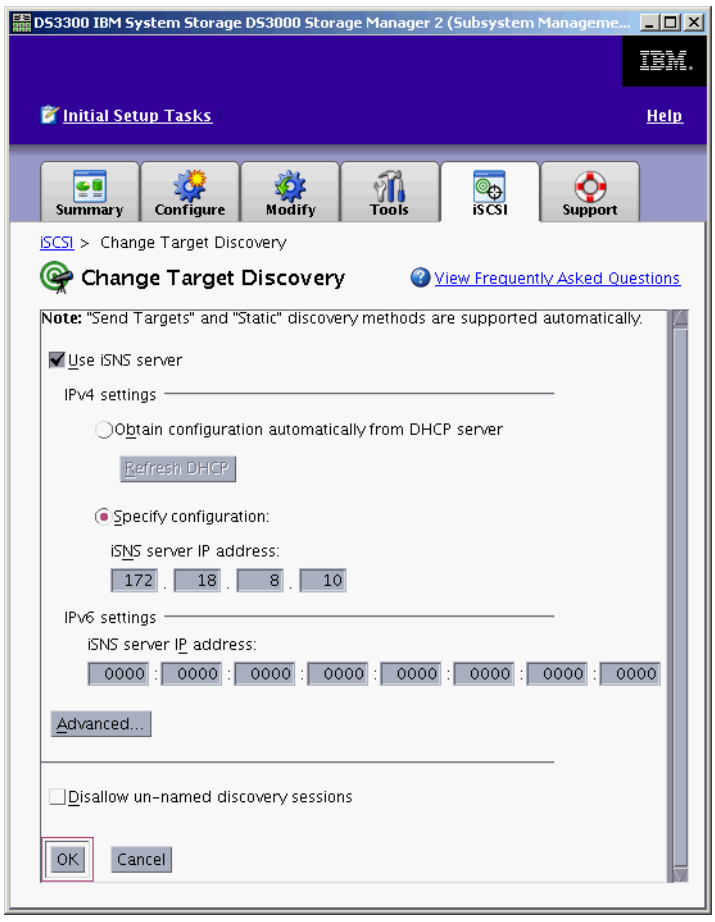


Figure 13-10 Storage Manager - Enable iSNS server

3. Enter the IPv4 or IPv6 address of the iSNS server or select **Obtain configuration automatically from DHCP server**.

Note: At the time of writing, if you want to use a DHCP-discovered iSNS server, option 43 must be configured on the DHCP server.

4. Click **Advanced** to define a custom port that is used to communicate with the iSNS server. The default port is 3205.
5. Select **Disallow un-named discovery sessions** if you want your storage subsystem only to respond to requests addressed specifically to this target; it will not respond to global requests.
6. Click **OK** to save the configuration modifications and return to the iSCSI management window.

13.2.3 Configure iSCSI Host Ports

iSCSI host ports are the ports on each DS3000 controller that are used to connect hosts to this storage subsystem. Each controller has two ports available on the iSCSI daughter card (see 3.2.2, “IBM System Storage DS3300” on page 21).

The Configure iSCSI Host Ports task is used to configure the network settings of the two iSCSI host ports on each controller. IPv4 or IPv6 addresses can be used along with native VLANs. Careful network planning must be done before the setup can be performed.

You need the following information before starting this configuration task. Your organization probably has networking standards that determine this information:

- Use of IPv4, IPv6, or both
- Use of native VLANs or not
- Use of static IPv4 addresses or addresses assigned by a DHCP/BOOTP server
- IP addresses used by the subsystem
- Whether to allow ICMP ping responses
- Requirement and size of jumbo frames

To configure the network settings of the iSCSI host ports, follow these steps:

1. Click **Configure iSCSI Host Ports** in the iSCSI tab. A window opens, as shown in Figure 13-11. Our example window shows an active IPv4 configuration that was assigned through DHCP.

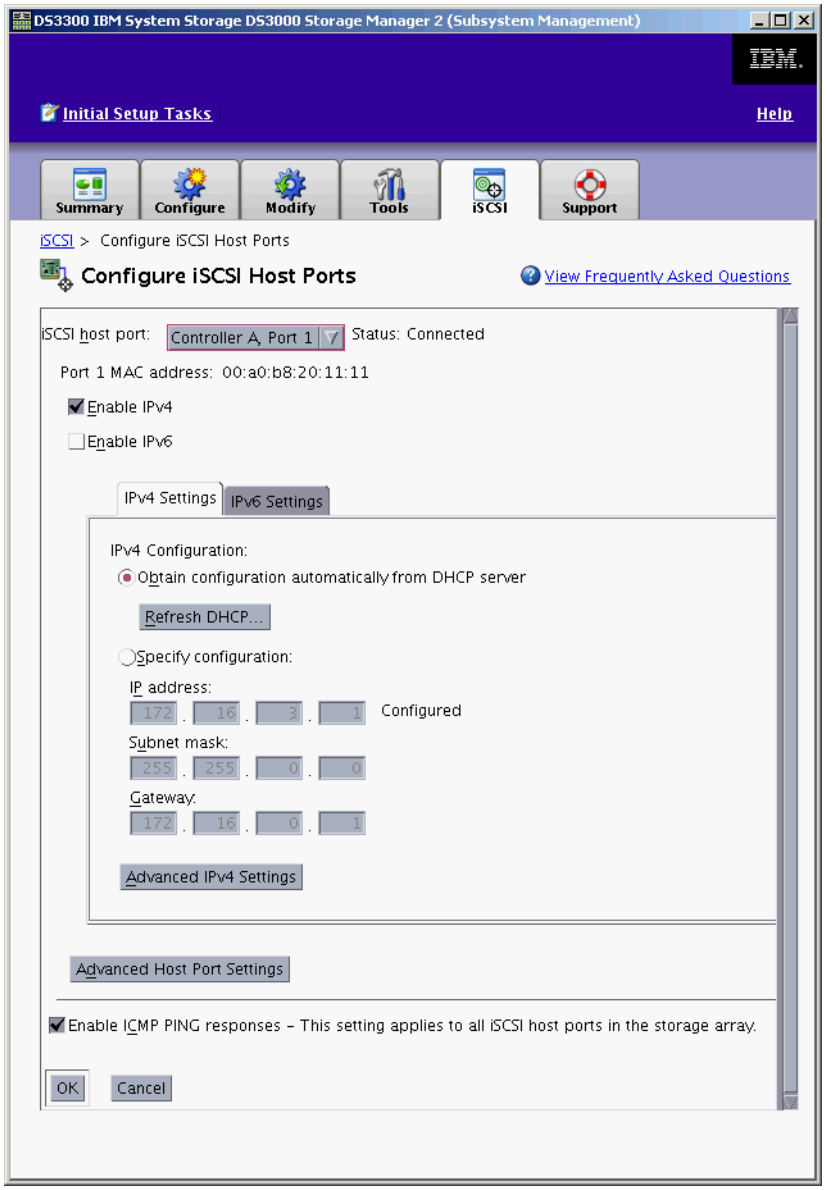


Figure 13-11 Configure iSCSI host port

2. Select the iSCSI host port to be configured from the **iSCSI host port** drop-down menu, as shown in Figure 13-12.

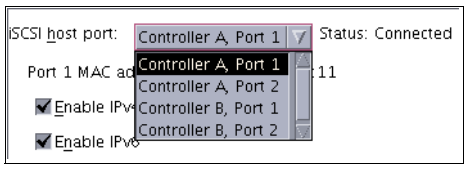


Figure 13-12 iSCSI host ports

3. IPv4 is enabled by default. To enable IPv6, select **Enable IPv6**, as shown in Figure 13-13



Figure 13-13 IP protocol versions

4. On the IPv4 Settings tab (Figure 13-14), click **Obtain configuration automatically from DHCP server** if the host port address will be assigned dynamically through DHCP. To ensure that the host ports always receive the same IP address, you should configure your DHCP server to use MAC addresses as the base for the fixed IP address assignment.

   Warning: Using DHCP for the target ports is generally not recommended. If you use DHCP, you should assign DHCP reservations so that leases are maintained consistently across restarts of the DS3300 storage subsystem. If static IP reservations are not provided or you lose the DHCP configuration, the initiator ports can lose communication to the DS3300 controller and may not be able to reconnect to the device.

   Alternatively, click **Specify configuration** if you want to enter a static IP address, and enter the address, subnet mask, and gateway.

   In Advanced IPv4 Settings, you can enable native VLAN support and specify a VLAN ID that will be used by the selected port.

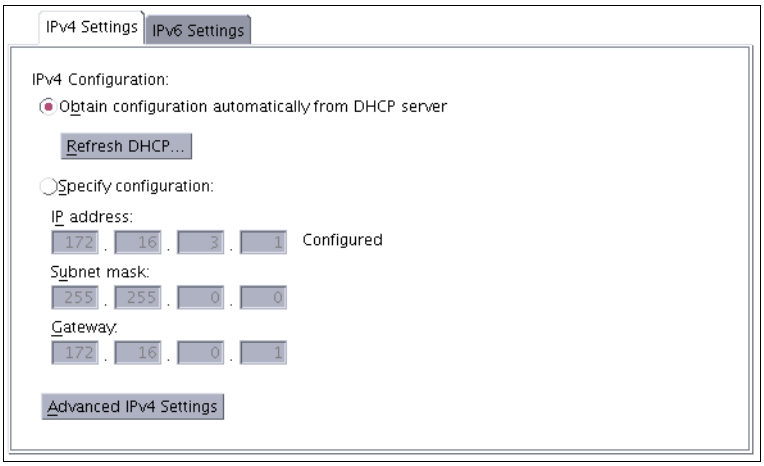


Figure 13-14 IPv4 address settings

5. Figure 13-15 shows the IPv6 settings.

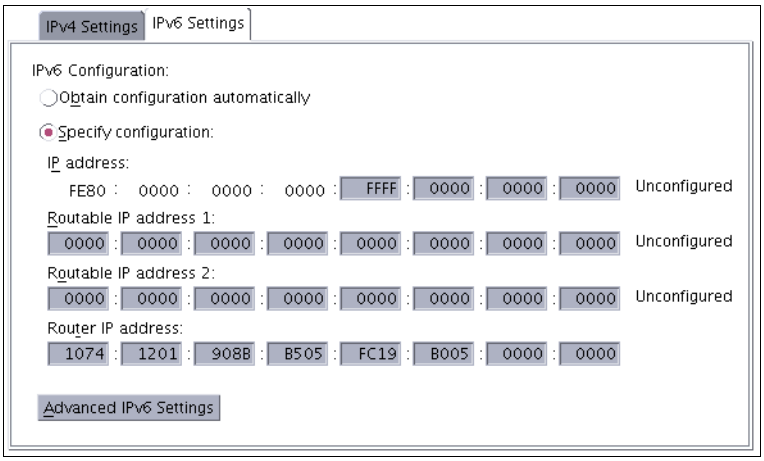


Figure 13-15 IPv6 address settings

Click **Advanced IPv6 Settings** to enable native VLAN support and specify a VLAN ID that will be used by the selected port.

6. After configuring IPv4, IPv6, or both settings, click **Advanced Host Port Settings** (shown in Figure 13-11 on page 292). These settings apply to both the IPv4 and IPv6 settings. The TCP listening port for iSCSI and jumbo frame support can be configured here. To use a TCP port other than 3260, check **Use custom listening por**t and specify the port. A custom port can be between port 49152 and 65535.

   Specify the jumbo frame size after activating jumbo frame support in the Advanced Host Port settings. The default frame size is 1500 bytes per frame. The MTU can be defined between 1501 and 9000 bytes per frame.

7. In the Configure iSCSI Host Ports window (Figure 13-11 on page 292), you have the option to disable ping (ICMP echo) requests for all iSCSI host ports. Figure 13-16 shows this option selected. A reset of all ports is required to apply this setting.

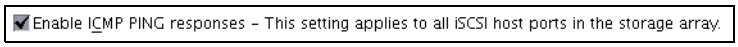


Figure 13-16 ICMP echo response setting

8. After setting all the options, click **OK** to save the modifications. In Figure 13-17 on page 295, you are advised which ports will be reset to apply the configuration. Active iSCSI sessions that might be connected to such ports will be closed. Click **Yes** to proceed.

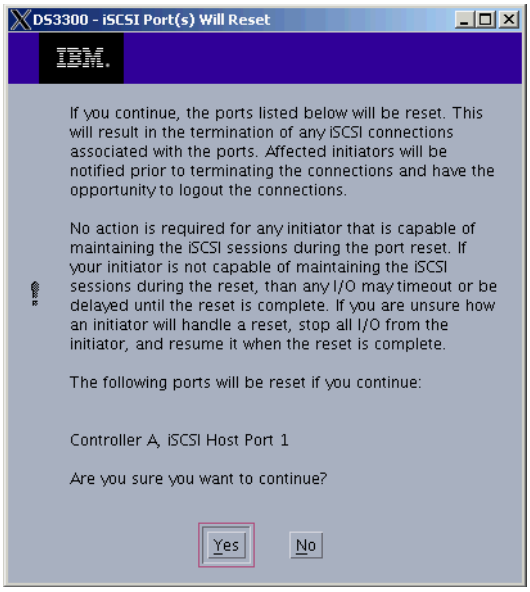


Figure 13-17 iSCSI port reset warning

9. Figure 13-18 shows the Storage Manager notification when the settings are applied. Click **OK** to return to the iSCSI tab, as shown in Figure 13-1 on page 284.

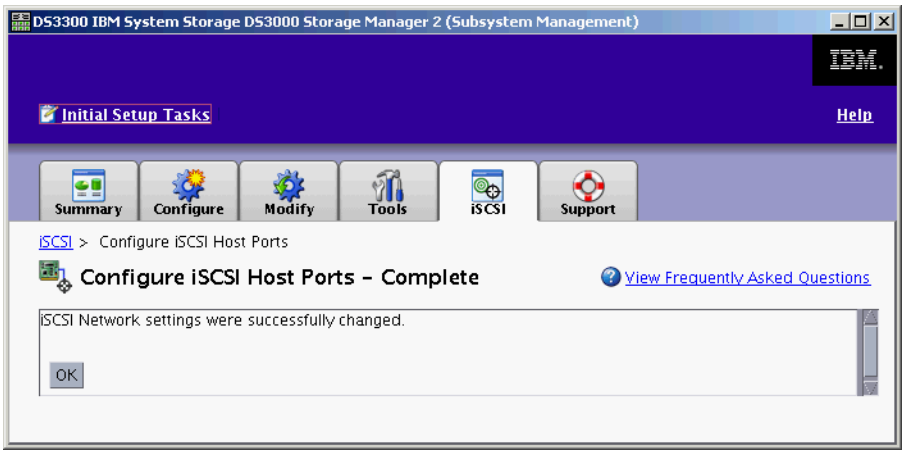


Figure 13-18 iSCSI host port setting confirmation

10. Configure the remaining iSCSI host ports in the same way.

13.3 Sessions and statistics

Sessions and statistics provide information about ongoing iSCSI connections to this target.

13.3.1 View/End iSCSI Sessions

The View/End iSCSI Sessions option shows all initiators that have active sessions. Details for each session can be viewed and saved. Do the following steps:

1. Click **View/End iSCSI Sessions** on the iSCSI tab. In the top pane you will see all the current sessions.

2. Select a session to see the details in the lower pan, as shown in Figure 13-19.

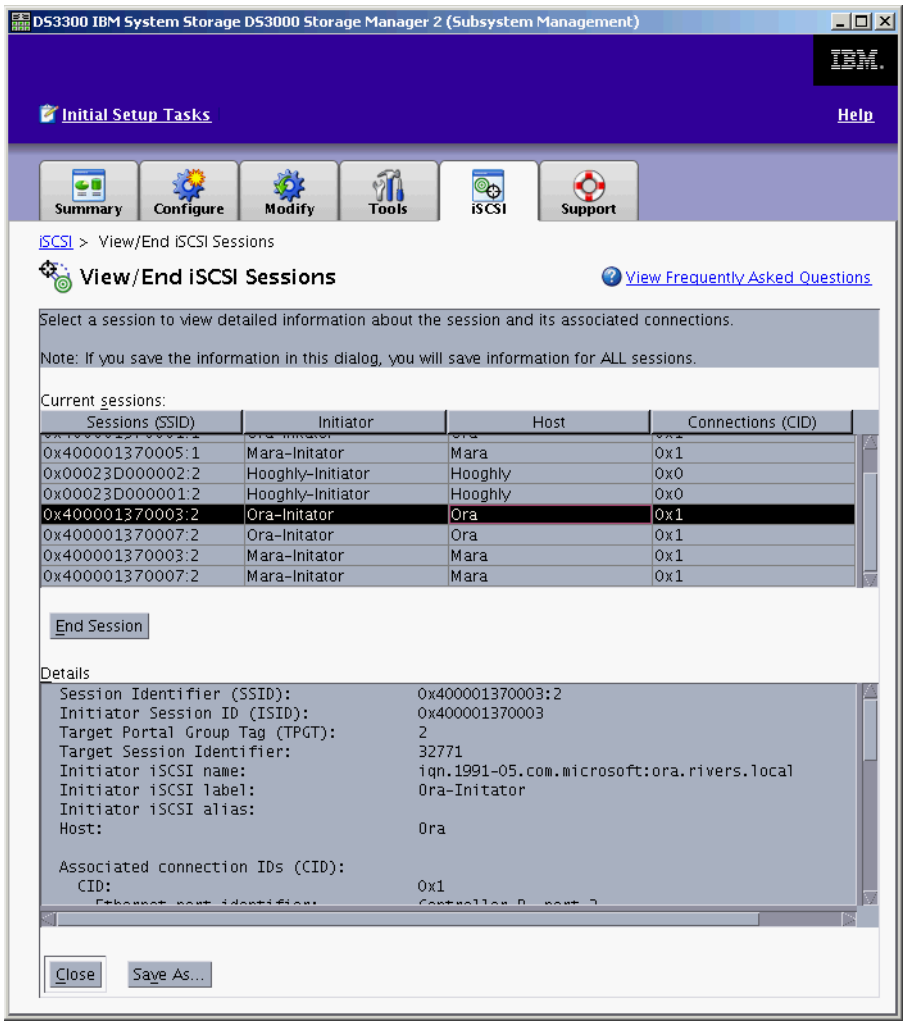


Figure 13-19 View sessions

3. To terminate a session, select it in the top pane and click **End Session**. If the initiator is still available, the session may be re-established.
4. In the bottom pane, click **Save As...** to copy all the session details into a plain text file.
5. Click **Close** to return to the iSCSI tab.
6. Example 13-1 on page 297 shows the complete session information for the selected initiator in Figure 13-19. This output can help verify that the host setup was done correctly.

Example 13-1 sample session information

```
Session Identifier (SSID):                0x400001370003:2
  Initiator Session ID (ISID):            0x400001370003
  Target Portal Group Tag (TPGT):         2
  Target Session Identifier:              32808
  Initiator iSCSI name:                   iqn.1991-05.com.microsoft:mara.rivers.local
  Initiator iSCSI label:                  Mara-Initiator
  Initiator iSCSI alias:
  Host:                                   Mara

  Associated connection IDs (CID):
    CID:                                  0x1
      Ethernet port identifier:           Controller B, port 1
      Initiator IP address:               172.16.200.91

  Negotiated login parameters:
    Authentication method:                None
    Header digest method:                 CRC32c
    Data digest method:                   CRC32c
    Maximum connections:                  4
    Target alias:
    Initiator alias:
    Target IP address:                    172.16.3.2
    Target portal group tag:              2
    Initial R2T:                          Yes
    Maximum burst length:                 262144
    First burst length:                   8192
    Default time to wait:                 0
    Default time to retain:               60
    Maximum outstanding R2T:              16
    Error recovery level:                 0
    Maximum receive data segment length:  65536

    Target name:                          iqn.1992-01.com.lsi:1535.00000000000000000000
    Initiator name:                       Mara-Initiator
```

13.3.2 View iSCSI Statistics

The View iSCSI Statistics link gives you information about packets in the OSI layers 2 to 5. Figure 13-20 on page 298 shows as an example of the TCP/IP statistics in base line mode.

Packets can be measured in RAW mode or base line mode; the difference between the two modes is the time since the packets were counted. RAW mode uses the time since the controller was started and base line mode use a time that was defined by the user by clicking **Set Baseline**. Using this button resets all base line counters to zero and starts counting from scratch.

Use the **Save As...** button to save all current statistics in a plain text file. RAW and baseline mode statistics for all levels are saved.

A link to the online help is provided to translate column names into more meaningful names for each statistic type.

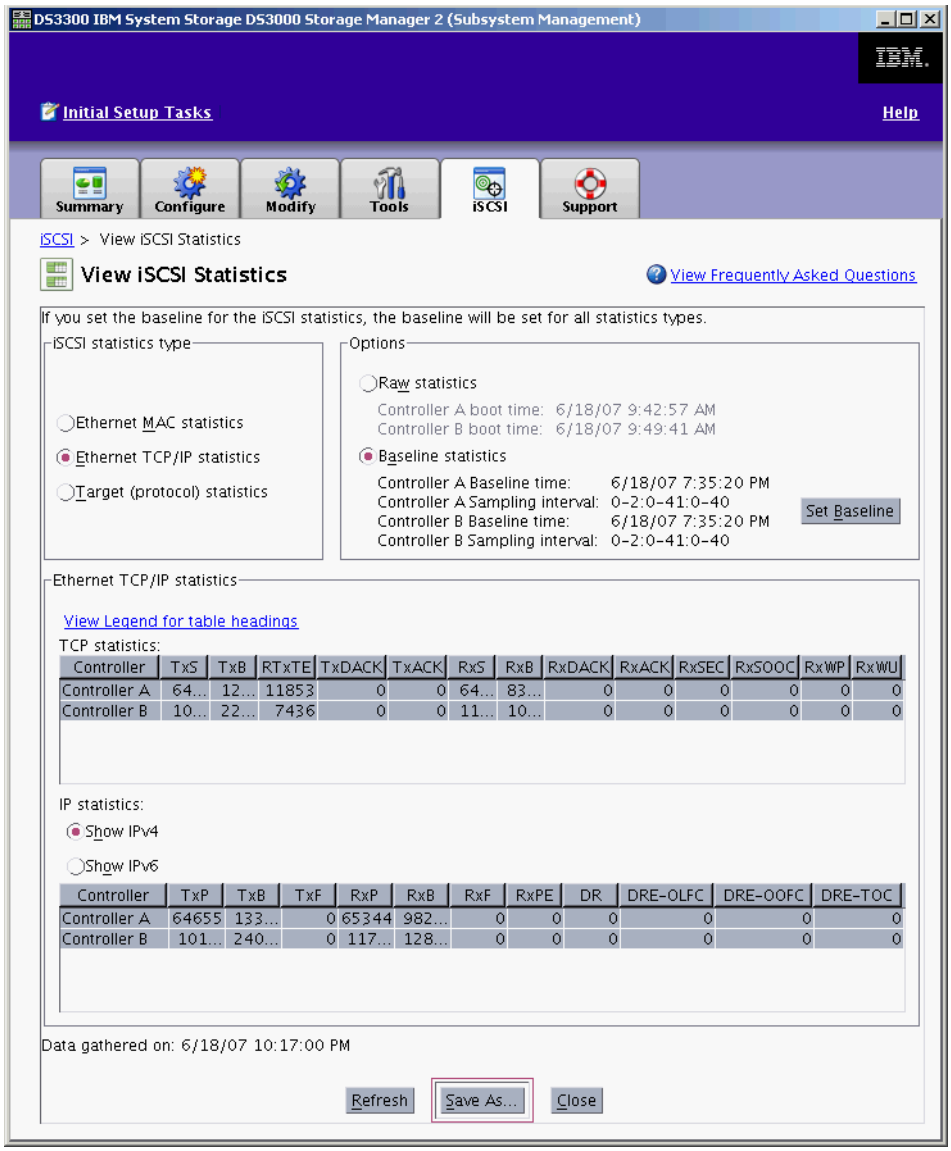


Figure 13-20 View iSCSI statistics for TCP/IP

14

Administration - Support

In this chapter, we describe some of the support functions of the IBM System Storage DS3000, which help with troubleshooting and maintenance tasks.

The Support tab (Figure 14-1 on page 300) covers the following links:

- Troubleshooting and maintenance
  - Gather Support Information
    - Save Support Information
    - Change Automatic Collection Setting
  - View Storage Subsystem Profile
  - Download Firmware
    - Download Controller Firmware
    - Download Controller NVSRAM
    - Download Drive Firmware
    - Download Environmental (ESM) Card Firmware
  - Manage Controllers
    - Place Controller Online or Offline
    - Redistribute Logical Drives
  - Advanced Support tasks
    - View Event Log
    - Import Array
    - Export Array
    - Change Remote Login Settings
  - Recovery Guru
- Help
  - View Online Help
  - About IBM System Storage DS3000we

© Copyright IBM Corp. 2007, 2009. All rights reserved.

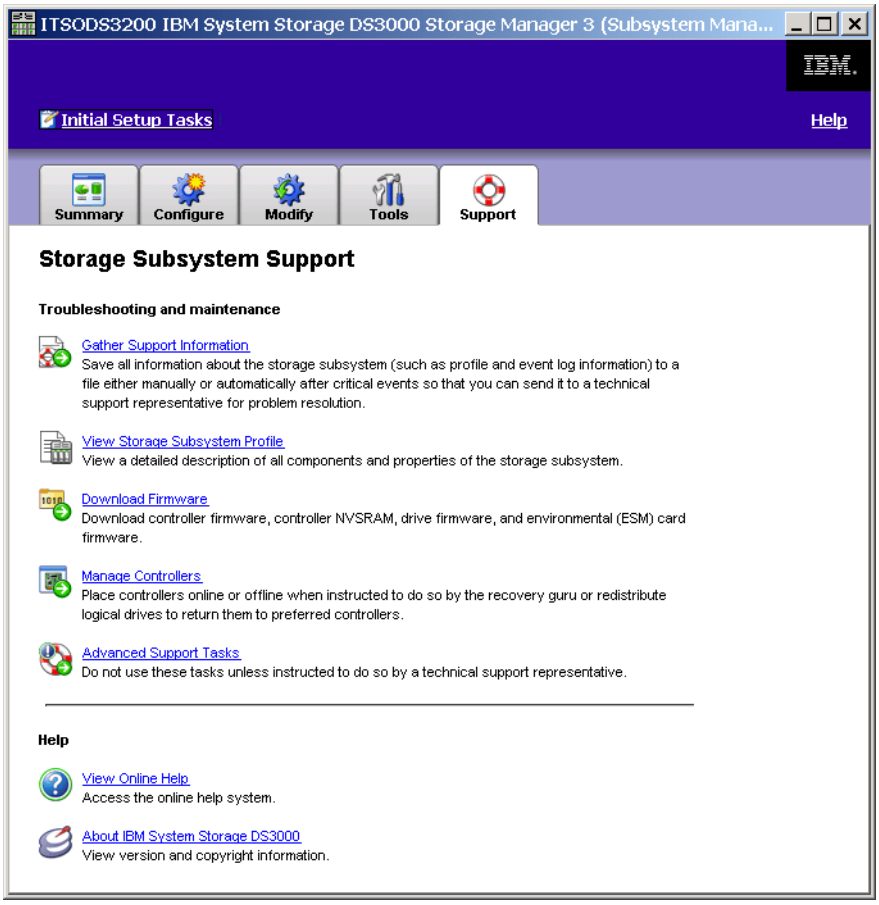


Figure 14-1 Storage Subsystem Support Tab

14.1 Gather Support Information

This wizard helps gather all storage subsystem support information, including profile and event log information, and stores it to a *.zip file. IBM technical support will typically request this file for analysis.

Click the link **Gather Support Information** in the main Support tab (Figure 14-1 on page 300) to display the following window (Figure 14-2) with two options.



Figure 14-2 Gather Support Information

14.1.1 Save Support Information

Follow these steps to manually collect all support information data to one *.zip file:

1. Select **Save Support Information** in the main window (Figure 14-2 on page 301). A new window appears (Figure 14-3), where you specify where to save support data.

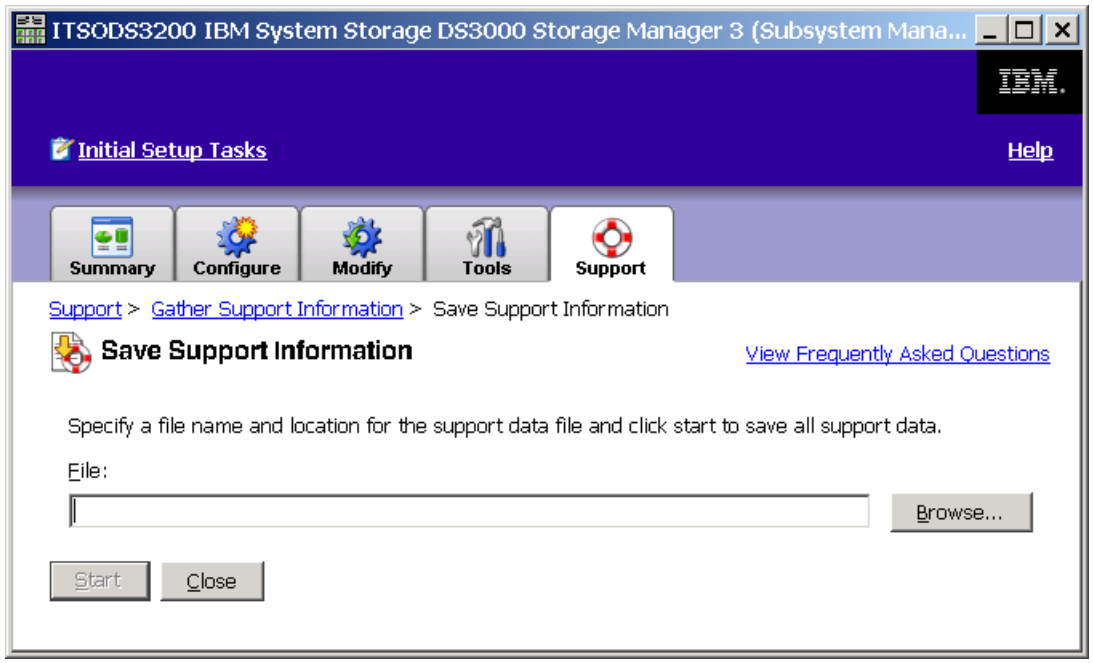


Figure 14-3 Save Support Information

2. Enter a directory and file name directly, or click **Browse** to select a name and location to store the file (Figure 14-4).

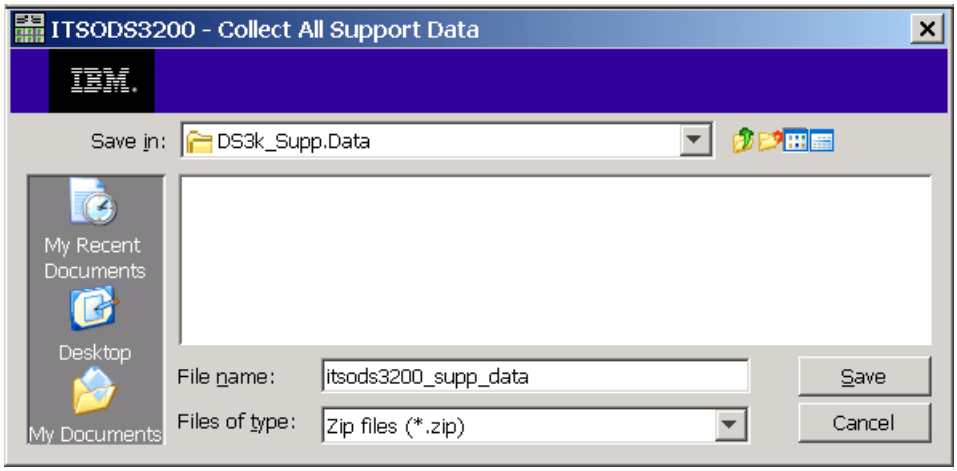


Figure 14-4 Collect All Support Data

3. Click **Start** and the system will start to collect all the data and store it to the selected location (Figure 14-5 on page 303). This might take a couple of minutes and it depends on number of hard disks and logical objects configured on the storage subsystem. Also, the

storage subsystem password is required during the “Controller state capture” data collection.

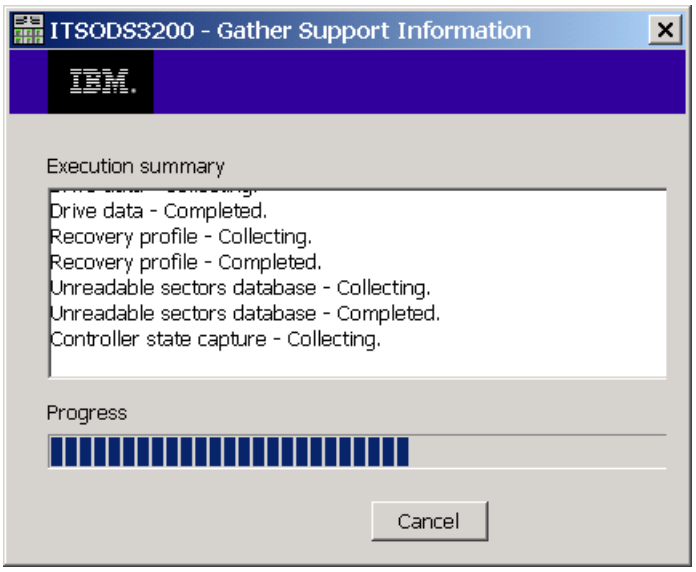


Figure 14-5 Gather Support Information

4. When the collection is done, click **Close** (Figure 14-6). The zip file will be generated, which you can send when requested.

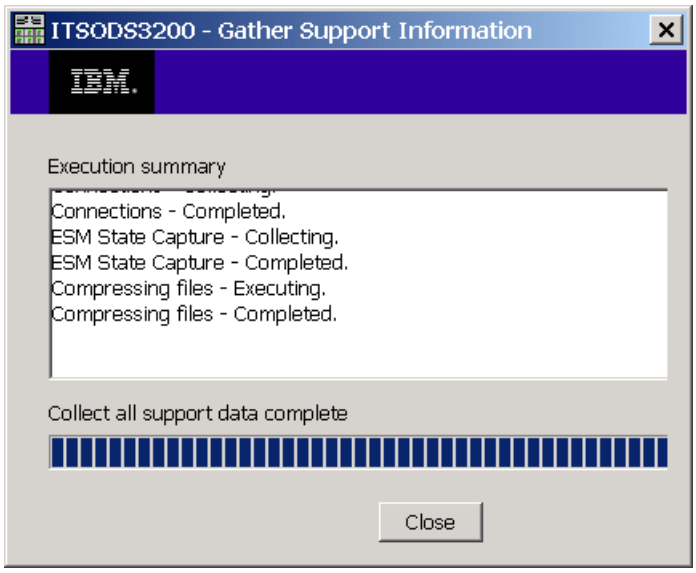


Figure 14-6 Gather Support Information - Complete

14.1.2 Change Automatic Collection Setting

The second option in the Gather Support Information window *(*Figure 14-2 on page 301*)* is Change Automatic Collection Setting.

DS3000 Storage Manager V10.35 incorporates a new option to enable automatic collection of support data. When enabled, in the case of a critical event, the support data file is collected and transferred to the directory specified and all the information relevant for troubleshooting

by your support representative is preserved. In case consecutive critical events occur, only the information related to the last event is stored in auto save support bundle.

Tip: We recommend that you enable the Automatic Collection option in order to have a support data file automatically generated and saved to the specified location after the occurrence of a critical event. Make sure to:

- Specify a directory outside your DS3000 system to collect the information.
- Have the SM Monitor process running on the workstation or host where you collect the logs.

To enable this feature:

1. Click **Change Automatic Collection Setting** in **Gather Support Information** (Figure 14-2 on page 301).
2. Click the **Automatically collects support data for critical events** check box in the new window (Figure 14-7) to enable the automatic data collection and, optionally, specify the destination folder for collected data.

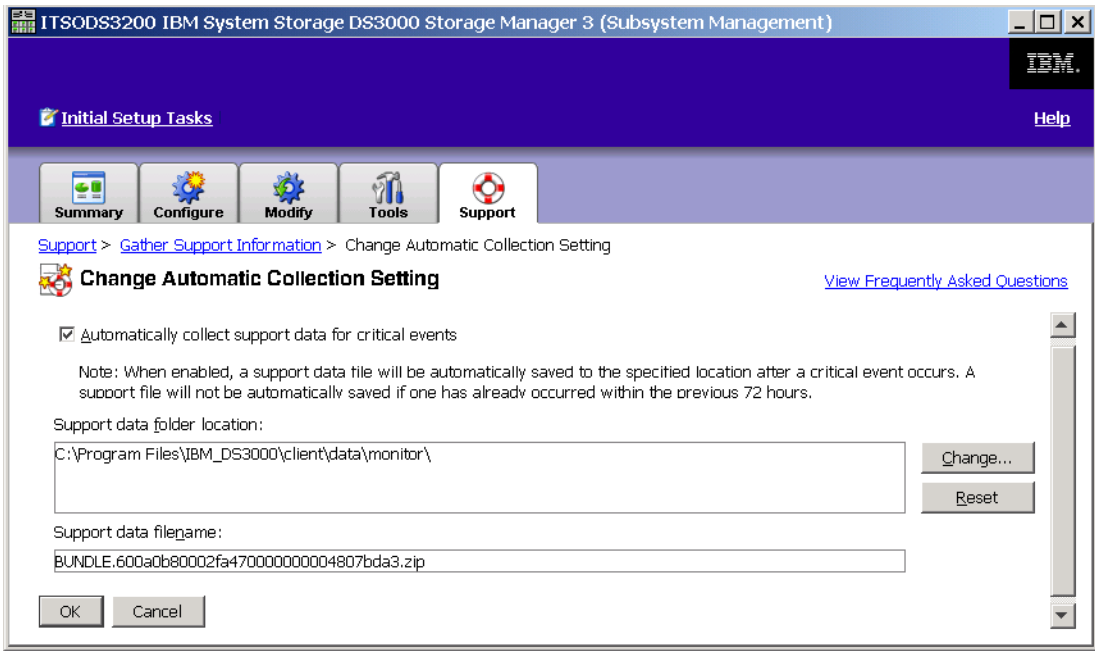


Figure 14-7 Change Automatic Collection Setting

14.2 View Storage Subsystem Profile

We introduced the Storage Subsystem Profile in 10.2.1, “Storage Subsystem Profile” on page 212. It shows detailed information about the configuration of the storage subsystem and should be a part of your system documentation. It includes information about:

- Summary
- Controllers
- Arrays
- Logical drives
- Drives
- Disk drive channels

- Enclosures
- Mappings

You can see all of the information about the storage subsystem at once, or filter information about a specified component. You can view the profile in the window or save it as a text file.

View Storage Subsystem Profile

Follow these steps:

1. Click **View Storage Subsystem Profile** in the Support tab (Figure 14-1 on page 300). The Storage Subsystem Profile window appears, as shown in the Figure 14-8.

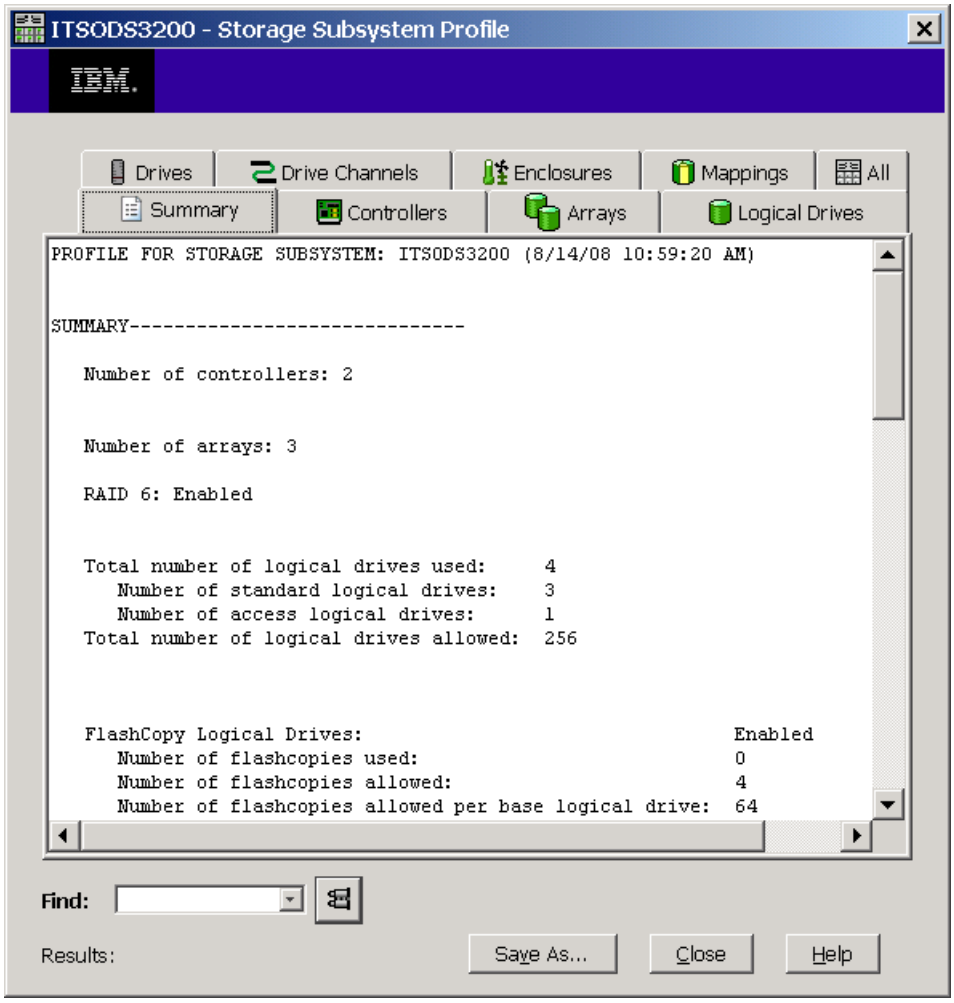


Figure 14-8 View Storage Subsystem Profile

2. Select the tab for the information you want to view. By default, it opens the Summary tab. Select a specific component or click the **All** tab to show everything.

3. Use the **Find** box (Figure 14-9) to search the profile. Enter the required text and click the “binoculars” icon to find the string in the current tab.



Figure 14-9 Find Box

4. To save the profile, click **Save as**. At the top of Figure 14-10, choose which information you want to save. Tabs that you have already viewed will be selected, and you can select any other tabs that you want.

   Select a location and file name for the saved profile and click **Save**. The selected information will be stored as a text file.

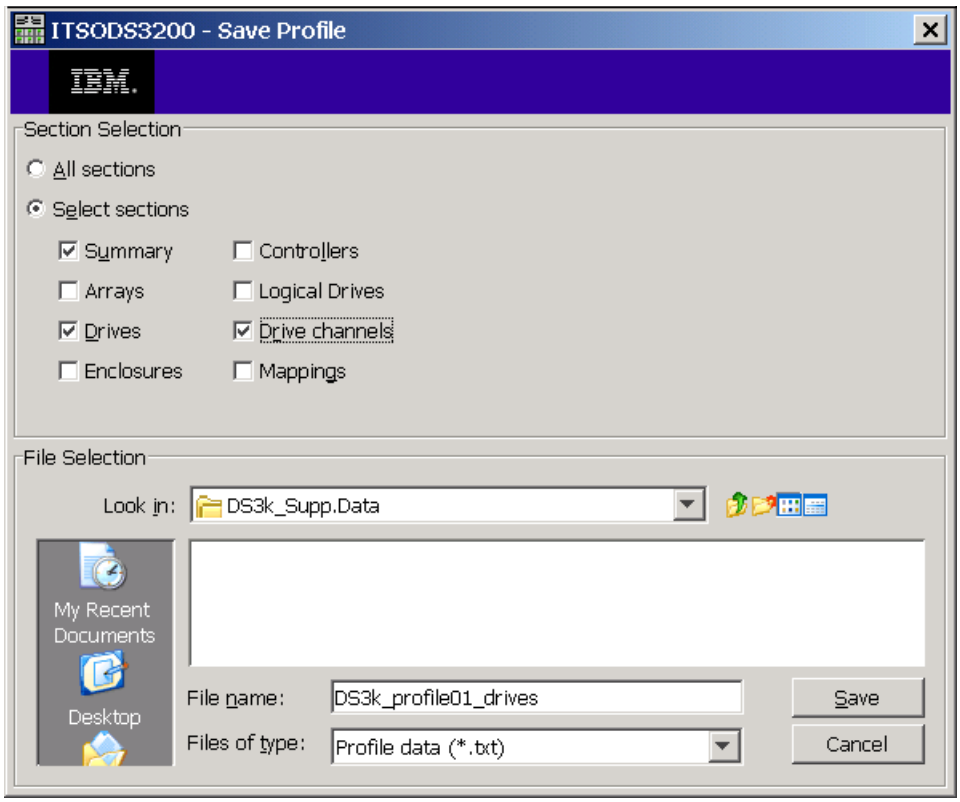


Figure 14-10 Save Profile

5. Click **Close** to return to the Support tab.

Note: We recommend that you save a new version of the profile and store it securely whenever the configuration changes. Even in the case of a complete configuration loss, you can restore the arrays and logical drives configuration as well as the mappings for the storage partitioning. The profile should be stored locally, and also be included in any offsite disaster recovery documentation.

Note: Much more information is saved if all support data is collected (see 14.1, “Gather Support Information” on page 301).

14.3 Download firmware

This section describes the steps required to upgrade the DS3000 storage server microcode. You can download the latest firmware at:

http://www.ibm.com/servers/storage/support/disk

Appendix B, “IBM Support Web site” on page 733 describes this Web site and its contents.

On the same Web page you can select (for example, as shown in Figure 14-11, for the DS3200) the firmware (controller, ESM, disk, and NVSRAM), HBA device drivers and firmware, and Storage Manager versions. Also, the links for the product documentation, troubleshooting tips, and all other information regarding DS3200 are placed on this site.

The DS3000 firmware is closely connected to the Storage Manager version and HBA device driver and firmware available for your hosts. Always check the DS3000 firmware readme and change-history files to see if there are any host software dependencies. If there is newer Storage Manager (including multi-path device drivers) or HBA software required for the hosts, download and install these first before starting the DS3000 firmware update. Also, ESM firmware should be upgraded to the highest level before starting the controller firmware upgrade. In general, we recommend running at the latest levels unless specifically advised otherwise.

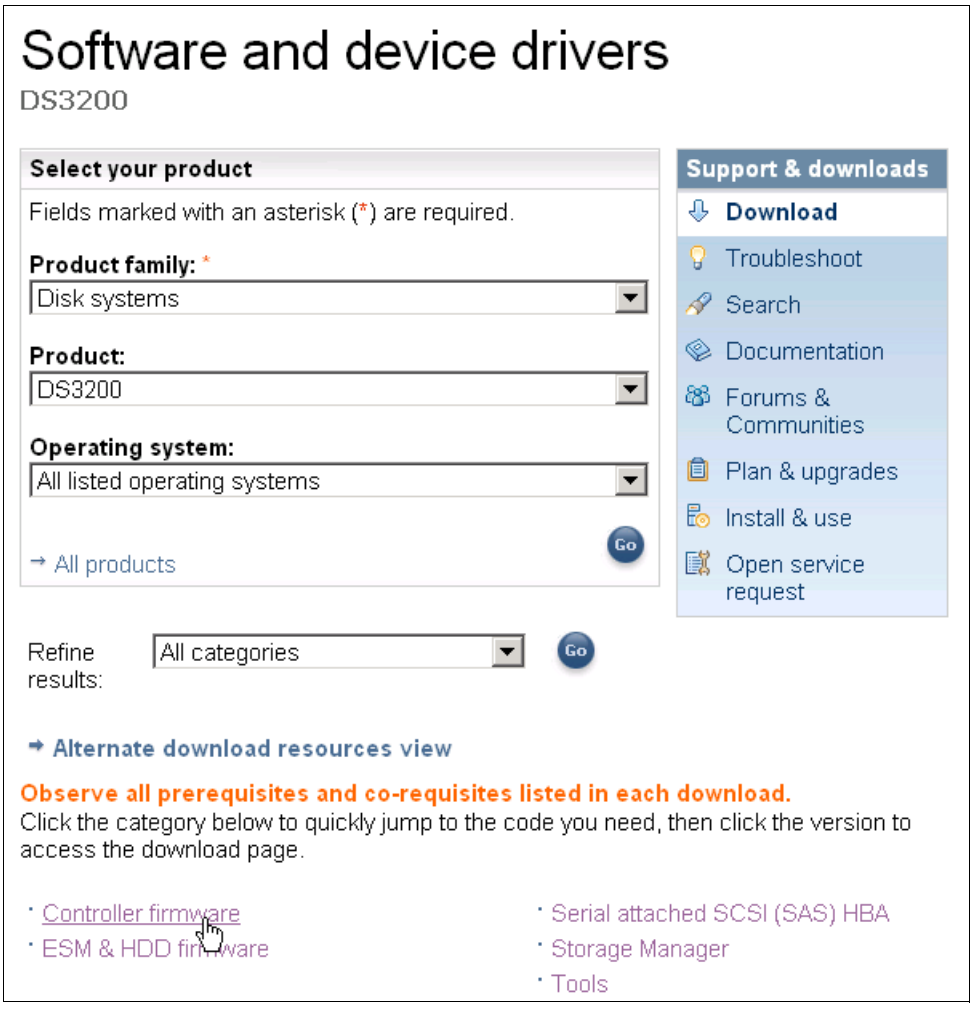


Figure 14-11 Code downloads for DS3200

Updating firmware should require a maintenance window in some cases, so it is better to plan an appropriate outage time in advance.

Download all the packages you will need from the IBM Support Web site and store them unzipped on a local disk.

Important: The procedures given in this section are provided as a guide; however, these can change. Always carefully check the readme files distributed with any firmware package and follow *those* procedures exactly.

14.3.1 Upgrade by DS3000 Storage Manager

This chapter describes how to upgrade DS3000 firmware from any 6.xx version to the latest 6.yy version, or from V7.35 to some future 7.xx versions by DS3000 Storage Manager.

Section 14.3.2, “Upgrade using the DS3000 Controller Firmware Upgrade Tool” on page 324 describes how to upgrade from any Version 6.xx (6.17.31.00 minimum) to Version 7.35 or some later 7.xx versions through a special firmware upgrade utility.

Updating the host

Before you update the firmware on your DS3000, the software on your hosts should also be up to date. Go to the IBM Support Web site listed at the beginning of 14.3, “Download firmware” on page 307 to check the levels of Storage Manager and HBA drivers. If there are newer versions available, download them and check the readme files for any special requirements for updating them.

Update the HBA driver and Storage Manager if necessary. Updating Storage Manager is described in Chapter 6, “Installing IBM System Storage DS3000 Storage Manager” on page 81 and Part 4, “Sample configurations” on page 389. Storage Manager also includes the associated multi path driver. This driver also has to be updated on all attached hosts when multiple HBAs are installed.

Note: Updating the host server software might require a system reboot. You may need to plan downtime.

Also check the HBA firmware; it should also be at the latest level, which you will find at the IBM Support Web site listed at the beginning of 14.3, “Download firmware” on page 307.

Updating the storage subsystem

After you have updated your hosts, you can update the code of your storage subsystem. To do this, click **Download Firmware**, as shown in Figure 14-1 on page 300. You will see the options for different types of firmware download, as shown in Figure 14-12 on page 309.

Tip: The current DS3000 firmware levels are listed in the Storage Subsystem Profile. Clicking each of the individual Download Firmware links in Figure 14-12 also displays the current firmware levels of the respective components.



Figure 14-12 Download Firmware

The four options are:

- Download Controller Firmware

  Includes the controller firmware of the storage subsystem.

- Download Controller NVSRAM

  The NVSRAM is similar to the settings in the BIOS of a host system. The controller firmware and the NVSRAM are closely related, so be sure to install the correct combination of the two packages. Also, be sure to use the single or dual controller version of NVSRAM. Both Controller Firmware and NVSRAM may be downloaded at the same time.

- Download Drive Firmware

  Includes the firmware of the hard disk drives. Host IO must be stopped during this upgrade.

- Download Environmental (ESM) Card Firmware

  Includes the firmware of expansion enclosures. Should be done before the controller firmware upgrade in some cases.

Note: The term “Download” in these links means to update the actual firmware on the hardware. Before starting any of these operations, you must have already obtained or downloaded the appropriate microcode from the IBM Support Web site.

Before you update a storage subsystem:

1. Make sure that the system is in the Optimal state (see 10.1, “Status box” on page 210). If not, run the Recovery Guru to diagnose and fix the problem before proceeding with the upgrade.
2. Always read the readme file before upgrading any firmware to check for compatibility with your system.
3. Take care with dual controller configurations. The upgrade procedure requires two connections to your storage server, one to each controller; it is not possible to perform a firmware upgrade with only one controller connected. Therefore, both controllers must be accessible, either in-band or out-of-band, and both controllers must be online. If you upgrade redundant configurations in-band, make sure that a multi-path driver is installed on the management host. This is necessary since the access logical drive moves from one controller to the other during this procedure and the DS3000 must be manageable during the entire time.

Note: To check the controller status, you can first view the overall system status on the Enterprise window. If the subsystem status is optimal, then both controllers are working. You can also go to the Storage Subsystem Profile in the Controllers section. This lists the status of each controller, as shown in Figure 14-13.

Note: It is also a best practice to do all support data collection (14.1, “Gather Support Information” on page 301) before storage subsystem upgrades.

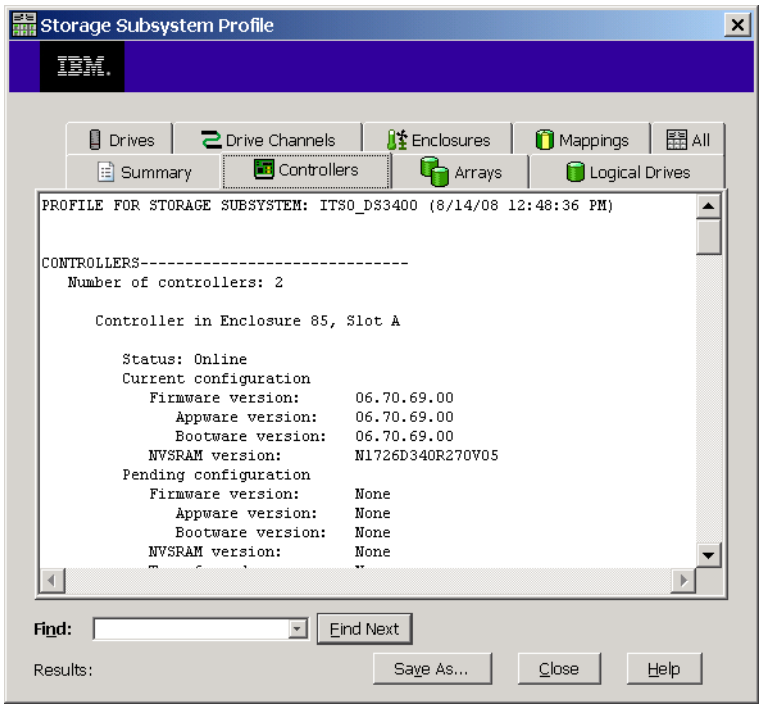


Figure 14-13 Display controller status

Note:

1. Update the ESMs first to the latest level, followed by the controller second. Outdated ESM firmware could make your expansion enclosures inaccessible after the DS3000 storage server firmware update. The required ESM firmware level for each Controller firmware level is mentioned in the readme for controller firmware.
2. Update the hosts, and if it is a redundant configuration, make sure that the latest multipath driver is installed.
3. Update the Controller firmware, followed by the NVSRAM; these two components must be updated together or immediately in this order.
4. Any power or IP network/SAN/SAS interruption during the update process may lead to configuration corruption. Therefore, do not power off the DS3000 storage server or the management station during the update. If you are using in-band management and have Fibre Channel hubs or managed hubs, then make sure no SAN connected devices are powered up during the update. Otherwise, this can cause a loop initialization process and interrupt the process.

Download ESM Card Firmware

Follow these steps:

1. To update the ESM code, click the **Download Environmental (ESM) Card Firmware** link (Figure 14-12 on page 309). Figure 14-14 shows the window.

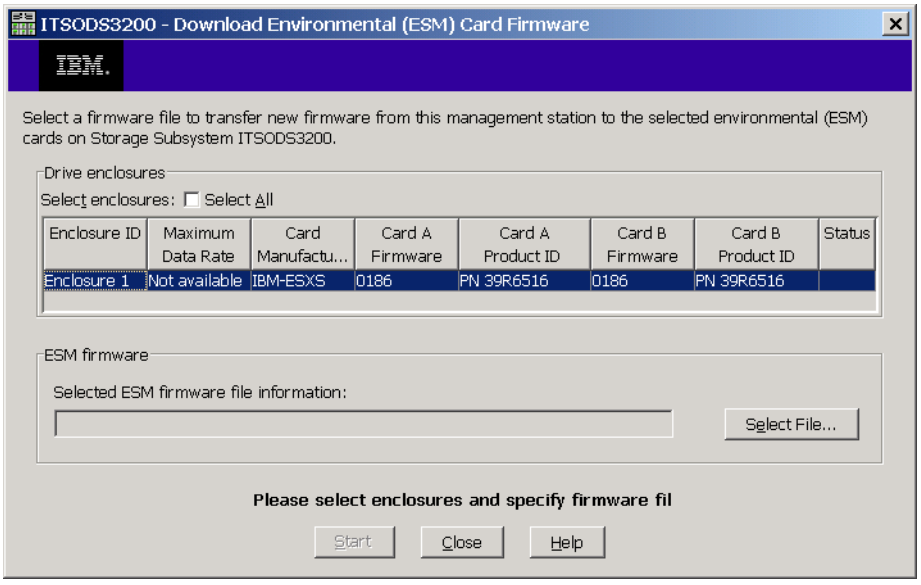


Figure 14-14 Download Environmental (ESM) Card Firmware

2. There are two sections in this view. The top section, *Drive enclosures*, lists all the installed enclosures. You select which enclosure(s) to update. Check the **Select All** check box to update all connected enclosures.

   In the second section, *ESM firmware*, select the ESM firmware file you have previously downloaded from the IBM Support Web page. Click **Select File** and select the firmware file.

   Click **Start** to update the ESM firmware (Figure 14-15).

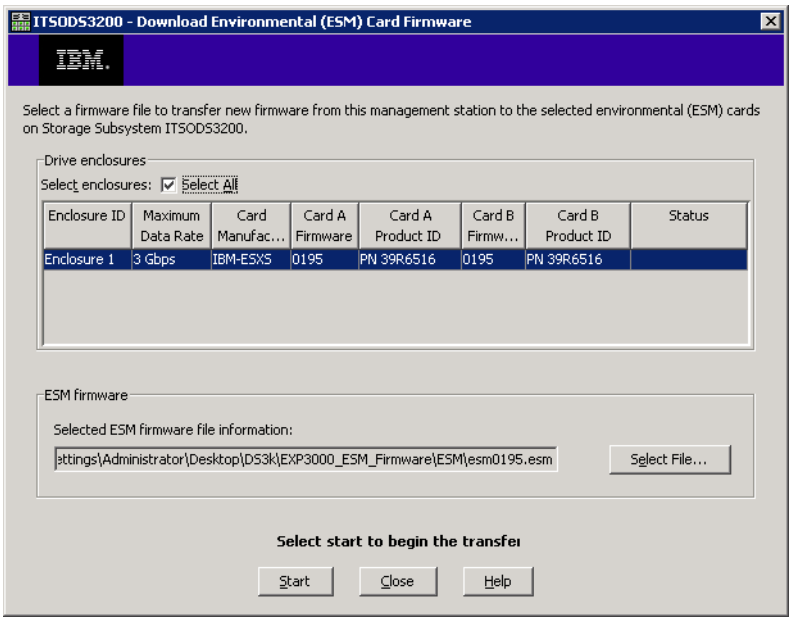


Figure 14-15 Download Environmental (ESM) Card Firmware - Start

3. The Confirm Download window appears. Read the information, confirm it by typing in yes, and click **OK** to start the update (Figure 14-16).

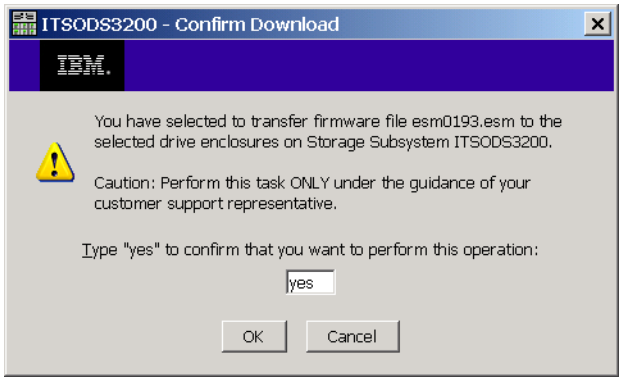


Figure 14-16 Confirm Download

4. The status field in the row of selected enclosures changes from pending to downloading while the ESM card firmware upgrade progresses. Monitor the progress and completion status of the download.

Download Controller Firmware

Note: If the Controller Firmware and NVSRAM upgrade from Version 6.xx to Version 7.xx is required, you must use the special utility (see 14.3.2, “Upgrade using the DS3000 Controller Firmware Upgrade Tool” on page 324). The utility is initially bundled in the DS3000 Storage Manager V10.35 software package as a separate application.

Do the following steps:

1. To update the controller firmware using the DS3000 Storage Manager, click **Download Controller Firmware**, as shown in Figure 14-12 on page 309. The window shown in Figure 14-17 appears.

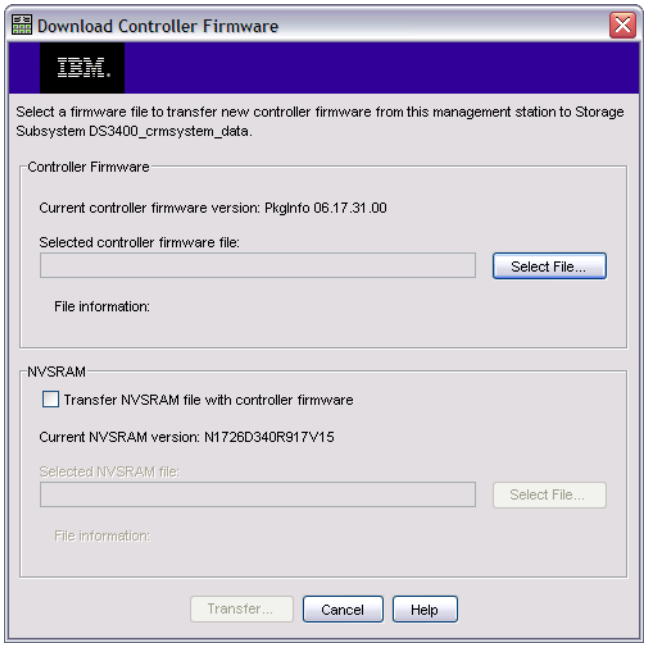


Figure 14-17 Download Controller Firmware

2. This window has two sections. The first, *Controller Firmware*, has a file check box for the controller firmware file. As previously mentioned, you must update NVSRAM after updating the controller firmware. In the second section, *NVSRAM*, you can also select the NVSRAM firmware file to update it together with the controller upgrade. Both upgrades will be done in one step.

3. Click **Select File** in the Controller Firmware section and enter or browse to the file name (Figure 14-18).

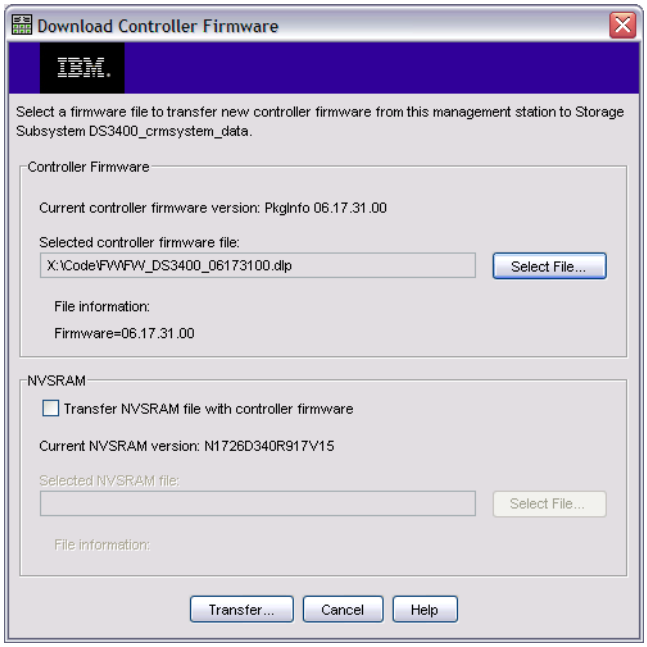


Figure 14-18 Download Controller Firmware - Select controller file

4. To update the NVSRAM in the same step, select the **Transfer NVSRAM file with controller firmware** check box and browse to the NVSRAM file (Figure 14-19 on page 315).

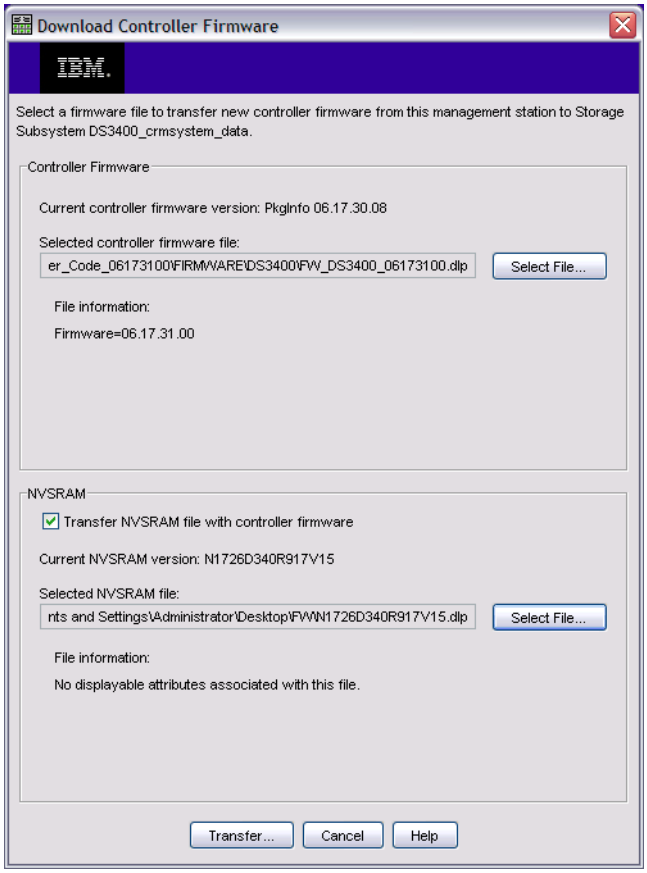


Figure 14-19 Download Controller Firmware - With NVSRAM firmware

5. Click **Transfer** to submit the new microcode to the storage subsystem. The Confirm Download window appears (Figure 14-20). Read it and click **Yes** to continue.

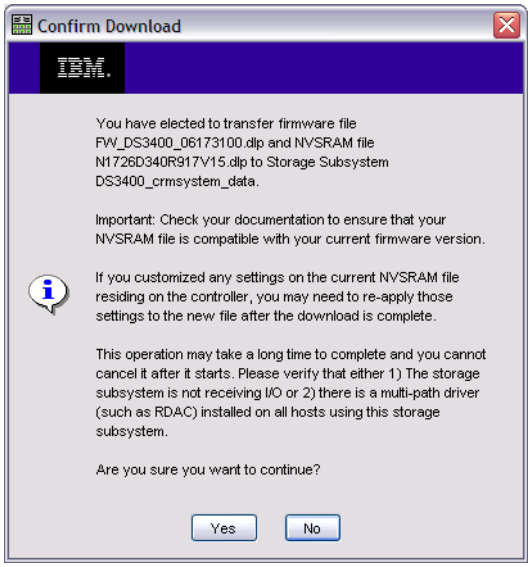


Figure 14-20 Confirm Download

6. The microcode will now be transferred to the storage subsystem and applied to both installed controllers. This may take some time and you can monitor the status in the Downloading window that appears (Figure 14-21).

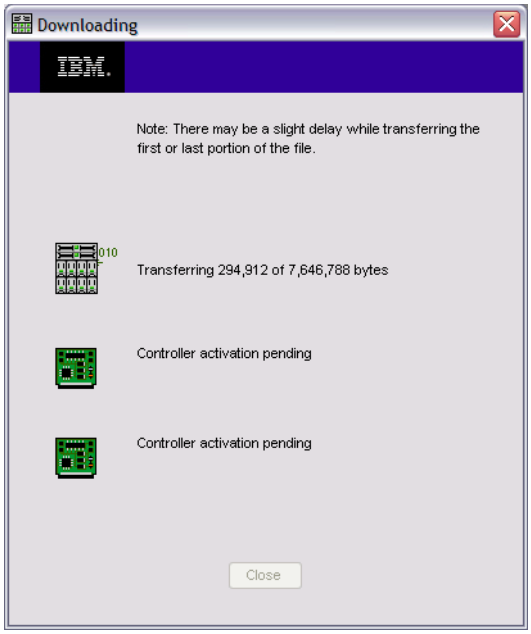


Figure 14-21 Downloading

7. After each task has completed successfully (Figure 14-22), click **Close** to return to the Download Firmware window (Figure 14-12 on page 309).

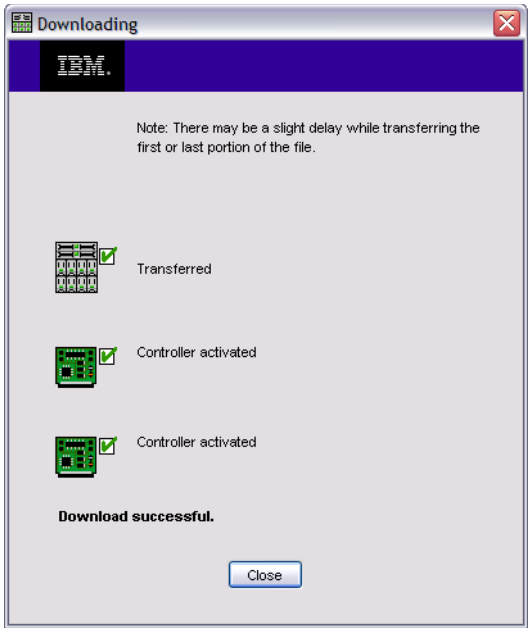


Figure 14-22 Download successful

Download Controller NVSRAM

Follow these steps:

1. To download a new NVSRAM file to your storage subsystem, click **Download Controller NVSRAM** (Figure 14-12 on page 309) to open the associated window (Figure 14-23). Click **Select File** to select the new NVSRAM file.

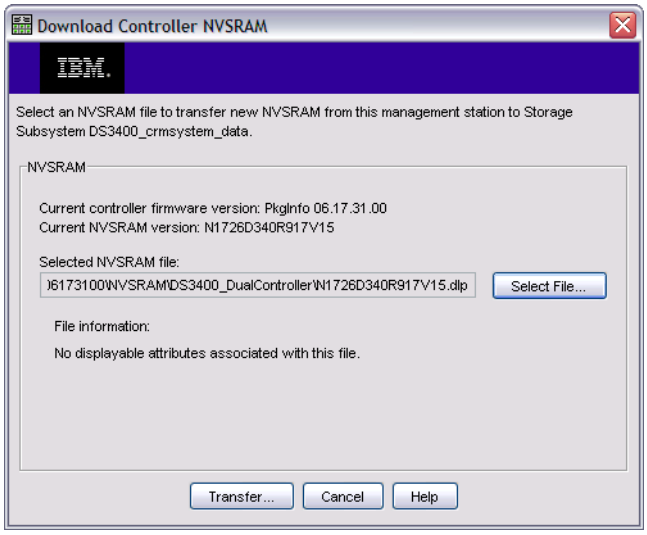


Figure 14-23 Download Controller NVSRAM

2. Click **Transfer** to submit the new NVSRAM file to the storage subsystem. The Confirm Download window appears (Figure 14-24). Read the information and click **Yes** to continue.

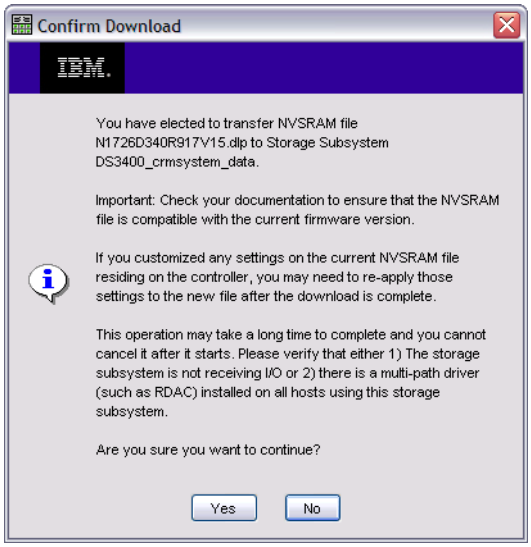


Figure 14-24 Confirm NVSRAM Download

3. The Downloading window appears and shows you the status of the transfer and the activation on the controller(s) (see Figure 14-25).

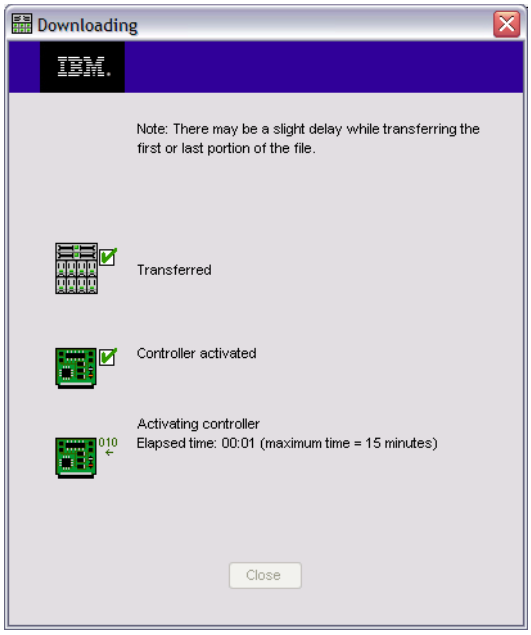


Figure 14-25 Download Progress NVSRAM

4. After the file is transferred and activated (Figure 14-26), click **Close** to return to the Download Firmware view (Figure 14-12 on page 309).

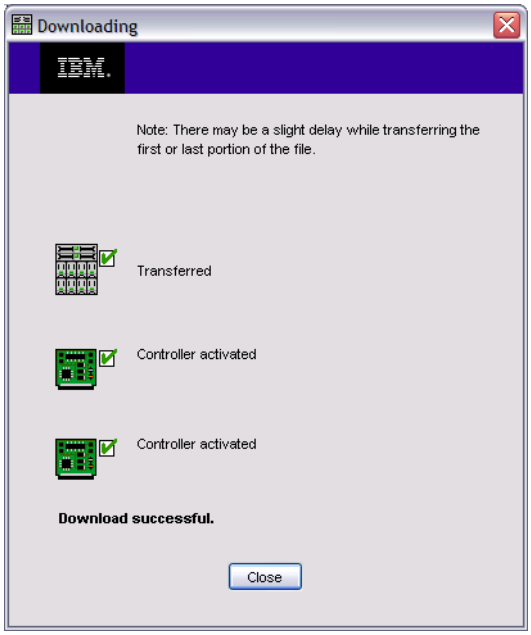


Figure 14-26 Download NVSRAM successful

Download Drive Firmware

Follow these steps when I/O operations are suspended:

Note: The hard drive upgrade process is disruptive, so be sure all host I/Os are suspended. Plan a production outage in advance.

1. Click **Download Drive Firmware** in the Download Firmware (Figure 14-12 on page 309) menu to load the new drive firmware. You can see the firmware levels of installed disks in the upper pane of the window, as shown in Figure 14-27. Click **Add** to select the appropriate drive firmware image file.

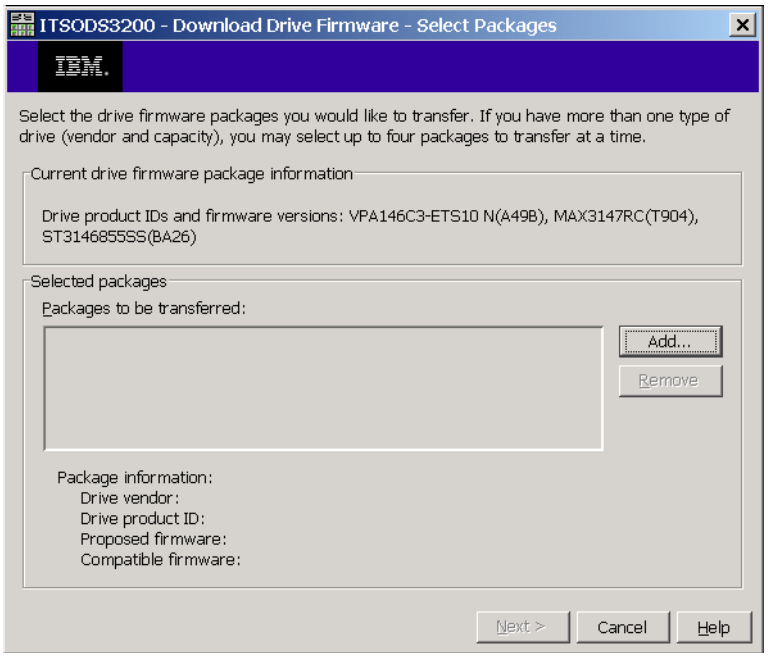


Figure 14-27 Download Driver Firmware

2. Select the drive firmware package (Figure 14-28). It can be found by the “Drive product ID” of the disk drive in the information pane of the window and by the appropriate name of the package file in the downloaded bundle (normally, just one big bundle for all hard disk types and all ESMs placed on the IBM Support Web site). Verify the package compatibility with the installed disk drives in the upper pane of the window (Figure 14-28).

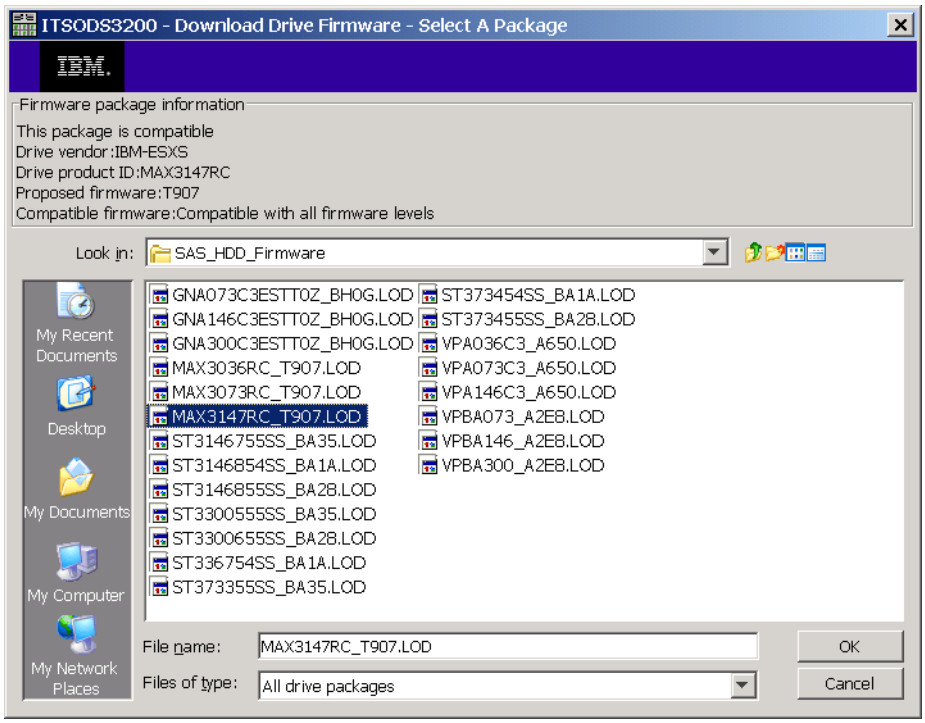


Figure 14-28 Select drive firmware package

3. Click **OK** to select the compatible firmware file only and return to the Download Firmware Window (Figure 14-27 on page 319).

4. Repeat these steps for every drive type you want to update. You can update up to four drive types at a time. Click **Next** in the Download Firmware window (Figure 14-29).

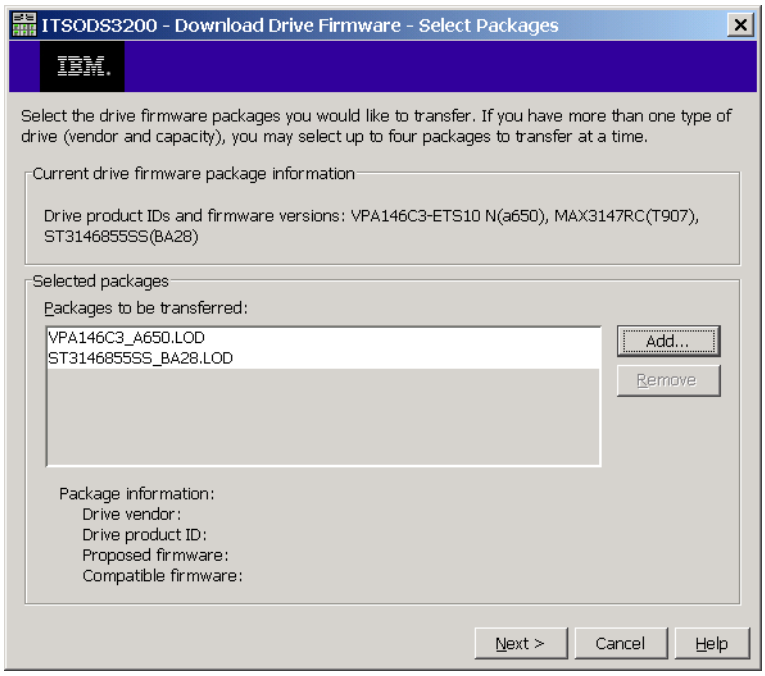


Figure 14-29 Drive packages added

5. Now you have to select the drives in the subsystem you want to update (Figure 14-30). You see the compatible drives only in this window. Check **Select all** to update all compatible drives.

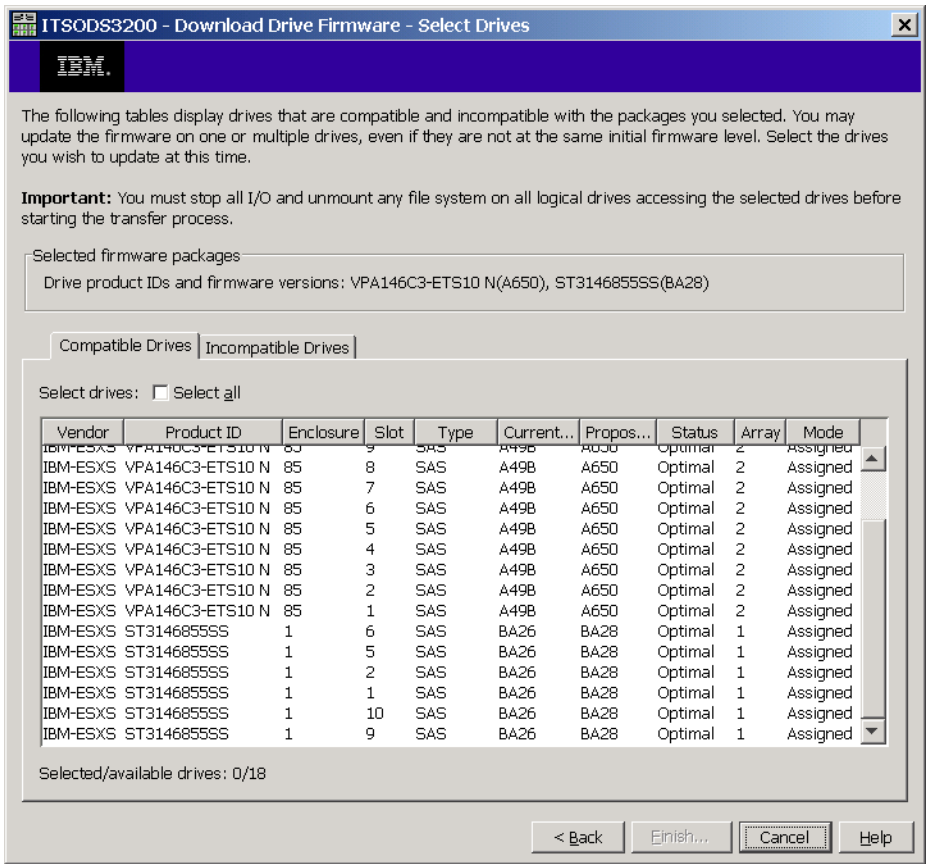


Figure 14-30 Compatible drives which can be updated

6. The drives that are not compatible for the selected updates are shown on the Incompatible Drives tab (Figure 14-31 on page 323). These are other drive types, and you will have to use other firmware packages.

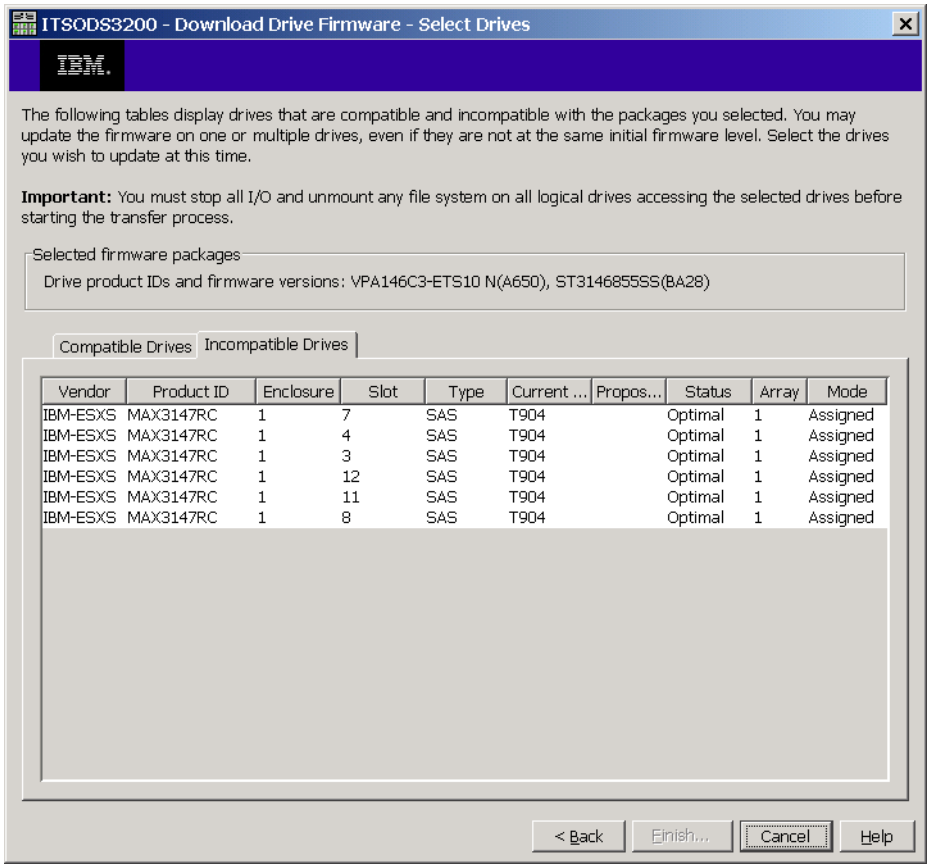


Figure 14-31 Incompatible drives

7. Click **Finish** to update the drive firmware. A confirmation window appears (Figure 14-32). Type yes and click **OK** to start the update.

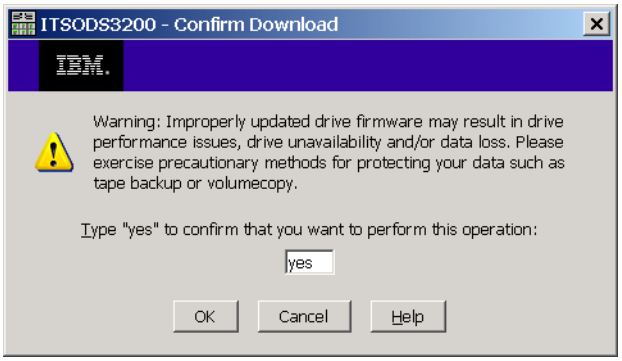


Figure 14-32 Confirm Download

8. After the update completes, a summary is displayed (Figure 14-33). If you want to update more drives, click **Transfer More**; otherwise, click **Close** to exit the drive update window.

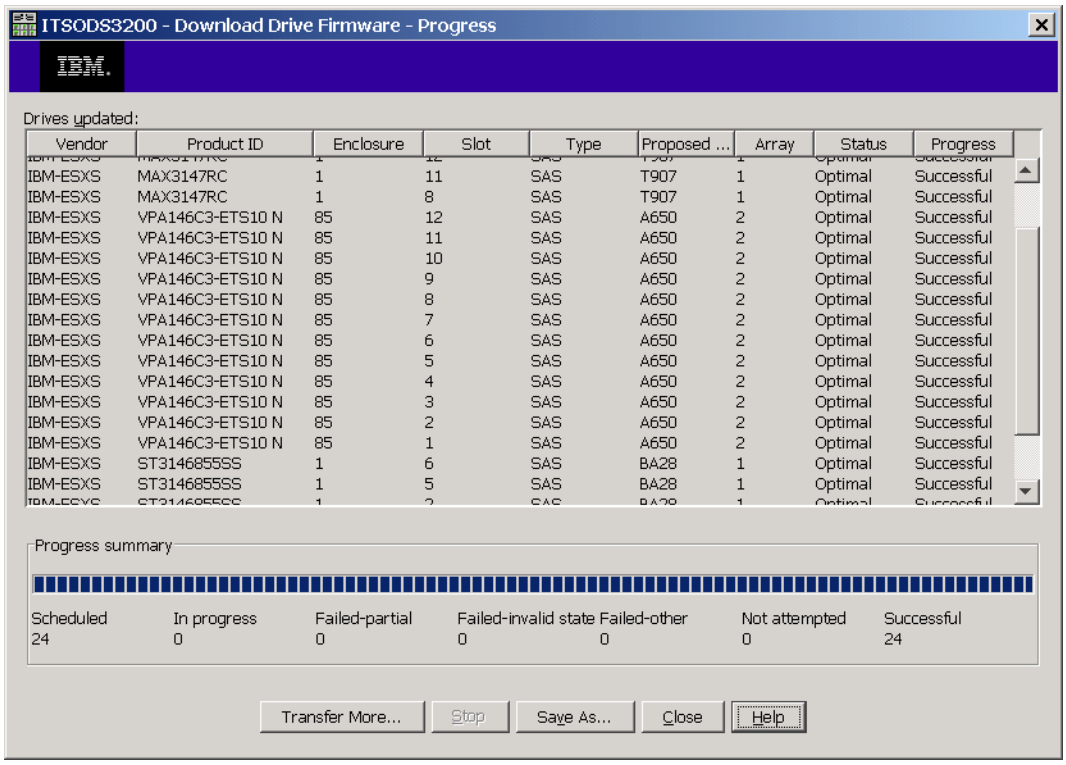


Figure 14-33 Drive update complete

14.3.2 Upgrade using the DS3000 Controller Firmware Upgrade Tool

Upgrading from firmware Version 6.xx to Version 7.35 or later cannot be done from the DS3000 Storage Manager client interface. You must use the provided special firmware upgrade tool. Before using the upgrade utility, make sure that you meet all of the following conditions:

- This utility is used to upgrade from a minimum 6.17.31.00 release to V7.35 firmware or later. If you are still using a version older than 6.17.31.00, you must first upgrade to the latest 6.xx code available from the IBM Support Web site.
- If you are already at any release of 7.xx firmware, further firmware upgrades should be done from within DS3000 Storage Manager as usual.

Note: First introduced with DS3000 Storage Manager V10.35, the automatic code synchronization feature ensures that both controllers are executing the same controller firmware version, preventing mismatched versions after controller replacement, or if an upgrade process is interrupted.

There are several important changes initially incorporated in controller firmware V7.35 that allow new functions and enhancements over previous versions. To accommodate those improvements, a rewrite of the DACstore region is necessary. This requires a non-concurrent update and the use of the special controller firmware upgrade utility.

Important: Upgrading from controller firmware V6.xx to V7.xx is a disruptive process (all host I/Os to the drives must be stopped, which means application downtime).

DS3000 Controller Firmware Upgrade Tool

This new utility was developed to upgrade DS3200, DS3300, and DS3400 systems running any controller firmware V6.xx later than 6.17.31.00 to the new firmware Version 7.35 or later.

The utility makes sure that the storage to upgrade is in an optimal state and that other conditions are met prior to performing the upgrade. The utility checks the compatibility for all files selected for download, and saves the recovery profile, the storage arrays configuration, the internal diagnostics output, and the Event log. This information is saved to the same directory where you chose to install the utility.

In addition to performing the upgrade of the controller firmware and NVSRAM, the utility also performs a migration of the drive configuration database region (DACstore) to a new format that can accommodate improvements introduced by firmware Version 7.35.

Installing the upgrade utility

To install:

1. Download the DS3000 Storage Manager V10.35 software bundle from the IBM Support Web site, along with the latest firmware (Versions 7.35). Note that new versions of Storage Manager can be introduced, so verify all information in the readme files on the IBM Support Web site.
2. Use the readme to also make sure that all restrictions, limitations, and recommendations for the utility (and DS3000 Storage Manager) are considered.
3. Start the installation of the DS3000 Storage Manger (as described in Chapter 6, “Installing IBM System Storage DS3000 Storage Manager” on page 81.) by running the installation package that you just downloaded. A window is displayed where you can select the language of your choice.
4. Read the Introduction and Copyright information window and accept the license agreement. You can select **Custom** installation and decide which DS3000 Storage Manager modules will be installed. The upgrade utility must be selected.
5. After a successful installation of the new DS3000 Storage Manager, the DS3000 Controller Firmware Upgrade Tool program is installed and accessible.

Upgrading the controller firmware and NVSRAM using the upgrade utility

Before you start, make sure that you have an adequate maintenance window to do the upgrade. The microcode for the DS3000 Storage Subsystem consists of two packages:

- Controller firmware
- NVSRAM

The NVSRAM is similar to the settings in the BIOS of a host system. The firmware and the NVSRAM are closely tied to each other and therefore are *not* independent. Be sure to install the correct combination of the two packages.

Important: Before upgrading the storage server firmware and NVSRAM, make sure that the system is in an optimal state. If not, run the Recovery Guru (see 14.6, “Recover from Failure” on page 347) to diagnose and fix the problem before you proceed with the upgrade. Always read the readme before upgrading any firmware to check for compatibility with your system.

Do the following steps:

1. Start the upgrade utility program. the utility asks for Automatic or Manual discovery of DS3000 subsystems. To add a subsystem manually, click **Add** and enter the DS3000 controller IP addresses or the IP address of the host used for in-band management of the DS3000. Once the DS3000 system is added successfully, the upgrade utility performs an automatic check for any non-optimal condition that would prevent the upgrade from succeeding. After the discovery process completes, a window similar to the one shown in Figure 14-34 appears.

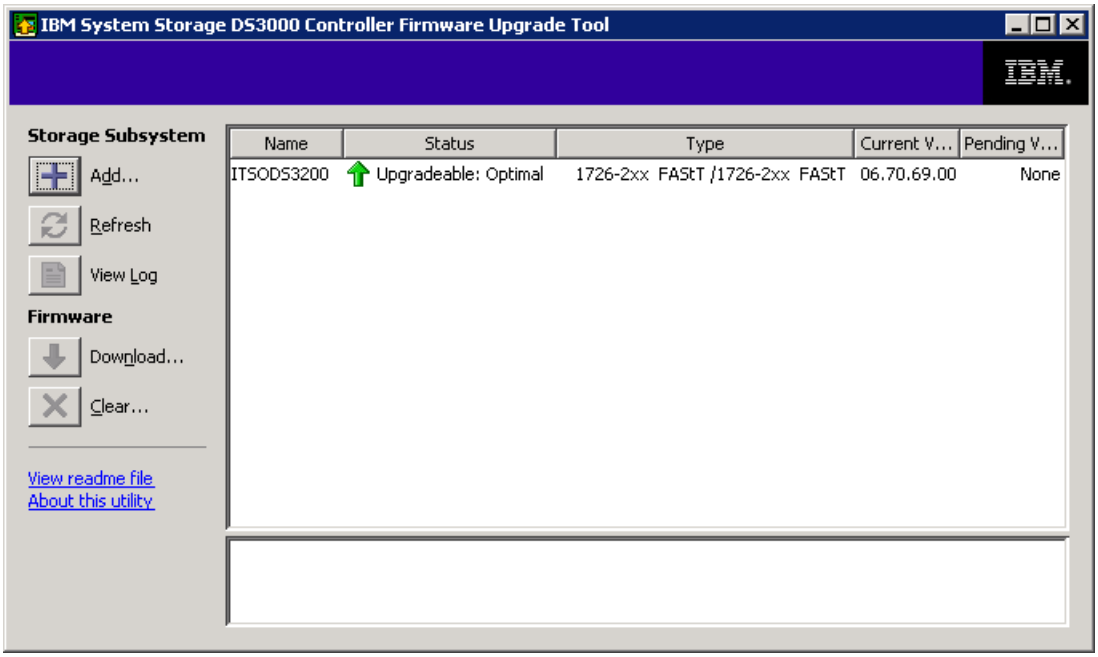


Figure 14-34 Firmware upgrade utility

You can add multiple DS3000 systems in the upgrade utility. However, we recommend installing just one at first, and then, if the operation was successful, continuing with others.

The Status column shows different values depending on the version found in the storage systems, since not all the automatically discovered models are compatible with DS3000 firmware V7.35 and later. Models not supported by this firmware continue to use the Storage Manager to install upgrades up to the latest level available for the particular model.

The only status that allows you to proceed with the firmware upgrade is Upgradeable:Optimal.

If the status is not optimal, you can use the **View Log** option to find the reason detected by the utility. The conditions checked by the utility for a system to be upgradeable and optimal are:

- Status upgradeable
  - If the DS3000 storage subsystem supports Version 7.xx.
  - If the current version is equal to or later than V6.17.31.00 and earlier than V7.35 (already upgraded to).
- Status optimal
  - If no failed assigned drives are found.
  - If no hot spares replacing failed drives are found.

- If there are no missing volumes.
- If controllers are in Optimal status.
- If there are no operations in progress.

If there is a problem with any of the above conditions, the utility program does not allow the upgrade of the affected DS3000 subsystem. Other non-optimal conditions are indicated, and we recommend fixing them before proceeding with the upgrade (not fixing those, however, does not prevent the upgrade utility from performing the upgrade).

Important: Make sure to install the firmware for each of the DS3000 components (ESM, drives, controller, and NVSRAM) in the sequence described in the readme file for that version.

Update the controller firmware and then the NVSRAM, or both at the same step.

Any power or IP-network/SAN interruption during the update process may lead to configuration corruption or extended downtime. Therefore, do not power off the DS3000 storage server or the management station during the update. If you are using in-band management and have Fibre Channel hubs or managed hubs, then make sure that no SAN-connected devices are powered up during the update; otherwise, this can cause a loop initialization process and interrupt the process.

2. Click the **Download** button in the left margin (Figure 14-34 on page 326). This brings up the windows shown in Figure 14-35 and Figure 14-36 on page 328, where you can select the controller firmware source file previously downloaded.

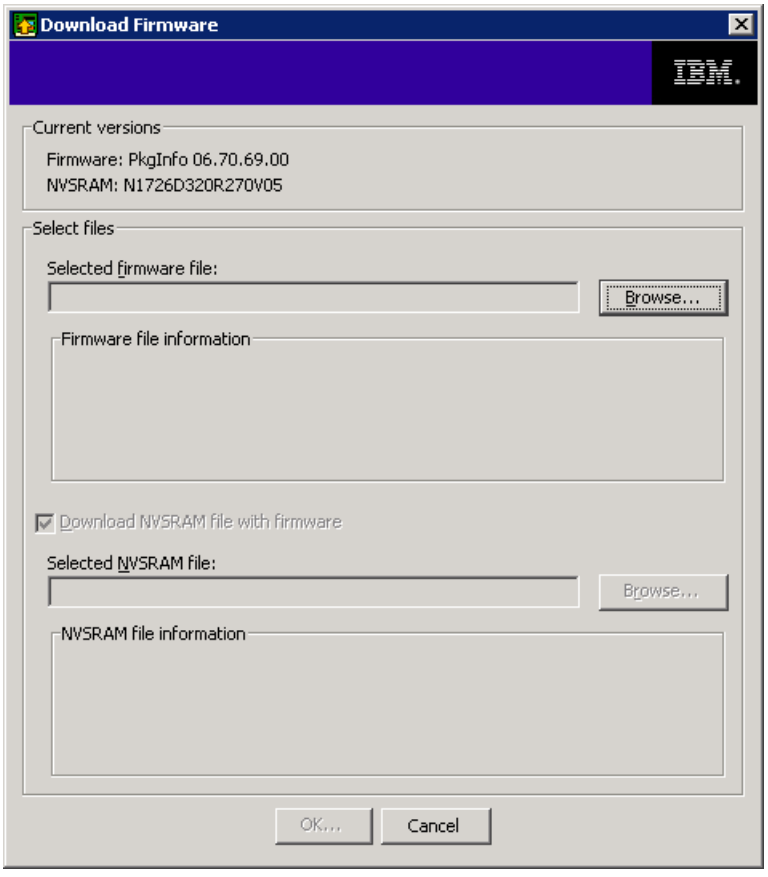


Figure 14-35 Click Browse to select files for download

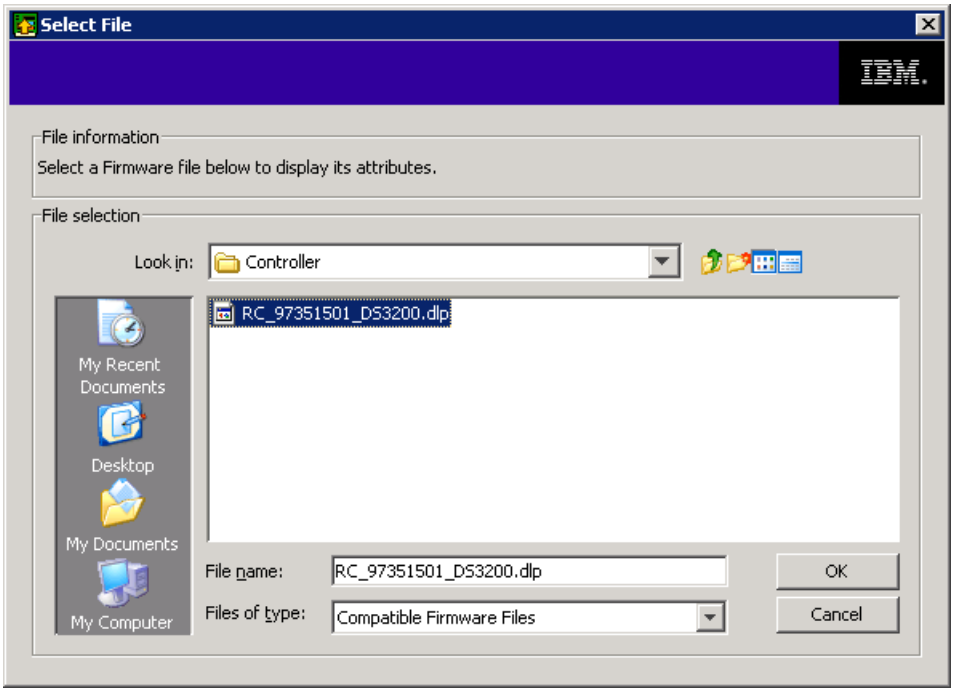


Figure 14-36 Select firmware source files

> **Note:** The upgrade utility windows shown are the same as the windows obtained when updating through the DS3000 Storage Manager, although the utility and background processes are different.

3. Check the **Download NVSRAM file with firmware** check box, and select the NVSRAM file downloaded previously using the same manner as for the controller firmware. Make sure that the NVSRAM selected is the one provided with the firmware file being downloaded, and for your specific DS3000 model. Notice that the NVSRAM for single and for dual controller models of DS3000 are different.

   You can do firmware and NVSRAM upgrades in separate instances, but since their versions need to match, we recommend installing both at the same time to make sure that you have a compatible pair of firmware and NVSRAM files.

   > **Note:** Remember that upgrading from firmware V6.xx to V7.xx requires no I/O traffic, because both controllers go offline at the same time to allow the upgrade (even if your hosts are configured with a driver providing redundancy).

4. Click **OK** to proceed when the controller firmware and NVSRAM files are selected in the window shown in Figure 14-35 on page 327. The warning window shown in Figure 14-37 appears, and you can confirm the upgrade by clicking **Yes**.

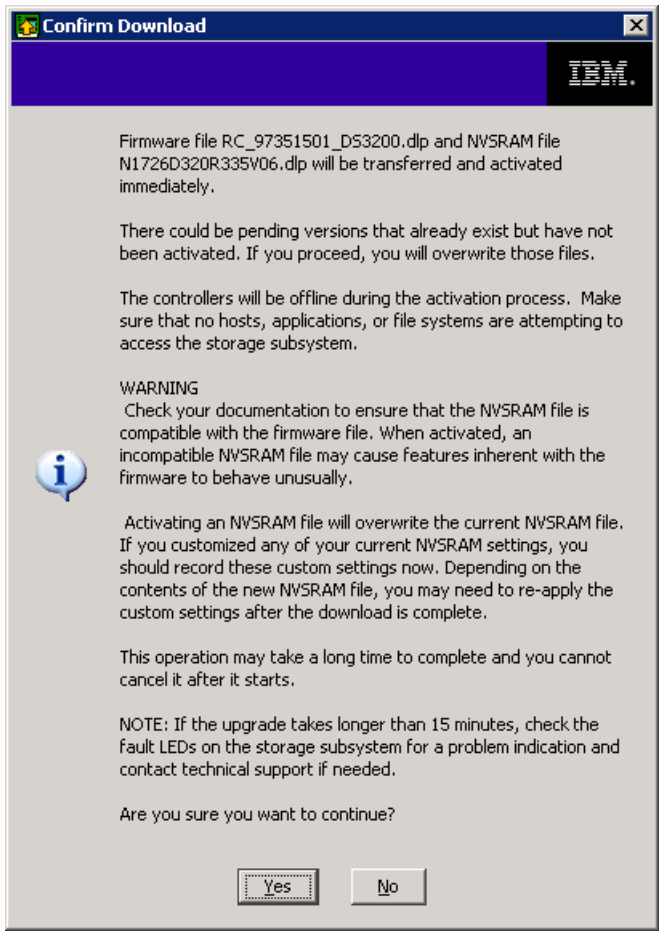


Figure 14-37 Confirming download

5. Observe the progress during the file transfer and activation. It can take a long time (25 to 30 minutes), but you can monitor the process in View Log. Once finished, the utility indicates the result of the operation. If the transfer and activation were successful, the window shown in Figure 14-38 appears.

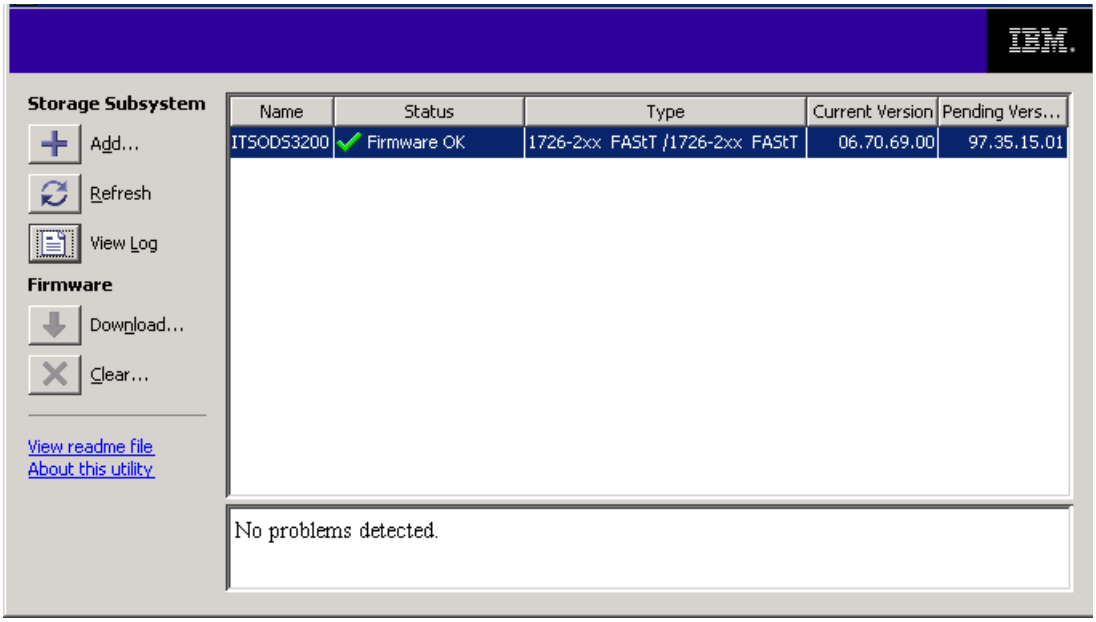


Figure 14-38 Firmware upgraded

If the operation finishes with an error, the utility displays the failure details in the notification field. Check for additional details for all processes using the **View Log** button.

14.4 Manage Controllers

In this menu, you can place DS3000 controllers online or offline, or redistribute logical drives to their preferred controller.

Click the **Manage Controllers** link in the Support tab in Figure 14-1 on page 300 to open the Manage Controllers window (Figure 14-39 on page 331).



Figure 14-39 Manage Controllers

14.4.1 Place Controller Online or Offline

Follow these steps:

1. Click **Place Controller Online or Offline**. The window in Figure 14-40 appears.

Note: You should only place a controller offline if advised by the Recovery Guru or by IBM technical support. Placing a controller offline makes it unavailable for I/O operations.



Figure 14-40 Place Controller Online or Offline

2. In the Controller drop-down menu, select the controller. The opposite status to the controller's current status will be activated. Click **OK** to change the status. Read the warning in the confirmation window (Figure 14-41) and click **Yes** to proceed.

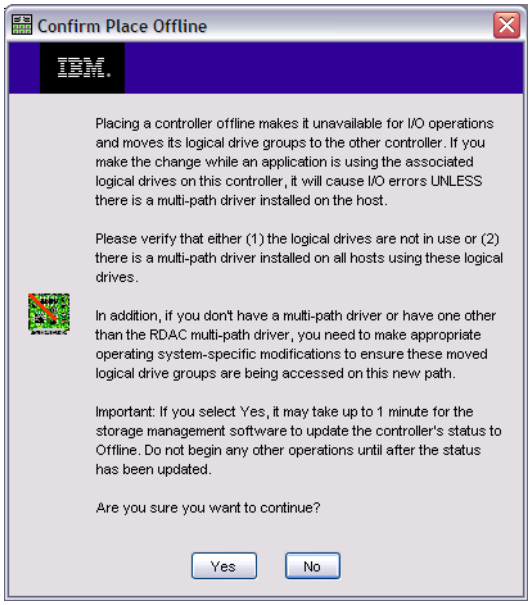


Figure 14-41 Confirm Place offline

3. In our example, Controller A is now set offline. Figure 14-42 appears when the operation completes. Click **OK** to return to the Manage Controllers window (Figure 14-39 on page 331).

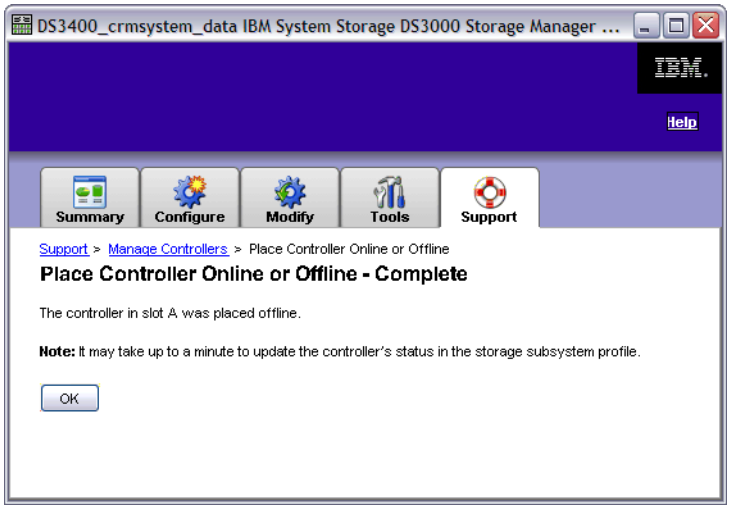


Figure 14-42 Place Controller Online or Offline - Complete

4. Follow the process to set the controller online again.

14.4.2 Redistribute Logical Drives

Each logical drive has its preferred path. If a problem occurs, the logical drives are moved automatically from the preferred controller to the other controller to maintain data accessibility. When the problem is corrected, you can move all the moved logical drives back to their preferred controller. Follow these steps:

1. Click **Redistribute Logical Drives** in the Manage Controllers window (Figure 14-39 on page 331) to start. The window shown in Figure 14-43 appears.

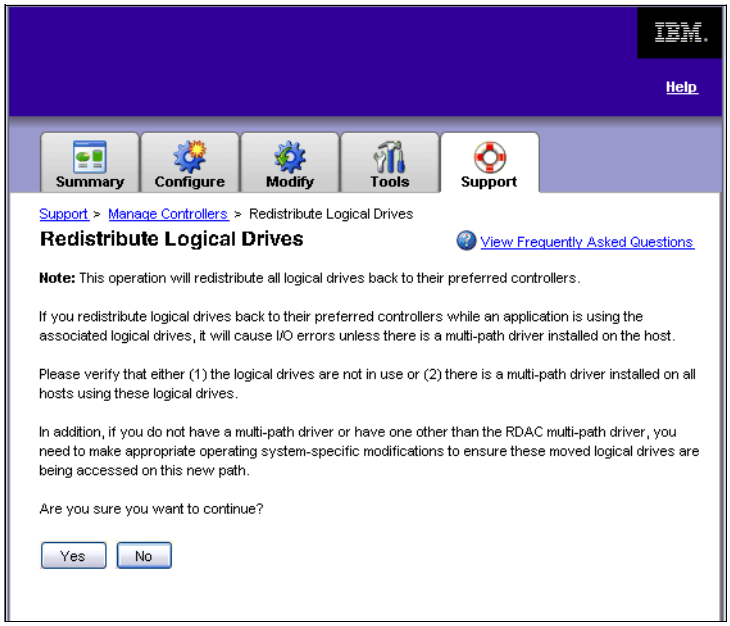


Figure 14-43 Redistribute Logical Drives

2. Read the information provided in this view and click **Yes** to start. The progress is displayed in the Redistribute Logical Drives window so that you can see when it is completed (Figure 14-44).

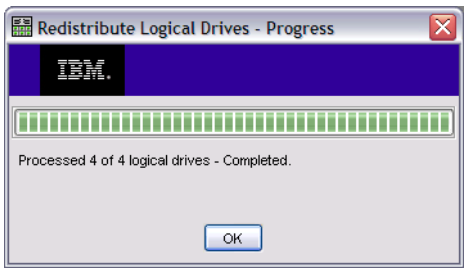


Figure 14-44 Redistribute Logical Drives - Progress

3. Click **OK** and you will get a confirmation that all logical drives have been redistributed successfully (Figure 14-45).

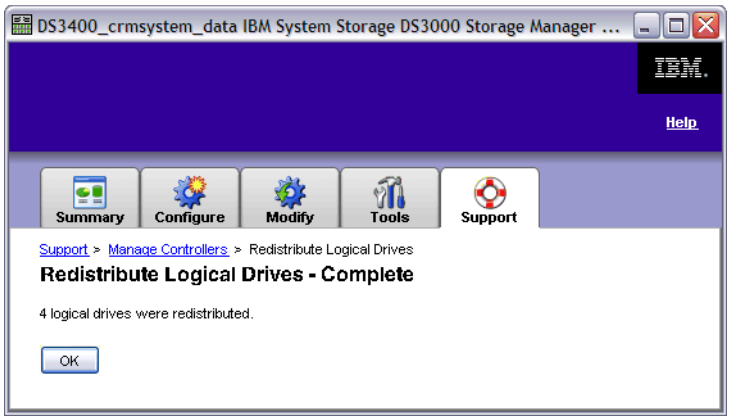


Figure 14-45 Redistribute Logical Drives - Complete

4. Click **OK** to return to the Manage Controllers window (Figure 14-39 on page 331).

14.5 Advanced Support Tasks

Warning: We recommend that you do not use these tasks unless instructed to do so by an IBM technical support representative.

The Advanced Support Tasks window (Figure 14-46 on page 335) offers the following options:

- View Event Log
- Import Array
- Export Array
- Change Remote Login Setting

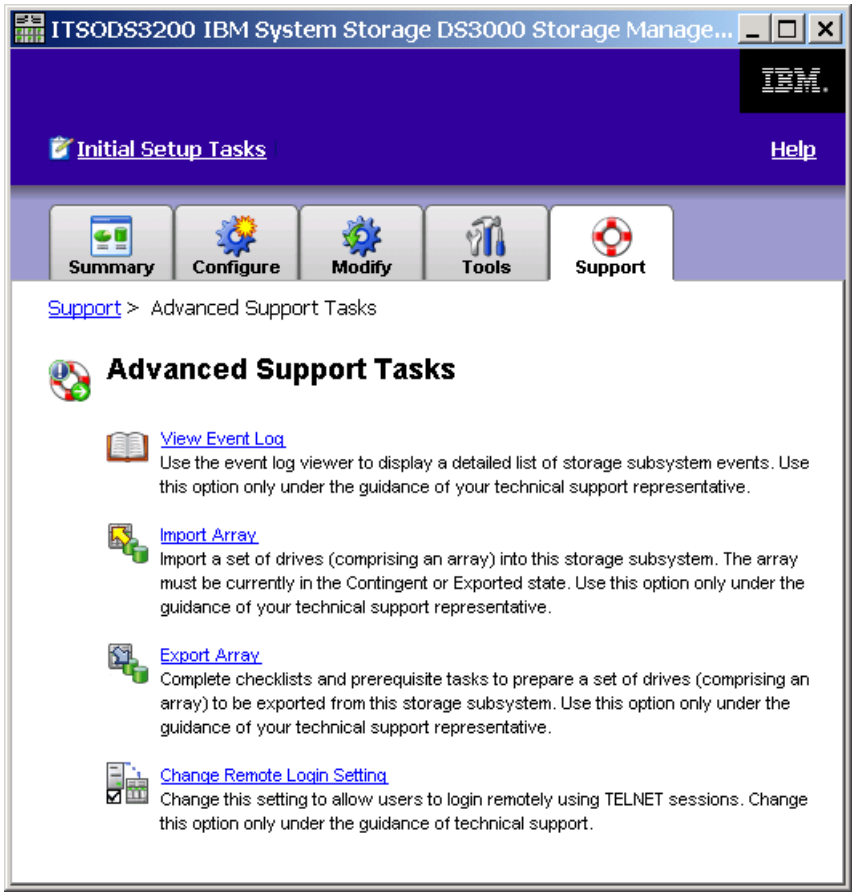


Figure 14-46 Advanced Support Tasks

14.5.1 View Event Log

Use this link to view a list of events that occurred on your storage subsystem. The event log is stored on reserved areas on the disk drives in the storage subsystem. It records configuration events and the failure of storage subsystem components. You can set the event log to show only critical events. All critical events are sent to the SNMP management console or to the e-mail recipient that you have configured to receive the alert notifications (as described in “Configure Alerts” on page 118).

The event log shows a summary for each event in the list, including:

- Date/Time

  The date and time that the event occurred. The date and time are set from the controllers’ clocks.

- Priority: Critical or Informational

  - A critical event shows a problem with the storage subsystem. Critical events should be investigated immediately, in order to maintain access to the data.
  - An informational event shows information about the storage subsystem. This information does not show a critical problem with the storage subsystem.

- Component Type

  The type of component that is affected by the event. The component type might indicate a physical component, such as a disk drive, or a logical component, such as a logical drive.

- Component Location

  The location of the component that is affected by the event. The location shows the enclosure or slot for the component.

- Description

  A short description of the event.

You can also view details for each event. The event details shows additional information about the event. The detail information covers the following information:

- Date/Time

  The date and time that the event occurred. The date and time are set from the controllers' clocks.

- Sequence Number

  A 64-bit number that is unique for each event in the log for a storage subsystem. This number increments by one with each new event.

- Event Type

  A four-digit number that identifies each type of event.

- Description

  A short description of the event.

- Event Specific Codes

  A specific code for the event, which could be requested by technical support.

- Event Category
  - Failure: A physical device in the storage subsystem has failed. For example, a disk drive or battery has failed.
  - State Change: A logical device in the storage subsystem has changed status. For example, an optimal logical drive has been marked offline.
  - Internal: These are internal controller operations that do not require your action.
  - Command: A command has been issued to the storage subsystem. For example, a hot spare disk drive has been assigned.
  - Host: The host software has posted the entry to the Event Log.
  - Error: An error condition has been found on the storage subsystem.
  - General: Any event that is not described by one of the other categories.

- Component Type

  The type of component that is affected by the event. The component type might indicate a physical component, such as a disk drive, or a logical component, such as a logical drive.

- Component Location

  The location of the component that is affected by the event. The location shows the enclosure or slot for the component.

- Logged by

  The name of the controller that logged the event.

- Raw Data

  Information that is used by a IBM Support representative.

View the Event Log

Click **View Event Log** (Figure 14-46 on page 335) to display the events in the internal log of DS3000 subsystem (Figure 14-47).

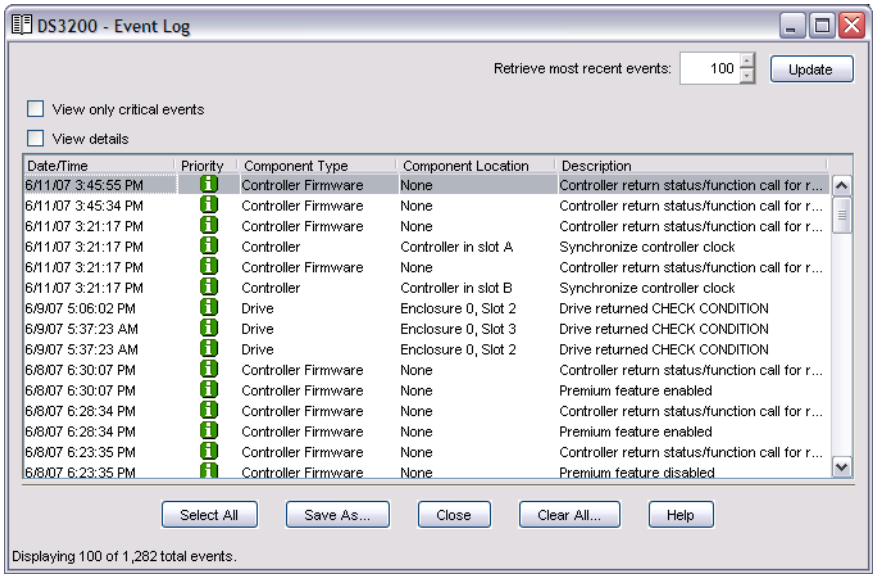


Figure 14-47 Event Log

In the right upper corner, you can set how many events to display; the default initial value is 100. To change this, enter another value and click **Update**. The bottom line of the window shows the total number of events stored on the subsystem. The maximum number of event is 8190. If more events occur, the old events are removed from the log on a first in, first out basis.

In the left upper corner are two check boxes:

- View only critical events

  Check this to display only the critical events, as shown in Figure 14-48.

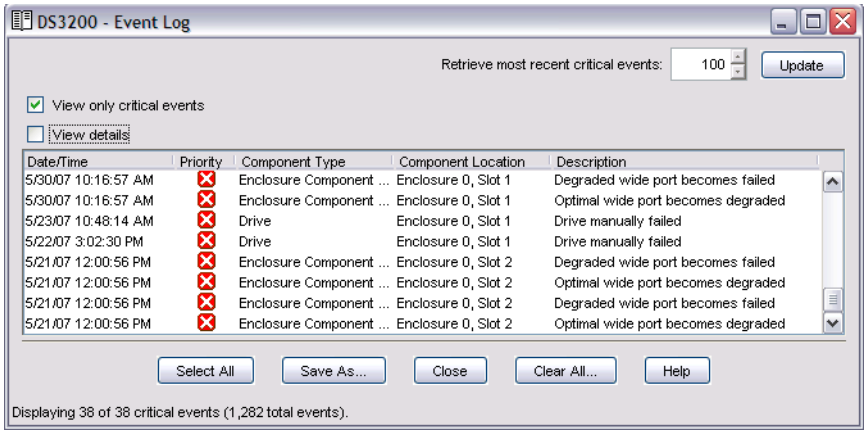


Figure 14-48 Event Log - Show critical events only

- View details

  Check this to show details of a selected event, as shown in Figure 14-49.

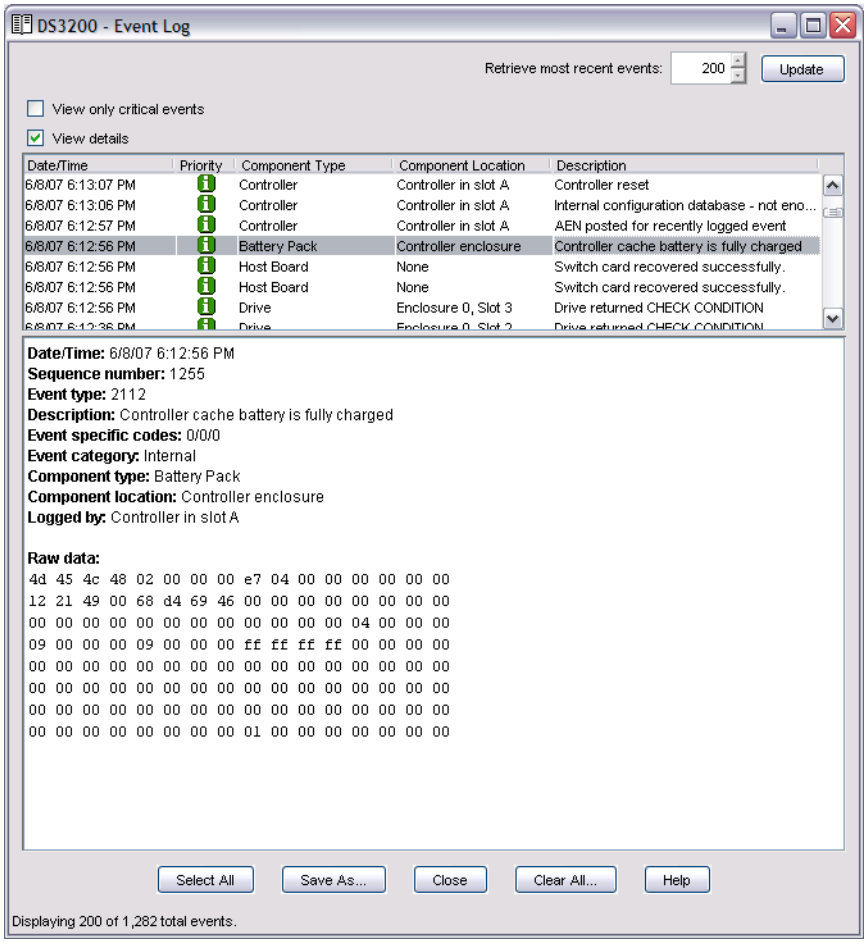


Figure 14-49 Event Log - Details

The buttons at the bottom of the window are:

- Select All

  Use this button to highlight all events.

- Save As

  This button saves the highlighted events. To save several entries, use the Ctrl key to select them; to save all entries, use the **Select All** button before you click **Save As**. You will be prompted for a location to save the events (the default extension is *.log).

Figure 14-50 on page 339 shows a sample log file.

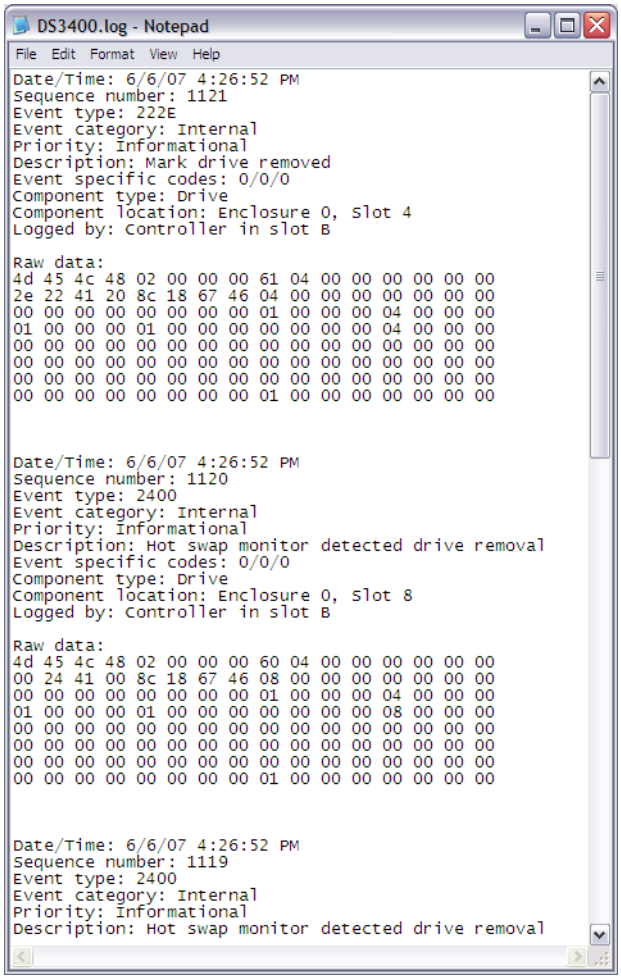


Figure 14-50 Event view - Log

- Close

  Closes the Event view

- Clear All

  Select this option to clear the event log. In the confirmation window (Figure 14-51), type yes to delete all the events. We recommend that you save all event logs before clearing the log.

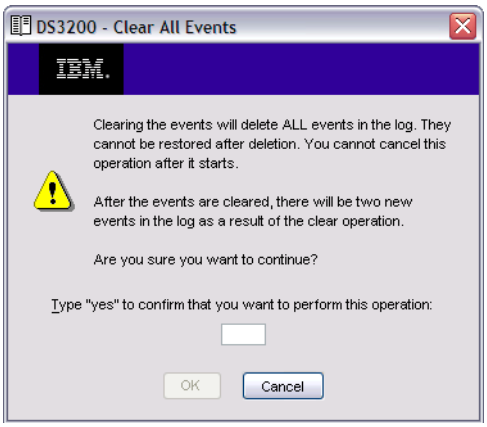


Figure 14-51 Event view - Confirm Clear All Events

- Help

  Opens the online help.

14.5.2 Import/Export Array

The DS3000 firmware 7.35 initially incorporates new export and import array options to safely move arrays (hard disk sets with a given RAID level and defined LUNs) between different DS3000 systems without losing customer application data.

This is very helpful when you have to upgrade or replace a DS3000 system with a new model or faster disks, but want to preserve the expansions and their data.

The export/import options check that all the conditions to support the disk migration are met before placing the disks offline and allowing removal of the disks.

Now, instead of using the option to place an array offline, as was the case with previous versions of DS3000 Storage Manager, just select the **Export Arra**y option from the source machine. Select the **Import Array** option on the destination machine to accept the exported disks with their data.

Important: We recommend upgrading both the source and the destination DS3000 subsystems to the latest firmware available before attempting any disk migration. Never attempt to migrate from a DS3000 running controller firmware Version 7.xx to a lower-level subsystem running Version 6.xx.

Remember that starting with Version 7.xx, the format of the DACstore region on the hard disks has changed significantly and may generate inconsistencies and loss of data if imported into a lower-level subsystem.

Export Array

The export array operation prepares the disk drives in the array for removal. You can remove the disk drives for offline storage, or you can import the array to a different storage subsystem. After you complete the export array operation, all of the disk drives are offline.

Any associated logical drives or free capacity are no longer shown in the storage management software.

Non-exportable components

You must remove any non-exportable components before you can complete the export array procedure.

Export Array procedure

Follow these steps:

1. You must remove any non-exportable components before you can start the export array procedure. You must remove these components:
   a. Host-to-logical drive mappings.
   b. VolumeCopy pairs.
   c. FlashCopy logical drives and FlashCopy repository logical drives.
2. Exporting an array includes following these steps on the source storage subsystem:
   a. Save the storage subsystem configuration.
   b. Back up the data on the logical drives in the array.
   c. Stop all I/O, and unmount or disconnect the file systems on all the logical drives in the exported array.
   d. Locate the array, and label the disk drives.
   e. Obtain blank disk drive CRUs or new disk drives.
3. Exporting an array includes these steps on the target storage subsystem:
   a. Make sure that the target storage subsystem has available disk drive slots.
   b. Make sure that the target storage subsystem supports the disk drives that you will import.
   c. Make sure that the target storage subsystem can support the new logical drives.
   d. Make sure that the latest version of firmware is installed on the controller.
   e. Make sure the target storage subsystem supports RAID 6 if you are exporting a RAID 6 array.
   f. Make sure that the latest version of the storage manager software is installed for the target storage subsystem.

4. Activate the **Export Array** task from Advanced Support menu (Figure 14-46 on page 335) and select the array that you want to export in the Introduction window (Figure 14-52).

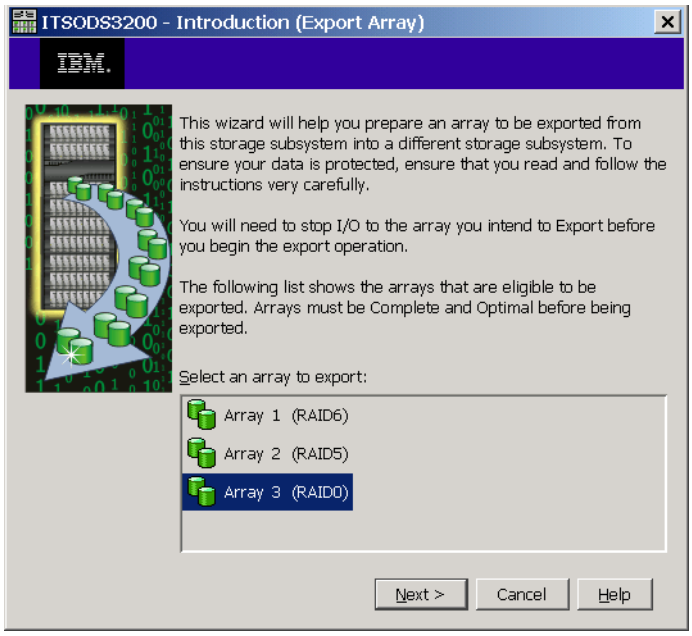


Figure 14-52 Export Array - Select array

5. All the previously mentioned requirements for export of an array (steps 2 and 3) must be confirmed by selecting all check boxes in the Preparation Checklist window (Figure 14-53), and then clicking **Finish**.

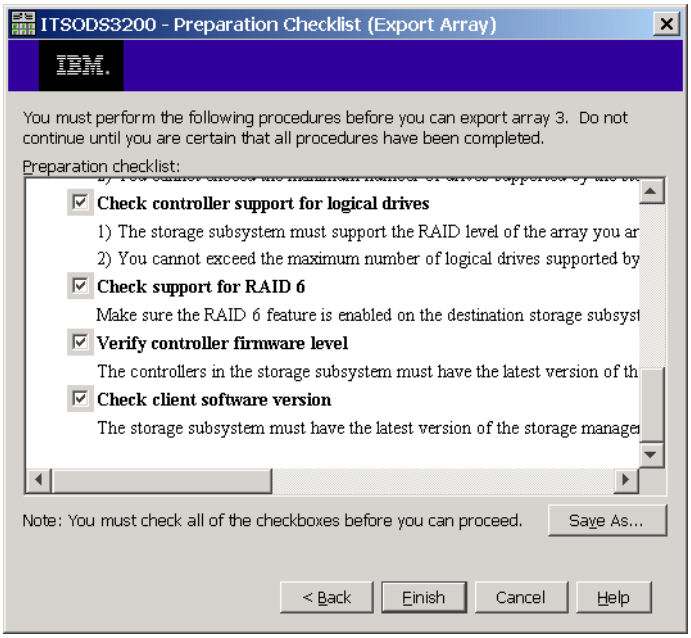


Figure 14-53 Export array - checklist

6. Confirm the Export Array procedure once again by typing yes in the confirmation window (Figure 14-54 on page 343) and clicking **OK**.



Figure 14-54 Export Array - Confirmation

7. After some period of time, the Export Completed window (Figure 14-55) confirms the successful completion of the Export Array procedure.

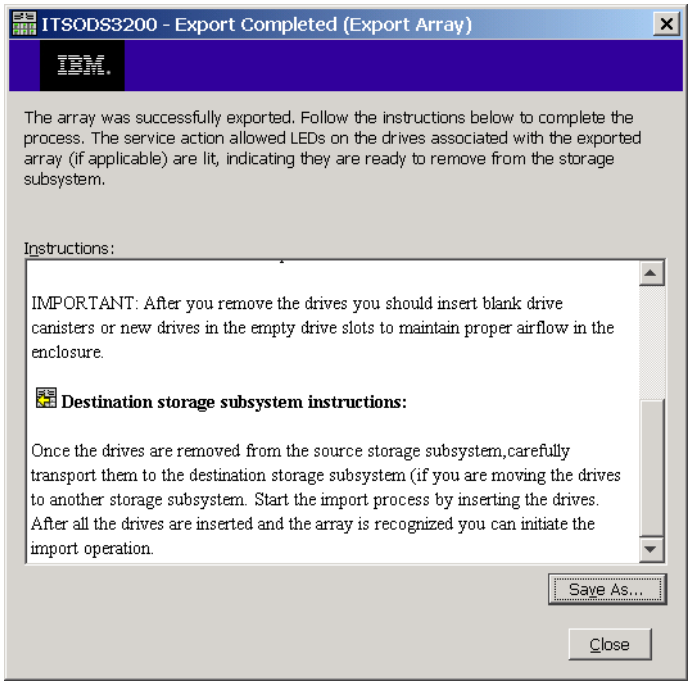


Figure 14-55 Export Array - Export Completed

8. The exported array in the source DS3000 has changed its icon, as shown in Figure 14-56. This window can be activated in the Summary tab by selecting the **Arrays & Logical Drives** link.

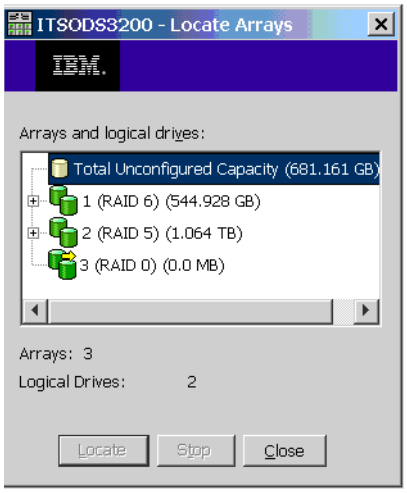


Figure 14-56 Show exported array

9. Finally, hard disks in the exported array may be carefully removed from the source DS3000 and may be inserted into the empty slots of the target DS3000. The imported array then will have the status “Exported - ready for import” in the target DS3000.

Import Array

The Import Array operation adds the imported array to the target storage subsystem. After you complete the import array operation, all of the disk drives have Optimal status. Any associated logical drives or free capacities are now shown in the storage management software installed for the target storage subsystem.

Important: You must insert *all* of the disk drives in the array into the target enclosure before the array can be imported.

Importing an array includes following these steps on the target storage subsystem:

1. Insert the exported disk drives into the available disk drive slots.
2. Review the Import Report for an overview of the array that you are importing.
3. Check for non-importable components.
4. Confirm that you want to proceed with the import procedure.

Non-importable components

Some components cannot be imported during the import array procedure. These components are removed during the import procedure:

- Host-to-logical drive mappings
- VolumeCopy pairs
- FlashCopy logical drives and FlashCopy repository logical drives

Import Array procedure

Do these steps:

1. Click the **Import Array** link in the Advanced Support Task window (Figure 14-46 on page 335). If imported disks are inserted, the Introduction window (Figure 14-57) appears.



Figure 14-57 Import Array

2. Click **Next** to start the Import Array procedure. The Import Report window (Figure 14-55 on page 343) shows some basic information about imported array and disks.

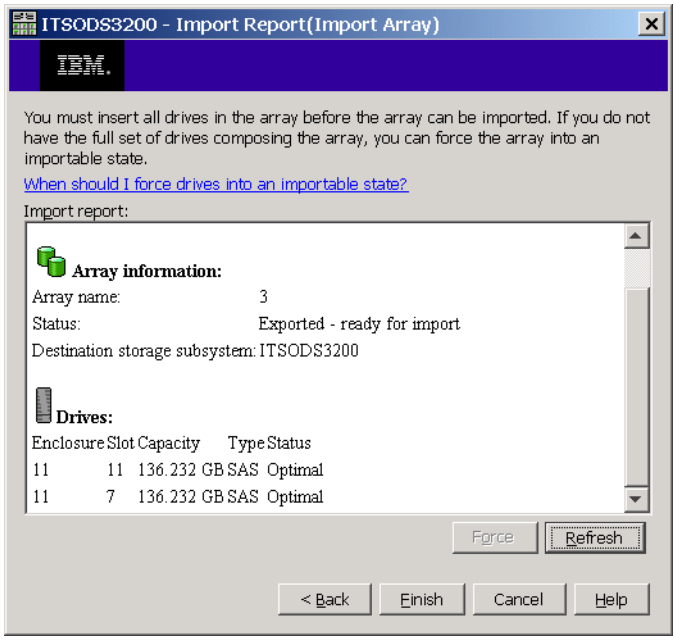


Figure 14-58 Import Array - Report

3. Click the **Finish** button. The window shown in Figure 14-59 appears. Give the final confirmation by typing yes and clicking **OK**.

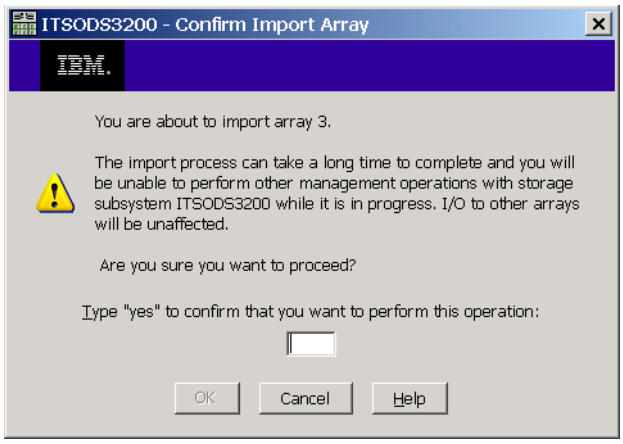


Figure 14-59 Import Array - Confirmation

4. The imported array is now visible in the window that appears when you click the **Arrays & Logical Drives** link in the Summary tab.

14.5.3 Change Remote Login Setting

This feature, shown in Figure 14-60, is initially introduced with DS3000 controller firmware Version 7.35 and can enable TELNET access to DS3000 management ports. TELNET access is disabled by default and, when required, it has to be manually enabled on both controllers.

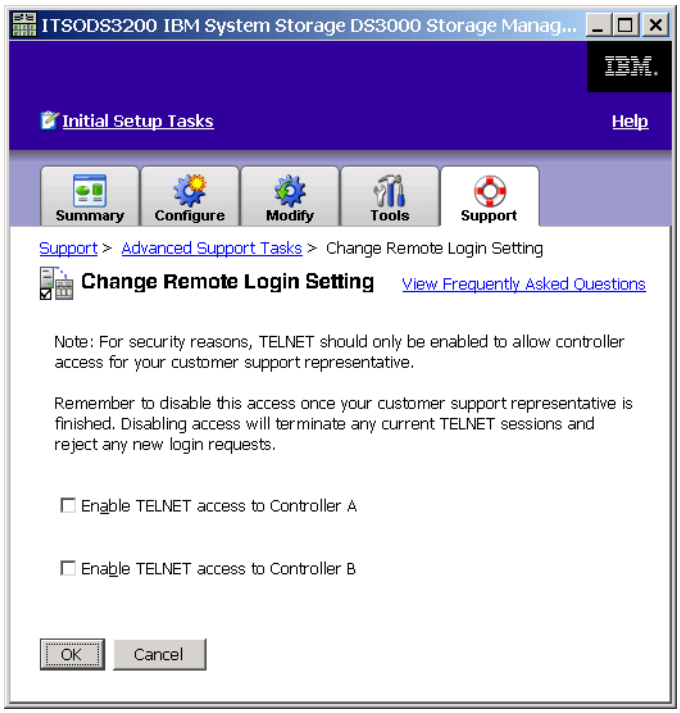


Figure 14-60 Change Remote Login Setting

Note: Leave the Remote Login option disabled if you do not have a specific need. This prevents unauthorized access to your DS3000. Normally, you only need to enable remote login to debug specific failures, and you will be directed by your Support Representative in that case.

14.6 Recover from Failure

If the storage subsystem is not in an optimal state, the **Gather Support Information** link in the Troubleshooting and maintenance section (Figure 14-1 on page 300) moves down and the **Recover from Failure** link appears first, as shown in Figure 14-61.

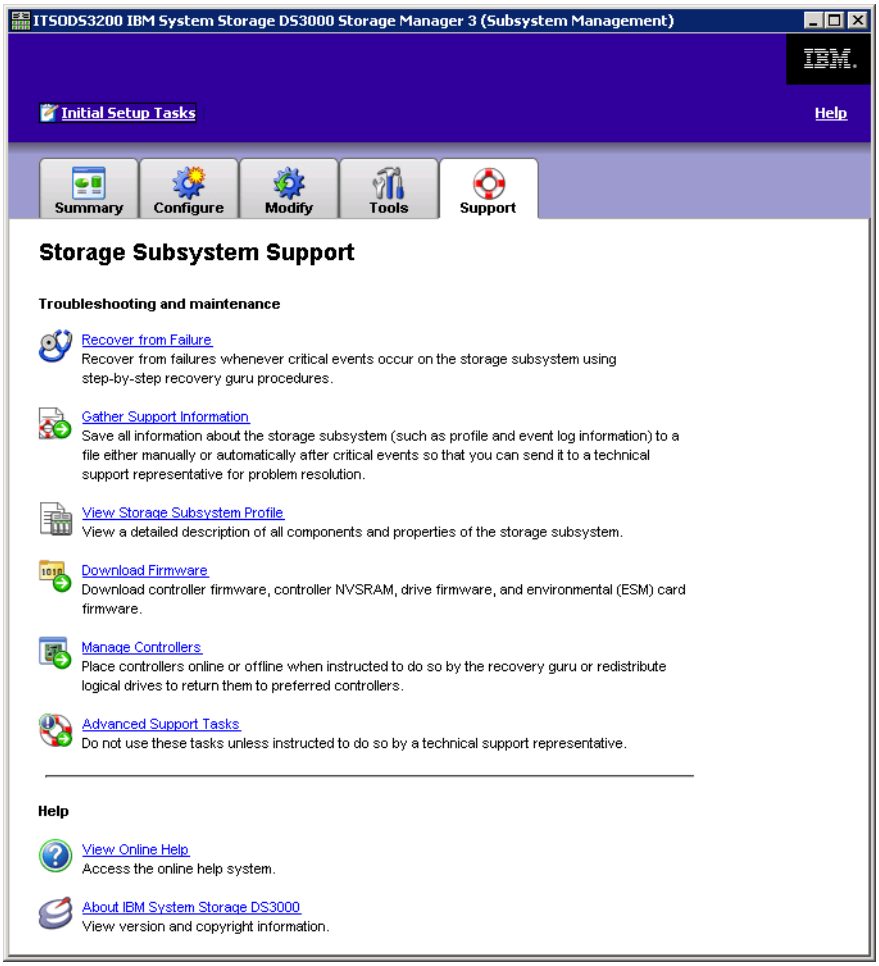


Figure 14-61 Support tab - Recover from Failure

Clicking this link starts the Recovery Guru (Figure 14-62). This is a very useful tool for troubleshooting.

The two upper panes (Summary and Details) give an overview of the problem. The larger pane below (Recovery Procedure) provides a detailed problem description and actions to resolve this problem.

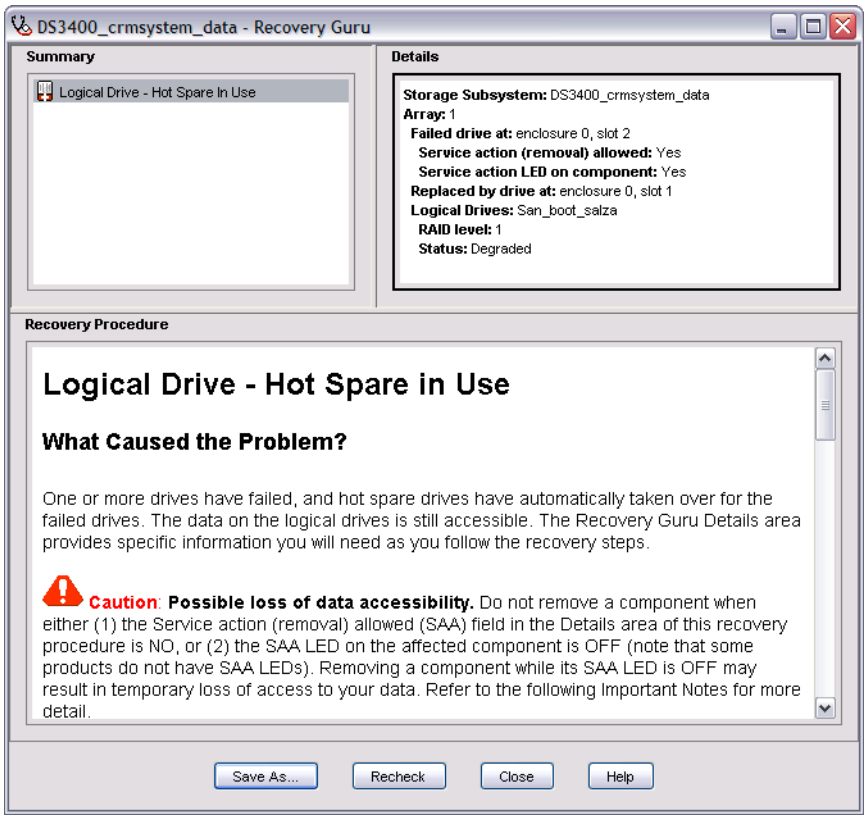


Figure 14-62 Recovery Guru

The example in Figure 14-62 is generated because of a failed hard drive. The Summary pane indicates that one logical drive has a problem, and that a hot spare drive is in use for this logical drive. The details box provides more details: the actual drive that has failed (enclosure 0 at slot 2), which hot spare is being used as the replacement active, and which logical drives are affected.

The Recovery Procedure box shows a detailed description of the problem. Read this information, then perform the recommended steps to resolve the problem (Figure 14-63).

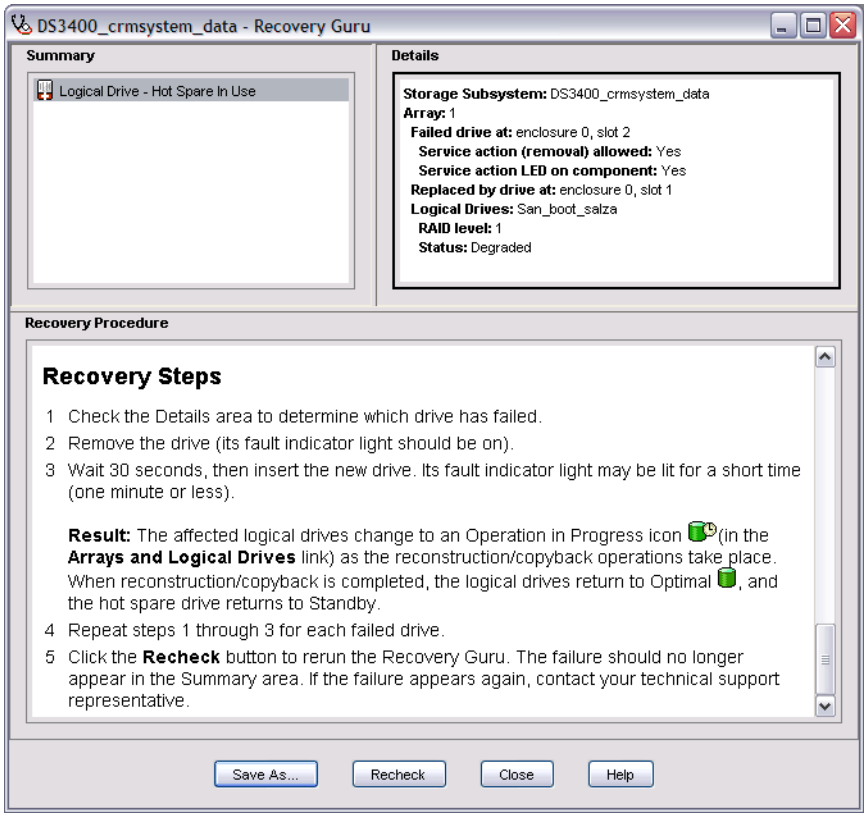


Figure 14-63 Recovery Guru - Recovery Steps

At the bottom of this window are four buttons:

- Save As

  Save the current Recovery Guru window as an HTML file.
- Recheck

  After you have completed the recovery steps, use this button to recheck the system.
- Close

  Close the Recovery Guru.
- Help

  Open the Recovery Guru Help.

14.7 View Online Help

Open this link in the Support tab to access the online help system. You will find a lot of useful information about your DS3000 storage subsystem in a new window. Select a topic from the right column. You can also search for text by clicking the magnifying glass icon.

Figure 14-64 shows an example of the Simplicity Online Help window.

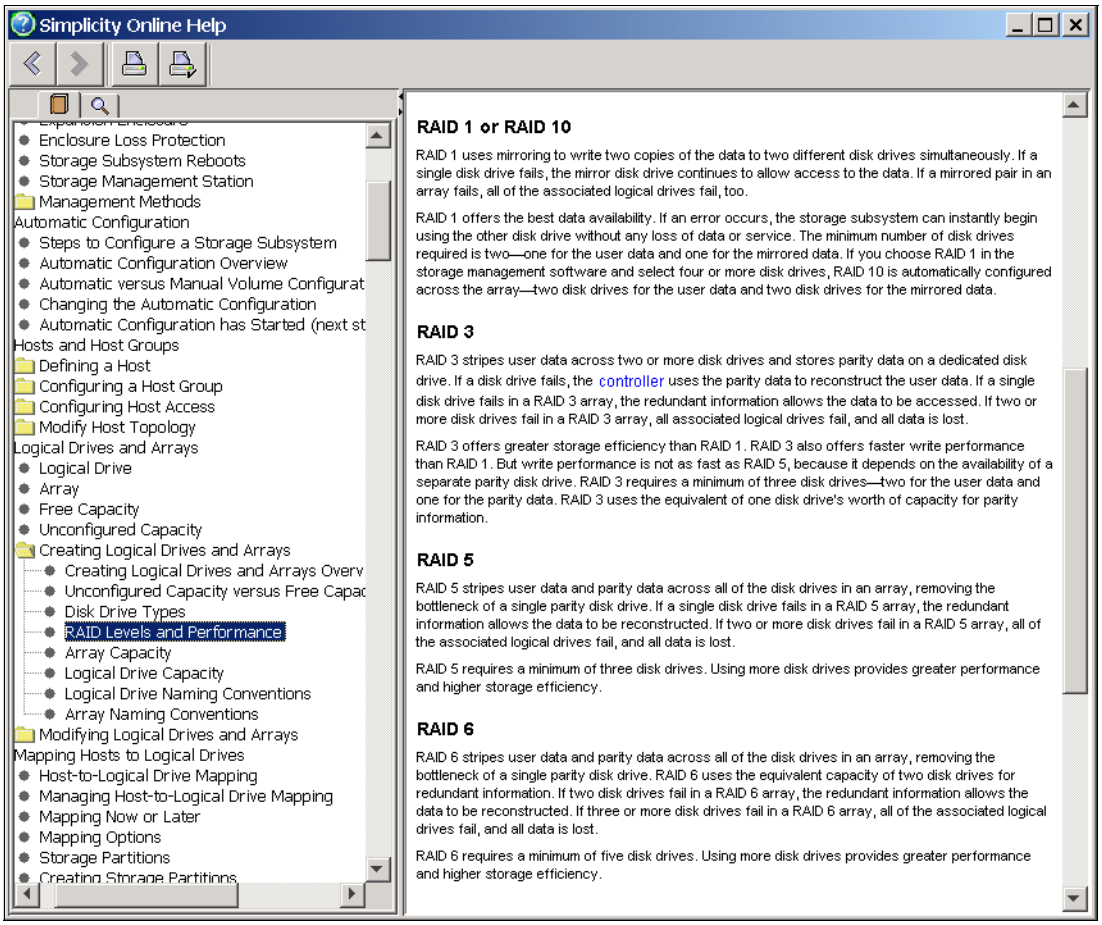


Figure 14-64 Online Help

14.8 About IBM System Storage 3000

Open this link in the Support tab to view the Storage Manager version and copyright information, as shown in Figure 14-65 on page 351. It shows the current version of the DS3000 Subsystem Manager window. It directly depends on the DS3000 Controller Firmware version of the managed storage subsystem.

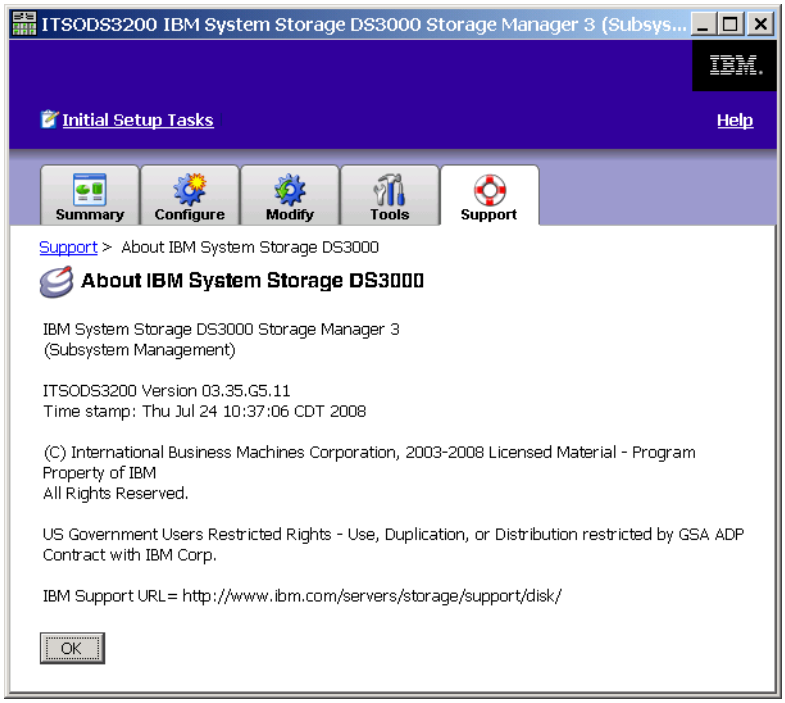


Figure 14-65 About DS3000

15

Remote Support Manager

In this chapter, we describe how to use Remote Support Manager (RSM™) with the IBM System Storage DS3000.

Note: Refer to the IBM Support Web site for the latest RSM and Storage Manager compatibility matrix:

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-66062&rs=594

© Copyright IBM Corp. 2007, 2009. All rights reserved.

15.1 DS3000 Remote Support Manager for Storage

The IBM Remote Support Manager for Storage (RSM for Storage) software is a no-charge software package that installs on an IBM System x server running Novell® SUSE® Linux Enterprise Server 9, SUSE Linux Enterprise Server 10, Red Hat Enterprise Linux 4 Advanced Server, or Red Hat Enterprise Linux 5, and provides problem reporting and remote access for IBM Service for the DS3000, DS4000, and DS5000 families of IBM storage subsystems.

The problem reporting provided by RSM for Storage automatically creates an entry in the IBM call management system for each subsystem that reports a problem. This is the equivalent of placing a voice call to IBM Service for a problem. Once in the system, problems are responded to with the same priority as specified by the maintenance agreement in place for the product.

RSM for Storage controls security for remote access by managing hardware and software components of the server it is installed on. Once installed, the server should be considered a single purpose appliance for problem reporting and remote access support for your storage subsystems. Only applications approved by IBM and specified in this document should be installed. (Management of the internal firewall and other configuration changes made by the software might prevent other applications from working.) There is no guarantee that applications that work with the current version of RSM for Storage will continue to work with future releases.

Remote access to the RSM for Storage system by IBM support is provided by either an external modem attached to the server or through an external SSH connection. This connection provides IBM Service with a command-line interface to the server. All bulk data transfers for logs and other problem determination files are sent to IBM through e-mail using the server's Ethernet interface. Isolation of remote and local users of the system from other devices on your intranet is performed by an internal firewall that is managed by the RSM for Storage software. Local and remote IBM users of the system do not have the ability to change any security features of the software.

Monitoring of storage subsystems is performed by your existing IBM DS Storage Manager software, which is configured to send SNMP traps to the Remote Support Manager when critical events are detected. Configuration of the management application is addressed later in this document.

The RSM for Storage user interface allows you to control and view the status of four management areas:

- System configuration
- Reporting
- Remote access
- Internal firewall

Your contact person for RSM for Storage will also receive e-mails of status changes in these areas if you have set up the application to do so.

One RSM for Storage server can support up to 50 DS5000, DS4000, DS3000, FAStT 200, and 500 storage subsystems. Only IP connectivity to the Ethernet management ports of the subsystems is required; serial cables are not needed for access to the disk subsystems. We recommend separating subsystems to be supported by an RSM for Storage system by country, and they must be covered under warranty or a current maintenance contract.

There are no annual fees for RSM.

15.1.1 Hardware and software requirements

RSM for Storage has the following hardware and software requirements.

Hardware requirements

RSM for Storage can be run on a dedicated IBM System x server or in a VMware client running on a System x server. Refer to the RSM for Storage *Compatibility Guide* for the minimum server requirements and a list of the specific servers that have been tested with the RSM software. The *Compatibility Guide* also contains the setup required for a VMware client that will host the Linux operating system running RSM for Storage. See the following URL for the RSM for Storage *Compatibility Guide;*

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-66062&rs=594

RSM has been tested on the following System x servers:

- x3250 4364
- x306m 8849

With these options:

- 512 MB memory.
- 20 GB disk space.
- Serial port: This must be on the system board. The serial port on the Remote Supervisor Adapter (RSA) cannot be accessed by RSM.
- Ethernet port: Note that if your SAN devices are on a private management LAN, a second Ethernet port for accessing your company's SMTP server and the Internet will be required if your selected server has only a single Ethernet port.

Software requirements

The RSM for Storage software requires the following prerequisite software:

- IBM DS3000 Storage Manager with Event Monitor installed in a management station in a different server.
- One of the following operating systems to install the RSM for Storage software:
  - Novell SLES 9 (Service Pack 3)
  - Novell SUSE SLES 10 (Base or SP1)
  - Red Hat RHEL 4 AS (Update 4 or 5)
  - Red Hat RHEL 5 (Base or update 1)

RSM for Storage software receives SNMP traps from the Event Monitor included with IBM DS Storage Manager. RSM for Storage software cannot be installed on the same system used to manage your storage network.

Note: Refer to the IBM RSM for Storage *Compatibility Guide* for the latest update of supported servers, modem, and operating systems. The document can be downloaded from the following Web page:

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-66062&rs=594

15.1.2 How RSM for Storage works

RSM for Storage uses an Ethernet connection for problem reporting and a modem or SSH connection for remote access by IBM Service, as shown in Figure 15-1. The DS3000 disk subsystem can be any supported disk subsystem.

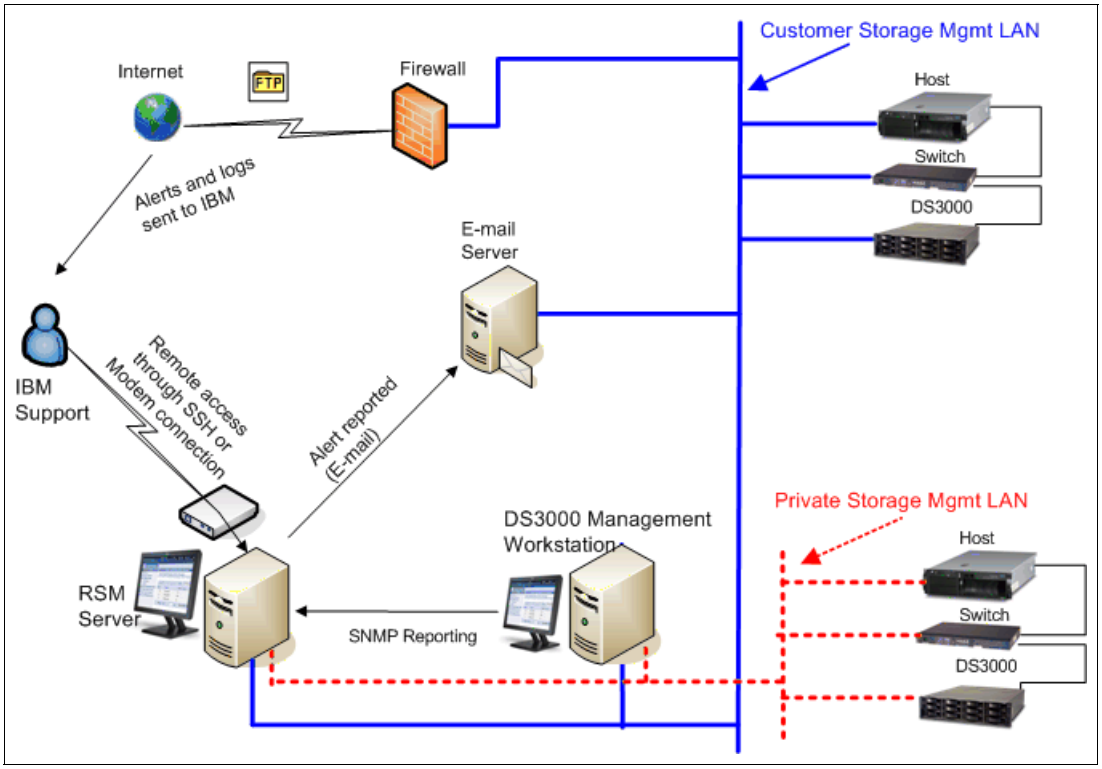


Figure 15-1 RSM for Storage connection diagram

The RSM for Storage server must have IP connectivity to the Ethernet management ports of the storage subsystems to be monitored and the management station running IBM Storage Manager's Event Monitor. It is also required that all storage subsystems, the management station, the e-mail server, and Internet gateway are accessible from the RSM server without requiring authorization through a firewall.

If your managed storage subsystems or other SAN devices are on a private management LAN, a second Ethernet port for accessing your company's SMTP server and the Internet will be required if your selected RSM server has only a single Ethernet port (see Figure 15-1).

This diagram shows that some storage subsystems and other SAN devices are managed through a private management LAN and the others are managed through the customer intranet. Therefore, the RSM server in this instance needs at least two network adapters.

After RSM is installed, configured, and activated, here are the steps in a sample scenario for RSM. Refer to Figure 15-1 to understand the flow:

1. For example, an error occurs in one of the storage subsystems and a critical event is logged in the management station (running Storage Manager).
2. The management station reports the critical event to the RSM server through an SNMP trap. The RSM system receives notification of the critical event and sends an alert to IBM Service.

When an alert is received from the management station, RSM downloads logs and other problem determination data, such as Major Event Log (MEL), Read Link Status (RLS), and the storage system profile of the storage subsystem's controllers that reports the problem using the out-of-band Ethernet interfaces, and sends them along with the alert to IBM Service by e-mail.

SNMP traps are sent by the IBM Storage Manager client or the IBM Storage Manager's Event Monitor service. As the Storage Manager Client may not always be running, we recommend that the Event Monitor be installed. Refer to Storage Manager documentation to check the installation of Event Monitor.

Refer to 15.1.3, "Notification e-mail and events filtering" on page 357 to configure the SNMP trap in Storage Manager.

3. IBM Support does problem determination based on information sent by the alert along with the problem determination data, such as MEL, RLS, and subsystem profile. If the problem can be fixed with existing information, IBM Support contacts the customer either by phone or e-mail to resolve the problem. After the problem is solved, either IBM Support or the customer must indicate *Service Complete* for all alerts in the RSM. IBM Support can dial to the RSM modem or use the SSH connection to **`acknowledge`** and indicate *Service Complete* for all alerts for the subsystem.
4. If the problem cannot be fixed with existing information, IBM Support dials the RSM modem or use an SSH connection, acknowledges the alert, and performs troubleshooting by connecting to the storage subsystem using the command-line interface (SMcli) or RLOGIN. IBM Support may need to contact the customer to obtain the password for the storage subsystem to use SMcli. IBM may also have to contact the customer to enable RLOGIN access. We recommend normally disabling RLOGIN.

   If IBM Support needs to send or upload logs or other information from the RSM server to IBM, they can use FTP or e-mail commands from the Linux shell at the RSM server while connected through the modem line or SSH connection. Any data connected is sent to an IBM server through a customer network, not through the modem line or SSH connection.
5. After the problem is resolved, all alerts must be closed either by IBM Support or the customer before reporting will resume for that subsystem.

Note: Once the RSM for Storage reports a problem to IBM for a given subsystem, no additional problems will be reported to IBM for that particular subsystem until all existing alerts are closed.

15.1.3 Notification e-mail and events filtering

The RSM for Storage software sends e-mails to notify of important changes in status. There are several types of notification sent by RSM to the primary contact or subsystem contact, as configured in the RSM:

- *Remote Access notifications*, sent when:
  - Remote access is enabled or disabled.
  - A remote user connects or disconnects from the system.
  - The remote access automatic time out is about to expire and the system has one or more active alerts.
- *Alerts Status notifications*, sent when an alert has been sent to IBM Service.
- *Daily Status e-mails*, which serve as a heartbeat notification that the RSM for Storage system is operational. This includes the summary status for the system and status for alerts that may be active for storage subsystems.

- *Firewall Status notifications*, sent when the internal firewall is enabled or disabled by a root or admin user.
- *Ignored Event notifications*, sent when an event is received that is configured to be ignored by the RSM for Storage system, and therefore not reported to IBM Service. These are events for which a response by IBM Service is not usually required, as listed in Table 15-1.

Table 15-1 Filtered events

Event code	Event text
6200	FlashCopy repository logical drive capacity - threshold.
6202	FlashCopy logical drive failed.
none	The persistent monitor running on Host xxxxxxx cannot reach the indicated storage system.
none	The persistent monitor running on Host xxxxxxx can now reach the indicated storage system.
4011	Logical Drive not on preferred path due to ADT/RDAC.

RSM and Storage Manager e-mail alerts

Storage Manager in the management station can be configured to send e-mail notifications when a problem is detected, as shown in “E-mail alerts” on page 120. However, this feature *should be disabled* when RSM for Storage is installed, if the e-mail contact configured in the Storage Manager is the same as the e-mail contact configured in the RSM Contact List. Otherwise, you will receive multiple notification e-mails about the same problem: one notification from RSM and another one from Storage Manager.

To disable e-mail alerts in Storage Manager, do the following steps:

1. Right-click your management station in the Storage Manager Enterprise window and select **Configure Alerts** to select All Storage Subsystems.
2. On the Email tab, delete any configured e-mail destinations.

If there are e-mail addresses already configured to receive e-mails from Storage Manager but are not listed in the RSM Contact List (see Figure 15-7 on page 365 for the Contact List), it is not necessary to delete them in Storage Manager.

15.1.4 Remote access methods

The required Remote Access for IBM support can be provided by one or both of two methods. An external modem can be attached to the server’s serial port, or remote access through an SSH client can be enabled.

Remote access through modem

The functional requirements for the modem are minimal; it is only used for remote access by IBM Support. Most "Hayes-compatible" external modems can be used. The modem should have a minimum speed of 56 Kbps supporting either V.90 or V.92.

The RSM for Storage software has been tested with the following modems:

- Multitech Multimodem II MT5600BA
- Multitech MultiModem ZBA MT5634ZBA

Refer to the RSM for Storage *Compatibility Guide* for updated information about which modems are supported. You will need to contact your IT support staff for installation and problem resolution related to the modem.

Remote access through SSH

Instead of using a modem for external access to the RSM for Storage system, you can allow remote access by an external SSH connection. To do this, you need to map a port on your external firewall to the IP address and SSH port (22) on the RSM for Storage system. While the RSM for Storage system has several layers of login protection on the SSH port, you can also require authentication before the external firewall makes a connection to the RSM for Storage system.

You can also choose to allow remote access by both methods. More information about setting up and using an SSH connection is available in the *IBM Remote Storage Manager for Storage Planing, Installation, and User's Guide*, GC26-7933, as well as the supplement "Installation Hints and Tips", found at:

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-66062&rs=594

15.1.5 RSM management interface

The RSM management interface can be accessed through a Web browser pointing to the IP address or host name of the RSM server using HTTPS. You can use the Web interface to check the status and configure RSM settings. For IBM Support, the interface is a command-line interface when connected to the RSM server through a modem.

Figure 15-2 is an example of a System Configuration menu for an already configured and activated RSM system.



Figure 15-2 RSM System Configuration menu

Under System Configuration, there are links at the top of the page that provide a summary status of the RSM system. Depending on the status, various icons may be displayed to the left of each link. The status values for each of these are as follows:

- System
  - OK: RSM is correctly configured.
  - Incomplete: Format or content of configuration fields is incorrect or missing.
  - Problem: Indicates that the configuration test has detected a problem.
- Reporting
  - Reporting: All subsystems are being monitored and no problems exist.
  - Suspended: Indicates that there is a problem with the configuration that makes the system status not OK. No events will be processed by the RSM if the reporting status is Suspended.
  - Partial: Reporting has been disabled for one or more subsystems.
  - Standby: Reporting has been disabled for all subsystems.

  - Storage Problem: One or more subsystems has a problem.
- Firewall
  - Disabled: No restrictions on access to the networks connected to the RSM.
  - Enabled:Closed: This is the normal state when there are no active alerts present and the system is waiting for notification of a problem from Storage Manager.
  - Enabled:Open: Firewall is allowing access to one or more subsystems or other configured SAN devices.
- Remote Access
  - Disabled: Remote access through the modem or SSH is not allowed.
  - Enabled: The modem will answer and the remote service (rservice) user ID is valid for login through the remote connection or SSH.
  - Active: The rservice remote user is logged into the RSM.

15.1.6 RSM security considerations

RSM for Storage controls security for remote access by managing the hardware and software components of the server it is installed on. Once installed, the server should be considered as a single purpose appliance for problem reporting and remote access support for your storage subsystems; it should not be used for other applications.

User ID

During installation, the RSM software creates three user IDs:

- admin: This is the administrative user that can perform management and configuration tasks.
- lservice: This is the local service user intended for use by IBM Service when on site. This user ID has some restrictions on directories it can access. This is to prevent any configuration change that might affect the security of the system.
- rservice: This is the remote service (IBM Service) user that is used exclusively for remote access to the system and only valid when Remote Access is enabled. This user ID also does not have the ability to change any of the RSM security features.

Passwords for user ID admin and lservice can be changed by the Linux root user using the command **`rsm-passwd admin`** or **`rsm-passwd lservice`**. We recommend setting a different password for each user ID.

For the remote user (rservice), the password is automatically generated by RSM and it is changed daily at midnight UTC. IBM Service has an internal tool that provides the current password, so you do not need to provide the current RSM password to IBM Service.

After validation of the initial login password, remote users are presented with a challenge string, which also requires access to an internal IBM tool in order to obtain the correct response. The response also includes an IBM employee user name that is recorded in the RSM for Storage Security Log.

15.2 Installing and setting up RSM

In this section, we show how to install and configure RSM. Before beginning, go to the RSM support Web page and carefully review the *IBM Remote Storage Manager for Storage Planing, Installation, and User's Guide*, GC26-7933, as well as the supplement "Installation Hints and Tips", found at:

http://www.ibm.com/support/docview.wss?uid=psg1MIGR-66062&rs=594

Tip: Do not use a remote connection when installing the RSM for Storage on the workstation. We recommend being logged on locally to the Graphical User Interface of the workstation, as RSM resets the firewall settings, preventing remote access to the Linux workstation.

15.2.1 Installing the host OS

There are various operating systems supported for the RSM host, as shown in "Software requirements" on page 355. In our example, we used a Red Hat Linux Enterprise License 5 as the host operating system on the RSM server. When installing RHEL5, we selected the following additional packages:

- expect
- mgetty
- Simple Web server (apache2)
- KDE desktop environment

See the *IBM Remote Storage Manager for Storage Planing, Installation, and User's Guide*, GC26-7933, for specific operating system installation instructions.

15.2.2 Installing RSM

The RSM software can be downloaded from

http://www.ibm.com/storage/disk/rsm

We installed RSM according to the instructions in *IBM Remote Storage Manager for Storage Planing, Installation, and User's Guide*, GC26-7933. After the installation, you have to define the admin and lservice user IDs.

15.2.3 Setting up RSM

After the installation is complete, we have to set up RSM.

1. On the login screen, do the following:
   a. Click **Session** and select KDE.
   b. Log in as the admin user.
   c. Click the **Manage** icon on the desktop to open the Web browser, which displays the main administration window shown in Figure 15-3 on page 363. Click **Login.**



Figure 15-3 RSM Administration window

2. Enter the user name and password of the RSM administrator. This account is named admin and the password was defined during the installation of RSM. See Figure 15-4.

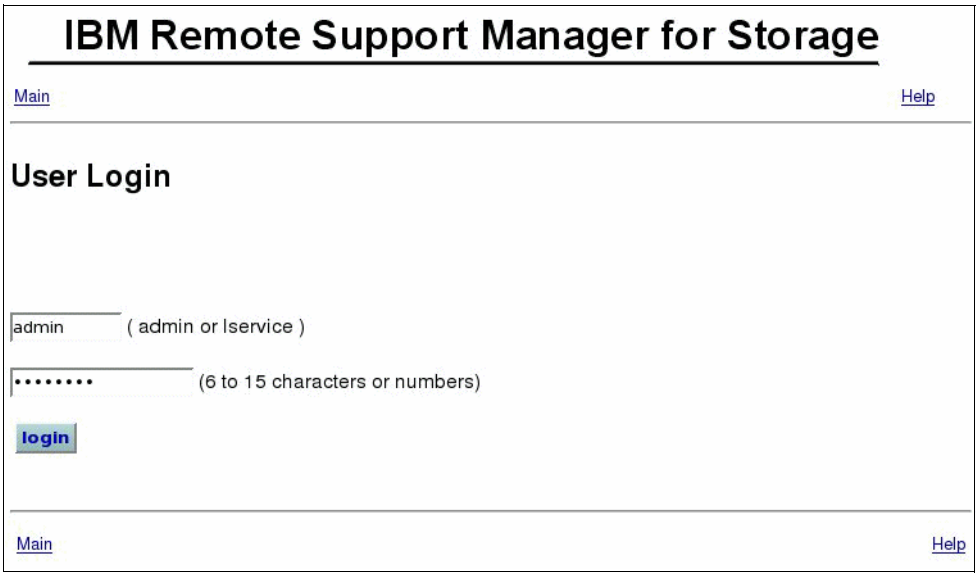


Figure 15-4 RSM logon window

3. You return to the main menu. The upper right menu contains a logout link.

   To start the setup, click **System**, as shown in Figure 15-5.



Figure 15-5 RSM Main Menu

4. The System Configuration window shows the incomplete tasks that need to be completed before the RSM system can be used (Figure 15-6).

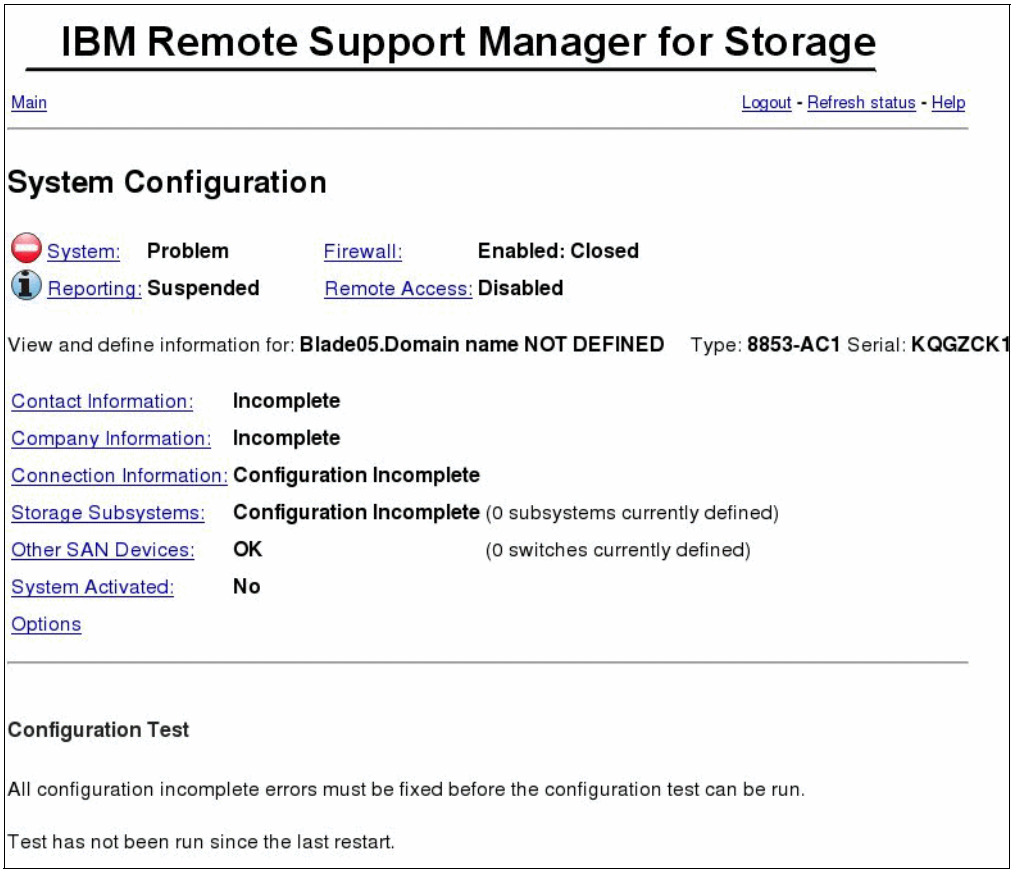


Figure 15-6 System configuration incomplete

5. Click **Contact Information** to access the contact list to add a new contact (Figure 15-7).



Figure 15-7 RSM Contact List

6. Click **Select to add**, then fill out the form and click **Update configuration**. This information is very important; when the DS3000 storage subsystem reports a problem, IBM Support will contact the person specified here. See Figure 15-8 on page 365.

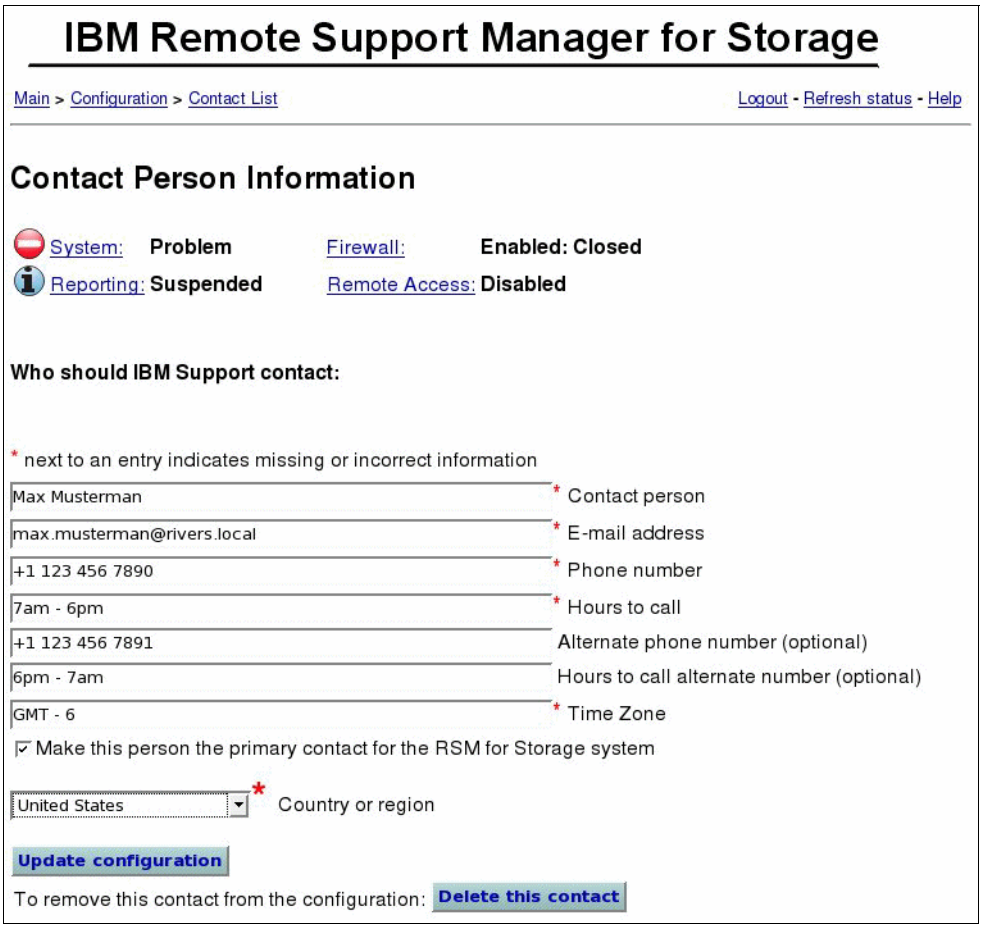


Figure 15-8 RSM Contact Person Information

7. The specified contact will be added to the contact list. Multiple contacts can be defined that may be assigned with different storage subsystems or SAN switches. Click **Configuration** at the bottom of the window to return to System Configuration. See Figure 15-9 on page 366.



Figure 15-9 RSM Contact List with contacts

8. The task Contact Information in System Configuration will be marked OK. Click **Company Information**, as shown in Figure 15-6 on page 364.
9. Fill out the form with the appropriate information and click **Update configuration**, as shown in Figure 15-10.

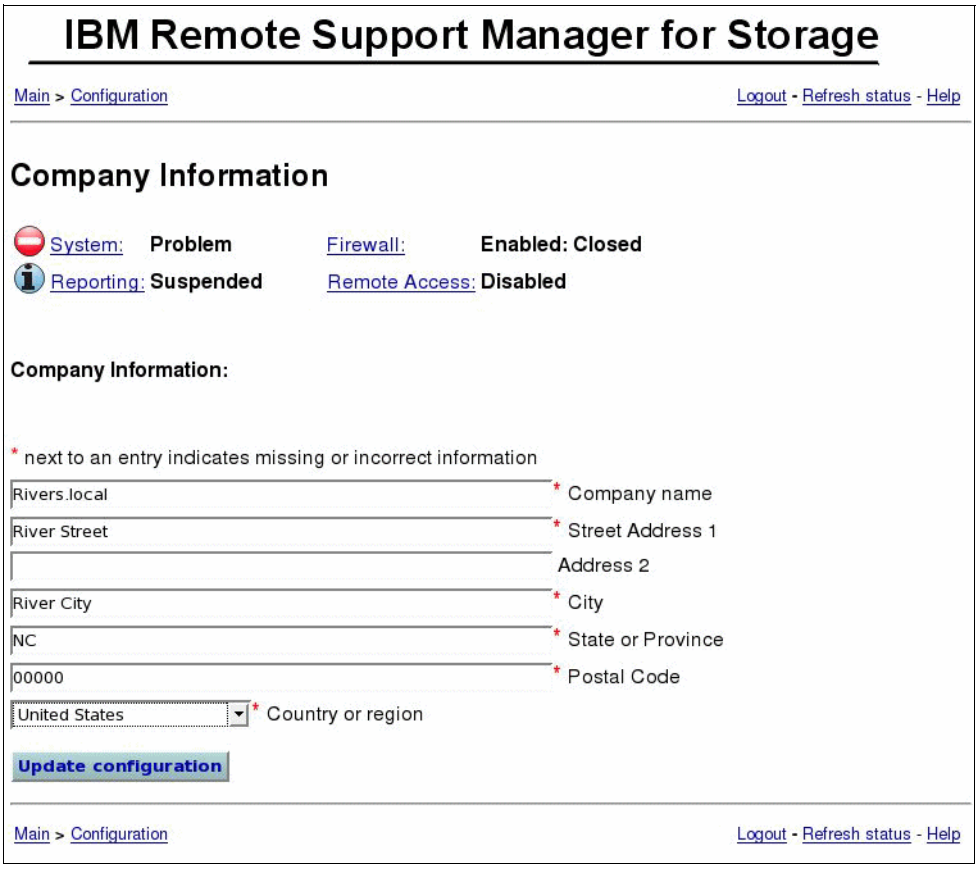


Figure 15-10 RSM Company Information

10. The task Company Information in System Configuration will be marked OK. Click **Connection Information**, as shown in Figure 15-6 on page 364.

11. Fill out the form with the appropriate information and click **Update configuration**, as shown in Figure 15-11.

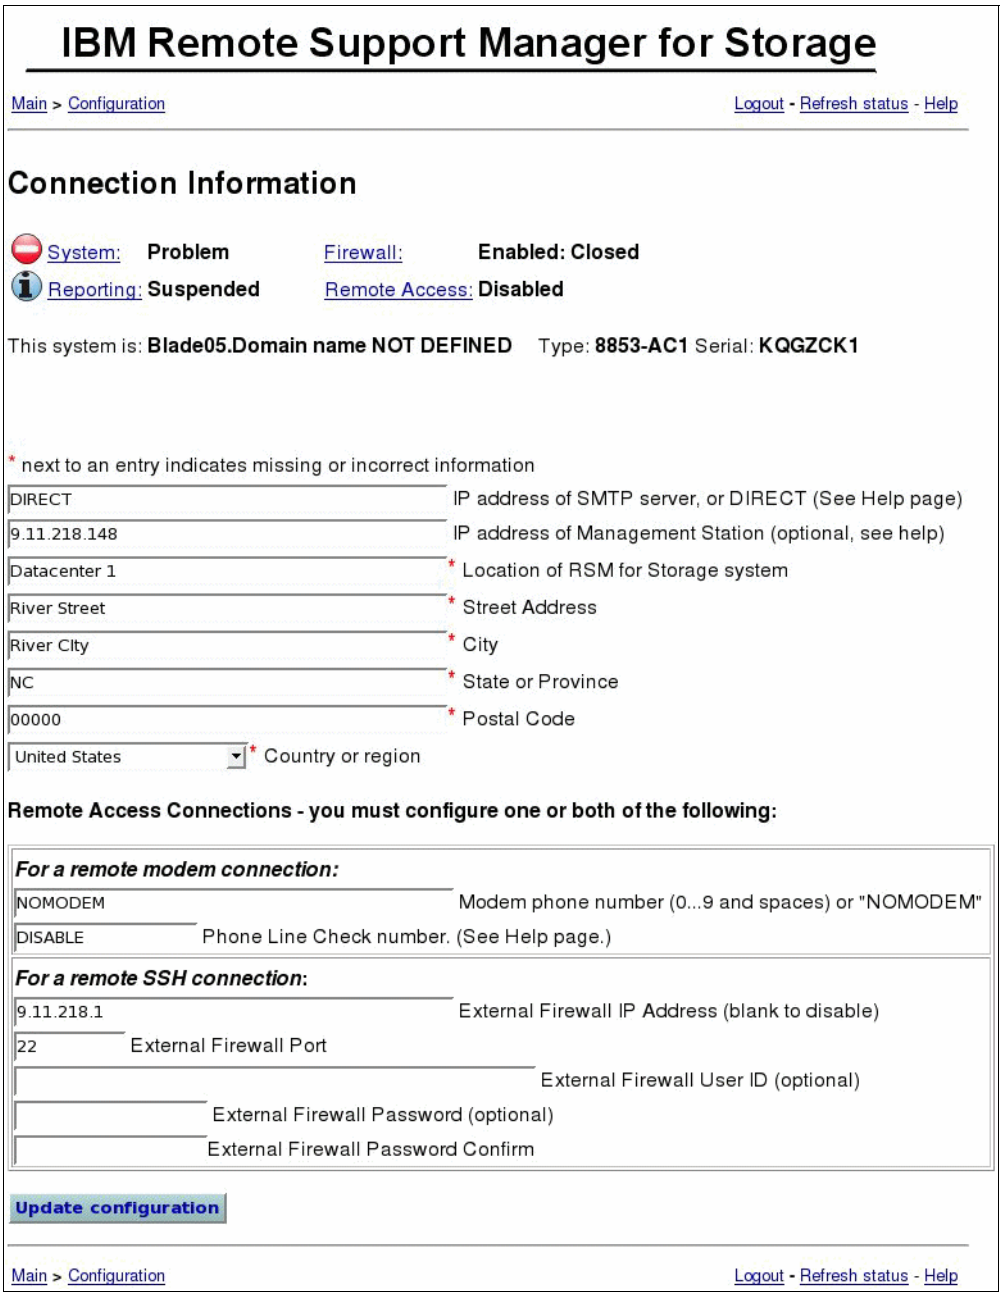


Figure 15-11 RSM Connection Information

12. The task Connection Information in System Configuration will be marked OK. Click **Storage Subsystems**, as shown in Figure 15-6 on page 364.
13. Click **Select to add** to define a storage subsystem in RSM (see Figure 15-12 on page 369).

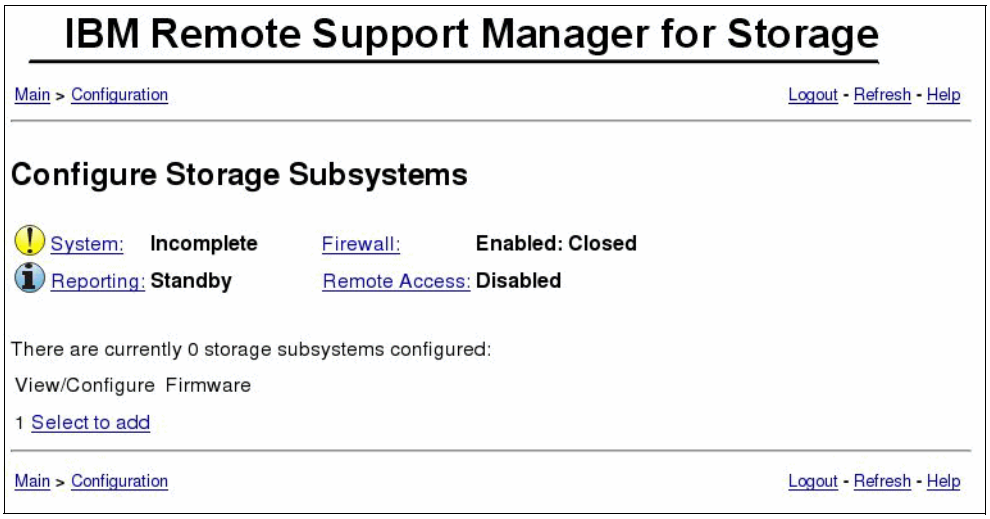


Figure 15-12 RSM Configure Storage Subsystems

RSM for Storage reports DS3000 Expansion units as well, so the configuration of drive expansion units is also required. During the configuration test (done later in the configuration steps), the profile for each subsystem will be downloaded. This will verify connectivity to the subsystem, verify that this version of RSM for Storage software is compatible with the firmware on the subsystem, and determine if there are any drive expansion units attached to the controller. If any drive expansion units are detected, the configuration status for the subsystem will change to Configuration Incomplete and additional configuration fields will now be available for setting the IBM machine type and serial numbers of each detected drive expansion unit.

Fill in the following fields in the form:

a. The name of a storage subsystem, as shown in the DS3000 Storage Manager.

b. The country where the storage subsystem is located.

c. A description where to find the storage subsystem, for example, room number or rack number in the location field.

d. Management interface IP addresses of all controllers in the subsystem. If only one controller is installed, use only one IP address.

e. Type, model and the serial number; this information can be found on the left front bezel or the rear side of the DS3000.

f. Select an entry from the contact list.

Click **Update configuration**. Figure 15-13 shows the form filled out with sample data.

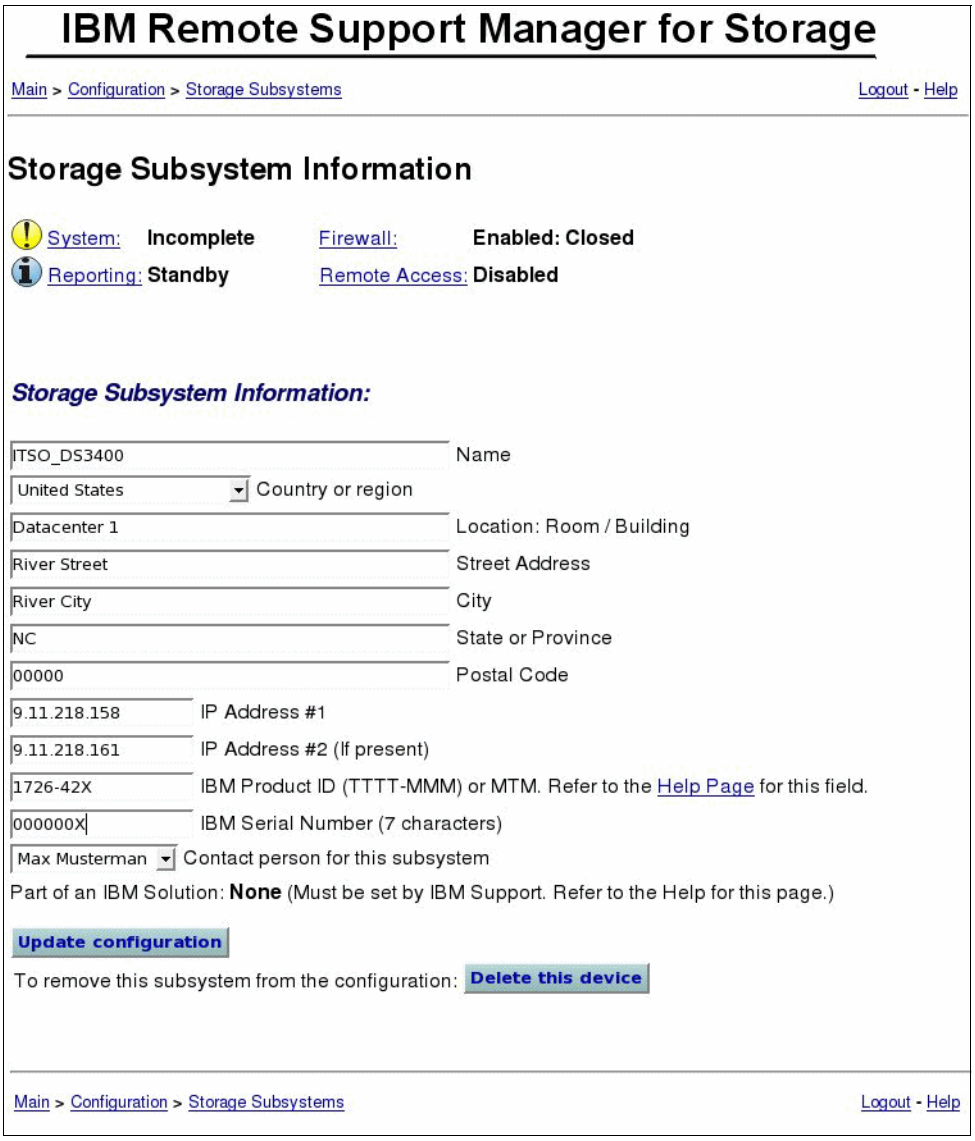


Figure 15-13 RSM Storage Subsystem Information

14. The storage subsystem will be added to the list of configured storage subsystems, as shown in Figure 15-14. Up to 50 storage subsystems and SAN switches can be added. Click **Configuration** to return to the System configuration.

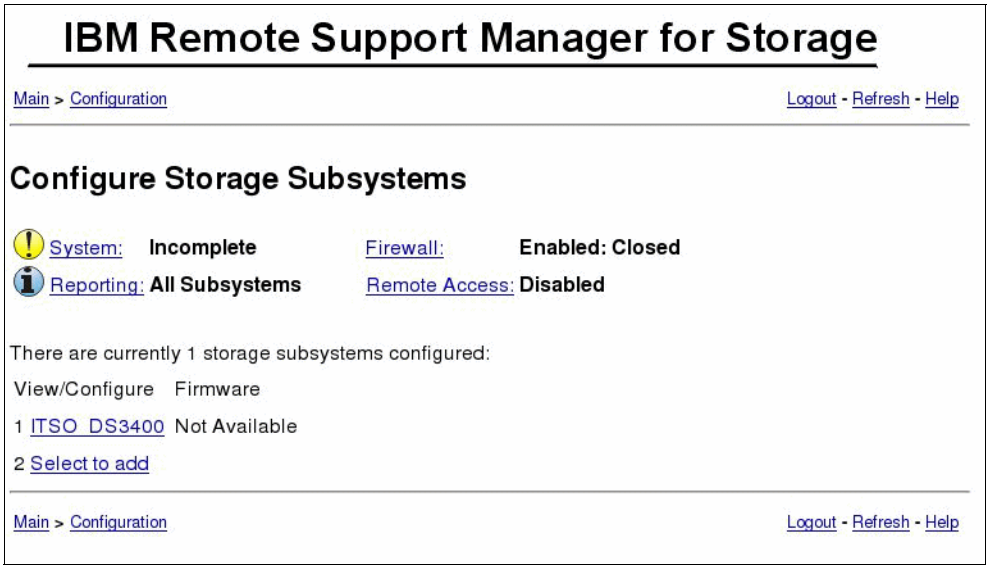


Figure 15-14 RSM Configure Storage Subsystems

15. When all tasks are completed correctly, run the configuration test. Click **Run Configuration Test**, as shown in Figure 15-15.

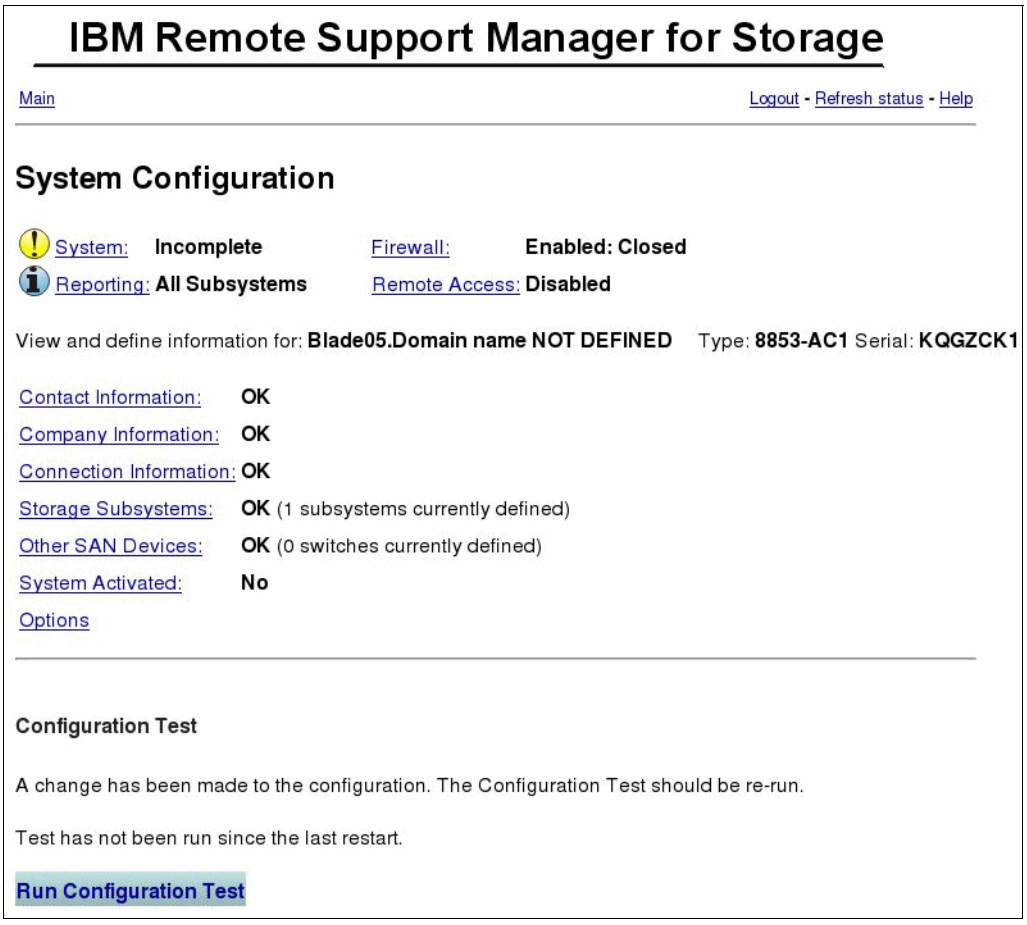


Figure 15-15 RSM Run Configuration Test

16. Click **Refresh status** to see the progress of the test (Figure 15-16).

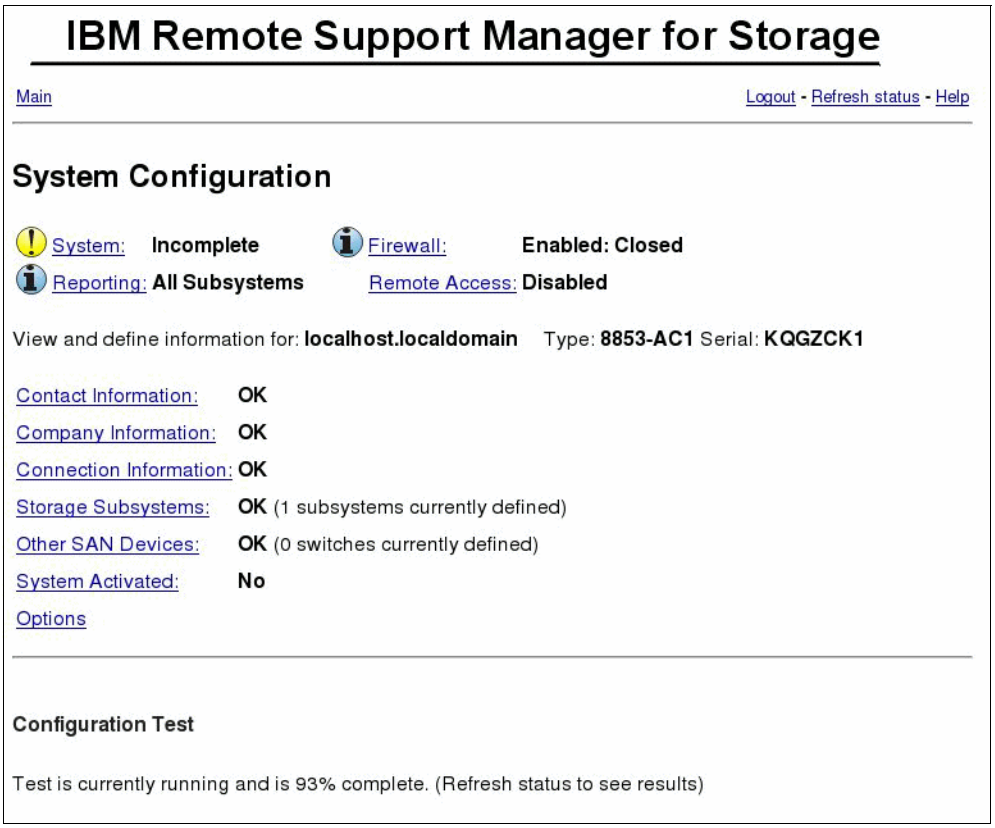


Figure 15-16 RSM configuration test

17. When the test is complete, the date and time of the last run is shown, as shown in Figure 15-17.



Figure 15-17 RSM Configuration Test complete

18. Results from this test are logged in the Activity log. You can access the Activity log through an icon on the KDE desktop, or you can use the command **`tail -fn 10000 /var/log/rsm/activity.txt`**. to see the contents of the log file, as shown in Figure 15-18.

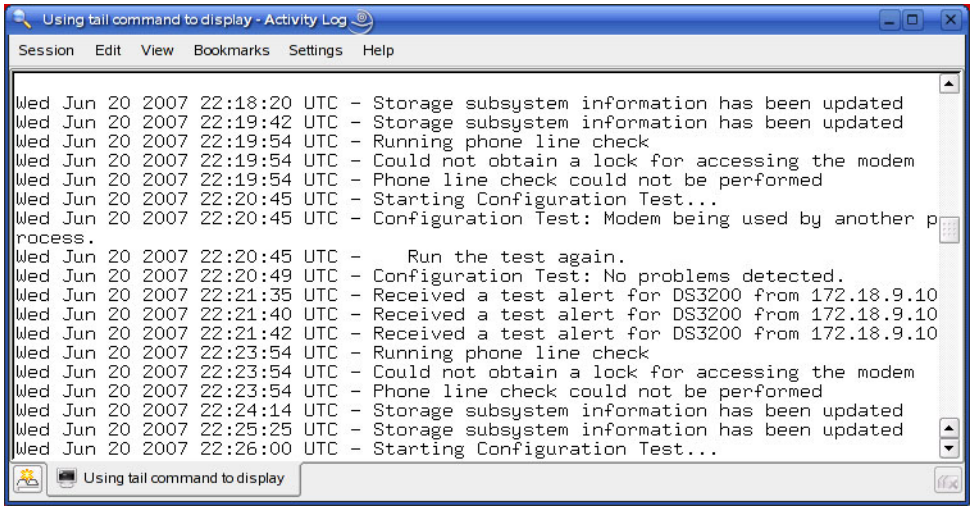


Figure 15-18 RSM Activity Log shown on a KDE desktop

15.2.4 Configuring SNMP traps in Storage Manager

To allow the DS3000 management workstation to send SNMP traps to the RSM server, set the RSM server as your SNMP traps destination in the Storage Manager client. We showed how to do this in “Configure Alerts” on page 118.

Click the **SNMP** tab and type the IP address of the RSM server in the Trap Destination field. Keep the default community name of public.

If you have an existing SNMP infrastructure and there is already an SNMP trap destination set, you can add the IP address of the RSM server as an additional SNMP trap destination without having to delete the existing SNMP trap destination setting.

Configure the DS3000 Storage Manager to send SNMP alerts for each defined storage subsystem in RSM to the RSM host. Open the Enterprise management window of the storage manager. Right-click a DS3000 storage subsystem. In the Connect menu, click **Configure Alerts...**. In the Configure Alerts dialog, click the **SNMP** tab. Enter the host name or the IP address of the RSM host in the trap destination field. Click **Add** (see Figure 15-19). Do not change the SNMP community name (the default is public).

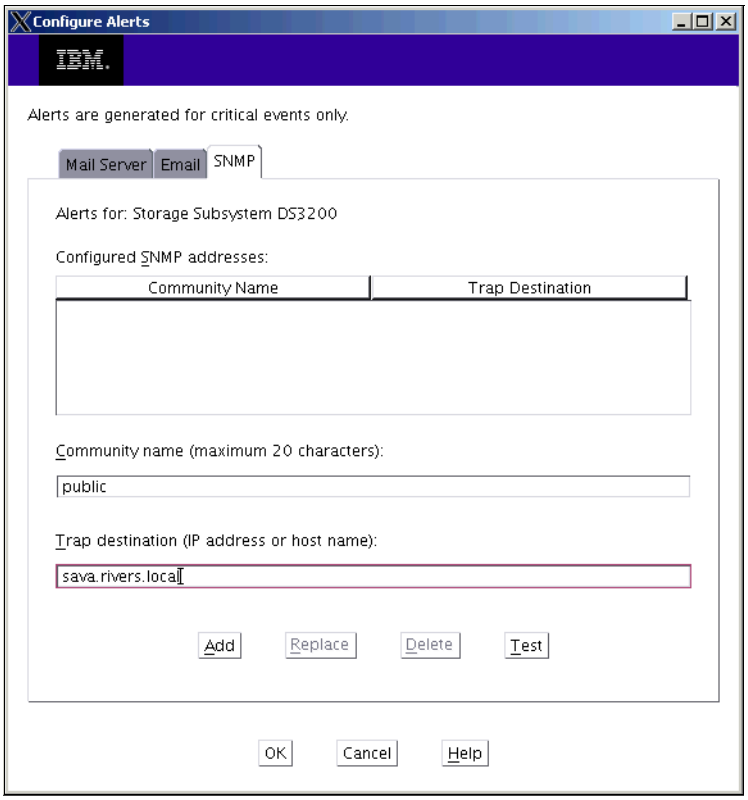


Figure 15-19 Add SNMP trap receiver

Validate SNMP configuration

Select the SNMP trap destination and click **Test** to send a test trap to the RSM host (Figure 15-20).

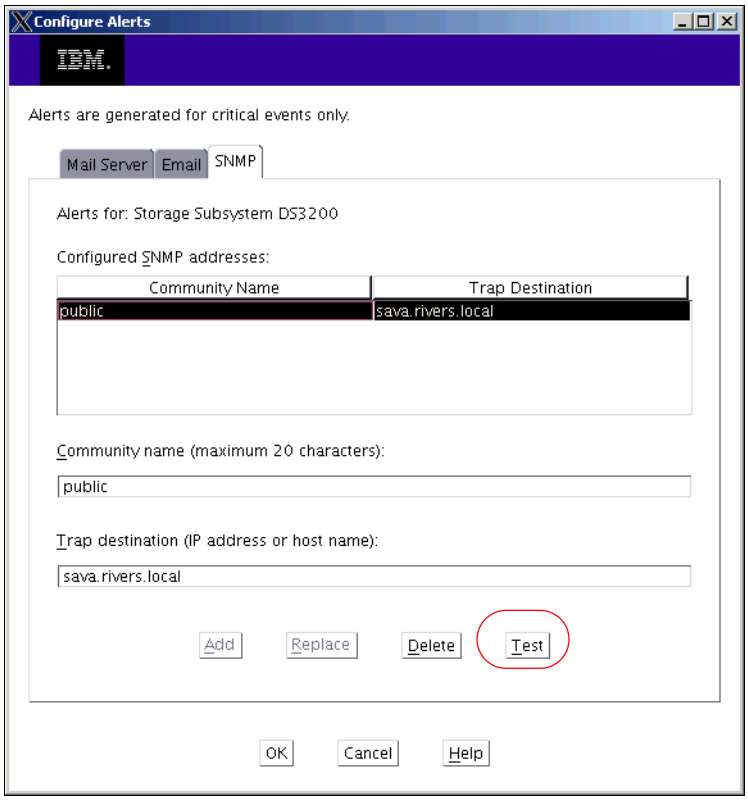


Figure 15-20 Send test trap

Check the Activity Log (as shown in step 18 on page 374) and verify that the trap was received. The activity log will contain a entry similar to what is shown in Example 15-1.

Example 15-1 Test alert received

```
Wed Aug 20 2008 22:21:35 UTC - Received a test alert for DS3400 from 9.11.218.148
```

15.2.5 Remote Access connectivity

Here we discuss the topic of Remote Access connectivity.

SSH connectivity

On the RSM for Storage Remote Access page, click **Enable Remote Access** (see “Remote Access policy” on page 377). This will reconfigure the RSM for Storage internal firewall to allow connections through SSH port 22.

Verify the connectivity with the following sequence:

1. From inside your network, open an SSH client and connect to the RSM for Storage system on port 22. (Remember, if you perform these connectivity checks over several days, that the RSM for Storage Remote Access control has a timeout that may need to be reset.) Verify that you are able to obtain a login prompt from the RSM for Storage system.

2. From outside your network, open an SSH client and connect to your external IP address port that has been assigned (port mapped) to the RSM for Storage system. If an authentication process has been put in place by your firewall administrator, verify that the user ID and password for the external firewall is specified in the RSM for Storage Connections configuration.

Modem connectivity

Adding a modem to one of your systems creates a potential entry point for unauthorized access to your network. RSM for Storage modifies many characteristics and behaviors of the system it is installed on to protect this entry point and to maximize the amount of control you have in managing remote access.

In RSM, the modem used for remote access by IBM Service will not answer unless one of the storage subsystems has an active alert or Remote Access has manually been enabled.

Normally, Remote Access is enabled automatically when an alert is sent to IBM, but you can choose to wait for IBM Service to contact you when an alert is received and manually enable Remote Access at that time.

On the RSM for Storage Remote Access page, click **Enable Remote Access**. This will enable the modem to answer when called. Verify the modem connectivity by calling the modem phone number from a voice phone:

1. Most modems will either flash an LED or you may hear a sound when a call is being received. If there is no indication that a call is being received:
   a. Plug the phone cable into an analog voice phone and verify that a dial tone is present.
   b. If a dial tone is present, hang up and then call this phone number from another phone and verify that the phone you just connected rings.
   c. Reconnect the phone cord to the modem connector labeled line #1.
2. Try again to verify the modem connectivity by calling the modem phone number from a voice phone.
3. Check the instructions that came with the modem and review any troubleshooting information.

Remote Access policy

To configure the Remote Access policy, click **Remote Access** from the Main Menu, as shown in Figure 15-21.

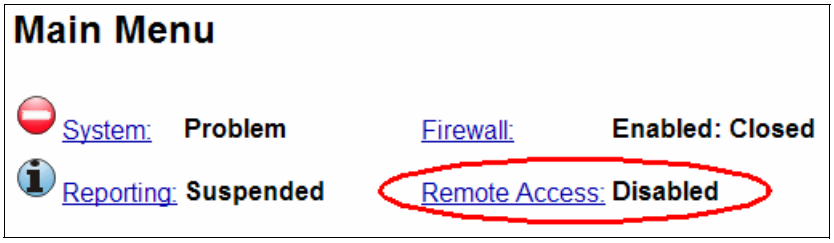


Figure 15-21 Accessing Remote Access Policy from Main Menu

In the Remote Access setting page, you can enable/disable the Remote Access service and enable/disable the option to automatically enable the Remote Access when an alert is sent to IBM. This is shown in Figure 15-22.

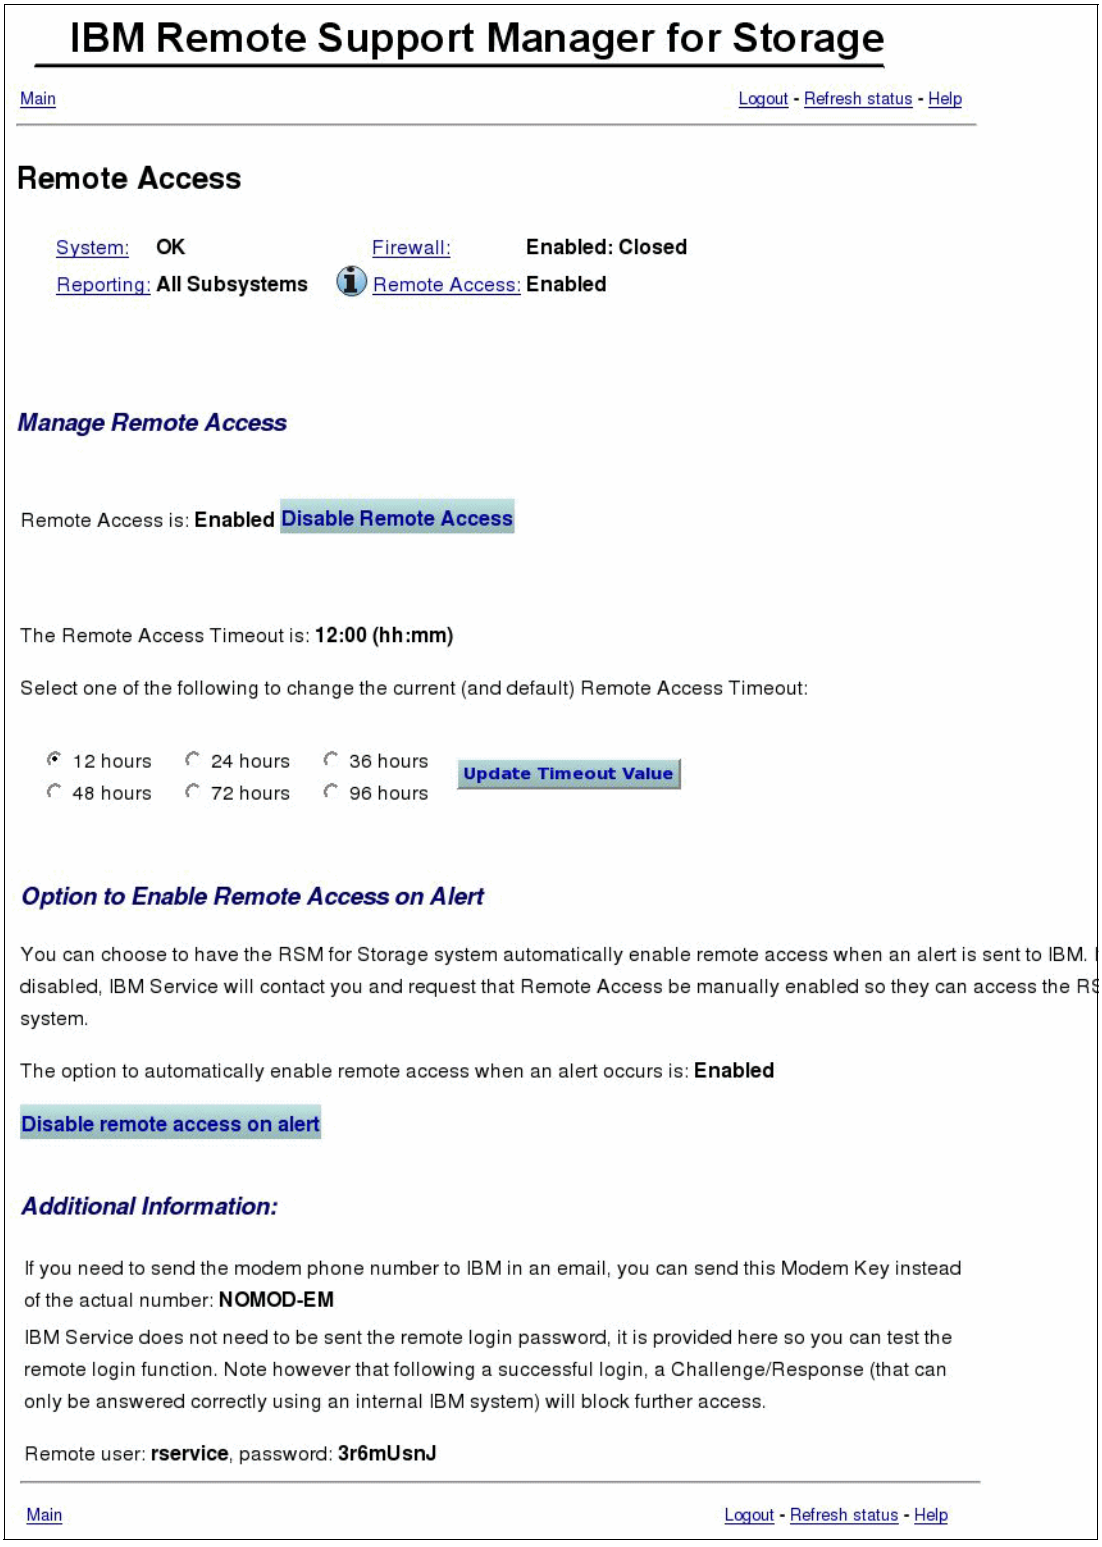


Figure 15-22 Remote Access settings

Remote Access also has a configurable timeout between 12 to 96 hours. You can manually disable remote access when the service is complete or allow it to time out. After the timeout period has elapsed, the system is guaranteed to return to a secure state without intervention.

To configure the timeout value, select the desired timeout value, and click **Update Timeout Value**, as shown in Figure 15-23.

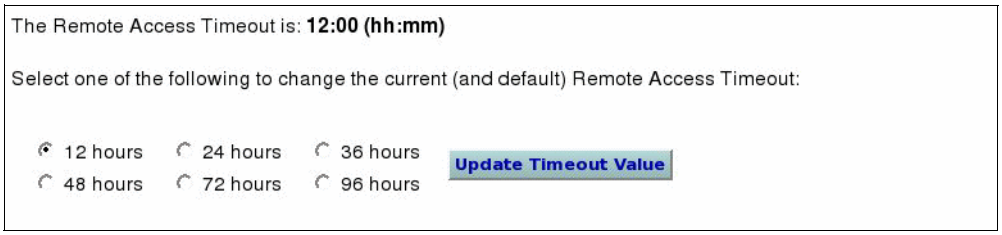


Figure 15-23 Configure Remote Access Timeout

Note: You do not need to provide the rservice user ID password to IBM Service, because IBM Service has an internal tool that provides the current rservice password. You only need to provide passwords of your storage subsystems or your other SAN devices, if required.

Internal firewall

RSM for Storage includes an internal firewall to limit the scope of access a remote user has to your network. It also limits the IP destinations that can be accessed by local and remote users of the system. The rules for inbound and outbound IP traffic that control the internal firewall are managed dynamically by the RSM for Storage software.

There are three possible states of the RSM firewall: Disabled, Enabled:Closed, and Enabled:Open.

- Disabled: There are no restrictions on access to the networks connected to the RSM. Remote access is not permitted if the firewall is in this state.
- Enabled:Closed: This is the normal state when there are no active alerts present and the system is waiting for notification of a problem from Storage Manager.

  The following rules are applied when RSM is in this state:

  - Inbound SNMP traps, ping, traceroute, and HTTPS requests are accepted.
  - Outbound traffic for DNS, ping, traceroute, IBM WWW and FTP sites, and port 25 of your SMTP server are accepted.
  - Access to any of your other SAN devices are blocked.
  - IP traffic forwarding is disabled.
  - All other connections are blocked.
- Enabled:Open: The firewall is allowing access to one or more subsystems or other configured SAN devices. Access is allowed only to those devices that have active alerts or those that you have placed in Service Access mode.

Note: Subsystems with active alerts are automatically allowed access from the Remote Support Manager while the alert is active and do not need to be enabled for Service Access.

To manage the RSM internal firewall and service access of your managed storage subsystems (and other SAN devices) from the Web interface, click **Firewall** on the Main Menu, as shown in Figure 15-24.



Figure 15-24 Accessing the firewall configuration from Main Menu

In the Internal Firewall and Service Access page, you can change the internal firewall status and service access mode of your managed storage subsystems, as shown in Figure 15-25 on page 381.

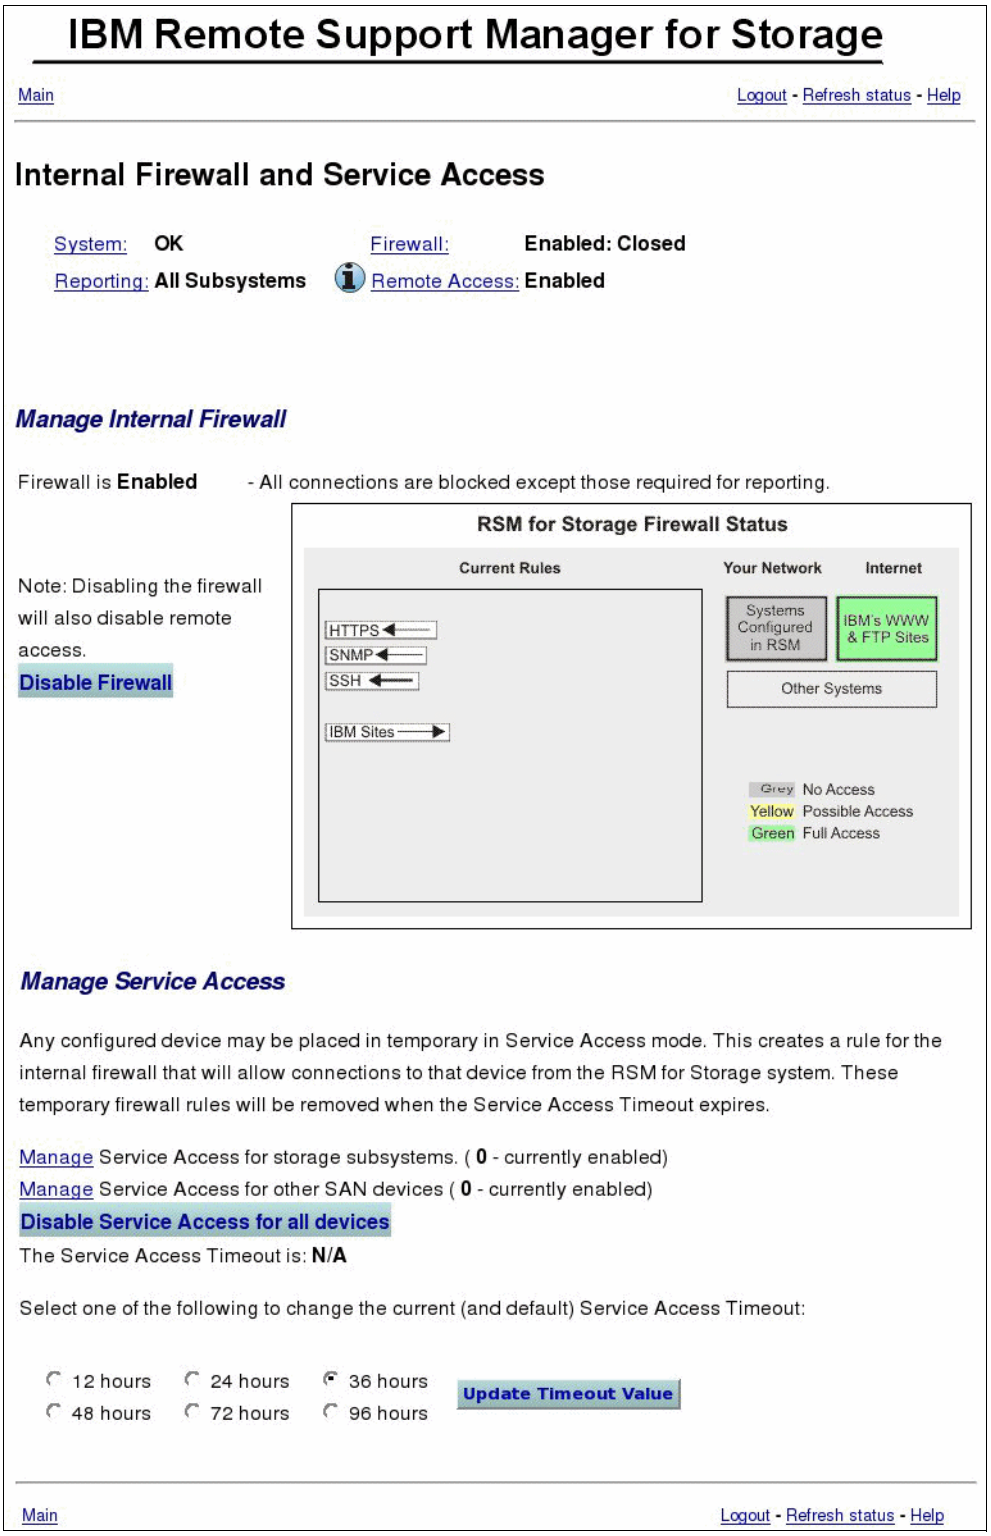


Figure 15-25 Internal Firewall and Service Access page

Placing a device into Service Access mode will create a rule for the internal firewall that will allow connections to that device from the RSM server. For subsystems with active alerts, they are automatically allowed access from the Remote Support Manager while the alert is active and do not need to be enabled for Service Access.

Similar to Remote Access, you can also modify the Service Access Timeout. To set the Service Access Timeout, go to the Manage Service Access section in the Internal Firewall and Service Access page, select the desired Service Access Timeout value, and click **Update Access Timeout**, as shown in Figure 15-26.

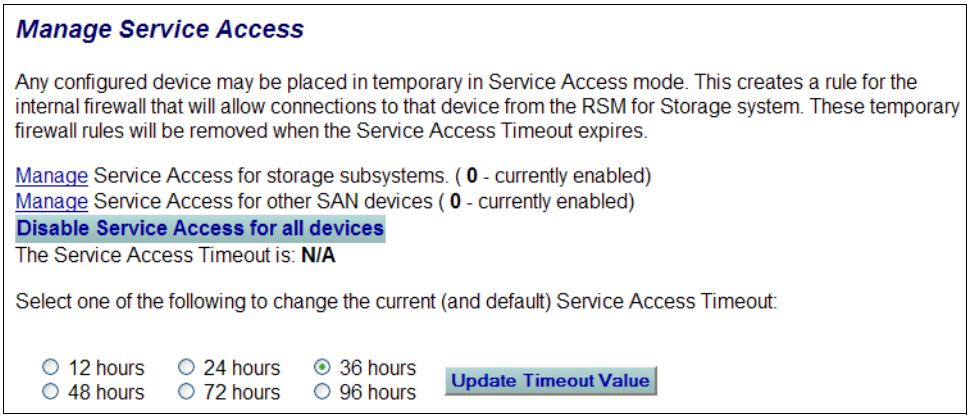


Figure 15-26 Manage Service Access

DS3000 subsystem security

The DS3000 Storage Manager has the ability to require an administrative password in order to make changes to the subsystem configuration. We recommend configuring this password.

The DS3000 also has a controller shell environment that is accessible using a remote login (RLOGIN) client. DS3000 Storage Manager has an option to disable RLOGIN, and we normally recommend disabling RLOGIN.

Refer to 5.4.3, “Setting a storage subsystem password” on page 73 and 5.4.5, “(Optional) Changing the network configuration” on page 76 on how to set the Storage Manager password and to disable remote login.

15.2.6 Activating RSM

The final step is to activate your system. You should complete all other configurations and run a successful Configuration Test before contacting IBM Service to activate RSM.

Make sure Remote Access is enabled by clicking **Remote Access** on the main RSM window to activate it, if required (as shown in 15.2.5, “Remote Access connectivity” on page 376).

Call the number for IBM Service for your region and give the IBM Machine Type and Serial Number of one of the DS3000 storage subsystems to be monitored by RSM for Storage. For support telephone numbers in your country or region, use the following link:

`http://www.ibm.com/planetwide`

Tell the disk support contact person that you are activating a RSM for Storage system.

Provide IBM Service with one or both of the following pieces of access information:

- The phone number of the modem attached to the RSM for Storage system.
- The IP address/port number and any required authentication information to be used for SSH access.

IBM Service will connect to the system, verify that the configuration is correct, send a test alert by e-mail, verify receipt of the alert and associated attachments, and then activate the system, as shown in Figure 15-27.

Note: If you decide to install the RSM for Storage software on another server, you will need to contact IBM Service to obtain a new activation key for the new server.



Figure 15-27 RSM activated

15.2.7 Managing alerts

You can manage alerts as follows:

1. When RSM receives SNMP alerts from one of the defined storage subsystems, an attention mark (!) is shown next to the reporting link, as shown in Figure 15-28. Click **Reporting** to see the alerts.



Figure 15-28 Main menu with storage problems

2. The reporting and alert page shows all those subsystems that have sent alerts sent to RSM (Figure 15-29). Click **View, Acknowledge or Close Alerts**.

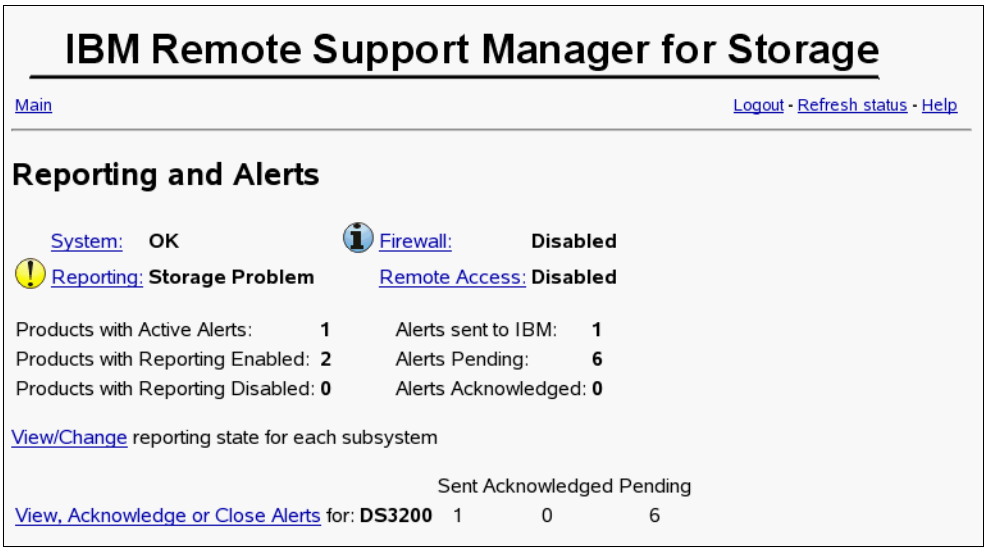


Figure 15-29 RSM storage subsystem with alerts

3. The alert list for the selected storage subsystem (Figure 15-30) shows all alerts that were received by RSM. Click **View** to see the details of an alert.



Figure 15-30 RSM list of alerts

4. The alert details and an error message are shown in Figure 15-31. Click **Acknowledge** to confirm this alert.



Figure 15-31 RSM alert details

5. The alert list of the selected storage subsystem shows the modified status of an alert, as shown in Figure 15-32. When the problem is solved, click **Close**.



Figure 15-32 RSM alert list with an acknowledged alert

6. After the alert is closed, it disappears from the alert list, as seen in Figure 15-33.



Figure 15-33 RSM alert list

7. The main menu status changes so that the attention mark disappears after all problems are solved (Figure 15-34).

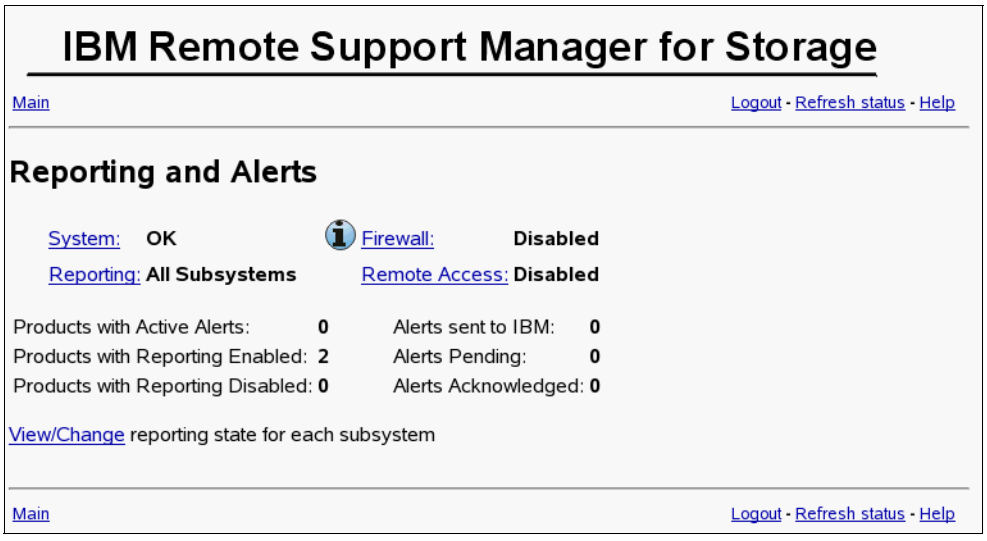


Figure 15-34 RSM main menu

Part 4

Sample configurations

In this part of the book, we show some sample configurations for the IBM System Storage DS3000. You can use these as a basis for setting up your own environment. They are divided up into the different models and attachment methods.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

16

SAS configuration 1 - Windows 2003 host

In this sample scenario, an IBM System x3755 server running Windows Server 2003 will be connected to a dual-controller DS3200 subsystem, using two PCI Express SAS HBAs for I/O path redundancy. You could connect up to three host servers in this manner, but in our sample configuration, we will use just one.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

16.1 Equipment required

We used the following hardware and software components:

- A System x3755 host server, with Windows Server 2003 installed. In addition, Service Pack 2 and the latest Storport fix (KB932755) are also required. The host name is COLORADO.
- Two PCI Express SAS HBAs, P/N 25R8060.
- A DS3200 subsystem with at least four disk drives.
- Two SAS cables.
- The latest SAS HBA drivers. You can download them from the IBM Systems support Web site.
- A LSI Logic SAS HBA Support CD.
- The latest version of DS3000 Storage Manager for Windows (V10.35 at the time of writing).

Figure 16-1 shows our sample hardware setup.

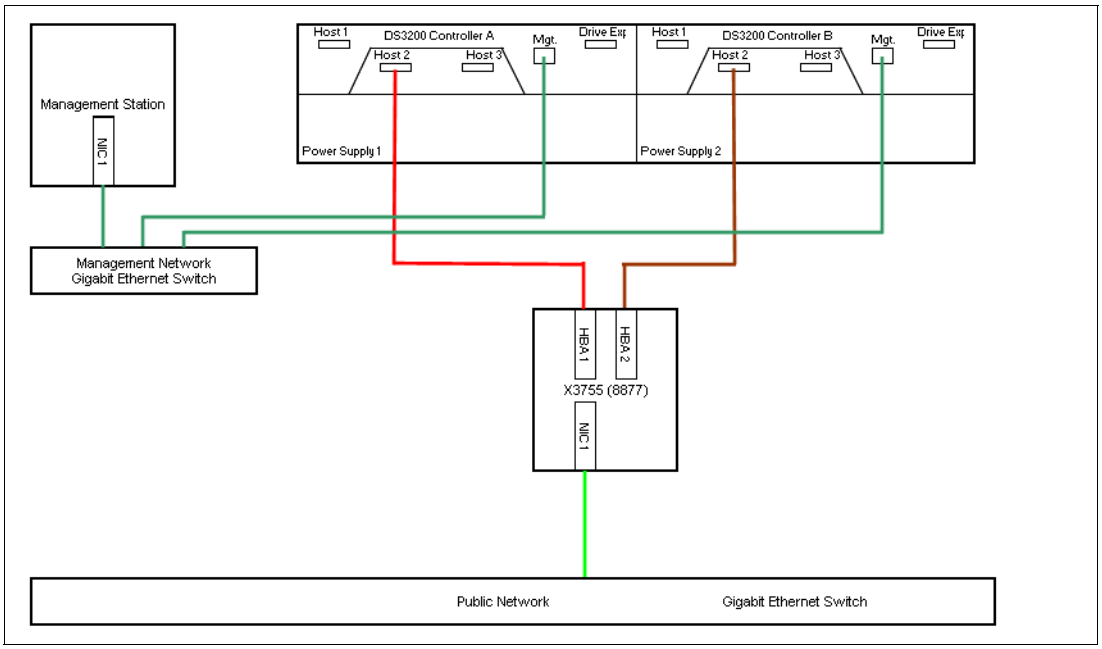


Figure 16-1 Hardware setup

16.2 Installing SAS HBAs

The first main step is to install our two PCI Express SAS HBAs in the x3755 server. The *IBM System x3755 Installation Guide*, found at https://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-65466&brandind=5000008, contains instructions for correct installation of these cards. When the HBAs are installed, we then use two SAS cables to attach each HBA to a separate controller on the DS3200.

We can power on the server and configure the DS3000 storage subsystem.

16.2.1 SAS HBA driver installation

We followed this procedure for the driver installation:

1. After the server boots Windows, the newly installed HBAs are detected, which launches the Found New Hardware Wizard. We downloaded the latest SAS HBA driver for Windows Server 2003 from the support site, so we do not want to connect to the Windows Update Web site (see Figure 16-2). We select **No, not this time**, and click **Next**.



Figure 16-2 Found New Hardware Wizard

2. We saved the downloaded HBA driver in the directory C:\x3755\SASHBAdrv\win2k3. We therefore select **Install from a list or specific location (Advanced)**, as shown in Figure 16-3.

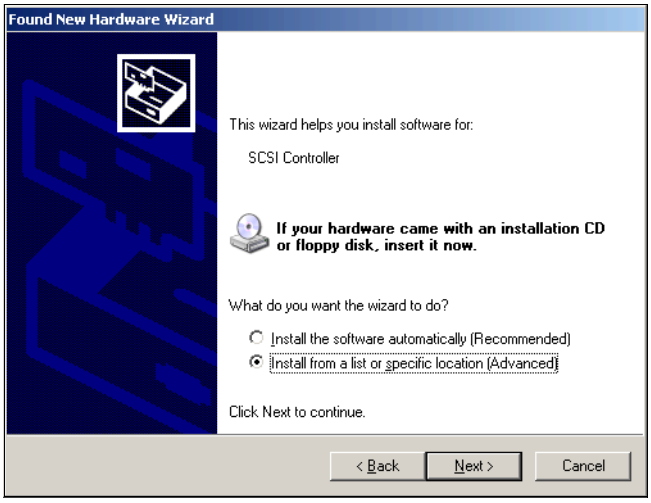


Figure 16-3 Install from specific location

3. In Figure 16-4, we select **Search for the best driver in these locations**, check the box **Include this location in the search**, and enter the directory where the driver code resides as the only location to be searched. This speeds up the installation, as it prevents other directory paths from being searched unnecessarily.

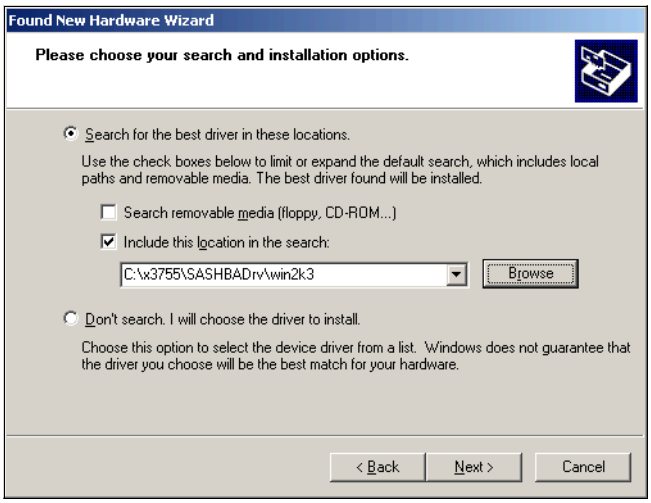


Figure 16-4 Specifying the directory location

4. The wizard now installs the driver, showing a progress bar. When complete, a window similar to Figure 16-5 displays.



Figure 16-5 SAS HBA driver installed

5. We now have to repeat the same driver installation steps for the second SAS HBA.

6. When this is done, verify in Device Manager that both SAS HBAs are listed and available, as shown in Figure 16-6.

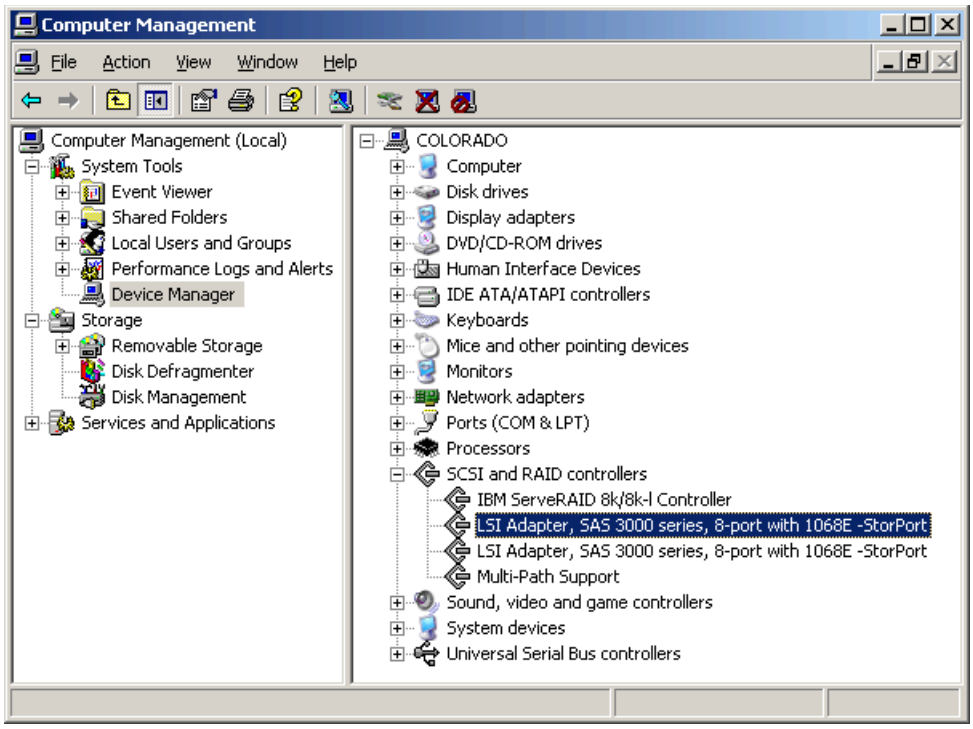


Figure 16-6 Checking the SAS HBA driver status in Device Manager

Both adapters are fine, so we can proceed with the next step.

16.3 Installing DS3000 Storage Manager host software

We will install the host server components of DS3000 Storage Manager. These components include:

- SMagent
- SMutil
- Multipath support
- SMclient (optional)

At the time of writing, in-band management is not supported on the DS3200, so we will not install the SMagent. We will be using out-of-band management.

We will install SMutil, because this will allow us to run the *hot_add* utility and hence avoid the necessity to reboot the host server when adding new logical drives.

Since the host server contains two SAS HBAs, we require the multipath support.

We will also install SMclient, although this component is usually not installed in the host server, but rather on a management workstation. In our case, we want to run the DS3000 Storage Manager GUI in the host server itself.

For the exact installation steps, see Chapter 6, “Installing IBM System Storage DS3000 Storage Manager” on page 81.

16.4 Host access configuration

Detailed steps for creating the logical drives are given in 9.3, “Configure storage” on page 155.

We used the DS3000 Storage Manager to create a sample logical drive with a size of 10 GB. We use the Locate Arrays task to identify our newly created logical drive, as shown in Figure 16-7.

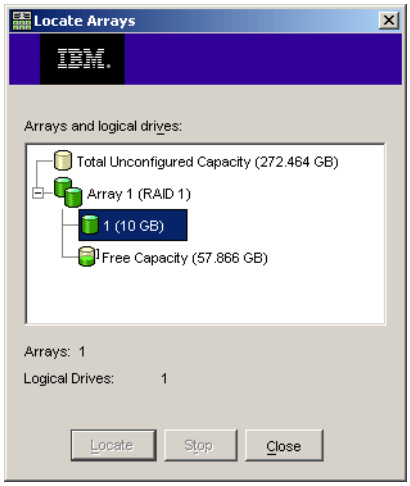


Figure 16-7 Logical drive 1 with a size of 10 GB

Our host server has to be given access to this logical drive. The first thing to do is to configure the host server in Storage Manager, so that we can map it to the logical drive later.

Automatic host access configuration

Follow these steps for the configuration:

1. The easiest way is to use the Automatic Host Access Configuration task. For this to work, the SMagent service must be installed and running in the host. An easy way to check this is by looking at the Application Event Log, as shown in Figure 16-8 on page 397.

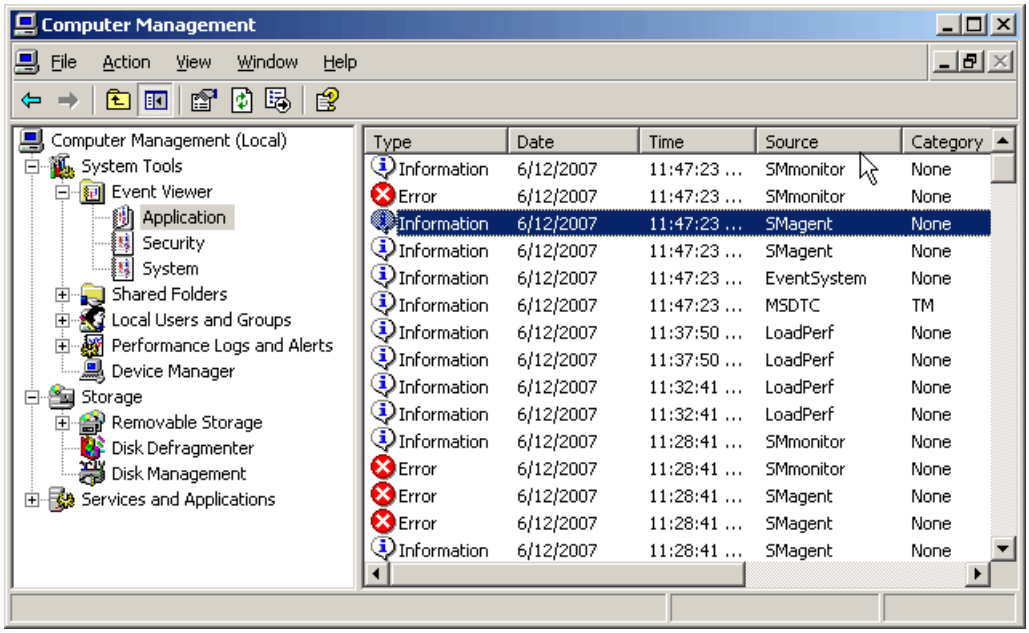


Figure 16-8 Application Event Log - SMagent service messages

When the SMagent service starts up correctly, the Application Event Log contains two messages from this service. The second message is the confirmation that the SMagent service is running. The first message contains confirmation that the SMagent service made contact with a DS3000 subsystem through the access logical drive (or UTM logical drive), as you can see in Figure 16-9.

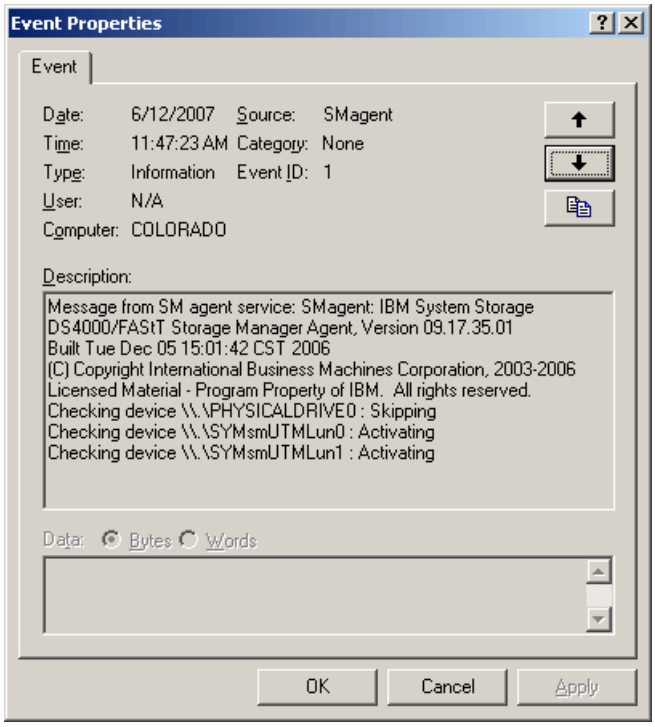


Figure 16-9 SMagent service identified the DS3000 through the access logical drive

This looks fine, so next we try the automatic host access configuration task.

2. As seen in Figure 16-10, the task correctly identified our host server named COLORADO and placed it in the Available Hosts list.

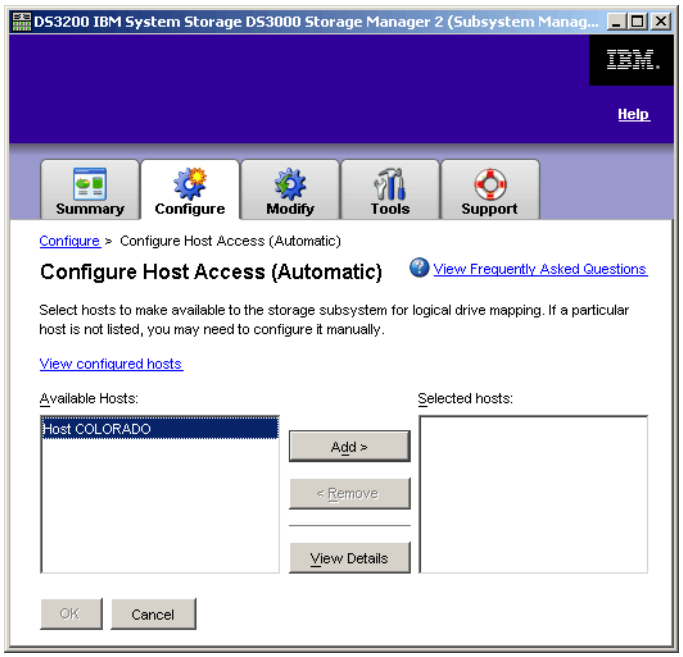


Figure 16-10 Configure Host Access (Automatic)

3. Click **Add** to place the host in the Selected Hosts window, then click **OK**. This completes the automatic host configuration. Figure 16-11 confirms that our configuration was successful.

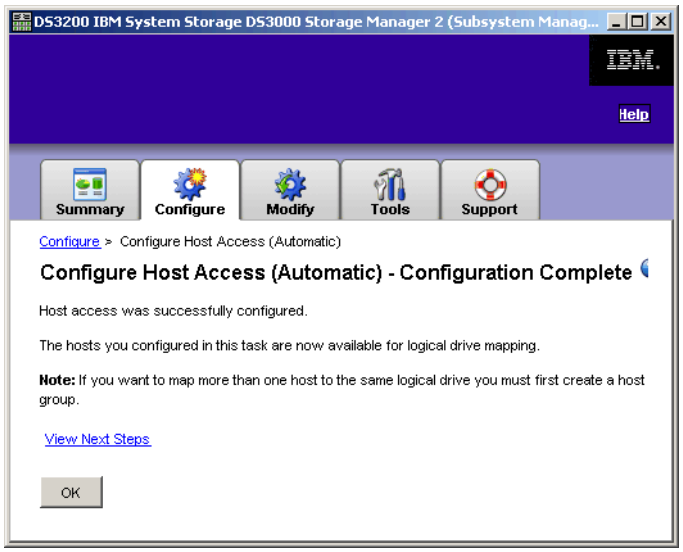


Figure 16-11 Automatic host configuration is complete

4. Let us now verify the host server definition. We expect to see two HBA host ports, as the server contains two SAS HBAs.

   We use the Edit Topology task. As you can see in Figure 16-12, the host server COLORADO definition indeed contains two SAS HBA host ports, with their identifiers. SAS HBAs use Physical SAS addresses as HBA identifiers.

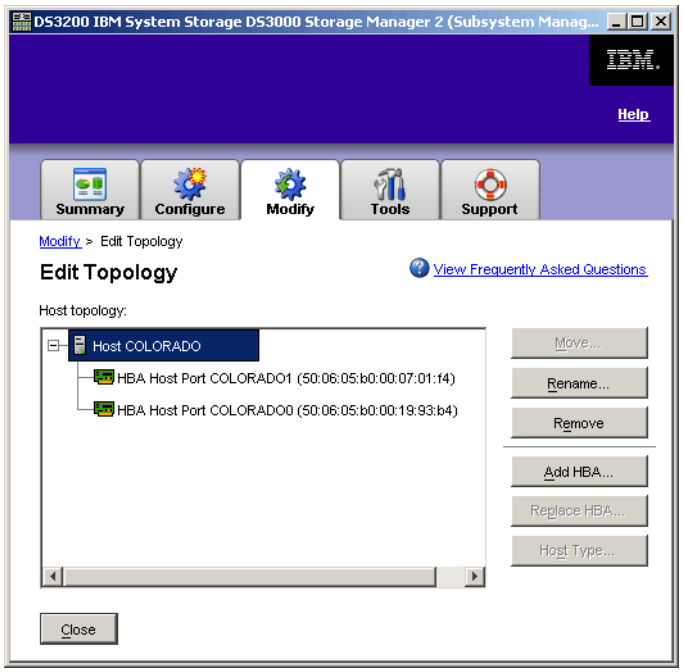


Figure 16-12 Host server COLORADO

The automatic host configuration was successful. We proceed with the host-to-logical-drive mapping.

Mapping the host server to the logical drive

All pieces are in place now. We have created a logical drive and defined a host server. Now we have to map the host to the logical drive. We use the Create Host-To-Logical-Drive Mappings task:

1. Select the host server. In our sample configuration, there is only one host, COLORADO (see Figure 16-13). Highlight the host and click **Next**.

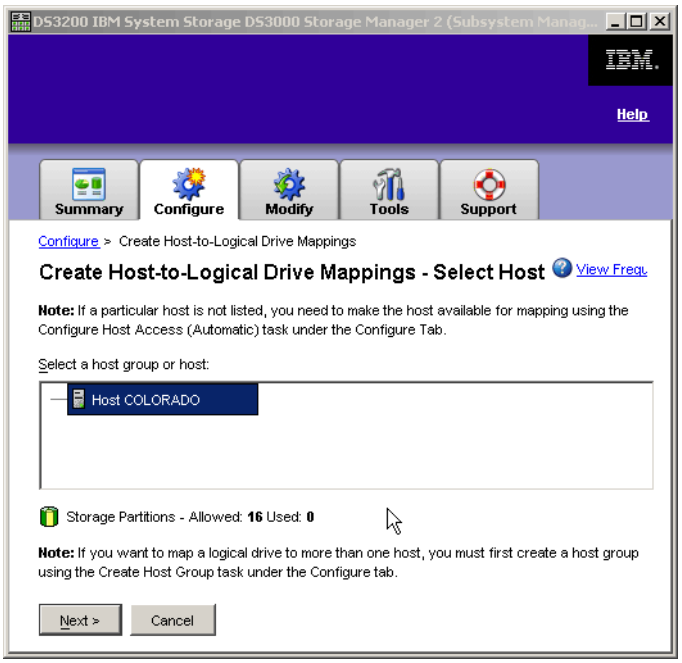


Figure 16-13 Select a host server

2. Select the logical drives that will be mapped to our host. In our case, we highlight logical drive 1 and assign LUN 0 to it, as shown in Figure 16-14.

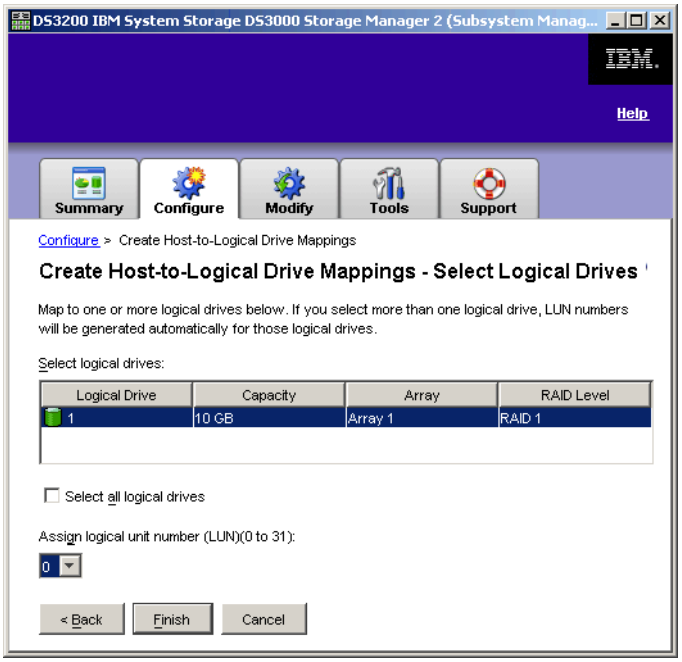


Figure 16-14 Select logical drives

Click **Finish** to complete the process.

3. You will see the progress bar while the operation takes place. Click **OK** in the progress window to display the completion window (see Figure 16-15).

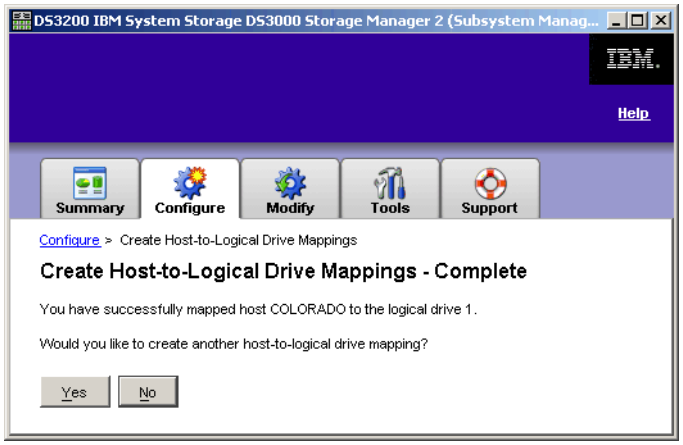


Figure 16-15 Host-to-Logical-Drive Mappings - Complete

4. Our host server COLORADO is now mapped to logical drive 1. To confirm this, we run the Host-to-Logical-Drive-Mappings task in the Summary tab. As you can see in Figure 16-16, the mapping has been set up correctly.

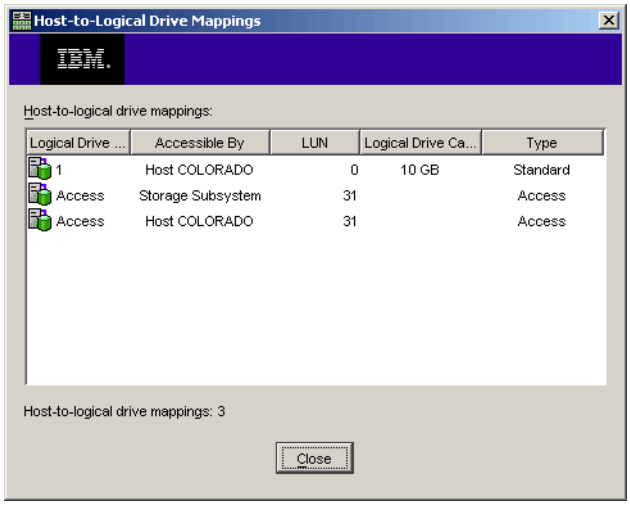


Figure 16-16 Host server COLORADO mapped to logical drive 1

Configure the disk space in Windows 2003

The logical drive now must be recognized in Windows before we can start using the disk space. One possibility is to shut down and reboot the host server. But there is a better alternative; We use the **`hot_add`** utility (a part of SMutil).

By default, the **`hot_add`** executable is located in C:\Program Files\IBM_DS3000\util. In a command prompt window, change to this directory. If you installed the DS3000 Storage Manager in a non-default directory, change to that directory instead, and then to subdirectory util.

Now run **`hot_add`**.

When the utility completes, you will be able to use the Disk Management applet in the Computer Management utility and configure the disk space for operating system use. This includes the following steps:

1. Initialize and optionally convert to dynamic disk.

   This task is done by the Initialize and Convert Disk Wizard. When the wizard finishes, the Disk Management applet looks similar to Figure 16-17.

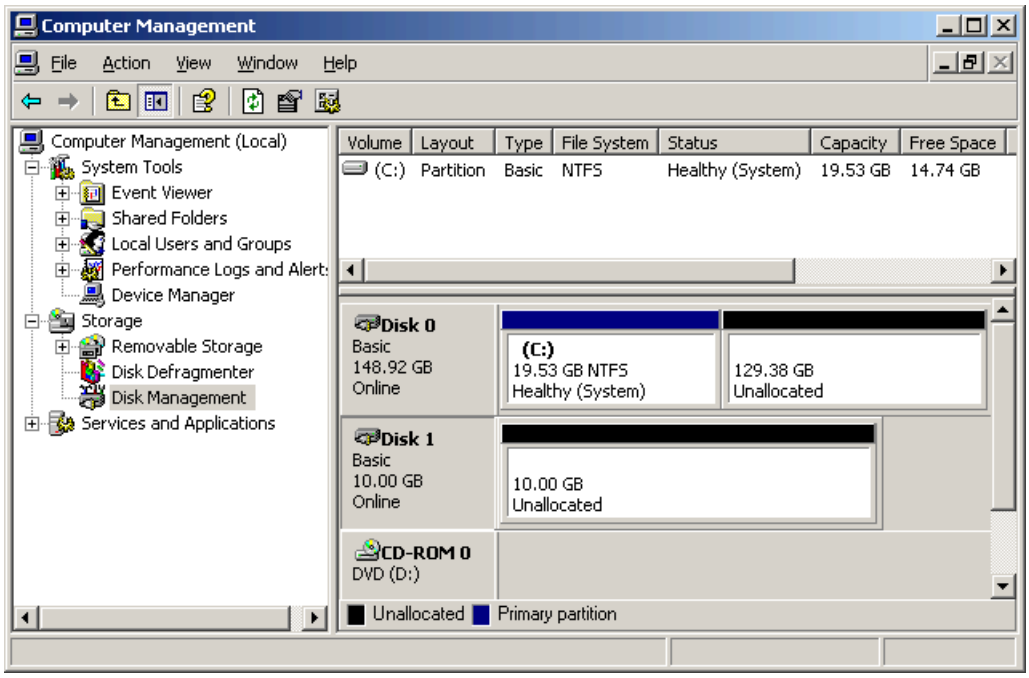


Figure 16-17 Disk Management applet - initialized disk

   As you can see, our 10 GB logical drive is recognized as Disk 1, but it does not contain any partition yet; it is unallocated.

2. Create and format a partition.

   Right-click in the unallocated disk space and select **New Partition...** from the context menu (Figure 16-18 on page 403).

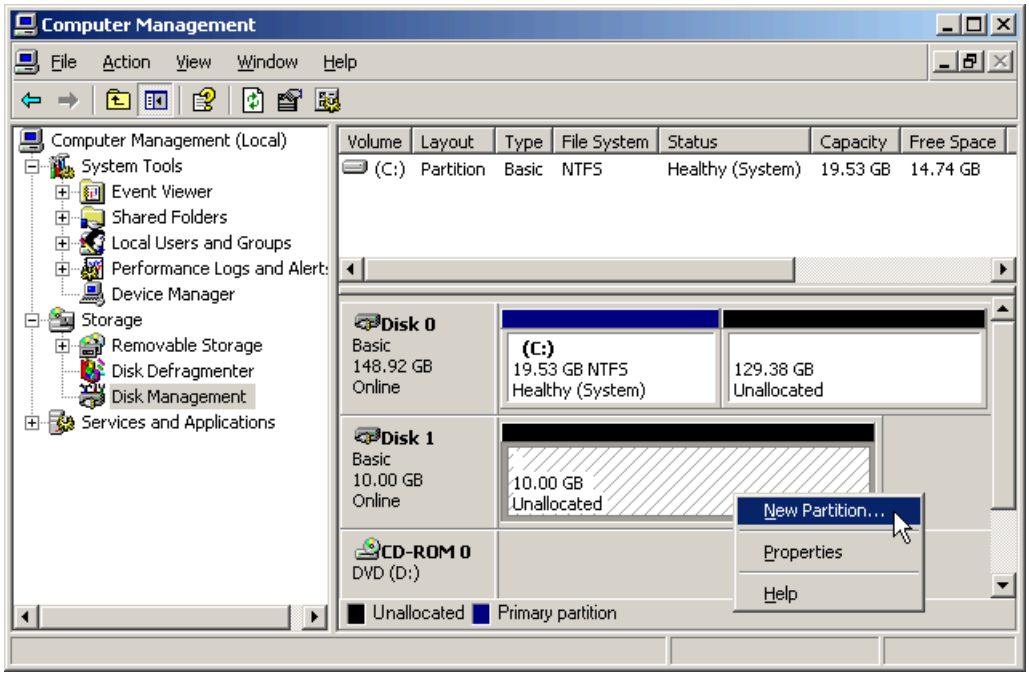


Figure 16-18 Creating a new partition

This launches the New Partition Wizard.

3. Follow the process under the guidance of the wizard to define the partition size, drive letter, file system, and volume label. The partition will be created and formatted. The Disk Management applet now displays the new partition (allocated drive E), which is ready to be used (see Figure 16-19).

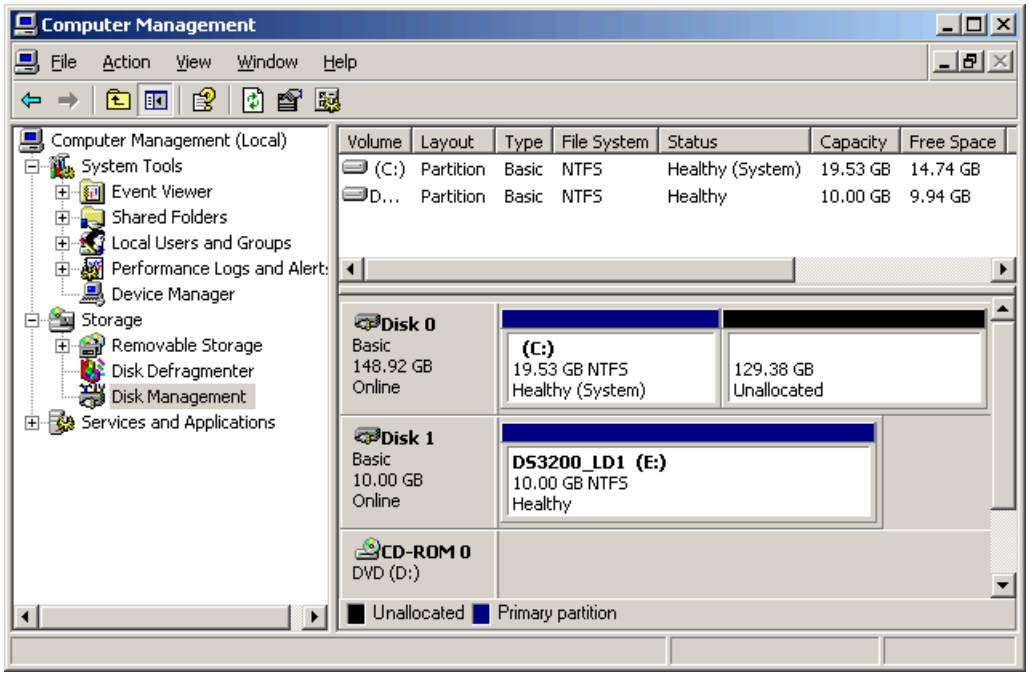


Figure 16-19 New partition on the logical drive

4. We can now access the logical drive on the DS3200 as drive letter E:. We show this in Figure 16-20.

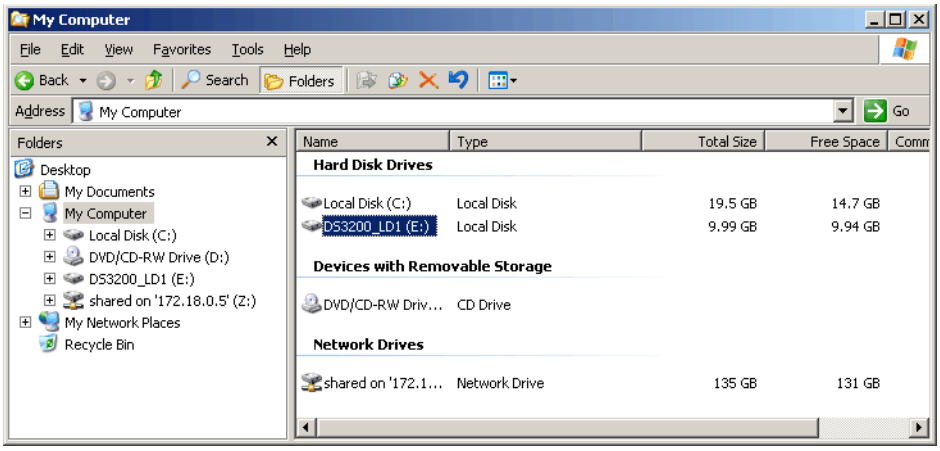


Figure 16-20 Logical drive is ready for use

17

SAS configuration 2 - Linux host

In this chapter, we describe explain how to connect to the DS3200 from a host server running Red Hat Enterprise Linux 4. The host server is IBM System x3755, and it contains two PCI Express SAS HBAs. It is attached to a dual controller model of DS3200.

Note: At the time of writing, RDAC Version 09.03.0B05.0023 was the used to illustrate the steps below. Refer to the IBM Systems support Web site, as described in 17.4, “Installing RDAC for Linux” on page 408, and download the latest available version of RDAC.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

17.1 Equipment required

We used the following hardware and software components:

- An IBM System x3755 server, installed with RHEL 4 U3 (or higher).
- Two PCI Express SAS HBAs, P/N 25R8060.
- Two SAS cables.
- The latest SAS HBA driver for Linux kernel 2.6. The driver is available for download on the IBM Systems support Web site.
- The latest version of DS3000 Storage Manager for Linux (at the time of writing, this is Version 10.35).
- The latest version RDAC for Linux (V09.03.0B05.0023 at the time of writing).

We show our sample hardware setup in Figure 17-1.

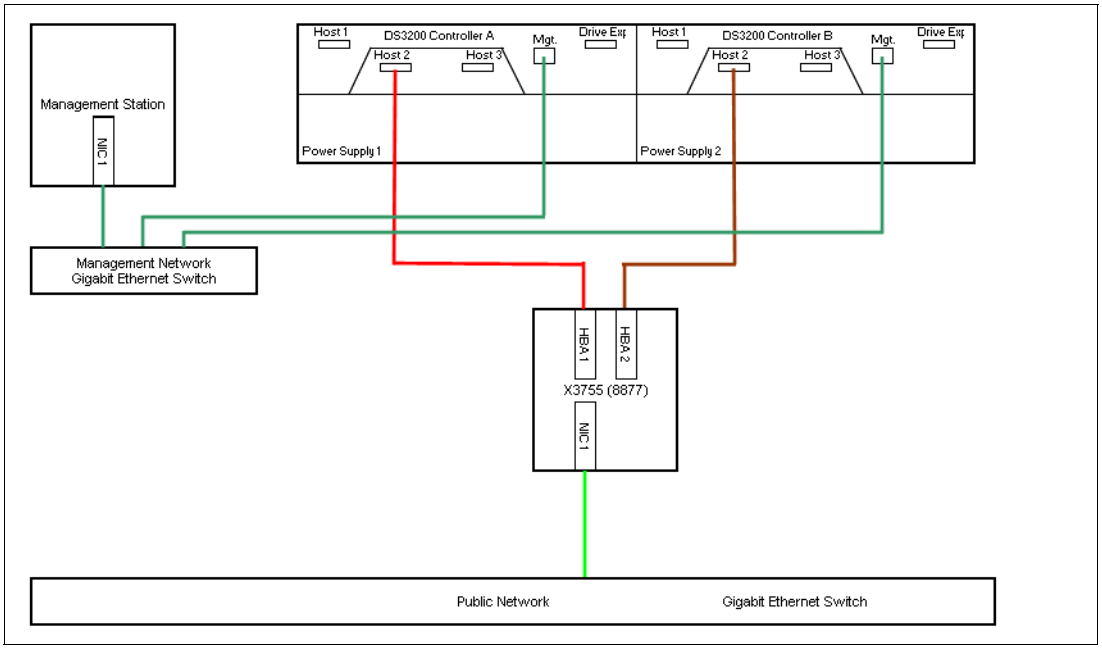


Figure 17-1 Hardware setup

17.2 Installing SAS HBAs

We have to install the two PCI Express SAS HBAs in the host server before we can proceed with other tasks. While this is not a difficult task, we do recommend that you consult the user's guide for the host server and follow the instructions for options installation. Before powering the host server on, we connect the SAS HBAs to the DS3200 controllers (each HBA to one controller). The next step is the SAS HBA driver installation.

17.2.1 SAS HBA driver installation

The latest Linux SAS HBA driver is available for download on the IBM Systems support Web site. A set of precompiled binary RPMs for different kernel flavors (default, SMP, and hugemem) is included, as well as the source RPM for use with kernel versions that do not match the precompiled driver versions.

In our case, we can simply install a precompiled driver. This can be done in two ways:

- Run the install.sh script, included with the driver:

  ```
  ./install.sh --update
  ```

- Alternatively, run the **rpm** command to install the driver package:

  ```
  rpm -i lsi-mptlinux-smp-3.02.83.05_2.6.9_22.EL-1.i686.rpm
  ```

There are several ways to verify that the driver package is correctly installed. One way is with the following command:

```
rpm -qa | grep lsi
```

The output should show the driver package information:

```
lsi-mptlinux-smp-3.02.83.05_2.6.9_22.EL-1
```

You could also run the **lsmod** command. The output should show that the following modules are running: mptsas, mptscsi, mptscsih, and mptbase.

17.3 Installing DS3000 Storage Manager software

A host server running Windows can have the following DS3000 Storage Manager components installed:

- SMagent, if in-band management is required.
- SMutil.
- SMclient, if we want to use the DS3000 Storage Manager GUI in the host itself.
- Multipath support.

On Linux, this is slightly different. The multipath driver, RDAC, is not included with the DS3000 Storage Manager for Linux; it is available as a separate package instead. We install this in 17.4, “Installing RDAC for Linux” on page 408.

We install SMclient, because we want to run the DS3000 Storage Manager GUI on the Linux server itself. We install SMutil as well.

We do not install SMagent, since at the time of writing, the DS3200 does not support in-band management. We will therefore not install the SMagent and our management will be performed out-of-band.

The latest DS3000 Storage Manager for Linux is available on the IBM Systems support Web site. At the time of writing, V10.35 is the current version for the Linux 2.6 kernel.

The versions reflected below might differ depending on the latest versions available for download from the IBM support Web site:

1. The Storage Manager package is available as a compressed tar file. Use the following command to uncompress it:

   ```
   tar -zxvf ibm_sw_ds3k_0217b505_linux2.6_anycpu.tar.gz
   ```

2. Now look for the Storage Manager InstallAnywhere (SMIA) installation script file. The file name will be something similar to SMIA-LINUX03.35.A5.11.bin.
3. Run this script file to start the installation script. This will launch the GUI installation process:

   ```
   sh ./SMIA-LINUX03.35.A5.11.bin
   ```

4. We do not show the installation windows and steps here, because the installation process is covered in detail in 6.2, “Installing Storage Manager on Linux” on page 90.

Once the DS3000 Storage Manager is installed, we have to add the DS3200 subsystem in the Enterprise Management Window and prepare the logical drives for our Linux server. For simplicity’s sake, we create just one logical drive of 10 GB.

We also have to define our Linux host server in the Storage Manager and map it to the logical drive. However, no I/O should be issued to the logical drive yet; we only start using it after the RDAC installation.

We already explained how to add storage subsystems, define hosts, prepare logical drives, and map them in Part 3, “Administration” on page 109, so we do not repeat these steps here. Figure 17-2 shows our 10 GB logical drive named 1 and mapped to the host server COLORADO.

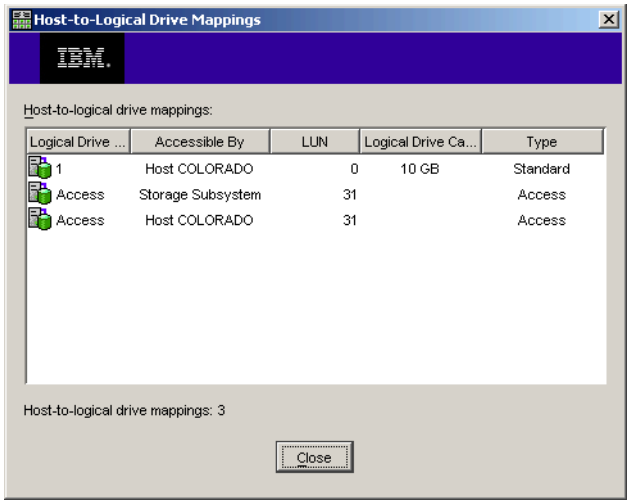


Figure 17-2 Drive mapping

17.4 Installing RDAC for Linux

RDAC for Linux is available for download on the IBM Systems support Web site. The IBM Web page leads you to the actual download Web page URL:

http://www.lsi.com/rdac/ds3000.html

Download the appropriate RDAC for your Linux kernel version (either 2.4 or 2.6) and continue with the next section. In our case, we use RHEL 4, so we downloaded the RDAC for kernel Version 2.6.

17.4.1 Building the RDAC driver

The RDAC package is available in a compressed tar file. Follow these steps to obtain it:

1. Unpack the file using the following command:

   ```
   tar -zxvf rdac-LINUX-09.03.0B05.0023-source.tar.gz
   ```

   The source files will be unpacked in the linuxrdac-09.03.0C05.0023 subdirectory of the current directory. When the uncompress is complete, change to this subdirectory.

2. Use the following command to remove old driver modules:

   ```
   make clean
   ```

3. Compile the driver modules and build the RDAC driver. Use this command:

   ```
   make
   ```

17.4.2 Installing the RDAC driver

We have to copy the driver modules to the kernel module tree and build the new RAMdisk image, which includes the RDAC driver modules. Enter the command:

```
make install
```

This will produce a lengthy output in the console window, as the driver modules are applied to the kernel tree. The last few lines will tell you how to add the new boot menu option to the boot loader, as shown in Example 17-1.

Example 17-1 RDAC installation

```
Creating new MPP initrd image...
You must now edit your boot loader configuration file, /boot/grub/menu.lst, to
add a new boot menu, which uses mpp-2.6.9-34.ELsmp.img as the initrd image.
Now Reboot the system for MPP to take effect.
The new boot menu entry should look something like this (note that it may
vary with different system configuration):
...
title Red Hat Linux (2.6.9-34.ELsmp) with MPP support
root (hd0,5)
kernel /vmlinuz-2.6.9-34.ELsmp ro root=LABEL=RH9
initrd /mpp-2.6.9-34.ELsmp.img
...
MPP driver package has been sucessfully installed on your system.
```

The new RAMdisk image filename is mpp-2.6.9-34.ELsmp.img. We have to add the following coding to the menu.lst file, as shown in Example 17-2.

Example 17-2 Modify menu.lst file

```
title Red Hat Linux (2.6.9-34.ELsmp) with MPP support
        root (hd0,1)
        kernel /vmlinuz-2.6.9-34.ELsmp ro root=LABEL=/1 pci=nommconf rhgb quiet
        initrd /mpp-2.6.9-34.ELsmp.img
```

We now restart the system using the new boot option.

17.4.3 Verification steps

When the server restarts, use the **lsmod** command to verify that the RDAC (MPP) modules are running. The following modules should be listed:

- mppUpper
- mppVhba
- sg
- sd_mod
- scsi_mod

If all these modules are running, the RDAC should now be managing access to the DS3200 logical drive. To verify this, use the following command:

```
ls -lR /proc/mpp
```

You should see an output similar to Example 17-3.

Example 17-3 Output of ls -lR /proc/mpp

```
/proc/mpp:
total 0
dr-xr-xr-x  4 root root 0 Jun 18 12:28 DS3200

/proc/mpp/DS3200:
total 0
dr-xr-xr-x  3 root root 0 Jun 18 12:28 controllerA
dr-xr-xr-x  3 root root 0 Jun 18 12:28 controllerB
-rw-r--r--  1 root root 0 Jun 18 12:28 virtualLun0

/proc/mpp/DS3200/controllerA:
total 0
dr-xr-xr-x  2 root root 0 Jun 18 12:28 mptsas_h0c0t17

/proc/mpp/DS3200/controllerA/mptsas_h0c0t17:
total 0
-rw-r--r--  1 root root 0 Jun 18 12:28 LUN0
-rw-r--r--  1 root root 0 Jun 18 12:28 UTM_LUN31

/proc/mpp/DS3200/controllerB:
total 0
dr-xr-xr-x  2 root root 0 Jun 18 12:28 mptsas_h1c0t33

/proc/mpp/DS3200/controllerB/mptsas_h1c0t33:
total 0
-rw-r--r--  1 root root 0 Jun 18 12:28 LUN0
-rw-r--r--  1 root root 0 Jun 18 12:28 UTM_LUN31
```

The Linux server can now access the logical drive, which is presented as LUN 0. The logical drive is accessible across both paths, controllerA and controllerB, and the RDAC (or MPP) makes sure that the dual path access is handled correctly.

You can also see the Access Logical Drive, presented as UTM_LUN31.

17.4.4 Configuring disk space in Linux

The logical drive is now visible in the host server. We can now create partitions and format them. Linux provides various tools for partition management, for example, **fdisk** or **parted**.

The logical drive is presented to Linux as /dev/sdb. We created two partitions, which appear under /dev/sdb1 and /dev/sdb2.

We can now mount these partitions and start using the disk space. For example, you could use the following commands to mount the partitions:

```
mount /dev/sdb1 /mnt/DS32000_1
mount /dev/sdb2 /mnt/DS3200_2
```

18

SAS configuration 3 - Windows 2008 on an IBM BladeCenter HS21 server

In this chapter, we describe and explain a sample configuration that shows how to connect logical drives configured on an IBM System Storage DS3200 to a Windows Server 2008 operating system running on an IBM BladeCenter HS21 server connected with SAS Connectivity Modules.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

18.1 Equipment required

We used the following hardware and software components:

- IBM BladeCenter E and HS21 BladeCenter installed with Windows Server 2008.
- IBM BladeCenter Advance Management Module (AMM) installed in BladeCenter.
- Two IBM BladeCenter SAS Connectivity modules.
- IBM BladeCenter SAS Expansion card (CFFv), P/N 43W3974.
- DS3200 System Storage subsystem.
- The latest version of DS3000 Storage Manager for Windows 2008 (at the time of writing, this is Version 10.35).

Figure 18-1 shows our sample hardware setup.

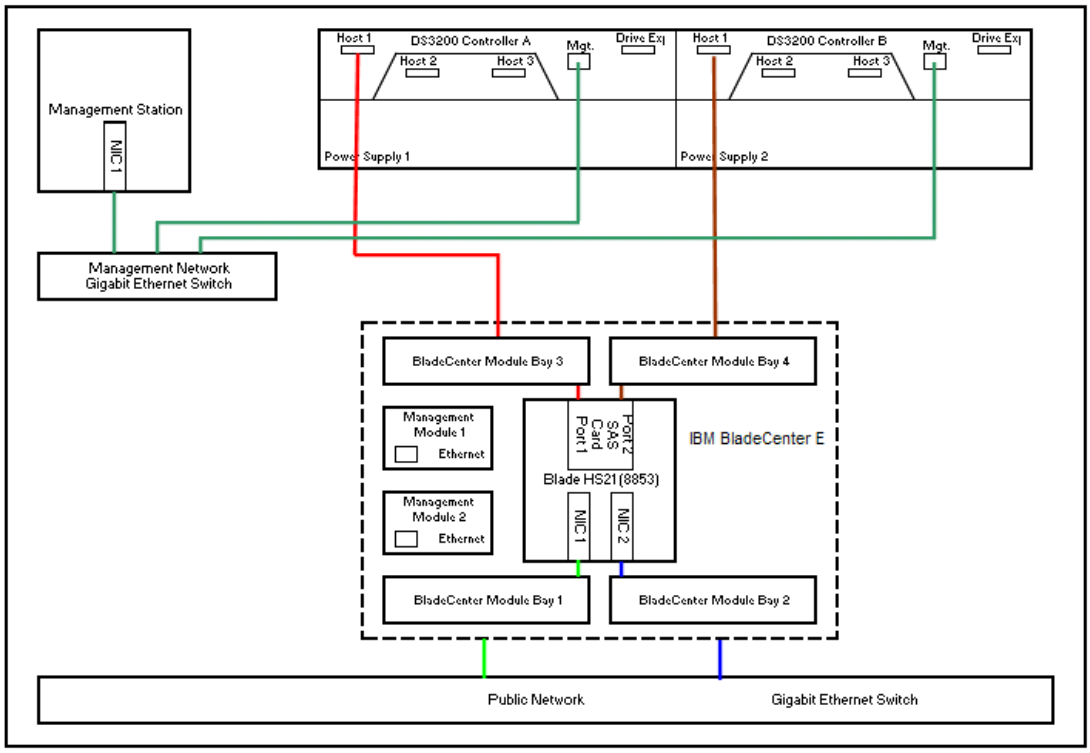


Figure 18-1 Hardware setup

18.2 IBM BladeCenter setup

Here we discuss the IBM BladeCenter setup.

18.2.1 Installing Windows Server 2008

Follow the operating system installation instructions that are available for each IBM BladeCenter blade and IBM System x server. The installation guides can be found in the “Install/use” section of each product’s support Web sites.

For IBM BladeCenter HS21 (8853) with Microsoft Windows Server 2008 setup, we used the instructions found at:

https://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-5074391&brandind=5000020

18.2.2 HS21 SAS Expansion Cards

We have to install the IBM BladeCenter SAS Expansion Card in the HS21 BladeCenter host before we can proceed with other tasks. While this is not a difficult task, we do recommend that you consult the user's guide for the host server and follow the instructions for options installation. The next step is the SAS Expansion Card driver installation.

Important: The connectivity modules in I/O module bay 3 and I/O module bay 4 and all expansion cards in the BladeCenter unit must use the same interface type. Therefore, you must install SAS expansion cards before you install connectivity modules in the blade servers in your BladeCenter unit. For more information about the SAS expansion card, see the *Installation and User's Guide* for the SAS expansion card at http://www.ibm.com/systems/support/.

Note: The BladeCenter SAS Expansion Card is a dual port card. Port #1 connects to SAS Connectivity Module in BladeCenter I/O module bay 3 and port #2 connects to SAS Connectivity Module in BladeCenter module bay 4 respectively.

18.2.3 Recording the SAS Expansion Card WWPN

The following example demonstrates how to record the SAS Expansion Card WWPN for later use in setting up the host-to-LUN mappings in the DS3200 System Storage Subsystem.

1. Turn on or restart the HS21 BladeCenter host.

2. Press <Ctrl + C> to enter the LSI Logic Configuration Utility, as shown in Figure 18-2.

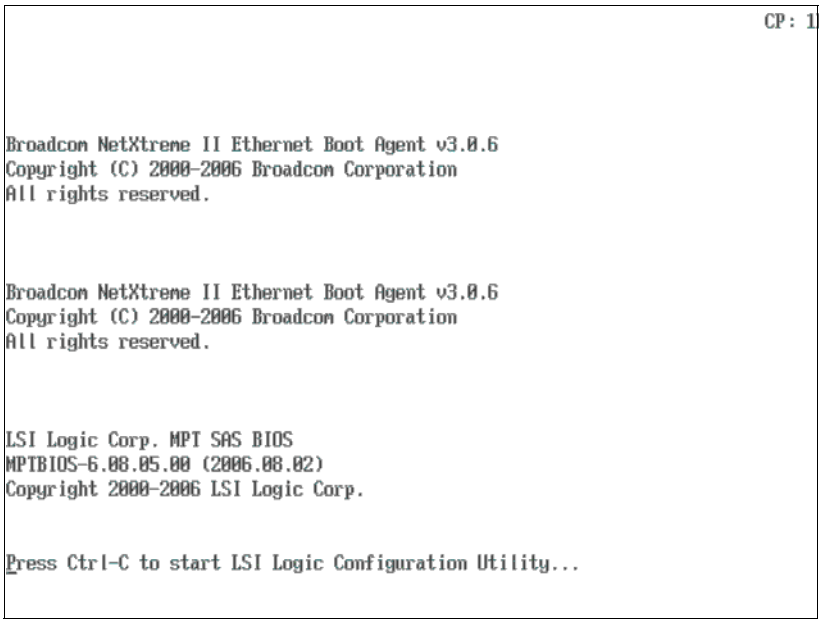


Figure 18-2 SAS Expansion Card - LSI Logic Configuration Utility

3. The following menu is displayed after pressing Crtl + C, as shown in Figure 18-3.

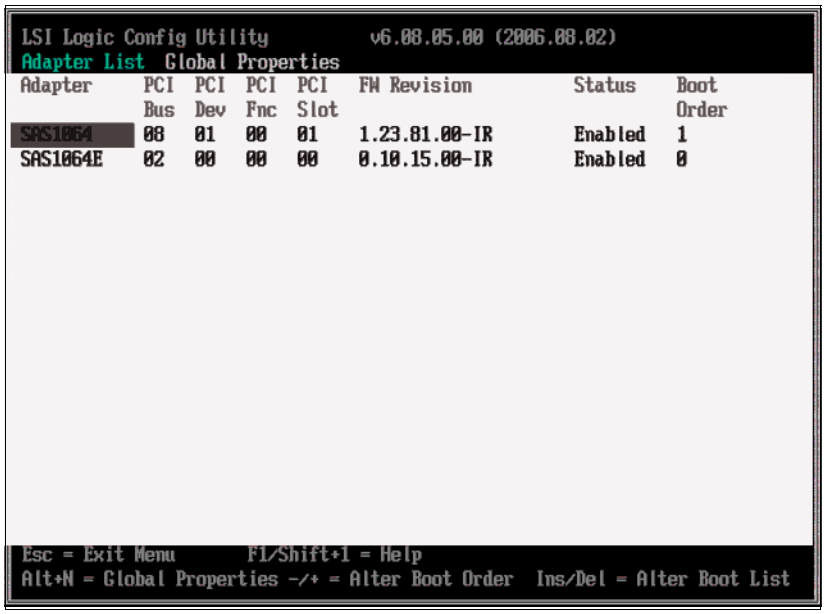


Figure 18-3 SAS Expansion Card - LSI Logic Configuration Utility Menu

4. Press the Enter key to select the SAS adapter internally connected to the SAS connectivity module in I/O Module bay 3, as shown in Figure 18-4 on page 417.



Figure 18-4 SAS Expansion Card - LSI Logic Config Utility adapter selected

Record the Worldwide Port Name (WWPN) of the first port on the SAS Expansion Card. The WWPN is needed for defining host ports on the DS3200 System Storage subsystem. The WWPN can also be retrieved from the SAS Connectivity Module Web interface.

The name of the SAS adapter for the expansion card is SAS1064 and is visible in the Adapter List screen. To determine if the SAS adapter is the expansion card, select a SAS adapter and use the View SAS Topology screen to display whether the SAS adapter is connected to the internal hard disk drives or to the SAS connectivity modules in your BladeCenter chassis, as shown in Figure 18-5.

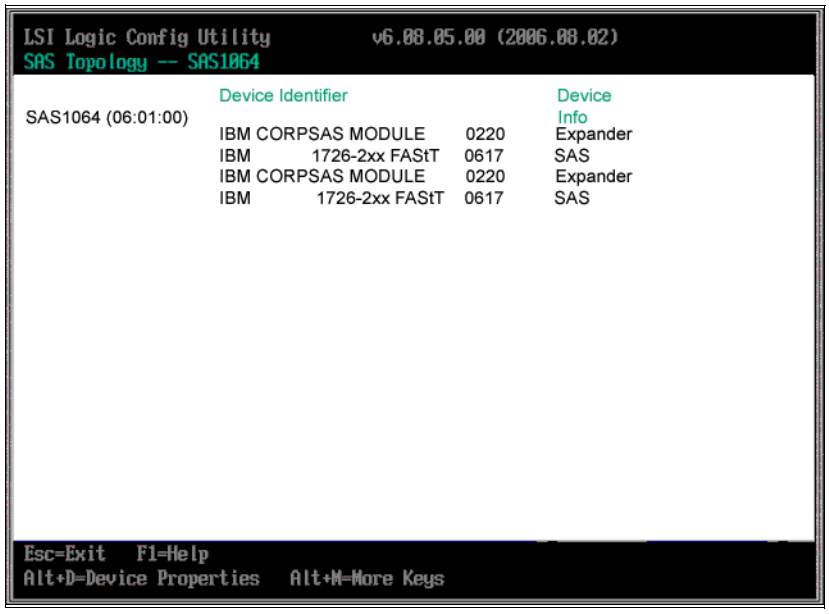


Figure 18-5 SAS Expansion Card - LSI Logic Config Utility adapter confirmation

18.2.4 HS21 SAS Expansion Card device driver

In this example, we use the SAS Expansion Card device driver that is part of the Windows Server 2008 installation. The SAS Expansion Card is installed and Windows Server 2008 recognizes the hardware and applies the appropriate device driver. At the time of writing, no SAS Expansion Card device driver for Windows Server 2008 was available from IBM, but may be available in future from the http://www.ibm.com/support Web site.

18.2.5 SAS Connectivity Modules

You must install the IBM BladeCenter SAS Connectivity modules only in BladeCenter I/O module bay 3 and I/O module bay 4 of the following supported BladeCenter units;

- BladeCenter Type 8677
- BladeCenter Types 8720 and 8730
- BladeCenter Type 8750
- BladeCenter Type 8852
- BladeCenter Type 8886

Installing a connectivity module in I/O module bay 3 or I/O module bay 4 provides connectivity to the SAS expansion card(s) installed in the blade servers in your BladeCenter unit. Installing two connectivity modules allows you to have two connections to the SAS expansion cards installed in the blade servers.

Important: The connectivity modules in I/O module bay 3 and I/O module bay 4 and all expansion cards in the BladeCenter unit must use the same interface type. Therefore, you must install SAS expansion cards before you install connectivity modules in the blade servers in your BladeCenter unit. For more information about the SAS expansion card, see the *Installation and User's Guide* for the SAS expansion card, found at http://www.ibm.com/systems/support/.

Connect the cables from the DS3200 Storage System controllers A and B to the external port# 1 of the two SAS connectivity modules, as shown in Figure 18-6 on page 419.

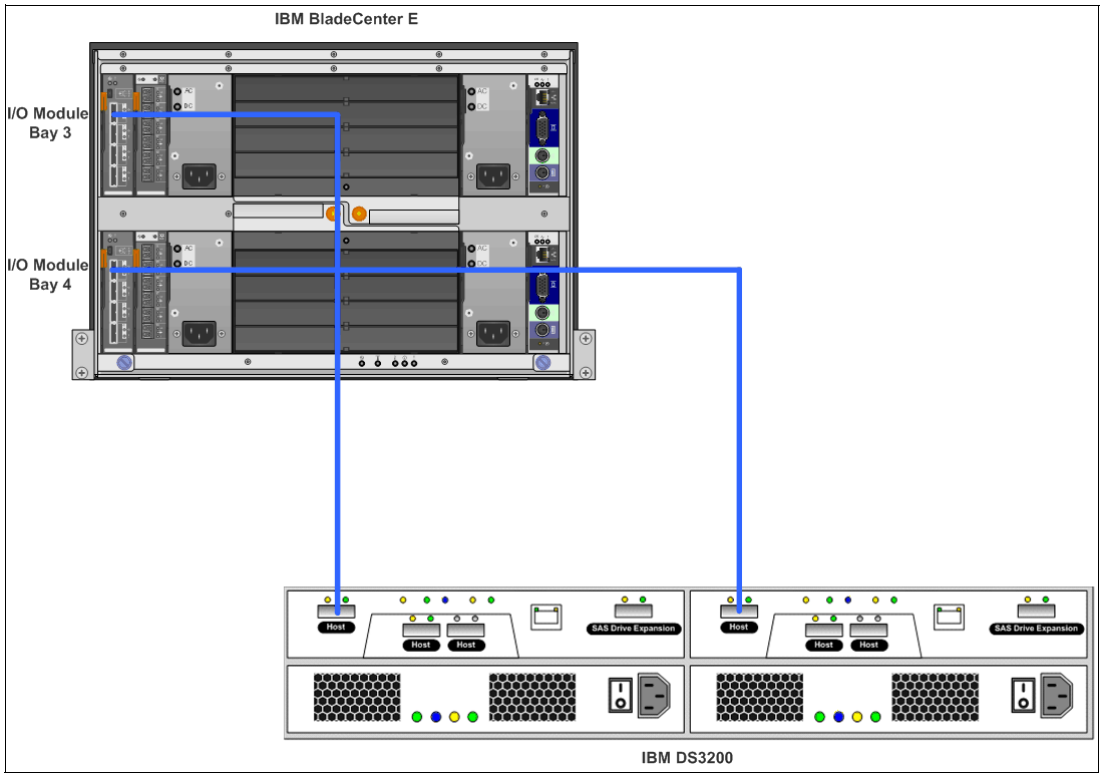


Figure 18-6 BladeCenter and DS3200 connectivity

18.2.6 SAS Connectivity Module firmware update

Ensure that your SAS Connectivity Module is updated with the latest firmware. For the latest firmware update, see http://www.ibm.com/systems/support/.

To update the connectivity-module firmware to the latest version, complete the following steps:

1. Log on to the SAS Connectivity Module using the Web interface with the IP address defined for the connectivity module in the BladeCenter Advance Management Module (AMM), as shown in Figure 18-7. Use USERID as the user Id and PASSW0RD as the password. You can change the password under the Administer Users menu option once you are logged on.

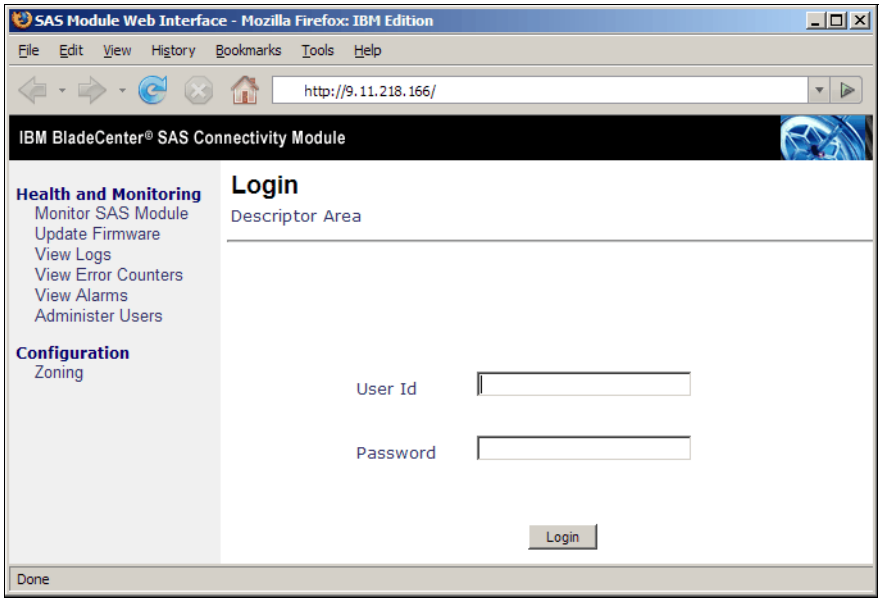


Figure 18-7 SAS Connectivity Module - Login

2. In the Monitor Module window, click **Update Firmware**. The Update Firmware window opens, as shown in Figure 18-8. The current firmware level is also displayed.

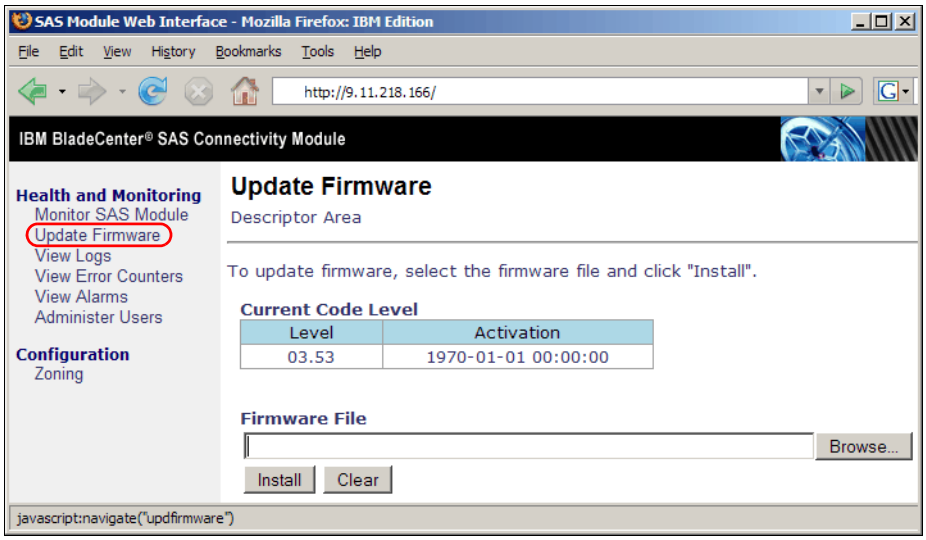


Figure 18-8 SAS Connectivity Module - Update Firmware

3. In the Firmware File field, enter the new firmware file name, or click **Browse** and locate the firmware file.

4. Click **Install** to install the new file. An installation confirmation window opens, as shown in Figure 18-9.

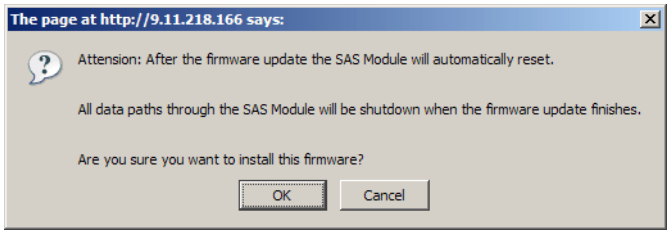


Figure 18-9 SAS Connectivity Module - Installation confirmation

5. Click **OK** or **Cancel**. If the installation of the new firmware file is successful, an installation confirmation window opens, as shown in Figure 18-10. If there are errors during the installation, an error message window opens.

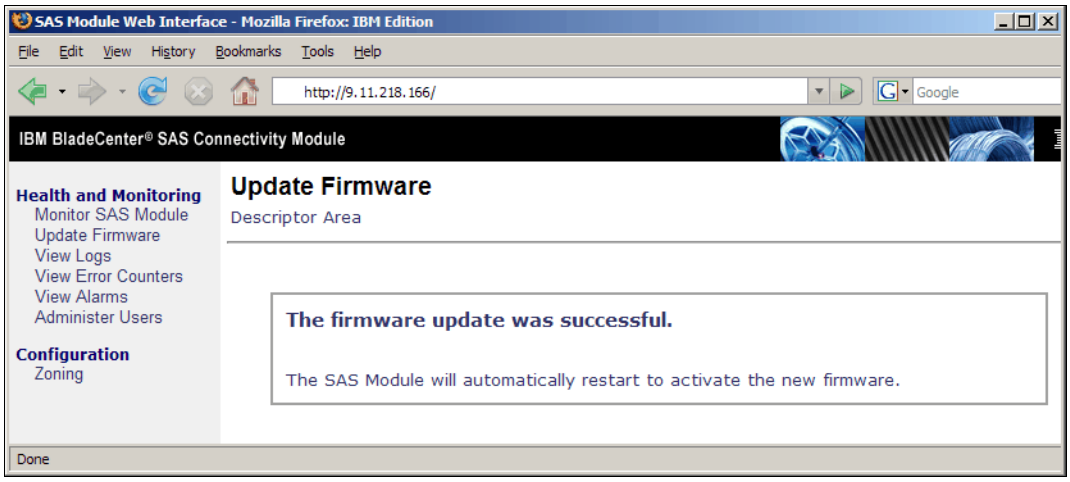


Figure 18-10 SAS Connectivity Module - Update successful

18.2.7 Configuring the SAS connectivity module

Ensure that the external ports on the SAS Connectivity Modules are enabled to allow connectivity from the DS3200 System Storage subsystem and I/O traffic to pass through the module.

From the BladeCenter Advance Management Module Web interface GUI, select **I/O Module Tasks** → **Admin/Power/Restart** to ensure that the external ports for the I/O modules in bays 3 and 4 are enabled, as shown in Figure 18-11.

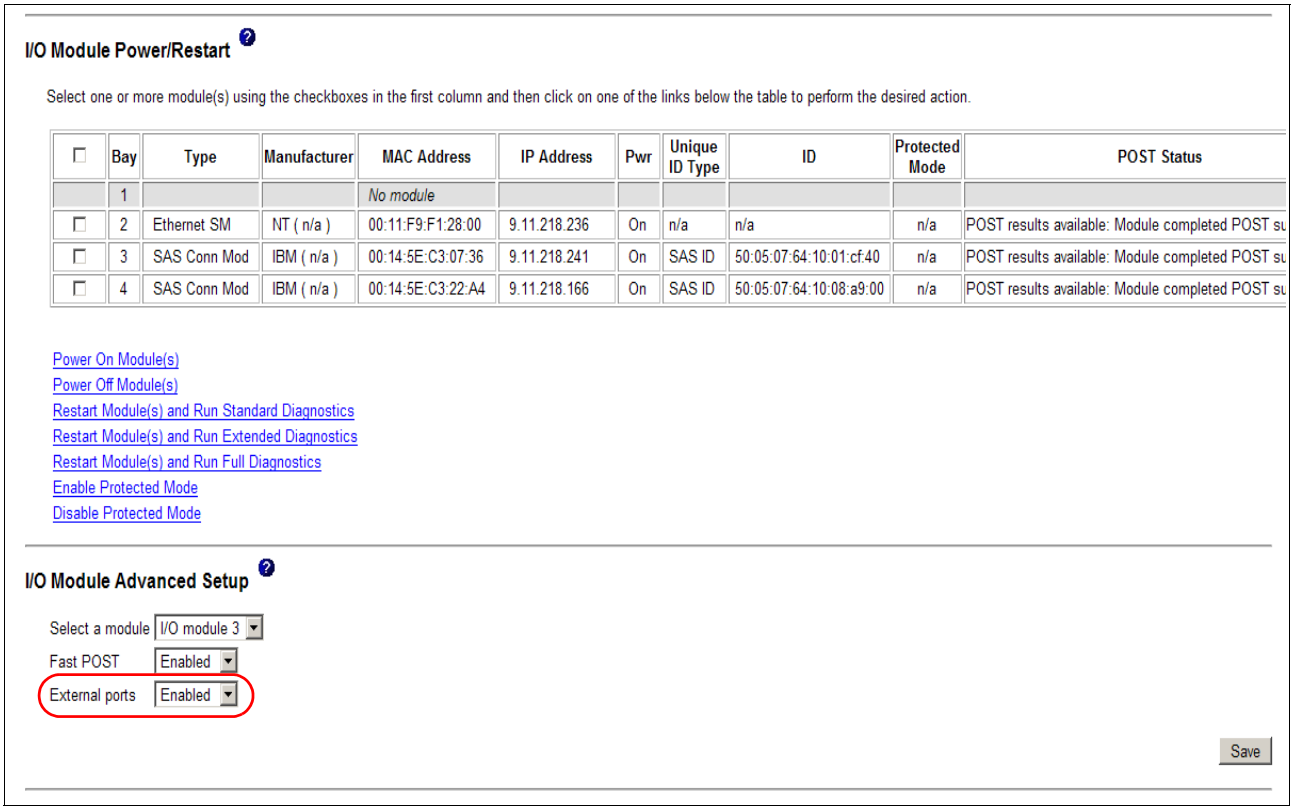


Figure 18-11 SAS Connectivity Module - Configuration

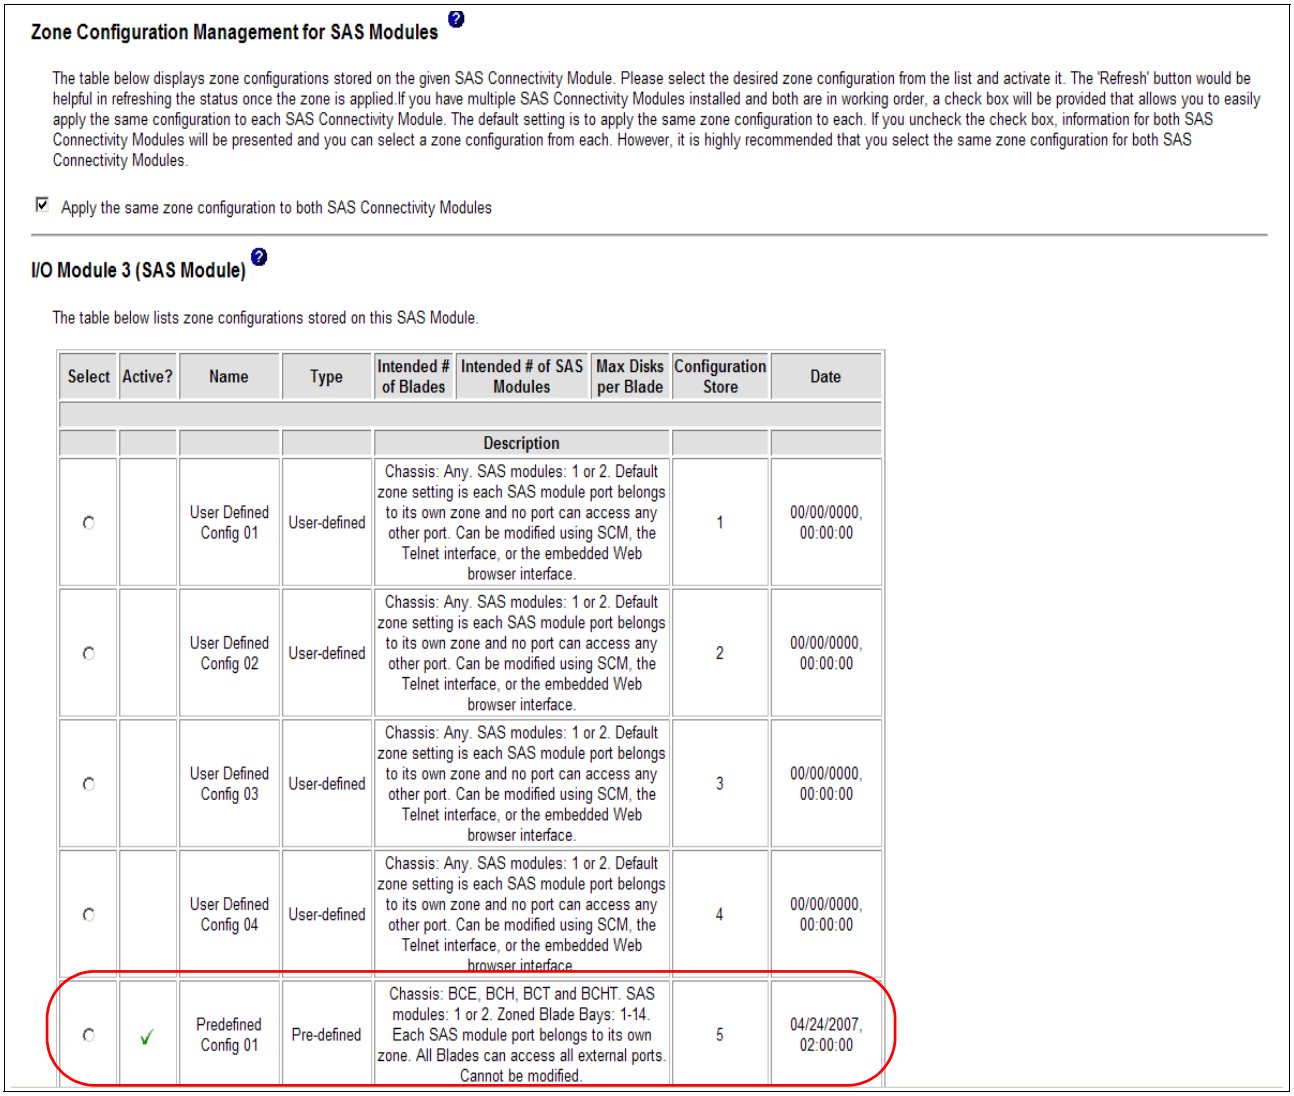
18.2.8 SAS Connectivity Module zoning

Zoning segregates devices at the fabric level by creating smaller virtual domains within the fabric. Zoning prohibits access between devices within the same logical fabric. For a host (initiator) to gain access to the storage subsystem (target), the initiator HBA WWPN or the switch port to which it is connected must be zoned with the corresponding target WWPN or the switch port and this zone should be a member of the active zoneset. Thus, although zoning is a tool to permit or deny access to the devices, it does not have the intelligence to apply controls beyond the fabric, that is, to present or hide the LUN to certain hosts.

Several predefined SAS fabric zone configurations are available from the factory and can be invoked by a simple selection from the Advanced Management Module (AMM). Zoning on the SAS Connectivity Module can be performed using the AMM I/O module configuration option, Web GUI, SAS, telnet, and SAS connectivity module (SCM) application.

Select **I/O Module Tasks** → **Configuration** in the BladeCenter AMM GUI Web interface window. For I/O Module 3 or I/O Module 4, select **Zone Configuration Management**. A window opens showing the predefined zone configuration options, as shown in Figure 18-12 on page 423.

You can select from five predefined zone configuration options. In this example, option 5 predefined zone configuration is currently active (indicated by a check mark), as shown in Figure 18-12. With this option, each server bay is exclusively zoned with all the external ports, thus allowing access to one or more storage controller ports connected to the SAS Connectivity Module.



Zone Configuration Management for SAS Modules

The table below displays zone configurations stored on the given SAS Connectivity Module. Please select the desired zone configuration from the list and activate it. The 'Refresh' button would be helpful in refreshing the status once the zone is applied.If you have multiple SAS Connectivity Modules installed and both are in working order, a check box will be provided that allows you to easily apply the same configuration to each SAS Connectivity Module. The default setting is to apply the same zone configuration to each. If you uncheck the check box, information for both SAS Connectivity Modules will be presented and you can select a zone configuration from each. However, it is highly recommended that you select the same zone configuration for both SAS Connectivity Modules.

☑ Apply the same zone configuration to both SAS Connectivity Modules

I/O Module 3 (SAS Module)

The table below lists zone configurations stored on this SAS Module.

Select	Active?	Name	Type	Intended # of Blades	Intended # of SAS Modules	Max Disks per Blade	Configuration Store	Date
				Description				
○		User Defined Config 01	User-defined	Chassis: Any. SAS modules: 1 or 2. Default zone setting is each SAS module port belongs to its own zone and no port can access any other port. Can be modified using SCM, the Telnet interface, or the embedded Web browser interface.			1	00/00/0000, 00:00:00
○		User Defined Config 02	User-defined	Chassis: Any. SAS modules: 1 or 2. Default zone setting is each SAS module port belongs to its own zone and no port can access any other port. Can be modified using SCM, the Telnet interface, or the embedded Web browser interface.			2	00/00/0000, 00:00:00
○		User Defined Config 03	User-defined	Chassis: Any. SAS modules: 1 or 2. Default zone setting is each SAS module port belongs to its own zone and no port can access any other port. Can be modified using SCM, the Telnet interface, or the embedded Web browser interface.			3	00/00/0000, 00:00:00
○		User Defined Config 04	User-defined	Chassis: Any. SAS modules: 1 or 2. Default zone setting is each SAS module port belongs to its own zone and no port can access any other port. Can be modified using SCM, the Telnet interface, or the embedded Web browser interface.			4	00/00/0000, 00:00:00
○	✓	Predefined Config 01	Pre-defined	Chassis: BCE, BCH, BCT and BCHT. SAS modules: 1 or 2. Zoned Blade Bays: 1-14. Each SAS module port belongs to its own zone. All Blades can access all external ports. Cannot be modified.			5	04/24/2007, 02:00:00

Figure 18-12 SAS Connectivity Module - Predefined Zones

The following example in Figure 18-13 was captured from the SAS Connectivity Module Web interface. It lists the Basic Zone Permission table for the HS21 blade in slot 6, because the blade in slot 6 is used for this example. The blade in slot 6 is zoned with four external ports. Notice that the External port is set to True under the Connected column and Normal under the Status column. This is because the DS3200 Storage System Controller A port is connected to the external port #1 for this example.

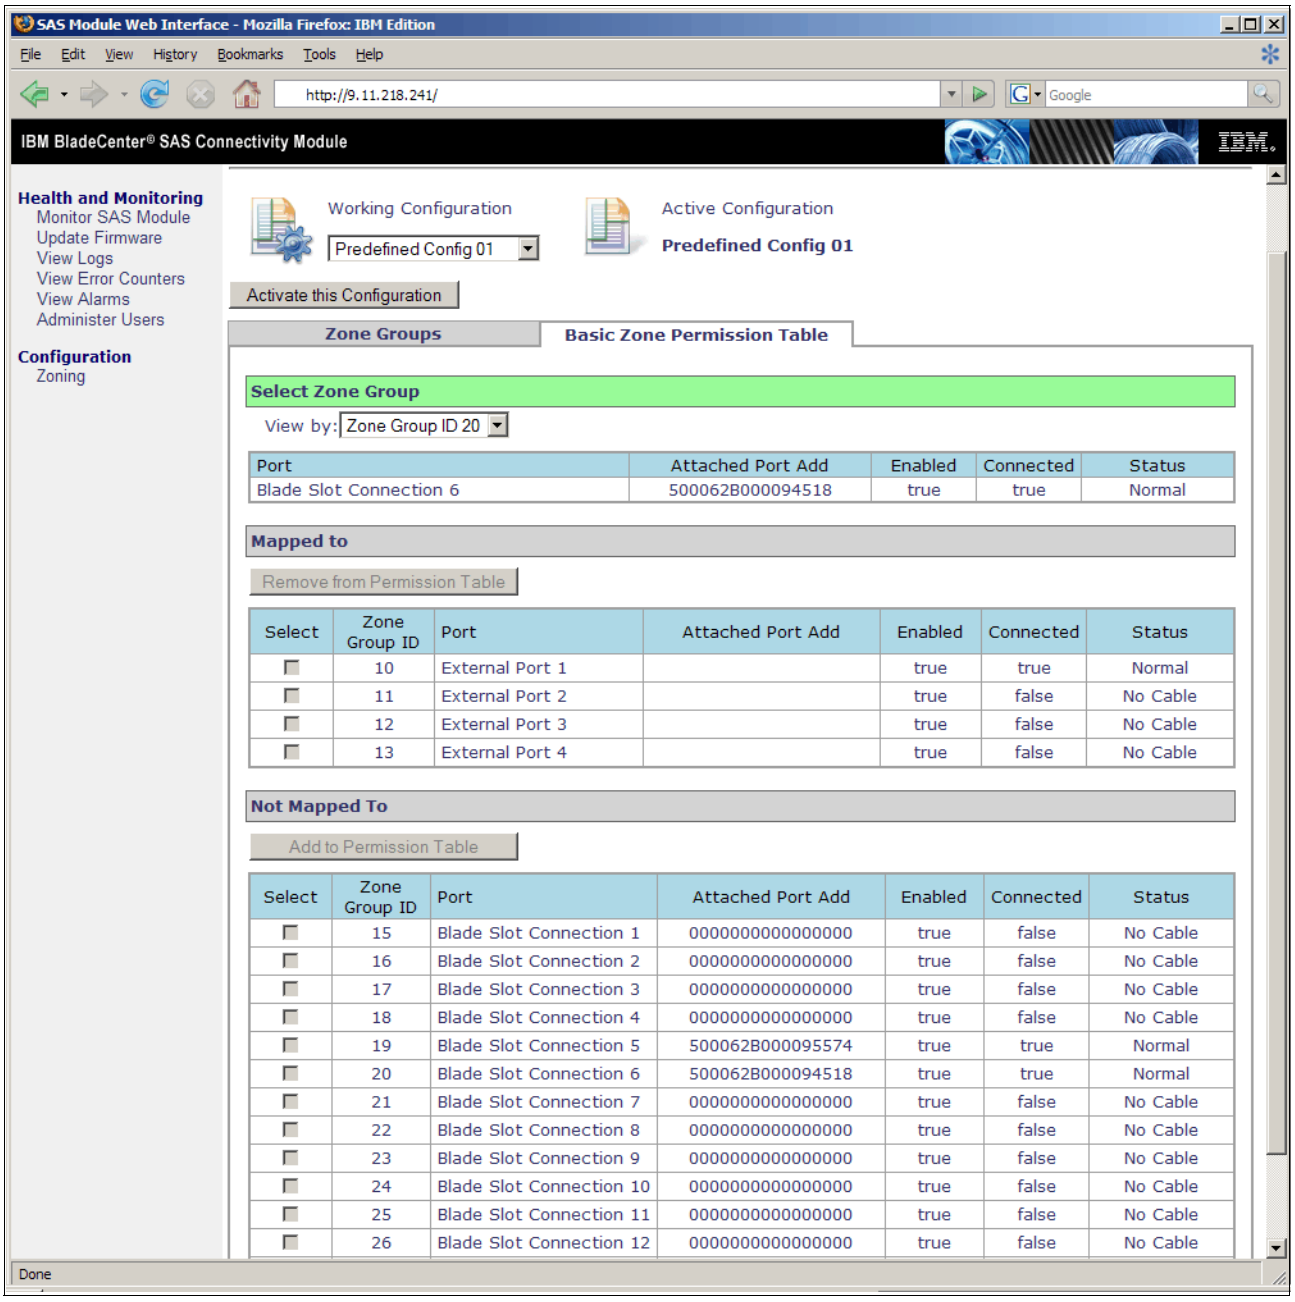


Figure 18-13 SAS Connectivity Module - Zone configuration

Note: There is only a single path active from the DS3200 Storage System controller A to the SAS Connectivity Module in BlaseCenter I/O slot bay 3. The second path from DS3200 Storage System controller B is to the SAS Connectivity Module in BladeCenter I/O slot bay 4.

18.3 Installing DS3000 Storage Manager host software

We will install the host server components of DS3000 Storage Manager. These components include:

- SMagent
- SMutil
- Multipath support
- SMclient (optional)

At the time of writing, in-band management is not supported on the DS3200, so we will not install the SMagent. We will be using out-of-band management.

We will install SMutil, because this will allow us to run the *hot_add* utility and hence avoid the necessity of rebooting the host server when adding new logical drives.

Since the host server contains a 2-port SAS Expansion Card, we require multipath support. This is installed as part of the Host selection when installing the DS3000 Storage Manager.

We will also install SMclient, although this component is usually not installed in the host server, but rather on a management workstation. In our case, we want to run the DS3000 Storage Manager GUI in the host server itself.

For the exact installation steps, see Chapter 6, “Installing IBM System Storage DS3000 Storage Manager” on page 81.

Once the DS3000 Storage Manager client is installed, we have to add the DS3200 subsystem in the Enterprise Management window and prepare the logical drives for our Windows server. For simplicity’s sake, we create just one logical drive of 10 GB.

We also have to define our Windows Server 2008 host in the Storage Manager and map it to the logical drive.

We have already explained how to add storage subsystems, define hosts, prepare logical drives, and map them in Part 3, “Administration” on page 109, so we do not repeat these steps here. Figure 18-14 shows our 10 GB logical drive named SAS_Windows_Vol1 and mapped to the host server Blade06.

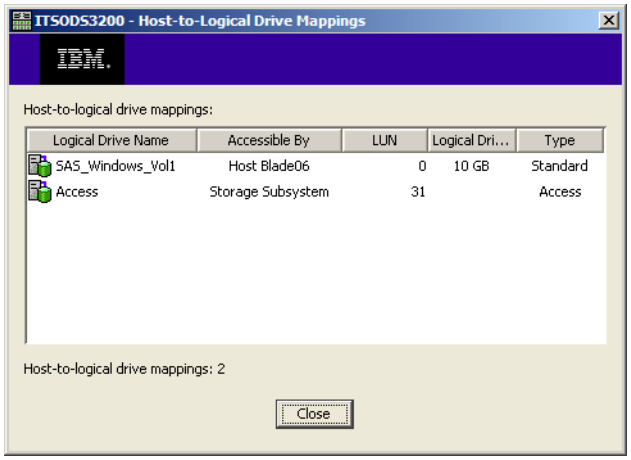


Figure 18-14 Drive mapping

18.4 Configure the disk space in Windows Server 2008

The logical drive now must be recognized in Windows, before we can start using the disk space. One possibility is to shut down and reboot the host server. But there is a better alternative: We will use the hot_add utility (a part of SMutil).

By default, the hot_add executable is located in C:\Program Files\IBM_DS3000\util. In a command prompt window, change to this directory. If you installed the DS3000 Storage Manager in a non-default directory, change to that directory instead, and then to subdirectory util.

Now run **`hot_add`**.

When the utility completes, you will be able to use the Disk Management applet in the Computer Management utility and configure the disk space for operating system use. This includes the following steps:

1. Initialize and optionally convert to dynamic disk.

   This task is done by the Initialize and Convert Disk wizard. When the wizard finishes, the Disk Management applet looks similar to Figure 18-15.

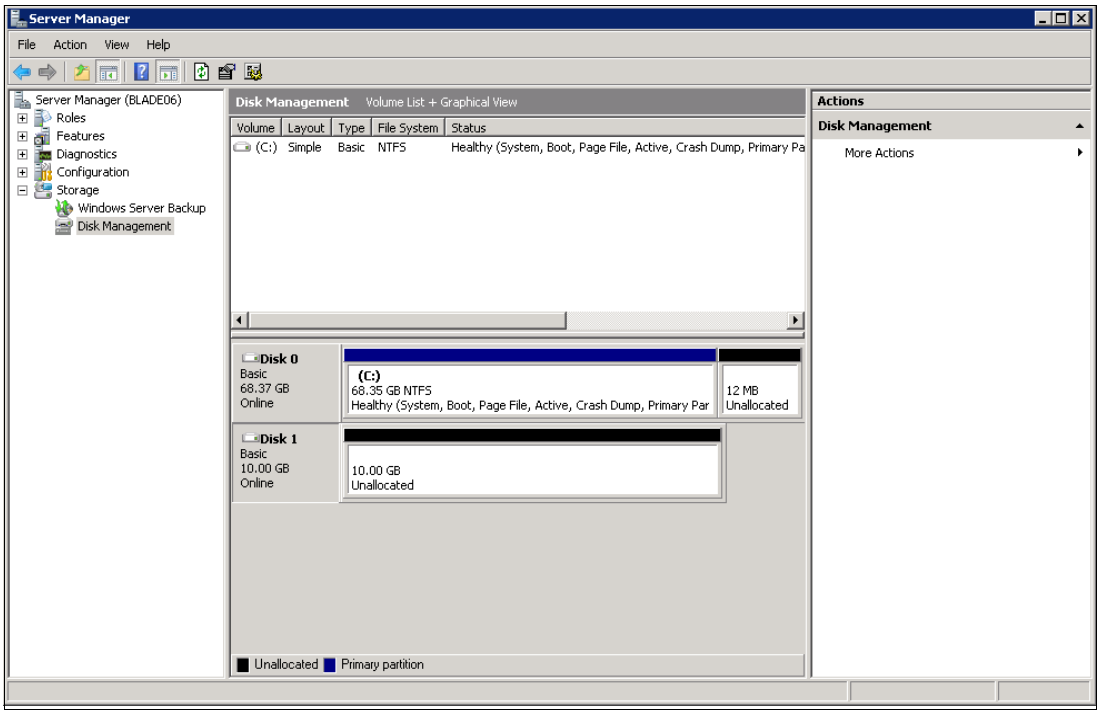


Figure 18-15 Disk Management applet - initialized disk

   As you can see, our 10 GB logical drive is recognized as Disk 1, but it does not contain any partition yet; it is unallocated.

2. Create and format a partition.

   Right-click in the unallocated disk space and select **New Simple Volume...** from the context menu (Figure 18-16 on page 427).

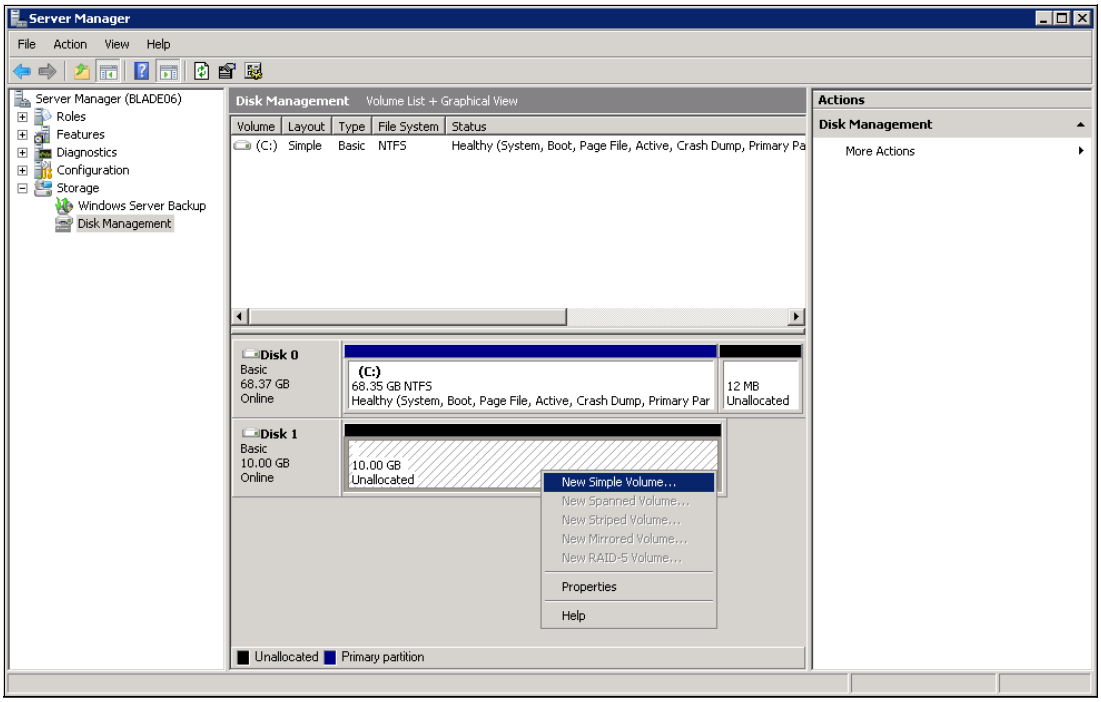


Figure 18-16 Creating a new partition

This launches the New Simple Volume Wizard.

3. Follow the process under guidance of the wizard to define the partition size, drive letter, file system, and volume label. The partition will be created and formatted. The Disk Management applet now displays the new partition (allocated drive D), which is ready to be used (see Figure 18-17).

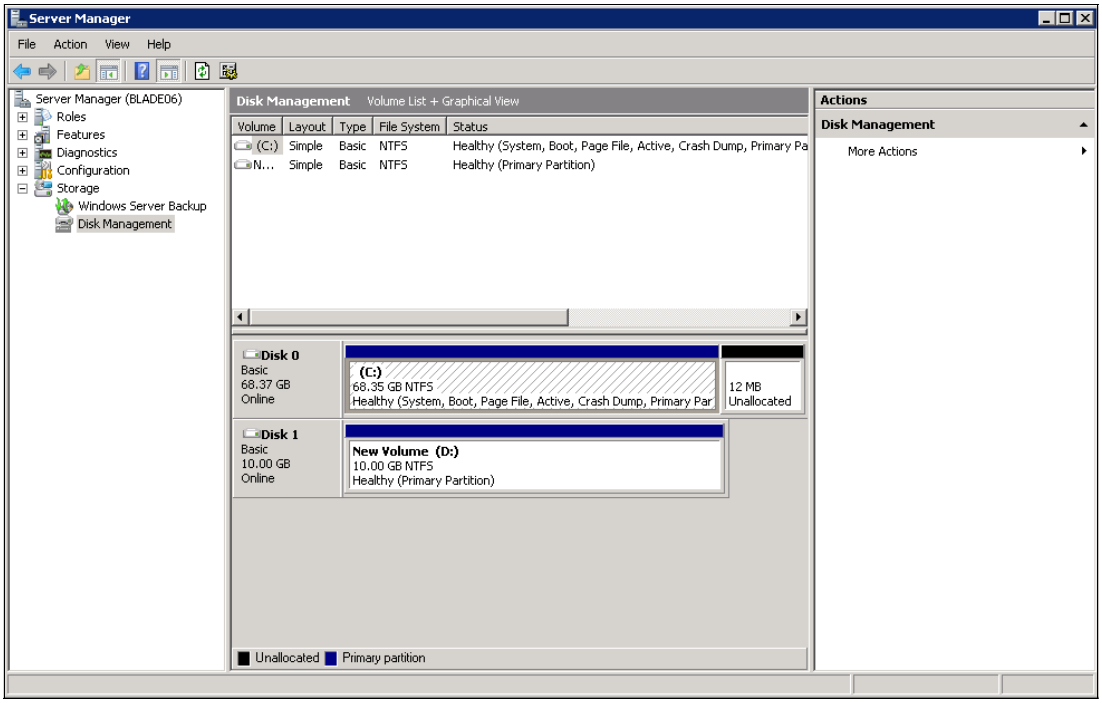


Figure 18-17 New partition on the logical drive

4. We can now access the logical drive on the DS3200 as drive letter D:. We show this in Figure 18-18.

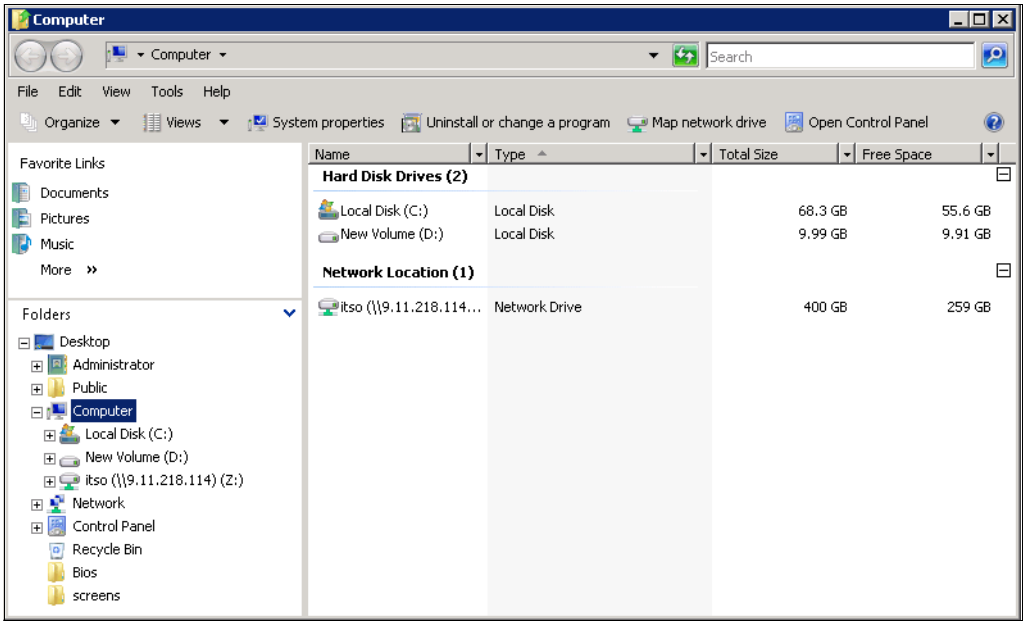


Figure 18-18 Logical drive is ready for use

19

SAS configuration 4 - RHEL V5.2 Linux on an IBM BladeCenter HS21 server

In this chapter, we describe and explain a sample configuration that shows how to connect logical drives configured on an IBM System Storage DS3200 to a Red Hat Enterprise Linux 5 Update 2 operating system running on an IBM BladeCenter HS21 server connected with SAS connectivity modules.

Note: At the time of writing, RDAC Version 09.03.0B05.0023 was the used to illustrate the steps below. Refer to the IBM Systems support Web site, as described in 17.4, “Installing RDAC for Linux” on page 408, and download the latest available version of RDAC.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

19.1 Equipment required

We use the following hardware and software components:

- IBM BladeCenter E and HS21 BladeCenter installed with RHEL 5 Update 2.
- IBM BladeCenter Advance Management Module (AMM) installed in BladeCenter.
- Two IBM BladeCenter SAS Connectivity Modules.
- IBM BladeCenter SAS Expansion card (CFFv), P/N 43W3974.
- The latest SAS Expansion Card device driver for Linux kernel 2.6. The driver is available for download on the IBM Systems support Web site.
- DS3200 System Storage subsystem.
- The latest version of DS3000 Storage Manager for Linux (at the time of writing, this is Version 10.35).
- The latest version of RDAC for Linux (V09.03.0B05.0023).

We show our sample hardware setup in Figure 19-1.

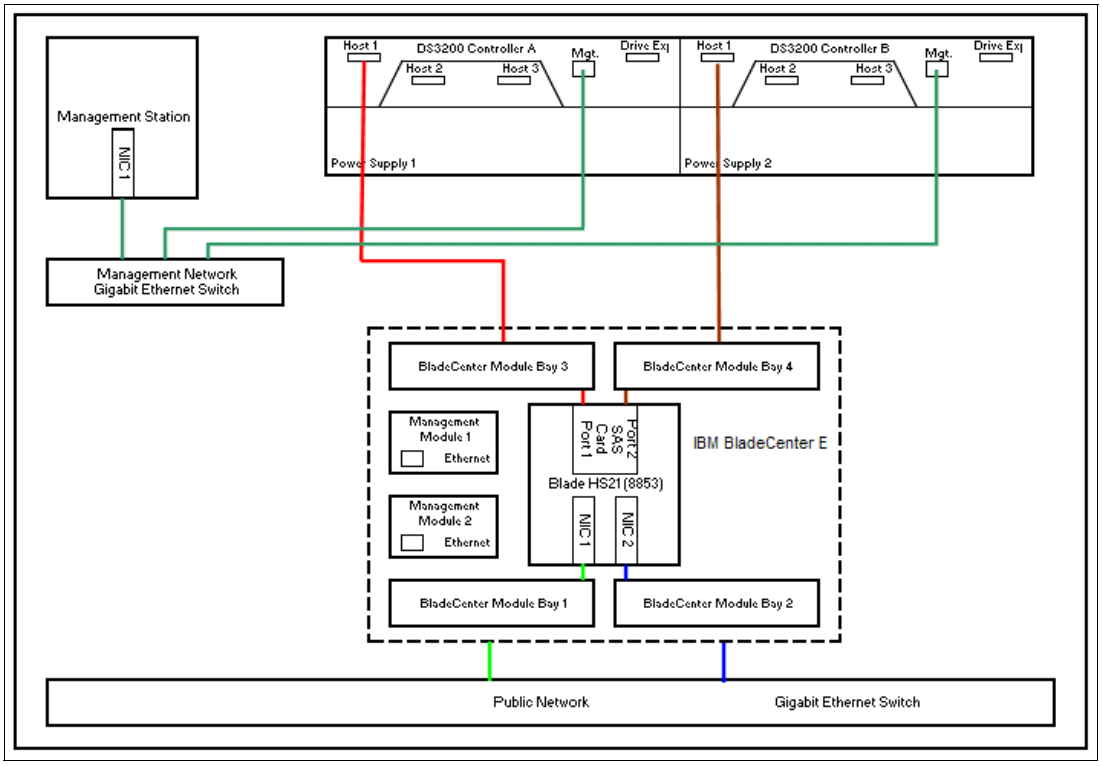


Figure 19-1 Hardware setup

19.2 IBM BladeCenter setup

Here we discuss the IBM BladeCenter setup.

19.2.1 Installing Red Hat Enterprise Linux 5 Update 2

Follow the operating system installation instructions that are available for each IBM BladeCenter blade and IBM System x server. The installation guides can be found in the “Install/use” section of each product’s support Web sites.

For our IBM BladeCenter HS21 (8853) with Red Hat Enterprise Linux 5, we used the instructions found at:

https://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-5070760&brandind=5000020

The installation was performed using Red Hat Enterprise Linux 5 with Update 2, although the instructions on the Web site are for Red Hat Enterprise Linux 5.

19.2.2 HS21 SAS Expansion Cards

We have to install the IBM BladeCenter SAS Expansion Card in the HS21 BladeCenter host before we can proceed with other tasks. While this is not a difficult task, we do recommend that you consult the user’s guide for the host server and follow the instructions for options installation. The next step is the SAS Expansion Card driver installation.

Important: The connectivity modules in I/O module bay 3 and I/O module bay 4 and all expansion cards in the BladeCenter unit must use the same interface type. Therefore, you must install SAS expansion cards before you install connectivity modules in the blade servers in your BladeCenter unit. For more information about the SAS expansion card, see the *Installation and User's Guide* for the SAS expansion card at http://www.ibm.com/systems/support/.

Note: The BladeCenter SAS Expansion Card is a dual port card. Port #1 connects to a SAS Connectivity Module in BladeCenter I/O module bay 3 and port #2 connects to a SAS Connectivity Module in BladeCenter I/O module bay 4 respectively.

19.2.3 Recording the SAS Expansion Card WWPN

The following example demonstrates how to record the SAS Expansion Card WWPN for later use in setting up the host-to-LUN mappings in the DS3200 System Storage Subsystem:

1. Turn on or restart the HS21 BladeCenter host.

2. Press <Ctrl + C> to enter the LSI Logic Configuration Utility, as shown in Figure 19-2.

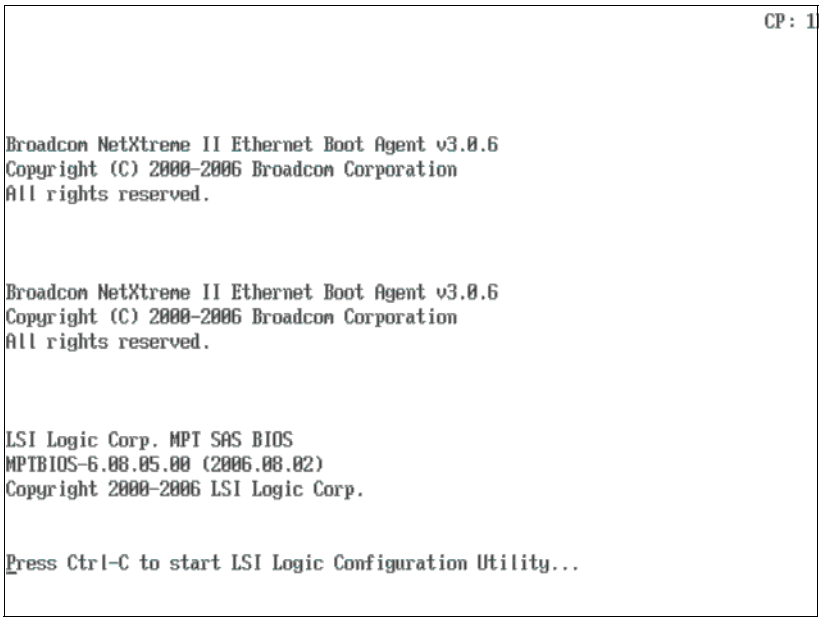


Figure 19-2 SAS Expansion Card - LSI Logic Config Utility

3. The menu shown in Figure 19-3 is displayed after pressing Crtl + C.

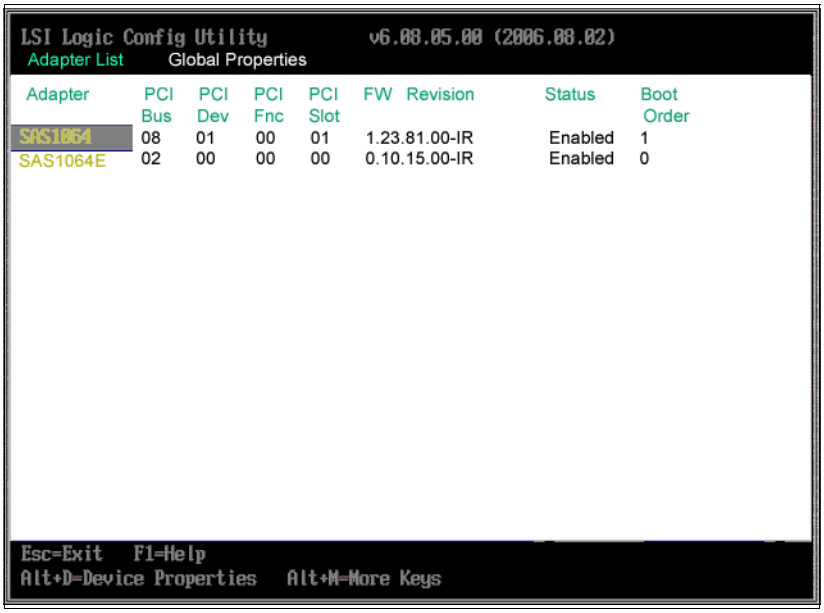


Figure 19-3 SAS Expansion Card - LSI Logic Config Utility menu

4. Press the Enter key to select the SAS adapter internally connected to the SAS connectivity module in I/O Module bay 3, as shown in Figure 19-4 on page 433.

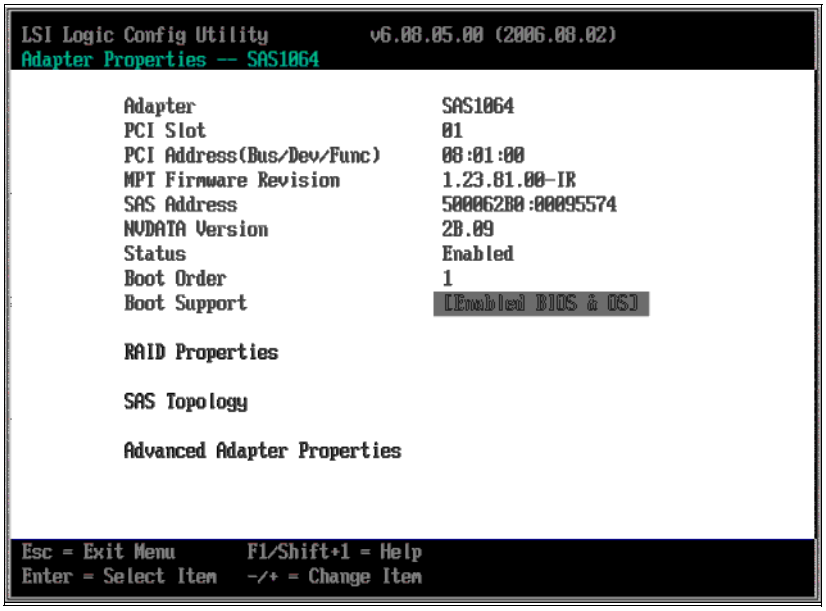


Figure 19-4 SAS Expansion Card - LSI Logic Config Utility adapter selected

Record the Worldwide Port Name (WWPN) of the first port on the SAS Expansion Card. The WWPN is needed for defining host ports on the DS3200 System Storage subsystem. The WWPN can also be retrieved from the SAS Connectivity Module Web interface.

The name of the SAS adapter for the expansion card is SAS1064 and is visible in the Adapter List screen. To determine if the SAS adapter is the expansion card, select a SAS adapter and use the View SAS Topology screen to display whether the SAS adapter is connected to the internal hard disk drives or to the SAS connectivity modules in your BladeCenter chassis, as shown in Figure 19-5.

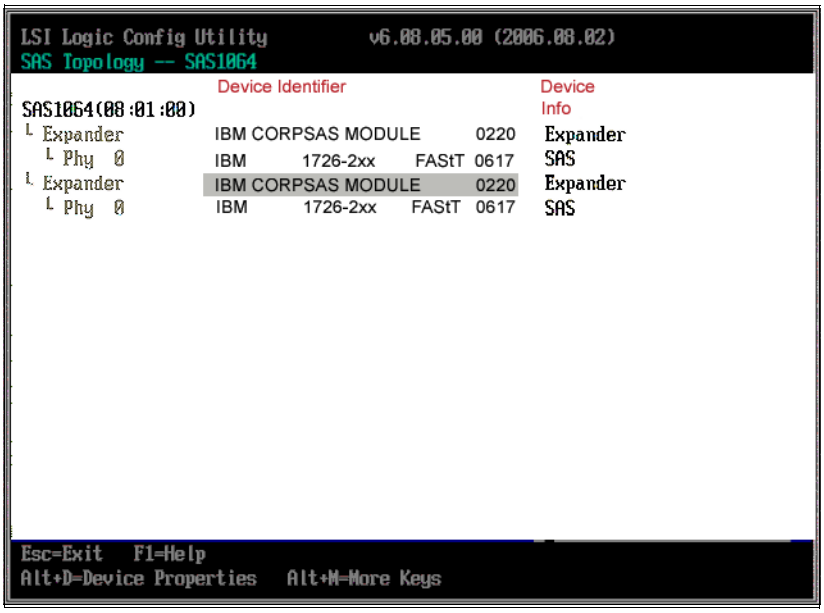


Figure 19-5 SAS Expansion Card - LSI Logic Config Utility adapter confirmation

19.2.4 HS21 SAS Expansion Card driver installation

The latest Linux SAS Expansion Card driver is available for download on the IBM Systems support Web site. A set of precompiled binary RPMs for different kernel flavors (default, SMP, and hugemem) is included, as well as the source RPM for use with kernel versions that do not match the precompiled driver versions.

In our case, we can simply install a precompiled driver. This can be done in two ways:

- Run the install.sh script, included with the driver:

  ```
  ./install.sh --update
  ```

- Alternatively, run the **rpm** command to install the driver package:

  ```
  rpm -ivh kmod-lsi-mptlinux-4.00.21.00-1.x86_64.rpm
  ```

There are several ways to verify that the driver package is correctly installed. One way is with the command:

```
rpm -qa | grep lsi
```

The output should show the driver package information, similar to this:

```
kmod-lsi-mptlinux-4.00.21.00-1
```

You could also run the **lsmod** command. The output should show that the following modules are running: mptsas, mptscsi, mptscsih, and mptbase.

19.2.5 SAS Connectivity Modules

You must install the IBM BladeCenter SAS Connectivity modules only in BladeCenter I/O module bay 3 and I/O module bay 4 of the following supported BladeCenter units:

- BladeCenter Type 8677
- BladeCenter Types 8720 and 8730
- BladeCenter Type 8750
- BladeCenter Type 8852
- BladeCenter Type 8886

Installing a connectivity module in I/O module bay 3 or I/O module bay 4 provides connectivity to the SAS expansion card(s) installed in the blade servers in your BladeCenter unit. Installing two connectivity modules allows you to have two connections to the SAS expansion cards installed in the blade servers.

Important: The connectivity modules in I/O module bay 3 and I/O module bay 4 and all expansion cards in the BladeCenter unit must use the same interface type. Therefore, you must install SAS expansion cards before you install connectivity modules in the blade servers in your BladeCenter unit. For more information about the SAS expansion card, see the *Installation and User's Guide* for the SAS expansion card at http://www.ibm.com/systems/support/.

Connect the cables from the DS3200 Storage System controllers A and B to the external port# 1 of the two SAS connectivity module modules, as shown in Figure 19-6 on page 435.

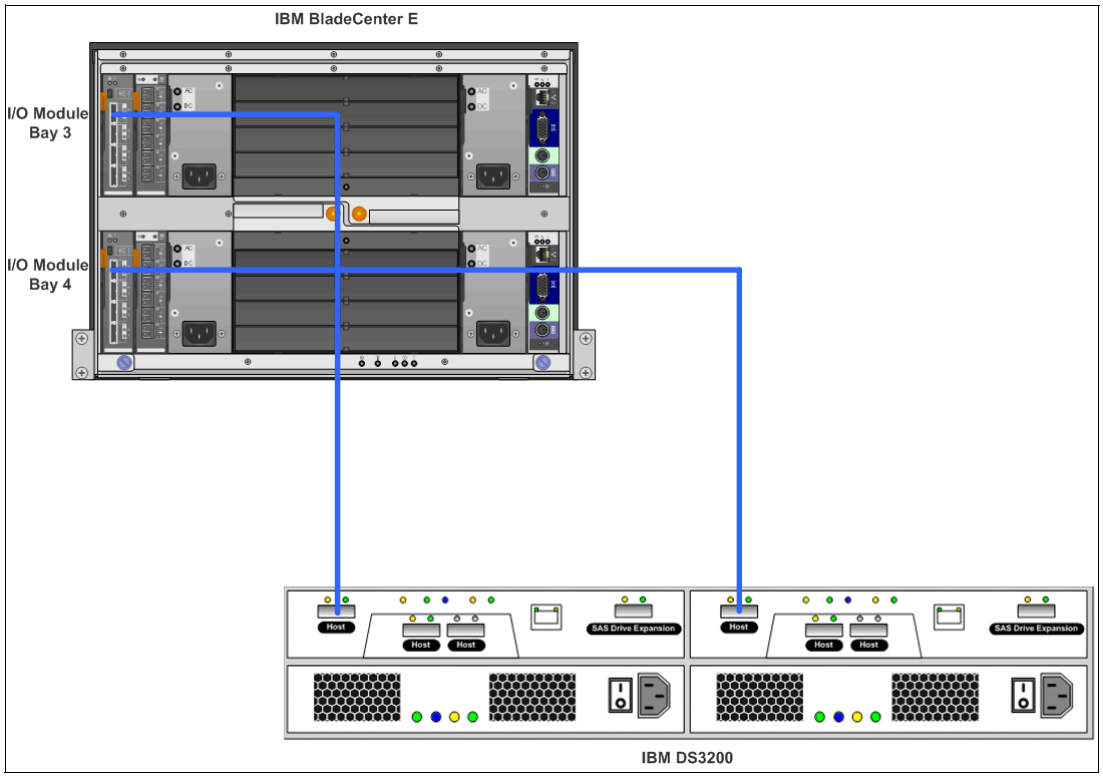


Figure 19-6 BladeCenter and DS3200 connectivity

19.2.6 SAS Connectivity Module firmware update

Ensure that your SAS Connectivity Module is updated with the latest firmware. For the latest firmware update, refer to http://www.ibm.com/systems/support/.

To update the connectivity module firmware to the latest version, complete the following steps:

1. Log on to the SAS Connectivity Module using the Web interface with the IP address defined for the connectivity module in the BladeCenter Advance Management Module (AMM), as shown in Figure 19-7. Use USERID as the user ID and PASSW0RD as the password. You can change the password under the Administer Users menu option once you are logged on.

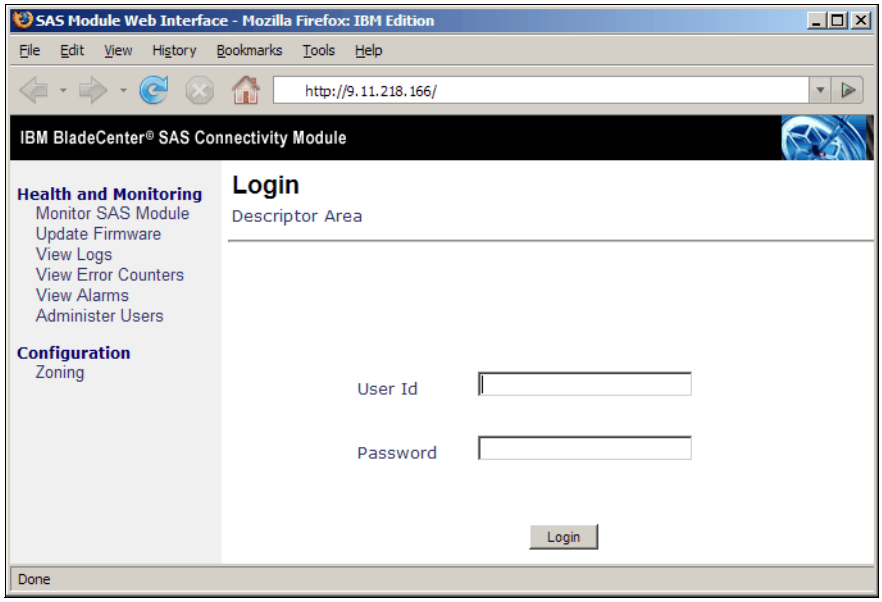


Figure 19-7 SAS Connectivity Module - Login

2. In the Monitor Module window, click **Update Firmware**. The Update Firmware window opens, as shown in Figure 19-8. The current firmware version is also displayed.

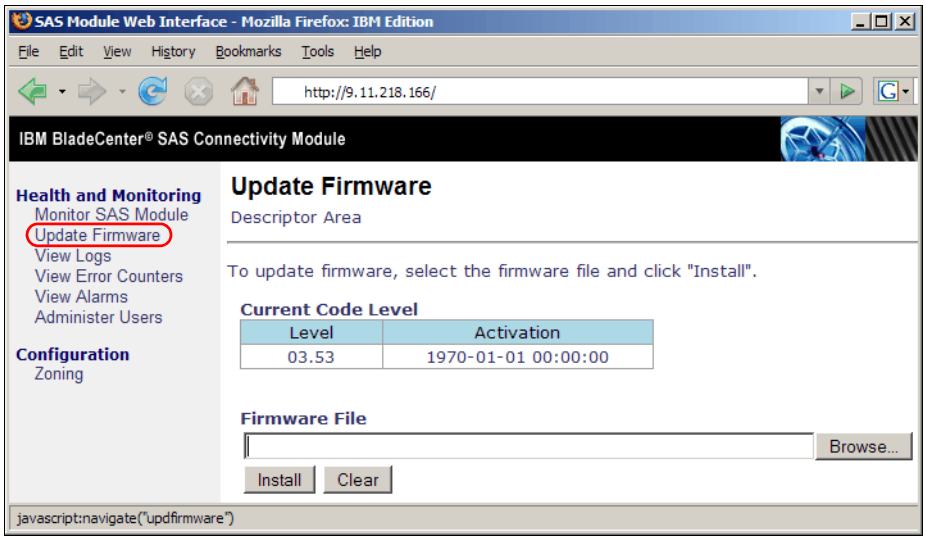


Figure 19-8 SAS Connectivity Module - Update Firmware

3. In the Firmware File field, enter the new firmware file name, or click **Browse** and locate the firmware file.

4. Click **Install** to install the new file. An installation confirmation window opens, as shown in Figure 19-9.

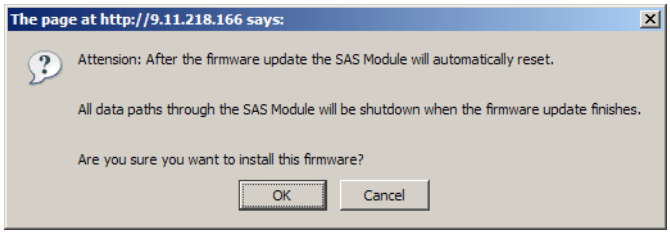


Figure 19-9 SAS Connectivity Module - Installation confirmation

5. Click **OK** or **Cancel**. If the installation of the new firmware file is successful, an installation confirmation window opens, as shown in Figure 19-10. If there are errors during the installation, an error message window opens.

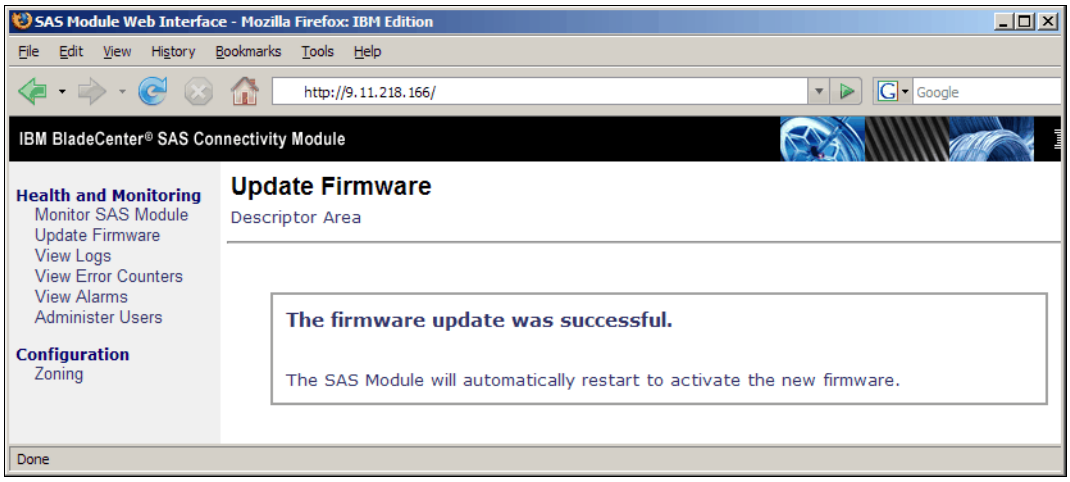


Figure 19-10 SAS Connectivity Module - Update successful

19.2.7 Configuring the SAS connectivity module

Ensure that the external ports on the SAS Connectivity Modules are enabled to allow connectivity from the DS3200 System Storage subsystem and I/O traffic to pass through the module.

From the BladeCenter Advance Management Module Web interface GUI, select **I/O Module Tasks** → **Admin/Power/Restart** to ensure that the external ports for the I/O modules in bays 3 and 4 are enabled, as shown in Figure 19-11.

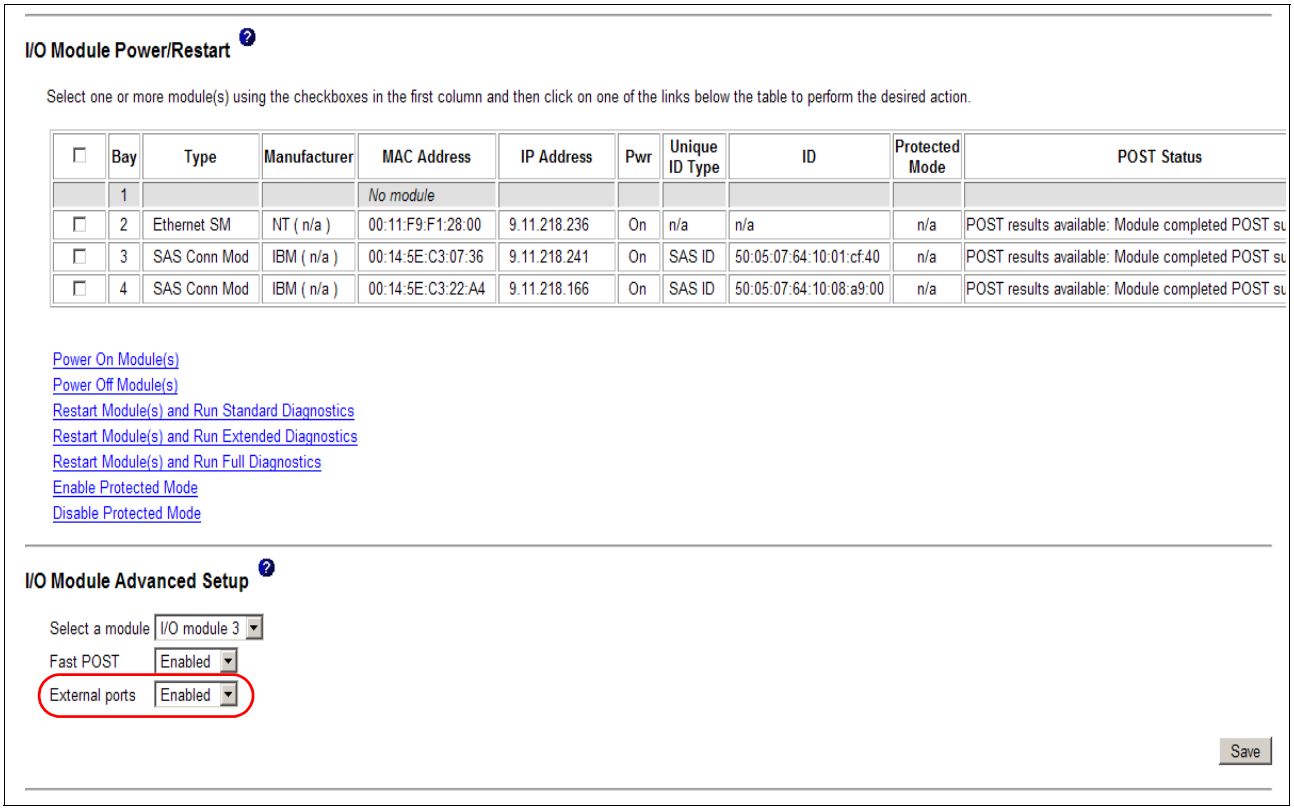


Figure 19-11 SAS Connectivity Module - Configuration

19.2.8 SAS Connectivity Module zoning

Zoning segregates devices at the fabric level by creating smaller virtual domains within the fabric. Zoning prohibits access between devices within the same logical fabric. For a host (initiator) to gain access to the storage subsystem (target), the initiator HBA WWPN or the switch port to which it is connected must be zoned with the corresponding target WWPN, or the switch port and this zone should be a member of the active zoneset. Thus, although zoning is a tool to permit or deny access to the devices, it does not have the intelligence to apply controls beyond the fabric, that is, to present or hide the LUN to certain hosts.

Several predefined SAS fabric zone configurations are available from the factory and can be invoked by a simple selection from the Advanced Management Module (AMM). Zoning on the SAS connectivity module can be performed using the AMM I/O module configuration option, Web GUI, SAS, telnet, and SAS Connectivity Module (SCM) application.

Select **I/O Module Tasks** → **Configuration** in the BladeCenter AMM GUI Web interface window. For I/O module 3 or I/O module 4, select **Zone Configuration Management**. A window opens showing the predefined zone configuration options, as shown in Figure 19-12 on page 439.

You can select from five predefined zone configuration options. In this example, option # 5 predefined zone configuration is currently active (indicated by the check mark), as shown in Figure 19-12 on page 439. With this option, each server bay is exclusively zoned with all the

external ports, thus allowing access to one or more storage controller ports connected to the SAS Connectivity Module.

Zone Configuration Management for SAS Modules

The table below displays zone configurations stored on the given SAS Connectivity Module. Please select the desired zone configuration from the list and activate it. The 'Refresh' button would be helpful in refreshing the status once the zone is applied.If you have multiple SAS Connectivity Modules installed and both are in working order, a check box will be provided that allows you to easily apply the same configuration to each SAS Connectivity Module. The default setting is to apply the same zone configuration to each. If you uncheck the check box, information for both SAS Connectivity Modules will be presented and you can select a zone configuration from each. However, it is highly recommended that you select the same zone configuration for both SAS Connectivity Modules.

☑ Apply the same zone configuration to both SAS Connectivity Modules

I/O Module 3 (SAS Module)

The table below lists zone configurations stored on this SAS Module.

Select	Active?	Name	Type	Intended # of Blades	Intended # of SAS Modules	Max Disks per Blade	Configuration Store	Date
				Description				
○		User Defined Config 01	User-defined	Chassis: Any. SAS modules: 1 or 2. Default zone setting is each SAS module port belongs to its own zone and no port can access any other port. Can be modified using SCM, the Telnet interface, or the embedded Web browser interface.			1	00/00/0000, 00:00:00
○		User Defined Config 02	User-defined	Chassis: Any. SAS modules: 1 or 2. Default zone setting is each SAS module port belongs to its own zone and no port can access any other port. Can be modified using SCM, the Telnet interface, or the embedded Web browser interface.			2	00/00/0000, 00:00:00
○		User Defined Config 03	User-defined	Chassis: Any. SAS modules: 1 or 2. Default zone setting is each SAS module port belongs to its own zone and no port can access any other port. Can be modified using SCM, the Telnet interface, or the embedded Web browser interface.			3	00/00/0000, 00:00:00
○		User Defined Config 04	User-defined	Chassis: Any. SAS modules: 1 or 2. Default zone setting is each SAS module port belongs to its own zone and no port can access any other port. Can be modified using SCM, the Telnet interface, or the embedded Web browser interface.			4	00/00/0000, 00:00:00
○	✓	Predefined Config 01	Pre-defined	Chassis: BCE, BCH, BCT and BCHT. SAS modules: 1 or 2. Zoned Blade Bays: 1-14. Each SAS module port belongs to its own zone. All Blades can access all external ports. Cannot be modified.			5	04/24/2007, 02:00:00

Figure 19-12 SAS Connectivity Module - Predefined zones

The following example in Figure 19-13 was captured from the SAS Connectivity Module Web interface. It lists the Basic Zone Permission table for the HS21 blade in slot 5, because the blade in slot 5 is used for this example. The blade in slot 5 is zoned with four external ports. Notice that the External port is set to True under the Connected column and Normal under the Status column. This is because the DS3200 Storage System Controller A port is connected to external port # 1 for this example.

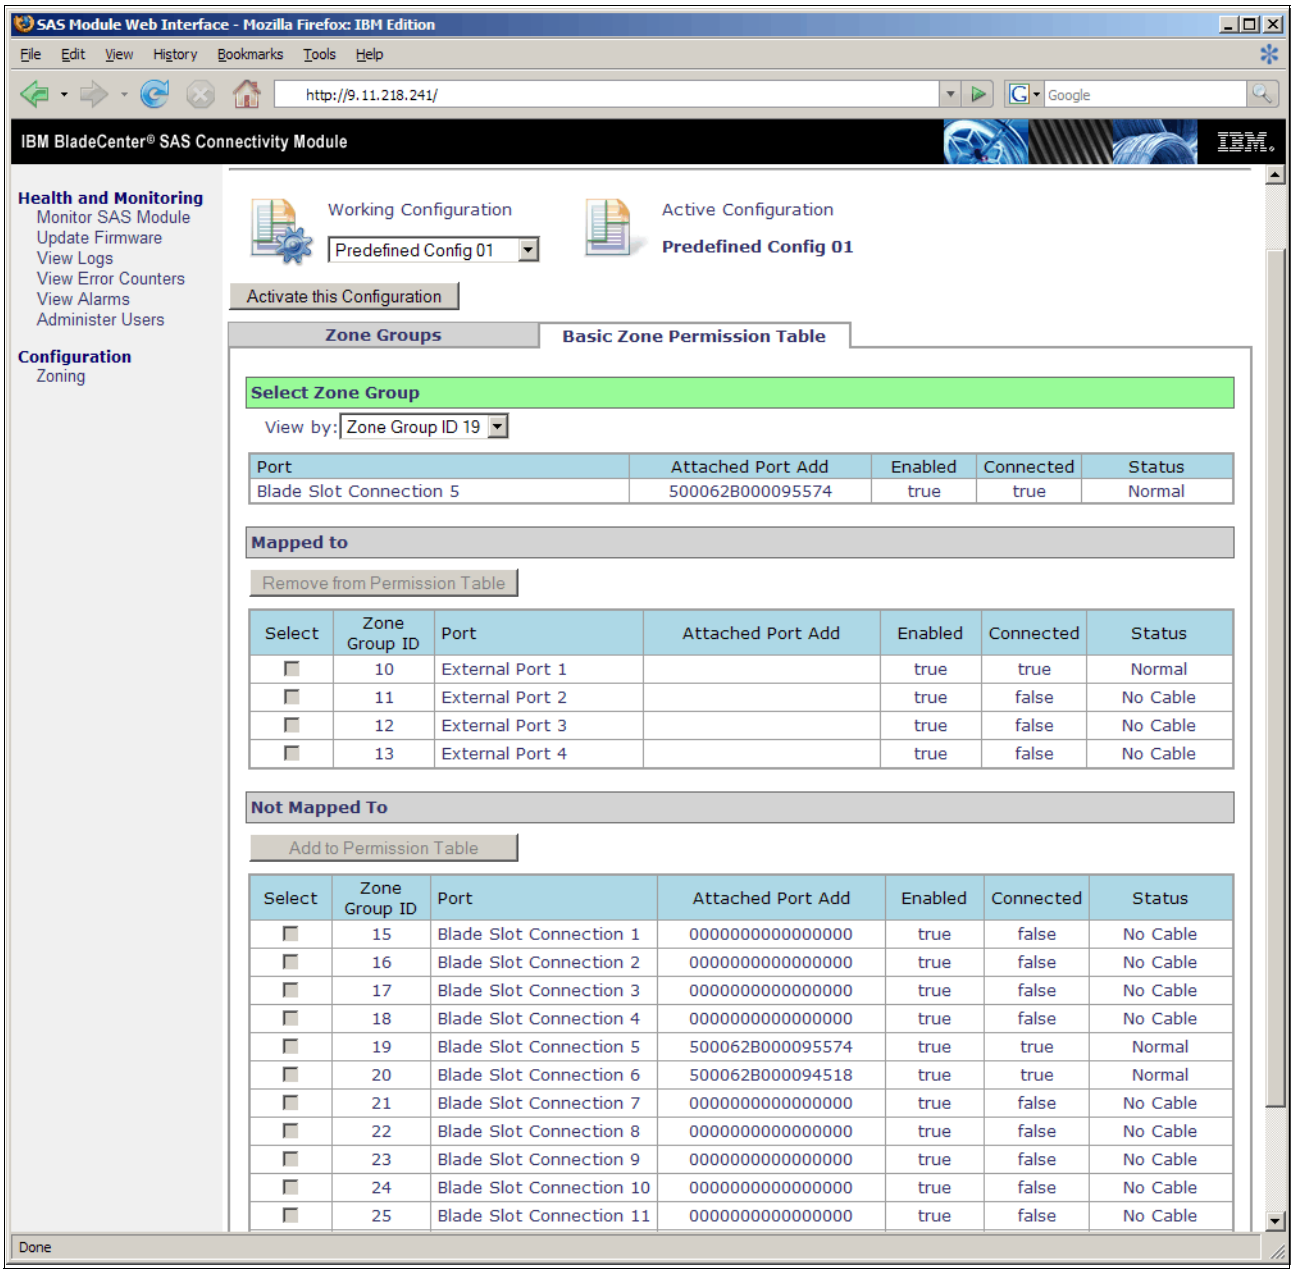


Figure 19-13 SAS Connectivity Module - Zone configuration

Note: There is only a single path active from the DS3200 Storage System controller A to the SAS Connectivity Module in BladeCenter I/O slot bay 3. The second path from DS3200 Storage System controller B is to the SAS Connectivity Module in BladeCenter I/O slot bay 4.

19.3 Installing DS3000 Storage Manager software

A host server running Windows can have the following DS3000 Storage Manager components installed:

- SMagent, if in-band management is required.
- SMutil.
- SMclient, if we want to use the DS3000 Storage Manager GUI in the host itself.
- Multipath support.

On Linux, this is slightly different. The multipath driver, RDAC, is not included with the DS3000 Storage Manager for Linux; it is available as a separate package instead. We will install this in 19.4, “Installing RDAC for Linux” on page 442.

We will install SMclient, because we want to run the DS3000 Storage Manager GUI on the Linux server itself. We will install SMutil as well.

We do not install SMagent, since at the time of writing, the DS3200 does not support in-band management. We will therefore not install the SMagent and our management will be performed out-of-band.

The versions reflected below might differ depending on the latest versions available for download from the IBM support Web site.

The latest DS3000 Storage Manager for Linux is available on the IBM Systems support Web site. At the time of writing, V10.35 is the current version for the Linux 2.6 kernel:

1. The Storage Manager package is available as a compressed tar file. Use the following command to uncompress it:

   ```
   tar -zxvf ibm_sw_ds3k_0217b505_linux2.6_anycpu.tar.gz
   ```

2. Now look for the Storage Manager InstallAnywhere (SMIA) installation script file. The file name will be something similar to this:

   ```
   SMIA-LINUX03.35.A5.11.bin
   ```

3. Run this script file to start the installation script. This will launch the GUI installation process:

   ```
   sh ./SMIA-LINUX03.35.A5.11.bin
   ```

4. We do not show the installation windows and steps here because the installation process is covered in detail in 6.2, “Installing Storage Manager on Linux” on page 90.

Once the DS3000 Storage Manager is installed, we have to add the DS3200 subsystem in the Enterprise Management Window and prepare the logical drives for our Linux server. For simplicity’s sake, we will create just one logical drive of 10 GB.

We also have to define our Linux host server in the Storage Manager and map it to the logical drive. However, no I/O should be issued to the logical drive yet. We will only start using it after the RDAC installation.

We have already explained how to add storage subsystems, define hosts, and prepare logical drives and map them in Part 3, “Administration” on page 109, so we do not repeat these steps here. Figure 19-14 shows our 10 GB logical drive named SAS_Linux_Vol1 and mapped to the host server Blade05.

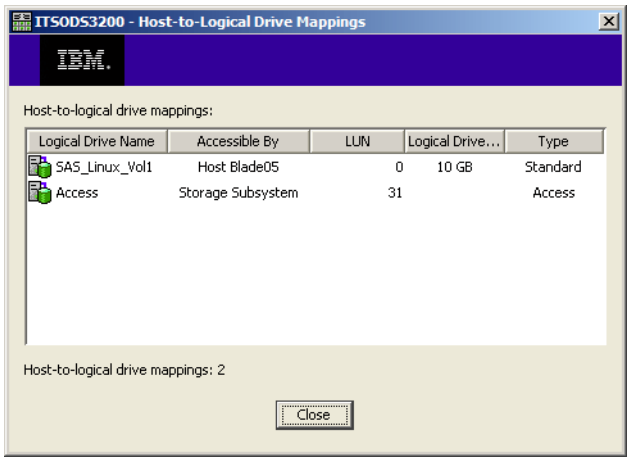


Figure 19-14 Drive mapping

19.4 Installing RDAC for Linux

RDAC for Linux is available for download on the IBM Systems support Web site. The IBM Web page leads you to the actual download Web page URL:

`http://www.lsi.com/rdac/ds3000.html`

Download the appropriate RDAC for your Linux kernel version (either 2.4 or 2.6) and continue with the next section. In our case, we use RHEL 5, so we downloaded the RDAC for kernel Version 2.6.

19.4.1 Building the RDAC driver

The RDAC package is available in a compressed tar file. Follow these steps to obtain it:

1. Unpack the file using the following command:

   ```
   tar -zxvf rdac-LINUX-09.03.0B05.0023-source.tar.gz dependant on the RDAC
   version of Exmoor
   ```

   The source files will be unpacked in the linuxrdac-09.03.0C05.0023 subdirectory of the current directory. When the uncompress is complete, change to this subdirectory.

2. Use the following command to remove old driver modules:

   ```
   make clean
   ```

3. Compile the driver modules and build the RDAC driver. Use this command:

   ```
   make
   ```

19.4.2 Installing the RDAC driver

We have to copy the driver modules to the kernel module tree and build the new RAMdisk image, which includes the RDAC driver modules. Enter the command:

```
make install
```

This will produce a lengthy output in the console window, as the driver modules are applied to the kernel tree. The last few lines will tell you how to add the new boot menu option to the boot loader, as shown in Example 19-1.

Example 19-1 RDAC installation

```
Creating new MPP initrd image...
You must now edit your boot loader configuration file, /boot/grub/menu.lst, to
add a new boot menu, which uses mpp-2.6.9-34.ELsmp.img as the initrd image.
Now Reboot the system for MPP to take effect.
The new boot menu entry should look something like this (note that it may
vary with different system configuration):
...
title Red Hat Linux (2.6.18-92.el5) with MPP support
root (hd0,5)
kernel /vmlinuz-2.6.18-92.el5 ro root=LABEL=RH9
initrd /mpp-2.6.18-92.el5.img
...
MPP driver package has been sucessfully installed on your system.
```

The new RAMdisk image filename is mpp-2.6.18-92.el5p.img. So we have to add the following code to the menu.lst file, as shown in Example 19-2.

Example 19-2 Modify menu.lst file

```
title Red Hat Linux (2.6.18-92.el5) with MPP support
        root (hd0,0)
         kernel /vmlinuz-2.6.18-92.el5 ro root=/dev/VolGroup00/LogVol00 rhgb quiet
        initrd /mpp-2.6.18-92.el5.img
```

We now restart the system using the new boot option.

19.4.3 Verification steps

When the server restarts, use the **lsmod** command to verify that the RDAC (MPP) modules are running. The following modules should be listed:

- mppUpper
- mppVhba
- sg
- sd_mod
- scsi_mod

If all these modules are running, the RDAC should now be managing access to the DS3200 logical drive. To verify this, use the following command:

```
ls -lR /proc/mpp
```

You should see output similar to Example 19-3.

Example 19-3 Output of ls -lR /proc/mpp

```
/proc/mpp:
total 0
dr-xr-xr-x  4 root root 0 Jun 18 12:28 DS3200

/proc/mpp/DS3200:
total 0
dr-xr-xr-x  3 root root 0 Jun 18 12:28 controllerA
dr-xr-xr-x  3 root root 0 Jun 18 12:28 controllerB
-rw-r--r--  1 root root 0 Jun 18 12:28 virtualLun0

/proc/mpp/DS3200/controllerA:
total 0
dr-xr-xr-x  2 root root 0 Jun 18 12:28 mptsas_h0c0t17

/proc/mpp/DS3200/controllerA/mptsas_h0c0t17:
total 0
-rw-r--r--  1 root root 0 Jun 18 12:28 LUN0
-rw-r--r--  1 root root 0 Jun 18 12:28 UTM_LUN31

/proc/mpp/DS3200/controllerB:
total 0
dr-xr-xr-x  2 root root 0 Jun 18 12:28 mptsas_h1c0t33

/proc/mpp/DS3200/controllerB/mptsas_h1c0t33:
total 0
-rw-r--r--  1 root root 0 Jun 18 12:28 LUN0
-rw-r--r--  1 root root 0 Jun 18 12:28 UTM_LUN31
```

The Linux server can now access the logical drive, which is presented as LUN 0. The logical drive is accessible across both paths, controllerA and controllerB, and the RDAC (or MPP) makes sure that the dual path access is handled correctly.

You can also see the Access Logical Drive, presented as UTM_LUN31.

19.4.4 Configuring disk space in Linux

The logical drive is now visible in the host server. We can now create partitions and format them. Linux provides various tools for partition management, for example, **fdisk** or **parted**.

The logical drive is presented to Linux as /dev/sdb. We created two partitions, which appear under /dev/sdb1 and /dev/sdb2.

We can now mount these partitions and start using the disk space. For example, you could use the following commands to mount the partitions:

```
mount /dev/sdb1 /mnt/DS3200_1
mount /dev/sdb2 /mnt/DS3200_2
```

20

SAS configuration 5 - Boot from SAN with Windows 2008 on an IBM BladeCenter HS21 server

In this chapter, we describe and explain a sample configuration that shows a Boot from SAN configuration with a Windows Server 2008 operating system running on an IBM BladeCenter HS21 server booting from a DS3200 System Storage subsystem.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

20.1 Equipment required

We use the following hardware and software components:

- IBM BladeCenter E and HS21 BladeCenter.
- IBM BladeCenter Advance Management Module (AMM) installed in BladeCenter.
- Two IBM BladeCenter SAS Connectivity Modules.
- IBM BladeCenter SAS Expansion card (CFFv), P/N 43W3974.
- DS3200 System Storage subsystem.
- The latest version of DS3000 Storage Manager running on an external management workstation (at the time of writing, this is Version 10.35).

Figure 20-1 shows our sample hardware setup.

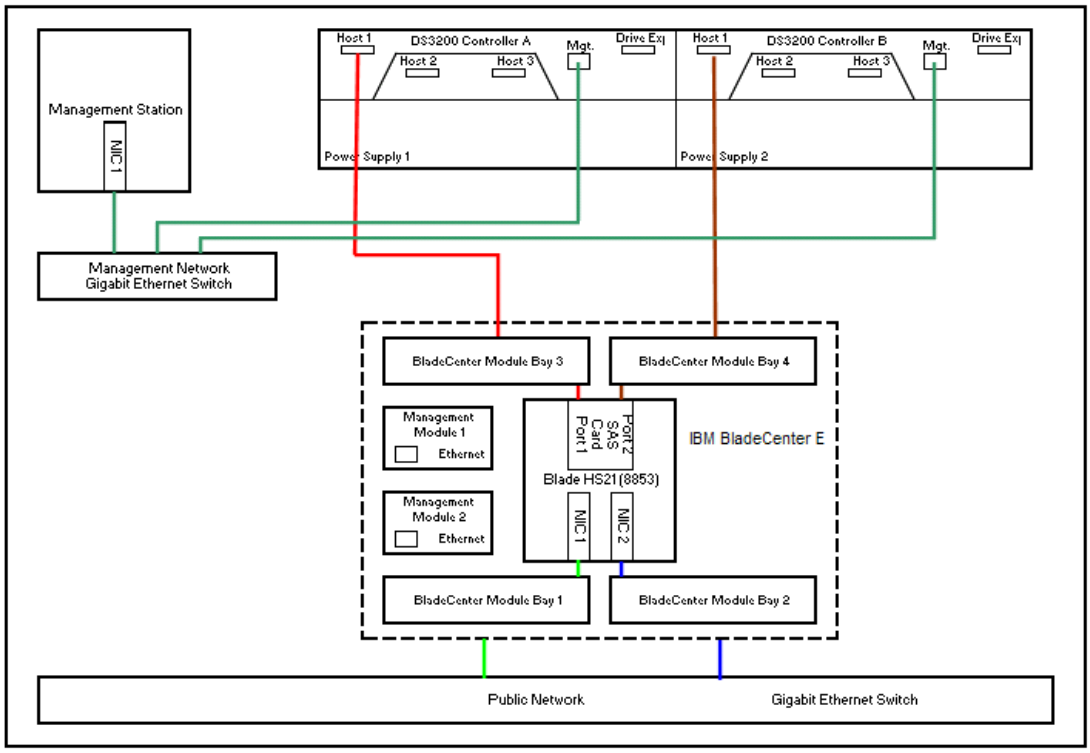


Figure 20-1 Hardware setup

20.2 IBM BladeCenter setup

Check the IBM support Web site for the latest DS3000 interoperability matrix to ensure that all hardware used for your Boot from SAN configuration are supported. Ensure that the BladeCenter, blade server, and modules are updated with the latest firmware. Consult the IBM support Web sites for the latest firmware.

20.2.1 HS21 SAS Expansion Card

We have to install the IBM BladeCenter SAS Expansion Card in the HS21 BladeCenter host before we can proceed with other tasks. While this is not a difficult task, we do recommend that you consult the user's guide for the host server and follow the instructions for options installation. Do not turn on the HS21 server at this time.

Important: The connectivity modules in I/O module bay 3 and I/O module bay 4 and all expansion cards in the BladeCenter unit must use the same interface type. Therefore, you must install SAS expansion cards before you install connectivity modules in the blade servers in your BladeCenter unit. For more information about the SAS expansion card, see the *Installation and User's Guide* for the SAS expansion card at http://www.ibm.com/systems/support/.

Note: The BladeCenter SAS Expansion Card is a dual port card. Port #1 connects to SAS Connectivity Module in BladeCenter I/O module bay 3 and port #2 connects to SAS Connectivity Module in BladeCenter module bay 4 respectively.

20.2.2 SAS Connectivity Modules

You must install the IBM BladeCenter SAS Connectivity Modules only in BladeCenter I/O module bay 3 and I/O module bay 4 of the following supported BladeCenter units:

- BladeCenter Type 8677
- BladeCenter Types 8720 and 8730
- BladeCenter Type 8750
- BladeCenter Type 8852
- BladeCenter Type 8886

Installing a connectivity module in I/O module bay 3 or I/O module bay 4 provides connectivity to the SAS expansion card(s) installed in the blade servers in your BladeCenter unit. Installing two connectivity modules allows you to have two connections to the SAS expansion cards installed in the blade servers.

Important: The connectivity modules in I/O module bay 3 and I/O module bay 4 and all expansion cards in the BladeCenter unit must use the same interface type. Therefore, you must install SAS expansion cards before you install connectivity modules in the blade servers in your BladeCenter unit. For more information about the SAS expansion card, see the *Installation and User's Guide* for the SAS expansion card at http://www.ibm.com/systems/support/.

Connect the cables from the DS3200 Storage System controllers A and B to the external port #3 of the two SAS Connectivity Modules, as shown in Figure 20-2.

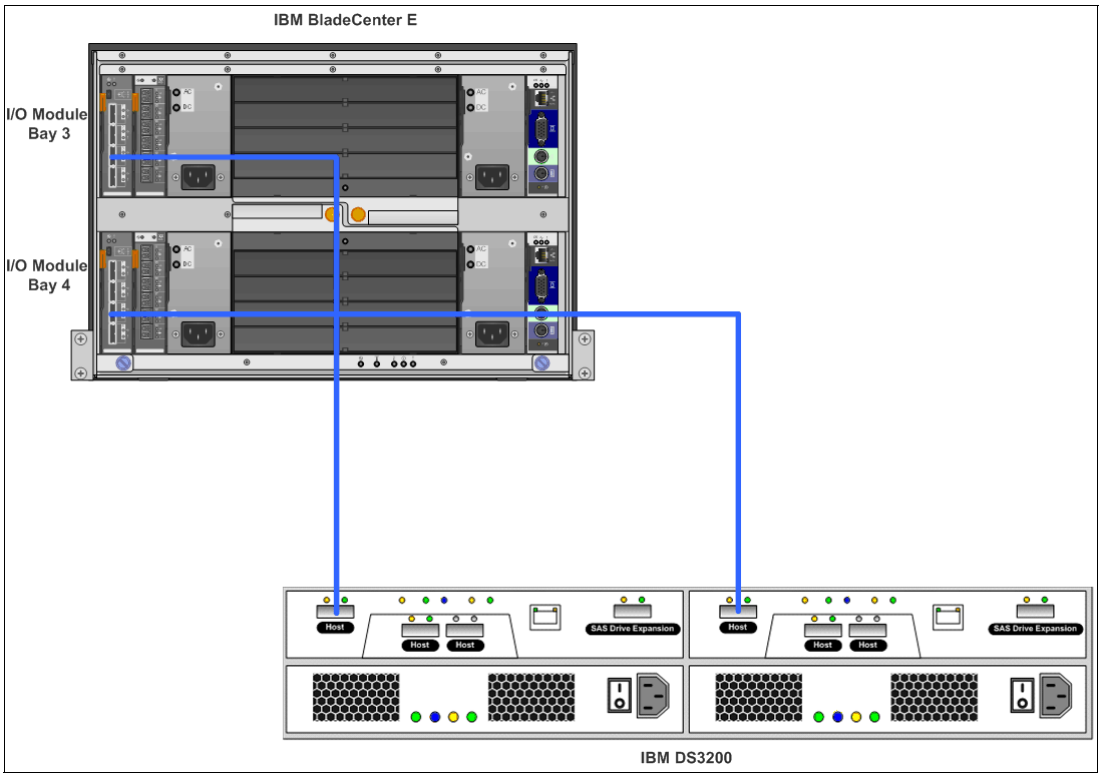


Figure 20-2 BladeCenter and DS3200 connectivity

Important: At the time of writing, only a connection to external port #3 of the SAS Connectivity Modules is supported on the DS32000 System Storage subsystem when booting from SAN.

20.2.3 SAS Connectivity Module firmware update

Ensure that your SAS Connectivity Module is updated with the latest firmware. For the latest firmware update, refer to `http://www.ibm.com/systems/support/`.

To update the connectivity module firmware to the latest version, complete the following steps:

1. Log on to the SAS Connectivity Module using the Web interface with the IP address defined for the connectivity module in the BladeCenter Advance Management Module (AMM), as shown in Figure 20-3 on page 449. Use USERID as the user ID and PASSW0RD as the password. You can change the password under the Administer Users menu option once logged on.

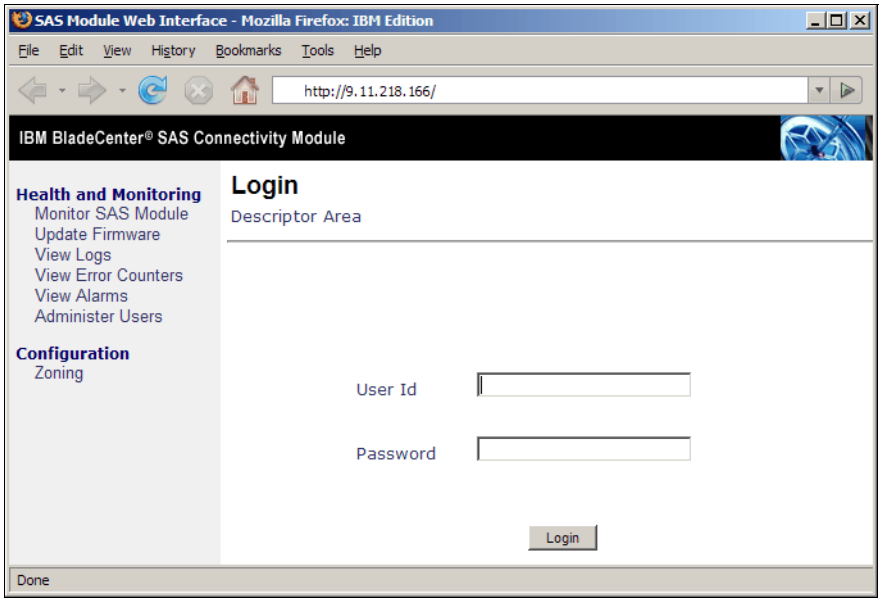


Figure 20-3 SAS Connectivity Module - Login

2. In the Monitor Module window, click **Update Firmware**. The Update Firmware window opens, as shown in Figure 20-4. The current firmware level is also displayed.

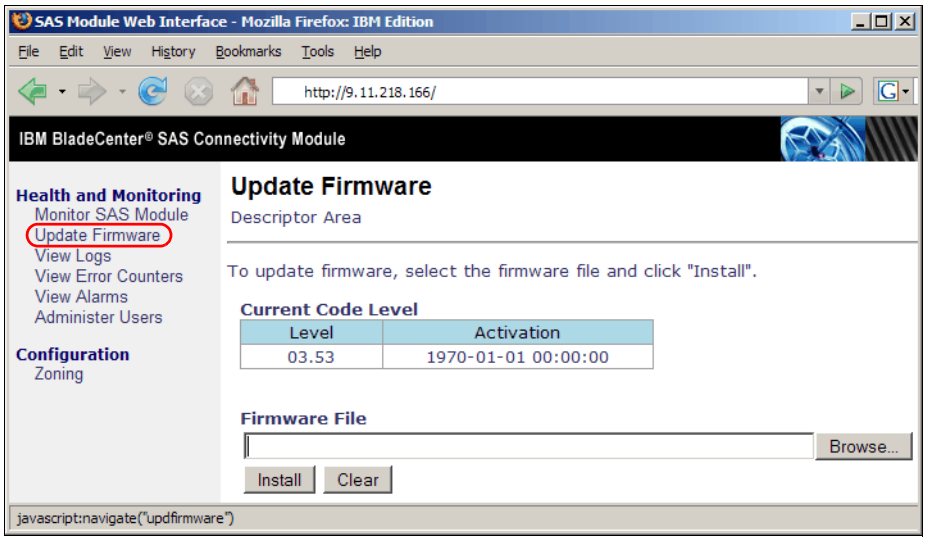


Figure 20-4 SAS Connectivity Module - Update Firmware

3. In the Firmware File field, enter the new firmware file name or click **Browse** and locate the firmware file.

4. Click **Install** to install the new file. An installation confirmation window opens, as shown in Figure 20-5.

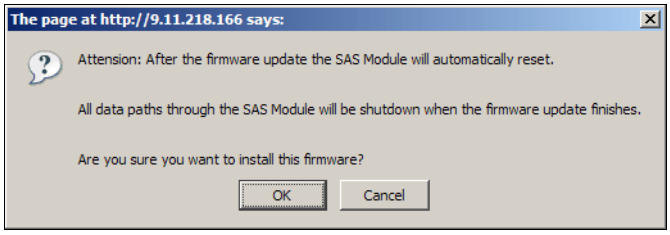


Figure 20-5 SAS Connectivity Module - Installation confirmation

5. Click **OK** or **Cancel**. If the installation of the new firmware file is successful, an installation confirmation window opens, as shown in Figure 20-6. If there are errors during the installation, an error message window opens.

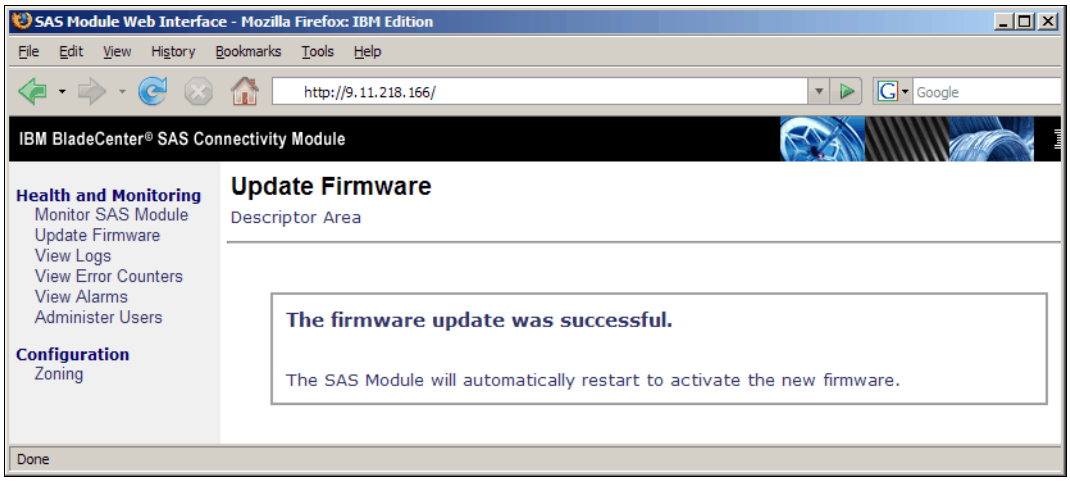


Figure 20-6 SAS Connectivity Module - Update successful

20.3 Boot from SAN setup

Here we discuss the Boot from SAN setup.

20.3.1 SAS Connectivity Module configuration

Ensure that the external ports on the SAS Connectivity Module in I/O slot bay 3 are enabled to allow connectivity from the DS3200 System Storage subsystem controller A to this module. Do not activate the external ports on the SAS Connectivity Module in I/O slot bay 4 that is connected to controller B at this time. Should it be enabled, disable it.

From the BladeCenter Advance Management Module Web interface GUI, select **I/O Module Tasks** → **Admin/Power/Restart** to ensure that the external ports for the I/O module in bay 3 are enabled, as shown in Figure 20-7.

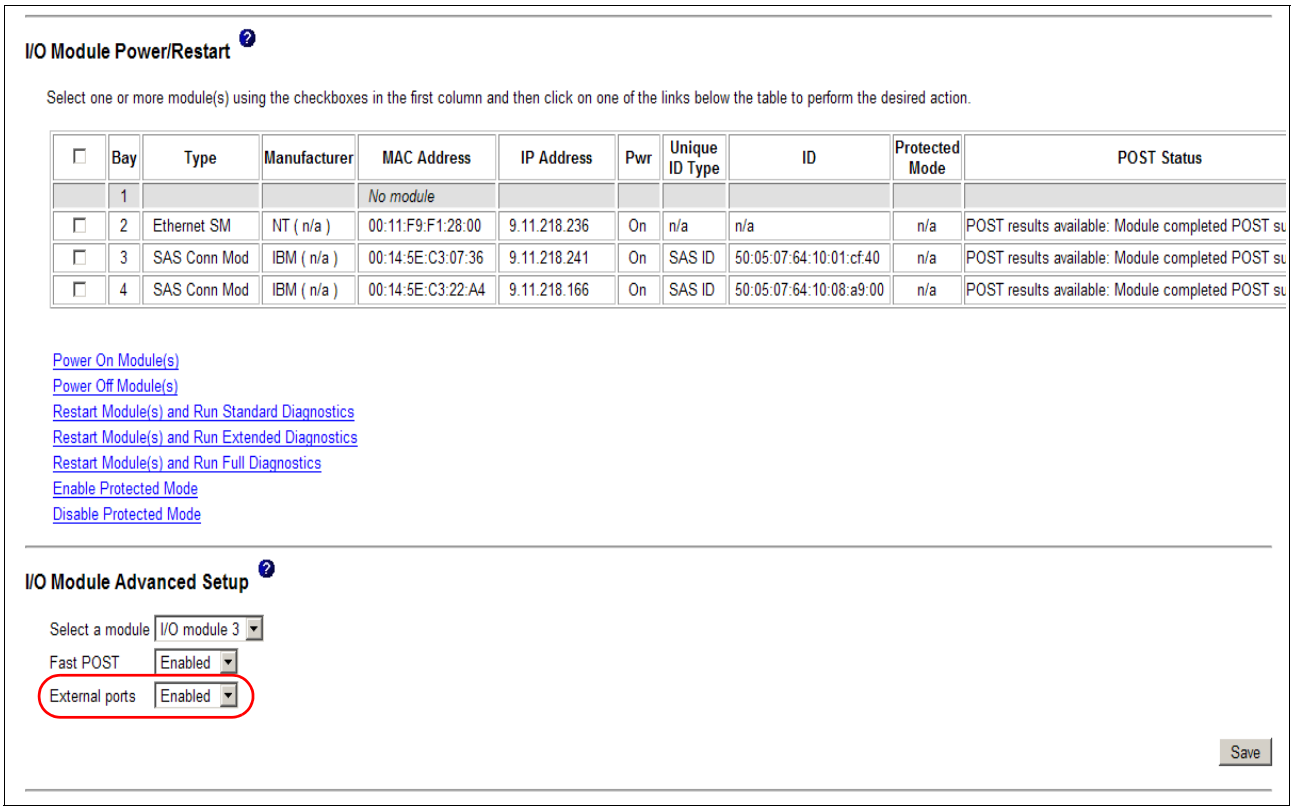


Figure 20-7 SAS Connectivity Module - Enabling external ports

20.3.2 SAS Connectivity Module zoning

Several predefined SAS fabric zone configurations are available from the factory and can be invoked by a simple selection from the Advanced Management Module (AMM). Ensure that the option #5 predefined config 01 is selected for the SAS Connectivity Modules in I/O module bay 3 and 4.

Select **I/O Module Tasks** → **Configuration** in the BladeCenter AMM GUI Web interface window. For I/O Module 3, select **Zone Configuration Management**. A window opens showing the zone configuration options, as shown in Figure 20-8.

In our example, option #5, predefined config 01 zone configuration, is currently active (indicated by a check mark), as shown in Figure 20-8. With this option, each server bay is exclusively zoned with all the external ports, thus allowing access to one or more storage controller ports connected to the SAS Connectivity Module.

Zone Configuration Management for SAS Modules

The table below displays zone configurations stored on the given SAS Connectivity Module. Please select the desired zone configuration from the list and activate it. The 'Refresh' button would be helpful in refreshing the status once the zone is applied.If you have multiple SAS Connectivity Modules installed and both are in working order, a check box will be provided that allows you to easily apply the same configuration to each SAS Connectivity Module. The default setting is to apply the same zone configuration to each. If you uncheck the check box, information for both SAS Connectivity Modules will be presented and you can select a zone configuration from each. However, it is highly recommended that you select the same zone configuration for both SAS Connectivity Modules.

☑ Apply the same zone configuration to both SAS Connectivity Modules

I/O Module 3 (SAS Module)

The table below lists zone configurations stored on this SAS Module.

Select	Active?	Name	Type	Intended # of Blades	Intended # of SAS Modules	Max Disks per Blade	Configuration Store	Date
				Description				
○		User Defined Config 01	User-defined	Chassis: Any. SAS modules: 1 or 2. Default zone setting is each SAS module port belongs to its own zone and no port can access any other port. Can be modified using SCM, the Telnet interface, or the embedded Web browser interface.			1	00/00/0000, 00:00:00
○		User Defined Config 02	User-defined	Chassis: Any. SAS modules: 1 or 2. Default zone setting is each SAS module port belongs to its own zone and no port can access any other port. Can be modified using SCM, the Telnet interface, or the embedded Web browser interface.			2	00/00/0000, 00:00:00
○		User Defined Config 03	User-defined	Chassis: Any. SAS modules: 1 or 2. Default zone setting is each SAS module port belongs to its own zone and no port can access any other port. Can be modified using SCM, the Telnet interface, or the embedded Web browser interface.			3	00/00/0000, 00:00:00
○		User Defined Config 04	User-defined	Chassis: Any. SAS modules: 1 or 2. Default zone setting is each SAS module port belongs to its own zone and no port can access any other port. Can be modified using SCM, the Telnet interface, or the embedded Web browser interface.			4	00/00/0000, 00:00:00
○	✓	Predefined Config 01	Pre-defined	Chassis: BCE, BCH, BCT and BCHT. SAS modules: 1 or 2. Zoned Blade Bays: 1-14. Each SAS module port belongs to its own zone. All Blades can access all external ports. Cannot be modified.			5	04/24/2007, 02:00:00

Figure 20-8 SAS Connectivity Module - Predefined zones

The following example in Figure 20-9 on page 453 was captured from the SAS connectivity module Web interface. It lists the Basic Zone Permission table for the HS21 blade in slot 6, because the blade in slot 6 is used for this example. The blade in slot 6 is zoned with four external ports. Notice that the External port is set to True under the Connected column and Normal under the Status column. This is because the DS3200 Storage System Controller A port is connected to the external port #3 for this example. There is only one external connection from Controller A to the SAS connectivity module in I/O bay 3.

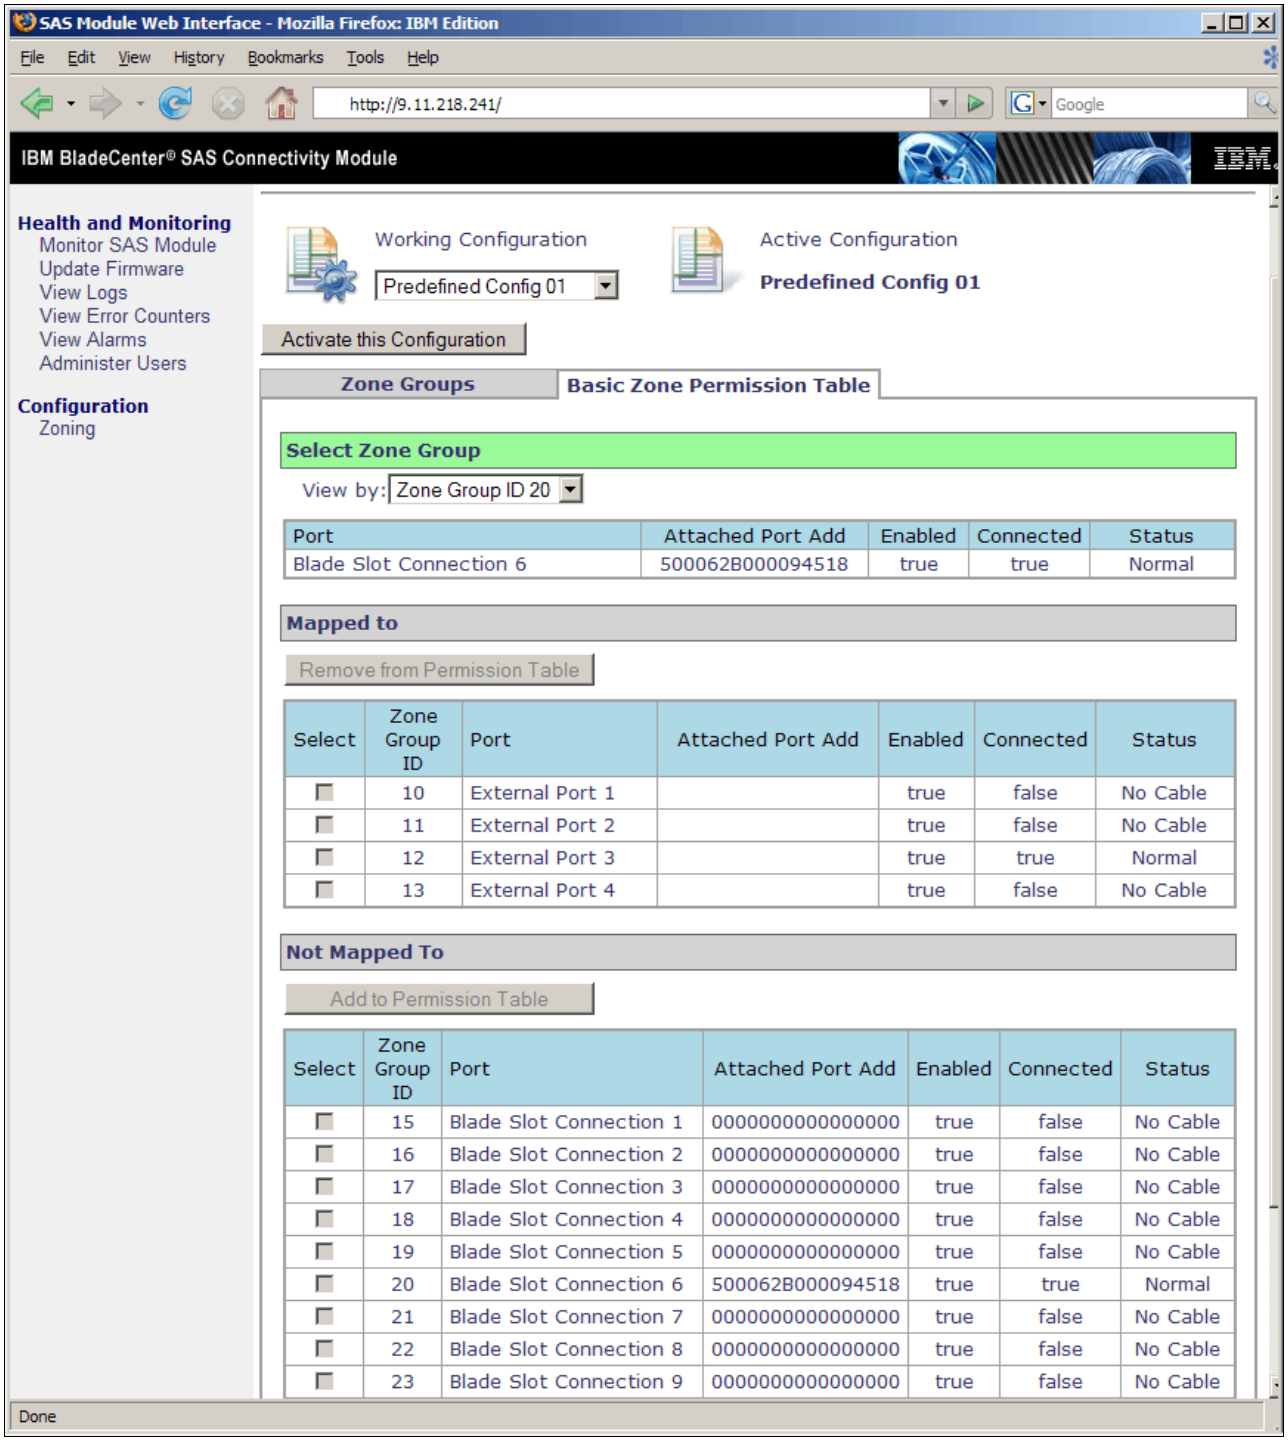


Figure 20-9 SAS Connectivity Module - Zone configuration

Note: There is only a single path active from the DS3200 Storage System controller A to the SAS Connectivity Module. The second path from DS3200 Storage System controller B to the SAS Connectivity Module is purposely disabled for the OS installation process to complete smoothly through the single path to the boot volume.

20.3.3 SAS controller and SAS Expansion Card configuration

The following example demonstrates how to disable the onboard SAS controller and the configuration of the SAS Expansion Card:

1. Power on the HS21 BladeCenter host.
2. Press F1 on bootup to enter the BIOS configuration menu, as shown in Figure 20-10.

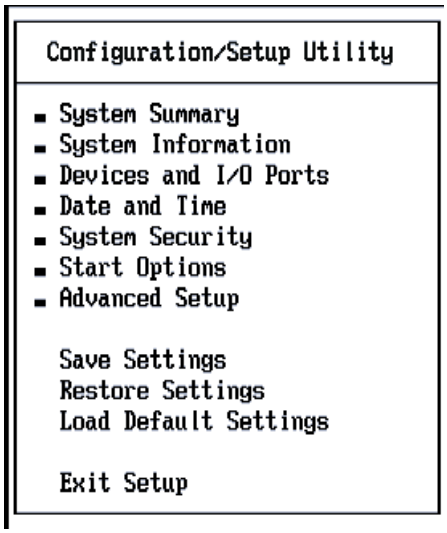


Figure 20-10 BIOS configuration menu

3. Select **Device and I/O Ports** from the menu.
4. Select the **Disable** option for the Planar SAS component, as shown in Figure 20-11. This option will disable the onboard SAS controller.

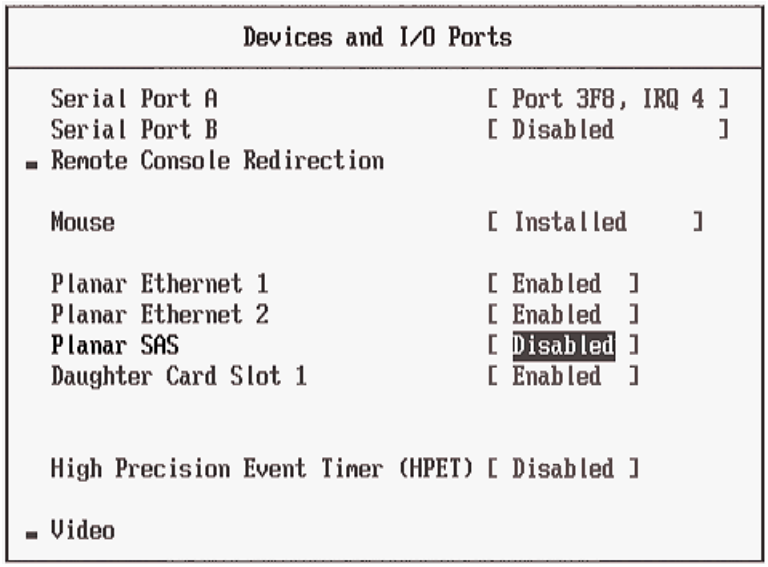


Figure 20-11 BIOS configuration - Planar SAS disable

5. Press Esc to exit the BIOS configuration utility and save the changes, as shown in Figure 20-12 on page 455.

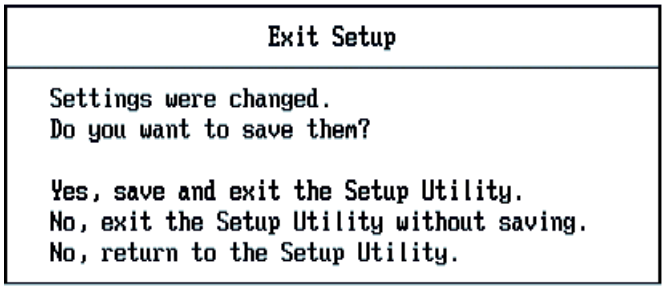


Figure 20-12 Exit BIOS configuration

6. The server will now continue to boot. Press <Ctrl + C> to enter the LSI Logic Configuration Utility when you reach the screen shown in Figure 20-13.

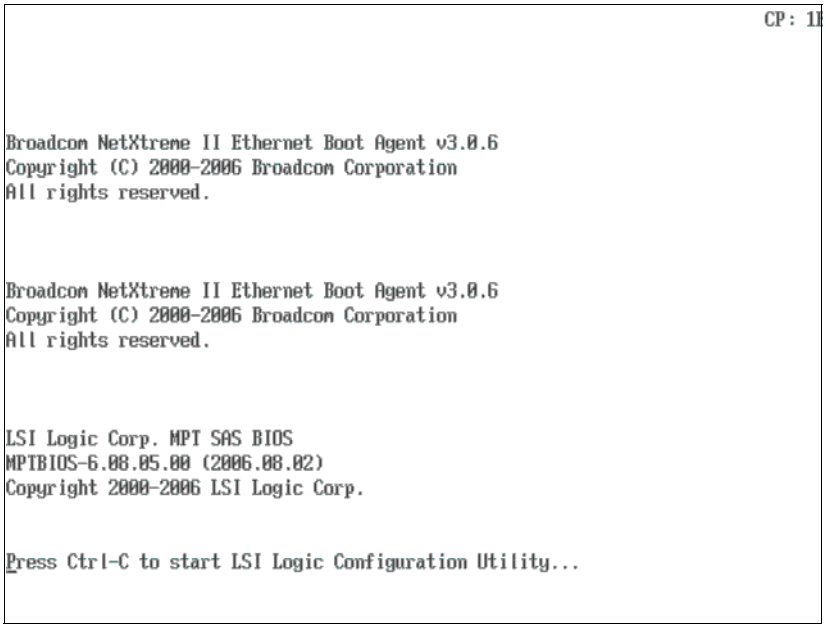


Figure 20-13 SAS Expansion Card - LSI Logic Configuration Utility

7. The following menu is displayed after pressing Crtl + C, as shown in Figure 20-14.

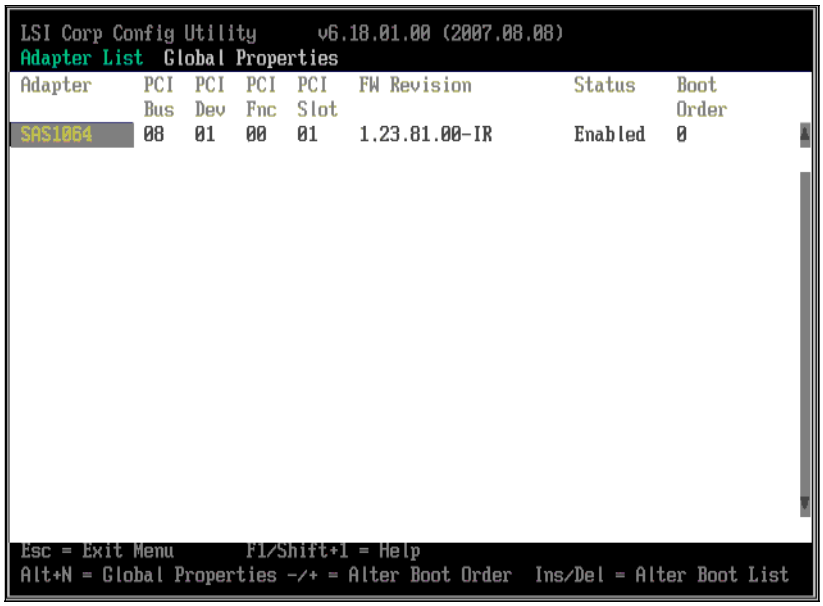


Figure 20-14 SAS Expansion Card - LSI Logic Config Utility menu

8. Press the Enter key to select the SAS adapter internally connected to the SAS connectivity module in I/O Module bay 3, as shown in Figure 20-15.

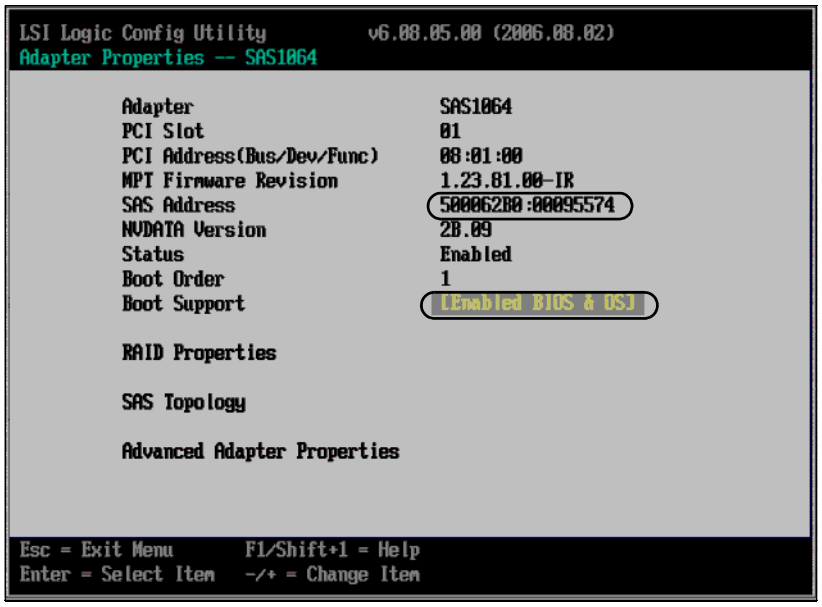


Figure 20-15 SAS Expansion Card - LSI Logic Config Utility adapter selected

Record the Worldwide Port Name (WWPN) of the first port on the SAS Expansion Card, reflected as the SAS Address. The WWPN is needed for defining host ports on the DS3200 System Storage subsystem when defining the host-to-LUN mapping. The WWPN can also be retrieved from the SAS Connectivity Module Web interface.

Ensure that the Boot Support parameter is set to Enabled BIOS & OS to allow the SAS Expansion Card to be activated during bootup.

9. Exit the LSI Configuration Utility and save the changes.

20.3.4 DS3200 boot volume setup

To complete the boot volume setup, we have to make use of an external management workstation running the DS3000 Storage Management software to create the volume and set up the host-to-LUN configuration.

For the exact installation steps for a DS3000 Storage Manager, see Chapter 6, “Installing IBM System Storage DS3000 Storage Manager” on page 81.

For our example, we create one logical volume for the HS21 BladeCenter to act as a boot volume for installation of the Windows Server 2008 operating system. The following steps show how to create such a volume:

1. In the DS3200 Subsystem Management window, click the **Configure** tab, as shown in Figure 20-16.

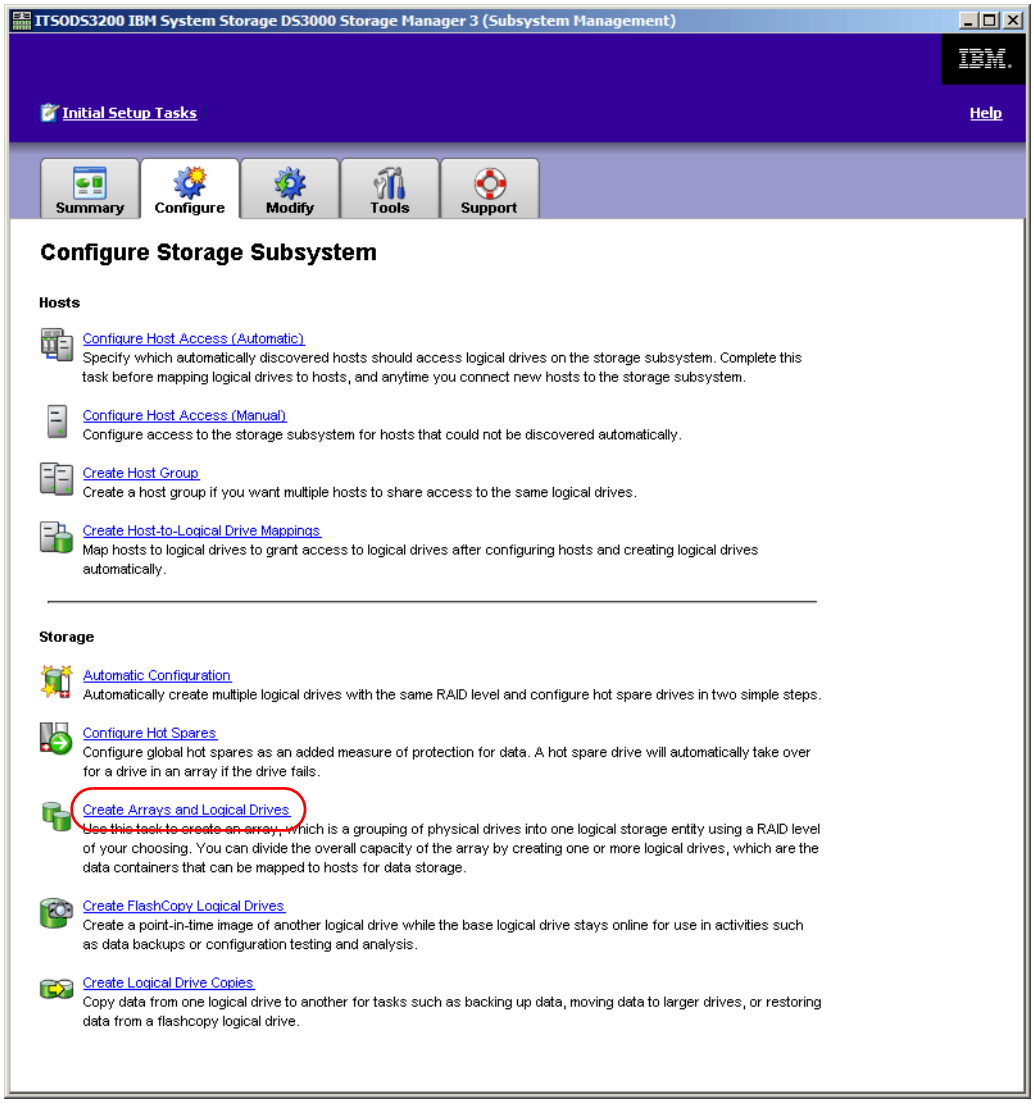


Figure 20-16 DS3200 Subsystem Management Window - Configure

2. Click the **Create Arrays and Logical Drives** option. Click **OK** if presented with the message shown in Figure 20-17. The host is configured after the boot logical drive setup is completed.

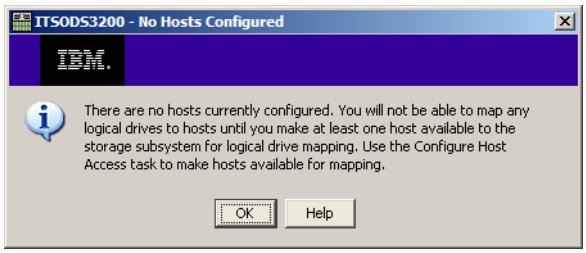


Figure 20-17 No Hosts Configured message

3. Select the **Array:Create new array using the unconfigured capacity in the storage subsystem** radio button and click **Next** (Figure 20-18).

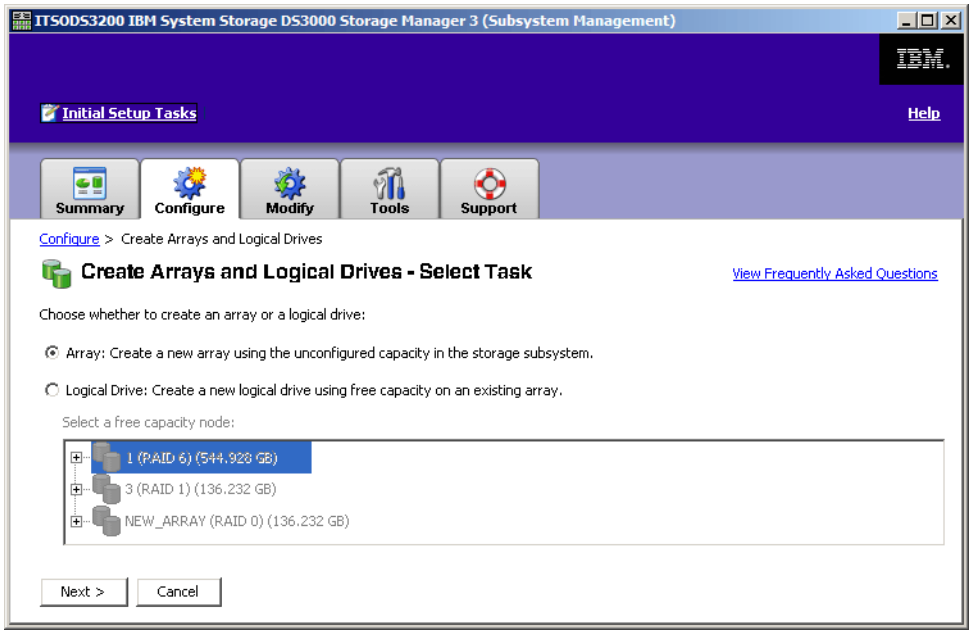


Figure 20-18 Create New Array

4. Type in an array name in the Array name section. In our example, we use Boot_Array as the array name. Select the **Automatic** radio button (recommended option). This will automatically select the physical drives to create the logical volume. Click **Next** to continue. See Figure 20-19 on page 459.

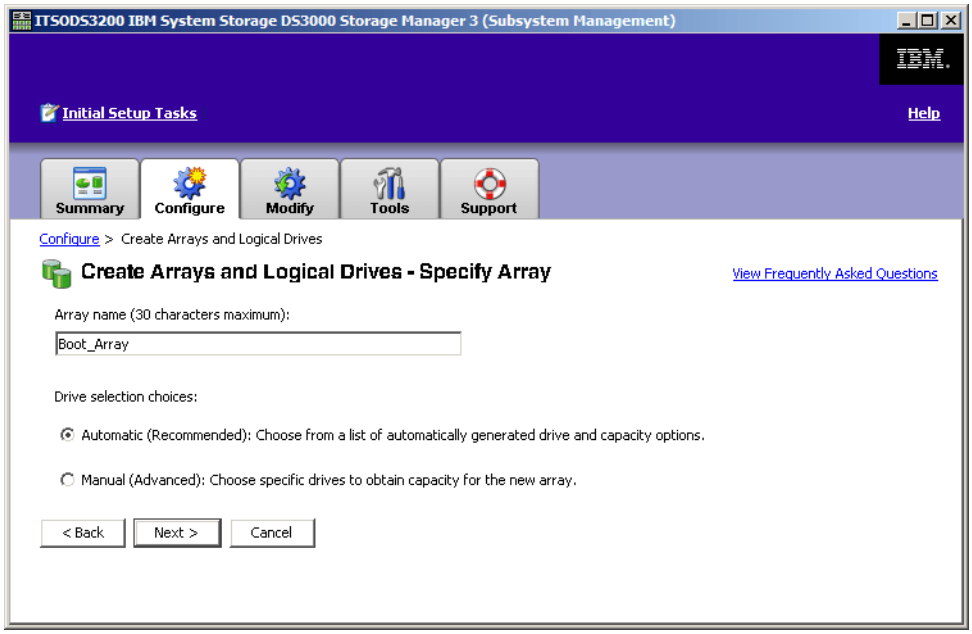


Figure 20-19 New array name

5. Select the RAID level from the drop-down list. In our example, we use RAID 5 for the boot volume, as shown in Figure 20-20. Click **Finish**.

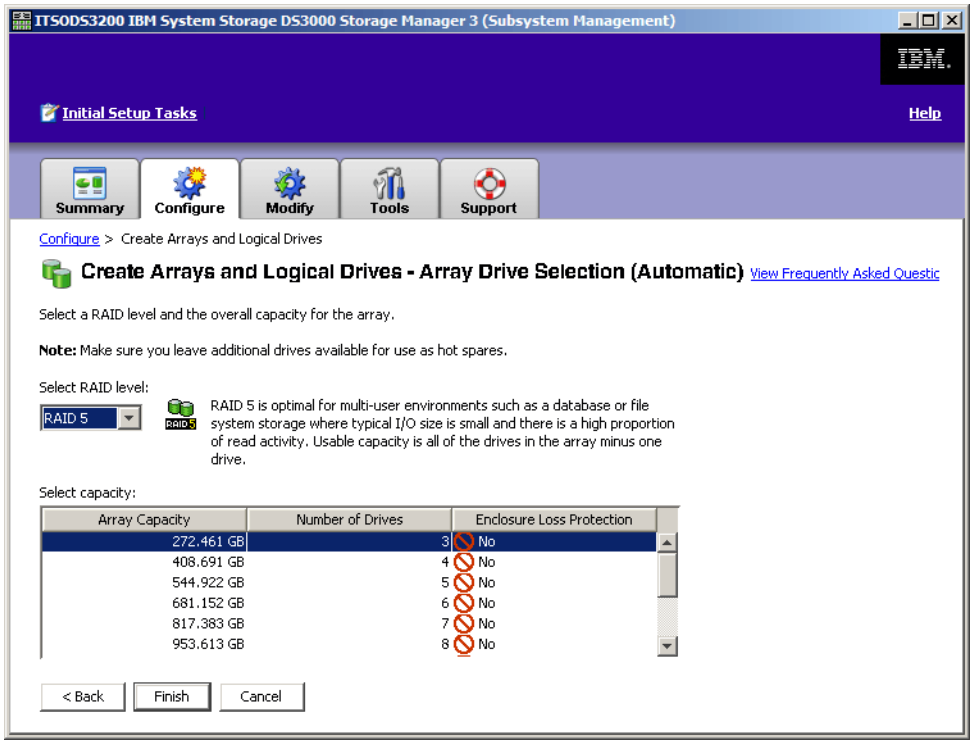


Figure 20-20 RAID Selection

6. A confirmation appears, as shown in Figure 20-21. Select the **Create a logical drive using the new array** radio button to create the new logical drive. Click **Yes**.

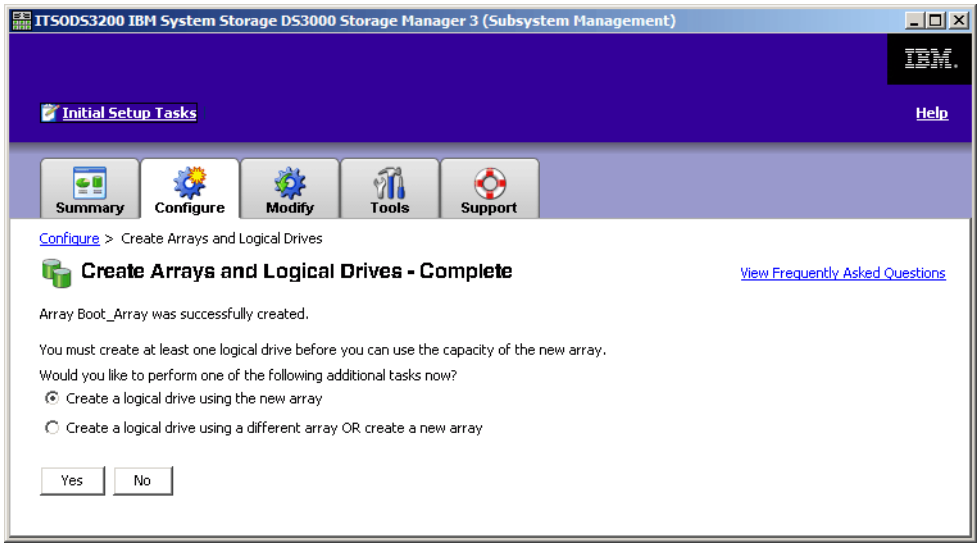


Figure 20-21 Create logical drive

7. Specify the size and name of the boot volume, as shown in Figure 20-22, and then click **Next**.

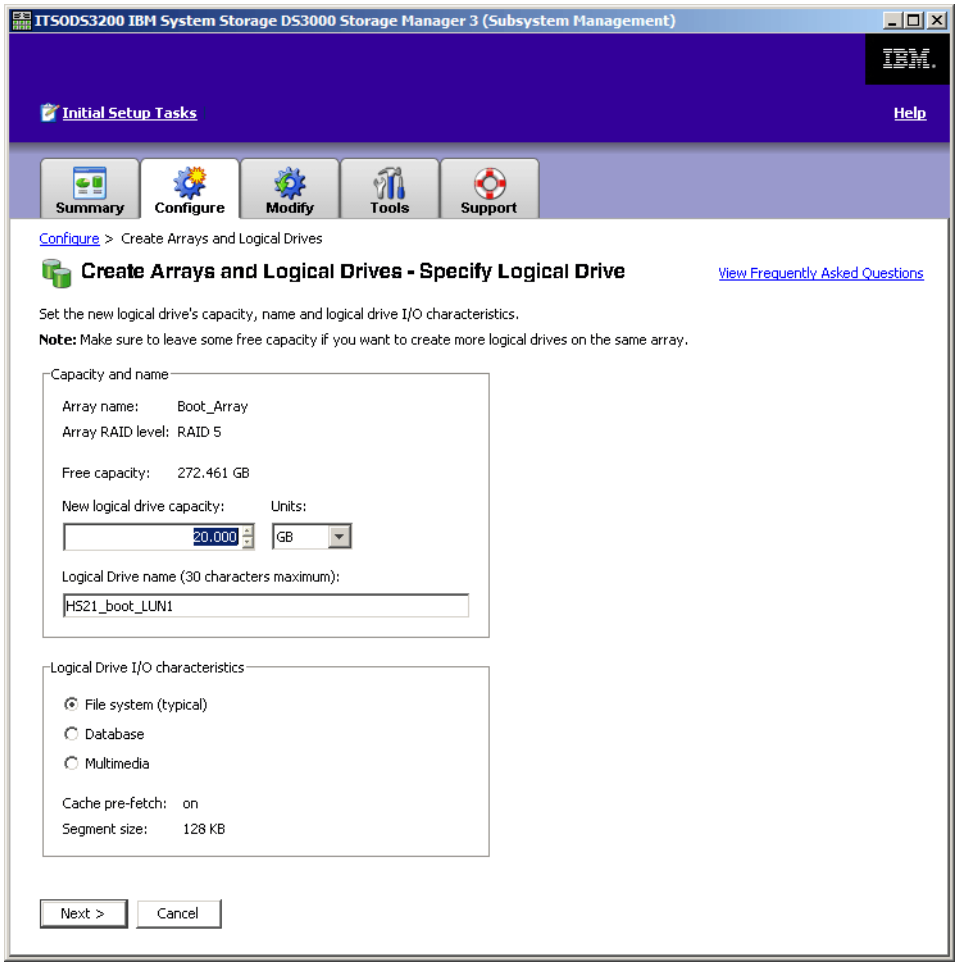


Figure 20-22 Logical drive size and name

8. Leave the **Map later** radio button selected, as shown in Figure 20-23. Click **Finish** to complete the step.

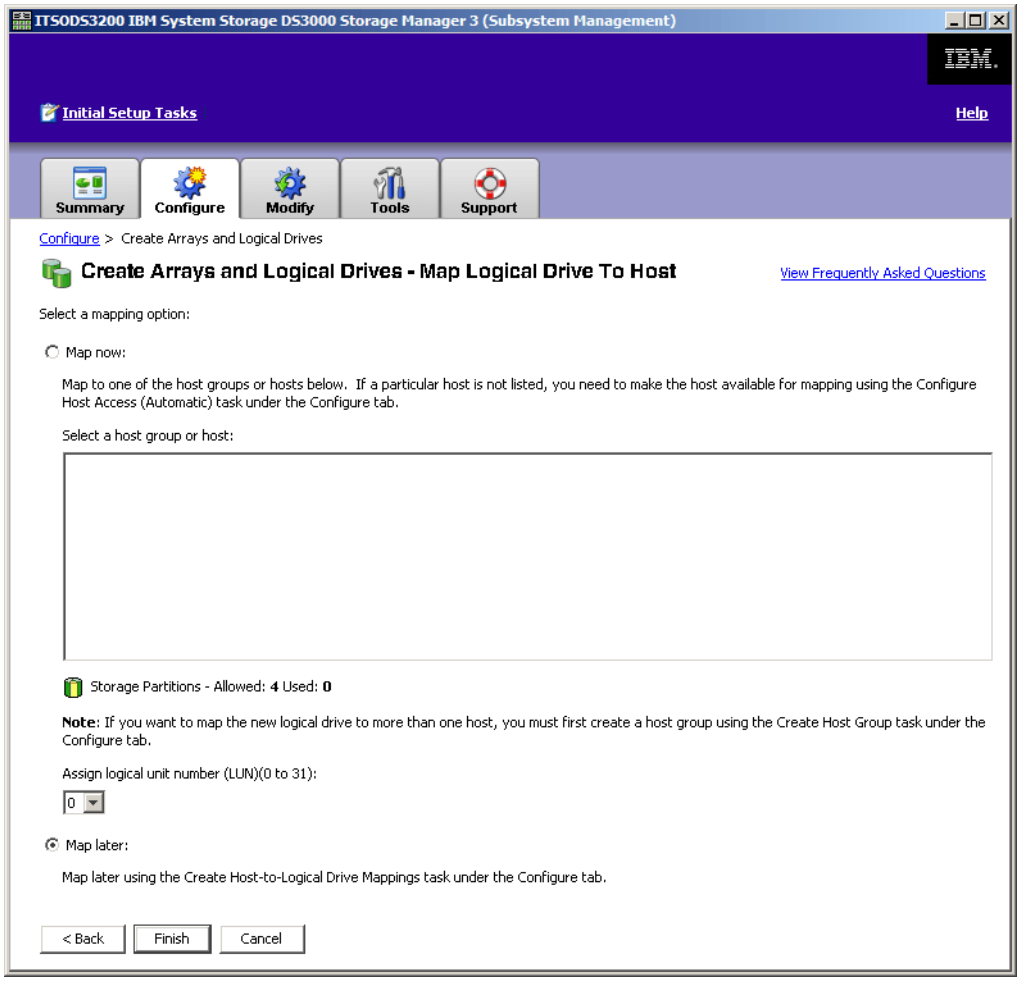


Figure 20-23 LUN ID

9. Click **No** to the question shown in Figure 20-24 on page 463 to complete the boot volume setup.

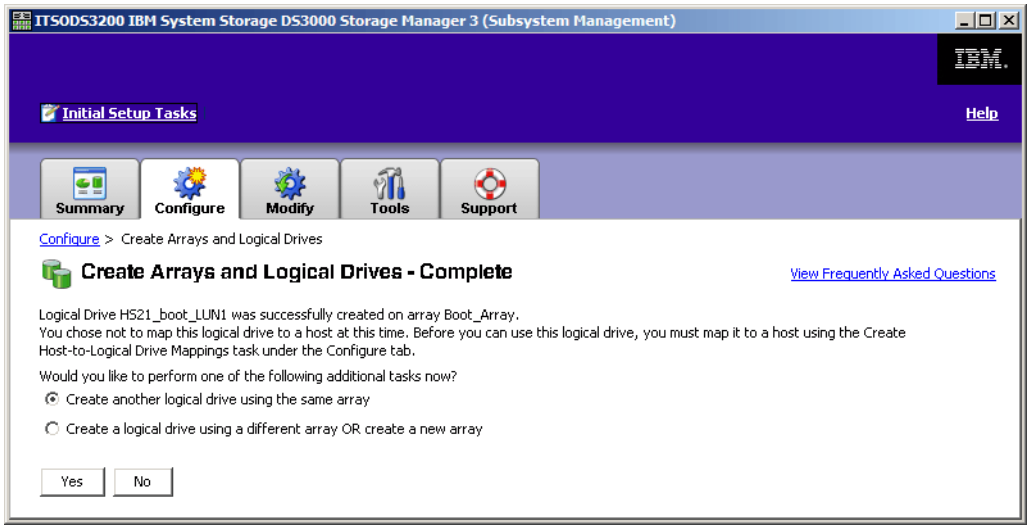


Figure 20-24 Logical drive completion

Note: When logical drives are created, Storage Manager will alternatively assign either controller A or controller B as the preferred owner. For this example, ensure that the boot volume created is assigned to controller A through the Change Logical Drive Ownership/Preferred Path option on the Modify tab of the DS3000 Storage Manager client.

20.3.5 DS3200 host mapping

Use the DS3000 Storage Manager client to define the host, host type, and host ports. Make sure that the host type defined matches the operating system you are installing, in our example, Windows Server 2008.

Do these steps;

1. Select **Configure Host Access (Manually)** from the Configure tab in the DS3200 Subsystem Management window, as shown in Figure 20-25.

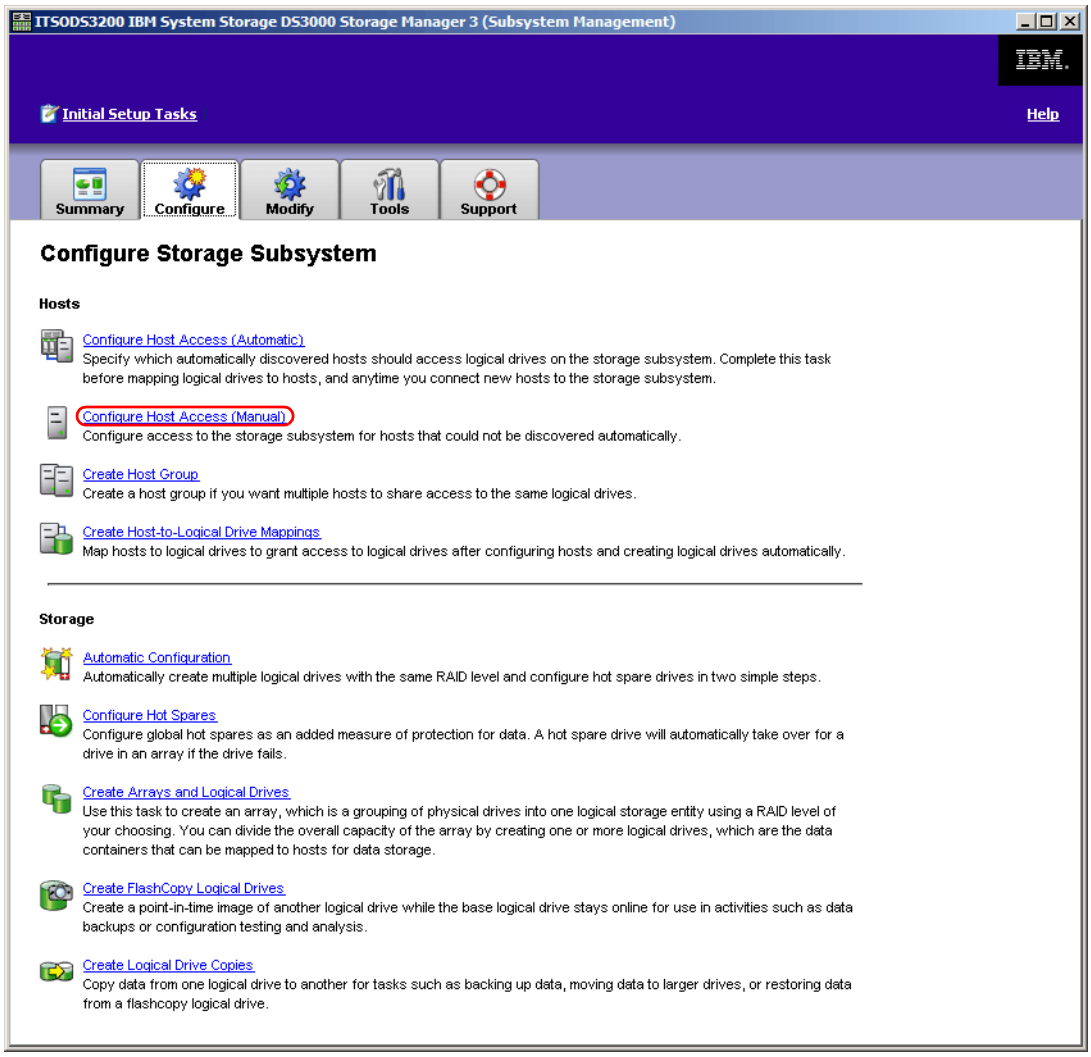


Figure 20-25 Configure Host Access

2. Specify the host name (in our example, HS21_Blade06) and select the appropriate host type (in our example, Windows Server 2008 Non-Clustered), as shown in Figure 20-26 on page 465.

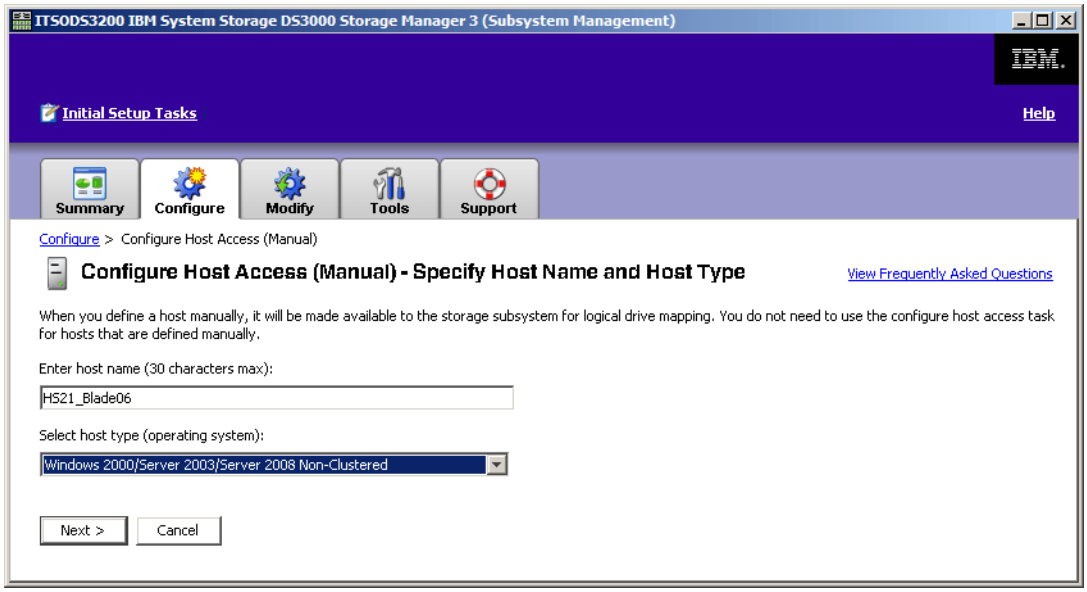


Figure 20-26 Host name and type selection

3. Select the WWPN of the first port of the SAS Expansion Card that was recorded in 20.3.3, “SAS controller and SAS Expansion Card configuration” on page 454. Click **Add** to add it to the selected host port identifiers section, as shown in Figure 20-27.

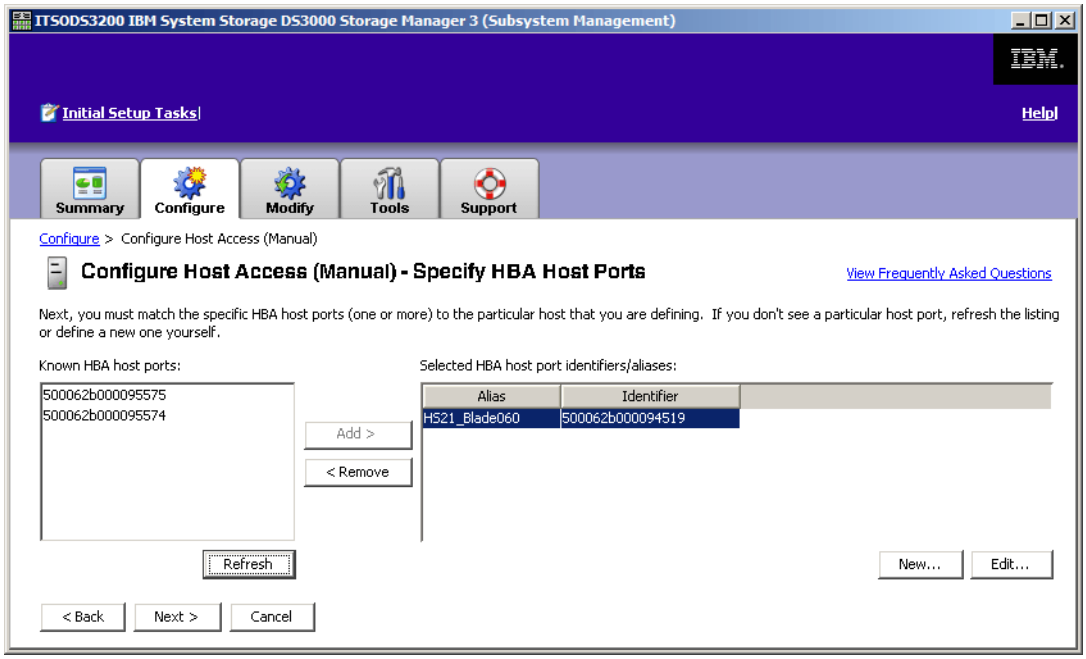


Figure 20-27 SAS Expansion Card - Host port selection

4. Click **Edit** to specify the alias of the selected host port of the SAS Expansion Card (in our example, we use HS21_Blade06_P1) and click **OK**, as shown in Figure 20-28.

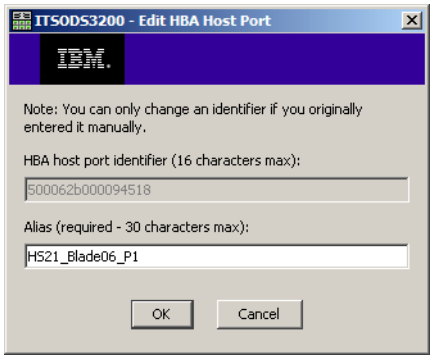


Figure 20-28 SAS Expansion Card - Host port alias

5. You will be taken back to the previous window. Click **Next** to continue.
6. Select the **No: This Host will NOT share access to the same logical drives with other hosts** radio button and click **Next,** as shown in Figure 20-29.

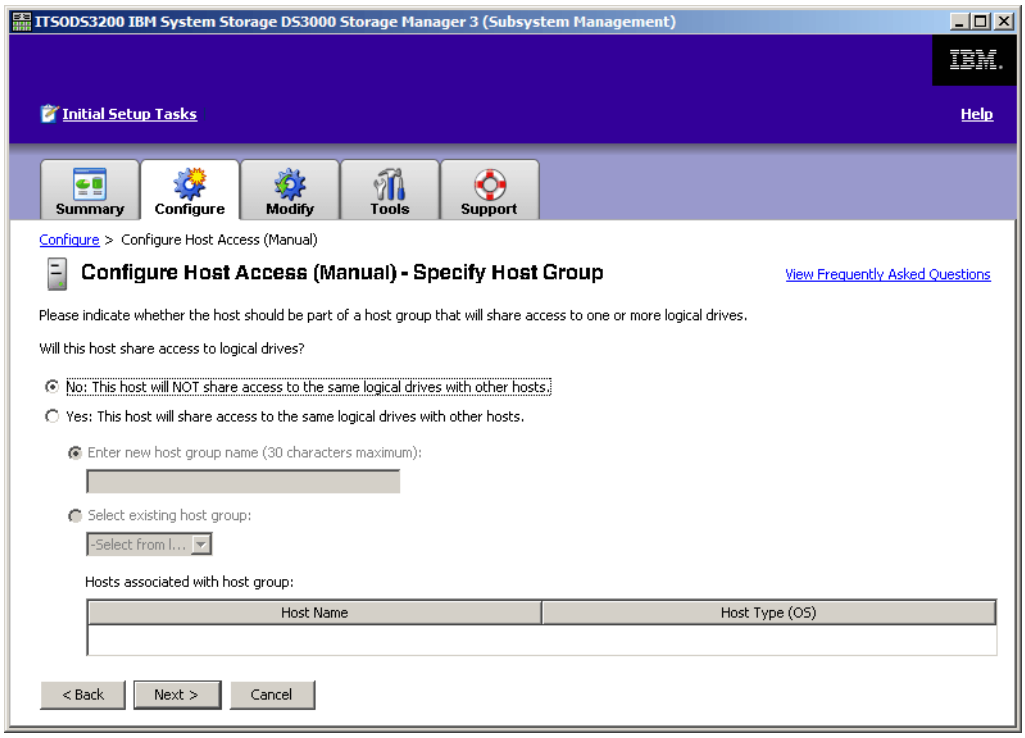


Figure 20-29 Logical Drive shared access

7. The next window shows all the selected parameters. Check to ensure that they are correct and click **Finish** to complete the step, as shown in Figure 20-30 on page 467.

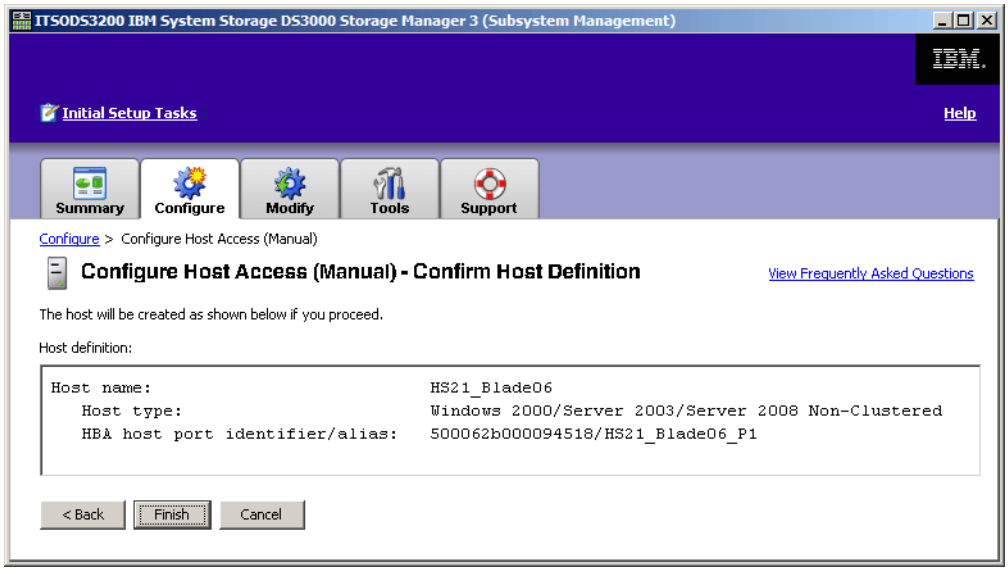


Figure 20-30 Host access confirmation

8. Click **No** to exit the host configuration when presented with the next window, as shown in Figure 20-31.

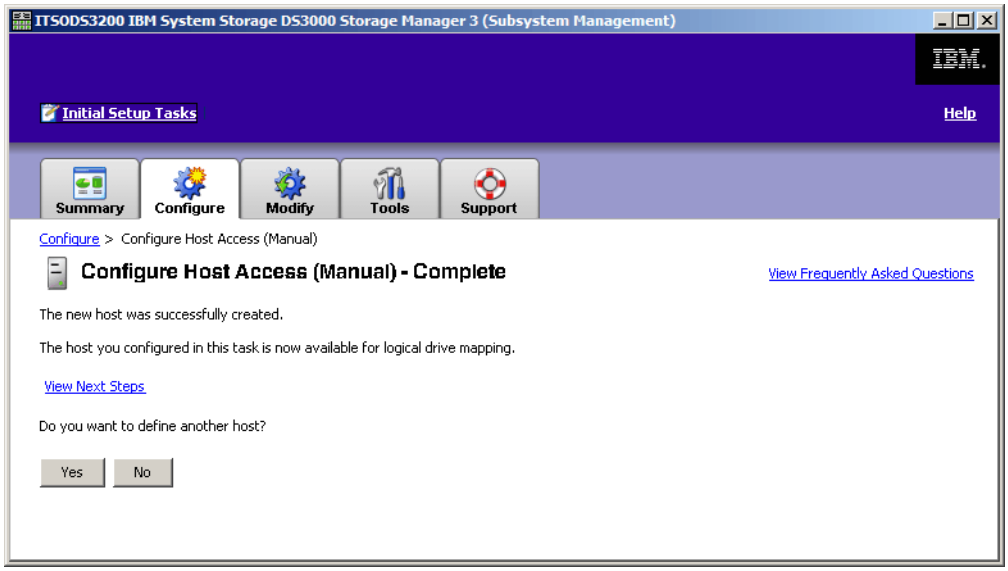


Figure 20-31 Host access completed

9. We now proceed with creating the host-to-logical drive mapping. Select **Create Host-to-Logical Drive Mappings** from the Configure tab in the DS3200 Subsystem Management window, as shown in Figure 20-32.

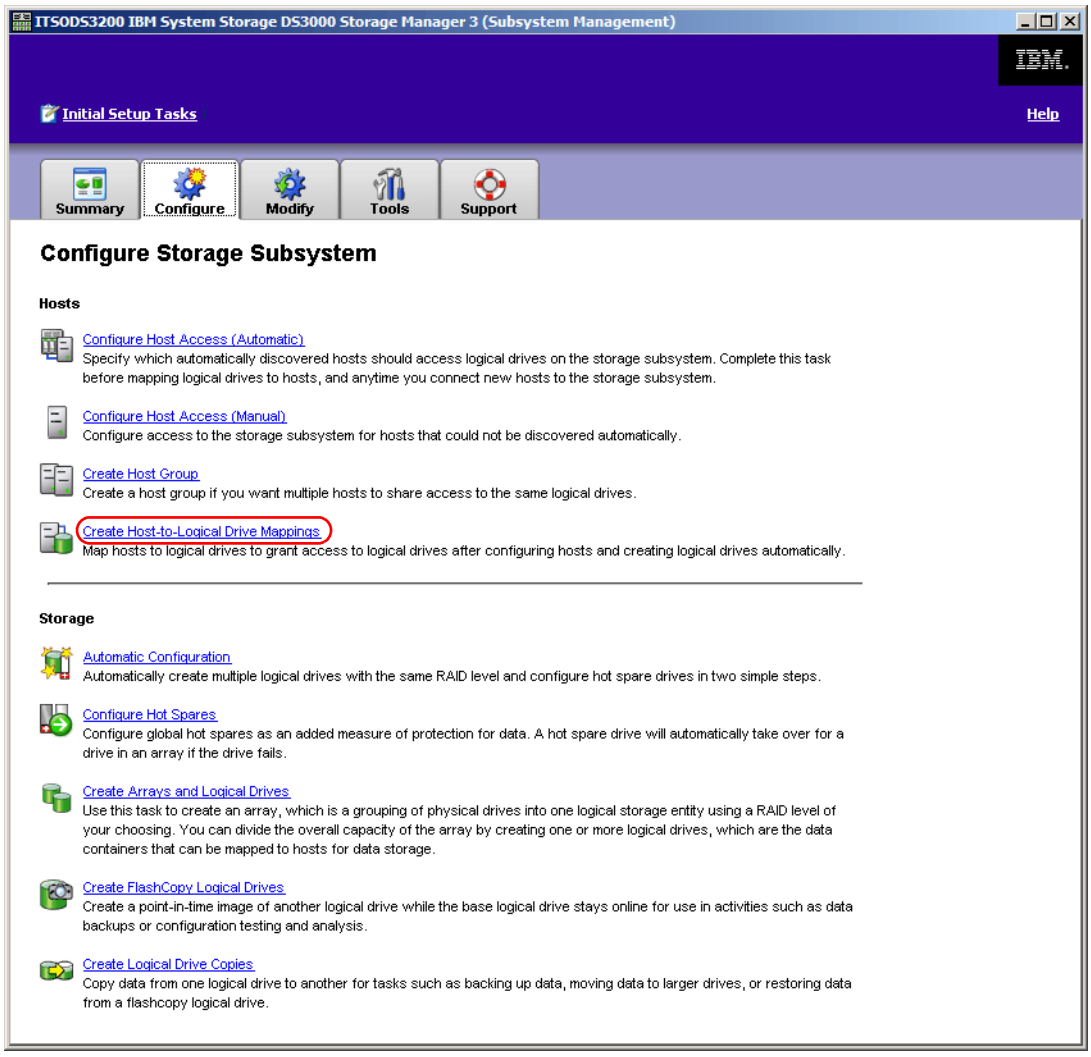


Figure 20-32 Create Host-to-Logical Drive Mappings

10. The next window presents the host previously created, as shown in Figure 20-33 on page 469. Select the host and click **Next**.

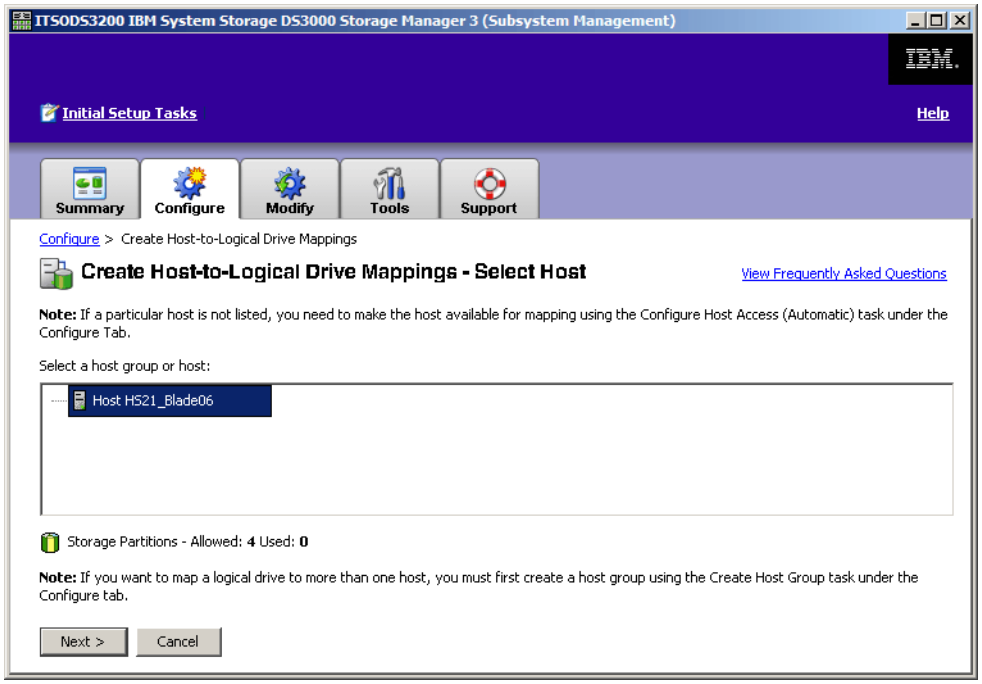


Figure 20-33 Host selection

11. Select the boot volume created earlier (in our example, HS21_boot_LUN), and assign LUN ID = 0 to it, as shown in Figure 20-34. Click **Finish** to complete the step.

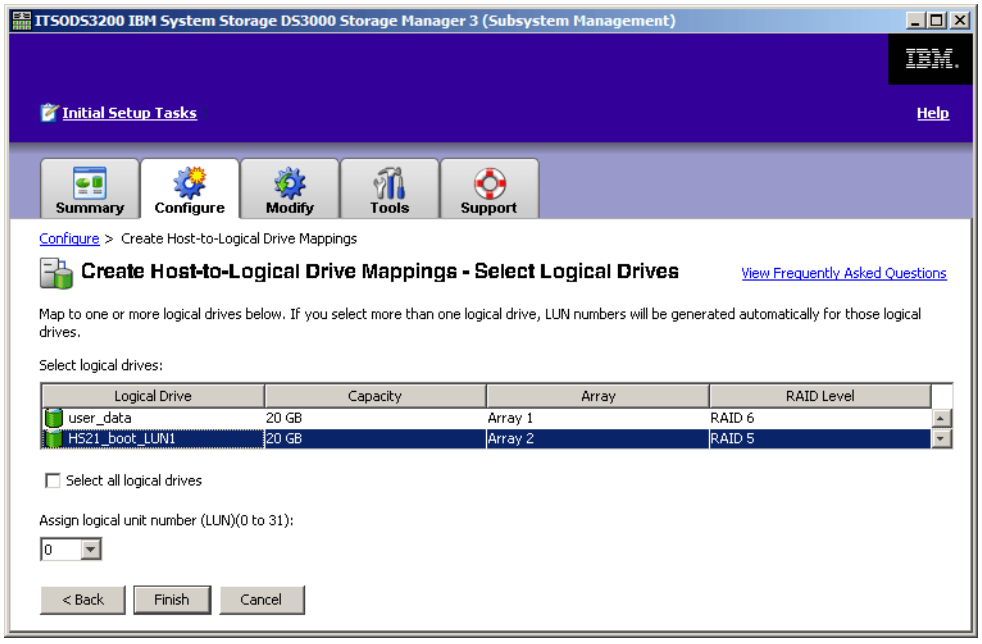


Figure 20-34 Logical volume selection

12. Figure 20-35 confirms the completion of this task. Click **OK** to proceed.

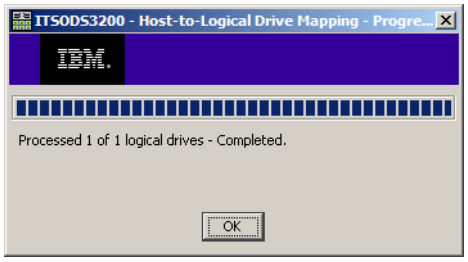


Figure 20-35 Confirmation

13. Select **No** to exit the Create Host-to-Logical Drive Mapping process, as shown in Figure 20-36.

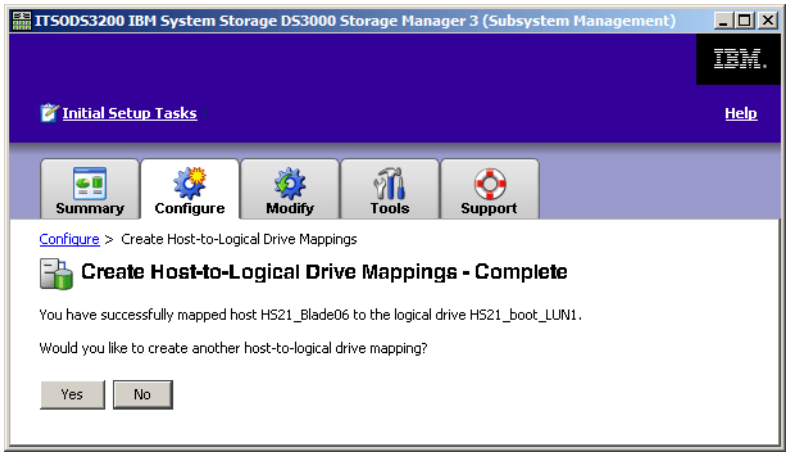


Figure 20-36 Logical Drive mapping completed

The necessary configuration to prepare the blade server to access the boot volume through a single path is complete.

20.3.6 Installing Windows Server 2008

We proceed with the Windows Server 2008 operating system installation steps for the boot logical volume created.

1. Ensure that the boot sequence is correctly setup to boot from CD-ROM as the first device. This can be set through the Web interface of the BladeCenter Advance Management Module (AMM). Select **Blade Tasks** → **Configuration** → **Boot Sequence** (we select **HS21Blade06** for our example). Select **Save**, as shown in Figure 20-37 on page 471.

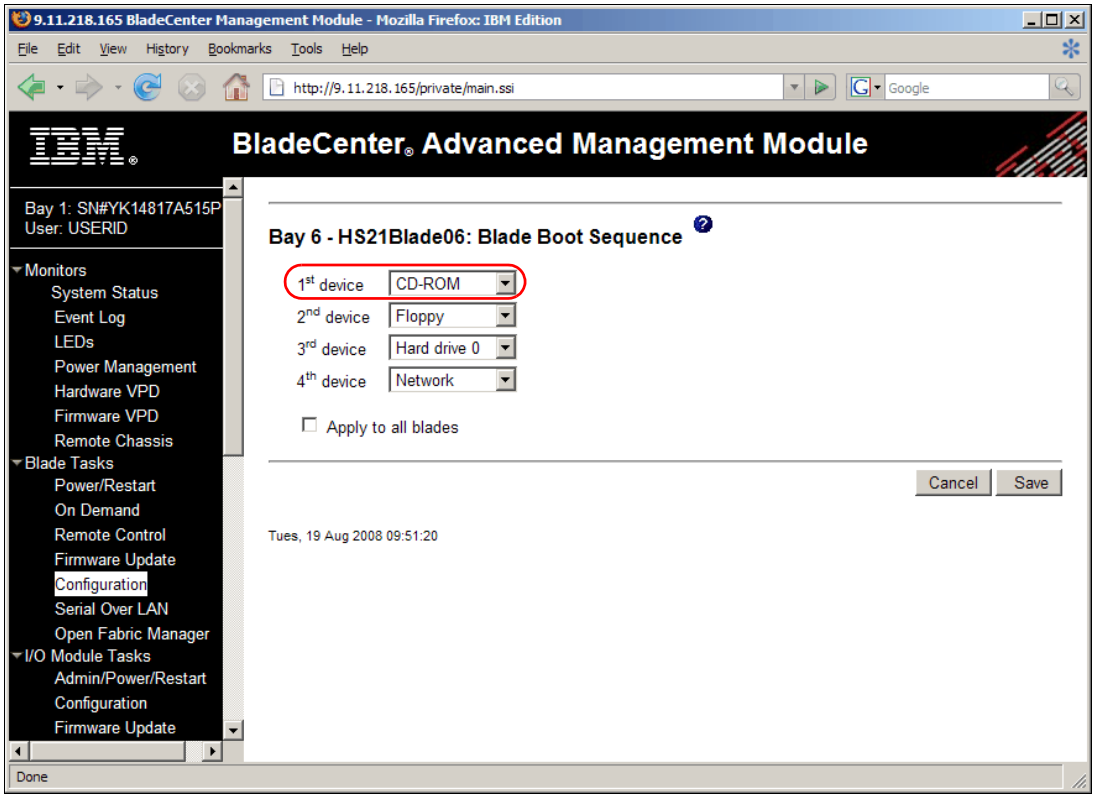


Figure 20-37 Boot device selection

2. Insert the Windows Server 2008 CD in the media tray and select the media tray for the appropriate blade server that is booting from SAN (in our example, Blade06).
3. Power on the HS21 BladeCenter and look at the LSI SAS BIOS, as shown in Figure 20-38. In our example, it shows that the 20 GB LUN with LUN ID=0 on DS3200 1726-2xx FastT is seen by the BladeCenter, which confirms that the boot LUN is visible through a single path.

```
LSI Corporation MPT SAS BIOS
MPTBIOS-6.18.01.00 (2007.08.08)
Copyright 2000-2007 LSI Corporation.


SLOT  ID  LUN VENDOR   PRODUCT          REVISION   INT13 SIZE \ NV
----  --- --- -------- ---------------- ---------- ----- ---------
  1    6  0  IBM      1726-2xx  FAStT  0617        Boot     20 GB
  1          LSILogic SAS1064-IR       1.23.81.00  NV 2B:09
```

Figure 20-38 Observe Boot LUN

4. We followed the operating system installation instructions that are available for the IBM BladeCenter HS21 (8853) with Microsoft Windows Server 2008 setup, found at:

   https://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-5074391&brandind=5000020

The installation guides can be found in the “Install/use” section of each product’s support Web sites.

5. When the installation starts, you will be presented with the window shown in Figure 20-39. It allows you to select the disk to be used for the installation. In our example, the 20 GB boot-LUN created in the DS3200 System Storage subsystem is listed. You have the option to load device drivers for external devices.

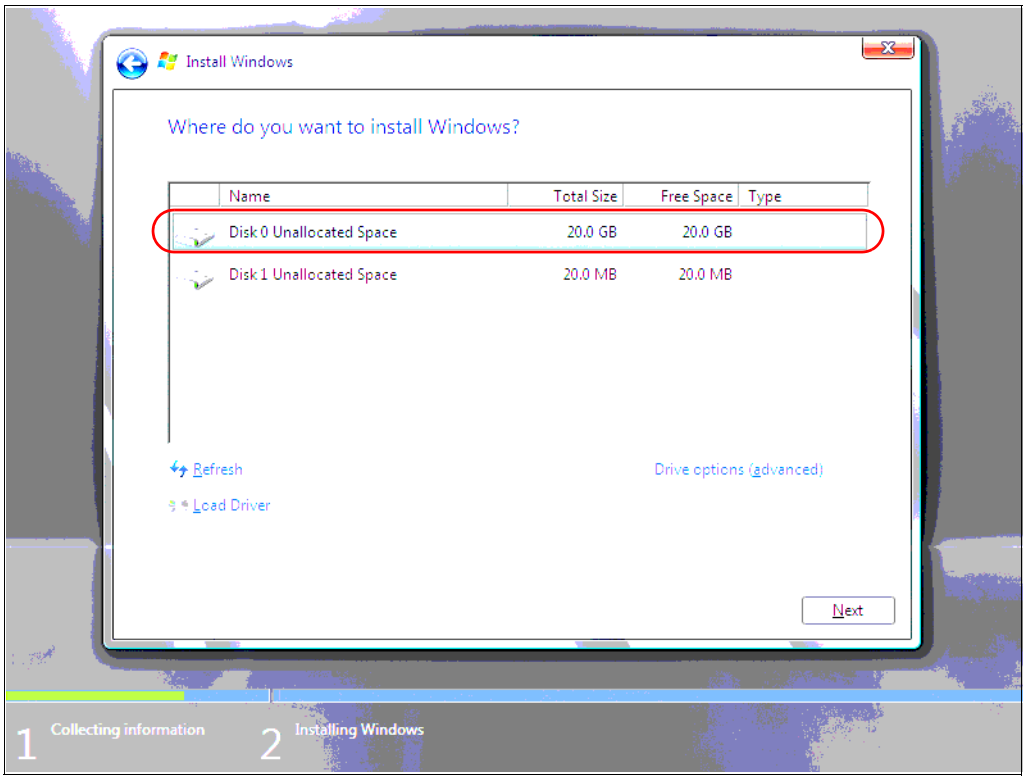


Figure 20-39 Installation disk selection

6. The system will restart after the Windows Server 2008 installation completes. Once Windows boots up, install the Multi-Path failover driver that is part of the DS3000 Storage Manager Host software package (refer to 20.3.7, “Installing DS3000 Storage Manager host software” on page 472).

20.3.7 Installing DS3000 Storage Manager host software

We will install the host server components of DS3000 Storage Manager. These components include:

- SMagent
- SMutil
- Multipath support
- SMclient (optional)

At the time of writing, in-band management is not supported on the DS3200, so we will not install the SMagent. We will be using out-of-band management.

We install SMutil, because it allows us to run the hot_add utility and avoid the necessity of rebooting the host server when adding new logical drives.

Since the host server contains a 2-port SAS Expansion Card, we require multipath support. This is installed as part of the host selection when installing the DS3000 Storage Manager. For exact installation steps, refer to Chapter 6, “Installing IBM System Storage DS3000 Storage Manager” on page 81.

20.3.8 Enabling a failover path for a Windows 2008 boot volume

We now enable the second path to the DS3200 System Storage subsystem to allow failover to a secondary path in the event of a failure.

Do these steps:

1. From the BladeCenter Advance Management Module Web interface GUI, select **I/O Module Tasks** → **Admin/Power/Restart** to ensure that the external ports for the I/O module in bay 4 are enabled, as shown in Figure 20-40. If not, select **Enabled** and click **Save**.

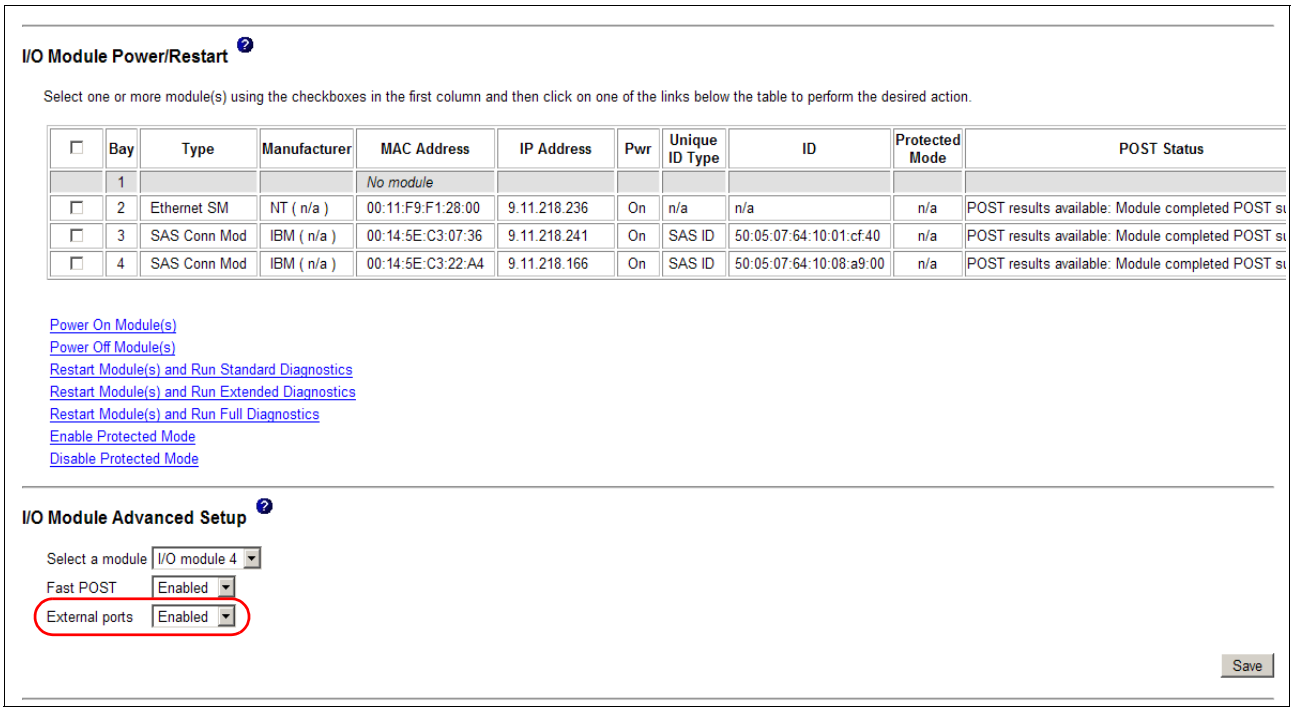


Figure 20-40 SAS Connectivity Module - Enabling external ports

2. Verify that the active zone configuration on the SAS Connectivity Module in I/O module bay 4 includes the host port in blade slot 6 and the external port to DS3200 System Storage subsystem controller B. This is done by logging on to the Web interface of the SAS Connectivity Module in bay 4 and viewing the “Predefined Config 01” zone configuration, as shown in Figure 20-41.

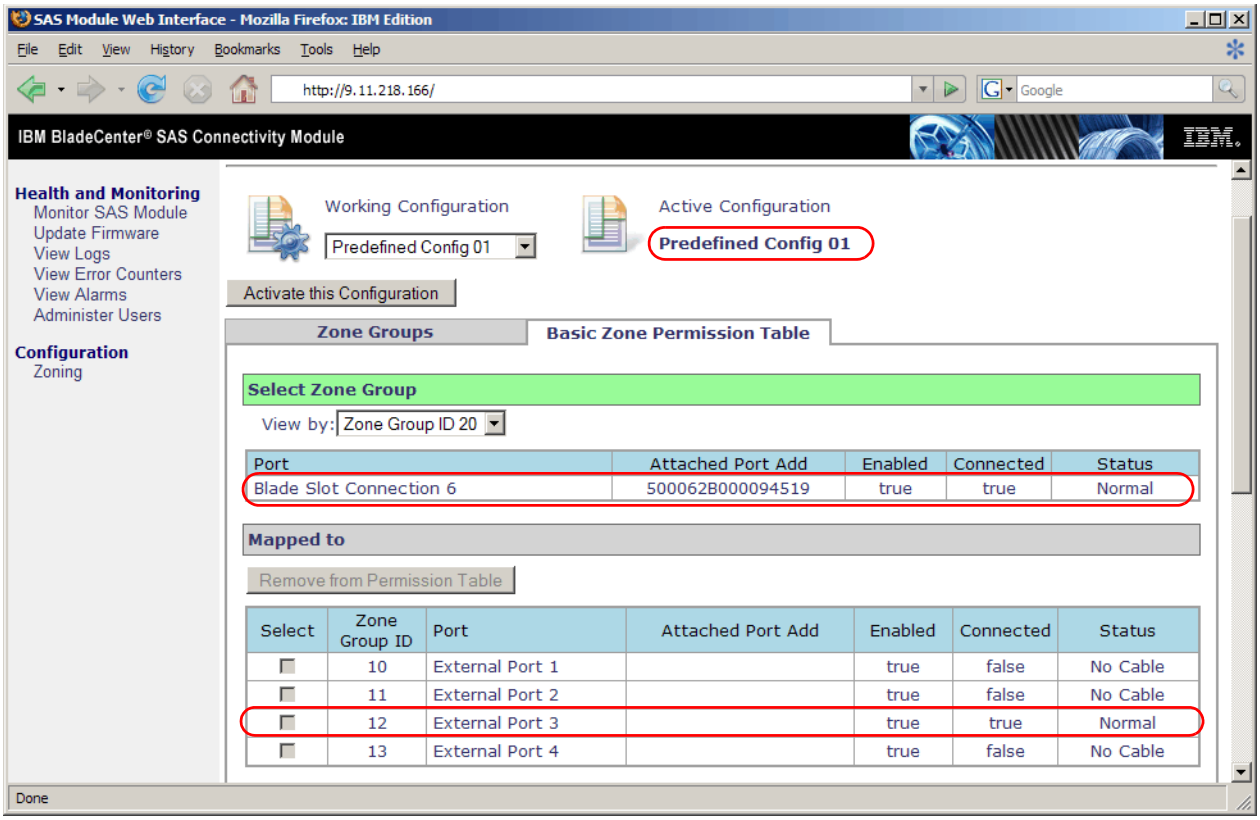


Figure 20-41 SAS Connectivity Module - Confirm zone configuration

3. In the DS3200 Storage Manager Subsystem Management window, on the Modify tab, select **Edit Host Topology**, as shown in Figure 20-41.

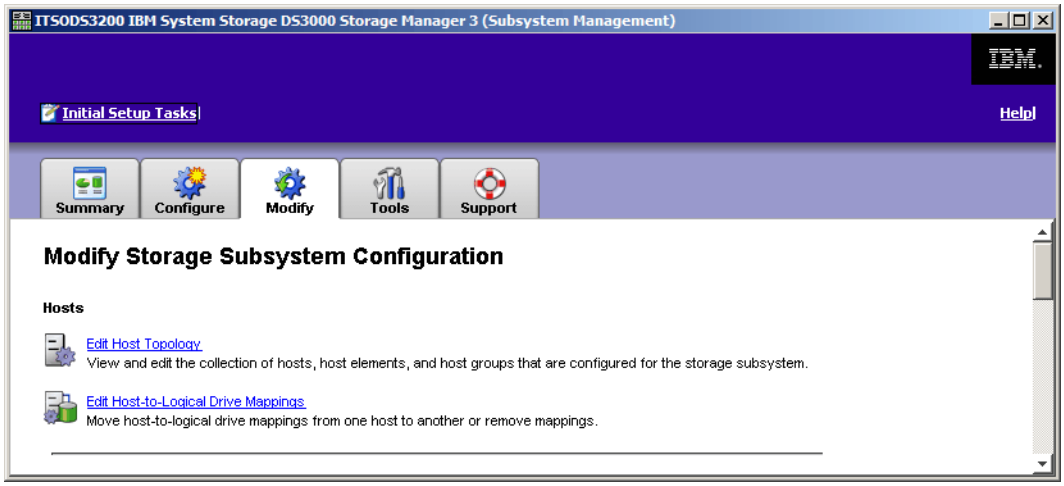


Figure 20-42 Edit host topology

4. Select the host **HS21_Blad06** and click **Add HBA...**, as shown in Figure 20-43 on page 475.

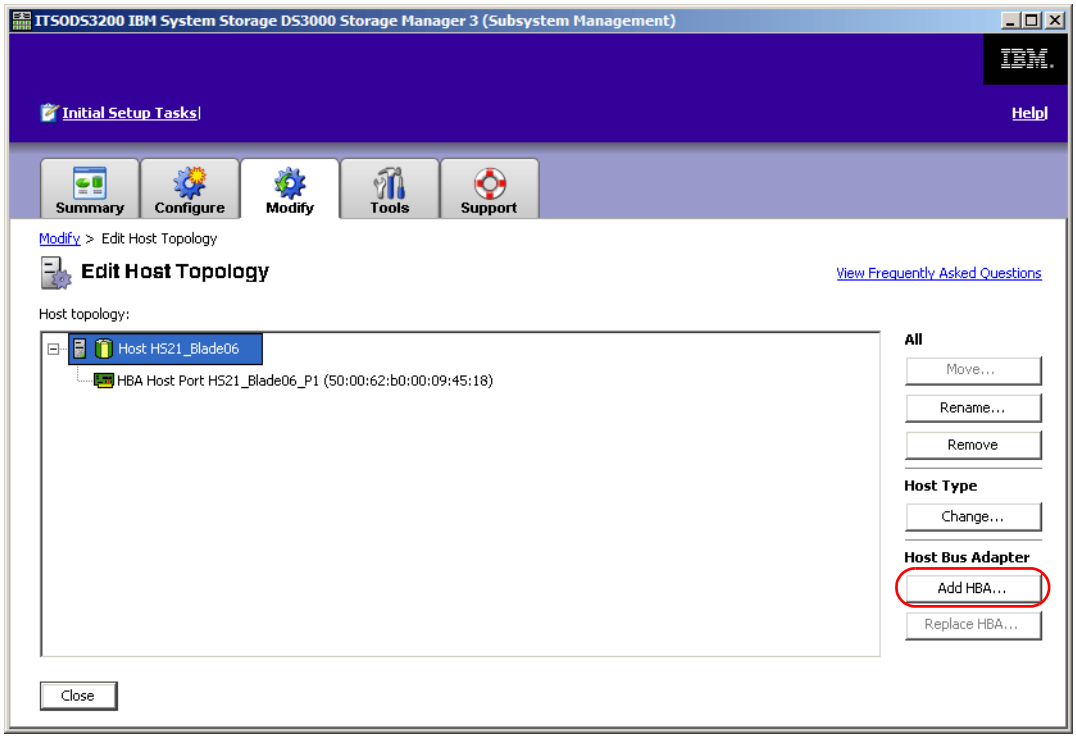


Figure 20-43 Add HBA

5. Select the second HBA port of the SAS Expansion Card from the drop-down list and specify an alias for the HBA, as shown in Figure 20-44. Click **Add** to continue.

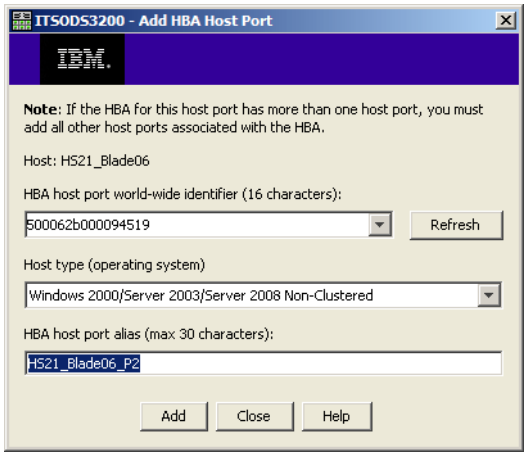


Figure 20-44 Selecting HBA

6. Figure 20-45 lists the HBA WWPN ports allocated to the host.

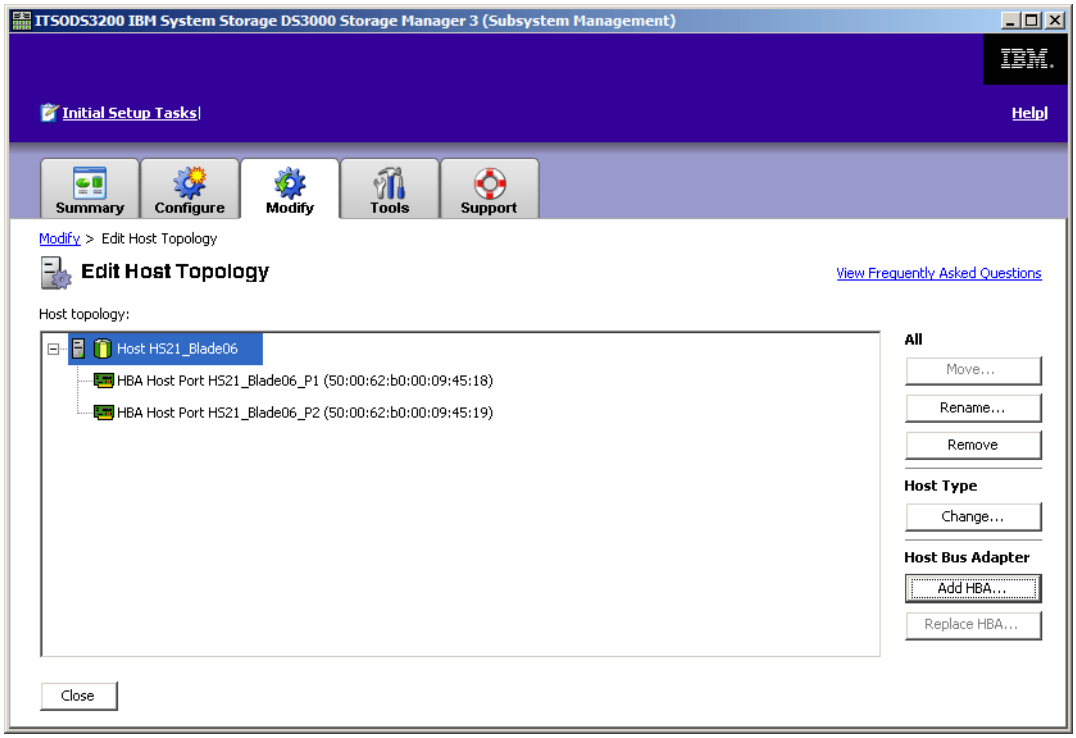


Figure 20-45 Host and HBA listing

7. Click **Close** to exit the Edit Host Topology section and complete the step.
8. Reboot the HS21 BladeCenter host to ensure that the Multi-path driver is loaded and the second path is activated.
9. Once Windows Server 2008 is booted up, verify in the Windows 2008 Disk Management window that only a single disk is seen as the boot volume, in our example, a 20 GB disk, as shown in Figure 20-46 on page 477. This confirms that the Multi-Path driver is installed successfully.

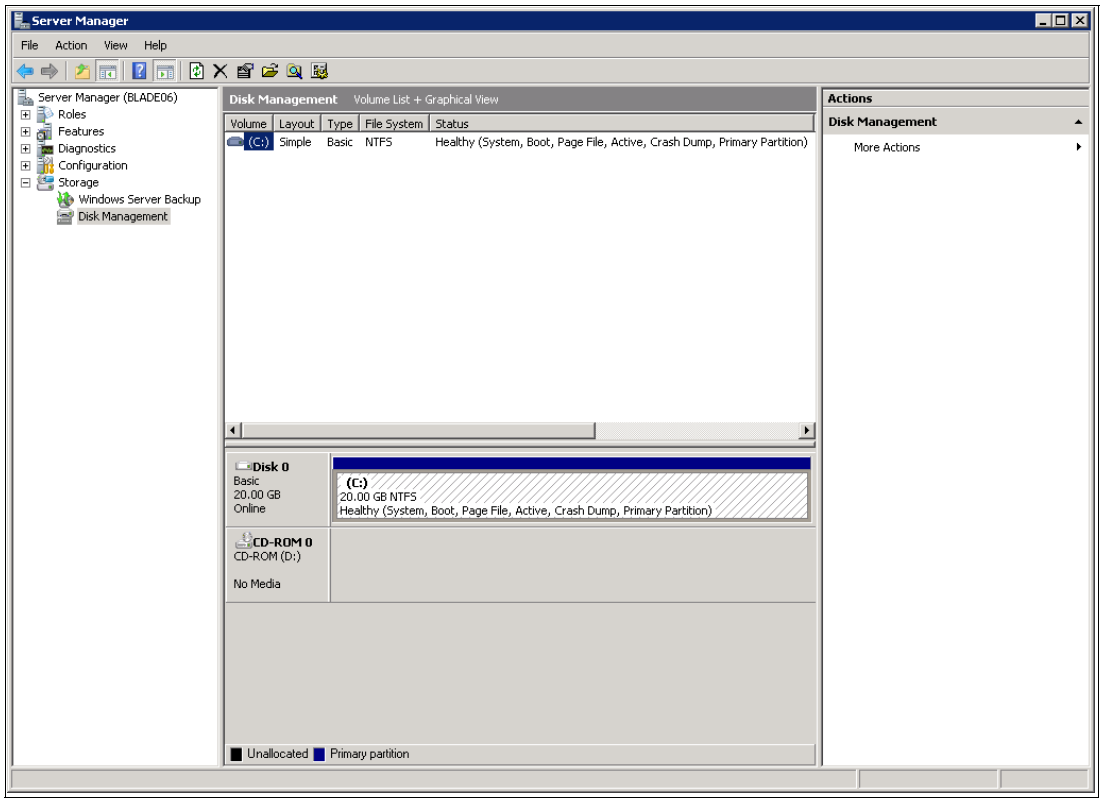


Figure 20-46 Confirming single boot volume

10. In the Windows Device Manager, we can confirm the IBM 1726-2xx Multi-Path disk device as well as the Multi-Path Bus Driver, as shown in Figure 20-47.

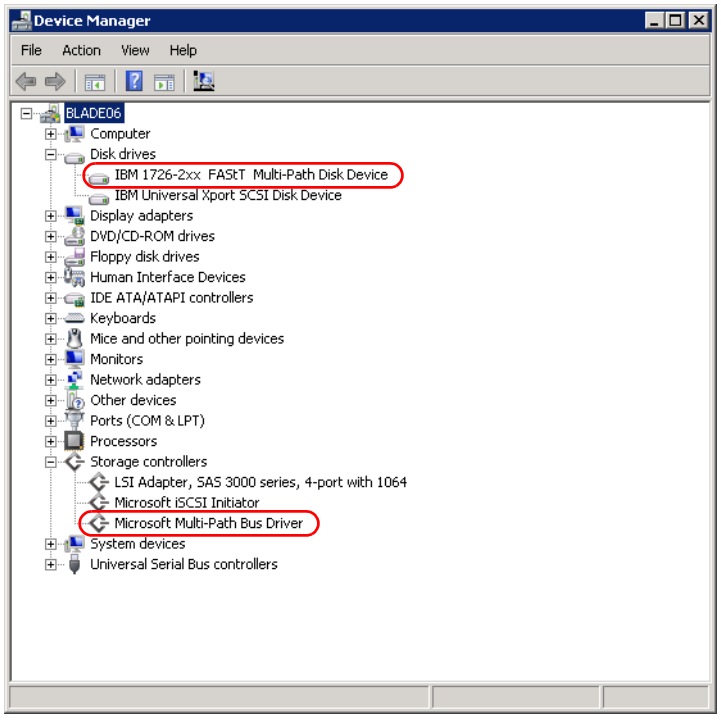


Figure 20-47 Confirming Multi-Path driver

20.3.9 Failover testing

1. Verify the current owner of the boot volume (in our example, Controller A) by viewing the logical drives section under the Subsystem Profile in the DS3200 Storage Subsystem Manager, as shown in Figure 20-48. Select **Storage Manager** → **Subsystem Profile** → **Logical Drives**.

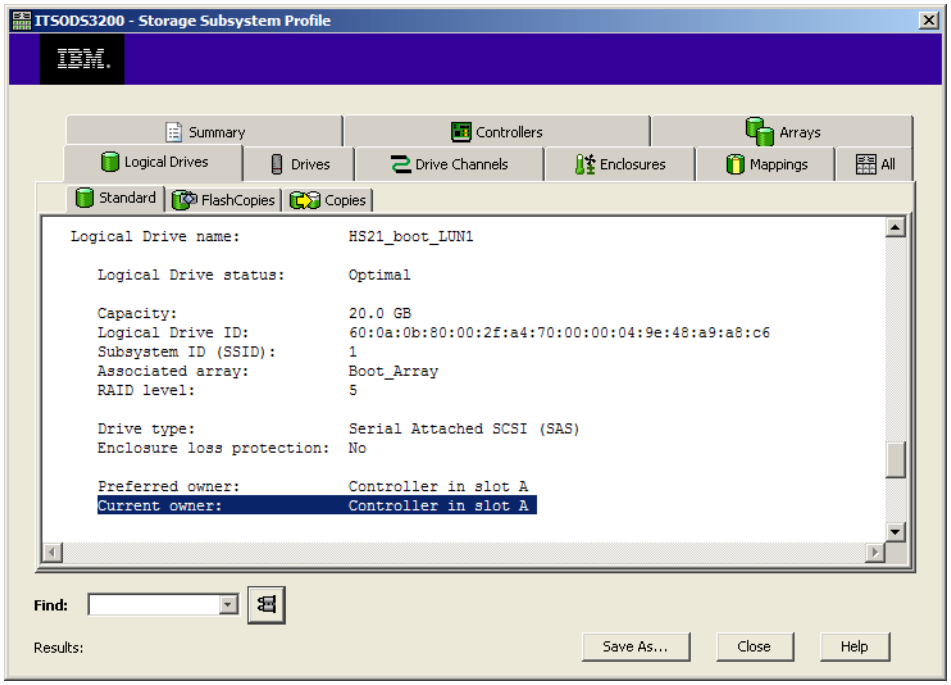


Figure 20-48 Boot volume owner - Controller A

2. We test the path failover by unplugging the cable connected to the BladeCenter I/O module bay 3 and controller A of the DS3200 System Storage subsystem. You can also disable the SAS HBA port connected to Controller A on the SAS Expansion Module itself to perform this test.
3. After the cable is unplugged, the Windows 2008 operating system will pause until the boot volume is failed over to the backup path, in our example, Controller B, as shown in Figure 20-49 on page 479.

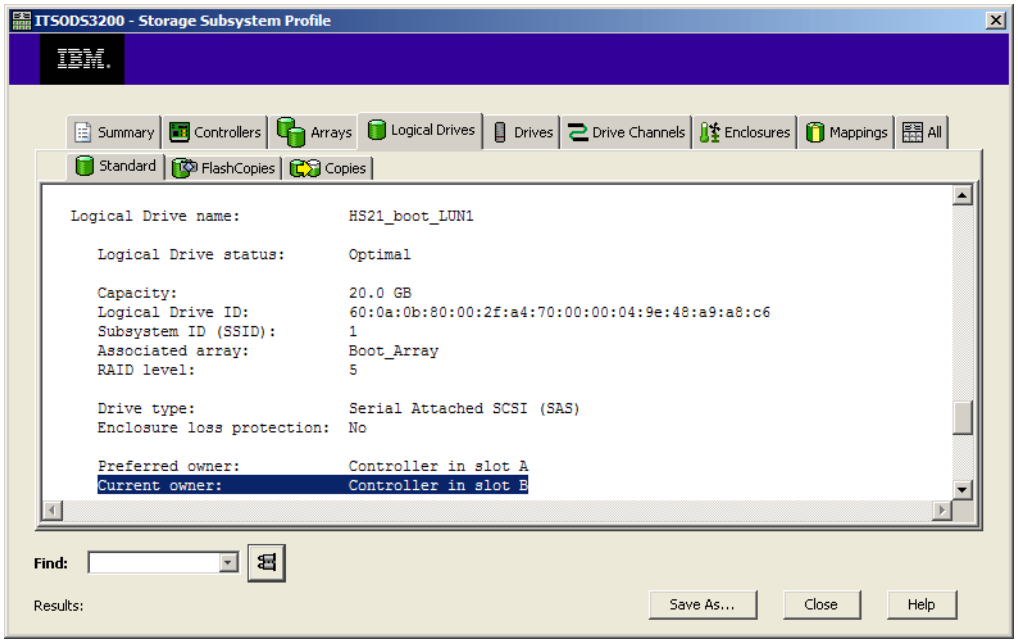


Figure 20-49 Boot volume owner - Controller B

A Recovery Guru alert will also be presented in the DS3200 Subsystem Management window, as shown in Figure 20-50, alerting you that the boot volume has moved ownership to controller B and is not on the preferred path.

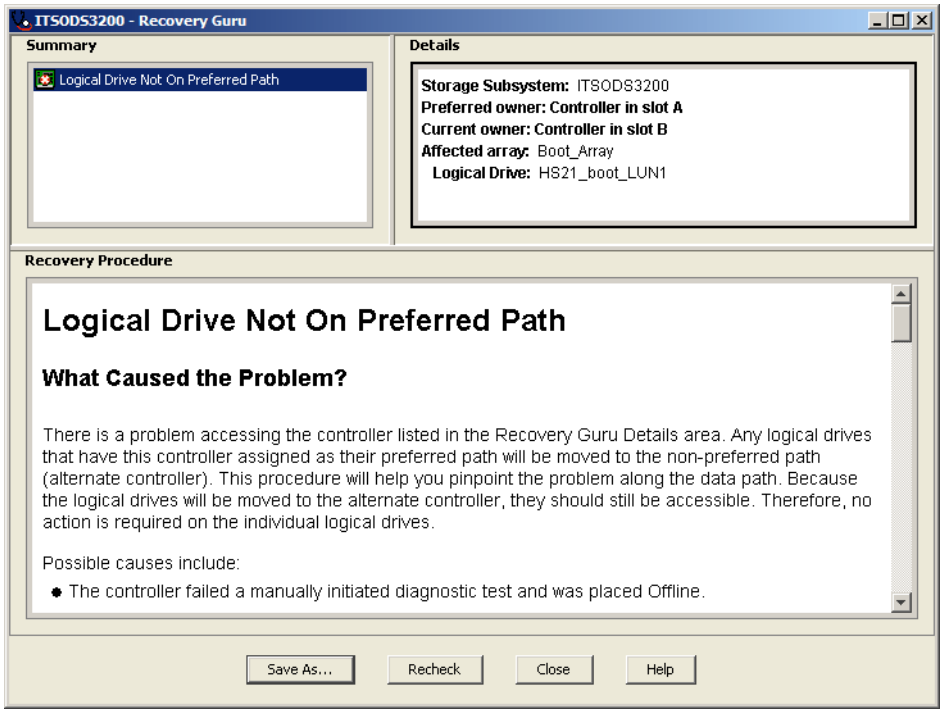


Figure 20-50 Recovery Guru alert

4. After verifying the successful failover, replug the cable between port #3 of the SAS Connectivity Module in I/O Module bay 3 and controller A of the DS3200 System Storage subsystem. The boot volume will now automatically move back to the preferred controller, (in our example, controller A) on the DS3200 System Storage subsystem.

21

iSCSI configuration 1 - Windows 2003 and 2008 with a QLogic QMC4052 HBA

In this chapter, we discuss a sample configuration that shows how to connect logical drives configured on an IBM System Storage DS3300 to a Windows Server 2003 running on an IBM BladeCenter HS21 with a QLogic iSCSI Expansion Card for BladeCenter. The majority of the setup tasks are related to Windows Server 2003, but references are also provided for setup with Windows Server 2008.

To be able to use the QLogic iSCSI Expansion Card for BladeCenter in an IBM BladeCenter HS21, a BladeCenter Storage IO Expansion 3 (BSE3) was used that allows standard and small form-factor daughter cards to be installed in addition to hard drives. The iSCSI HBA was installed in the BSE3 in the slot that connects to module bays 3 and 4 of the BladeCenter chassis. The decision to use module bays 3 and 4 was made to have dedicated network equipment for the iSCSI traffic. That increases network stability and provides a safer SCSI connection later to the host.

Nortel Layer 2/3 Gigabit Ethernet switches are installed in module Bays 1, 3, and 4. Bay 1, which connects to the first onboard NIC of the blades, and is used for the public network access. Bay 3 and 4, which connect to an expansion slot on the blade and the BSE3, are used to connect to the installed iSCSI HBA, as well as in a later sample configuration to a regular network card to handle iSCSI traffic over a software initiator instead of hardware initiation.

Note: iSCSI boot from SAN is not covered at this time. Configuration information for iSCSI boot from SAN for IBM BladeCenter HS21 can be found in the *iSCSI Boot from SAN Installation and Setup Guide*. This guide can be found at the following URL:

https://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-64880&brandind=5000020

© Copyright IBM Corp. 2007, 2009. All rights reserved.

A basic network diagram that provides an overview of the cabling between the components is shown in Figure 21-1 on page 483.

21.1 Network configuration

A diagram of the network setup used is shown in Figure 21-1.

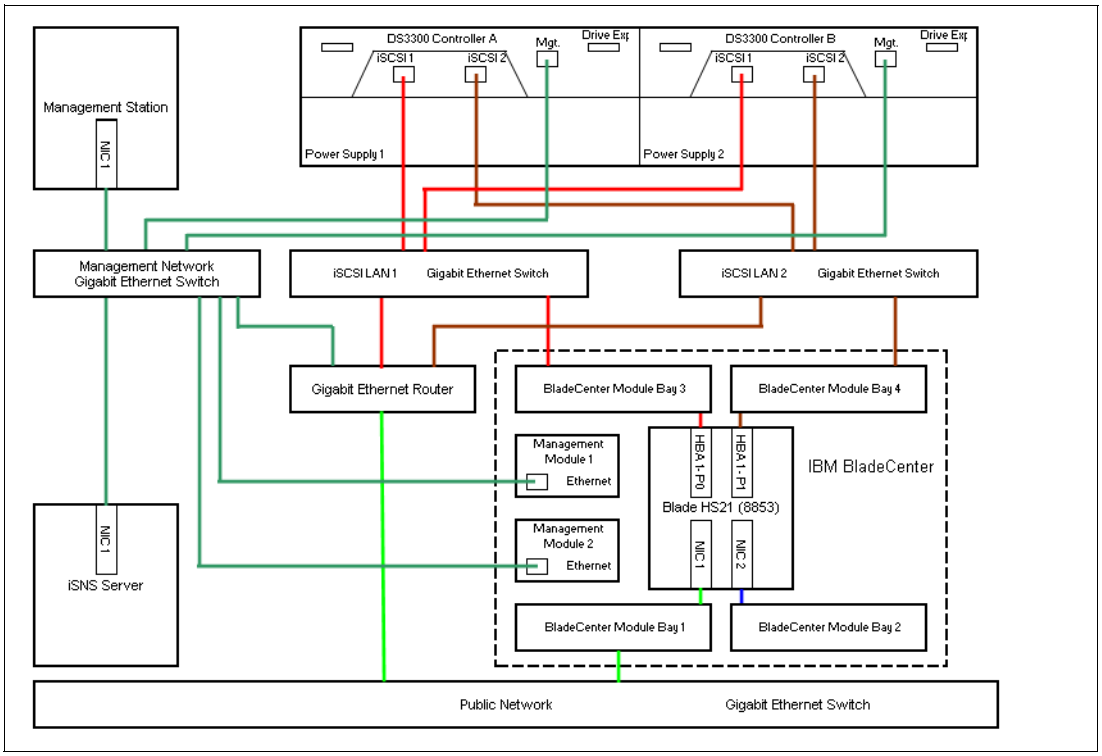


Figure 21-1 Network diagram

Table 21-1 shows a list of network addresses used in this sample configuration.

Table 21-1 Network configuration details

Device	Description	Value	Comments
BladeCenter	Management interface	172.18.0.125/16	Management LAN
	Nortel Layer 2/3 Gigabit Ethernet Switch Bay 1	172.18.0.127/16	
	Nortel Layer 2/3 Gigabit Ethernet Switch Bay 2	172.18.0.128/16	
	Nortel Layer 2/3 Gigabit Ethernet Switch Bay 3	172.18.0.129/16	
	Nortel Layer 2/3 Gigabit Ethernet Switch Bay 4	172.18.0.130/16	
Blade HS21 (8853)	Host name	hudson.rivers.local	
	onboard NIC 1	172.19.200.60/16	

Device	Description	Value	Comments
	iSCSI HBA 1 - NIC Port 0	172.16.200.62/16	
	iSCSI HBA 1 - NIC Port 1	172.17.200.62/16	
	iSCSI HBA 1 -iSCSI Port 0	172.16.200.63/16	
	iSCSI HBA 1 -iSCSI Port 1	172.17.200.63/16	
System Storage DS3300	Controller A - Mgt.	172.18.3.1/16	Management LAN
	Controller A - iSCSI Port 1	172.16.3.1/16	iSCSI LAN 1
	Controller A - iSCSI Port 2	172.17.3.1/16	iSCSI LAN 2
	Controller B - Mgt.	172.18.3.2/16	Management LAN
	Controller B- iSCSI Port 1	172.16.3.2/16	iSCSI LAN 1
	Controller B - iSCSI Port 2	172.17.3.2/16	iSCSI LAN 2
iSNS	Service IP	172.18.0.9/16	Management LAN

21.2 Installing Windows Server 2003

Follow the operating system installation instructions that are available for each IBM BladeCenter and IBM System x server. The installation guides can be found in the “Install/use” section of each product’s support Web sites.

For our IBM BladeCenter HS21 (8853) with Microsoft Windows Server 2003, we used the instructions found at:

http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-66309&brandind=5000020

21.3 Installing Windows Server 2008

Follow the operating system installation instructions that are available for each IBM BladeCenter and IBM System x server. The installation guides can be found in the “Install/use” section of each product’s support Web sites.

For the IBM BladeCenter HS21 (8853) with Microsoft Windows Server 2008 setup, use the instructions found at:

https://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-5074391&brandind=5000020

21.4 Installing the iSCSI HBA driver

Microsoft Windows Server 2003 and Windows Server 2008 does not include the device driver for the QLogic iSCSI HBA. In this section, we will show how to install this driver.

21.4.1 Preparation

Open Device Manager by selecting **Start** → **Programs** → **Administrative Tools** → **Computer Management.** Click **Device Manager**. You should see two unknown Ethernet Controllers and two Network Controllers under Other devices, as shown in Figure 21-2.

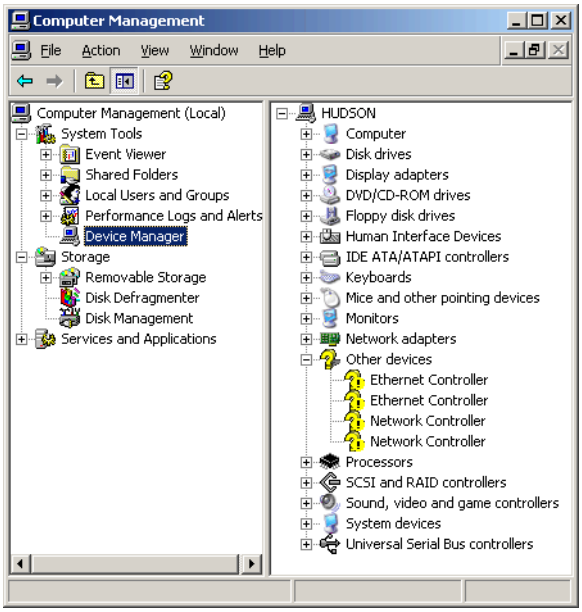


Figure 21-2 Device Manager

Right-click each one of these devices and select **Properties**. The Properties window for each of those four unknown devices shows that they all belong to the same device, representing different functions. In this configuration, these are:

- PCI Slot 6 on PCI Bus 13, device 1 and function 0 for the first Ethernet Controller
- PCI Slot 6 on PCI Bus 13, device 1 and function 2 for the second Ethernet Controller
- PCI Slot 6 on PCI Bus 13, device 1 and function 1 for the first Network Controller
- PCI Slot 6 on PCI Bus 13, device 1 and function 3 for the second Network Controller

The QLogic iSCSI Expansion Card for BladeCenter is a dual port adapter. Each port has two different functions that are visible here as Ethernet Controller and Network Controller. Port 0 has function 0 and 1 and port 1 has function 2 and 3. The Ethernet Controller Function is a regular Ethernet interface that can be used just like any other Ethernet NIC. This interface does not benefit from the offload capabilities of this adapter.

The Network Controller function that each of the ports owns is an iSCSI port.

Note: When both functions of a port are used, different IP addresses must be assigned.

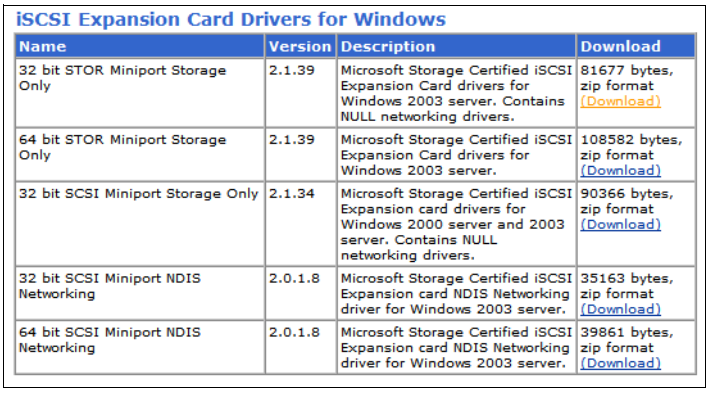
The Ethernet controller requires an NDIS driver and the Network Controller requires an iSCSI driver. These driver packages for Windows Server 2003 can be downloaded from:

http://support.qlogic.com/support/oem_detail_all.asp?oemid=369

Attention: To support Storport miniport device drivers, Service Pack 2 and the latest Storport miniport hotfix KB932755 must be installed in the Windows Server 2003 operating system. You can download the latest Storport hotfix from:

http://www.microsoft.com/downloads/details.aspx?FamilyID=dead950b-da47-4dc2-abcd-40f61bd4d1aa&DisplayLang=en

At this Web site, partially shown in Figure 21-3, you will find packages containing the iSCSI driver but no NDIS driver (Storage only) and packages that contain both drivers. In our configuration, we use the driver package without the NDIS driver because we do not need the NDIS function to connect to a DS3300.



iSCSI Expansion Card Drivers for Windows

Name	Version	Description	Download
32 bit STOR Miniport Storage Only	2.1.39	Microsoft Storage Certified iSCSI Expansion Card drivers for Windows 2003 server. Contains NULL networking drivers.	81677 bytes, zip format (Download)
64 bit STOR Miniport Storage Only	2.1.39	Microsoft Storage Certified iSCSI Expansion Card drivers for Windows 2003 server.	108582 bytes, zip format (Download)
32 bit SCSI Miniport Storage Only	2.1.34	Microsoft Storage Certified iSCSI Expansion card drivers for Windows 2000 server and 2003 server. Contains NULL networking drivers.	90366 bytes, zip format (Download)
32 bit SCSI Miniport NDIS Networking	2.0.1.8	Microsoft Storage Certified iSCSI Expansion card NDIS Networking driver for Windows 2003 server.	35163 bytes, zip format (Download)
64 bit SCSI Miniport NDIS Networking	2.0.1.8	Microsoft Storage Certified iSCSI Expansion card NDIS Networking driver for Windows 2003 server.	39861 bytes, zip format (Download)

Figure 21-3 QLogic iSCSI Expansion Card downloadable drivers for Windows Server 2003

To prevent unknown devices from being left in Device Manager, the Storage only driver installs a *Null* driver for the Ethernet Controllers. These controllers will then appear as System devices after the driver package is installed, as shown in Figure 21-8 on page 489.

STOR Miniport driver packages for Windows Server 2008 (32-bit) covering general hardware support can be downloaded from:

http://driverdownloads.qlogic.com/QLogicDriverDownloads_UI/SearchByOs.aspx?ProductCategory=82&OsCategory=1&Os=172&OsCategoryName=Windows&ProductCategoryName=iSCSI%20HBAs&OSName=Windows%20Server%202008%20(32-bit)

21.4.2 Installing the driver

Follow these steps to install the driver:

1. Download the driver for the iSCSI HBA for Windows Server 2003 from:

   http://support.qlogic.com/support/oem_detail_all.asp?oemid=369

   Download the driver for the iSCSI HBA for Windows Server 2008 (32-bit) covering general hardware support from:

   http://driverdownloads.qlogic.com/QLogicDriverDownloads_UI/SearchByOs.aspx?ProductCategory=82&OsCategory=1&Os=172&OsCategoryName=Windows&ProductCategoryName=iSCSI%20HBAs&OSName=Windows%20Server%202008%20(32-bit)

Extract the archive into a local directory. Be sure that the STOR miniport driver is used and not the SCSI miniport driver.

2. Right-click the first Network Controller in Device Manager and select **Update Driver....** from the context menu, as shown in Figure 21-4.

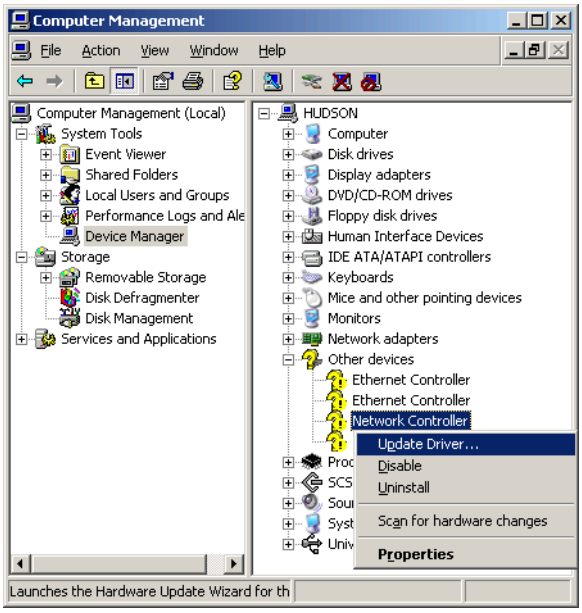


Figure 21-4 Device Manager - Update driver

3. Select **Install from a list or specific location (Advanced)** and click **Next**. See Figure 21-5.

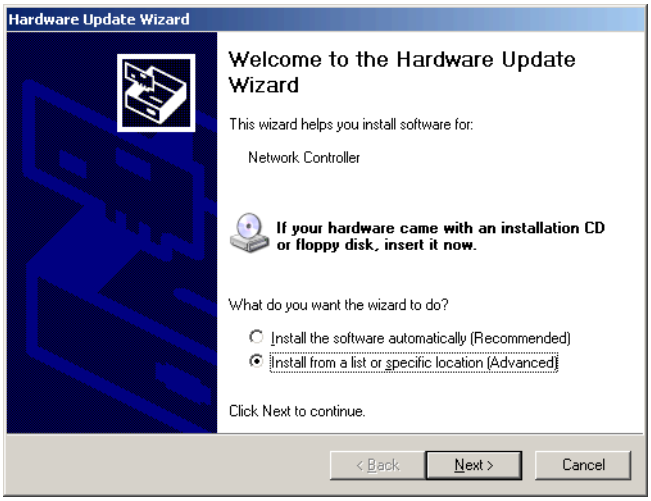


Figure 21-5 Hardware Update Wizard - Welcome

4. Uncheck **Search removable media (floppy, CDROM...)** and check **Include this location in the search**. Enter the location where the driver was unpacked in the previous step. Click **Next**. See Figure 21-6.

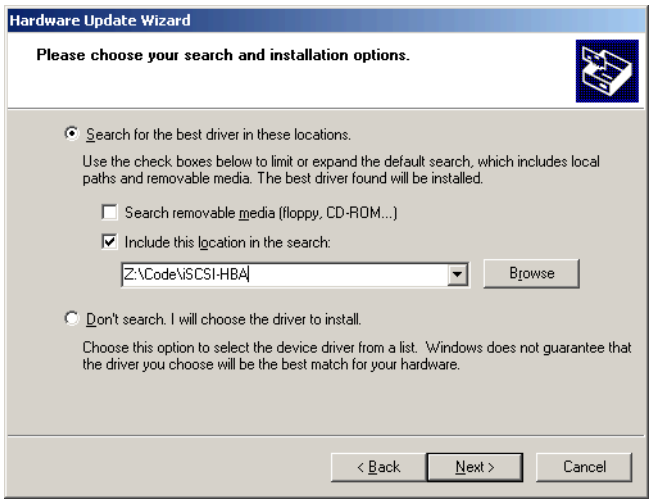


Figure 21-6 Hardware Update Wizard - Search and install options

5. The driver installs now. When you see the completion window, click **Finish** to close the Hardware Update wizard.
6. Repeat step 2 on page 487 to step 5 for the second Network Controller in Device Manager.
7. Both Network Controllers are now installed and can be found in the SCSI and RAID controllers section of Device Manager as a QLogic iSCSI Adapter. See Figure 21-7.

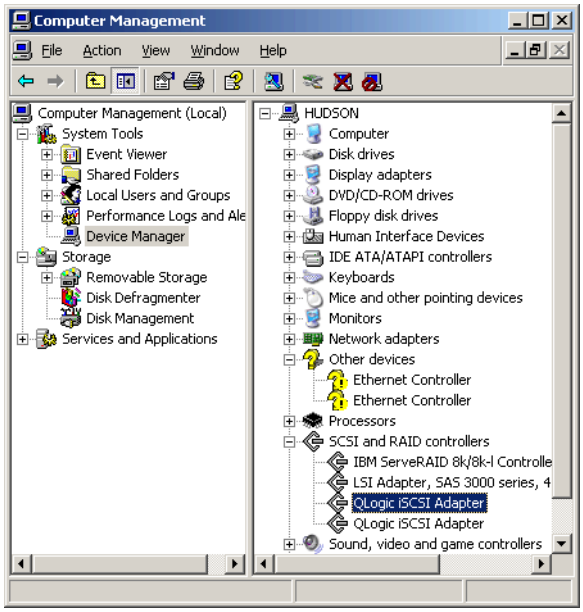


Figure 21-7 Device Manager - Installed QLogic iSCSI Adapter

8. The NULL driver for the Ethernet controller is installed in the same way: Repeat step 2 on page 487 to step 5 for both Ethernet Controllers in Device Manager. You will then see the two Ethernet Controllers listed under System Devices (QLogic NULL Driver), as shown in Figure 21-8 on page 489.

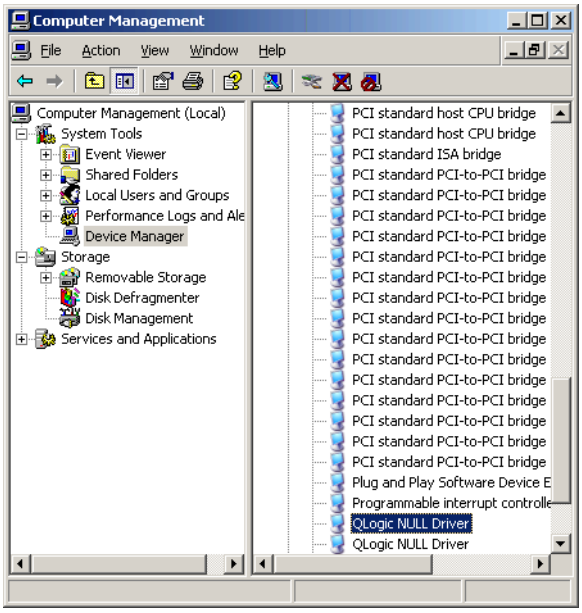


Figure 21-8 Device Manager with QLogic NULL Driver

Driver installation is now complete. Next, we install the HBA management application to configure the iSCSI HBA.

21.5 Installing SANsurfer

QLogic SANsurfer is a client-server based management application for QLogic Fibre Channel and iSCSI HBAs as well as QLogic Fibre Channel Switches. Only one client is required: this can be installed either on a separate management station or on one of the hosts with an HBA installed. Every host with an HBA installed requires an agent to communicate with the client and the installed HBA drivers.

There are different SANsurfer packages available that provide support for the various Fibre Channel devices. Here, since we are only using the iSCSI HBA, we download the package that contains only iSCSI support. Follow these steps:

1. Download SANsurfer for Windows Server 2003 from:

   `http://support.qlogic.com/support/oem_detail_all.asp?oemid=369`

   Download SANsurfer for Windows Server 2008 from:

   `http://driverdownloads.qlogic.com/QLogicDriverDownloads_UI/SearchByOs.aspx?ProductCategory=82&OsCategory=1&Os=172&OsCategoryName=Windows&ProductCategoryName=iSCSI%20HBAs&OSName=Windows%20Server%202008%20(32-bit)`

2. Execute the self-extracting archive package iSCSI_SANsurfer_4_03_17_windows_package.exe. Click **Run** to start the execution, as shown in Figure 21-9.



Figure 21-9 Open File- Run

3. Click **Unzip** to extract the content of the archive into the suggested location (Figure 21-10). The setup will start automatically after extraction is complete.

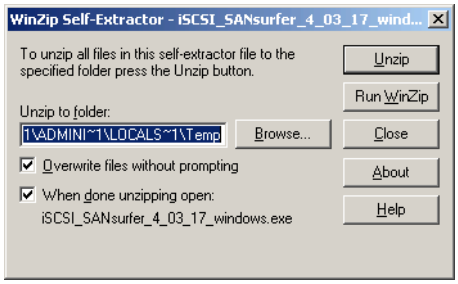


Figure 21-10 WinZip Self-Extractor

4. Click **OK** to confirm that all files are extracted (Figure 21-11).



Figure 21-11 WinZip complete

5. SANsurfer starts to install automatically, as shown in Figure 21-12 on page 491. Click **Next**.

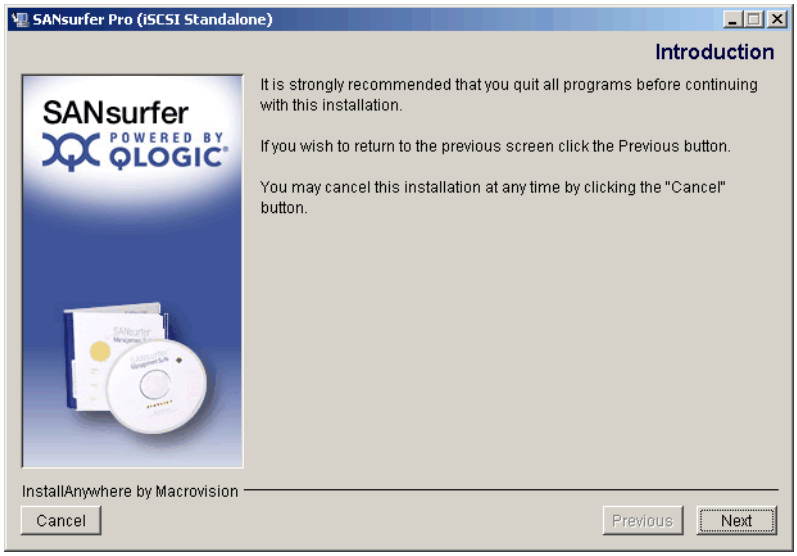


Figure 21-12 SANsurfer installation - Introduction

6. Read the important information about SANsurfer and click **Next.** See Figure 21-13.

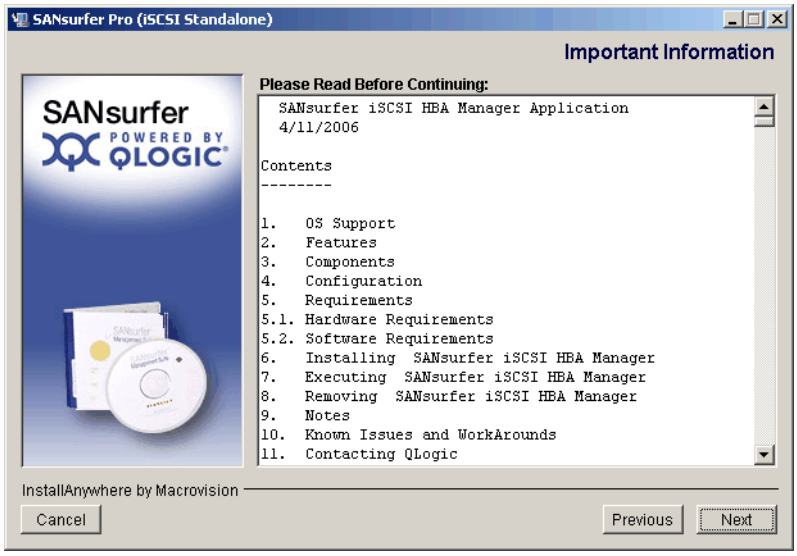


Figure 21-13 SANsurfer installation - Important information

7. Choose the installation type in Figure 21-14. To manage the devices directly from the host, install the iSCSI host bus adapter GUI and agent. This is the default selection. Click **Next**. Alternatively, you could choose to install only the agent in the host and the GUI on a separate management station.



Figure 21-14 SANsurfer installation - Choose product features

8. Choose the installation folder and click **Next**. See Figure 21-15.

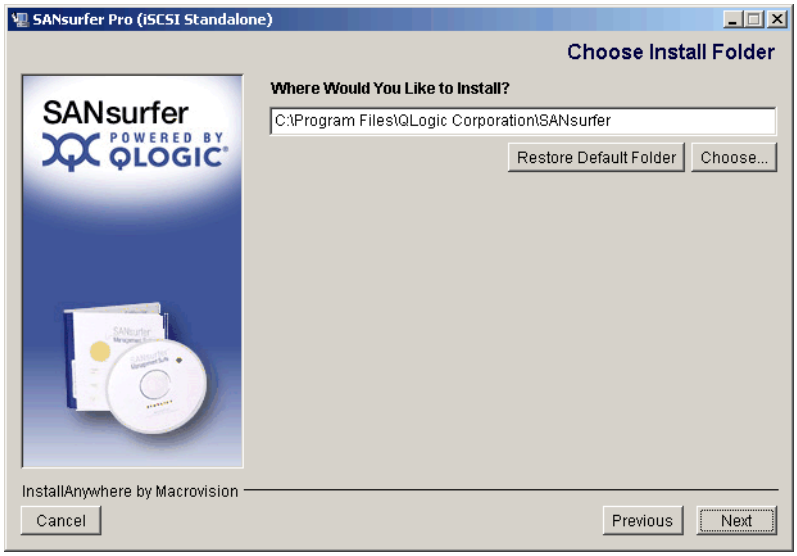


Figure 21-15 SANsurfer installation - Choose installation folder

9. In Figure 21-16, choose a profile to install the program shortcuts. In this example, we use All Users Profile.

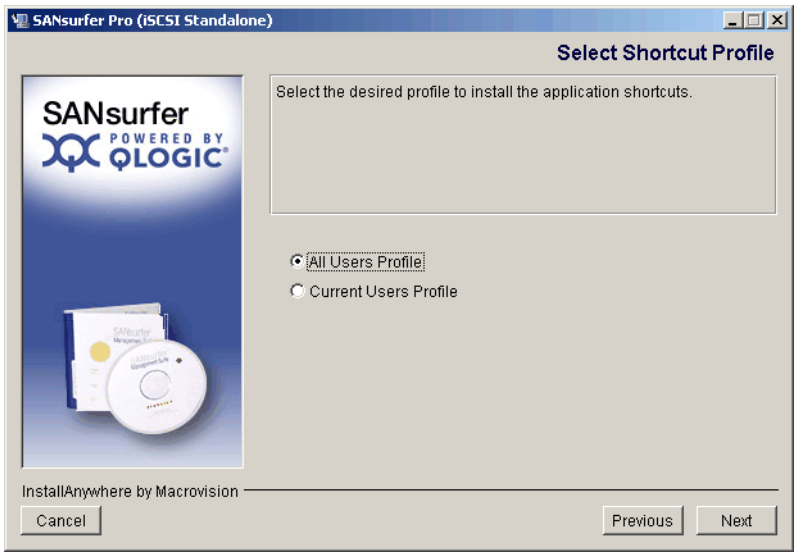


Figure 21-16 SANsurfer installation - Select shortcut profile

10. Click **Next** to create an icon for the program on the desktop (Figure 21-17).

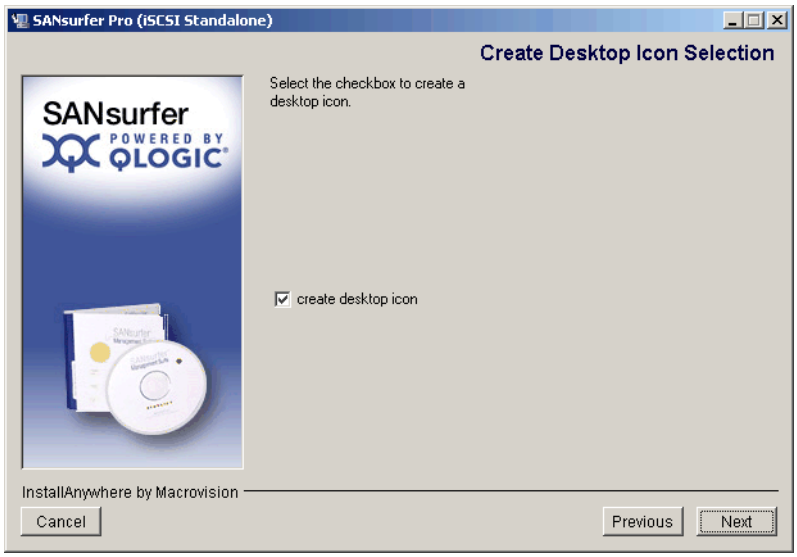


Figure 21-17 SANsurfer installation - Create desktop icon

11. Verify the pre-installation summary shown in Figure 21-18, and click **Install**.

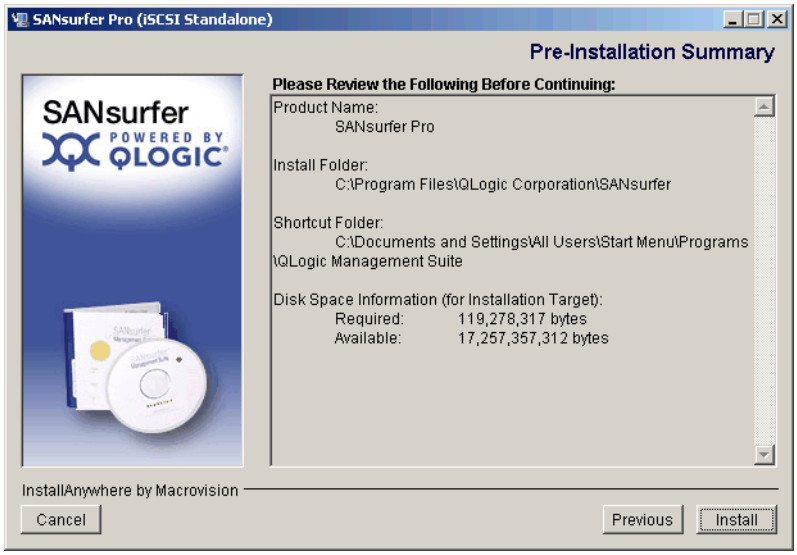


Figure 21-18 SANsurfer installation - Pre-installation summary

12. After installation has completed, click **Done (**Figure 21-19).

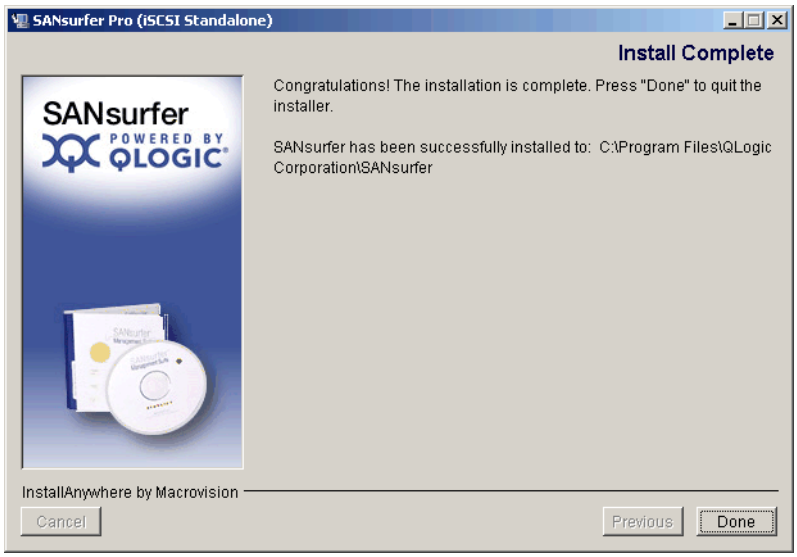


Figure 21-19 SANsurfer installation -Installation complete

SANsurfer is installed; we can now configure the iSCSI HBA.

21.6 Configuring the iSCSI HBA network settings with SANsurfer

The basic configuration of the iSCSI HBA consists of the following tasks:

- Connect SANsurfer to the host with the HBA installed.
- Gather data about the HBA file:
  - BIOS and firmware code level
  - iSCSI qualified name
  - MAC address, in case DHCP is planned to be used for address assignment
- Flash the HBA with the latest supported code level supported by the storage product.
- Configure the network settings of the iSCSI HBA ports.

21.6.1 Connecting with SANsurfer to the host where the HBA is installed

Use a SANsurfer console in the host to start SANsurfer:

1. Select **Start** → **Programs** → **QLogic Management Suite** → **SANsurfer** or click the icon on the desktop.
2. SANsurfer starts and by default connects to the local host, when the GUI and agent are both installed in the host where the HBA is installed. When the GUI is installed on a separate management station, specify the DNS name of the host in the field Enter host name or IP address. Click **Connect** in the Connect to Host dialog. See Figure 21-20.

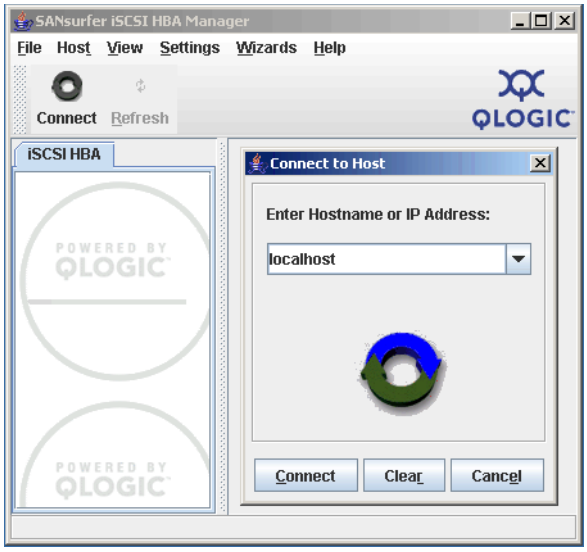


Figure 21-20 SANsurfer - Connect to host

3. Since our iSCSI adapter has not yet been configured, a message appears that an unconfigured port was detected. Check **Do not show this dialog again for this session** and click **No**. See Figure 21-21.

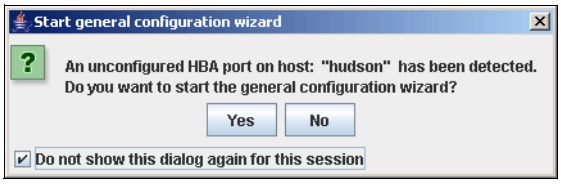


Figure 21-21 SANsurfer - Start general configuration wizard

4. After successful login, some general information is shown about the host and a list of the HBAs and ports with their iSCSI Qualified Names (IQNs) (see 1.3.1, “iSCSI initiators and targets” on page 9). See Figure 21-22.

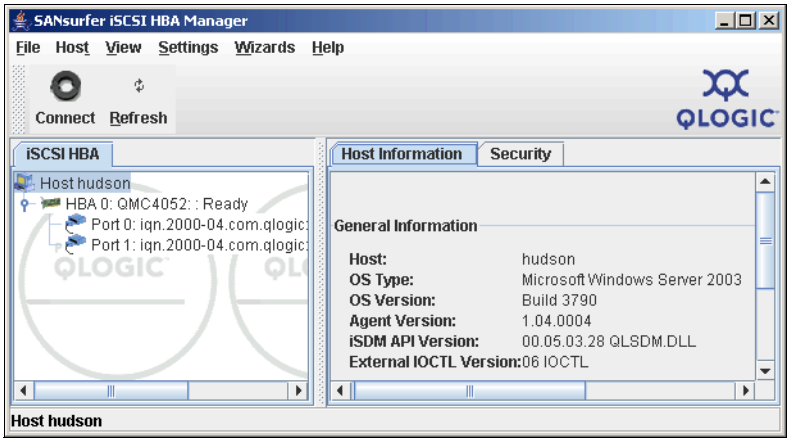


Figure 21-22 SANsurfer - General information

We have successfully connected to this host’s iSCSI HBAs. Multiple hosts can be managed from a single SANsurfer client. We can now proceed with the HBA data collection.

21.6.2 HBA data collection

This section explains how to gather information that is required in later steps during the setup. The following details will be gathered:

- Adapter BIOS and firmware level
- iSCSI qualified name of the adapter ports
- MAC addresses of the iSCSI port

Do these steps:

1. Connect to the SANsurfer GUI host, as described in 21.6.1, “Connecting with SANsurfer to the host where the HBA is installed” on page 495. Select the HBA from the iSCSI HBA listing and click the HBA information tab. See Figure 21-23.

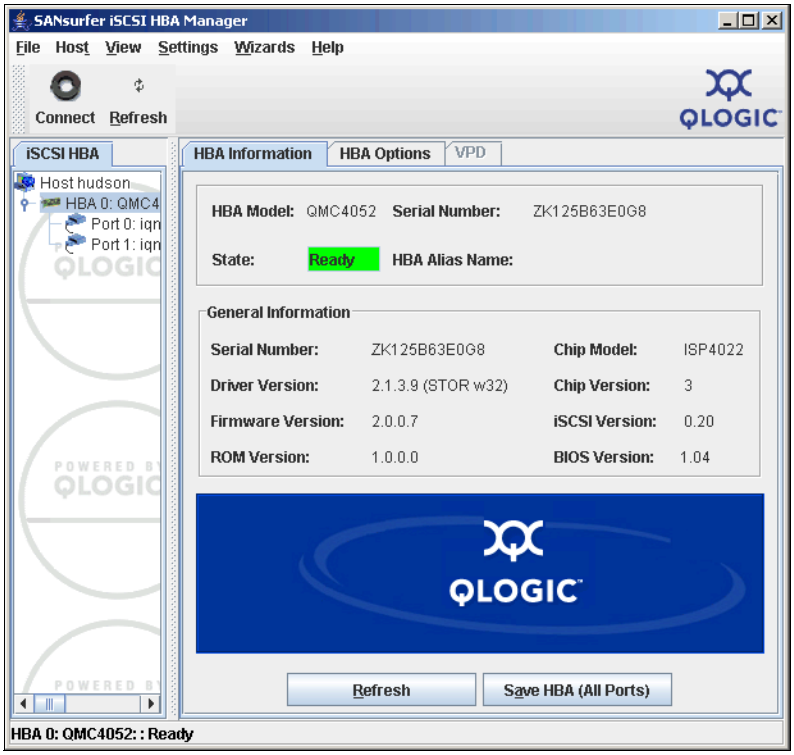


Figure 21-23 SANsurfer - HBA information

This tab shows the Adapter BIOS is at V1.04 and firmware at V2.0.0.7. Since this is back-level code, we will upgrade these levels in 21.6.3, “Flash iSCSI HBA to latest level” on page 499.

2. To start the data collection, select the first port of the HBA, and click **Port Information**, as shown in Figure 21-24.

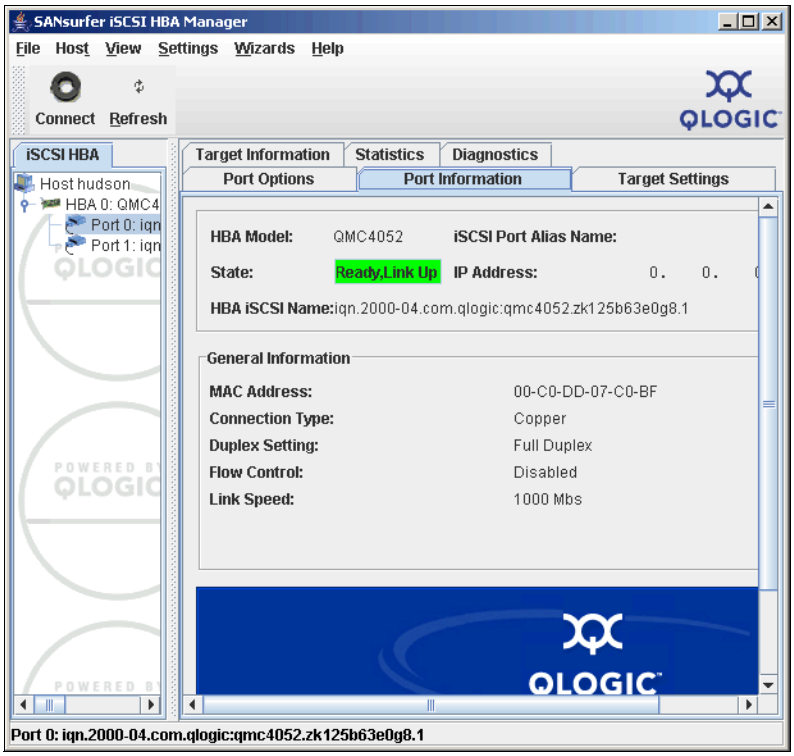


Figure 21-24 SANsurfer - Port information

The iSCSI qualified name (called HBA iSCSI Name) and MAC address of the iSCSI port are displayed here. Select these values and record them for later use.

3. Repeat the previous step for the second port.
4. Example 21-1 shows a file with the collected information.

Example 21-1 Collected HBA information

```
Host HUDSON:
  iSCSI HBA1:
    Firmware: 2.0.0.7    ==>  2.0.0.45
    BIOS:     1.04       ==>  1.09

    Port 0 MAC: 00-C0-DD-07-C0-BF
    Port 0 IQN: iqn.2000-04.com.qlogic:qmc4052.zk125b63e0g8.1

    Port 1 MAC: 00-C0-DD-07-C0-C1
    Port 1 IQN: iqn.2000-04.com.qlogic:qmc4052.zk125b63e0g8.2
```

Note: Adapter MAC addresses are also printed on the under side of the HBA.

After gathering the data, we will update the HBA code.

21.6.3 Flash iSCSI HBA to latest level

The HBA should have the latest supported code level installed that is supported by the adapter itself and the disk subsystem to which it will connect. The QLogic iSCSI Expansion Card for BladeCenter can be updated using either DOS or SANsurfer. To update through DOS, you need to first create a bootable DOS floppy disk, since the bootable image is not provided with the code download.

Updating the code using SANsurfer can be done using the code files directly, as we will show in this section. The update package contains a readme file that describes in detail how to perform the update under DOS.

Do these steps:

1. Download the BIOS and firmware files from:

   http://support.qlogic.com/support/oem_detail_all.asp?oemid=369

   At the time of writing, the latest BIOS was V1.09A, and the HBA firmware was V2.00.00.62

2. Unpack both downloaded archives into a local directory in the host running the SANsurfer GUI. As well as the readmes and release notes, you should see files similar to:

   ql4022rm.BIN Adapter BIOS

   qla4022.dl Adapter firmware

3. Start SANsurfer. Select the HBA to be updated from the iSCSI HBA listing and click the HBA options tab, as shown in Figure 21-25.

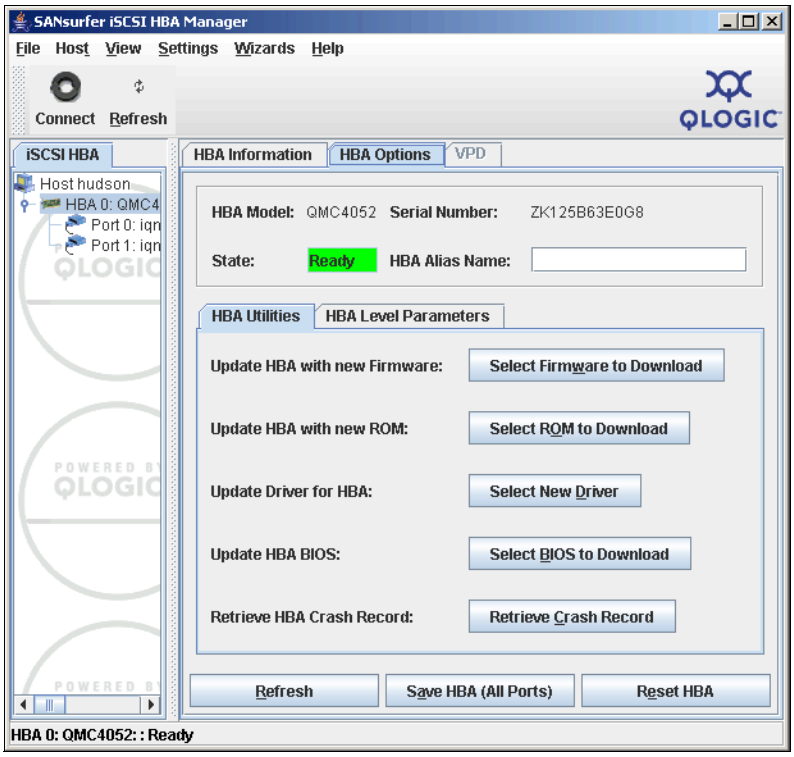


Figure 21-25 SANsurfer - HBA utilities

4. In HBA Utilities on the HBA options tab, click **Select BIOS to Download**. Select the adapter BIOS update file from its extracted location in step 2 on page 499 and click **Save**. See Figure 21-26.

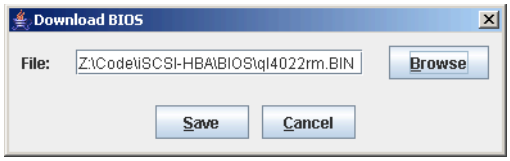


Figure 21-26 SANsurfer - Download BIOS

5. In Figure 21-27, enter the SANsurfer password and click **OK**. The default SANsurfer password is config.



Figure 21-27 SANsurfer - Security check

6. The BIOS update is now performed. You will see a message similar to Figure 21-28 when the update has finished successfully. Click **OK**.



Figure 21-28 SANsurfer - BIOS download status

7. After updating the iSCSI HBA BIOS, a refresh of the configuration might be requested. Click **Yes** to refresh the configuration, as shown in Figure 21-29.



Figure 21-29 SANsurfer - Configuration change

8. The window shown in Figure 21-25 on page 499 reappears. This time, click **Select Firmware to Download** from HBA Utilities.

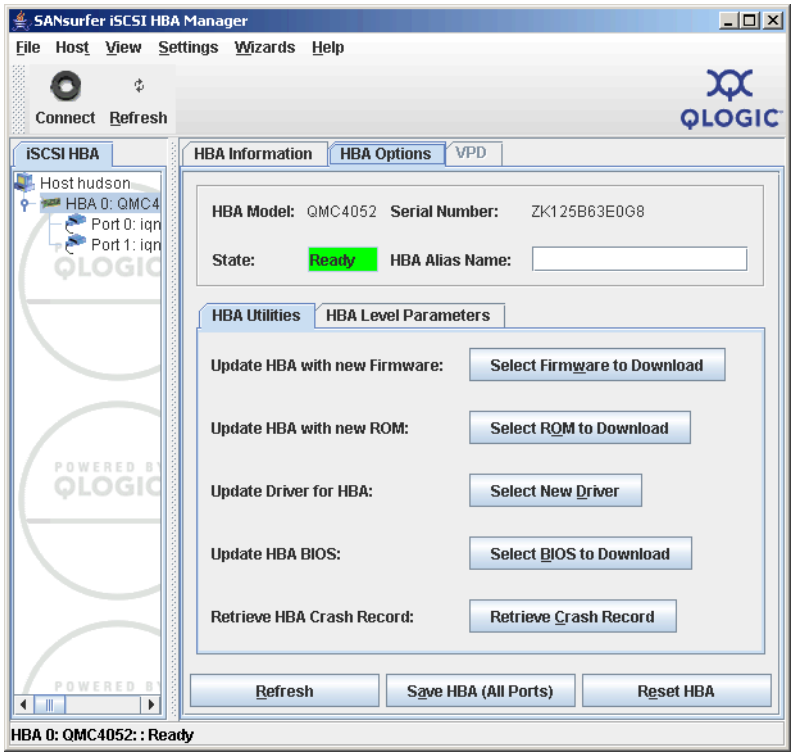


Figure 21-30 SANsurfer - HBA utilities

9. A firmware download warning appears (Figure 21-31) about actions that need to be done before the update can be performed. At this stage of the setup, no further actions are required, so click **Yes**.

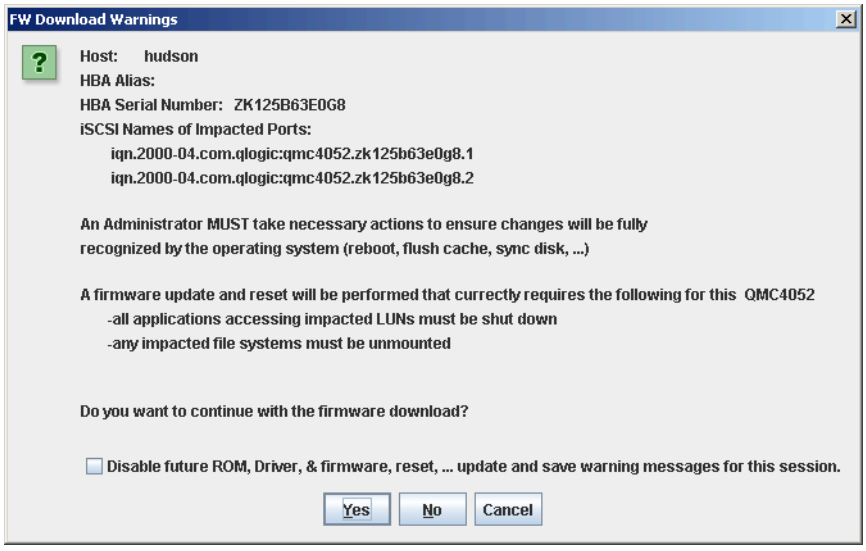


Figure 21-31 SANsurfer - Firmware download warning

10. Select the adapter firmware update file from the location where it was extracted to in step 2 on page 499. Click **Save**. See Figure 21-32.

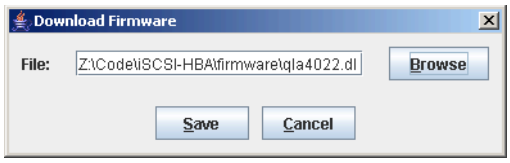


Figure 21-32 SANsurfer - Download firmware

11. Enter the SANsurfer password and click **OK**. The default password of SANsurfer is config. See Figure 21-33.



Figure 21-33 SANsurfer - Security check

12. The firmware update is performed. A message similar to Figure 21-34 displays when the update has finished successfully. Click **OK**.



Figure 21-34 SANsurfer - Firmware download status

13. After updating the iSCSI HBA firmware, a refresh of the configuration might be requested. Click **Yes** to refresh the configuration, as shown in Figure 21-35.



Figure 21-35 SANsurfer - Configuration change

14. The HBA information tab shows now the latest code level, as shown in Figure 21-36.

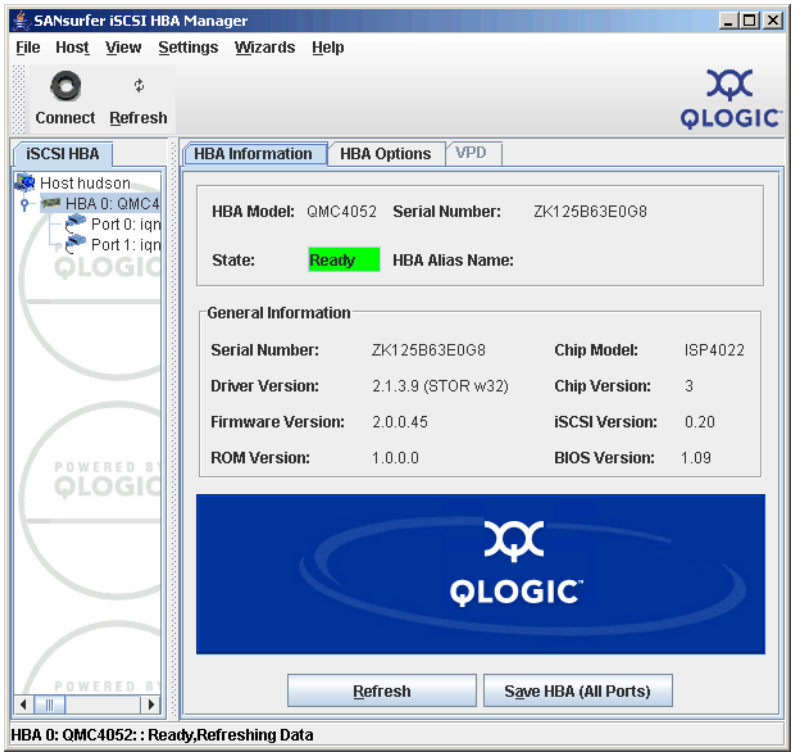


Figure 21-36 SANsurfer - HBA information with latest code level

The HBA is now running the latest level. Proceed with the HBA network configuration.

21.6.4 Configure network settings

The iSCSI HBA can use either a static IP address or a dynamic address, requested through DHCP. When using a DHCP address, we recommend configuring the DHCP server to assign the IP address based on the initiator's MAC addresses. This will allow the initiator to always receive the same address.

Note: A current adapter limitation is that a DHCP address can be assigned to only one of the ports of a QLogic multiport iSCSI HBA.

This example will use DHCP for the first port and a fixed address for the second port.

Do these steps:

1. Start SANsurfer, select HBA Port 0, click the **Port Options** tab, and then click **Network**. You should see a window similar to Figure 21-37.

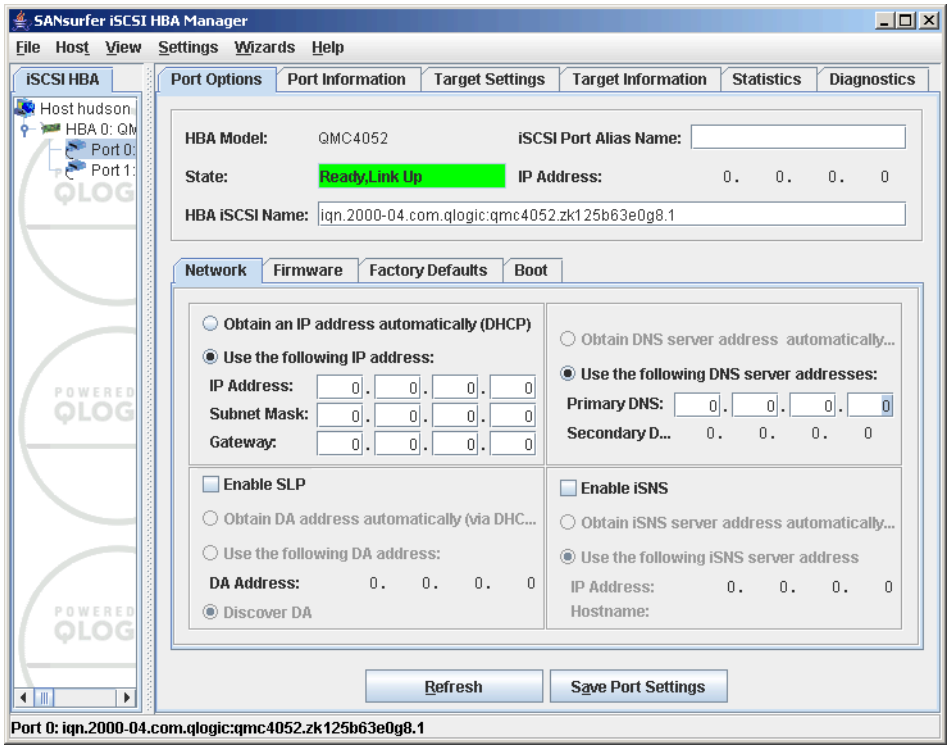


Figure 21-37 SANsurfer - Port Options - Network

2. Click **Obtain an IP address automatically**. Click **Obtain DNS server address automatically...**, and then click **Save Port Settings**. See Figure 21-38 on page 505.

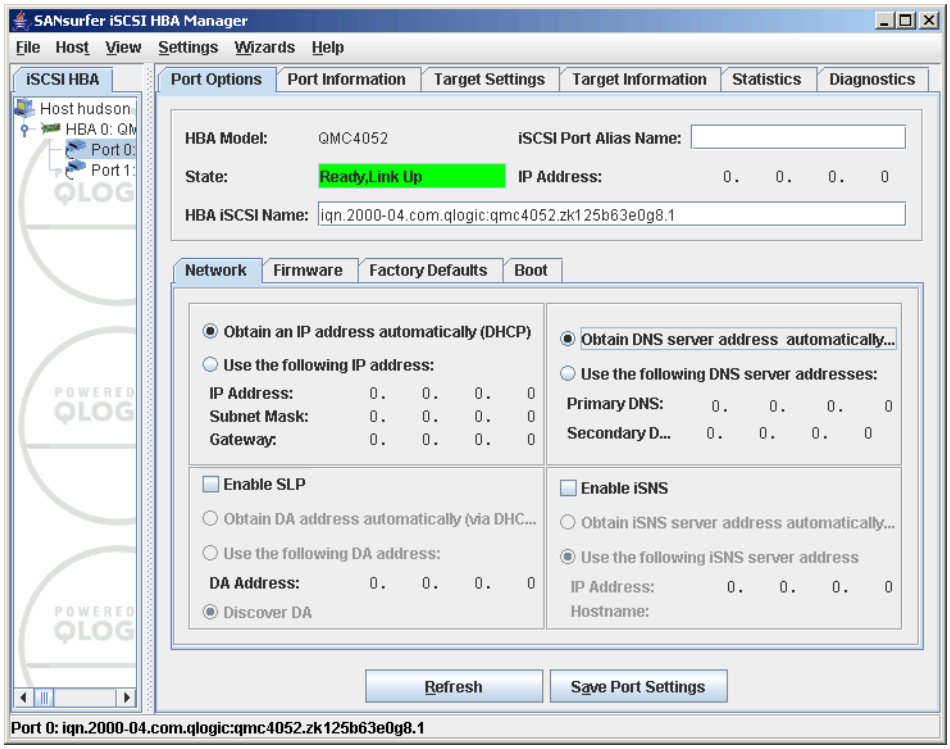


Figure 21-38 SANsurfer - Port Options - DHCP settings

3. An address configuration requires the adapter to be reset. Figure 21-39 shows the warning that applications accessing logical drives attached to this adapter need to be unmounted first. Click **Yes** to acknowledge.

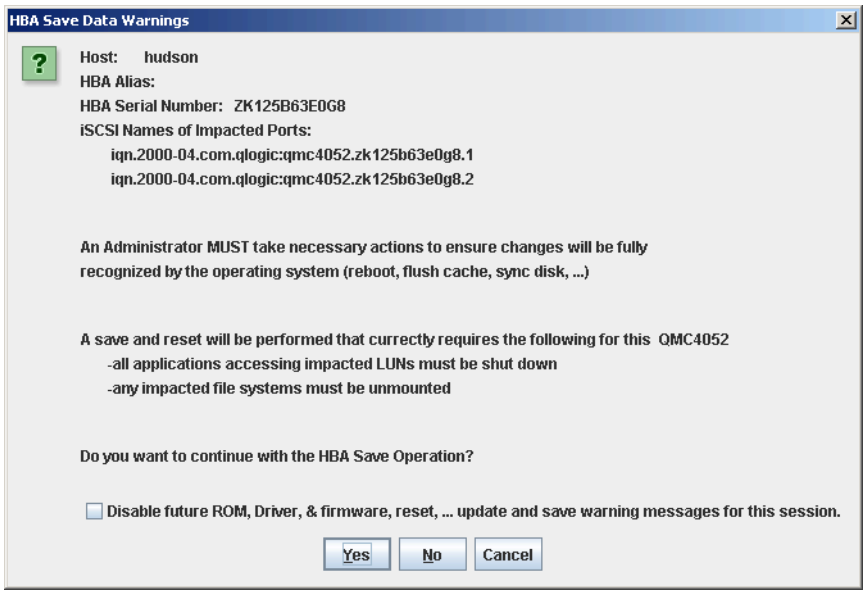


Figure 21-39 SANsurfer - HBA save data warning

4. Enter the SANsurfer password and click **OK**. The default password of SANsurfer is config.
5. A configuration refresh is requested after the network settings were saved. Click **Yes** to refresh the configuration.

6. After the configuration refresh is done, the IP address assigned by the DHCP server is visible. See Figure 21-40.

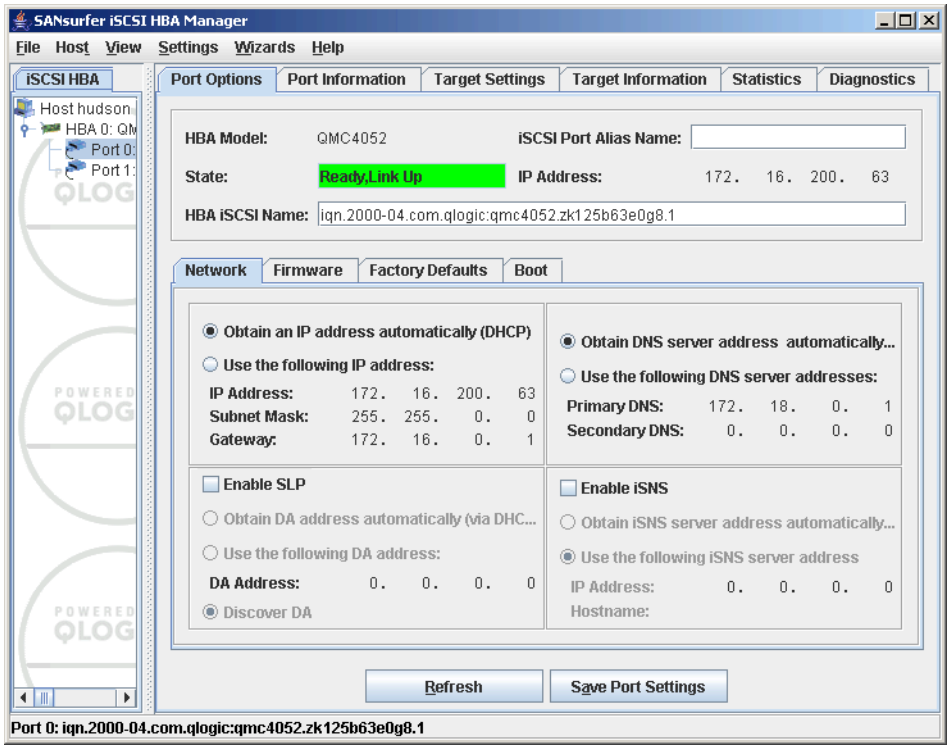


Figure 21-40 SANsurfer - Assigned DHCP address to an iSCSI port

7. Configure the IP address of the second HBA port to use later redundant paths to the storage subsystem. Select the second port of the HBA on the iSCSI HBA tab. See Figure 21-41 on page 507.

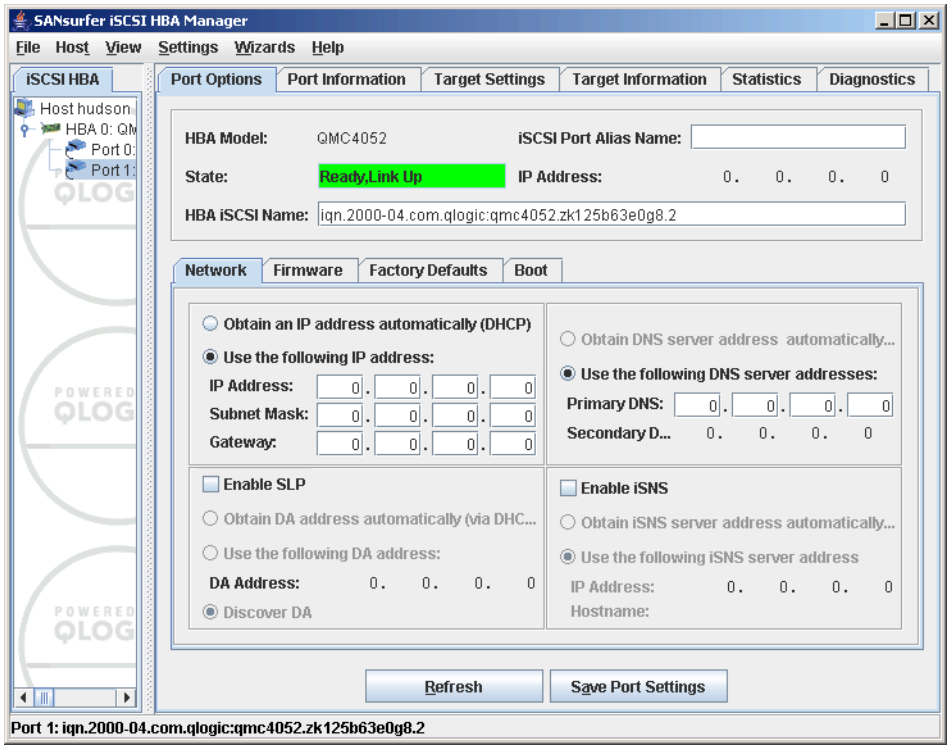


Figure 21-41 SANsurfer - Port option - Network -Second port

8. Enter the IP address, subnet mask, default gateway, and DNS server, and then click **Save Port Settings**, as shown in Figure 21-42.

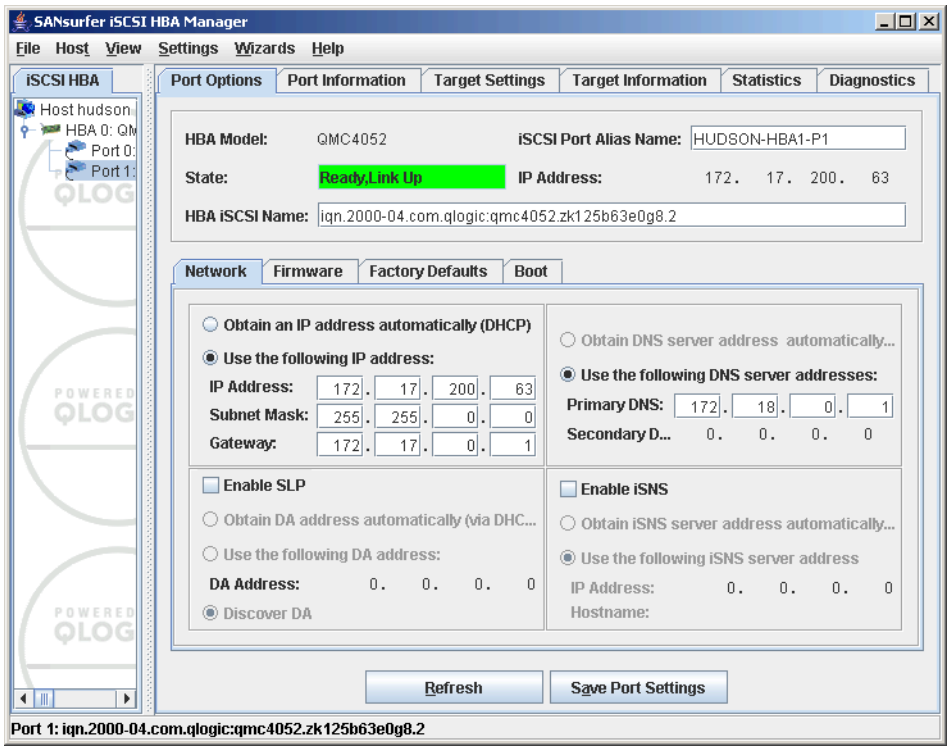


Figure 21-42 SANsurfer - Port Options - Fixed IP address settings

9. An address configuration requires a reset of the adapter. A warning is shown that applications accessing logical drives that might be attached to this adapter need to be unmounted first. Click **Yes**. See Figure 21-39 on page 505.
10. Enter the SANsurfer password and click **OK**. The default SANsurfer password is config.
11. A configuration refresh is requested after the network settings were saved. Click **Yes** to refresh the configuration.

All network settings for the iSCSI HBA are now done.

21.6.5 Test network settings

Since the IP interface from the Network Controller is not available to the running operating system, tests of the IP connection must be done from an external system or using SANsurfer. SANsurfer provides a diagnostic capability that, among other things, can be used to perform a ping test to verify the network settings.

Do these steps:

1. Connect with the SANsurfer GUI to the host, select the port from the HBA that you want to test, and click the **Diagnostic** tab. You will see the ping test, as shown in Figure 21-43.

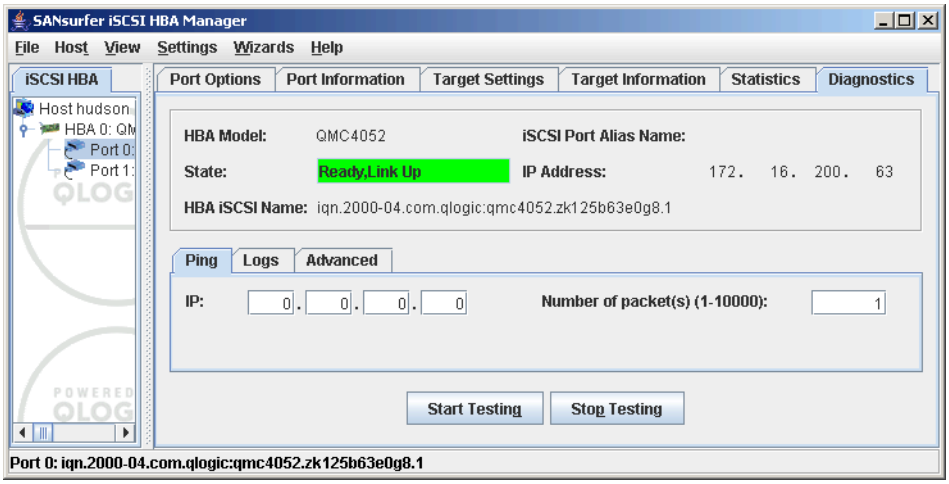


Figure 21-43 SANsurfer - Diagnostic - Ping

2. Enter the IP address of a remote host, for example, the IP address of the default gateway. Specify the number of ICMP packets to be sent and click **Start Testing**, as shown in Figure 21-44 on page 509.

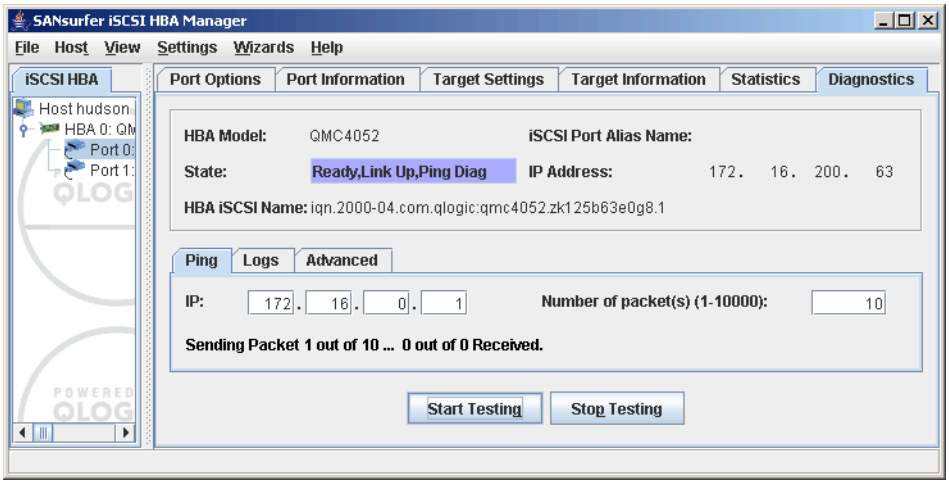


Figure 21-44 SANsurfer - Diagnostic ping ongoing

The line below the IP address shows the status of the packets sent.

3. After all packets are sent, you will see a summary of the test, similar to Figure 21-45. Click **OK**.

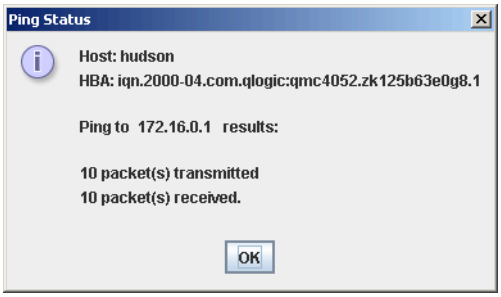


Figure 21-45 SANsurfer - Ping summary

4. Repeat this diagnostic task for the remaining iSCSI ports.

Attention: After the host is prepared, the storage subsystem must be configured.

21.7 Setting up the DS3300 logical drives and host mapping

The basic network setup needs to be done as described in 5.4.5, “(Optional) Changing the network configuration” on page 76. Also, a management station is required to run the DS3000 Storage Manager. See Chapter 6, “Installing IBM System Storage DS3000 Storage Manager” on page 81 for details about installing Storage Manager.

In this section, we describe the CLI configuration that was done to provide two logical drives configured on an IBM System Storage DS3300 to the host HUDSON running Windows Server 2003. The same script can also be used to create logical drives for a host running Windows Server 2008; however, the hostType needs to be set to 2 for the **`create host`** command.

Do these steps:

1. A command script file, shown in Example 21-2, was created containing the commands to create logical drives, a hot spare, and set up logical drive mapping. We named the file *hudson.cmd*. For formatting reasons in this book, we use backslashes to show line breaks; however, you should remove these line break characters; every line is a single command.

   Example 21-2 HUDSON command file

```
create hostGroup userLabel="North-America";

create host userLabel="Hudson" hostGroup="North-America" hostType=0;

create iscsiInitiator iscsiName="iqn.2000-04.com.qlogic:qmc4052.zk125b63e0g8.1" \
userLabel="Hudson-Port0" host="Seine";

create iscsiInitiator iscsiName="iqn.2000-04.com.qlogic:qmc4052.zk125b63e0g8.2" \
userLabel="Hudson-Port1" host="Seine";

set drive [7,1] hotspare=true;

create logicalDrive drives=(7,2 7,3) RAIDLevel=1 userLabel="Hudson-1" segmentSize=64 \
capacity=10 GB owner=a usageHint=fileSystem  mapping=none;

create logicalDrive drive [7,4] RAIDLevel=0 userLabel="Hudson-2" segmentSize=64 \
capacity=8 GB owner=b usageHint=fileSystem  mapping=none;

set logicalDrive ["Hudson-1"] logicalUnitNumber=0 host="Hudson";

set logicalDrive ["Hudson-2"] logicalUnitNumber=1 host="Hudson";
```

2. The script file is executed by SMcli using the syntax shown in Example 21-3. The output is directed to a file called hudson.out for later verification. When this file has a size of zero, this means that every command in the command file was executed without errors.

   Example 21-3 SMcli executes the script

```
amazon:~ # SMcli -n DS3300 -f hudson.cmd -o hudson.out -S
amazon:~ #
```

Now the logical drives and host mapping are done. Later, we define security for accessing the DS3300 after the basic access is working.

21.8 Configuring iSCSI targets in the host

The DS3300 is now configured. Next, we want to configure access for the iSCSI HBA to the defined logical drives. To set up this access, do the following tasks:

1. Start SANsurfer, select the first port from the HBA, as shown in Figure 21-37 on page 504, and click **Target Settings**. A window similar to Figure 21-46 on page 511 will display, showing a list of all discovered targets.

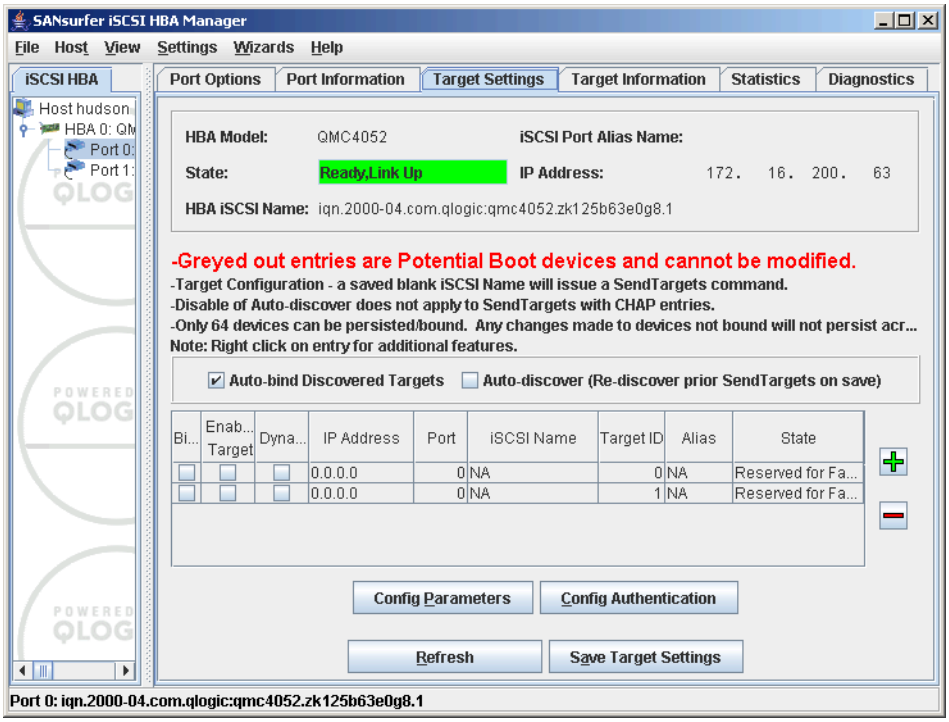


Figure 21-46 SANsurfer - Target settings

2. There are currently no targets defined in the target list. Click the green plus sign to add a new target. A dialog opens to specify the IP address of an iSCSI portal. The example here uses the first iSCSI port IP address of the first controller in the DS3300, as shown in Table 21-1 on page 483. This was configured on the DS3000 in 13.2.3, “Configure iSCSI Host Ports” on page 291. Enter the address and click **OK**. See Figure 21-47.

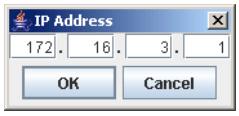


Figure 21-47 SANsurfer - iSCSI portal IP address

3. A target is added to the list, with the specified address, which will be used to discover the specified target and return a list of LUNs. The text `Not saved` in the State column for the new target indicates that this configuration modification was not yet saved. Click **Save Target Settings** to save the target settings. See Figure 21-48.

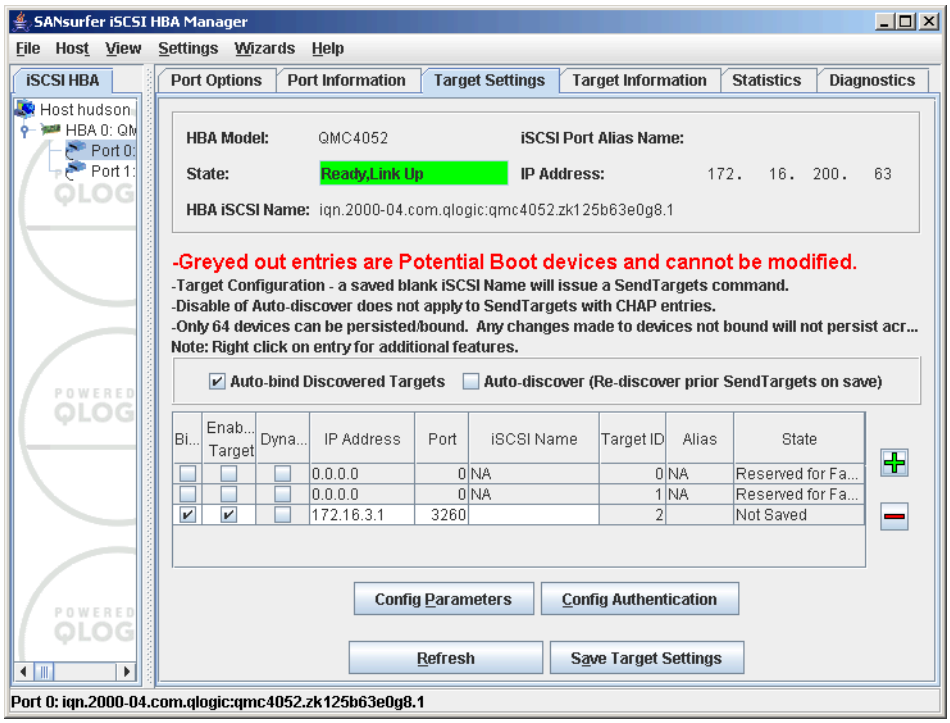


Figure 21-48 SANsurfer - Configured target

4. Click **Yes** to acknowledge the warning in Figure 21-49.

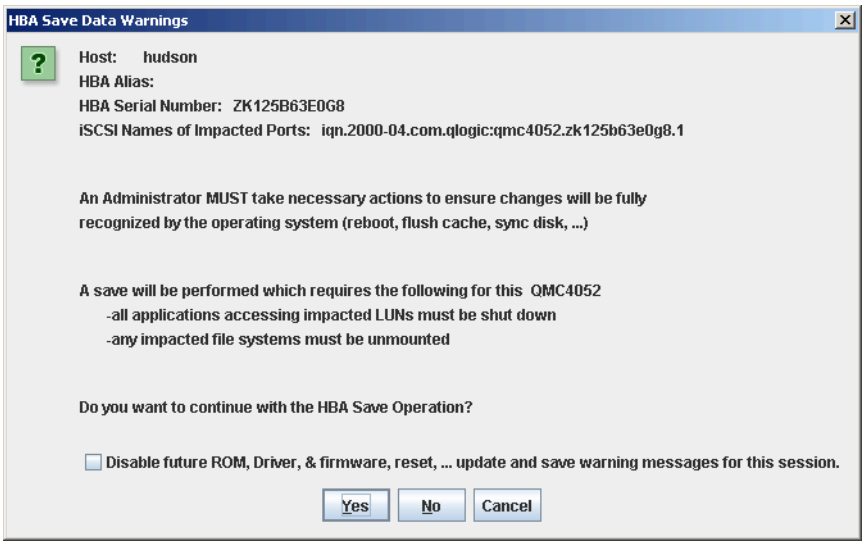


Figure 21-49 SANsurfer - HBA save data warning

5. Enter the SANsurfer password (Figure 21-50 on page 513) (the default password is config) and click **OK**.



Figure 21-50 SANsurfer - Security Check

6. After the configuration is saved, a configuration refresh is requested. Click **Yes** to accept this configuration refresh.
7. After the configuration has refreshed, the specified target is listed under the host port in the left hand pane. Figure 21-51 shows this as the blue entry in the iSCSI HBA tab. This blue target is a special entry called Sent Target Hosts that is used to get available LUNs from the target. There will never be an active session for this kind of target and it is only used for discovery purposes.

   A second line in the target list (for address 172.16.3.1) has also appeared that represents the target that was specified in the previous step and all available LUNS. For this target, unlike the previous one, an iSCSI qualified name is included in the target list. The state shows an active connection (Status “Session Active”).

   The iSCSI HBA list also now includes a green Device listing that represents the newly discovered target. All LUNs that are available are also visible. The target list shows a target ID of 64 for this target. That means that this target is a dynamic target, which is not persistently bound to this initiator. All dynamic targets start with the target ID of 64. Persistent targets start with ID 0, but ID 0 and 1 are reserved for boot entries. Therefore, since the highest persistent target is ID 63, there are 62 maximum targets that can be configured.

   The dynamically discovered target must be saved to persistently bind it. Click **Save Target Settings**.

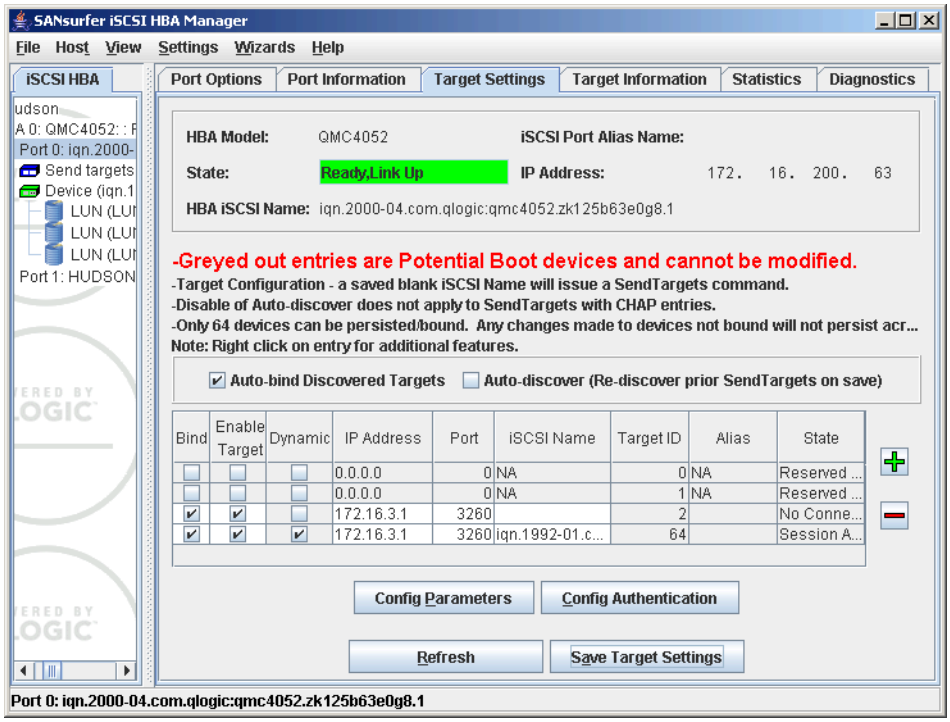


Figure 21-51 SANsurfer - Persistent and dynamic target

8. Repeat steps 3 to 7 for HBA port 1 to the first iSCSI port of the second DS3300 controller using the address specified in Table 21-1 on page 483.
9. Repeat steps 3 to 8 for Port 0 and Port 1 to the second iSCSI port of both controllers in the DS3300, using the addresses specified in Table 21-1 on page 483.
10. After adding all the iSCSI portal IP addresses to the iSCSI initiator ports, the configuration should look like Figure 21-52.

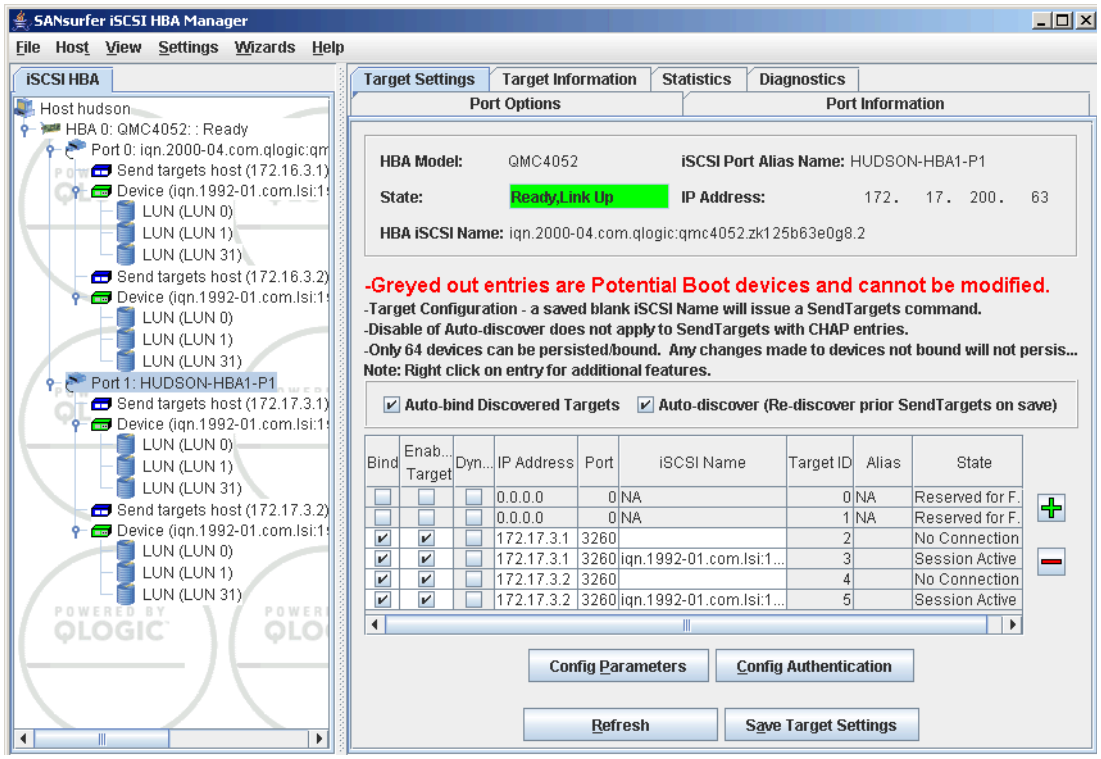


Figure 21-52 SANsurfer - targets discovered

The configuration of the targets is done. Now we install the failover driver on the attached host.

21.9 Installing RDAC driver for DS3000

We need to install the RDAC driver in the host that will access the iSCSI LUNs to handle multiple paths to the same logical drive. For Windows Server 2003, Windows 2003 Service Pack 2 and Microsoft hotfix KB932755 must be installed first before the RDAC component of Storage Manager can be installed. Install hotfix KB932755, which can be downloaded from:

http://www.microsoft.com/downloads/details.aspx?FamilyID=dead950b-da47-4dc2-abcd-40f61bd4d1aa&DisplayLang=en

Note: Here, RDAC means a device specific module (DSM) for Microsoft MPIO. Do not confuse this with the RDAC that is included in the DS4000 Storage Manager (it is a completely separate package); make sure to use the DS3000 version.

The Windows RDAC driver is packaged with DS3000 Storage Manager; when installing RDAC, you must also select the utilities. Those components are in the host selection. For more details about installing Storage Manger, see Chapter 6, “Installing IBM System Storage DS3000 Storage Manager” on page 81.

After installing RDAC, you have to reboot the host to activate the driver.

Note: For more information about this topic, refer to:

http://support.microsoft.com/kb/923801

21.10 Test configuration

After you have installed RDAC, your disks should be visible. You should test to verify proper operation of the access to the storage subsystem and correct function of the failover driver. To simulate network outages between the storage subsystem and a host, we recommend detaching a network cable; merely disabling the port through the software is typically not a good enough test.

21.11 Enhancing iSCSI connection security

After everything is working and tested, you should implement security for the iSCSI connection; essentially, this means configuring the initiator and target authentication. Initiator authentication means that an initiator must prove its identity with a password that is known by the target, when the initiator attempts access. Target authentication is the opposite: The target authenticates itself to the initiator with a password.

By default, the only security enabled is the iSCSI Qualified Name (IQN) that is used to allow access to the logical drives (which we configured). Since an IQN can be modified within Storage Manager, this does not protect against spoofing of the qualified name, and security can be compromised. This section describes how to set up initiator and target authentication.

To implement security, all accessed logical drives should be unmounted first in the operating system. You receive a notification when changing the SANsurfer configuration as a reminder that unmounting is necessary.

21.11.1 Setting up initiator authentication

Initiator authentication is called *mutual authentication permissions* in DS3000 Storage Manager.

Do these steps:

1. Stop any applications that may be accessing the logical drives of the DS3300. Log off from all targets connected to an initiator port. To log off, open SANsurfer and select the initiator port 0. The Target Settings tab shows two iSCSI targets with an status of Session Active. Right-click the first one and select **Logout/Reconnect Device** from the context menu, as shown in Figure 21-53.

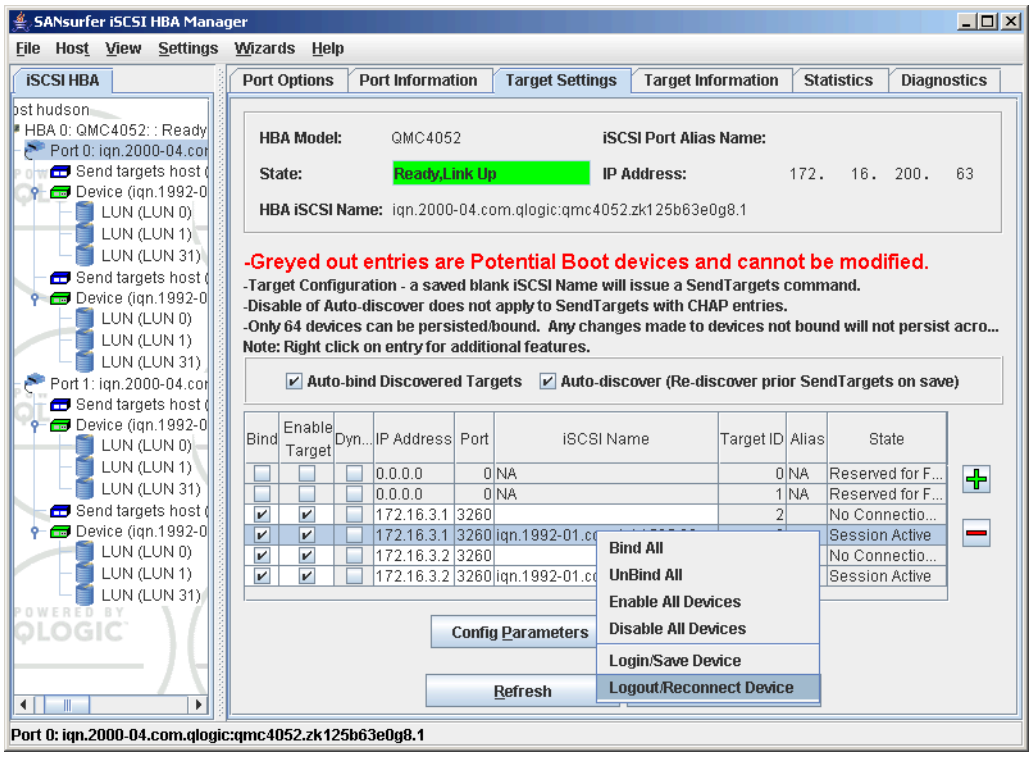


Figure 21-53 SANsurfer - Logout from target

2. Enter the SANsurfer password and click **OK**.
3. Save the configuration and refresh the configuration as prompted (click **Yes** to accept the configuration refresh). The configuration should now appear similar to Figure 21-54 on page 517, with no sessions active for the chosen target.

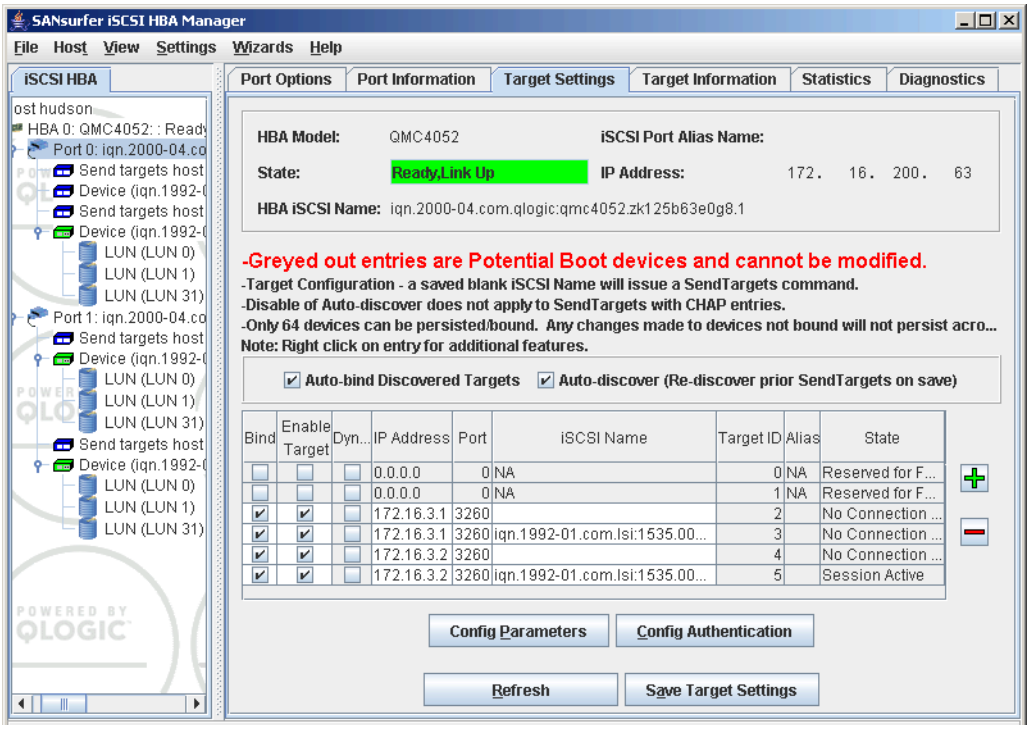


Figure 21-54 SANsurfer - Disconnect target

4. Repeat the logout for all the remaining targets and initiator ports. Leave SANsurfer open; you will need it again soon.
5. Use the Storage Manager CLI commands (**`set iscsiInitiator`**) shown in Example 21-4 to set up the CHAP secret (Challenge Handshake Authentication Protocol) for the already defined host ports of host HUDSON. See Example 21-4.

Example 21-4 SMcli commands to set CHAP secrets for initiators

```
amazon:~ # SMcli -n DS3300 -c "set iscsiInitiator [\"Hudson-Port0\"]
host=\"Hudson\" chapSecret=\"a1234567890z\"; " -S

amazon:~ # SMcli -n DS3300 -c "set iscsiInitiator [\"Hudson-Port1\"]
host=\"Hudson\" chapSecret=\"a1234567890y\"; " -S

amazon:~ #
```

Note: CHAP (RFC1944) is the most basic level of iSCSI security available.

6. After setting the CHAP secret on the DS3300, the HBA initiators need to know the CHAP secret and the associated targets to use the secret for. In SANsurfer, click **Configure Authentication**. See Figure 21-55.

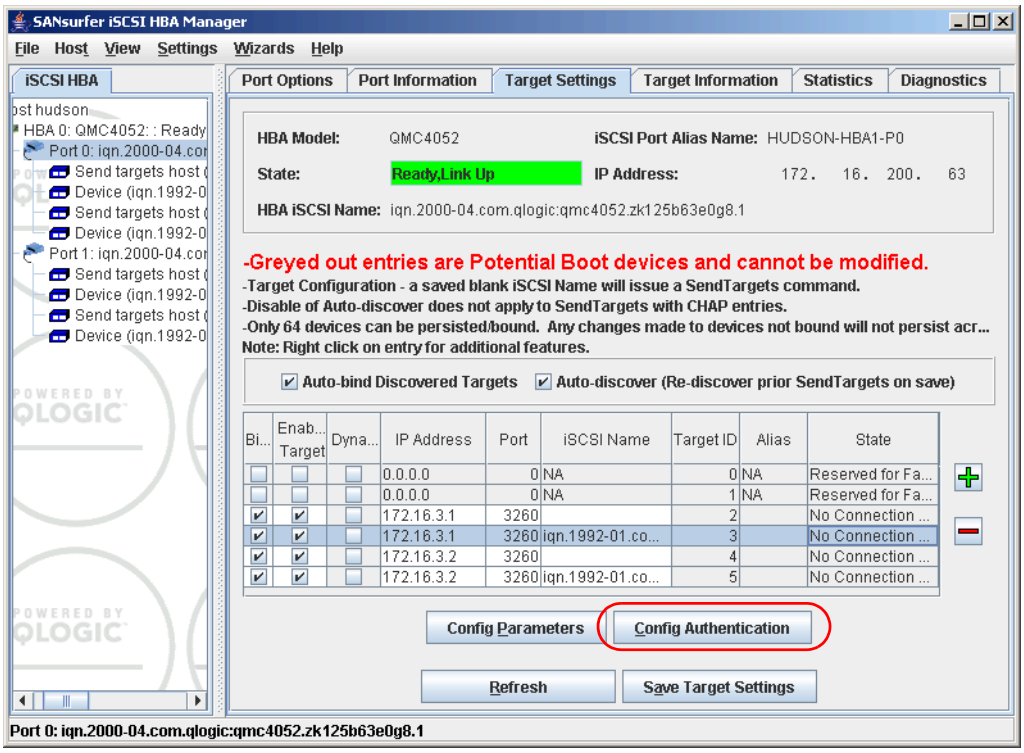


Figure 21-55 SANsurfer - Configure initiator authentication

7. Enter the SANsurfer password and click **OK**.

8. In Figure 21-56, click the green plus sign beside the CHAP entries table to add a new line in the table. Enter the IQN of the iSCSI initiator port to be configured in the Initiator Name column. Enter the CHAP secret as it was defined on the DS3300 in Example 21-4 on page 517 in the initiator Secret column.

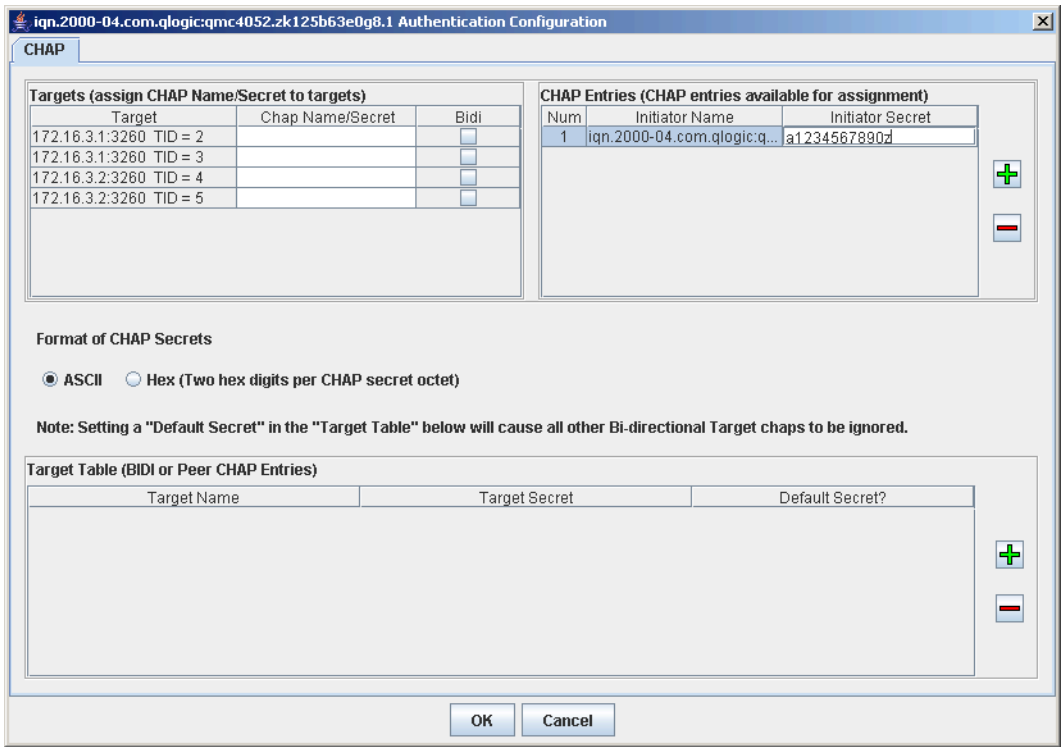


Figure 21-56 SANsurfer - Initiator CHAP entry

9. In the target table in Figure 21-57, select from the drop-down menu in the CHAP Name/Secret column the CHAP entry that was just created. The line number of the CHAP entries table helps to identify the entries that are selected for a specific target. Click **OK** to exit the authentication settings. Do not refresh the configuration until the modifications are saved.

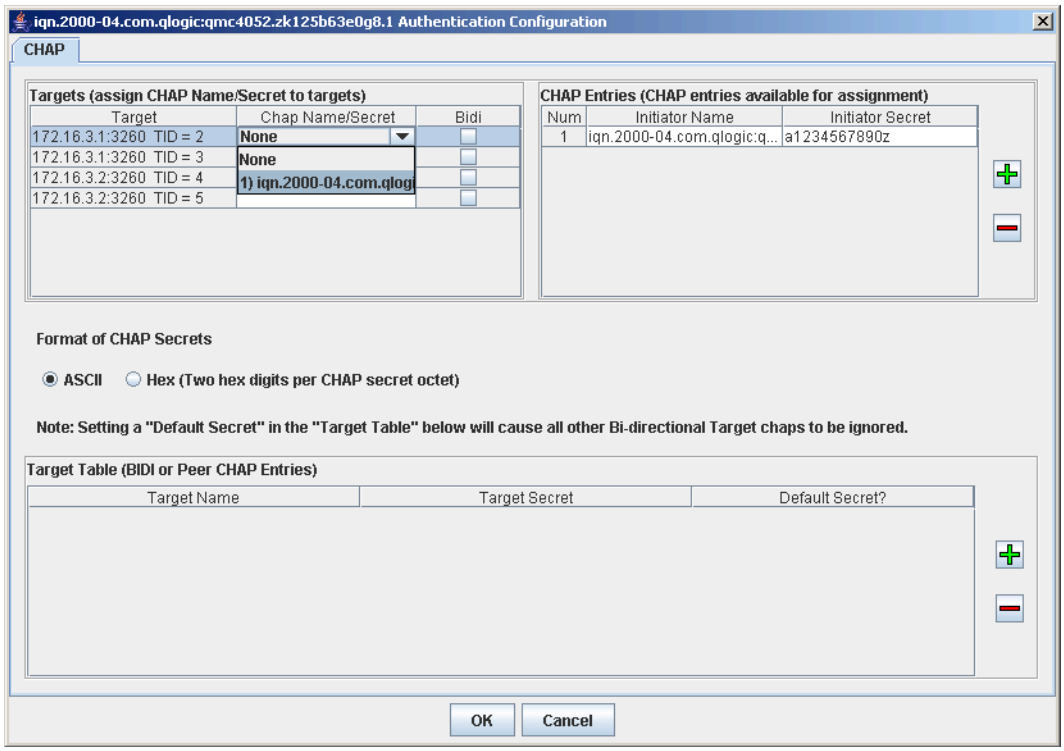


Figure 21-57 SANsurfer - CHAP entries assigned to targets

10. In Figure 21-58 on page 521, click **Save Target Settings** to save the authentication modifications.

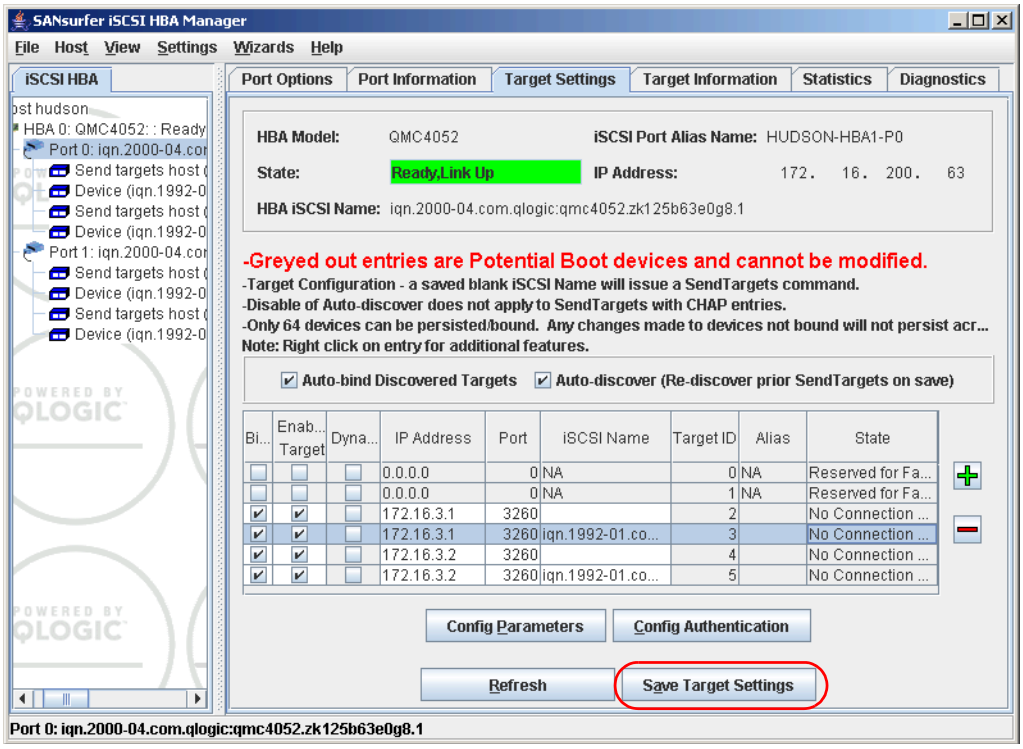


Figure 21-58 SANSurfer - iSCSI targets

11. Enter the SANsurfer password and click **OK**.
12. Confirm the warning by clicking **Yes**. See Figure 21-59.

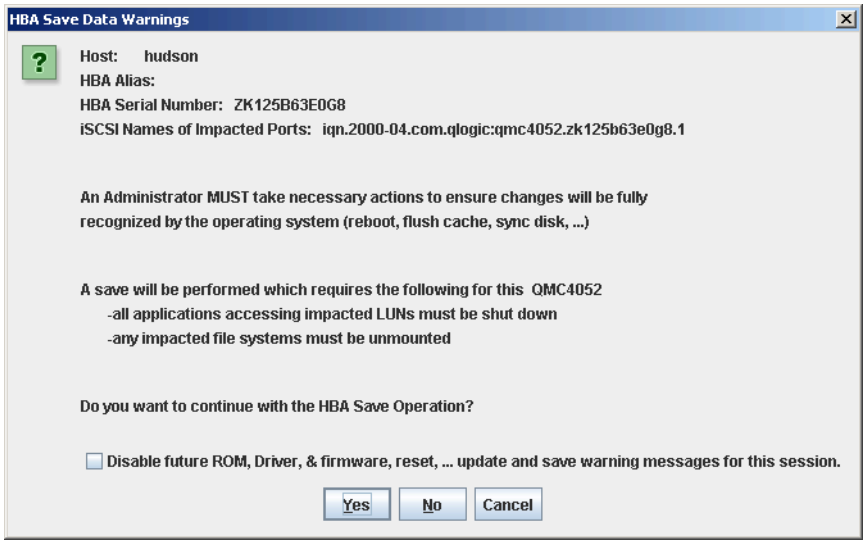


Figure 21-59 SANsurfer - Save Data Warning

13. A configuration refresh might be requested by SANsurfer. Click **OK** to refresh the configuration.
14. Repeat step 6 on page 518 to step 13 to configure authentication settings for the other iSCSI host port using the other CHAP secret defined in Example 21-4 on page 517 and the targets configured on this host port.

15. After configuring the authentication settings on the DS3300 and the HBA, a login to each target can be performed. In SANsurfer, select the first iSCSI port of the HBA. On the Target Settings tab, right-click a line that contains an IQN. From the context menu, select **Login/Save Device**, as shown in Figure 21-60. This will establish a new iSCSI session with the target using the new security settings.

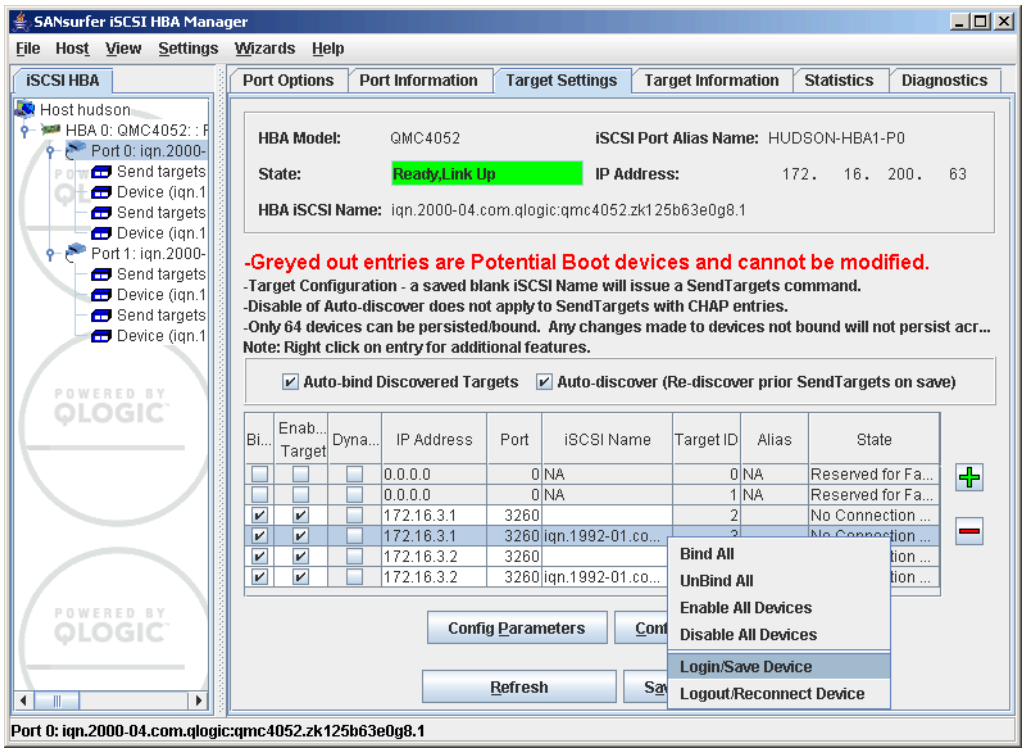


Figure 21-60 SANsurfer - Login into iSCSI target

16. A configuration refresh might be requested by SANsurfer. Click **OK** to refresh the configuration.
17. Repeat step 15 to step 16 for each remaining target of all the iSCSI host ports.
18. After all sessions are active (as shown in Figure 21-61 on page 523), the applications that access these logical drives can be re-started.
19. The authentication will now be performed automatically whenever the system re-starts.

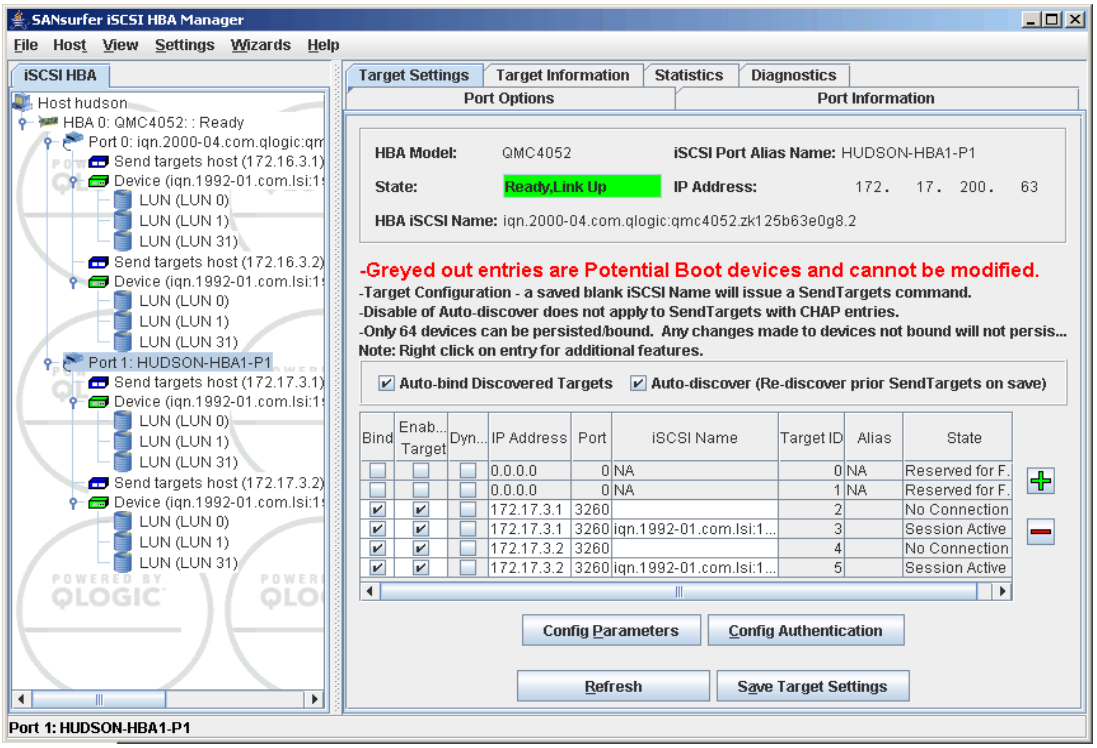


Figure 21-61 SANsurfer - All targets reconnected

For even more security in the iSCSI connection, target authentication can also be enabled.

21.11.2 Setting up target authentication

Target authentication with the QLogic iSCSI HBAs requires initiator authentication to be already configured.

Follow these steps:

1. Stop applications that may be accessing the logical drives on the DS3300 and log out from all the targets connected to each initiator port, as described in step 1 on page 516 to step 4 on page 517. SANsurfer should show no sessions active on the targets. Keep SANsurfer open, because it will be required again soon.
2. Use the Storage Manager CLI (**`set iscsiTarget`**) command, as shown in one of the following examples. In Example 21-5, the command forces iSCSI sessions to use target authentication with the correct CHAP secret specified, as it by default disables authenticationMethod none. In Example 21-6 on page 524, iSCSI sessions are permitted both with and without target authentication, since authenticationMethod=none is specified along with authenticationMethod=chap. You should use the appropriate format of the command for your configuration, depending on whether you want target authentication to be compulsory. See 13.1, “Authentication” on page 285 for more information.

Example 21-5 SMcli - set and enforce target authentication

```
amazon:~ # SMcli -n DS3300 -c "set iscsiTarget
<\"iqn.1992-01.com.lsi:1535.000000000000000000000000000000\">
authenticationMethod=chap chapSecret=\"a01234567890t\";" -S
amazon:~ #
```

Example 21-6 SMcli - set target authentication as optional

```
amazon:~ # SMcli -n DS3300 -c "set iscsiTarget
<\"iqn.1992-01.com.lsi:1535.000000000000000000000000000000\">
authenticationMethod=none authenticationMethod=chap
chapSecret=\"a01234567890t\";" -S
amazon:~ #
```

3. After modifying the configuration on the DS3300, the initiators need to know the CHAP secret of the iSCSI target to be able to identify the subsystem. In SANsurfer, click **Configure Authentication**, as shown in Figure 21-55 on page 518.
4. Enter the SANsurfer password and click **OK**.
5. Click the green plus sign beside the Targets table (BIDI or Peer CHAP Entries) to add a new line in the table. Enter the IQN of the iSCSI target that will be configured with target authentication. Enter the CHAP secret (as defined on the DS3300 in Example 21-5 on page 523) in the Target Secret column. **Default Secret** could be checked when only one iSCSI target is connected to this initiator, which uses target authentication. This entry will be used for every target that is listed in the upper left Targets table (Assign CHAP Name/Secret to targets), which has the Bidi check box checked. See Figure 21-62.

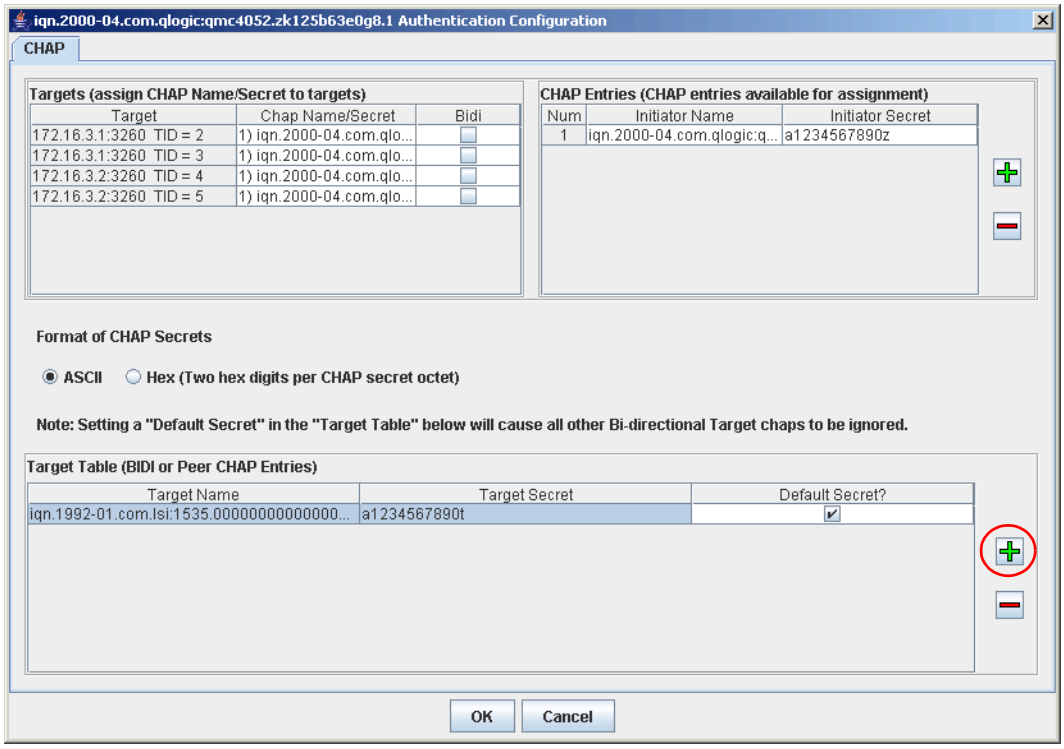


Figure 21-62 SANsurfer - Configure target secrets

6. Check the **Bidi** check box for each target that requires target authentication (Figure 21-63). Click **OK** to leave the authentication settings dialog.

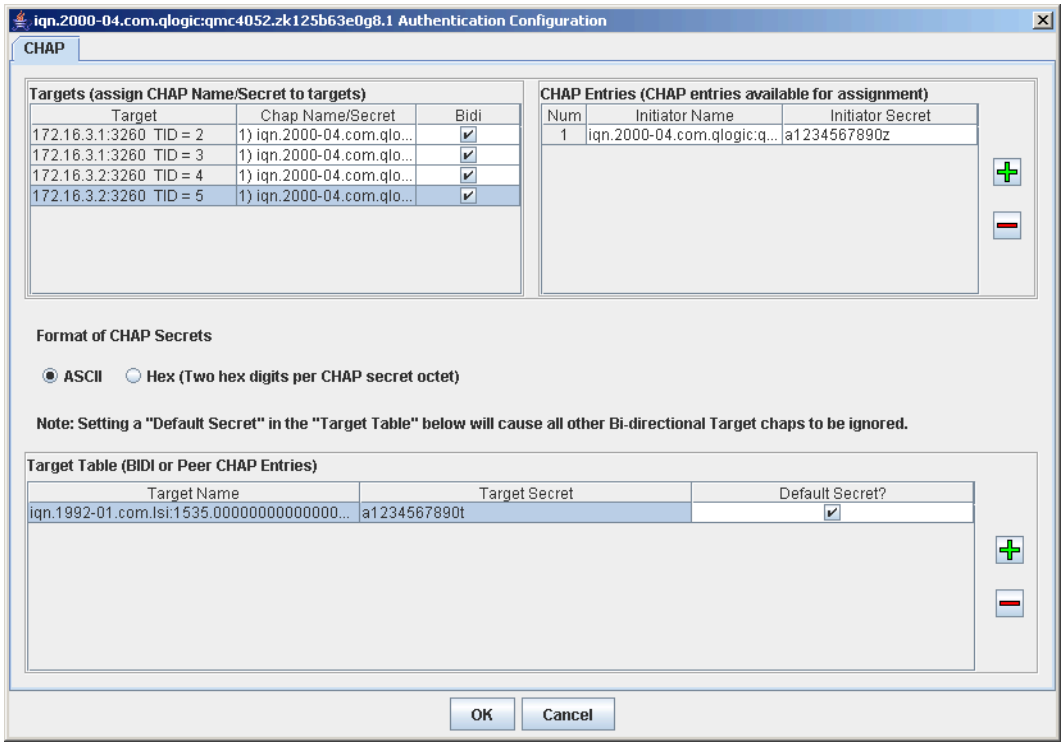


Figure 21-63 SANsurfer - Use target authentication with specified targets

7. Click **Save Target Settings** to save the authentication modifications just made (Figure 21-64).

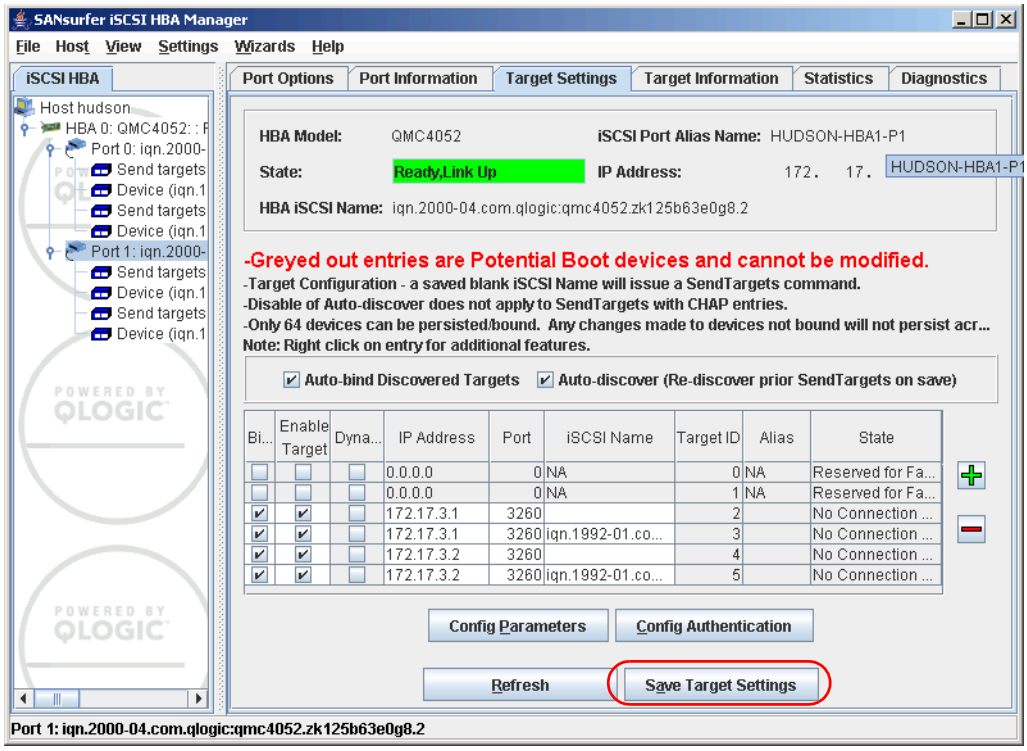


Figure 21-64 SANSurfer - iSCSI targets

8. Enter the SANsurfer password and click **OK**. Confirm the warning by clicking **Yes**, as shown in Figure 21-65.

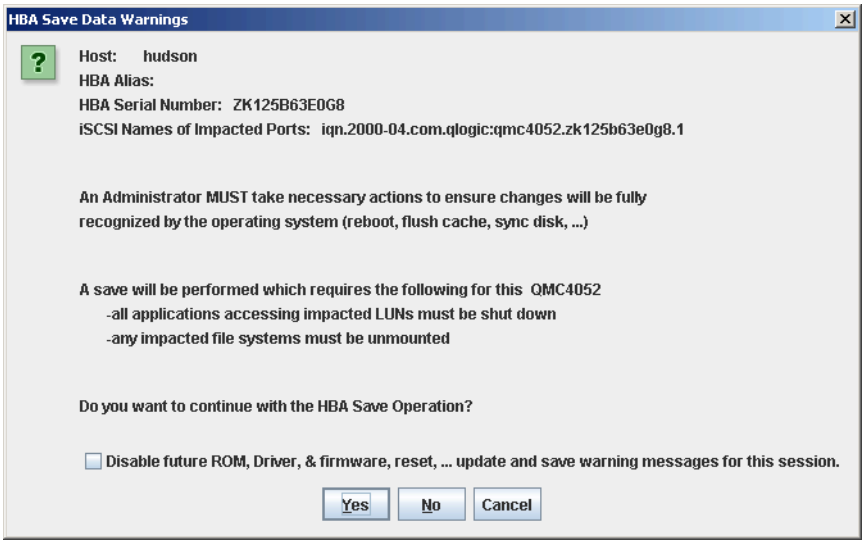


Figure 21-65 SANsurfer - Save data warning

9. A configuration refresh might be requested by SANsurfer. Click **OK** to refresh the configuration.
10. Repeat step 3 on page 524 to step 9 to configure authentication settings for the other iSCSI host port and the targets configured on this host port.

11. After configuring the authentication settings on the DS3300 and the HBA, a login to each target can be performed. In SANsurfer, select the first iSCSI port of the HBA. On the Target Settings tab, right-click a line that contains an IQN. From the context menu, select **Login/Save Device** (Figure 21-66). This will establish a new iSCSI session with the target using the new security settings.

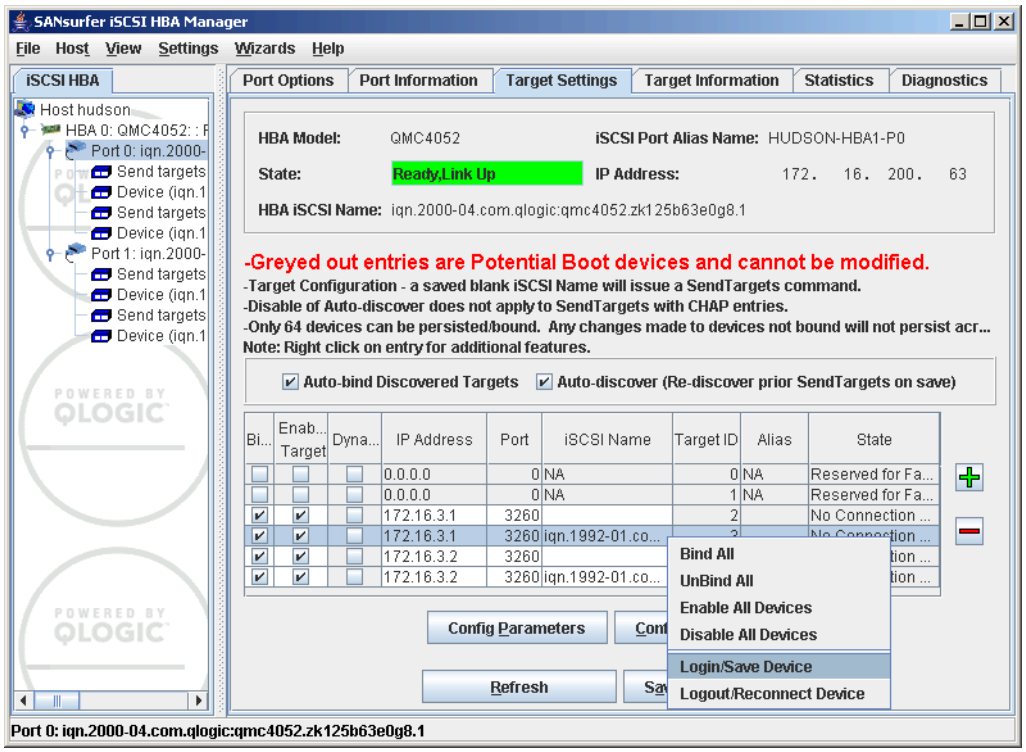


Figure 21-66 SANsurfer - Login into iSCSI target

12. Enter the SANsurfer password and click **OK**. A configuration refresh might be requested by SANsurfer. Click **OK** to refresh the configuration.
13. Repeat step 11 to step 12 for each remaining target of all iSCSI host ports.

14. After all sessions are active (as shown in Figure 21-67), the applications that access these logical drives can be re-started.

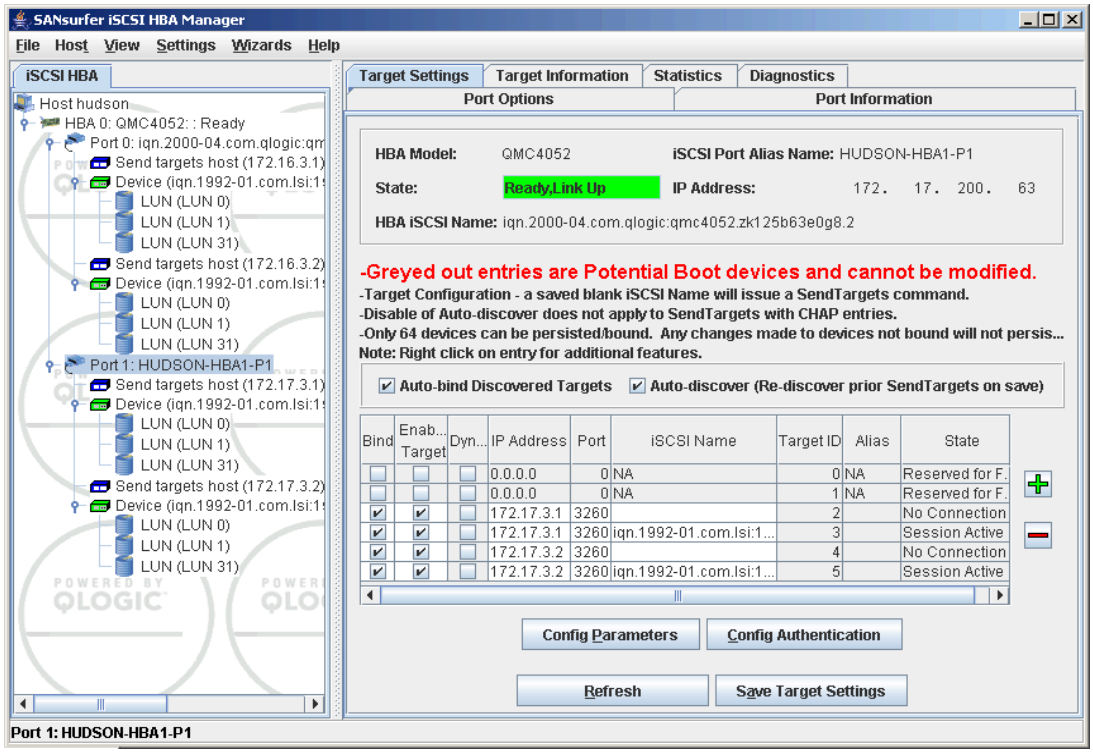


Figure 21-67 SANsurfer - All targets reconnected

22

iSCSI configuration 2 - Software Initiator on Red Hat Linux V5.2

In this chapter, we discuss a sample configuration that shows how to connect logical drives configured on an IBM System Storage DS3300 to a Red Hat Enterprise Linux 5 update 2 operating system running on an IBM BladeCenter HS21 with included Open iSCSI software initiator.

Note: iSCSI boot from SAN is not covered at this time. Configuration information for iSCSI boot from SAN for IBM BladeCenter HS21 can be found in the *iSCSI Boot from SAN Installation And Setup Guide*. This guide can be found at the following URL:

`https://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-64880&brandind=5000020`

The blade server has an Ethernet daughter card installed in the slot that connects to module bays 3 and 4 of the BladeCenter chassis. The decision to use module bays 3 and 4 was made to have dedicated network equipment for the iSCSI traffic. That increases network stability and provides a safer SCSI connection later to the host.

Nortel Layer 2/3 Gigabit Ethernet switches are installed in module Bays 1, 3, and 4. Bay 1, which connects to the first onboard NIC of the blades, is used for the public network access. Bay 3 and 4, which connect to an expansion slot on the blade server, are used to connect to the installed gigabit Ethernet daughter card.

We assume that the DS3300 is set up with IP addresses, as described in 13.2.3, “Configure iSCSI Host Ports” on page 291, and is ready to have the logical drives and host mapping configured.

Note: At the time of writing, RDAC V09.03.0B05.0023 was used to illustrate the following steps. Refer to the IBM Systems support Web site, as described in 17.4, “Installing RDAC for Linux” on page 408, and download the latest available version of RDAC

© Copyright IBM Corp. 2007, 2009. All rights reserved.

22.1 Network configuration

Figure 22-1 shows a basic network diagram of the connections between the components.

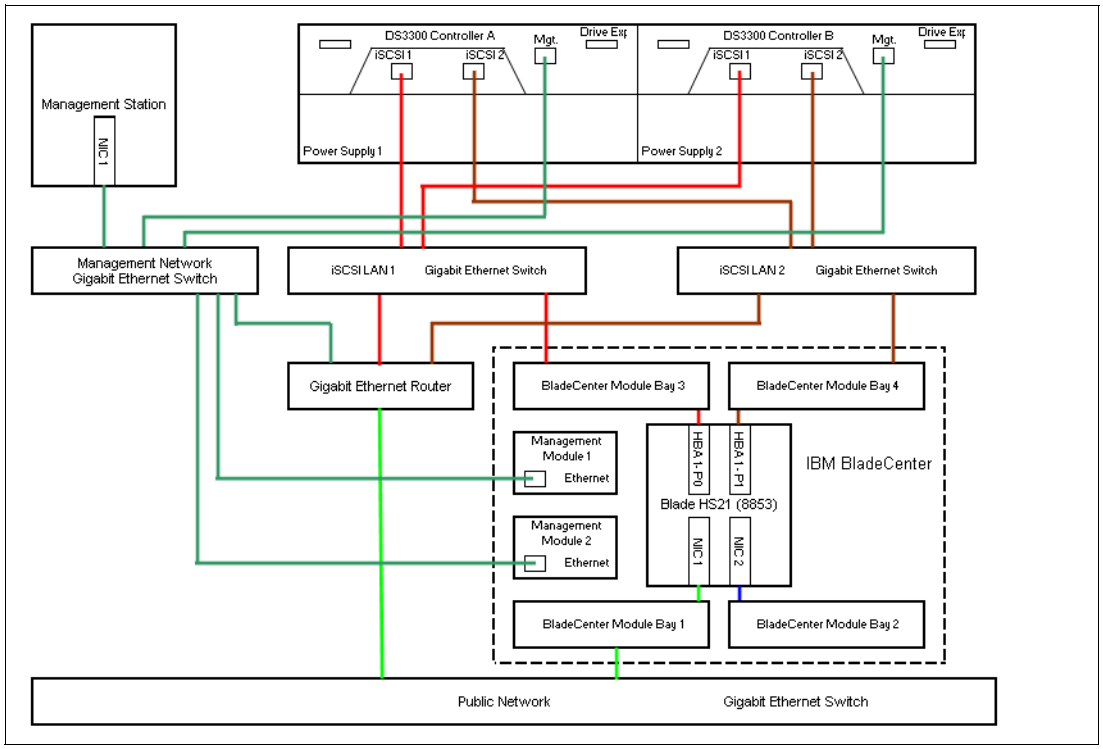


Figure 22-1 Network diagram

Table 22-1 shows a list of network addresses used in this sample configuration.

Table 22-1 Network configuration details

Device	Description	Value	Comments
BladeCenter	Management Module	172.18.0.125/16	
	Nortel Layer 2/3 Gigabit Ethernet Switch Bay 1	172.18.0.127/16	
	Nortel Layer 2/3 Gigabit Ethernet Switch Bay 2	172.18.0.128/16	
	Nortel Layer 2/3 Gigabit Ethernet Switch Bay 3	172.18.0.129/16	
	Nortel Layer 2/3 Gigabit Ethernet Switch Bay 4	172.18.0.130/16	
Blade HS21 (8853)	Host name	hooghly.rivers.local	
	Onboard NIC 1	172.19.200.70/16	eth0 - Public LAN
	Onboard NIC 2	Not used	

Device	Description	Value	Comments
	Gigabit Ethernet daughter card 1- Port 0	172.16.200.70/16	eth2 - iSCSI LAN 1
	Gigabit Ethernet daughter card 1- Port 1	172.17.200.70/16	eth3 - iSCSI LAN 2
System Storage DS3300	Controller A - Mgt.	172.18.3.1/16	Management LAN
	Controller A - iSCSI Port 1	172.16.3.1/16	iSCSI LAN 1
	Controller A - iSCSI Port 2	172.17.3.1/16	iSCSI LAN 2
	Controller B - Mgt.	172.18.3.2/16	Management LAN
	Controller B- iSCSI Port 1	172.16.3.2/16	iSCSI LAN 1
	Controller B - iSCSI Port 2	172.17.3.2/16	iSCSI LAN 2
Management Station	Host name	amazon.rivers.local	
	NIC 1	172.18.9.10/16	

22.2 Installing Red Hat Enterprise Linux 5 Update 2

Follow the operating system installation instructions that are available for each IBM BladeCenter blade and IBM System x server. The installation guides can be found in the “Install/use” section of each product’s support Web sites.

For our IBM BladeCenter HS21 (8853) with Red Hat Enterprise Linux, we used the instructions found at:

https://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-5070760&brandind=5000020

Although the instructions on the Web site are for Red Hat Enterprise Linux 5, the installation was performed using Red Hat Enterprise Linux 5 with Update 2. No additional packages were chosen during the installation.

In Table 22-1 on page 530, we can see that the BladeCenter HS21 has three network interfaces. Interface eth0 is used to access the public network, and eth2 and 3 are dedicated for iSCSI connections to the System Storage DS3300.

22.3 Installing and configuring the Open iSCSI software initiator

Red Hat Enterprise Linux includes an iSCSI software initiator. The software initiator is included in the installation media in the package directory as an RPM file.

To install the software initiator, mount the installation media and install the iSCSI software initiator RPM file. Example 22-1 shows the package installed from an NFS mounted file system.

Example 22-1 Installation of the iSCSI software initiator RPM file

```
[root@hooghly ~]# mount -t nfs 172.18.0.1:/srv /mnt
[root@hooghly ~]# cd /mnt/ftp/rhel_5.2x86_64/server/
[root@hooghly RPMS]# rpm -ivh iscsi-initiator-utils-6.2.0.868-0.7.el5.x86_64.rpm
Preparing...                ########################################### [100%]
   1:iscsi-initiator-utils  ########################################### [100%]
[root@hooghly RPMS]#
```

After the iSCSI software initiator is installed, we need to check two configuration files in the /etc/iscsi directory.

/etc/iscsi/initiatorname.iscsi

The first file is /etc/iscsi/initiatorname.iscsi. After the iSCSI daemon is started for the first time, this file will be populated with the iSCSI qualified name of the host.

Example 22-2 shows the file with an automatically generated IQN. This IQN is regenerated to the same value when the software initiator gets reinstalled.

Example 22-2 /etc/iscsi/nitiatorname.iscsi after first iSCSI daemon start

```
[root@hooghly ~]# cat /etc/initiatorname.iscsi
InitiatorName=iqn.1994-05.com.redhat:e61e98c51e1
[root@hooghly ~]#
```

An alternative way to generate the IQN is with the **iscsi-iname** command. This command will generate a unique IQN and output it to the console. This IQN might be used to define the IQN of a given host. Example 22-3 shows how to use the **iscsi-iname** command. The output can be redirected into /etc/initiatorname.iscsi.

Example 22-3 iscsi-iname

```
[root@hooghly ~]# iscsi-iname
iqn.1994-05.com.redhat:e61e98c51e1
[root@hooghly ~]#
```

/etc/iscsi/iscsi.conf

The second configuration file is /etc/iscsi/iscsi.conf and is used to configure the iSCSI software initiator itself. This configuration file is well documented in its man page and with comments in the configuration file itself. After installation, the file contains only comments and needs to be modified at least with a discovery address. Without the discovery address, the daemon will fail to start. Edit /etc/iscsi/iscsi.conf and make the modifications shown in Example 22-4 on page 533.

Example 22-4 Configuration of the iSCSI software initiator

```
DiscoveryAddress=172.16.3.1:3260

HeaderDigest=always
DataDigest=always

Targetname=iqn.1992-01.com.lsi:1535.00000000000000000000000000000000
        Enabled=yes
        ConnFailTimeout=15
```

The field DiscoveryAddress contains the IP address of one of the iSCSI interfaces of the DS3300; only one address has to be specified. The response of the sendTargets host command will contain the addresses of all four iSCSI interfaces.

HeaderDigest and DataDigest are both set to always. This option requires the use of CRC checks to detect corrupt iSCSI PDUs. This will increase the stability of the iSCSI connection.

Targetname specifies the IQN of the DS3300. The last two options (Enabled and ConnFailTimeout) are local options that are only valid for the specific Targetname (IQN iqn.1992-01.com.lsi:1535.00000000000000000000000000000000 in this case). Global settings start at the beginning of a line; local settings are always indented with a space or tab. A local setting follows a global setting and overwrites an already defined global setting.

Start iSCSI daemon

After the configuration modifications are done, start the iSCSI daemon to generate the iSCSI qualified name of the host. This is only required if the IQN was not set manually.

Example 22-5 shows the output of the command **`service iscsi start`** that is used to load the kernel modules for the iSCSI support.

Example 22-5 Load iSCSI driver

```
[root@hooghly ~]# service iscsi start
Checking iscsi config:                                     [  OK  ]
Loading iscsi driver:                                      [  OK  ]
Starting iscsid:                                           [  OK  ]
[root@hooghly ~]#
```

Some status and error information is logged to stdout in the host as well as to the standard log files. The IQN can be retrieved from /etc/iscsi/initiatorname.iscsi to be used during host mapping on the DS3300.

22.4 Installing the RDAC failover driver

The RDAC failover driver is used. At the time of writing, in iSCSI environments, this driver is not supported if an HBA of any kind (iSCSI, Fibre Channel, or SAS) and the iSCSI software initiator are installed on the same host, since the RDAC driver installation will detect the HBA first, and not use the software initiator. Therefore, there must not be any HBAs installed.

22.4.1 Preparing the host to install the RDAC driver

The RDAC driver has some prerequisite packages that must be installed first. These packages can be installed as part of the Linux operating system installation or can be installed afterwards. These are the kernel source, compiler, and libraries:

- glibc-headers
- glibc-devel
- gcc
- kernel-devel
- kernel-headers

Ensure that these packages are installed before attempting the RDAC driver installation. If not installed, the MPP driver installation will fail. Check the kernel used on the system and install the appropriate kernel source. Most systems will use the smp kernel, but the uniprocessor kernel or the huge memory kernel are also possible. Choose the appropriate package and install the kernel source.

22.4.2 Installing the RDAC driver

Do these steps:

1. Download the driver from `http://www.lsi.com/rdac/ds3000.html` and refer to 26.7, “Installing Linux RDAC in RHEL5” on page 663 and 17.4, “Installing RDAC for Linux” on page 408 for more details about installing this driver.
2. Unpack the driver archive into a local directory, for example, /root, as shown in Example 22-6. Note that the actual archive name changes to match the actual version.

   Example 22-6 Unpack archive

   ```
   [root@hooghly ~]# tar -zxvf rdac-LINUX-09.03.0B05.0023-source.tar.gz
   ```

3. A subdirectory linuxrdac-<*version number*> is created. Change to this directory (Example 22-7).

   Example 22-7 Change to unpacked directory

   ```
   [root@hooghly ~]# cd linuxrdac-09.03.0C05.0023
   [root@hooghly linuxrdac-09.03.0C05.0023]#
   ```

4. If an older RDAC driver version is installed, run **`make uninstall`** in this directory. A message is output by **`make`** to notify you if there is an old version installed on the system when an operation other than **`make uninstall`** or **`make clean`** is executed.
5. Run the **`make`** command followed by **`make install`** in the RDAC driver source directory to compile the driver, as shown in Example 22-8 on page 535. The compiler output has been abbreviated, as indicated by the three dots shown in this example.

   The **`make install`** phase scans for available HBAs. As already mentioned, no HBAs can be installed when RDAC is used to provide failover functionality for the iSCSI software initiator together with a DS3300. It is important that the iSCSI software adapter is found and the question is asked: **“**`Do you want MPP to manage an iSCSI storage array?`**”** Answer **“**`Do you want to continue (yes or no)?`**”** with `yes` to proceed. If an HBA is found, the RDAC will detect this instead of the iSCSI software initiator; remove the HBA and try again.

Example 22-8 Install RDAC driver

```
[root@hooghly linuxrdac-09.01.C5.19]# make && make install
make V=1 -C/lib/modules/2.6.18-92.el5/build  M=/root/linuxrdac-09.03.0C05.0023
MODVERDIR=/lib/modules/2.6.18-92.el5/build/.tmp_versions
SUBDIRS=/root/linuxrdac-09.03.0C05.0023 modules
make[1]: Entering directory `/usr/src/kernels/2.6.18-92.el5-x86_64'
.
.
.
Checking Host Adapter Configuration...
iSCSI software initiator found. Do you want MPP to manage an iSCSI storage
array?
Do you want to continue (yes or no) ? y
Wait while we modify the system configuration files.
Your kernel version is 2.6.18-92.el5
Preparing to install MPP driver against this kernel version...
Generating module dependencies...
Creating new MPP initrd image...
You must now edit your boot loader configuration file, /boot/grub/menu.lst, to
add a new boot menu, which uses mpp-2.6.18-92.el5.img as the initrd image.
Now Reboot the system for MPP to take effect.
The new boot menu entry should look something like this (note that it
may vary with different system configuration):

        ...

                title Red Hat Linux (2.6.18-92.el5) with MPP support
                root (hd0,5)
                kernel /vmlinuz-2.6.18-92.el5 ro root=LABEL=RH9
                initrd /mpp-2.6.18-92.el5.imgmg
        ...
MPP driver package has been sucessfully installed on your system.
[root@hooghly linuxrdac-09.01.C5.19]#
```

6. As shown in the output, the boot configuration needs to be modified to use the new created MPP initial RAM disk.

   Edit /boot/grub/menu.lst and make a copy of the existing boot entry to start the operating system. Modify the name of the initial RAMdisk to the name of the RAMdisk that contains the RDAC driver support. Usually the word initrd is replaced with RDAC. Example 22-9 shows a partial content of the boot loader GRUB config. The first section is used to boot the system with the initial RAMdisk that contains the RDAC driver support, and the second section uses the standard initial RAMdisk. We recommend using the MPP initial RAMdisk entry as the default boot entry in the configuration. Usually this is achieved by having this boot configuration as the first boot configuration defined in the config file.

Example 22-9 GRUB configuration file menu.lst

```
title Red Hat Enterprise Linux (2.6.18-92.el5) [MPP]
        root (hd0,2)
        kernel /vmlinuz-2.6.18-92.el5 ro root=LABEL=/ rhgb quiet
        initrd /mpp-2.6.18-92.el5.img

title Red Hat Enterprise Linux (2.6.18-92.el5)
        root (hd0,2)
        kernel /vmlinuz-2.6.18-92.el5 ro root=LABEL=/ rhgb quiet
```

```
initrd /initrd-2.6.18-92.el5.img
```

7. Do not reboot until after the DS3300 is configured, as described in the next section.

22.5 Setting up the DS3300 logical drives and host mapping

The basic network setup needs to be done, as described in 5.4.5, "(Optional) Changing the network configuration" on page 76. The iSCSI interface ports of the DS3300 must be configured with the correct IP addresses, as described in 13.2.3, "Configure iSCSI Host Ports" on page 291. You might also want to install a separate management station to run the DS3000 Storage Manager if Storage Manager is not installed in the host. See Chapter 6, "Installing IBM System Storage DS3000 Storage Manager" on page 81 for details on installing Storage Manager.

In this section, we describe the configuration that was done to provide two logical drives configured on an IBM System Storage DS3300 to the host HOOGHLY running Red Hat Enterprise Linux 5. We show these steps using the CLI; however, the same setup can be also done with the GUI.

1. A command script file, called hooghly.cmd, shown in Example 22-10, contains the commands to create logical drives in existing arrays and set up logical drive mapping. For formatting reasons in this book, we use backslashes to show line breaks; however, you should remove these line break characters; every line is a single command.

Example 22-10 hooghly.cmd command file

```
show "Create a host group with the name Asia";
create hostGroup userLabel="Asia";

show "Create a host Hooghly that is member of host group Asia";
create host userLabel="Hooghly" hostGroup="Asia";

show "Create an iSCSI host port alias initiator";
create iscsiInitiator iscsiName="iqn.1994-05.com.redhat:e61e98c51e1"
userLabel="Hooghly-Initiator" host="Hooghly";

show "Create a logical drive in an existing array and do not map that logical
drive";
create logicaldrive array=1 userLabel="Hooghly-1" freeCapacityArea=1 capacity=5
GB \ mapping=none;
create logicaldrive array=2 userLabel="Hooghly-2" freeCapacityArea=1 capacity=7
GB \ mapping=none;

show "Map logical drives to an host";
set logicalDrive ["Hooghly-1"] logicalUnitNumber=0 host="Hooghly";
set logicalDrive ["Hooghly-2"] logicalUnitNumber=1 host="Hooghly";
```

2. The script file is executed by **SMcli** using the syntax shown in Example 22-11. The output is directed to a file called hooghly.out; check it for errors. It will contain the comments specified with the **show** command to help debug any errors.

Example 22-11 SMcli executes the script

```
amazon:~ # SMcli -n DS3300 -f hooghly.cmd -o hooghly.out -S
amazon:~ #
```

Now the logical drives and host mapping configuration is done. We will later define security for access once the basic connection is working.

22.5.1 Finishing the configuration

To load the iSCSI driver during reboot, it is necessary to modify the service database. Run the command **chkconfig iscsi on** to enable the load of the iSCSI driver during system start, as shown in Example 22-12.

Example 22-12 Enable iSCSI during system start

```
[root@hooghly ~]# chkconfig iscsi on
[root@hooghly ~]#
```

Reboot the server to activate the MPP driver and get access to the logical drives configured on the DS3300.

Now the iSCSI driver and MPP should be loaded. To verify that the drivers are working, check these commands and files:

- **iscsi-ls**
- /proc/scsi/scsi
- /proc/mpp/*
- **fdisk -l**

iscsi-ls

The command **iscsi-ls** shows all detected iSCSI target portals in the host. Example 22-13 shows the output of the command with all four iSCSI portal IP addresses of the DS3300 discovered. Each portal shows an active iSCSI session (SESSION STATUS: ESTABLISHED).

Example 22-13 iscsi-ls

```
[root@hooghly ~]# iscsi-ls
*******************************************************************************
SFNet iSCSI Driver Version ...4:0.1.11-4(15-Jan-2007)
*******************************************************************************
TARGET NAME    : iqn.1992-01.com.lsi:1535.000000000000000000000000000
000
TARGET ALIAS            :
HOST ID                 : 14
BUS ID                  : 0
TARGET ID               : 0
TARGET ADDRESS          : 172.16.3.2:3260,2
SESSION STATUS          : ESTABLISHED AT Mon Jun 11 19:20:36 EDT 2007
SESSION ID              : ISID 00023d000001 TSIH 8210
*******************************************************************************
TARGET NAME    : iqn.1992-01.com.lsi:1535.000000000000000000000000000
000
TARGET ALIAS            :
HOST ID                 : 15
BUS ID                  : 0
TARGET ID               : 0
TARGET ADDRESS          : 172.17.3.2:3260,2
SESSION STATUS          : ESTABLISHED AT Mon Jun 11 19:20:36 EDT 2007
SESSION ID              : ISID 00023d000002 TSIH 820f
*******************************************************************************
```

```
TARGET NAME    : iqn.1992-01.com.lsi:1535.000000000000000000000000000
000
TARGET ALIAS            :
HOST ID                 : 16
BUS ID                  : 0
TARGET ID               : 0
TARGET ADDRESS          : 172.16.3.1:3260,1
SESSION STATUS          : ESTABLISHED AT Mon Jun 11 19:20:36 EDT 2007
SESSION ID              : ISID 00023d000003 TSIH 1e8
************************************************************************
TARGET NAME    : iqn.1992-01.com.lsi:1535.000000000000000000000000000
000
TARGET ALIAS            :
HOST ID                 : 17
BUS ID                  : 0
TARGET ID               : 0
TARGET ADDRESS          : 172.17.3.1:3260,1
SESSION STATUS          : ESTABLISHED AT Mon Jun 11 19:20:36 EDT 2007
SESSION ID              : ISID 00023d000004 TSIH 1e9
************************************************************************
[root@hooghly ~]#
```

/proc/scsi/scsi

If all four iSCSI ports of the DS3300 are used to connect to servers, every logical drive that is mapped to the host will initially be seen four times before the MPP driver combines them and shows only one disk to the user. This information can be found in the file /proc/scsi/scsi, which shows SCSI devices that are detected by the Linux kernel. See Example 22-14. SCSI port 5, 6, 7, and 8 are the iSCSI connections to the iSCSI ports of the DS3300. Both LUNs (0 and 1) appear on each SCSI port as well as the Access LUN (also called Universal Xport). SCSI port 9 represents the two logical drives configured on the DS3300 that are seen by the user as physical disks.

Example 22-14 /proc/scsi/scsi

```
[root@hooghly ~]# cat /proc/scsi/scsi
Attached devices:
Host: scsi0 Channel: 00 Id: 00 Lun: 00
  Vendor: TEAC     Model: FD-05PUB         Rev: 2000
  Type:   Direct-Access                    ANSI SCSI revision: 02
Host: scsi1 Channel: 00 Id: 00 Lun: 00
  Vendor: TEAC     Model: CD-224E          Rev: 2.9B
  Type:   CD-ROM                           ANSI SCSI revision: 02
Host: scsi2 Channel: 00 Id: 00 Lun: 00
  Vendor: ServeRA  Model: Drive 1          Rev: V1.0
  Type:   Direct-Access                    ANSI SCSI revision: 02
Host: scsi2 Channel: 00 Id: 01 Lun: 00
  Vendor: ServeRA  Model: Drive 2          Rev: V1.0
  Type:   Direct-Access                    ANSI SCSI revision: 02
Host: scsi5 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM      Model: 1726-3xx  FAStT  Rev: 0250
  Type:   Direct-Access                    ANSI SCSI revision: 05
Host: scsi6 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM      Model: 1726-3xx  FAStT  Rev: 0250
  Type:   Direct-Access                    ANSI SCSI revision: 05
Host: scsi8 Channel: 00 Id: 00 Lun: 00
```

```
  Vendor: IBM      Model: 1726-3xx  FAStT  Rev: 0250
  Type:   Direct-Access                    ANSI SCSI revision: 05
Host: scsi7 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM      Model: 1726-3xx  FAStT  Rev: 0250
  Type:   Direct-Access                    ANSI SCSI revision: 05
Host: scsi5 Channel: 00 Id: 00 Lun: 01
  Vendor: IBM      Model: 1726-3xx  FAStT  Rev: 0250
  Type:   Direct-Access                    ANSI SCSI revision: 05
Host: scsi6 Channel: 00 Id: 00 Lun: 01
  Vendor: IBM      Model: 1726-3xx  FAStT  Rev: 0250
  Type:   Direct-Access                    ANSI SCSI revision: 05
Host: scsi8 Channel: 00 Id: 00 Lun: 01
  Vendor: IBM      Model: 1726-3xx  FAStT  Rev: 0250
  Type:   Direct-Access                    ANSI SCSI revision: 05
Host: scsi7 Channel: 00 Id: 00 Lun: 01
  Vendor: IBM      Model: 1726-3xx  FAStT  Rev: 0250
  Type:   Direct-Access                    ANSI SCSI revision: 05
Host: scsi5 Channel: 00 Id: 00 Lun: 31
  Vendor: IBM      Model: Universal Xport  Rev: 0250
  Type:   Direct-Access                    ANSI SCSI revision: 05
Host: scsi6 Channel: 00 Id: 00 Lun: 31
  Vendor: IBM      Model: Universal Xport  Rev: 0250
  Type:   Direct-Access                    ANSI SCSI revision: 05
Host: scsi8 Channel: 00 Id: 00 Lun: 31
  Vendor: IBM      Model: Universal Xport  Rev: 0250
  Type:   Direct-Access                    ANSI SCSI revision: 05
Host: scsi7 Channel: 00 Id: 00 Lun: 31
  Vendor: IBM      Model: Universal Xport  Rev: 0250
  Type:   Direct-Access                    ANSI SCSI revision: 05
Host: scsi9 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM      Model: VirtualDisk      Rev: 0250
  Type:   Direct-Access                    ANSI SCSI revision: 05
Host: scsi9 Channel: 00 Id: 00 Lun: 01
  Vendor: IBM      Model: VirtualDisk      Rev: 0250
  Type:   Direct-Access                    ANSI SCSI revision: 05
[root@hooghly ~]#
```

/proc/mpp

The MPP driver adds a directory in the proc file system that allows us to verify what paths are active and which devices and LUNs are discovered. Example 22-15 shows the directory structure of /proc/mpp with targets and LUNs that are handled by the MPP driver.

Example 22-15 /proc/mpp/ directory structure

```
[root@hooghly ~]# ls -R /proc/mpp/
/proc/mpp/:
DS3300

/proc/mpp/DS3300:
controllerA  controllerB  virtualLun0  virtualLun1

/proc/mpp/DS3300/controllerA:
iscsi_sfnet_h7c0t0  iscsi_sfnet_h8c0t0

/proc/mpp/DS3300/controllerA/iscsi_sfnet_h7c0t0:
```

```
LUN0  LUN1  UTM_LUN31

/proc/mpp/DS3300/controllerA/iscsi_sfnet_h8c0t0:
LUN0  LUN1  UTM_LUN31

/proc/mpp/DS3300/controllerB:
iscsi_sfnet_h5c0t0  iscsi_sfnet_h6c0t0

/proc/mpp/DS3300/controllerB/iscsi_sfnet_h5c0t0:
LUN0  LUN1  UTM_LUN31

/proc/mpp/DS3300/controllerB/iscsi_sfnet_h6c0t0:
LUN0  LUN1  UTM_LUN31
[root@hooghly ~]#
```

fdisk -l

Use the **fdisk** command to show a list of available devices, as shown in Example 22-16. The two logical drives accessed over iSCSI appear here as /dev/sdd and /dev/sde with no partition information.

Example 22-16 List of drives

```
[root@hooghly ~]# fdisk -l

Disk /dev/sdb: 85.8 GB, 85899345920 bytes
255 heads, 63 sectors/track, 10443 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/sdb1               1        2550    20482843+   7  HPFS/NTFS
/dev/sdb2   *        2551        2567      136552+  83  Linux
/dev/sdb3            2568        2586      152617+  83  Linux
/dev/sdb4            2587       10443    63111352+   f  W95 Ext'd (LBA)
/dev/sdb5            2587        3370     6297448+  83  Linux
/dev/sdb6            3371        3501     1052226   82  Linux swap
/dev/sdb7            3502        4776    10241406   83  Linux

Disk /dev/sdc: 60.6 GB, 60689481728 bytes
255 heads, 63 sectors/track, 7378 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/sdc1   *           1        7378    59263753+   6  FAT16

Disk /dev/sdd: 5368 MB, 5368709120 bytes
166 heads, 62 sectors/track, 1018 cylinders
Units = cylinders of 10292 * 512 = 5269504 bytes

Disk /dev/sdd doesn't contain a valid partition table

Disk /dev/sde: 7516 MB, 7516192768 bytes
232 heads, 62 sectors/track, 1020 cylinders
Units = cylinders of 14384 * 512 = 7364608 bytes

Disk /dev/sde doesn't contain a valid partition table
```

```
[root@hooghly ~]#
```

Now you can partition and format the drives for use.

22.6 Changing the configuration at run time

In this section, we briefly explain two MPP commands that are useful under some circumstances. These are:

- **mppBusRescan**
- **mppUpdate**

mppBusRescan

When new logical drives are mapped to the host at run time, use the **mppBusRescan** command to scan for new LUNs available to this host. This command is similar to the Windows **hotadd** command. See Example 22-17.

Example 22-17 mppBusRescan command

```
[root@hooghly ~]# mppBusRescan
Starting new devices re-scan...
scan iSCSI software initiator host /sys/class/scsi_host/host8...
        found 8:0:0:2
scan iSCSI software initiator host /sys/class/scsi_host/host7...
        found 7:0:0:2
scan iSCSI software initiator host /sys/class/scsi_host/host5...
        found 5:0:0:2
scan iSCSI software initiator host /sys/class/scsi_host/host6...
        found 6:0:0:2
run /usr/sbin/mppUtil -s busscan...
scan mpp virtual host /sys/class/scsi_host/host9...
        found 9:0:0:2->/dev/sdf
/usr/sbin/mppBusRescan is completed.
[root@hooghly ~]#
```

mppUpdate

Use the **mppUpdate** command to create a new MPP initial RAM disc. This could be required in case a driver for the SCSI controller, where the disk with the root file system is attached, requires an update. See Example 22-18.

Example 22-18 mppUpdate

```
[root@hooghly ~]# mppUpdate
Creating new MPP initrd image...
[root@hooghly ~]#
```

22.7 Enhancing connection security

After everything is working and tested, you should implement security for the iSCSI connection; essentially, this means configuring initiator and target authentication. Initiator authentication means that an initiator must prove its identity with a password that is known by the target when the initiator attempts access. Target authentication is the opposite; the target authenticates itself to the initiator with a password.

Since an iSCSI qualified name can be modified within Storage Manager, this does not protect against spoofing of the qualified name, and hence security can be compromised. We describe the steps required to set up initiator and target authentication.

Like the QLogic iSCSI HBA described in Chapter 21, “iSCSI configuration 1 - Windows 2003 and 2008 with a QLogic QMC4052 HBA” on page 481, the iSCSI software initiator requires initiator authentication as a prerequisite for target authentication. Follow these steps:

1. Edit the configuration file /etc/iscsi/iscsi.conf of the iSCSI software initiator. Define the incoming and outgoing user names and passwords in this file. Incoming means the target has to authenticate itself against the initiator, and is also called target authentication. Outgoing represents the initiator authentication. The initiator has to authenticate against a target.

   Note: The initiator and target CHAP secrets cannot be identical.

   Incoming and outgoing user names are limited to valid IQNs by the DS3300 defined as host ports.

   Example 22-19 shows the /etc/iscsi/iscsi.conf file with the incoming and outgoing account details. The incoming account was configured as a local option for the target and not a global option. Other subsystems may use a different password.

Example 22-19 Configuration of the iSCSI software initiator

```
HeaderDigest=always
DataDigest=always
OutgoingUsername=iqn.1994-05.com.redhat:e61e98c51e1
OutgoingPassword=b1234567890h

Targetname=iqn.1992-01.com.lsi:1535.00000000000000000000000000000000
        Enabled=yes
        IncomingUsername=iqn.1992-01.com.lsi:1535.000000000000000000000000000000000
        IncomingPassword=a1234567890t
        ConnFailTimeout=15
```

2. Shut down the server until the DS3300 is also configured.
3. Use the Storage Manager CLI commands (**`set iscsiInitiator`**) shown in Example 22-20 to set up the CHAP secret (Challenge Handshake Authentication Protocol) for the already defined host ports of host HOOGHLY.

   Example 22-20 SMcli commands to set CHAP secrets for initiators

```
amazon:~ # SMcli -n DS3300 -c "set iscsiInitiator [\"Hooghly-Initiator\"]
host=\"Hooghly\" chapSecret=\"b1234567890h\"; " -S
amazon:~ #
```

Note: CHAP (RFC1944) is the most basic level of iSCSI security available.

4. Clarify if there are any initiators without target authentication configured that access the DS3300. In that case, use the command in Example 22-21; otherwise, use the command shown in Example 22-22.

Example 22-21 SMcli - Set target authentication - CHAP only

```
amazon:~ # SMcli -n DS3300 -c "set iscsiTarget
<\"iqn.1992-01.com.lsi:1535.000000000000000000000000000000000\">
authenticationMethod=chap chapSecret=\"a01234567890t\";" -S
amazon:~ #
```

Example 22-22 SMcli - Set target authentication - CHAP and no CHAP

```
amazon:~ # SMcli -n DS3300 -c "set iscsiTarget
<\"iqn.1992-01.com.lsi:1535.000000000000000000000000000000000\">
authenticationMethod=none authenticationMethod=chap
chapSecret=\"a01234567890t\";" -S
amazon:~ #
```

5. After modifying the configuration on the DS3300, reboot the server to get access to the logical drives by using initiator and target authentication.

23

iSCSI configuration 3 - Windows 2003 cluster

In this chapter, we discuss a sample configuration that shows how to connect logical drives configured on an IBM System Storage DS3000 to clustered servers running the Windows Server 2003 operating system with the Microsoft iSCSI Software initiator installed.

The server is an IBM System x3850 running VMware V3.1. At the time of writing, native failover support with the DS3300 is not available, therefore we use the iSCSI software initiator, along with the DS3000 failover driver.

Our cluster will consist of two virtual machines with the host names Mara and Ora.

Figure 23-1 on page 546 gives the cabling and networking details of the configuration.

At the time of writing, clustered servers running Windows Server 2008 is supported, but it is beyond the intended scope of this book to demonstrate the steps involved in setting this up. For more information, see Chapter 4 of *Installation and Support Guide for Microsoft Windows Server 2003, Linux, Novell NetWare, and VMware ESX - IBM System Storage DS3200, DS3300, DS3400, BladeCenter Boot Disk (Type 1726) Guide*, which can be found at:

https://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-5069917&brandind=5000028

© Copyright IBM Corp. 2007, 2009. All rights reserved.

23.1 Network configuration

A diagram of the general network setup used is shown in Figure 23-1.

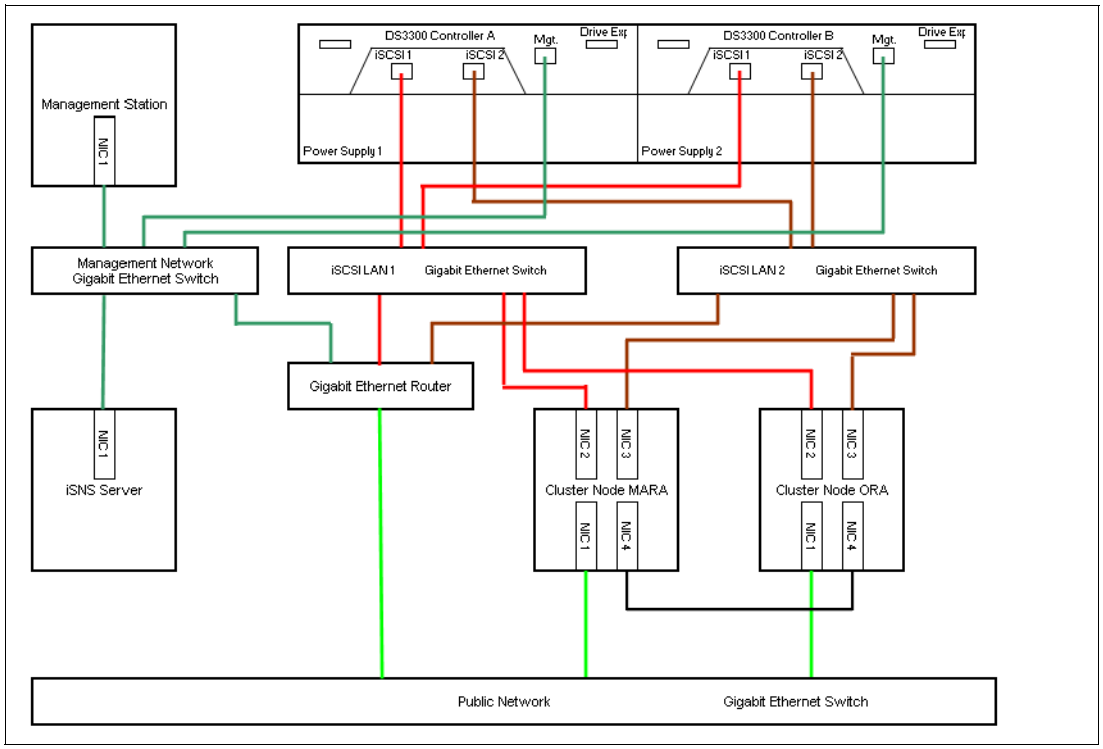


Figure 23-1 Network diagram

Table 23-1 shows a list of network addresses used in this sample configuration.

Table 23-1 Network configuration details

Device	Description	Value	Comments
VMware server	Management interface	172.18.0.55/16	Management LAN
Virtual machine 1	Host name	mara.rivers.local	
	NIC 1	172.19.200.91/16	Public LAN
	NIC 2	172.16.200.91/16	iSCSI LAN 1
	NIC 3	172.17.200.91/16	iSCSI LAN 2
	NIC 4	10.0.0.1/8	Private LAN
Virtual machine 2	Host name	ora.rivers.local	
	NIC 1	172.19.200.92/16	Public LAN
	NIC 2	172.16.200.92/16	iSCSI LAN 1
	NIC 3	172.17.200.92/16	iSCSI LAN 2
	NIC 4	10.0.0.2/8	Private LAN
Microsoft Cluster	Cluster name	south-africa.rivers.lo-cal	

Device	Description	Value	Comments
	Cluster IP	172.19.200.90	Public LAN
System Storage DS3300	Controller A - Mgt.	172.18.3.1/16	Management LAN
	Controller A - iSCSI Port 1	172.16.3.1/16	iSCSI LAN 1
	Controller A - iSCSI Port 2	172.17.3.1/16	iSCSI LAN 2
	Controller B - Mgt.	172.18.3.2/16	Management LAN
	Controller B- iSCSI Port 1	172.16.3.2/16	iSCSI LAN 1
	Controller B - iSCSI Port 2	172.17.3.2/16	iSCSI LAN 2
iSNS	Service IP	172.18.8.10/16	Management LAN

23.2 Operating system installation

Perform a standard installation of Microsoft Windows Server 2003 Enterprise Edition on each cluster node. Apply Service Pack 2, as it is required as a prerequisite for hotfix KB932755. This hotfix is a requirement for the RDAC failover driver. Install hotfix KB932755, which can be downloaded from:

http://www.microsoft.com/downloads/details.aspx?FamilyID=dead950b-da47-4dc2-abcd-40f61bd4d1aa&DisplayLang=en

When installing on actual System x servers, refer to the installation guide for that server for more details about installing the operating system.

23.3 Preparing the cluster installation

Follow these steps:

1. After installing the operating system, review the checklists for Microsoft Server Cluster Services (MSCS) for specific recommendations. The checklists are available in the Windows 2003 online help. We assume you are familiar with MSCS, and provide summary information only here.
2. Configure the network interfaces that are required. Our configuration is shown in 23.1, “Network configuration” on page 546. We recommend using static IP addresses. After setting the IP addresses, go to My Network Places and rename the network connections to represent their functions. To do this, right-click each interface in turn and select **Rename**. This makes identifying each network easier.
3. Join both nodes to the same Windows domain. It is not supported to have Active Directory® installed on the cluster nodes; a non-clustered node is required. The domain membership is used to provide a domain account available for the cluster service account on all nodes. That is the account under which the cluster service will operate.

4. Create a domain account to be used later on the cluster nodes to run the cluster service. There are no special requirements for this account; all necessary rights will be granted to this account when the cluster service is set up on the first node.
5. Download the Microsoft Software initiator from:

   http://www.microsoft.com/downloads/details.aspx?FamilyID=12cb3c1a-15d6-4585-b385-befd1319f825&DisplayLang=en

   The user guide and release notes can be found in the same location.

23.4 Installing iSCSI software initiator

Install the Microsoft Software initiator onto each node by doing the following steps:

1. Run the installer package downloaded for your platform during preparation. Figure 23-2 shows an example package.

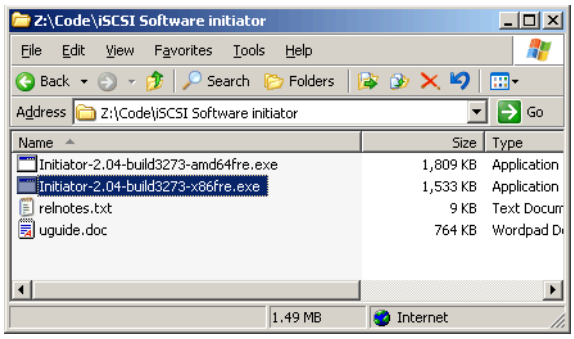


Figure 23-2 iSCSI software initiator install package

2. In the Welcome window (Figure 23-3), click **Next**.

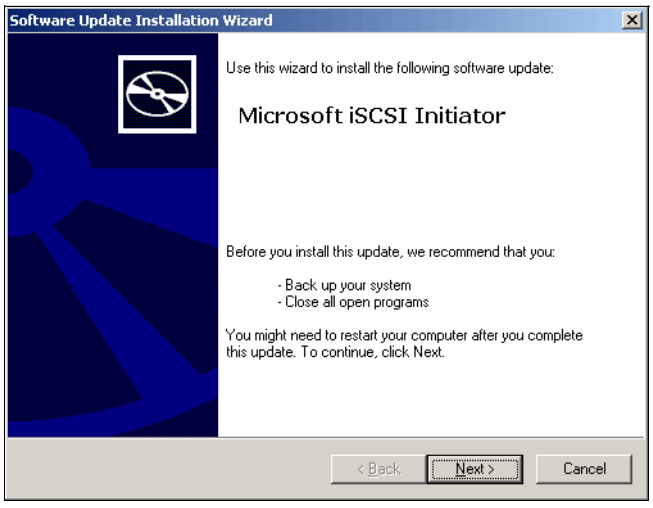


Figure 23-3 Software update installation wizard

3. Select **Initiator Service** and **Software Initiator**. See Figure 23-4 on page 549.

> **Note:** Do not install MPIO at this time. MPIO is included in the RDAC package together with the device specific module for the DS3300.

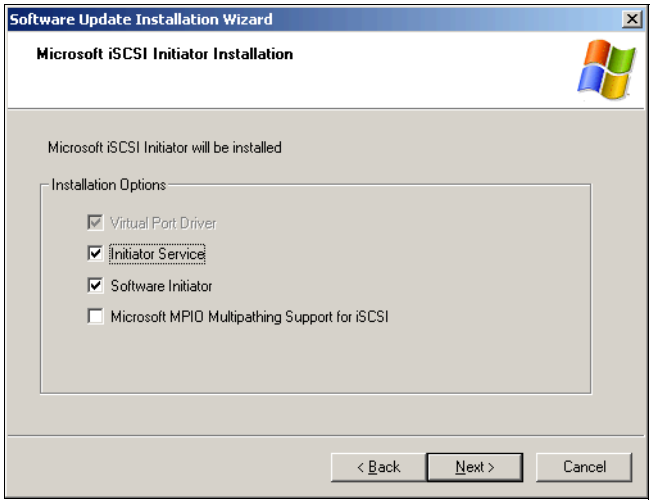


Figure 23-4 Installation options

4. Read the license agreement, select **I Agree**, and click **Next**.
5. The iSCSI software initiator will be installed now. When installation is complete, click **Finish** to close the installer. See Figure 23-5.

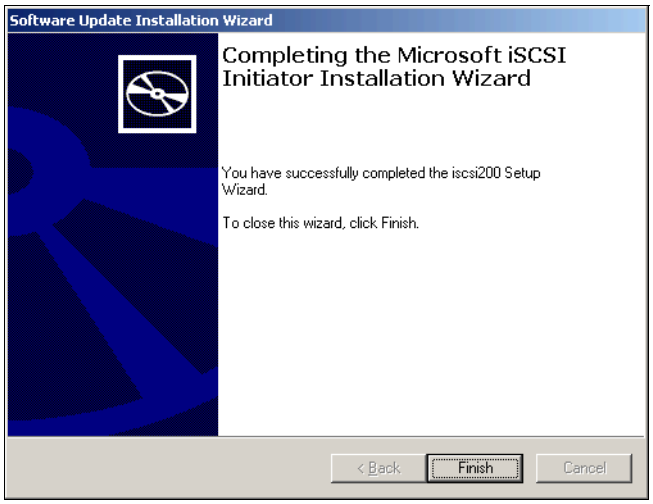


Figure 23-5 iSCSI software initiator installation complete

23.5 Collecting the iSCSI qualified name of each cluster node

An IQN is generated by the software initiator on each node. This IQN is required later to configure the DS3300. Do the following steps:

1. Start the iSCSI software initiator to collect this IQN by selecting **Start** → **Programs** → **Microsoft iSCSI Initiator** → **Microsoft iSCSI Initiator**. The software initiator opens, as shown in Figure 23-6. Click **Change**.

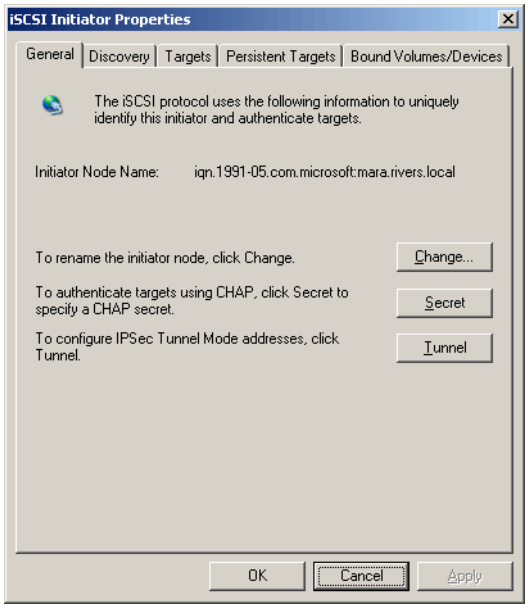


Figure 23-6 iSCSI software initiator properties

2. The IQN is displayed in Figure 23-7. Copy the IQN to the clipboard and click **Cancel** to close this dialog.

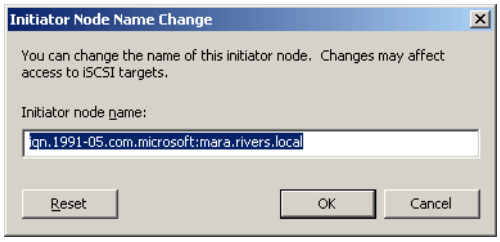


Figure 23-7 Initiator node name change

3. Click **OK** to close the iSCSI software initiator properties window.
4. Paste the IQN string into a text file for later use. You will have two IQN entries, one for each node.

23.6 Installing IBM System Storage DS3000 Storage Manager

Install IBM System Storage DS3000 Storage Manager on each node. Storage Manager contains the MPIO failover driver and the device specific module to handle multiple paths to the DS3300. In the Storage Manager setup, this component is included in the Host selection and is called RDAC. This is the minimum components that must be installed.

This section covers the installation of the host components only. Client software is available on a management system that was used to prepare the DS3300 and perform the initial setup. See 6.1, “Installing DS3000 Storage Manager on Microsoft Windows 2003 and 2008” on page 82 for full details about how to install Storage Manager. Do the following steps:

1. Run the DS3000 Storage Manager installation package.
2. Select the language to display the license agreement and click **OK**.
3. Read the introduction and click **Next**.
4. Read the copyright statement and click **Next**.
5. Read the license agreement and click **I accept the terms of this license agreement**. Click **Next**.
6. Choose a directory location to install Storage Manager. The default location C:\Program Files\IBM_DS3000 should work in most situations. Click **Next**.
7. In Figure 23-8, select the **Host** installation type or optionally choose a Full or Custom installation. Be sure that the RDAC and host utilities are included in the selection. Click **Next**.

Note: The RDAC component of Storage Manager requires both Windows Server 2003 Service Pack 2 and hotfix KB932755.



Figure 23-8 Storage manager installation type

8. Read the pre-installation summary (Figure 23-9) and click **Install**.

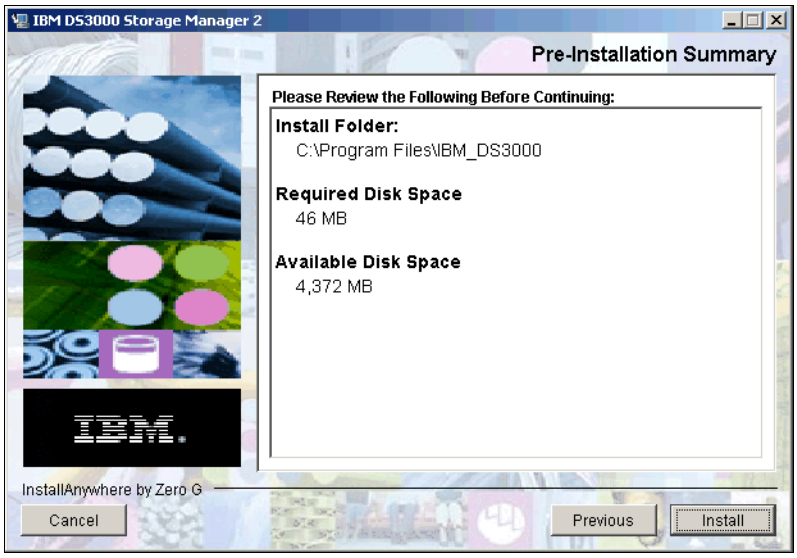


Figure 23-9 Storage manager pre-installation summary

9. The files will be copied to the system; you can monitor the progress as shown in Figure 23-10.

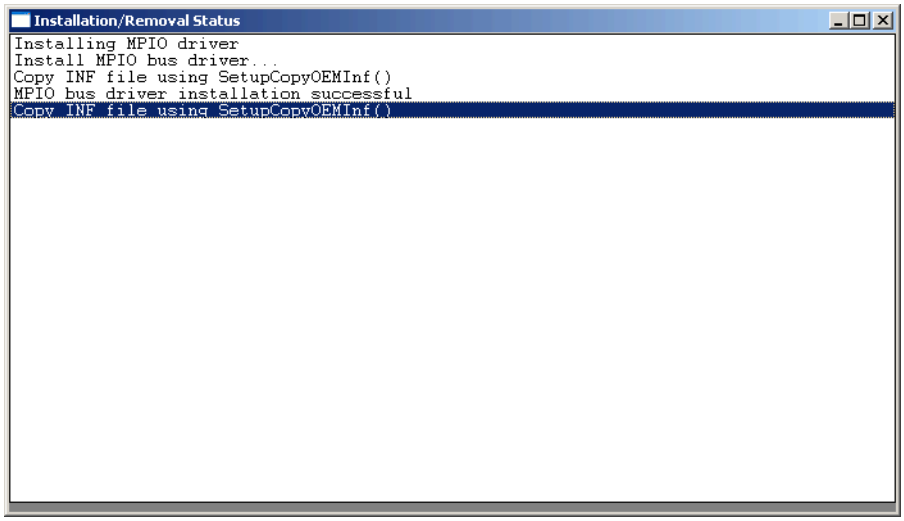


Figure 23-10 Installation progress

10. During the installation, you will might see a message (Figure 23-11) indicating that the device specific module for MPIO is not yet digitally signed. Click **Yes** to continue.

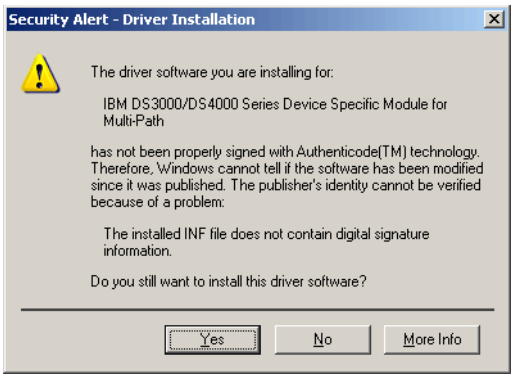


Figure 23-11 Security alert

11. After the MPIO installation is complete, the driver for the Universal Xport or UTM (access) LUN is installed. See Figure 23-12.

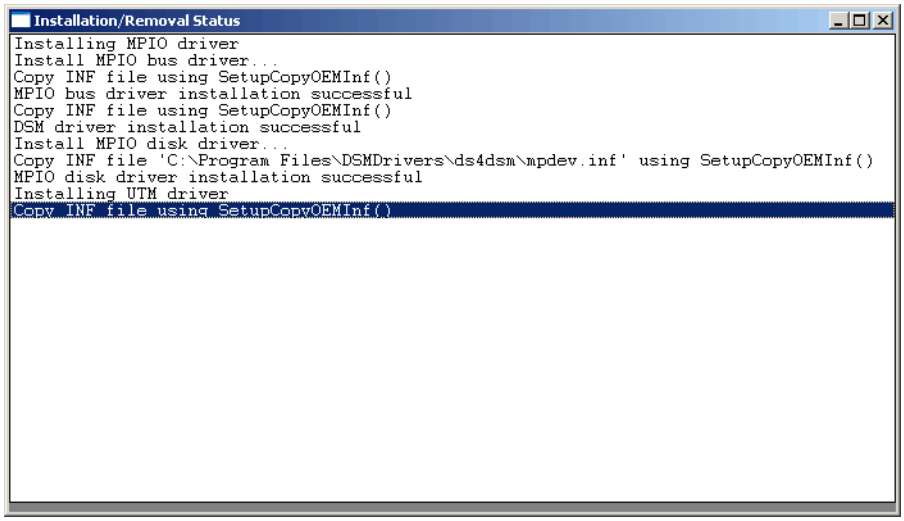


Figure 23-12 Installation progress

12. You will might see a message (Figure 23-13) indicating that the driver for the UTM LUN is not yet digitally signed. Click **Yes** to continue.

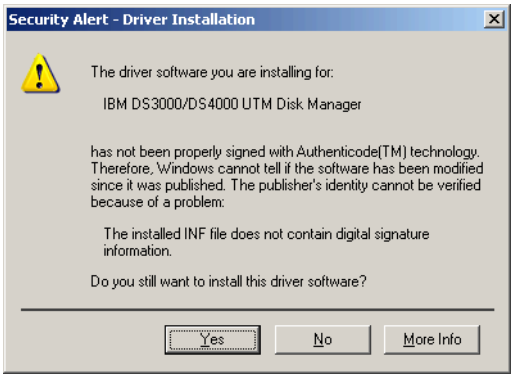


Figure 23-13 Security alert

13. After the installation is completed, you will be prompted to reboot, as shown in Figure 23-14. Click **Done**. The server will reboot to activate all installed drivers.

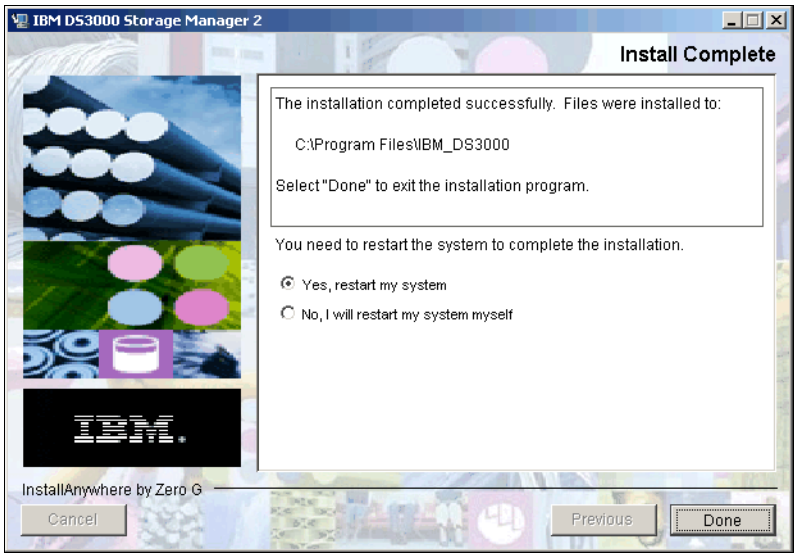


Figure 23-14 Storage manager install complete

23.7 Configuring the IBM System Storage DS3300

Configure the DS3300 to make the LUNs available that can be used as shared storage in the cluster, as follows:

1. The command script in Example 23-1 performs the initial configuration. There are two hosts, Mara and Ora in the host group, South-Africa, and each host has one iSCSI initiator defined. Four logical drives are then created in existing arrays, and they are mapped into the created host group in turn. The command script file is called south-africa.cmd.

   For formatting reasons in this book, we use backslashes to show line breaks; however, you should remove these line break characters; every line is a single command. The **show** commands are included to comment the actions that are performed; they are not necessary for the configuration.

Example 23-1 South-Africa command script

```
show "Create a host group with the name South-Africa";
create hostGroup userLabel="South-Africa";

show "Create a host Mara and Ora that are member of host group South-Africa";
create host userLabel="Mara" hostGroup="South-Africa";
create host userLabel="Ora" hostGroup="South-Africa";

show "Creata a iSCSI host port for each node";
create iscsiInitiator iscsiName="iqn.1991-05.com.microsoft:mara.rivers.local" \
userLabel="Mara-Iniatator" host="Mara";
create iscsiInitiator iscsiName="iqn.1991-05.com.microsoft:ora.rivers.local" \
userLabel="Ora-Initiator" host="Ora";

show "Create a logical drive in an existing array, no mapping done";
create logicaldrive array=1 userLabel="South-Africa-Quorum" freeCapacityArea=1 \
capacity=10 GB owner=a mapping=none;
```

```
create logicaldrive array=1 userLabel="South-Africa-eMail" freeCapacityArea=1 \
capacity=20 GB owner=b mapping=none;
create logicaldrive array=2 userLabel="South-Africa-DB" freeCapacityArea=1 capacity=50 \
GB owner=a mapping=none;

show "Create a logical drive in a new array";
create logicalDrive drive [7,12] RAIDLevel=0 \ userLabel="South-Africa-Temp-Spool" \
capacity=30 GB owner=b mapping=none;

show "Map logical drives to an host";
set logicalDrive ["South-Africa-Quorum"] logicalUnitNumber=0 hostGroup="South-Africa";
set logicalDrive ["South-Africa-eMail"] logicalUnitNumber=1 hostGroup="South-Africa";
set logicalDrive ["South-Africa-DB"] logicalUnitNumber=2 hostGroup="South-Africa";
set logicalDrive ["South-Africa-Temp-Spool"] logicalUnitNumber=3 \
hostGroup="South-Africa";

show “Define iSNS Server IP v4 Address”;
set storageSubsystem isnsIPV4ConfigurationMethod=static isnsIPV4Address=172.18.8.10;

Show “Activate iSNS server usage”;
set iscsiTarget <"iqn.1992-01.com.lsi:1535.000000000000000000000000000000000"> \
isnsRegistration=true;
```

2. Run the command script shown in Example 23-2. The output is directed to a file called south-africa.out for later verification. When this file has a size of zero, this means that every command in the command file was executed without errors. However, since we included show commands to provide some explanation of what the commands are doing, the output file will not be zero, because this output goes into the output file.

Example 23-2 SMcli executes the script

```
amazon:~ # SMcli -n DS3300 -f south-africa.cmd -o south-africa.out -S
amazon:~ #
```

23.8 Configuring the Microsoft iSCSI software initiator

The software initiator must be configured to access the created logical drives. In this section, we describe two different variants of target discovery. On node Mara, we show target discovery using iSNS, and on node Ora, target discovery will use the **sendTargets** command for discovery. Once the target discovery is done, the remaining configuration tasks are the same. If an iSNS server is available, we recommend that you use it; if no iSNS server is available, the sendTargets method should be used.

We are describing both methods for completeness; normally you would use the same method on all cluster nodes.

23.8.1 Configuring target discovery using iSNS server

Follow these steps (on node Mara in our example) to configure target discovery using iSNS:

1. Open the iSCSI software initiator by selecting **Start** → **Program Files** → **Microsoft Software Initiator** → **Microsoft Software Initiator**. A window opens, as shown in Figure 23-15. Click the **Discovery** tab.

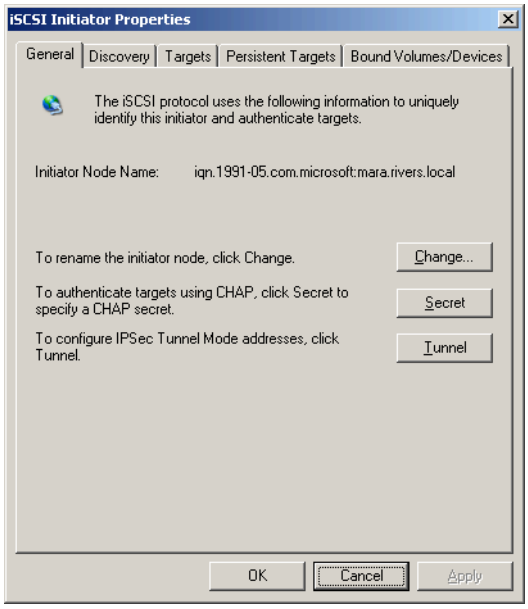


Figure 23-15 iSCSI initiator properties - General

2. In Figure 23-16, click **Add** in the iSNS server pane to add a new iSNS server.

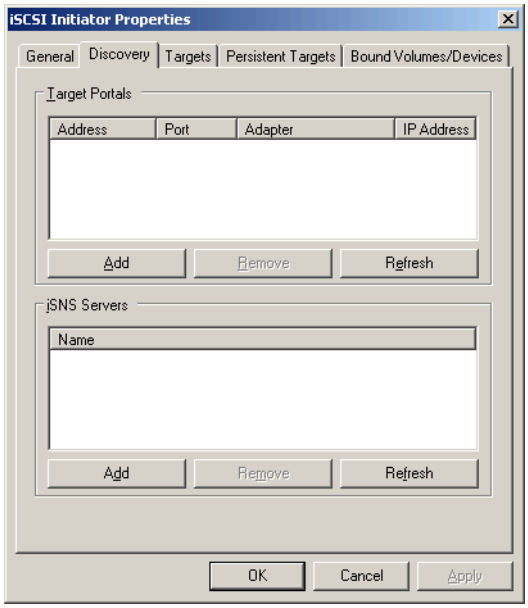


Figure 23-16 iSCSI initiator properties - Discovery

3. In Figure 23-17 on page 557, you can enter the IP address or the host name of the iSNS server you will use. Click **OK**.

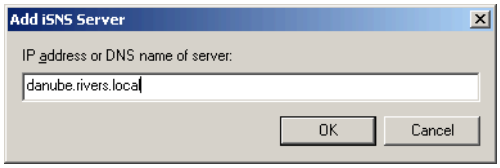


Figure 23-17 Add iSNS server

4. The specified iSNS server is contacted. If it is available, it is added to the list of iSNS servers, as shown in Figure 23-18.

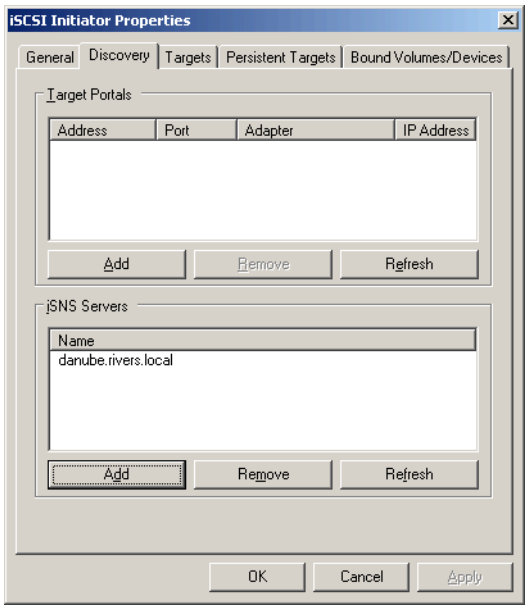


Figure 23-18 iSCSI initiator properties - Added iSNS server

5. Click the **Targets** tab. This tab, shown in Figure 23-19, contains all iSCSI targets that are registered at the iSNS server. The current status of this target is inactive. Select a target from the target list and click **Log On...**.

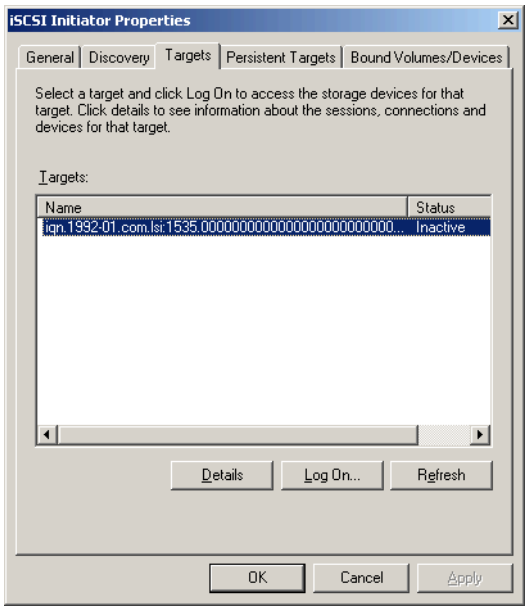


Figure 23-19 iSCSI initiator properties - Targets

6. Proceed with the steps shown in 23.8.3, “Configuring login to discovered targets” on page 562 to log into the target.

23.8.2 Configuring target discovery using sendTargets

Follow these steps (on node Ora in our example) to configure target discovery using iSNS:

1. Open the iSCSI software initiator by selecting **Start** → **Program Files** → **Microsoft Software Initiator** → **Microsoft Software Initiator**. In the window shown in Figure 23-20 on page 559, click the **Discovery** tab.

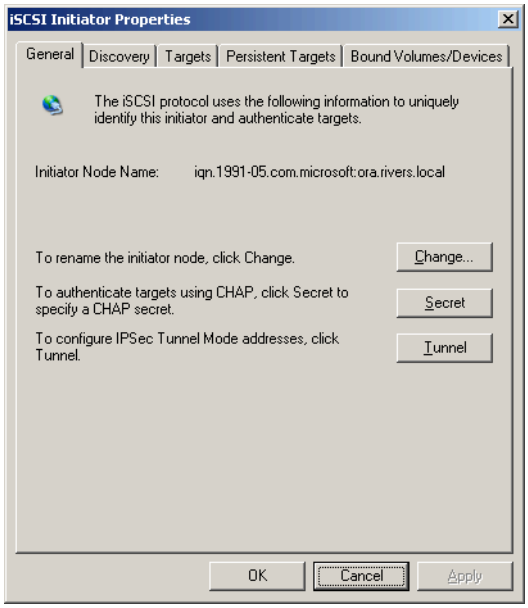


Figure 23-20 iSCSI initiator properties - General

2. In the window shown in Figure 23-21, click **Add** in the Target Portals pane to add one of the iSCSI interfaces of the DS3300.

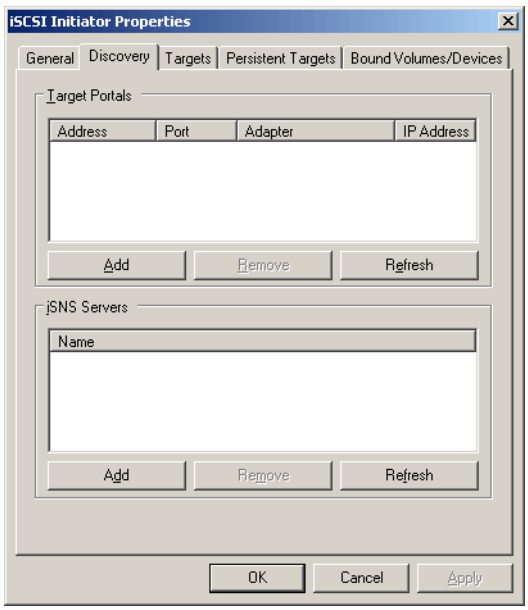


Figure 23-21 iSCSI initiator properties - Discovery

3. In Figure 23-22, enter the IP address or host name of one of the iSCSI interfaces of the DS3300. These are listed for our configuration in Table 23-1 on page 546.

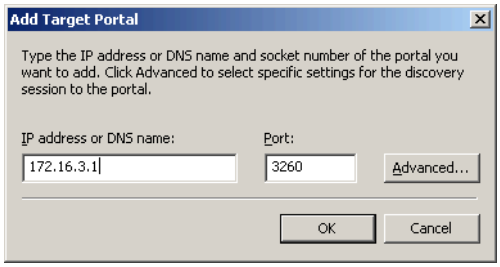


Figure 23-22 Add target portal

4. Click **Advanced...**. Select the Microsoft iSCSI initiator from the **Local adapter** drop-down menu. In the **Source IP** drop-down menu, select the IP address of the node that is connected to the same network as the specified iSCSI portal IP address. Check **Data digest** and **Header digest**, as shown in Figure 23-23, and click **OK**.

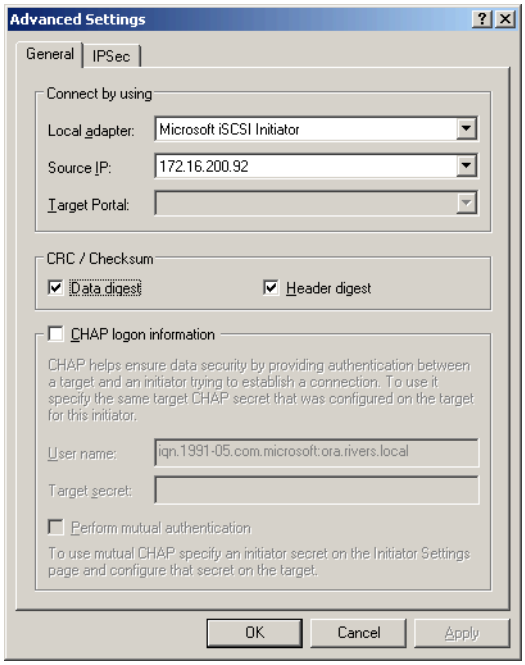


Figure 23-23 Advanced settings

5. The Add Target Portal window (Figure 23-23) appears. Click **OK**.
6. A connection is made to the specified target portal IP address, and the **`sendTargets`** command is issued to discover the storage subsystem, as shown in Figure 23-24 on page 561.

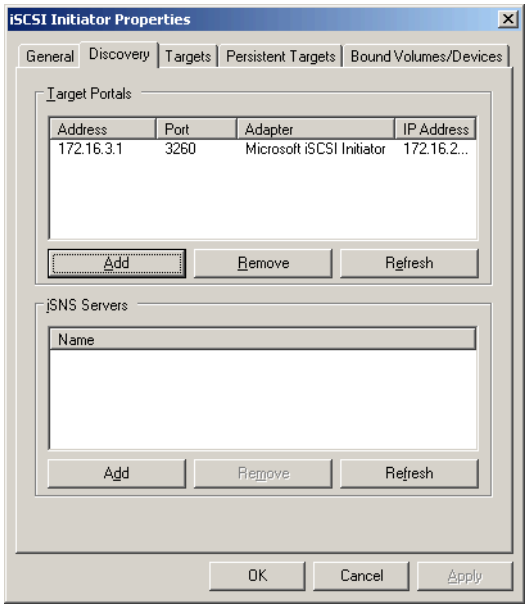


Figure 23-24 iSCSI initiator properties - Discovery address

7. Click the **Target** tab to see the discovered storage subsystems (Figure 23-25). This is obviously the same subsystem that was discovered in 23.8.1, “Configuring target discovery using iSNS server” on page 556.

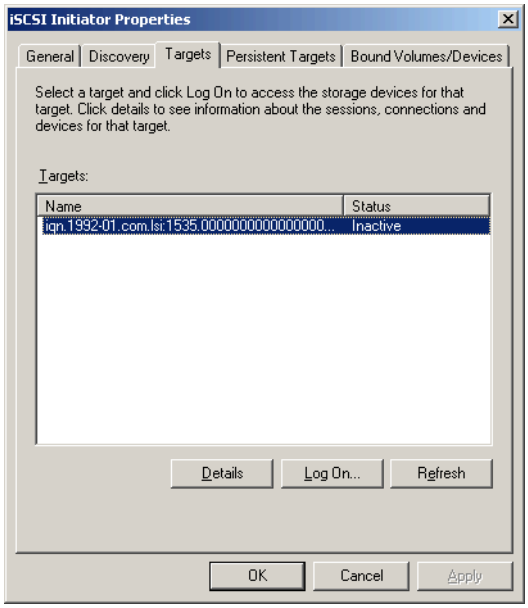


Figure 23-25 iSCSI initiator properties - Discovered targets

8. Proceed with the steps in 23.8.3, “Configuring login to discovered targets” on page 562 to log into the target.

23.8.3 Configuring login to discovered targets

This section describes the configuration tasks required after a target is discovered. We will first configure node Mara and then repeat these steps on Ora.

1. The Target tab contains all iSCSI targets that were discovered, in our case, our DS3300. The current status of this target is inactive. Select a target from the target list and click **Log On...**.See Figure 23-26.

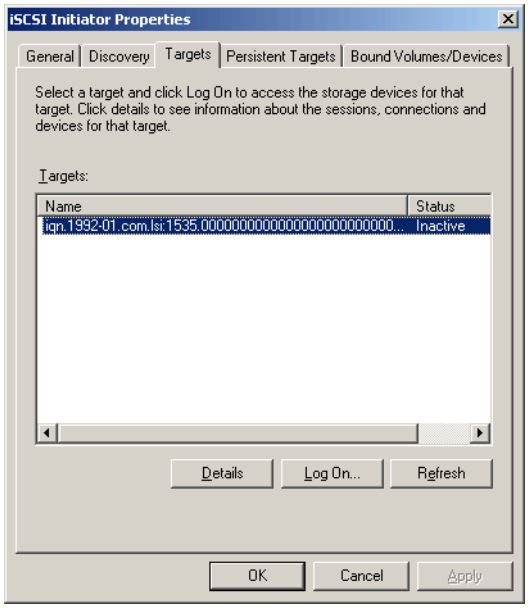


Figure 23-26 iSCSI initiator properties - Targets

2. Check **Automatically restore this connection when the system boots** and **Enable multi-path**, as shown in Figure 23-27. This makes sure the login will be done each time the server reboots. Multi-path must be enabled to allow multiple logons to the same target. In step 9 on page 566, we show how to enable multiple connections.

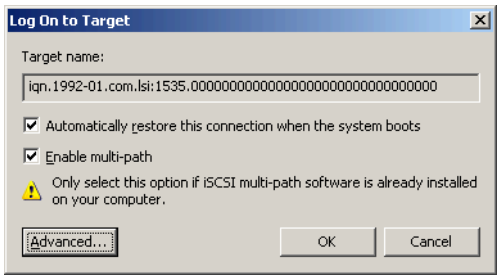


Figure 23-27 Log on to target

3. Click **Advanced...**. Select the Microsoft iSCSI initiator in the **Local adapter** drop-down menu. In the **Source IP** drop-down menu, click the IP address that is connected to the first iSCSI LAN. See Figure 23-28.

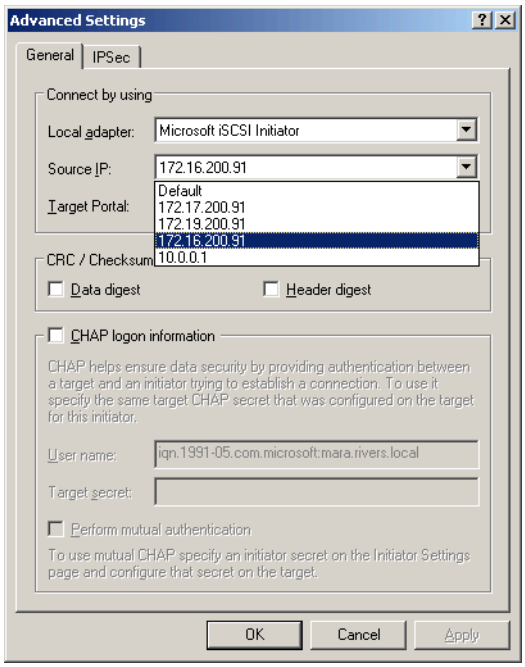


Figure 23-28 Advanced settings - Source IP address

4. Select the IP address of the first controller in the first iSCSI network from the **Target Portal** drop-down menu, as shown in Figure 23-29.

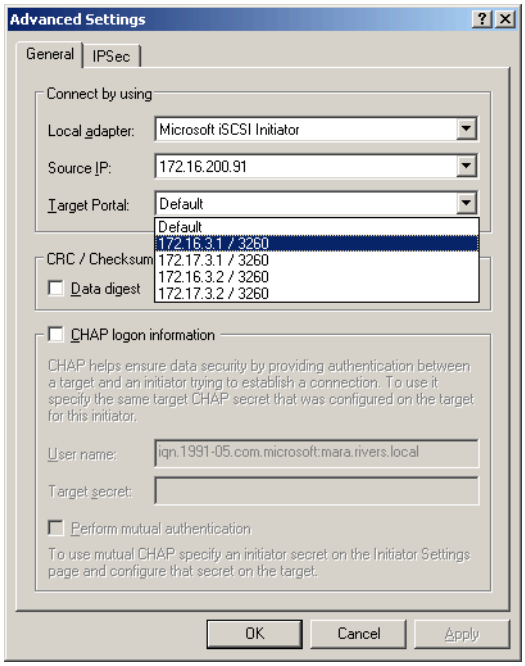


Figure 23-29 Advanced settings - Target portal

5. In Figure 23-30, check **Data digest** and **Header digest** and click **OK**.

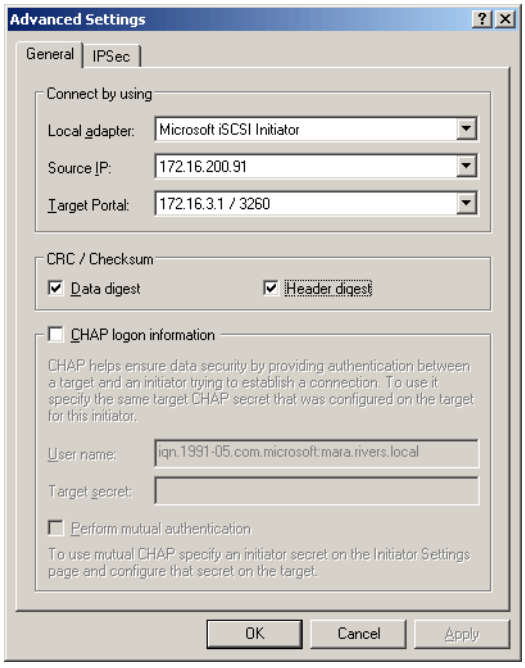


Figure 23-30 Advanced settings - finally settings

6. Back in the Log On to Target window (Figure 23-31), click **OK** to complete the logon.

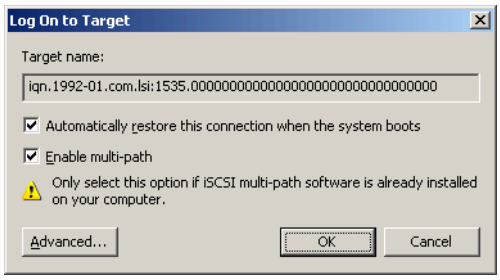


Figure 23-31 Log on to target

7. The iSCSI initiator now has the status of Connected, as shown in Figure 23-32.

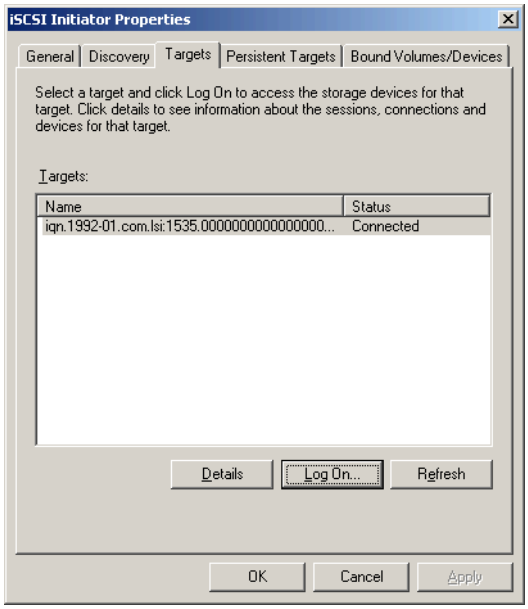


Figure 23-32 iSCSI initiator properties with connected target

8. Open Device Manager; we can see the four logical drives configured on the DS3300 (IBM 1726-3xx FAStT SCSI Disk Device) the Universal Xport SCSI Disk Device, and four IBM 1726-3xx FAStT Multipath Disk Devices, as shown in Figure 23-33. The redundant paths are not yet configured or used.

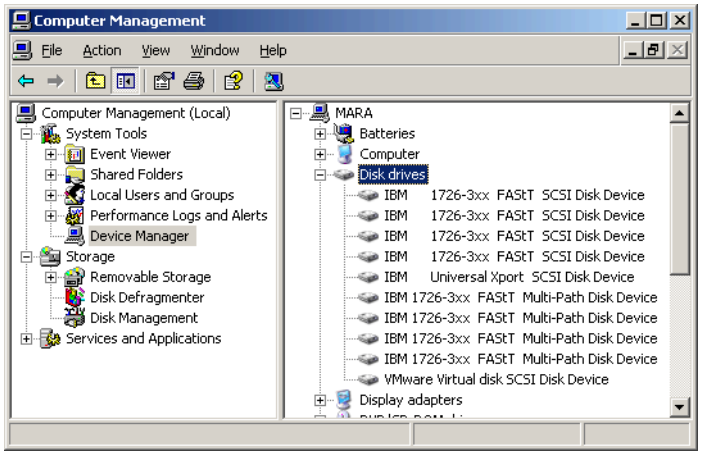


Figure 23-33 Windows Device Manager

9. We need to log on three more times to establish iSCSI sessions to each iSCSI portal or interface of the DS3300. So far, only one session from the NIC interface with IP 172.16.200.91 to 172.16.3.1 was established. To establish the remaining session, select the iSCSI target from the target list and click **Log on**, as shown in Figure 23-34.

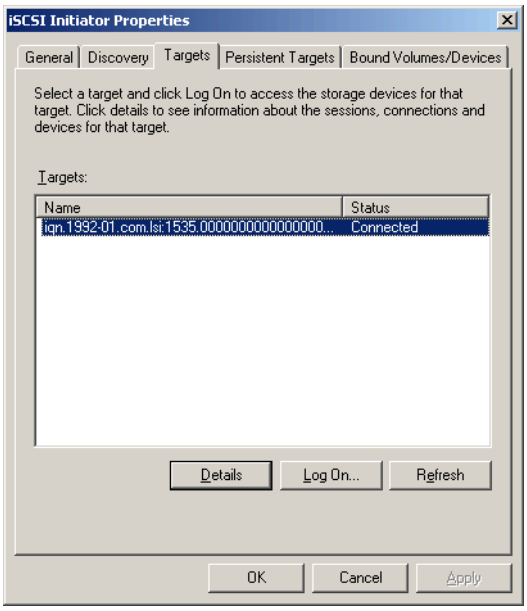


Figure 23-34 iSCSI initiator properties

10. Follow steps 2 on page 562 to 6 on page 564 to configure the advanced properties for the remaining iSCSI sessions. Use the following combinations of source IP and target Portal:
    - Source IP: 172.16.200.91: Target Portal 172.16.3.2
    - Source IP: 172.17.200.91: Target Portal 172.17.3.1
    - Source IP: 172.17.200.91: Target Portal 172.17.3.2
11. After the iSCSI sessions are created to each target portal, you can see details of the sessions. Select the iSCSI target in the target list (Figure 23-34) and click **Details**. See Figure 23-35 on page 567.

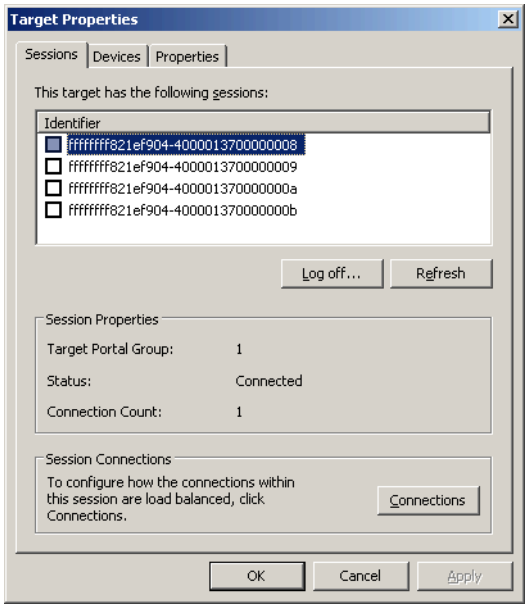


Figure 23-35 Target properties

12. The Devices and Properties tabs can be used to get more details about configured LUNs available to the configured iSCSI sessions. This is useful in case the configuration did not succeed and the disks do not appear in Device Manager.

13. Figure 23-36 shows the disks as they appear in Device Manager after all connections are made. The two network interface cards are connected each to two iSCSI ports of the DS3300. Four logical drives are configured on the DS3300 that are visible over each port. That adds up to 16 IBM 1726-3xx FASt SCSI Disk Devices in the Device Manager. Each iSCSI port of the DS3300 also provides, by default, the Universal Xport SCSI disk device, which is also seen four times. The multipath driver presents four IBM 1726-3xx FAStT Multipath Disk Devices to the host that can be used to create partitions in the next section.

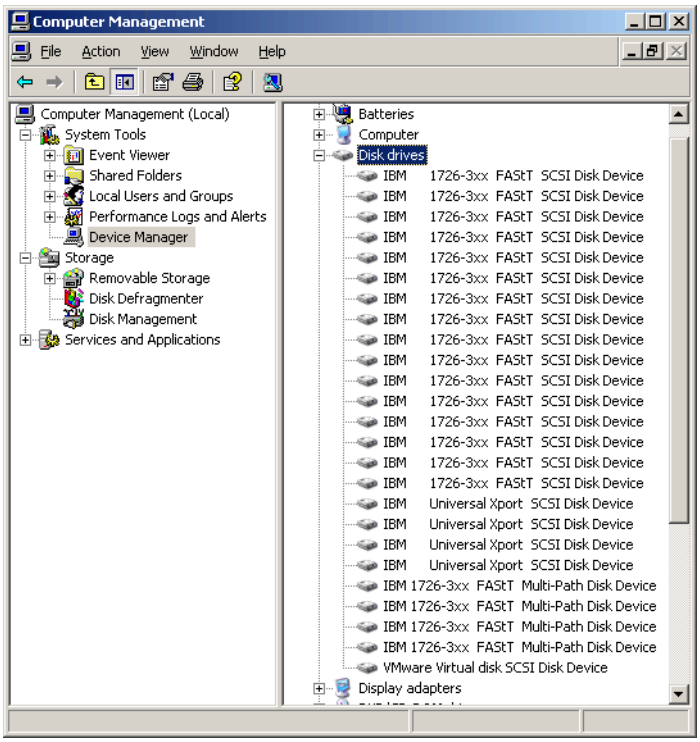


Figure 23-36 Windows Device Manager

Repeat this section for host Ora; the connections to be made are:

- Source IP: 172.17.200.92: Target Portal 172.17.3.2
- Source IP: 172.16.200.92: Target Portal 172.16.3.2
- Source IP: 172.16.200.92: Target Portal 172.16.3.2
- Source IP: 172.17.200.92: Target Portal 172.17.3.1

Proceed with the preparation of the shared disks.

23.9 Preparing shared drives on a cluster node

After the software initiator is configured, the logical drives of the DS3300 are available to the cluster nodes as physical disks. These disks need to be initialized and a partition created and formatted with an NTFS file system on each. The cluster service requires that the drives are basic drives, not dynamic drives. We recommend using labels on each disk to identify it by its drive letter and usage. An example would be Q_Quorum for a drive that is mounted as Drive Q: and used as a Quorum Disk resource.

The drive partitioning and drive letter assignment has to be done only on the first node. After you finish this task, the configuration should look as shown in Figure 23-37 on page 569.

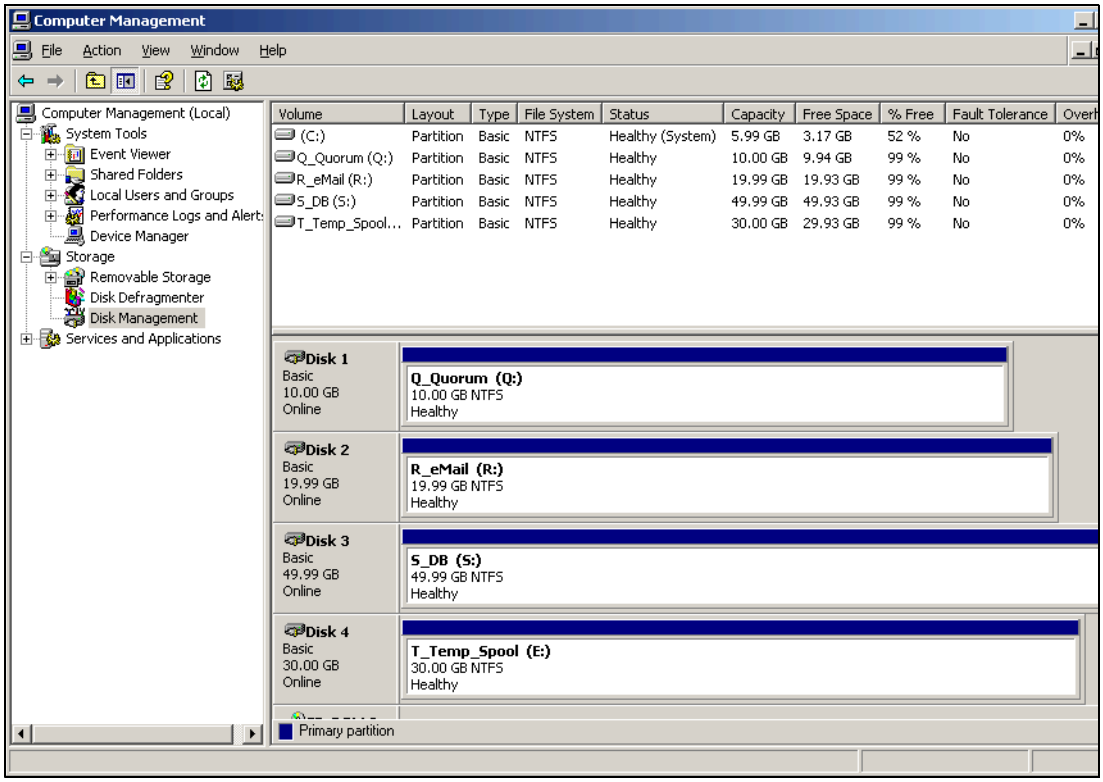


Figure 23-37 Shared drive setup in the operating system

23.10 Installing the cluster service on the first cluster node

The cluster service is now installed on the first node, Mara. This can be done remotely or locally. The remaining nodes will be joined to the cluster, as described in 23.11, “Installing cluster service on the remaining nodes” on page 575. Do these steps:

1. Select **Start** → **Program Files** → **Administrative Tools** → **Cluster Administrator**. The Cluster Administrator starts and the Open Connection To Cluster window appears, as shown in Figure 23-38.

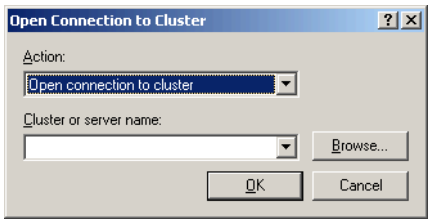


Figure 23-38 Open connection to cluster

2. Select **Create new cluster** (Figure 23-39) from the drop-down menu and click **OK**.

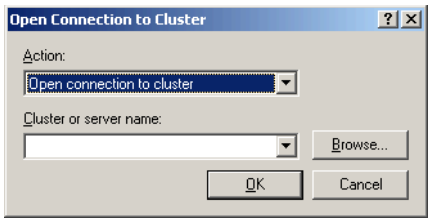


Figure 23-39 Create new cluster

3. The **New Server Cluster** wizard starts, as shown in Figure 23-40. Click **Next**.



Figure 23-40 New server cluster wizard

4. Enter the name for your cluster and click **Next**. The domain is detected automatically. See Figure 23-41.

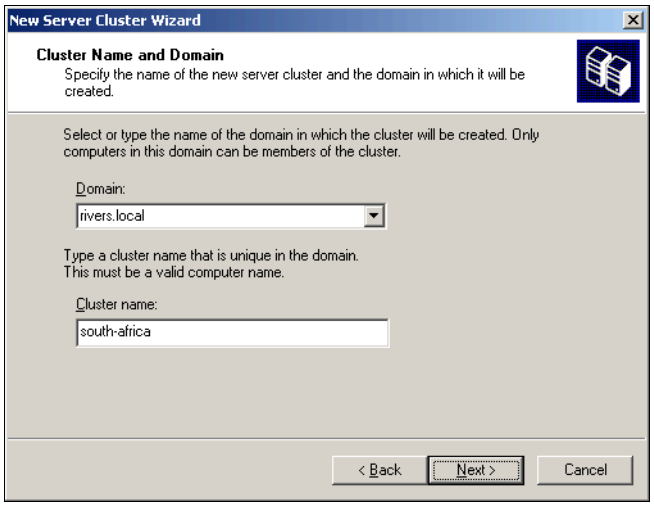


Figure 23-41 Cluster name and domain

5. In Figure 23-42, enter the name of the computer where the cluster service needs to be installed. This is only required when the cluster administrator is running on a management station system and not directly on the first cluster node. The server cluster wizard will add the local host name already in this dialog. Click **Next**.

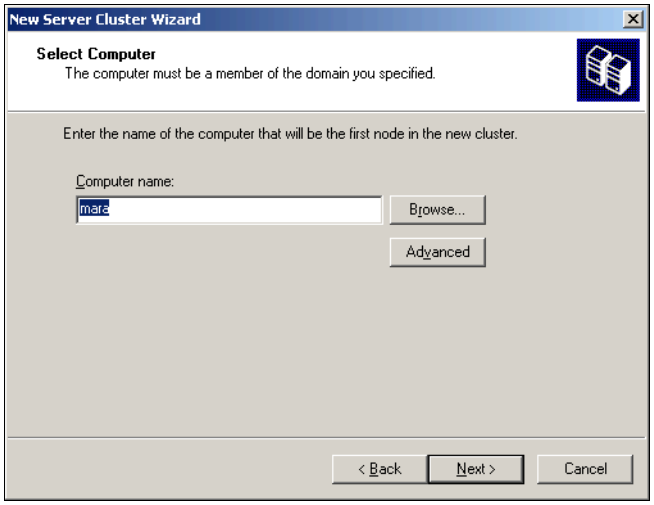


Figure 23-42 Select computer

6. The specified host and configuration will be checked if it is suitable for cluster usage. Any problems that are detected will be logged in a separate log file as well as in the current dialog. You must rectify any issues reported before continuing with the installation.

   Figure 23-43 shows a successful analysis of the host. Click **Next**.

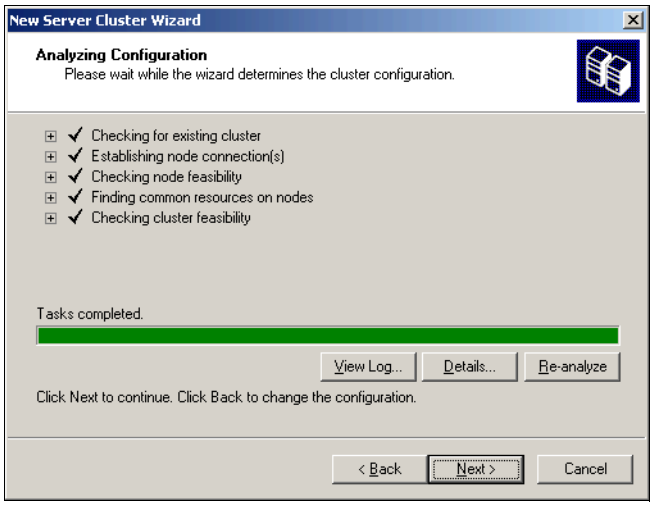


Figure 23-43 Configuration analysis

7. In Figure 23-44, enter the IP address of the cluster (from Table 23-1 on page 546) and click **Next**.

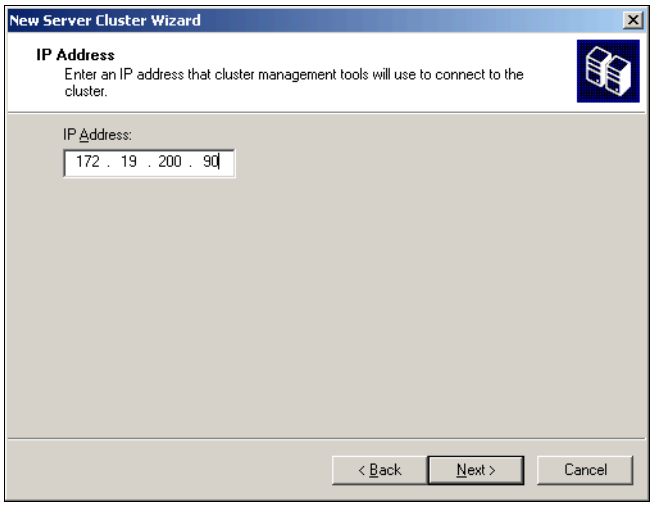


Figure 23-44 Cluster IP address

8. Enter the user name and password of the domain account that was created in step 4 on page 548 to be used as a cluster service account. Click **Next**. See Figure 23-45.

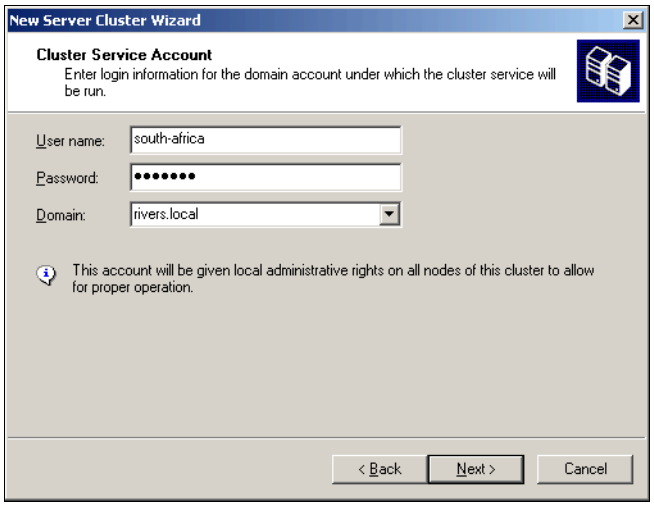


Figure 23-45 Cluster service account

9. Check the summary of the proposed cluster configuration, as shown in Figure 23-46. By default, the first shared disk that is found is used in the configuration as a quorum resource. Click **Quorum** if you want to use another disk as the quorum resource. Because our example is a two-node cluster, it is not possible to use the Majority Node Set; this requires at least three nodes in the cluster. Click **Next** to start the creation of the cluster.

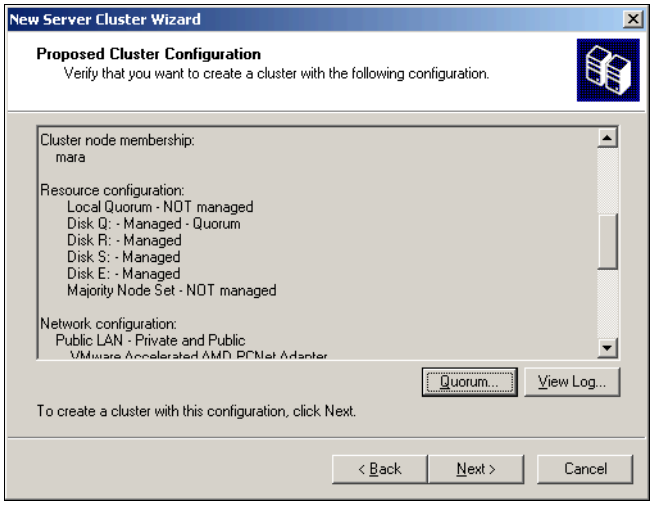


Figure 23-46 Cluster configuration proposal

10. The creation of the cluster takes some time. When it is finished, the green line shown in Figure 23-47 is visible. If something went wrong, this line is red and the cluster is not installed. Click **Next** to finish the wizard.

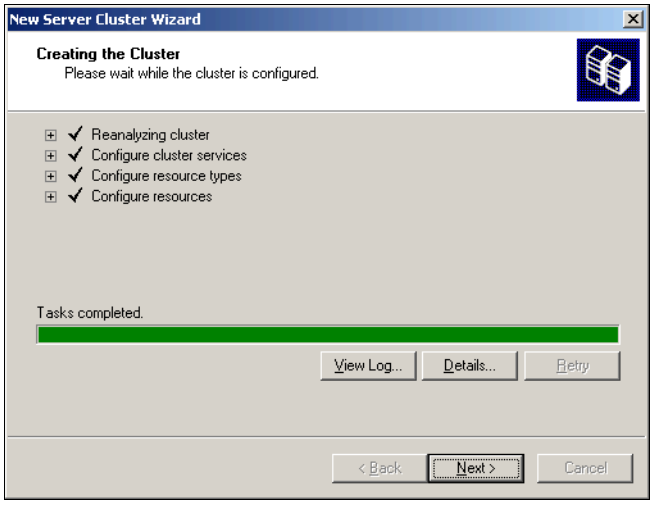


Figure 23-47 Create cluster

11. Click **Finish** to close the wizard. A log is available that can be stored for future reference. See Figure 23-48.

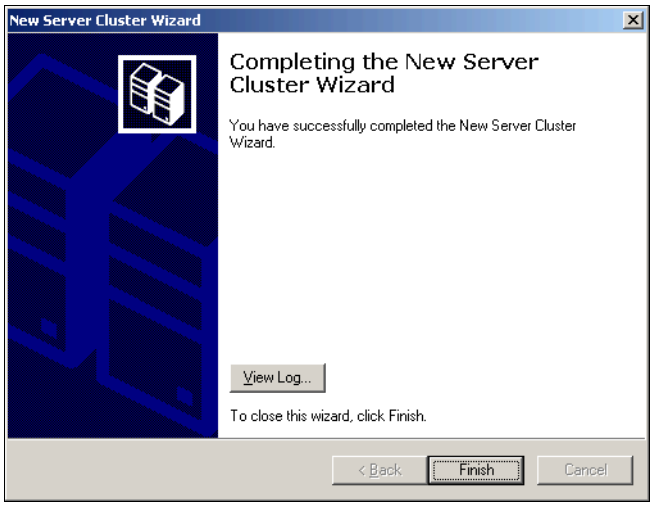


Figure 23-48 Finish server cluster wizard

12. Now the Cluster Administrator is visible, showing a connection to the newly created cluster. The cluster contains our node Mara, as shown in Figure 23-49.

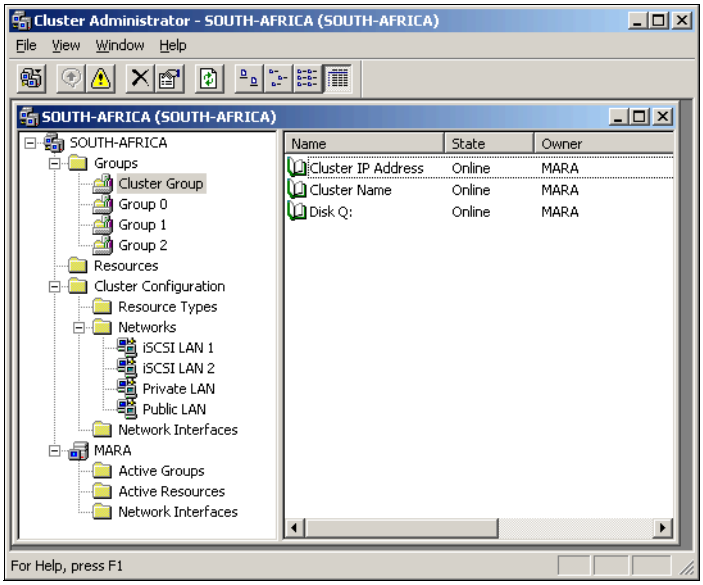


Figure 23-49 Single node cluster on host MARA

Now we will install the cluster service on the remaining nodes of the cluster.

23.11 Installing cluster service on the remaining nodes

The cluster currently contains only the first node. Now we need to install the cluster service on the remaining nodes (Ora in our configuration). This action can be done from a management station where a Cluster Administrator is available, from the first node, or any other node. This section uses the first node to install the cluster service on the remaining nodes. Do these steps:

1. Use the running Cluster Administrator on the first node. Select **File** → **Open Connection**. In the Open Connection to Cluster dialog, select **Add nodes to cluster** from the **Action** drop-down menu. Ensure that the Cluster or Server name list contains the name of the cluster, as shown in Figure 23-50. Click **OK**.

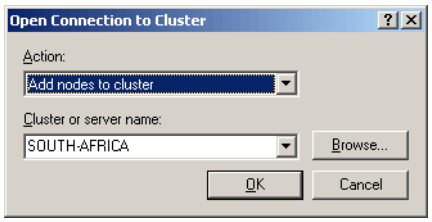


Figure 23-50 Add nodes to cluster

2. The Add Nodes Wizard starts (Figure 23-51). Click **Next**.



Figure 23-51 Add Nodes Wizard

3. Enter a host name of one of the remaining cluster nodes and click **Add**. The **Browse** button can be used to find the nodes. Our example is a two-node cluster, so enter the other node name, Ora, as shown in Figure 23-52.

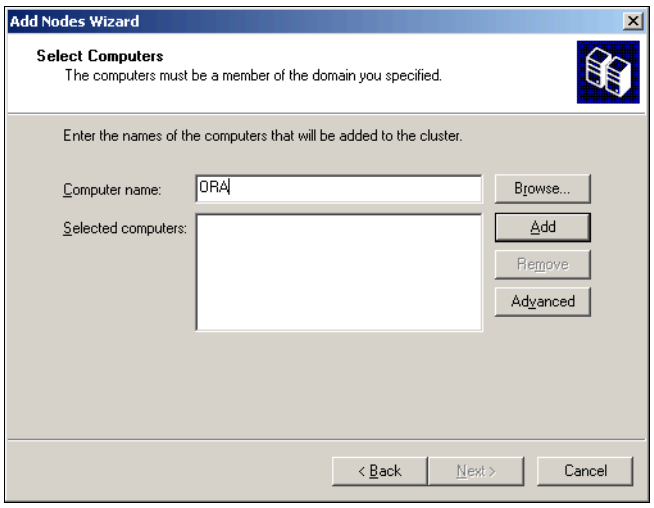


Figure 23-52 Select computers

4. Add more nodes if required and click **Advanced**. See Figure 23-53.

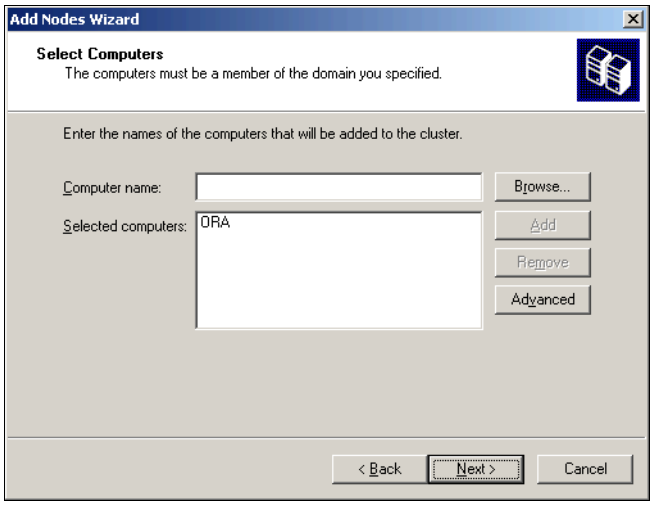


Figure 23-53 Selected computers

5. Select **Advanced (minimum) configuration** and click **OK**. See Figure 23-54 on page 577. Not all storage will be included in the configuration. This is not necessary, because the cluster is already running and has access to the drives. The remaining nodes currently cannot access these drives.

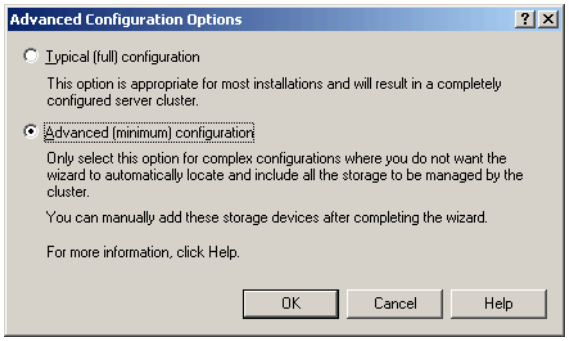


Figure 23-54 Advanced configuration options

6. In the Select Computers dialog of the Add Nodes Wizard (Figure 23-53 on page 576), click **Next**.
7. The cluster configuration will be analyzed on the added computers, as shown in Figure 23-55. Click **Next** when the analysis is complete.

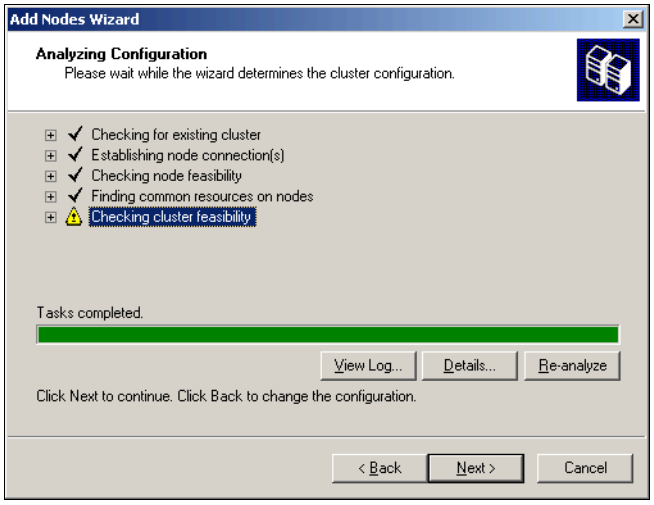


Figure 23-55 Analyzing configuration

8. In Figure 23-56, enter the password of the cluster service account as it was created on the first node in step 4 on page 548, and click **Next**.

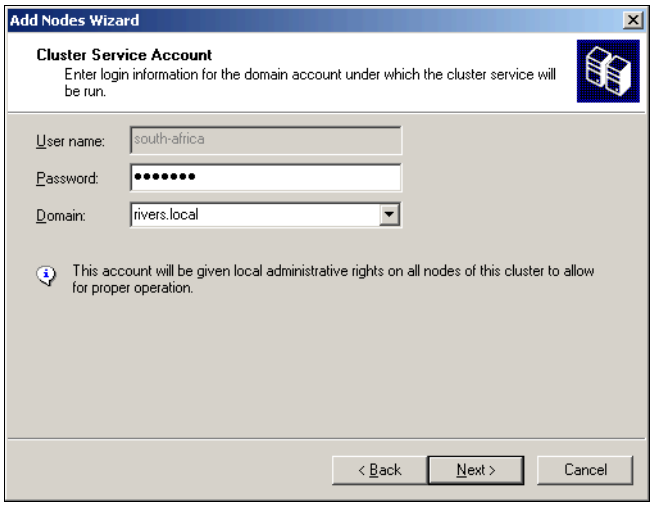


Figure 23-56 Cluster service account password

9. Verify the proposed cluster configuration shown in Figure 23-57 and click **Next**.

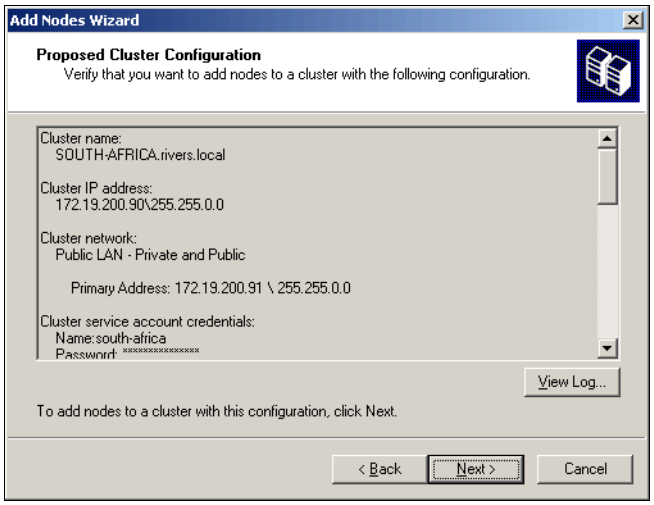


Figure 23-57 Cluster configuration proposal

10. The cluster configuration is analyzed again, and the cluster service is installed on the selected nodes. The attention mark in Figure 23-58 on page 579 is there because the storage is not detected, because of the Advanced (minimum) configuration selection in Figure 23-54 on page 577. Click **Next** when the installation is complete.

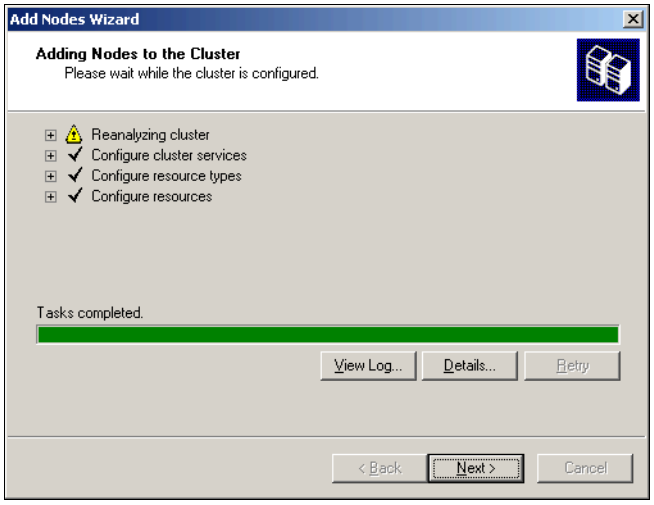


Figure 23-58 Adding nodes to the cluster

11. Click **Finish** to close the Add Nodes Wizard (Figure 23-59).

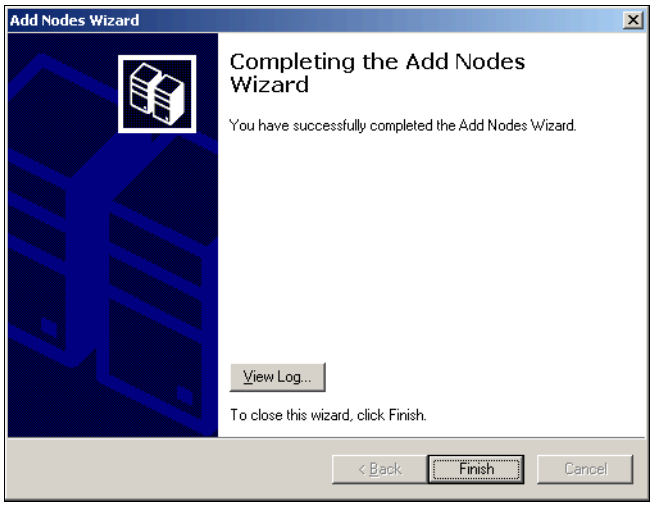


Figure 23-59 Add nodes wizard complete

12. The Cluster Administrator in Figure 23-60 shows that the cluster is now running with two cluster nodes.

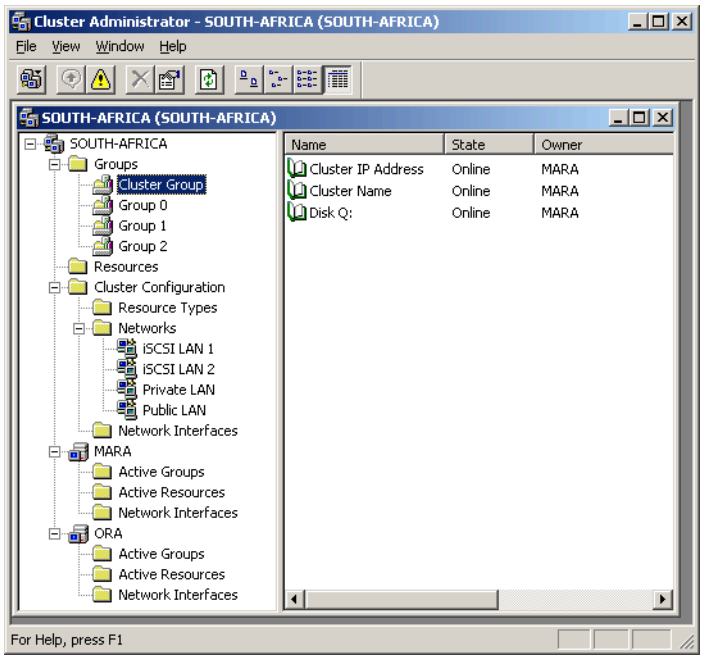


Figure 23-60 Cluster with both nodes running

The cluster is now running with the default network usage. Now we configure this for our requirements.

23.12 Configuring the cluster network usage

By default, all networks found during the cluster installation are used for internal and client communication. One network must be configured for internal communication only. To have redundant internal network communication, it is possible to run internal and client communication on the same network. This section describes the setup of the network usage in the cluster. The Public LAN will be used for internal and client communication, the Private LAN will be used only for internal communication, and the two iSCSI networks are not used in the cluster, since they are reserved for iSCSI traffic only.

Do these steps:

1. In Cluster Administrator, select **Cluster Configuration** → **Networks**. The Role column in Figure 23-61 on page 581 shows that all networks are currently used for both internal and client communication.

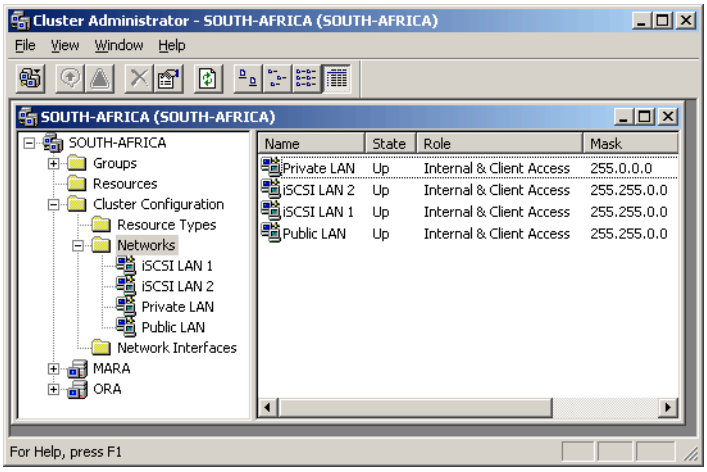


Figure 23-61 Network usage after installation

2. Right-click **iSCSI LAN 1** and select **Properties**, as shown in Figure 23-62.

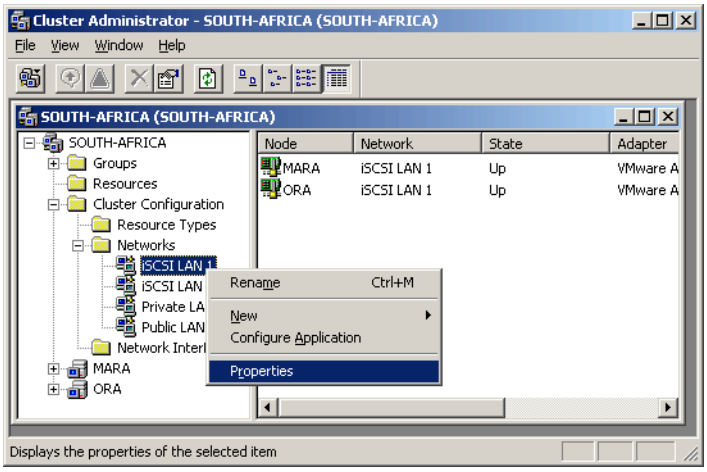


Figure 23-62 iSCSI LAN 1

3. In Figure 23-63, uncheck **Enable this network for cluster use** and click **OK**.

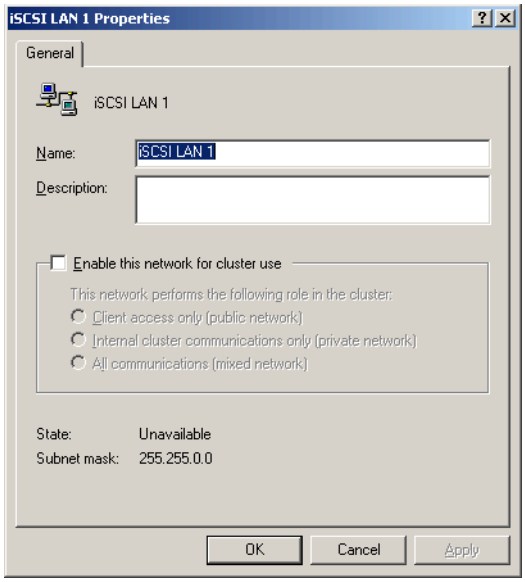


Figure 23-63 iSCSI LAN 1 properties

4. Similarly, disable iSCSI LAN 2 network use in the cluster.
5. Right-click **Private LAN** and select **Properties**, as shown in Figure 23-64.

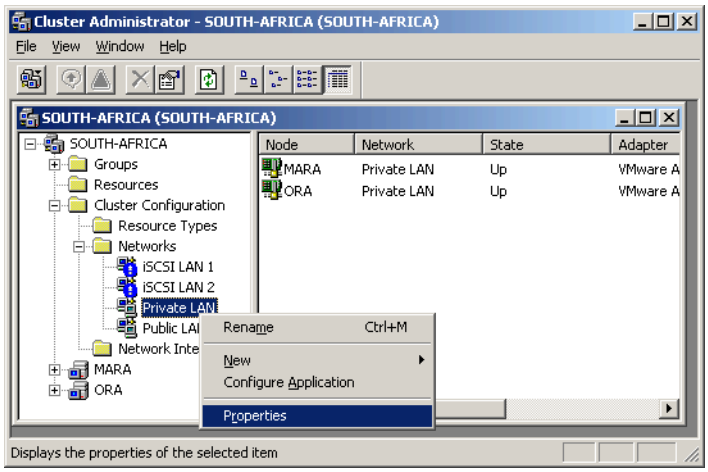


Figure 23-64 Private LAN

6. Select **Internal cluster communications only (private network)** and click **OK**. See Figure 23-65.

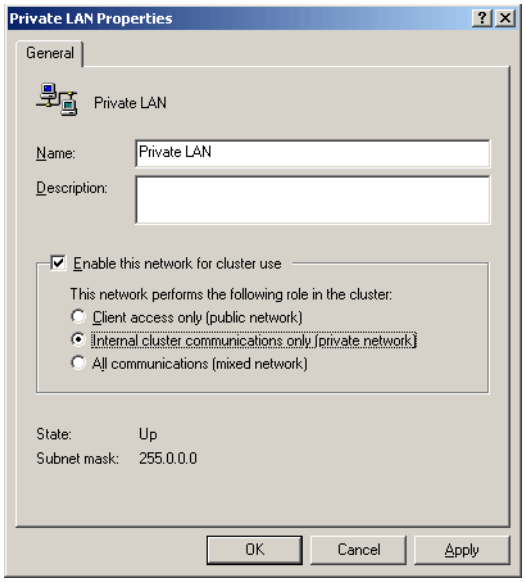


Figure 23-65 Private LAN properties

7. Verify that the private LAN has the highest priority for internal cluster communication. To do this, right-click the cluster name in the Cluster Administrator and select **Properties**. See Figure 23-66.

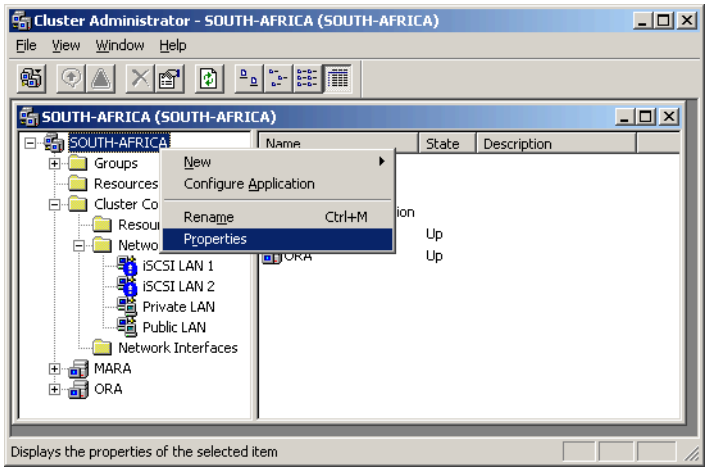


Figure 23-66 Cluster context menu

8. In the Properties dialog, select the **Network Priority** tab, as shown in Figure 23-67. Verify that the network that is used for internal communication only is the first network in the list. Use the **Move up** or **Move down** buttons to change this if required. Click **OK**.

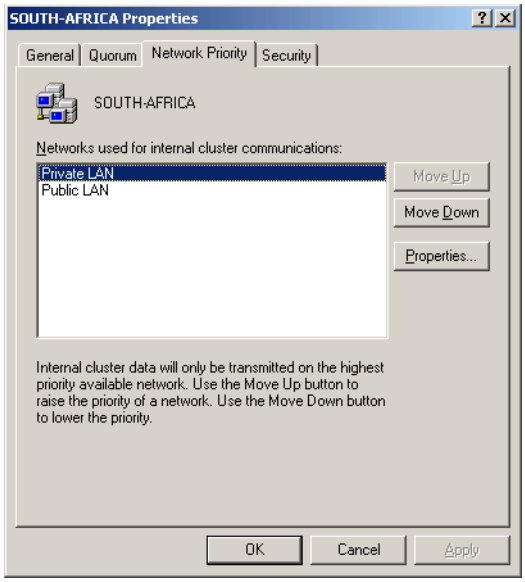


Figure 23-67 Cluster network priority

9. The network usage configuration should now look as shown in Figure 23-67.

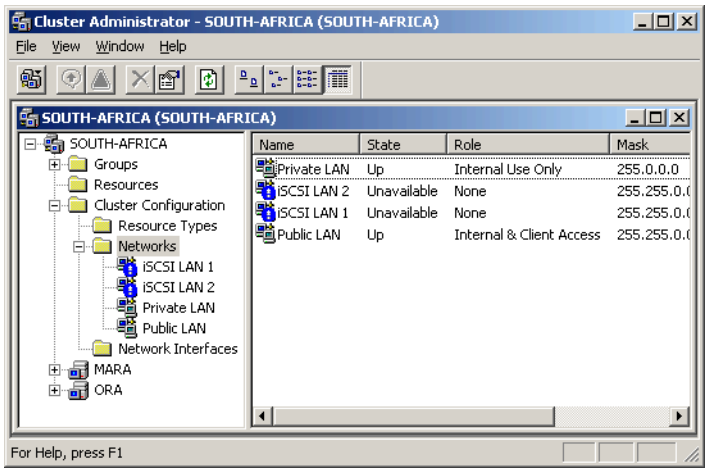


Figure 23-68 Networks reconfigured

The cluster is now ready for testing.

23.13 Testing the cluster configuration

The cluster installation is done and should be tested. You should test that you can move resources between the cluster nodes and that there is proper access to the shared drives in the cluster on each node. You should also verify that failover works when each node goes down. After all tests are successful, you should install applications on the drives.

24

FC configuration 1 - Emulex HBA boot blade server from SAN

In this chapter, we discuss a sample configuration and describe how to boot a blade server with an Emulex HBA from a DS3400. For this example, we use two FC Blade Switches (McData), directly attached to the DS3400. Before you start, make sure that all the firmware on the storage subsystem and in the host is flashed to the latest level, as described in 14.3, “Download firmware” on page 307.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

24.1 Cabling and zoning

In 4.5, “Host attachment” on page 40, we discuss the most common host cabling scenarios for Fibre Channel. In this sample configuration, we use a blade as the host technology. We use two FC Switches in the BladeCenter to provide redundant attachment between the blades and the DS3400. Our topology, even though we are using an IBM BladeCenter, is as shown in Figure 4-9 on page 45. The blade has one dual port HBA installed. Each HBA port is connected to one of the BladeCenter switches. One switch is connected to Controller A of the DS3400, and the other switch is connected to Controller B. The two FC switches are not interlinked. Figure 24-1 shows our specific configuration. Figure 24-2 shows the FC connections in the BladeCenter.

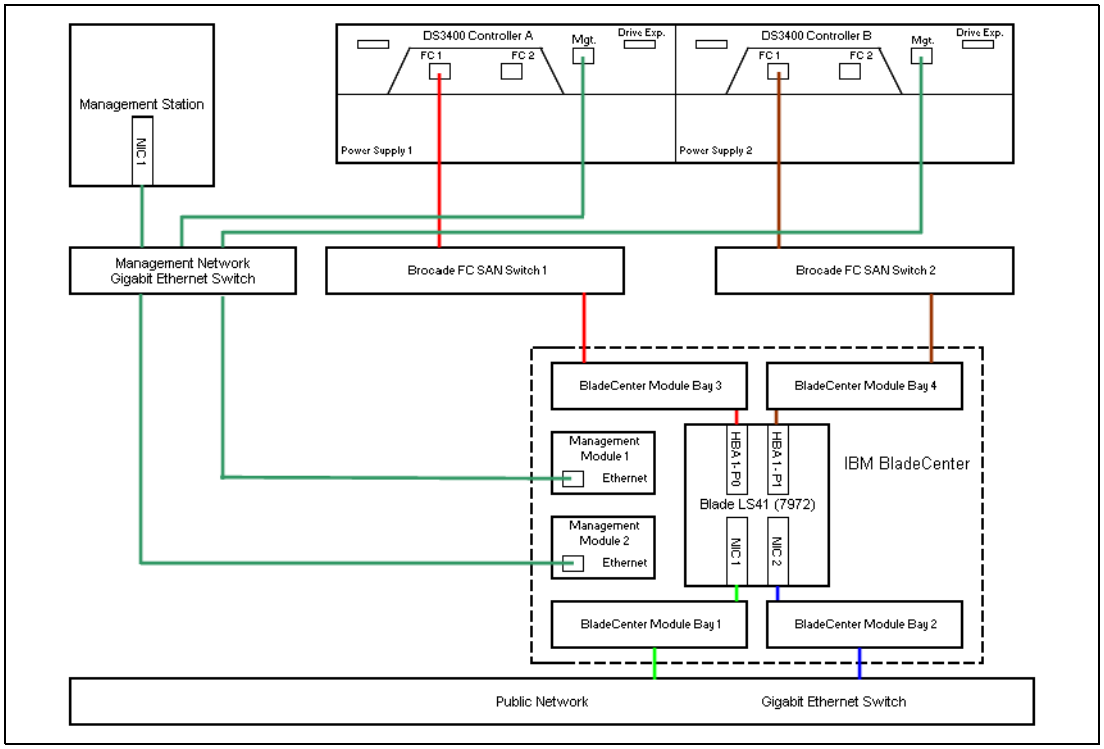


Figure 24-1 Our sample environment

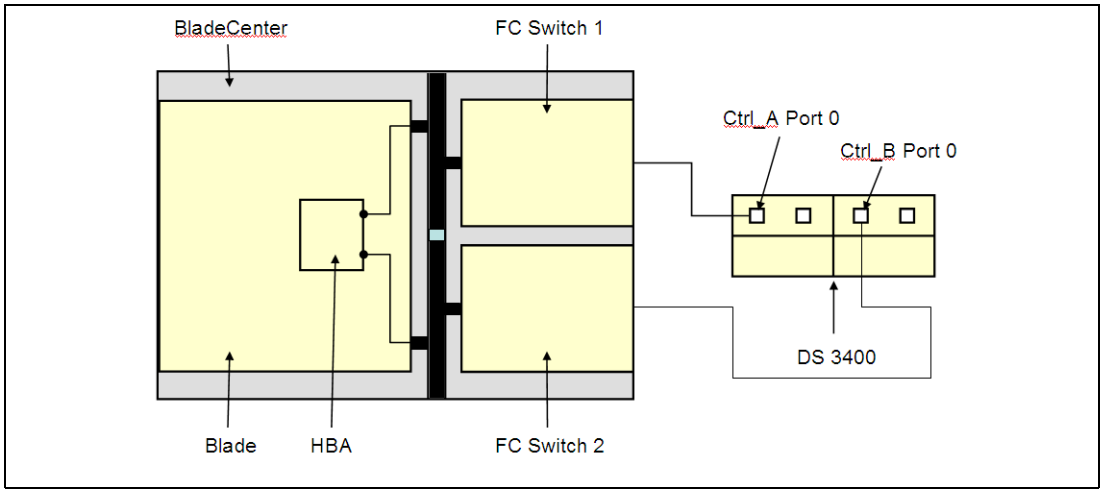


Figure 24-2 BladeCenter sample configuration

First, we cable our environment as shown in Figure 24-2 on page 586 so that it provides redundancy.

After the cabling, we zone the switches for our SAN. Our sample configuration uses McData 4 GB BladeCenter FC Switches. Log on to one of the BladeCenter FC Switches and highlight the switch to display the associated view (Figure 24-3).

Note: A detailed description of blade technology and how to manage it is beyond the scope of this book. You can find information about the IBM BladeCenter and how to access the switch modules in *Advanced Management Module / Management Module User's Guide - IBM BladeCenter E, H, T, HT, S,* MIGR-5073887, at the following site:

http://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-5073887&brandind=5000020

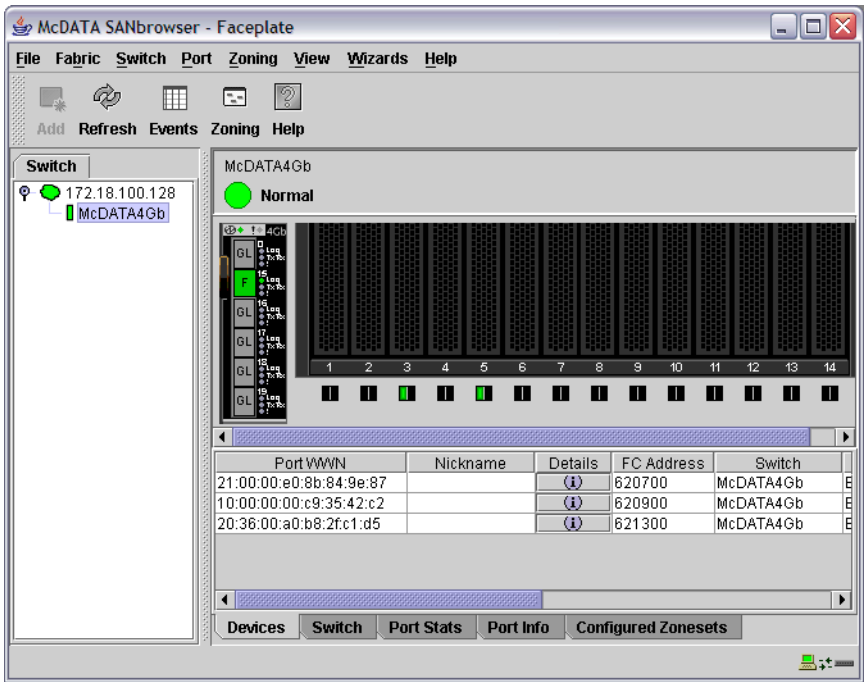


Figure 24-3 McData SANbrowser - Faceplate

This view displays all the WWNs attached to this switch. If you do not already know the WWN of your blade HBA, refer to 9.1.2, “Preparations in the host server” on page 132 to discover it.

Do these steps:

1. Select **Zoning** → **Edit Zoning** from the menu bar to open the Zoning configuration (Figure 24-4).

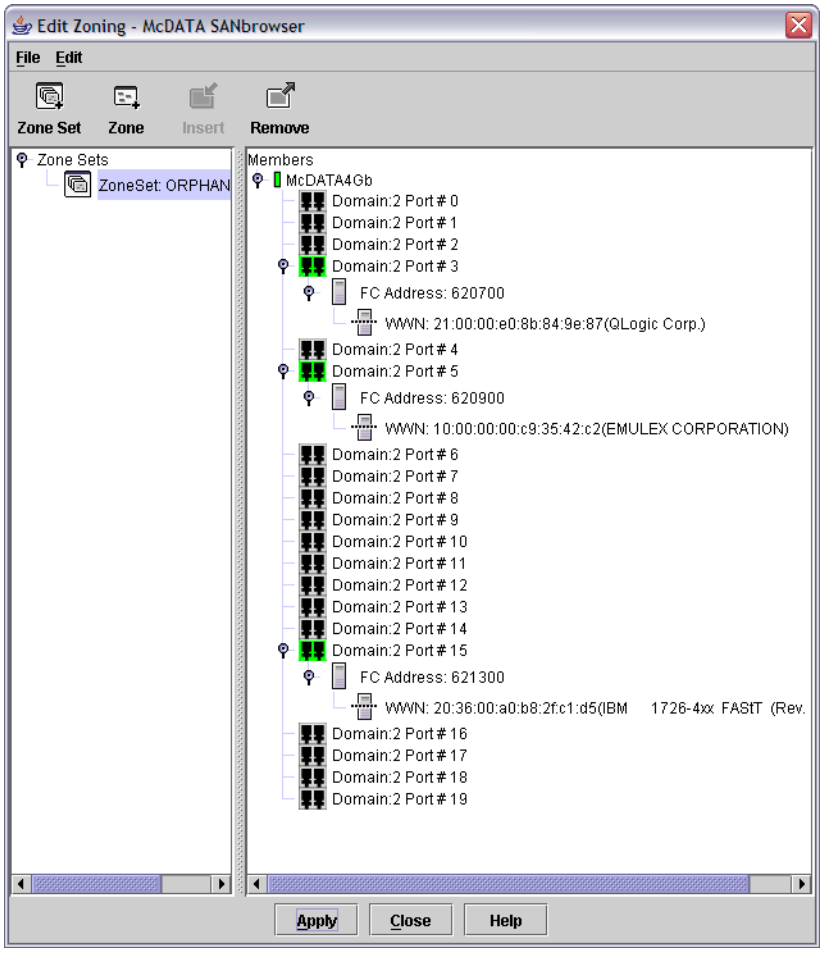


Figure 24-4 Edit Zoning

2. Select **Edit** → **Create Zone Set** from the menu bar. The Create Zone Set window appears (Figure 24-5). Enter a name for the new zone set and click **OK**.

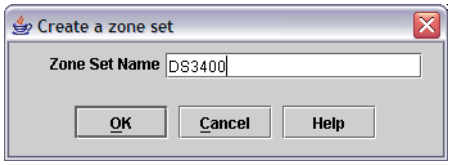


Figure 24-5 Create a zone set

3. The zone set appears in the right column (Figure 24-6). Right-click it and select **Create a Zone**.

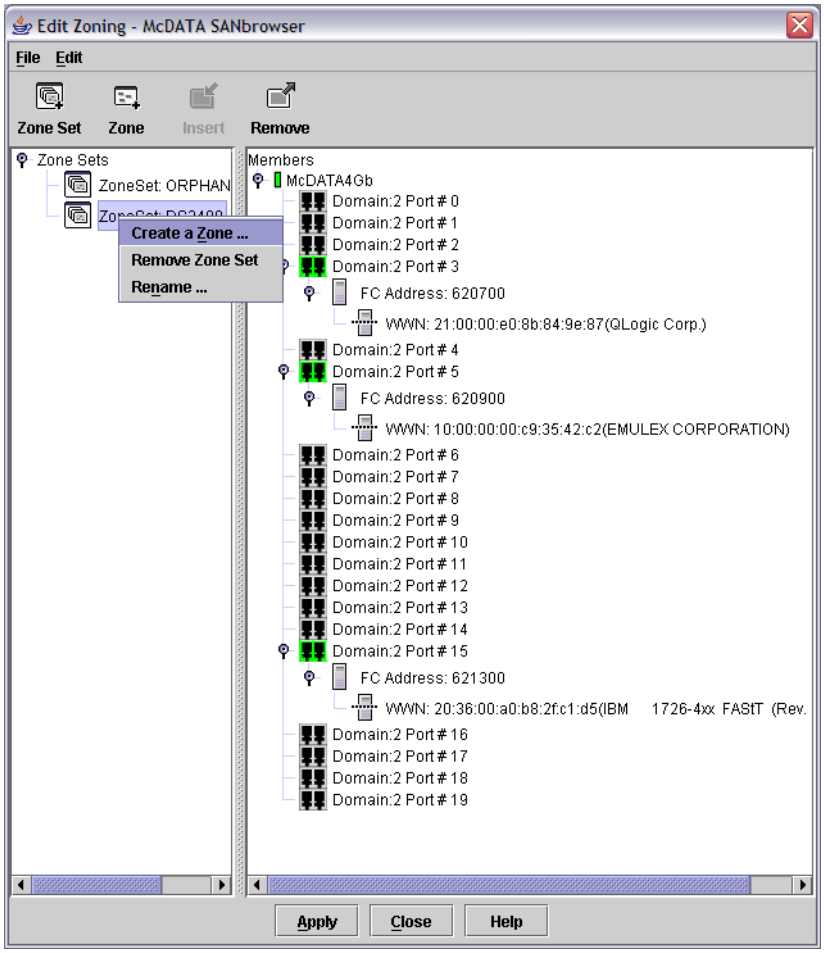


Figure 24-6 Create a zone

4. Enter a zone name (Figure 24-7) and click **OK**.

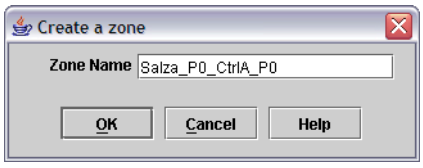


Figure 24-7 Zone Name

5. The new zone now appears in the right column if you expand the zone set (Figure 24-8). Highlight the zone, then drag and drop the WWNs of the members to the zone. Click the WWN of the BladeCenter HBA and drop it on the zone, and repeat this for the DS3400 port.

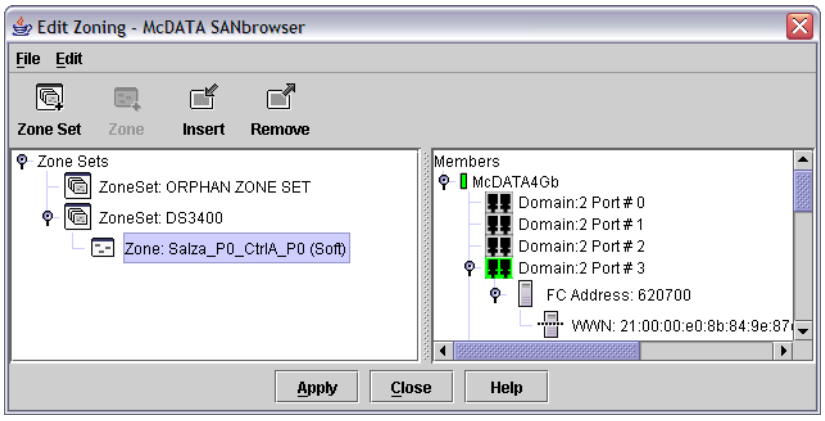


Figure 24-8 Zone created

6. After adding the zone members, the view should look similar to Figure 24-9.

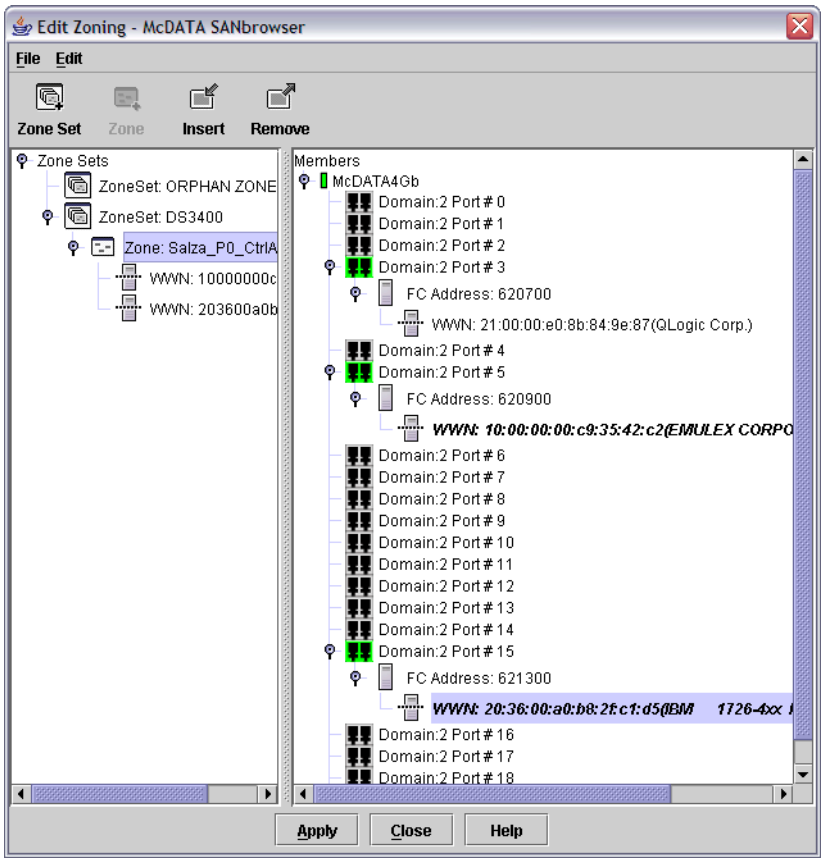


Figure 24-9 Members added to zone

7. Save the new zoning; click **Apply** and the Save Zoning & Error Check window appears (Figure 24-10). Click **Perform Error Check**, and if no errors are found, click **Save Zoning**.

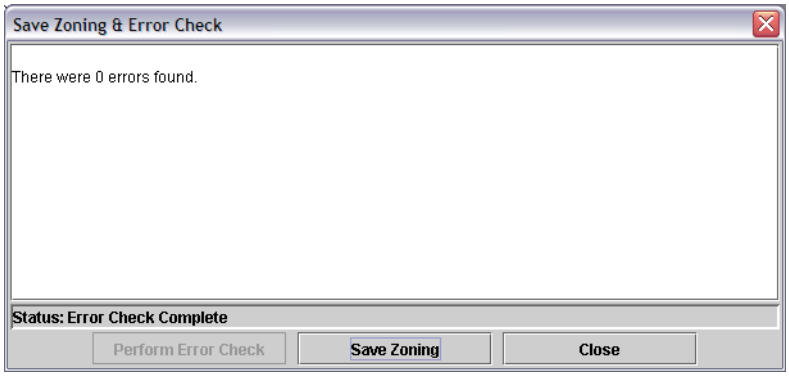


Figure 24-10 Error Check

8. You will be asked if you want to activate this zone after you have created it (Figure 24-11).

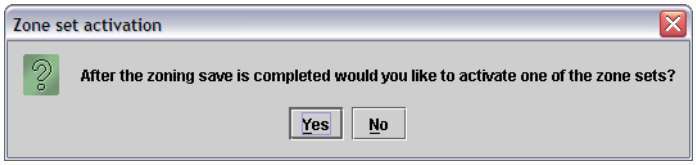


Figure 24-11 Activate Zoning

9. Select **Yes** and choose the zone set you want to enable (Figure 24-12). Select the zone set you have currently created and click **OK**.

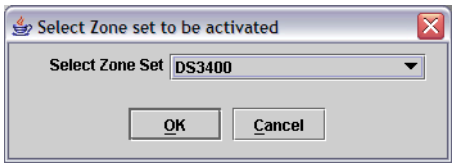


Figure 24-12 Select Zone Set to be activated

You have now successfully created the zoning on the first BladeCenter switch. The first HBA can now communicate over one connection to the DS3400.

Repeat these steps on the second FC switch to enable the second redundant connection.

24.2 Configuring the DS3400

Now you have to configure the IBM DS3400 storage server as we described in Chapter 9, “Administration - Configure” on page 131. Since in this configuration we are using boot from SAN, we will only initially add one HBA WWN. The second WWN will be added to the host after we have installed the operating system on the disk and we have added the MPIO multipath driver.

On the DS3400, create the host and the host port (HBA), create a logical drive of the size required for the boot partition (or use an existing free logical drive), and assign the logical drive to the host, specifying LUN 0.

Notes:

- In order to successfully install the operating system, there should only be a single unique path from the blade server to the boot LUN through the storage controller that currently owns the boot LUN; otherwise, the operating system installation will fail due to lack of multipath support in the OS. Additional paths can be enabled after the operating system has been successfully installed, including the multipath driver.
- The boot device always has to be LUN=0.

24.3 Enabling boot from SAN with Emulex adapters

Now we will configure the HBA for booting from SAN. In this sample configuration, we show this on the Emulex BladeCenter 4 GB HBA.

Note: In our configuration, we are using Emulex HBAs, however, HBAs from QLogic are also supported (See Chapter 27, “FC configuration 4 - QLogic HBA boot for IBM System x from SAN” on page 671). The BIOS utility procedures for QLogic cards are also described in Chapter 26, “FC configuration 3 - Linux SAN boot from a DS3400 with an IBM System x server” on page 639.

24.3.1 Set port switch speed

Depending on the switch module used in the BladeCenter, you should set the switch ports associated with the Emulex HBAs to a fixed negotiation speed of 4 GBps, rather than the default of Auto. This is required on the QLogic or McData 4 GBps 10/20 port switch modules. Consult your switch documentation for information about how to do this. On our McData switch, use **`telnet`** to open a connection to the switch. We then use the commands shown in Example 24-1 to change the port speed.

Example 24-1 Set fixed speed on port

```
admin start
config edit
set port 5 speed 4
config save
```

24.3.2 Enable adapter BIOS

By default, the adapter BIOS is disabled. For boot from SAN, the BIOS must be enabled so that the boot devices can be properly selected. Do these steps:

1. Enter the Adapter BIOS (as described in “FC Emulex HBA” on page 136). From the screen in Figure 9-9 on page 137, select the WWN that was assigned to the LUN=0 logical drive in 24.2, “Configuring the DS3400” on page 591.
2. From the port menu (Figure 9-10 on page 137), type 2 and press Enter to open Configure This Adapter’s Parameters (Figure 24-13).

```
    Adapter 01:             PCI  Bus:0B Device:05 Function:00

    LP11000:        I/O Base: 5000  Firmware Version: BS2.10A5
    Port Name: 10000000 C93542C2   Node Name: 20000000 C93542C2
    Topology: Auto Topology: Loop First (Default)
    The BIOS of this Adapter is Disabled

1.  Enable or Disable BIOS
2.  Change Default ALPA Of This Adapter
3.  Change PLOGI Retry Timer (+Advanced Option+)
4.  Topology Selection (+Advanced Option+)
5.  Enable or Disable Spinup Delay (+Advanced Option+)
6.  Auto Scan Setting (+Advanced Option+)
7.  Enable or Disable EDD 3.0 (+Advanced Option+)
8.  Enable or Disable Start Unit Command (+Advanced Option+)
9.  Enable or Disable Environment Variable (+Advanced Option+)
A.  Enable or Disable Auto Boot Sector (+Advanced Option+)
B.  Link Speed Selection (+Advanced Option+)

  Enter a Selection: 1_

Enter <x> to Exit         <Esc> to Previous Menu
```

Figure 24-13 Emulex BIOS - Configure This Adapter’s Parameters

3. Type 1 and press Enter to “Enable or Disable BIOS”. In Figure 24-14, you can see that the BIOS is currently disabled.

```
    Adapter 01:             PCI  Bus:0B Device:05 Function:00


    The BIOS is Disabled!!


    Enable Press 1, Disable Press 2:_









Enter <x> to Exit         <Esc> to Previous Menu
```

Figure 24-14 Emulex BIOS - Enable Adapter BIOS

4. Type 1 and press Enter again to enable the BIOS. You can see in Figure 24-15 that the BIOS is now enabled.

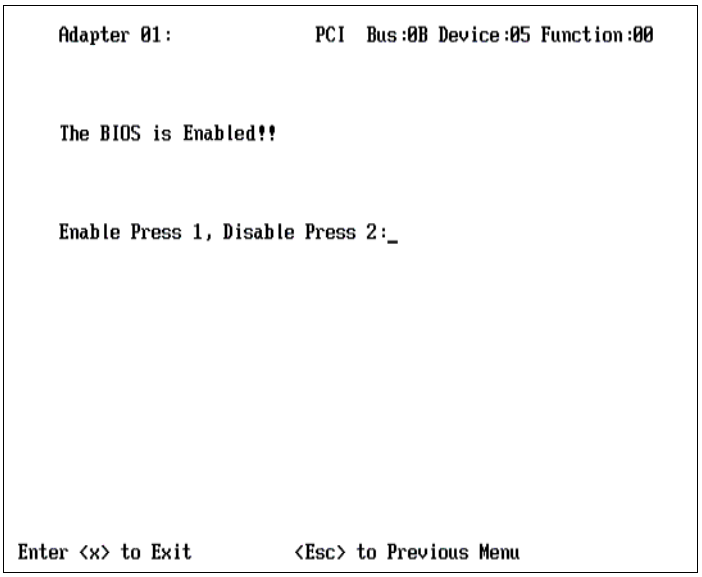


Figure 24-15 Emulex BIOS - Adapter BIOS enabled

5. Press Esc to exit the Emulex BIOS and return to the Configure Adapter's Parameters menu (Figure 24-13 on page 593). Press Esc again to return to the Main Configuration menus (Figure 9-10 on page 137).

Now we have to configure the boot device.

24.3.3 Select boot device

Do these steps:

1. From the port menu (Figure 9-10 on page 137), type 1 and press Enter to go to Configure Boot Devices. This lists all the saved boot devices (Figure 24-16 on page 595). No device is configured yet. Now we add the DS3400 to the list. Type 1 and press Enter to configure the primary boot device.

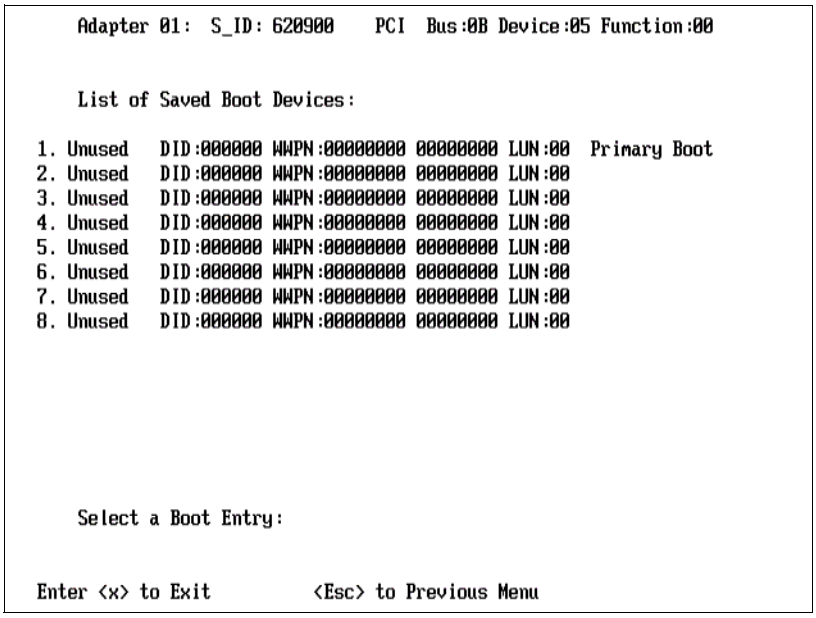


Figure 24-16 Emulex BIOS - List of saved boot devices

2. The next view shows all the possible boot devices that were scanned and detected during initialization. Since the DS3400 (1726-4xx) is attached to the HBA, it is shown in the list. Type the number of the device you want to select and press Enter. In our sample configuration, we have to type 01 and press Enter to select the DS3400 (Figure 24-17).

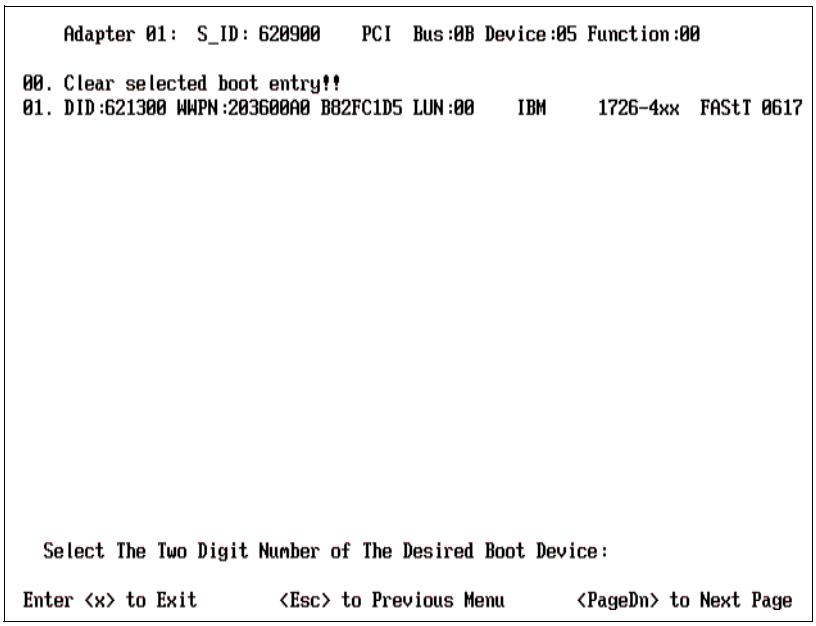


Figure 24-17 Emulex BIOS - Select device

3. You are prompted to specify the starting LUN number to display; a maximum of 16 LUNs are displayed on the screen, so you need to specify a starting point. In our case, we have configured the LUN 0 to be the boot logical drive. Type 00 and press Enter (Figure 24-18).

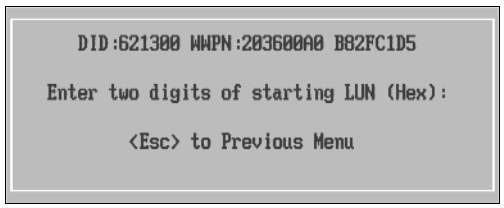


Figure 24-18 Emulex BIOS - Enter LUN

4. You will see a list of LUNs that are configured in the storage device. We only configured one LUN, which is LUN 0, as shown in Figure 24-19. Enter the number corresponding to the LUN that you want to be the boot logical drive; in our case, we type 01 and press Enter.

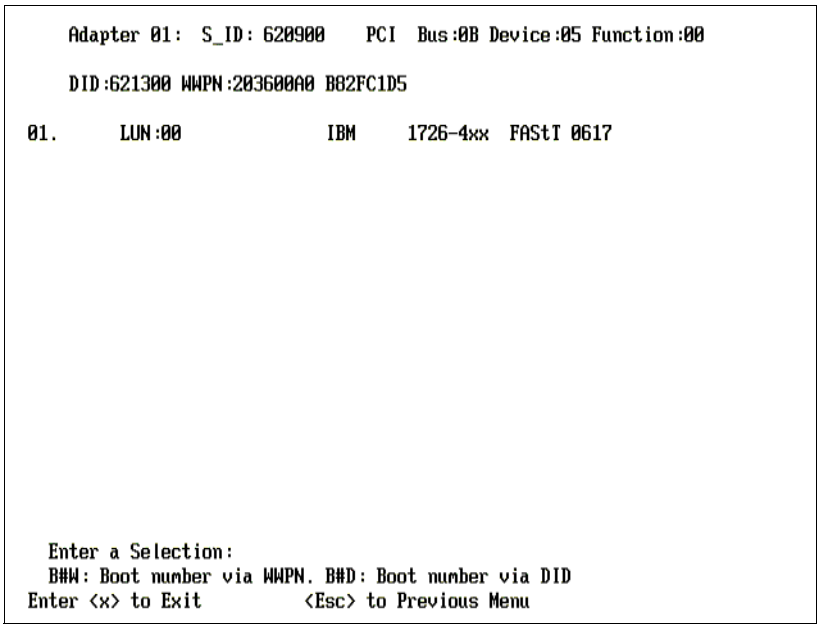


Figure 24-19 Emulex BIOS - Select identifier

5. In the pop-up screen, specify how to identify the boot device. We recommend that you use the WWPN for all boot from SAN configurations. Type 1 and press Enter (see Figure 24-20).



Figure 24-20 Emulex BIOS - Select identifier WWPN

6. Figure 24-16 on page 595 reappears, this time showing the boot device just configured. Press x to save the configuration and exit.

7. Type Y to reboot the system with the changed settings (Figure 24-21).

```
Reboot the System to Make All the Changes to Take Effect!

REBOOT THE SYSTEM (Y/N):
```

Figure 24-21 Emulex BIOS - Reboot

The configuration of the HBA for booting from SAN is now complete.

24.4 Installing Windows 2003 with the Emulex HBA

In the previous section, we set up our environment. Now we need to install Windows 2003 on our boot disk.

24.4.1 Creating the Emulex HBA driver diskette

First, we have to create a driver diskette for the Emulex HBA; you can do this on another server running Windows 2003:

1. Find and download the appropriate Emulex HBA driver kit from the IBM Support site:

   http://www.ibm.com/servers/storage/support/disk

   or directly from Emulex site dedicated to IBM:

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

2. Start the downloaded file to extract the drivers (Figure 24-22). Click **Next**.

> **Note:** The version of current Emulex StorportMiniport Driver will be different than in the following examples.

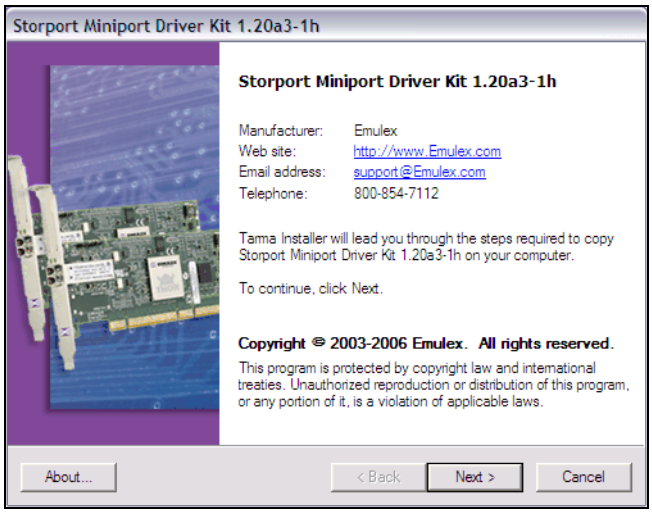


Figure 24-22 Extract Emulex Drivers

3. Type or browse the directory and click **Install** (Figure 24-23 on page 598) to extract the files.

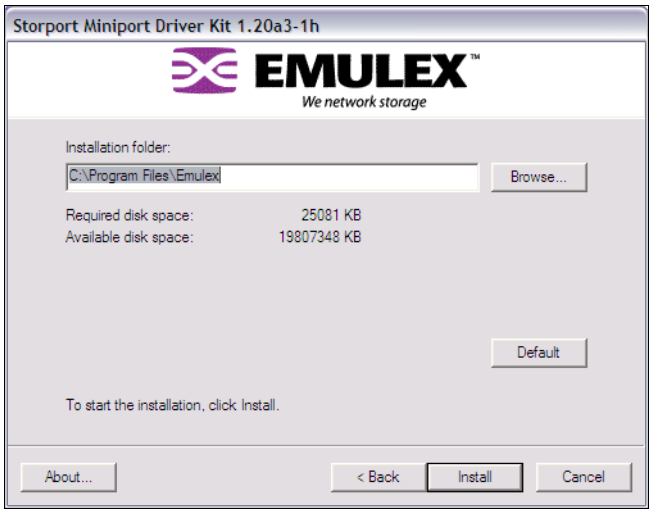


Figure 24-23 Extract Emulex Drivers - Install

4. When the installation is complete, (Figure 24-24), deselect **Start Auto Pilot Installer** and click **Finish**.

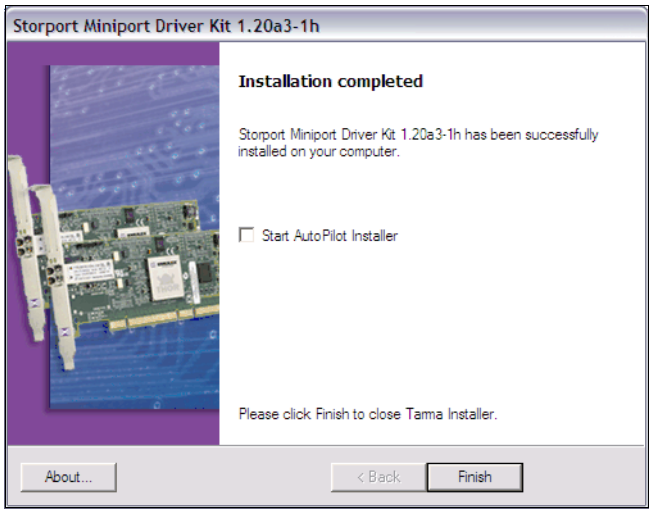


Figure 24-24 Extract Emulex Driver - Finish

5. Open Windows Explorer and navigate to the extracted files. If you have chosen the default location, they are stored in C:\Program Files\Emulex\AutoPilot Installer\Drivers\Storport. Open the directory for the driver you want to use (x86, x64, or IA64) and copy all the files to a diskette.

24.4.2 Installing Windows 2003

Do the following steps:

1. Attach the diskette device to your BladeCenter (with the diskette you just created in it) and boot the blade from a Windows 2003 CD. You can do this over the media tray, or directly from your management station using the Management Module Web interface. At the start of the Windows boot process from the CD, you will be asked if you want to install an additional driver. When you see this message, press F6 (Figure 24-25 on page 599).



Figure 24-25 Press F6 during windows installation

2. The setup asks if you want to specify an additional driver; press S (Figure 24-26) to confirm this.

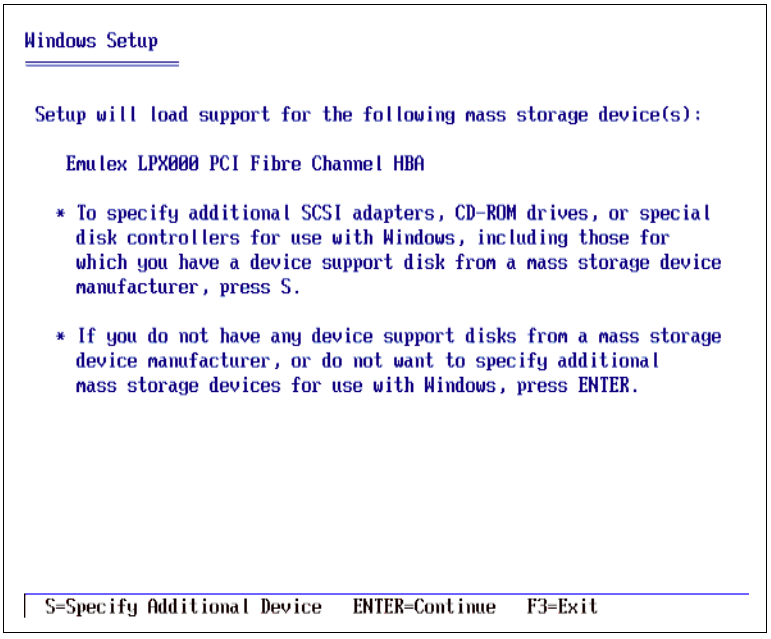


Figure 24-26 Specify driver

3. The driver will load from the diskette drive and you are asked if you want to use it (Figure 24-27). Press Enter to confirm and load the driver. When it is done, it returns to the view shown in Figure 24-26 on page 599. Press Enter to continue the setup.



Figure 24-27 Confirm Emulex driver

4. The installation proceeds as usual. When you have to choose the disk where you want to install Windows (Figure 24-28), the LUN=0 logical drive and an Access LUN can display. You can distinguish the Access LUN because it has a 16 MB size only. Select the logical drive you created for the installation/boot, create a partition on it, and continue installing Windows. The rest of the setup will be like a normal installation. When complete, Windows will boot from the DS3400 logical drive.

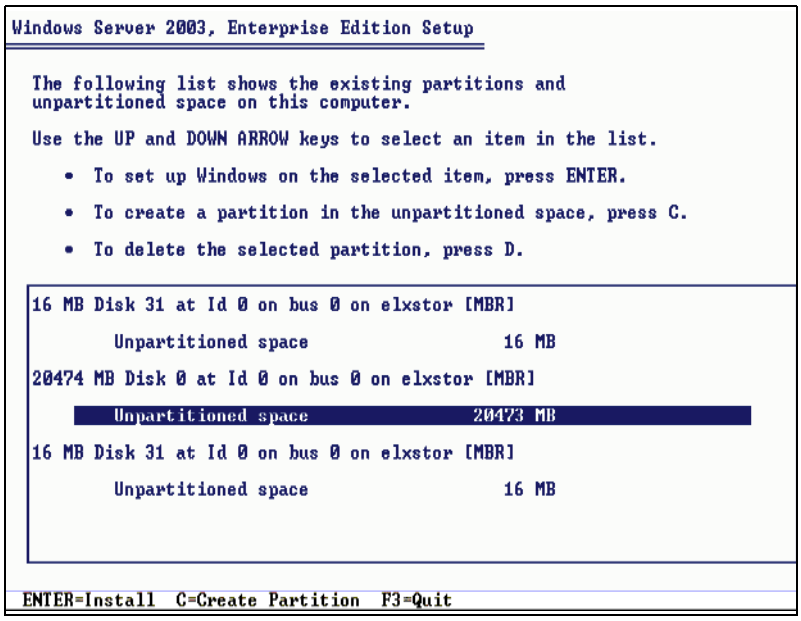


Figure 24-28 Create Partition

From now on, perform a standard installation of Microsoft Windows Server 2003. Apply the latest Service Pack and patches. At a minimum, hotfix KB932755 is required. This hotfix is a requirement for MPIO failover drivers of DS3000. It upgrades MS® Storport driver as required. Next, install hotfix KB932755, which can be downloaded from Microsoft:

http://support.microsoft.com/kb/932755

24.4.3 Installing multipath driver

After the Windows setup is complete, we install the multipath driver on the blade and enable the second HBA to access the logical drive to provide redundancy.

Install IBM DS3000 Storage Manager on the blade, as shown in 6.1, “Installing DS3000 Storage Manager on Microsoft Windows 2003 and 2008” on page 82. If you do not want to use the newly installed blade as a DS3000 management station, you can just install the host part of Storage Manager, which includes at least the multipath driver. You need to reboot after installing the driver.

After the reboot, open DS3000 Storage Manager, and add the second WWN of the HBA to the host that you have created for this blade. Now the blade can access the logical drive over both paths. On the blade, open MS Windows Device Manager and expand the disk drives.

You should now see the DS3400 logical drive two times (redundant connection to “FAStT SCSI Device”) and the multipath drive (“FAStT Multi Path Disk Device”), as shown in Figure 24-29.

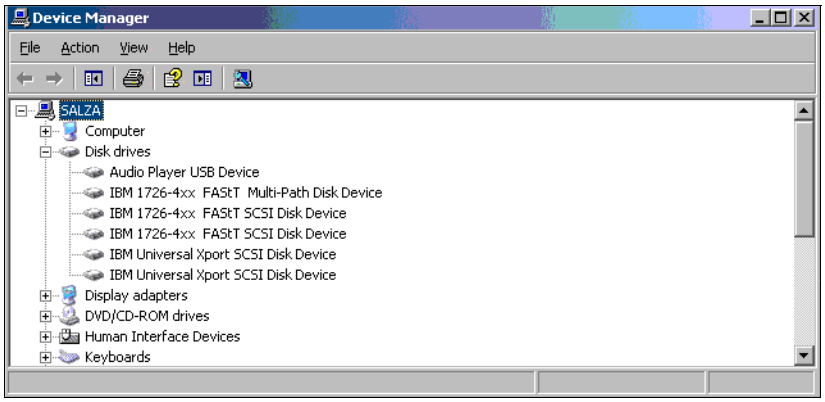


Figure 24-29 Device Manager in Windows 2003 server

Note that in Windows 2008 it looks different. You should see the logical disk drive on the DS3400 (IBM 1726-4xx FAStT Multi-Path Disk Device) as a single device, similar to Figure 24-30. Note that two “IBM Universal Xport” devices in our figure are DS3000 access logical drives, which do not necessarily need to be seen redundantly, as they are used just for in-band management by DS3000 Storage Manager Agent (as described in 3.4.2, “Management methods” on page 30).

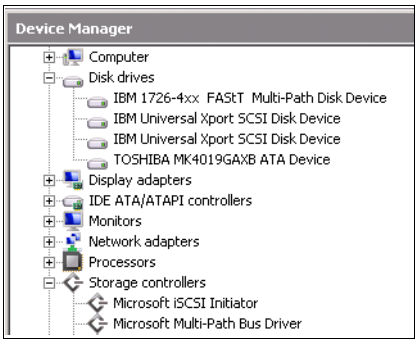


Figure 24-30 Device Manager in Windows 2008

If you select the DS3400 disk and display its properties, you can verify both paths to the LUN in the MPIO tab (Figure 24-31).

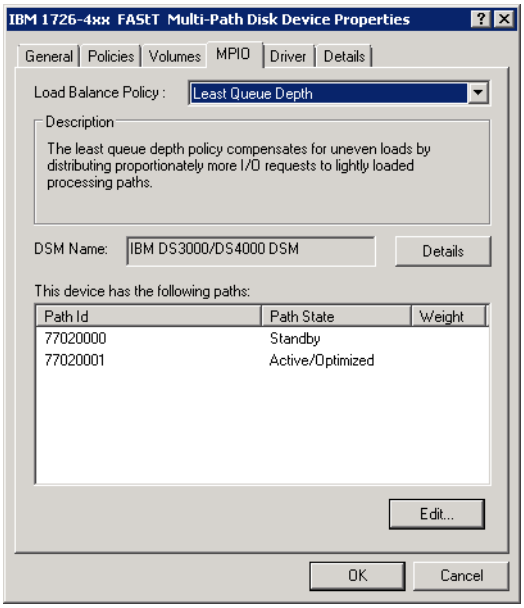


Figure 24-31 MPIO data paths

You have now completed the basic installation of MS Windows Server for SAN boot. You can now download and install additional required drivers for other devices on the blade at:

`http://www.ibm.com/systems/support/supportsite.wss/brandmain?brandind=5000020`

24.5 Installing Emulex HBAnyware

Emulex HBAnyware is a product for managing Emulex HBAs installed in your environment. It is a very extensive program, and describing all the functions of HBAnyware is beyond the scope of this book. However, we want to show some basic functions of this tool.

Do these steps:

1. Copy the Emulex HBA driver package you downloaded in 24.4, “Installing Windows 2003 with the Emulex HBA” on page 597 to your host and start the installation wizard.

   Follow the same steps as before, but this time on the final window (Figure 24-24 on page 598) make sure to check **Start Auto Pilot Installer**. The Emulex Auto Pilot Installer will start (Figure 24-32 on page 603).

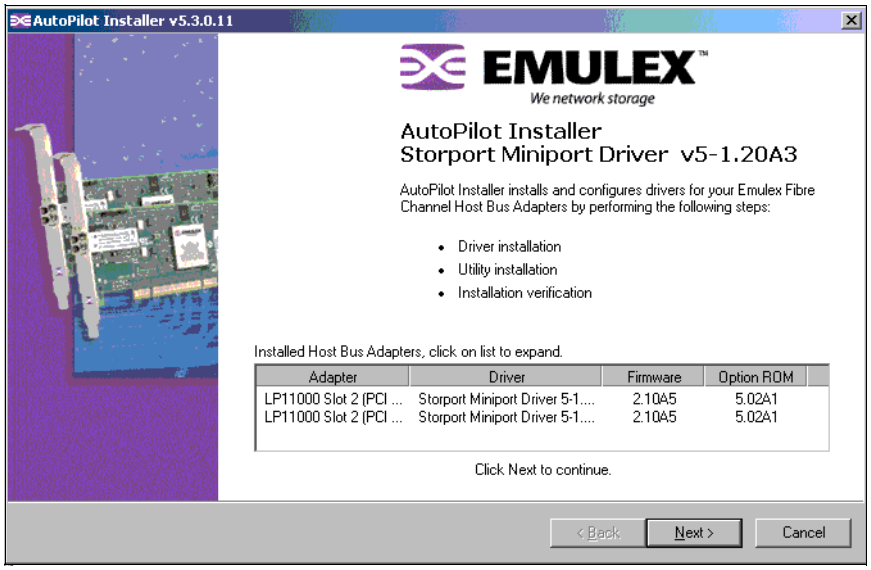


Figure 24-32 Emulex Auto Pilot Installer

2. Click **Next** to display the available HBA drivers (Figure 24-33).

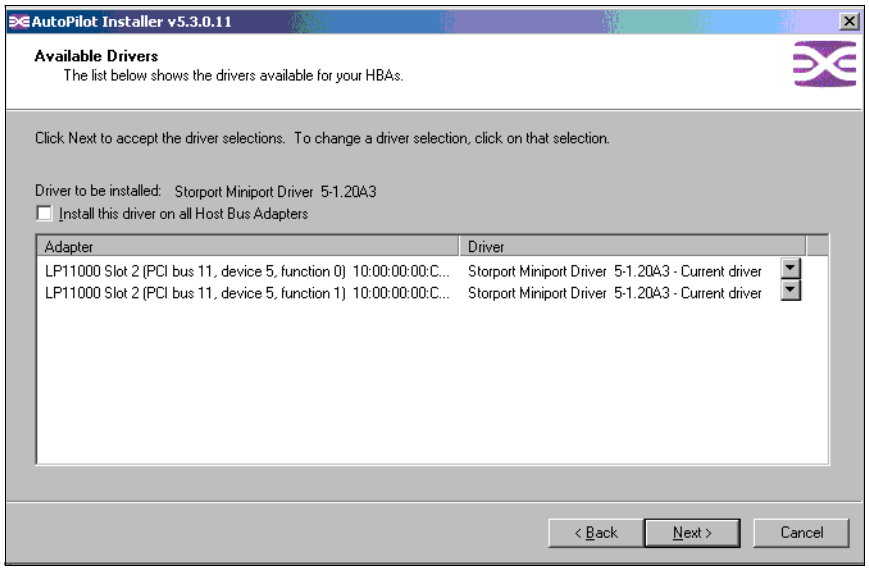


Figure 24-33 Emulex - Available Drivers

3. Click **Next** again. In the pop-up window, you will see that the driver is already installed. To install HBAnyware, you have to reinstall the driver, so select **Yes** (Figure 24-34) to do this.

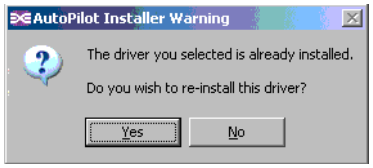


Figure 24-34 Emulex - Already installed

4. You are advised that another restart is necessary (Figure 24-35). Click **OK** to confirm.

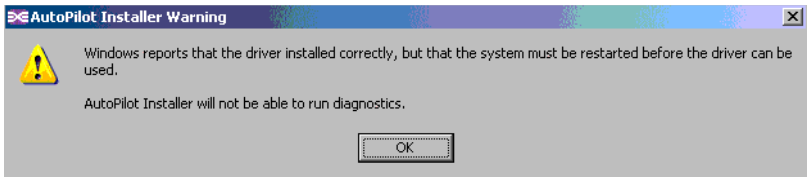


Figure 24-35 Emulex - AutoPilot Installer Warning - Restart

5. The installation history is displayed (Figure 24-36). Click **Next**.

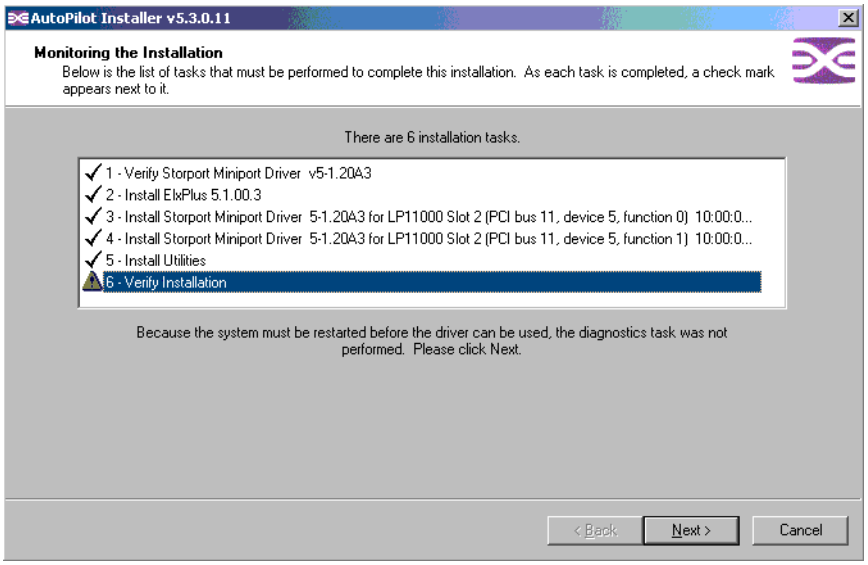


Figure 24-36 Emulex - Monitoring the installation

6. On Figure 24-37, click **Finish** to complete the installation and reboot.

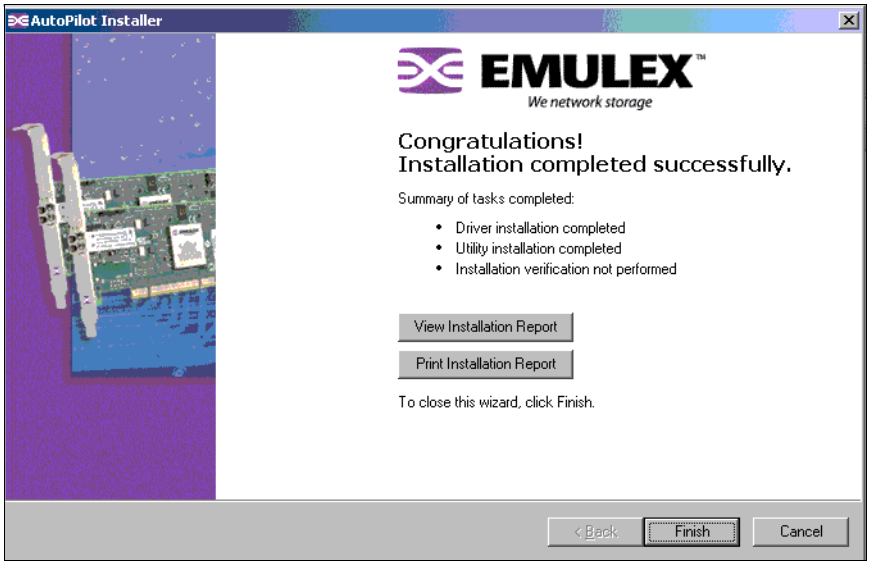


Figure 24-37 Emulex - HBAnyware installed successfully

7. After rebooting, start HBAnyware by selecting **Start** → **All Programs** → **Emulex** → **HBAnyware**. Figure 24-38 shows the opening window. The left pane shows the hosts with the Emulex HBAs that were discovered. The application allows you to monitor both the local and other hosts, but we will only cover local monitoring.

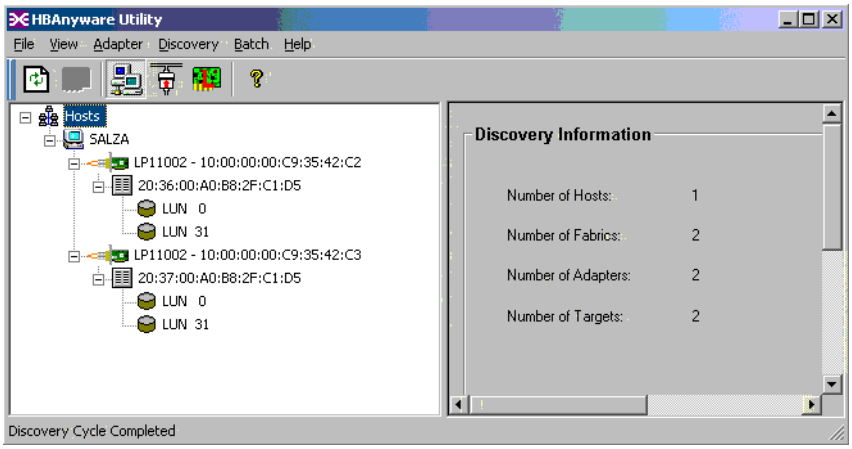


Figure 24-38 HBAnyware

Expand your host to see all the HBA ports of the host. Expand the HBA port to show the WWNs of all the connected targets. In our sample configuration, one HBA port is only able to see one DS3400 controller port. Figure 24-38 shows that the HBA port with a WWN ending in *C2* can see a target with the WWN 20:36:00:A0:B8:2F:C1:D5. This is the WWN of the used port on the DS3400 used in the sample configuration. By expanding this port, you also see the LUNs, which the HBA can access over this path.

8. The right pane displays a summary of a highlighted item in the left pane. For example, in Figure 24-39, “LUN 0” is highlighted, and the summary shows that this is a DS3400 device.

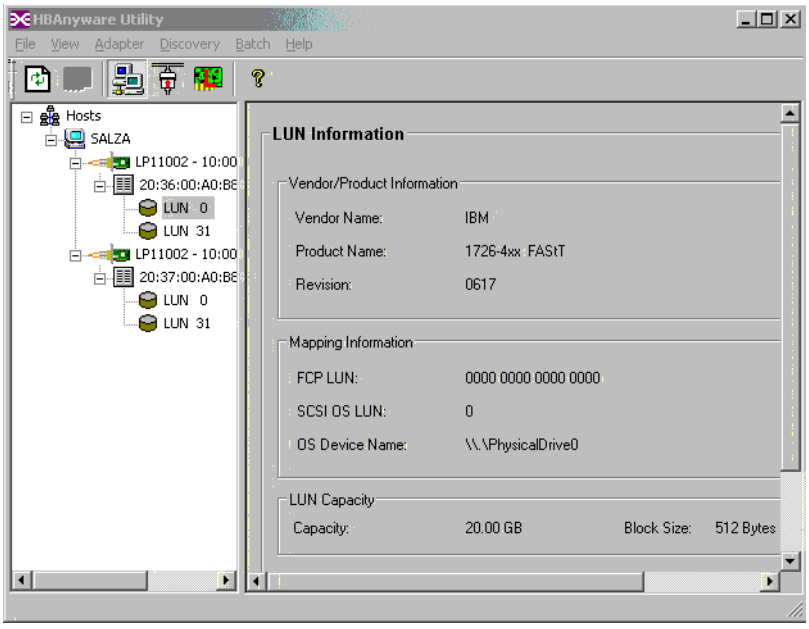


Figure 24-39 Emulex - HBAnyware - LUN Information

9. You can update the firmware and boot code of an Emulex HBA using a boot diskette, or with HBAnyware. Download the latest Emulex HBA firmware and boot code from the IBM Support Web page and extract it to your hard drive. In HBAnyware, select the host and select **Batch** → **Download firmware** from the menu bar (Figure 24-40).

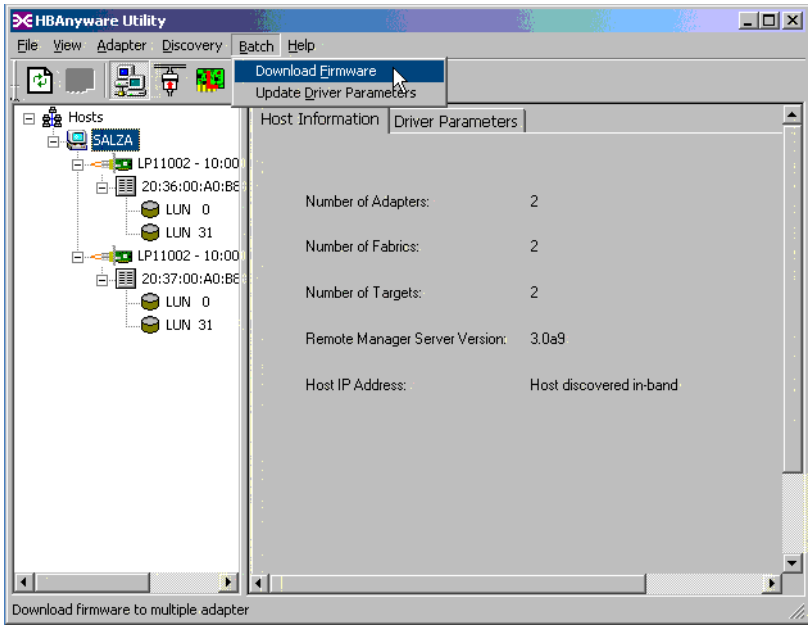


Figure 24-40 Emulex - HBAnyware - Download Firmware

10. Navigate to the extracted firmware file, select the firmware image, and click **Open** (Figure 24-41).

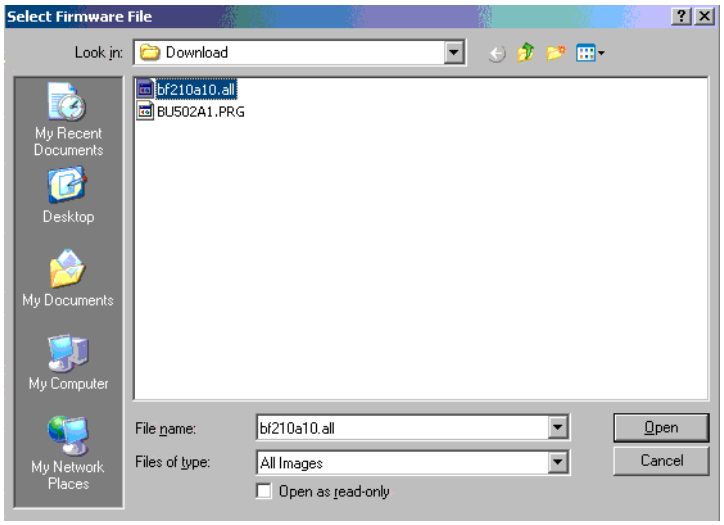


Figure 24-41 Emulex - HBAnyware - Select Firmware File

11. The HBAnyware Batch Firmware Download window appears (Figure 24-42). Mark all the HBAs you want to update and click **Start Download**. The image will be transferred to the HBA; this may take a few seconds.

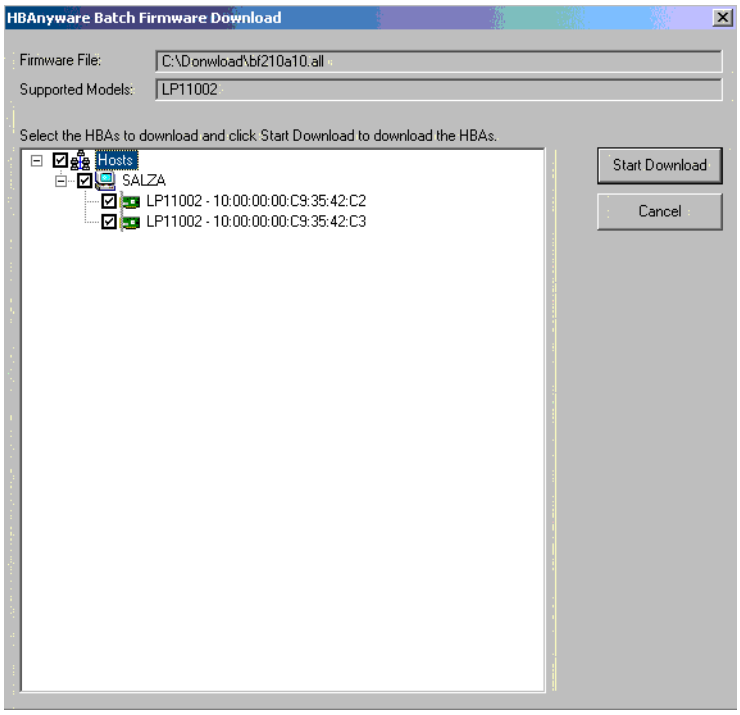


Figure 24-42 Emulex - HBAnyware - Batch Firmware Download

12. Figure 24-43 shows that the download was successful. Review the summary and click **Close** to exit the window.

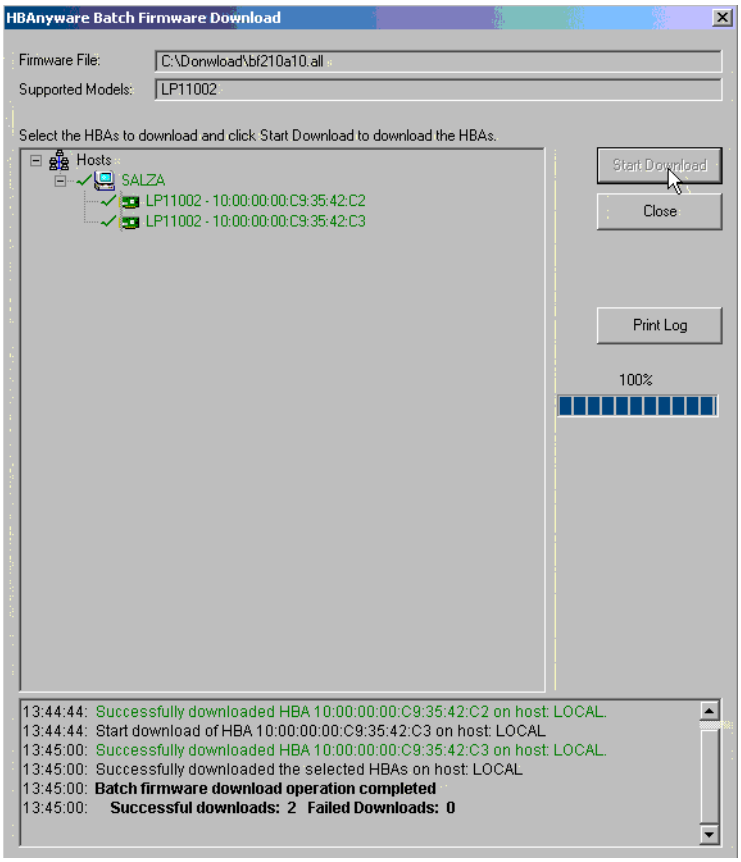


Figure 24-43 Emulex - HBAnyware - Batch firmware download successful

The firmware of the HBA is now updated successfully; repeat the same steps, but this time select the downloaded boot code image. The Emulex HBA does not need to reboot to update the firmware or the boot code.

24.6 Determining what to do if the boot path fails

If the boot path fails, follow the procedures in this section to correct the problem.

24.6.1 Boot from SAN with one HBA

Remember that while configuring the HBA, we enabled the Adapter BIOS on one adapter port (24.3.2, “Enable adapter BIOS” on page 593). The blade is always only able to boot from one port, since at the start of the boot process, the multipath driver is not loaded, so only one device is allowed to access the logical drive. This is normally not a problem, but if the boot path fails, you will not be able to reboot the blade.

Figure 24-44 shows a situation where the boot path has failed. This can occur, for example, because of a broken cable. So long as the system is still running, there is no problem, because the redundant connection can be used to access the logical drive over the second path.

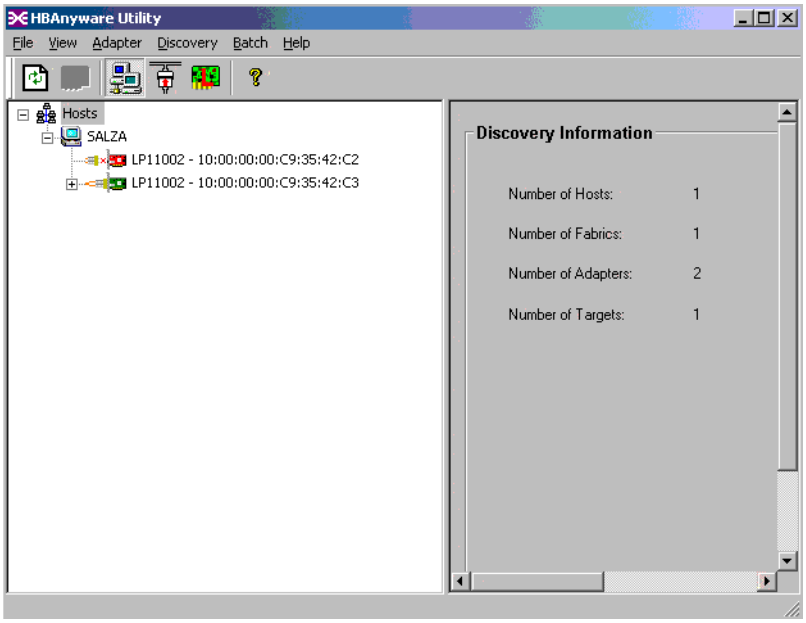


Figure 24-44 Emulex - HBAnyware - Link down

If you have to reboot the blade, you have two possibilities:

- Fix the path failure before rebooting the blade.
- Edit some settings before rebooting the blade.

If you choose the second way, you have to do the following tasks:

1. Reboot the server and enter the Emulex BIOS (refer to “FC Emulex HBA” on page 136).
2. Disable the Adapter BIOS on the Adapter where it is currently enabled, as described in 24.3.2, “Enable adapter BIOS” on page 593, but this time choose 2 to disable the BIOS.
3. Now configure the boot device and enable the BIOS on the port that is connected to the second path, following the instructions in 24.3, “Enabling boot from SAN with Emulex adapters” on page 592.
4. Save the configuration changes on the adapter and reboot the server. Now the operating system will boot up using the other HBA path.

24.6.2 Multi-port failover Boot from SAN

For high availability of the SAN boot solution, you can set both ports on the FC expansion card to be capable of Boot from SAN. You can boot your blade server even if the first FC switch module in the BladeCenter is temporarily removed, offline, or broken.

Emulex FC expansion cards for blades can support path failover also at boot time. It is not maintained by the MPIO driver in MS Windows, as this driver is not loaded at boot time, but SAN boot failover is supported by features of the Blade System BIOS. This enable the use of each FC port for booting from the same boot logical drive in DS3000.

You have to customize the boot sequence in the blade's BIOS configuration. If the first FC port loses access (FC link) to the previously configured boot LUN, you can set the other FC port of the Emulex card as the next (alternate) boot device. You have to set all the parameters of the second port in the same way, as described in 24.3, "Enabling boot from SAN with Emulex adapters" on page 592.

During the setting of the alternate SAN boot path to second port of the Emulex Blade FC expansion card, you must temporarily move the boot logical drive controller ownership to the second DS3000 controller, which is zoned to the second FC port.

Note: The operating system has to be already installed when you add a second booting path. The OS installation is possible with only one FC path to the DS3000 server.

Depending on your configuration of the SAN environment, it should be useful to create an FC zone with both DS3000 controllers visible by the first FC port, which has to be done after a successful OS installation. In this situation, the boot is provided by the first FC port when the Blade FC switch and FC links to the first FC port are correct, but failover happened somewhere else in the SAN or in the DS3000 internally. It means that the first controller is not accessible. In this situation, your server can boot using the first FC port even though the boot logical drive is owned by the second DS3000 controller.

25

FC configuration 2 - Blade server Microsoft Cluster

In this chapter, we describe a sample configuration to attach a DS3400 server to Microsoft Windows Server 2003 and how to create a Microsoft Cluster. In our sample configuration, we use an IBM BladeCenter, but from a storage perspective, it is the same as installing it on an IBM System x server.

This example is prepared for MS Windows Server 2003, but very similar steps have to be done on a DS3400 server in case you want to install an MS Windows Server 2008 cluster.

We also provide two simple examples of SMcli scripts for some configuration steps on a DS3400 Storage Subsystem.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

25.1 Preparing the environment

First, we perform some initial setup: the infrastructure cabling and zoning, and DS3000 basic configuration.

25.1.1 Cabling and zoning

In 4.5, “Host attachment” on page 40, we show some possible host attachments. In our sample configuration, we use the recommended redundant attachment (Figure 4-9 on page 45). In a BladeCenter environment, this will appear as shown in Figure 24-2 on page 586. Prepare the environment so that every SAN HBA port of each blade can see one controller of your storage subsystem. How to create the zoning on the IBM FC switches is shown in 24.1, “Cabling and zoning” on page 586.

25.1.2 Creating logical drives and host mapping with an SMcli script

After you have cabled and zoned your SAN, configure your DS3400 storage subsystem. This includes the following steps:

- Create a host: Refer to “Configure host access (manual)” on page 145. Add both HBAs (Ports) to the host.
- Create a logical drive(s) and map them to the host: Refer to “Create logical drives” on page 168. Create one or more logical drives and map them to the host after creating them.

These referenced sections describe how to accomplish our task using the GUI (you can also use a script). Example 25-1 shows a CLI script that creates two logical drives and maps them to a host.

Example 25-1 Script - Create hosts - Logical drives and map them

```
show "Setting the Storage Subsystem user label to DS3400_Thailand.";
set storagesubsystem userLabel="DS3400_Thailand";

show "Create a host named PaSak";
create host userLabel="PaSak";

show "Add host ports to host";
create hostPort identifier="210000e08b849e87" userLabel="PaSak_Port0" host="PaSak"
hostType="Windows 2000/Server 2003 Non-Clustered";
create hostPort identifier="210100e08ba49e87" userLabel="PaSak_Port1" host="PaSak"
hostType="Windows 2000/Server 2003 Non-Clustered";
show allHostPorts;

show "Creating hot spare at Enclosure 0 Slot 1.";
set drive[0,1] hotSpare=true;

show "Creating Raid 1 Logical Drive PaSak_Raid1 on new Array 1.";
// This command creates the Array and the initial Logical Drive on that array.
// NOTE: For Arrays that use all available capacity, the last Logical Drive on
this array is
// created using all remaining capacity by comitting the capacity= logical drive
creation parameter.
create logicaldrive drives[0,2 0,3] RaidLevel=1 userLabel="PaSak_Raid1" owner=A
segmentSize=128 capacity=60 GB;
```

```
show "Setting additional attributes for Logical Drive PaSak_Raid1.";
// Configuration settings that can not be set during Logical Drive creation.
set logicaldrive["PaSak_Raid1"] cacheFlushModifier=10;
set logicaldrive["PaSak_Raid1"] cacheWithoutBatteryEnabled=false;
set logicaldrive["PaSak_Raid1"] mirrorEnabled=true;
set logicaldrive["PaSak_Raid1"] readCacheEnabled=true;
set logicaldrive["PaSak_Raid1"] writeCacheEnabled=true;
set logicaldrive["PaSak_Raid1"] mediaScanEnabled=true;
set logicaldrive["PaSak_Raid1"] redundancyCheckEnabled=false;
set logicaldrive["PaSak_Raid1"] readAheadMultiplier=1;
set logicaldrive["PaSak_Raid1"] modificationPriority=high;

show "Creating Raid 5 Logical Drive PaSak_Raid5 on new Array 2 / using the whole
capacity.";
create logicaldrive drives[0,5 0,6 0,8 0,11] RaidLevel=5 userLabel="PaSak_Raid5"
owner=A segmentSize=128;
show "Setting additional attributes for Logical Drive PaSak_Raid5.";
set logicaldrive["PaSak_Raid5"] cacheFlushModifier=10;
set logicaldrive["PaSak_Raid5"] cacheWithoutBatteryEnabled=false;
set logicaldrive["PaSak_Raid5"] mirrorEnabled=true;
set logicaldrive["PaSak_Raid5"] readCacheEnabled=true;
set logicaldrive["PaSak_Raid5"] writeCacheEnabled=true;
set logicaldrive["PaSak_Raid5"] mediaScanEnabled=true;
set logicaldrive["PaSak_Raid5"] redundancyCheckEnabled=false;
set logicaldrive["PaSak_Raid5"] readAheadMultiplier=1;
set logicaldrive["PaSak_Raid5"] modificationPriority=high;


show "Map logical drives to an host";
set logicalDrive ["PaSak_Raid1"] logicalUnitNumber=1 host="PaSak";
set logicalDrive ["PaSak_Raid5"] logicalUnitNumber=2 host="PaSak";
```

25.2 Installing Microsoft Windows and device drivers

Perform a standard installation of Microsoft Windows Server 2003 on your host server. Apply the latest Service Pack and patches. At a minimum, hotfix KB932755 is required. This hotfix is a requirement for MPIO failover drivers of DS3000. It upgrades MS Storport driver.

Install hotfix KB932755, which can be downloaded from Microsoft at:

http://support.microsoft.com/kb/932755

When installing Windows Server 2003 on a physical System x or BladeCenters, refer to the installation guide for that server for more details about installing the operating system.

You can find all drivers for completing the installation at:

http://www-304.ibm.com/jct01004c/systems/support/

In the case of a Microsoft Windows Server 2008 installation, all the necessary drivers for MPIO are included, but verify the recommendation for MPIO in the readme file of the DS3000 Storage Manager package.

25.2.1 Installing HBA drivers

The IBM Support Web site also provides the QLogic HBA drivers.

Note: In our configuration, we are using QLogic HBAs; however, HBAs from Emulex are also supported. Download the drivers from the Emulex Web site and install them as directed. The other instructions and procedures should be similar.

Once the drivers are installed, the subsequent procedure is identical.

Do these steps:

1. Download the correct drivers for your system and extract the zip file to the hard drive (Figure 25-1).

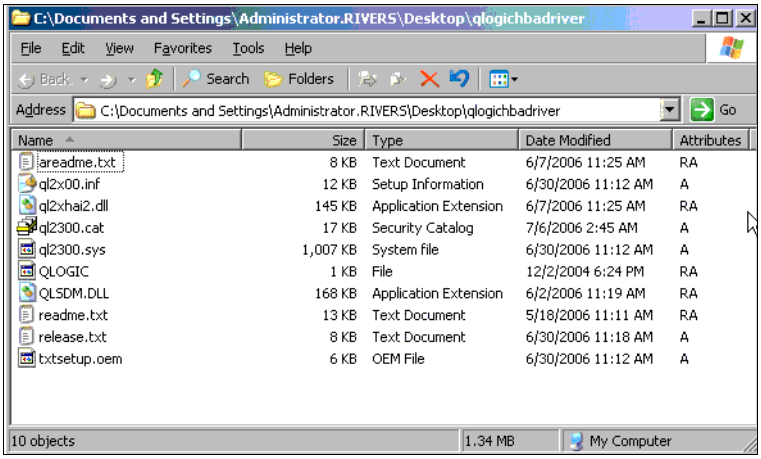


Figure 25-1 Extracted QLogic Drivers

2. Open Windows Device Manager. Right-click one of the yellow question marks and select **Update Driver** (Figure 25-2).

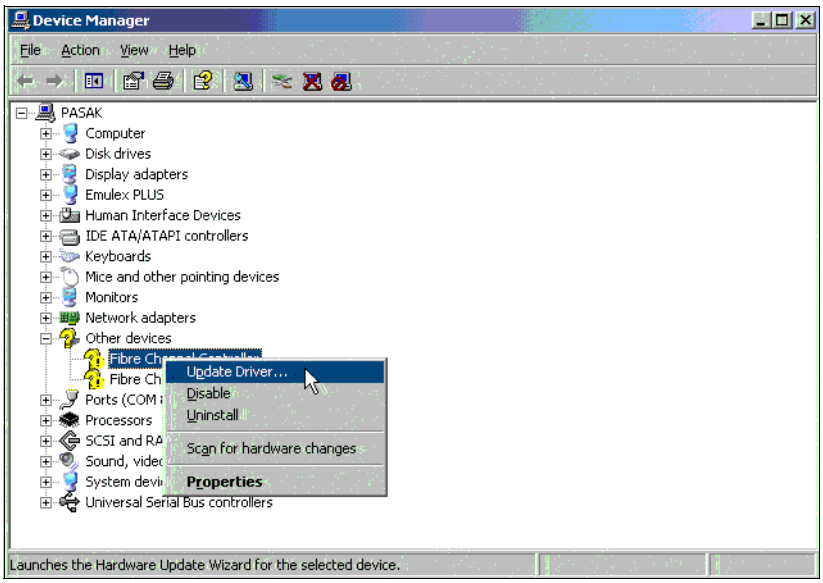


Figure 25-2 Update Driver

3. In the Hardware Update Wizard (Figure 25-3), check **No, not this time** and click **Next**.



Figure 25-3 Hardware Update Wizard

4. In the next window, check **Install from a list or specific location (Advanced)** and click **Next**.
5. In Figure 25-4, check only **Include this location in the search** and click **Browse**.

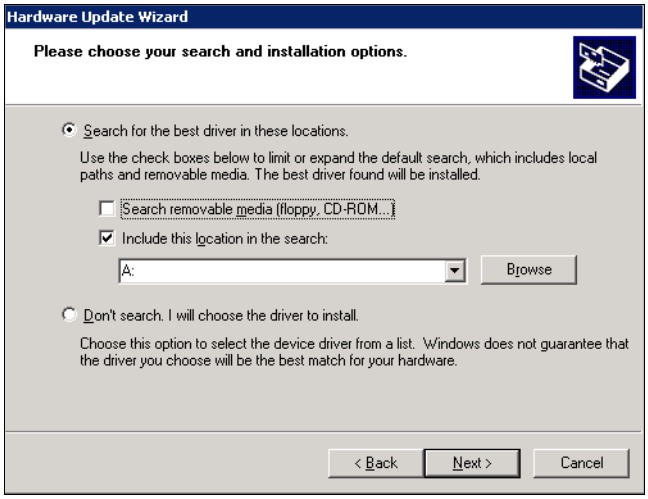


Figure 25-4 Include this location in the search

6. Highlight the folder where you have extracted the HBA driver and click **OK** (Figure 25-5).



Figure 25-5 Browse for Folder

7. the window shown in Figure 25-4 on page 615 reappears. Click **Next** to install the HBA driver (Figure 25-6).

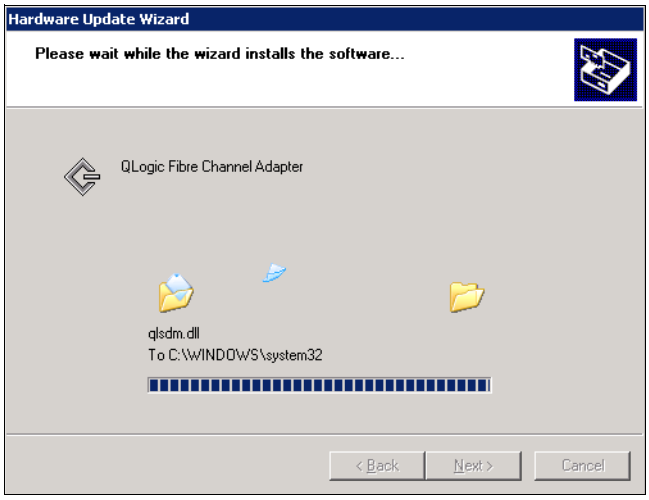


Figure 25-6 QLogic driver installation

8. When the installation is complete, click **Finish** (Figure 25-7 on page 617) to close the wizard. The HBA drivers have been installed successfully.



Figure 25-7 QLogic driver installation complete

25.2.2 Installing the multipath driver

The next step is to install the multipath driver. Do these steps:

1. Start Storage Manager on your host server. A detailed description for installing the Storage Manager is given in 6.1, “Installing DS3000 Storage Manager on Microsoft Windows 2003 and 2008” on page 82. Follow the steps, select either the **Host** installation type or **Custom** installation (Figure 25-8) and click **Next**.



Figure 25-8 Multipath driver installation

2. You will be asked to decide which modules to install for the Custom installation. Do not forget to select the MPIO/RDAC drivers (in Figure 25-9). Click **Next** and then **Install**.
3. To complete the installation, you have to reboot your server.

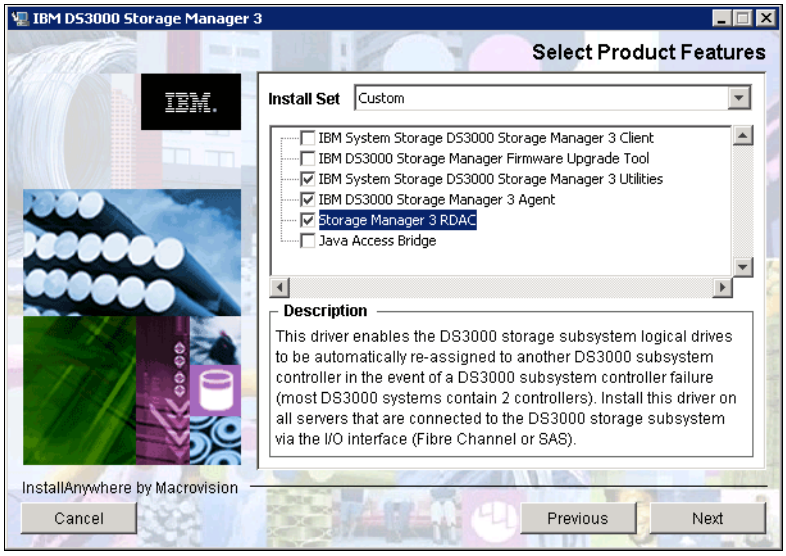


Figure 25-9 Install MPIO/RDAC driver

4. After the reboot, open Device Manager and expand **Disk drives**. Remember that, in the script (Example 25-1 on page 612), we created two logical drives. In Device Manager (Figure 25-10), you now see all the disks twice (redundant connections) and also two Multi-Path Disk Devices. This verifies that the operating system is seeing all paths, and the two created logical drives are presented to the OS by the multipath driver.

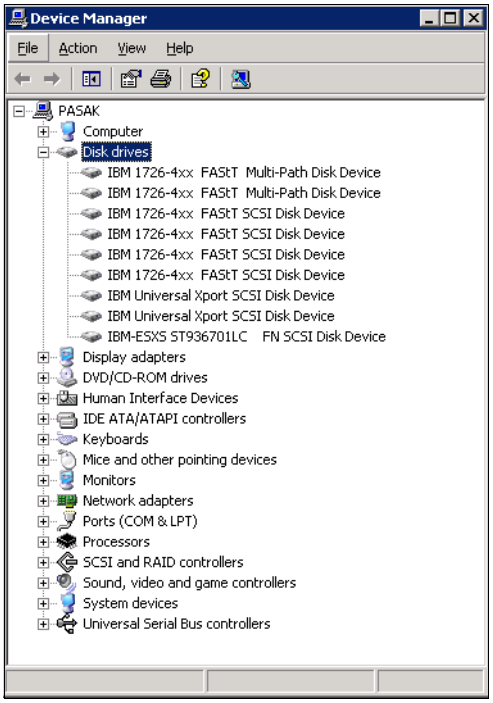


Figure 25-10 Device Manager in Windows 2003 Server

Note that in Windows 2008 Server you see your two previous configured logical disk drives on the DS3400 as single devices only (IBM 1726-4xx FAStT Multi-Path Disk Device), as shown in Figure 25-11.

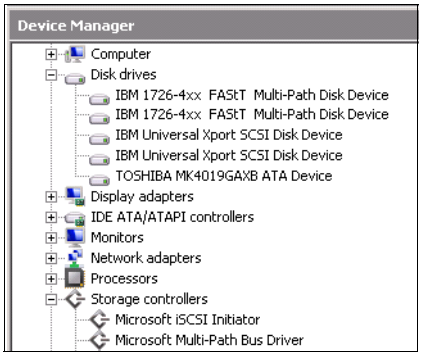


Figure 25-11 Device Manager - Disk Drives

You can select the properties of these logical drives and verify that you can see valid redundant paths in the MPIO tab, as shown in Figure 25-12.

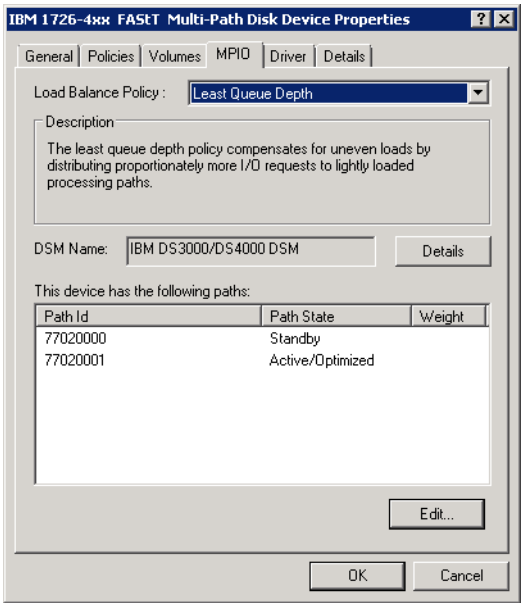


Figure 25-12 MPIO tab in Windows 2008 Server

25.2.3 Windows Disk Manager

At this point, we have done the basic steps to attach the SAN disks to the operating system. Now we have to prepare the disks in Microsoft Disk Manager so they can be used.

Do these steps:

1. Start Disk Manager by selecting **Start** → **Run** and typing in diskmgmt.msc. The disk management window appears, and because it is discovering new disks, the Initialize and Convert Disk Wizard will appear in the foreground (Figure 25-13).



Figure 25-13 Initialize and Convert Disk Wizard

2. Click **Next** and you are prompted for which disks to initialize (Figure 25-14). By default, all newly discovered disks are displayed. Mark all disks you want to initialize and click **Next**.

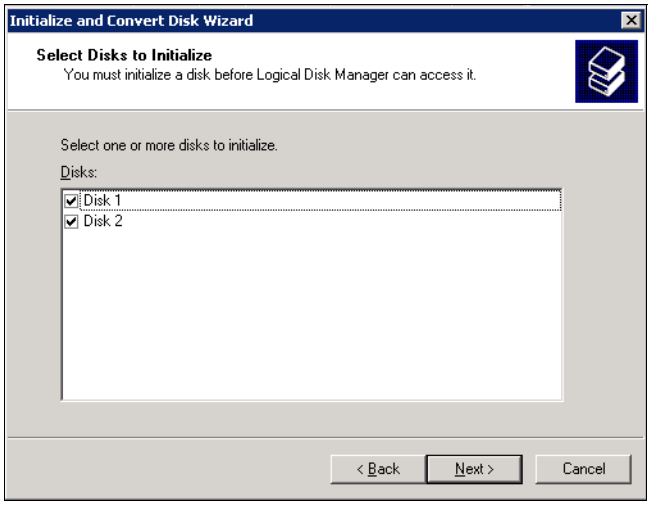


Figure 25-14 Select disks to initialize

3. If you want to convert one of the disks to a dynamic disk, mark it here and click **Next** (Figure 25-15 on page 621). However, dynamic disks are not supported for Microsoft Cluster, so we will not do this.

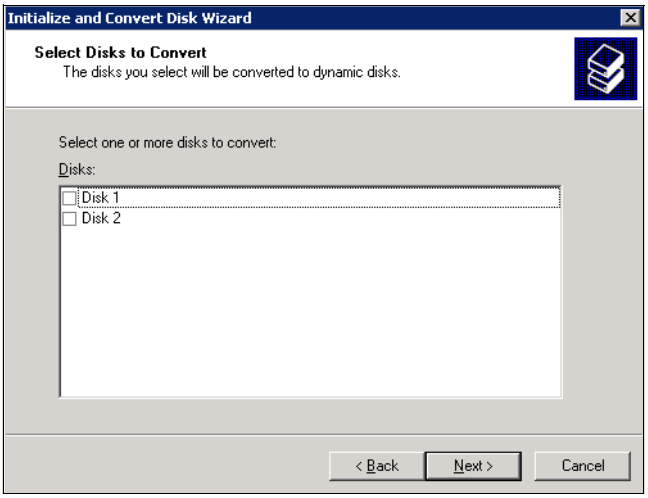


Figure 25-15 Select disks to convert

4. The Completing the Initialize and Convert Disk Wizard shows you a summary of your selections (Figure 25-16). Review them and click **Finish** to complete the tasks.

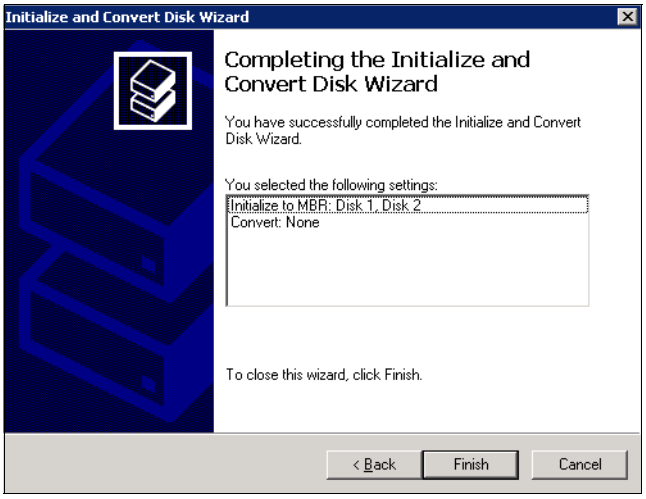


Figure 25-16 Completing the Initialize and Convert Disk Wizard

5. The wizard closes and the new disks are available in Disk Manager (Figure 25-17).

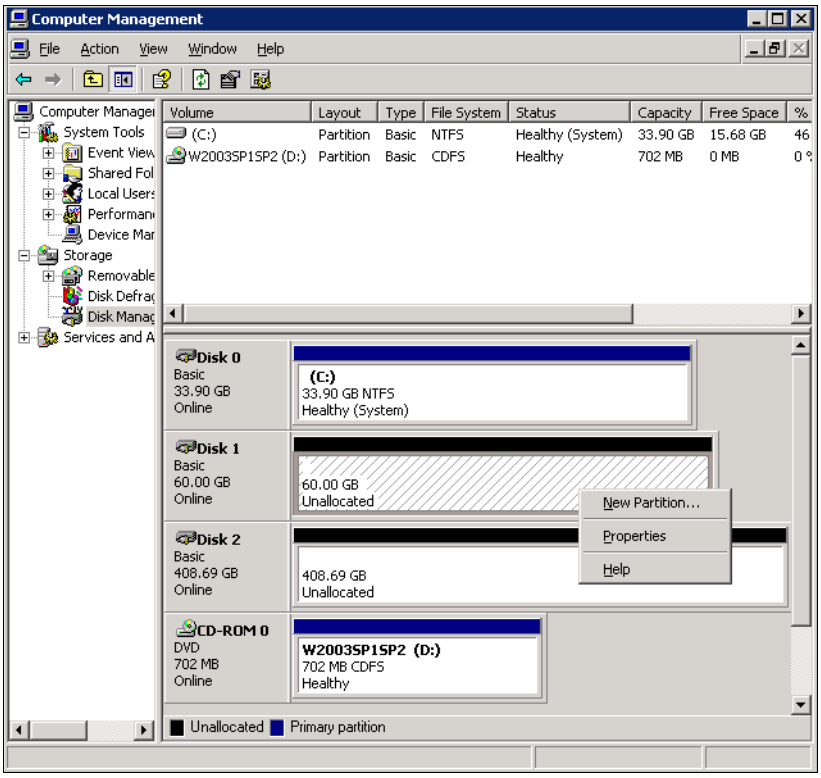


Figure 25-17 Disk Management

6. To use the disks in Windows, you have to create a partition on them and format them. To do this, right-click one disk and select **New Partition** (as shown in Figure 25-17).
7. The New Partition Wizard (Figure 25-18) appears. Click **Next**.

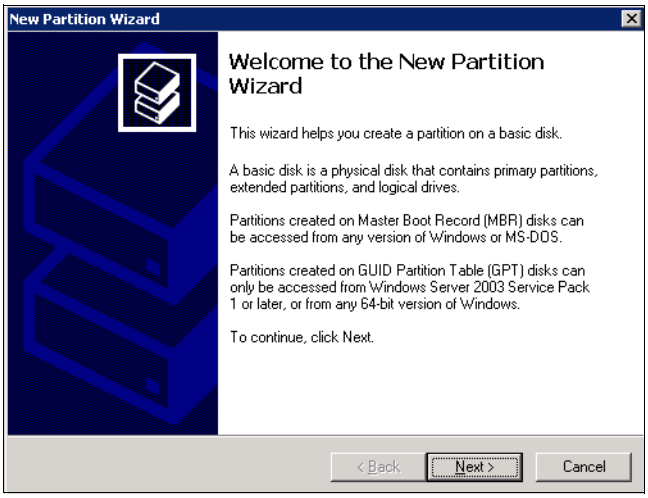


Figure 25-18 New Partition Wizard

8. In Figure 25-19 on page 623, mark the partition type you want to create and click **Next**.

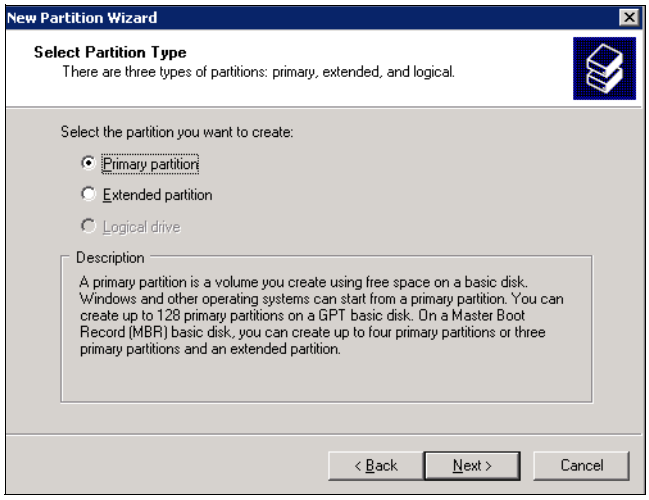


Figure 25-19 Select Partition you want to create

9. Enter the size of the new partition in MB and click **Next** (Figure 25-20).

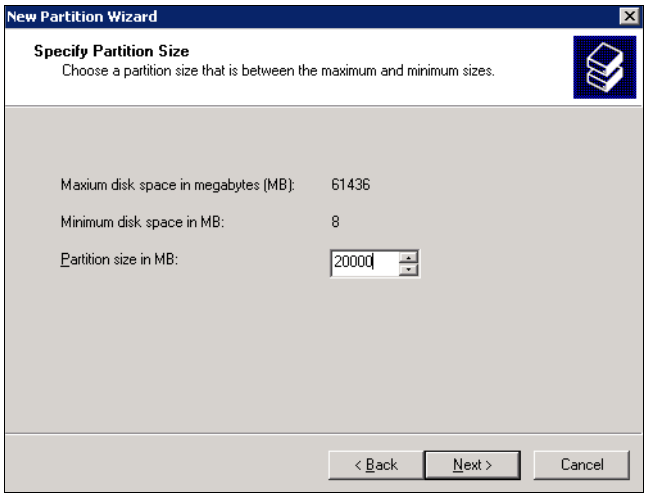


Figure 25-20 Enter disk space

10. Assign a drive letter or mount this space into an existing NTFS Folder and click **Next** (Figure 25-21).

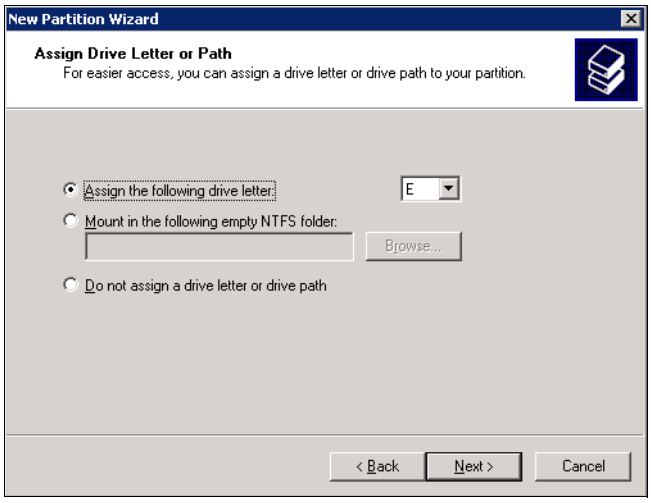


Figure 25-21 Assign drive letter or path

11. Format the partition. Select your requirements and click **Next** (Figure 25-22).

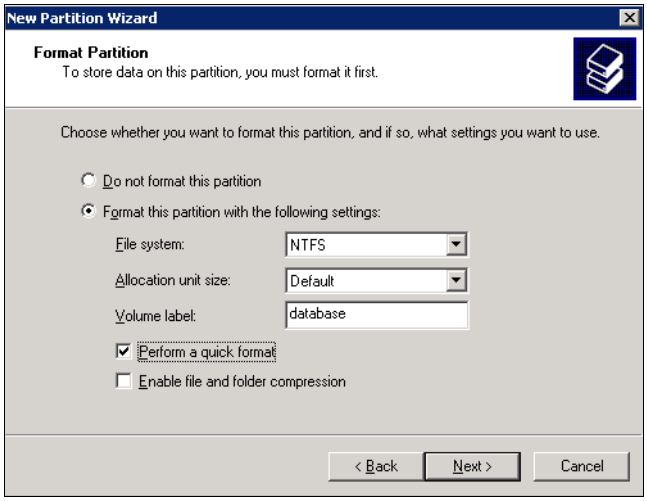


Figure 25-22 Format Partition

12. The Completing the New Partition Wizard view appears (Figure 25-23 on page 625). Review your selections and click **Finish** to complete the tasks.

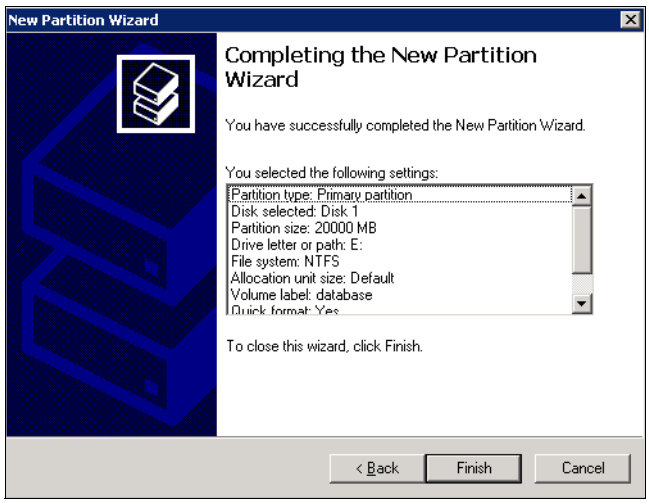


Figure 25-23 Completing the new partition

Now the disk will appear in Windows Explorer and can be used. Repeat these steps on each new disk.

25.3 Creating a Microsoft Cluster

In the first part of this chapter, we attach the DS3400 to a host server named PaSak. Now we will attach a second host server, and use the DS3400 as a shared cluster device.

25.3.1 Installing the second node

Do these steps:

1. Install the operating system on a second identical host. After installing the operating system, review the checklists for Microsoft Cluster for specific recommendations. The checklists are available in the online help of Windows Server 2003 or 2008. We assume you are familiar with Microsoft Cluster installation, and provide summary information only here. Our cluster environment is shown in Figure 25-24.

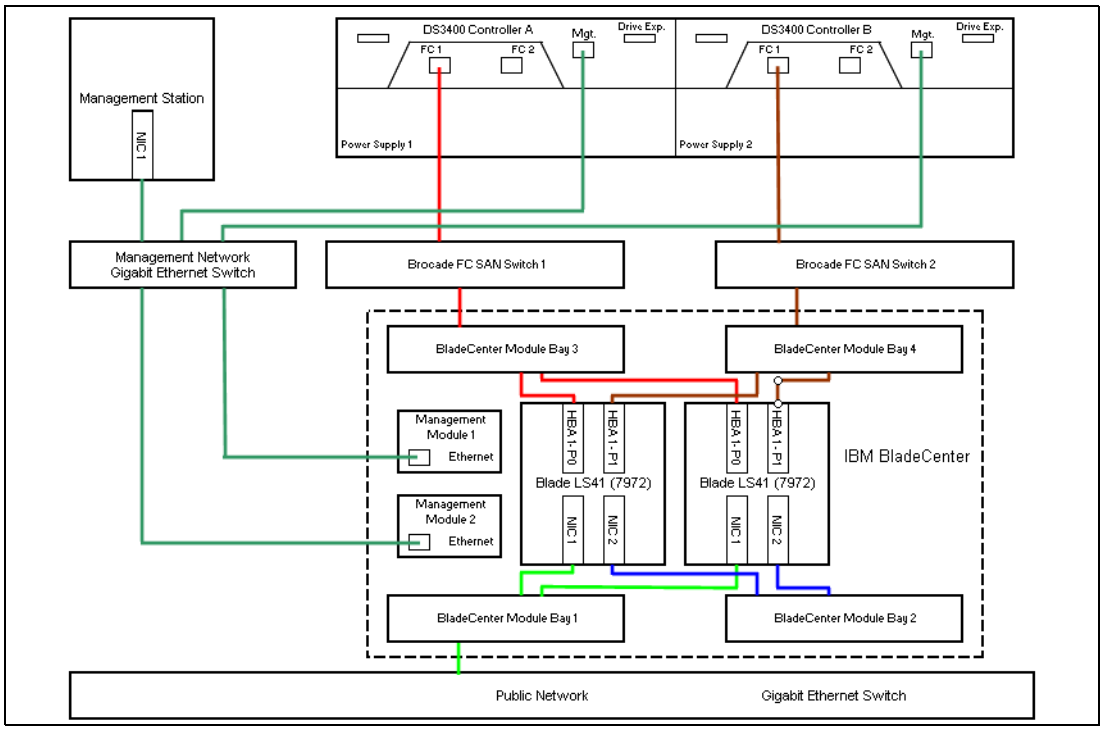


Figure 25-24 FC example environment

2. Configure the network interfaces that are required. Our configuration is shown in Table 25-1 on page 627. We recommend that you use static IP addresses. After setting the IP addresses, go to My Network Places and rename the network connections to represent their functions. To do this, right-click each interface in turn and select **Rename**. This makes identifying each network easier.
3. In Table 25-1 on page 627, you see that card NIC 0 can access the public LAN. Configure this IP address so that it suits your environment. The second network card NIC 1 cannot access the public LAN. This NIC will only be used for the cluster heartbeat. Assign an IP address from another range to this NIC. In our sample configuration, we use the IP addresses in Table 25-1 on page 627.

Table 25-1 IP Config in Sample Configuration

Component	NIC 0	NIC 1
Host PaSak (Node 1)	172.19.0.135	10.0.0.1
Host Yom (Node 2)	172.19.0.136	10.0.0.2
Cluster Address	172.19.0.140	
Domain Controller	172.19.0.8	

4. Join both nodes into the same domain. Active Directory installed on the cluster nodes is not supported; a non-clustered node is required. The domain membership is used to provide a domain account available for the cluster service account on all nodes. That is the account under which the cluster service will operate.
5. Create a domain account that can be used later on the cluster nodes to run the cluster service. There are no special requirements for this account; all necessary rights will be granted to this account when the cluster service is set up on the first node.
6. Modify your SAN zoning, if required, so that both hosts are able to connect to the DS3400.

Note: The second node will be able to see the same logical drives as the first node, although the cluster service is not installed yet. Do not initialize the logical drives on the second node!

25.3.2 Editing the storage subsystem configuration for clustering by SMcli

Now we have to do some preparations on the storage subsystem. Previously, we configured it so that only one host can access the logical drives, but in a clustered configuration, at least two hosts will access the same logical drive. This requires the following steps:

1. Create the second host: Refer to 9.2, “Configure hosts” on page 140. Create a new host (Yom in our configuration), add its WWNs, and select the host type “W2k3 Clustered”.
2. Create a host group: Refer to “Host group” on page 149. Create a host group and add the two hosts that will join the cluster to this group.
3. Edit the host type of the first host created in 25.1.2, “Creating logical drives and host mapping with an SMcli script” on page 612. Refer to 11.1.1, “Edit Host Topology” on page 223. In our sample configuration, we initially defined the first host with type “not clustered”; change this to “clustered”.
4. Create a cluster quorum drive: Refer to “Create logical drives” on page 168. Create a RAID 1 Cluster Quorum logical drive.
5. Map logical drives: Refer to “Change mapping” on page 229. Map the Quorum drive and the previous generated RAID 5 logical drive to the host group.

These steps show how to perform these steps in the Storage Manager GUI. Example 25-2 shows a CLI script to do the same steps.

Example 25-2 Change storage configuration

```
show "Create a host group";
create hostGroup userLabel="Thailand_Cluster";

show "Create a host named Yom";
create host userLabel="Yom";

show "Add host ports to host";
create hostPort identifier="10000000c93542c2" userLabel="Yom_Port_0" host="Yom"
hostType="Windows 2000/Server 2003 Clustered";
create hostPort identifier="10000000c93542c3" userLabel="Yom_Port_1" host="Yom"
hostType="Windows 2000/Server 2003 Clustered";
show allHostPorts;

show "Change PaSak host ports to clustered host type";
set hostPort [PaSak_Port0] hostType="Windows 2000/Server 2003 Clustered";
set hostPort [PaSak_Port1] hostType="Windows 2000/Server 2003 Clustered";

show "Add hosts to host group";
set Host [Yom] hostGroup="Thailand_Cluster";
set Host [PaSak] hostGroup="Thailand_Cluster";


show "Creating Logical Drive Quorum on existing Array 1.";
create LogicalDrive array=1 userLabel="Quorum" owner=A segmentSize=128 capacity=1
GB;
show "Setting additional attributes for Logical Drive Quorum.";
// Configuration settings that can not be set during Logical Drive creation.
set logicaldrive["Quorum"] cacheFlushModifier=10;
set logicaldrive["Quorum"] cacheWithoutBatteryEnabled=false;
set logicaldrive["Quorum"] mirrorEnabled=true;
set logicaldrive["Quorum"] readCacheEnabled=true;
set logicaldrive["Quorum"] writeCacheEnabled=true;
set logicaldrive["Quorum"] mediaScanEnabled=true;
set logicaldrive["Quorum"] redundancyCheckEnabled=false;
set logicaldrive["Quorum"] readAheadMultiplier=1;
set logicaldrive["Quorum"] modificationPriority=high;

show "Map logical drives to the host group";
set logicalDrive ["Quorum"] logicalUnitNumber=0 hostGroup="Thailand_Cluster";
remove logicalDrives [PaSak_Raid5] lunMapping;
set logicalDrive ["PaSak_Raid5"] logicalUnitNumber=2 hostGroup="Thailand_Cluster";
remove logicalDrives [PaSak_Raid1] lunMapping;
set logicalDrive ["PaSak_Raid1"] logicalUnitNumber=1 hostGroup="Thailand_Cluster";

show "Rename logical drive that it fits to the cluster naming";
set logicalDrive ["PaSak_Raid1"] userLabel="Thailand_Cl_Raid1";
set logicalDrive ["PaSak_Raid5"] userLabel="Thailand_ClRaid5";
```

25.3.3 Enabling the Cluster Service

Open the disk management of your first cluster node and create a NTFS file system on the newly generated Quorum Logical drive. Assign drive letter Q to this drive. The required steps are shown in 25.2.3, “Windows Disk Manager” on page 620.

Initiating the Cluster Wizard

Follow these steps:

1. To start the Cluster Wizard, click **Start** → **Run** and type cluadmin.
2. From the Action menu, select **Create new cluster** and click **OK** (Figure 25-25).

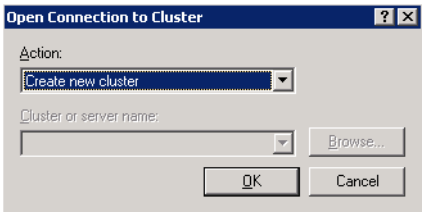


Figure 25-25 Create new cluster

3. The New Server Cluster Wizard starts (Figure 25-26). Click **Next**.



Figure 25-26 New Server Cluster Wizard

4. Select the domain in which the cluster will be created (if it is not detected automatically) and assign a unique name to the cluster (Figure 25-27). Click **Next**.

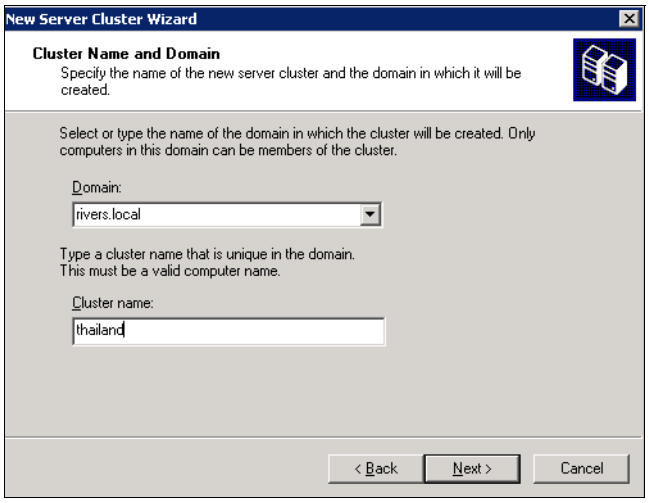


Figure 25-27 Cluster Name and Domain

5. In Figure 25-28, select the computer name that will be defined as the first node in the cluster or click **Browse** to perform a search. Click **Advanced**.

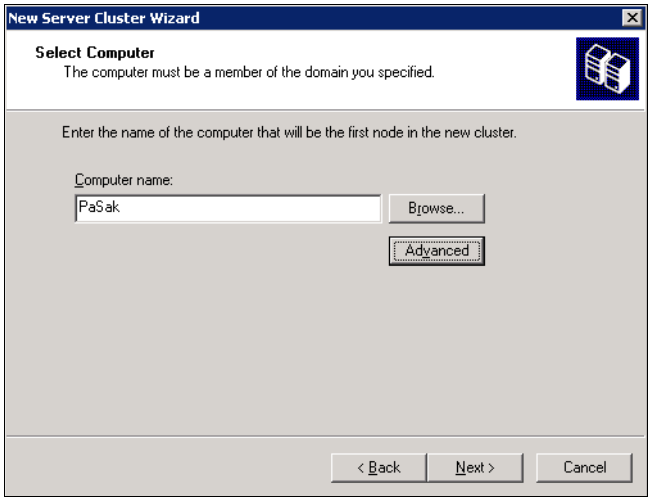


Figure 25-28 Select Computer

6. You can choose from a Typical Configuration or an Advanced (minimum) Configuration (Figure 25-29 on page 631). The Advanced Configuration should be used for complex shared-storage configurations. We select **Typical (Full) Configuration** and click **Next**.

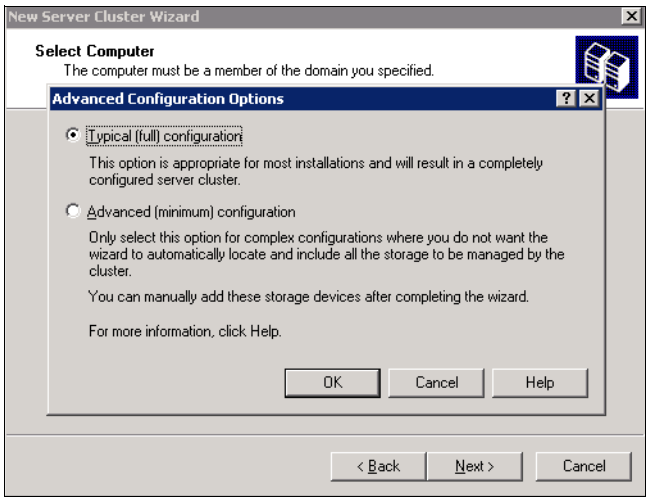


Figure 25-29 Advanced Configuration Options

7. The Cluster Wizard will now perform a diagnostic check on the system configuration to ensure that the hardware components meet the Windows 2003 Server cluster requirements (Figure 25-30). If any problems arise, they will be highlighted.

   If required, make any necessary changes and click **Re-analyze** to re-run the diagnostics. Click **Viewlog** to see, in a text file, the steps performed by the cluster wizard for the diagnostic check. Alternatively, click **Details** to comparable data from a task viewpoint; this allows you to step through the tasks individually. Click **Next** to continue.

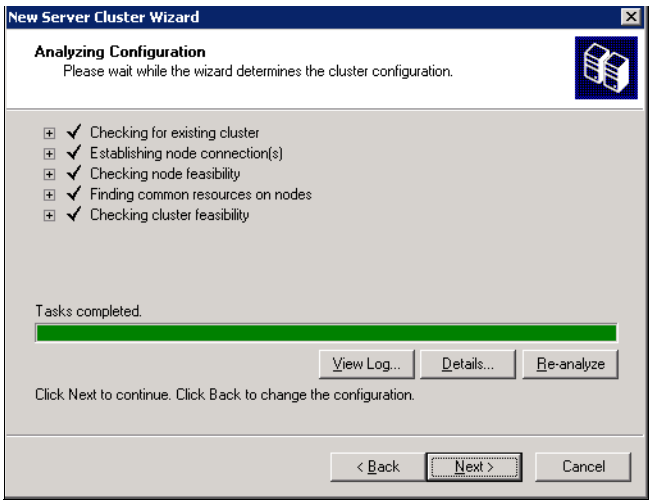


Figure 25-30 Analyzing Configuration

8. Enter the IP address you will use for the cluster (from Table 25-1 on page 627) and click **Next** (Figure 25-31).

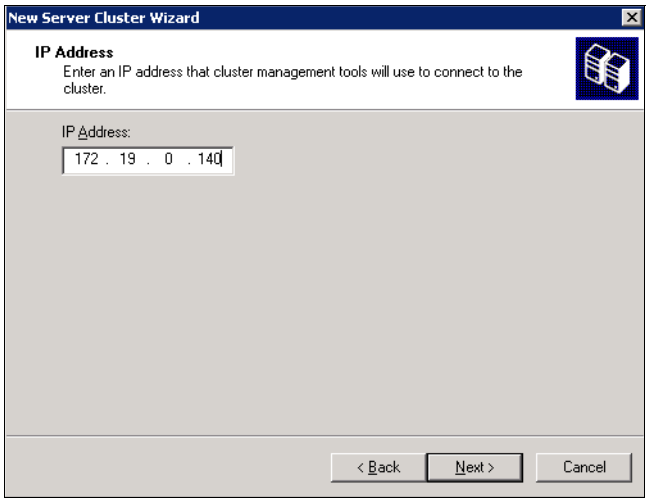


Figure 25-31 Enter cluster IP address

9. In Figure 25-32, enter the user name and password of the domain account that was created in step 5 on page 627 to be used as a cluster service account. Click **Next**.

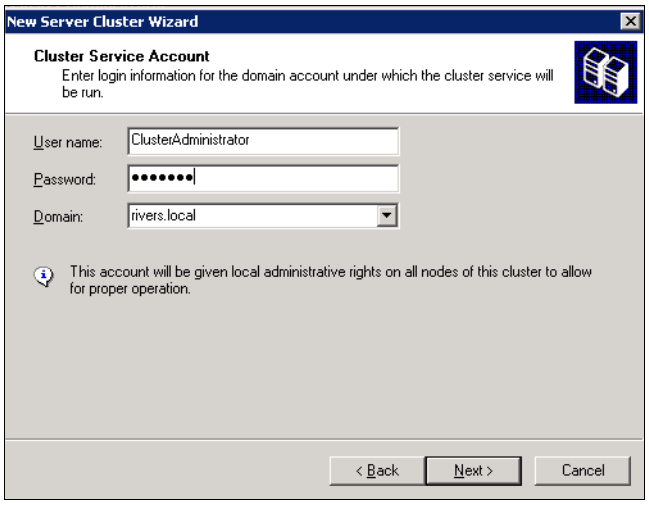


Figure 25-32 Cluster Service Account

10. Read the summary of the Proposed Cluster Configuration shown in Figure 25-33 on page 633. Click **Next** to start the cluster configuration.

> **Note:** Click **Quorum** (Figure 25-34) to see which disk has been selected for the quorum disk. It must point to the previously generated Quorum disk on your DS3000 storage. A local quorum and a Majority Node Set are not suitable for a fibre cluster.

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

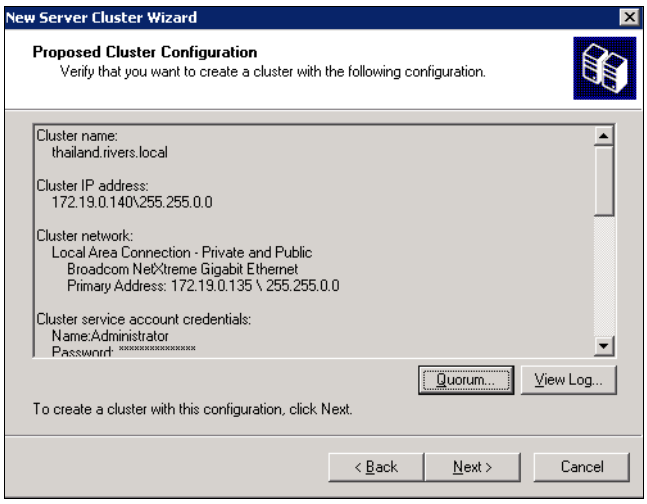


Figure 25-33 Proposed Cluster Configuration

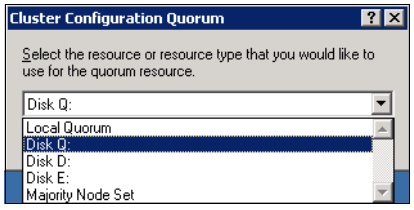


Figure 25-34 Cluster Configuration Quorum

11. The cluster wizard now starts to install the cluster software and services (Figure 25-35).

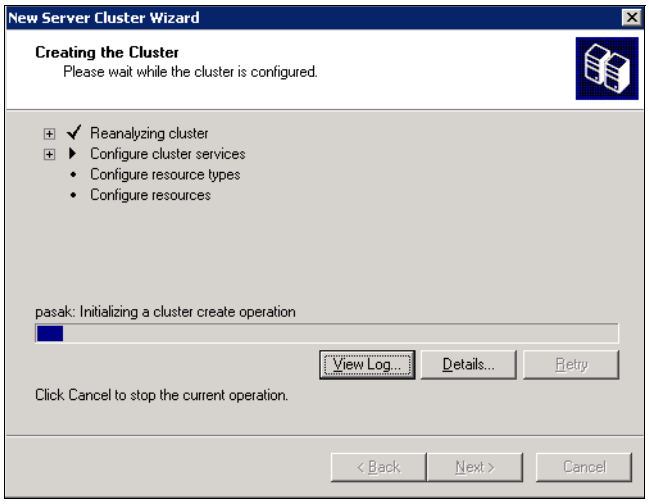


Figure 25-35 Creating the cluster

12. Once the tasks are completed, click **Next** and **Finish** (Figure 25-36) to complete the cluster installation on the first cluster node.



Figure 25-36 Completing the New Server Cluster Wizard

13. The single node cluster is now visible through the Cluster Administrator (Figure 25-37). Double check the groups and resources to ensure that the system configuration matches the cluster configuration.

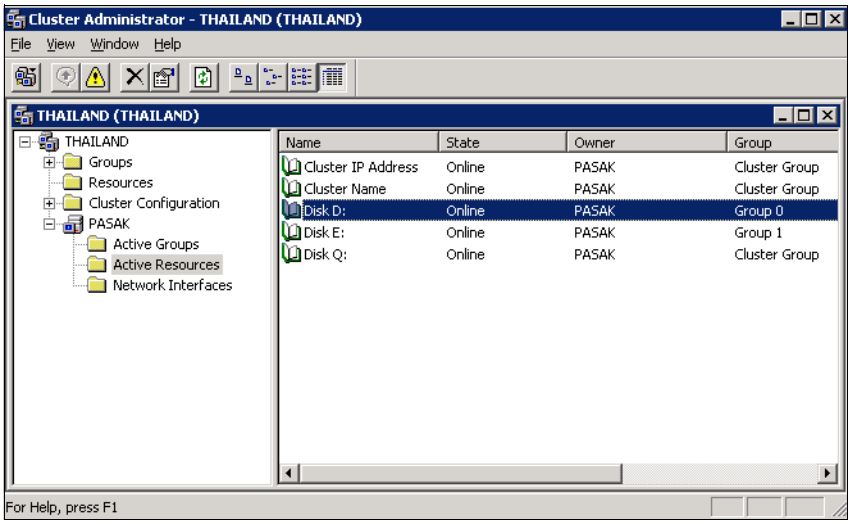


Figure 25-37 Cluster Administrator

25.3.4 Joining the second node to the cluster

Our cluster currently contains only the node PaSak. We need to add node Yom to the cluster. Windows Server 2003 allows the cluster service to be installed from either node.

Do these steps:

1. Start the Cluster Administrator (either select **Start** → **Programs** → **Administrative Tools** → **Cluster Administrator** or **Run** → **Cluadmin**).
2. If this is not the first time you have run Cluster Administrator on this host, select f **File** → **Open Connection** in the Cluster Administrator window.

3. Select **Action** → **Add nodes to cluster**, as shown in Figure 25-38.

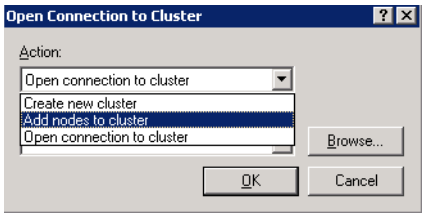


Figure 25-38 Add node to cluster

4. Click **Browse** and you will see the clusters that are available in the domain. Select your cluster from the list (Thailand, in our configuration) and click **OK** (Figure 25-39).

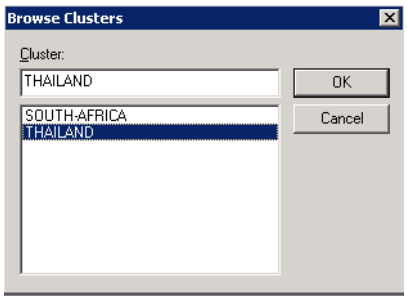


Figure 25-39 Browse Clusters

5. Click **OK** again.
6. The **Add Nodes Wizard** starts (Figure 25-40). Click **Next**.



Figure 25-40 Add Nodes Wizard

7. In the Computer Name field, enter the computer name of your second node (Figure 25-41), click **Add**, and then **Next**.

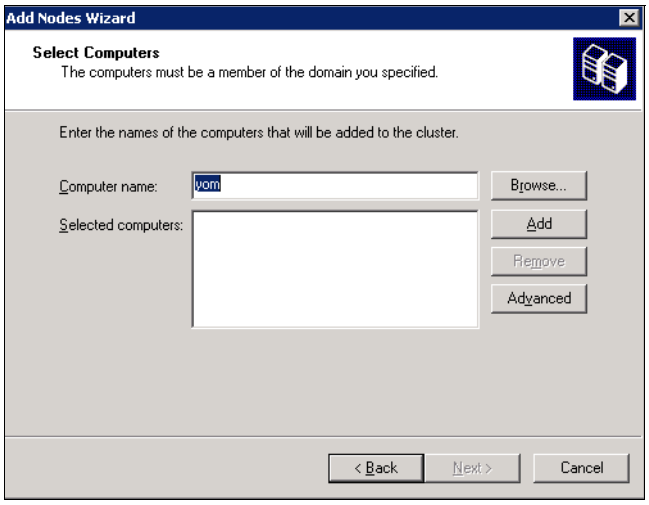


Figure 25-41 Select Computers

8. The wizard now performs the cluster diagnostic test for this node (Figure 25-42). Click **Next** when the test has completed successfully.

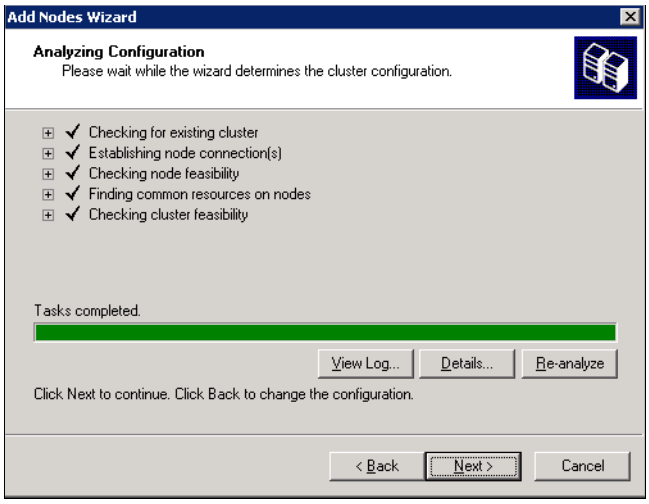


Figure 25-42 Analyzing Configuration

9. Enter the cluster service account (Figure 25-32 on page 632) and click **Next**.
10. Review the proposed cluster configuration and click **Next** to complete (Figure 25-43 on page 637). The wizard will now add the node to the cluster.

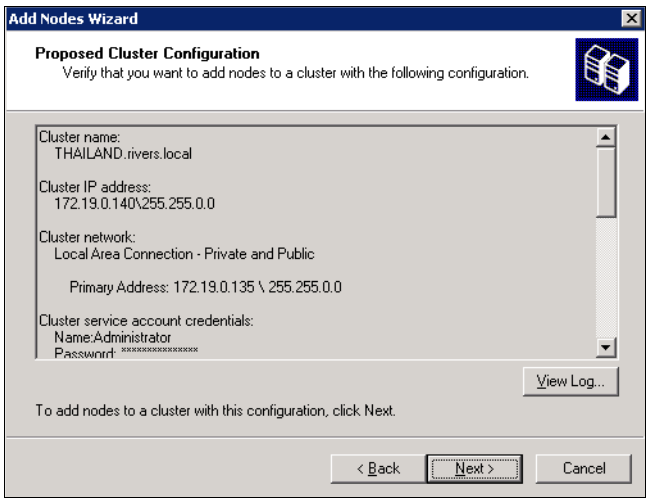


Figure 25-43 Proposed Cluster Configuration

11. Once the wizard has finished, click **Next** (Figure 25-44) and **Finish** (Figure 25-45 on page 638) to complete the addition of the second node.

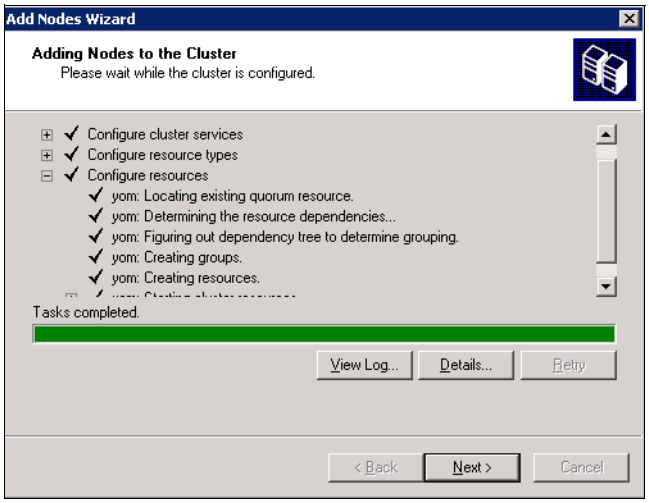


Figure 25-44 Adding Nodes to Cluster Complete

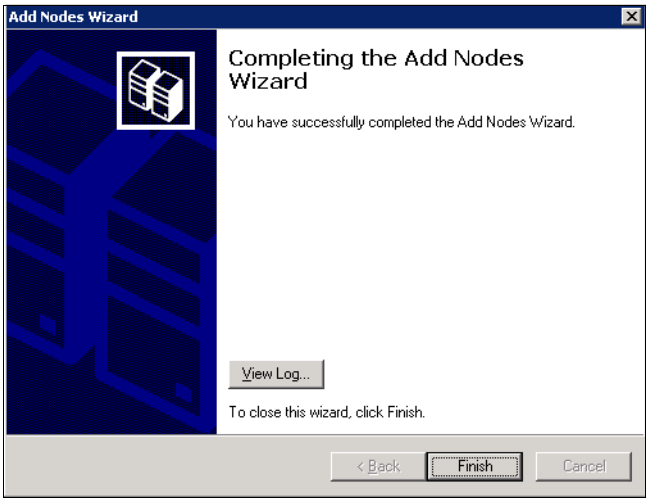


Figure 25-45 Completing the Add Nodes Wizard

Now the basic configuration of your cluster is done and you see both nodes in the cluster administrator (Figure 25-46). Add more nodes to the cluster, or start testing for failover and install your applications.

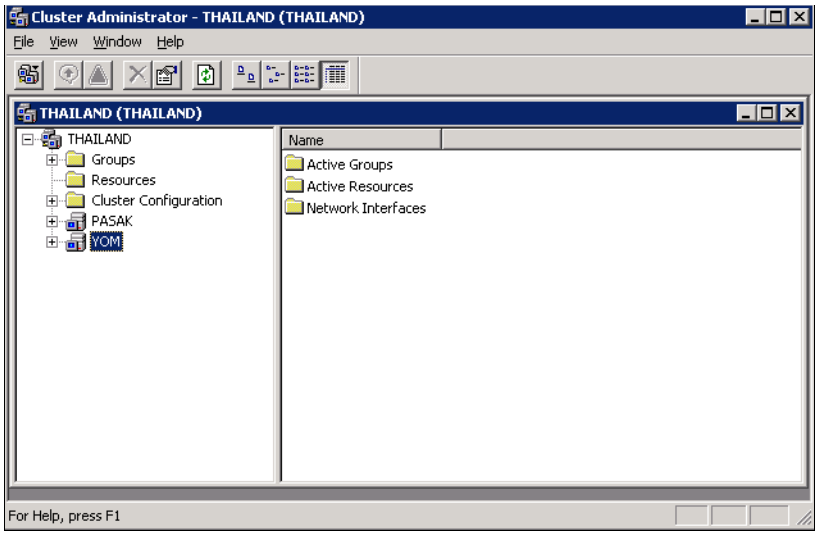


Figure 25-46 Cluster Administrator

26

FC configuration 3 - Linux SAN boot from a DS3400 with an IBM System x server

In this chapter, we discuss a sample configuration that shows how to configure Linux SAN boot from the DS3400 subsystem. A host server, an IBM System x server, with two FC ports on a QLogic-based FC HBA will be connected to a dual-controller DS3400 across an FC SAN. We will define a logical drive on the DS3400 and map our host server to it. Our goal is to configure the server so that it will boot into Red Hat Enterprise Linux 5 from the logical drive on the DS3400.

The installation steps are as follows:

1. Configure the first FC HBA port. Note its WWPN and enable its BIOS.
2. Create a zone on the SAN switch for the first HBA port and controller A.
3. Create a logical drive for the operating system on the DS3400 server. This logical drive will be mapped to our host server as LUN 0. At this stage, you should define the host server with one host port only (first HBA port WWPN).
4. Use the FC HBA BIOS utility to define the logical drive as boot device.
5. Install Red Hat Enterprise Linux 5 on the logical drive. When Linux installation completes, update the FC HBA driver and install Linux RDAC (MPP).
6. Create a zone on the SAN switch for the second FC HBA port and controller B.
7. Add a second host port (second FC HBA port WWPN) to the host server definition in the DS3000 Storage Manager.
8. Verify that both controllers are accessible from the host server and test I/O path redundancy.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

26.1 What is required

For our setup, we used the following hardware and software components:

- IBM System x server.
- Dual Port 4 Gbps FC HBA.
- SAN switches (we will use one 2005-B32).
- A DS3400 server with enough free disk space for the logical drive.
- The latest FC HBA driver for Linux kernel 2.6. The driver is available for download on the IBM Systems support Web site.
- A management workstation running the latest version of DS3000 Storage Manager 3.
- The latest version IBM RDAC drivers for Linux.
- Red Hat Enterprise Linux 5 installation media.

Figure 26-1 shows our sample hardware setup.

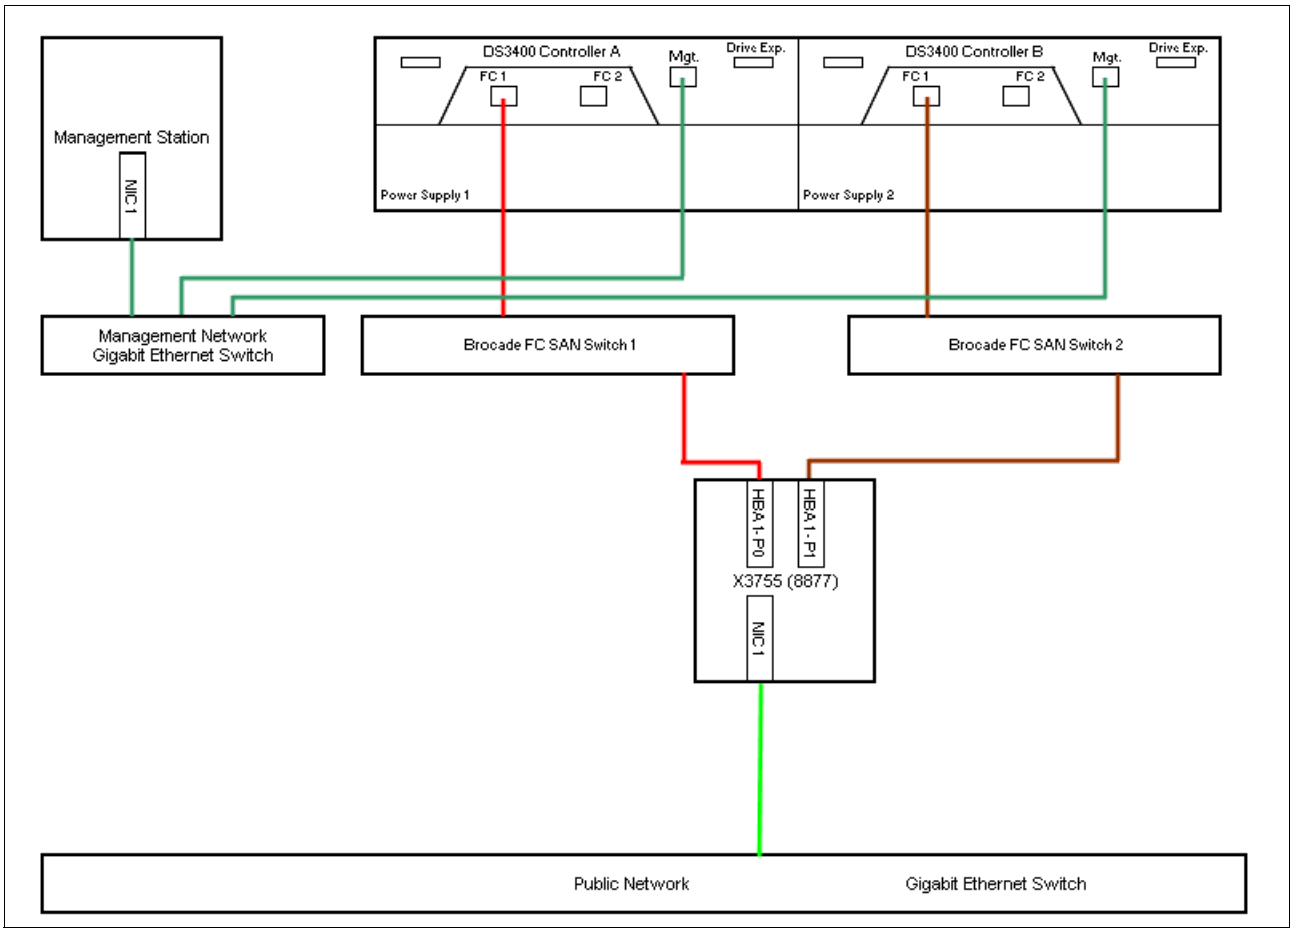


Figure 26-1 Sample hardware setup

The figure shows two SAN switches, however, for simplicity, we use just one in our real configuration. It has to be zoned appropriately so that two separate I/O paths between the server and the DS3400 server are available. In real life scenarios, we recommend using two separate switches for full redundancy.

For the main part of the installation process, we will only use single path (first FC HBA port connected to controller A). Only after the Linux RDAC is installed onto the logical drive will we define and enable the second /IO path (second HBA port connected to controller B). This is to avoid problems and conflicts that could occur if the logical drive was seen across both I/O paths before having the Linux RDAC installed and operational.

Note: You have to disable internal SCSI or SAS on-board adapters in the server BIOS if you want to set FC HBA adapters to be able boot from SAN.

26.2 First FC HBA port configuration

In this section, we use the QLogic HBA BIOS utility (named Fast!UTIL) to note the WWPNs of both HBA ports and to enable the BIOS on the first HBA port.

Note: In our configuration, we are using QLogic HBA; however, the Emulex PCI Express 4 Gbps FC HBA is also supported, as per the IBM Support Web site. The BIOS utility procedures are similar to those described in Chapter 24, “FC configuration 1 - Emulex HBA boot blade server from SAN” on page 585.

Do these steps:

1. To enter Fast!UTIL, power the host server on and press Ctrl+Q when prompted. The first step is HBA port selection. as shown in Figure 26-2.

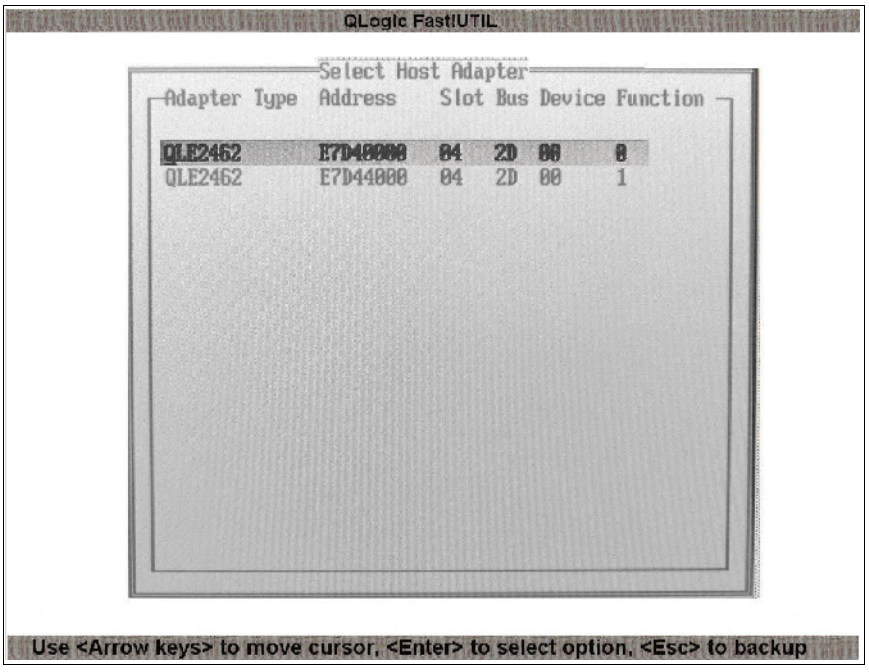


Figure 26-2 Select the FC HBA port

2. Select the first HBA port and press Enter. The Fast!UTIL Options menu displays (Figure 26-3). Select **Configuration Settings** and press Enter.

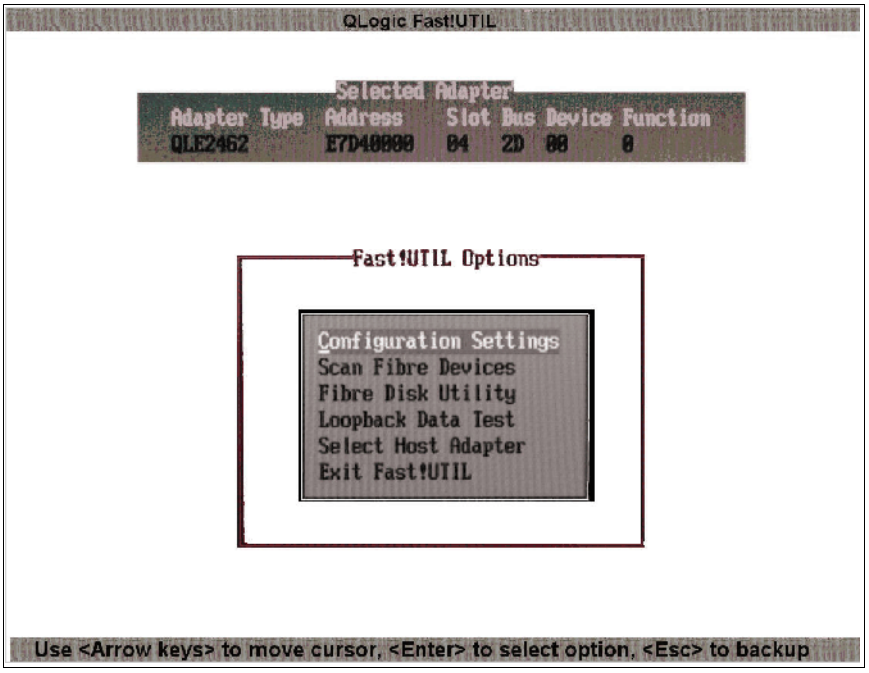


Figure 26-3 Fast!UTIL Options

3. In the Configuration Settings menu, shown in Figure 26-4, select **Adapter Settings**.

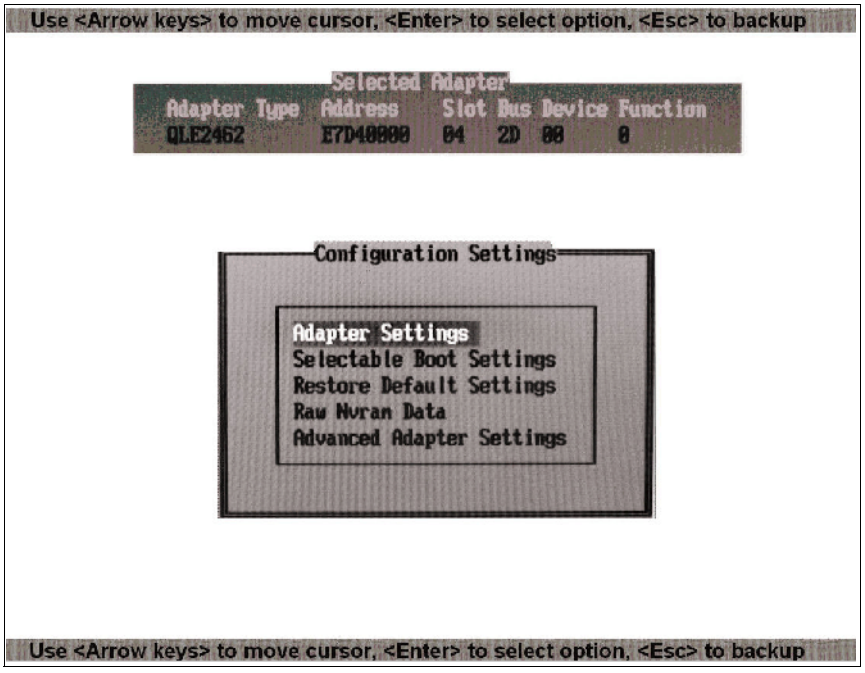


Figure 26-4 Configuration Settings menu

4. The Adapter Port Name field shows the HBA port WWN. In our case, the WWN is 21:00:00:E0:8B:85:41:60 (see Figure 26-5). Record this value.

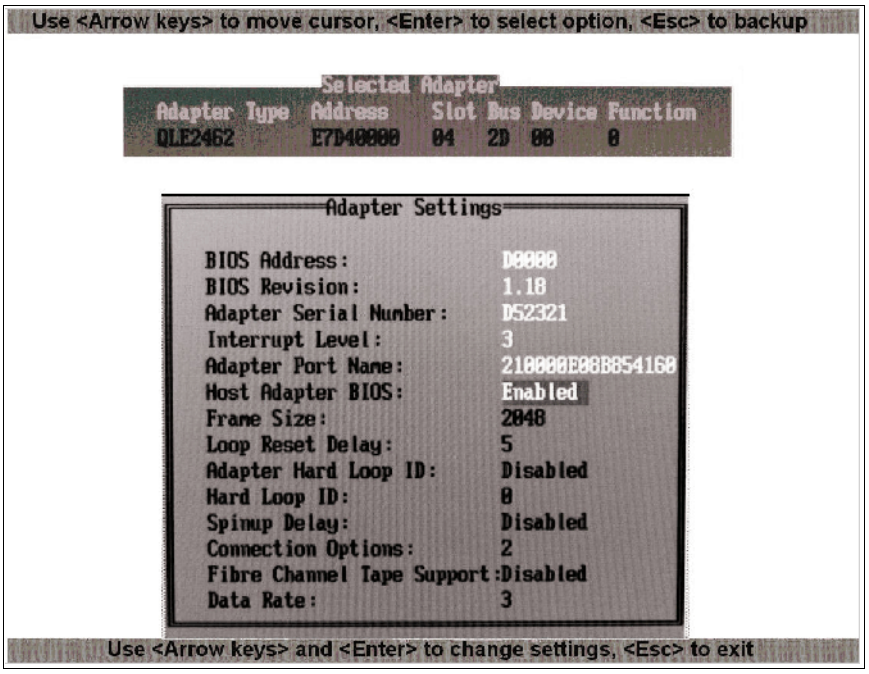


Figure 26-5 Adapter Settings

5. Here we also enable the BIOS for this HBA port. Remember, the BIOS must be enabled if we want to boot through this HBA port.

   Change the Host Adapter BIOS setting from **Disabled** to **Enabled**, then press Esc. Select **Save changes** to confirm.

6. Later in the procedure, we also enable the second HBA port for booting, but it will be after successful installation of operating system on the first port. We also need the WWPN of the second port. Display and record the WWPN of the second HBA port now.

7. Switch to the second HBA port. Return to the Fast!UTIL Options menu (Figure 26-3 on page 642) and choose **Select Host Adapter**. You will see two QLE2462 adapter ports displayed, as shown in Figure 26-2 on page 641. Select the second line and press Enter.

8. Select **Configuration Settings** → **Adapter Settings** and look for the WWPN of the second HBA port. The WWN is 21:01:00:E0:8B:A5:41:60. Note that BIOS will be enabled on the second HBA later, not in this step.

In this section, we enabled the BIOS on the first HBA port and noted down the WWPNs of both HBA ports:

- HBA port 1: 21:00:00:E0:8B:85:41:60
- HBA port 2: 21:01:00:E0:8B:A5:41:60

26.3 SAN switch zoning for the first I/O path

For our example, we use just one IBM 2005-B32 SAN switch. We always recommend using two redundant FC switches for the production environment. Our switch now needs to be zoned.

Do these steps:

1. First we create the following aliases:
   - x3755_QL_P0 for the first HBA port WWPN (21:00:00:E0:8B:85:41:60).
   - x3755_QL_P1 for the second HBA port WWPN (21:01:00:E0:8B:A5:41:60)
   - DS3400_CA_P0 for the DS3400 controller A WWN (20:26:00:A0:B8:2F:C1:D5)
   - DS3400_CB_P0 for the DS3400 controller B WWN (20:27:00:A0:B8:2F:C1:D5)

   The WWNs of host ports on DS3400 controllers A and B were determined from the DS3000 storage subsystem profile (see 10.2.1, “Storage Subsystem Profile” on page 212).
2. We create two zones that will be based on the aliases listed above.

   Zone1: x3755_P0_CtrlA, with members
   - First HBA port (alias x3755_QL_P0)
   - DS3400 controller A (alias DS3400_CA_P0) WWN

   This zone defines the I/O path between the first HBA port and DS3400 controller A.

   Zone2: x3755_P1_CtrlB, with members
   - Second HBA port (alias x3755_QL_P1)
   - DS3400 controller B (alias DS3400_CB_P0)

   This zone links second HBA port with controller B and thus represents the alternative I/O path.
3. We create a configuration based on the zones. Actually, we will create two configurations:
   - x3755_CtrlA_only

     This configuration contains zone x3755_P0_CtrlA only. This means only the primary I/O path between the first HBA and controller A is enabled. The initial SAN boot installation steps require that the host server and the DS3400 only use one I/O path; this will therefore be the configuration used in this part of the process.
   - x3755_bothCtlr

     We use this configuration to enable both I/O paths. Therefore, both zones (x3755_P0_CtrlA and x3755_P1_CtrlB) will be members. We enable this configuration after the Linux RDAC is installed.

Figure 26-6 shows our zoning. As stated above, the effective configuration must be x3755_CtrlA_only at this time.

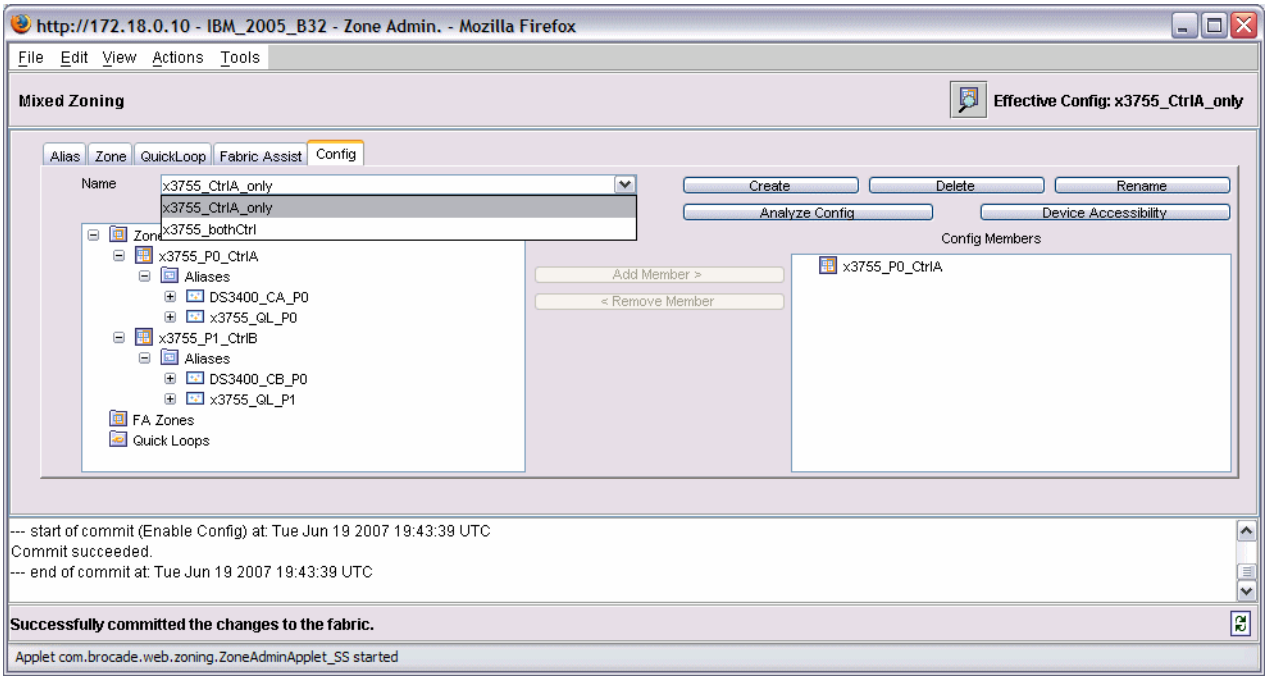


Figure 26-6 Zoning configuration

26.4 Creating a logical drive

We use the DS3000 Storage Manager running on a management workstation to create a logical drive that will be used to install and later boot Linux. But before creating the logical drive, we have to define our host server and its host port (first HBA port).

26.4.1 Host server definition

Do these steps:

1. In the DS3000 Storage Manager subsystem management window, select the **Configure** tab, and click **Configure Host Access (Manual)**, as shown in Figure 26-7.

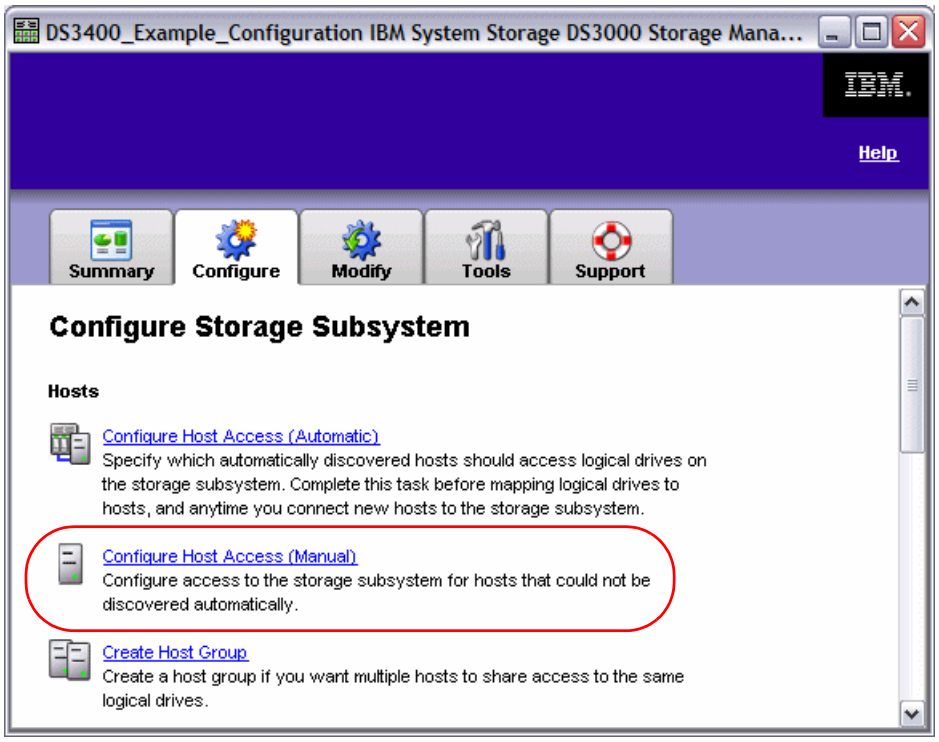


Figure 26-7 Configure Host Access (Manual)

2. Specify the host name and type. In our sample case, the host server name will be x3755_SAN_Boot_Linux and the type will be Linux (see Figure 26-8 on page 647).

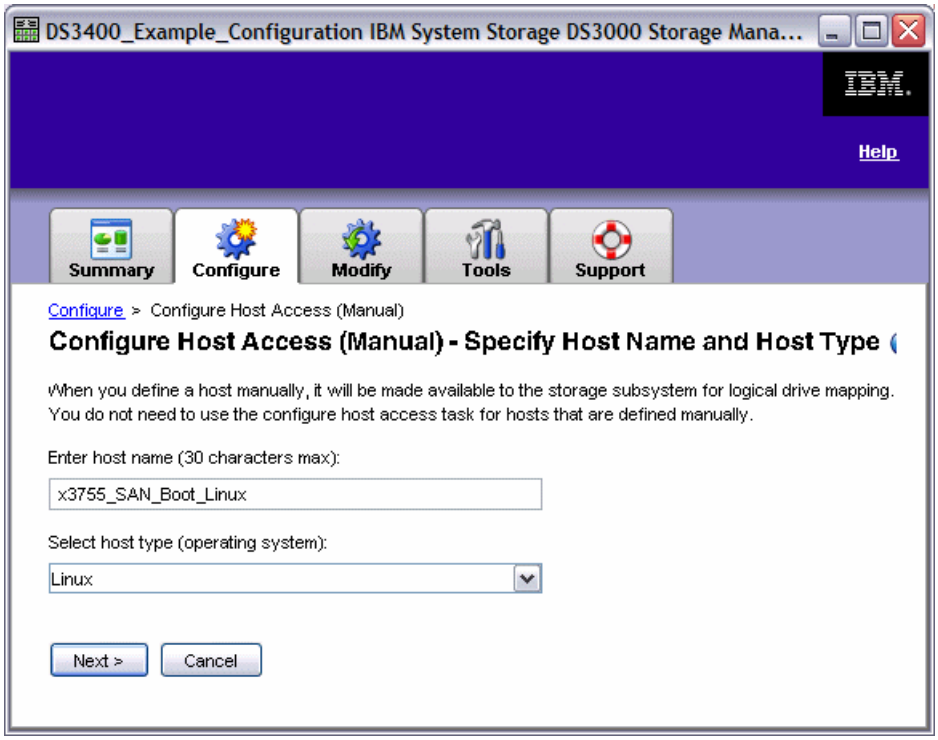


Figure 26-8 Specify Host Name and Host Type

3. Specify the HBA host ports. Remember, we have to specify the first HBA port only. Its WWN is 21:00:00:E0:8B:85:41.60. Select this WWN and click **Add**, as shown in Figure 26-9.

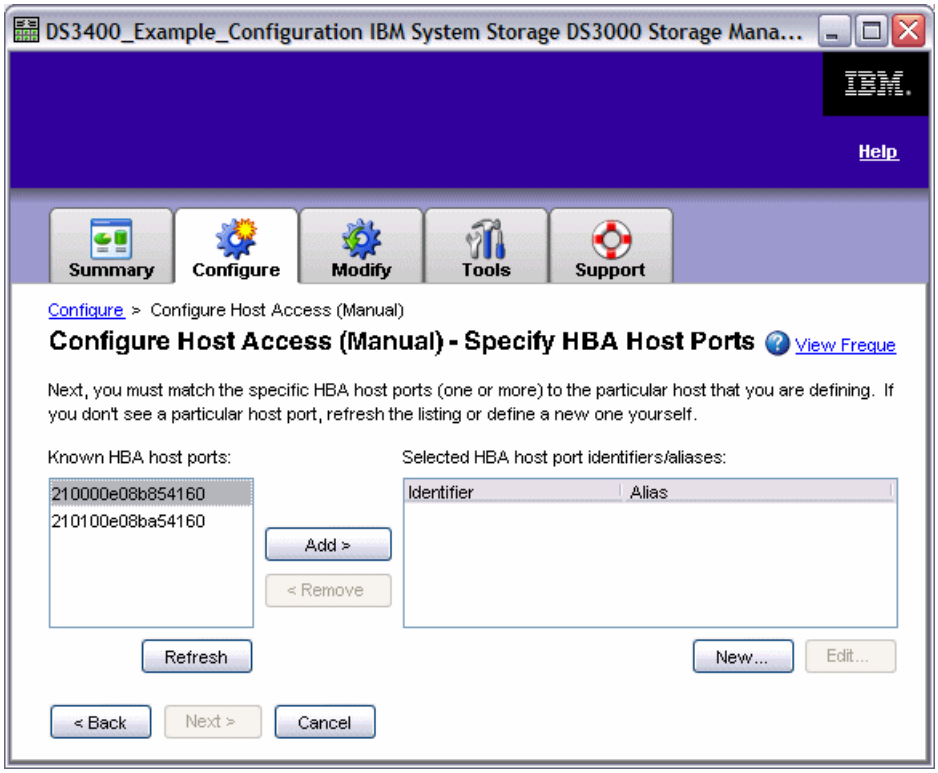


Figure 26-9 Specify HBA Host Port

4. Next you can specify the host group. Since our host server is not clustered, we will not create any host group. Select **No: This host will NOT share access to the same logical drives with other hosts** and click **Next**, as shown in Figure 26-10.

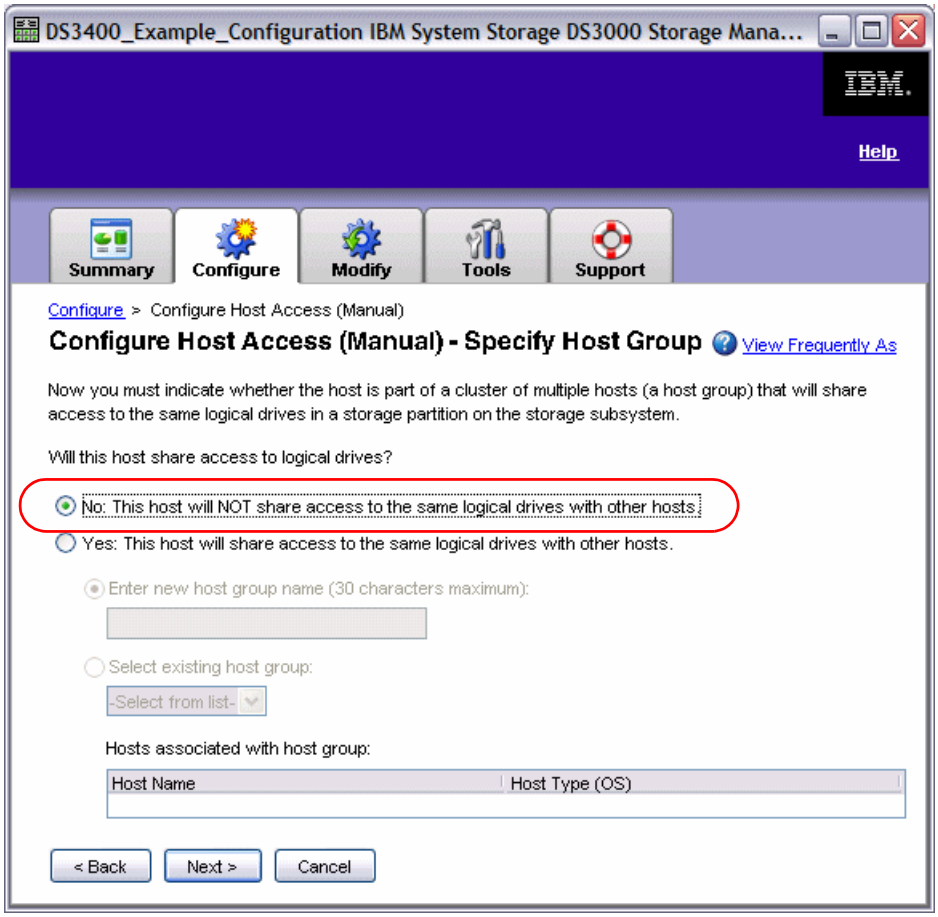


Figure 26-10 Specify Host Group

5. The confirmation window displays, as shown in Figure 26-11 on page 649. Review the host definition and click **Finish**.

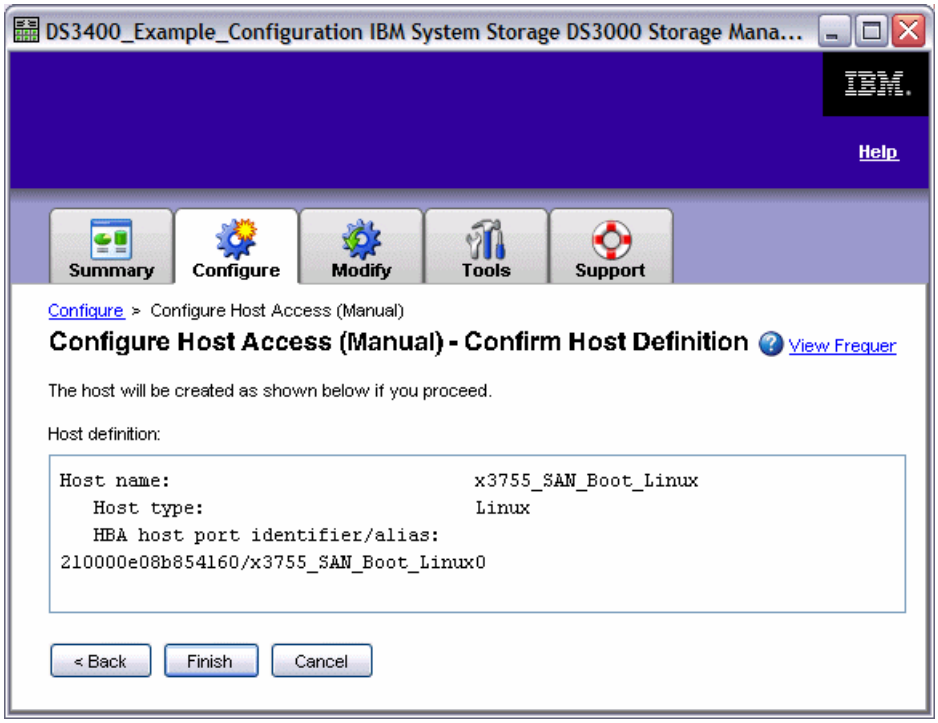


Figure 26-11 Confirm Host Definition

6. On the completion window (see Figure 26-12), you can define another host. As we will use just this host server, we do not want to define another one at this time. Click **No**.

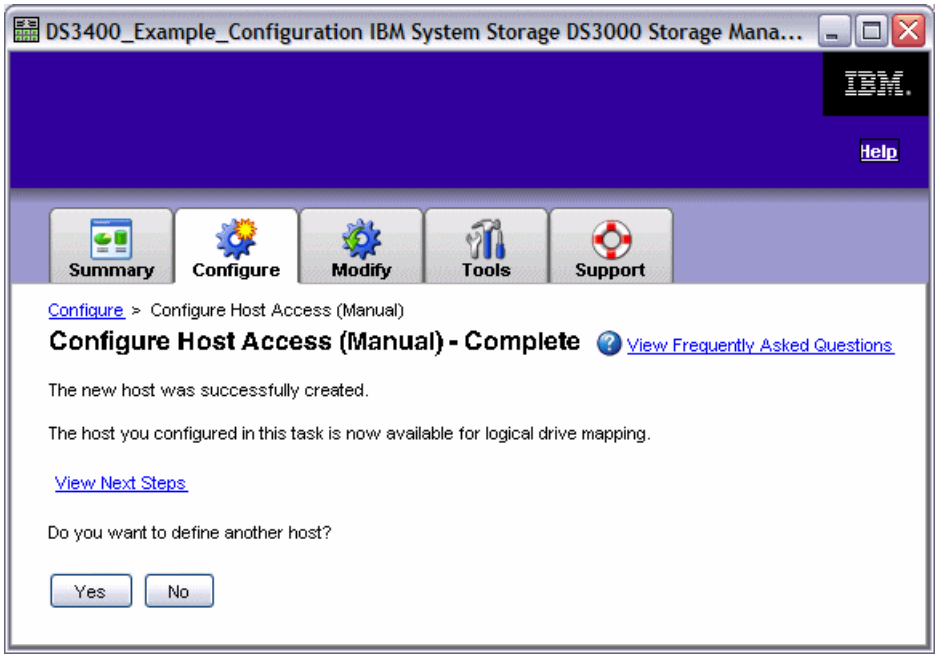


Figure 26-12 Host definition is complete

26.4.2 Creating a logical drive

Follow these steps:

1. To define the logical drive, select the **Configure** tab and click **Create Logical Drives**.
2. Select the capacity type, as shown in Figure 26-13.

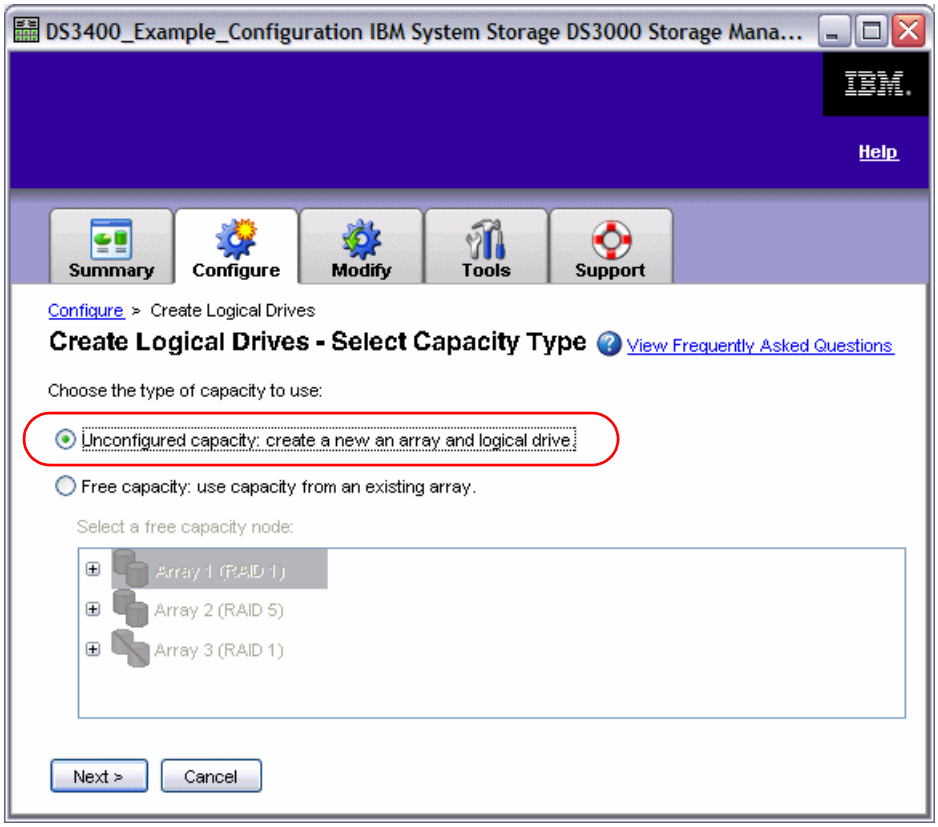


Figure 26-13 Select Capacity Type

3. We will use two unassigned disk drives to create a new array. Therefore, we select **Unconfigured capacity: create a new array and logical drive**.
4. In Figure 26-14, there are two drive selection choices. We select **Automatic**.

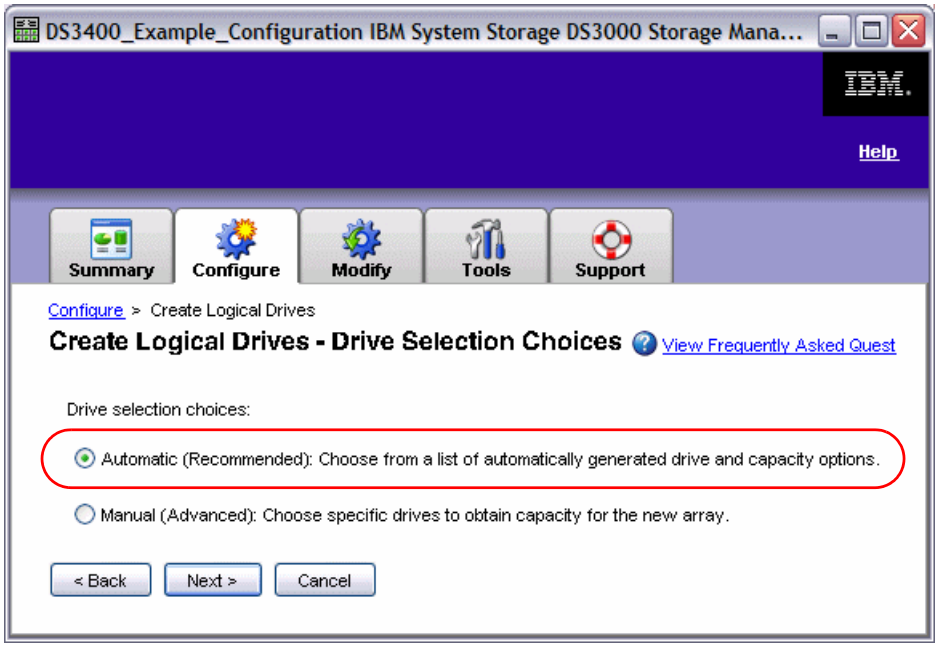


Figure 26-14 Drive Selection Choices

5. In the Specify Capacity window, we specify the RAID level and then select one of the presented capacity options. We have two unassigned 68 GB disk drives installed and we want to create a RAID 1 array. So we need to select RAID level 1; at this point, the window displays an entry showing two drives (Figure 26-15). Highlight this line and click **Next**.

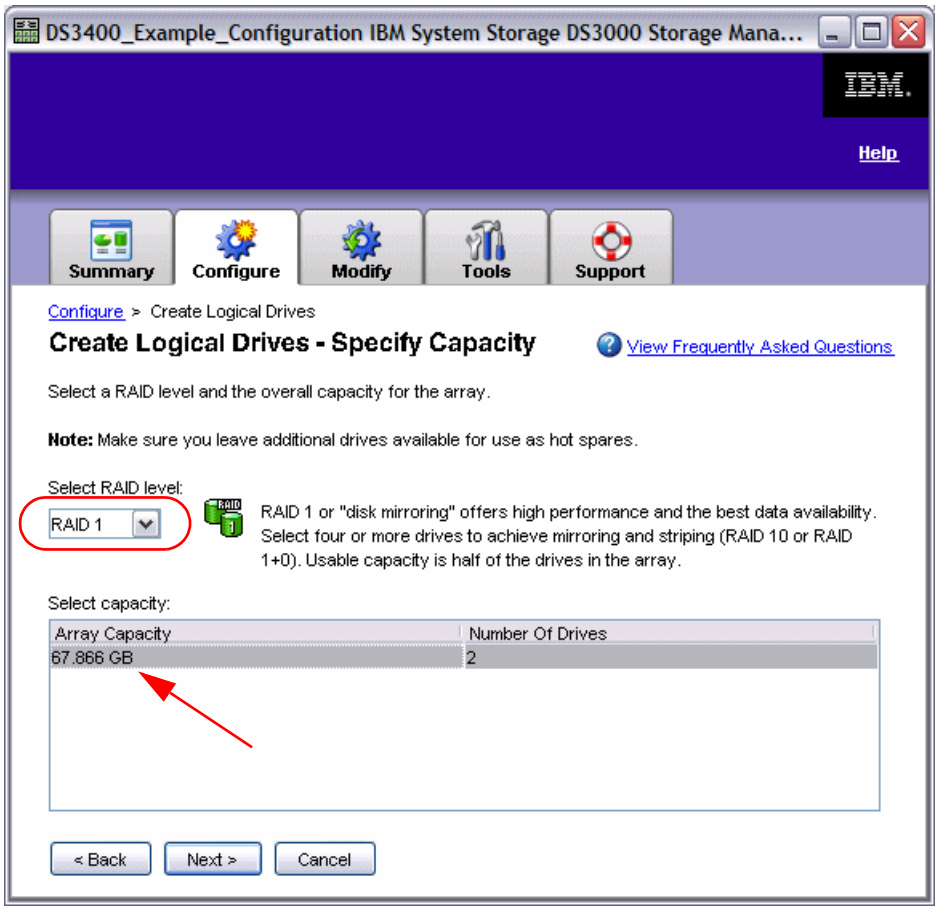


Figure 26-15 Specify Capacity

6. Now we need to specify the logical drive parameters: capacity, name, and I/O characteristics. We set 32 GB for capacity, the name will be SAN_boot_Linux and the I/O characteristics will be left at File system (typical) (see Figure 26-16).

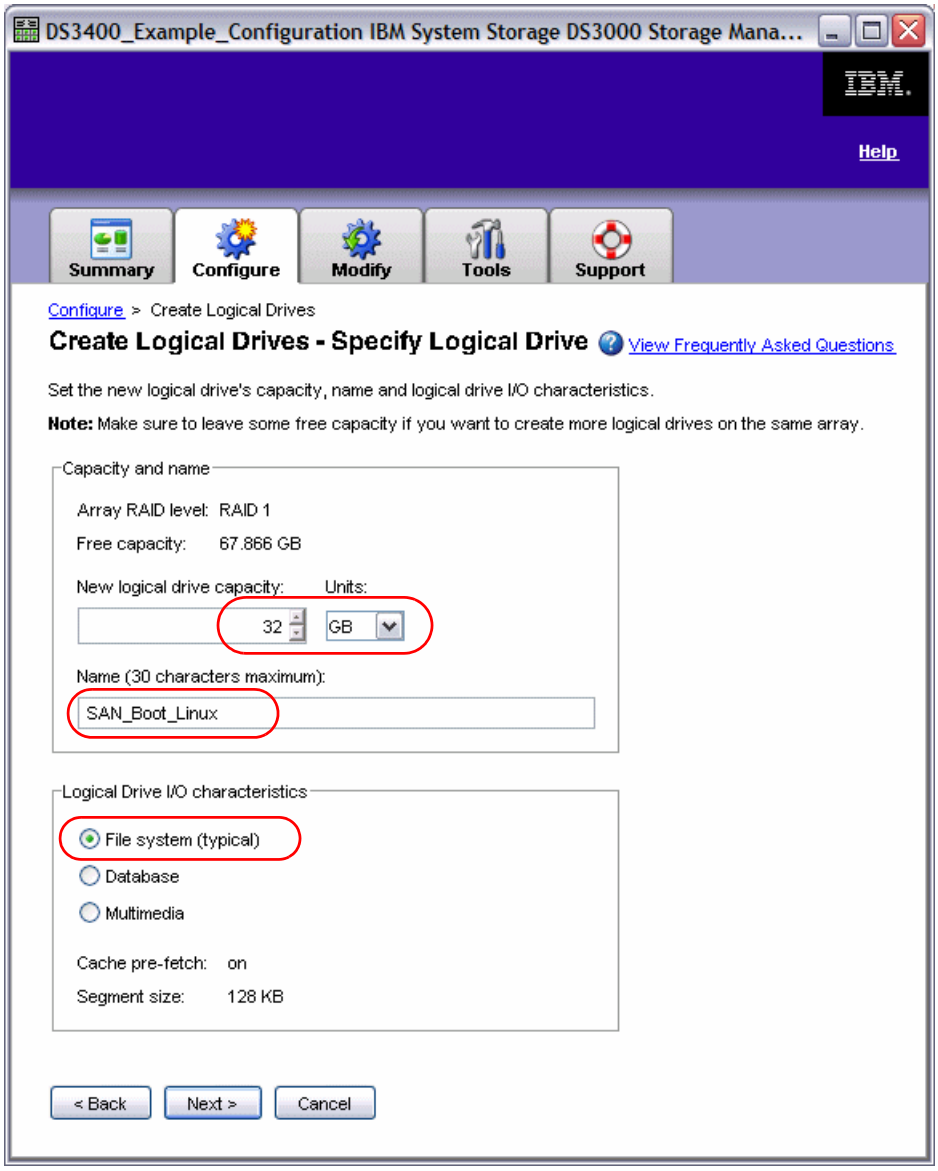


Figure 26-16 Specify Logical Drive

7. Now we map the host to the logical drive. Select the newly defined host **x3755_SAN_Boot_Linux** and assign the logical drive as LUN 0, as shown in Figure 26-17.

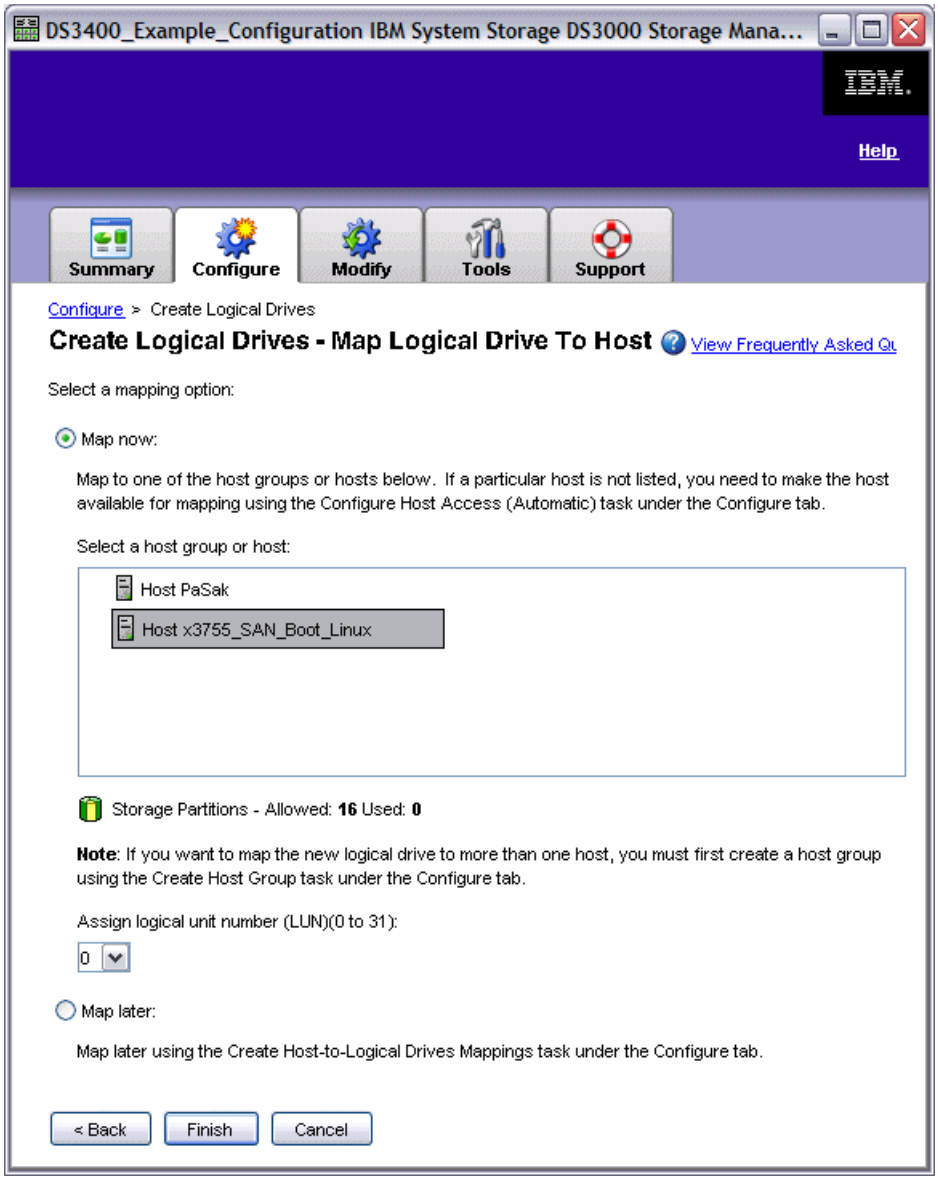


Figure 26-17 Map Logical Drive To Host

8. The logical drive is created and assigned to our host. The completion window displays (see Figure 26-18) and asks whether we want to create another logical drive. Click **No** to finish this section.

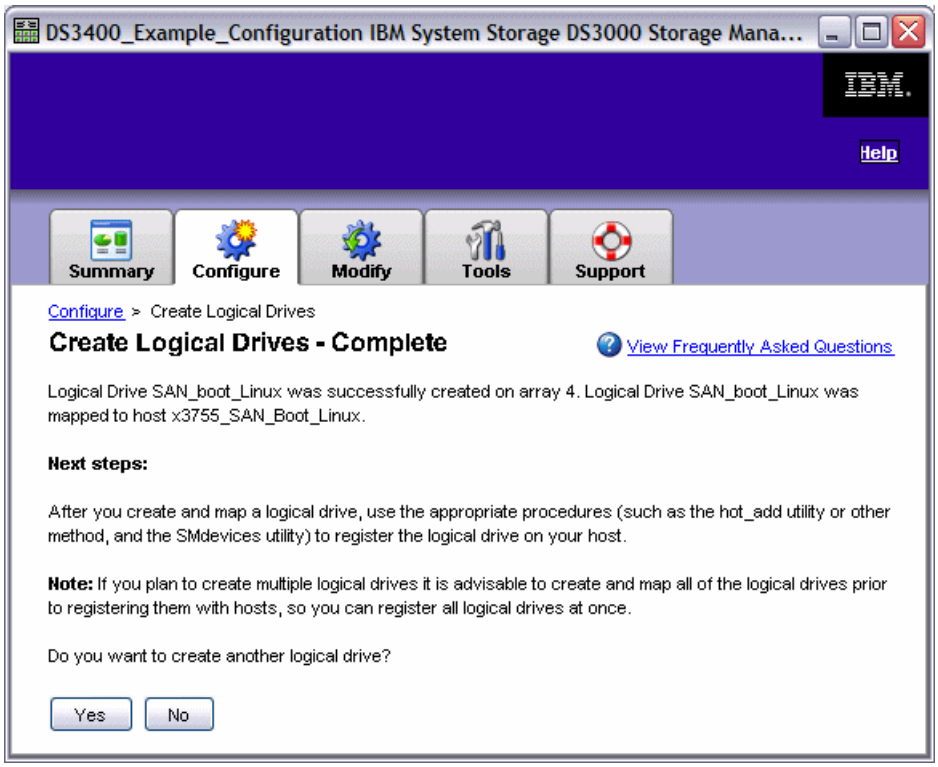


Figure 26-18 Logical drive creation complete

26.4.3 Logical drive ownership

The host server can only talk to controller A at this point. This means controller A must be the preferred and current owner of the logical drive named SAN_boot_Linux. To verify this, use the Change Logical Drive Ownership/Preferred Path task in the Modify tab. As shown in Figure 26-19, our logical drive is currently owned by controller B, which is also the preferred controller. Click **Change** to reset it to controller A.

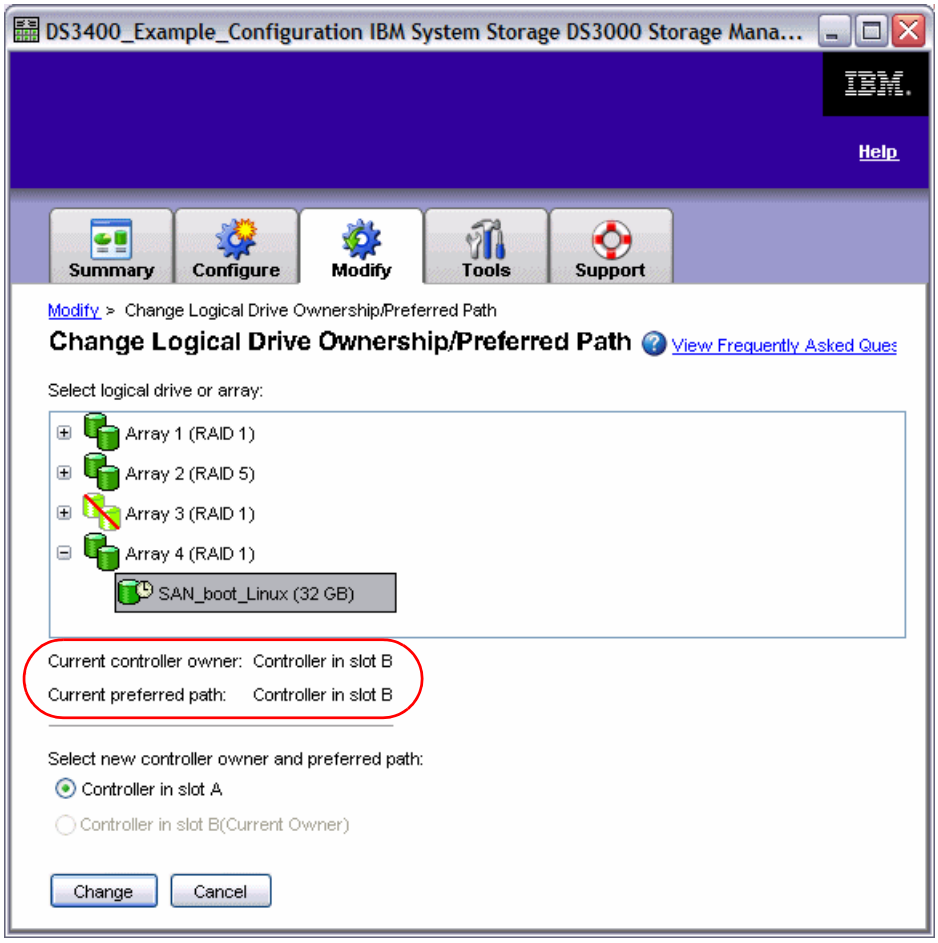


Figure 26-19 Change Logical Drive Ownership/Preferred Path

You are asked to confirm the ownership and preferred path change (see Figure 26-20). Click **Yes**.

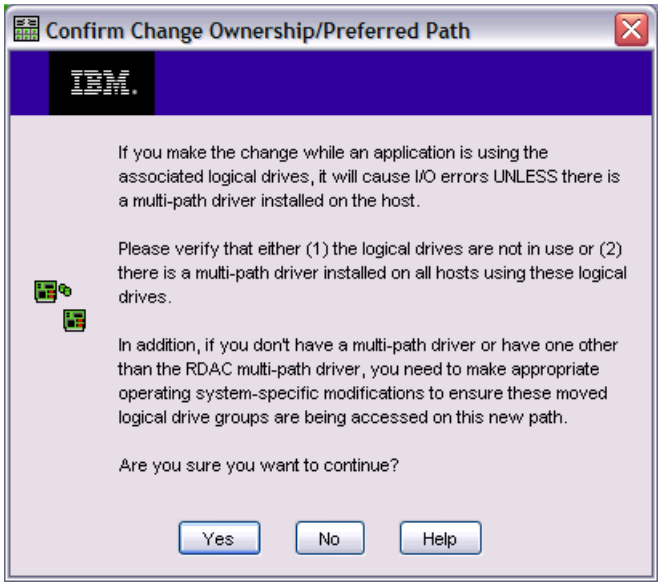


Figure 26-20 Confirm Change Ownership/Preferred Path

The task starts on the DS3400, as indicated in Figure 26-21. Click **OK** to return to the subsystem management window.

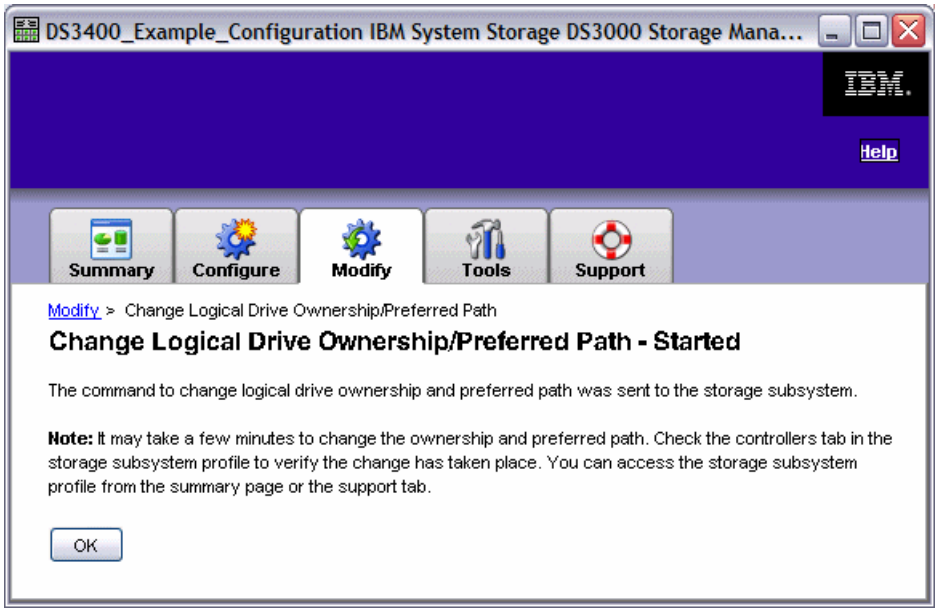


Figure 26-21 Change Logical Drive Ownership/Preferred Path - Started

We now have a logical drive available and assigned to our host server. The logical drive is owned by controller A, so we can proceed to the next section.

26.5 Configuring logical drive as boot device

We will now use Fast!UTIL to specify that our logical drive is the bootable device. Reboot the host server and press Ctrl+Q when prompted to launch Fast!UTIL.

Do these steps:

1. Select the first FC HBA port and then **Configuration Settings** → **Selectable Boot Settings**. This is shown in Figure 26-22.

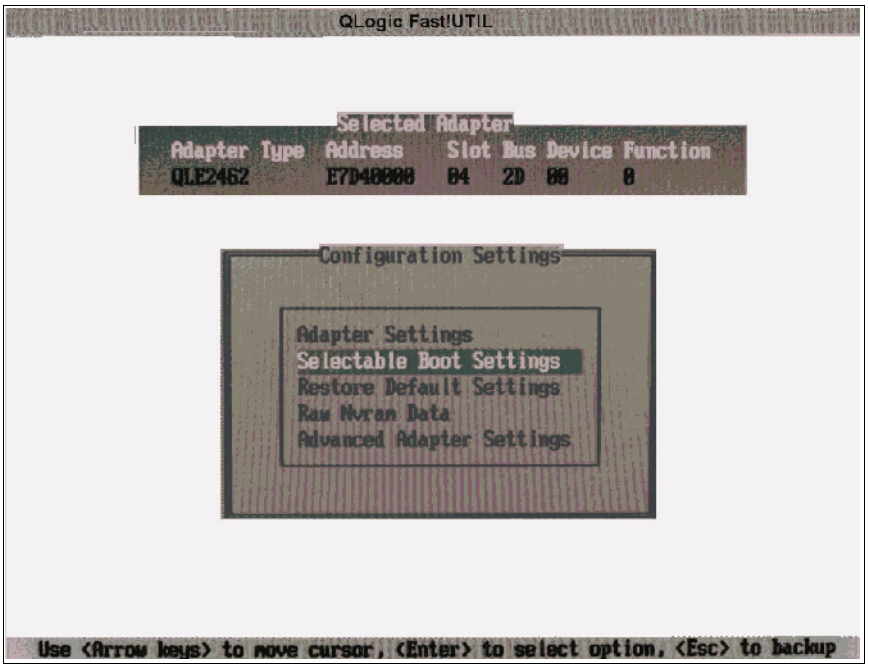


Figure 26-22 Selectable Boot Settings menu option

The Selectable Boot Settings menu displays. You need to do the following now:

1. In the Selectable Boot Settings menu, make sure the Selectable Boot field is set to **Enabled**.
2. Move the cursor to (Primary) Boot Port Name, Lun field and press Enter. Now select the DS3400 entry as the boot device. The controller and FC host port WWN should appear and the LUN should be 0 (see Figure 26-23 on page 659).

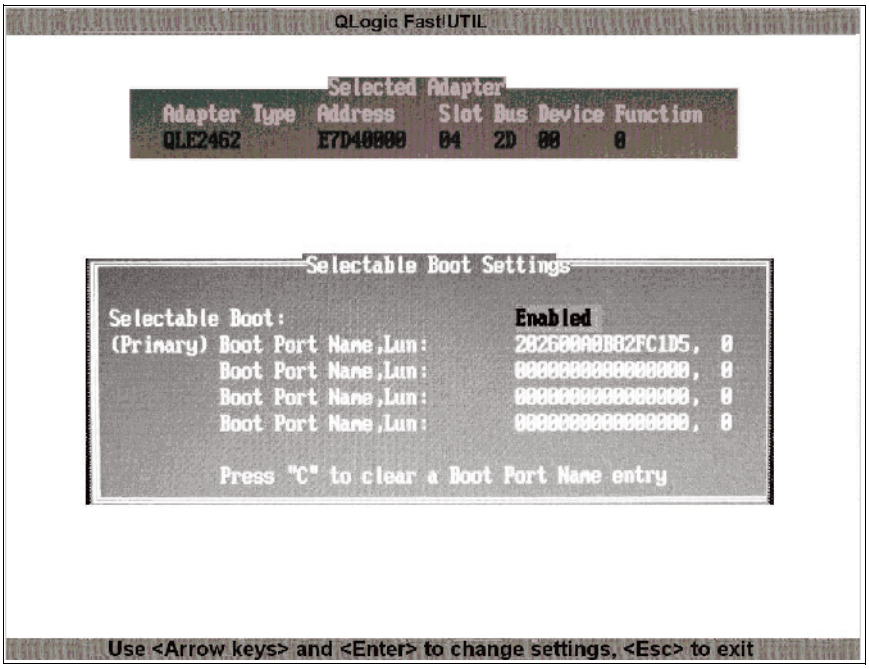


Figure 26-23 Selectable Boot Settings

3. Press Esc twice to save changes and exit Fast!UTIL.

26.6 Installing Linux

Now we will install Red Hat Enterprise Linux 5 (RHEL5) on our SAN-bootable disk. The installation DVD of RHEL5 already contains a driver for all current QLogic FC HBA adapters, so we do not need to provide this driver separately. However, we will update the driver to the latest level after installing if necessary.

The installation process is the same as installing RHEL on a local disk drive. The logical drive on the DS3400 server is presented to Linux as /dev/sda. We can create three partitions on the logical drive, as shown in Table 26-1.

Table 26-1 Partition layout

Partition	Mount point	File system	Size
/dev/sda1	/boot	ext3	100 MB
/dev/sda2	/	ext3	7 GB
/dev/sda3		swap	2 GB

We will install Red Hat Enterprise Linux 5, kernel 2.6.18-92.el5. The installation process is guided by the Anaconda installer and is quite straightforward and self explanatory. We will therefore not document it in detail.

Note: It is possible to install Red Hat 4 using similar steps and requires two post-installation tasks:

- Install the kernel source tree.
- Update the QLA2400 driver.

For Red Hat 5, you do not need to do these steps.

26.6.1 Installing the kernel source tree (RHEL4)

Once RHEL4 is installed and running, we need to install the kernel source. We obtain the kernel source RPM file (for example, from Red Hat Network at http://rhn.redhat.com). For our particular kernel version in our example, the file name is kernel-2.6.9-34.EL.src.rpm.

Install the kernel source by following these steps:

1. Install the source RPM with the following command:

   ```
   rpm -Uvh kernel-2.6.9-34.EL.src.rpm
   ```

2. Change to the following directory with this command:

   ```
   cd /usr/src/redhat/SPECS
   ```

3. Build the kernel source tree with this command:

   ```
   rpmbuild -bp kernel-2.6.spec --target=i686
   ```

4. The kernel source tree is now in the following directory:

   /usr/src/redhat/BUILD/kernel-2.6.9/linux-2.6.9

26.6.2 Updating the QLA2400 driver (RHEL4)

The RHEL4 installation CDs include the QLA2400 driver, but not the latest version. So we need to update the driver from the old version to the latest version (at the time of writing, this version was V8.01.06).

Download the latest QLogic driver from the QLogic Web site (linked from the IBM Support site):

http://driverdownloads.qlogic.com/QLogicDriverDownloads_UI/IBM.aspx?companyid=6

The file name is qla2xxx-8.01.06-dist.tgz.

To update the driver, do these steps:

1. Extract the contents from the driver file:

   ```
   tar -xvzf qla2xxx-8.01.06-dist.tgz
   ```

   This will create the subdirectory qlogic and put the driver files into it.

2. Change to the subdirectory qlogic and extract the driver source files by executing **drvrsetup**:

   ```
   cd qlogic
   ./drvrsetup
   ```

3. Use the **libinstall** command to set up the SNIA API library:

   ```
   ./libinstall
   ```

4. Run the script build.sh with the install parameter. This will build and install the driver module files:

   ```
   cd qla2xxx-8.01.06
   ./extras/build.sh install
   ```

   The QLA2400 driver modules will be installed to the directory /lib/modules/2.6.9-34.ELsmp/kernel/drivers/scsi/qla2xxx.

5. We need to edit the /etc/modprobe.conf file to add the following three lines:

   ```
   alias scsi_hostadapter0 qla2xxx_conf
   alias scsi_hostadapter1 qla2xxx
   alias scsi_hostadapter2 qla2400
   ```

   We also need to specify that the QLA2400 driver should run in non-failover mode, because I/O path failover will be provided by the Linux RDAC. Add the following line to /etc/modprobe.conf:

   ```
   options qla2xxx ql2xfailover=0
   ```

   The /etc/modprobe.conf file should look similar to Example 26-1.

Example 26-1 Updated /etc/modprobe.conf

```
alias eth0 bnx2
alias eth1 bnx2
alias scsi_hostadapter aacraid
alias usb-controller ehci-hcd
alias usb-controller1 ohci-hcd
install qla2xxx /sbin/modprobe qla2xxx_conf; /sbin/modprobe --ignore-install
qla2xxx
remove qla2xxx /sbin/modprobe -r --first-time --ignore-remove qla2xxx && {
/sbin/modprobe -r --ignore-remove qla2xxx_conf; }
alias scsi_hostadapter0 qla2xxx_conf
alias scsi_hostadapter1 qla2xxx
alias scsi_hostadapter2 qla2400
options qla2xxx ql2xfailover=0
```

6. Create a new RAMdisk image to be used at boot time. We do this with the following three commands:

   ```
   cd /boot
   depmod -a
   mkinitrd -f new_initrd-2.6.9-34.ELsmp.img 2.6.9-34.ELsmp
   ```

   The new RAMdisk image file name will be new_initrd-2.6.9-34.ELsmp.img.

7. Finally, edit the /boot/grub/menu.lst file to specify that the new RAMdisk image should be used. We recommend adding a new boot menu entry for the new RAMdisk file. Example 26-2 shows modified /boot/grub/menu.lst file.

Example 26-2 Modified /boot/grub/menu.lst

```
# grub.conf generated by anaconda
#
# Note that you do not have to rerun grub after making changes to this file
# NOTICE:  You have a /boot partition.  This means that
#          all kernel and initrd paths are relative to /boot/, eg.
#          root (hd0,0)
```

```
#          kernel /vmlinuz-version ro root=/dev/sda2
#          initrd /initrd-version.img
#boot=/dev/sda
default=0
timeout=10
splashimage=(hd0,0)/grub/splash.xpm.gz
hiddenmenu
title Red Hat Enterprise Linux AS (2.6.9-34.ELsmp)
   root (hd0,0)
   kernel /vmlinuz-2.6.9-34.ELsmp ro root=LABEL=/ pci=nommconf rhgb quiet
   initrd /initrd-2.6.9-34.ELsmp.img
title Red Hat Enterprise Linux AS (2.6.9-34.ELsmp) New QLA2400 Driver
   root (hd0,0)
   kernel /vmlinuz-2.6.9-34.ELsmp ro root=LABEL=/ pci=nommconf rhgb quiet
   initrd /new_initrd-2.6.9-34.ELsmp.img
title Red Hat Enterprise Linux AS-up (2.6.9-34.EL)
   root (hd0,0)
   kernel /vmlinuz-2.6.9-34.EL ro root=LABEL=/ pci=nommconf rhgb quiet
   initrd /initrd-2.6.9-34.EL.img
title Other
   rootnoverify (hd2,0)
   chainloader +1
```

8. Restart the server and select the boot menu option with the updated RAMdisk image. When Linux boots up, we want to verify that the QLA2400 driver is properly updated:

   a. Run this command:

      ```
      ls -al /proc/scsi/qla2xxx
      ```

   Two files should be shown, one for each FC HBA port, as you can see in Example 26-3.

Example 26-3 Two files representing two FC HBA ports

```
total 0
dr-xr-xr-x  2 root root 0 Jun 20 14:43 .
dr-xr-xr-x  5 root root 0 Jun 20 14:43 ..
-rw-r--r--  1 root root 0 Jun 20 14:43 1
-rw-r--r--  1 root root 0 Jun 20 14:43 2
```

   b. Display the contents of these two files with the **cat** command:

      ```
      cat /proc/scsi/qla2xxx/1
      cat /proc/scsi/qla2xxx/2
      ```

      The driver version will be shown in the second line (see Example 26-4).

Example 26-4 QLA2400 driver version

```
QLogic PCI to Fibre Channel Host Adapter for QLE2462:
        Firmware version 4.00.23 [IP] , Driver version 8.01.06
ISP: ISP2432
...
...
...
```

      As you can see, the driver version is now 8.01.06, so the update procedure was successful.

26.7 Installing Linux RDAC in RHEL5

RDAC for Linux can be downloaded from the IBM Support Web site. The IBM Web page leads you to the actual download Web page URL:

http://www.lsi.com/rdac/ds3000.html

Download the appropriate RDAC for your Linux kernel version (either 2.4 or 2.6). In our case, we are using RHEL5, so we download the latest RDAC for kernel version 2.6. For previous versions of Red Hat, the procedure is almost identical.

26.7.1 Building the RDAC driver

Do these steps:

1. The RDAC package comes as a compressed tar file. Unpack the file using the following command:

   ```
   tar -zxvf rdac-LINUX-09.03.C5.23-source.tar.gz
   ```

2. The source files will be unpacked in the linuxrdac-09.03.C5.23 subdirectory of the current directory. Change to this subdirectory. Note that the version of RDAC can be different; the name/version we mentioned is used in this example.

3. Use the following command to remove old driver modules:

   ```
   make clean
   ```

4. Compile the driver modules and build the RDAC driver. Use this command:

   ```
   make
   ```

26.7.2 Installing the RDAC driver

We need to copy the driver modules to the kernel module tree and build the new RAMdisk image which includes the RDAC driver modules.

Do these steps:

1. Run **`make install`**. This will produce a lengthy output in the console window, as the driver modules are applied to the kernel tree. The last few lines will tell you how to add the new boot menu option to the boot loader, as shown in Example 26-5.

Example 26-5 RDAC installation

```
Creating new MPP initrd image...
You must now edit your boot loader configuration file, /boot/grub/menu.lst, to
add a new boot menu, which uses mpp-2.6.18-92.el5.img as the initrd image.
Now Reboot the system for MPP to take effect.
The new boot menu entry should look something like this (note that it may
vary with different system configuration):
...
title Red Hat Linux (2.6.18-92.el5) with MPP support
root (hd0,5)
kernel /vmlinuz-2.6.18-92.el5 ro root=LABEL=RH9
initrd /mpp-2.6.18-92.el5.img
...
MPP driver package has been sucessfully installed on your system.
```

2. The RDAC installation created a new RAMdisk image file in the /boot directory named mpp-2.6.18-92.el5.img. Modify the /boot/grub/menu.lst file to add the new boot menu option; the extra lines are shown in Example 26-6.

Example 26-6 Boot menu option for /boot/grub/menu

```
title Red Hat Enterprise Linux AS (2.6.18-92.el5) RDAC (MPP)
        root (hd0,0)
        kernel /vmlinuz-2.6.18-92.el5 ro root=LABEL=/ pci=nommconf rhgb quiet
        initrd /mpp-2.6.18-92.el5.img
```

3. Restart the system using the new boot option.

26.7.3 Verification

When the server restarts, use the **lsmod** command to verify that the RDAC (MPP) modules are running. The following modules should be listed:

- mppUpper
- mppVhba
- sg
- sd_mod
- scsi_mod

Run the following command:

```
ls -lR /proc/mpp
```

This command produces output similar to Example 26-7.

Example 26-7 ls -lR /proc/mpp output

```
# ls -lR /proc/mpp
/proc/mpp:
total 0
dr-xr-xr-x 3 root root 0 Sep  5 12:15 DS3400_Thailand

/proc/mpp/DS3400_Thailand:
total 0
dr-xr-xr-x 3 root root 0 Sep  5 12:15 controllerA
-rw-r--r-- 1 root root 0 Sep  5 12:15 virtualLun0

/proc/mpp/DS3400_Thailand/controllerA:
total 0
dr-xr-xr-x 2 root root 0 Sep  5 12:15 qla2xxx_h0c0t0

/proc/mpp/DS3400_Thailand/controllerA/qla2xxx_h0c0t0:
total 0
-rw-r--r-- 1 root root 0 Sep  5 12:15 LUN0
```

As you can see, RDAC is communicating with the DS3400 subsystem; however, it can only talk to controller A at this time, because only one I/O path is configured. To enable multipath, configure the I/O path between FC HBA port 1 and DS3400 controller B.

26.8 Enabling the second I/O path

Change the effective SAN switch configuration. Up to now, the effective configuration was x3755_CtrlA_only, as defined in 26.3, “SAN switch zoning for the first I/O path” on page 644.

Enable **x3755_bothCtlr** as the effective configuration now. This will allow communication across both I/O paths.

26.8.1 Adding a second FC HBA port to the host server definition

Our host server is still defined with one host port only. Now we can add the second host port, which will communicate with controller B.

Do these steps:

1. In the Storage Manager Modify tab, click **Edit Topology**. Click **Add HBA** to specify the new host port, as shown in Figure 26-24.

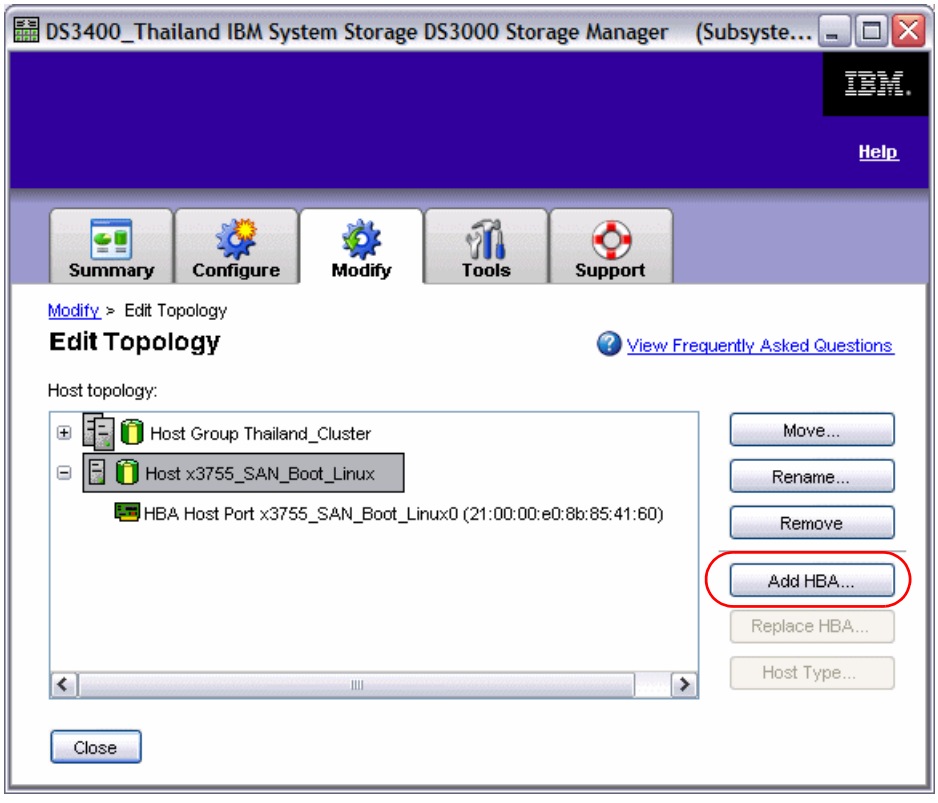


Figure 26-24 Edit Topology - Add HBA

2. This launches the Add HBA Host Port window, as shown in Figure 26-25.

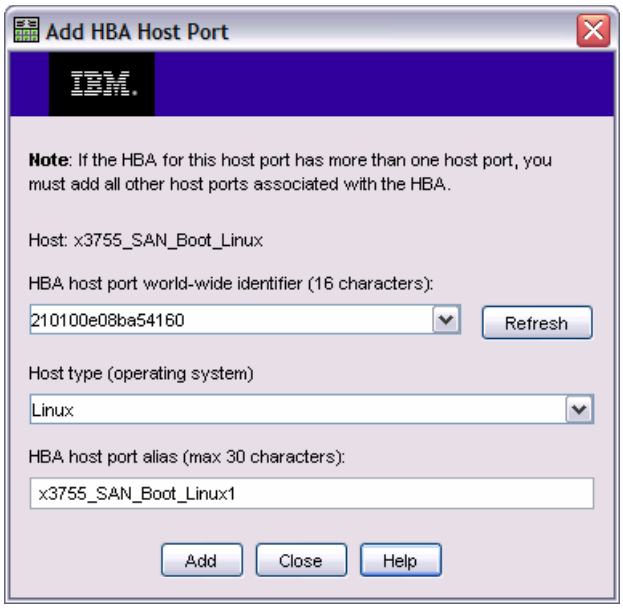


Figure 26-25 Add HBA Host Port

3. Set the correct values for WWPN, host type, and host port alias, and then click **Add**. As you can see in Figure 26-26, our host server now contains both host port definitions.

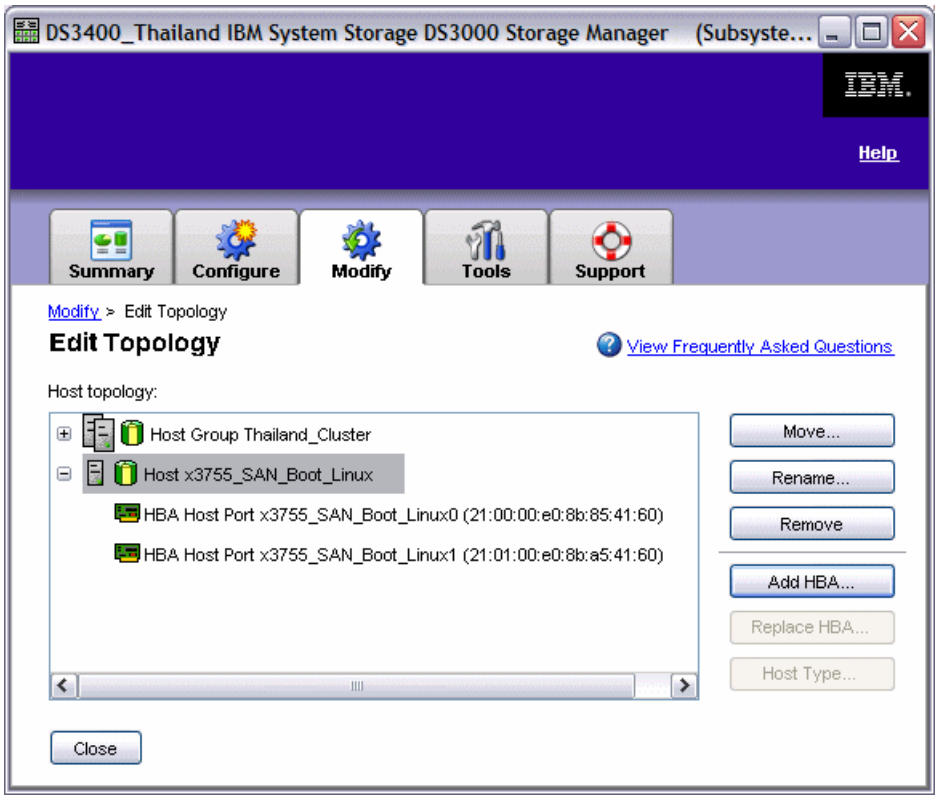


Figure 26-26 Both host ports are defined

26.8.2 Enabling a second FC HBA port for SAN boot

We have the second FC path to the server ready. Now, we can enable the second FC port for boot from SAN.

Note: SAN boot from the alternate HBA port is possible in case you have the first HBA port unavailable, for example, the FC cable is unplugged, or if controller A of the DS3400 server is not online or is disconnected. The boot logical drive has to be owned by the alternate controller of DS3400.

Warning: We recommend having AVT/ADT disabled for the RDAC drivers, and if RDAC is not active during boot time yet, you have to have the correct logical drive ownership set during the boot time. If the boot drive is owned by controller A and the first HBA port is unplugged or controller A is unplugged, the boot will hang, as no boot logical drive is visible on the SAN. You must manually switch ownership to controller B to successfully boot using the second HBA port.

To enable the second FC port for boot from SAN, follow these steps:

1. Power off the server and change ownership of the boot logical drive in the DS3400 server. The procedure is described in 26.4.3, “Logical drive ownership” on page 656. At this time, you have to change the preferred ownership of boot LUN=0 to controller B.
2. Disconnect the FC cable from the first HBA port.
3. Power on the server and enter Fast!UTIL by pressing Ctrl+Q when prompted. Select the second HBA port (Figure 26-27) and press Enter.

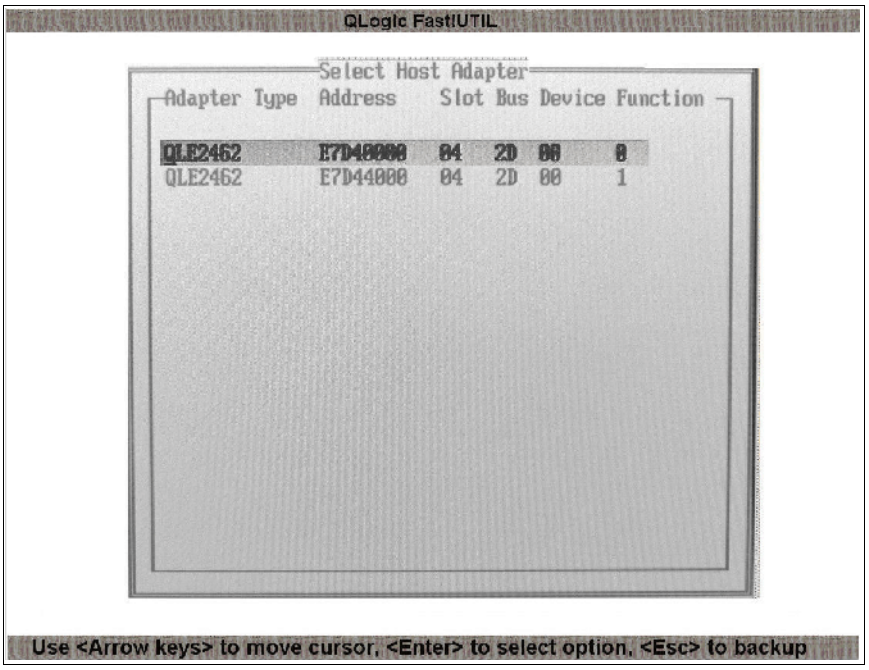


Figure 26-27 Select the second FC HBA port

4. Select **Configuration Settings** and press Enter.
5. In the Configuration Settings menu, select **Adapter Settings** and press Enter.
6. Change the Host Adapter BIOS setting from **Disabled** to **Enabled** (see Figure 26-28) and then press Esc to go back.

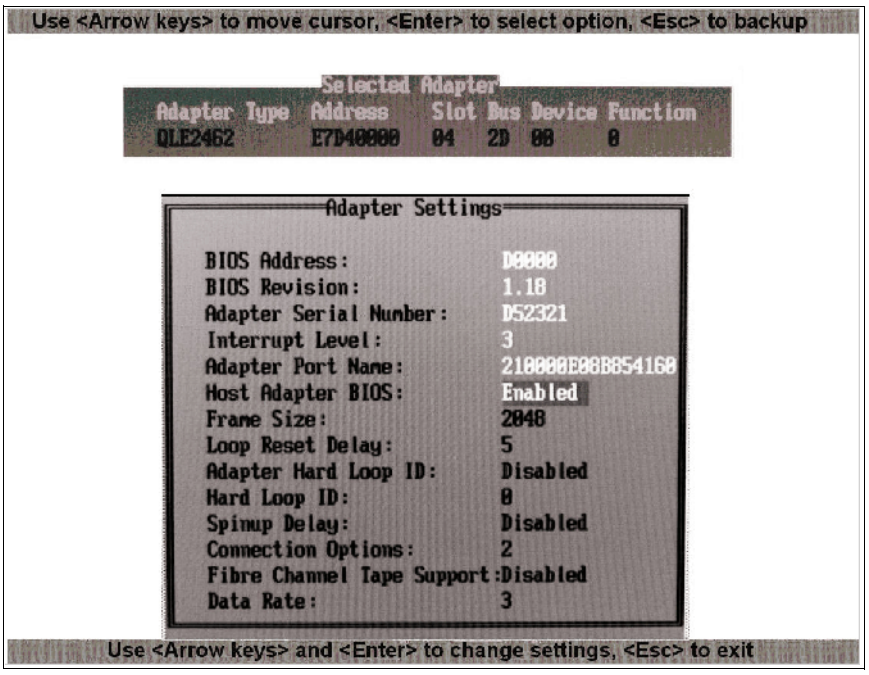


Figure 26-28 Adapter settings

7. Now select **Selectable Boot Settings**. This is shown in Figure 26-22 on page 658.
8. In the Selectable Boot Settings menu, make sure the Selectable Boot field is set to **Enabled**.
9. Move the cursor to (Primary) Boot Port Name, LUN field and press Enter. Now select the DS3400 entry as the boot device. The controller B FC host port WWN should appear and the LUN should be 0 (see Figure 26-29 on page 669). If you do not see an entry, verify that the ownership of LUN=0 is on controller B and zoning is done correctly.

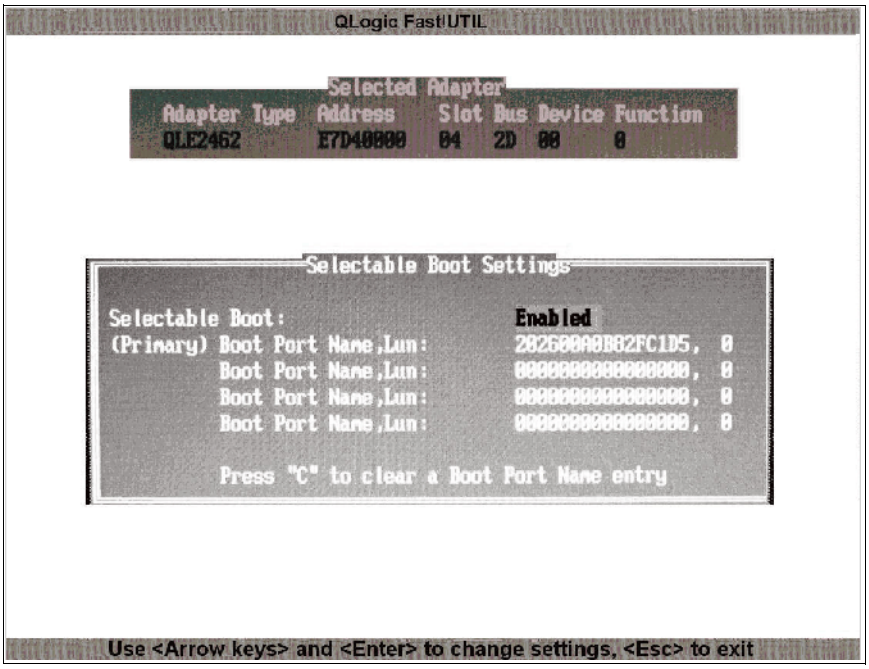


Figure 26-29 Selectable Boot Settings

10. Press Esc twice, save your changes, and exit Fast!UTIL.
11. Your server reboots again. Remember that you disconnected the first HBA port before, so the server boot using the second HBA port. Controller B is still the owner of the boot LUN.
12. After you verify the boot was successful, you have to connect the FC cable to the first FC port and change ownership back to controller A. Your failover FC boot from SAN solution is done.

26.9 Verifying the installation

Run the following command to verify that RDAC installation was successful:

```
ls -lR /proc/mpp
```

The output of this command is shown in Example 26-8.

Example 26-8 Both I/O paths are enabled

```
/proc/mpp:
total 0
dr-xr-xr-x  4 root root 0 Jun 5 17:23 DS3400_Thailand

/proc/mpp/DS3400_Thailand:
total 0
dr-xr-xr-x  3 root root 0 Sep 5 17:23 controllerA
dr-xr-xr-x  3 root root 0 Jun 5 17:23 controllerB
-rw-r--r--  1 root root 0 Jun 5 17:23 virtualLun0

/proc/mpp/DS3400_Thailand/controllerA:
total 0
dr-xr-xr-x  2 root root 0 Jun 5 17:23 qla2xxx_h1c0t0
```

```
/proc/mpp/DS3400_Thailand/controllerA/qla2xxx_h1c0t0:
total 0
-rw-r--r--  1 root root 0 Jun 5 17:23 LUN0

/proc/mpp/DS3400_Thailand/controllerB:
total 0
dr-xr-xr-x  2 root root 0 Jun 5 17:23 qla2xxx_h2c0t0

/proc/mpp/DS3400_Thailand/controllerB/qla2xxx_h2c0t0:
total 0
-rw-r--r--  1 root root 0 Jun 5 17:23 LUN0
```

As you can see in Example 26-8 on page 669, the RDAC now works with both controllers and manages both I/O paths. This completes the redundant SAN boot configuration and installation of Red Hat Enterprise Linux 5.

26.9.1 RDAC commands

RDAC drivers allow for the addition of new LUNs without needing to reboot the server. You can create new logical drives, map them to the server, and discover them in the Red Hat operating system by using two equivalent commands:

- **hot_add**
- **mppBusRescan**

Another useful command is **mppUtil**. It displays all detailed information regarding logical disk drives (LUNs), controllers, and disk subsystems in the operating system. Here are two examples of this command:

- **mppUtil -a** (This command shows the DS3400 subsystems.)
- **mppUtil -g 0** (This command shows all objects in DS3400 subsystem ID=0.)

27

FC configuration 4 - QLogic HBA boot for IBM System x from SAN

In this chapter, we discuss a sample configuration, describing how to boot a System x with an QLogic HBA, from a DS3400 System Storage subsystem, using the Windows Server 2008 operating system. Before you start, make sure that all the firmware on the storage subsystem and in the host is flashed to the latest level, as described in 14.3, “Download firmware” on page 307.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

27.1 What is required

For our setup, we used the following hardware and software components:

- IBM System x server
- Two PCI Express FC HBAs, P/N 24P0962 (QLA2340)
- SAN switch
- DS3400 System Storage subsystem with enough free disk space for the logical drive
- A management workstation running the latest version of DS3000 Storage Manager (at the time of writing, Version 10.35)
- Windows Server 2008 operating system installation media

In 4.5, “Host attachment” on page 40, we discuss the most common host cabling scenarios for Fibre Channel. In this sample configuration, we use a System x server as the host, connected through a single Fibre Channel (FC) switch to the DS3400 System Storage subsystem, even though we recommend having two FC switches to provide redundancy on the FC switch layer. Our topology, even though we are only using one System x host, is shown in Figure 4-8 on page 44. The System x server has two separate Fibre Channel (FC) HBAs installed. Each HBA port is connected to the switch. One connection on the switch is connected to Controller A of the DS3400 server, and another connection on the switch is connected to Controller B. Figure 27-1 shows our specific configuration.

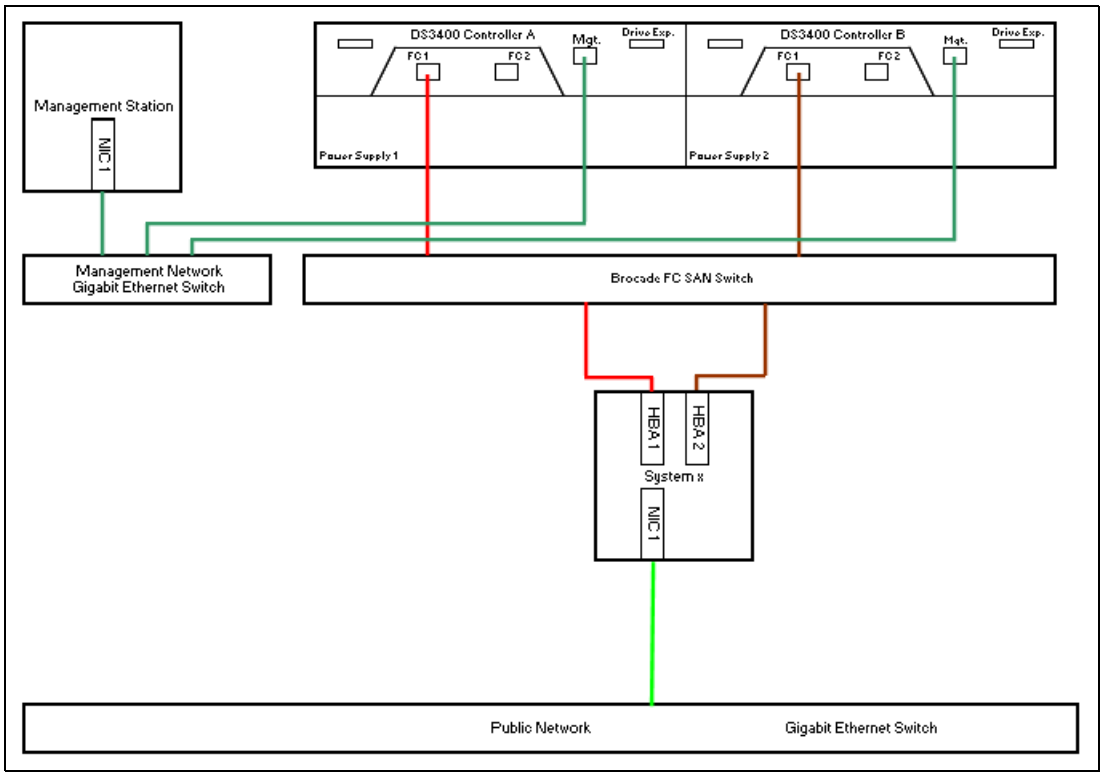


Figure 27-1 Our sample environment

27.2 First FC HBA adapter configuration

In this section, we use the QLogic HBA BIOS utility (named Fast!UTIL) to note the WWNs of both HBA adapters and to enable the BIOS on the first HBA adapter.

Do these steps:

1. To enter Fast!UTIL, power the host server on and press Ctrl+Q when prompted. The first step is HBA adapter selection. as shown in Figure 27-2.

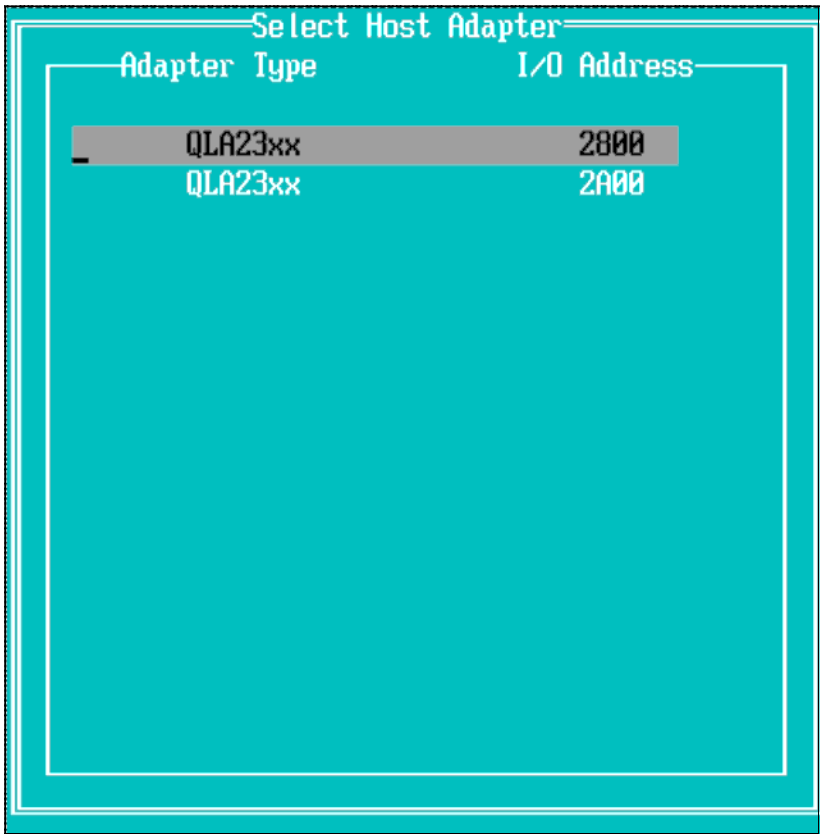


Figure 27-2 Select the FC HBA adapter

2. Select the first HBA adapter and press Enter. The Fast!UTIL Options menu displays (Figure 27-3). Select **Configuration Settings** and press Enter.

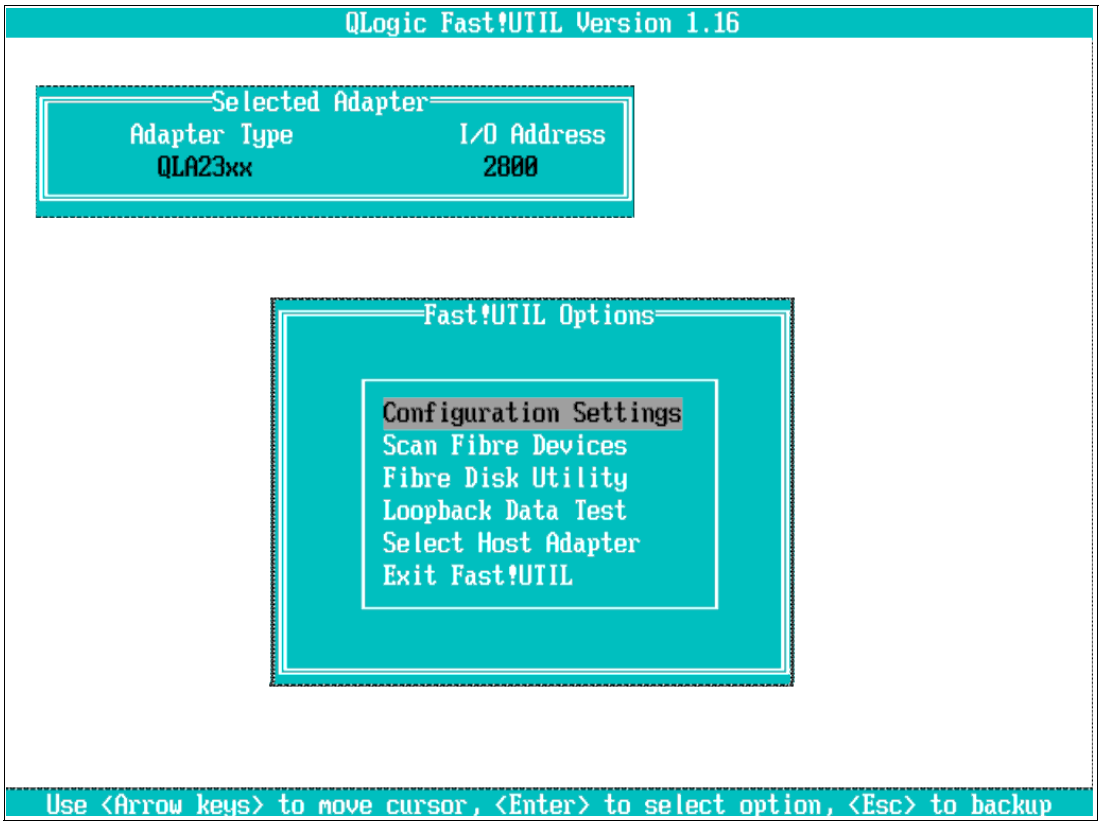


Figure 27-3 Fast!UTIL Options

3. In the Configuration Settings menu, shown in Figure 27-4, select **Adapter Settings**.



Figure 27-4 Configuration Settings menu

4. The Adapter Port Name field shows the HBA port WWN. In our case, the WWN is 21:00:00:E0:8B:89:04:C0 (see Figure 27-5). Record this value.

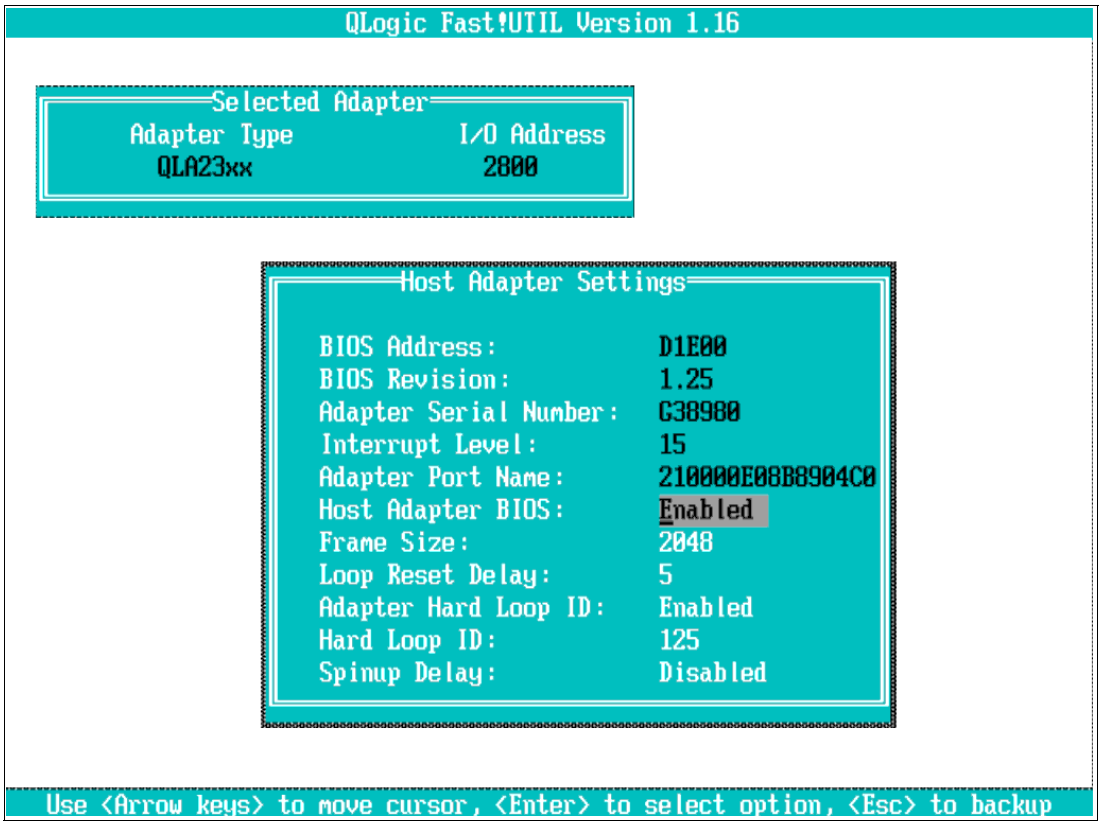


Figure 27-5 Adapter settings

5. Here we also enable the BIOS for this HBA adapter. Remember, the BIOS must be enabled if we want to boot through this HBA adapter.

   Change the Host Adapter BIOS setting from **Disabled** to **Enabled**, then press Esc. Select **Save changes** to confirm.

6. Later in the procedure, we also need the WWN of the second HBA adapter. Display and record this information now.

7. Switch to the second HBA adapter. Return to the Fast!UTIL Options menu (Figure 27-3 on page 674) and choose Select **Host Adapter**. You see a second QLE2340 HBA adapter displayed, as shown in Figure 27-2 on page 673. Select the second line and press Enter.

8. Select **Configuration Settings** → **Adapter Settings** and look for the WWN of the second HBA adapter. As you can see in Figure 27-6 on page 677, the WWN is 21:01:00:E0:8B:0A:18:27.

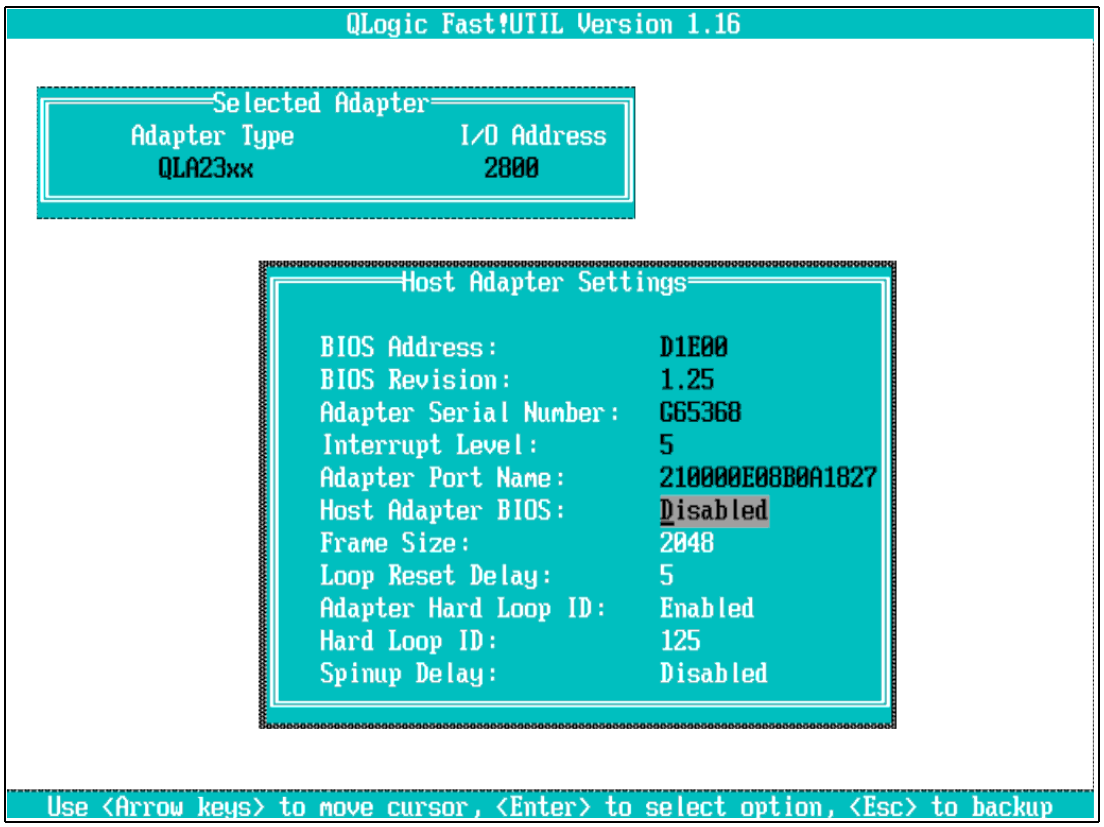


Figure 27-6 Second FC HBA port WWN

In this section, we enabled the BIOS on the first HBA adapter and noted the WWNs of both HBA adapters:

- HBA port 1: 21:00:00:E0:8B:89:04:C0
- HBA port 2: 21:01:00:E0:8B:0A:18:27

27.3 Configuring the system BIOS

After Boot BIOS is enabled on the System x server for the QLogic HBA, you must set up the system BIOS within the server. This can be accomplished by powering off or rebooting the server blade and pressing the F1 key.

Note: There may be slight differences within the System BIOS Configuration and Setup Utility, depending on the System x model and BIOS version being used. Knowledge of BIOS and ROM memory space usage may be required in certain situations.

1. In our example, the screen shown in Figure 27-7 appears. Select **Devices and I/O ports.**

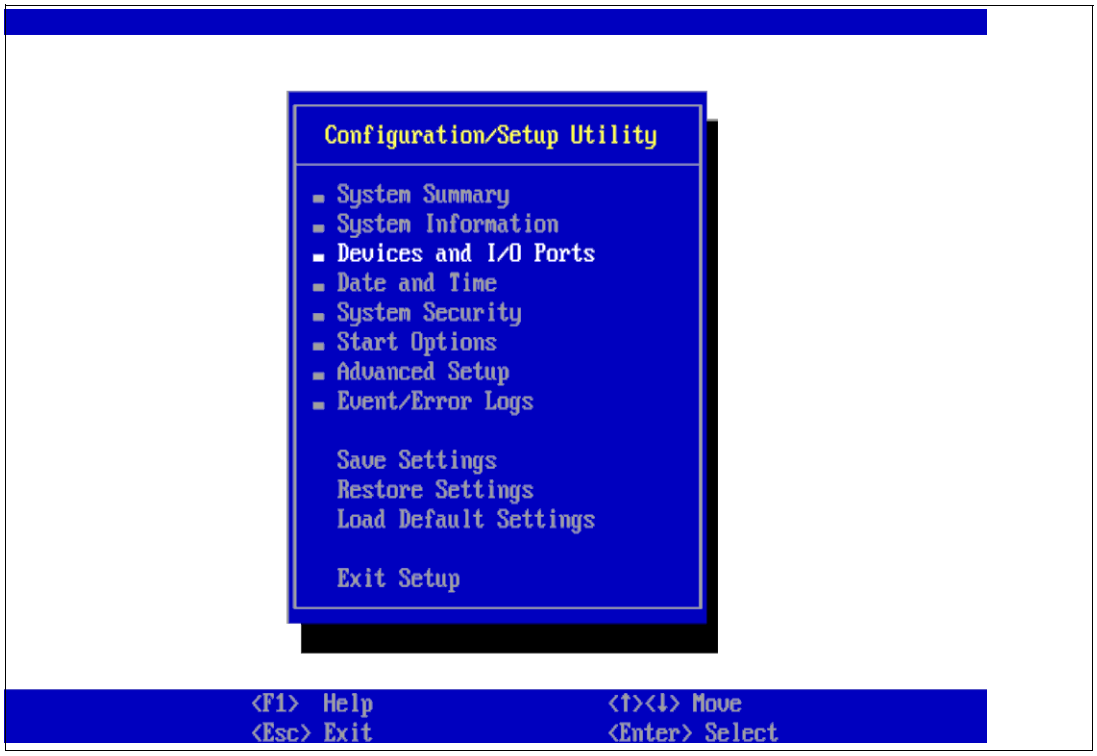


Figure 27-7 System BIOS

2. Depending on the internal disk technology of the host server, you have to disable this internal boot device. In our example, the System x server uses internal SAS disks. We disable the onboard Planar SAS, as shown in Figure 27-8. If available, select the IDE/SAS/SATA configuration menu and disable the primary and secondary.

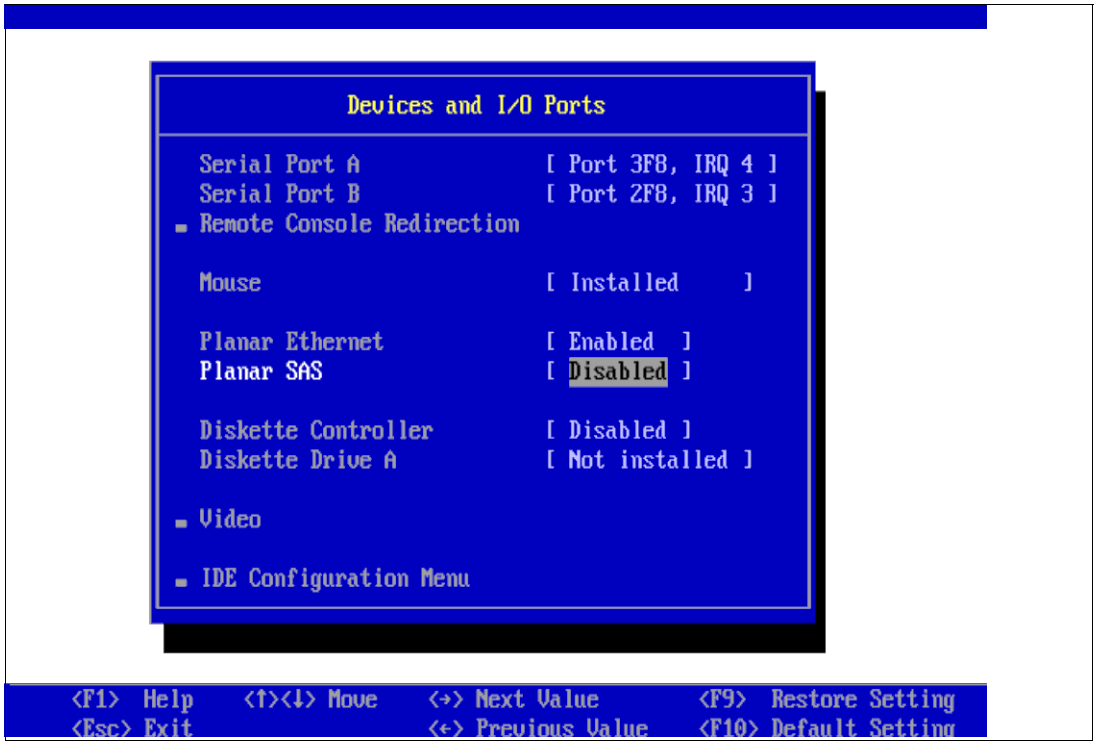


Figure 27-8 Disabling internal boot device

3. From the main Configuration/Setup Utility window, select **Start Options**. Verify the Startup Sequence Option is configured, as shown in Figure 27-9.

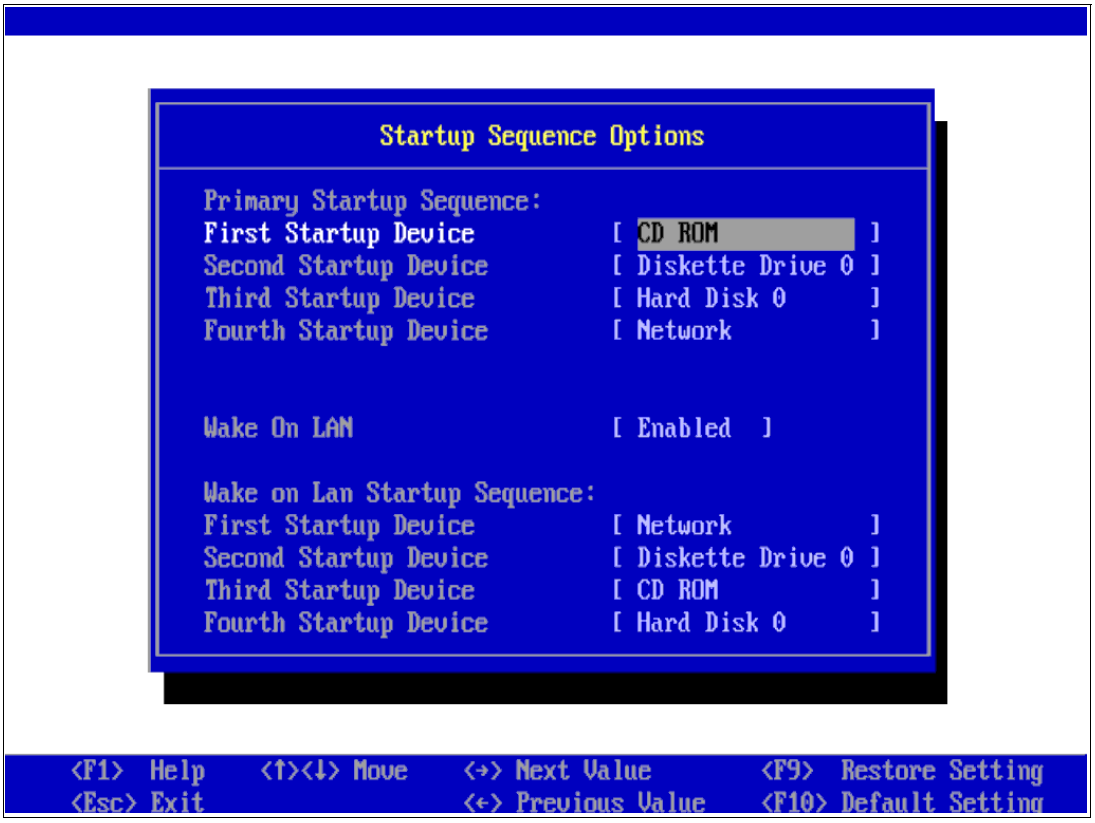


Figure 27-9 Startup Sequence Options

4. Depending on which PCI slot your first HBA adapter is placed in, select the appropriate boot priority for PCI devices. From the Start Options menu, select the PCI Device Boot Priority; for our example, the first HBA card is in slot 1, as shown in Figure 27-9.

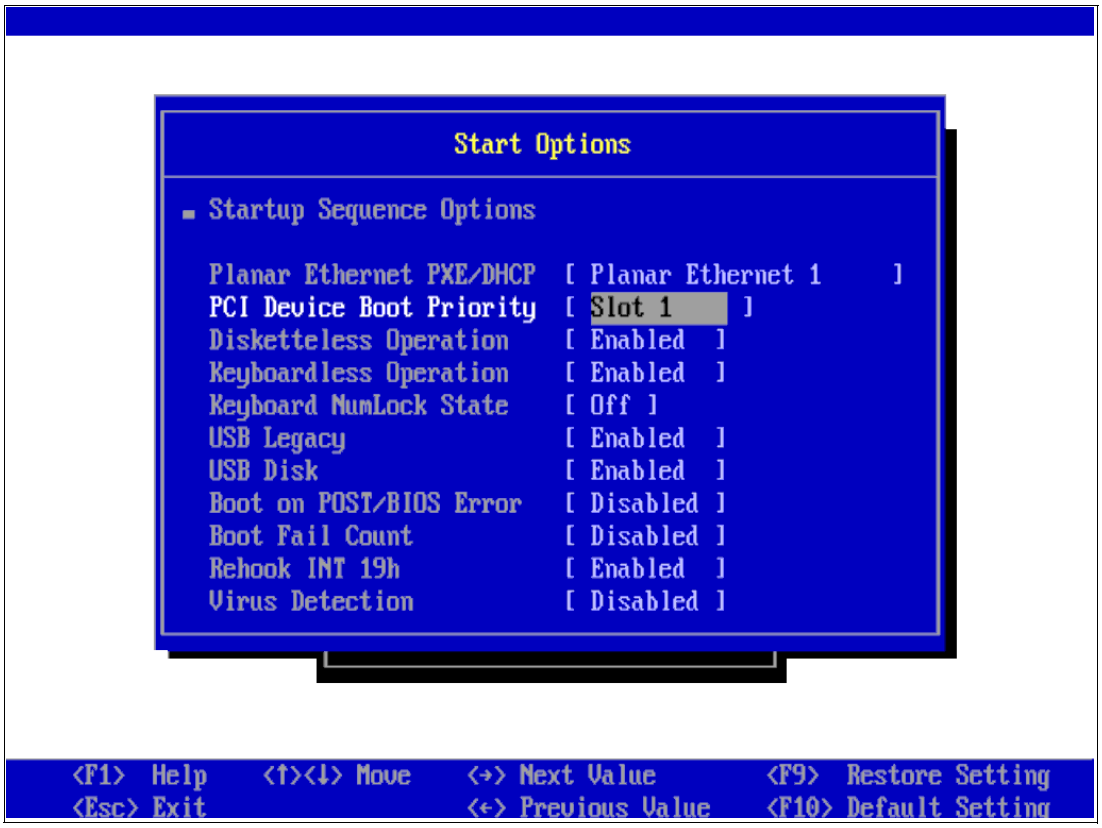


Figure 27-10 PCI Device Boot Priority

5. Exit the system BIOS and reboot the System x server. Boot into the QLogic HBA BIOS utility (named Fast!UTIL) by pressing Ctrl+Q when prompted.

27.4 SAN switch zoning for the first I/O path

To do zoning properly, you should have the System x server turned on, and be in the QLogic HBA BIOS utility so that the HBA adapter will log in to the switch. For most of the installation process, we only use single path (the first FC HBA port connected to controller A). Only after the Windows Multi-Path driver is installed will we define and enable the second I/O path (the second HBA adapter connected to controller B). This avoids problems and conflicts that could occur if the logical drive was seen across both I/O paths before having the Windows Server 2008 installed and operational.

We zone the fabric so that only the first HBA adapter of the System x server is zoned to controller A of the DS3400 System Storage subsystem that we are going to boot from, as shown in Figure 27-11. For the OS installation, the OS can only see one path to the boot disk.

We recommend that you put all the boot volumes on one controller (A).

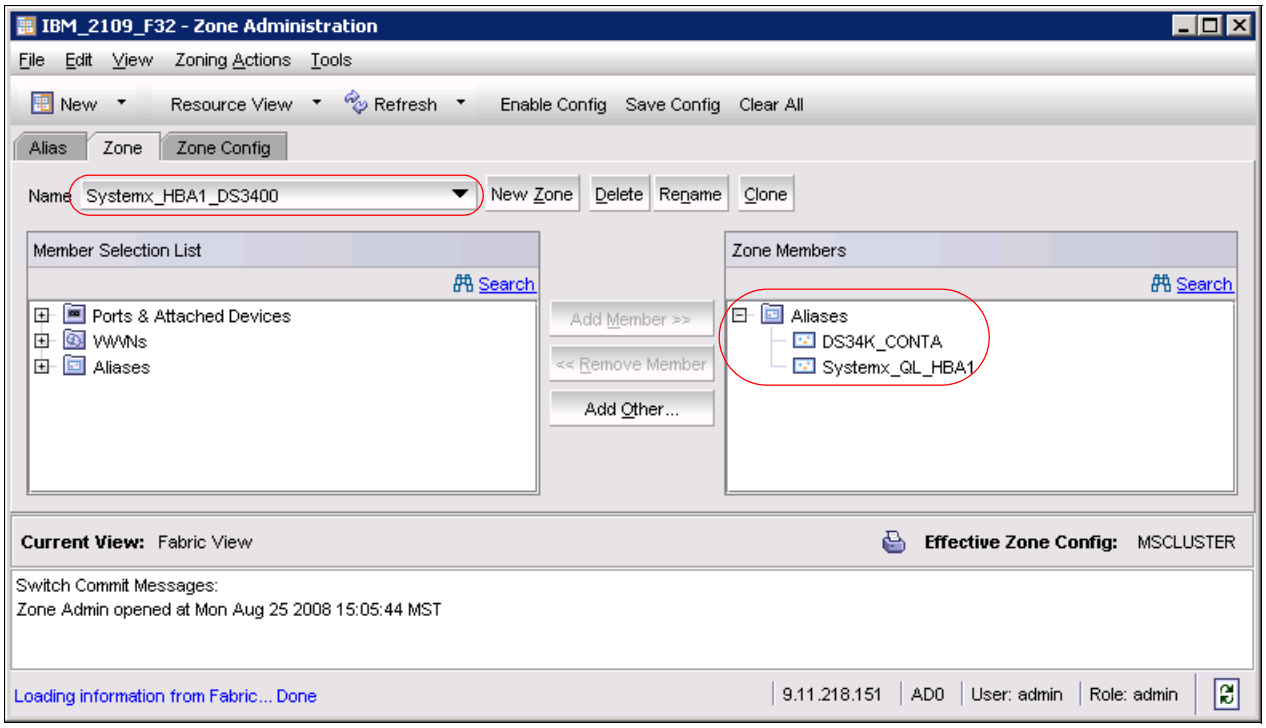


Figure 27-11 Zoning configuration

27.5 Creating a logical drive

We will use the DS3000 Storage Manager running on a management workstation to create a boot logical drive that will be used to install and later boot Windows Server 2008. But before creating the logical drive, we have to define our host server and its host port (the first HBA adapter).

27.5.1 Host server definition

1. In the DS3000 Storage Manager subsystem management window, select the **Configure** tab, and click **Configure Host Access (Manual)**, as shown in Figure 27-12 on page 683.

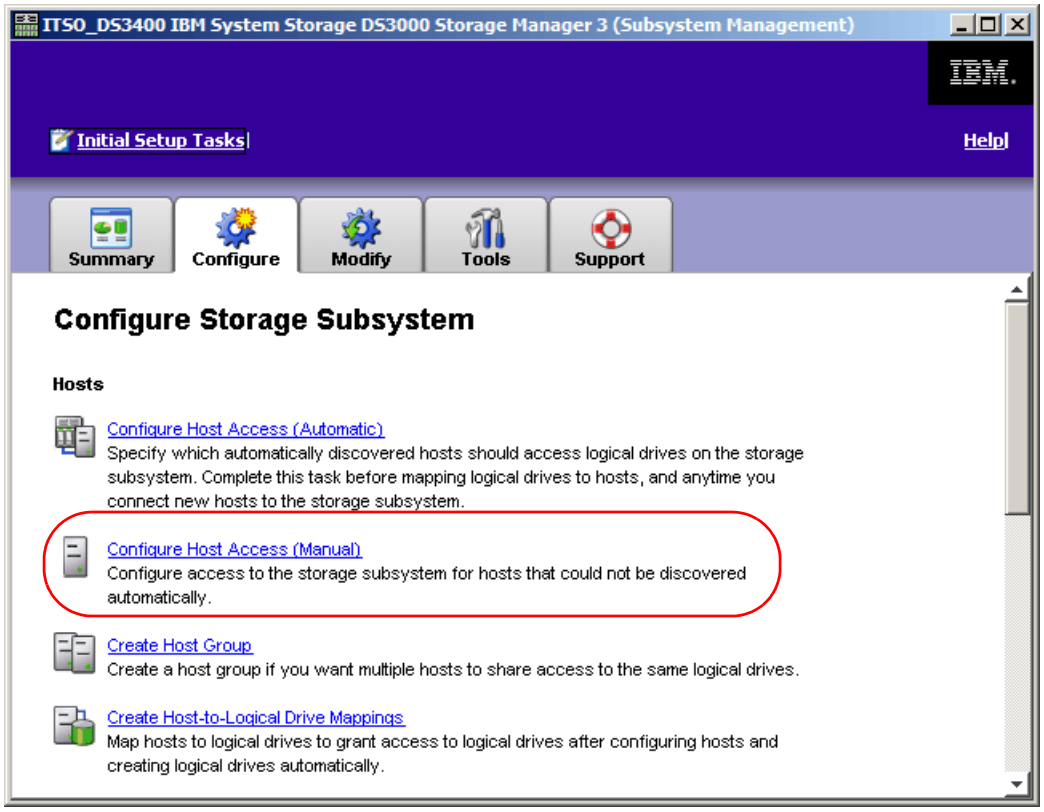


Figure 27-12 Configure Host Access (Manual)

2. Specify the host name and type. In our example, the host server name will be Systemx_SAN_Boot and the type will be Windows 2000/Server 2003/Server2008 Non-Clustered (supports DMP) host type, because we want to have auto volume transfer (AVT) enabled for boot path protection (see Figure 27-13).

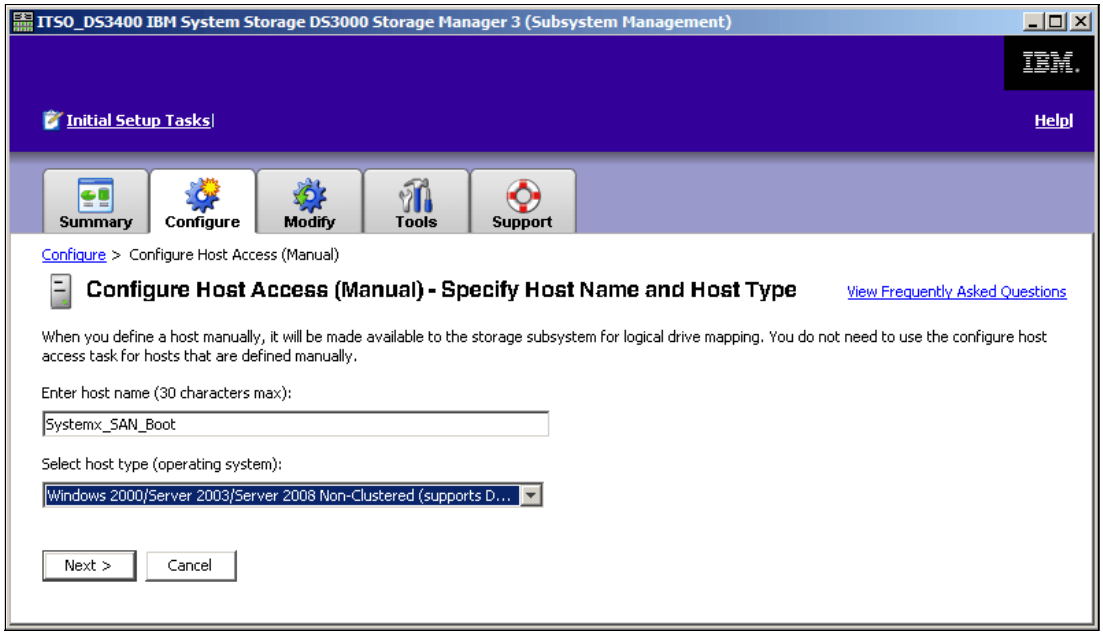


Figure 27-13 Specify Host Name and Host Type

3. Specify the HBA host port. Remember, we have to specify the first HBA adapter only. Its WWN is 21:00:00:E0:8B:89:04:C0. Select this WWN and click **Add**, as shown in Figure 27-14.

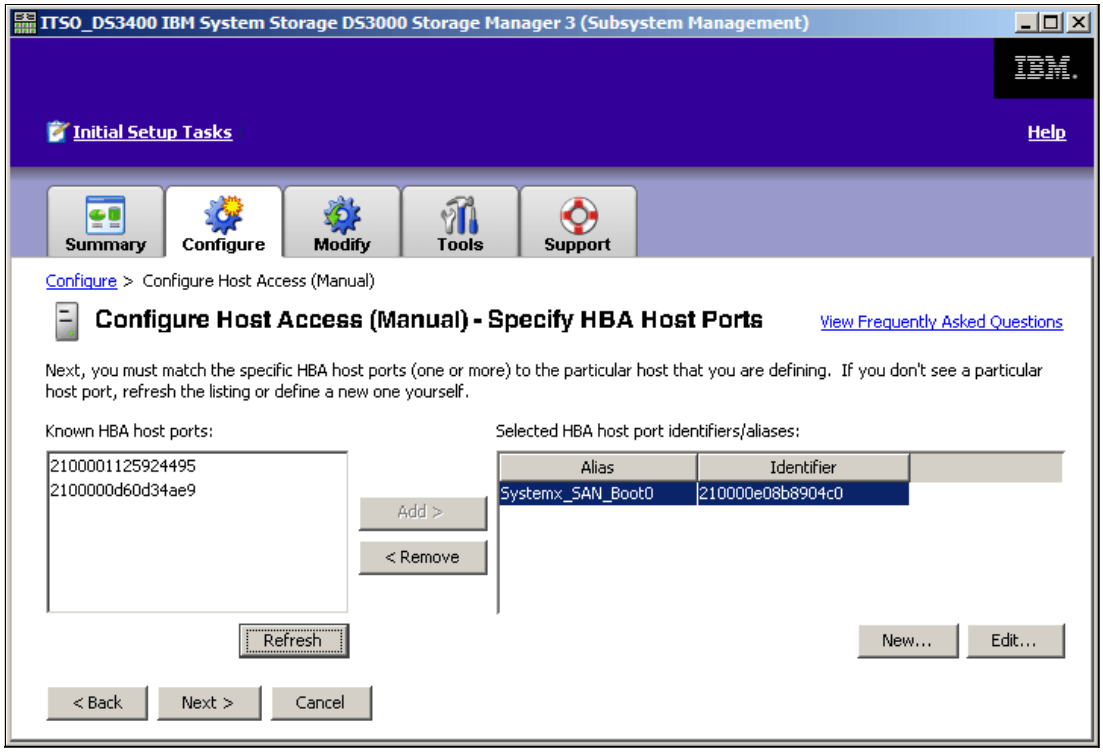


Figure 27-14 Specify HBA Host Port

4. Next, you can specify the host group. Since our host server is not clustered, we will not create any host group. Select **No: This host will NOT share access to the same logical drives with other hosts** and click **Next**, as shown in Figure 27-15 on page 685.

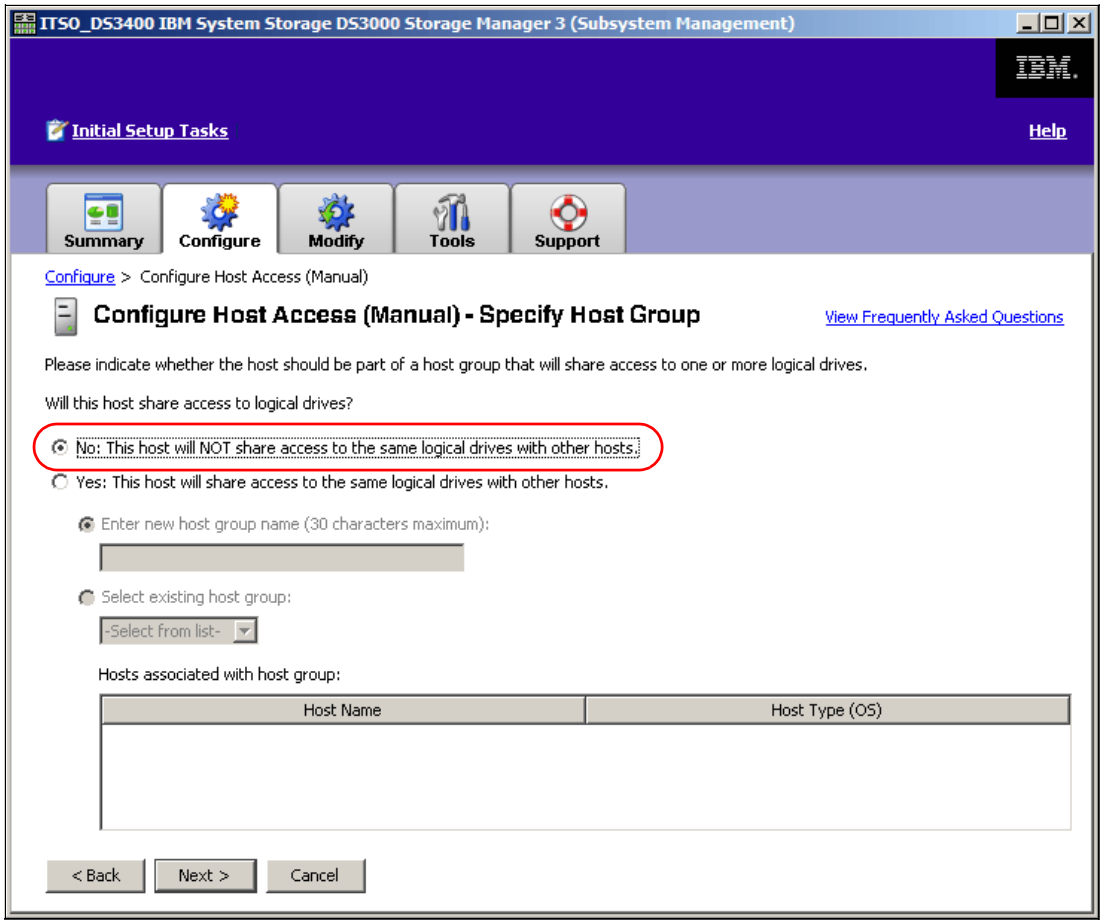


Figure 27-15 Specify Host Group

5. The confirmation window displays, as shown in Figure 27-16. Review the host definition and click **Finish**.

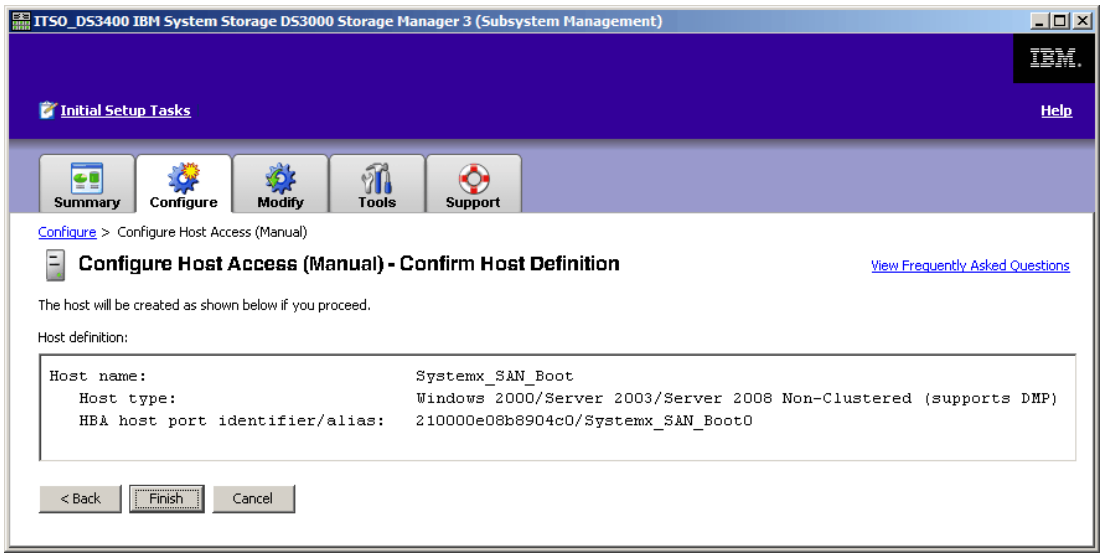


Figure 27-16 Confirm Host Definition

6. In the completion window (see Figure 27-17), you can define another host. As we will use just this host server, we do not want to define another one at this time. Click **No**.

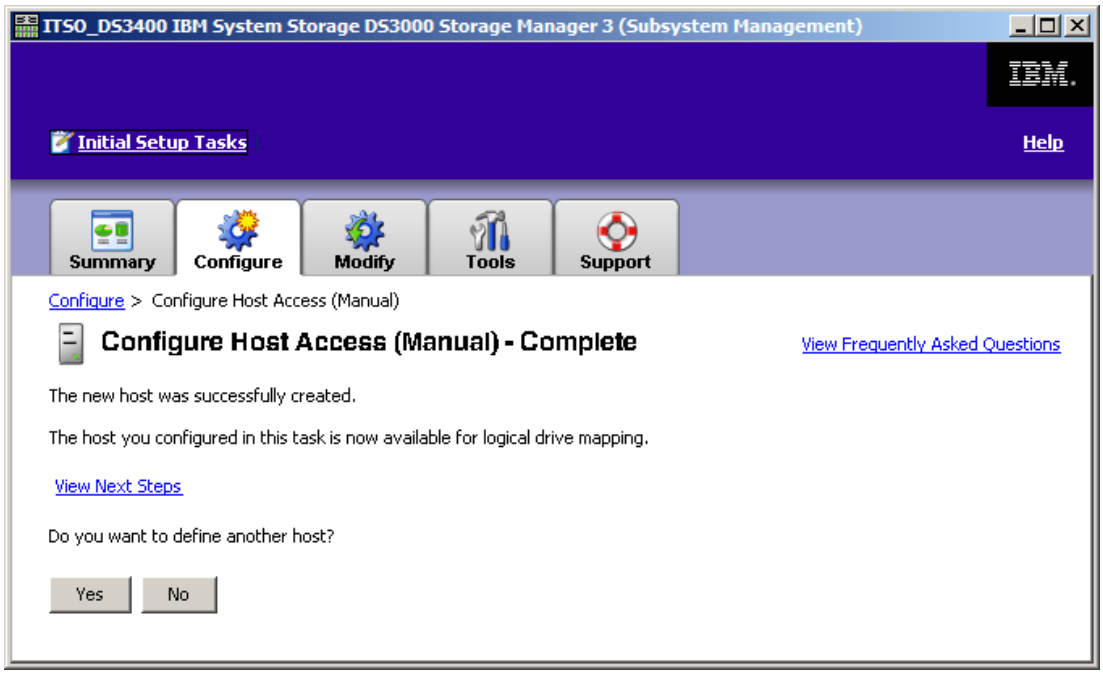


Figure 27-17 Host definition is complete

27.5.2 Creating a logical drive

Follow these steps:

1. To define the logical drive, select the **Configure** tab and click **Create Logical Drives**.
2. Select the capacity type, as shown in Figure 27-18 on page 687.

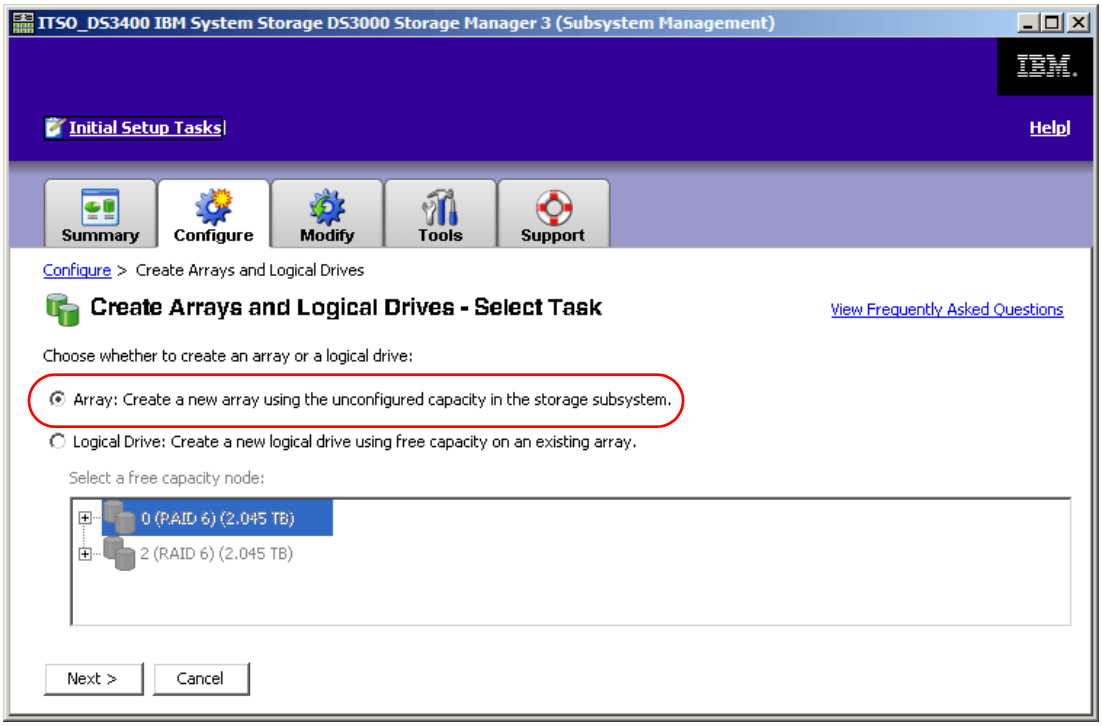


Figure 27-18 Create Arrays and Logical Drives

3. We will use two unassigned disk drives to create a new array. Therefore, we select **Unconfigured capacity: create a new array and logical drive**.
4. In Figure 27-19, specify the array name, in our example FC_Boot_Array. There are two drive selection choices. We select **Automatic**.

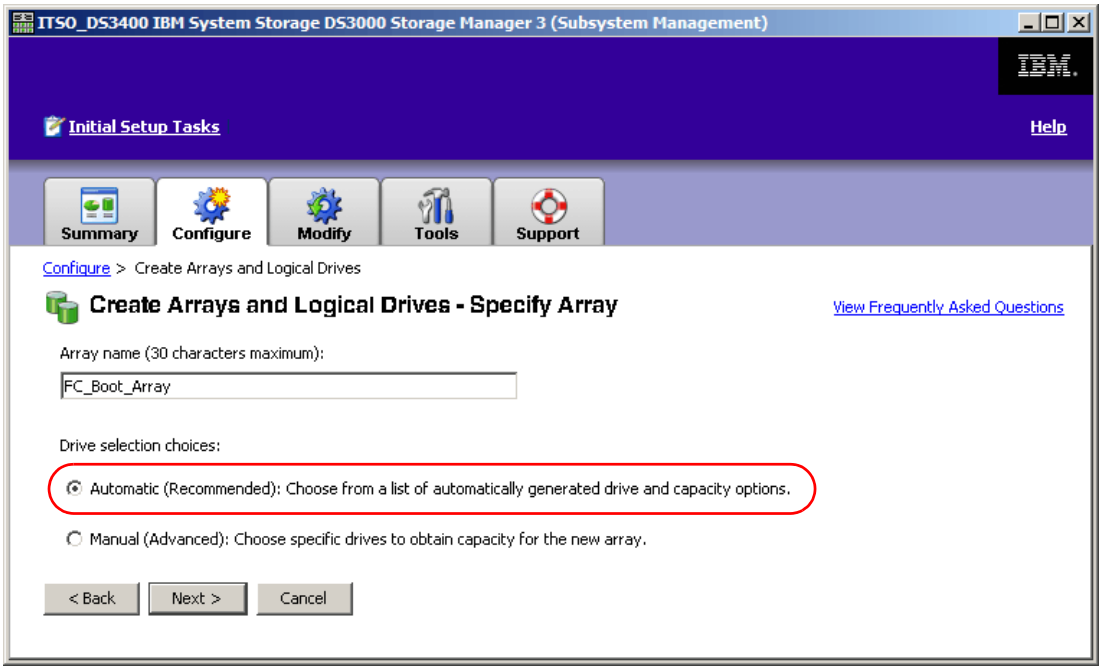


Figure 27-19 Drive selection choices

5. In the Array Drive Selection window, we specify the RAID level and then select one of the presented capacity options. We need to select RAID level 1; at this point, the window displays three entries, as shown in Figure 27-20. We select the first entry with two drives. Highlight this line and click **Finish**.

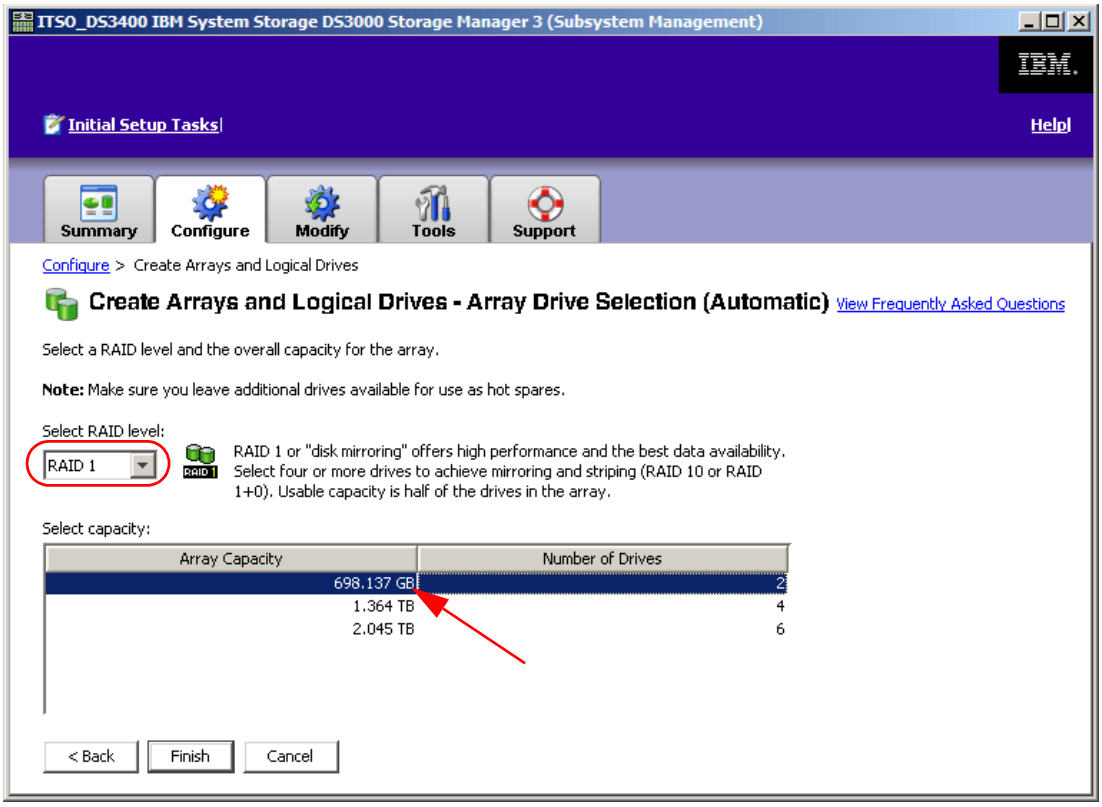


Figure 27-20 Array Drive Selection (Automatic)

6. A window appears, as shown in Figure 27-21, confirming that the array was completed successfully. We now continue creating the boot volume. Select the **Create a logical drive using the new array** radio button. Click **Yes** to continue.

.

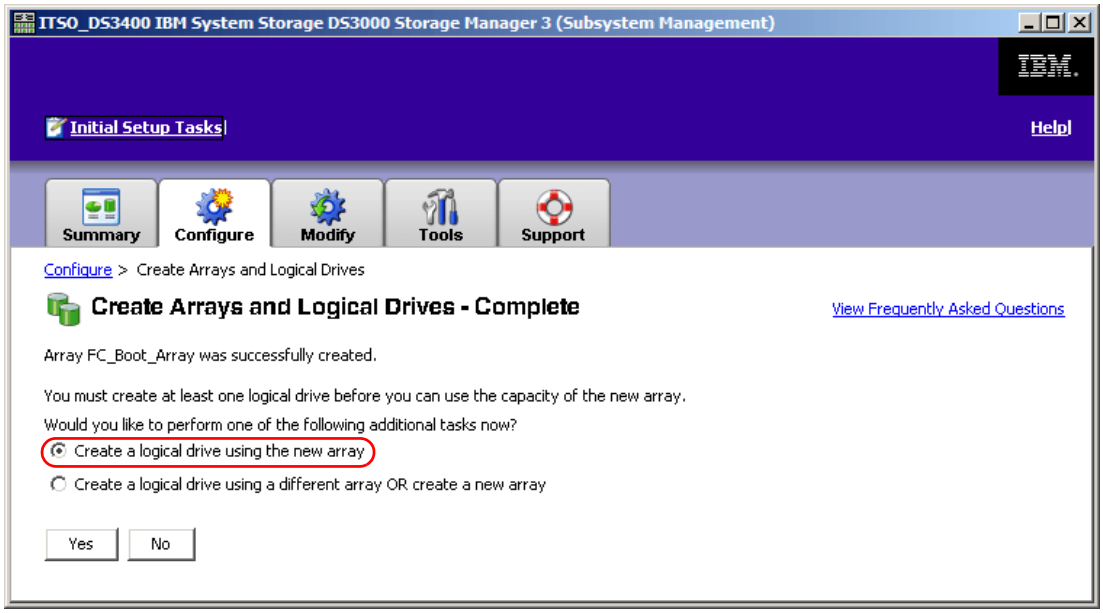


Figure 27-21 Specify Logical Drive

7. Now we need to specify the logical drive parameters: capacity, name, and I/O characteristics. We set 20 GB for capacity, the name will be SAN_Boot_Win2008, and the I/O characteristics will be left at File system (typical) (see Figure 27-22).

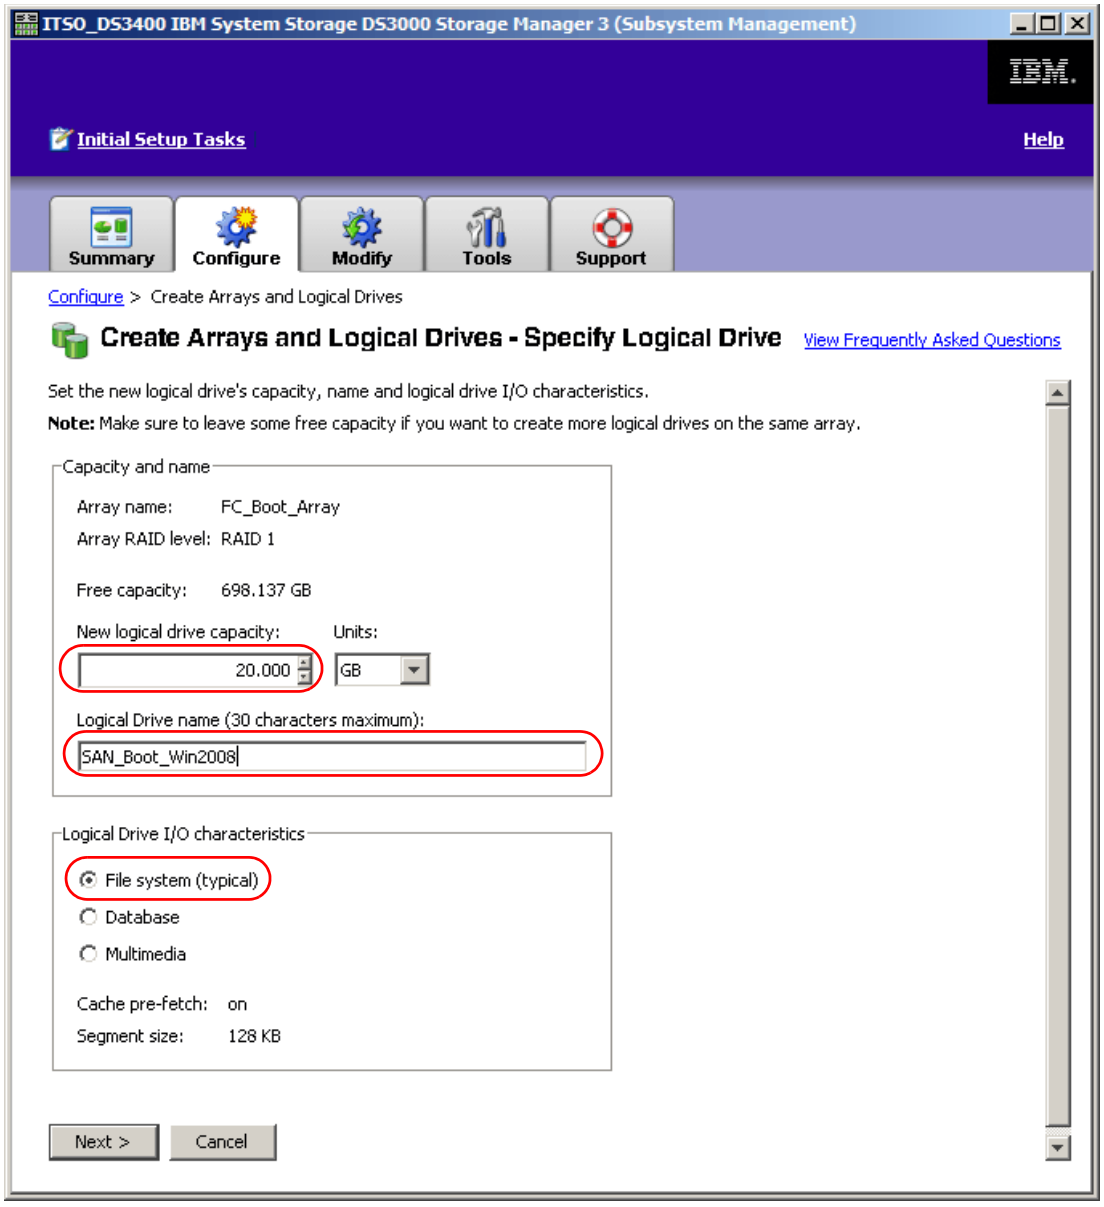


Figure 27-22 Specify Logical Drive

8. Now we map the host to the logical drive. Select the newly defined host SAN_Boot_Linux and assign the logical drive as LUN 0, as shown in Figure 27-23. Click **Finish** to continue.

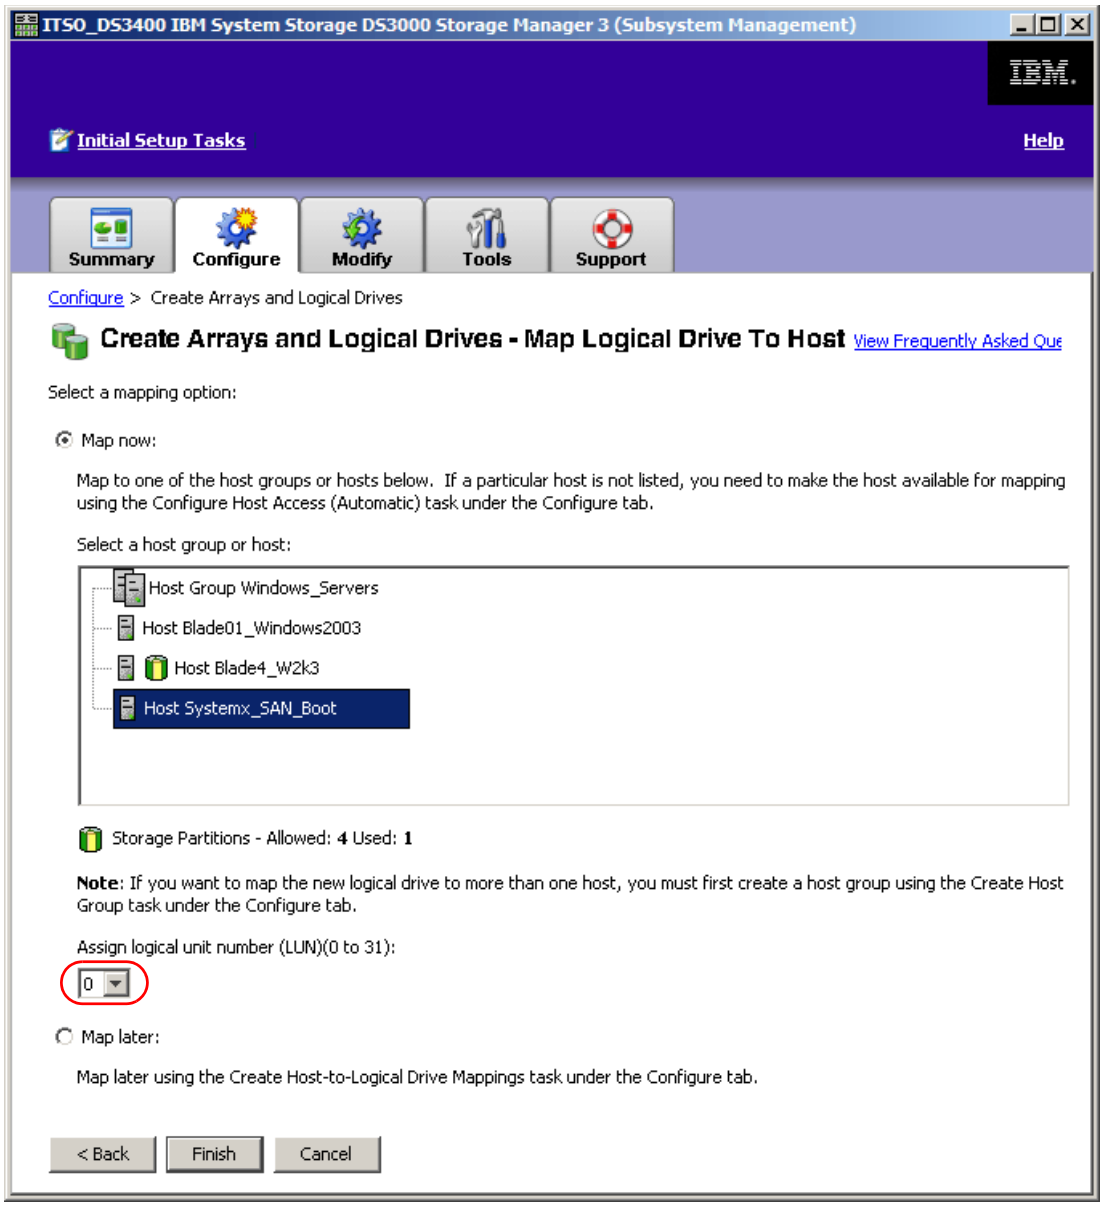


Figure 27-23 Map Logical Drive To Host

9. The logical drive is created and assigned to our host. The completion window displays (see Figure 27-24) and asks whether we want to create another logical drive. Click **No** to finish this section.

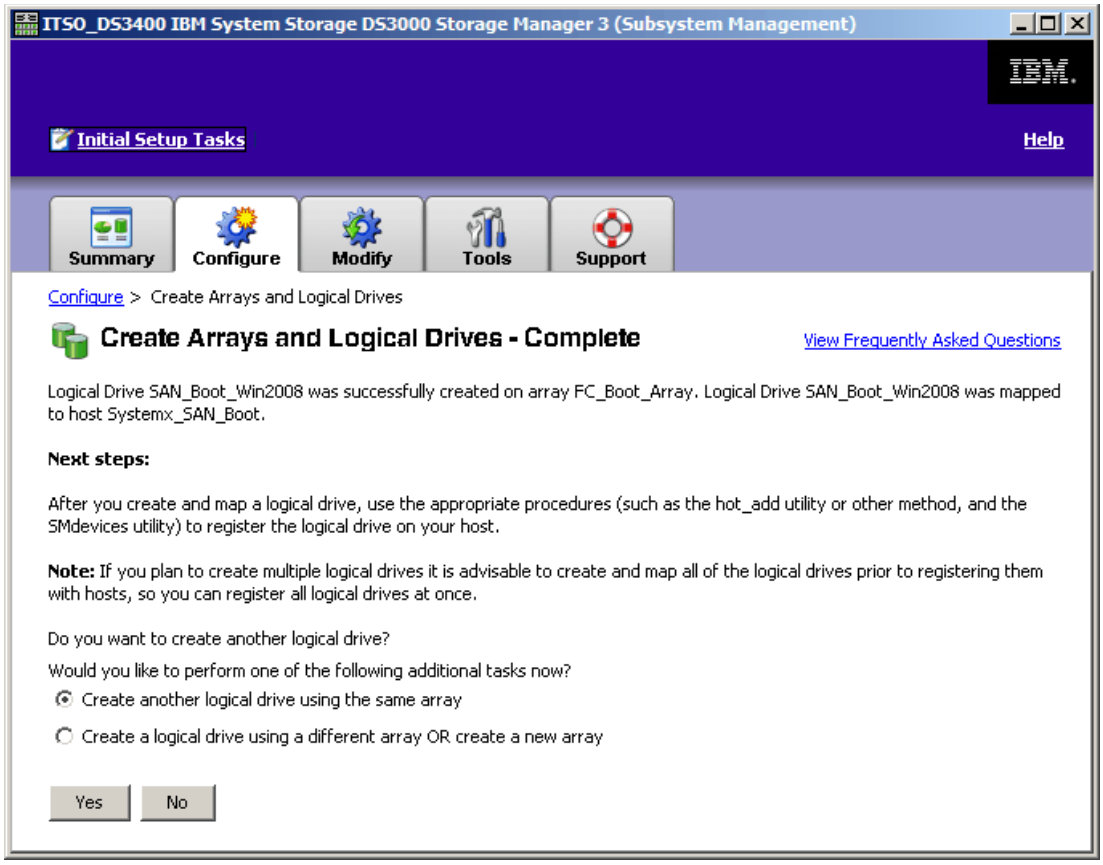


Figure 27-24 Logical drive creation complete

27.5.3 Logical drive ownership

The IBM System x host server can only talk to controller A at this point. This means controller A must be the preferred and current owner of the logical drive named SAN_Boot_Win2008. To verify this, use the **Change Logical Drive Ownership/Preferred Path** task in the Modify tab. As shown in Figure 27-25 on page 693, our logical drive is currently owned by controller A, which is also the preferred controller.

Should it be owned by controller B, click **Change** to reset it to controller A.

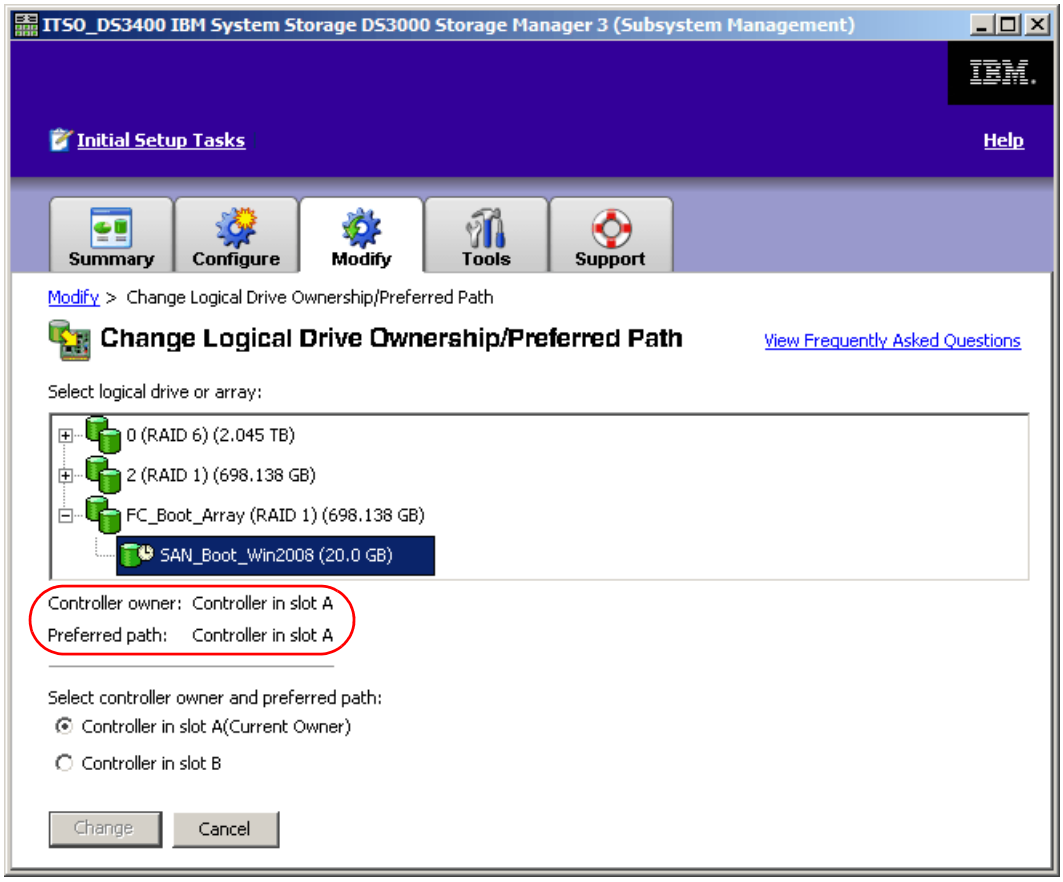


Figure 27-25 Change Logical Drive Ownership/Preferred Path

You are then asked to confirm the ownership and preferred path change (see Figure 27-26). Click **Yes**.

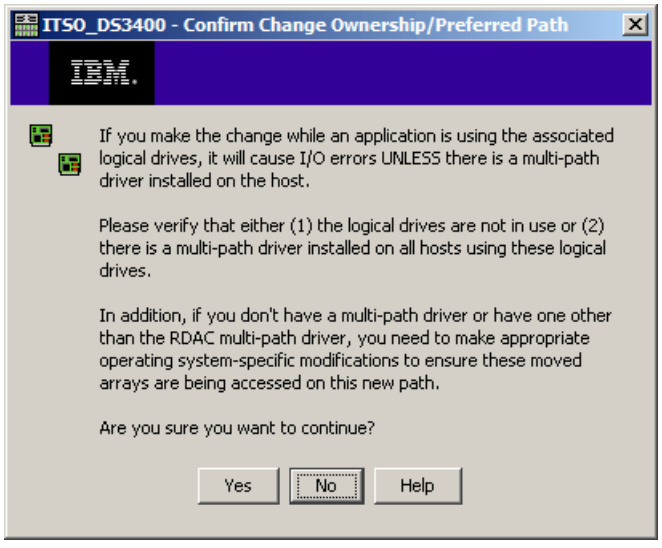


Figure 27-26 Confirm Change Ownership/Preferred Path

The task starts on the DS3400, as shown in Figure 27-27. Click **OK** to return to the subsystem management window.

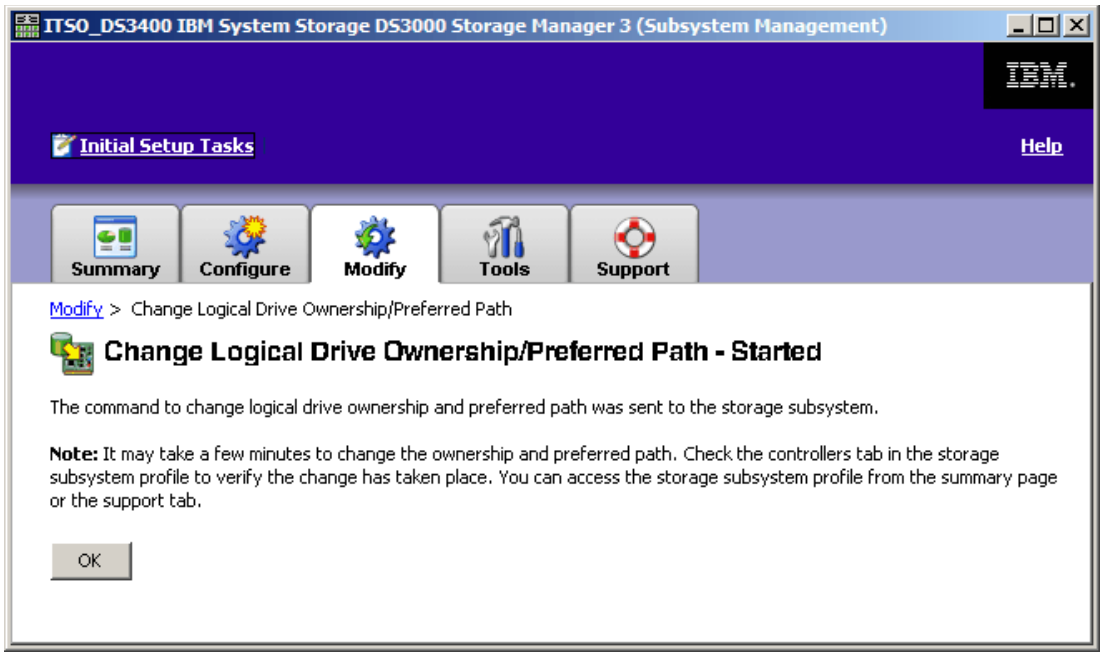


Figure 27-27 Change Logical Drive Ownership/Preferred Path - Started

We now have a logical drive available and assigned to our host server. The logical drive is owned by controller A, so we can proceed to the next section.

27.6 Configuring a logical drive as a boot device

We will now use QLogic HBA BIOS Fast!UTIL to specify that our logical drive is the bootable device. If you are not already in QLogic HBA BIOS Fast!UTIL, reboot the host server and press Ctrl+Q when prompted to launch Fast!UTIL.

1. Select the first FC HBA adapter and then **Configuration Settings** → **Selectable Boot Settings**. This is shown in Figure 27-28.

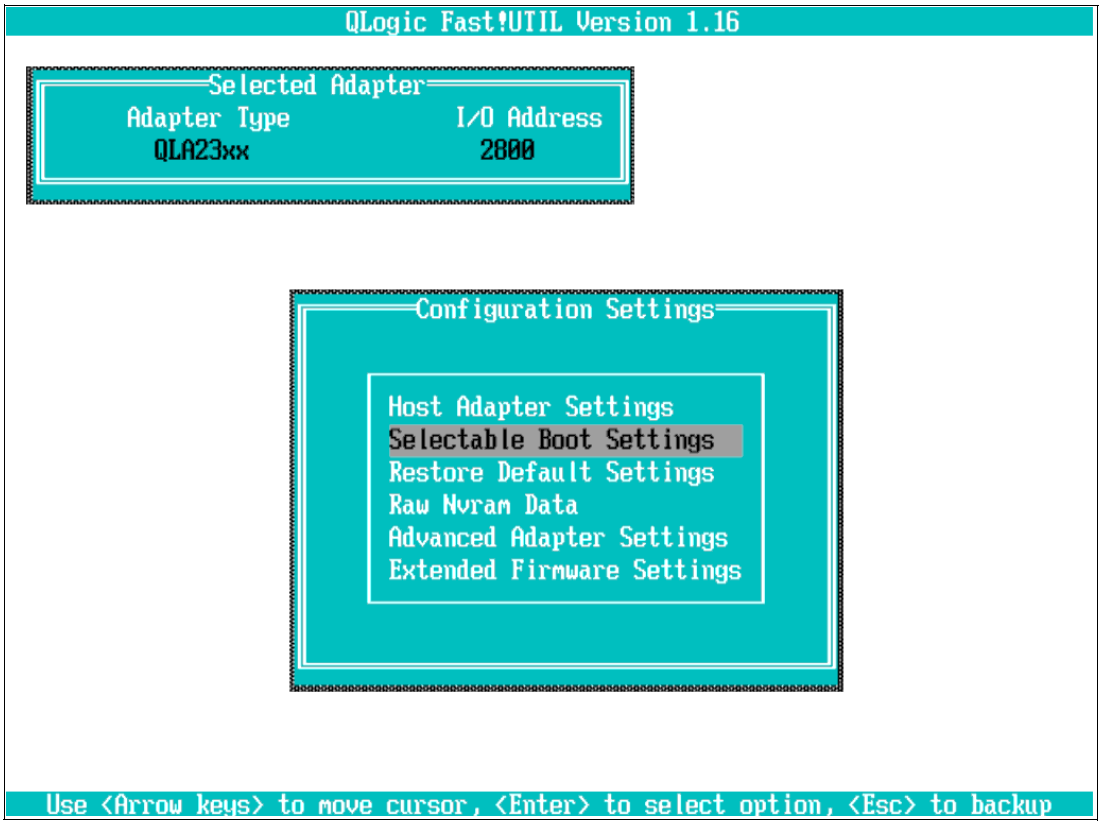


Figure 27-28 Selectable Boot Settings menu option

The Selectable Boot Settings menu displays. You need to do the following:

1. In the Selectable Boot Settings menu, make sure the Selectable Boot field is set to **Enabled**.

2. Move the cursor to (Primary) Boot Port Name, Lun field and press Enter. Now select the DS3400 entry as the boot device. The controller an FC host port WWN should appear and the LUN should be 0 (see Figure 27-29).

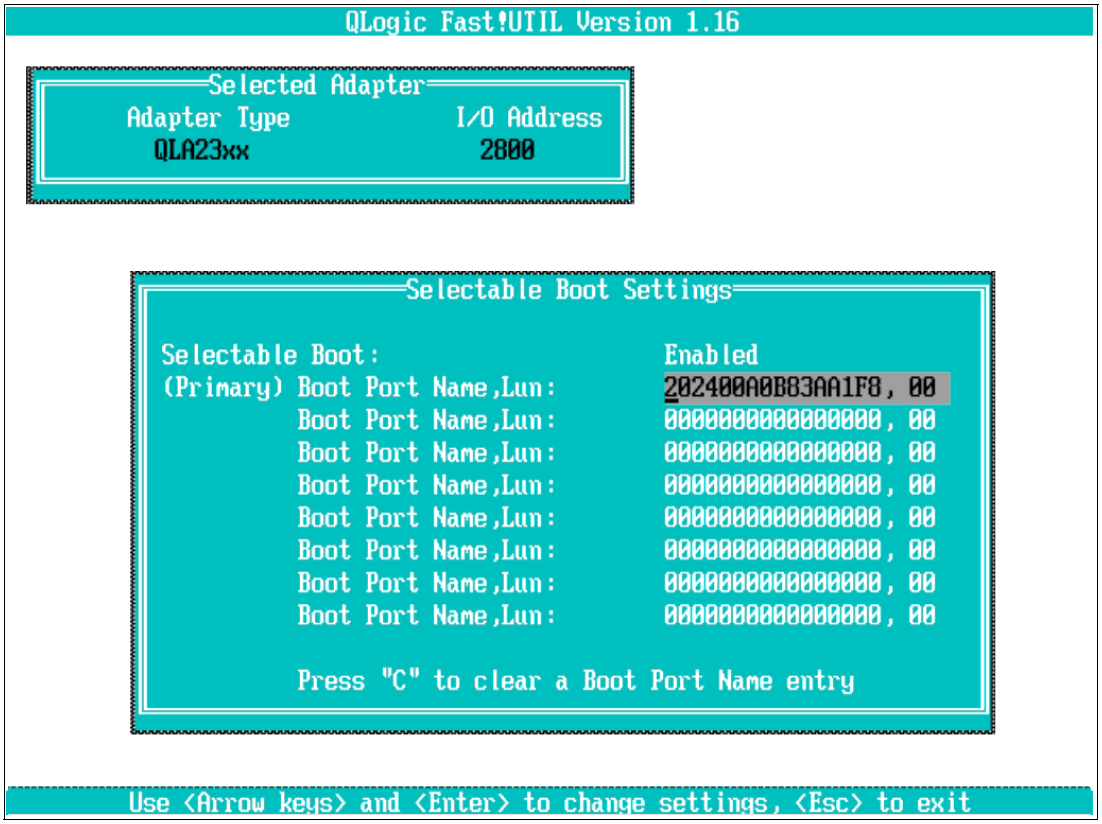


Figure 27-29 Selectable Boot Settings

3. Press Esc twice to save changes and exit Fast!UTIL.

27.7 Installing Windows Server 2008

We proceed with the Windows Server 2008 operating system installation steps for the boot logical volume created. Insert the Windows Server 2008 Media in the media tray and boot up the IBM System x server.

Specific installation instruction for your server can be found on the IBM support Web site. The installation guides can be found in the “Install/use” section of each product’s support Web sites.

Do these steps:

1. When the installation starts you will be presented with the window shown in Figure 27-30 on page 697. You can select the disk to be used for the installation here. In our example, the 20 GB boot-LUN created in the DS3400 System Storage subsystem is listed. You have the option to load device drivers for external devices.

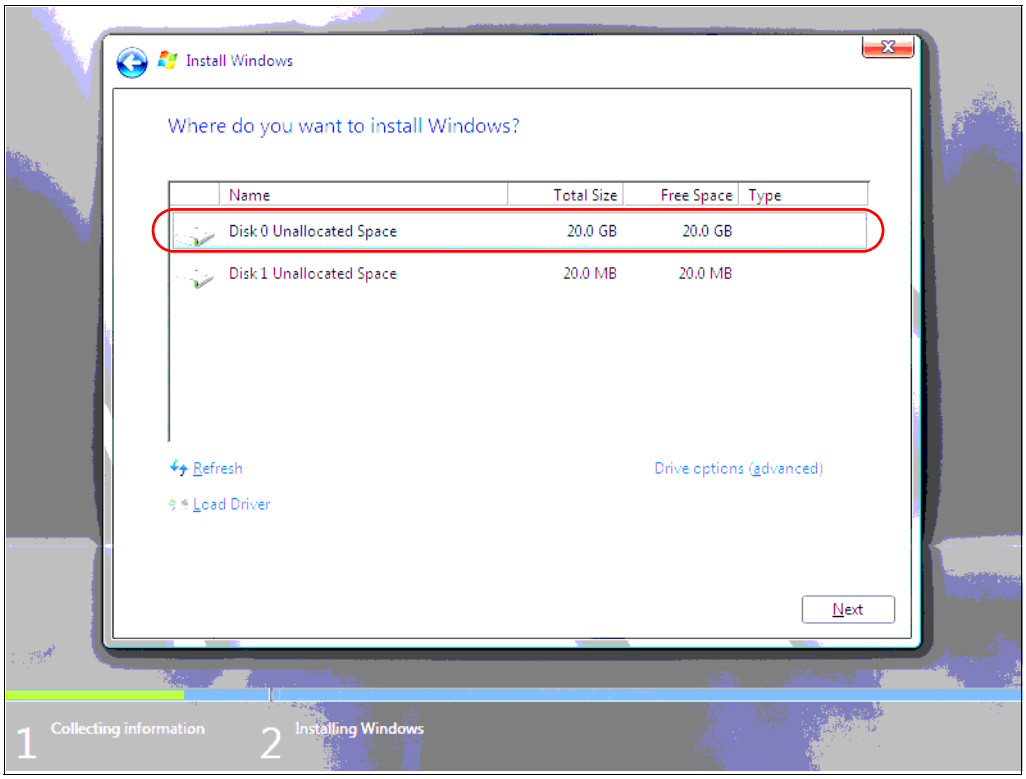


Figure 27-30 Installation disk selection

2. The system will restart after the Windows Server 2008 installation is complete. Once Windows boots up, install the Multi-Path failover driver that is part of the DS3000 Storage Manager Host software package, as described in 27.7.1, “Installing DS3000 Storage Manager host software” on page 697.

27.7.1 Installing DS3000 Storage Manager host software

We install the host server components of DS3000 Storage Manager. These components include:

- SMagent
- SMutil
- Multipath support (MPIO)
- SMclient (optional)

We install SMutil, because this will allow us to run the hot_add utility and avoid the necessity to reboot the host server when adding new logical drives.

Since the host server contains two Fibre Channel HBAs, we need multipath support. This is installed as part of the host selection when installing the DS3000 Storage Manager.

For the exact installation steps, see Chapter 6, “Installing IBM System Storage DS3000 Storage Manager” on page 81.

27.8 Enabling the second I/O path

Enabling the second path provides the capability to recover from a path failure at initial boot of the server from SAN as well as path failures during normal operation.

27.8.1 Zoning configuration for second path

Up to now, zoning was done to allow only the first HBA adapter to connect to controller A of the DS3400 System Storage subsystem.

We now zone the fabric to include a zone that allows the second HBA adapter to connect to controller B of the DS3400 System Storage subsystem that we are going to boot from, as shown in Figure 27-31. This will allow for redundant paths from the System x server to the DS3400 System Storage subsystem.

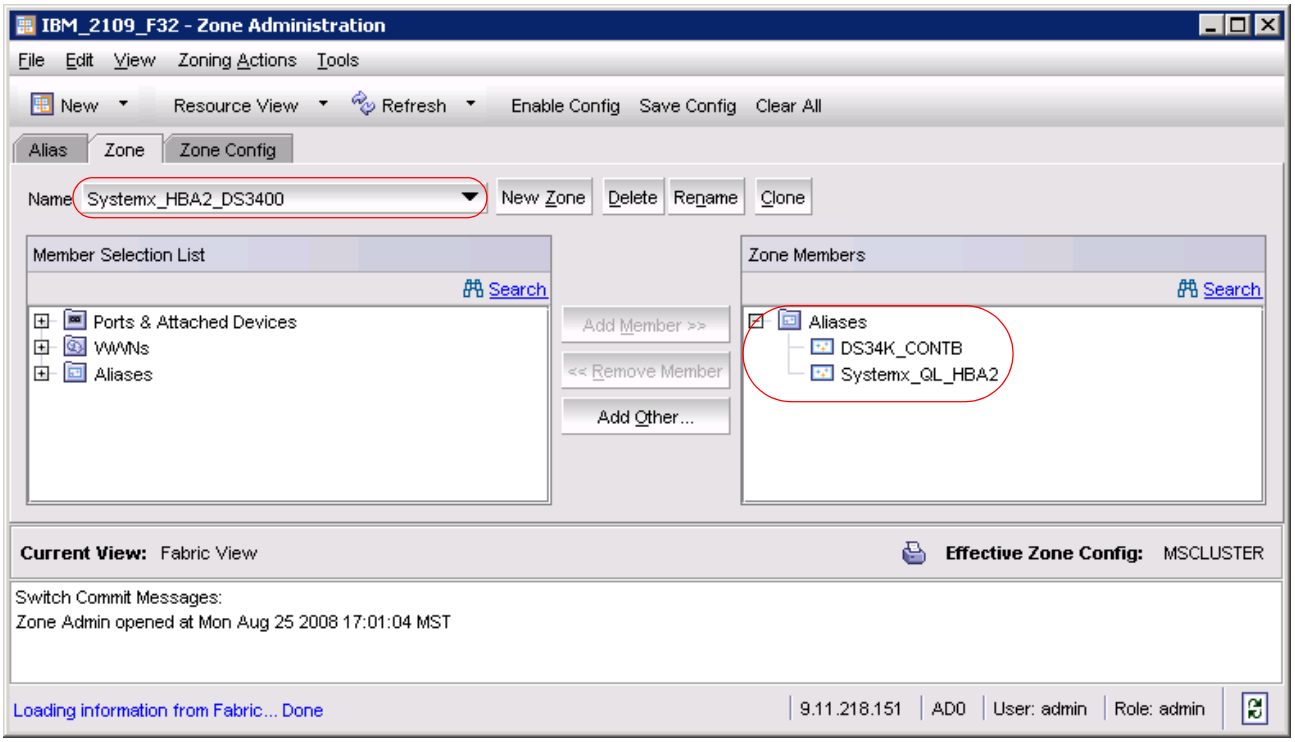


Figure 27-31 Zoning configuration for second path

27.8.2 Adding a second FC HBA port to the DS3400 host server definition

Our host server is still defined with one host port only. Now we can add the second host port, which will communicate with controller B, by doing these steps:

1. In the Storage Manager Modify tab, click **Edit Topology**. Click **Add HBA** to specify the new host port, as shown in Figure 27-32 on page 699.

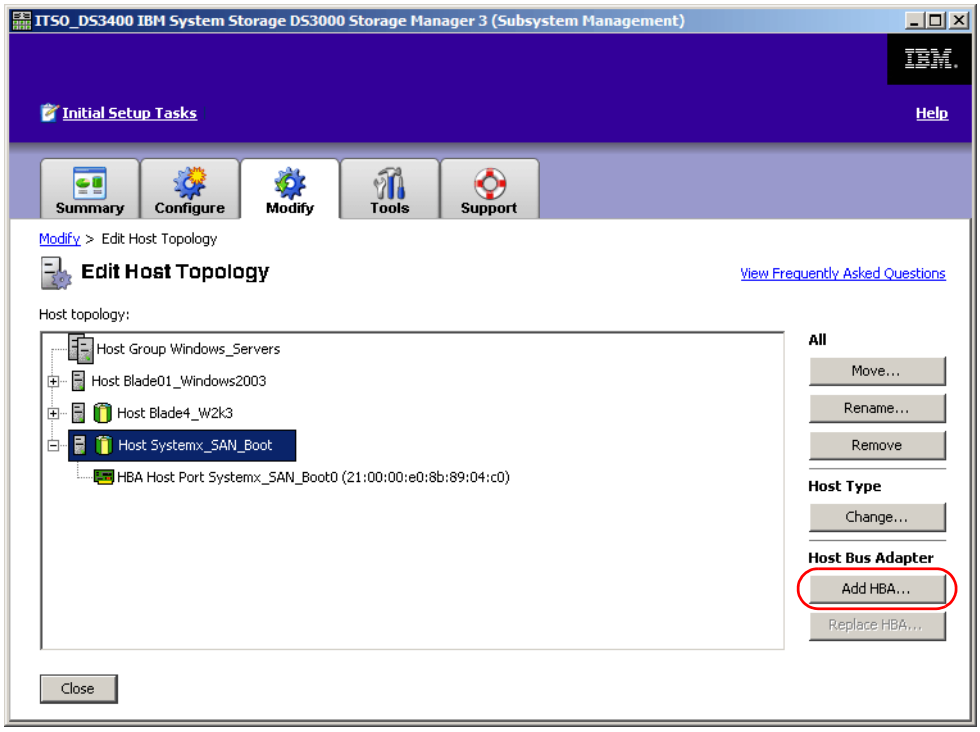


Figure 27-32 Edit Host Topology - Add HBA

2. This launches the Add HBA Host Port window, as shown in Figure 27-33.



Figure 27-33 Add HBA Host Port

3. Set the correct values for WWN, host type, and host port alias, and then click **Add**. As you can see in Figure 27-34, our host server now contains both host port definitions.

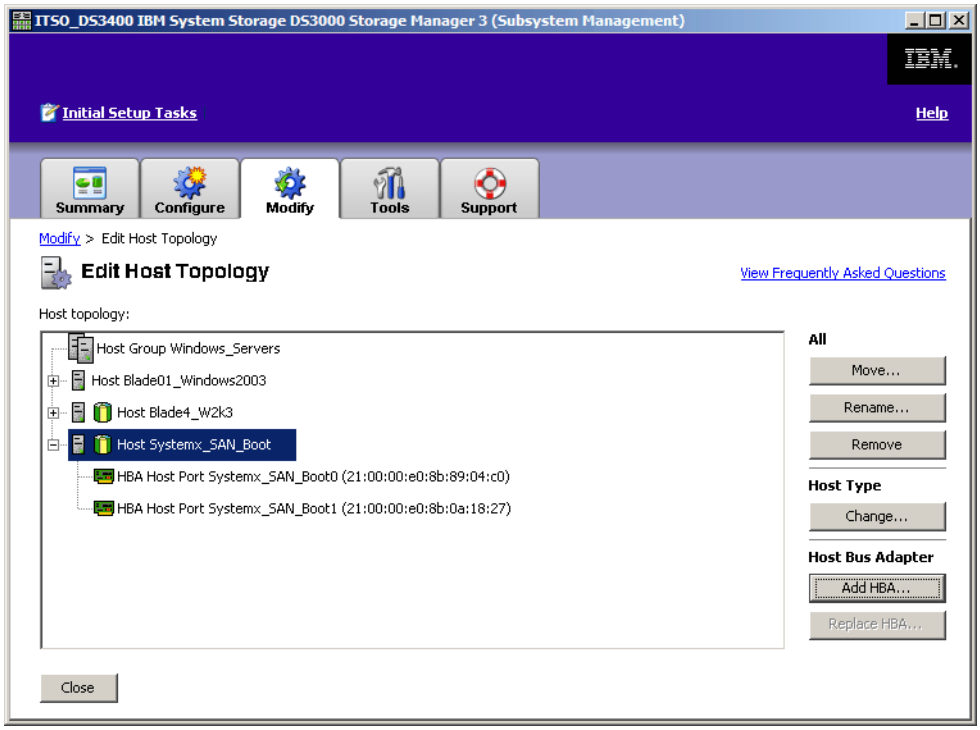


Figure 27-34 Both host ports are defined

27.8.3 Second HBA adapter configuration

Do these steps:

1. Reboot the IBM System x server. Enter Fast!UTIL by pressing Ctrl+Q when prompted. Select the second HBA adapter and press Enter. The Fast!UTIL Options menu displays (Figure 27-35). Select **Configuration Settings** and press Enter.

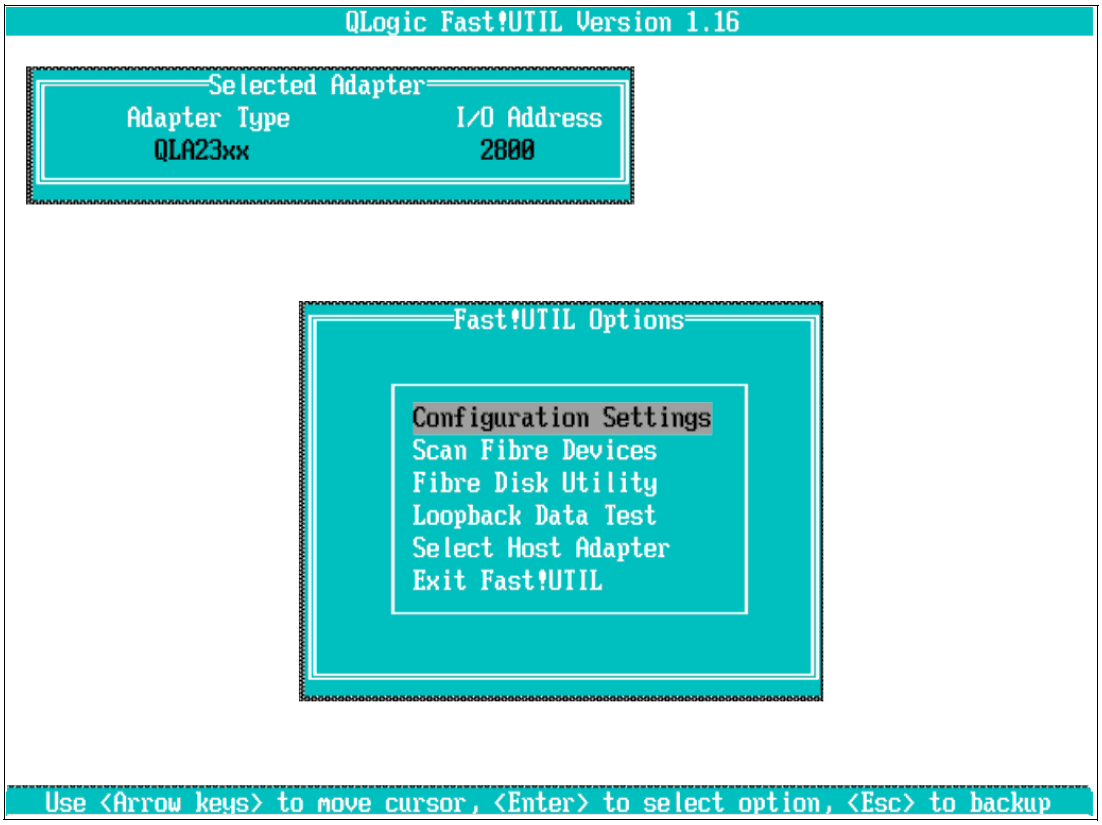


Figure 27-35 Fast!UTIL Options for second HBA adapter

2. In the Configuration Settings menu, shown in Figure 27-36, select **Adapter Settings**.

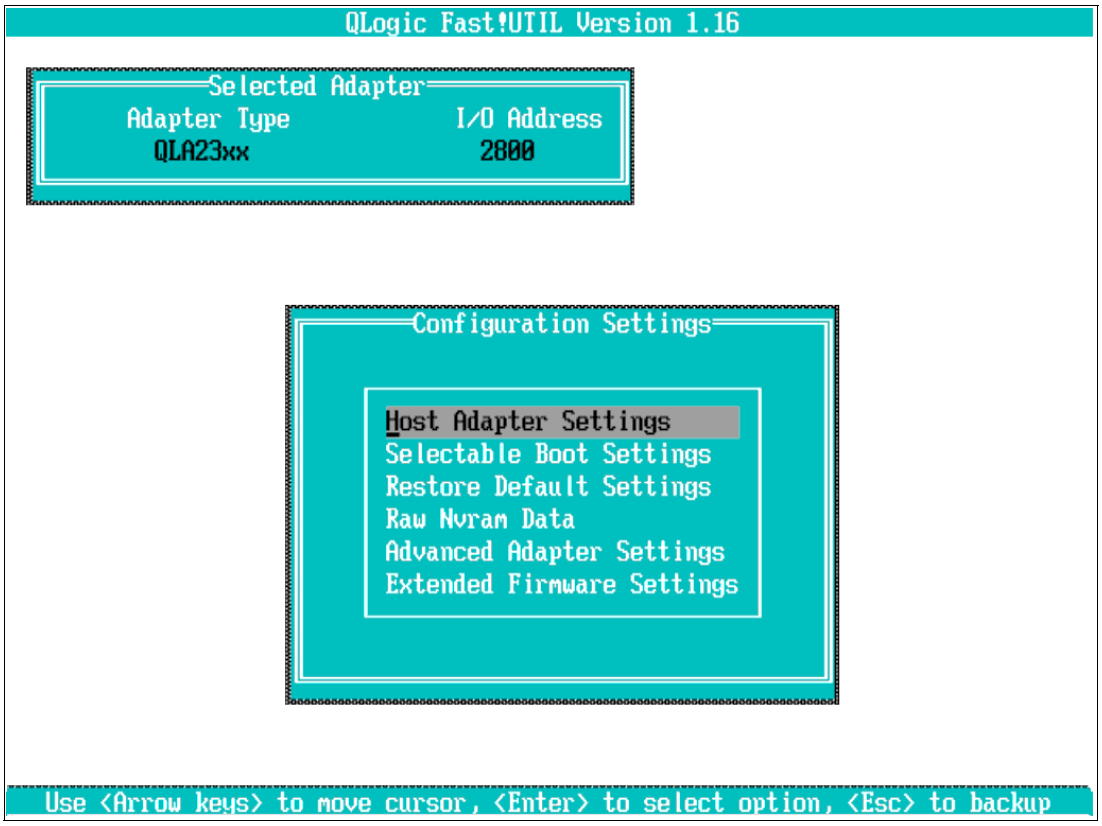


Figure 27-36 Configuration Settings menu for second HBA adapter

3. Enable the BIOS for this HBA adapter (see Figure 27-37). Remember, the BIOS must be enabled if we want to boot through this HBA adapter.

   Change the Host Adapter BIOS setting from **Disabled** to **Enabled**, then press Esc. Select **Save changes** to confirm.

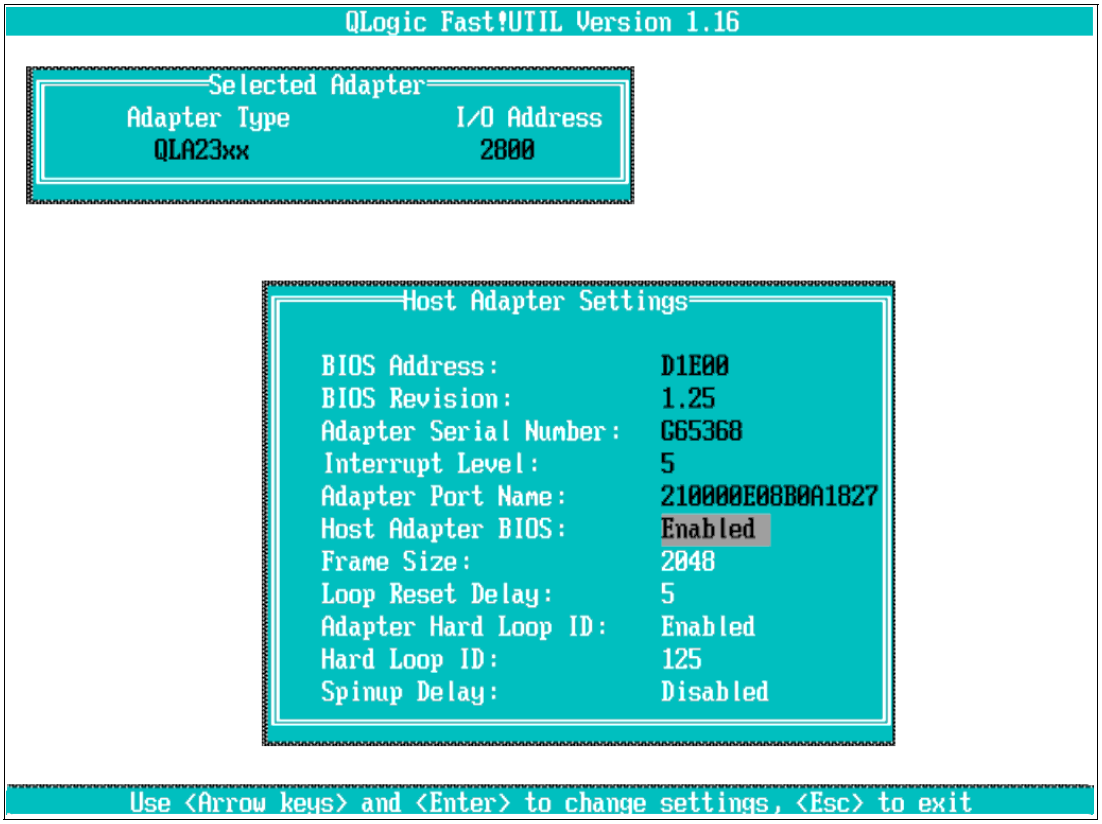


Figure 27-37 Adapter settings for second HBA adapter

4. Exit this menu and **Selectable Boot Settings** from the main Configuration Settings window, as shown in Figure 27-38.

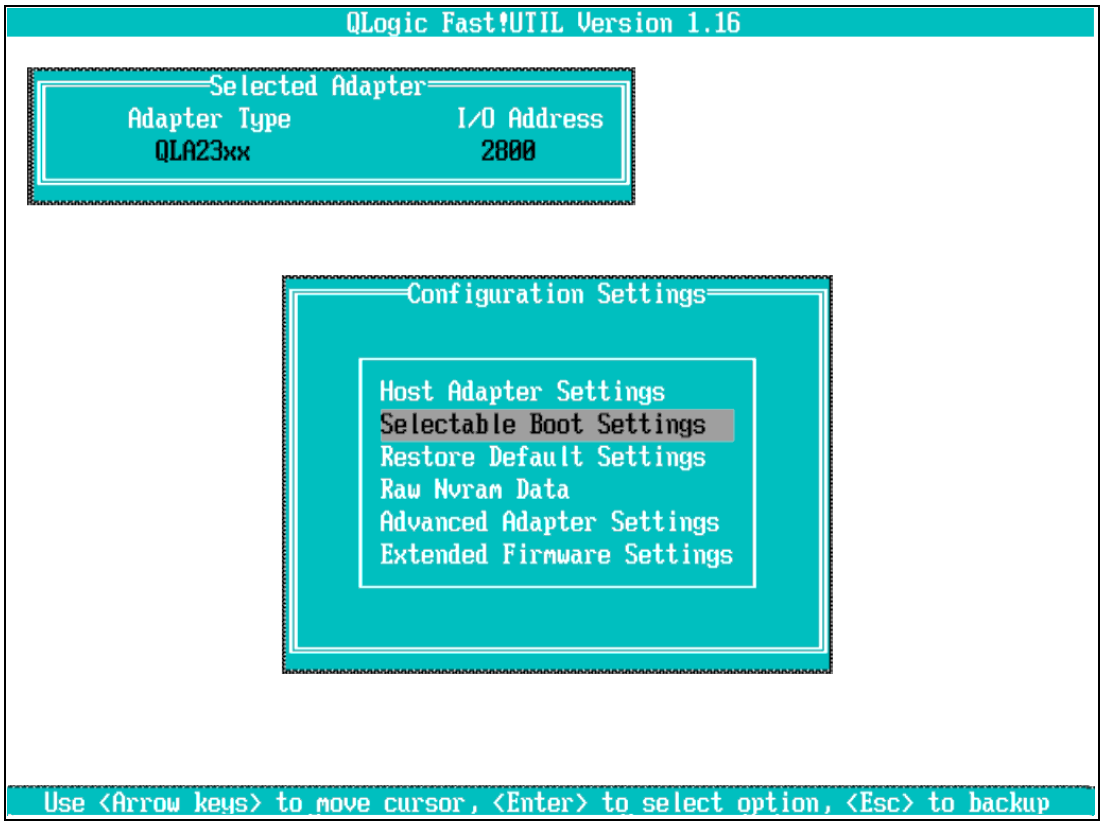


Figure 27-38 Selectable Boot Settings menu option for second HBA

5. The Selectable Boot Settings menu displays. You need to do the following:
   a. In the Selectable Boot Settings menu, make sure the Selectable Boot field is set to **Enabled**.
   b. Move the cursor to (Primary) Boot Port Name, Lun field and press Enter. Now select the DS3400 Controller B entry as the boot device. The Controller B FC host port WWN should appear and the LUN should be 0 (see Figure 27-39 on page 705).

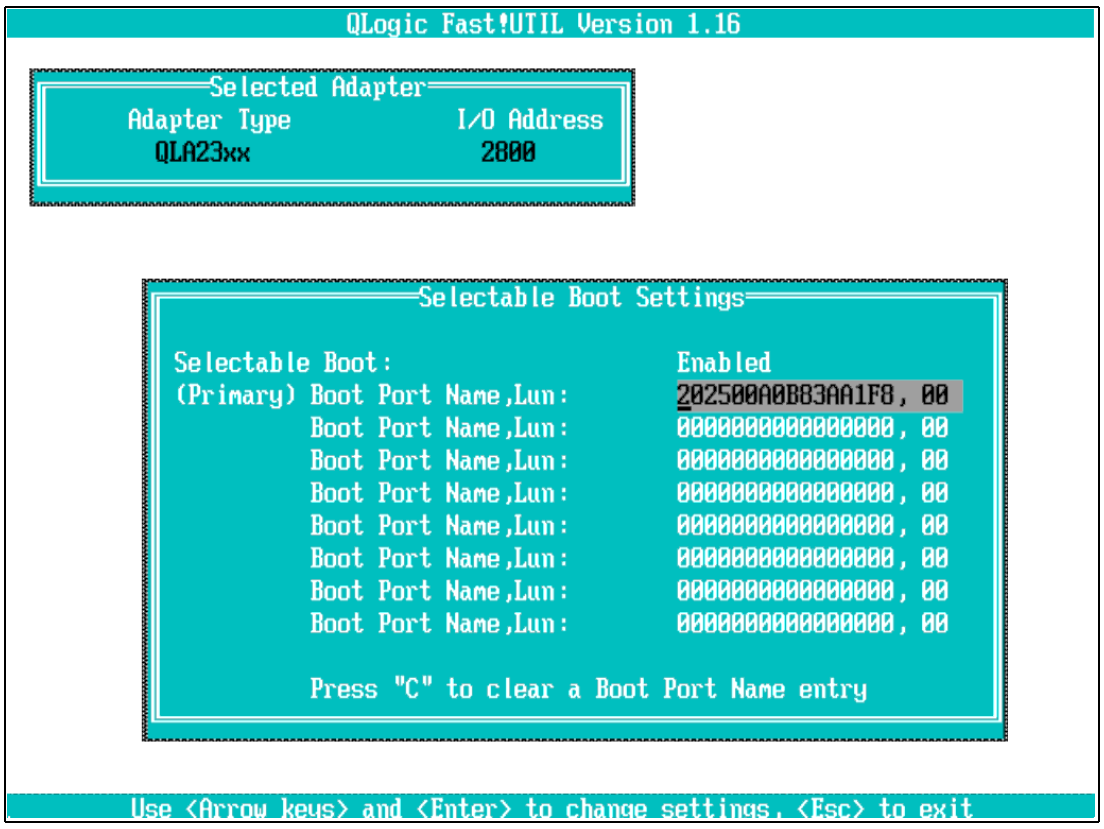


Figure 27-39 Selectable Boot Settings for second HBA

6. Press Esc twice to save changes and exit Fast!UTIL.

27.8.4 Verifying Multi-path in Windows

Do these steps:

1. Reboot the System x host to ensure that the Multi-path driver is loaded and the second path is activated.

2. Once Windows Server 2008 is booted up, verify in Windows 2008 Disk Management that only a single disk is seen as the boot volume, in our example, a 20 GB disk, as shown in Figure 27-40. This confirms that the Multi-Path driver is installed successfully.

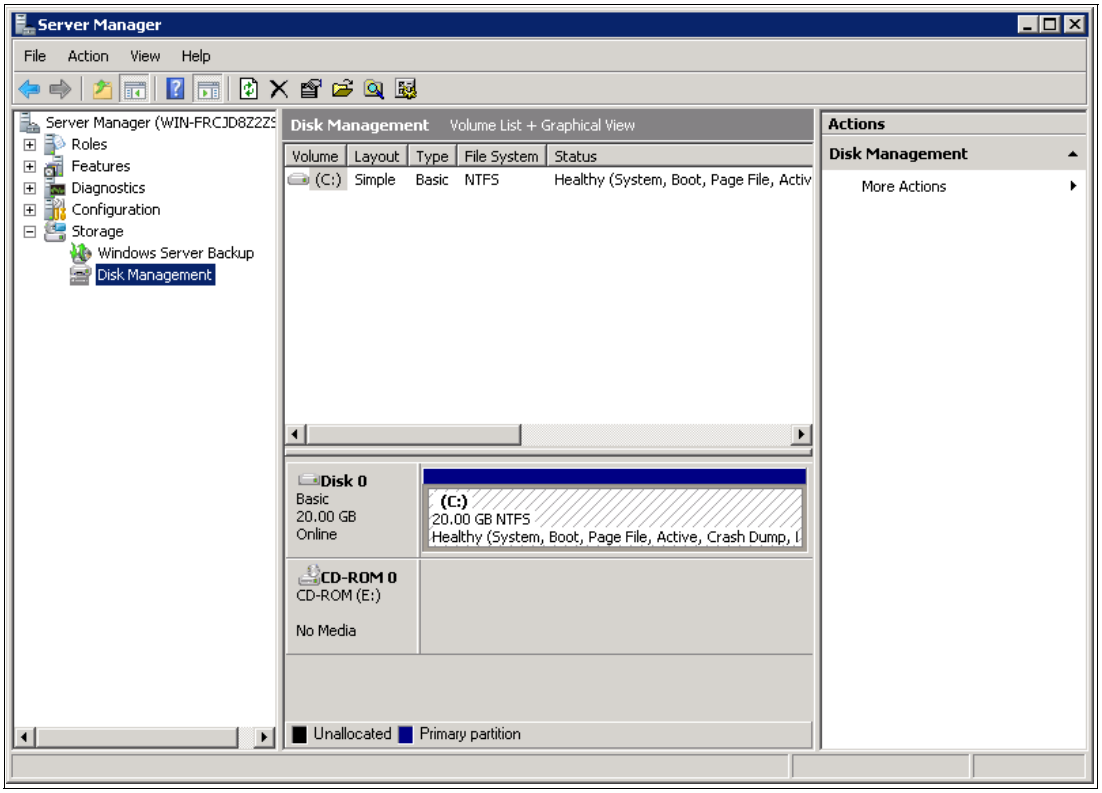


Figure 27-40 Confirming single boot volume

3. In the Windows Device Manager, we can now also confirm the presence of the IBM 1726-4xx Multi-Path disk device as well as the Multi-Path Bus Driver, as shown in Figure 27-41 on page 707.

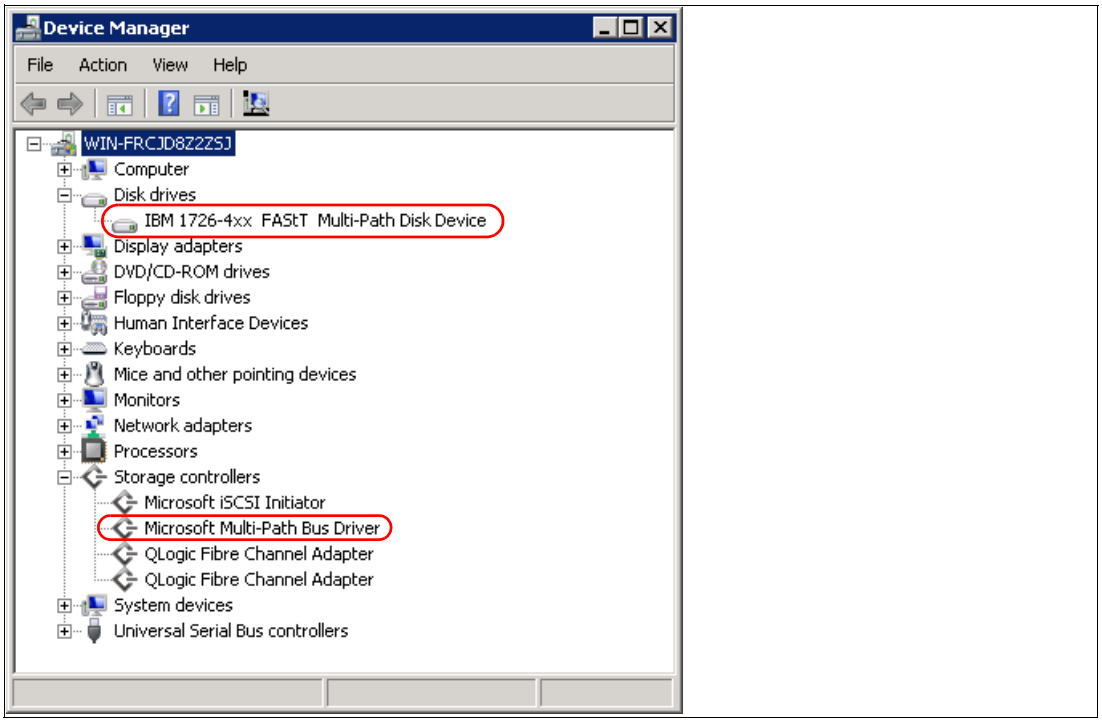


Figure 27-41 Confirming Multi-Path driver

27.8.5 Failover testing

Do the following steps to test for failover recovery:

1. Verify the current owner of the boot volume, in our example, Controller A, by viewing the logical drives section under the Subsystem Profile in the DS3400 Storage Subsystem Manager, as shown in Figure 27-42. Do this by selecting **Storage Manager** → **Subsystem Profile** → **Logical Drives**.

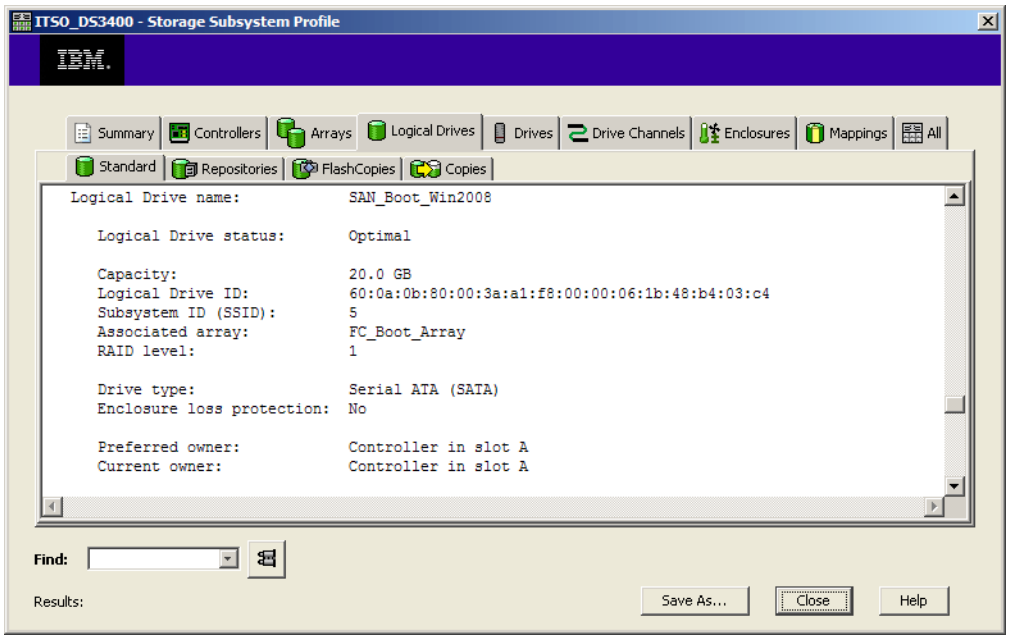


Figure 27-42 Boot volume owner - Controller A

2. We test the path failover by unplugging the cable connected from the first HBA adapter on the System x server to the SAN switch.
3. After the cable is unplugged, the Windows 2008 operating system will pause until the boot volume is failed over to the backup path, in our example, Controller B, as shown in Figure 27-43.

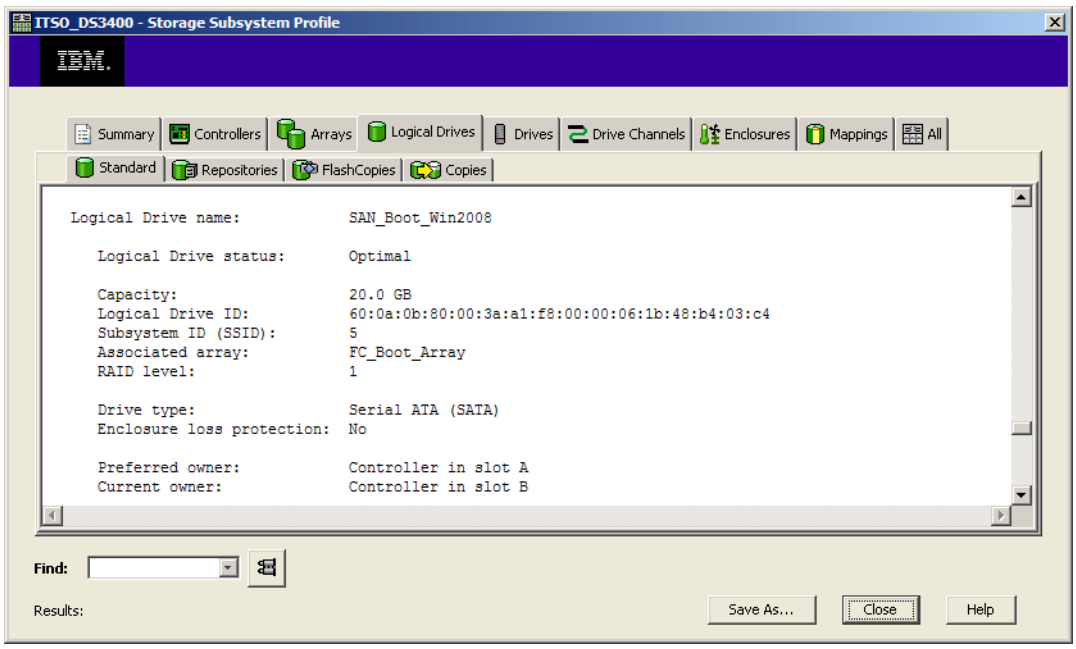


Figure 27-43 Boot volume owner - Controller B

A Recovery Guru alert will also be presented in the DS3400 Subsystem Management window, as shown in Figure 27-44, alerting you that the boot volume has moved ownership to Controller B and is not on the preferred path.

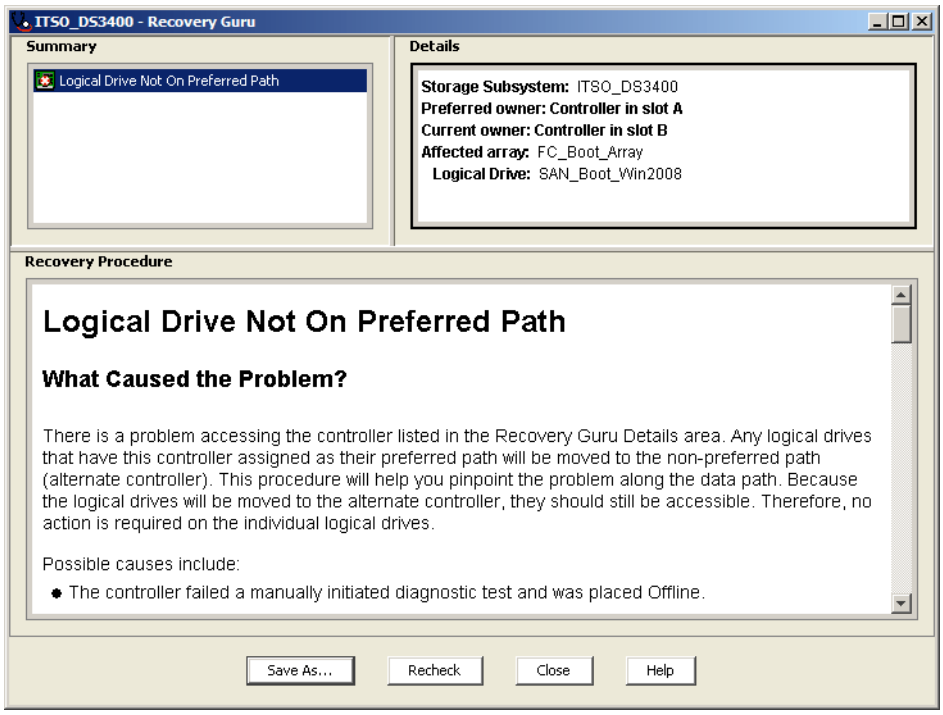


Figure 27-44 Recovery Guru alert

4. After verifying the successful failover, replug the cable between the first HBA adapter on the System x server and the SAN switch. The boot volume will now automatically move back to the preferred controller, in our example, Controller A on the DS3400 System Storage subsystem.

28

FC configuration 5 - Boot from SAN with VMware ESX Server

In this chapter, we discuss a sample configuration, describing how to boot an IBM System x running VMware ESX Server V3.5.0, from a DS3400 System Storage subsystem. Before you start, make sure that all the firmware on the storage subsystem and in the host is flashed to the latest level, as described in 14.3, “Download firmware” on page 307.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

28.1 What is required

For our setup, we used the following hardware and software components:

- IBM System x366 server
- Two PCI Express FC HBAs, P/N 24P0962 (QLA2340)
- SAN switch
- DS3400 System Storage subsystem with enough free disk space for the logical drive
- A management workstation running the latest version of DS3000 Storage Manager (at the time of writing, Version 10.35)
- VMware ESX Server V3.5.0

VMware ESX Server is a virtualization solution offered by VMware, Inc. ESX Server is a component of VMware Infrastructure and can be installed on “bare metal“.

28.1.1 Installing ESX Server on a DS3400 server - configuration procedure overview

In summary, the general steps required to set up an ESX Server SAN boot on a DS3400 server are:

1. Server HBA configuration:
   a. Power on the server.
   b. Disable the internal SCSI/SAS devices in the x366.
   c. Record WWPN and enable BIOS of the FC HBAs using the FastT!Util menu.
2. SAN switches configuration:

   Ensure that the Fibre Channel SAN switches are correctly set, according to the following guidelines:
   a. Complete the necessary fibre cable connections with the DS3400 controllers, SAN switches, and server’s HBAs, and ensure that all the devices are properly connected.
   b. Verify that all the switches in the fabric are configured with unique domain ID and IP addresses.
   c. Verify and confirm that all of the switches are running the latest and compatible firmware version.
   d. Define and activate zoning.
3. DS3400 storage configuration for the primary path:
   a. Create a logical drive to be used as the operating system disk for the server.
   b. Complete the host group, host, and host port definitions.
   c. Define the storage partition and map the logical drive with LUN ID 0 to the first host port in the host group.
4. Fibre Channel host configuration for primary path:

   Configure the boot device (FC BIOS parameter) for the primary path from the server to the DS3400 Controller A by identifying and selecting the boot device.

   Note: Verify that the displayed WWPN matches the WWPN of the DS3400 Controller A zoned with the first FC HBA and the LUN ID=0.

5. Operating system installation:
   a. Install the ESX Server software.
   b. Verify that the server successfully boots from the logical drive on the primary path.
6. Verify and test access to the secondary path from the host:
   a. Check whether both controllers are visible from the host.
   b. Test path redundancy.

28.1.2 Installing ESX Server SAN boot on a DS3400 - step-by-step procedure

Unlike other operating systems, the ESX Server loads its multipathing driver during the installation. See Figure 28-1 for an example configuration with two paths to the DS3400 with the appropriate zoning.

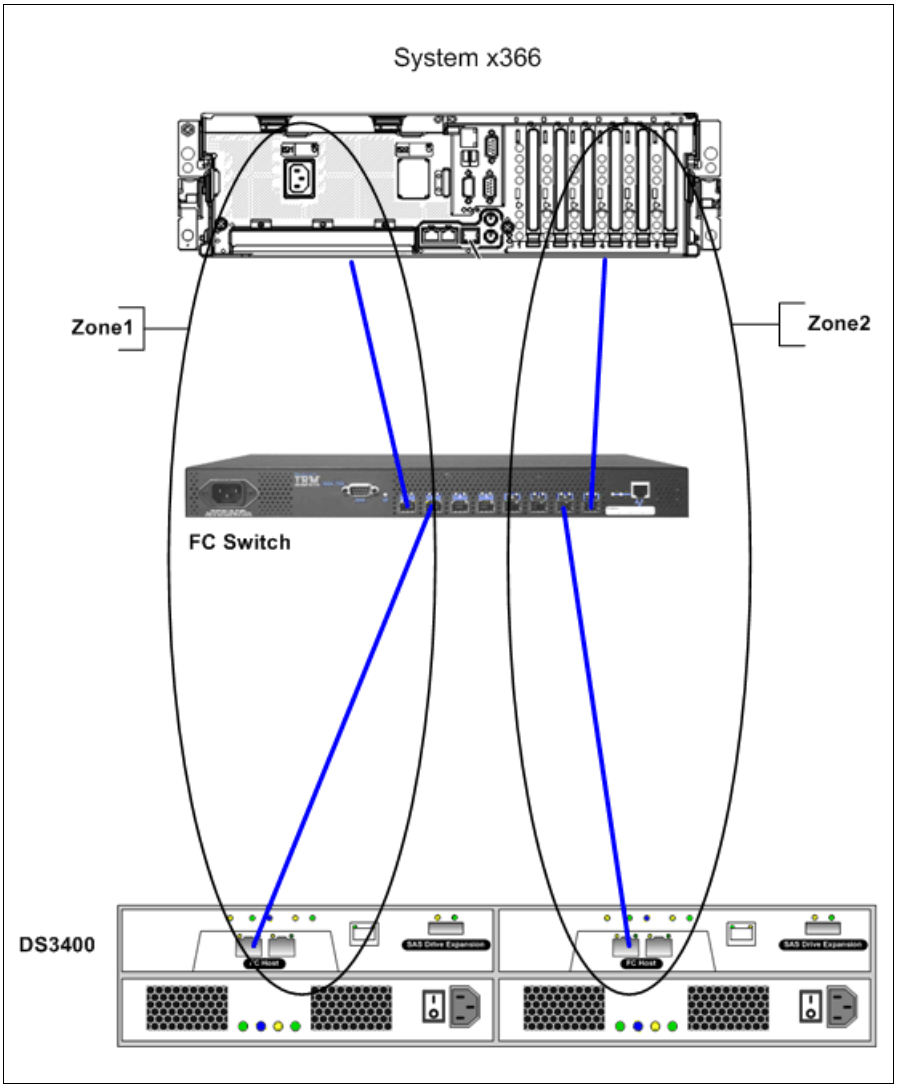


Figure 28-1 Zoning

Step 1 - System x BIOS configuration

For this step:

1. Power on the server and interrupt the boot sequence to enter the system BIOS by pressing the F1 key.
2. Select devices and I/O ports from the main menu.
3. Select **Planar SAS** from the next menu and disable the internal SAS controller, as shown in Figure 28-2.
4. Press Esc to exit out of this menu and save the changes.

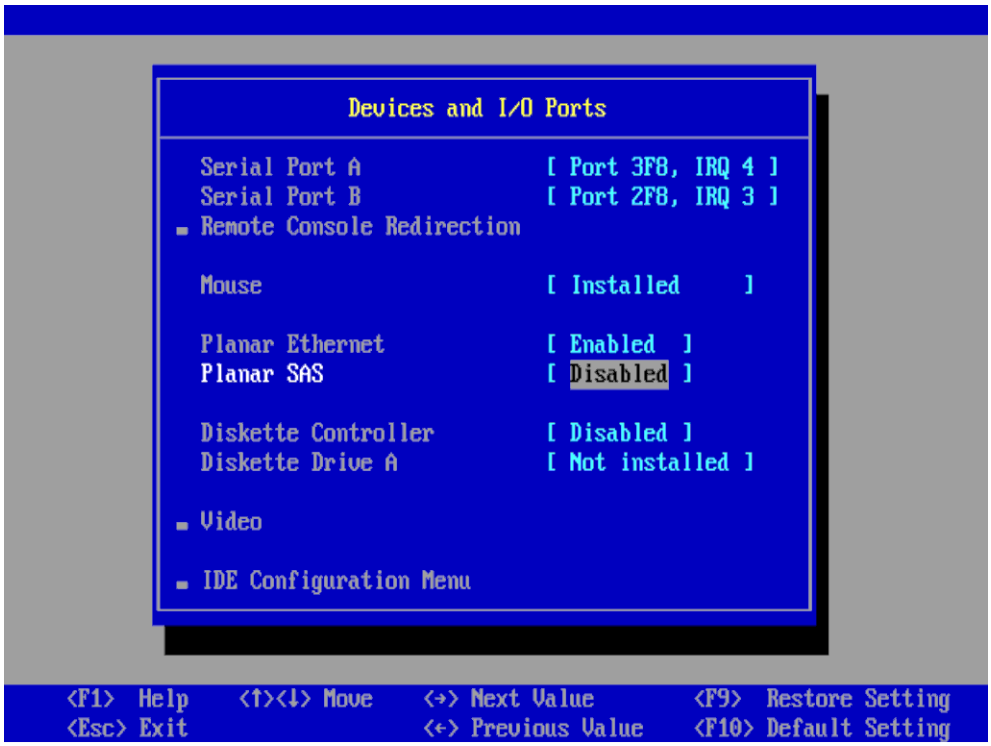


Figure 28-2 Disabling SAS controller

5. Power on or restart the server.

Step 2 - Record WWPN of HBAs and enable BIOS of HBA

Do the following to record the WWPN and to enable the HBAs BIOS of the first HBA:

1. Power on the server and press Ctrl+Q when prompted to enter the QLogic Fast!Util (Figure 28-3 on page 715).

```
Broadcom NetXtreme Ethernet Boot Agent v9.0.12
Copyright (C) 2000-2006 Broadcom Corporation
All rights reserved.


Broadcom NetXtreme Ethernet Boot Agent v9.0.12
Copyright (C) 2000-2006 Broadcom Corporation
All rights reserved.

QLogic Corporation
QLA2312 PCI Fibre Channel ROM BIOS Version 1.43      Subsystem Vendor ID 1014
Copyright (C) QLogic Corporation 1993-2004. All rights reserved.
www.qlogic.com

Press <CTRL-Q> for Fast!UTIL

BIOS for Adapter 0 is disabled

BIOS for Adapter 1 is disabled

<CTRL-Q> Detected, Initialization in progress, Please wait...

ROM BIOS NOT INSTALLED
_
```

Figure 28-3 Invoke FC HBA BIOS

2. In the QLogic Fast!Util, highlight the first HBA and press Enter, as shown in Figure 28-4.

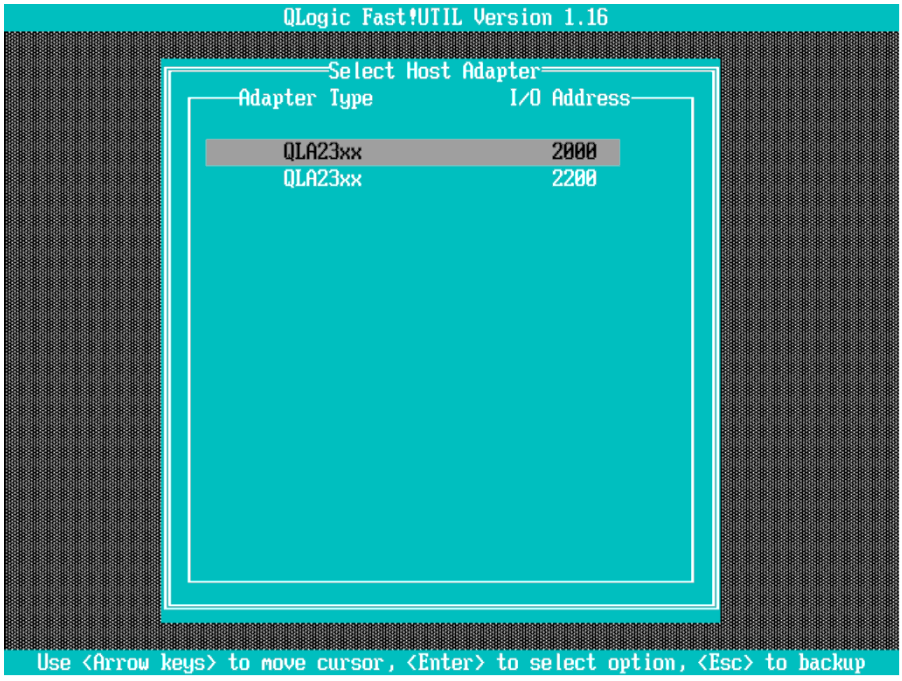


Figure 28-4 Select the first HBA

3. Select **Configuration Settings** and then **Host Adapter Settings** to bring up the Adapter Setting window, as shown in Figure 28-5.

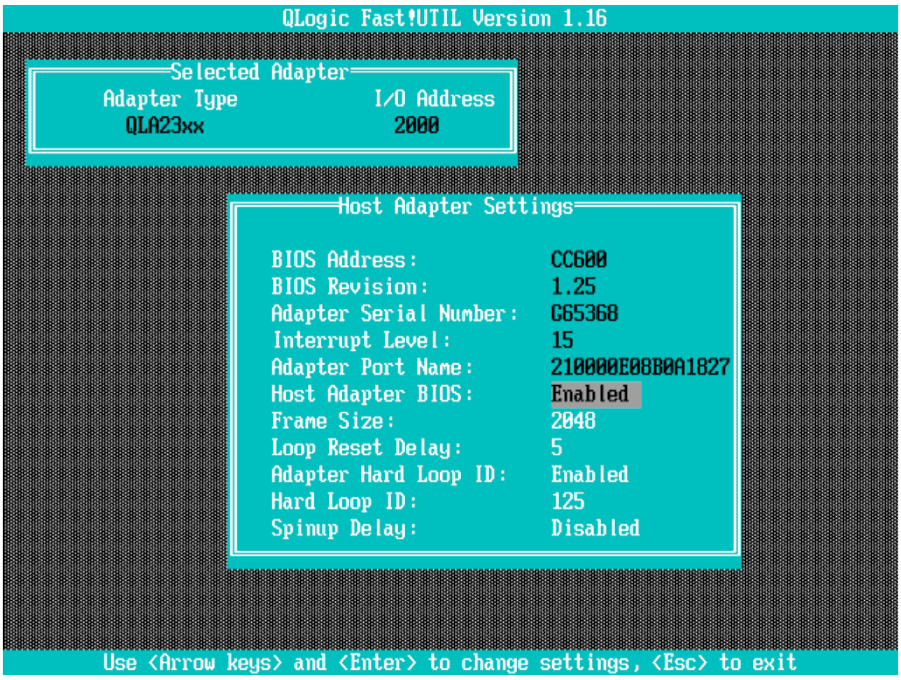


Figure 28-5 Host Adapter Settings

4. Change the Host Adapter BIOS to enabled (the default value is disabled).
5. Record the Adapter Port Name. This is the World Wide Port Name (WWPN) of the HBA. We will need this number when configuring the zone in the SAN switches and storage partition in Storage Manager.
6. Record the system's second HBA WWPN as well.
7. Save the configuration by pressing Esc twice, and leave the Fast!Util open. We do not need to reboot the server at this stage since we will come back to Fast!Util after we configure the SAN Switch and storage.

Step 3 - Zone configuration on the SAN switch

Configure zoning on the FC SAN switch, as shown in Figure 28-1 on page 713. Refer to the vendor specific SAN switch guide for detailed instructions on SAN zoning.

Step 4 - DS3400 storage subsystem configuration

The DS3400 storage subsystem configuration should be performed from the Remote Management workstation using the Storage Manger Client utility. Create a logical drive to be used as the boot disk using the Create Logical Drive Wizard and map it to the host using the Mappings View.

Important: The logical volume must have LUN ID 0 to be able to boot the operating system and the host type should be defined as LNXCLVMware to disable the AVT feature, because failover will be handled on the ESX Server operating system level.

Refer to Chapter 9, "Administration - Configure" on page 131 for detailed step-by-step procedures on creating logical drives and mapping of hosts.

At this point, the storage configuration process is complete.

Step 5 - Fibre Channel host configuration

Back in Fast!Util, we continue to configure the server's first HBA to enable and map the boot device.

In the Fast!Util menu, select **Configuration Settings** and then select **Selectable Boot Settings**. In the Selectable Boot Settings, change the value of Selectable Boot to **Enabled**. Move the cursor down to (Primary) Boot Port Name and press Enter to select the boot device to the DS3400 controller WWPN and ensure that the boot LUN is 0, as shown in Figure 28-6.

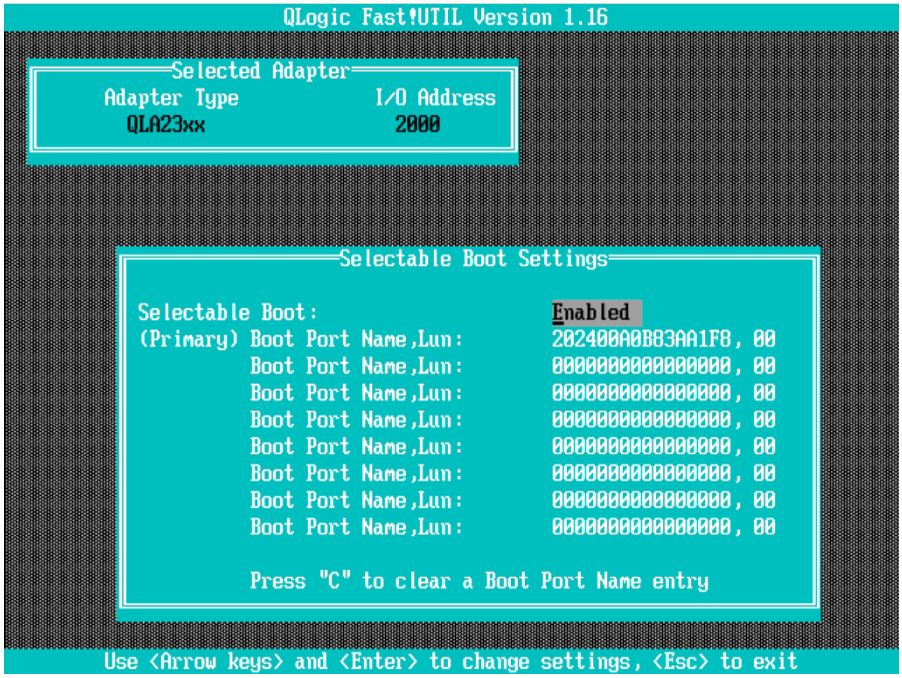


Figure 28-6 Selectable Boot Settings

The server should now detect a new DS3400 logical disk. We can proceed with the ESX Server installation.

Step 6 - ESX Server software installation

For this step:

1. Boot the server using the VMware ESX Server installation CD.
2. The FC device driver is already included in the server software, so loading a specific driver during the installation is not necessary.
3. Click **Yes** on the drive initalization.

4. Make sure that the drive is correctly recognized (Figure 28-7).

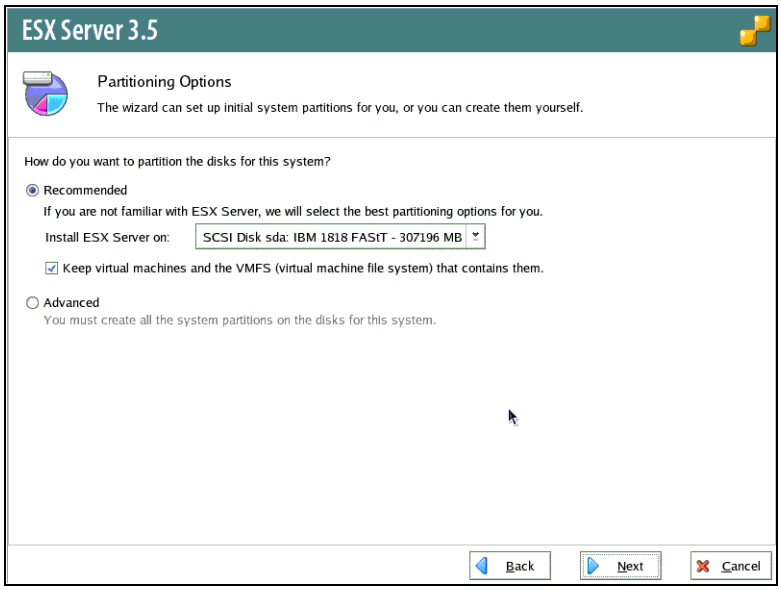


Figure 28-7 SAN Disk recognition

5. Proceed as in a normal ESX server installation and restart the server.
6. The server should display a similar message on your local console, as shown in Figure 28-8.



Figure 28-8 Finished VMware ESX server installation

Step 7 - Verify the ESX Server paths

As the VMware ESX Server already has the multipathing drivers included, no further configuration is needed. To check the available paths, the VMware Infrastructure Client is needed. The client can be downloaded directly from the freshly installed server using the Web interface address showed in Figure 28-8.

Follow these steps to check the available paths:

1. Open the VMware Infrastructure Client and connect to the ESX Server.
2. Click the **Configuration** tab.
3. Click **Storage** in the left navigation field.
4. Right-click the specific storage and choose **Properties** (Figure 28-9 on page 719).

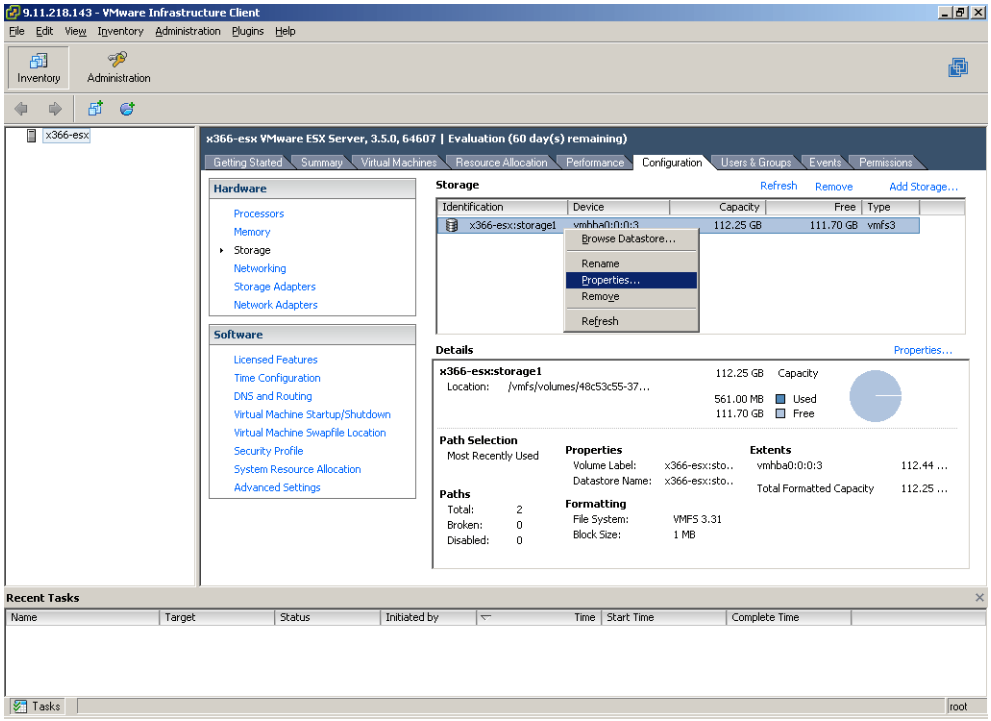


Figure 28-9 ESX Server storage properties

5. Click **Manage Paths** (Figure 28-10).

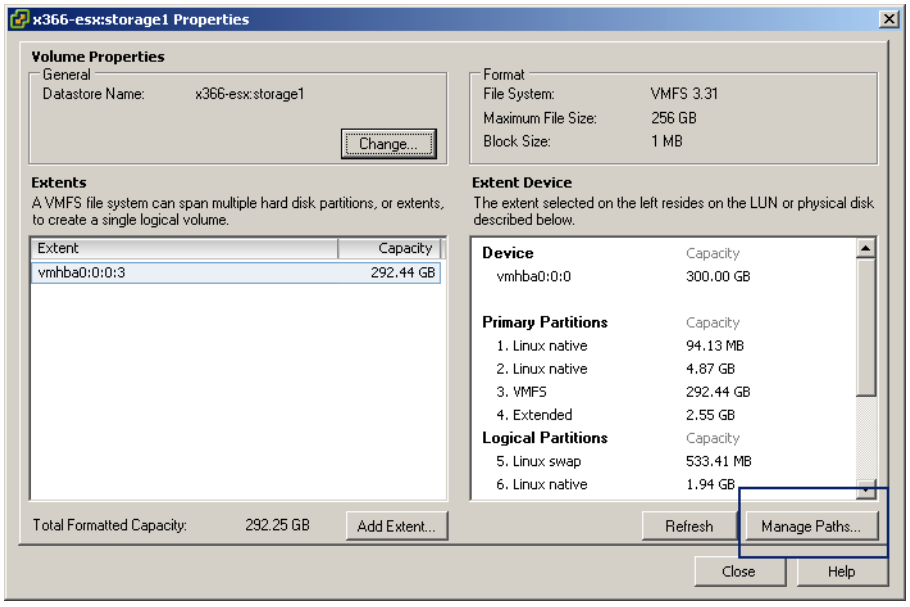


Figure 28-10 Open Manage Paths window

6. In the window shown in Figure 28-11, all paths to the specific storage are shown. The Device field displays your HBA and the SAN Identifier your corresponding Controller on the DS3400. Make sure that the failover policy is set to **Most Recently Used**. This is the default and should be chosen for the DS3400 storage subsystem.

Note: Using active/passive arrays with a fixed path policy can potentially lead to path thrashing. For more information about Active/Active and Active/Passive Disk Arrays and path thrashing, see the *SAN System Design and Deployment Guide*, found at:

`http://www.vmware.com/pdf/vi3_san_design_deploy.pdf`

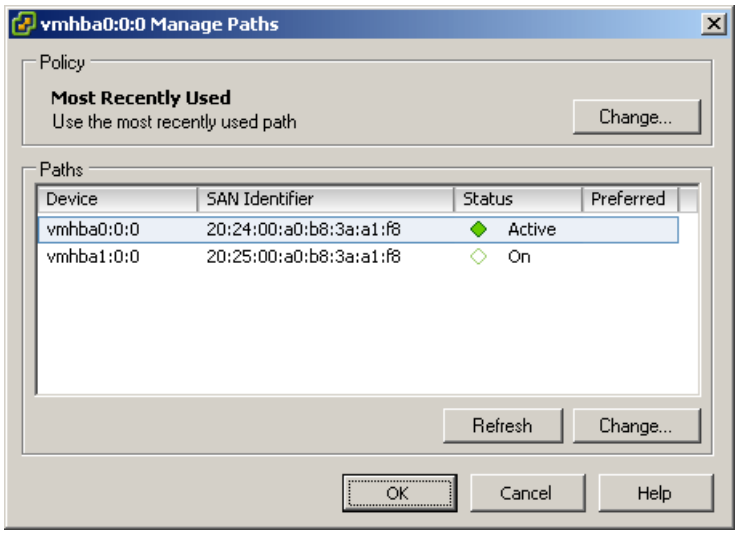


Figure 28-11 Manage Paths

Step 8 - Controller failover simulation

In this section, we simulate a controller failure and discuss how the ESX Server does the failover to the other path. As seen in Figure 28-11, the active and preferred path is Zone1, as shown in Figure 28-1 on page 713. This is the normal and expected situation.

Now we simulate a controller failure by deleting Zone1. The path to Controller A immediately becomes unavailable and the failover takes place. Eventually, you will see that the second path to Controller B becomes active. We do not recommend doing this in your production environment.

The failover should take approximately 30 seconds to 1 minute to complete. To check the current path, check the Manage Paths window. It should look similar to Figure 28-12 on page 721.

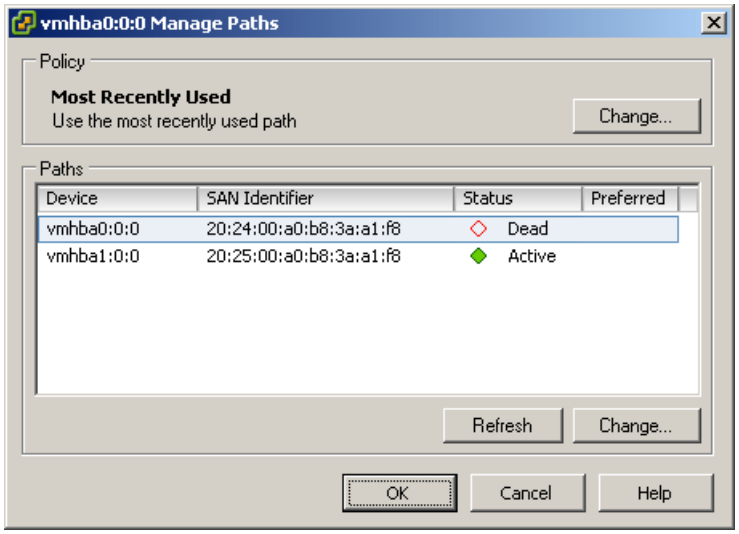


Figure 28-12 ESX path failover

Note: ESX Server will periodically check the failed path but will not move over to the preferred path automatically once the failed path is active again. This is achieved by redistributing the LUNs to the preferred path on the DS3400 storage subsystem.

The status will change from Dead to On should the failed path become active again (Figure 28-13).

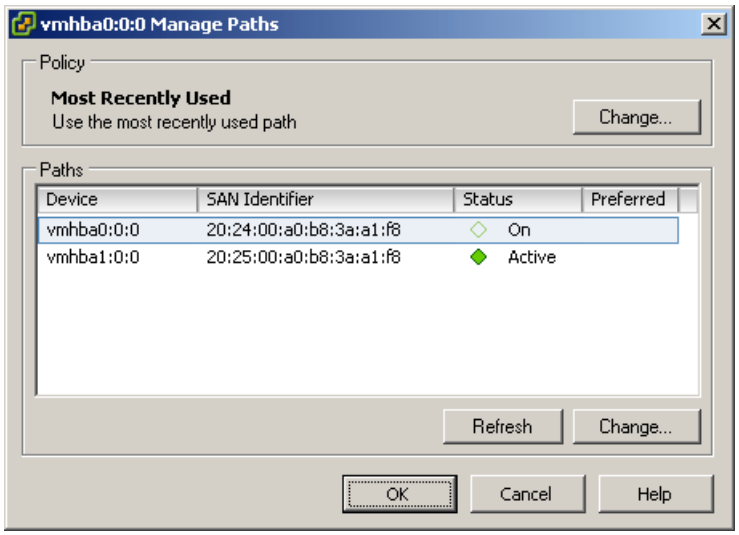


Figure 28-13 Preferred path up and running again

Note: Check the preferred path and redistribute the LUN before restarting a server after a failover has occurred.

You can now proceed with the guest host operating system installation as you normally would in an ESX Server environment. See *VMware Basic System Administration for ESX Server 3.5 Guide* for more information, which can be found at:

http://www.VMware.com/pdf/vi3_35/esx_3/r35/vi3_35_25_admin_guide.pdf

See the *IBM System Storage DS3000 Interoperability Matrix* on the IBM support Web site for the complete list of the supported VMware guest operating systems.

Part 5

Appendices

In this part of the book, we provide some additional information about the IBM System Storage DS3000 server:

- An example of configuring and using FlashCopy on Windows 2003
- Information about how to navigate the IBM Support Web site for the DS3000 server
- More details about using the DS3000 command-line interface (CLI)

© Copyright IBM Corp. 2007, 2009. All rights reserved.

A

FlashCopy example - Windows 2003 and 2008

In this appendix, we provide an overview and detailed information about using FlashCopy logical drives with specific operating systems and disk file systems.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

Windows 2003 and 2008 - basic/regular disks

In this section, we discuss the procedures required for creating FlashCopy logical drives on the Windows 2003 or 2008 operating systems using basic/regular disks.

Attention: Failure to complete the steps listed for your host operating system can result in a loss of FlashCopy data consistency.

We discuss the procedures using the following outline:

- Process overview
- Additional instructions for Windows Server 2003 and 2008 - basic disks

Process overview

In the following process overviews, we outline the key steps required to:

- Create a FlashCopy logical drive on Windows Server 2003 or 2008 (using basic or regular disks). See Figure A-1.
- Reuse a disabled FlashCopy logical drive on Windows server 2003 or 2008. See Figure A-2 on page 728.

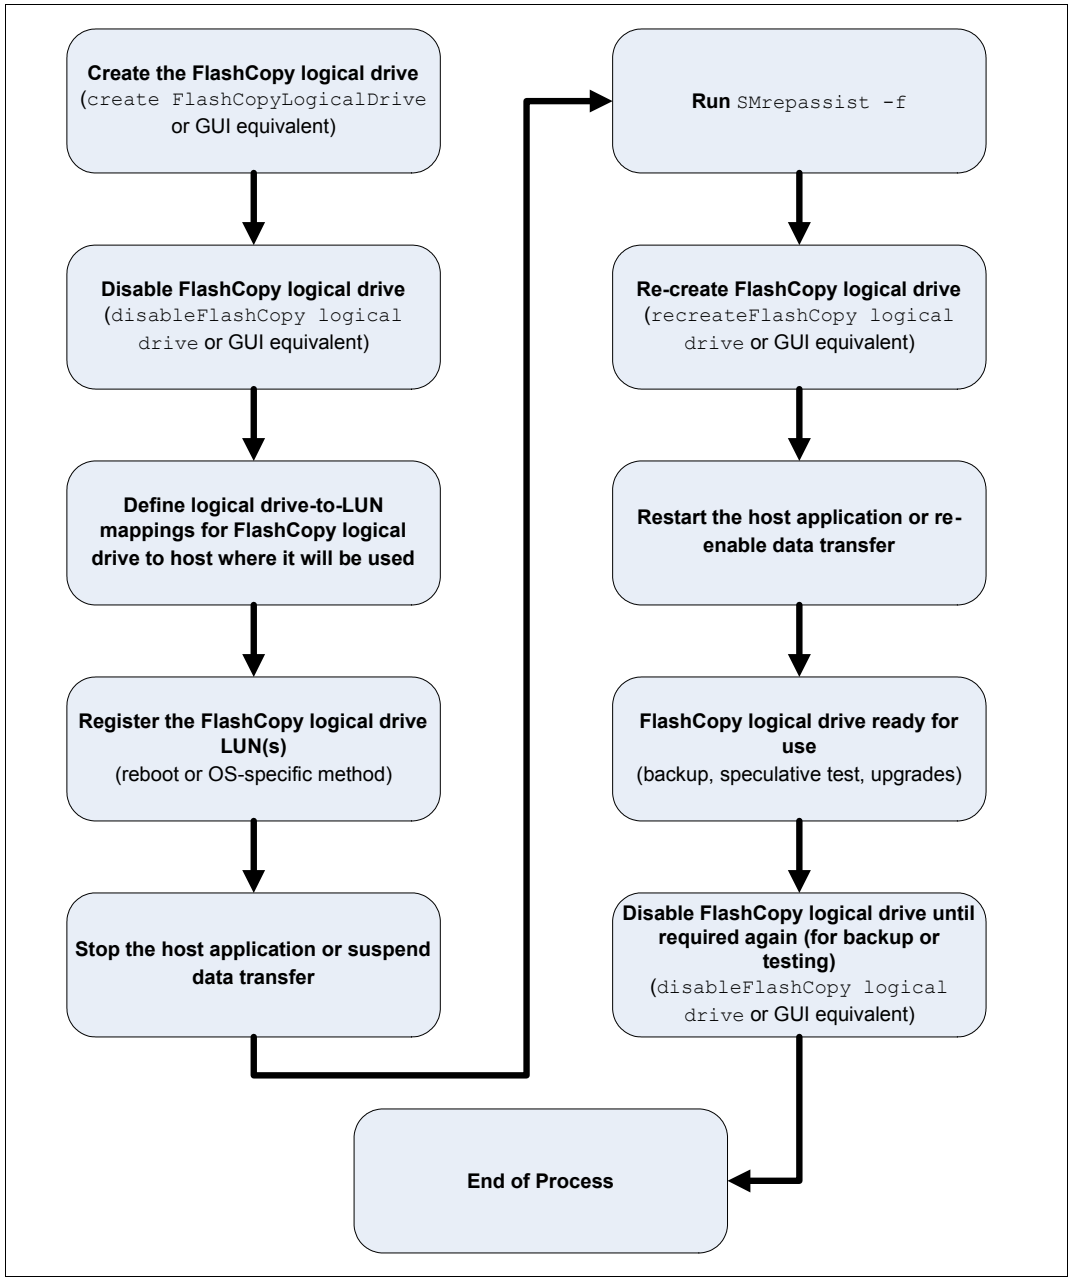


Figure A-1 Creating a FlashCopy logical drive - Windows 2003 or 2008 basic/regular disks

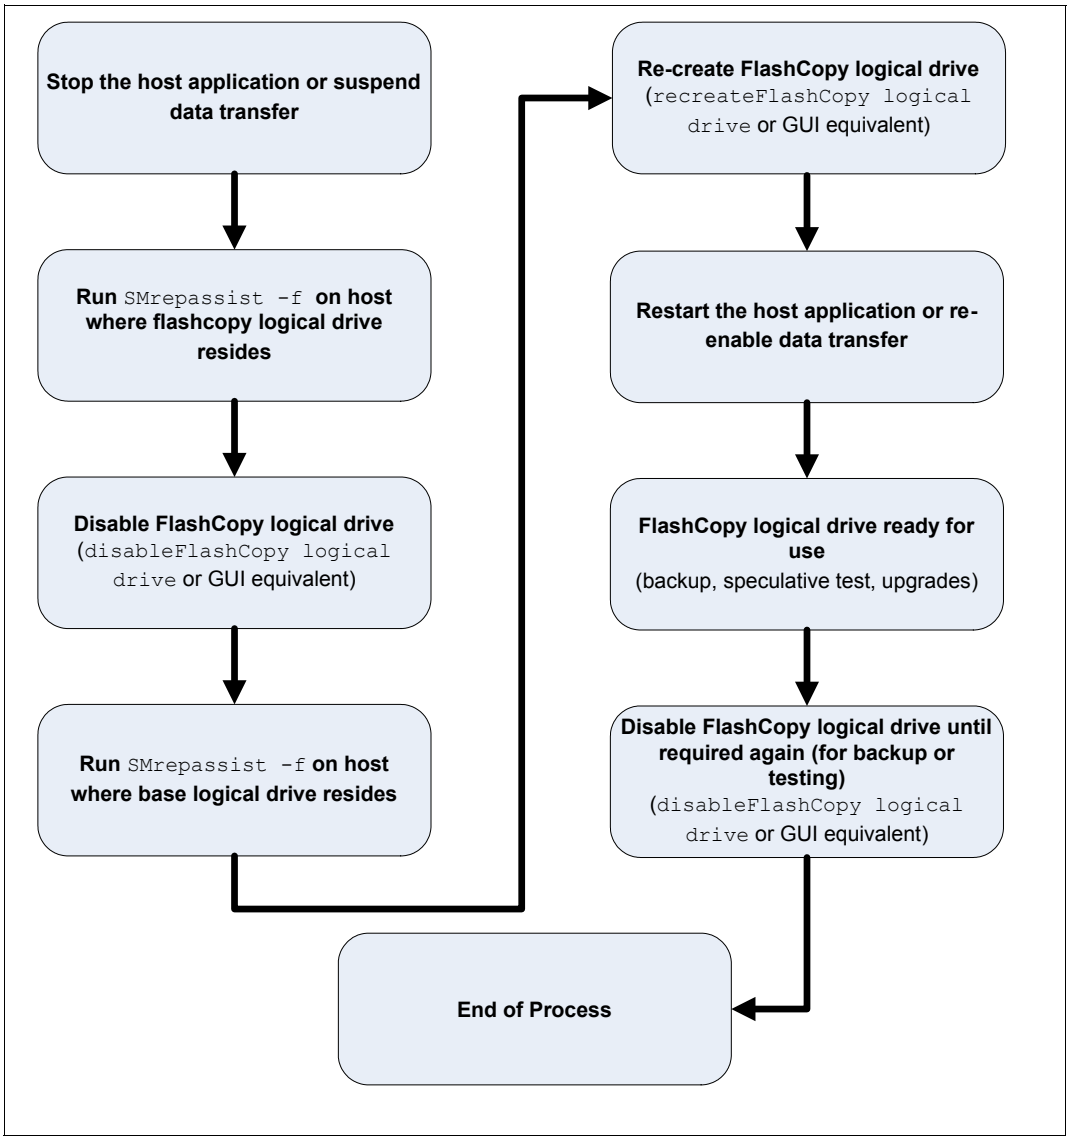


Figure A-2 Re-using disabled FlashCopy logical drives - Windows basic/regular disks

Detailed instructions for Windows 2003 and 2008 basic disks

Use the following procedure when creating FlashCopy logical drives on a host running Windows Server 2003 or 2008, using *basic disks*. Failure to complete the steps listed can result in an inaccurate point-in-time image of the base logical drive.

FlashCopy logical drives can be reused (for frequent or nightly backups) or created for one-time usage (speculative change or upgrade testing). For instructions about how to reuse a disabled FlashCopy logical drive, see “Reusing FlashCopy logical drives” on page 730.

Tip: For command references and information about the use of the Script Editor and CLI, see Appendix C, “CLI” on page 747.

Attention: Risk of Application Errors: Windows operating systems traditionally support a Master Boot Record (MBR) disk partition format. 64-bit Windows Server 2003 and 32- and 64-bit Windows Server 2008 support the new GUID Partition Table (GPT) disk partition format. The 32-bit Windows Server 2003 operating system does not have support for this new partition format and does not know how to properly access a GPT disk.

When a base logical drive mapped to a 64-bit Windows Server 2003 or 32- or 64-bit Windows Server 2008 host is initialized with the GPT disk partition format, any FlashCopy logical drives created from this base logical drive must also be mapped to a 64-bit Windows Server 2003 or an 32- or 64-bit Windows Server 2008 host. If you map the FlashCopy logical drive to a 32-bit Windows Server 2003, the operating system will *not* recognize the partition format of the FlashCopy logical drive.

Alternatively, you can initialize the base logical drive on the 64-bit Windows Server 2003 and 32- or 64-bit Windows Server 2008 host with the older MBR disk partition format, and any FlashCopies of that base logical drive will be accessible by other 32-bit Windows hosts.

Creating a FlashCopy logical drive

We assume that you will use the DS3000 CLI to create the FlashCopy logical drive. In most cases, we just name the command to use. For examples of the precise syntax of the commands, see the online help for the DS3000 CLI.

Tip: For information about how to create FlashCopy logical drives using the DS3000 Storage Manager GUI, refer to 9.4, “Advanced functions - FlashCopy” on page 183.

To create a FlashCopy logical drive, follow these steps:

1. Initialize the FlashCopy logical drive using the following command:

   ```
   create flashcopyLogicalDrive
   ```

2. Disable the FlashCopy logical drive using the following command:

   ```
   stop flashCopy LogicalDrive
   ```

3. Assign a logical drive-to-LUN mapping between the FlashCopy logical drive and the host that will access the FlashCopy logical drive using the following command:

   ```
   set LogicalDrive
   ```

4. Stop the host application that is accessing the base logical drive.
5. Ensure that all data that is related to the base logical drive has been written to disk. In Windows environments, use the SMrepassist utility to flush all the write buffers from the base logical disk drive. This utility is part of the utilities installed during the host installation section of the DS3000 Storage Manager package. Use the following command:

   ```
   SMrepassist -f <filesystem-identifier>
   ```

   where <filesystem-identifier> is the drive letter assigned to the base logical drive.

6. Re-create the FlashCopy logical drive using the following command:

   ```
   recreate flashcopy logicalDrive
   ```

7. Run the hot_add utility or reboot the host where the FlashCopy will be used. This ensures that the host operating system recognizes the FlashCopy logical drive. This utility is part of the utilities installed during the host installation section of the DS3000 Storage Manager package.

8. Open Disk Management in the host:

   Select **Start** → **Settings** → **Control Panel**, double-click the Administrative Tools icon, and then double-click **Computer Management**. In the console tree under Storage, select **Disk Management**. The Disk Management window is displayed with a graphical representation of all the physical disks connected to the host and their associated partitions.

   In the Disk Management window, locate the disk and logical drive definition that represents the FlashCopy logical drive you re-created and ensure that a new drive letter has automatically been assigned.

9. Use the FlashCopy logical drive in conjunction with your backup application (reusing a FlashCopy logical drive) or for speculative change and upgrade testing (one-time usage).

Once the FlashCopy logical drive is no longer required, disable or delete the FlashCopy logical drive.

Reusing FlashCopy logical drives

Complete the following steps to create a new point-in-time image of the same base logical drive:

1. Stop the host application that is accessing the base logical drive.
2. Ensure that all data that is related to the base logical drive has been written to disk. In Windows environments, use the SMrepassist utility to flush all the write buffers from the base logical disk drive. This utility is part of the utilities installed during the host installation section of the DS3000 Storage Manager package. Use the following command:

   ```
   SMrepassist -f <filesystem-identifier>
   ```

   where <filesystem-identifier> is the drive letter assigned to the base logical drive.

3. List all of the mounted logical drives, and locate the disk representing the FlashCopy logical drive in the host. At the command prompt, type the following command:

   ```
   mountvol [drive:]path /l
   ```

   where [drive:]path is the NTFS folder where the mount point resides.

   All of the mount logical drives are listed. For more information about using **mountvol**, refer to your operating system documentation.

4. Unmount the disk representing the FlashCopy logical drive in the host. At the command prompt, type the following command and then press Enter:

   ```
   mountvol [drive:]path /p
   ```

   where [drive:]path is the NTFS folder where the mount point resides.

   The disk is unmounted. For more information about using the **mountvol** utility, refer to your operating system documentation.

5. Remove any logical drive-to-LUN mappings for the FlashCopy logical drive. Use the command:

   ```
   remove LogicaDrive ["logical drive name"]
   ```

   where ["logical drive name"] is the FlashCopy logical drive name.

6. Disable the FlashCopy logical drive by using the following command:

   ```
   stop flashcopy logicalDrive
   ```

7. Run the SMrepassist utility in the host where the base logical drive is mounted to flush all the write buffers from the new disk drive. This utility is part of the utilities installed during the host installation section of the DS3000 Storage Manager package. Run the following command:

   ```
   SMrepassist -f <filesystem-identifier>
   ```

   where <filesystem-identifier> is the drive letter assigned to the base logical drive.

8. Re-create the FlashCopy logical drive by running the following command:

   ```
   recreate FlashCopy logicalDrive
   ```

9. Assign a logical drive-to-LUN mapping between the FlashCopy logical drive and the host that will access the FlashCopy logical drive by running the following command:

   ```
   set LogicalDrive
   ```

10. Run the hot_add utility (or operating system-specific utility) or reboot the host where the FlashCopy will be used. This will ensure that the host operating system recognizes the FlashCopy logical drive. This utility is part of the utilities installed during the host installation section of the DS3000 Storage Manager package.

11. Mount the disk representing the FlashCopy logical drive by running the following command:

    ```
    mountvol [drive:]path Logical Drive{GUID}
    ```

    where [drive:]path is the NTFS folder where the mount point will reside, and Logical Drive{GUID} is the name of the logical drive that is the target of the mount point in Logical DriveGUID format.

    The disk representing the FlashCopy logical drive is mounted. For more information about using **mountvol**, refer to your operating system documentation.

12. This step is optional. If you are configuring your storage system using auto scripting operations, we recommend that you complete this step.

    Assign the original logical drive label to the disk representing the FlashCopy logical drive. At the command prompt, type the following command and then press Enter:

    ```
    label [drive:][label]
    ```

    where [drive:] is the location of the disk you want to name, and [label] is the new logical drive label.

    The disk representing the FlashCopy logical drive is renamed. For more information about using the **label** command, refer to your operating system documentation.

13. This step is optional. If you are configuring your storage system using auto scripting operations, we recommend that you complete this step.

    Run the **chkdsk** command on the FlashCopy logical drive to create and display a status report for the disk. Run:

    ```
    chkdsk [drive:]
    ```

    where [drive:] is the drive that contains the disk that you want to check.

    A status report for the FlashCopy logical drive is displayed. For more information about using the **chkdsk** command, refer to your operating system documentation.

14. Use the FlashCopy logical drive in conjunction with your backup application (or with another application).

When the FlashCopy logical drive is no longer required, disable the FlashCopy logical drive.

B

IBM Support Web site

In this appendix, we explain how and where to find information about the IBM Support Web site for the IBM System Storage DS3000 server.

© Copyright IBM Corp. 2007, 2009. All rights reserved.

Sample navigation procedure

The following steps show, as an example, how to navigate the IBM Support page to find updated firmware code for a DS3400 server. Also, Storage Manager updates, HBA firmware and drivers, as well as product documentation, can be found using similar steps.

1. In your Web browser, go to:

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

   A window similar to Figure B-1 appears. Select **System Storage** from the Choose support type box. Click the right arrow.



Figure B-1 IBM Support - Main entry

2. Select **Disk systems** from the Product family drop-down menu, as shown in Figure B-2.

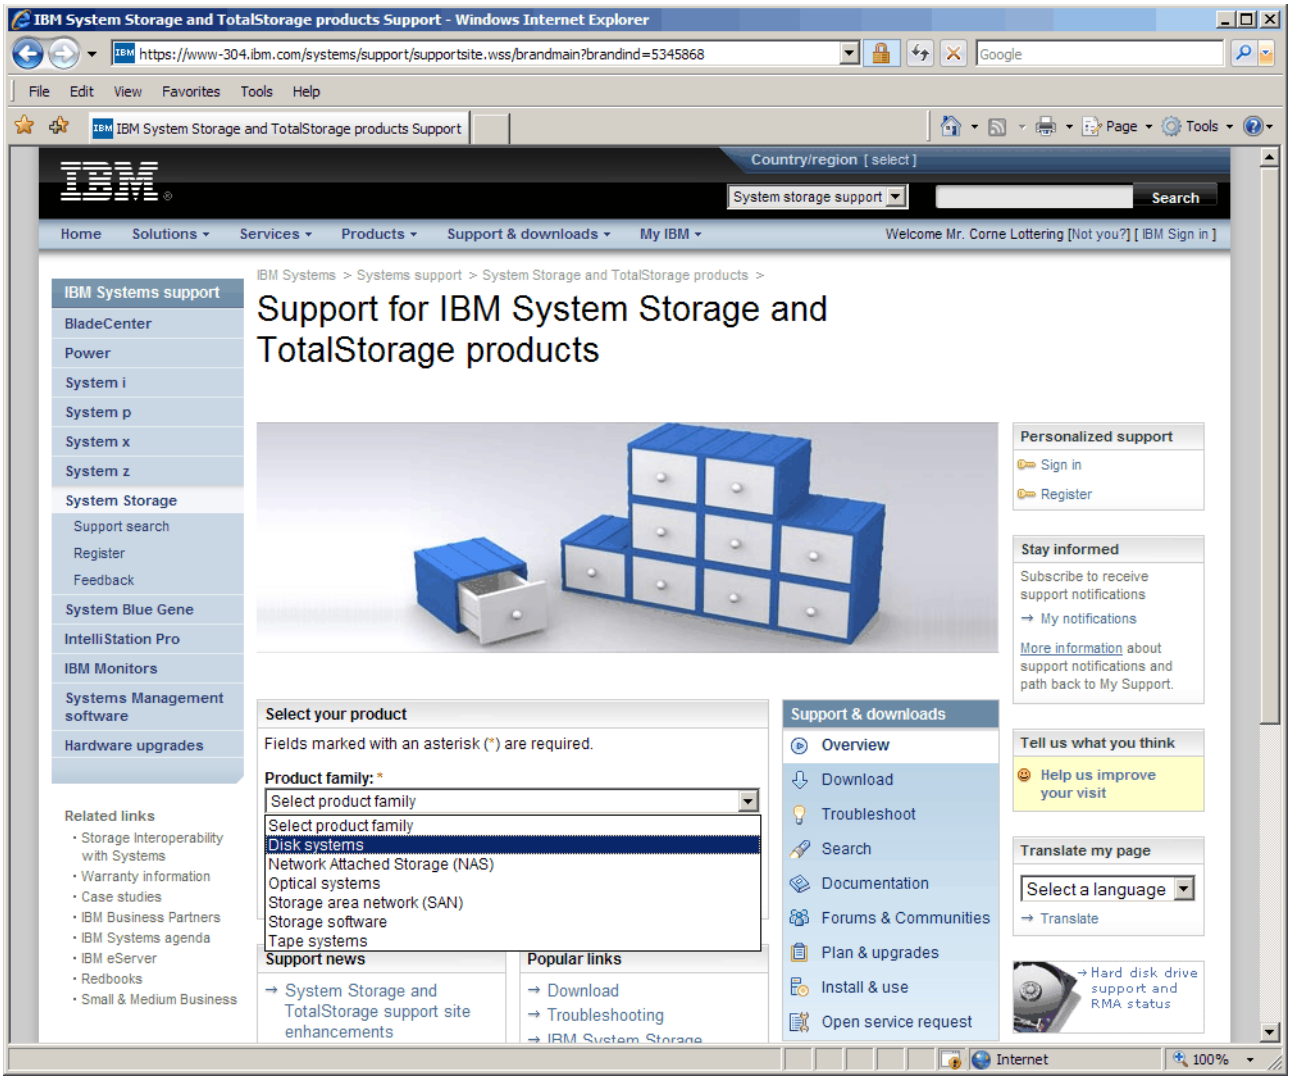


Figure B-2 IBM Support - Select Product family

3. From the Product drop-down menu, select the appropriate DS3000 storage product. Figure B-3 uses the DS3400 server as an example.

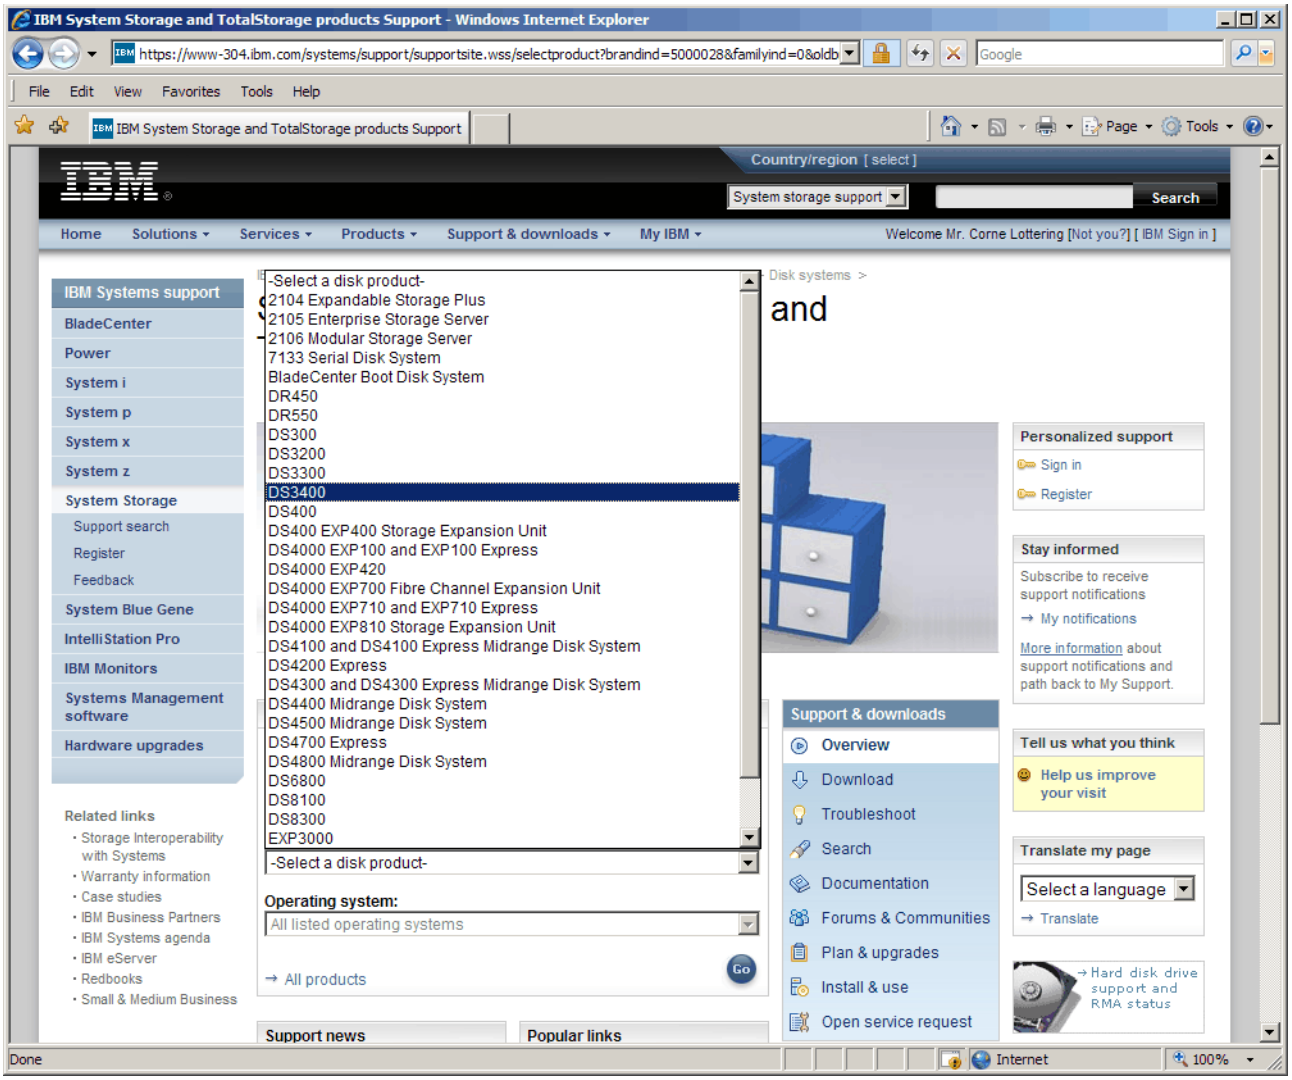


Figure B-3 IBM Support - Select a product

4. Click the **Go** button, as shown in Figure B-4.

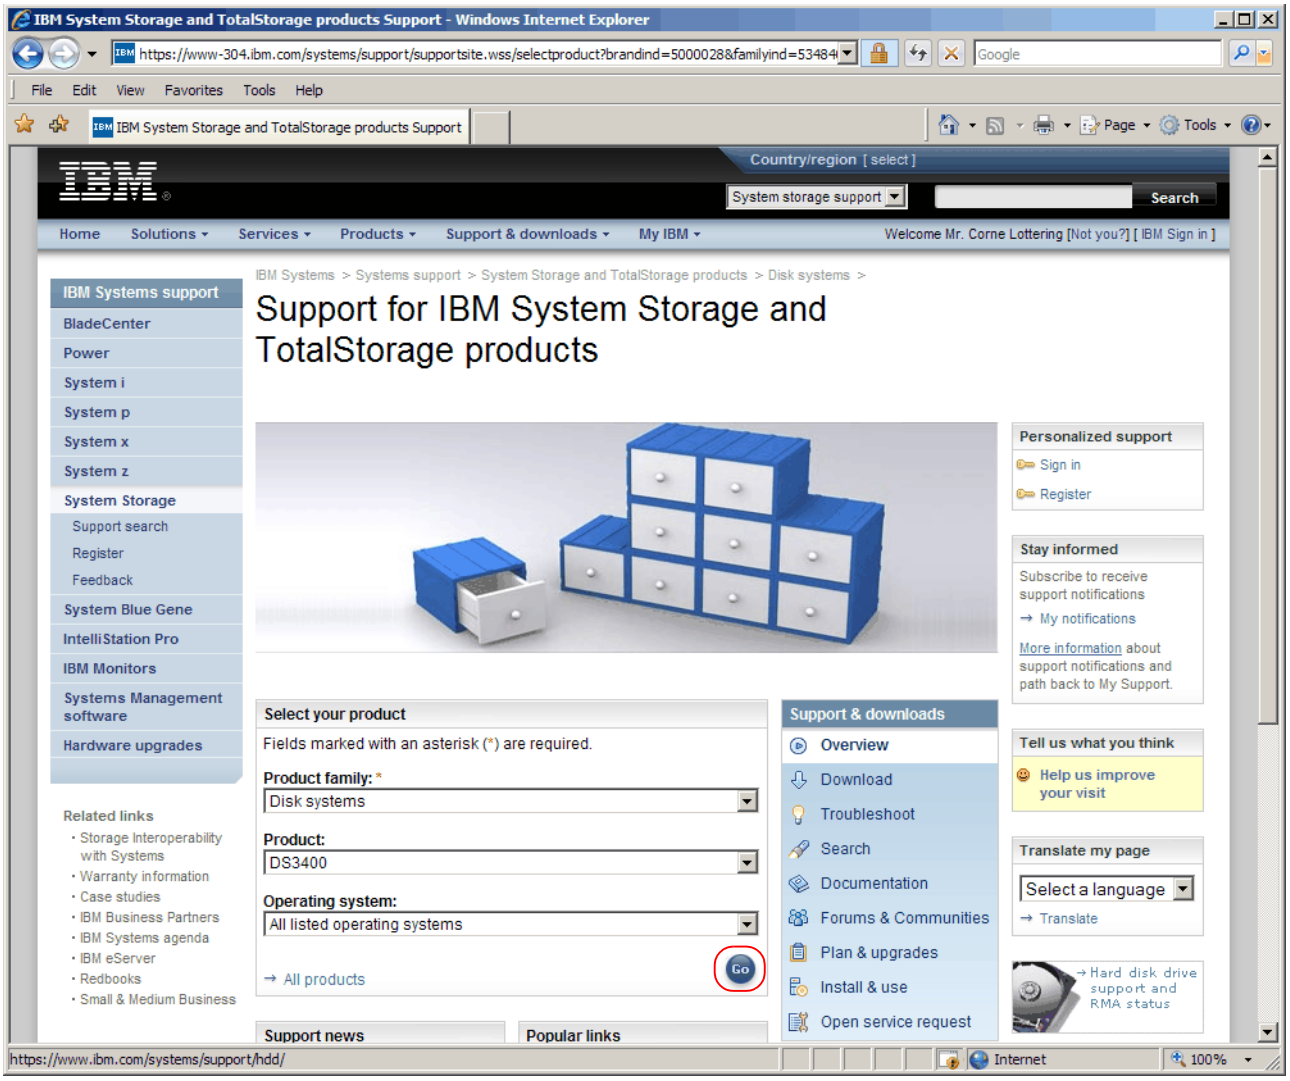


Figure B-4 IBM Support - Go

5. The DS3400 support page displays, as shown in Figure B-5. The Support & downloads box shows the available categories. Microcode, firmware, and device drives are found in the **Download** section. Click it now.

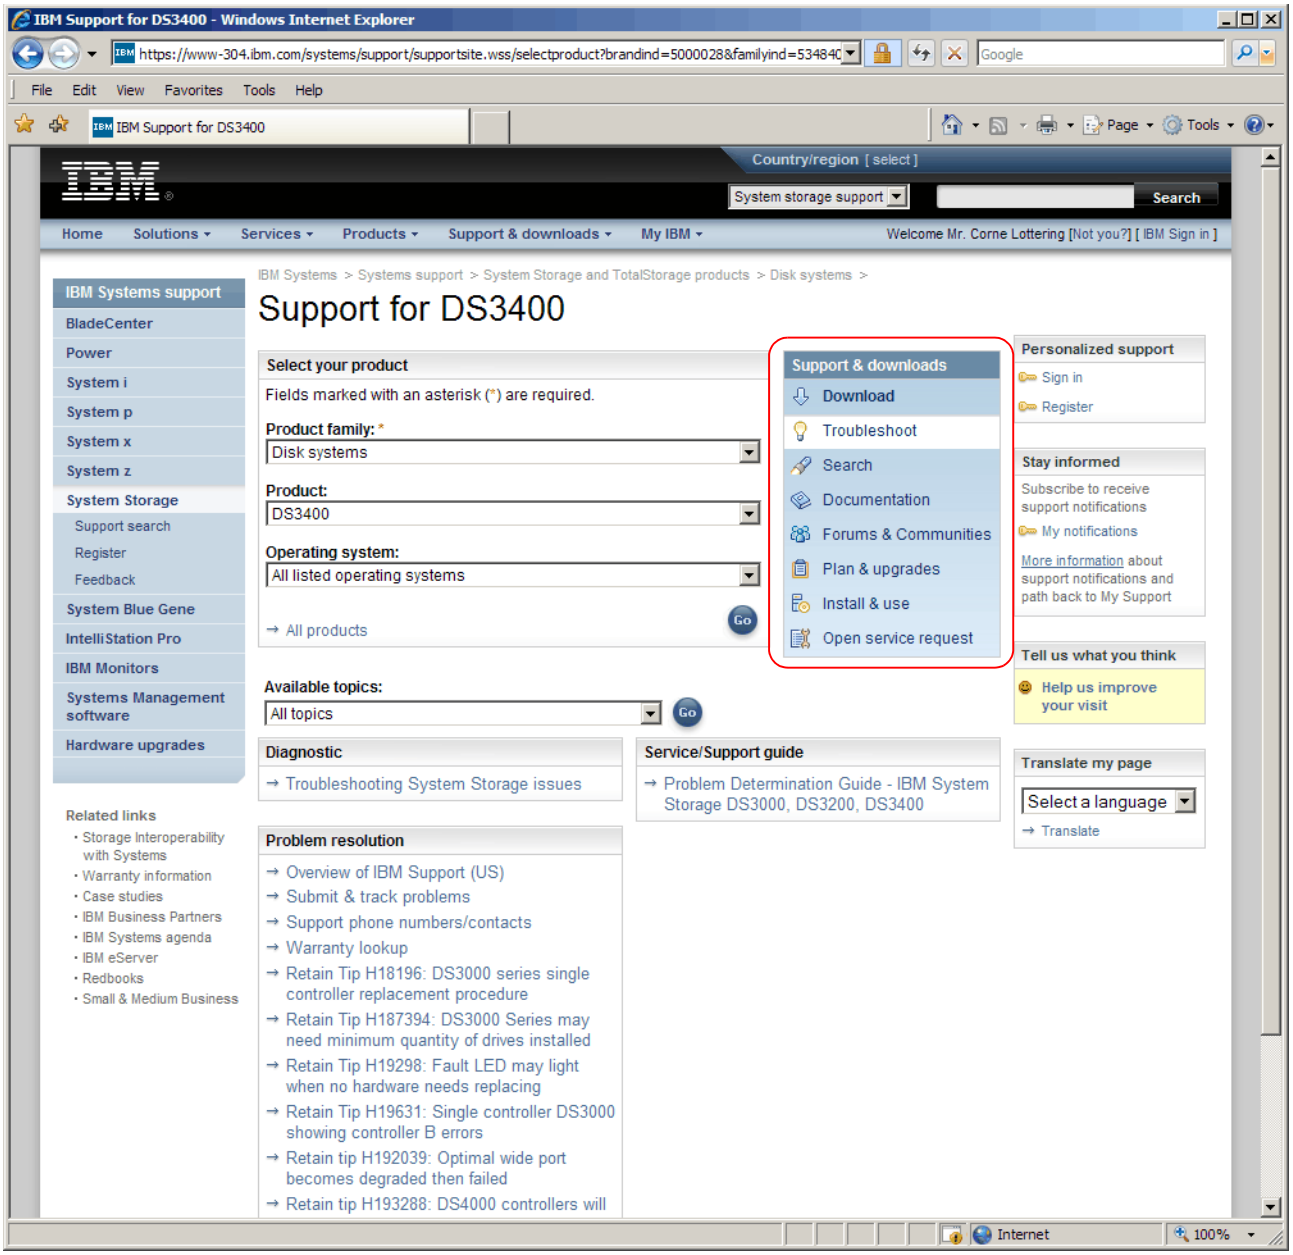


Figure B-5 IBM Support - DS3400 support area

6. Now you can download all the code necessary to operate a DS3000 server. Click the appropriate section for the component you require.

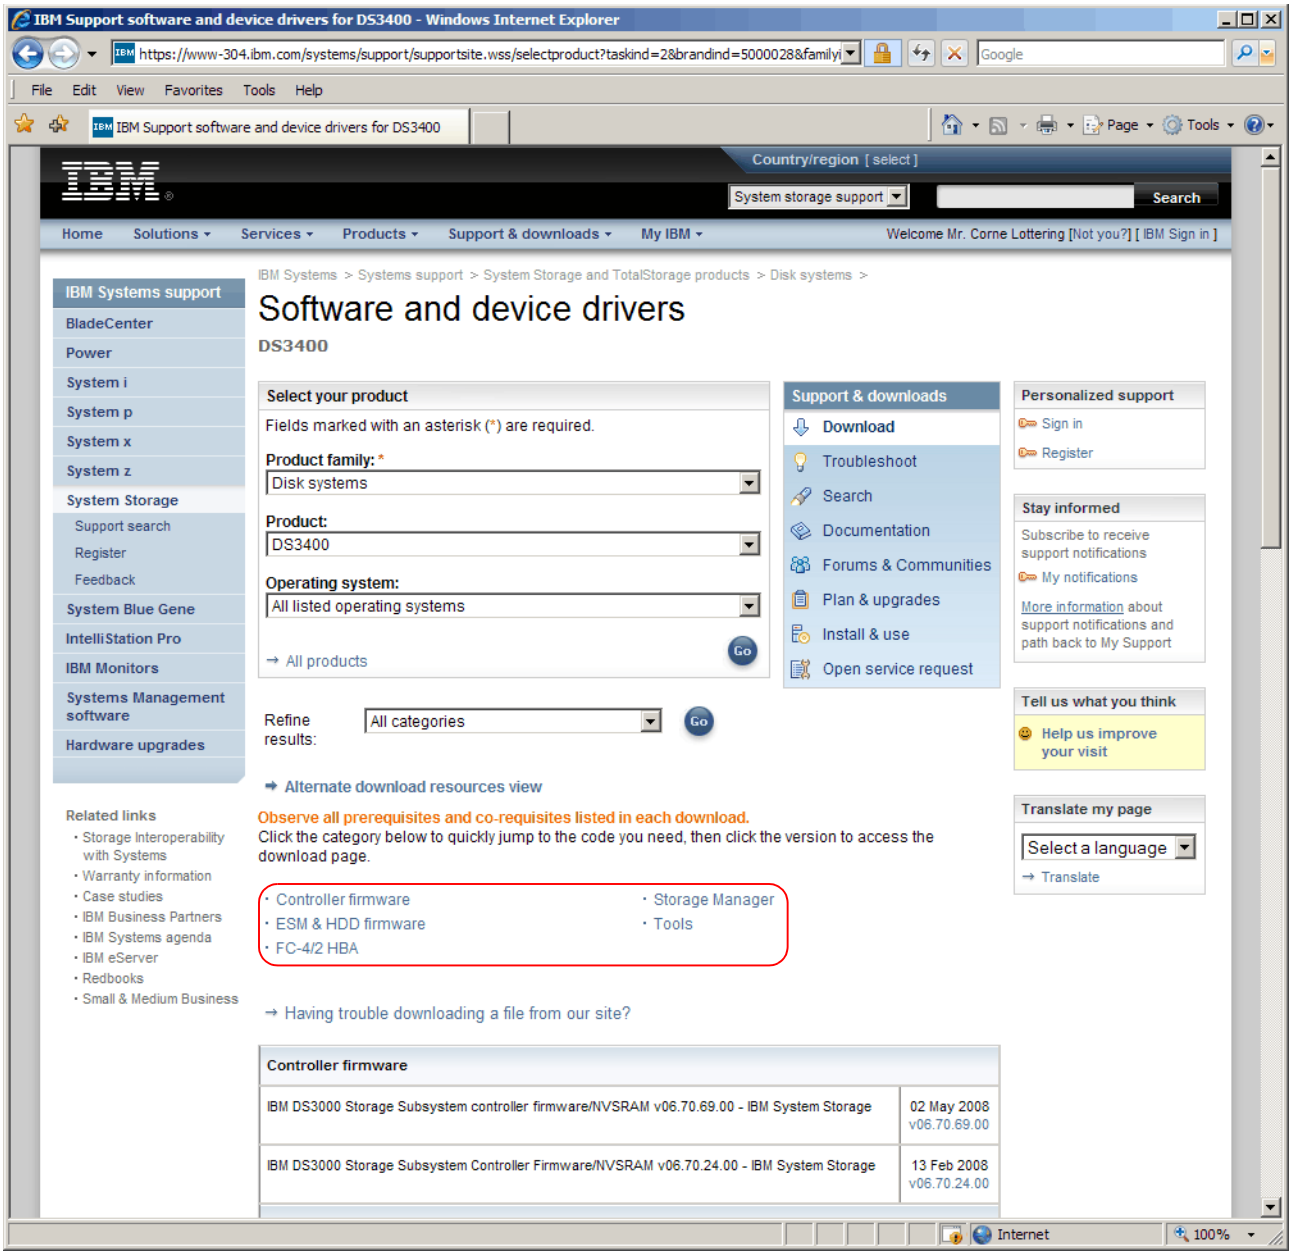


Figure B-6 IBM Support - DS3400 download area

My notifications

The IBM My notifications Web site can be used to keep you informed of the latest firmware and other important product updates. The following steps show, as an example, how to set up notifications to receive product updates for the DS3400 System Storage subsystem.

1. In your Web browser, go to:

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

A window similar to Figure B-7 appears. Click **My notifications** in the Stay informed box on the right side of the support Web page.

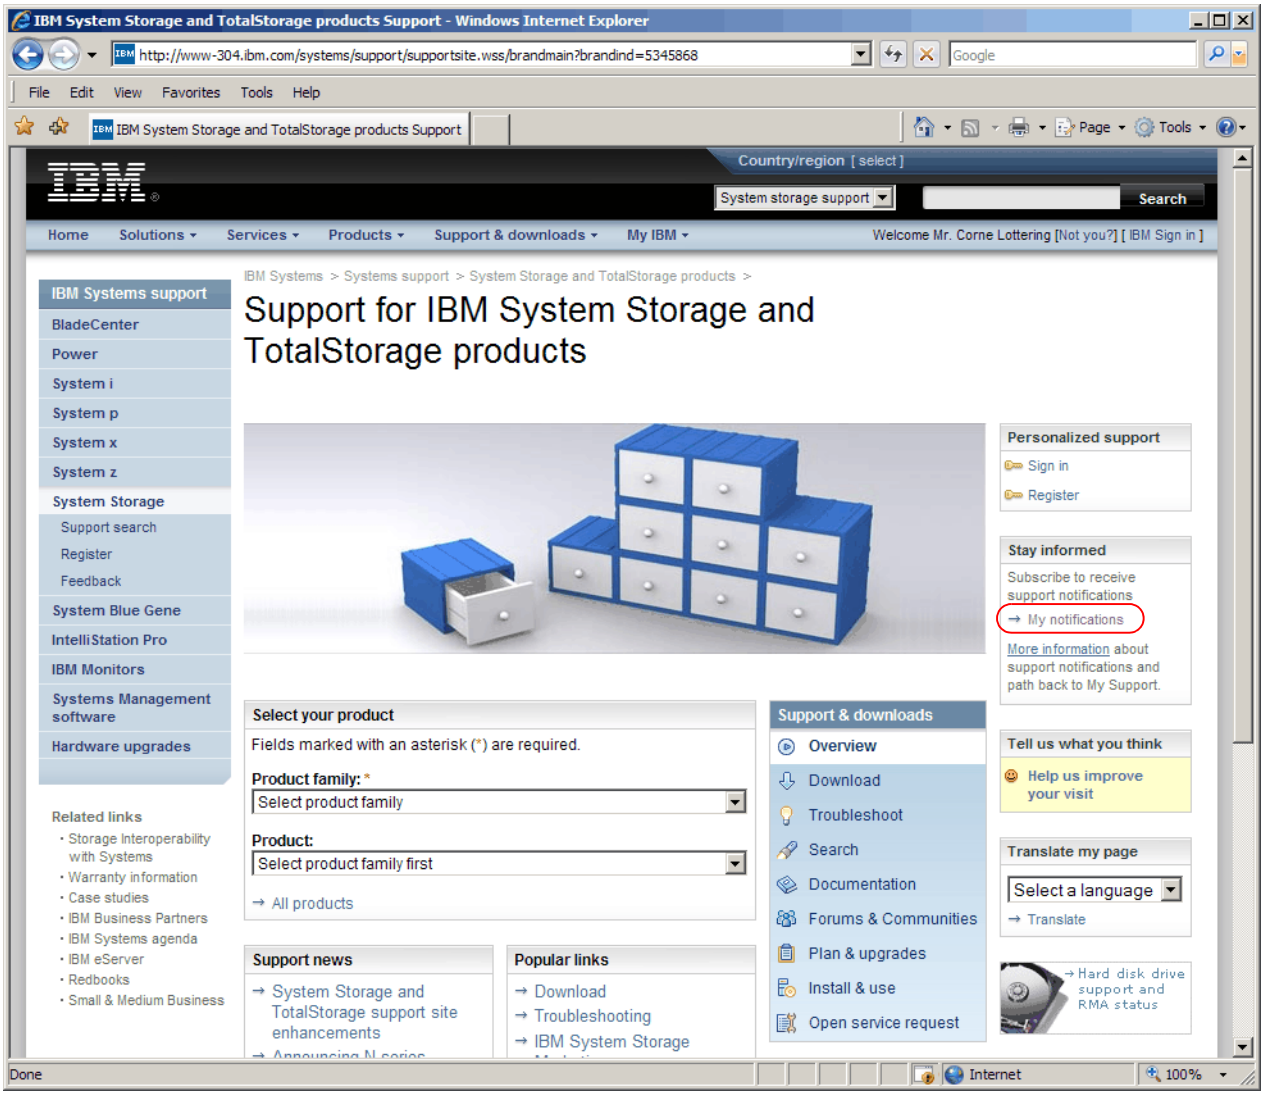


Figure B-7 IBM Storage Support - Main Entry

2. The Sign In window appears, as shown in Figure B-8. Enter your IBM ID and password in the sign in area to sign in and proceed. If you are not currently registered with the Web site, click **Register Now**.

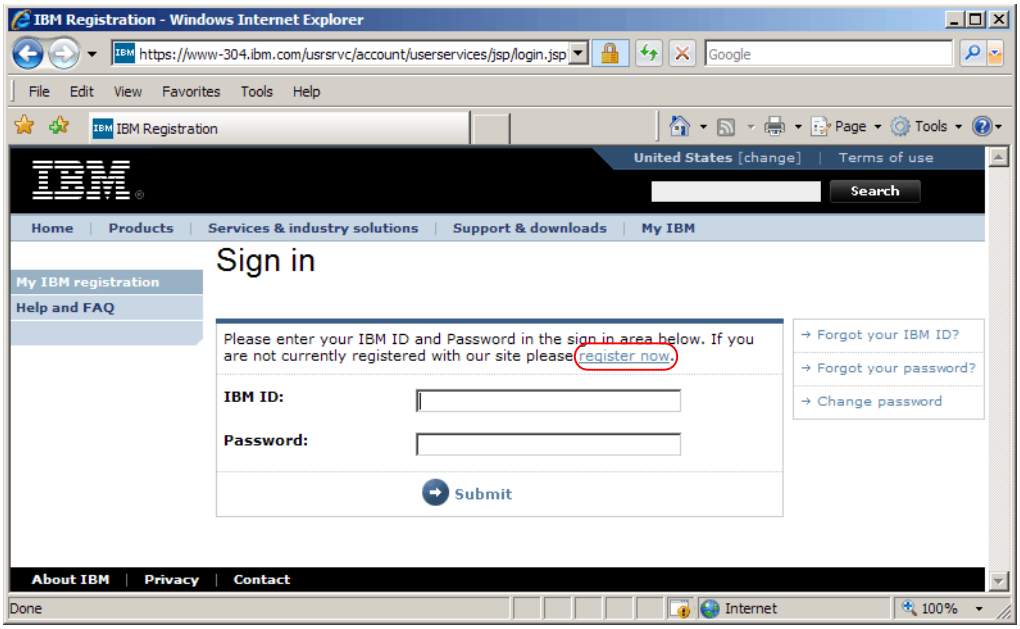


Figure B-8 IBM MySupport - Sign In window

3. The My notifications main page opens, as shown in Figure B-9. Click the **Disk systems** link.

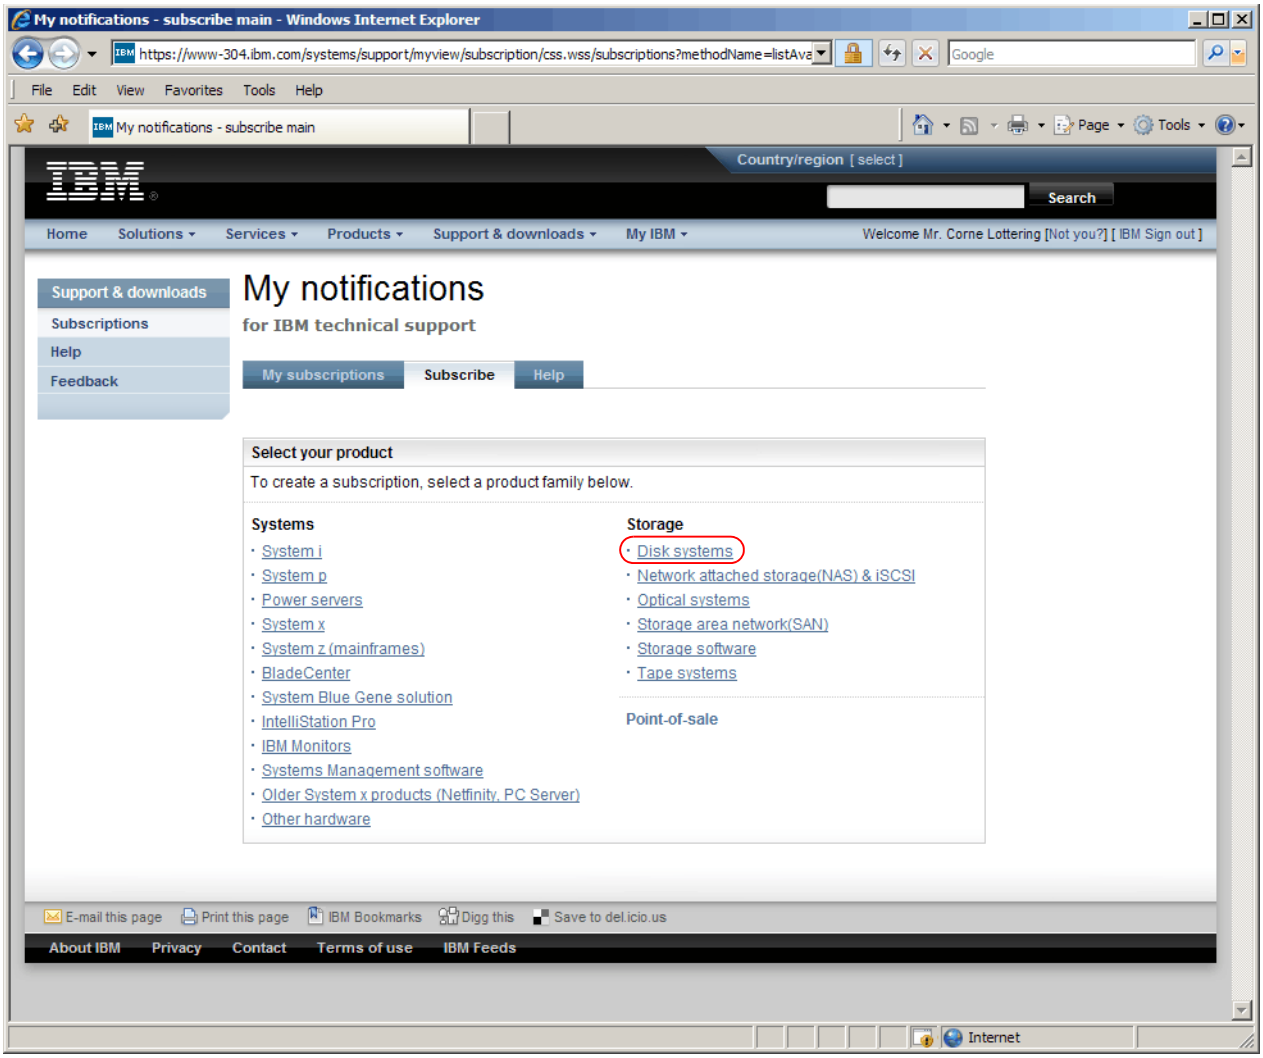


Figure B-9 IBM My notifications - Main page

4. A window appears, as shown in Figure B-10, listing the Disk Storage products that are available for notifications. In our example, we put a tick in the **DS3400** check box to allow us to receive notifications on the DS3400 System Storage subsystem product. Click **Continue** to proceed.

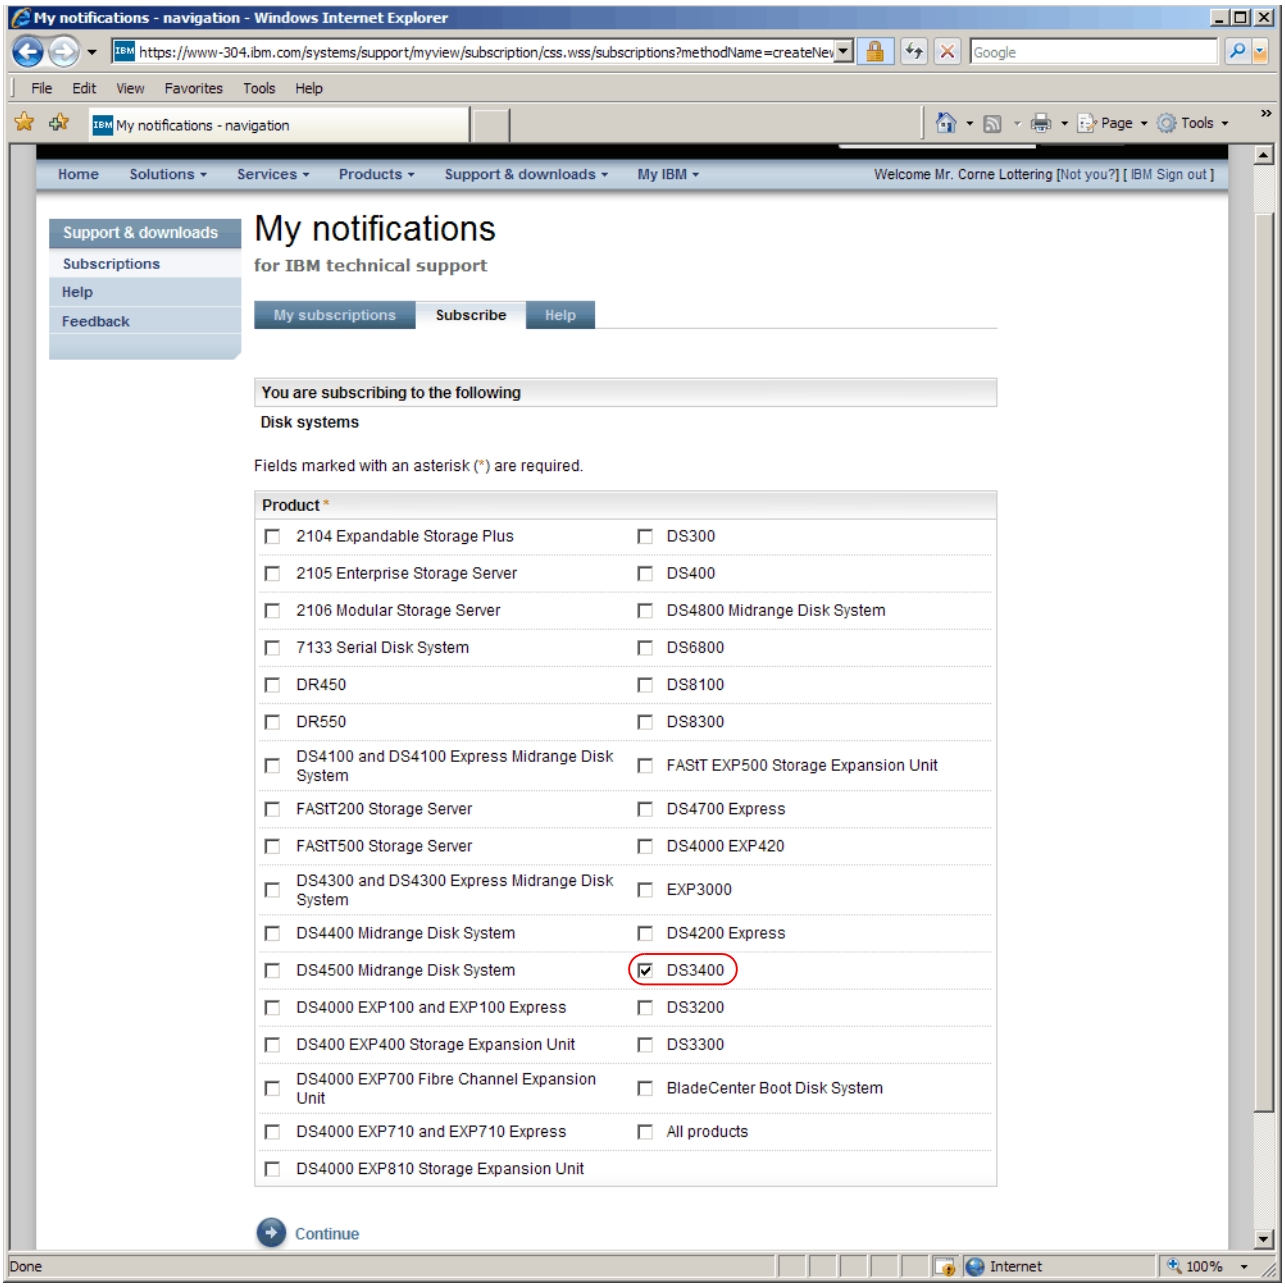


Figure B-10 IBM My notifications - Product notification selection

5. The following window, shown in Figure B-11, allows you to select a name for the notification, save the notification in an existing or new folder, and select the notification e-mail configuration. We keep the name as DS3400, select a new folder called **IBM DS3000**, and accept the default e-mail configurations. Click **Submit** to proceed.

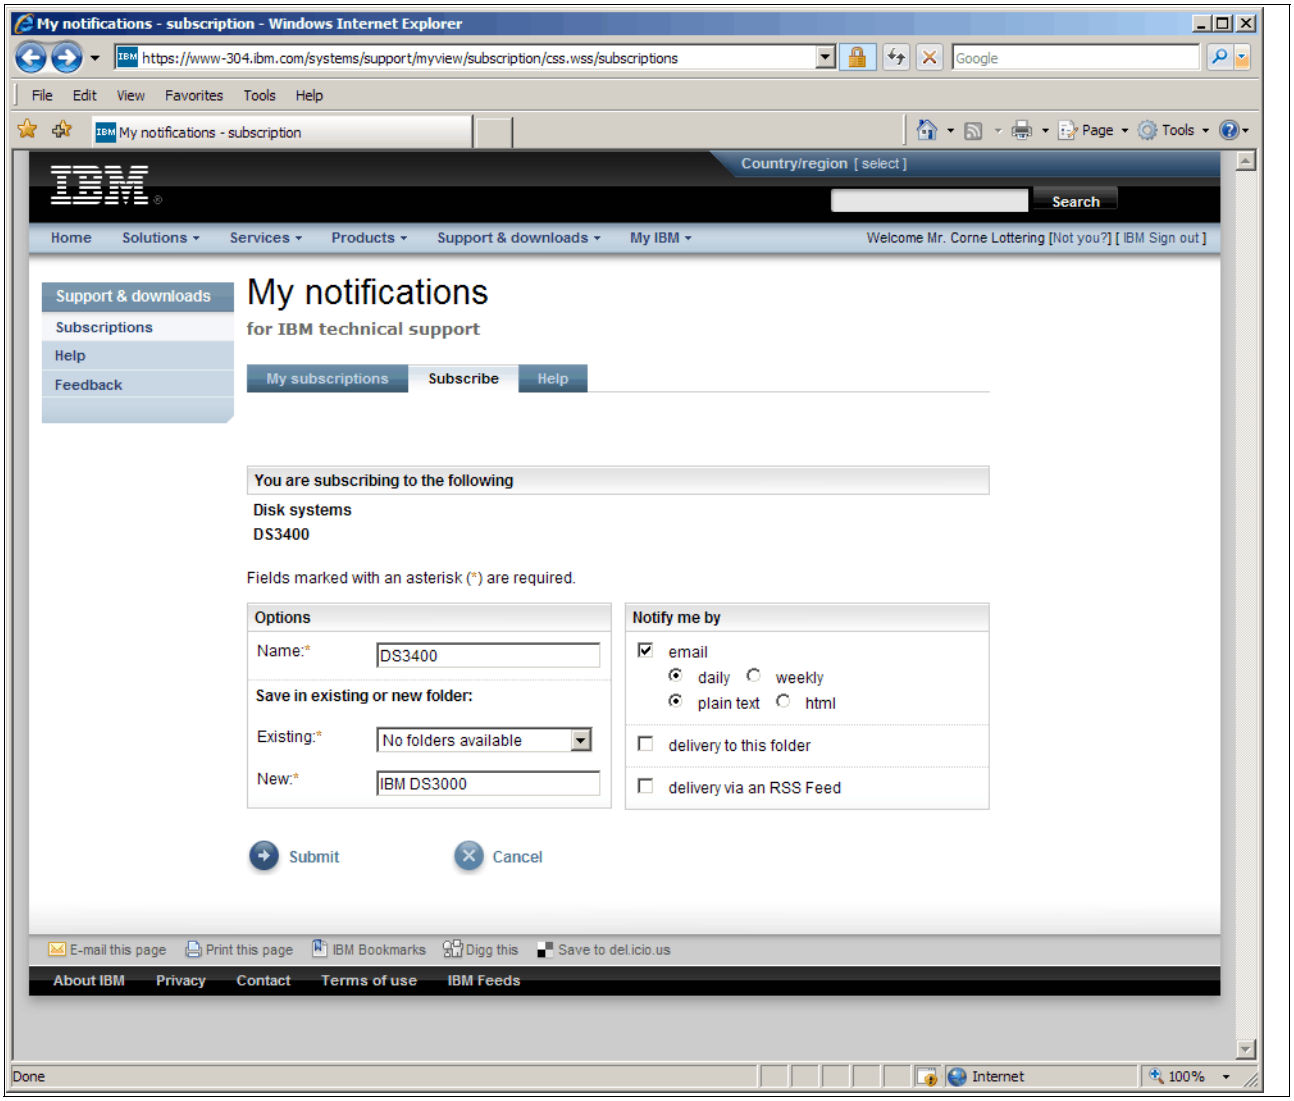


Figure B-11 IBM My notifications - Setup notification

6. The following window, shown in Figure B-12 on page 745, lists the notification setup for the DS3400 System Storage subsystem. This completes the setup of the notifications for the DS3400 System Storage subsystem. You can repeat the above steps to add additional DS3000 System Storage products to your notification list.

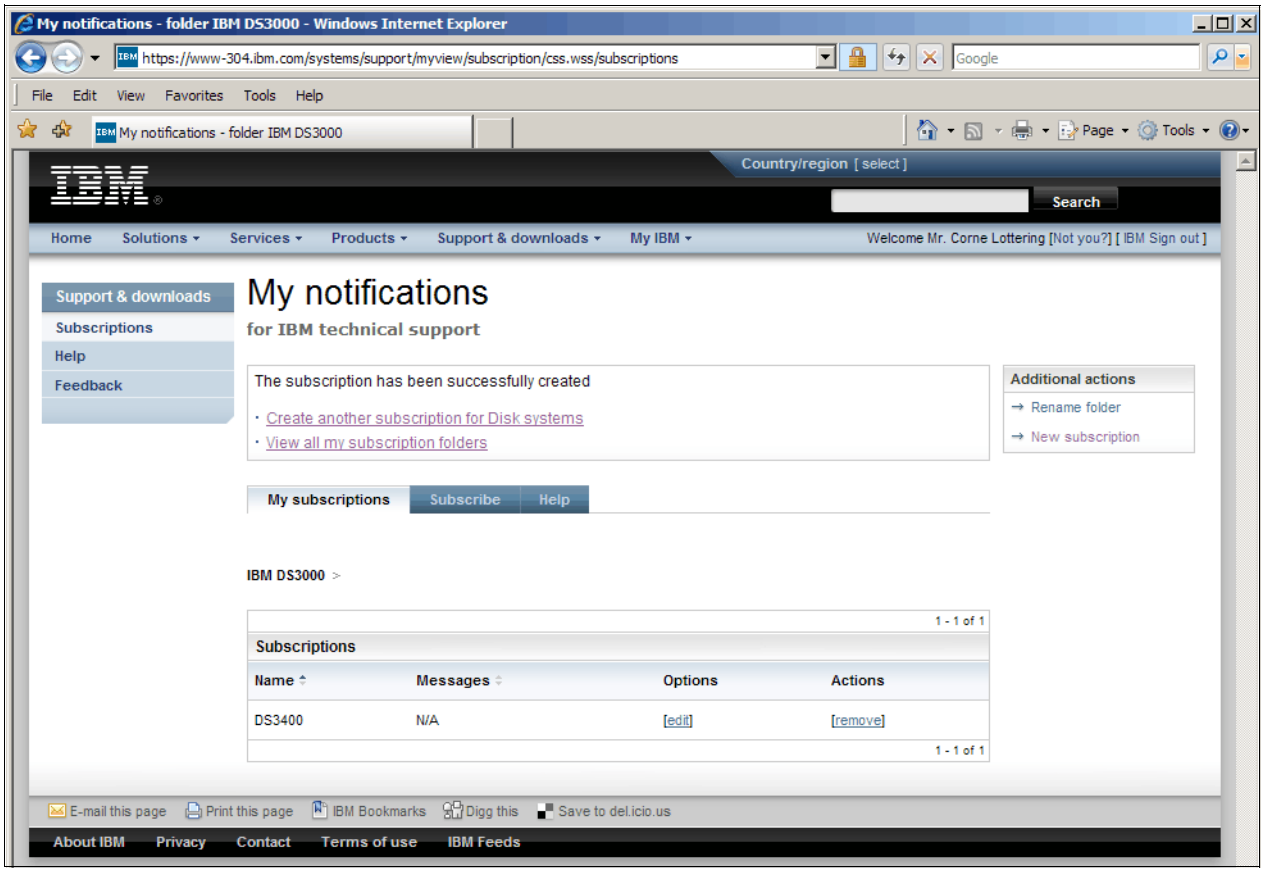


Figure B-12 IBM My notifications - Listing notification

C

CLI

In this appendix, we explain the basics of the Command-Line Interface (CLI) that can be used to manage an IBM System Storage DS3000. Unlike the DS3000 Storage Manager GUI, the CLI can manage a DS4000 server as well.

The CLI is installed together with the DS3000 Storage Manager Client. It can be found in the client subdirectory of the location where the Storage Manager was installed. In Microsoft Windows, that is usually C:\Program Files\IBM_DS3000\client and in Linux /opt/IBM_DS3000/client/. The CLI program is called SMcli in Windows and Linux. Commands that can be used on the CLI to administer a DS3000 or DS4000 storage subsystem are identical to the commands used in the Script Editor. Not all commands apply to each model of the storage subsystem.

The CLI gives access to all the functions provided in the Storage Manager GUI, as well as some additional management functions. For example, it is only possible to modify the blocksize of a logical drive using the CLI. Similarly, the CLI is required to save the configuration file.

This appendix is organized in the following sections:

- "Running the CLI" on page 748
- "General SMcli syntax" on page 750
- "Adding a storage subsystem to the Storage Manager configuration" on page 753
- "Showing defined subsystems in the Storage Manager configuration" on page 754
- "Configuring alerts" on page 755
- "Issuing commands to the storage subsystem" on page 758

© Copyright IBM Corp. 2007, 2009. All rights reserved.

Running the CLI

There are two ways to enter commands into the DS3000: interactively, using SMcli, or by creating saved scripts with the Scripting Editor. We describe SMcli in “General SMcli syntax” on page 750.

Script Editor

The Script Editor allows you to create and save files of commands from the CLI so you can batch execute or schedule them.

Do these steps:

1. To start the Script Editor, from Storage Manager, select **Tools** → **Execute Script**, as shown in Figure C-1.

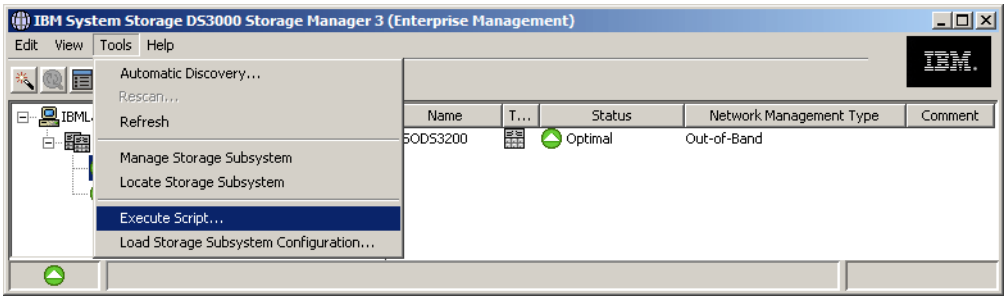


Figure C-1 Invoke Script Editor

2. A blank Script Editor appears, as shown in Figure C-2. From here you can enter commands and save them for later execution.



Figure C-2 Script Editor window

3. Select **File** → **Save Script** and **File** → **Load Script** to save and retrieve script files that you create, respectively. The Tools menu includes options to verify the syntax and run the scripts, as shown in Figure C-3 on page 749.

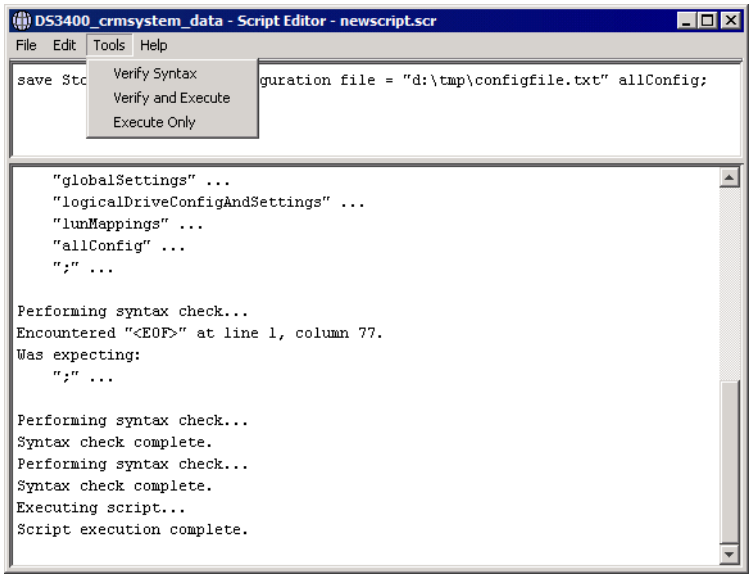


Figure C-3 Tools menu

4. The Script Editor contains an online help that describes the use of each of the available commands. Select **Help** → **Command Reference** to display the online help (Figure C-4) or **Help** → **Overview** for help on using the Script Editor itself.

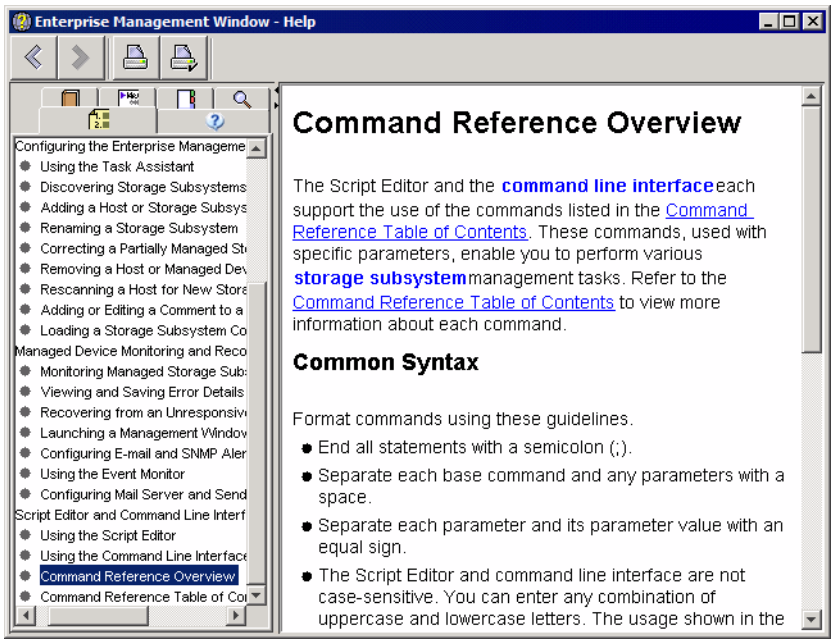


Figure C-4 Command reference

General SMcli syntax

You can either start SMcli from the client subdirectory, or add the directory to your path to execute SMcli from any window.

To list the full SMcli syntax, run **SMcli -?**, as shown in Example C-1.

Note: The PATH environment variable in Linux is modified during the installation of the Storage Manager to allow the execution of CLI and the GUI from everywhere in the system. On Microsoft Windows Systems, this is not the case.

Example: C-1 SMcli syntax

```
amazon:~ # SMcli -?
SMcli <DNS-network-name-or-IP-address>
        [<DNS-network-name-or-IP-address>]
        [-c "<command>;[<command2>;...]"]
        [-n <storage-array-name> | -w <WWID>]
        [-o <outputfile>] [-p <password>] [-e] [-S] [-quick]
SMcli <DNS-network-name-or-IP-address>
        [<DNS-network-name-or-IP-address>]
        [-f <scriptfile>]
        [-n <storage-array-name> | -w <WWID>]
        [-o <outputfile>] [-p <password>] [-e] [-S] [-quick]
SMcli {-n <storage-array-name> | -w <WWID>}
        [-c "<command>;[<command2>;...]"]
        [-o <outputfile>] [-p <password>] [-e] [-S] [-quick]
SMcli {-n <storage-array-name> | -w <WWID>}
        [-f <scriptfile>]
        [-o <outputfile>] [-p <password>] [-e] [-S] [-quick]
SMcli -d [-i] [-s] [-w] [-v] [-S]
SMcli -A [<DNS-network-name-or-IP-address1> [DNS-network-name-or-IP-address2]]
[-S]
SMcli -X (-n <storage-array-name> | -w <WWID> | -h <hostName>)
SMcli -m <ip address> -F <email address> [-g <contactInfoFile>] [-S]
SMcli -x email:<email address>
        [<hostname or IP address1> [<hostname or IP address2>]]
        [-n <storage-array-name> | -w WWID | -h <hostName> | -r (direct_sa |
host_sa)]}
        [-S]
SMcli -a email:<email address>
        [<hostname or IP address1> [<hostname or IP address2>]]
        [-n <storage-array-name> | -w <WWID> | -h <hostName> | -r (direct_sa |
host_sa)]}
        [-I <informationToInclude>] [-q <frequency>] [-S]
SMcli {-a | -x} trap:<community>,<hostname or IP address>
        [<hostname or IP address1> [<hostname or IP address2>]]
        [-n <storage-array-name> | -w <WWID> | -h <hostName> | -r (direct_sa |
host_sa)}
        [-S]
SMcli -?

For additional information, refer to your Command Line Interface documentation

SMcli completed successfully.
amazon:~ #
```

The various syntax options shown by **SMcli -?** perform the following tasks:

- Execute commands on a given storage subsystem that are specified on the command line over an out-of-band management connection.
- Execute commands that are specified in a command script over an out-of-band management connection.
- Execute commands on a given storage subsystem that are specified on the command line over an in-band management connection.
- Execute commands that are specified in a command script over an in-band management connection.
- List defined storage subsystems.
- Add storage subsystems to the configuration file of the Enterprise Management Window.
- Remove an already defined e-mail alert.
- Configure e-mail alerts.
- Configure or remove SNMP alerts.

As we said, the CLI can perform all the functions provided by the Storage Manager GUI. In the GUI, the first four tasks are performed in the subsystem management windows of a specific subsystem and the last five tasks (“List defined storage subsystems” and onwards) can be performed in the Enterprise Management Windows of the client

Note: Always specify the IP addresses or host names from the management interface of all installed DS3000 controllers.

Table C-1 gives an explanation of the various parameters that can be specified with SMcli commands.

Table C-1 SMcli parameters

Option	Description
<DNS-network-name-or-IP-address>	IP address or fully qualified domain name of the first and second controller management port. IP address fully qualified domain name of the host running the Storage Manager agent for in-band management. Use either IP addresses or host names, the subsystem name, or WWID
-A -A <IP C1> <IP C2>	Use this parameter to add a storage array to the configuration files. If you do not follow the -A parameter with a host name or IP address, auto-discovery scans the local subnet for storage arrays.
-X	Remove a storage subsystem specified by the name (-n) or the WWID (-w) or the host or IP address.
-a	Use this parameter to add a Simple Network Management Protocol (SNMP) trap destination or an e-mail address alert destination. - When adding an SNMP trap destination, the SNMP community is automatically defined as the community name for the trap and the host is the IP address or Domain Name Server (DNS) host name of the system to which the trap should be sent. - When adding an e-mail address for an alert destination, the e-mail address is the e-mail address where you want the alert message to be sent.

Option	Description
-g <filename>	Specify a plain text file that contains information about the subsystem or contact person with a phone or pager number.
-h <hostname>	Host name of the host running the storage manager agent for in-band management. Use this parameter when alert destinations are defined.
-I event Only -I profile -I supportBundle	Define what information is included in an alert that gets sent through e-mail. Use one of these options: ► eventOnly ► profile ► supportBundle By default, only the event data is sent.
-r <direct_sa | host_sa>	Define the alert recipient for a group of DS3000 servers depending on their management type. direct_sa represents those subsystems that are managed over an outbound connection (Ethernet) and host_sa subsystems represents those subsystems that are managed over a host agent (I/O path).
-x	Use this option to remove already defined SNMP or e-mail alert recipients.
-m <mail server>	Specify the mail server to be used to send out e-mails. Use this parameter together with the -F parameter.
-F <email addr>	The e-mail address that will be used in the sent by field of an e-mail alert. Responses to an alert will be sent to this address. Will be used for all e-mails. Use this parameter together with the parameter -m.
-n <subsystem name>	The name of the storage subsystem shown by the -d options. Use either IP or host names, subsystem name, or WWID.
-w <WWID>	The world wide identifier of the subsystem. Use option -d -w to discover it.
-c “<command;>”	Specify one or more commands separated by colons. Each command needs to closed by a semicolon. Use either option -c or -f.
-f <scriptfile>	Specify a script that contains the commands instead of entering them on the command line. Each line of the script can contain one command closed by a semicolon. Use either option -c or -f.
-o <output file>	Specify the file that will contain the output. Errors will be sent to the standard error output, which is usually the console. Use redirection when required. The output will not appended to the output file. This is an optional option.
-p <password>	The password of the storage subsystem to perform management tasks. This is an optional option that is only required when a password was defined for the specified storage subsystem. By default, no password is defined and this option is not required.
-d	Show the currently defined storage subsystems. More details about defined storage subsystems can be discovered with the parameters -i, -s, -w and -v.
-i	This option should only be used with parameter -d. It shows IP addresses instead of host names.
-s	This option should only be used with parameter -d. It shows the defined alert recipients.
-w	This option should only be used with parameter -d. It shows the world wide identifier of the defined storage subsystems.
-v	This option should only be used with parameter -d. It shows the health status of the defined storage subsystems.

Option	Description
-e	Do not check script syntax. This is an optional option.
-S	Do not list progress information. This is an optional option.
-quick	Use this parameter to reduce the amount of time that is required to run a single-line operation. An example of a single-line operation is the **recreate snapshot volume** command. This parameter reduces time by not running background processes for the duration of the command. Do not use this parameter for operations involving more than one single-line operation. Extensive use of this command can overrun the controller with more commands than the controller can process, which causes operational failure. Also, status and configuration updates normally collected from background processes will not be available to the CLI. This parameter causes operations that depend on background information to fail.

We now show some specific examples of using the CLI for tasks that have already been demonstrated using the GUI earlier in this book.

Adding a storage subsystem to the Storage Manager configuration

Adding storage subsystems to the Storage Manager configuration can be done manually or with an automated discovery as follows:

1. Upon installation, Storage Manager has no subsystems defined. Use option -d of SMcli to list the defined subsystems. See “Showing defined subsystems in the Storage Manager configuration” on page 754 for a detailed description of the command shown in Example C-2.

Example: C-2 SMcli - List storage subsystems defined in the Storage Manager configuration

```
amazon:/ # SMcli -d
There are currently no storage subsystems listed in the configuration file. Add
storage subsystems using the Add Storage Subsystem option in the storage
management
software or by command line.

SMcli failed.

amazon:/ #
```

2. To add a storage subsystem, use either option -A alone, to perform automatic discovery of available subsystems, or specify the controller’s management interface IP addresses, to perform manual discovery. Example C-3 shows both methods.

Example: C-3 SMcli - Add storage subsystems to the Storage Manager configuration

```
amazon:/ # SMcli -A 172.18.4.1 172.18.4.2
New storage subsystem was discovered at address 172.18.4.1.

New storage subsystem was discovered at address 172.18.4.2.

SMcli completed successfully.
```

```
amazon:~ # SMcli -A
Starting auto discovery.
.....................................................................
.........
Auto discovery operation successful.
SMcli completed successfully.

amazon:~ #
```

Showing defined subsystems in the Storage Manager configuration

Option -d of the SMcli gives a list of defined storage subsystems.There are additional parameters that can be used together with option -d to display more details about the defined storage subsystems. These are:

[-i] Show IP addresses instead of host names.

[-s] Show alert settings.

[-w] Show world wide identifier.

[-v] Show status of the defined subsystems in the configuration file.

[-S] Do not list progress information.

We use all these parameters in Example C-4. The output shows that there is a global alert in place that sends e-mail alerts to e-mail address dsx000rivers.local. The mail server nile.rivers.local is used and a return address ds3000@nile.rivers.local is shown in the [MAIL SERVER] section of the output. In addition to that alert, alerts are enabled for the in-band managed storage subsystems DS3200i and DS3400i.

SNMP alerts are shown for the DS3200. Public is used as the community string and the SNMP receiver is nile.rivers.local.

The storage array table contains a list of all defined storage subsystems in the configuration file with their name, the world wide identifier, host names, or IP addresses of the controllers and the status of the subsystem. This information can be found as well in the Enterprise Management window of the Storage Manager GUI.

Example: C-4 Displaying details about defined storage subsystems

```
amazon:~ # SMcli -d -i -s -w -v -S
[MAIL SERVER]
nile.rivers.local ds3000@nile.rivers.local

[ALERT SETTINGS - DEVICE GROUPS]

All storage subsystems
        DSx000@nile.rivers.local

All out-of-band storage subsystems
        <None>

All in-band storage subsystems
        colorado@nile.rivers.local

[ALERT SETTINGS - SPECIFIC DEVICES]
```

```
Storage arrays:

DS3200  600a0b80001d2bf40000000046265018        172.18.2.1      172.18.2.2      Optimal
        public,nile.rivers.local

DS3400  600a0b80002fc1d5000000004654a5a1        172.18.4.1      172.18.4.2 Unresponsive
        <None>

DS3400i 092ca896030303030303030303030303        172.18.0.2      Optimal
        <None>

DS3200i 092ca896040404040404040404040404        172.18.0.2      Needs Attention
        <None>


amazon:~ #
```

Configuring alerts

This section describes how to manage alert recipients.

Defining the mail server and e-mail address to send out the e-mail alerts

To define the global setting for the e-mail server and e-mail return address, use the following command:

```
SMcli -m <IP or host name of mail server> -F <email address>
```

See Example C-5 for more details.

To verify the setting, use **`SMcli -d -s`** to see the current settings.

Example: C-5 SMcli - define email server and return email address

```
amazon:~ # SMcli -m nile.rivers.local -F DSx000@nile.rivers.local
SMcli completed successfully.


amazon:~ #
```

To delete these settings, use empty strings, as shown in Example C-6.

Example: C-6 SMcli - delete mail server and return email address setting

```
amazon:~ # SMcli -m "" -F ""
SMcli completed successfully.


amazon:~ #
```

Defining email alert recipients

Alerts can be defined for a single DS3000 storage subsystem, all out-of-band managed storage subsystems, all in-band managed storage subsystems, or all subsystems managed by a single management station. This section describes those settings.

Note: An e-mail server should be defined, as shown in “Defining the mail server and e-mail address to send out the e-mail alerts” on page 755. Without a mail server, the local system running the Storage Manager will be used to send e-mail alerts, which might not always be possible.

Example C-7 shows how to define an e-mail alert recipient for a single DS3200 server.

Example: C-7 SMcli - define e-mail alert recipient for single system

```
amazon:~ # SMcli -a email:ds3200@nile.rivers.local -n "DS3200"
SMcli completed successfully.


amazon:~ #
```

To define an e-mail recipient for an in-band or out-of-band managed storage subsystem, specify the parameter -r followed by direct_sa or host_sa for out-of-band or in-band managed systems, respectively. It is not necessary to use a parameter to specify a storage subsystem here. Example C-8 shows an e-mail recipient for all in-band managed storage subsystems.

Example: C-8 SMcli - E-mail recipient for a group of storage subsystems

```
amazon:~ # SMcli -a email:ds3000-inband@nile.rivers.local -r host_sa
SMcli completed successfully.


amazon:~ #
```

To send alerts generated by any single subsystem that is defined in the configuration file of Storage Manager, use the command shown in Example C-9.

Example: C-9 SMcli - E-mail alert recipient for all subsystems

```
amazon:~ # SMcli -a email:ds3000@nile.rivers.local -r host_sa
SMcli completed successfully.


amazon:~ #
```

Global e-mail alert recipients are configured by omitting the target specification like -r host_sa or a host name. Example C-10 Shows a global e-mail recipient configuration.

Example: C-10 SMcli - configure global email recipient

```
amazon:~ # SMcli -a email:dsx000@nile.rivers.local
SMcli completed successfully.

amazon:~ #
```

Deleting e-mail alert recipients

To delete an e-mail alert recipient, use option -x and the same syntax that was used to add the e-mail alert recipient. Example C-11 shows the commands to remove e-mail recipients from the configuration. The first command will show an example of a configured management station system and all the following commands will remove the e-mail recipients that are defined in this sample configuration beginning with a dedicated storage subsystem, followed by in-band and out-of-band managed storage subsystems, and last but not least, the global e-mail recipient.

Example: C-11 SMcli - Delete e-mail recipients

```
amazon:~ # SMcli -d -s
[MAIL SERVER]
nile.rivers.local ds3000@nile.rivers.local

[ALERT SETTINGS - DEVICE GROUPS]

All storage subsystems
        global@nile.rivers.local

All out-of-band storage subsystems
        out-band@nile.rivers.local

All in-band storage subsystems
        in-band@nile.rivers.local

[ALERT SETTINGS - SPECIFIC DEVICES]

Storage arrays:

DS3200  ds3200-a.rivers.local   ds3200-b.rivers.local
        <None>
DS3200p ds3300-a.rivers.local   ds3300-b.rivers.local
        DS3300@nile.rivers.local

DS3400  ds3400-a.rivers.local   ds3400-b.rivers.local
        <None>

DS3400i colorado.rivers.local
        host@nile.rivers.local

DS3200i colorado.rivers.local
        host@nile.rivers.local


SMcli completed successfully.

amazon:~ # SMcli -x email:ds3300@nile.rivers.local -n DS3300
SMcli completed successfully.

amazon:~ # SMcli -x email:in-band@nile.rivers.local -r host_sa
SMcli completed successfully.

amazon:~ # SMcli -x email:out-band@nile.rivers.local -r direct_sa
SMcli completed successfully.
```

```
amazon:~ # SMcli -x email:host@nile.rivers.local -h colorado.rivers.local
SMcli completed successfully.

amazon:~ # SMcli -x email:global@nile.rivers.local
SMcli completed successfully.

amazon:~ #
```

SNMP alert recipients

SNMP alert recipients are handled very similarly to e-mail alert recipients, but the trap parameter is used instead of the e-mail parameter. To specify the recipient, use an SNMP community string followed by the SNMP trap receiver systems IP address or host name. Example C-12 shows how to configure SNMP trap recipients for a single DS3400 system, in-band and out-of-band managed DS3000 systems, and a global trap recipient.

Example: C-12 SMcli - SNMP alert recipient

```
amazon:~ # SMcli -a trap:ds3400,nile.rivers.local -n ds3400
SMcli completed successfully.

amazon:~ # SMcli -a trap:in-band,nile.rivers.local -r host_sa
SMcli completed successfully.

amazon:~ # SMcli -a trap:out-band,nile.rivers.local -r direct_sa
SMcli completed successfully.

amazon:~ # SMcli -a trap:global,nile.rivers.local
SMcli completed successfully.

amazon:~ #
```

Issuing commands to the storage subsystem

Commands can be issued to one of the defined storage subsystems by using the -c option. Use the following syntax to issue the commands:

```
SMcli {<DNS-network-name-or-IP-address> [<DNS-network-name-or-IP-address>] | -n
<Subsystem Name> | -w <WWID> }
{[-c "<command>;[<command2>;...]"] | [-f <scriptfile>]}
[-o <outputfile>] [-p <password>] [-e] [-S] [-quick]
```

See “General SMcli syntax” on page 750 for information about the meaning of each option.

The parameter -c requires that all commands be enclosed in single or double quotes. Each command must be terminated by a semicolon. Multiple commands can follow on one command line. Example C-13 on page 759 shows a command to a storage subsystem to get the health status of that storage subsystem.

Example: C-13 SMcli - Command

```
amazon:/ # SMcli -n DS3200 -c "show storagesubsystem healthstatus;" -S
The following failures have been found:
Failed Drive - Unassigned or Hot Spare
Storage Subsystem: DS3200
  Enclosure: Controller/Drive enclosure
    Affected drive slot(s): 1
    Service action (removal) allowed: Yes
    Service action LED on component: Yes

amazon:/ #
```

An alternative to specifying multiple commands in one line is to create a script file that contains each command on a single line. You can do this using the Script Editor (which allows you to validate and execute the commands), or in your favorite text editor.

In the script file, each command must be terminated with a semicolon. Quotation marks are not necessary. Example C-14 shows a script file that creates some logical drives and assigns them to a host.

Example: C-14 Example of a command script file

```
create logicalDrive drives=(7,2 7,3) RAIDLevel=1 userLabel="Hudson-1"
segmentSize=64 \ capacity=10 GB owner=a usageHint=fileSystem  mapping=none;

create logicalDrive drive [7,4] RAIDLevel=0 userLabel="Hudson-2" segmentSize=64 \
capacity=8 GB owner=b usageHint=fileSystem  mapping=none;

set logicalDrive ["Hudson-1"] logicalUnitNumber=0 host="Hudson";

set logicalDrive ["Hudson-2"] logicalUnitNumber=1 host="Hudson";
```

To execute this script file (outside of Script Editor), save it (for example, createDrive.cmd) and run it as shown in Example C-15.

Example: C-15 SMcli executes the script

```
amazon:~ # SMcli -n DS3200 -f createDrive.cmd -o hudson.out -S
amazon:~ #
```

Some commands require the input of string values, such as the command shown in Example C-16.

Example: C-16 SMcli - Character masking

```
amazon:~ # SMcli -n DS3400 -c "set logicalDrive [\"Orinoco-1\"]
userLabel=\"Orinoco-2\";" -S -e
amazon:~ #
```

This command will rename the logical drive Orinoco-1 into Orinoco-2. The parameter logicalDrive requires the name of a logical drive specified in square brackets. Because the actual logical drive name contains a number, double quotation marks are required. To avoid conflicts with the shell running this command, it is sometimes necessary to mask some characters. In this example, the double quotation marks of the logical drive name are masked with a backslash.

How you mask special characters depends on the shell and the operating system being used. The shell documentation contains more details about character masking. An alternative that is not available for Microsoft Windows would be the use of single quotation marks around the complete command and double quotation marks without the backslashes.

Here is a list of commands that can be used to manage a DS3000 storage subsystem:

- **accept**
- **activate**
- **autoConfigure**
- **check**
- **clear**
- **create**
- **deactivate**
- **delete**
- **diagnose**
- **disable**
- **download**
- **enable**
- **recopy**
- **recover**
- **recreate**
- **remove**
- **repair**
- **reset**
- **resume**
- **revive**
- **save**
- **set**
- **show**
- **start**
- **stop**
- **suspend**

The CLI provides help when a command is not entered completely. Follow the Script Editor online help or the CLI online help to get details about each command.

Sample command: Save configuration script file

Use the **save StorageSubsystem configuration** command to create a script containing the configuration statements for your environment. A sample run of this command is shown in Example C-17.

Example: C-17 CLI command to save the storage subsystem configuration

```
save StorageSubsystem configuration file="DS3400_1.cmds" allConfig;
```

This creates a file of CLI commands that create the array, logical drives, hosts, and host mappings. Our file is shown in Example C-18.

Example: C-18 Sample storage subsystem configuration file

```
// Logical configuration information from Storage Subsystem DS3400_crmsystem_data.
// Saved on 21 August, 2008
// Firmware package version for Storage Subsystem DS3400_crmsystem_data =
07.35.00.00
```

```
// NVSRAM package version for Storage Subsystem DS3400_crmsystem_data =
N1726D340R917V15

//on error stop;

// Uncomment the two lines below to delete the existing configuration.
//show "Deleting the existing configuration.";
//set storagesubsystem resetConfiguration = true;



// Storage Subsystem global logical configuration script commands
show "Setting the Storage Subsystem user label to DS3400_crmsystem_data.";
set storagesubsystem userLabel="DS3400_crmsystem_data";

show "Setting the Storage Subsystem media scan rate to 18.";
set storagesubsystem mediaScanRate=18;

// Uncomment the three lines below to remove default logical drive, if exists,
script command. NOTE: Default logical drive name is always = "" (blank).
//on error continue;
//show "Deleting the default logical drive created during the removal of the
existing configuration.";
//delete logicaldrive[""];
//on error stop;

// Copies the hot spare settings
// NOTE: These statements are wrapped in on-error continue and on-error stop
statements to
// account for minor differences in capacity from the drive of the Storage
Subsystem on which the
// configuration was saved to that of the drives on which the configuration will
be copied.
//on error continue;
show "Creating hot spare at Enclosure 0 Slot 1.";
set drive[0,1] hotSpare=true;
//on error stop;

show "Setting the Storage Subsystem cache block size to 4.";
set storagesubsystem cacheBlockSize=4;

show "Setting the Storage Subsystem to begin cache flush at 80% full.";
set storagesubsystem cacheFlushStart=80;

show "Setting the Storage Subsystem to end cache flush at 80% full.";
set storagesubsystem cacheFlushStop=80;

// Creating Host Topology
show "Creating Host Salza.";
create host userLabel="Salza";

show "Creating Host Port Salza0 on Host Salza with WWN 10000000c93542c2 and Host
Type Index 2.";
// This Host Type Index corresponds to Type Windows 2000/Server 2003 Non-Clustered
```

```
create hostPort host="Salza" userLabel="Salza0" identifier="10000000c93542c2"
hostType=2;

show "Creating Host Port Salza1 on Host Salza with WWN 10000000c93542c3 and Host
Type Index 2.";
// This Host Type Index corresponds to Type Windows 2000/Server 2003 Non-Clustered
create hostPort host="Salza" userLabel="Salza1" identifier="10000000c93542c3"
hostType=2;


show "Creating RAID 1 Logical Drive San_boot_salza on new Array 1.";
// This command creates the Array and the initial Logical Drive on that array.
// NOTE: For Arrays that use all available capacity, the last Logical Drive on
this array is
// created using all remaining capacity by omitting the capacity= logical drive
creation parameter.
create logicaldrive drives[0,2 1,1] raidLevel=1 userLabel="San_boot_salza" owner=A
segmentSize=128 capacity=21474836480 Bytes;
show "Setting additional attributes for Logical Drive San_boot_salza.";
// Configuration settings that can not be set during Logical Drive creation.
set logicaldrive["San_boot_salza"] cacheFlushModifier=10;
set logicaldrive["San_boot_salza"] cacheWithoutBatteryEnabled=false;
set logicaldrive["San_boot_salza"] mirrorEnabled=true;
set logicaldrive["San_boot_salza"] readCacheEnabled=true;
set logicaldrive["San_boot_salza"] writeCacheEnabled=true;
set logicaldrive["San_boot_salza"] mediaScanEnabled=true;
set logicaldrive["San_boot_salza"] redundancyCheckEnabled=false;
set logicaldrive["San_boot_salza"] readAheadMultiplier=1;
set logicaldrive["San_boot_salza"] modificationPriority=high;
show "Creating Logical Drive-to-LUN Mapping for Logical Drive San_boot_salza to
LUN 0 under Host Salza.";
set logicaldrive ["San_boot_salza"] logicalUnitNumber=0 host="Salza";
```

We show how to load this file to create the configuration on another system in “Load Storage Subsystem Configuration option” on page 116. Note that loading the configuration file will overwrite any existing data on the array; you should only do this on a new array.

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

For information about ordering these publications, see “How to get IBM Redbooks publications” on page 764. Note that some of the documents referenced here may be available in softcopy only.

- *IBM BladeCenter Products and Technology, SG24-7523*
- *IBM i and Midrange External Storage,* SG24-7668
- *IBM System Storage DS4000 and Storage Manager V10.10, SG24-7010*
- *Implementing IBM Director 5.20, SG24-6188*
- *Introduction to Storage Area Networks*, SG24-5470

Other publications

These publications are also relevant as further information sources:

- *Advanced Management Module and Management Module User's Guide - IBM BladeCenter*, MIGR-45153

Online resources

These Web sites are also relevant as further information sources:

- Brocade

  http://www.brocade.com
- Emulex

  http://www.emulex.com/
- Download IBM System Storage DS3000 Storage Manager

  http://www.ibm.com/servers/storage/support/disk
- IBM Disk Support Web site

  http://www.ibm.com/servers/storage/support/disk
- IBM Server Support Web site

  http://www.ibm.com/systems/support/
- Implementing the IBM BladeCenter Boot Disk System Solution

  http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101198

© Copyright IBM Corp. 2007, 2009. All rights reserved.

- Implementing Microsoft Cluster Server (MSCS) and Windows 2000/2003 Boot from IBM DS4000 SAN with Blade Servers

  http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101052

- Interoperability Guide - IBM BladeCenter

  https://www-304.ibm.com/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-5073016&brandind=5000020

- Linux RDAC download

  http://www.lsi.com/rdac/ds3000.html

- Microsoft Support Web site for hotfixes

  http://support.microsoft.com

- QLogic

  http://www.qlogic.com

- Red Hat network

  https://rhn.redhat.com

- VMware

  http://www.vmware.com

How to get IBM Redbooks publications

You can search for, view, or download IBM Redbooks publications, IBM Redpaper publications, Technotes, draft publications, and additional materials, as well as order hardcopy Redbooks, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

A

access logical drive 30, 397
accessibility 281
Active Directory® 627
add drives 234
add free capacity 234
add HBA 227
add storage subsystem 127, 753
Advanced Support Tasks 334
alert notifications 335
alerts 211, 335, 755
arbitrated loop 4
array 55, 172, 218
  locate 264
automatic configuration 156
automatic host access configuration 142
automatic subsystem discovery 68, 126
automatically discover new storage subsystems 125

B

base logical drive 183, 185, 190, 198, 200, 254, 728
  data blocks 186
  full box 186
  inaccurate point-in-time image 728
  multiple FlashCopies 183
basic disk 728
Basic/Regular Disc 726, 728
battery age 281
bidirectional authentication 285
BladeCenter
  boot from SAN 592
BladeServer
  DS3400 clustered configuration 611
  DS3400 configuration 585, 671, 711

C

cabling 35
capacity 215
Change Array RAID Level 237
Change Automatic Collection Setting 303
change enclosure ID 272
change media scan settings 272–273
change network configuration 278
change VolumeCopy priority 255
CLI 198, 728, 747
CLI script 612
click 202
Cluster Administrator 634
clustered disk access 58
color and font settings 281
compatible drive 322
concurrent disk access 58
configure alerts 118
configure host access 132, 396
Configure tab 131
configured hosts 216
controller
  manage 330
  online/offline 331
  ownership 241
controller firmware 196, 313, 327
controller settings 278
copy 186
Copy Manager 200, 205, 250–251, 254–255
  remove copy pair 255
copy pair 250
  copy priority 255
  remove 255
copy priority 202
copy priority rate 200
Copy wizard 201–202
copy-on-write 184, 186, 243, 245
create host group 150
create host-to-logical-drive mappings 160
create logical drive 650, 686
create VolumeCopy 201

D

DACstore 324
Data 190
data access pattern 56
data block 184
data striping 51
dedicated parity drive 52
defragment 186, 199
delete logical drive 232
Detailed 728
Device Manager 395
Device-Specific Module 104
Diagnostic Data Capture 15
different FlashCopy logical drive 245
direct-attached management 26
directory agent 11
disable FlashCopy logical drive 243
disk attachment 3
  FC 4
  SAS 5
disk partition 403, 427
disk signature 192
disk system 18
distributed parity drive 53
drive capacity 215
drive mapping 177
drive mode 213
drives 213
  add 234
  firmware 319

© Copyright IBM Corp. 2007, 2009. All rights reserved.

locate 265
DS family 18
positioning 18
DS3000 18
access logical drive 30, 397
accessibility 281
add drives 234
add free capacity 234
add storage subsystem 68, 753
Administration Configure tab 131
alerts 118, 211, 335
array 55, 172, 218
automatic configuration 156
automatic host access configuration 142
automatic subsystem discovery 68
battery 281
boot from SAN 593
cabling 35
capacity 215
CLI 747
CLI script 612
clustered disk access 58
command-line interface 28, 747
concurrent disk access 58
configure alerts 118, 755
configure disk access for Linux 411, 444
configure host access 132, 141
controller 330
controller clock 278
controller firmware 313
controller settings 278
create FlashCopy 187
create host-to-logical-drive mappings 160
define host server 646, 682
direct-attached management 26
disk expansion 22
drive capacity 215
drive firmware 319
drive mapping 177
drives 213
DSM 104
dual controller 38, 241
dynamically add logical drive 29
e-mail alerts 120
enable second I/O path 665, 698
enclosure ID 36, 272
enclosure loss protection 57, 172
ESM firmware 311
Ethernet management port 19, 26, 40, 76
event log 335, 337
Event Monitor Service 118
EXP3000 attachment 37
EXP3000 Expansion License 26
FC attachment 22
FlashCopy 25, 183, 243
FlashCopy Expansion License 25
FlashCopy premium feature 187
FlashCopy repository 25, 247
FlashCopy tips 190
FlashCopy VolumeCopy License 26
free capacity 234
hardware components 212
help 349
heterogeneous hosts 153
host 216
host access 141, 396
host attachment 40
host definition 132
host group 132, 149–150
host to logical drive mapping 151, 177
host-attached management 26
hot spare 57, 164, 215
hot_add 29, 401, 426
IBM Support Web site 734
in-band management 26, 30
incorrect SAS cabling 39
Information Center 219
Initial Setup 112
Initial Setup Tasks wizard 50, 71
install multipath driver 600
install Storage Manager 81
issue commands 758
Linux disk access 411, 444
Linux RDAC 663
list storage subsystems 754
load configuration 116
locate array 218
locate components 262–263
locate storage subsystem 72
locate subsystem 115
logical drive 51, 55, 168, 646, 682
logical drive capacity 215
logical drive copies 250
logical drive I/O characteristics 176
logical drive preferred path 333
lost connection 211
manage controllers 330
management methods 30
manual configuration 163
manual host access configuration 145
map host 400
map host to logical drive 160
media scan 272–273
modify storage subsystem 231
MPP 29
multipath driver 600, 617
multipath support 29
network configuration 76, 278
NVSRAM 313, 317
online help 219, 349
out-of-band management 26, 31
partially managed storage subsystems 125
Partition Expansion License 25, 59
partitioning 58
password 73, 260
password protection 40
power cables 45
powering off 48
powering on 47
preferred controller 241

preferred controller ownership 199
preferred path 656, 692
premium features 74, 266
products 19
RAID controller 19–20
RAID level 51, 55, 172
RDAC 29, 90, 408, 442, 663
Recovery Guru 39, 347–348
redundancy check 274
remove mapping 230
remove storage partitions 229–230
rename storage subsystem 72
rename subsystem 259
sample configurations 389
SAS attachment 20
SAS cabling 23, 36
save configuration script file 760
Script Editor 115, 748
script file 759
second I/O path 665, 698
security 40, 260
segment size 56
serial management port 19
set password 73
single controller 38
SMagent service 396
SMcli 750
SNMP 335
SNMP alerts 122
SNMP traps 118
Software Feature Pack 267
status 210
status icons 114
Storage Manager see DS3000 Storage Manager
storage partitions 25, 58, 152, 217, 229–230
storage subsystem profile 212, 304
storage subsystem status 210
subsystem status 114
support information 301
Universal XPort device 31
VolumeCopy 25, 196
VolumeCopy permission 253
VolumeCopy Premium Feature 196, 201
VolumeCopy priority 255
DS3000 Controller Firmware Upgrade Tool 324
DS3000 Storage Manager 26–27, 201, 335
accessibility 281
add drives 234
add free capacity 234
add HBA 227
add storage subsystem 68, 127
alert status 211
alerts 118
array 172
automatic subsystem discovery 126
automatically discover subsystems 125
battery age 281
change enclosure ID 36, 272
change media scan settings 272–273
change network configuration 76, 278
change password 260
change preferred path 242
change VolumeCopy priority 255
check premium features 201
CLI 28
color and font settings 281
components 27
configure alerts 118
configure host access 141, 396
configure host group 150
configure hot spare 164
configure storage subsystem 156, 163
Configure tab 131, 221
configured hosts 216
connection lost 211
controller firmware 313
controller settings 278
controllers 330
Copy Manager 250
create FlashCopy 187
create host group 150
create logical drive 168, 650, 686
create VolumeCopy 201
define a host 145
define array 172
define host group 149
delete logical drive 232
disable FlashCopy logical drive 243
drive capacity 215
drive firmware 319
drive mapping 177
drive mode 213
drives 213
DSM 104
edit topology 223
e-mail alerts 120
enable Premium Features 266
enclosure ID 272
Enterprise Management Window 27, 112–113, 199
ESM firmware 311
event log 335, 337
Event Monitor Service 118
execute script 115
FlashCopy repository 247
FlashCopy wizard 188
free capacity 234
gather support information 301
GUI 27, 198
hardware components 212
help 349
host access 141, 396
host group 150
host server 26
host to logical drive mapping 228
host type 228
hosts and mappings 216
hot spare 215
hot_add 29
in-band management 26, 29–30
Information Center 219

Initial Setup 112
Initial Setup Tasks 208
Initial Setup Tasks wizard 50, 71
install 81, 395, 407, 425, 441, 472, 697
install on Linux 90
install on Windows 82
installation 59
Java Access Bridge 64
load storage subsystem configuration 116
locate array 218
locate components 262–263
locate drive 214
locate hot spare 215
locate storage subsystem 115
logical drive capacity 215
logical drive preferred path 333
lost connection 211
manage controllers 330
manage logical drive copies 250
management station 26
map drive to host 177
map host to logical drive 153, 160
media scan settings 272–273
modify FlashCopy logical drives 243
modify host configuration 223
move HBA 225
move host 224
network configuration 76, 278
NVSRAM 313, 317
online help 219, 349
operations in progress 210
out-of-band management 26, 31
partially managed storage subsystems 125
password 260
preferred path 333
Premium Features 266
quick access buttons 125
RDAC 105
Recovery Guru 48, 199, 347–348
recreate FlashCopy logical drive 245
redistribute logical drives 333
redundancy check 274
remove host 226
remove mapping 230
remove storage partitions 229–230
remove storage subsystem 117, 128
remove VolumeCopy pair 255
rename host 226
rename logical drive 231
rename storage subsystem 122, 259
replace HBA 227
reset battery age 281
resize FlashCopy repository 247
Script Editor 28, 115
security 260
set password 73, 260
SMagent 29, 144
SMcli 28
SMdevices 29
SMrepassist 29, 188
SMutil 29
SNMP alerts 122
SNMP traps 118
status 210
Status Box 210
status icons 114
stop VolumeCopy 252
Storage Manager Agent 29
Storage Manager Utility 29
storage partitions 229–230
storage subsystem profile 212, 305
storage subsystem status 210
Subsystem Management 113
Subsystem Management Window 28
subsystem status 114
Summary tab 207
support information 301
suspend media scan 274
synchronize controller clock 278
Task Assistant 112, 127
Tools tab 258
topology 223
view drives 213
view event log 335, 337
view storage subsystem profile 305
View/Enable Premium Features 74
VolumeCopy permission 253
VolumeCopy priority 255
DS3200 20
Linux configuration 405
SAS 2-Port Daughter Card 21
SAS host attachment 40
Windows configuration 391
DS3300
iSCSI host attachment 41
DS3400 22
boot from SAN 593
configure 591
FC host attachment 41
in-band management 29
Linux SAN boot configuration 639
Windows clustered configuration 611
Windows configuration 585, 671, 711
DS4000 18
overall positioning 18
RDAC 105
DS4000 Series
overall positioning 18
DS6000 18
DS8000 18
dual controller 241
Dynamic Capacity Expansion (DCE) 199
dynamic disk 192, 402, 426
Dynamic Logical Drive Expansion (DVE) 199
Dynamic RAID Level Migration (DRM) 199
Dynamic Reconstruction Rate (DRR) 199
Dynamic Segment Sizing (DSS) 199

E

edge expander 6

edit topology 223
e-mail alerts 120, 755
Emulex
  boot from SAN 592
Emulex BIOS 593
Emulex HBA BIOS 136
Emulex HBA device driver 597
Emulex HBA firmware 606
Emulex HBAnyware 602
enable Premium Features 266
enclosure
  locate 263
enclosure ID 36, 272
enclosure loss protection 57, 172
Enterprise Management Window 27, 112–113, 199
Environmental Services Module see ESM
ESM 22, 36
ESM firmware 311
ESS 18
Ethernet management port 40, 76
event log 199, 241, 335, 337
  critical entry 241
  informational entry 241
Event Monitor Service 118
existing FlashCopy repository logical drive
  storage capacity 247
EXP3000 22
  attachment 37
  enclosure ID 36
  enclosure loss protection 57
  ESM 22, 36
  expansion license 26
  power supply module 23
  SAS expansion enclosure 19
EXP3000 Expansion License 25
Export Array 340

F

failover 241
failure policy 186
fan-out expander 6
Fast!UTIL 641, 673, 681
FAStT Storage Server 327
FC
  disk attachment 4
  DS3400 cluster configuration 611
  DS3400 configuration 585, 671, 711
  DS3400 SAN boot configuration 639
  host attachment 41
  SFP 22
FC cables 5
FC protocol layers 5
FC topologies 4
FC World Wide Name 5
FC WWN 134
FC zoning 612
FCP 4
feature key 196
Fibre Channel
  hub 327
Fibre Channel see FC
file system 192, 202
firewall 379
FlashCopy 25, 183, 233, 243
  base logical drive 183
  create 187
  disable 243
  failure policy 186
  management overhead 185
  metadata 184
  recreate 245
  repository 25
  repository logical drive 185, 247
  suspend host I/O 187
  threshold exceeded 186
  tips 190
  using 195
  wizard 188
FlashCopy logical drive 183–186, 190, 192, 198, 200, 244–245, 254, 726, 728
  creating on Windows 729
  logical drive-to-LUN mappings 191
  original names 245
FlashCopy logical drive reuse 730
FlashCopy premium feature 187
FlashCopy repository logical drive 184–186, 190, 193, 245, 247
free capacity 199, 234, 247

G

gather support information 301

H

hardware address 36
hardware components 212
HBA BIOS 102
HBA configuration 102
HBA device driver 104
HBA firmware 132
HBA install 132
HBA multipath support 104
HBA port settings 134
HBA settings 643, 676
HBAnyware 602
help 349
heterogeneous hosts 153
host 132, 216, 228
  update DS3000 software 308
host access 141
host attachment 40
host bus adapter (HBA) 152, 241
host group 132, 149–150
host operating system 191, 195
  FlashCopy logical drive 195
host port 152
host server 26, 646, 682
host system 152, 325
host to logical drive mapping 151, 177, 228
host type 228

host-attached management 26
Host-to-Logical Drive Mappings 217
hot spare 57, 164, 215
hot_add 29, 401, 426

I

I/O activity 197, 199–200, 202
I/O path 26, 241
I/O request 202, 241
IBM My notifications 740
IBM Support Web site 733
IBM System Storage DS3000 see DS3000
IBM System Storage EXP3000see EXP3000
Import Array 344
in-band management 26, 29–30
in-band management (IBM) 327
incompatible drive 322
Information Center 219
Initial Setup Tasks 208
install DS3000 Storage Manager 81, 395, 407, 425, 441, 472, 697
install HBA 132
Internet Storage Name Server 11
IPv6 16
IQN 10
is 192
iSCSI
  host attachment 41
  IQN 10
  target authentication 285
iSCSI software initiator 532, 548
iSNS 290
  See Internet Storage Name Server

J

Java Access Bridge 29, 64

K

KB 185

L

Linux
  configure disk space 411, 444
  DS3000 Storage Manager 407, 441
  DS3200 disk access 408, 425, 441
  FC SAN boot configuration 639
  install DS3000 Storage Manager 90
  install for SAN boot 659
  MPP 29, 105
  RDAC 29, 105, 407–408, 441–442, 663
  SAS disk configuration 405
  SAS HBA driver 407, 434
list storage subsystems 754
load storage subsystem configuration 116
locate array 218
locate drive 214
locate DS3000 components 262–263
locate hot spare 215
locate storage subsystem 72, 115
Logical Disk Manager (LDM) 192
logical drive 51, 55, 152, 168, 185, 196, 198, 200, 202, 241, 251, 255, 646, 682, 728–729
  configure for boot 658, 695
  current owner 241
  delete 232
  I/O requests 241
  manage copies 250
  map host 400
  preferred controller ownership 241
  preferred ownership 241
  preferred path 333
  primary 196
  re-create option 244
  redistribute 333
  redundancy check 274
  rename 231
  stop VolumeCopy 252
logical drive capacity 215
logical drive I/O characteristics 176
Logical Drive Manager 726
logical drive mapping 153
logical drive ownership 656, 692
logical unit number (LUN) 152
Logical View 206
lost connection 211
lservice 361
LSI SAS HBA BIOS 138
LUN 31 31
LUN assignment 152
LUNs 152

M

Major Event Log (MEL) 357
management interface 359
management station 26, 327
manual configuration 163
manual host access configuration 145
map drive 177
map host 400
Mappings View 191
MBR 729
media scan 272–273
Microsoft Cluster Server 58
Microsoft iSCSI Software Initiator 9
Microsoft MPIO 104
modification operation 198, 200, 203, 251
modify FlashCopy logical drives 243
modify host configuration 223
modify storage subsystem 231
Monitor process 304
monitoring 241
move HBA 225
move host 224
MPIO 104, 618
MPP 29, 105, 410, 443
MSCS
  add node to cluster 634
  cluster service 629

create cluster 626
DS3400 configuration 611
multidrop cable 7
multi-mode fiber 5
multipath driver 600, 617
multipath support 29, 104

N

network configuration 76, 278
NVSRAM 313, 317, 325

O

one-time usage 728
online help 219, 349
operating system 195, 725
FlashCopy logical drives 725
operations in progress 210
out-of-band management 26, 31

P

p+q 14
parity 52
partially managed storage subsystems 125
partitions 15
password 73, 260
password protection 40
path 241
Persistent reservation 203
place controller online 331
point-in-time copy 197
point-in-time image 183
point-in-time replication 196
point-to-point 4, 7
power cables 45
preferred controller 241
preferred controller ownership 199
preferred path 656, 692
Premium Features 74, 266
Pre-Read Redundancy Check 16, 276
primary 196
Primary Contact 357
Process 726–727
profile 212

Q

QLogic
BIOS utility 133
HBA settings 643, 676
QLogic Fast!UTIL 102, 133
QLogic HBA BIOS utility 641, 673
QLogic HBA driver 614, 660
QLogic SANSurfer 102
quick access buttons 125
quiesce database application 188
Quorum Logical drive 629

R

RAID 0 51, 173
RAID 1 52, 173, 273
RAID 10 53–54, 172–173
RAID 3 52, 173, 273
RAID 5 53, 173, 273
RAID level 51, 172
RAID parity 52
RDAC 16, 29, 90, 105, 407–408, 441–442, 663
RDAC commands 670
Read Link Status (RLS) 357
read-only attribute 205, 250
Recovery Guru 48, 199, 325, 347–348
recreate FlashCopy logical drive 245
Redbooks Web site 764
Contact us xviii
redundancy check 274
remove host 226
remove mapping 230
remove storage partitions 229–230
remove storage subsystem 117, 128
remove VolumeCopy pair 255
rename host 226
rename logical drive 231
rename storage subsystem 72, 122, 259
Replace Drives 238
replace HBA 227
repository logical drive 247
reset battery age 281
resize FlashCopy repository 247
reuse FlashCopy logical drive 730
RLOGIN 357, 382
rservice 361
RSM
change password 361
firewall 379
management interface 359
remote access security 377
security 361
user ID 361
rsm-passwd 361

S

same base logical drive 244–245
sample configurations 389
SAN 4
boot operating system 593
SAN boot 639
SAN cabling 586, 612, 672
SAN switch 22
SAN zoning 132
SANsurfer 489
SAS 5
DS3200 configuration 391, 405
host attachment 40
incorrect cabling 39
SAS cabling 23, 36
SAS domain 6
SAS drive technology 8

SAS expander 6
SAS HBA 392
SAS HBA driver 393, 407, 434
SAS protocol layers 6
SAS wide port 7
SATA 7
save configuration script file 760
Save Support Information 302
Script Editor 28, 115, 728, 748
script file 759
SCSI
  compared to SAS 5
SCSI commands 31
second I/O path 665, 698
security 260
segment size 56
selected source logical drive
  new VolumeCopy 251
Serial Attached SCSI see SAS
Service Access Timeout 382
service agent 11
Service Location Protocol 11
  directory agent 11
  service agent 11
  user agent 11
SFP 22, 41
signalling overhead 7
signature 192
single-mode fibre 5
SLP
  See Service Location Protocol
SMagent 29, 144, 395, 425, 472
SMagent service 396
SMcli 28, 747–748, 750
SMclient 27, 395, 407, 425, 441
SMdevices 29
SMrepassist 29, 188
SMutil 29, 395, 407, 425, 441, 472, 697
SNMP 335, 357
SNMP alert recipient 758
SNMP alerts 122
SNMP trap 379
SNMP traps 118
Software Feature Pack 14, 267
source logical drive 25, 196, 198–200, 202, 206, 250, 255
  file systems 202
  preferred controller 199
spare drive 247
sparse LUNs 106
speculative change 728
standard logical drive 200
Status Box 210
status icons 114
stop VolumeCopy 252
storage capacity 247
Storage Manager Agent 29
Storage Manager Utility 29
storage partition 152, 217
storage partitioning 58
storage partitions 25, 152, 217, 229–230
  host systems 152
storage server 325
storage subsystem 186, 196, 198–200, 206, 243, 250, 252
  locate 263
  logical drive copies 250
  manage logical drive copies 250
storage subsystem password 260
storage subsystem profile 212, 304
storage subsystem status 210
Subsystem Contact 357
Subsystem Management Window 28
support bundle 15
support information 301
supports 152
suspend host I/O 187
suspend media scan 274
switch 241
  zoning 587, 612, 644
synchronize controller clock 278

T

target authentication 285
target logical drive 196, 198–200, 203, 250–251
  read-only attribute 205, 254
Task Assistant 112, 127
termination 7
threshold exceeded 186
topology 223

U

unconfigured capacity 199, 247
Universal XPort device 31
UNIX
  Logical Drive Manager 726
update host DS3000 software 308
user agent 11
using FlashCopy logical drive 195
UTM 31, 397

V

view event log 335, 337
view storage subsystem profile 305
View/Enable Premium Features 74
VMware ESX 711
VolumeCopy 25, 196, 198–202, 204, 250, 252, 255
  change priority 255
  controller ownership 199
  Copy Manager 250
  copy priority rate 200
  create 201
  failed 199
  logical drive 255
  modification operations 198, 200
  pair 256
  performance considerations 200
  permission 253

Premium Feature 201
primary logical drive 196
priority 204, 255
priority rates 200
remove copy pair 255
stop 252
target logical drive 196
Wizard 202
VolumeCopy Premium Feature 196

W

Windows
basic/regular disks 726
configure DS3400 disk 620
creating a FlashCopy logical drive 729
Device Manager 395
Disk Manager 620
disk signature 192
DS3000 Storage Manager 395, 407, 425, 441, 472, 697
dynamic disk 192
FC clustered disk configuration 611
FC disk configuration 585, 671, 711
install DS3000 Storage Manager 82
Java Access Bridge 29
RDAC 29
SAS HBA driver 393
Windows 2003
basic disks 728
configure disk space 401, 426
create FlashCopy logical drive 729
disabled FlashCopy logical drive 727
disk partition 622
DS3000 multipath driver 617
Emulex HBA device driver 597
FlashCopy 726
FlashCopy logical drive 727
install on BladeCenter 598
QLogic HBA driver 614
SAS disk configuration 391
Windows Server 2003 391
WWN 5, 134, 643, 676

Z

zoning 587, 612, 644





Redbooks

IBM System Storage DS3000: Introduction and Implementation Guide

(1.0” spine)
0.875”<->1.498”
460 <-> 788 pages

IBM System Storage DS3000: Introduction and Implementation Guide



Sample configurations with step by step instructions

Configuration and administration with Storage Manager

DS3200, DS3300, DS3400, and EXP3000

This IBM Redbooks publication introduces the IBM System Storage DS3000, providing an overview of its design and specifications, and describing in detail how to set up, configure, and administer it. Since the DS3000 has different disk attachment options, we describe these different methods, including SAS and Fibre Channel. This edition covers updates and additional functions available with the DS3000 Storage Manager Version 10.35 (firmware level 7.35).

You learn how to install the DS3000 Storage Manager, and how to use its GUI and command-line options. We cover quick setup of the DS3000: creating arrays and logical drives and making the storage available to a hosts. You will also find advanced management tasks, including setting up and managing FlashCopy and Volume Copy. Finally, we provide several configuration scenarios. The scenarios describe different methods for attaching various DS3000 models under Windows and Linux, including clustered and boot-from-SAN environments. These examples include step by step instructions using both the GUI as well as command-line scripts.

This book is intended for customers, IBM Business Partners, and IBM technical professionals who want to learn more about the capabilities and advanced functions of the DS3000 Series of storage servers with Storage Manager Software. It also targets those who have a DS3000 storage system and need detailed advice on how to configure and manage it.

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

SG24-7065-01 ISBN 073843213X