


IBM System Storage DS3500 Introduction and Implementation Guide

Sample configurations with
step-by-step instructions

Configuration and administration
with Storage Manager

Troubleshooting and
maintenance



Sangam Racherla
Reza Fanaei Aghdam
Hartmut Lonzer
L G (Len) O'Neill
Mario Rodriguez
Vaclav Sindelar
Alexander (Al) Watson

Redbooks



International Technical Support Organization

**IBM System Storage DS3500 Introduction and
Implementation Guide**

May 2011

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

First Edition (May 2011)

This edition applies to IBM System Storage DS3500 running:

- ▶ Firmware version 7.70
- ▶ IBM System Storage DS Storage Manager version 10.70

© Copyright International Business Machines Corporation 2011. 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 who wrote this book	xv
Now you can become a published author, too!	xviii
Comments welcome	xviii
Stay connected to IBM Redbooks	xviii
Chapter 1. Disk attachment technology	1
1.1 Fibre Channel disk attachment	2
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	11
1.3.3 iSCSI security considerations	12
Chapter 2. Introduction to IBM System Storage DS3500	13
2.1 IBM System Storage Portfolio	14
2.2 DS3500 product models	14
2.2.1 DS3512 and DS3524 Components	15
2.3 EXP3512 and EXP3524	18
2.4 Premium Features	20
2.5 DS3500 and DS3950 Comparisons	22
2.6 IBM System Storage DS Storage Manager	22
Chapter 3. IBM System Storage DS3500 Storage System planning tasks	29
3.1 Planning your SAN and storage server	30
3.1.1 SAN zoning for the DS3500 Storage System	31
3.1.2 Zoning considerations for Enhanced Remote Mirroring	33
3.2 Planning for physical components	33
3.2.1 Rack considerations	33
3.2.2 SAS cables and connectors	35
3.2.3 Ethernet cable and connections	37
3.2.4 Fibre Channel cables and connectors	38
3.2.5 Fibre Channel adapters	44
3.2.6 Disk expansion enclosures	46
3.3 Planning your storage structure	48
3.3.1 Selecting drives	49
3.3.2 Understanding RAID types	50
3.3.3 Array configuration	57
3.3.4 Hot spare drives	59
3.3.5 Enclosure loss protection planning	61
3.3.6 Logical drives and controller ownership	63
3.3.7 Storage partitioning	63
3.3.8 Segment size	67
3.3.9 Media scan	68
3.3.10 Cache parameters	69
3.4 Planning for premium features	70

3.4.1	FlashCopy	71
3.4.2	VolumeCopy	71
3.4.3	Enhanced Remote Mirroring	71
3.4.4	Drive Security	73
3.4.5	Obtaining premium features key	73
3.5	Additional planning considerations	73
3.5.1	Planning for systems with LVM: AIX example	74
3.5.2	Planning for systems without LVM: Windows example	76
3.5.3	Virtualization	78
3.5.4	IBM System Storage SAN Volume Controller overview	78
3.6	Host support and multipathing	79
3.6.1	Supported server platforms	79
3.6.2	Supported operating systems	79
3.6.3	Clustering support	79
3.6.4	Multipathing	80
3.6.5	Microsoft Windows MPIO	80
3.6.6	AIX MPIO	81
3.6.7	AIX Subsystem Device Driver Path Control Module	81
3.6.8	HP-UX IBM Subsystem Device Driver	81
3.6.9	Linux: RHEL/SLES	82
3.6.10	Function of the Auto-Logical Drive Transfer feature	83
3.7	Operating system restrictions	85
3.7.1	Maximum capacity for a logical drive	86
3.7.2	Maximum number of LUNs per host	86
	Chapter 4. IBM System Storage DS3500 and EXP3500 Cabling	87
4.1	DS3500 controller connectors	88
4.1.1	DS3500 controller with standard port configuration	88
4.1.2	DS3500 controller with optional SAS host port adapter	88
4.1.3	DS3500 controller with optional Fibre Channel host port adapter	89
4.1.4	DS3500 controller with optional iSCSI host port adapter	89
4.1.5	EXP3500 ports	89
4.2	Enclosure ID settings	90
4.3	SAS cables	91
4.4	Fibre Channel cabling	93
4.4.1	SFP transceiver modules	93
4.4.2	Fibre Channel cables	96
4.4.3	Interoperability of 2 Gbps, 4 Gbps, and 8 Gbps devices	97
4.5	iSCSI Ethernet cables	97
4.6	EXP3500 attachment	98
4.6.1	Redundant drive channels	98
4.6.2	Drive channel cabling rules	99
4.6.3	Single controller DS3500 with one or more EXP3500 enclosures	99
4.6.4	Dual controller DS3500 with one EXP3500 enclosure	100
4.6.5	Dual Controller DS3500 with two or more EXP3500 enclosures	101
4.6.6	Adding an EXP3500 enclosure to a running dual-controller configuration	103
4.6.7	SAS drive channel miswires	105
4.7	Management connections	105
4.7.1	Out-of-band management	105
4.7.2	In-band management	107
4.8	Host attachment	108
4.8.1	SAS attachment	109
4.8.2	iSCSI attachment	114

4.8.3 Direct attached Fibre Channel	118
4.8.4 SAN fabric-attached DS3500	123
4.9 Power Cabling	129
4.9.1 The DS3500 power supply	129
4.9.2 Powering on and off	130
Chapter 5. Installing IBM System Storage DS Storage Manager	131
5.1 Installing DS Storage Manager on Microsoft Windows 2008	132
5.1.1 Installation preparation	132
5.1.2 Installing the Storage Manager Client on Microsoft Windows 2008	132
5.2 Installing DS Storage Manager on Linux	140
5.2.1 Preparing for the installation	140
5.2.2 Installing Storage Manager using the GUI	141
5.2.3 Installing DS Storage Manager using a text console	146
5.2.4 Uninstalling DS Storage Manager on Linux	148
5.3 Installing DS Storage Manager on AIX	149
5.3.1 Preparing for the installation	150
5.4 Completing the DS Storage Manager installation	151
5.4.1 Performing an automatic discovery of storage subsystems	151
5.4.2 Performing a manual discovery of storage subsystems	152
5.4.3 Add Storage Subsystem verification	155
Chapter 6. Administration - Enterprise Management	157
6.1 Enterprise Management window overview	158
6.1.1 Initial Setup Tasks	158
6.1.2 Enterprise Management window	159
6.2 Functions in the Enterprise Management window	160
6.2.1 Subsystem context menu	160
6.2.2 The Enterprise Management window menu bar	170
6.2.3 The Quick Access buttons	172
Chapter 7. Administration - Summary Tab	177
7.1 Status	178
7.1.1 Storage Subsystem Profile	178
7.1.2 Storage subsystem status	179
7.1.3 Operations in Progress	180
7.1.4 Connection lost	180
7.2 Hardware Components	181
7.3 Capacity	181
7.4 Hosts & Mappings	182
7.4.1 Configured Hosts	182
7.4.2 Host-to-Logical Drive Mappings	183
7.4.3 Storage partitions	183
7.5 Arrays & Logical Drives	184
7.6 Information Center	184
Chapter 8. Administration - Subsystem Management	187
8.1 DS Storage Manager - Subsystem Manager window	188
8.2 Pull-Down Menu	189
8.2.1 Storage Subsystem Menu	189
8.2.2 View menu	205
8.2.3 Mappings menu	207
8.2.4 Array menu	208
8.2.5 Logical Drive menu	208

8.2.6	Controller menu	208
8.2.7	Drive menu	209
8.2.8	Advanced menu	209
8.2.9	Help menu	209
8.3	Toolbar	212
8.3.1	Create new logical drives and arrays	212
8.3.2	View diagnostic event log	212
8.3.3	Monitor Performance	213
8.3.4	Recover from failures	214
8.3.5	Manage enclosure alarm	214
8.3.6	Find in tree	214
8.3.7	Launch copy manager	214
8.4	Status bar	215
8.5	Tabs	216
Chapter 9. Administration - Logical Tab		219
9.1	Logical tab	220
9.2	Working with unconfigured capacity	222
9.2.1	View Associated Physical Components	222
9.2.2	Create array	223
9.3	Working with arrays	225
9.3.1	Locate and View Associated Components	226
9.3.2	Change Ownership and RAID level	227
9.3.3	Add Free Capacity (Drive)	230
9.3.4	Secure Drive	231
9.3.5	Delete and Rename	231
9.3.6	Replace Drive	232
9.4	Working with Free Capacity	232
9.4.1	Create logical drive	233
9.5	Working with logical drives	237
9.5.1	Change Modification Priority	239
9.5.2	Change Cache Settings	241
9.5.3	Change media scan settings	243
9.5.4	Change Pre-Read Redundancy Check	246
9.5.5	Change Ownership/Preferred Path	247
9.5.6	Change Segment Size	248
9.5.7	Increase Capacity	251
9.5.8	Copy Services operations	254
9.5.9	Delete and Rename	254
Chapter 10. Administration - Physical Tab		257
10.1	Physical tab	258
10.2	Discover component properties and location	259
10.2.1	Show disks type	259
10.2.2	View Enclosure Components	259
10.2.3	Disk Drive menu	260
10.2.4	Controller menu	261
10.3	Set hot spare drive	262
10.4	Failed disk drive replacement	266
10.5	Set preferred loop ID	268
10.6	Set remote access	270
10.7	Set Ethernet management ports	270
10.7.1	Initial settings for default IP addresses in DS Storage Manager	271

10.7.2 Configure Ethernet Management Ports.	272
10.8 Configure iSCSI Ports.	273
Chapter 11. Administration - Mappings Tab	275
11.1 Mappings tab.	276
11.2 Defining Host.	277
11.2.1 Adding a new Host to existing Host Group	285
11.3 Defining Storage Partitioning	286
11.4 Defining Host Group	289
11.5 Manage Host Port Identifiers.	290
11.6 Define Additional Mapping	291
11.7 View Unassociated Ports	293
11.8 Move, Remove and Rename Host	293
11.9 Change Host Operating System	294
11.10 Change and Remove Mapping	295
Chapter 12. Administration - Setup tab	297
12.1 Setup tab.	298
12.2 Locate Storage Subsystem.	298
12.3 Rename Storage Subsystem	299
12.4 Set a Storage Subsystem Password.	299
12.5 Configure iSCSI Host Ports.	300
12.6 Configure Storage Subsystem	301
12.6.1 Automatic configuration	302
12.6.2 Configure hot spare drives	305
12.6.3 Create arrays and logical drives	305
12.7 Map Logical Drives	306
12.8 Save Configuration	306
12.9 Manually Define Hosts	307
12.10 Configure Ethernet Management Ports.	308
12.11 View/Enable Premium Features	308
12.12 Manage iSCSI Settings.	308
Chapter 13. Administration - iSCSI	309
13.1 Planning for iSCSI attachment	310
13.2 iSCSI Configuration summary.	311
13.3 Manage iSCSI protocol settings	312
13.3.1 Target Authentication	312
13.3.2 Mutual Authentication	314
13.3.3 Target Identification	315
13.3.4 Target Discovery.	316
13.4 Configure iSCSI Host Ports.	317
13.5 View/End iSCSI Sessions	321
13.6 View iSCSI Statistics.	323
13.7 Defining iSCSI hosts	325
13.7.1 View Unassociated iSCSI initiators.	325
13.7.2 Defining new iSCSI host	326
13.7.3 Manage iSCSI host ports	327
Chapter 14. Administration - Support	329
14.1 The Subsystem Management support tab	330
14.2 Gather Support Information.	331
14.2.1 Saving the Support Data.	332
14.2.2 Support Data Automatic Settings	334

14.2.3 Collect drive data	335
14.3 View Storage Subsystem Profile	338
14.4 Storage Manager Support Monitor	341
14.4.1 Support Monitor installation	341
14.4.2 Support Monitor overview	341
14.4.3 The Support Monitor Profiler console	342
14.4.4 Support Monitor functions	344
14.4.5 Support Monitor - View Module Logs	347
14.5 Download firmware	349
14.5.1 Before you upgrade	350
14.5.2 Updating the host	352
14.5.3 Upgrading the DS3500 controller firmware	353
14.5.4 Using the Enterprise Management upgrade tool	355
14.5.5 Using the DS3500 Storage Manager (Subsystem Management)	368
14.6 View Event Log	385
14.7 Performance Monitor	390
14.8 Import/Export array	392
14.8.1 Export array	392
14.8.2 Import Array procedure	397
14.9 Maintenance - Persistent reservations	401
14.10 Troubleshooting - Drive channels	403
14.11 Troubleshooting - Run Diagnostics	405
14.12 Troubleshooting - Prepare for removal	408
14.13 Recovery Guru - Recover from Failure	409
14.14 Common Recovery Commands	411
14.14.1 Initialize	412
14.14.2 Revive drive	415
14.14.3 Recovery - Clear Configuration	416
14.14.4 Recovery - Place controller	418
14.14.5 Recovery - Reset controller	422
14.14.6 Recovery - Enable controller data transfer	423
14.14.7 Recovery - Place Logical drives online	424
14.14.8 Recovery - Redistribute Logical Drives	424
14.14.9 Recovery - Fail drive	426
14.14.10 Recovery - reconstruct drive	428
14.14.11 Recovery - Defragment Array	429
14.14.12 Recovery - check array redundancy	432
14.14.13 Recovery - Unreadable sectors	434
14.15 View Online Help	436
14.16 About IBM System Storage DS Storage Manager	436
Chapter 15. Disk Security with Full Disk Encryption drives	439
15.1 The need for encryption	440
15.1.1 Encryption method used	440
15.2 Disk Security components	442
15.2.1 DS3500 Disk Encryption Manager	442
15.2.2 Full Data Encryption (FDE) disks	443
15.2.3 Premium feature license	443
15.2.4 Keys	443
15.2.5 Security key identifier	443
15.2.6 Passwords	444
15.3 Setting up and enabling a secure disk	445
15.3.1 FDE and premium feature check	445

15.3.2	Secure key creation	448
15.3.3	Enable Disk Security on array	454
15.4	Additional secure disk functions	456
15.4.1	Changing the security key	456
15.4.2	Save security key file	458
15.4.3	Secure erase	459
15.4.4	FDE drive status	460
15.4.5	Hot spare drive	460
15.5	Migrating secure disk arrays	461
15.5.1	Planning checklist	461
15.5.2	Export the array	461
15.6	Import secure drive array	465
15.6.1	Unlock drives	467
15.6.2	Import array	468
Chapter 16.	IBM Remote Support Manager for Storage	473
16.1	IBM Remote Support Manager for Storage	474
16.1.1	Hardware and software requirements	475
16.1.2	DS-RSM Model RS3	477
16.1.3	Installation choices for RSM for Storage	478
16.1.4	How RSM for Storage works	479
16.1.5	Notification email and events filtering	480
16.1.6	Remote access methods	485
16.1.7	RSM management interface	486
16.1.8	RSM security considerations	487
16.2	Installing and setting up RSM	489
16.2.1	Installing the host OS	489
16.2.2	Installing RSM	490
16.2.3	Setting up RSM	490
16.2.4	Configuring SNMP traps in Storage Manager	506
16.2.5	Activating RSM	507
16.2.6	Remote access security	509
16.2.7	Managing alerts	514
Chapter 17.	Command-Line Interface (CLI)	519
17.1	How to Use the Command Line Interface	520
17.1.1	Usage Notes	520
17.2	Running the CLI	521
17.2.1	Script Editor	521
17.3	General SMcli syntax	523
17.4	Adding a storage subsystem to the Storage Manager configuration	527
17.5	Showing defined subsystems in the Storage Manager configuration	528
17.6	Configuring alerts	529
17.6.1	Defining the mail server and email address to send out the email alerts	529
17.6.2	Defining email alert recipients	529
17.6.3	Deleting email alert recipients	530
17.6.4	SNMP alert recipients	531
17.7	Issuing commands to the storage subsystem	532
17.7.1	Sample command: Save configuration script file	534
17.8	More Information	536
Chapter 18.	Windows SAS configuration guide for IBM BladeCenter	537
18.1	Equipment required	538
18.2	IBM BladeCenter setup	538

18.2.1	Installing Windows Server 2008	539
18.2.2	HS21 SAS Expansion Cards	539
18.2.3	Recording the SAS Expansion Card WWPN	539
18.2.4	HS21 SAS Expansion Card device driver	542
18.2.5	SAS Connectivity modules	542
18.2.6	SAS Connectivity Module firmware update	543
18.2.7	Configuring the SAS connectivity module	546
18.2.8	SAS Connectivity Module zoning	547
18.3	Installing DS Storage Manager host software	549
18.4	Configure the disk space in Windows Server 2008	550
Chapter 19.	Microsoft Cluster configuration with DS3500	555
19.1	Overview of a failover cluster	556
19.1.1	Hardware requirements for a two-node failover cluster	556
19.2	Preparing the environment	557
19.2.1	SAN Zoning configuration	557
19.2.2	DS3500 Storage configuration	557
19.3	Installing DS Storage Manager host software	562
19.3.1	Installing the multipath driver	562
19.4	Windows Server 2008 Failover Clustering	563
19.4.1	Installing the Failover Clustering Feature	563
19.4.2	Validate a Configuration	566
19.4.3	Create a cluster	570
19.4.4	Quorum configuration	572
19.4.5	Steps for configuring a two-node file server cluster	578
Chapter 20.	SuSE Linux configuration guide	587
20.1	DS3500 SAS storage configuration on SLES 11 using RDAC	588
20.1.1	Preparing for the installation	588
20.1.2	Installing the RDAC Multipath Driver	593
20.1.3	Setting up the DS3500 logical drives and host mapping	594
20.1.4	Scan and verify the storage logical drive	595
20.1.5	Configuring RDAC (MPP)	599
20.2	DS3500 iSCSI storage configuration on SLES 11 using RDAC	600
20.2.1	Preparing for the installation	600
20.2.2	Configuring iSCSI software initiator with YaST	602
20.2.3	Configuring iSCSI software initiator manually	607
20.3	DS3500 FC SAN boot configuration for SLES 11 server using RDAC	609
20.3.1	Preparing for the installation	609
20.3.2	SuSE Linux Enterprise 11 installation	615
20.3.3	SuSE Linux final zoning topology	616
20.4	DS3500 FC storage configuration on SLES 11 using DMM	616
20.4.1	DMM Overview	616
20.4.2	Comparing RDAC (MPP) to DMM	618
20.4.3	Planning for the installation	618
20.4.4	Installing the DMM multipath driver	619
20.5	Scan and manage the storage logical drive	620
Chapter 21.	AIX 6.1 configuration guide	625
21.1	Planning for the installation	626
21.1.1	Zoning considerations	627
21.1.2	SAN Boot implementation possibilities	627
21.1.3	Planning MPIO	628
21.2	Configuring MPIO	632

21.3	Setting up the DS3500 logical drives and host mapping.	633
21.4	Scan and manage the storage logical drive from AIX.	634
21.4.1	Ways to manage the paths	636
21.5	AIX SAN Boot with the IBM System Storage DS3500	637
21.5.1	Creating a boot disk with alt_disk_install	637
21.5.2	AIX SAN installation with NIM.	638
21.5.3	AIX SAN installation with CD-ROM.	642
21.5.4	AIX Operating System Installation	642
Chapter 22.	VMware ESX Server and DS3500 Storage Configuration	647
22.1	Introduction to IBM VMware Storage Solutions	648
22.1.1	VMware installation prerequisites	648
22.2	SAN Zoning configuration	649
22.3	DS3500 Storage configuration	649
22.3.1	Mapping LUNs to a storage partition.	650
22.3.2	Steps for verifying the storage configuration for VMware	651
22.4	Installing the VMware ESX Server	652
22.4.1	Configuring the hardware	652
22.4.2	Configuring the software on the VMware ESX Server host	656
22.4.3	Connecting to the VMware vSphere Server	680
22.4.4	Post-Install Server configuration	687
22.4.5	Configuring VMware ESX Server Storage	689
22.4.6	Creating additional virtual switches for guests' connectivity	699
22.4.7	Creating virtual machines	703
22.4.8	Additional VMware ESX Server Storage configuration	717
Appendix A.	IBM Support Portal website	719
	Sample navigation procedure	720
	Download code updates.	723
	My notifications	727
Related publications	731
IBM Redbooks		731
Other publications		731
Online resources		731
How to get Redbooks.		732
Help from IBM		732
Index		733

Notices

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

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

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

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

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

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

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

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


COPYRIGHT LICENSE:

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

Trademarks

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

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

AIX 5L™	FlashCopy®	System i®
AIX®	GPFS™	System p®
Application Advantage™	IBM®	System Storage DS®
BladeCenter®	Netfinity®	System Storage®
DB2®	Power Systems™	System x®
DS4000®	POWER6®	Tivoli®
DS6000™	POWER®	TotalStorage®
DS8000®	Redbooks®	
Express Storage™	Redbooks (logo)  ®	

The following terms are trademarks of other companies:

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

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

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

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

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

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

Preface

This IBM® Redbooks® publication introduces the IBM System Storage® DS3500, providing an overview of its design and specifications, and describing in detail how to set up, configure, and administer it. This edition covers updates and functions available with the DS3500 Storage Manager Version 10.70 (firmware level 7.70).

IBM has combined best-of-breed development with leading 6 Gbps host interface and drive technology in the IBM System Storage DS3500 Express. With its simple, efficient, and flexible approach to storage, the DS3500 is a cost-effective, fully integrated complement to System x servers, BladeCenter®, and Power Systems™. Offering substantial improvements at a price that will fit most budgets, the DS3500 delivers superior price to performance ratios, functionality, scalability, and ease-of-use for the entry-level storage user.

The DS3500 supports intermixing four 1 Gbps iSCSI or four 8 Gbps FC host ports with its native 6 Gbps SAS interfaces. This flexible and multi-purpose dual protocol approach allows organizations to implement a single storage system to support all of their shared storage requirements, there by maximizing productivity, reliability, and cost.

Delivering solid input/output per second (IOPS) and throughput, the DS3500 controllers offer balanced and sustainable performance. The DS3500 can effectively double the performance of the previous DS3000 series of storage systems in both throughput and IOPS.

The DS3500 DS Storage Manager is the same management software offered with the DS5000 and DS4000® series. Now, any of these storage systems can be viewed and managed from a single interface. This allows for consolidated management of these various storage systems and a reduced learning curve. The DS3500 also supports enhanced remote mirroring over FC host ports, which is also compatible with the DS5000 and DS4000 series. This allows for low-cost backup and recovery with a DS5000 and DS4000 at a production site and a DS3500 at the secondary site.

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 IBM System Storage DS3500 with Storage Manager Software. It also targets those who have a DS3500 storage system and need detailed advice on how to configure and manage it.

The team who wrote this book

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



Sangam Racherla is an IT Specialist and Project Leader working at the International Technical Support Organization (ITSO), San Jose Center. He holds a degree in electronics and communication engineering and has ten years of experience in the IT field. He has been with the ITSO for the past seven years and has extensive experience installing and supporting the ITSO lab equipment for various Redbooks publication projects. His areas of expertise include Microsoft® Windows®, Linux®, AIX®, System x®, and System p® servers, and various SAN and storage products.



Reza Fanaei Aghdam is a Senior IT Specialist working in Zurich, Switzerland. He has 17 years of professional experience with x86-based hardware, storage technologies, and systems management, with more than 12 of them at IBM. He instructs Business Partners and customers on how to configure and install System x, BladeCenter, Systems Director, Storage, VMware, and Hyper-V. He is an IBM Certified Systems Expert - System x BladeCenter, IBM Certified Specialist - Midrange Storage Technical Support, and VMware Certified Professional.



Hartmut Lonzer is a Technical Consultant in the Partnership Solution Center Southwest / Germany. As a former Storage FTSS member, his main focus is on Storage and System x. Today, he is responsible for educating and supporting the Business Partners and customers in technical matters. His experience regarding the DS Storage goes back to the beginning of this product. He has been with IBM for 33 years in various technical roles.




L G (Len) O'Neill is a Product Field Engineer (PFE) for IBM System x hardware support based at IBM Greenock in the UK. The PFE team in IBM Greenock provides post-sales technical support for all IBM System x and IBM BladeCenter products for the EMEA (Europe, Middle-East and Africa) region. He has been with IBM for 12 years and in his current role for 11 years. He specializes in providing post-sales technical support for the IBM DS3000 storage products, and previously specialized in supporting IBM SCSI, ServerAID, and Microsoft Windows clustering products within the System x product range. He holds a degree in Physics from Trinity College Dublin.



Mario Rodriguez is an IT Specialist in IBM Uruguay since 2001. He holds MCSE, AIX, LPI, and other Comptia certifications. His areas of expertise include SAN switches (Brocade, Cisco MDS), SAN Storage (DS3000, DS4000, DS6000™, and DS8000®), Linux, AIX, TSM, and VMware. His role in IBM Uruguay is to provide technical support services for virtualization and storage products.



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 7 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.



Alexander (Al) Watson is an ATS Specialist for Storage Advanced Technical Skills (ATS) Americas in the United States. He is a Subject Matter Expert on SAN switches and the IBM Midrange system storage products. He has over fifteen years of experience in planning, managing, designing, implementing, problem analysis, and tuning of SAN environments and storage systems. He has worked at IBM for eleven years. His areas of expertise include SAN fabric networking, Open System Storage IO, and the IBM Midrange Storage solutions.

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

Tamikia Barrow
Margaret Ticknor
David Watts
International Technical Support Organization, Raleigh Center

Doris Konieczny
Harold Pike
Tony Iles
Pete Urbisci
John Fasano
Roger Bullard
Danh T Le
Raul A Gallardo
Paul Goetz
Gene Cullum
James I (Jim) Kennish
David Bennin
Richard Conway

Donald Brennan
IBM

David Worley
Stacey Dershem
Jamal Boudi
LSI Corporation

Brian Steffler
Yong Choi
Alan Hicks
Brocade Communications Systems, Inc.

Now you can become a published author, too!

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

ibm.com/redbooks

- ▶ Send your comments in an email to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

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

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:

<http://www.facebook.com/IBMRedbooks>

- ▶ Follow us on Twitter:

<http://twitter.com/ibmredbooks>

- Look for us on LinkedIn:

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

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

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

- Stay current on recent Redbooks publications with RSS Feeds:

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



Disk attachment technology

In this chapter, we describe basic disk attachment methods in the context of the IBM System Storage DS3500. 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, because the DS3500 products are geared towards Small and Medium Business (SMB) and departmental environments, SAS and iSCSI attachment technologies are supported as well.

1.1 Fibre Channel disk attachment

Fibre Channel (FC) is a high-speed disk attachment technology primarily used for storage networking. It is designed to connect a large number of storage devices to a number of host servers across a Storage Area Network (SAN). Fibre Channel is a transport Protocol (FCP) which transfers SCSI commands and data over Fibre Channel networks.

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 data center solutions.

At the time of writing, the DS3500 Storage maximum FC throughput is 8 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

There are three major Fibre Channel topologies, describing how a number of ports are connected together. A port in Fibre Channel terminology is any entity that actively communicates over the network, not necessarily a hardware port. This port is usually implemented in a device such as disk storage, an HBA on a server, or a Fibre Channel switch.

- Point-to-point

Two devices are connected directly to each other. This is the simplest topology and provides a direct link between an FC HBA inside a host server and a storage device, providing limited connectivity.

- 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 127 devices.

A minimal loop containing only two ports, although appearing to be similar to FC-P2P, differs considerably in terms of the protocol. Only one pair of ports can communicate concurrently on a loop. 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 used these days because 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 complex in large scenarios, connecting hundreds of host servers to a 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.

All devices or loops of devices are connected to Fibre Channel switches, similar conceptually to modern Ethernet implementations. Advantages of this topology over FC-P2P or FC-AL include:

- The switches manage the state of the fabric, providing optimized interconnections.
- The traffic between two ports flows through the switches only, and is not transmitted to any other port.
- Failure of a port is isolated and should not affect operation of other ports.
- Multiple pairs of ports can communicate simultaneously in a fabric.

FC protocol layers

Fibre Channel does not follow the OSI model layering, but is split similarly into five layers as shown in Figure 1-1. These layers include:

- ▶ 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, which implements line coding of signals. 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.

Layers FC0 through FC2 are also known as FC-PH, the physical layers of Fibre Channel.

Fibre Channel routers operate up to FC4 level (i.e. they can operate as SCSI routers), switches up to FC2, and hubs on FC0 only.

Fibre Channel products are available at 1, 2, 4, 8, 10 and 20 Gbit/s. Products based on the 2, 4 and 8 Gbit/s standards should be interoperable and backward compatible. The 10 Gbit/s standard and its 20 Gbit/s derivative, however, are not backward compatible with any of the slower speed devices because they differ considerably on FC1 level in using 64b/66b encoding instead of 8b/10b encoding, and are primarily used as inter-switch links.

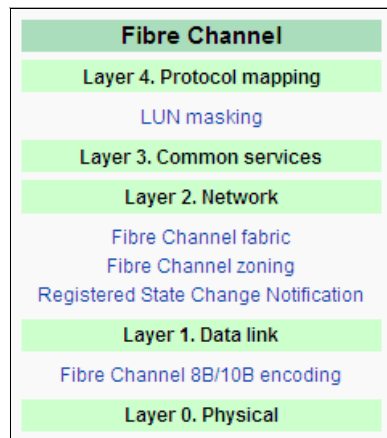


Figure 1-1 Fibre Channel layers

FC cable types

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

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 low. However, single-mode fiber cables can be much longer. Cable length can exceed 50 km.

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 300 m.

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 long cable lengths are needed, single-mode cabling will be required.

Despite its name, Fibre Channel signaling can run on both copper wire and fiber-optic cables as shown in Figure 1-2.

FC Media Type	Speed MB/s	Transmitter	Variant	Distance
Copper	400	4 Pairs RJ45	GF GBase-T	40/100m
	200	4 Pairs RJ45	GF GBase-T	60/70/100m
	100	4 Pairs RJ45	GF GBase-T	100m
Single- Mode Fibre	800	1310 nm Longwave Laser	800-SM-LC-L	2 m – 10 km
	400	1310 nm Longwave Laser	400-SM-LC-L	2 m - 10 km
	200	1550 nm Longwave Laser	200-SM-LL-V	2 m - >50 km
		1310 nm Longwave Laser	200-SM-LC-L	2 m - 10 km
	100	1550 nm Longwave Laser	100-SM-LL-V	2 m - >50 km
		1310 nm Longwave Laser	100-SM-LC-L	2 m - 10 km
		1310 nm Longwave Laser	100-SM-LL-I	2 m - 2 km
Multi- Mode Fibre	800	850 nm Shortwave Laser	800-M5E-SA-I	0.5 m – 300m
	400		400-M5/6-SN-I	0.5 m - 150m
	200		200-M5/6-SN-I	0.5 m - 300m
	100		100-M5/6-SN-I	0.5 m - 500m
			100-M6-SL-I	2 m - 175m

Figure 1-2 FC cable types

Small form-factor pluggable (SFP) transceiver

The small form-factor pluggable (SFP) or Mini-GBIC is a compact, hot-pluggable transceiver used for both telecommunication and data communications applications. It interfaces a network device mother board (for a switch, router, media converter, or similar device) to a fiber optic or copper networking cable. SFP transceivers are designed to support SONET, Gigabit Ethernet, Fibre Channel, and other communications standards.

SFP transceivers are available with a variety of transmitter and receiver types, allowing users to select the appropriate transceiver for each link to provide the required optical reach over the available optical fiber type (for example multi-mode fiber or single-mode fiber).

Optical SFP modules are commonly available in several categories:

- ▶ 850 nm 550m MMF (SX)
- ▶ 1310 nm 10 km SMF (LX)
- ▶ 1550 nm [40 km (XD)]
- ▶ 80 km (ZX)
- ▶ 120 km (EX or EZX)
- ▶ DWDM

SFP transceivers are also available with a copper cable interface, allowing a host device designed primarily for optical fiber communications to also communicate over unshielded twisted pair networking cable.

SFP transceivers are commercially available with capability for data rates up to 4.25 Gbit/s. The standard is expanding to SFP+ which supports data rates up to 10.0 Gbit/s (that will include the data rates for 8 gigabit Fibre Channel, 10 GbE, and OTU2).

FC World Wide Names (WWN)

A World Wide Name (WWN) or World Wide Identifier (WWID) is a unique identifier that identifies a particular Fibre Channel, Advanced Technology Attachment (ATA) or Serial Attached SCSI (SAS) target. Each WWN is an 8 byte number derived from an IEEE OUI and vendor-supplied information.

There are two formats of WWN defined by the IEEE:

- ▶ Original format: addresses are assigned to manufacturers by the IEEE standards committee, and are built into the device at build time, similar to an Ethernet MAC address. The first 2 bytes are either hex 10:00 or 2x:xx (where the x's are vendor-specified), followed by the 3-byte vendor identifier and 3 bytes for a vendor-specified serial number
- ▶ New addressing schema: first nibble is either hex 5 or 6 followed by a 3-byte vendor identifier and 36 bits for a vendor-specified serial number

1.2 Serial Attached SCSI (SAS) disk attachment

SAS is a computer bus used to move data to and from computer storage devices such as hard drives and tape drives. SAS depends on a point-to-point serial protocol that replaces the parallel SCSI bus technology, and uses the standard SCSI command set.

At the time of writing, typical SAS throughput is 6 Gbps full duplex. SAS has the capability to reach 24 Gbps if the host can drive it at that speed. When the first 6 Gbps connection is full, the next 6 Gbps connection is used, and so on, up to four connections.

Figure 1-3 on page 6 shows the SAS technical specifications.

Performance	Full-duplex with link aggregation (4-ports wide at 24 Gbit/s)
	3.0 Gbit/s at introduction, 6.0 Gbit/s available February 2009
Connectivity	10 m external cable
	255 device port expanders (>65k total devices)
Availability	SAS-to-SATA compatibility
	Dual-port HDDs
Driver	Multi-initiator point-to-point
	Software-transparent with SCSI

Figure 1-3 SAS Technical Specifications

A SAS Domain, an I/O system, consists of a set of SAS devices that communicate with one another by means of a service delivery subsystem. Each SAS device in a SAS domain has a globally unique identifier called a World Wide Name (WWN or SAS address). The WWN uniquely identifies the device in the SAS domain just as a SCSI ID identifies a device in a parallel SCSI bus. A SAS domain can contain up to a total of 65,535 devices.

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 be greatly increased. There are two types of expanders:

- ▶ Fan-out expanders

A fanout expander can connect up to 255 sets of edge expanders, known as an edge expander device set, allowing for even more SAS devices to be addressed. A fanout expander cannot do subtractive routing: it can only forward subtractive routing requests to the connected edge expanders.

- ▶ Edge expanders

An edge expander allows for communication with up to 255 SAS addresses, allowing the SAS initiator to communicate with these additional devices. Edge expanders can do direct table routing and subtractive routing.

In the current DS3500 implementation, up to 96 drives can be configured in a single DS3500 using three EXP3500 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, that 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

Serial Attached SCSI comprises three transport protocols:

- ▶ Serial SCSI Protocol (SSP): supports SAS disk drives.
- ▶ Serial ATA Tunneling Protocol (STP): supports SATA disks.
- ▶ Serial Management Protocol (SMP): manages SAS Expanders.

At the physical layer, the SAS standard defines connectors and voltage levels. The physical characteristics of the SAS wiring and signaling are compatible with and have loosely tracked that of SATA up to the present 6 Gbit/s rate, although SAS defines more rigorous physical signaling specifications and a wider allowable differential voltage swing intended to support longer cabling. Although SAS-1.0/SAS-1.1 adopted the physical signaling characteristics of SATA at the 1.5 Gbit/s and 3 Gbit/s rates, SAS-2.0 development of a 6 Gbit/s physical rate led the development of an equivalent SATA speed. According to the SCSI Trade Association, 12 Gbit/s is slated to follow 6 Gbit/s in a future SAS-3.0 specification.

SAS wide ports

Each SAS port includes four full duplex links or *lanes* within a single connector, as shown in Figure 1-4, with each lane running a speed of 6 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 6 gigabits per second. If another block of data then needs to be written to disk while the first link 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 24 Gbps.




Figure 1-4 SAS wide ports

SAS drive technology

Figure 1-5 shows how SAS drives are attached to the controllers. The point-to-point topology used in SAS configurations 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 1-5 Point to Point SAS Topology

1.3 iSCSI disk attachment

iSCSI stands for *Internet Small Computer System Interface*, an Internet Protocol (IP)-based storage networking standard for linking data storage facilities. By carrying SCSI commands over IP networks, iSCSI is used to facilitate data transfers over intranets and to manage storage over long distances. iSCSI can be used to transmit data over local area networks (LANs), wide area networks (WANs), or the Internet, and can enable location-independent data storage and retrieval.

iSCSI uses TCP/IP (typically TCP ports 860 and 3260). In essence, iSCSI simply allows two hosts to negotiate and then exchange SCSI commands using IP networks. By doing this iSCSI takes a popular high-performance local storage bus and emulates it over wide-area networks, creating a storage area network (SAN).

Unlike certain SAN protocols, iSCSI requires no dedicated cabling: it can be run over existing switching and IP infrastructure. However, the performance of an iSCSI SAN deployment can be severely degraded if not operated on a dedicated network or subnet (LAN or VLAN). As a result, iSCSI is often seen as a low-cost alternative to Fibre Channel, which requires dedicated infrastructure except in its Fibre Channel over Ethernet (FCoE) form.

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, whereas an FC SAN can run up to 8 Gbps. However, there are ways to address iSCSI performance:

- ▶ Although the host servers can use almost any Ethernet network interface card for iSCSI traffic, this does mean that the CPUs on 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 on 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 and it gains wider commercial acceptance. Migrating to 10 GbE can significantly increase the performance of an iSCSI infrastructure.

1.3.1 iSCSI initiators and targets

iSCSI uses the concept of initiators and targets, as shown in Figure 1-6.

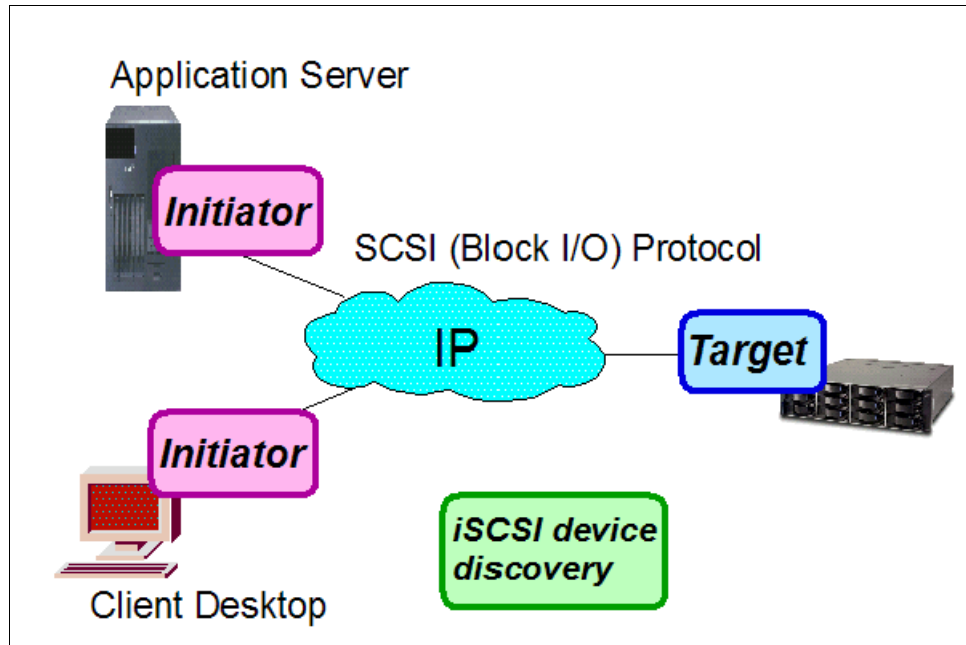


Figure 1-6 iSCSI components

The protocol allows clients (called initiators) to send SCSI commands (CDBs) to SCSI storage devices (targets) on remote servers.

An initiator functions as an iSCSI client. An initiator typically serves the same purpose to a computer as a SCSI bus adapter would, except that instead of physically cabling SCSI devices (like hard drives and tape changers), an iSCSI initiator sends SCSI commands over an IP network. An initiator falls into two broad types:

- Software initiator

A software initiator uses code to implement iSCSI. Typically, this happens in a kernel-resident device driver that uses the existing network card (NIC) and network stack to emulate SCSI devices by utilizing the iSCSI protocol. Software initiators are available for most operating systems, and are the most common method of deploying iSCSI.

An example of an iSCSI stack is the Microsoft iSCSI Software Initiator, which runs on Windows Server 2003, Windows Server 2008, and VMware ESX4. At the time of writing, the current version is V2.08 and is available for download from the Microsoft website. For Windows Server 2008 and ESX4, the iSCSI initiator is included in the box.

Note: Refer to System Storage Operation Center (SSIC) for the complete list of the supported operating systems. The SSIC can be found at:

http://www-03.ibm.com/systems/support/storage/config/ssic/displayessearchwithoutjs.wss?start_over=yes

For the IBM AIX operating system, refer to the “iSCSI software initiator and software target” topic at the following URL:

<http://publib.boulder.ibm.com/infocenter/systems/index.jsp>

► Hardware initiator

A hardware initiator uses dedicated hardware, typically in combination with software (firmware) running on that hardware, to implement iSCSI. A hardware initiator mitigates the overhead of iSCSI and TCP processing and Ethernet interrupts, and therefore might improve the performance of servers that use iSCSI.

An iSCSI host bus adapter (more commonly, HBA) implements a hardware initiator. A typical HBA is packaged as a combination of a Gigabit (or 10 Gigabit) Ethernet NIC, a TCP/IP offload engine (TOE) technology, and a SCSI bus adapter, which is how it appears to the operating system. Inside the operating system, the iSCSI HBAs are classified as storage adapters.

An iSCSI HBA can include PCI option ROM to allow booting from an iSCSI target.

A TCP Offload Engine, or “TOE Card”, offers an alternative to a full iSCSI HBA. A TOE “offloads” the TCP/IP operations for this particular network interface from the host processor, freeing up CPU cycles for the main host applications. When a TOE is used rather than an HBA, the host processor still has to perform the processing of the iSCSI protocol layer itself, but the CPU overhead for that task is low.

iSCSI HBAs or TOEs are used when the additional performance enhancement justifies the additional expense of using an HBA for iSCSI, rather than using a software-based iSCSI client (initiator).

An iSCSI target usually represents hard disk storage that works over the IP or Ethernet networks. DS3500. 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, and as part of all sessions established between initiators and targets. The types of iSCSI names are:

- IQN
- EUI
- NASA

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.2010-08.local.rivers.rhine`

The DS3500 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`

NAA

NAA name formats were added to iSCSI to provide compatibility with naming conventions used in Fibre Channel and Serial Attached SCSI (SAS) storage technologies.

The format of a NAA is naa plus NAA 64 or 128 bit identifier. For example:

naa.52004567BA64678D

iSCSI addressing

Usually an iSCSI participant can be defined by three or four fields:

1. Hostname or IP Address (for example, "iscsi.example.com")
2. Port Number (for example, 3260)
3. iSCSI Name (for example, the IQN "iqn.2003-01.com.ibm:00.fcd0ab21.shark128")
4. An optional CHAP Secret (for example, "secrets")

The iSCSI address can have 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 most common 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 can 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” on page 12)

Service Location Protocol

The Service Location Protocol (SLP) 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 that is not accessible to Ethernet network users. iSCSI is a SAN technology that uses the Ethernet network, which is a lot more vulnerable to intrusion. Therefore, iSCSI security is important.

iSCSI connection authentication

iSCSI initiators and targets prove their identity to each other using the Challenge Handshake Authentication Protocol (CHAP), which includes a mechanism to prevent cleartext passwords from appearing on the wire. 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. In addition to CHAP, several authentication methods can be used:

- ▶ 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

Because iSCSI relies on TCP/IP communication, IP Security (IPSec) can be used to achieve increased security. 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)



Introduction to IBM System Storage DS3500

In this chapter, we introduce the new IBM System Storage DS3500 Storage Subsystem offerings and functionality. These products consists of models of storage subsystems that provide a variety of environments to meet various user needs. We describe the EXP3512 and EXP3524 SAS disk drive enclosures as well.

We also explain the Premium Features philosophy and how the Storage Manager utility works with these new products.

2.1 IBM System Storage Portfolio

IBM has brought together into one family, known as the DS family, a broad range of disk systems to help small and large enterprises select the right solutions for their needs. The DS family combines the high-performance IBM System Storage DS8000 Series of enterprise servers with the DS5000 series of mid-range systems and the DS3000 entry level systems.

The DS3000 series consist of two new major products: the DS3500 and the DS3950. Both of these products are a good fit for the entry to mid-range SAN and direct-attach market space. With the common Storage Manager shared by these new DS3000 storage systems and the DS5000 storage systems, there is a smooth link into the DS5000 series systems, with remote mirroring and copy services features being shared by these two platforms. The DS3500 and the DS3950 offer robust functionality, exceptional reliability, and availability with the common ease of storage management being shared by all. The overall positioning of these new DS3000 series products within the IBM System Storage DS® family is shown in Figure 2-1.

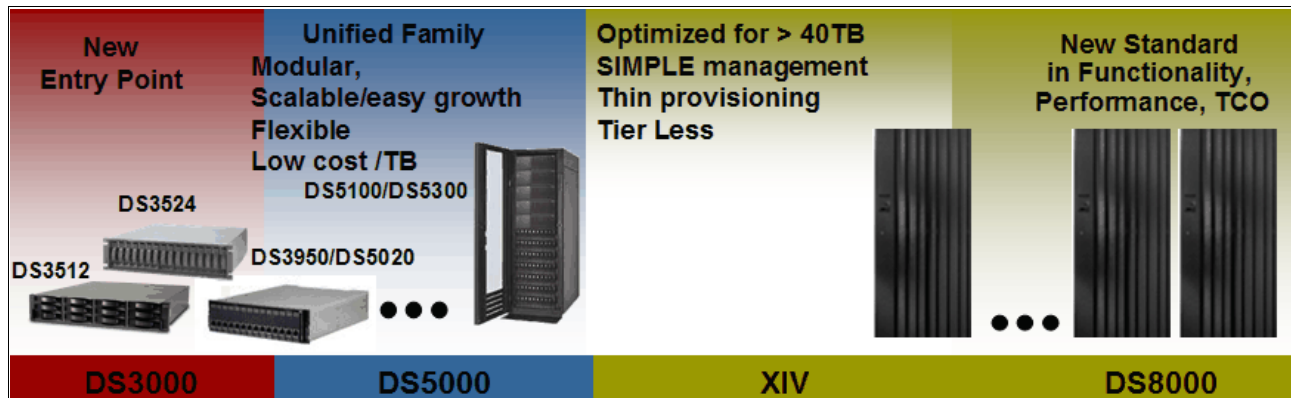


Figure 2-1 IBM System Storage family

2.2 DS3500 product models

The new IBM System Storage DS3500 series storage subsystems support up to two redundant RAID controllers in either a 12 or 24 drive configuration. The models for the storage servers are DS3512 and DS3524. There are also two models of drive expansion chassis (a 12 and a 24 drive) that can be attached to either of the storage subsystems. The models for these are EXP3512 and EXP3524. The new DS3500 models provides a number of new capabilities from the previous generations. The enhancements are:

- ▶ Allows for one storage subsystem to be able to perform in the environments of the three older DS3000 family members, with support options for SAS, iSCSI, and Fibre Channel host connections.
- ▶ With this new generation we have the marriage of the DS3000 and the DS5000 Storage Manager and firmware releases, allowing for a common management console to support the entry and midrange DS families.
- ▶ Adds the Enhanced Remote Mirroring (ERM) Premium Feature to the DS3000 line¹.
- ▶ New 6 Gbps SAS technology for host and drive attachments.
- ▶ Support for greater capacity with new larger capacity SAS drive offerings.

¹ Only available in DS3500 and DS3950

Figure 2-2 and Figure 2-3 show the front view of both chassis models of these subsystems.



Figure 2-2 DS3512 and EXP3512 subsystem assembly from the front view



Figure 2-3 DS3524 and EXP3524 servers assembly from the front view

2.2.1 DS3512 and DS3524 Components

The DS3500 storage server is a 2U rack mountable enclosure, containing either one or two RAID controller modules, two power supplies, and up to 12 or 24 disk modules. See Figure 2-4 for the component layouts.




Figure 2-4 DS 3500 components

RAID controller

RAID controllers support RAID levels 0, 1, 3, 5, 6, and 10. Each controller has 1 GB (upgradeable to 2 GB) of user data cache with battery backup. The battery provides power if the cache needs to be destaged to the SD flash card if power is disrupted.

In dual controller configurations, the controller on the left is A and the right is B, when viewed from the rear 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.

All DS3500 RAID controllers have connectors for the following ports built into them:

- ▶ Two 6 Gbps SAS host server attachment ports
- ▶ Drive side 6 Gbps SAS expansion port
- ▶ Ethernet management port
- ▶ Serial management port

The RAID controllers and two redundant power supply modules are installed in the rear of the subsystem as shown in Figure 2-5.




Figure 2-5 DS3500 controller subsystem rear view

In Figure 2-5, the controller modules are in the upper half of the subsystem and the power supply modules are in the lower half.

Power Supply

The DS3500 power supply module is a 585 Watt DC power supply. It is auto ranging, 100-240VAC input capable. As shown in Figure 2-5, the power supply provides LED indicators for the following states (starting from left):

- ▶ Standby power LED (green): Currently this LED is not used.
- ▶ DC power LED (green): When this LED is lit, it indicates that the DS3500 is turned on and is supplying both 5-volt and 12-volt dc power.
- ▶ OK to remove LED (blue): When this blue LED is lit, it indicates that it is safe to remove the power supply.
- ▶ 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 storage subsystem is receiving ac power

Host interface cards

As mentioned earlier, the DS3500 comes with two SAS host attachment ports built into the controller modules. Additional host server connectivity is supported through the use of an optional daughter card (shown in Figure 2-6 on page 17). This interface card can provide for one of the following to be added to the DS3500:

- ▶ Additional four SAS ports
- ▶ Eight 1 Gbit iSCSI ports (four per controller)
- ▶ Eight FC ports (four per controller)



Figure 2-6 Example host interface daughter card module

Both the single and the dual controller models of the DS3500 storage servers can be upgraded to include an optional host interface daughter card. When dual controllers are installed, both controllers must be equipped with the same daughter card option to enable the support of the controller failover functions.

Figure 2-7 shows the SAS optional daughter card installed in the controller. With this option the subsystem will have up to eight 6 Gbps SAS connections for host attachments. For details on the cabling and use of this configuration with the BladeCenter and stand-alone environments see 4.8, “Host attachment” on page 108.




Figure 2-7 Controller module with optional SAS host interface daughter card

Figure 2-8 shows the iSCSI optional daughter card installed in the controller. With this option the subsystem will have up to eight iSCSI connections for host attachments. For details on the cabling and use of this configuration with the BladeCenter and stand-alone environments see 4.8, “Host attachment” on page 108.



Figure 2-8 Controller module with optional iSCSI host interface daughter card

Figure 2-9 on page 18 shows the Fibre Channel optional daughter card installed in the controller. With this option the subsystem will have up to eight 8 Gbps Fibre Channel connections for host attachments. For details on the cabling and use of this configuration with the BladeCenter and stand-alone environments see 4.8, “Host attachment” on page 108.



Figure 2-9 Controller module with optional FC host interface daughter card

Note: Only one type of optional interface can be added to any one DS3500 storage server. Mixing interface daughter cards between controllers in the same DS3500 is not supported.

Disk drives

The most important difference between the DS3512 and the DS3524 product models and their equivalent expansion models are the hard disks that are supported with them. The difference starts with the physical drive size and extends to their speeds and storage capacities. The DS3512 and EXP3512 support 12 drives in the 3.5 inch format; the DS3524 and EXP3524 supports 24 drives in the 2.5 inch format. The disk drives are installed at the front, as shown in Figure 2-2 on page 15 and Figure 2-3 on page 15. Available drive types for each of these subsystems are shown in Table 2-1.

Table 2-1 DS3500 families HDD support

Drives Supported	DS3512/EXP3512	DS3524/EXP3524
SAS 15KRPM	300 GB, 450 GB, 600 GB	73 GB, 146 GB
SAS 15KRPM SED	600 GB	None
SAS 10KRPM	None	300 GB
SAS 10KRPM SED	None	300 GB
Nearline SAS 7.5KRPM	1 TB, 2 TB	500 GB
Maximum drives	12/96	24/96
Storage system capacity (max)	450 GB SAS / 1 TB SATA	450 GB SAS / 1 TB SATA

Note: In DS3500 family, you can add a mix of EXP3512 or EXP3524 expansion units to attain a maximum capacity of 190 TB per subsystem.

2.3 EXP3512 and EXP3524

The EXP3512 and EXP3524 expansion subsystems allow for the growth of the DS3500 storage subsystem up to the 96 drive maximum, by adding either the 12 or 24 drive chassis to the storage server's SAS drive expansion port. Any mix of the expansion models can be added up to the maximum allowed drive count. The EXP3512 and EXP3524 differ from the DS3512 and DS3524 in that in place of the controller module they are equipped with an Environmental Services Module (ESM). As with the DS3500 controllers the expansions can

be optionally upgraded with a second ESM module for redundant paths. Each ESM has a 6 Gbps SAS connection providing 600 MBps throughput. Figure 2-10 shows the rear view of the EXP35xx with its port connections for cabling.




Figure 2-10 EXP3512 and EXP3524 rear view with SAS port connections

With the EXP 3512 and EXP3524, only one of the two IN ports are used on each ESM to connect expansions together into a cascaded loop configuration. As shown in Figure 2-11 on page 20, the cabling scheme used for connecting these expansions follows what is known as a *top down, bottom up* method. This provides the expansion loops with redundant paths to the enclosures, and in the event of one expansion encountering a catastrophic failure, the others are still able to continue to run. With a proper RAID layout this can provide for uninterrupted operations.

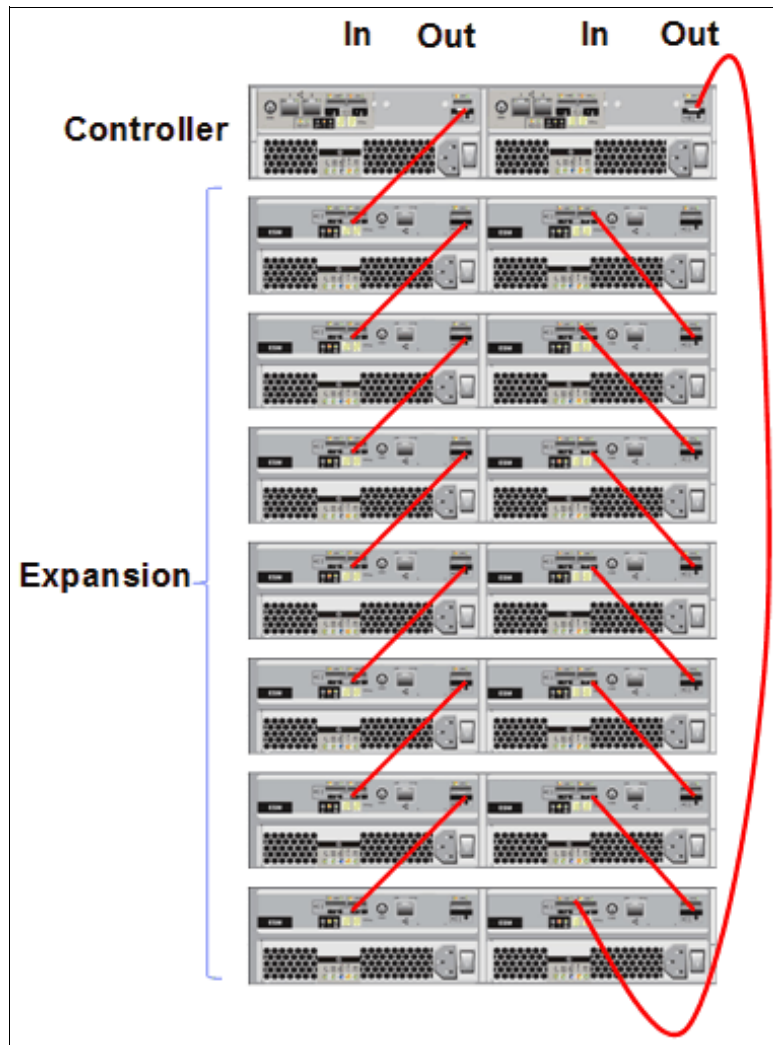


Figure 2-11 EXP35XX expansion cascaded loop

2.4 Premium Features

Standard configurations of DS3500 storage subsystems come with the following Premium Features.

DS3000 Partition Expansion License

As part of the standard configuration, four storage partitions are enabled on the DS3500 with Windows and Linux on Intel® host attach license (this can be increased to 8, 16, 32, or 64). The maximum number of storage partitions is 64 on all DS3500 products. Use the DS3500 Partition Expansion License to enable all 64 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 *IBM Midrange System Storage Copy Services Guide*, SG24-7822.

Note: Flashcopy does not provide a permanent full image copy for recovery use. For this functionality you must use the VolumeCopy feature.

As part of the standard configuration, two FlashCopies are enabled on every DS3500 storage subsystem and this Premium Feature enables up to 64 FlashCopies. A maximum of eight flashcopies can be created on a single logical drive.

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 the 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 depends on the size of the source data and the operation priority settings. While establishing the copy, the source logical drive will be in a read-only state.

After 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 *IBM Midrange System Storage Copy Services Guide*, SG24-7822.

The VolumeCopy Premium Feature allows for up to 128 VolumeCopies.

Important: 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 Enhanced Remote Mirroring

With the DS3500 and DS3950, Enhanced Remote Mirroring (ERM) provides a way to create a remote image copy of a source logical drive. This capability is frequently used for the creation of a disaster recovery site located in a separate location some distance from the primary location. ERM provides three mirroring modes to choose from:

- ▶ **synchronous:** Provides mirroring capability in the metro (campus) environments, generally within 10 mile radius. Requires a low latency network as mirroring must complete before the primary completes the acknowledgement to the host.
- ▶ **asynchronous with write order consistency group:** Provides mirroring capability in the global environments, generally over 10 mile distances, where latency might exceed acceptable response times for the primary host IO operation. IO operations are handled in the order of their original request across the selected logical drives who are grouped together through the use of a consistency group.
- ▶ **asynchronous copy:** Provides copying capability in the global environments, generally over 10 mile distances, where latency is high and bandwidth is limited causing a high backup of IO that might impact the principle operation site. IOs in this mode are not guaranteed to be processed in the order received across a number of logical drives.

This capability is shared across the DS3500 and DS3950 and the DS5000 midrange products that allows the remote location to be a consolidated collection site.

Up to eight mirrored pairs can be supported on the DS3500 and 64 mirrored pairs are supported by the DS3950. For more information about Enhanced Remote Mirroring, see *IBM Midrange System Storage Copy Services Guide*, SG24-7822.

2.5 DS3500 and DS3950 Comparisons

Table 2-2 shows the capabilities of the new DS3000 product offerings for ease of comparison.

Table 2-2 DS3500 and DS3950 Specification comparisons

Fully configured systems	DS3500	DS3950
Host interfaces	Eight 8 Gbps FC and four 6 Gbps SAS Eight 1 Gbps iSCSI and four 6 Gbps SAS Four or eight 6 Gbps SAS	Four 8 Gbps FC, or Four 8 Gbps FC and four 1 Gbps iSCSI
Redundant drive channels	Two 6 Gbps SAS	Four 4 Gbps FC
Drive enclosures	12-drive (3.5-in) 2U 24-drive (2.5-in) 2U	16-drive (3.5-in) 3U
Max drives	48 standard 96 premium feature	112
Drives supported	SAS, NL SAS, FDE	FC, SATA, FDE
XOR engine	Integrated XOR / P+Q	Integrated XOR / P+Q
Dual Ethernet	Yes	Yes
Data cache / memory	2 or 4 GB	2 or 4 GB
Max logical drives	256	1,024
Max LUN size	> 2 TB	> 2 TB
RAID levels	1, 3, 5, 6, 10	1, 3, 5, 6, 10
Max RAID 10 / 0 size	96 drives	112 drives
Max RAID 3 / 5 / 6 size	30 drives	30 drives
Max partitions	64	128
Max snapshots per base volume	8	8
Maximum remote mirrors (across FC)	8	64

2.6 IBM System Storage DS Storage Manager

The new DS3500 and DS3950 storage subsystem uses the same Storage Manager as the DS5000 product line and the older DS3000 models. At the time of writing the current version

of the software was 10.70.G5.08. When you receive your DS3500 storage subsystem, you might also receive with it a copy of the IBM System Storage DS Storage Manager Software and Host Kit CDs. If you do not receive this software, or the version you receive is not compatible with your DS3500's firmware release, you can download it from the IBM support website.

Using IBM System Storage DS 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. We discuss the IBM System Storage DS Storage Manager software and its use in great detail later in this book, so we will just briefly describe some of the main points here.

When discussing the IBM System Storage DS Storage Manager, it is important to differentiate between the following two terms:

- ▶ Host server

This is a server attached to the DS3500 storage subsystem through the I/O path (SAS, iSCSI, or Fibre Channel). The host server has access to the logical drives that are defined on the DS3500 storage server for its storage use.

- ▶ Management station

The management station is the system that is responsible for managing all, or a portion of, a storage network. The IBM System Storage DS Storage Manager provides a GUI that runs on the management station. This management station can be based on Windows, Linux, AIX, or Solaris operating systems. There might be slight shading differences between the operating system version of the displays, but the fit and functions are the same for all. You need to establish a management connection between the management station and the DS3500 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 DS3500 RAID controller. All management communication flows across the TCP/IP connection between the management station and the DS3500. We also call this method *direct-attached management*. The management station in this case only requires an Ethernet connection to the DS3500.

- In-band

This method utilizes the I/O path between a host server and the DS3500. In this case, the management station does not have direct TCP/IP connection to the DS3500, but rather communicates with the DS3500 through an HBA that acts as a gateway to the DS3500 storage subsystem. That is, communication between the management station and the host server is across the Fibre Channel, SAS or iSCSI I/O path.

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

Because each method has associated advantages and disadvantages we will discuss them both to help you select which is more appropriate for your environment. 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 DS3500 to transfer management commands and information.

This method does not use the management Ethernet ports on DS3500 RAID controllers and does not require a management TCP/IP network. However, it does require a special *access logical drive* to manage the DS3500 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 2-12. Two host servers are attached to the DS3500 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 DS3500 across the FC I/O path.




Figure 2-12 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 because you should not use the production LAN or VLAN for management network traffic.

A separate management workstation is another requirement. Typically system administrators use their own workstation for this purpose. Figure 2-13 shows the management workstation and the DS3500 subsystem connected on the Ethernet management network.




Figure 2-13 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 DS3500 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.

Generally, set up and use both methods because this will introduce redundancy to your management setup and provide management access to the DS3500 subsystem even in the case of I/O path or Ethernet management network failure.

The IBM System Storage DS 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)

Storage Manager Client (SMclient)

This is the actual graphical user interface (GUI) that you use to manage the DS3500 subsystems. 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. There are two high level management screens that are used for these functions:

- ▶ Enterprise Management window
- ▶ Subsystem Management window

Enterprise Management window

This window opens when the DS3500 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. A sample of the Linux based Enterprise Management window is shown in Figure 2-14.

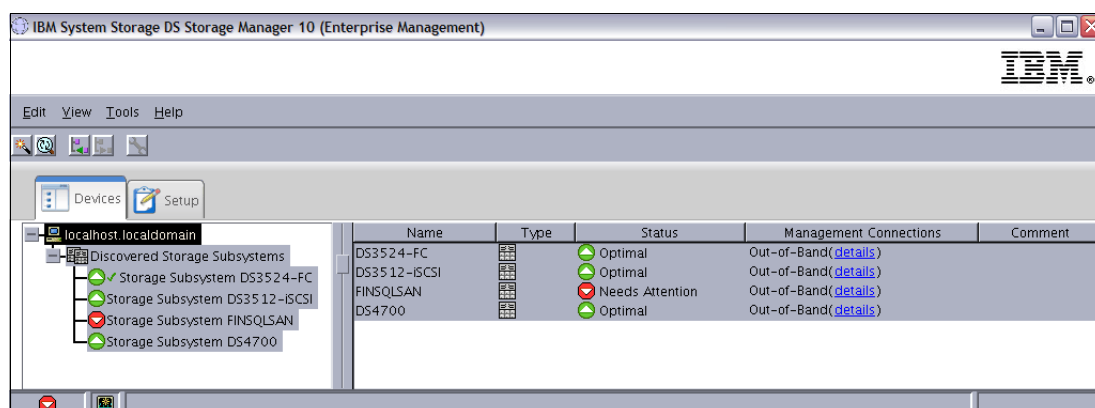


Figure 2-14 Enterprise Management window

Subsystem Management window

This window allows you to manage a particular DS3500 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 2-15 on page 27 for an example of this window again based on the Linux based version.

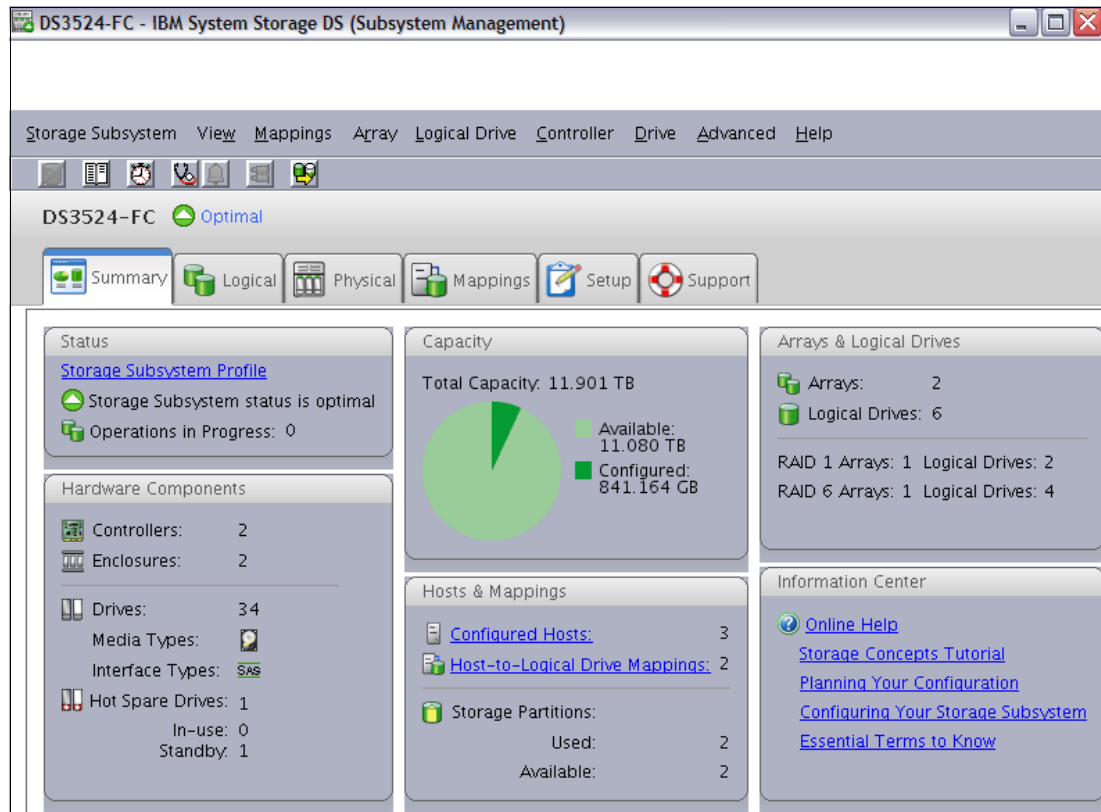


Figure 2-15 DS3500 IBM System Storage DS Subsystem Management window

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.

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. 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 DS3500 Storage Manager, launched from the Enterprise Management window.

You can either run a single command at a time or execute pre-written scripts. See Chapter 17, "Command-Line Interface (CLI)" on page 519 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 on the host server and allows the SMclient to communicate with DS3500 across the I/O path. At the time of writing, only the FC I/O path is supported for management communication.

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 on 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 on the host server for ease of troubleshooting and problem analysis. With these utilities there are many quick changes that can be implemented and a great deal of detail on the current configuration settings and their impacts.

Multipath driver support

In general, install two HBAs on the host servers and use the dual controller DS3500 subsystems. This provides the ability to create completely redundant paths to the host server's storage, therefore improving the availability through multiple I/O paths to the controllers and full storage redundancy out the backend to the disks. However, host dual path configurations can only work properly if you install the appropriate multipath driver for your particular operating system into the host server.

The IBM System Storage DS Storage Manager Software and Host Kit for Windows includes multipath support based on the Microsoft MPIO framework. The IBM System Storage DS Storage Manager Software and Host Kit for Linux does not include any multipath support. RDAC for Linux is available as a separate open source package called MPP. For Solaris there is a version of DMP that can be implemented for multi-path support. In AIX there is an MPIO driver available to support the DS family of storage subsystems. All of these drivers require special configurations and setups to be followed when implementing them. A simple parameter setting can be the difference between a successful failover to an alternate path and an application outage.

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.



IBM System Storage DS3500 Storage System planning tasks

Careful planning is essential to any new storage installation. This chapter provides guidelines to help you with the planning process.

Choosing the right equipment and software, and also knowing what the right settings are for a particular installation, can be challenging. Every installation has to answer these questions and accommodate specific requirements, and there can be many variations in the solution.

Having a well thought-out design and plan prior to the implementation will help you get the most out of your investment for the present and protect it for the future.

During the planning process, you need to answer numerous questions about your environment:

- ▶ What are my SAN requirements?
- ▶ What hardware do I need to buy?
- ▶ What reliability do I require?
- ▶ What redundancy do I need? (For example, do I need off-site mirroring?)
- ▶ What compatibility issues do I need to address?
- ▶ Will I use any storage virtualization product such as IBM SAN Volume controller?
- ▶ Will I use any unified storage product such as the IBM System Storage N series?
- ▶ What operating system am I going to use (existing or new installation)?
- ▶ What applications will access the storage subsystem?
- ▶ What are the hardware and software requirements of these applications?
- ▶ What will be the physical layout of the installation? Only local site, or remote sites as well?
- ▶ What level of performance do I need?
- ▶ How much does it cost?

This list of questions is not exhaustive, and as you can see, certain questions go beyond simply configuring the storage system. But the focus in this chapter is to help with the creation of a successful solution and that frequently extends beyond a single subsystem.

Use this chapter as a reference to help you gather the information for the statements.

3.1 Planning your SAN and storage server

When planning the setup of a Storage Area Network (SAN), you want the solution to answer your current requirements and fulfill your future needs.

First, the SAN fabric must be able to accommodate a growing demand in storage (generally storage needs double every two years). Second, the SAN must be able to keep up with the constant evolution of technology and resulting hardware upgrades and improvements. It is estimated that a storage installation needs to be upgraded every 2 to 3 years.

Ensuring compatibility among various pieces of equipment is crucial when planning the installation. The important question is what device works with what, and also who has tested and certified that equipment.

When designing a SAN storage solution, complete the following steps:

1. Produce a statement outlining the solution requirements that can be used to determine the type of configuration you need. Then use this statement to cross-check that the solution design delivers the basic requirements. The statement must have easily defined bullet points covering the requirements, for example:
 - New installation or upgrade of existing infrastructure
 - Infrastructure type(s) to be used: SAS, iSCSI, Fibre Channel (direct or fabric)
 - Host Bus Adapter (HBA) selection
 - HBA driver type selection: SCSIPort or StorPort
 - Multipath Driver selection: RDAC, DMMP, MPIO, or SDDPCM
 - Types of applications accessing the SAN (whether transaction or throughput intensive)
 - Required capacity
 - Required redundancy levels
 - Type of data protection needed
 - Current data growth patterns for your environment
 - Whether current data is more read or write based
 - Backup strategies in use: Network, LAN-free, or Server-less
 - Premium features required: Partitioning, FlashCopy, Volume Copy, or Enhanced Remote Mirroring
 - Number of host connections required
 - Types of hosts and operating systems that will connect to the SAN
 - Zoning required
 - Distances between equipment and sites (if there is there more than one site)
2. Produce a hardware checklist. It must cover such items that require you to:
 - Make an inventory of existing hardware infrastructure. Ensure that any existing hardware meets the minimum hardware requirements and is supported with the DS3500 Storage System.
 - Make a complete list of the planned hardware requirements.
 - Ensure that you have enough rack space for future capacity expansion.
 - Ensure that the power and environmental requirements are met.

- Ensure that your existing network of SAS, Ethernet or Fibre Channel switches and cables are properly configured.
3. Produce a software checklist to cover all the required items that need to be certified and checked. It must include such items that require you to:
- Ensure that the existing versions of firmware and storage management software are up to date.
 - Ensure that host operating systems are supported with the storage system. Check the IBM System Storage interoperation Center (SSIC) for your specific storage system available from this website for more information:

<http://www.ibm.com/servers/storage/disk>

These lists are not exhaustive, but the creation of the statements is an exercise in information gathering and planning; it gives you a greater understanding of what your needs are in your current environment and creates a clear picture of your future requirements. The goal ought to be quality rather than quantity of information.

Use this chapter as a reference to help you gather the information for the statements.

Understanding the applications is another important consideration in planning for your DS3500 Storage System setup. Applications can typically either be I/O intensive, such as high number of I/O per second (IOPS); or characterized by large I/O requests, that is, high throughput or MBps.

- ▶ Typical examples of high IOPS environments are Online Transaction Processing (OLTP), databases, and Microsoft Exchange servers. These have random writes and fewer reads.
- ▶ Typical examples of high throughput applications are data mining, imaging, and backup storage pools. These have large sequential reads and writes.

By understanding your data and applications, you can also better understand growth patterns. Being able to estimate an expected growth is vital for the capacity planning of your DS3500 Storage System installation. Clearly indicate the expected growth in the planning documents, although the actual patterns might differ from the plan according to the dynamics of your environment.

Selecting the right DS Storage System model for your current and perceived future needs is one of the most crucial decisions you will make. The good side, however, is that the DS3500 platforms offer scalability and expansion flexibility. Premium features can be purchased and installed at a later time to add more functionality to the storage server as well.

In any case, it is perhaps better to purchase a higher model than one strictly dictated by your current requirements and expectations, which will allow for greater performance and scalability as your needs and data grow. Starting out with a maximum configured storage solution to meet your current needs with no room to expand can save money initially, but could quickly become too small for your business needs and near term growth.

3.1.1 SAN zoning for the DS3500 Storage System

Zoning is an important part of integrating a DS3500 Storage System in a SAN. When done correctly, it can eliminate many common problems. Zoning also helps with creating paths that can be used by the multi-path drivers for better failover protection. Understanding the capabilities of the multi-path driver is important when designing the paths it will use.

Important: Disk and tape must be on separate HBAs, and so the disk and tape access will also be in separate zones. With certain UNIX® systems, this arrangement is supported by the DS3500 Storage System, but might have hardware limitations. Therefore, *do not share HBAs between disk storage and tape*.

Zoning with MPIO

MPIO is a multipath driver that provides for a host to be able to recognize multiple paths to the attached storage device. This is done by utilizing multiple HBA ports or devices on the host server connected to SAN fabric switches, which are also connected to the multiple ports on the storage devices. The DS3000, DS4000, and DS5000 storage products have two controllers within the subsystem that manage and control the disk drives. These controllers behave in an active/passive fashion. Ownership and control of a particular LUN is done by one controller. The other controller is in a passive mode until a failure occurs, at which time the LUN ownership is transferred to that controller. Each controller can have more than one fabric port for connectivity to the SAN fabric. This helps to provide faster recovery for path failures that can be resolved through the use of an alternate connection to the same controller rather than requiring the full failover process to be run.

The DS3500 models only support drivers that by design follow the rules for this type of path management. See Figure 3-1 for an example of how the configuration should be implemented.

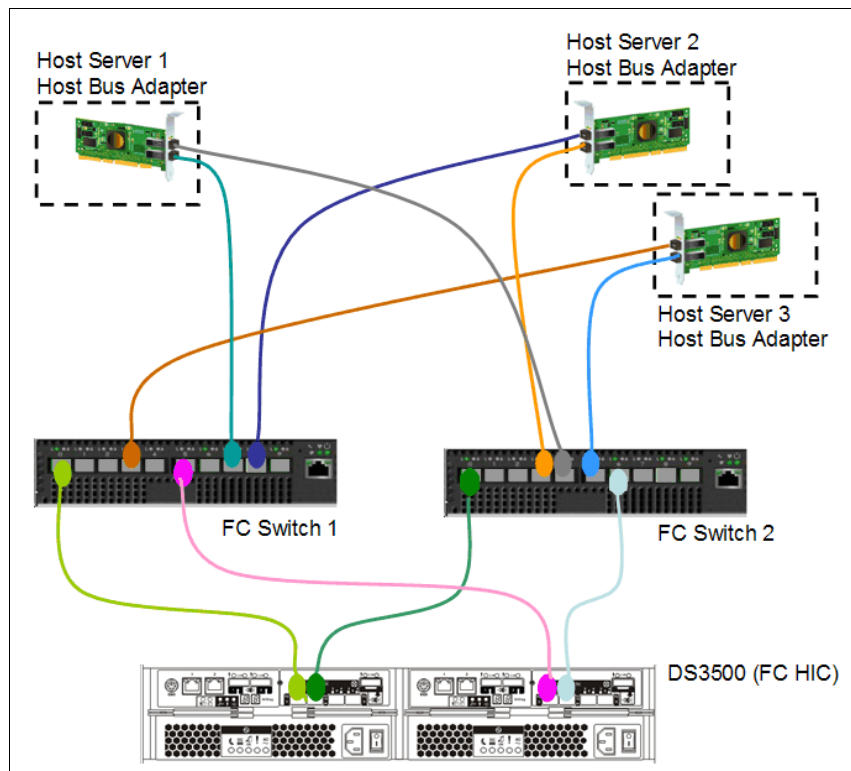


Figure 3-1 Host HBA to storage subsystem controller multipath sample configuration

With MPIO, it is better to create a zone for each HBA port to be able to see both controllers to help decrease failover operations for network related failures. In the above example, to do this you would create zones for the three host servers (two zones each) with one being from HBA port 1 to both controller A port 1 and controller B port 1 of the DS3500 storage server, and the

second zone from HBA port 2 of the host to controller A and Controller B through port 2 of the controllers.

With this configuration, if a single path has a fault another path can be used to access the devices without the need of a controller failover.

General rule: For MPIO driver solutions, create separate zones for each HBA port connection from the host to both controllers (one zone for HBA port 1 and one zone for HBA port 2), which isolates each initiator (HBA port) from the other.

3.1.2 Zoning considerations for Enhanced Remote Mirroring

The Enhanced Remote Mirroring (ERM) can only be connected through a Fibre Channel (FC) connection, which must be dedicated for data replication between the subsystems. The ports used for ERM cannot be used to send or receive I/Os from any host. These requirements are addressed by defining SAN zoning. There must be two zones defined for the ERM network, one for controller A and one for controller B. Any zones defined must separate the host ports from the storage system mirroring ports, and also separate the mirroring ports between the controllers.

When using ERM, you must create two additional zones:

- ▶ The first zone contains the ERM source DS3500 controller A and ERM target DS3500 controller A.
- ▶ The second zone contains the ERM source DS3500 controller B and ERM target DS3500 controller B.

ERM is detailed further in *IBM Midrange System Storage Copy Services Guide*, SG24-7822.

3.2 Planning for physical components

In this section, we review elements related to physical characteristics of an installation, such as rack considerations, fiber cables, Fibre Channel adapters, and other elements related to the structure of the storage system and disks, including enclosures, arrays, controller ownership, segment size, storage partitioning, caching, hot spare drives, and Enhanced Remote Mirroring.

3.2.1 Rack considerations

The DS3500 Storage System and possible expansions are mounted in rack enclosures.

General planning

Consider the following general planning guidelines. Determine the following:

- ▶ The size of the floor area required by the equipment:
 - Floor-load capacity
 - Space needed for expansion
 - Location of columns
- ▶ The power and environmental requirements

Create a floor plan to check for clearance problems. Be sure to include the following considerations in the layout plan:

- ▶ Determine service clearances required for each rack or suite of racks.
- ▶ If the equipment is on a raised floor, determine:
 - The height of the raised floor
 - Anything that might obstruct cable routing
- ▶ If the equipment is not on a raised floor, determine:
 - The placement of cables to minimize obstruction
 - If the cable routing is indirectly between racks (such as along walls or suspended), the amount of additional cable needed
 - Cleanliness of floors, so that the fan units will not attract foreign material such as dust or carpet fibers
- ▶ Determine the location of:
 - Power receptacles
 - Air conditioning equipment, placement of grilles, and controls
 - File cabinets, desks, and other office equipment
 - Room emergency power-off controls
 - All entrances, exits, windows, columns, and pillars
 - Fire control systems
- ▶ Check access routes for potential clearance problems through doorways and passage ways, around corners, and in elevators for racks and additional hardware that will require installation.
- ▶ Store all flammable spare materials in properly designed and protected areas.

Rack layout

To be sure that you have enough space for the racks, create a floor plan before installing the racks. You might need to prepare and analyze several layouts before choosing the final plan.

If you are installing the racks in two or more stages, prepare a separate layout for each stage.

The following considerations apply when you make a layout:

- ▶ The flow of work and personnel within the area
- ▶ Operator access to units, as required
- ▶ If the rack is on a raised floor, determine:
 - The need for adequate cooling and ventilation
- ▶ If the rack is not on a raised floor, determine:
 - The maximum cable lengths
 - The need for cable guards, ramps, and so on to protect equipment and personnel
- ▶ Location of any planned safety equipment
- ▶ Future expansion

Review the final layout to ensure that cable lengths are not too long and that the racks have enough clearance.

Allow at least 76 cm (30 in.) of clearance in front of the controller-drive tray and 61 cm (24 in.) behind the controller-drive tray for service clearance, ventilation, and heat dissipation.

Airflow Through the Controller-Drive Tray with 12 Drives:

- ▶ 76 cm (30 in.) clearance in front of the cabinet
- ▶ 61 cm (24 in.) clearance behind the cabinet

Airflow Through the Controller-Drive Tray with 24 Drives:

- ▶ 76 cm (30 in.) clearance in front of the cabinet
- ▶ 61 cm (24 in.) clearance behind the cabinet

Important: All vertical rack measurements are given in rack units (U). One U is equal to 4.45 cm (1.75 in.). The U levels are marked on labels on one front mounting rail and one rear mounting rail. All DS3500 storage systems that are in a chassis are 2U in height.

Figure 3-2 shows an example of the required service clearances for a 9306-900 42U rack. Check the documentation for the specific rack model that you will use for a statement on the required clearances.




Figure 3-2 9306 enterprise rack space requirements

The rack space requirements will vary depending on the model of your rack. Always refer to the manufacturer's documentation for information specific to your rack.

3.2.2 SAS cables and connectors

All DS3500 storage subsystems, regardless of the optional host attachment feature added, will have SAS host connections available. Also, with the addition of any expansion subsystems they use SAS cabling to attach expansion enclosures. Therefore, let us consider the SAS cables and connectors used.

The SAS ports on the DS3500 controller and EXP3500 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 3-3. IBM provides SAS cables in two cable lengths: 1 and 3 meters.




Figure 3-3 SAS cable

Careful planning should be done 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

Attention: Care should be taken to not use the cable for leverage when removing the cable from the mini-SAS port.

In Figure 3-4 on page 37 we show an example of SAS connections being made to direct attached hosts and a BladeCenter through an internal SAS switch module. The DS3500 comes with two standard built-in SAS ports per controller. The figure shows the system with the additional option daughter card for the additional two ports per controller.

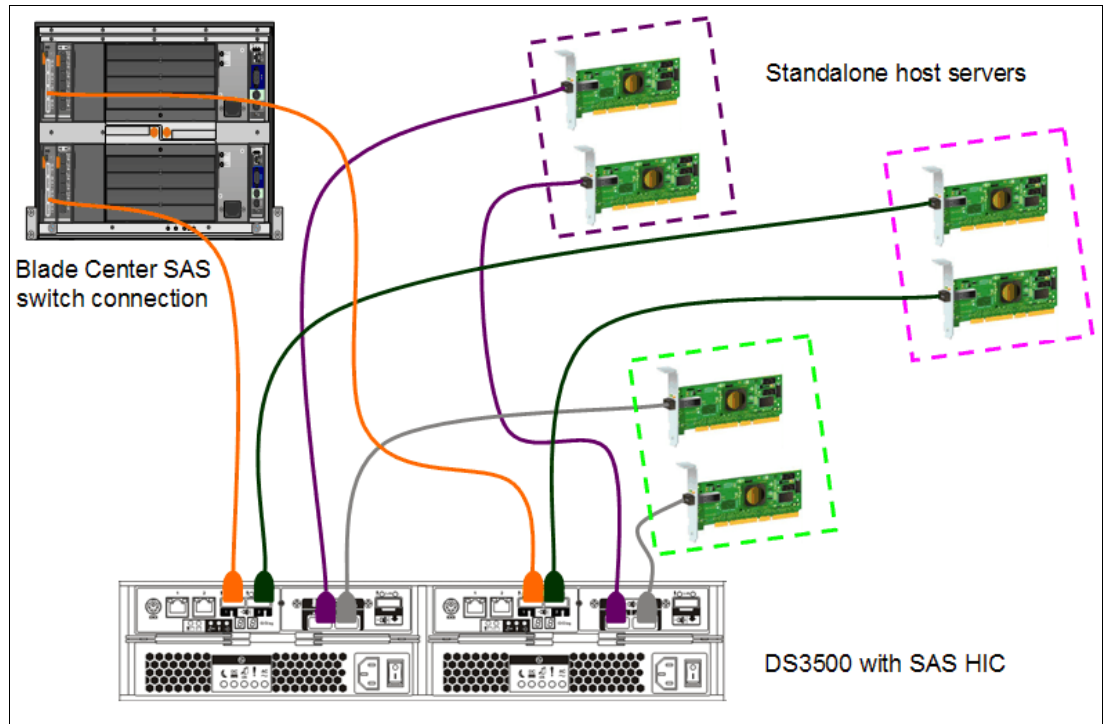


Figure 3-4 DS3500 with optional added SAS host connections

3.2.3 Ethernet cable and connections

Each DS3500 RAID controller contains an Ethernet management port that you can use for out-of-band management. If you have a dual controller DS3500 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 DS3500 Storage Manager will not be able to manage the enclosure.

Avoid using your public LAN for DS3500 out-of-band management. Instead, set up a dedicated LAN or VLAN just for management purposes. This will provide increased security of your DS3500 storage system. If the DS3500 RAID controllers are on a public LAN, a knowledgeable user could install the DS3500 Storage Manager on a separate workstation, or use the CLI to run potentially destructive tasks. For an additional layer of security, also enable password protection on the DS3500 storage system. Refer to “Set Password” on page 204 for more details.

The DS3500 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 3-5 on page 38 for an example.




Figure 3-5 Four host servers attached to a DS3500 subsystem with iSCSI

3.2.4 Fibre Channel cables and connectors

In this section, we discuss various essential characteristics of fiber cables and connectors. This information can help you understand the options you have for connecting and cabling the DS3500 Storage System.

Cable types (shortwave or longwave)

Fiber cables are basically available in multi-mode fiber (MMF) or single-mode fiber (SMF).

Multi-mode fiber allows light to disperse in the fiber so that it takes many paths, bouncing off the edge of the fiber repeatedly to finally get to the other end (multi-mode means multiple paths for the light). The light taking these various paths gets to the other end of the cable at slightly separate times (separate paths, separate distances, and separate times). The receiver has to determine which incoming signals go together.

The maximum distance is limited by how “blurry” the original signal has become. The thinner the glass, the less the signals “spread out,” and the further you can go and still determine what is what on the receiving end. This dispersion (called modal dispersion) is the critical factor in determining the maximum distance a high-speed signal can travel. It is more relevant than the attenuation of the signal (from an engineering standpoint, it is easy enough to increase the power level of the transmitter or the sensitivity of your receiver, or both, but too much dispersion cannot be decoded no matter how strong the incoming signals are).

There are two core sizes of multi-mode cabling available: 50 micron and 62.5 micron. The intermixing of the two core sizes can produce unpredictable and unreliable operation. Therefore, core size mixing is not supported by IBM. Users with an existing optical fiber infrastructure are advised to ensure that it meets Fibre Channel specifications and is a consistent size between pairs of FC transceivers.

Single-mode fiber (SMF) is so thin (9 microns) that the light can barely “squeeze” through and it tunnels through the center of the fiber using only one path (or mode). This behavior can be explained (although not simply) through the laws of optics and physics. The result is that because there is only one path that the light takes to the receiver, there is no “dispersion confusion” at the receiver. However, the concern with single mode fiber is attenuation of the signal. Table 3-1 lists the supported distances.

Table 3-1 Cable type overview

Fiber type	Speed	Maximum distance
9 micron SMF (longwave)	1 Gbps	10 km
9 micron SMF (longwave)	2 Gbps	2 km
50 micron MMF (shortwave)	1 Gbps	500 m
50 micron MMF (shortwave)	2 Gbps	300 m
50 micron MMF (shortwave)	4 Gbps	150 m
50 micron MMF (shortwave)	8 Gbps	50 m
62.5 micron MMF (shortwave)	1 Gbps	300 m
62.5 micron MMF (shortwave)	2 Gbps	150 m
62.5 micron MMF (shortwave)	4 Gbps	70 m
62.5 micron MMF (shortwave)	8 Gbps	21 m

Note that the “maximum distance” shown in Table 3-1 is just that, a maximum. Low quality fiber, poor terminations, excessive numbers of patch panels, and so on, can cause these maximums to be far shorter.

All IBM fiber feature codes that are orderable with the DS3500 Storage System will meet these standards.

Interfaces, connectors, and adapters

In Fibre Channel technology, frames are moved from source to destination using gigabit transport, which is a requirement to achieve fast transfer rates. To communicate with gigabit transport, both sides have to support this type of communication, which is accomplished by using specially designed interfaces that can convert other types of communication transport into gigabit transport.

The interfaces that are used to convert the internal communication transport of gigabit transport are Small Form Factor Transceivers (SFF), also often called Small Form Pluggable (SFP). See Figure 3-6 on page 40. Gigabit Interface Converters (GBIC) are no longer used on current models, although the term GBIC is still sometimes incorrectly used to describe these connections.



Figure 3-6 Small Form Pluggable (SFP) with LC connector fiber cable

Obviously, the particular connectors used to connect a fiber cable to a component will depend upon the receptacle into which they are being plugged.

LC connector

Connectors that plug into SFF or SFP devices are called LC connectors. The two fibers each have their own part of the connector. The connector is keyed to ensure correct polarization when connected, that is, transmit to receive and vice-versa.

The main advantage that these LC connectors have over the SC connectors is that they are of a smaller form factor, and so manufacturers of Fibre Channel components are able to provide more connections in the same amount of space.

All DS3500 series products use SFP transceivers and LC fiber cables. See Figure 3-7.




Figure 3-7 LC fiber cable connector

General Rule: When you are not using an SFP, it is best to remove it from the port on the DS3500 storage controller and replace it with a cover. Similarly, unused cables must be stored with ends covered to help eliminate risk of dirt or particles contaminating the connection while not in use.

Interoperability of 2 Gbps, 4 Gbps, and 8 Gbps devices

The Fibre Channel standard specifies a procedure for speedy auto-detection. Therefore, if a 4 Gbps port on a switch or device is connected to a 2 Gbps port, it must negotiate down and the link will run at 2 Gbps. If there are two 8 Gbps ports on either end of a link, the negotiation runs the link at 8 Gbps if the link is up to specifications. A link that is too long or “dirty” can end up running at 4 Gbps, even with 8 Gbps ports at either end, so care must be taken with cable lengths distances and connector quality.

The same rules apply to 8 Gbps devices relative to 4 Gbps and 2 Gbps environments. The 8 Gbps devices have the ability to automatically negotiate back down to either 4 Gbps or 2 Gbps, depending upon the attached device and the link quality. A 4 Gbps device has the ability to automatically negotiate back down to either 2 Gbps or 1 Gbps. If the link does unexpectedly negotiate to a slower speed than expected, then the causes or reasons for this ought to be investigated and remedied.

The DS3500 Storage System now has 8 Gbps functionality: there are several switches and directors that operate at this speed.

Note: On certain fiber switch vendor models, it might be necessary to configure the port to a specific speed of 2, 4, or 8 Gbps to obtain the required speed instead of leaving “auto-detection” on the port.

FC host attachment methods

With FC SAN environments there are two methods of host attachment to the storage system that are commonly used: direct-attached and switch-attached (fabric).

Direct-attached DS3500

Figure 3-8 shows a simple direct-attached configuration. Two host servers are connected to a dual-controller DS3500. 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.

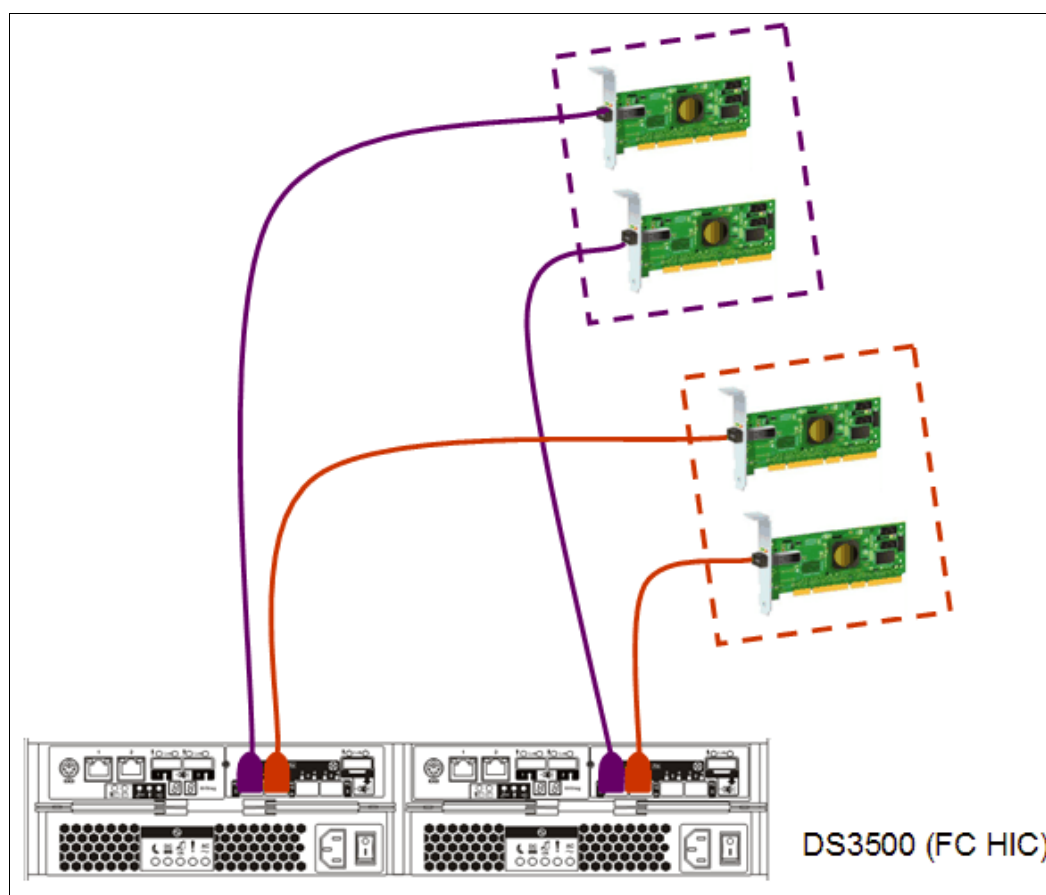


Figure 3-8 Host servers attached to DS3500 FC direct connect

Switch-attached DS3500

Because the DS3500 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 3-9 displays a sample configuration with four host servers attached to a dual controller DS3500. Each host has two FC HBAs for redundant I/O path support.

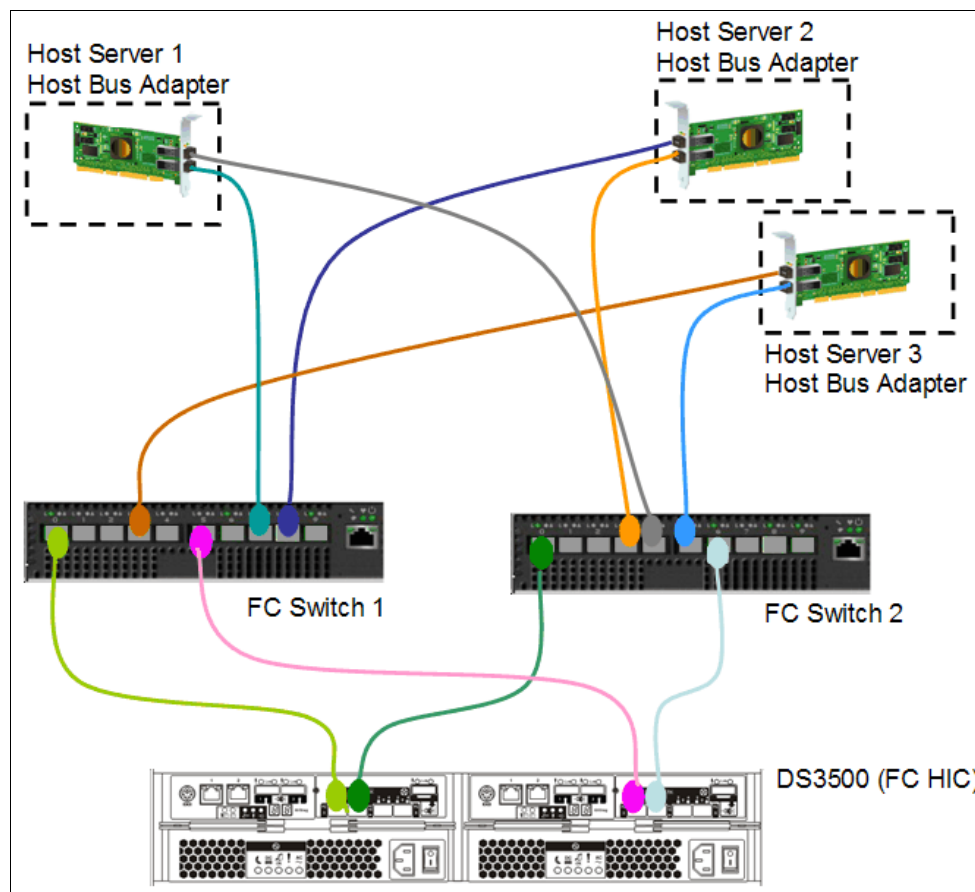


Figure 3-9 DS3500 connected to hosts through FC switch SAN

If you plan to use the Enhanced Remote Mirroring (ERM) feature, the fourth FC port on each controller must be dedicated to this connection. No host servers should be zoned or configured to use these ports. ERM can only be used in a SAN switch environment.

Cable management and labeling

Cable management and labeling for solutions should include the installation, routing, and labeling of all cables including power, Ethernet LAN network, SAS for both host attachments and backend expansions, and Fibre Channel. Cable management and labeling needs have expanded from the traditional labeling of network connections to management and labeling of most cable connections between your servers, disk subsystems, multiple network connections, and power for all the subsystems involved in your solution. In many cases components of the solution have differing cable needs. Multiple unique systems could be located in the same rack, or a system could span across multiple racks. In certain cases the solution could be configured where components are not be physically located in the same room, building, or site.

Cable planning

Successful cable management planning includes three basic activities: site planning (before your solution is installed), cable routing, and cable labeling.

Site planning

Having adequate site planning completed before your solution is installed will result in a reduced chance of installation problems. Significant attributes covered by site planning are location specifications, electrical considerations, raised/non-raised floor determinations, and cable lengths. Consult the documentation for each component of your solution for special site planning considerations.

Cable routing

With effective cable routing, you can keep your solution's cables organized, reduce the risk of damaging cables, and allow for effective service and support. Use the following guidelines to assist with cable routing:

- ▶ When installing cables to devices mounted on sliding rails:
 - Run the cables neatly along equipment cable-management arms and tie the cables to the arms. Obtain the cable ties locally.

Note: Do not use cable-management arms for fiber cables.

- Take particular care when attaching fiber optic cables to the rack. See the instructions included with your fiber optic cables for guidance on minimum radius, handling, and care of fiber optic cables.
 - Run the cables neatly along the rack rear corner posts.
 - Use cable ties to secure the cables to the corner posts.
 - Make sure the cables cannot be pinched or cut by the rack rear door.
 - Run internal cables that connect devices in adjoining racks through the open rack sides.
 - Run external cables through the open rack bottom.
 - Leave enough slack so that the device can be fully extended without putting a strain on the cables.
 - Tie the cables so that the device can be retracted without pinching or cutting the cables.
- ▶ To avoid damage to your fiber optic cables, follow these guidelines:
 - Use great care when utilizing cable management arms.
 - When attaching to a device on slides, leave enough slack in the cable so that it does not bend to a radius smaller than is advised by your fiber optic cable guide when extended or become pinched when retracted.
 - Route the cable away from places where it can be snagged by other devices in the rack.
 - Do not overtighten the cable straps or bend the cables to a radius smaller than is advised by your fiber optic cable guide.
 - Do not put excess weight on the cable at the connection point and be sure that it is well supported. For example, a cable that goes from the top of the rack to the bottom *must* have a method of support other than the strain relief boots built into the cable.

- For long cable runs, ensure that enough slack is made for rack movement in accordance with your computer room standards for earthquake proofing.

Additional information for routing cables for IBM Netfinity® Rack products can be found in *IBM Netfinity Rack Planning and Installation Guide*, part number 24L8055. This publication includes pictures providing more details about how to set up the cable routing.

Cable labeling

When labeling your solution, follow these tips:

- ▶ As you install cables in the rack, label each cable with the appropriate identification.
- ▶ Remember to attach labels to any cables that you replace.
- ▶ Document deviations from the label scheme you use. Keep a copy in your Change Control Log book.
- ▶ Comply with an existing cable naming convention, or define and adhere to a simple logical naming convention.

An example of label naming convention might include these attributes:

- ▶ The function, to help identify the purpose of the cable
- ▶ Location information must be broad to specific (for example, the site/building to a specific port on a server or hub)

Other cabling mistakes

Avoid making these common mistakes in cabling:

- ▶ Leaving cables hanging from connections with no support.
- ▶ Not using dust caps.
- ▶ Not keeping connectors clean. Certain cable manufacturers require the use of lint-free alcohol wipes to maintain the cable warranty.
- ▶ Leaving cables on the floor where people might kick or trip over them.
- ▶ Not removing old cables when they are no longer needed or planned for future use.

Tip: Collect all SFP, HBA, and cable dust caps, and store them in a dust-free container, to be used for future cabling work. Do not re-use dust caps that have been left loose in the rack or computer room.

3.2.5 Fibre Channel adapters

We now review topics related to Fibre Channel adapters:

- ▶ Placement on the host system bus
- ▶ Distributing the load among several adapters
- ▶ Queue depth
- ▶ Driver selection

Host system bus

Today, there is a choice of high-speed adapters for connecting disk drives. Fast adapters can provide better performance. The HBA must be placed in the fastest supported slot available.

Important: Do not place all the high-speed Host Bus Adapters (HBAs) on a single system bus. If you do, the computer bus will become a performance bottleneck.

It is always a general rule to distribute high-speed adapters across several buses. When you use PCI adapters, make sure that you first review your system specifications. Certain systems include a PCI adapter placement guide.

The number of adapters you can install depends on the number of PCI slots available on your server, but also on what traffic volume you expect on your SAN. The rationale behind multiple adapters is either redundancy (failover) or load sharing.

Failover

When multiple adapters are installed on the host system and used with a multipath driver, the multipath driver checks to see if all the available paths to the storage server are still functioning. In the event of an HBA or cabling failure, the path is changed to the other HBA, and the host continues to function without loss of data or functionality.

In general, all operating systems support two paths to the DS3500 Storage System. Microsoft Windows 2003 and 2008, and Linux support up to four paths to the storage controller. AIX with MPIO can support more than four paths to the controller modules. However, too many paths can delay failover when it actually is needed due to a controller failure. Care should be taken to ensure that there are enough redundant paths without having too many. In general, having four paths provides a good level of redundancy for load balancing, and still allows for timely controller failover when the need arises.

Load balancing

Load balancing or load sharing means distributing I/O requests from the hosts between multiple adapters, which can be done by assigning LUNs to DS3500 controllers A and B alternatively (see also 3.3.6, “Logical drives and controller ownership” on page 63).

Figure 3-10 shows the principle for a load-sharing setup. A multipath driver checks all available paths to the controller. In Figure 3-10, that is two paths (red and blue). The driver forces the data down all paths in a *round-robin* scheme, which means that it does not really check for the workload on a single path, but moves the data down in a *rotational manner* (round-robin).

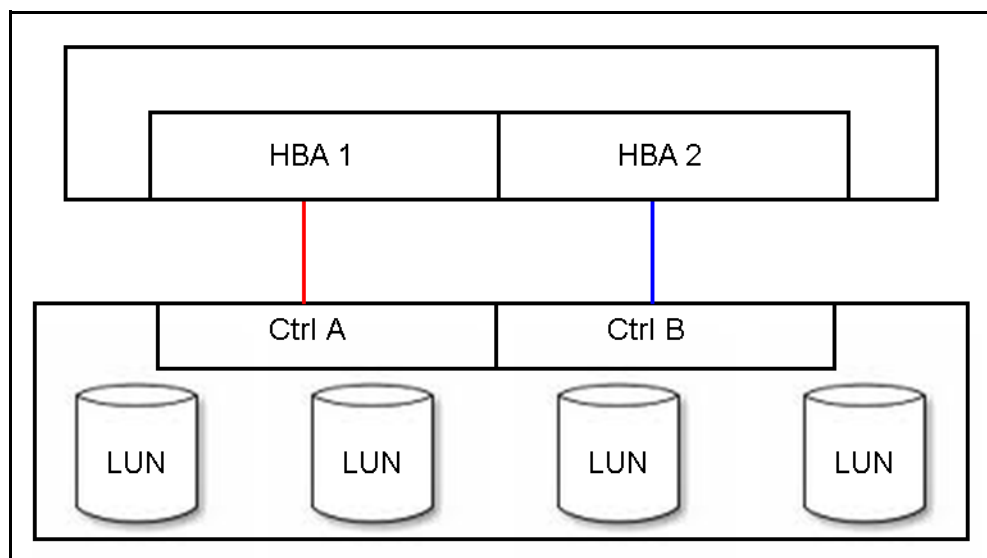


Figure 3-10 Load sharing approach for multiple HBAs

The RDAC drivers for Linux supports round-robin load balancing.

Note: In a cluster environment, you need a single path to each of the controllers (A and B) of the DS3500 Storage System. However, if the cluster software and host application can do persistent reservations, you can keep multiple paths and the multipath driver will route the I/O request using the appropriate path to the reserved logical drive.

Queue depth

There are a number of locations in the SAN solution that can have queue depth settings within their control. The queue depth is the maximum number of IO commands that can be queued for processing at the system, SAN network, or the storage at one point in time.

For the host this value might be adjustable at the device level and HBA level. Certain host operating systems only use the HBA level setting. Care should be taken when setting these levels because response time performance can be impacted by too high a value, and *device busy* status being returned by the storage when its capabilities are exceeded.

For QLogic based HBAs, the queue depth is known as *execution throttle*, which can be set with either QLogic SANsurfer or in the BIOS of the QLogic-based HBA by pressing Ctrl+Q during the boot process.

For the storage there are values at the controller and drive level. These values can vary between code levels and performance enhancement features.

For the latest firmware for the DS3500 controller, see the following website:

<http://www.ibm.com/systems/storage/disk/>

3.2.6 Disk expansion enclosures

The DS3500 Storage Systems offer the EXP3500 expansion enclosures for expanding the subsystem beyond the internal drive count of the DS3500. There are two models of expansions to choose from: the EXP3512 and the EXP3524. These expansions connect to the DS3500 storage system through the 6 Gbps SAS expansion ports on the controller modules. Figure 3-11 and Figure 3-12 on page 47 are examples of the cable connections needed to attach the DS3500 with dual controllers to an EXP3500.




Figure 3-11 DS3500 and EXP3500 cable attachments for dual controllers




Figure 3-12 DS3500 and multiple EXP3500 cable attachments for dual controllers

In looking at Figure 3-12 it is important to notice the direction of the cabling scheme used when adding the additional expansions to the loops. The DS3500 uses a “top down-bottom up” cabling configuration. This aids in ensuring the availability of the arrays when they are configured in the “zig-zag” pattern in the event of an enclosure being lost in the configuration. See 3.3.3, “Array configuration” on page 57 for further details on array layouts.

When planning for which enclosure to use with your DS3500, you must look at your applications and the workload type that you will need to provide. With transaction-based applications, the workload will need to be able to perform a high number of IO, and therefore having the greater number of spindles provided by the EXP3524 can prove to be a great advantage. With throughput-based applications, the need is generally for higher capacity drives and more bandwidth. With the new larger drives, spindle count does not need to be high to meet this goal. For this environment the EXP3512 can provide the disk capacity at a lower price point. Careful planning can allow for the mixing of these two expansions in the storage system to meet a mixed environment need.

The EXP3512 and EXP3524 enclosures support both SAS and SAS nearline 6 Gbps drives. The EXP3512 drives have a 3.5 inch footprint, whereas the EXP3524 drives have a 2.5 inch.

Important: There is no migration path of disks from the previous DS3000 products to the new DS3500 because the 6 Gbps hardware will not support the earlier 3 Gbps hardware.

Enclosure IDs

It is important to correctly set the tray (enclosure) IDs. They are used to differentiate multiple EXP enclosures that are connected to the same DS3500 Storage System. Each EXP enclosure must use a unique value. The DS3500 Storage Manager (SM) uses the tray IDs to identify each EXP enclosure.

For the EXP3500, the enclosure ID is indicated by a two seven-segment LEDs located on the back of each ESM next to the other ESM indicator lights. The storage server firmware automatically sets the enclosure ID number. If needed, you can change the enclosure ID

setting through the DS3500 storage management software only. There are no switches on the EXP3500 or EXP810 chassis to manually set the enclosure ID.

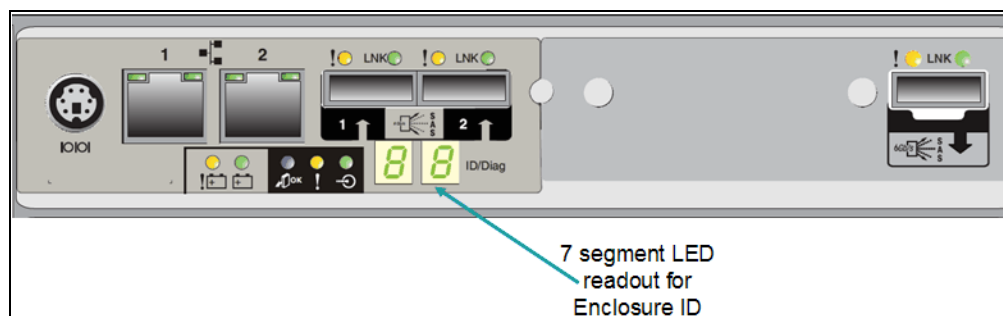


Figure 3-13 Enclosure ID LEDs for EXP3500

Enclosure guidelines

The base controller unit and each expansion enclosure has an ID number associated with it. The ID allows each enclosure to be identified properly to the base controller unit.

Because the base and all enclosures are connected by Fibre Channel loop, each ID address must be distinct and unique for I/O to flow properly. An ID address is composed of two digits: a tens digit (x10) and a ones digit (x1). Enclosure IDs are typically between the values of x00 and x77.

The EXP3500 is automatically assigned enclosure IDs.

Because the storage system follows a specific address assignment scheme for the drives, you have to observe a few guidelines when assigning enclosure IDs to expansion units if you were to assign these manually. Failure to adhere to these guidelines can cause issues with I/O error recovery, and make the troubleshooting of certain drive communication issues more difficult:

- ▶ Whenever possible, maintain an even number of expansion units on the DS3500 Storage System and configure it for enclosure loss protection. Add EXP expansion units in pairs and for the DS3500 Storage System (ideally) in groups of four.
- ▶ Add drives into an expansion enclosure in pairs.

3.3 Planning your storage structure

It is important to configure a storage system in accordance to the needs of the user workload. An important question and primary concern for most users or storage administrators is how to configure the storage subsystem to achieve the best performance. There is no simple answer to fit all cases, and no single guideline for storage performance optimization that is valid in every environment and for every particular situation. In this section we provide a general (and less detailed) performance discussion.

Also, in this section, we discuss the use of RAID protection, and review other aspects of the system configuration that can help optimize the storage capacity and resilience of the system. In particular, we review array configuration and enclosure loss protection.

3.3.1 Selecting drives

The speed and the type of the drives used will impact the performance. Typically, the faster the drive, the higher the performance. This increase in performance comes at a cost: the faster drives typically cost more than the lower performance drives.

The DS3500 Storage System currently supports the following types of SAS drives for the two models of chassis as shown in Table 3-2.

Table 3-2 DS3500 families HDD support

Drives Supported	DS3512/EXP3512	DS3524/EXP3524
SAS 15KRPM	300 GB, 450 GB, 600 GB	73 GB, 146 GB
SAS 15KRPM FDE	600 GB	None
SAS 10KRPM	None	300 GB
SAS 10KRPM FDE	None	300 GB
Nearline SAS 7.5KRPM	1 TB, 2 TB	500 GB
Maximum drives	12/96	24/96
Storage system capacity (max)	450 GB SAS / 1 TB SATA	450 GB SAS / 1 TB SATA

The Full Disk Encryption (FDE) drives are drives with built-in disk encryption hardware that prevents unauthorized access to the data on a drive that is physically removed from the storage subsystem.

General Rule: Generally it is best to use the fastest drives available for best performance. This can be critical to transaction-based (high IOPS) workloads.

The speed of the drive is measured by the number of revolutions per minute (RPM). A 15K drive rotates 15,000 times per minute. With higher speeds, the drives tend to be denser because a large diameter plate driving at such speeds is likely to wobble. With faster speeds, greater throughput is possible.

Seek time is the measure of how long it takes for the drive head to move to the correct sectors on the drive to either read or write data. It is measured in thousands of a second (milliseconds or ms). The faster the seek time, the quicker data can be read from or written to the drive. The average seek time decreases as the speed of the drive increases. Typically, a 7.2K drive will have an average seek time of around 9 ms, a 10K drive will have an average seek time of around 5.5 ms, and a 15K drive will have an average seek time of around 3.5 ms.

Command queuing (or queue depth) allows for multiple commands to be outstanding to the disk drive at the same time. The drives have a queue where outstanding commands can be dynamically rescheduled or re-ordered, along with the necessary tracking mechanisms for outstanding and completed portions of workload. The DS3500 provides a drive command queue depth of four operations per disk. The *High Performance Tier* increases the queue depth for all drives to 16.

Avoid using the SAS nearline drives for high IOPS operations. SAS nearline can, however, be used for streaming and archiving applications. These are both good uses for the slower RPM drives, where high throughput rates are required at a lower cost. If properly configured, these

drives can push high throughput numbers with large host IO block sizes and sequential workloads.

3.3.2 Understanding RAID types

In this section we introduce arrays, logical drives and associated terminology, and then describe the various RAID levels that are supported by the IBM System Storage DS3500 storage subsystem. RAID is an acronym for Redundant Array of Independent Disks, and is a storage solution in which part of the total storage capacity is used to store redundant information about user data stored on the remainder of the storage capacity.

RAID relies on a series of configurations, called levels, to determine how user data and redundancy data are written to and retrieved from the drives. RAID Level 1, RAID Level 10, RAID Level 3, RAID Level 5, and RAID Level 6 write redundancy data to the drive media for fault tolerance. The redundancy data might be an exact copy of the data (mirrored) or an error correcting code derived from the data. If a drive fails, you can use the redundancy data to quickly reconstruct information on a replacement drive.

DS3500 arrays and RAID levels

An array is a set of drives that the system logically groups together to provide one or more logical drives to an application host or cluster. The DS3500 storage subsystem supports RAID levels 0, 1, 10, 3, 5 and 6. Each of these RAID levels offers a compromise between capacity, performance, and data redundancy. The attributes of each of these RAID levels is described in more detail over the following pages of this book.

Note that the maximum number of physical drives in a RAID 0, 1 or 10 array is limited by the maximum number of physical drives that can be installed in a fully populated DS3500 storage subsystem, which is 96 drives. The maximum number of physical drives in a RAID 3, 5, or 6 array is always 30 drives.

The DS3500 storage subsystem is able to dynamically change the RAID level without requiring downtime. This feature is called Dynamic RAID Migration (DRM).

Each RAID array contains one or more associated logical drives. A logical drive is the basic structure that you create to store data on the storage subsystem. Each logical drive appears as a separate physical drive to the operating system on the host server. The DS3500 storage subsystem supports a maximum of 256 logical drives for the entire storage subsystem.

We now briefly describe the main features of each of the various RAID levels that are supported by the DS3500 storage subsystem.

RAID 0: Data striping

RAID 0 (Figure 3-14 on page 51) is also known as *data striping*. In this RAID level, the data is striped sequentially across all participating physical drives. RAID Level 0 is only designed to increase performance and has no data redundancy. The DS3500 storage subsystem supports a minimum of one drive and a maximum of 96 drives in a RAID 0 array.

RAID 0 is well-suited for applications that require fast access to non-critical data, or data that can be easily restored from backup. Increasing the number of disk drives in the array will increase the data access performance of the array.

Attention: The failure of a single disk in a RAID 0 array will cause the failure of the entire array and all of the associated logical drives, and access to all data on the array will be lost. For this reason, never use RAID Level 0 for critical applications that require high availability.

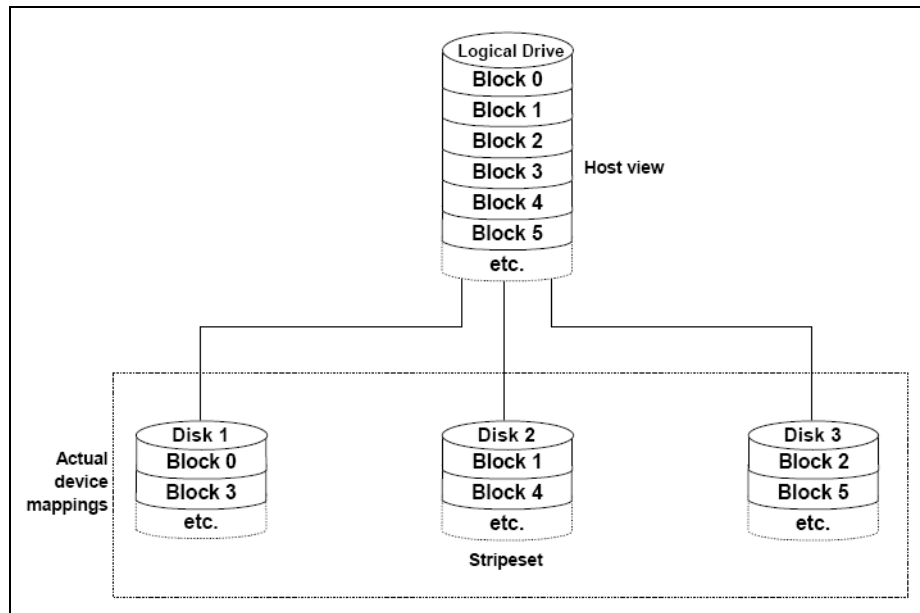


Figure 3-14 RAID 0

RAID 1 and RAID 10: Disk mirroring and disk mirroring with striping

RAID 1 is also known as *disk mirroring* and is a mirrored pair of drives without parity. RAID Level 1 uses exactly two drives to mirror the data between them. A RAID 10 (Figure 3-15 on page 52) array is automatically created when you create a RAID 1 array with four or more drives (two pairs of drives). RAID 10 is also known as RAID 1+0 and *disk mirroring with striping*, and it implements block interleave data striping and mirroring. In RAID 10, data is striped across the physical disk drives, and each of those drives is then mirrored to a second drive. The DS3500 storage subsystem supports a maximum of 96 drives in a RAID 10 array.

Note: RAID Level 1 is a specific implementation of RAID Level 10 that uses exactly two drives to mirror the data between them. A RAID Level 10 array is automatically created when you select four or more drives in a RAID 1 array.

RAID Levels 1/10 provide good redundancy; in the case of a single disk failure in each mirrored pair, the array and associated logical drives become degraded but all the data is still available and accessible from the second drive of the mirrored pair.

For each pair of mirrored drives, read operations can be performed from either physical disk of the mirrored pair. Write operations are performed by writing to both physical disks of the mirrored pair. In this manner small blocksize writes can be completed quickly, making this RAID type a great solution for a high write-intensive application when data protection is desired. For this type of application, this RAID type is generally preferred by database administrators.

However, because the data is mirrored, the capacity of the associated logical drives on a RAID 1 or RAID 10 array is 50% of the physical capacity of the hard disk drives in the array.

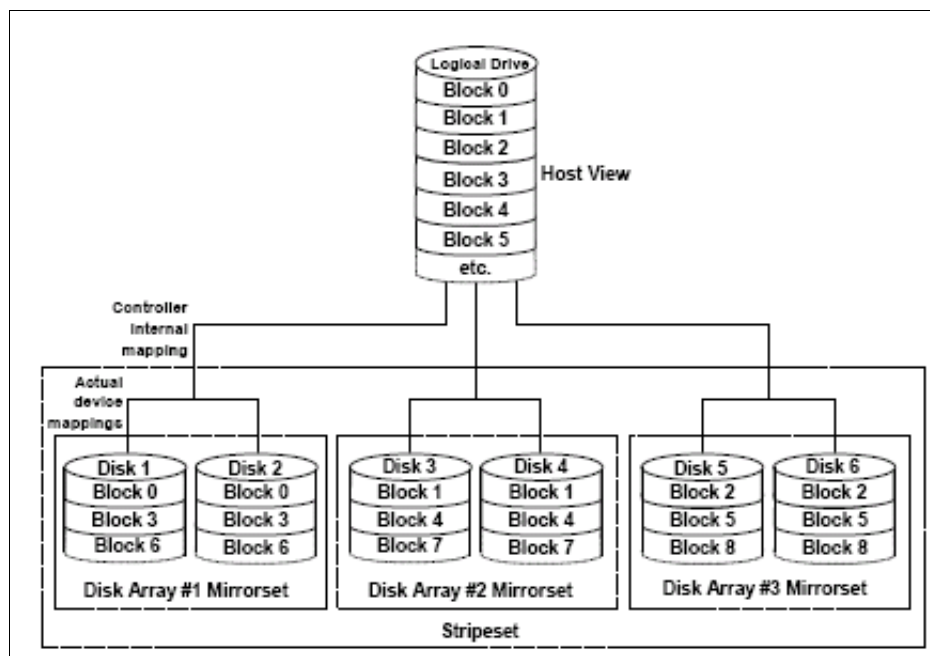


Figure 3-15 RAID 10. A RAID 1 array will only consist of one mirrored pair (#1 Mirrorset)

When using RAID 1/10 arrays, do the following:

- ▶ Use a two drive RAID 1 array for the disks that contain your operating system. It is a good choice because the operating system can usually fit on one disk.
- ▶ Use RAID 1 for transaction logs. Typically, the database server transaction log can fit on one disk drive. In addition, the transaction log performs mostly sequential writes. Only rollback operations cause reads from the transaction logs. Therefore, you can achieve a high rate of performance by isolating the transaction log on its own RAID 1 array.
- ▶ Use write caching on RAID 1/10 arrays. Because a RAID 1/10 write will not complete until both writes have been completed on both disks in a mirrored pair, the performance of writes can be improved through the use of a write cache. Always ensure that you use battery-backed write cache.
- ▶ The performance of RAID 10 is comparable to RAID 0 for sequential I/Os, but RAID 10 provides data redundancy through disk mirroring.

Note: There are no guaranteed choices as to which type of RAID to use because this is dependent on the workload read and write activity. A good general guide might be to consider using RAID 1 if random writes exceed about 25%, with a peak sustained I/O rate that exceeds 50% of the storage subsystem's capacity

When comparing RAID Level 10 with RAID Level 5:

- ▶ RAID 10 uses two write I/Os to write a single block of data (one write I/O to each drive in the mirrored pair). RAID 5 requires two read I/Os (read original data and parity) and then two write I/Os to write the same block of data. For this reason, random writes are significantly faster on RAID 10 compared to RAID 5.
- ▶ RAID 10 rebuilds take less time than RAID 5 rebuilds. If one drive fails, RAID 10 rebuilds it by copying all the data on the mirrored drive to a replacement/hotspare drive. RAID 5

rebuilds a failed disk by merging the contents of the surviving disks in an array and writing the result to a spare.

RAID 3: Data striping with a dedicated parity drive

A RAID 3 array uses data striping with a dedicated parity drive. Similar to RAID 0 data striping, information written to disk is split into chunks (a fixed amount of data), and each chunk is written out to the same physical position on separate disks (in parallel). This architecture requires parity information to be written for each stripe of data. RAID 3 uses a dedicated physical drive for storing parity data. If any one disk drive in the array fails, the array and associated logical drives become degraded, but all data is still accessible by the host application.

However, with RAID 3, the dedicated parity drive is a performance bottleneck during writes. Because each write operation requires the parity to be re-computed and updated, this means that the parity drive is accessed every time a block of data is written to the array. Because of this, RAID 3 is rarely used today in the industry and RAID 5 has taken its place. The DS3500 storage subsystem supports a maximum of 30 drives in a RAID 3 array.

RAID 5: Data striping with distributed parity

Like RAID Level 3, RAID Level 5 also uses parity for data protection but unlike RAID 3 it does not use a dedicated parity drive. Instead, the parity blocks are evenly distributed across all physical disk drives in the array, as shown in Figure 3-16. The failure of a single physical drive in a RAID 5 array will cause the array and associated logical drives to be degraded, but all the data will remain accessible to the host application. This level of data redundancy is known as $n+1$ redundancy because the data remains accessible after a single drive failure. When you create a RAID 5 array, the capacity of the array is reduced by the equivalent capacity of one drive (for parity storage). The DS3500 storage subsystem requires a minimum of three and supports a maximum of 30 drives in a RAID 5 array.

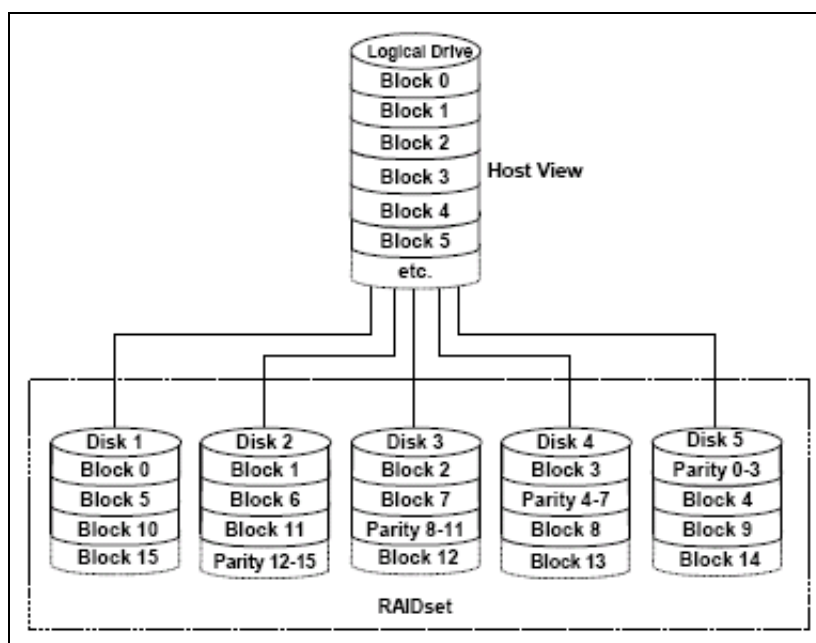


Figure 3-16 RAID 5

RAID Level 5 is best used in environments requiring high availability and fewer writes than reads.

RAID Level 5 can be 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. Applications with a low read percentage and a high random write percentage might not perform as well on RAID 5 logical drives because parity data must be recalculated for each write operation and then written to each drive in the array.

Use write caching on RAID 5 arrays, because each RAID 5 write will not be completed until at least two read I/Os (one data one parity) and two write I/Os (one data and one parity) have occurred. The write performance penalty can be mitigated by using battery-backed write cache. RAID 5 arrays with caching can give as good as performance as any other RAID level, and with certain workloads, the striping effect can provide better performance than RAID 1/10. Applications that require high throughput sequential write I/Os are an example of one such workload. In this situation, a RAID Level 5 array can be configured to perform just one additional parity write when using “full stripe writes” (also known as “full stride writes”) to perform a large write I/O when compared to the two writes per data drive (self and its mirror) that are needed for each write I/O with a RAID 1 array.

You must configure the RAID 5 array with certain number of physical drives to take advantage of full stripe writes. This is illustrated for the case of a RAID 5 array with 8 total drives (7 data + 1 parity) in the Figure 3-17.

A: Total # of drives written	B: Total # of RAID 5 Read Data I/Os	C: Total # of RAID 5 Read Parity I/Os	D: Total # of RAID 5 Write Data I/Os	E: Total # of RAID 5 Write Parity I/Os	X: Total # of I/Os for RAID 5 (= B+C+D+E)	Y: Total # of I/Os for RAID 1 (= Ax2)
1	1	1	1	1	4	2
2	2	1	2	1	6	4
3	3	1	3	1	8	6
4	3	0	4	1	8	8
5	2	0	5	1	8	10
6	1	0	6	1	8	12
7	0	0	7	1	8	14

Figure 3-17 Table illustrating the potential performance advantages of RAID 5 full stripe writes

Column A lists the number of drives that are being written to. Column Y is the number of write I/Os for a RAID 1 and will always be twice the value of A. Columns B, C, D, and E contain the numbers of read data/parity and write data/parity I/Os required for the number of drives that are being written to. You can see that for seven drives, no read I/Os are required for RAID 5 arrays because the full stripe is being written at once. This substantially reduces the total number of I/Os (column X) required for each write operation.

The decrease in the overhead read operations with the full stripe write operation is the advantage you are looking for. You must be careful when implementing this type of layout to ensure that your data pattern does not change, which might decrease its effectiveness. However, this layout might work well for you in a large sequential write environment. Due to the small size of segments, reads might suffer, so mixed I/O environments might not fare well, which might be worth testing if your writes are high.

When the IBM DS system storage detects that it is receiving contiguous full stripe writes, it will switch internally to an even faster write capability known as Fast Stripe Write Through. In this method of writing, the DS system storage uses the disk as the mirror device for the cache

write and shortens the write process. This method of writing can increase throughput as much as 30% on the system storage.

This does however, require that the following rules are being met by the IO pattern:

- All write IO's are full stripe writes (no partial stripe writes can be requested)
- Write IO's are sequential and contiguous in nature so no seeks are required.

If any interruptions in this pattern are detected, the writes will revert back to the standard full stripe write model, which nevertheless still gives a benefit to the RAID5 for large sequential writes over the RAID1.

RAID 6: Data striping with dual distributed parity

RAID 6 (see Figure 3-18) is a RAID level with dual rotational parity that is distributed across the drives in the array. A RAID 6 array has $n+2$ redundancy, which means that data remains accessible to the host application after two concurrent disk drives failures in the array. RAID 6 achieves $n+2$ redundancy because it calculates two sets of parity information for each block of data (P+Q) that is striped across the disks. The DS3500 storage subsystem performs the P+Q parity calculations in hardware.

There is no performance penalty for read operations from a RAID 6 array, but there is a performance penalty for write operations because two sets of parity information (P+Q) must be calculated for each write operation. The write penalty in RAID Level 6 can be mitigated by using battery-backed write caching.

The DS3500 storage subsystem requires a minimum of five drives and supports a maximum of 30 drives in a RAID Level 6 array.




Figure 3-18 RAID 6

The general characteristics of a RAID Level 6 array are:

- It uses disk striping with dual distributed parity: two independent parity blocks per stripe (P+Q) that are calculated in hardware for each stripe of data.
- It has high data ($n+2$) redundancy: the array and associated logical drives can survive the simultaneous loss of two physical disks without losing data.

- ▶ It requires two sets of parity data for each write operation, resulting in a significant decrease in write performance. Using battery-backed write cache can mitigate the impact of this write penalty.
- ▶ It has additional costs because of the extra disk capacity required by using two parity blocks per stripe: compared to a RAID 5 array: additional drives are required to achieve the same usable logical drive capacity.
- ▶ Any application that has high read request rates and average write request rates such as Transaction servers, web servers, data mining applications, and Exchange servers will benefit from RAID 6.

RAID Levels summary

In this section we summarize the general characteristics of the various RAID levels supported by the DS3500 storage subsystem. The following note and Table 3-3 summarize this information.

Summary: The general performance characteristics of the various RAID levels are:

- ▶ RAID 0 offers high performance, but does not provide any data redundancy.
- ▶ RAID 1/10 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 offers high availability with performance slightly lower than RAID 5.

Table 3-3 RAID levels comparison

RAID	Description	Application	Advantage	Disadvantage
0	Stripes data across multiple drives.	IOPS Mbps	Performance, due to parallel operation of the access.	No redundancy. If one drive fails, the data is lost.
1/10	The drive data is mirrored to another drive.	IOPS	Performance, as multiple requests can be fulfilled simultaneously.	Storage costs are doubled.
3	Drives operate independently with data blocks distributed among all drives. Parity is written to a dedicated drive.	Mbps	High performance for large, sequentially accessed files (image, video, and graphics).	Degraded performance with 8-9 I/O threads, random IOPS, and smaller, more numerous IOPS.
5	Drives operate independently with data and parity blocks distributed across all drives in the group.	IOPS Mbps	Good for reads, small IOPS, many concurrent IOPS, and random I/Os.	Writes are particularly demanding.

RAID	Description	Application	Advantage	Disadvantage
6	Stripes blocks of data and parity across an array of drives, and calculates two sets of parity information for each block of data.	IOPS Mbps	Good for multi-user environments, such as database or file system storage, where typical I/O size is small, and in situations where additional fault tolerance is required. This is the most reliable RAID level on the DS5000 storage subsystem	Slower in writing data, complex RAID controller architecture.

RAID Levels reliability considerations

At first glance, both RAID 3 and RAID 5 appear to provide good protection against drive failure. With today's high-reliability drives, it seems unlikely that a second drive in an array will fail (causing data loss) before an initial failed drive can be replaced. But if you look at RAID 6 and calculate the possibility of data loss, the chance to loose data is theoretically much less than on RAID 3 and RAID 5.

However, field experience has shown that when a RAID 3 or RAID 5 array fails, it is not usually due to two drives in the array experiencing complete failure. Instead, most failures are caused by one drive going bad, and a single block somewhere else in the array that cannot be read reliably.

This problem is exacerbated by using large arrays with RAID 5. This *stripe kill* can lead to data loss when the information to rebuild the stripe is not available. The end effect of this issue will of course depend on the type of data and how sensitive it is to corruption. Although most storage subsystems (including the DS3500 storage subsystem) have mechanisms in place to try to prevent this from happening, they cannot work 100% of the time.

Any selection of RAID type should take into account the cost of downtime. Simple calculations tell us that RAID 3 and RAID 5 are going to suffer from failures more often than RAID 10. Exactly how often is subject to many variables and is beyond the scope of this book. The money saved by economizing on drives can be easily overwhelmed by the business cost of a crucial application going down until it can be restored from backup.

No data protection method is 100% reliable, and even if RAID were faultless, it will not protect your data from accidental corruption or deletion by program error or operator error. Therefore, all crucial data should be backed up by the appropriate software, according to business needs.

3.3.3 Array configuration

Before you can start using the physical disk space, you must configure it. Divide your (physical) disk drives into arrays and create one or more logical drives inside each array.

In simple configurations, you can use all of your drive capacity in a single array and create all of your logical drives in that particular array. However, this method has following drawbacks:

- ▶ If you experience a (physical) drive failure, the rebuild process affects all logical drives, and the overall system performance goes down.
- ▶ Read/write operations to various logical drives are still being made to the same set of physical hard drives.

The array configuration is crucial to performance. You must take into account all the logical drives inside the array, because all logical drives inside the array will impact the same physical disks. If you have two logical drives inside an array and they both are high throughput, then there can be contention for access to the physical drives as large read or write requests are serviced. It is crucial to know the type of data that each logical drive is used for and try to balance the load so contention for the physical drives is minimized. Contention is impossible to eliminate unless the array only contains one logical drive.

Number of drives

In the transaction intense environment, it is more important to ensure that there are enough disk drives configured to perform the I/Os demanded by the host application, than to focus on the amount of possible storage space on the storage subsystem.

Use Table 3-2 on page 49 to determine what drive type, speed, and capacity is best suited for your workload and environment needs.

Transaction intensive workload

In a transaction intense environment, you want to have higher drive numbers. This can be done by creating larger arrays with more disks. The storage subsystems can have a maximum of 30 drives per RAID 5 array/logical drive, and 96 drives for a full subsystem RAID 10 array.

With a RAID 10 array, the logical drive size can be configured to encompass the entire array, although operating system limitations on maximum logical drive size might restrict the usefulness of this capability. Although there are circumstances where this model can work well, it is best that smaller arrays be used to better balance the drive usage and access patterns across the controllers and to avoid contention and intermixed IO patterns.

Configuring multiple LUNs striped across a single large array can be used to make use of all the capacity. However, with this layout, consideration should be given to the workload types for which these LUNs will be used, so as not to mix throughput and transaction based IO on the same array.

Another factor to consider is congestion when accessing the drives on the back-end loops. This situation can be avoided by using multiple arrays.

Generally, an array of 8 to 16 disks provides the best performance for RAID 5 workloads that are OLTP based.

General rule: For high transaction environments requiring highest redundancy and protection, logical drives should be built on arrays with 8 to 12 disks when using RAID 5 or RAID 10. Spreading the arrays evenly across the two controllers in a dual controller environment will give you the best workload balance.

For large databases, consider using the host volume management software to spread the workload evenly across multiple arrays/LUNs to evenly balance the workload. Build the volume across sets of logical drives laid out per the RAID type in the previous discussion. By using multiple arrays, you will also be able to increase the number of controllers that are involved in handling the load, therefore getting full use of the storage subsystems resources.

For example, if you need to build a database that is 1 TB in size, you can use five 300 GB drives in a 4+1 parity RAID 5 single array/logical drive, or you can create two RAID 5 arrays of 8+1 parity using 73 GB drives, giving two 584 GB logical drives on which to build the 1 TB database. In most cases the second method for large databases will work best, as it brings twelve more disks into play for handling the high host application transaction workload.

Large throughput workload

In the large throughput environment, it typically does not take high numbers of disks to reach the maximum sustained throughput. Considering that this type of workload is usually made of sequential I/O, which reduces disk latency, in most cases about 20 to 28 drives are enough to reach the maximum throughput.

This does, however, require that the drives be spread evenly across the DS3500 to best utilize the system bandwidth. The storage subsystem is optimized in its firmware to give increased throughput when the load is spread across all parts. Here, bringing all the DS3500 resources into play is extremely important. Keeping the drive loops and bus busy with high data throughput is the winning answer. This is also the perfect model for using the high capacity drives because we are looking to push a large volume of data and it will likely be large blocks of sequential reads and writes.

Consider building smaller arrays that are 4+P or 8+P in size with single logical drives for higher combined throughput. If multiple logical drives are to be created on the array, do not exceed the number of data drives in the array. The higher the number of logical drives, the greater the chance for contention for the drives.

General rule: For high throughput, logical drives should be built on arrays with 4+1, or 8+1 drives in them when using RAID 5. Data drive number and *segment size* must equal host I/O blocksize for full stripe write. Use multiple logical drives on separate arrays for maximum throughput.

An example configuration for this environment is to have a single logical drive /array with 16+1 parity 300 GB disks doing all the transfers through one single path and controller. An alternative consists of two 8+1 parity defined to the two controllers using separate paths, creating two separate streams of heavy throughput in parallel and filling all the channels and resources at the same time, which keeps the whole server busy with a cost of one additional drive.

Further improvements can be gained by splitting the two 8+1 parity into four 4+1 parity arrays giving four streams, but the addition of three drives is needed. A main consideration here is to plan for the array data drive count to be a number such that the host I/O blocksize can be evenly spread using one of the storage subsystem's segment size selections, which will enable the full stripe write capability discussed in the next section.

Regardless of the workload type or the RAID type being used for the array group, in all cases building the array using an equal number of odd and even drive slots is advantageous to the performance of the array and its LUNs. This is frequently done by using a "diagonal" or "orthogonal" layout across all the expansions used to attain enclosure loss protection as well as shown in Figure 3-20 on page 62.

3.3.4 Hot spare drives

A hot spare drive is like a replacement drive installed in advance. Hot spare disk drives provide additional protection that might prove to be essential in case of a disk drive failure in a fault tolerant array.

When possible, split the hot spares so that they are in separate enclosures and are not on the same drive loops (see Figure 3-19 on page 60).

Attention: SAS nearline drives cannot be used as hot spares for the SAS 10KRPM or SAS 15KRPM drives. When intermixing SAS and SAS nearline, create hot spares for both types of drives.

General rule: When assigning disks as hot spares, make sure that they have enough storage capacity. If the failed disk drive is larger than the hot spare, reconstruction is not possible. Use the largest drive size for the drive type installed as hot spare drives. See 10.3, “Set hot spare drive” on page 262 for more details on defining hot spare drives.

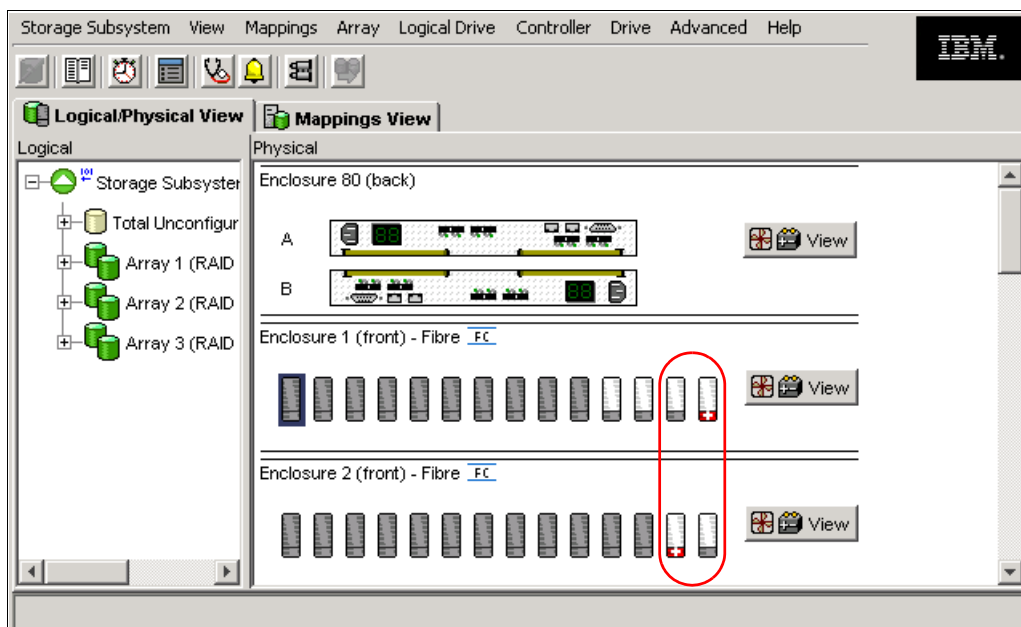


Figure 3-19 Hot spare coverage with alternating loops

Planning for 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, 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 hot 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 DS3500 configurations with a relatively low number of disk drives, one hot spare might suffice. But for larger configurations, define several hot spares.

General rule: For best configuration define one hotspare for every 18 - 20 disk drives in the system.

When defining hot spare drives, you need to consider the following rules:

- ▶ Hot spares must be of the same drive type (SAS or SAS nearline).
- ▶ Hot spares must be of the same or larger size as the failed drive.

For example, if you use a mixture of 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 of the same type in the system that it is protecting. This also means that in a mixed drive type environment you would need to have a minimum of two hot spares, one SAS and one SAS nearline.

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 IBM DS Storage Manager to configure these options.

3.3.5 Enclosure loss protection planning

Enclosure loss protection will enable your system to be more resilient against hardware failures. Enclosure loss protection means that you spread your protection arrays across multiple enclosures rather than in one enclosure so that a failure of a single enclosure does not take an array offline.

By default, the automatic configuration is enabled. However, we will use the manual method because it allows for more configuration options to be available at creation time.

General rule: Manual array configuration allows for greater control over the creation of arrays.

Figure 3-20 on page 62 shows an example of the enclosure loss protection. If enclosure number 2 fails, the array with the enclosure loss protection can still function (in a degraded state), because the other drives are not affected by the failure.

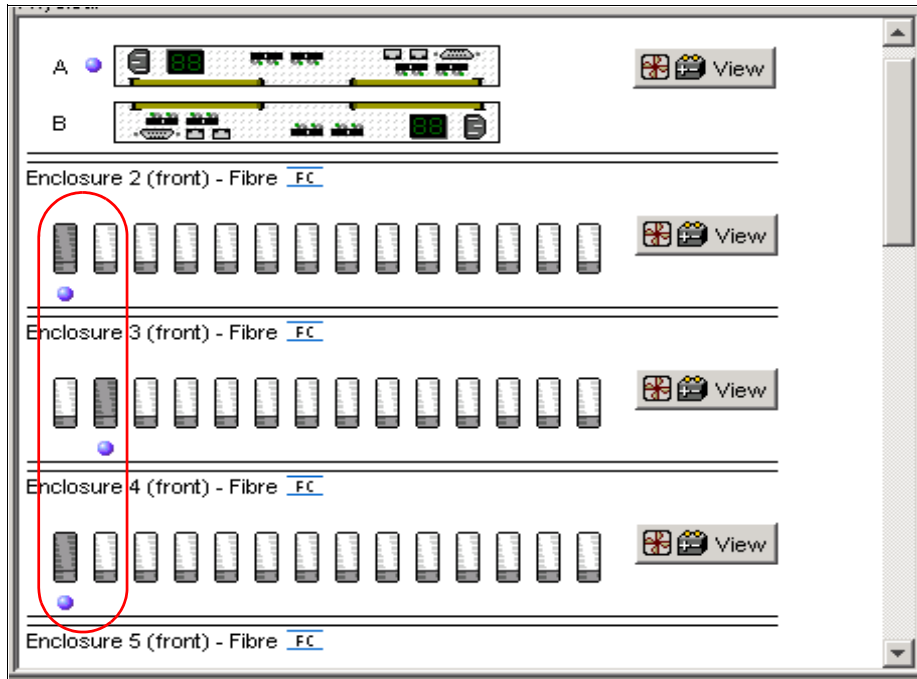


Figure 3-20 Example of enclosure loss protection

In the example shown in Figure 3-21, if enclosure number 1 fails, the entire array becomes inaccessible.




Figure 3-21 An array without enclosure loss protection

General rule: When possible, use enclosure loss protection for your arrays.

3.3.6 Logical drives and controller ownership

Logical drives, sometimes simply referred to as logical volumes or Logical Unit Numbers (LUNs), are the logical segmentation of arrays. A logical drive or LUN is a logical structure that you create on a storage subsystem for data storage. A logical drive is defined over a set of drives called an array group that has a defined RAID level and capacity. The attributes of the array are hidden from the host computer, which is only presented with logical drive details including capacity.

The DS3500 provides great flexibility in terms of configuring arrays and logical drives. However, when assigning logical drives to the systems, it is important to remember that the DS3500 uses a preferred controller ownership approach for communicating with logical drives, which means that every logical drive is owned by only one controller at a time. It is, therefore, important at the system level to make sure that traffic is correctly balanced between controllers with a dual controller DS3500 system, which is a fundamental principle for a correct setup of the storage system.

Balancing traffic

Balancing traffic is unfortunately not always a trivial task. For example, if an application requires a large disk space to be located and accessed in one chunk, it becomes harder to balance traffic by spreading the smaller volumes among controllers. This can sometimes be aided through the use of host based volume management software products.

In addition, typically, 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, hence the importance of monitoring the system.

General rule: Here are guidelines for logical drive assignment and storage partitioning:

- ▶ Assign logical drives across both controllers to balance controller utilization.
- ▶ Use the manual method of creating logical drives, which allows greater flexibility for configuration settings, such as enclosure loss protection and utilizing all drive loops.
- ▶ Avoid mixing workload types (transaction based and throughput based) on the same array of disks.
- ▶ Always leave a small amount of free space in the array after the logical drives have been created.

Enhanced remote mirror (ERM) considerations

A secondary logical drive in a remote mirror will always have the same preferred owner as the associated primary logical drive. For example, if controller A owns the primary logical drive in the primary storage subsystem, controller A owns the associated secondary logical drive in the secondary storage subsystem. If controller ownership changes on the primary logical drive, then this will cause a corresponding controller ownership change of the secondary logical drive also. For more details on ERM see the *IBM Midrange System Storage Copy Services Guide*, SG24-7822.

3.3.7 Storage partitioning

Storage partitioning adds a high level of flexibility to the DS3500 Storage System. The DS3500 comes with four partitions by default, which can be expanded to a maximum of 64. See 3.4, “Planning for premium features” on page 70 on for discussion on premium feature

needs. It enables you to connect multiple and heterogeneous host systems to the same storage server, either in stand-alone or clustered mode. The term storage partitioning is somewhat misleading because it actually represents a host or a group of hosts and the logical drives they access.

Without storage partitioning, the logical drives configured on a DS3500 Storage System can only be accessed by a single host system or by a single cluster, which can lead to inefficient use of storage server hardware unless the use of the DS3500 Storage System is dedicated to a single host (for example, SVC attachment, where it is seen as a single host).

Storage partitioning, however, allows the creation of “sets” containing the hosts with their host bus adapters and the logical drives. We call these sets *storage partitions*. 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 allows mapping and masks the logical drive or LUN (that is why it is also referred to as “LUN masking”), which means that after the logical drive is assigned to a host, it is hidden from all other hosts connected to the same storage server. Therefore, access to that logical drive is exclusively reserved for that host.

As a general rule, configure storage partitioning prior to connecting multiple hosts. Operating systems such as Windows will write their signatures to any device it can access.

Heterogeneous host support means that the host systems can run various operating systems. But be aware that all host systems within a particular storage partition have unlimited access to all logical drives assigned to the partition. Therefore, file systems or disk structure on these logical drives must be compatible with host systems. To ensure this, run the same operating system on all hosts within the same partition. In certain special cases it might be necessary to run operating systems that will be able to mount foreign file systems and share logical drives. In these cases care should be taken to ensure the partition members are properly defined and configured to share.

Storage partition topology is a collection of topological elements (default group, host groups, hosts, and host ports) shown as nodes in the topology view of the mappings view. To map a logical drive or LUN to a specific host server or group of hosts, each component of the storage partition must be defined.

A storage partition contains several components:

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

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

A *host* is a single system that can be mapped to a *logical drive*.

A *host port* is the FC port of the host bus adapter (HBA) on the host system. The host port is identified by its world-wide name (WWN). A single host can contain more than one host port. If the servers are attached using full redundancy, each server will have two host bus adapters, that is, it needs two host ports within the same host system. It is possible to have a host with a single HBA, but for redundancy, it must be able to access both DS3500 controllers, which can be achieved by SAN zoning.

The DS3500 Storage System only communicates through the use of the WWN. The storage system is not aware of which host bus adapters are in the same server or in servers that have a certain relationship, such as a cluster. The host groups, the hosts, and their host ports reflect a logical view of the physical connections of the SAN, and the logical connection between servers, such as clusters.

With the logical setup defined as previously described, mappings are specific assignments of logical drives to particular host groups or hosts.

The storage partition is the combination of all these components. It ensures correct access to the various logical drives even if there are several hosts or clusters connected.

The default host group is a placeholder for hosts that are defined but have not been mapped. The default host group is also normally used when storage partitioning is not enabled. If this is the case, then only one type of operating system can be sharing the logical drives.

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

With Storage Manager, it is possible to have up to 64 storage partitions on a DS3500, which allows the storage subsystem to provide storage capacity to a number of heterogeneous hosts, allowing for greater flexibility and scalability. Note that on DS3500 models, the number of partitions also depends on the premium feature licence you have purchased.

For the maximum number of logical drives per partitions for a specific host type, see Table 3-4.

Table 3-4 Host logical drive restrictions

Operating system	Maximum number of logical drives supported per partition
Windows Server 2003	32
Windows Server 2008	32
HP-UX	32
AIX	32
Linux	32

Every mapping of a logical drive to a new host or host group creates a new storage partition. If additional logical drives are required for an existing host or host group, a new storage partition is not required. For example, a cluster with two nodes with redundant I/O paths gets configured as one host group with two hosts. Each host then has two host ports for redundancy, and several logical drives are mapped to this host group. All of these components represent one storage partition. If another single host system is attached to the same storage system and other logical drives are mapped to that host, another storage partition then has to be created for it. If a new logical drive is created and mapped to either the cluster or the single host, it uses an existing storage partition.

Note: There are limitations as to how many logical drives you can map per host. DS3500 series storage servers will support up to 256 logical drives (including the “access” logical drive) per partition (although there are also restrictions, depending on the host operating system) and a maximum of two partitions per host. Keep these limitations in mind when planning the DS3500 Storage System installation.

Storage partitioning considerations

By default an “access” logical drive is mapped to each partition created. If in-band management is being used, then the access logical drive must be mapped to the storage partition for the managing host. For all others this logical drive can be removed.

In a security-sensitive environment, you can assign a password to the storage system as well.

Note: Each host with separately assigned storage will use a storage partition. Each host group with separately assigned storage to be shared by the group will also use a storage partition.

In order to configure the storage partitioning correctly, you need the WWN of your host HBAs. Mapping is done on a WWN basis. Depending on your HBA, you can obtain the WWN either from the BIOS or the QLogic SANsurfer tool if you have QLogic cards. Emulex adapters and IBM adapters for IBM System p and IBM System i® servers have a sticker on the back of the card. The WWN is also usually printed on the adapter itself or the box that the adapter was shipped in.

If you are connected to a hub or switch, check the Name Server Table of the hub or switch to identify the WWN of the HBAs.

When planning your partitioning, keep in mind that:

- ▶ By default the DS3500 comes with four partitions
- ▶ In a cluster environment, you need to use host groups.
- ▶ You can optionally purchase additional partitions (maximum of 64).

When planning for your storage partitioning, create a table of planned partitions and groups so that you can clearly map out and define your environment.

General Rule: If you have a single server in a host group that has one or more LUNs assigned to it, do the mapping to the host and not the host group. All servers with the same host type (for example, Windows servers) can be in the same group if you want, but by mapping the storage at the host level, you can define what specific server accesses which specific logical drives.

However, if you have a host cluster, you will need to assign the shared logical drives at a host group level so that all of the host servers in that host group have access to these logical drives for host cluster failover to work. This practice is also used when creating a partition for shared filesystem servers (GPFS™ for example).

Table 3-5 shows an example of a storage partitioning plan that shows the host groups, hosts, port names, WWN of the ports, and the operating systems used in that environment. Other columns can be added to the table for future references, such as HBA BIOS levels, driver revisions, and switch ports used, all of which can then form the basis of a change control log.

Table 3-5 Sample plan for storage partitioning

Host group	Host name	Port name	WWN	OS type
Windows 2003	Windows Host	MailAdp_A	200000E08B28773C	Windows 2003 Non-Clustered
		MailAdp_B	200000E08B08773C	
Linux	Linux_Host	LinAdp_A	200100E08B27986D	Linux
		LinAdp_B	200000E08B07986D	

POWER6®	AIX_Host	AIXAdp_A	20000000C926B6D2	AIX
		AIXAdp_B	20000000C926B08	

Heterogeneous hosts

When implementing a DS3500 Storage System solution, a mixture of host servers with various operating systems can be supported, and both clustered and non-clustered variants of the same operating systems. In general, all host accessing logical drives in a single storage partition would want to be configured for the same operating system. Very seldom would this not be the case, and great care must be taken when creating those environments.

Important: Non-shared storage can only be supported with storage partitioning enabled.

Delete the access logical drive (31)

The DS3500 Storage System will automatically create a logical drive for each host attached (logical drive id 31). This drive is used for in-band management, so if you do not plan to manage the DS3500 Storage System from that host, you should delete this logical drive. This will also give you one more logical drive to use for host user data.

Important: In all cases when not using in-band management you should delete the access logical driver for the host partition.

3.3.8 Segment size

The segment size is the maximum amount of data that is written or read from a disk per operation before the next disk in the array is used. Segment size can be quite different for different workload types. It is important to know the host IO blocksize that will be used when deciding on the segment size.

For small host I/Os, as mentioned earlier, set the segment size larger than the host I/O size. Doing this makes it unnecessary to access a second drive for a single small host I/O. For certain storage subsystems, having the segment size equal to the host I/O size is preferred, which is *not* the case with the Midrange Storage Subsystems. In many cases, using the segment size that is equal to the host IO size can result in multiple disk operations due to alignment differences. This can be easily remedied with a larger segment size selection for small IO block sizes.

There is no advantage in using a smaller segment size with RAID 1. Only in instances of large throughput block sizes does this help with RAID 5 (which we discuss later). Because only the data that is to be written to the disk is written to cache for an I/O, there is no cache penalty encountered either. As mentioned earlier on the host sections, aligning data on segment boundaries is important for performance. With larger segment sizes, there are less occasions to have misaligned boundaries impacting your performance, as more small I/O boundaries reside within a single segment, decreasing the chance of a host I/O spanning multiple drives. This technique can be used to help mask the effect of poor layout of the host data on the disks due to boundary differences.

General rule: For most high transaction workloads with the Midrange Storage Subsystems, the segment size of 128 KB (default) works best.

With high throughput workload, the focus is on moving larger but fewer I/Os. This workload is generally sequential in nature.

General rule: In the high throughput environment, you want the *stripe size* to be equal to, or an even multiple of, the host I/O size.

The total of all the segments for one pass of all the back-end data disks is a *stripe*. Therefore, large segment sizes that can equal the I/O size might be desired to accomplish the higher throughput you are looking for. For high read throughput, you want to have large segments (128 K or higher) to get the most from each stripe. For example if the host I/O is 512 KB, and the write is to a RAID 10 array, you want to use a segment size of 512 KB to limit the disk operations as much as possible.

However, when the workload consists of high writes and you are using RAID 5, use a method known as *full stripe (stride) write*, which can work well to improve your performance. With RAID 5 the parity is based on the value calculated for a stripe. As discussed earlier in 3.3.2, “Understanding RAID types” on page 50, when the I/O being written is spread across the entire stripe width, no reads are required to calculate the parity, and the I/O completes with fewer back-end I/Os being required.

This design can use a smaller segment size to align the host I/O size with the size of the stripe width. For instance with our previous example, if the array is a 4+P RAID 5 you want to use a 128 KB segment size to achieve a full stripe write. This type of management requires that a few host I/Os not equal a full stripe width.

Tips: The possible segment sizes available are 8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, and 512 KB:

- ▶ Storage Manager sets a default segment size to 128 KB.
- ▶ For database applications, a segment size of 128 KB or 256 KB (for data warehousing) have shown to be more effective.
- ▶ In a large file environment, such as media streaming or CAD, 128 KB or more are preferable with a focus on full stripe writes.
- ▶ For a web server or file and print server, a range of 64 KB to 128 KB should provide best results.

A performance testing schedule must be undertaken in the environment before going into production with a given segment size. Segment size can be dynamically changed, but only by rewriting the data, which consumes bandwidth and impacts performance. Plan this configuration carefully to avoid having to reconfigure the options chosen.

3.3.9 Media scan

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, which 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*.

Media scan continuously runs in the background, using spare cycles to complete its work. The default media scan is for a scan every 30 days, that is, the maximum time media scan will have to complete the task. During the scan process, the DS3500 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. After the media scan has completed, it will start over again 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 media scan

over host activity to ensure that the scan completes in the allocated time. This scan can impact on performance, but improve data integrity.

Media scan must be enabled for the entire storage subsystem. The 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 and media scan.

A media scan can be considered a surface scan of the hard drives, whereas a redundancy check scans the blocks of a RAID 3, 5, or 6 logical drive and compares it against the redundancy data. In the case of a RAID 1 logical drive, the redundancy scan compares blocks between copies on mirrored drives.

We have seen no effect on I/O with a 30 day setting unless the processor is utilized in excess of 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.

Note: If you change the media scan duration setting, the changes will not take effect until the current media scan cycle completes or the controller is reset.

3.3.10 Cache parameters

Cache memory is an area of temporary volatile storage (RAM) on the controller that has a faster access time than the drive media. This cache memory is shared for read and write operations.

Efficient use of the RAID controller cache is essential for good performance of the DS3500. See 9.5.2, “Change Cache Settings” on page 241 for detailed explanation about the cache parameter settings.

Cache blocksize selection

On the DS3500, the cache blocksize is a variable value that can be set to 4 K, 8 K, or 16 K. The general default setting is 8 K. The main goals with setting this value are to minimize the number of cache IO's needed, and at the same time, not waste space. This value is a storage subsystem wide parameter, and when set, the value is used by all cache operations.

For example, if the I/O of greatest interest is taken from your database operations during the day rather than from your weekly backups, you want to tune this value to handle the high transactions better. Knowing that the higher transactions will have smaller I/O sizes, using the 4 K settings is generally best for transaction intense environments.

General rule: Set the cache blocksize to 4 K for the DS3500 for transaction intense environments.

In a throughput intense environment, as we discussed earlier, you want to get as much data into cache as possible. In this environment it is generally best to use the 16 K blocksize for the cache.

General rule: Set the cache blocksize to 16 K for the DS3500 system for throughput intense environments.

In mixed workload environments, you must decide which workload type is most critical and set the system-wide settings to best handle your business needs.

General Rule: Set the cache blocksize to 8 K for the DS3500 for mixed workload environments.

Tip: Throughput operations, although impacted by smaller cache blocksize, can still perform reasonably if all other effects have been accounted for. Transaction based operations are normally the higher concern, and therefore must be the focus for setting the server wide values if applicable.

Cache flush control settings

In addition to the cache blocksize, the DS3500 also has a cache control that determines the amount of data that can be held in write cache. With the *cache flush* settings, you can determine what level of write cache usage will be reached before the server will start to flush the data to disk, and at what level the flushing will stop.

By default, these parameters are set to the value of “80” for each, which means that the server will wait until 80% of the write cache is used before it will flush the data to disk. In a fairly active write environment, this value might be far too high. You can adjust these settings up and down until you find a particular value that best suits your environment. If the values are not the same, then back-end drive inactive time increases, and you have surging with peaks and valleys occurring instead of a steady usage of back-end disks.

You can also vary the maximum amount of time the write data can remain in cache prior to being forced out, and written to disks. This value by default is set to ten seconds but can be changed by using the following Storage Manager (SM) command:

```
'set logical Drive [LUN] cacheflushModifier=[new_value];'
```

General rule: Begin with Start/Stop flush settings of 50/50, and adjust from there. Always keep the values for Start and Stop equal to each other.

3.4 Planning for premium features

The premium features that come with the DS3500 Storage Manager software are enabled by purchasing a premium feature license key for each feature. When planning for any of the premium features, it is a good idea to document what the goals and rationale for purchasing the feature are, which clearly defines from the outset what you want to achieve and why. There are many premium features available for different types of expanded capabilities. The following is a list of the currently available features for the DS3500:

- ▶ Storage Partitioning
- ▶ High Performance Tier
- ▶ Maximum 96 drive support
- ▶ RAID 6 logical drive support
- ▶ Full Disk Encryption
- ▶ External Key Management
- ▶ FlashCopy
- ▶ VolumeCopy
- ▶ Enhanced Remote Mirroring

Points that might play an important part with premium features and should be considered are the following:

- ▶ Which premium features will you need?
- ▶ What additional resources will they require?
- ▶ Additional configuration needs
- ▶ Amount of free space to support
- ▶ Retention of copies needed
- ▶ Automated or manual copies
- ▶ Disaster recovery or backup operation needs

Document the needs and requirements to support all the features being implemented. Many of these premium features are purely to enable special capabilities within the DS3500. However, features like the Copy Services can require additional resources to be able to implement them that need to be considered when planning and implementing your environment.

3.4.1 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 you can create it much more quickly than a physical copy. Additionally, it requires less disk space. In DS3500 Storage Manager, the logical drive from which you are basing the FlashCopy, called the base logical drive, must be a standard logical drive in the storage server. Typically, you create a FlashCopy so that an application (for example, an application to take backups) can access the FlashCopy and read the data while the base logical drive remains online and user-accessible. The DS3500 can have up to eight FlashCopy volumes.

Plan carefully with regard to the space available to make a FlashCopy of a logical drive even though FlashCopy takes only a small amount of space compared to the base image.

3.4.2 VolumeCopy

The VolumeCopy feature is a firmware-based mechanism for replicating logical drive data within a storage subsystem. This feature is designed as a system management tool for tasks, such as relocating data to other drives for hardware upgrades or performance management, data backup, and restoring snapshot logical drive data.

A VolumeCopy creates a complete physical replication of one logical drive (source) to another (target) within the same storage subsystem. The target logical drive is an exact copy or clone of the source logical drive. VolumeCopy can be used to clone logical drives to other arrays inside the DS3500 Storage System.

The VolumeCopy premium feature must be enabled by purchasing a Feature Key. For efficient use of VolumeCopy, FlashCopy must be installed as well.

3.4.3 Enhanced Remote Mirroring

The Enhanced Remote Mirroring (ERM) option is used for online, real-time replication of data between storage subsystems over a long distance. ERM mirrors by creating a full logical drive image on a secondary of the logical drive that is being mirrored from the primary. The DS3500 can have up to eight mirrored pair relationships. The DS3500 can mirror to another DS3500 or another member of the IBM Midrange DS Storage System family running 7.70.xx code or higher. At the current time these are the only supported configurations.

Operating modes for ERM

Enhanced Remote Mirroring (formerly named “Remote Volume Mirroring”) offers three operating modes:

- ▶ **Metro mirroring**

Metro mirroring is a synchronous mirroring mode. Any host write requests are written to the primary (local) storage subsystem and then to the secondary (remote) storage subsystem. The remote storage controller acknowledges the write request operation to the local storage controller, which reports a write completion to the host. This mode is called synchronous. The host application does not get the write request result until the write request has been executed on both (local and remote) storage controllers.

- ▶ **Global copy**

This mode copies a non-synchronous, remote copy function designed to complete write operations on the primary storage system before they are received by the secondary storage system. This capability is designed to prevent primary performance from being affected by wait time from writes on the secondary system. Therefore, the primary and secondary copies can be separated by long distances. This function is appropriate for remote data migration, offsite backups, and transmission of inactive database logs at virtually unlimited distances.

- ▶ **Global mirroring**

This mode is a two-site remote data mirroring function designed to maintain a complete and consistent remote mirror of data asynchronously at virtually unlimited distances with virtually no degradation of application response time. Separating data centers by longer distances helps provide protection from regional outages. This asynchronous technique can help achieve better performance for unlimited distances by allowing the secondary site to trail in data currency a few seconds behind the primary site. With Global Mirror, currency can be configured to be as little as three to five seconds with respect to host I/O. This two-site data mirroring function is designed to provide a high-performance, cost-effective global distance data replication and disaster recovery solution.

The Enhanced Remote Mirroring has also been equipped with new functions for better business continuance solution design and maintenance tasks.

A minimum of two storage subsystems is required. One storage subsystem can have primary volumes being mirrored to arrays on other storage subsystems and hold secondary volumes from other storage subsystems. Also note that because replication is managed on a per-logical drive basis, you can mirror multiple individual logical drives from a primary storage subsystem to appropriate secondary logical drives that are located on several separate remote storage subsystems. However, only one primary (source) and one secondary (target) member can exist in any mirrored pair relationship, and a logical drive cannot be a member of more than one mirror relationship at a time.

Planning considerations for ERM

The following are planning considerations to keep in mind:

- ▶ DS3500 Storage Systems (minimum of two)
- ▶ Switched fabric networks between sites
- ▶ Distances between sites (ensure that it is supported)
- ▶ Switches, directors or multiprotocol routers used for connections
- ▶ Redundancy
- ▶ Additional storage space requirements

Note: ERM requires a dedicated *switched fabric* connection per controller to be attached to Host port 4 on both A and B controllers of the DS3500 FC HIC option.

This same dedication is required at both the source and target ends of the ERM solution.

3.4.4 Drive Security

This is a new premium feature where Full Disk Encryption (FDE) protects the data on the disks only when the drives are removed from storage subsystem enclosures. Drive Security requires security capable drives (FDE) and provide access to data only through a controller that has the correct security key when Drive Security is enabled.

Requirements

Businesses must comply with a growing number of corporate standards and government regulations. Drive security is one tool that can enhance security to help comply with these new standards and regulations.

Full Disk Encryption

Full Disk Encryption (FDE) does not prevent someone from copying the data in the storage subsystems through Fibre Channel host port connections when the drives are unlocked and operating. FDE also does not prevent unauthorized management access. A security capable drive encrypts data during writes and decrypts data during reads. FDE prevents the physical removal of the disk from the DS3500 system and interpreting data it contained. The FDE drive with Drive Security enabled will be locked on power up and will only unlock after successful authentication with the DS3500 system.

The Encryption Key is generated by the drive and never leaves the drive, so it always stays secure. It is stored in encrypted form performing symmetric encryption and decryption of data at full disk speed with no impact on disk performance. Each FDE drive uses its own unique encryption key which is generated when the disk is manufactured and regenerated when required by the storage administrator using the DS3500 Disk Encryption Manager.

The security enabled drives can be used as normal drives and intermixed in an array with drives of equal type and capacity when this feature is not enabled. This new feature is detailed in *IBM Midrange System Storage Hardware Guide*, SG24-7676.

3.4.5 Obtaining premium features key

You can generate the feature key file by using the premium feature activation tool that is located at the following website:

<http://www-912.ibm.com/PremiumFeatures/jsp/keyInput.jsp>

The key can then be added to your DS3500 system as detailed in “Premium Features” on page 194.

3.5 Additional planning considerations

In this section, we review additional elements to consider when planning your DS3500 Storage Systems using a Logical Volume Manager and virtualization options.

3.5.1 Planning for systems with LVM: AIX example

Many modern operating systems implement the concept of a Logical Volume Manager (LVM) that can be used to manage the distribution of data on physical disk devices.

The Logical Volume Manager controls disk resources by mapping data between a simple and flexible logical view of storage space and the actual physical disks. The Logical Volume Manager does this by using a layer of device driver code that runs above the traditional physical device drivers. This logical view of the disk storage is provided to applications and is independent of the underlying physical disk structure.

Figure 3-22 illustrates the layout of those components in the case of the AIX Logical Volume Manager.

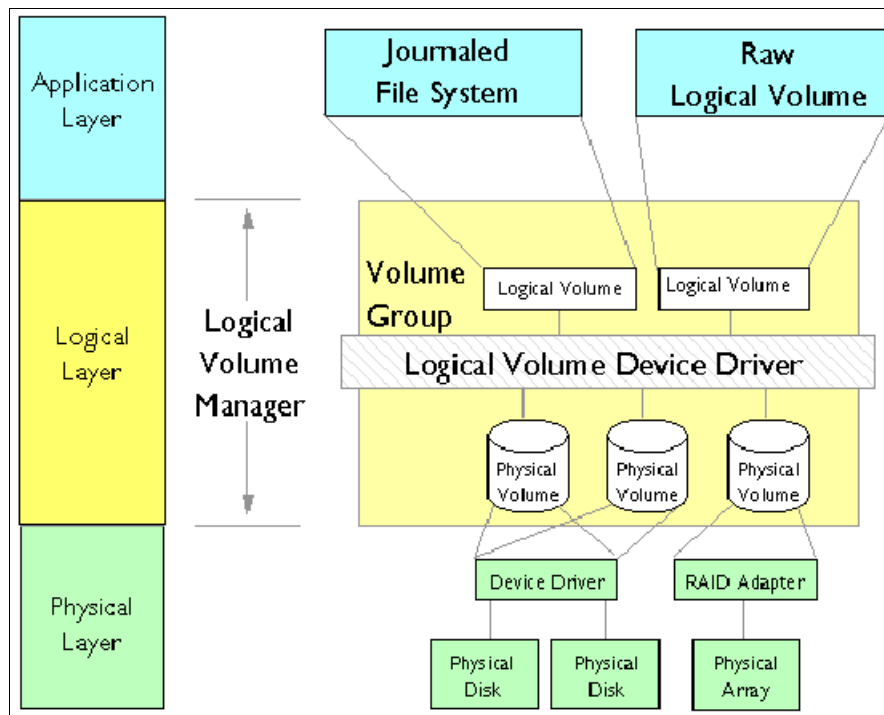


Figure 3-22 AIX Logical Volume Manager

Hierarchy of structures in disk storage

A hierarchy of structures is used to manage the actual disk storage, and there is a well defined relationship among these structures.

In AIX, each individual disk drive is called a physical volume (PV) and has a name, usually /dev/hdiskx where x is a unique integer on the system. In the case of the DS3500 Storage System, such physical volumes correspond to a LUN.

- ▶ Every physical volume in use belongs to a volume group (VG) unless it is being used as a raw storage device.
- ▶ Each physical volume is divided into physical partitions (PPs) of a fixed size for that physical volume.
- ▶ Within each volume group, one or more logical volumes (LVs) are defined. Logical volumes are groups of information located on physical volumes. Data on logical volumes appear contiguous to the user, but can be spread (striped) on multiple physical volumes.

- ▶ Each logical volume consists of one or more logical partitions (LPs). Each logical partition corresponds to at least one physical partition (see Figure 3-23). If mirroring is specified for the logical volume, additional physical partitions are allocated to store the additional copies of each logical partition.
- ▶ Logical volumes can serve a number of system purposes (paging, for example), but each logical volume that holds ordinary systems, user data, or programs contains a single journaled file system (JFS or JFS2). Each file system consists of a pool of page-size blocks. In AIX Version 4.1 and later, a given file system can be defined as having a fragment size of less than 4 KB (512 bytes, 1 KB, or 2 KB).





Figure 3-23 Relationships between LP and PP

The Logical Volume Manager controls disk resources by mapping data between a simple and flexible logical view of storage space and the actual physical disks. The Logical Volume Manager does this by using a layer of device driver code that runs above the traditional physical device drivers. This logical view of the disk storage is provided to applications and is independent of the underlying physical disk structure (Figure 3-24).



General rule: When using a DS3500 Storage System with operating systems that have a built-in LVM, or if a LVM is available, you need to make use of the LVM.

The AIX LVM provides a number of facilities or policies for managing both the performance and availability characteristics of logical volumes. The policies that have the greatest impact on performance in general disk environment are the intra-disk allocation, inter-disk allocation, write scheduling, and write-verify policies.

Because a DS3500 Storage System has its own RAID arrays and logical volumes, we do not work with real physical disks in the system. Functions such as intra-disk allocation, write scheduling, and write-verify policies do not help much, and it is hard to determine the performance benefits when using them. They must only be used after additional testing, and it is not unusual that trying to use these functions will lead to worse results.

On the other hand, we must not forget about the important inter-disk allocation policy.

Inter-disk allocation policy

The inter-disk allocation policy is used to specify the number of disks and how the logical partitions (LPs) are placed on specified physical volumes, which is also referred to as *range of physical volumes* in the **smitty mklv** panel:

- ▶ With an inter-disk allocation policy of minimum, LPs are placed on the first PV until it is full, then on the second PV, and so on.
- ▶ With an inter-disk allocation policy of maximum, the first LP is placed on the first PV listed, the second LP is placed on the second PV listed and so on, in a round robin fashion.

By setting the inter-physical volume allocation policy to maximum, you also ensure that the reads and writes are shared among PVs, and in systems such as a DS3500 Storage System, also among controllers and communication paths.

General rule: For random I/O, the general rule is to create arrays of the same type and size. For applications that do not spread I/Os equally across containers, create VGs that consist of one logical drive from every array, use a maximum inter-disk allocation policy for all LVs in the VG, and use a random disk order for each LV. Check that the ownership of logical drives selected are spread evenly across DS3500 controllers. Applications that spread their I/Os equally across containers, such as DB2®, use another layout.

If systems are using only one big volume, it is owned by one controller, and all the traffic goes through one path only, which happens because of the static load balancing that DS3500 controllers use.

3.5.2 Planning for systems without LVM: Windows example

Today, the Microsoft Windows operating system does not have a powerful LVM like certain UNIX systems. Distributing the traffic among controllers in such an environment might be a little bit harder. Actually, Windows systems have an integrated reduced version of Veritas Volume Manager (also known as Veritas Foundation Suite) called Logical Disk Manager (LDM), but it does not offer the same flexibility as regular LVM products. The integrated LDM version in Windows is used for the creation and use of *dynamic disks*.

With Windows 2003 and 2008, there are two types of disks: basic disks and dynamic disks. By default, when a Windows system is installed, the basic disk system is used.

Basic disks and basic volumes are the storage types most often used with Microsoft Windows operating systems. A basic disk refers to a disk that contains basic volumes, such as primary partitions and logical drives. A basic volume refers to a partition on a basic disk. For Windows 2003 and 2008, a primary partition on a basic disk can be extended using the **extend** command in the diskpart.exe utility.

Dynamic disks provide features that basic disks do not, such as the ability to create volumes that span multiple disks (spanned and striped volumes), and the ability to create software level fault tolerant volumes (mirrored and RAID 5 volumes). All volumes on dynamic disks are known as *dynamic volumes*.

With the DS3500 Storage System, you can use either basic or dynamic disks, depending on your needs and requirements (certain features might not be supported when using dynamic disks). There are cases for both disk types: the decision depends on your individual circumstances. In certain large installations, where you might have the requirement to span or stripe logical drives and controllers to balance the work load, dynamic disk might be your only choice. For smaller to mid-size installations, you might be able to simplify and just use basic disks.

When using the DS3500 Storage System, the use of software mirroring and software RAID 5 is not required. Instead, configure the storage on the DS3500 Storage System for the redundancy level required.

If you need greater performance and more balanced systems, you have two options:

- If you want to have the UNIX-like capabilities of LVM, you can purchase and use the Veritas Storage Foundation (from Symantec) suite or a similar product. With this product, you get several features that go beyond LDM. Volume Manager does not just replace the Microsoft Management Console (MMC) snap-in: it adds a much more sophisticated set of storage services to Windows 2003 and 2008. After Windows is upgraded with Volume Manager, you are able to manage better multidisk direct server-attached (DAS) storage, JBODs (just a bunch of disks), Storage Area Networks (SANs), and RAID.

The main benefit is the ability to define sub-disks and disk groups. You can divide a dynamic disk into one or more sub-disks. A sub-disk is a set of contiguous disk blocks that represent a specific portion of a dynamic disk, that is mapped to a specific region of a physical disk. A sub-disk is a portion of a dynamic disk's public region.

A sub-disk is the smallest unit of storage in Volume Manager. Therefore, sub-disks are the building blocks for Volume Manager arrays. A sub-disk can be compared to a physical partition. With disk groups, you can organize disks into logical collections.

Assign disks to disk groups for management purposes, such as to hold the data for a specific application or set of applications. A disk group can be compared to a volume group. By using these concepts, you can make a disk group with more LUNs that are spread among the controllers.

Using Veritas Volume Manager and tuning the databases and applications goes beyond the scope of this guide. Browse the application vendor sites and vendor documentation for more information.

For Veritas Volume Manager (VxVM), see:

http://www.symantec.com/enterprise/products/overview.jsp?pcid=1020&pvid=203_1

- You can use the DS3500 Storage System and Windows dynamic disks to spread the workload between multiple logical drives and controllers, which can be achieved with spanned, striped, mirrored, or RAID 5:
 - Spanned volumes combine areas of un-allocated space from multiple disks into one logical volume. The areas of un-allocated space can be various sizes. Spanned

volumes require two disks, and supports up to 32 disks. If one of the disks containing a spanned volume fails, the entire volume fails, and all data on the spanned volume becomes inaccessible.

- Striped volumes can be used to distribute I/O requests across multiple disks. Striped volumes are composed of stripes of data of equal size written across each disk in the volume. They are created from equally sized, un-allocated areas on two or more disks. The size of each stripe is 64 KB and cannot be changed. Striped volumes cannot be extended and do not offer fault tolerance. If one of the disks containing a striped volume fails, the entire volume fails, and all data on the striped volume becomes inaccessible.
- Mirrored and RAID 5 options are software implementations that have an additional impact on top of the existing underlying fault tolerance level configured on the DS3500 Storage System. They can be employed to spread the workload between multiple disks, but there are two lots of redundancy happening at two separate levels.

These possibilities must be tested in your environment to ensure that the solution chosen suits your needs and requirements.

3.5.3 Virtualization

With the growth and popularity of storage area networks, storage environments are getting more and more complex. Storage virtualization reduces the complexity and costs of managing storage environments and optimizes storage utilization in a heterogeneous environment.

The IBM Storage Area Network Volume Controller (SVC) and IBM TotalStorage® Productivity Center products addresses these needs.

3.5.4 IBM System Storage SAN Volume Controller overview

The IBM System Storage SAN Volume Controller is a scalable hardware and software solution to allow aggregation of storage from various disk subsystems. It provides storage virtualization and a consistent view of storage across a Storage Area Network (SAN).

The SAN Volume Controller provides in-band storage virtualization by creating a pool of managed disks from attached back-end disk storage systems. These managed disks are then mapped to a set of virtual disks for use by various host systems.

In conjunction with the DS3500 Storage System family, the SAN Volume Controller (SVC) can increase the storage copy services functionality and also the flexibility of SAN-based storage. The SVC is flexible in its use. It can manage all host storage requirements or just part of them. A DS3500 Storage System can still be used to allocate storage to hosts or use the SVC, which is dependent upon various needs and requirements.

The SVC also offers an alternative to FlashCopy, VolumeCopy, and Enhanced Remote Mirroring for disaster recovery, high availability, and maintenance. If the use of SVC with a DS3500 Storage System is planned, then these premium features will not be required.

The SVC can also reduce the requirement for additional partitions. The SVC only consumes one storage partition. If you plan to use the SVC for all of your hosts, then a storage partition upgrade might not be required.

SVC is licensed by the capacity that is being managed. This capacity also includes the capacity used by the copy services.

For detailed information about SVC implementation, see *Implementing the IBM System Storage SAN Volume Controller V4.3*, SG24-6423.

For more information about Storage Virtualization, see the IBM TotalStorage Virtualization website at:

<http://www.ibm.com/servers/storage/software/virtualization/index.html>

For more information about SAN Volume Controller, see its home page at:

<http://www.ibm.com/servers/storage/software/virtualization/svc/index.html>

3.6 Host support and multipathing

Our intent in this section is to list the most popular supported operating system platforms and topologies. For a complete and up-to-date list, see the DS3500 series interoperability matrix, available at:

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

In 3.6.4, “Multipathing” on page 80, we discuss the available multipathing drivers and their supported operating systems.

3.6.1 Supported server platforms

The following server platforms are supported:

- ▶ IBM System x
- ▶ IBM System p
- ▶ IBM BladeCenter
- ▶ HP (Intel)

3.6.2 Supported operating systems

At the time of publication, the following operating systems are supported:

- ▶ Microsoft Windows Server 2003 SP2
- ▶ Microsoft Windows Server 2008 SP2
- ▶ Red Hat Enterprise Linux 5.4
- ▶ Red Hat Enterprise Linux 5.1
- ▶ Novell SUSE SLES 10 SP3
- ▶ Novell SUSE SLES 11
- ▶ VMware ESX 4.0
- ▶ IBM AIX V6.1 TL5
- ▶ HP-UX 11iv2 (11.23) and HP-UX 11iv3 (11.31)

3.6.3 Clustering support

The following clustering services are supported:

- ▶ Microsoft Cluster Services
- ▶ HP Service Guard manager 11.18.00

3.6.4 Multipathing

IBM offers various multipath drivers that you can use with your DS3500 Storage System. Only one of these drivers is required. Each driver offers multipath support, I/O load balancing, and automatic path failover.

The multipath driver is a proxy for the real, physical-level HBA drivers. Each multipath driver hides from the application the fact that there are redundant connections by creating a virtual device. The application uses this virtual device, and the multipath driver will connect the application to the correct physical path.

When you create a logical drive, you assign one of the two active controllers to own the logical drive (called *preferred controller ownership*, as described in 3.3.6, “Logical drives and controller ownership” on page 63) and to control the I/O between the logical drive and the application host along the I/O path. The preferred controller normally receives the I/O requests from the logical drive. If a problem along the data path (such as a component failure) causes an I/O to fail, the multipath driver issues the I/O to the alternate controller.

A multipath device driver is not required when the host operating system has its own mechanism to handle multiple I/O paths.

Always refer to the System Storage Interoperation Center (SSIC) for the latest supported operating systems and the available multipathing drivers at the following URL:

<http://www.ibm.com/systems/support/storage/config/ssic/index.jsp>

3.6.5 Microsoft Windows MPIO

Microsoft Windows MPIO is a Driver Development Kit (DDK) from Microsoft for developing code that manages multipath devices. It contains a core set of binary drivers that are installed with the DS3500 Device Specific Module (DSM) to provide a transparent system architecture that relies on Microsoft Plug and Play to provide LUN multipath functionality and maintain compatibility with existing Microsoft Windows device driver stacks.

Note: The MPIO Driver is included in the Storage Manager software package for Windows and supports Microsoft Windows 2003 and 2008 on 32-bit and x64 systems. In Windows 2008, MPIO is already part of the operating system.

The MPIO driver performs the following tasks:

- ▶ Detects and claims the physical disk devices presented by the DS3500 Storage Systems based on vendor/product ID strings and manages the logical paths to the physical devices.
- ▶ Presents a single instance of each LUN to the rest of the Windows operating system.
- ▶ Provides an optional interface through WMI for use by user-mode applications.
- ▶ Relies on the vendor's (IBM) customized Device Specific Module (DSM) for the information about the behavior of storage subsystem devices in regard to:
 - I/O routing information
 - Conditions requiring a request to be retried, failed, failed over, or failed back (for example, vendor-specific errors)
 - Handling of miscellaneous functions, such as release/reservation commands
- ▶ Allows multiple Device Specific Modules (DSMs) for various disk storage subsystems to be installed in the same host server.

See 3.6.10, “Function of the Auto-Logical Drive Transfer feature” on page 83 for more discussion about multipathing and failover considerations.

For compatibility information, always see the current DS3500 interoperability matrix or the readme file for your HBA driver. See the following website:

<http://www.ibm.com/systems/support/storage/config/ssic/index.jsp>

3.6.6 AIX MPIO

With Multiple Path I/O (MPIO), a device can be individually detected through one or more physical connections, or paths. A path-control module (PCM) provides the path management functions.

An MPIO-capable device driver can control more than one type of target device. A PCM can support one or more specific devices. Therefore, one device driver can be interfaced to multiple PCMs that control the I/O across the paths to each of the target devices.

The AIX PCM has a health-check capability that can be used to do the following:

- ▶ Check the paths and determine which paths are currently usable for sending I/O.
- ▶ Enable a path that was previously marked failed because of a temporary path fault (for example, when a cable to a device was removed and then reconnected).
- ▶ Check currently unused paths that might be used if a failover occurred (for example, when the algorithm attribute value is failover, the health check can test the alternate paths).

MPIO is part of the AIX operating system and does not need to be installed separately.

3.6.7 AIX Subsystem Device Driver Path Control Module

The Subsystem Device Driver Path Control Module (SDDPCM) is a loadable path control module for supported storage devices to supply path management functions and error recovery algorithms. When the supported storage devices are configured as Multipath I/O (MPIO) devices, SDDPCM is loaded as part of the AIX MPIO FCP (Fibre Channel Protocol) device driver during the configuration. The AIX MPIO-capable device driver with the supported storage devices SDDPCM module enhances the data availability and I/O load balancing.

Note: Only certain versions of the SDDPCM are supported through request for price quotation (RPQ)

3.6.8 HP-UX IBM Subsystem Device Driver

IBM now offers a Subsystem Device Driver (SDD) for HP-UX. The SDD is only supported through RPQ and provides the following functions:

- ▶ Enhanced data availability
- ▶ Dynamic input/output (I/O) load balancing across multiple paths
- ▶ Automatic path failover protection
- ▶ Concurrent download and installation of licensed machine code

See Figure 3-25 on page 82 for information about the flow of an I/O operation.




Figure 3-25 SDD on HP-UX

SDD in combination with a DS3500 Storage System only HP-UX 11iv2 (11.23) is supported. LVM is supported on HP-UX 11iv2 (11.23) and 11iv3 (11.31) as a multipathing driver.

More information about the architecture and installation of SDD can be found in the latest *Multipath Subsystem Device Driver User's Guide* (S7000303), or at the following URL:

<http://www.ibm.com/servers/storage/support/software/sdd>

3.6.9 Linux: RHEL/SLES

In this section we discuss the Linux multipath options.

Redundant Disk Array Controller

The Redundant Disk Array Controller (RDAC), also known as Multi-Path Proxy (MPP), is the preferred multipathing driver for Linux based operating systems such as Red Hat Enterprise Linux (RHEL) or SUSE Linux Enterprise Server (SLES).

The current RDAC driver implementation performs the following tasks:

- ▶ Detects and claims the physical devices (LUNs) presented from the DS3500 Storage Systems (*hides* them) based on vendor/product ID strings and manages all of the paths to the physical devices.
- ▶ Presents a single instance of each LUN to the rest of the Linux operating system components.
- ▶ Manages all of the Plug and Play interactions.
- ▶ Provides I/O routing information.
- ▶ Identifies conditions requiring a request to be retried, failed, or failed over.
- ▶ Automatically fails over the LUNs to their alternate controller when detecting problems in sending I/Os to the LUNs in their preferred controller, and fails back the LUNs to their preferred controller when detecting the problems in the preferred path have been fixed.
- ▶ Handles miscellaneous functions such as persistent reservation translation.
- ▶ Uses a round robin (load distribution or load balancing) model.

RDAC implementation

RDAC is implemented between the HBA driver and the operating system disk driver, operating as a low-level filter driver. It has the following advantages:

- ▶ It is much more transparent to the OS and applications.
- ▶ I/O controls at the HBA driver level are not as tightly coupled to the OS as those at the disk driver level. Consequently, it is easier to implement I/O control functionality in the MPP-based RDAC driver for routing functions.

Because the driver is positioned at the HBA level (see Figure 3-26 on page 83), it has access to the SCSI command and sense data interface of the HBA driver, and therefore can make more informed decisions about what to do in the case of path failures.




Figure 3-26 Linux RDAC/MPP driver

Note: In Linux, RDAC cannot be installed with the Installation Wizard. If you need RDAC, you have to download and install it separately.

DMM for Linux

At the time of writing DMM for Linux is only supported with the SLES 11 release of Linux. There is currently no Redhat version of this driver supported. With the SLES 11 the following requirements must be adhered to for the driver to function:

- ▶ OS version: SLES11.0 with the following patch installed
- ▶ Patch installed: lsi-scsi_dh_rdac-kmp-default-01.00_2.6.27.19_5-1
- ▶ Kernel upgraded to 2.6.31.5-0.16-default
- ▶ DMM version: device-mapper-1.02.27-8.6

For detailed instructions on how to implement and use DMM with SLES 11, see Chapter 20, “SuSE Linux configuration guide” on page 587.

3.6.10 Function of the Auto-Logical Drive Transfer feature

In a DS3500 Storage System equipped with two controllers, you can provide redundant I/O paths with the host systems. There are two separate components that provide this redundancy: a multipathing driver and Auto-Logical Drive Transfer (ADT).

ADT for logical drive-level failover

ADT is a built-in feature of controller firmware that allows logical drive-level failover rather than controller-level failover (as is the case with RDAC).

Note: ADT is not a failover driver. ADT provides storage systems with the flexibility to work with certain third-party failover software.

- ▶ ADT-disabled failover:

The multi-path software will send a SCSI Mode Select command to cause a change in volume ownership before using the alternate path. All logical drives on the preferred controller are transferred to the alternate controller, which is the configuration setting for Microsoft Windows, IBM AIX, and Linux (when using the RDAC or SDD driver and non-failover Fibre Channel HBA driver) systems.

When ADT is disabled, the I/O data path is still protected as long as you use a multi-path driver. After the I/O data path problem is corrected, the preferred controller does not automatically reestablish ownership of the logical drive. You must open a storage management window, select **Redistribute Logical Drives** from the Advanced menu, and perform the Redistribute Logical Drives task.

Figure 3-27 shows the ADT-disabled failover mode phases.

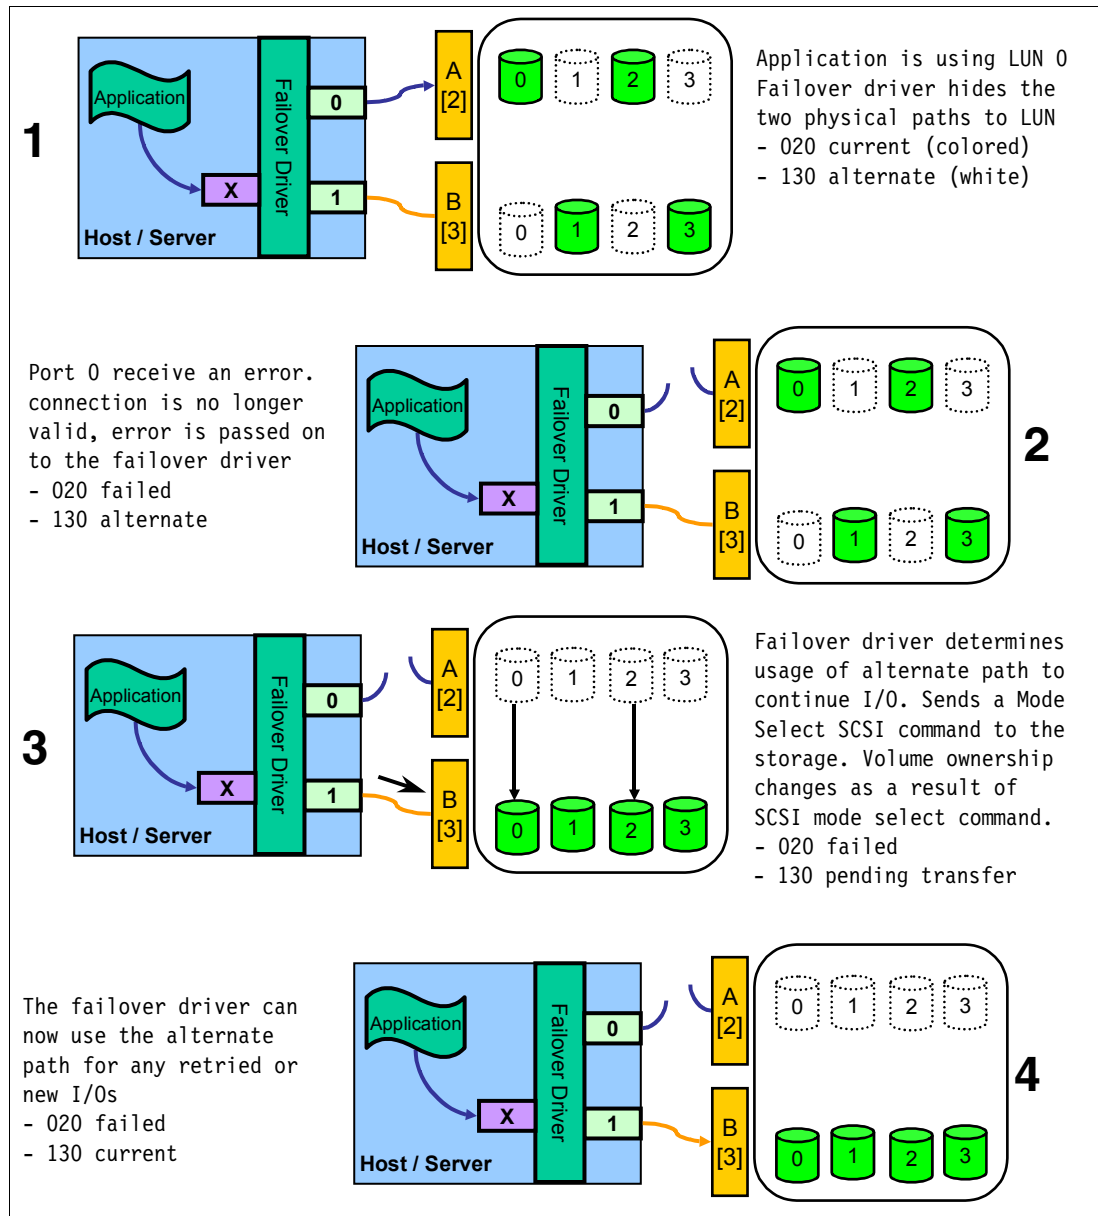


Figure 3-27 ADT-disabled mode path failover

Note: In ADT-disabled mode, you are required to issue a redistribution command manually to get the LUNs balanced across the controllers.

► ADT-enabled failover:

The multi-path driver starts using the alternate path by sending the I/O down the path it chooses and lets the ADT react, which is the normal configuration setting for Novell NetWare, Linux (when using FC HBA failover driver instead of RDAC), and Hewlett Packard HP-UX systems. After the I/O data path problem is corrected, the preferred controller automatically reestablishes ownership of the logical drive as soon as the multipath driver detects that the path is normal again.

Phases of failover

Figure 3-28 shows the phases of failover in an ADT-enabled case.

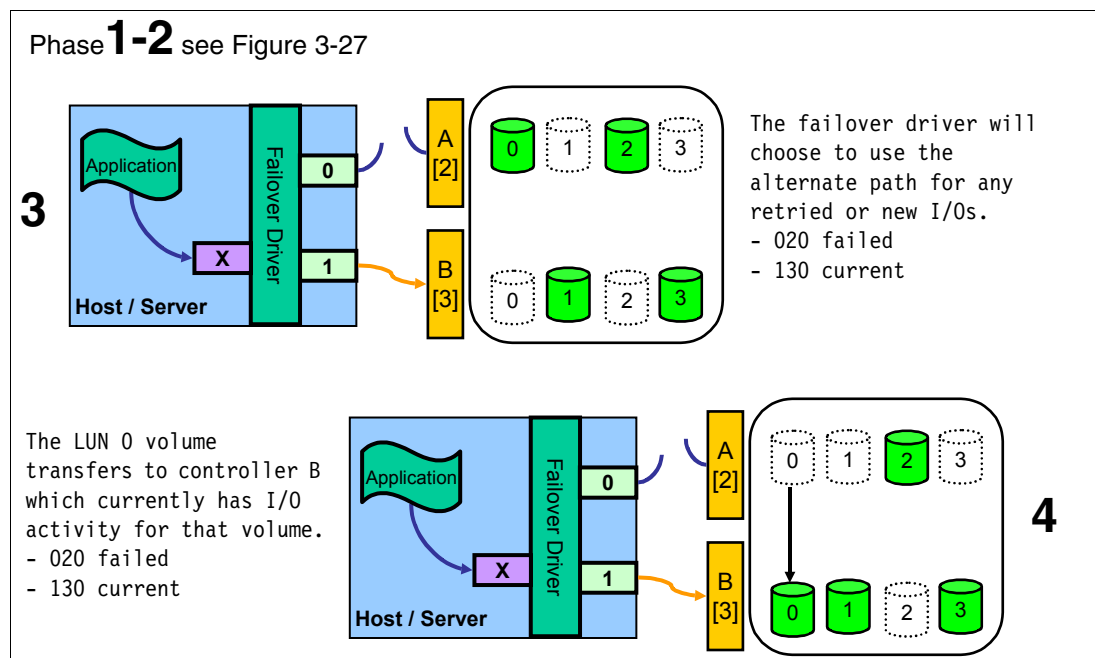


Figure 3-28 ADT-enabled mode path failover

Note: In ADT mode, RDAC automatically redistributes the LUNs to their preferred path after the failed path is again operational.

3.7 Operating system restrictions

In your planning, you must also consider the restrictions of your operating system. We have covered the maximum file system size ordered by the operating system and their commonly used file systems, the maximum number of LUNs that can be assigned to one host, and the maximum logical volume size supported by various operating systems.

3.7.1 Maximum capacity for a logical drive

Table 3-6 shows the maximum capacity a logical drive can have and be formatted with a specific file system.

Table 3-6 Maximum file system size supported by various operating systems

Operating System	File system	Size
Linux	ext2	32 TB
	ext3	32 TB
AIX	JFS2	64 PB
Windows	NTFS	16 EB
	FAT32	8 TB
HP-UX 11i v2	HFS	128 GB
	JFS 3.5	2 TB / 32 TB ^{a1}
ESX Server 3.5U5, 4.0U1 and 4.1	VMFS	2TB / 64TB per VMFS

a. VxVM 3.5 only and OnlineJFS license for greater than 2 TB file systems (only HP-UX 11iV2 B. 11.23)

3.7.2 Maximum number of LUNs per host

The maximum number of LUNs that can be assigned to a host is dependent on the host operating system type and mapping limit within a storage partition on the DS3500 Storage System. At the time of writing, the DS3500 Storage System supports 32 LUNs per single storage partition, and 256 LUNs support will soon follow. You can find more information about storage partitioning in 3.3.7, "Storage partitioning" on page 63. Table 3-7 lists the operating system limit and the limit of the DS3500 Storage System.

Table 3-7 Maximum number of LUNs per host

Host operating system type	Maximum Number of LUNs per host (theoretical OS limit)	Maximum Number of logical drives or LUNs per host (DS3500 Storage System limit)
AIX	64000	255
HP-UX	16000	255
Linux	256	255
Windows	256	255
ESX Server 3.5U5, 4.0U1 and 4.1	256	255

Note: The total number of logical drives on a DS3500 Storage System cannot exceed 2048.



IBM System Storage DS3500 and EXP3500 Cabling

In this chapter we describe the guidelines for the cabling of the DS3500 storage subsystem and the EXP3500 expansion enclosure. We introduce the various connector ports located on the rear of the DS3500 controller and the EXP3500 enclosure, and then describe the guidelines for:

- ▶ SAS expansion connections
- ▶ Host-side connections (SAS, iSCSI, and Fibre Channel)
- ▶ Ethernet management connections
- ▶ Power cabling

We also describe cable management and the correct power-on/off procedures for the IBM System Storage DS3500.

For more information about all the topics discussed in this chapter, you can refer to the *Installation, User's, and Maintenance Guide - IBM System Storage DS3500 and EXP3500* available at the following URL:

<http://www.ibm.com/systems/support/supportsite.wss/docdisplay?ln docid=MIGR-5084732&brandind=5000028>

4.1 DS3500 controller connectors

In this section we describe the various connectors that are located on the rear panel of each of the four DS3500 controller models and the EXP3500 enclosure.

4.1.1 DS3500 controller with standard port configuration

Figure 4-1 shows the rear panel of a DS3500 controller with the standard port layout.



Figure 4-1 The rear panel of a DS3500 controller

From the left, the ports are:

- ▶ A serial port: This port is reserved for use by IBM Service personnel only.
- ▶ Two Ethernet management ports labelled 1 and 2 for Out-Of-Band (OOB) management.
- ▶ SAS Host Ports labelled 1 and 2: These are two SAS 2.0 ports for host-side connections. The DS3500 host ports are x4 multilane, 6 Gbps universal mini-SAS ports.
- ▶ The blank space on the rear panel is where the additional host ports on one of the optional host adapter cards would be located. You can install one of three types of optional host adapter cards if you need additional SAS host port support or concurrent host connectivity with either SAS and FC, or SAS and iSCSI.
- ▶ The port on the extreme right is a x4 multilane mini-SAS port for connecting to EXP3500 drive expansion enclosures.

Below SAS ports 1 and 2 there is a 7-segment numeric LED display that provides information about enclosure identification and diagnostics. The functions of this and the other status LEDs located on the rear panel are described in more detail in the Chapter 14, “Administration - Support” on page 329.

4.1.2 DS3500 controller with optional SAS host port adapter

Figure 4-2 shows the rear panel of a DS3500 controller with the optional SAS host port adapter installed. In addition to the standard ports described in 4.1.1, “DS3500 controller with standard port configuration” on page 88, the optional SAS host port adapter adds an additional two SAS 2.0 host ports for a total of four per controller.




Figure 4-2 The rear panel of a DS3500 controller with the optional SAS host port adapter installed

The SAS Host Ports labelled 1, 2, 3, and 4 are SAS 2.0 ports for host-side connections. SAS ports 3 and 4 are located on the optional SAS host port adapter. Each DS3500 host port is a x4 multilane, 6 Gbps universal mini-SAS port.

4.1.3 DS3500 controller with optional Fibre Channel host port adapter

Figure 4-3 shows the rear panel of a DS3500 controller with the optional four port Fibre Channel host adapter installed.



Figure 4-3 The rear panel of a DS3500 controller with the optional Fibre Channel host port adapter installed

SAS host ports are labelled 1 and 2, and the Fibre Channel host ports are labelled FC 3, 4, 5, and 6. These four ports are located on the optional Fibre Channel host adapter. Each of these ports supports a small-form-factor pluggable (SFP) transceiver and is capable of operating at 8 Gbps, 4 Gbps, or 2 Gbps.

4.1.4 DS3500 controller with optional iSCSI host port adapter

Figure 4-4 shows the rear panel of a DS3500 controller with the optional four port iSCSI host adapter installed.



Figure 4-4 The rear panel of a DS3500 controller with the optional iSCSI host port adapter installed

SAS host ports are labelled 1 and 2, and iSCSI host ports labelled Ch 3, 4, 5, and 6. These four ports are located on the optional iSCSI host adapter. The iSCSI host ports are auto-negotiated 100/1000 Mbps Ethernet ports that support iSCSI off load. The host ports also provide support for both Internet Protocol version 4 (IPv4) and IPv6, and have the advanced networking functions of jumbo frames, IEEE 802.1p, and virtual local area network (VLAN) support.

4.1.5 EXP3500 ports

The EXP3500 expansion enclosure comes with one ESM, which enables you to connect the EXP3500 expansion enclosure to a DS3500 storage subsystem.

Note: If you are attaching the EXP3500 to a dual controller DS3500 storage subsystem, you must install a second ESM in the EXP3500 expansion enclosure to provide a redundant drive path to the DS3500 storage subsystem.

Figure 4-5 shows the port layout on the rear of an EXP3500 enclosure with two ESMs installed.



Figure 4-5 The rear panel of an EXP3500 enclosure with dual ESMs installed.

From the left, these ports are available on every ESM:

- ▶ Two SAS expansion ports labelled 1 and 2. These two ports are x4 multilane, 6 Gbps universal mini-SAS ports. They are both IN ports but only one multiprotocol can be used at any time. Trunking is not supported on these ports.
- ▶ A Serial port reserved for use by IBM Service personnel only.
- ▶ An Ethernet port reserved for use by IBM Service personnel only.
- ▶ A SAS expansion port. This port is for connecting to other EXP3500 enclosures only and is a x4 multilane, 6 Gbps universal mini-SAS port.

Below SAS ports 1 and 2 there is a 7-segment numeric LED display that provides information about enclosure identification and diagnostics. The functions of this and the other status LEDs located on the rear panel are described in more detail in the Chapter 14, “Administration - Support” on page 329.

4.2 Enclosure ID settings

The enclosure ID is a unique two-digit identifier for DS3500 and EXP3500 enclosure in the storage subsystem configuration. In a dual controller DS3400 subsystem, both controllers should have the same enclosure ID. Similarly, on an EXP3000 with two ESMs, both ESMs must have the same enclosure ID under normal operating conditions.

The default enclosure ID of the DS3500 controller is 85, and generally you should leave this value unchanged. The enclosure ID of each EXP3500 is normally set to a value of 00 at the factory. The DS3500 controller firmware will automatically assign a unique enclosure ID to each EXP3500 enclosure that you connect to it. The automatically assigned EXP3500 enclosure IDs will begin with enclosure ID 0, then 1, 2 and so on. The range of supported enclosure ID values is 0 through 99.

You can view and change the enclosure IDs through the DS Storage Manager software. See “Enclosure ID” on page 202 for information about how to change the enclosure ID.

In normal operation, the enclosure ID is also displayed on the seven-segment numeric LED display located on the back of each controller and ESM. Figure 4-6 on page 91 shows the 7 segment numeric LED displaying the enclosure ID.




Figure 4-6 The 7 segment numeric LED displaying the enclosure ID

Note: The 7 segment numeric LED display is also used for diagnostic purposes.

These enclosure ID values are important because the unique hardware address of each drive in the storage subsystem is determined by the unique enclosure ID and the slot number of the drive in that enclosure. If two enclosures are configured with the same enclosure ID, the disk drives will have conflicting hardware addresses.

4.3 SAS cables

In this section we describe the SAS cables that you can use to connect the DS3500 to hosts and to EXP3500 expansion enclosures.

Each DS3500 RAID controller has up to four x4 multilane SAS host ports and a single x4 multilane SAS port for drive channel connections. IBM provides SAS cables in both 1 metre and 3 metre cable lengths for connecting to EXP3500 enclosures and for SAS connections to hosts. These SAS cables are also used to connect EXP3500 enclosures together. Figure 4-7 shows a typical SAS cable with mini-SAS 4x multilane SAS connectors on both ends of the cable.




Figure 4-7 SAS cable

Follow these precautions to avoid damage to the SAS cables:

- ▶ 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 as shown in Figure 4-8. Make sure that it locks into place.




Figure 4-8 Connecting a mini-SAS cable

Follow these steps to remove a mini-SAS cable as shown in Figure 4-9:

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.




Figure 4-9 Removing a mini-SAS cable

4.4 Fibre Channel cabling

In this section we describe the features of Fibre Channel cabling that you use to connect the DS3500 controllers to hosts with Fibre-Channel host adapters or to Fibre-Channel SAN switches. We also describe general guidelines for cable management of the fibre optical cables.

The DS3500 supports the use of Small Form-factor Pluggable (SFP) transceiver modules and Fibre Channel cables with LC connectors at speeds of 2 Gbps, 4 Gbps, and 8 Gbps.

Figure 4-10 shows a Fibre Channel cable with LC connector attached to an SFP module.



Figure 4-10 Small Form Pluggable (SFP) module with LC connector fiber cable

4.4.1 SFP transceiver modules

Each DS3500 controller can be installed with a optional Fibre-Channel host adapter, and can support up four Fibre Channel connections as shown in Figure 4-3 on page 89. You must install a Small- Form-factor Pluggable (SFP) module into each fibre-channel port that you want to use for host connections. The SFP module is a laser product and its function is to convert electrical signals to and from optical signals for transmission to and from the controllers. You then connect the fibre optic cable to the SFP module.

SFP modules are available in both long-wave and shortwave types, but only shortwave SFP modules are supported in the DS3500 storage subsystem. You can use the DS Storage Manager Client to view the Storage Subsystem Profile to confirm that only shortwave SFP modules are installed in the DS3500 storage subsystem. SFP modules are also available in 2-Gbps, 4-Gbps, and 8-Gbps models. We recommend that you only use 8-Gbps SFP modules for optimum performance.

Attention: You must not use long-wave SFP modules or Gigabit Interface Converters (GBICs) in any of the Fibre Channel ports in the DS3500 storage subsystem controllers. Long-wave SFP modules and GBICs are not supported for use in the DS3500 storage subsystem.

General guidelines for using SFP modules

You should follow these general guidelines when working with SFP modules:

- ▶ The SFP module housing has an integral guide key that is designed to prevent you from inserting the SFP module incorrectly.
- ▶ Use minimal pressure when inserting an SFP module into a Fibre Channel port. Forcing the SFP module into a port could cause damage to the SFP module or to the port.

- ▶ You can insert or remove the SFP module while the port is powered on.
- ▶ The operational or redundant loop performance is not affected when you install or remove an SFP module.
- ▶ You must insert the SFP module into a port before you connect the fiber optic cable.
- ▶ You must remove the fiber optic cable from the SFP module before you remove the SFP module from the port.
- ▶ The speed of the SFP module determines the maximum operating speed of the Fibre Channel port in which the SFP module is installed. For example, a 2-Gbps SFP module that is connected to a 4-Gbps-capable port will limit the speed of that port to a maximum of 2 Gbps.

Attention: Because there are no physical features that distinguish an 8 Gbps, 4 Gbps, or 2 Gbps SFP module from each other, you should always carefully check the IBM part number, option number, and FRU part number of the SFP module so that you can correctly identify its speed.

Installing SFP modules

The storage subsystem requires SFP modules. SFP modules convert electrical signals to the optical signals that are required for Fibre Channel transmission to and from RAID controllers. After you install the SFP modules, you use fiber optic cables to connect the storage subsystem to other Fibre Channel devices.

To install an SFP module, complete the following steps:

1. Remove the SFP module from its static-protective package.
2. Remove the protective cap from the SFP module, as shown in Figure 4-11. Save the protective cap in a dust-free place for future use.




Figure 4-11 SFP module and protective cap

3. Remove the protective cap from the SFP port. Save the protective cap in a dust-free place for future use.
4. Insert the SFP module into the host port until it clicks into place as shown in Figure 4-12 on page 95.




Figure 4-12 Installing an SFP module into the host port

5. Connect an LC-LC Fibre Channel cable.

Removing SFP modules

To remove the SFP module from the host port, complete the following steps:

1. Remove the LC-LC Fibre Channel cable from the SFP module.
2. Unlock the SFP module latch:
 - For SFP modules that contain plastic tabs, unlock the SFP module latch by pulling the plastic tab outward 10°, as shown in Figure 4-13.




Figure 4-13 Unlocking the SFP module latch - plastic variety

- For SFP modules that contain wire tabs, unlock the SFP module latch by pulling the wire latch outward 90°, as shown in Figure 4-14 on page 96.




Figure 4-14 Unlocking the SFP module latch - wire variety

3. With the SFP module latch in the unlocked position, remove the SFP module.
 - For SFP modules that contain plastic tabs, slide the SFP module out of the port.
 - For SFP modules that contain wire tabs, grasp the wire latch and pull the SFP module out of the port.
4. Replace the protective cap on the SFP module.
5. Place the SFP module into a static-protective package.
6. Replace the protective cap on the host port.

4.4.2 Fibre Channel cables

In this section we describe the fibre optical cables that you can use for connecting to the small form-factor pluggable (SFP) modules in the DS3500 storage subsystem. We also describe general rules for cable management of fibre optical cables.

Because the DS3500 only supports SFP modules, Fibre Channel cables with LC connectors are required. Connectors that plug into SFF or SFP devices are called LC connectors. The two fibres in the cable (transmit and receive) each have their own part of the connector. The connector is keyed to ensure correct polarization when connected, that is, transmit to receive and vice-versa. Figure 4-15 shows an LC connector at one end of a Fibre Channel cable.




Figure 4-15 An LC connector

Most Fibre Channel cables have an LC connector at both ends, and we call these LC-LC cables. You will use LC-LC cables when you are connecting the DS3500 host ports to a port on a SAN switch or a Host Bus Adapter that supports SFP modules.

General guidelines when using Fibre Channel cables

You should always follow these general guidelines when using Fibre Channel cables:

- ▶ Do not route the cable along a folding cable-management arm.
- ▶ For devices on slide rails, leave enough slack in the cables so they do not bend to a diameter of less than 76 mm (3 in.), or a radius less than 38 mm (1.5 in.) when extended or become pinched when retracted.
- ▶ Route the cable away from places where it can be damaged by other devices in the rack cabinet.
- ▶ Do not use plastic cable ties in place of the provided cable straps.
- ▶ Do not overtighten the cable straps or bend the cables to a diameter of less than 76 mm (3 in.), or a radius less than 38 mm (1.5 in.).
- ▶ Do not put excess weight on the cable at the connection point. Be sure that the cable is well supported.
- ▶ The following are the recommended maximum cable lengths.
 - 2 Gbps: 300 m (984 ft.) of 50/125 um fiber, 150 m (492 ft.) of 62.5/125 um fiber
 - 4 Gbps: 150 m (492 ft.) of 50/125 um fiber, 70 m (230 ft.) of 62.5/125 um fiber
 - 8 Gbps: 50 m (164 ft.) of 50/125 um fiber, 21 m (115 ft.) of 62.5/125 um fiber

Note: Refer to Table 3-1 on page 39 for the more information about the supported distances for the OM1 and OM2 category Fibre Channel cables.

4.4.3 Interoperability of 2 Gbps, 4 Gbps, and 8 Gbps devices

The Fibre Channel standard specifies a procedure for speedy auto-detection. For example, if a 4 Gbps port on a switch or device is connected to a 2 Gbps port, it must negotiate down and the link will run at 2 Gbps. If there are two 8 Gbps ports on either end of a link, the negotiation runs the link at 8 Gbps if the link is up to specifications. A link that is too long or “dirty” can end up running at 4 Gbps, even with 8 Gbps ports at either end, so care must be taken with cable lengths and that connector quality is sound.

The same rules apply to 8 Gbps devices relative to 4 Gbps and 2 Gbps environments. The 8 Gbps and 4 Gbps devices have the ability to automatically negotiate back down to either 4 Gbps, 2 Gbps, or 1 Gbps depending on the attached device and the link quality. If the link does unexpectedly negotiate to a slower speed than expected, then the causes or reasons for this ought to be investigated and remedied.

The DS3500 Storage System has 8 Gbps functionality: there are several SAN switches, directors, and Host Bus Adapters that can operate at this speed.

Note: On certain Fibre Channel switch vendor models, it might be necessary to manually configure the port to a specific speed of 2, 4, or 8 Gbps to obtain the required speed instead of leaving “auto-detection” on the port.

4.5 iSCSI Ethernet cables

In this section we describe the iSCSI Ethernet cabling that you use for iSCSI connections to hosts.

You must use either a Category 5e or a Category 6 Ethernet cable when you are connecting a DS3500 controller iSCSI host port to a gigabit Ethernet switch, and then connect a similar Ethernet cable from the switch to the host iSCSI initiator port.

4.6 EXP3500 attachment

In this section we describe the requirements and general guidelines for connecting EXP3500 enclosures to a DS3500 storage subsystem.

The DS3500 supports attachment of up to 96 drives. The DS3512 and EXP3512 enclosures support 12 drives, and the DS3524 and EXP3524 enclosures support 24 drives. The DS3500 storage subsystem supports intermixing of EXP3512 and EXP3524 enclosures so long as the total number of disk drives in the storage configuration does not exceed 96.

Table 4-1 shows the various combinations of DS3500 storage subsystem and EXP3500 enclosures that can be configured so that you do not exceed the limit of 96 drives.

Table 4-1 DS3500 storage subsystem and EXP3500 enclosures limits

DS3512		DS3524	
Enclosure Type	Enclosure quantity	Enclosure type	Enclosure quantity
EXP3512	7	EXP3512	6
EXP3524	0	EXP3524	0
or		or	
EXP3512	5	EXP3512	4
EXP3524	1	EXP3524	1
or		or	
EXP3512	3	EXP3512	2
EXP3524	2	EXP3524	2
or		or	
EXP3512	1	EXP3512	0
EXP3524	3	EXP3524	3

4.6.1 Redundant drive channels

Each DS3500 controller has a drive expansion channel with one x4 SAS port. The EXP3500 expansion enclosures that are connected to this port from one drive channel. Each drive channel supports a maximum of 96 disk drives. In a dual-controller DS3500, one drive channel from each controller combine to form a redundant drive-channel pair. If any

component in the drive channel fails, the DS3500 controllers can still access the EXP3500 expansion enclosures in the redundant drive-channel pair.

Attention: The EXP3500 ships with one ESM installed. If you are going to attach the EXP3500 to a dual controller DS3500 storage subsystem then you must install a second ESM in the EXP3500 expansion enclosure to provide redundant drive paths to the EXP3500 enclosures. Connecting an EXP3500 with one ESM installed to a DS3500 with two controllers installed is not a supported configuration.

4.6.2 Drive channel cabling rules

You should always follow these two basic connection rules when attaching EXP3500 enclosures to a DS3500 storage subsystem:

- ▶ Connect the drive expansion SAS port on the DS3500 controller to the SAS IN port on the EXP3500.
- ▶ Connect the SAS OUT port on the EXP3500 to the SAS IN port on the next EXP3500.

4.6.3 Single controller DS3500 with one or more EXP3500 enclosures

Figure 4-16 on page 100 shows the SAS connections between a single controller DS3500 subsystem and three EXP3500 enclosures, with a single ESM installed in each EXP3500. You should use the two connection rules above to connect the SAS ports on the DS3500 and the EXP3500s as shown.




Figure 4-16 EXP3500 attachment - single controller and single ESM configuration

Note: The configuration shown in Figure 4-16 is a fully functional configuration but there is no redundancy on the drive side connections. If there is a fault with any one of the EXP3500 enclosures or the ESMs or a SAS cable, then there will be a loss of connection to the drives in all the EXP3500s that are connected below the point of failure. This situation can be avoided by using a configuration with two controllers in the DS3500 and two ESMs installed in each EXP3500 that is connected to the DS3500.

4.6.4 Dual controller DS3500 with one EXP3500 enclosure

In this section we describe how to correctly cable a dual controller DS3500 to a single EXP3500 enclosure with two ESMs installed. This configuration provides redundancy on the drive-side connections. Figure 4-17 on page 101 shows the SAS cable connections in this configuration.




Figure 4-17 DS3500 storage subsystem (top) with one EXP3500 attached (bottom)

Follow this sequence of steps to connect a single EXP3500 with two ESMs installed to a DS3500 with two controller installed:

1. Connect the EXP3500 expansion enclosure to RAID controller A:
 - a. Connect one end of a SAS cable to the SAS OUT port of DS3500 controller A.
 - b. Connect the other end of the SAS cable to one of the IN SAS connectors on the left ESM in the EXP3500 expansion enclosure.
2. Connect the EXP3500 expansion enclosure to RAID controller B:
 - a. Connect one end of a SAS cable to the SAS OUT port of DS3500 controller B.
 - b. Connect the other end of the SAS cable to one of the SAS IN ports on the right ESM in the EXP3500 expansion enclosure.

4.6.5 Dual Controller DS3500 with two or more EXP3500 enclosures

In this section we describe how to cable a DS3500 to two or more EXP3500 enclosures using what is known as the top-down/bottom-up cable layout. The top-down/bottom-up cable layout is preferable in configurations with two or more EXP3500 enclosures because it maximizes redundancy of the drive-side SAS connections.

Use top-to-bottom order on the left-side controller and EXP3500 ESMs. The right-side controller and EXP3500 ESMs should be connected in bottom-to-top order. This will eliminate a single point of failure on the drive side, such as a SAS cable or ESM fault. If any EXP3500 fails completely, the remaining EXP3500s will still be accessible through the alternative drive path.

Figure 4-18 on page 102 shows the SAS cable connections when using the top-down/bottom up cabling layout with the DS3500 storage subsystem at the top. This is the preferred SAS expansion cabling for when you are connecting two or more EXP3500 enclosures, up to the supported maximum of eight EXP3500 enclosures.




Figure 4-18 The top-down/bottom-up SAS drive-side connections

You should follow this sequence of steps to connect DS3500 controllers A and B to multiple EXP3500 expansion enclosures that each have two ESMs installed:

1. Connect the first EXP3500 expansion enclosure to DS3500 controller A:
 - a. Connect one end of a SAS cable to the drive expansion OUT port on the DS3500 controller A.
 - b. Connect the other end of the SAS cable to one of the IN SAS connectors on the left ESM in the EXP3500 expansion enclosure.
2. Connect the left ESM of the EXP3500 expansion enclosure to the next EXP3500 expansion enclosure in the chain:
 - a. Connect one end of a SAS cable to the SAS OUT connector on the left ESM of the EXP3500 expansion enclosure that you just connected.
 - b. Connect the other end of the SAS cable to one of the IN SAS connectors on the left ESM in the next EXP3500 expansion enclosure in the chain.
 - c. Repeat steps 2a and 2b for each EXP3500 expansion enclosure that you add.
3. Connect the last EXP3500 expansion enclosure in the chain to DS3500 controller B:
 - a. Connect one end of a SAS cable to the drive expansion port on the DS3500 controller B.
 - b. Connect the other end of the SAS cable to one of the IN SAS connectors on the right ESM in the last EXP3500 expansion enclosure of the chain that you created.
4. In reverse order (last in chain to first in chain), connect the right ESMs in the EXP3500 expansion enclosures in the chain:
 - a. Connect one end of a SAS cable to the OUT SAS connector on the right ESM of the EXP3500 expansion enclosure that you just connected.
 - b. Connect the other end of the SAS cable to one of the IN SAS connectors on the right ESM of the previous EXP3500 expansion enclosure in the chain.
 - c. Repeat steps 4a and 4b until you connect the right ESM of the first EXP3500 expansion enclosure in step 2 to the right ESM of the EXP3500 expansion enclosure in step 1.

4.6.6 Adding an EXP3500 enclosure to a running dual-controller configuration

In this section we describe the steps required to attach an EXP3500 enclosure to a running DS3500 configuration. There are two possible scenarios when attaching an EXP3500 to a running DS3500 configuration:

- There are no EXP3500 enclosures already attached.
- There are EXP3500 enclosures already attached.

Attention: You must ensure that each EXP3500 expansion enclosure has two ESMs installed before you connect it to a dual-controller DS3500 storage subsystem.

Adding one EXP3500 to a running DS3500 configuration with no EXP3500 enclosures attached

In this scenario, there are no EXP3500 enclosures already attached to the DS3500. We use this sequence of steps to attach the EXP3500 to the DS3500:

1. You must ensure that the DS3500 storage subsystem status is Optimal in the DS Storage Manager software.

2. Connect the power cables to the new expansion enclosure and power-on the expansion enclosure.
3. Connect one end of a Serial Attached SCSI (SAS) cable to the Drive expansion port on Controller A in the DS3500 storage subsystem.
4. Connect the other end of the SAS cable to either of the SAS IN ports on the left ESM in the EXP3500 expansion enclosure.
5. Connect one end of a second SAS cable to the Drive expansion port on Controller B in the DS3500 storage subsystem.
6. Connect the other end of the second SAS cable to the IN port on the right ESM in the expansion enclosure.
7. Wait until the storage subsystem status is Optimal in the DS Storage Manager software before proceeding to step 8.
8. Add the hard disk drives to the expansion enclosure one at a time.

Attention: Make sure that you wait at least 30 seconds for the DS Storage Manager software to recognize each newly added drive before you install the next drive.

Adding an EXP3500 to a running DS3500 configuration with EXP3500 attached

In this scenario, there are already EXP3500 enclosures attached to a running DS3500 configuration. We use the following steps to attach the EXP3500 to the DS3500:

1. Make sure that the storage subsystem status is Optimal in the DS Storage Manager software.
2. Connect the power cables to the new expansion enclosure and power on the expansion enclosure.
3. Connect one end of a Serial Attached SCSI (SAS) cable to the IN port on the left ESM in the expansion enclosure that is being added to the configuration.
4. Connect the other end of the SAS cable to the OUT port on the left ESM in the last expansion enclosure in the configuration (this expansion enclosure is directly connected to the Drive expansion port on controller A of the DS3500).
5. On the last expansion enclosure in the configuration, disconnect the SAS cable from the IN port on the right ESM and connect it to the IN port on the right ESM on the new expansion enclosure that is being added to the configuration.

Note: When you disconnect the SAS cable in step 5 there will be a loss of drive path redundancy and the DS Storage Manager Recovery Guru will report an error message. You can safely disregard this error message because you will restore a redundant drive path in step 6 when you connect the SAS cable again.

6. Connect a SAS cable to the OUT port on the right ESM of the new expansion enclosure and connect the other end to either of the IN port on the right ESM in the expansion enclosure that was previously connected to controller B of the DS3500.
7. Wait until the storage subsystem status is Optimal in the DS Storage Manager software before proceeding to step 10.
8. Add the hard disk drives to the expansion enclosure.

Attention: Make sure that you wait at least 30 seconds for the DS Storage Manager software to recognize each newly added drive before you install another drive.

4.6.7 SAS drive channel miswires

If the SAS cabling on the drive side is incorrect, it will be detected by the DS3500 controller firmware. The Storage Manager application will alert you about the condition by logging a critical Major Event Log event. In addition, the Storage Manager Recovery Guru will point to the miswire condition and advise you of the steps to correct the problem. See 14.13, “Recovery Guru - Recover from Failure” on page 409 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 EXP3500 enclosures in a technically correct manner that is not 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:

- ▶ EXP3500 enclosures are connected to SAS host ports instead of to SAS drive ports on the DS3500 controller.
- ▶ The SAS IN port is connected to the SAS IN port on another EXP3500.
- ▶ The SAS OUT port is connected to the SAS OUT port on another EXP3500.
- ▶ SAS ports are interconnected within a single EXP3500.

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. Although this configuration does work and does not result in a miswire condition, it does not provide redundancy. For example, if an entire EXP3500 fails, the EXP3500 enclosures beneath it will all lose access to the DS3500 subsystem.

4.7 Management connections

There are two connection methods that you can use to manage the DS3500 storage subsystem: host-agent (in-band) management and direct (out-of-band) management.

Note: If you have a dual controller DS3500 subsystem, you must ensure that the management station has a management connection to both controllers. If only one controller is accessible by the management station, then Storage Manager will not be able to correctly manage the enclosure.

The installation and configuration of in-band and out-of-band management connections is described in detail in the Chapter 5, “Installing IBM System Storage DS Storage Manager” on page 131.

4.7.1 Out-of-band management

Out-of-band management uses a network connection to the TCP/IP Ethernet ports on the rear panels of each DS3500 controller (see Figure 4-19 on page 106). Typically, you will only

need to connect one Ethernet port from each controller to your management LAN. You can either specify a fixed TCP/IP configuration or you can use DHCP/BOOTP settings. You can use the host server or a dedicated workstation as your OOB management station.

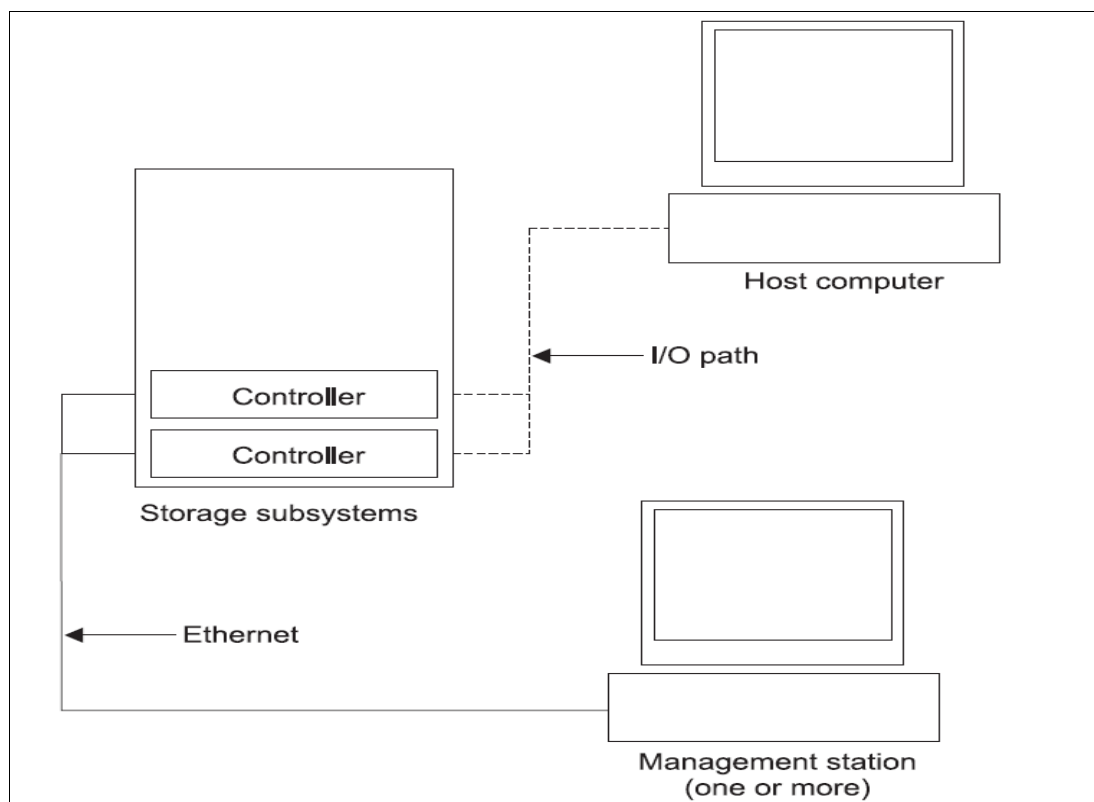


Figure 4-19 Out of band management connections

General rule: We recommend that you set up a dedicated management LAN or VLAN for OOB management of the DS3500, and avoid using your public LAN for OOB management of the DS3500. This will provide increased security of your DS3500 subsystem

Advantages and disadvantages of out-of-band management

In this section we discuss the advantages and disadvantages of using the out-of-band management method.

Managing storage systems using the direct (out-of-band) management method has the following advantages:

- ▶ The Ethernet connections to the controllers enable a management station running SMclient to manage storage systems that are connected to a host running any of the operating systems that are supported by the current level of Storage Manager.
- ▶ In-band management (described in 4.7.2, “In-band management” on page 107) requires an access logical drive to communicate with the controllers. Out-of-band management does not require that you use an access logical drive to communicate with the controllers. The maximum number of logical drives that are supported by the host operating system can be configured and used.
- ▶ The storage system can be managed when there are problems with the Fibre Channel links.

- ▶ Security is enhanced when management LANs/VLANs and more advanced solutions, such as VPN, can be used to manage the system remotely.
- ▶ You can use one DS Storage Manager client to manage more than one DS3500 storage subsystems in the network.

Managing storage systems using the direct (out-of-band) management method has the following disadvantages:

- ▶ Two Ethernet cables are required to connect the storage system controllers to a network.
- ▶ When adding devices, the IP address or host name for each controller must be provided.
- ▶ DHCP/BOOTP server and network preparation tasks are required. This can be avoided by assigning static IP addresses to the controller, or by using the default IP address.

Note: To manage storage subsystems through a firewall, configure the firewall to open port 2463 for TCP and UDP data.

4.7.2 In-band management

When you use the in-band management method (also known as the host-agent method), the DS3500 storage system is managed through the host I/O connections to the host server (see Figure 4-20). You must install the Storage Manager agent (SMagent) software on each host server that you want to configure for in-band management. You can use the host server as the management station, or you can use a dedicated management station that connects to the host server over a network connection.

Note: In-band management is not supported with all of the possible combinations of operating system and host bus adapters. You should always refer to the latest DS Storage Manager documentation for information about any in-band limitations or restrictions that are associated with your particular combination of controller, operating system, and host bus adapter.




Figure 4-20 In-band management connections

Advantages and disadvantages of in-band management

In this section we discuss the advantages and disadvantage of using the in-band management method.

Using in-band management through the host-agent has the following advantages:

- ▶ Ethernet cables do not need to be run to the controllers' management ports.
- ▶ A host name or IP address must only be specified for the host instead of for the individual controllers in a storage system. Storage systems that are attached to the host can be automatically discovered.

Using in-band management through the host agent has the following disadvantages:

- ▶ The host agent uses a special logical drive, called the *access logical drive*, to communicate with the controllers in the storage system. The access logical drive is mapped to the host server that you use for in-band management. Therefore, the host server is limited to one less logical drive than the maximum number that is allowed by the operating system and the host adapter that is being used.
- ▶ Not all operating systems support the access logical drive. In-band management is not supported on these systems.
- ▶ You will be unable to manage or monitor the DS3500 storage subsystem if there is a fault with the host I/O connection between the DS3500 storage subsystem and the host server.

4.8 Host attachment

In this section, we describe the general rules that you should follow when cabling the hosts and DS3500 storage subsystem. We also describe the general rules for Fibre Channel switch SAN zoning and iSCSI LAN connections. We provide examples of the various host attachment implementations with both single and dual controller DS3500 storage subsystems using SAS, Fibre Channel, and iSCSI.

The specific implementation details for the various host connections, including SAN zoning requirements and associated prerequisites, are described as required in the specific host configuration chapters later in this book.

Note: If you have a dual controller DS3500 storage subsystem, you must ensure that each host server has a path to both controllers. A configuration where a host server is only connected to one controller of a dual controller DS3500 storage subsystem is not a supported configuration. This requirement applies to SAS, FC, and iSCSI attachment.

The DS3500 supports multiple-host, multiple-port, and multiple-fabric (Fibre Channel or iSCSI, and SAS) configurations. In other words, the DS3500 storage subsystem supports simultaneous SAS and either iSCSI or Fibre Channel host connections with the optional iSCSI or Fibre Channel host adapter installed, up to a maximum of 64 hosts. Figure 4-21 on page 109 shows an example of this type of configuration with one direct-attached SAS host, and one fabric-attached iSCSI or FC host.

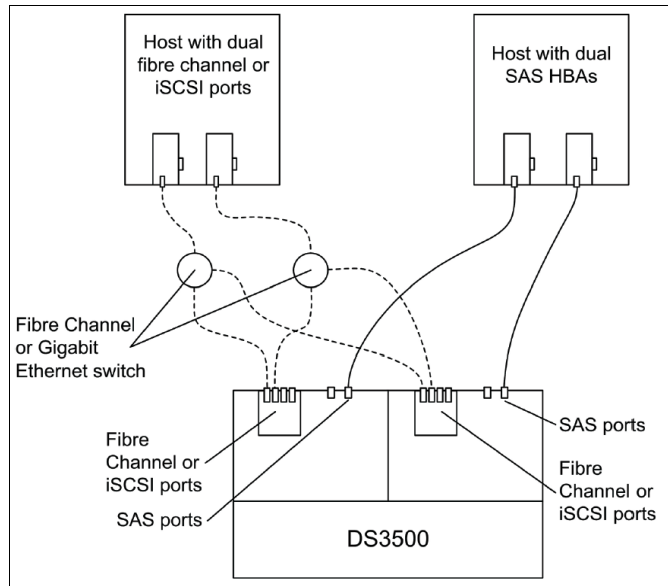


Figure 4-21 An example of a multi-host, multi-port and multi-fabric host configuration

Before you attach hosts to the storage subsystem, you should always ensure that the host systems and host bus adapters (HBAs) are installed correctly and updated with the supported firmware and drivers levels. The supported levels are documented in the readme files that come with the version of DS3500 firmware and DS Storage Manager that you want to use.

4.8.1 SAS attachment

Host server SAS HBAs are used to connect to the DS3500 SAS ports. Each DS3500 controller can have up to four SAS host ports when you install the optional SAS host interface adapter. When the optional SAS host interface adapter is installed, the DS3500 can support up to 28 Blade Center hosts when attached to the SAS Connectivity Module for IBM Blade Center and up to four hosts in a direct attached host environment.

Always use redundant host connections. The use of redundant host connections with a single controller DS3500 storage subsystem is supported but limits the number of hosts to a maximum of two if the optional SAS host card is installed in the DS3500.

SAS host attachment examples

In this section we show examples of SAS host attachment with single and dual controller DS3500 storage subsystems.

Figure 4-22 on page 110 shows the simplest host attachment method: a single host server with one SAS host bus adapter connecting to a single controller DS3500 storage subsystem.

General rule: Although the example shown in Figure 4-22 on page 110 is a valid supported configuration, there is no redundancy of host side connections. Always use redundant host side connections instead.




Figure 4-22 A single host with one SAS connection to the DS3500

Figure 4-23 on page 111 shows a single host with two host adapters connected to a single controller DS3500 storage subsystem. In this example there are redundant host side connections.




Figure 4-23 Single host with two SAS connections to a DS3500 storage subsystem

Figure 4-24 on page 112 shows a host server with two SAS host bus adapters connecting to a dual controller DS3500 storage subsystem. This particular configuration has redundant host side connections and redundant controllers in the DS3500 storage subsystem.




Figure 4-24 Host attachment with two SAS host bus adapters and dual DS3500 controllers

Figure 4-25 on page 113 show two host servers, each with one SAS host adapter, connecting to a single controller DS3500 storage subsystem. In this example, there are no redundant paths and no redundant controller.

General rule: Although the example shown in Figure 4-25 on page 113 is a valid supported configuration, it does not have redundancy of either host-side connections or controllers. Always use redundant host connections.

Note: For a single controller DS3500 storage subsystem, you will need to install the optional SAS host interface card in the controller to attach two hosts using redundant SAS connections.




Figure 4-25 Two hosts connected to a single controller DS3500 using single SAS host adapters

Figure 4-26 on page 114 shows a configuration with two hosts connected using redundant SAS connections to a dual controller DS3500 storage subsystem. This is a good configuration because it has redundant paths to redundant controllers.




Figure 4-26 Two hosts with redundant SAS connections to a dual controller DS3500

4.8.2 iSCSI attachment

The DS3500 supports iSCSI SAN, which utilizes the standard Ethernet infrastructure, using standard Cat 5E and Cat 6 Ethernet cables on the host side. The DS3500 supports a maximum of 64 hosts connecting through a gigabit Ethernet network.

Attention: You must always connect the DS3500 iSCSI host ports to a gigabit Ethernet switch. Direct attachment to host servers using iSCSI is not a supported configuration.

You should always follow these guidelines when connecting iSCSI hosts:

- ▶ Do not configure any two iSCSI interfaces on the same subnet. Each iSCSI interface must be on a separate subnet.
- ▶ If you are using DHCP on more than one interface, multiple interfaces can end up on the same subnet. We recommend that you configure fixed IP addresses for the iSCSI interfaces.
- ▶ Direct iSCSI connections from the host to the storage subsystem are not supported.
- ▶ You should always isolate iSCSI network traffic from your business network traffic by using VLANs or separate Ethernet network hardware.

iSCSI host attachment examples

Here we show examples of the iSCSI host attachment

Figure 4-27 shows the simplest iSCSI host connection with one host connecting to a single controller DS3500 storage subsystem through a gigabit Ethernet switch.

General rule: Although this is a valid supported configuration, there is no redundancy on the host-side connections. Always use redundant host side connections.

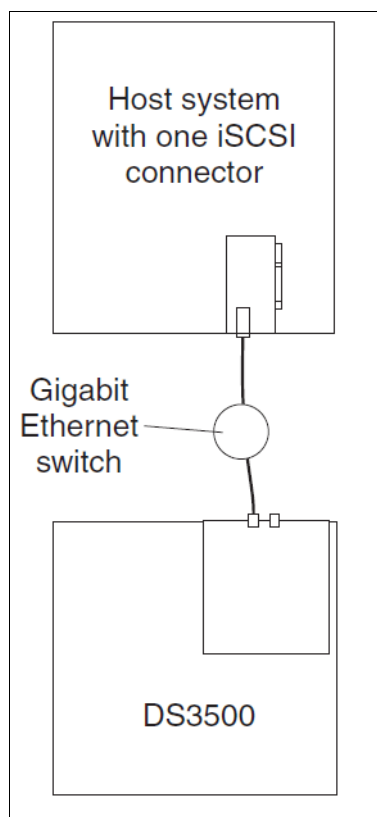


Figure 4-27 One host with a single iSCSI connection to a single controller DS3500

Figure 4-28 on page 116 shows a configuration with redundant host-side connections. There are redundant paths between the host server and the gigabit Ethernet switch, and redundant paths between the switch and the single controller DS3500 storage subsystem. In this example, the DS3500 controller and the gigabit switch are the two single points of failure.

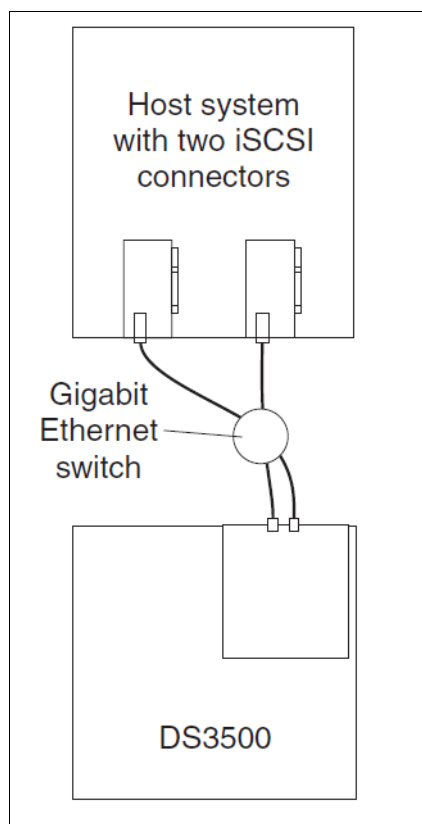


Figure 4-28 Single host with redundant iSCSI connections to a single controller DS3500

Figure 4-29 on page 117 shows a configuration with redundant iSCSI connections to a dual controller DS3500 storage subsystem. In this example, the gigabit Ethernet switch represents a single point of failure. You can achieve additional redundancy of the host side connections by using two gigabit switches: you connect each iSCSI host bus adapter to a separate gigabit Ethernet switch. We show an example of the use of redundant gigabit Ethernet switches in Figure 4-30 on page 118.




Figure 4-29 Single host with redundant iSCSI connections to a dual controller DS3500

Figure 4-30 on page 118 shows a configuration with two hosts connecting to a dual controller DS3500 with redundant iSCSI connections to two gigabit Ethernet switches, and redundant iSCSI connections between the switch and the DS3500 controllers. There is no single point of failure in this example configuration.

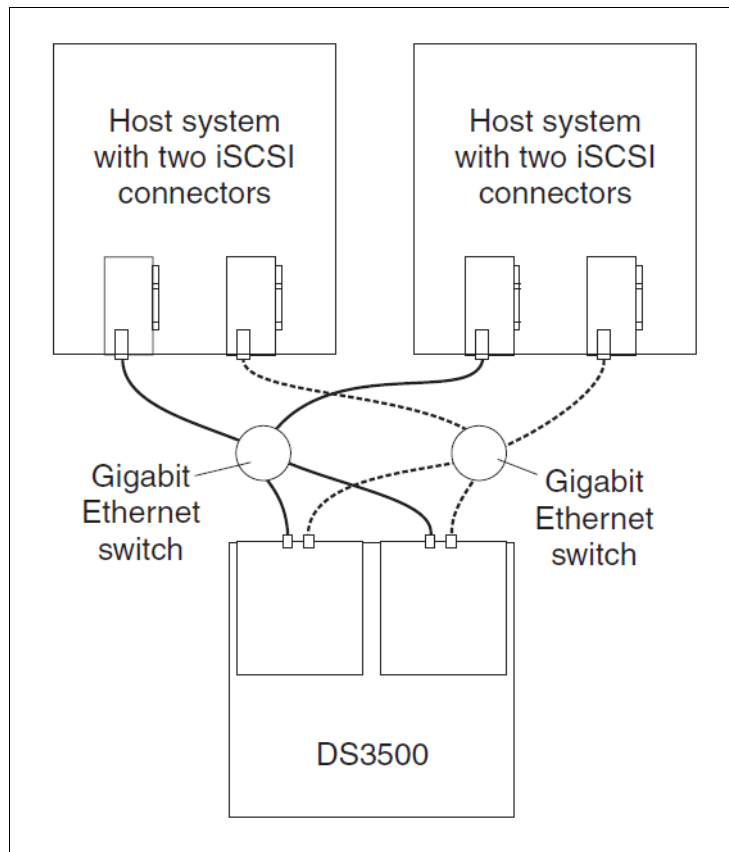


Figure 4-30 Two hosts with redundant iSCSI host connections to a dual controller DS3500

4.8.3 Direct attached Fibre Channel

The DS3500 storage subsystem supports both direct attached and SAN fabric attached hosts. You can attach up to four direct-connected hosts or up to 64 host servers in a switched fabric SAN infrastructure. In this section we describe example configurations of direct-attached Fibre Channel host servers.

Direct-attached DS3500 examples

If you are connecting four or less host servers using Fibre Channel, you can directly connect the four servers to the DS3500.

Figure 4-31 on page 119 shows the simplest configuration: one host server connected to a single controller DS3500 by a single Fibre Channel link.




Figure 4-31 Single direct-attached host server

General Rule: Although this particular configuration is a valid supported configuration, it does not have redundant host connections. Always use redundant host connections as shown in Figure 4-32 on page 120.

Figure 4-32 on page 120 shows a single host server with two host bus adapters and redundant paths to a single controller DS3500.




Figure 4-32 Direct attached host with redundant paths

Although the sample configuration shown in Figure 4-32 has redundant paths to the single controller, the controller is the single point of failure.

Figure 4-33 on page 121 shows a configuration with one host and redundant paths to a dual controller DS3500. This configuration offers both redundant Fibre Channel host connections and redundant controllers, and offers improved redundancy compared to the sample configuration shown in Figure 4-32.




Figure 4-33 Direct attached host with redundant host connections and controllers

Figure 4-34 on page 122 is an example of a configuration with two direct-attached host servers with non-redundant connections to a DS3500 storage subsystem. This sample configuration does not have either redundant host connections or redundant controllers.

Note: Although the example shown in Figure 4-34 on page 122 is a valid supported configuration, it does not have redundant host-side connections. Always use redundant host-side connections.

Note: Because there are four Fibre Channel host ports on the DS3500 (when the optional Fibre Channel host interface card is installed) a maximum of two direct attached hosts with redundant paths can be connected to a single controller DS3500 storage subsystem.




Figure 4-34 Two host serves with non-redundant Fibre Channel connections

The sample configuration shown in Figure 4-35 on page 123 has two direct-attached host servers with redundant connections to a dual controller DS3500. This particular configuration has both redundant host-side connections and redundant controllers.

Note: A maximum of four direct-attached host servers with redundant Fibre Channel connections can be connected to a dual controller DS3500 storage subsystem.

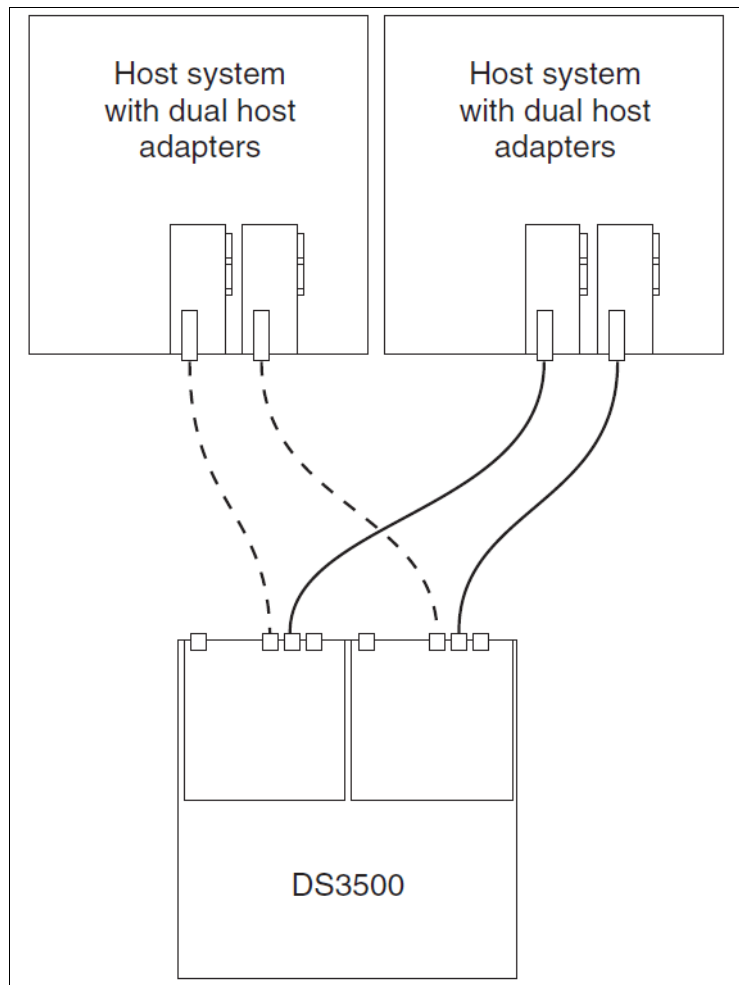


Figure 4-35 Two hosts with redundant connections to a dual controller DS3500

If you need to connect more than four Fibre Channel hosts to a DS3500 storage subsystem using redundant host connections, you will need to use a SAN fabric infrastructure. This is discussed in the next section.

4.8.4 SAN fabric-attached DS3500

If you want to connect more than four host servers using the Fibre Channel interface, you will need to use a switched fabric SAN infrastructure. In this section we describe the general characteristics and requirements of a SAN infrastructure, introduce SAN zoning requirements, and show examples of SAN fabric-attached configurations.

SAN interconnects

A Fibre Channel SAN fabric is used to connect devices. A SAN fabric can be as simple as a single cable connecting two devices. However, the term is most often used to describe a more complex network using cables and interface connectors, HBAs, extenders, and switches.

SAN switches

Fibre Channel switches function in a manner similar to traditional network switches to provide increased bandwidth, scalable performance, an increased number of devices, and in certain

cases, increased redundancy. Fibre Channel switches vary from simple edge switches to enterprise-scalable core switches and Fibre Channel directors.

A switch enables multiple concurrent connections between nodes in the SAN. There are two types of switches: circuit switches and frame switches. Circuit switches establish a dedicated connection between two nodes, whereas frame switches route frames between nodes and establish the connection only when needed, which is also known as a switched fabric.

Inter-Switch Links (ISLs)

SAN switches can be linked together using either standard connections or a type of connection called an *Inter-Switch Link*. Under normal circumstances, traffic moves around a SAN fabric using the Fabric Shortest Path First (FSPF) protocol, which allows data to move around a SAN from initiator to target using the quickest of all available routes. However, it is possible to implement a direct, high-speed path between switches in the form of dedicated ISL connections.

ISL Trunking

Inter-Switch Links (ISLs) can be combined into logical groups to form trunks. In IBM TotalStorage switches, trunks can be groups of up to four ports on a switch connected to four ports on a second switch. At the outset, a trunk master is defined, and subsequent trunk slaves can be added, which has the effect of aggregating the throughput across all links. Therefore, in the case of switches with 8 Gbps ports, if we trunk up to four ports, we allow for a 32 Gbps Inter-Switch Link.

SAN zoning

When you connect a DS3500 storage subsystem to a switched fabric SAN infrastructure, you will normally need to implement a *SAN zoning* configuration on the SAN switch. In this section we provide a brief overview of SAN zoning, and describe general guidelines for implementing a zoning configuration in your SAN infrastructure.

A SAN zoning configuration is implemented on the SAN switch, and it is a method to partition a Fibre-Channel SAN infrastructure into smaller logical subsets. This is done by assigning each Fibre Channel device in the SAN to a specific subgroup called a zone. Zones can vary in size depending on the number of fabric-connected devices, and devices can belong to more than one zone.

Attention: Each SAN zone must contain only one initiator (host bus adapter).

Typically, you can use zones to do the following tasks:

- ▶ Provide security: Use zones to provide controlled access to fabric segments and to establish barriers between operating environments. For example, isolate systems with various uses or protect systems in a heterogeneous environment.
- ▶ Customize environments: Use zones to create logical subsets of the fabric to accommodate closed user groups or to create functional areas within the fabric. For example, include selected devices within a zone for the exclusive use of zone members, or create separate test or maintenance areas within the fabric.
- ▶ Optimize IT resources: Use zones to consolidate equipment logically for IT efficiency, or to facilitate time-sensitive functions. For example, create a temporary zone to back up non-member devices.

Attention: You must always implement a SAN zoning configuration when connecting two or more host bus adapters to a DS3500 storage subsystem through a SAN switched fabric infrastructure. SAN zoning is mandatory when a SAN includes more than one operating system type, or when both tape and disk devices are attached.

Because different applications and host operating systems can require different SAN zoning configurations, the specific SAN zoning requirements for each host operating system are described in the prerequisite sections of the host implementation chapters later in this book.

A step-by-step description of implementing a specific SAN zoning configuration is beyond the scope of this book because the method to implement a SAN zoning configuration can differ greatly between the various switch models supported with DS3500 storage subsystem.

Zone Types

A zone member can be specified using one of the following zone types:

Port Level Zone	A type of zone containing members specified by switch ports (domain ID, port number) only. Port level zoning is enforced by hardware in the switch.
WWPN Zone	A type of zone containing members specified by device World Wide Port Name (WWPN) only. WWPN zones are hardware enforced in the switch.
Mixed Zone	A zone containing various members specified by WWPN and certain members specified by switch port. Mixed zones are software enforced through the fabric name server.

Zones can be hardware enforced or software enforced:

- ▶ In a hardware-enforced zone, zone members can be specified by physical port number, or by WWPN, but not within the same zone.
- ▶ A software-enforced zone is created when a port member and WWPN members are in the same zone.

Note: You do not normally need to explicitly specify a type of enforcement for a zone. The type of zone enforcement (hardware or software) depends on the type of member it contains (WWPNs, ports, or both).

For more complete information regarding Storage Area Networks, see the following IBM Redbooks publications:

- ▶ *Introduction to Storage Area Networks*, SG24-5470
- ▶ *IBM SAN Survival Guide*, SG24-6143

Zoning configuration

Zoning is configured using your switch management software. You can use WWPN zoning to set up each zone so that it contains one server port, and whatever storage device ports that host port requires access to. You do not need to create a separate zone for each source/destination pair. Do not put disk and tape access in the same zone. Also avoid using the same HBA for disk and tape access.

We cannot stress enough that you ensure that all zoning information is fully documented and that the documentation is kept up to date. This information must be kept in a safe location for

reference, documentation, and planning purposes. If done correctly, the document can be used to assist in diagnosing zoning problems.

When configuring World Wide Name (WWN) based zoning, it is important to always use the World Wide Port Name (WWPN), not the World Wide Node Name (WWNN). With many systems, the WWNN is based on the Port WWN of the first adapter detected by the HBA driver. If the adapter that the WWNN is based on fails, and you based your zoning on that WWNN, then your zoning configuration becomes invalid. Subsequently, the host with the failing adapter then completely loses access to the storage attached to that switch.

Keep in mind that you will need to update the zoning information if you ever need to replace a Fibre Channel adapter in one of your servers. Most storage systems such as the DS3500, storage subsystem, and IBM Tape Libraries have a WWN tied to the Vital Product Data of the system unit, so individual parts can usually be replaced with no effect on zoning.

For more details on configuring zoning with your particular switch, see the following IBM Redbooks publications:

- ▶ *Implementing an IBM/Brocade SAN with 8 Gbps Directors and Switches*, SG24-6116
- ▶ *Implementing an IBM/Cisco SAN*, SG24-7545

Multiple fabrics

Depending on the size of your infrastructure, levels of required redundancy, and your budget, you might want more than one switched fabric. Multiple fabrics increase the redundancy and resilience of your SAN by duplicating the fabric infrastructure. With multiple fabrics, the hosts and the resources have simultaneous access to both fabrics, and have zoning to allow multiple paths over each fabric. There are many advantages to using multiple SAN fabrics:

- ▶ Each server can have two or more HBAs. In a two-HBA configuration, each HBA can connect to a separate fabric.
- ▶ Each DS3500 storage subsystem can use separate host ports to connect to multiple SAN fabrics, thus giving it a presence in each fabric.
- ▶ Zoning in each fabric means that the host servers can have many paths to the DS3500 storage subsystem resources, which also means that the zoning has to be done in each fabric separately.
- ▶ The complete loss of one fabric means that the host can still access the resources through the other fabric.

The multiple fabric increases the complexity, resiliency, and redundancy of the SAN infrastructure. This, however, comes at a larger cost due to the duplication of switches, HBAs, zoning administration, and fiber connections. This trade-off has to be carefully examined to see whether your SAN infrastructure requirements justify multiple fabrics.

Sample configurations

In this section we show three sample configurations with two host servers connecting to a dual controller DS3500 storage subsystem through one or more SAN fabrics.

Dual SAN fabric configuration

Figure 4-36 on page 127 shows a configuration with two hosts connecting through two SAN switches to a dual controller DS3500 storage subsystem. Because the SAN switches are not directly connected together, each SAN switch creates its own separate SAN fabric (domain).

In this example, each host bus adapter has a single connection to only one controller in the DS3500. A zoning configuration where each host bus adapter has a single connection to only one controller in a dual controller DS3500 is known as a *one-to-one* (1:1) zoning

configuration. Certain host operating systems and multipath device drivers require a one-to-one zoning configuration. In this example, the two SAN switches can be replaced by a single switch with a one-to-one zoning configuration for the same effective connections, but generally use two switches for path redundancy.




Figure 4-36 Example of a dual Fibre Channel SAN fabric configuration

SAN fabric with Inter-Switch Links

A similar configuration to the configuration in Figure 4-36 but using an *Inter-Switch Link (ISL)* between the two switches is shown in Figure 4-37 on page 128.

In this example, the two SAN switches form a single SAN fabric because of the ISL connection between them.




Figure 4-37 Example of a single Fibre Channel SAN fabric configuration

Cross connected dual SAN fabric

Figure 4-38 on page 129 is an example of a configuration using a cross-connected dual SAN fabric (also known as *mesh*). In a cross-connected SAN fabric, each host bus adapter has a connection to both controllers in a dual controller DS3500 storage subsystem. Certain applications and host operating systems require a cross-connected SAN fabric.




Figure 4-38 Cross-connected SAN fabric connection

4.9 Power Cabling

In this section we describe the power cabling and the correct power on/off sequence for the DS3500 storage subsystem.

4.9.1 The DS3500 power supply

Each DS3500 and EXP3500 enclosure comes with two 585 Watt power supplies as standard. Each power supply has two internal fans for cooling the enclosure.

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

Connecting the power cables of DS3500 subsystems and EXP3500 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, you should consider connecting each power supply to a separate power source so that the failure

of one power source does not cause an unexpected power down event of the entire enclosure.

Note: For controller LED and other hardware information, refer to the appropriate System Storage DS3500 Storage Subsystem: Installation, User's, and Maintenance Guide manual, which can be found at the following URL.

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

4.9.2 Powering on and off

A typical DS3500 configuration includes the following components:

- ▶ Host servers
- ▶ DS3500 storage subsystem
- ▶ EXP3500 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 powered down in case of emergency.
2. EXP3500 enclosures are next. Turn on the switches on both power supplies on each EXP3500 enclosure. Allow enough time for the enclosures to fully power up, which can take several minutes.
3. When all EXP3500 enclosures are up and running, then power on the DS3500. Be aware that large configurations can take up to ten minutes to fully power up.
4. Finally, power on all the host servers when the DS3500 is fully online. You should check the status of the DS3500 by using Storage Manager from your out-of-band management station.

Note: A DS3500 subsystem that powers up with no drives installed will get a dummy FEI. Also, the controllers are accessible through DS Storage Manager and new (different) firmware can be loaded. This means that certain maintenance actions can be taken. Drives can be re-added or powered up and the controllers will reboot to load information out of the DACSTORE on the drives.

Powering off

In normal circumstances, the DS3500 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 DS3500 logical drives are unmounted at the operating system level.

Important: Unless it is an emergency, the DS3500 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.



Installing IBM System Storage DS Storage Manager

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

- ▶ 5.1, “Installing DS Storage Manager on Microsoft Windows 2008” on page 132
- ▶ 5.2, “Installing DS Storage Manager on Linux” on page 140
 - 5.2.2, “Installing Storage Manager using the GUI” on page 141
 - 5.2.3, “Installing DS Storage Manager using a text console” on page 146
- ▶ 5.3, “Installing DS Storage Manager on AIX” on page 149

The sections here describe installing DS3500 Storage Manager on a system that is used only for management purposes. Depending on the usage of the system, additional Storage Manager components might need to be installed. For more details about Storage Manager components and their use, see 2.6, “IBM System Storage DS Storage Manager” on page 22.

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.

5.1 Installing DS Storage Manager on Microsoft Windows 2008

This section describes how to install DS Storage Manager on a Windows 2008 system. Depending on the usage of the system, different components need to be installed, as described in 2.6, “IBM System Storage DS Storage Manager” on page 22.

5.1.1 Installation preparation

Prior the installation, you can follow these steps:

1. Download the Storage Manager installation package from the IBM Support website. 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 DS3500 fixes, code, and documentation, see Appendix A, “IBM Support Portal website” on page 719.
2. Unpack the compressed archive into a local directory. The compressed archive contains:
 - A README file with information about the package and the installation
 - 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.

These steps are applicable for Manager and Client options.

5.1.2 Installing the Storage Manager Client on Microsoft Windows 2008

To install the Storage Manager client on a Microsoft Windows 2008 server, follow these steps:

1. Run the installer file. Our version is SMIA-WS32-10.70.35.08.exe. The files unpack as shown in Figure 5-1.

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




Figure 5-1 IBM DS Storage Manager InstallAnywhere

2. Select the language to display the license agreement, and click **OK** (see Figure 5-2).

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




Figure 5-2 IBM DS Storage Manager Language Selection

3. Read the introduction and click **Next**, as shown in Figure 5-3.

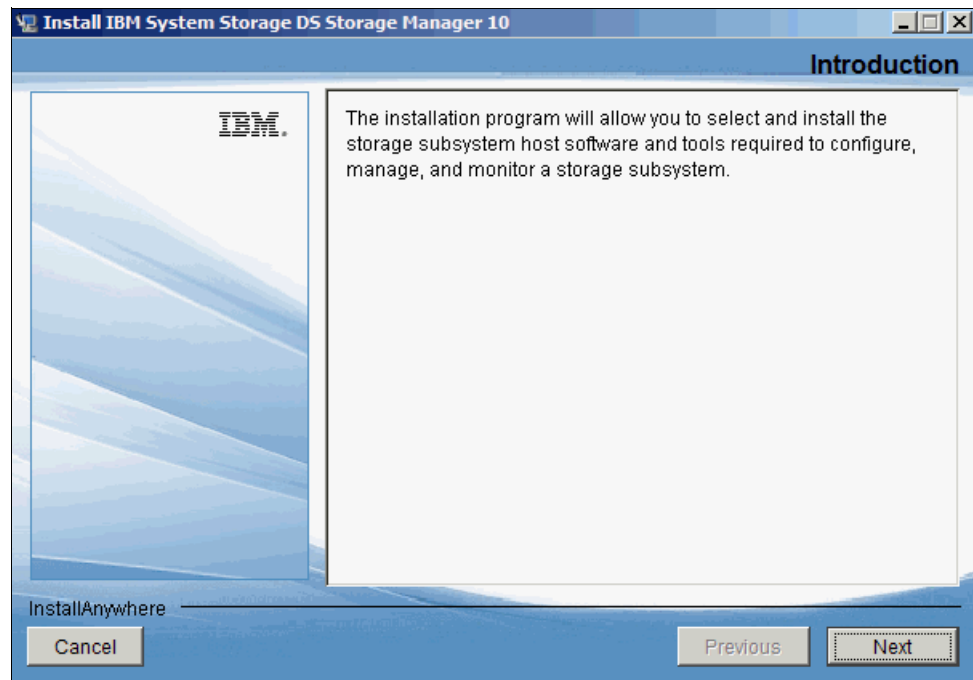


Figure 5-3 IBM DS Storage Manager Introduction

4. Read the copyright statement and click **Next** as shown in Figure 5-4.



Figure 5-4 IBM DS Storage Manager Copyright Statement

5. Read the license agreement and click **Next** as shown in Figure 5-5 on page 135.

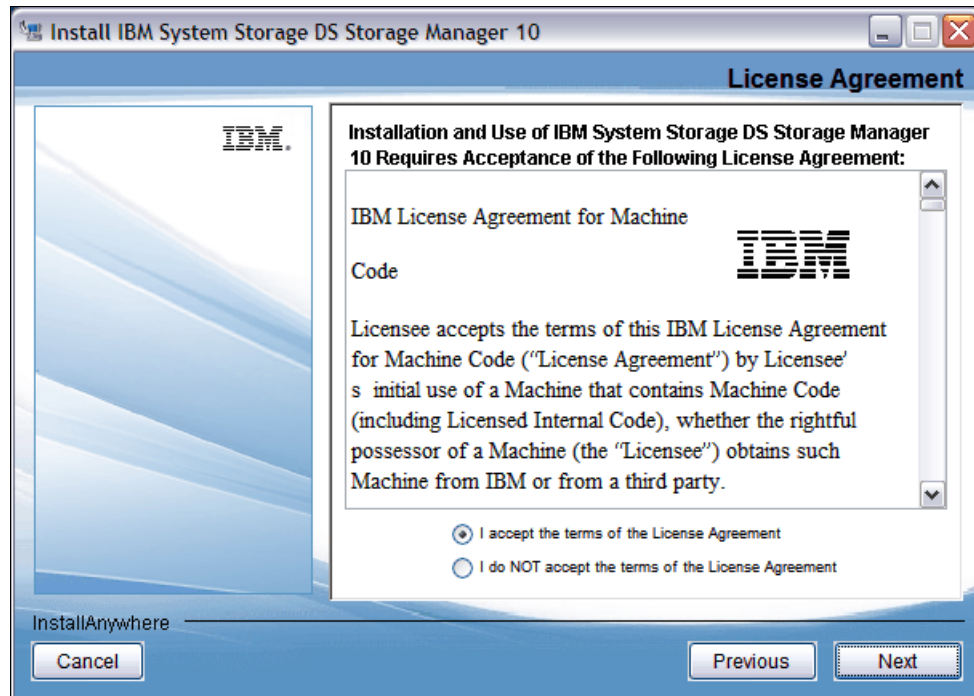


Figure 5-5 IBM DS Storage Manager License Agreement

6. Specify the directory where you want to install as shown in Figure 5-6. The default is C:\Program Files (x86)\IBM_DS.




Figure 5-6 IBM DS Storage Manager Installation Folder

7. For Installation Type (Figure 5-7), 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 5-7 IBM DS Storage Manager Select Installation Type

8. In Figure 5-8 on page 137, select **Install Set** from the drop-down menu. The options in the drop-down menu are same as shown in Figure 5-7. When you select an Install Set, you can see what components will be installed and modify them if you desire. See 2.6, "IBM System Storage DS Storage Manager" on page 22 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.

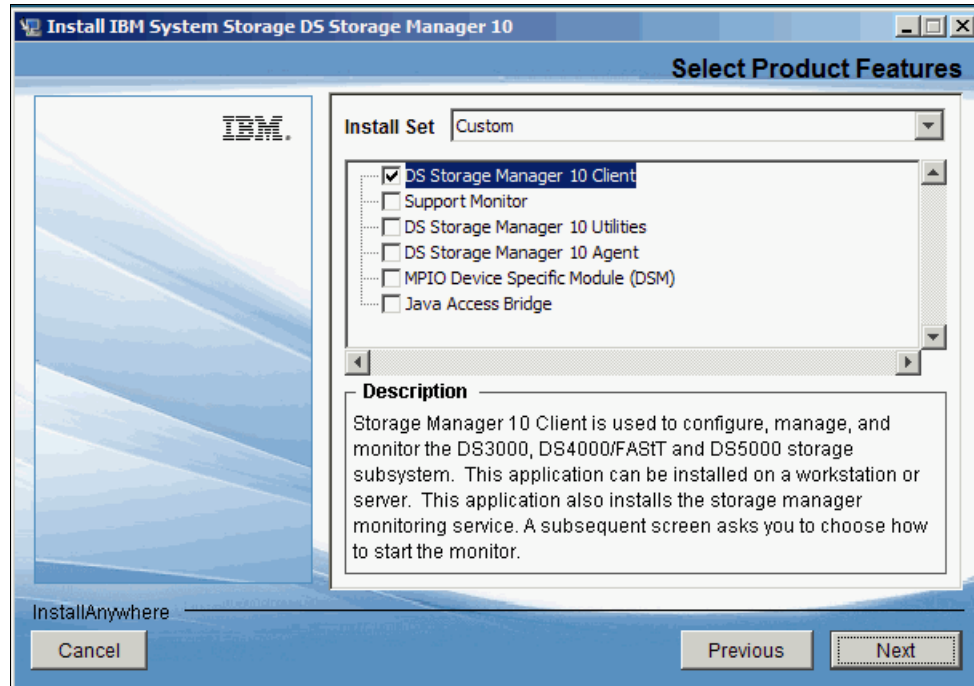


Figure 5-8 Select product features to install

9. If a version of Storage Manager is already installed, you will see the warning shown in Figure 5-9. 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 will not see this warning.

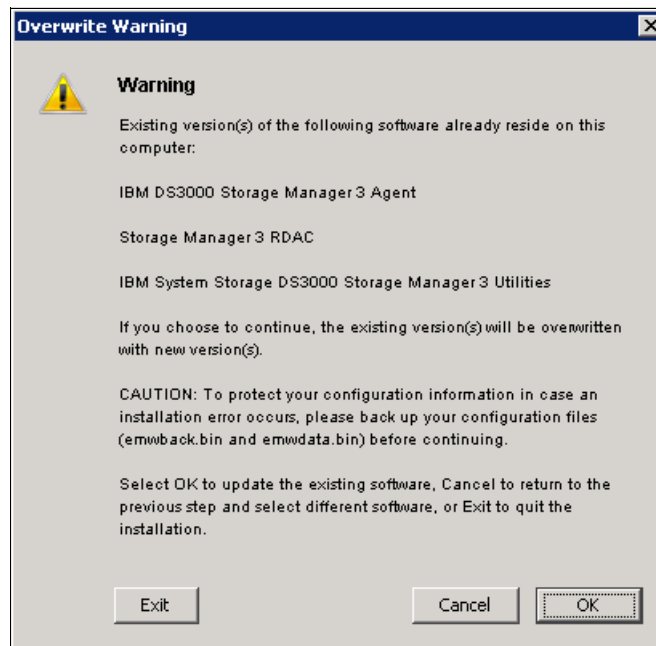


Figure 5-9 Overwrite Warning

10. In Figure 5-10, you can select whether to automatically start the Storage Manager Monitor. Click the **Next** button.

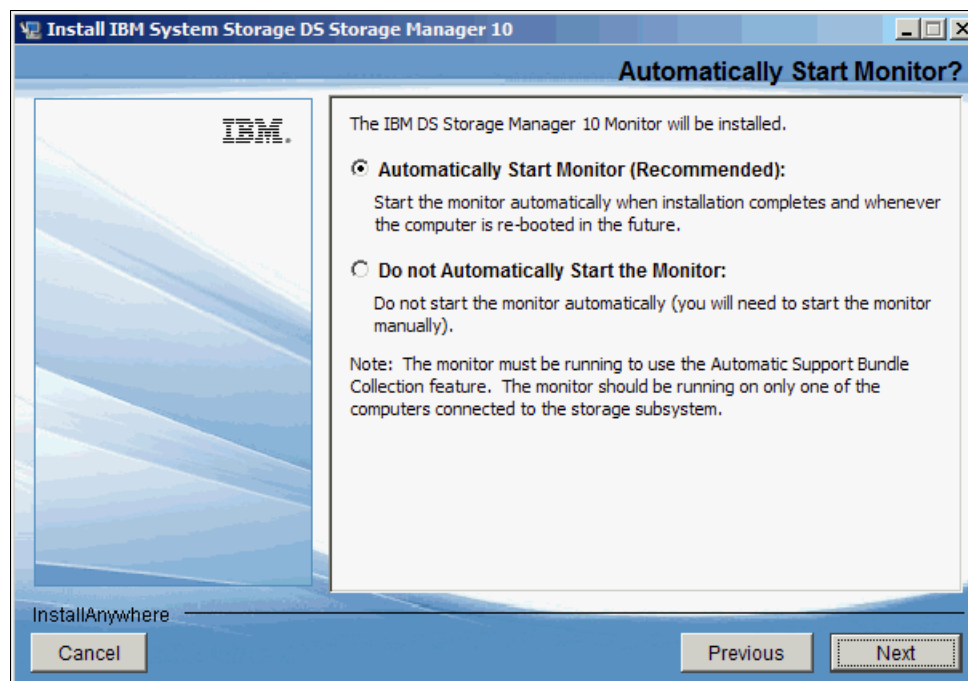


Figure 5-10 IBM DS Storage Manager Automatically Start Monitor

11. Verify the pre-installation summary and click **Install**. See Figure 5-11.




Figure 5-11 IBM DS Storage Manager Pre-Installation Summary

12. The installation status window will appear during the installation as shown in Figure 5-12.

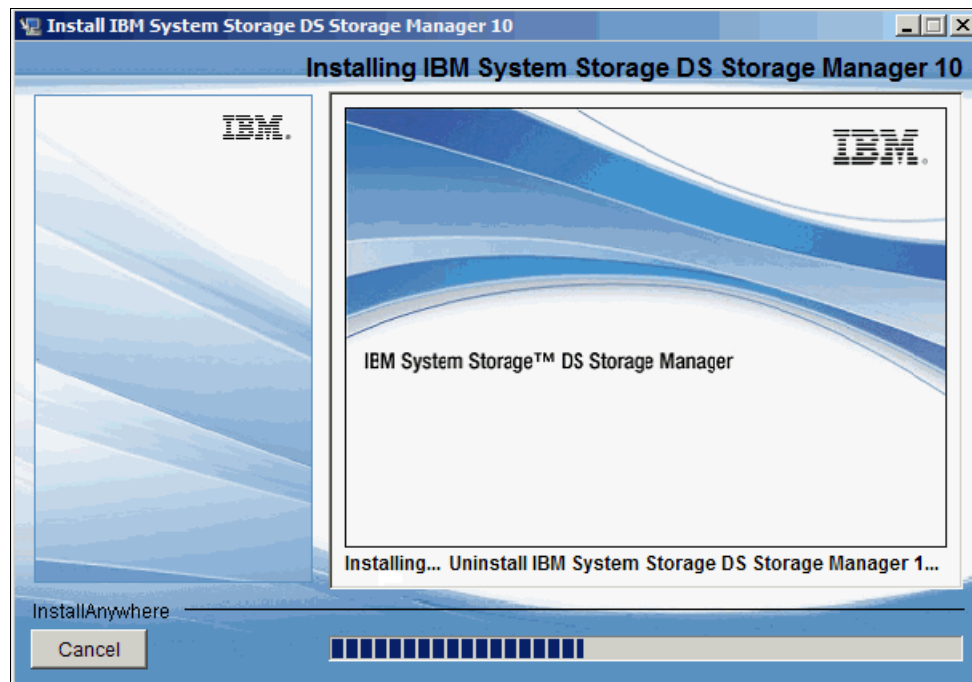


Figure 5-12 IBM DS Storage Manager Installation status window

After the installation has finished, verify that there were no errors and click **Done**. See Figure 5-13.




Figure 5-13 IBM DS Storage Manager Install Complete

5.2 Installing DS Storage Manager on Linux

The procedure to install DS Storage Manager on Linux is similar to Windows, with one major difference: the multipath driver is not included in the Storage Manager installation package. It must be installed separately from the Storage Manager package. Refer to 20.1.2, “Installing the RDAC Multipath Driver” on page 593 for details about installing the multipath drivers.

The IBM Systems Storage DS 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:

- ▶ 5.2.1, “Preparing for the installation” on page 140
- ▶ 5.2.2, “Installing Storage Manager using the GUI” on page 141
- ▶ 5.2.3, “Installing DS Storage Manager using a text console” on page 146

After following these steps you will be able to start IBM Storage Manager by running *SMclient* from a terminal window.

5.2.1 Preparing for the installation

To prepare for the installation, use the following procedure:

1. Download the latest version of the IBM Systems Storage DS Storage Manager Client to the management station or host connected to the DS3500 for inband or outband connections from the IBM Support website. You can find more information about the IBM Support website in Appendix A, “IBM Support Portal website” on page 719.
2. Verify that there is enough space available in the /opt (1.6 MB for the complete installation) and /tmp (1 GB) directories to perform this installation. It is possible to set the IATEMPDIR environment variable to change the temporary directory to extract the files on a disk partition with enough free disk space instead of using /tmp.

Note: The setting of this variable ONLY changes the temporary location where the IBM Storage Manager will extract the files for the installation.

To set the variable enter one of the following commands at the UNIX command line prompt before running this installer again:

- for Bourne shell (sh), ksh, bash and zsh:

```
$ IATEMPDIR=/your/free/space/directory
```

```
$ export IATEMPDIR
```

- for C shell (csh) and tcsh:

```
$ setenv IATEMPDIR /your/free/space/directory
```

3. Extract the file to a local directory as shown in Example 5-1.

Example 5-1 Extract package archive

```
# tar -zxvf SM10.70_Linux_SMIA-10.70.xx.08.tgz
Linux10p70/Linux/
Linux10p70/Linux/SMIA-LINUX-10.70.A5.08.bin
Linux10p70/SM10R2MIB/
Linux10p70/SM10R2MIB/SM10_R2.MIB
Linux10p70/SM10R2MIB/README_First.txt
```

4. The following files are unpacked:
 - The installation package and executable
 - The readme file that describes the package
 - A subdirectory SM10R2MIB that contains the MIB files

Note: This MIB file will only work with the IBM DS4000/DS5000 TotalStorage Storage Manager.

5. The file SMIA-LINUX-10.70.A5.08.bin is the installer that contains the installation files. To be able to execute this installer, you need to type `sh` before the file or change the file permissions to be able to execute.
6. By default, the installer requires a graphical console to launch, but can be also installed on a text console or left unattended.

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

The options are:

- **swing:** 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.
 - **console:** Use a text based console for the installation.
 - **silent:** Perform an unattended installation.
7. Choose which method you want to use, and launch the installation file, specifying your option with the `-i` parameter. For example:

```
$sh SMIA-LINUX-10.70.A5.08.bin -i console
```

5.2.2 Installing Storage Manager using the GUI

Note: It is possible to run the GUI installation remotely exporting the display by using X11 forwarding.

To install Storage Manager, perform the following steps:

1. Log on to the Linux server with administrative rights, that is, usually the root user, on a graphical console such as KDE, CDE, or Gnome.
2. Start the installer that was extracted in step 3 on page 140 in a terminal window by running the following command:

```
sh SMIA-LINUX-10.70.A5.08.bin
```
3. A window will open. Select the language to display the license agreement and click **OK** as shown in Figure 5-14 on page 142.

Note: The installation will always be performed in English.

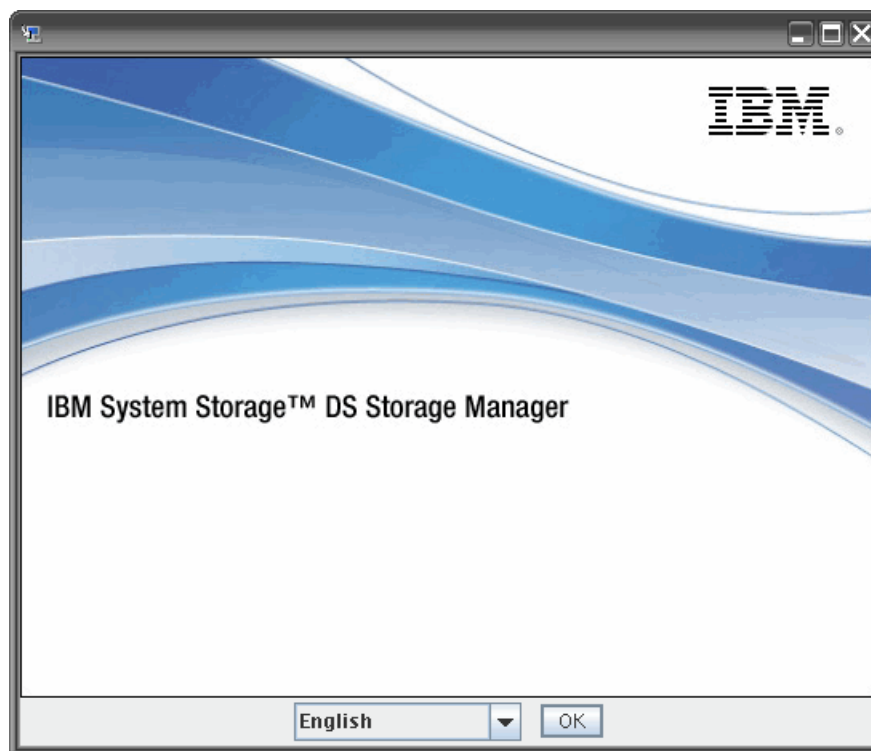


Figure 5-14 Select the language to display license agreement

4. Read the introduction and click **Next** as shown in Figure 5-15



Figure 5-15 Introduction

5. Read the copyright statement and click **Next** as shown in Figure 5-16

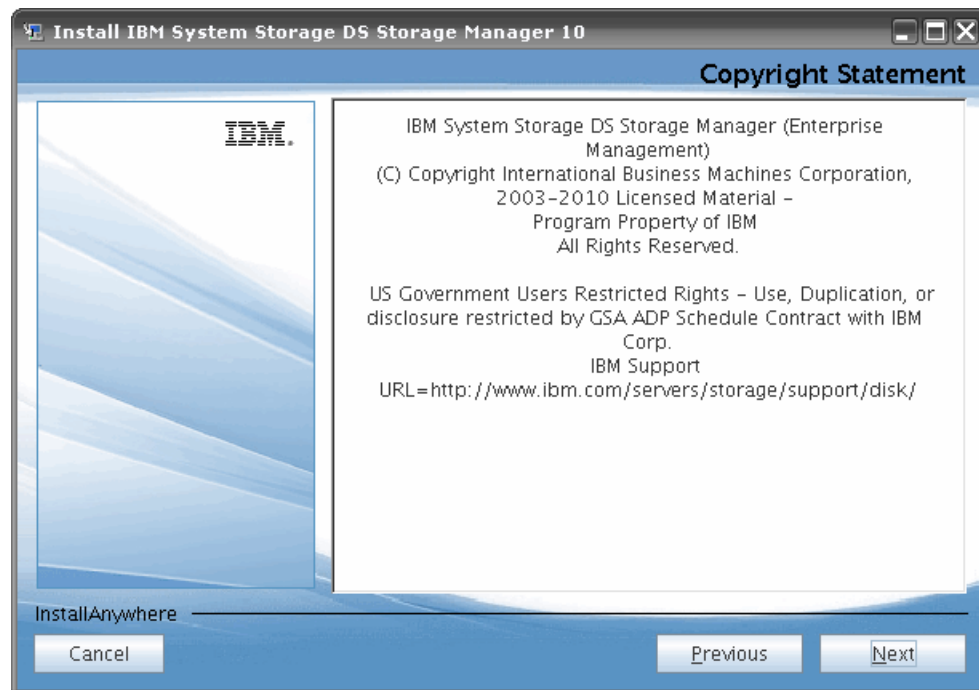


Figure 5-16 Copyright statement

6. Read and accept the license agreement. Click **Next** as shown in Figure 5-17



Figure 5-17 License agreement

7. Select **Custom** installation type and click **Next** as shown in Figure 5-18. 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.

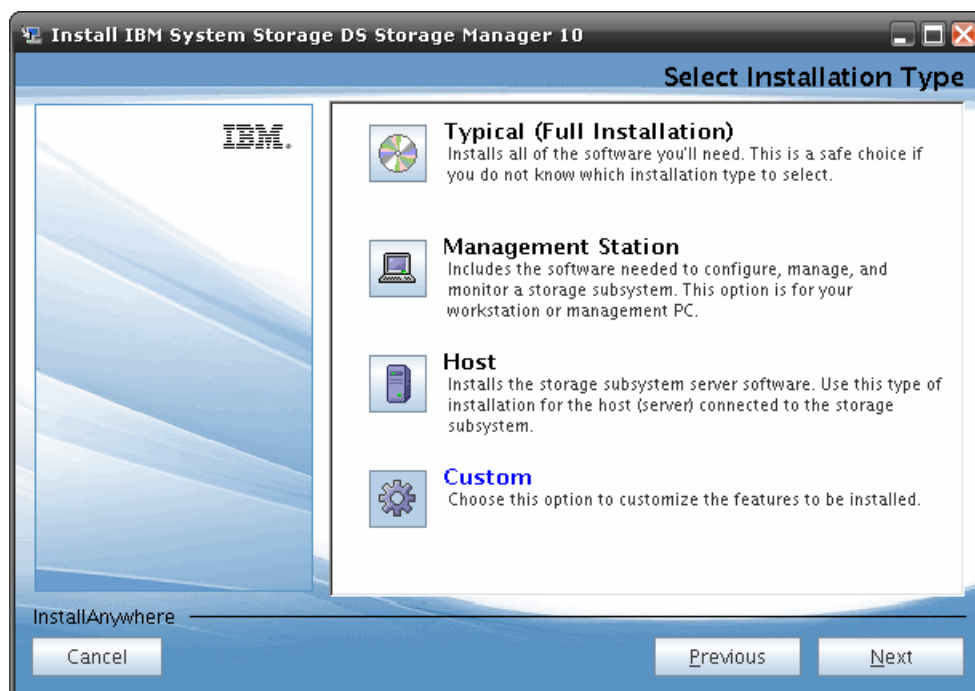


Figure 5-18 Installation type

8. Select the Install set to be installed (DS Storage Manager 10 Client in our case). Click **Next** as shown in Figure 5-19.

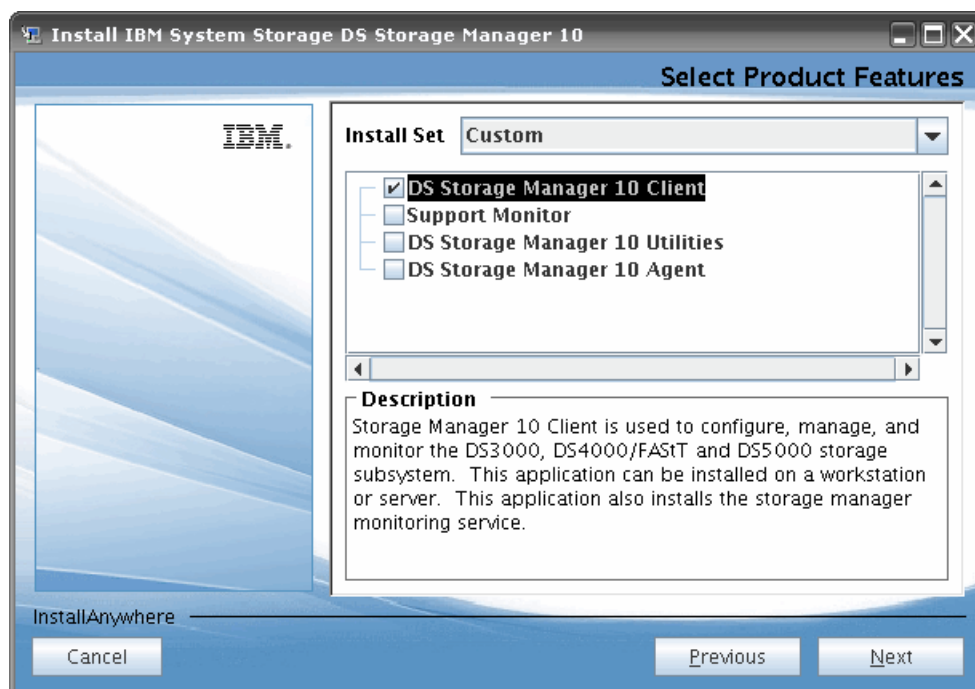


Figure 5-19 Select product features

9. Verify the installation options and click **Next** as shown in Figure 5-20



Figure 5-20 Pre-installation summary

10. After a successful installation, a similar window as in Figure 5-21 will be shown. If there are no errors, click **Done** to exit.



Figure 5-21 Install complete

11. Optionally, verify the installation by running the command `rpm -qa |grep SM` and look for the following RPM packages in the installed package list:

```
SMclient-10.70.G5.08-1.noarch
SMesm-10.70.G5.02-1.noarch
SMruntime-10.70.A5.00-1.i586
```

5.2.3 Installing DS Storage Manager using a text console

It is possible to install the IBM Systems Storage DS Storage Manager without any graphical interface by using the console mode installation. 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.

To install using a text console, perform the following steps:

1. Start the installer with the option **-i console**. Select the locale (the default is English), and press Enter as shown in Example 5-2

Example 5-2 Text based installation

```
# sh SMIA-LINUX-10.70.A5.08.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...
```

```
Launching installer...
Preparing CONSOLE Mode Installation...
```

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

```
1- Deutsch
->2- English
3- Español
4- Francais
5- Italiano
6- Portugues (Brasil)
```

```
CHOOSE LOCALE BY NUMBER:
```

2. Read the introduction and press Enter to proceed.
3. Read the copyright statement, press Enter to continue to the following page, and press Y to accept it.
4. Choose the type of installation to perform as shown in Example 5-3. 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. Press 4 for Custom and press Enter.

Example 5-3 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. In this custom installation panel, we can select or deselect individual components. By default, all components are selected and you can deselect the component by entering the number associated with the component you do not want to install. Here we only choose the DS Storage Manager 10 Client as shown in Example 5-4.

Example 5-4 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- ☒ DS Storage Manager 10 Client
- 2- ☐ Support Monitor
- 3- ☐ DS Storage Manager 10 Utilities
- 4- ☐ DS Storage Manager 10 Agent

Choose the Features to be installed by this installer:

6. Confirm the pre-installation summary and press Enter as shown in Example 5-5.

Example 5-5 Storage Manager - Pre-installation summary

Pre-Installation Summary

Review the Following Before Continuing:

Install Folder:
/opt/IBM_DS

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

Required Disk Space
632 MB

Available Disk Space
506 MB

PRESS <ENTER> TO CONTINUE:

7. The installation is now performed into the directory /opt/IBM_DS. After successful installation, a message like Example 5-6 will be shown.

Example 5-6 Installation complete

```
Installation Complete
-----

Congratulations. IBM_DS3000 has been successfully installed to:

/opt/IBM_DS

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

5.2.4 Uninstalling DS Storage Manager on Linux

Use the following steps to uninstall DS Storage Manager:

1. Run the following at the command prompt:

```
# /Storage_Manager_dir/Uninstall\ IBM\ System\ Storage\ DS\ Storage\ Manager\ 10
```

Note: Normally, the IBM Storage Manager directory by default is /opt/IBM_DS. It is possible to discover which is the current directory by running the following command:

```
rpm -qil SMclient-10.70.G5.08-1.noarch
```

At the time of writing, that is the name of the package. The current name of the RPM console package should appear with the following command:

```
rpm -qa |grep SM
```

2. Read the introduction and press Enter to continue as shown in Example 5-7.

Example 5-7 Press Enter to continue

```
Preparing CONSOLE Mode Uninstallation...
=====
IBM System Storage DS Storage Manager 10          (created with InstallAnywhere)
-----
=====
Uninstall IBM System Storage DS Storage Manager 10
-----

About to uninstall...

IBM System Storage DS Storage Manager 10

This will remove features installed by InstallAnywhere. It will not remove
files and folders created after the installation.

PRESS <ENTER> TO CONTINUE:
```

3. Select the features to uninstall as shown in Example 5-8

Example 5-8 Features to uninstall

```
=====
Uninstall Options
-----

ENTER THE NUMBER FOR YOUR CHOICE, OR PRESS <ENTER> TO ACCEPT THE DEFAULT:

->1- Completely remove all features and components.
    2- Choose specific features that were installed by InstallAnywhere.

Please choose one of the following options:: 1
```

4. If the uninstallation is successful, a message similar to the one in Example 5-9 will be shown.

Example 5-9 Complete uninstallation

```
=====
Uninstall Complete
-----

All items were successfully uninstalled.

linux:/opt/IBM_DS/Uninstall IBM System Storage DS Storage Manager 10 #
```

5.3 Installing DS Storage Manager on AIX

The procedure to install DS Storage Manager on AIX is similar to Linux, but with some differences. The size requirements for the installation and the package manager are different. As on Linux the multipath drivers is not included in the installation package. It must be installed separately from the Storage Manager package. Refer to 21.2, “Configuring MPIO” on page 632 for more details about the AIX multipath drivers.

As on Linux, the IBM Storage Manager can be installed using a GUI or a text-based interface. Because we already described the graphic and console mode in 5.2.2, “Installing Storage Manager using the GUI” on page 141 and in 5.2.3, “Installing DS Storage Manager using a text console” on page 146, we only are going to include the preparation for the installation.

After following this steps you should have installed the packages detailed below and will able to start IBM Storage Manager by running *SMclient* from a terminal window.

The packages are shown by running the command `ls1pp -a1 |grep SM`

```
► SMclient.aix.rte      10.70.6508.0  COMMITTED  IBM DS Storage Manager 10
► SMesm.aix.rte         10.70.6502.0  COMMITTED  IBM DS Storage Manager ESM
► SMruntime.aix.rte     10.70.6500.0  COMMITTED  IBM DS Storage Manager 10
► SMutil.aix.rte        10.0.6514.0   COMMITTED  IBM DS Storage Manager 10
```

5.3.1 Preparing for the installation

To prepare for the installation, use the following procedure:

1. Download the latest version of the IBM Systems Storage DS Storage Manager Client to the management station or host connected to the DS3500 for inband or outband connections from the IBM Support website. Refer to Appendix A, “IBM Support Portal website” on page 719 for more information about the IBM Support website.
2. Verify there is enough space available in the /usr (615 MB), /opt (450 MB), /var (31 MB) and /tmp (1 GB) directories to perform this installation. It is possible to set the IATEMPDIR environment variable to change the temporary directory to extract the files on a disk partition with enough free disk space instead of using /tmp.

Note: The setting of this variable ONLY changes the temporary location where the IBM Storage Manager will extract the files for the installation.

To set the variable enter one of the following commands at the UNIX command line prompt before running this installer again:

- for Bourne shell (sh), ksh, bash and zsh:

```
$ IATEMPDIR=/your/free/space/directory
```

```
$ export IATEMPDIR
```

- for C shell (csh) and tcsh:

```
$ setenv IATEMPDIR /your/free/space/directory
```

3. Extract the file to a local directory as shown in Example 5-10.

Example 5-10 Extract package archive

```
# gzip -d SM10.70_Aix_SMIA-10.70.xx.08.tgz
# tar -xvf SM10.70_Aix_SMIA-10.70.xx.08.tar
x Aix10p70/Aix
x Aix10p70/Aix/SMIA-AIX-10.70.65.08.bin, 707704559 bytes, 1382236 tape blocks
x Aix10p70/SM10R2MIB
x Aix10p70/SM10R2MIB/SM10_R2.MIB, 5298 bytes, 11 tape blocks
x Aix10p70/SM10R2MIB/README_First.txt, 126 bytes, 1 tape blocks
```

4. The following files are unpacked:
 - The installation package and executable
 - The readme file that describes the package
 - A subdirectory SM10R2MIB that contains the MIB files

Note: This MIB file will only work with the IBM DS4000/DS5000 TotalStorage Storage Manager.

5. The file SMIA-AIX-10.70.65.08.bin is the installer that contains the installation files. To be able to execute this installer, you need to type sh before the file or change the file permissions to be able to execute.
6. By default, the installer requires a graphical console to launch, but can also be installed on a text console or left unattended.

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

The options are:

- **swing**: 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.
 - **console**: Use a text based console for the installation.
 - **silent**: Perform an unattended installation.
7. Choose which method you want to use, and launch the installation file, specifying your option with the -i parameter. For example:
- ```
$sh SMIA-AIX-10.70.65.08.bin -i console
```

## 5.4 Completing the DS Storage Manager installation

This section contains procedures for using the Enterprise Management and Subsystem Management features of DS Storage Manager to complete the storage management installation tasks for all host operating systems.

To complete a Storage Manager installation, the following tasks must be performed:

- ▶ Initial automatic discovery of storage subsystems
- ▶ Initial manual discovery of storage subsystems

The Enterprise Management window opens when you start the DS Storage Manager. You can use the Enterprise Management window to do the following:

- ▶ Add and discover the storage subsystems
- ▶ View all storage subsystems in your management domain

The other options of the Enterprise Management window are described in Chapter 6, “Administration - Enterprise Management” on page 157.

### 5.4.1 Performing an automatic discovery of storage subsystems

Complete the following steps to perform an initial automatic discovery of storage subsystems:

1. From the Setup tab DS Storage Manager Enterprise Management window, click the **Add Storage Subsystem** link as shown in Figure 5-22




Figure 5-22 Add Storage Subsystem

2. Select Automatic discovery as shown in Figure 5-23 and click **OK**.

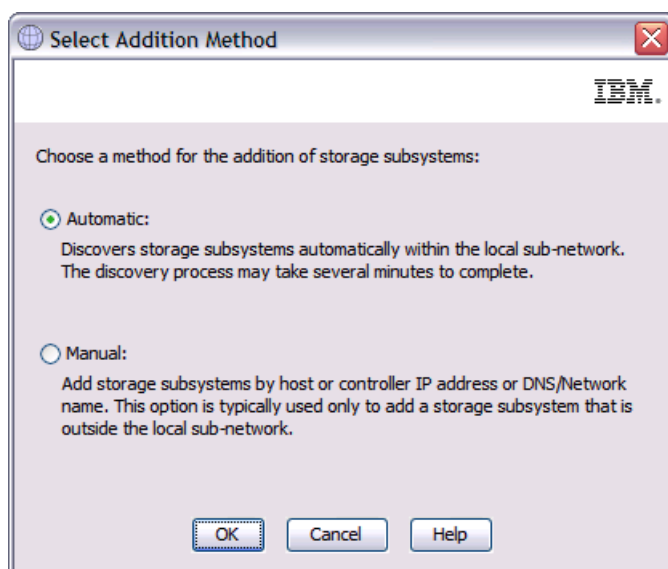


Figure 5-23 Select Addition Method - Automatic

3. Click **OK** to begin an initial automatic discovery of hosts and storage subsystems attached to the local subnetwork as shown in Figure 5-24.

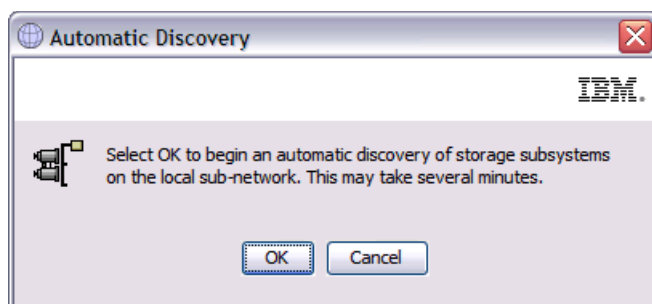


Figure 5-24 Automatic Discovery

After the initial automatic discovery is complete, the Devices tab of the Enterprise Management window displays all hosts and storage subsystems attached to the local subnetwork.

**Note:** The Enterprise Management window can take up to a minute to refresh after an initial automatic discovery.

## 5.4.2 Performing a manual discovery of storage subsystems

You can manually add more hosts or storage subsystems. You can use this option to selectively manage a group of storage subsystems from an SMclient. You can also use this option to add additional devices to be managed that were not discovered during the SMclient initial discovery.

**Note:**

- ▶ When you add new storage subsystems to the existing storage subsystems in a SAN that are managed through the host-agent software, you must stop and restart the host-agent service. When the host-agent service restarts, the new storage subsystem is detected. Then go to the Enterprise Management window and click **Tools** → **Rescan** to add the new storage subsystems to the management domain.
- ▶ When you add new storage subsystems to existing storage subsystems that are managed using the direct-management method, be sure to specify the IP addresses for both controllers.

To add the Storage Subsystem manually, perform these steps:

1. From the Setup tab DS Storage Manager Enterprise Management window, click the **Add Storage Subsystem** link as shown in Figure 5-25.

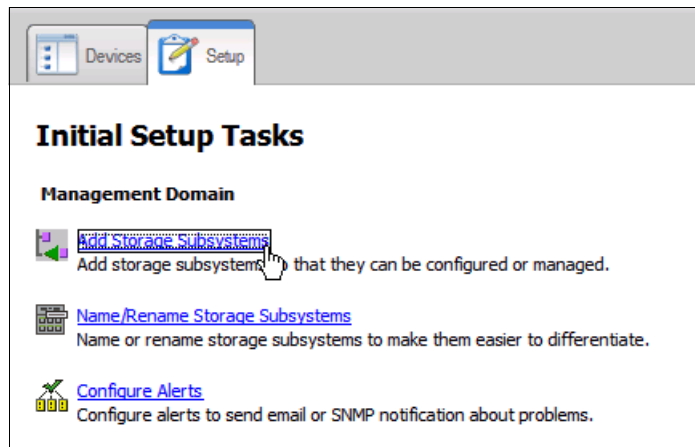


Figure 5-25 Add Storage Subsystem

2. Select Manual as shown in Figure 5-26 and click **OK**.

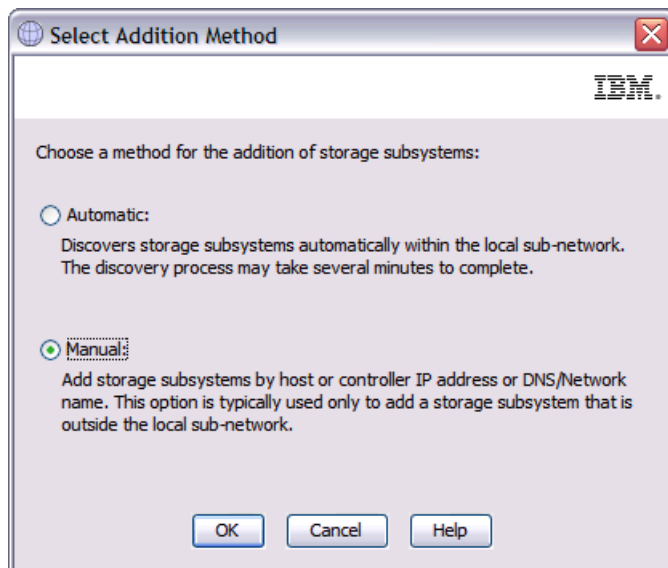


Figure 5-26 Select Addition Method-Manual

3. Select a management method and add the Storage Subsystem:
  - When using the Out-of-band management, select Out-of-band management and enter the IP addresses of both the controllers as shown in Figure 5-27.

The screenshot shows a Windows-style dialog box titled "Add New Storage Subsystem - Manual" with an IBM logo in the top right corner. The dialog contains several links: "What are in-band and out-of-band management connections?", "Adding controllers with more than one Ethernet port", and "What if my system only has one controller?". Below these links, the text "Select a management method:" is followed by two radio button options. The first option, "Out-of-band management:", is selected and includes the instruction "Manage the storage subsystem using the controller Ethernet connections." Below this, there are two text input fields for controller addresses, both labeled "Controller (DNS/Network name, IPv4 address, or IPv6 address):". The first field contains "192.168.1.10" and the second field contains "192.168.1.11". The second option, "In-band management:", is unselected and includes the instruction "Manage the storage subsystem through an attached host." Below this, there is a text input field for a host address, labeled "Host (DNS/Network name, IPv4 address, or IPv6 address):", which is currently empty. At the bottom of the dialog are three buttons: "Add", "Cancel", and "Help".

Figure 5-27 Out-of-band management

- For In-band management, select the In-band management option and enter the IP address of attached host as shown in Figure 5-28 on page 155.

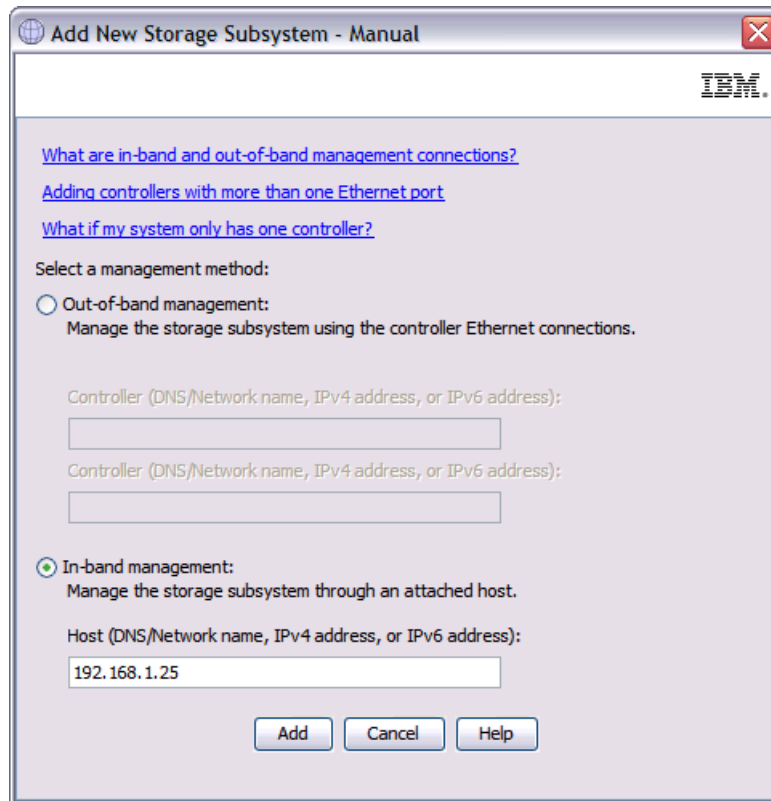


Figure 5-28 In-band management

Click **Add** to complete the process. After this is successfully added, you will have an option to add another Storage Subsystem as shown in Figure 5-29, and you can repeat the process until all the desired Storage Subsystems are added.

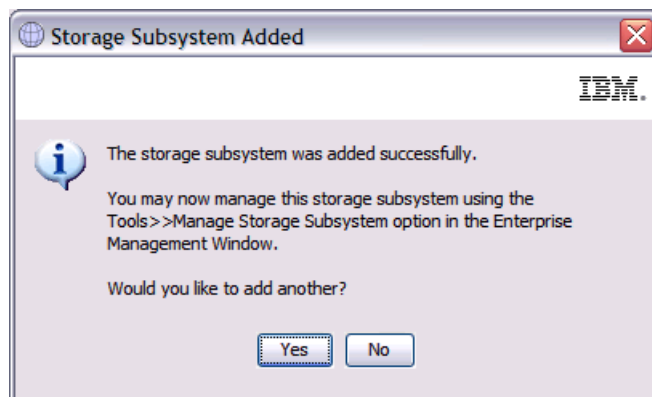


Figure 5-29 Storage Subsystem Added

### 5.4.3 Add Storage Subsystem verification

Verify that each host and storage subsystem displays in the Enterprise Management window. If a host or storage subsystem is not displayed, perform the following tasks:

- ▶ Check the hardware and hardware connections for possible problems.
- ▶ If you are using the network management method (commonly known as out-of-band management), verify that all hosts and storage subsystems are connected to the same

subnet network. If you are using the host-agent method (commonly known as in-band management), ensure that the Fibre Channel connection between the host and storage subsystems is made.

- ▶ Make sure that all of the preparation steps for setting up the storage subsystem for a network managed system are completed. Use the **Add Device** option to add the IP addresses of the storage subsystem. Add both IP addresses of the controller. Otherwise, you will get a "partially-managed device" error message when you try to manage the storage subsystem.

**Note:** To use the auto-discovery method, the storage subsystem and this host must be on the same subnet. Otherwise use the manual method to add a subsystem.

- ▶ If you are using the host-agent management method, perform the following steps:
  - Make sure that the SMagent is installed on the host.
  - Verify that you have a Fibre Channel connection from the storage subsystems to the host that has the SMagent installed.
  - Verify that all of the preparation steps are complete, then perform the following steps:
    - Run the hot\_add utility.
    - Restart the SMagent.
    - Right-click the host, and click **Tools** → **Rescan** in the Enterprise Management window.

**Note:** In certain situations, a storage subsystem might be duplicated in the device tree after an automatic discovery. You can remove a duplicate storage management icon from the device tree by using the **Remove Device** option in the Enterprise Management window.

- ▶ Verify that the status of each storage subsystem is Optimal. If a device shows a status of Unresponsive, right-click the device and select **Remove Device** to delete it from the management domain. Then use the **Add Device** option to add it to the management domain again.



# Administration - Enterprise Management

In this chapter, we describe the IBM System Storage DS3500 Storage Manager Enterprise Management window, including its functions and how to use it.

## 6.1 Enterprise Management window overview

When you launch Storage Manager, you get the Enterprise Management window (shown in Figure 6-2 on page 159), and also, by default, the Task Assistant.

### 6.1.1 Initial Setup Tasks

Initially, the Initial Setup Tasks window displays as shown in Figure 6-1. The Task Assistant gives you a quick way to access common tasks.

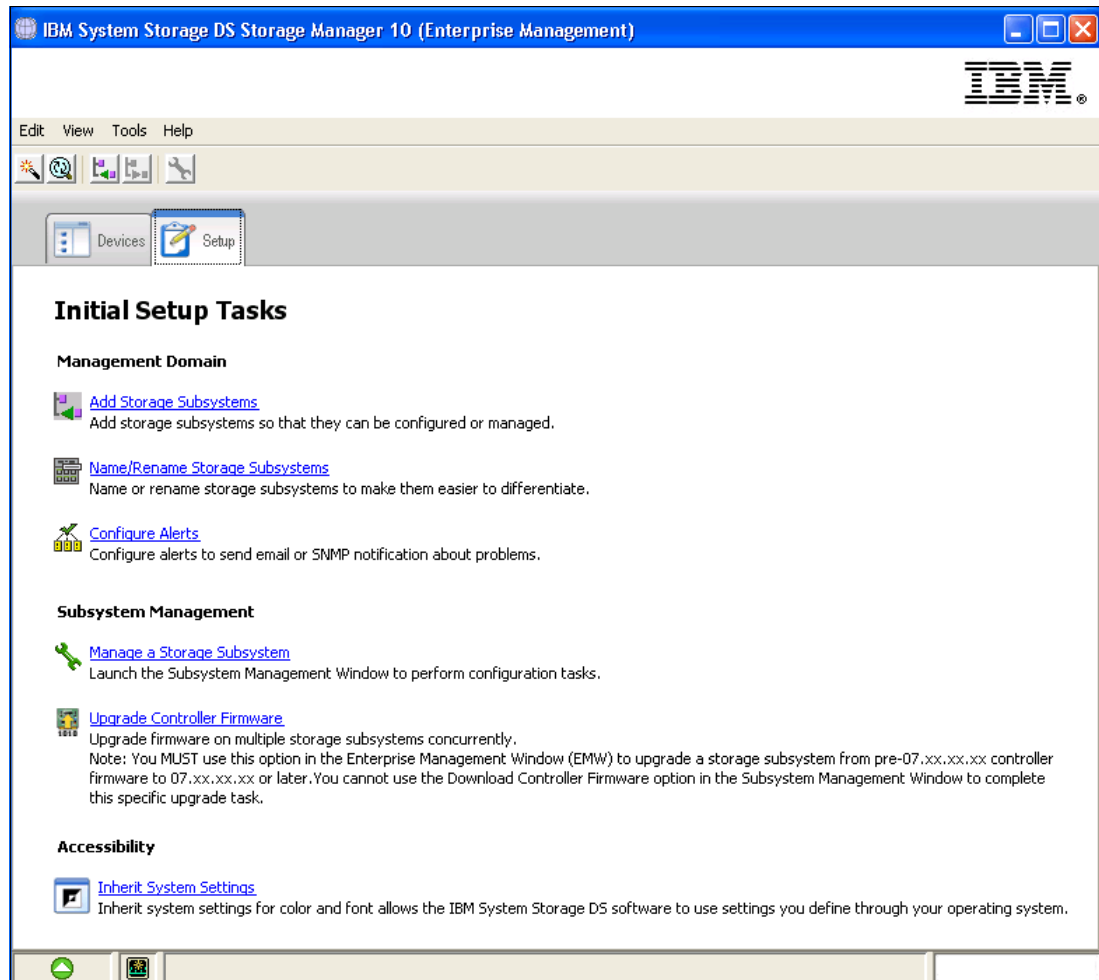


Figure 6-1 Setup Tasks

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 Setup Tasks, with information about where the use of these functions is described in this book:

- ▶ Initial Setup Tasks:
  - Add Storage Subsystems: See 5.2, “Installing DS Storage Manager on Linux” on page 140 or “Add Storage Subsystem to view” on page 173.
  - Name/Rename Storage Subsystems: See 12.3, “Rename Storage Subsystem” on page 299.
  - Configure Alerts: See “Configure Alerts” on page 164.



- ▶ Subsystem Management:
  - Manage a Storage Subsystem: Starts the main Storage Manager application, as described in the following chapter.
  - Upgrade Controller Firmware: Upgrades firmware on multiple storage subsystems, as described in Chapter 14, “Administration - Support” on page 329.
- ▶ Accessibility:
  - Inherit System Settings: Inherits system settings for color and fonts, which allows the IBM System Storage DS software to use settings defined through operating system

## 6.1.2 Enterprise Management window

The Enterprise Management window (Figure 6-2) is the entry point to manage each DS3500 storage subsystem. We described how to add storage subsystems in 5.4, “Completing the DS Storage Manager installation” on page 151. After they have been added, they will appear every time you start Storage Manager.

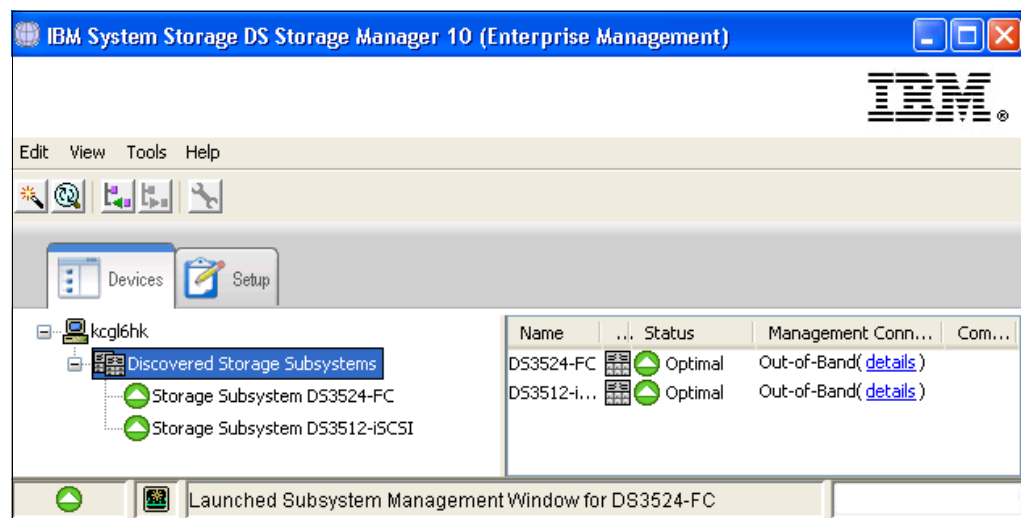


Figure 6-2 Enterprise Management window

The Enterprise Management window displays a list of all DS3500 storage subsystems that you 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 of this system in the right window pane. If the subsystem appears red in this view, then it needs attention. What you do in this case is described in Chapter 14, “Administration - Support” on page 329.

Figure 6-3 shows the various status icons for a storage subsystem. Table 6-1 explains the meaning of each icon.

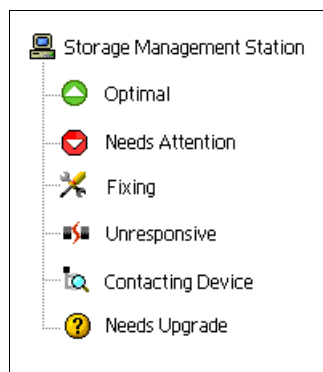


Figure 6-3 Status icons

Table 6-1 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 and 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.                 |

## 6.2 Functions in the Enterprise Management window

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

### 6.2.1 Subsystem context menu

Right-click one of the subsystems to see the context menu with the tasks as shown in Figure 6-4 on page 161.

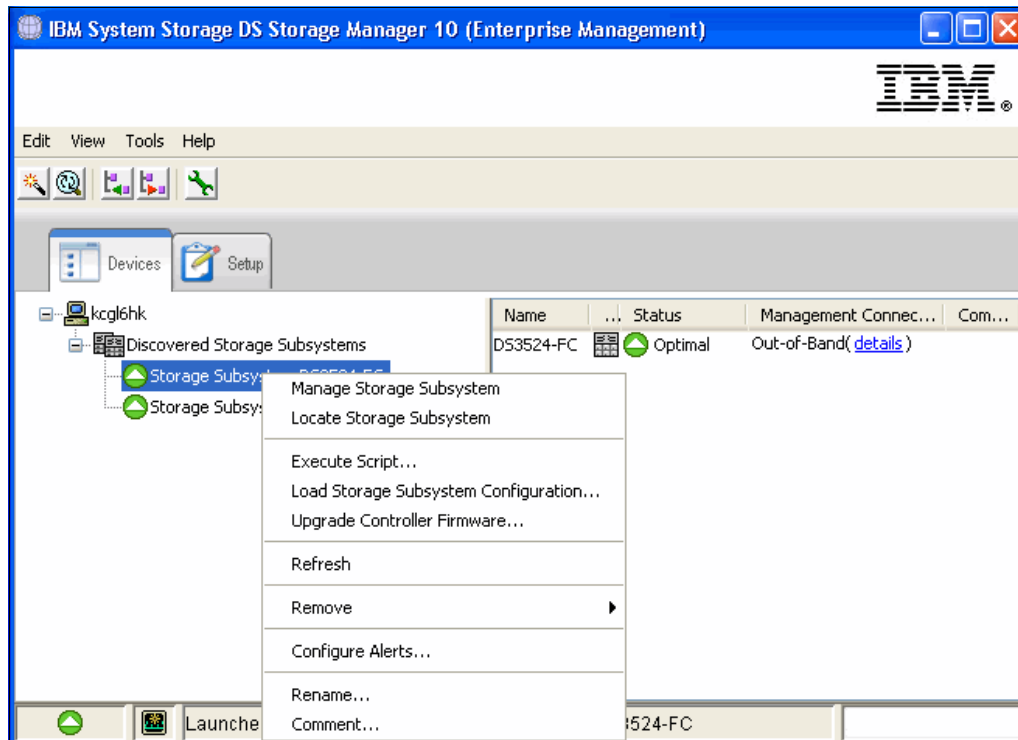


Figure 6-4 Context menu

## 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 6-5).

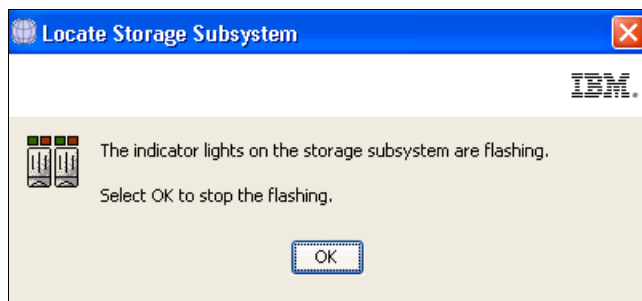


Figure 6-5 Indicator lights are flashing

## Execute Script

Here you can open the DS3500 Script Editor. This Editor can be an effective and useful tool to manage your subsystem. For more information about how to use it, see 17.2.1, “Script Editor” on page 521.

## Load Storage Subsystem Configuration option

In 17.7.1, “Sample command: Save configuration script file” on page 534, 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 storage subsystem 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. 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 lines shown in Figure 6-6 and uncomment them. After uncommenting these two lines, select **Tools** → **Verify** → **Execute** from the Menu bar to execute the script and upload the new configuration.
    - show "Deleting the existing configuration.";
    - set storagesubsystem resetConfiguration = true;

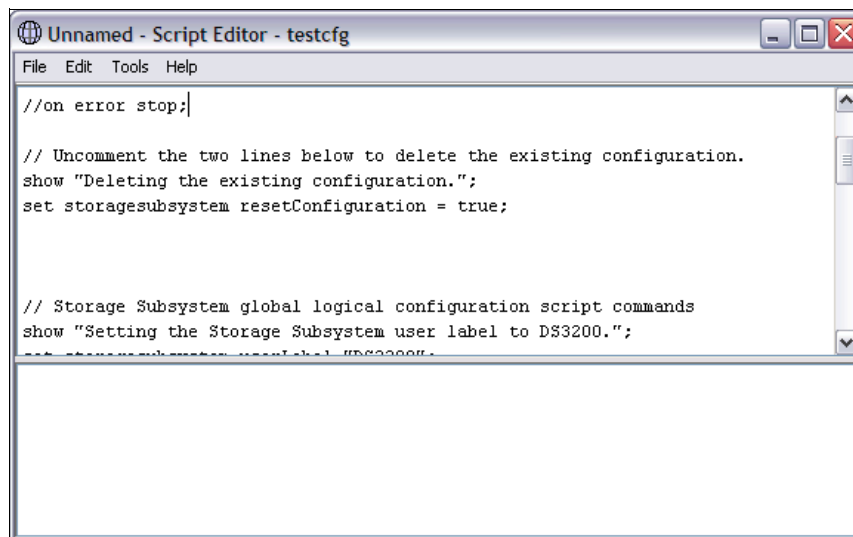


Figure 6-6 Uncomment lines to delete the existing configuration before uploading a new one

**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 because it will initialize the drives and delete the data. Instead, you should contact IBM technical support to see if your data can be recovered.

**Note:** To save time, use the load configuration file option for new DS3500s using a previously created file for an existing DS3500. 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 separate 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 DS3500s if your original hardware was destroyed in a disaster.

## Upgrade Controller Firmware

Select this option to upgrade the firmware of the controllers in the storage subsystem. In the process of upgrading the firmware, the firmware file is downloaded from the host to the controller. After downloading the firmware file, you can upgrade the controllers in the storage subsystem to the new firmware immediately. Optionally, you can download the firmware file to the controller and upgrade the firmware at a more convenient time.

The process of upgrading the firmware after downloading the firmware file is known as activation. During activation, the existing firmware file in the memory of the controller is replaced with the new firmware file. The firmware upgrade process requires that the controllers have enough free memory space in which the firmware file can reside until it is activated.

The Upgrade Controller Firmware window and the Download Firmware dialog are shown in Figure 6-7.

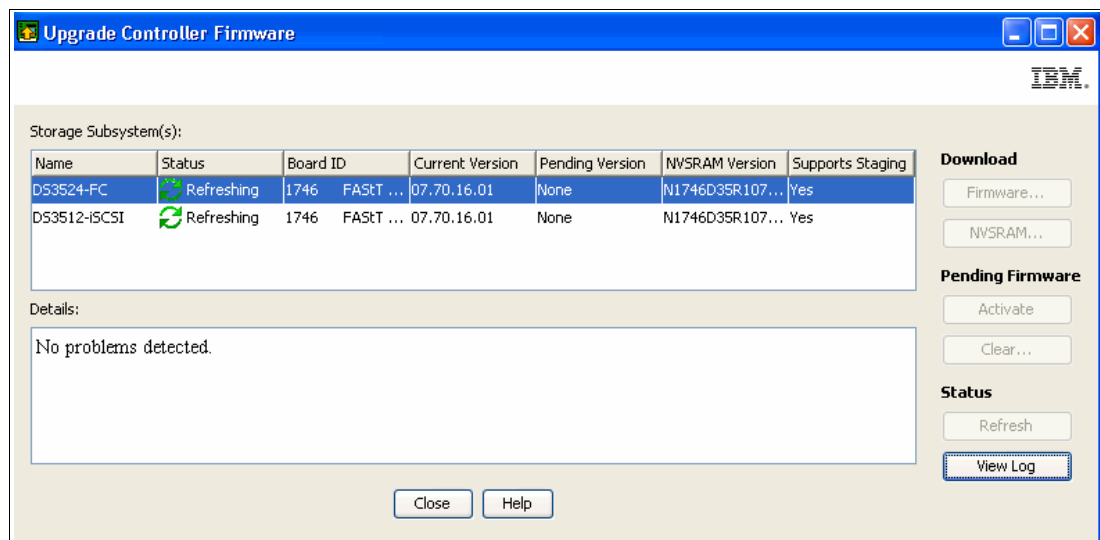


Figure 6-7 Upgrade Controller Firmware

You also can use the command line interface (CLI) to download and activate firmware to several storage subsystems.

## Refresh

Select this option to refresh the view of the Enterprise Management window (Figure 6-8). You might use this if a system's status has changed, for example, after an error recovery.

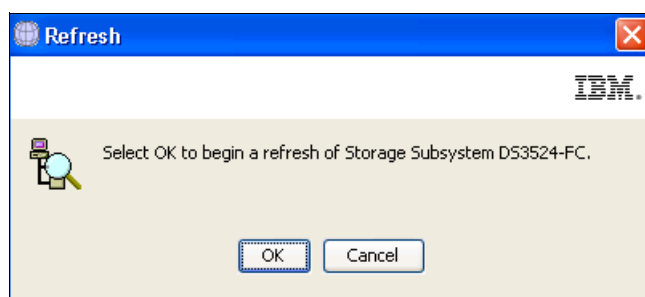


Figure 6-8 Refresh View

## 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 6-9).

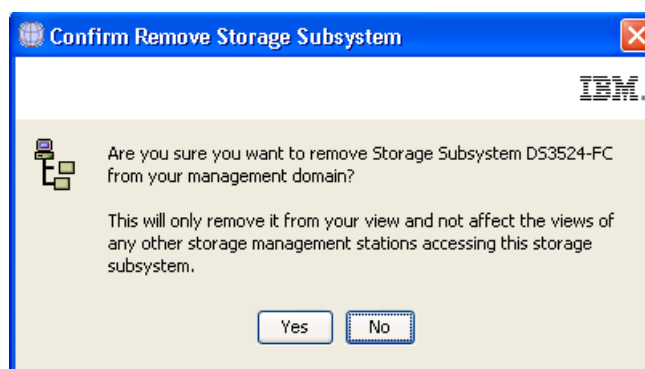


Figure 6-9 Remove Subsystem

## Configure Alerts

The IBM System Storage DS3500 Manager includes the Event Monitor Service, which enables the host running this monitor to send out alerts through email (SMTP) or traps (SNMP). The Event Monitor can be used to alert you of critical events for any of the DS3500 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 DS3500 does not send the email 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 alerts will be sent.

### ***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 select **Configure Alerts**, as shown in Figure 6-10.

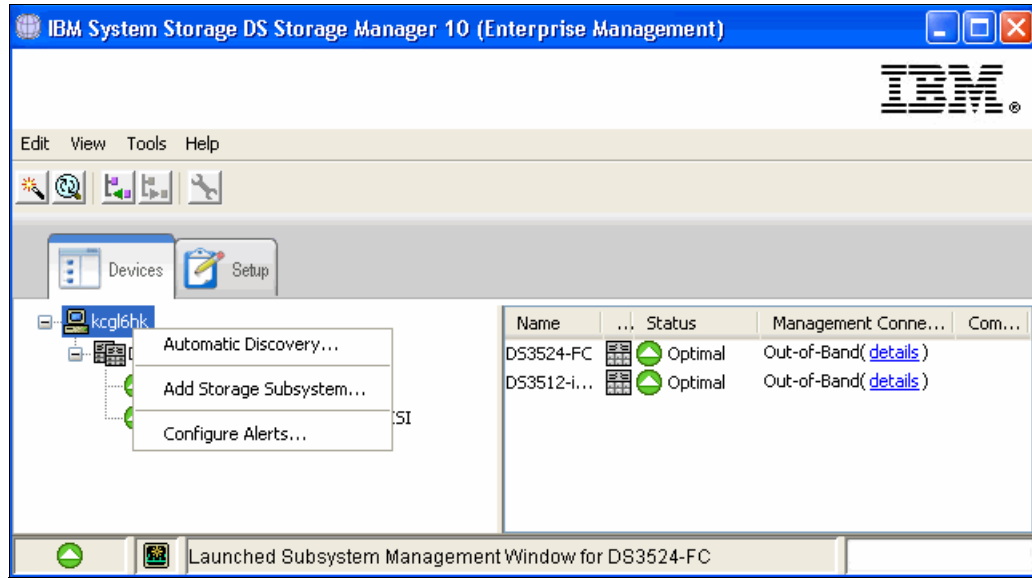


Figure 6-10 Configure Alerts for all systems

- If you can see the same storage subsystem through separate 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 6-11 on page 166.

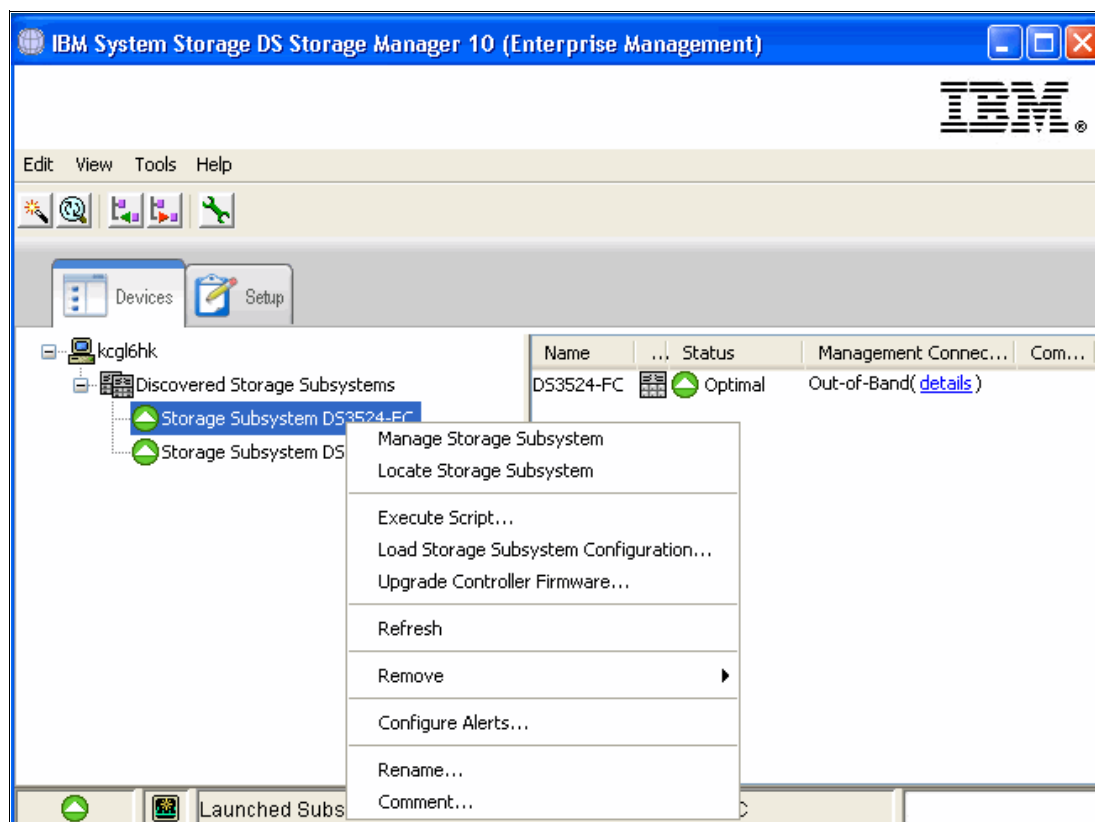


Figure 6-11 Configure Alerts for particular subsystem

- Select **Configure Alerts** for one or more systems and the alert window as shown in Figure 2 on page 167 appears.

### ***E-mail alerts***

Here we discuss how to work with email alerts:

1. To send email alerts, you have to first define an SMTP server. Enter the name or IP address of your mail server and the sender email address as shown in Figure 2 on page 167.



**Configure Alerts**

Alerts are generated for critical events only.

Mail Server | Email | SNMP

Mail server (SMTP server name, IPv4 address, or IPv6 address):  
amazon.rivers.local

Email sender address:  
DS3524-FC@rivers.local

☒ Include contact information with alerts:

Name:  
Title:  
Company:  
Phone:  
Cell phone:  
Pager:  
Email:  
Fax:  
Additional info:

OK Cancel Help

Figure 6-12 Configure Alerts - define mailserver

2. Select the **Email** tab to configure the email address to which the alerts are sent. Enter the email addresses and click **Add** as shown in Figure 6-13 on page 168.

- Now your window should look similar to Figure 6-13. 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 email to validate that your configuration is working.

Configure Alerts

Alerts are generated for critical events only.

Mail Server Email **SNMP**

Alerts for: DS3524-FC

Configured email addresses:

| Address                     | Information To Send | Frequency   |
|-----------------------------|---------------------|-------------|
| support@amazon.rivers.local | Event Only          | Every event |

Email address:

support@amazon.rivers.local

Add Replace Delete Test

OK Cancel Help

Figure 6-13 Configure Alerts - Defining email 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 6-14 on page 169.

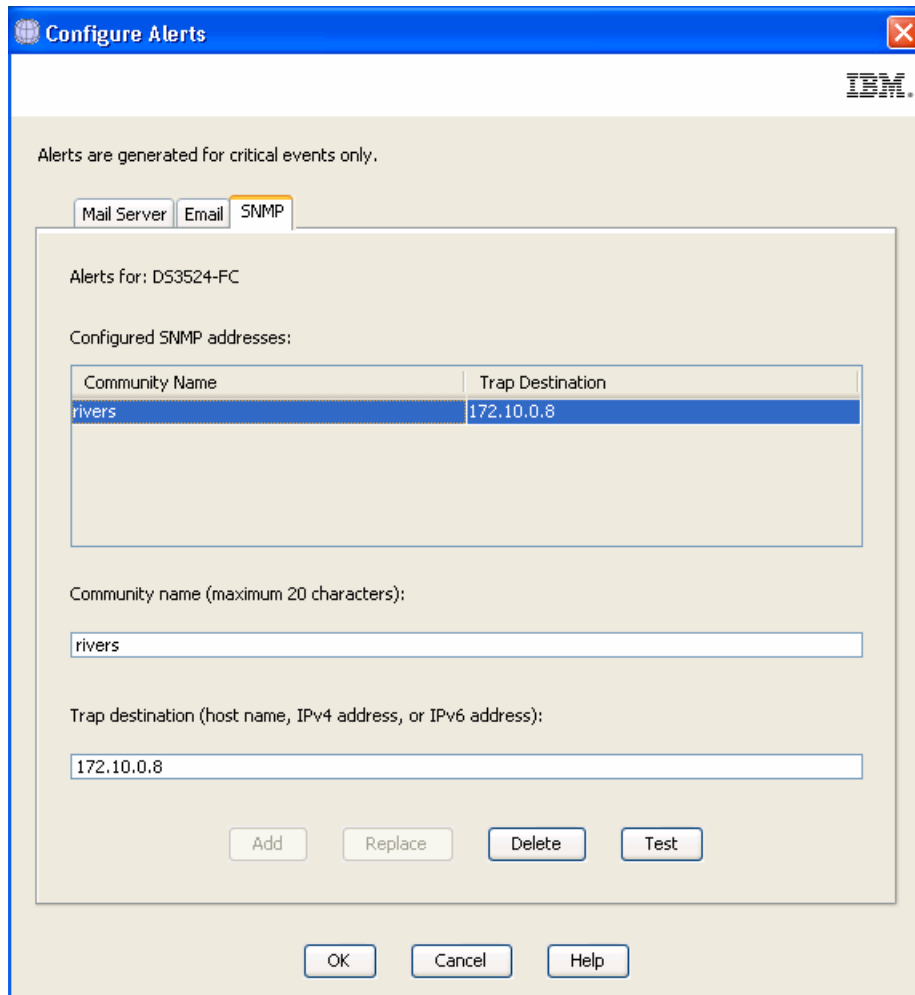


Figure 6-14 Configure Alerts - SNMP

**Note:** An SNMP manager is required to receive and translate your SNMP traps, for example, IBM Systems Director. For more information about IBM Systems Director and how to configure it, see *Implementing IBM Systems Director 6.1*, SG24-7694.

## Rename

If you have multiple DS3500s installed, give each one a user-specified name for easier identification. To rename a DS3500, select this option, enter a new name for the subsystem, and click **OK** as shown in Figure 6-15.

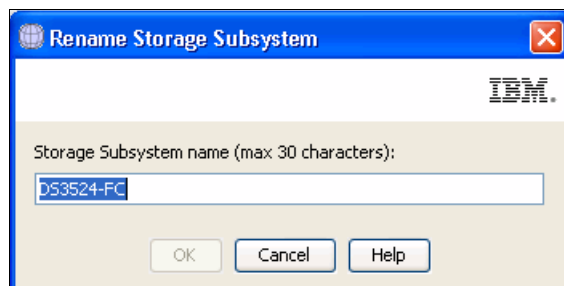


Figure 6-15 Rename Storage Subsystem

**Note:** Set the name only during the initial setup of the box and do not change it later because 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. Select this option, enter your comment, and click **OK** as shown in Figure 6-16.




Figure 6-16 Edit comment

Your comment will appear in the right pane of the Enterprise Management window as shown in Figure 6-17. Your comment could be a description of the disk subsystem.

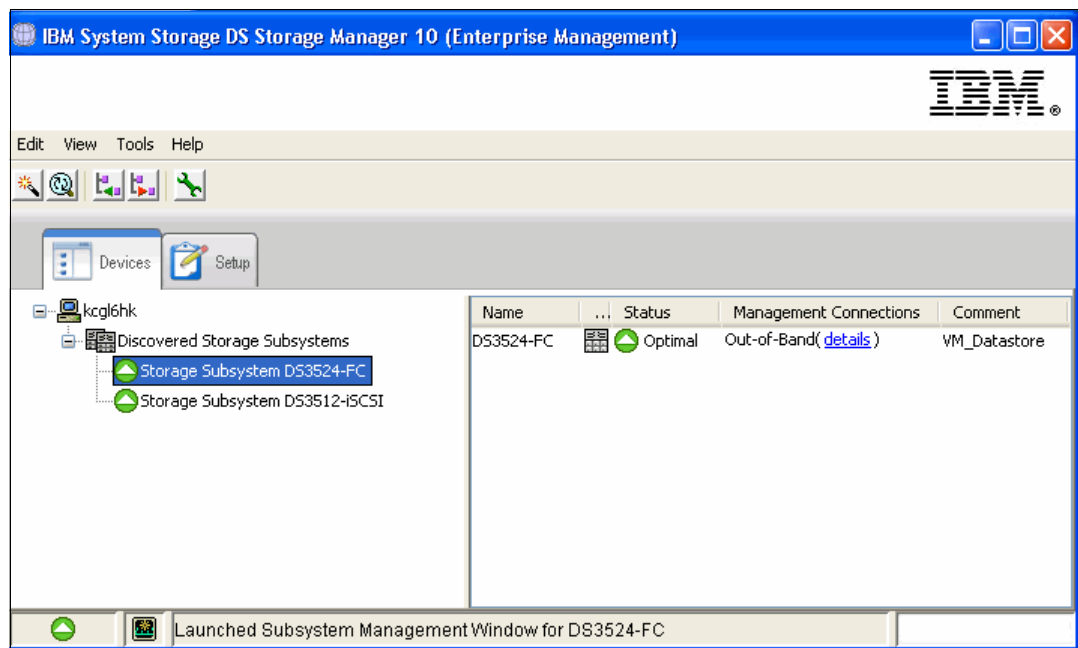


Figure 6-17 Add comment to Storage Subsystem

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

## 6.2.2 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 6.2.1, “Subsystem context menu” on page 160 can also be accessed from the task menu bar. To do this, highlight a subsystem and open the appropriate menu. The **Edit** (Figure 6-18) and **Tools** (Figure 6-19) submenus contain all of the tasks. Performing these tasks is done exactly as described in the previous section.

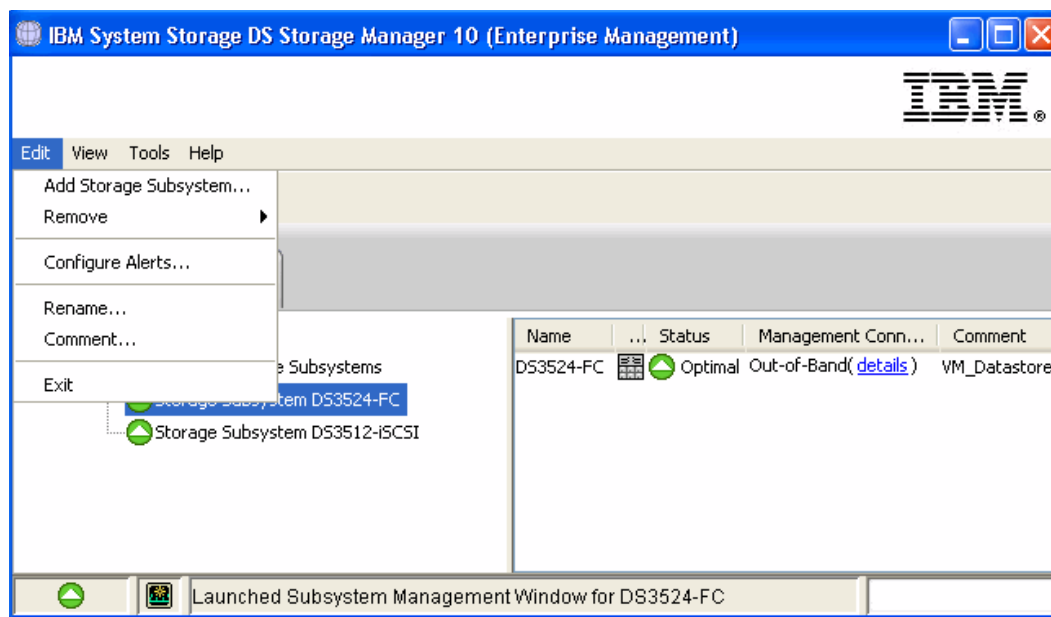


Figure 6-18 Edit menu

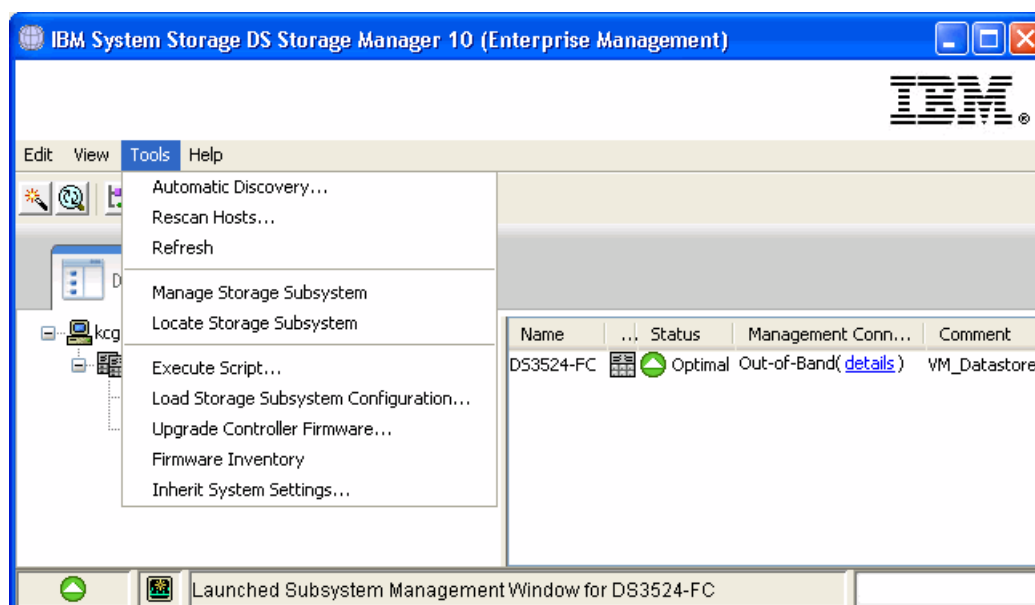


Figure 6-19 Tools menu

## The View submenu

Use this menu to start the Enterprise Management window Task Assistant (Figure 6-20 on page 172) or to customize the view of your enterprise window by selecting one of the sort options. The Task Assistant is described in 6.1.1, “Initial Setup Tasks” on page 158.

### 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. Certain maintenance tasks require a connection to both controllers.

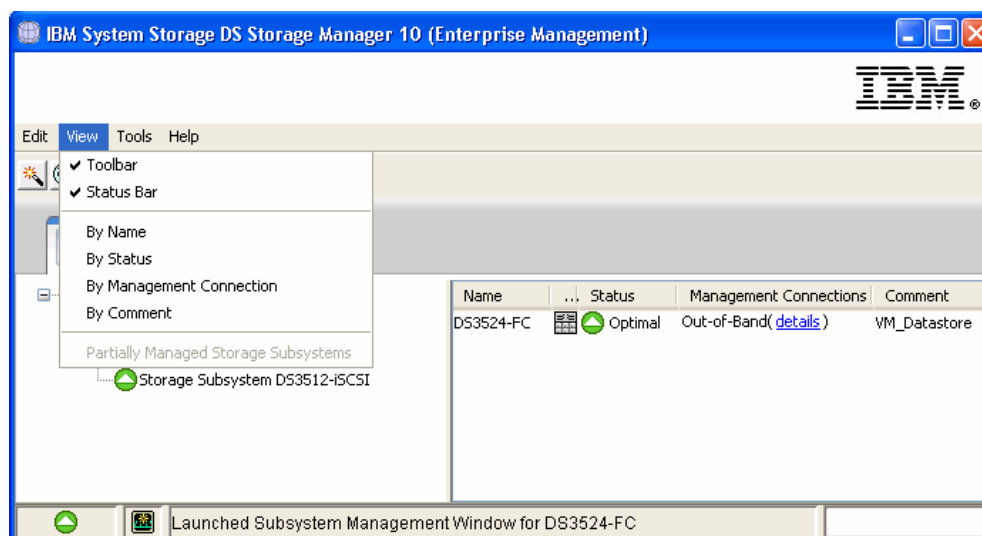


Figure 6-20 View submenu

## 6.2.3 The Quick Access buttons

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

### Automatically discover new storage subsystems



Figure 6-21 Automatically discover new storage systems button

This button starts an automatic discovery of new storage subsystems (Figure 6-22 on page 173) by sending in-band and out-of-band broadcasts. If it finds directly attached storage systems or hosts running the DS3500 Storage Manager agent (with an attached Storage subsystem), it adds these storage subsystems to the Enterprise Management window.

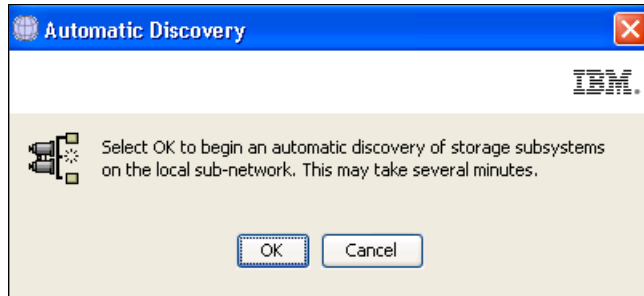


Figure 6-22 Automatic Discovery

**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 173.

## Rescan selected host for new storage subsystems



Figure 6-23 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. Click **OK** and the host will scan for newly attached storage devices.

## Add Storage Subsystem to view



Figure 6-24 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 6-25 on page 174). If the storage subsystem has dual controllers, you must provide the IP address of both controllers.

**Add New Storage Subsystem - Manual**

IBM

[What are in-band and out-of-band management connections?](#)

[Adding controllers with more than one Ethernet port](#)

[What if my system only has one controller?](#)

Select a management method:

☒ Out-of-band management:  
Manage the storage subsystem using the controller Ethernet connections.

Controller (DNS/Network name, IPv4 address, or IPv6 address):  
1.1.1.2

Controller (DNS/Network name, IPv4 address, or IPv6 address):  
1.1.1.3

☐ In-band management:  
Manage the storage subsystem through an attached host.

Host (DNS/Network name, IPv4 address, or IPv6 address):

Add Cancel Help

Figure 6-25 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 6-26).

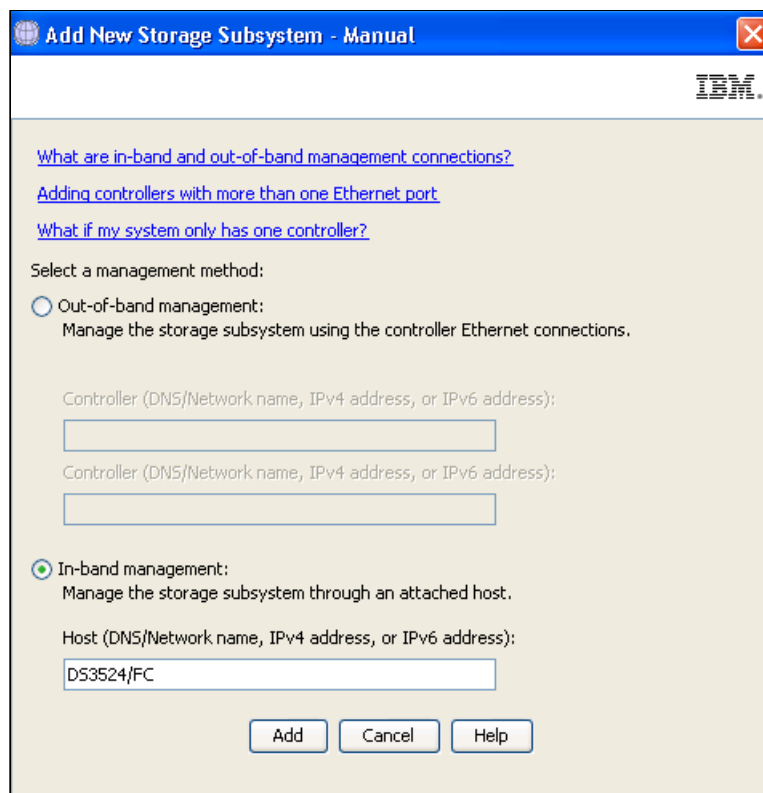


Figure 6-26 Manually add in-band system

## Remove Storage Subsystem from View



Figure 6-27 Remove Storage Subsystem from View

To remove a storage subsystem from your view, highlight it, click this button, and click **OK** to confirm (Figure 6-28 on page 176). This subsystem is only removed from your view: the views of other management stations are not affected.

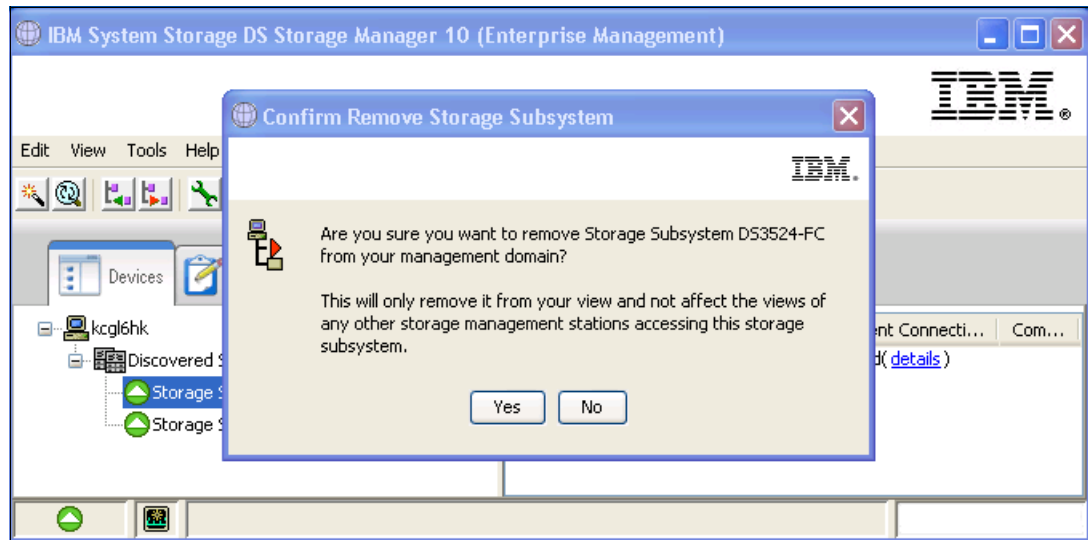



Figure 6-28 Remove Subsystem from view

## Launch Management window



Figure 6-29 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.



## Administration - Summary Tab

In this chapter, we give you an overview of the **Summary** tab of the IBM System Storage DS (Subsystem Management) window.

The following topics are discussed:

- ▶ How to monitor the basic status of your storage subsystem
- ▶ How to get a basic overview of your installed hardware components
- ▶ How to get a basic information about the current logical configuration

The first tab on the left side of Subsystem Manager in Tab area is the Summary tab. Figure 7-1 on page 178 shows a simply configured system that is running in an optimal status.

The Summary tab is divided into six boxes, which we describe in the following sections:

- ▶ 7.1, “Status” on page 178
- ▶ 7.2, “Hardware Components” on page 181
- ▶ 7.3, “Capacity” on page 181
- ▶ 7.4, “Hosts & Mappings” on page 182
- ▶ 7.5, “Arrays & Logical Drives” on page 184
- ▶ 7.6, “Information Center” on page 184

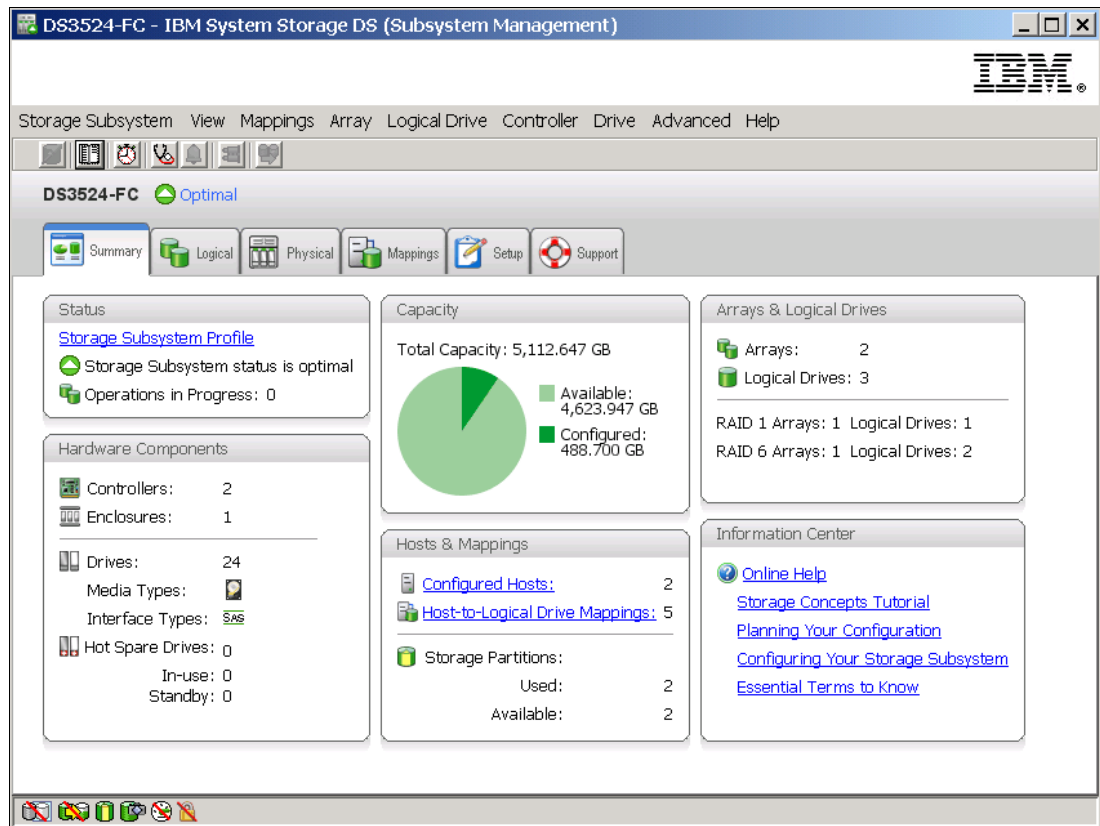


Figure 7-1 IBM System Storage DS Storage Manager Summary tab

## 7.1 Status

The first box is the **Status** box, which displays the basic status of your storage subsystem (Figure 7-2).




Figure 7-2 Status box

### 7.1.1 Storage Subsystem Profile

Configuring a storage server can be a complex task. It is essential to document the initial configuration and every later configuration change. The configuration documentation can be saved in a text file known as a subsystem profile. This profile stores information (properties) about the controllers, attached drives, enclosures, microcode levels, arrays, logical drives, storage partitioning and others.

To generate and open the subsystem profile, click the **Storage Subsystem Profile** link in the Status box (Figure 7-2 on page 178). The information will be gathered from the various components and the profile window will open (Figure 7-3). The various tabs show the properties of various objects in storage subsystem. Use the **Save As** button to save the whole profile locally to disk, to help document a change history of your storage subsystem. It will be saved as a text file. You can look at this file later, or send it to somebody who has no direct access to your storage subsystem to analyze it. For more information about the profile, refer to Chapter 14, “Administration - Support” on page 329.




Figure 7-3 Subsystem Profile window

**Note:** You should save a new version of the profile (or collect all support data (CASD) as described in Chapter 14, “Administration - Support” on page 329) and store it securely whenever the configuration changes. Even in the case of a complete configuration loss, your data on disks might not be affected, and in this case you can restore the arrays and logical drives configuration and the mappings for the storage partitioning. The profile should be stored locally, and also can be included in any offsite disaster recovery documentation.

### 7.1.2 Storage subsystem status

The second line in the **Status** box shows the current subsystem status. If the status is green as on Figure 7-2 on page 178, the subsystem is running in an optimal status. If there is any problem with the subsystem, the status icon will change to red (Figure 7-4). 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 329. The same red icon is also visible in upper part of the Subsystem Manger window, close to the name of subsystem, and the Recovery Guru icon in the toolbar is flashing.




Figure 7-4 Status box - Storage Subsystem Needs Attention

### 7.1.3 Operations in Progress

The third line shows operations that are currently in progress (Figure 7-5). If there is at least one operation in progress, you can click the link to display a window with more detailed information (Figure 7-6). You can find more information about operations in progress in Chapter 9, “Administration - Logical Tab” on page 219.




Figure 7-5 Status Box - Operations in Progress

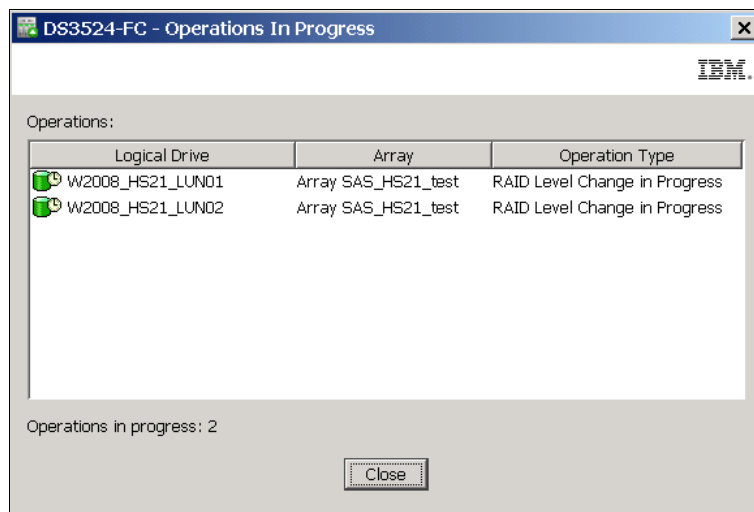


Figure 7-6 Detailed In Progress information

### 7.1.4 Connection lost

If the management connection from the management station to the storage subsystem is lost, the subsystem icon and text is changed to notify you of this communication error (Figure 7-7). 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 Management window will then display the last known state of the subsystem only, and you will not be able to perform any administrative operations.



Figure 7-7 Status box - connection lost

## 7.2 Hardware Components

Below the **Status** box is an overview of the **Hardware Components** (Figure 7-8) that are physically installed in this storage subsystem. The upper part of the box shows the number of installed controllers (two in this case) and the number of connected enclosures (one in this case). The lower part of the box shows information about installed hard drives, number of drives, the drive type, and an overview of the hot spare drives.




Figure 7-8 Hardware Components

To find out how to get more information about HW component, go to Chapter 10, “Administration - Physical Tab” on page 257.

## 7.3 Capacity

The **Capacity** box is located at the top of the middle column and is shown in Figure 7-9. It displays the overall capacity of your subsystem: that is, the total capacity of the installed disk drives. Darker green color is used to show already configured capacity in logical volumes, and light green color is for the capacity that can be used for configuration of new arrays and Logical Drives.




Figure 7-9 Capacity - Configured and unconfigured capacity

## 7.4 Hosts & Mappings

Below the **Capacity** box is an overview of your **Hosts & Mappings** as shown in Figure 7-10.




Figure 7-10 Hosts & Mappings

Hosts are single server 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, Fibre Channel, or iSCSI). If multiple hosts access one or a group of logical drives, they all would have equal access to these drives. Because most operating systems (without clustering software) do not usually allow shared access to the same logical drives, you must create storage partitions. Do this by mapping specific logical drives to the host ports of the host systems. Configuring host access is described in details in Chapter 11, “Administration - Mappings Tab” on page 275.

### 7.4.1 Configured Hosts

This line simply shows you how many hosts are currently configured on this storage subsystem. Click the link to display the configured host names (Figure 7-11). This view is only an overview of the configured hosts. The storage subsystem only communicates with hosts by the use of Word Wide Port Name (for example: 50:00:62:b0:00:09:40:51) or by iSCSI Initiators name (for example: iqn.1996-04.de.suse:01:27dc9458a936). The storage subsystem is not aware of which Host Bus Adapters or iSCSI Initiators are in the same server or cluster unless you configure it. More information about host ports configuration is covered in Chapter 11, “Administration - Mappings Tab” on page 275.

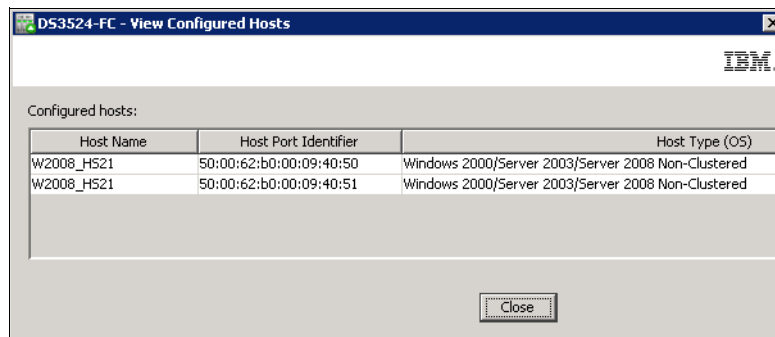


Figure 7-11 Configured Host window



## 7.4.2 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 logical drives that are already mapped. This means the number of logical drives that are allowed to be seen by a specific host. Opening the link displays a more detailed view (Figure 7-12).

In this window, 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 window also includes the LUN Number, the Logical Drive Capacity, and the Type. More information about drive mappings is provided in Chapter 11, “Administration - Mappings Tab” on page 275.




Figure 7-12 Hosts & Mappings - Host-to-Logical Drive Mappings

## 7.4.3 Storage partitions

The lower part of the box in Figure 7-10 on page 182 displays information about the number of storage partitions that are used and available on this storage subsystem. Storage partitions are enabled by a Premium Feature License, 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 when the number of your servers will increase.

All DS3500 storage subsystems have four storage partitions as standard. It can be optionally upgraded to 8, 16, 32, or 64 storage partitions. More information about storage partitions is provided in 3.3.7, “Storage partitioning” on page 63

The **Storage Partition** part of the box shows how many storage partitions are already used on your subsystem and how many partitions is available to use. Sum of these two numbers shows how many storage partitions is the DS3500 subsystem licensed for.

## 7.5 Arrays & Logical Drives

At the top of the right column of the Summary tab is the Arrays & Logical Drives pane. This box shows how many arrays and logical drives are configured on your system (Figure 7-13). An array is a set of physical disk 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 level. For more details about the RAID levels, refer to 3.3.2, "Understanding RAID types" on page 50. Depending on your configuration, this affects the system's fault tolerance and performance. The bottom part of the box shows which RAID levels are configured on your storage subsystems, how many arrays of this level exist, and how many logical drives are in each level.

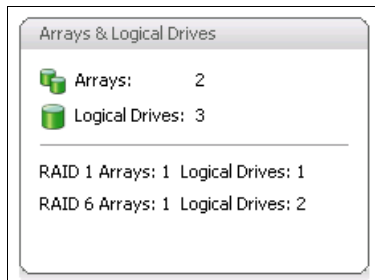


Figure 7-13 Arrays & Logical Drives

**Note:** A RAID1 level array on more than two disks is equivalent to RAID10. However, in this case you will see RAID1 displayed in this box.

## 7.6 Information Center

At the right bottom edge of the Summary tab is the **Information Center** box (Figure 7-14). It includes several links to help you understand and set up your storage subsystem, and answer any questions you can have.




Figure 7-14 Information Center

Each of the five links opens the same Help application window but with different content. Click one of the links to navigate in the online help directly to the related content:

- ▶ Online Help
- ▶ Storage Concepts Tutorial
- ▶ Planning your Configuration
- ▶ Configuring Your Storage Subsystem
- ▶ Essential Terms to Know

As shown in Figure 7-15, topics on how to configure DS3500 storage subsystem display if the **Storage Concepts Tutorial** link is selected.

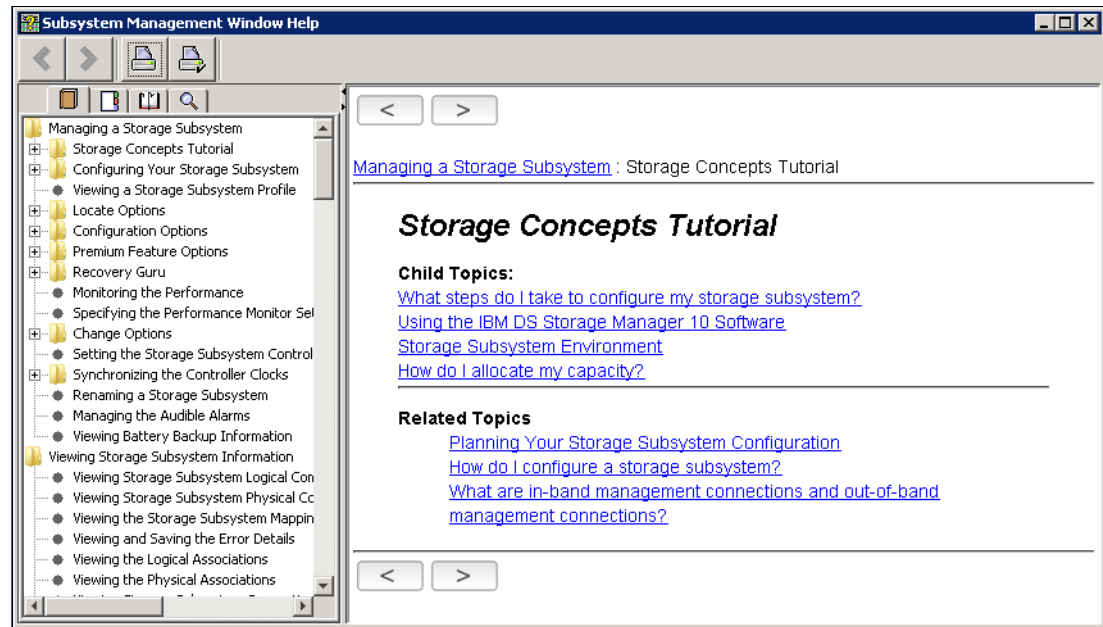


Figure 7-15 Information Center - Storage Concepts Tutorial





## Administration - Subsystem Management

Chapter 6, “Administration - Enterprise Management” on page 157 describes how to start managing a storage subsystem from the **IBM System Storage DS Storage Manager 10 (Enterprise Management)** window.

When you choose to manage a specific storage subsystem, the **IBM System Storage DS (Subsystem Management)** window opens, along with the Initial Setup Task on background and a small window prompting for the password (Figure 8-1 on page 188).

**Note:** Starting with DS Storage Manger version 10.70, all managed subsystems must have a password set. Setting the password and a description of the new strong password rules are described in “Set Password” on page 204

After entering the correct password, all six tabs of Subsystem Manager are accessible. The last used tab from previous session is selected automatically and the size of window is also restored from previous session. See Chapter 12, “Administration - Setup tab” on page 297 for the Initial Setup Tasks for the first configuration steps on a new subsystem.

## 8.1 DS Storage Manager - Subsystem Manger window

The IBM System Storage DS Storage Manager Subsystem Management window has four main areas as shown in Figure 8-1.

We cover these areas in four sections:

- ▶ 8.2, “Pull-Down Menu” on page 189
- ▶ 8.3, “Toolbar” on page 212
- ▶ 8.4, “Status bar” on page 215
- ▶ 8.5, “Tabs” on page 216

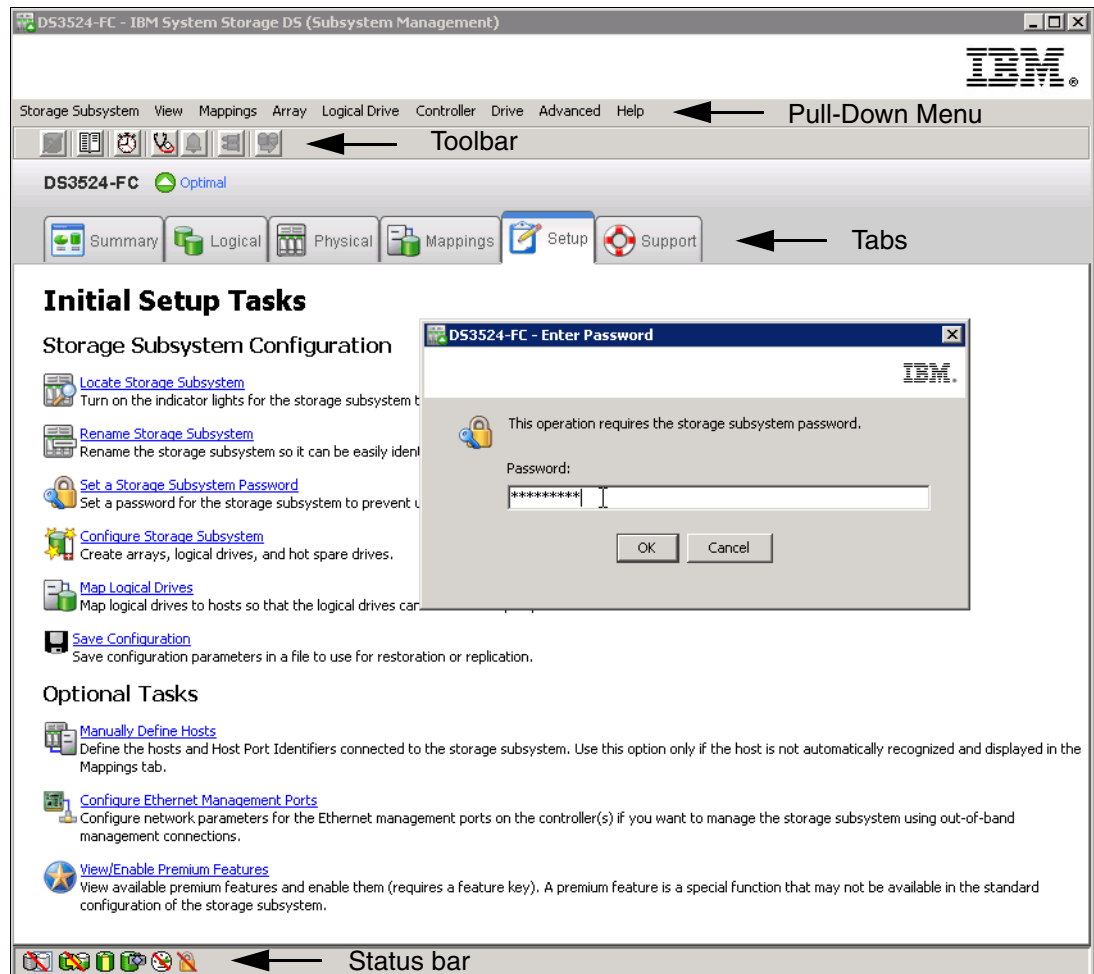


Figure 8-1 Subsystem Manger window and Enter Password window

## 8.2 Pull-Down Menu

The area at the top of Subsystem Management window is the pull-down menu bar. Although it is possible to do almost all configuration tasks from this menu, sometimes it is necessary to select the logical or physical components in Tab area first before you can perform the desired task. Another option is to configure the storage subsystem directly in the Tab area. We describe configuration tasks that are the easiest to execute from the pull-down menu in this section. Other configuration tasks are described in chapters dedicated to a specific Tab in sections.

In this section we cover the following pull-down menus:

- ▶ 8.2.1, “Storage Subsystem Menu” on page 189
- ▶ 8.2.2, “View menu” on page 205
- ▶ 8.2.3, “Mappings menu” on page 207
- ▶ 8.2.4, “Array menu” on page 208
- ▶ 8.2.5, “Logical Drive menu” on page 208
- ▶ 8.2.6, “Controller menu” on page 208
- ▶ 8.2.7, “Drive menu” on page 209
- ▶ 8.2.8, “Advanced menu” on page 209
- ▶ 8.2.9, “Help menu” on page 209

### 8.2.1 Storage Subsystem Menu

This menu is used to display or configure the general properties and features of the managed storage subsystem. It is used often and is independent of the tab that is currently selected in the Tab area.

Many items can be seen in this menu:

- ▶ “View” on page 190
- ▶ “Locate” on page 192
- ▶ “Configuration” on page 192
- ▶ “Premium Features” on page 194
- ▶ “Remote Logical Drive Mirroring” on page 197
- ▶ “iSCSI” on page 198
- ▶ “Drive Security” on page 198
- ▶ “Recovery Guru” on page 198
- ▶ “Monitor Performance” on page 198
- ▶ “Change” on page 198
- ▶ “Synchronize Controller Clocks” on page 203
- ▶ “Manage Enclosure Alarm” on page 204
- ▶ “Rename Subsystem” on page 204
- ▶ “Set Password” on page 204
- ▶ “Exit” - to close the Subsystem Manager window

Click **Storage Subsystem** to open the associated menu items as shown in Figure 8-2 on page 190. If you select items with the small right arrow on the right side of the menu item, it opens another sub-menu. The other items (with three dot characters on the end of menu text) directly open a window with requested task.




Figure 8-2 Storage Subsystem Menu

Here we describe the most common items of this pull-down menu. The others (iSCSI, Drive Security, Recovery Guru, and Monitor Performance) advanced options are covered in separate chapters of this book.

## View

Click **Storage Subsystem** → **View**. This menu item offers several option for configuration view and detailed information about the storage subsystem.

### **View** → **Profile**

This opens a window with a Storage Subsystem Profile. More detailed explanation of subsystem profile content is in Chapter 14, “Administration - Support” on page 329.

**Note:** All the information and views offered by Storage Manager originate from the subsystem profile. Usually all the component properties windows in Storage Manager just show information from the subsystem profile more comfortably.

### **View** → **Enclosure Components** → **Controller Enclosure**

Because the DS3500 only has Controller/Drive Enclosures and Drive Enclosures, this item is not available for DS3500.

### **View** → **Enclosure Components** → **Controller/Drive Enclosure**

This shows properties of the DS3500 Controller Enclosure (Figure 8-3 on page 191). All displayed information is taken from the subsystem profile. You can also change Enclosure ID here or locate enclosure in the rack by blinking its blue LED diodes on the left side of the DS3500 enclosure.



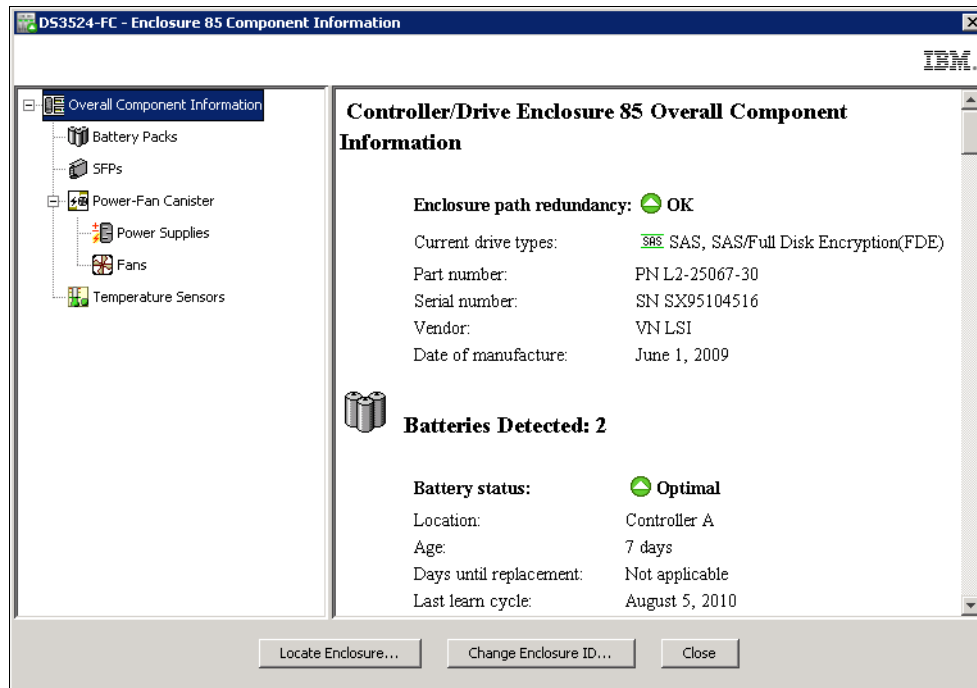


Figure 8-3 View Controller/Drive Enclosure

### **View → Enclosure Components → Drive Enclosure**

Almost the same type of information can also be displayed for EXP3500 Drive Enclosures connected to the managed DS3500 subsystem.

### **View → Connections**

This opens a window with information about the attachment of expansion enclosures (their ESM boards) to controllers of storage subsystem as shown in Figure 8-4.

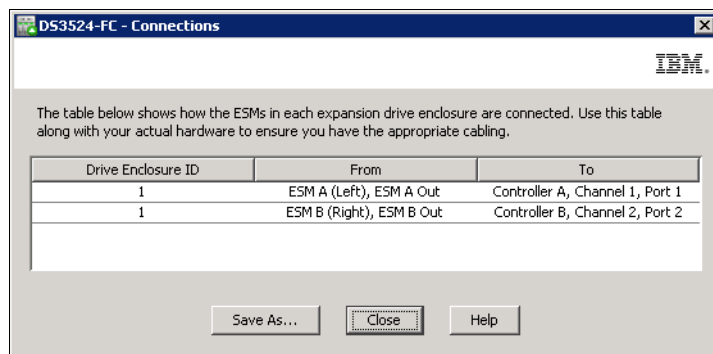


Figure 8-4 Connections of ESM

### **View → Unassociated Host Port Identifiers**

It shows the list of ports (HBA) or iSCSI Initiators on servers that are available, and can be attached to DS3500 subsystem (Figure 8-5 on page 192). It shows only port identifiers (WWPN or iSCSI Initiators) and the type of interface: it does not show names of hosts. Association can be done in Mapping Tab of Subsystem Manager as described in 11.2, “Defining Host” on page 277. You can define a name of the host and assign host ports.

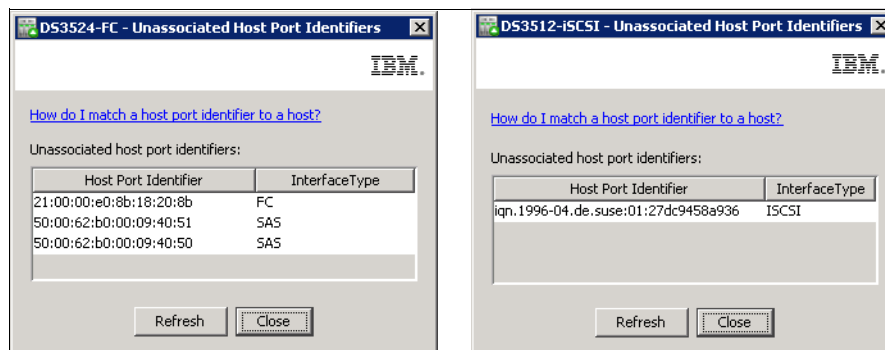


Figure 8-5 Unassociated Ports - FC and SAS HBA's and iSCSI Initiator

## Locate

This menu item helps you to find the DS3500 subsystem in rack. When **Storage Subsystem → Locate → Storage Subsystem** is activated, blue LED diodes on the controller enclosure and on all expansion enclosures attached to this DS3500 subsystem start to flash.

## Configuration

Click **Storage Subsystem → Configuration**. This menu item offers two tasks. You can start the wizard for **Automatic** configuration of available unconfigured disk drives in managed subsystem, and **Save** the current logical configuration of storage subsystem in form of a script file.

### Automatic Configuration

This option creates arrays and logical drives on all unconfigured disk drives of subsystem. Generally, do not use this fully automatic configuration of DS3500 (by selecting the **Choose a suggested configuration** option) because the results are not optimal for most cases. A slightly better result can be achieved if you select semi-automatic mode (**Create your own configuration** option) as shown in Figure 8-6 on page 193. Here you will have a choice to set RAID level, number of arrays, number of disk drive per array, and number of logical drives. You will still have no control over the size of the logical drive and selection of disks for array because this is set by subsystem automatically. Even in this semi-automatic mode, the results still can be slightly different than what you expect. For most environments, configure DS3500 subsystem using manual configuration in the Logical tab as documented in Chapter 9, "Administration - Logical Tab" on page 219. Using this option you can set all parameters for created arrays and logical drives there.

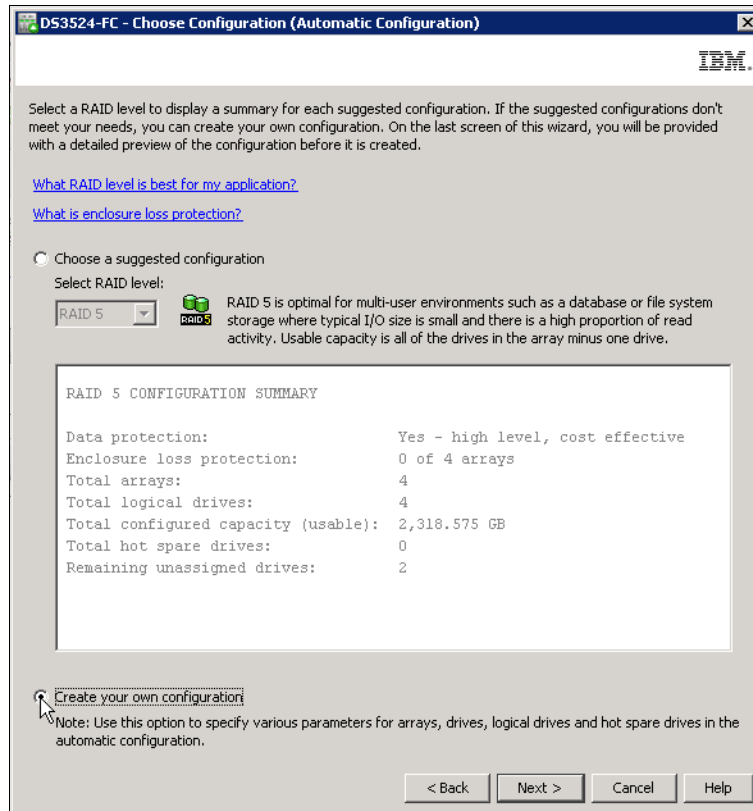


Figure 8-6 Automatic configuration

Detailed description of automatic configuration of unconfigured disk drives on DS3500 Storage Subsystem is in 12.6.1, "Automatic configuration" on page 302.

### Save Configuration

Use the **Storage Subsystem** → **Configuration** → **Save Configuration** option (Figure 8-7 on page 194) to save the logical configuration settings of a storage subsystem to a script file.

Later, you can select **Tools** → **Load Storage Subsystem Configuration** in the **Enterprise Management** window to copy the configuration data back from the saved file to a storage subsystem with the same hardware configuration. For more information about the **Load Storage Subsystem Configuration** option, refer to Chapter 6, "Administration - Enterprise Management" on page 157.

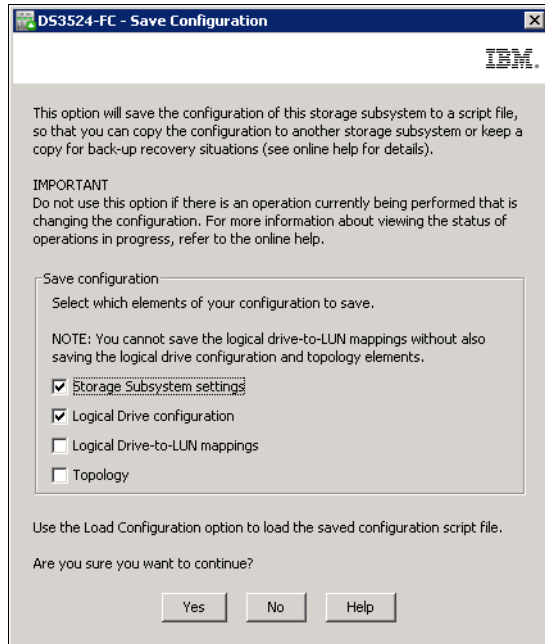


Figure 8-7 Save Configuration

The **Configuration** → **Save** option lets you save storage subsystem settings, logical drive configuration, topology, and mappings to the text file in form of a script file that can also be used by SMcli command line interface. If there are individual topology items or logical drives that you do not want to include in the saved configuration, you can edit the script file after it is saved. You can also modify certain SMcli commands from this file and use them in another script that can create a new logical devices in your storage subsystem. For more information about SMcli, see Chapter 17, “Command-Line Interface (CLI)” on page 519.

The **Configuration** → **Save** option does not save these settings:

- ▶ The age of the battery
- ▶ The controller time of day
- ▶ The nonvolatile static random access memory (NVS RAM) settings
- ▶ Any premium features
- ▶ The storage subsystem password
- ▶ The operating status and states of the hardware components
- ▶ The operating status (except Optimal) and states of the arrays

## Premium Features

This is the task used for viewing or activation of software licenses - Premium Features. You might have some of them already activated because certain basic licenses are delivered with each DS3500 as a standard, and you can buy more as needed. Types of licenses are described in 3.4, “Planning for premium features” on page 70.

### Obtaining the license \*.key file

You need a special license \*.key file to activate additional Premium Features in this step. This file must be downloaded from the IBM feature activation web site (Figure 8-8 on page 195), where each Premium Feature is registered based on the serial number of your DS3500 enclosure and Feature Enable Identifier.

The IBM activation website is located at the following URL:

<http://www-912.ibm.com/PremiumFeatures/>

|                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IBM Systems support           | IBM Systems > Support >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| BladeCenter                   | <b>IBM DS3000, DS4000, DS5000 and BladeCenter Boot Disk System activation and registration</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Power                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| System i                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| System p                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| System x                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| System z                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| System Storage                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Support search                | This website is intended for owners of IBM System Storage DS3000, DS4000, DS5000, FASTT and BladeCenter Boot Disk System storage controller units who have purchased one or more Premium Features or Software Feature Packs. To complete your purchase you must have received your Premium Feature or Software Feature Pack web activation instruction card from IBM.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Register                      | Effective February 12, 2008, the DS3200, DS3300 and DS3400 Software Feature Pack and EXP3000 expansion license will be included in base controller firmware for these products at no additional cost to users. To utilize the inclusion of these features, users must be at controller firmware version 6.70.24.00 or higher.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Feedback                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Systems networking            | <b>Please select from the following options:</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| System Blue Gene              | <ul style="list-style-type: none"> <li>• <a href="#">Activate a Premium Feature or Software Feature Pack</a><br/>Download a key to assign to a controller. A Premium Feature or Software Feature Pack web activation instruction card is required for this transaction. If you do not have a web activation instruction card, please contact your local IBM support.</li> <li>• <a href="#">Reactivate Premium Feature(s) or Software Feature Pack(s)</a><br/>Download all keys that have been previously assigned to a controller. Premium Feature or Software Feature Pack web activation instruction cards are not required for this transaction.</li> <li>• <a href="#">Register an expansion module</a><br/>Register an expansion module which establishes assignment to a controller used on some DS3000, DS4000 and DS5000 products. A Premium Feature web activation instruction card is required for this transaction.</li> </ul> |
| IntelliStation Pro            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| IBM Monitors                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Systems Management software   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Hardware options and upgrades |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

Figure 8-8 Premium Feature activation web site

The following Premium Features can be delivered with your DS3500 Storage Subsystem or you can buy these as optional (additional charge) Premium Features:

- ▶ More than four Storage Partitions
- ▶ More than two FlashCopies
- ▶ VolumeCopy
- ▶ Enhanced Remote Mirroring
- ▶ Encryption
- ▶ High Performance Tier

You can find instructions and the actual procedure to get the license \*.key file in the Premium Feature Certificate, which is included in the license envelope and is part of your DS3500 Premium Feature delivery. You can also find a unique access code for the IBM activation website where the \*.key file must be registered, generated, and downloaded.

**Warning:** Store the paper 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 when you 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) from the same IBM activation website because it was already registered and generated. You will only need the serial number of the DS3500.

### View/Enable Premium Feature

Select **Storage Subsystem** → **Premium Feature** in pull-down menu to open a window where the list of possible features is displayed in the first pane as shown in Figure 8-9 on page 196. Some features are enabled at the plant, whereas some are disabled and can be enabled/activated by a \*.key file that has to be downloaded from IBM web site (Figure 8-8).

The Feature Enable Identifier displayed at the bottom of this window is necessary (together with DS3500 enclosure serial number) during \*.key file generation at IBM activation web site.

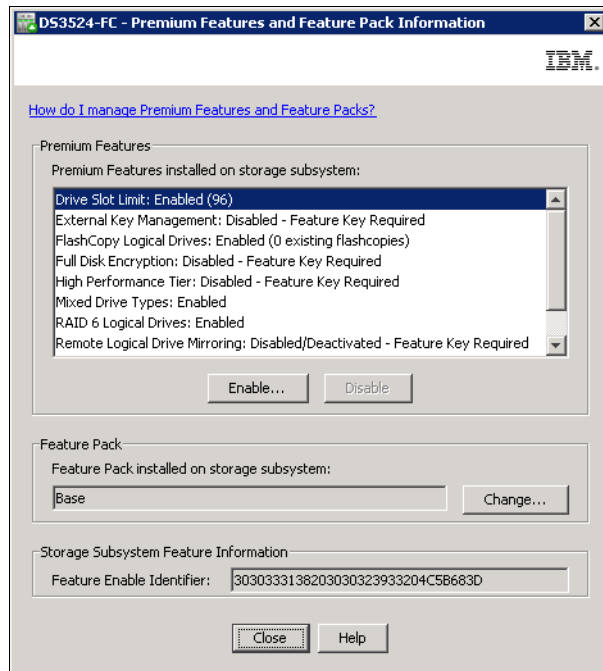


Figure 8-9 Premium Feature window

Depending on the selected Premium Feature, you can **Enable** a disabled feature, or **Disable** an enabled feature. To enable a feature, select **Enable** button and a window for selecting the \*.key file on your management station opens. When key file is accepted, a window with additional information as in Figure 8-10 appears.




Figure 8-10 Premium Feature Enable - information

To disable a feature, select that feature and click **Disable**. A warning window displays as shown in Figure 8-11 on page 197.




Figure 8-11 Disable Premium Feature

### Feature Pack

The second pane in the Premium Feature window (Figure 8-9 on page 196) is Feature Pack. It is not possible to change the Feature Pack on any DS3500 subsystem.

**Note:** Starting February 2008, the Feature Pack has been incorporated into all base DS3000 and DS3500 models, and is no longer needed as a prerequisite for the additional Premium Feature activation.

### Remote Logical Drive Mirroring

By selecting the **Storage Subsystem** → **Remote Logical Drive Mirroring** option, you can **Activate** or **Deactivate** this functionality. The Remote Logical Drive Mirroring Premium Feature must already be enabled (see “Premium Features” on page 194) to activate mirroring in this step. Only DS3500 with optional FC daughter cards can be used for remote mirroring because mirrored subsystems can mirror only on FC links. Activation of remote mirroring will create two special Repository logical drives dedicated to internal use of remote mirroring. The “Free capacity” on an existing array or the “Unconfigured capacity” must be selected in the window shown in Figure 8-12. The mirroring repository logical drives will be placed there.

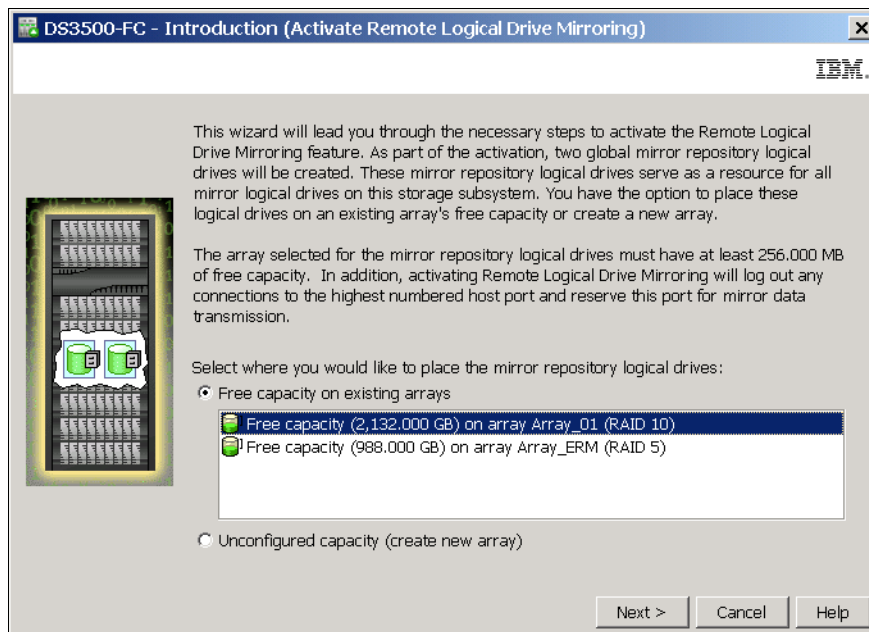


Figure 8-12 Activation of Remote Mirroring

The array containing the mirror repositories must not be deleted until the remote mirroring is deactivated again. These two repositories are automatically deleted during deactivation of Remote Logical Drive Mirroring.




Figure 8-13 Mirror Repositories

All details about remote mirroring are fully described in *IBM Midrange System Storage Copy Services Guide*, SG24-7822.

## iSCSI

This option in the **Storage Subsystem** pull-down menu is available and visible only if iSCSI daughter cards are installed in DS3500 subsystem. For a detailed description of how to set attachment to iSCSI hosts, see Chapter 13, “Administration - iSCSI” on page 309.

## Drive Security

This option in the **Storage Subsystem** pull-down menu is available and visible only if Full Disk Encryption premium feature is enabled (see “Premium Features” on page 194) and the SED disk drives are installed. Details on how to configure encrypted disk drives are covered in Chapter 15, “Disk Security with Full Disk Encryption drives” on page 439.

## Recovery Guru

The **Recovery Guru** option in the **Storage Subsystem** pull-down menu opens the same window just like if you click the **Recovery Guru** icon in the Toolbar area. More details about Recovery Guru can be found in Chapter 14, “Administration - Support” on page 329.

## Monitor Performance

The **Monitor Performance** option in the **Storage Subsystem** pull-down menu opens the same window just like if you click the **Monitor Performance** icon in Toolbar area. More details about Monitor Performance can be found in Chapter 14, “Administration - Support” on page 329.

## Change

The **Change** menu item in **Storage Subsystem** menu is one of the most important for setting global parameters of the DS3500 storage subsystem. Many changes done by this menu can also be done directly in tabs, but we describe most of them in this section.

Click **Storage Subsystem** → **Change**. The menu as shown in Figure 8-14 on page 199 offers several options to configure:

- ▶ Battery Settings
- ▶ Cache Settings
- ▶ Default Host Type
- ▶ Failover Alert Delay
- ▶ Time-out Setting
- ▶ Enclosure ID
- ▶ Enclosure Order






Figure 8-14 Change subsystem menu

### Battery Settings

The DS3500 storage subsystem batteries do not maintain expiration dates. Batteries can be replaced only when the LEDs indicates that they are failed, or Recovery Guru recommends you do so. The battery is periodically checked by the DS3500 controller firmware every eight weeks. You have only an option (see Figure 8-15) to set the time and day of week when the check (learn cycle) will start.

Select **Storage Subsystem** → **Change** → **Battery Settings** to set the scheduled day and time for battery check.

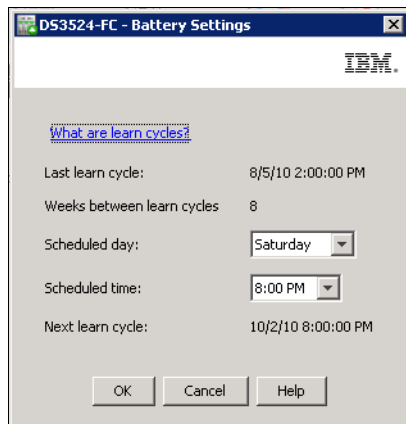


Figure 8-15 Battery Settings

**Learn Cycle:** The learn cycle for the DS3500 is much shorter than previous DS3000 models and takes about 15 minutes to complete. During this period the write caching on all of the logical drives can be disabled if battery low threshold is reached. As a general rule, schedule the learn cycle during a low usage period on the storage subsystem.

### Cache Settings

This menu item sets global cache settings of DS3500 storage subsystem. It is possible to set other cache setting as dedicated for each logical drive. It is described in 9.5.2, “Change Cache Settings” on page 241. Recommended cache settings depend on required host load and type of load, and can vary. General rules for cache settings are described in 3.3.10, “Cache parameters” on page 69.

Click **Storage Subsystem** → **Change** → **Cache Settings**. The window shown in Figure 8-16 on page 200 opens.

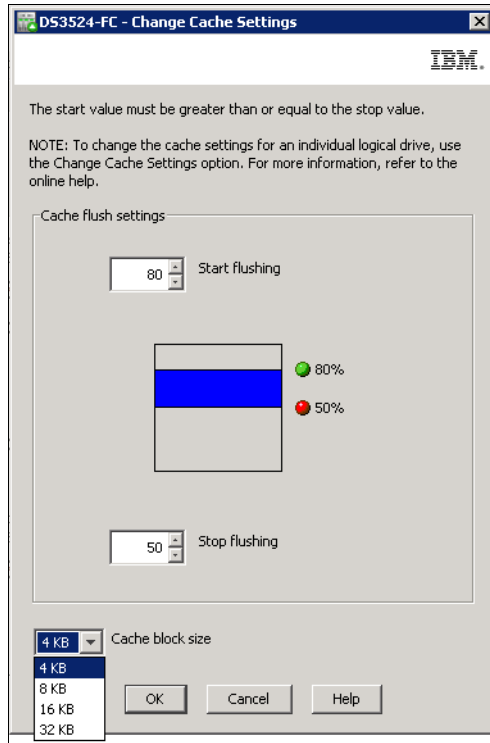


Figure 8-16 Cache Settings for Storage Subsystem

You have a choice to set the Start and Stop level for cache flushing, and the Cache block size. For the majority of installations, set 50%/50% for Start/Stop cache flushing and a 16 kB cache block size. But in certain cases a different setting can achieve better performance results. All settings in this window are immediately accepted if **OK** is selected and do not need any reset of DS3500 storage subsystem, so it is possible to test your settings during regular production on the subsystem.

### Default Host Type

The default host type can be changed in window as in Figure 8-17, after you click **Storage Subsystem** → **Change** → **Default Host Type**. The default host type is used as initial host type during new host definition in the Mapping tab.




Figure 8-17 Default Host Type

**Note:** The default host type is also host type for the Default Host Group, that is used for all host ports which are not assigned to storage partitions yet (see unassociated host ports in section “View ® Unassociated Host Port Identifiers” on page 191. It means no hosts are defined yet for these host ports in Mapping Tab. We do not recommend to use the Default Host Group for regular logical drive mappings as more hosts (with different operating systems) can have shared access to LUNs assigned to default host group and this can result in data loss on these LUNs.

### ***Failover Alert Delay***

A failover alert delay can be specified on DS3500 storage subsystem. It can delay the logging of a critical event if the multipath failover driver transfers the logical drive to the non-preferred controller of the subsystem. If the multipath driver transfers the logical drives back to the preferred controller within the specified delay period, no critical event is logged. If the transfer exceeds this delay period, then an alert is issued as a critical event.

Click **Storage Subsystem** → **Change** → **Failover Alert Delay**, and in the window as in Figure 8-18, you can set a range of 0-60 minutes. The default failover delay is 5 minutes.




Figure 8-18 Failover Alert Delay

**Note:** If the ownership transfers back to preferred controller within the delay period, the critical alert notification is not sent. However, a “Needs Attention” status is displayed immediately and the Recovery Guru shows details about that. See Chapter 14, “Administration - Support” on page 329 for more details about alert sending.

### ***Time-out Setting***

Select this menu item to change password inactivity time-out. A new, stronger rule for Storage Manager password was introduced in version 10.70 where the DS Storage Manager Subsystem Manager is locked after a selected period of inactivity. Default time-out is 30 minutes. If no activity in Storage Manager is recognized in this period, the password must be entered again to continue with configuration. It can be changed to anywhere from 1 minute to 2 hours. You can also set the time-out period to “Never” to disable the time-out feature.

Click **Storage Subsystem** → **Change** → **Time-out Setting**. Click and slide the pointer shown in Figure 8-19 on page 202 to the required value to set a new time-out period for the password.

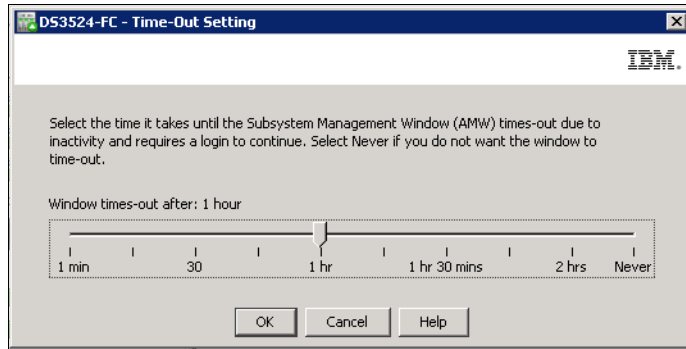


Figure 8-19 Time-out for password re-entering

### Enclosure ID

Each DS3500 Enclosure (including the controller enclosure) must have a unique Enclosure ID. To set this, click **Storage Subsystem** → **Change** → **Enclosure ID** and in the window that opens (Figure 8-20), select required enclosure based on current ID and set a new ID number. Do not use 0 (zero) as the ID because when a new empty and unused expansion enclosure is attached to the running storage subsystem, it has ID = 0, and if an existing enclosure has also ID = 0, it can have unexpected results.

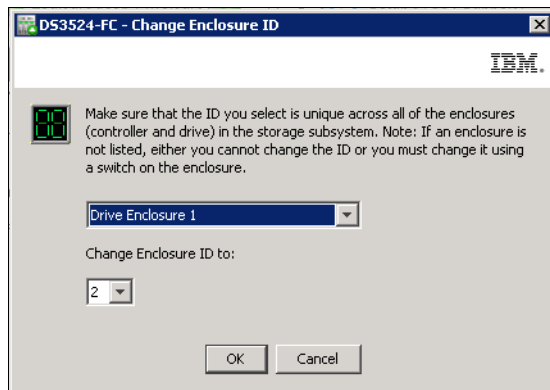


Figure 8-20 Change Enclosure ID window

### Enclosure Order

You can see enclosure layout in the left pane of the Physical tab as seen in Figure 8-21 on page 203. In this example the controller enclosure is below the drive enclosure. If you have the controller enclosure physically placed above the drive enclosure, you can change the layout to be the same as it is in your real rack.

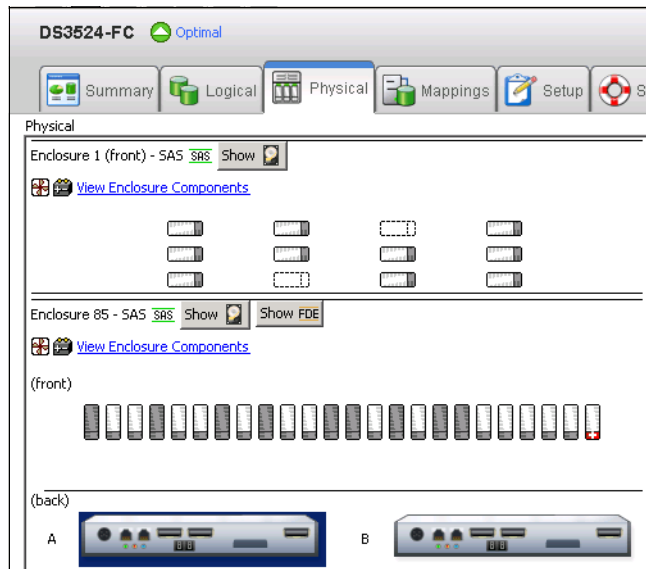


Figure 8-21 Enclosures Layout

Click **Storage Subsystem** → **Change** → **Enclosure Order** to change the layout. In the window shown in Figure 8-22, you can select the enclosure that has to be moved and y move it **Up** or **Down** using the buttons in the right side of the window.

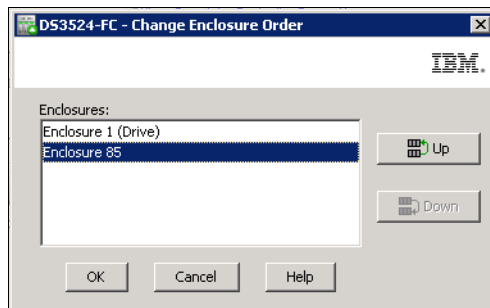


Figure 8-22 Enclosure Order change

## Synchronize Controller Clocks

Click **Storage Subsystem** → **Synchronize Controller Clocks** to set the correct time in DS3500 controllers. As shown in Figure 8-23, the time of controllers and the time on management station (where DS Storage Manager is running) are displayed. If you select the **Synchronize** button, than time, date, and time zone of the management station is used as the new time for controllers.




Figure 8-23 Synchronize Clocks

## Manage Enclosure Alarm

Because no DS3500 storage subsystem uses an acoustic alarm feature, this menu cannot be used.

## Rename Subsystem

When you add a new storage subsystem to your Storage Manager, it is identified as “Unnamed”. If you have several storage subsystems installed in your environment, you can give each subsystem a meaningful name.

Click **Storage Subsystem** → **Rename** to change the name of subsystem as shown in Figure 8-24.




Figure 8-24 Rename Subsystem

The subsystem name is also immediately updated in Enterprise Manager.

**Note:** Renaming the storage subsystem can have influence on services running outside the Storage Manager. For example, you can run some customized SMcli scripts where the subsystem is referenced by its name. Also, some types of multipath failover drivers on servers use the name of the attached subsystem. Rename a new subsystem as soon as possible before it is in production.

## Set Password

Starting with Storage Manager version 10.70, new stronger password rules have taken effect, so you must have passwords set for all DS3500-managed subsystems. You must enter the password when Subsystem Manager is started for given storage subsystem, even if you just want to see the current configuration. If you start Subsystem Manager for the first time, you must set a new password. In previous versions, the password was only requested in case of configuration changes such as deleting objects, and creating a new array or Logical Volume.

Now a strong password must meet the following requirements:

- ▶ Must be between 8 and 30 characters in length
- ▶ Must contain at least one number (0–9)
- ▶ Must contain at least one lowercase letter (a–z)
- ▶ Must contain at least one uppercase letter (A–Z)
- ▶ Must contain at least one non-alphanumeric character (!, @, #, \$, etc.)
- ▶ Must not contain any spaces, including trailing spaces

Select **Storage Subsystem** → **Set Password** to define a new password. You must set the current password and enter it twice in window like in Figure 8-25 on page 205.



Figure 8-25 Set Password

If you have forgotten the current password, contact your IBM Customer Support Representative. If a password was not set previously (such as if you have got a new subsystem), the current password is not required to establish a new password.

If the inactivity time-out is not set to “Never”, the window shown in Figure 8-26 will open after the defined period of inactivity time. Setting the inactivity time-out is described in “Time-out Setting” on page 201.

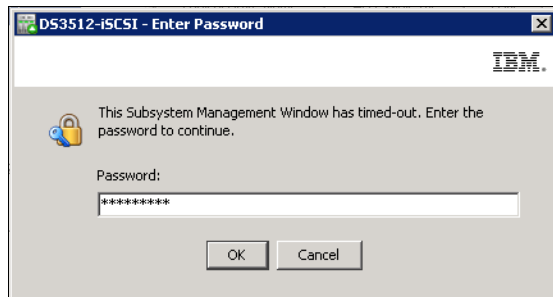


Figure 8-26 Re-enter password

## 8.2.2 View menu

The **View** pull-down menu (Figure 8-27) can be used for four relatively independent types of tasks.




Figure 8-27 View Menu

**View** menu items can be used:

1. For customizing the basic Storage Subsystem Manager window: By clicking **View** → **Toolbar** and **View** → **Status Bar**, you can hide or show these two areas in Storage

Manger window. Visibility of these bars is indicated by the check mark on the left side of these two menu items.

2. For viewing the associated logical elements and physical components: Use of these two menu items depends on the component that is selected in Logical, Physical, or Mappings tab.

If you select logical object in Logical tab or Mapping tab, clicking **View → Associated Physical Components** opens the window shown in Figure 8-28. Associated physical disk drives and their owning controller are marked by blue bullets.

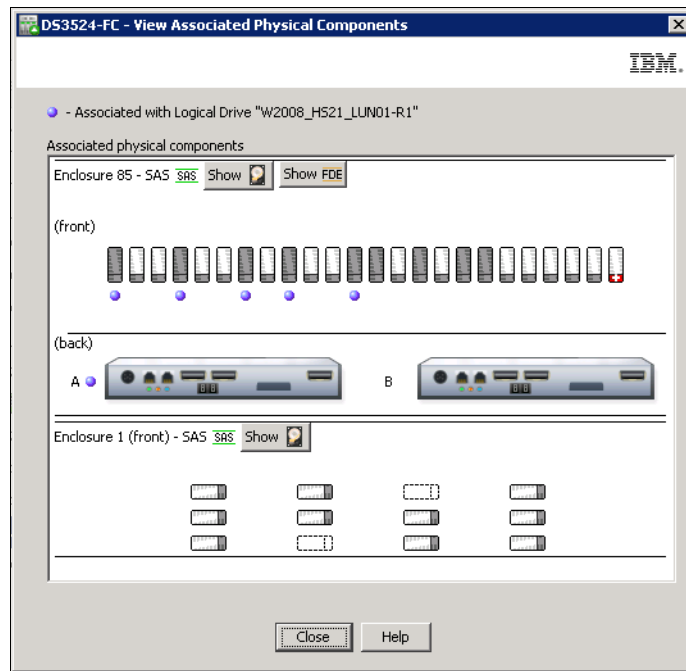


Figure 8-28 Associated Physical Components

Similarly, if you select a disk drive in the Physical tab, you can get related logical elements (Figure 8-29) if you click **View → Associated Logical Elements**.

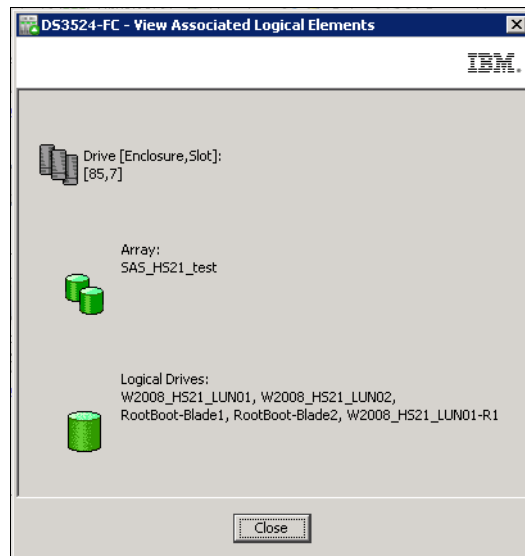


Figure 8-29 Associated Logical Elements



**Note:** The same windows can be opened directly in Logical, Physical, and Mappings tabs. Right-click the object and select **View Associated Component**. The same windows like in Figure 8-28 on page 206 or Figure 8-29 display.

3. For finding logical objects in the Logical and the Mapping tabs: You can search an array or a logical drive in the Logical Tab, or you can search Hosts and Host Groups in the Mapping Tab.

Click **View** → **Find** to open the window shown in Figure 8-30. You can specify the search object as text or just a part of searched text.

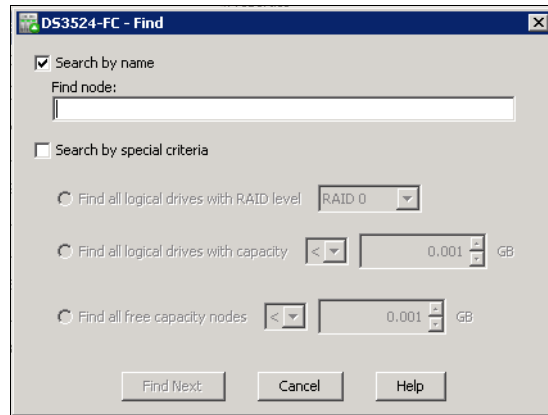


Figure 8-30 Find object

**Note:** The same function can be activated also in the Toolbar by clicking the **Find node in Tree** icon.

4. Skipping between FlashCopy or VolumeCopy objects: If you select FlashCopy or FlashCopy Repository, or alternatively VolumeCopy Source or VolumeCopy Target logical drives in Logical Tab, you can skip to associated logical drive in the copy pair. Click **View** → **Go to** and the associated logical drive will be highlighted.

**Note:** The same option is also available when you right-click the selected FlashCopy or VolumeCopy object. Select **Go to FlashCopy** or **Go to VolumeCopy** in the menu that appears.

### 8.2.3 Mappings menu

The **Mappings** pull-down menu (Figure 8-31 on page 208) is associated with Mappings tab. All items of this menu can be activated by right-clicking the appropriate object in the Mappings tab. For information on these settings, see Chapter 11, “Administration - Mappings Tab” on page 275.




Figure 8-31 Mappings menu

## 8.2.4 Array menu

The **Array** pull-down menu (Figure 8-32) is associated with the Logical tab. All items in this menu can be activated also by right-clicking the appropriate array (or unconfigured capacity, if you can create a new array) that is selected in the Logical tab. If another tab is selected, no menu item can be selected. For information on these settings, see Chapter 9, “Administration - Logical Tab” on page 219.




Figure 8-32 Array menu

## 8.2.5 Logical Drive menu

The **Logical Drive** pull-down menu (Figure 8-33) is associated with the Logical tab. All items of this menu can be activated also by right-clicking the appropriate logical drive (or free capacity, if you can create a new logical drive) that is selected in the Logical tab. If another tab is selected, no menu item can be selected. For information on these settings, see Chapter 9, “Administration - Logical Tab” on page 219.




Figure 8-33 Logical Drive menu

## 8.2.6 Controller menu

The **Controller** pull-down menu (Figure 8-34 on page 209) is associated with the Physical tab. All items of this menu can be activated by right-clicking the appropriate controller in the Physical tab. If another tab or object in the Physical tab is selected, no menu item can be

selected. For information on these settings, see Chapter 10, “Administration - Physical Tab” on page 257.




Figure 8-34 Controller menu

## 8.2.7 Drive menu

The **Drive** menu (Figure 8-35) is associated with the Physical tab. All items of this menu can be activated also by right-clicking the appropriate disk drive in the Physical tab. If another tab or object in Physical tab is selected, no menu item can be selected. For information on these settings, see Chapter 10, “Administration - Physical Tab” on page 257.




Figure 8-35 Drive menu

## 8.2.8 Advanced menu

The **Advanced** pull-down menu (Figure 8-36) is used for advanced configuration, troubleshooting, and maintenance. For information on these settings, see Chapter 14, “Administration - Support” on page 329.




Figure 8-36 Advanced menu

## 8.2.9 Help menu

The last pull-down menu is the **Help** menu (Figure 8-37). Select **Contents** or **Reference Recovery Procedures** to access Online Help. You can also access embedded help pages with content related to the information about almost every window in Storage Manager.




Figure 8-37 Help menu

The Online Help window (Figure 8-38 on page 210) has four basic tabs.




Figure 8-38 Help window

From left to right are the following tabs:

1. Content

In the Content tab you can find all the chapters of the Online Help in the left part of the window. Select the chapter you are interested in, and its content is displayed in the right pane.

2. Index

3. Glossary

The Index and the Glossary tabs offer an alphabetic list of most important word and terms used in online help. If the word you are looking for is not found, use the Search tab.

4. Search

Type a term or phrase into the Find tab, and a list of results will be in the left pane. The number next to each line shows how many times the word is used on the page. Click a topic link in the left pane and the content of that page is displayed in the right pane as shown in Figure 8-39 on page 211.

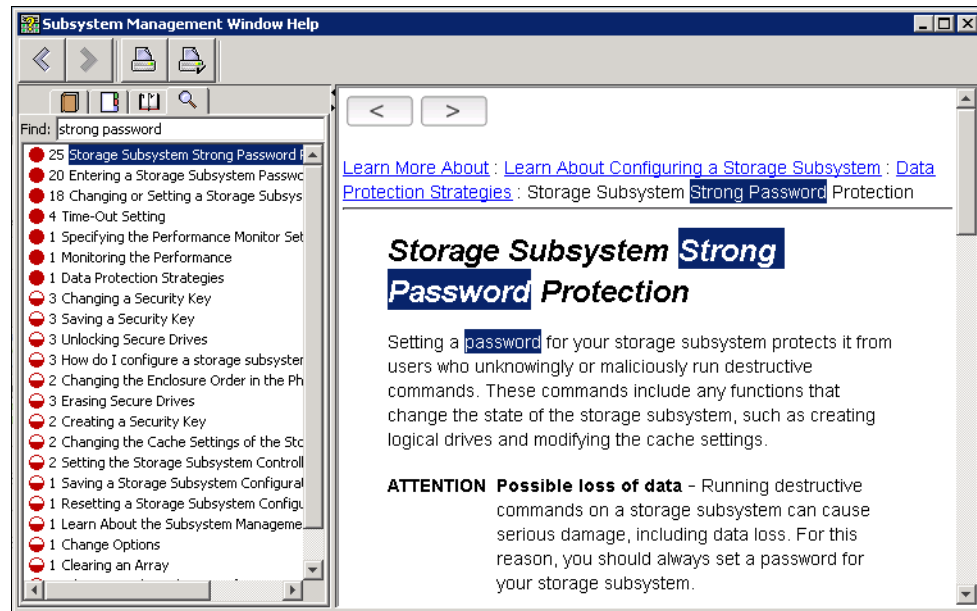


Figure 8-39 Help - Find

Click **Help** → **About** to get information (Figure 8-40) about the version of IBM System Storage DS Storage Manager that you are using. It is important to know that the version of the Enterprise Manager can be higher than the version of Subsystem manager because the Subsystem Manager follows the version of the firmware level of the managed subsystem.

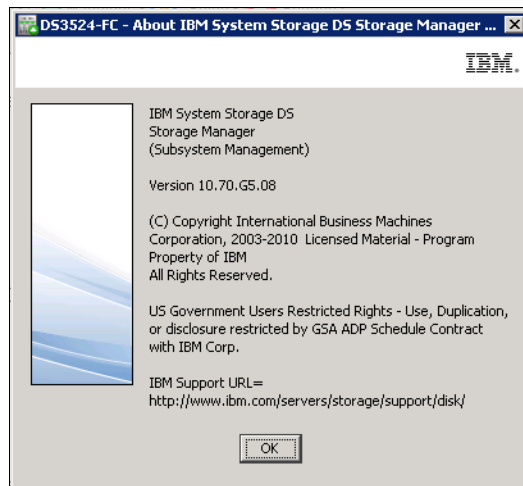


Figure 8-40 About DS Subsystem Manager

## 8.3 Toolbar

Several icons are in the **Toolbar** area and shown in Figure 8-41. These icons can help you to easily navigate to certain management tasks.

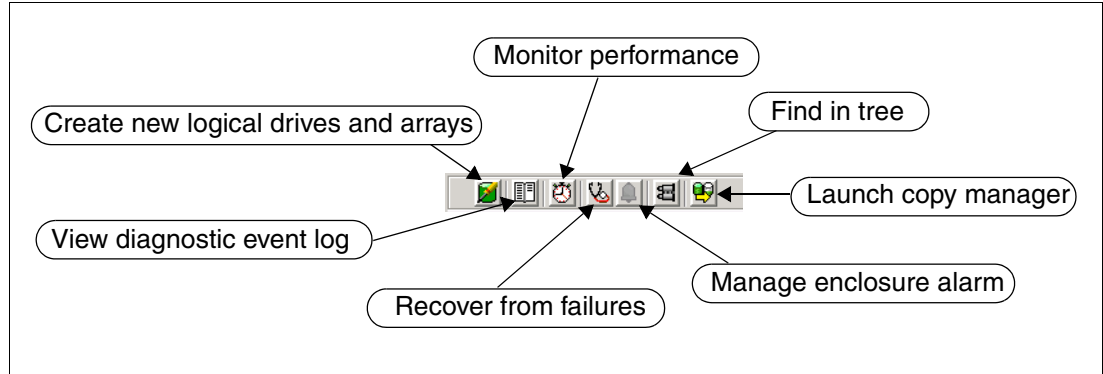


Figure 8-41 Toolbar icons

The seven icons in the Toolbar are (from left to right):

- ▶ 8.3.1, “Create new logical drives and arrays” on page 212
- ▶ 8.3.2, “View diagnostic event log” on page 212
- ▶ 8.3.3, “Monitor Performance” on page 213
- ▶ 8.3.4, “Recover from failures” on page 214
- ▶ 8.3.5, “Manage enclosure alarm” on page 214
- ▶ 8.3.6, “Find in tree” on page 214
- ▶ 8.3.7, “Launch copy manager” on page 214

### 8.3.1 Create new logical drives and arrays

You can select this icon only if the Logical tab is activated and an “Unconfigured Capacity” or “Free Capacity” item is selected. If you click it, the window with the wizard opens. The same wizard opens if you right-click an unconfigured or free capacity item in the Logical tab. You can easily configure logical drives directly from there. Creating a new logical drive or array is further described in Chapter 9, “Administration - Logical Tab” on page 219.

### 8.3.2 View diagnostic event log

The diagnostic event log, also called the Major Event Log (MEL), is one of the basic sources for maintenance and troubleshooting. Click this icon to open the window shown in Figure 8-42 on page 213. All the details about the event log is described in Chapter 14, “Administration - Support” on page 329.

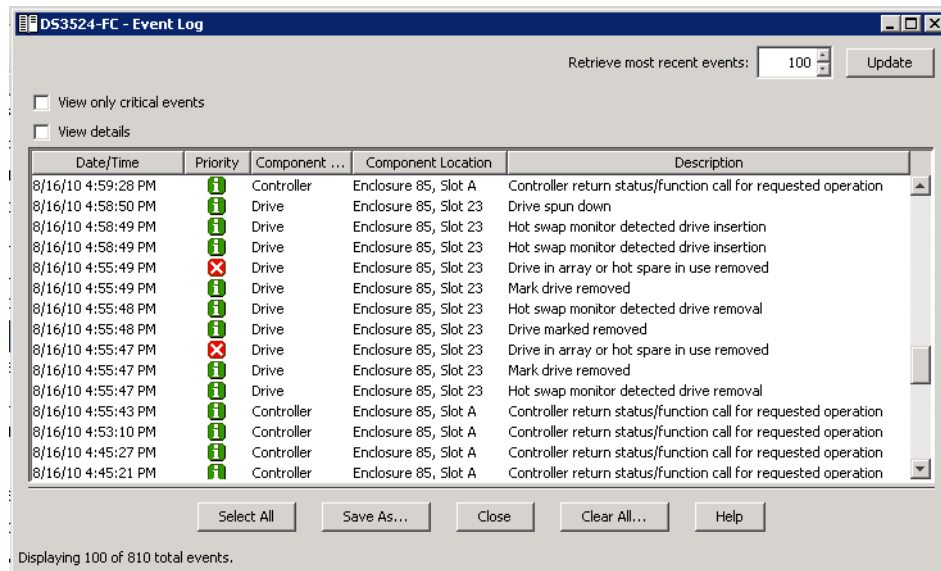


Figure 8-42 Event Log

The Event Log window remains open even if you go back to the subsystem manager window. You must click the **Update** button to get fresh events displayed when you come back.

### 8.3.3 Monitor Performance

Click this icon to open the Performance Monitor window as shown in Figure 8-43. This monitoring tool is important for performance analysis of managed storage subsystem. How to understand displayed data and set performance monitoring is described in Chapter 14, “Administration - Support” on page 329.

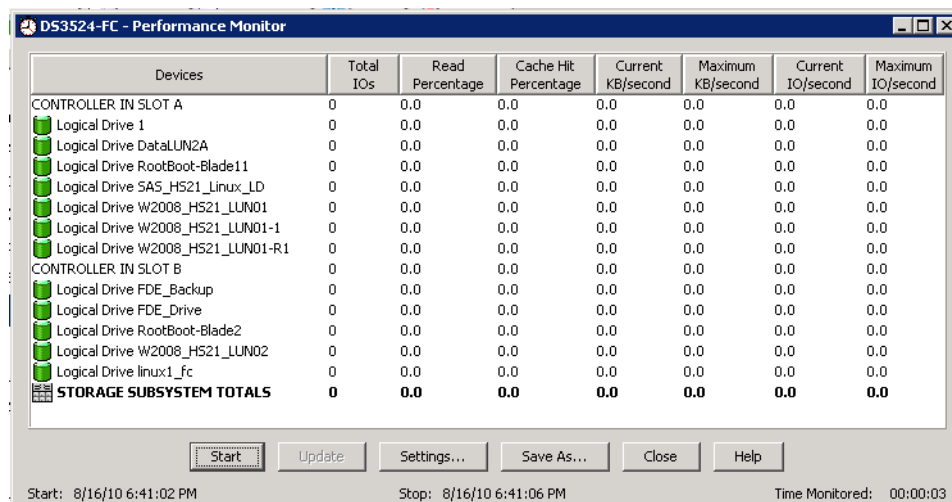


Figure 8-43 Performance Monitor window

This window remains open and will continue to collect performance data even if you go back to Subsystem Manager window.

### 8.3.4 Recover from failures

This icon (also called the Recovery Guru) has dynamic behavior. If the whole storage subsystem is in optimal status, it is stable and clicking it opens just window with the information No failures were detected on the storage subsystem. If a failure happens, the small red shape starts flash on the Recovery Guru icon. Select the icon to get window similar to Figure 8-44 where the failure explanation and recommended actions are displayed.

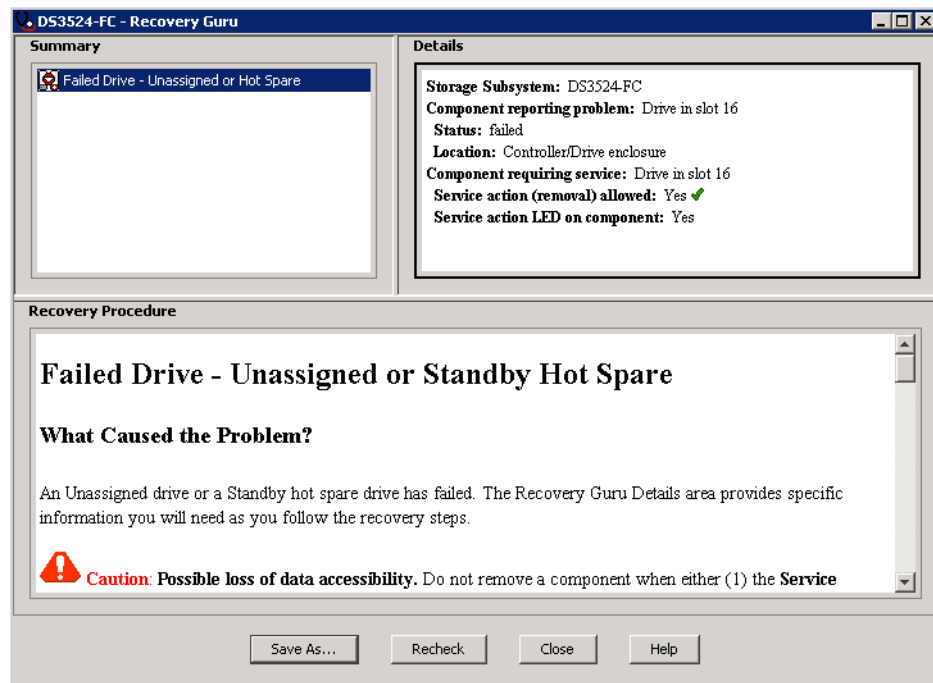


Figure 8-44 Recovery Guru window

The Recovery Guru window remains open if you go back to Subsystem Manager window. If the failure is corrected, it will close after you click **Recheck**.

More detailed description of the Recovery Guru can be found in Chapter 14, “Administration - Support” on page 329.

### 8.3.5 Manage enclosure alarm

This icon normally helps to set an acoustic alarm on managed subsystem, but because DS3500 storage subsystem has no beeper for acoustic alarms, it is not possible to use it.

### 8.3.6 Find in tree

This icon works only when the Logical or Mapping tab is activated. It is used to find logical objects in a selected tab, and opens a window like the one shown in Figure 8-30 on page 207.

### 8.3.7 Launch copy manager

The Copy Manger is a tool (see Figure 8-45 on page 215) that can help to monitor progress and status of VolumeCopy operations. Also, configuration changes can be done on existing VolumeCopy pairs. Find detailed information about the Copy Manager, VolumeCopy, and



about all other copy services in *IBM Midrange System Storage Copy Services Guide*, SG24-7822.




Figure 8-45 Copy Manager window

The Copy Manger window can stay open while you configure other things in the Subsystem Manager window. Copy progress is refreshed automatically, and copy operations are not affected if the window is closed.

## 8.4 Status bar

Several icons are displayed in the Status bar area, and they can easily show the types of license that are available on the managed subsystem. Some of them are shown activated. If the icon is crossed by red line as shown in Figure 8-46, this premium feature is not activated.

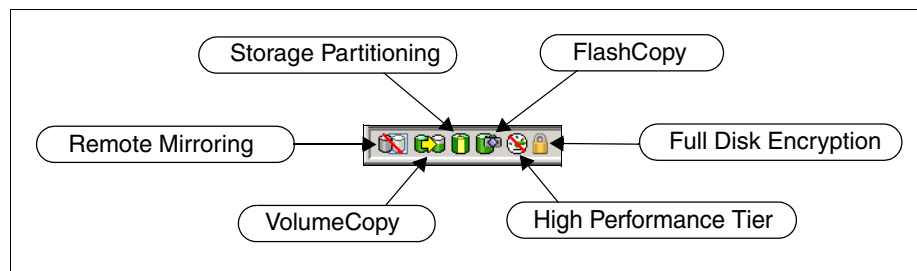


Figure 8-46 Status bar icons

If you move your cursor onto one of the icons, a small text box (Figure 8-47) is shown after about one second. It shows the description of selected Premium Feature and its enablement status.



Figure 8-47 Premium Feature icon - description text

To activate a new licenses (Premium Features) you can click **Storage Subsystem** → **Premium Features**, as described in “View/Enable Premium Feature” on page 195.

## 8.5 Tabs

This is the main management area of the **IBM System Storage DS (Subsystem Management)**. In this area (see Figure 8-48) you can see the highlighted name of managed storage subsystem in upper part, and close to the name is an icon showing current status of the whole storage subsystem. It changes from green to red if the subsystem has degraded components, or when something has failed. The description changes from Optimal to Needs Attention in this case.

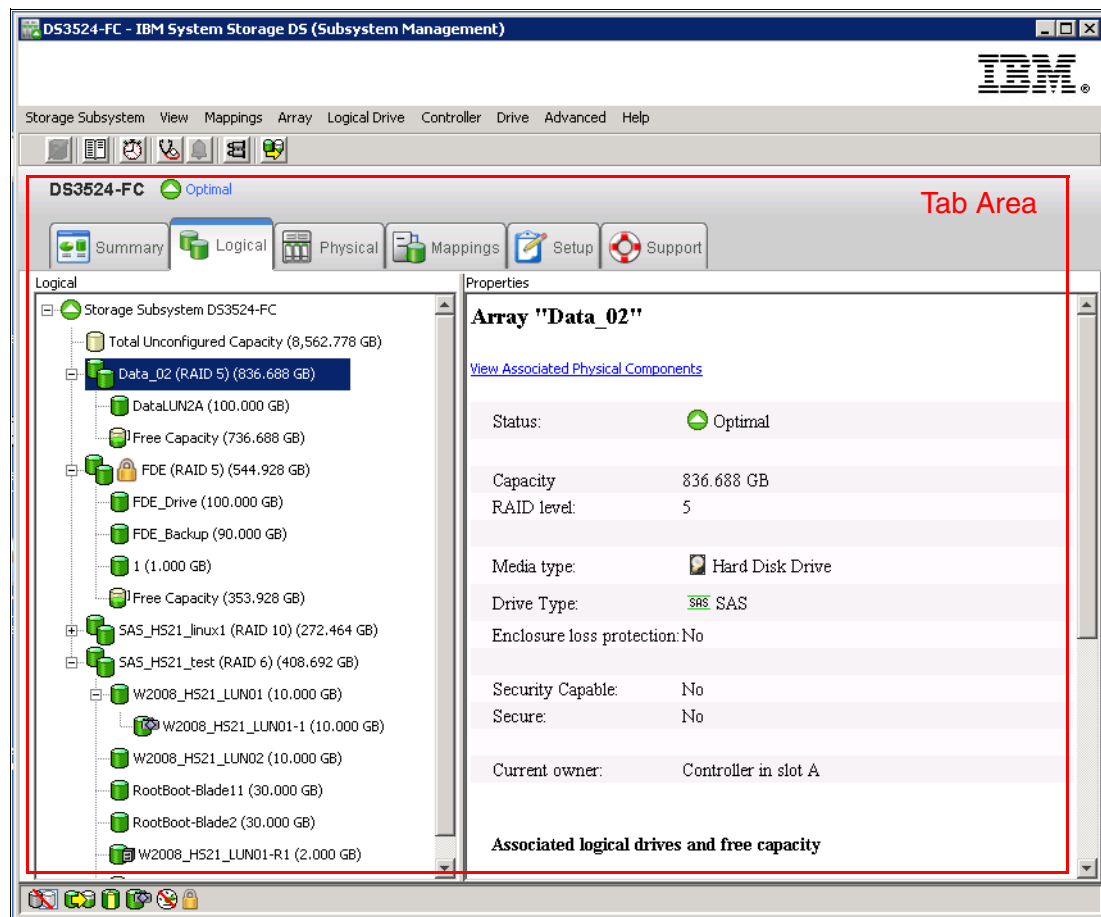


Figure 8-48 Tab area

Each of these six tabs are described in the following chapters in this book:

- ▶ Summary tab in Chapter 7, "Administration - Summary Tab" on page 177.  
Shows the status of the storage subsystem. Information about the hardware components, capacity, hosts, mappings, and storage partitions, arrays, and logical drives in the storage subsystem appears. Additionally, the link to the Recovery Guru window is shown when the storage subsystem needs attention. Links to the Storage Subsystem Profile window, relevant online help topics, and the storage concepts tutorial also appear in the Summary tab.
- ▶ Logical tab in Chapter 9, "Administration - Logical Tab" on page 219.  
Shows the organization of the storage subsystem by logical drives, arrays, free capacity nodes, and any unconfigured capacity for the storage subsystem. This tab is divided into two panes: the Logical pane on the left and the Properties pane on the right. The Logical

pane provides a view of the logical components in the storage subsystem in a tree structure. The Properties pane displays detailed information about the component that is selected in the Logical pane.

- Physical tab in Chapter 10, “Administration - Physical Tab” on page 257.

Shows the organization of the storage subsystem by controllers, drives, and other hardware components. This tab is divided into two panes: the Physical pane on the left and the Properties pane on the right. The Physical pane shows a graphical representation of the physical components in the storage subsystem, such as the controller modules and the drive expansion enclosures. The Properties pane displays detailed information about the component that is selected in the Physical pane.

- Mappings tab in Chapter 11, “Administration - Mappings Tab” on page 275.

Contains the Topology pane and the Defined Mappings pane. In the Topology pane, you can define the hosts, host groups, and host ports. The Topology pane shows a tree structured view of the logical nodes related to storage partitions in the storage subsystem. The Defined Mappings pane shows the mapping that is associated with the selected logical node in the Topology pane. You can also change the mappings to grant logical drive access to host groups and hosts and create storage partitions.

- Setup tab in Chapter 12, “Administration - Setup tab” on page 297.

Contains links to tasks that you can perform when setting up a storage subsystem such as configuring the storage subsystem, setting the storage subsystem password, and other tasks. Additionally, the links to configure the Ethernet management ports in the storage subsystem, view and enable the premium features, and configure the iSCSI settings are shown in the Setup tab. You can click a link to open the corresponding window.

- Support tab in Chapter 14, “Administration - Support” on page 329.

Contains links to tasks, such as recovering from storage subsystem failure; gathering support information, such as the Event Log and a description of a storage subsystem; and sending the information to your IBM Customer Support representative. The Support tab also contains the link to download the controller firmware, the NVSRAM, the drive firmware, the ESM firmware, and the ESM configuration settings. Additionally, the Support tab has links to view the description of all components and properties of a storage subsystem, the Event Log of a storage subsystem, online help, and the version and copyright information of the storage management software. You can click a link to open the corresponding window.





## Administration - Logical Tab

In this chapter, we describe the Logical tab of the DS Storage Manager. In this tab, you can logically configure a DS3500 Storage Subsystem and these configuration options include important tasks like:

- ▶ Creation of arrays and logical drives
- ▶ Creation of FlashCopy, VolumeCopy, and Remote Mirror Drives
- ▶ Display of properties of these logical objects
- ▶ Usage of dynamic functions
- ▶ Change of controller ownership
- ▶ Change of modification priority, cache settings, and media scanning

## 9.1 Logical tab

The Logical tab (Figure 9-1) shows the organization of the storage subsystem by logical drives, arrays, free capacity nodes, and any unconfigured capacity for the storage subsystem. This tab is divided into two panes: the *Logical pane* on the left and the *Properties pane* on the right. The Logical pane provides a view of the logical components in the storage subsystem in a tree structure. The Properties pane displays detailed information about the component that is selected in the Logical pane.

**Note:** The Properties pane is not used for any configuration tasks. It just shows properties of selected objects in the same form as it is in the subsystem profile. You can find more information about the profile in 14.3, “View Storage Subsystem Profile” on page 338.

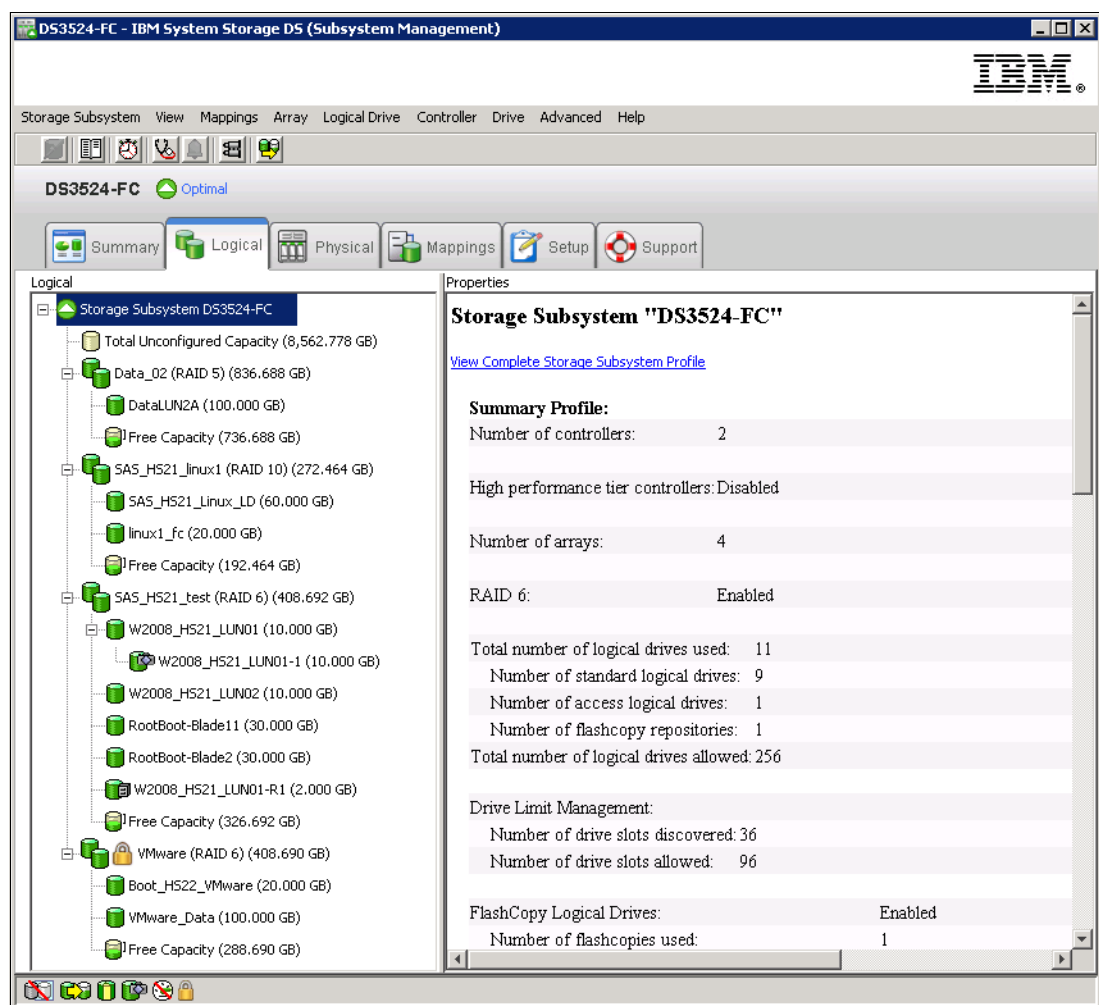


Figure 9-1 Logical tab

The first line in the Properties pane is the **View Associated Components** link. When selected, it opens a window with the physical components belonging to the selected logical object in the left Logical pane. The associated components are marked by blue bullets in this window.

If the first line of Logical pane (Storage Subsystem) is selected, the link opens a window with Subsystem Storage profile. As usual, you can also get required information from menus, tabs, and icons in DS Storage Manager.

Your options to configure the storage subsystem depend on the object selected in the Logical pane. All configuration options are shown as a menu when you right-click a particular object.

We have split this chapter into four basic sections to explain the following configuration options:

- ▶ 9.2, “Working with unconfigured capacity” on page 222
- ▶ 9.3, “Working with arrays” on page 225
- ▶ 9.4, “Working with Free Capacity” on page 232
- ▶ 9.5, “Working with logical drives” on page 237

The first line in the Logical pane (the root of the logical directory tree) is **Storage Subsystem** (already selected in Figure 9-1 on page 220). The name of subsystem and basic status (green - Optimal, or red - Needs Attention) of this managed subsystem is shown there. In the right Properties pane you can see overall information about the whole DS3500 Storage Subsystem.

Right-click the Storage Subsystem in the left pane and a full list of possible configurations opens in menu as shown in Figure 9-2. The first menu item is **Copy Manager**. Click it and a new window opens that lists the VolumeCopy pairs. You can find all details about VolumeCopy, Copy Manger, and more copy services information in *IBM Midrange System Storage Copy Services Guide*, SG24-7822. All other menu items are the same as when you click the **Storage Subsystem** menu in pull-down menu bar of DS Storage Manager. This menu is described in 8.2.1, “Storage Subsystem Menu” on page 189.

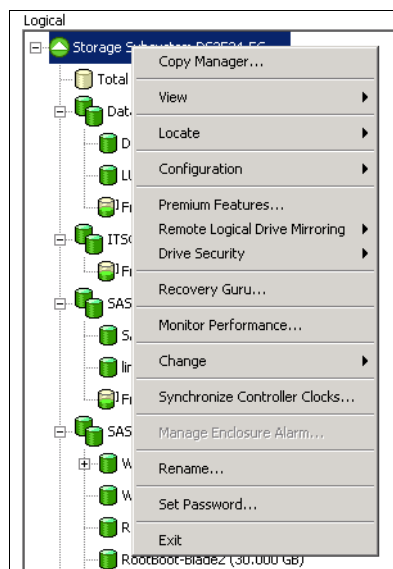


Figure 9-2 Storage Subsystem menu

Several other objects in the left Logical pane include:

- ▶ Unconfigured capacity
- ▶ Array
- ▶ Free capacity
- ▶ Standard Logical Drive
- ▶ Special Logical Drive (FlashCopy, FlashCopy Repository, Mirror Repository, and so on)

In the following sections we will describe most of them. others are described in chapters or sections dedicated to specific topics.

## 9.2 Working with unconfigured capacity

If you get a new DS3500 subsystem, it probably has all unconfigured capacity. As you start to configure physical disk drives into Logical arrays, this unconfigured capacity will decrease. When the whole subsystem is configured and no disks remains, this object disappears from the Logical pane.

The main task that can be performed on the unconfigured capacity is to create an array. In Figure 9-3, you can see an example of a partially configured DS3524 Storage Subsystem, where about 7 TB unconfigured capacity still remains.

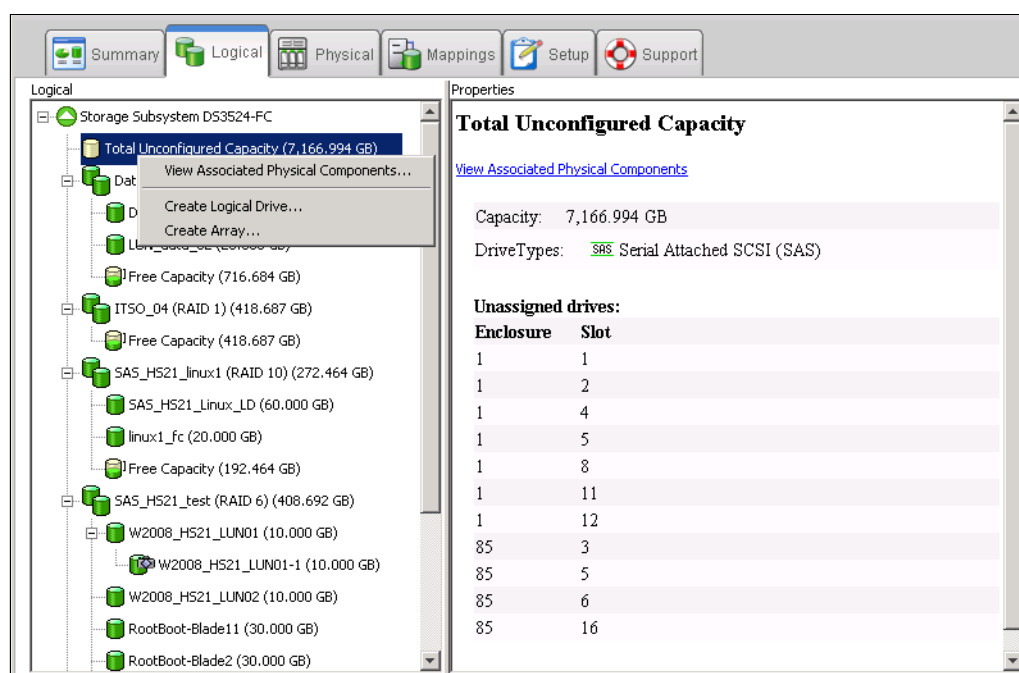


Figure 9-3 Unconfigured Capacity example

### 9.2.1 View Associated Physical Components

Select unconfigured capacity and in the Properties pane on the right side, you can see the list of associated disk drives that are not configured yet. The **View Associated Physical Components** link can also be used to open a window where the unconfigured disks are placed (see Figure 9-4 on page 223 for an example).

The same window opens if you right-click the Unconfigured Capacity and select **View Associated Physical Components**.



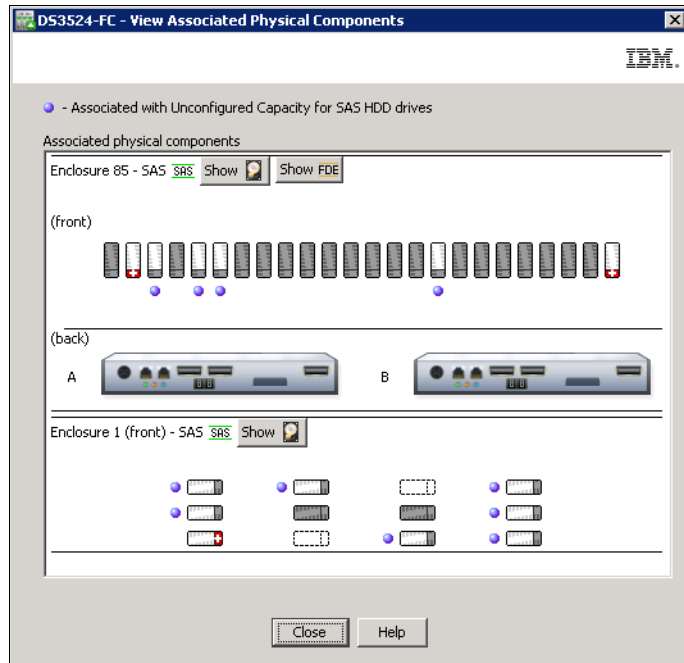


Figure 9-4 Unconfigured Disks

The disks marked with the blue bullets are unconfigured. You can also recognize this by the icon for unused disks.

Right-click the Unconfigured Capacity in the Logical pane, and start configuring an array by selecting **Create Array**.

If you can create a logical drive directly from Unconfigured Capacity (see the menu in Figure 9-3 on page 222), then you must create an array by first choosing the RAID type, disk drives, and so on, and then continue with the logical drive configuration process. If an array is not created, you can still create a logical drive, but the Storage Manager will direct you to create the array first.

## 9.2.2 Create array

The example that follows assumes that there is unconfigured capacity on the DS3500 storage subsystem. This procedure illustrates the most common steps to follow in setting up an array from unconfigured drives. We first start with the creation of an empty array, and then the logical drive creation is described in 9.4.1, “Create logical drive” on page 233.

Perform the following steps to create an array:

1. In the Logical pane (Figure 9-5), right-click the Total Unconfigured Capacity and select **Create Array**.




Figure 9-5 Create Array from Unconfigured Capacity

This action starts the wizard for creating the array. The first window of the wizard is an introduction to the process. It displays the available unconfigured capacity. Read the introduction and then click **Next** to proceed.

2. In the Create Array window, type a meaningful array name that describes your set of disks, and then select either **Automatic mode** or **Manual mode**. Automatic mode is the default option as shown in Figure 9-6. By selecting the Automatic option, the Storage Manager software selects a combination of available drives to optimize performance and availability, and attempts to select physical drives from separate enclosures to provide enclosure protection whenever possible. If some SED disk drives are used in DS3500, you will also see the check box for selecting the encrypted disks for array creation. For more information about disk encryption, refer to Chapter 15, “Disk Security with Full Disk Encryption drives” on page 439.

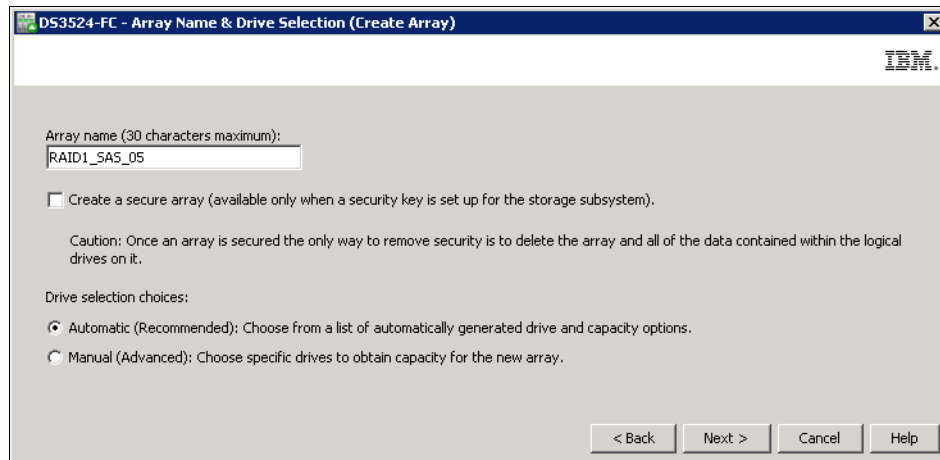


Figure 9-6 Create Array and assign a name and mode

The Automatic (Recommended) creation of the array selects the drives in the following order:

- Same capacity-same speed enclosure redundancy
- Same capacity-mixed speed enclosure redundancy
- Same capacity-same speed no enclosure redundancy
- Same capacity-mixed speed no enclosure redundancy
- Mixed capacity-same or mixed speed-no enclosure redundancy

In the Manual (Advanced) mode, you have to select all the drives individually. Make sure that you select them to maximize performance and availability.

You can also perform manual selection from the Physical tab. You can select multiple unconfigured disk drives in Physical pane using the Ctrl key before selecting to create a new array. If you right-click one of preselected disks and select **Create Array**, the same initiation window will open as shown in Figure 9-6. In this window you can see manual mode is preselected. With this manual mode, follow the instructions in this chapter, and when you reach the window for disk drives selection, you will recognize, your planned disks are already selected for new array.

3. Click **Next** to select the RAID level.
  - Select the desired RAID level. The window shown in Figure 9-7 on page 225 now displays the capacity options depending on the unconfigured drives available in your subsystem. If you have different disk sizes, you have more than one option for the same number of disks. See Chapter 3, “IBM System Storage DS3500 Storage System

planning tasks” on page 29 to determine the best RAID level for your specific environment and application.

- Select the line with required capacity in the list and click **Finish**. In our example (Figure 9-7), we create a RAID1 array with two 15kRPM SAS drives.

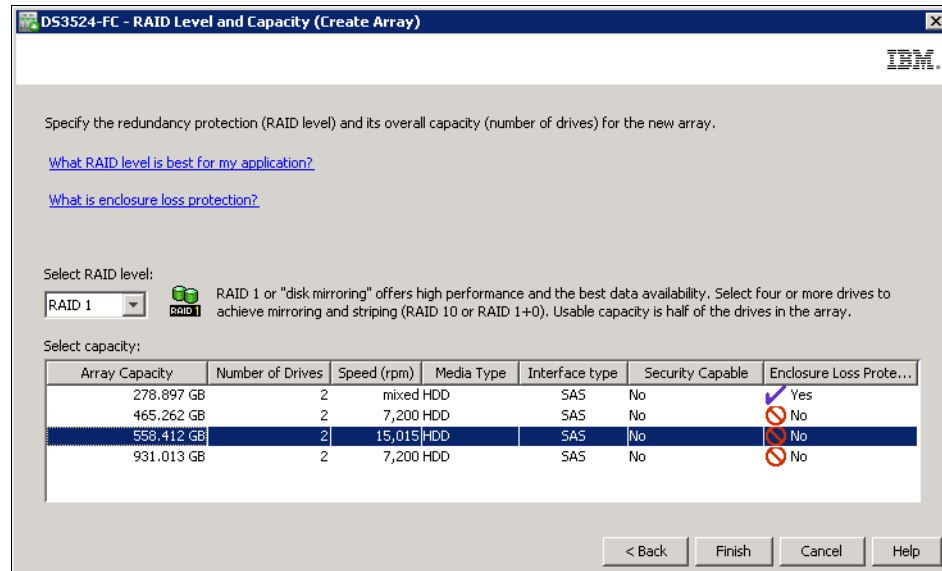


Figure 9-7 RAID level and capacity

- The Array Success window (Figure 9-8) appears and confirms that the array is now created. You can click **Yes** to continue with the creation of a logical drive (see 9.4.1, “Create logical drive” on page 233). Select **No** to finish for now.

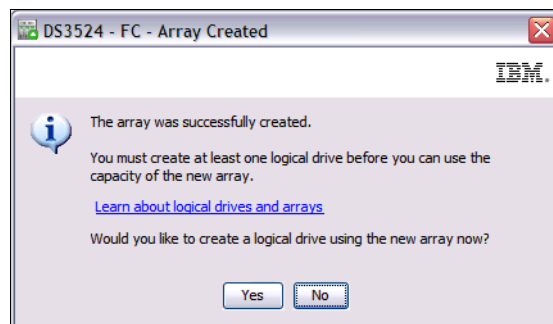


Figure 9-8 Array Created

- Your empty array is created immediately because it does not need to format any logical drive capacity. You can repeat the same process for creating another array.

## 9.3 Working with arrays

The array (sometimes referenced also as Disk Group) is a set of disks with a defined RAID level. You can configure one or more logical drives on it. Available capacity of the array is displayed in the Logical pane as Free Capacity. If no logical drives are configured, the free capacity displayed is also total capacity of the array.

Actions that you can perform on the array can be accessed by right-clicking the array and selecting from the menu (Figure 9-9).

These options are described in the following sections:

- ▶ 9.3.1, “Locate and View Associated Components” on page 226
- ▶ 9.3.2, “Change Ownership and RAID level” on page 227
- ▶ 9.3.3, “Add Free Capacity (Drive)” on page 230
- ▶ 9.3.4, “Secure Drive” on page 231
- ▶ 9.3.5, “Delete and Rename” on page 231
- ▶ 9.3.6, “Replace Drive” on page 232

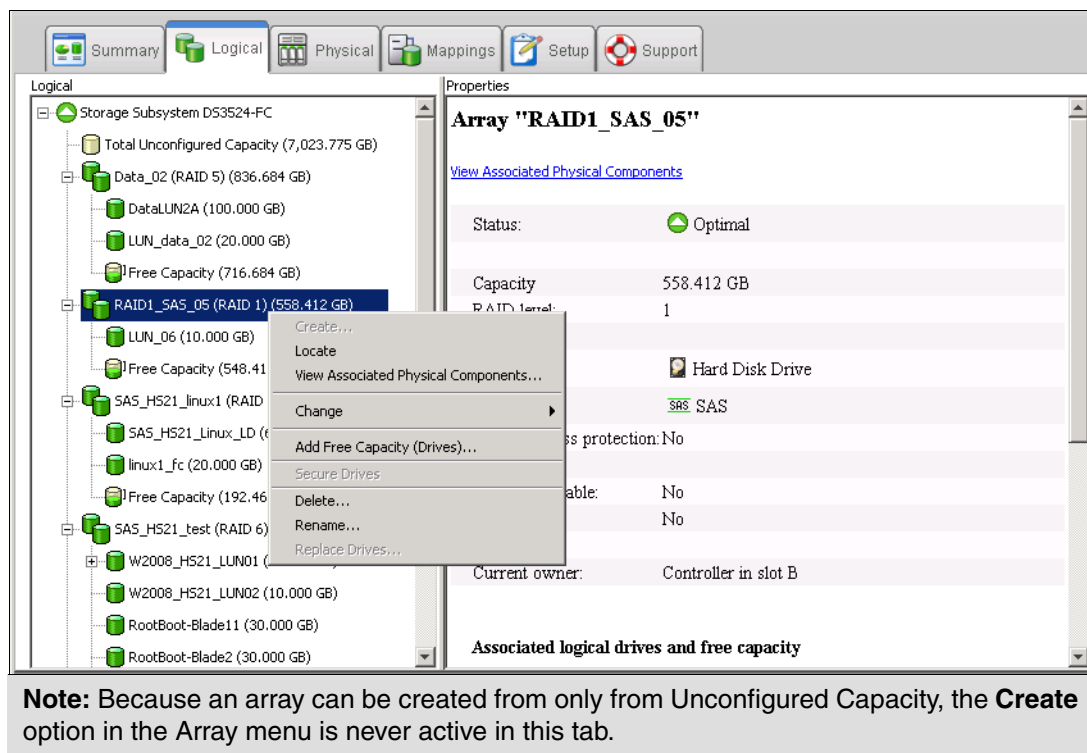


Figure 9-9 Array menu

### 9.3.1 Locate and View Associated Components

Select **Locate** to find the disk drives in the rack from which the array is configured. The orange LED diodes on these disks blink.

Select **View Associated Physical Components** to see them in window as shown in Figure 9-10 on page 227. The disks marked by blue bullets are members of the selected array. The blue bullet can mark also one or both controllers, if the logical drives are owned by both controllers.

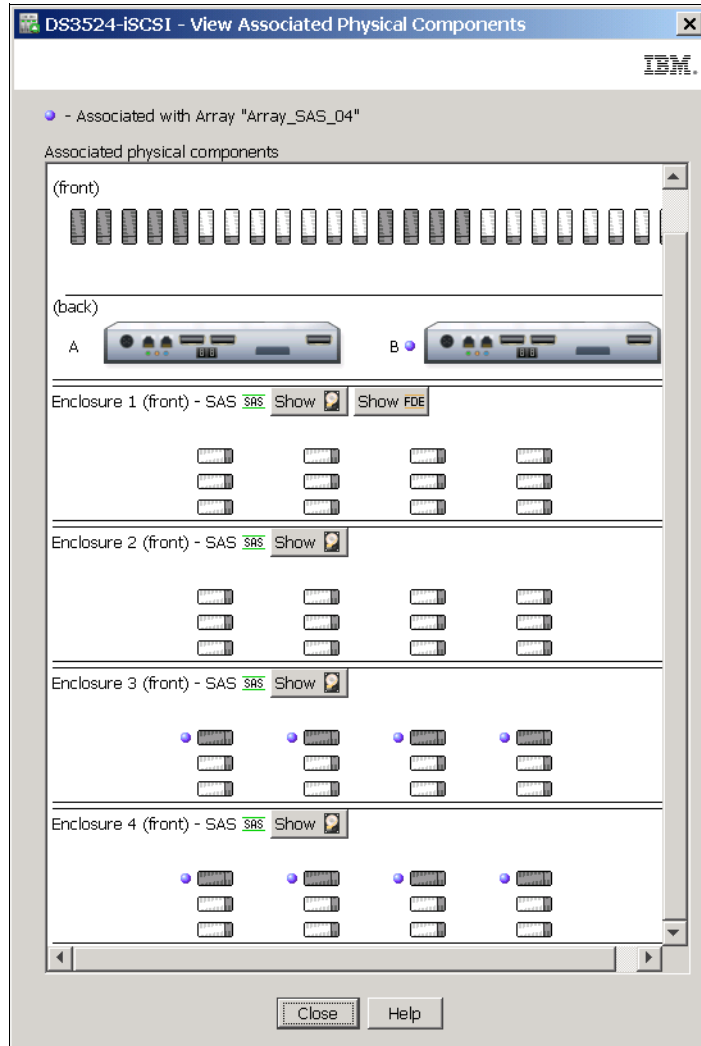


Figure 9-10 Associated components for array

### 9.3.2 Change Ownership and RAID level

In this section we discuss changing the controller ownership and the RAID level.

#### Ownership

Click **Change** → **Ownership/Preferred Path** → **Controller A/B** to change the logical drive controller ownership of all logical drives in the array to one selected controller. You can be sure about multipath drivers and redundant paths in your environment. The change of controller ownership must be confirmed as shown in Figure 9-11 on page 228.

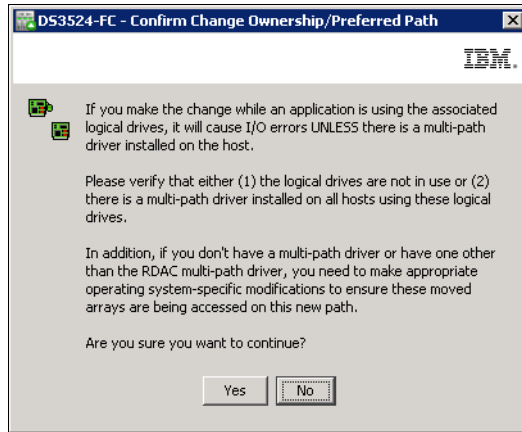


Figure 9-11 Change Ownership confirmation

## RAID level

Changing the RAID level of an array is performed in a non disruptive manner. This feature is called *Dynamic RAID Migration (DRM)*. The system remains fully operational while the process takes place. You might want to perform this operation for the following reasons:

- ▶ The storage requirements changed over time and existing RAID levels are no longer optimal for a particular environment.
- ▶ The performance tuning process indicates that a different RAID level is more appropriate than the existing one.

It is possible to change any RAID level depending on the following restrictions that apply to the new arrays:

- ▶ RAID 1 or 10 requires an even number of disk drives.
- ▶ RAID 3 and 5 require at least three drives.
- ▶ RAID 6 requires at least five drives.
- ▶ There is a limit of 30 drives per array for RAID 3, 5, and 6 arrays.

There are also other limitations when there is not enough free space in the array. For example, a RAID 5 array of four disk drives with no free space cannot be migrated directly to RAID 1. If this migration is attempted, an error message will be displayed stating that there is not enough free space. There must be enough free capacity to change the RAID level. Also, if the array has an odd number of drives and a migration to RAID 1 is required, a disk must be added to the array prior to performing the procedure.

When changing from RAID 1 to RAID 5, free space in the array can be gained that can be used to define new logical drives or expand existing ones.

When the procedure starts, it reorganizes the data segments in the array according to the new RAID level, and a large amount of I/O happens, so there will be an impact on performance while the migration happens. The performance impact can be controlled to a certain extent by changing the value of the modification priority. This parameter is set on a logical drive basis, and it can be changed for one logical drive, for a set of them, or for all logical drives in the array. See 9.5.1, “Change Modification Priority” on page 239 for more details.

Changing the modification priority to a low value during the migration process minimizes performance degradation. When the migration finishes, the value can be increased to reduce the time for a rebuild in case of a drive failure. This minimizes the critical time of non-redundant operation caused by a disk drive fault.

**Attention:** After the migration starts, it cannot be stopped. During the migration process, the array cannot be deleted, logical drives cannot be deleted, and new logical drives cannot be created in the migrated array.

Certain configuration changes that do not change data structures on the array can be performed on logical drives during the migration process. These configuration changes include changing the modification priority, logical drive controller ownership, name change. Also it's possible to define the mappings during RAID level migration.

**Note:** Even though RAID migration is a non-disruptive process, carry out this migration when I/O activity is at a minimum.

To change the RAID level, right-click the array and select **Change** → **RAID Level** → **n** as in Figure 9-12. The current RAID level is marked by black bullet, and you can select your new RAID level.

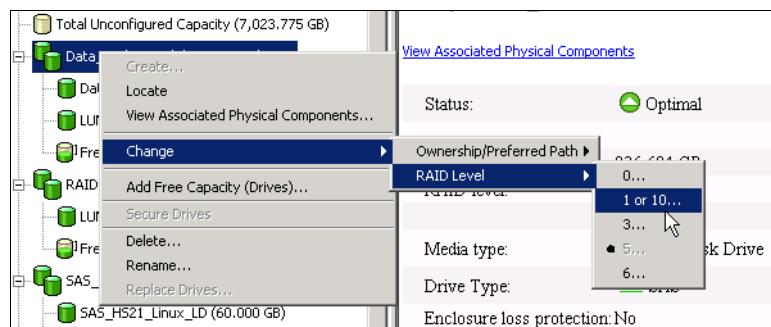


Figure 9-12 Change RAID level menu

The process of RAID level migration is non-disruptive and no redundancy is lost, but it can take a long time. A warning as shown in Figure 9-13 appears.

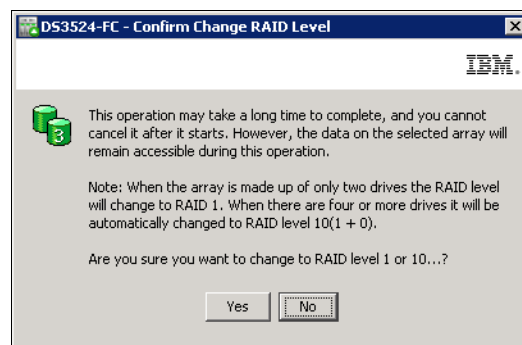


Figure 9-13 RAID level change confirmation

If you confirm, the process is started and cannot be stopped. All the logical drives of the array in the Logical pane tree are marked with clock icon. You cannot do certain other configuration changes on the array and logical drives until the process is completed.

### 9.3.3 Add Free Capacity (Drive)

In today's IT environment, the need for storage space grows constantly. Because it is essential that this process be non-disruptive and not cause any downtime, the ability to increase the available free capacity in an array without needing to restart the host system is an important feature. This feature is called Dynamic Capacity Expansion (DCE).

With DS Storage Manager, it is possible to add new disk drives to the storage subsystem and start the expansion procedure while the system remains fully operational. After the procedure starts, it cannot be stopped. This procedure might have a performance impact because the expansion process competes with normal disk access. When possible, this type of activity should be performed when I/O activity is at a minimum. The new free capacity can be used to create additional logical drives. Existing logical drives in the array do not increase in size as a result of this operation. Logical drive increasing is called *Dynamic Volume Expansion* (DVE), and is described in chapter 9.5.7, "Increase Capacity" on page 251.

**Note:** Storage Manager supports RAID 0 and 1 arrays with more than 30 drives. In certain DS3500 storage subsystem configurations, this can improve performance, provided that the system is optimally tuned. It also improves the data capacities of these arrays. RAID 1 or 10 requires an even number of disk drives.

**Attention:** It is still not possible to use more than 30 drives in RAID 3, 5, and 6 arrays. After the maximum number of drives is reached, you cannot add new drives.

To add new drives into an array, right-click the array and select **Add Free Capacity (Drives)**. In the Add Drives window (Figure 9-14), select one or two drives to be added to the array, depending on whether RAID 1 or 10 is being used by the array. The controller firmware enables a maximum of two drives to be added at one time, but this operation can be repeated later to add more drives to an array.

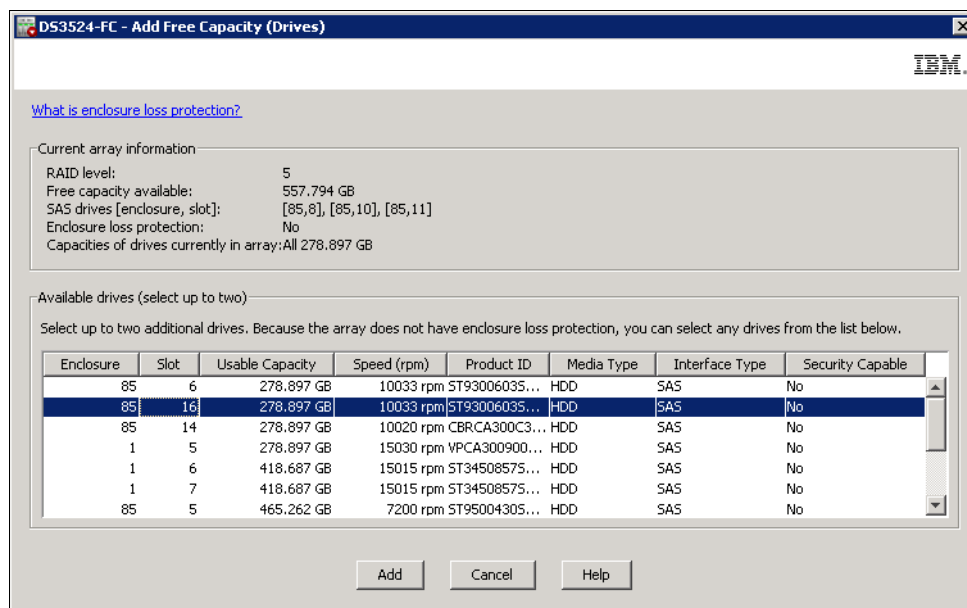


Figure 9-14 Adding a new drive to array



For RAID levels 3, 5, and 6, select one or two drives, and for RAID levels 1 and 10, select two drives. Pay attention to the following considerations:

- Only SAS drives are supported on DS3500, so every unconfigured disk drive with appropriate capacity can be listed as candidate to add.

**Note:** It is not a general rule to add a new SAS drive with slower rotational speed, as performance of the a new disk can slow down existing faster disks in the array. In the other case, a new disk can be under utilized if its speed is higher then speed of current disks in the array.

- Select drives that have a capacity equal to the current drive capacities in the array.  
Drives larger than the other drives participating in the array can be added, but avoid this because their usable capacity will be reduced so that they match the current drives capacities in the array.
- Do not mix regular drives with SED (encryptable) drives.

**Note:** It can be possible to mix standard and secure drives when encryption is not configured on the subsystem and SED disks are already inserted in DS3500 subsystem. However, this could later cause problems when Secure Drive functionality is activated.

After the procedure for adding drive is started, it cannot be stopped because the subsystem needs to redistribute the data contained in the array to all drives, including the new ones. There is a performance impact during this operation, but the logical drives of the array remain available to the host systems.

### 9.3.4 Secure Drive

The **Secure Drive** menu item is used for drive encryption functionality, and is fully described in Chapter 15, “Disk Security with Full Disk Encryption drives” on page 439.

### 9.3.5 Delete and Rename

Deleting the array will automatically delete all logical drives in this array (see Figure 9-15) if they are configured.

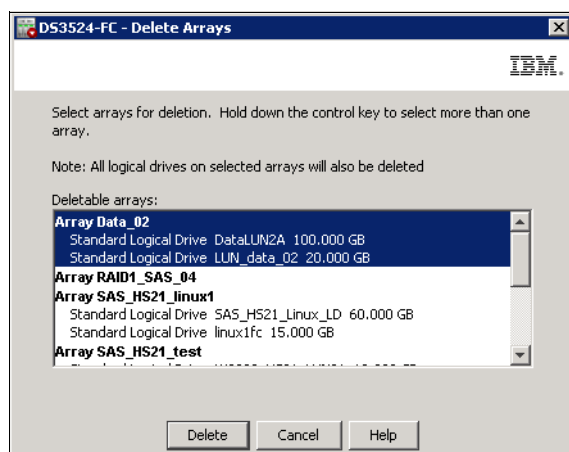


Figure 9-15 Array and logical drives deletion

If some of the LogicalDrive cannot be deleted (because of any reason), the array cannot be deleted.

If you are renaming the selected array (Figure 9-16), you must use a unused name. Names in DS Storage Manager are case sensitive, and an array name can consist of letters, numbers, and the special characters underscore (\_), hyphen (-), and pound (#). If you choose any other characters, an error message appears.




Figure 9-16 Rename array

### 9.3.6 Replace Drive

This menu item is active in the case that a disk drive in the array has failed or is missing. In this case you can replace the failed or missed drive with another one. This procedure is described in 10.4, “Failed disk drive replacement” on page 266.

## 9.4 Working with Free Capacity

The only way to use the Free Capacity of a specified array is to configure a new logical drive. The theoretical limit of logical drives per array is 256, but the practical limit is much fewer.

If you select Free Capacity line in the Logical pane tree as shown in Figure 9-17 on page 233, you can see the basic information about this capacity on the Properties pane. You can right-click the Free Capacity and select **View Associated Physical Components** to see which disk drives are associated with that free capacity.

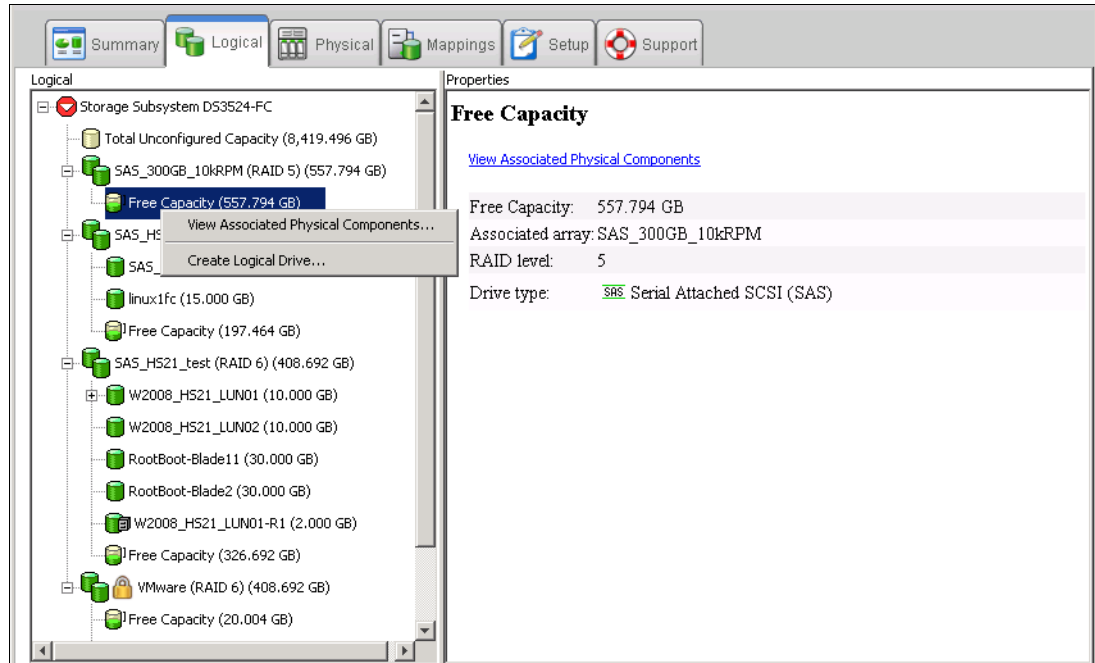


Figure 9-17 Free Capacity menu

### 9.4.1 Create logical drive

To create a logical drive (sometimes referred to as a Volume), perform the following steps:

1. Select required **Free Capacity** line in the array, where you want to create a new logical drive. Right-click the Free Capacity and select **Create Logical Drive**.
2. In the initial Create Logical Drive wizard window, read the introduction and click **Next** to start configuring the logical drive.
3. In the Specify Capacity/Name window as shown in Figure 9-18 on page 234, perform the following steps:
  - a. Enter the desired size capacity, which must be less than the available free capacity of the array. Choose this number appropriately if you want to define more than one logical drive in the array.

**Note:** Specify size by typing the required number even if you can use all the capacity of an array. The offered number could be rounded, and you might have troubles later when you need to define other logical drives with the same capacity (for example for copy services).

- b. Assign a unique name to the logical drive.
- c. If you want to change advanced logical drive settings, such as the segment size or cache settings, select the Customize settings check box and click **Next**.

Figure 9-18 Specifying logical drive capacity and name

4. The Customize Advanced Logical Drive Parameters window opens (Figure 9-19). You can set your new logical drive using any of the predefined I/O types listed, or manually set the cache read ahead multiplier, segment size, and controller ownership:

Figure 9-19 Customize logical drive parameters

- a. For logical drive I/O characteristics, you can specify file system, database, or multimedia defaults. The **Custom** option allows you to disable or enable the dynamic

cache read prefetch and the segment size. Table 9-1 shows the defaults predefined for each I/O type.

*Table 9-1 Logical drive defaults I/O characteristics*

| I/O type                         | File system | Database   | Multimedia |
|----------------------------------|-------------|------------|------------|
| Segment size                     | 128 K       | 128 K      | 256 K      |
| Modification priority            | High        | High       | High       |
| Read cache                       | Enable      | Enable     | Enable     |
| Write cache                      | Enable      | Enable     | Enable     |
| Write cache without batteries    | Disable     | Disable    | Disable    |
| Write cache with mirroring       | Enable      | Enable     | Enable     |
| Flush write cache after          | 10 seconds  | 10 seconds | 10 seconds |
| Dynamic cache prefetch           | Enable      | Enable     | Enable     |
| Enable background media scan     | Enable      | Enable     | Enable     |
| Media scan with redundancy check | Disable     | Disable    | Disable    |
| Pre-read redundancy check        | Disabled    | Disabled   | Disabled   |

- b. The segment size is chosen according to the usage pattern. For custom settings, you can directly define the segment size.
- c. Dynamic cache read prefetch is important for multimedia applications that use sequential I/O because it allows the controller, while it is reading and copying host-requested data blocks from disk, to copy additional data blocks into the cache.

**Note:** Generally, use the default settings: dynamic cache read prefetch enabled for all RAID types, and segment size 128 kB.

- d. The preferred controller handles the logical drive normally if both controllers and I/O paths are online. You can distribute your logical drives between both controllers to provide better load balancing between them. The default is to alternate the logical drives on the two controllers.

It is better to spread the logical drives by the load that they cause on the controller. If you do not know the expected access pattern for each logical drive, you can evaluate it by using the performance monitor option integrated with the Storage Manager. Based on data gathered from the Performance Monitor, move logical drives to the other preferred controller to balance the load if required. See Chapter 3, “IBM System Storage DS3500 Storage System planning tasks” on page 29 for more information.

- 5. Click **Next**, and the Specify Logical Drive-to-LUN Mapping window (Figure 9-20 on page 236) opens. This window allows you to choose between mapping your created logical drive to the default host group or to **Map later using the Mappings View**. If you select **Default mapping**, your new logical drive is mapped to the Default Group and is available to any host zoned to the DS3500 Storage Subsystem.

**Important:** Manually mapping logical drives to hosts prevents unwanted mapping and is always the best choice.

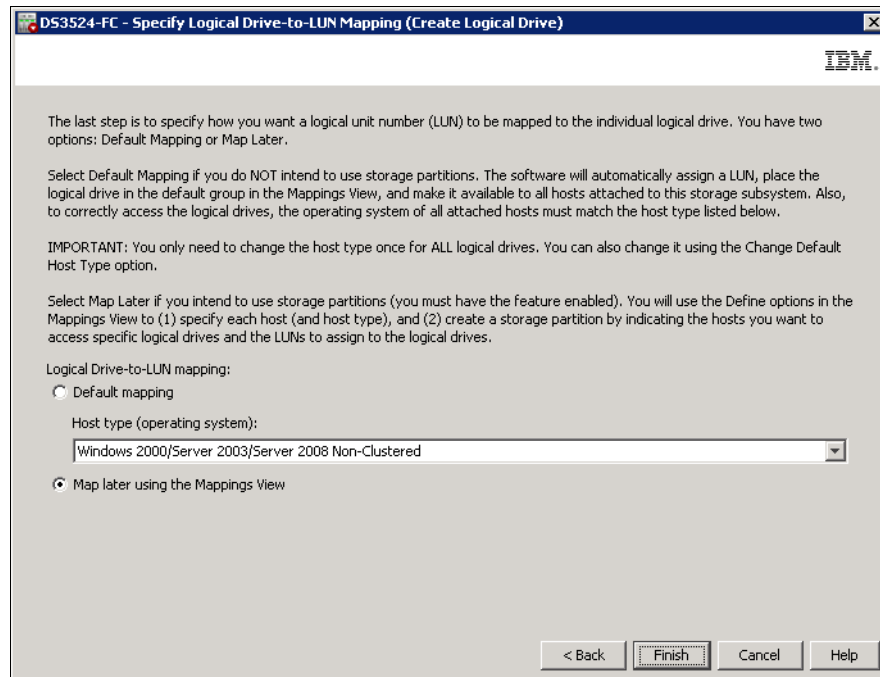


Figure 9-20 New logical drive mapping

- Click **Finish** and the Creation Successful window opens as shown in Figure 9-21. You are prompted with option to create another logical drive. Click **No** if you are finished creating logical drives or want to continue at a later time.




Figure 9-21 Create logical drive - Creation Successful

- The Completed window opens as shown in Figure 9-22 on page 237. Click **OK** to finish.

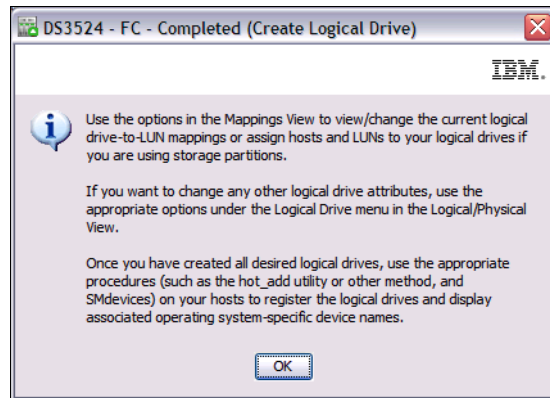


Figure 9-22 Create logical drive - Completed

8. The new logical drive will be formatted. In the Properties pane of Storage Manager you can check the progress of this operation (see Figure 9-23). The Logical Drive icon in the Logical pane is marked by clock symbol during this process.




Figure 9-23 Logical drive formatted

9. After the format is complete, the new logical drive is ready to be mapped to host. You can repeat the same process to create other arrays or logical drives.

## 9.5 Working with logical drives

The logical drive is the main building block of the DS3500 Storage Subsystem. Your host servers can only see these logical drives. The logical drive can be tuned using many parameters described in the following sections.

The basic information about a logical drive is displayed in the Properties pane of the Logical tab (see Figure 9-24 on page 238).

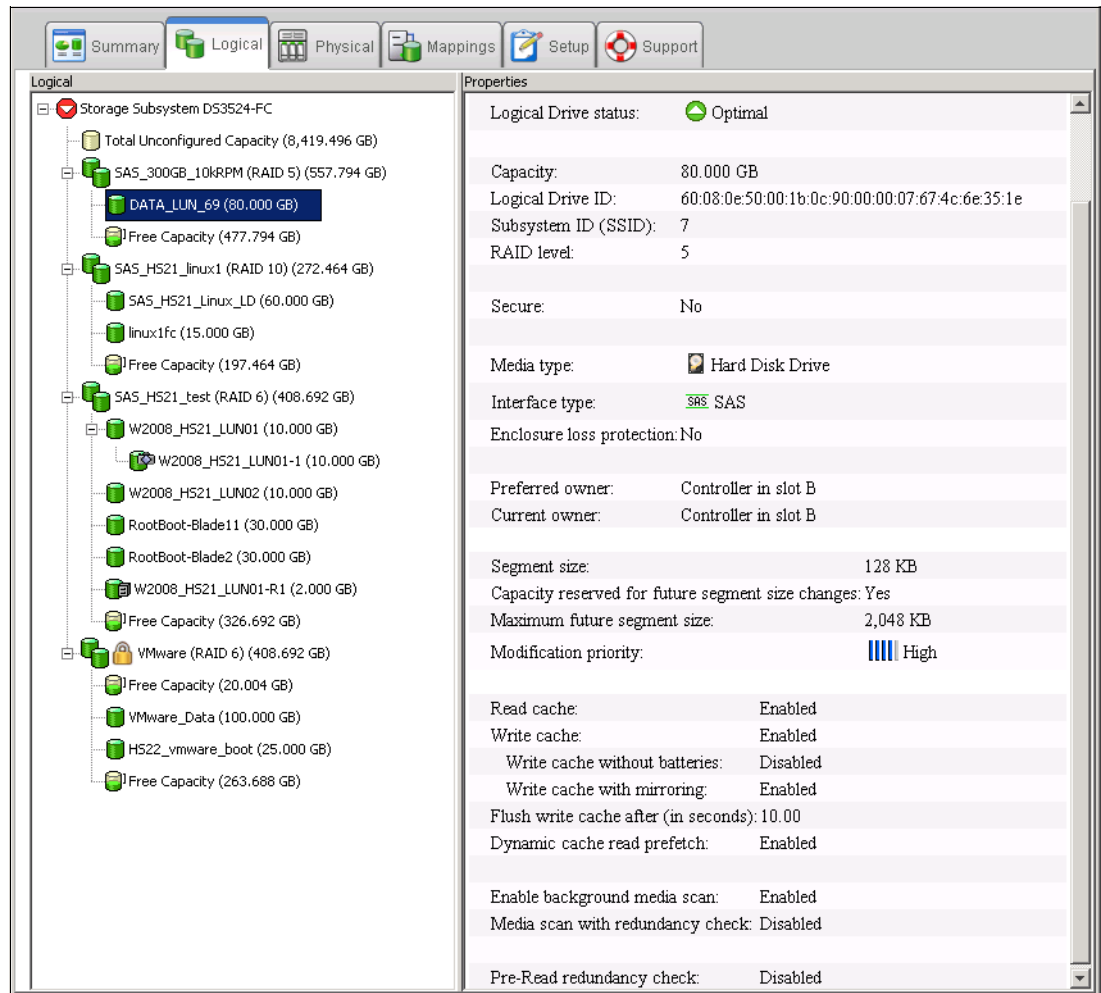


Figure 9-24 Logical Drive - Logical and Properties pane

Most of the following parameters can be modified for each regular logical drive separately:

- ▶ Capacity (can be increased)
- ▶ Controller ownership
- ▶ Segment size
- ▶ Modification priority
- ▶ Cache settings:
  - Read cache
  - Write cache
  - Write cache without batteries
  - Write cache with mirroring
  - Dynamic cache read prefetch
- ▶ Media scan
- ▶ Media scan with redundancy check
- ▶ Pre-read redundancy check

You can also set your regular logical drive as a source of advanced copy services, such as FlashCopy, VolumeCopy, and Enhanced Remote Mirroring in the Logical pane (Figure 9-24). These services are described in *IBM Midrange System Storage Copy Services Guide*, SG24-7822.



To start configuring the logical drive, right-click the logical drive you want to modify in the Logical pane tree and select **View Associated Physical Components**. You can also select the link on second line of the Properties pane. A window displays showing the components associated with the logical drive like the disk drives, location of the drives, the controller, and the controller that currently owns selected the logical drive. These physical components are marked by blue bullets.

The second option in the menu is **Change**. Using this option, you can modify all important parameters of the logical drive. Select this line and next menu opens as shown in Figure 9-25.

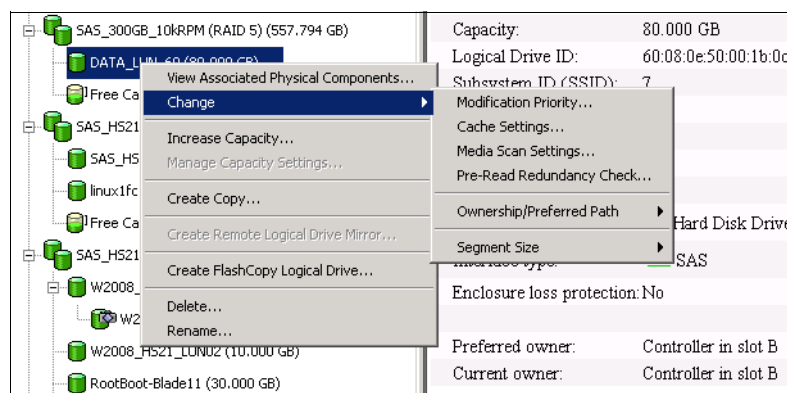


Figure 9-25 Logical Drive - Change menu

The possible changes of a logical drive parameters are described in the following sections:

- ▶ 9.5.1, “Change Modification Priority” on page 239
- ▶ 9.5.2, “Change Cache Settings” on page 241
- ▶ 9.5.3, “Change media scan settings” on page 243
- ▶ 9.5.4, “Change Pre-Read Redundancy Check” on page 246
- ▶ 9.5.5, “Change Ownership/Preferred Path” on page 247
- ▶ 9.5.6, “Change Segment Size” on page 248

## 9.5.1 Change Modification Priority

The Modification Priority defines how much processing time is allocated for operations modifying the logical drive relative to the system performance. Operations that cause a logical drive modification are:

- ▶ Initializing a logical drive
- ▶ Reconstructing after a disk failure
- ▶ Copying back from a hot spare drive
- ▶ Changing the segment size of a logical drive
- ▶ Expanding a logical drive
- ▶ Adding free capacity to an array
- ▶ Defragmenting an array
- ▶ Changing the RAID level of an array

If the logical drive contains critical data, you can set a high modification priority to keep the time of a critical state (for example, after losing a disk) as short as possible, even if this affects the system performance during the modification process.

The following modification priority rates are available:

- ▶ Lowest
- ▶ Low

- ▶ Medium
- ▶ High
- ▶ Highest

**Note:** The lowest priority rate favors system performance, but the modification operation takes longer. The highest priority rate favors the modification operation, but system performance might be compromised.

To change the modification priority, perform the following steps:

1. Select a logical drive, right-click it, and select **Change** → **Modification Priority**.
2. Make sure that the correct logical drive is selected, and set the new Modification Priority value as shown in Figure 9-26.




Figure 9-26 Modification priority for a logical drive

3. Click **OK** and then **Yes** in next window to confirm this change.
4. You will then see the completion window as shown in Figure 9-27

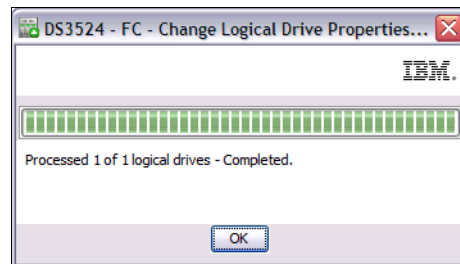


Figure 9-27 Change Modification priority - Completed

You can select more drives in Figure 9-26 by using the **Ctrl** or **Shift** key during selection or by clicking the **Select All** button. All of the selected drives will have the same modification priority after the change.

## 9.5.2 Change Cache Settings

Using the next option in the menu (Figure 9-25 on page 239), you can change the cache settings on the logical drive.

- ▶ You can specify the DS3500 subsystem global cache settings as described in “Cache Settings” on page 199. These settings will affect all arrays and logical drives created on the subsystem:
  - Start and stop cache flushing levels
  - Cache block size
- ▶ You can also specify the following settings per logical drive:
  - Read cache
  - Dynamic cache read prefetch
  - Write cache
  - Write cache with mirroring
  - Write cache without battery

Figure 9-28 shows the typical values when you create a logical drive with default values. With the Storage Manager, cache settings can be specified for each logical drive independently, giving greater flexibility.




Figure 9-28 Default values for logical drives

**Note:** We recommend that the values are manually set during creation to suit the performance needs of the logical drive. These settings can be changed after logical drive creation for tuning purposes.

To change this setting, highlight your logical drive in the Logical pane tree, right-click it, and select **Change** → **Cache Settings**. The window shown in Figure 9-29 on page 242 appears. The logical drive you selected is highlighted.

You can select more logical drives in Figure 9-29 on page 242, using the **Ctrl** or **Shift** key during selection or by clicking the **Select All** button. All of the selected drives will have the same modification priority after the change.

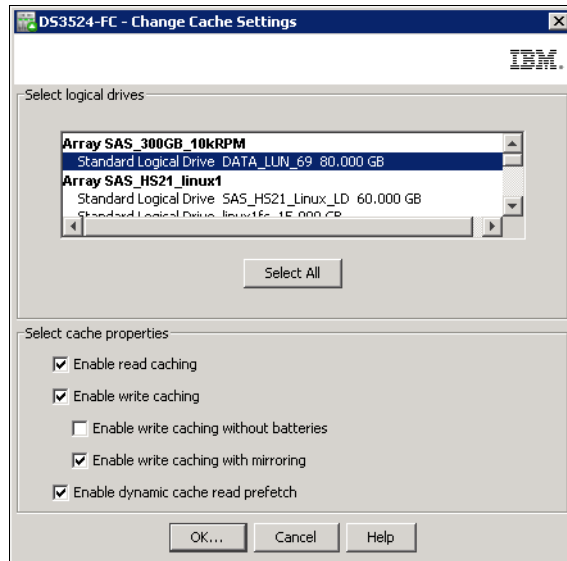


Figure 9-29 Logical Drive cache settings - defaults

These cache settings can have a large impact on the performance of the DS3500 storage subsystem and on the availability of data. Be aware that performance and availability will conflict with each other. If maximum performance is required, in most cases availability might have to be compromised and vice versa.

The default settings are read and write cache for all logical drives, with cache mirroring to the alternate controller for all write data. The write cache is only used if the cache battery is fully charged. Dynamic cache read prefetch is also used on new logical drives by default.

## Read caching

Read caching allows you to keep data in controller cache memory even if they were already sent to host. If a host initially requests data that are not in the cache, the controller reads the needed data blocks from the disks and places them in the cache. Until the cache is flushed, any other requests for this data are fulfilled with the cache data instead of initiating another read operation to the disks.

## Write caching

The write caching enables the storage subsystem to cache write data instead of writing it directly to the disks. This can improve performance significantly, especially for environments with random writes such as databases. For sequential writes, the performance gain varies with the size of the data written. If the logical drive is only used for read access, it might improve overall performance to disable the write cache for this logical drive.

## Write cache with mirroring

The DS3500 storage subsystem write cache mirroring provides the integrity of cached data if a RAID controller fails. This is an excellent feature from a high availability perspective, but it decreases performance because cache mirroring occupies twice as much cache memory for writing. Logical drives that are used to store temporary data can have cache mirroring disabled, but be sure you can regenerate the data again easily because you can lose them if a controller fails.

By default, the write cache is always mirrored to the other controller to ensure proper contents, even if the logical drive moves to the other controller. Otherwise, the data of the logical drive can be corrupted if the logical drive is shifted to the other controller and the

cache still contains unwritten data. If you turn off this parameter, you risk data loss in the case of a controller failover, which could also be caused by a path failure in your fabric.

The cache of the DS3500 storage subsystem is protected by a battery against power loss. If the batteries fail or are not fully charged, for example, just after powering on, the controllers automatically disable the write cache. If you enable the parameter, the write cache is used even if no battery backup is available, resulting in a higher risk of data loss. Controller write cache mirroring should be left enabled for data integrity reasons in case of a controller failure.

### **Write-through**

Write-through means that writing operations do not use cache at all. The data is always going to be written directly to the disk drives. Disabling write caching frees up cache for reading because the cache is shared for read and write operations.

Write caching can increase the performance of write operations. The data is not written straight to the disk drives: it is only written to the cache. From an application perspective, this is much faster than waiting for the disk write operation to complete. Therefore, a significant gain in application writing performance can be expected. It is the responsibility of the cache controller to eventually flush the unwritten cache entries to the disk drives.

Write cache mode (Write-Back) appears to be faster than Write-Through mode because it increases the performance of both reads and writes. But this is not always true because it depends on the disk access pattern and workload.

A lightly loaded disk subsystem usually works faster in write-back cache mode, but when the workload is high, the write cache can become inefficient. As soon as the data is written to the cache, it has to be flushed to the disks to make room for new data arriving into cache. The controller performs faster if the data goes directly to the disks. In this case, writing data to the cache can be an unnecessary step that decreases throughput.

### **Dynamic cache read prefetch**

Cache read-ahead, or “prefetch,” allows the controller, while it is reading and copying host-requested data blocks from disks into the cache, to copy additional data blocks into the cache. This increases the chance that a future request for data will be fulfilled from the cache. Cache read prefetch is important for multimedia applications that use sequential I/O.

This feature uses an automatic pre-fetching multiplier to maximize its cache hit efficiency and system performance. This will turn on monitoring of the I/O to the logical drive and enable the new algorithm to dynamically choose how much to read ahead. This simplifies the process for the administrator because there is no need to manually set a specific value for the read ahead multiplier. The system will tune itself depending on the I/O characteristics. When sequential access is detected, the controller will automatically start using read ahead buffering. When random or non-sequential I/O is used, then it will stop using the read ahead buffer. To disable this feature, simply click to clear the **Dynamic Cache Read Prefetch** check box for the relevant logical drive.

## **9.5.3 Change media scan settings**

Media scan is a background process, enabled by default, that checks logical drives over hard disk drives 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 could eventually disrupt normal data reads or writes. This process is sometimes known as *data scrubbing*.

Media scan is an option available for logical drive space configured on hard disk drives. Unused hard disks or hot spares are not scanned.

The media scan runs on all logical drives in the storage subsystem that meet the following conditions:

1. The logical drive is in an optimal status.
2. There are no modification operations in progress.
3. The Media Scan parameter is enabled.

The media scan continuously runs in the background using spare cycles to complete its work. The default media scan is for a scan every 30 days, that is, the maximum time the media scan has to complete the task. During the scan process, the DS3500 storage subsystem 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. After the media scan has completed, it starts over again and resets its time for completion to the current setting. This media scan setting can be reduced. However, if the setting is too low, priority is given to the media scan over host activity to ensure that the scan completes in the allocated time. This scan can impact performance, but will improve data integrity in the long term.

The media scan is enabled for the entire storage subsystem. The system-wide enabling specifies the duration over which the media scan runs. By default, the media scan process runs without checking redundancy data. You can optionally specify whether to do a redundancy check or to stop the media scan.

A **Media Scan** can be considered a surface scan of the hard drives, and a **Redundancy Check** scans the data blocks of a RAID 3, 5, or 6 logical drive and compares it against the redundancy data calculated. In the case of a RAID 1 logical drive, the redundancy scan compares blocks between copies on mirrored drives.

**Note:** A media scan is only capable of resolving media errors and data or parity mismatches. A media scan does not attempt to resolve any other sort of error occurring during I/O operations.

We have seen no effect on I/O when we use a 30-day setting unless the processor is utilized in excess of 95%. The length of time that it takes to scan the logical drives depends on the capacity of all the logical drives on the system and the utilization of the controller.

**Important:** The media scan must be enabled for the entire storage subsystem and enabled on each logical drive within the storage subsystem to protect the logical drive from failure due to media errors. This is the default and recommended configuration.

Table 9-2 on page 245 shows the errors and describes several of the actions that the DS3500 storage subsystem takes as a result of a media scan and redundancy check operations.

Table 9-2 Media scan errors

| Reported error          | Description                                                                                           | Result                                                                                                                                  |
|-------------------------|-------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| Unrecovered media error | The data cannot 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 cannot 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 cannot be read, and parity or redundancy information cannot be used to regenerate it.        | An error is reported.                                                                                                                   |

A media scan is a long running process with the lowest priority, and so you can only run a media scan on logical drives that meet the following conditions:

- ▶ Optimal status
- ▶ No modification operation in progress

To locate this setting, highlight the logical drive in the Logical pane, right-click it and select **Change** → **Media Scan Settings**. The window shown in Figure 9-30 opens.

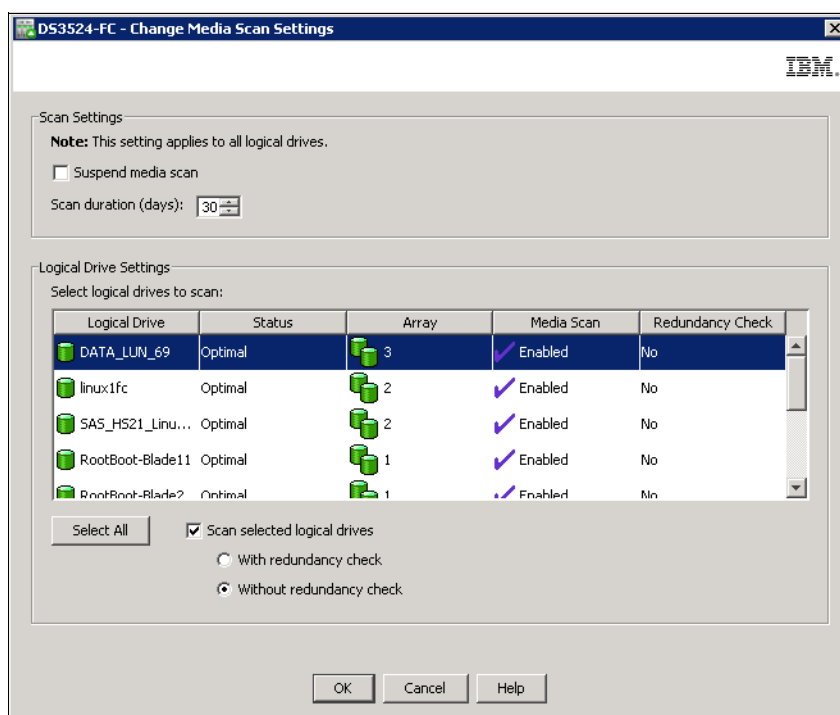


Figure 9-30 Media Scan Settings for Logical Drive

The Change Media Scan Settings window has two sections. In the first, you can suspend media scan or set the scan duration for entire storage subsystem, 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 (with help of **Ctrl** or **Shift** keys) one or more logical drives you want to set.
2. To enable the media scan, select the **Scan selected logical drives** check box, select with or without redundancy check, and click **OK**.
3. The changes will be saved on the storage subsystem, and you are informed when it is completed.

## 9.5.4 Change Pre-Read Redundancy Check

You can use this option to define the capability of the DS3500 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 or Remote Mirror target 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, FlashCopy repository drives, and Remote Mirror source drive.

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 9.5.3, “Change media scan settings” on page 243). 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.

To locate this setting, perform the following steps:

1. Highlight the logical drive in the Logical pane tree, and right-click it and select **Change** → **Media Scan Settings**. The window shown in Figure 9-31 opens.

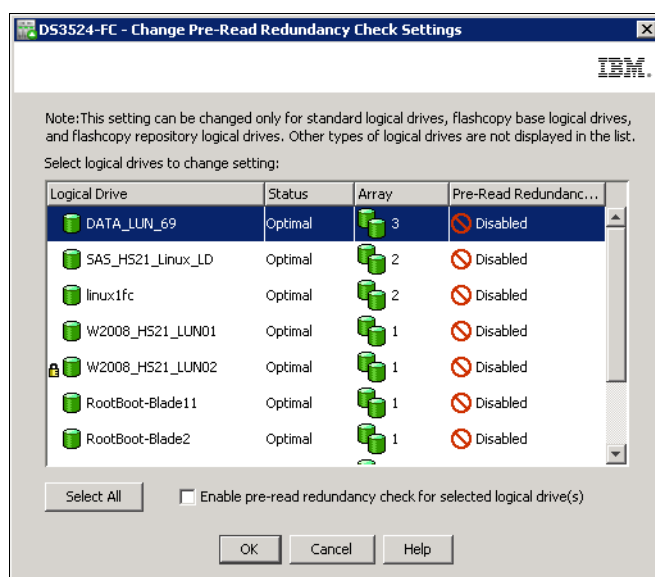


Figure 9-31 Change Pre-Read Redundancy Check Settings

2. Select the logical drives that are required for a pre-read redundancy check. You can use the **Ctrl** or **Shift** key to select part of the list, or click **Select All** to select all of them. Select the Enable pre-redundancy check box and click **OK** to activate this feature.
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.

### 9.5.5 Change Ownership/Preferred Path

If you are using a dual controller version of the DS3500, 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. The alternate controller only takes over and handles the I/O requests in case of a failure along the I/O path, for example, a defective host bus adapter or switch. When defining logical drives, the system normally alternates ownership between the two controllers as they are defined.

Situations can occur when all heavily stressed logical drives reside on only one controller and the other one handles only a small amount of all I/O requests. To balance the workload between the controllers, the preferred ownership of a logical drive can be changed to the other controller.

**Important:** Be sure that the operating system using the logical drive uses a multipath I/O driver; otherwise, it will lose access to the logical drive if ownership is changed.

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, typically, the load across controllers and logical drives are constantly changing. The logical drives and data can be accessed at any time depending on which applications and users are active during that time period, which is why monitoring the system is important.

The Performance Monitor provides data that is useful for monitoring the I/O activity of a specific controller and a specific logical drive, which can help identify high-traffic I/O areas. Identify actual I/O patterns in the individual logical drives and compare them with the expectations for an application. If a particular controller has considerably more I/O activity, consider moving logical drives to the other controller in the storage system.

You might find a disparity in the total I/Os (workload) of controllers. For example, the workload of one controller is heavy or is increasing over time, and that of the other controller is lighter or more stable. In this case, consider changing the controller ownership of one or more logical drives to the controller with the lighter workload.

**Tip:** Here are guidelines for logical drives assignment and storage partitioning:

- ▶ Assign defined logical drives evenly across all controllers to balance controller utilization.
- ▶ Use the manual method of creating logical drives. This allows for greater flexibility of configuration settings, such as enclosure loss protection and utilizing both drive loops.
- ▶ If some logical drives are highly utilized, when possible, separate them by putting them on their own or another array. This will reduce disk contention for that 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 has to be replaced, the controller should be disabled first. This will intentionally cause a failover of LUNs to the other controller and allow the removal and replacement of the preferred controller. 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 email and SNMP alert destinations.

To change the preferred ownership from one controller to the other, highlight the logical drive, right-click it, and select **Change** → **Ownership/Preferred Path**. Then select the controller to which the logical drive is to be moved. Depending on the current workload, the operation can take a long time to finish.

## 9.5.6 Change Segment Size

A segment, in a logical drive, is the amount of data, in kilobytes, that the controller writes on a single physical drive before writing data on the next physical drive. Depending on your data structure, you can optimize the creation of your logical drives for this value.

When you create a logical drive, the default segment size choose is 128 KB unless you are creating logical drives under a RAID3, or choosing the multimedia type for the logical drive, in which case a segment size of 256 KB is used.

The choice of a segment size can have a major influence on performance in both IOPS and throughput. You can choose the defaults, which are a good choice for general usage, or consider the following to properly select a logical drive segment size customized for your needs:

- ▶ The characteristics of your application, IOPS demanding, or Throughput (MBps) demanding, whether random or sequential.
- ▶ The I/O block size of the host that will use the logical drive.
  - Small IO sizes allow for greater transaction (IOPS) rates.
  - Large IO sizes allows a better transfer rate in MBps.
- ▶ The RAID protection and number of disk drives that are participating in the logical drive array.

To calculate the best segment size in this environment, consider the RAID protection and number of drives participating in the array.

### ***Segment Size > I/O block size request***

This is the best option for high IOPS and random requests.

- ▶ Several transactions are satisfied in a single operation to a disk.
- ▶ There is higher IOPS.
- ▶ It is ideal for random I/O requests, such as the database file system.
- ▶ The best option is to start with Segment size  $\geq 2 \times$  Block Size for high IOPS.

### ***Segment Size = I/O block size request***

It is difficult to align segment size with block size, so this option is not practical to implement, so use the other two scenarios.

- ▶ Every request transaction uses exactly one disk operation.
- ▶ There is high IOPS.
- ▶ This option is ideal for random I/O requests, such as the database file system.

### ***Segment Size < I/O block size request***

This is the best option to obtain MBps performance and low IOPS, which is the norm for multimedia applications.

- ▶ More disks are used or requested.
- ▶ There is higher throughput (MBps).
- ▶ It is ideal for sequential writes, such as a large multimedia application.
- ▶ It is optimized when a single I/O request can be serviced with a single or exact multiple data stripes (the segment size multiplied by the number of drives in the array that are used for I/O). In this case, multiple disks are used for the same request, but each disk is only accessed once, or the exact number of times if the block size is too big and you have to use multiple stripes.

To calculate the best segment size in this environment, consider the RAID protection and number of drives participating in the array by using the following formula:

Segment size = (Block size / X), where X is the number of drives used in the array for I/O and where:

- RAID 0: Number of drives

- RAID 1 or RAID 10: Number of drives / 2
- RAID 5: Number of drives - 1
- RAID 6: Number of drives - 2

If the resulting Segment Size is >512 KB, then divide it by an integer number starting with 2 to obtain multiple data stripes of each block requested.

**Tips:** The possible segment sizes available are 8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, and 512 KB.

- Storage Manager sets a default block size of 128 KB for every logical volume except for RAID3 volumes, which are set to 256 KB.
- For database applications, block sizes between 32–128 KB have shown to be more effective.
- In a large file environment, such as media streaming or CAD, block sizes of 128 KB and more work best.
- For a web server or file and print server, the range should be between 16–64 KB.

The performance monitor can be used to evaluate how a given segment size affects the workload.

**Note:** A performance testing schedule should be undertaken in the environment before going into production with a given segment size. Segment size can be dynamically changed, but only by rewriting the data, which consumes bandwidth and impacts performance. Plan this configuration carefully to avoid having to reconfigure the options.

To initiate segment size change, follow these steps:

1. Highlight the logical drive in the Logical pane, right-click it, and select **Change** → **Segment Size** → **value**. The *value* highlighted in menu is either twice or half of the current value. See example in Figure 9-32.

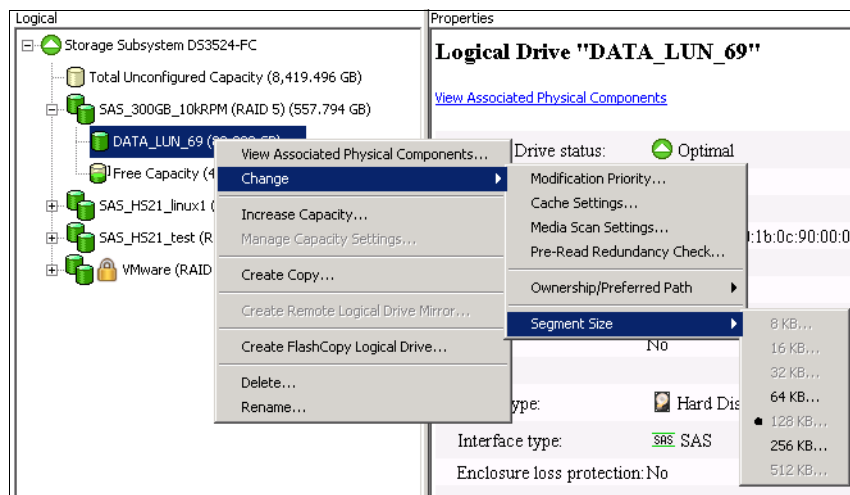


Figure 9-32 Change Segment Size menu

2. If you select a new segment size, a warning window shown in Figure 9-33 on page 251 opens.




Figure 9-33 Segment Size warning

Segment size change is long term process that can also influence the IO operations of attached hosts. The process cannot be stopped. If you can change segment size in multiplies of four or more, you must do it in more steps. Sometimes it can be faster to delete the logical drive and recreate it with your required segment size.

3. Select **Yes** and the logical drive starts to be modified. You can see the clock icon on it and progress bar in the Properties pane.

## 9.5.7 Increase Capacity

It is also possible to increase the size of the logical drive. This action is called *Dynamic Volume Expansion* (DVE). The capacity of standard logical drives and FlashCopy repository logical drives can be increased using one or both of the following capacities:

- ▶ Free capacity available on the array
- ▶ Unconfigured capacity (in the form of unused drives) on the storage subsystem

Storage Manager provides support for logical drives larger than 2 TB. This improves the capacity requirement for applications that require large capacity logical drives.

**Note:** Increasing the capacity of a standard logical drive is only supported on certain operating systems. If the logical drive capacity is increased on a host operating system that is not supported, the expanded capacity will be unusable and the original logical drive capacity will not be able to be restored.

The operating systems that support a dynamic increase of capacity in a mapped logical drive are:

- ▶ AIX
- ▶ Linux
- ▶ NetWare
- ▶ Solaris
- ▶ Windows Dynamic Disks
- ▶ Windows Basic Disks

**Tip:** If a Logical Drive-to-LUN mapping has not yet been defined, it is possible to increase the capacity for a standard logical drive on any host operating system because the logical drive is not formatted yet in operating system.

The storage capacity of a standard logical drive cannot be increased if:

- ▶ One or more hot spare drives are in use in the array that it resides on.
- ▶ The logical drive has *Non-Optimal* status.

- ▶ Any logical drive in the array is in any state of modification.
- ▶ The controller that owns this logical drive is in the process of adding capacity to another logical drive (each controller can add capacity to only one logical drive at a time).
- ▶ No free capacity exists in the array and no unconfigured capacity (in the form of drives) is available to be added to the array.

Increase the logical drive capacity using the following steps:

1. Highlight the logical drive to be expanded, right-click it, and select **Increase Capacity**. A window with instructions (Figure 9-34) opens.

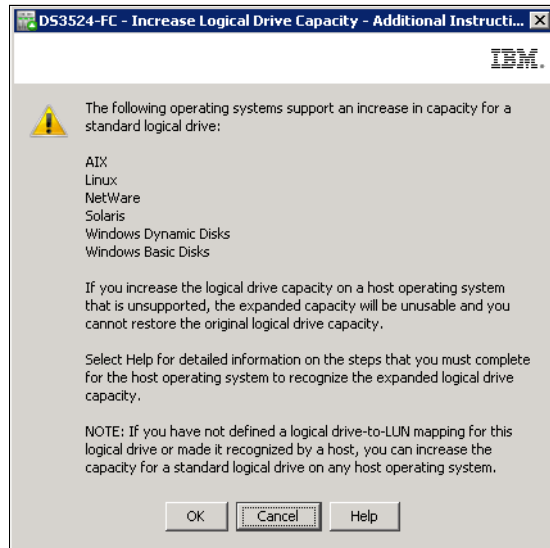


Figure 9-34 Dynamic logical drive expansion

2. Click **OK** and in the Increase Logical Drive Capacity window (Figure 9-35 on page 253), enter the amount of space by which the logical drive will be enlarged.

In the top part of the window, the current size of the logical drive and the available free capacity in the array is displayed. If no free configured space is available, but there are unassigned drives, these can be added from this same window by clicking **Add Drives** before proceeding to enlarge the logical drive. This will perform the same process described in 9.3.3, “Add Free Capacity (Drive)” on page 230.

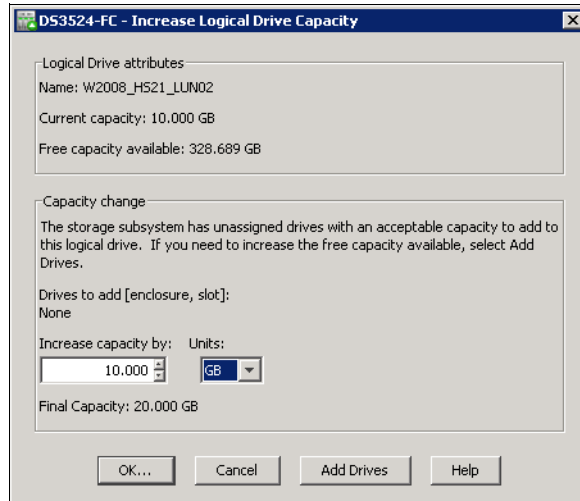


Figure 9-35 Increase Capacity - set amount

**Note:** If the RAID level of the array is 3, 5, or 6, and the drive enclosure has enclosure loss protection, the **Add Drives** option displays only drives that ensure enclosure loss protection. If the RAID level is 1 or 10, a minimum of two drives must be added.

3. Click **OK** after selecting the capacity to add. A warning message appears indicating that this operation cannot be stopped after it is started and that it might take a long time to complete. However, the data on the selected logical drive and all other logical drives on this array (if new drives have been added) remains accessible during this time. As with all operations requiring a redistribution of the data on the physical disks, the procedure might affect the performance. From the host operating system, the administrator will then have to perform a procedure to utilize the newly allocated space.
4. After a couple of seconds, the operation starts, the data on all following logical drives in array are redistributed, and the affected logical drives are marked with a small clock icon. You can see progress of operation in the Properties pane (Figure 9-36).




Figure 9-36 Increase Capacity - progress

## 9.5.8 Copy Services operations

If you select a logical drive and right-click, you can see four menu items related to Copy Services (Figure 9-37). Items might be unavailable depending on the current Premium Feature enablement.

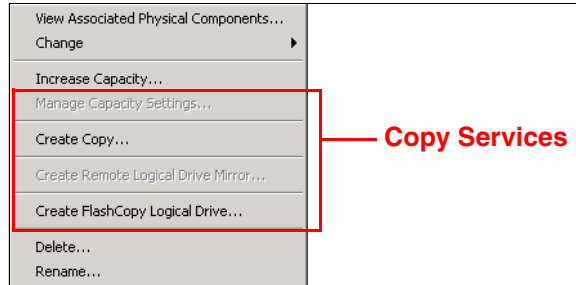


Figure 9-37 Logical Drive - Copy Services

The four options are:

- ▶ Manage Capacity Settings: Sets FlashCopy repository parameters
- ▶ Create Copy: Creates a VolumeCopy pair
- ▶ Create Remote Logical Drive Mirror: Creates a Remote Mirror pair
- ▶ Create FlashCopy Logical Drive: Creates a FlashCopy

In this book we describe operations and configurations related to standard logical drives, but in Logical pane you can find also special logical drives used for Advanced Copy Services. All the settings related to Copy Services can be found in *IBM Midrange System Storage Copy Services Guide*, SG24-7822.

## 9.5.9 Delete and Rename

To rename a logical drive you must use a unique name that has not been used for any other object in the managed DS3500 Storage Subsystem. Renaming has no influence on IO operations and connection to hosts. However, renaming will impact the SMcli scripts that might be used because logical drive related commands uses the name of the logical drive in syntax.

When you delete a logical drive, the free space remains as a separate island as shown in Figure 9-38 on page 255. This free space can be used to create a new logical drive of equal or smaller size than this free space. Defragmenting the array can eliminate these free capacity islands. See Chapter 14, “Administration - Support” on page 329 for more information.



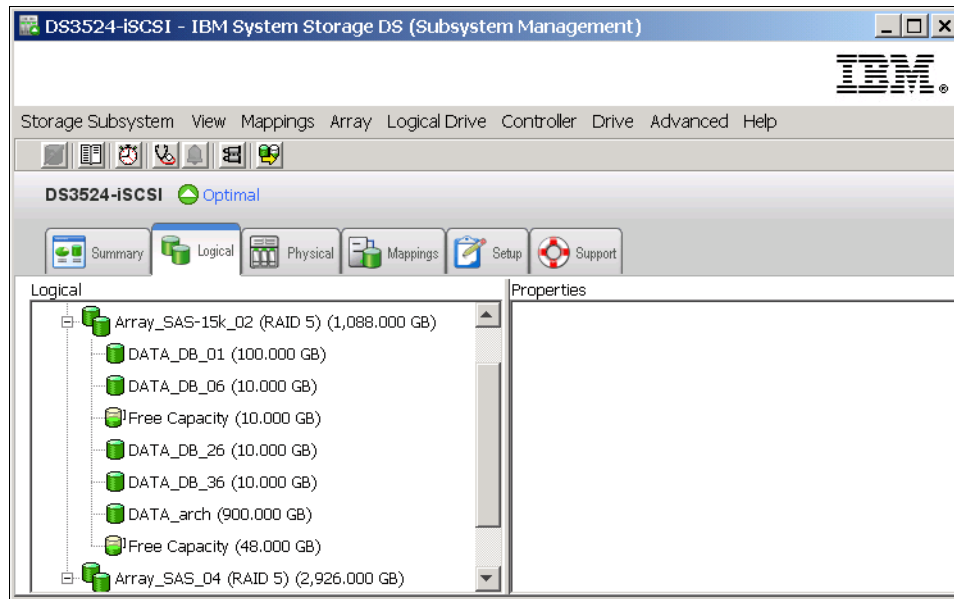


Figure 9-38 Free Capacity island in array





## Administration - Physical Tab

We have described how to configure logical objects (arrays and logical drives) of the DS3500 Storage Subsystem in Chapter 9, “Administration - Logical Tab” on page 219. The next tab in DS Storage Manger is the Physical tab (Figure 10-1 on page 258). Here you can find options to configure selected hardware components and see their properties.

## 10.1 Physical tab

The Physical tab (Figure 10-1) shows the organization of the storage subsystem by controllers, disk drives, and other hardware components. This tab is divided into two panes: the Physical pane on the left and the Properties pane on the right.

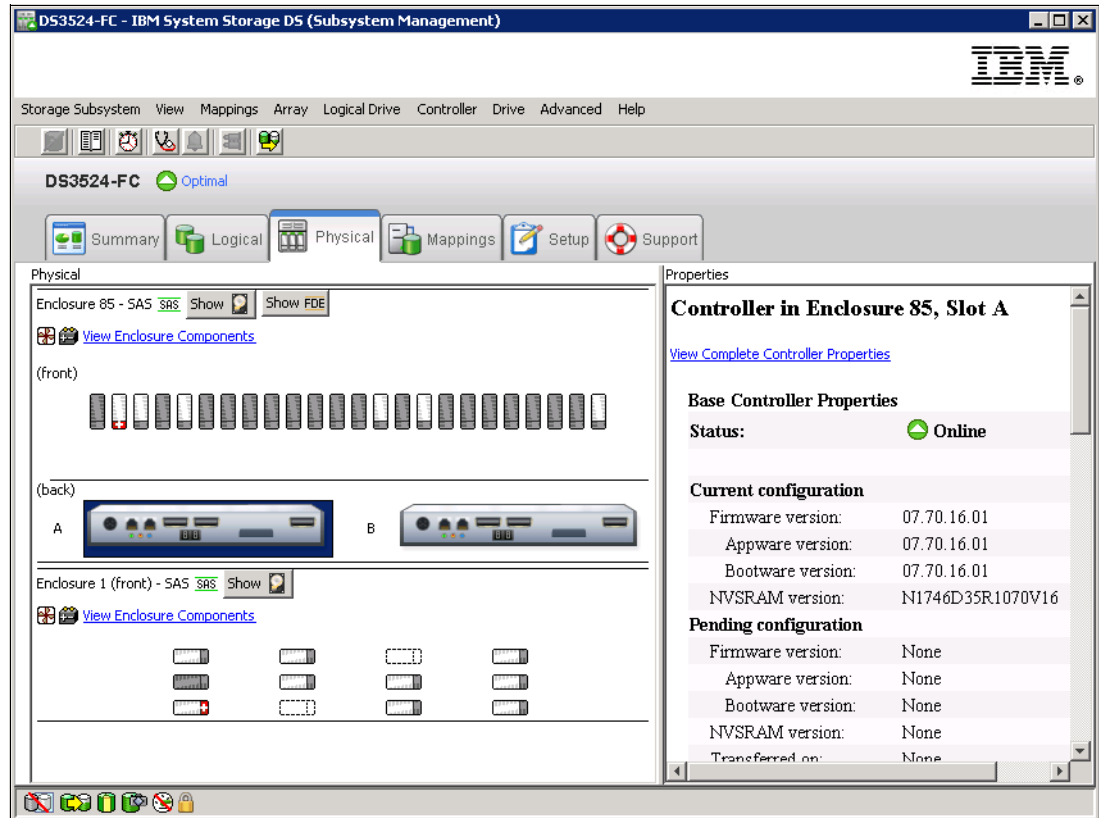


Figure 10-1 The Physical tab

The *Physical* pane shows a graphical representation of the physical components in the storage subsystem such as the controller modules and the drive expansion enclosures. You can configure hardware related parameters here.

The *Properties* pane displays detailed information about the component that is selected in the Physical pane. As in Logical tab, the Properties pane shows the same information that you can find in the storage subsystem profile.

Because there is nothing to configure in the Properties pane, we will focus on the Physical pane in the following sections:

- ▶ 10.2, "Discover component properties and location" on page 259
- ▶ 10.3, "Set hot spare drive" on page 262
- ▶ 10.4, "Failed disk drive replacement" on page 266
- ▶ 10.5, "Set preferred loop ID" on page 268
- ▶ 10.6, "Set remote access" on page 270
- ▶ 10.7, "Set Ethernet management ports" on page 270
- ▶ 13.4, "Configure iSCSI Host Ports" on page 317

The other configuration options in the Physical pane are already described in previous chapters and we will provide the references as needed.

## 10.2 Discover component properties and location

In the left Physical pane you can see schematic layout of physical enclosures of DS3500 Storage Subsystem. In Figure 10-1 on page 258, a subsystem with DS3524 controller enclosure for 24x 2.5 inch disk drives, and one attached EXP3512 expansion enclosure for 12x 3.5 inch disk drives are displayed. You can align this layout to the actual physical placement in your rack as described in “Enclosure Order” on page 202.

### 10.2.1 Show disks type

In the upper part of the Physical pane you can see one or two buttons. The first one is for selection of SAS disks, and the second one is for SED (Self Encrypted Disk) SAS disks. Encrypted disks are also known as FDE (Full Disk Encryption), so the button name is **Show FDE**. At the time of writing this book, only SAS drives are available for DS3500, so if you click the **Show SAS disk** button, all disks are marked as in Figure 10-2.




Figure 10-2 SAS disks selection

If you click the **Show FDE** button, disks that were selected as SAS are marked again as you can see in Figure 10-3.

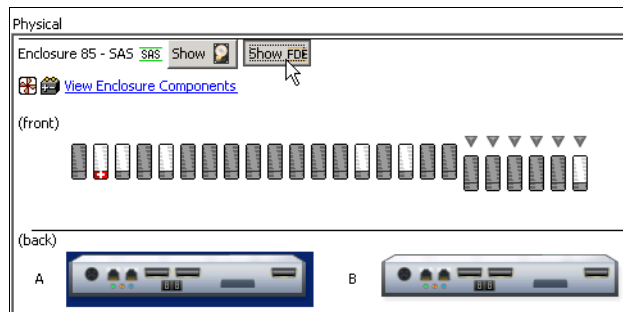


Figure 10-3 SED disks selection

More information about disk encryption can be found in Chapter 15, “Disk Security with Full Disk Encryption drives” on page 439.

### 10.2.2 View Enclosure Components

The second line in the Physical pane is a link that opens a new window with detailed information about hardware components of specific enclosure (Figure 10-4 on page 260). If you select the buttons on the bottom of the window, you can also locate this enclosure in rack (blue LED diode starts to blink) or you can change the enclosure ID. Refer to “Enclosure ID” on page 202 for more details.

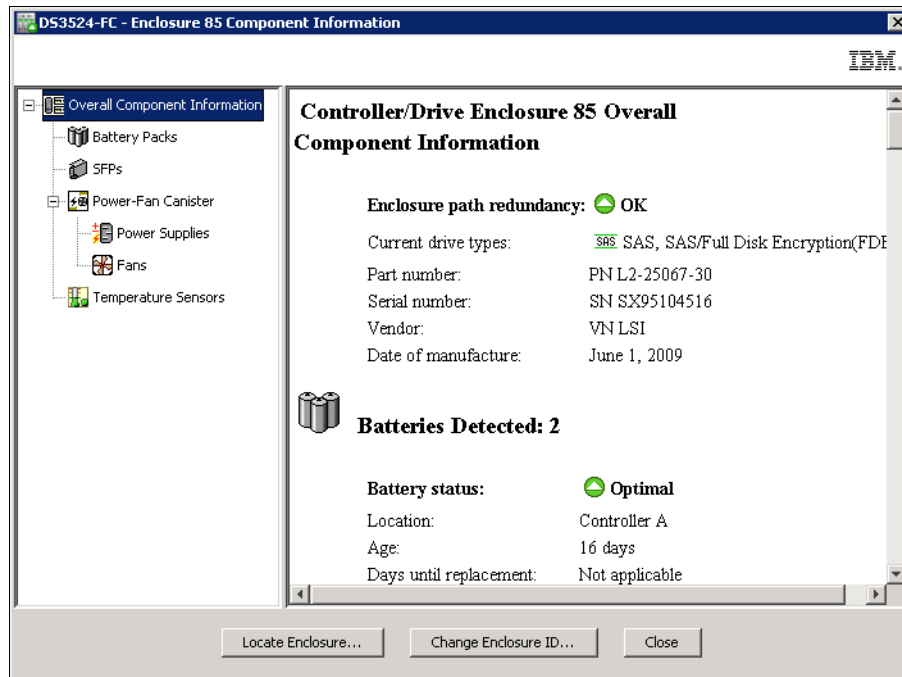


Figure 10-4 Enclosure Properties

The remaining physical components in the Physical pane are disks and controllers. If you select a disk or a controller, you will get relevant information about that disk or controller in the Properties pane.

### 10.2.3 Disk Drive menu

If you right-click a disk drive, the menu for other configuration options appears. Certain items in menu (Figure 10-5) might not be available because not all options are possible for all disk drive types. This also depends on status of disk drive and how the disk is configured.




Figure 10-5 Disk drive menu

The two menu options we describe in separate sections in this chapter are:

- ▶ Hot Spare Coverage in 10.3, “Set hot spare drive” on page 262
- ▶ Replace Drive in 10.4, “Failed disk drive replacement” on page 266

The other menu items will be covered in the following sections.

#### View Associated Logical Element

If you select this menu item, you can see related logical objects as shown in Figure 10-6 on page 261 such as arrays and logical drives that are using the selected disk drive.

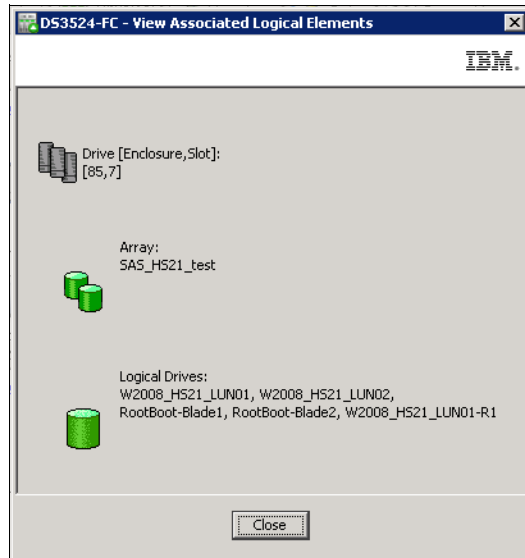


Figure 10-6 Associated Logical Elements

### Create Array and Create Logical Drive (manual mode)

These two menu options take you to a window for creating an array or a logical drive in manual mode. You can further select multiple unconfigured disk drives in the Physical pane using the **Ctrl** key before selecting this option. If you right-click one of selected disks, and select the **Create Array** menu item, the same initiation window will open as shown in Figure 9-6 on page 224. In this window you can see, manual mode is preselected. Follow the instructions described in 9.2.2, “Create array” on page 223, and when you reach the window for disk drives selection, you will recognize your planned disks are already selected for the new array.

The **Create Logical Drive** menu option must start with creating an array. When a new array is created as described above, you can define a new logical drive using the same procedure as described in Chapter 9, “Administration - Logical Tab” on page 219.

### Locate

This menu item is used to find where the selected disk is placed in your enclosure in the rack. You also have a choice to locate the enclosure where the disk is placed.

### Secure Erase and Import Security File Key

If the selected disk drive in Physical pane is a SED (Self Encrypted Disk) disk drive, you will have a choice to securely erase an unconfigured SED disk drive, or if you just inserted SED disks from a different DS3500 Storage Subsystem, you can import a security key to enable their use in your subsystem. For more information, refer to Chapter 15, “Disk Security with Full Disk Encryption drives” on page 439.

## 10.2.4 Controller menu

If you right-click one of the two controllers in the Physical pane, the menu (Figure 10-7 on page 262) for other controller configuration options appears.

|                                  |   |
|----------------------------------|---|
| View Associated Logical Elements |   |
| Locate Controller Enclosure      |   |
| Change                           | ▶ |
| Configure                        | ▶ |

Figure 10-7 Controller menu

## View Associated Logical Element

If you select this menu item, you can see a list of related logical objects, like array and logical drives, that are using the selected controller.

## Locate

This is used to find location of the selected controller in your enclosure in the rack.

## Change

Select this menu item and configuration options that are displayed include:

- ▶ Preferred Loop ID: This is discussed in detail in 10.5, “Set preferred loop ID” on page 268 (FC version only).
- ▶ Remote Login: This is discussed in detail in 10.6, “Set remote access” on page 270

## Configure

Select this menu item and configuration options that are displayed include:

- ▶ Ethernet Management Ports: This is discussed in detail in 10.7, “Set Ethernet management ports” on page 270
- ▶ iSCSI Ports: This is discussed in detail in Chapter 13, “Administration - iSCSI” on page 309, and this menu is available only if subsystem has iSCSI daughter card installed.

## 10.3 Set hot spare drive

A hot spare drive is like a replacement disk drive installed in advance. Hot spare disk drives provide additional protection that can be essential if there is a disk drive failure in a fault-tolerant RAID array.

As DS3500 Storage Subsystem configuration can only have SAS disk types, but with different capacities and speeds, you have to plan in advance how to provide the best hot spare protection for your storage subsystem.

Follow these guidelines to plan your hot spare coverage properly:

- ▶ Hot spare disk drives must be of the same media type and interface type as the disk drives that they are protecting.
- ▶ Hot spare disk drives must have capacities equal to or larger than the used capacity on the disk drives that they are protecting. The DS3500 storage subsystem can use a larger drive to recover a smaller failed drive to. It cannot use smaller drives to recover a larger failed drive. If a larger drive is used, the remaining excess capacity is blocked from use.
- ▶ FDE (SED) disk drives (encryptable) provide coverage for both security capable and non-security capable disk drives. Non-security capable disk drives can provide coverage only for other non-security capable disk drives.
  - For an array that has secured FDE drives (security key activated), the hot-spare drive should be a secured or unsecured FDE drive only, with the same or greater capacity.



- For an array that has FDE drives that are not secured (security key not activated yet), the hot-spare drive can be either a FDE drive or a non-FDE drive.

In a mixed disk environment that includes NearLine SAS drives, high speed SAS drives, and FDE (SED) SAS drives (either with security enabled or not enabled), use at least one disk drive of each type at the largest capacity within the array.

**Important:** When assigning disks as hot spares, make sure they have enough storage capacity. If the failed disk's used capacity is larger than the hot spare capacity, reconstruction is not possible. Ensure that you have at least one of each size or all larger drives configured as hot spares.

### ***Hot spares locations***

Distribute the hot spare drives evenly across the expansions of your storage subsystem, but avoid having multiple ones in a single enclosure. Because hot spare drives are in standby, without traffic or I/O until a drive fails, you want to maximize the overall performance of your system by evenly distributing your production drives across the different expansions. At the same time, this avoids the risk of a single disk drive channel, or expansion enclosure failure causing loss of access to all hot spare drives in the storage subsystem.

### ***Quantity and type of hot-spares drives***

There is no fixed rule about the quantity of disk drives to assign as hot spares, but as a general rule, define at least one of every 18 - 20 disk drives of a particular media type and interface type, or one for every two fully populated enclosures. Because of disk sizes, and especially in large configurations with arrays containing numerous drives, the reconstruction of a failed drive to a hot spare drive can take a long time proportional to the quantity of drives in the array and the size of the disks. If in addition to that time, you have to wait to have a new disk available on-site to replace the failed drive, then the probability of having another disk failure increases. Having multiple spare drives will not mitigate the reconstruction time, but at least an array will be prepared sooner for a second disk failure.

**Note:** There is no definitive rule about how many hot spares you should install, but use a ratio of one hot spare for about 18 - 20 identical disk drives in large installations. In small installations with just a couple of disks in DS3500 Controller enclosure, use one hot spare disk of largest capacity for high speed SAS and one for NearLine SAS, if both these types are used.

In the following section we cover the available methods of defining hot spare drives with Storage Manager.

To start the configuration of a new hot spares from the Physical pane, select an unconfigured disk drive that you plan to use as hot spare disk, right-click it, and select **Hot Spare Coverage**. The window shown in Figure 10-8 on page 264 opens.

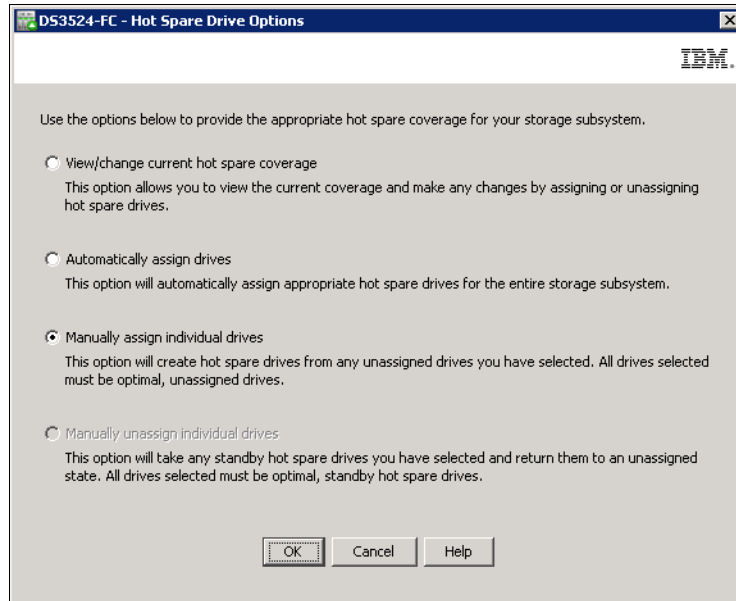


Figure 10-8 Hot Spare Drive options

You have four options to define hot spare drives:

- ▶ View/Change coverage
- ▶ Automatic assignment
- ▶ Manual assignment (if unassigned disk was selected in the Physical pane)
- ▶ Manual unassignment (if hot spare disk was selected in the Physical pane)

### **Automatic assignment**

For automatic hot spare assignment, follow these steps:

1. To automatically create the hot spare coverage using the drives that are available, select **Automatically assign drives**. The recommended quantity of spares drives needed for your configuration are created automatically.
2. Select the **Physical** tab to view the results of the automatic hot spare creation.

This automatic hot spare creation function on DS3500 creates one hot spare drive for every 10 unassigned disk drives. Results can be unexpected in environment with mixed drive types and if disk drives are already assigned to arrays. In such scenarios, we recommend using the manual assignment, instead of the automatic method.

### **Manual assignment and unassignment**

To perform manual hot spare assignment, follow these steps:

- ▶ Select an unconfigured disk drive that you plan to use as hotspare disk, right-click it, and select **Hot Spare Coverage**. The window shown in Figure 10-8 displays with the option **Manually assign individual drives** already preselected. Click **OK** and the unassigned drive is defined as a hot spare.
- ▶ Select a hot spare disk drive that you plan not to use as hotspare disk any more, right-click it, and select **Hot Spare Coverage**. The window shown in Figure 10-8 display with the option **Manually unassign individual drives** already preselected. Click **OK**, confirm by typing yes in next window, and the former hot spare drive is an unassigned disk drive again.

#### Tip for hot spare selection:

- ▶ Select drives of equal or greater size than the capacity of the largest disk in the storage subsystem.
- ▶ In large configurations with arrays containing many drives, it is necessary to define multiple hot spares as the reconstruction of a failed drive to a hot spare can take a long time.
- ▶ Consider having hot spares for each speed and type of disk drives, that is SAS, NearLine SAS, and SED SAS (encryptable).

### View/change hot spare coverage

In the Hot Spare Drive Option window (Figure 10-8 on page 264) select the **View/change hot spare coverage** option. The Hot Spare Coverage window similar to Figure 10-9 appears.

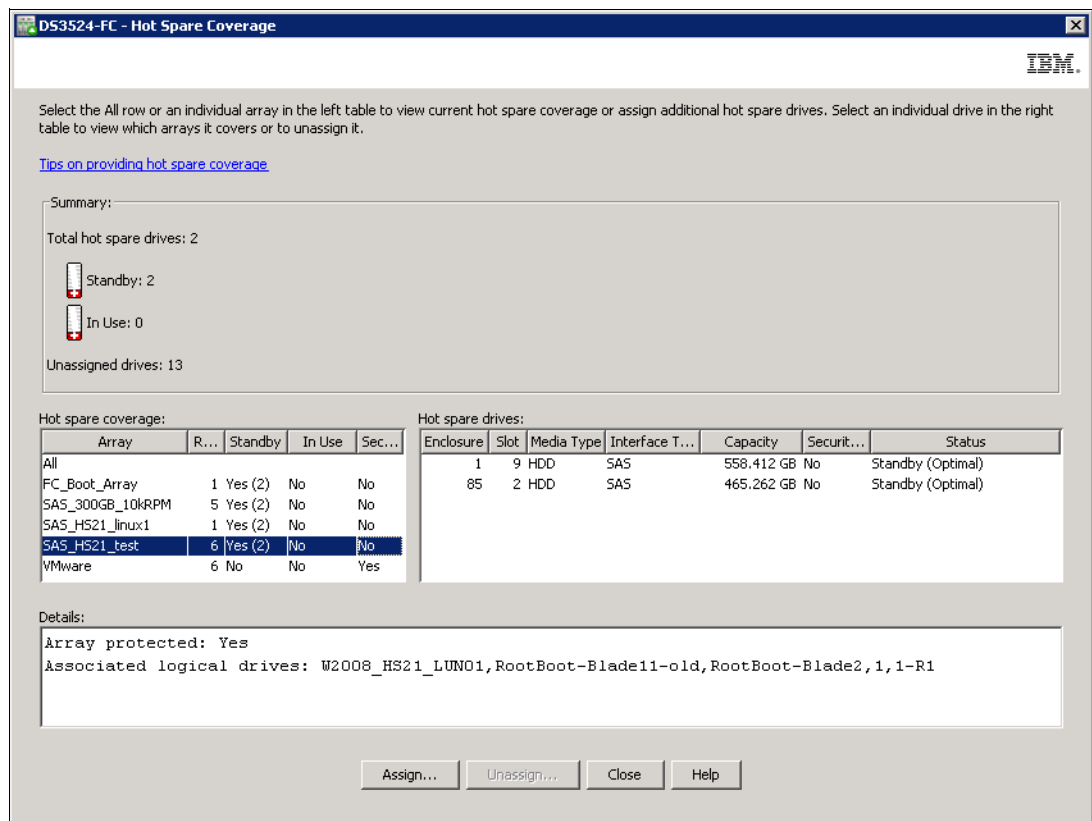


Figure 10-9 View/Change hot spare coverage

This window has four panes:

- ▶ The Summary pane, which displays:
  - Numbers of hot spares
  - Number of hot spares in use
  - Number of unassigned disks that can be used as hot spares
- ▶ The Hot spare coverage pane shows all arrays and whether they are protected by hot spares. If an array is selected in this left pane, in the right **Hot spare drives** pane you see the hot spare assigned. If you select an array that has no hot spare defined, you can click **Assign** in the bottom part of the window, and a new window like in Figure 10-10 on page 266 opens.

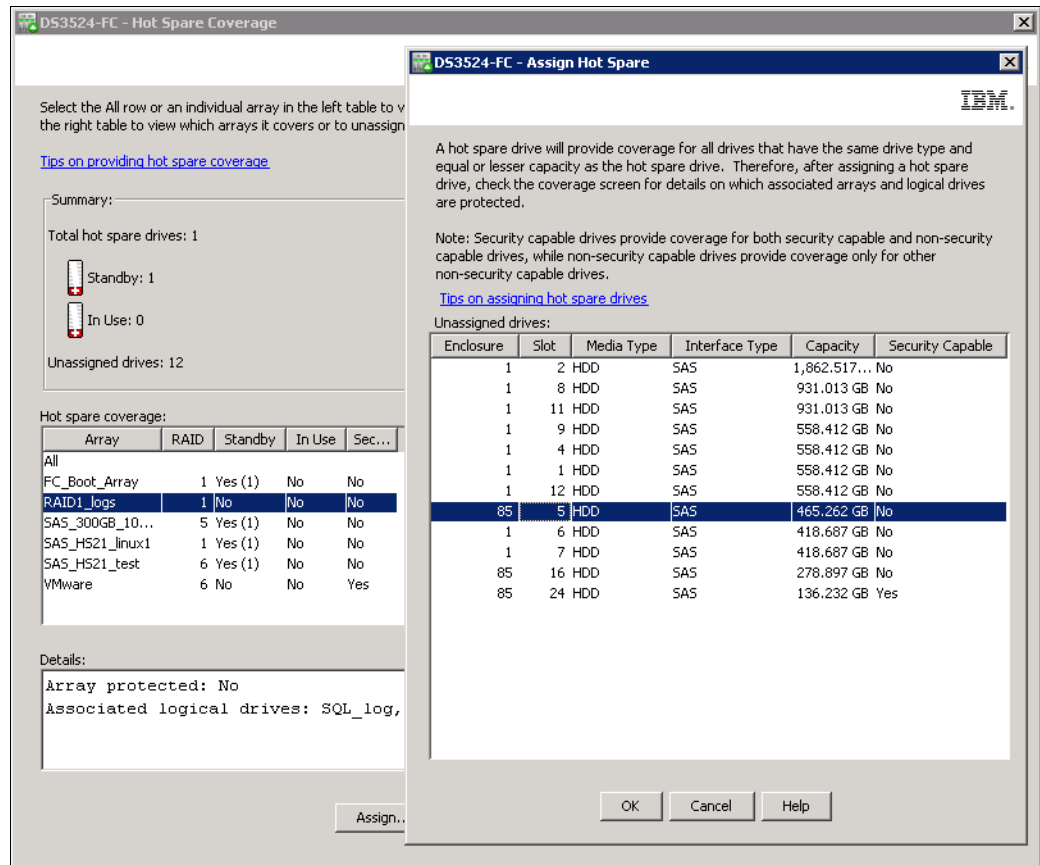


Figure 10-10 Assigning hot spare

Here you can select new hot spare for your unprotected array.

- ▶ The Hot spare drives pane shows hot spare drives used. If you have more hot spare drives then you need, you can select hot spare drive here, and click **Unassign** to unassign it. The **Unassign** button will be unavailable if you can lose protection by unassigning this hot spare disk. If you still need to unassign this hot spare drive, do it manually as described in “Manual assignment and unassignment” on page 264.
- ▶ The Details pane shows logical objects that belong to selected array or hot spare in panes above.

Make sure to use this window to check that all your arrays are protected. If they are not protected, implement complete hot spare protection for all arrays as soon as possible.

## 10.4 Failed disk drive replacement

You can replace a failed or missing disk drive that is used by an array. Logical drives are in “Degraded” status if parity for the failed disk is still calculated to the hot spare drive (see the example of the Recovery Guru window in Figure 10-11 on page 267).

You can assign a working unassigned disk drive that is already inserted in the DS3500 Storage Subsystem as a replacement drive. In addition, the hot spare in use can be used as replacement drive but only after parity recalculation is completed.

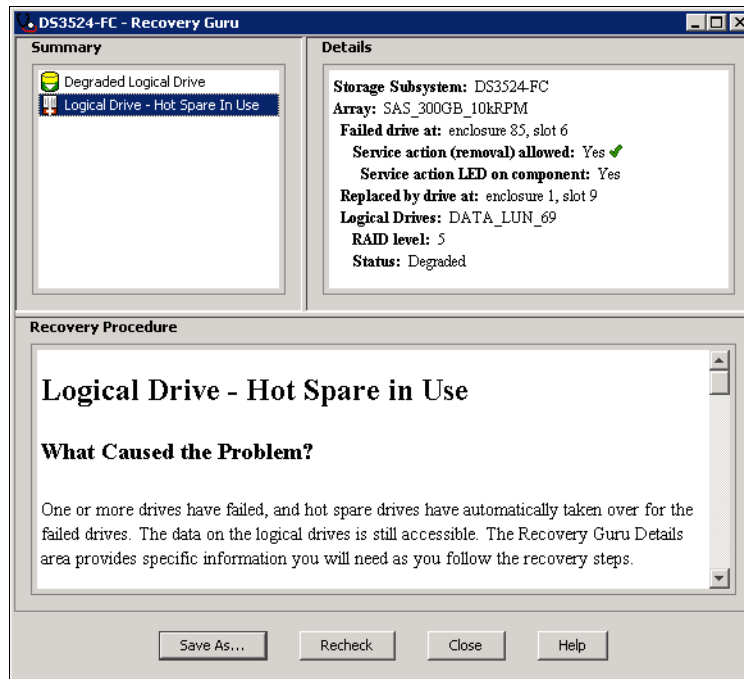


Figure 10-11 Degraded Logical Drive + Hot Spare In Use

To replace a failed disk drive, perform the following steps:

1. Select the failed disk drive or missing disk drive icon in the Physical pane, right-click it, and select **Replace Drives** in the menu that appears (Figure 10-12).

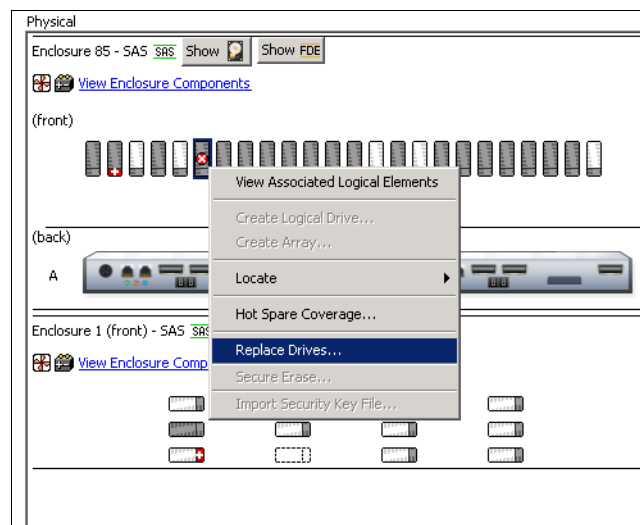


Figure 10-12 Replace Drive

2. The list of eligible replacement disk drives appears in the window (Figure 10-13 on page 268).

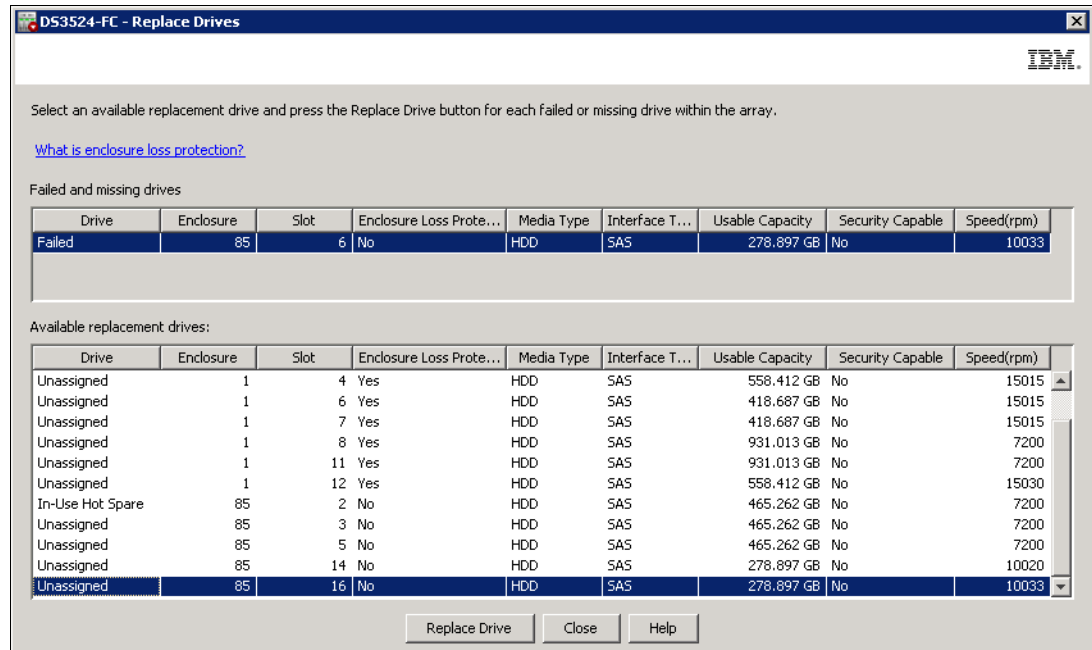


Figure 10-13 Replace Drives - Select disk

3. Select the missing or failed drive from the upper list (if more than one drive is missing or failed). Select the appropriate replacement drive from the lower list of available drives. The available replacement disk drives are equal to or greater in size, and have the disk drive type as the failed or missing disk drive. Select the drive that has the same or higher RPM speed, and that has a Yes entry in the column "Enclosure Loss Protection", if such a disk exists in the list.

You also have a choice to select the In-Use Hot Spare drive as a replacement drive (see Figure 10-13). If this hot spare disk drive is defined and is in use by the array with a missing or failed disk, you have to wait until the DS3500 controller reconstructs parity of all the logical drives in that array.

4. Click the **Replace Drive** button to start the operation. The confirmation window appears. Click **OK** and a new updated window like Figure 10-13 shows any remaining missed or failed drives to replace. Continue with the replacement or click **Close** to finish.

**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 might damage the array and the data on the drive, and the data might not be recovered.

If an In-Use Hot Spare disk drive is used as replacement drive, define another disk drive as a hot spare. Returning a missing disk or repaired disk does not set the Hot Spare setting automatically. A repaired or reinserted disk drive becomes a regular unassigned disk drive.

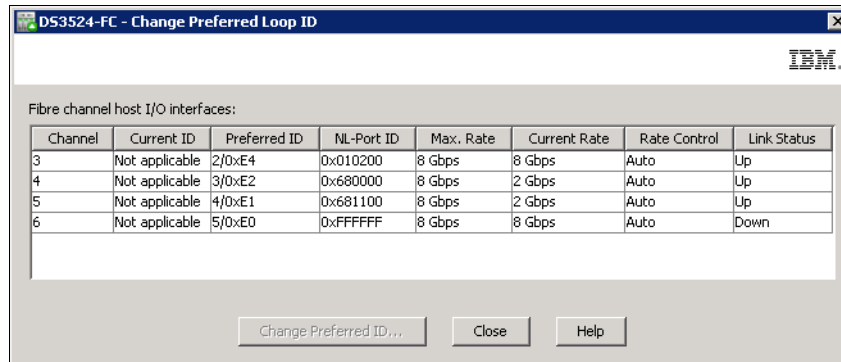
## 10.5 Set preferred loop ID

Loop ID is unique ID of a port in fibre-channel loop topology, sometimes is referred as a loop address. The **Preferred Loop ID** menu is not available for DS3500 without a FC daughter card installed.

**Attention:** If your Fibre Channel host adapter driver manages storage subsystems based on a loop ID (or arbitrated loop physical address - AL\_PA), do not change the value of the preferred loop ID unless you have a special reason to do so.

To change the preferred loop ID, follow these steps:

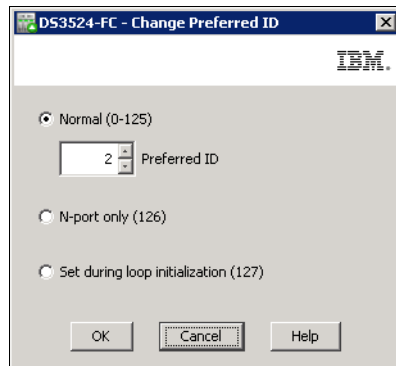
1. Select the controller in the Physical pane, right-click it, and select **Preferred Loop ID** in the menu. A list of FC ports displays in a new window (Figure 10-14).



| Channel | Current ID     | Preferred ID | NL-Port ID | Max. Rate | Current Rate | Rate Control | Link Status |
|---------|----------------|--------------|------------|-----------|--------------|--------------|-------------|
| 3       | Not applicable | 2/0xE4       | 0x010200   | 8 Gbps    | 8 Gbps       | Auto         | Up          |
| 4       | Not applicable | 3/0xE2       | 0x680000   | 8 Gbps    | 2 Gbps       | Auto         | Up          |
| 5       | Not applicable | 4/0xE1       | 0x681100   | 8 Gbps    | 2 Gbps       | Auto         | Up          |
| 6       | Not applicable | 5/0xE0       | 0xFFFFFFFF | 8 Gbps    | 8 Gbps       | Auto         | Down        |

Figure 10-14 FC ports - Loop ID

2. Select the line with port you want to change and click **Change Preferred ID** at the bottom of the window. The window shown in Figure 10-15 displays.



☒ Normal (0-125)  
 Preferred ID: 2  
☐ N-port only (126)  
☐ Set during loop initialization (127)

Figure 10-15 Loop ID setting

3. Select one of the following options:
  - **Normal (0-125):** This option assigns a normal preferred loop ID value to the controller port. Select a value from the spinner box. If the controller port currently has a preferred loop ID between 0 and 125, this radio button is selected by default. The current ID is shown in the spinner box.
  - **N-port only (126):** This option configures only the controller host interface as a point-to-point connection. This option forces the controller to operate as an N\_Port, which disables the loop mode. If the controller port currently has a preferred loop ID of 126, this radio button is selected by default.
  - **Set during loop initialization (127):** Use this option to assign the controller port to accept any available address during loop initialization. If the controller port currently has a preferred loop ID of 127, this radio button is selected by default.
4. Click **OK** to save the setting.

## 10.6 Set remote access

This option lets IBM engineers access the subsystem controller by TELNET. It can be done from outside of the server room where DS3500 Storage Subsystem is installed, allowing the IBM engineer to service the subsystem without having to use the special serial management port of DS3500, which is used for maintenance only.

**Note:** Leave the Remote Login option disabled if you do not have a specific need or directions from IBM support. This prevents unauthorized access to your DS3500. Normally, you only need to enable remote login to debug specific failures, and you will be directed by your IBM Support Representative in that case.

To set remote access, perform the following steps:

1. Select the controller in the Physical pane, right-click it, and select **Remote Login** in the menu. The window shown in Figure 10-16 appears.

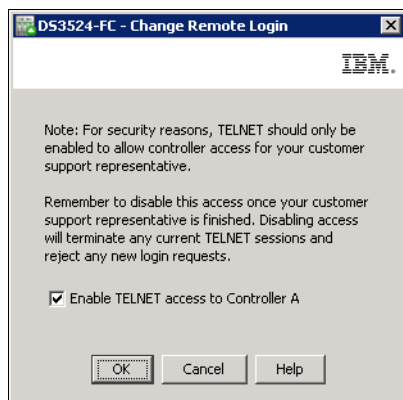


Figure 10-16 Remote login set

2. Read warning in the window, select **Enable TELNET** to set direct remote access to selected controller, and click **OK**. You will need to do the same procedure for the second controller.

## 10.7 Set Ethernet management ports

Setting the IP addresses of Ethernet management ports can be done in the configuration window. By default, the DS3500 Storage Subsystem will try the BOOTP/DHCP service to request an IP address. If no BOOTP/DHCP service is found in the network, the controllers revert to fixed default IP addresses.

The DS3500 storage subsystem has two management Ethernet ports per controller. Use one Ethernet port for daily management of your DS3500 storage subsystem. Reserve the other port for use by service personnel or for subsystem monitoring hardware that might be available in the future.

The default IP addresses of the first management Ethernet ports are:

- ▶ Controller A: 192.168.128.101
- ▶ Controller B: 192.168.128.102



The default IP addresses of the second management Ethernet ports are:

- ▶ Controller A: 192.168.129.101
- ▶ Controller B: 192.168.129.102

**Note:** Do not connect the second Ethernet ports to your LAN, and keep the default IP addresses on them so that service personnel can attach an isolated Ethernet switch if needed.

### 10.7.1 Initial settings for default IP addresses in DS Storage Manager

In order to change the default IP addresses, you need to connect to the DS3500 storage subsystem. Perform the following steps:

1. Connect the Ethernet ports to an isolated Ethernet switch.
2. Set the IP address in your management machine to match the default IP address subnet on the DS3500 storage subsystem (192.168.128.xxx).
3. Connect the management machine to the same Ethernet switch.
4. Use the Storage Manager client interface to change the IP addresses on both controllers:
  - a. Start the Storage Manager client interface. The DS Storage Manager - Enterprise Management window opens. For more information, see Chapter 6, "Administration - Enterprise Management" on page 157.

Select the autodiscovery icon or click **Tools** → **Automatic Discovery**. If this is the first time you are using the Storage Manager software after it is launched, the program presents a window to start discovery of the attached devices automatically.

- b. Select **OK** to initiate automatic discovery.

If the DS3500 storage subsystem does not appear after this action, check the network connectivity to both default IP addresses of the storage subsystem using the **ping** command. Each time a network cable is plugged to an Ethernet port of a controller, it detects the linkup and initiates a request for a dynamic address. If it does not find one, it assigns the default IP address. If you try to connect to the storage during that time, you might receive an error, so wait at least 5 minutes after plugging in the cable to attempt either the automatic discovery or manually adding the controller IP addresses.

**Tip:** Before starting the automatic discovery or adding the DS3500 storage subsystem manually, wait for the controller to finish its boot process and then another 5 minutes after connecting the network cable to allow for the DHCP process to complete. If this process is interrupted unexpectedly, the controller might assign IP addresses 0.0.0.0 by error, and you will have to reboot the storage subsystem to repeat the procedure again.

Alternatively, you can add your storage manually from the Enterprise Management selecting menu by clicking **Edit** → **Add Storage Subsystem**, and then completing the fields with the default IP addresses, as shown in Figure 10-17.

Figure 10-17 Manually adding a storage system

- c. A confirmation message appears and prompts you for additional storage subsystems. Click **No** to finish adding storage subsystems.
- d. After successfully adding a storage subsystem, the Devices tab of Enterprise Management shows the new system.  
Double-click the storage subsystem in Devices tab to open the Subsystem Management window. If not already set, you will be prompted to set the mandatory password for your storage subsystem.
- e. The storage subsystem might also ask you to synchronize the clock with the management station if the controllers are out of synchronization. Click **OK** to accept synchronization if prompted to do so.
- f. Now you are connected to the DS3500 Storage Subsystem, but with the default IP addresses. Change the address to your specific ones. Follow the standard procedure described in 10.7.2, “Configure Ethernet Management Ports” on page 272 for configuring the Ethernet ports on DS3500 controllers.

## 10.7.2 Configure Ethernet Management Ports

To configure the Ethernet management ports of the DS3500, perform the following steps:

1. Select one controller in Physical pane, right-click it, and select **Configure** → **Ethernet Management Ports**. A new window like in Figure 10-18 on page 273 appears. In this window you can set IP addresses, IP masks, default Gateway and the link speed of Ethernet port, for all four Ethernet ports on both controllers.

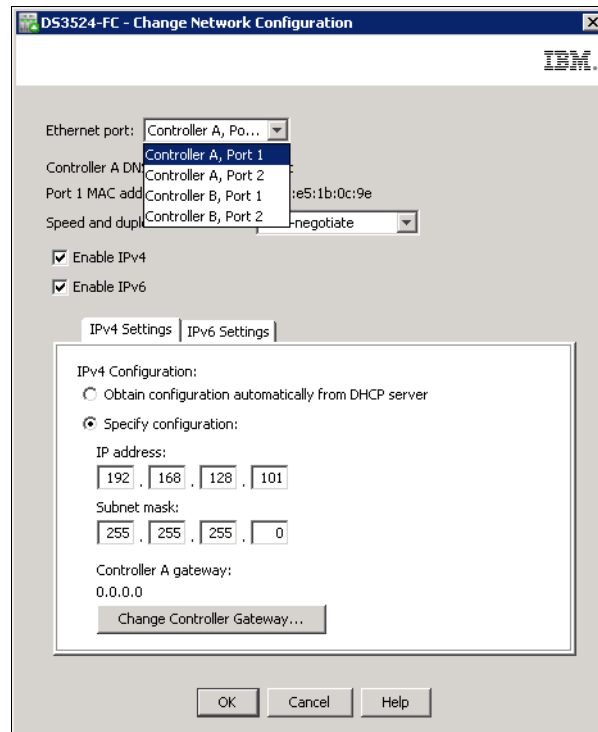


Figure 10-18 Change Ethernet settings

2. Input the required network information and click **OK**. You can configure both controllers from this window by selecting the appropriate Ethernet port from the drop down list. Make sure to click **OK** after configuring both of them. You must set the same default gateway on both controllers.

The DS3500 has two Ethernet ports per controller. Make sure to select the port of each controller you already cabled. If you are planning to have the system connected to only one network, select port 1, leaving port 2 for service.

**Note:** After changing the IP addresses, DS Storage Manager can freeze, because you have lost management connection to subsystem. Close the Storage Subsystem window, and remove the subsystem with changed IP addresses from the Enterprise Management window. Then add the subsystem by automatic discovery or manually by entering your newly assigned IP addresses.

**Tip:** To manage storage subsystems through a firewall using out-of-band management, configure the firewall to open port 2463 to TCP data.

## 10.8 Configure iSCSI Ports

If your DS3500 has an optional iSCSI daughter card installed, select one controller in Physical pane, right-click it, and select **Configure** → **iSCSI Ports**

It opens a window as in Figure 10-19 on page 274 used for setting parameters of iSCSI ports on daughter cards in DS3500. controllers. All iSCSI configuration options and iSCSI port settings are described in Chapter 13, “Administration - iSCSI” on page 309.





Figure 10-19 iSCSI port settings on DS3500



## Administration - Mappings Tab

In this chapter, we describe the *Mappings* tab of the IBM System Storage DS Storage Manager. This tab is primarily used for assigning logical drives to host servers by storage partitioning, which is also called “mapping”. This functionality is also known as “LUN masking”. If a logical drive is mapped to a host, it is seen there as a LUN (Logical Unit Number).

In 3.3.7, “Storage partitioning” on page 63 we explained the concept of storage partitioning. In this chapter we show practical examples of how to configure storage partitioning for SAS or FC connected hosts. If you need a specific procedure for iSCSI host attachment, see Chapter 13, “Administration - iSCSI” on page 309.

Because heterogeneous hosts can be attached to the DS3500 storage subsystem, you need to configure storage partitioning for two reasons:

- ▶ Each host operating system requires slightly different settings on the DS3500 storage subsystem. You need to tell the storage subsystem the host type to which it is attached.
- ▶ There is interference between the hosts if every host has access to every logical drive. By using storage partitioning (LUN masking), you ensure that each host or host group only has access to its assigned logical drives.

You can have up to 256 LUNs assigned to a single storage partition, and a maximum of 256 LUNs configured per DS3500 storage subsystem.

## 11.1 Mappings tab

The Mappings tab contains the Topology pane and the Defined Mappings pane. In the Topology pane, you can define the host groups, hosts, and host ports. The Topology pane shows a tree structured view of the logical nodes related to storage partitions in the storage subsystem. The Defined Mappings pane shows the mappings (logical drives) that are associated with a selected host group or host in the Topology pane. You can also change the mappings to grant logical drive access to host groups or hosts, and create storage partitions.

The overall process of defining the storage partitions is described in this chapter in the following topics:

- ▶ 11.2, “Defining Host” on page 277
- ▶ 11.3, “Defining Storage Partitioning” on page 286
- ▶ 11.4, “Defining Host Group” on page 289
- ▶ 11.5, “Manage Host Port Identifiers” on page 290
- ▶ 11.6, “Define Additional Mapping” on page 291
- ▶ 11.7, “View Unassociated Ports” on page 293
- ▶ 11.8, “Move, Remove and Rename Host” on page 293
- ▶ 11.9, “Change Host Operating System” on page 294
- ▶ 11.10, “Change and Remove Mapping” on page 295

To start with mapping, select the **Mappings** tab in DS Storage Manger subsystem window. If you have not defined any storage partitions yet, the Mapping Start-Up Help window (shown in Figure 11-1) opens. The information in the window advises you to only create host groups if your plan includes sharing logical volumes across hosts, normally a cluster.

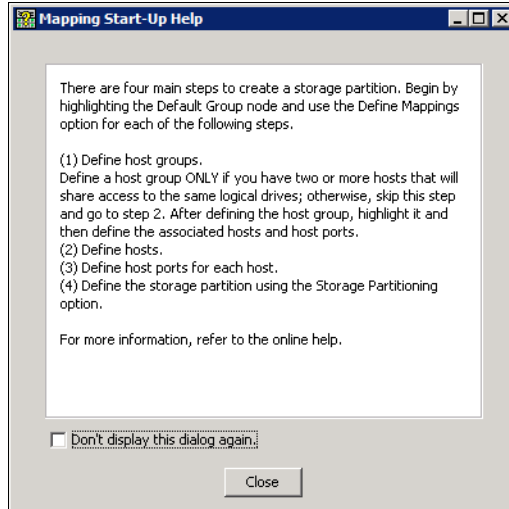


Figure 11-1 Mappings Start-Up Help

Figure 11-2 shows an example of the **Mappings** tab. The right side of the window lists all mappings that are owned by the object you selected on the left side. If you highlight the Storage Subsystem, you see a list of all defined mappings. If you highlight a specific host group or host, only the associated mappings are listed.

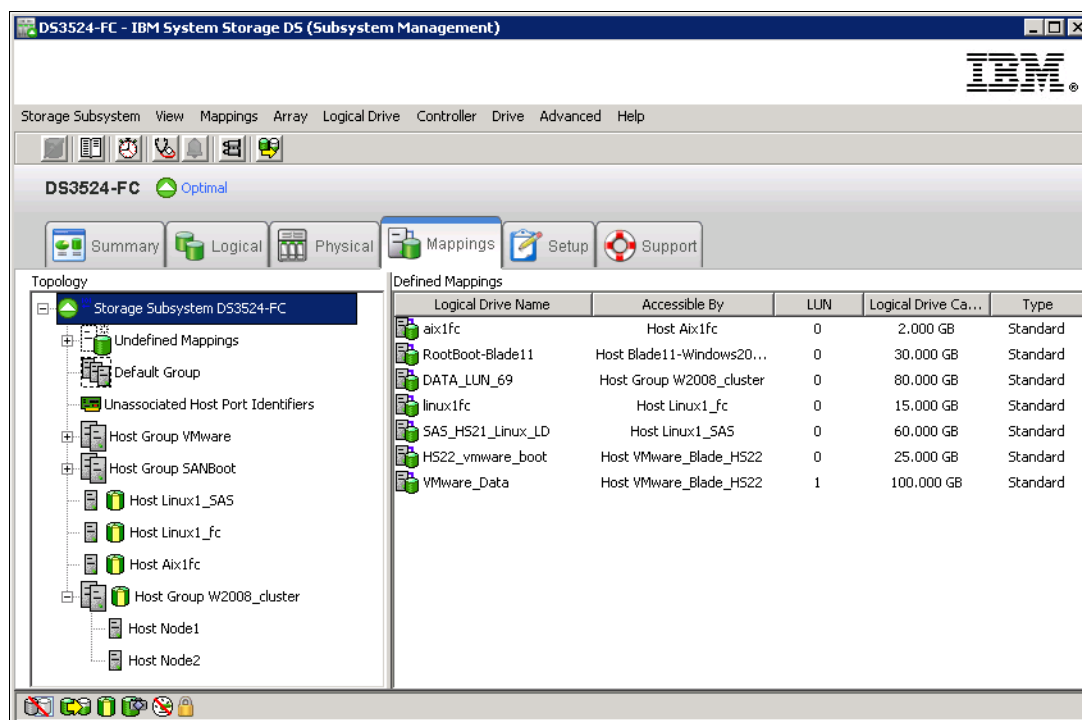


Figure 11-2 Mappings tab in DS Subsystem Management window

In the following sections, we show how to configure the storage partitioning.

**Note:**

- ▶ Before configuring the storage partitions, if you have your storage connected to a FC switch, make sure to first define the SAN zones appropriately.
- ▶ Perform a hardware rescan on each host that has to be mapped to the DS3500 storage subsystem to reconfigure the FC or SAS devices (HBA) and allow the WWPN to be presented in the Storage Manager.
- ▶ Generally, configure only one host HBA per SAN zone along with one DS3500 storage subsystem controller.

## 11.2 Defining Host

The Host can be defined as stand-alone server, or as a cluster member. Perform the following steps to define host:

1. Right-click your Storage Subsystem icon, right-click the Default Group icon, or right-click an existing Host Group. In the menu that appears, select **Define** → **Host** as shown in Figure 11-3 on page 278. Even with no host group created, it is easier to create the hosts first and then the host groups. In the Figure 11-3 on page 278 you can also see that all

general configuration options for the whole Storage Subsystem are accessible from the Mappings tab. These options are described in Chapter 8, “Administration - Subsystem Management” on page 187.

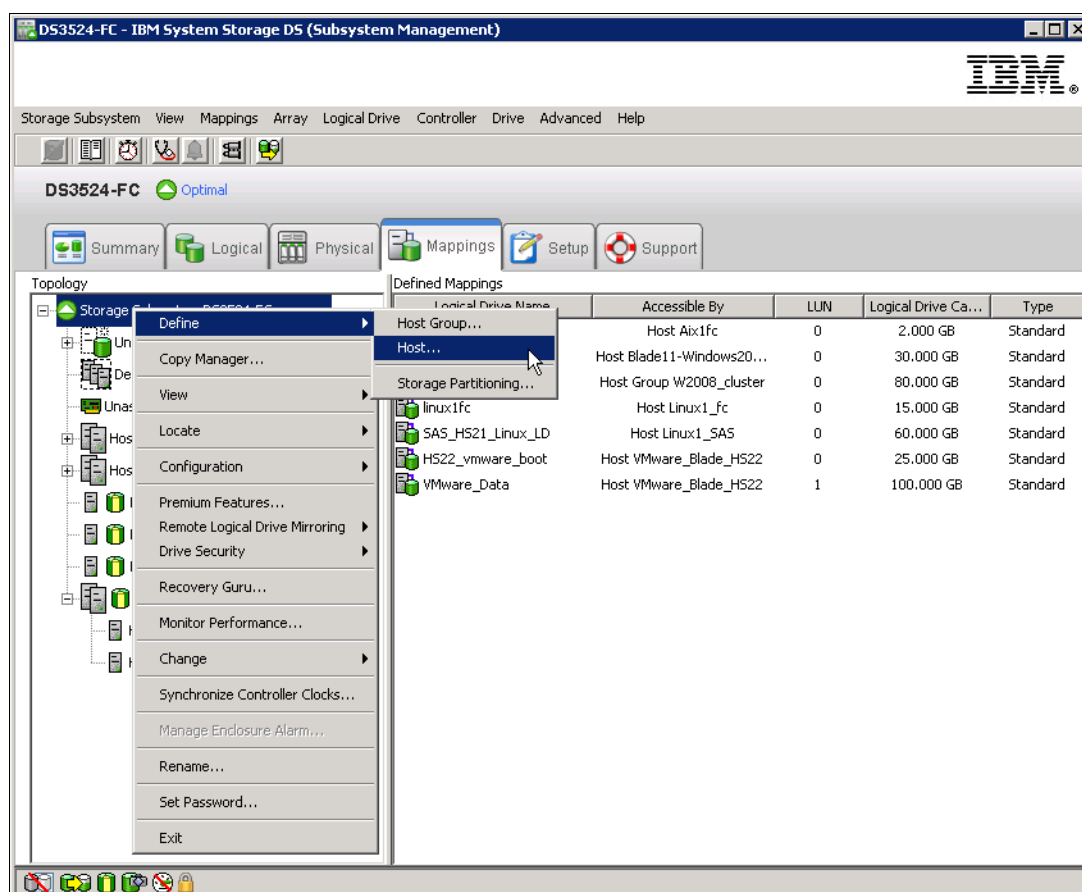


Figure 11-3 Define Host menu

When **Define** → **Host** option is selected, you will be presented with a wizard for defining a new host. You have to provide the following information during the process:

- Name of the new host
  - Protocol attachment method:
    - SAS
    - FC (if FC daughter card installed)
    - iSCSI (if iSCSI daughter card installed)
  - Host port identifiers with their Alias
  - Host type (the operating system that runs on the host)
  - Whether the host is going to participate in a cluster
  - Host Group name if the host is defined as clustered
2. The first window of the wizard is an introduction. In the next window (Figure 11-4 on page 279) you are asked to assign a **Host name** to the host being defined, and whether you plan to use storage partitioning.



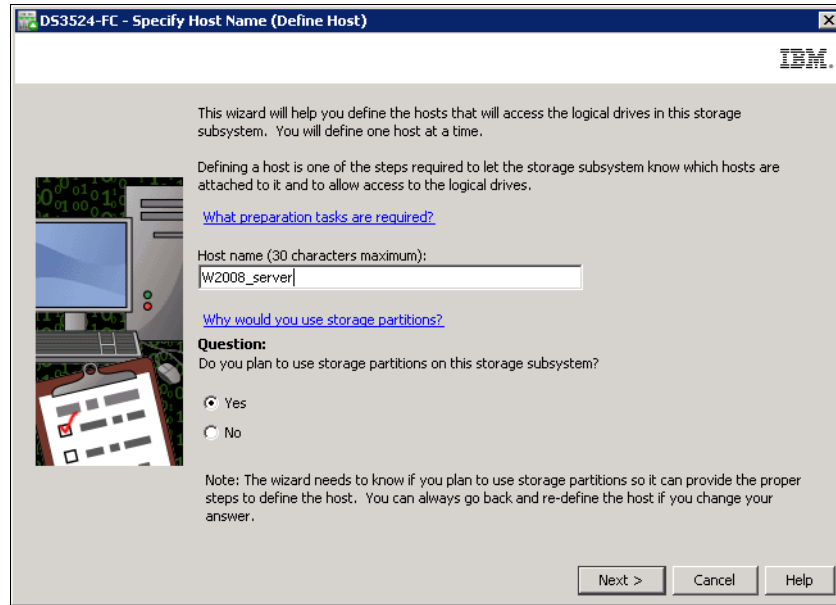


Figure 11-4 Define Host wizard: specify Host name

3. If you only plan to attach a single host server, then you do not need storage partitioning. But if you plan to group host environments together, or if clusters are sharing logical drives, click **Yes** to create different hosts groups at a later time. Select **Next** to continue.
4. In the next window (Figure 11-5 on page 280), you need the attachment protocol to specify the host interface type (SAS, FC, or iSCSI) and the HBA host port information (host port identifier and Alias). Remember that the HBA host port identifier is the world-wide port name of the particular HBA port. To make sure that you know what your server's HBA WWPN is, use the appropriate management tool for your HBA (for example SANsurfer management tool for Qlogic HBA), or use Fabric Name Server in FC fabric management tool, or you can note WWPN from BIOS during boot process of the server where HBA is placed.

DS3524-FC - Specify Host Port Identifiers (Define Host)

IBM.

The host communicates with the storage subsystem through its host bus adapters (HBAs) or its iSCSI initiators where each physical port has a unique host port identifier. In this step, select or create an identifier, give it an alias or user label, then add it to the list to be associated with host W2008\_server.

[How do I match a host port identifier to a host?](#)

Choose a host interface type:  
SAS

Choose a method for adding a host port identifier to a host:

☒ Add by selecting a known unassociated host port identifier

Known unassociated host port identifier:

- Select Identifier -

- Select Identifier -

50:00:62:b0:00:09:40:51

50:00:62:b0:00:09:40:50

Refresh

Alias (30 characters maximum):

Add Remove

Host port identifiers to be associated with the host:

| Host Port Identifier | Alias / User Label |
|----------------------|--------------------|
|----------------------|--------------------|

< Back Next > Cancel Help

Figure 11-5 Define Host: Interface type, known port identifier and Alias

- Select your interface type (SAS in our example)
- Select an active port already detected by the storage subsystem in the Known unassociated host port identifier menu as shown in Figure 11-5, or type the world-wide port name directly into the field, if the server is not connected or zoned yet, but you already know its WWPN (see Figure 11-6 on page 281).

The host communicates with the storage subsystem through its host bus adapters (HBAs) or its iSCSI initiators where each physical port has a unique host port identifier. In this step, select or create an identifier, give it an alias or user label, then add it to the list to be associated with host W2008\_server.

[How do I match a host port identifier to a host?](#)

Choose a host interface type:  
SAS

Choose a method for adding a host port identifier to a host:

☐ Add by selecting a known unassociated host port identifier

Known unassociated host port identifier:  
50:00:62:b0:00:09:40:51 Refresh

☒ Add by creating a new host port identifier

New host port identifier (16 characters required):  
500062b00094050

Alias (30 characters maximum):  
port\_2

Add Remove

Host port identifiers to be associated with the host:

| Host Port Identifier    | Alias / User Label |
|-------------------------|--------------------|
| 50:00:62:b0:00:09:40:51 | port_1             |

< Back Next > Cancel Help

Figure 11-6 Define Host: Interface type, new port identifier and Alias

- Type an Alias for the specified host port identifier and click **Add**.
  - Repeat the same process until you define all the HBAs. Remember that if you define only one HBA, the host can lose access if there is a problem with that data path. Select **Next** after defining all the HBAs for given host.
5. The next window requires you to specify the host type (Figure 11-7 on page 282). This is the operating system running on the host. It is vital that you select the appropriate host type because the RDAC and ADT/AVT settings rely on this setting. In addition, this is the part of the configuration where you configure the heterogeneous host support. Each operating system expects slightly different settings and handles SCSI commands differently. If you make a wrong choice, your operating system might not recognize LUNs anymore or path failover might not work. Select a host type version that supports clustering if you plan to configure the host as a cluster node.

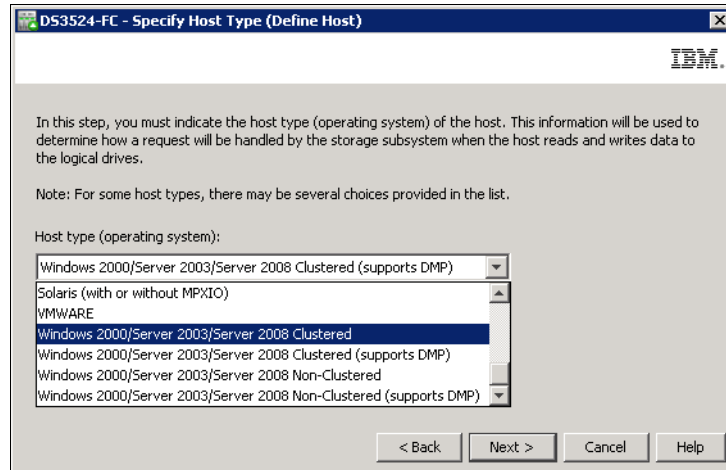


Figure 11-7 Define Host: Specify Host type

6. In the next step, you are asked whether the host is a part of a cluster (Figure 11-8). Choose **Yes** if it will be clustered, or **No**, if it will be a stand-alone server. Click **Next** to continue.

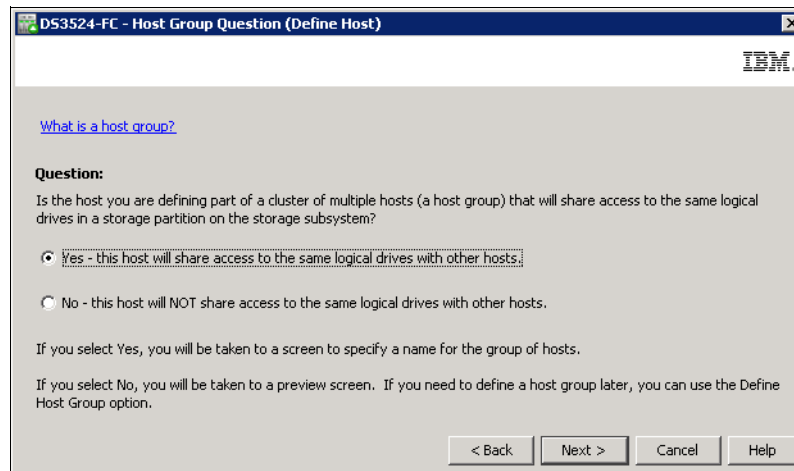


Figure 11-8 Define Host: Specify sharing (clustering)

7. (Optional) If the answer is **Yes**, then you need to specify a host group. The host group can be either a new or an existing one. If you select an existing host group, hosts that are already associated with that group are displayed in the right pane of this window. If you enter a new name, as shown in Figure 11-9 on page 283, the new host group will be created with this new host as a member. Select **Next** to continue.

**DS3524-FC - Specify Host Group (Define Host)**

IBM.

[What is a host group?](#)

Because you specified on the previous screen that the host you are defining will share access to logical drives with one or more other hosts, you must indicate the name of the host group that this host will be associated with.

You can either (1) manually enter a new host group name or (2) select an existing host group. If you select an existing one, you will be shown the hosts currently associated with it.

☒ Enter name (30 characters maximum)

W2k\_Cluster\_02

☐ Select existing host group

-Select from list-

Associated hosts in host group:

| Name | Host Type |
|------|-----------|
|------|-----------|

< Back   Next >   Cancel   Help

Figure 11-9 Define Host: Specify new Host Group

8. Finally, you have a chance to preview the new host definition (Figure 11-10 on page 284). If all the selections are correct, click **Finish** to define the new host. Note that you can save the host definition as a script file for SMcli by clicking the **Save As Script** button in the upper right corner of this window. You can use this script command to easily change the host definition for mapping copies of logical drives in high level scripts for disaster recovery.

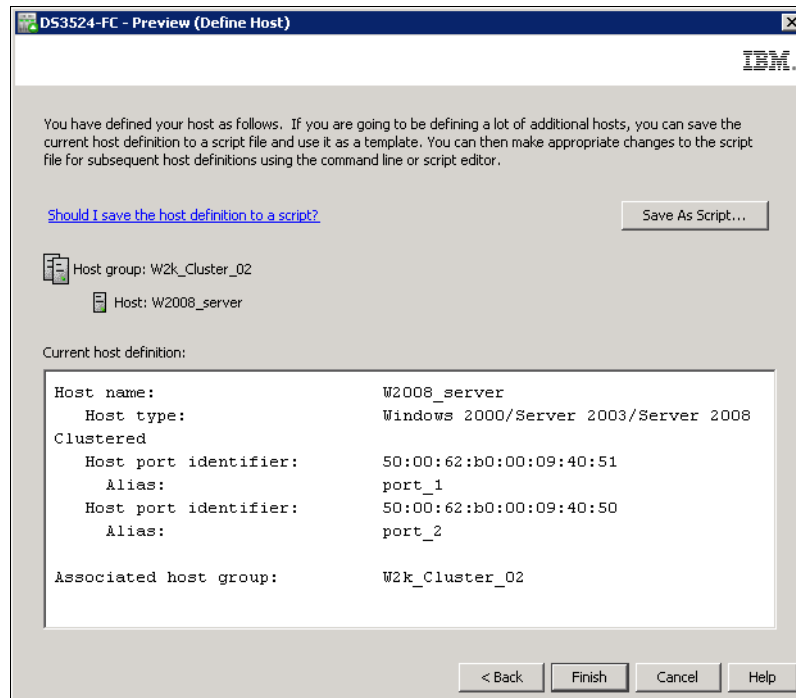


Figure 11-10 Define Host: new Host definition Preview

- After the process is completed, the new host (and the host group, if it was also defined) is placed in the Default Group in the Topology tree (Figure 11-11). It will stay there until you actually create a storage partition by assigning the logical drives to that host (or host group).




Figure 11-11 Host defined

- Repeat these steps to create more hosts and host groups.

## 11.2.1 Adding a new Host to existing Host Group

You can decide to incorporate a new host to the host group, and make it work together with the other host or hosts as a cluster. Remember that when you put the new host into the host group, the new host can access all logical drives mapped to this host group immediately.

Perform the following steps to add Host to Host Group:

1. Define the new host by performing the steps shown in 11.2, “Defining Host” on page 277.
2. Specify the host type to be the same as other hosts in the target host group, and select **Yes - this host will share access to** in the window shown in Figure 11-8 on page 282,
3. Incorporate the host into the existing host group where your other host or hosts are already defined. Select **Select existing host group** as shown in Figure 11-12 and click **Next** to finish.

DS3524-FC - Specify Host Group (Define Host)

[What is a host group?](#)

Because you specified on the previous screen that the host you are defining will share access to logical drives with one or more other hosts, you must indicate the name of the host group that this host will be associated with.

You can either (1) manually enter a new host group name or (2) select an existing host group. If you select an existing one, you will be shown the hosts currently associated with it.

☐ Enter name (30 characters maximum)

☒ Select existing host group

W2k\_Cluster\_02

Associated hosts in host group:

| Name         | Host Type           |
|--------------|---------------------|
| W2008_server | Windows Server 2008 |

< Back   Next >   Cancel   Help

Figure 11-12 Add Host to existing Host Group

After performing these steps, the Mappings tab appears as shown in Figure 11-13 on page 286, showing the correct configuration for a cluster environment.

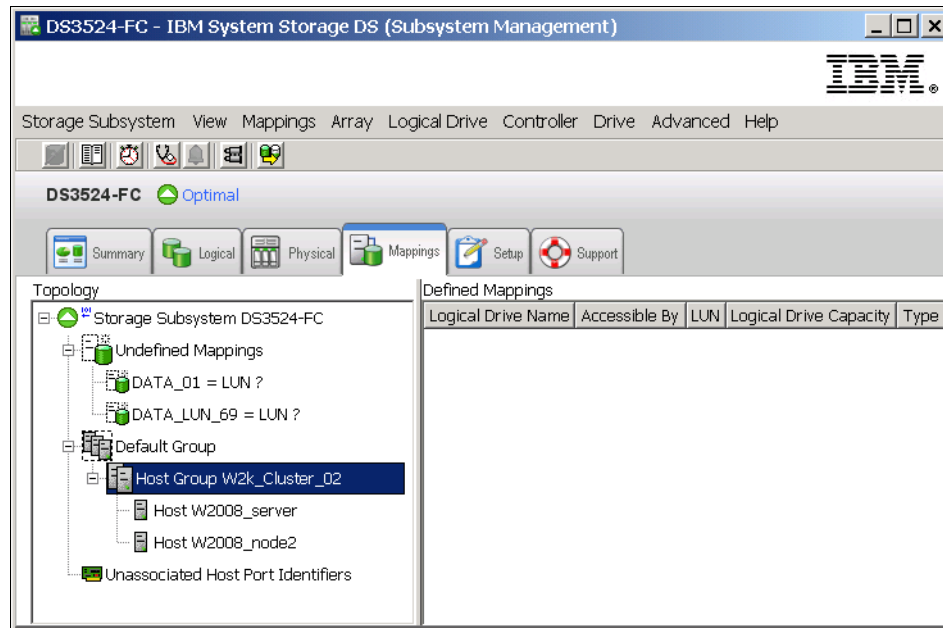


Figure 11-13 Second Host Group added for clustering

## 11.3 Defining Storage Partitioning

Next we describe how to set a new storage partition definition for Host or Host Group that has no logical drives mapped yet. You need to use the **Additional Mapping** menu to map a new logical drive to a host or host group where logical drives are already mapped.

Creating a storage partition by assigning the logical drives to new hosts or host groups using the following steps:

1. Right-click Default Group or any unmapped host or host group icon in the Topology tree below the Default Group (Figure 11-14 on page 287) and select **Define** → **Storage Partitioning**. The Storage Partitioning wizard will lead you through the process. Note that an unmapped logical drive must exist to start this wizard.



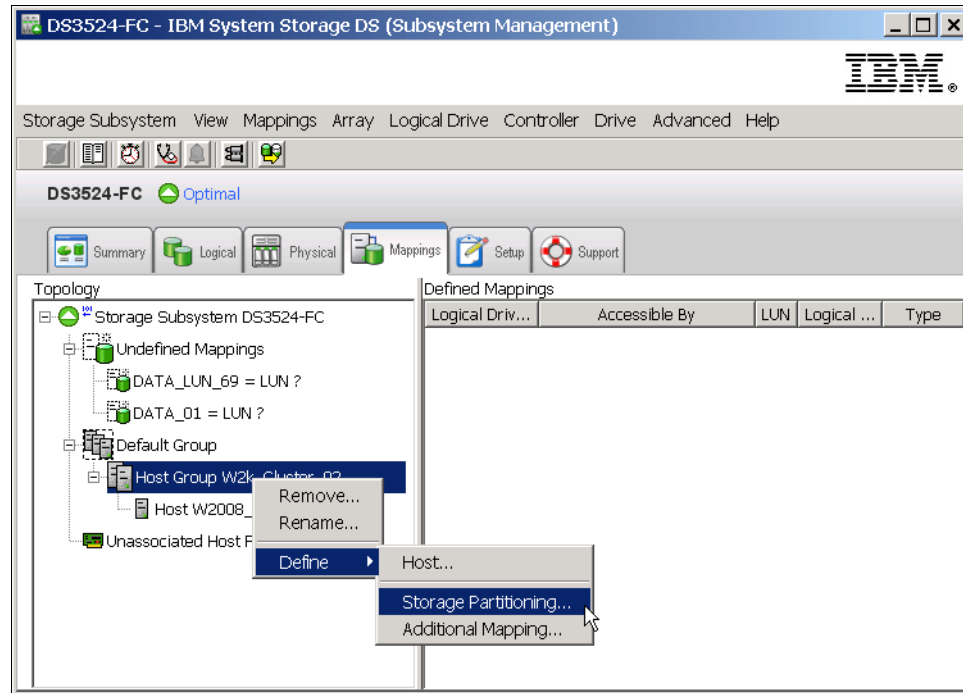


Figure 11-14 Storage Partitioning menu

- After the introductory window, the wizard asks you to select either a host or a host group (Figure 11-15). If you are creating a storage partition for clustered host servers, you need to specify the appropriate host group. Otherwise you can select an individual host. In this window, you can only see hosts and host groups with no logical drives mapped. If you selected a host or host group before you started Storage Partitioning wizard, it is already selected in this window.




Figure 11-15 Storage Partitioning - select Host or Host Group

- The next window (Figure 11-16 on page 288) allows you to select the logical drives that are going to be mapped to the host or the host group. Specify a LUN number that will be assigned to selected logical drive and click **Add**. This LUN number is visible as identification in operating system of mapped host. The Access logical drive as shown in

this window is used for In-Band management (by Fibre Channel) only. See “In-band management” on page 23 for more information. Do not map the access LUN if you do not plan on using In-Band management.

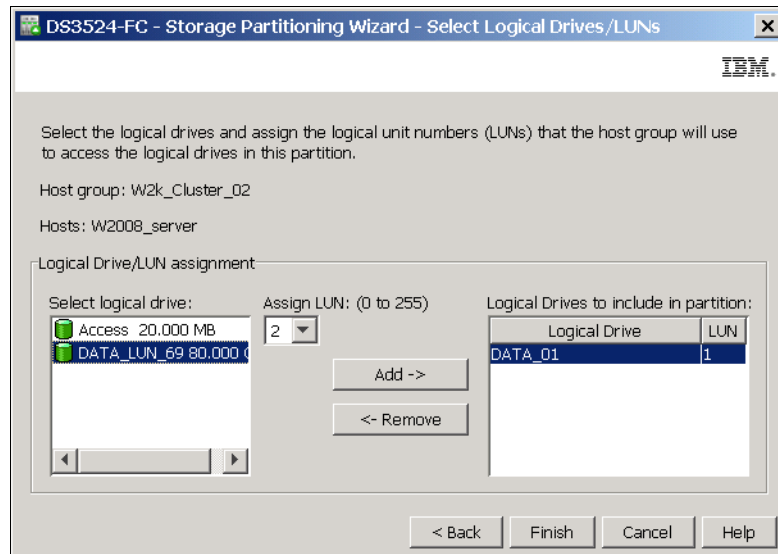


Figure 11-16 Storage Partitioning - select logical drives/LUNs

4. Click **Finish** when you are done selecting the logical drives and assigning the LUNs.
5. Your host group, host, and your new mapping are visible immediately in the Topology tree of the Mappings tab, and the host group and the host are now removed from the Default Group. The Storage Partitioning icon shows mapping is there as seen in Figure 11-17. In the Defined Mappings pane you can see the logical drive with assigned LUN number.

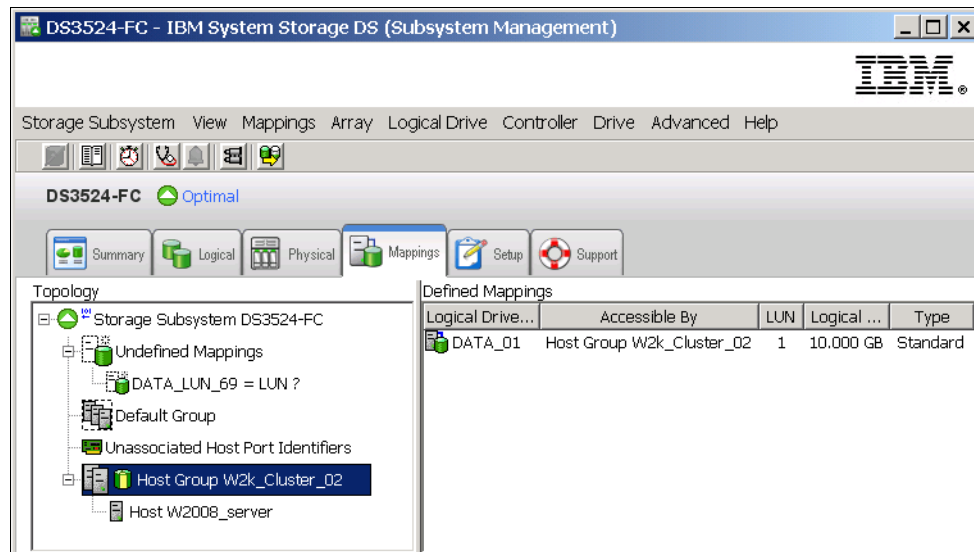


Figure 11-17 Displaying created mappings

Notice that the host **W2008\_server** is defined as a member of the host group **W2k\_Cluster\_02** in our example, and the logical drive is mapped to the host group. As mentioned before, the logical drive mapped to the host group is automatically mapped to all hosts included in that host group. You can check it by selecting the host icon in the tree.

You can continue creating other hosts or host groups and mapping other logical drives to them. Although you can create multiple partitions, you can only map logical drives to partitions up to the maximum number allowed by your specific storage subsystem and the premium feature configuration. However, even if a logical drive is mapped to host group with several hosts, only one storage partition is needed for the host group.

**Note:** You need to use only one storage partition per host group even if you have more hosts in this host group. However, if some of the hosts exclusively assigned one particular logical drive could be not shared by other hosts in this host group (for example SAN boot drive), you need to use another storage partition for this logical drive and host.

## 11.4 Defining Host Group

Because the host groups and hosts can be created together with the Define Host wizard, there might not be a need to define the host groups from outside the wizard. But you still have an option to do it and it is possible to start by creating a empty host group definition, then continue with the host definition inside the host group. The opposite sequence is also possible: you can first define stand-alone hosts and later define a host group and add the defined hosts into it during host definition.

It is also possible to move already defined hosts into the already defined host groups, and then move mapping from host to the host group. We describe moving hosts and mapping change on the end of this chapter (11.8, “Move, Remove and Rename Host” on page 293 and 11.10, “Change and Remove Mapping” on page 295).

Define a host group by performing the following steps:

1. Right-click Default Group and select **Define** → **Host Group**, as shown in Figure 11-18.

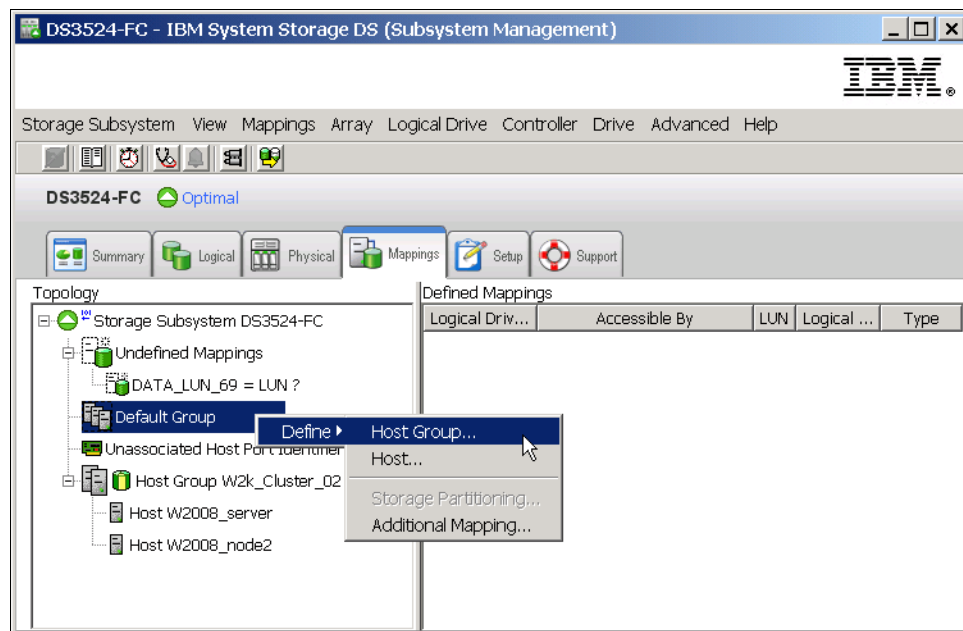


Figure 11-18 Define Host Group menu

- The Define Host Group window (shown in Figure 11-19) opens. Enter the name of the host group you want to define. Select every host you want to add to the group and click **Add**. If you select no hosts to add, the host group will stay empty and you can move existing hosts into it later, or you can add a host to this empty host group during host definition. When you are done, exit the window by clicking **OK**.

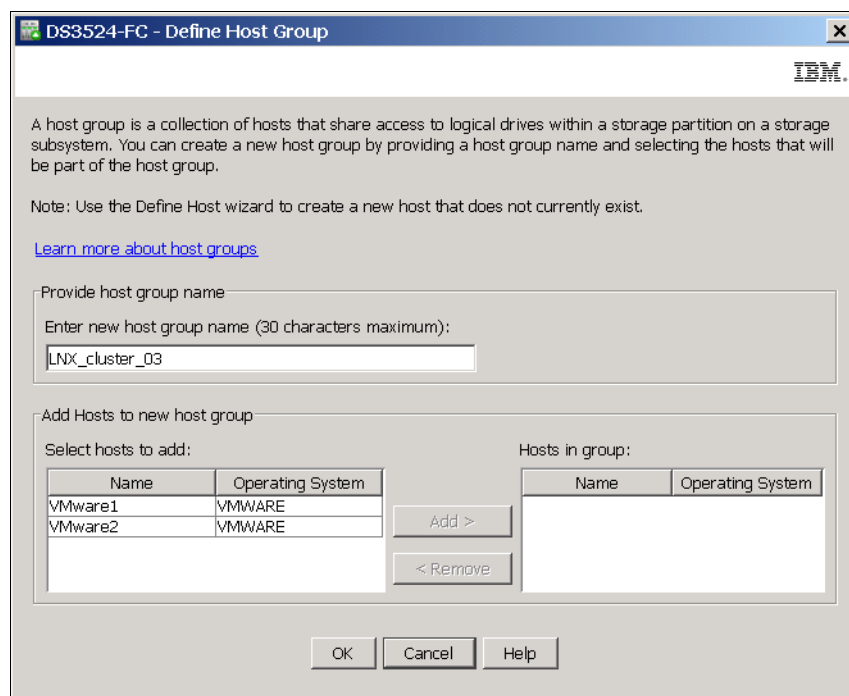


Figure 11-19 Define Host Group name

Because storage partitioning of the DS3500 storage subsystem is based on the world wide port names of the host ports, the definitions for the host groups and the hosts only represent a view of the physical and logical setup of your fabric. When this structure is available, it is much easier to identify which host ports are allowed to see the same logical drives and which are in separate storage partitions.

## 11.5 Manage Host Port Identifiers

For environments with multiple hosts and attachment types, you can use this option to have a single source of information about the available host ports on your servers.

Click **Subsystem Management** → **Mappings** → **Manage Host Port Identifiers**. Another option is to select any host in the Topology tree, right-click it, and select **Manage Host Port Identifiers**. The Manage Host Port Identifiers window as shown in Figure 11-20 on page 291 displays.

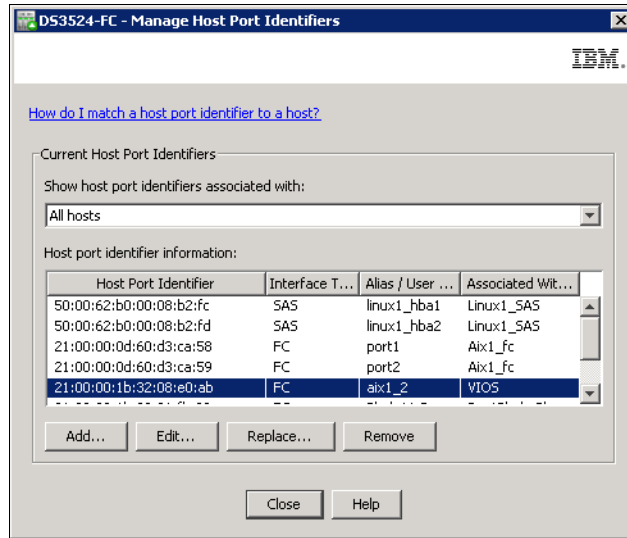


Figure 11-20 Manage Host Port Identifiers

Select this option if you need to review your port configuration assignment; add, remove, and change port settings; and whenever you need to replace a host interface card after a hardware replacement to continue presenting the logical volumes to the new host port identifier (WWPN, or iSCSI Initiator).

## 11.6 Define Additional Mapping

A particular host or a host group can be part of a certain storage partition. This means the logical drives are already mapped to this host or host group. If you need to map additional logical drives to this host or host group, use the Define Additional Mapping option by performing the following steps:

1. Right-click the host or host group to which you want to map a new logical drive and select **Define** → **Additional Mapping**, as shown in Figure 11-21 on page 292. You can also select a Logical drive in Undefined mappings section of the Topology tree, right-click it, and select **Define Additional Mapping** to continue.

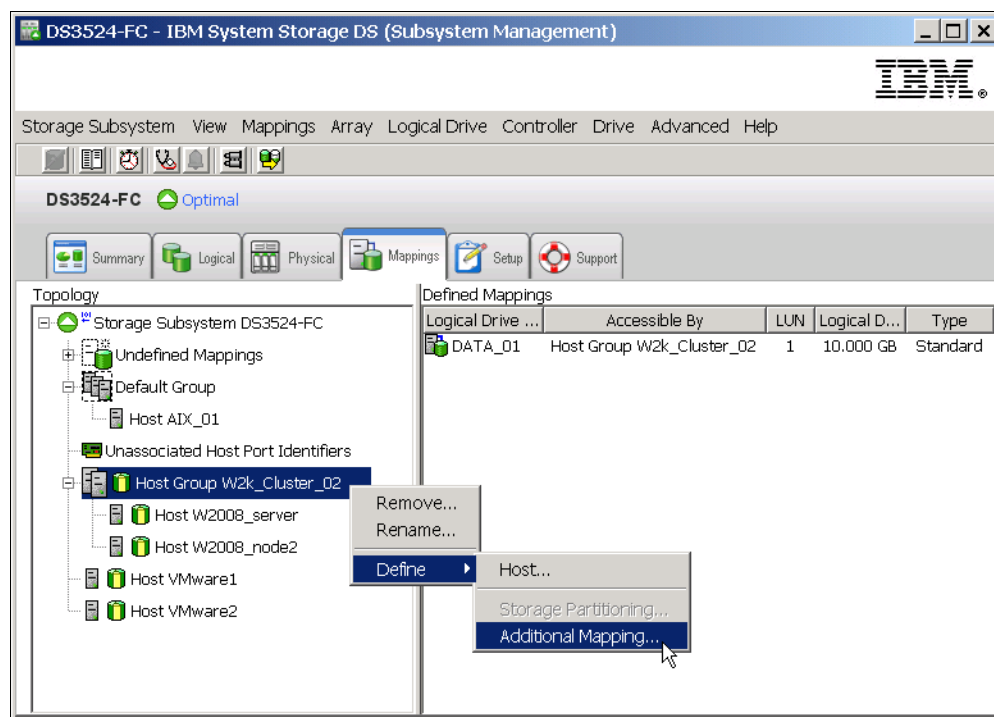


Figure 11-21 Define Additional Mapping

2. In the Define Additional Mapping window (Figure 11-22), select the logical drive you want to map to this host group or host, and assign the correct LUN number. Your host group or host is already preselected. In case you selected a logical drive in Undefined Mapping section for additional mapping, you need to select the host or the host group where it has to be mapped, and the correct LUN number. In our example the logical drive is preselected.

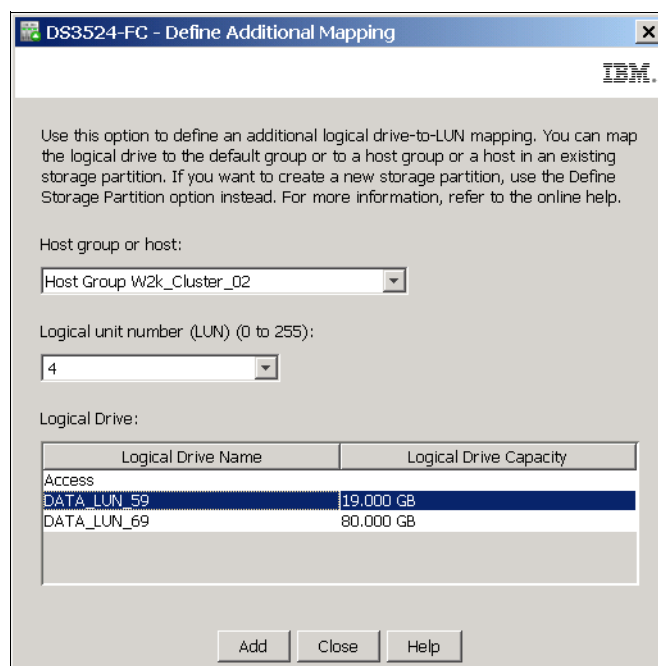


Figure 11-22 Additional Mapping - select logical drive

3. Click **Add**. The selected mapping is active immediately. You can repeat this selection and add processes until no unmapped logical drive are left. All hosts and all host groups are capable for additional mapping, and only unmapped logical drives can be selected in the lower part for mapping.
4. To make this new added logical drive (or drives, if you select more) available to the host systems without rebooting, you need to ensure that all host bus adapters are rescanned for new devices, and the devices are assigned within the operating system. This can be done using tools based on operating system. This procedure is documented in the sample configuration chapters later in this book.

## 11.7 View Unassociated Ports

In previous sections we described how to assign host ports (HBA in servers) to the Host. To check what other host ports are already recognized by your DS3500 Storage Subsystem but are not assigned to a Host yet, click **Mappings** → **View Unassociated Host Port Identifiers**, or select the **Unassociated Host Port Identifiers** icon in Topology tree, right-click it, and select **View**. If there are any, they are shown in a window as shown in Figure 11-23.

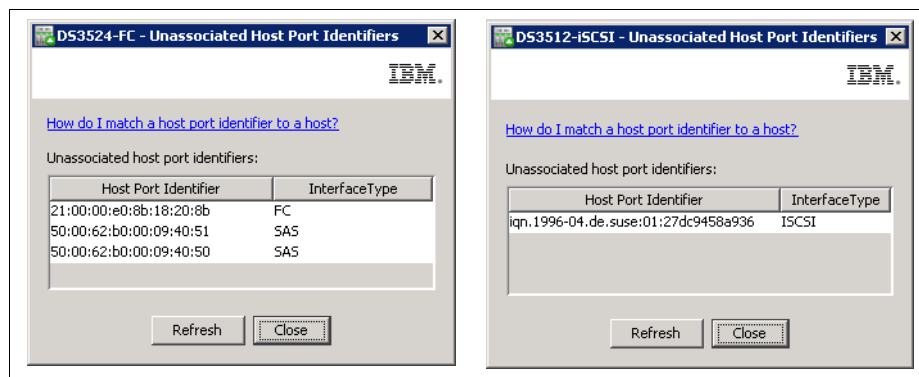


Figure 11-23 Unassociated port identifiers (FC, SAS and iSCSI Initiators)

The only difference for the iSCSI host attachment in the Mapping tab of DS Storage Manager, is that the hosts are defined by its iSCSI Qualified Name (IQN) instead of the WWPN as in Fibre Channel or SAS. Detailed information about how to configure iSCSI hosts is available in Chapter 13, “Administration - iSCSI” on page 309.

## 11.8 Move, Remove and Rename Host

If you right-click any **Host** in the Topology tree, it opens the menu, where you can see several options to configure:

### **Move**

Using this option, you can move the selected server from the current host group to another host group, or outside of the current group (Figure 11-24 on page 294).




Figure 11-24 Move host from host group or to another group

If you have a stand-alone host, you can move it to any appropriate host group for clustering.

**Note:** Be careful if the stand-alone host already has logical drives mapped. If the host group and the host use the same LUN number for different logical drives, you cannot add the host into the host group. You need to change LUN number for the host and rediscover logical drive in this host first.

### **Remove**

Use this option to delete host or host group definitions. Host port identifiers assigned to deleted host are moved and they are still visible in Unassociated host port identifiers window.

### **Rename**

This option is used to rename a host or host group.

**Note:** Be careful if SMcli scripts are applied on your DS3500. If you rename your host and its name is used in SMcli script, the script will no longer work. Changes of the host mapping are used almost in every DR process.

## **11.9 Change Host Operating System**

It is not difficult to reinstall the operating system on server hardware. If OS is changed, you do not have to define this server as a new host. It is possible to just change the host type (operating system), leaving the other definitions like host ports and mapped LUNs to stay unchanged. To change the host operating system, perform the following steps:

1. Right-click the host on which operating system has to be changed in the Topology tree and select **Change Operating System**. The window shown in Figure 11-25 on page 295 opens.



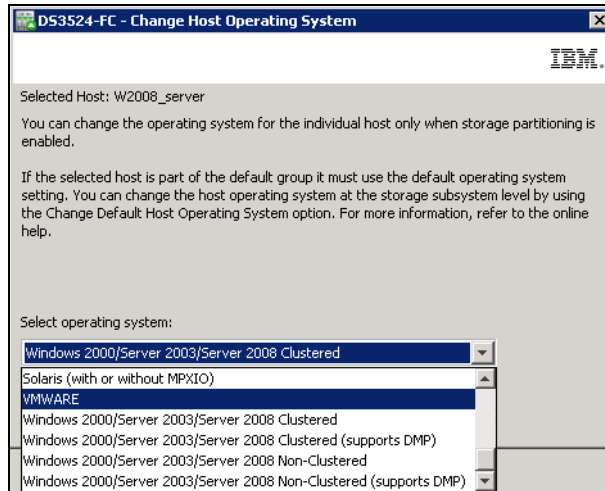


Figure 11-25 Change host type

2. Select the required operating system and click **OK**. The change is accepted immediately.

## 11.10 Change and Remove Mapping

Until now, all described mapping configuration changes were provided in the **Topology** pane. In the right **Defined Mapping** pane, you can see mapped logical drives. The logical drives displayed there are mapped to the host that is selected in Topology pane. Select a logical drive, right-click it, and two options appear:

- **Change Mapping**
- **Remove Mapping**

### **Change Mapping**

Select this option to change host or host group, where the logical drive will be newly mapped, or you can change LUN number. the change mapping window is shown in Figure 11-26.




Figure 11-26 Change Mapping - Host and LUN number

Click **OK** and a warning window opens (Figure 11-27 on page 296).




Figure 11-27 Change Mapping - warning

**Note:** Even if only LUN number is changed, your host operating system can lose access to already saved data on the mapped logical drive. Be careful if you change mapping on running hosts.

### ***Remove Mapping***

If you remove mapping, the selected logical drive is no longer accessible by host server. However, the data is still on the logical drive and it can be used later for other purposes.



## Administration - Setup tab

In this chapter, we describe the Setup tab of the DS Storage Manager. This tab can be used for initial configuration of a new DS3500 Storage Subsystem. The Setup tab contains links to tasks that you can perform when setting up a storage subsystem such as setting the storage subsystem name and password, configuring the Ethernet management ports, enabling the premium features, configuring the iSCSI settings, configuring capacities on the storage subsystem, and others. You can click a link to open the corresponding dialog or wizard.

We cover the following tasks:

- ▶ 12.2, “Locate Storage Subsystem” on page 298
- ▶ 12.3, “Rename Storage Subsystem” on page 299
- ▶ 12.4, “Set a Storage Subsystem Password” on page 299
- ▶ 12.5, “Configure iSCSI Host Ports” on page 300
- ▶ 12.6, “Configure Storage Subsystem” on page 301
- ▶ 12.7, “Map Logical Drives” on page 306
- ▶ 12.8, “Save Configuration” on page 306
- ▶ 12.9, “Manually Define Hosts” on page 307
- ▶ 12.10, “Configure Ethernet Management Ports” on page 308
- ▶ 12.11, “View/Enable Premium Features” on page 308
- ▶ 12.12, “Manage iSCSI Settings” on page 308

If you successfully provide all the base configurations that are referenced in this chapter, you can be almost certain that your subsystem is properly configured.

We will describe all these tasks and provide you with appropriate links for more information about a particular task as necessary.

## 12.1 Setup tab

Figure 12-1 shows the Setup tab in IBM System Storage DS Storage Manager.

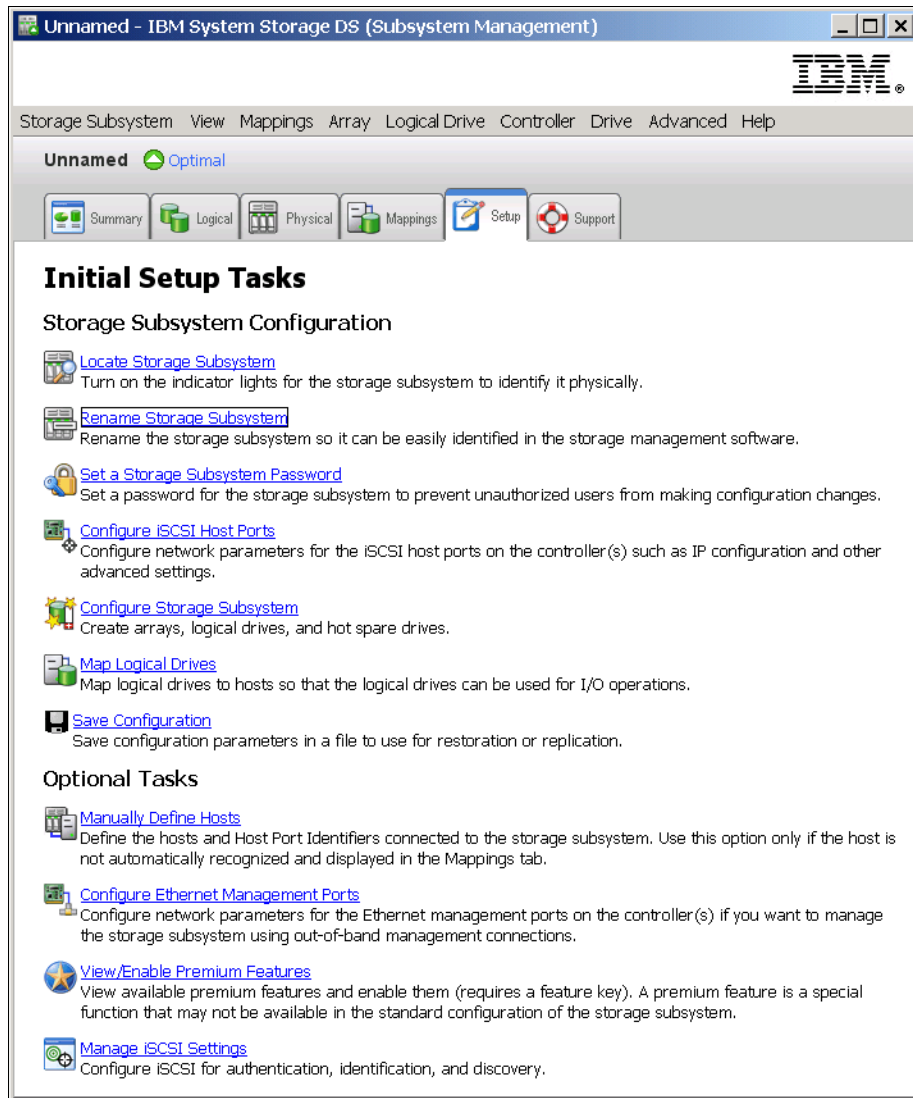


Figure 12-1 Setup tab of DS Storage Manager

When the iSCSI daughter card is not installed in DS3500 Storage Subsystem, the two iSCSI task links will not be displayed in the Setup tab.

We describe the links from top to bottom in the following sections.

## 12.2 Locate Storage Subsystem

The Locate Storage Subsystem link causes the blue LEDs of your DS3500 Storage Subsystem and all its EXP3500 expansion enclosures to flash so you can find where your subsystem is placed in the rack. Select **OK** as shown in Figure 12-2 on page 299 to stop flashing.

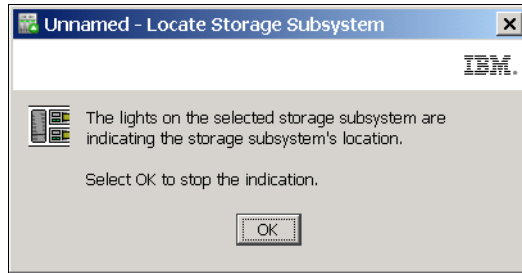


Figure 12-2 Locate Storage Subsystem

## 12.3 Rename Storage Subsystem

The Rename Storage Subsystem link opens a window where you can change the name of DS3500 Storage Subsystem. When you add a new storage subsystem to your Storage Manager, it is identified as Unnamed. If you have several storage subsystems installed in your environment, give each subsystem a meaningful name as shown in Figure 12-3.




Figure 12-3 Rename Storage Subsystem

The subsystem name is immediately updated in the Enterprise Manger window.

**Note:** It is important to know that renaming the storage subsystem can influence certain services running outside the Storage Manger. For example, if you run customized SMcli scripts where the subsystem is referenced by its name, the script will no longer work properly. Also, certain types of multipath failover drivers on servers use the name of attached subsystem. Rename a new subsystem before it is in production.

## 12.4 Set a Storage Subsystem Password

This link opens a window (Figure 12-4 on page 300) where you can change password. For DS3500, strong password rules are valid, so the password must be changed the first time you open DS Subsystem Manager for a new DS3500. You can find more details about DS3500 password settings in “Set Password” on page 204.




Figure 12-4 Password change

If you are setting the password for the first time on a new DS3500, the **Current password** field remains empty.

## 12.5 Configure iSCSI Host Ports

This link is shown in the Setup tab only if optional iSCSI daughter card is installed. It opens a window shown in Figure 12-5 on page 301 that is used for setting parameters of iSCSI ports on daughter card in DS3500. All iSCSI configuration options and iSCSI port settings are described in Chapter 13, “Administration - iSCSI” on page 309.

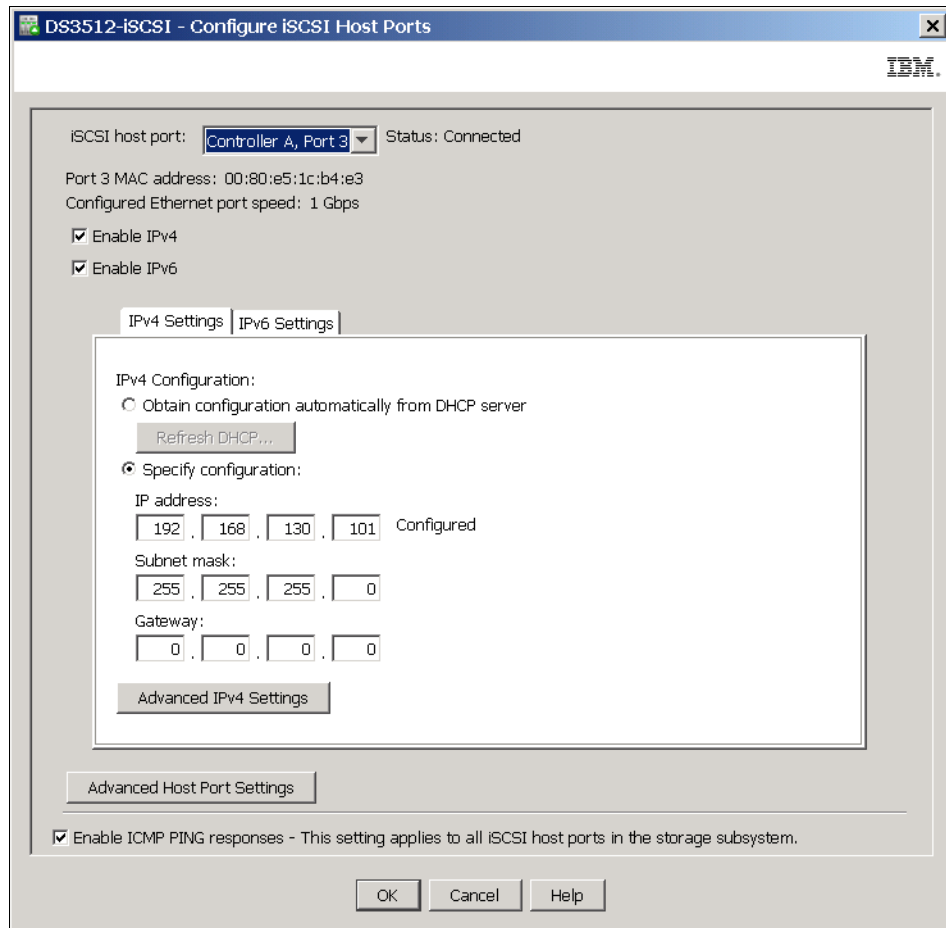


Figure 12-5 iSCSI port settings on DS3500

## 12.6 Configure Storage Subsystem

If you select this link, the wizard opens as shown in Figure 12-6 on page 302. This can be used to configure logical drives and arrays on installed disk drives in the storage subsystem.

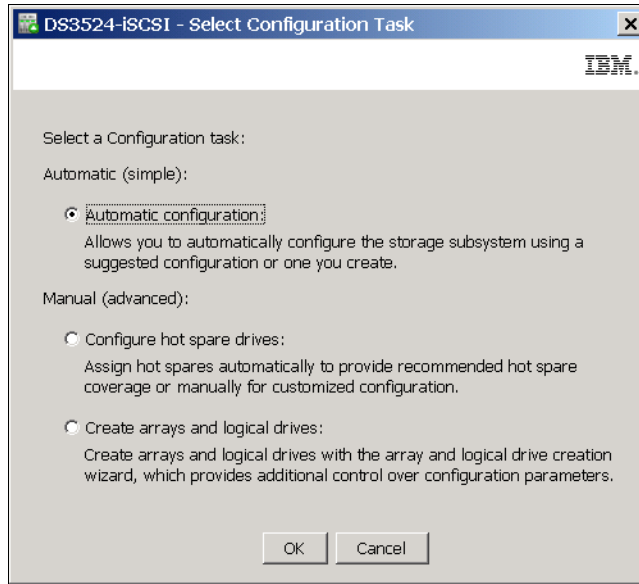


Figure 12-6 Configuration task

This wizard has three options:

- ▶ 12.6.1, “Automatic configuration” on page 302
- ▶ 12.6.2, “Configure hot spare drives” on page 305
- ▶ 12.6.3, “Create arrays and logical drives” on page 305

Select the required option and click **OK**.

## 12.6.1 Automatic configuration

If the Automatic configuration option is selected, a new window with an introduction to automatic configuration opens that uses of dedicated options to create well tuned logical drives. Click **Next** and a window as shown in Figure 12-7 on page 303 opens.

**Note:** If your storage subsystem contains drives with different media types and different interface types, or has no unassigned drives available for configuration, the Automatic Configuration window does not appear.



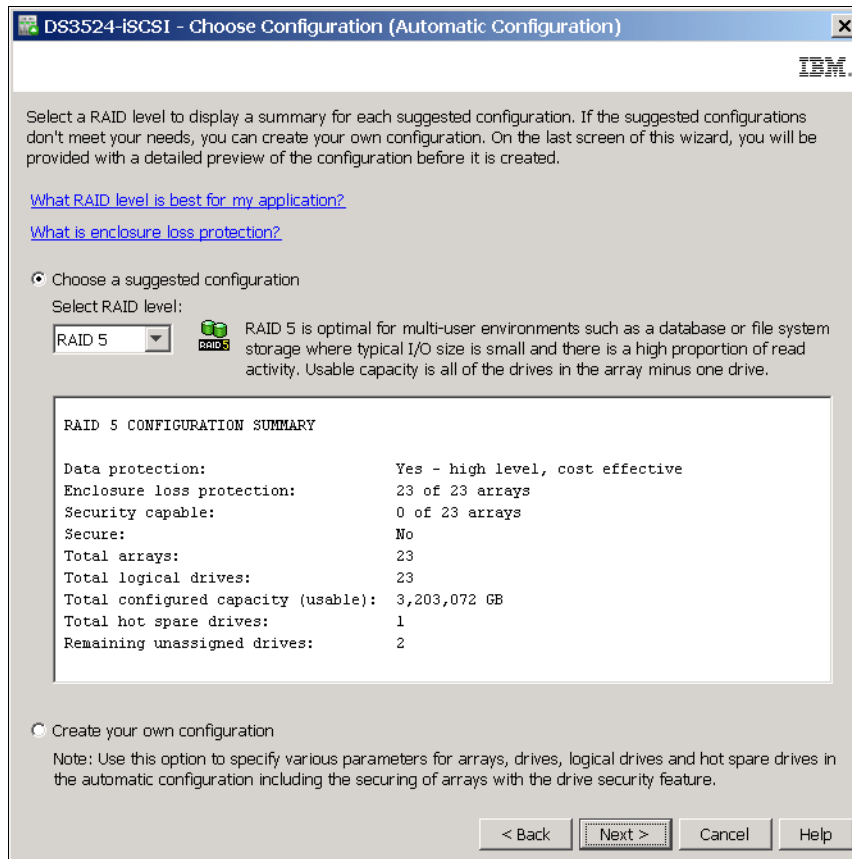


Figure 12-7 Fully automatic configuration of arrays

Do not use the Choose a suggested configuration option because it creates an impractical set of many arrays with a minimal number of disks in each array. Because you are only setting the RAID level, you have no control on the number and size of logical disks and number of arrays created using this option.

Instead, select the option to create your own configuration and click **Next**. A new wizard window opens. This wizard offers more parameters to modify for automatic configuration.

You can preset:

- ▶ RAID level
- ▶ Number of disk drives per array
- ▶ Number of Hot spare drives
- ▶ Number of logical drives per array
- ▶ If you select Change I/O Type you can preset:
  - Dynamic cache pre-fetch
  - Array segment size

Unfortunately you still have no influence on the size of logical drives, which is set automatically. You also cannot select which disk drives are used for arrays and hot spares, and selection of disk drives from separate enclosures is impossible.

Set your required values in the wizard window as shown in Figure 12-8 on page 304, and click **Next**.

DS3524-iSCSI - Create Your Own (Automatic Configuration)

IBM.

Use the following worksheet to define the appropriate number of arrays, drives, and logical drives you want to create.

[Tips and examples on allocating capacity.](#)

[What RAID level is best for my application?](#)

[What is the I/O characteristic type?](#)

Array parameters

Drives available for configuration: 72 (5 of 72 are security capable)

Select RAID level:

RAID 1 RAID 1 or "disk mirroring" offers high performance and the best data availability. Select four or more drives to achieve mirroring and striping (RAID 10 or RAID 1+0). Usable capacity is half of the drives in the array.

☐ Create a secure array (available only when a security key is set up for the storage subsystem).

Caution: Once an array is secured the only way to remove security is to delete the array and all of the data contained within the logical drives on it.

|                                      |                   |                         |
|--------------------------------------|-------------------|-------------------------|
| Number of arrays:                    | Drives per array: | <b>Allocated Drives</b> |
| 4                                    | 16                | 64                      |
|                                      | Hot spare drives: |                         |
|                                      | 4                 | 4                       |
| Total allocated drives (72 maximum): |                   | 68                      |
| Remaining unassigned drives:         |                   | 4                       |

Number of logical drives

Number of equal-sized capacity logical drives (per array):

8

I/O characteristic type: File system (typical)

Dynamic cache read prefetch: Enabled

Segment size: 128 KB

Change I/O Type...

Show Summary

< Back Next > Cancel Help

Figure 12-8 Automatic configuration with more parameters

A Summary window and a window with additional information about the automatic configuration process (Figure 12-9) open.

DS3524-iSCSI - Automatic Configuration

IBM.

The automatic configuration request has been successfully received by the storage subsystem. Please wait until this configuration has completed before submitting another automatic configuration request.

The management window will be updated with the new configuration. The time it takes to complete the request depends on the number of logical drives being created. A large configuration will take a long time to complete. Check the Event Log to determine the success or failure of the automatic creation.

OK

Figure 12-9 Automatic Configuration - final information

Because automatic configuration can create unexpected results, use manual configuration tasks to configure logical drives and arrays as described in 9.5, "Working with logical drives" on page 237 and 9.2.2, "Create array" on page 223.

## 12.6.2 Configure hot spare drives

If you select the Configure hot spare drives option in the Select Configuration Task window (Figure 12-6 on page 302), a new window opens for hot spare configuration. In this window (Figure 12-10) you have two choices: Automatic or manual configuration of hot spares.

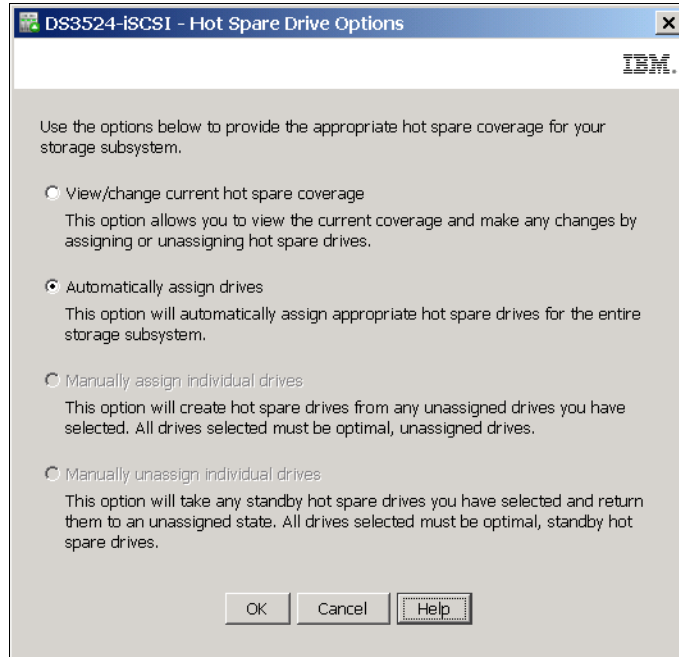


Figure 12-10 Hot spare drive configuration option

The concept of a hot spare drive is explained in 10.3, “Set hot spare drive” on page 262. There you can also find detailed description of the same “Hot Spare Drive Option” window, and how to set hot spare drives.

## 12.6.3 Create arrays and logical drives

The third option in Figure 12-6 on page 302 is **Create Arrays and Logical Drives**. It is for manual definition of arrays and logical drives.

When you select this option, the window as shown in Figure 12-11 on page 306 opens. This wizard will guide you through the definition process. If no array has been created already, you will only have last option (Create a new array using unconfigured capacity) available. After you create an array, you can continue with a logical drive creation on this array immediately.

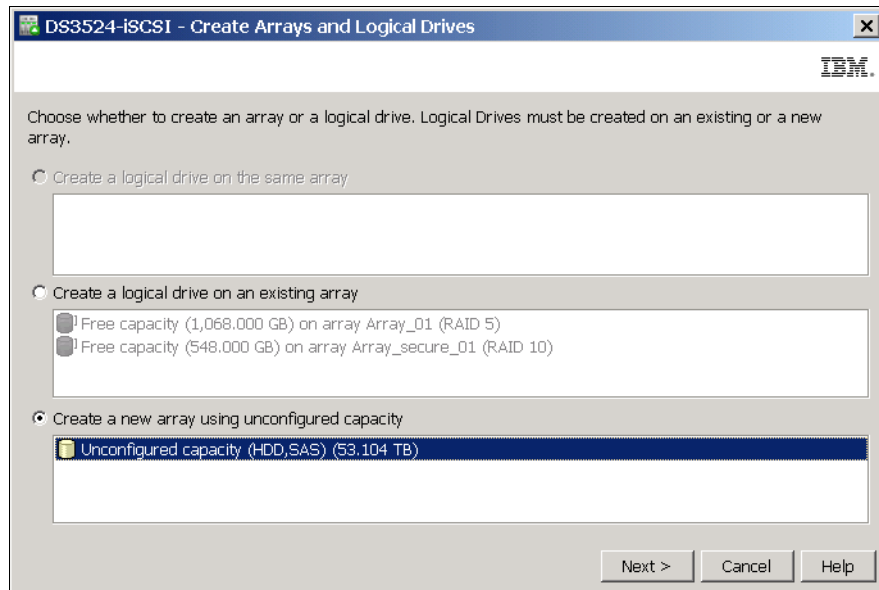


Figure 12-11 Configure new array

Detailed description of arrays and logical drive definitions are in Chapter 9, “Administration - Logical Tab” on page 219. This chapter describes all manual configurations needed for an array and a logical drive definition.

## 12.7 Map Logical Drives

This link in the Setup tab of DS Storage Manager navigates you to a window that is used to assign existing logical drive to existing hosts. This window is described in 11.6, “Define Additional Mapping” on page 291. All other configuration steps needed for assigning logical drives to hosts are described in Chapter 11, “Administration - Mappings Tab” on page 275.

## 12.8 Save Configuration

Use the **Save Configuration** link to open a window (Figure 12-12 on page 307) where you can save the logical configuration settings of a storage subsystem to a script file. You can then click **Tools** → **Load Storage Subsystem Configuration** in the **Enterprise Management** window to copy the configuration data back from the saved file to a storage subsystem with the exact hardware configuration. For more information about the **Load Storage Subsystem Configuration** option, refer to Chapter 6, “Administration - Enterprise Management” on page 157.




Figure 12-12 Save Configuration

When your DS3500 Storage Subsystem is fully configured, the **Save Configuration** option lets you save storage subsystem logical settings, logical drive configuration, topology, and mappings to the text file in the form of script file that also can be used by SMcli command line interface. If there are individual topology items or logical drives that you do not want to include in the saved configuration, you can edit the script file after it is saved. You can also use modified SMcli commands from this file to create a new logical devices in your storage subsystem. More information about SMcli are in Chapter 17, “Command-Line Interface (CLI)” on page 519.

The **Configuration** → **Save** option does not save the following settings:

- ▶ The age of the battery
- ▶ The controller time of day
- ▶ The nonvolatile static random access memory (NVS RAM) settings
- ▶ Any premium features
- ▶ The storage subsystem password
- ▶ The operating status and states of the hardware components
- ▶ The operating status and states of the arrays and logical drives

## 12.9 Manually Define Hosts

This link in the Setup tab of DS Storage Manager opens a window in which you can define a new host. This wizard is described in 11.2, “Defining Host” on page 277. All other configuration steps needed for assigning logical drives to hosts are described in Chapter 11, “Administration - Mappings Tab” on page 275.

## 12.10 Configure Ethernet Management Ports

This link enables you to change current IP addresses of your storage subsystem. Follow the steps and comments in 10.7, “Set Ethernet management ports” on page 270.

## 12.11 View/Enable Premium Features

Select this link to activate your optional licenses for Premium Features such as storage partitioning, copy services, encryption, and so on. This procedure is described in “Premium Features” on page 194.

## 12.12 Manage iSCSI Settings

This link is shown in the Setup tab only if an optional iSCSI daughter card is installed. It opens a window as shown in Figure 12-13 used for setting parameters of iSCSI protocol. All iSCSI configuration options and iSCSI port settings are described in Chapter 13, “Administration - iSCSI” on page 309.

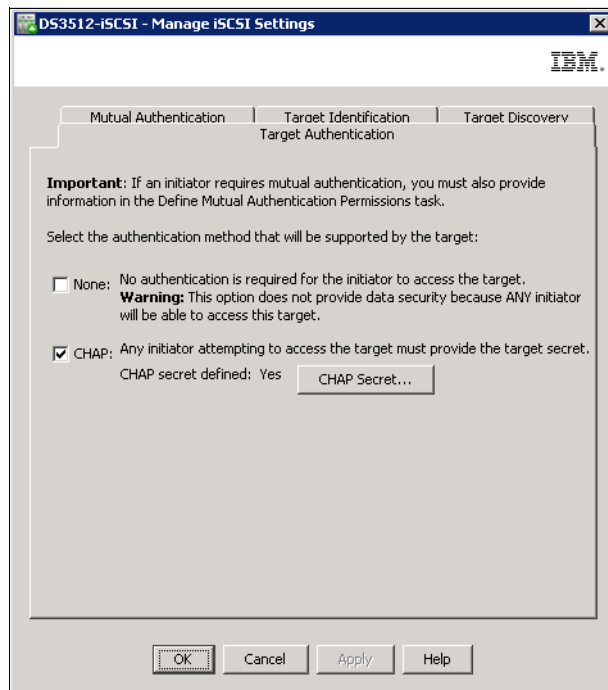


Figure 12-13 iSCSI protocol settings



## Administration - iSCSI

In this chapter we describe iSCSI management options that are available in the DS Storage Manager. These options are only for IBM System Storage DS3500 with optional iSCSI daughter cards installed.

In contrast with previous versions of DS Storage Manager for DS3300 Storage Subsystem (iSCSI model), there is no special iSCSI tab to manage all iSCSI settings. The most of iSCSI settings are done from the Storage Subsystem pull-down menu.

Theoretical information about iSCSI protocol are described in Chapter 1, “Disk attachment technology” on page 1.

## 13.1 Planning for iSCSI attachment

In addition to using SAS and Fibre Channel as the interface connection method, iSCSI host interfaces of DS3500 allow servers with specialized hardware iSCSI cards, or with software running over regular Ethernet cards, to run the iSCSI protocol as the connection method to attach the DS3500 Storage Subsystems. As with SAS or FC, before beginning an iSCSI deployment, plan in advance, understand, and document the network topology that will be used.

### Hardware initiators

As certain iSCSI hardware initiators might be supported for attachment to DS3500 in the future, check currently supported types and versions at the IBM interoperability web site:

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

**Note:** At the time of writing this book there is no iSCSI HW Initiator supported with DS3500.

Before you install and configure these adapters, make sure that you have installed the latest management application and the latest firmware code supported by IBM.

### Software initiators

Check currently supported iSCSI Software Initiators for various operating systems on the IBM interoperability web site:

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

**Note:** The native MPIO that is provided with the Microsoft iSCSI Software Initiator is supported at the time of writing this book, but use the DSM (MPIO Device Specific Module) that is provided with the DS Storage Manager software to make sure that failover and I/O access are correct. If the native MPIO is used for the Microsoft iSCSI Software Initiator, it can have unwanted effects in certain configurations.

### Storage Partitioning

Because the SAS/FC and iSCSI protocols provide radically different latency and throughput capabilities, and this mixture within a server might cause failover driver conflict, performance degradation, or potential data loss, you have to plan ahead using the following recommendations:

- ▶ Define separated partitions for SAS or FC based hosts from iSCSI based hosts. Avoid mixing them in same storage partition.
- ▶ A single host should not be configured for both iSCSI connections and SAS or FC connections to the storage system.

In order to define hosts and storage partitions in iSCSI, remember that you have to plan for iSCSI addresses for the host ports, and use the iSCSI Qualified Name (IQN) of the host you want to map.

### Network settings

Unlike traditional Fibre Channel, which requires special cabling and SAN switches, iSCSI can be run over an existing network infrastructure. However, in complex network environments, and to protect the integrity of the data in your DS3500 Storage Subsystem, and its continuous access, we suggest that, whenever possible, to try to isolate the iSCSI traffic in a dedicated network. The iSCSI multipathing architecture provides failover to the alternate controller in the



event of an outage situation. With MPIO, IBM provides DSM, which also offers load-balancing algorithms.

For better redundancy, you can increase the availability of your connections using redundant networks so a failure in one does not interrupt the remaining redundant connection.

Aside from the basic iSCSI connectivity parameters, such as IP address per target Ethernet port and associated iSCSI Qualified Names, you could plan in advance several optional configuration parameters, including enabling jumbo frames, configuring of a VLAN, and setting a specific Ethernet priority:

- ▶ Jumbo frames are created when the MTU is adjusted above 1500 bytes per frame, and they are set by port. The frame sizes supported are between 1501 and 9000 bytes. When using jumbo frames, ensure that all of the devices on your iSCSI network, including switches, initiators, and targets, are configured to use the same maximum jumbo frame size.
- ▶ VLAN: As previously mentioned, for performance and availability reasons have separate networks for redundant interfaces. If it is not possible to segregate an iSCSI storage system onto a physically separate LAN with the IBM DS3500 storage subsystems that are connected by iSCSI, you can use VLANs to maximize the potential performance.
- ▶ Ethernet priority: Ethernet priority, sometimes referred to as quality of service or class of service, is supported in the DS3500 series of storage systems. You can set the Ethernet priority of the target iSCSI interfaces to increase the class of service received within the network itself.

## Security

Unlike FC SANs or direct SAS connection, Ethernet networks can be more open, so to provide additional security, configure the following additional authentication protocols on the DS5000 storage subsystems:

- ▶ The Internet Storage Name Service (iSNS) protocol allows for automated discovery, management, and configuration of iSCSI devices on a TCP/IP network. iSNS servers offer additional security services through explicitly defined initiator-to-target mappings and simplified asset locators similar to that provided by DNS and WINS for IP address lookup facilities
- ▶ Challenge Handshake Authentication Protocol (CHAP) provides an additional security layer within the iSCSI network on the IBM Storage System DS3500 subsystem.

## 13.2 iSCSI Configuration summary

We cover settings of all mentioned parameters above in the following topics:

- ▶ Managing iSCSI Settings
  - Target Authentication
  - Mutual Authentication
  - Target Identification
  - Target Discovery
- ▶ Configuring iSCSI host ports on DS3500
- ▶ Session and statistics
  - View iSCSI Statistics
  - View/End iSCSI Sessions

We also cover differences in iSCSI host server attachment:

- ▶ View unassociated iSCSI Initiators
- ▶ Define new iSCSI Host
- ▶ Manage iSCSI Initiators

The main menu for iSCSI configuration can be found in the **Storage Subsystem** pull-down menu as shown in Figure 13-1. Links for configuration of iSCSI ports on DS3500 and for configuration of iSCSI protocol settings are also accessible from the Setup tab. These two settings are the most important, and should be done on a new DS3500 during initial configuration.

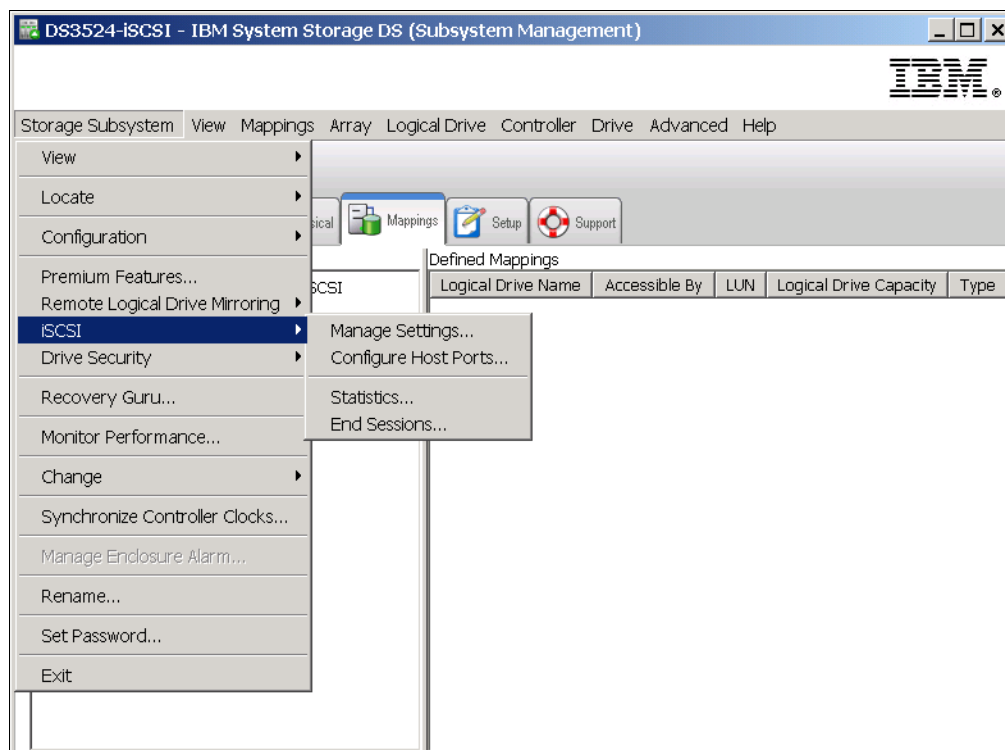


Figure 13-1 Basic iSCSI configuration setting menu

## 13.3 Manage iSCSI protocol settings

Settings that affect the security and discovery of iSCSI connections can be done here. We also discuss iSCSI security in 1.3.3, “iSCSI security considerations” on page 12.

### 13.3.1 Target Authentication

The Target Authentication (also called *uni-directional authentication*) is used to allow only configured iSCSI Initiators in servers to access the storage subsystem. When target authentication is enabled, the target (DS3500 subsystem) needs to authenticate itself against the initiator that attempts to access the storage subsystem. One Target CHAP secret must be defined for all such initiators.

Another type of authentication - the Mutual Authentication (also called *bidirectional authentication*) can only be configured on the iSCSI Initiator, when the initiator itself

authenticates against the storage subsystem. If this is required, you have to also configure an Initiator CHAP secret on the DS3500 by following the steps in 13.3.2, “Mutual Authentication” on page 314.

To configure target authentication, follow these steps:

1. Click **Storage Subsystem** → **iSCSI** → **Manage Settings**, and the Target Authentication tab (Figure 13-2) appears in a new window.

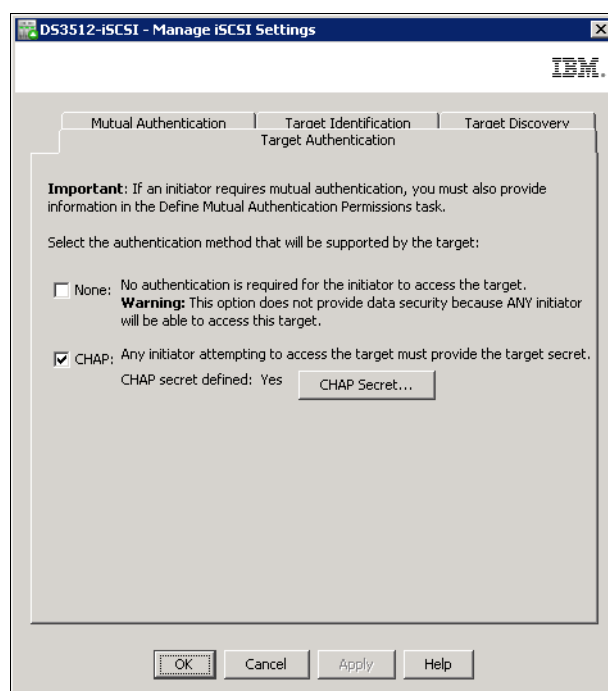


Figure 13-2 Manage iSCSI Settings - Target Authentication

2. This window, with the Target Authentication tab selected, offers two options:

- None
- CHAP

This setting affects the connection between an iSCSI Initiator and DS3500 iSCSI ports. The 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, you should enable CHAP. However, because the configuration of CHAP adds some complexity, you should set up and test all connections with no CHAP, and later implement the security.

3. If CHAP is enabled, define the CHAP target secret by clicking **CHAP Secret**. The Enter Target CHAP Secret window shown in Figure 13-3 on page 314 displays.
4. Type a CHAP secret that is at least 12 characters long. Repeat the text in second line to confirm.

Alternatively, click **Generate Random Secret** to generate a 16 character long secret.

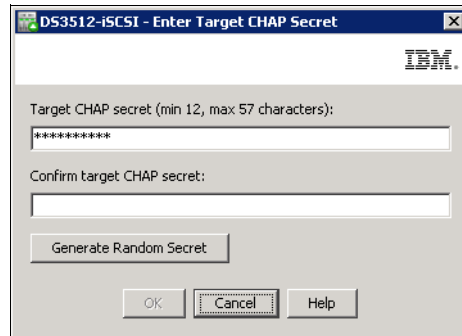


Figure 13-3 Target CHAP secret

5. Click **OK** to return to the Target Authentication window (Figure 13-2 on page 313). Click **OK** to save the modifications.

### 13.3.2 Mutual Authentication

With Mutual Authentication, a CHAP secret is configured that a specific initiator must use to get access to the storage subsystem. That is also called initiator authentication or bidirectional authentication. Each initiator can have its own CHAP secret.

To activate Mutual Authentication, perform the following steps:

1. In the Manage iSCSI Settings window (Figure 13-2 on page 313), select the Mutual Authentication tab.
2. In the tab shown in Figure 13-4 on page 315, select a host port that has already been defined as described in 13.7.2, “Defining new iSCSI host” on page 326. The host port is referred to as the Initiator Label. In this example, there is no CHAP secret defined (No permissions entered) for this initiator as shown in the Permission Set? column. Click **CHAP Secret**.

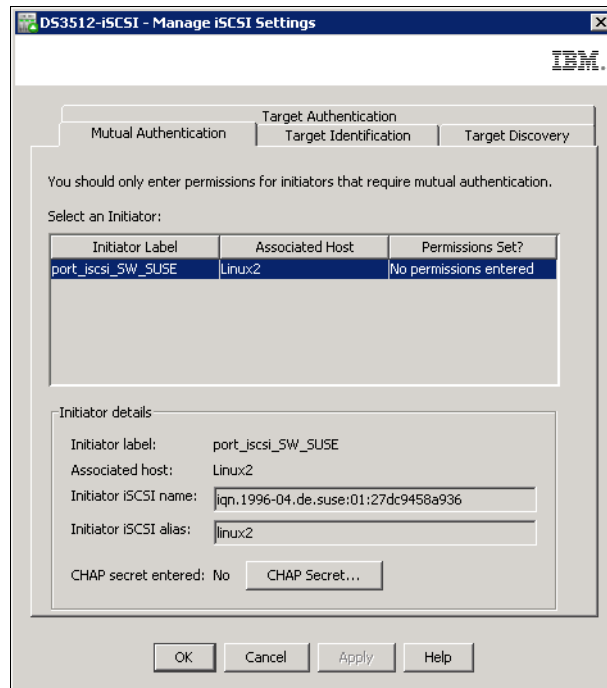


Figure 13-4 iSCSI Mutual Authentication

3. In the new window as shown in Figure 13-5, type an Initiator CHAP Secret with at least 12 characters twice. The same Initiator CHAP secret must be set for iSCSI Initiator on any given host server. Click **OK** to continue.

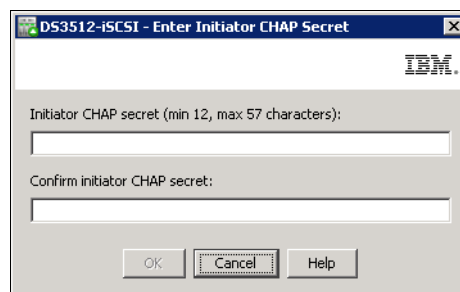


Figure 13-5 Storage Manager - Initiator CHAP Secret

4. The Permission Set column in Figure 13-4 now indicates that the port has a CHAP secret defined. Define CHAP secrets for all remaining host ports, if needed, using the same procedure. Each host with iSCSI Initiator can have a separate Initiator CHAP secret. When this is complete, click **OK** to exit this task.

### 13.3.3 Target Identification

The Target Identification tab allows you to set up an alias name for the iSCSI Target device (your DS3500) that can be discovered by iSCSI Initiators on hosts. In iSCSI protocol, the storage is considered the iSCSI Target, whereas the host server is the iSCSI Initiator. You can check your DS3500 Storage Subsystem IQN name in this tab.

Set up an iSCSI target alias using the following steps:

1. Click the **Target Identification** tab.

2. Enter an alias name in the second line in this tab as shown in Figure 13-6, and click **OK**.

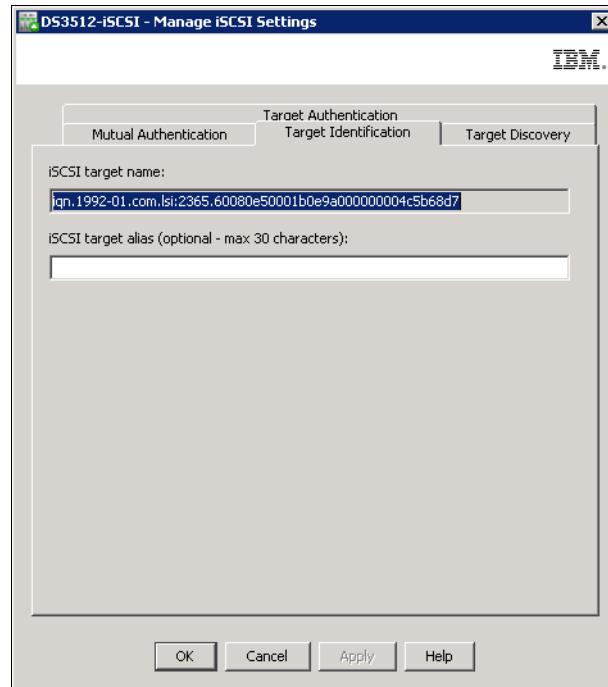


Figure 13-6 iSCSI Target Identification

### 13.3.4 Target Discovery

We discussed techniques for target discovery in 1.3.2, “iSCSI discovery” on page 11. The Target Discovery tab in the Manage iSCSI Setting window (Figure 13-7 on page 317) can be used to specify if an Internet Storage Name Service (iSNS) will be used to register this iSCSI Target for easier location by iSCSI Initiators.

To change default target discovery, perform these steps:

1. Select the **Target Discovery** tab.
2. Select **Use iSNS server** as shown in Figure 13-7 on page 317.

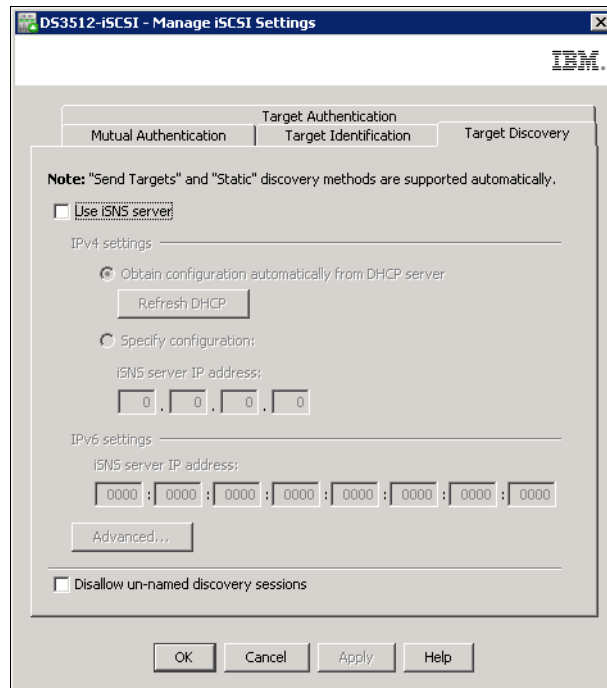


Figure 13-7 Target Discovery - Enable iSNS server

3. Enter the IPv4 or IPv6 address of the iSNS server, or select **Obtain configuration automatically from DHCP server**.
4. Click **Advanced** to define a custom port (if needed) 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.

## 13.4 Configure iSCSI Host Ports

The iSCSI host ports are the 1 Gb Ethernet ports on each DS3500 controller with a iSCSI daughter card that are used to connect hosts to this storage subsystem by iSCSI protocol. Each controller has four RJ45 ports available on the iSCSI daughter card (see 4.1.4, “DS3500 controller with optional iSCSI host port adapter” on page 89).

The Configure iSCSI Host Ports task (see Figure 13-8 on page 318) is activated by clicking **Storage Subsystem** → **iSCSI** → **Configure Host Ports**. This task can also be activated from the Setup tab, and from the Physical tab if any controller is selected. It is used to configure the network settings of the four 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
- Whether to allow ICMP ping responses
- Requirement and size of jumbo frames

**Note:** Do not use the same network subnet for more than one iSCSI port in one controller. Although, for data paths high availability, the same network subnet is necessary for the pair of iSCSI ports when each port is on different controller. Also, do not use the same network subnet for iSCSI port as for the special management port.

To configure the network settings of the iSCSI host ports, follow these steps:

1. Click **Storage Subsystem** → **iSCSI** → **Configure Host Ports**. A window opens as shown in Figure 13-8. Our example window shows an active IPv4 configuration with the default IP address for the iSCSI port.

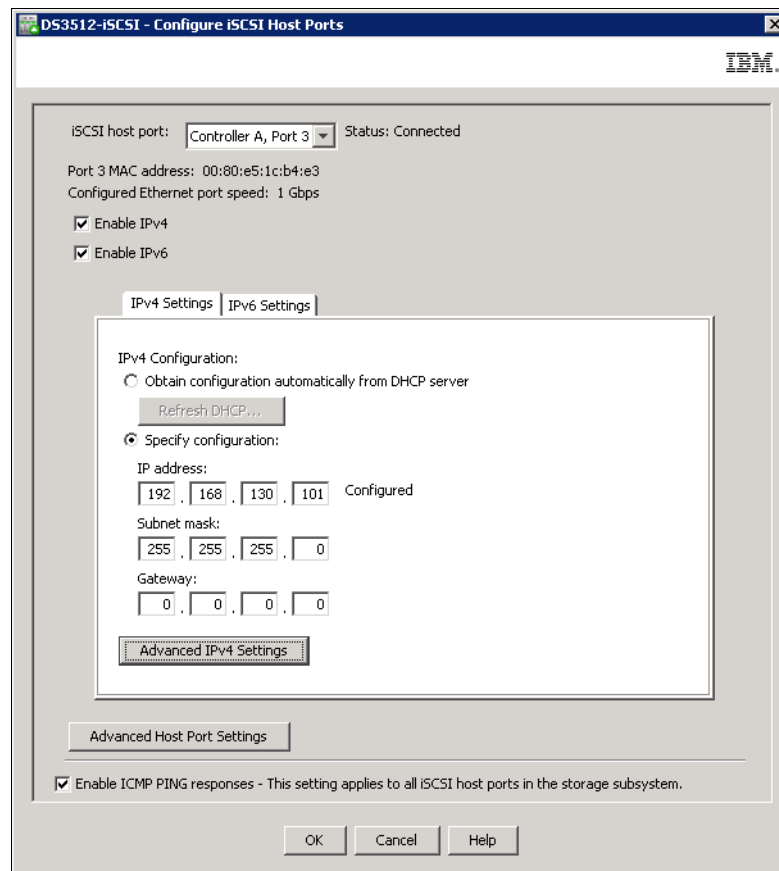


Figure 13-8 Configure iSCSI Host Ports

2. Select the iSCSI host port to be configured from the **iSCSI host port** drop-down menu as shown in Figure 13-9 on page 319.






Figure 13-9 iSCSI host port selection

Default subnet mask is 255.255.255.0 and default IP addresses for iSCSI ports are:

A3: 192.168.130.101  
A4: 192.168.131.101  
A5: 192.168.132.101  
A6: 192.168.133.101

B3: 192.168.130.102  
B4: 192.168.131.102  
B5: 192.168.132.102  
B6: 192.168.133.102

3. IPv4 and IPv6 are both enabled by default. If you do not need to use IPv6, click to clear the **Enable IPv6** check-box.
4. On the IPv4 Settings tab, 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, configure your DHCP server to use MAC addresses as the base for the fixed IP address assignment.

**Warning:** Avoid using DHCP for the target ports. If you use DHCP, you should assign DHCP reservations so that leases are maintained consistently across restarts of the DS3500 storage subsystem. If static IP reservations are not provided or you lose the DHCP configuration, the iSCSI Initiators can lose connection to the DS3500 controller iSCSI ports and will 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.

5. In **Advanced IPv4 Settings** window (Figure 13-10 on page 320), you can enable native VLAN support and specify a VLAN ID that will be used by the selected port. In addition, Ethernet priority can be set here.

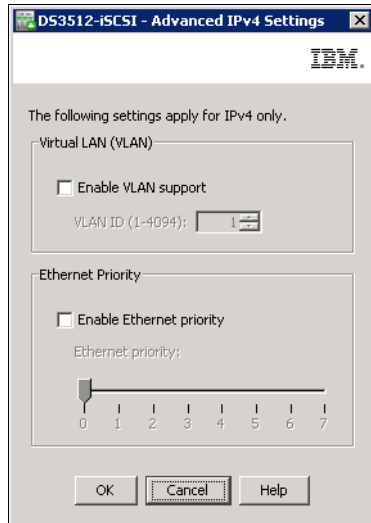


Figure 13-10 Advanced IPv4 Settings

Similar configurations are available for the IPv6 settings.

6. After configuring IPv4 and /or IPv6, you can select **Advanced Host Port Settings** as shown in Figure 13-11. 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 the default (3260), select the **Use custom listening port** check box and specify the port. A custom port can be between port 49152 and 65535.

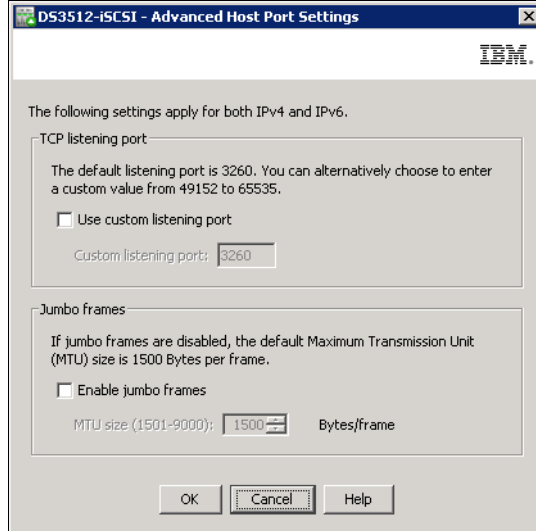


Figure 13-11 Advanced iSCSI Host Port Settings

Specify the jumbo frame size after enabling jumbo frames 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. Configure the remaining iSCSI host ports in the **iSCSI host port** drop-down menu (Figure 13-9 on page 319) in the same way.
8. In bottom part of the Configure iSCSI Host Ports window (Figure 13-8 on page 318), you have the option to disable ping (ICMP echo) requests for all iSCSI host ports. A reset of all ports is required to apply this setting.

9. After setting all the options, click **OK** to save the modifications. In Figure 13-12, you are advised which ports will be reset to apply the configuration. Active iSCSI sessions that can be connected to this ports will be closed. Click **Yes** to proceed.

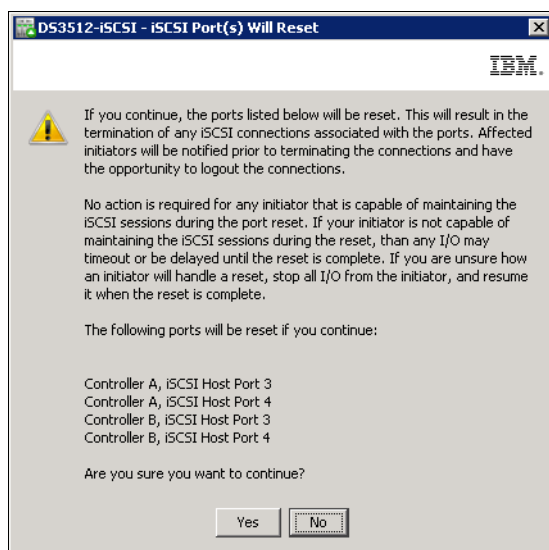


Figure 13-12 iSCSI Port Reset warning

## 13.5 View/End iSCSI Sessions

The View/End iSCSI Sessions window shows all initiators that have active sessions. Details for each session can be viewed and saved using the following steps:

1. Select **Storage Subsystem** → **iSCSI** → **End iSCSI Sessions** in Subsystem Management pull-down menu, and a window as in Figure 13-13 on page 322 appears. In the top pane you can see all the current sessions.
2. Select a session to see the details in the lower pane.

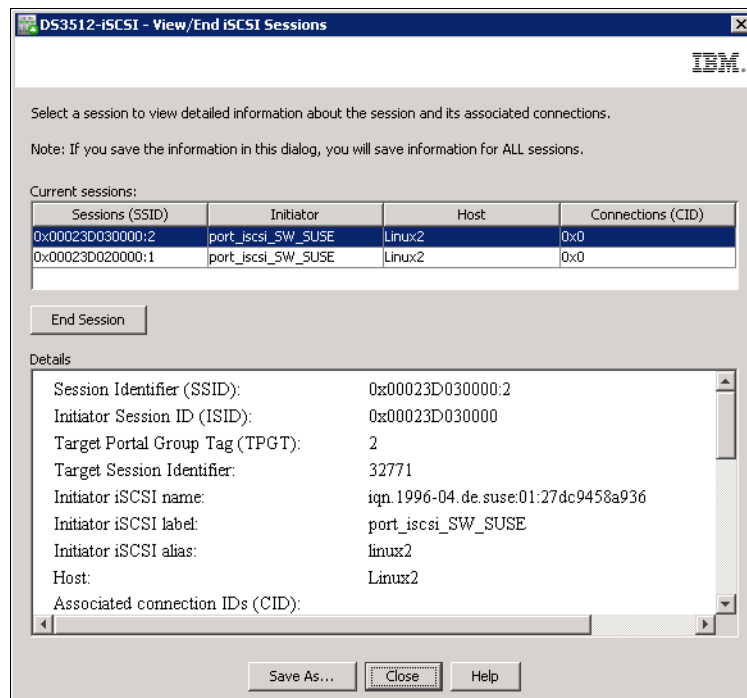


Figure 13-13 View/End iSCSI Sessions

3. To terminate a session, select it in the top pane and click **End Session**. If the initiator is still available and active, the session can be re-established.
  4. In the bottom pane, click **Save As** to copy the session details into a text file.
- Example 13-1 on page 323 shows the complete session information for the selected initiator in Figure 13-13. This output can help verify that the host setup was done correctly.

### Example 13-1 sample session information

---

|                                      |                                           |
|--------------------------------------|-------------------------------------------|
| Session Identifier (SSID):           | 0x00023D030000:2                          |
| Initiator Session ID (ISID):         | 0x00023D030000                            |
| Target Portal Group Tag (TPGT):      | 2                                         |
| Target Session Identifier:           | 32771                                     |
| Initiator iSCSI name:                | iqn.1996-04.de.suse:01:27dc9458a936       |
| Initiator iSCSI label:               | port_iscsi_SW_SUSE                        |
| Initiator iSCSI alias:               | linux2                                    |
| Host:                                | Linux2                                    |
| Associated connection IDs (CID):     |                                           |
| CID:                                 | 0x0                                       |
| Ethernet port identifier:            | Controller B, port 3                      |
| Initiator IP address:                | 192.168.130.120                           |
| Negotiated login parameters:         |                                           |
| Authentication method:               | None                                      |
| Header digest method:                | None                                      |
| Data digest method:                  | None                                      |
| Maximum connections:                 | 1                                         |
| Target alias:                        | DS3500-iSCSI                              |
| Initiator alias:                     | linux2                                    |
| Target IP address:                   | 192.168.130.102                           |
| Target portal group tag:             | 2                                         |
| Initial R2T:                         | Yes                                       |
| Maximum burst length:                | 65536                                     |
| First burst length:                  | 8192                                      |
| Default time to wait:                | 2                                         |
| Default time to retain:              | 0                                         |
| Maximum outstanding R2T:             | 1                                         |
| Error recovery level:                | 0                                         |
| Maximum receive data segment length: | 8192                                      |
| Target name:                         | iqn.1992-01.com.lsi:2365.60080e50001b0e9a |
| Initiator name:                      | iqn.1996-04.de.suse:01:27dc9458a936       |

---

5. Click **Close** to return.

## 13.6 View iSCSI Statistics

Click **Storage Subsystem** → **iSCSI** → **Statistics** to get information about transfers in the OSI layers 2 to 5. Figure 13-14 on page 324 shows an example of the Ethernet MAC statistics in Raw mode.

Transfers can be measured in Raw mode or Baseline mode. The difference between the two modes is the time since the transfers 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 using **Set Baseline**. This button resets all base line counters to zero and starts counting from the beginning again.

Use the **Save As** button to save all current statistics in a \*.csv file. Raw and baseline mode statistics for all levels are saved.

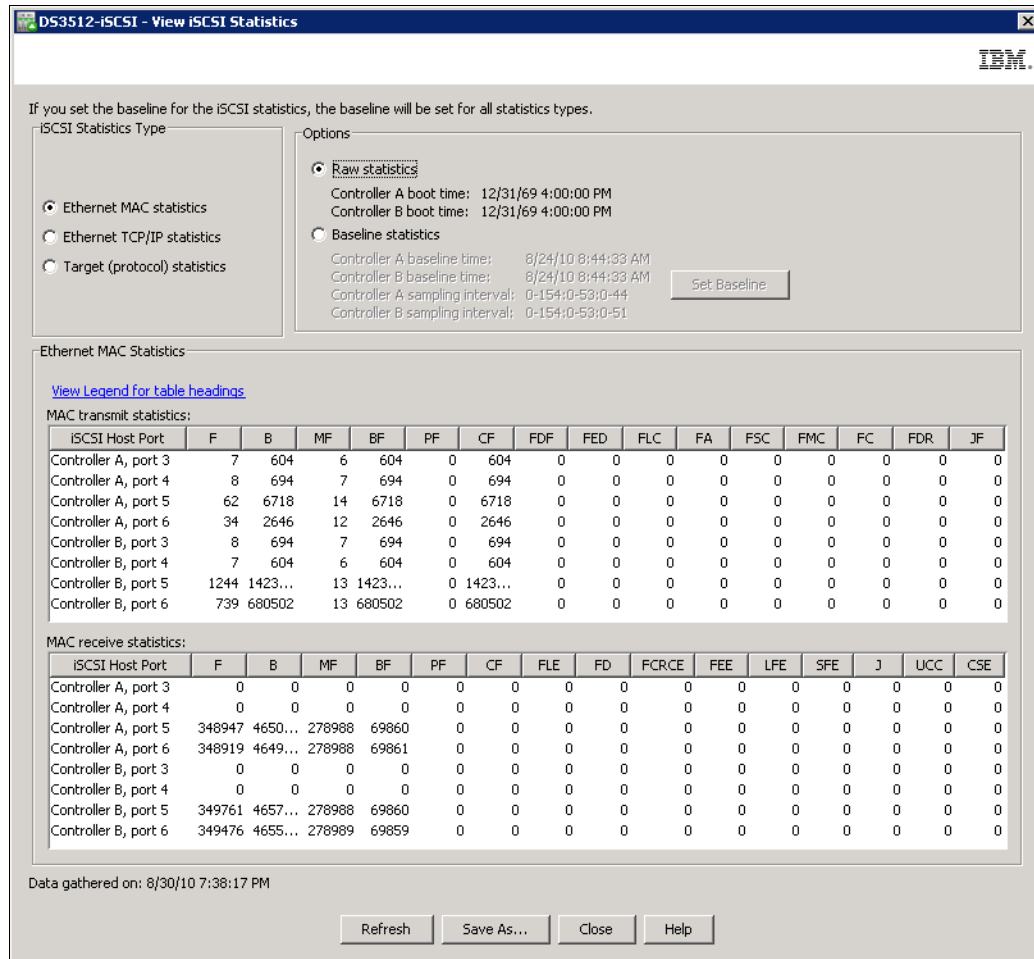


Figure 13-14 View iSCSI statistics for TCP/IP

The **View legend for table headings** link opens the Online Help window where description is provided to translate column names into more meaningful names for each statistic type. See Figure 13-15 on page 325.

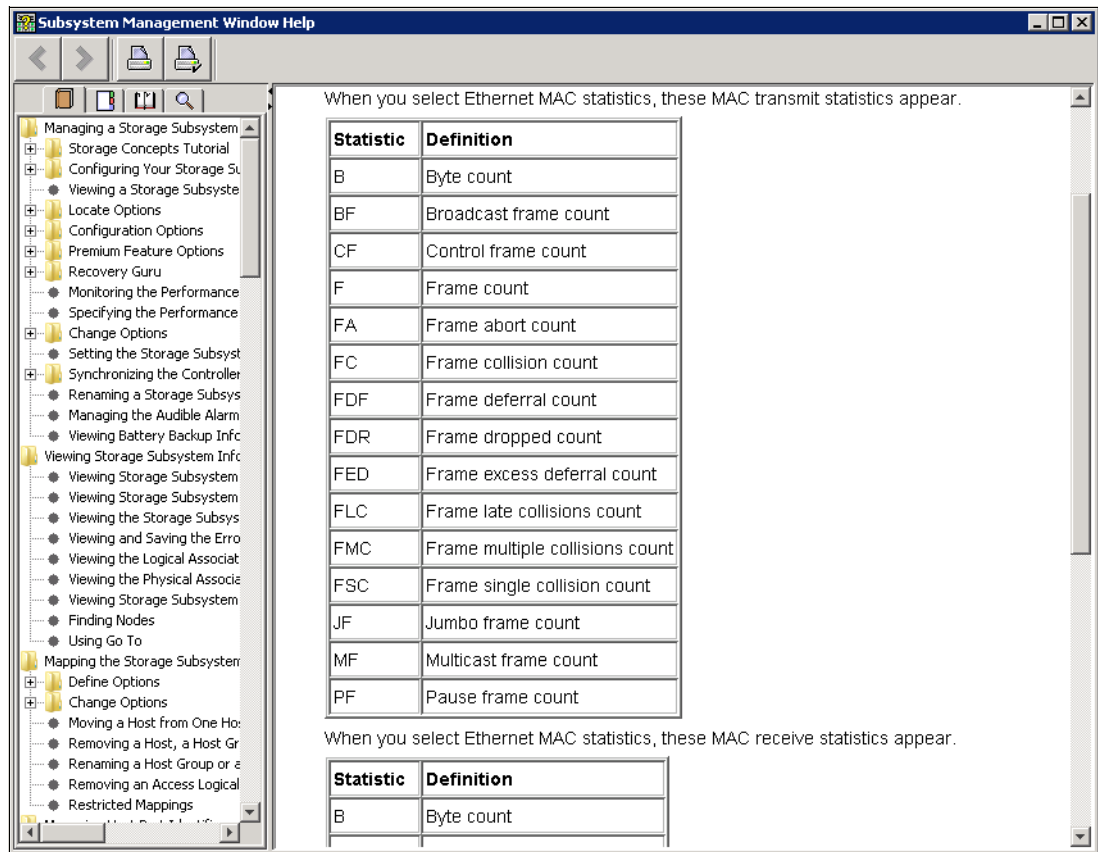


Figure 13-15 iSCSI context Help

## 13.7 Defining iSCSI hosts

Defining host using an iSCSI Initiator is slightly different than definition of host with FC or SAS HBAs. In case of iSCSI HW Initiators you can see as many iSCSI host port identifiers as iSCSI HBAs installed in such a host. In case of iSCSI SW Initiator, you see just one iSCSI host port identifier, even if you have more Ethernet cards used for iSCSI in your host server. Multiple redundant data paths are presented as multiple iSCSI sessions in this case as described in 13.5, “View/End iSCSI Sessions” on page 321.

Definition of a Host with traditional SAS or FC HBAs is described in 11.2, “Defining Host” on page 277 in detail. In the following sections we show iSCSI host definition differences.

### 13.7.1 View Unassociated iSCSI initiators

A Host definition is provided in Mapping tab of DS Storage Manager. To see if servers with active iSCSI Initiator are available to be defined as a host, you can view unassociated host port identifiers.

Several methods can be used to show unassociated host ports. For example, you can select **Mappings** → **View Unassociated Host Port Identifiers** in the Subsystem Management pull-down menu. A new window as shown in Figure 13-16 on page 326 appears. You can see World Wide Port Names for SAS or FC HBAs, and if an iSCSI Initiator is already configured and authenticated for access to your DS3500 subsystem, you will see the IQN name of this iSCSI Initiator. An iSCSI SW initiator in SUSE Linux is shown in Figure 13-16 on page 326.

The IQN name offers an easier way to recognize the operating system of the host. If you see an iSCSI initiator in this window, you can continue to define a new iSCSI Host.




Figure 13-16 Unassociated iSCSI Software Initiator

**Note:** You can see iSCSI Initiator IQN name of your server in this window only if you already discovered and authenticated your DS3500 subsystem as an iSCSI Target on your host server.

### 13.7.2 Defining new iSCSI host

A new iSCSI host can be defined by the same procedure as described in 11.2, “Defining Host” on page 277. The only different settings have to be done in **Specify Host Port Identifier** window as shown in Figure 13-17 on page 327:

1. Select **iSCSI** as host interface type, and if you select the **Known unassociated host port identifier** selection box, you can find your iSCSI Initiator IQN name in the corresponding pull-down selection menu.
2. Type an **User Label** (Alias), and select **Add** to add this iSCSI Initiator as a host port to your newly defined host. For iSCSI software initiator you add only one IQN name per host.
3. Click **Next** to continue and finalize the host definition.



**DS3512-iSCSI - Specify Host Port Identifiers (Define Host)**

The host communicates with the storage subsystem through its host bus adapters (HBAs) or its iSCSI initiators where each physical port has a unique host port identifier. In this step, select or create an identifier, give it an alias or user label, then add it to the list to be associated with host Linux2.

[How do I match a host port identifier to a host?](#)

Choose a host interface type:  
 iSCSI

Choose a method for adding a host port identifier to a host:

☒ Add by selecting a known unassociated host port identifier

Known unassociated host port identifier:  
 iqn.1996-04.de.suse:01:27dc9458a936 Refresh

☐ Add by creating a new host port identifier

New host port identifier (max 223 characters):

User Label (30 characters maximum):  
 port\_iscsi\_SW\_SUSE

Add Remove

Host port identifiers to be associated with the host:

| Host Port Identifier | Alias / User Label |
|----------------------|--------------------|
|                      |                    |

< Back Next > Cancel Help

Figure 13-17 Define iSCSI Host: Host Port Identifier specification

Mapping of logical drives to iSCSI host is the same as for every other host. It does not depend on the type of host attachment interface. It is described in Chapter 11, “Administration - Mappings Tab” on page 275.

### 13.7.3 Manage iSCSI host ports

You can check or modify your current definition of iSCSI hosts in the **Manage Host Port Identifier** menu of the DS Storage Manager.

1. Click **Mappings** → **Manage Host Port Identifiers**. A new window as in Figure 13-18 on page 328 opens.
2. Select your host in **Show host port identifiers associated with:** field, and you can see details about your iSCSI host port identifier in the **Host port identifier information** field.

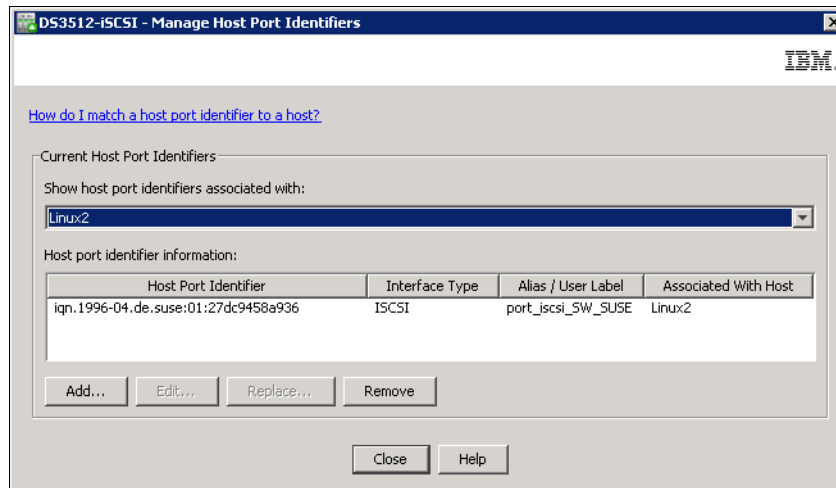


Figure 13-18 iSCSI Host Port information

3. Theoretically, you can **Add** or **Replace** the iSCSI port identifier, but practically, if you have iSCSI software initiator tightly coupled with your server operating system, these options cannot be used.



## Administration - Support

In this chapter, we describe the support functions of the IBM System Storage DS3500. These functions can help with troubleshooting and maintenance tasks. Specifically, we describe the functions that are accessed through the DS Storage Manager (Subsystem Management) tab, and the maintenance, troubleshooting and recovery functions that are accessed from the DS Storage Manager (Subsystem Management) Advanced menu. We also describe the tools available in Storage Manager to update the controller firmware and NVSRAM, and the drive and ESM firmware.

## 14.1 The Subsystem Management support tab

In this section we introduce the main support functions that are accessed from the DS Storage Manager (Subsystem Management) support tab (Figure 14-1).

**Note:** The support tab only contains the four most commonly used support tasks. More advanced support functions are accessed from the **Advanced** menu in the Subsystem Management Window.

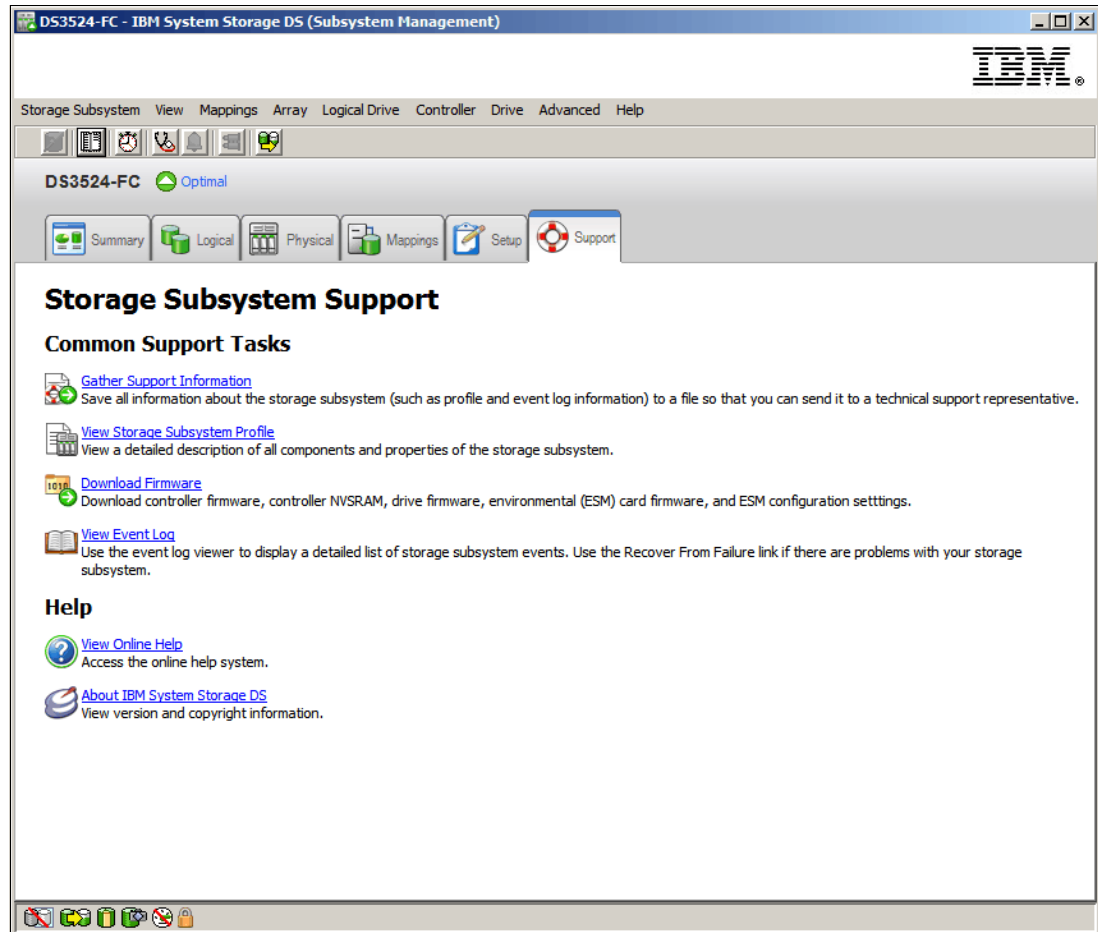


Figure 14-1 The DS Storage Manager (Subsystem Management) Support Tab

The Support tab (Figure 14-1) provides links to the following support functions:

- ▶ **Gather Support Information:** Clicking this link will open a wizard to save the DS3500 storage subsystem support information. This function is described in 14.2, “Gather Support Information” on page 331.
- ▶ **View Storage Subsystem Profile:** Clicking this link will open the DS3500 profile form viewing. This function is described in 14.3, “View Storage Subsystem Profile” on page 338.
- ▶ **Download Firmware:** Clicking this link will start the firmware update wizard. This function is described in 14.5, “Download firmware” on page 349
- ▶ **View Event Log:** Clicking this link will open the DS3500 storage subsystem event log for viewing. This function is described in 14.6, “View Event Log” on page 385.

- **View Online Help:** Clicking this link displays the DS Storage Manager Online Help and is further discussed in 14.15, “View Online Help” on page 436.
- **About IBM System Storage DS:** Clicking this link displays the version information of the DS Storage Manager software as described in 14.16, “About IBM System Storage DS Storage Manager” on page 436.

## 14.2 Gather Support Information

This wizard is an automated way to collect data for remote troubleshooting and analysis of problems with the storage management software and hardware. All of the files that are gathered are compressed into a single archive that is sometimes referred to as the CASD log (Collect All Support Data). You can then send the compressed archive to your IBM Customer Support representative for troubleshooting and further analysis.

**Attention:** In situations where a critical event has occurred that results in the loss of access to physical drives, logical drives, or arrays, do not attempt to resolve the problem yourself. Instead you should always save the CASD log and contact your IBM Customer Support representative for assistance.

There are two ways to open the Collect All Support Data window:

- Click the **Gather Support Information** link in the main Support tab (Figure 14-1 on page 330) to open the **Collect All Support Data** window (Figure 14-3 on page 332).
- Click **Advanced** → **Troubleshooting** → **Support Data** → **Collect** as shown in Figure 14-2 to open the **Collect All Support Data** window (Figure 14-3 on page 332).




Figure 14-2 Open the Collect All Support Data window from the Advanced menu

The **Collect All Support Data** window allows you to specify the filename and directory location where the DS3500 support data will be saved. The **Which data am I collecting?** link in the window displays the relevant section of the Storage Manager help file that contains the complete list of files collected by the wizard.

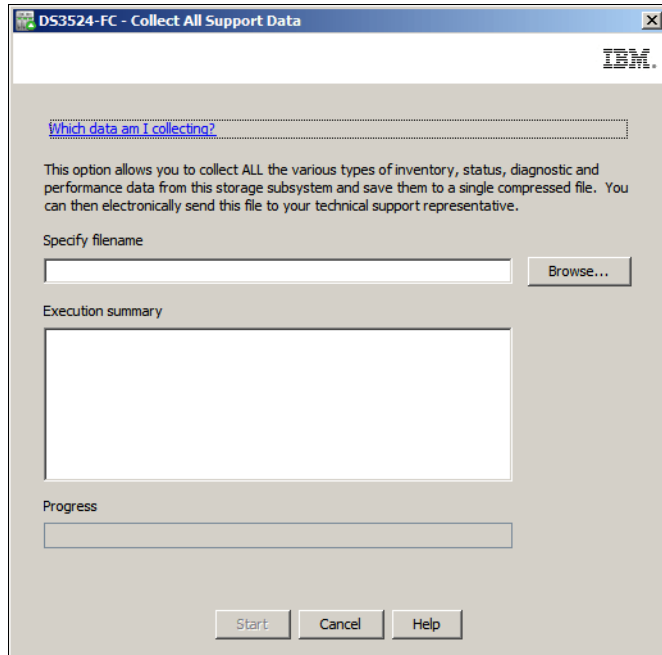


Figure 14-3 Collect All Support Data window

## 14.2.1 Saving the Support Data

Use the following procedure to save the support data:

1. Open the Collect All Support Data window and click the **Browse** button to specify the file name and folder location where the wizard will save the collected support data as shown in Figure 14-4.

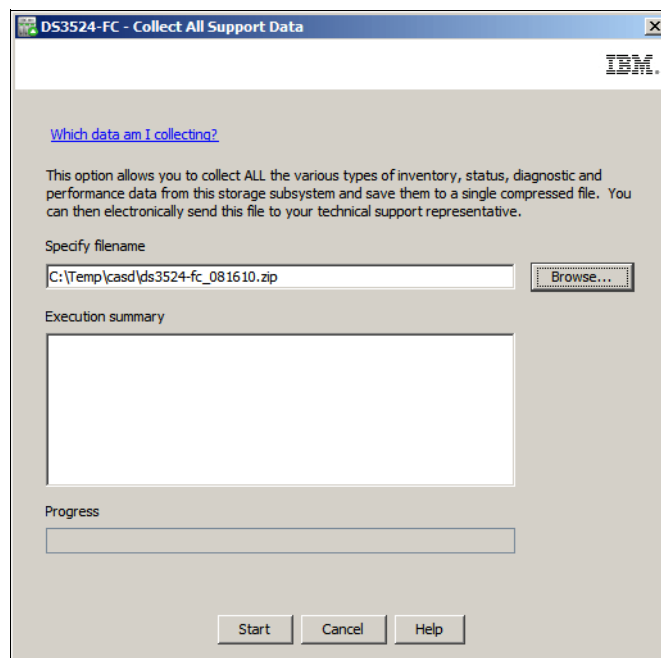


Figure 14-4 Collect All Support Data - specify the filename and folder location

2. Click the **Start** button to begin the data collection as shown in Figure 14-5 on page 333.

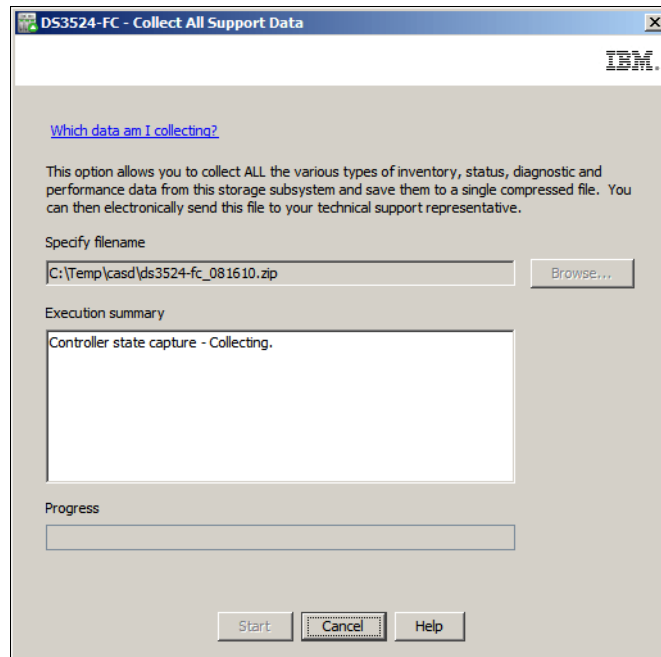


Figure 14-5 Collect All Support Data - click the Start button to begin the data collection

3. The Collect All Support Data window will display a progress bar while the data is being collected as shown in Figure 14-6.

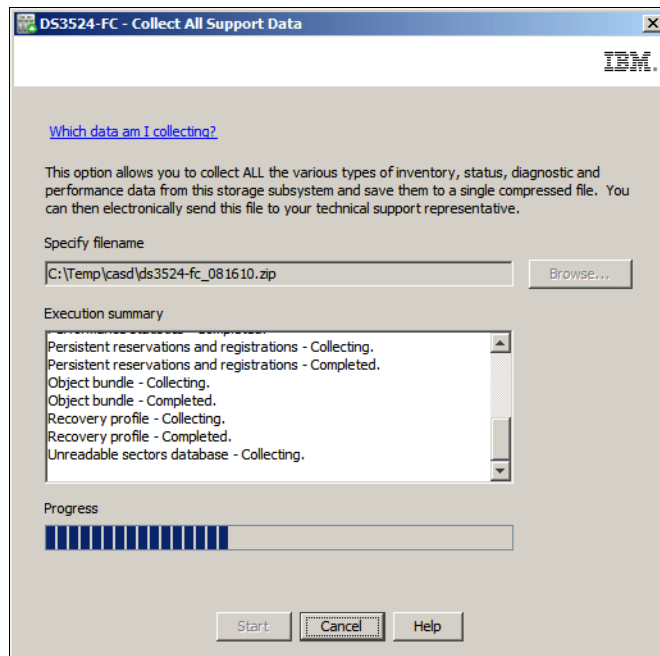


Figure 14-6 Collect All Support Data - data collection in progress

4. The window displays a message indicating when the support data has been collected and saved to the file that you specified in step 1, see Figure 14-7 on page 334.

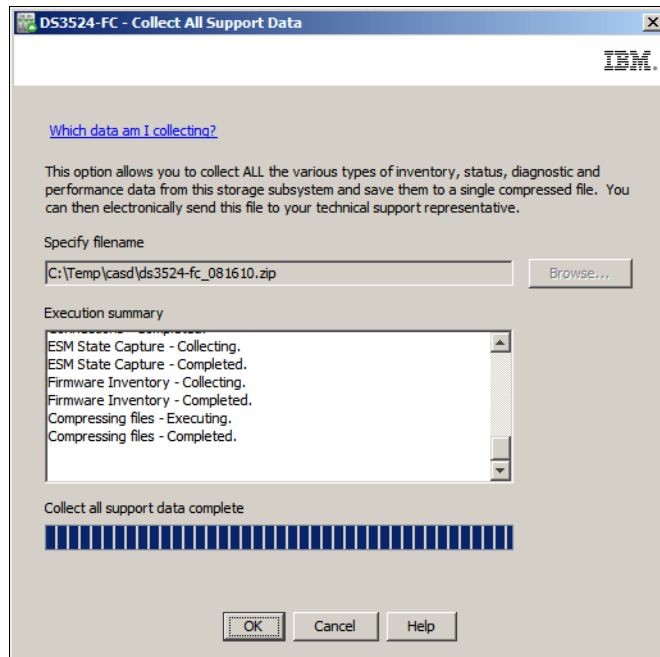


Figure 14-7 Collect All Support Data - Save Completed

5. Send the compressed file containing the collected support data to your IBM Customer Support representative for troubleshooting and further analysis.

## 14.2.2 Support Data Automatic Settings

DS Storage Manager incorporates the option to enable automatic collection of the support data. When enabled in the case of a critical event, the support data file is collected and automatically transferred to the directory specified and all the information relevant for troubleshooting by your support representative is preserved. In case of consecutive critical events occur, only the information related to the last event is stored in the auto save support bundle.

**Tip:** Enable the Automatic Collection option 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 DS3500 system to collect the information.
- Have the SM Monitor process running on the workstation or host where you want to collect the logs.

The settings for this function are accessed through the Storage Manager menu:

1. Click **Advanced** → **Troubleshooting** → **Support Data** → **Automatic Settings** (Figure 14-8) to open the Automatic Support Data Collection window (see Figure 14-9 on page 335).




Figure 14-8 Open the Automatic Settings function from the Advanced menu



2. Click the **Automatically collect support data for critical events** check box in the new window (Figure 14-9) to enable the automatic data collection and, optionally, you can use the **Change** button to specify the destination directory for the collected data. The value of the **Support data filename** cannot be changed.

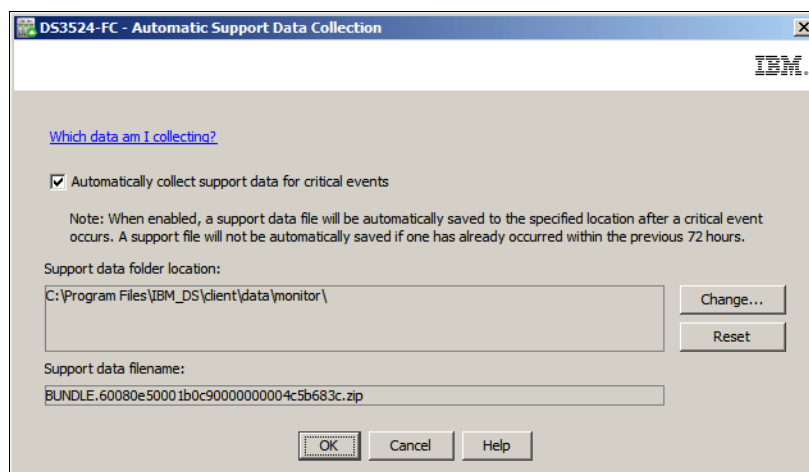


Figure 14-9 The Automatic Support Data Collection window

### 14.2.3 Collect drive data

You can save diagnostic data from all of the physical drives installed in the DS3500 storage subsystem. This drive diagnostic data is only readable by your IBM Customer Support representative, and this option is normally only used at their specific request.

**Note:** This log is saved as part of the Collect All Support Data log (14.2.1, “Saving the Support Data” on page 332).

Each drive in the DS3500 storage subsystem keeps a log of statistical information that can be used by your IBM Customer Support representative to analyze the performance of your drives and for troubleshooting any drive problems that might exist.

1. In the DS Storage Manager (Subsystem Management) window, click **Advanced** → **Troubleshooting** → **Collect Drive Data** to start the **Collect Drive Data** window as shown in Figure 14-10.

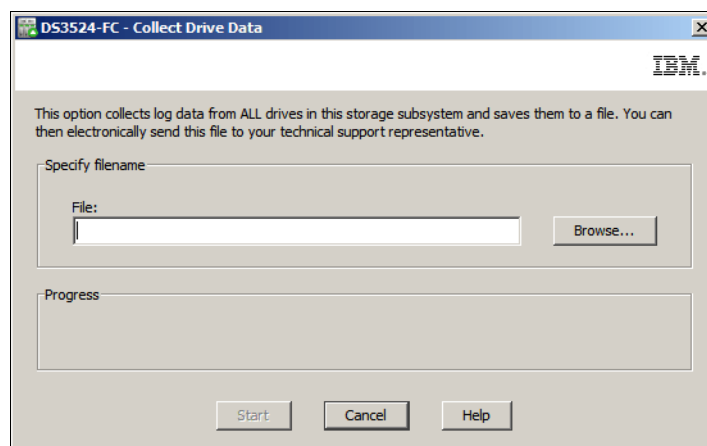


Figure 14-10 Collect Drive Data

2. Click the **Browse** button to specify the location and filename of the compressed file where you will save the drive diagnostic data. Then click the **Start** button to save the file. Figure 14-11 shows the window after the drive diagnostic data has been saved.

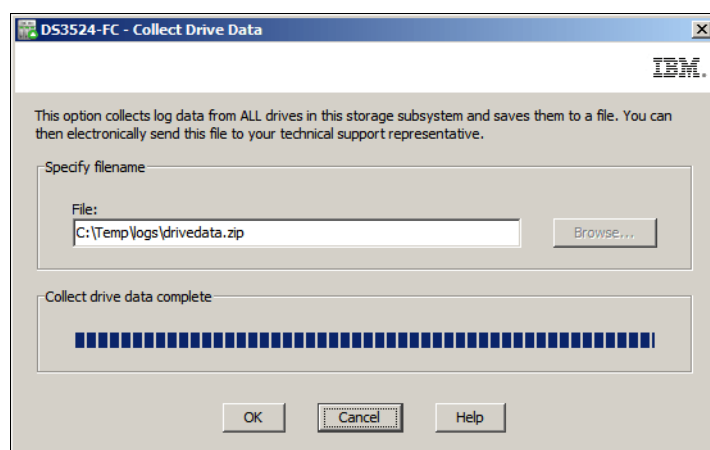


Figure 14-11 Collect drive data complete

## Collect state capture

You use this option to save detailed diagnostic information about the current state of your storage subsystem to a file. You can then send the captured information to your IBM Customer Support representative for analysis.

**Note:** This log is saved as part of the Collect All Support Data log (14.2.1, “Saving the Support Data” on page 332).

**Attention:** Using this option has the potential to cause the storage subsystem to become unresponsive to both the host and the storage management station. Use this option only under the guidance of your IBM Customer Support representative.

1. In the DS Storage Manager (Subsystem Management) window, click **Advanced** → **Troubleshooting** → **Capture State Information** to open the **State Capture** window as shown in Figure 14-12 on page 337.

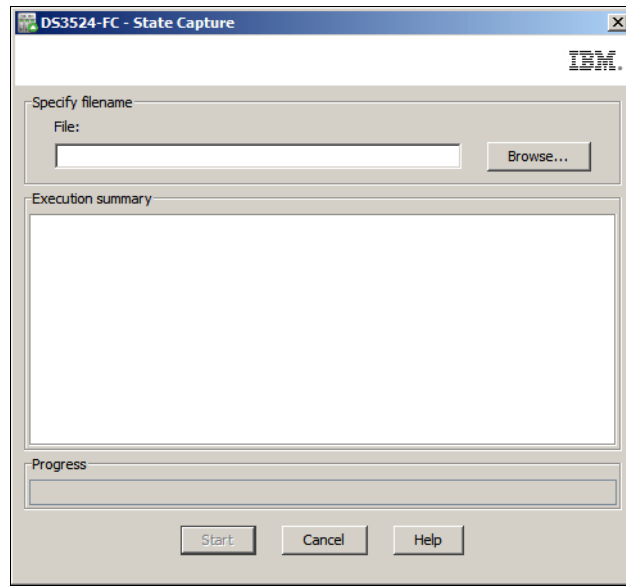


Figure 14-12 State Capture

2. Click the **Browse** button to specify the location and filename of the state capture information file (the default file extension is .dmp). Then click the **Start** button to begin saving the file. You will be asked to confirm that you want to save the state capture data. Type yes in the box to confirm and then click **OK** to continue.

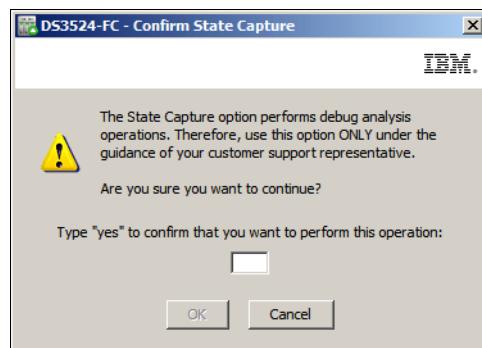


Figure 14-13 Confirm State Capture

3. Figure 14-14 on page 338 shows the State Capture window after the state capture information has been saved to a file.

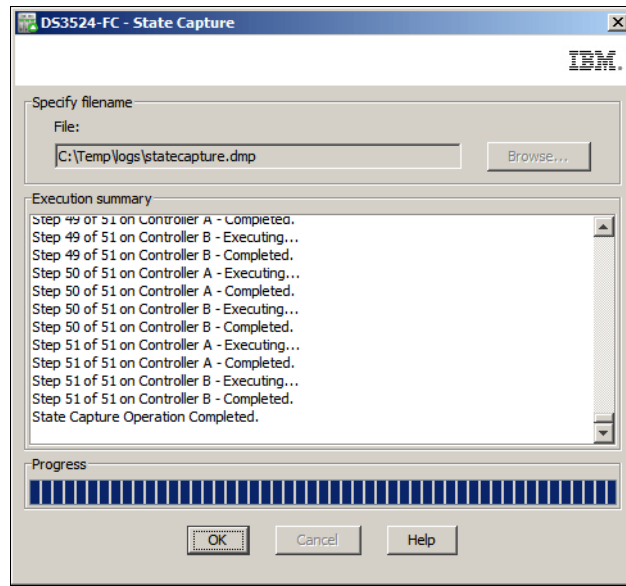


Figure 14-14 State Capture completed

## 14.3 View Storage Subsystem Profile

We introduced the Storage Subsystem Profile in 7.1.1, “Storage Subsystem Profile” on page 178. The profile 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

Here 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.

Follow these steps to view the DS3500 storage subsystem profile:

1. Click **View Storage Subsystem Profile** in the Support tab (Figure 14-1 on page 330). The Storage Subsystem Profile window appears as shown in Figure 14-15 on page 339.

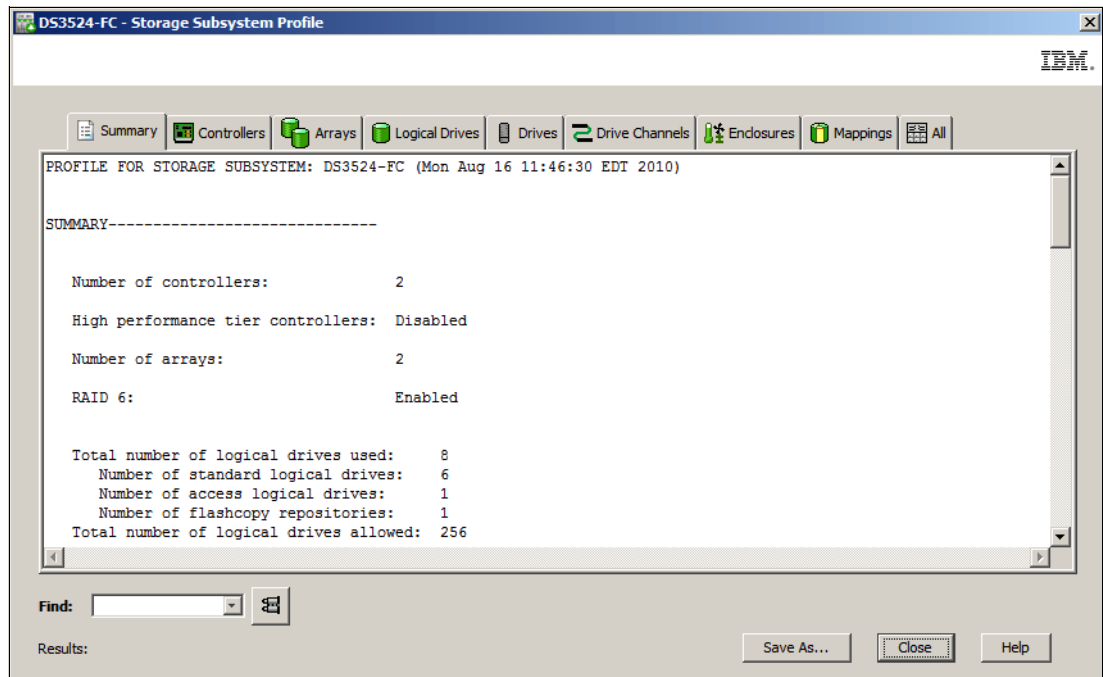


Figure 14-15 The default view when viewing the DS3500 Storage Subsystem Profile

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 as in Figure 14-16.




Figure 14-16 The tab bar of the Profile window with the All tab selected

3. Use the **Find** box (Figure 14-17) in the lower left corner of the profile window to search the profile. Enter the required text and click the Binoculars icon to find the string in the current tab.




Figure 14-17 The Find box

4. To save the profile, click the **Save As** button. At the top of Figure 14-18, choose which sections of the profile that you want to save. By default, the tab that you were viewing will be selected, and you can select any other tabs that you want. The general rule is to select the **All Sections** option.

Select a location and file name for the saved profile and click **Save**. The selected information will be stored as a text file.

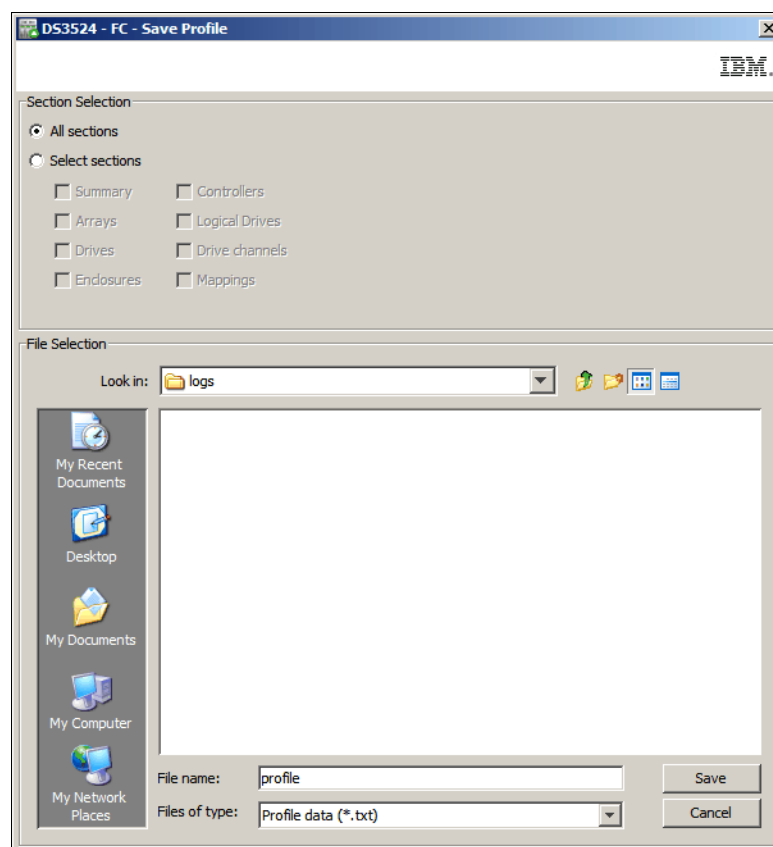


Figure 14-18 Save Profile

5. After the profile is saved, click **Close** (in Figure 14-15 on page 339) to return to the Support tab.

**Note:** Much more information is saved if the Collect All Support Data log is gathered (see 14.2, “Gather Support Information” on page 331).

## 14.4 Storage Manager Support Monitor

In this section we discuss the configuration and use of the DS Storage Manager Support Monitor Profiler software.

Additional information about the Support Monitor is available from these places:

- ▶ The Support Monitor online help which is accessed from the Support Monitor web-interface, or directly at this URL on any system running the Support Monitor software:  
[http://localhost:9000/help/Profiler\\_OLH\\_IBM\\_conditionalized\\_version.html](http://localhost:9000/help/Profiler_OLH_IBM_conditionalized_version.html)
- ▶ DS Storage Manager Version 10: Installation and Host Support Guide available at the following URL:  
<http://www-947.ibm.com/systems/support/supportsite.wss/docdisplay?lnocid=MIGR-5075652&brandind=5000028>

### 14.4.1 Support Monitor installation

The installation process for the DS Storage Manager software automatically installs the Support Monitor when the installation type is either **Typical** or **Management Station**. The **Custom** installation option allows you to select whether or not you want to install the Support Monitor. Installing the Support Monitor will also install the Apache Tomcat web server and MySQL database software packages. Support Monitor will also install a number of required services as shown in Figure 14-19.

|                                                                                        |                                        |         |           |              |
|----------------------------------------------------------------------------------------|----------------------------------------|---------|-----------|--------------|
|     | IBMStorageManagerProfilerCollector     | Started | Automatic | Local System |
|     | IBMStorageManagerProfilerEventReceiver | Started | Automatic | Local System |
|     | IBMStorageManagerProfilerMaintenance   | Started | Automatic | Local System |
|     | IBMStorageManagerProfilerPoller        | Started | Automatic | Local System |
|     | IBMStorageManagerProfilerWebServer     | Started | Automatic | Local System |
| Apache Tomcat Server - <a href="http://tomcat.apache.org">http://tomcat.apache.org</a> |                                        | Started | Automatic | Local System |

Figure 14-19 Support Monitor services

### 14.4.2 Support Monitor overview

The IBM DS Storage Manager Profiler Support Monitor tool performs the following functions:

- ▶ It will automatically collect a support bundle from each managed DS storage subsystem through the out-of-band management connection. The default is to collect the support bundle daily at 2 a.m. The support data bundle is a compressed file of the following items:
  - Collect all support data (CASD) bundle
  - Storage subsystem configuration file
  - ESM state capture
  - RLS logs and SOC logs from DS storage subsystems that have Fibre-Channel drive channels connections such as the IBM System Storage DS5020.

**Note:** The DS3500 only has SAS drive channel connections, and therefore RLS and SOC data is not collected for DS3500 systems.

- ▶ It automatically saves only the last five collected support data bundles and deletes older support data bundles.
- ▶ It provides a web-based interface for selecting the appropriate support data bundle to send to IBM Support.

**Important:** If you have multiple management stations running the Support Monitor software, problems can arise if each instance of Support Monitor requests a log collection from the same DS3500 storage subsystem at the same time (2 a.m.). There are three workarounds to prevent this from occurring:

1. Stagger the times that the various Support Monitors are scheduled to collect CASDs from the same DS storage subsystem. The default time to collect a CASD is 2 a.m. For each instance of Support Monitor, ensure that there is at least a 15 minute gap between the scheduled collection times.
2. Ensure that each DS3500 storage subsystem is managed by only one DS Storage Manager Client.
3. Change the frequency that the Support Monitor collects a CASD to “never”.

See Retain Tip H196538 at the following URL for more information:

<https://www-947.ibm.com/support/entry/myportal/docdisplay?indocid=MIGR-5083168&brandind=5000028>

### 14.4.3 The Support Monitor Profiler console

The Support Monitor Profiler console is a web-based interface for configuring and administrating the Support Monitor functions.

1. You can start the Support Monitor Profiler console in one of four ways:
  - Type `http://localhost:9000/` in the address field of a web browser. Note that you must replace the value `localhost` with the specific host name or IP address of the management station.
  - Double-click the **IBMStorageManagerProfiler Console** icon on the desktop. Note that this icon might not be present on every installation.
  - On the Classic Start menu, click **Start** → **Program** → **IBMStorageManagerProfiler Console**.
  - On the Start menu, click **Start** → **All Programs** → **IBMStorageManagerProfiler Console**.
2. The profiler console opens in a web browser as shown in Figure 14-20 on page 343.



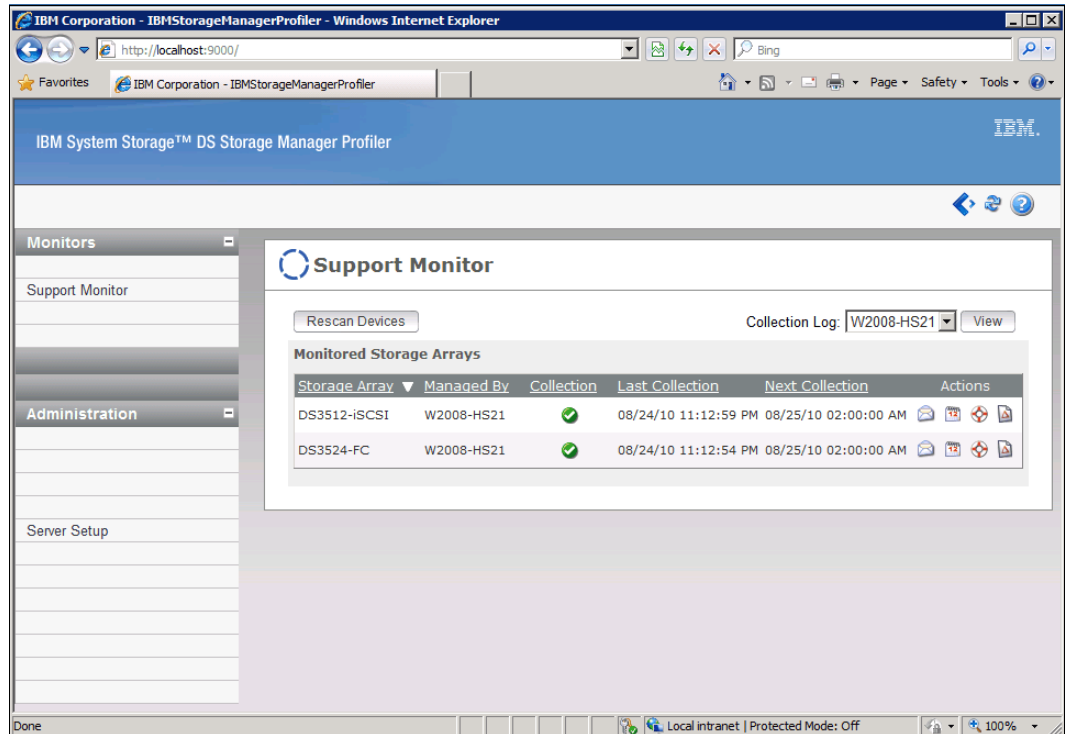


Figure 14-20 Support Monitor Profiler Console

Any DS storage systems that are being managed by the DS Storage Manager instance on the management system running Support Monitor will automatically appear in the Support Monitor list. If you add additional DS storage systems to DS Storage Manager, you might need to click the **Rescan Devices** button in the **Support Monitor** view for these additional systems to appear in the list.

- The navigation bar (Figure 14-21) on the left hand side of the console window allows you to select either the Support Monitor view or the Server Setup view.

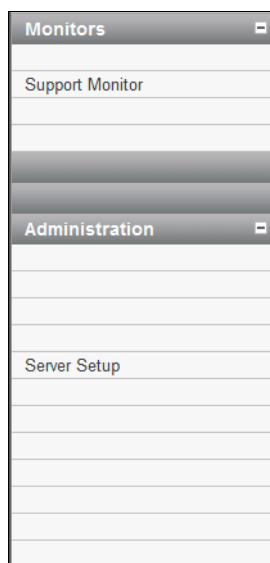


Figure 14-21 Console navigation bar

4. Click **Monitors** → **Support Monitor** in the navigation bar to display the Support Monitor window as shown in Figure 14-22.

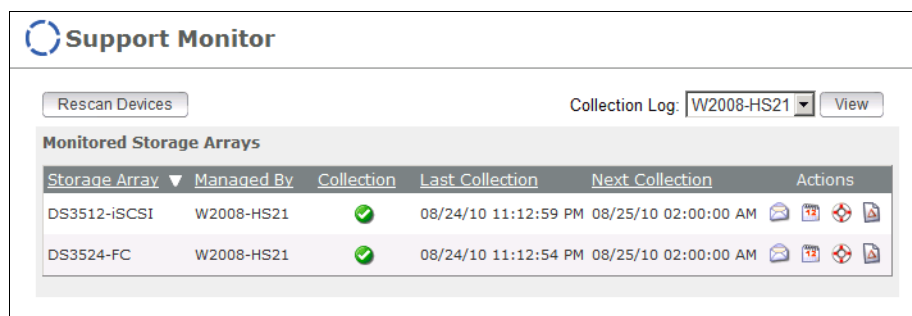


Figure 14-22 Support Monitor

5. Click **Administration** → **Server Setup** in the navigation bar to display the Server Setup window as shown in Figure 14-23.

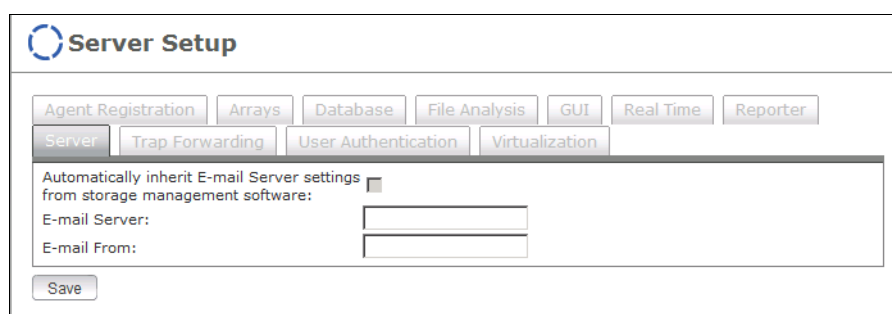


Figure 14-23 Server Setup

Here you configure the E-mail settings that the Support Monitor will use:

- a. In the **E-mail Server** text box, type the IP address or host name of the email server.
- b. In the **E-mail From** text box, type the email address from which support personnel will receive support data.
- c. Click the **Save** button.

#### 14.4.4 Support Monitor functions

In the **Support Monitor** view, each DS storage subsystem that appears in the list has four icons on the right hand side as shown in Figure 14-24. In this section we describe the functions associated with these icons.




Figure 14-24 The Support Monitor entry for a DS3524 storage subsystem

1. Click the Mail icon to open the Profiler - Send Support Data window as shown in Figure 14-25 on page 345. In this window you can specify which support bundles to send to the IBM support personnel for review.

**Send Support Data**

Storage Array: DS3524-FC

**Support Bundle Files:**

| Available                    |   | Selected |
|------------------------------|---|----------|
| DS3524-FC_d082410-231254.zip | » |          |
| DS3524-FC_p081910-020025.zip | » |          |
| DS3524-FC_p082010-020024.zip | » |          |
| DS3524-FC_p082110-020023.zip | » |          |
| DS3524-FC_p082210-020022.zip | » |          |
| DS3524-FC_p082310-020026.zip | » |          |

**SOC Change Log Files:**

| Available |   | Selected |
|-----------|---|----------|
|           | » |          |
|           | » |          |
|           | » |          |
|           | » |          |

**RLS Change Log Files:**

| Available |   | Selected |
|-----------|---|----------|
|           | » |          |
|           | » |          |
|           | » |          |
|           | » |          |

**Send To:**

☐ Send change log files to a repository address:

**CC:**  Note: Use a semi-colon to separate multiple addresses.

**Subject:** DS3524-FC Support Data

**Details:**

Note: Details are limited to 1000 characters.

Figure 14-25 Profiler - Send Support Data

Do the following to select and send collected support logs to your support personnel for review:

- Highlight one or more of the **Support Bundle Files** listed in the **Available** box and click the arrow to move them to the **Selected** box.  
The **SOC Change Log Files** and **RLS Change Log Files** are not collected by the DS3500 storage subsystem.
  - Enter the email address of the support personnel that you are sending the logs to in the **Send To** box. You can optionally include additional email recipients in the **CC** box. Each email address must be separated by a semi-colon.
  - The **Subject** box is automatically populated.
  - In the **Details** box, you can also optionally include additional information regarding the support bundles that you are submitting for review such as a description of the problem that you want investigated.
  - Click the **Send** button after you have completed all the required fields.
- Click the calendar icon to schedule the times that the support data collection will occur, as shown in Figure 14-26 on page 346.




Figure 14-26 Schedule Support Data Collection

3. Click the **Life Preserver** icon to save a data collection immediately. You might see a message alerting you that an older support bundle will be over-written, as shown in Figure 14-27. Click the **Yes** button to continue.




Figure 14-27 Collect Support Data immediately

4. Click the fourth and final icon on the extreme right hand side to open the Profiler - Compute Change Log Files window as shown in Figure 14-28 on page 347.

**Note:** SOC and RLS data are only collected by DS storage subsystem with Fibre-Channel drive channel connections (such as the DS5020) and are not collected for DS storage subsystems that have SAS drive channel connections (such as the DS3500). As such, the functions provided by this window are not relevant to the DS3500 storage subsystem.

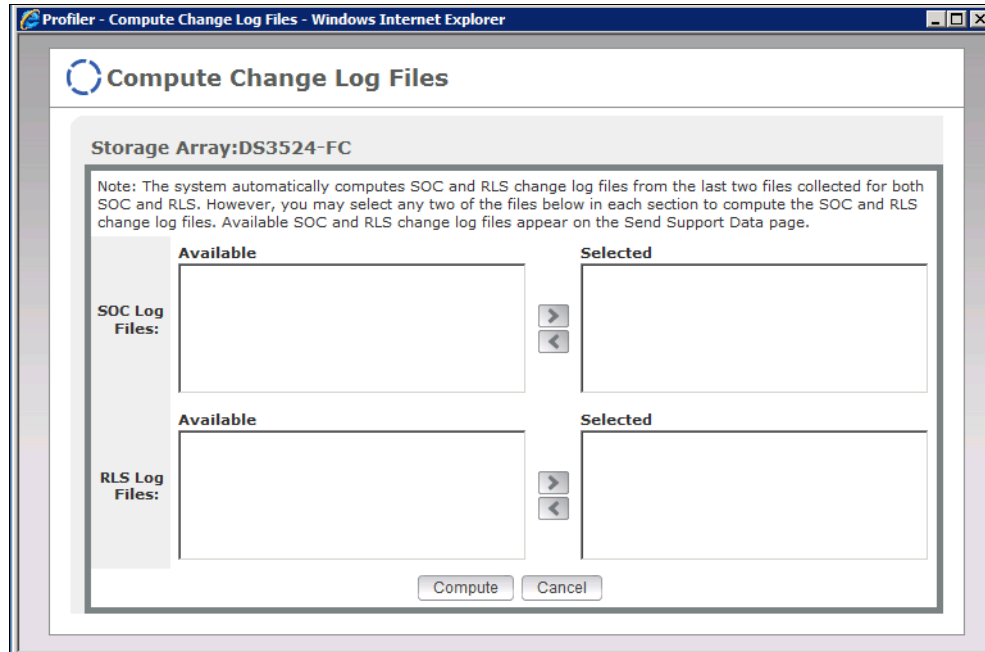


Figure 14-28 Profiler - Compute Change Log Files

### 14.4.5 Support Monitor - View Module Logs

In this section we describe how to view the running log of the DS Storage Manager Profiler Support Monitor. The module logs can be used to identify and resolve problems with the Support Monitor profiler.

1. In the Support Monitor view, click the **View** button to the left of the host name listed in the **Collection Log** box as shown in Figure 14-29.




Figure 14-29 View Collection Log

2. The **View Module Log File** window is displayed as shown in Figure 14-30 on page 348. The name `mod.sys.support.Support` is displayed in the **Existing Modules** field. In this window, you can specify the number of log entries to display (up to 1,000 lines) and the frequency of the data refresh.

**View Module Log File - W2008-HS21** (123.456.789.123 )

[Download Module Logs and Data](#)

Module: **mod.sys.support.Support** Number of lines: **100** Refresh Rate: **10 Min** **Refresh**

[Previous Page](#) [Next Page](#)

```

Aug 2010 01:40:15 [WARN]: discovery(2): discovery process completed in 0 secs
25 Aug 2010 01:44:59 [WARN]: registration - creating array monitor registration bean for Support Array Monitor
25 Aug 2010 01:44:59 [WARN]: Data ready for mod.sys.support.Support_1.0, sending notification to server
25 Aug 2010 01:45:15 [WARN]: discovery(2): discovering arrays/smtp on 300 sec intervals
25 Aug 2010 01:45:16 [WARN]: discovery(2): no new arrays discovered
25 Aug 2010 01:45:16 [WARN]: discovery(2): no unmanaged arrays detected
25 Aug 2010 01:45:16 [WARN]: discovery(2): discovery process completed in 0 secs
25 Aug 2010 01:50:16 [WARN]: discovery(2): discovering arrays/smtp on 300 sec intervals
25 Aug 2010 01:50:16 [WARN]: discovery(2): no new arrays discovered
25 Aug 2010 01:50:16 [WARN]: discovery(2): no unmanaged arrays detected
25 Aug 2010 01:50:17 [WARN]: discovery(2): discovery process completed in 0 secs
25 Aug 2010 01:55:17 [WARN]: discovery(2): discovering arrays/smtp on 300 sec intervals
25 Aug 2010 01:55:17 [WARN]: discovery(2): no new arrays discovered
25 Aug 2010 01:55:17 [WARN]: discovery(2): no unmanaged arrays detected
25 Aug 2010 01:55:18 [WARN]: discovery(2): discovery process completed in 0 secs
25 Aug 2010 02:00:01 [WARN]: DS3524-FC started periodic support data capture
25 Aug 2010 02:00:01 [WARN]: Data ready for mod.sys.support.Support_1.0, sending notification to server
25 Aug 2010 02:00:01 [WARN]: DS3512-iSCSI started periodic support data capture
25 Aug 2010 02:00:01 [WARN]: DS3524-FC checking array firmware
25 Aug 2010 02:00:01 [WARN]: DS3512-iSCSI checking array firmware
25 Aug 2010 02:00:18 [WARN]: discovery(2): discovering arrays/smtp on 300 sec intervals
25 Aug 2010 02:00:18 [WARN]: discovery(2): no new arrays discovered
25 Aug 2010 02:00:18 [WARN]: discovery(2): no unmanaged arrays detected
25 Aug 2010 02:00:18 [WARN]: discovery(2): discovery process completed in 0 secs
25 Aug 2010 02:00:27 [WARN]: DS3512-iSCSI invalid firmware (0.0), not capturing support data. Firmware must be 6.19.x.x, 6.23.x.x, 6.
25 Aug 2010 02:00:27 [ERROR]: DS3512-iSCSI support data capture failed: command failed in getting firmware version
25 Aug 2010 02:00:53 [WARN]: DS3524-FC invalid firmware (0.0), not capturing support data. Firmware must be 6.19.x.x, 6.23.x.x, 6.60.
25 Aug 2010 02:00:53 [ERROR]: DS3524-FC support data capture failed: command failed in getting firmware version
25 Aug 2010 02:00:53 [WARN]: Data ready for mod.sys.support.Support_1.0, sending notification to server

```

Figure 14-30 View Module Log File

- Click the **Download Module Logs and Data** link to open the window shown in Figure 14-31. Here you can specify the module logs and additional Configuration File and Data Logs that you want to save to the local system by selecting the appropriate check boxes.

**Download Module Logs and Data - W2008-HS21** (123.456.789.123 )

**Modules:**

**Available:**  
mod.sys.support.Support

**Selected:**

Get Selected Files (selected) Get All Files

**Files:**

☒ Log Files ☐ Configuration File ☐ Data Files

**Download** **Cancel**

Figure 14-31 Download Module Logs and Data

## 14.5 Download firmware

In this section we describe the steps to upgrade the DS3500 storage subsystem firmware and NVSRAM microcode. You can download the latest firmware at:

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

Appendix A, “IBM Support Portal website” on page 719 describes this website and its contents.

On the same website you can select (for example, as shown in Figure 14-32, for the DS3500) the firmware (controller firmware and NVSRAM, ESM and disk firmware), HBA device drivers and firmware, RDAC multipath drivers, and the DS Storage Manager application. Also, the links for the product documentation, troubleshooting tips, and all other information regarding DS3500 are available on this site.

IBM Systems > Systems support > System Storage and TotalStorage products > Disk systems >

### Software and device drivers

DS3500 (DS3512, DS3524)

**Select your product**

Fields marked with an asterisk (\*) are required.

**Product family: \***

Disk systems

**Product:**

DS3500 (DS3512, DS3524)

**Operating system:**

All listed operating systems

→ All products

Go

Refine results: All categories

Go

**Support & downloads**

Download

Troubleshoot

Search

Documentation

Forums & Communities

Plan & upgrades

Install

Use

Open service request

Figure 14-32 Code downloads for DS3500

The DS3500 firmware version is closely connected to the Storage Manager version, and the HBA device driver and firmware available for your hosts. You should always check the DS3500 firmware readme and change-history files to see if there are any host software dependencies. If there is a newer version of Storage Manager (including multi-path device drivers) or HBA software required for the hosts, download and install these first before updating the DS3500 firmware. ESM firmware should also be upgraded to the required level before starting the controller firmware upgrade. In general, run at the latest levels unless specifically advised otherwise.

**Note:** Updating firmware will normally require a maintenance window, so it is better to plan an appropriate outage time in advance.

Download all the packages you will need from the IBM Support website 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.

## 14.5.1 Before you upgrade

There are precautions you should take before you update a storage subsystem:

- ▶ Make sure that the overall system status is Optimal by checking the Summary tab of the Storage Manager Subsystem Management Window as shown in Figure 14-33. If not, run the Recovery Guru as described in 14.13, “Recovery Guru - Recover from Failure” on page 409 to diagnose and fix the problem before proceeding with the upgrade.
- ▶ Always read the appropriate readme file before upgrading any firmware to check for compatibility with your system.
- ▶ Take care with dual controller configurations. The upgrade procedure requires a management connection to both DS3500 controllers, and it is not possible to perform a firmware upgrade with only one controller connected. Therefore, both controllers must be online and accessible, either through an in-band or out-of-band management connection. If you upgrade redundant configurations in-band, make sure that a multi-path driver is installed on the management host. This is necessary because the access logical drive moves from one controller to the other during this procedure and the DS3500 must be manageable during the entire time.

**Note:** The overall DS3500 status can be checked in the Summary tab of the Subsystem Management Window as shown in Figure 14-33.




Figure 14-33 Viewing the DS3500 Storage Subsystem status in DS Storage Manager

**General rule:** Save the All Support Data collection (14.2, “Gather Support Information” on page 331) before starting any storage subsystem upgrades.



### Notes:

1. Update the hosts, and ensure that the correct version of multipath driver is installed.
2. Update the ESMs first to the latest level, followed by the controller firmware. Outdated ESM firmware could make your expansion enclosures inaccessible after the DS3500 storage server firmware update. The required ESM firmware level for each Controller firmware level is mentioned in the readme for controller firmware.
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 might lead to configuration corruption. Therefore, do not power off the DS3500 storage system or the management station during the update.

## Firmware Inventory

Before and after updating the firmware levels of the DS3500 storage subsystem, you can use the **Firmware Inventory** to view the firmware levels for your storage subsystems.

To view the firmware inventory for one storage subsystem:

1. Click **Advanced** → **Maintenance** → **Firmware Inventory** in the Storage Subsystem window. Figure 14-34 shows the Firmware Inventory window that displays.
2. Save the firmware inventory to a text file by clicking the **Save As** button.

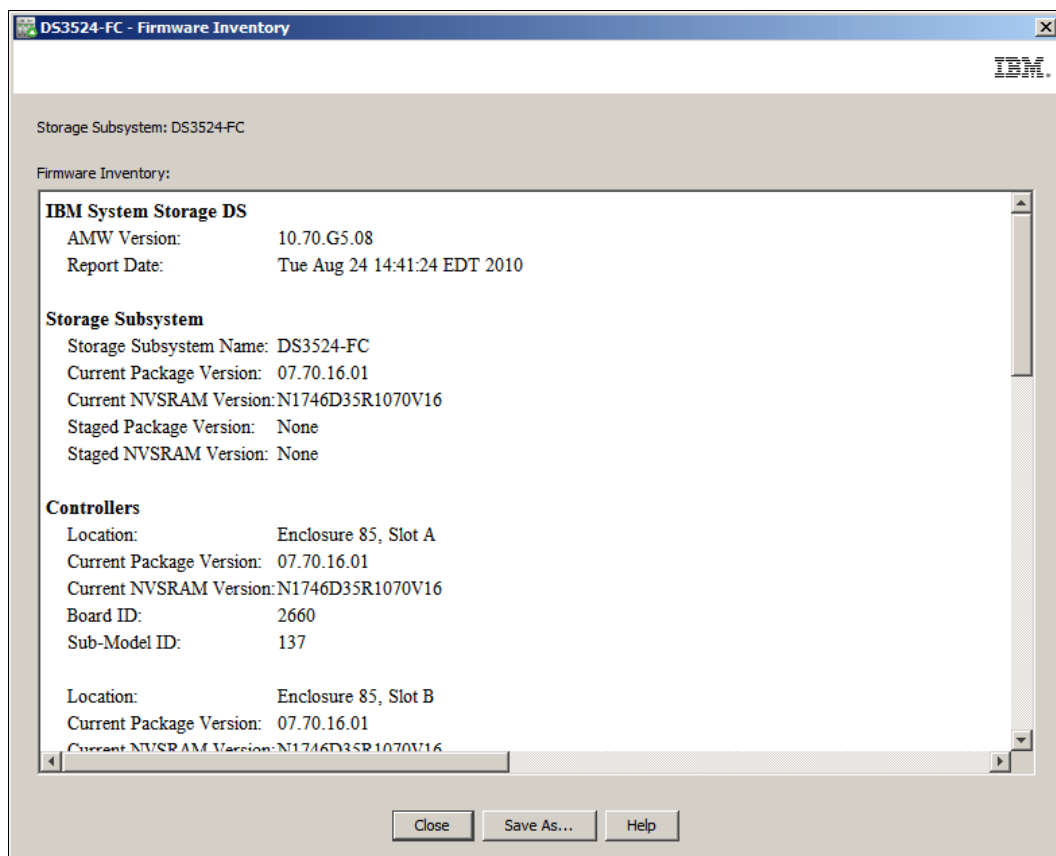


Figure 14-34 Storage Subsystem Firmware Inventory

To view the firmware inventory for more than one storage subsystem:

1. Click the **Tools** → **Firmware Inventory** menu command from the Enterprise Management window. Figure 14-35 shows the Firmware Inventory window that displays.
2. Save the firmware inventory to a text file by clicking the **Save As** button.

**Note:** If you want to save the Firmware Inventory for more than one storage subsystem, you will need to select each of the storage subsystems in the drop-down box and then click **Save As**.

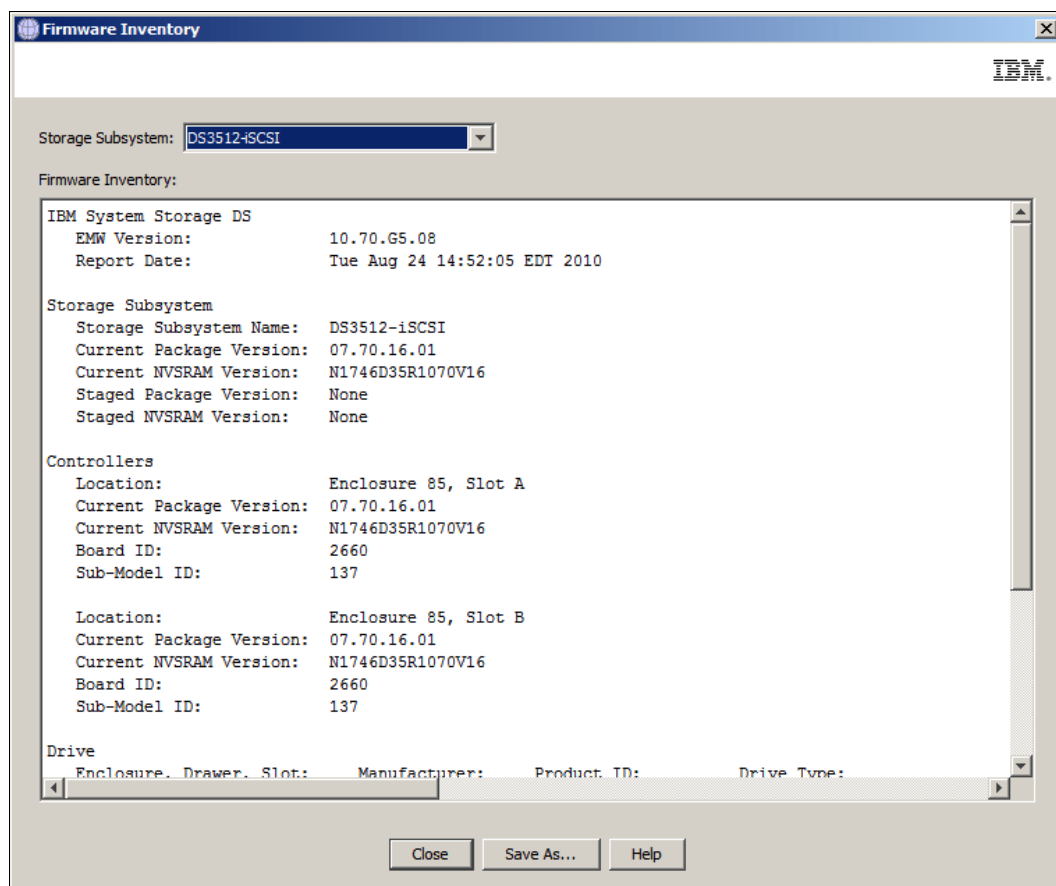


Figure 14-35 Enterprise Management firmware inventory

## 14.5.2 Updating the host

Before you update the firmware on your DS3500, the software on your hosts should also be up to date. Go to the IBM Support website listed at the following URL to check the levels of Storage Manager and HBA drivers:

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

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 5, “Installing IBM System Storage DS Storage Manager” on page 131. Storage Manager for Windows also includes the associated multi path driver. The multipath

RDAC driver is a separate package for Linux. Other host operating systems provide their own multipath drivers. The required multipath driver 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 might need to plan downtime.

Also check the HBA firmware to make sure it is at the latest supported level. You can also find this information at the IBM Support website.

### 14.5.3 Upgrading the DS3500 controller firmware

In this section we describe how you can upgrade the firmware of the controllers in the storage subsystem by using the DS Storage Manager software.

#### Staged upgrades and activation

The process of upgrading the firmware is a two-stage process:

1. The firmware file is transferred from the management host to the DS3500 storage system.
2. The DS3500 is upgraded with the new firmware file. The process of upgrading the firmware after downloading the firmware file is known as activation. During activation, the existing firmware file in the memory of the controller is replaced with the new firmware file.

After transferring the firmware file, you can choose to automatically activate the controllers in the storage subsystem to the new firmware immediately. Optionally, you can transfer the firmware file to the controller and activate the firmware later at a more convenient time. This is known as a *staged* upgrade.

#### Pre-upgrade checks

The storage management software checks for existing conditions in the storage subsystem that will prevent a successful firmware upgrade before upgrading the firmware. If it detects one of these conditions, the firmware upgrade tool reports the status of the DS3500 storage system as *non-upgradeable* (see Figure 14-37 on page 356). Any of the following conditions can prevent the firmware upgrade:

- ▶ An unsupported controller type or controllers of different types that are in the storage subsystem that cannot be upgraded
- ▶ One or more failed drives
- ▶ One or more hot spare drives that are in use
- ▶ One or more arrays that are incomplete
- ▶ Operations, such as defragmenting an array, downloading of drive firmware, and others, that are in progress
- ▶ Missing logical drives that are in the storage subsystem
- ▶ Controllers that have a status other than Optimal
- ▶ The storage partitioning database is corrupt
- ▶ A data validation error occurred in the storage subsystem
- ▶ The storage subsystem has a Needs Attention status
- ▶ The storage subsystem is unresponsive, and the storage management software cannot communicate with the storage subsystem

- The Event Log entries are not cleared

You can correct some of these conditions by using the functions accessible through the Subsystem Management Window. However, for some of these conditions, you might need to contact your IBM Customer Support representative. The storage management software saves the information about the firmware upgrade process in log files. This action helps the IBM Customer Support representative to understand the conditions that prevented the firmware upgrade.

**Note:** The location where the firmware upgrade tool log files are saved on your system is reported at the top of the firmware upgrade tool View Log window. See Figure 14-38 on page 356.

## Major and Minor firmware upgrades

A version number exists for each firmware file. For example, at the time of writing 07.70.16.01 is the current version number for the DS3500 firmware file. The first two digits indicate the major revision of the firmware file. The remaining digits indicate the minor revision of the firmware file. You can view the version number of a firmware file in the Upgrade Controller Firmware and the Download Firmware windows. For more information, refer to the Downloading the Firmware online help topic in the Enterprise Management window.

The process of upgrading the firmware can be either a major upgrade or a minor upgrade depending on the version of the firmware. For example, the process of upgrading the firmware is major if the version of the current firmware is 07.70.16.01, and you want to upgrade the firmware to a theoretical version 08.ab.cd.ef. In this example, the first two digits of the version numbers are different and indicate a major upgrade. In a minor upgrade, the first two digits of the version numbers are the same. For example, the process of upgrading the firmware is minor if the version of the current firmware is 07.70.16.01, and you want to upgrade the firmware to version 07.ab.cd.ef or any other minor revision of the firmware.

## DS Storage Manager upgrade software

There are two separate tools included in the DS Storage Manager software to update the controller firmware and NVSRAM:

- The firmware upgrade tool accessed from the DS Storage Manager (Enterprise Management) window. This tool is able to perform staged firmware upgrades in which the firmware is transferred to the DS3500 but the upgrade is not applied. This tool is able to perform both minor firmware upgrades and major firmware upgrades. This tool is described in 14.5.4, “Using the Enterprise Management upgrade tool” on page 355.
- The firmware upgrade tool accessed from the DS Storage Manager (Subsystem Management) window. In addition to the controller firmware and NVSRAM, this tool is also used to update the ESM and drive firmware. This tool is not able to perform either staged firmware upgrades or major firmware upgrades. This tool is described in 14.5.5, “Using the DS3500 Storage Manager (Subsystem Management)” on page 368.

**Note:** You also can use the command line interface (CLI) to download and activate firmware to the DS3500 storage subsystems. For more information, refer to the About the Command Line Interface online help topic in the Enterprise Management window.

Table 14-1 on page 355 lists the major differences in functionality between the DS Storage Manager firmware update tools.

Table 14-1 Comparison of the firmware updater functions

| Update Type                | Enterprise Management | Subsystem Management |
|----------------------------|-----------------------|----------------------|
| Minor Controller Firmware  | Yes                   | Yes                  |
| Major Controller Firmware  | Yes                   | No                   |
| Staged Controller Firmware | Yes                   | No                   |
| NVSRAM Firmware            | Yes                   | Yes                  |
| ESM Firmware               | No                    | Yes                  |
| ESM Configuration          | No                    | Yes                  |
| Drive Firmware             | No                    | Yes                  |

## 14.5.4 Using the Enterprise Management upgrade tool

In this section we describe the procedure to update the DS3500 controller firmware using the firmware upgrade tool accessible from the DS Storage Manager Enterprise Management window.

**Note:** The upgrade tool allows you to perform both minor and major controller firmware upgrades.

1. From the DS Storage Manager Enterprise Management window, right-click a selected DS3500 storage subsystem and select the **Upgrade Controller Firmware** option on the context menu as shown in Figure 14-36.

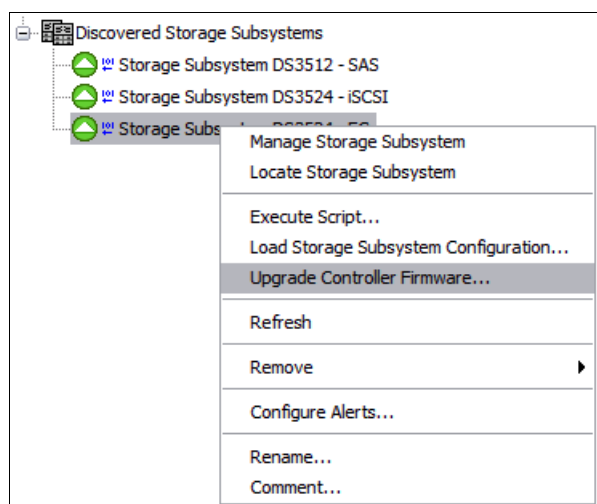


Figure 14-36 Upgrade Controller Firmware option in the Enterprise Management window

**Note:** This function is also accessible from the **Tools** → **Upgrade Controller Firmware...** menu in the DS Storage Manager (Enterprise Management) window.

2. The **Upgrade Controller Firmware** window appears (Figure 14-37 on page 356). In the example displayed in Figure 14-37 on page 356, the upgrade tool reports that the

controller firmware is not presently upgradeable. The reason or reasons for this status is reported in the **Details** section of the upgrade tool window.

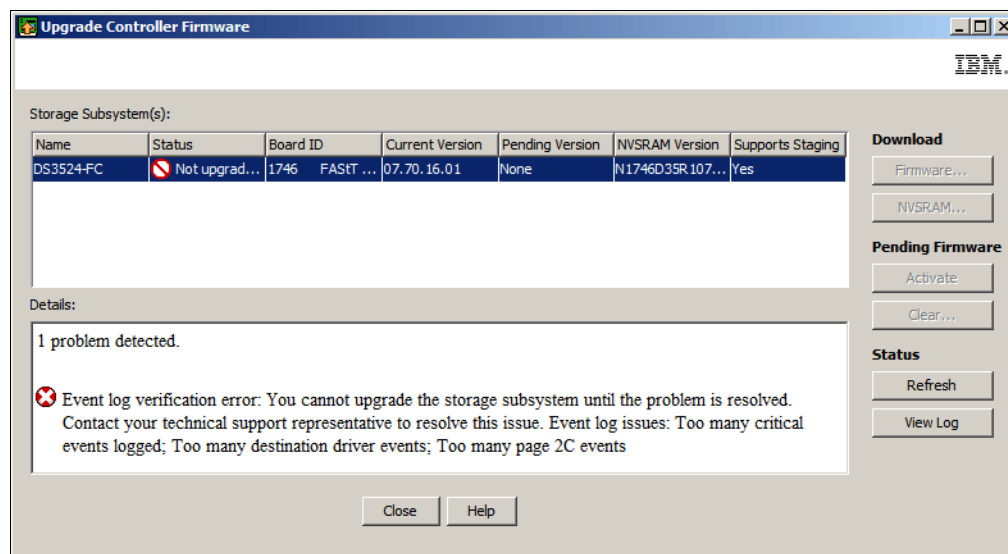


Figure 14-37 Upgrade Controller Firmware

If the firmware upgrade tool reports a status of *non-upgradeable* as shown in Figure 14-37 you can click the **View Log** button to review the cause of the non-upgradeable status. Figure 14-38 shows the upgrade log.

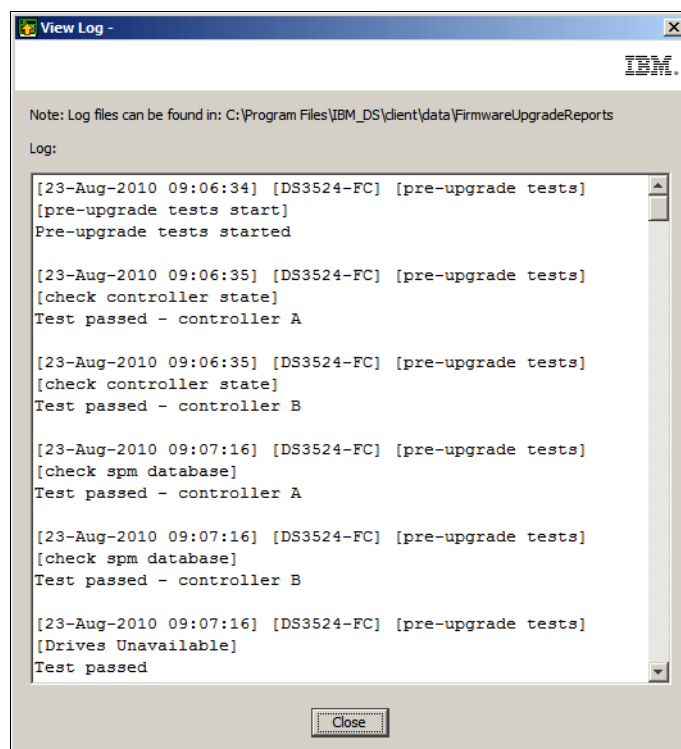


Figure 14-38 Viewing the firmware upgrade tool log

In this example, the problem that was preventing a firmware upgrade was fixed by clearing the DS3500 storage subsystem event log. See 14.6, “View Event Log” on page 385 for

instructions on how to clear the subsystem event log. However, certain types of problems require that you contact your local IBM Support Representative.

3. If the firmware upgrade tool does not detect any problems that would prevent the upgrade, the tool reports a status of *Upgradeable*, as shown in Figure 14-39.




Figure 14-39 Upgrade Controller Firmware

4. Click the **Firmware** button to open the **Download Firmware** window as shown in Figure 14-40 on page 358.

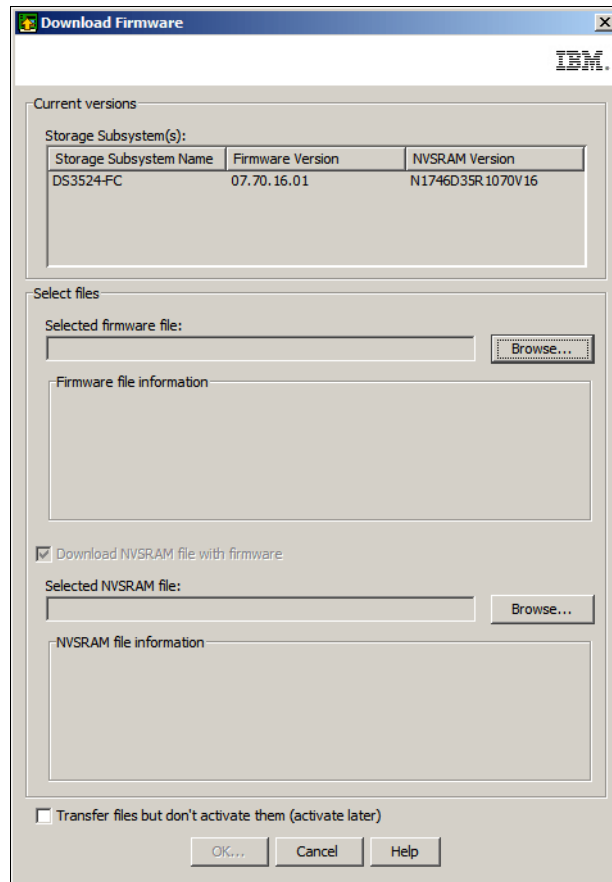


Figure 14-40 Download Firmware window

There are two main sections to this window:

- **Current versions:** This section reports the versions of controller firmware and NVSRAM that are currently installed.
- **Select files:** This section is where you select the firmware and NVSRAM files that you want to transfer to the DS3500 storage system. Information about the firmware and NVSRAM file packages that you select is also displayed in this section.

**Note:** The Download NVSRAM file with firmware check box is enabled by default. Always transfer the firmware and NVSRAM at the same time.

- a. In the **Selected firmware file** pane, click the **Browse** button and select the firmware upgrade file as shown in Figure 14-41 on page 359.



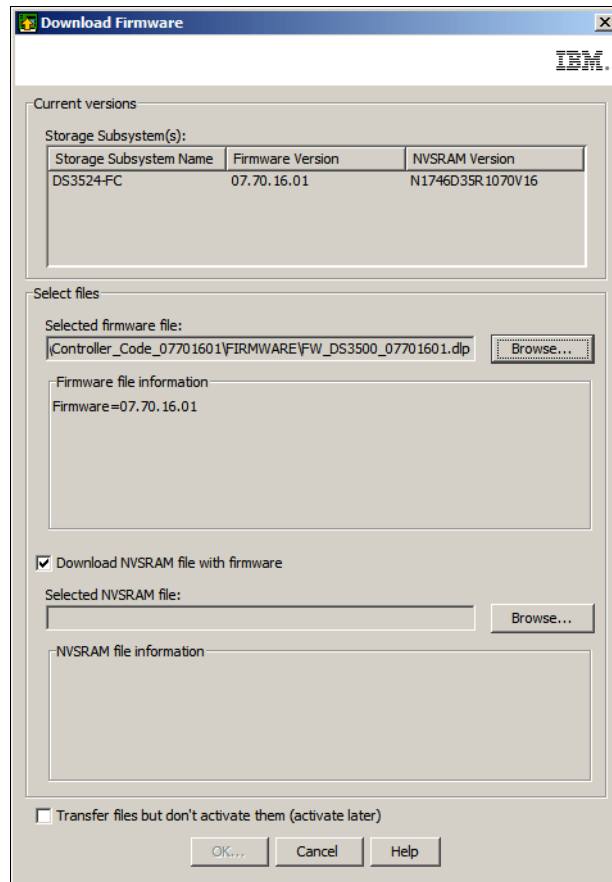


Figure 14-41 Select firmware file

- b. In the **Selected NVSRAM file** pane, click the **Browse** button and select the NVSRAM file that you want to transfer (Figure 14-42 on page 360).

**Attention:** Ensure that you choose the version of the NVSRAM file that is appropriate for your DS3500 storage system. For example, if your DS3500 storage system has two controllers, you must select the dual controller NVSRAM file. If your DS3500 storage system has one controller, you must select the single controller NVSRAM file.

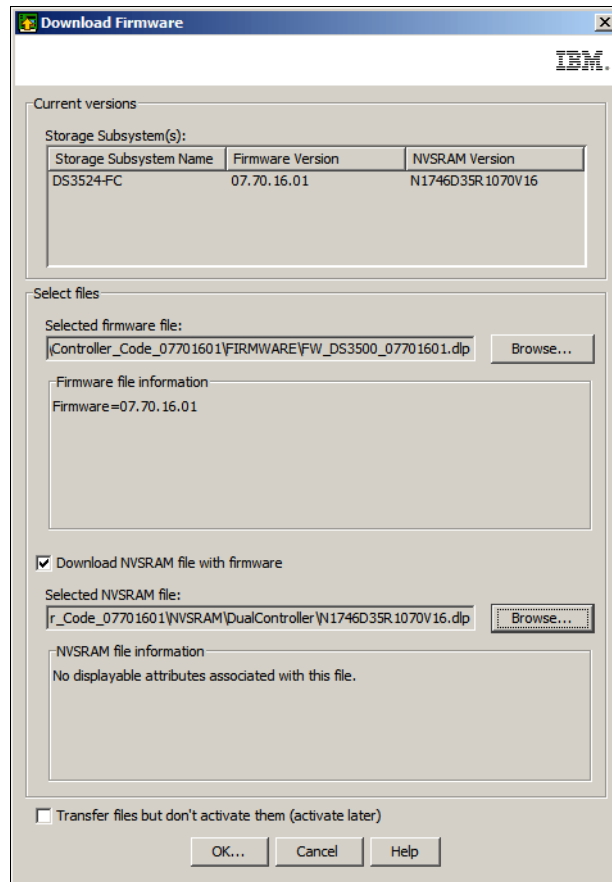


Figure 14-42 Select the NVSRAM file

5. If you want to transfer the firmware and NVSRAM and activate these immediately, then click the **OK** button. You will be asked to confirm that you want to transfer and activate the firmware and NVSRAM immediately. See Figure 14-43 on page 361.

**Note:** The process to transfer the firmware and NVSRAM images to the DS3500, but activate them later is described in step 6.

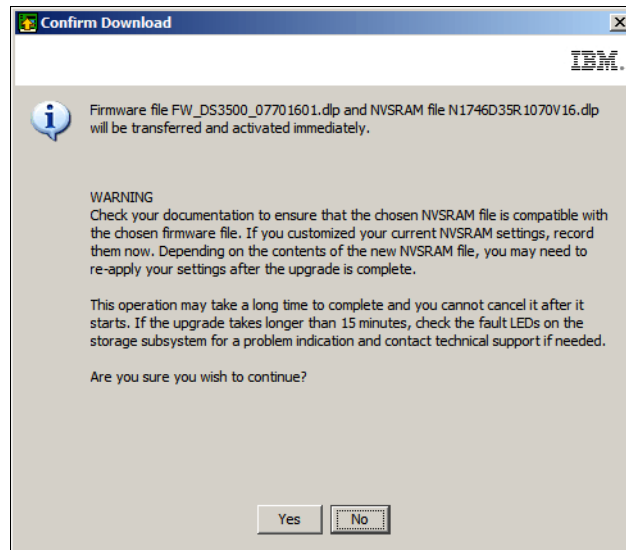


Figure 14-43 Transfer and activate

- Click the **Yes** button to start the transfer and activation.
- The firmware upgrade tool will begin transferring the firmware images to the controllers. The status can be monitored in the tool (Figure 14-44).




Figure 14-44 Firmware upgrade in progress

- Figure 14-45 on page 362 shows the status after the firmware has been transferred and activated.




Figure 14-45 Successful firmware transfer and activation

6. As described above, the default procedure is to transfer the firmware and NVSRAM and activate the new versions immediately. However, you can decide to only transfer the firmware and NVSRAM at this time, and then activate the firmware and NVSRAM at a later time using these steps:
  - a. To transfer the firmware and NVSRAM, select the check box option at the bottom of the **Download Firmware** window. See Figure 14-46.

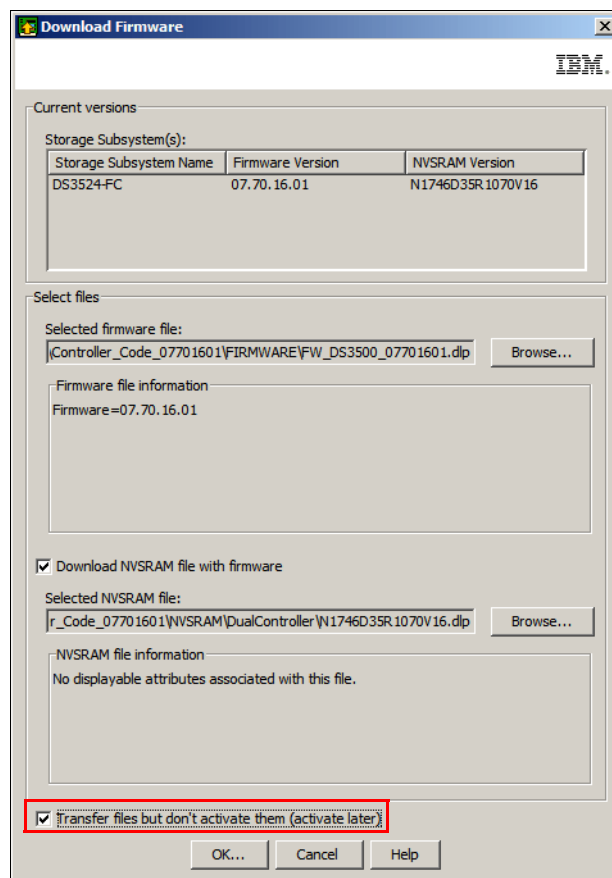


Figure 14-46 Transfer firmware and NVSRAM only (activate later)

- b. Click the **OK** button. You will be asked to confirm that you want to transfer the firmware and NVSRAM now but activate the new versions later (Figure 14-47).

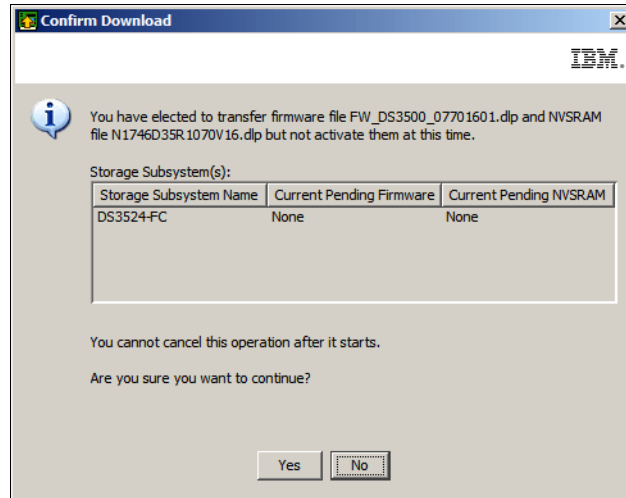


Figure 14-47 Confirm Download

- c. The status of the transfer of the firmware and NVSRAM can be monitored in the tool as shown in Figure 14-48.




Figure 14-48 Transferring the firmware and NVSRAM

- d. Figure 14-49 on page 364 shows the status after the firmware and NVSRAM have been transferred. The status is reported as *Firmware pending* and the **Activate** button can now be clicked.




Figure 14-49 Firmware pending

**Note:** A pending firmware upgrade can also be activated in the DS Storage Manager (Subsystem Management) window by clicking **Advanced** → **Maintenance** → **Activate Controller Firmware**.

- e. Click the **Activate** button to activate the pending firmware. The window shown in Figure 14-50 will appear.

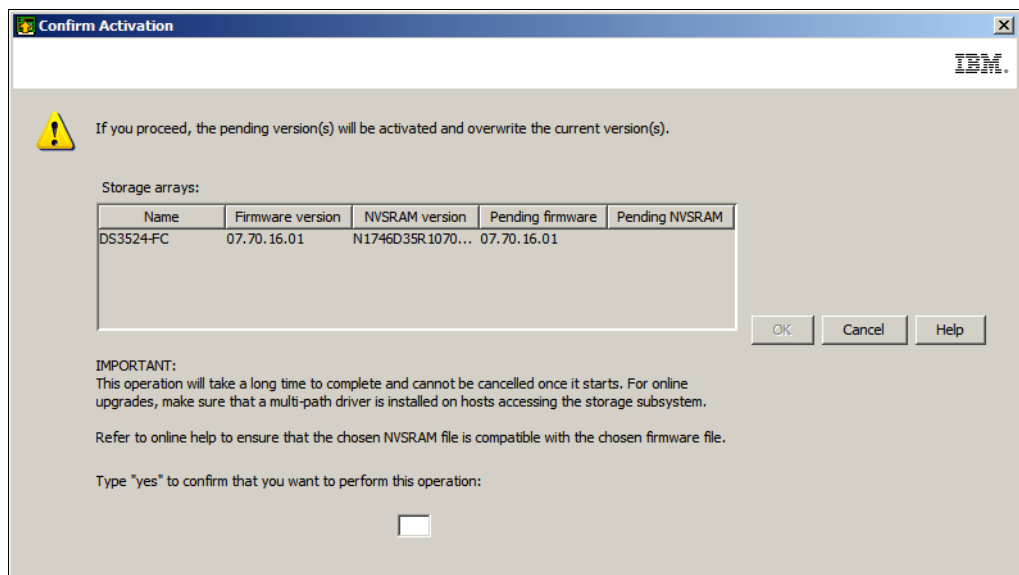


Figure 14-50 Confirm activation

- f. Type yes in the box to confirm that you want to activate the pending firmware and click **OK**.
- g. The activation process can take some time to complete, but the status of the process can be monitored in the tool. Figure 14-51 on page 365 shows the status after the firmware has been successfully activated.




Figure 14-51 Firmware activated successfully

7. It is also possible to clear a pending firmware and NVSRAM upgrade image that has been transferred to the DS3500 storage system, but not yet activated.

**Note:** You can also clear a pending firmware upgrade in the DS Storage Manager (Subsystem Management) window by clicking **Advanced** → **Maintenance** → **Clear Controller Firmware**.

- a. Click the **Clear** button in the **Upgrade Controller Firmware** window shown in Figure 14-52.

**Note:** You might need to click the **Refresh** button before the **Clear** button is enabled.




Figure 14-52 Pending firmware upgrade

- b. You will be asked to confirm that you want to clear the pending firmware upgrade as shown in Figure 14-53 on page 366.




Figure 14-53 Clear pending firmware upgrade

- c. Click the **Yes** button to clear the pending firmware upgrade and the **Clear Pending Firmware Progress** window shown in Figure 14-54 will appear.




Figure 14-54 Clear Pending Firmware Progress

- d. The **Upgrade Controller Firmware** window reports the status shown in Figure 14-55 after the firmware has been cleared.




Figure 14-55 Status after pending firmware has been cleared



8. You can also use the upgrade tool to install only an updated NVSRAM file:
- Click the **NVSRAM** button in Figure 14-55 on page 366. The **Download NVSRAM** window appears as shown in Figure 14-56.

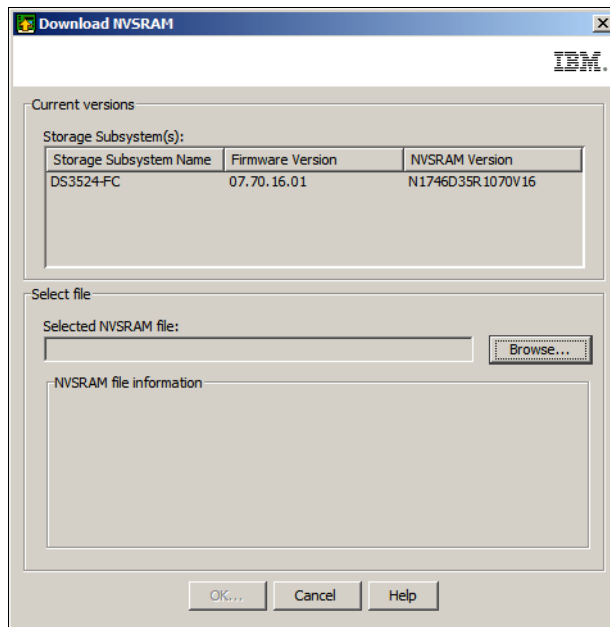


Figure 14-56 Download NVSRAM

- Click the **Browse** button to select the appropriate NVSRAM file (either single controller or dual controller NVSRAM) for your DS3500 storage system (Figure 14-57).

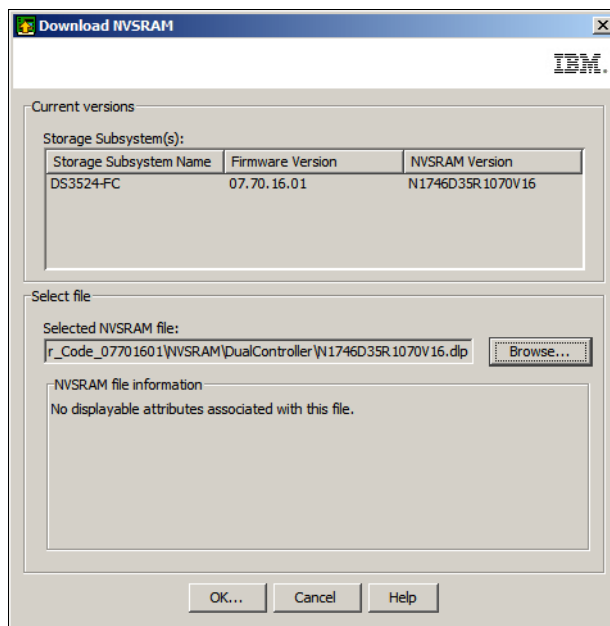


Figure 14-57 Download NVSRAM with NVSRAM file selected

- Click the **OK** button to continue and the **Confirm Download** window shown in Figure 14-58 on page 368 will appear.

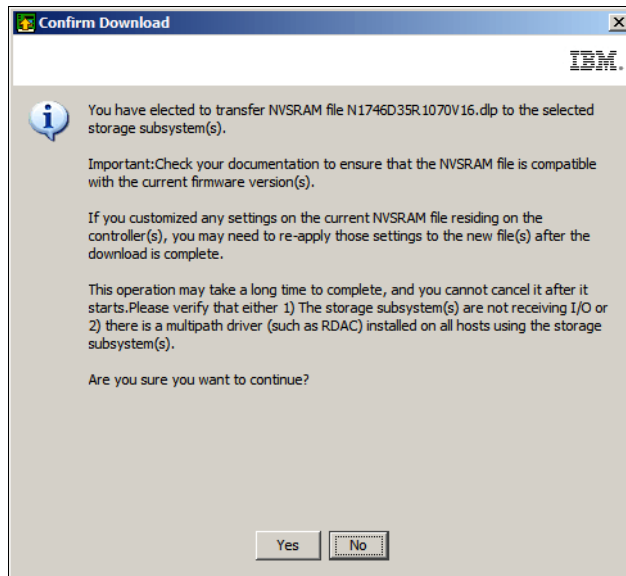


Figure 14-58 Confirm Download

- d. You can monitor the status of the NVSRAM upgrade in the **Upgrade Controller Firmware** window. Figure 14-59 shows the status after the NVSRAM has been transferred and activated successfully.




Figure 14-59 NVSRAM upgrade successful

### 14.5.5 Using the DS3500 Storage Manager (Subsystem Management)

Here now describe upgrade procedure for the DS3500 firmware and NVSRAM using the DS Storage Manager (Subsystem Management) application.

#### Updating the storage subsystem

After you have updated your hosts (see 14.5.2, “Updating the host” on page 352), you can update the code of your storage subsystem. To do this, click **Download Firmware** as shown in Figure 14-1 on page 330. You will see the options for different types of firmware to download as shown in Figure 14-60 on page 369.

**Tip:** The current DS3500 firmware levels are listed in the Storage Subsystem Profile. Clicking each of the individual Download Firmware links in Figure 14-60 also displays the current firmware levels of the respective components.

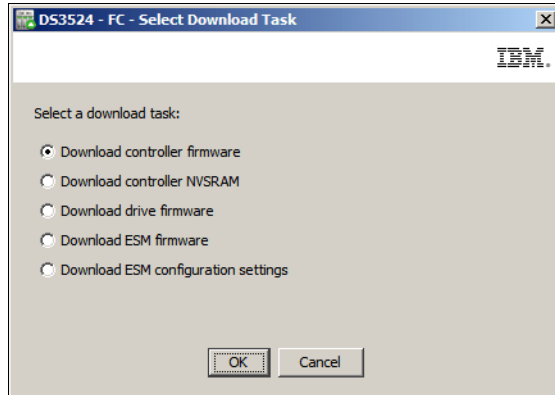


Figure 14-60 Select Download Task Firmware window

**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.

The five download options are:

- ▶ Download controller firmware: The controller firmware of the DS3500 storage subsystem.
- ▶ Download controller NVSRAM: The controller 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. Although the controller FW is the same for both single and dual controller models, there is different version of NVSRAM, so make sure to install the correct single or dual controller version of NVSRAM. Both Controller Firmware and NVSRAM can be downloaded at the same time.
- ▶ Download drive firmware: The firmware for the hard disk drives. Host IO must be stopped during this upgrade.
- ▶ Download ESM firmware: The firmware of the ESM modules in the EXP3500 expansion enclosures. This update should be installed before the controller firmware upgrade in certain cases.
- ▶ Download ESM configuration settings: A file that contains configuration settings for the drive enclosure ESMs.

These code download options are also accessible by clicking **Advanced** → **Maintenance** → **Download** as shown in Figure 14-61.




Figure 14-61 Download code options accessed from the Advanced menu

## Download ESM Card Firmware

Follow the steps described in this section to install the ESM firmware for the EXP3500 enclosure(s).

**Note:** The **Download ESM firmware** function is also accessible by clicking **Advanced** → **Maintenance** → **Download** → **ESM Firmware**.

1. To update the ESM code, select the **Download ESM firmware** option (Figure 14-60 on page 369) and then click **OK**. Figure 14-62 shows the resulting window.

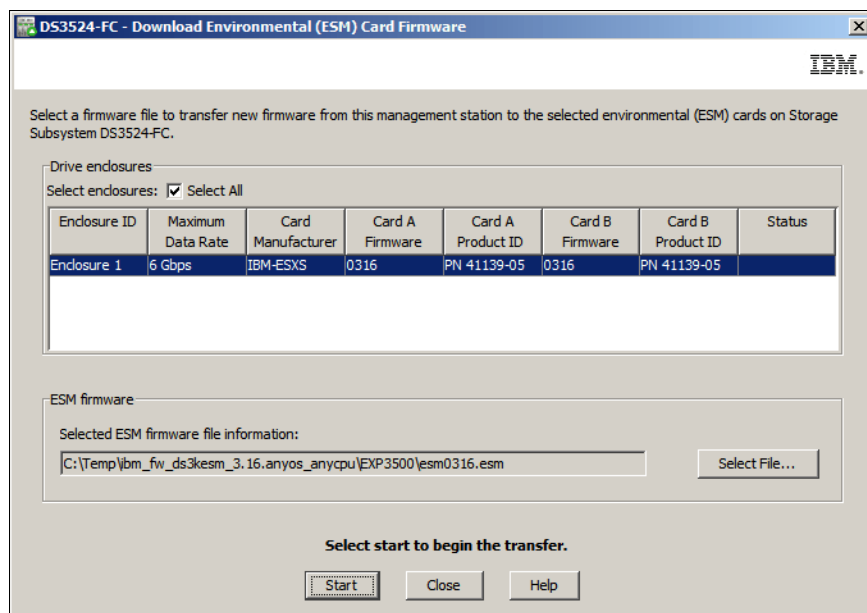


Figure 14-62 Download Environmental (ESM) Card Firmware

2. There are two sections in this view. The top section, *Drive enclosures*, lists all the installed enclosures. You can either select which enclosure(s) to update or select 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 website. Click **Select File** and select the firmware file.

Click **Start** to update the ESM firmware.

3. The Confirm Download window appears. Read the information, confirm it by typing yes, and click **OK** to start the update (Figure 14-63 on page 371).

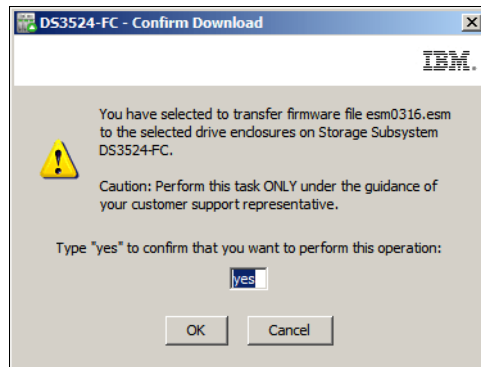


Figure 14-63 Confirm Download

4. The status field in the row of selected enclosures changes from pending to transferring while the ESM card firmware upgrade progresses. Monitor the progress and completion status of the download as shown in Figure 14-64.




Figure 14-64 transferring the ESM firmware

5. The status will change to completed when the ESM firmware has been downloaded as shown in Figure 14-65 on page 372.




Figure 14-65 ESM firmware transfer completed

## Download ESM configuration settings

Follow the steps described in this section to install new ESM configuration settings from a file.

**Note:** The **Download ESM configuration settings** function is also accessible by clicking **Advanced** → **Maintenance** → **Download** → **ESM Configuration Settings**.

1. To download ESM configuration settings, select the Download ESM Configuration Settings option as shown in Figure 14-60 on page 369. The window shown in Figure 14-66 will appear.

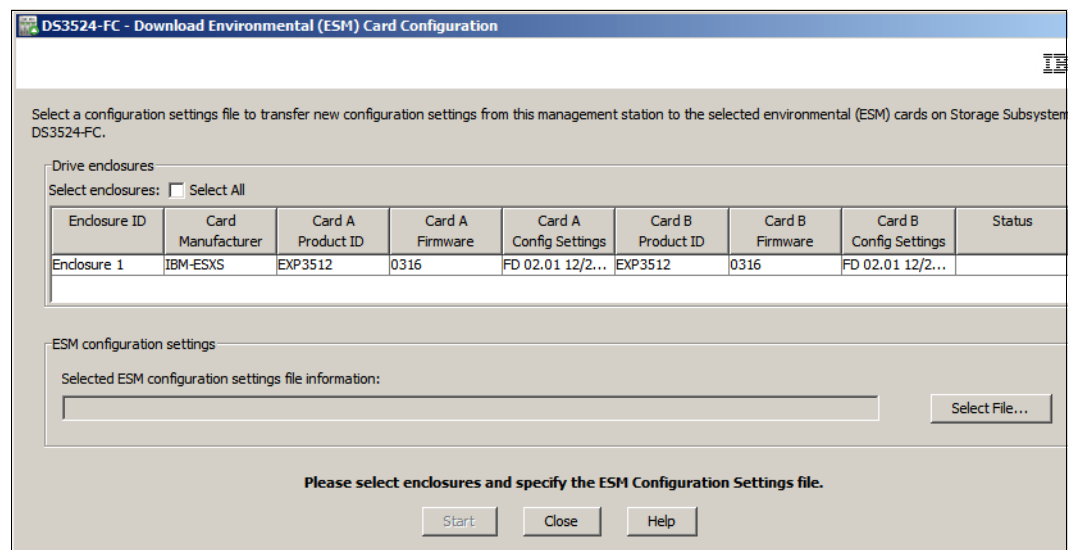


Figure 14-66 Download ESM Card Configuration

2. If there is more than one drive enclosure connected to the DS3500, select the **Select All** check box to install the ESM configuration file to all of the enclosures. Alternatively, if you

want to install the ESM configuration file to just one enclosure, then select just that enclosure in the list of enclosures.

3. Click the **Select File** button and specify the ESM configuration file that you want to transfer.
4. Click **Start**. You will be asked to confirm that you want to proceed by typing yes in the text box and clicking **OK** to continue.

## Download Controller Firmware

Follow the steps described in this section to install the DS3500 storage subsystem firmware.

**Note:** The **Download Controller Firmware** function is also accessible by clicking **Advanced** → **Maintenance** → **Download** → **Controller Firmware**.

**Attention:** Each firmware file has a version number. The version number indicates whether the firmware is a major version or a minor version. The procedure described in this section describes how to download and activate only the minor firmware versions. However, you can use the Enterprise Management window to download and activate both major and minor firmware versions. The procedure using the Enterprise Management window to download and activate the controller firmware is described in 14.5.4, “Using the Enterprise Management upgrade tool” on page 355.

1. To update the controller firmware using the DS Storage Manager, select the **Download controller firmware** option as shown in Figure 14-60 on page 369, and click **OK**.
2. Sometimes after selecting the **Download controller firmware** option, a window similar to the one shown in Figure 14-67 will appear. As in the example shown below, this can occur if the DS3500 has logged too many critical events.

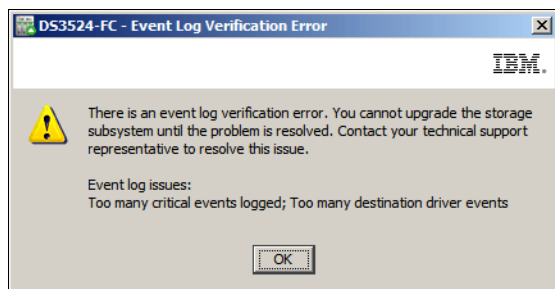


Figure 14-67 Event Log Verification Error

You will need to view the DS3500 event log and resolve the errors (see 14.6, “View Event Log” on page 385). In this example, clearing the DS3500 storage system event log fixed the condition causing the verification error. If you are unable to fix the condition that prevents the firmware upgrade, contact your local IBM Technical Support representative.

3. If there are no event log verification errors then the **Download Controller Firmware** window shown in Figure 14-68 on page 374 will appear.

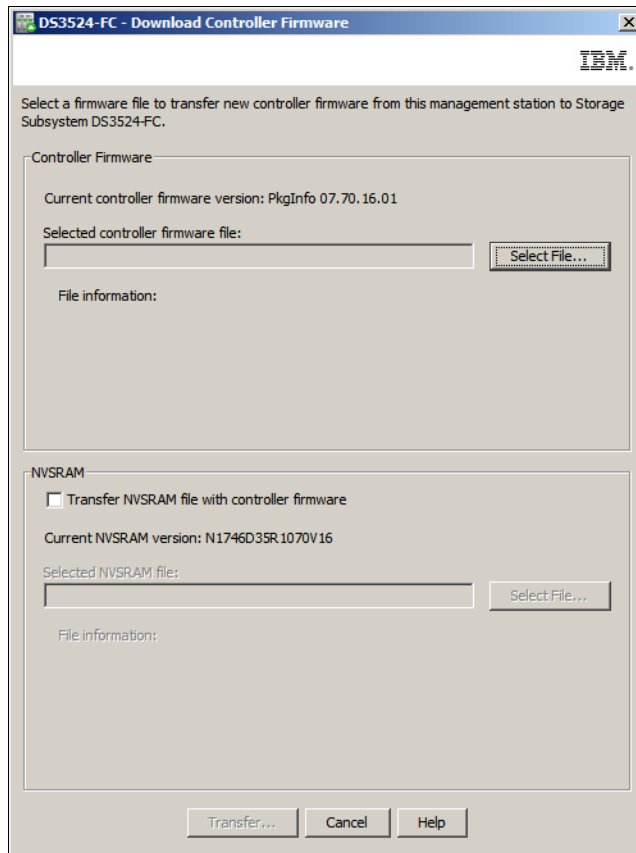


Figure 14-68 Download Controller Firmware

The **Download Controller Firmware** window has two sections. The first section, *Controller Firmware*, is where you select and specify the controller firmware file to be transferred. In the second section, *NVSRAM*, you select and specify the NVSRAM firmware file to transfer together with the controller firmware upgrade. When you select this option, both the controller firmware and controller NVSRAM upgrades will be done in one step.

**Note:** Install the controller firmware and controller NVSRAM updates in one step.

4. Click **Select File** in the Controller Firmware section and type or browse to the file name (Figure 14-69 on page 375).



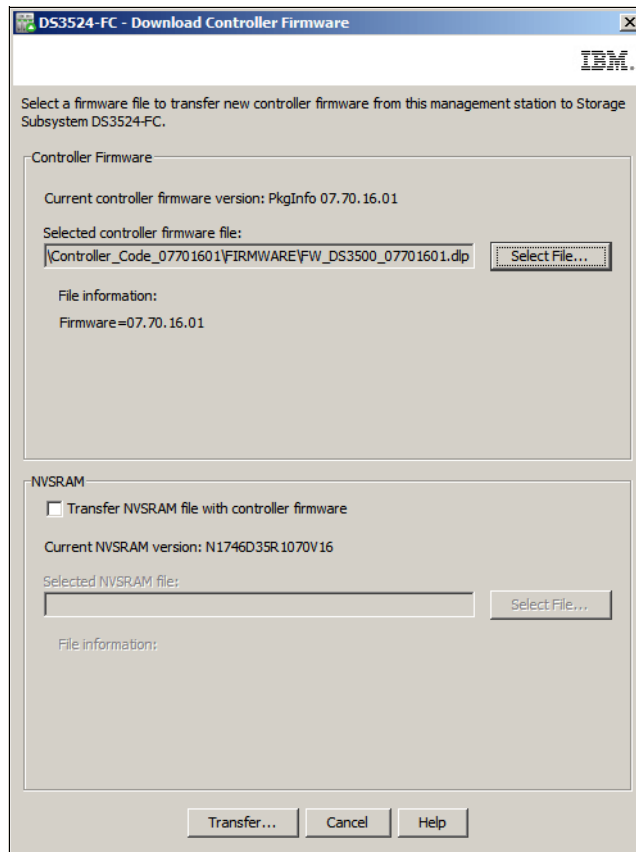


Figure 14-69 Download Controller Firmware - Select controller file

5. 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 14-70 on page 376).

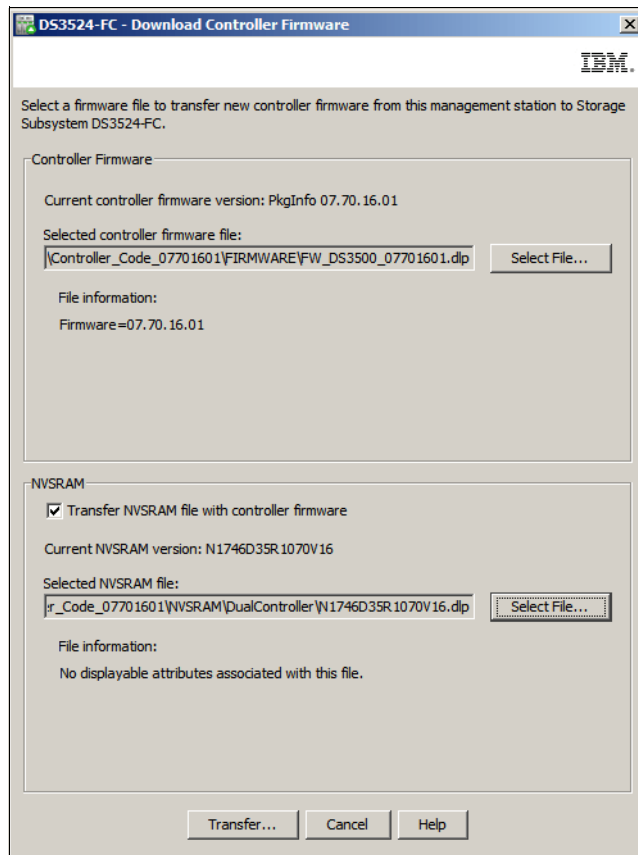


Figure 14-70 Download Controller Firmware - With NVSRAM firmware

6. Click **Transfer** to submit the new microcode to the storage subsystem. The Confirm Download window appears (Figure 14-71). Read it and click **Yes** to continue.

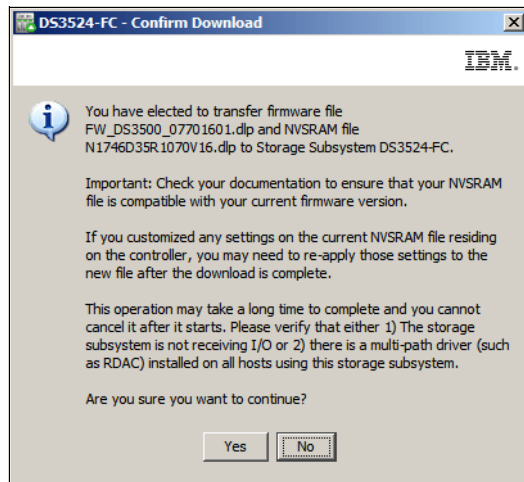


Figure 14-71 Confirm Download

7. The microcode will now be transferred to the storage subsystem and applied to both installed controllers. This can take some time. You can monitor the status in the Downloading window that appears (Figure 14-72 on page 377).

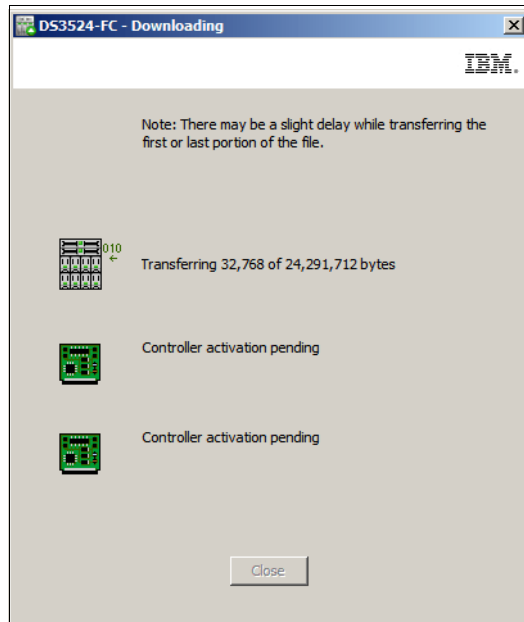


Figure 14-72 Downloading the firmware

After the controller firmware has been transferred, it is activated on both controllers. This step can also take some time. The activation status is reported as shown in Figure 14-73.

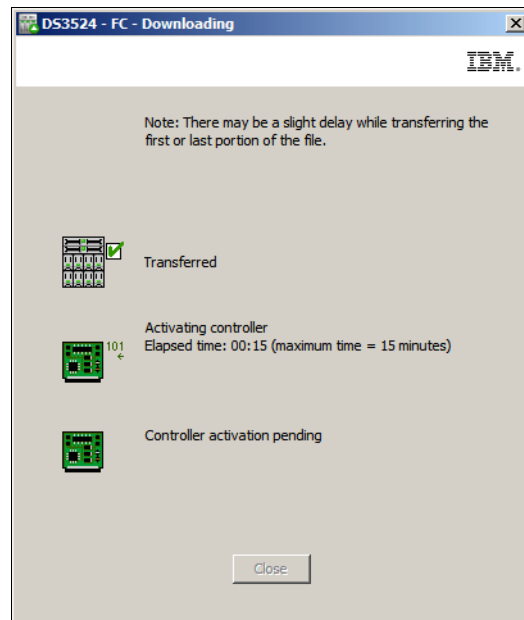


Figure 14-73 Activating the firmware

8. After the firmware upgrade has completed successfully (Figure 14-74 on page 378), click **Close** to return to the Download Firmware window (Figure 14-60 on page 369).




Figure 14-74 Download successful

## Download Controller NVSRAM

Follow the steps described in this section to install the controller NVSRAM.

**Note:** The **Download Controller NVSRAM** function is also accessible from the **Advanced** → **Maintenance** → **Download** → **Controller NVSRAM** menu.

1. To download a new NVSRAM file to your storage subsystem, click **Download Controller NVSRAM** (Figure 14-60 on page 369) to open the associated window (Figure 14-75). Click the **Select File** button to select the NVSRAM file.




Figure 14-75 Download Controller NVSRAM

2. Click the **Select File** button to select the NVSRAM file as shown in Figure 14-76 on page 379.




Figure 14-76 Select the NVSRAM file to install

3. After you have selected the NVSRAM file that you want to install, click the **Transfer** button to begin installing the new NVSRAM file on the storage subsystem. The **Confirm Download** window appears (Figure 14-77). Read the information and click **Yes** to start the NVSRAM download.

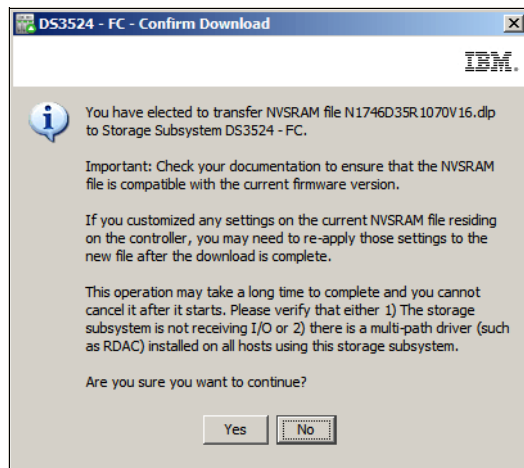


Figure 14-77 Confirm NVSRAM Download

4. The **Downloading** window appears and shows you the status of the transfer and the activation on the controller(s) as shown in Figure 14-78.

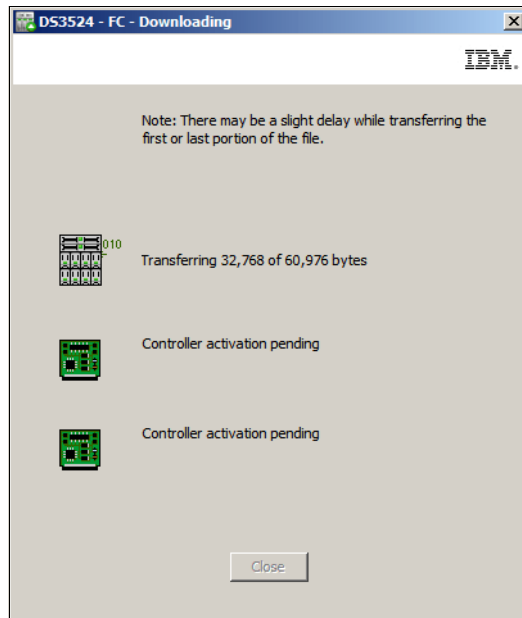


Figure 14-78 Download Progress NVSRAM

5. After the NVSRAM file is transferred and activated successfully (Figure 14-79), click **Close** to return to the Download Firmware window (Figure 14-60 on page 369).

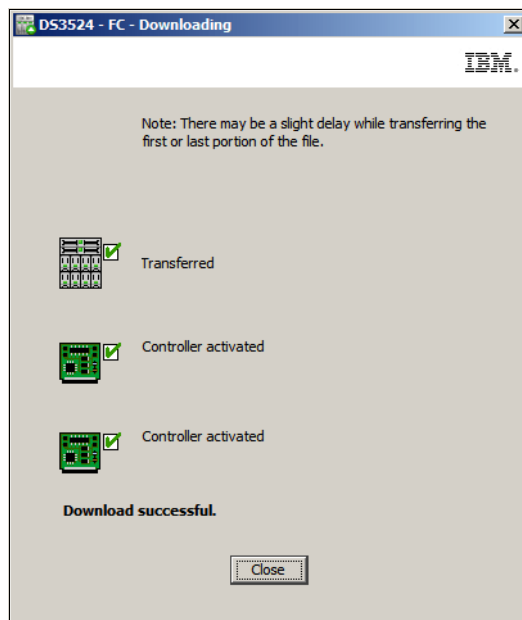


Figure 14-79 Download NVSRAM successful

## Download Drive Firmware

The hard drive firmware upgrade process is disruptive and requires that all host I/Os are suspended for the duration of the update. Plan a production outage in advance and follow these steps (when I/O operations are suspended) to perform the drive firmware upgrade:

**Note:** The **Download Drive Firmware** function is also accessible by clicking **Advanced** → **Maintenance** → **Download** → **Download Drive Firmware**.

1. Click **Download Drive Firmware** in the Download Firmware (Figure 14-60 on page 369) menu to open the Download Drive Firmware wizard. Read the information that is displayed on the introduction window (Figure 14-80) and then click **Next**.

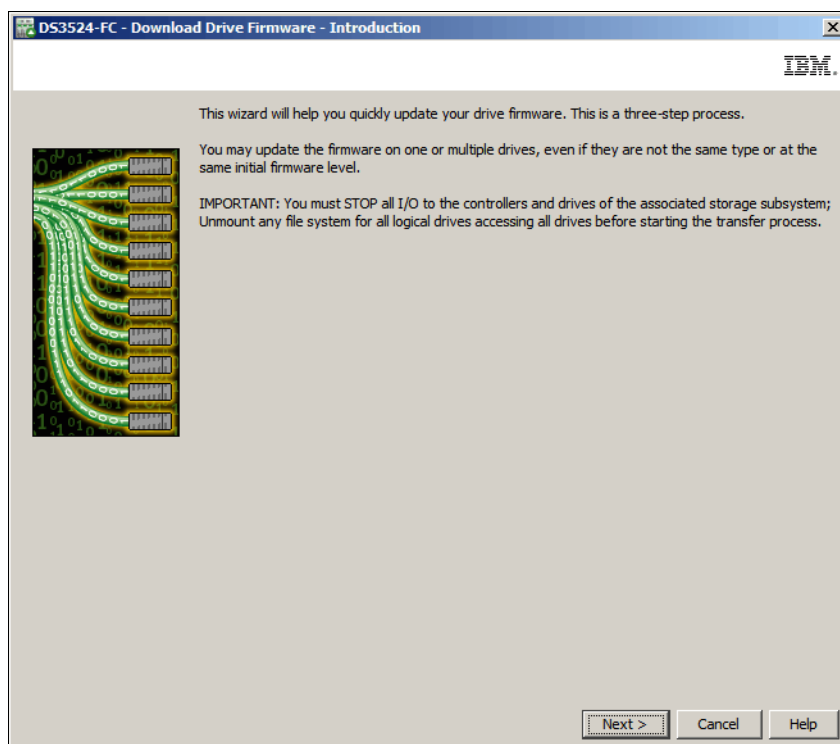


Figure 14-80 Download Drive Firmware - Introduction window

2. The **Download Drive Firmware - Add Packages** window appears (Figure 14-81 on page 382). This window has two main sections:
  - The *Current package* section contains a list of all the drive models and the firmware levels that are currently installed on each of these drive models.
  - The *Selected packages* section contains the list of firmware package files to be transferred after you have added them.

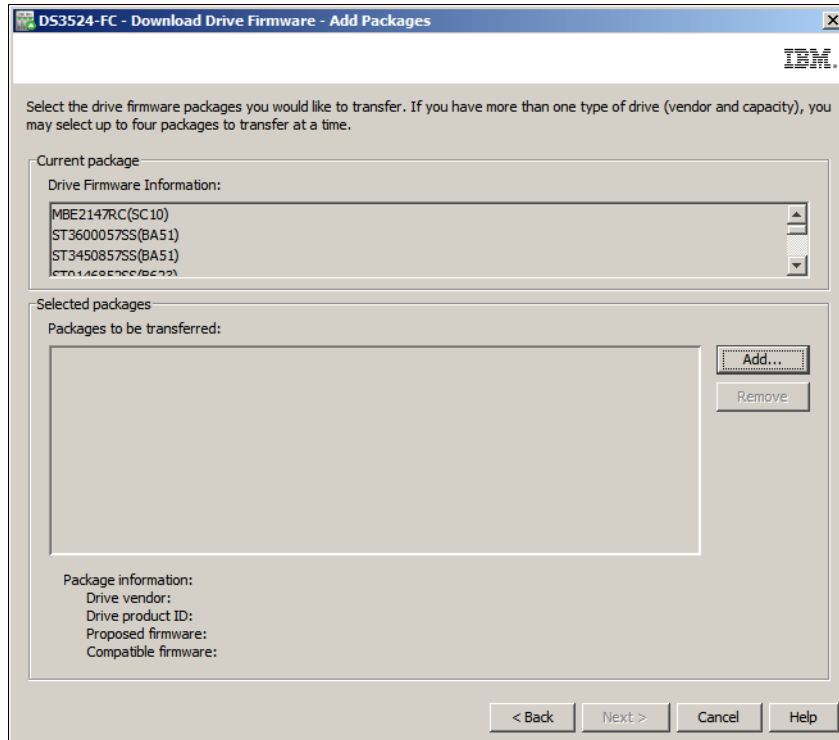


Figure 14-81 Select drive firmware package

3. Click **Add** to open the **Select A Package** window (Figure 14-82) and browse to the folder that contains the drive package files.

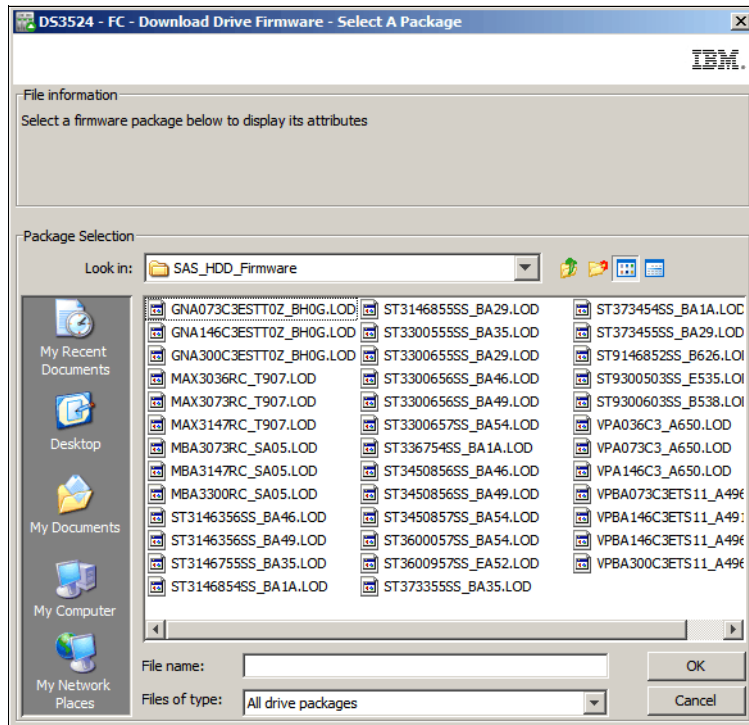


Figure 14-82 Download Drive Firmware - Select A Package



As you highlight each package listed in the window shown in Figure 14-82 on page 382, the **File Information** section will report if the selected drive firmware package is Compatible (Figure 14-83) or Not Compatible (Figure 14-84) with the types of drives that are installed in your system.




Figure 14-83 Drive firmware is compatible




Figure 14-84 Drive firmware is not compatible

4. Repeat these steps for every drive type that you want to update. You can update up to four drive types at a time. After you have selected the compatible drive firmware packages, click **Next** in the Download Firmware window (Figure 14-85) to continue.

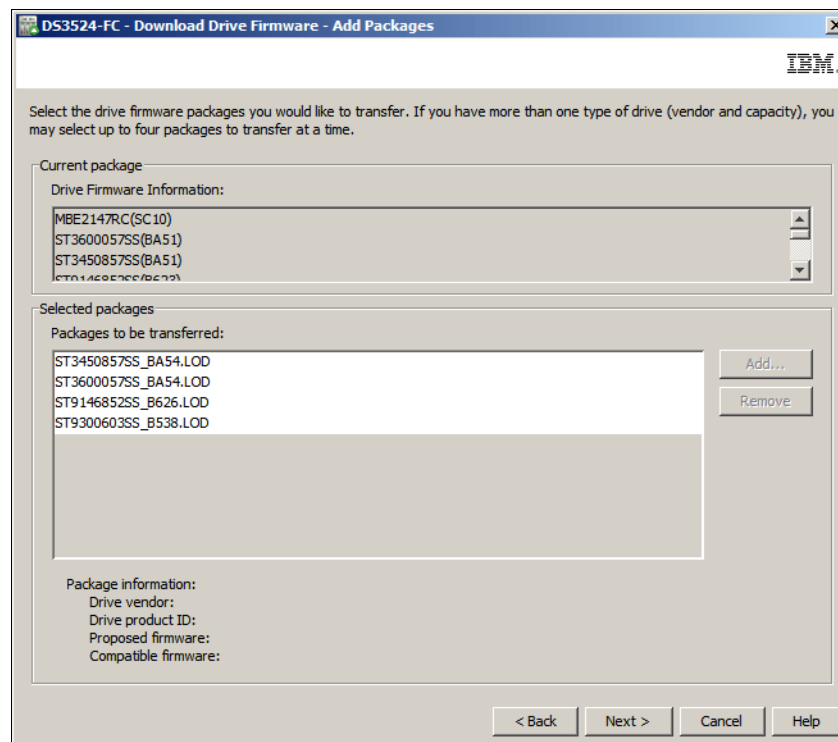


Figure 14-85 Drive packages added

5. Now you have to select the drives in the subsystem you want to update (Figure 14-86 on page 384). You only see the compatible drives in this window. Check **Select all** to update all compatible drives.

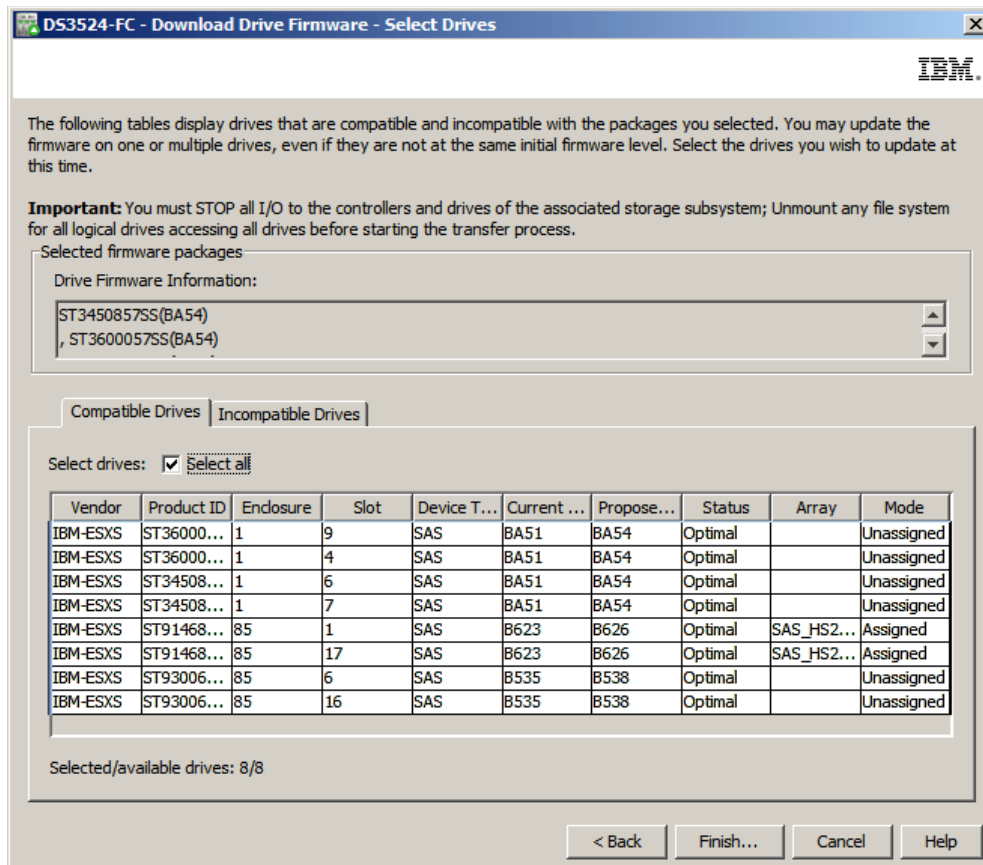


Figure 14-86 Compatible drives which can be updated

- The drives that are not compatible with the selected updates are shown on the Incompatible Drives tab (Figure 14-87). These are other drive types, and you will have to use other firmware packages.

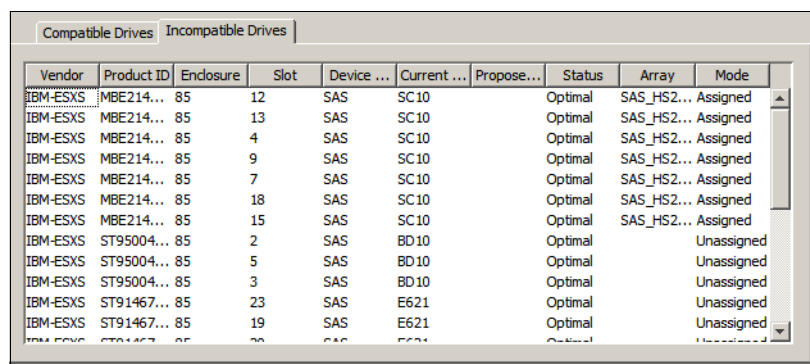


Figure 14-87 Incompatible drives tab

- Click **Finish** to update the drive firmware. A confirmation window appears (Figure 14-88 on page 385). Type yes and click **OK** to start the update.




Figure 14-88 Confirm Download

8. After the update completes, a summary is displayed (Figure 14-89). If you want to update more drives, click **Transfer More** to select the compatible packages. Otherwise, click **Close** to exit the drive update window.




Figure 14-89 Drive update complete

## 14.6 View Event Log

You use this link to view a list of events that have occurred on your DS3500 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. These events are assigned either Informational or Critical priority. The default view is to display the last 100 events of all types, but you can change the view to show only critical

events, or up to 8191 events of all types. All critical events are sent to the SNMP management console or to the email recipient that you have configured to receive the alert notifications as described in “Configure Alerts” on page 164.

**Note:** The DS3500 **Event Log** can also be viewed by clicking **Advanced** → **Troubleshooting** → **View Event Log**.

The initial Event Log view is shown in Figure 14-90.

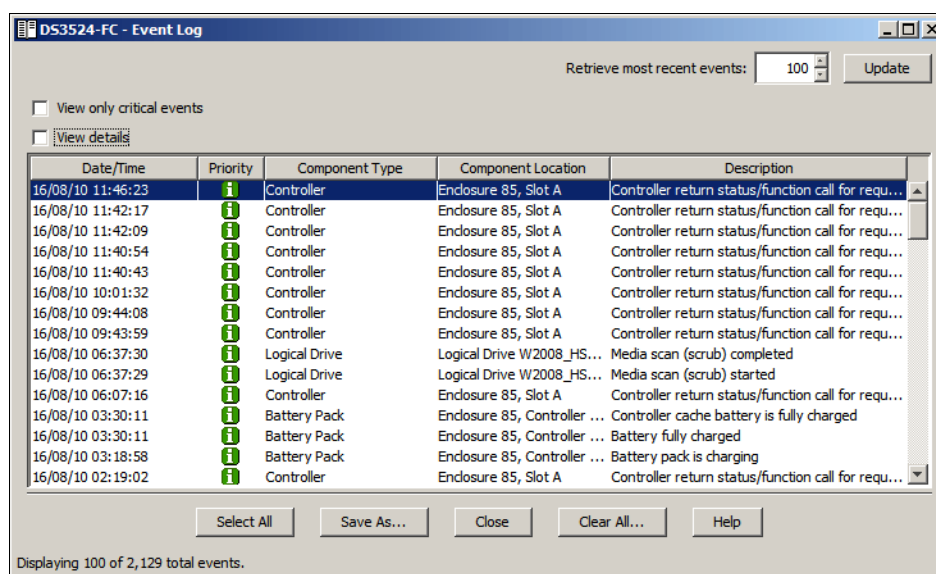


Figure 14-90 Event Log

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.
  - 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 by selecting the **View Details** check box (Figure 14-91) that is located in the upper left side of the Event Log window. The event details shows additional information about the event (Figure 14-92 on page 387).




Figure 14-91 View details check box

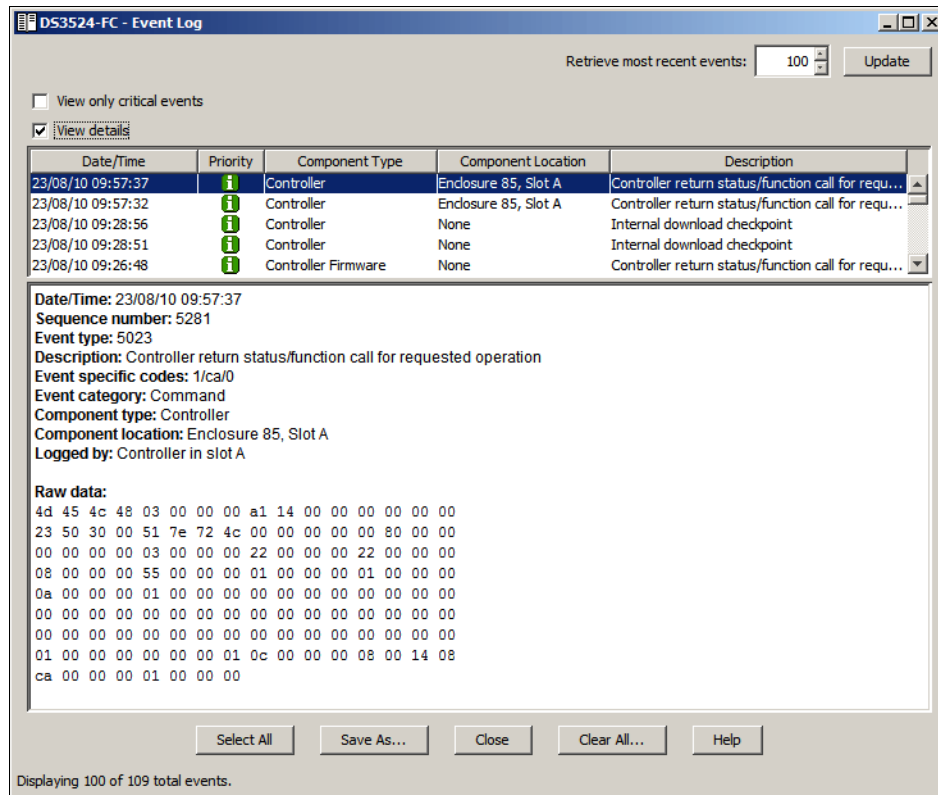


Figure 14-92 Viewing the Event Log details

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 that 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 an IBM Support representative.

In the right upper corner, you can set how many events to display; the default initial value is 100. To change this, type another value and click **Update** (Figure 14-93). The bottom left hand part of the window shows the total number of events stored on the subsystem. The maximum number of event is 8191. If more events occur, the old events are removed from the log on a first in, first out basis.




Figure 14-93 Change the number of recent events to display

In the left upper corner are two check boxes:

- **View only critical events:** Select this to display only the critical events as shown in Figure 14-94.
- **View details:** Select this check box to view the details of a highlighted event in the lower half of the window as shown in Figure 14-94.

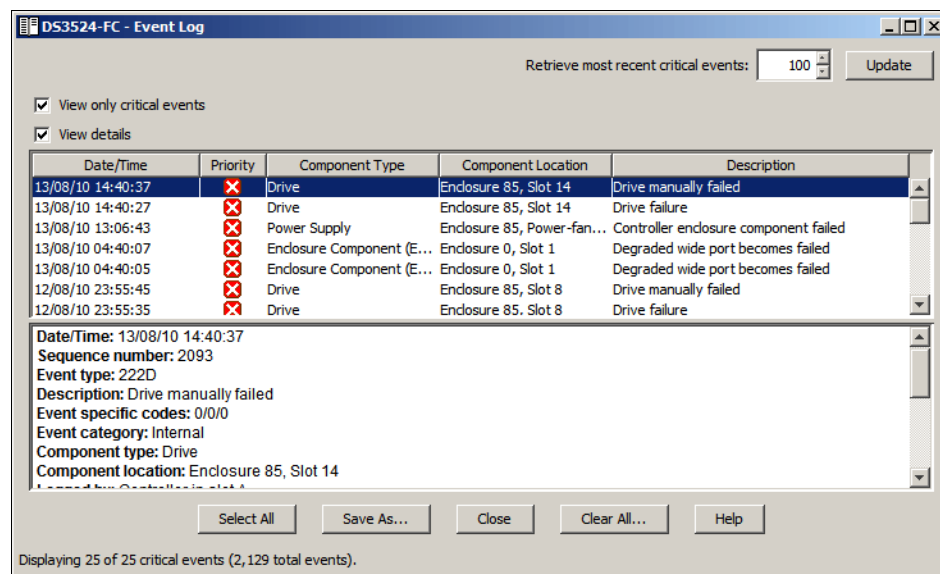


Figure 14-94 View only Critical Events

The buttons at the bottom of the window are:

- **Select All:** Use this button to highlight all events (Figure 14-95 on page 389).




Figure 14-95 Select 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, click the **Select All** button and then click **Save As**. You will be prompted for a location to save the events (the default extension is \*.log).

Figure 14-96 shows a sample log file viewed in Notepad.

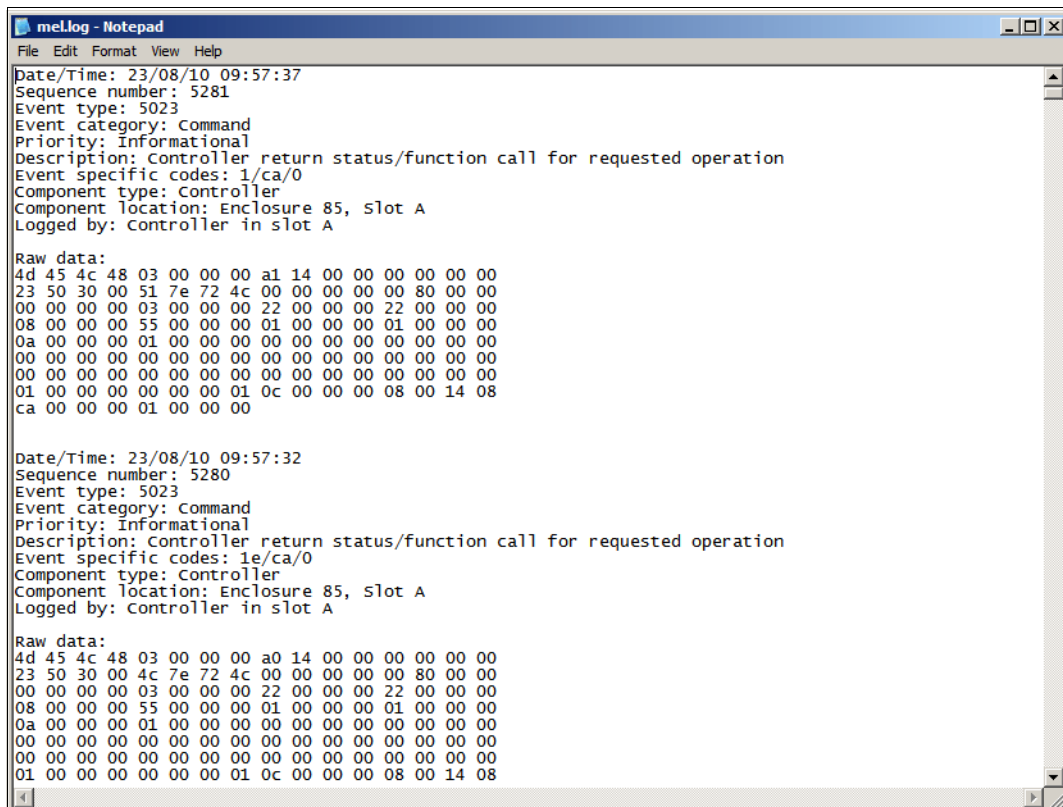


Figure 14-96 Event view - Log

- **Close:** Closes the Event viewer
- **Clear All:** Select this option to clear the event log. In the confirmation window (Figure 14-97), type yes to delete all the events. Save all event logs before clearing the log.

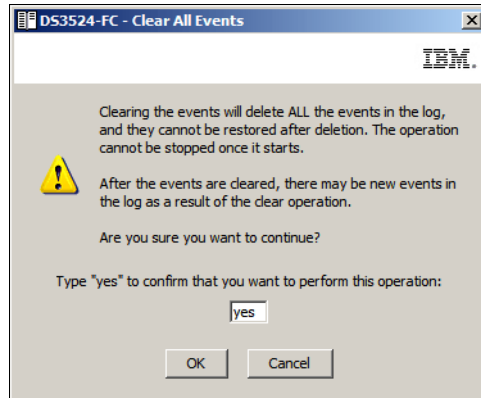



Figure 14-97 Event view - Confirm Clear All Events

- **Help:** Opens the online help.

## 14.7 Performance Monitor

In this section we discuss the DS Storage Manager Performance Monitor software. You can use the Performance Monitor to monitor the performance of various aspects of all of the subsystem logical drives and controllers, or you can select specific logical drives and controllers to monitor.

1. Open the **Performance Monitor** by clicking **Storage Subsystem** → **Performance Monitor** in the DS Storage Manager (Subsystem Management) window or by clicking the Stopwatch icon  in the Subsystem Management icon bar. The Performance Monitor window is shown in Figure 14-98.

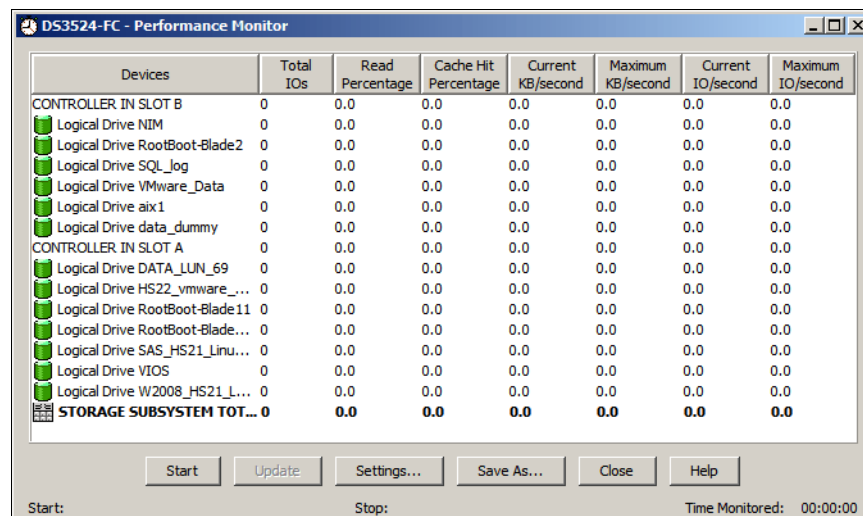


Figure 14-98 Performance Monitor



- Click the **Settings** button to configure the Performance Monitor (Figure 14-99). In this window you select the specific logical drives and controllers that you want to collect performance counters for. Click the **Select All** button to collect performance counters for all logical drives and controllers in the subsystem. You also specify how often the performance counters should be updated by specifying the **Polling interval** in seconds. Click **OK** to exit the settings window after you have selected the subsystem logical drives and controllers that you want to monitor.

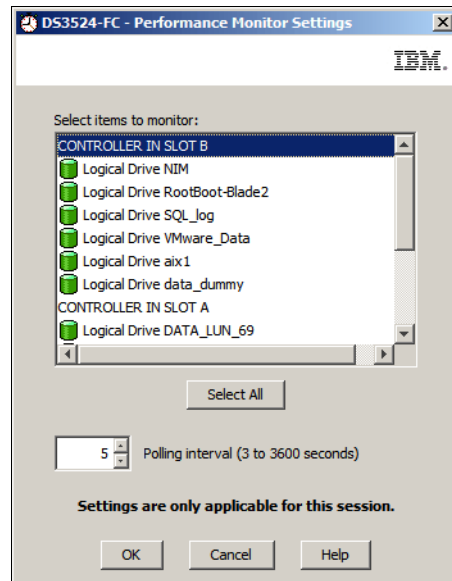


Figure 14-99 Performance Monitor Settings

- Click the **Start** button to begin collecting the performance counters (Figure 14-100). The performance counter values that are listed in the table are updated at the time interval that you specified in the Polling Interval setting in step 3. Click the **Update** button to force an immediate update of the performance counters.

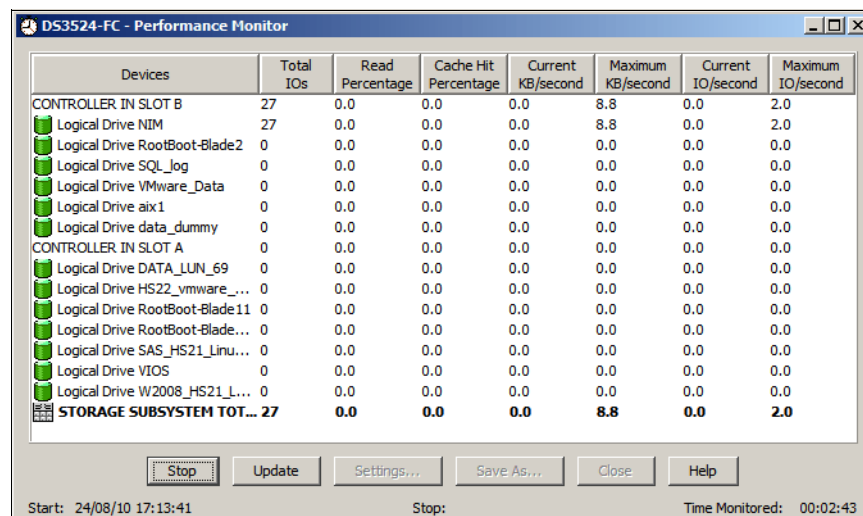


Figure 14-100 Performance Monitor collecting data

- When you have collected the performance data, click the **Save As** button to save it to a file. The default file extension is \*.perf but you can also save the performance data as a comma separated variable (\*.csv) file for viewing in commercial spreadsheet software.

## 14.8 Import/Export array

The DS3500 storage subsystem has the ability to export and import arrays, which enables you to safely move arrays between separate DS3500 systems without losing application data.

This is helpful when you have to upgrade or replace a DS3500 system with a newer model, or faster disks, but want to preserve the expansions and their data.

The export/import function checks that all the conditions to support the disk migration are met before placing the disks offline and allowing removal of the physical disks from the enclosure.

**Important:** You should always ensure that the source and destination DS3500 storage subsystems are running the same firmware levels before starting this procedure. Generally, upgrade both the source and the destination DS3500 subsystems to the latest firmware available before attempting any disk migration.

### 14.8.1 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 separate DS3500 storage subsystem. After you complete the export array operation, all of the disk drives are offline. Any logical drives or free capacity associated with the exported array are no longer shown in the storage management software.

#### Export array prerequisites

In this section we describe all of the prerequisites that must be completed before you can export an array.

- ▶ You must remove any non-exportable components before you can start the export array procedure. You must remove these components for each array that you want to export:
  - Host-to-logical drive mappings
  - VolumeCopy pairs
  - FlashCopy logical drives and FlashCopy repository logical drives
- ▶ Exporting an array includes following these steps on the source storage subsystem:
  - Save the storage subsystem configuration. This is a precaution to help you restore your configuration in the event of a problem during the export.
  - Back up the data on the logical drives in the array. Stop all I/O, and unmount or disconnect the file systems on all the logical drives in the exported array.
  - Locate the array, and label the disk drives. Use the Locate Array function to flash the LEDs on the drives in the array, and then label them with the source and destination storage subsystem names, array name, and total number of drives in the array.
  - Obtain blank disk drive CRUs or new disk drives. Upon removing the drives in the array, you must replace them with blank drive canisters or new drives to maintain proper airflow within the enclosure
- ▶ Exporting an array includes these steps on the target storage subsystem:
  - Make sure that the target storage subsystem has enough available empty disk drive slots for the number of drives that you will be importing.
  - Make sure that the target storage subsystem supports the type of disk drives that you will import. You must also ensure that you do not exceed the maximum number of drives supported by the storage subsystem.
  - Make sure that the target storage subsystem can support the new logical drives. The storage subsystem must support the RAID level of the array you are exporting and you

cannot exceed the maximum number of logical drives supported by the storage subsystem.

- Make sure that the latest version of firmware is installed on the controller.
- Make sure the target storage subsystem supports RAID 6 if you are exporting a RAID 6 array.
- Make sure that the latest version of the storage manager software is installed for the target storage subsystem.

## Export Array procedure

The procedure to export an array consists of the following steps:

1. Open the Storage Manager logical tab and select the array to be exported. Click **Advanced** → **Maintenance** → **Export Array** as shown in Figure 14-101.

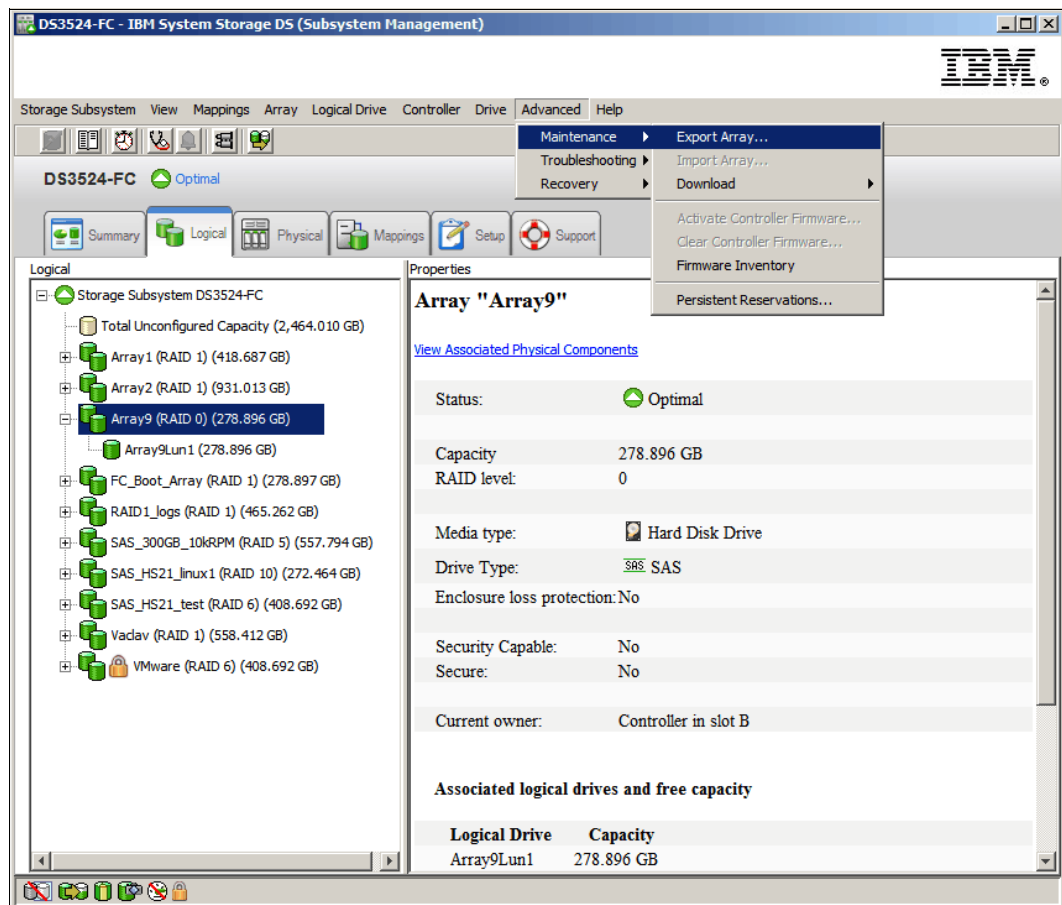


Figure 14-101 Select an array to be exported

2. The Introduction (Export Array) window opens as shown in Figure 14-102 on page 394. Read the information and click the **Next** button to continue.




Figure 14-102 Introduction (Export Array)

3. The **Preparation Checklist (Export Array)** window is displayed as shown in Figure 14-103 on page 395. This checklist lists the preparation instructions for both the source storage subsystem and destination storage subsystem. You can save the checklist to a text file for reference by clicking the **Save As** button. Select each procedure in the checklist as you complete the procedure. When all of the items in the checklist are completed, click the **Finish** button.

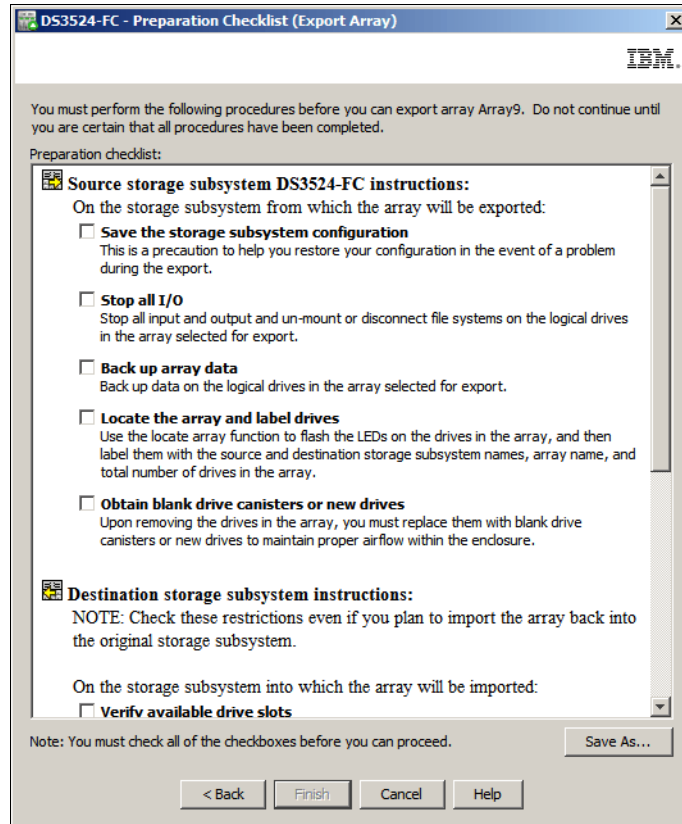


Figure 14-103 Preparation Checklist (Export Array)

4. All the previously mentioned requirements for export of an array (steps 1, 2 and 3) must be confirmed by selecting all check boxes in the Preparation Checklist window (Figure 14-103), and then clicking **Next**. You will then be asked to confirm that you want to export the array, as shown in Figure 14-104. Type **yes** and click **OK** to export the array.

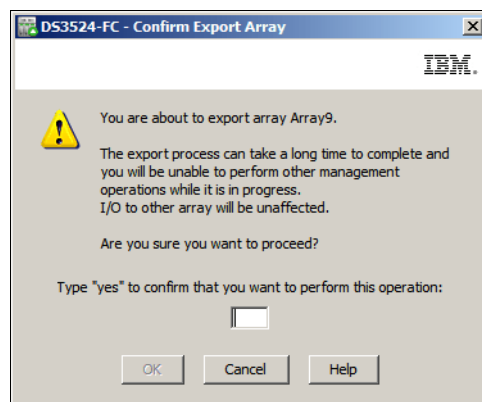


Figure 14-104 Confirm Export Array

5. The export array wizard then begins the export procedure as shown in Figure 14-105 on page 396.




Figure 14-105 Processing array export

6. The window shown in Figure 14-106 is displayed when the array has been exported. You can also confirm that the array has been exported by checking the Storage Manager logical tab as shown in Figure 14-107 on page 397. The status of the array is reported as Exported - ready for import.

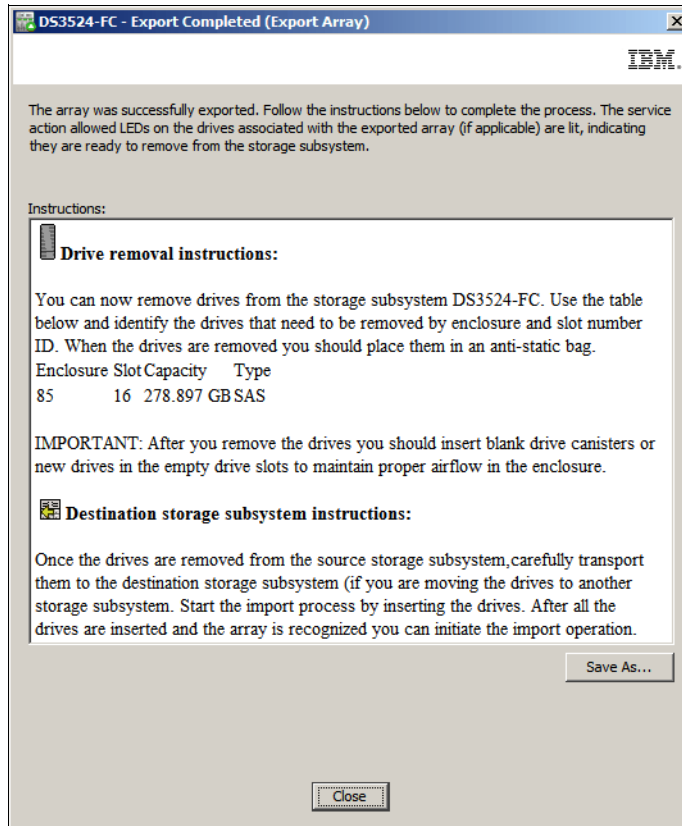


Figure 14-106 Export Completed (Export Array)

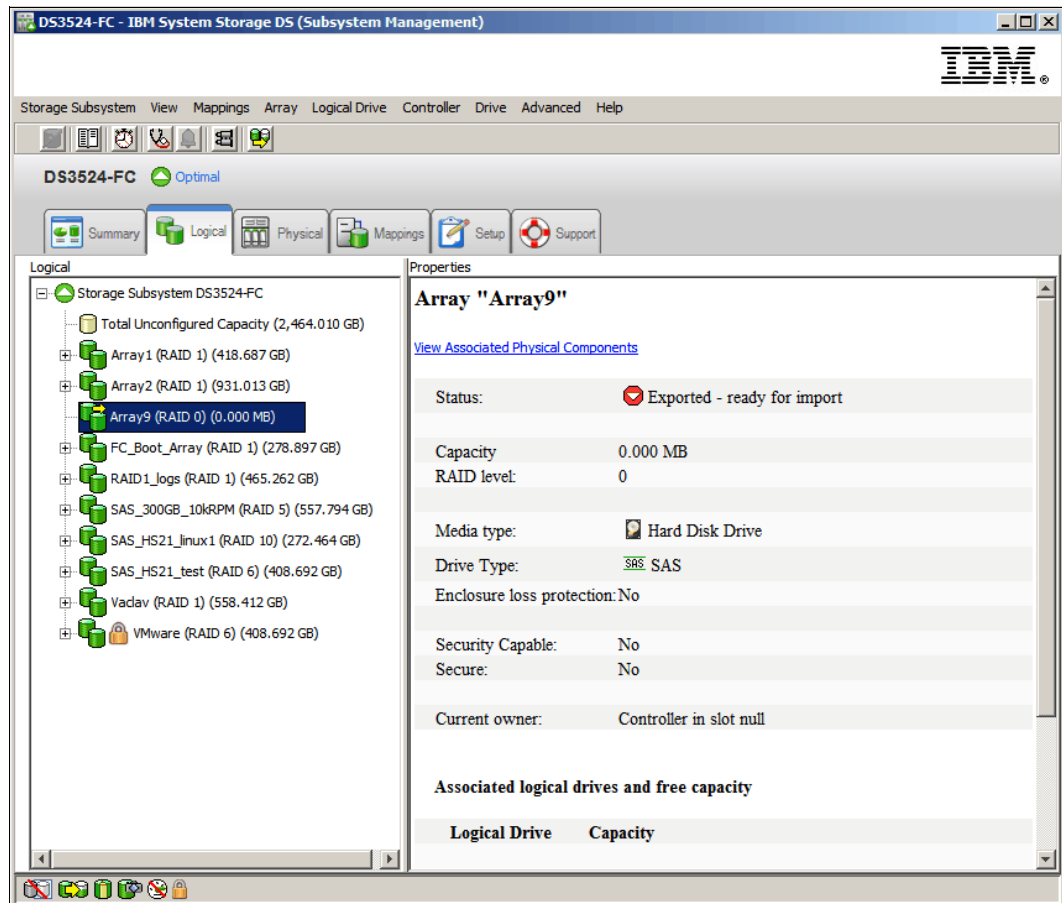


Figure 14-107 Exported - ready for import

7. Finally, the hard disks in the exported array can be carefully removed from the source DS3500 and installed into the empty slots of the target DS3500. The imported array then will have the status Exported - ready for import in the target DS3500, as shown in Figure 14-107.

## 14.8.2 Import Array procedure

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 will 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 involves performing 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

Certain 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

In this section we describe the steps to import an array:

1. Open the Storage Manager logical tab and select the array to be imported. The status of the array should be Exported - ready for import. Click **Advanced** → **Maintenance** → **Import Array** as shown in Figure 14-108.

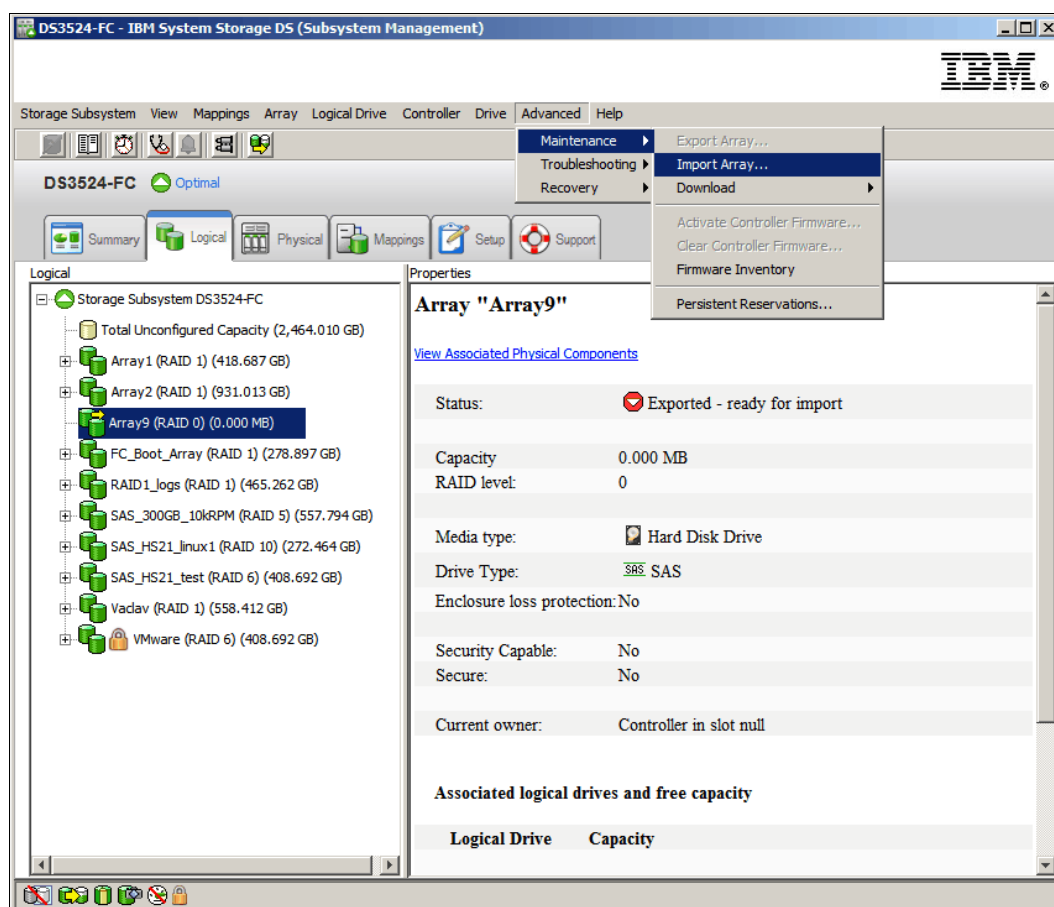


Figure 14-108 Import Array

2. The **Introduction (Import Array)** window is displayed as shown in Figure 14-109 on page 399. Read the information and then click the **Next** button to continue.



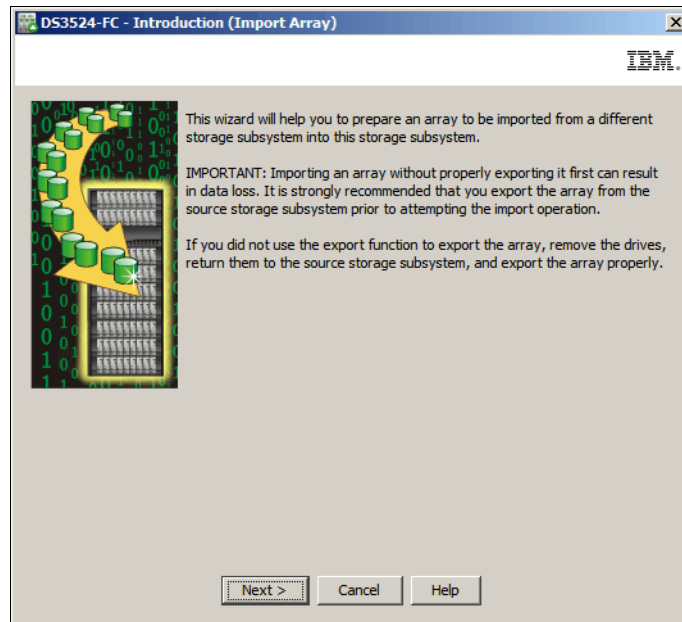


Figure 14-109 Introduction (Import Array)

3. The **Import Array (Import Array)** window is displayed as shown in Figure 14-110. This window displays the array information and associated disk drives.
  - Click the **Refresh** button to update the displayed information after inserting additional drives.
  - The **Force** button option allows you to import an array that has not been correctly exported, or when an exported array incurs lost data or damaged data in transit. This option can be used for disaster recovery purposes.

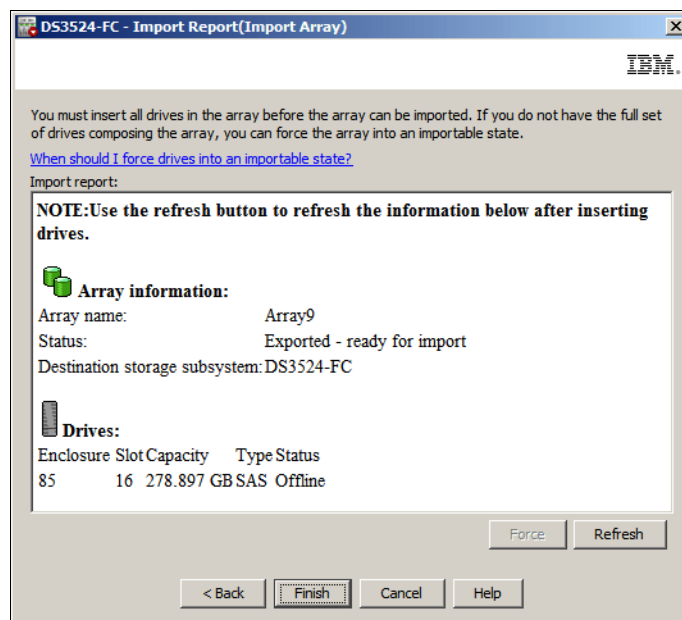


Figure 14-110 Import Array (Import Array)

- Click the **Finish** button in Figure 14-110 on page 399 to import the array. The **Confirm Import Array** window as shown in Figure 14-111 will appear. Type yes and then click **OK** to continue.

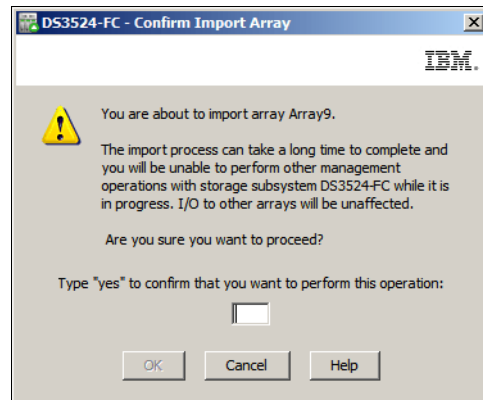


Figure 14-111 Confirm Import Array

- The Array Import wizard will start importing the array and display the status as shown in Figure 14-112.




Figure 14-112 Processing array import

- The Import Completed window shown in Figure 14-113 will be displayed when the array has been successfully imported. You can save the array migration to a text file by clicking the **Save As** button.




Figure 14-113 Import Completed

- The status of the newly imported array is reported as Optimal in the Storage Manager logical tab as shown in Figure 14-114.

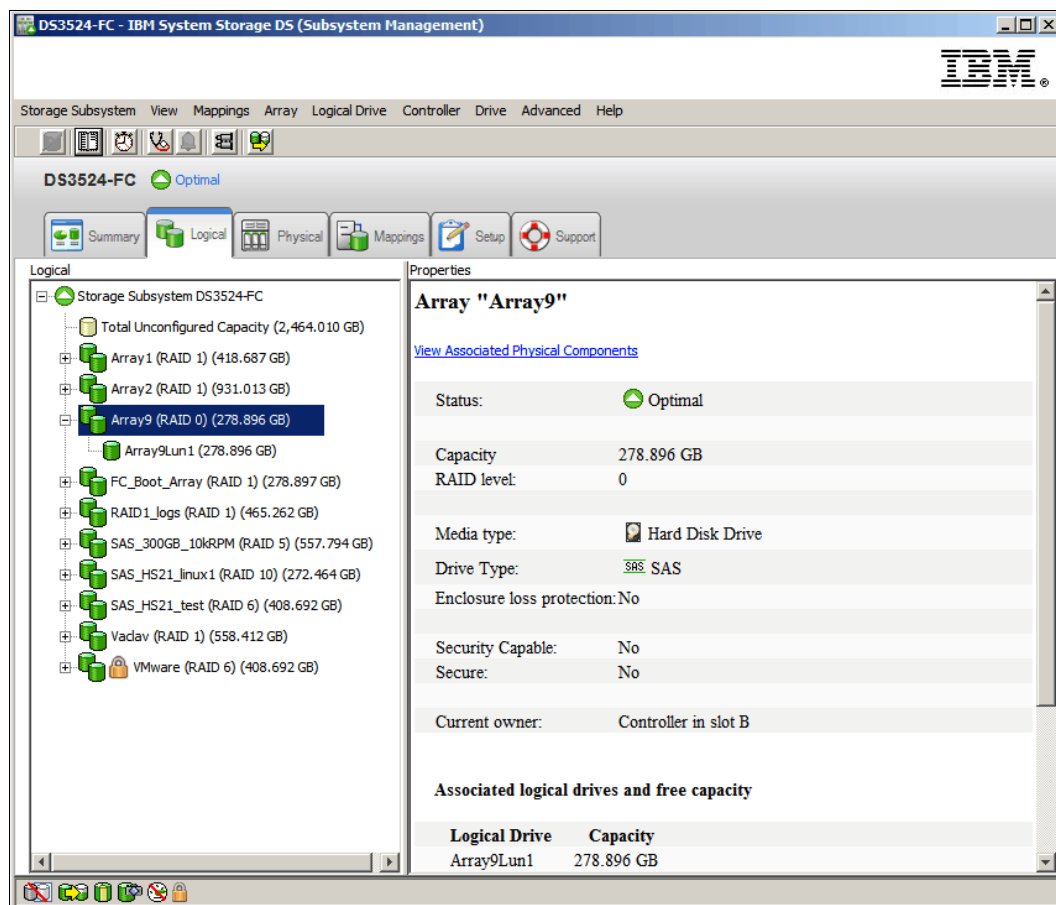


Figure 14-114 Imported array

- You will now need to create host to logical drive mappings for the logical drives associated with the newly imported array.

## 14.9 Maintenance - Persistent reservations

The Persistent Reservations option lets you view and clear logical drive reservations and associated logical drive registrations. Persistent reservations, which the cluster server software configures and manages, preserve logical drive reservations and logical drive registrations. Persistent reservations also prevent other hosts from accessing the logical drive. Unlike other types of reservations, a persistent reservation is used to perform the following functions:

- ▶ Reserve access across multiple host ports
- ▶ Provide various levels of access control
- ▶ Query the storage array about registered ports and reservations
- ▶ Provide for persistence of reservations in the event of a storage system power loss

**Note:** This function should normally only be used under the guidance of your IBM Customer Support representative.

The DS Storage Manager software allows you to manage persistent reservations from the Subsystem Management Window. The **Persistent Reservation** option enables you to perform the following tasks:

- ▶ View registration and reservation information for all volumes in the storage array
- ▶ Save detailed information about volume reservations and registrations
- ▶ Clear all registrations and reservations for a single volume or for all volumes in the storage array

**Note:** You can also manage persistent reservations through the script engine and the command line interface.

Perform the following steps to view the Persistent Reservations:

1. Click **Advanced** → **Maintenance** → **Persistent Reservations** to open the **Persistent Reservations** window as shown in Figure 14-115.

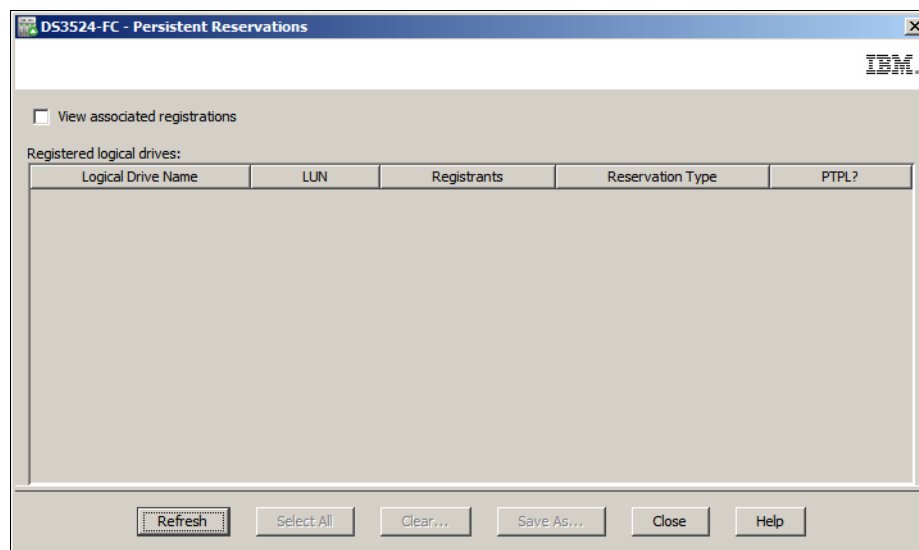


Figure 14-115 Persistent Reservations

2. To view the registrations that are associated with the logical drive reservation, either double-click the desired logical drive or highlight the logical drive and select the **View Associated Registrations** check box in the upper left of the dialog. You can view only one registration for a logical drive at a time.
3. Use the **Refresh** option to update the reservation and registration view.
4. Use the **Save As** option to save all reservation and registration information as a text file. All reservation and registration information for each logical drive is saved, no matter how many logical drives you have highlighted in the upper pane.
5. Use the **Clear** option to perform these functions:
  - Release logical drive reservations and remove associated registrations when the host port is no longer part of the configuration.
  - Clear old data that is associated with a logical drive that no longer has a reservation.

You can clear persistent reservations using this procedure:

1. In the upper-left corner of the Persistent Reservations dialog, make sure that the **View Associated Registrations** check box is not selected.
2. Click one or more desired logical drives. To select all of the logical drives, click **Select All**.
3. Click **Clear**.
4. In the text box in the Clear Registrations/Reservations dialog, type yes and click **OK**.

The reservation and registrations that were associated with the logical drives that you highlighted in the upper pane are cleared.

## 14.10 Troubleshooting - Drive channels

You can view detailed information about the DS3500 drive channels by clicking **Advanced** → **Troubleshooting** → **Drive Channels**. The Drive Channels Summary window displays as shown in Figure 14-116. From here you can manage the status of the two drive channels on the DS3500 storage subsystem. The ability to effectively manage the drive channel status on your storage subsystem lets you diagnose any problems on drive channels.

**Note:** This function should normally only be used under the guidance of your IBM Customer Support representative.

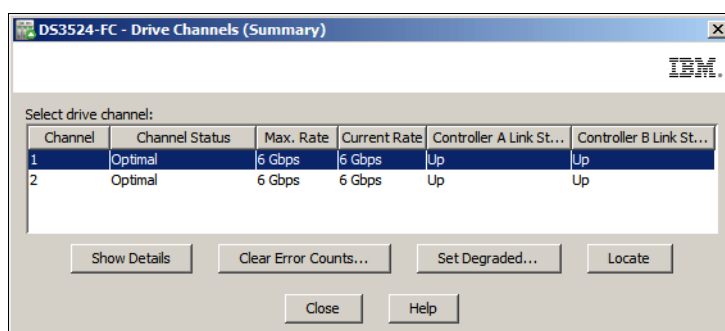


Figure 14-116 Drive Channels Summary

1. You can view details for a specific Drive Channel by selecting either Drive Channel 1 or 2 in the Summary window and then clicking the **Show Details** button to open the Drive Channel Status Details window as shown in Figure 14-117 on page 404. Click the **Save As** button to save these details as a text file.

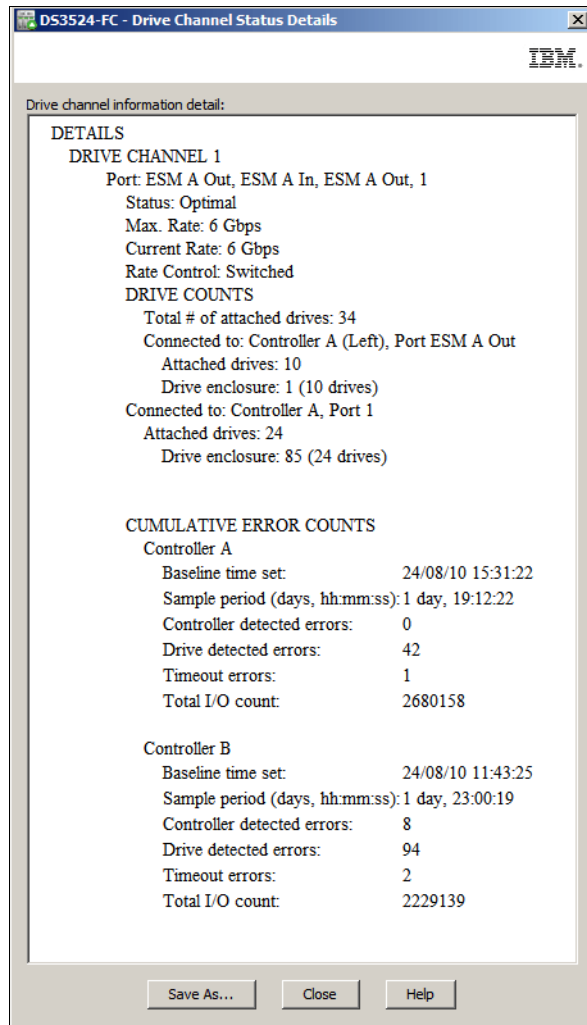


Figure 14-117 Drive Channel Status Details

- You can clear the cumulative error counts for both drive channels by selecting either Drive Channel 1 or 2 in the Summary view and then clicking the **Clear Error Counts** button.
- You can manually set a drive channel to the *Degraded* state to force all I/Os to use a different drive channel by selecting either Drive Channel 1 or 2 in the Summary view and clicking the **Set Degraded** button. You will be asked to confirm that you want to perform this operation. Type yes in the text box and click **OK** to continue.

**Important:** This procedure is an advanced troubleshooting technique and should not be attempted if there are no problems with your DS3500 storage system

- If the status of either of the two Drive Channels is Degraded, the Summary view reports the status as shown in Figure 14-118 on page 405. To set the Drive Channel to the *Optimal* status, select the Drive Channel in the Summary view and click the **Set Optimal** button. You will be asked to confirm that you want to perform this operation, Click **Yes** to continue.

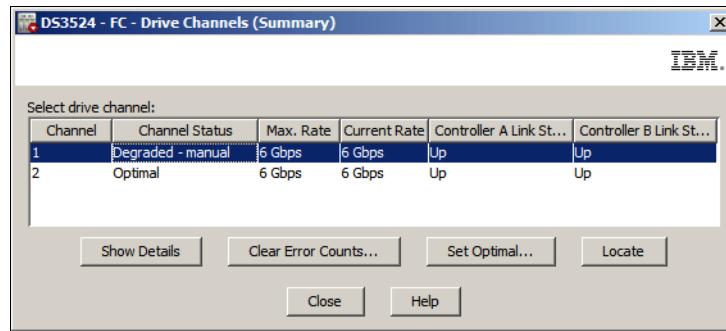


Figure 14-118 Drive Channel 1 in Degraded status

- To locate drive expansion enclosures that are connected to a drive channel, use the **Locate** option to physically locate and identify all drive expansion enclosures that are attached to that particular drive channel. An LED comes on for each drive expansion enclosure in the storage subsystem that is connected to the particular drive channel. Select either Drive Channel 1 or 2 in the Summary view and click the **Locate** button. If you select Drive Channel 2 the window shown in Figure 14-119 will be displayed. After you have located the drive enclosures on the Drive Channel, click **OK** to stop the LEDs from blinking.




Figure 14-119 Locate Drive Channel

## 14.11 Troubleshooting - Run Diagnostics

You can use the controller diagnostics function to run various internal tests that verify if a selected controller is functioning correctly.

**Important:** The controller is not available for I/O when the controller diagnostics are running.

**Note:** You should normally only use this option only under the guidance of your IBM Customer Support representative

Follow these steps to use the run the controller diagnostics:

- Open the Physical tab of the DS Storage Manager (Subsystem Management) and select the controller that you want to test (either controller A or controller B). Then click **Advanced** → **Troubleshooting** → **Run Diagnostics** → **Controller** as shown in Figure 14-120 on page 406. In this example, we have selected controller A.

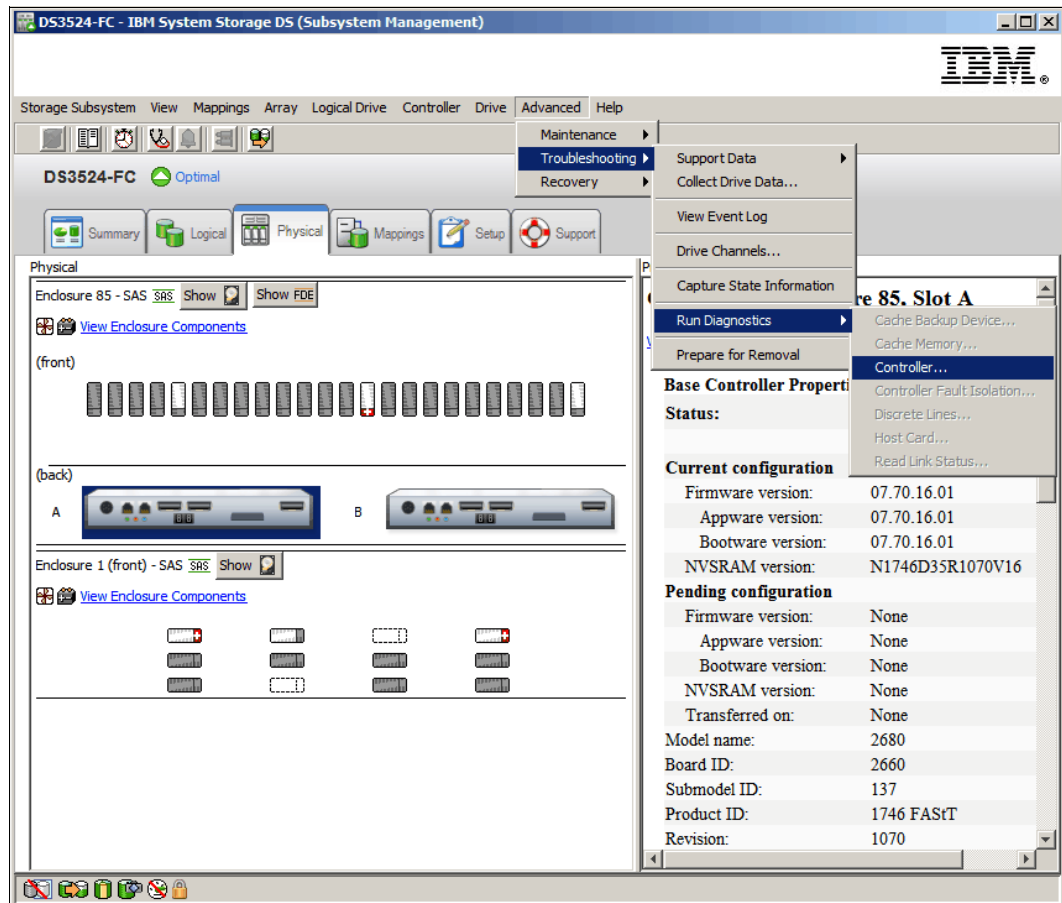


Figure 14-120 Run Controller Diagnostics

2. The **Run Diagnostics** window shown in Figure 14-121 on page 407 will open. In this window you can select the tests that you want to run on the selected controller. Select the check boxes for the diagnostic tests to be run. There are two tests applicable to the DS3500:
  - Read test: This test initiates a read command as it would be sent over an I/O data path. The read test compares data with a known, specific data pattern, and the test checks for data integrity and redundancy errors. If the read command is unsuccessful or the data compared is not correct, the test considers that the controller is in error. The test fails the controller.
  - Write test: This test initiates a write command as it would be sent over an I/O data path to the diagnostics region on a specified drive. This diagnostics region is then read and compared to a specific data pattern. If the write fails or the data compared is not correct, the test fails the controller and places the controller offline.

**Note:** The data loopback test only runs on controllers that have Fibre Channel connections between the controller and the drives. Because the DS3500 has SAS connections between the controller and the drives, this test does not apply to the DS3500.



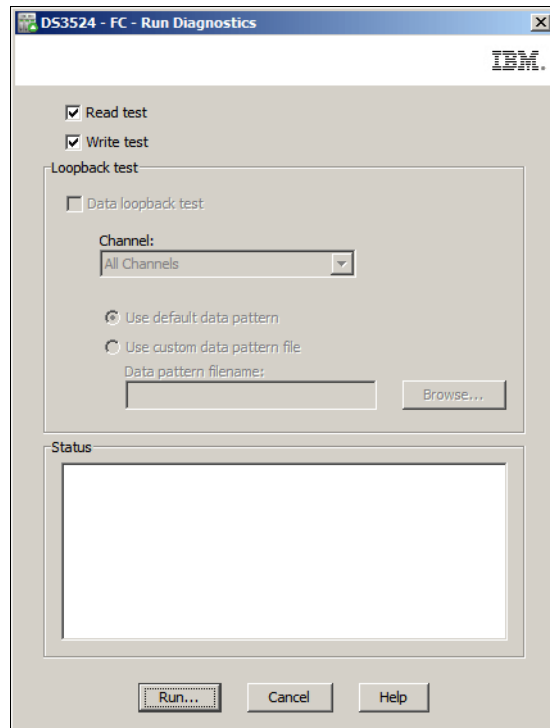


Figure 14-121 Run Diagnostics

3. After you have selected the tests to run, click the **Run** button. You will be asked to confirm that you want to run the tests as shown in Figure 14-122. Read the important information, then type yes in the text box and click **OK** to continue.

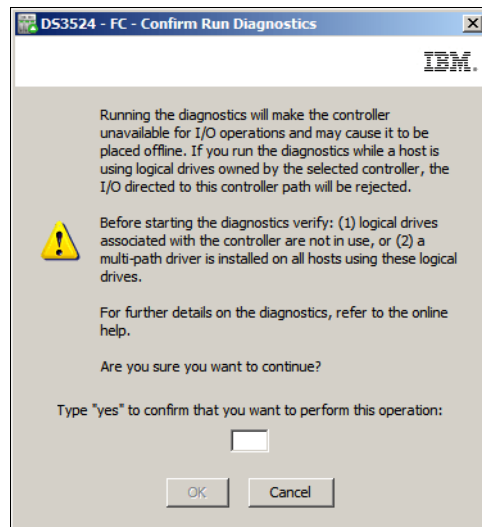


Figure 14-122 Confirm Run Diagnostics

4. The tests will run and the test report will be displayed in the **Status** section of the Run Diagnostics window as shown in Figure 14-123 on page 408.

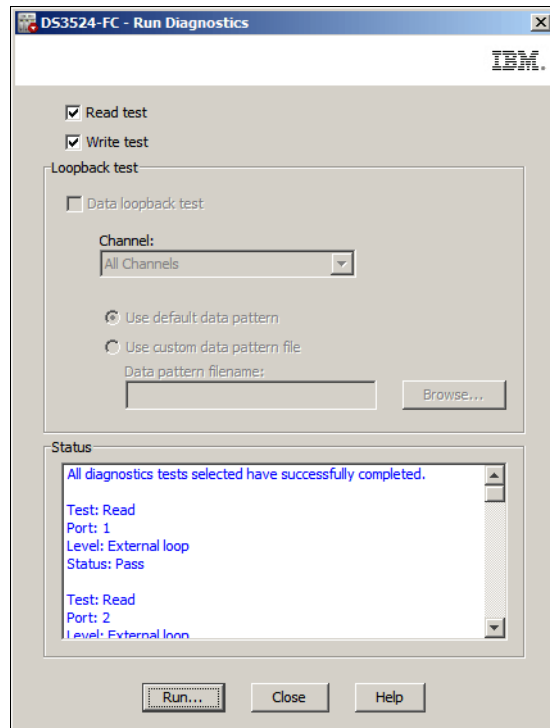


Figure 14-123 Run Diagnostics - completed

## 14.12 Troubleshooting - Prepare for removal

Before you remove a component, the DS Storage Manager software checks to see if another component should be serviced first. If you can safely remove the component, the Service Action Allowed (SAA) LED on the component comes on. The types of components that can be prepared for removal are:

- ▶ DS controllers
- ▶ Drives
- ▶ DS storage enclosure and expansion enclosure power supplies
- ▶ Expansion enclosure ESMs

**Attention:** Use this option only when instructed to do so by the **Recovery Guru** (see 14.13, “Recovery Guru - Recover from Failure” on page 409) or by your IBM Customer Support representative. For CRUs that have an SAA LED, after the SAA LED comes on, you cannot turn it off without removing the CRU.

To prepare a component for removal, perform the following steps:

1. To access the software tool, click **Advanced** → **Troubleshooting** → **Prepare for removal**. The window shown in Figure 14-124 on page 409 will open.

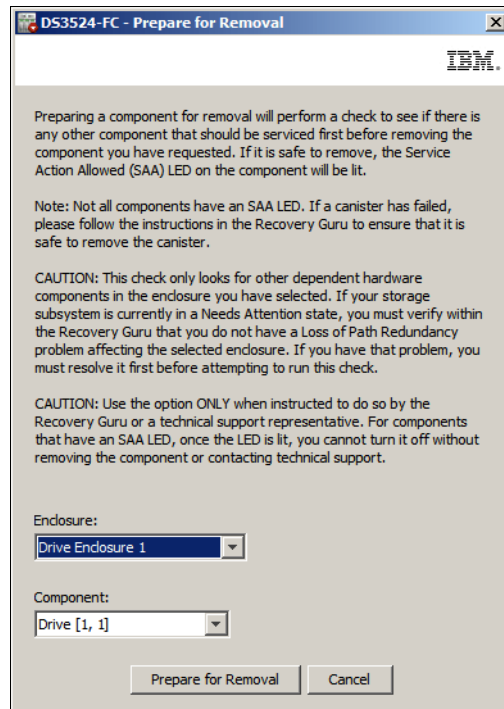


Figure 14-124 Prepare for removal

2. In this window, select the **Enclosure** and then the **Component** within that enclosure from the two drop-down menus. Then click the **Prepare for Removal** button.

## 14.13 Recovery Guru - Recover from Failure

The Recovery Guru is a component of the DS Storage Manager (Subsystem Management Window) (SMW) that diagnoses storage subsystem problems and recommends recovery procedures to fix those problems. Problems in the storage subsystem are indicated as follows:

- ▶ A *Needs Attention* storage subsystem icon appears within the following elements of the storage management software:
  - In the Status column on the Devices tab in the Enterprise Management window.
  - In the Storage Subsystem Status area below the menu bar in the SMW
  - In the title bar of the SMW
  - In the Status portlet on the Summary tab in the SMW
- ▶ The Recovery Guru Optimal toolbar button in the SMW changes to a blinking Recovery Guru Needs Attention toolbar button.
- ▶ Non-Optimal component icons appear on the Logical tab and the Physical tab of the SMW.
- ▶ Critical Simple Network Mail Protocol (SNMP) notifications or email error messages are sent.
- ▶ The fault LEDs on the hardware come on.
- ▶ The **Recover from Failure** link appears in the DS Storage Manager (Subsystem Management) Support tab as shown in Figure 14-125 on page 410.

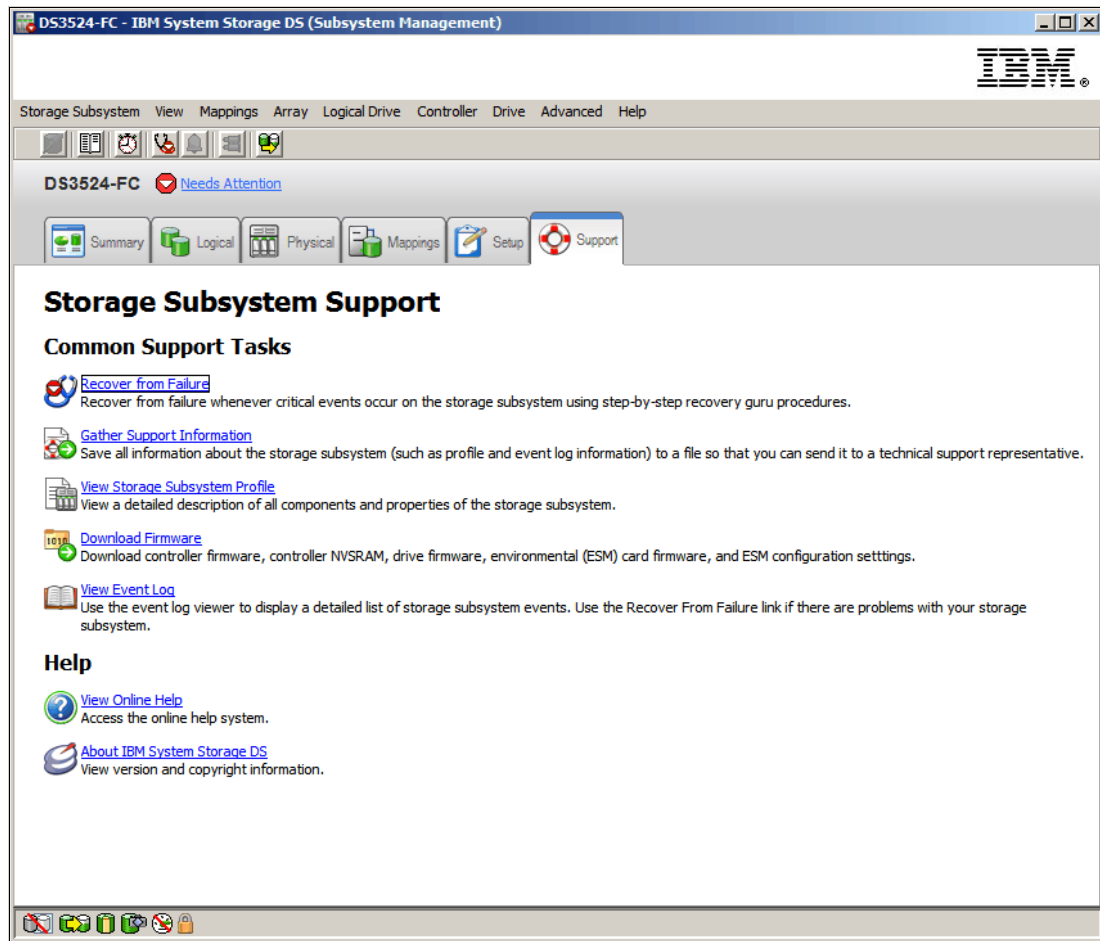


Figure 14-125 Support tab - Recover from Failure

Click the **Recover from failure** link to launch the **Recovery Guru** as shown in Figure 14-126 on page 411.

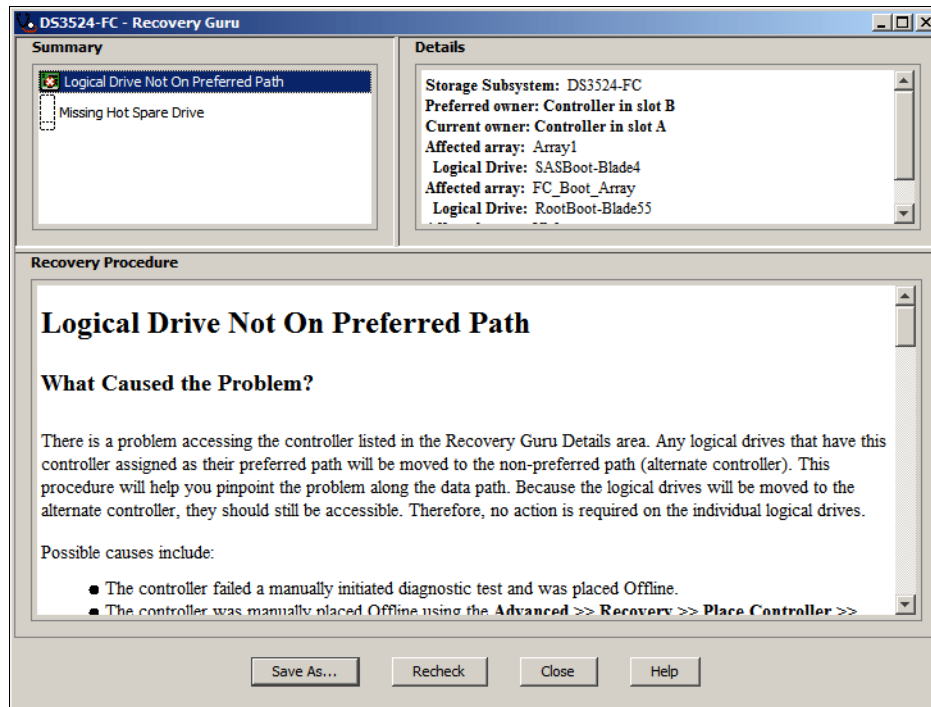


Figure 14-126 Recovery Guru

The Recovery Guru windows has three main sections, known as panes:

- ▶ **Summary pane:** All detected problems are listed here. There are two separate problems listed in the example shown in Figure 14-126.
- ▶ **Details pane:** Lists specific details for the problem selected in the Summary section. If there is more than one problem listed in the Summary section, select each problem to view the relevant details.
- ▶ **Recovery Procedure pane:** Provides troubleshooting and recovery steps for the specific problem that is selected in the Summary section.

At the bottom of the Recovery Guru window, there 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.

## 14.14 Common Recovery Commands

In this section we describe the common recovery commands that are accessible from the Advanced Recovery menu in the DS Storage Manager (Subsystem Management) window as shown in Figure 14-127 on page 412.

**Attention:** You should normally only use these functions if instructed to do so by the Recovery Guru or your IBM Technical Support representative.




Figure 14-127 Recovery commands

### 14.14.1 Initialize

When you select **Advanced** → **Recovery** → **Initialize** in the SMW, the following submenu options appear:

- ▶ **Drive**
- ▶ **Logical Drive**
- ▶ **Array**

#### Initialize - Drive

The initialize drive command is only used when only when you have moved a subset of drives that were previously part of a multi-drive array from one storage subsystem to another. Because you did not move the entire set of drives, the logical drive and array information about the drives is incomplete.

Therefore, to erase all logical drive and array information about these drives and make them available again for new logical drive creation, you must first initialize these drives. This action erases the logical drive and array information and returns the selected drives to an Unassigned state, which adds new or additional unconfigured capacity to the storage subsystem. You can then use this capacity to create additional logical drives.

**Important:** When you initialize a drive, all data on the drive is lost.

Perform these steps to initialize drives:

1. Open the physical tab in the DS Storage Manager (Subsystem Management) window and select one or more drives to be initialized. You can select more than one drive by pressing the CTRL key and clicking each drive.
2. Click **Advanced** → **Recovery** → **Initialize** → **Drive**.
3. You will be asked to confirm that you want to proceed. Type yes in the text box and click **OK** to continue.

## Initialize - Logical Drive

A logical drive is automatically initialized when it is first created. However, the Recovery Guru might advise you to manually initialize a logical drive to recover from certain failure conditions.

**Attention:** Initializing a logical drive, will destroy all the data on those drives. If the Recovery Guru advises that you initialize a logical drive, you should contact your IBM Customer Support representative before you begin this procedure.

Keep these guidelines in mind before you initialize a logical drive:

- ▶ Use this option only under the guidance of your IBM Customer Support representative.
- ▶ You cannot cancel the operation after it starts.
- ▶ Do not use this option if any modification operations are in progress on the logical drive or the array.
- ▶ Do not change the cache parameters of the logical drive while the Initialization operation is in progress.

Perform these steps to initialize a logical drive:

1. Select the logical drive to be initialized in the Logical tab of the DS Storage Manager (Subsystem Management) window.
2. Click **Advanced** → **Recovery** → **Initialize** → **Logical drive**.
3. You will be asked to confirm that you want to proceed. Type yes in the text box and click **OK** to continue.
4. You can monitor the progress of the logical drive initialization by selecting the logical drive in the Logical tab as shown in Figure 14-128 on page 414.

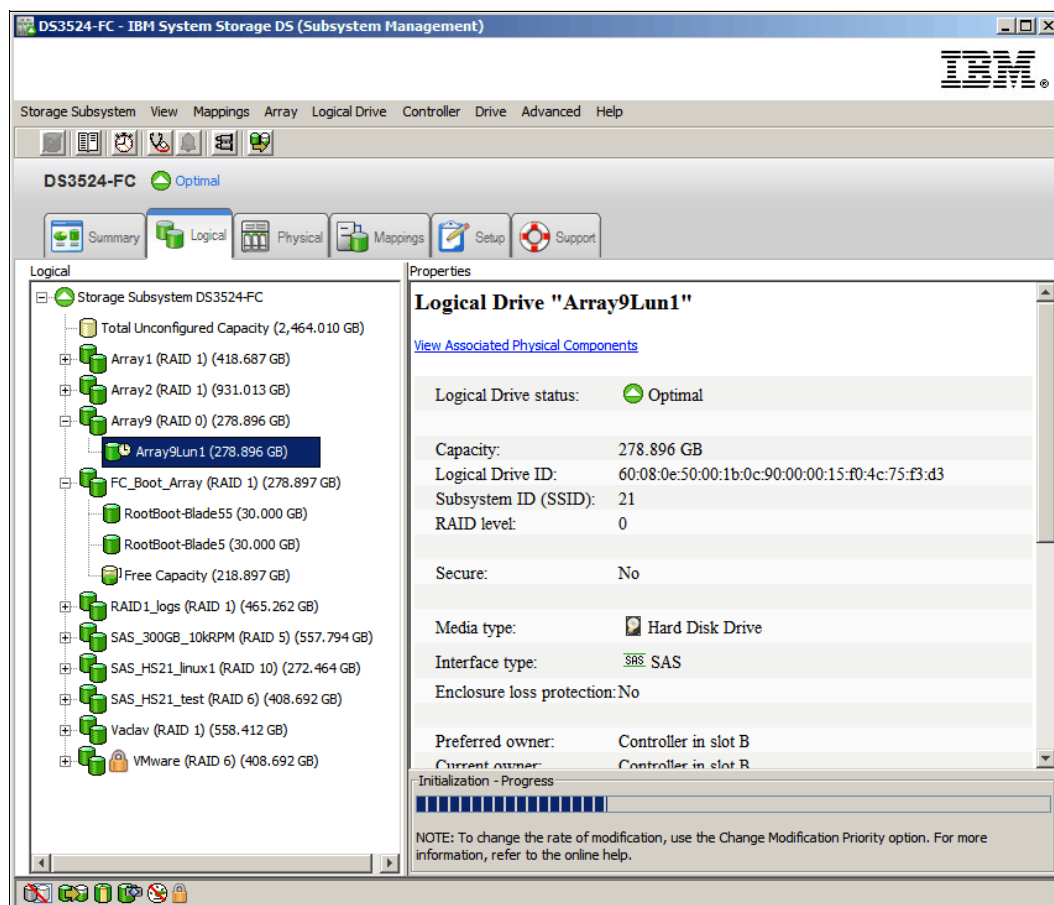


Figure 14-128 Initialize logical drive progress

## Initialize - Array

Use the **Initialize Array** option when the Recovery Guru advises that you manually initialize all of the logical drives in an array.

**Attention:** All data on all logical drives in the array will be destroyed. If the Recovery Guru advises that you initialize an array, you should contact your IBM Customer Support representative before you start this procedure.

Keep these guidelines in mind before you initialize an array:

- ▶ Use this option only under the guidance of your IBM Customer Support representative.
- ▶ You cannot cancel the operation after it starts.
- ▶ Do not use this option if the logical drives in the array do not have Optimal status and if any modification operations are in progress on the logical drive or the array.
- ▶ Do not change the cache parameters of the logical drive while the initialize operation is in progress.

Perform these steps to initialize an array:

1. Select the array to be initialized in the Logical tab of the DS Storage Manager (Subsystem Management) window.
2. Click **Advanced** → **Recovery** → **Initialize** → **Array**.



3. You will be asked to confirm that you want to proceed. Type yes in the text box, and then click **OK** to continue.
4. You can monitor the progress of each logical drive initialization in that array by selecting the logical drives in the Logical tab as shown in Figure 14-128 on page 414.

## 14.14.2 Revive drive

The Revive drive function will revive (attempt to spin up) a selected failed drive.

**Important:** Correct use of this option depends on the data configuration on all of the drives in the array. Never attempt to revive a drive unless you are supervised by your IBM Customer Support representative.

**Attention:** If you have multiple failed drives, do not attempt to revive them. Do not attempt to reseat the drives in their slots. Always save a Collect All Support Data log and contact your IBM Technical Support representative for guidance on recovery.

Perform these steps to revive a failed drive:

1. In the Physical tab of the DS Storage Manager (Subsystem Management) window, select the failed drive.
2. Click **Advanced** → **Recovery** → **Revive** → **Drive** as shown in Figure 14-129 on page 416.

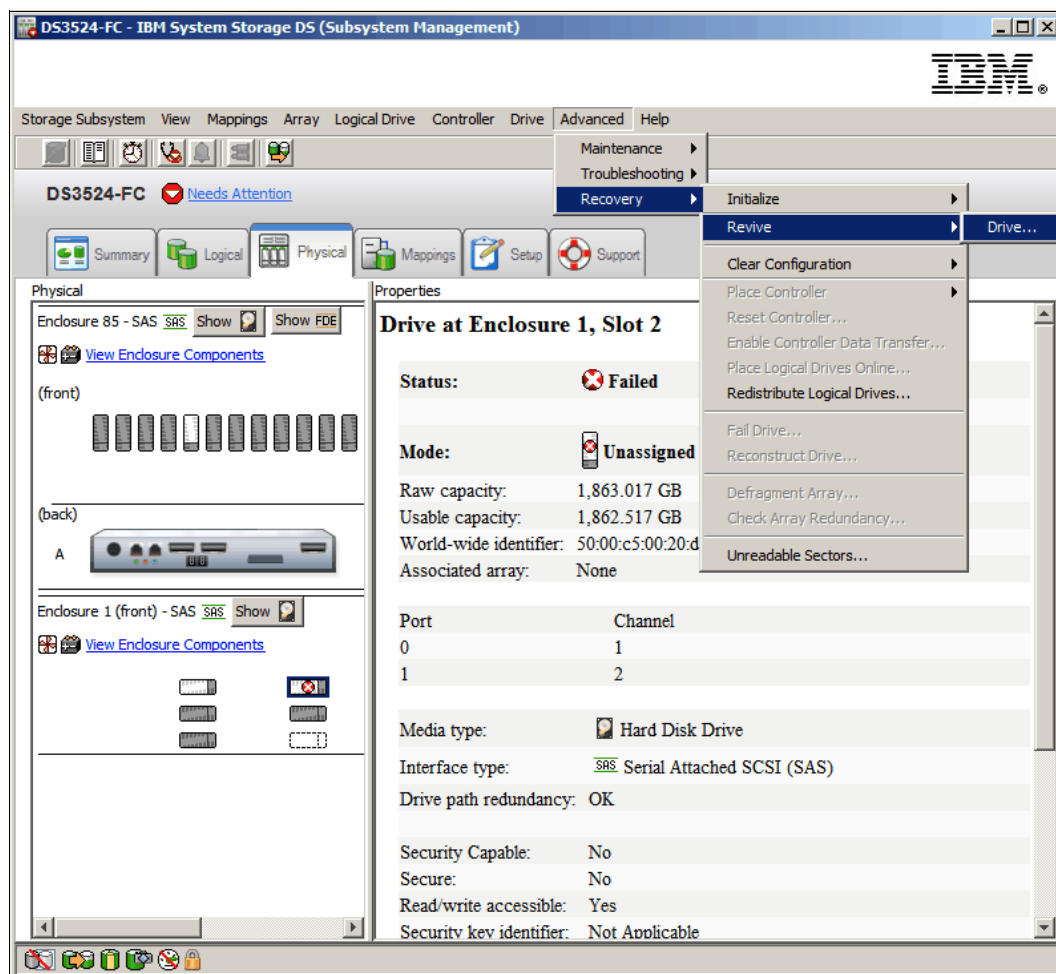


Figure 14-129 Revive drive

3. You will be asked to confirm that you want to proceed. Type yes in the text box and click **OK** to continue.

### 14.14.3 Recovery - Clear Configuration

The **Recovery** → **Clear Configuration** option from the SMW has the following submenu options:

#### Clear Storage Subsystem

Use the **Clear Storage System Configuration** option to permanently delete the current logical drive configuration. Using this function will cause these conditions to occur:

- ▶ All logical drives and arrays are deleted. This will destroy all data on these logical drives and arrays.
- ▶ The user-supplied DS3500 storage subsystem name is deleted.
- ▶ If you have password-protected your storage subsystem for destructive operations, that password protection is lost.

**Attention:** When you use this option, all data currently stored on the storage subsystem is lost. Do not attempt this operation unless you have a verified current backup of all data on the storage subsystem.

Perform these steps to clear the storage subsystem configuration:

1. In the DS Storage Manager (Subsystem Management) window, click **Advanced** → **Recovery** → **Clear Configuration** → **Storage Subsystem**.

You will be asked to confirm that you want to proceed as shown in Figure 14-130.

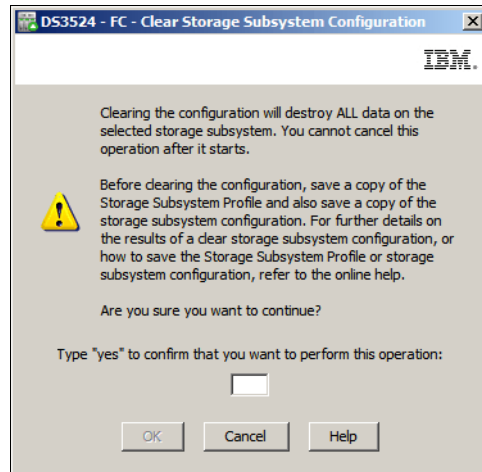


Figure 14-130 Clear Storage Subsystem Configuration

2. Type yes in the text box and click **OK** to clear the storage subsystem configuration.

## Clear Array Configuration

You can use the **Clear Array** option to permanently delete the current array configuration on the DS3500 storage subsystem.

**Attention:** All data currently stored on the array is lost when you use this option. Do not attempt this operation unless you have a verified current backup of all data on the logical drives.

Perform these steps to clear the storage subsystem array configuration:

1. In the DS Storage Manager (Subsystem Management) window, click **Advanced** → **Recovery** → **Clear Configuration** → **Array**.

You will be asked to confirm that you want to clear the array configuration as shown in Figure 14-131 on page 418.

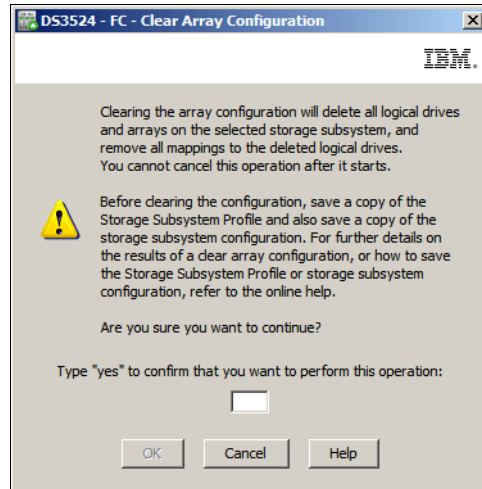


Figure 14-131 Clear Array Configuration

2. Type yes in the text box and click **OK** to clear the array configuration.

#### 14.14.4 Recovery - Place controller

You can use the **Advanced** → **Recovery** → **Place controller** submenu functions to set the selected controller state to online, offline, or service mode. In this section we describe how to use these options.

**Note:** These three functions should normally only be used under the guidance of your IBM technical support representative.

##### Place controller online

Placing a controller online sets it to Optimal status and places it in an active mode, available for I/O operations. Under normal operating conditions, both controllers will be online and in the optimal state. This function is only used if one of the two controllers is in either the offline state, or in service mode.

Perform the following steps to use this option:

1. Select the controller in the Physical tab of the DS Storage Manager (Subsystem Management) window and then click **Advanced** → **Recovery** → **Place controller** → **Online** as shown in Figure 14-132 on page 419.

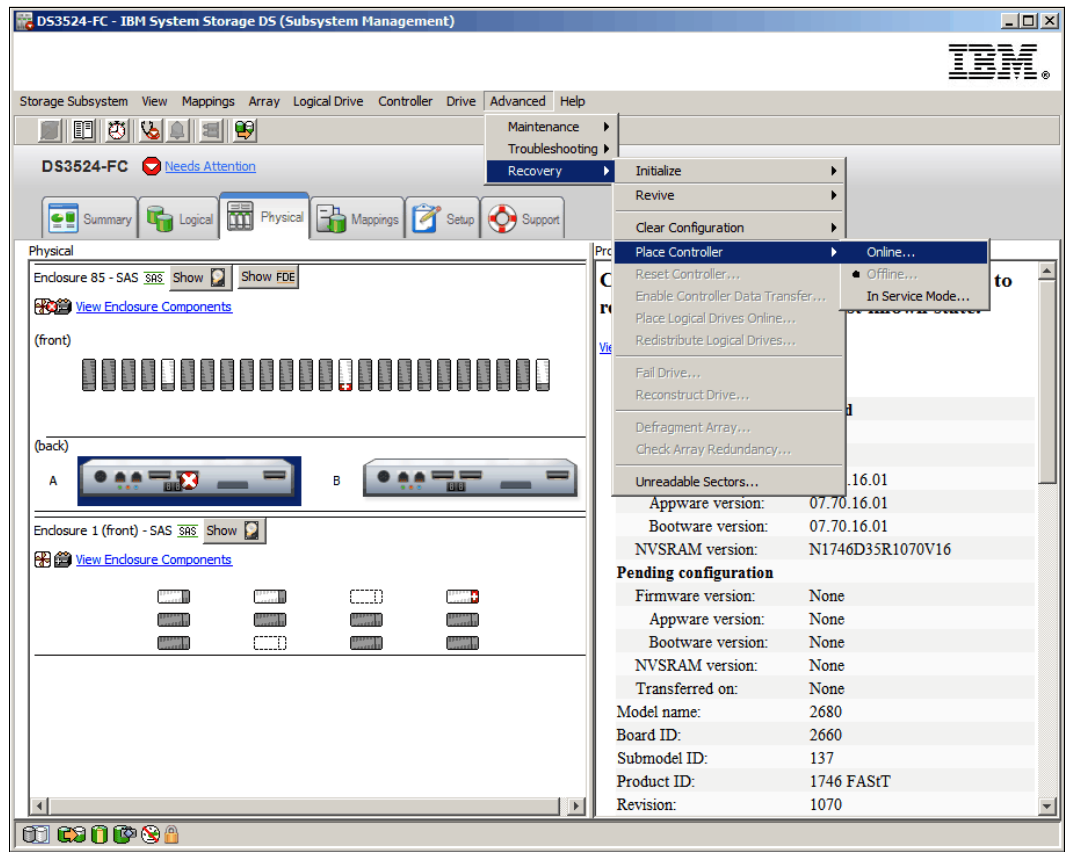


Figure 14-132 Place controller online

2. The window shown in Figure 14-133 will be displayed. Read the information and then click the **Yes** button to confirm that you want to place the controller online.

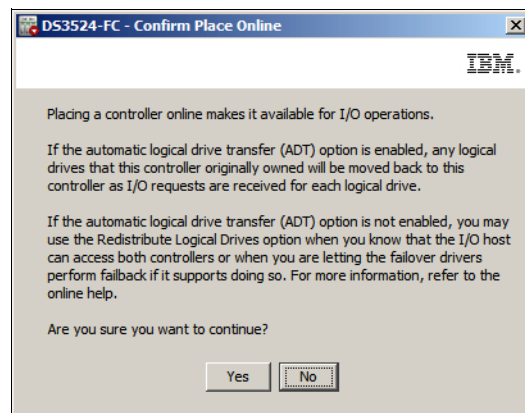


Figure 14-133 Confirm Place Online

## Place controller Offline

Placing a controller offline makes it unavailable for I/O operations and moves its arrays to the other controller if failover protection is enabled.

Perform the following steps to use this function:

1. Select the controller in the Physical tab of the DS Storage Manager (Subsystem Management) window and then click **Advanced** → **Recovery** → **Place controller** → **Offline** as shown in Figure 14-134.

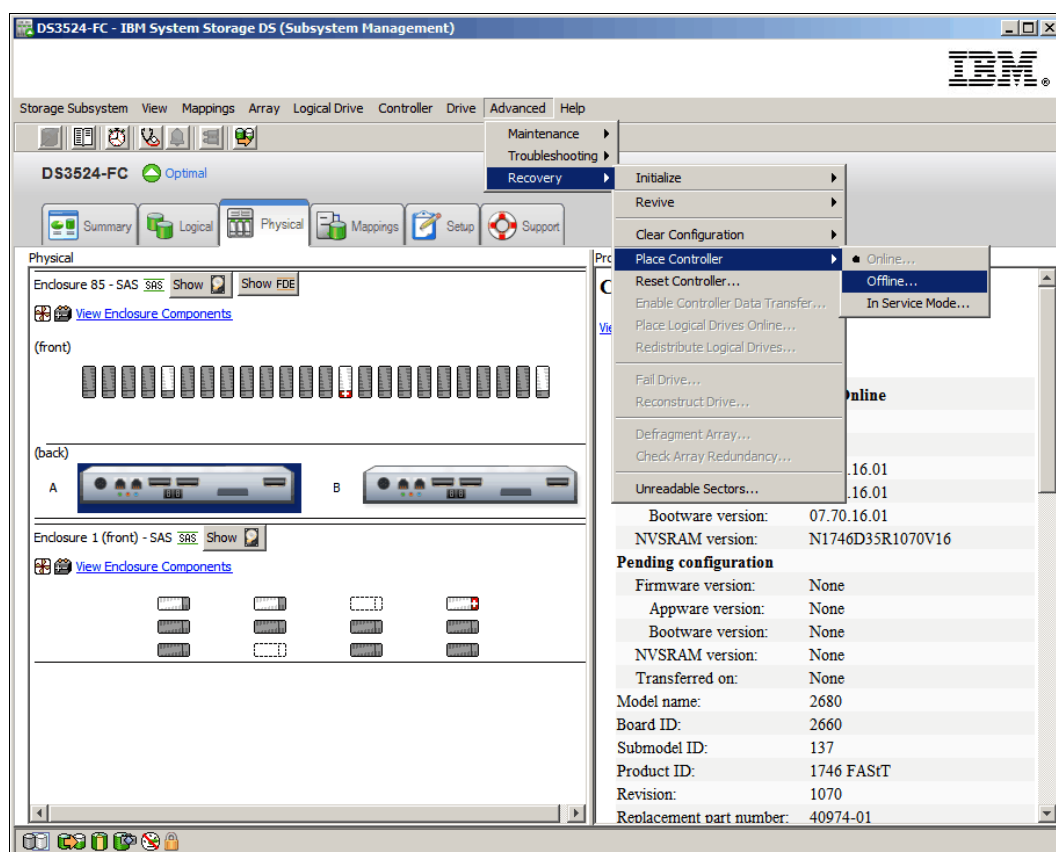


Figure 14-134 Place controller offline

2. The window shown in Figure 14-135 will be displayed. Read the information and then click the **Yes** button to confirm that you want to place the controller offline.

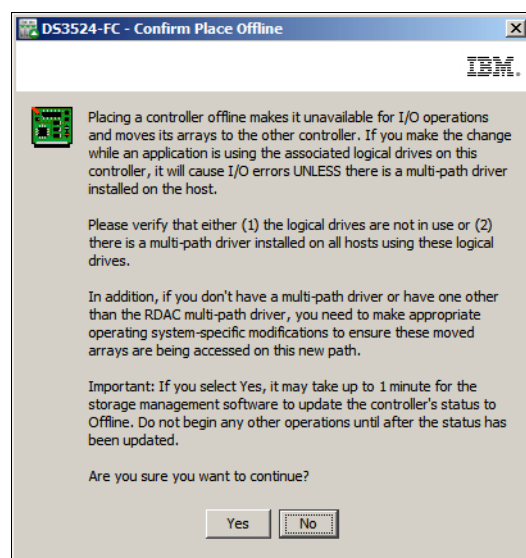


Figure 14-135 Confirm Place Offline

## Place controller in service mode

You can use the place controller in Service Mode option to safely run controller diagnostics without affecting data availability. These conditions apply when you use Service mode:

- ▶ When you place a controller into Service mode, all of the logical drives are forced to the active controller. This mode provides for continuous data access through the active controller. The preferred logical drive owner does not change. When the controller is removed from Service mode, the logical drives are redistributed back to their preferred owner.
- ▶ You can place a controller into Service mode only if two controllers are in the storage subsystem and they are both active.
- ▶ Service mode is a persistent state. If a controller is reset, goes through a power failure, or is removed and reinstalled, the controller remains in Service mode.
- ▶ You cannot make a controller that is in Service mode the owner of any logical drive.

Use the following procedure to set the controller in or out of service mode:

1. Select the controller in the Physical tab of the DS Storage Manager (Subsystem Management) window and then click **Advanced** → **Recovery** → **Place controller** → **In Service Mode** as shown in Figure 14-136.

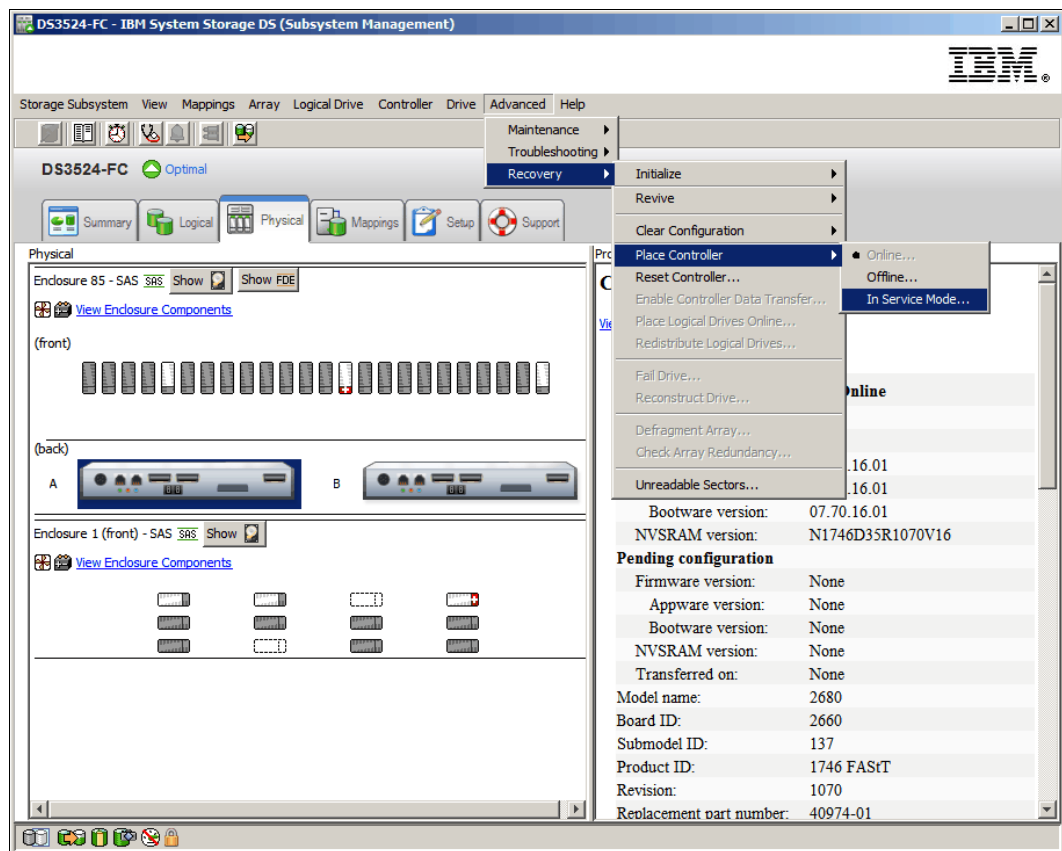


Figure 14-136 Place controller In Service Mode

2. The window shown in Figure 14-137 on page 422 will be displayed. Read the information and then click the **Yes** button to confirm that you want to place the controller in service mode.

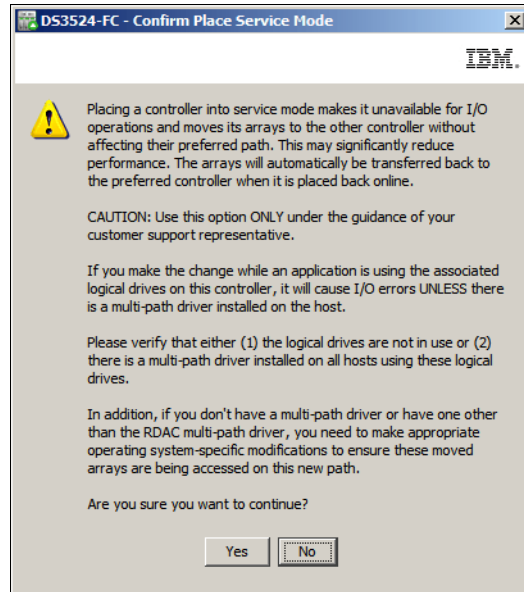


Figure 14-137 Confirm Place In Service Mode

3. A controller in Service mode can be set to either the online or offline state using the procedures described previously.

### 14.14.5 Recovery - Reset controller

This function causes the selected controller to perform a restart.

**Note:** Resetting a controller makes the controller temporarily unavailable for I/O until the reset is complete. If a host is using logical drives owned by the controller being reset, the I/O directed to the controller is rejected. Before resetting the controller, either make sure that the logical drives owned by the controller are not in use, or make sure that a multi-path driver is installed on all of the hosts using these logical drives. If a multipath driver is installed, the I/Os will be processed through the alternate controller.

Use this procedure to reset a controller:

1. Select the controller in the Physical tab of the DS Storage Manager (Subsystem Management) window and then click **Advanced** → **Recovery** → **Reset Controller** as shown in Figure 14-138 on page 423.



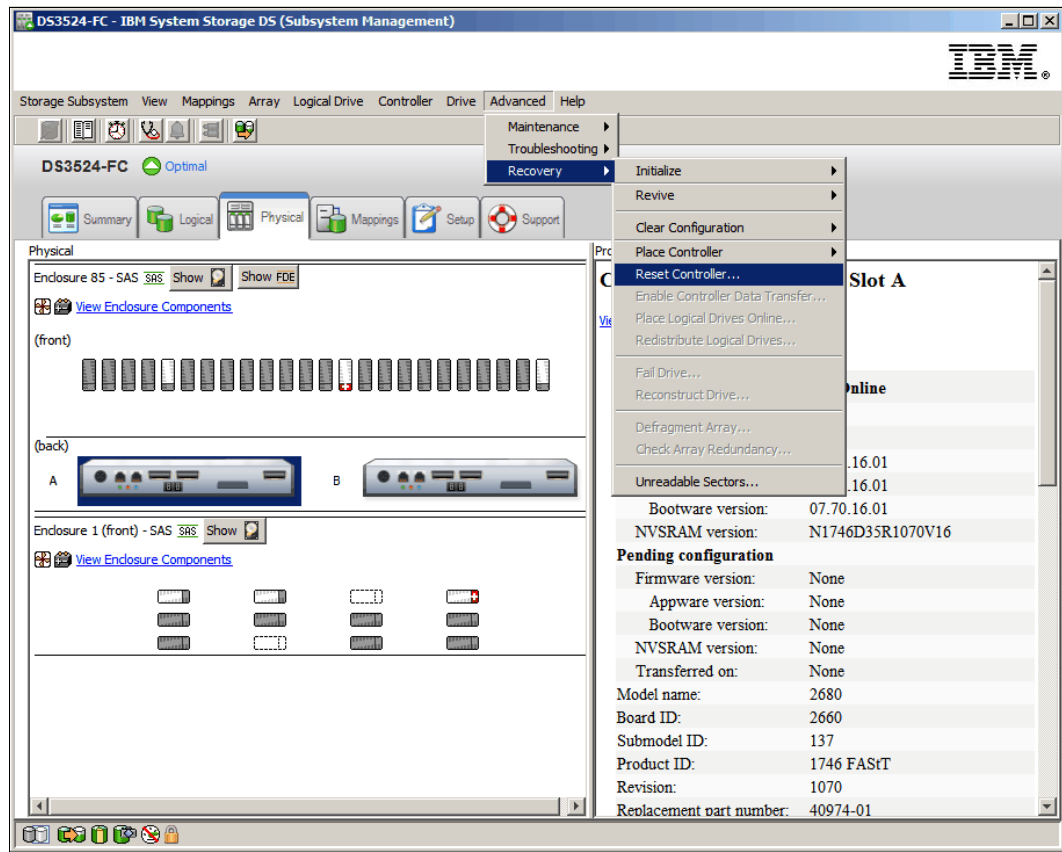


Figure 14-138 Reset Controller

2. The window shown in Figure 14-139 will be displayed. Read the information and click the **Yes** button to reset the selected controller.

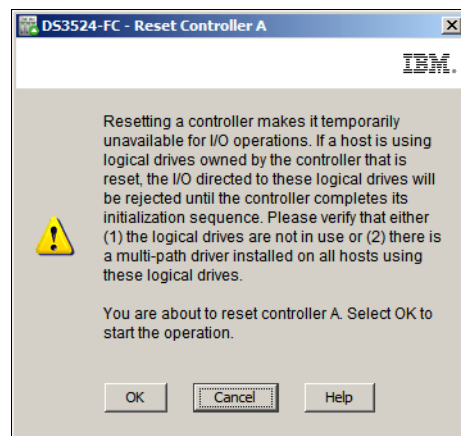


Figure 14-139 Confirm Reset Controller

## 14.14.6 Recovery - Enable controller data transfer

A controller can be in the disabled state due to diagnostics being run or firmware being downloaded. A controller in the disabled state no longer accepts I/O and is marked with a dashed red border in the Physical tab of the DS Storage Manager (Subsystem Management) window. If a multi-path driver is in use, I/O will be sent to the other controller path, initiating a

logical drive ownership transfer. When the operation that caused the controller to be disabled is finished, you can perform these steps to re-enable data transfer:

1. Select the disabled controller in the Physical tab of the DS Storage Manager (Subsystem Management) window and click **Advanced** → **Recovery** → **Enable Controller Data Transfer**.
2. The controller resumes normal operation and will accept host I/O.

### 14.14.7 Recovery - Place Logical drives online

If the DS3500 storage subsystem experiences an interruption of power, then cache data loss might occur. If the cache data cannot be recovered, the affected logical drives are automatically placed offline. You can use the Place Logical drives Online function to return logical drives that have been placed offline to the online status.

**Attention:** You should only use this function under the guidance of your IBM Technical Support representative.

To open the Place Logical drives Online dialog, select the **Advanced** → **Recovery** → **Place Logical Drives Online** menu command from the DS Storage Manager (Subsystem Management) window.

### 14.14.8 Recovery - Redistribute Logical Drives

You can use the Redistribute Logical Drives function to move logical drives back to their preferred owners. Usually, multipath drivers move logical drives from their non-preferred controller owners when a problem occurs along the preferred data path between the host and the storage subsystem.

Use the following procedure to restore the logical drives to their preferred controllers if the Recovery Guru reports that there are logical drives on a non-preferred path:

1. In the DS Storage Manager (Subsystem Management) window, click **Advanced** → **Recovery** → **Redistribute Logical Drives** as shown in Figure 14-140 on page 425.

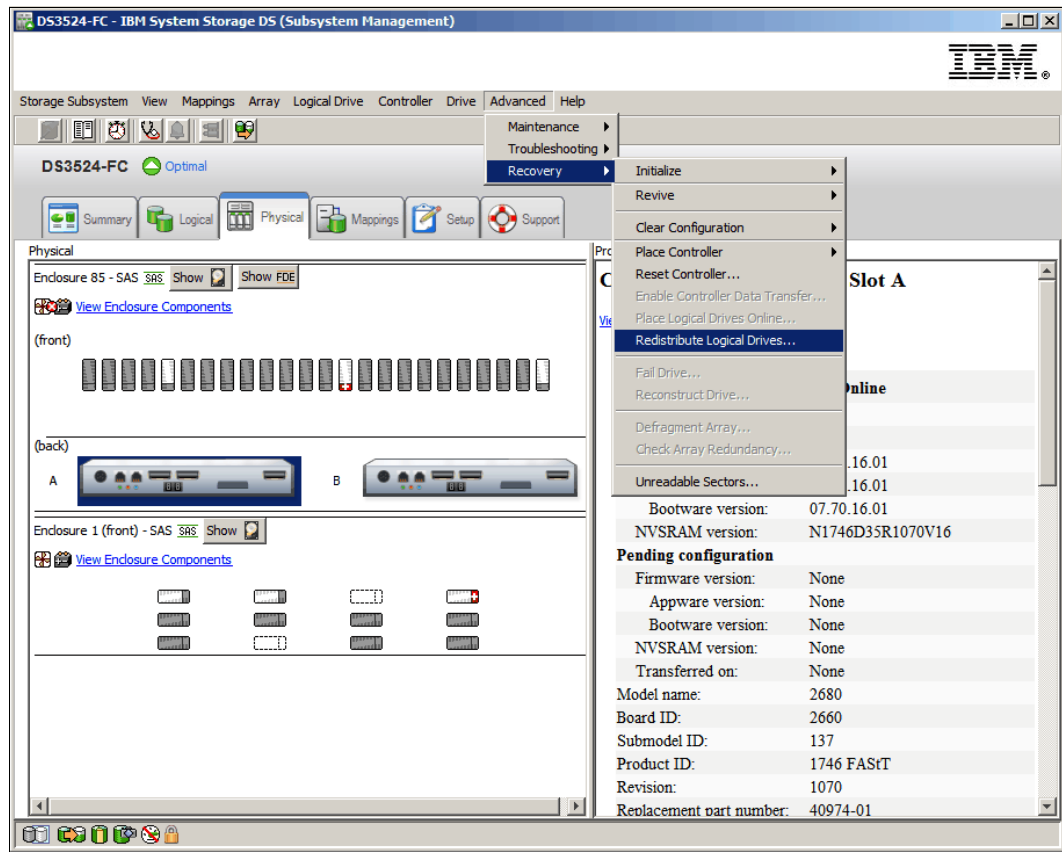


Figure 14-140 Redistribute Logical Drives

2. The window shown in Figure 14-141 will be displayed. Read the information and then click the **Yes** button to redistribute the logical drives to their preferred controller.

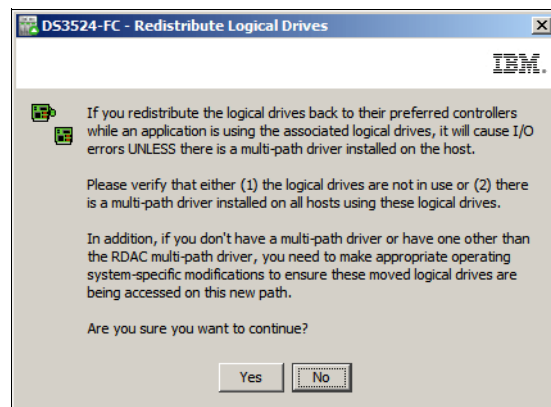


Figure 14-141 Redistribute Logical Drives confirmation

3. The progress of the logical drive redistribution will be displayed. The window shown in Figure 14-142 on page 426 will be displayed when all the logical drives have been redistributed to their preferred owners.




Figure 14-142 Redistribute Logical Drives - Progress

### 14.14.9 Recovery - Fail drive

You can use this function to set the state of a drive to failed.

**Note:** You should normally only use this function under the guidance of your IBM Technical Support representative. If you are failing a drive to replace it, unassign all hotspare drives first. This will prevent hotspare reconstruct and copyback operations. The hotspare drives can be reassigned after the new drives has been installed.

Use the following procedure to fail a drive:

1. Select the drive to be failed in the DS Storage Manager (Subsystem Management) window, and then click **Advanced** → **Recovery** → **Fail Drive** as shown in Figure 14-143 on page 427.

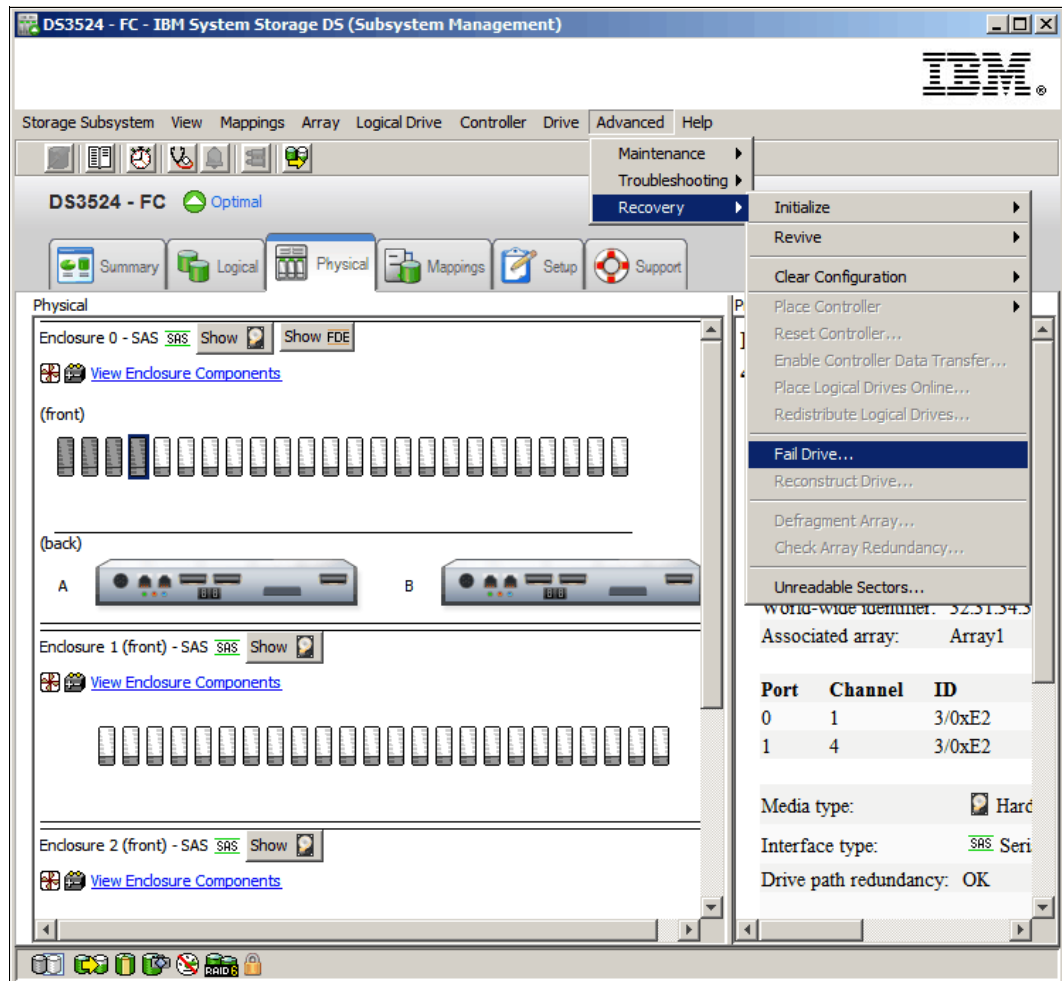


Figure 14-143 Fail Drive

You will be asked to confirm that you want to fail the selected drive as shown in Figure 14-144.




Figure 14-144 Confirm Fail Drive

2. Read the information and then click the **OK** button to fail the drive.

## 14.14.10 Recovery - reconstruct drive

You can use the Reconstruct Drive function to manually start reconstruction of a drive.

**Note:** This option is available only when a drive is assigned to a RAID Level 1/10, RAID Level 3, RAID Level 5, or RAID Level 6 array with either Failed status or Replaced status, and reconstruction does not automatically start. Use this option only under the guidance of your IBM Customer Support representative.

Use the following procedure to manually start reconstruction of a failed drive:

1. Select the failed drive in the Physical tab of the DS Storage Manager (Subsystem Management) window.
2. Click **Advanced** → **Recovery** → **Reconstruct Drive** as shown in Figure 14-145.




Figure 14-145 Reconstruct Drive

3. The window shown in Figure 14-146 on page 429 will be displayed. Read the information and then click the **Yes** button to start reconstruction of the drive.




Figure 14-146 Confirm Reconstruct

4. The drive will have a status of Replaced while the reconstruction is occurring. The array and logical drive(s) will have a status of Degraded until the reconstruction has completed. You can view the progress of a logical drive reconstruction by selecting the logical drive in the Logical tab of the DS Storage Manager (Subsystem Management) window as shown in Figure 14-147.

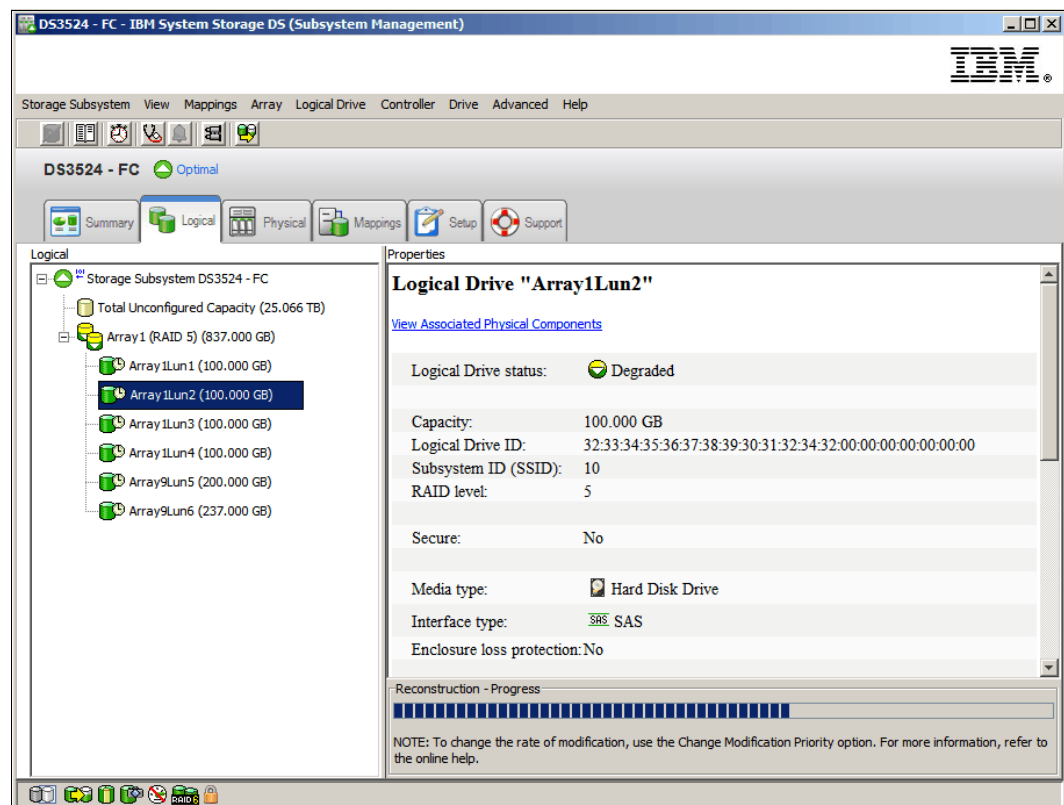


Figure 14-147 Reconstruction - Progress

### 14.14.11 Recovery - Defragment Array

You can use the Defragment Array function to consolidate all free capacity on a selected array. Using this option lets you create additional logical drives from the maximum amount of free capacity. A fragmented array can result from logical drive deletion or from not using all of the available free capacity in a Free Capacity node during logical drive creation.

Keep these guidelines in mind when you defragment an array:

- You cannot cancel the operation after it begins.

- ▶ Your data remains accessible during the defragment operation.
- ▶ To use this option, all of the logical drives in the array must be online and in Optimal status, and no logical drive modification operations, such as changing the segment size of a logical drive, can be in progress.

Use the following procedure to defragment an array:

1. In the Logical tab of the DS Storage Manager (Subsystem Management) window, select the array and click **Advanced** → **Recovery** → **Defragment Array** as shown in Figure 14-148.

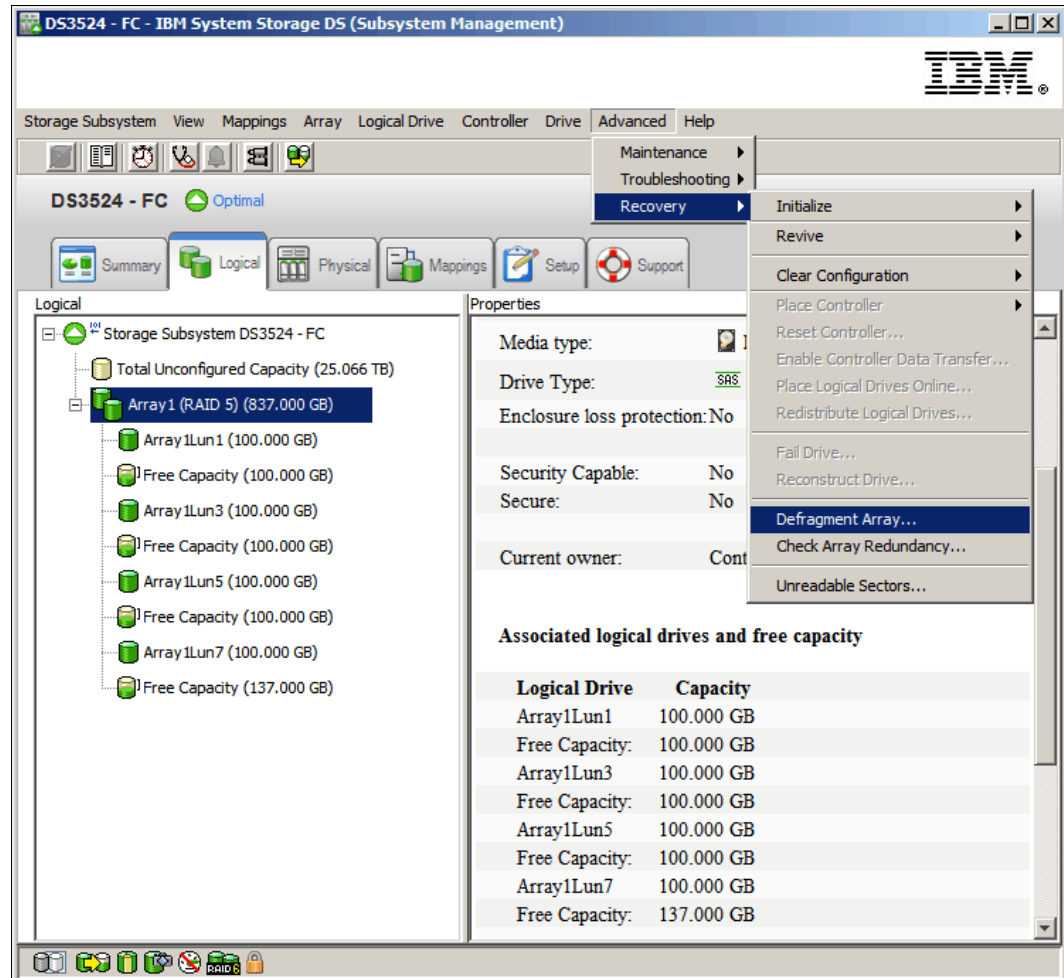


Figure 14-148 Defragment Array

2. You will be asked to confirm that you want to defragment the array as shown in Figure 14-149 on page 431.






Figure 14-149 Defragment Array - confirm

3. Click the **Yes** button. All of the free data blocks on the Free Capacity nodes are consolidated into one Free Capacity node that contains a larger set of contiguous data blocks. The Logical tab is updated to show the larger, single Free Capacity node as shown in Figure 14-150. You can select this Free Capacity node for logical drive creation.

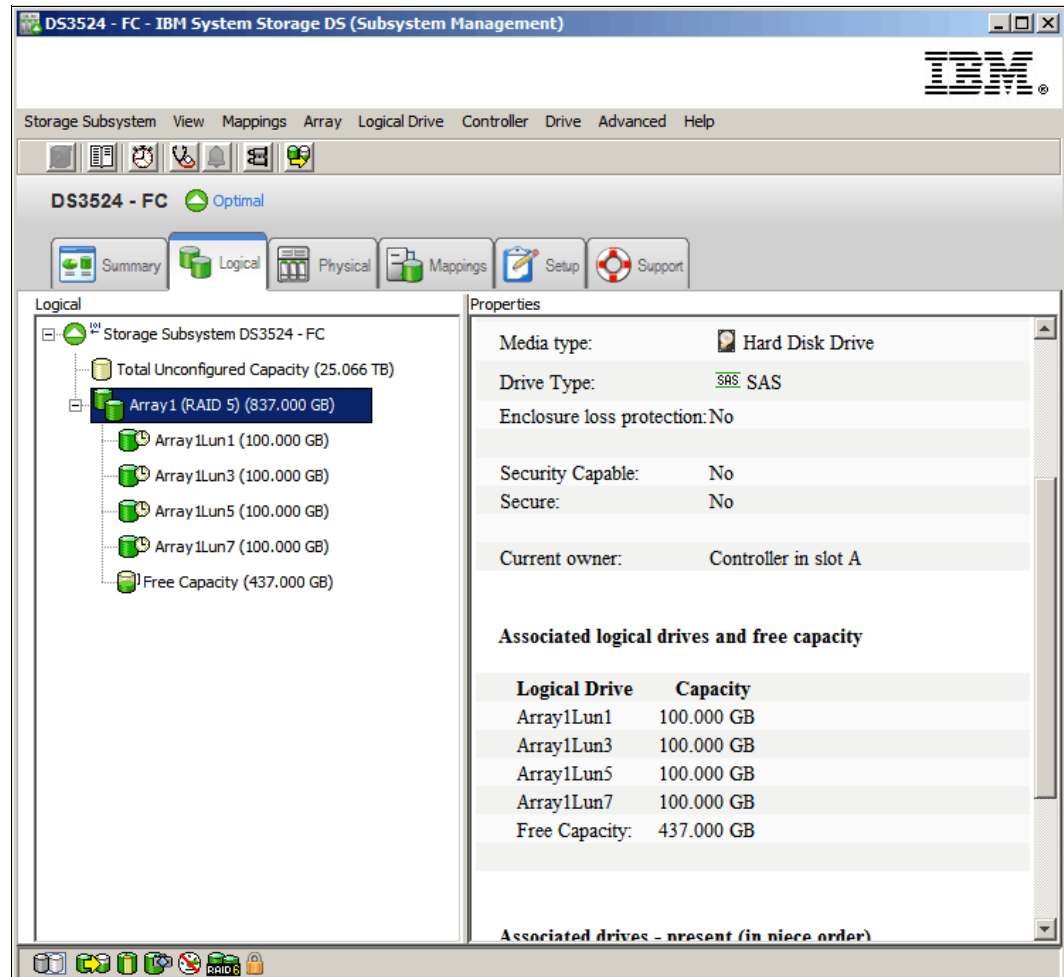


Figure 14-150 Defragment Array completed

### 14.14.12 Recovery - check array redundancy

You can use the Check Array Redundancy function to check the redundancy on a selected array.

**Note:** You should normally only use this option under the guidance of your IBM Customer Support representative.

Keep these important guidelines in mind before you check data redundancy on an array:

- ▶ You cannot use this option on RAID Level 0 arrays that have no redundancy.
- ▶ If you use this option on a RAID Level 1/10 array, the redundancy check compares the data on the mirrored drives.
- ▶ If you perform this operation on a RAID Level 3, RAID Level 5, or RAID Level 6 array, the redundancy check inspects the parity information that is striped across the drives.

To successfully perform this operation, all of these conditions must be true:

- ▶ The logical drives in the array must be in Optimal status
- ▶ The array must have no logical drive modification operations in progress.
- ▶ You can perform this operation only on one array at a time. However, you can perform a redundancy check on selected logical drives during a media scan operation. You can enable a media scan redundancy check on one or more logical drives in the storage subsystem.

Follow this procedure to check the redundancy of an array:

1. Select the array in the Logical tab of the DS Storage Manager (subsystem Management) window and then select the **Advanced** → **Recovery** → **Check Array Redundancy** menu command as shown in Figure 14-151 on page 433.

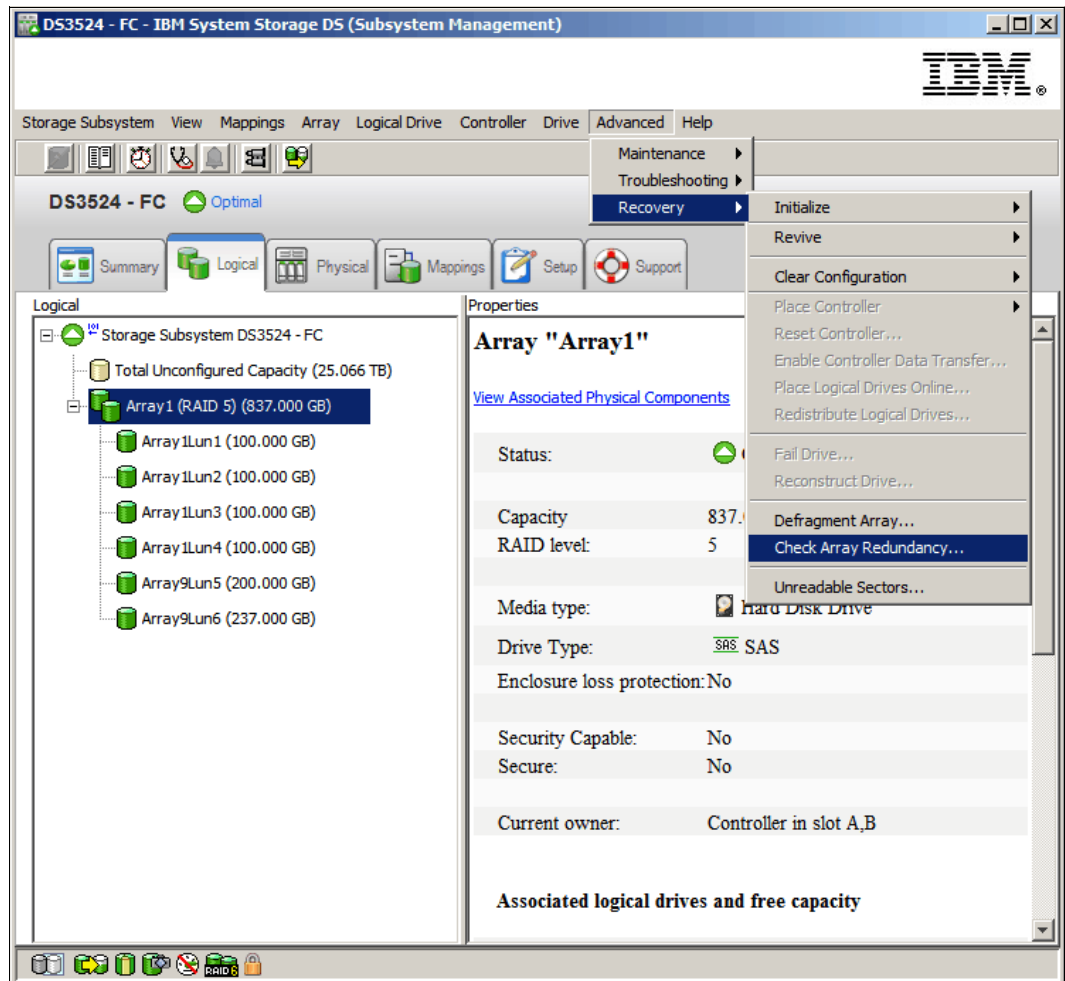


Figure 14-151 Check Array Redundancy

- The window shown in Figure 14-152 will be displayed.

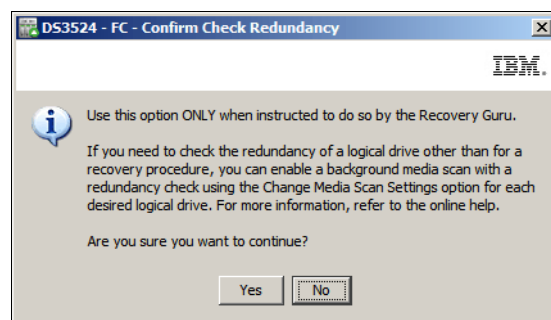


Figure 14-152 Confirm Check Redundancy

- Read the information and click the **Yes** button to display the Check Redundancy window as shown in Figure 14-153 on page 434.

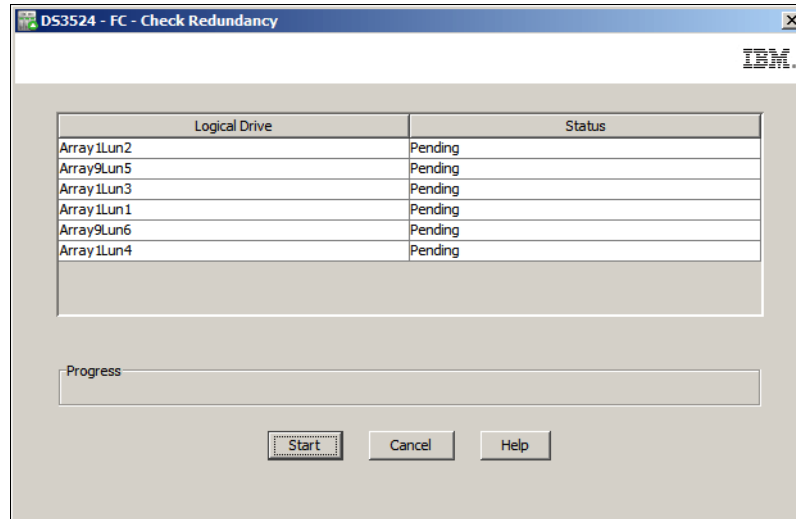


Figure 14-153 Check Redundancy

- Click the **Start** button and the check redundancy operation will begin. The logical drives in the array are sequentially scanned, starting from the top of the table in the logical drive dialog. These actions occur as each logical drive is scanned.

The progress bar shows the completion percentage of the redundancy check for a single logical drive, with 100 percent being complete. If an operation error occurs before the redundancy check completes, an error message appears.

- The window shown in Figure 14-154 is displayed after the last logical drive in the array has been checked, click **Done** to exit.

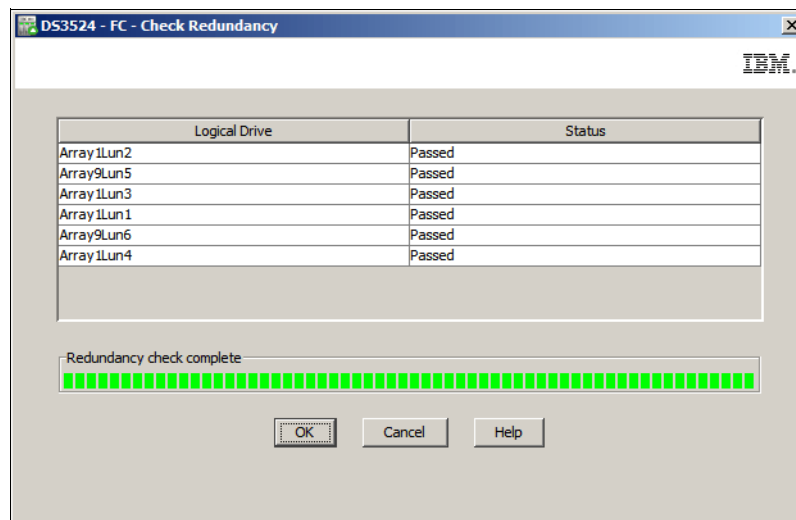


Figure 14-154 Check Redundancy completed

### 14.14.13 Recovery - Unreadable sectors

You can use the Unreadable Sectors Log View dialog to show a detailed list of all of the known unreadable sectors detected on the storage subsystem. This window lets you clear and save unreadable sectors. Information on unreadable sectors is maintained for both redundant (RAID Level 1/10, RAID Level 3/5, and RAID Level 6) logical drives and non-redundant (RAID Level 0) logical drives. The Unreadable Sectors Log can store up to

1000 unreadable sectors. When the Unreadable Sectors Log reaches 1000 entries, the following conditions apply:

- ▶ If new unreadable sectors are detected during reconstruction, the reconstruction fails, and no entry is logged.
- ▶ If new unreadable sectors are detected during I/O, the I/O fails, and no entry is logged.

## View the Unreadable Sectors Log

Use this procedure to view the Unreadable Sectors Log:

1. Click **Advanced** → **Recovery** → **Unreadable Sectors** to view the Unreadable Sectors Log as shown in Figure 14-155.




Figure 14-155 Unreadable Sectors Log

2. The Unreadable Sectors Log has two main panes:
  - The left pane lists identifying information about each logical drive that contains an unreadable sector.
  - The right pane lists details about each unreadable sector on the logical drive that is selected in the left pane.
3. The five buttons along the bottom have these functions:
  - **Select All:** Select at least one unreadable sector in the right window pane, and then click this button to select all unreadable sectors in the log.
  - **Clear:** Clicking this button will clear all unreadable sectors that have been selected.
  - **Save As:** Clicking this button will save the Unreadable Sectors log to a \*.txt file.

**Note:** The unreadable sectors log is also contained in the CASD support bundle.

- **Close:** This closes the Unreadable Sectors Log.
- **Help:** This opens the relevant section of the DS Storage Manager online help.

## 14.15 View Online Help

Click this link in the Support tab to access the Subsystem Management online help system, as shown in Figure 14-156. You will find a lot of useful information about your DS3500 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.

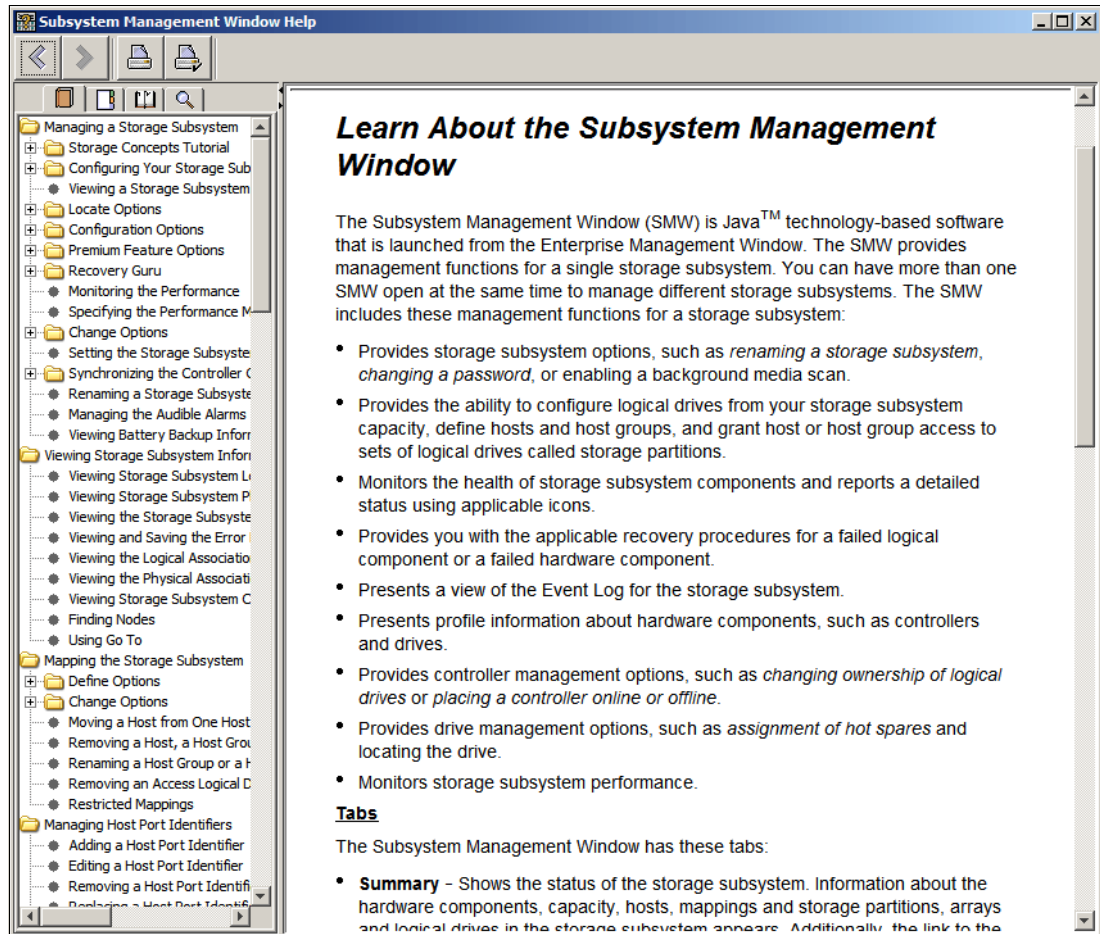


Figure 14-156 Subsystem Management Window Help

## 14.16 About IBM System Storage DS Storage Manager

Open this link in the Support tab to view the Storage Manager version and copyright information as shown in Figure 14-157 on page 437. It shows the current version of the DS Storage Manager (Subsystem Management) window.

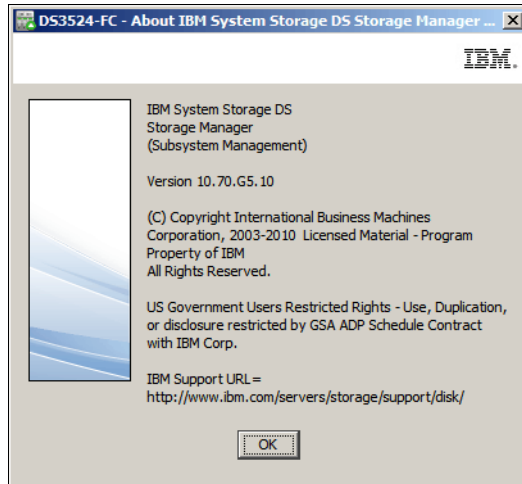


Figure 14-157 About IBM System Storage DS Storage Manager







## Disk Security with Full Disk Encryption drives

Disk Security is a new feature that is now also available for the IBM System Storage DS3500 storage subsystem that uses the newly available Full Disk Encryption (FDE) disk drives. It is supported by the latest level of the DS3500 firmware (Version 7.70) and IBM System Storage Manager V10.70 (Storage Manager). This chapter discusses how this new feature can add a greater level of security while your data resides on disk, what it does, the various components of the feature, and how to implement it.

The Disk Security premium feature requires security capable drives. A security capable drive encrypts data during writes, and decrypts data during reads. Each security capable drive has a unique drive encryption key. When a security capable drive has the security enabled, the drive requires the correct security key from the DS3500 for authentication before allowing reading or writing the data. This is managed on each of the DS3500 controllers by the IBM Disk Encryption Storage Manager, which is only available on RPQ basis at time writing this book. All of the drives in the DS3500 share the same security key, and the shared security key provides read and write access to the drives. The drive encryption key on each drive is used to encrypt the data.

## 15.1 The need for encryption

Data security breaches are becoming increasingly common, and the threat of unauthorized access to sensitive data and intellectual property is growing. Data security is now a common component of the corporate landscape, mostly driven by regulatory compliance, and security efforts can often be fragmented.

At some point, all drives are out of an organization's control and vulnerable to a breach. For example, with "re-purposing", decommissioning, or disposal of individual disks, common methods of security are often insufficient when:

- ▶ Deleted files can be reinstated.
- ▶ Drive disposal methods are subject to human error.
- ▶ Password protected data can still be accessed by a skilled individual.
- ▶ Disk reformatting erases data on disks but information can still be recovered.<sup>1</sup>

In each case, a risk is introduced where legible data might be recovered from disk. This can be made significantly more difficult if Disk Security on a DS3500 storage subsystem is employed.

**Important:** FDE and drive-level encryption technology is a new and additional level of data protection, but it does not replace the access controls and security processes that exist: rather, it complements them.

### 15.1.1 Encryption method used

The FDE drives have encryption hardware, and can perform symmetric encryption and decrypting of data at full disk speed with no impact on performance. The disk encryption hardware is used in conjunction with IBM Disk Encryption Storage Manager on the DS3500 storage subsystem. It uses asymmetric encryption to encrypt and decrypt the data key. IBM Disk Encryption Storage Manager will generate encryption and decryption keys that are used to lock each FDE drive.

Without these IBM Disk Encryption Storage Manager managed keys, the user (authorized or unauthorized) can no longer decrypt the data on disk.

**Important:** Should these keys and all copies be lost, then the encrypted data on disk cannot be decrypted and is therefore considered lost.

---

<sup>1</sup> Some utilities used to "erase" data from disks and arrays are not fully successful and therefore the data can still be recovered using forensic techniques

Figure 15-1 shows the relationship of the IBM Disk Encryption Manager and an individual FDE drive with encryption enabled.

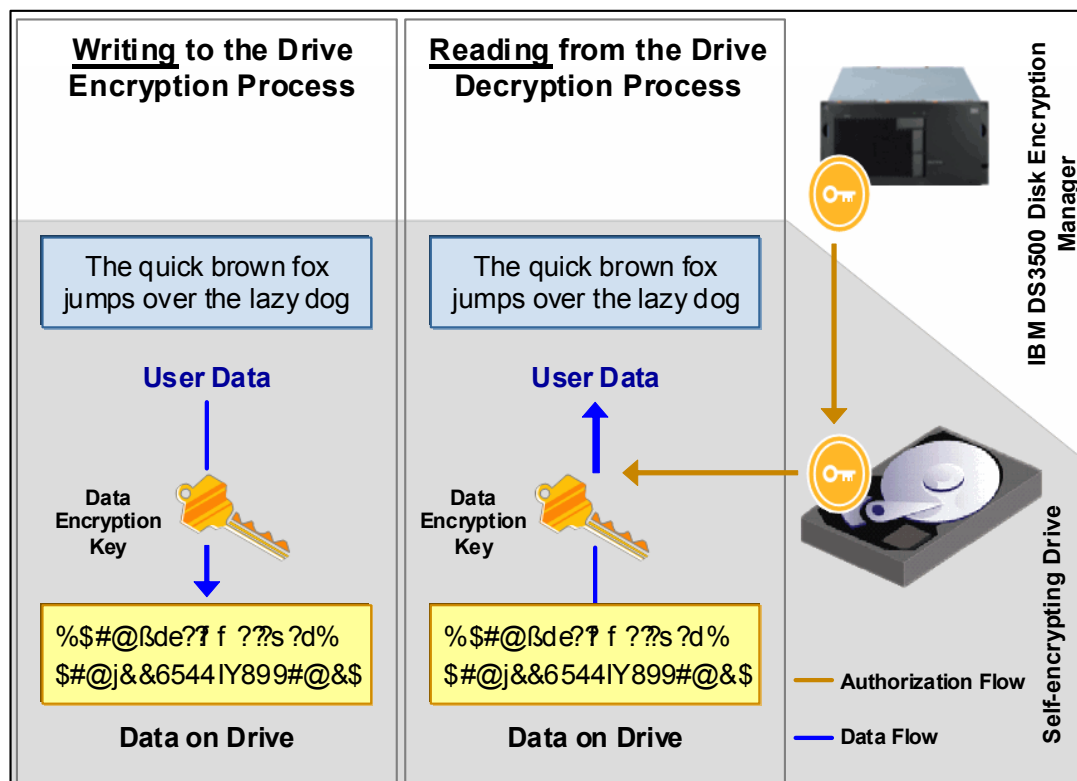


Figure 15-1 DS3500 Disk Encryption Manager using self encrypting FDE drives

With this relationship, the correct keys, and authentication, the FDE drive will encrypt data written and decrypt data read from it. But if the disk is removed and someone attempts to read the data on the disk, as shown in Figure 15-2 on page 442, the user will not have the appropriate authorizations, and so the data cannot be read from or written to the drive without authenticating with the DS3500 Disk Encryption Manager, which will unlock the drive.

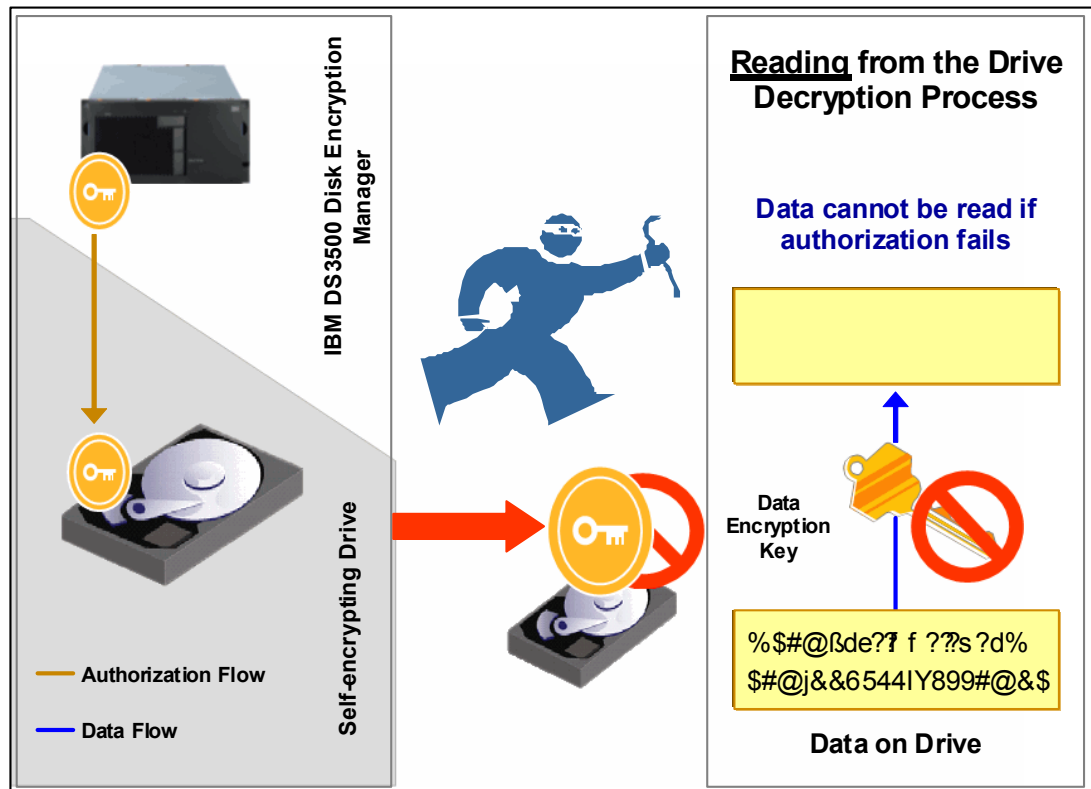


Figure 15-2 Unauthorized access to the drive results in the data remaining encrypted

## 15.2 Disk Security components

There are a number of new components to this new feature that are detailed in this section. All of these features are managed by the Storage Manager (V10.70.x and higher).

### 15.2.1 DS3500 Disk Encryption Manager

The Disk Encryption Manager on the DS3500 system maintains and controls the key linkage and communication with FDE drives. It will be included with the firmware and Storage Manager. It has the following functions:

- ▶ Provides all the management tools necessary to quickly and simply enable and secure FDE drives.
- ▶ Establishes and manages a single authorization scheme for all the FDE drives in a DS3500 storage subsystem.
  - Places FDE drives in a secured state.
  - Defines secure arrays.
  - Supports the decommissioning or re-purposing of drives with Instant Secure Erase.

With this function you can record both the security key ID, pass phrase, and the secure file location in a safe place.

- ▶ Using the FDE drive, it generates and encrypts a security key:
  - Creates a unique security key ID that is paired with the security key.
  - Adds a randomly generated number.

- The security key ID is saved. This folder location will be needed whenever a security operation requires the key ID (for example, when a drive powers up).
- Creates a backup of the security key and the security key identifier.
- A secure backup is provided in which the security key and the security key identifier are encrypted utilizing a user-selected pass phrase.

## 15.2.2 Full Data Encryption (FDE) disks

FDE drives are required to enable Disk Security. You must use Serial Attached SCSI (SAS) disks with a speed of up to 15,000 rpm. These disks include:

- ▶ ST9146852SS 6 Gbps SAS 2.0 2.5-inch 147 GB 15k
- ▶ ST9300503SS 6 Gbps SAS 2.0 2.5-inch 300 GB 10k
- ▶ ST3300657SS 6 Gbps SAS 2.0 3.5-inch 300 GB 15k
- ▶ ST3600957SS 6 Gbps SAS 2.0 3.5-inch 600 GB 15k

## 15.2.3 Premium feature license

The DS3500 requires that the Full Disk Encryption premium feature be installed and enabled for Disk Security to function. See 3.4, “Planning for premium features” on page 70 for details about this topic.

## 15.2.4 Keys

There are two types of keys that are used with Drive Security and FDE drives:

- ▶ The *encryption key* is generated by the drive and never leaves the drive, so it always stays secure. It is stored in encrypted form and performs symmetric encryption and decryption of data at full disk speed with no impact on disk performance. Each FDE drive uses its own unique encryption key that is generated when the disk is manufactured and regenerated when required by the storage administrator using the DS3500 Disk Encryption Manager.
- ▶ The *lock key* or *security key* is a 32 byte random number that authenticates the drive with the DS3500 Disk Encryption Manager using asymmetric encryption for authentication. When the FDE drive is secure “enabled”, it must authenticate with the Disk Encryption Manager or it will not return any data and remains locked. After the drive has been authenticated, access to the drive operates like any other disk drive. One security key is created for all FDE drives on the DS3500 storage subsystem, where it is generated, encrypted, and hidden in the subsystem (NVSRAM). The authentication typically only occurs after the FDE has powered up, where it will be in a “locked” state.

If the lock key is not initially established between the DS3500 Disk Encryption Manager and the disk, then the disk is considered unlocked with access unlimited as per a non-FDE drive.

## 15.2.5 Security key identifier

For additional protection, the security key that is used to unlock FDE drives is not visible to the user. The security key identifier is used to refer to a security key instead. You can see the security key identifier during operations that involve the drive security key backup file, such as creating or changing the security key. The security key identifier is stored in a special area of the disk; it can always be read from the disk and can be written to the disk only if security has been enabled and the drive is unlocked.

The security key identifier field in the FDE Drive Properties window, shown in Figure 15-3, includes a random number that is generated by the controller when you create or change the security key. One security key is created for all FDE drives on the storage subsystem.

Note that the Security Capable and Secure fields in the Drive Properties window show whether the drive is secure capable, and whether it is in Secure (Yes) or Unsecured (No) state. The example shows that the drive is both capable (FDE) and enabled.

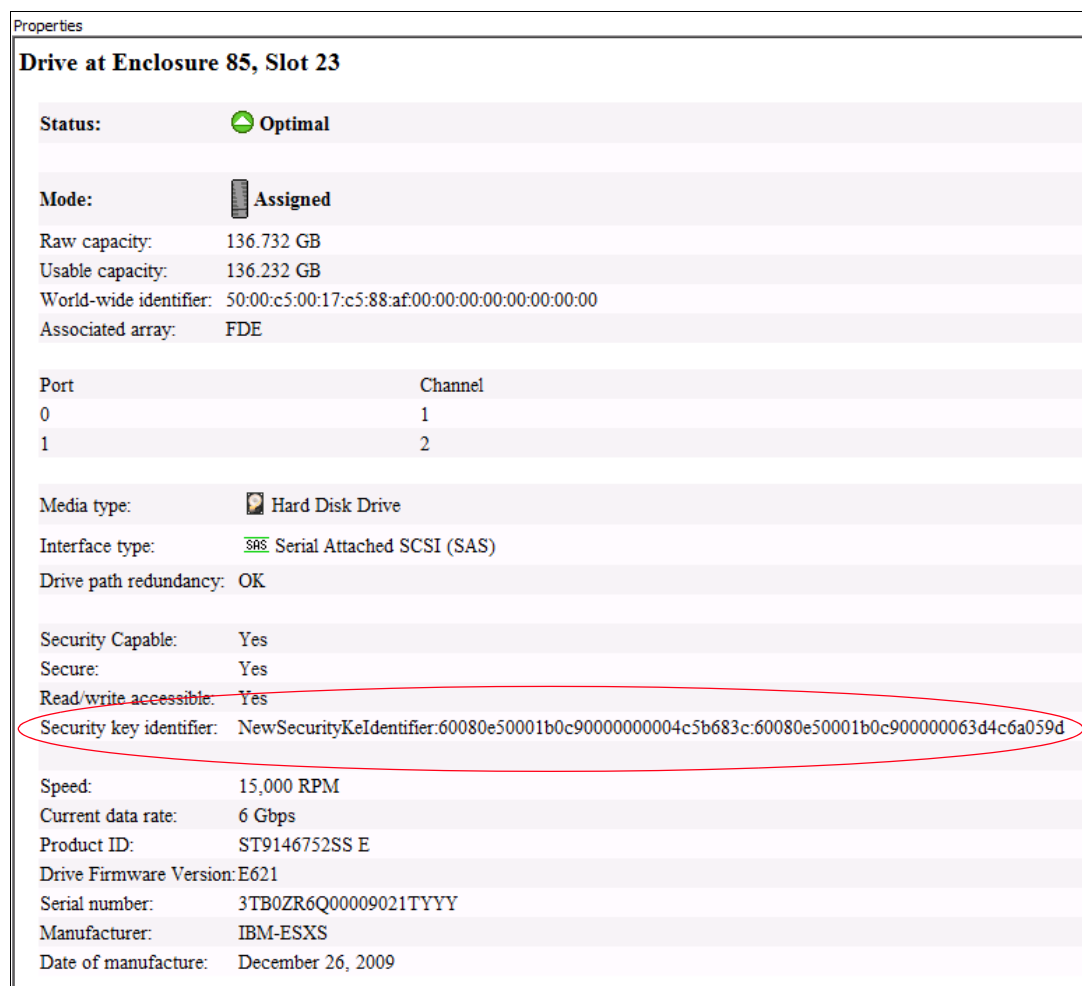


Figure 15-3 FDE drive properties showing security ID and status

## 15.2.6 Passwords

For Disk Security to be enabled, the DS3500 has to have the administration pass phrase or password set. The password must be “strong” (not easy to guess). A check is made on the password and if the system does not consider it to be strong enough when you log in or are prompted for the password, the message **Strong password requirement not met** will appear as shown in Figure 15-9 on page 449. The description of the warning will include suggestions about how the password can be made stronger.

The security key and the security key identifier are encrypted using a different password or pass phrase when the key is created or changed (see 15.3.2, “Secure key creation” on page 448 and 15.4.1, “Changing the security key” on page 456). The array then returns a file that is called a *blob*, or key backup. If the array needs that key later, you give the blob and pass phrase to the GUI, which sends it down to the array where the original key is decrypted.

The user-specified alphanumeric character string is not stored anywhere on the DS3500 or in the security key backup file.

## 15.3 Setting up and enabling a secure disk

This section shows a step-by-step process to create a key and file on the IBM Disk Encryption Storage Manager of the DS3500. It will then show how to enable a previously configured array that has FDE drives. Figure 15-4 shows a system without FDE enabled

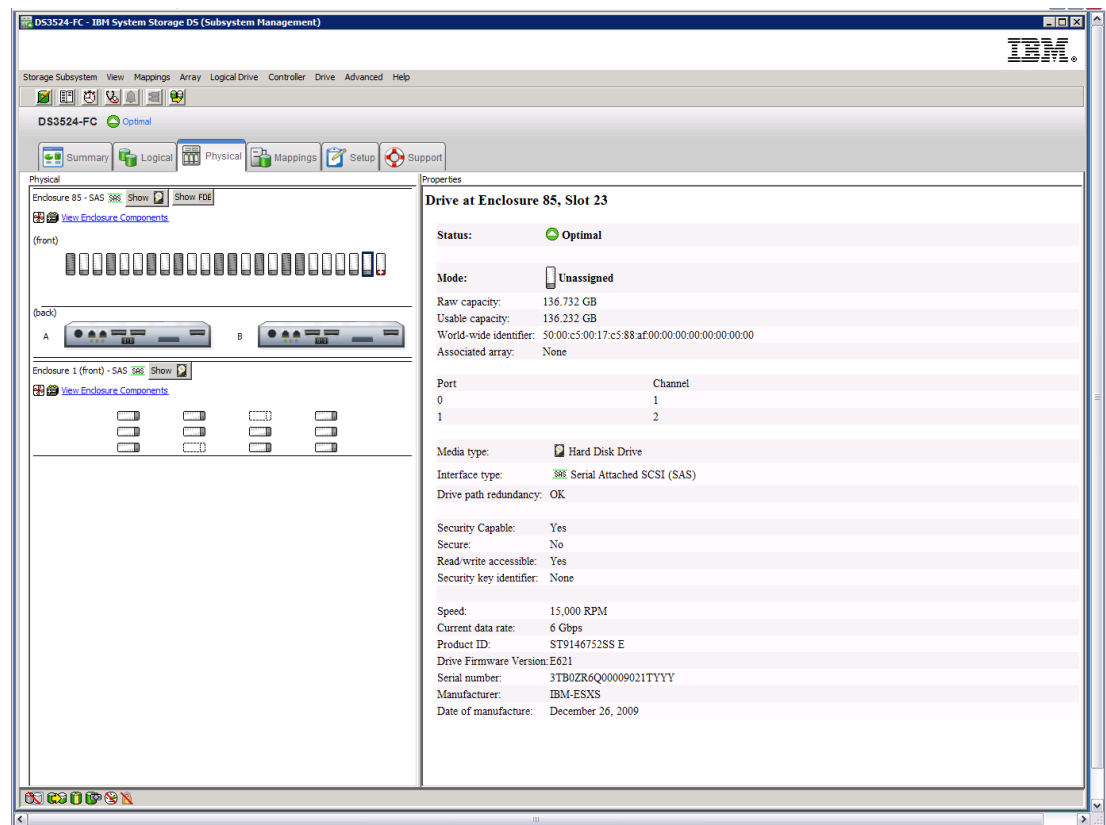


Figure 15-4 System without FDE enabled

### 15.3.1 FDE and premium feature check

There are a number of checks to make prior to key creation. First, you must check that the premium feature key has been applied to the system. To do this task, from the Storage Manager window, select **Storage Subsystem** → **Premium Features**.

Figure 15-5 on page 446 shows that the Drive Security premium feature key has been obtained and successfully installed. This premium feature key is installed the same as any other premium feature key (see “Premium Features” on page 194 for details about how to install premium feature keys).




Figure 15-5 The correct premium feature has been obtained and installed

Now you will have the possibility to identify the FDE drives by pressing the **Show FDE** button. Figure 15-6 on page 447 shows an example



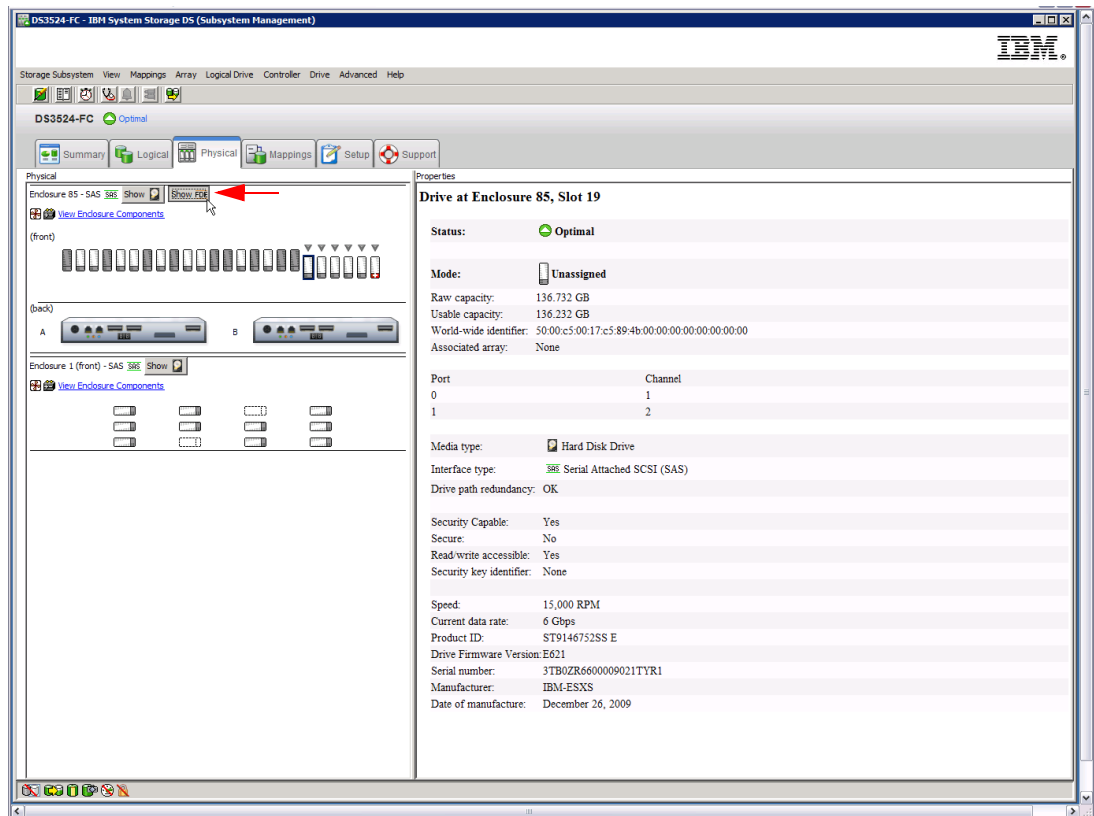


Figure 15-6 Selection of the FDE Drives

The next check is to ensure that the physical disks are FDE drives and are secure capable. In Figure 15-7 on page 448, you can see that an array has been created where disk security is not enabled, as indicated by the unlocked padlock. If you click the array and then right-click and select **View Associated Physical Components**, you can view the disks. If you click the **Show FDE** button, it will display all the FDE drives and confirm that all the drives in the array are secure capable.

The array is already being used, has several logical drives configured, and consists of all FDE drives. Therefore, it can be enabled.

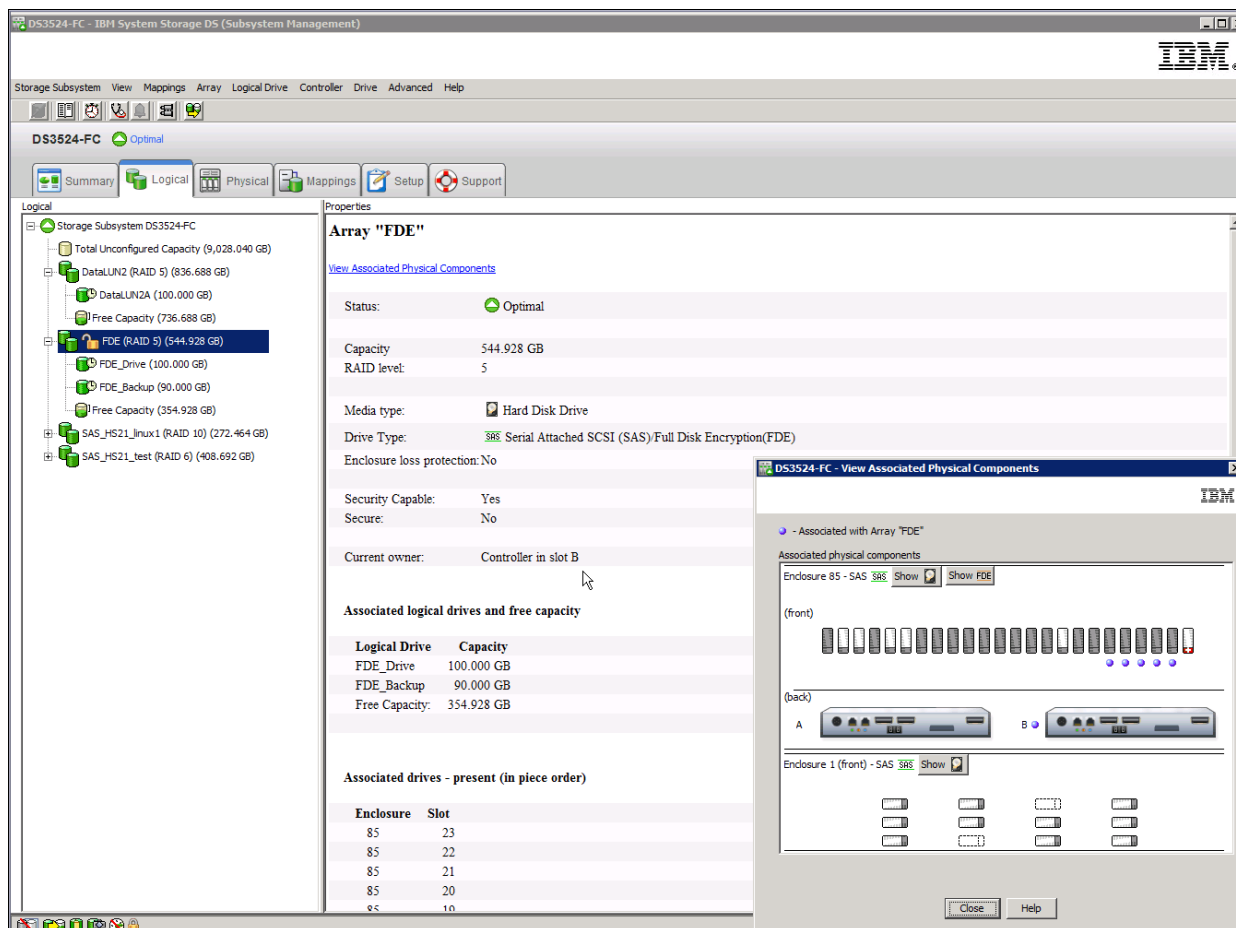


Figure 15-7 Array created with FDE drives with Disk Security disabled

## 15.3.2 Secure key creation

To create the secure key, follow these steps:

1. Click **Storage Subsystem** → **Drive Security** → **Create Security Key** in the top left corner of the IBM System Storage DS Storage Manager Subsystem Management window as shown in Figure 15-8 on page 449.

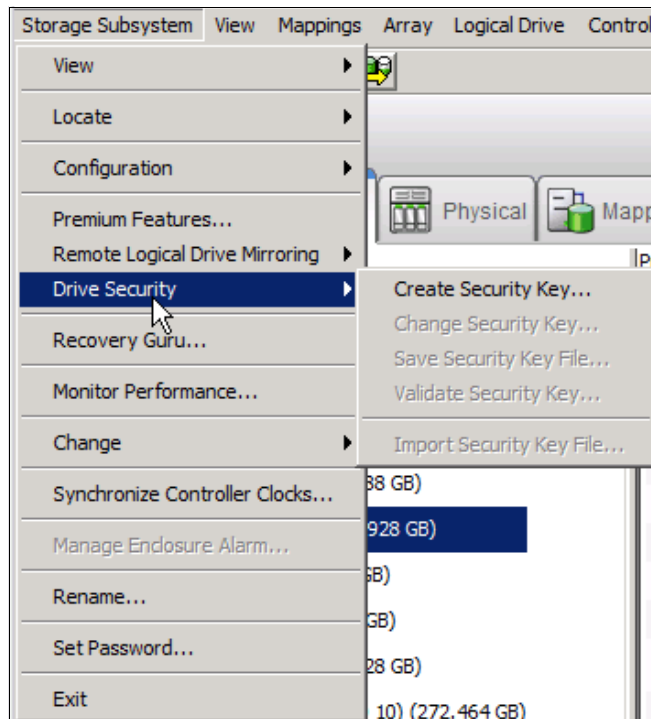


Figure 15-8 Create Security Key

You must have your DS3500 subsystem password set to proceed, and you might see the message shown in Figure 15-9 if the existing password be considered is too weak.

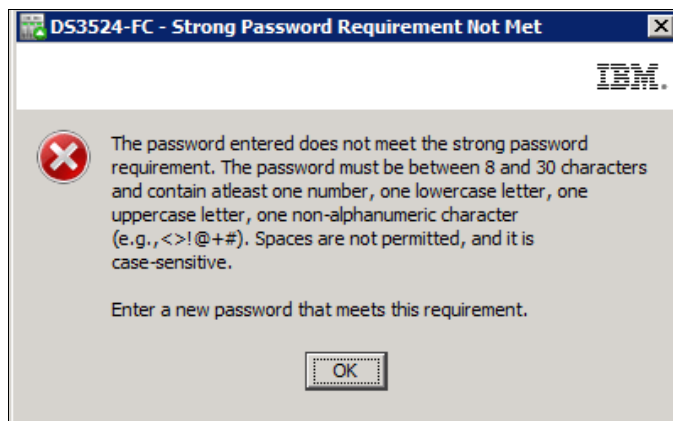


Figure 15-9 Warning regarding the strong password requirements

2. Enter the password as shown in Figure 15-10 on page 450.




Figure 15-10 Enter your password to continue with the creation of the security key

3. Fill in the empty form as shown in Figure 15-11 on page 451

DS3524-FC - Create Security Key

IBM.

Security key identifier

The security key identifier is paired with the security key to help you remember which key to use for secured operations. The security key identifier can be left blank or you may type up to 189 alphanumeric characters. The system will add the storage subsystem world-wide identifier and a randomly generated number to what you type in the field. You will have the opportunity to record the final security key identifier later.

Security key identifier:

File save location

File name:

C:\Program Files (x86)\IBM\_DS\client\data\securityLockKey\

Browse...

Pass phrase

The pass phrase is required to perform security operations. It must be between eight and thirty-two characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g., < > @ +). Spaces are not permitted, and it is case-sensitive.

Pass phrase:

Confirm pass phrase:

Create Key Cancel Help

Figure 15-11 Empty form for security key generation

4. The window shown in Figure 15-12 on page 452 opens. Complete the fields, which includes a file name and pass phrase. The Security key identifier is optional. If nothing is entered the system will add the storage subsystem world wide identifier and a randomly generated number.

**DS3524-FC - Create Security Key**

IBM

**Security key identifier**

The security key identifier is paired with the security key to help you remember which key to use for secured operations. The security key identifier can be left blank or you may type up to 189 alphanumeric characters. The system will add the storage subsystem world-wide identifier and a randomly generated number to what you type in the field. You will have the opportunity to record the final security key identifier later.

Security key identifier:

**File save location**

File name:

**Pass phrase**

The pass phrase is required to perform security operations. It must be between eight and thirty-two characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g., < > @ +). Spaces are not permitted, and it is case-sensitive.

Pass phrase:

Confirm pass phrase:

Figure 15-12 Requirements displayed for the security key creation

5. Click **Storage Subsystem** → **Drive Security** → **Save Security Key File** as shown in Figure 15-13 on page 453. The key location default is in the user's local PC directory. Copy the key and keep it in a safe location.

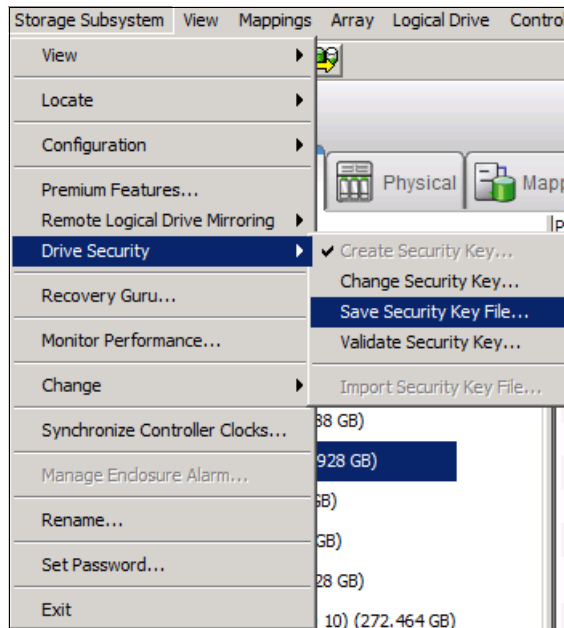


Figure 15-13 Save Security Key file

**Tip:** The general rule is to store the security key file with your key management policies along with the pass phrase. It is important to record and remember where this file is stored because the security key file is required along with pass phrase when a drive is moved from one storage subsystem to another, or when both controllers in a storage subsystem are replaced at the same time.

6. When the process is complete, the key is created and located as specified; also, as shown in Figure 15-14 on page 454 the security identifier is displayed. There are three items that must be kept secure to manage any changes to the FDE drives when they are encryption enabled:
  - Security key file created by the process
  - Security key identifier created by the process
  - Password used



Figure 15-14 Security key creation process complete

The key authorizations now generated are synchronized between both controllers in the DS3500 storage subsystem. With these authorizations in place, arrays on the FDE drives in the storage subsystem can be secured.

### 15.3.3 Enable Disk Security on array

Now that the keys have been created, Disk Security can now be enabled on the array. Use the following steps:

1. Right-click the array and select **Secure Drives** as shown in Figure 15-15 on page 455.



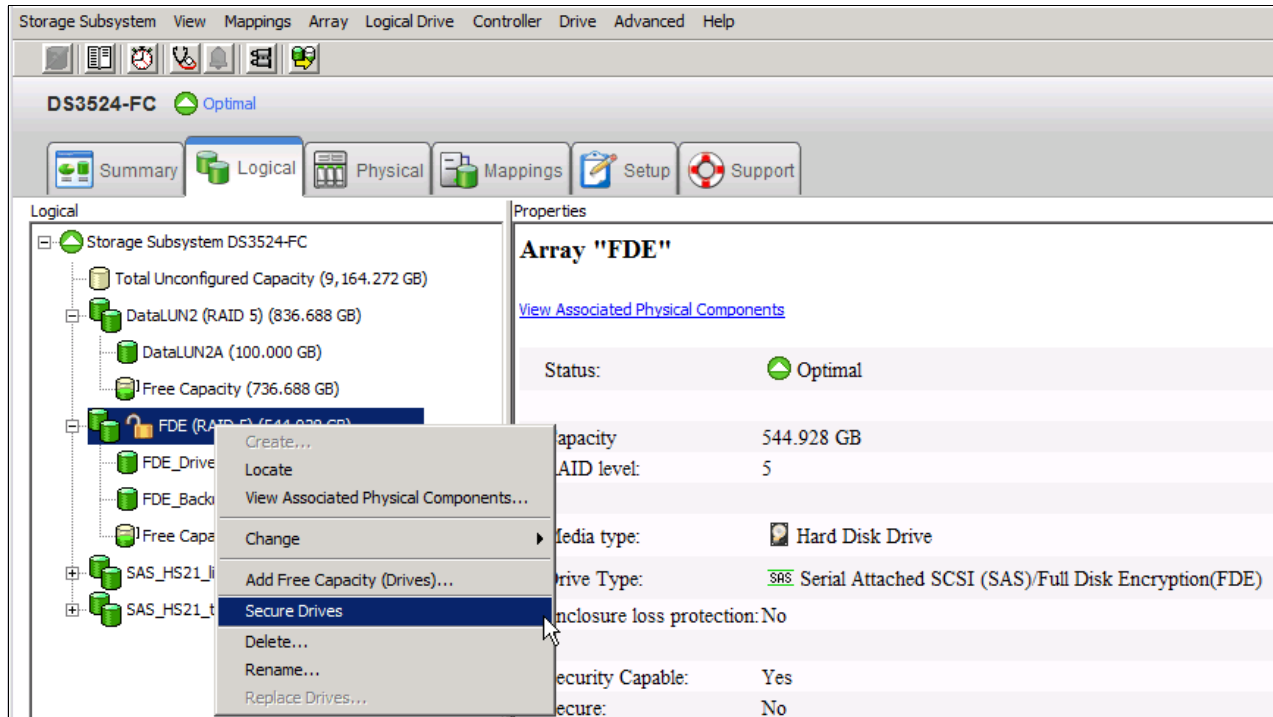


Figure 15-15 Secure all drives in the array

2. You will then be prompted to confirm the Secure Drives on the array, as shown in Figure 15-16.




Figure 15-16 Confirm Array Drive Security

The array is now secured, as indicated by the padlock in a locked position, as shown in Figure 15-17 on page 456.




Figure 15-17 Array is now secured with Disk Security enabled

## 15.4 Additional secure disk functions

In the following sections, we discuss the following functions:

- ▶ Changing the security key
- ▶ Saving the security key file
- ▶ Secure disk erase
- ▶ FDE drive status
- ▶ Hot spare drives

### 15.4.1 Changing the security key

The security key can be changed if the details of the existing key be corrupted or the pass phrase forgotten, provided that there are no outstanding Secure Disk communications between the FDE drives and Disk Encryption Manager (for example, if a disk is in a “locked” state). Because the disk encryption key never leaves the disk, you might want to periodically change the encryption key, the way a user might periodically change the administrative password to an operating system. This depends on the organization’s security guidelines.

The process to change the security key is similar to that of creating it initially. To change the key, perform the following steps:

1. Click **Storage Subsystem** → **Drive Security** → **Change Security Key**. The confirmation window to change the security opens as shown in Figure 15-18.

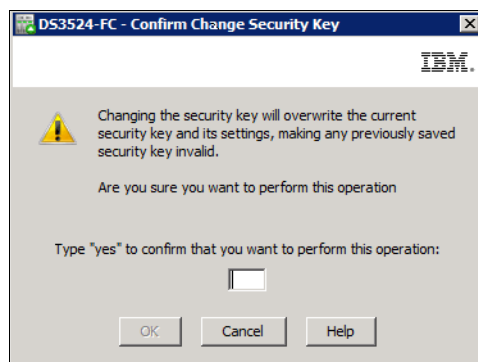


Figure 15-18 Confirm change security key

2. The window shown in Figure 15-19 opens and you are prompted to add a security identifier (optional), a location to store the key file in, and a pass phrase.

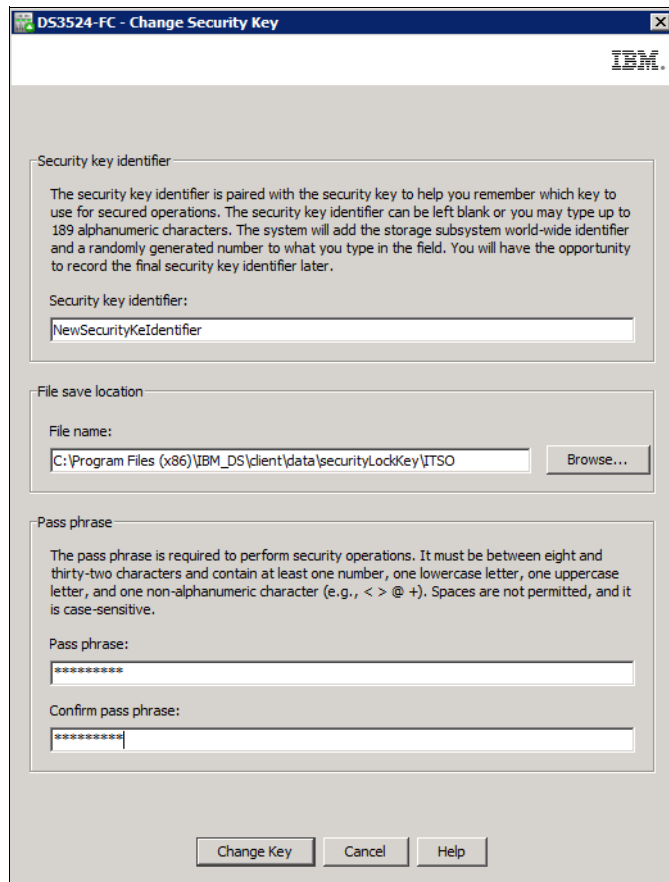


Figure 15-19 Change Security Key options

The new security key is generated by the controller firmware and is hidden in the storage subsystem. The new security key replaces the previous key that was used to unlock the security-enabled FDE drives in the storage subsystem. The controller negotiates with all of the security-enabled FDE drives for the new key.

The original security key is also stored in the storage subsystem for protection in case something prevents the controllers from completing the negotiation of the new security key with the security-enabled FDE drives (for example, loss of storage subsystem power during the key change process). If this happens, you must change the security key so that only one version of the security key is used to unlock all drives in a storage subsystem. The original key is stored in the storage subsystem only. It cannot be changed directly or exported to a security key backup file.

3. When the security key has been successfully changed, a confirmation window opens, as shown in Figure 15-20 on page 458, where the new key file location and security key identifier are shown. Click **OK**.



Figure 15-20 Change Security Key Complete confirmation window

### 15.4.2 Save security key file

This action will save a backup of the security key file and will require the original pass phrase to copy it. It can therefore also be used to verify that the pass phrase stored is correct. To save the security key file, perform the following steps:

1. Click **Storage Subsystem** → **Drive Security** → **Save Security Key File**.
2. You will be prompted for the location to store the file and the pass phrase used to create or change the existing security key file, as shown in Figure 15-21. The DS3500 Disk Encryption Manager uses the pass phrase to encrypt the security key before it exports the security key to the security key backup file.

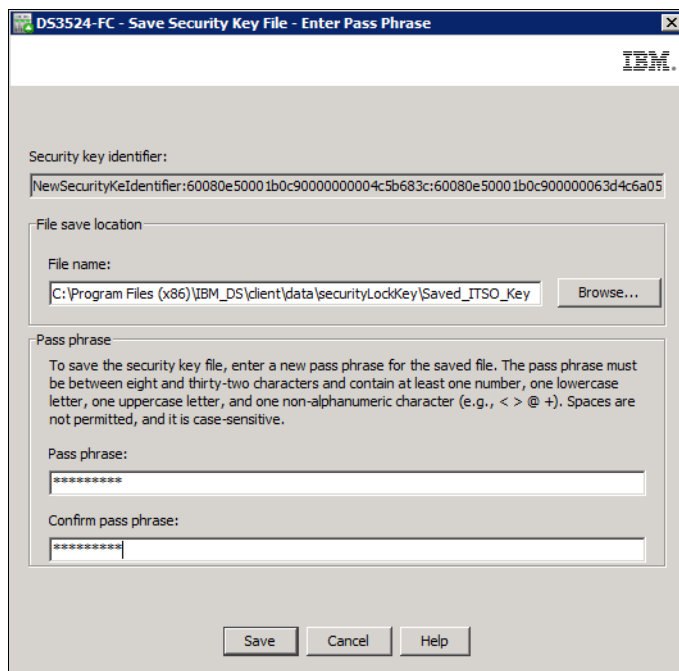


Figure 15-21 Save Security Key File window

### 15.4.3 Secure erase

Secure erase provides a higher level of data erasure than traditional methods. When you initiate secure erase with the DS3500 Storage Manager, a command is sent to the FDE drive to perform a “cryptographic erase”. This erases the existing data encryption key and then generates a new encryption key inside the drive, making it impossible to decrypt the data. Drive security becomes disabled and must be re-enabled if it is required again. To perform a Secure erase of the drive, right-click the drive and select **Secure Erase** as shown in Figure 15-22.

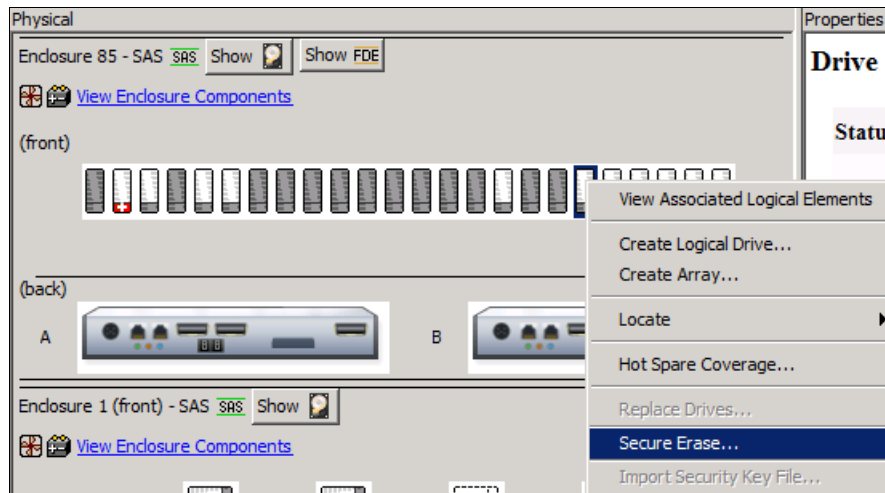


Figure 15-22 Secure Erase of a FDE drive

The secure erase process is shown in Figure 15-23.

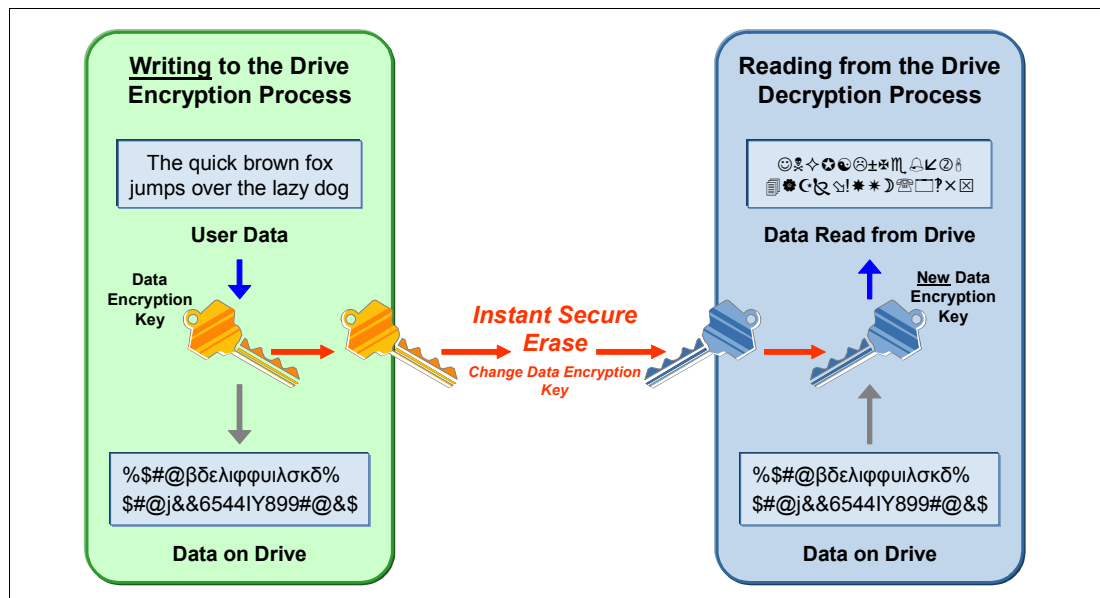


Figure 15-23 Secure erase process

**Warning:** All data on the disk will be permanently and irrevocably erased when the secure erase operation is completed for a security-enabled FDE drive. Do not perform this action unless you are sure that you want to erase the data.

Secure erase can only be performed on drives that are not allocated to an array. The process is also referred to as re-provisioning, where:

- ▶ The FDE drive becomes fully reusable.
- ▶ The drive can be reused in secure or non-secure applications.
- ▶ Previous data and keys are not accessible.
- ▶ It executes in less than a second.
- ▶ It returns the drive to the original factory state.

#### 15.4.4 FDE drive status

The FDE drives have a status indicating whether the disk can be accessed. The statuses are:

- ▶ Locked
  - The drive is security capable.
  - The drive is security enabled.
  - The lock key has not been supplied to the drive.
  - Data cannot be read or written from drive.
- ▶ Unlocked
  - The drive is security capable.
  - The drive is security enabled.
  - The lock key has been supplied to the drive.
  - Data can be read or written from drive.

The locked state will rarely be seen, usually only when the array containing the disks has been moved to another DS3500 or controllers have been replaced. The drive also becomes locked whenever the disk is powered down. The drive will remain unlocked during firmware upgrades or while other components are being replaced. When the drive is powered on, the status will be locked. If it detects a security key identifier, then it will remain locked until it has successfully authenticated with the DS3500 Disk Encryption Manager.

#### 15.4.5 Hot spare drive

If a disk drive fails in the DS3500 storage subsystem, the controller uses redundant data to reconstruct the data on the failed drive on a global hot-spare drive. The global hot-spare drive is automatically substituted for the failed drive without intervention. When the failed drive is eventually replaced, the data from the hot-spare drive is copied back to the replacement drive.

Hot-spare drives must meet the array hot-spare requirements. The following drive types are required for hot-spare drives when secure-capable arrays are configured. If a drive does fail, the Storage Manager automatically determines which hot-spare drive to substitute according to the type of the failed drive:

- ▶ For an array that has secured FDE drives, the hot-spare drive must be an unsecured FDE drive of the same or greater capacity. After the unsecured FDE hot-spare drive is used as a spare for a failed drive in the secured RAID array, it is security enabled.
- ▶ For an array that has FDE drives that are not secured, the hot-spare drive can be either an unsecured FDE drive or a non-FDE drive.
- ▶ An unconfigured secured FDE drive cannot be used as a global hot-spare drive. If a global hot spare is a secured FDE drive, it can be used as a spare drive only in secured arrays.

## 15.5 Migrating secure disk arrays

This section will detail how to migrate an array with Disk Security enabled to another DS3500. The process consists of exporting the array on the source DS3500 and importing the array on the target DS3500 after the disk have been physically moved. User data remains intact on the disks because configuration metadata is stored on every drive in the DS3500. The export process is the same as an ordinary array, but there are extra steps to be carried out for the locked drives of the secure array when importing.

### 15.5.1 Planning checklist

This list is displayed in a window of the migration wizard (Figure 15-25 on page 463). It is important to follow this list when planning an export and import of an array.

On the source DS3500, ensure the following action have been completed:

- ▶ Save the DS3500 configuration by clicking **Storage Subsystem** → **Configuration** → **Save**.
- ▶ Back up all data at the host level on logical drives of the array.
- ▶ Ensure that all I/O has been stopped to the array, and unmap any logical drives that are mapped on the array, ensuring that each host no longer requires the disk presented to it by the array to be migrated.
- ▶ Locate the array disks on the DS3500, noting the enclosure ID and location slot of each disk member.
- ▶ Save the security key file as discussed in 15.4.2, “Save security key file” on page 458.
- ▶ Obtain blank canisters or new drives to be inserted in the DS3500 when the drives are removed.

On the target DS3500, ensure the following action have been completed:

- ▶ Verify that there are available drive slots to host the disk drives.
- ▶ Check that the drives that are being moved are supported. The Drive Security premium feature should be already enabled as shown in Figure 15-5 on page 446.
- ▶ Verify that the firmware is the same on the target DS3500 as the source DS3500, and is up to date.
- ▶ Unlock drives before importing the array.

### 15.5.2 Export the array

To export the array, perform the following steps:

1. From the Storage Manager window, in the Logical tab, select the array that you want to export and then select **Advanced** → **Maintenance** → **Export Array**.
2. The Export Array wizard will guide you through the export. Click **Next**.




Figure 15-24 Export Array wizard

3. The window with the preparation checklist opens as shown in Figure 15-25 on page 463. Check that no planning steps have been missed. All check boxes must be checked before you can proceed.



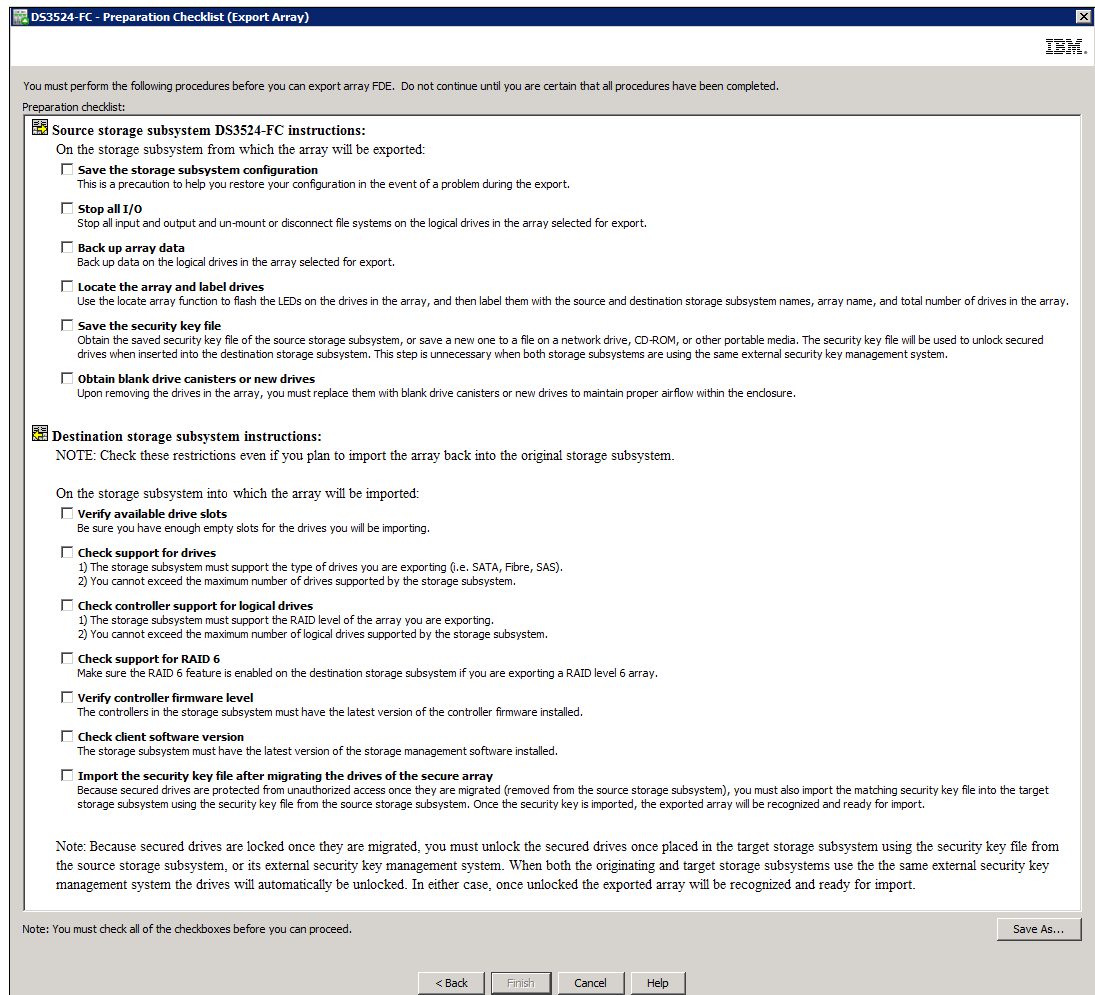


Figure 15-25 Preparation checklist

4. For security reason you will be asked to verify this selection as shown in Figure 15-26.




Figure 15-26 Confirmation to export array

5. During execution of the export command the array shows an degraded status as shown in Figure 15-27 on page 464




Figure 15-27 Degraded Status of the array

6. When the process has completed, the wizard will indicate that the drives can be removed as shown in Figure 15-28.

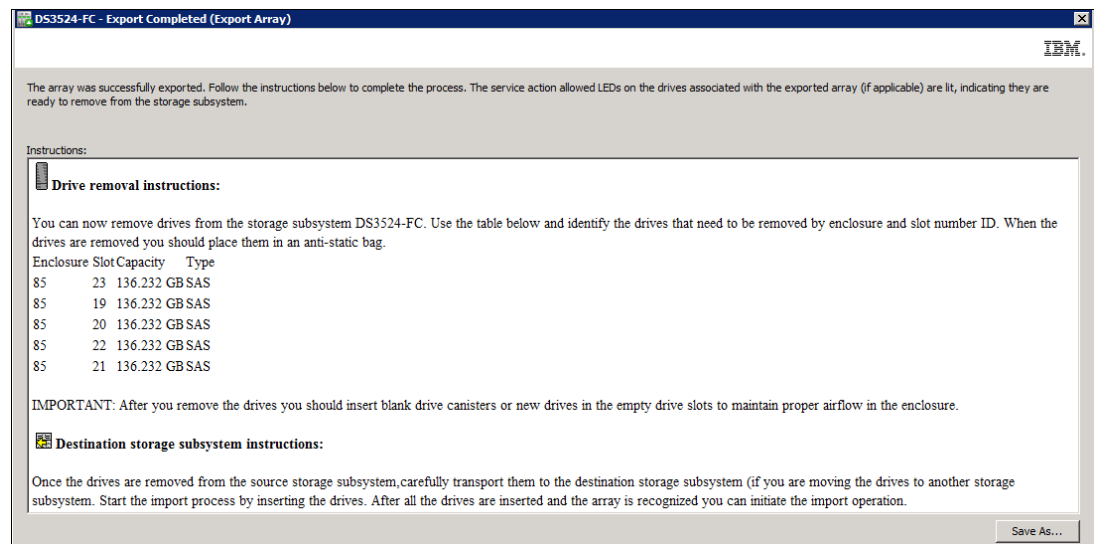


Figure 15-28 Export completed

The Storage Manager will now indicate that the array has been exported and the physical disks are offline. The windows shown in Figure 15-29 on page 465 and Figure 15-30 on page 465 confirm this information.




Figure 15-29 Array now exported

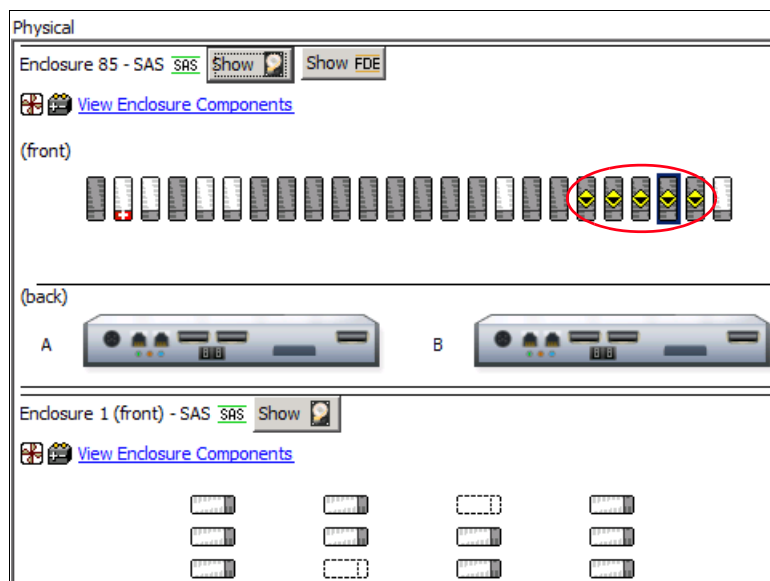


Figure 15-30 Array disks now offline

- Now you can carefully remove the exported drive array. It might be a good idea to switch the locate lights on for this array by clicking **Storage Subsystem** → select the exported array → **Array** and then click the **Locate** button. Now the yellow lights on the disk drives will blink.

## 15.6 Import secure drive array

After all the disks have been physically moved and are seated correctly in the target DS3500 enclosure, the drives can be unlocked and then the array can be imported. Make sure all the drives are seen by the DS3500 and that there is not a problem with them, as any problems will affect the array after it has been imported.

All drives will be in a locked state as shown in Figure 15-31 on page 466. In our example, we reinserted the array in the same Storage Subsystem so the key details are still ok. When you place the exported array in a separate Storage Subsystem you will get the message in the Status window Access is not allowed due to invalid key details.

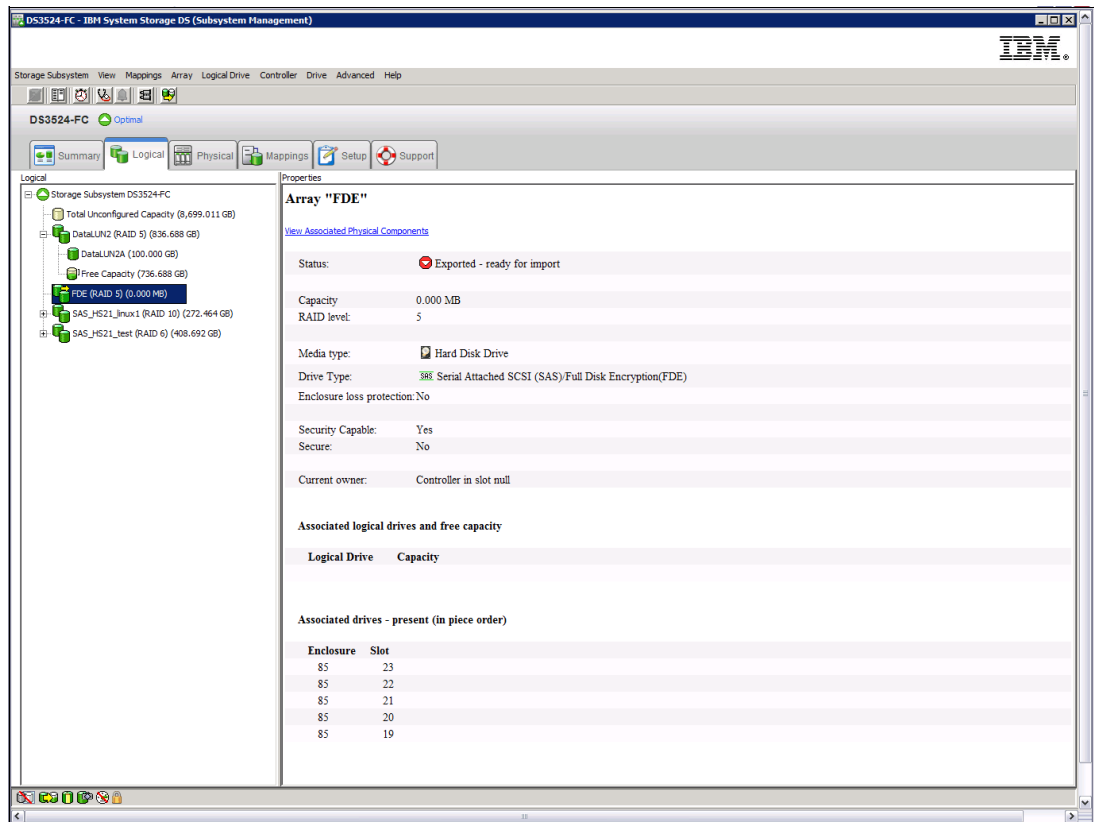


Figure 15-31 Drive properties show no key details and no security feature enabled

The DS3500 storage subsystem will also indicate that it needs attention due to the locked state of the drives. You can see the details about this situation in the Recovery Guru as shown in Figure 15-32.

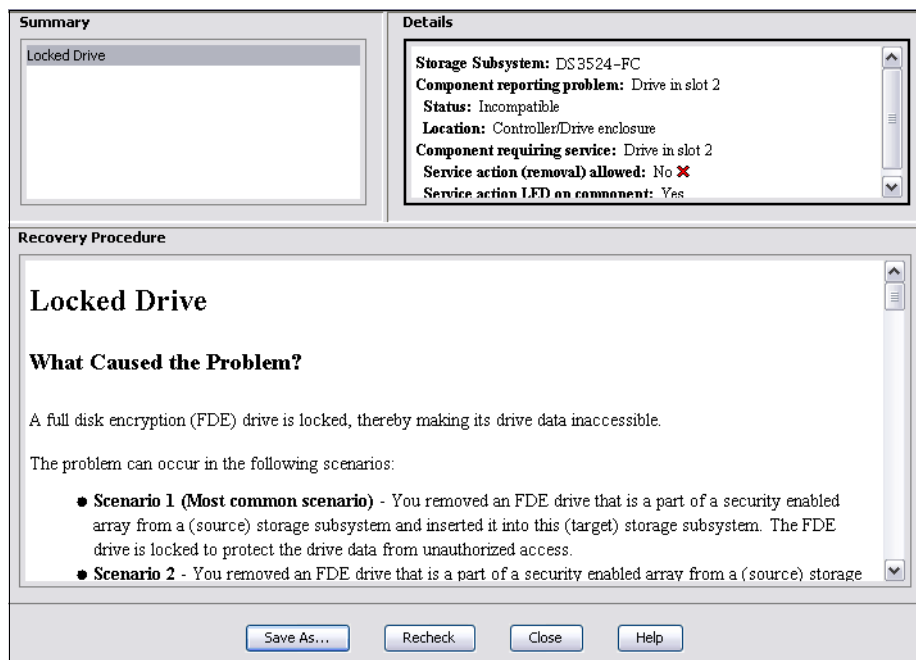


Figure 15-32 Drive locked message in the Recovery Guru

## 15.6.1 Unlock drives

Prior to importing the array, all drives must be unlocked. This procedure is performed on one disk, but will then be applied to all disks. The examples are taken on a different machine.

1. From the Physical tab in the Storage Manager window, select one of the FDE drives that has just been installed (it will be marked as offline). Right-click the drive and select **Unlock** as shown in Figure 15-33.

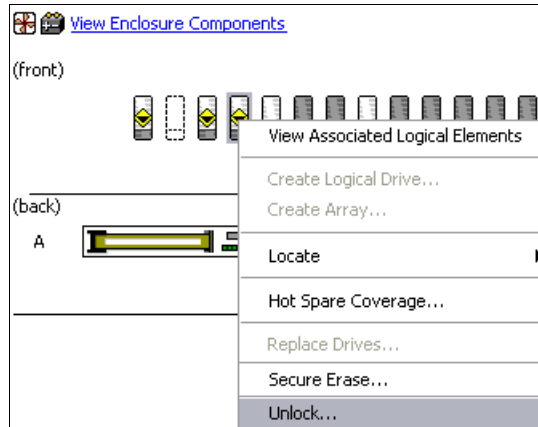


Figure 15-33 Select the Unlock option on the FDE drive

2. The DS3500 storage subsystem will recognize the FDE drives that are locked and will prompt for the key file and pass phrase, as shown in Figure 15-34.




Figure 15-34 FDE drives locked with the key file and pass phrase to unlock

3. When the correct key file is selected and the pass phrase is entered, all the drives will be successfully unlocked, as shown in Figure 15-35 on page 468.




Figure 15-35 FDE drive unlock completed successfully

## 15.6.2 Import array

To import the array, follow these procedures:

1. The array will be displayed in the Logical tab of the Storage Manager window, and is ready to import, as shown in Figure 15-36.

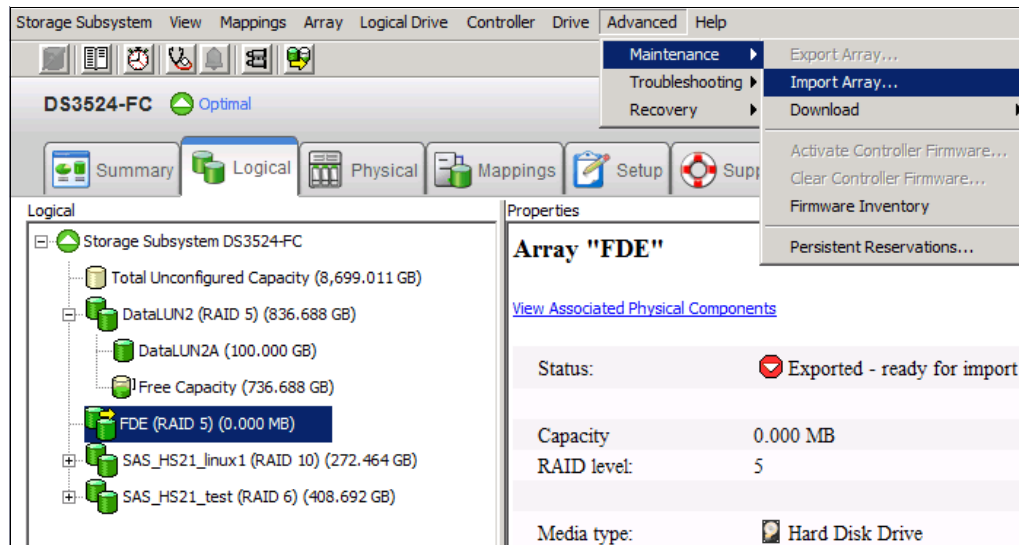


Figure 15-36 Array ready to be imported

Use the Import Array option to import an array that you have previously exported. From the Storage Manager window in the Logical view, select the array and click **Advanced** → **Maintenance** → **Import Array**.

2. The wizard opens and guide you through the procedure as shown in Figure 15-37. Click **Next**.

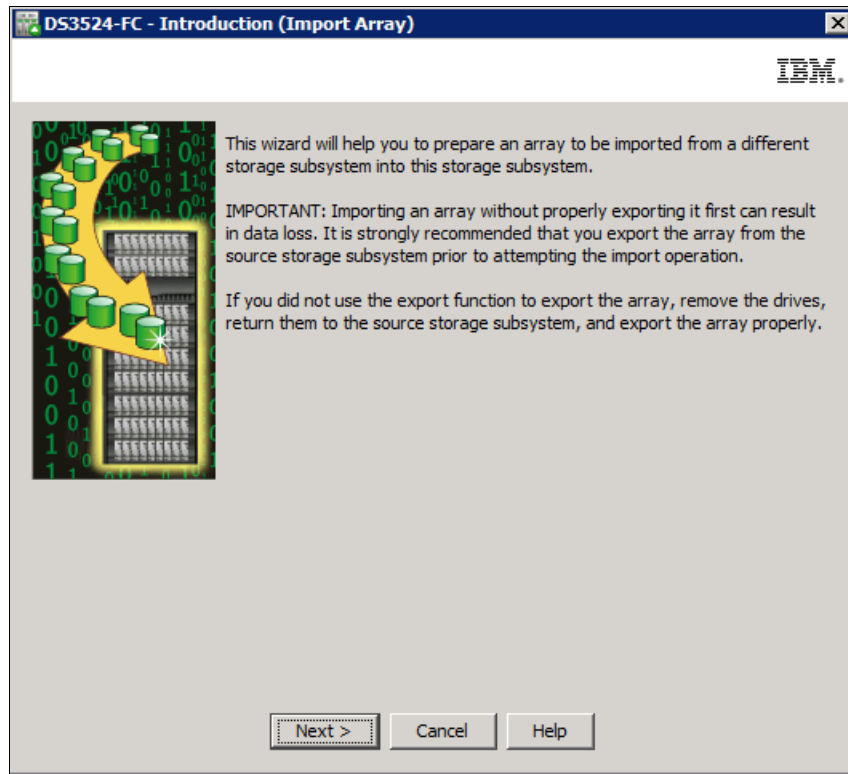


Figure 15-37 Introduction to import the array

3. The import array panel, shown in Figure 15-38, gives details about the disks that will be imported. Check that all the disks that are displayed are correct and that none of them are missing. If they are all correct, click **Finish**.

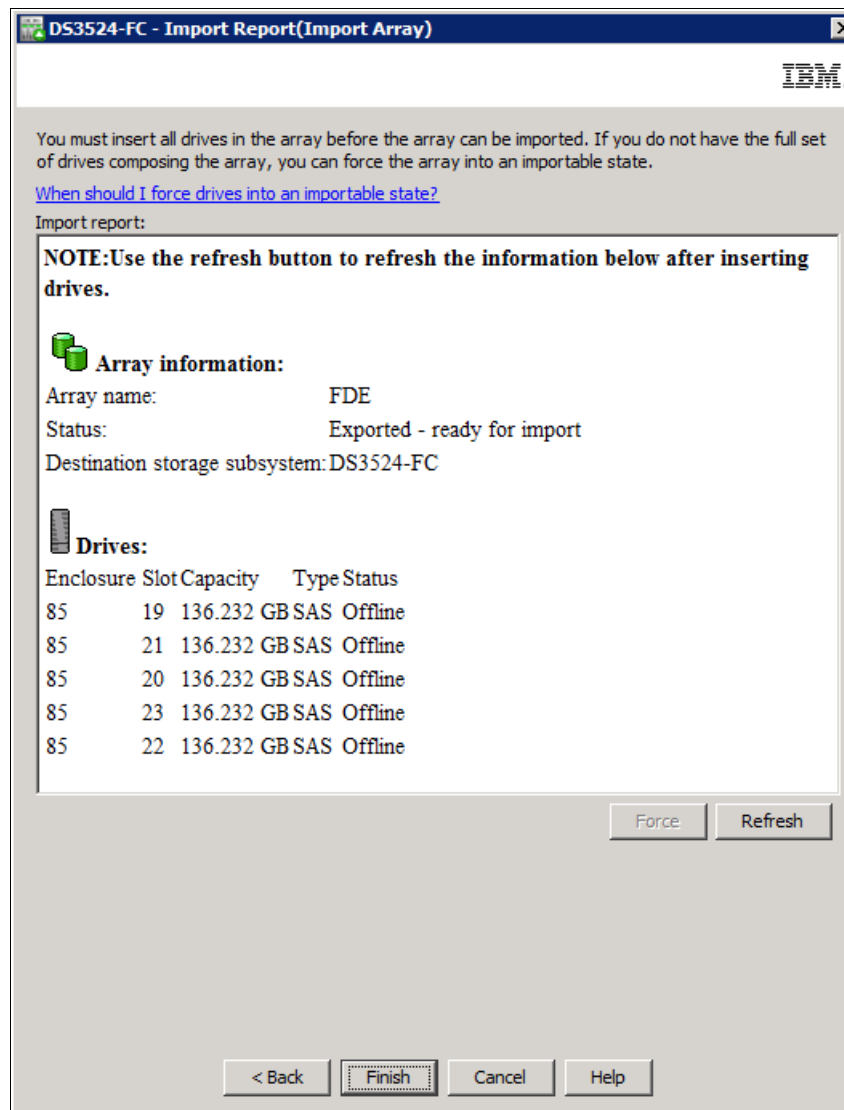


Figure 15-38 Import report indicating all that drives to be imported



4. Figure 15-39 confirms the importation of the drives.

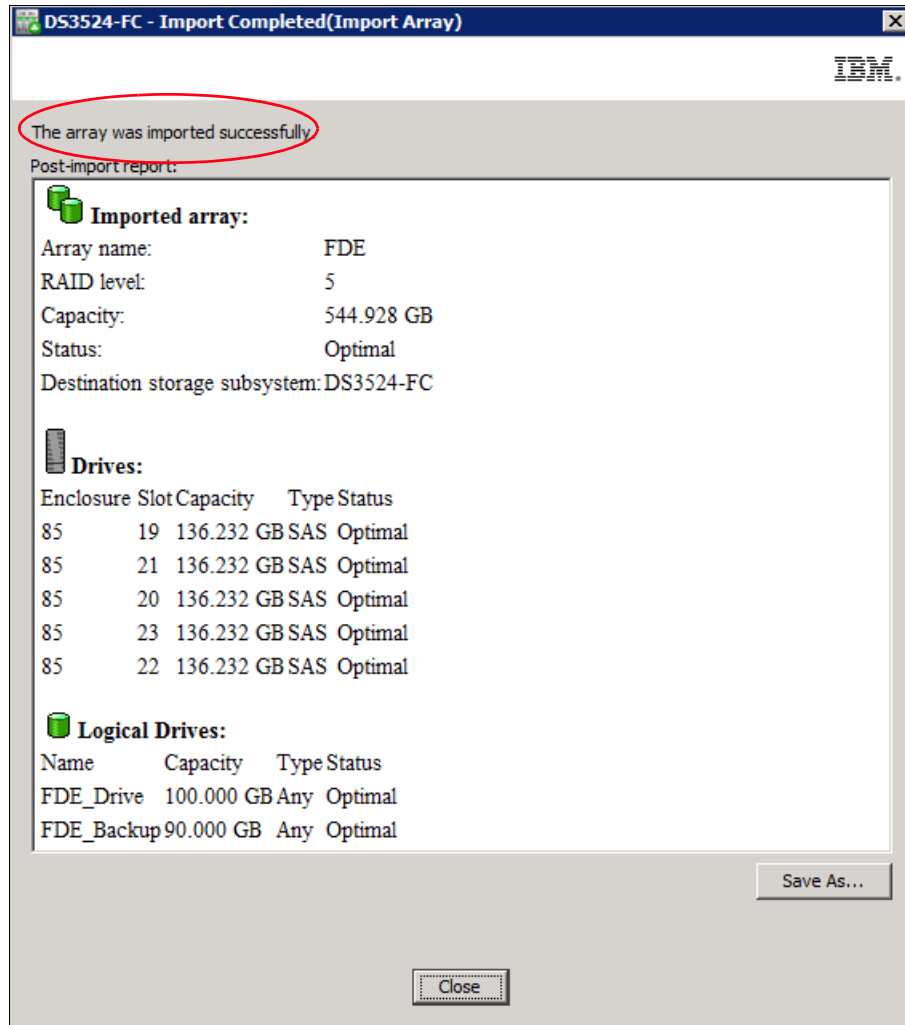


Figure 15-39 Import completed

The logical drives in the newly imported array are now ready to be mapped to hosts, where data can be accessed for read and write operations. After the Import o secure the drives again by right-clicking array and selecting **Secure array** as shown in Figure 15-40.

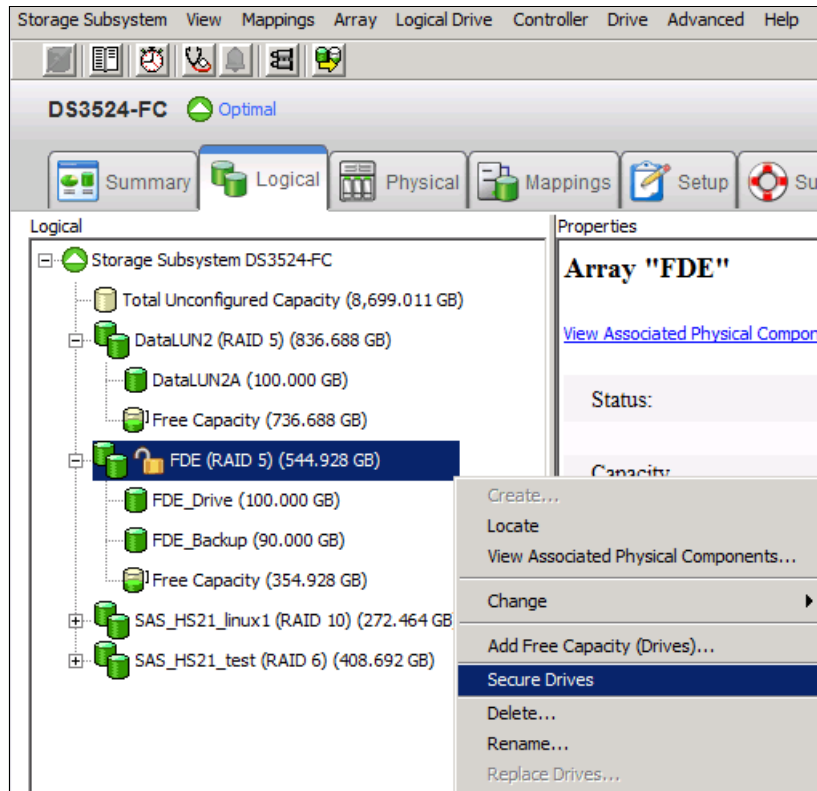


Figure 15-40 Secure drives again

Change the key after a secure array is imported to a DS3500 storage subsystem when it already has secure arrays. This will ensure that all secure enabled FDE drives will use a common key for authentication with the DS3500 storage subsystem when they are powered on. Follow the procedure described in 15.3.3, “Enable Disk Security on array” on page 454.



# IBM Remote Support Manager for Storage

In this chapter, we describe how to use IBM Remote Support Manager (RSM) for Storage with the IBM DS3500 Storage subsystem.

## 16.1 IBM Remote Support Manager for Storage

The IBM Remote Support Manager for Storage (RSM for Storage) software is a no-charge software package that is installed 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. It provides problem reporting and remote access for IBM Service for the IBM System Storage DS3000, DS4000, and DS5000 families.

The problem reporting utility 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 priority 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 on which it is installed. After it is 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 Service 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 email using the server's Ethernet interface as shown in Figure 16-1. An internal firewall that is managed by the RSM for Storage software keeps remote and local users of the system from accessing other devices on your intranet. Local and remote IBM users of the system do not have the ability to change any security features of the software.

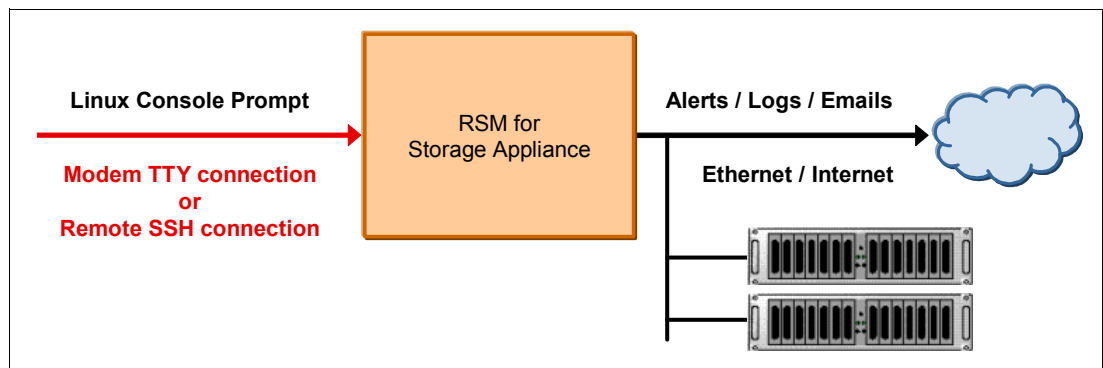


Figure 16-1 RSM layout

Your existing IBM System Storage DS Storage Manager v10, also known as Storage Manager, software monitors the storage subsystems. This software is configured to send SNMP traps to the Remote Support Manager when critical events are detected. Configuration of the management application is discussed in 16.2.4, "Configuring SNMP traps in Storage Manager" on page 506.

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 emails about status changes in these areas.

One RSM for Storage server can support up to fifty DS5000, DS4000, DS3000, FASiT 200, and FASiT 500 storage servers. Only IP connectivity to the Ethernet management ports of the subsystems is required: serial cables are not needed for access to the disk subsystems. The storage subsystems must be under a warranty or a current IBM maintenance agreement to be eligible to use the RSM for Storage software. There are no annual fees for RSM.

**Note:** Documentation and installation code is available at the following URL:

<http://www.ibm.com/storage/disk/rsm>

## 16.1.1 Hardware and software requirements

RSM for Storage has the following hardware and software requirements.

### Hardware requirements

IBM internal databases that support RSM for Storage use the IBM Machine Type and Serial Numbers of the servers, and therefore an IBM System x server must be used. This is required to properly register heartbeat records from the RSM for Storage system in IBM tracking databases. RSM for Storage can be run on a dedicated System x server or in a VMware client running on a System x server. See the *IBM RSM for Storage Compatibility Guide* (found at <http://www.ibm.com/support/docview.wss?uid=psg1MIGR-66062&rs=594>) for the minimum server requirements and a list of the specific servers that have been tested with the RSM software.

### *Suggested servers*

The following servers have been tested with the RSM for Storage software, and installation instructions are available for configuring and using them with RSM for Storage software:

- ▶ IBM System x3250 4364
- ▶ IBM System x306m 8849
- ▶ IBM System x3550 7978
- ▶ IBM System x3550m2 7946
- ▶ IBM System x3550m3 7944
- ▶ IBM DS-RSM Model RS3 (1818-RS3)

The *IBM Remote Support Manager for Storage: Installation Hints and Tips* document (found at <http://www.ibm.com/support/docview.wss?uid=psg1MIGR-66062&rs=594>) contains specific information about the BIOS configuration and device drivers that might be required to install the Linux OS on the above servers.

### **Other IBM servers**

You can use other System x servers that meet the following requirements:

- ▶ 512 MB memory.
- ▶ 20 GB disk space.
- ▶ Serial port: The serial port must be on the server system board. Some servers have a build-in Remote Supervisor Adapter (RSA) that also includes a serial port. The serial port on the RSA cannot be accessed by the RSM for Storage software.
- ▶ Ethernet port: 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 only has a single Ethernet port.

**Note:** To use servers other than those specifically tested with RSM for Storage, you will need to see the System x server support website for technical information about BIOS configuration and device drivers that might be required to install the Linux OS or configure the server's internal RAID capability.

The *IBM RSM for Storage Compatibility Guide* also contains the setup required for a VMware client that will host the Linux operating system running RSM for Storage.

### **RSM for Storage in a VMware virtual client**

RSM for Storage has been tested for operation in a VMware virtual client. See the *IBM RSM for Storage Compatibility Guide* for specific configurations.

In setting up the virtual client, allocate the following resources:

- ▶ Storage: 20 GB HDD.
- ▶ Memory: 512 MB.
- ▶ Assign exclusive physical ownership of the host system's first serial port to the virtual client as /dev/ttyS0.

The client must be configured to automatically start when the host reboots.

RSM for Storage has been tested on both SLES 10 and RHEL 5 running on an System x server with:

- ▶ VMware Server 1.05, 2.0
- ▶ VMware ESX Server 3

### **Software requirements**

The RSM for Storage software requires the following prerequisite software to monitor an IBM System Storage DS subsystem:

- ▶ IBM System Storage DS Storage Manager V10.70 or later (the latest version is preferable) with Event Monitor installed on a management station in a separate server.
- ▶ Storage subsystems with controller firmware supported by Storage Manager V10.70 or later. The latest supported firmware version is suggested.
- ▶ One of the following operating systems to install the RSM for Storage software:
  - Novell SUSE Linux Enterprise Server 9 (SP3 or SP4)
  - Novell SUSE Linux Enterprise Server 10 (Base, SP1, SP2, SP3)
  - Red Hat Enterprise Linux 4 AS (Update 4, 5, or 5)
  - Red Hat Enterprise Linux 5 (Base, Update 1, 2, 3, 4)

RSM for Storage software receives SNMP traps from the Event Monitor included with the Storage Manager. RSM for Storage software cannot be installed on the same system used to manage your storage network.

**Note:** See 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 website:

<http://www.ibm.com/support/docview.wss?uid=psg1MIGR-66062&rs=594>

### 16.1.2 DS-RSM Model RS3

The IBM System Storage DS Remote Support Manager for Storage (DS-RSM) Model RS3 server is available to simplify the task of obtaining the hardware and software components required for running the RSM for Storage software. The DS-RSM Model RS3 (Figure 16-2) can be ordered as an appliance with your IBM DS3500 System Storage storage subsystem.

The DS-RSM Model RS3 is designed to help manage problem reporting and analyze alert notifications for IBM System Storage DS subsystems. The DS-RSM Model RS3 contains more advanced processor and memory technology for enhanced processing capabilities. The DS-RSM Model RS3 can attach to the LAN, and collects alerts from IBM System Storage DS subsystems within the data center and expedites the forwarding of those alerts to an IBM Service center.



Figure 16-2 DS-RSM Model RS3

Additionally, the DS-RSM Model RS3 is designed to enable:

- ▶ IBM to establish communication with the attached supported storage systems
- ▶ Faster response time on processing alerts
- ▶ The ability to transfer detailed information to assist error analysis
- ▶ Enhanced problem resolution

The DS-RSM Model RS3 can help improve and facilitate service and support for the DS storage subsystems. For each data center with a DS storage subsystem present, install at least one DS-RSM Model RS3 that can interoperate with all of the DS5000/DS4000/DS3000 controllers located at the site. Each DS-RSM Model RS3 unit has the capability to manage up to fifty DS5000/DS4000/DS3000 storage systems.

DS-RSM Model RS3 is compatible with the following IBM System Storage DS subsystems:

- ▶ DS3000 series
- ▶ DS4000 series
- ▶ DS5000 series

The DS-RSM model RS3 is based on the System x3550 M3 Type 7944 server. The server is preloaded with Novell SLES 10 and RSM for Storage software. An internal modem is available as an option if you choose to use a modem to provide remote access for IBM Service. The internal modem might not be compatible in all countries or regions. In that case, you can provide an external modem appropriate for your phone system. As an alternative to using a modem, you can allow remote access using an external SSH connection. A keyboard, mouse, and monitor are required by IBM Service for installation and troubleshooting of the DS-RSM Model RS3 server. An optional, rack-mounted KVM drawer can be ordered with the DS-RSM Model RS3 server to meet this requirement.

### 16.1.3 Installation choices for RSM for Storage

Various options exist for how RSM for Storage can be set up, as shown in Figure 16-3.

See the *IBM RSM for Storage Compatibility Guide* and the *IBM Remote Storage Manager for Storage Planning, Installation and Users Guide* found at the following address for the latest information related to these options:

<http://www.ibm.com/support/docview.wss?uid=psg1MIGR-66062&rs=594>

| Installation Choices |                                                                                       |                                |                                                                                                                                                            |                                                                                                                                                                                            |
|----------------------|---------------------------------------------------------------------------------------|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hardware choices     | Build your own                                                                        | Use existing hardware          | Use your own server and obtain a supported version of the Linux. The system act as a gateway for IBM Support with limited scope of access to your network. | RSM documentation has step by step procedures for installing the Linux OS and the RSM for Storage software. Decide if you will allow IBM support to connect via a modem or SSH             |
|                      |                                                                                       | Order a new System x server    | Order a new System x server for use with RSM for Storage and get the Linux OS as a drop in the box option                                                  |                                                                                                                                                                                            |
|                      |                                                                                       | Run in a VMware guest OS       | Install a supported OS distribution in a VMware system and reduce your hardware costs.                                                                     | Consider hosting your SAN Management Station and the RSM for Storage on the same VMware platform.                                                                                          |
|                      | DS-RSM                                                                                | Purchase a RSM ready appliance | DS-RSM is reloaded with SLES 10 and RSM software and optional internal modem.                                                                              | DS-RSM is installed and supported by IBM. You will need to complete the installation worksheet and have networking infrastructure in place in order for IBM to complete the configuration. |
| Remote Access        | Problem reports are sent to IBM via email. You must choose a method of remote access. | SSH                            | Requires configuration of your external firewall                                                                                                           | The use of an authentication server adds another layer of access control.                                                                                                                  |
|                      |                                                                                       | Modem                          | Requires an analog phone line and modem.                                                                                                                   | You can use most external modems or if using the DS-RSM, order it with an internal modem adapter.                                                                                          |
| OS                   | RSM for Storage is a collection of software elements that integrates with a Linux OS  | SLES                           | 9.3 & 9.4<br>10.0, 10.1, 10.2 & 10.3                                                                                                                       | The DS-RSM is preloaded with Novell SLES 10.3                                                                                                                                              |
|                      |                                                                                       | RHEL                           | 4.4, 4.5 & 4.6<br>5.0, 5.1, 5.2, 5.3 & 5.4                                                                                                                 | RHEL is only available on "build your own" systems.                                                                                                                                        |
| KVM                  | RSM requires an attached monitor, mouse and keyboard for some tasks                   |                                | As the system can be anywhere in your network, you can connect to a KVM switch with other servers.                                                         | The DS-RSM can be ordered with a rack mounted KVM kit.                                                                                                                                     |

Figure 16-3 RSM installation choices



### 16.1.4 How RSM for Storage works

RSM for Storage uses an Ethernet connection for problem reporting and a modem (optional) for remote access by IBM Service as shown in Figure 16-4. The storage system can be any supported IBM System Storage DS subsystem.

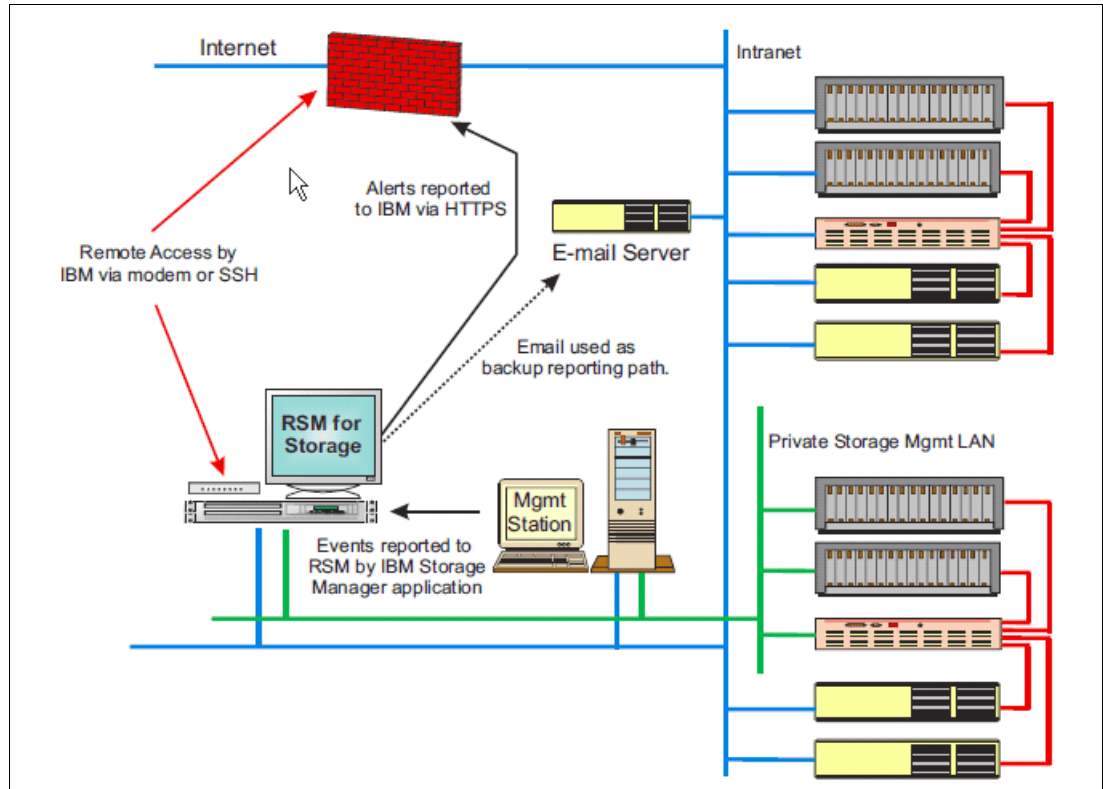


Figure 16-4 RSM for Storage connection

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 must be running IBM Storage Manager's Event Monitor. All the storage subsystems, the management station, the email server, and Internet gateway must be 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 16-4).

Figure 16-4 shows that some storage subsystems and other SAN devices are managed through a private management LAN and others that 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 is an example scenario for RSM. See Figure 16-4 on page 479 to understand the flow:

1. 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 the Major Event Log (MEL), Read Link Status (RLS), and 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 email.

SNMP traps are sent by the Storage Manager client or the Storage Manager's Event Monitor service. As the Storage Manager Client might not always be running, Event Monitor should be installed. See the Storage Manager documentation to learn about the installation of Event Monitor.

See 16.1.5, "Notification email and events filtering" on page 480 to configure the SNMP trap in Storage Manager.

3. IBM Service does problem determination based on information sent by the alert along with the problem determination data, such as MEL, RLS, and the subsystem profile. If the problem can be fixed with existing information, IBM Service contacts the customer either by phone or email to resolve the problem. After the problem is solved, either IBM Service or the customer must indicate *Service Complete* for all alerts in the RSM. IBM Service can dial to the RSM modem or use an 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 Service dials the RSM modem or uses an SSH connection, acknowledges the alert, and performs troubleshooting by connecting to the storage subsystem using the Storage Manager command-line interface (SMcli) or RLOGIN. IBM Service might need to contact the customer to obtain the password for the storage subsystem to use SMcli. IBM might also have to contact the customer to enable RLOGIN access.

If IBM Service needs to send or upload logs or other information from the RSM server to IBM, they can use FTP or email commands from the Linux shell at the RSM server while connected through the modem line or using SSH connection. Any data connected is sent to an IBM server through a customer network, not through the modem line (the modem connection is a TTY serial console, not an IP connection).

5. After the problem is resolved, all alerts must be closed either by IBM Service or the customer before reporting will resume for that subsystem.

**Note:** After 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.

### 16.1.5 Notification email and events filtering

The RSM for Storage software sends emails to notify you about important changes in status. Up to 20 contact people can be configured. One will be the primary contact for RSM related events and will receive all notifications about RSM system operation and problems. The primary contact will receive all notifications about RSM system operation and problems. Additional contacts can be configured for subsystems located at other sites. They only

receive notifications related to subsystem problems. See Table 16-1 for information about the notifications.

Table 16-1 E-Mail notifications

| Notification                                                                                                                 | Primary contact | Subsystem contact |
|------------------------------------------------------------------------------------------------------------------------------|-----------------|-------------------|
| Remote Access is enabled or disabled.                                                                                        | x               |                   |
| Remote User connects and logs into the RSM for Storage system or logs out.                                                   | x               |                   |
| Firewall is enabled or disabled.                                                                                             | x               |                   |
| Remote Access is about to expire or has expired.                                                                             | x               | x                 |
| An alert occurs for a subsystem.                                                                                             | x               | x                 |
| Daily status updates at local noon.                                                                                          | x               |                   |
| Daily status updates at local noon when a subsystem has an active alert.                                                     | x               | x                 |
| Daily phone line check fails.                                                                                                | x               |                   |
| Ping Check fails.                                                                                                            | x               |                   |
| RSM receives an event it is configured to ignore.                                                                            | x               | x                 |
| Notifications from IBM Support or RSM for Storage Development about the availability of updates to RSM for Storage software. | x               |                   |

There are several types of notifications sent by RSM to the primary contact or subsystem contact, as configured in the RSM:

- ▶ *Remote Access notifications* are sent when:
  - Remote access is enabled or disabled as shown in Example 16-1.

*Example 16-1 Remote access notification*

---

Remote access to the RSM for Storage system has been enabled.

The remote access timeout for localhost is now 35:59 (HH:MM)

There are no remote users connected to localhost at this time.

```
#
RSM system type: 1746-A4D
RSM system serial: KQDFGV8
RSM system name: localhost.localdomain
RSM modem Key:
K7EuTWmKlqoDuIsgP0xNYiVI13tDt1qoDuZbIvKXYI9pFXy5Le9qDMbewKFXkn8l093NAZxi96T3ZYr
b21zFSRUkaB00wEGcJ6NfmYvasl2w
RSM build: T1.15.4
RSM location: Datacenter 1
RSM address: River Street
RSM city: River City
RSM state: NC
RSM zip: 00000
```

```
RSM country: United States
RSM contact name: Max Musterman
RSM contact phone: 11234567890
RSM contact hours: 7am - 6pm
RSM contact timezone: GMT - 6
RSM contact alt. phone: 1 123 456 7891
RSM contact alt. hours: 6pm - 7am
RSM contact email: max.musterman@rivers.local
RSM contact country: United States
```

---

- A remote user connects or disconnects from the system.
  - The remote access automatic timeout is about to expire and the system has one or more active alerts. Timeout warnings are sent at 4 hours, 2 hours, and 1 hour before the timeout occurs.
- *Alerts Status notifications* are sent when an alert has been sent to IBM Service as shown in Example 16-2.

---

*Example 16-2 Alert Status notification*

---

Notification from RSM for Storage system: localhost

An alert was sent to IBM Service at Thu Aug 28 22:26:58 2010 UTC.

Remote Access to the RSM for Storage system is enabled for the next 12 hours

Check the RSM for Storage console for additional information related to this problem.

```
#
Machine type: 1746-A4D
Machine serial: 1234567
Machine name: ITS03500
Machine location: Datacenter 1
Machine address: River Street
Machine city: River City
Machine state: NC
Machine zip: 00000
Machine country: United States
Machine country code: 000
Machine info: None
Machine info name: None
Machine software id:
Machine contract:
#
Contact name: Max Musterman
Contact phone: 12234567890
Contact hours: 7am - 6pm
Contact timeZone: GMT - 6
Contact alt. phone: 1 223 456 7891
Contact alt. hours: 6pm - 7am
Contact email: max.musterman@rivers.local
Contact country: United States
```

The event code was: 5005

The event message was:

Place controller offline.  
Controller.  
Controller in slot B

---

- *Daily Status emails* 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 might be active for storage subsystems as shown in Example 16-3.

---

*Example 16-3 Daily Status email*

---

Status report for IBM Remote Support Manager running on: localhost.localdomain

Application has been running for: 0 days, 21 hours, 41 minutes

System status: OK  
One or more subsystems has reported a problem.  
The internal firewall is enabled.  
Remote access is disabled.  
1 storage subsystem has an active alert.  
0 alerts have been acknowledged by IBM Service.  
1 alert has been sent to IBM, but has not yet been acknowledged.  
0 additional alerts are pending.

Subsystem: ITS03500, located at Datacenter 1

0 alerts acknowledged.  
1 alert sent to IBM.  
0 additional alerts pending.

Because IBM Service has already been alerted to a problem  
with this subsystem, no additional alerts will be sent to IBM  
for this subsystem until existing alerts are closed.

This problem has been open for 19 hours.

```
Mini Inventory : Family IBM MTM IBM Serial Firmware Hostname
Mini Inventory : DS3000 1746-A4D 1234567 07.70.35.08 ITS03500
#
RSM system type: 1746A4d
RSM system serial: KQDFGV8
RSM system name: localhost.localdomain
RSM modem Key:
K7EuTWmKlqoDuIsgP0xNYiVI13tDt1qoDuZbIvKXYI9pFXy5Le9qDMbeWKFXkn8l093NAZxi96T3ZYr
b21zFSRUkaB00wEGcJ6NfmYvasl2w
RSM build: T1.15.4
RSM location: Datacenter 1
RSM address: River Street
RSM city: River City
RSM state: NC
RSM zip: 00000
RSM country: United States
RSM contact name: Max Musterman
RSM contact phone: 12234567890
RSM contact hours: 7am - 6pm
```

```
RSM contact timezone: GMT - 6
RSM contact alt. phone: 1 223 456 7891
RSM contact alt. hours: 6pm - 7am
RSM contact email: max.musterman@rivers.local
RSM contact country: United States
```

---

- *Firewall Status notifications* are sent when the internal firewall is enabled or disabled by a root or admin user as shown in Example 16-4.

**Example 16-4 Firewall notification**

---

RSM for Storage (localhost) internal firewall has been disabled.

```
#
RSM system type: 1746A4D
RSM system serial: KQDFGV8
RSM system name: localhost.localdomain
RSM modem Key:
K7EuTWmKlqoDuIsgP0xNYiVI13tDt1qoDuZbIvKXYI9pFXY5Le9qDMbeWKFxkn8l093NAZxi96T3ZYr
b21zFSRUkaB00wEGcJ6NfmYvasl2w
RSM build: T1.15.4
RSM location: Datacenter 1
RSM address: River Street
RSM city: River City
RSM state: NC
RSM zip: 00000
RSM country: United States
RSM contact name: Max Musterman
RSM contact phone: 12234567890
RSM contact hours: 7am - 6pm
RSM contact timezone: GMT - 6
RSM contact alt. phone: 1 223 456 7891
RSM contact alt. hours: 6pm - 7am
RSM contact email: max.musterman@rivers.local
RSM contact country: United States
```

---

- *Ping check notifications* are sent when two ping checks in a row fail. Ping checks of the configured SMTP sever and Management Station are performed every half hour.
- *Ignored Event notifications* are sent when an event is received that is configured to be ignored by the RSM for Storage system, and is therefore not reported to IBM Service. These are events for which a response by IBM Service is not usually required as shown in Table 16-2.

**Table 16-2 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 xxxxxxxx cannot reach the indicated storage system.  |
| None       | The persistent monitor running on Host xxxxxxxx can now reach the indicated storage system. |
| 4011       | The logical drive is not on a preferred path due to ADT/RDAC.                               |

- *Phone Line Check notification* is sent if the RSM for Storage system cannot verify that an active phone line is connected to the modem. The check is run daily after midnight local time.

### **RSM and Storage Manager email alerts**

Storage Manager in the management station can be configured to send email notifications when a problem is detected. However, this feature *must be disabled* when RSM for Storage is installed if the email contact configured in the Storage Manager is the same as the email contact configured in the RSM Contact List. Otherwise, you will receive multiple notification emails about the same problem: one notification from RSM and another one from Storage Manager.

To disable email alerts in Storage Manager, perform 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 email destinations.

If there are email addresses already configured to receive emails from Storage Manager but are not listed in the RSM Contact List (see Figure 16-11 on page 492 for the Contact List), it is not necessary to delete them in Storage Manager.

## **16.1.6 Remote access methods**

The required Remote Access for IBM Service 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 by modem**

The functional requirements for the modem are minimal: it is only used for remote access by IBM Service. Most "Hayes-compatible" external modems can be used. Any V.92 56K modem supporting the common AT command set will work.

The RSM for Storage software has been tested with the following modems:

- Multitech MultiModem II MT5600BA
- Multitech MultiModem ZBA MT5634ZBA

See the *IBM RSM for Storage Compatibility Guide* and *IBM Remote Support Manager for Storage: Installation Hints and Tips* for updated information about which modems are supported. You will need to contact your IT support staff for installation and problem resolutions related to the modem.

The optional internal modem that can be ordered with the DS- RSM Model RS3 server is supplied with an RJ-11 phone cord. You might need to supply an adapter to connect the RJ-11 phone cord to your phone system.

### **Remote access by 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 will need to map a port on your external firewall to the IP address and SSH port (22) on the RSM for Storage system. Although 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, and the supplement *IBM Remote Support Manager for Storage: Installation Hints and Tips*, found at the following address:

<http://www.ibm.com/support/docview.wss?uid=psg1MIGR-66062&rs=594>

## 16.1.7 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 Service, the interface is a command-line interface when connected to the RSM server through a modem or SSH connection.

Figure 16-5 shows an example of a System Configuration menu for an already configured and not activated RSM system.

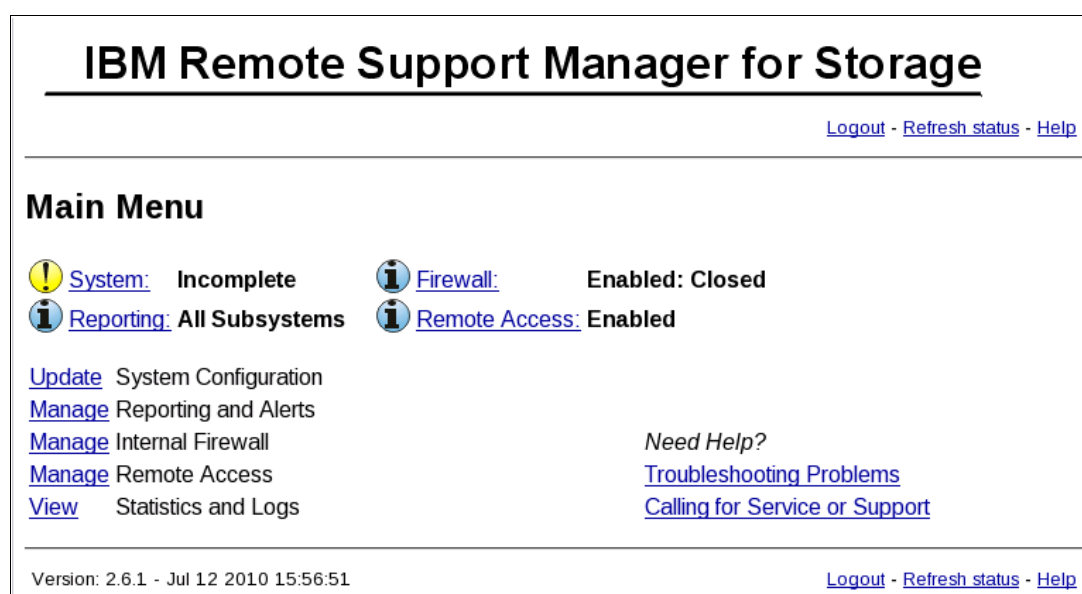


Figure 16-5 RSM System Configuration menu

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 might be displayed to the left of each link. The status values for each of these are as follows:

- System:
  - OK: Remote Support Manager is operating properly.
  - Incomplete: One or more required configuration settings are missing or incorrect.
  - Problem: There is a problem that is preventing correct operation.

**Note:** Reporting is disabled until all configuration problems are fixed.

- Reporting:
  - All Subsystems: Reporting is enabled for all configured subsystems.



- Standby: Reporting has been disabled for all subsystems.
- Partial: Reporting has been disabled for some but not all subsystems.
- Suspended: Reporting is not performed while a configuration problem exists.
- Storage Problem: A problem has been reported by one or more subsystems.
- ▶ Firewall:
  - Enabled:Closed: The firewall is enabled. Only connections required for reporting are allowed.
  - Enabled:Open: The firewall is enabled. Connections are open to one or more devices being serviced.
  - Enabled:Custom: The firewall is enabled. Only connections required for reporting and permitted by the custom rules defined in `/etc/rsm/rsm-firewall.conf` are allowed.
  - Disabled: The firewall is disabled. There are no restrictions on access to the networks connected to the RSM.
- ▶ Remote Access:
  - Disabled: Modem answer and remote user login is disabled.
  - Enabled: Modem is enabled and remote user login is allowed.
  - Active: A remote user is logged into the system.

### 16.1.8 RSM security considerations

RSM for Storage controls security for remote access by managing the hardware and software components of the server on which it is installed. After it is installed, the server should be considered a single purpose appliance for problem reporting and remote access support for your storage subsystems. Do not use it for other applications.

Remote access to your system has the following four layers of control:

- ▶ The modem is configured to only answer when Remote Access is enabled by the RSM for Storage software. Likewise, the SSH daemon is only allowed to respond to connection attempts when Remote Access is enabled.

You can manually enable and disable remote access, or you can choose to have remote access automatically enabled when a storage subsystem reports a problem. When remote access is enabled, a timer is started that will automatically disable remote access when it expires. You do not have to remember to make the system secure after service has been completed.

The person identified as the primary contact for the RSM for Storage system is notified by email whenever a change in the remote access settings occurs, and all state changes are written to the security log.

- ▶ The user ID reserved for remote access (`rservice`) is only valid when Remote Access is enabled. Attempts to log in using the `root`, `admin`, or `lservice` user IDs are rejected.

**Note:** For this reason, do not create additional users on this system.

- ▶ The initial login password is changed daily at midnight UTC. IBM Service has an internal tool that provides the current password for RSM for Storage systems.
- ▶ After validation of the initial login password, remote users are presented with a challenge string, which also requires access to an internal IBM tool to obtain the correct response.

The response also includes an IBM employee user name that is recorded in the RSM for Storage security log.

## 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 restrictions regarding the 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* for the RSM for Storage browser user interface can be changed by the Linux *root* user using the command **rsm-passwd admin** or **rsm-passwd lservice**. For security reasons, set a separate 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.

The Switch User (**su**) command is disabled to prevent a normal user from attempting to become “root” and have unrestricted access to the system. The RSM for Storage software makes other changes in program and directory permissions to limit what programs and files these users can access.

## Internal firewall

RSM for Storage includes an internal firewall to limit the scope of access a remote user has to your network. Without an internal firewall, the remote user will have unrestricted access to your network. The RSM software configures an internal firewall on the RSM system to limit the scope of access that users of the RSM system have to your network, as shown in Figure 16-6 on page 489. When no alerts are active, the firewall only allows incoming SNMP traps and outbound SMTP email. When an alert occurs, a rule is automatically added to the firewall to allow access to the configured controllers for the storage subsystem reporting the problem. There might be times when you want to allow IBM Service to be able to access a device to troubleshoot a problem (such as a performance issue) for a subsystem that is not reporting a failure. You can manually enable “service access” for any configured storage subsystem. Service Access settings have a configurable timeout from 12 to 96 hours, after which the firewall rules for access are removed.

Firewall rules that are added for a device that is reporting an alert are removed when the alert is closed on the RSM for Storage system.

Although the internal firewall effectively limits the scope of access for remote users, at the same time it also limits the access of any program running on the server. Because management applications such as Storage Manager or IBM Director require access to all devices being managed, the presence of this internal firewall prevents the use of the RSM server as a management station. Therefore, the RSM for Storage system should be considered to be a single purpose appliance dedicated to problem reporting and remote access for IBM Service.

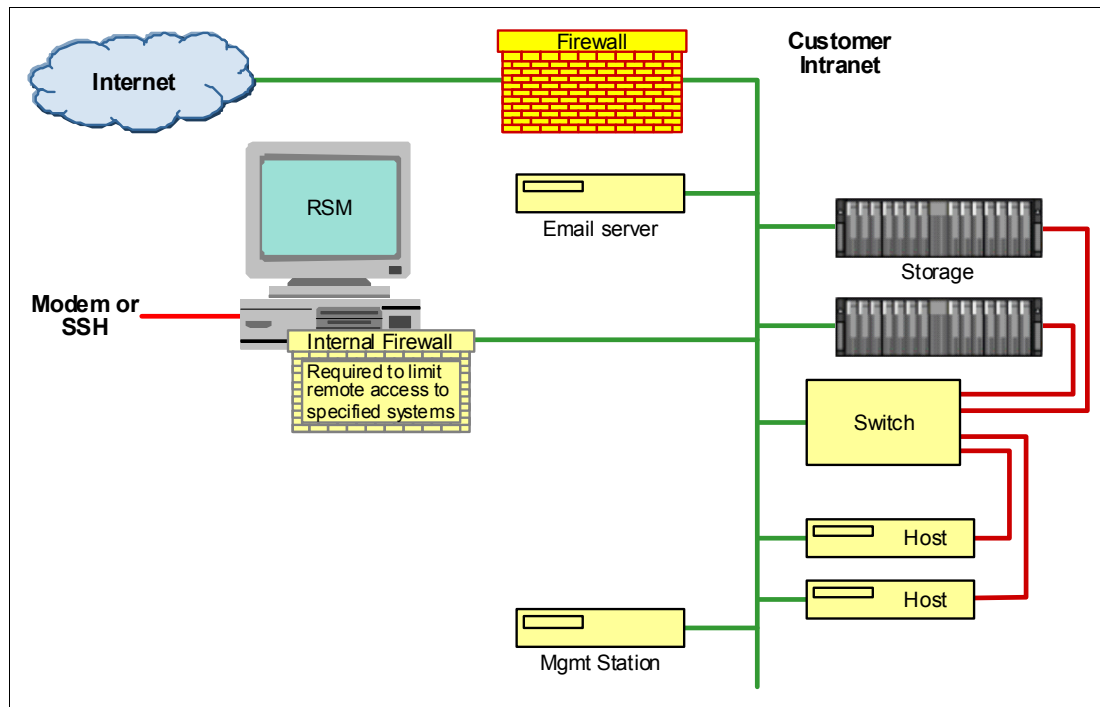


Figure 16-6 RSM internal firewall

## 16.2 Installing and setting up RSM

In this section, we show how to install and configure RSM. Before beginning this task, go to the RSM support website and carefully review the *IBM Remote Storage Manager for Storage* documents *Planning, Installation and Users Guide*, GC26-7933, and *Installation Hints and Tips*, found at the following address:

<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. Log on locally to the Graphical User Interface of the workstation because RSM resets the firewall settings to prevent remote access to the Linux workstation and will be configured later.

### 16.2.1 Installing the host OS

There are various operating systems supported for the RSM host as shown in “Software requirements” on page 476. 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
- ▶ libXp-devel
- ▶ Postfix
- ▶ net-snmp-utils
- ▶ KDE desktop environment

See *IBM Remote Storage Manager for Storage: Planning, Installation, and User's Guide*, GC26-7933 for specific operating system installation instructions.

## 16.2.2 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.

## 16.2.3 Setting up RSM

After the installation is complete, we have to set up RSM by performing the following steps:

1. On the Linux login window, perform the following steps:
  - a. Click **Session** and select **KDE**.
  - b. Log in as the admin user.
  - c. Click the Manage icon to open a web browser that shows the main administration window as shown in Figure 16-7. Click **Login**.

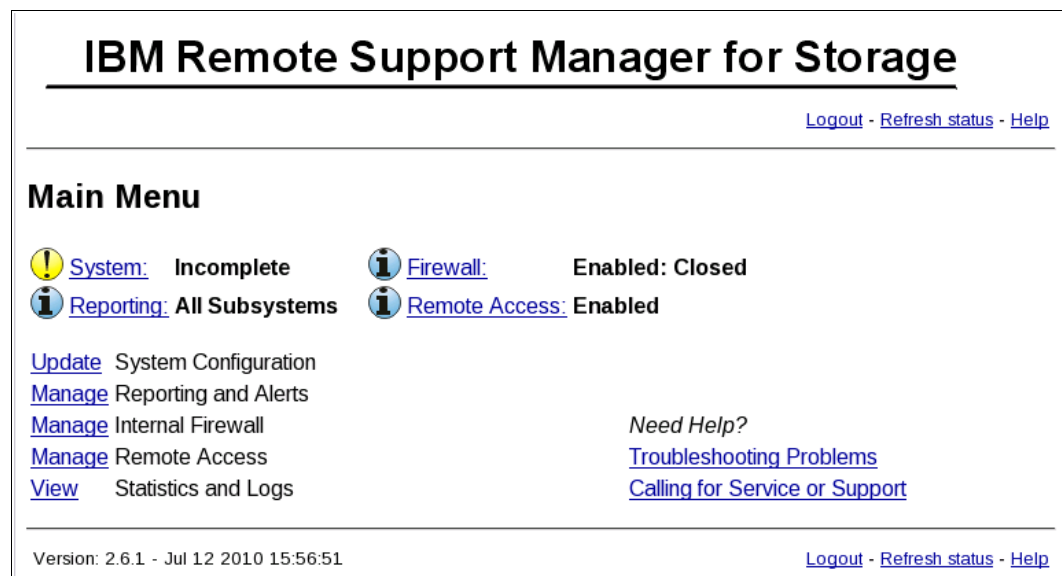


Figure 16-7 RSM administration window

2. Click **Login** and 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 16-8 on page 491.

Figure 16-8 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 16-9.

Figure 16-9 RSM main menu

4. The System Configuration window shows incomplete tasks that need to be accomplished before the RSM system can be used as shown in Figure 16-10 on page 492.

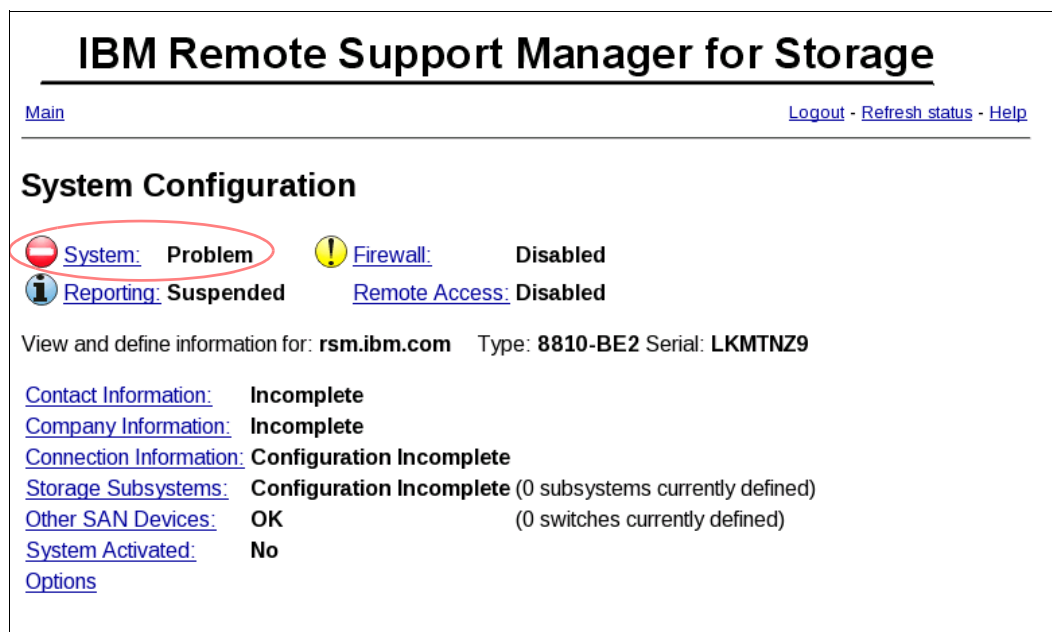


Figure 16-10 System Configuration incomplete

- Click **Contact Information** to access the contact list to add a new contact (Figure 16-11).

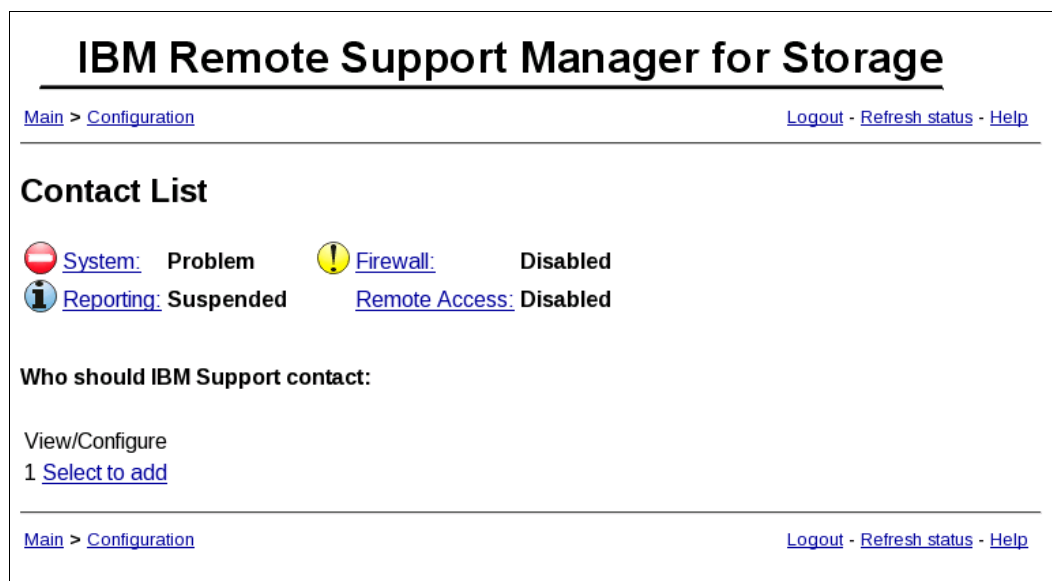




Figure 16-11 RSM contact list


- Click **Select to add**, fill out the form, and click **Update configuration**. This information is important: when a DS Storage Subsystem reports a problem, IBM Service will contact the person specified here. See Figure 16-12 on page 493.

IBM Remote Support Manager for Storage

[Main](#) > [Configuration](#) > [Contact List](#)
[Logout](#) - [Refresh status](#) - [Help](#)

Contact Person Information

 [System](#): Problem
  [Firewall](#): Disabled

 [Reporting](#): Suspended
 [Remote Access](#): Disabled

Who should IBM Support contact:

\* next to an entry indicates missing or incorrect information

\* Contact person
  \* E-mail address
  \* Phone number
  \* Hours to call or other information (see Help)
  \* Time Zone
  \* Alternate phone number (optional)
  Hours to call or other information for alternate number (optional)
 ☒ Make this person the primary contact for the RSM for Storage system
  \* Country or region

To remove this contact from the configuration:



Figure 16-12 RSM contact person information


- The specified contact will be added to the contact list. Multiple contacts can be defined and assigned to different storage subsystems. Click **Configuration** to return to System Configuration as shown in Figure 16-13.

IBM Remote Support Manager for Storage

[Main](#) > [Configuration](#)
[Logout](#) - [Refresh status](#) - [Help](#)

Contact List

 [System](#): Problem
  [Firewall](#): Disabled

 [Reporting](#): Suspended
 [Remote Access](#): Disabled

Who should IBM Support contact:

View/Configure

1 [Max Mustermann](#) United States (Primary contact for RSM for Storage)
 2 [Select to add](#)

[Main](#) > [Configuration](#)
[Logout](#) - [Refresh status](#) - [Help](#)

Figure 16-13 RSM contact list with contacts

- The task Contact Information in System Configuration will be marked OK. Click **Company Information** as shown in Figure 16-10 on page 492.


9. Complete the form with the appropriate information and click **Update configuration** as shown in Figure 16-14.


## IBM Remote Support Manager for Storage


[Main](#) > [Configuration](#)

[Logout](#) - [Refresh status](#) - [Help](#)

### Company Information

 [System:](#) Problem

 [Firewall:](#) Disabled

 [Reporting:](#) Suspended

[Remote Access:](#) Disabled

Company Information:

\* next to an entry indicates missing or incorrect information

|                                            |                     |
|--------------------------------------------|---------------------|
| <input type="text" value="IBM"/>           | * Company name      |
| <input type="text" value="My Street"/>     | * Street Address 1  |
| <input type="text"/>                       | Address 2           |
| <input type="text" value="Local City"/>    | * City              |
| <input type="text" value="No State"/>      | * State or Province |
| <input type="text" value="12345"/>         | * Postal Code       |
| <input type="text" value="United States"/> | Country or region   |

Figure 16-14 RSM company information

10. The task Company Information in System Configuration will be marked OK. Click **Connection Information** as shown in Figure 16-10 on page 492.




11. Complete the form with the appropriate information and click **Update configuration** as shown in Figure 16-15.


## IBM Remote Support Manager for Storage


[Main](#) > [Configuration](#)

[Logout](#) - [Refresh status](#) - [Help](#)

### Connection Information

 [System](#): Problem

 [Firewall](#): Disabled

 [Reporting](#): Suspended

[Remote Access](#): Disabled

This system is: **rsm.ibm.com**    Type: **8810-BE2**    Serial: **LKMTN29**

\* next to an entry indicates missing or incorrect information

|                                            |                                                       |
|--------------------------------------------|-------------------------------------------------------|
| <input type="text" value="DIRECT"/>        | IP address of SMTP server, or DIRECT (See Help page)  |
| <input type="text" value="9.42.171.25"/>   | IP address of Management Station (optional, see help) |
| <input type="text" value="LAB"/>           | * Location of RSM for Storage system                  |
| <input type="text" value="Lab Street"/>    | * Street Address                                      |
| <input type="text" value="Downtown Lab"/>  | * City                                                |
| <input type="text" value="LAB State"/>     | * State or Province                                   |
| <input type="text" value="12345"/>         | * Postal Code                                         |
| <input type="text" value="United States"/> | Country or region                                     |

#### Reporting Connection to IBM

The current Reporting Transmit Method is: **ECC with Email backup**.

Note: You can change the method of sending reports to IBM on the [Options configuration page](#).

IBM ECC Client Status: **Unknown, assumed OK** (Run test to update).

A change has been made to the configuration. Run this (or the full configuration test again).

The ECC Client Connection Test is a shorter version of the Configuration Test. Refresh the main configuration page until the test results are displayed.  
[ECC Client Connection Test](#)

**ECC HTTPS Proxy Configuration (if required by your network)**  
 Proxy IP Address (blank if not required)  
 Proxy Port Number  
 Proxy User ID (optional)  
 Proxy Password (optional)  
 Proxy Password Confirm

#### Remote Access Connections - you must configure one or both of the following:

**For a remote modem connection:**  
 \* Modem phone number (0...9 and spaces) or "NOMODEM"  
 Phone Line Check number. (See Help page.)

**For a remote SSH connection:**  
 \* External IP Address IBM should use (blank to disable)  
 External Port IBM should use  
 External Firewall User ID (optional)  
 External Firewall Password (optional)  
 External Firewall Password Confirm  
 SSH Port RSM will listen on for remote access (default: 22)

Figure 16-15 RSM connection information

12. The task Connection Information in System Configuration will be marked OK. Click **Storage Subsystems** as shown in Figure 16-10 on page 492.
13. Click **Select to add** to define a storage subsystem in RSM (Figure 16-16).

**IBM Remote Support Manager for Storage**

[Main > Configuration](#) [Logout](#) - [Refresh](#) - [Help](#)

---

### Configure Storage Subsystems

[System:](#) **Incomplete**
[Firewall:](#) **Disabled**  
[Reporting:](#) **Standby**
[Remote Access:](#) **Disabled**

There are currently 0 storage subsystems configured:  
 View/Configure Firmware  
 1 [Select to add](#)

---

[Main > Configuration](#) [Logout](#) - [Refresh](#) - [Help](#)

Figure 16-16 RSM configure storage subsystem

RSM for Storage reports DS Storage System Expansion units as well, and therefore the configuration of drive expansion units is required. During the configuration test, 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 be available for setting the IBM machine type and serial numbers of each detected drive expansion unit.

Complete the form with the following fields:


- a. The name of a storage subsystem, as shown in the DS Storage Manager.
- b. The country where the storage subsystem is located.
- c. A description of 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 top of the DS storage subsystem. Select an entry from the contact list.


Click **Update configuration**. Figure 16-17 shows the completed form with sample data.


## IBM Remote Support Manager for Storage

[Main](#) > [Configuration](#) > [Storage Subsystems](#)[Logout - Help](#)

### Storage Subsystem Information

 [System](#): Incomplete

 [Firewall](#): Disabled

 [Reporting](#): Standby

[Remote Access](#): Disabled

#### Storage Subsystem Information:

\* next to an entry indicates missing or incorrect information

DS3524-FC

\* Hostname (Defined in Storage Manager)

Firmware: **Not available at this time**

United States

\* Country or region

LAB

\* Location: Room / Building

LAB Street

\* Street Address

Downtown LAB

\* City

LAB State

\* State or Province

12345

\* Postal Code

1.1.1.2

\* IP Address #1

1.1.1.3

\* IP Address #2 (If present)

1746-C4A

\* IBM Product ID (TTTT-MMM) or MTM. Refer to the [Help Page](#) for this field.

nonumb

\* IBM Serial Number (7 characters)

Max Mustermann

\* Contact person for this subsystem

.....

\* Storage Manager Password

.....

\* Confirm Password

(May be required in order for RSM for Storage to be able to automatically collect a full set of log files. See [Help](#) for more information).

Part of an IBM Solution: **None** (Must be set by IBM Support. Refer to the Help for this page.)

[Update configuration](#)

To remove this subsystem from the configuration: [Delete this device](#)

To define a new subsystem based on this configuration: [Copy this form](#)


Figure 16-17 RSM storage subsystem information


- If there is an Expansion cabinet attached, the System Configuration shows that the Configuration is still incomplete. The Storage Subsystem Information must be added for the enclosure too as shown in Figure 16-18 on page 498.


IBM Remote Support Manager for Storage

[Main](#)
[Logout](#) - [Refresh status](#) - [Help](#)

System Configuration


[System:](#) Incomplete


[Firewall:](#) Disabled


[Reporting:](#) Configuration Problem

[Remote Access:](#) Disabled

View and define information for: **rsm.ibm.com**    Type: **8810-BE2** Serial: **LKMTNZ9**

[Contact Information:](#)    OK

[Company Information:](#)    OK

[Connection Information:](#)    OK

[Storage Subsystems:](#)    **Configuration Incomplete** (1 subsystems currently defined)

[Other SAN Devices:](#)    OK    (0 switches currently defined)

[System Activated:](#)    No

[Options](#)


Figure 16-18 RSM Configure Storage Subsystem continued


15. Select the Storage Subsystem again as shown in Figure 16-19.


IBM Remote Support Manager for Storage

[Main](#) > [Configuration](#)
[Logout](#) - [Refresh](#) - [Help](#)

Configure Storage Subsystems


[System:](#) Incomplete


[Firewall:](#) Disabled


[Reporting:](#) Configuration Problem

[Remote Access:](#) Disabled

There are currently 1 storage subsystems configured:

View/Configure    Firmware

1 [DS3524-FC](#)    07.70.16.01

2 [Select to add](#)

Figure 16-19 RSM Configure Storage Subsystem continued

16. Now you will be able to fill in the necessary information for the additional enclosure as shown in Figure 16-20 on page 499. Click the **Copy this form** button to add more enclosures.

## IBM Remote Support Manager for Storage

[Main](#) > [Configuration](#) > [Storage Subsystems](#)
[Logout](#) - [Help](#)

---

### Storage Subsystem Information

[System:](#) Incomplete

[Reporting:](#) Configuration Problem

[Firewall:](#) Disabled

[Remote Access:](#) Disabled

**Storage Subsystem Information:**

Hostname (Defined in Storage Manager)

Firmware: **07.70.16.01** (as of Mon Aug 9, 2010 15:31:14)

Country or region

Location: Room / Building

Street Address

City

State or Province

Postal Code

IP Address #1

IP Address #2 (If present)

IBM Product ID (TTTT-MMM) or MTM. Refer to the [Help Page](#) for this field.

IBM Serial Number (7 characters)

Contact person for this subsystem

Storage Manager Password

Confirm Password

(May be required in order for RSM for Storage to be able to automatically collect a full set of log files. See [Help](#) for more information).

Part of an IBM Solution: **None** (Must be set by IBM Support. Refer to the Help for this page.)

### Drive Expansion Units

Expansion Unit ID: 0

\* IBM Product ID (TTTT-MMM) or MTM.

\* IBM Serial Number (7 characters)

[Update configuration](#)

To remove this subsystem from the configuration: [Delete this device](#)

To define a new subsystem based on this configuration: [Copy this form](#)

Figure 16-20 RSM Configure Storage Subsystem continued

Table 16-3 shows a possible selection of the IBM Product ID.

Table 16-3 RSM Product IDs used for RSM

| Machine Type | Model | HVEC PN | Description                                                        |
|--------------|-------|---------|--------------------------------------------------------------------|
| 1746         | C2A   | 1746A2S | IBM System Storage DS3512 Express Single Controller Storage System |



| Machine Type | Model | HVEC PN | Description                                                        |
|--------------|-------|---------|--------------------------------------------------------------------|
| 1746         | C2A   | 1746A2D | IBM System Storage DS3512 Express Dual Controller Storage System   |
| 1746         | E2A   | 1746A2E | IBM System Storage EXP3512 Express Storage™ Expansion Unit         |
| 1746         | C4A   | 1746A4S | IBM System Storage DS3524 Express Single Controller Storage System |
| 1746         | C4A   | 1746A4D | IBM System Storage DS3524 Express Dual Controller Storage System   |
| 1746         | E4A   | 1746A4E | IBM System Storage EXP3524 Express Storage Expansion Unit          |


17. The storage subsystem will be added to the list of configured storage subsystems as shown in Figure 16-21. Up to 50 storage subsystems and SAN switches can be added. Click **Configuration** to return to the System Configuration.

## IBM Remote Support Manager for Storage

[Main](#) > [Configuration](#)
[Logout](#) - [Refresh](#) - [Help](#)

### Configure Storage Subsystems

 [System:](#) **Incomplete**
 [Firewall:](#) **Disabled**

 [Reporting:](#) **All Subsystems**
[Remote Access:](#) **Disabled**

There are currently 1 storage subsystems configured:

View/Configure Firmware

- 1 [DS3524-FC](#) 07.70.16.01
- 2 [Select to add](#)

Figure 16-21 RSM Configure Storage Subsystem continued

18. Now we want to add other SAN Devices. Click the **Other SAN Devices** button as seen in Figure 16-22 on page 501.

**IBM Remote Support Manager for Storage**

[Main](#) [Logout](#) - [Refresh status](#) - [Help](#)

---

### System Configuration

[System](#): Incomplete     
 [Firewall](#): Disabled  
[Reporting](#): All Subsystems     
 [Remote Access](#): Disabled

View and define information for: **rsm.ibm.com**    Type: **8810-BE2** Serial: **LKMTN29**

[Contact Information](#): OK  
[Company Information](#): OK  
[Connection Information](#): OK  
[Storage Subsystems](#): OK (1 subsystems currently defined)  
[Other SAN Devices](#): OK (0 switches currently defined)  
[System Activated](#): No  
[Options](#)

Figure 16-22 RSM Add other SAN Devices

19. Fill in the description and the Management IP Address for the SAN Device as shown in Figure 16-23. After entering the required Information, press **Update configuration**.

**IBM Remote Support Manager for Storage**

[Main](#) > [Configuration](#) > [Other SAN Devices](#) [Logout](#) - [Help](#)

---

### Other SAN Devices Information

[System](#): Incomplete     
 [Firewall](#): Disabled  
[Reporting](#): All Subsystems     
 [Remote Access](#): Disabled

\* Name  
 \* Management IP Address  
  
 To remove this switch from the configuration:

Figure 16-23 RSM Add other SAN Devices continued...

20. The newly added SAN Device is now added as shown Figure 16-24 on page 502.

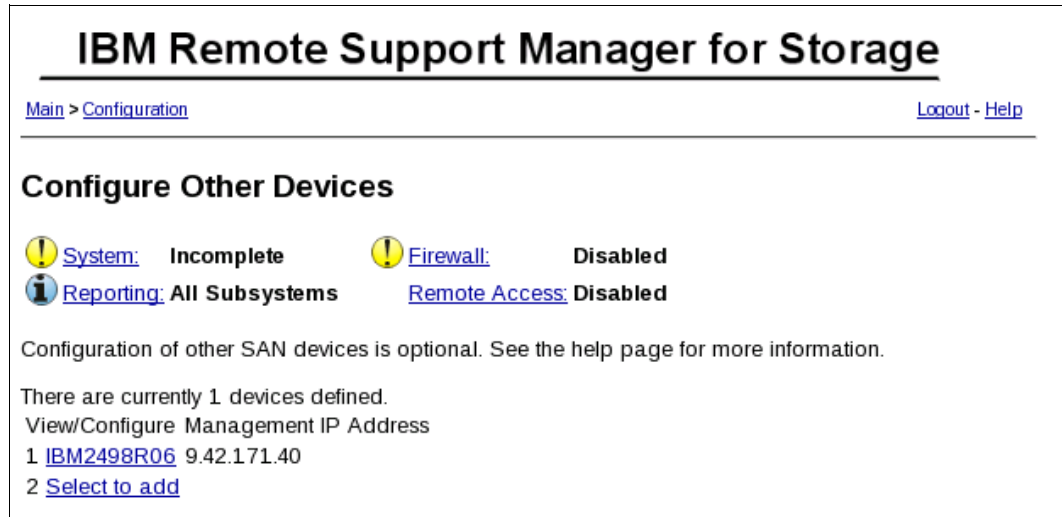


Figure 16-24 RSM Add other SAN Devices continued

21. After entering all the desired information, you can verify the information by clicking the **Run Configuration Test** button as shown in Figure 16-25 on page 503.



[Logout](#) - [Refresh status](#) - [Help](#)

## System Configuration

 [System:](#) Incomplete
  [Firewall:](#) Disabled  
 [Reporting:](#) All Subsystems
 [Remote Access:](#) Disabled

View and define information for: [rsm.ibm.com](http://rsm.ibm.com) Type: 8810-BE2 Serial: LKMTNZ9

|                                         |                                     |
|-----------------------------------------|-------------------------------------|
| <a href="#">Contact Information:</a>    | OK                                  |
| <a href="#">Company Information:</a>    | OK                                  |
| <a href="#">Connection Information:</a> | OK                                  |
| <a href="#">Storage Subsystems:</a>     | OK (1 subsystems currently defined) |
| <a href="#">Other SAN Devices:</a>      | OK (1 switches currently defined)   |
| <a href="#">System Activated:</a>       | No                                  |
| <a href="#">Options</a>                 |                                     |

Export a configuration for review and verification:

[Export to CSV](#)

Select "Export to CSV" to create a spreadsheet file of key configuration values in CSV format.  
This file will be created as `/var/log/rsm/rsm-config-info.csv`.

You can also email the file to someone by specifying their email address below:

## Configuration Test

Test was last run at: Tue Aug 10, 2010 18:05:51 UTC

A change has been made to the configuration. The Configuration Test should be re-run.

The last test found one problem .

```
2010 08 10 17:58:53 UTC - Configuration Test: Checking for connection to IBM ECC Problem Reporting s
2010 08 10 17:58:53 UTC - Configuration Test: Waiting for ECC Connection test to complete.
2010 08 10 18:05:51 UTC - Configuration Test: ECC Connection test timeout.
```

### Run Configuration Test

Figure 16-25 RSM Run Configuration Test


22. Click **Refresh status** to see the progress of the test as shown in Figure 16-26 on page 504.


# IBM Remote Support Manager for Storage


[Main](#)
[Logout](#) - [Refresh status](#) - [Help](#)

---

## System Configuration

 [System:](#) **Incomplete**

 [Reporting:](#) **All Subsystems**

 [Firewall:](#) **Disabled**

[Remote Access:](#) **Disabled**

View and define information for: **rsm.ibm.com**    Type: **8810-BE2** Serial: **LKMTN29**

[Contact Information:](#)    **OK**

[Company Information:](#)    **OK**

[Connection Information:](#) **OK**

[Storage Subsystems:](#)    **OK** (1 subsystems currently defined)

[Other SAN Devices:](#)      **OK** (1 switches currently defined)

[System Activated:](#)        **No**

[Options](#)

Export a configuration for review and verification:

[Export to CSV](#)

Select "Export to CSV" to create a spreadsheet file of key configuration values in CSV format. This file will be created as `/var/log/rsm/rsm-config-info.csv`.

You can also email the file to someone by specifying their email address below:

---

## Configuration Test

Test is currently running and is 27% complete. ([Refresh status](#) to see results)

```

2010 08 10 21:53:40 UTC - Configuration Test: NOMODEM specified, skipping modem checks
2010 08 10 21:53:40 UTC - Configuration Test: Checking for connection to IBM ECC Problem Reporting s
2010 08 10 21:53:40 UTC - Configuration Test: Waiting for ECC Connection test to complete.

```

Figure 16-26 RSM Run configuration Test

23. The results from this test are logged in the activity log. You can access the activity log by using 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 Example 16-5.

*Example 16-5 Activity log example*

```

2010 09 04 18:28:55 UTC - Starting Configuration Test...
2010 09 04 18:28:55 UTC - Configuration Test: Checking for Modem
2010 09 04 18:28:55 UTC - Configuration Test: NOMODEM specified, skipping modem
checks
2010 09 04 18:28:55 UTC - Configuration Test: Setting all devices in Temporary
Service Access mode for duration of test.
2010 09 04 18:28:55 UTC - Configuration Test: Checking connectivity to
Management Station at 9.42.171.25

```


2010 09 04 18:28:55 UTC - Configuration Test: Checking connectivity to DS3524\_FC  
 2010 09 04 18:28:55 UTC - Configuration Test: Validating subsystem name for DS3524\_FC  
 2010 09 04 18:28:55 UTC - Configuration Test: Waiting for results of SMcli: "show storagesubsystem profile"  
 2010 09 04 18:29:03 UTC - Configuration Test: Checking results of SMcli command  
 2010 09 04 18:29:03 UTC - History file has been updated  
 2010 09 04 18:29:03 UTC - Found 1 drive expansion boxes for DS3524\_FC  
 2010 09 04 18:29:03 UTC - Adding new expansion box (0) to configuration  
 2010 09 04 18:29:03 UTC - Configuration Test: No problems detected.  
 2010 09 04 18:29:03 UTC - Configuration Test: Restoring firewall to original state.


24. After the configuration is updated, go to the Main menu, select **System**, and run the configuration test. When the test is complete, the date and time of the last run is shown as shown in Figure 16-27.


# IBM Remote Support Manager for Storage

[Main](#)
[Logout](#) - [Refresh status](#) - [Help](#)

## System Configuration


[System:](#) Incomplete


[Firewall:](#) Disabled


[Reporting:](#) All Subsystems

[Remote Access:](#) Disabled

View and define information for: **rsm.ibm.com**    Type: **8810-BE2** Serial: **LKMTN29**

[Contact Information:](#) OK

[Company Information:](#) OK

[Connection Information:](#) OK

[Storage Subsystems:](#) OK (1 subsystems currently defined)

[Other SAN Devices:](#) OK (1 switches currently defined)

[System Activated:](#) No

[Options](#)

Export a configuration for review and verification:
 

Export to CSV

Select "Export to CSV" to create a spreadsheet file of key configuration values in CSV format. This file will be created as /var/log/rsm/rsm-config-info.csv.  
  
 You can also email the file to someone by specifying their email address below:

## Configuration Test

Test was last run at: Wed Aug 11, 2010 21:42:05 UTC

A change has been made to the configuration. The Configuration Test should be re-run.

Figure 16-27 RSM configuration test complete

## 16.2.4 Configuring SNMP traps in Storage Manager

To allow the DS storage subsystem management workstation to send SNMP traps to the RSM server, set the RSM server as your SNMP traps destination in the Storage Manager client.

To configure the Storage Manager to send SNMP alerts for each defined storage subsystem in RSM to the RSM host, perform these steps:

1. Open the Enterprise management window of the Storage Manager.
2. Right-click a DS storage subsystem. In the Connect menu, select **Configure Alerts**.
3. In the Configure Alerts window, click the **SNMP** tab.
4. Type the host name or the IP address of the RSM host in the trap destination field and click **Add** (Figure 16-28). Do not change the SNMP community name (the default is *public*).

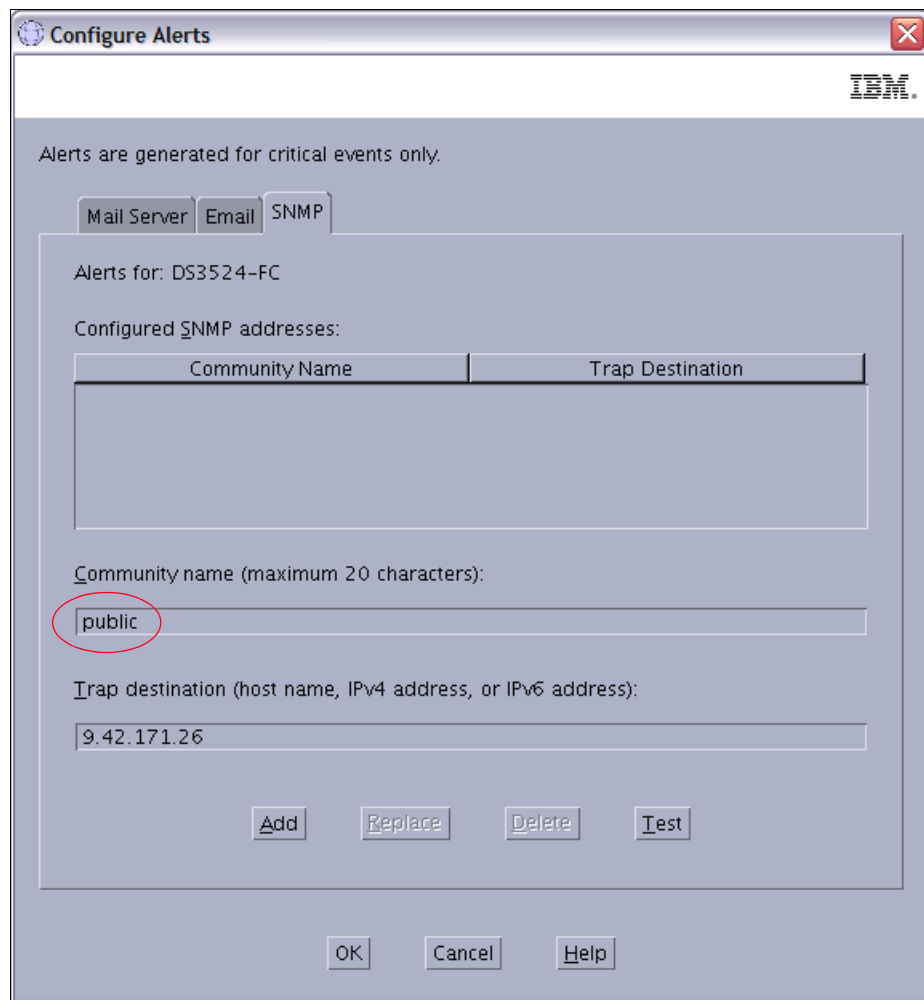


Figure 16-28 RSM Add SNMP trap receiver

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.

## Validate SNMP configuration

To validate the SNMP configuration, select the SNMP trap destination and click **Test** to send a test trap to the RSM host (Figure 16-29).

Configure Alerts

Alerts are generated for critical events only.

Mail Server Email **SNMP**

Alerts for: DS3524-FC

Configured SNMP addresses:

| Community Name | Trap Destination |
|----------------|------------------|
| public         | 9.42.171.26      |

Community name (maximum 20 characters):

public

Trap destination (host name, IPv4 address, or IPv6 address):

9.42.171.26

Add Replace Delete **Test**

OK Cancel Help

Figure 16-29 RSM Send test trap

Check the activity log (as shown in Figure 16-27 on page 505) and verify that the trap was received. The activity log will contain a entry as shown in Example 16-6.

### Example 16-6 Test alert received

---

2010 09 04 22:05:08 UTC - Received a test alert for DS3524-FC from 9.42.171.25

---

## 16.2.5 Activating RSM

The final step is to activate your system. Complete all the other configurations and run a successful Configuration Test before contacting IBM Service to activate RSM.

To activate RSM, perform these steps:

1. Make sure Remote Access is enabled. Click **Remote Access** in the main RSM window to activate it, if required as shown in Figure 16-30. For a complete description, see 16.2.6, “Remote access security” on page 509.

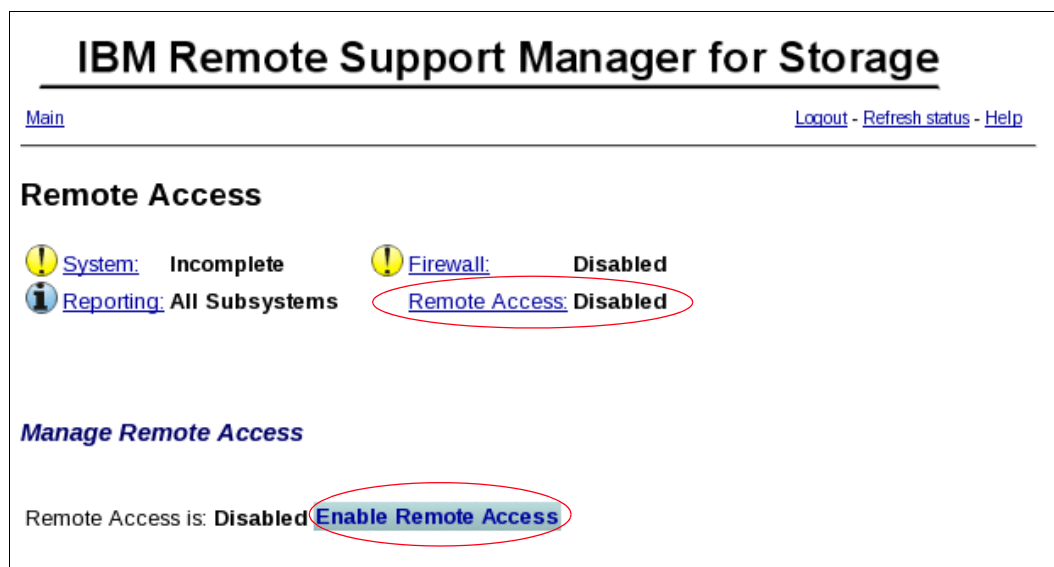


Figure 16-30 RSM Remote access

2. Call the number for IBM Service for your region and give the IBM Machine Type and Serial Number of one of the DS storage subsystems to be monitored by RSM for Storage. For support telephone numbers in your country or region, navigate to the following URL:  
<http://www.ibm.com/planetwide>
3. Tell the disk support contact person that you are activating an RSM for Storage system.
4. Provide IBM Service with the phone number of the modem attached to the RSM for Storage system (if used). IBM Service will connect to the system, verify that the configuration is correct, send a test alert through email, verify receipt of the alert and associated attachments, and then activate the system.

When the RSM for Storage system is ready to receive events, the System Status will be OK. See Figure 16-31 on page 509.

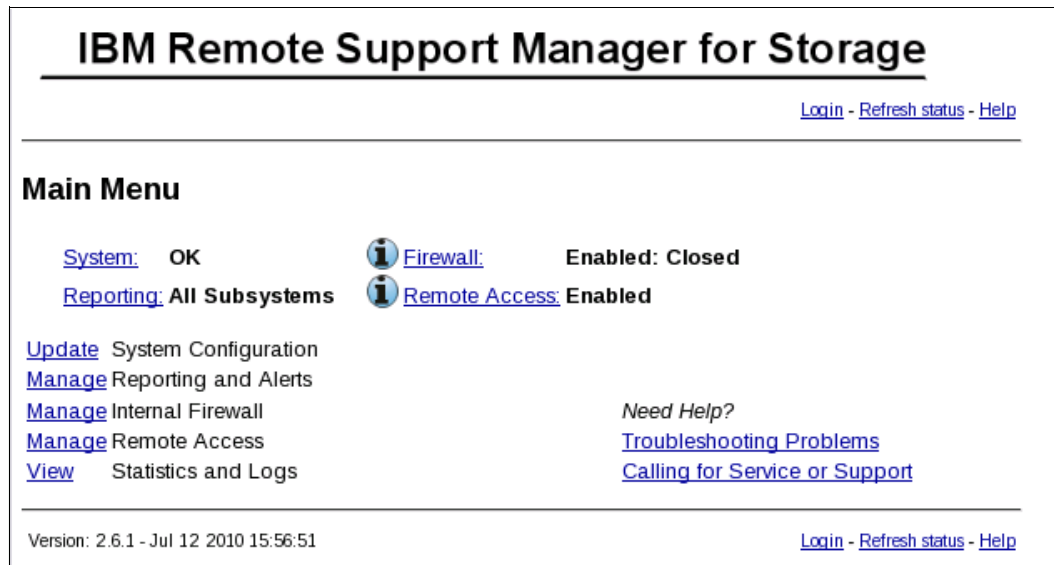


Figure 16-31 RSM main menu: Configured

**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.

## 16.2.6 Remote access security

This section will cover how to secure your IBM System Storage DS subsystem from unauthorized access.

### SSH connectivity

In the RSM for Storage Remote Access window, click **Enable Remote Access**. This will reconfigure the RSM for Storage internal firewall to allow connections through SSH port 22.

Verify connectivity using the following steps:

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 might 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.

You should be connected to the RSM for Storage system and receive a login prompt.

**Note:** You will not be able to complete the login. Authentication requires a special tool only available to IBM Remote Support.

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.

3. Verify that the SSH connection is correctly configured on the Connections Configuration page. This information will be encrypted and sent with each alert to IBM.

## 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 window, click **Enable Remote Access**. This will enable the modem to answer when called. Verify modem connectivity by calling the modem phone number from a voice phone:

1. Most modems will either flash an LED or you might 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.

To configure the Remote Access policy, perform the following steps:

1. Click **Remote Access** from the Main Menu as shown in Figure 16-31 on page 509.
2. In the Remote Access setting window, 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 16-32 on page 511.



# IBM Remote Support Manager for Storage

[Main](#)

[Logout](#) - [Refresh status](#) - [Help](#)

## Remote Access

[System:](#) OK



[Firewall:](#)

Enabled: Closed

[Reporting:](#) All Subsystems



[Remote Access:](#)

Enabled

### Manage Remote Access

Remote Access is: **Enabled** [Disable Remote Access](#)

The Remote Access Timeout is: **11:09 (hh:mm)**

Select one of the following to change the current (and default) Remote Access Timeout:

☐ 12 hours ☐ 24 hours ☒ 36 hours  
☐ 48 hours ☐ 72 hours ☐ 96 hours

[Update Timeout Value](#)

### Option to Enable Remote Access on Alert

You can choose to have the RSM for Storage system automatically enable remote access when an alert is sent to IBM. If this option is disabled, IBM Service will contact you and request that Remote Access be manually enabled so they can access the RSM for Storage system.

The option to automatically enable remote access when an alert occurs is: **Enabled**

[Disable remote access on alert](#)

### Additional Information:

If you need to send the modem phone number or SSH connection information to IBM in an email, you can send this string instead of the actual information:

**DD06F-1FEE1-6DD11-AA001-111EE-16D5A-A6742-B9D56-22A3A-DBDB2-A24E5-DB3A9-C7E2D-3A916-242B0-F8669-8C9EF-A446E-F218F-16D5A-A45A3-C3CBD-62E51-F37C9-00**

IBM Service does not need to be sent the remote login password, it is provided here so you can test the remote login function. Note however that following a successful login, a Challenge/Response (that can only be answered correctly using an internal IBM system) will block further access.

Remote user: **rservice**, password: **3vn37GTe**

Figure 16-32 Remote Access settings

3. 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, scroll down the Remote Access settings window, select the desired timeout value, and click **Update Timeout Value**, as shown in Figure 16-32 on page 511.

**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.

- ▶ The normal state for the firewall is Enabled:Closed, which means that the firewall is operational and configured to allow SNMP traps to be received and emails to be sent. However, access to other devices on your network is not allowed.
- ▶ The Enabled: Custom state indicates that one or more custom rules have been added to `/etc/rsm/rsm-firewall.ibm.conf`. These rules will be active any time the firewall is enabled.
- ▶ The Enabled: Open state means that access to one or more other devices has been enabled. The firewall allows access to any storage subsystem that has an active alert, and also storage subsystems and other SAN devices that have been placed in Service Access mode.

Service Access mode allows you to manually allow access to a device from the RSM for Storage system. You can select storage subsystems that you have previously configured.

Disabling the firewall allows unrestricted access from the RSM for Storage system to your network. To maintain the security of your network, disabling the firewall also disables remote access. Likewise, enabling Remote Access will automatically enable the firewall.

**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, perform the following steps:

1. Click **Firewall** on the Main Menu, as shown in Figure 16-31 on page 509.
2. In the Internal Firewall and Service Access window, you can change the internal firewall status and service access mode of your managed storage subsystems as shown in Figure 16-33 on page 513.

# IBM Remote Support Manager for Storage

[Main](#)

[Logout](#) - [Refresh status](#) - [Help](#)

## Internal Firewall and Service Access

[System:](#) OK



[Firewall:](#)

Enabled: Closed

[Reporting:](#) All Subsystems



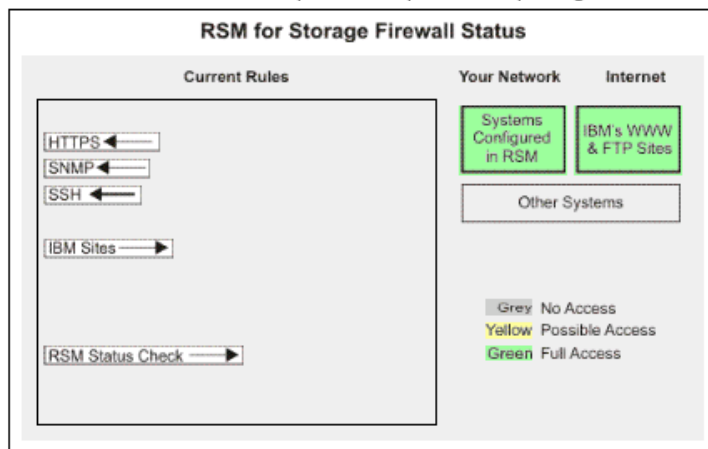
[Remote Access:](#) Enabled

### Manage Internal Firewall

Firewall is **Enabled** - All connections are blocked except those required for reporting.

Note: Disabling the firewall will also disable remote access.

[Disable Firewall](#)



### Manage Service Access

Any configured device may be placed in temporary in Service Access mode. This creates a rule for the internal firewall that will allow connections to that device from the RSM for Storage system. These temporary firewall rules will be removed when the Service Access Timeout expires.

[Manage](#) Service Access for storage subsystems. (0 - currently enabled)

[Manage](#) Service Access for other SAN devices (0 - currently enabled)

[Disable Service Access for all devices](#)

The Service Access Timeout is: **N/A**

Select one of the following to change the current (and default) Service Access Timeout:

- ☐ 12 hours   ☐ 24 hours   ☒ 36 hours   ☐ 48 hours   ☐ 72 hours   ☐ 96 hours   [Update Timeout Value](#)

Figure 16-33 Internal Firewall and Service Access window

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.

3. 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 window, select the desired Service Access Timeout value, and click **Update Timeout Value** as shown in Figure 16-33 on page 513.

### IBM System Storage DS security

The Storage Manager has the ability to require an administrative password to make changes to the subsystem configuration. We suggest configuring this password.

The IBM System Storage DS subsystems also have a controller shell environment that is accessible using a remote login (RLOGIN) client. Storage Manager has an option to disable RLOGIN, and we normally suggest disabling RLOGIN.

## 16.2.7 Managing alerts

To manage alerts, perform these steps:

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 16-34. Click **Reporting** to see the alerts.

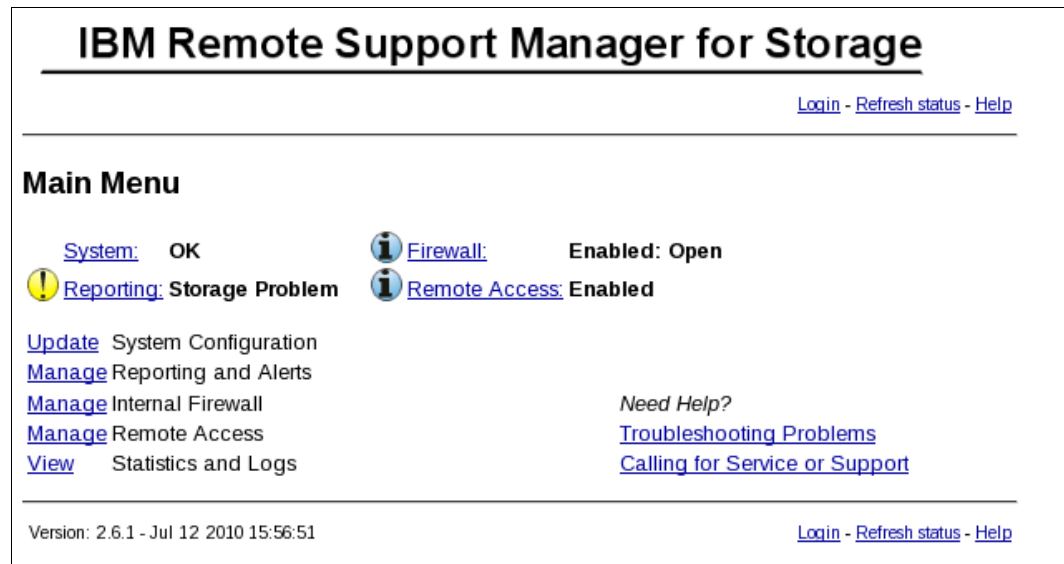


Figure 16-34 Main Menu: Storage Problem


2. The reporting and alert window shows all those subsystems that have sent alerts sent to RSM (Figure 16-35 on page 515). Click **View or Close Alerts**.

IBM Remote Support Manager for Storage


[Main](#)
[Logout](#) - [Refresh status](#) - [Help](#)

Reporting and Alerts


System: OK



Reporting: Storage Problem



Firewall: Enabled: Open



Remote Access: Enabled

Products with Active Alerts: 1

Alerts sent to IBM: 1

Products with Reporting Enabled: 1

Alerts Pending: 0

Products with Reporting Disabled: 0

Alerts Acknowledged: 0

[View/Change](#) reporting state for each subsystem
   
[View Closed Alerts](#) ( 0 )

Active Alerts:

Sent Acknowledged Pending Firmware

[View or Close Alerts](#) for: **DS3524-FC**

10007.70.16.01

[Main](#)
[Logout](#) - [Refresh status](#) - [Help](#)

Figure 16-35 RSM storage subsystem with alerts


- The alert list for the selected storage subsystem (Figure 16-36) shows all alerts that were received by RSM. Click **View** to see the details of an alert.

IBM Remote Support Manager for Storage


[Main](#) > [Reporting and Alerts](#)
[Logout](#) - [Refresh Status](#) - [Help](#)

Alert List for DS3524-FC


System: OK



Reporting: Storage Problem



Firewall: Enabled: Open



Remote Access: Enabled

Subsystem: DS3524-FC, 1.1.1.2 1.1.1.3

IBM Type/Serial: 1746-C4A / NONUMBE

Firmware Version: 07.70.16.01 (as of Tue Aug 10, 2010 18:00:45)

Location: LAB, LAB Street, Downtown LAB, 12345, United States

Contact: Max Mustermann

Total alerts for this subsystem: 1

NOTE: Do not close alerts unless IBM Service has contacted you about the problem. See the help page for more information.

|                                              | Date and Time                 | State     | Duplicates | Event Code |
|----------------------------------------------|-------------------------------|-----------|------------|------------|
| <a href="#">View</a> - <a href="#">Close</a> | Fri Aug 13, 2010 17:06:57 UTC | Submitted | 0          | 280b       |

[Close all active alerts for DS3524-FC](#)

[Main](#) > [Reporting and Alerts](#)
[Refresh Status](#) - [Help](#)

Figure 16-36 RSM list of alerts

Chapter 16. IBM Remote Support Manager for Storage 515

4. The alert details and an error message are shown in Figure 16-37.

## IBM Remote Support Manager for Storage

[Main](#) > [Reporting and Alerts](#) > [DS3524-FC](#)[Logout](#) - [Refresh Status](#) [Help](#)

---

### Alert Details

[System:](#) OK

 [Firewall:](#) Enabled: Open

 [Reporting:](#) Storage Problem

 [Remote Access:](#) Enabled

Subsystem Name:

DS3524-FC

Type and Model:

1746-C4A

Serial number:

NONUMBE

Alert Status:

Submitted

IBM Service logged in?

No

Time of alert:

Fri Aug 13, 2010 17:06:57 UTC

Sent to IBM at

not available

Duplicates:

0

Time of last duplicate:

N/A

Alert Event Code:

280b

Alert Description:

Controller enclosure component failed.  
Power Supply.  
Enclosure 85, Power-fan canister 2, Slo

Log File Available:

Log files could not be downloaded

NOTE: Do not close alerts unless you have talked with IBM Service

[Close this alert](#) (Service is complete)

[Main](#) > [Reporting and Alerts](#) > [DS3524-FC](#) >

[Refresh Status](#) [Help](#)

Figure 16-37 RSM alert details

5. The alert list of the selected storage subsystem shows the modified status of an alert as shown in Figure 16-36 on page 515. When the problem is solved, click **Close**.
- After the alert is closed, it disappears from the alert list. The main menu status changes so that the attention mark disappears after all problems are solved (Figure 16-38 on page 517).


# IBM Remote Support Manager for Storage

[Main](#)[Logout](#) - [Refresh status](#) - [Help](#)


---

## Reporting and Alerts

[System:](#) OK

 [Firewall:](#) Enabled: Closed

[Reporting:](#) All Subsystems

 [Remote Access:](#) Enabled

Products with Active Alerts: 0

Alerts sent to IBM: 0

Products with Reporting Enabled: 1

Alerts Pending: 0

Products with Reporting Disabled: 0

Alerts Acknowledged: 0

[View/Change](#) reporting state for each subsystem

There are no active alerts.

[View Closed Alerts \( 1 \)](#)

[Main](#)[Logout](#) - [Refresh status](#) - [Help](#)

Figure 16-38 RSM main menu







## Command-Line Interface (CLI)

In this chapter, we explain the basics of the Command-Line Interface (CLI) that can be used to manage an IBM System Storage DS3500. Along with the DS3500, the CLI can manage older DS3000 models, and DS4000 and DS5000 storage subsystems as well.

The command-line interface (CLI) is a software tool that lets storage subsystem installers and engineers configure and monitor storage subsystems using script commands. Using the CLI, you can run commands from an operating system prompt, such as the Windows command prompt, a Linux operating system console, or a Solaris operating system console.

You have to install the IBM DS Storage Manager client to run the script commands either through the script window, which is invoked from the IBM DS Storage Manager client Enterprise window, or through the command line interface using the SMcli program. The script command engine is automatically installed as part of the IBM DS Storage Manager client installation.

The CLI 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_DS\client` and in Linux `/opt/IBM_DS/client/`. The CLI program is called SMcli in Windows and Linux. Commands that can be used on the CLI to administer DS3500, DS4000, and DS5000 storage subsystems are identical to the commands used in the Script Editor. Not all commands apply to all models of the storage subsystem.

The CLI gives access to all the functions provided in the Storage Manager GUI, and some additional management functions. Each command performs a specific action for managing a storage subsystem or returning information about the status of a storage subsystem. You can enter individual commands, or you can run script files when you need to perform operations more than once. For example, you can run script files when you want to install the same configuration on several storage subsystems. The CLI allows you to load a script file from a disk and run the script file. The CLI provides a way to run storage management commands on more than one network storage subsystem. 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 chapter is organized in the following sections:

- 17.1, “How to Use the Command Line Interface” on page 520

- ▶ 17.2, “Running the CLI” on page 521
- ▶ “General SMcli syntax” on page 523
- ▶ “Adding a storage subsystem to the Storage Manager configuration” on page 527
- ▶ “Showing defined subsystems in the Storage Manager configuration” on page 528
- ▶ “Configuring alerts” on page 529
- ▶ “Issuing commands to the storage subsystem” on page 532

## 17.1 How to Use the Command Line Interface

The commands that you run on the CLI provide access to the script engine, specify the storage subsystem to receive the script commands, and set operation environment parameters.

A CLI command consists of these elements:

- ▶ The term **SMcli**
- ▶ The storage subsystem identifier
- ▶ Parameters
- ▶ Script commands

A CLI command takes this form:

**SMcli storageSubsystem parameters script-commands;**

- ▶ *SMcli* invokes the command line interface.
- ▶ *storageSubsystem* is the name or the IP address of the storage subsystem.
- ▶ *parameters* are CLI parameters that define the environment and the purpose for the command.
- ▶ *script-commands* are one or more script commands or the name of a script file that contains script commands. The script commands are the storage subsystem configuration commands.

### 17.1.1 Usage Notes

If you enter **SMcli** and a storage subsystem name but do not specify CLI parameters, script commands, or a script file, the command line interface runs in interactive mode. Interactive mode lets you run individual commands without prefixing the commands with **SMcli**. In interactive mode, you can enter a single command, view the results, and enter the next command without typing the complete **SMcli** string. Interactive mode is useful for determining configuration errors and quickly testing configuration changes.

To end an interactive mode session, press the termination command (Ctrl+C) while in interactive mode.

If you enter an incomplete or inaccurate **SMcli** string that does not have the correct syntax, parameter names, options, or terminals, the script engine returns usage information.

## 17.2 Running the CLI

There are two ways to enter commands into the DS3500: interactively using the SMcli, or by creating saved scripts with the Scripting Editor. We describe SMcli in 17.3, “General SMcli syntax” on page 523.

### 17.2.1 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.

To create a file using the Script Editor, perform these steps:

1. To start the Script Editor, from Storage Manager, click **Tools** → **Execute Script** as shown in Figure 17-1.

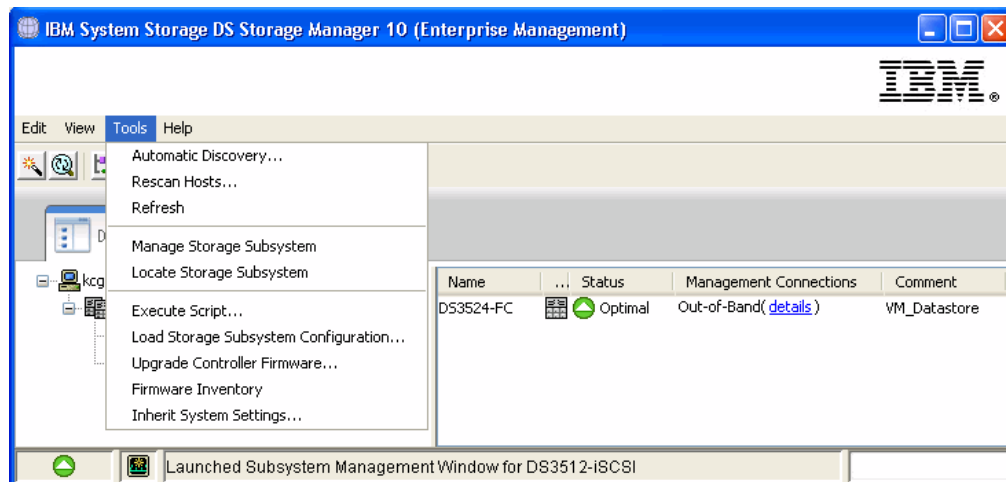


Figure 17-1 Invoke Script Editor

2. A blank Script Editor appears as shown in Figure 17-2. From here you can enter commands and save them for later execution.




Figure 17-2 Script Editor window

3. Click **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 17-3 on page 522.

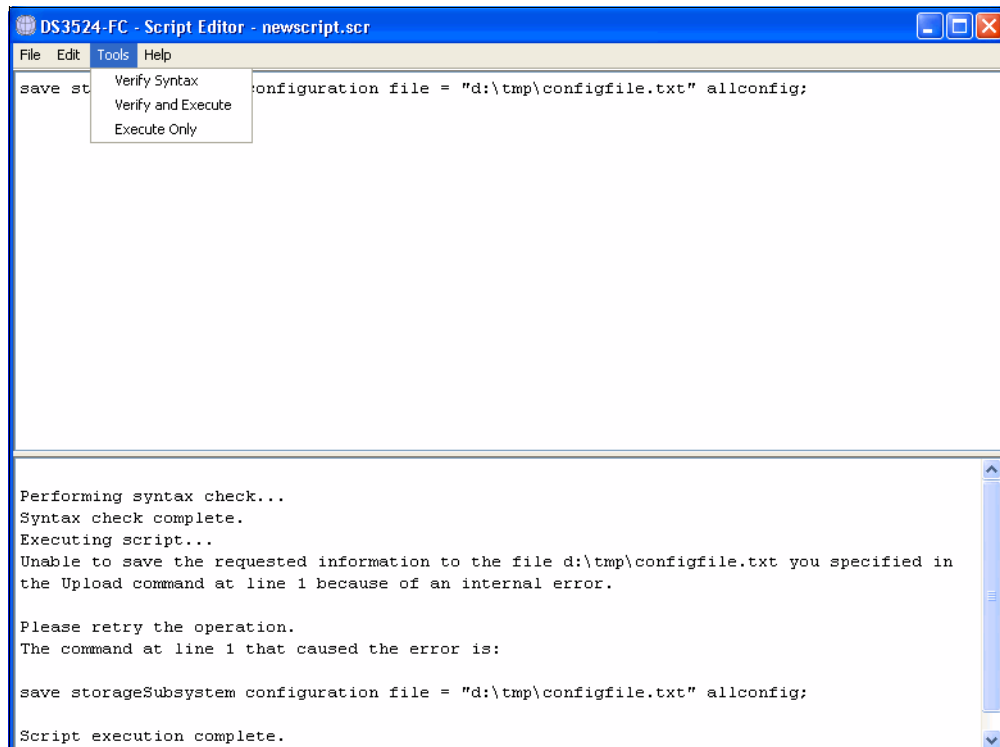


Figure 17-3 Tools menu

4. The Script Editor contains an online help that describes the use of each of the available commands. Click **Help** → **Command Reference** to display the online help (Figure 17-4) or **Help** → **Overview** for help on using the Script Editor itself.

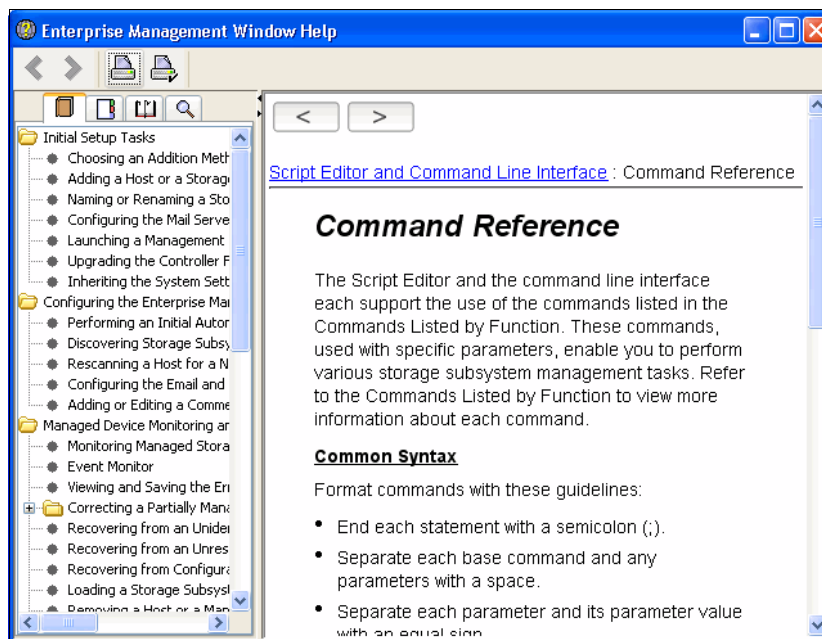


Figure 17-4 Command reference

## 17.3 General SMcli syntax

The following table lists the conventions used in the command statement syntax to manage the storage subsystem

Table 17-1 Command Name Syntax Conventions

| Convention                | Definition                                                                                        |
|---------------------------|---------------------------------------------------------------------------------------------------|
| <b>(a   b   c)</b>        | Alternative ("a" or "b" or "c")                                                                   |
| <i>italicized-words</i>   | A terminal that needs user input to fulfill a parameter (a response to a variable).               |
| [ ... ] (square brackets) | Zero or one occurrence (square brackets are also used as a delimiter for some command parameters) |
| { ... } (curly braces)    | Zero or more occurrences.                                                                         |
| <b>bold</b>               | A terminal that needs a command parameter entered to start an action                              |

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 17-1.

**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 17-1 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 email alert.
- ▶ Configure email alerts.
- ▶ Configure or remove SNMP alerts.

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 DS3500 controllers.

Table 17-2 on page 525 gives an explanation of the various parameters that can be specified with SMcli commands.

Table 17-2 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.<br>IP address or fully qualified domain name of the host running the Storage Manager agent for in-band management.<br>Use either IP addresses or host names, the subsystem name, or WWID                                                                                                                                                                                                                                                                                               |
| -A<br>-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.                                                                                                                                                                                                                                                                                                                                                                                |
| -a                                             | Use this parameter to add a Simple Network Management Protocol (SNMP) trap destination or an email address alert destination.<br><ul style="list-style-type: none"> <li>▶ 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.</li> <li>▶ When adding an email address for an alert destination, the email address is the email address where you want the alert message to be sent.</li> </ul> |
| -c "<command;>"                                | Specify one or more commands separated by colons. Each command needs to be closed by a semicolon. Use either option -c or -f.                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| -d                                             | Shows the contents of the script configuration file. The file content takes this format:<br>storage-system-name host-name1 host-name2                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| -e                                             | Runs the commands without performing a syntax check first.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| -F <email addr>                                | The email address that will be used in the sent by field of an email alert. Responses to an alert will be sent to this address. Will be used for all emails. Use this parameter together with the parameter -m.                                                                                                                                                                                                                                                                                                                                                                      |
| -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.                                                                                                                                                                                                                                                                                                                                                                                         |
| -g                                             | Specify a plain text file that contains email sender contact information that will be included in all email alert notifications.                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| -h                                             | Specifies the host name that is running the SNMP agent to which the storage subsystem is connected.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| -l eventOnly<br>-l profile<br>-l supportBundle | Define what information is included in an alert that gets sent through email. Use one of these options:<br><ul style="list-style-type: none"> <li>▶ eventOnly</li> <li>▶ profile</li> <li>▶ supportBundle</li> </ul> By default, only the event data is sent.                                                                                                                                                                                                                                                                                                                        |
| -i                                             | Shows the IP address of the known storage subsystems. This option should only be used with parameter -d. The file contents takes this format:<br>storage-system-name IP-address1 IP-address2                                                                                                                                                                                                                                                                                                                                                                                         |
| -m <mail server>                               | Specifies the host name or the IP address of the email server from which email alert notifications are sent.                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| -n <subsystem name>                            | The name of the storage subsystem shown by the -d options. Use either IP or host names, subsystem name, or WWID.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

| Option                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| -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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| -q                       | Specifies the frequency that you want to include additional profile or support bundle information in the email alert notifications. An email alert notification containing at least the basic event information is always generated for every critical event. If you set the -l terminal to eventOnly, the only valid value for the -q terminal is every Event. If you set the -l terminal to either the profile value or the supportBundle value, this information is included with the emails with the frequency specified by the -q terminal.                                                                                                                                                                                                          |
| -r <direct_sa   host_sa> | Sets or changes the alert notifications for all of the storage subsystems that are managed by a specific storage management station. For storage subsystems that are managed out-of-band, use the direct_sa value. For storage subsystems that are managed in-band, use the host_sa value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| -S (uppercase)           | Suppresses informational messages describing the command progress that appear when you run script commands                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| -s (lowercase)           | Shows the alert settings in the configuration file when used with the -d terminal.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| -v                       | Show the current global status of the known devices in a configuration file when used with the -d.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| -w <WWID>                | Specifies the WWID of the storage subsystem. This terminal is an alternate to the -n terminal. Use the -w terminal with the -d terminal to show the WWIDs of the known storage subsystems. The file content take this format: storage-system-name world-wide-ID IPaddress1 IP-address2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| -w                       | This option should only be used with parameter -d. It shows the world wide identifier of the defined storage subsystems.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| -X (uppercase)           | Deletes a storage subsystem from a configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| -x (lowercase)           | Removes an SNMP trap destination or an email address alert destination. The community is the SNMP community name for the trap, and the host is the IP address or DNS host name of the system to which you want the trap sent.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| -quick                   | <p>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 <b>recreate snapshot volume</b> command. This parameter reduces time by not running background processes for the duration of the command.</p> <p>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.</p> <p>This parameter causes operations that depend on background information to fail.</p> |
| -?                       | Shows usage information about the CLI commands.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |



We now show specific examples of using the CLI for tasks that have already been demonstrated using the GUI earlier in this book.

## 17.4 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 528 for a detailed description of the command shown in Example 17-2.

*Example 17-2 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 17-3 shows both methods.

*Example 17-3 SMcli - Add storage subsystems to the Storage Manager configuration*

---

```
amazon:/ # SMcli -A 1.1.1.2 1.1.1.3
New storage subsystem was discovered at address 1.1.1.2

New storage subsystem was discovered at address 1.1.1.3

SMcli completed successfully.

amazon:~ # SMcli -A
Starting auto discovery.
.....
.....
Auto discovery operation successful.
SMcli completed successfully.

amazon:~ #
```

---

## 17.5 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 17-4. The output shows that there is a global alert in place that sends email alerts to email address dsx000rivers.local. The mail server nile.rivers.local is used and a return address ds3500@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 DS3512i.

SNMP alerts are shown for the DS3524-fc. 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 17-4 Displaying details about defined storage subsystems*

---

```
amazon:~ # SMcli -d -i -s -w -v -S
[MAIL SERVER]
nile.rivers.local ds3500@nile.rivers.local

[ALERT SETTINGS - DEVICE GROUPS]

All storage subsystems
 DS3500@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:

DS3524-fc 600a0b80001d2bf40000000046265018 1.1.1.2 Optimal
 public,nile.rivers.local

DS3512i 092ca896040404040404040404040404 1.1.1.4 Needs Attention
 <None>

amazon:~ #
```

---

## 17.6 Configuring alerts

This section describes how to manage alert recipients.

### 17.6.1 Defining the mail server and email address to send out the email alerts

To define the global setting for the email server and email return address, use the following command:

```
SMcli -m <IP or host name of mail server> -F <email address>
```

See Example 17-5 for more details.

To verify the setting, use **SMcli -d -s** to see the current settings.

*Example 17-5 SMcli - define email server and return email address*

---

```
amazon:~ # SMcli -m nile.rivers.local -F DS3500@nile.rivers.local
SMcli completed successfully.
```

```
amazon:~ #
```

---

To delete these settings, use empty strings as shown in Example 17-6.

*Example 17-6 SMcli - delete mail server and return email address setting*

---

```
amazon:~ # SMcli -m "" -F ""
SMcli completed successfully.
```

```
amazon:~ #
```

---

### 17.6.2 Defining email alert recipients

Alerts can be defined for a single DS3500 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 email server should be defined as explained in 17.6.1, “Defining the mail server and email address to send out the email alerts” on page 529. Without a mail server, the local system running the Storage Manager will be used to send email alerts, which might not always be possible.

Example 17-7 shows how to define an email alert recipient for a single DS3500 server.

*Example 17-7 SMcli - define email alert recipient for single system*

---

```
amazon:~ # SMcli -a email:ds3500@nile.rivers.local -n "DS3500"
SMcli completed successfully.
```

```
amazon:~ #
```

---

To define an email 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 17-8 shows an email recipient for all in-band managed storage subsystems.

---

*Example 17-8 SMcli - E-mail recipient for a group of storage subsystems*

---

```
amazon:~ # SMcli -a email:ds3500-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 17-9.

---

*Example 17-9 SMcli - E-mail alert recipient for all subsystems*

---

```
amazon:~ # SMcli -a email:ds3500@nile.rivers.local -r host_sa
SMcli completed successfully.
```

```
amazon:~ #
```

---

Global email alert recipients are configured by omitting the target specification like **-r host\_sa** or a host name. Example 17-10 shows a global email recipient configuration.

---

*Example 17-10 SMcli - configure global email recipient*

---

```
amazon:~ # SMcli -a email:ds3500@nile.rivers.local
SMcli completed successfully.
```

```
amazon:~ #
```

---

### 17.6.3 Deleting email alert recipients

To delete an email alert recipient, use the option **-x** and the same syntax that was used to add the email alert recipient. Example 17-11 shows the commands to remove email 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 email 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 email recipient.

---

*Example 17-11 SMcli - Delete email recipients*

---

```
amazon:~ # SMcli -d -s
[MAIL SERVER]
nile.rivers.local ds3500@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:

DS3524 ds3524-a.rivers.local ds3524-b.rivers.local
 <None>
DS3512 ds3512-a.rivers.local ds35120-b.rivers.local
 <None>

DS3524i colorado.rivers.local
 host@nile.rivers.local

DS3512i colorado.rivers.local
 host@nile.rivers.local

SMcli completed successfully.

amazon:~ # SMcli -x email:ds3524@nile.rivers.local -n DS3524
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:~ #

```

---

## 17.6.4 SNMP alert recipients

SNMP alert recipients are handled similarly to email alert recipients, but the trap parameter is used instead of the email parameter. To specify the recipient, use an SNMP community string followed by the SNMP trap receiver systems IP address or host name. Example 17-12 shows how to configure SNMP trap recipients for a single DS3524 system, in-band and out-of-band managed DS3500 systems, and a global trap recipient.

*Example 17-12 SMcli - SNMP alert recipient*

---

```

amazon:~ # SMcli -a trap:ds3524,nile.rivers.local -n ds3524
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:~ #
```

---

## 17.7 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 17.3, “General SMcli syntax” on page 523 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 17-13 shows a command to a storage subsystem to get the health status of that storage subsystem.

### *Example 17-13 SMcli - Command*

---

```
amazon:/ # SMcli -n DS3524 -c "show storagesubsystem healthstatus;" -S
The following failures have been found:
Failed Drive - Unassigned or Hot Spare
Storage Subsystem: DS3524
 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 17-14 shows a script file that creates logical drives and assigns them to a host.

### *Example 17-14 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 (in our example, as `createDrive.cmd`) and run it as shown in Example 17-15.

*Example 17-15 SMcli executes the script*

---

```
amazon:~ # SMcli -n DS3524 -f createDrive.cmd -o hudson.out -S
amazon:~ #
```

---

Certain commands require the input of string values, such as the command shown in Example 17-16.

*Example 17-16 SMcli - Character masking*

---

```
amazon:~ # SMcli -n DS3524 -c "set logicalDrive [\"0rinoco-1\"]
userLabel=\"0rinoco-2\";" -S -e
amazon:~ #
```

---

This command will rename the logical drive Orinoco-1 to 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 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 DS3500 storage subsystem:

- ▶ **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.

### 17.7.1 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 17-17.

*Example 17-17 CLI command to save the storage subsystem configuration*

---

```
save StorageSubsystem configuration file="DS3524_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 17-18.

*Example 17-18 Sample storage subsystem configuration file*

---

```
// Logical configuration information from Storage Subsystem DS3524_crmsystem_data.
// Saved on 21 August, 2010
// Firmware package version for Storage Subsystem DS3400_crmsystem_data =
07.70.16.01
// NVSRAM package version for Storage Subsystem DS3524_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 DS3524_crmsystem_data.";
set storagesubsystem userLabel="DS3524_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 2003/Server 2008 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 2003/Server 2008 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 161. Note that loading the configuration file will overwrite any existing data on the array. You should only do this on a new array.

## 17.8 More Information

For detailed information about all the CLI parameters, consult the Command Line reference included in the SMclient online help or the guide *IBM System Storage DS3000, DS4000, and DS5000 Command Line Interface and Script Commands Programming Guide* available at the below IBM support website.

<http://www-947.ibm.com/systems/support/supportsite.wss/docdisplay?ln docid=MIGR-5076792&brandind=5000028>



## Windows SAS configuration guide for IBM BladeCenter

In this chapter, we describe and explain a sample configuration that shows how to connect logical drives configured on an IBM System Storage DS3500 to a Windows Server 2008 operating system running on an IBM BladeCenter HS21 server connected with SAS Connectivity Modules.

Refer to the following documents for the FC, and iSCSI Windows host attachment guides.

- ▶ *IBM Midrange System Storage Implementation and Best Practices Guide*, SG24-6363
- ▶ IBM System Storage DS Storage Manager Version 10 - Installation and Host Support Guide available at the following URL:

<http://www-947.ibm.com/support/entry/portal/docdisplay?brand=5000028&indocid=MI-GR-5075652>

## 18.1 Equipment required

We used the following hardware and software components for this configuration:

- ▶ IBM BladeCenter E
- ▶ IBM BladeCenter HS21 with Windows Server 2008 installed on local SAS disk
- ▶ TIBM BladeCenter Advance Management Module (AMM) installed in BladeCenter
- ▶ Two IBM BladeCenter SAS Connectivity Modules
- ▶ IBM BladeCenter dual port SAS Expansion card (CFFv)
- ▶ A DS3500 System Storage subsystem
- ▶ Two SAS cables
- ▶ The latest version of DS Storage Manager running on an external management workstation (at the time of writing, this is Version 10.70)

Figure 18-1 shows our sample hardware setup.




Figure 18-1 Hardware setup

## 18.2 IBM BladeCenter setup

Check the IBM support website for the latest DS3500 interoperability matrix to ensure that all hardware used for this configuration are supported. Ensure that the BladeCenter, blade server, and modules are updated with the latest firmware. Consult the IBM support websites for the latest firmware.

## 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 websites.

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?ln docid=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. Although this is not a difficult task, be sure to 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 the following URL:

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

**Note:** The BladeCenter SAS Expansion Card is a dual port card. Port #1 connects to the SAS Connectivity Module in BladeCenter I/O module bay 3 and port #2 connects to the SAS Connectivity Module in BladeCenter module bay 4 respectively.

## 18.2.3 Recording the SAS Expansion Card WWPN

The following steps demonstrate how to record the SAS Expansion Card WWPN for later use in setting up the host-to-LUN mappings in the DS3500 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 on page 540.

```
Broadcom NetXtreme II Ethernet Boot Agent v3.4.8
Copyright (C) 2000-2007 Broadcom Corporation
All rights reserved.

Broadcom NetXtreme II Ethernet Boot Agent v3.4.8
Copyright (C) 2000-2007 Broadcom Corporation
All rights reserved.

LSI Corporation MPT SAS BIOS
MPTBIOS-6.26.02.00 (2009.06.08)
Copyright 2000-2009 LSI Corporation.

Press Ctrl-C to start LSI Corp Configuration Utility...
```

Figure 18-2 SAS Expansion Card - LSI Logic Configuration Utility

3. The menu shown in Figure 18-3 is displayed after pressing Ctrl + C.

|                                                                            |         |         |         |          |               |         |            |
|----------------------------------------------------------------------------|---------|---------|---------|----------|---------------|---------|------------|
| LSI Corp Config Utility v6.26.02.00 (2009.06.08)                           |         |         |         |          |               |         |            |
| Adapter List Global Properties                                             |         |         |         |          |               |         |            |
| Adapter                                                                    | PCI Bus | PCI Dev | PCI Fnc | PCI Slot | FW Revision   | Status  | Boot Order |
| SAS3020XB                                                                  | 00      | 01      | 00      | 01       | 1.30.03.00-IR | Enabled | 0          |
| C1064E                                                                     | 02      | 00      | 00      | 00       | 0.10.15.00-IR | Enabled |            |
|                                                                            |         |         |         |          |               |         |            |
| Esc = Exit Menu F1/Shift+l = Help                                          |         |         |         |          |               |         |            |
| Alt+N = Global Properties -/+ = Alter Boot Order Ins/Del = Alter Boot List |         |         |         |          |               |         |            |

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 541.

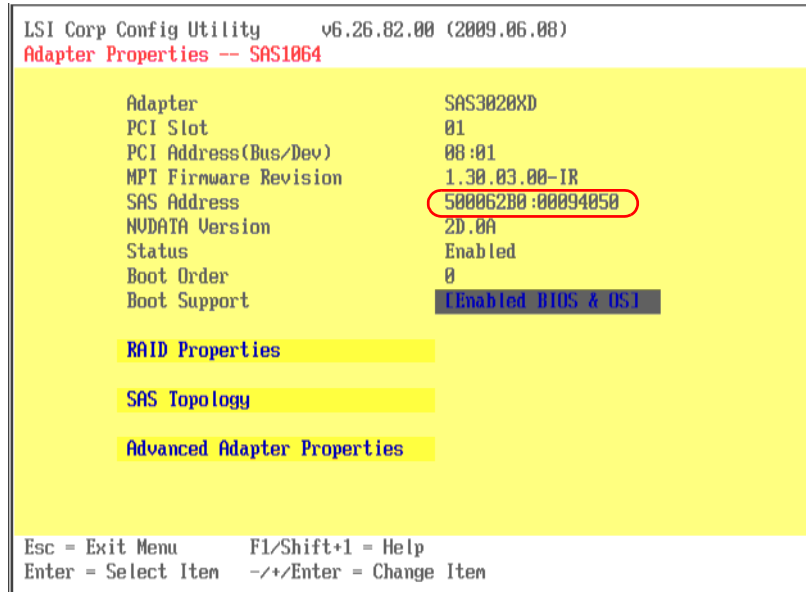


Figure 18-4 SAS Expansion Card - LSI Logic Config Utility adapter selected

- 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 DS3500 System Storage subsystem. The WWPN can also be retrieved from the SAS Connectivity Module web interface.

The WWPN can also be recorded by clicking the **Hardware VPD** → **Blade Bay** (in our example Blade4) → **Expansion Card**. This opens BladeCenter Vital Product Data window. Click the **Ports** tab to record the WWPNs as shown in Figure 18-5.

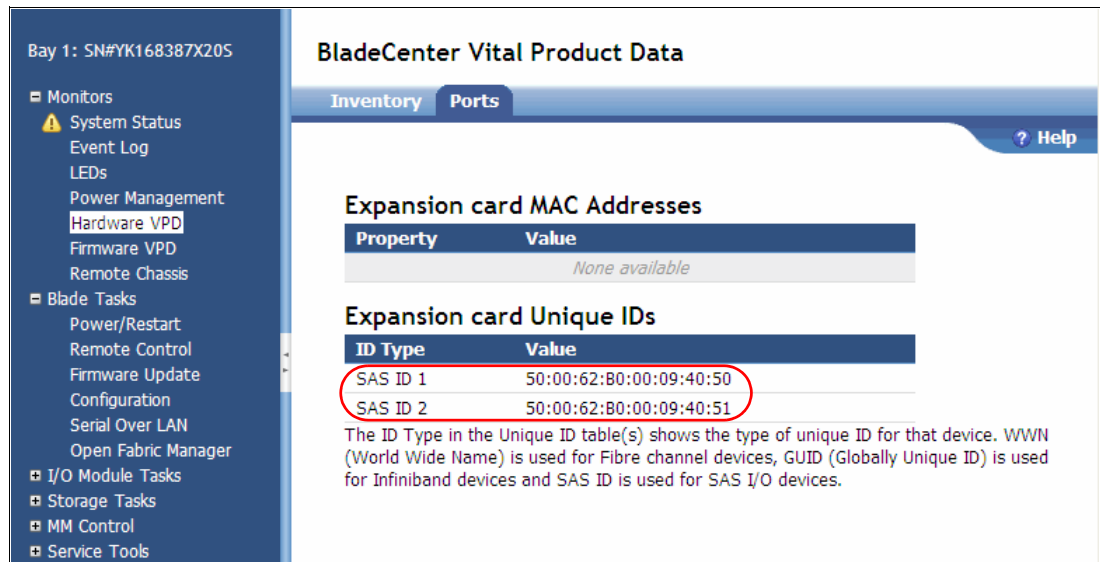


Figure 18-5 Expansion Card Unique IDs

The name of the SAS adapter for the expansion card is SAS3020XD and is visible in the Adapter List window. To determine if the SAS adapter is the expansion card, select a SAS adapter and use the View SAS Topology window 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 18-6 on page 542.

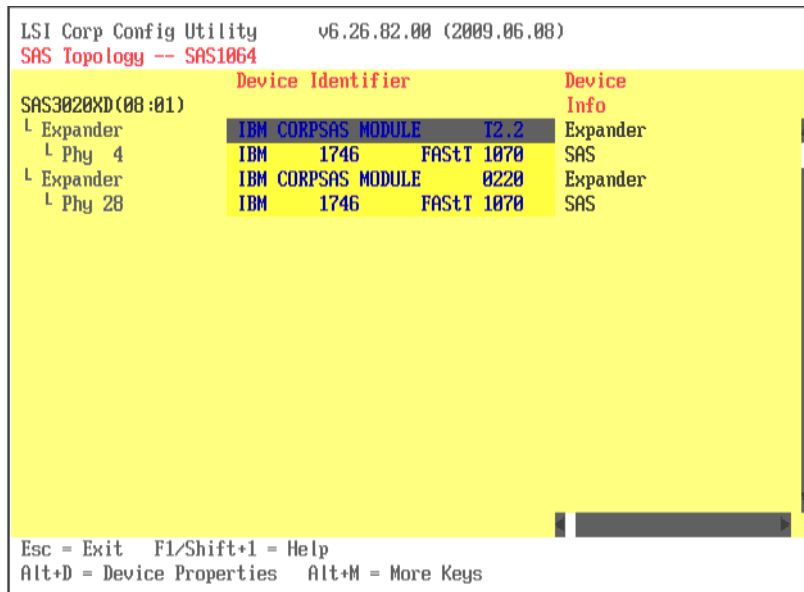


Figure 18-6 SAS Expansion Card - LSI Logic Config Utility adapter confirmation

## 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, the latest SAS Expansion Card device driver for Windows Server 2008 was available from following IBM website:

<http://www-947.ibm.com/support/entry/portal/docdisplay?lnodocid=MIGR-5077094&brandid=5000020>

## 18.2.5 SAS Connectivity modules

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 8740
- ▶ 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 the following URL:

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

Connect the cables from the DS3500 Storage System controllers A and B to external ports # 4 and 1 of the two SAS connectivity modules, as shown in Figure 18-7.

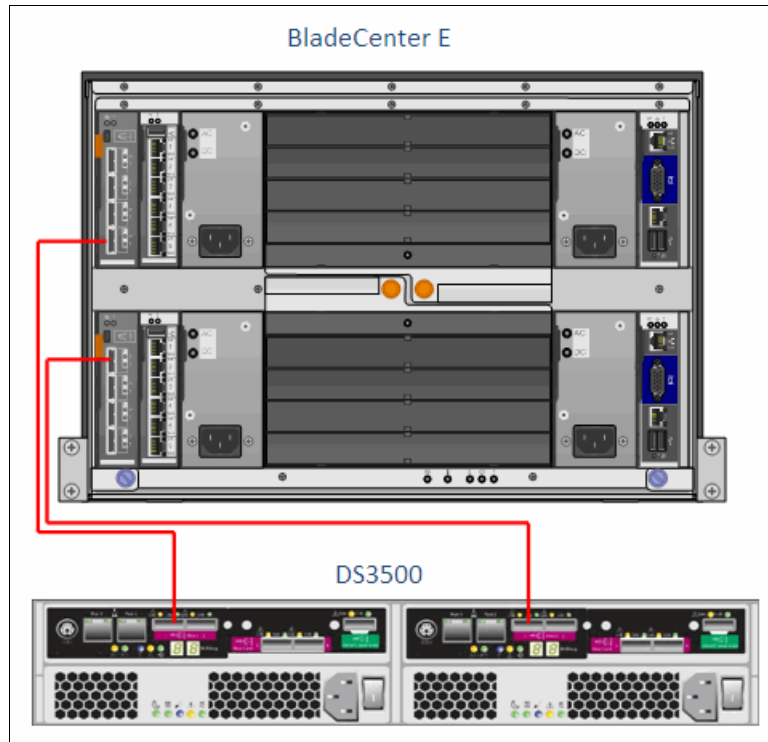


Figure 18-7 BladeCenter and DS3200 connectivity

## 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 18-8. Use USERID as the user Id and PASSWORD as the password. You can change the password under the Administer Users menu option after you are logged on.

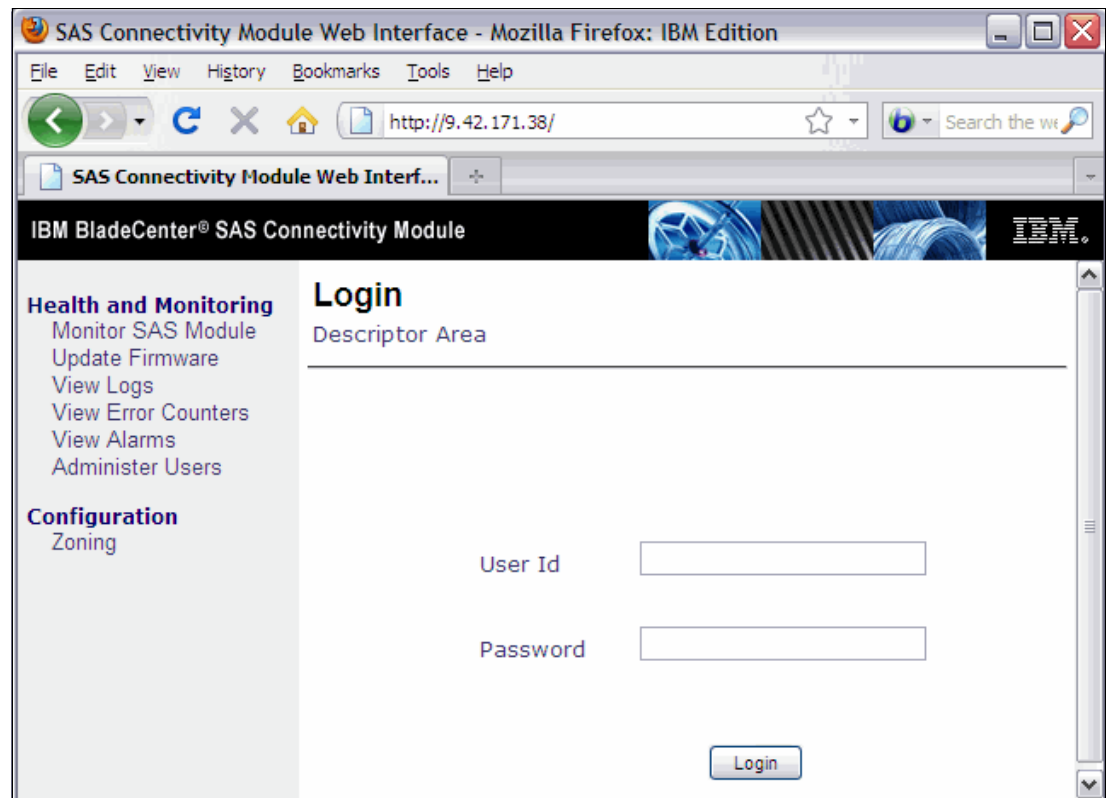


Figure 18-8 SAS Connectivity Module - Login

2. In the Monitor Module window, click **Update Firmware**. The Update Firmware window opens as shown in Figure 18-9 on page 545. The current firmware level is also displayed.

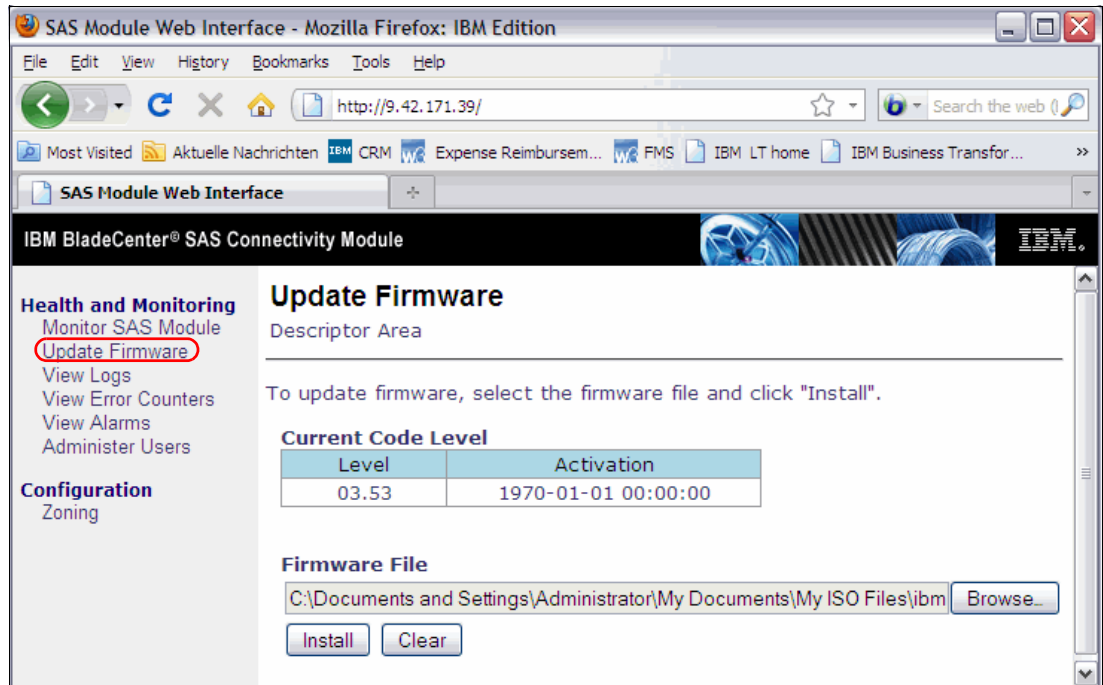


Figure 18-9 SAS Connectivity Module - Update Firmware

3. In the Firmware File field, type the new firmware file name, or click **Browse** and locate the firmware file.
4. Click **Install** to install the new file. A firmware update progress window opens as shown in Figure 18-10.




Figure 18-10 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 18-11 on page 546. If there are errors during the installation, an error message window displays.




Figure 18-11 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 DS3500 System Storage subsystem and I/O traffic to pass through the module.

From the BladeCenter Advance Management Module web interface GUI, click **I/O Module Tasks** → **Admin/Power/Restart** and ensure that the external ports for the I/O modules in bays 3 and 4 are enabled as shown in Figure 18-12.

| Bay                        | Type         | Manufacturer | MAC Address       | IP Address           | Pwr | Unique ID Type | ID                      |
|----------------------------|--------------|--------------|-------------------|----------------------|-----|----------------|-------------------------|
| <input type="checkbox"/> 1 | Ethernet SM  | DLINK (n/a)  | 00:05:5D:71:82:04 | <a href="#">View</a> | On  | n/a            | n/a                     |
| <input type="checkbox"/> 2 | Ethernet SM  | NT (BNT)     | 00:16:60:FD:5B:00 | <a href="#">View</a> | On  | n/a            | n/a                     |
| <input type="checkbox"/> 3 | SAS Conn Mod | IBM (n/a)    | 00:14:5E:C3:0D:2A | <a href="#">View</a> | On  | n/a            | 50:05:07:64:10:03:4c:40 |
| <input type="checkbox"/> 4 | SAS Conn Mod | IBM (n/a)    | 00:14:5E:C3:23:EA | <a href="#">View</a> | On  | n/a            | 50:05:07:64:10:08:fa:80 |
| 5                          | No Module    |              |                   |                      |     |                |                         |
| 6                          | No Module    |              |                   |                      |     |                |                         |
| 7                          | No Module    |              |                   |                      |     |                |                         |
| 8                          | No Module    |              |                   |                      |     |                |                         |
| 9                          | No Module    |              |                   |                      |     |                |                         |
| 10                         | No Module    |              |                   |                      |     |                |                         |

† If this notation is shown next to an IP address, it means the address is the external stack management address.

**Available actions**

Power On Module(s)

---

**I/O Module Advanced Setup** ?

Select a module: I/O module 4

Fast POST: Enabled

External ports: Enabled

Figure 18-12 SAS Connectivity Module - Configuration

## 18.2.8 SAS Connectivity Module zoning

Zoning is the partitioning of a Fibre Channel fabric into smaller subsets to restrict interference, add security, and to simplify management. Although a SAN makes available several virtual disks (LUNs), each system connected to the SAN should only be allowed access to a controlled subset of the LUNs. 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 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.

Click **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 18-13 on page 548.

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 18-13 on page 548. 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.

### I/O Module 3 ( SAS Conn Mod ) ?

The table below lists zone configurations stored on this I/O Module.

| Select                | Active? | Name                   | Type         | Description                                                                                                                                                                                                                            | Configuration Store |
|-----------------------|---------|------------------------|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|
| <input type="radio"/> |         | 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                   |
| <input type="radio"/> |         | 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                   |
| <input type="radio"/> |         | 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                   |
| <input type="radio"/> |         | 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                   |
| <input type="radio"/> | ✓       | 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                   |

Figure 18-13 SAS Connectivity Module - Predefined Zones

The following example in Figure 18-14 on page 549 was captured from the SAS Connectivity Module web interface. It lists the Basic Zone Permission table for the HS21 blade in slot 4, because the blade in slot 4 is used for this example. The blade in slot 4 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 DS3500 Storage System Controller A port is connected to the external port #4 for this example.

IBM BladeCenter® SAS Connectivity Module

Health and Monitoring

Monitor SAS Module  
Update Firmware  
View Logs  
View Error Counters  
View Alarms  
Administer Users

Configuration

Zoning

Manage Fabric

Descriptor Area

Working Configuration

Predefined Config 01

Activate this Configuration

Active Configuration

Predefined Config 01

Zone Groups

Basic Zone Permission Table

Select Zone Group

View by: Zone Group ID 10

| Port            | Attached Port Add | Enabled | Connected | Status   |
|-----------------|-------------------|---------|-----------|----------|
| External Port 1 |                   | true    | false     | No Cable |

Mapped to

Remove from Permission Table

| Select                   | Zone Group ID | Port                    | Attached Port Add | Enabled | Connected | Status   |
|--------------------------|---------------|-------------------------|-------------------|---------|-----------|----------|
| <input type="checkbox"/> | 10            | External Port 1         |                   | true    | false     | No Cable |
| <input type="checkbox"/> | 11            | External Port 2         |                   | true    | false     | No Cable |
| <input type="checkbox"/> | 12            | External Port 3         |                   | true    | false     | No Cable |
| <input type="checkbox"/> | 13            | External Port 4         |                   | true    | true      | Normal   |
| <input type="checkbox"/> | 15            | Blade Slot Connection 1 | 0000000000000000  | true    | false     | No Cable |
| <input type="checkbox"/> | 16            | Blade Slot Connection 2 | 0000000000000000  | true    | false     | No Cable |
| <input type="checkbox"/> | 17            | Blade Slot Connection 3 | 0000000000000000  | true    | false     | No Cable |
| <input type="checkbox"/> | 18            | Blade Slot Connection 4 | 500062B000094050  | true    | true      | Normal   |

Figure 18-14 SAS Connectivity Module - Zone configuration

**Note:** There is only a single path active from the DS3500 Storage System controller A to the SAS Connectivity Module in BladeCenter I/O slot bay 3. The second path from DS3500 Storage System controller B is to the SAS Connectivity Module in BladeCenter I/O slot bay 4.

## 18.3 Installing DS Storage Manager host software

We will install the host server components of DS3500 Storage Manager. These components include:

- ▶ DS Storage Manager 10 Utilities
- ▶ DS Storage Manager 10 Agent
- ▶ MPIO Device Specific Module (DSM)

We will install DS Storage Manager 10 Utilities because this will allow us to run the *hot\_add* utility so that we do not need to reboot the host server when adding new logical drives.

Because 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 DS Storage Manager.

For the exact installation steps, see Chapter 5, “Installing IBM System Storage DS Storage Manager” on page 131.

After the DS3500 Storage Manager client is installed, we add the DS3500 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 20 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 Chapter 11, “Administration - Mappings Tab” on page 275, so we do not repeat these steps here. Figure 18-15 shows our 20 GB logical drive named Data and mapped to the host server Blade04.

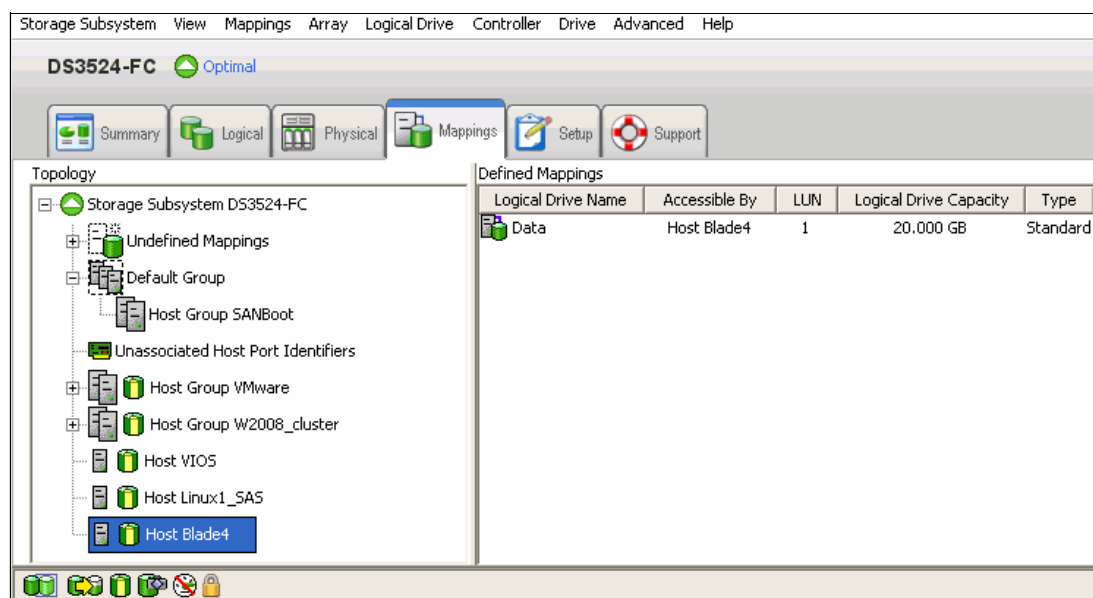


Figure 18-15 Drive mapping

## 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 or to select **Rescan Disks** from **Action** menu bar of **Disk Management** window. But the a better alternative is to use the **hot\_add** utility (a part of SMutil).

By default, the **hot\_add** executable is located in C:\Program Files (x86)\IBM\_DS\util. In a command prompt window, change to this directory. If you installed the DS3500 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 18-16.

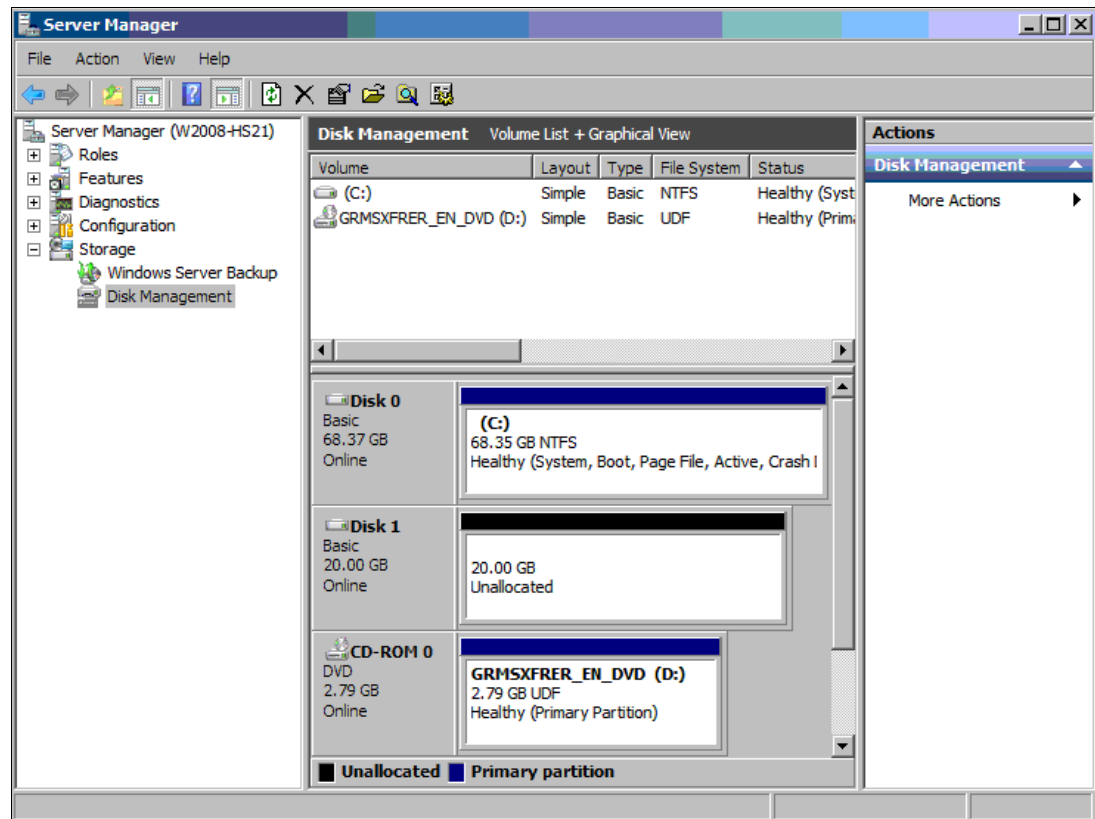


Figure 18-16 Disk Management applet - initialized disk

As you can see, our 20 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 as shown in Figure 18-17 on page 552.

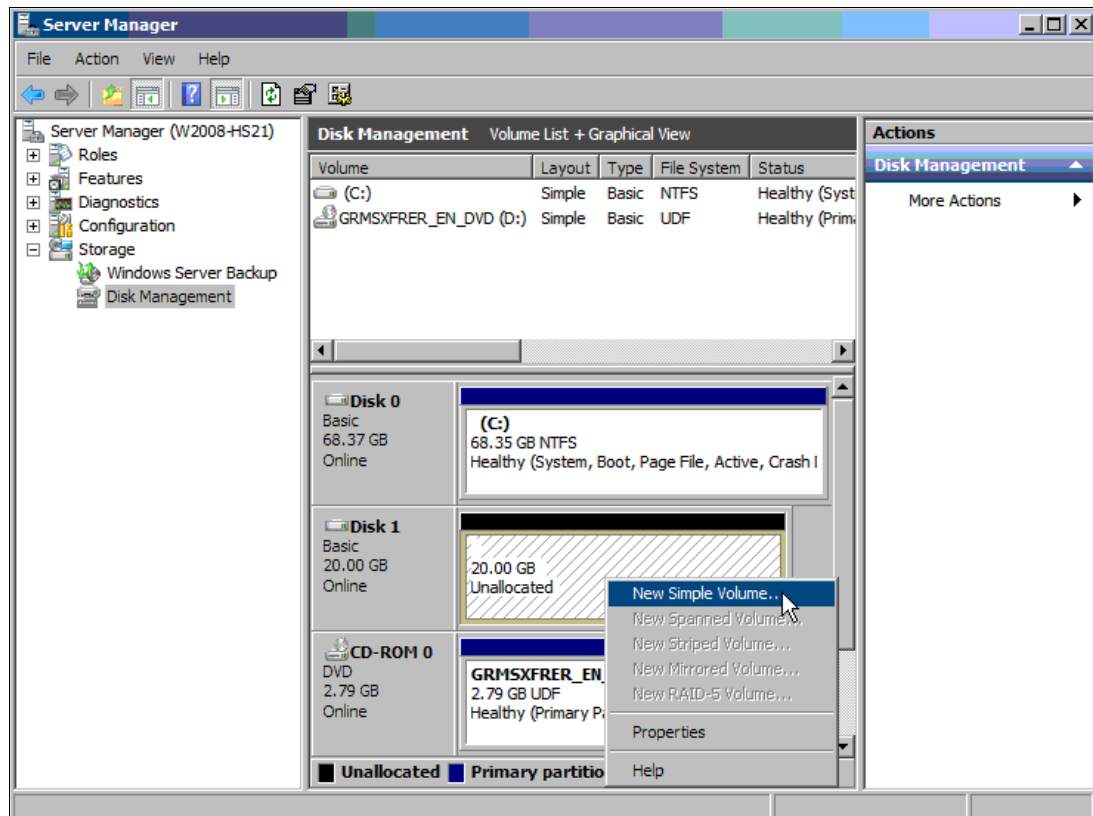


Figure 18-17 Creating a new partition

This launches the New Simple Volume Wizard.

- 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 E) that is ready to be used (see Figure 18-18 on page 553).




Figure 18-18 New partition on the logical drive

4. We can now access the logical drive on the DS3500 as drive letter E: as shown in Figure 18-19.




Figure 18-19 Logical drive is ready for use





## Microsoft Cluster configuration with DS3500

In this chapter, we describe a sample configuration to attach a DS3500 server to Microsoft Windows Server 2008 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.

For more information about Microsoft Cluster creation, refer to the appropriate documents at the following URL:

<http://www.microsoft.com/downloads/details.aspx?FamilyID=75566F16-627D-4DD3-97CB-83909D3C722B&displaylang=en&displaylang=en>

## 19.1 Overview of a failover cluster

A failover cluster is a group of independent servers, or nodes, that are physically connected by a local area network (LAN) or a wide area network (WAN) and that are connected by cluster software. The group of nodes is managed as a single system and shares a common namespace. The group usually includes multiple network connections and data storage connected to the nodes through storage area networks (SANs). The failover cluster operates by moving resources between nodes to provide service if system components fail.

Normally, if a server that is running a particular application crashes, the application will be unavailable until the server is fixed. Failover clustering addresses this situation by detecting hardware or software faults and immediately restarting the application on another node without requiring administrative intervention, a process known as failover. Users can continue to access the service and might be completely unaware that it is now being provided from a separate server.

Servers in a failover cluster can function in a variety of roles, including the roles of file server, print server, mail server, and database server, and they can provide high availability for a variety of other services and applications. This guide describes how to configure a two-node file server cluster.

A failover cluster usually includes a storage unit that is physically connected to all the servers in the cluster, although any given volume in the storage is only accessed by one server at a time. The Figure 19-1 shows a two-node failover cluster connected to a storage unit.




Figure 19-1 Failover clustering

### 19.1.1 Hardware requirements for a two-node failover cluster

You will need the following hardware for a two-node failover cluster:

1. **Servers with a minimum of two network adapters:** Use a set of matching computers that contain the same or similar components. In our example, we used two IBM Blade Server HS22 connected through two Fibre Channel switches to a DS3500 storage subsystem.
2. **Device controllers or appropriate adapters for the storage:** Using Serial Attached SCSI, iSCSI or Fibre Channel, in all clustered servers, all components of the storage stack

should be identical. The multipath I/O (MPIO) software and Device Specific Module (DSM) software components must be identical. In addition, the mass-storage device controllers, that is the host bus adapter (HBA), HBA drivers, and HBA firmware that are attached to cluster storage, should also be identical.

3. **Storage:** You must use shared storage that is compatible with Windows Server 2008. For a two-node failover cluster, the storage should contain at least two separate volumes (LUNs) configured at the hardware level. One volume will function as the witness disk (also known as Quorum drive). One volume will contain the files that are being shared to users.

## 19.2 Preparing the environment

First, we perform initial setup that includes:

- ▶ Infrastructure cabling and zoning
- ▶ DS3500 basic configuration

### 19.2.1 SAN Zoning configuration

We set up four paths between the server and the storage system. At least two HBA ports should be used, and both HBA ports must see both controllers.

Refer to the following materials for more information about zoning:

- ▶ *Implementing an IBM b-type SAN with 8 Gbps Directors and Switches*, SG24-6116
- ▶ *Implementing an IBM/Cisco SAN*, SG24-7545

### 19.2.2 DS3500 Storage configuration

Use the DS Storage Manager client to define the logical drive, host group, host, host type, and host ports. Make sure that the host type defined matches the operating system you are installing (in our example, Windows 2000/Server 2003/Server 2008 Clustered).

Follow these steps for the configuration:

1. Zone your Windows Server to your IBM DS3500 Storage Subsystem. Ensure that your environment has sufficient paths and connections for redundancy and high availability.
2. We used the Logical tab in DS Subsystem Management window to create a sample logical drives as Quorum(witness) and shared data disks
3. Define a host group by selecting the **Mapping** menu in Mapping tab as shown in Figure 19-2 on page 558. Specify a name for the host group and follow the wizard's instructions to complete the task. In our example, we named our host group MSCluster.

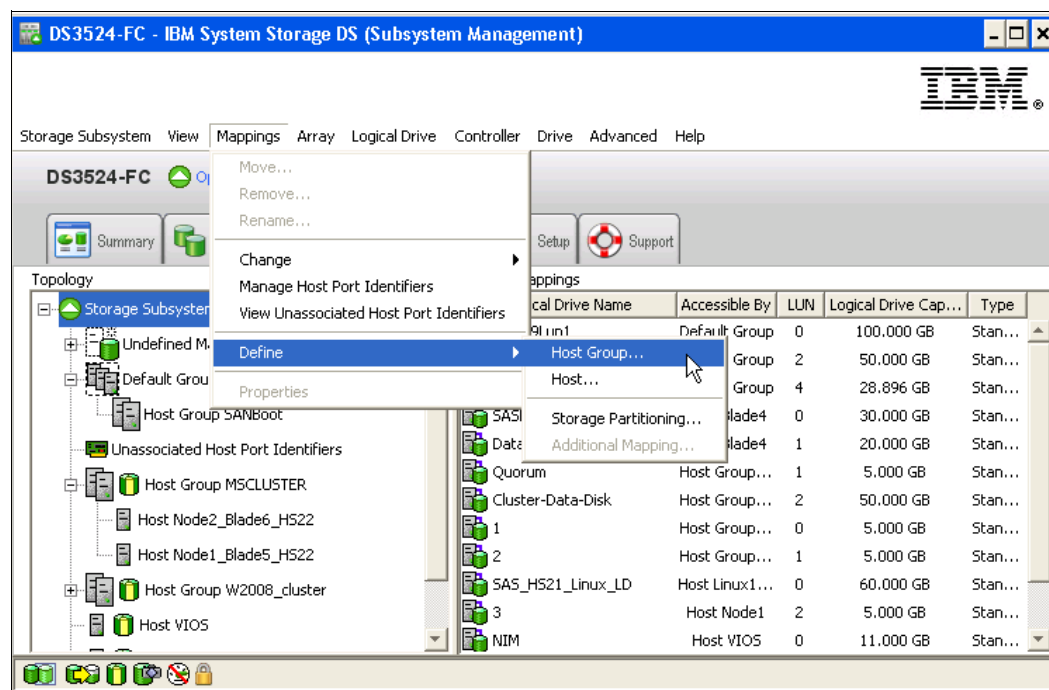


Figure 19-2 Define Host Group

- Click the Mapping tab, and then click **Mapping** → **Define** → **Host** as shown in Figure 19-3 on page 559. Specify the host name and the host interface type. Add the host port identifiers and select the appropriate host type. In our example, we defined Node1-Blade5 and Node2-Blade6 as our two cluster nodes.



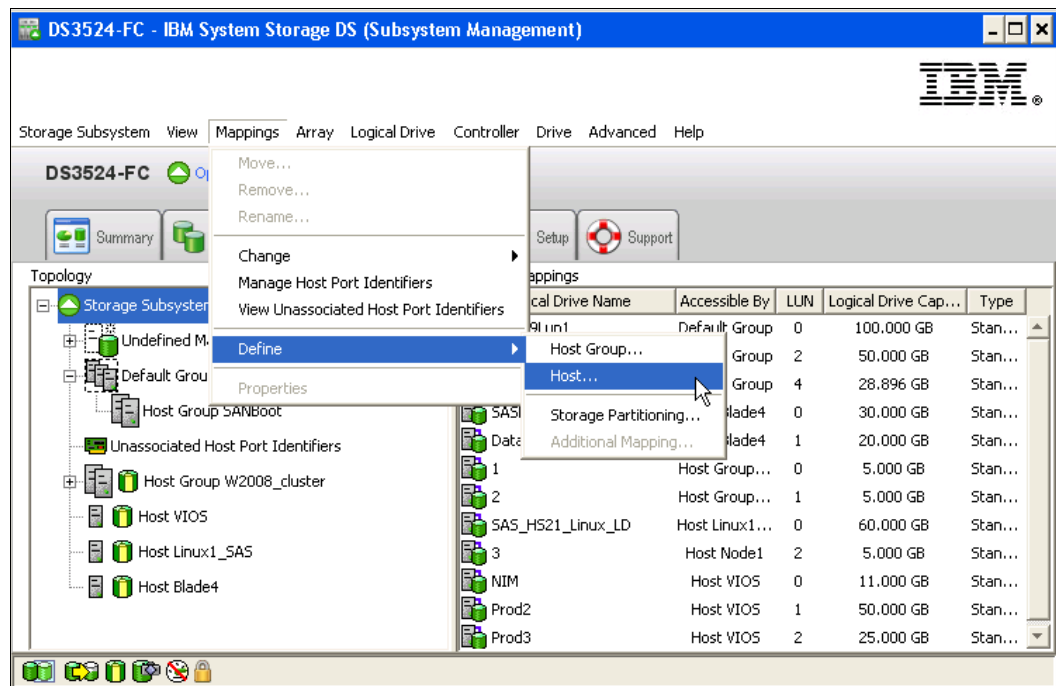


Figure 19-3 Define Hosts

**Note:** In our example, we selected Windows 2000/Server 2003/Server 2008 Clustered as the host operating system.

5. Add the defined hosts to the host group to share the storage resources between the hosts and for failover as shown in Figure 19-4 on page 560




Figure 19-4 Add Hosts to the Host Group

6. Create the host group-to-logical drive mapping. Select the host group created in previous step (in our example, MSCluster) and define storage partitioning by clicking the **Mapping** → **Define** → **Storage Partitioning** as shown in Figure 19-5 on page 561.




Figure 19-5 Create Host Group-to-Logical Drive Mapping

7. Select the logical drives and assign the logical unit numbers (LUNs) that the host group will use to access the logical drives in the partition as shown in Figure 19-6

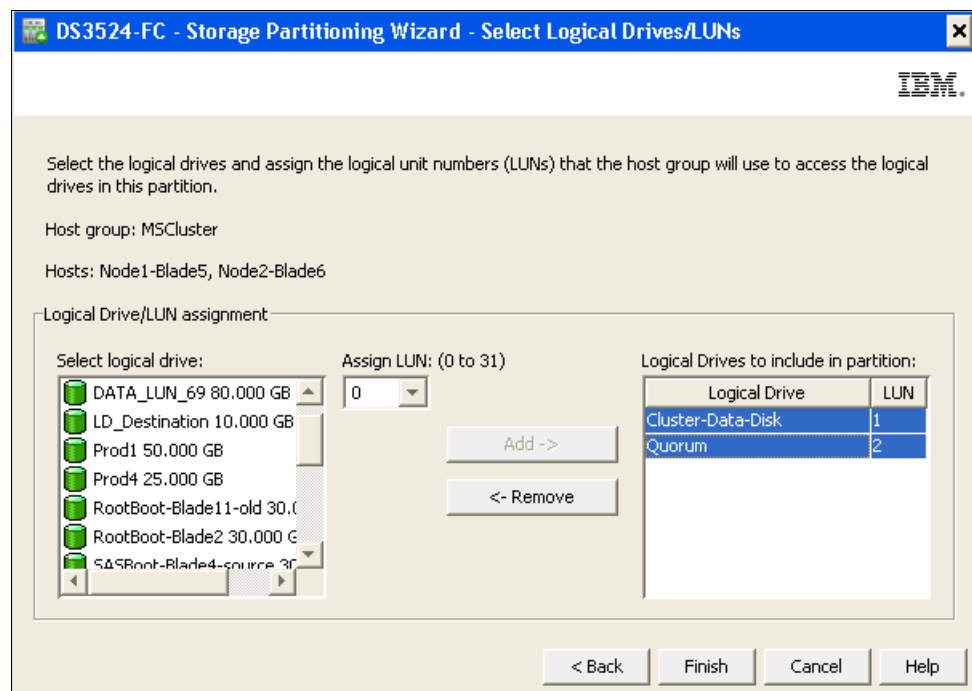


Figure 19-6 Storage Partitioning Wizard

Figure 19-7 on page 562 shows an example of a valid LUN mapping for installation purposes. Check to ensure that the logical drives are mapped to the host group and the hosts are the members of the host group.

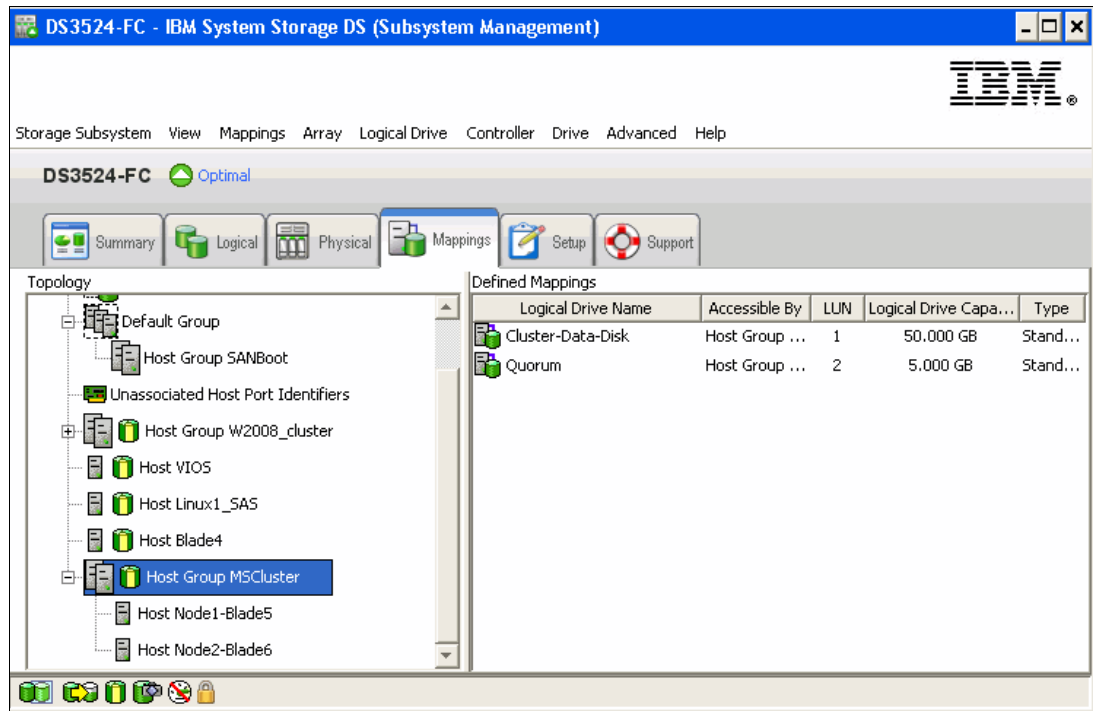


Figure 19-7 Host Group-to-Logical Drive Mapping

## 19.3 Installing DS Storage Manager host software

We will install the host server components of DS Storage Manager. These components include:

- ▶ DS Storage Manager 10 Utilities
- ▶ DS Storage Manager 10 Agent
- ▶ MPIO Device Specific Module (DSM)

We will install DS Storage Manager 10 Utilities, 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.

For the exact installation steps, see Chapter 5, “Installing IBM System Storage DS Storage Manager” on page 131.

### 19.3.1 Installing the multipath driver

The next step is to install the multipath driver. Perform the following steps:

1. Start Storage Manager on your host server. A detailed description for installing the Storage Manager is given in Chapter 5, “Installing IBM System Storage DS Storage Manager” on page 131. Follow the steps, select either the **Host** installation type or **Custom** installation (Figure 19-8 on page 563), and click **Next**.




Figure 19-8 Multipath driver installation

2. Follow the wizard instructions to complete the installation. Reboot your server
3. After the reboot, open Disk Management to create the partitions and ensure that the partitions are accessible from both cluster nodes.

## 19.4 Windows Server 2008 Failover Clustering

This guide describes the steps for installing and configuring a failover cluster that has two nodes. By creating the configuration in this guide, you can learn about failover clusters and familiarize yourself with the Failover Cluster Management snap-in.

### 19.4.1 Installing the Failover Clustering Feature

To set up the cluster, we have to add the server clustering feature. The failover clustering is a feature that can be added using the Initial Configuration Tasks Wizard or using the Server Manager console. Follow the steps to install the failover clustering.

1. In Server Manager, expand **Features** and select **Add Features** as shown in Figure 19-9 on page 564.




Figure 19-9 Add Features

2. In the Select Features window, select the feature **Failover Clustering** as shown in Figure 19-10 and click **Next**.

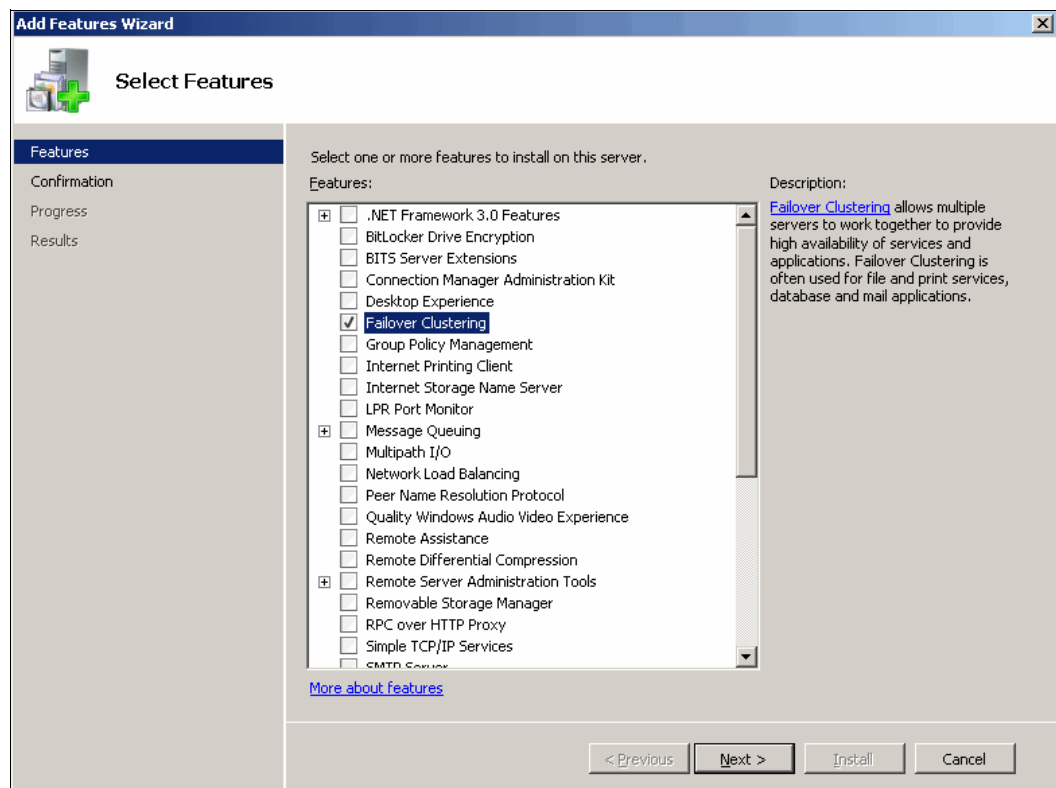


Figure 19-10 Enabling the Failover Clustering

**Note:** The Failover Clustering allows multiple servers to work together to provide high availability of services and applications.

3. Confirm installation selection and click **Install** to enable the failover clustering feature as shown in Figure 19-11.

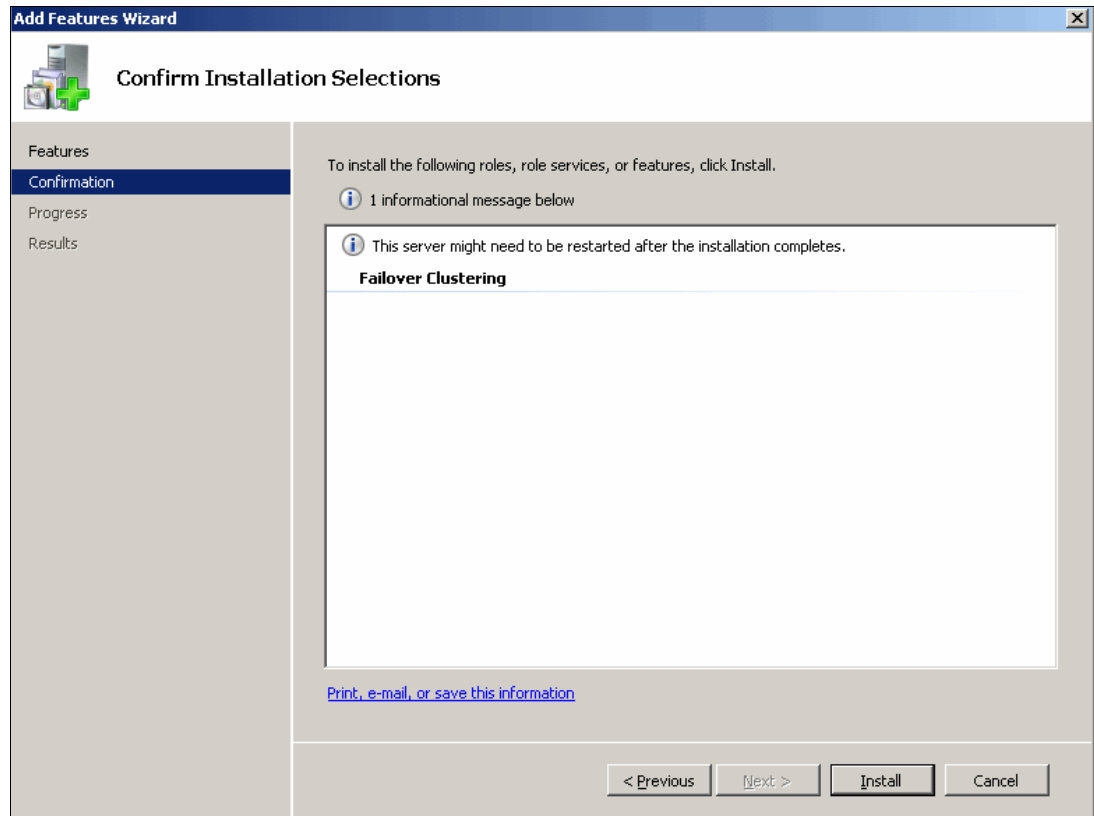


Figure 19-11 Confirm Installation Selection

4. Ensure that the failover clustering feature is installed successfully as shown in Figure 19-12 on page 566.




Figure 19-12 Installation Succeeded

## 19.4.2 Validate a Configuration

Before creating a cluster, validate your configuration. Validation helps you confirm that the configuration of your servers, network, and storage meets a set of specific requirements for failover clusters.

To validate the failover cluster configuration

1. To open the failover cluster snap-in, click **Start**, click **Administrative Tools**, and then click Failover Cluster Management as shown in Figure 19-13 on page 567.



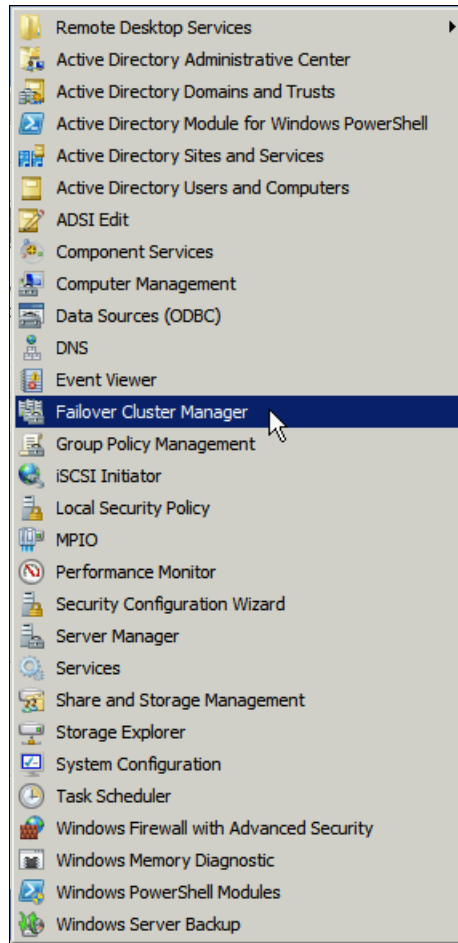


Figure 19-13 The Failover Cluster Manager

2. Confirm that **Failover Cluster Management** is selected and then, in the center pane under **Management**, click **Validate a Configuration** as shown in Figure 19-14 on page 568 and click **Next** to continue.

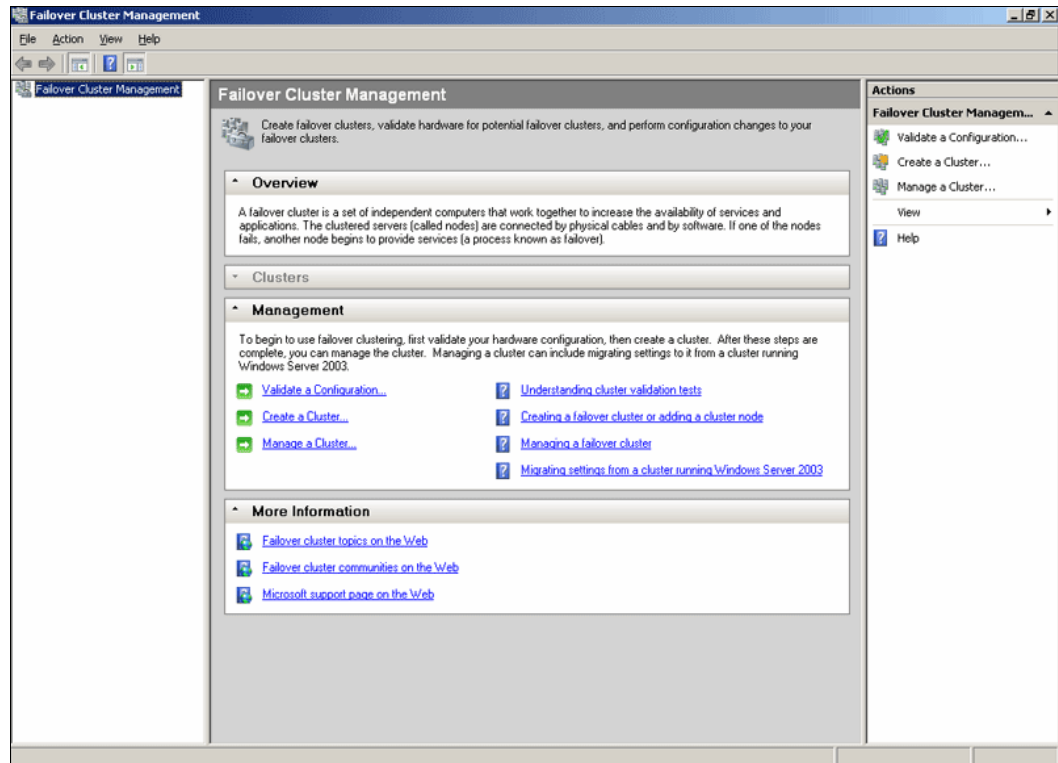


Figure 19-14 Validate a Configuration

3. Follow the instructions in the wizard to specify the two servers and the tests, and then run the tests. To fully validate your configuration, run all tests before creating a cluster. Select the Testing Options (in our example, we selected Run all tests) to run the tests as shown in Figure 19-15 on page 569 and click **Next**.

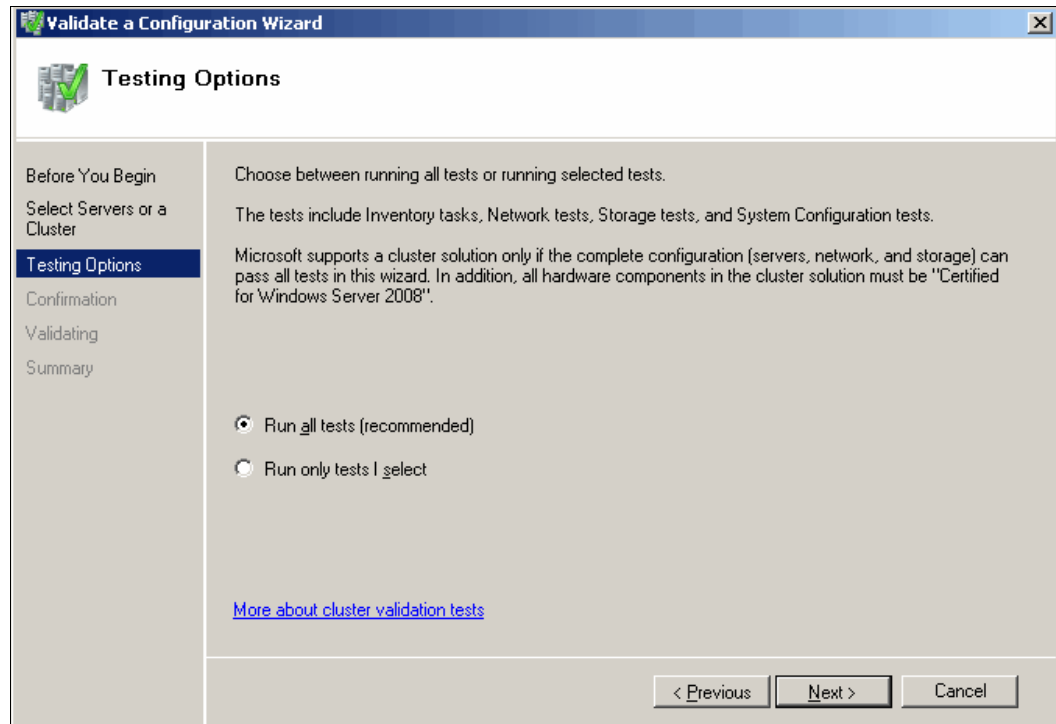


Figure 19-15 Testing Options

**Note:** Microsoft supports a cluster solution only if the complete configuration can pass all tests in the Validate a Configuration Wizard.

4. The Failover Cluster Validation Report shows the successful validation of our configuration as shown in Figure 19-16 on page 570. Click **Finish** to complete the task.

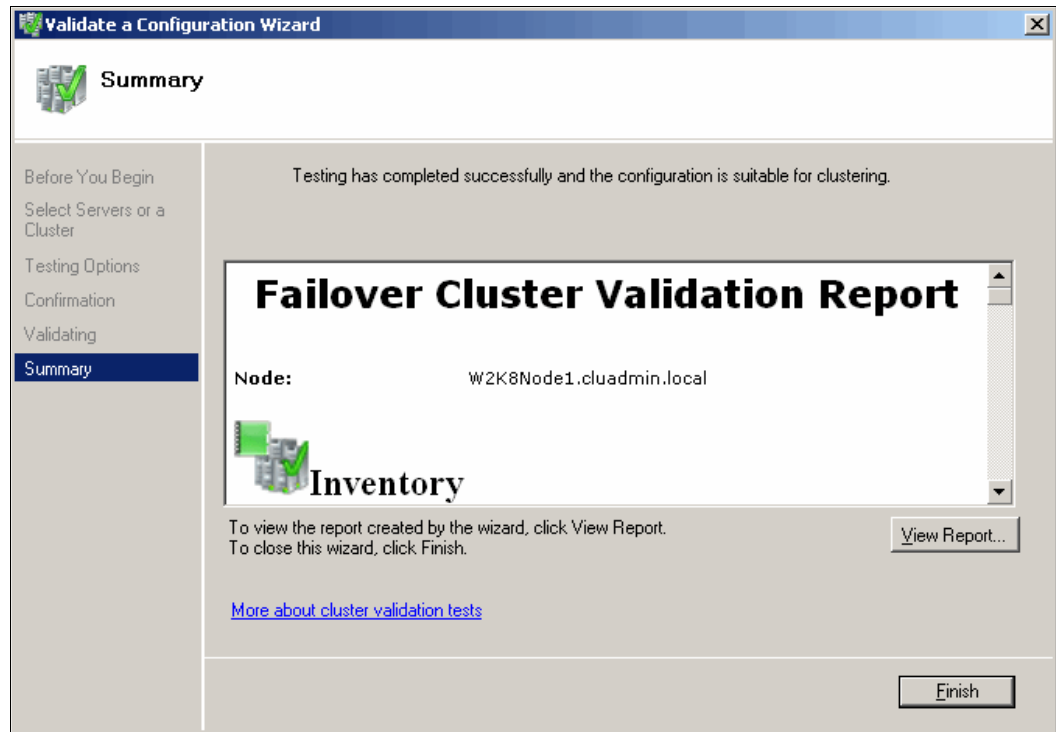


Figure 19-16 Failover Cluster Validation Report

### 19.4.3 Create a cluster

After a successful validation of the cluster configuration, start the cluster installation by performing the following steps:

1. Confirm that **Failover Cluster Management** is selected and then, in the center pane under **Management**, click **Create a cluster** as shown in Figure 19-17 on page 571 and click **Next**.

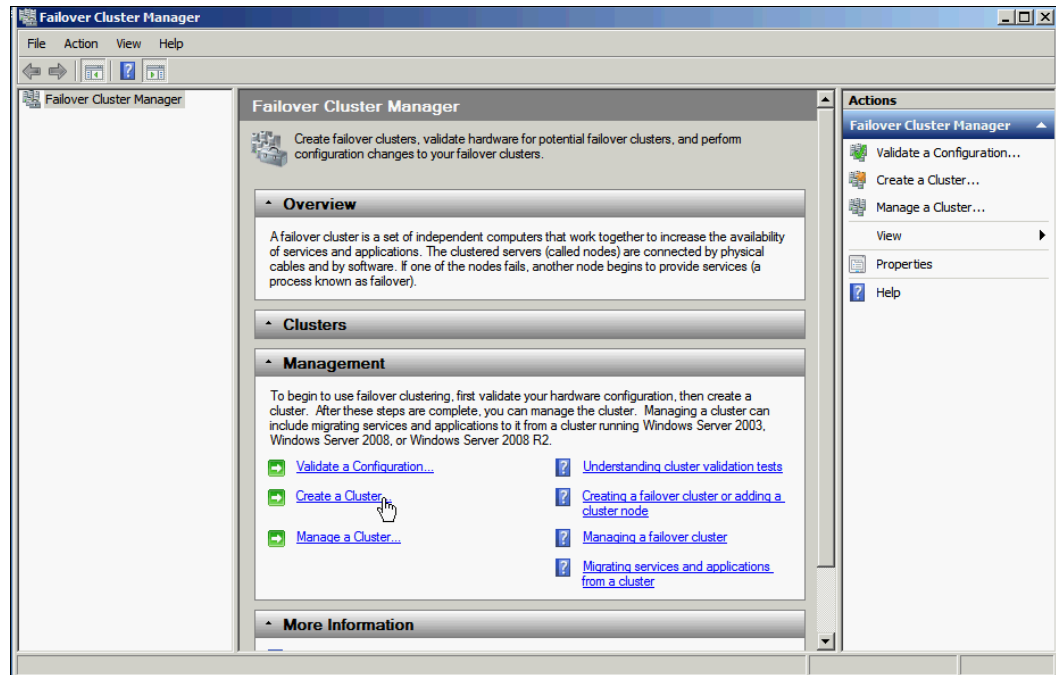


Figure 19-17 Create a Cluster

2. Follow the instructions in the wizard to specify:
  - The servers to include in the cluster
  - The name of the cluster
  - Any IP address information that is not automatically supplied by your DHCP settings
3. After the wizard runs and the Summary page appears, ensure that the cluster is installed successfully as shown in Figure 19-18.

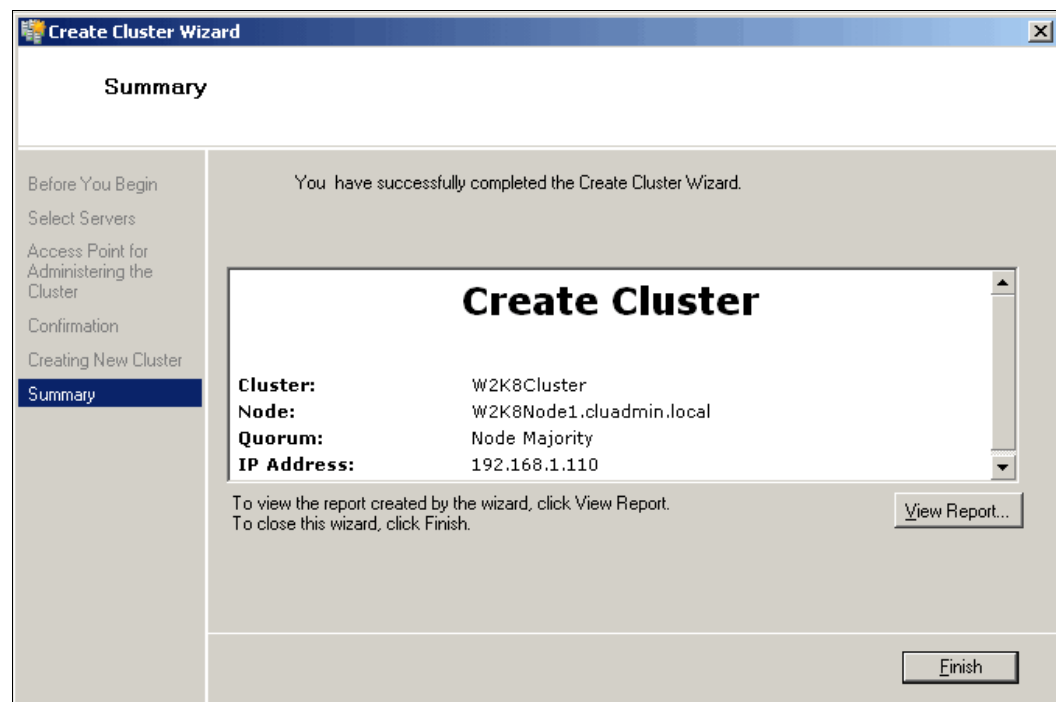


Figure 19-18 Cluster Installation Completed

## 19.4.4 Quorum configuration

In the following we will configure the quorum options, add the second node, and configure the first cluster resources.

Follow these steps to complete these tasks.

1. Add a Disk by selecting the **Storage** in the Failover Cluster Management window as shown in Figure 19-19.

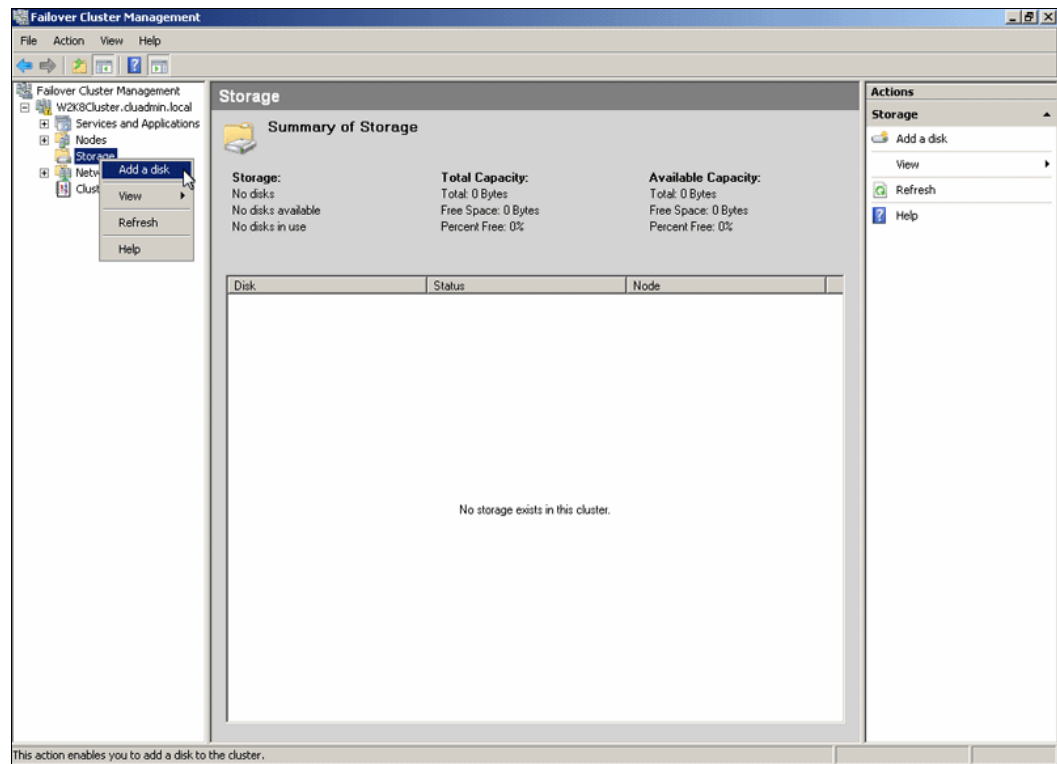


Figure 19-19 Add a Disk

2. Select the **Available Disk** resources to add to the cluster as shown in Figure 19-20 on page 573 and click **OK** to launch the summary window as shown in Figure 19-21 on page 573.

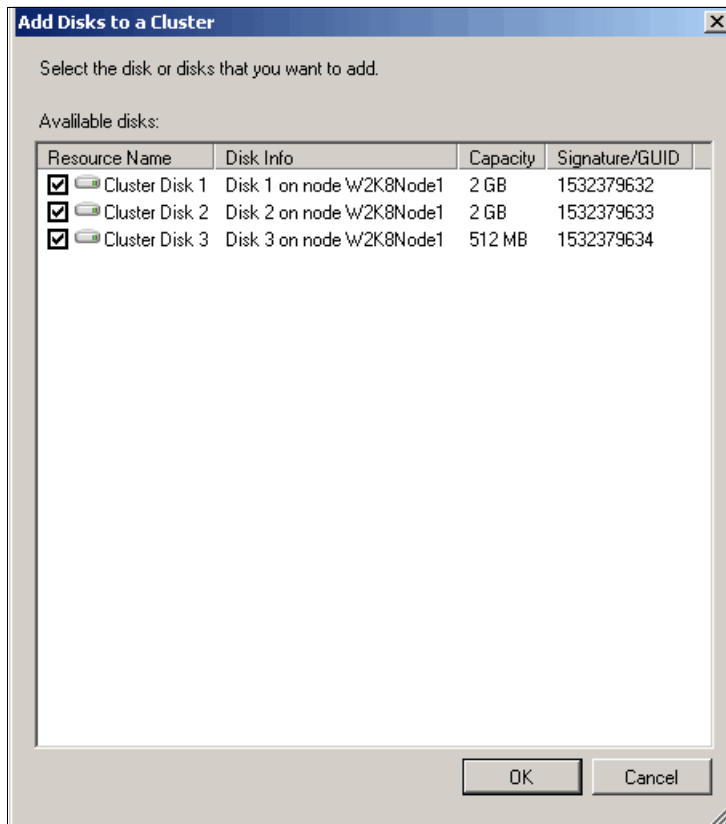


Figure 19-20 Add Disk Resources

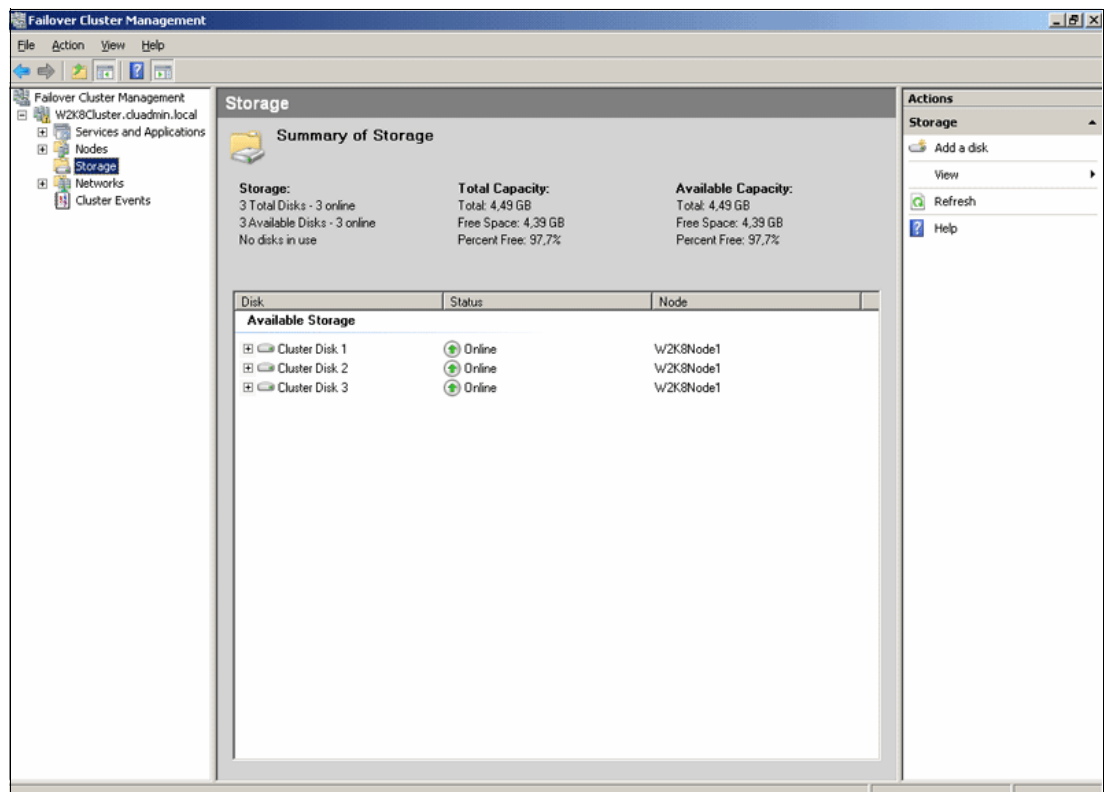


Figure 19-21 Summary of Storage

3. Select **Configure Cluster Quorum Settings** by right-clicking the cluster (in our example **W2K8Cluster.cluadmin.local**) and selecting **More Actions** in the Failover Cluster Management window as shown in Figure 19-22 and click **Next** to continue.

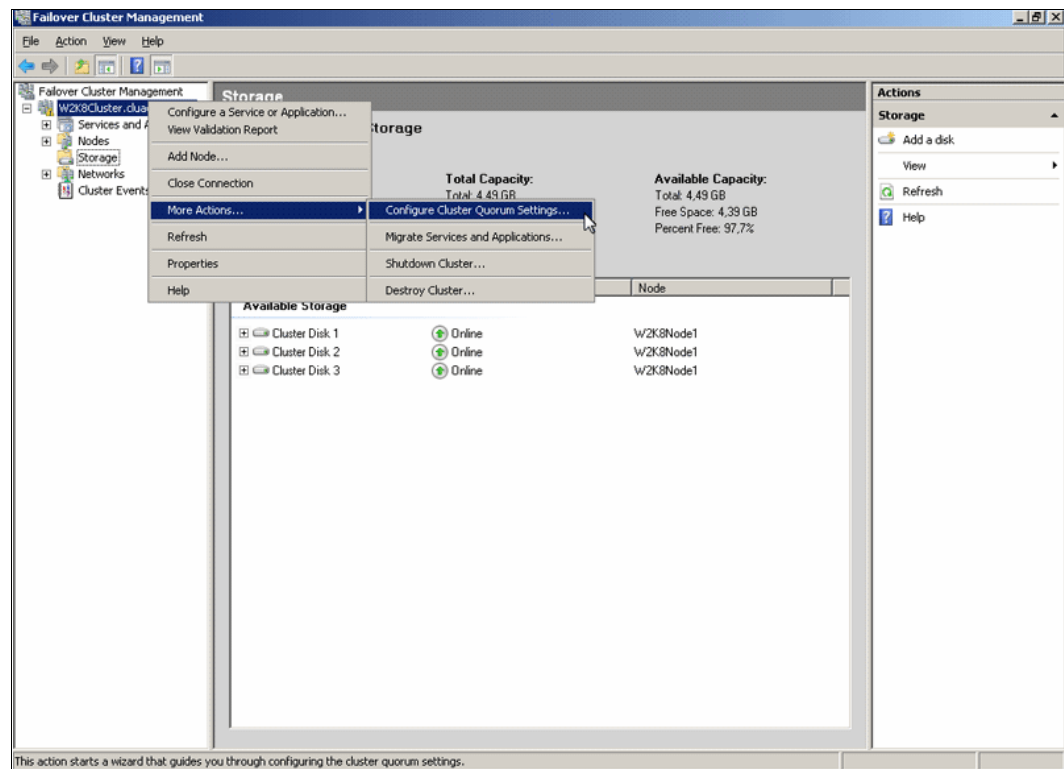


Figure 19-22 Configure Cluster Quorum Settings

4. Select the **Quorum Configuration** as shown in Figure 19-23 on page 575 and click **Next** to launch Configure Cluster Quorum Wizard as shown in Figure 19-24 on page 576.



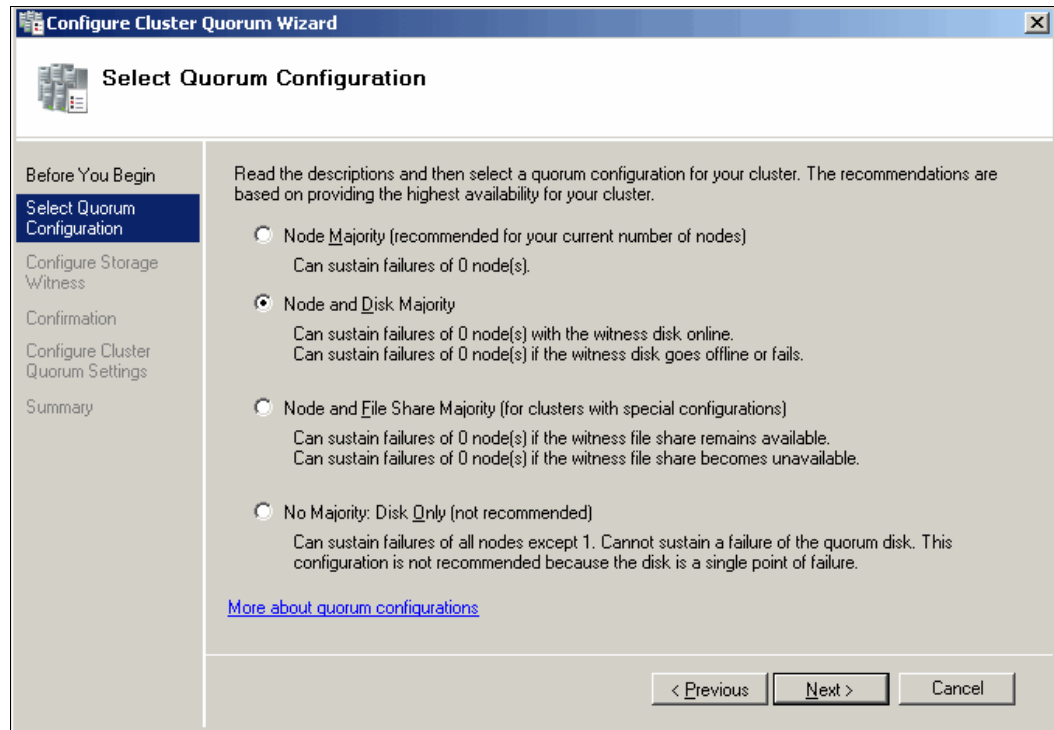


Figure 19-23 Select Quorum Configuration

**Note:** For this two-node cluster, the quorum configuration will be Node and Disk Majority, the default for a cluster with an even number of nodes. Node and Disk Majority means that the nodes and the witness disk each contain copies of the cluster configuration, and the cluster has quorum as long as a majority (two out of three) of these copies are available.

5. Select the storage volume (in our example, **Cluster Disk 3**) that you want to assign as the disk witness (Quorum). Click **Next**.

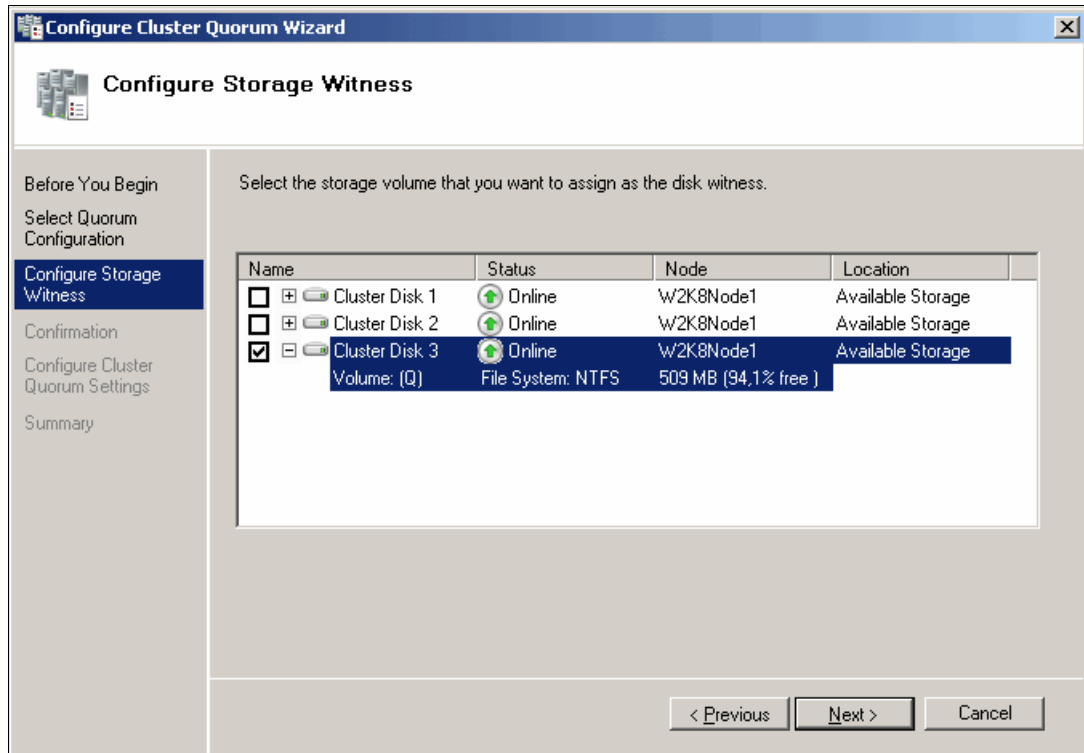


Figure 19-24 Configure Storage Witness

**Note:** The witness disk is a disk in the cluster storage that is designated to hold a copy of the cluster configuration database.

6. Confirm the Quorum Configuration as shown in Figure 19-25 on page 577 and click **Next**.

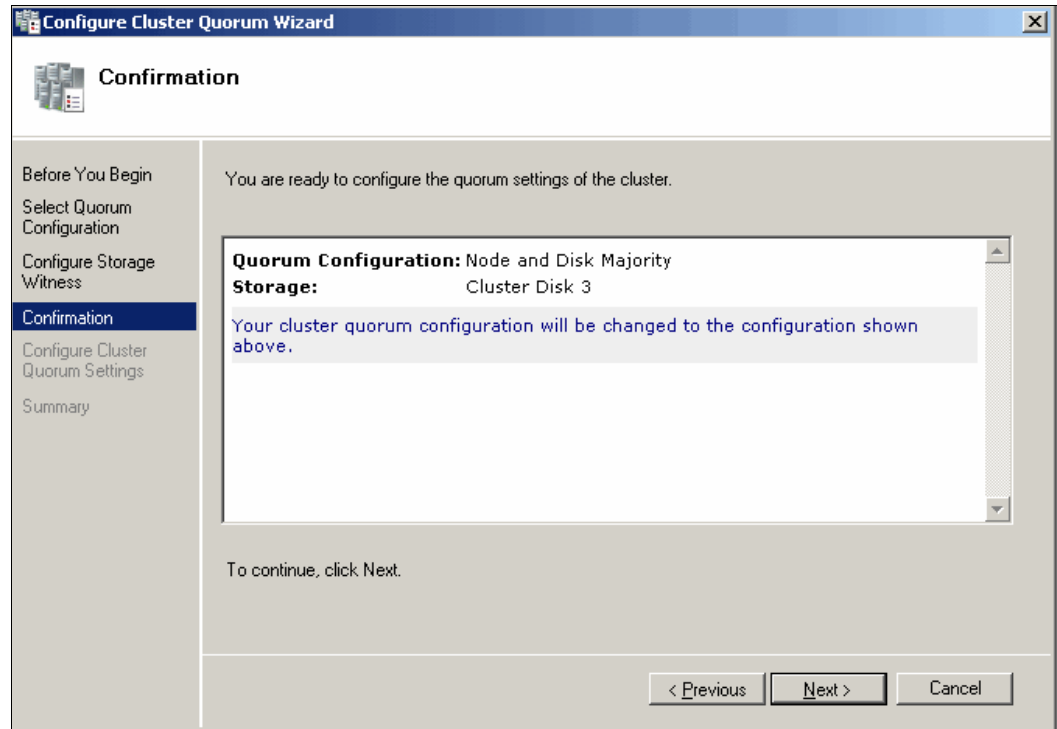


Figure 19-25 Confirmation of Quorum Configuration

7. Verify the Cluster Quorum Settings as shown in Figure 19-26 and click **Finish** to complete the task.

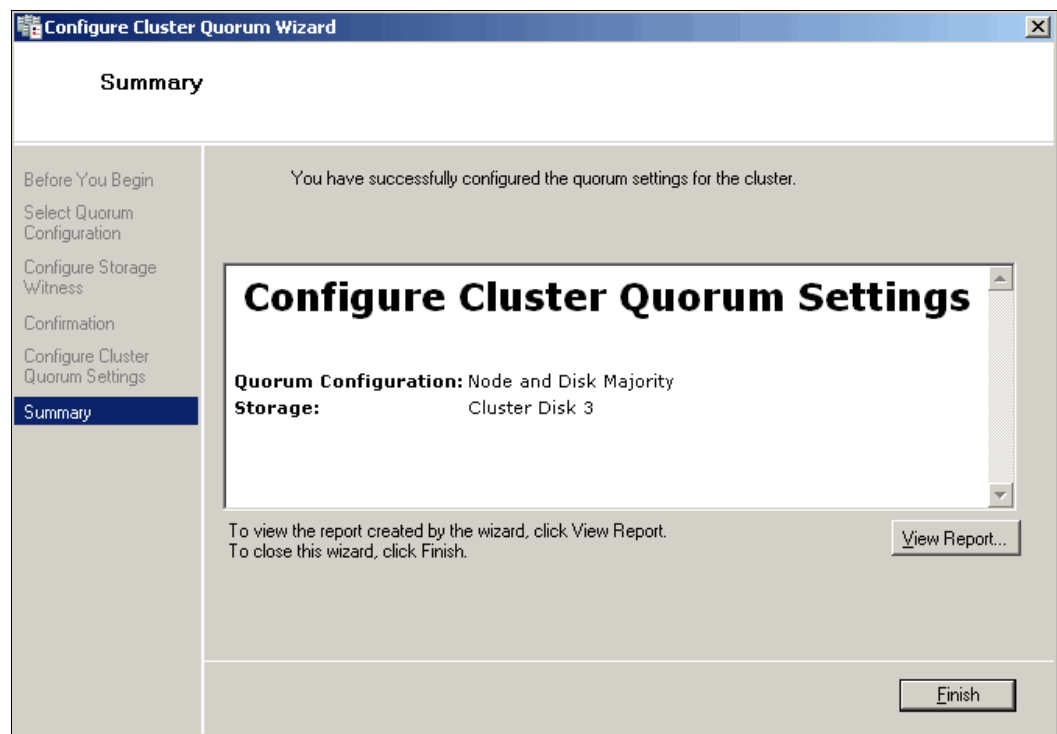


Figure 19-26 Summary of the Cluster Quorum Settings

8. Now add the second node as shown in Figure 19-27. Only then we will have formed true redundancy and a valid cluster.




Figure 19-27 Add the Second Node

The wizard for adding a node is similar to the dialogues of the installation process of a cluster and as soon as the second cluster node is added, the setup of the cluster is complete.

### 19.4.5 Steps for configuring a two-node file server cluster

To configure a two-node file server failover cluster, follow these steps:

1. To open the failover cluster snap-in, click **Start**, click **Administrative Tools**, and then click **Failover Cluster Management**.
2. In the console tree, if the cluster that you created is not displayed, right-click **Failover Cluster Management**, click **Manage a Cluster**, and then select the cluster you want to configure.
3. In the console tree, click the plus sign next to the cluster that you created to expand the items underneath it.
4. If the clustered servers are connected to a network that is not to be used for network communication in the cluster (for example, a network intended only for iSCSI), then under **Networks**, right-click that network, click **Properties**, and then click **Do not allow the cluster to use this network**. Click **OK**.
5. Click **Services and Applications**. Under **Actions** (on the right), click **Configure a Service or Application**.
6. Review the text on the first page of the wizard as shown in Figure 19-28 on page 579, and then click **Next**

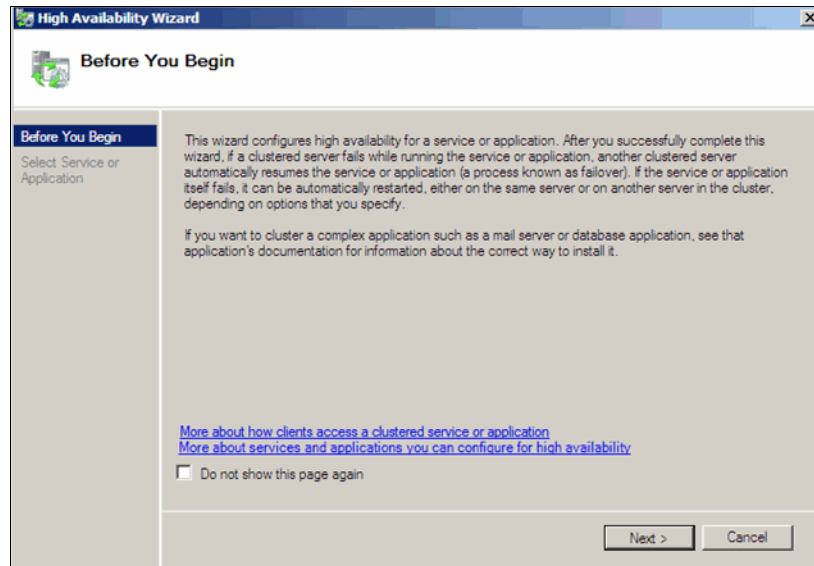


Figure 19-28 HA Wizard

7. Select **File Server** as shown in Figure 19-29, and then click **Next**.

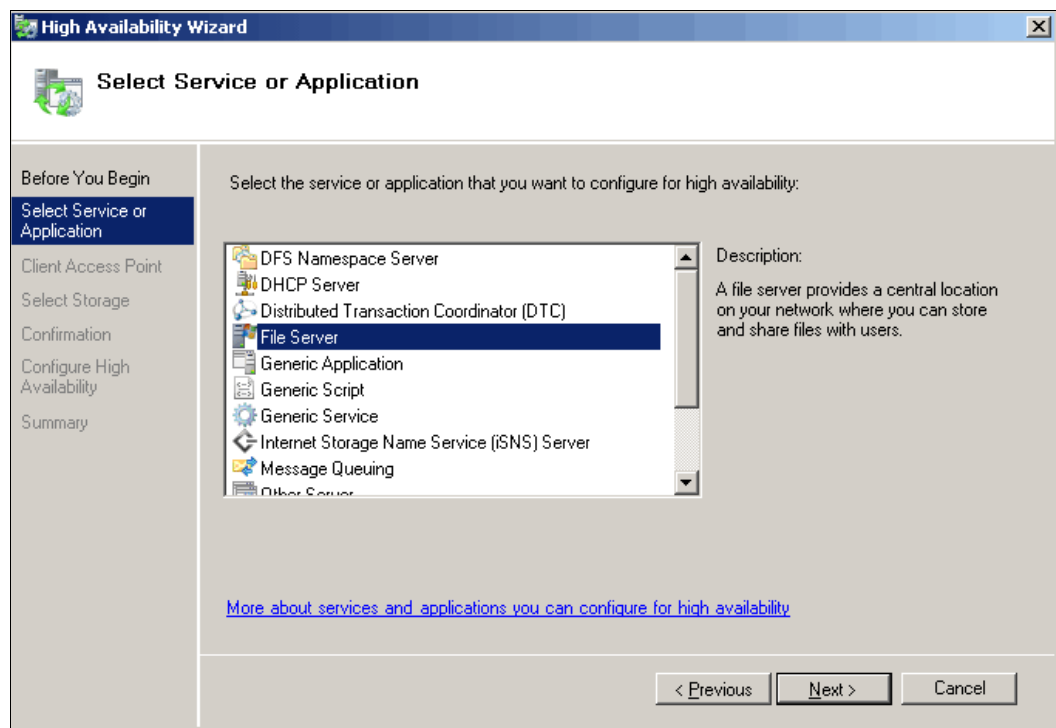


Figure 19-29 Select Service or Application: File Server

8. Follow the instructions in the wizard to specify a name for the clustered file server and an IP address that is not automatically supplied by your DHCP settings, for example, a static IPv4 address for this clustered file server as shown in Figure 19-30 on page 580.

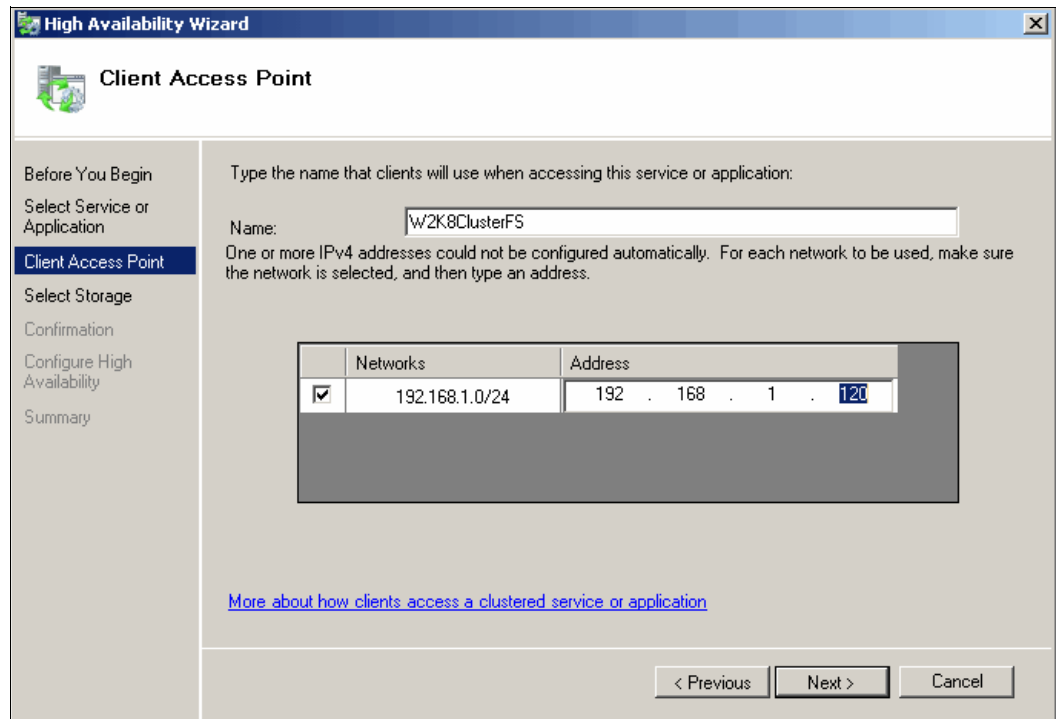


Figure 19-30 Specify a Name and an IP Address

9. Select the storage volume or volumes that the clustered file server should use as shown in Figure 19-31.




Figure 19-31 Select Storage Volume

10. After the wizard runs, the **Summary** page appears as shown in Figure 19-32. To close the wizard, click **Finish**.

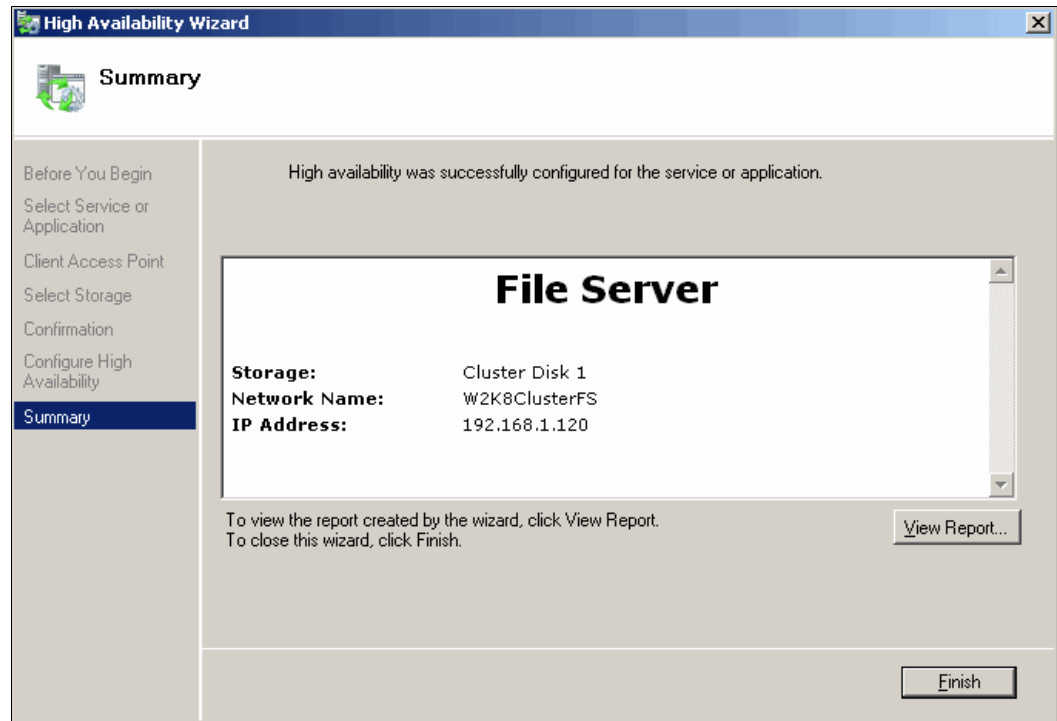


Figure 19-32 File Server Configured Successfully

11. Figure 19-33 on page 582 shows our just-created **File Server cluster** in **Failover Cluster Management** console.

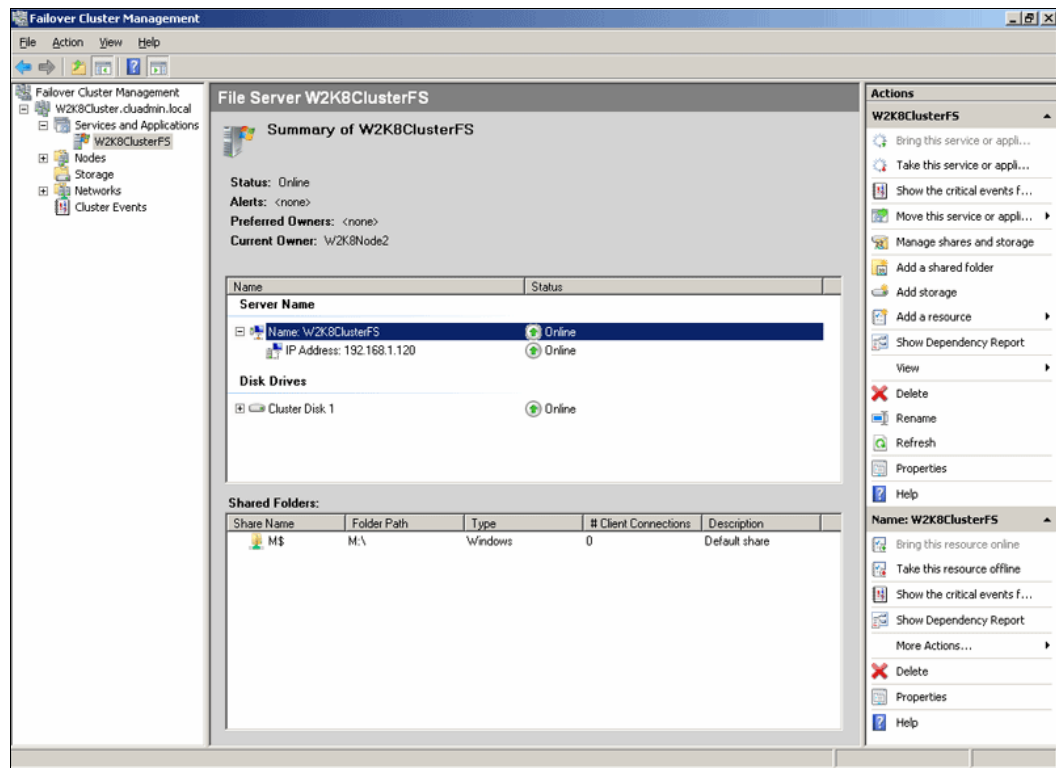


Figure 19-33 File Server Cluster Configuration Summary

- 12..In the console tree, make sure **Services and Applications** is expanded, and then select the clustered file server that you just created.
- 13.Under **Actions**, click **Add a shared folder**. The **Provision a Shared Folder Wizard** appears as shown in Figure 19-34 on page 583.



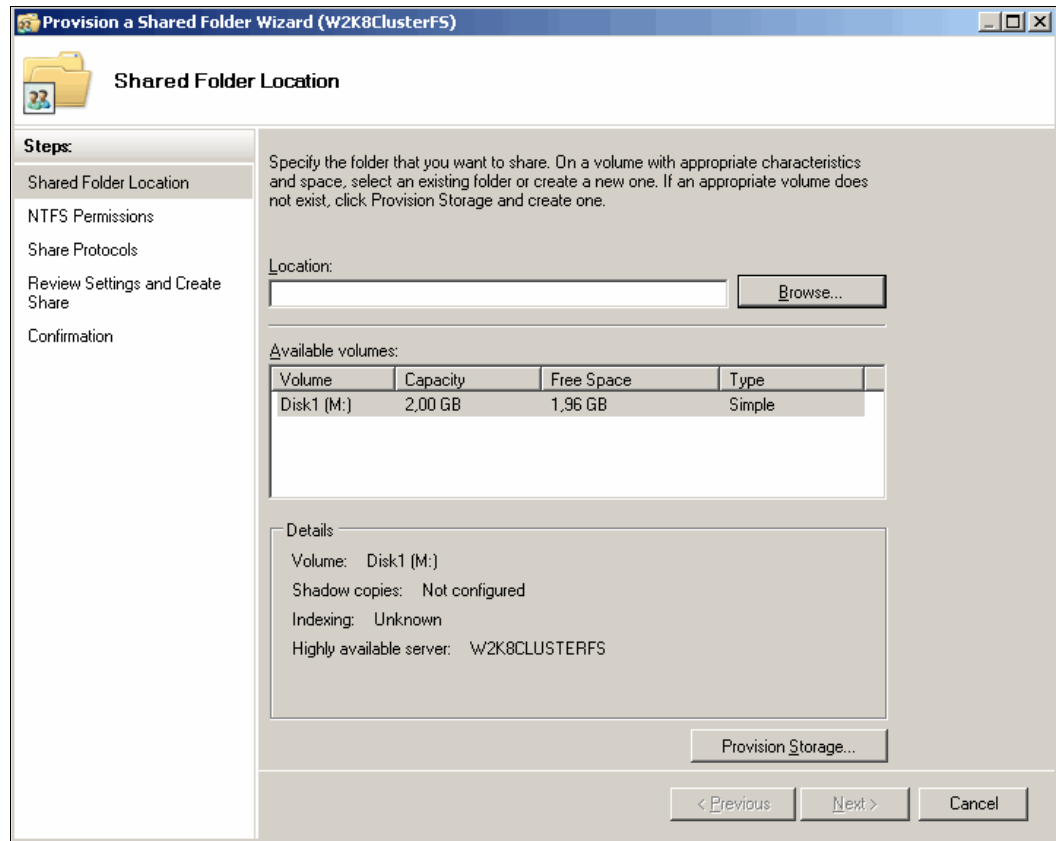


Figure 19-34 Specify the Share Folder Location

14. Follow the instructions in the wizard to specify the following settings for the shared folder:

- Path and name.
- NTFS permissions.
- Share name and path.
- Advanced settings for the SMB protocol (optional). SMB is used by Windows-based clients.
- SMB share permissions.
- DFS namespace publishing.

15. Review the settings and click **Create** to create the share as shown in Figure 19-35 on page 584.

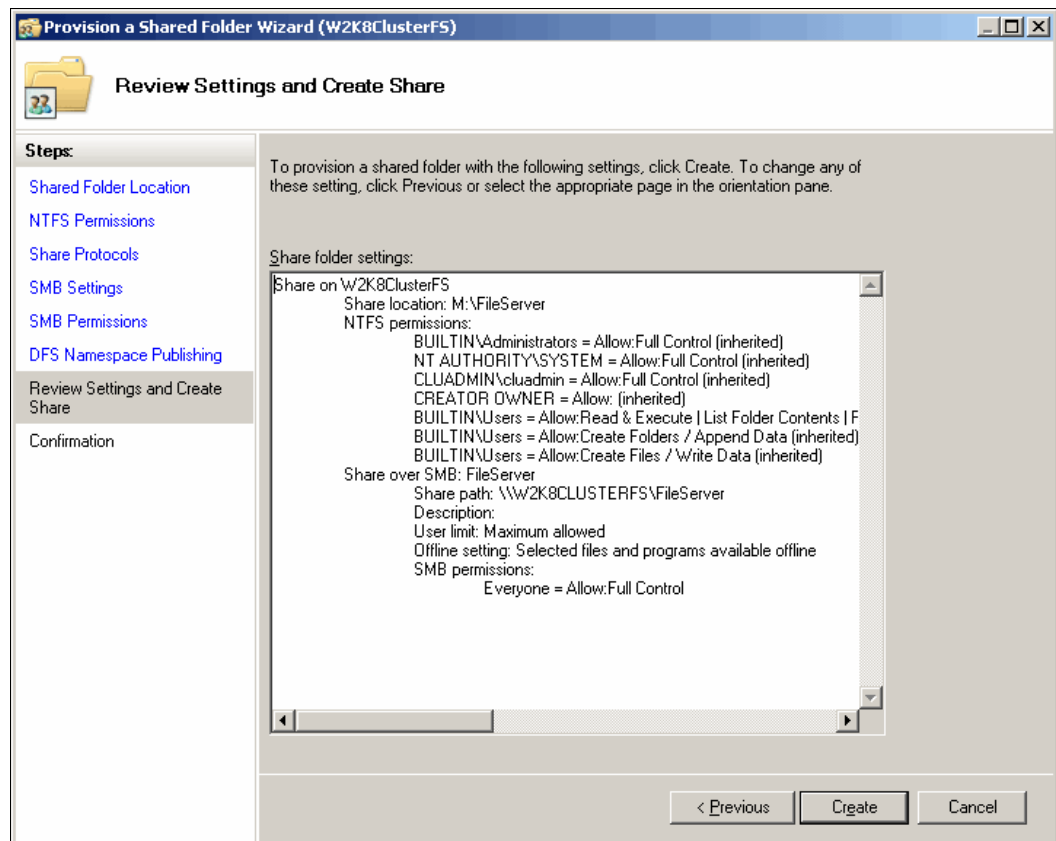


Figure 19-35 Create Share

16. Figure 19-36 on page 585 shows the Summary window of the file server cluster with configured share folders.

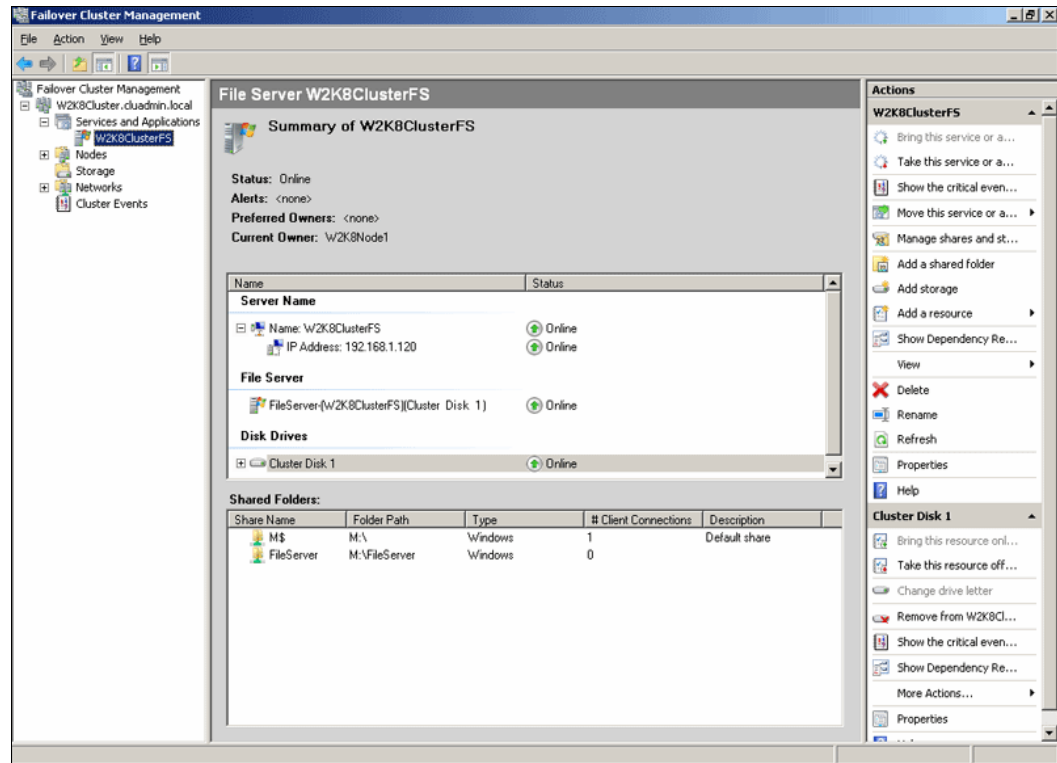


Figure 19-36 File server cluster with configured share folders

17. File server comes online. If it does not, review the state of the networks and storage and correct any issues. Then right-click the new clustered file server and click **Bring this service or application online**.

18. To perform a basic test of failover, right-click the clustered file server, click **Move this service or application to another node**, and click the node you want from the list of available nodes. When prompted, confirm your choice.

You can observe the status changes in the center pane of the snap-in as the clustered file server instance is moved.

We have installed Windows Server 2008 Failover Clustering on two nodes and configured a file server cluster with share folders.





## SuSE Linux configuration guide

In this chapter, we discuss sample configurations to show how to configure SuSE Linux Enterprise Server 11 to access a logical drive from the DS3500 storage subsystem. In the SAS sample configuration, we are using a HS21 with two SAS adapters and two SAS connectivity modules connected to a DS3500. For the iSCSI sample configuration we are using an HS22 with Ethernet software initiator HBAs and two Ethernet switches connected to DS3500. And for the FC sample configuration we are using an HS21 with two FC adapters, and two Brocade FC switches connected to a DS3500.

This chapter will include the following:

- ▶ 20.1, “DS3500 SAS storage configuration on SLES 11 using RDAC” on page 588
- ▶ 20.2, “DS3500 iSCSI storage configuration on SLES 11 using RDAC” on page 600
- ▶ 20.3, “DS3500 FC SAN boot configuration for SLES 11 server using RDAC” on page 609
- ▶ 20.4, “DS3500 FC storage configuration on SLES 11 using DMM” on page 616

**Note:** At the time of writing, iSCSI hardware initiators not supported with DS3500. Check the SSIC website for future enhancements.

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

## 20.1 DS3500 SAS storage configuration on SLES 11 using RDAC

In this section we discuss how to add a logical drive to an SLES 11 using DS3500 SAS storage using RDAC multipath driver

### 20.1.1 Preparing for the installation

For our setup, we used the following hardware and software components:

- ▶ HS21 BladeServer
- ▶ Dual Port 6 Gbps SAS HBA
- ▶ Two SAS Connectivity Modules
- ▶ DS3500 Storage Subsystem
- ▶ A management workstation running the latest version of DS Storage Manager
- ▶ The latest version IBM RDAC drivers for Linux
- ▶ SuSE Linux Enterprise 11




Figure 20-1 Hardware setup

In SAS configurations only, zoning configuration is not necessary except when booting from a SAS Storage. In our example, zoning is not required and we are attaching a logical drive to an operating system that is already installed on local hard disks.

The following are recommendations before mapping the logical drive(s) to the server:

- ▶ Check the compatibility in IBM System Storage Interoperation Center (SSIC) web site (<http://www-03.ibm.com/systems/support/storage/config/ssic>) to be sure it is a supported environment and also to be sure which multipath driver is needed (RDAC, MPIO, DMM, and so on). The SSIC can provide the interoperability matrix, and is a selection driven tool where you provide the criteria for your configuration.

- ▶ Download the latest multipath driver from the IBM support web site.
- ▶ Read carefully the multipath driver README file to see which kernel version is supported with that driver.
- ▶ Install the necessary packages in the Linux Server to build the multipath driver:
  - gcc-4.3-62.198
  - kernel-source-2.6.27.19-5.1
  - linux-kernel-headers-2.6.27-2.22
  - make-3.81-128.20
  - glibc-devel(64 bits) (only needed for 64 bit systems)
- ▶ For the systems that have LSI Logic HBAs (like in this case), INITRD\_MODULES should include mptbase and mptscsi (or mptscsih) for pre-3.02.xx version of LSI mpt driver. The INITRD\_MODULES should include mptbase, mptscsi, mptspi, mptsas, and mptfc for 3.02.xx or later version of LSI mpt driver.
- ▶ Although the system can have Fibre Channel HBAs from multiple vendors or multiple models of HBAs from the same vendor, only HBAs from same vendor can be connected to each storage array.
- ▶ Use udev by-id specification in the /etc/fstab file. The ID paths are device specific and will not change as a result of hardware changes. If you use specifications like /dev/sda for partitions, after mapping the logical drive, /dev/sda could probably not be the same hard disk. In this case, the operating system will try to boot from the storage instead of booting from the local hard disks, resulting in an error.

An example of an fstab file using by-id specification is shown in Example 20-1.

*Example 20-1 Example of fstab using by-id specification*

---

```
cat /etc/fstab
/dev/disk/by-id/scsi-35000c500009c6dcb-part1 swap swap defaults 0 0
/dev/disk/by-id/scsi-35000c500009c6dcb-part2 / ext3 acl,user_xattr 1 1
/dev/disk/by-id/scsi-35000c50002ebe6bf-part1 /windows/C ntfs-3g
users,gid=users,fmask=133,dmask=022,locale=en_US.UTF-8 0 0

proc /proc proc defaults 0 0
sysfs /sys sysfs noauto 0 0
debugfs /sys/kernel/debug debugfs noauto 0 0
usbfs /proc/bus/usb usbfs noauto 0 0
devpts /dev/pts devpts mode=0620,gid=5 0 0
```

---

It is possible to modify the udev specification by running the following steps (only applicable to SuSE):

- a. Run the following command to configure hard disk definitions:

```
yast disk
```

A window similar to that shown in Figure 20-2 on page 590 will appear.

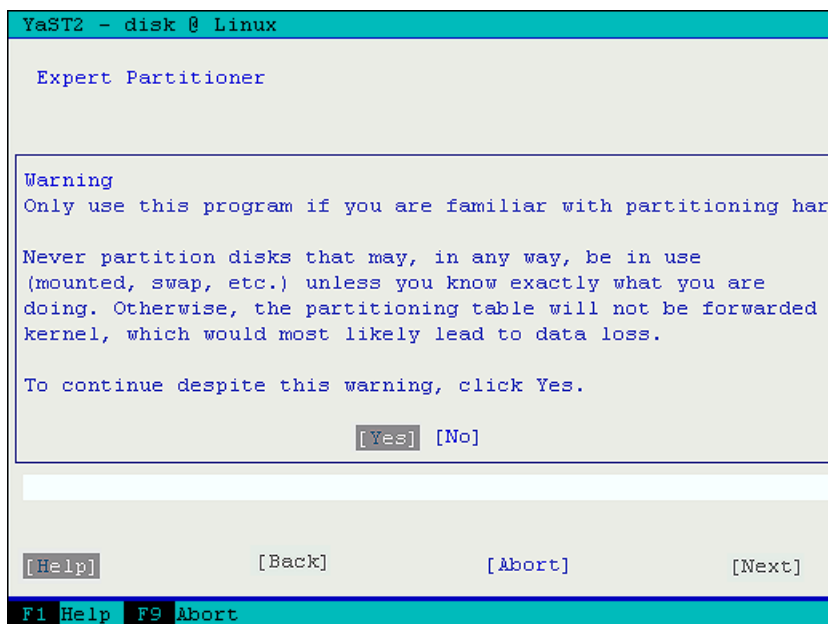


Figure 20-2 Enter YES to access Expert partitioner

- b. Press Enter on **Hard Disks** → Select a partition → **Edit** as shown in Figure 20-3. It is necessary to repeat this step and the next step on every partition (/dev/sda1 and /dev/sda2 are shown here).




Figure 20-3

- c. Select Fstab Options and press Enter as shown in Figure 20-4 on page 591.






Figure 20-4 Select Fstab Options

- d. Select Device ID and press F10 to finish as shown in Figure 20-5.




Figure 20-5 Select Device ID

- Verify the boot order in MPT SAS BIOS to boot from the local hard disks using the following procedure. If this setting are not correctly configured, the system will try to boot from the logical drive instead of the local hard disk.
  - a. Reboot the server and press Ctrl+C at MPT SAS POST as shown in Figure 20-6 on page 592.

```

LSI Logic Corp. MPT SAS BIOS
MPTBIOS-6.00.05.00 (2006.00.02)
Copyright 2000-2006 LSI Logic Corp.

Press Ctrl-C to start LSI Logic Configuration Utility...

```

Figure 20-6 Press Ctrl+C

- b. In this case we have a SAS1064E SAS adapter with the local hard disks and a SAS1064 SAS adapter with the external storage. Select the SAS adapter where the local hard disks are connected (SAS1064E). Alter the boot order as shown in Figure 20-7 by pressing Insert and Delete keys on your keyboard. The PCI Slot should have the lowest number.

| Adapter List Global Properties |         |         |         |          |               |         |            |
|--------------------------------|---------|---------|---------|----------|---------------|---------|------------|
| Adapter                        | PCI Bus | PCI Dev | PCI Fnc | PCI Slot | FW Revision   | Status  | Boot Order |
| SAS1064                        | 00      | 01      | 00      | 01       | 1.23.01.00-IR | Enabled | 1          |
| SAS1064E                       | 02      | 00      | 00      | 00       | 0.10.15.00-IR | Enabled | 0          |

Figure 20-7 Press Insert and Delete to alter the boot order

- Record the SAS Expansion Card (SAS1064) ID. The SAS ID is the unique address that identifies each device port. We need this address to map the logical drive to this specific server. It is possible to find it in the SAS adapter BIOS (Figure 20-7) by selecting the adapter and pressing Enter. It is also possible to get find the SAS ID from the BladeCenter AMM using the following procedure:
  - a. In the AMM console, select Hardware VPD as shown in Figure 20-8.

```

Monitors
 System Status
 Event Log
 LEDs
 Power Management
 Hardware VPD
 Firmware VPD
 Remote Chassis
Blade Tasks
 Power/Restart
 Remote Control
 Firmware Update
 Configuration
 Serial Over LAN
 Open Fabric Manager
I/O Module Tasks
Storage Tasks
MM Control
Service Tools

```

Figure 20-8 AMM main menu

- b. Click **Expansion Card** as shown in Figure 20-9 on page 593.

|                                    |                      |           |
|------------------------------------|----------------------|-----------|
| <a href="#">[6] Blade40</a>        | HS21 (Type 8853)     | Installed |
| <a href="#">[1] Expansion Card</a> | SAS Expansion Option | Installed |
| <a href="#">[1] Add-in Card</a>    | cKVM                 | Installed |

Figure 20-9 Click Expansion card

- c. Click the Ports tab to see the SAS address for each SAS adapter port as shown in Figure 20-10.

| BladeCenter Vital Product Data                                                                                                              |                         |
|---------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| Inventory                                                                                                                                   | Ports                   |
| <b>Expansion card MAC Addresses</b>                                                                                                         |                         |
| Property                                                                                                                                    | Value                   |
| None available                                                                                                                              |                         |
| <b>Expansion card Unique IDs</b>                                                                                                            |                         |
| ID Type                                                                                                                                     | Value                   |
| SAS ID 1                                                                                                                                    | 50:00:62:B0:00:08:B2:FC |
| SAS ID 2                                                                                                                                    | 50:00:62:B0:00:08:B2:FD |
| The ID Type in the Unique ID table(s) shows the type of unit used for Fibre channel devices, GUID (Globally Unique ID) for SAS I/O devices. |                         |

Figure 20-10 Click the Ports tab to see the SAS ID's values

## 20.1.2 Installing the RDAC Multipath Driver

In this section we discuss how to install the RDAC multipath driver in SuSE Linux Enterprise Server 11 we downloaded in 20.1.1, “Preparing for the installation” on page 588.

According to the multipath driver README file, the driver version to install according to our kernel version (2.6.27.19-5) is 09.03.0C05.0331.

Because SAS configuration does not need zoning, we must install the RDAC driver before mapping the logical drive to the operating system to avoid buffer I/O errors on disks.

### RDAC driver building

The RDAC drivers are released as a source-code package in the format of a compressed tar file. To unpack the source package and build the source-code, follow these steps:

1. The following command is used to unpack the file.

```
tar -zxvf rdac-LINUX-09.03.0C05.0331-source.tar.gz
```

The source files will be unpacked in the `linuxrdac-09.03.0C05.0331` subdirectory. When it finishes, change to this directory.

2. The following command is used to remove old driver modules:

```
make clean
```

3. The following command is used to compile the driver modules and build the RDAC driver:

```
make
```

## RDAC driver installation

Perform the following steps to install the RDAC drivers:

1. The following command is used to uninstall the old RDAC driver if previously installed:

```
make uninstall
```

2. The following command is used to install the new RDAC driver:

```
make install
```

This command will copy the driver modules to the kernel module tree and build the new RAMdisk image (mpp-2.6.27.19-5-default.img) that includes the RDAC driver modules and all driver modules that are needed at boot. Without these modules the operating system will not be able to find the disks.

In certain cases, for advanced troubleshooting, it is possible to first copy the mpp-2.6.27.19-5-default.img file to another directory and unpack it with the following command:

```
gunzip -c /new_directory/mpp-2.6.27.19-5-default.img |cpio -i --make directories
```

The following command will pack the image file again:

```
find . | cpio -H newc --quiet -o |gzip -9 > mpp-2.6.27.19-5-default.img ||exit 1
```

3. Now that we have the initrd image, it is necessary to modify the boot loader (grub) to load the image when the server starts. For that, we need to edit the /boot/grub/menu.lst file and add the line shown in Example 20-2.

*Example 20-2 Edit the /boot/grub/menu.lst*

---

```
RDAC Driver
title SUSE Linux (2.6.27.19-5-default) with MPP support
 root (hd0,1)
 kernel /boot/vmlinuz root=/dev/disk/by-id/scsi-35000c500009c6dcb-part2
 splash=silent resume=/dev/disk/by-id/scsi-35000c500009c6dcb-part1 elevator=cfq
 showopts vga=0x317
 initrd /boot/mpp-2.6.27.19-5-default.img
```

---

Take the information for the root and resume parameters from the previous line by default.

4. Reboot the server and when the grub menu appears select the SUSE Linux (2.6.27.19-5-default) with MPP support option as shown in Figure 20-11. It is possible to make this selection as default by changing the default parameter in the menu.lst file.




Figure 20-11 Boot menu when server starts after RDAC installation

### 20.1.3 Setting up the DS3500 logical drives and host mapping

Map the logical drive(s) to the SuSE operating system by using the normal procedure as described in Chapter 11, “Administration - Mappings Tab” on page 275.

One of the important things to consider is using the correct Host type when mapping the logical drive to the host from the IBM Storage Manager as shown in Figure 20-12.

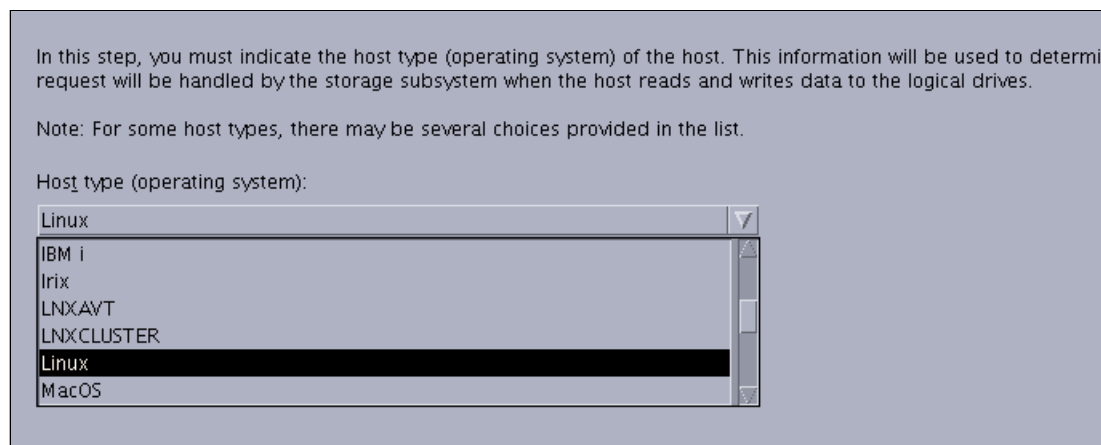


Figure 20-12 Defining the Host Type of the host

When you define the host port, the host type of the attached host is defined as well. Through this information, the DS Storage System modifies the NVSRAM settings to behave according to the host type defined.

The most common of the specifications that differ according to the host type setting is the Auto Volume Transfer function (AVT), which is set as enable for certain host types, and disable for others. It is important to choose the correct host type from the list of available types because this is the part of the configuration dealing with heterogeneous host support. Each operating system expects slight variations in settings and can handle SCSI commands separately. Incorrect selections can result in failure to boot, or loss of path failover function when attached to the storage server.

For RDAC and DMM, always use the Linux host type. When you are using clustering services, the correct host type is LNXCluster.

The LNXAVT option is used when third-party multipath drivers are installed. For example, to use Dynamic Multi-pathing (DMP) as the default multipath driver. DMP is part of VERITAS Volume Manager.

#### 20.1.4 Scan and verify the storage logical drive

During the server restart, you will see a list of devices listed below the corresponding LSI SAS adapter as shown in Figure 20-13 on page 596. This shows that the SAS adapter installed in slot 1 is accessing the same logical drive (LUN 0) from the IBM System Storage DS3500 by four different paths. One path for each SAS port connected to the Storage is shown in Figure 20-1 on page 588.

```

LSI Logic Corp. MPT SAS BIOS
MPTBIOS-6.08.05.00 (2006.08.02)
Copyright 2000-2006 LSI Logic Corp.

```

| SLOT | ID | LUN | VENDOR   | PRODUCT     | REVISION   | INT13    | SIZE \ NV |
|------|----|-----|----------|-------------|------------|----------|-----------|
| 0    | 0  | 0   | IBM-ESXS | ST973401SS  | B51D       | Boot     | 68 GB     |
| 0    |    |     | LSILogic | SAS1064E-IR | 0.10.15.00 | NV 25:00 |           |
| 1    | 23 | 0   | IBM      | 1746 FASTT  | 1070       | 81h      | 60 GB     |
| 1    | 24 | 0   | IBM      | 1746 FASTT  | 1070       | 82h      | 60 GB     |
| 1    | 25 | 0   | IBM      | 1746 FASTT  | 1070       | 83h      | 60 GB     |
| 1    | 26 | 0   | IBM      | 1746 FASTT  | 1070       | 84h      | 60 GB     |
| 1    |    |     | LSILogic | SAS1064-IR  | 1.23.01.00 | NV 2B:00 |           |

```

LSI Logic Corp. MPT boot ROM successfully installed!

```

Figure 20-13 SAS adapters device list

At this point, we already should have the necessary modules in the Linux OS to scan the logical drives. The `lsmod` command can be used to verify that the RDAC (MPP) modules are running. The output of the command will list the following modules:

- ▶ `mppUpper`
- ▶ `mppVhba`
- ▶ `sg`
- ▶ `sd_mod`
- ▶ `scsi_mod`

If all of these modules are running, the RDAC should be able to manage access to the DS3500 logical drive. It is possible to verify this using the following procedure.

A similar directory structure as shown in Example 20-3 is created by the RDAC multipath driver. We will have a directory for each Storage Controller (A and B), one directory for each controller port and one file for each LUN showing its status similar to Example 20-4.

*Example 20-3 Directory structure created by RDAC multipath driver*

```

/proc/mpp/DS3524-FC/controllerA/mptsas_h0c0t2/LUN0
/proc/mpp/DS3524-FC/controllerA/mptsas_h0c0t1/LUN0

/proc/mpp/DS3524-FC/controllerB/mptsas_h0c0t3/LUN0
/proc/mpp/DS3524-FC/controllerB/mptsas_h0c0t0/LUN0

```

*Example 20-4 LUN0 file content*

```

Linux MPP driver. Version:09.03.0C05.0331 Build:Wed Apr 14 11:37:55 CDT 2010
Lun WWN:60080e50001b0c9e00000f694c60d855
Physical HBA driver: mptsas
Device Scsi Address: host_no:0 channel:0 target:0 Lun:0
Queue Depth = 64
I/O Statistics:
 Number of I/Os:0
 Longest trip of all I/Os:0
 Shortest trip of all I/Os:0

```

```

 Number of occurrences of IO failed events:0
Device state: [0] OPTIMAL
Device state: [1] OPTIMAL
Device state: [2] OPTIMAL
Device state: [3] OPTIMAL
Device state: [4] OPTIMAL
Device state: [5] OPTIMAL
Device state: [6] OPTIMAL
Device state: [7] OPTIMAL
Device state: [8] OPTIMAL
Device state: [9] OPTIMAL
Path state:[9] OPTIMAL
Path state:[0] OPTIMAL
Path state:[1] OPTIMAL_NEED_CHECK
Path state:[2] OPTIMAL_CHECKING
Path state:[3] OPTIMAL
Path state:[4] OPTIMAL_NEED_CHECK
Path state:[5] OPTIMAL_CHECKING
Path state:[6] OPTIMAL
Path state:[7] OPTIMAL_NEED_CHECK
Path state:[8] OPTIMAL_CHECKING
Controller Failed? 0
Outstanding IOs on this device:
total size:1136

```

---

The WWN LUN shown in Example 20-4 on page 596 (in this case 60080e50001b0c9e00000f694c60d855) should be listed in the output of `ls /dev/disk/by-id` as shown in Example 20-5. The word `scsi-3` before the LUN number is the iSCSI standard.

*Example 20-5* `ls /dev/disk/by-id` output

---

```

drwxr-xr-x 2 root root 180 Aug 12 07:08 .
drwxr-xr-x 5 root root 100 Aug 12 07:07 ..
lrwxrwxrwx 1 root root 9 Aug 12 07:08 edd-int13_dev80 -> ../../sda
lrwxrwxrwx 1 root root 10 Aug 12 07:08 edd-int13_dev80-part1 -> ../../sda1
lrwxrwxrwx 1 root root 10 Aug 12 07:08 edd-int13_dev80-part2 -> ../../sda2
lrwxrwxrwx 1 root root 9 Aug 12 07:07 scsi-35000c500009c6dcb -> ../../sda
lrwxrwxrwx 1 root root 10 Aug 12 07:07 scsi-35000c500009c6dcb-part1 -> ../../sda1
lrwxrwxrwx 1 root root 10 Aug 12 07:07 scsi-35000c500009c6dcb-part2 -> ../../sda2
lrwxrwxrwx 1 root root 9 Aug 12 07:07 scsi-360080e50001b0c9e00000f694c60d855 ->
../../sdb

```

---

The RDAC driver installs three additional scripts to interact with the logical drives:

- ▶ **mppBusRescan**

This script allows the addition of new logical drives without having to reboot the server. It is possible to discover new array devices, to delete unmapped or disconnected devices, and to update the size of the devices.

- ▶ **mppUpdate**

This script uses the `mppUtil` utility to discover the RDAC assigned virtual target IDs of the attached arrays. It then updates the RDAC driver configuration file (`/var/mpp/devicemapping`) so the entries will be persistent after a reboot. The first time the RDAC driver sees a storage array, it will arbitrarily assign a target ID for the virtual target that represents the storage array. At this point the target ID assignment is not persistent

and it could change on a reboot. The mppUpdate utility updates the RDAC driver configuration files so that these target ID assignments are persistent and will not change across reboots. After it is made persistent, the user does not have to worry about the possibility of device names changing and invalidating mount points. mppUpdate needs to be executed after an array is added to or removed from the system. If an array is added to the system, mppBusRescan also needs to be run before mppUpdate is run. It also needs to be executed whenever user changes the RDAC configuration file /etc/mpp.conf to rebuild the new configuration file into RAM Disk image that will be used for the next reboot.

► mppUtil

This general purpose script is used with the RDAC driver to perform the various functions provided by the driver. With this script it is possible to:

- Display the information about all the arrays currently detected by the host.
- Set the current debug reporting level to debug.
- Report the Up/Down state of the controllers and paths for each LUN in real-time.

Probably, the most used parameters with this script is in the Example 20-6. It is showing the same info as displayed in Figure 20-13 on page 596 but from an operating system perspective.

*Example 20-6 mppUtil -g 0 output to see the available paths to the logical drive*

---

```
mppUtil -g 0

Controller 'A' Path

NumLunObjects: 2 RoundRobinIndex: 0
 Path #1: LunPathDevice: present
 DevState: OPTIMAL
 RemoveState: 0x0 StartState: 0x1 PowerState: 0x0
 Path #2: LunPathDevice: present
 DevState: OPTIMAL
 RemoveState: 0x0 StartState: 0x1 PowerState: 0x0

Controller 'B' Path

NumLunObjects: 2 RoundRobinIndex: 0
 Path #1: LunPathDevice: present
 DevState: OPTIMAL
 RemoveState: 0x0 StartState: 0x1 PowerState: 0x0
 Path #2: LunPathDevice: present
 DevState: OPTIMAL
 RemoveState: 0x0 StartState: 0x1 PowerState: 0x0
```

---

Another way to list the all available SCSI devices is by using the **lsscsi** command as shown in Example 20-7. This will show each path and the Virtual Disk created by RDAC.

*Example 20-7 lsscsi command*

---

```
lsscsi
[0:0:0:0] disk IBM-ESXS ST9500430SS BD11 /dev/sda
[1:0:0:0] cd/dvd HL-DT-ST RW/DVD GCC-T10N 1.00 /dev/sr0
[2:0:0:0] disk IBM 1746 FASTT 1070 -
[3:0:0:0] disk IBM 1746 FASTT 1070 -
[4:0:0:0] disk IBM 1746 FASTT 1070 -
[5:0:0:0] disk IBM VirtualDisk 1070 /dev/sdb
```

---



```
[6:0:0:0] disk IBM 1746 FAStT 1070 -
```

---

Now the SuSE OS should see the new /dev/sdb disk as shown in Example 20-8.

*Example 20-8 fdisk -l Output*

---

```
fdisk -l
```

```
Disk /dev/sda: 73.4 GB, 73407488000 bytes
255 heads, 63 sectors/track, 8924 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Disk identifier: 0x04b643b3
```

| Device    | Boot | Start | End  | Blocks   | Id | System               |
|-----------|------|-------|------|----------|----|----------------------|
| /dev/sda1 |      | 1     | 262  | 2104483+ | 82 | Linux swap / Solaris |
| /dev/sda2 | *    | 263   | 8924 | 69577515 | 83 | Linux                |

```
Disk /dev/sdb: 64.4 GB, 64424509440 bytes
64 heads, 32 sectors/track, 61440 cylinders
Units = cylinders of 2048 * 512 = 1048576 bytes
Disk identifier: 0x00000000
```

```
Disk /dev/sdb doesn't contain a valid partition table
```

---

## 20.1.5 Configuring RDAC (MPP)

The Redundant Disk Array Controller Driver (a.k.a. Multi-Path Proxy Driver or MPP) adds a new layer in the Linux OS driver stack. It is a combination of two drivers, one an upper level driver, and the other a virtual Host Bus Adapter (HBA) driver:

- ▶ The upper level driver has the same design role as an "sd" or "sg" driver on Linux. It is responsible for the management of multiple paths to the disk array.
- ▶ The virtual HBA driver is a Linux SCSI lower level driver. It sits at the same level of the driver stack as any driver for a physical QLogic or LSI Fibre Channel HBA.

The virtual HBA driver creates virtual disks to represent the array LUNs, which the Linux OS disk target driver "sd" will attach to. The array devices are then addressed through the virtual host/LUN nodes. When an array/LUN has multiple possible paths with which it can be accessed, the RDAC driver picks the correct path and routes I/O to it, performing any necessary failover. This is transparent to the OS. This driver provides such capabilities as I/O path load balancing and transparent failover in the event of a path failure. The virtual HBA driver also handles support in a server clustering environment, by translating SCSI-2 reservation/release commands to SCSI-3 persistent reservations.

UTM (Universal Transport Mechanism) LUNs are not represented as virtual disks. Thus the RDAC driver does not create any special device nodes of its own for the UTM devices. Instead, it allows the sg driver (the generic scsi driver on Linux) to claim the UTM physical Luns, thus creating /dev/sg\* entries, which can be discovered and used by the SMagent or SANtricity UTM agent.

The following command shows a description of each parameter that the user can alter to affect the RDAC execution:

```
man rdac
```

The system-specific mechanism for reconfiguring the RDAC driver is the following:

1. Change the necessary parameters located in `/etc/mpp.conf`.
2. Run `mppUpdate` to build the new `initrd` images with the changes.
3. A driver reload or system reboot is necessary after the changes are made to the configuration file

## **20.2 DS3500 iSCSI storage configuration on SLES 11 using RDAC**

In this section we discuss how to add a logical drive to an SLES 11 using DS3500 iSCSI storage using RDAC multipath driver.

### **20.2.1 Preparing for the installation**

For our setup, we used the following hardware and software components:

- ▶ HS22 BladeServer.
- ▶ Dual Port 1 Gbps Ethernet HBA.
- ▶ Two additional Ethernet switches for BladeCenters
- ▶ DS3500 Storage Subsystem
- ▶ A management workstation running the latest version of DS Storage Manager
- ▶ The latest version IBM RDAC drivers for Linux
- ▶ SuSE Linux Enterprise 11
- ▶ External Ethernet switch




Figure 20-14 Hardware setup

The following are recommendations to do before mapping the logical drives.

- ▶ Check the compatibility in IBM System Storage Interoperation Center (SSIC) web site <http://www.ibm.com/systems/support/storage/config/ssic> to be sure it is a supported environment and to find out which multipath driver is needed (RDAC, MPIO, DMM, and so on).
- ▶ Download the latest multipath driver from the IBM support web site.
- ▶ Read carefully the multipath driver README file to see which kernel version is supported with that driver.
- ▶ For iSCSI there is no specific zoning configuration to be done on the switches. Isolate the iSCSI traffic from the regular Ethernet traffic by using separate physical switches or by configuring VLANs.

**Note:** Storage subsystems with iSCSI ports do not support direct-attached connections from the host systems to the storage subsystem iSCSI ports.

- ▶ Use different subnets for each Storage port in the same controller as shown in Figure 20-15 on page 602. In our case we use network 192.168.130.0/24 for subnet 1 and 192.168.131.0/24 for subnet 2.




Figure 20-15 iSCSI network subnets

- ▶ Use udev by-id specification in `/etc/fstab` file. A detailed procedure is shown in 20.1.1, “Preparing for the installation” on page 588
- ▶ Install the RDAC multipath driver before mapping the logical drive(s) to the server. The complete procedure is in 20.1.2, “Installing the RDAC Multipath Driver” on page 593.

## 20.2.2 Configuring iSCSI software initiator with YaST

iSCSI initiator (also called client) can be used to connect to any iSCSI target. The configuration of iSCSI initiator includes the discovery of available iSCSI targets and the setup of an iSCSI session. Both can be done with YaST.

1. Configure each Ethernet adapter in the blade server to access each iSCSI host port on the DS3500. In our case we have two Ethernet adapters installed on the server and configured to the 192.168.130.0 and 192.168.131.0 subnets. Verify that all the iSCSI host ports are accessible by using the **ping** command from the server.
2. Now it is necessary to configure the iSCSI initiator that is also called client. On a Linux terminal type `yast iscsi-client`. A panel for installing the necessary packages will appear as shown in Figure 20-16 on page 603. Press Enter to select Install. You need to have the SLES CD-ROM installation media or an online software repository previously configured.

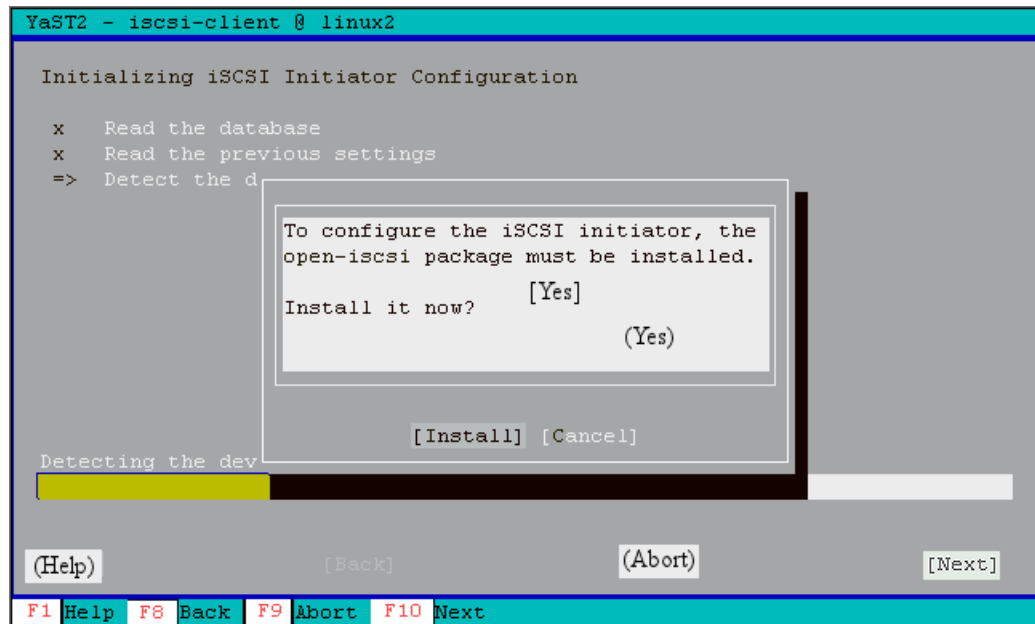


Figure 20-16 Press Enter key to install the necessary packages

3. The initiator name is automatically generated by SuSE. This is the Initiator name that Storage Manager will discover. Select the option **When booting** to automatically start the service. The panel should be similar to Figure 20-17.




Figure 20-17 Select When Booting

4. Map the logical drive to the server using the procedure documented in Chapter 11, "Administration - Mappings Tab" on page 275.
5. Although it is possible to discover all the targets, by discovering only the first address of iSCSI host port, we will type each ip address manually in Connected targets tab. We will demonstrate the other procedure to configure it manually by using the command line on 20.2.3, "Configuring iSCSI software initiator manually" on page 607. Select Connected

Targets and select Add as shown in Figure 20-18 to add each controller port. Each port will be a path to the logical drive.

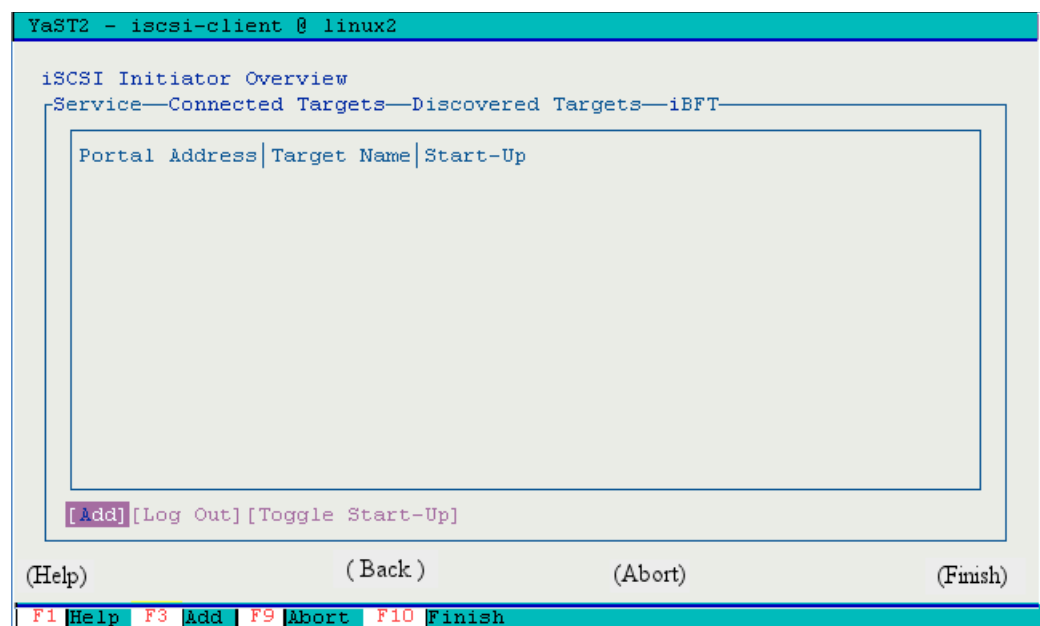


Figure 20-18 Select Add to add each Storage controller port

6. Type the IP address of one of the Storage iSCSI host port as shown in Figure 20-19. The TCP listening port by default is 3260. In our case we had configured previously a target authentication in the IBM Storage Manager. The Username is always the same as the initiator name (in this case iqn.1996-04.de.suse:01:27dc9458a936) and the password is the same as we provided in the IBM Storage Manager as shown in Figure 20-20 on page 605. For more information refer to Chapter 13, “Administration - iSCSI” on page 309.




Figure 20-19 IP address of the Storage iSCSI host port and authentication




Figure 20-20 iSCSI Target authentication on IBM Storage Manager

7. Select Connect to log in to the iSCSI storage as shown in Figure 20-21.

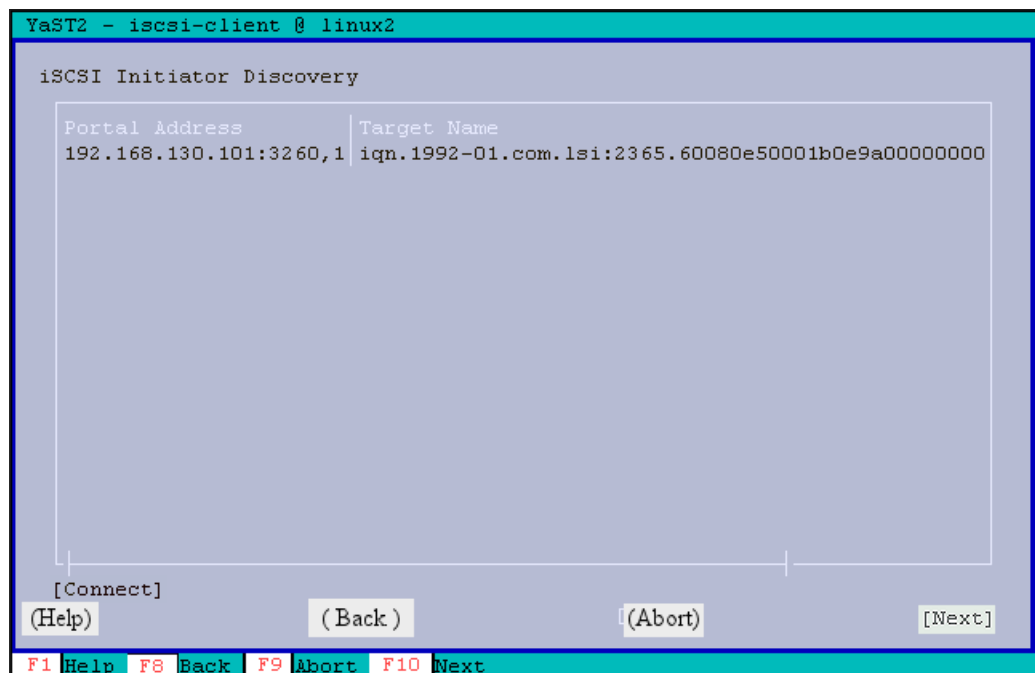


Figure 20-21 Select Connect to login

8. Configure the authentication settings as in step 6) and select Next to continue as shown in Figure 20-22.

```

YaST2 - iscsi-client @ linux2

iSCSI Initiator Discovery

[] No Authentication

[] Incoming Authentication
Username Password

[x] Outgoing Authentication
Username Password
iqn.1996-04.de.suse:01:27dc9458a9><*****

[Help] [Back] [Abort] [Next]

F1 Help F8 Back F9 Abort F10 Next

```

Figure 20-22 Select Next

9. It is necessary to repeat all the previous steps for configure each iSCSI controller port. At the end, the panel should be similar to Figure 20-23.

```

YaST2 - iscsi-client @ linux2

iSCSI Initiator Overview
Service—Connected Targets—Discovered Targets—iBFT—

Portal Address | Target Name
192.168.130.101:3260,1 | iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000
192.168.130.102:3260,2 | iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000
192.168.131.101:3260,1 | iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000
192.168.131.102:3260,2 | iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000

[Add] [Log Out] [Toggle Start-Up]

[Help] [Back] [Abort] [Finish]

F1 Help F3 Add F9 Abort F10 Finish

```

Figure 20-23 All the iSCSI host ports connected

10. For each connected target select Toggle Start-up to keep the connection after rebooting the server as shown in Figure 20-24 on page 607. It means that the target will be automatically connected when the service is started.



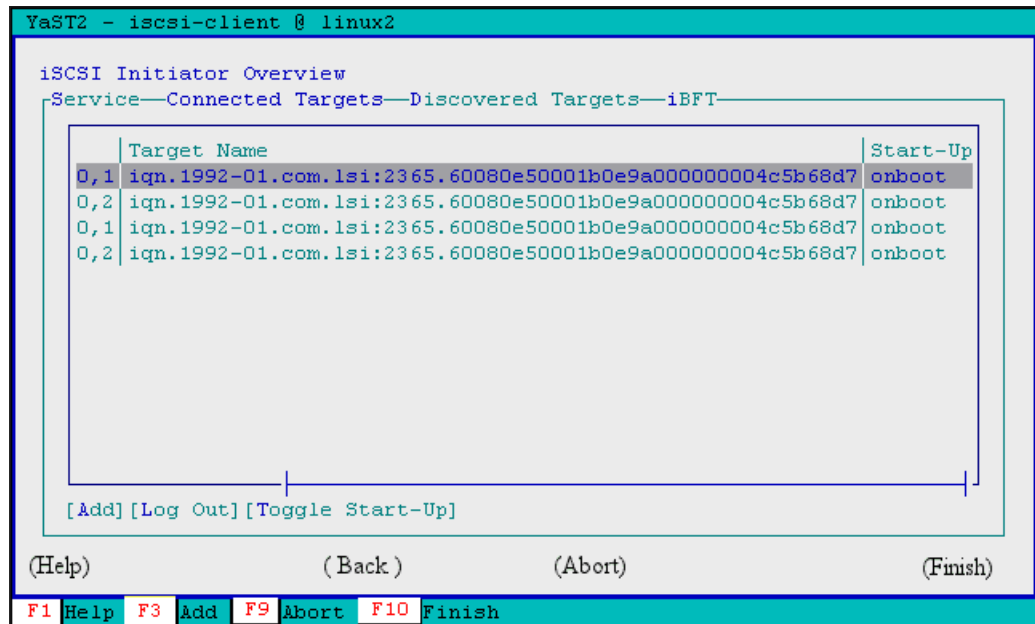


Figure 20-24 Be sure to select startup on boot for each target

11. At this point it is possible to map the logical drive to the SuSE server. A complete procedure is in Chapter 11, “Administration - Mappings Tab” on page 275. Because this procedure is for software initiator, it is the operating system that assigns an initiator address, in this case `iqn.1996-04.de.suse:01:27dc9458a936`. This address is the one that identifies the host and this will be used for mapping the LUN in IBM System Storage Manager. This name is specified also in the `/etc/iscsi/initiatorname.iscsi` file.
12. After following your normal procedure to create the filesystems, edit the `/etc/fstab` file and add the following line to mount the iSCSI target automatically:  
`/dev/sdb1 /mnt/ ext3 acl,user,nofail 0 0`

**Note:** In previous versions SLES versions, you could add the `hotplug` option to the device in the `/etc/fstab` file to mount the iSCSI targets. For SLES 11, the `hotplug` option is no longer available and you have to use the `nofail` option instead.

13. Verify the correct disk configuration. A detailed verification procedure is shown in 20.1.4, “Scan and verify the storage logical drive” on page 595.

### 20.2.3 Configuring iSCSI software initiator manually

YaST is a configuration tool only featured in the SuSE Linux distributions. In this section, we describe the manual procedure, which is also applicable to other distributions.

1. If the target is password protected, it is necessary to provide this information to `iscsid` by adding the following lines to the end of the `/etc/iscsid.conf`. The username should be the same as the `iqn` name.  

```
discovery.sendtargets.auth.authmethod = CHAP
discovery.sendtargets.auth.username = <username>
discovery.sendtargets.auth.password = <password>
```
2. To discover all the available targets, we only have to type one of the target ip addresses, and the rest of the ip addresses will be found during the discovery. These values are also

saved in an internal persistent database. Type the following command to run the discovery:

```
iscsiadm -m discovery -t st -p 192.168.130.101
```

In the output of the previous command, a line for each target defined will be shown as in Example 20-9. These targets could be reachable or not.

*Example 20-9 Discovery output*

---

```
192.168.132.102:3260,2iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7
192.168.130.101:3260,1iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7
192.168.130.102:3260,2iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7
192.168.131.102:3260,2iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7
192.168.131.101:3260,1iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7
192.168.133.102:3260,2iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7
192.168.133.101:3260,1iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7
192.168.132.101:3260,1iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7
```

---

3. By typing the **iscsiadm** command with the **--login** parameter, the initiator will try to log in to each of the targets as shown in Example 20-10.

*Example 20-10 Login to each target*

---

```
iscsiadm -m node --login
Logging in to [iface: default, target:
iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7, portal:
192.168.132.102,3260]
Logging in to [iface: default, target:
iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7, portal:
192.168.130.101,3260]
Logging in to [iface: default, target:
iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7, portal:
192.168.130.102,3260]
Logging in to [iface: default, target:
iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7, portal:
192.168.131.102,3260]
Logging in to [iface: default, target:
iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7, portal:
192.168.131.101,3260]
Logging in to [iface: default, target:
iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7, portal:
192.168.133.102,3260]
Logging in to [iface: default, target:
iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7, portal:
192.168.133.101,3260]
Logging in to [iface: default, target:
iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7, portal:
192.168.132.101,3260]
iscsiadm: Could not login to [iface: default, target:
iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7, portal:
192.168.132.102,3260]:
iscsiadm: initiator reported error (8 - connection timed out)
Login to [iface: default, target:
iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7, portal:
192.168.130.101,3260]: successful
```

```

Login to [iface: default, target:
iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7, portal:
192.168.130.102,3260]: successful
Login to [iface: default, target:
iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7, portal:
192.168.131.102,3260]: successful
Login to [iface: default, target:
iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7, portal:
192.168.131.101,3260]: successful
iscsiadm: Could not login to [iface: default, target:
iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7, portal:
192.168.133.102,3260]:
iscsiadm: initiator reported error (8 - connection timed out)
iscsiadm: Could not login to [iface: default, target:
iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7, portal:
192.168.133.101,3260]:
iscsiadm: initiator reported error (8 - connection timed out)
iscsiadm: Could not login to [iface: default, target:
iqn.1992-01.com.lsi:2365.60080e50001b0e9a000000004c5b68d7, portal:
192.168.132.101,3260]:
iscsiadm: initiator reported error (8 - connection timed out)

```

---

4. After following your normal procedure to create the filesystems, edit the `/etc/fstab` file and add the following line to mount the iSCSI target automatically:

```
/dev/sdb1 /mnt/ ext3 acl,user,nofail 0 0
```

**Note:** In previous versions SLES versions, you could add the `hotplug` option to the device in the `/etc/fstab` file to mount the iSCSI targets. For SLES 11, the `hotplug` option is no longer available and you have to use the `nofail` option instead.

5. Now verify that the disk configuration is correct. A detailed verification procedure is shown in 20.1.4, “Scan and verify the storage logical drive” on page 595.

## 20.3 DS3500 FC SAN boot configuration for SLES 11 server using RDAC

In this section we discuss how to configure an HS21 BladeServer to be able to boot from a DS3500 storage subsystem using FC HBAs on a SuSE Enterprise Server environment.

### 20.3.1 Preparing for the installation

For our setup, we used the following hardware and software components:

- ▶ HS21 BladeServer
- ▶ Dual Port 4Gbps FiberChannel HBA
- ▶ Two FiberChannel switches
- ▶ DS3500 Storage Subsystem
- ▶ A management workstation running the latest version of DS3500 Storage Manager
- ▶ The latest version IBM RDAC drivers for Linux

► SuSE Linux Enterprise 11

The installation can be done by performing the following steps:

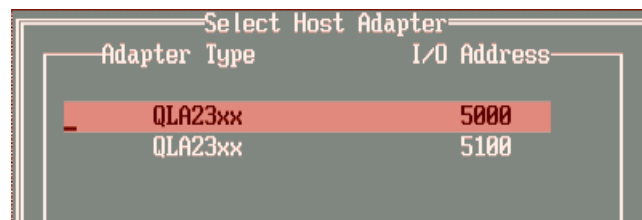
1. Check the compatibility in IBM System Storage Interoperation Center web site (<http://www-03.ibm.com/systems/support/storage/config/ssic>) to be sure it is a supported environment and also to be sure which multipath driver is needed (RDAC, MPIO, DMM, and so on). The SSIC can provide the interoperability matrix, and is a selection driven tool where you provide the criteria for your configuration.
2. Download the latest multipath driver from the IBM support web site and read carefully the multipath driver README file to see which kernel version is supported with that driver.
3. The first thing to do is to enable the Host Adapter Bios setting. This parameter enables the ROM Bios in the Qlogic adapter to search for SCSI devices. After this, the switch will be able to see the WWN of the Qlogic if it is already connected.
  - a. While booting the server, press Ctrl+Q to enter the Fast!UTIL configuration utility as show in Figure 20-25.

```
QLogic Corporation
QLA2312 PCI Fibre Channel ROM BIOS Version 1.38 Subsystem Vendor ID 1014
Copyright (C) QLogic Corporation 1993-2003. All rights reserved.
www.qlogic.com

Press <CTRL-Q> for Fast!UTIL
```

Figure 20-25 Press Ctrl-q to enter Fast!UTIL configuration Utility

- b. Use the cursor to select the HBA and press Enter as shown in Figure 20-26. We need to repeat these following steps on both available HBAs.



| Adapter Type | I/O Address |
|--------------|-------------|
| QLA23xx      | 5000        |
| QLA23xx      | 5100        |

Figure 20-26 Use the cursor to select the HBA

- c. In the Fast!UTIL Options panel, select Configuration Settings and press Enter.
    - d. In the Host Adapter Settings panel, select Host Adapter BIOS and press Enter to set the value to Enabled as shown in Figure 20-27 on page 611. While on this panel, you can make a note of the Adapter Port Name. It will be necessary to do the zoning configuration and do the mapping in Storage Manager.

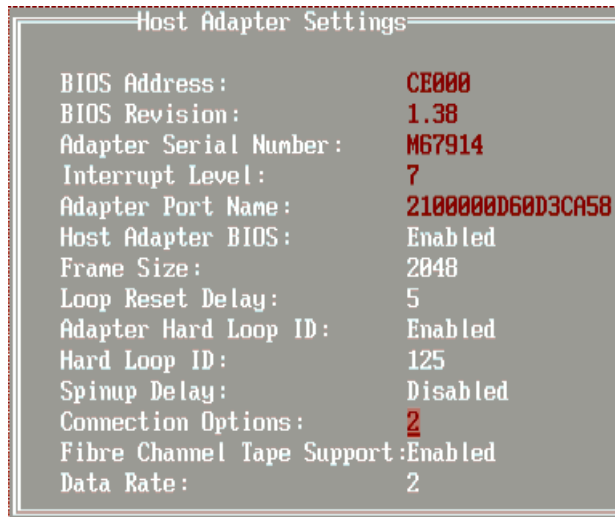


Figure 20-27 Select Enabled for Host Adapter Bios

- e. Repeat the steps from step c on page 610 for the second adapter with the I/O Address 5100.

**Note:** Boot level failover is not supported in the HBA Bios. Booting from SAN requires selecting the logical drive with the lowest LUN number. This storage subsystem is active-passive type. Because of that, the selected logical drive must also be in its preferred controller.

4. Consider to disable the planar SAS if it will not be used anymore. To do that, the procedure is the following:
  - a. Press F1 when the system is booting up as shown in Figure 20-28 to enter BIOS configuration menu.




Figure 20-28 Press F1 to enter BIOS configuration menu

- b. Select Devices and I/O Ports as shown Figure 20-29.




Figure 20-29 Select Devices and I/O Ports

- c. Disable Planar SAS as shown in Figure 20-30.




Figure 20-30 Disable Planar SAS

5. One of the most important things to consider when booting from a FC Storage Subsystem, in this case DS3500, is that when installing the OS, there must be only one path from the server to the storage. In other words, the Operating System at the time of being installed should see only one disk. The reason is that at this time the OS does not have any multipath driver installed to manage the different paths to the same disk.

**Note:** After the multipath driver is installed in the OS, the zoning configuration needs to be modified again to enable all the available paths.

To be able to do that, zoning must be properly configured. This means that there must be a zoneset including a zone only with the WWN of one of the Storage controller ports and the WWN with one of the host HBAs. There are several models of SAN switches (Brocade, Qlogic, Cisco, and so on). So we will show how to do the configuration on only one Qlogic switch. The SAN topology with the zoning configuration should be similar to that shown in Figure 20-31.

For more information about zoning the switches of other models, refer to:

- *Implementing an IBM/Brocade SAN with 8 Gbps Directors and Switches*, SG24-6116
- *Implementing an IBM/Cisco SAN*, SG24-7545




Figure 20-31 Temporary zoning configuration to boot from storage

Follow these steps to do the proper SAN zoning configuration in a Qlogic switch. This will enable us to install a SuSE Linux Enterprise server in the Storage subsystem:

- a. Log in to the Qlogic Switch and a window similar to Figure 20-32 on page 614 will be displayed. The window shows that the blade server in slot 6 and the blade server in slot 11 are connected to the internal ports of the Qlogic switch. On the left, it is possible to see the WWN that identifies each HBA host port. On the other hand, it also shows that two ports of the storage that are connected to the external ports of the switch (port 5 and port 6). We need to create a zone (BootZone) on this switch to include the WWN from the blade in slot 6 (that is our server for this sample configuration) and one port of the storage controller as indicated by the Red line in Figure 20-32 on page 614.

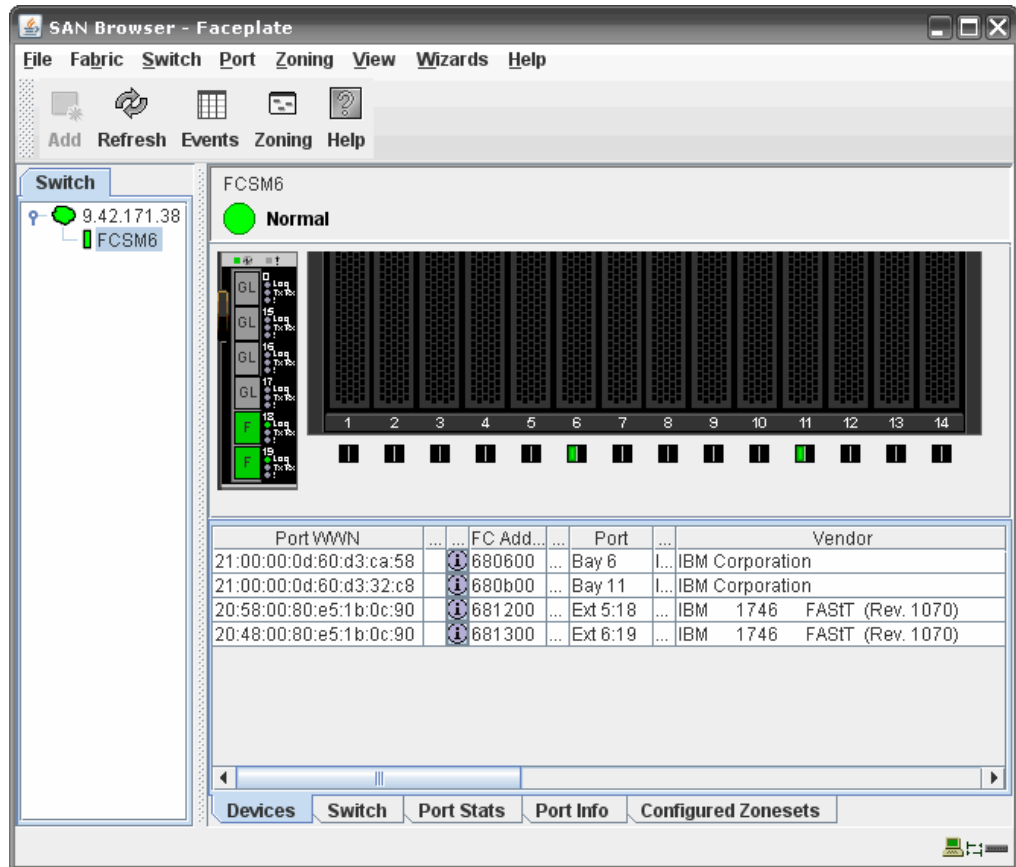


Figure 20-32 Qlogic San switch

- b. Click **Zoning** to edit the zoning configuration.
- c. Click **ZoneSet** to create a new ZoneSet called BootZone.
- d. Right-click the new ZoneSet BootZone and select **Create a zone** called Boot.
- e. Click **Alias** to create an alias called Blade6\_port1. Select the WWN from the right panel and drag and drop it over the new alias.
- f. Click **Alias** to create an alias called Storage\_port1. Select the WWN from the right panel and drag and drop it on the last created alias.
- g. Drag and drop each alias to the Boot zone and the final configuration should be similar as shown in Figure 20-33 on page 615.
- h. Click **Apply** → **Save Zoning** → **YES** to activate one of the zone sets and select **BootZone**.






Figure 20-33 Final zoning configuration

- When the zoning is correctly done, you will be able to map the logical drive using the procedure in Chapter 11, “Administration - Mappings Tab” on page 275 and continue with the SuSE Linux Server installation.

## 20.3.2 SuSE Linux Enterprise 11 installation

Using the normal procedure for the SuSE Linux Enterprise Server installation, start the installation. The window showing the available hard disks for the installation should be similar to Figure 20-34. Only one hard disk corresponding to the Storage Subsystem should be shown after the zoning configuration, and because we already disabled the local harddrive in this blade server, we only see the hard drive that is configured from the DS3500 in Figure 20-34.




Figure 20-34 Available Storage Logical Drive

After finishing the installation, install the multipath driver as shown in 20.1.2, “Installing the RDAC Multipath Driver” on page 593 and modify the zoning configuration as shown in 20.3.3, “SuSE Linux final zoning topology” on page 616.

### 20.3.3 SuSE Linux final zoning topology

The final zoning configuration for RDAC for Linux should include the following considerations:

- Multiple Fibre Channel HBAs within the same server must not be able to discover the same DS3500 controller port.
- Fibre Channel HBAs are isolated from each other.
- Each Fibre Channel HBA must be able to discover only one DS3500 controller, not both.

An example of this is shown in Figure 20-35. Follow the procedure in 20.3.1, “Preparing for the installation” on page 609 to create Zone 1 and Zone 2.




Figure 20-35 Definitive zoning configuration for linux RDAC

## 20.4 DS3500 FC storage configuration on SLES 11 using DMM

In this section we discuss how to attach a DS3500 logical drive to a SuSE Linux Enterprise Server 11 using Device Mapper Multipath (DMM) module. This multipath driver is an alternative to RDAC when it is not supported by the Linux kernel.

Like RDAC, Device Mapper Multipath allows you to configure multiple I/O paths between server nodes and storage arrays creating a new virtual device that manages the paths to provide failover in an active/passive storage subsystem.

### 20.4.1 DMM Overview

Device mapper is a block subsystem that provides layering mechanism for block devices. Device mapper can be written to provide a specific functionality on top of a block device.

DMM resolves all the issues that arise in accessing a multipathed device in Linux. It also provides a consistent user interface for storage devices provided by multiple vendors. There is only one block device (/dev/mapper/XXX) for a LUN and this is the device created by device mapper. Paths are grouped into priority groups, and one of the priority groups will be used for I/O, and is called active. A path selector selects a path in the priority group to be used for an I/O based on a load balancing algorithm (for example round-robin).

When a I/O fails in a path, that path gets disabled and the I/O is re-tried in a different path in the same priority group. If all paths in a priority group fails, a different priority group that is enabled will be selected to send I/O.

DMM consists of four components:

1. dm-multipath: This module is responsible for the multipathing decisions in event of a failure.
2. multipath: This command is used to detect multiple devices for failover, to list the available paths, and to remove virtual devices.
3. multipathd daemon: This daemon constantly checks the paths to mark them as failed when it finds a path is faulty and switches the I/O to the next enabled priority group. It also keeps checking the failed path, and, after the failed path comes alive based on the failback policy, reactivates the path.
4. multipathd command: When used with -k parameter it is possible to run a command line interface as show in Example 20-11.

*Example 20-11 multipathd command line*

---

```
multipathd -k
multipathd> help
multipath-tools v0.4.8 (08/02, 2007)
CLI commands reference:
list|show paths
list|show paths format $format
list|show status
list|show maps|multipaths
list|show maps|multipaths status
list|show maps|multipaths stats
list|show maps|multipaths topology
list|show topology
list|show map|multipath $map topology
list|show config
list|show blacklist
list|show devices
list|show wildcards
add path $path
remove|del path $path
add map|multipath $map
remove|del map|multipath $map
switch|switchgroup map|multipath $map group $group
reconfigure
suspend map|multipath $map
resume map|multipath $map
resize map|multipath $map
reload map|multipath $map
reinstate path $path
fail path $path
```

## 20.4.2 Comparing RDAC (MPP) to DMM

Although MPP is actually a driver, and performs its virtualization functions at a lower level, much of Device Mapper's virtualization occurs at a higher level using a service.

If you run **SMdevices**, the system will show every path to each device. These will usually be named something like `/dev/sdb`, `/dev/sdc`, and so on. Each of these is called a "block device", to which smash can write IO.

Device mapper takes these underlying block devices and maps all the paths leading to a single volume to a single virtualized block device. These will usually be named something like `/dev/dm-0`, `/dev/dm-1`, and so on. These are also block devices, and will be our smash targets instead of the `/dev/sd*` devices.

Because device mapper performs its virtualization above the SCSI midlayer (which acts as an intermediary layer between the lower and upper layers), the system will still see each and every path to the volume. Therefore, if a path that is running IO is lost, the host will report an "IO error". This, in itself, does not indicate a test failure, but rather that the host tried to send IO down the path and the path was not present. The SCSI midlayer will show that there was an IO error, and will report that to the device mapper multipath service, which will fail the path and transfer IO to the alternate path.

This is different from MPP where two paths are virtualized into a single device in such a way that the SCSI midlayer only sees one device. In that case, the SCSI midlayer would never report the IO error if the system was running MPP.

## 20.4.3 Planning for the installation

Follow the following guidelines for planning:

- ▶ Multiple combinations of systems, AIX versions, Host Bus Adapters (HBAs), firmware levels, and multipath drivers can be used to connect a POWER® system to a DS3500 storage subsystem. For detailed HBA and systems information regarding combinations, see the following website to access the System Storage Interoperation Center (SSIC):  
<http://www.ibm.com/systems/support/storage/config/ssic>
- ▶ Regarding zoning configuration, ensure that:
  - Multiple Fibre Channel HBAs within the same server must not be able to discover the same DS3500 controller port.
  - Fibre Channel HBAs are isolated from each other.
  - Each Fibre Channel HBA must be able to discover only one DS3500 controller.

An example of a zoning configuration for this environment is shown in Figure 20-36 on page 619.




Figure 20-36 Zoning configuration sample applicable for this environment

- In order to map the logical drive to the host and configure the zoning, it is necessary to know the WWN of each device. To be able to get it, follow the procedure in the 20.3.1, “Preparing for the installation” on page 609.

## 20.4.4 Installing the DMM multipath driver

Perform the following steps to install the DMM multipath driver.

### Setting HBA timeouts

Modify the `/etc/modprobe.conf.local` file and add the following lines to this file:

```
options lpfc_nodev_tmo 10
options qla2xxx qlport_down_retry=10
```

### Setup multipath.conf

It is important to set up this file because it is consulted by the `multipath` command for the configuration of multipath devices. It is essentially the configuration file for devices virtualized by device mapper.

The file has four sections:

- Defaults: Specifies system level default override
- Blacklist: Specifies a list of devices to not be multipathed
- Devices: Specifies the list of devices that the setting will be applied to
- Multipaths: Specifies configuration settings for tuning

Create the `/etc/multipath.conf` file as shown in Example 20-12.

*Example 20-12 Multipath.conf file sample*

```
multipaths {
 multipath {
 wwid 360080e50001b0c90000007504c6e060a
```

```

 alias mpath0
 path_grouping_policy multibus
 path_checker readsector0
 path_selector "round-robin 0"
 failback "5"
 rr_weight priorities
 no_path_retry "5"
 }
}
devices {
 device {
 vendor "IBM"
 product "1746"
 hardware_handler "1 rdac"
 path_checker rdac
 failback 0
 path_grouping_policy multibus
 prio_callout "/sbin/mpath_prio_tpc /dev/%n"
 }
}

```

---

The vendor field must be the same as in the `/sys/block/sdb/device/vendor` file and the product field must be the same as in the `/sys/block/sdb/device/model` file.

**Note:** In this case `/dev/sdb` is one of the paths to the logical drive.

Use an alias for naming the devices instead of using the uid of the device. In this case we choose to use the name `mpath`. To choose the name, consider that the devices will be created as `/dev/mapper/mpath[n]`. This name will be created under the `/dev/mapper` directory.

## Configure Services

Configure the `multipathd` service to start on boot with the following command:

```
chkconfig multipathd on
```

The `hal` daemon can sometimes cause problems by automatically mounting the underlying `/dev/sd*` paths after a failback. To work around this, the `hal` daemon needs to be disabled in SLES11 base. This issue is fixed in service pack 1.

If running SLES11 base, disable the `haldaemon` from starting on boot with the following commands:

```
chkconfig network-remotefs off
chkconfig haldaemon off
```

Create and map LUNs to the host(s) as shown in Chapter 11, “Administration - Mappings Tab” on page 275.

## 20.5 Scan and manage the storage logical drive

To scan for new devices run the following commands:

```
rescan-scsi-bus.sh
multipath
```

To list all the available paths use the following command (output shown in Example 20-13):

```
#multipath -ll
```

This command is used to list all of the “virtual devices” created by device mapper, and shows the physical paths associated with each.

*Example 20-13 Output of multipath -ll command*

---

```
multipath -ll
mpath0 (360080e50001b0c90000007504c6e060a) dm-0 IBM,1746 FAStT
[size=15G][features=1 queue_if_no_path][hwhandler=1 rdac][rw]
_ round-robin 0 [prio=4][active]
 _ 4:0:0:0 sdb 8:16 [active][ready]
 _ 5:0:0:0 sdc 8:32 [active][ghost]
 _ 5:0:1:0 sdd 8:48 [active][ghost]
 _ 4:0:1:0 sde 8:64 [active][ready]
```

---

Example 20-13 shows that the disks, in this case paths `/dev/sdb` and `/dev/sdc`, are ready. The reason is that those paths are connected to the controller that is owner of the logical drive in an active/passive array storage.

The other two paths (`/dev/sdc` and `/dev/sdd`) are in standby mode (ghost). In the event when both actives paths fail, device mapper will fail over to the active ghost paths. Afterwards, the Storage subsystem will change the logical drive ownership to the secondary Storage controller and there will be a message in Storage Recovery Guru saying that logical drive is not in the preferred path.

## Partitioning the virtual device

According to the Example 20-13, we have four paths to the logical drive and we can create our partition on one of the paths as shown in Example 20-14.

*Example 20-14 Creating a 100MB partition with fdisk*

---

```
Linux:/dev/mapper # fdisk /dev/sdb
Device contains neither a valid DOS partition table, nor Sun, SGI or OSF disklabel
Building a new DOS disklabel with disk identifier 0xebba465c.
Changes will remain in memory only, until you decide to write them.
After that, of course, the previous content won't be recoverable.
The number of cylinders for this disk is set to 15360.
There is nothing wrong with that, but this is larger than 1024,
and could in certain setups cause problems with:
 1) software that runs at boot time (e.g., old versions of LILO)
 2) booting and partitioning software from other OSs
 (e.g., DOS FDISK, OS/2 FDISK)
Warning: invalid flag 0x0000 of partition table 4 will be corrected by w(rite)

Command (m for help): n
Command action
 e extended
 p primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-15360, default 1):
Using default value 1
Last cylinder, +cylinders or +size{K,M,G} (1-15360, default 15360): +100M
```

Command (m for help): w  
The partition table has been altered!

Calling ioctl() to re-read partition table.  
Syncing disks.

---

We now need to create a map from the previous partition table. In order to do that, run the following command:

```
kpartx -a /dev/mapper/mpath0
```

List the /dev/mapper directory and the partitions on /dev/dm-0 to verify the correct partitioning as shown in Example 20-15.

*Example 20-15 Verify the partitioning*

---

```
ls /dev/mapper/
control mpath0 mpath0p1
```

```
fdisk -l /dev/dm-0
Disk /dev/dm-0: 16.1 GB, 16106127360 bytes
64 heads, 32 sectors/track, 15360 cylinders
Units = cylinders of 2048 * 512 = 1048576 bytes
Disk identifier: 0xebba465c
```

| Device      | Boot | Start | End | Blocks | Id | System |
|-------------|------|-------|-----|--------|----|--------|
| /dev/dm-0p1 |      | 1     | 101 | 103408 | 83 | Linux  |

---

The last step is to format the partition and mount it with the following commands:

```
mkfs.ext3 /dev/mapper/mpath0p1
mount /dev/mapper/mpath0p1 /mnt
```

**Important:** You must make sure you are running IO to the virtual devices created by device mapper, and not the physical device paths. If you are running IO to the physical paths, device mapper will not be able to manage a failover and the IO will fail.

## Removing Virtualized Devices

The following multipath commands will “flush” all unused device mappings.

```
multipath -F
```

After flushing the paths, use the **multipath** command with no options to add back mappings for the devices.

**Note:** It is not possible to flush the mappings while the device is being used by the system. Unmount any filesystems for /dev/dm-\* devices before running this command.

The same command with the **-f** parameter can be used to remove a single device mapping as shown in Example 20-16 on page 623. Run **multipath -ll** and find the “dm-\*” device you want to remove. Find the scsi id of the device to select the device.



```
multipath -f 360080e50001b0c90000007504c6e060a
```

---

## Troubleshooting Guide

Here we present a few common problems and ways to troubleshoot them.

### ***rescan\_scsi\_bus.sh is not showing any new device***

If the server did not recognize any new mapped LUNs after running the `rescan_scsi_bus.sh` command a few times, reboot the server.

### ***multipath -ll output does not show all the LUNs.***

This most probably is because Device Mapper is not seeing all the LUNs mapped to it. This can happen due to various reasons,

If mappings are edited with the host connected to the array, the host might fail to monitor for newly added LUNs and there will not be a scsi block device created for these newly added LUNs. To fix this issue, rescan for the scsi devices using the command `rescan_scsi_bus`. This should create scsi device for the newly added LUNs. Run `multipath` to add the scsi devices to device mapper table.

### ***multipath does not show paths equal to the number of LUNs mapped***

The first thing that you should check is the 'lsscsi' output to see if the SCSI midlayer is seeing (# of paths x # of LUNs) number of block devices.

If not run the command `rescan-scsi-bus.sh`. This should fix the issue. If not, reboot the host.

If "lsscsi" is seeing all the devices, then perform the following steps:

- ▶ Issue `multipath` so that device mapper can add the devices seen by lsscsi.
- ▶ Flush all the mappings and re-add them using the commands `multipath -F` and `multipath`.
- ▶ Restart multipath daemon using the `service multipathd restart` command

### ***multipath -ll output is showing the paths as [failed][faulty]***

This could happen if the physical path to a controller fails, or if the controller itself fails and failover happens.

- ▶ Verify the disk uid and parameters in `multipath.conf` file.
- ▶ Verify cables and zoning configuration.
- ▶ If this happens after an error in test and you want to make the configuration optimal, save the logs first, then flush and re-add multipath devices, and restart the multipath daemon. If the issue still persists, reboot the system.





## AIX 6.1 configuration guide

In this chapter, we discuss sample configurations to show how to configure AIX 6.1 to access a logical drive from a DS3500 Storage subsystem. For this sample FC configuration, we are using an JS22 blade server with dual port FC adapter and two FC switches connected to a DS3500 storage subsystem.

This chapter will include the following:

- ▶ 21.1, “Planning for the installation” on page 626
- ▶ 21.2, “Configuring MPIO” on page 632
- ▶ 21.4, “Scan and manage the storage logical drive from AIX” on page 634
- ▶ 21.5, “AIX SAN Boot with the IBM System Storage DS3500” on page 637

**Note:** At the time of writing, only FC configuration was supported on AIX.

## 21.1 Planning for the installation

Proper planning is always a crucial step in every configuration, and here we list some important considerations:

- ▶ AIX host attachment to the DS3500 storage subsystems requires the additional purchase of the AIX operating system Host Kit Option or Feature Code. The AIX Host Kit options contain the required IBM licensing to attach an IBM AIX Host System to the DS3500. Contact your IBM service representative or IBM reseller for purchasing information.
- ▶ Multiple combinations of servers, AIX versions, Host Bus Adapters (HBAs), firmware levels, and multipath drivers can be used to connect a POWER system to a DS3500 storage subsystem. For detailed HBA and systems information regarding combinations, see the following website to access the System Storage Interoperation Center (SSIC):

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

The SSIC can provide the interoperability matrix, and is a selection driven tool where you provide the criteria for your configuration.

To verify the microcode level on the HBAs, use the following commands:

- `lscfg -v1 fcs0` (The Z9 or ZA fields contains the firmware level)
- `lsmcode -d -d fcs0`
- `fcstat fcs0`
- ▶ To be able to map the logical drive to the host and to configure the zoning, get the WWN of each HBA port (fcsX). The Example 21-1 shows how to get the WWN in VIOS, and the Example 21-2 shows how to get it on AIX.

### *Example 21-1 WWN from VIOS*

---

```
$ lsdev -dev fcs0 -vpd
fcs0 U78A5.001.WIH1752-P1-C10-T1 FC Adapter

Network Address.....2100001B3208E0AB
ROS Level and ID.....49535020
Device Specific.(Z0).....0004001B
Device Specific.(Z1).....02040083
Device Specific.(Z2).....49535020
Device Specific.(Z3).....00002422
Device Specific.(Z4).....000E0002
Device Specific.(Z5).....00010009
Device Specific.(Z6).....00010002
Device Specific.(Z7).....00000000
Device Specific.(Z8).....2100001B3208E0AB
Device Specific.(Z9).....ISP
Hardware Location Code.....U78A5.001.WIH1752-P1-C10-T1
```

### PLATFORM SPECIFIC

```
Name: fibre-channel
Model: QMI2472
Node: fibre-channel@1
Device Type: fcp
Physical Location: U78A5.001.WIH1752-P1-C10-T1
```

---

#### Example 21-2 WWN from AIX

```
lscfg -vl fcs0
fcs0 U78A5.001.WIH1752-P1-C10-T1 FC Adapter

Network Address.....2100001B3208E0AB
ROS Level and ID.....49535020
Device Specific.(Z0).....0004001B
Device Specific.(Z1).....02040083
Device Specific.(Z2).....49535020
Device Specific.(Z3).....00002422
Device Specific.(Z4).....000E0002
Device Specific.(Z5).....00010009
Device Specific.(Z6).....00010002
Device Specific.(Z7).....00000000
Device Specific.(Z8).....2100001B3208E0AB
Device Specific.(Z9).....ISP
Hardware Location Code.....U78A5.001.WIH1752-P1-C10-T1
```

- Unlike System x (enable HBA BIOS), Power Systems do not require any particular configuration to be able to discovery the logical drive.

### 21.1.1 Zoning considerations

A zone is a group of fabric-connected devices arranged into a specified grouping. Zones can vary in size depending on the number of fabric-connected devices, and devices can belong to more than one zone.

**Note:** Utilizing zoning is always a good idea with SANs that include more than one host. With SANs that include more than one operating system, or SANs that contain both tape and disk devices, it is mandatory.

Without zoning, failing devices that are no longer following the defined rules of fabric behavior might attempt to interact with other devices in the fabric. This type of event is similar to an Ethernet device causing broadcast storms or collisions on the whole network, instead of being restricted to one single segment or switch port. With zoning, these failing devices cannot affect devices outside of their zone.

A general rule is to create a zone for the connection between the host bus adapter (HBA1) and controller A and a separate zone that contains the other HBA2 to controller B, and then create additional zones for access to other resources. For a single HBA server, use a separate zone to controller A and a separate zone to controller B.

### 21.1.2 SAN Boot implementation possibilities

Implementations of SAN boot with AIX include the following possibilities:

- To implement SAN boot on a system with an already installed AIX operating system, you can:
  - Use the `alt_disk_install` system utility.
  - Mirror an existing SAN Installation to several other LUNs using Volume Copy.
- To implement SAN boot on a new system, you can:
  - Start the AIX installation from a bootable AIX CD install package.

- Use the Network Installation Manager (NIM).

The methods known as `alt_disk_install` or mirroring are simpler to implement than the more complete and more sophisticated method using the Network Installation Manager (NIM).

### 21.1.3 Planning MPIO

With Multiple Path I/O (MPIO), a device can be individually detected through one or more physical connections, or paths. A path-control module (PCM) provides the path management functions.

**Note:** At the time of writing this Redbook publication, even though both MPIO and RDAC were available, RDAC is being replaced by MPIO. So we will only cover the MPIO multipath driver.

An MPIO-capable device driver can control more than one type of target device. A PCM can support one or more specific devices. Therefore, one device driver can be interfaced to multiple PCMs that control the I/O across the paths to each of the target devices. The AIX PCM has a health-check capability that can be used to do the following:

- ▶ Check the paths and determine which paths are currently usable for sending I/O.
- ▶ Enable a path that was previously marked failed because of a temporary path fault (for example, when a cable to a device was removed and then reconnected).
- ▶ Check currently unused paths that might be used if a failover occurred (for example, when the algorithm attribute value is failover, the health check can test the alternate paths).

This device driver is packaged with the AIX set of installation media and will automatically be installed with AIX if the POWER system has HBAs physically on the machine or LPAR when it was installed. If the HBAs are physically installed afterwards or dynamically associated with the LPAR then there is a possibility that all License Program Products (LPPs) and prerequisites for the MPIO (and the HBA device drivers) are not installed. In this case the AIX installation media might be required when the command `cfgmgr` is run. The command will automatically configure the HBAs to AIX and, with the installation media, install all the required HBA and MPIO LPPs. For example, to use the media on device `/dev/cd0`, insert the installation media CD 1 in the CDRom, and enter the command `cfgmgr -i /dev/cd0`.

You might be prompted to insert another CD media if needed to complete the configuration, after which you can re-apply the Technology Level (TL) package again, containing updates to the base level file sets just installed on the CD media. To do this, follow your normal procedure of applying TL maintenance to AIX.

It is possible to verify that the MPIO driver packages are correctly installed as shown in Example 21-3.

*Example 21-3 Verify that the following packages are installed*

---

```
lspp -a | grep -i mpio
devices.common.IBM.mpio.rte
6.1.4.0 COMMITTED MPIO Disk Path Control Module
```

---

#### Planning for single Virtual I/O Server

Figure 21-1 on page 629 shows the recommended scenario for a single Virtual I/O Server solution. This scenario was not supported previously for RDAC because each HBA controller could not have access to both storage controllers. The Virtual I/O Server can provide a virtual

disk for each AIX client partition by sharing a single physical disk with the clients, or assign an entire physical disk dedicated to each client partition.

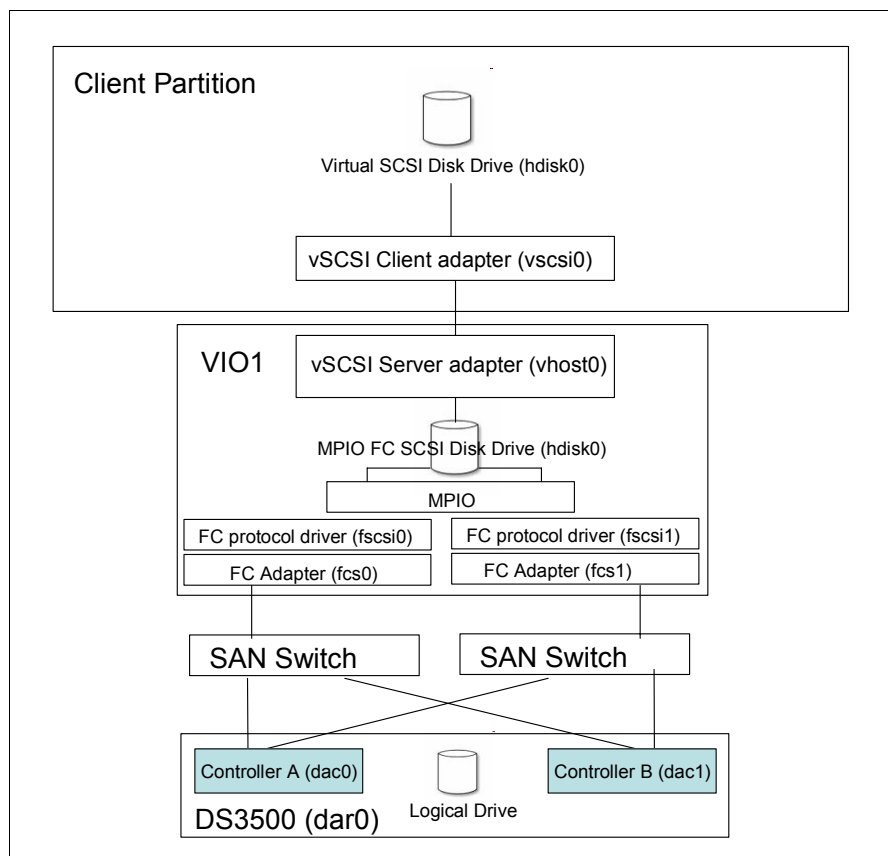


Figure 21-1 Single VIO configuration with MPIO

The Virtual I/O Server gives the ability to create multiple (boot or data) disks from a single physical disk to provide virtual disks for the logical partitions. The backing devices for the virtual disks can be logical volumes and SAN logical drive devices. These choices can be mixed and matched. The Virtual I/O Server is not limited to using just one type of backing storage device.

If using logical volumes for the backing storage device on the Virtual I/O Server, a volumegroup (or storage pool in the case of Integrated Virtualization Manager systems) is created from one or more available physical volumes. Logical volumes are created on the volume group. Those logical volumes get mapped to the client partitions as virtual disks, one logical volume for each virtual disk. When creating the logical volume, the size specified for the logical volume will become the size of the virtual disk. This gives the flexibility to effectively use a few disk drives to support multiple partitions.

Physical volumes are preferred when there is enough space, and performance becomes a priority. If using a whole physical disk for a backing device for your virtual disks, it is necessary to use a disk that the Virtual I/O Server has not assigned to a volume group. This allows client partitions to share the physical storage controller and also use full performance of a physical disk. Because of this, a failure of a single disk will only impact the single client partition using that disk.

If using SAN storage, the Virtual I/O Server can use a LUN for the storage backing device. The SAN environment can be used to provision and allocate LUNs, and the Virtual I/O Server can map the LUN to a virtual disk as though it were a physical disk. If the LUN is configured

with RAID, the disk being presented to the client partition is RAID protected by the SAN storage. If the client partitions require data protection, using a LUN configured as a RAID array can be an effective solution. It is preferred to map SAN-attached LUNs as physical volume backing storage rather than configuring volume groups and logical volumes on top of them.

If using the Integrated Virtualization Manager instead of the Hardware Management Console, the concept of logical volumes is replaced with storage pools. Like logical volumes, the storage pools can subdivide physical disks and present them to different clients as virtual disks. This technology can be used in entry-level server environments or in BladeCenter environments.

## Planning for dual Virtual I/O Servers

Figure 21-2 shows the optimal scenario for dual Virtual I/O Server solution. This configuration provides multiple paths to the data at the client partition and at the Virtual I/O Server. In addition, two SAN switches are included in the configuration to provide redundancy in the SAN.

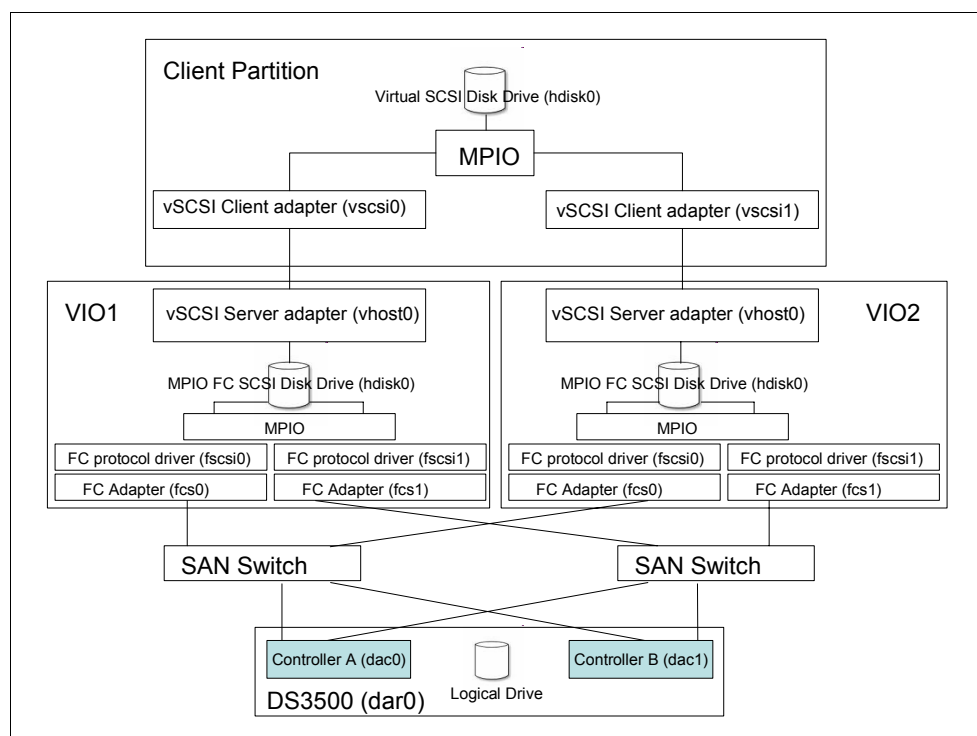


Figure 21-2 Dual VIO Configuration with MPIO

Although this configuration illustrates the virtualization of a single client disk, it can be extended to multiple disks if required. Each Virtual I/O Server uses multiple paths to the SAN LUNs through two host bus adapters, each connected to a different SAN switch. This allows the Virtual I/O Server to reach the data if one of the SAN switches fail. The SAN LUNs are treated as physical volumes and are mapped to two different virtual SCSI adapters to support two different client partitions. Each client partition has two virtual SCSI adapters. Each client virtual SCSI adapter connects to a different Virtual I/O Server. The client partition uses MPIO to provide separate paths to the data disk.

This configuration is very robust, but is more complicated than other configurations. It is necessary to set up MPIO on the client. Two Virtual I/O Servers need to be installed and configured. Two or more host bus adapters are needed in each Virtual I/O Server partition.



The Virtual I/O Server is also configured for multipath I/O. The SAN LUNs are mapped as physical volumes, not logical volumes.

This architecture provides protection from various component outages whether they are planned (such as for maintenance) or unplanned. For example, it is possible to shut down one Virtual I/O Server for a code update without interrupting the client partitions.

Even a SAN switch failure would not prevent the client partitions from getting to their data. Finally, by using RAID in the SAN LUNs, there is redundancy to a disk failure. Attach an FC Storage Subsystem to an AIX 6.1 Server.

In this scenario, because the disks are accessed through both Virtual I/O Servers concurrently, the `reserve_policy` for each disk must be set to `no_reserve` on both Virtual I/O Servers. Change the `reserve_policy` attribute to `no_reserve` using the `chdev` command as shown in Example 21-4.

*Example 21-4 Change the `reserve_policy` attribute to `no_reserve`*

```
chdev -dev hdisk2 -attr reserve_policy=no_reserve
hdisk2 changed
```

**Note:** If the `vscsi` device is in the active state, it will be necessary to add the parameter `-perm` at the end. This will update ODM and the setting will take effect on the next reboot.

Use the `lsdev` command to make sure `reserve_policy` attribute is set to `no_reserve` as shown in Example 21-5. If you used the `-perm` parameter, the `reserve_policy` will be set as `no_reserve` but a reboot will be needed to apply the change.

*Example 21-5 `reserve_policy` is set as `no_reserve`*

```
$ lsdev -dev hdisk2 -attr
```

| attribute       | value                  | description                      | user_settable |
|-----------------|------------------------|----------------------------------|---------------|
| PCM             | PCM/friend/otherapdisk | Path Control Module              | False         |
| PR_key_value    | none                   | Persistent Reserve Key Value     | True          |
| algorithm       | fail_over              | Algorithm                        | True          |
| autorecovery    | no                     | Path/Ownership Autorecovery      | True          |
| clr_q           | no                     | Device CLEARs its Queue on error | True          |
| cntl_delay_time | 0                      | Controller Delay Time            | True          |
| cntl_hcheck_int | 0                      | Controller Health Check Interval | True          |
| dist_err_pcmt   | 0                      | Distributed Error Percentage     | True          |
| dist_tw_width   | 50                     | Distributed Error Sample Time    | True          |
| hcheck_cmd      | inquiry                | Health Check Command             | True          |
| hcheck_interval | 60                     | Health Check Interval            | True          |
| hcheck_mode     | nonactive              | Health Check Mode                | True          |
| location        |                        | Location Label                   | True          |
| lun_id          | 0x0                    | Logical Unit Number ID           | False         |
| max_transfer    | 0x40000                | Maximum TRANSFER Size            | True          |
| node_name       | 0x20080080e51b0c90     | FC Node Name                     | False         |
| pvid            | none                   | Physical volume identifier       | False         |
| q_err           | yes                    | Use QERR bit                     | True          |
| q_type          | simple                 | Queuing TYPE                     | True          |
| queue_depth     | 10                     | Queue DEPTH                      | True          |
| reassign_to     | 120                    | REASSIGN time out value          | True          |
| reserve_policy  | <b>no_reserve</b>      | Reserve Policy                   | True          |
| rw_timeout      | 30                     | READ/WRITE time out value        | True          |
| scsi_id         | 0x680000               | SCSI ID                          | False         |
| start_timeout   | 60                     | START unit time out value        | True          |

## 21.2 Configuring MPIO

The following parameters need to be changed when using MPIO with physical devices:

**Important:** Because these attributes cannot be changed while the device is in active state, the **-perm** option for VIOS and **-P** for AIX are used so that the change is made in the ODM only. The changes are applied to the device when the system is rebooted.

- Change the `fc_err_recov` attribute to `fast_fail` and `dyntrk` attribute to `yes` on all the Fibre Channel adapters. It is possible to use the `lsdev -type adapter` command to find the number of that Fibre Channel adapter. Use the `chdev` command as shown in Example 21-6 and Example 21-7 to change the attributes.

*Example 21-6 Changing the attributes on VIOS*

```
$ chdev -dev fscsi0 -attr fc_err_recov=fast_fail dyntrk=yes -perm
fscsi0 changed
$ lsdev -dev fscsi0 -attr
attribute value description
user_settable
attach switch How this adapter is CONNECTED False
dyntrk yes Dynamic Tracking of FC Devices True
fc_err_recov fast_fail FC Fabric Event Error RECOVERY Policy True
scsi_id 0x660c00 Adapter SCSI ID False
sw_fc_class 3 FC Class for Fabric True
```

*Example 21-7 Changing the attributes on AIX*

```
chdev -l fscsi0 -a fc_err_recov=fast_fail -a dyntrk=yes -P
fscsi0 changed
```

It is very important to change these two parameters because of the following:

- `Dyntrk`: The default for this setting is `no`. When set to `yes`, it will enable dynamic tracking of FC devices, and the FC adapter driver detects when the Fibre Channel N\_Port ID of a device changes. The FC adapter driver then reroutes traffic destined for that device to the new address while the devices are still online. Events that can cause an N\_Port ID to change include moving a cable between a switch and storage device from one switch port to another, connecting two separate switches using an inter-switch link (ISL), and possibly rebooting a switch.
- `fc_err_recov`: This parameter controllers the fast I/O failure and is useful in situations where multipathing software is used. Setting the `fc_err_recov` attribute to `fast_fail` can decrease the I/O fail times because of link loss between the storage device and switch. This technique supports faster failover to alternate paths when FC adapter driver detects a link event, such as a lost link between a storage device and a switch. The FC adapter driver waits a short period of time, approximately 15 seconds, so that the fabric can stabilize. This occurs when the default setting of `delayed_fail` is set. At that point, if the FC adapter driver detects that the device is not on the fabric, it begins failing all I/Os at the adapter driver. Any new I/O or future retries of the failed I/Os are failed immediately by the adapter until the adapter driver detects that the device has rejoined the fabric. In single-path configurations, especially configurations with a single

path to a paging device, the `delayed_fail` default setting is preferred, one such scenario is AIX server direct attachment to the DS3500 controller host port.

- To enable load balancing across multiple Fibre Channel adapters within the Virtual I/O Servers when using the base AIX MPIO support, set the algorithm to `round_robin` for each physical disk as shown in Example 21-8 and Example 21-9.

**Note:** It is necessary to set the `reserve_policy` to `no_reserve` (as shown in Example 21-4 and Example 21-5) before setting the algorithm to `round_robin`.

---

*Example 21-8 Setting the `round_robin` parameter on VIOS*

---

```
$ chdev -dev hdisk2 -attr algorithm=round_robin -perm
hdisk2 changed
```

---

---

*Example 21-9 Setting the `round_robin` parameter on AIX*

---

```
chdev -l hdisk2 -a algorithm=round_robin -P
hdisk2 changed
```

---

**Note:** MPIO for virtual SCSI devices currently only support failover mode in AIX.

- Enable the health check mode for the disk so that the status of the disks is automatically updated. Health check mode is disabled by default (`hcheck_interval=0`). As long as it is disabled the client partition does not update the path status in case of a failure of the active path. To activate the health check function, use the `chdev` command, as shown in Example 21-10 and Example 21-11. In this example, we use a health check interval of 60 seconds. To check for the attribute setting, use the `lsattr` command.

**Note:** The path switching also works if the `hcheck_interval` attribute is disabled, but it is recommended to have the status updated automatically.

---

*Example 21-10 Changing the health interval parameter on VIOS*

---

```
$ chdev -dev hdisk2 -attr hcheck_interval=60 -perm
hdisk2 changed
```

---

---

*Example 21-11 Changing the health interval parameter on AIX*

---

```
chdev -l hdisk2 -a hcheck_interval=60 -P
hdisk2 changed
```

---

All these parameters can also be changed from Smitty `mpio` to MPIO Device Management to Change/Show MPIO Devices Characteristics.

## 21.3 Setting up the DS3500 logical drives and host mapping

Map the logical drive(s) to the AIX operating system by using the normal procedure as described in Chapter 11, “Administration - Mappings Tab” on page 275.

One of the important things to consider is using the correct Host type when mapping the logical drive to the host from the IBM Storage Manager as shown in Figure 21-3 on page 634.

When you define the host port, the host type of the attached host is defined as well. Through this information, the DS Storage System modifies the NVSRAM settings to behave according to the host type defined.

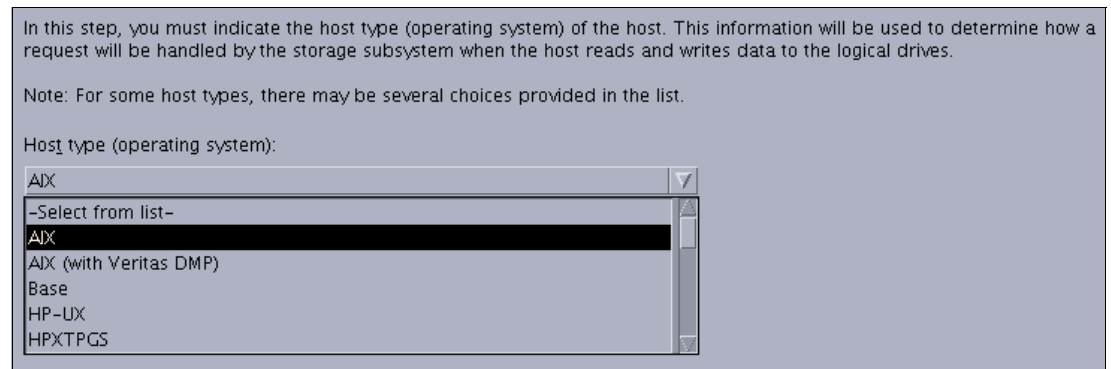


Figure 21-3 Defining the Host Type of the host

The most common of the specifications that differ according to the host type setting is the Auto Volume Transfer function (AVT), which is set as enable for certain host types, and disable for others. It is important to carefully choose the correct host type from the list of available types because this is the part of the configuration dealing with heterogeneous host support. Each operating system expects slight variations in settings, and can handle SCSI commands separately. Incorrect selections can result in failure to boot, or loss of path failover function when attached to the storage server.

For RDAC and MPIO, always use the AIX host type.

AIX (with Veritas DMP) is used when third-party multipath drivers are installed. For example to use Dynamic Multi-pathing (DMP) as the default multipath driver. DMP is part of VERITAS Volume Manager.

## 21.4 Scan and manage the storage logical drive from AIX

In this section it is assumed that the logical drive is already mapped to the system. For a detailed procedure, see 9.4.1, "Create logical drive" on page 233.

In order to discover new devices on AIX, the **cfgmgr** command must be used. This command configures devices and optionally installs devices software into the system.

In order to discover new devices on Virtual I/O Server, the **cfgdev** command must be used. This command does the same function as **cfgmgr** and is only applicable to VIOS.

Example 21-12 shows the disk that we configured for the DS3500 Storage Subsystem. The third column shows the location code. In this example is 03-09-01. Each AIX system has its own set of location codes that describe the internal path of that device, including bus and host adapter locations. See the service manual for your system type to identify device locations.

Example 21-12 Disk status and location codes

```
lsdev |grep -i mpio
hdisk2 Available 03-09-01 MPIO Other FC SCSI Disk Drive
```

The **mpio\_get\_config** command must be used to see all the available paths to the logical drive as shown in Example 21-13. This command displays the following information about the hdisks that are associated with the subsystem:

- ▶ hdisk name
- ▶ LUN number
- ▶ Current ownership
- ▶ Preferred path
- ▶ User-assigned label for the volume

**Note:** This command is not available in the Virtual I/O Server IOS. In order to use the command, run the **oem\_setup\_env** command first to start a root AIX shell. Afterwards, the **mpio\_get\_config -Av** command will be available to use it.

**Important:** To protect the configuration database, do not interrupt the **mpio\_get\_config** command. Stopping the command before it is complete can result in a corrupted database.

*Example 21-13 mpio\_get\_config output*

---

```
mpio_get_config -Av
Frame id 0:
 Storage Subsystem worldwide name: 608e5001bc9000004c5b683c
 Controller count: 2
 Partition count: 1
 Partition 0:
 Storage Subsystem Name = 'DS3524-FC'
 hdisk LUN # Ownership User Label
 hdisk2 0 A (preferred) aix1fc
```

---

The other command to list the available paths is **lspath**. This command displays one of three types of information about paths to an MPIO capable device. It either displays the operational status for one or more paths to a single device, or it displays one or more attributes for a single path to a single MPIO capable device.

By default, this command will display the information in columnar form. When no flags are specified that qualify the paths to display, the three columns are status, device and parent as shown in Example 21-14. This example shows the output for a dual controller DS3500 storage subsystem with a dual controller host configuration with one LUN attached. It also shows that the paths to the SAS1 hard disk is in the missing state.

*Example 21-14 lspath command output*

---

```
lspath
Enabled hdisk0 sas0
Enabled hdisk1 sas0
Enabled hdisk2 fscsi1
Enabled hdisk2 fscsi1
Enabled hdisk2 fscsi0
Missing ses1 sas1
Enabled hdisk2 fscsi0
Missing ses2 sas1
```

---

Possible values that can appear for the status column are:

- ▶ **enabled**: Indicates that the path is configured and operational. It will be considered when paths are selected for IO.
- ▶ **disabled**: Indicates that the path is configured, but not currently operational. It has been manually disabled and will not be considered when paths are selected for IO.
- ▶ **failed**: Indicates that the path is configured, but it has had IO failures that have rendered it unusable. It will not be considered when paths are selected for IO.
- ▶ **defined**: Indicates that the path has not been configured into the device driver.
- ▶ **missing**: Indicates that the path was defined in a previous boot, but it was not detected in the most recent boot of the system.
- ▶ **detected**: Indicates that the path was detected in the most recent boot of the system, but for some reason it was not configured. A path should only have this status during boot, and so this status should never appear as a result of the **lspath** command.

## 21.4.1 Ways to manage the paths

There are three commands available to manage the paths:

1. **chpath**: The **chpath** command either changes the operational status of paths to the specified device or it changes one, or more, attributes associated with a specific path to the specified device. The required syntax is slightly different depending upon the change being made.
2. **mkpath**: The **mkpath** command defines, and possibly configures, one or more paths to the target device. The paths are identified by a combination of the **-l Name**, **-p Parent**, and **-w Connection** flags. Both the target device and parent must be previously defined in the system to define a path. They both must be AVAILABLE to configure a path.
3. **rmpath**: The **rmpath** command unconfigures, and possibly undefines, one or more paths associated with the specified target device. The set of paths that are removed are determined by the combination of the **-l Name**, **-p Parent** and **-w Connection** flags. If the command will result in all paths associated with the device being unconfigured or undefined, the command will exit with an error and without unconfiguring or undefining any path. In this situation, **rmdev** command must be used instead to unconfigure or undefine the target device itself. The default action unconfigures each specified path, but does not completely remove it from the system. If the **-d** flag is specified, the **rmpath** command unconfigures (if necessary) and removes, or deletes, the path definition(s) from the system.

All this commands including **lspath** are combined in the following menu:

```
smitty mpio
```

In fact, with this menu is also possible to set the priority for each device's path with the following sequence:

```
smitty mpio → Path Management → Change/Show Path Characteristics → Change/Show Characteristics for Device's Path.
```

Type a valid MPIO device name or select it from the list and afterwards, type or select a valid MPIO path. A similar windows will display as shown in Example 21-15.

*Example 21-15 Change the priority for a specific path*

---

AIX MPIO Active/Passive Disk PCM Change Path Characteristics

Type or select values in entry fields.  
 Press Enter AFTER making all desired changes.

|             |                    |
|-------------|--------------------|
|             | [Entry Fields]     |
| Device Name | hdisk2             |
| Path Id     | 0                  |
| Parent Name | fscsi1             |
| Connection  | 20490080e51b0c90,0 |
| State       | Enabled            |
| Priority    | [1]                |

|          |            |           |          |
|----------|------------|-----------|----------|
| F1=Help  | F2=Refresh | F3=Cancel | F4=List  |
| F5=Reset | F6=Command | F7=Edit   | F8=Image |
| F9=Shell | F10=Exit   | Enter=Do  |          |

---

## 21.5 AIX SAN Boot with the IBM System Storage DS3500

This section contains a step-by-step illustration of SAN boot implementation using AIX 6.1 on a POWER System environment. We will discuss the procedure for all the possible implementations mentioned in 21.1.2, “SAN Boot implementation possibilities” on page 627.

**Note:** At the time of writing, Fibre Channel is the only supported environment for POWER Systems and the IBM Systems Storage DS3500. Because of that we will only include the SAN boot configuration.

When installing the operating system, you can only have one path to the storage device (LUN). Because the servers are generally equipped with two HBAs, and because most DS3500 storage subsystems have two RAID controllers, you must isolate (disconnect) the second HBA or do an appropriate zoning.

### 21.5.1 Creating a boot disk with alt\_disk\_install

The following procedure is based on the `alt_disk_install` command, cloning the current operating system to the DS3500 storage subsystem SAN disk. With this procedure, it is possible to copy the operating system from an internal hard disk to a logical drive in the storage or copy the operating system from another server to the logical drive and then remap it to the new server.

Perform the follow steps:

1. Create a logical drive on the DS3500 storage subsystem having the same or greater capacity than the source disk.
2. Map logical drive from the DS3500 to the AIX server
3. Type the following command from a terminal window on the AIX server:  

```
smitty alt_clone
```

 Afterwards, type the target disk in the Target Disk(s) to install field or select it from the list. In this case the target disk is `hdisk1`.
4. Ensure the proper boot logical setting for the target disk:

```
/usr/lpp/bosinst/blvset -d /dev/hdisk1 -g level
```

5. Update the cloned boot logical volume manually with the following command:

```
echo y | /usr/lpp/bosinst/blvset -d /dev/hdisk3 -plevel
```

## 21.5.2 AIX SAN installation with NIM

Network Installation Manager (NIM) is a client server infrastructure and service that allows remote installation of the operating system, manages software updates, and can be configured to install and update third-party applications. Although both the NIM server and client file sets are part of the operating system, a separate NIM server has to be configured that will keep the configuration data and the installable product file sets.

**Note:** For a complete information regarding NIM configurations, refer to *NIM from A to Z in AIX 5L*, SG24-7296.

### NIM preparations

The following preparations must have been completed:

- ▶ The NIM environment is deployed and all of the necessary configuration on the NIM master is already done.
- ▶ The NIM server is properly configured as the NIM master, and the basic NIM resources have been defined.
- ▶ The Fibre Channel Adapters are already installed in the machine onto which AIX is to be installed.
- ▶ The Fibre Channel Adapters are connected to a SAN, and, on the DS3500 storage subsystem, have the designated logical drive defined and mapped to the host or NIM client.
- ▶ The target machine (NIM client) currently has no operating system installed and is already define in the NIM server.

The procedure consists of setting up a remote IPL in the client to use the bootpd (boot protocol) and tftp (trivial file transfer protocol) services provided by the NIM server. With this procedure, the NIM server will be able to deliver the OS image to the client using the network.

Perform the following procedure:

1. Select the SMS Menu to enter the configuration menu as shown in Example 21-16.

*Example 21-16 Initial Power System window*

```
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
```



```

IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM
IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM IBM

```

```

1 = SMS Menu
8 = Open Firmware Prompt
5 = Default Boot List
6 = Stored Boot List

```

2. Select Setup Remote IPL (Initial Program Load) as shown in Example 21-17.

*Example 21-17 Select "Setup Remote IPL"*

Main Menu

1. Select Language
2. **Setup Remote IPL (Initial Program Load)**
3. Change SCSI Settings
4. Select Console
5. Select Boot Options

3. Select the host Ethernet adapter to boot from. In this case it is the second one as shown in Example 21-18.

*Example 21-18 Select the host Ethernet adapter to boot from*

| NIC Adapters | Device                                | Location Code                  | Hardware Address    |
|--------------|---------------------------------------|--------------------------------|---------------------|
| 1.           | Interpartition Logical LAN            | U7998.60X.100EB0A-V3-C4-T1     | 22f206db9504        |
| 2.           | <b>PORT - 5 IBM Host Ethernet Ada</b> | <b>U78A5.001.WIH1752-P1-T6</b> | <b>001a648449fe</b> |

4. Select the internet protocol version. In this case it is IPV4 as shown in Example 21-19.

*Example 21-19 Select the internet protocol version*

Select Internet Protocol Version.

1. **IPv4 - Address Format 123.231.111.222**
2. IPv6 - Address Format 1234:5678:90ab:cdef:1234:5678:90ab:cdef

5. Select the BOOTP protocol as shown in Example 21-20.

*Example 21-20 Select the BOOTP protocol*

Select Network Service.

PORT - 5 IBM Host Ethernet Adapter: U78A5.001.WIH1752-P1-T6

1. **BOOTP**
2. ISCSI

6. Select the IP Parameters option as shown in Example 21-21.

*Example 21-21 Select the IP Parameters menu option*

Network Parameters

PORT - 5 IBM Host Ethernet Adapter: U78A5.001.WIH1752-P1-T6

1. **IP Parameters**
2. Adapter Configuration
3. Ping Test
4. Advanced Setup: BOOTP

- 
7. Configure the IP Parameters as shown in Example 21-22. In the Example 21-21 on page 639 menu, it is also possible to verify the connectivity by pinging the server. Do this to make sure the connection is properly established and then return to the main menu.

*Example 21-22 Enter the necessary data to boot*

---

IP Parameters

PORT - 5 IBM Host Ethernet Adapter: U78A5.001.WIH1752-P1-T6

1. Client IP Address [9.42.171.30]
2. Server IP Address [9.42.171.26]
3. Gateway IP Address [000.000.000.000]
4. Subnet Mask [255.255.254.0]

- 
8. From the main menu, select the boot options. After that, select Select Install/Boot Device as shown in Example 21-23.

*Example 21-23 Select the Install/Boot Device menu option*

---

Multiboot

1. **Select Install/Boot Device**
2. Configure Boot Device Order
3. Multiboot Startup <OFF>

- 
9. Select Network as the device type as shown in Example 21-24.

*Example 21-24 Select the Network menu option*

---

Select Device Type

1. Diskette
2. Tape
3. CD/DVD
4. IDE
5. Hard Drive
6. **Network**
7. List all Devices

- 
10. Select the BOOTP protocol as shown in Example 21-25.

*Example 21-25 Select the BOOTP menu option*

---

Select Network Service.

PORT - 5 IBM Host Ethernet Adapter: U78A5.001.WIH1752-P1-T6

1. **BOOTP**

## 2. ISCSI

11. Select the Host Ethernet adapter port to boot from as shown in Example 21-26.

*Example 21-26 Select the host Ethernet adapter to boot from*

Select Device

| Device Number | Current Position | Device Name                                                                         |
|---------------|------------------|-------------------------------------------------------------------------------------|
| 1.            | 4                | Interpartition Logical LAN<br>( loc=U7998.60X.100EB0A-V3-C4-T1 )                    |
| 2.            | -                | <b>PORT - 5 IBM Host Ethernet Adapter</b><br><b>( loc=U78A5.001.WIH1752-P1-T6 )</b> |

12. Select Normal Mode Boot to start the installation as shown in Example 21-27. Select Yes in the following menu to exit System Management Services.

*Example 21-27 Select the Normal Mode menu option*

Select Task

PORT - 5 IBM Host Ethernet Adapter  
( loc=U78A5.001.WIH1752-P1-T6 )

1. Information
2. **Normal Mode Boot**
3. Service Mode Boot

13. A message similar to the one shown in Example 21-28 will appear. This means that the TFTP transfer has started.

*Example 21-28 TFTP transfer*

```
TFTP BOOT -----
Server IP.....9.42.171.26
Client IP.....9.42.171.30
Subnet Mask.....255.255.254.0
(1) Filename...../tftpboot/aix1
TFTP Retries.....5
Block Size.....512
```

```
FILE : /tftpboot/aix1
BLKSIZE : 512
TFTP-RETRIES: 5
PACKET COUNT = 9000
```

14. After the setting these parameters, continue with the actual AIX installation as documented in 21.5.4, “AIX Operating System Installation” on page 642.

### 21.5.3 AIX SAN installation with CD-ROM

To install AIX from a CD-ROM, perform the following steps:

1. Insert disk 1 of the AIX installation CD set (this has a bootable image) into the CD-ROM.
2. For a new AIX system, select CD-ROM as the Install Device to make the system boot from the CD using the System Management Services (SMS) menu.
3. Let the system boot from the AIX CD image after you have left the SMS menu.
4. Continue with the actual AIX installation as documented in 21.5.4, "AIX Operating System Installation".

### 21.5.4 AIX Operating System Installation

After following the appropriate procedures either in 21.5.2, "AIX SAN installation with NIM" on page 638 or 21.5.3, "AIX SAN installation with CD-ROM" on page 642, continue with the AIX operating system installation using these steps:

1. Press "1" and Enter to use the current terminal as the system console as shown in Example 21-29.

---

*Example 21-29 Press "1" and Enter to select the system console*

---

\*\*\*\*\* Please define the System Console. \*\*\*\*\*

Type a 1 and press Enter to use this terminal as the system console.

Pour definir ce terminal comme console systeme, appuyez sur 1 puis sur Entree.

Taste 1 und anschliessend die Eingabetaste druecken, um diese Datenstation als Systemkonsole zu verwenden.

Premere il tasto 1 ed Invio per usare questo terminal come console.

Escriba 1 y pulse Intro para utilizar esta terminal como consola del sistema.

Escreveu 1 e pulse Intro para utilizar este terminal como consola do sistema.

Digite um 1 e pressione Enter para utilizar este terminal como console do sistema.

---

2. Press "1" and Enter to select the English language during the installation as shown in Example 21-30.

---

*Example 21-30 Press "1" to select the English language during the installation*

---

>>> 1 Type 1 and press Enter to have English during install.

---

3. Press "2" to change the installation settings as shown in Example 21-31.

---

*Example 21-31 Press "2" to change the installation settings*

---

Welcome to Base Operating System  
Installation and Maintenance

Type the number of your choice and press Enter. Choice is indicated by >>>.

```
>>> 1 Start Install Now with Default Settings
 2 Change/Show Installation Settings and Install
 3 Start Maintenance Mode for System Recovery
 4 Configure Network Disks (iSCSI)
 5 Select Storage Adapters

 88 Help ?
 99 Previous Menu
```

---

4. Press “1” to select System Settings as shown in Example 21-32.

*Example 21-32 Press “1” to select system settings*

---

#### Installation and Settings

Either type 0 and press Enter to install with current settings, or type the number of the setting you want to change and press Enter.

```
1 System Settings:
 Method of Installation.....New and Complete Overwrite
 Disk Where You Want to Install....hdisk0

2 Primary Language Environment Settings (AFTER Install):
 Cultural Convention.....C (POSIX)
 Language.....C (POSIX)
 Keyboard.....C (POSIX)

3 Security Model.....Default
4 More Options (Software install options)
```

```
>>> 0 Install with the settings listed above.
```

---

5. Select the installation method. In this case we want a new installation. For that reason, we selected New and Complete Overwrite as shown in Example 21-33.

*Example 21-33 Select the installation method*

---

#### Change Method of Installation

Type the number of the installation method and press Enter.

```
>>> 1 New and Complete Overwrite
 Overwrites EVERYTHING on the disk selected for installation.
 Warning: Only use this method if the disk is totally empty or if there
 is nothing on the disk you want to preserve.

2 Preservation Install
 Preserves SOME of the existing data on the disk selected for
 installation. Warning: This method overwrites the usr (/usr),
 variable (/var), temporary (/tmp), and root (/) file systems. Other
 product (applications) files and configuration data will be destroyed.
```

---

6. Verify the disk in which the AIX operating system is going to be installed. Press “0” to accept the current disk as shown in Example 21-34.

---

*Example 21-34 Verify the disk*

---

Change Disk(s) Where You Want to Install

Type one or more numbers for the disk(s) to be used for installation and press Enter. To cancel a choice, type the corresponding number and Press Enter. At least one bootable disk must be selected. The current choice is indicated by >>>.

|       | Name                                                  | Location Code | Size(MB) | VG Status | Bootable |
|-------|-------------------------------------------------------|---------------|----------|-----------|----------|
| >>> 1 | hdisk0                                                | none          | 11264    | rootvg    | Yes No   |
| >>> 0 | Continue with choices indicated above                 |               |          |           |          |
| 55    | More Disk Options                                     |               |          |           |          |
| 66    | Disks not known to Base Operating System Installation |               |          |           |          |
| 77    | Display Alternative Disk Attributes                   |               |          |           |          |
| 88    | Help ?                                                |               |          |           |          |
| 99    | Previous Menu                                         |               |          |           |          |

---

7. The next menu is the same as in step 3. Press “1” to confirm the settings shown in the menu as shown in Example 21-35.

---

*Example 21-35 Confirm the installation*

---

Overwrite Installation Summary

Disks: hdisk0  
 Cultural Convention: C  
 Language: C  
 Keyboard: C  
 JFS2 File Systems Created: yes  
 Graphics Software: no  
 System Management Client Software: yes  
 Enable System Backups to install any system: yes

Optional Software being installed:

>>> 1 Continue with Install

---

8. A window similar to the one shown in Example 21-36 will appear. The file copy has begun. When the process completes, the system will be restarted.

---

*Example 21-36 AIX Installation*

---

Installing Base Operating System

Please wait...


| Approximate<br>% tasks complete | Elapsed time<br>(in minutes) |
|---------------------------------|------------------------------|
|---------------------------------|------------------------------|

---

9. Change the zoning configuration to make all the paths available and run **cfgmgr** to discover and enable all the paths from the AIX perspective.
10. Follow the procedures in 21.4, “Scan and manage the storage logical drive from AIX” on page 634.







## VMware ESX Server and DS3500 Storage Configuration

The IBM DS3500 Storage System is designed to deliver reliable performance for mixed applications including transaction and sequential workloads. These workloads include applications that are typical of a virtual infrastructure including email, database, web server, file server, data warehouse, and backup profiles.

The IBM DS3500 Storage System offers shared storage that enable all of VMware's advanced functionality such as VMware Distributed Resource Scheduler (DRS), VMware vCenter Site Recovery Manager (SRM), and VMware High Availability (HA).

In this chapter, we outline the process for installing the VMware ESX Server and configuring settings. All of the examples are based on Fibre Channel. Detailed information about iSCSI can be found here:

[http://www.vmware.com/pdf/vsphere4/r41/vsp\\_41\\_iscsi\\_san\\_cfg.pdf](http://www.vmware.com/pdf/vsphere4/r41/vsp_41_iscsi_san_cfg.pdf)

Because we only focus on the actual implementation in this chapter, refer to *VMware Implementation with IBM System Storage DS4000/DS5000*, REDP-4609 for detailed information.

## 22.1 Introduction to IBM VMware Storage Solutions

Many companies employ VMware virtualization solutions to reduce IT costs while increasing the efficiency, utilization, and flexibility of their hardware. In fact, 190,000 customers have deployed VMware, including 96% of Fortune 1000 businesses. However, maximizing the operational benefits from virtualization requires network storage that helps optimize the VMware infrastructure.

The IBM Storage solutions for VMware offer customers:

- ▶ Flexibility: Support for iSCSI and Fibre Channel shared storage, plus HBA and storage port multi-pathing and boot from SAN.
- ▶ Performance: Outstanding high-performance block-level storage that scales with VMware's VMFS file system; independently verified high performance by the SPC-1, SPC-2, and SPCvirt\_sc2010 (Storage Performance Council) benchmarks; and balanced performance delivered by the IBM DS3500 Storage System for mixed applications running in a virtual infrastructure.
- ▶ Horizontal scalability: From entry-level through midrange to enterprise class network storage with commonality of platform and storage management.
- ▶ Hot Backup and Quick recovery: Non-disruptive backup solutions using Tivoli® and NetBackup with and without VCB (VMware Consolidated Backup). Quick recovery at the file or virtual machine level.
- ▶ Disaster recovery: DS3500 Enhanced Remote Mirror offers affordable disaster recovery with automatic failover in conjunction with VMware vCenter Site Recovery Manager (SRM).
- ▶ Affordability: Low TCO shared storage with included IBM Storage System DS Storage Manager Software (Storage Manager) and no separate software maintenance fees, cost-effective tiered storage within the same storage system, leveraging SAS drives for high performance, and SAS Nearline drives for economical capacity.
- ▶ Efficiency: Data Services features such as FlashCopy and VolumeCopy enable VMware Centralized Backup to disk and eliminate backup windows, and provide required network storage for VMware ESX Server features such as VMware vMotion, VMware Storage vMotion, Resource Pools, VMware Distributed Resource Scheduler (DRS), and VMware High Availability. Storage I/O Control gives you the chance to adjust your Storage I/O to match the mixed Workload in your environment.

VMware vSphere includes components and operations that are essential for managing virtual machines. The following components form part of the new VMware vSphere suite:

- ▶ VMware ESX/ESXi server
- ▶ VMware vCenter Server
- ▶ Storage I/O Control (SIOC)
- ▶ Datastore
- ▶ Host Agent

### 22.1.1 VMware installation prerequisites

- ▶ Refer to the following VMware website for a complete up to date list of the prerequisites for installing VMware ESX server:  
[http://www.vmware.com.tt/pdf/vsphere4/r41/vsp\\_41\\_esx\\_get\\_start.pdf](http://www.vmware.com.tt/pdf/vsphere4/r41/vsp_41_esx_get_start.pdf)
- ▶ Refer to the following VMware website for more information about the maximum hardware capabilities of the VMware ESX Server:  
[http://www.vmware.com/pdf/vsphere4/r41/vsp\\_41\\_config\\_max.pdf](http://www.vmware.com/pdf/vsphere4/r41/vsp_41_config_max.pdf)

## 22.2 SAN Zoning configuration

As recommended by VMware, we set up four paths between the server and the storage system as shown in Figure 22-1. At least two VMware ESX Server HBA ports must be used, and both HBA ports must see both controllers.




Figure 22-1 Paths Between the VMware ESX Server and the DS3500 Storage System

Refer to the following materials for more information about zoning:

- ▶ *Implementing an IBM/Brocade SAN with 8 Gbps Directors and Switches*, SG24-6116
- ▶ *Implementing an IBM/Cisco SAN*, SG24-7545

## 22.3 DS3500 Storage configuration

As a first step, you need to configure your storage on the IBM DS3500 Storage Subsystem. Complete the following steps as part of the storage configurations:

1. Zoning: Zone your VMware ESX Server to your IBM DS3500 Storage Subsystem. Ensure that your VMware environment has sufficient paths and connections for redundancy and high availability. See Figure 22-1.
2. Create a LUN: Create a LUN of size that fits your VMware partition requirements.
3. Storage partition: From the IBM DS Storage Manager mapping window, create a VMware host and define the host ports of the Fibre Channel HBAs as shown in Figure 22-2 on page 650.

**Note:** A VMware ESX Server-specific host type is now available for DS3500 Storage Subsystems. If you are using the default host group, ensure that the default host type is VMWARE.

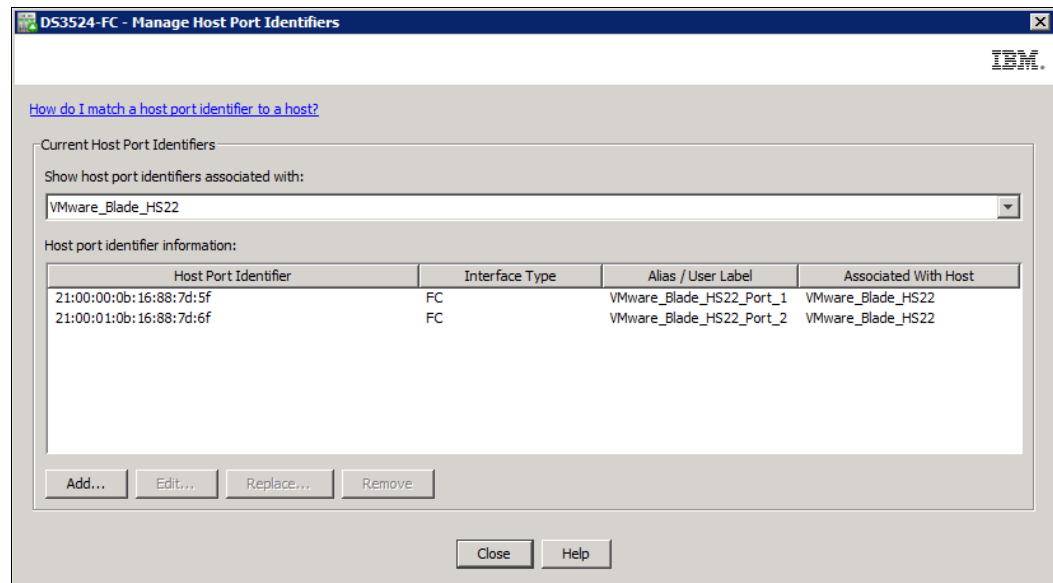


Figure 22-2 Storage partitioning

4. LUN mapping: Map the LUN that you created in number 2 to the host partition you created in the preceding step. Figure 22-3 shows an example of a valid LUN mapping for installation purposes.




Figure 22-3 LUN Mapping for VMware installation

**Note:** The IBM System Storage DS Storage Manager cannot be installed on the VMware ESX Server. Instead, it must be installed on a Linux or Windows management workstation. However, it can be the same management workstation that you used for the browser-based VMware Management Interface.

### 22.3.1 Mapping LUNs to a storage partition

In this section, we provide notes about LUN mapping that are specific to VMware ESX Servers.

When you map your LUNs on VMware ESX Server, consider the following guidelines:

- ▶ Map the LUNs using consecutive numbers, starting with LUN 0. For example, map LUNs to numbers 0, 1, 2, 3, 4, 5, and so on, without skipping any numbers.
- ▶ On each partition, you must map a LUN 0.
- ▶ If your configuration does not require LUN sharing (single or multiple independent VMware ESX Servers, local virtual cluster), each logical drive must be mapped either directly to a host, or to a host group with a single host as a member.
- ▶ LUN sharing across multiple VMware ESX Servers is only supported when you are configuring VMware vMotion enabled hosts or Microsoft Cluster nodes. On LUNs that are mapped to multiple VMware ESX Servers, you must change the access mode to Shared.

You can map the LUNs to a host group for the VMware ESX Servers, so they are available to all members of the host group. For additional information about Windows Clustering with VMware ESX Server, see the VMware ESX Server Installation Guide at the following URL:

<http://www.vmware.com/support/pubs/>

## 22.3.2 Steps for verifying the storage configuration for VMware

Perform the following steps to verify that your storage setup is fundamentally correct and that you can see the IBM Midrange Storage Subsystem on your VMware ESX Server:

1. Boot the server.
2. On initialization of the QLogic BIOS, press Ctrl+Q to enter the Fast!UTIL setup program.
3. Select the first host bus adapter that is displayed in the Fast!UTIL panel.
4. Select **Host Adapter Settings**, and press Enter.
5. Select **Scan Fibre Devices**, and press Enter. Figure 22-4 shows the resulting output.




Figure 22-4 Scanning for Fiber Devices

**Note:** Depending on how the configuration is cabled, you might see multiple instances.

If you do not see a DS3500 controller, verify the cabling, switch zoning, and LUN mapping.

## 22.4 Installing the VMware ESX Server

In this section, we provide the procedure and details that are related to installing the VMware ESX server.

You need a VMWare ESX Server CD. In addition, you will need the following information in Table 22-1 before you begin. The Licensing model is different in VI3 and vSphere implementation.

Table 22-1 VMware ESX Server information

|                                    |                     |
|------------------------------------|---------------------|
| Server Name (FQDN)                 | _____.<domain>.com  |
| IP Address                         | ____.____.____.____ |
| Subnet Mask                        | ____.____.____.____ |
| Default Gateway                    | ____.____.____.____ |
| DNS Server IP Addresses            | ____.____.____.____ |
| Location of vSphere Server License |                     |

**Note:** If this is a re-build of the host, remove it from the vCenter before proceeding.

### 22.4.1 Configuring the hardware

Power off the server hardware if it is powered on. Perform the following steps:

1. If needed, install additional network adapters.
2. If needed, install Fibre Channel HBA card(s).
3. After the chassis is closed and the machine is re-racked, plug in all associated cables EXCEPT SAN Fibre Channel cables. (this is only valid, when you do not plan to install ESX on a SAN disk. In our example we have the SAN Fibre Channel cables connected to the Switch)
4. Configure the BIOS and RAID, as described in the following section.

#### Configuring the server BIOS

1. Check all firmware and update as necessary (BIOS, HBA, Internal RAID).
2. Ensure that your server BIOS is set up to accommodate virtualization technology. Refer to your server vendor specific documentation for guidelines.

#### Configuring the server HBA

To enable the HBA BIOS:

**Note:** In this example, we used a QLogic QMI2572 card.

1. Press Ctrt-Q or Alt-Q when prompted during the boot process to configure the QLogic BIOS as shown in Figure 22-5 on page 653.

```

Searching for devices at HBA 0...

SLOT ID LUN VENDOR PRODUCT REVISION INT13 SIZE \ NV

 0 1 0 IBM-ESXS ST9500430SS BD11 Boot 465 GB
 0 0 0 LSILogic 385-064E-JR 1.30.03.00 NV 2D:19

LSI Corporation MPT host ROM successfully installed!

QLogic Corporation
QMI2572 PCI3.0 Fibre Channel ROM BIOS Version 2.08
Copyright (C) QLogic Corporation 1993-2009. All rights reserved.
www.qlogic.com

Press <CTRL-Q> or <ALT-Q> for Fast!UTIL

BIOS for Adapter 0 is disabled

BIOS for Adapter 1 is disabled
ROM BIOS NOT INSTALLED

<Alt-Q> Detected, Initialization in progress, Please wait...

```

Figure 22-5 Press CTRL + Q or ALT + Q to get Qlogic Fast!UTIL Menu

2. Select the first QLogic card entry. Press Enter as shown in Figure 22-6.

```

QLogic Fast!UTIL

Select Host Adapter
Adapter Type Address Slot Bus Device Function

QMI2572 1200 01 24 00 0
QMI2572 1000 01 24 00 1

Use <Arrow keys> to move cursor, <Enter> to select option, <Esc> to backup

```

Figure 22-6 Selecting HBA adapter

3. Select **Configure Settings** for the selected HBA card and press Enter as shown in Figure 22-7 on page 654.

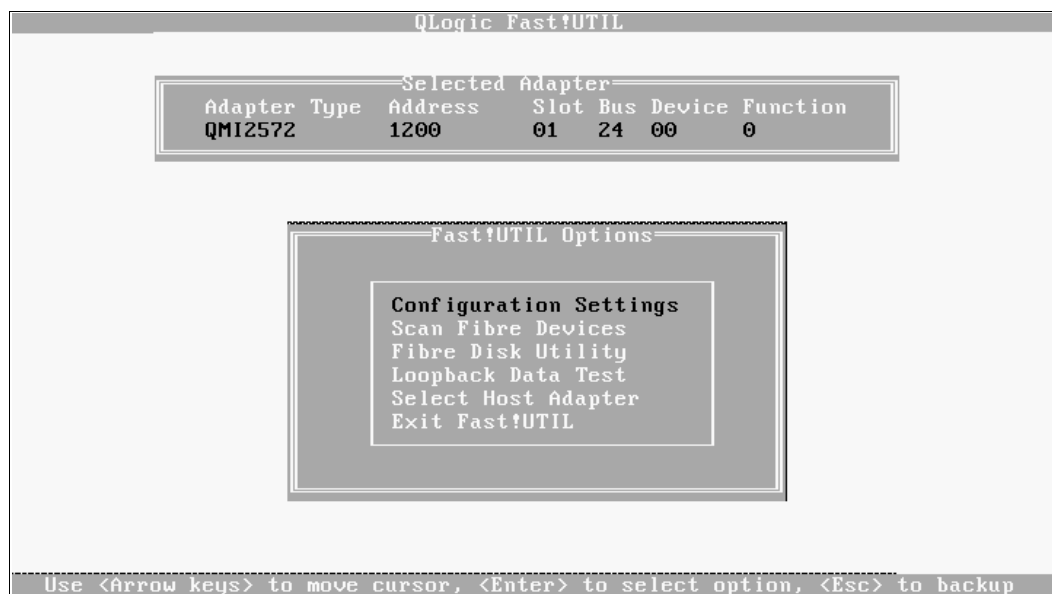


Figure 22-7 HBA Configure Settings

4. Select **Adapter Settings** and press Enter as shown in Figure 22-8.




Figure 22-8 HBA Adapter settings

5. Highlight the Host Adapter BIOS entry and ensure that it is set to Enabled. If not, press Enter to set it to Enabled as shown in Figure 22-9 on page 655.





Figure 22-9 HBA Adapter BIOS Enabled

- As seen in Figure 22-10, press the Esc key twice to exit, and select **Save Changes** when prompted. Press Enter.



Figure 22-10 Saving Changes

- If additional HBAs are present, highlight **Select Host Adapter** as shown in Figure 22-11 on page 656, press Enter, and repeat steps 2 through 6 for each additional adapter.



Figure 22-11 Selecting additional HBA adapters

#### SAS HBA BIOS timeout settings for correct failover:

To make sure that the DS3500 storage subsystem provides correct path failover during heavy I/O in a VMware environment, you must reconfigure two timeout values in the BIOS of the SAS HBA. Failure to reconfigure the two timeout values in the BIOS of the SAS HBA can cause I/O timeouts. The values that you must change are Report Device Missing Delay and I/O Device Missing Delay.

For more information about this setting, refer to *System Storage DS3000 Storage Manager Version 10 Installation and Support Guide for Microsoft Windows Server, Linux, Novell NetWare, and VMware ESX Server* available at the following URL:

[ftp://ftp.software.ibm.com/systems/support/system\\_x\\_pdf/46m1364.pdf](ftp://ftp.software.ibm.com/systems/support/system_x_pdf/46m1364.pdf)

#### Configuring the server RAID

To configure the server RAID, at the prompt during reboot, press Ctrl+A to enter the Controller Configuration menu. Configure the Internal server RAID controller for RAID1 or RAID 10 configuration, which allows you to preserve a working OS set in case of a drive failure. Performance on the local drives is not as critical as compared to the actual Virtual Machines Data stores.

### 22.4.2 Configuring the software on the VMware ESX Server host

Continue with the procedure to install the VMware vSphere software on the VMware ESX Server host using these steps:

1. On the initial panel, select Install ESX in graphical mode as shown in Figure 22-12.




Figure 22-12 ESX Install initial panel

2. At the Welcome panel as shown in Figure 22-13, click **Next**.




Figure 22-13 ESX Welcome Screen

3. On the End User License Agreement window (Figure 22-14), accept the license terms and click **Next**.

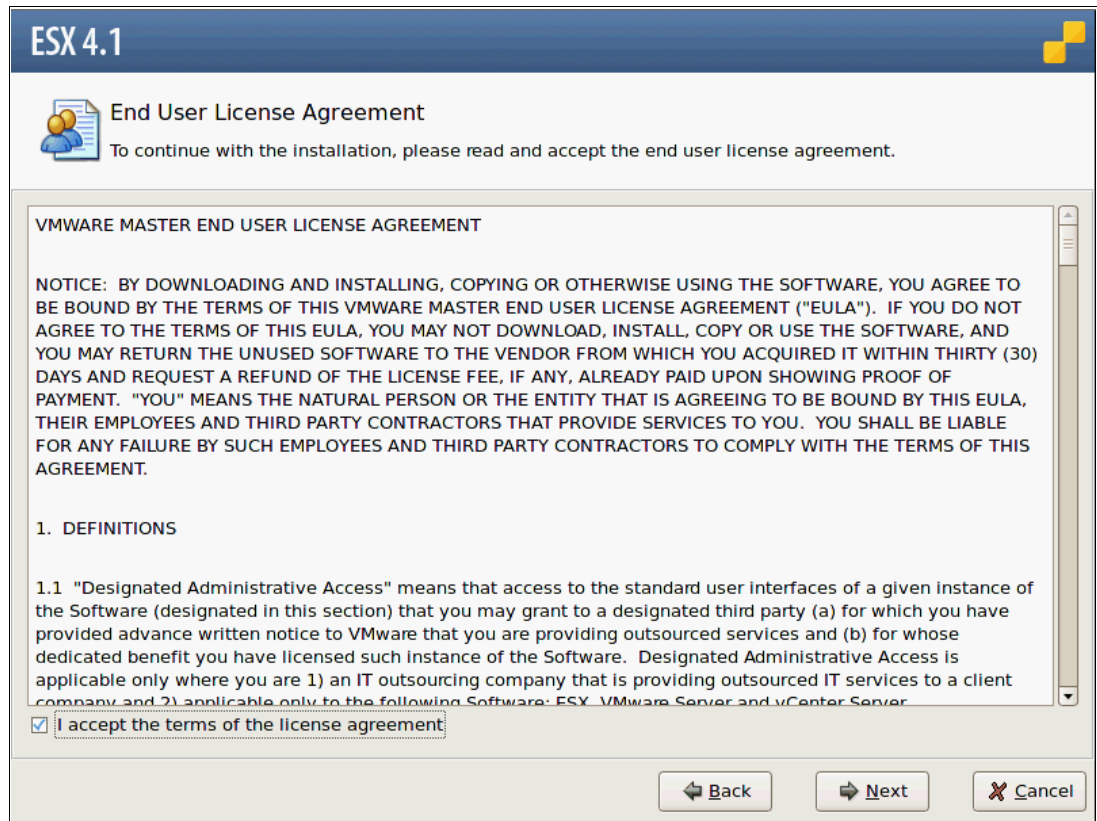


Figure 22-14 ESX End User License Agreement

4. On the Keyboard Selection window (Figure 22-15), select the Keyboard type you have and click **Next**.




Figure 22-15 ESX keyboard selection

5. In the Custom Drivers window (Figure 22-16), if needed, install any custom device drivers. If not, accept the default option (No) and click **Next**.




Figure 22-16 ESX Custom Drivers

6. In the Load Drivers window (Figure 22-17), click **Yes**.

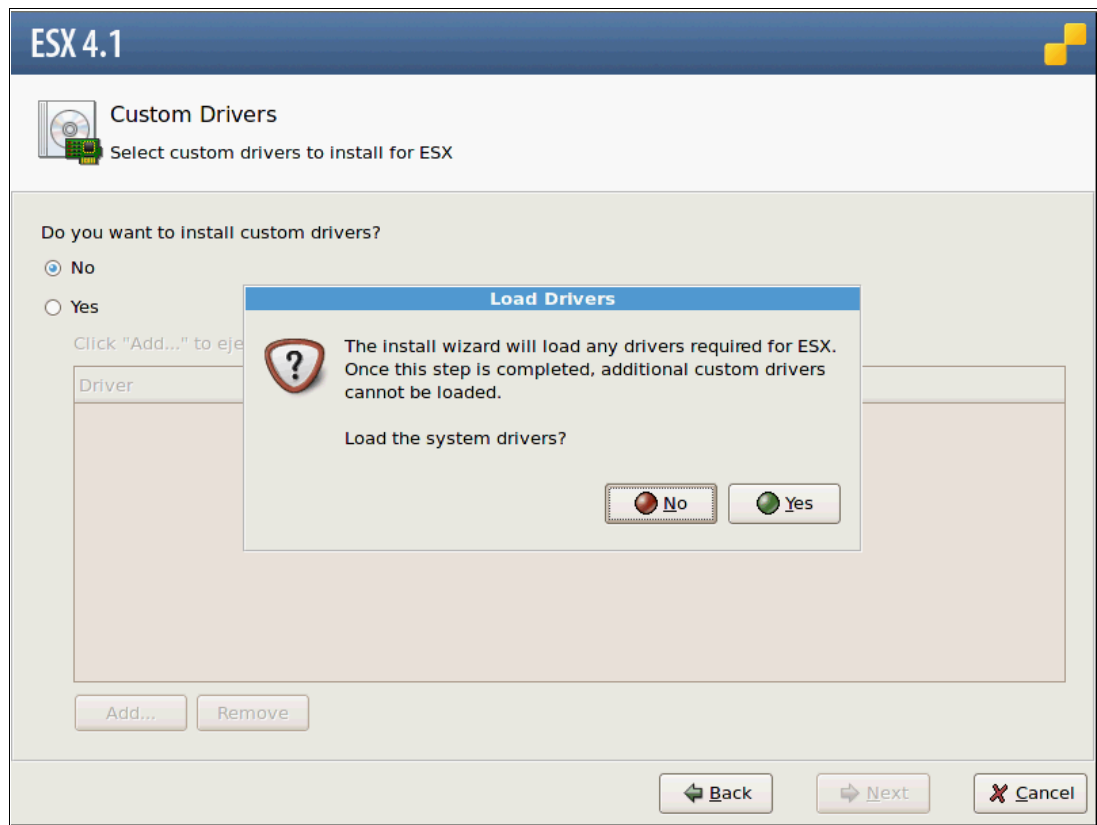


Figure 22-17 ESX Load Drivers Confirmation

The Drivers load process continues. Click **Next** to continue.



7. During driver load you will see a progress window like Figure 22-18. Wait until the progress bar shows the load is complete, then click **Next**.

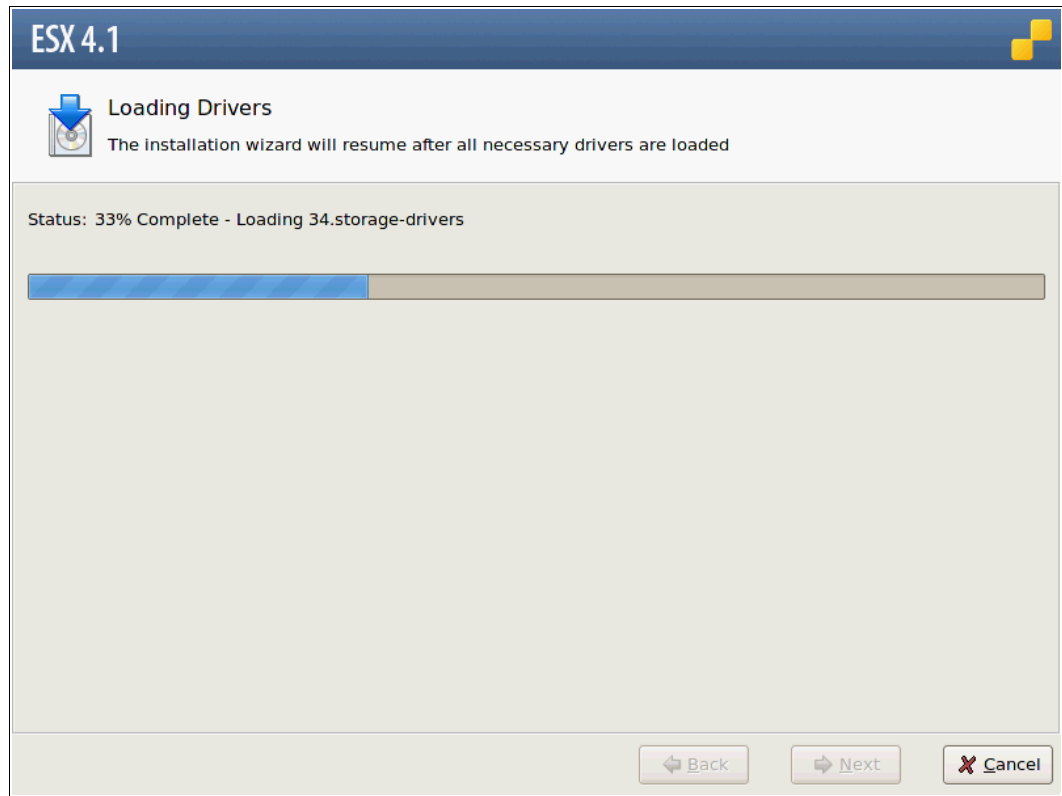
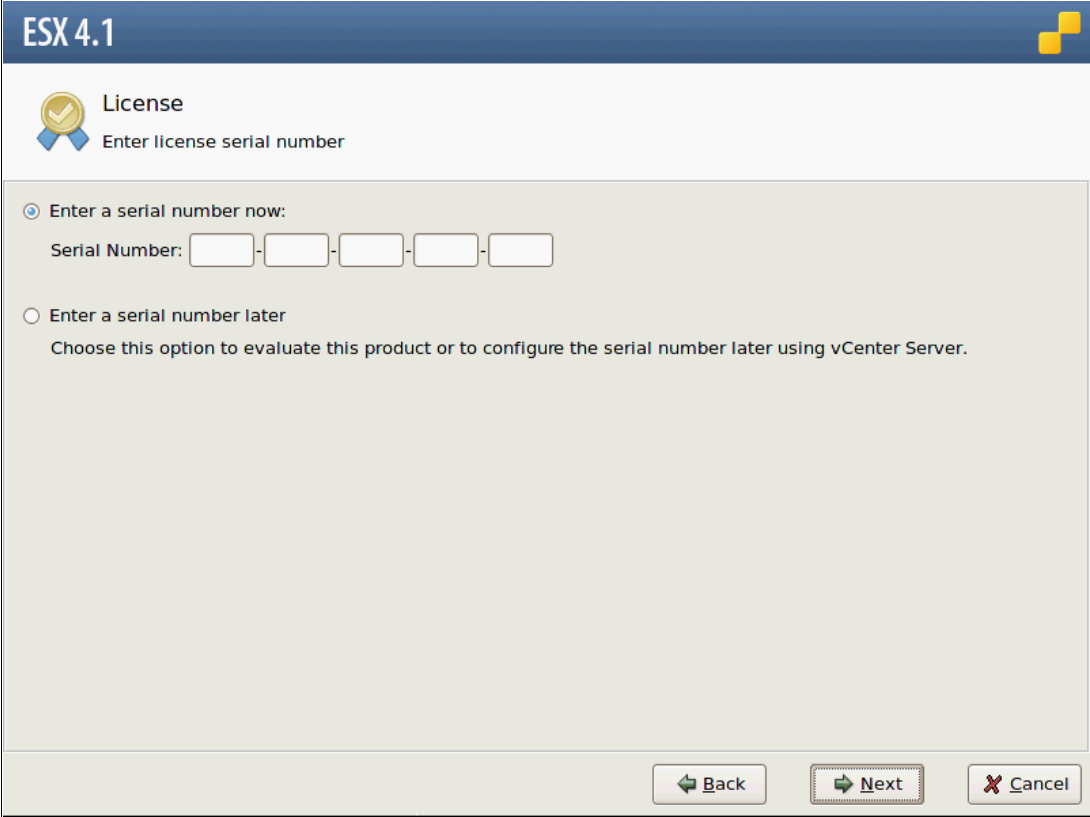



Figure 22-18 ESX Driver load

8. On the Licensing panel (Figure 22-19), type the serial number now, or you can enter the serial number later using vCenter Server. Click **Next**.



The image shows a software window titled "ESX 4.1" with a blue header bar. Below the header, there is a section titled "License" with a gold checkmark icon and the text "Enter license serial number". The main area contains two radio button options. The first option, "Enter a serial number now:", is selected and followed by a "Serial Number:" label and five empty text boxes separated by hyphens. The second option, "Enter a serial number later", is unselected and followed by the text "Choose this option to evaluate this product or to configure the serial number later using vCenter Server." At the bottom right, there are three buttons: "Back" with a left arrow, "Next" with a right arrow, and "Cancel" with a red X.

ESX 4.1

 **License**  
Enter license serial number

☒ Enter a serial number now:  
Serial Number:  -  -  -  -

☐ Enter a serial number later  
Choose this option to evaluate this product or to configure the serial number later using vCenter Server.

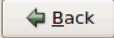
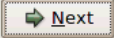
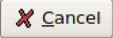
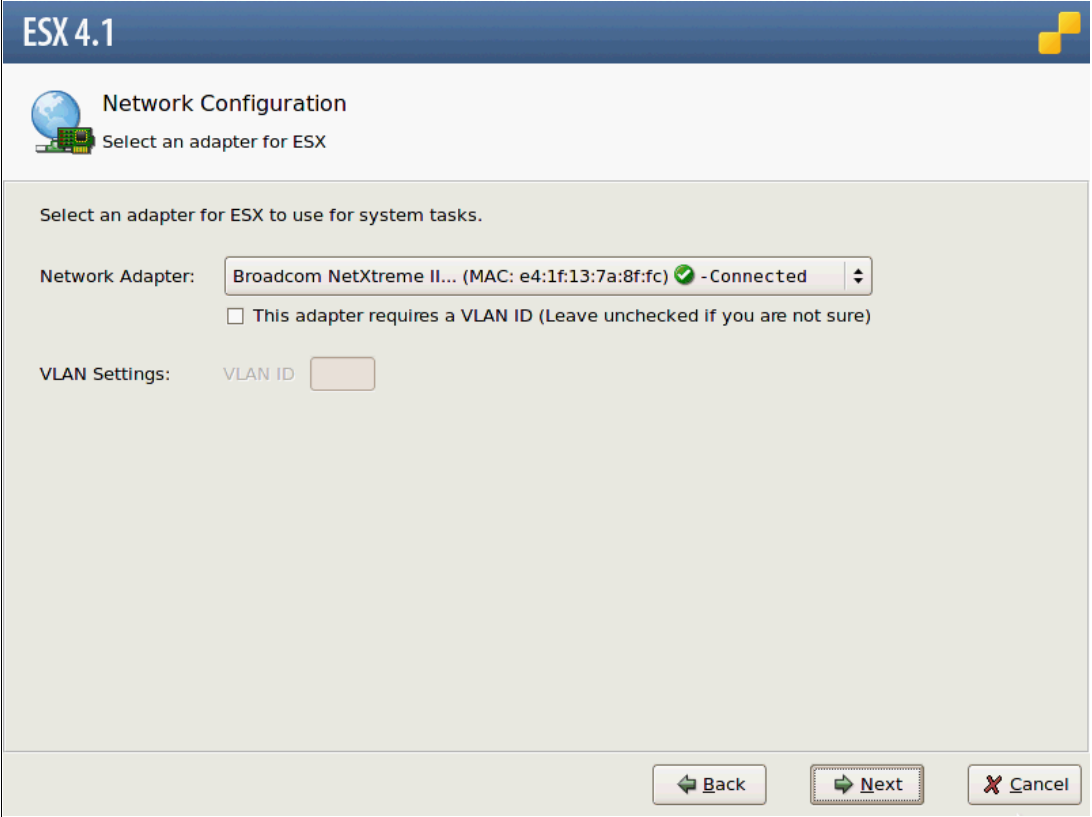
  

Figure 22-19 ESX License information

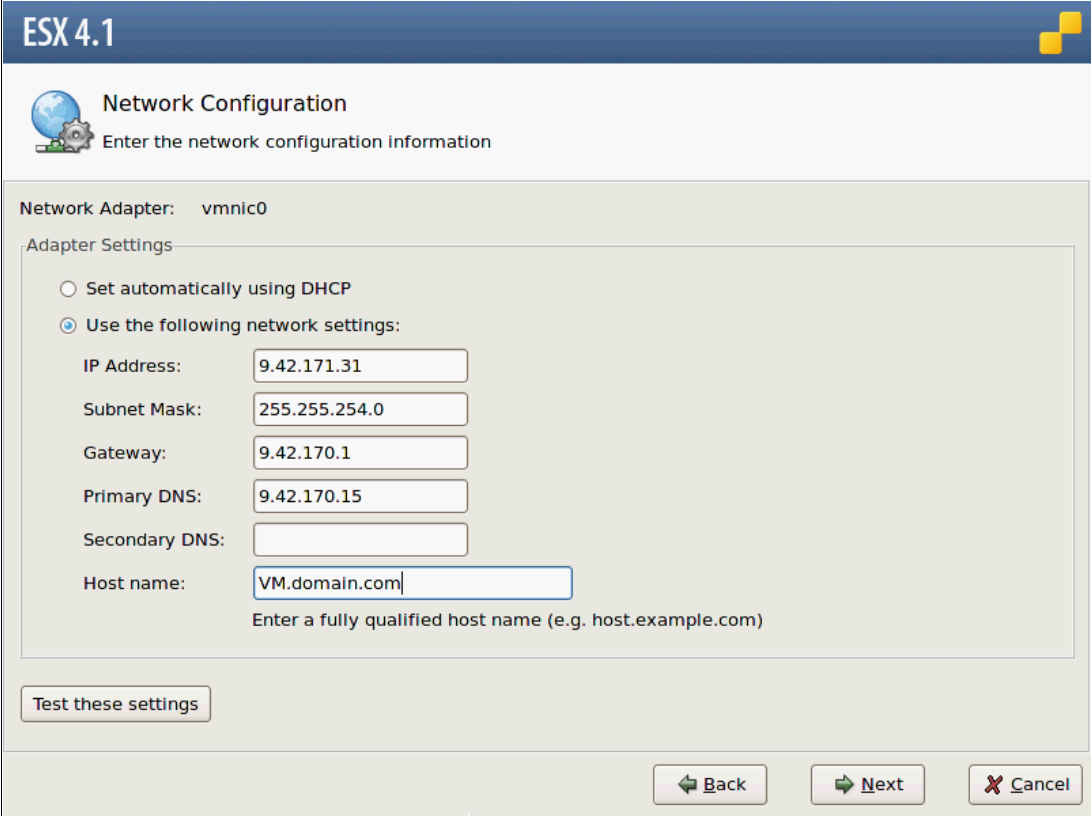
9. In the Network Configuration window (Figure 22-20), select an Adapter for the ESX Console Default switch and click **Next**.



The image shows the ESX 4.1 Network Configuration window. The title bar is blue with 'ESX 4.1' on the left and a yellow VMware logo on the right. Below the title bar, there's a header area with a globe icon and the text 'Network Configuration' and 'Select an adapter for ESX'. The main content area has a light beige background and contains the instruction 'Select an adapter for ESX to use for system tasks.' Below this, there's a 'Network Adapter:' label followed by a dropdown menu showing 'Broadcom NetXtreme II... (MAC: e4:1f:13:7a:8f:fc) - Connected' with a green checkmark and a dropdown arrow. Below the dropdown is a checkbox labeled 'This adapter requires a VLAN ID (Leave unchecked if you are not sure)'. Underneath, there's a 'VLAN Settings:' label followed by 'VLAN ID' and an empty text input field. At the bottom right, there are three buttons: 'Back' with a left arrow, 'Next' with a right arrow, and 'Cancel' with a red X.


Figure 22-20 ESX Console Network Adapter

10. In the Network Configuration window (Figure 22-21), type the Network parameters using the information you recorded in Table 22-1 on page 652. The host name must be a valid DNS Host name. Click **Next**.



The image shows the ESX 4.1 Network Configuration window. The title bar is blue with 'ESX 4.1' on the left and a yellow square icon on the right. Below the title bar, there is a globe icon and the text 'Network Configuration' followed by 'Enter the network configuration information'. The main area is titled 'Network Adapter: vmnic0' and 'Adapter Settings'. It contains two radio buttons: 'Set automatically using DHCP' (unselected) and 'Use the following network settings:' (selected). Below these are text boxes for 'IP Address' (9.42.171.31), 'Subnet Mask' (255.255.254.0), 'Gateway' (9.42.170.1), 'Primary DNS' (9.42.170.15), 'Secondary DNS' (empty), and 'Host name' (VM.domain.com). A note below the host name box says 'Enter a fully qualified host name (e.g. host.example.com)'. At the bottom left is a 'Test these settings' button. At the bottom right are three buttons: 'Back' (with a left arrow), 'Next' (with a right arrow), and 'Cancel' (with a red X).

ESX 4.1

 **Network Configuration**  
Enter the network configuration information

Network Adapter: vmnic0

Adapter Settings

☐ Set automatically using DHCP

☒ Use the following network settings:

IP Address:

Subnet Mask:

Gateway:

Primary DNS:

Secondary DNS:

Host name:

Enter a fully qualified host name (e.g. host.example.com)

Figure 22-21 ESX Network Configuration

11. Click **Test these settings** to confirm correct network connectivity. See Figure 22-22. Click **OK** and then **Next** to continue.

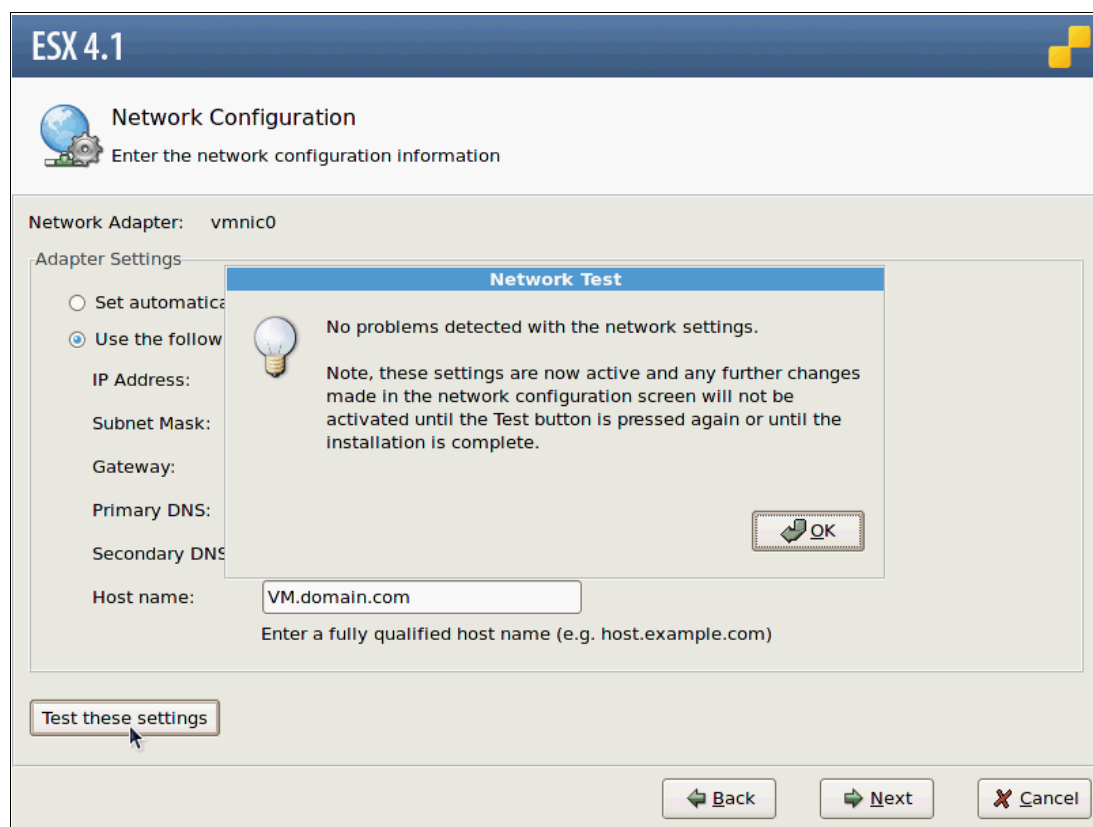


Figure 22-22 ESX Test Network settings

12. In the Setup Type window (Figure 22-23), if custom partitions configurations are not required, select **Standard setup** and click **Next**.

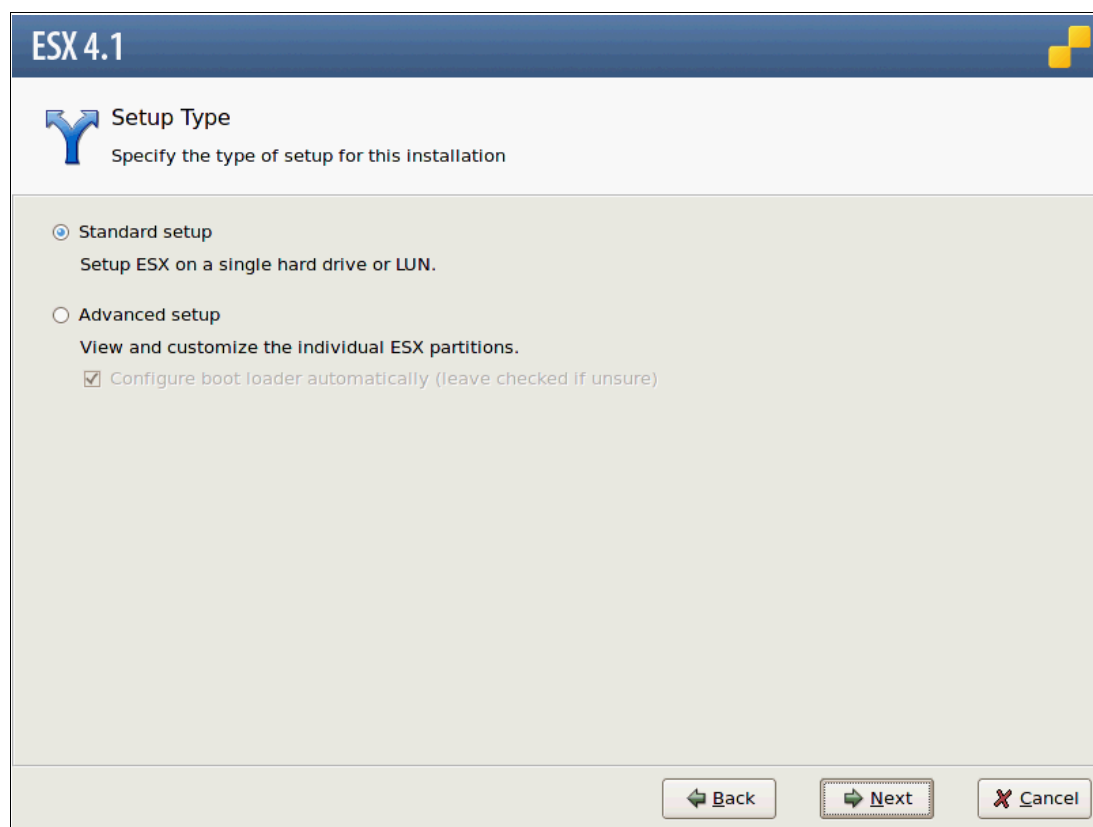


Figure 22-23 ESX Setup Type

13. Select the correct Storage location for the ESX binaries installation from the list in the ESX Storage Device window shown in Figure 22-24. In our example we have one local Harddisk and two SAN attached disk drives. Click **Next**.

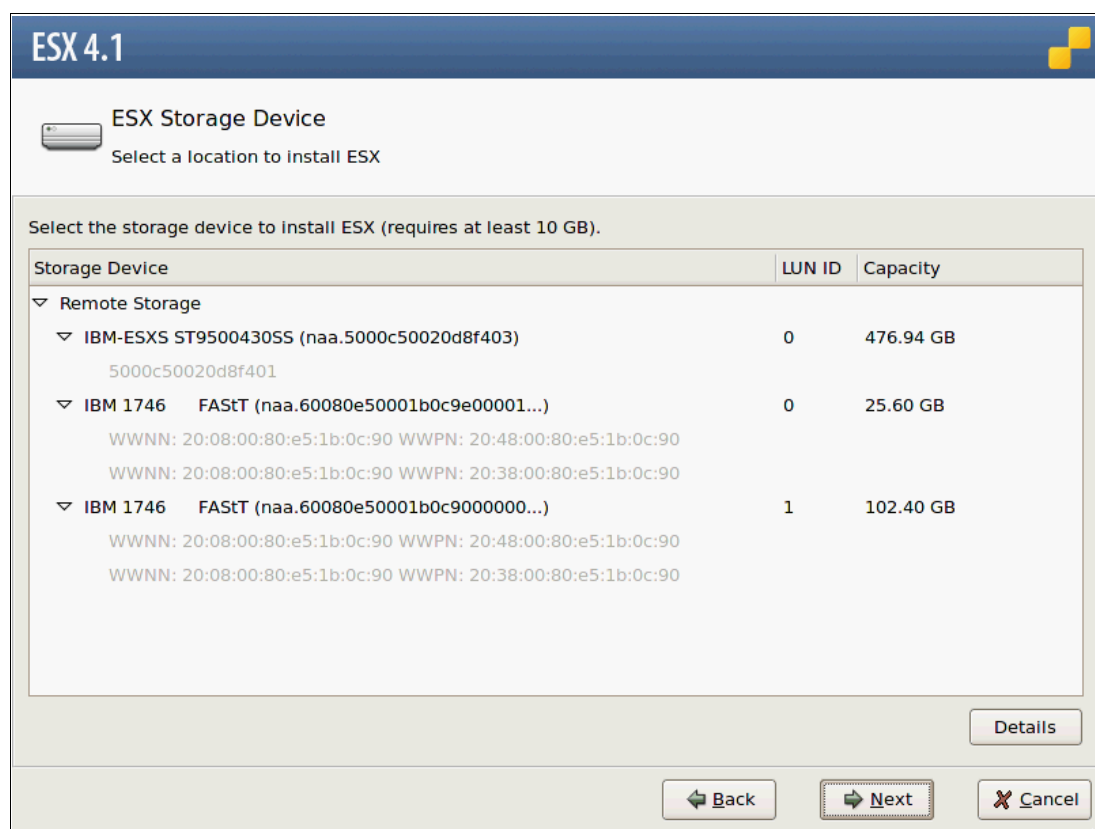


Figure 22-24 ESX Storage Device

14.If VMware finds an existing Installation, you will get a warning about it. You can choose to preserve your VMFS datastore on the Delete Device Contents warning panel as shown in Figure 22-25. Click **OK** to continue.

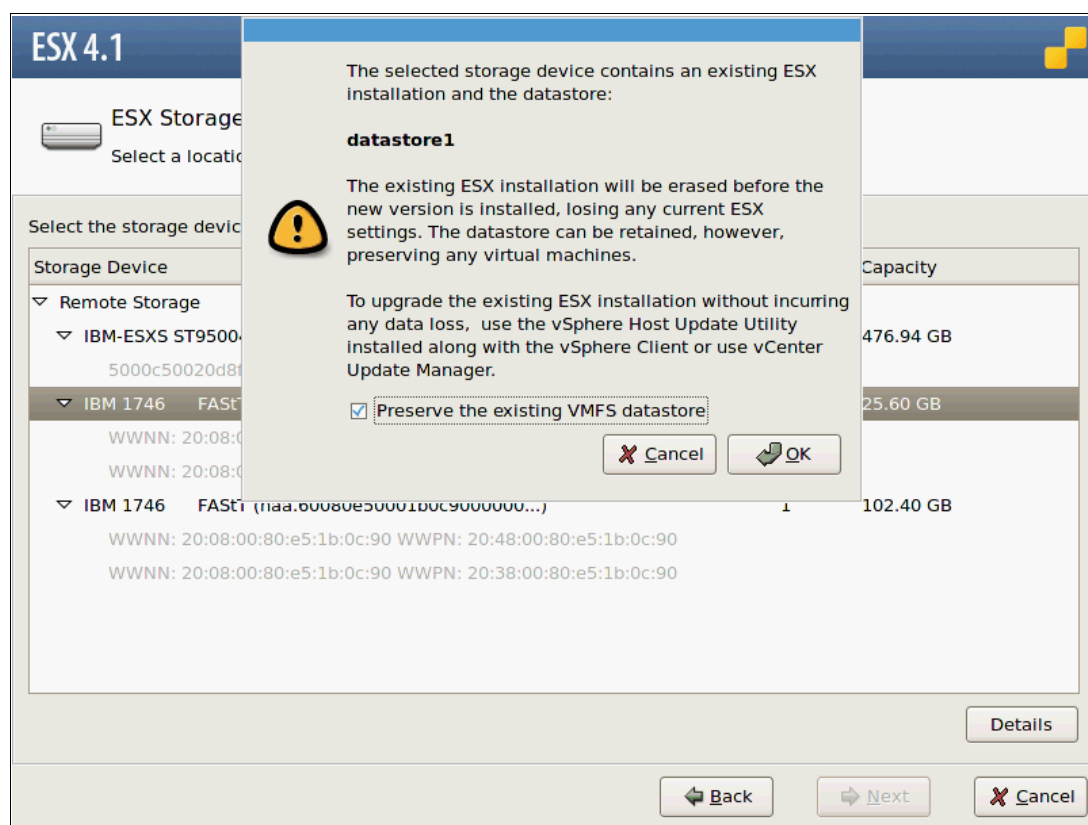


Figure 22-25 ESX Device content warning



15. In the Time Zone Settings window (Figure 22-26), select the appropriate time zone and click **Next**.

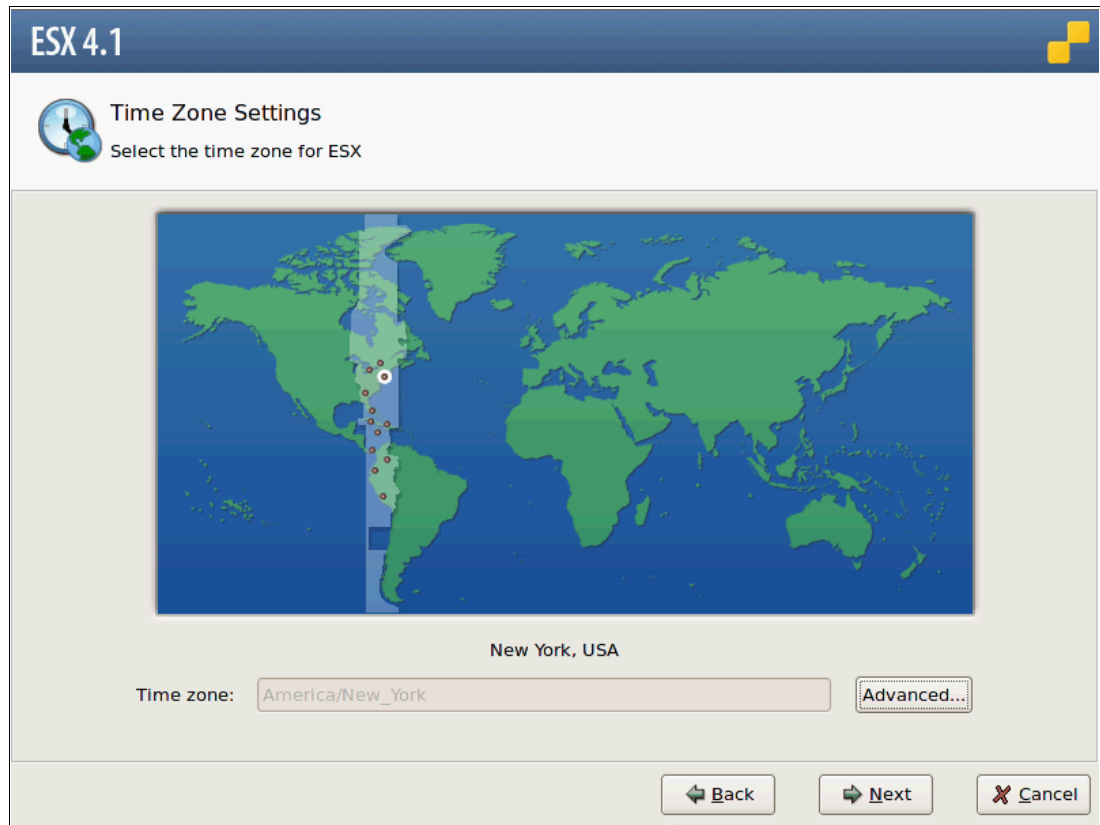



Figure 22-26 ESX Time Zone settings

16.In the Date and Time window (Figure 22-27), you can select the date and time manually.

ESX 4.1

 **Date and Time**  
Specify the date and time for ESX

How should the date and time for ESX be configured?

☐ Automatically

NTP Server:

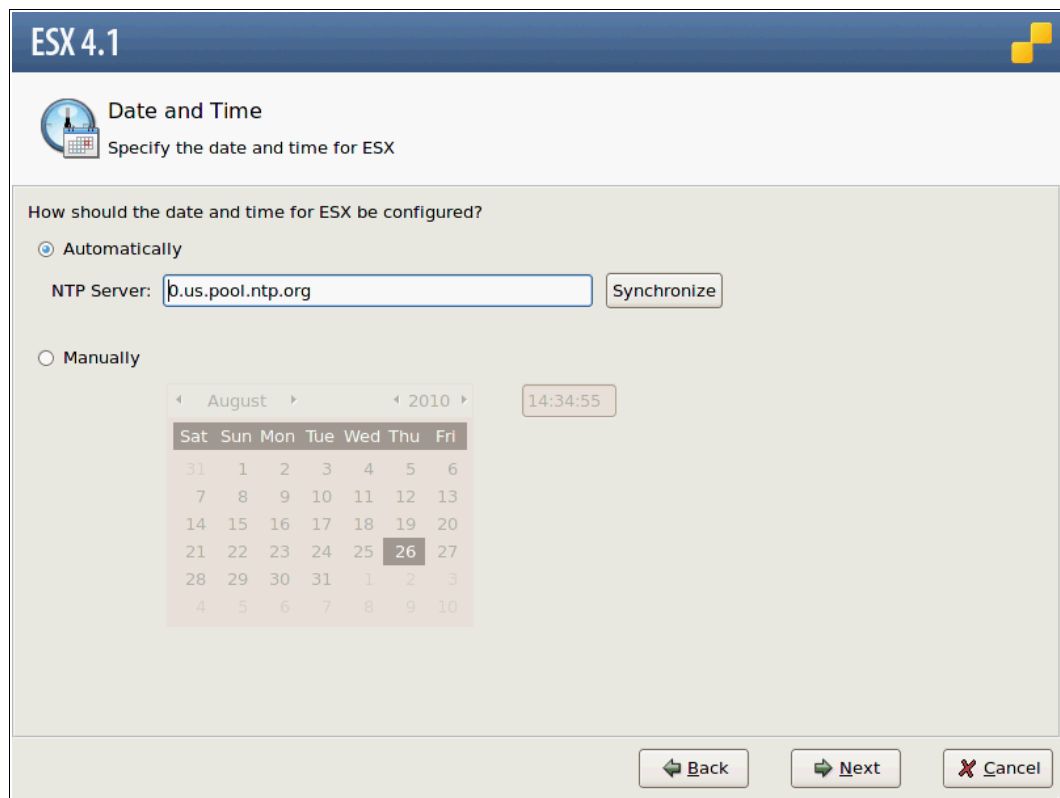
☒ Manually

◀ August ▶      ▶ 2010 ▶

| Sat | Sun | Mon | Tue | Wed | Thu | Fri |
|-----|-----|-----|-----|-----|-----|-----|
| 31  | 1   | 2   | 3   | 4   | 5   | 6   |
| 7   | 8   | 9   | 10  | 11  | 12  | 13  |
| 14  | 15  | 16  | 17  | 18  | 19  | 20  |
| 21  | 22  | 23  | 24  | 25  | 26  | 27  |
| 28  | 29  | 30  | 31  | 1   | 2   | 3   |
| 4   | 5   | 6   | 7   | 8   | 9   | 10  |


Figure 22-27 ESX Date and Time settings

You can also enter an NTP Server as shown in Figure 22-28. At the URL <http://www.pool.ntp.org/en/>, you will find a summary of available, active NTP servers worldwide. Click **Next** to continue.



The image shows the ESX 4.1 'Date and Time' configuration window. The title bar says 'ESX 4.1'. Below the title, there's a clock icon and the text 'Date and Time' and 'Specify the date and time for ESX'. The main area asks 'How should the date and time for ESX be configured?'. There are two radio buttons: 'Automatically' (selected) and 'Manually'. Under 'Automatically', there is an 'NTP Server:' label, a text box containing '0.us.pool.ntp.org', and a 'Synchronize' button. Under 'Manually', there is a calendar widget showing 'August' and '2010'. The calendar has a table of dates with '26' highlighted. To the right of the calendar is a time input box showing '14:34:55'. At the bottom right, there are three buttons: 'Back', 'Next', and 'Cancel'.

ESX 4.1

 **Date and Time**  
Specify the date and time for ESX

How should the date and time for ESX be configured?

☒ Automatically

NTP Server:

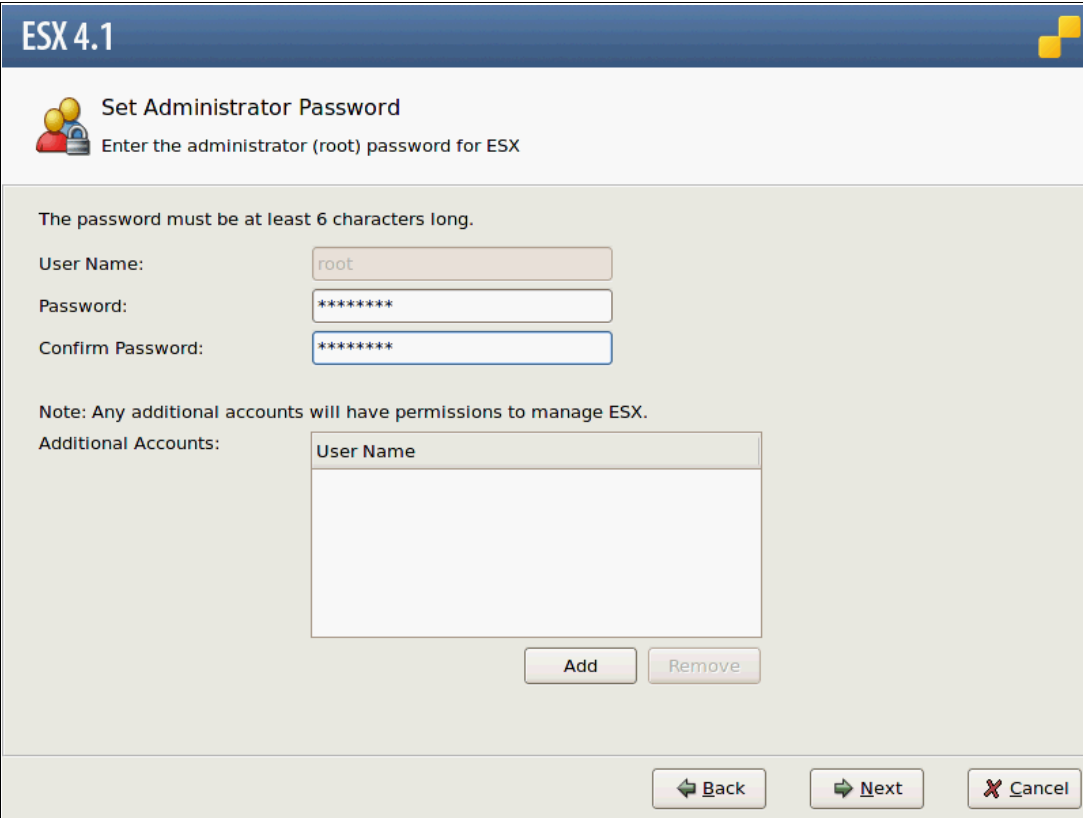
☐ Manually

August 2010 14:34:55

| Sat | Sun | Mon | Tue | Wed | Thu | Fri |
|-----|-----|-----|-----|-----|-----|-----|
| 31  | 1   | 2   | 3   | 4   | 5   | 6   |
| 7   | 8   | 9   | 10  | 11  | 12  | 13  |
| 14  | 15  | 16  | 17  | 18  | 19  | 20  |
| 21  | 22  | 23  | 24  | 25  | 26  | 27  |
| 28  | 29  | 30  | 31  | 1   | 2   | 3   |
| 4   | 5   | 6   | 7   | 8   | 9   | 10  |


Figure 22-28 ESX NTP Server selection

17. In the Set Administrator Password window (Figure 22-29), type the administrator password. You also have the option to create additional user accounts. Click **Next** to continue.



The screenshot shows the 'Set Administrator Password' window for ESX 4.1. The window has a blue header bar with 'ESX 4.1' and a yellow logo. Below the header, there's a title bar with a lock icon and the text 'Set Administrator Password' and 'Enter the administrator (root) password for ESX'. The main area is light gray and contains a message: 'The password must be at least 6 characters long.' Below this, there are three input fields: 'User Name:' with 'root' entered, 'Password:' with '\*\*\*\*\*' entered, and 'Confirm Password:' with '\*\*\*\*\*' entered. A note states: 'Note: Any additional accounts will have permissions to manage ESX.' Below the note, there's a section for 'Additional Accounts' with a table header 'User Name' and an empty table body. At the bottom of the table are 'Add' and 'Remove' buttons. At the very bottom of the window are three buttons: 'Back', 'Next', and 'Cancel'.

ESX 4.1

 **Set Administrator Password**  
Enter the administrator (root) password for ESX

The password must be at least 6 characters long.

User Name:

Password:

Confirm Password:

Note: Any additional accounts will have permissions to manage ESX.

Additional Accounts:

| User Name |
|-----------|
|-----------|

Figure 22-29 ESX Administrator Password

18. In the Summary panel (Figure 22-30), review the configuration. Click **Next**.

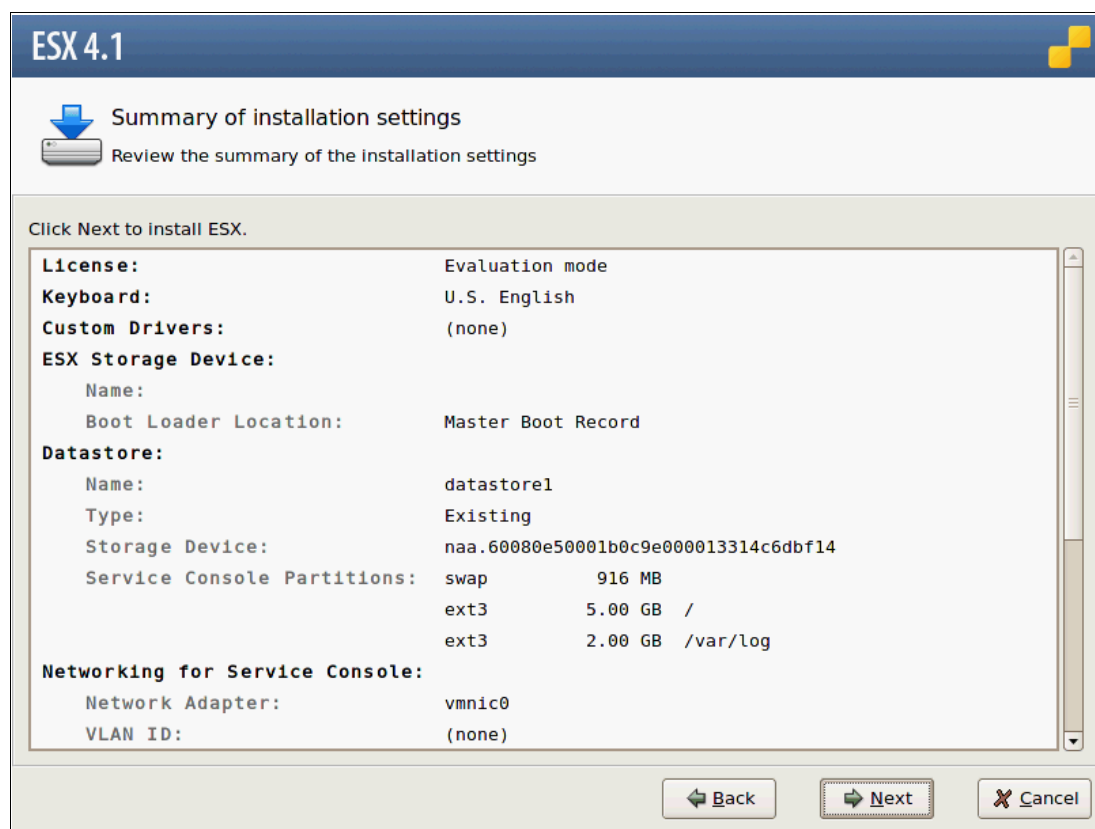


Figure 22-30 ESX Summary Installation window

19. The ESX Installation starts copying files as shown in Figure 22-31.




Figure 22-31 ESX installing

20. After the installation is complete, the window shown in Figure 22-32 displays. Click **Next**.




Figure 22-32 ESX Installing Completed

21. In the ESX 4.1 Installation Complete window (Figure 22-33), click **Finish**.

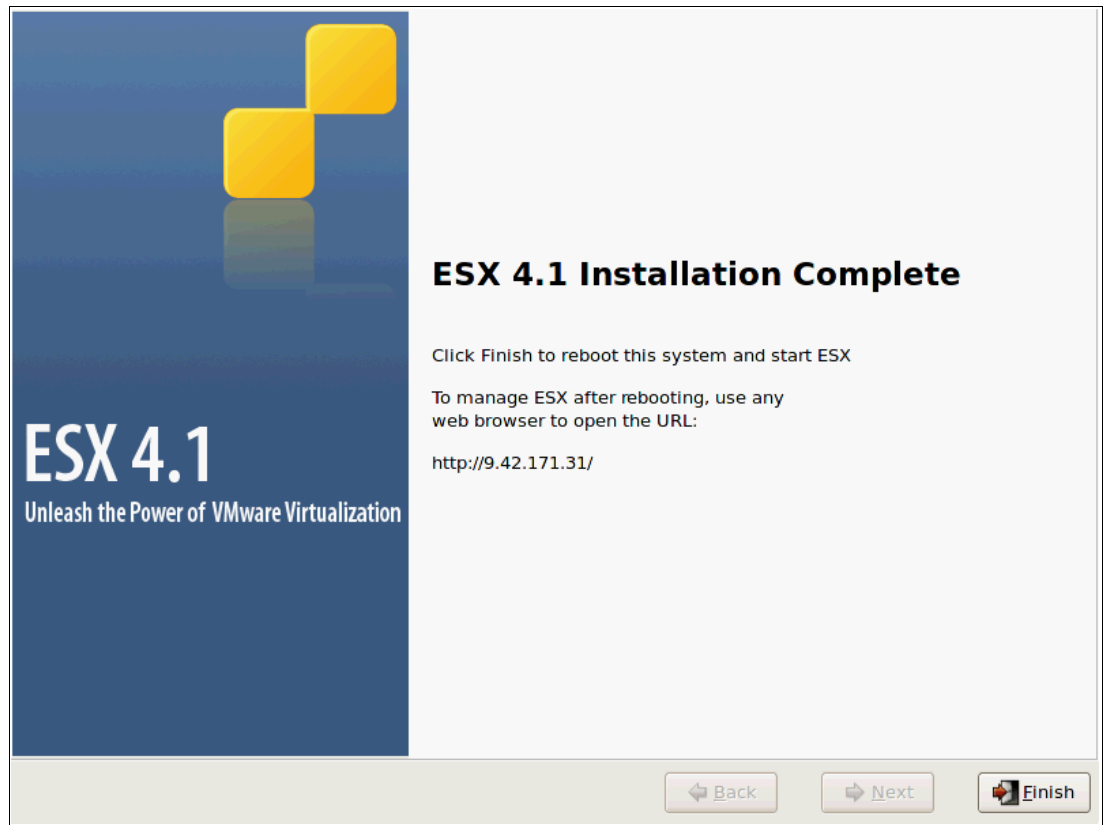


Figure 22-33 ESX Installation Completed

The vSphere Server reboots. You can now plug the Fibre Channel SAN connections back in if you installed VMware on a local disk.



After the server reboots, you are presented with the initial startup panel as shown in Figure 22-34.

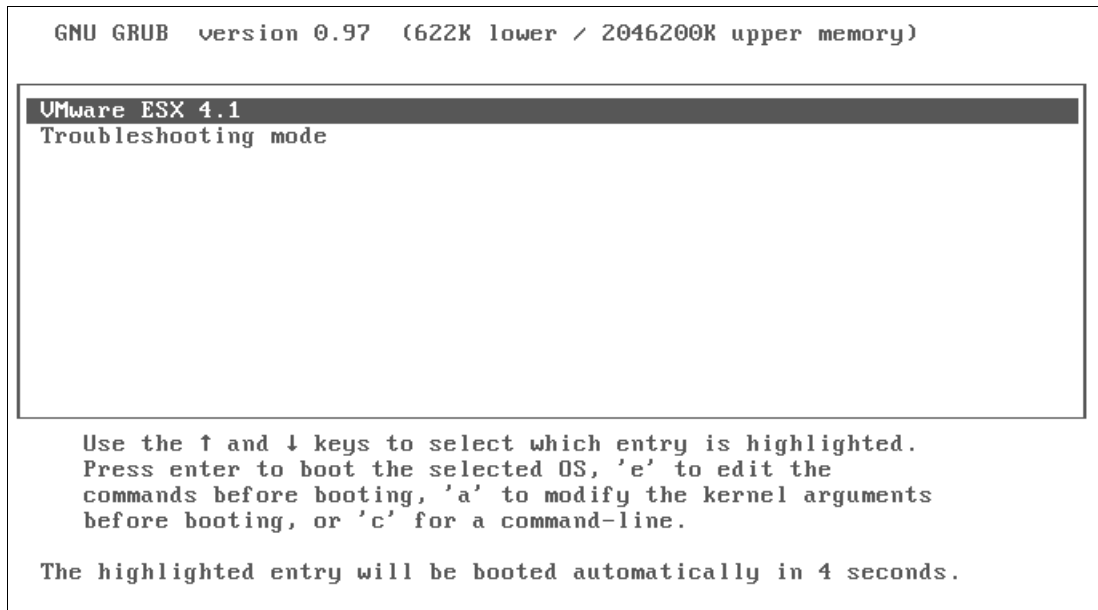


Figure 22-34 ESX Initial Startup Screen

22. The boot panel follows as shown in Figure 22-35.




Figure 22-35 ESX Boot sequence

23. After the VMware ESX Server startup is complete, the logon panel for VMware ESX Server 4.1 server (Figure 22-36) is displayed.




Figure 22-36 ESX Server Logon panel

### 22.4.3 Connecting to the VMware vSphere Server

At this stage, we are now ready to connect to the VMware ESX Server Host. Perform the following procedure on the management workstation that will be used for administration of the VMware ESX Server:

1. Using your web browser, connect to the Hostname or IP address of the newly created VMware ESX Server. We used Firefox as the Internet browser application. The error in Figure 22-37 is normal behavior.

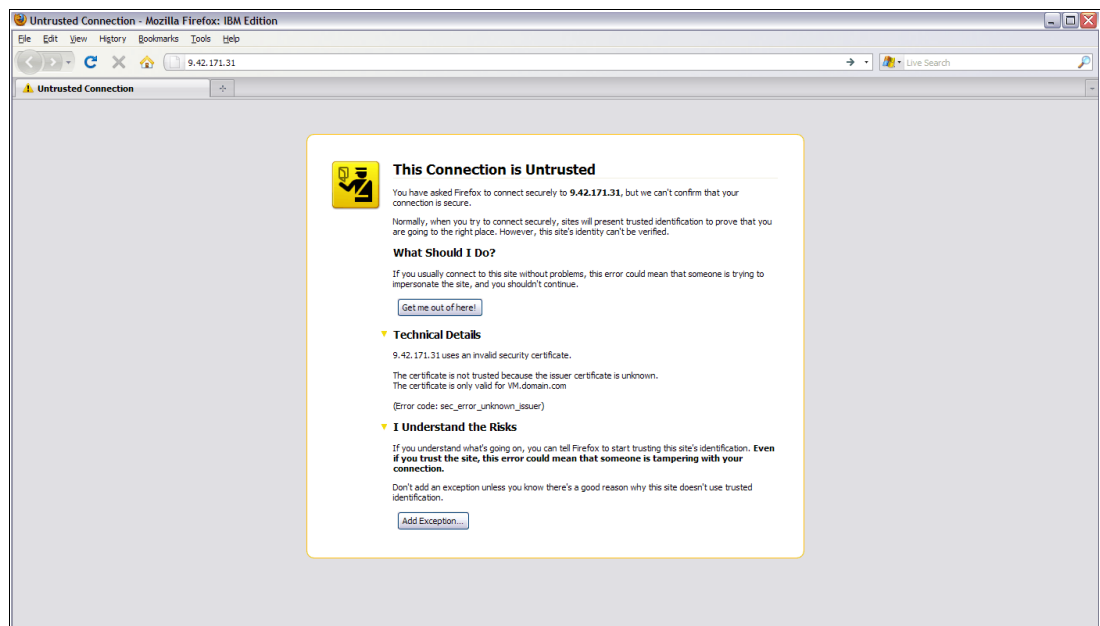


Figure 22-37 Connecting to ESX Server

2. Click **Add an Extension**. Figure 22-38 is presented to acquire the SSL Certificate. Click **Get Certificate** and then **Confirm Security Exception**.

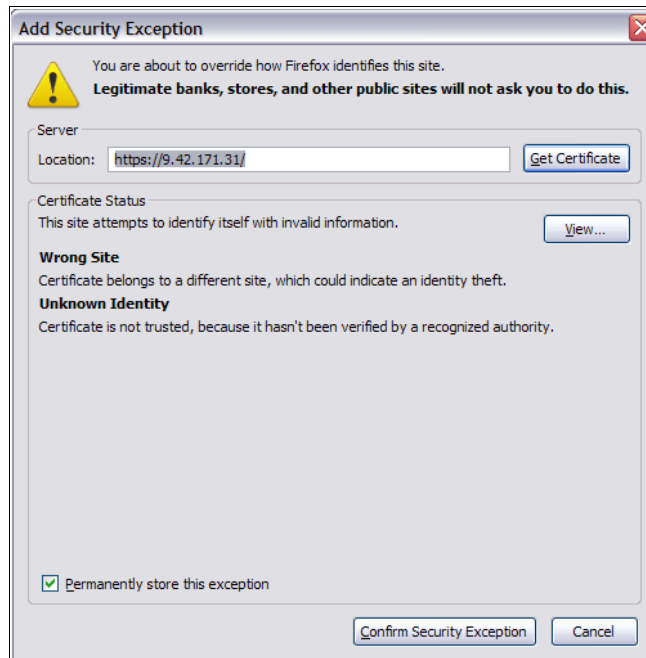


Figure 22-38 Add Security Exception

You are now presented with the VMware ESX Server as shown in Figure 22-39.

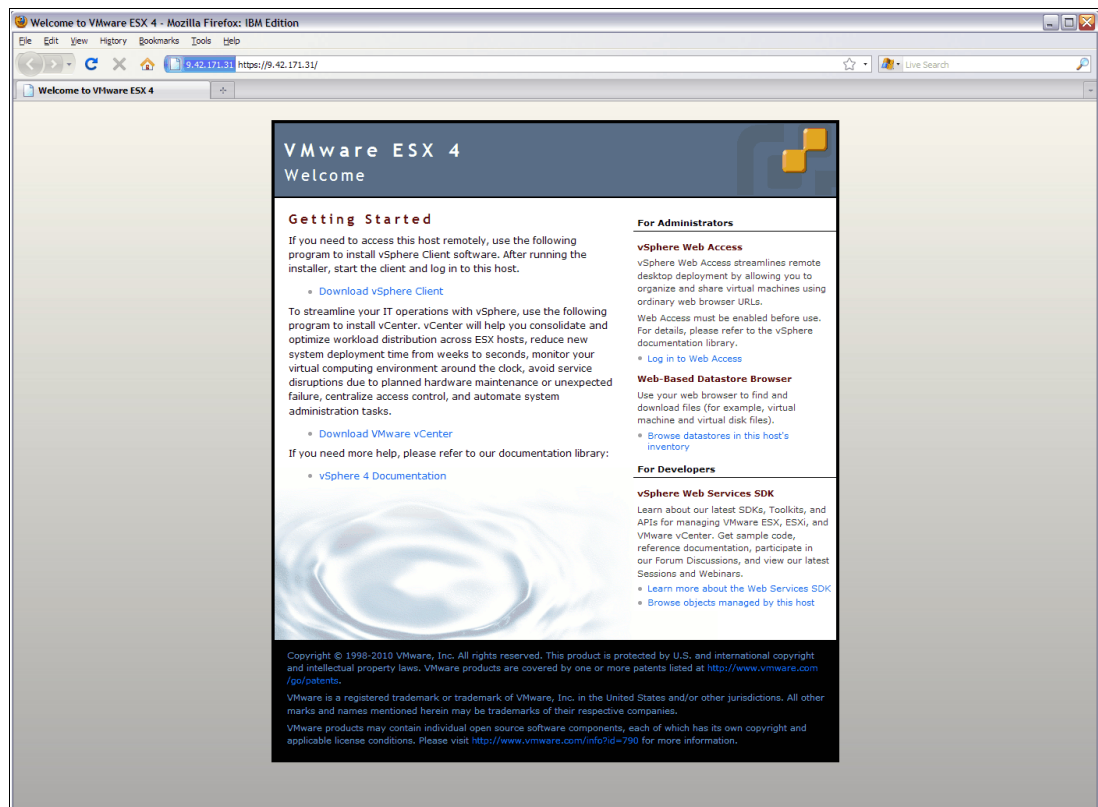


Figure 22-39 ESX vSphere Server Webpage

3. Click **Download vSphere Client** to start the setup package on the system that will be used as the initial administrative workstation.
4. Using Add/Remove Programs, install the VMware vSphere client. Figure 22-40 is the first panel of the Client Setup process. Select your language and click **OK**.




Figure 22-40 ESX client installation: Select language

5. In the Welcome window (Figure 22-41), click **Next**.

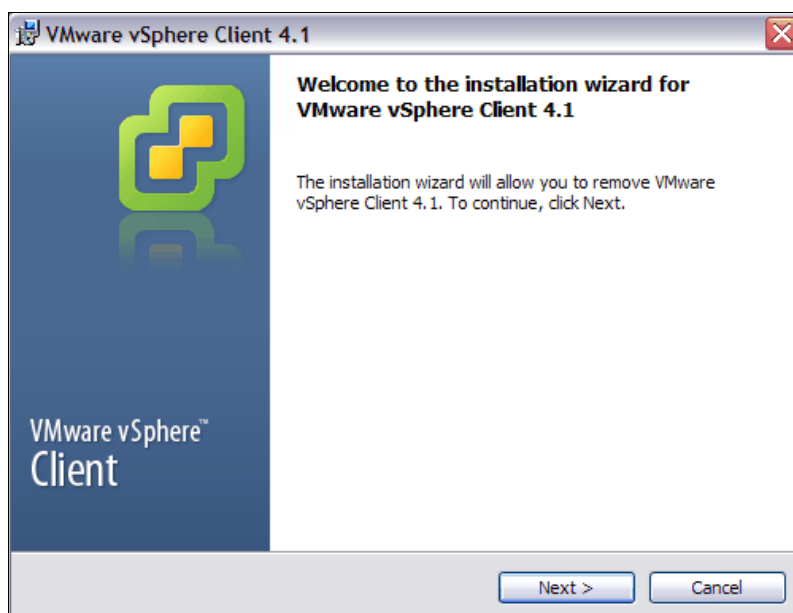


Figure 22-41 ESX Client Installation - Welcome panel

6. In the End/User Patent Agreement window (Figure 22-42 on page 683), click **Next**

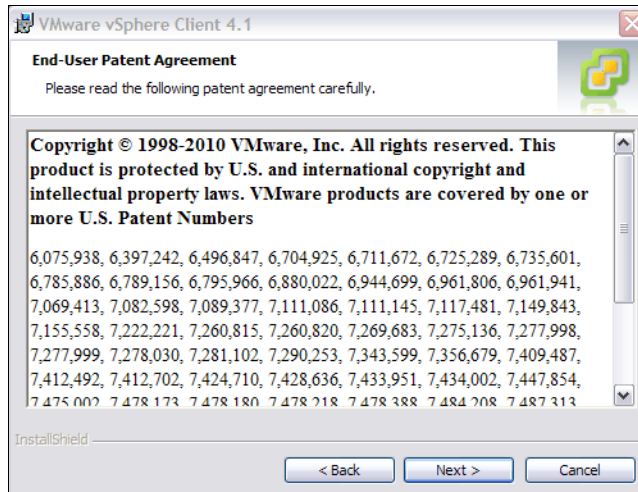


Figure 22-42 VMware vSphere Client End-User Agreement

7. In the License Agreement window (Figure 22-43), select **I agree** and click **Next**.

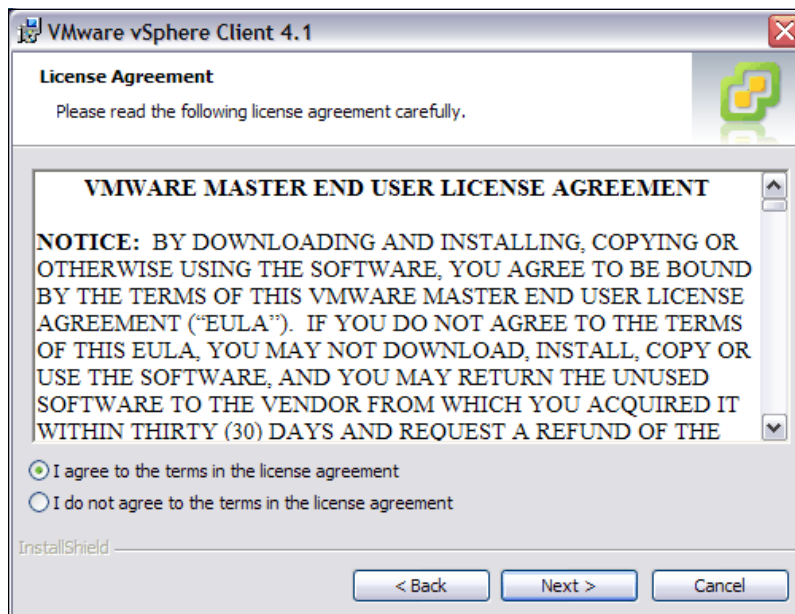


Figure 22-43 ESX Client Installation - License Agreement

8. In the Customer Information window (Figure 22-44 on page 684), type your organizational information and click **Next**.




Figure 22-44 ESX Client Installation - Organizational Information

9. In the Ready to Install window (Figure 22-45), click **Install**. The installation starts.




Figure 22-45 ESX Client Installation - Installing

10. After the installation is completed, click **Finish** as shown in Figure 22-46 on page 685.

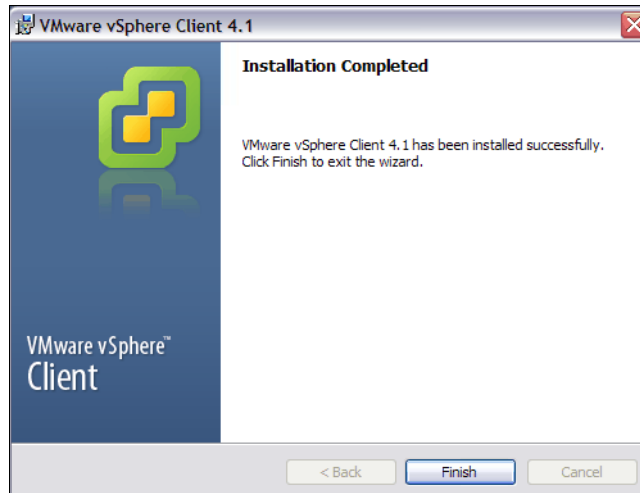


Figure 22-46 ESX Client Installation - Completed

11. Using the vSphere Client shortcut, launch the application. Type the IP address or the Hostname of the newly created VMware ESX Server as shown in Figure 22-47 and click **Login**.




Figure 22-47 ESX vSphere Client

12. For now, ignore the Certificate warning (Figure 22-48 on page 686) by clicking **Ignore**.



Figure 22-48 ESX vSphere client - Certificate warning

13. Figure 22-49 is the initial administrative interface for the VMware ESX Server. Click the **Inventory** icon.




Figure 22-49 ESX vSphere Client - initial administrative interface

14. The summary window (Figure 22-50 on page 687) for the newly created VMware ESX Server opens. Your VMware ESX Server installation is now completed. A tutorial page will come up on the first tab.



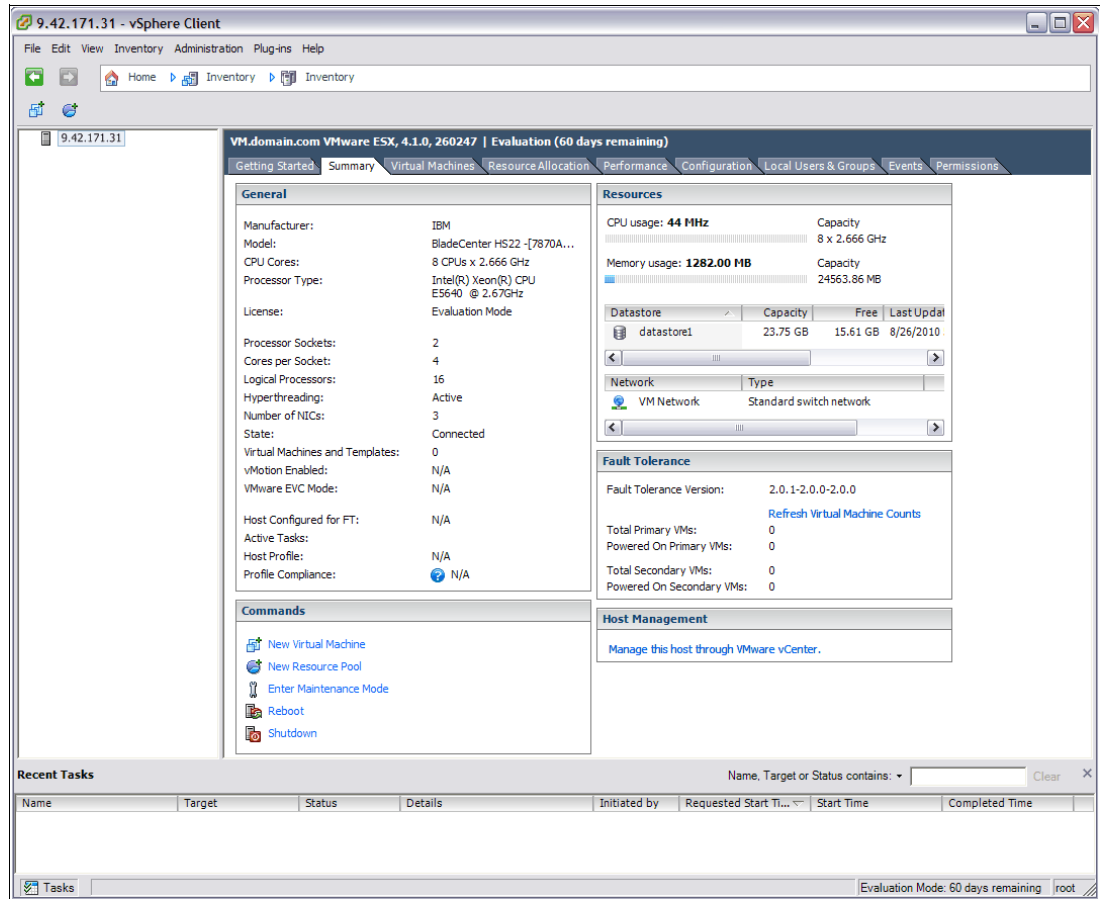


Figure 22-50 ESX vSphere Client - Summary panel

## 22.4.4 Post-Install Server configuration

The procedure in this section must be performed after the VMware ESX Server software installation is completed.

### Enabling SSH connectivity

This procedure describes the steps to follow to enable SSH connectivity for the VMware ESX Server:

1. Using the ESX server console, log into the VMware ESX Server.
2. Change directory to `/etc/ssh/` and copy the `sshd_config` file into `sshd_config.orig`.
3. vi the `sshd_config` file.
4. Scroll down until you find `PermitRootLogin no`, and change it to `PermitRootLogin yes`. See Figure 22-51 on page 688

```

#ListenAddress ::

HostKey for protocol version 1
#HostKey /etc/ssh/ssh_host_key
HostKeys for protocol version 2
#HostKey /etc/ssh/ssh_host_rsa_key
#HostKey /etc/ssh/ssh_host_dsa_key

Lifetime and size of ephemeral version 1 server key
#KeyRegenerationInterval 1h
#ServerKeyBits 768

Logging
obsoletes QuietMode and FascistLogging
#SyslogFacility AUTH
SyslogFacility AUTHPRIV
#LogLevel INFO

Authentication:

#LoginGraceTime 2m
PermitRootLogin yes_
#StrictModes yes
#MaxAuthTries 6
-- INSERT --

```

Figure 22-51 Permit Root Login

- Exit by pressing ESC and save your changes using the command **wq** for write/quit as shown in Figure 22-52.

```

#X11Forwarding no
X11Forwarding yes
#X11DisplayOffset 10
#X11UseLocalhost yes
#PrintMotd yes
#PrintLastLog yes
#TCPKeepAlive yes
#UseLogin no
#UsePrivilegeSeparation yes
#PermitUserEnvironment no
#Compression delayed
#ClientAliveInterval 0
#ClientAliveCountMax 3
#ShowPatchLevel no
#UseDNS yes
#PidFile /var/run/sshd.pid
#MaxStartups 10
#PermitTunnel no

no default banner path
#Banner /some/path

override default of no subsystems
Subsystem sftp /usr/libexec/openssh/sftp-server
:wq

```

Figure 22-52 Exit of vi

- Type **service sshd restart** as shown in Figure 22-53 on page 689.

```

override default of no subsystems
Subsystem sftp /usr/libexec/openssh/sftp-server
"sshd_config" 118L, 3300C written
[root@VM ssh]# service sshd restart
Stopping sshd: [OK]
Starting sshd: [OK]
[root@VM ssh]# _

```

Figure 22-53 SSHD restart

7. Type **logout**, and press Alt+F11 to return to your VMware Screen.

## 22.4.5 Configuring VMware ESX Server Storage

The following procedure demonstrates a basic configuration of SAN storage for a VMware ESX Server guest VM. This configuration might differ depending on your specific setup, for example clustered or shared. Refer to VMware documentation for more information.

1. Using the VMware vSphere Client, connect to the new VMware ESX Server as root as shown in Figure 22-54.




Figure 22-54 VMware vSphere Client Logon

2. Click the **Configuration** tab as shown in Figure 22-55 on page 690.

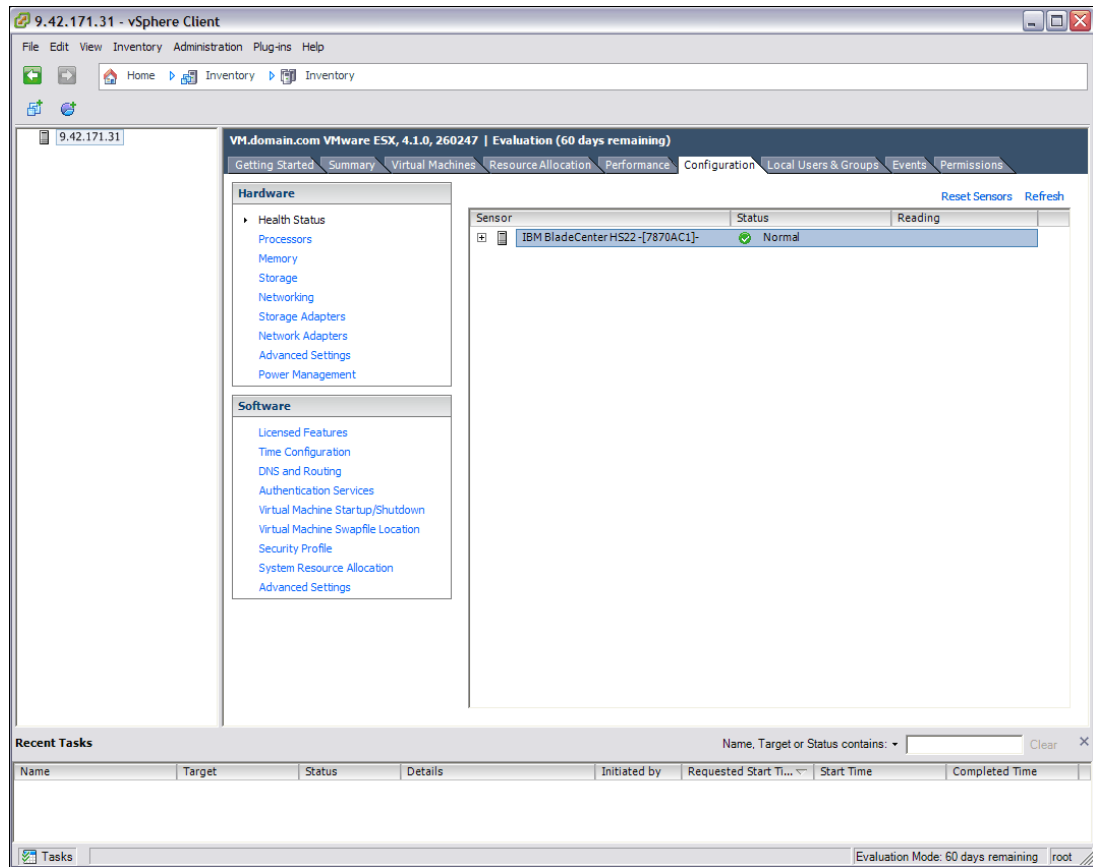


Figure 22-55 ESX Server - Configuration tab

3. In the Hardware section, select **Storage** and click **Add Storage** as shown in Figure 22-56 on page 691.

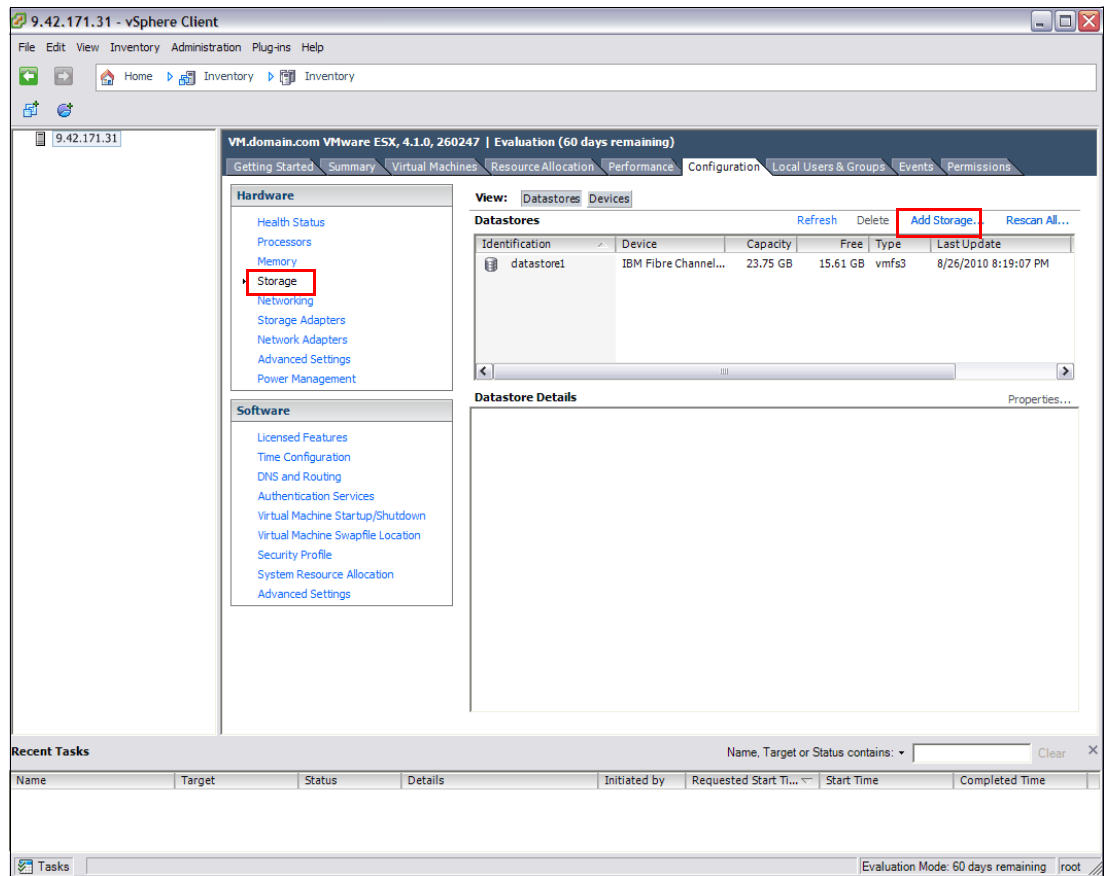


Figure 22-56 ESX Server - Storage selection

4. You are presented with the Storage Type selection pane. Select **Disk/LUN** to create a new datastore on the Fibre Channel SAN disks as shown in Figure 22-57 on page 692. Click **Next**.

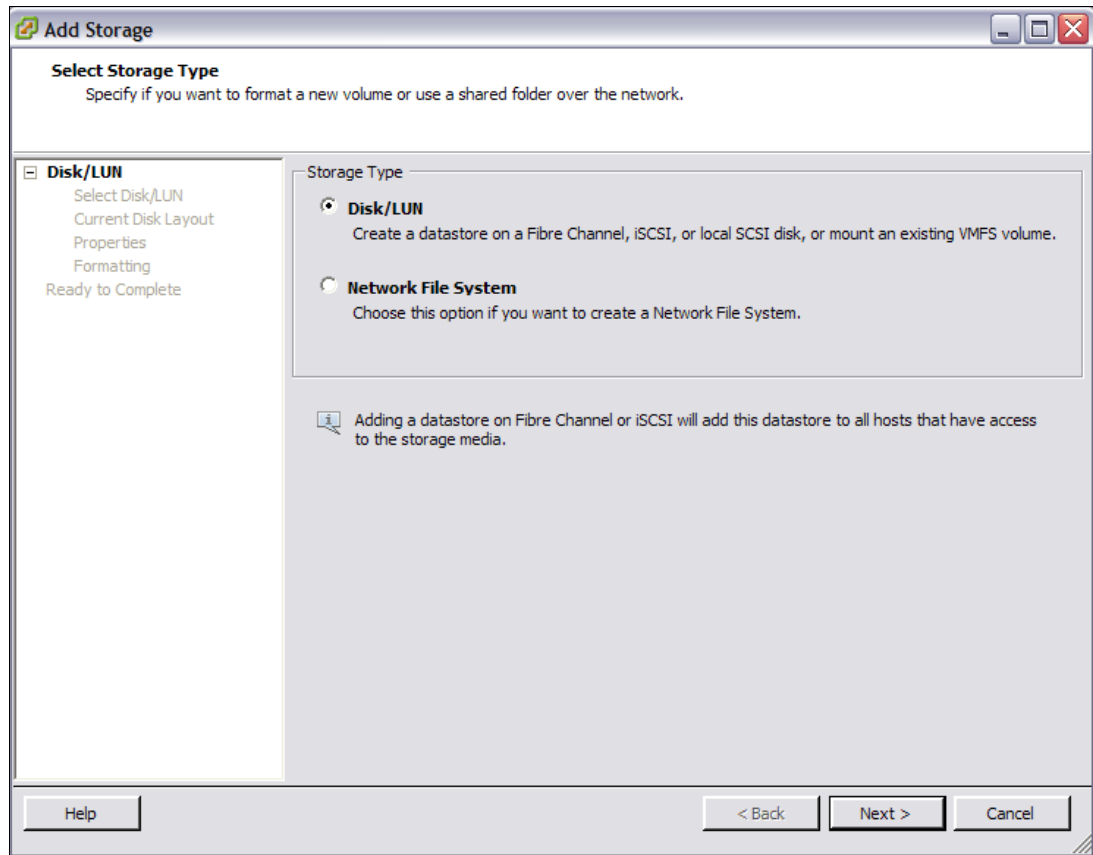


Figure 22-57 Storage Type selection

- As shown in Figure 22-58, select the SAN Fibre Channel LUN on which you want to create the Datastore VMFS partition. Click **Next**.




Figure 22-58 Select LUN

- Figure 22-59 on page 693 shows the disk layout of the LUN. Click **Next**.

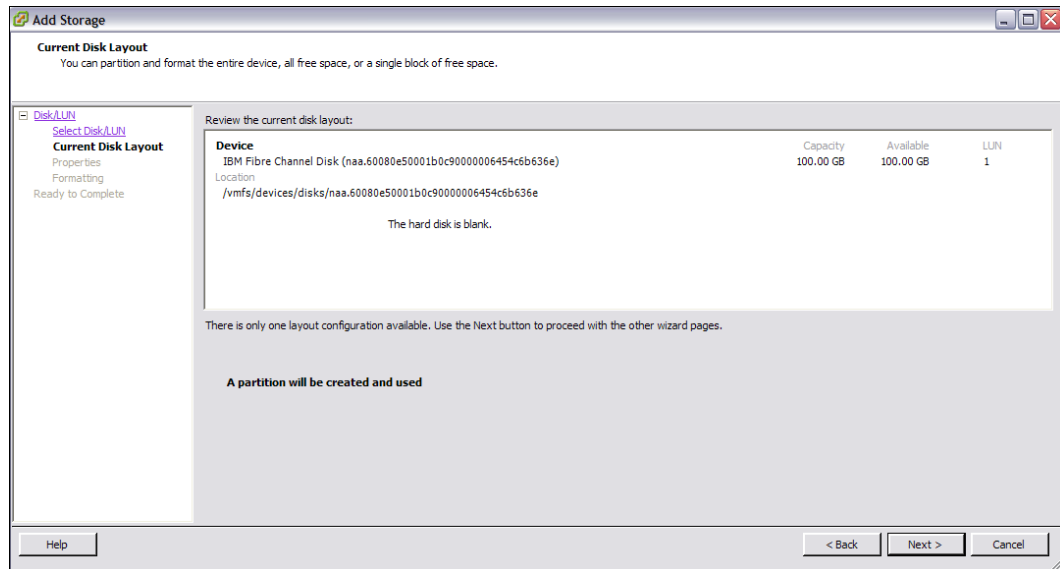


Figure 22-59 Disk layout

- As shown in Figure 22-60, type a descriptive name for the Datastore, for example, VMware\_Datastore, and click **Next**.




Figure 22-60 Datastore name

- As shown in Figure 22-61 on page 694, select the appropriate LUN format settings and click **Next**.

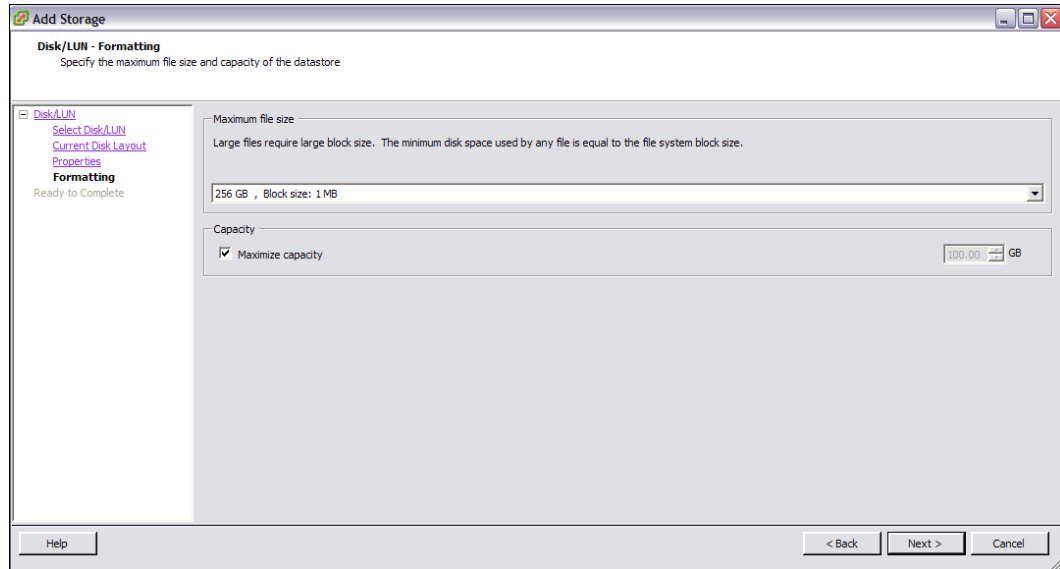


Figure 22-61 LUN formatting

9. Figure 22-62 is a summary panel for adding the storage. Click **Finish** to proceed.

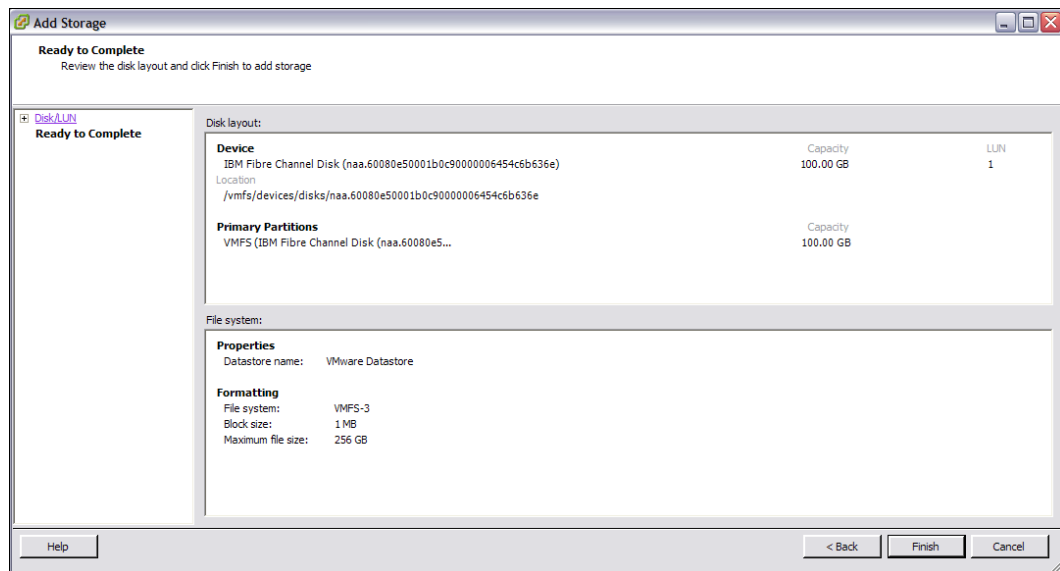


Figure 22-62 Adding Storage - Summary panel

10. As shown in Figure 22-63 on page 695, click **Refresh** to refresh the window to show the newly created Datstore.



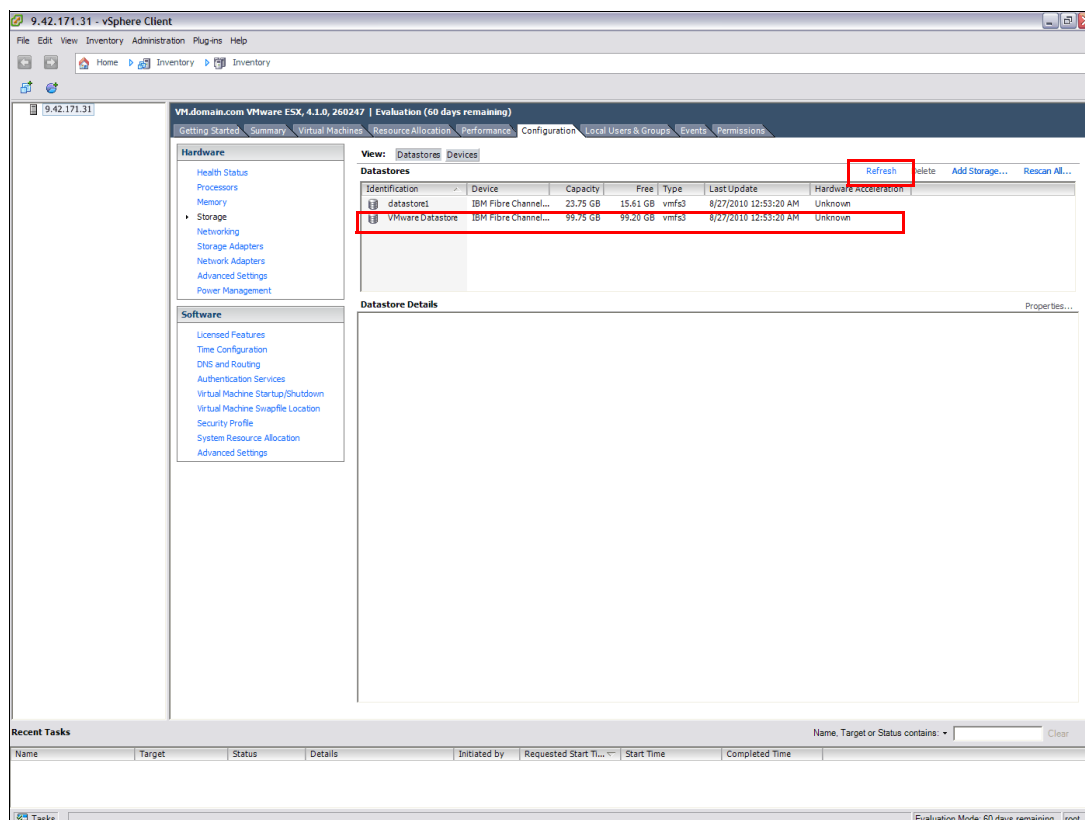


Figure 22-63 vSphere client window

11. Repeat these steps for additional SAN Fibre Channel LUNs.

## Assigning additional physical network adapters

To assign additional physical network adapters, perform the following steps:

1. In the Hardware section, select **Networking**. In our example, we have one virtual switch, vSwitch0, as shown in Figure 22-64 on page 696.




Figure 22-64 ESX Virtual Switch

2. Select **Properties** for the vSwitch0, and then click the **Network Adapters** tab as shown in Figure 22-65 on page 697. Click **Add**.




Figure 22-65 Adding Network Adapters

3. Add the adapter (vmnic1 in our example), and click **Next**, **Next**, and then **Finish**. See Figure 22-66 on page 698.

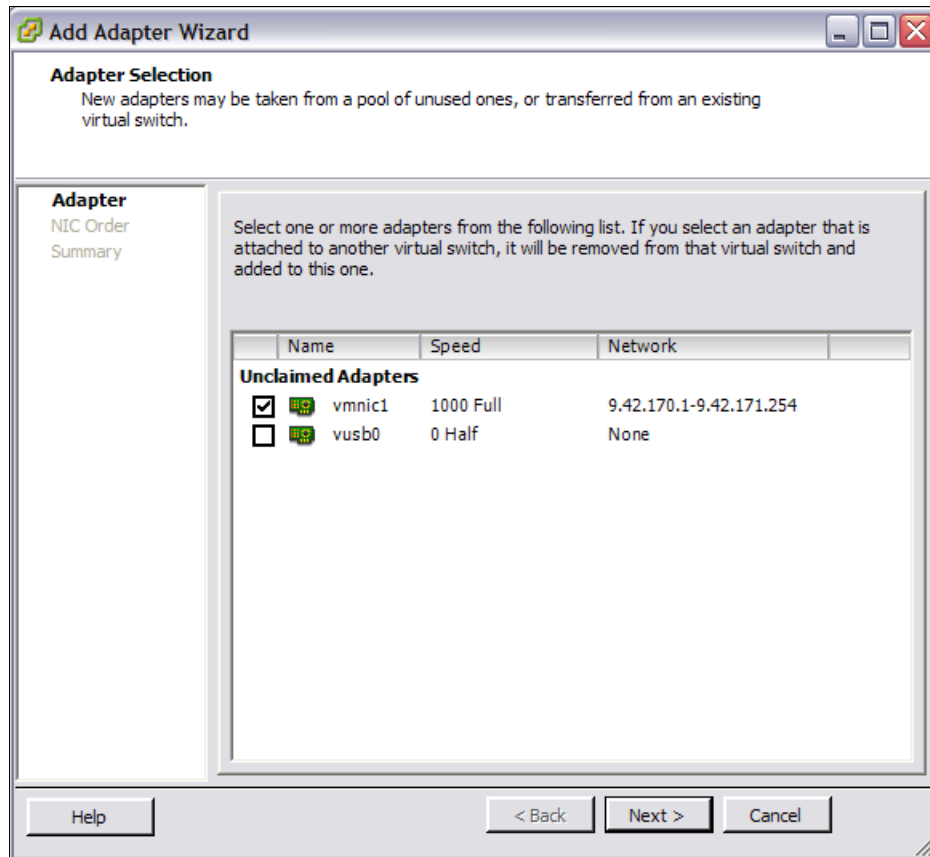


Figure 22-66 Adding second network adapter

4. You are presented with the main vSwitch0 Properties window. Highlight the first adapter, click **Edit**, and set the Configured Speed, Duplex to 100 Mbps, Full Duplex or 1000 Mbps, as shown in Figure 22-67 on page 699. Click **OK**.




Figure 22-67 Setting Network Speed

Note that there might be a refresh problem, and you will get extra LAN adapter entries. Close the window and open it again.

5. Perform these steps for the second adapter

## 22.4.6 Creating additional virtual switches for guests' connectivity

To create extra virtual switches to allow guests to connect, perform the following steps:

1. In the upper right corner, click **Add Networking** as shown in Figure 22-68 on page 700.

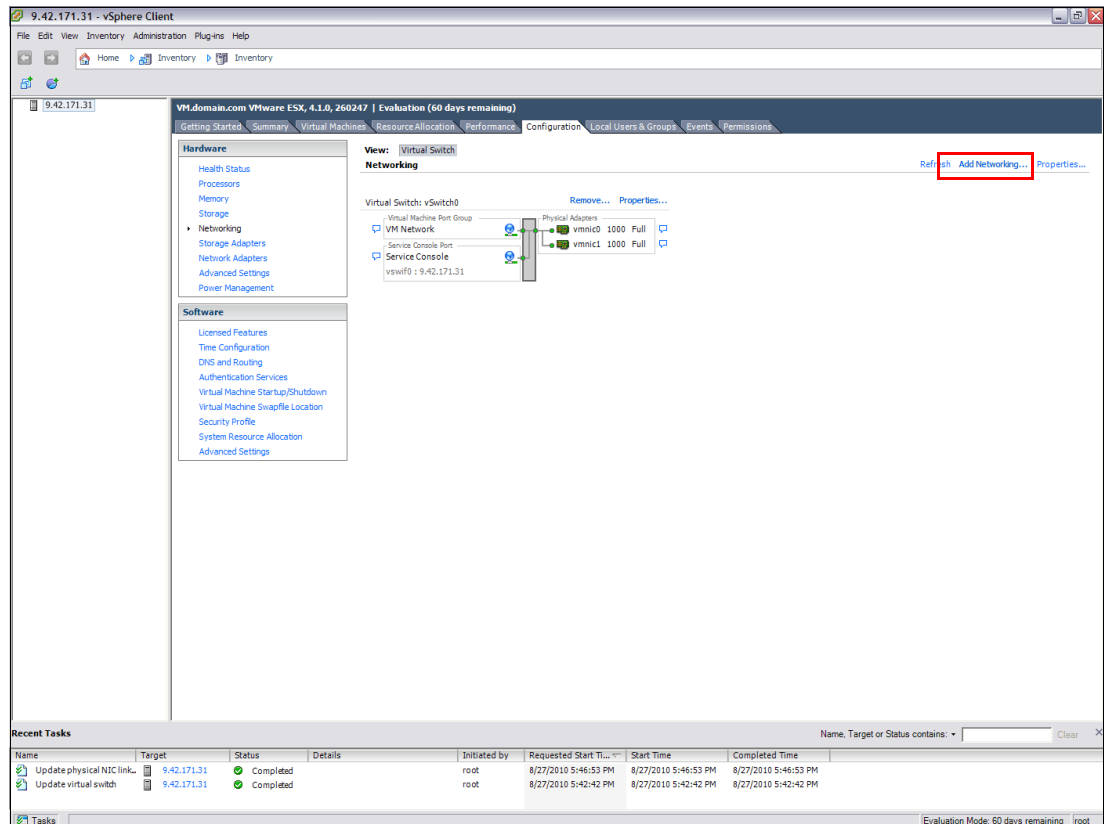


Figure 22-68 Add Networking

2. Select Virtual Machine as the Connection Type, and click **Next** as shown in Figure 22-69 on page 701.

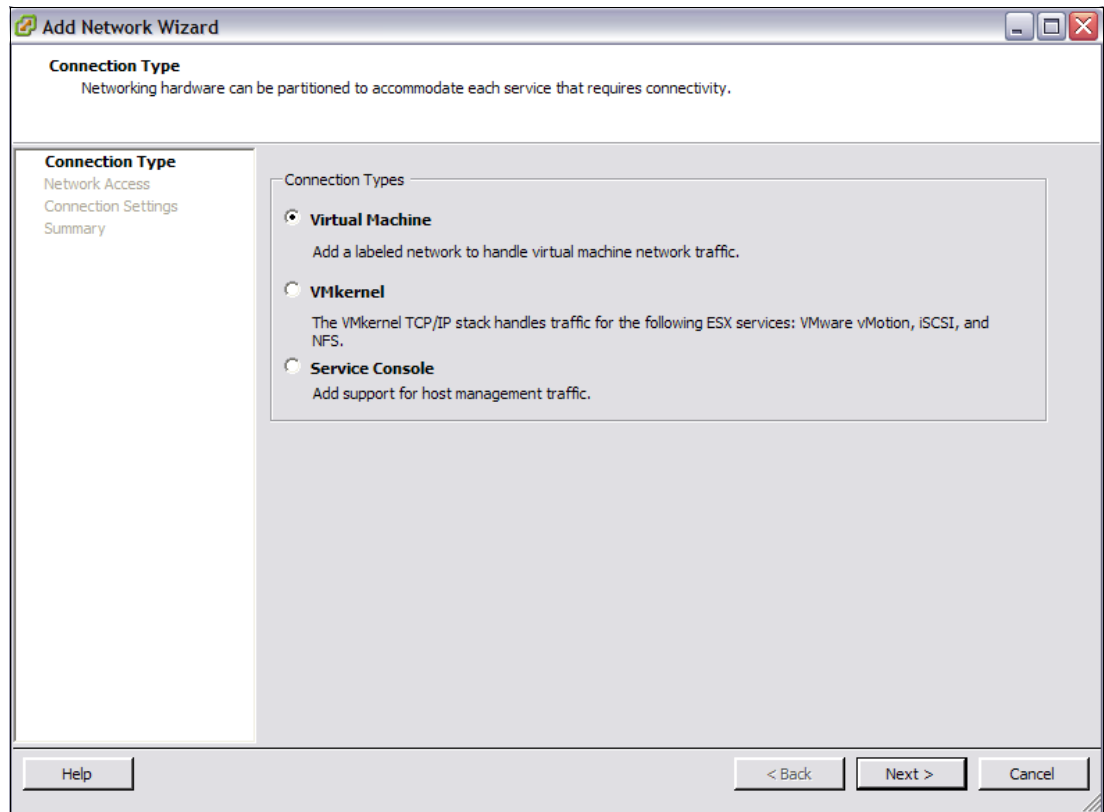


Figure 22-69 Connection Type

3. Pick one of the remaining LAN adapters (in this case vmnic0) as shown in Figure 22-70 on page 702. Click **Next**.




Figure 22-70 Selecting Network Access

- On the Network Label, type Network 0 as shown in Figure 22-71, and click **Next** and then **Finish**.

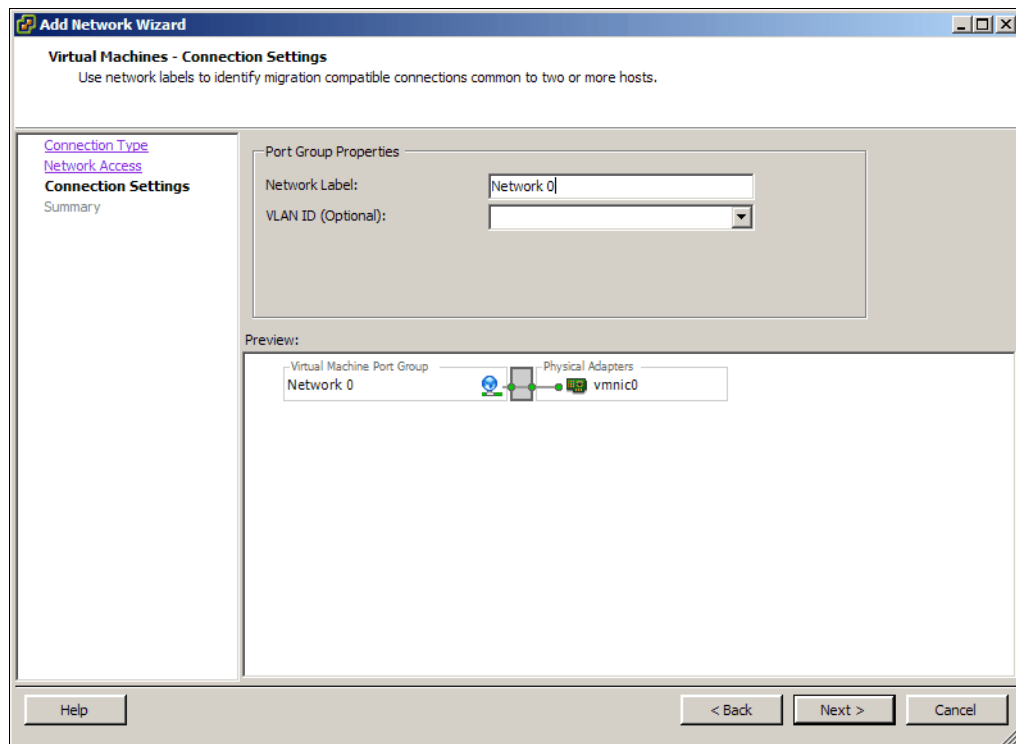


Figure 22-71 Virtual Network Label



5. Use the Properties option of each switch to configure the ports to 100 Mbps Full Duplex or 1000 Mbps.

## 22.4.7 Creating virtual machines

In this section, we discuss how to create virtual machines.

### Prerequisites

The prerequisites for creating virtual machines are:

- ▶ VMware VI ESX Host or vSphere Host with all current patches
- ▶ VMware vCenter Server access
- ▶ Available space on SAN or Local VMFS-3 volumes
- ▶ Resource availability in the Virtual Datacenter
- ▶ Administrative privilege on the VMware ESX Server host where the guest will be deployed
- ▶ OS License for the “to-be deployed” guest

In this section, we explain how to create a virtual machine. How you configure your virtual machines is dictated by your requirements (guest operating system, virtual hardware requirements, function, and so on).

We selected the creation of a virtual machine running Windows 2008 as an example.

To create a virtual machine:

1. Log into the vCenter Client with appropriate credentials.
2. Expand the left pane to locate the VM Host that the Guest will reside on. Right-click the Host, and choose **New Virtual Machine**. A wizard starts. See Figure 22-72 on page 704.

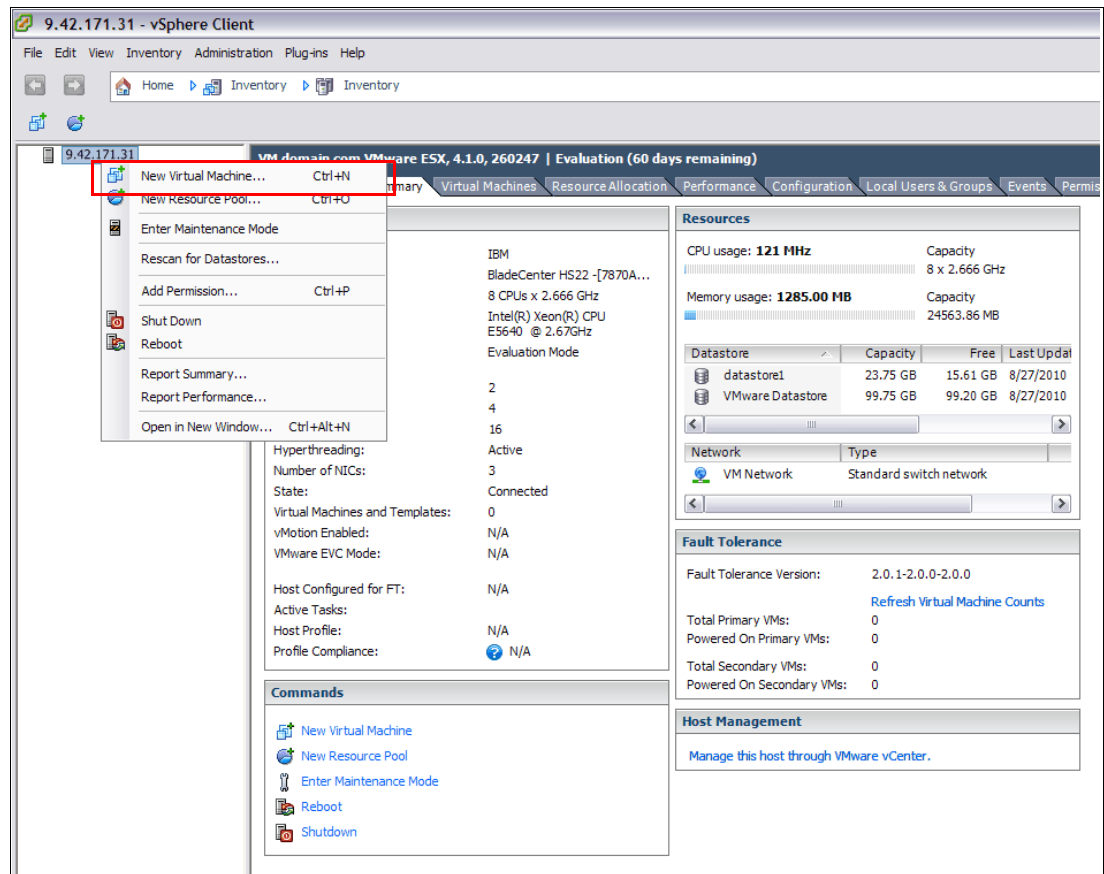


Figure 22-72 New Virtual Machine

3. In the Configuration window, select **Custom** as shown in Figure 22-73 on page 705. Click **Next**.

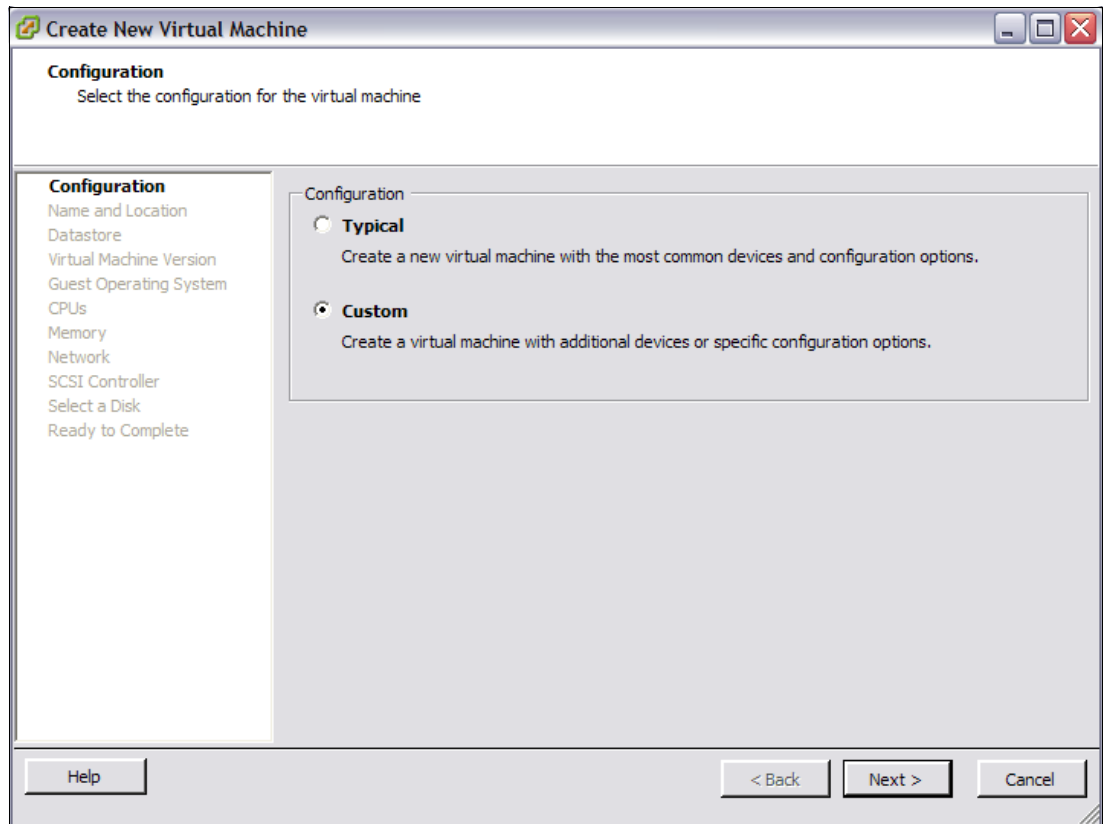


Figure 22-73 Configuration type

4. In the Virtual Machine Name field, type the name of the guest OS as it is registered in the DNS table (FQDN is not required) as shown in Figure 22-74 on page 706. Click **Next**.

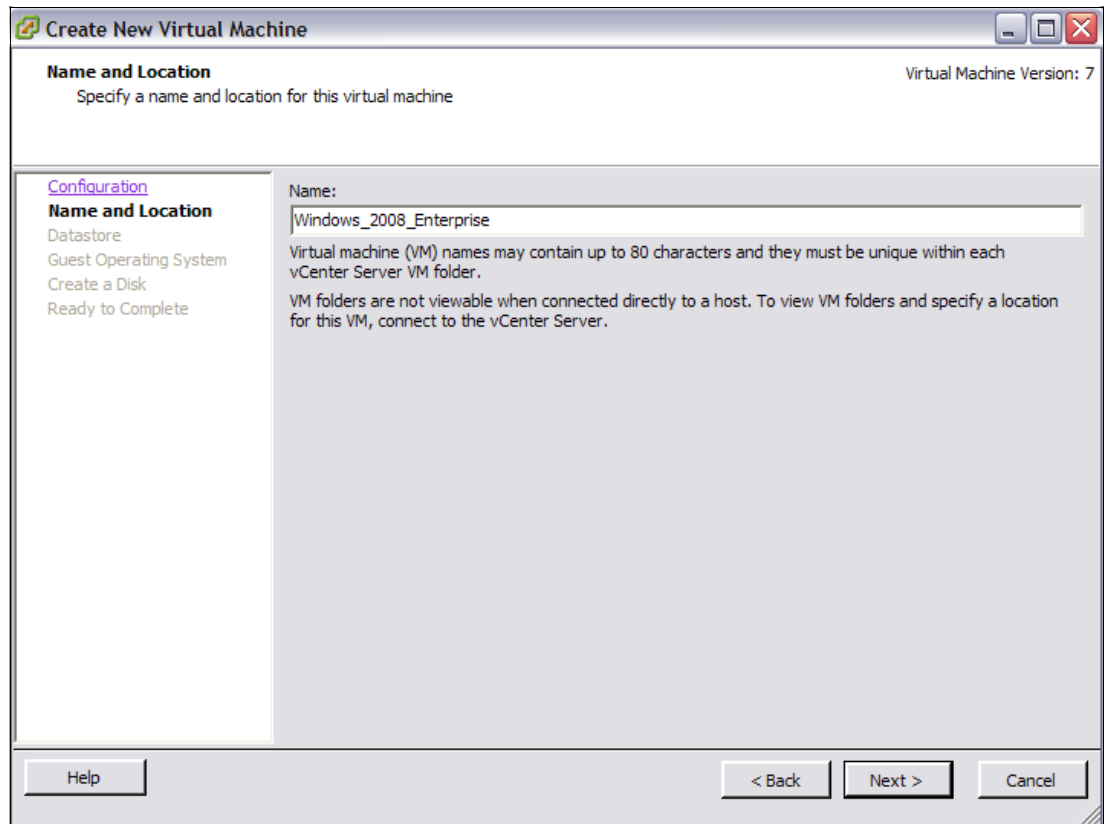


Figure 22-74 Virtual Machine Name

5. Select the Datastore VMFS partition where the Guest files reside (note that all configuration files including the disk file(s) are residing in that location), which can be the VM\_Datastore partition created in the earlier steps as shown in Figure 22-75 on page 707. Click **Next**.




Figure 22-75 Datastore VMFS Partition

6. Select the Virtual Machine Version based on your requirements, as shown in Figure 22-76 on page 708. Click **Next**.

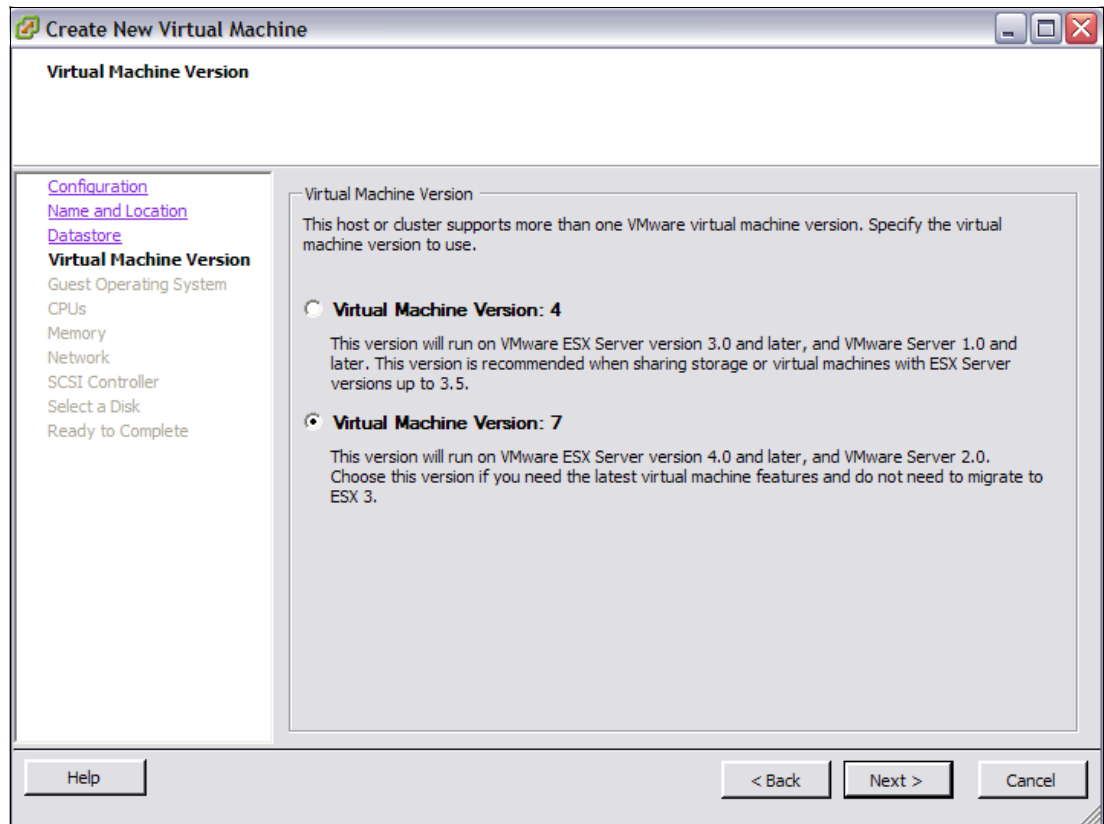


Figure 22-76 Virtual machine version

7. Select the Guest Operating System vendor and the corresponding Version. In our example, we used Microsoft Windows 2008 R2 64-bit as shown in Figure 22-77 on page 709. Click **Next**.

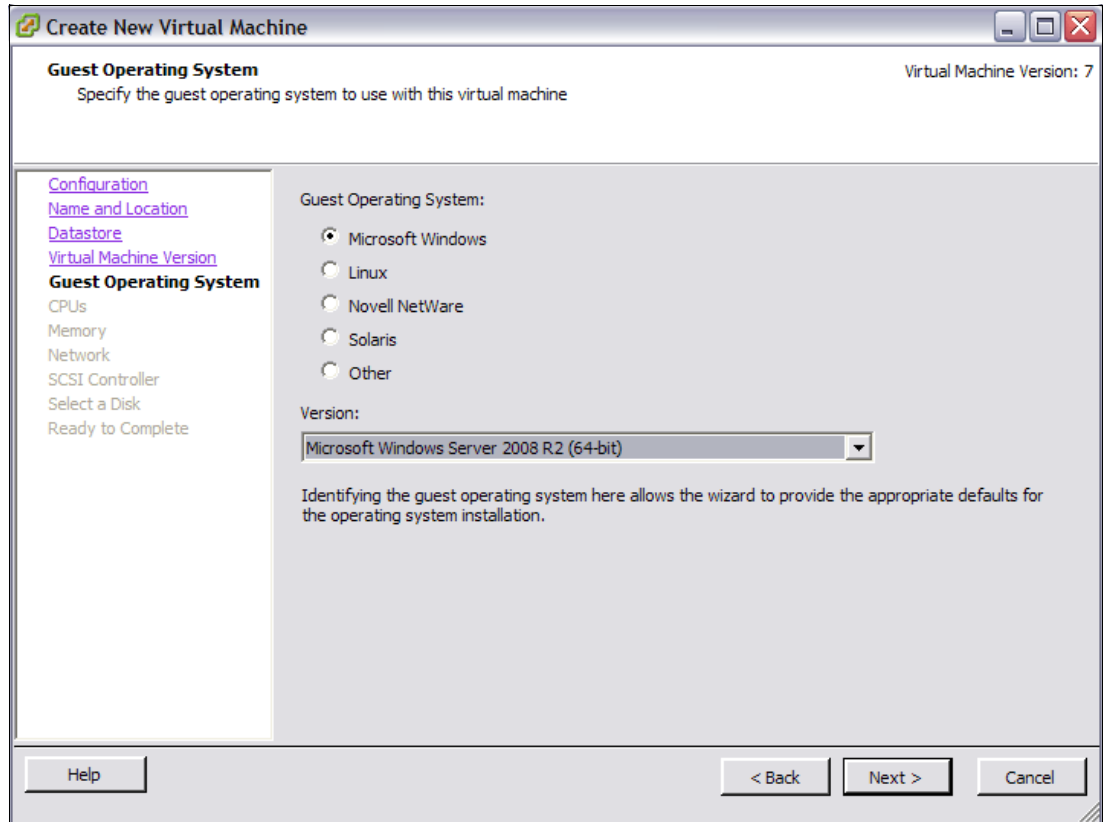


Figure 22-77 Guest Operating System selection

8. On the Number of vCPU, select **1** unless you are requiring Multi-CPU kernel (this has an impact on the overall performance of the guest) as shown in Figure 22-78 on page 710. Click **Next**.

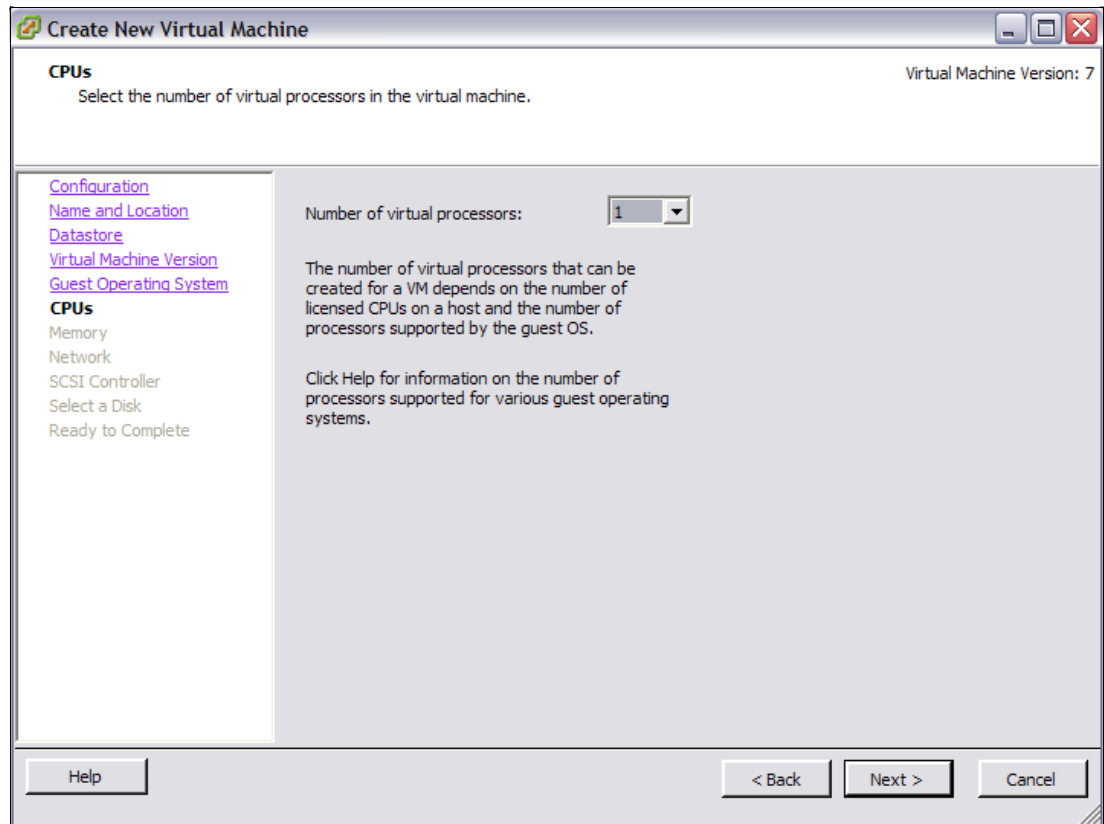


Figure 22-78 vCPU selection

9. In the Memory allocation window (Figure 22-79 on page 711), provide the guest with the necessary amount of memory as required. Click **Next**.



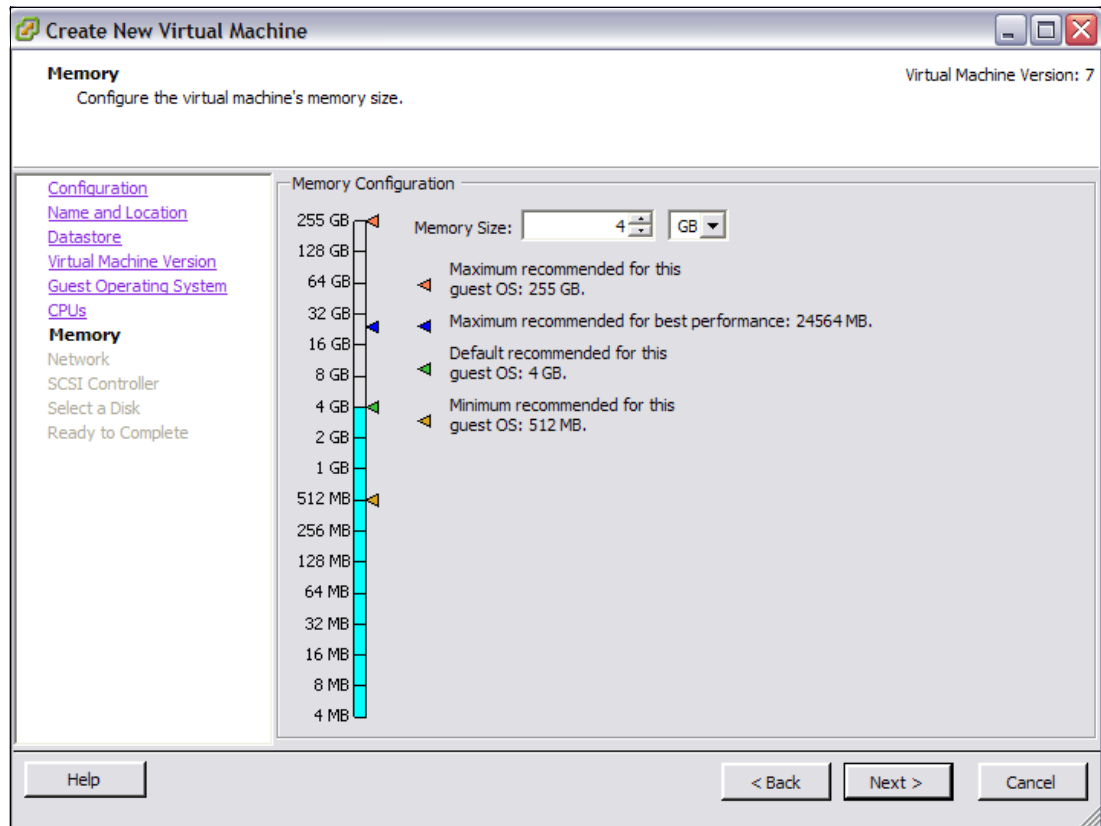


Figure 22-79 Memory allocation

10. In the Choose Networks window (Figure 22-80 on page 712), select the appropriate number of Network Adapters that the guest will operate with (Default is 1). Select the appropriate Network Label (ensure that the host is not overloading one particular network), which can be the VM network that we defined in earlier steps. Click **Next**.

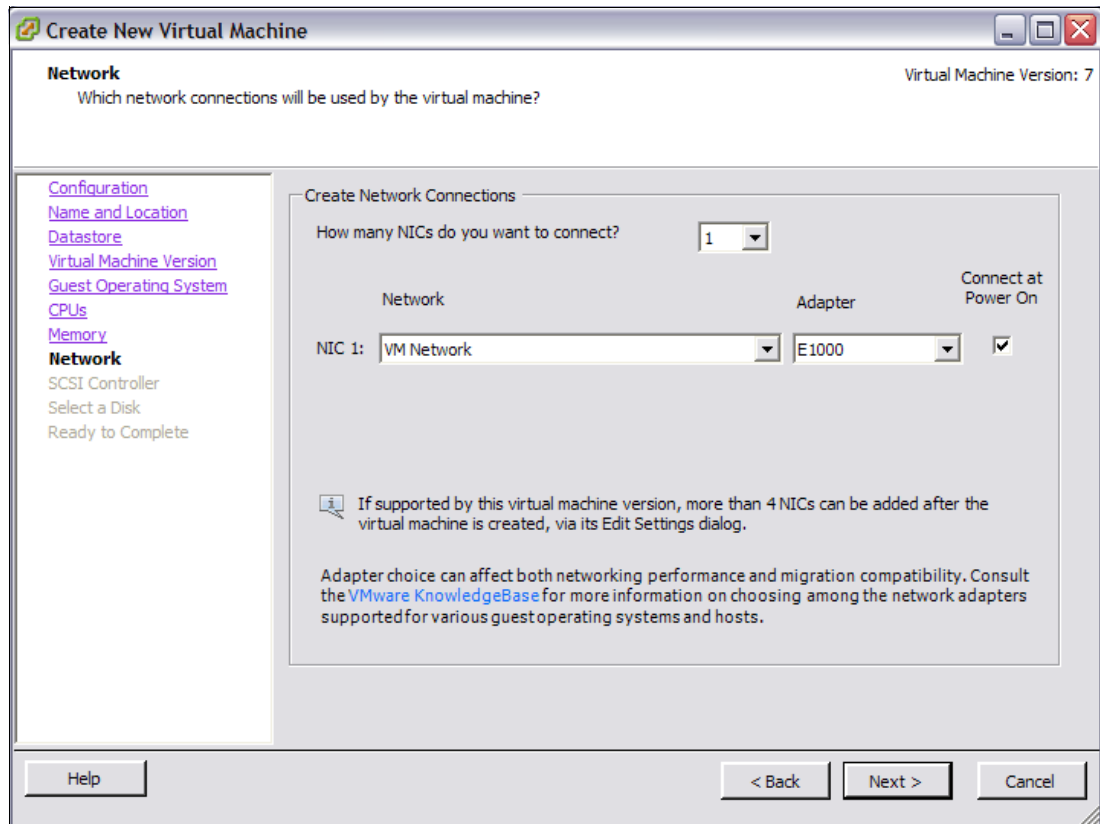


Figure 22-80 Network selection

11. In the SCSI Controller Types window (Figure 22-81 on page 713), select the controller based on the OS requirement. In our example, we select the LSI Logic SAS SCSI controller. Click **Next**.

For additional information about the types of SCSI controllers that are available, refer to the VMware Administration Guide and the Guest Operating System Installation Guide. You can find the guides at:

<http://www.vmware.com/support/pubs/>

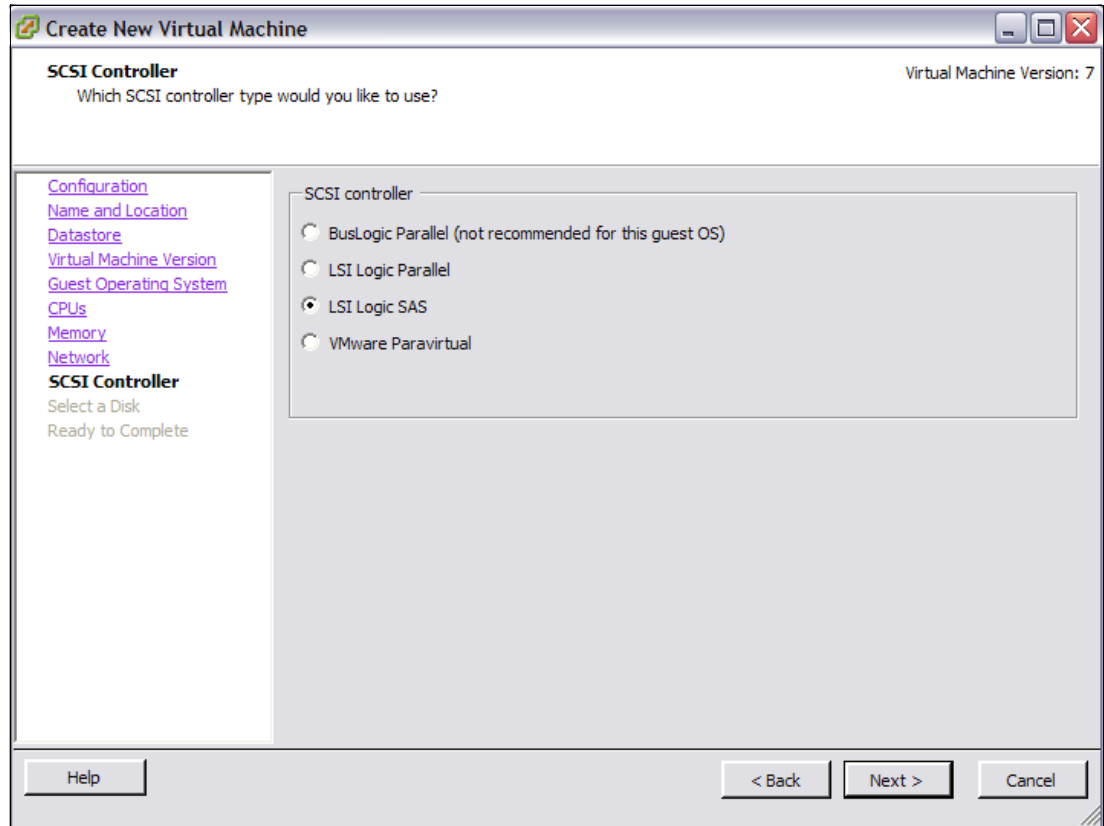


Figure 22-81 VM SCSI Controller type

12. In the Select a Disk window (Figure 22-82 on page 714), select:

- Create a new virtual disk: Use this option if there is no existing disk.
- Use an existing virtual disk: If you are connecting the guest to a previously built .vmdk file.
- Raw Device Mappings: Direct access to Fibre Channel SAN disks.

Click **Next**.

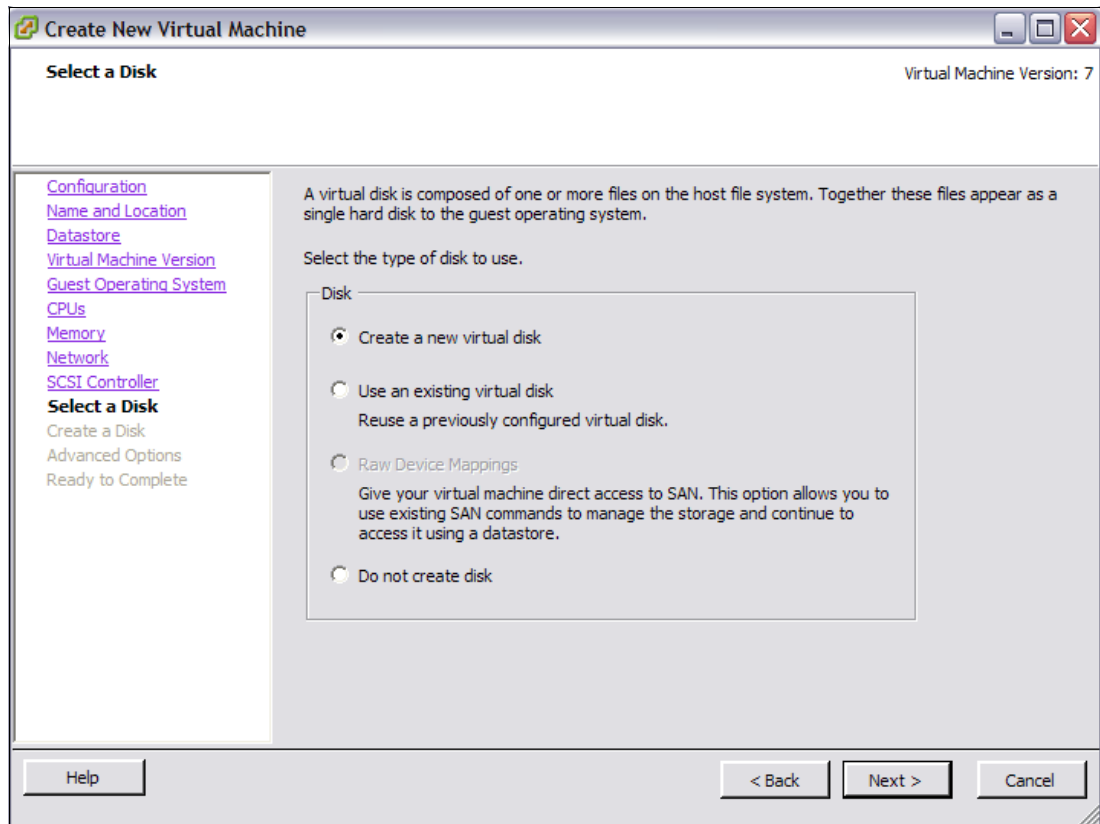


Figure 22-82 Select a disk

13. If you selected the Create a new virtual disk option, allocate the Disk size (this will be the size of the .vmdk file that represents the hardware disk in the virtual machine's configuration) as shown in Figure 22-83 on page 715. Specify the Datastore to be used by clicking **Browse**. In our example, we use the VM\_Datastore that we created in earlier steps. Click **Next**.

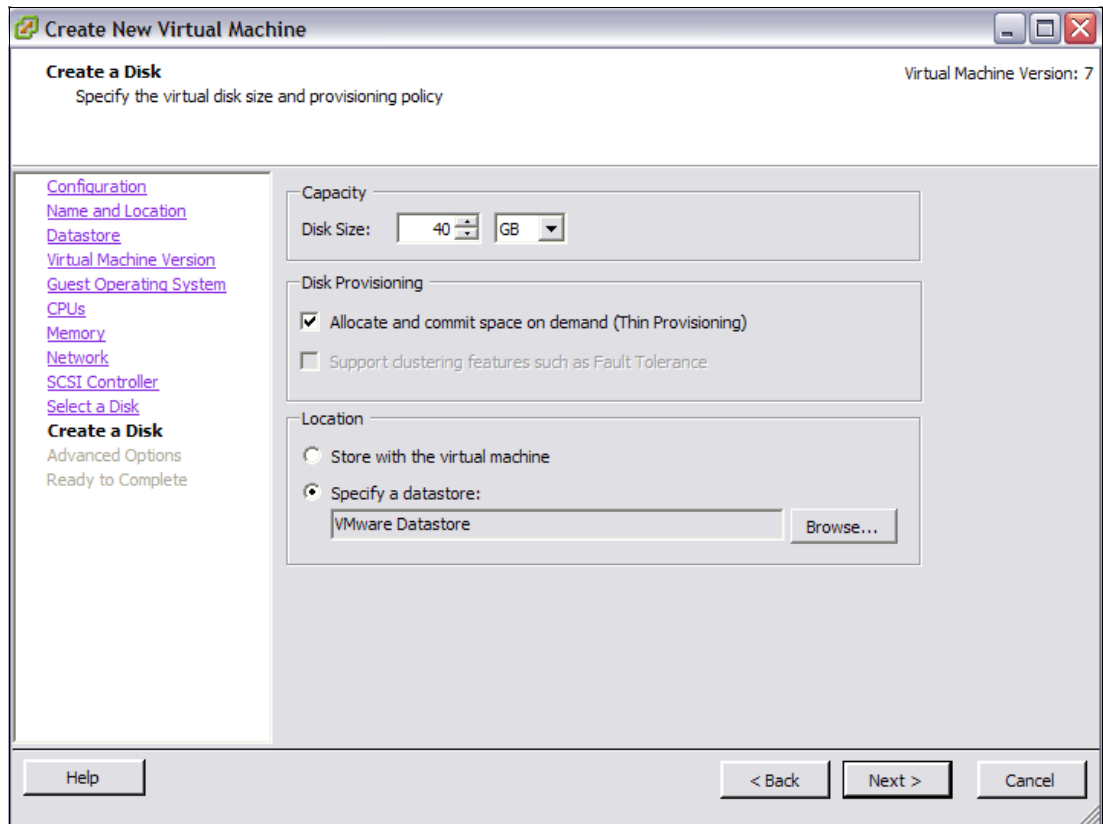


Figure 22-83 Create a disk

14. In the Specify Advanced Options window (Figure 22-84 on page 716), click **Next**.

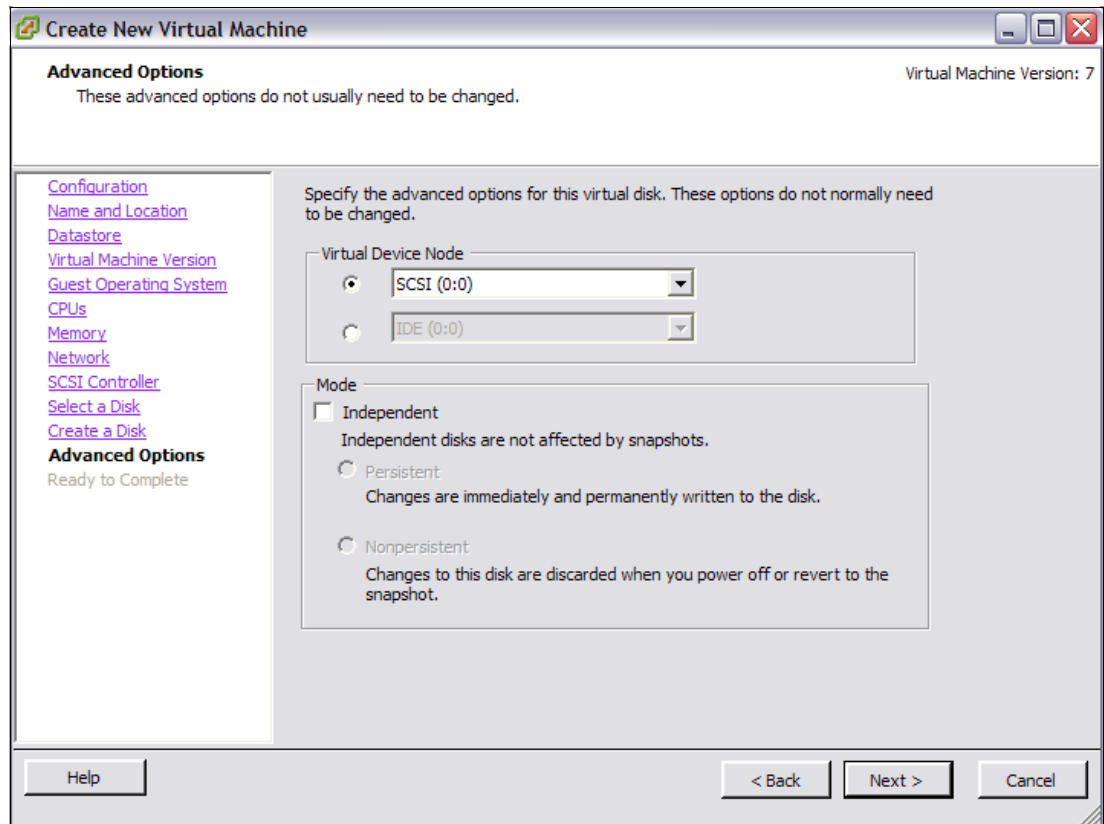


Figure 22-84 Advance Options

15. In the Summary panel (Figure 22-85 on page 717), click **Finish**. You can see the progress in the Recent Tasks pane at the bottom of the vCenter GUI.

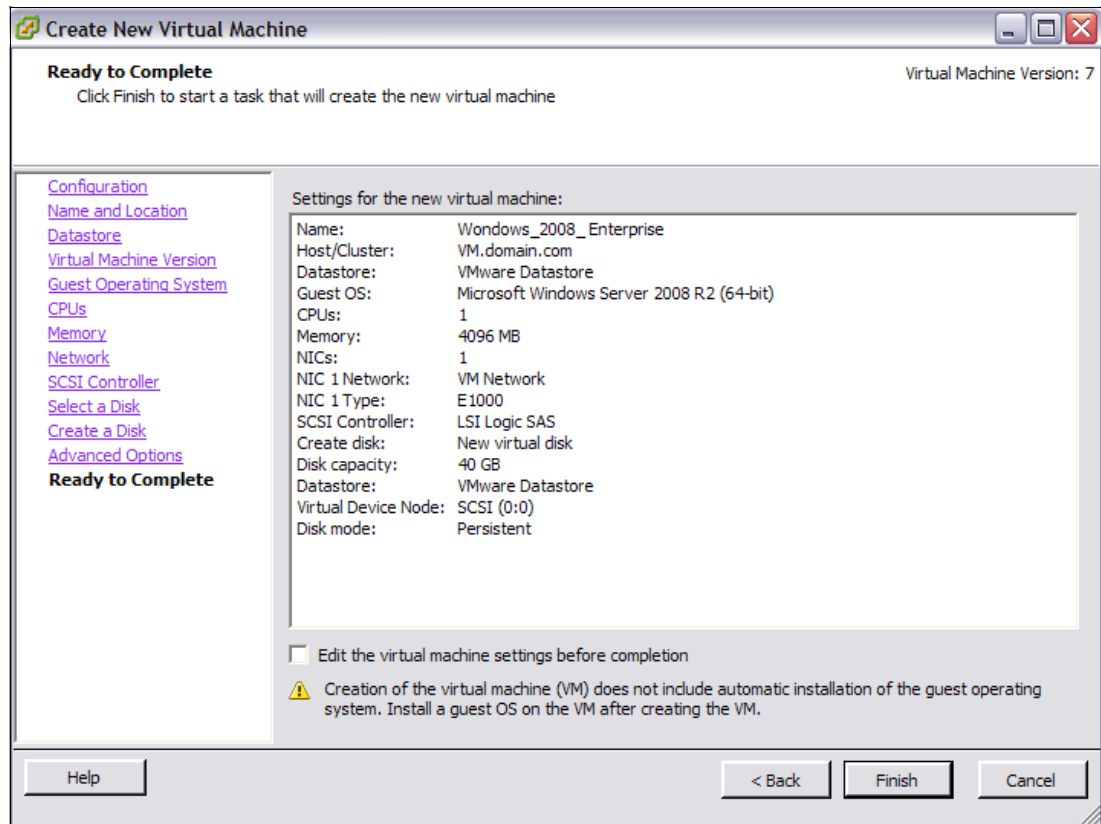


Figure 22-85 Summary screen

16. You are now ready to perform the Guest Operating System installation.

## 22.4.8 Additional VMware ESX Server Storage configuration

Refer to the following VMware documents for more information about ESX Server Storage configuration.

[http://www.vmware.com/pdf/vsphere4/r41/vsp\\_41\\_san\\_cfg.pdf](http://www.vmware.com/pdf/vsphere4/r41/vsp_41_san_cfg.pdf)

[http://www.vmware.com/pdf/vsphere4/r41/vsp\\_41\\_iscsi\\_san\\_cfg.pdf](http://www.vmware.com/pdf/vsphere4/r41/vsp_41_iscsi_san_cfg.pdf)

[http://www.vmware.com/pdf/vsphere4/r41/vsp\\_41\\_esx\\_server\\_config.pdf](http://www.vmware.com/pdf/vsphere4/r41/vsp_41_esx_server_config.pdf)







# A

## **IBM Support Portal website**

In this appendix, we explain how to use the IBM Support Portal web site to obtain downloads and support information for the IBM System Storage DS3500 storage subsystem.

## Sample navigation procedure

The following steps show, as an example, how to navigate the IBM Support Portal page to find software downloads for a DS3500 server.

1. In your web browser, go to this URL:

<http://www.ibm.com/support/entry/portal>

A window similar to Figure A-1 appears.

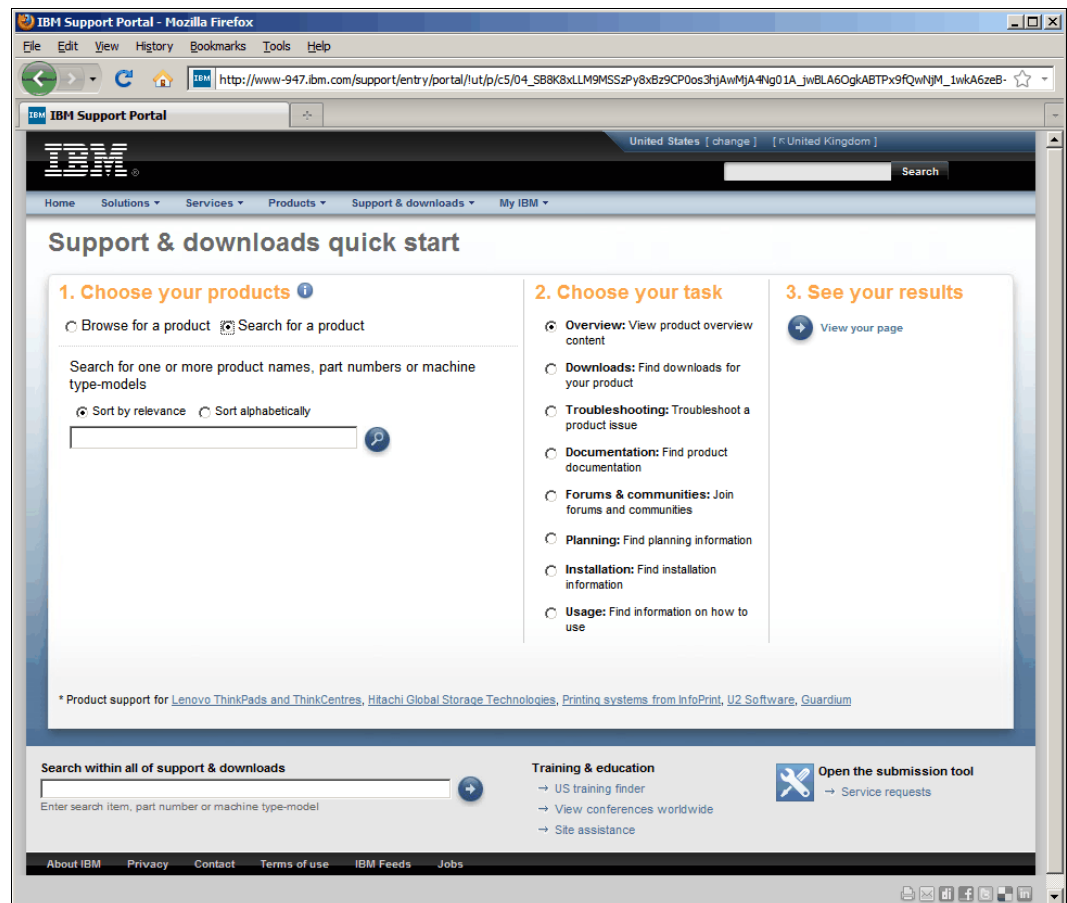


Figure A-1 IBM Support Portal

2. Type DS3500 in the search box and press Enter. A window similar to Figure A-2 on page 721 will appear.

**Support & downloads quick start**

**1. Choose your products**

☐ Browse for a product ☒ Search for a product

Search for one or more product names, part numbers or machine type-models

☒ Sort by relevance ☐ Sort alphabetically

ds3500

☐ DS3500 (DS3512, DS3524)

**2. Choose your task**

☒ Overview: View product overview content

☐ Downloads: Find downloads for your product

☐ Troubleshooting: Troubleshoot a product issue

☐ Documentation: Find product documentation

☐ Forums & communities: Join forums and communities

☐ Planning: Find planning information

☐ Installation: Find installation information

☐ Usage: Find information on how to use

**3. See your results**

[View your page](#)

\* Product support for [Lenovo ThinkPads and ThinkCentres](#), [Hitachi Global Storage Technologies](#), [Printing systems from InfoPrint](#), [U2 Software](#), [Guardium](#)

Figure A-2 Search for DS3500

3. Select the check box beside the **DS3500 (DS3512, DS3524)** and then click the **View your page** link as shown in Figure A-3.

**Support & downloads quick start**

**1. Choose your products**

Selected products (1)

☐ Browse for a product ☒ Search for a product

Search for one or more product names, part numbers or machine type-models

☒ Sort by relevance ☐ Sort alphabetically

ds3500

☒ DS3500 (DS3512, DS3524) | [Select OS](#)

**2. Choose your task**

☒ Overview: View product overview content

☐ Downloads: Find downloads for your product

☐ Troubleshooting: Troubleshoot a product issue

☐ Documentation: Find product documentation

☐ Forums & communities: Join forums and communities

☐ Planning: Find planning information

☐ Installation: Find installation information

☐ Usage: Find information on how to use

**3. See your results**

[View your page](#)

\* Product support for [Lenovo ThinkPads and ThinkCentres](#), [Hitachi Global Storage Technologies](#), [Printing systems from InfoPrint](#), [U2 Software](#), [Guardium](#)

Figure A-3 Select the DS3500

4. A support web site similar to the one shown in will be displayed. This is a customizable view, and you can add or remove information that is displayed.

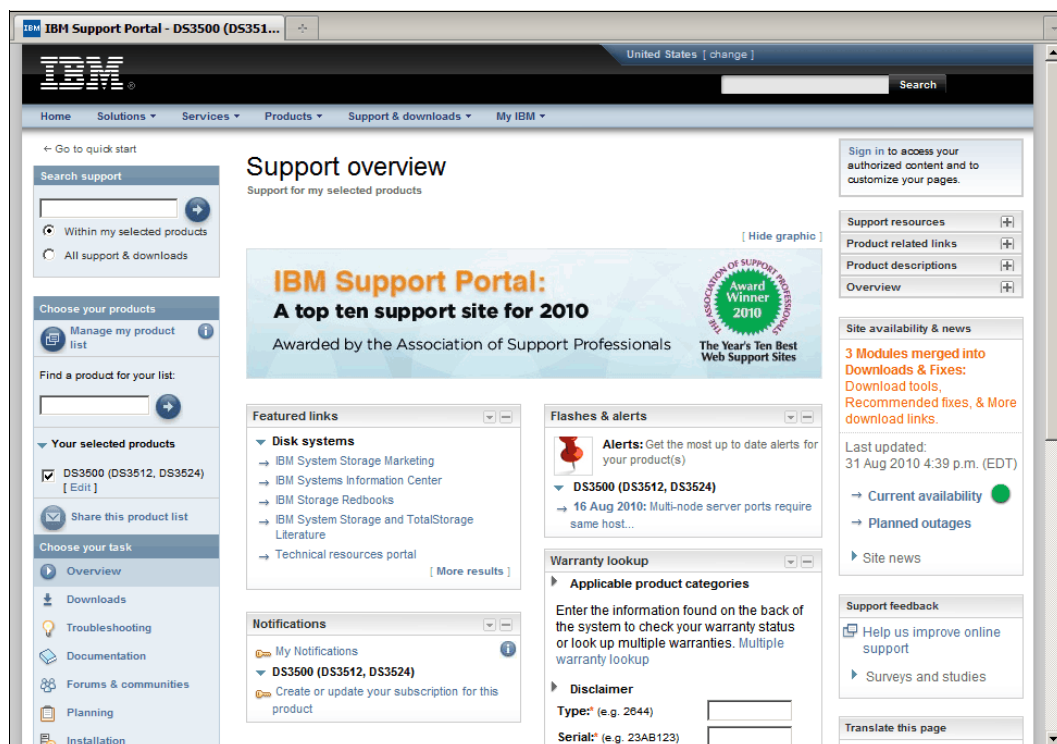


Figure A-4 IBM Support Portal DS3500

- Down the left side of this page, there is a **Choose your task** function as shown in Figure A-5. You use these links to select the support option that you want to access.




Figure A-5 Choose your task

# Download code updates

To download code updates, perform the following steps:

1. From the **Choose your task** list (Figure A-5 on page 722), click the **Downloads** link. A webpage similar to the one shown in Figure A-6 will be displayed.

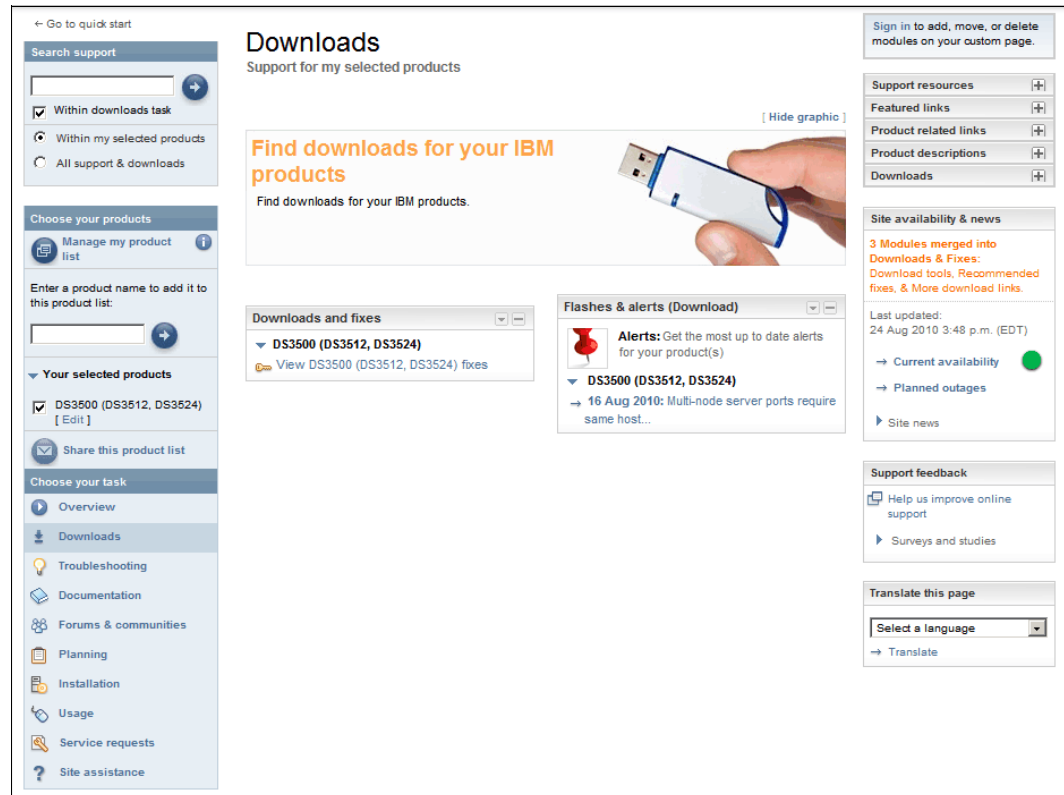


Figure A-6 Support Portal - Downloads page

2. Click the link called **View DS3500 (DS3512, DS3524) fixes**.
  - a. If you have an *IBM ID* and are already signed in to the IBM Support Portal, you will see the available downloads as shown in Figure A-7 on page 724. Here you can select which fixes and updates to download.

## Select fixes

Disk Systems, DS3500 (DS3512, DS3524) (All releases, All platforms)

**Select fixes category view**

The following results match your request. Select the fixes you wish to download.  
To try a different query, go to the [Identify fixes](#) page.

[Generate a link to this page](#)

Refine results:

All fixes

[Select all](#)
[Deselect all](#)
[Show fix details](#)

[Controller firmware](#)
[Others](#)

[ESM & HDD firmware](#)
[Storage Manager](#)

[Linux RDAC Driver](#)

**Controller firmware**

☐ **fix pack: Disk-CntlFW-June-2010-version-7.70.16.01** 15-Jun-2009  
 IBM DS Controller firmware version 7.70.16.01 code package

[Back to top](#)

**ESM & HDD firmware**

☐ **fix pack: Disk-HDD-June-2010-version-4.7** 15-Jun-2010  
 IBM DS3K HDD firmware package version 4.7

☐ **fix pack: Disk-ESM-June-2010-version-3.16** 15-Jun-2010  
 IBM DS ESM firmware package version 3.16

[Back to top](#)

**Linux RDAC Driver**

☐ **refresh pack: LnxRDAC\_DD-June-2010-version-09.03.0C05.0331** 15-Jun-2010  
 IBM DS Linux RDAC version 09.03.0C05.0331

Figure A-7 Select fixes

- b. If you do not have an IBM ID, or you are not currently signed in to the IBM Support Portal, you will see the website as shown in Figure A-8.

United States [change]

Home Solutions Services Products Support & downloads My IBM

My IBM registration  
Help and FAQ

## Sign in

Please enter your IBM ID and Password in the sign in area below. If you are not currently registered with our site please [register now](#).

Tip: To easily return to your destination, bookmark the page after this page. Do not bookmark the sign in page.

IBM ID:

Password:

[Submit](#)

**Additional resources**

- [Forgot your IBM ID?](#)
- [Forgot your password?](#)
- [Change password](#)

About IBM Privacy Contact Terms of use IBM Feeds Jobs

Figure A-8 IBM ID - Sign in

Here you have the option to Sign in to the IBM Support Portal using an existing IBM ID, or create a new IBM ID by clicking the **Register now** link. If you choose the **Register now** option, then you will be presented with the IBM registration window as shown in Figure A-9 on page 725.

IBM registration

Step 1 of 3

Asterisks (\*) indicate fields required to complete this transaction.

Email address \* (Will be your IBM ID)

Password \* (8 Characters min) Re-enter password \*

? Help with IDs and passwords?

Security question \*

Answer to security question \*

Country of residence \*

Language \*

Continue to the next page to complete your registration.

Continue Cancel

All information submitted is secure.

Figure A-9 IBM registration

- When you are signed in to the IBM Support Portal using your IBM ID, you will see the code download options as shown in Figure A-7 on page 724. Select the check box for each code update that you want to download. Then click the **Continue** button located at the bottom of the page (Figure A-10).

Select all Deselect all Hide fix details

Continue Back

Figure A-10 Download page buttons

- The Download Options website is displayed as shown in Figure A-11.

Select download options

Select the download method to be used to download fixes.

☒ Download using Download Director [What is this?](#)

☐ Download using your browser (HTTP)

**CAUTION: Do not assume that Fix Central will show you all the prerequisites you need.** Be sure to always click the [More information](#) link for additional prerequisite and other important fix information. Click [here](#) for an explanation of what prerequisites you can expect Fix Central to provide.

☒ Include prerequisites and co-requisite fixes (you can de-select the ones you do not need later)

Continue Back

Figure A-11 Support Portal - Download options

You have the option to download using your web browser (HTTP) or download using the Download Director, which is the default selection. The Download Director is a Java applet that provides enhanced download capability and reliable downloads with pause and resume capability. It starts in a separate browser window and provides a graphical user interface (GUI) to specify the download location, and view download status.

Click the **What is this?** link to learn more details about the Download Director.

5. If you choose to download using Download Director, you will see the Download Director window as shown in Figure A-12.

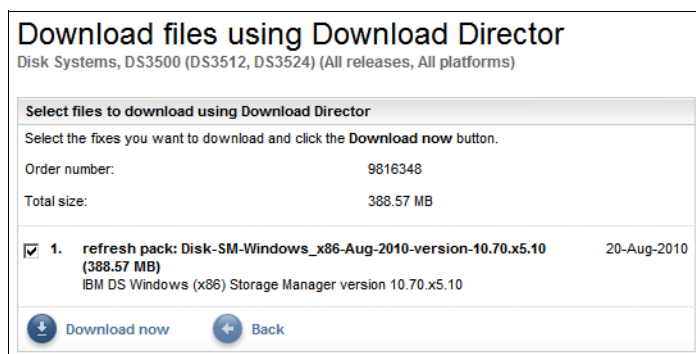


Figure A-12 Download files using Download Director

- a. Ensure the check box beside the code package is selected and then click the **Download now** button.

The first time that you use the Download Director, a Java applet will download and install.

- b. You will then be asked to specify a default location for all downloads using the Download Director as shown in Figure A-13. Click the **Browse** button to change the folder location. Click **OK** when you have chosen the download location.




Figure A-13 Download Director - Initial Setup

- c. The Download Director will then download the selected files as shown in Figure A-14.




Figure A-14 Download Director - download progress

6. If you choose to download using your browser (HTTP), a website similar to the one shown in Figure A-15 on page 727 displays.






Figure A-15 Download files using HTTP

When using this option, you will need to click each link to save the file to your local disk.

## My notifications

The IBM My notifications web site can be used to keep you informed of the latest firmware updates and other important product updates. The following steps show, as an example, how to set up notifications to receive product updates for the DS3500 System Storage subsystem.

1. From the main support portal window as shown in Figure A-16, click the **My notifications** link in the Stay informed pane on the right side of the support website.

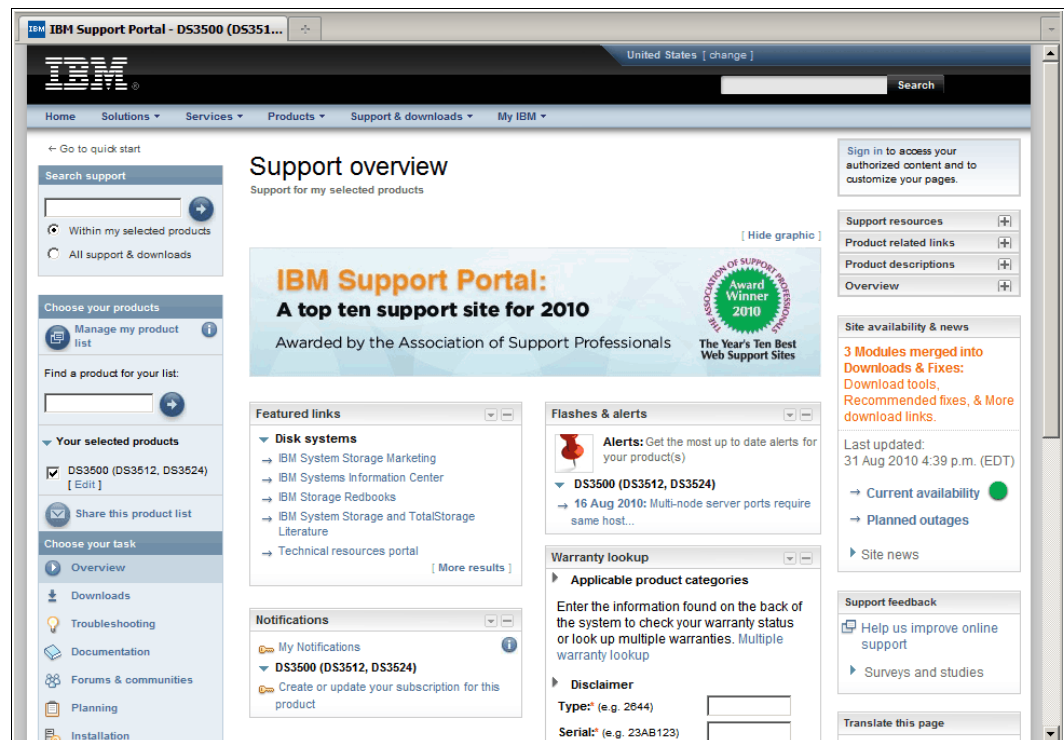


Figure A-16 IBM Support site - Overview

- The Sign In window appears as shown in Figure A-8 on page 724. Enter your IBM ID and password in the sign in area to sign in and proceed. If you are not currently registered with the website, click **Register Now**.
- The My notifications main page opens as shown in Figure A-17.




Figure A-17 My notifications

- Select the **Subscribe** tab, and you will see the list of available product types as shown in Figure A-18.

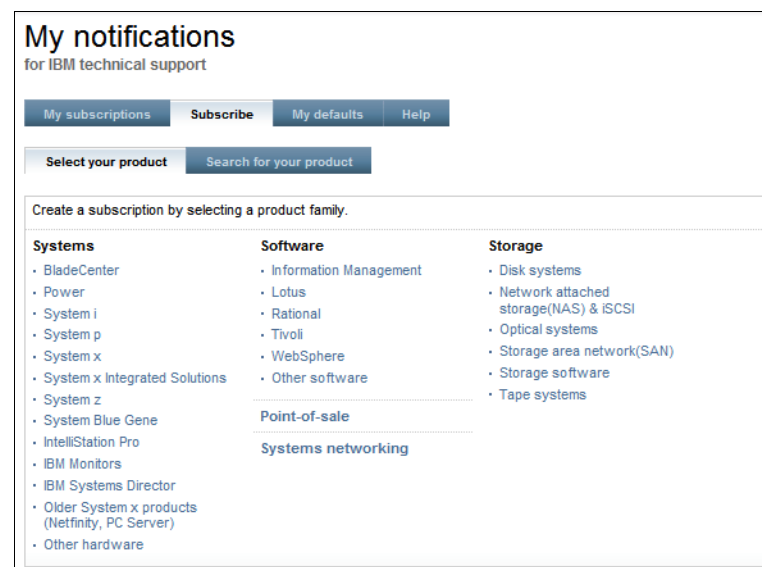


Figure A-18 My notifications - subscribe

5. Click Disk Systems and you will be presented with a list of the IBM Disk Storage Systems as shown in Figure A-19.

The screenshot shows the IBM My notifications - navigation page. The page has a navigation bar with links: Home, Solutions, Services, Products, Support & downloads, and My IBM. Below the navigation bar, there is a search bar and a welcome message. The main content area is titled "My notifications for IBM technical support". On the left, there is a sidebar with links: Return to IBM Support Portal, IBM Support Portal, Subscriptions, Help, and Feedback. The main content area shows a section titled "You are subscribing to the following Disk systems". Below this, there is a "Continue" button. A table lists various IBM Disk Storage Systems with checkboxes for selection. The table has two columns: Product and Name. The products listed are: All products, 2104 Expandable Storage Plus, 2105 Enterprise Storage Server, 2106 Modular Storage Server, 7133 Serial Disk System, BladeCenter Boot Disk System, DR450, DS300, DS3200, DS3300, DS4400 Midrange Disk System, DS4500 Midrange Disk System, DS4700 Express, DS4800 Midrange Disk System, DS5020, DS5100, DS5300, DS6800, DS8100, and DS8300.

| Product *                                               |                                                      |
|---------------------------------------------------------|------------------------------------------------------|
| <input type="checkbox"/> All products                   | <input type="checkbox"/> DS4400 Midrange Disk System |
| <input type="checkbox"/> 2104 Expandable Storage Plus   | <input type="checkbox"/> DS4500 Midrange Disk System |
| <input type="checkbox"/> 2105 Enterprise Storage Server | <input type="checkbox"/> DS4700 Express              |
| <input type="checkbox"/> 2106 Modular Storage Server    | <input type="checkbox"/> DS4800 Midrange Disk System |
| <input type="checkbox"/> 7133 Serial Disk System        | <input type="checkbox"/> DS5020                      |
| <input type="checkbox"/> BladeCenter Boot Disk System   | <input type="checkbox"/> DS5100                      |
| <input type="checkbox"/> DR450                          | <input type="checkbox"/> DS5300                      |
| <input type="checkbox"/> DS300                          | <input type="checkbox"/> DS6800                      |
| <input type="checkbox"/> DS3200                         | <input type="checkbox"/> DS8100                      |
| <input type="checkbox"/> DS3300                         | <input type="checkbox"/> DS8300                      |

Figure A-19 My notifications - Disk systems

6. Select the check box beside all the Disk systems that want to receive notifications for, and then click the **Continue** button. You will see the website as shown in Figure A-20 on page 730. Here you can manage when and how you receive your product notifications.

**Note:** At the time of writing, it was not possible to select to receive notifications for the DS3500 and EXP3500 products. Alternative IBM midrange disk systems have been used in this example to illustrate the procedure.

In this example, we have chosen "Midrange disk systems" as the name of the folder where the notifications will be saved.

After you have configured your subscriptions, click the **Submit** button at the bottom of the page.

**My notifications - subscription**

Country/region [ select ]

Home Solutions Services Products Support & downloads My IBM

Welcome username [Not you?] [ IBM Sign out ]

← Return to IBM Support Portal

**IBM Support Portal**

Subscriptions

Help

Feedback

Related links

- Lenovo support
- InfoPrint support

## My notifications for IBM technical support

My subscriptions Subscribe My defaults Help

You are subscribing to the following

**Disk systems**

DS5020, EXP3000, DS3400, DS3200, DS3300, BladeCenter Boot Disk System, DS5100, DS5300, DS3950

Fields marked with an asterisk (\*) are required.

**Options**

Name:

Save in existing or new folder:

Existing:

New:

**Notify me by**

☒ e-mail

☐ daily e-mail ☐ weekly e-mail

☒ plain text e-mail ☐ html e-mail

☐ delivery to this folder

☐ delivery via syndication feed (RSS,Atom)

[what's this?](#)

About IBM Privacy Contact Terms of use Accessibility IBM Feeds Jobs

Figure A-20 My notifications - setup

- You will see the window shown in Figure A-21 when you have completed the My notifications setup.

**My notifications - action successful**

Country/region [ select ]

Home Solutions Services Products Support & downloads My IBM

Welcome username [Not you?] [ IBM Sign out ]

← Return to IBM Support Portal

**IBM Support Portal**

Subscriptions

Help

Feedback

Related links

- Lenovo support
- InfoPrint support

## Subscription action was successful

The subscription has been successfully created

**Please select an option:**

- [Return to IBM Support Portal](#)
- [View all my subscription folders](#)

Figure A-21 My notifications - Successful

# 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

For information about ordering these publications, see “How to get Redbooks” on page 732. Note that some of the documents referenced here might be available in softcopy only.

- ▶ *IBM System Storage DS Storage Manager Copy Services Guide*, SG24-7822
- ▶ *IBM System Storage DS4000 and Storage Manager V10.30*, SG24-7010
- ▶ *IBM Midrange System Storage Hardware Guide*, SG24-7676
- ▶ *IBM System Storage DS3000: Introduction and Implementation Guide*, SG24-7065
- ▶ *IBM Midrange System Storage Implementation and Best Practices Guide*, SG24-6363
- ▶ *Implementing an IBM/Brocade SAN with 8 Gbps Directors and Switches*, SG24-6116
- ▶ *Implementing an IBM/Cisco SAN*, SG24-7545

## Other publications

These publications are also relevant as further information sources:

- ▶ *IBM System Storage DS Storage Manager Version 10 - Installation and Host Support Guide*
- ▶ *IBM System Storage DS3000, DS4000, and DS5000 Command Line Interface and Script Commands Programming Guide*
- ▶ *Copy Services User's Guide - IBM System Storage DS Storage Manager v10.70*
- ▶ *IBM System Storage DS3000, DS4000, and DS5000 - Hard Disk Drive and Storage Expansion Enclosure Installation and Migration Guide*
- ▶ *System Storage DS3500 and EXP3500 Rack Installation and Quick Start Guide*
- ▶ *System Storage DS3500 and EXP3500 Installation, User's, and Maintenance Guide*
- ▶ *System Storage DS3000 Storage Manager Version 10 Installation and Support Guide for Microsoft Windows Server, Linux, Novell NetWare, and VMware ESX Server*

All the above publications are available for download from the IBM support web site.

## Online resources

These websites are also relevant as further information sources:

- ▶ System Storage Interoperation Center (SSIC):  
<http://www.ibm.com/systems/support/storage/ssic/>

- ▶ Support for IBM Disk systems:  
<http://www.ibm.com/systems/support/storage/disk>
- ▶ IBM DS3000, DS4000, DS5000, and BladeCenter Boot Disk System Premium Feature Activation  
<http://www.ibm.com/storage/fasttkeys>
- ▶ IBM Glossary of Computing Terms  
<http://www.ibm.com/ibm/terminology>

## How to get Redbooks

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, and order hardcopy Redbooks publications, at this website:

[ibm.com/redbooks](http://ibm.com/redbooks)

## Help from IBM

IBM Support and downloads

[ibm.com/support](http://ibm.com/support)

IBM Global Services

[ibm.com/services](http://ibm.com/services)

# Index

## A

- About IBM System Storage DS 331
- AC power LED 16
- access logical drive 23–24, 350
- access LUN 288
- Activate Controller Firmware 364
- activate RSM 507
- activate the firmware 360
- activate the pending firmware 364
- activating RSM 507
- active zoneset 547
- Adapter Port Name 610
- Adapter Settings 654
- Add Free Capacity 230
- Add Storage Subsystem 272
- Adding free capacity 239
- Adding storage subsystem 527
- Additional Mapping 286, 291
- ADT 83
- ADT See *Auto-Logical Drive Transfer (ADT)*
- ADT/AVT 281
- ADT-disabled failover 83
- ADT-enabled failover 84
- Advance Management Module (AMM) 544
  - console 592
- Advanced 209
- Advanced Host Port Settings 320
- Advanced Recovery menu 411
- Advanced Technology Attachment (ATA) 5
- age of the battery 194, 307
- AIX 6.1 625
- AIX host attachment 626
- AIX Host Kit 626
- alert settings 528
- Alerts Status notifications 482
- alt\_disk\_install command 637
- AMM See *Advance Management Module (AMM)*
- Arbitrated loop 2
- Array Import wizard 400
- Array segment size 303
- Arrays 219, 221, 226, 338
- Arrays & Logical Drives pane 184
- Associated Logical Elements 206
- Associated Physical Components 206
- asynchronous copy 21
- asynchronous with write order consistency group 21
- ATA See *Advanced Technology Attachment (ATA)*
- Auto Volume Transfer function (AVT) 595, 634
- Auto-Logical Drive Transfer (ADT) 83
- automatic assignment 61
- Automatic configuration 302
- automatic configuration 192
- automatic discovery 151, 271
- automatic hot spare assignment 264

- Automatic mode 224
- Automatic Support Data Collection 334
- AVT See *Auto Volume Transfer function (AVT)*

## B

- Baseline mode 323
- Basic Zone Permission 548
- Battery Settings 198
- bidirectional authentication 312, 314
- BIOS configuration 611
- BladeCenter Vital Product Data 541
- blob 444
- block device 618
- boot order 591
- BOOTP/DHCP service 270

## C

- cable routing 43
- cache 70
  - block size 69, 241
  - flush 70
  - flushing levels 241
  - memory 69
  - read ahead multiplier 234
  - read-ahead 243
  - settings 198–199, 219, 233, 238, 241
- capacity 181, 238
- Capture State Information 336
- cascaded loop configuration 19
- CASD See *Collect All Support Data (CASD)*
- Cat 5E 114
- Cat 6 114
- cfgmgr 634
- Challenge Handshake Authentication Protocol (CHAP) 12, 311, 313
  - Secret 313
- Change I/O Type 303
- Change Mapping 295
- Change Operating System 294
- Change Security Key 456
- Changing the RAID level 239
- changing the security key 456
- changing the segment size 239
- CHAP See *Challenge Handshake Authentication Protocol (CHAP)*
- Check Array Redundancy 432
- chpath command 636
- Clear Array Configuration 417
- Clear Configuration 416
- Clear Controller Firmware 365
- Clear Error Counts 404
- Clear Pending Firmware Progress 366
- Clear Registrations/Reservations 403
- Clear Storage Subsystem 416

- CLI See *command-line interface (CLI)*
- cluster installation 570
- cluster node 281
- Collect All Support Data (CASD) 179, 331, 341
  - log 331
- Collect Drive Data 335
- collect state capture 336
- command line interface (CLI) 354
- Command queuing 49
- Command Reference 522
- Command-Line Interface 519
- command-line interface (CLI) 27, 519
  - online help 534
- communication error 180
- Company Information 493
- Component Location 386, 388
- Component Type 386, 388
- Compute Change Log Files 346
- Configure Alerts 158, 485, 506
- Configure Cluster Quorum Wizard 574
- Configure e-mail alerts 524
- Configure hot spare drives 305
- Configure iSCSI Ports 273
- configure logical drives 301
- Configured Hosts 182
- connection lost 180
- Contact Information 492
- Contacting Storage Subsystem 160
- contiguous data blocks 431
- controller 208, 338
  - diagnostics 405, 421
  - firmware 329, 349, 373–374
  - ownership 219, 227, 234, 238
  - time of day 194, 307
- Controller NVSRAM 378
- controllers' clocks 386
- Copy Manager 214, 221
- Copy Services 254
- copying back from a hot spare drive 239
- correct syntax 520
- create a cluster 570
- create a script 534
- create an array 222
- Create Array 223, 261
- Create Arrays and Logical Drives 305
- Create Copy 254
- Create FlashCopy Logical Drive 254
- Create Logical Drive 233, 261
- Create Remote Logical Drive Mirror 254
- Create Security Key 448
- critical event 335, 386, 474, 480
- cryptographic erase 459
- current password 300
- current versions 358
- custom listening port 320
- customize settings 233

## D

- Daily Status e-mails 483
- data loopback 406

- data scrubbing 68
- data security 440
- data striping 50
- database 234
- Datastore 648, 693
  - VMFS partition 706
- Date/Time 386
- daughter card 17
- DC power LED 16
- DCE 230
- DDK See *Driver Development Kit (DDK)*
- decrypt data 441
- default frame size 320
- Default Group 284, 286, 288–289
- Default Host Type 198, 200
- default IP addresses 271
- default subnet mask 319
- Define Additional Mapping 291
- define host 277
- Define Host Group 290
- Define Host wizard 289
- Defined Mapping 288, 295
- defined storage subsystems 528
- Defragment Array function 429
- defragmenting an array 239
- Degraded 404
- deleting Logical Drive 254
- description of the event 386–387
- device mapper 616
- Device Mapper Multipath (DMM) 588, 595, 616–617
  - for Linux 83
  - multipath driver 619
- Device Specific Module (DSM) 80, 311, 557
- DHCP server 317
- DHCP/BOOTP server 318
- diagnostic event log 212
- direct-attached configuration 41
- disable a feature 196
- Disabled 487
- disallow un-named discovery sessions 317
- disk drive channels 338
- Disk Encryption Manager 442
- disk firmware 349
- Disk Management 550, 563
- disk mirroring 51
- disk mirroring with striping 51
- Disk Security enabled 461
- Disk Security premium feature 439
- DMM See *Device Mapper Multipath (DMM)*
- DMP See *Dynamic Multi-pathing (DMP)*
- DNS 311
  - Host name 666
- download controller firmware 369, 373
- download controller NVSRAM 369
- Download Director 725
- download drive firmware 369, 381
- download ESM configuration settings 369
- download ESM firmware 369
- Download Firmware 330, 354, 357, 362, 368
- Download NVSRAM 367



- Download NVSRAM file with firmware 358
- Drive Channel Status Details 403
- Drive Channels 403
- drive firmware 354–355, 381
- Drive menu 209
- Drive Security premium feature key 445
- drive-level encryption 440
- Driver Development Kit (DDK) 80
- drives 338
- DRM See *Dynamic RAID Migration (DRM)*
- DS controllers 408
- DS Storage Manager 132, 140, 221
- DS Storage Manager 10 Agent 549, 562
- DS Storage Manager 10 Utilities 549, 562
- DS3000 Enhanced Remote Mirroring ERM 21
- DS3000 FlashCopy Expansion License 20
- DS3000 FlashCopy VolumeCopy License 21
- DS3000 Partition Expansion License 20
- DS3000 VolumeCopy License 21
- DS3500 1
- DS3512 14
- DS3524 14
- DSM See *Device Specific Module (DSM)*
- DS-RSM 477
  - Model RS3 477
- DVE See *Dynamic Volume Expansion (DVE)*
- dynamic cache prefetch 235, 303
- dynamic cache read prefetch 234–235, 238, 241, 243
- Dynamic Capacity Expansion (DCE) 230
- dynamic disks 76
- Dynamic Multi-pathing (DMP) 595, 634
- Dynamic RAID Migration (DRM) 50, 228
- dynamic tracking 632
- Dynamic Volume Expansion (DVE) 230, 251

## E

- edge expanders 6
- edit scan duration days 246
- eligible replacement disk 267
- e-mail alert recipient 529–530
- enable background media scan 235
- Enable Controller Data Transfer 424
- enable feature 196
- Enable Remote Access 509
- Enabled
  - Closed 487
  - Custom 487
  - Open 487
- enclosure ID 47, 90, 198, 202, 259
- enclosure loss protection 61
- Enclosure Order 198
- Enclosures 338
- encrypt data 441
- encrypted data 440
- encryption and decryption keys 440
- encryption key 443
- End Session 322
- Enhanced Remote Mirroring (ERM) 14, 21, 33, 63, 71, 238
- Enterprise Management 157

- Enterprise Management window 26
- Enterprise Manager window 299
- Enterprise Unique Identifier (EUI) 10
- Environmental Services Module (ESM) 18, 349, 354
  - configuration 355
  - configuration settings 372
  - firmware 329, 349, 351, 355, 370
  - state capture 341
- ERM See *Enhanced Remote Mirroring (ERM)*
- ESM See *Environmental Services Module (ESM)*
- Ethernet management 87
- Ethernet management port 16, 37, 88, 262, 270, 272, 297, 479
- Ethernet priority 311, 319
- EUI See <Index>Enterprise Unique Identifier (EUI)
- Event Category 387
- event log 385
- Event Monitor Service 164
- Event Specific Codes 387
- Event Type 387
- Execute Script 161, 521
- execution throttle 46
- EXP3500 expansion enclosure 46, 87
- EXP3512 13–14, 18, 46
- EXP3524 13–14, 18, 46
- expanding a logical drive 239
- expansion enclosure ESMs 408
- explicit assignment 61
- export and import arrays 392
- export array 392–393, 461
- Export Array wizard 395
- Exported - ready for import 396, 398

## F

- fail a drive 426
- Fail Drive 426
- Failed status 428
- Failover 45
- Failover Alert Delay 198, 201
- failover cluster 556, 563
- Failover Cluster Management 563, 567, 581
- Failover Cluster Validation Report 569
- Fan-out expanders 6
- Fast!UTIL 610, 651
- Fault LED 16
- FC cable types 4
- FC protocol layers 3
- FDE See *Full Disk Encryption (FDE)*
- Feature Pack 197
- fiber cables 33
- Fibre Channel (FC) 1
- Fibre Channel adapters 33
- Fibre Channel over Ethernet (FCoE) 8
- Fibre Channel topologies 2
- file system 234
- firewall 487
- Firewall Status notifications 484
- firmware 349
- Firmware Inventory 351
- firmware pending 363

- firmware upgrade tool 354
- first management Ethernet ports 270
- Fixing status 160
- FlashCopy 20, 71, 219, 221, 238
  - logical drives 392, 398
  - repository 221
  - repository logical drives 392, 398
- FlashCopy repository 20
- flush write cache after 235
- FQDN 705
- frame sizes 311
- free capacity 221, 232–233, 251
- Full Disk Encryption (FDE) 49, 73, 259, 439
  - drive status 456, 460
  - drives 440, 443
  - hot-spare drive 460
  - premium feature 198, 443

## G

- Gather Support Information 330–331
- GBIC See *Gigabit Interface Converters (GBIC)*
- Generate Random Secret 314
- Gigabit Interface Converters (GBIC) 39
- global copy 72
- global e-mail recipient 530
- global mirroring 72
- global setting 529
- Guest Operating System 708

## H

- haldaemon 620
- hard drive firmware 381
- Hardware Components 181
- Hardware initiator 10, 310
- Hardware Management Console 630
- Hardware VPD 592
- hardware-enforced zone 125
- HBA See *Host Bus Adapters (HBA)*
- health check mode 633
- High Performance Tier 49
- Host Adapter BIOS 610
- Host Agent 648
- Host Bus Adapters (HBA) 2
  - driver 352
  - firmware 353
  - host port identifier 279
  - timeouts 619
- Host Group 289
- host group 64, 285–286, 289
- host interface card 16
- host interface daughter card 17
- host interface type 279
- host name 525
- host operating system 559
- host port identifier 327, 558
- Host ports 64
- host server 23
- Host system bus 44
- Host type 595, 633

- host-agent method 107
- host-attached management 23
- Hosts 64, 182
- Host-side connections 87
- host-to-logical drive mappings 183, 392, 398
- hot spare coverage 260
- hot spare drive 59, 181, 262, 456, 460
- hot spare location 263
- hot\_add 549–550
- Hot-add utility 28

## I

- I/O characteristics 234
- I/O data path 406
- I/O paths 616
- I/O per second (IOPS) 31
- IBM 22
- IBM activation 194
- IBM activation web site 195
- IBM Director 488
- IBM Disk Encryption Manager 441
- IBM Disk Encryption Storage Manager 440
- IBM My notifications 727
- IBM Product ID 499
- IBM Remote Support Manager 473
- IBM Support Portal page 720
- IBM Support Portal web site 719
- IBM System Storage DS Storage Manager 131
- IBM System Storage DS3500 1
- IBM System Storage DS3500 Storage Subsystem 13
- IBM System Storage Interoperation Center (SSIC) 588
- IBM System Storage interoperation Center (SSIC) 31
- ICMP 318
  - echo 320
- Ignored Event notifications 484
- Import Array 397–398, 468
- Import Completed 400
- import secure drive array 465
- Import Security File Key 261
- importing an array 397
- inactivity time-out 201
- in-band 23, 530
- in-band management 23, 107, 288
- Incompatible Drives 384
- Information Center 184
- informational event 386
- Inherit System Settings 159
- Initial Setup Task 187
- initialize a logical drive 413
- Initialize and Convert Disk wizard 551
- Initialize Array 414
- initialize drive 412
- initializing a logical drive 239
- initiator 547
- initiator authentication 314
- initiators 9
- Instant Secure Erase 442
- interactive mode 520
- inter-disk allocation 76
  - policy 76

- internal firewall 475, 488, 512
- Internet SCSI (iSCSI) 1
- Internet Small Computer System Interface 8
- Internet Storage Name Service (iSNS) 11, 311
- interoperability matrix 538
- Inter-Switch Link (ISL) 632
- Inter-Switch Links (ISL) 124
  - Trunking 124
- intra-disk allocation 76
- invalid key details 465
- IP address 308, 525, 528
- IP Security (IPSec) 12
- IPv4 89, 317
- IPv6 89, 317
- IPv6 settings 320
- IQN See *iSCSI Qualified Name (IQN)*
- iSCSI 8
  - daughter card 300, 308
  - discovery 11
  - hardware initiators 310
  - host 312
  - host port 604
  - host port identifiers 325
  - Host Ports 300, 317
  - initiator 291, 325, 602
  - initiator IQN name 326
  - initiators 312
  - management 309
  - multipathing architecture 310
  - naming 10
  - port settings 300
  - ports 262
  - protocol 309, 312
  - Qualified Names 311
  - sessions 311, 321
  - settings 297, 311
  - Software Initiators 310
  - statistics 311
  - SW initiator 325
  - target 602, 609
- iSCSI Qualified Name (IQN) 10, 325
- ISL See *Inter-Switch Links (ISL)*
- iSNS 311
- iSNS server 316

## J

- Java Access Bridge 28
- jumbo frames 311, 318

## K

- Kerberos V5 (KRB5) 12
- key backup 444
- known unassociated host port identifier 280, 326

## L

- large throughput 59
- LC connector 40
- LC-LC cables 96

- LDM See *Logical Disk Manager (LDM)*
- License Program Product (LPP) 628
- Life Preserver 346
- list of commands 533
- LNNAVIT 595
- load balancing 45, 633
  - algorithm 617
- Load Script 521
- Load Storage Subsystem Configuration 193, 306
- load-balancing algorithms 311
- Locate 192, 226, 261–262
- Locate Storage Subsystem 161, 298
- lock key 443
- Locked 460
- Logical Disk Manager (LDM) 76
- Logical Drive 63, 208, 219, 237, 338
  - mappings 64
- Logical Drive Capacity 183
- logical partition (LP) 75
- Logical Tab 208
- Logical tab 208, 219–220
- Logical Unit Number (LUN) 63, 275, 561
  - format settings 693
  - mapping 650
  - masking 64, 275
  - number 287, 292
- logical volume (LV) 74
- Logical Volume Manager (LVM) 74
- LP See *logical partition (LP)*
- LPP See *License Program Product (LPP)*
- LSI Logic Configuration Utility 539
- LUN See *Logical Unit Number (LUN)*
- LV See *logical volume (LV)*
- LVM See *Logical Volume Manager (LVM)*

## M

- Major Controller Firmware 355
- Major Event Log (MEL) 212, 480
- major firmware upgrades 354
- major revision 354
- major upgrade 354
- Manage 204, 214
- Manage a Storage Subsystem 159
- manage alert 529
- Manage Capacity Settings 254
- Manage Host Port Identifier 290, 327
- Manage Storage Subsystem 161
- management Ethernet ports 23
- management network 24
- management ports 24
- management station 23
- managing alerts 514
- manual discovery 152
- manual hot spare assignment 264
- Manual mode 224
- mapping change 289
- mappings 338
- Mappings menu 207
- Mappings tab 207, 275
- maximum number of LUNs 86

- media scan 68, 219, 238, 244
  - settings 243
  - with redundancy check 235, 238
- MEL See *Major Event Log (MEL)*
- memory allocation 710
- metro mirroring 72
- Microsoft Cluster 555
- Microsoft Windows MPIO 80
- mini-GBIC 4
- mini-SAS cable 36
- mini-SAS connector 36
- mini-SAS port 36
- Minor Controller Firmware 355
- minor firmware upgrades 354
- minor revision 354
- minor upgrade 354
- mirror repository 198, 221
- Mixed Zone 125
- mkpath command 636
- Modem connectivity 510
- modification priority 219, 235, 238–239
  - rates 239
- module logs 347
- Monitor Performance 198
- Move 293
- moving hosts 289
- MPIO Device Specific Module (DSM) 549, 562
- MPIO See *multipath I/O (MPIO)*
- mpio\_get\_config -Av command 635
- MPP See *Multi-Path Proxy (MPP)*
- mppBusRescan 597
- mppUpdate 597
- mppUtil 598
- MTU 320
- multi path driver 352
- multimedia 234
- multi-mode 4
- multi-mode fiber (MMF) 4, 38
- multipath driver 28, 80, 140, 562, 588
- multipath I/O (MPIO) 32, 81, 311, 557, 588, 628, 632
- Multi-Path Proxy (MPP) 28, 82, 599, 618
  - Driver 599
- multipath RDAC driver 352
- multipath support 550
- multipath.conf 619
- multipathd service 620
- multiple commands 532
- Mutual Authentication 311–312, 314
- My notifications 727

## N

- native VLANs 318
- Needs Attention 160, 353, 409
- Needs Upgrade 160
- Network Adapters 696
- Network Installation Manager (NIM) 638
  - server 638
- Network settings 310
- New Simple Volume 551
- New Virtual Machine 703

- NIM See *Network Installation Manager (NIM)*
- non-upgradeable 353, 356
- nonvolatile static random access memory (NVSRAM) 194, 329, 349, 374, 378, 595
  - file 359
  - firmware 355
  - microcode 349
  - settings 307
- Normal (0-125) 269
- N-port only (126) 269
- NTP server 673
- number of disk drives per array 303
- number of hot spare drives 303
- number of logical drives per array 303
- NVSRAM See *nonvolatile static random access memory (NVSRAM)*

## O

- offline state 418
- OK to remove LED 16
- Online Help 209, 436
- Online Transaction Processing (OLTP) 31
- operating status 194
- operations in progress 180
- Optimal 160, 404
- out-of-band 23, 530
  - management 24, 105
- output file 526

## P

- partition table 622
- pass phrase 442
- password 187, 297, 526
- path-control module (PCM) 81, 628
- PCM See *path-control module (PCM)*
- pending firmware 364
- performance counters 391
- Performance Monitor 213, 235, 248, 390
- Permission Set 314
- Persistent Reservation 401–402
- Phone Line Check notification 485
- physical partition (PP) 74
- Physical tab 208–209, 257
- physical volume (PV) 74
- ping 320
- ping check notifications 484
- place controller in Service Mode 421
- place controller Offline 419
- place controller Online 418
- place Logical drives Online 424
- point-in-time image 71
- point-to-point 2
- polling interval 391
- Port Level Zone 125
- power cabling 129
- power on/off sequence 129
- power supply 16, 129
- powering off 130
- powering on 130

- PP See *physical partition (PP)*
- preferred controller ownership 80
- Preferred Loop ID 262, 268
- prefetch 243
- Premium Feature Certificate 195
- Premium Feature License 183
- premium features 20, 70, 297, 307–308
- preparation checklist 395, 462
- Preparation Checklist (Export Array) window 394
- prepare for removal 408–409
- pre-read redundancy check 235, 238, 246
- pull-down menu 189
- PV See *physical volume (PV)*

## Q

- QLogic BIOS 651
- queue depth 46
- Quick Access buttons 172
- quorum 557, 572
- Quorum Configuration 574
- quorum drive 557

## R

- rack considerations 33
- RAID 0 50
- RAID 1 51
- RAID 10 51
- RAID 3 53
- RAID 5 53
- RAID 6 55
- RAID controller 15
- RAID level 184, 224, 227–228, 303
  - migration 229
- RAID levels 50
- RAID type 223
- Raw Data 388
- Raw mode 323
- RDAC 28, 281, 588, 595, 628
  - driver 594, 599
  - driver building 593
  - driver installation 594
  - multipath driver 349, 588, 593
- RDAC See *Redundant Disk Array Controller (RDAC)*
- read cache 235, 238, 241
- read caching 242
- Read Link Status (RLS) 341, 480
  - Change Log Files 345
  - logs 341
- Read test 406
- Reconstruct Drive 428
- reconstructing after a disk failure 239
- Recover from Failure 409
- recovered media error 245
- recovery commands 411
- Recovery Guru 179, 198, 214, 350, 409
  - icon 179
- Recovery Procedure 411
- Redistribute Logical Drives 424
- Redundancy Check 244

- Redundancy mismatches 245
- Redundant Disk Array Controller (RDAC) 82
  - Driver 599
- Reference Recovery Procedures 209
- Refresh 164
- registered ports 401
- remote access 270, 475, 487
  - by modem 485
  - by SSH 485
  - notifications 481
- Remote Logical Drive Mirroring 197
  - Premium Feature 197
- Remote Login 262, 270, 514
- Remote Mirror Drives 219
- Remote Supervisor Adapter (RSA) 476
- Remove 164, 294
- Remove Mapping 296
- Rename 204, 294
- rename a Logical Drive 254
- Rename Storage Subsystem 299
- Replace 232
- Replace Drive 260, 267
- replace iSCSI port identifier 328
- Replaced status 428
- Reporting 475, 486
- Rescan Devices 343
- Rescan Disks 550
- reservations 401
- reserve\_policy 631
- Reset Controller 422
- resetting a controller 422
- Revive drive 415
- RLOGIN 480, 514
- RLS See *Read Link Status (RLS)*
- rmpath command 636
- round-robin 45
- RPQ 439
- RSA See *Remote Supervisor Adapter (RSA)*
- rservice 487–488
- RSM 473, 490
  - administrator 490
  - for storage 474
  - management interface 486
  - software 490
- Run Configuration Test 502
- Run Diagnostics 405

## S

- sample configuration 537
- SAN Volume Controller (SVC) 78
- SANtricity UTM agent 599
- SAS cables 35, 91
- SAS Connectivity module 542
  - firmware 543
- SAS connectivity module 540
- SAS drive channel 341
- SAS Expansion Card 539
  - device driver 542
  - WWPN 539
- SAS Host Ports 88

- SAS ID 592
- SAS IN 99
- SAS OUT 99
- SAS protocol layers 6
- SAS See <Index>Serial-Attached SCSI (SAS)
- SAS wide ports 7
- Save As Script 283
- Save Configuration 193
  - link 306
- Save Security Key File 458
- save the support data 332
- saving the security key file 456
- script commands 520
- Script Editor 521
  - online help 534
- script file 533
- Scripting Editor 521
- SDD See <Index>Subsystem Device Driver (SDD)
- SDDP-CM See <Index>Subsystem Device Driver Path Control Module (SDDPCM)
- second management Ethernet ports 271
- secure array 472
- secure disk erase 456
- Secure Drive 231, 454
- Secure Erase 261, 459
- Secure fields 444
- secure file 442
- Secure Remote Password (SRP) 12
- secured state 442
- Security Capable 444
- security key 442–443
  - backup file 457
  - file 453
  - generation 451
  - ID 442
  - identifier 443, 451, 453, 457
- segment size 67, 233–235, 238, 249
- select files 358
- Self Encrypted Disk 259
- Sequence Number 387
- Serial ATA Tunneling Protocol (STP) 7
- Serial Attached SCSI (SAS) 5
- serial management port 16
- Serial Management Protocol (SMP) 7
- serial port 88
- Serial SCSI Protocol (SSP) 7
- Serial-Attached SCSI (SAS) 1, 5
- server BIOS 652
- server HBA 652
- server RAID 656
- Service 12
- Service Access settings 488
- Service Access Timeout value 514
- Service Access window 512
- service mode 418
- Set Baseline 323
- Set Degraded 404
- set during loop initialization 269
- Set Password 204
- Setup tab 297–298
- seven-segment LED 47
- SFF See <Index>Small Form Factor Transceivers (SFF)
- SFP See <Index>small-form-factor pluggable (SFP)
- Show FDE 259, 446
- Simple Network Mail Protocol (SNMP) 164, 168, 409
  - alert recipients 531
  - alerts 524
  - management 386
  - traps 474, 506
  - traps destination 506
- Simple Public-Key generic security service API Mechanism (SPKM1) 12
- Simple Public-Key generic security service API Mechanism (SPKM2) 12
- single Free Capacity node 431
- single-mode 4
- single-mode fiber (SMF) 4, 38
- SIOC See <Index>Storage I/O Control (SIOC)
- SM Monitor 334
- SMagent 27, 107, 599
- Small Form Factor Transceivers (SFF) 39
- small-form-factor pluggable (SFP) 4–5, 89, 93
  - module 93
  - transceivers 4
- SMcli 27, 283, 480, 519–520, 524
  - command-line interface 27
- SMdevices utility 28
- SMF See <Index>single-mode fiber (SMF)
- SMP See <Index>Serial Management Protocol (SMP)
- SMrepassist utility 28
- SMTP 164
- SMTP server 479
- SMutil 28
- SNMP See *Simple Network Mail Protocol (SNMP)*
- SOC 341
  - Change Log Files 345
  - logs 341
- software initiator 9, 310, 607
- software-enforced zone 125
- source logical drive 20
- spanned volumes 77
- Special Logical Drive 221
- Specify Capacity/Name 233
- Specify Host Port Identifier 326
- SRP See *Secure Remote Password (SRP)*
- SSH 474
- SSH connectivity 509, 687
- SSIC See *IBM System Storage interoperation Center (SSIC)*
- SSP See *Serial SCSI Protocol (SSP)*
- Staged Controller Firmware 355
- staged upgrades and activation 353
- Standard Logical Drive 221
- standby power LED 16
- State Capture 336
- state capture data 337
- State Change 387
- Status 178
- Status bar 215
- Storage Area Network (SAN) 2, 30

- boot 627
- boot implementation 637
- topology 613
- zoning 31, 124
- zoning configuration 613
- Storage Concepts Tutorial 185
- Storage I/O Control (SIOC) 648
- Storage Manager 132
- Storage Manager Agent (SMagent) 27
- Storage Manager Client (SMclient) 26
- Storage Manager Utility (SMutil) 28
- storage partition 183, 649
- storage partitioning 63, 286, 290, 310
- Storage Partitioning wizard 286
- storage subsystem
  - configuration file 341
  - identifier 520
  - name 297
  - password 194, 299, 307
  - profile 179, 338
- Storage Subsystem Needs Attention 179
- STP See <Index>Serial ATA Tunneling Protocol (STP)
- stripe size 68
- striped volumes 78
- strong password 204
  - rules 204, 299
- Strong password requirement not met 444
- Subsystem Device Driver (SDD) 81
- Subsystem Device Driver Path Control Module (SD-DP-CM) 81
- Subsystem Management window 26
- Subsystem Manager 177
- subsystem name 299, 525
- subsystem profile 178, 190
- Summary 338
- Summary tab 177
- support bundle files 345
- support data filename 335
- Support Monitor 341, 343–344
- Support Monitor Profiler 341–342
- Support tab 330, 436
- SuSE Linux 587
- SuSE Linux Enterprise Server 593, 615
- Suspend media scan 246
- SVC 78
- SVC See *SAN Volume Controller (SVC)*
- Switch-attached DS3500 42
- switched fabric 2
- Synchronize Controller Clocks 203
- synchronize the clock 272
- synchronous 21
- syntax 523
- system configuration 475

## T

- Target Authentication 311–312
- Target CHAP Secret 313
- target discovery 311, 316
- target Ethernet port 311
- Target Identification 311, 315

- target WWPN 547
- targets 9
- Task Assistant 158
- TCP listening port 604
- TCP Offload Engine (TOE) 8
- TELNET 270
- termination command 520
- Time-Out setting 198
- TOE See *TCP Offload Engine (TOE)*
- toolbar 212
- transfer NVSRAM file with controller firmware 375
- Transport mode 12
- traps 164
- TTY serial console 480
- Tunnel mode 12
- two-node file server failover cluster 578

## U

- Unassociated Host Port Identifiers 293
- unconfigured capacity 221–222, 251
- Undefined Mapping 292
- unfixable error 245
- uni-directional authentication 312
- unique security key ID 442
- Universal Transport Mechanism (UTM) 599
- unlock 467
- unlock drives 467
- Unlocked 460
- unmapped host 286
- Unreadable Sectors 434
  - Log 435
- unrecovered media error 245
- Unresponsive 160
- Update Timeout Value 512, 514
- Upgrade Controller Firmware 159, 163, 355, 365–366
- Upgradeable 357
- user IDs 488
- User Label 326
- UTM See *Universal Transport Mechanism (UTM)*

## V

- Validate a Configuration 567
- Validate SNMP configuration 507
- Veritas Volume Manager (VxVM) 77, 595
- VG See *volume group (VG)*
- View Associated Logical Element 260, 262
- View Associated Physical Components 222, 226, 239, 447
- View Associated Registrations 402
- View Event Log 330, 386
- View legend for table headings 324
- View Log 356
- View Module Log File 347
- View Online Help 331
- View Storage Subsystem Profile 330, 338
- View Unassociated Host Port Identifiers 293, 325
- view/change hot spare coverage 265
- VIOS 632
- Virtual I/O Server 628, 630

- virtual machine 700, 703
- Virtual Private Networks (VPN) 12
- VLAN 311
- VLAN ID 319
- VMFS datastore 670
- VMware Distributed Resource Scheduler (DRS) 647–648
- VMware ESX server 647, 652
- VMware ESX/ESXi server 648
- VMware High Availability 647
- VMware vCenter Server 648
- VMware vCenter Site Recovery Manager (SRM) 647–648
- VMware vSphere 656
- volume group (VG) 74
- VolumeCopy 21, 71, 219, 238
- VolumeCopy pairs 392, 398
- VPN See *Virtual Private Networks (VPN)*
- vSphere Client 682
- vSphere Server 678
- VxVM See *Veritas Volume Manager (VxVM)*

## W

- Windows Server 2008 537
- WINS 311
- witness disk 557, 576
- World Wide Identifier (WWID) 5, 528
- World Wide Name (WWN) 5, 126
- World Wide Node Name (WWNN) 126
- World Wide Port Name (WWPN) 126, 279–280, 291, 539, 541
  - Zone 125
- write cache 235, 238, 241
- Write cache with mirroring 235, 241–242
- write cache with mirroring 238
- write cache without batteries 235, 238, 241
- write caching 242
- write scheduling 76
- write test 406
- write-through 243
- write-verify policies 76
- WWID See *World Wide Identifier (WWID)*
- WWN See *World Wide Name (WWN)*
- WWNN See *World Wide Node Name (WWNN)*

## Y

- YaST 602, 607

## Z

- zoning 31, 547, 627, 649
- zoning configuration 588, 618, 645
- zoning with MPIO 32





# IBM System Storage DS3500 Introduction and Implementation Guide

(1.0" spine)

0.875" <-> 1.498"

460 <-> 788 pages







# IBM System Storage DS3500 Introduction and Implementation Guide

## Sample configurations with step-by-step instructions

## Configuration and administration with Storage Manager

## Troubleshooting and maintenance

This IBM Redbooks publication introduces the IBM System Storage DS3500, providing an overview of its design and specifications, and describing in detail how to set up, configure, and administer it. This edition covers updates and functions available with the DS3500 Storage Manager Version 10.70 (firmware level 7.70).

IBM has combined best-of-breed development with leading 6 Gbps host interface and drive technology in the IBM System Storage DS3500 Express. With its simple, efficient, and flexible approach to storage, the DS3500 is a cost-effective, fully integrated complement to System x servers, BladeCenter, and Power Systems. Offering substantial improvements at a price that will fit most budgets, the DS3500 delivers superior price to performance ratios, functionality, scalability, and ease-of-use for the entry-level storage user.

The DS3500 supports intermixing four 1 Gbps iSCSI or four 8 Gbps FC host ports with its native 6 Gbps SAS interfaces. This flexible and multi-purpose dual protocol approach allows organizations to implement a single storage system to support all of their shared storage requirements, thereby maximizing productivity, reliability, and cost.

Delivering solid input/output per second (IOPS) and throughput, the DS3500 controllers offer balanced and sustainable performance. The DS3500 can effectively double the performance of the previous DS3000 series of storage systems in both throughput and IOPS.

The DS3500 DS Storage Manager is the same management software offered with the DS5000 and DS4000 series. Now, any of these storage systems can be viewed and managed from a single interface. This allows for consolidated management of these various storage systems and a reduced learning curve. The DS3500 also supports enhanced remote mirroring over FC host ports, which is also compatible with the DS5000 and DS4000 series. This allows for low-cost backup and recovery with a DS5000 and DS4000 at a production site and a DS3500 at the secondary site.

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 IBM System Storage DS3500 with Storage Manager Software. It also targets those who have a DS3500 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](http://ibm.com/redbooks)

SG24-7914-00

ISBN 0738435414