

# IBM System Storage DCS3700 Introduction and Implementation Guide

IBM System Storage DS Storage  
Manager V10.8x

Hardware planning, installation,  
and maintenance

Host configuration and  
troubleshooting



Sangam Racherla  
Diego Silva Cogo  
Mario David Ganem  
Hartmut Lonzer  
Minja Obradovic  
Gabriel Requejo

**Redbooks**





International Technical Support Organization

**IBM System Storage DCS3700 Introduction and  
Implementation Guide**

July 2013

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

**First Edition (July 2013)**

This edition applies to IBM System Storage DCS3700 running V7.8x firmware and IBM System Storage DS Storage Manager V10.8x.

© Copyright International Business Machines Corporation 2013. 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

<b>Notices</b> .....	xi
Trademarks .....	xii
<b>Preface</b> .....	xiii
Authors .....	xiv
Now you can become a published author, too! .....	xv
Comments welcome .....	xv
Stay connected to IBM Redbooks .....	xv
<b>Chapter 1. Introduction to IBM System Storage DCS3700</b> .....	1
1.1 IBM System Storage Portfolio .....	2
1.2 IBM System Storage DCS3700 .....	3
1.2.1 DCS3700 components .....	5
1.3 IBM System Storage DCS3700 expansion unit .....	10
1.4 Premium Features .....	12
1.4.1 Partition Expansion License .....	12
1.4.2 FlashCopy expansion license .....	12
1.4.3 VolumeCopy license .....	12
1.4.4 FlashCopy VolumeCopy license .....	12
1.4.5 Try and Buy Premium Features .....	13
1.5 DCS3700 and DS3700P comparison .....	14
1.6 IBM System Storage DS Storage Manager .....	14
1.6.1 In-band management .....	15
1.6.2 Out-of-band management .....	17
1.6.3 Storage Manager Client (SMclient) .....	18
1.6.4 Storage Manager Agent (SMagent) .....	20
1.6.5 Storage Manager Utility (SMutil) .....	20
1.6.6 Multipath driver support .....	20
1.6.7 Java Access Bridge .....	21
<b>Chapter 2. IBM System Storage DCS3700 storage subsystem planning tasks</b> .....	23
2.1 Planning your SAN and storage server .....	24
2.1.1 SAN zoning for the DCS3700 storage subsystem .....	26
2.1.2 Zoning considerations for Enhanced Remote Mirroring .....	27
2.2 Planning for physical components .....	28
2.2.1 Rack considerations .....	28
2.2.2 SAS cables and connectors .....	30
2.2.3 Ethernet cable and connections .....	32
2.2.4 Fibre Channel cables and connectors .....	33
2.2.5 Fibre Channel adapters .....	40
2.2.6 Disk expansion units .....	43
2.3 High Performance Module (DCS3700P) .....	46
2.4 Planning your storage structure .....	46
2.4.1 Selecting drives .....	47
2.4.2 Understanding RAID types .....	48
2.4.3 Array configuration .....	57
2.4.4 Dynamic Disk Pooling .....	59
2.4.5 Thin provisioning .....	61
2.4.6 Hot spare drives .....	62

2.4.7	Data loss protection planning . . . . .	64
2.4.8	Logical drives and controller ownership . . . . .	66
2.4.9	Storage partitioning . . . . .	67
2.4.10	Segment size . . . . .	71
2.4.11	Media scan . . . . .	72
2.4.12	Cache parameters . . . . .	73
2.5	Planning for Premium Features . . . . .	74
2.5.1	VolumeCopy . . . . .	74
2.5.2	FlashCopy . . . . .	76
2.5.3	Enhanced FlashCopy . . . . .	77
2.5.4	Enhanced Remote Mirroring . . . . .	79
2.5.5	Obtaining Premium Features key . . . . .	82
2.6	Additional planning considerations . . . . .	82
2.6.1	Planning for systems with Logical Volume Manager: AIX example . . . . .	82
2.6.2	Planning for systems without LVM: Windows example . . . . .	85
2.6.3	Virtualization . . . . .	87
2.6.4	SAN Volume Controller overview . . . . .	87
2.7	Host support and multipathing . . . . .	88
2.7.1	Supported server platforms . . . . .	88
2.7.2	Supported operating systems . . . . .	88
2.7.3	Clustering support . . . . .	89
2.7.4	Multipathing . . . . .	89
2.7.5	Microsoft Windows MPIO . . . . .	89
2.7.6	AIX MPIO . . . . .	90
2.7.7	AIX Subsystem Device Driver Path Control Module . . . . .	90
2.7.8	Linux: RHEL/SLES . . . . .	90
2.7.9	Function of the Auto-Logical Drive Transfer feature . . . . .	92
2.7.10	Asymmetric logical unit access . . . . .	94
2.8	Operating system considerations . . . . .	95
2.8.1	Maximum capacity for a logical drive . . . . .	95
2.8.2	Maximum number of LUNs per host . . . . .	96
	<b>Chapter 3. IBM System Storage DCS3700 and DCS3700 expansion unit cabling . . . . .</b>	<b>97</b>
3.1	DCS3700 controller connectors . . . . .	98
3.1.1	DCS3700 controller with standard port configuration . . . . .	98
3.1.2	The DCS3700 controller with an optional SAS host port adapter . . . . .	98
3.1.3	DCS3700 controller with an optional Fibre Channel host port adapter . . . . .	99
3.1.4	DCS3700 Performance Module . . . . .	99
3.1.5	DCS3700 expansion unit ports . . . . .	100
3.2	Enclosure ID settings . . . . .	101
3.3	SAS cables . . . . .	101
3.4	Fibre Channel cabling . . . . .	103
3.4.1	SFP transceiver modules . . . . .	104
3.4.2	Fibre Channel cables . . . . .	106
3.4.3	Interoperability of 2 Gbps, 4 Gbps, 8 Gbps, and 16 Gbps devices . . . . .	108
3.5	iSCSI Ethernet cabling . . . . .	108
3.6	DCS3700 expansion unit attachment . . . . .	108
3.6.1	Redundant drive channels . . . . .	108
3.6.2	Drive channel cabling rules . . . . .	109
3.6.3	Connecting the DCS3700 expansion unit . . . . .	109
3.6.4	Adding a DCS3700 expansion unit to a running dual-controller configuration . . . . .	111
3.6.5	SAS drive channel miswires . . . . .	113
3.7	Management connections . . . . .	113

3.7.1	Out-of-band management . . . . .	114
3.7.2	In-band management . . . . .	115
3.8	Host attachment . . . . .	116
3.8.1	SAS attachment . . . . .	117
3.8.2	Direct attached Fibre Channel . . . . .	118
3.8.3	SAN fabric-attached DCS3700 . . . . .	120
3.8.4	iSCSI attachment . . . . .	125
3.9	Power cabling . . . . .	128
3.9.1	The DCS3700 power supply . . . . .	128
3.9.2	Powering on and off . . . . .	129
<b>Chapter 4. IBM System Storage DS Storage Manager installation . . . . .</b>		<b>131</b>
4.1	Installing DS Storage Manager on Microsoft Windows 2008 . . . . .	132
4.1.1	Installation preparation . . . . .	132
4.1.2	Installing the Storage Manager Client on Microsoft Windows 2008 . . . . .	132
4.2	Installing DS Storage Manager on Linux . . . . .	140
4.2.1	Preparing for the installation . . . . .	140
4.2.2	Installing Storage Manager by using the GUI . . . . .	141
4.2.3	Installing DS Storage Manager by using a text console . . . . .	146
4.2.4	Uninstalling DS Storage Manager on Linux . . . . .	148
4.3	Installing DS Storage Manager on AIX . . . . .	150
4.3.1	Preparing for the installation . . . . .	150
4.4	Completing the DS Storage Manager installation . . . . .	151
4.4.1	Performing an automatic discovery of storage subsystems . . . . .	152
4.4.2	Performing a manual discovery of storage subsystems . . . . .	153
4.4.3	Adding a Storage Subsystem verification . . . . .	157
<b>Chapter 5. Administration: Enterprise Management . . . . .</b>		<b>159</b>
5.1	Enterprise Management window overview . . . . .	160
5.1.1	Initial Setup Tasks . . . . .	160
5.1.2	Enterprise Management window . . . . .	162
5.2	Functions in the Enterprise Management window . . . . .	163
5.2.1	Subsystem menu . . . . .	164
5.2.2	The Enterprise Management window menu bar . . . . .	179
5.2.3	The Quick Access icons . . . . .	181
<b>Chapter 6. Administration: Summary tab . . . . .</b>		<b>187</b>
6.1	DS Storage Manager: Summary tab . . . . .	188
6.2	Monitor . . . . .	189
6.2.1	Storage subsystem status . . . . .	189
6.2.2	Storage Subsystem communication error . . . . .	189
6.2.3	Operations in Progress . . . . .	190
6.2.4	View Firmware Inventory . . . . .	190
6.2.5	View Storage Subsystem Profile . . . . .	192
6.2.6	View Event Log . . . . .	193
6.2.7	Collect Support Data . . . . .	194
6.3	Capacity . . . . .	195
6.3.1	Create Storage . . . . .	196
6.3.2	Create Logical Drive . . . . .	197
6.4	Storage & Copy Services . . . . .	198
6.5	Host Mappings . . . . .	198
6.5.1	Host-to-Logical Drive Mappings . . . . .	199
6.6	Premium Features . . . . .	199
6.6.1	Manage Premium Features . . . . .	200

6.7 Hardware . . . . .	200
6.8 Information Center . . . . .	201
<b>Chapter 7. Administration: Subsystem Management . . . . .</b>	<b>203</b>
7.1 DS Storage Manager: Subsystem Management. . . . .	204
7.2 Drop-down menu. . . . .	205
7.2.1 Storage Subsystem menu. . . . .	205
7.3 Toolbar . . . . .	221
7.4 Status bar . . . . .	222
7.5 Tabs . . . . .	222
<b>Chapter 8. Administration: Storage &amp; Copy Services tab . . . . .</b>	<b>225</b>
8.1 Storage & Copy Services tab . . . . .	226
8.2 Working with unconfigured capacity . . . . .	228
8.2.1 View Associated Physical Components . . . . .	228
8.2.2 Disk pool . . . . .	229
8.2.3 Create Array . . . . .	232
8.3 Working with arrays. . . . .	235
8.3.1 Locate and View Associated Components . . . . .	236
8.3.2 Add Free Capacity (Drives). . . . .	236
8.3.3 Replace Drive . . . . .	238
8.3.4 Change Ownership and RAID level . . . . .	238
8.3.5 Deleting and renaming arrays. . . . .	241
8.4 Creating a disk pool . . . . .	242
8.5 Working with a disk pool . . . . .	244
8.5.1 Locate and View Associated Components . . . . .	244
8.5.2 Add Free Capacity (Drive). . . . .	245
8.6 Working with free capacity in arrays . . . . .	246
8.6.1 Creating a logical drive . . . . .	246
8.7 Working with free capacity in a disk pool . . . . .	249
8.7.1 Creating a standard logical drive. . . . .	249
8.7.2 Thin logical drive . . . . .	250
8.7.3 Creating a thin logical drive. . . . .	252
8.8 Working with logical drives . . . . .	255
8.8.1 Create option . . . . .	257
8.8.2 Increase Capacity option . . . . .	258
8.8.3 Increase Repository Capacity . . . . .	260
8.8.4 Change . . . . .	262
8.8.5 Change Modification Priority . . . . .	262
8.8.6 Change Cache Settings . . . . .	264
8.8.7 Change Media Scan Settings . . . . .	268
8.8.8 Change Pre-Read Redundancy Check. . . . .	271
8.8.9 Change Ownership/Preferred Path. . . . .	272
8.8.10 Change Segment Size . . . . .	273
8.8.11 Change Repository settings . . . . .	277
8.8.12 Add to Consistency Group and Remove from Consistency Group . . . . .	278
8.8.13 Convert FlashCopy to Enhanced FlashCopy Group. . . . .	278
8.8.14 View Associated Physical Component . . . . .	281
8.8.15 Delete and Rename . . . . .	281
8.8.16 Advanced . . . . .	281
8.9 Consistency groups. . . . .	285
8.9.1 Creating a consistency group . . . . .	288
8.9.2 Work with Consistency Group. . . . .	293

<b>Chapter 9. Administration: Hardware tab</b> . . . . .	305
9.1 Hardware tab . . . . .	306
9.2 Discovering component properties and location . . . . .	307
9.2.1 Showing the disk type . . . . .	307
9.2.2 View Enclosure Components . . . . .	308
9.2.3 Disk Drive menu . . . . .	311
9.2.4 Controller menu . . . . .	312
9.3 Setting a hot spare drive . . . . .	314
9.4 Failed disk drive replacement . . . . .	319
9.5 Setting a preferred loop ID . . . . .	321
9.6 Setting remote access . . . . .	323
9.7 Setting the Ethernet management ports . . . . .	323
9.7.1 Initial settings for the default IP addresses in DS Storage Manager . . . . .	324
9.7.2 Configuring the Ethernet management ports . . . . .	326
<b>Chapter 10. Administration: Mappings tab</b> . . . . .	327
10.1 Mappings tab . . . . .	328
10.2 Defining a host . . . . .	330
10.2.1 Adding a host to an existing host group . . . . .	338
10.3 Defining storage partitioning . . . . .	339
10.4 Defining a host group . . . . .	343
10.5 Manage Host Port Identifiers option . . . . .	346
10.6 Define Additional Mapping option . . . . .	346
10.7 Viewing unassociated ports . . . . .	348
10.8 Moving, removing, and renaming a host . . . . .	348
10.8.1 Move option . . . . .	348
10.8.2 Remove option . . . . .	350
10.8.3 Rename option . . . . .	350
10.9 Changing a host type . . . . .	350
10.10 Changing and removing mapping . . . . .	351
10.10.1 Change Mapping option . . . . .	351
10.10.2 Remove Mapping option . . . . .	352
<b>Chapter 11. Administration: Setup tab</b> . . . . .	353
11.1 Setup tab . . . . .	354
11.1.1 Locate Storage Subsystem . . . . .	354
11.1.2 Rename Storage Subsystem . . . . .	355
11.1.3 Change Hardware View Order . . . . .	356
11.1.4 Set a Storage Subsystem Password . . . . .	356
11.1.5 Manage Premium Features . . . . .	357
11.1.6 Configure iSCSI Host Ports . . . . .	358
11.1.7 Create Storage . . . . .	358
11.1.8 Save Configuration . . . . .	360
11.1.9 Manually Define Hosts . . . . .	361
11.1.10 Map Logical Drives . . . . .	361
11.1.11 Configure Ethernet Management Ports . . . . .	361
11.1.12 Manage iSCSI Setting . . . . .	362
<b>Chapter 12. Administration: iSCSI</b> . . . . .	363
12.1 Planning for iSCSI attachment . . . . .	364
12.1.1 Hardware initiators . . . . .	364
12.1.2 Software initiators . . . . .	364
12.1.3 Storage partitioning . . . . .	364
12.1.4 Network settings . . . . .	365

12.1.5 Security . . . . .	365
12.2 iSCSI configuration summary . . . . .	365
12.3 Manage iSCSI protocol settings . . . . .	367
12.3.1 Target authentication . . . . .	367
12.3.2 Mutual authentication . . . . .	369
12.3.3 Target identification . . . . .	371
12.3.4 Target discovery . . . . .	371
12.4 Configuring iSCSI host ports . . . . .	372
12.5 View/End iSCSI Sessions . . . . .	377
12.6 Viewing iSCSI statistics . . . . .	379
12.7 Defining iSCSI hosts . . . . .	381
12.7.1 Viewing unassociated iSCSI initiators . . . . .	381
12.7.2 Defining a new iSCSI host . . . . .	382
12.7.3 Managing iSCSI host ports . . . . .	384
<b>Chapter 13. Administration: Monitor . . . . .</b>	<b>385</b>
13.1 The Subsystem Management Monitor menu . . . . .	386
13.2 View Health (Recovery Guru) . . . . .	387
13.3 Performance Monitor . . . . .	390
13.3.1 Performance Monitor: Table View . . . . .	390
13.3.2 Graphical Performance Monitor . . . . .	392
13.4 Collect Support Data . . . . .	393
13.4.1 Saving the support data . . . . .	394
13.4.2 Automatic Support Data Collection settings . . . . .	397
13.5 Retrieve Trace Buffers . . . . .	399
13.6 Storage subsystem diagnostic tests . . . . .	400
13.6.1 Read Link Status . . . . .	401
13.6.2 Collecting I/O Path Statistics . . . . .	401
13.6.3 Validating the configuration database . . . . .	403
13.7 Collecting drive data . . . . .	403
13.8 Collecting state capture . . . . .	405
13.9 Operations in Progress . . . . .	407
13.10 Storage Subsystem Profile . . . . .	407
13.11 Cable connection . . . . .	409
13.12 Event Log . . . . .	410
13.13 Unreadable Sectors Log . . . . .	417
13.14 Maintenance: Persistent Reservations . . . . .	417
13.15 IBM Remote Support Manager for Storage . . . . .	419
<b>Chapter 14. Administration: Upgrade . . . . .</b>	<b>421</b>
14.1 The Subsystem Management Upgrade menu . . . . .	422
14.2 Downloading firmware . . . . .	423
14.2.1 Before you upgrade . . . . .	424
14.2.2 Firmware inventory . . . . .	425
14.2.3 Updating the host . . . . .	427
14.3 Upgrading the DCS3700 controller firmware . . . . .	428
14.3.1 Staged upgrades and activation . . . . .	428
14.3.2 Pre-upgrade checks . . . . .	428
14.3.3 Major and minor firmware upgrades . . . . .	429
14.3.4 Using the Enterprise Management upgrade tool . . . . .	430
14.3.5 Using the DCS3700 Storage Manager (Subsystem Management) . . . . .	448
14.4 Controller NVSRAM . . . . .	453
14.5 Drive firmware . . . . .	455

14.6 ESM firmware . . . . .	460
14.7 Enclosure configuration settings . . . . .	462
<b>Chapter 15. Command-line interface (CLI) . . . . .</b>	<b>463</b>
15.1 Introduction to the command-line interface . . . . .	464
15.2 How to use the command-line interface . . . . .	464
15.2.1 Usage notes . . . . .	465
15.3 Running the command-line interface . . . . .	465
15.3.1 Script Editor . . . . .	465
15.4 General SMcli syntax . . . . .	468
15.5 Adding a storage subsystem to the Storage Manager configuration . . . . .	472
15.6 Showing the defined subsystems in the Storage Manager configuration . . . . .	473
15.7 Configuring alerts . . . . .	474
15.7.1 Defining the mail server and email address to send out the email alerts . . . . .	474
15.7.2 Defining email alert recipients . . . . .	474
15.7.3 Deleting email alert recipients . . . . .	475
15.7.4 SNMP alert recipients . . . . .	476
15.8 Issuing commands to the storage subsystem . . . . .	476
15.8.1 Sample command: Save configuration script file . . . . .	478
15.9 More information . . . . .	480
<b>Chapter 16. Windows SAS configuration guide for IBM BladeCenter . . . . .</b>	<b>481</b>
16.1 Required equipment . . . . .	482
16.2 IBM BladeCenter setup . . . . .	483
16.2.1 Installing Windows Server 2008 R2 . . . . .	483
16.2.2 HS23 SAS Expansion Cards . . . . .	483
16.2.3 Recording the SAS Expansion Card WWPN . . . . .	483
16.2.4 HS23 SAS Expansion Card device driver . . . . .	486
16.2.5 SAS Connectivity Modules . . . . .	486
16.2.6 SAS Connectivity Module firmware update . . . . .	487
16.2.7 Configuring the SAS Connectivity Module . . . . .	493
16.2.8 SAS Connectivity Module zoning . . . . .	493
16.3 Installing the DS Storage Manager host software . . . . .	495
16.4 Configuring the disk space in Windows Server 2008 R2 . . . . .	496
<b>Chapter 17. Windows Fibre Channel configuration guide . . . . .</b>	<b>503</b>
17.1 Fibre Channel storage configuration . . . . .	504
17.1.1 Planning for the installation . . . . .	504
17.1.2 Installing the HBA drivers . . . . .	505
17.1.3 Installing QLogic SANSurfer . . . . .	506
17.1.4 Installing the Windows multipath driver . . . . .	511
17.1.5 Discovering devices (logical disk drives) . . . . .	514
<b>Chapter 18. Linux configuration guide . . . . .</b>	<b>521</b>
18.1 Fibre Channel storage configuration . . . . .	522
18.1.1 Planning for the installation . . . . .	522
18.1.2 Installing HBA drivers . . . . .	523
18.1.3 Installing QLogic SANSurfer . . . . .	524
18.1.4 Installing the Linux multipath driver . . . . .	528
18.2 SAS storage configuration . . . . .	536
18.2.1 Preparing for the installation . . . . .	538
18.2.2 Getting SAS Adapters address . . . . .	540
18.2.3 Installing the Linux multipath driver . . . . .	542
18.2.4 Discovering devices (disk drives) . . . . .	542

18.3	Managing the disk space by using IBM DS Storage Manager and LVM. . . . .	543
18.4	Troubleshooting guide. . . . .	548
18.4.1	rescan_scsi_bus.sh is not showing any new device. . . . .	548
18.4.2	multipath -ll output does not show all the LUNs. . . . .	548
18.4.3	multipath does not show paths that are equal to the number of LUNs that are mapped. . . . .	548
18.4.4	multipath -ll output is showing the paths as [failed][faulty] . . . . .	548
<b>Chapter 19. IBM AIX V7.1 configuration guide</b> . . . . .		549
19.1	Planning for the installation. . . . .	550
19.1.1	Zoning considerations. . . . .	552
19.1.2	Planning MPIO . . . . .	553
19.2	Configuring MPIO . . . . .	558
19.3	Setting up the DCS3700 logical drives and host mapping . . . . .	560
19.4	Scanning and managing the storage logical drive from AIX . . . . .	561
19.4.1	Ways to manage the paths . . . . .	562
19.5	SAS attachment . . . . .	563
19.5.1	Equipment required. . . . .	563
19.5.2	IBM BladeCenter setup. . . . .	564
19.5.3	SAS Connectivity Modules . . . . .	564
19.5.4	SAS card installation. . . . .	564
19.5.5	SAS Connectivity Module firmware update. . . . .	566
19.5.6	SAS Connectivity Module zoning . . . . .	568
19.5.7	Mapping a LUN from IBM DS Storage Manager. . . . .	571
19.5.8	Adding a LUN in AIX. . . . .	577
<b>Appendix A. IBM Support Portal website</b> . . . . .		583
	Sample navigation procedure . . . . .	584
	Downloading code updates . . . . .	587
	My notifications . . . . .	595
<b>Related publications</b> . . . . .		601
	IBM Redbooks . . . . .	601
	Other publications . . . . .	601
	Online resources . . . . .	602
	Help from IBM . . . . .	602



# Notices

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

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

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

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

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

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

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

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

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

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

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

## COPYRIGHT LICENSE:

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

# Trademarks

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

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

AIX®	Netfinity®	System i®
AIX 5L™	POWER®	System p®
BladeCenter®	Power Systems™	System Storage®
DB2®	POWER6®	System Storage DS®
DS4000®	POWER7®	System x®
DS8000®	PowerHA®	Tivoli®
FlashCopy®	Redbooks®	XIV®
GPFST™	Redbooks (logo)  ®	
IBM®	Storwize®	

The following terms are trademarks of other companies:

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

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

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

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

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

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

# Preface

The IBM® System Storage® DCS3700 is designed to meet the storage needs of highly scalable, data streaming applications in high-performance computing environments. The DCS3700 offers optimal space usage, low power consumption, and high performance. The DCS3700 is equally adept at delivering throughput to bandwidth-intensive applications and I/O operations to transactional applications. By combining proven IBM storage controllers with up to 60 drives in just 4U of rack space, the DCS3700 can reduce operational costs for capacity-intensive applications.

The DCS3700 is composed of an IBM System Storage DCS3700 storage subsystem and IBM System Storage DCS3700 expansion unit. The DCS3700 features the latest technologies, including 6 Gbps SAS and 8 Gbps Fibre Channel host interfaces, along with 6 Gbps SAS drives. The DCS3700 also has a 10 Gbps iSCSI host interface with an optional Performance Modules system.

The DCS3700 provides a simple, efficient, and flexible approach to storage that is based on seven generations of design knowledge and firmware development. The DCS3700 can act as a cost-effective, fully integrated complement to IBM System x® servers, IBM BladeCenter®, and IBM Power Systems™ for a wide variety of intensive computing environments.

This IBM Redbooks® publication specifically addresses the hardware features, configuration, and implementation of the DCS3700. We present detailed descriptions of the hardware configurations and options that are offered with the DCS3700. We then present the concepts and functions that are used in planning and managing the storage servers, such as multipathing and path failover. This book offers a step-by-step guide to using the Storage Manager to create arrays, logical drives, and other basic (and advanced) management tasks.

This book also contains practical information about diagnostic tests and troubleshooting, and includes practical examples of how to use scripts and the command-line interface (CLI).

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 DCS3700 storage servers with Storage Manager Software V10.8x. It also targets those who have a DCS3700 storage subsystem and need detailed advice about how to configure it.

**Important:** The iSCSI option for DCS3700 is only available starting with firmware level Version 7.84. Although this book was written before the release of the iSCSI option, we have since added the necessary information for iSCSI to the best of our ability.

## Authors

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 ITSO in San Jose, CA. He has 12 years of experience in the IT field and has been with the ITSO for the past eight years. Sangam has extensive experience in installing and supporting the ITSO lab equipment for various IBM Redbooks publications projects. He has expertise in working with Microsoft Windows, Linux, IBM AIX®, IBM System x, and IBM System p® servers, and various SAN and storage products. Sangam holds a degree in Electronics and Communication Engineering.

**Diego Silva Cogo** is an IT professional that specializes in cloud and virtualization solutions. He has more than 11 years of experience in the IT industry. Diego resides in Porto Alegre, Brazil, where he works as a VMware IT Specialist for the Gerdau Long Steel North America Strategic Outsourcing account. He holds a Bachelor's degree in Network Technology and more than 20 Microsoft certifications.

**Mario David Ganem** is an IT professional that specializes in cloud computing and storage solutions. He has 15 years of experience in the IT industry. Mario resides in Buenos Aires, Argentina, where he works as Infrastructure IT Architect in the IBM Global Delivery Center. Before he started his career with IBM in 2006, Mario worked for many companies, such as Hewlett Packard, Compaq, and Unisys. He developed the internal virtualization products curricula training that he teaches to DCA professionals. He holds, among others, many industry certifications from Microsoft, Red Hat, VMware, Novell, Cisco, CompTIA, Hewlett Packard, and Compaq.

**Hartmut Lonzer** is a Technical Consultant at the Partnership Solution Center Southwest / Germany. As a former Storage FTSS member, his main focus is on Storage and System x. He is responsible for educating and supporting the Business Partners and customers in his area in technical matters. His experience regarding DS Storage goes back to the beginning of this product. He has been with IBM in various technical roles for 35 years.

**Minja Obradovic** is an IT Specialist in Technical Support Services at IBM Serbia since 2007. He is a member of the CEEMEA Virtual Front End Midrange Storage Support team, which provides storage subsystems support in Central Eastern Europe, the Middle East, and Africa. His primary focus is post-sales Storage Virtualization, SAN, and Midrange Storage support. His areas of expertise include SAN switches (Brocade and Cisco MDS), SAN Storage (IBM System Storage DS3000, IBM System Storage DS4000®, IBM System Storage V7000, and IBM XIV®), IBM System x, and Power Systems.

**Gabriel Requejo** is an IT professional with more than 15 years of experience in the IT field. For the last 4 years, he has worked on several virtualization projects for different companies around the world by using VMware products with different IBM Storage solutions. Gabriel worked for several domestic and international customers providing Citrix and Windows oriented solutions using System x servers. Gabriel has experience with AIX, Linux, Novell, and Microsoft products.

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

Jon Tate, David Bennin, David Watts, Tamikia Barrow, Shari Deiana  
**IBM International Technical Support Organization**

Dahn Le, Richard Mancini, Gene Cullum, Alexander Watson, John Sexton, Mousa Hammad, Kerry Anders, Jason M Schwarz, Patrick Caporale, Jim Gallagher, Syman N Lee, Nelson YH Hsu,  
**IBM**

Doug Merrill, John Bish, David Worley, Stephen Carl, Michael Tischendorf  
**NetApp**

Brian Steffler and Steven Tong  
**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](http://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](http://ibm.com/redbooks)

- ▶ Send your comments in an email to:

[redbooks@us.ibm.com](mailto: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>



# **Introduction to IBM System Storage DCS3700**

This chapter introduces the IBM System Storage DCS3700 offerings and functionality. These products consist of models of storage subsystems that provide various environments to meet various user needs. This chapter describes the DCS3700 expansion unit as well.

This chapter also explains the Premium Features philosophy and how the Storage Manager utility works with the DCS3700 storage subsystem.

# 1.1 IBM System Storage Portfolio

IBM has brought together into one family, which is 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 IBM System Storage DS5000 series of mid-range systems and the IBM System Storage DS3000 entry level systems. The whole range of IBM storage subsystems is shown in Figure 1-1.

IBM System Storage Portfolio		Optimized for 'open systems'	Optimized for z/OS and IBM i
	Block		File
<b>Enterprise</b>	<b>DS8000</b> For clients requiring: <ul style="list-style-type: none"> <li>Advanced disaster recovery with 3-way mirroring and System z GDPS support</li> <li>Continuous availability, no downtime for upgrades</li> <li>Best-in-class response time for OLTP or transaction workloads, flash optimization</li> <li>Single system capacity scalable to the PB range</li> </ul>	<b>XIV</b> For clients requiring: <ul style="list-style-type: none"> <li>Virtually unlimited snapshots</li> <li>High utilization with automated performance ("hot spot") management</li> <li>Advanced self-healing architecture</li> <li>Breakthrough ease of use and management</li> </ul>	<b>SONAS</b> For clients requiring: <ul style="list-style-type: none"> <li>Massive I/O, backup, or restores</li> <li>Consolidation, or scale large numbers of clusters</li> </ul>
<b>Midrange</b>	<b>DS5000</b> For clients requiring: <ul style="list-style-type: none"> <li>Good cost/performance, general-purpose storage</li> <li>Need to add capacity to existing configuration</li> <li>10-100s of TB of capacity</li> </ul>	<b>Storwize V7000</b> For clients requiring: <ul style="list-style-type: none"> <li>10s of TB of rack-mounted storage with sophisticated software functions</li> <li>Breakthrough ease of use and management</li> <li>Non-disruptive migration from, or virtualization of, existing disk</li> </ul>	<b>N series</b> For clients requiring: <ul style="list-style-type: none"> <li>NAS storage</li> <li>Simple two-site high availability</li> </ul> <b>Storwize V7000 Unified</b> For clients requiring: <ul style="list-style-type: none"> <li>NAS storage virtualization</li> <li>Breakthrough ease of use and management</li> </ul>
<b>Specialized Performance/Density</b>	<b>New DCS3700 / DCS3700P</b> <b>High density and performance to support intensive computational applications</b> <b>High sequential bandwidth - for Digital Media, Financial Mkts, Telecomm</b>		<b>DCS3700 + GPFS</b>
<b>Entry</b>	<b>DS3000</b> For clients requiring: <ul style="list-style-type: none"> <li>SMBs or branch office locations; cost sensitive; start at as small as 100s GB, up to 10s TB capacity</li> </ul>		<b>N series</b> For clients requiring: <ul style="list-style-type: none"> <li>NAS storage</li> <li>Simple two-site high availability</li> </ul>
<b>Storage Optimizers</b>	<b>SVC, Easy Tier, ProtecTIER, Information Archive, Real-time Compression Appliances</b> For clients requiring optimization of storage costs: <ul style="list-style-type: none"> <li>Storage technologies that add function, performance, ease of use, or efficiency to new or existing storage</li> </ul>		

Figure 1-1 IBM System Storage Portfolio

With the announcement of IBM System Storage DCS3700, IBM introduced a new member of the DS family that has the characteristics of entry and mid-range models.

IBM System Storage DCS3700 is good fit for the entry to mid-range SAN and direct-attach market space. With a common Storage Manager shared by DS3000 storage subsystems and the DS5000 storage subsystems, there is a smooth link between DS3000 and DS5000 series systems, with remote mirroring and copy services features being shared by these two platforms.



The overall positioning of the DCS3700 within the entry and mid-range IBM System Storage DS® family is shown in Figure 1-2.

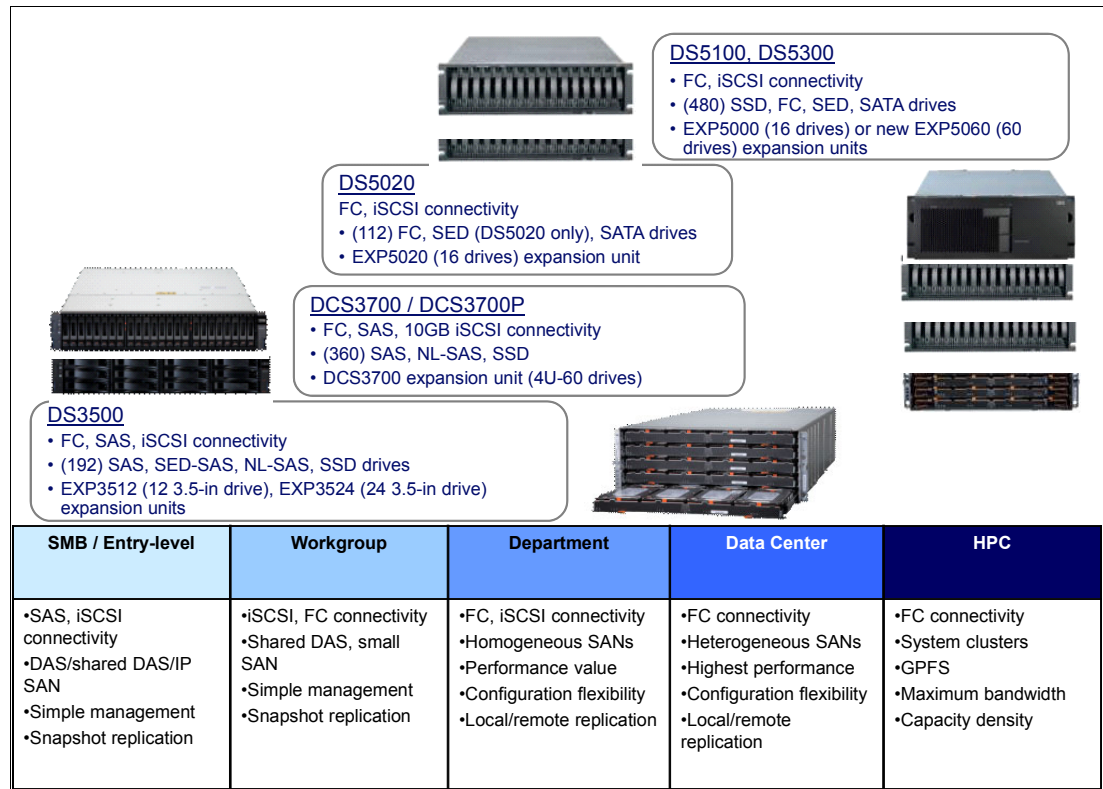


Figure 1-2 IBM DS3000/5000 entry level and mid-range storage positioning

## 1.2 IBM System Storage DCS3700

The era of smarter computing is here and requires more data with faster data access. Low-latency application performance in applications such as life sciences, real-time analytics, rich media, seismic processing, weather forecasting, telecommunications, and financial markets require high-performance storage architectures. In addition, organizations are often looking for a better way to improve operational efficiency while maintaining the same data center footprint, quality of service (QoS), and high availability.

The IBM System Storage DCS3700 meets an organization's need for more capacity within their space constraints. Designed for applications with high-performance streaming data requirements, the DCS3700 offers optimal space usage, low power consumption, and high performance. Organizations can now have a storage solution to maximize storage density, reduce operational expenditures, and ensure high productivity.

The DCS3700 is designed to meet the storage needs of highly scalable, data streaming applications in high-performance computing environments. *DCS* stands for *Deep Computing Storage*. IBM has the IBM System Storage DCS9900 for large enterprise deployments, so the smaller DCS3700 is targeted at mid-range deployments. When combined with the best-in-class IBM General Parallel File System (GPFS™), the DCS3700 storage subsystem can help organizations optimize the flow and management of large file-based data while retaining ease-of-data access. Combining GPFS clustered file management software and the DCS3700 creates a scalable and dense file-based management system.

The DCS3700 is the first mid-range dual active storage subsystem, with two intelligent array controllers and up to 180 or 360 drives (with Performance Module Controllers). The machine type for DCS3700 storage is 1818-80C. The machine type for the IBM System Storage DCS3700 expansion unit is 1818-80E. The new DCS3700 model provides a number of new capabilities with, at time of the writing of this book, the latest firmware (Version 07.83.22.00):

- ▶ The DCS3700 supports attachment for SAS and Fibre Channel host connections.

**Note:** A new Performance Module was announced after we wrote this book. With the new Performance Module, hosts can connect to the DCS3700 with a 6 Gb SAS HIC or a 10 Gb iSCSI HIC, which requires firmware Version 07.84 or later.

For more information, see the following website:

<http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=AN&subtype=CA&htmlfid=897/ENUS112-223&appname=USN>

- ▶ This new generation combines the DS3000, DCS3700, and DS5000 Storage Manager and firmware releases, allowing for a common management console to support the entry and mid-range DS families.
- ▶ The Dynamic Disk Pooling feature enables easy to configure Worry-Free storage, which reduces maintenance requirements.
- ▶ The Thin Provisioning, ALUA, VAAI, and IBM Enhanced FlashCopy® features deliver increased usage, higher efficiency, and performance.
- ▶ The 6 Gbps SAS technology for host and drive attachments and 8 Gbps Fibre Channel Protocol (FCP) allows the DCS3700 to be managed by an IBM SAN Volume Controller or IBM Storwize® V7000.
- ▶ Try and Buy Licensing for Enhanced FlashCopy and Enhanced Remote Mirroring.
- ▶ Support for greater capacity with new, larger capacity SAS drive offerings.
- ▶ Support for low-power, highly reliable, and high-performance solid-state drives (SSDs).

**Note:** You can install up to 24 SSD drives in the DCS3700.

Figure 1-3 show the front view of both the chassis and expansion models.



Figure 1-3 DCS3700 and EXP3700 subsystem assembly from the front view

## 1.2.1 DCS3700 components

The DCS3700 storage server is a 4U rack-mountable enclosure, containing two RAID controller modules, two power supplies, and up to 60 disk modules in one enclosure. A minimum of 20 drives are required in each enclosure.

For the component layouts, see Figure 1-4.

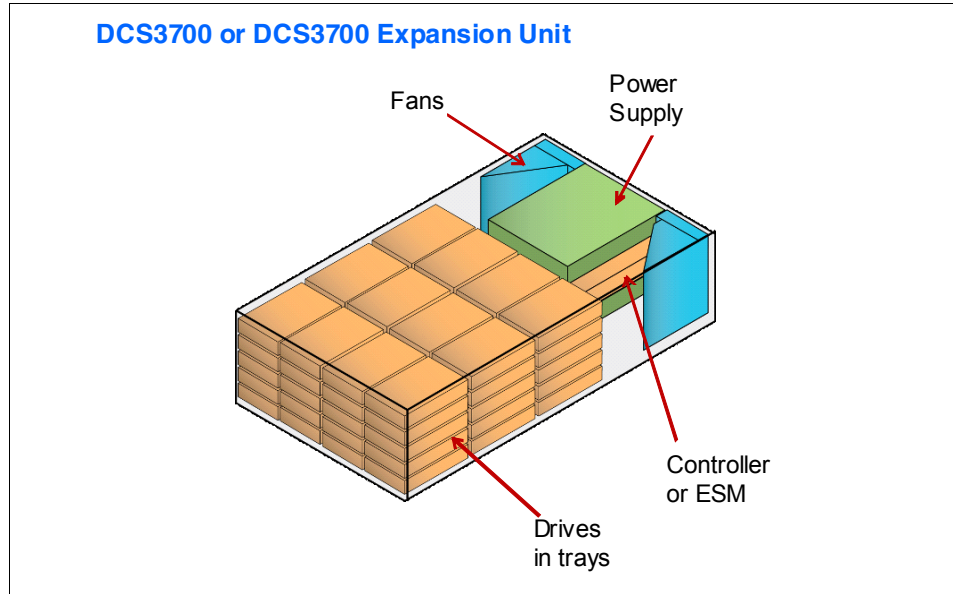


Figure 1-4 DCS3700 and DCS3700 expansion unit components

### RAID controllers

With DCS3700, IBM provides two different RAID controllers:

- ▶ Dual-active, intelligent array controllers with 2 GB (upgradeable to 4 GB), providing 8 GB of total cache for the subsystem. These controllers support up to 180 disk drives in three enclosures. These controllers have connections for the following ports built into them:
  - Two 6 Gbps SAS host server attachment ports
  - Drive side 6 Gbps SAS expansion ports
  - Ethernet Management ports
  - Serial management port
- ▶ Performance Module Controllers, dual-active, intelligent array controllers with up to 24 GB, providing 48 GB of total cache for the subsystem. The controllers support up to 360 disk drives in six enclosures. These controllers have connections for the following ports built into them:
  - 4-Port FC 8Gb/s Host connector on base board
  - 10 GB iSCSI expansion ports<sup>1</sup>
  - Drive side 6 Gbps SAS expansion ports
  - Ethernet Management Ports
  - Serial management port

RAID controllers support RAID levels 0, 1, 3, 5, 6, 10 and, from firmware release Version 7.83.xx.xx onward, Dynamic Disk Pools (DDP). The battery provides power if the cache needs to be destaged to the SD flash card if power is disrupted.

<sup>1</sup> Compatible with firmware release Version 7.84 or later.

Dual controller configurations are the only supported configuration on the DCS3700. The upper controller is controller A and the lower controller is controller B when viewed from the rear of the subsystem, as shown in Figure 1-5. Dual controller configurations offer redundant access to disk storage. If there is a controller or I/O path failure, the other controller continues to provide access to disk drives.

The RAID controllers, two redundant power supplies, and two fan modules are installed in the rear of the subsystem, as shown in Figure 1-5.




Figure 1-5 DCS3700 rear view

As shown in Figure 1-5, the controller modules are in the upper half of the subsystem and the power supply modules are in the lower half.

Figure 1-6 shows the controller module in detail.

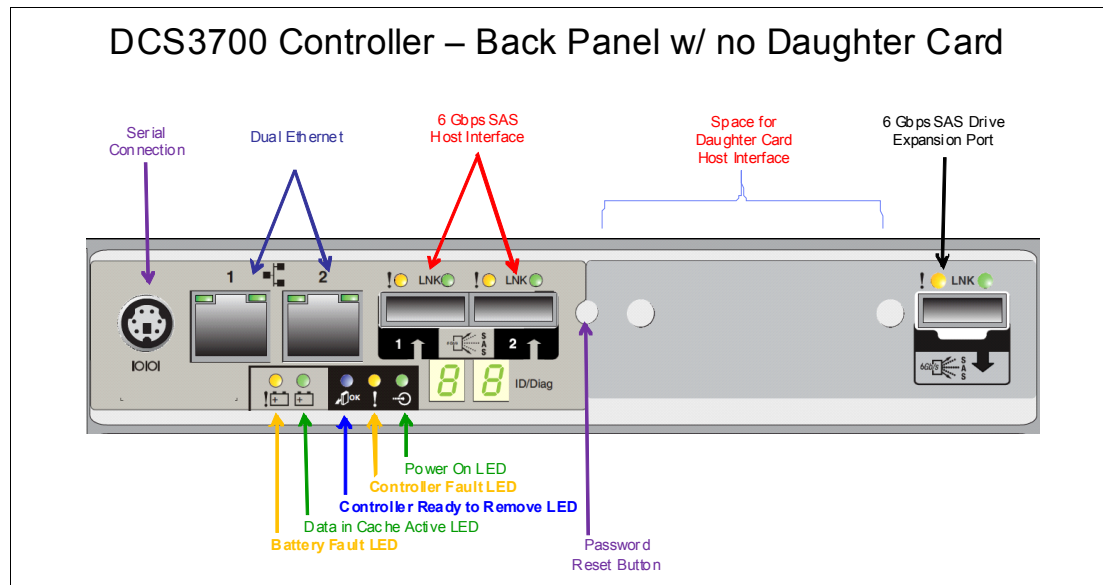


Figure 1-6 DCS3700 controller rear view

### Power supply

The DCS3700 power supply module is a 1755 watt DC power supply. It is auto ranging and only 200 - 240 VAC input capable.

As shown in Figure 1-7, the power supply provides LED indicators for the following states (starting from the left):

- ▶ Standby power LED (green): Currently, this LED is not used.
- ▶ DC power LED (green): When this LED is lit, it indicates that the DCS3700 is turned on and is supplying both 5 V and 12 V 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




Figure 1-7 Power supply LED identification



## Host interface cards

As mentioned earlier, the DCS3700 can be ordered with two different controller types. Depending on which ones you order, you can add different host interface cards. Additional host server connectivity is supported through the usage of an optional daughter card (shown in Figure 1-8). This interface card can provide support for one of the following ports to be added to the DCS3700:

- ▶ For standard controllers:
  - An additional four SAS ports
  - Eight FC ports (four per controller)
- ▶ For Performance Module Controllers:
  - Eight FC ports (4-port 8Gb FC ports per controller)
  - Dual port 10 Gb iSCSI
  - An additional four SAS ports



Figure 1-8 Example host interface daughter card module

Both controllers must be equipped with the same daughter card option to enable the support of the controller failover functions.

Figure 1-9 shows the SAS optional daughter card that is installed in the controller. With this option, the subsystem has up to eight 6 Gbps SAS connections for host attachments. For more information about the cabling and usage of this configuration with the BladeCenter and stand-alone environments, see 3.8, “Host attachment” on page 116.




Figure 1-9 Controller module with an optional SAS host interface daughter card

Figure 1-10 on page 9 shows the Fibre Channel optional daughter card that is installed in the controller. With this option, the subsystem has up to eight 8 Gbps Fibre Channel connections for host attachments. For more information about the cabling and usage of this configuration with the BladeCenter and stand-alone environments, see 3.8, “Host attachment” on page 116.




Figure 1-10 Controller module with an optional FC host interface daughter card

Figure 1-11 shows the 10 Gb iSCSI optional daughter card that is installed in the DCS3700 Performance controller. With this option, the subsystem has up to four 10 Gb iSCSI connections for host attachments. For more information about the cabling and usage of this configuration in your environment, see 3.8, “Host attachment” on page 116.



Figure 1-11 Controller module with an optional iSCSI host interface daughter card

**Note:** Only one type of optional interface can be added to any one DCS3700 storage server. Mixing interface daughter cards between controllers in the same DCS3700 is not supported.

Figure 1-12 shows the Fibre Channel optional daughter card that is installed in the Performance Module controller. With this option, the subsystem has up to sixteen 8 Gbps Fibre Channel connections for host attachments. For more information about the cabling and usage of this configuration with the BladeCenter and stand-alone environments, see 3.8, “Host attachment” on page 116.

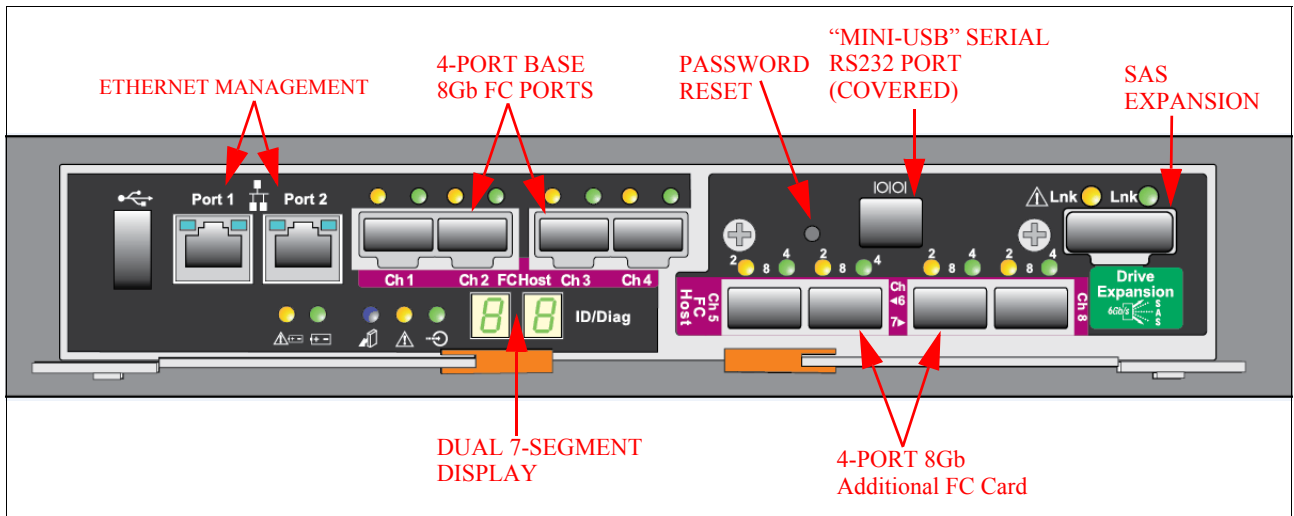


Figure 1-12 Performance Module Controller

## Disk drives

The main advantage of the DCS3700 enclosure is that it supports up to 60 drives in 4U of space. The disk drives are installed at the front in five drawers, as shown in Figure 1-3 on page 4. The available drive types for each of these subsystems at the time of the writing of this book are shown in Table 1-1.

Table 1-1 DCS3700 families HDD support

Drives supported	DCS3700
SAS 15 K RPM	300 GB
Nearline SAS 7.5 K RPM	2 TB, 3 TB
Maximum drives	180 or 360
Storage system capacity (max)	108 TB SAS / 540 TB SATA or 216 TB SAS / 1080 TB SATA

For the latest list of the drives that are supported by the DCS3700, refer to the specifications page at the following link:

<http://www-03.ibm.com/systems/storage/disk/dcs3700/specifications.html>

The population order in the drawer is shown in Figure 1-13.

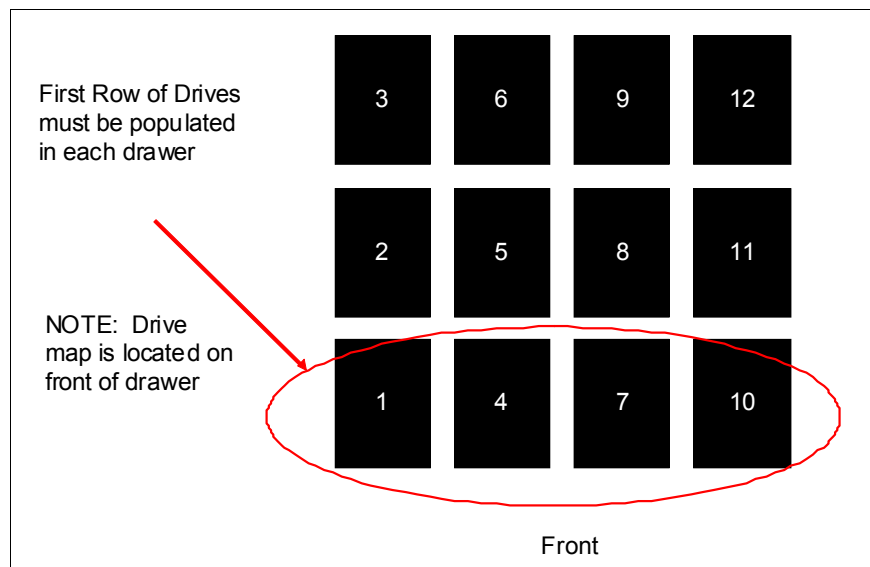


Figure 1-13 Drive orientation/population

## 1.3 IBM System Storage DCS3700 expansion unit

The IBM System Storage DCS3700 expansion unit allows the DCS3700 storage subsystem to grow up to the 180 or 360 drive maximum (depending on the controller type). The IBM System Storage DCS3700 expansion unit differs from the DCS3700 in that instead of the controller module, they are equipped with an Environmental Services Module (ESM). Each ESM has a 6 Gbps SAS connection (with four lanes per port) providing 2400 MBps throughput.



Figure 1-14 shows a view of the IBM System Storage DCS3700 expansion unit ESM with its port connections for cabling.




Figure 1-14 DCS3700 expansion unit ESM

With the IBM System Storage DCS3700 expansion unit, only one of the two IN ports are used on each ESM to connect expansions together in to a cascaded loop configuration. As shown in Figure 1-15, the cabling scheme that is 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 correct RAID layout, this can provide for uninterrupted operations.




Figure 1-15 DCS3700 expansion unit cascaded loop configuration

## 1.4 Premium Features

Standard configurations of DCS3700 storage subsystems come with the following Premium Features.

### 1.4.1 Partition Expansion License

As part of the standard configuration, 128 storage partitions are enabled on the DCS3700 and 512 storage partitions on the DCS3700 Performance product. This requires Firmware Code Level Version 7.84 or higher.

### 1.4.2 FlashCopy expansion license

As part of the standard configuration, 32 Enhanced FlashCopy images are enabled. This requires Firmware Code Level Version 7.84 or higher.

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 usage of a defined *FlashCopy repository*, which contains the original content of the data. 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 is lost as well. For more information about FlashCopy, see *IBM System Storage DS Storage Manager Copy Services Guide*, SG24-7822.

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

### 1.4.3 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 that is 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 remains available if the source logical drive is lost. For more information about VolumeCopy, see *IBM System Storage DS Storage Manager Copy Services Guide*, SG24-7822.

**Important:** FlashCopy is a prerequisite for VolumeCopy.

### 1.4.4 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, which is a bundle of both Premium Features.

## 1.4.5 Try and Buy Premium Features

With Firmware Code Level V7.83.xx.xx, IBM introduced the Trial Version feature for Premium Feature. A trial version of a premium feature allows you to use it for a limited amount of time and for a limited allowable value so that you can try it before you buy it. Each trial version of a premium feature requires a trial license. You can obtain a trial license for one or more of the following premium features:

- ▶ Enhanced FlashCopy

An Enhanced FlashCopy image is a logical image of the content of an associated base logical drive that is created at a specific moment. An Enhanced FlashCopy image can be thought of as a restore point. It is not directly read or write accessible to hosts because the Enhanced FlashCopy image is used to save only the data that is captured from the base logical drive.

You must create an Enhanced FlashCopy logical drive to give the host access to a copy of the data that is contained in the Enhanced FlashCopy image. The Enhanced FlashCopy logical drive contains its own repository, which is used to save any subsequent modifications that are made by the host application to the base logical drive without affecting the referenced Enhanced FlashCopy image.

- ▶ Enhanced Remote Mirroring

With the DCS3700, 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 that is in a separate location that is some distance from the primary location. ERM provides three mirroring modes from which to choose:

- Synchronous: Provides mirroring capability in metro (campus) environments, generally within a 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 global environments, generally over 10 miles in distance, where latency might exceed acceptable response times for the primary host I/O operation. I/O operations are handled in the order of their original request across the selected logical drives that are grouped through the usage of a consistency group.
- Asynchronous copy: Provides copying capability in global environments, generally over 10 miles in distance, where latency is high and bandwidth is limited, which causes a high backup of I/O that might impact the principle operation site. I/Os in this mode are not ensured to be processed in the order they are received across a number of logical drives.

This capability is shared across the DS3000, DCS3700, and the DS5000 mid-range products, which allows the remote location to be a consolidated collection site.

For more information about Enhanced Remote Mirroring, see *IBM System Storage DS Storage Manager Copy Services Guide*, SG24-7822.

## 1.5 DCS3700 and DS3700P comparison

Table 1-2 shows the comparison of DCS3700 and DCS3700P (with Performance Module Controllers).

Table 1-2 DCS3700 and DCS3700P specification comparisons

Fully configured systems	DCS3700	DCS3700P
Host interfaces	<ul style="list-style-type: none"> <li>▶ Two 6 Gbps SAS host ports per controller standard</li> <li>▶ Two 6 Gbps SAS ports per optional host interface card</li> <li>▶ Four 8 Gbps Fibre Channel ports per optional HIC.</li> </ul>	<ul style="list-style-type: none"> <li>▶ Four 8 Gbps Fibre Channel host ports per controller</li> <li>▶ Four 8 Gbps Fibre Channel ports per optional HIC</li> <li>▶ Four 6 Gbps SAS ports per optional HIC</li> <li>▶ Two 10 Gbps iSCSI ports per optional HIC</li> </ul>
Redundant drive channels	Two 6 Gbps SAS	Two 6 Gbps SAS
Drive enclosures	3	6
Max drives	180	360
Drives supported	SAS, NL SAS, and SSD	SAS, NL SAS, and SSD
XOR engine	Integrated XOR / P+Q	Integrated XOR / P+Q
Dual Ethernet	Yes	Yes
Data cache / memory	4 or 8 GB	12, 24, or 48 GB
Max logical drives	512	2048
Max LUN size	> 2 TB	> 2 TB
RAID levels	1, 3, 5, 6, 10, or DDP	1, 3, 5, 6, 10, or DDP
Enhanced FlashCopy copies (PiTs)	512	2048
Volume Copy	511	1024
Enhanced Remote Mirror	16	16
Enhanced Global Mirror	32	128
Max Thin Provisions	512	2048
Maximum remote mirrors (across FC)	8	64

## 1.6 IBM System Storage DS Storage Manager

The new DCS3700 uses the DS Storage Manager, which is similar to DS Storage Manager that is used for the DS3000, DS4000, and DS5000 product lines. At the time of the writing of this book, the current version of the software was Version 10.83.G5.18. When you receive your DCS3700 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 that you receive is not compatible with your DCS3700 firmware release, you can download it from the IBM support website, found at:

<http://www.ibm.com/support/entry/portal/overview>

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 Copy Services and Enhanced Remote Mirroring, and capturing logs for troubleshooting. This section describes the IBM System Storage DS Storage Manager software and its usage in great detail later in this book, so this section briefly describes some of the main points.

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

► Host server

This is a server that is attached to the DCS3700 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 DCS3700 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 the 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 DCS3700 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 DCS3700 RAID controller. All management communication flows across the TCP/IP connection between the management station and the DCS3700. This method also is called *direct-attached management*. The management station in this case requires only an Ethernet connection to the DCS3700.

– In-band

This method uses the I/O path between a host server and the DCS3700. In this case, the management station does not have direct TCP/IP connection to the DCS3700, but rather communicates with the DCS3700 through an HBA that acts as a gateway to the DCS3700 storage subsystem. Communication between the management station and the host server is across the Fibre Channel or SAS I/O path.

This method also is called *host-attached management*.

Because each method has associated advantages and disadvantages, this section describes 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.

## 1.6.1 In-band management

The in-band management method uses the I/O path between the host server and the DCS3700 to transfer management commands and information.

This method does not use the management Ethernet ports on DCS3700 RAID controllers and does not require a management Internet Protocol network. However, it does require a special *access logical drive* to manage the DCS3700 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 not a problem because virtually all customers find the maximum number of logical drives more than sufficient.

An example of in-band management is shown in Figure 1-16. Two host servers are attached to the DCS3700 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 DCS3700 across the FC I/O path.

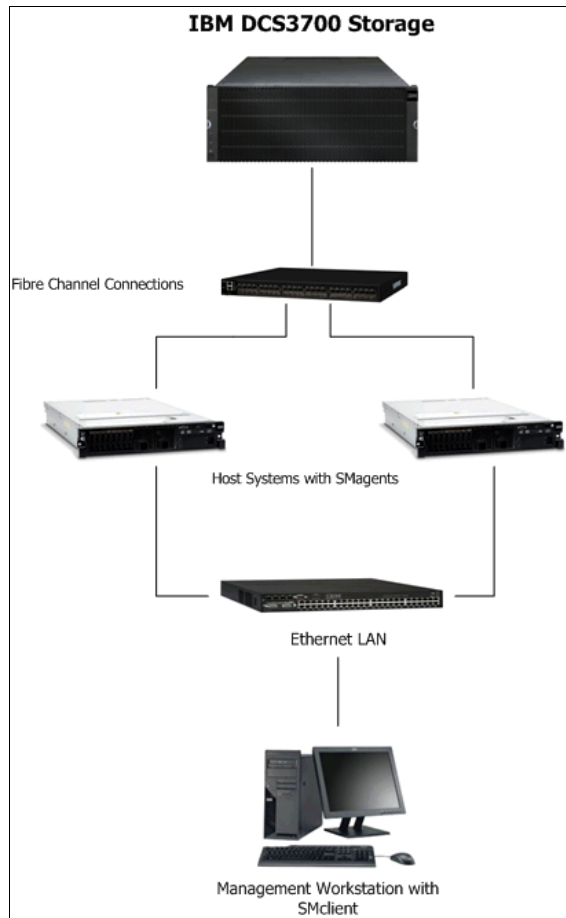


Figure 1-16 In-band management

### Access logical drive

The access logical drive exists in each storage partition by default (no manual configuration is required). This is not a real logical drive, although it is presented to the operating system as a drive on LUN 31. To allow for in-band management communication across the I/O path, you 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.

## 1.6.2 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 1-17 shows the management workstation and the DCS3700 subsystem that is connected on the Ethernet management network.




Figure 1-17 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 that are supported by the host servers' operating system.
- ▶ If the I/O paths fail, you can still access the DCS3700 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 setup introduces redundancy to your management setup and provide management access to the DCS3700 subsystem even if there is 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)

### 1.6.3 Storage Manager Client (SMclient)

This is the actual graphical user interface (GUI) that you use to manage the DCS3700. You install SMclient on the management station. This is 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 DCS3700 is started. It lists the storage subsystems that it knows about. You can add new storage subsystems and perform various tasks on the enterprise level. A sample of the Windows based Enterprise Management window is shown in Figure 1-18.

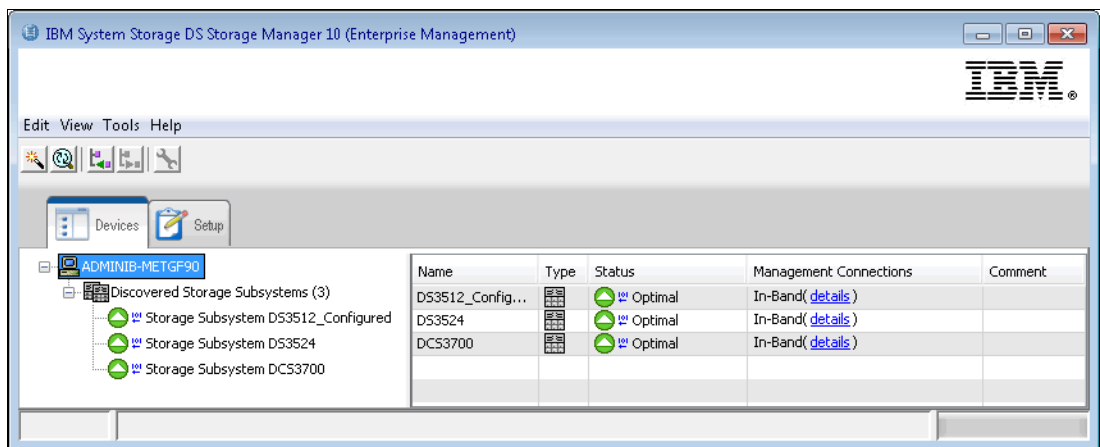


Figure 1-18 Enterprise Management window



## Subsystem Management window

This window allows you to manage a particular DCS3700. Management tasks such as creating arrays, logical drives, storage partitions, FlashCopy, and VolumeCopy are all performed from within the Subsystem Management window. For an example of this window for the Windows based version, see Figure 1-19.

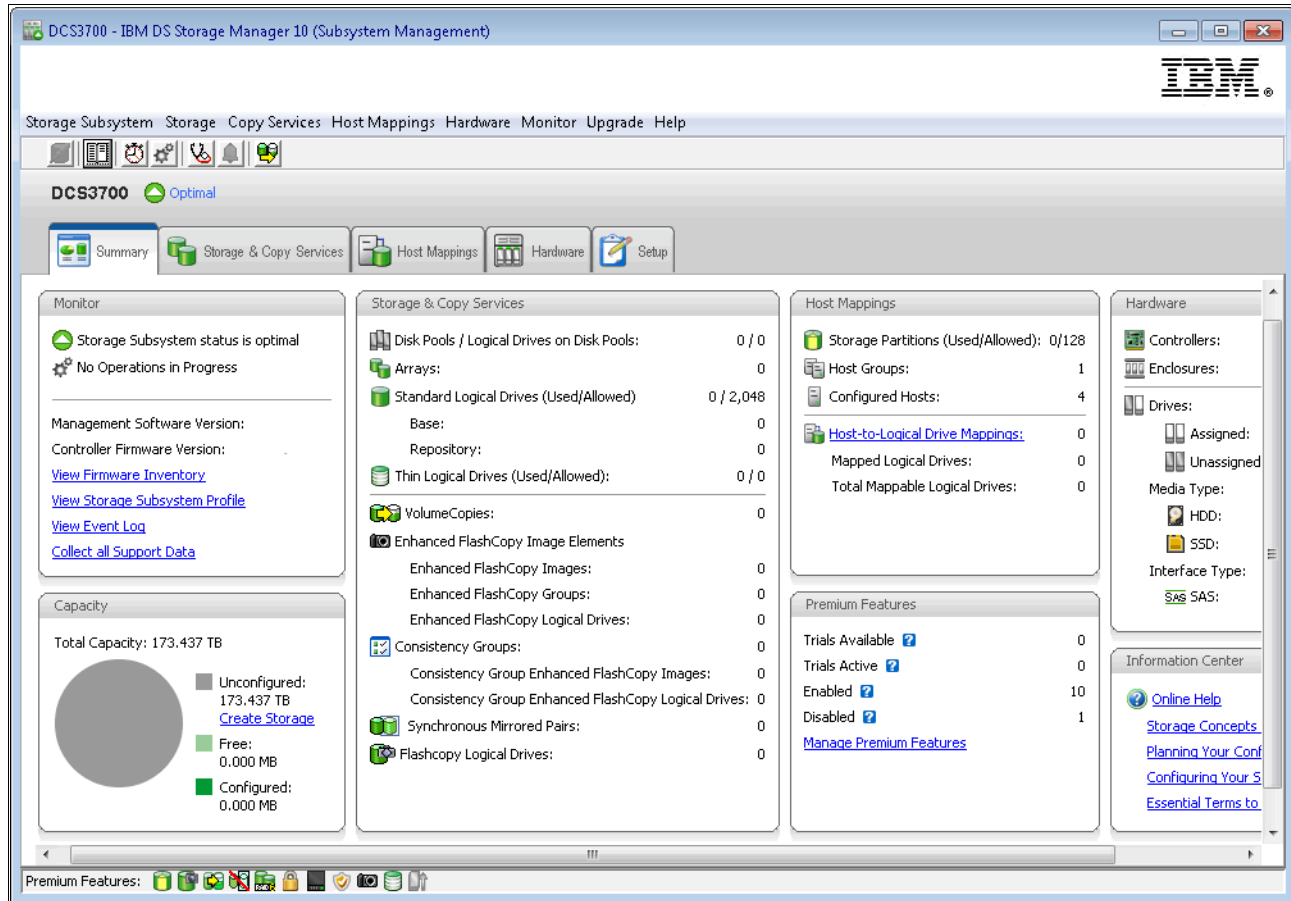


Figure 1-19 Subsystem Management Window

Install SMclient on the management station. This station is 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 that is called SMcli. SMcli provides a powerful command-line interface (CLI). All the tasks available in the Storage Manager GUI can also be run by 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 DCS3700 Storage Manager, which is started from the Enterprise Management window.

You can either run a single command at a time or run pre-written scripts. For more information about the CLI and Script Editor, see Chapter 15, “Command-line interface (CLI)” on page 463.

## 1.6.4 Storage Manager Agent (SMagent)

SMagent is an optional component that is required only for in-band management. SMagent is installed on the host server and allows the SMclient to communicate with the DCS3700 across the I/O path. At the time of the writing of this book, only the FC I/O path is supported for management communication.

## 1.6.5 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 that is 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 that is used with FlashCopy and VolumeCopy. The utility allows you to flush cached data before you create 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 can be found about the current configuration settings and their impacts.

## 1.6.6 Multipath driver support

In general, install two HBAs on the host servers and use the dual controller DCS3700 subsystems. This action provides you with the ability to create 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 back end to the disks. However, host dual path configurations can work correctly only if you install the appropriate multipath driver for your particular operating system in to the host server.

The IBM System Storage DS Storage Manager Software and Host Kit for Windows includes multipath support that is 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 multipath support. In AIX, there is an MPIO driver that is 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 alternative path and an application outage.

## **1.6.7 Java Access Bridge**

Java Access Bridge 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 DCS3700 storage subsystem planning tasks

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

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

Having a thought-out design and plan before the implementation helps you get the most out of your investment for the present and protect it for the future.

During the planning process, you must 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 subsystem. 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.

## 2.1 Planning your SAN and storage server

When you plan 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 - 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 you design 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 that 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 types to be used: SAS or Fibre Channel (direct or fabric), iSCSI.
  - Host Bus Adapter (HBA) selection.
  - HBA driver type selection: SCSIPort or StorPort.
  - Multipath Driver selection: RDAC, 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 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 do the following tasks:
  - Make an inventory of existing hardware infrastructure. Ensure that any existing hardware meets the minimum hardware requirements and is supported by the DCS3700 storage subsystem.
  - Make a complete list of the planned hardware requirements.
  - Ensure that you have enough rack space for future capacity expansion.
  - Assess the weight capacity of your raised floor according to your configuration options:
    - DCS3700 storage subsystem: 46.0 kg or 101.4 lbs (Empty) and 102.3 kg or 225.5 lbs (fully configured)
    - DCS3700 expansion unit: 46.0 kg or 101.4 lbs (Empty) and 104.1 kg or 222.6 lbs (fully configured)
  - Ensure that the power and environmental requirements are met.
  - Ensure that your existing network of SAS, Ethernet, or Fibre Channel switches and cables are correctly 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 do the following tasks:
  - Ensure that the existing versions of firmware and storage management software are up to date.
  - Ensure that host operating systems are supported by the storage subsystem. For more information, check the IBM System Storage Interoperation Center (SSIC) for your specific storage subsystem, which is available at this website:  
<http://www.ibm.com/systems/support/storage/ssic/interoperability.wss>

These steps 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 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 DCS3700 storage subsystem 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 expected growth is vital for the capacity planning of your DCS3700 storage subsystem 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 correct DS/DCS storage subsystem 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 DCS3700 platforms offer scalability and expansion flexibility. Premium features can be purchased and installed later 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.

### 2.1.1 SAN zoning for the DCS3700 storage subsystem

Zoning is an important part of integrating a DCS3700 storage subsystem in a SAN. When done correctly, it can eliminate many common problems. Zoning also helps with creating paths that can be used by the multipath drivers for better failover protection. Understanding the capabilities of the multipath driver is important when designing the paths it uses.

**Important:** Disk and tape must be on separate HBAs, so the disk and tape access are also in separate zones. *HBA sharing between disk storage and tape is not recommended.*

#### 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 using multiple HBA ports or devices on the host server that is connected to SAN fabric switches, which are also connected to the multiple ports on the storage devices. The DCS3700 has two controllers within the subsystem that manage and control the disk drives. These controllers behave in an active/active fashion called Asymmetric Logical Unit Access (ALUA). For more information about ALUA, see 2.7.10, “Asymmetric logical unit access” on page 94.

Each controller can have more than one fabric port for connectivity to the SAN fabric. This helps provide faster recovery for path failures that can be resolved through the usage of an alternative connection to the same controller rather than requiring the full failover process to be run. The DCS3700 models support only drivers that by design follow the rules for this type of path management. For an example of how the configuration should be implemented, see Figure 2-1 on page 27.



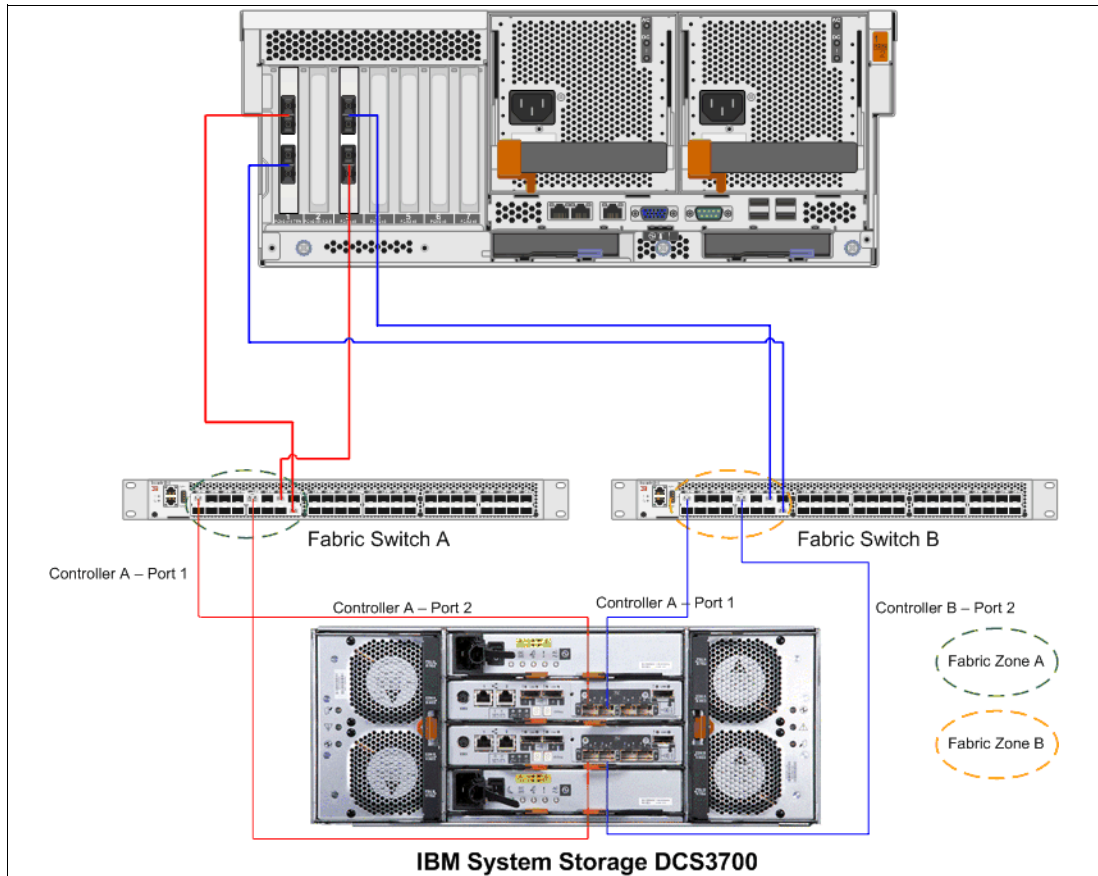


Figure 2-1 Host HBA to storage subsystem controller multipath sample configuration

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 Figure 2-1, to accomplish this task, 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 DCS3700 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. It is also named as Single Initiator Zoning

## 2.1.2 Zoning considerations for Enhanced Remote Mirroring

Enhanced Remote Mirroring (ERM) is now supported across Fibre Channel (FC) or IP host interfaces. When using an IP host interface for ERM, the port can be shared with normal traffic. Direct IP connection is not possible, as you need the name services from the switches.

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

When using FC ERM, you must create two additional zones:

- ▶ The first zone contains the ERM source DCS3700 controller A and ERM target DCS3700 controller A.
- ▶ The second zone contains the ERM source DCS3700 controller B and ERM target DCS3700 controller B.

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

## 2.2 Planning for physical components

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

### 2.2.1 Rack considerations

The DCS3700 storage subsystem and possible expansions are mounted in rack enclosures.

#### General planning

Consider the following general planning guidelines. Determine the following items:

- ▶ The size of the floor area that is required by the equipment:
  - Floor-load capacity
  - Space that is 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 the service clearances that are required for each rack or suite of racks.
- ▶ If the equipment is on a raised floor, determine the following items:
  - The height of the raised floor
  - Anything that might obstruct cable routing
- ▶ If the equipment is not on a raised floor, determine the following items:
  - 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 the following items:
  - 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 the 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 correctly designed and protected areas.

### **Rack layout**

To be sure that you have enough space for the racks, create a floor plan before you install the racks. You might need to prepare and analyze several layouts before you choose 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 following items:
  - 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 the 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.

**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 DCS3700 storage subsystems that are in a chassis are 4U in height.

Figure 2-2 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 2-2 9306 enterprise rack space requirements

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

## 2.2.2 SAS cables and connectors

All DCS3700 storage subsystems, regardless of the optional host attachment feature that is added, will have SAS host connections available. Also, with the addition of any expansion subsystems, they use SAS cabling to attach expansion units. Therefore, let us consider the SAS cables and connectors that are used.

The SAS ports on the DCS3700 controller and DCS3700 Expansion ESM all support mini-SAS 4x multilane connectors. SAS cables with mini-SAS connectors that fit in these ports are required, as shown in Figure 2-3 on page 31. IBM provides SAS cables in three cable lengths: 1, 3, and 8 meters.

**Caution:** Longer SAS cables (at least 6 m) are nonstandard. Alternatively, the DCS3700s can be configured with an FC daughter card where longer cables are standard.




Figure 2-3 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 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.

Figure 2-4 shows an example of SAS connections being made to direct attached hosts and a BladeCenter through an internal SAS switch module. The DCS3700 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 2-4 DCS3700 with optional added SAS host connections

### 2.2.3 Ethernet cable and connections

Each DCS3700 RAID controller contains an Ethernet management port that you can use for out-of-band management. If you have a dual controller DCS3700 subsystem, make sure that the management workstation can access the management port on each controller. If only one controller is accessible by the management machine, the DCS3700 Storage Manager will not be able to manage the enclosure.

Avoid using your public LAN for DCS3700 out-of-band management. Instead, set up a dedicated LAN or VLAN just for management purposes. This provides increased security for your DCS3700 storage subsystem. If the DCS3700 RAID controllers are on a public LAN, a knowledgeable user could install the DCS3700 Storage Manager on a separate workstation, or use the CLI to perform potentially destructive tasks. For an additional layer of security, also enable password protection on the DCS3700 storage subsystem. For more information, see “Security: Set Password” on page 208.

## 2.2.4 Fibre Channel cables and connectors

This section describes various essential characteristics of fiber cables and connectors. This information can help you understand the options that you have for connecting and cabling the DCS3700 storage subsystem.

### Cable types (shortwave or longwave)

Fiber cables are 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 use the signal on the receiving end. This dispersion (called modal dispersion) is the critical factor in determining the maximum distance that 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 that there 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, so there is no “dispersion confusion” at the receiver. However, the concern with single mode fiber is attenuation of the signal. Table 2-1 lists the supported distances for Optical Mode-1 (OM-1)

Table 2-1 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

Fiber type	Speed	Maximum distance
50 micron MMF (shortwave)	8 Gbps	50 m
50 micron MMF (shortwave)	16 Gbps	35 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
62.5 micron MMF (shortwave)	16 Gbps	15 m

The “maximum distance” shown in Table 2-1 on page 33 is a maximum. Low quality fiber, poor terminations, excessive numbers of patch panels, and so on, can cause these maximums to be far shorter.

Table 2-2 shows the maximum distances for other optical modes.

*Table 2-2 Maximum distances for all optical modes*

Speed	Wavelength (nm)	Maximum Channel Length (meters)			
		OM1	OM2	OM3	OM4
4 Gbps	850	70	150	380	400
8 Gbps	850	21	50	150	200
16 Gbps	850	15	35	100	130

All IBM fiber feature codes that are orderable with the DCS3700 storage subsystem meet these standards.

### **Interfaces, connectors, and adapters**

In Fibre Channel technology, frames are moved from source to destination by using gigabit transport, which is needed 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), as shown in Figure 2-5. 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 2-5 Small Form Pluggable (SFP) with LC connector fiber cable

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

### LC connector

Connectors that plug in to 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 DCS3700 series products use SFP transceivers and LC fiber cables. Figure 2-6 shows an LC fiber cable connector.




Figure 2-6 LC fiber cable connector

**General rule:** When you are not using an SFP, it is best to remove it from the port on the DCS3700 storage controller and replace it with a cover. Similarly, unused cables must be stored with their ends covered to help eliminate the risk of dirt or particles contaminating the connection while it is not in use.

## Interoperability of 2 Gbps, 4 Gbps, 8 Gbps, and 16 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 runs 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 can automatically negotiate back down to either 4 Gbps or 2 Gbps, depending upon the attached device and the link quality. A 4 Gbps device can 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 DCS3700 storage subsystem 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 subsystem that are commonly used: direct-attached and switch-attached (fabric). A mix of direct and switched attached is possible.

## Direct-attached DCS3700

Figure 2-7 shows a simple direct-attached configuration. Two host servers are connected to a dual-controller DCS3700. 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 still has access to logical drives. This type of setup is suitable for a two-node Microsoft Cluster Server configuration.




Figure 2-7 Host servers that are attached to DCS3700 FC direct connect

## Switch-attached DCS3700

Because the DCS3700 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 2-8 displays a sample configuration with four host servers that are attached to a dual controller DCS3700. Each host has two FC HBAs for redundant I/O path support.




Figure 2-8 DCS3700 that are 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 be used only 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 back-end expansions, and Fibre Channel. Cable management and labeling needs have expanded from the traditional labeling of network connections to the management and labeling of most cable connections between your servers, disk subsystems, multiple network connections, and power for all the subsystems that are involved in your solution. In many cases, components of the solution have different cable needs. Multiple unique systems could be in the same rack, or a system could span across multiple racks. In certain cases, the solution could be configured where components are not 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 that is completed before your solution is installed results in a reduced chance of installation problems. Significant attributes that are 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 you install cables to devices that are mounted on sliding rails, consider the following items:
  - 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. For guidance about the minimum radius, handling, and care of fiber optic cables, see the instructions that are included with your 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 that 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 using cable management arms.
    - When attaching a cable 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 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 in to 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.

More 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 that provide 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 that 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 a 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 reuse dust caps that have been left loose in the rack or computer room.

## **2.2.5 Fibre Channel adapters**

This section reviews topics that are 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 that is available.

**Important:** Do not place all the high-speed Host Bus Adapters (HBAs) on a single system bus. If you do, the computer bus becomes 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 that you can install depends on the number of PCI slots that are 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 whether all the available paths to the storage server are still functioning. If there is 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 DCS3700 storage subsystem. Microsoft Windows 2003 and 2008, and Linux support up to four paths to the storage controller. The latest VMware version supports up to four paths as well. AIX with MPIO can support more than four paths to the controller modules. However, too many paths can delay failover when it is needed because of a controller failure. Take care to ensure that there are enough redundant paths without having too many paths. 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 DCS3700 controllers A and B alternatively (for more information, see 2.4.8, “Logical drives and controller ownership” on page 66).

Figure 2-9 shows the principle for a load-sharing setup. A multipath driver checks all the available paths to the controller. In Figure 2-9, there are 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 2-9 Load sharing approach for multiple HBAs

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

**Note:** In a cluster environment, you need a single path to each of the controllers (A and B) of the DCS3700 storage subsystem. However, if the cluster software and host application can do persistent reservations, you can keep multiple paths and the multipath driver routes the I/O request by 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 use only 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 a *device busy* status is returned by the system 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 subsystem, 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 DCS3700 controller, see the following website:

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



## 2.2.6 Disk expansion units

The DCS3700 storage subsystem offers the DCS3700 expansion units for expanding the subsystem beyond the internal drive count of the DCS3700. These expansions connect to the DCS3700 storage subsystem through the 6 Gbps SAS expansion ports on the controller modules. Up two expansion units can be added. Figure 2-10 and Figure 2-11 on page 44 are examples of the cable connections that are needed to attach the DCS3700 with dual controllers to a DCS3700 expansion box.




Figure 2-10 DCS3700 with DCS3700 expansion unit




Figure 2-11 DCS3700 with dual DCS3700 expansion units

Looking at Figure 2-11, it is important to notice the direction of the cabling scheme that is used when you add the additional expansions to the loops. The DCS3700 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. For more information about array layouts, see 2.4.3, “Array configuration” on page 57.

When you plan for which enclosure to use with your DCS3700, you must look at your applications and the workload type that you need to provide. With transaction-based applications, the workload needs to be able to perform a high number of I/Os, and therefore having the greater number of spindles that are provided by the DCS3700 expansion unit can prove to be a great advantage. With throughput-based applications, the need is generally for higher capacity drives and more bandwidth.

For this environment, the DCS3700 expansion unit can provide the disk capacity at a lower price point. Careful planning can allow for the mixing of these two expansions in the storage subsystem to meet a mixed environment need.

The DCS3700 expansion unit supports the same disk drives that are supported by the DCS3700 storage subsystem.

### Enclosure IDs

Each controller and ESM contain two 7-segment numeric LED displays. These LEDs provide enclosure identification and diagnostic data.

The two digits that comprise the enclosure ID are referred to as x10 and x1 digits. The enclosure ID provides a unique identifier for each enclosure in the storage subsystem.

The storage management software automatically sets the enclosure ID for each controller. You can change the enclosure ID setting through the Storage Manager software only. There are no switches on the enclosure to manually set the enclosure ID. Both controllers or ESMs have enclosure IDs that are identical under normal operating conditions.

### Enclosure guidelines

The controller automatically sets the enclosure ID. You can change the setting through the DS Storage Manager software, if necessary.

The range of supported enclosure ID settings is 0 - 99. The enclosure ID is normally set to a value of 00 at the factory.

The enclosure ID is displayed on the seven-segment numeric display at the back of each controller and ESM, as shown in Figure 2-12.




Figure 2-12 Enclosure ID LEDs on the DCS3700

## 2.3 High Performance Module (DCS3700P)

The new IBM Performance Module features provide new controller canisters for the DCS3700 with increased processor speeds, larger cache memory capacities, and double the scalability of the original offering. Designed for use with HPC computing environments, the DCS3700 with Performance Modules increases the processing capability and storage capacity for HPC uses in the oil and gas industries, media applications, and life sciences that have installations in large corporate data centers and government labs.

Each Performance Module brings four 8 Gb Fibre Channel (FC) host ports, dual gigabit Ethernet ports, and a 6 Gb SAS expansion connection to the DCS3700. A slot is available for additional host interface cards, such as a new 4 port 8 Gb FC host interface card (HIC), 10 GB iSCSI, or SAS HIC. There are three slots for DDR3 memory DIMMs in 2, 4, and 8 GB capacities, for a total of up to 6, 12, or 24 GB of memory for the canister, for a total of up to 12, 24, or 48 GB for the DCS3700 with Performance Modules.

The recently released DCS3700 with Performance Modules (DCS3700+) is being updated with new hardware options, such as dual port 10 Gb iSCSI and 4-port 6 Gb SAS HICs for optional use in the available slot in the Performance Module. The new CFW V7.84 and the new 600 GB and 900 GB 10 K drive options are also being released for the DCS3700+.

In addition to the increased memory scalability, the DCS3700 with Performance Modules has increased scalability in drive capacity, with two 6 Gb SAS ports to redundantly connect to up to five EXP3700 expansion units, doubling the number of attached disks to 360. Using the 3 TB NL SAS drive options yields over 1 PB of raw disk capacity.

Increased scalability applies as well to the Premium Features that bring exceptional functionality. Loaded with the latest code, the DCS3700 with Performance Modules uses Dynamic Disk Pooling to provide thin provisioning and larger volumes. Disk pooling produces more efficient FlashCopy copies, rolls back FlashCopy copies, and consistency groups. The Performance Modules allow double the number of FlashCopy instances, up to 2048. With the Performance Module, a special bundle of FlashCopy and VolumeCopy offers up to 2048 FlashCopy copies and up to 511 volume copies per array.

DCS3700 with Performance Modules is supported for attachment to certain of the following systems:

- ▶ IBM System x, IBM BladeCenter, and IBM Power System servers and blades
- ▶ Intel and AMD processor-based servers

**Note:** For more information, see the IBM System Storage Interoperation Center (SSIC), found at:

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

## 2.4 Planning your storage structure

It is important to configure a storage subsystem in accordance with 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. This section provides a general (and less detailed) performance discussion.

This section also describes the use of RAID protection, and reviews other aspects of the system configuration that can help optimize the storage capacity and resilience of the system. In particular, this section reviews array configuration and enclosure loss protection.

## 2.4.1 Selecting drives

The speed and the type of the drives that are used 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 DCS3700 storage subsystem supports the following types of drives for the two models of chassis, as shown in Table 2-3.

Table 2-3 DCS3700 families HDD support

Drives supported	DCS3700
6 Gbps SAS 3.5-inch drives	2 TB 7.2 K RPM nearline 3 TB 7.2 K RPM nearline
6 Gbps SAS 2.5-inch drives	300 GB 15 K RPM 600 GB 10 K RPM 900 GB 10 K RPM
SAS 2.5-inch SSDs	200 GB SSD 400 GB SSD
Maximum drives	180 - DCS3700 360 - DCS3700P
Storage subsystem capacity (maximum)	1080 TB - DCS3700 1.080 PB - DCS3700P

**General rule:** Generally, it is best to use the fastest drives that are 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 15 K 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 7.2 K drive has an average seek time of around 9 ms, a 10 K drive has an average seek time of around 5.5 ms, and a 15 K drive has 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 reordered, along with the necessary tracking mechanisms for outstanding and completed portions of workload. The DCS3700 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 correctly configured, these drives can push high throughput numbers with large host IO block sizes and sequential workloads.

## 2.4.2 Understanding RAID types

This section introduces arrays, logical drives, and associated terminology, and then describes the various RAID levels that are supported by the DCS3700 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 that is 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 error correcting code that is derived from the data. If a drive fails, you can use the redundancy data to quickly reconstruct information on a replacement drive.

### DCS3700 arrays and RAID levels

An array is a set of drives that the system logically groups to provide one or more logical drives to an application host or cluster. The DCS3700 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 are described in more detail in the following sections of this book.

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 DCS3700 storage subsystem, which is 360 drives. The maximum number of physical drives in a RAID 3, 5, or 6 array is always 30 drives. The 30 drives in an RAID array limitation is specific to DS/DCS products.

The DCS3700 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 DCS3700 storage subsystem supports a maximum of 2048 logical drives for the entire storage subsystem, and 256 logical drives per partition.

This section now briefly describes the main features of each of the various RAID levels that are supported by the DCS3700 storage subsystem.

## RAID 0: Data striping

RAID 0 (Figure 2-13) 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 DCS3700 storage subsystem supports a minimum of one drive and a maximum of 360 drives in a RAID 0 array.

RAID 0 is 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 increases the data access performance of the array.

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

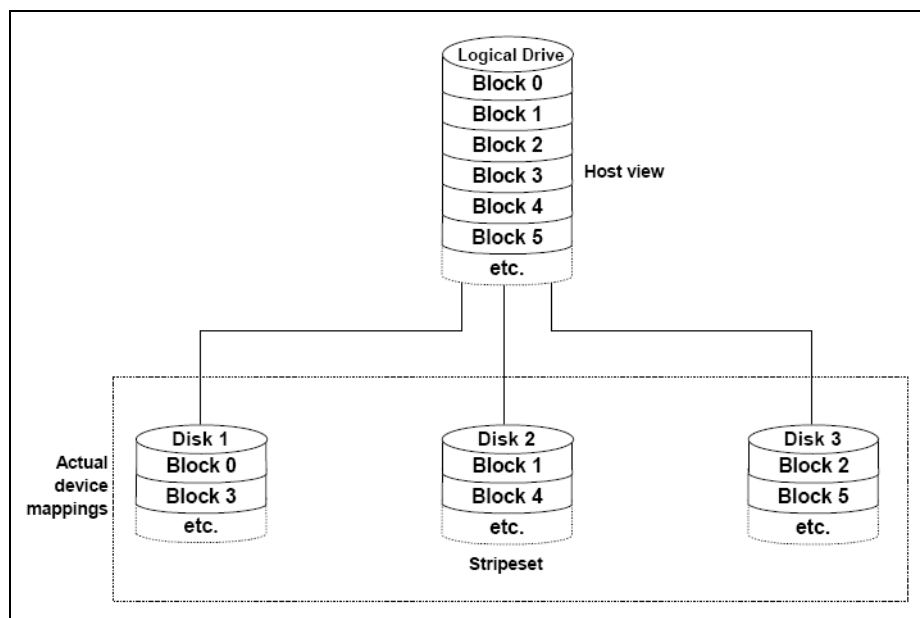


Figure 2-13 RAID 0

## 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 2-14 on page 50) 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 DCS3700 storage subsystem supports a maximum of 360 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 wanted. 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.

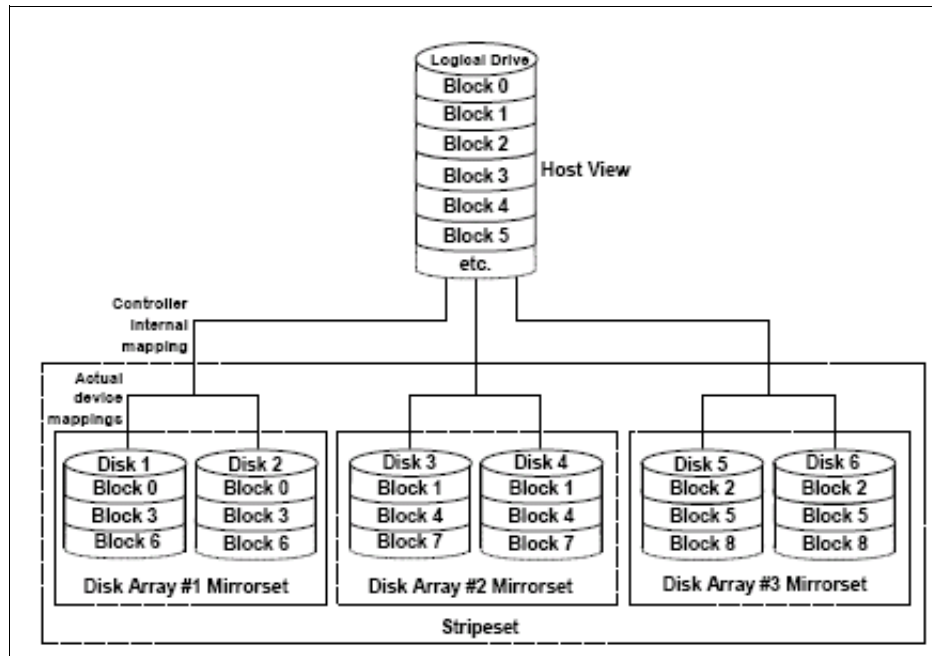


Figure 2-14 RAID 10; a RAID 1 array consists of only one mirrored pair (#1 Mirrorset)

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

- ▶ 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 does not complete until both writes complete on both disks in a mirrored pair, the performance of writes can be improved through the usage of a write cache. Always ensure that you use a 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 10 with RAID 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 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 that is 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 recomputed 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 DCS3700 storage subsystem supports a maximum of 30 drives in a RAID 3 array.

## RAID 5: Data striping with distributed parity

Like RAID 3, RAID 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 2-15. The failure of a single physical drive in a RAID 5 array causes the array and associated logical drives to be degraded, but all the data remains 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 DCS3700 storage subsystem requires a minimum of three and supports a maximum of 30 drives in a RAID 5 array.

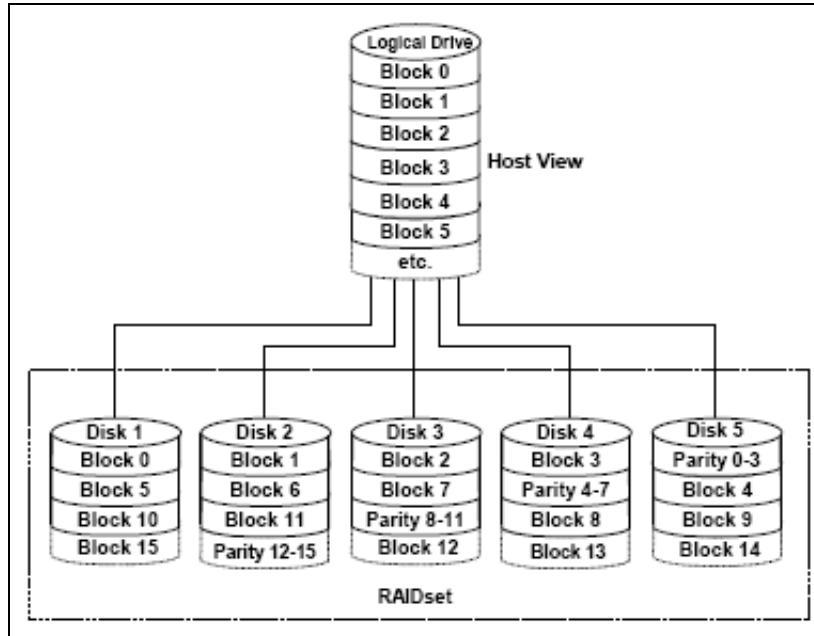


Figure 2-15 RAID 5

RAID Level 5 is best used in environments that require 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 is not complete 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 a 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 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.

With the newest firmware, RAID 5 with nine drives (eight data and one parity) is recommended.

You must configure the RAID 5 array with a certain number of physical drives to take advantage of full stripe writes. This is illustrated for the case of a RAID 5 array with eight total drives (seven data + one parity), as shown in Figure 2-16.

A: Total # of drives written	B: Total # of RAID 5 Read Data IOs	C: Total # of RAID 5 Read Parity IOs	D: Total # of RAID 5 Write Data IOs	E: Total # of RAID 5 Write Parity IOs	X: Total # of IOs for RAID 5 ( = B+C+D+E )	Y: Total # of IOs 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 2-16 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 is always twice the value of A. Columns B, C, D, and E contain the numbers of read data/parity and write data/parity I/Os that are 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) that are required for each write operation.

The decrease in the overhead read operations with the full stripe write operation is the advantage that 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 DS storage subsystem detects that it is receiving contiguous full stripe writes, it switches internally to an even faster write capability that is known as Fast Stripe Write Through. In this method of writing, the DS storage subsystem 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 storage subsystem.

This situation requires that the following rules are being met by the I/O pattern:

- ▶ All write I/Os are full stripe writes (no partial stripe writes can be requested).
- ▶ Write I/Os are sequential and contiguous in nature, so no seeks are required.

If any interruptions in this pattern are detected, the writes revert to the standard full stripe write model, which nevertheless still gives a benefit to the RAID 5 for large sequential writes over the RAID 1.

## RAID 6: Data striping with dual distributed parity

RAID 6 (Figure 2-17) 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 the 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 DCS3700 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 6 can be mitigated by using battery-backed write caching.

The DCS3700 requires a minimum of five drives and supports a maximum of 30 drives in a RAID 6 array.

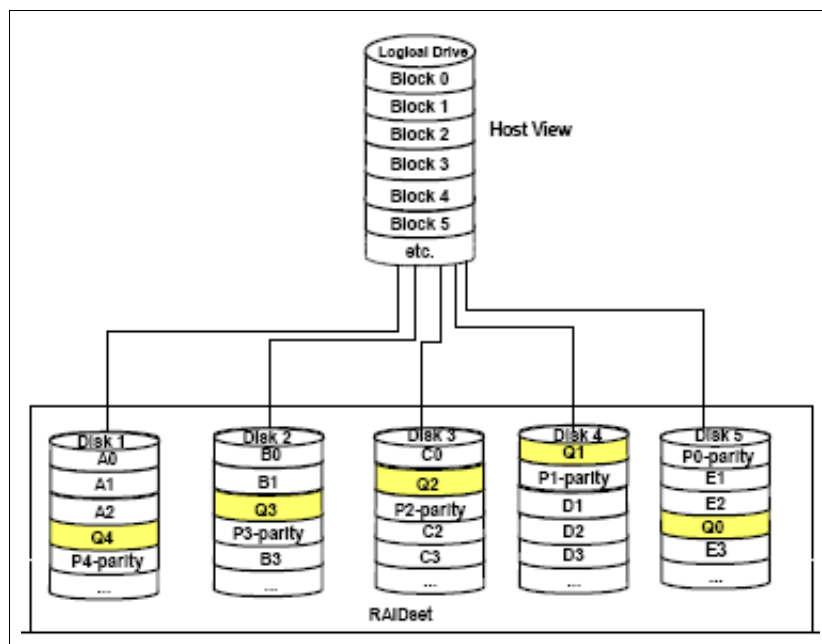


Figure 2-17 RAID 6

Here are the general characteristics of a RAID 6 array:

- ▶ It uses disk striping with dual distributed parity. Two independent parity blocks per stripe ( $P+Q$ ) are calculated in the 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 a battery-backed write cache can mitigate the impact of this write penalty.
- ▶ It has additional costs because of the extra disk capacity that is 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, benefit from RAID 6.

## RAID levels summary

This section summarizes the general characteristics of the various RAID levels that are supported by the DCS3700 storage subsystem. The following note and Table 2-4 summarize this information.

**Summary:** Here are the general performance characteristics of the various RAID levels:

- ▶ 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 2-4 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 that are 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.
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 DS storage subsystem	Slower in writing data, complex RAID controller architecture.

**Note:** Starting with firmware release V7.83.xx.xx, you can also use Dynamic Disk Pools (DDP). For more information about DDP, see 2.4.4, “Dynamic Disk Pooling” on page 59.

## **Throughput-based processes (Mbps)**

Throughput-based workloads are seen with applications or processes that must send massive amounts of data, and frequently use large sequential blocks to reduce disk latency. Throughput rates are heavily dependent on the storage subsystem's internal bandwidth. Generally, fewer numbers of drives are needed to reach maximum throughput rates with the storage subsystem. This number can vary between the various members of the mid-range family and their respective bandwidth handling capabilities. Newer storage subsystems with broader bandwidths are able to reach higher numbers and bring higher rates to bear. In this environment, read operations use the cache to stage greater chunks of data at a time to improve the overall performance.

## **Transaction-based processes (IOPS)**

High performance in transaction-based environments cannot be created with a low-cost model (with a small quantity of physical drives) of a storage subsystem. Transaction-intense applications frequently use a small random data block pattern to transfer data. A read cache is far less effective, and the misses need to be retrieved from disk.

This situation results in these applications being heavily dependent on the number of back-end drives that are available for parallel processing of the host's workload. With too few drives, high queuing with longer response times is encountered.

Determining at what point the performance can be accepted is important in deciding how large a storage subsystem is needed. The willingness to have higher response times can help in reducing the cost, but might not be acceptable by the user community. Cost and response time frequently determines how many drives are enough.

In many cases, slow transaction performance problems can be traced directly to "hot" files that cause a bottleneck on a particular critical component (such as a single physical disk). This situation can occur even when the overall storage subsystem is seeing a fairly light workload. When bottlenecks occur, they can present a difficult and frustrating task to resolve.

Because workload content can be continually changing throughout the course of the day, these bottlenecks can be mysterious in nature and appear and disappear, or move from one location to another location over time.

## **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 lose 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 because of 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 depends on the type of data and how sensitive it is to corruption. Although most storage subsystems (including the DCS3700) 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 that is 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 does 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.

### 2.4.3 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 the 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 disk 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 impact the same physical disks. If you have two logical drives inside an array and they are both 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 contains only 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 2-3 on page 47 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 task 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 360 drives for a full system 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 use all the capacity. However, with this layout, consideration should be given to the workload types for which these LUNs are 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 - 16 disks provides the best performance for RAID 5 workloads that are OLTP-based.

**General rule:** For high-transaction environments that require the highest redundancy and protection, logical drives should be built on arrays with 8 - 12 disks when you use RAID 5 or RAID 10. Spreading the arrays evenly across the two controllers in a dual controller environment provides 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 that are laid out per the RAID type in the previous discussion. By using multiple arrays, you can increase the number of controllers that are involved in handling the load, therefore getting full usage of the storage subsystems resources.

### ***Large throughput workload***

In a 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 - 28 drives are enough to reach the maximum throughput.

This situation does, however, require that the drives be spread evenly across the DCS3700 to best use 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 DCS3700 resources into play is important. Keeping the drive loops and bus busy with high data throughput is the ideal situation. This is also the perfect model for using the high capacity drives because you are looking to push a large volume of data, and this data 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 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 way consists of two 8+1 parity disks that are 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 at the cost of one additional drive.

Transfer time can be gained by splitting the two 8+1 parity into four 4+1 parity arrays giving four streams, but the addition of two 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 by using one of the storage subsystem's segment size selection, which enables the full stripe write capability that is described in 2.4.4, "Dynamic Disk Pooling" on page 59.

Regardless of the workload type or the RAID type being used for the array group, in all cases, building the array by 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 that are used to attain enclosure loss protection.

## 2.4.4 Dynamic Disk Pooling

This disk pool feature is a new way to deliver RAID protection and consistent performance. A disk pool is a set of drives that is logically grouped in the storage subsystem. The drives in each disk pool must be of the same drive type and drive media type, and they must be similar in size. As with an array, you can create one or more logical drives in the disk pool. However, the disk pool is different from the array by the way the data is distributed across the drives that comprise the disk pool.

In an array, the data is distributed across the drives based on a RAID level. You can specify the RAID level when you create the array. The data for each logical drive is written sequentially across the set of drives that comprise the array.

In a disk pool, the storage management software distributes the data for each logical drive randomly across a set of drives that comprise the disk pool. Each disk pool must have a minimum of 11 drives. Although there is no limit to the maximum number of drives that can comprise a disk pool, the disk pool cannot contain more drives than the maximum limit for each storage subsystem. The storage management software automatically configures the RAID level when you create the disk pool. You cannot set or change the RAID level of disk pools or the logical drives in the disk pools.

**Important:** Disk pools require a minimum of 11 drives to automatically distribute logical drives (or volumes) throughout the pool.

### Disk pool benefits

Here are some of the benefits of disk pools:

- ▶ **Easy to create:** It is easy to create a disk pool by using the storage management software. To create a disk pool, you select the drives from a list of eligible drive candidates. After a disk pool is created, you create logical drives. When you create disk pool logical drives, the only attribute you must specify is the logical drive capacity.

- ▶ **Better utilization of drives:** When you add drives to a storage subsystem, the storage management software automatically detects the drives and prompts you to create a single disk pool or multiple disk pools that are based on the drive type and the current configuration. If disk pools were previously defined, the storage management software provides the option of adding the compatible drives to an existing disk pool. When new drives are added to an existing disk pool, the storage management software automatically redistributes the data across the new capacity, which now includes the new drives that you added. The data in the logical drives remain accessible when you add the drives to the disk pool. When you delete disk pool logical drives, the capacity of those logical drives is added to the total usable capacity of the disk pool and, therefore, can be reused.

**Note:** You can manually create a disk pool if you prefer not to proceed with the automatic disk pool creation process.

- ▶ **Reduced hot spots:** A host might access some drives in the array for data more frequently than other drives because of the sequential manner in which the data is written to the drives. This frequency of access to drives creates hot spots in the array. In a disk pool, the hot spots are reduced because of the random manner in which the data is spread across many drives. The reduction of hot spots in the disk pool improves performance of the storage subsystem.
- ▶ **Faster reconstruction of data:** Disk pools do not use hot spare drives for data protection like an array does. Instead of hot spare drives, disk pools use spare capacity within each drive that comprises the disk pool.

In hot spare drive coverage, the maximum drive IOPS limits the speed of reconstruction of data from the failed drive to the hot spare drive. In a disk pool, the reconstruction of data is much faster because the spare capacity in all of the drives that comprise the disk pool is used. Additionally, the data to reconstruct after a drive failure is reduced because the data is spread randomly across more drives in a disk pool.

Faster reconstruction of data in a disk pool also reduces the risk of additional drive failures during a reconstruction operation. For example, consider a drive failure in a RAID 5 array that is composed of three drives. The time that it takes to reconstruct the data from the failed drive is relatively longer for an array. During the reconstruction of data, if another drive fails in this array, data loss occurs. Unlike arrays, the time period for which the disk pool is exposed to multiple drive failures during a reconstruction operation is reduced.

- ▶ **Reduced maintenance:** You can configure the storage management software to send alert notifications when the configured capacity of a disk pool is reaching a specified percentage of free capacity. Additionally, you do not need to manage any hot spare drives. You can replace a set of drives during a scheduled maintenance of the storage subsystem.
- ▶ **Consistent performance:** Dynamic Disk Pools delivers and maintains exceptional performance under all conditions, whether optimal or under the stress of a drive failure. DDP minimizes the performance impact of a drive failure in multiple dimensions. By distributing parity information and spare capacity throughout the disk pool, DDP is able to use every drive in the pool for the intensive process of rebuilding a failed drive. This dynamic rebuild process can return the system to optimal condition faster than traditional RAID.

## Disk pool considerations

Here are some things to consider when you use disk pools:

- ▶ Dynamic Segment Sizing (DSS) is not supported for disk pools.

- ▶ You cannot change the RAID level of a disk pool. The storage management software automatically configures disk pools as RAID 6.
- ▶ You cannot export a disk pool from a storage subsystem or import the disk pool to a different storage subsystem.
- ▶ All drive types (Fibre Channel, SATA, and SAS) in a disk pool must be the same type.
- ▶ All drive media types in a disk pool must be the same type. Solid-state drives (SSDs) are not supported.
- ▶ You can use the T10 Protection Information (T10 PI) capabilities of a drive set in a disk pool if all drives match in their T10 PI capabilities. However, you can use a drive set with mixed attributes, but the T10 PI capabilities of the drive cannot be used.
- ▶ If you downgrade the controller firmware version of a storage subsystem that is configured with a disk pool to a firmware version that does not support disk pools, the logical drives are lost and the drives are treated as unaffiliated with a disk pool.

### 2.4.5 Thin provisioning

When you create logical drives from a disk pool, you can create them as either standard logical drives or thin logical drives. For a standard logical drive, all of the capacity is allocated when the logical drive is created. For a thin logical drive, the capacity is allocated as the data is being written. Thin logical drives have two types of capacity: virtual capacity and physical capacity. Virtual capacity is the capacity that is reported to the hosts, and physical capacity is the amount of drive space that is allocated for writing data. Thin provisioning allows logical drives to be created with a large virtual capacity and relatively small physical capacity, which is beneficial for storage usage and efficiency. Thin logical drives can help simplify storage administration because the physical capacity can increase as the application’s needs change, without disrupting the application, allowing for better storage usage.

**Note:** You can create a thin logical drive only from a disk pool, not from an array.

**Note:** Because the storage management software does not allocate the full capacity when it creates a thin logical drive, insufficient free capacity might exist in the disk pool. Insufficient space can block writes to the disk pool not only for the thin logical drives, but also for other operations that require capacity from the disk pool, for example, FlashCopy logical drives, Enhanced FlashCopy images, or Enhanced FlashCopy logical drives. If this situation occurs, you receive a free-capacity alert threshold warning.

Table 2-5 lists the properties of a thin logical drive.

*Table 2-5 Properties of a thin logical drive*

Properties of a thin logical drive	Description
Virtual capacity	The capacity that is reported to the hosts.
Physical capacity	The amount of drive space that is allocated for writing data.
Controller ownership	Defines the controller that is designated to be the owning, or primary, controller of the logical drive. Controller ownership is important and should be planned carefully. Make sure that the controllers are balanced as closely as possible for total I/Os.

Properties of a thin logical drive	Description
Mapping	How host LUNs are mapped to a logical drive.
Name	Descriptive name indicating the type of data that is stored in the logical drive.

Table 2-6 lists the capacity limits for a thin logical drive.

*Table 2-6 Capacity limits for a thin logical drive*

Type of capacity	Size
Minimum virtual capacity	32 MB
Maximum virtual capacity	64 TB
Minimum physical capacity	4 GB
Maximum physical capacity	64 TB

Thin logical drives support all of the operations that standard logical drives do with the following exceptions:

- ▶ You cannot change the segment size of a thin logical drive.
- ▶ You cannot enable the pre-read redundancy check for a thin logical drive.
- ▶ You cannot use a thin logical drive as the target logical drive in a VolumeCopy.
- ▶ You cannot use a thin logical drive in a FlashCopy operation.
- ▶ You cannot use a thin logical drive in an Enhanced Remote Mirroring operation.

If you want to change a thin logical drive to a standard logical drive, use the VolumeCopy operation to create a copy of the thin logical drive. The target of a VolumeCopy is always a standard logical drive.

## 2.4.6 Hot spare drives

A hot spare drive is like a replacement drive that is installed in advance. Hot spare disk drives provide additional protection that might prove to be essential if there is 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 2-18).

**General rule:** When assigning disks as hot spares, make sure that they have enough storage capacity. If the failed disk drive is larger, and the used capacity is more than the size of the hot spare, reconstruction is not possible. Use the largest drive size for the drive type that is installed as hot spare drives. For more information about defining hot spare drives, see 9.3, “Setting a hot spare drive” on page 314.

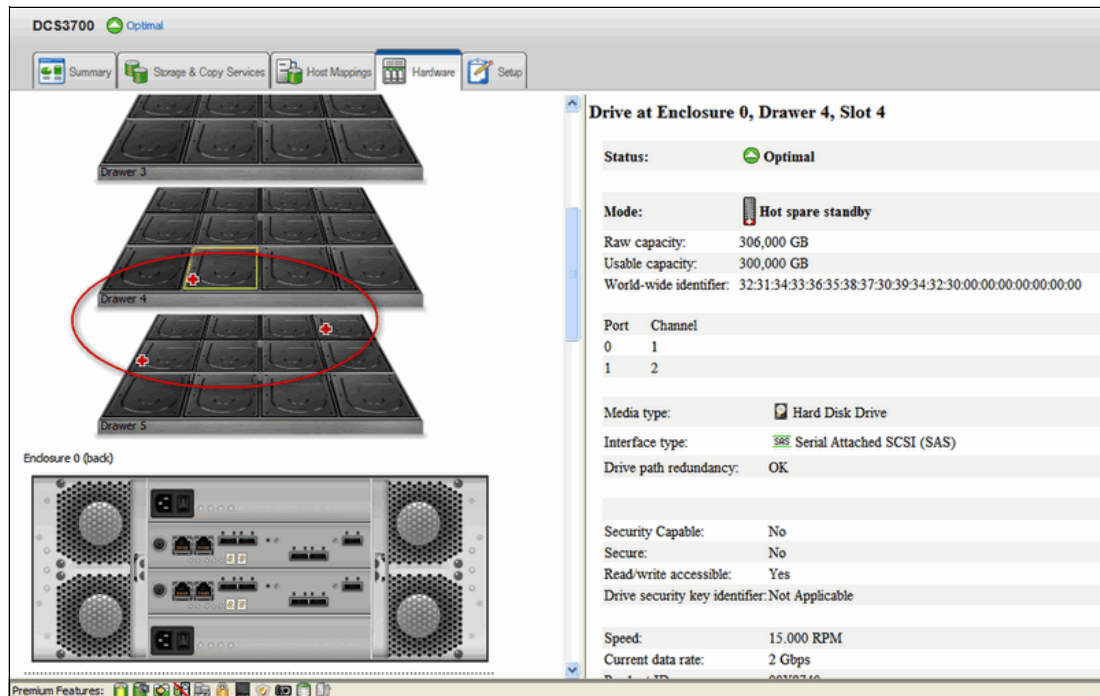


Figure 2-18 Hot spare coverage with alternating loops

## Planning for hot spares

Even though disk drives are becoming increasingly more reliable, disk failures can occur. To protect your data if there is a drive failure, you should primarily use a redundant RAID level, for example, RAID 1, 5, 6, or 10. This way, your data remains intact should a hard disk drive fail. However, the array then operates in a degraded state until the failed disk drive is replaced and data reconstruction completes on the new disk drive.

To ensure as high an 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 automatically is used as a replacement drive. Data reconstruction on to the hot spare takes place immediately. When reconstruction finishes, the array is again in a fully operational status and the hot spare becomes a member of the array.

Depending on your configuration, it might be wise to use more than just one hot spare. In small DCS3700 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 the best configuration, define one hot spare for every 20 disk drives in the system.

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

- ▶ Hot spares must be of the same drive type (SSD, SAS, or SAS nearline) as the failed drive.
- ▶ Hot spares must be of the same or larger size as the failed drive.
- ▶ A minimum of three hot spares (one SSD, one SAS, and one SAS nearline).

For example, if you use a disk drive with different sizes in your array, 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 that are 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.

## 2.4.7 Data loss protection planning

The data loss protection in a DCS3700 is further classified into drawer loss protection and enclosure loss protection.

### Drawer loss protection

Drawer loss protection is a characteristic of an array. In drive expansion units that contain drives in drawers, a drawer failure can lead to inaccessibility of data on the logical drives in an array. A drawer might fail because of a loss of power, a failure of an environmental services module (ESM), or a failure of an internal component within the drawer.

The availability of drawer loss protection for an array is based on the location of the drives that make up the array. If there is a single drawer failure, data on the logical drives in an array remains accessible if the array has drawer loss protection. If a drawer fails and the array is drawer loss protected, the array changes to Degraded status, and the data remains accessible.

To achieve drawer loss protection, make sure that the drives that make up an array is in different drawers with respect to their RAID levels, as shown in Table 2-7.

Table 2-7 Criteria for Drawer Loss Protection

RAID level	Criteria for drawer loss protection
RAID 6	RAID 6 requires a minimum of five drives. Place all of the drives in different drawers or place a maximum of two drives in the same drawer and the remaining drives in different drawers to achieve drawer loss protection for a RAID 6 array.
RAID Level 3 and RAID Level 5	RAID 3 and RAID 5 require a minimum of three drives. Place all of the drives in different drawers for a RAID 3 array and for a RAID 5 array to achieve drawer loss protection. Drawer loss protection cannot be achieved for RAID 3 and RAID 5 if more than one drive is placed in the same drawer.

RAID level	Criteria for drawer loss protection
RAID 1 and RAID 10	<p>RAID 1 requires a minimum of two drives. Make sure that each drive in a mirrored pair is in a different drawer.</p> <p>If you make sure that each drive in a mirrored pair is in a different drawer, you can have more than two drives of the array within the same drawer. For example, if you create a RAID 1 array with six drives (three mirrored pairs), you can achieve the drawer loss protection for the array with only two drawers, as shown in this example:</p> <p>Six-drive RAID 1 array:</p> <ul style="list-style-type: none"> <li>▶ Mirror pair 1 = Drive in enclosure 1, drawer 1, slot 1, and drive in enclosure 1, drawer 2, slot 1</li> <li>▶ Mirror pair 2 = Drive in enclosure 1, drawer 1, slot 2, and drive in enclosure 1, drawer 2, slot 2</li> <li>▶ Mirror pair 3 = Drive in enclosure 1, drawer 1, slot 3, and drive in enclosure 2, drawer 2, slot 3</li> </ul> <p>RAID 10 requires a minimum of four drives. Make sure that each drive in a mirrored pair is in a different drawer.</p>
RAID 0	You cannot achieve drawer loss protection because the RAID 0 array does not have redundancy.

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

**Note:** If you create an array by using the Automatic drive selection method, the storage management software attempts to choose drives that provide drawer loss protection. If you create an array by using the Manual drive selection method, you must use the criteria that are specified in the previous table.

### Enclosure loss protection

Enclosure loss protection is an attribute of an array. Enclosure loss protection ensures accessibility to the data on the logical drives in an array if a total loss of communication occurs with a single drive expansion unit. An example of total loss of communication might be a loss of power to the drive expansion unit or failure of both environmental services modules (ESMs).

**Note:** Enclosure loss protection is not ensured if a drive has already failed in the array. In this situation, losing access to a drive expansion unit and another drive in the array causes a double drive failure and loss of data.

Enclosure loss protection is achieved when you create an array where all of the drives that make up the array are in different drive expansion units. This distinction depends on the RAID level, as shown in Table 2-8. If you choose to create an array by using the Automatic method, the software attempts to choose drives that provide enclosure loss protection. If you choose to create an array by using the Manual method, you must use the criteria that are specified in Table 2-8.

Table 2-8 Enclosure loss protection

RAID level	Criteria for enclosure loss protection
RAID 3, RAID 5, or RAID 6	<p>Make sure that all of the drives in the array are in different drive expansion units.</p> <p>Because a RAID 3 array or a RAID 5 array requires a minimum of three drives, enclosure loss protection cannot be achieved if your storage subsystem has less than three drive expansion units. Because a RAID 6 array requires a minimum of five drives, enclosure loss protection cannot be achieved if your storage subsystem has less than five drive expansion units.</p>
RAID 1	<p>Make sure that each drive in a mirrored pair is in a different drive expansion unit.</p> <p>If you make sure that each drive in a mirrored pair is in a different drive expansion unit, you can have more than two drives in the array within the same drive expansion unit. For example, if you were creating a six-drive array (three mirrored pairs), you still could achieve enclosure loss protection with only two drive expansion units by specifying that the drives in each mirrored pair are in different drive expansion units. This example shows this concept.</p> <p>Six-drive RAID 1 array:</p> <ul style="list-style-type: none"> <li>▶ Mirror pair 1 = Drive in enclosure 1, slot 1, and drive in enclosure 2, slot 1</li> <li>▶ Mirror pair 2 = Drive in enclosure 1, slot 2, and drive in enclosure 2, slot 2</li> <li>▶ Mirror pair 3 = Drive in enclosure 1, slot 3, and drive in enclosure 2, slot 3</li> </ul> <p>Because a RAID 1 array requires a minimum of two drives, you cannot achieve enclosure loss protection if your storage subsystem has less than two drive expansion units.</p>
RAID Level 0	<p>Because a RAID 0 array does not have redundancy, you cannot achieve enclosure loss protection.</p>

## 2.4.8 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 that are called an *array group*, which 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.

### Balancing traffic

Balancing traffic is 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 process can sometimes be helped by using host-based volume management software products.



In addition, typically, the load across controllers and logical drives is constantly changing. The logical drives and data that is accessed depends on which applications and users are active during that time period, which is why it is important to monitor the system.

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

- ▶ Assign logical drives across both controllers to balance controller usage.
- ▶ Use the manual method of creating logical drives, which allows greater flexibility for configuration settings, such as enclosure loss protection and using 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 are created.

### Enhanced remote mirror (ERM) considerations

A secondary logical drive in a remote mirror always has 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 causes a corresponding controller ownership change of the secondary logical drive as well. For more information about ERM, see the *IBM System Storage DS Storage Manager Copy Services Guide*, SG24-7822.

## 2.4.9 Storage partitioning

Storage partitioning adds a high level of flexibility to the DCS3700 storage subsystem.

Depending on the firmware, DCS3700 comes with 4 - 128 partitions as the default, and this can be expanded to a maximum of 512. For more information about Premium Feature needs, see 2.5, “Planning for Premium Features” on page 74. Storage partitioning 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 misleading because it represents a host or a group of hosts and the logical drives they access.

Without storage partitioning, the logical drives that are configured on a DCS3700 storage subsystem can be accessed only by a single host system or by a single cluster, which can lead to the inefficient usage of storage server hardware unless the usage of the DCS3700 storage subsystem is dedicated to a single host (for example, a SAN Volume Controller attachment, where it is seen as a single host).

Storage partitioning, however, allows the creation of “sets” that contain the hosts with their host bus adapters and the logical drives. These sets are called *storage partitions*. The host systems can access only their assigned logical drives, 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 the other hosts that are 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 before connecting multiple hosts. Operating systems such as Windows 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 that are assigned to the partition. Therefore, file systems or disk structures 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 can mount other file system types and share logical drives. In these cases, care should be taken to ensure that the partition members are correctly defined and configured to share.

Storage partition topology is a collection of topological elements (default group, host groups, hosts, and host ports) that is 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 Fibre Channel port of the host bus adapter (HBA) on the host system.

**Note:** The iSCSI and SAS HBAs are not covered in this section.

The host port is identified by its worldwide name (WWN). A single host can contain more than one host port. If the servers are attached using full redundancy, each server has 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 DCS3700 controllers, which can be achieved by SAN zoning.

The DCS3700 storage subsystem communicates only by using a WWN. The storage subsystem 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 that are 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 2048 storage partitions on a DCS3700, which allows the storage subsystem to provide storage capacity to a number of heterogeneous hosts, allowing for greater flexibility and scalability. On DCS3700 models, the number of partitions also depends on the Premium Feature license you have purchased.

For the maximum number of logical drives per partitions for a specific host type, see Table 2-9.

Table 2-9 Host logical drive restrictions

Operating system	Maximum number of logical drives that are supported per partition
Windows Server 2003	256
Windows Server 2008	256
AIX	256
Linux	256

Every mapping of a logical drive to a new host or host group creates a 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 subsystem and other logical drives are mapped to that host, another storage partition must be created for it. If a 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. DCS3700 storage servers support up to 256 logical drives (including the “access” logical drive) per partition (although there are also limitations, depending on the host operating system) and a maximum of two partitions per host. Keep in mind these limitations when planning the DCS3700 storage subsystem installation.

## Storage partitioning considerations

By default, an “access” logical drive is mapped to each partition that is 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 other environments, this logical drive can be removed.

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

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

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:

- ▶ By default, the DCS3700 comes, depending of the firmware, with 8 or 128 partitions.
- ▶ In a cluster environment, you need to use host groups.
- ▶ You can optionally purchase additional partitions (maximum of 512).

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 that are 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 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 file system servers (for example, GPFS).

Table 2-10 shows an example of a storage partitioning plan that shows the host groups, hosts, port names, WWN of the ports, and the operating systems that are used in that environment. Other columns can be added to the table for future reference, such as HBA BIOS levels, driver revisions, and switch ports that are used, all of which can then form the basis of a change control log.

*Table 2-10 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	
Power Systems	AIX_Host	AIXAdp_A	20000000C926B6D2	AIX
		AIXAdp_B	20000000C926B08	

### Heterogeneous hosts

When you implement a DCS3700 storage subsystem 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 that access logical drives in a single storage partition should 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.

### Deleting the access logical drive (31)

The DCS3700 storage subsystem automatically creates a logical drive for each host that is attached (logical drive ID 31). This drive (also known as access LUN) is used for in-band management, so if you do not plan to manage the DCS3700 storage subsystem from that host, you should delete this logical drive. This gives you one more logical drive to use for host user data.

**Important:** In all cases where you are not using in-band management, you should delete the access logical driver for the host partition. After it is deleted, the access LUN cannot be re-created.

## 2.4.10 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 different for different workload types. It is important to know the host I/O blocksize that will be used when you decide on the segment size.

For small host I/Os, as mentioned earlier, set the segment size larger than the host I/O size. Doing so 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 mid-range storage subsystems. In many cases, using the segment size that is equal to the host IO size can result in multiple disk operations because of alignment differences. This can be easily remedied with a larger segment size selection for small I/O 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 are described 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. Aligning data on segment boundaries is important for performance. With larger segment sizes, there are fewer boundaries 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 a poor layout of the host data on the disks because of boundary differences.

**General rule:** For most high transaction workloads with the mid-range storage subsystems, the segment size of 128 KB (the default) works best.

With a 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 wanted to accomplish the higher throughput you are looking for. For high read throughput, you want to have large segments (128 KB 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 that is known as *full stripe (stride) write*, which can work well to improve your performance. With RAID 5, the parity is based on the value that calculated for a stripe. As described in 2.4.2, “Understanding RAID types” on page 48, when the I/O that is 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 example, 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 been 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, 64 - 128 KB should provide the best results.

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

## 2.4.11 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 that are 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 has to complete the task. During the scan process, the DCS3700 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 that is allocated. After the media scan completes, 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 media scan over host activity to ensure that the scan completes in the allocated time. This scan can impact performance, but improve data integrity.

Media scan must be enabled for the entire storage subsystem. This system-wide enablement specifies the duration over which the media scan runs. The logical drive enablement specifies whether to do a redundancy check and media scan.

A media scan can be considered a surface scan of the hard disk drives, and 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 used in excess of 95%. The length of time that it takes to scan the LUNs depends on the capacity of all the LUNs on the system and the usage of the controller.

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

## 2.4.12 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 DCS3700. For a detailed explanation about the cache parameter settings, see 8.8.6, “Change Cache Settings” on page 264.

### Cache blocksize selection

On the DCS3700, the cache blocksize is a variable value that can be set to 4K, 8K, 16K, or 32 K. The general default setting is 8 K. The main goals with setting this value are to minimize the number of cache I/Os that are needed and 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 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 DCS3700 for transaction-intensive environments.

In a throughput-intensive environment, you want to get as much data in to 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 DCS3700 for throughput-intensive 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 DCS3700 for mixed workload environments.

**Tip:** Throughput-based operations, although impacted by smaller cache blocksize, can still perform reasonably if all other effects are 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 DCS3700 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 is reached before the server starts to flush the data to disk, and at what level the flushing stops.

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 flushes 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 that the write data can remain in cache before it is forced out and written to disks. This value by default is set to 10 seconds but can be changed by running 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.

## 2.5 Planning for Premium Features

The Premium Features that come with the DCS3700 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 that are available for different types of expanded capabilities. Here is a list of the available Premium Features for the DCS3700:

- ▶ VolumeCopy
- ▶ FlashCopy
- ▶ Enhanced FlashCopy (up to 256)
- ▶ Enhanced FlashCopy Upgrade (up to 2048)
- ▶ Enhanced FlashCopy (512) / VolumeCopy (256) Bundle
- ▶ Enhanced Remote Mirroring
- ▶ Storage Partitioning (up to 512 partitions)
- ▶ Super Key

Super Key for DCS3700 enables all the functionality in the Disaster Recovery Option, the Backup and Restore Option, and the Performance Read Cache.

Here are points that might play an important part with Premium Features and should be considered:

- ▶ 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 DCS3700. 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.

### 2.5.1 VolumeCopy

The VolumeCopy Premium Feature enables you to create a point-in-time copy of a logical drive by creating two separate logical drives, the source logical drive and the target logical drive, on the same storage subsystem. VolumeCopy performs a byte-by-byte copy from the source logical drive to the target logical drive; therefore, the data on the target logical drive is identical to the data on the source logical drive.



Use the VolumeCopy Premium Feature to perform these tasks:

- ▶ Copy data from one logical drive (the source logical drive) to another logical drive (the target logical drive) in a single storage subsystem.
- ▶ Copy data from logical drive groups that use smaller capacity drives to logical drive groups that use larger capacity drives.
- ▶ Create an online copy of data from a logical drive within a storage subsystem while still being able to write to the logical drive with the copy in progress.
- ▶ Back up data or restore FlashCopy logical drive data to the base logical drive.

We now highlight the VolumeCopy Features and benefits.

### **Data copying for greater access**

As your storage requirements for a logical drive change, use the VolumeCopy Premium Feature to copy data to a logical drive in an array that uses larger capacity drives within the same storage subsystem. This premium feature lets you perform these functions:

- ▶ Move data to larger drives, for example, 73 GB to 146 GB.
- ▶ Change to drives with a higher data transfer rate, for example, 2 Gbps to 4 Gbps.
- ▶ Change to drives using new technologies for higher performance.

### **Data backup**

The VolumeCopy Premium Feature lets you back up a logical drive by copying data from one logical drive to another logical drive in the same storage subsystem. You can use the target logical drive as a backup for the source logical drive, for system testing, or to back up to another device, such as a tape drive.

### **Enhanced FlashCopy logical drive data restoration to the base logical drive**

If you need to restore data to the base logical drive from its associated Enhanced FlashCopy logical drive, use the VolumeCopy Premium Feature to copy data from the Enhanced FlashCopy logical drive to the base logical drive. You can create a VolumeCopy of the data on the Enhanced FlashCopy logical drive, and then copy the data to the base logical drive.

### **Types of VolumeCopy copies**

You can perform either an offline VolumeCopy or an online VolumeCopy. To ensure data integrity, all I/O to the target logical drive is suspended during either VolumeCopy operation. This suspension occurs because the state of data on the target logical drive is inconsistent until the procedure is complete. After the VolumeCopy operation is complete, the target logical drive automatically becomes read-only to the hosts.

The offline and online VolumeCopy operations are described below.

#### ***Offline copy***

An offline copy reads data from the source logical drive and copies it to a target logical drive while suspending all updates to the source logical drive with the copy in progress. All updates to the source logical drive are suspended to prevent chronological inconsistencies from being created on the target logical drive. The offline VolumeCopy relationship is between a source logical drive and a target logical drive.

Source logical drives that are participating in an offline copy are available for read I/O activity only while a VolumeCopy has a status of In Progress or Pending. Write requests are allowed after the offline copy has completed. If the source logical drive has been formatted with a journaling file system, any attempt to issue a read request to the source logical drive might be rejected by the storage subsystem controllers, and an error message might appear. The journaling file system driver issues a write request before it attempts to issue the read request. The controller rejects the write request, and the read request might not be issued due to the rejected write request. This condition might result in an error message appearing, which indicates that the source logical drive is write protected. To prevent this issue from occurring, do not attempt to access a source logical drive that is participating in an offline copy while the VolumeCopy has a status of In Progress. Also, make sure that the Read-Only attribute for the target logical drive is disabled after the VolumeCopy has completed to prevent error messages from appearing.

### ***Online copy***

An online copy creates a point-in-time copy of any logical drive within a storage subsystem while still being able to write to the logical drive with the copy in progress. This function is achieved by creating an Enhanced FlashCopy of the logical drive and using the Enhanced FlashCopy as the actual source logical drive for the copy. The online VolumeCopy relationship is between an Enhanced FlashCopy logical drive and a target logical drive. The logical drive for which the point-in-time image is created is known as the base logical drive and must be a standard logical drive or a thin logical drive in the storage subsystem.

An Enhanced FlashCopy logical drive and a repository logical drive are created during the online copy operation. The Enhanced FlashCopy logical drive is not an actual logical drive containing data; rather, it is a reference to the data that was contained on a logical drive at a specific time. For each Enhanced FlashCopy that is taken, a repository logical drive is created to hold the data for the Enhanced FlashCopy. The repository logical drive is used only to manage the Enhanced FlashCopy image.

Before a data block on the source logical drive is modified, the contents of the block to be modified are copied to the repository logical drive for safekeeping. Because the repository logical drive stores copies of the original data in those data blocks, further changes to those data blocks write only to the source logical drive. The online copy operation uses less disk space than a full physical copy because only data blocks that are stored in the repository logical drive are those that have changed since the time of the Enhanced FlashCopy.

## **2.5.2 FlashCopy**

A FlashCopy logical drive is a logical point-in-time image of a logical drive, which called a base logical drive. A FlashCopy logical drive has the following features:

- ▶ It is created quickly and requires less disk space than an actual logical drive.
- ▶ It can be assigned a host address so that you can perform backups with the FlashCopy logical drive while the base logical drive is online and accessible.
- ▶ You can use the FlashCopy logical drive to perform application testing or both scenario development and analysis. This does not affect the actual production environment.
- ▶ The maximum number of allowed FlashCopy logical drives is one-half of the total logical drives that are supported by your controller model.

## 2.5.3 Enhanced FlashCopy

An Enhanced FlashCopy image is a logical image of the content of an associated base logical drive that is created at a specific moment. An Enhanced FlashCopy image can be thought of as a restore point. It is not directly read or write accessible to hosts, as the Enhanced FlashCopy image is used to save only the data that is captured from the base logical drive.

You must create an Enhanced FlashCopy logical drive to allow the host access to a copy of the data that is contained in the Enhanced FlashCopy image. The Enhanced FlashCopy logical drive contains its own repository, which is used to save any subsequent modifications that are made by the host application to the base logical drive without affecting the referenced Enhanced FlashCopy image.

### Enhanced FlashCopy images usage

Enhanced FlashCopy images are useful when you must roll back to a known good data set at a specific point in time. For example, before you perform a risky operation on a logical drive, you can create an Enhanced FlashCopy image to enable the “undo” capability for the entire logical drive. An Enhanced FlashCopy image is created almost instantaneously, and initially uses no disk space, because it stores only the incremental changes that are needed to roll the logical drive back to the point-in-time when the Enhanced FlashCopy image was created.

You can roll back data by performing either of the following actions:

- ▶ Creating an Enhanced FlashCopy logical drive of an Enhanced FlashCopy image, which allows you to retrieve deleted files from that Enhanced FlashCopy logical drive (the base logical drive remains undisturbed).
- ▶ Restoring an Enhanced FlashCopy image to the base logical drive, which allows you to roll back the base logical drive to a previous state.

### Enhanced FlashCopy image operations

You can create Enhanced FlashCopy images manually or automate the process with a schedule. An Enhanced FlashCopy image is generated automatically, based on the date and time when the Enhanced FlashCopy image was created, and is displayed with its associated Enhanced FlashCopy group.

You can create Enhanced FlashCopy images for these storage objects:

- ▶ Standard logical drives
- ▶ Thin logical drives
- ▶ Consistency groups

To create an Enhanced FlashCopy image, you must first create an Enhanced FlashCopy group and reserve Enhanced FlashCopy repository space for the logical drive. The repository space is based on a percentage of the current logical drive reserve.

You can delete the oldest Enhanced FlashCopy image in an Enhanced FlashCopy group either manually or you can automate the process by enabling the Auto-Delete setting for the Enhanced FlashCopy group. When an Enhanced FlashCopy image is deleted, its definition is removed from the system, and the space that is occupied by the Enhanced FlashCopy image in the repository is released and made available for reuse within the Enhanced FlashCopy group.

## Characteristics of Enhanced FlashCopy images

Here are the characteristics of Enhanced FlashCopy images:

- ▶ Enhanced FlashCopy images are always created inside Enhanced FlashCopy groups.
- ▶ Each Enhanced FlashCopy image is associated with exactly one Enhanced FlashCopy group.
- ▶ There is a maximum limit of Enhanced FlashCopy groups for a single associated base logical drive (depending on your configuration).
- ▶ There is a maximum limit of Enhanced FlashCopy images per Enhanced FlashCopy group (depending on your configuration).

## Differences between FlashCopy copies and Enhanced FlashCopy images

An Enhanced FlashCopy image is a logical point-in-time image of a logical drive. The Enhanced FlashCopy image feature is similar to the FlashCopy feature, with the following differences:

- ▶ When there are multiple point-in-time images for a base logical drive, the new Enhanced FlashCopy image feature offers improved performance. The FlashCopy feature uses one data repository for each FlashCopy logical drive. The new Enhanced FlashCopy image feature uses one data repository for all of the Enhanced FlashCopy images that are associated with a base logical drive. Therefore, when a base logical drive is written to, the new Enhanced FlashCopy image feature requires only one write operation instead of multiple sequential write operations.
- ▶ The new Enhanced FlashCopy image feature adds the concept of an Enhanced FlashCopy group. Because there is only one repository for multiple Enhanced FlashCopy images, the repository is associated with the Enhanced FlashCopy group instead of with the Enhanced FlashCopy image as it is with the FlashCopy feature.
- ▶ Unlike a FlashCopy logical drive, a new Enhanced FlashCopy image is not directly read/write accessible by hosts because the Enhanced FlashCopy image is used only to save the changed data for a base logical drive. To provide hosts with read/write access to an Enhanced FlashCopy image, you must first create an Enhanced FlashCopy logical drive.
- ▶ You can create either FlashCopy copies or Enhanced FlashCopy images from a base logical drive, but not both.

## Enhanced FlashCopy consistency groups

If you frequently want to perform the same Enhanced FlashCopy image operations on multiple logical drives, you can create a consistency group. A consistency group is a group of logical drives that you treat as one entity. An operation that you perform on the consistency group is performed simultaneously on all of the logical drives in the consistency group. Some of the Enhanced FlashCopy image operations that you can perform on a consistency group include creating, scheduling, and rolling back.

## 2.5.4 Enhanced Remote Mirroring

The Enhanced Remote Mirroring Premium Feature is used for online, real-time data replication between storage subsystems over a remote distance. Storage system controllers manage the mirror operation, which is transparent to host machines and software applications. You create one or more mirrored pairs that consist of a primary logical drive at the primary site and a secondary logical drive at a secondary, remote site. After you create the mirror relationship between the two logical drives, the current owner of the primary logical drive copies all of the data from the primary logical drive to the secondary logical drive. This process is called a full synchronization.

**Caution:** There is a possible loss of data access. You cannot create a mirror relationship if the primary logical drive contains unreadable sectors. Furthermore, if an unreadable sector is discovered during a mirror operation, the mirror relationship fails.

**Note:** Because replication is managed on a per-logical drive basis, you can mirror individual logical drives in a primary storage subsystem to the appropriate secondary logical drives in several different remote storage subsystems.

### Disaster recovery

The secondary, remote logical drive is unavailable to secondary host applications while the mirror operation is in progress. If there is a disaster at the primary site, you can fail over to the secondary site. To fail over, perform a role reversal to promote the secondary logical drive to a primary logical drive. Then, the recovery host is able to access the newly promoted logical drive, and business operations can continue.

### Data replication

When the current owner of the primary logical drive receives a write request from a host, the controller first logs information about the write to a special logical drive. This logical drive is called a mirror repository logical drive. It writes the data to the primary logical drive. Next, the controller initiates a remote write operation to copy the affected data blocks to the secondary logical drive at the remote site.

Finally, the controller sends an I/O completion indication back to the host system to confirm that the data was copied successfully to the secondary storage subsystem. The write mode that you selected when you first created a remote mirror determines when the I/O completion indication is sent to the host system.

The storage management software provides two write modes:

- ▶ **Metro Mirror** 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 ran on both (local and remote) storage controllers.
- ▶ **Global copy mode** copies a non-synchronous, remote copy function that completes write operations on the primary storage subsystem before they are received by the secondary storage subsystem. This capability prevents 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.

When write caching is enabled on either the primary logical drive or the secondary logical drive, the I/O completion is sent when data is in the cache on the side (primary or secondary) where write caching is enabled. When write caching is disabled on either the primary logical drive or the secondary logical drive, the I/O completion is not sent until the data has been stored on the physical media on that side.

Host write requests that are received by the controller are handled normally. No communication takes place between the primary storage subsystem and the secondary storage subsystem.

### **Link interruptions or secondary logical drive errors**

When processing write requests, the primary controller might be able to write to the primary logical drive, but a link interruption prevents communication with the remote secondary controller.

In this case, the remote write cannot complete to the secondary logical drive. The primary logical drive and the secondary logical drive are no longer appropriately mirrored. The primary controller changes the mirrored pair in to Unsynchronized status and sends an I/O completion to the primary host. The primary host can continue to write to the primary logical drive, but remote writes do not take place.

When connectivity is restored between the current owner of the primary logical drive and the current owner of the secondary logical drive, a full synchronization takes place. Only the blocks of data that have changed on the primary logical drive during the link interruption are copied to the secondary logical drive. The mirrored pair changes from an Unsynchronized state to Mirror Synchronization in Progress status.

The primary controller also marks the mirrored pair as Unsynchronized when a logical drive error on the secondary side prevents the remote write from completing. For example, an offline secondary logical drive or a failed secondary logical drive can cause the remote mirror to become unsynchronized. When the logical drive error is corrected (the secondary logical drive is placed online or is recovered to Optimal status), a full synchronization automatically begins. The mirrored pair then changes to Synchronization in Progress status.

### **Connectivity and logical drive ownership**

A primary controller attempts to communicate only with its matching controller in the secondary storage subsystem. For example, controller A in the primary storage subsystem attempts communication only with controller A in the secondary storage subsystem. The controller (A or B) that owns the primary logical drive determines the current owner of the secondary logical drive. If the primary logical drive is owned by controller A on the primary side, the secondary logical drive is owned by controller A on the secondary side. If primary controller A cannot communicate with secondary controller A, controller ownership changes do not take place.

The next remote write processed automatically triggers a matching ownership change on the secondary side if one of these conditions exists:

- ▶ When an I/O path error causes a logical drive ownership change on the primary side
- ▶ If the storage administrator changes the current owner of the primary logical drive

For example, a primary logical drive is owned by controller A, and then you change the controller owner to controller B. In this case, the next remote write changes the controller owner of the secondary logical drive from controller A to controller B. Because controller ownership changes on the secondary side are controlled by the primary side, they do not require any special intervention by the storage administrator.

## Controller Resets and storage subsystem power cycles

Sometimes a remote write is interrupted by a controller reset or a storage subsystem power cycle before it can be written to the secondary logical drive. The storage subsystem controller does not need to perform a full synchronization of the mirrored pair in this case. A controller reset causes a controller ownership change on the primary side from the preferred controller owner to the alternative controller in the storage subsystem. When a remote write is interrupted during a controller reset, the new controller owner on the primary side reads information that is stored in a log file in the mirror repository logical drive of the preferred controller owner. It then copies the affected data blocks from the primary logical drive to the secondary logical drive, eliminating the need for a full synchronization of the mirrored logical drives.

## Enhanced Remote Mirroring Premium Feature activation

Like other Premium Features, you enable the Enhanced Remote Mirroring Premium Feature by purchasing a feature key file from your storage supplier. You must enable the Premium Feature on both the primary storage subsystem and the secondary storage subsystem.

Unlike other Premium Features, you also must activate the Premium Feature after you enable it. To activate the Premium Feature, use the Activate Enhanced Remote Mirroring wizard in the Subsystem Management Window (SMW). Each controller in the storage subsystem must have its own mirror repository logical drive for logging write information to recover from controller resets and other temporary interruptions. The Activate Enhanced Remote Mirroring wizard guides you to specify the placement of the two mirror repository logical drives (on newly created free capacity or existing free capacity in the storage subsystem).

After you activate the PREMIUM FEATURE, one Fibre Channel (FC) host side I/O port on each controller is solely dedicated to Enhanced Remote Mirroring operations. Host-initiated I/O operations are not accepted by the dedicated port. I/O requests received on this port are accepted only from remote controllers that are participating in Enhanced Remote Mirroring operations with the controller.

## Connectivity requirements

You must attach dedicated Enhanced Remote Mirroring ports to a Fibre Channel fabric environment. In addition, these ports must support the Directory Service interface and the Name Service.

You can use a fabric configuration that is dedicated solely to the Enhanced Remote Mirroring ports on each controller. In this case, host systems can connect to the storage subsystems by using fabric, Fibre Channel Arbitrated Loop (FC-AL), or point-to-point configurations. These configurations are independent of the dedicated Enhanced Remote Mirroring fabric.

Alternatively, you can use a single Fibre Channel fabric configuration for both the Enhanced Remote Mirroring connectivity and for the host I/O paths to the controllers.

The maximum distance between the primary site and the secondary site is 10 km (6.2 miles), using single-mode fiber gigabit interface converters (GBICs) and optical long-wave GBICs.

**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 DCS3700 FC HIC option.

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

## Considerations

These considerations apply to mirrored logical drive candidates and storage subsystem mirrors:

- ▶ RAID level, caching parameters, and segment size can be different on the two mirrored logical drives.
- ▶ The secondary logical drive must be at least as large as the primary logical drive.
- ▶ The only type of logical drive that might participate in a mirror relationship is a standard logical drive. FlashCopy logical drives cannot participate.
- ▶ You can create a FlashCopy logical drive by using either a primary logical drive or a secondary logical drive as the base logical drive.
- ▶ A primary logical drive can be a source logical drive or a target logical drive in a VolumeCopy. A secondary logical drive cannot be a source logical drive or a target logical drive unless a role reversal was initiated after the copy has completed. If a role reversal is initiated during a Copy in Progress status, the copy fails and cannot be restarted.
- ▶ A given logical drive might participate in only one mirror relationship.

### 2.5.5 Obtaining Premium Features key

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

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

The key can then be added to your DCS3700 system as detailed in “Premium Features” on page 205.

## 2.6 Additional planning considerations

This section reviews additional elements to consider when planning your DCS3700 storage subsystems that use a Logical Volume Manager and virtualization options.

### 2.6.1 Planning for systems with Logical Volume Manager: 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 LVM 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 2-19 illustrates the layout of those components in the case of the AIX LVM.

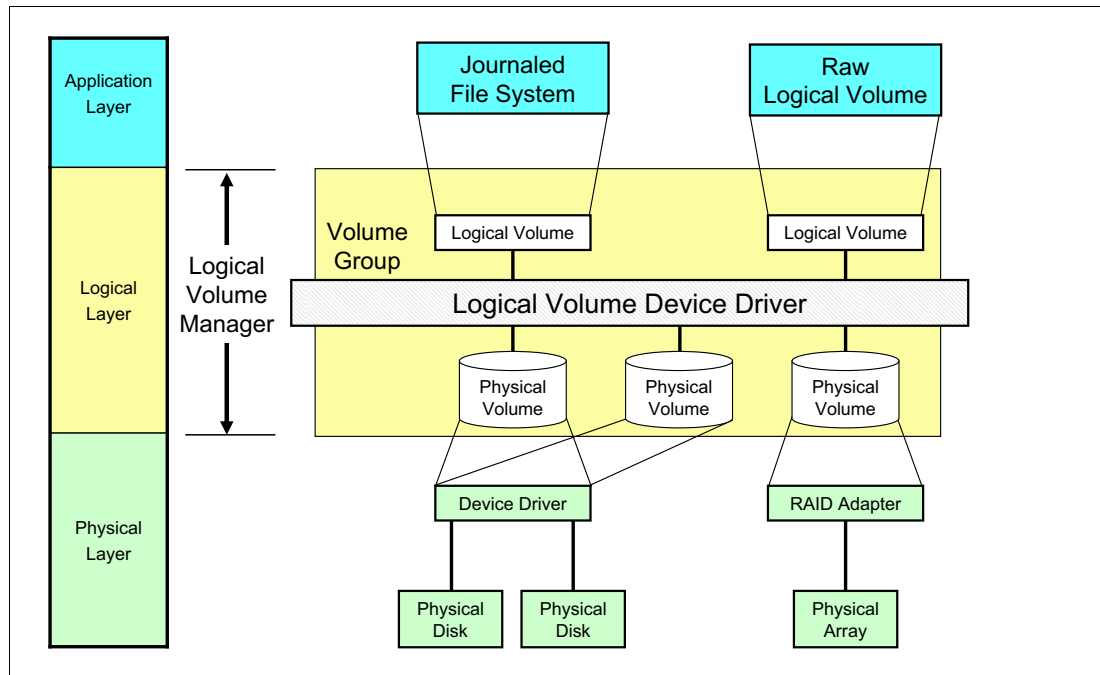


Figure 2-19 AIX Logical Volume Manager

## Hierarchy of structures in disk storage

A hierarchy of structures is used to manage the actual disk storage, and there is a 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 DCS3700 storage subsystem, such physical volumes correspond to a LUN.

Here is the hierarchy of the structures:

- ▶ 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 that are 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 2-20). If mirroring is specified for the logical volume, additional physical partitions are allocated to store the additional copies of each logical partition.

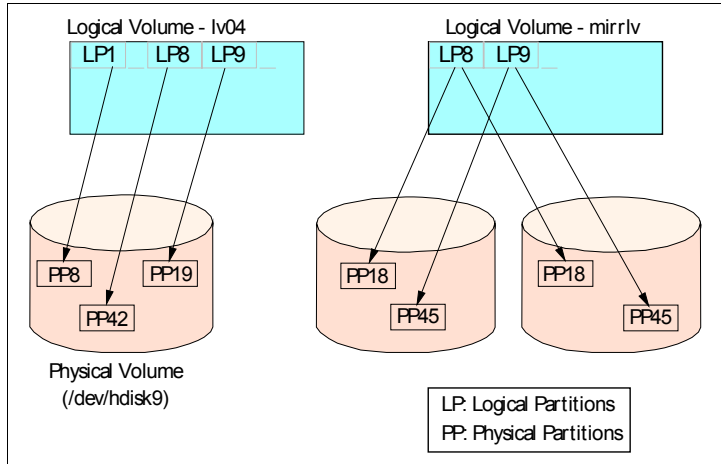


Figure 2-20 Relationships between LPs and PPs

- ▶ 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).

The LVM controls disk resources by mapping data between a simple and flexible logical view of storage space and the actual physical disks. The LVM provides this view 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 2-21).

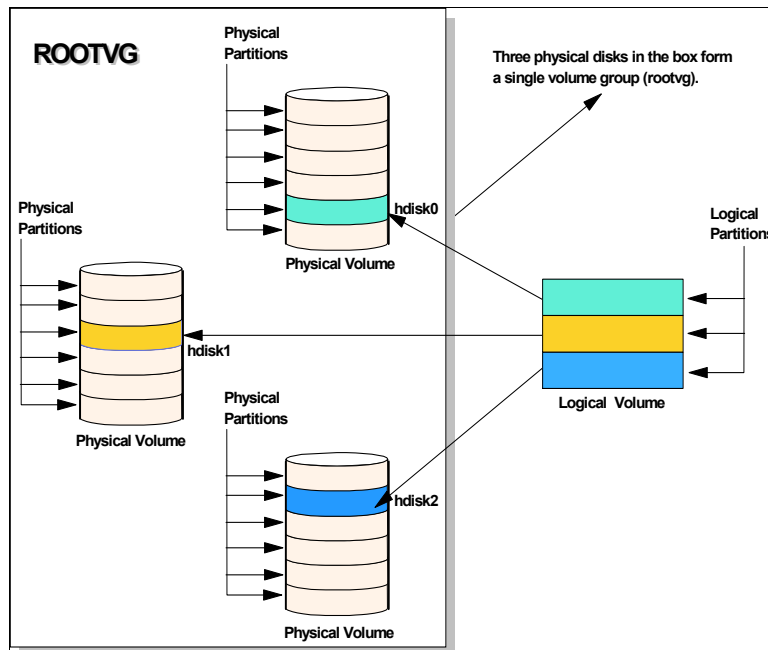


Figure 2-21 AIX Logical Volume Manager conceptual view

**General rule:** When you use a DCS3700 storage subsystem with operating systems that have a built-in LVM, or if an LVM is available, you need to use 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 DCS3700 storage subsystem has its own RAID arrays and logical volumes, you 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 be used only after additional testing, and it is not unusual that trying to use these functions leads to worse results.

However, do 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 DCS3700 storage subsystem, 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 is spread evenly across DCS3700 controllers. Applications that spread their I/Os equally across containers, such as IBM DB2®, use another layout.

If systems are using only one large 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 DCS3700 controllers use.

## 2.6.2 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 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 that are 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 DCS3700 storage subsystem, 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 use basic disks.

When using the DCS3700 storage subsystem, the usage of software mirroring and software RAID 5 is not required. Instead, configure the storage on the DCS3700 storage subsystem for the redundancy level that is 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 multidisk direct server-attached (DAS) storage, JBODs (just a bunch of disks), storage area networks (SANs), and RAID better.

The main benefit is the ability to define subdisks and disk groups. You can divide a dynamic disk into one or more subdisks. A subdisk 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 subdisk is a portion of a dynamic disk's public region.

A subdisk is the smallest unit of storage in Volume Manager. Therefore, subdisks are the building blocks for Volume Manager arrays. A subdisk 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 book. 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](http://www.symantec.com/enterprise/products/overview.jsp?pcid=1020&pvid=203_1)

- ▶ You can use the DCS3700 storage subsystem 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 options:
  - Spanned volumes combine areas of unallocated space from multiple disks into one logical volume. The areas of unallocated 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 that is written across each disk in the volume. They are created from equally sized and unallocated 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 that is configured on the DCS3700 storage subsystem. 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 that you choose suits your needs and requirements.

### 2.6.3 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 usage in a heterogeneous environment.

The IBM System Storage SAN Volume Controller and IBM Tivoli® Storage Productivity Center products address these needs.

### 2.6.4 SAN Volume Controller overview

The SAN Volume Controller is a scalable hardware and software solution to allow aggregation of storage from various disk systems. 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 subsystems. These managed disks are then mapped to a set of virtual disks for use by various host systems.

With the DCS3700 storage subsystem family, the SAN Volume Controller can increase the storage copy services functionality and also the flexibility of SAN-based storage. The SAN Volume Controller is flexible in its usage. It can manage all host storage requirements or just part of them. A DCS3700 storage subsystem can still be used to allocate storage to hosts or use the SAN Volume Controller, which is dependent upon various needs and requirements.

The SAN Volume Controller also offers an alternative to FlashCopy, VolumeCopy, and Enhanced Remote Mirroring for disaster recovery, high availability, and maintenance. If the use of SAN Volume Controller with a DCS3700 storage subsystem is planned, then these premium features are not required.

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

SAN Volume Controller is licensed by the capacity that is being managed. This capacity also includes the capacity that is used by the copy services.

For detailed information about the SAN Volume Controller implementation, see *Implementing the IBM System Storage SAN Volume Controller V6.3*, SG24-7933.

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:

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

## 2.7 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 DCS3700 series interoperability matrix, available at:

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

Section 2.7.4, “Multipathing” on page 89 describes the available multipathing drivers and their supported operating systems.

### 2.7.1 Supported server platforms

The following server platforms are supported:

- ▶ IBM System x
- ▶ IBM System p
- ▶ IBM BladeCenter Servers

### 2.7.2 Supported operating systems

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

- ▶ Microsoft Windows Server 2003
- ▶ Microsoft Windows Server 2008
- ▶ Microsoft Windows Server 2008 R2
- ▶ Red Hat Enterprise Linux 6.1
- ▶ Red Hat Enterprise Linux 6.2
- ▶ Novell SUSE SLES 10 SP4
- ▶ Novell SUSE SLES 11 SP2
- ▶ VMware ESX Server 3.5, 4.0 and 4.1
- ▶ VMware ESXi 4.0, 4.1, and 5.0
- ▶ IBM AIX V6.1 TL7
- ▶ IBM AIX V7.1 TL1

## 2.7.3 Clustering support

The following clustering services are supported:

- ▶ Microsoft Cluster Services
- ▶ Red Hat Cluster Suite
- ▶ VMware Fault Tolerance (FT)
- ▶ VMware High Availability (HA)

## 2.7.4 Multipathing

IBM offers various multipath drivers that you can use with your DCS3700 storage subsystem. 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 connects 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 2.4.8, “Logical drives and controller ownership” on page 66) 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 alternative 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/ssic/interoperability.wss>

## 2.7.5 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 DCS3700 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 that are presented by the DCS3700 storage subsystems 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 the following items:
  - I/O routing information
  - Conditions requiring a request to be tried again, 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.

For more information about multipathing and failover considerations, see 2.7.9, "Function of the Auto-Logical Drive Transfer feature" on page 92.

For compatibility information, always see the current DCS3700 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>

## 2.7.6 AIX MPIO

With multipath 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 perform the following actions:

- ▶ Check the paths and determine which paths are 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 unused paths that might be used if a failover occurs (for example, when the algorithm attribute value is failover, the health check can test the alternative paths).

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

## 2.7.7 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 MPIO devices, SDDPCM is loaded as part of the AIX MPIO Fibre Channel Protocol (FCP) 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)

## 2.7.8 Linux: RHEL/SLES

This section describes 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 DCS3700 storage subsystems (*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 that require a request to be tried again, failed, or failed over.
- ▶ Automatically fails over the LUNs to their alternative controller when problems in sending I/Os to the LUNs in their preferred controller are detected, and fails back the LUNs to their preferred controller when the problems in the preferred path are 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 2-22), 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 2-22 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.

## Device Mapper Multipathing for Linux

Device Mapper Multipathing (DM-Multipath) allows you to configure multiple I/O paths between server nodes and storage arrays into a single device. These I/O paths are physical SAN connections that can include separate cables, switches, and controllers. Multipathing aggregates the I/O paths, creating a new device that consists of the aggregated paths, providing redundant failover/failback support for the logical drives in the DCS3700 storage subsystems that are mapped to the Linux host server.

DMM for Linux is supported by Red Hat Enterprise Linux (RHEL) or SUSE Linux Enterprise Server (SLES).

For detailed instructions about how to implement and use DMM, see “Linux Device mapper multipathing driver (DMM / DM-Multipath)” on page 533.

## 2.7.9 Function of the Auto-Logical Drive Transfer feature

In a DCS3700 storage subsystem that is 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 subsystems with the flexibility to work with certain third-party failover software.

► ADT-disabled failover:

The multipath software sends a SCSI Mode Select command to cause a change in volume ownership before it uses the alternative path. All logical drives on the preferred controller are transferred to the alternative 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 if you use a multipath 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 2-23 shows the ADT-disabled failover mode phases.

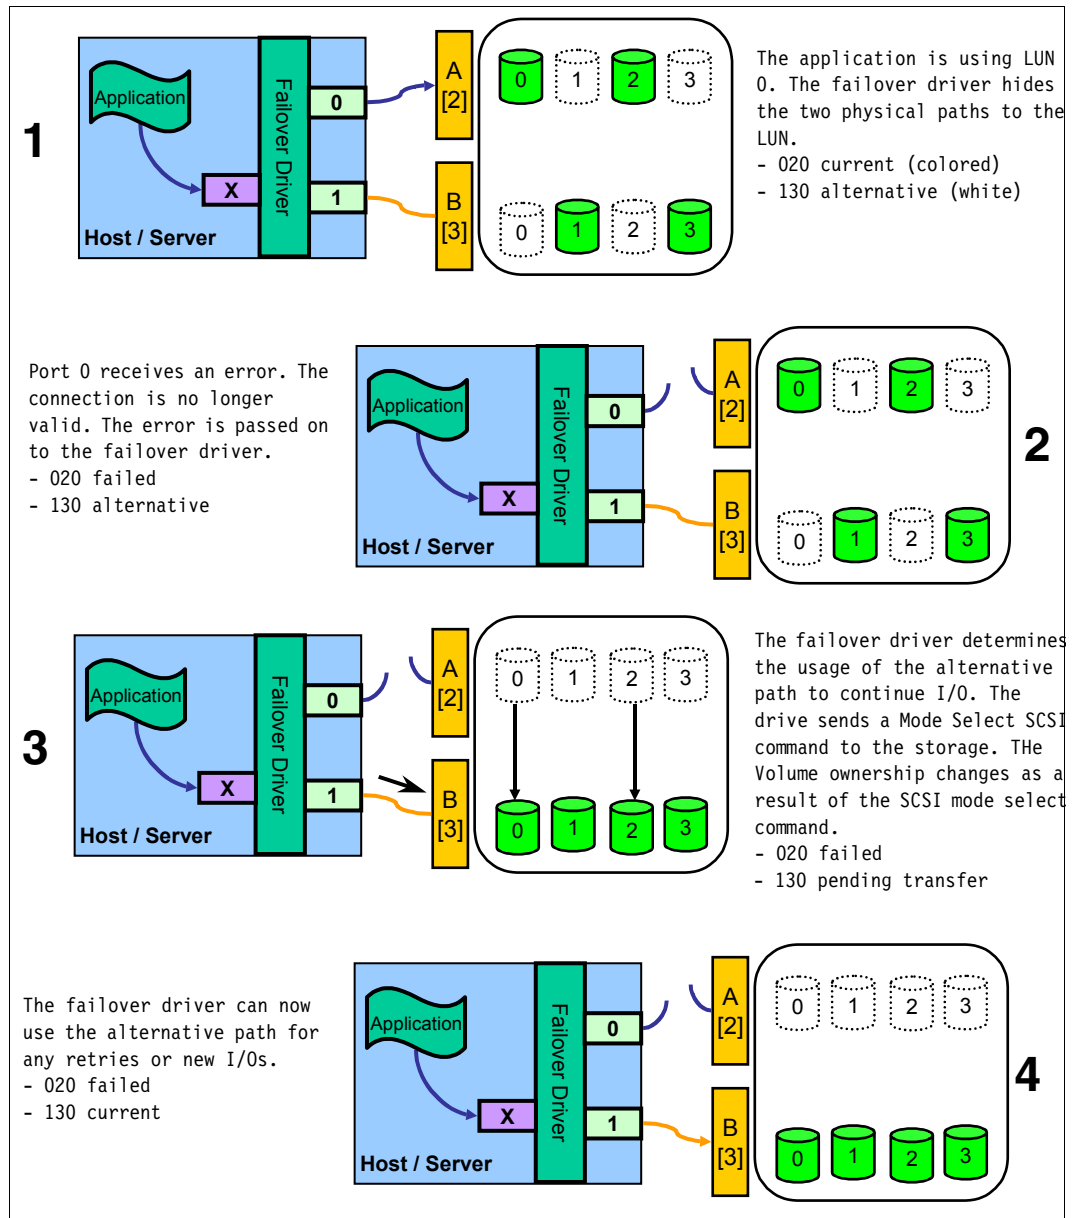


Figure 2-23 ADT-disabled mode path failover

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

► ADT-enabled failover:

The multipath driver starts using the alternative 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 when the multipath driver detects that the path is normal again.

## Phases of failover

Figure 2-24 shows the phases of a failover in an ADT-enabled case.

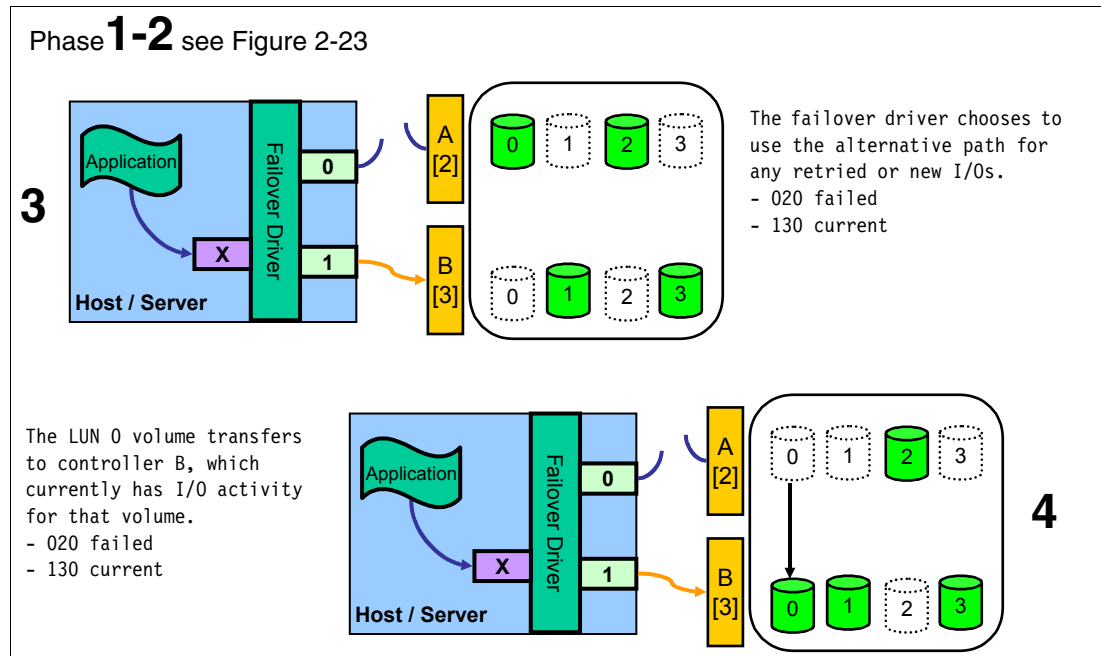


Figure 2-24 ADT-enabled mode path failover

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

### 2.7.10 Asymmetric logical unit access

Asymmetric logical unit access (ALUA) is an industry standard protocol that enables the communication of storage paths and path characteristics between an initiator port and a target port. This process occurs when the access characteristics of one port might differ from those of another port. A logical unit may be accessed from more than one target port. One target port might provide full performance access to a logical unit, while another target port, possibly on a different physical controller, either might provide lower performance access or might support a subset of the available SCSI commands to the same logical unit.

Using ALUA, you can make these improvements:

- ▶ Flexible configuration accommodation
- ▶ Automatic interface / ownership optimization
- ▶ Eliminates environmental complexity
- ▶ Faster deployment of new OS / HBA drivers

ALUA provides the following benefits:

- ▶ SAN boot is reliable.
- ▶ Volume ownership thrashing is eliminated.
- ▶ Failovers are smoother.
  - Especially on RVM systems
  - Firmware Download

- ▶ Clusters work better.
  - Faster boots
  - Less affected by connectivity problems
- ▶ Prevents failover/failback for transient errors.
- ▶ More usage of industry standard multipath drivers.

Figure 2-25 shows the differences between using multipathing and ALUA.

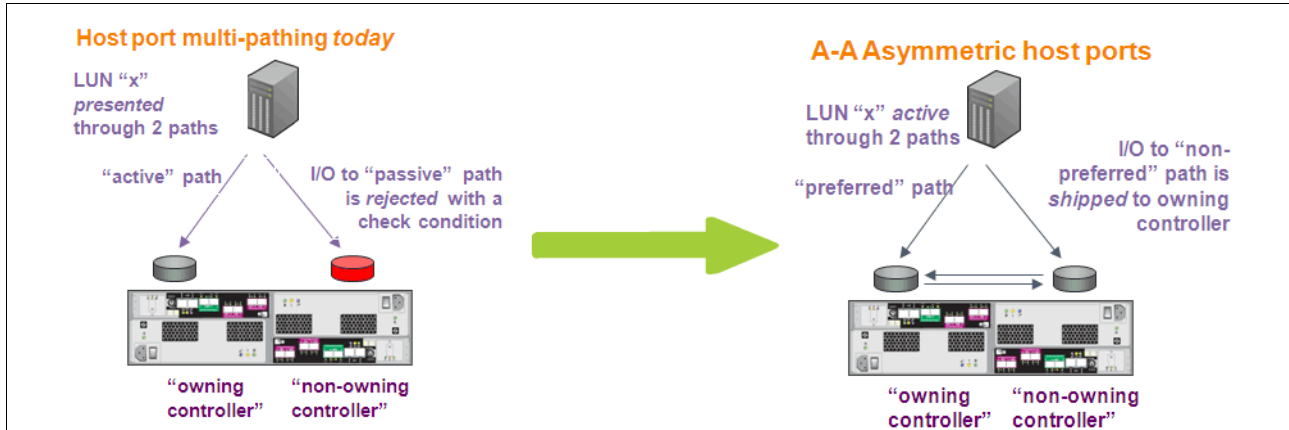


Figure 2-25 Multipathing versus ALUA

## 2.8 Operating system considerations

In your planning, you must also consider the limitations of your operating system. This section covers the maximum file system size ordered by 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 that is supported by various operating systems.

### 2.8.1 Maximum capacity for a logical drive

Table 2-11 shows the maximum capacity that a logical drive can have and be formatted with a specific file system.

Table 2-11 Maximum file system size that is supported by various operating systems

Operating System	File system	Size
Linux	ext2	32 TB
	ext3	32 TB
	ext4	1 EB
AIX	JFS2	64 PB
Windows	NTFS	16 EB
	FAT32	8 TB
ESX Server 3.5, 4.0, and 4.1 / ESXi 4.1	VMFS3	2TB
ESXi 5.0	VMFS5	64 TB

## 2.8.2 Maximum number of LUNs per host

The maximum number of LUNs that can be assigned to a host depends on the host operating system type and mapping limit within a storage partition on the DCS3700 storage subsystem. The DCS3700 storage subsystem supports 256 volumes per partition. You can find more information about storage partitioning in 2.4.9, “Storage partitioning” on page 67. Table 2-12 lists the operating system limit and the limits of the DCS3700 storage subsystem.

Table 2-12 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 (DCS3700 storage subsystem limit)
AIX	64000	255
Linux	256	255
Windows	256	255
ESX Server 3.5, 4.0 and 4.1 / ESXi 4.1 and 5.0	256	255

**Note:** The maximum number of logical drives on a DCS3700 storage subsystem is 512, and 2048 with Performance Module Controller DCS3700P.



# IBM System Storage DCS3700 and DCS3700 expansion unit cabling

This chapter describes the guidelines for the cabling of the DCS3700 storage subsystem and the DCS3700 expansion unit. It introduces the various connector ports that are on the rear of the DCS3700 controller and the DCS3700 expansion unit, and then describe the following guidelines:

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

This chapter also describes cable management and the correct power-on/off procedures for the DCS3700 storage subsystem.

For more information about all the topics that are described in this chapter, see the *Installation, User's, and Maintenance Guide - IBM System Storage DCS3700*, found at:

<http://www.ibm.com/support/entry/portal/docdisplay?lnodocid=MIGR-5087760>

## 3.1 DCS3700 controller connectors

This section describes the various connectors that are on the rear panel of the DCS3700 controller model and expansion.

### 3.1.1 DCS3700 controller with standard port configuration

Figure 3-1 shows the rear panel of a DCS3700 controller with the standard port layout.



Figure 3-1 The rear panel of a DCS3700 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 that are labelled 1 and 2 for Out-Of-Band (OOB) management.
- ▶ SAS Host Ports labeled 1 and 2: These are two SAS 2.0 ports for host-side connections. The DCS3700 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 adapters would be located. You can install one of two types of optional host adapters if you need additional SAS host port support or concurrent host connectivity with SAS and FC.
- ▶ The hole between the two plates is the password reset button. Use a thin metal tip to press the button. It is only necessary to do this on one controller.
- ▶ The port on the extreme right is a x4 multilane mini-SAS port for connecting expansion units.

Below SAS ports 1 and 2 there is a 7-segment numeric LED display that provides information about enclosure identification and diagnostic information.

### 3.1.2 The DCS3700 controller with an optional SAS host port adapter

Figure 3-2 on page 99 shows the rear panel of a DCS3700 controller with an optional SAS host port adapter installed. In addition to the standard ports that are described in Figure 3-1, the optional SAS host port adapter adds an additional two SAS 2.0 host ports for a total of four per controller.





Figure 3-2 The rear panel of a DCS3700 controller with the optional SAS host port adapter installed

The SAS host ports labeled 1, 2, 3, and 4 are SAS 2.0 ports for host-side connections. SAS ports 3 and 4 are on the optional SAS host port adapter. Each DCS3700 host port is a x4 multilane, 6 Gbps universal mini-SAS port.

### 3.1.3 DCS3700 controller with an optional Fibre Channel host port adapter

Figure 3-3 shows the rear panel of a DCS3700 controller with an optional four port Fibre Channel host adapter installed.




Figure 3-3 The rear panel of a DCS3700 controller an optional Fibre Channel host port adapter installed

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

### 3.1.4 DCS3700 Performance Module

Figure 3-4 shows the rear panel of a DCS3700 Performance Module.




Figure 3-4 DCS3700 Performance Module rear panel with an optional FC host port adapter

The ports from left to right are USB for Development and debugging, Ethernet Management ports, 4-Port Base 8 Gb FC ports, an optional 4-Port 8 Gb Fibre Channel Host Interface Card, and the SAS Expansion port. On the lower left side, there are the Controller LEDs and a dual 7-segment numeric LED display that provides information about enclosure identification and diagnostic information. On the upper right side is the hole for password reset and a Mini-USB Serial RS232 port, which is covered by default. The LEDs from left to right are Battery fault, Battery charging, Ready to remove, Controller fault, and Cache active.

Figure 3-5 shows the 10 Gb iSCSI optional daughter card that is installed in the DCS3700 Performance controller. With this option, the system has up to four 10 Gb iSCSI connections for host attachments. An optional SAS HIC is also available to use in the Performance Module.



Figure 3-5 Controller module with an optional iSCSI host interface daughter card

### 3.1.5 DCS3700 expansion unit ports

The DCS3700 expansion unit comes with two ESMs, which enable you to connect the DCS3700 expansion unit to a DCS3700 storage subsystem.

Figure 3-6 shows the port layout on the rear of an expansion unit.



Figure 3-6 The rear panel of a DCS3700 expansion unit.

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

- ▶ Two SAS expansion ports labeled 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 that is reserved for use by IBM Service personnel only.
- ▶ An Ethernet port that is reserved for use by IBM Service personnel only.
- ▶ A SAS expansion port. This port is for connecting to other DCS3700 expansion units 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 diagnostic information.

## 3.2 Enclosure ID settings

The enclosure ID is a unique two-digit identifier for a DCS3700 and its expansion unit in the storage subsystem configuration. In a dual controller DCS3700 system, both controllers should have the same enclosure ID. Similarly, on an expansion unit, both ESMs must have the same enclosure ID under normal operating conditions.

The default enclosure ID of the DCS3700 controller is 99, and you should leave this value unchanged. The enclosure ID of each expansion is set to a value of 00 at the factory. The DCS3700 controller firmware automatically assigns a unique enclosure ID to each expansion unit that you connect to it. The automatically assigned expansion unit IDs begin with enclosure ID 0, then 1, 2, and so on. The range of supported enclosure ID values is 0 - 99.

You can view and change the enclosure IDs by using the DS Storage Manager software.

In normal operation, the enclosure ID is also displayed on the seven-segment numeric LED display on the back of each controller and ESM. Figure 3-7 shows the seven-segment numeric LED, displaying the enclosure ID.




Figure 3-7 The seven-segment numeric LED displaying the enclosure ID

**Note:** The seven-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 have conflicting hardware addresses.

## 3.3 SAS cables

This section describes the SAS cables that you can use to connect the DCS3700 to hosts and to DCS3700 expansion units.

Each DCS3700 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 m and 3 m cable lengths for connecting to expansion units and for SAS connections to hosts. These SAS cables are also used to connect expansion units together.

Figure 3-8 shows a typical SAS cable with mini-SAS 4x multilane SAS connectors on both ends of the cable.




Figure 3-8 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 supported.

To connect a mini-SAS cable, insert the mini-SAS connector into a mini-SAS port, as shown in Figure 3-9. Make sure that it locks into place.




Figure 3-9 Connecting a mini-SAS cable

To remove a mini-SAS cable, as shown in Figure 3-10, 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.




Figure 3-10 Removing a mini-SAS cable

## 3.4 Fibre Channel cabling

This section describes the features of Fibre Channel cabling that you use to connect the DCS3700 controllers to hosts with Fibre Channel host adapters or to Fibre Channel SAN switches. This section also describes general guidelines for cable management of the fiber optical cables.

The DCS3700 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 3-11 shows a Fibre Channel cable with an LC connector that is attached to an SFP module.



Figure 3-11 Small Form Pluggable (SFP) module with an LC connector fiber cable

### 3.4.1 SFP transceiver modules

Each DCS3700 controller can be installed with an optional Fibre Channel host adapter, and can support up to four Fibre Channel connections as shown in Figure 3-3 on page 99. The DCS3700 Performance Module has up to eight Fibre Channel connections. 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 fiber 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 DCS3700 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 DCS3700 storage subsystem. SFP modules are also available in 2 Gbps, 4 Gbps, and 8 Gbps models. We recommend that you use only 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 DCS3700 storage subsystem controllers. Long-wave SFP modules and GBICs are not supported for use in the DCS3700 Storage System.

#### 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 prevents 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 damage the SFP module or 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 4 Gbps SFP module that is connected to an 8 Gbps-capable port limits the speed of that port to a maximum of 4 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 3-12. Save the protective cap in a dust-free place for future use.




Figure 3-12 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 3-13.




Figure 3-13 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 3-14.




Figure 3-14 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 3-15.




Figure 3-15 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.

### 3.4.2 Fibre Channel cables

This section describes the fiber optic cables that you can use for connecting to the small form-factor pluggable (SFP) modules in the DCS3700 storage subsystem. This section also describes general rules for cable management of fiber optic cables.



Because the DCS3700 supports only SFP modules, Fibre Channel cables with LC connectors are required. Connectors that plug into SFP or SFP devices are called LC connectors. The two fibers 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 3-16 shows an LC connector at one end of a Fibre Channel cable.




Figure 3-16 An LC connector

Most Fibre Channel cables have an LC connector at both ends, which are called LC-LC cables. You use LC-LC cables when you are connecting the DCS3700 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 supported.
- ▶ The following are the recommended maximum cable lengths:
  - 2 Gbps: 300 m (984 ft.) of 50/125  $\mu$ m fiber, 150 m (492 ft.) of 62.5/125  $\mu$ m fiber
  - 4 Gbps: 150 m (492 ft.) of 50/125  $\mu$ m fiber, 50 m (164 ft.) of 62.5/125  $\mu$ m fiber
  - 8 Gbps: 50 m (164 ft.) of 50/125  $\mu$ m fiber, 21 m (69 ft.) of 62.5/125  $\mu$ m fiber
  - 16 Gbps: 35m (115 ft.) of 50/125  $\mu$ m fiber, 15 m (49 ft.) of 62.5/125  $\mu$ m fiber

**Note:** For more information about the supported distances for the OM1, OM2 and OM3 category Fibre Channel cables, see Table 2-1 on page 33.

### 3.4.3 Interoperability of 2 Gbps, 4 Gbps, 8 Gbps, and 16 Gbps devices

The Fibre Channel standard specifies a procedure for speedy auto-detection. For example, if an 8 Gbps port on a switch or device is connected to a 4 Gbps port, it must negotiate down and the link runs at 4 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 16 Gbps devices relative to 8 Gbps, 4 Gbps, and 2 Gbps environments. The 16 Gbps, 8 Gbps, and 4 Gbps devices automatically can negotiate back down to either 8 Gbps, 4 Gbps, or 2 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 matter ought to be investigated and remedied.

The DCS3700 storage subsystem 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, 8 or 16 Gbps to obtain the required speed instead of leaving “auto-detection” on the port.

## 3.5 iSCSI Ethernet cabling

This section describes 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 DCS3700 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.

## 3.6 DCS3700 expansion unit attachment

This section describes the requirements and general guidelines for connecting DCS3700 expansion units to a DCS3700 storage subsystem.

The DCS3700 supports attachment of up to 300 drives in five DCS3700 expansion units. So, the total number of disk drives in the storage configuration does not exceed 360.

### 3.6.1 Redundant drive channels

Each DCS3700 controller has a drive expansion channel with one x4 SAS port. The DCS3700 expansion units that are connected to this port from one drive channel. Each drive channel supports a maximum of 360 disk drives. In a dual-controller DCS3700, one drive channel from each controller combines to form a redundant drive-channel pair. If any component in the drive channel fails, the DCS3700 controllers can still access the DCS3700 expansion units in the redundant drive-channel pair.

## 3.6.2 Drive channel cabling rules

You should always follow these two basic connection rules when attaching DCS3700 expansion units to a DCS3700 storage subsystem:

- ▶ Connect the drive expansion SAS port on the DCS3700 controller to the SAS IN port on the DCS3700 expansion unit.
- ▶ Connect the SAS OUT port on the DCS3700 expansion unit to the SAS IN port on the next DCS3700 expansion unit.

## 3.6.3 Connecting the DCS3700 expansion unit

The DCS3700 storage subsystem supports up to five DCS3700 expansion units per physical expansion port, so multiple DCS3700 expansion units can be connected by chaining them together.

To connect the DCS3700 storage subsystem RAID controllers A and B to a single DCS3700 expansion unit, complete the following steps:

1. Connect the DCS3700 expansion unit to controller A:
  - a. Connect one end of a SAS cable to the drive expansion port on controller A.
  - b. Connect the other end of the SAS cable to one of the In (ff) SAS connectors on the ESM in slot A in the DCS3700 expansion unit.
2. Connect the DCS3700 expansion unit to controller B:
  - a. Connect one end of a SAS cable to the drive expansion port on controller B.
  - b. Connect the other end of the SAS cable to one of the In (ff) SAS connectors on the ESM in slot B in the DCS3700 expansion unit.

Figure 3-17 shows a DCS3700 storage subsystem that is configured with a single DCS3700 expansion unit.

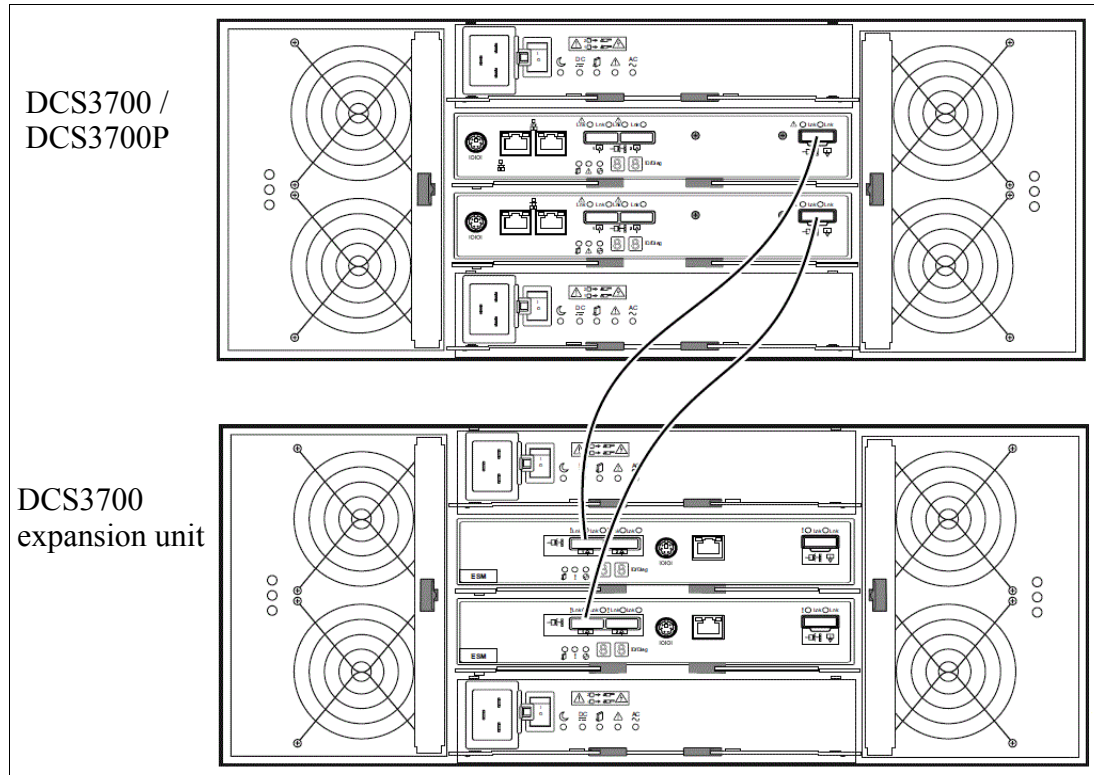


Figure 3-17 Single expansion units

To connect the DCS3700 storage subsystem RAID controllers A and B to multiple DCS3700 expansion enclosures, complete the following steps:

1. Connect the DCS3700 expansion unit to controller A:
  - a. Connect one end of a SAS cable to the drive expansion port on controller A.
  - b. Connect the other end of the SAS cable to one of the In (ff) SAS connectors on the ESM in slot A in the first DCS3700 expansion unit.
2. Connect the DCS3700 expansion unit to the next DCS3700 expansion unit in the chain:
  - a. Connect one end of a SAS cable to the Out (ff) SAS connector on the ESM in slot A of the DCS3700 expansion unit that you connected.
  - b. Connect the other end of the SAS cable to one of the In (ff) SAS connectors on the ESM in slot A in the next DCS3700 expansion unit in the chain.
3. Connect the last DCS3700 expansion unit in the chain to controller B:
  - a. Connect one end of a SAS cable to the drive expansion port on controller B.
  - b. Connect the other end of the SAS cable to one of the In (ff) SAS connectors on the ESM in slot B in the last DCS3700 expansion units of the chain that you created.

4. Connect the second DCS3700 expansion unit to the first DCS3700 expansion unit.
  - a. Connect one end of a SAS cable to the Out (ff) SAS connector on the ESM in slot B of the second DCS3700 expansion unit.
  - b. Connect the other end of the SAS cable to one of the In (ff) SAS connectors on the ESM in slot A of the first DCS3700 expansion unit.

Figure 3-18 shows the cabling for dual expansion units




Figure 3-18 DCS3700 with two DCS3700 expansion units

### 3.6.4 Adding a DCS3700 expansion unit to a running dual-controller configuration

This section describes the steps that are required to attach a DCS3700 expansion unit to a running DCS3700 configuration. There are two possible scenarios when attaching a DCS3700 expansion unit to a running DCS3700 configuration:

- ▶ There are no DCS3700 expansion unit enclosures already attached.
- ▶ There are DCS3700 expansion unit enclosures already attached.

## Adding one DCS3700 expansion unit to a running DCS3700 configuration with no DCS3700 expansion unit attached

In this scenario, there are no DCS3700 expansion units that are already attached to the DCS3700. To attach the DCS3700 expansion unit to the DCS3700, complete the following steps:

1. You must ensure that the DCS3700 storage subsystem status is Optimal in the DS Storage Manager software.
2. Connect the power cables to the new expansion unit and power-on the expansion unit.
3. Connect one end of a Serial Attached SCSI (SAS) cable to the drive expansion port on Controller A in the DCS3700 storage subsystem.
4. Connect the other end of the SAS cable to either of the SAS IN ports on the left ESM in the DCS3700 expansion unit.
5. Connect one end of a second SAS cable to the drive expansion port on Controller B in the DCS3700 storage subsystem.
6. Connect the other end of the second SAS cable to the IN port on the right ESM in the expansion unit.
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 unit 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 a DCS3700 expansion unit to a running DCS3700 configuration with DCS3700 expansion unit attached

In this scenario, there are already DCS3700 expansion unit enclosures that are attached to a running DCS3700 configuration. To attach the DCS3700 expansion unit to the DCS3700, complete the following steps:

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 unit and power on the expansion unit.
3. Connect one end of a Serial Attached SCSI (SAS) cable to the IN port on the left ESM in the expansion unit 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 unit in the configuration (this expansion unit is directly connected to the drive expansion port on controller A of the DCS3700).
5. On the last expansion unit 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 unit that is being added to the configuration.

**Note:** When you disconnect the SAS cable in step 5, there is a loss of drive path redundancy and the DS Storage Manager Recovery Guru reports an error message. You can safely disregard this error message because you restore a redundant drive path in step 6 on page 113 when you connect the SAS cable again.

6. Connect a SAS cable to the OUT port on the right ESM of the new expansion unit and connect the other end to either of the IN ports on the right ESM in the expansion unit that was previously connected to controller B of the DCS3700.
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 unit.

**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.

### 3.6.5 SAS drive channel miswires

If the SAS cabling on the drive side is incorrect, it is detected by the DCS3700 controller firmware. The Storage Manager application alerts you about the condition by logging a critical Major Event Log event. In addition, the Storage Manager Recovery Guru points to the miswire condition and advises you of the steps to take to correct the problem. For more information about failure recovery, see 13.2, “View Health (Recovery Guru)” on page 387.

A miswire condition is reported only if the wrong SAS connections result in a non-working configuration. It is also possible to attach the DCS3700 expansion unit in a technically correct manner that is not optimal. Such configurations do not produce a miswire event. Therefore, you should carefully follow the recommendations here for best results.

Here are a few examples of miswired configurations:

- ▶ DCS3700 expansion units are connected to SAS host ports instead of to SAS drive ports on the DCS3700 controller.
- ▶ The SAS IN port is connected to the SAS IN port on another DCS3700 expansion unit.
- ▶ The SAS OUT port is connected to the SAS OUT port on another DCS3700 expansion unit.
- ▶ SAS ports are interconnected within a single DCS3700 expansion unit.

An example of a non-optimal, but technically correct configuration, is connecting both left and right ESMs in a 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 DCS3700 expansion unit fails, the DCS3700 expansion units beneath it all lose access to the DCS3700.

## 3.7 Management connections

There are two connection methods that you can use to manage the DCS3700 storage subsystem: host-agent (in-band) management and direct (out-of-band) management.

**Note:** 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 cannot correctly manage the enclosure.

The installation and configuration of in-band and out-of-band management connections is described in Chapter 4, “IBM System Storage DS Storage Manager installation” on page 131.

### 3.7.1 Out-of-band management

Out-of-band (OOB) management uses a network connection to the TCP/IP Ethernet ports on the rear panels of each DCS3700 controller (see Figure 3-19). Typically, you need to connect only 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.

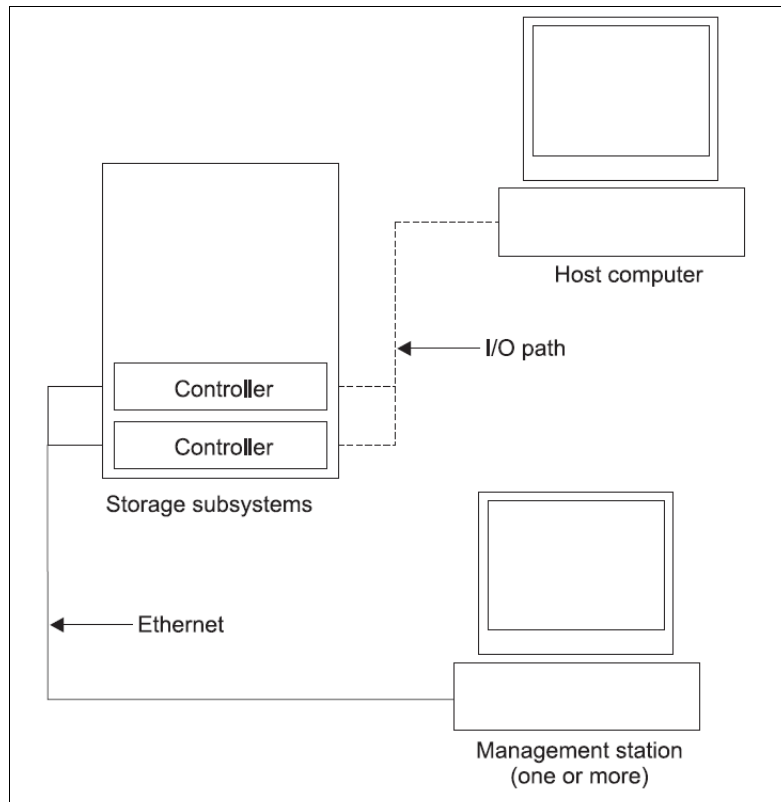


Figure 3-19 Out of band management connections

**General rule:** We recommend that you set up a dedicated management LAN or VLAN for OOB management of the DCS3700, and avoid using your public LAN for OOB management of the DCS3700. This configuration provides increased security for your DCS3700.

#### Advantages and disadvantages of out-of-band management

This section describes the advantages and disadvantages of using the out-of-band management method.

Managing storage subsystems by 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 subsystems 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 3.7.2, “In-band management” on page 115) 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 subsystem 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 DCS3700 storage subsystem in the network.

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

- ▶ Two Ethernet cables are required to connect the storage subsystem 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. These tasks 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.

### 3.7.2 In-band management

When you use the in-band management method (also known as the host-agent method), the DCS3700 storage subsystem is managed through the host I/O connections to the host server (see Figure 3-20). 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 by 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 that are associated with your particular combination of controller, operating system, and host bus adapter.




Figure 3-20 In-band management connections

## Advantages and disadvantages of in-band management

This section describes the advantages and disadvantages 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 be specified only for the host instead of for the individual controllers in a storage subsystem. 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 subsystem. 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 cannot manage or monitor the DCS3700 storage subsystem if there is a fault with the host I/O connection between the DCS3700 storage subsystem and the host server.

## 3.8 Host attachment

This section describes the general rules that you should follow when cabling the hosts and DCS3700 storage subsystem. It also describes the general rules for Fibre Channel switch SAN zoning. We provide examples of the various host attachment implementations with both single and dual controller DCS3700 storage subsystems using SAS and Fibre Channel.

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:** For a DCS3700 storage subsystem, you must ensure that each host server has a path to both controllers. A configuration where a host server is connected only to one controller of a dual controller DCS3700 storage subsystem is not a supported configuration. This requirement applies to SAS and FC attachment.

The DCS3700 supports multiple-host, multiple-port, and multiple-fabric (Fibre Channel or SAS) configurations. The DCS3700 storage subsystem supports simultaneous SAS and Fibre Channel host connections with the optional Fibre Channel host adapter installed, up to a maximum of 512 hosts.

Figure 3-21 shows an example of this type of configuration with one direct-attached SAS host, and one fabric-attached FC host.

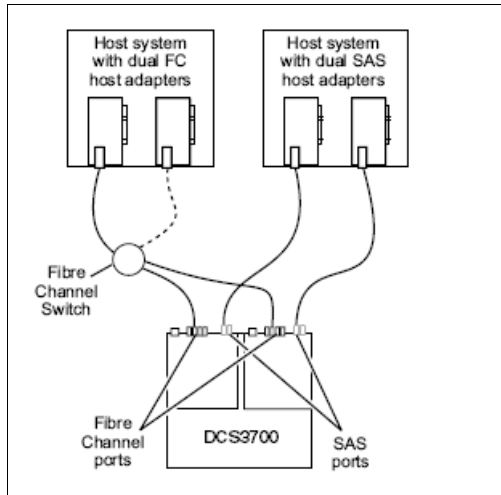


Figure 3-21 An example of a multi-host, multi-port, and multi-fabric host configuration

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 DCS3700 firmware and DS Storage Manager that you want to use.

### 3.8.1 SAS attachment

Host server SAS HBAs are used to connect to the DCS3700 SAS ports. Each DCS3700 controller can have up to two SAS host ports. Always use redundant host connections. This section shows examples of SAS host attachment with dual controller DCS3700 storage subsystems.

Figure 3-22 shows a host server with two SAS host bus adapters connecting to a dual controller DCS3700 storage subsystem. This particular configuration has redundant host side connections and redundant controllers in the DCS3700 storage subsystem.




Figure 3-22 Single host with two SAS connections to a DCS3700 storage subsystem

Figure 3-23 shows a configuration with two hosts that are connected by using redundant SAS connections to a dual controller DCS3700 storage subsystem. This is a good configuration because it has redundant paths to redundant controllers.

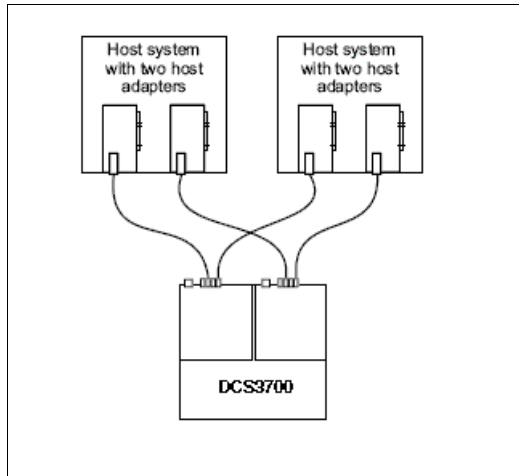


Figure 3-23 Two hosts with redundant SAS connections to a dual controller DCS3700

### 3.8.2 Direct attached Fibre Channel

The DCS3700 storage subsystem supports both direct attached and SAN fabric attached hosts. You can attach up to four direct-connected hosts or up to 512 host servers in a switched fabric SAN infrastructure. This section describes example configurations of direct-attached Fibre Channel host servers.

#### Direct-attached DCS3700 examples

If you are connecting four or less host servers by using Fibre Channel, you can directly connect the four servers to the DCS3700.

Figure 3-24 shows a configuration with one host and redundant paths to a dual controller DCS3700. This configuration offers both redundant Fibre Channel host connections and redundant controllers, and offers improved redundancy.




Figure 3-24 Direct attached host with redundant host connections and controllers

The sample configuration that is shown in Figure 3-25 has two direct-attached host servers with redundant connections to a dual controller DCS3700. 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 DCS3700 storage subsystem.

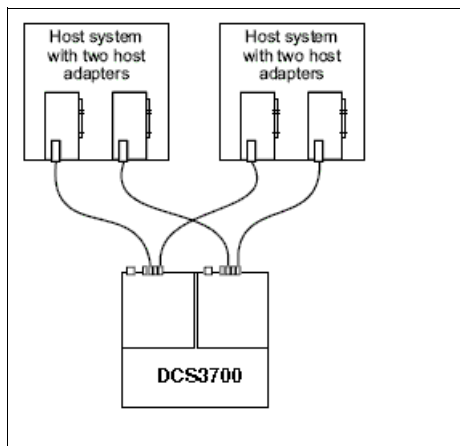


Figure 3-25 Two hosts with redundant connections to a dual controller DCS3700

If you need to connect more than four Fibre Channel hosts to a DCS3700 storage subsystem by using redundant host connections, you need to use a SAN fabric infrastructure, as described in 3.8.3, "SAN fabric-attached DCS3700" on page 120.

### 3.8.3 SAN fabric-attached DCS3700

If you want to connect more than four host servers by using the Fibre Channel interface, you must use a switched fabric SAN infrastructure. This section describes the general characteristics and requirements of a SAN infrastructure, introduces SAN zoning requirements, and shows 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, and 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 by using either standard connections or a type of connection that is 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 that is connected to four ports on a second switch. At the outset, a trunk master is defined, and subsequent trunk subordinates can be added, which aggregates the throughput across all links. Therefore, in the case of switches with 8 Gbps ports, if you trunk up to four ports, you allow for a 32 Gbps Inter-Switch Link.

#### **SAN zoning**

When you connect a DCS3700 storage subsystem to a switched fabric SAN infrastructure, you normally need to implement a *SAN zoning* configuration on the SAN switch. This section provides a brief overview of SAN zoning, and describes 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 accomplish the following tasks:

- ▶ Provide security: Use zones to provide controlled access to fabric segments and to establish barriers between operating environments. For example, you can 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, you can 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, you can 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 DCS3700 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 that are supported by the DCS3700 storage subsystem.

### **Zone types**

A zone member can be specified by using one of the following zone types:

- |                        |   |
|------------------------|---|
| <b>Port Level Zone</b> | A type of zone containing members that are specified by switch ports (domain ID or port number) only. Port level zoning is enforced by hardware in the switch.                        |
| <b>WWPN Zone</b>       | A type of zone containing members that are specified by device worldwide port name (WWPN) only. WWPN zones are hardware enforced in the switch.                                       |
| <b>Mixed Zone</b>      | A zone containing various members that are specified by WWPN and certain members that are 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 information about storage area networks, see the following publications:

- ▶ *Introduction to Storage Area Networks and System Networking*, SG24-5470
- ▶ *IBM SAN Survival Guide*, SG24-6143

## **Zoning configuration**

Zoning is configured by 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 help diagnose zoning problems.

When configuring worldwide name (WWN) based zoning, it is important to always use the worldwide port name (WWPN), not the worldwide node name (WWNN). With many systems, the WWNN is based on the Port WWN of the first adapter that is 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. Then, the host with the failing adapter then completely loses access to the storage attached to that switch.

You must update the zoning information if you ever need to replace a Fibre Channel adapter in one of your servers. Most storage subsystems such as the DCS3700 storage subsystem and IBM Tape Libraries have a WWN that is tied to the Vital Product Data of the system unit, so individual parts can usually be replaced with no effect on zoning.

For more information about configuring zoning with your particular switch, see the following publications:

- ▶ *Implementing an IBM b-type 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 DCS3700 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 DCS3700 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

This section shows three sample configurations with two host servers connecting to a dual controller DCS3700 storage subsystem through one or more SAN fabrics.

#### Dual SAN fabric configuration

Figure 3-26 shows a configuration with two hosts connecting through two SAN switches to a dual controller DCS3700 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 DCS3700. A zoning configuration where each host bus adapter has a single connection to only one controller in a dual controller DCS3700 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 3-26 Example of a dual Fibre Channel SAN fabric configuration

### **SAN fabric with Inter-Switch Links**

A similar configuration to the configuration shown in Figure 3-26 on page 123 but using an *Inter-Switch Link (ISL)* between the two switches is shown in Figure 3-27.

In this example, the two SAN switches form a single SAN fabric because of the ISL connection between them.




Figure 3-27 Example of a single Fibre Channel SAN fabric configuration

### **Cross-connected dual SAN fabric**

Figure 3-28 on page 125 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 DCS3700 storage subsystem. Certain applications and host operating systems require a cross-connected SAN fabric.




Figure 3-28 Cross-connected SAN fabric connection

### 3.8.4 iSCSI attachment

The DCS3700 supports iSCSI SAN, which uses the standard Ethernet infrastructure, using standard Cat 5E and Cat 6 Ethernet cables on the host side. The DCS3700 supports a maximum of 64 hosts connecting through a gigabit Ethernet network.

**Attention:** You must always connect the DCS3700 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

This section shows examples of the iSCSI host attachment.

Figure 3-29 shows the simplest iSCSI host connection with one host connecting to a single controller DS3500 Storage System 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 3-29 One host with a single iSCSI connection to a single controller DSC3700

Figure 3-30 shows a configuration with redundant iSCSI connections to a dual controller DCS3700 storage subsystem. In this example, the gigabit Ethernet switch represents a single point of failure. You can achieve more 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 3-31 on page 128.



*Figure 3-30 Single host with redundant iSCSI connections to a dual controller DCS3700*

Figure 3-31 shows a configuration with two hosts connecting to a dual controller DCS3700 with redundant iSCSI connections to two gigabit Ethernet switches, and redundant iSCSI connections between the switch and the DCS3700 controllers. There is no single point of failure in this example configuration.

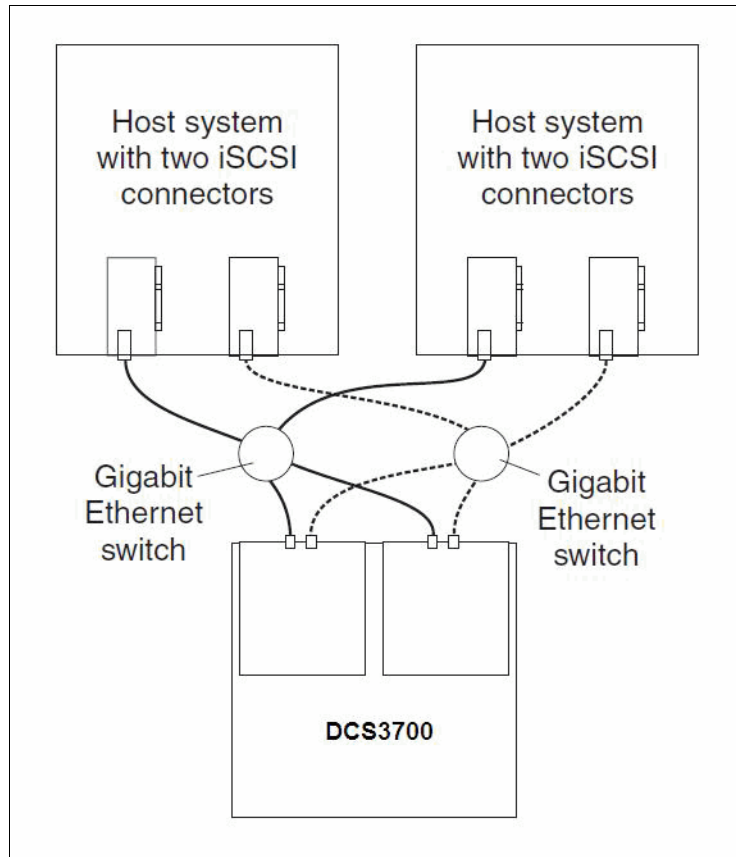


Figure 3-31 Two hosts with redundant iSCSI host connections to a dual controller DCS3700

## 3.9 Power cabling

This section describes the power cabling and the correct power on/off sequence for the DCS3700 storage subsystem.

### 3.9.1 The DCS3700 power supply

Each DCS3700 and DCS3700 expansion unit comes with two 1755 W power supplies as standard.

**Note:** The DCS3700 requires 200 - 240 V AC power. Ensure that it is connected only to the correct AC supply voltage. If there are multiple DCS3700 units to be connected, ensure that the AC source and the rack Power Distribution Unit current ratings are sufficient to support the DCS3700 operating and surge current requirements

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

Connecting the power cables of DCS3700 storage subsystems and DCS3700 expansion units 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, see the appropriate *System Storage DCS3700 storage subsystem: Installation, User's, and Maintenance Guide* manual, which can be found at the following URL.

<http://www.ibm.com/support/entry/portal/docdisplay?lnocid=MIGR-5087760>

### 3.9.2 Powering on and off

A typical DCS3700 configuration includes the following components:

- ▶ Host servers
- ▶ DCS3700 storage subsystem
- ▶ DCS3700 expansion units
- ▶ (Optional) SAN infrastructure

#### Powering on

Follow the correct power-on sequence, or the system might not work correctly. Power the components on in this order:

1. SAN switches should be powered on first. In a typical SAN environment, the SAN switches are running all the time, and are typically only powered down in case of an emergency.
2. DCS3700 expansion units are next. Turn on the switches on both power supplies on each DCS3700 expansion unit. Allow enough time for the enclosures to fully power up, which can take several minutes.
3. When all DCS3700 expansion units are up and running, power on the DCS3700. Large configurations can take up to ten minutes to fully power up.
4. Finally, power on all the host servers when the DCS3700 is fully online. You should check the status of the DCS3700 by using Storage Manager from your out-of-band management station.

**Note:** A DCS3700 system that powers up with no drives installed gets a dummy Feature Enable Identifier (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 readded or powered up and the controllers reboot to load information out of the DACSTORE on the drives.

#### Powering off

In normal circumstances, the DCS3700 storage subsystem rarely is 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 must power the subsystem off. Use the same steps in “Powering on”, but in the reverse order. If you cannot shut down and power off the host servers, then make sure the DCS3700 logical drives are unmounted at the operating system level.

**Important:** Unless it is an emergency, the DCS3700 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.





# IBM System Storage DS Storage Manager installation

This chapter describes how to install the IBM System Storage DS Storage Manager on Microsoft Windows, Linux, and AIX. This chapter is organized in to the following sections:

- ▶ Installing DS Storage Manager on Microsoft Windows 2008
- ▶ Installing DS Storage Manager on Linux
  - Installing Storage Manager by using the GUI
  - Installing DS Storage Manager by using a text console
- ▶ Installing DS Storage Manager on AIX

The sections describe installing DS 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 information about Storage Manager components and their usage, see 1.6, “IBM System Storage DS Storage Manager” on page 14.

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

## 4.1 Installing DS Storage Manager on Microsoft Windows 2008

This section describes how to install DS Storage Manager on a Windows 2008 R2 system. Depending on the usage of the system, different components need to be installed, as described in 1.6, “IBM System Storage DS Storage Manager” on page 14.

### 4.1.1 Installation preparation

Before you perform the installation, complete the following steps:

1. Download the Storage Manager installation package from the IBM Support website. There are 32-bit (x86) and 64-bit (x86-64) versions. Make sure to download the package that corresponds to the Windows version of the host or management station. For more information about how to download DCS3700 fixes, code, and documentation, see Appendix A, “IBM Support Portal website” on page 583.
2. Extract the compressed archive into a local directory. The compressed archive contains the following items:
  - 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
  - A scripts subdirectory that contains different scripts, such as Disable AVT
3. Decide on the installation type and proceed with the installation.

These steps are applicable for Manager and Client options.

### 4.1.2 Installing the Storage Manager Client on Microsoft Windows 2008

To install the Storage Manager client on a Microsoft Windows 2008 server, complete the following steps:

1. Run the installer file. Our version is SMIA-WS64-10.83.35.18.exe. The files extract themselves as shown in Figure 4-1.

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




Figure 4-1 IBM DS Storage Manager InstallAnywhere

2. Select the language to display the license agreement and click **OK** (see Figure 4-2).

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



Figure 4-2 IBM DS Storage Manager Language Selection

3. Read the introduction and click **Next**, as shown in Figure 4-3.

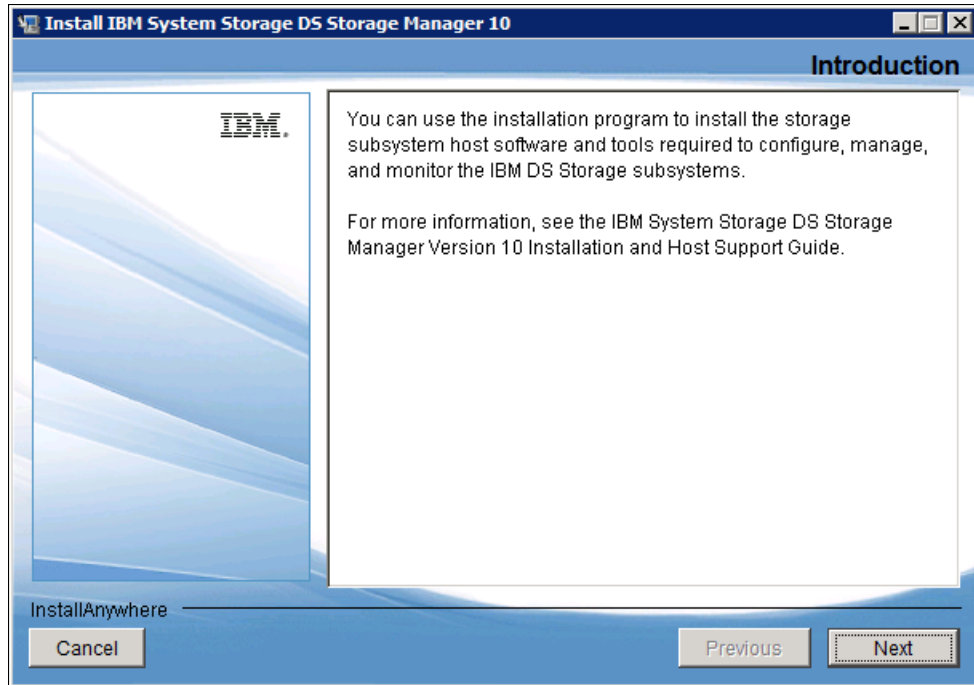


Figure 4-3 IBM DS Storage Manager Introduction

4. Read the copyright statement and click **Next**, as shown in Figure 4-4.

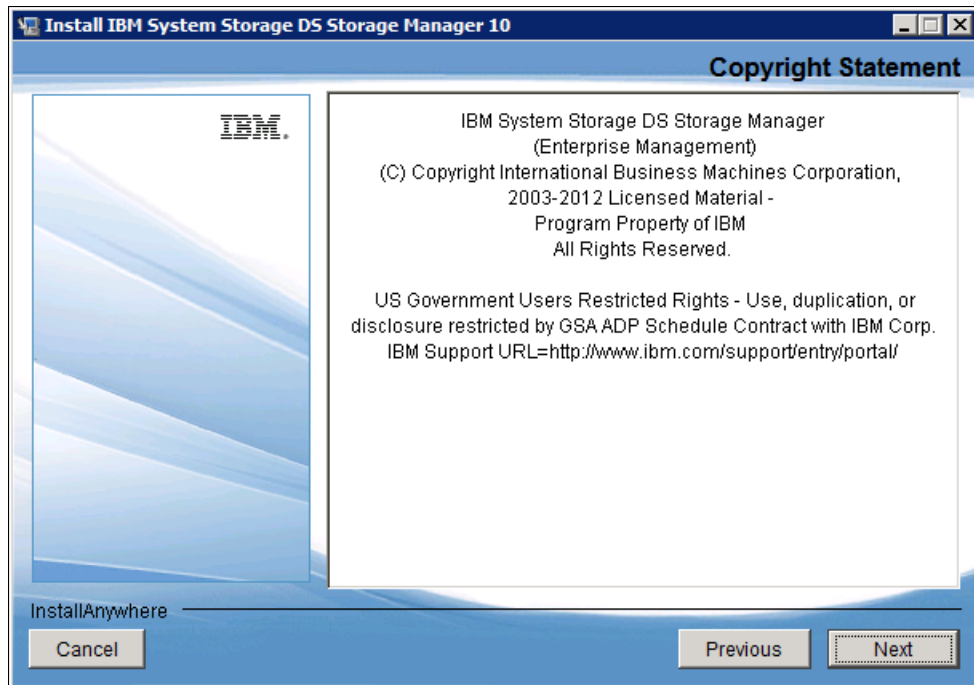


Figure 4-4 IBM DS Storage Manager Copyright Statement

5. Read the license agreement and click **Next**, as shown in Figure 4-5.

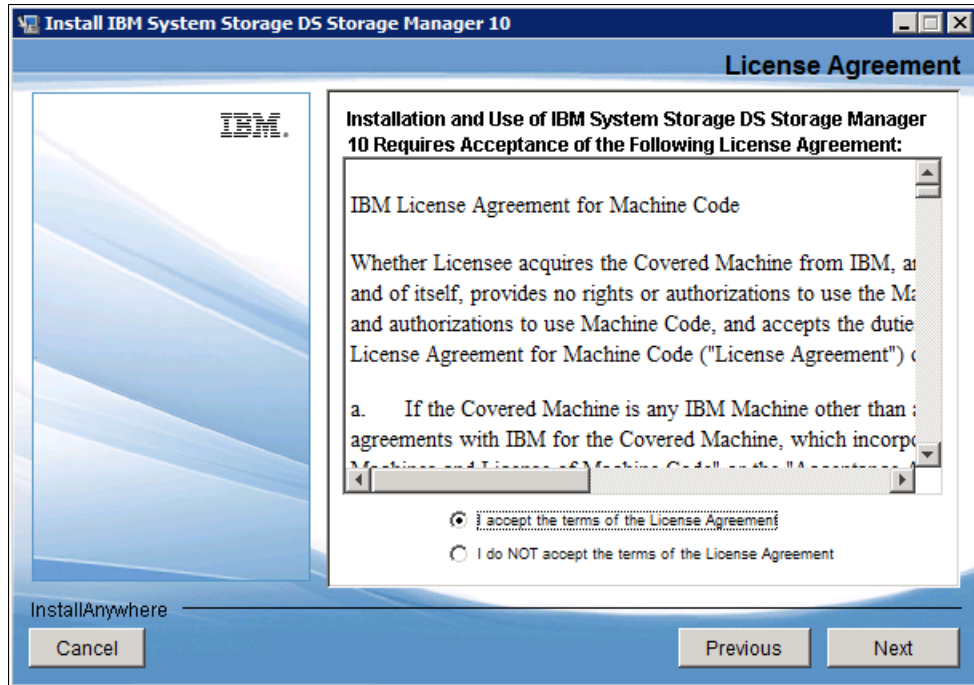


Figure 4-5 IBM DS Storage Manager License Agreement

6. Specify the directory where you want to install, as shown in Figure 4-6. The default is C:\Program Files (x86)\IBM\_DS.




Figure 4-6 IBM DS Storage Manager Choose Install Folder

7. For Installation Type (Figure 4-7), 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.

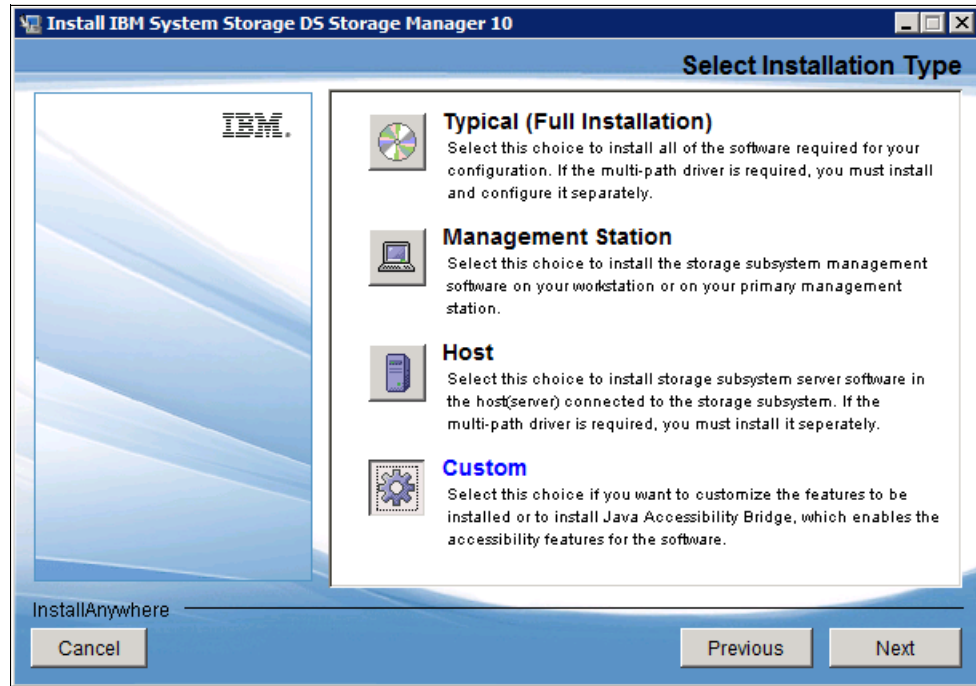


Figure 4-7 IBM DS Storage Manager Select Installation Type

8. In Figure 4-8 on page 137, select **Install Set** from the drop-down menu. The options in the drop-down menu are same as shown in Figure 4-7. When you select an Install Set, you can see what components will be installed and modify them if you want. For a description of each component, see 1.6, "IBM System Storage DS Storage Manager" on page 14. Our example shows the client selection because the system being installed is for management purposes and the other components are not required for it.

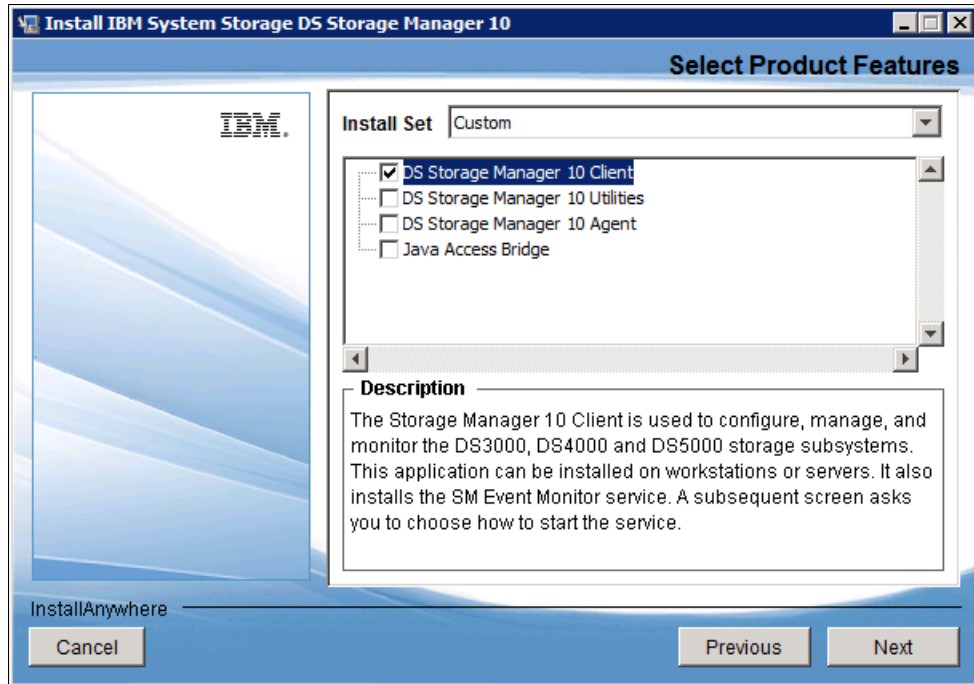


Figure 4-8 Select Product Features

9. If a version of Storage Manager is already installed, you see the warning that is shown in Figure 4-9. You can select to overwrite the existing installation or to cancel. Click **OK** to continue. If you do not have Storage Manager already installed, you do not see this warning.

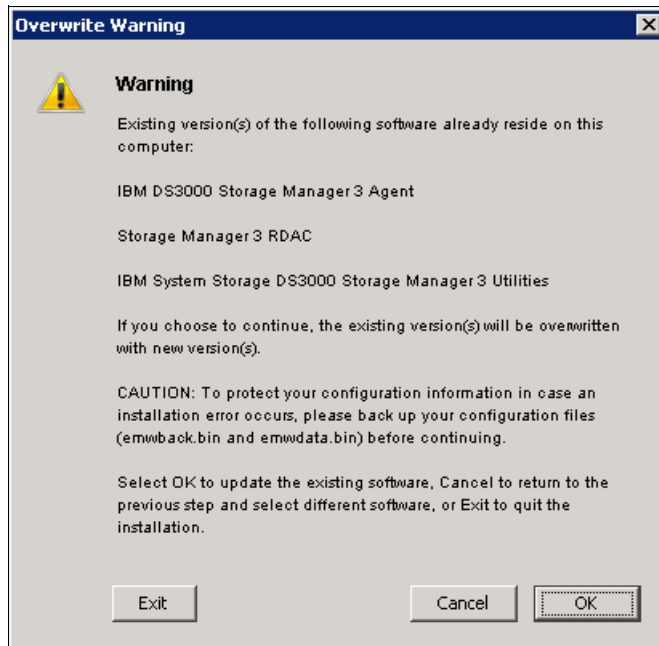


Figure 4-9 Overwrite Warning

10. In Figure 4-10, you can select whether to automatically start the Storage Manager Monitor. Click **Next**.

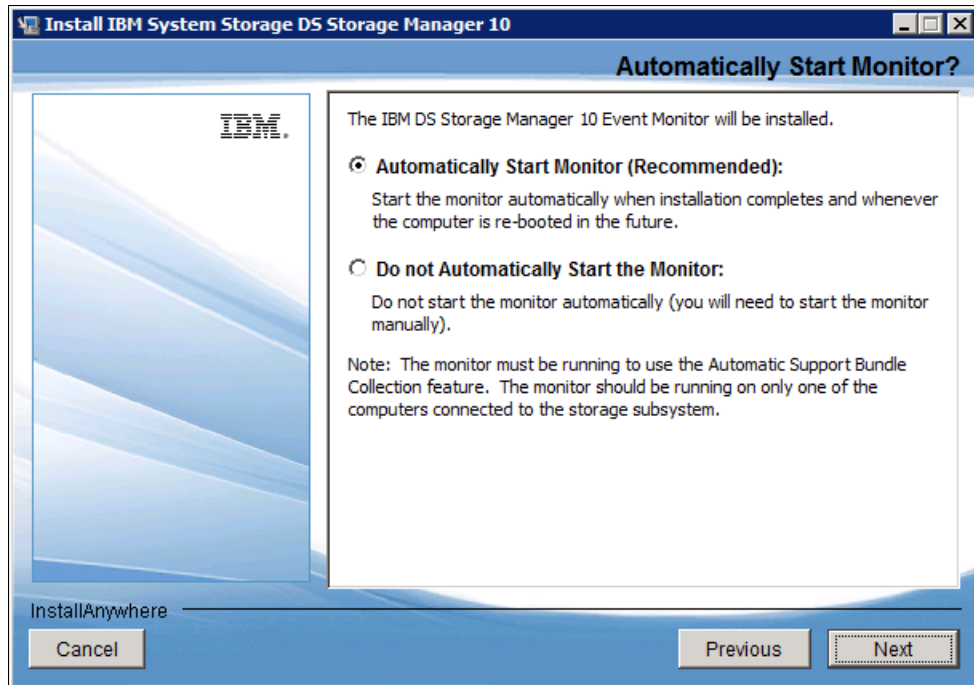


Figure 4-10 IBM DS Storage Manager Automatically Start Monitor

11. Verify the preinstallation summary and click **Install** (Figure 4-11).



Figure 4-11 IBM DS Storage Manager Pre-Installation Summary



12. The installation status window opens during the installation, as shown in Figure 4-12.

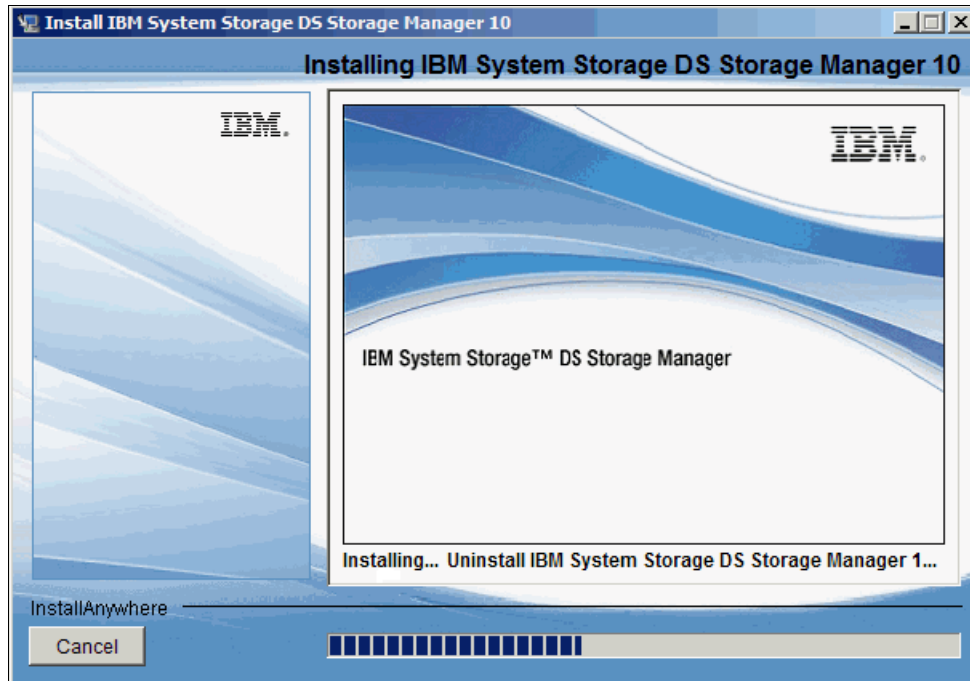


Figure 4-12 IBM DS Storage Manager installation status window

After the installation finishes, verify that there are no errors and click **Done** (Figure 4-13).

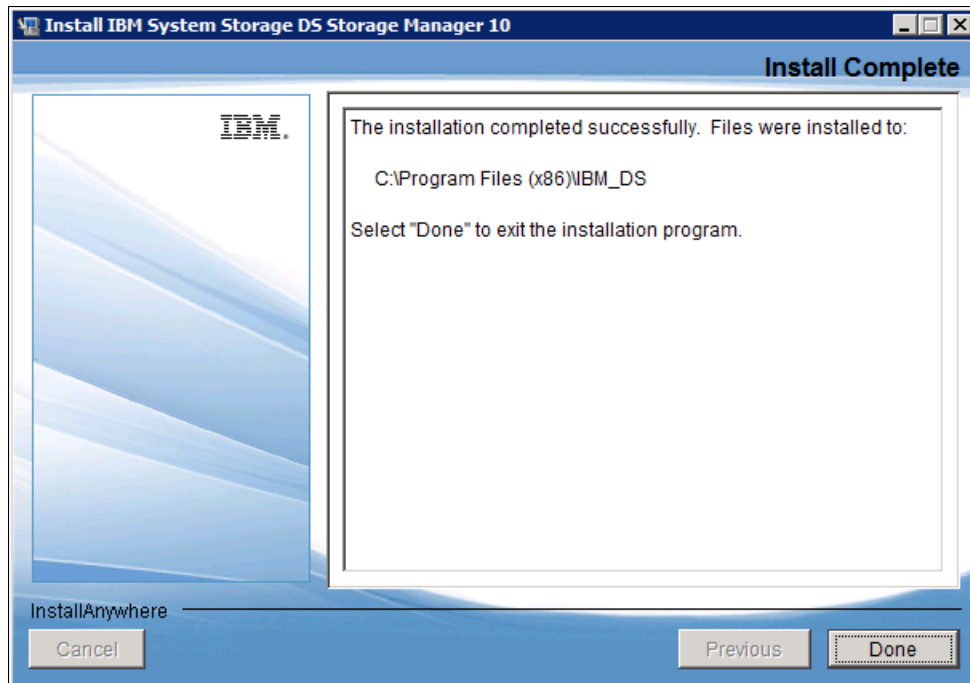


Figure 4-13 IBM DS Storage Manager Install Complete

## 4.2 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. For more information about installing the multipath drivers, see 18.2.3, “Installing the Linux multipath driver” on page 542.

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

- ▶ Preparing for the installation
- ▶ Installing Storage Manager by using the GUI
- ▶ Installing DS Storage Manager by using a text console

After you complete these steps, you can start IBM Storage Manager by running **SMclient** from a terminal window.

### 4.2.1 Preparing for the installation

To prepare for the installation, complete the following steps:

1. Download the latest version of the IBM Systems Storage DS Storage Manager Client to the management station or host that is connected to the DCS3700 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 583.
2. Verify that there is enough space available in the `/opt` (450 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 changes *only* the temporary location where the IBM Storage Manager extracts the files for the installation.

To set the variable, run the following commands at the Linux command-line interface (CLI) before running this installer again:

- For Bourne shell (sh), ksh, bash, and zsh, run the following commands:

```
$ IATEMPDIR=/your/free/space/directory
$ export IATEMPDIR
```

- For C shell (csh) and tcsh, run the following commands:

```
$ setenv IATEMPDIR /your/free/space/directory
```

3. Extract the file to a local directory, as shown in Example 4-1.

*Example 4-1 Extract the package archive*

---

```
# tar -zxvf SM10.83_Linux_64bit_x86-64_SMIA-10.83.x5.18.tgz
Linux_x64_10p83/Linux_x86-64/
Linux_x64_10p83/Linux_x86-64/SMIA-LINUX64-10.83.A5.18.bin
Linux_x64_10p83/SM10R2MIB/
Linux_x64_10p83/SM10R2MIB/SM10_R2.MIB
Linux_x64_10p83/SM10R2MIB/README_First.txt
Linux_x64_10p83/scripts
```

---

The following files are unpacked:

- The installation package and executable
- The readme file that describes the package
- A subdirectory that is named SM10R2MIB that contains the MIB files
- A subdirectory that is named scripts that contains different scripts, such as Disable AVT

**Note:** This MIB file work only with the IBM System Storage DS3500, DCS3700, DS4000, and DS5000 Storage Manager.

4. The SMIA-LINUX-10.83.A5.18.bin file is the installer that contains the installation files. To be able to run this installer, you need to type `sh` before the file or change the file permissions to be able to execute.
5. By default, the installer requires a graphical console to start, but can be also installed on a text console or left unattended.

Run this file with the `-?` option to display information about the available options that can be used to install Storage Manager.

Here are the options:

- **swing:** A GUI installation that allows you to select the language of the license agreement. 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.
6. Choose which method you want to use, and start the installation file, specifying your option with the `-i` parameter. For example:

```
$sh SMIA-LINUX-10.83.A5.18.bin -i console
```

## 4.2.2 Installing Storage Manager by using the GUI

**Note:** It is possible to run the GUI installation remotely, where you export the display by using X11 forwarding.

To install Storage Manager, complete 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-10.83.A5.18.bin
```

3. A window opens. Select the language to display the license agreement and click **OK**, as shown in Figure 4-14.

**Note:** The installation is performed in English always.

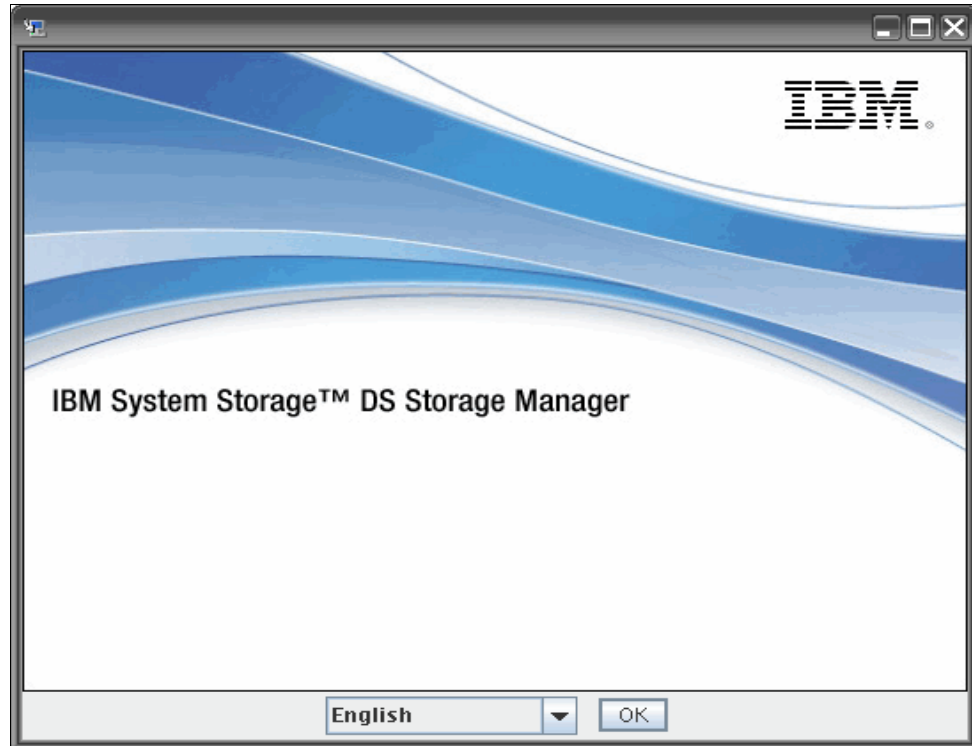


Figure 4-14 Select the language to display license agreement

4. Read the introduction and click **Next**, as shown in Figure 4-15.




Figure 4-15 Introduction

5. Read the copyright statement and click **Next**, as shown in Figure 4-16.

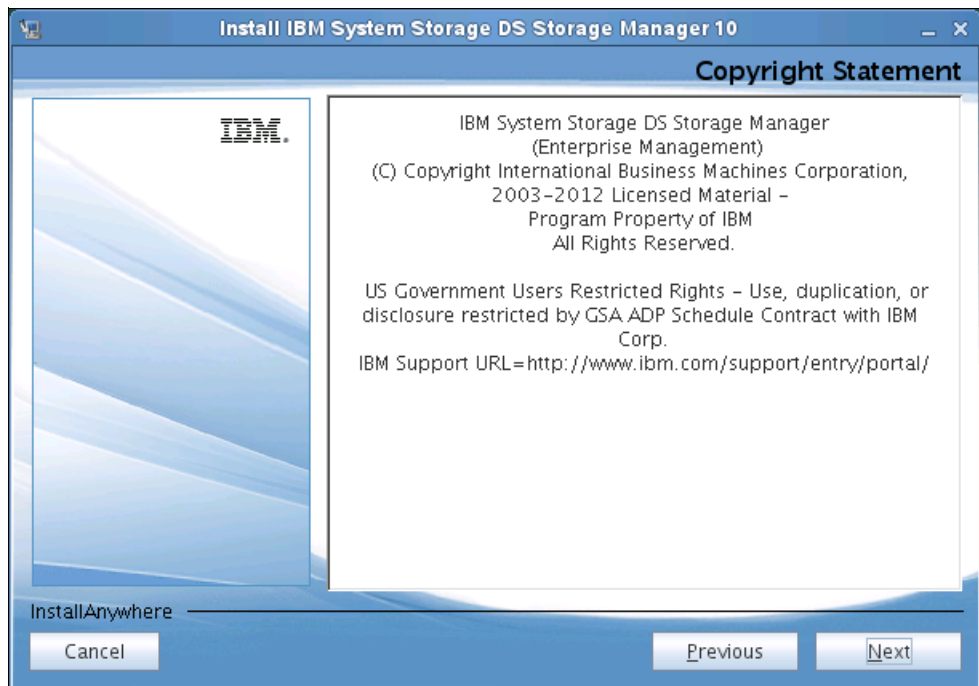


Figure 4-16 Copyright statement

6. Read and accept the license agreement. Click **Next**, as shown in Figure 4-17.

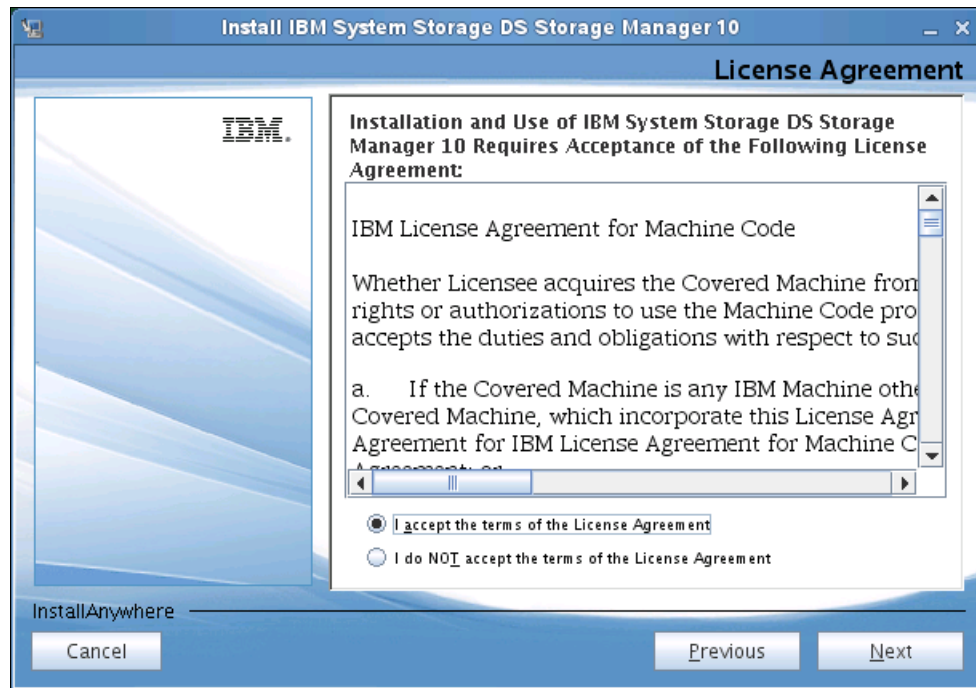


Figure 4-17 License agreement

7. Select the **Custom** installation type and click **Next**, as shown in Figure 4-18. 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.

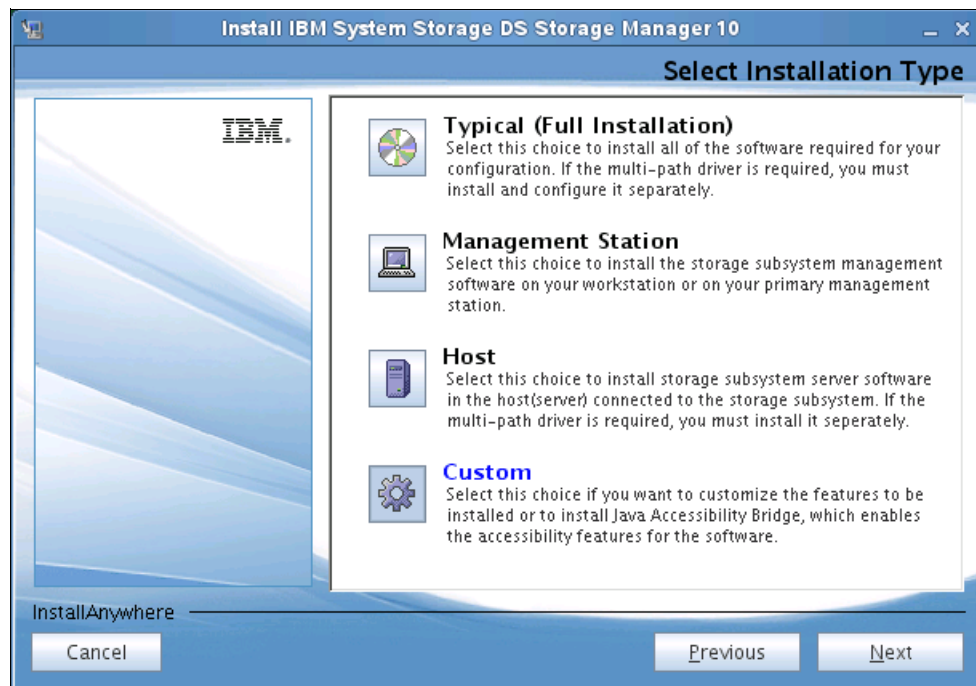


Figure 4-18 Installation type

8. Select the Install Set to be installed (DS Storage Manager 10 Client in our case). Click **Next**, as shown in Figure 4-19.

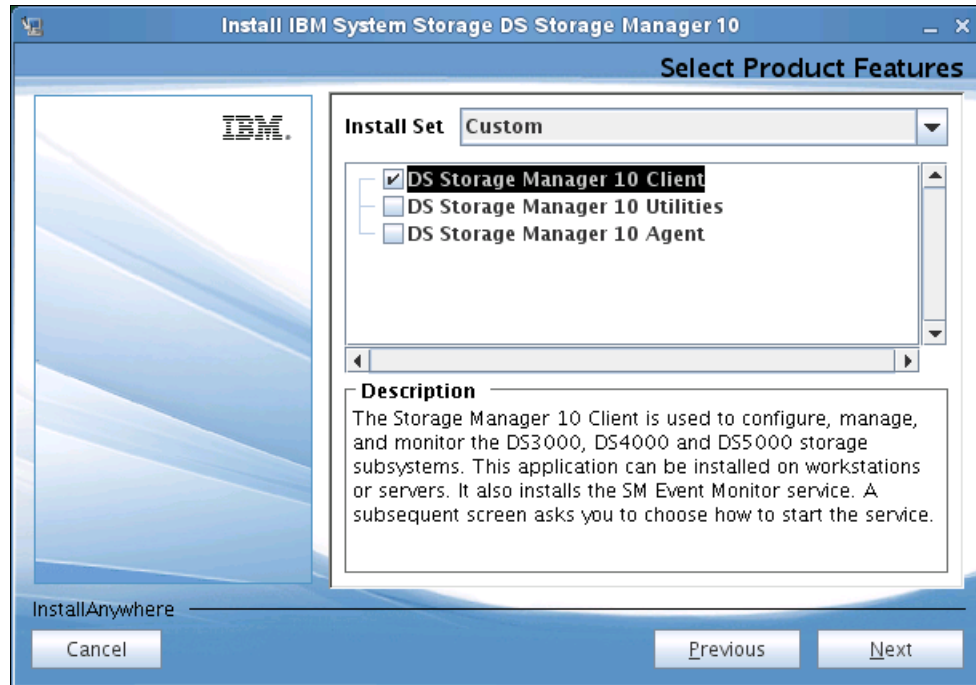


Figure 4-19 Select Product Features

9. Verify the installation options and click **Next**, as shown in Figure 4-20.



Figure 4-20 Pre-Installation Summary

10. After a successful installation, a window similar to Figure 4-21 opens. If there are no errors, click **Done** to exit.

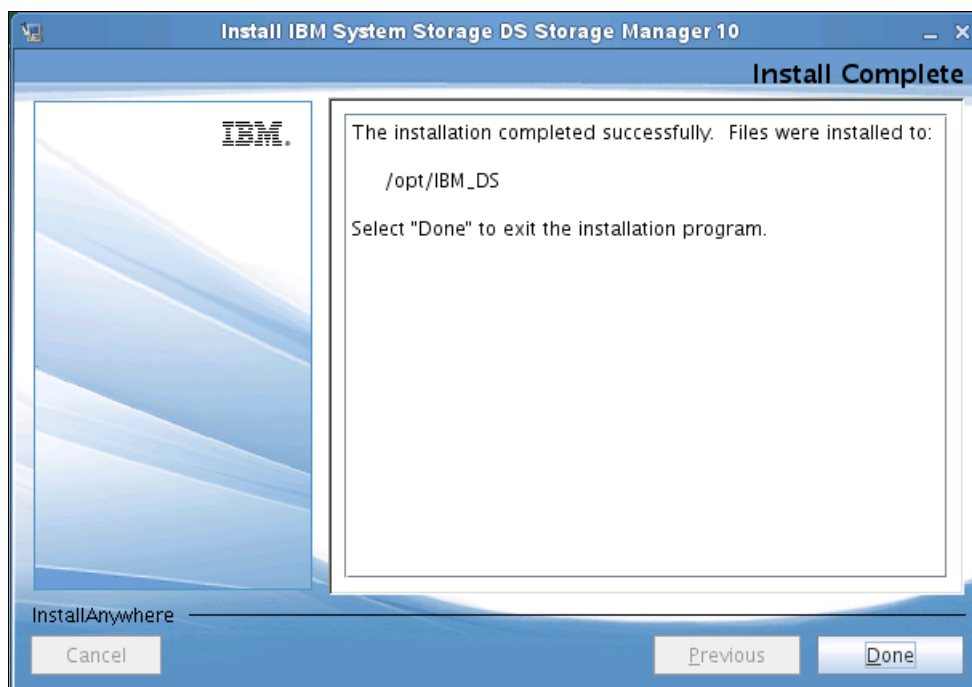


Figure 4-21 Install Complete

11. Optionally, verify the installation by running the `rpm -qa |grep SM` command and look for the following RPM packages in the installed package list:
- SMclient-10.83.G5.18-1
  - SMesm-10.83.G5.02-1
  - SMruntime-10.83.A5.00-1

### 4.2.3 Installing DS Storage Manager by 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 the DS Storage Manager by using a text console, complete the following steps:

1. Start the installer with the `-i console` option. Select the locale (the default is English), and press Enter, as shown in Example 4-2.

*Example 4-2 Text-based installation*

```
# sh SMIA-LINUX-10.83.A5.18.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- Bahasa Indonesia
- 2- Deutsch
- >3- English
- 4- Español
- 5- Francais
- 6- Italiano
- 7- Português (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 4-3. A full installation installs the Storage Manager Client, Utilities, and Agent. The option Management Station installs the Storage Manager Client and Utilities, and the option Host installs only the Storage Manager Agent. Press 4 for Custom and press Enter.

*Example 4-3 Installation type*

---

Select Installation Type  
-----

Please 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, you can select or clear individual components. By default, all components are selected and you can clear a component by entering the number that is associated with the component that you do not want to install. Here we choose only the DS Storage Manager 10 Client, as shown in Example 4-4.

*Example 4-4 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- [X] DS Storage Manager 10 Client
- 2- [ ] DS Storage Manager 10 Utilities
- 3- [ ] 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 4-5.

*Example 4-5 Storage Manager - Pre-Installation Summary*

---

Pre-Installation Summary

-----

Review the Following Before Continuing:

Install Folder:

/opt/IBM\_DS

Product Components:

DS Storage Manager 10 Client,  
DS Storage Manager 10 Utilities

Required Disk Space

714 MB

Available Disk Space

8,735 MB

PRESS <ENTER> TO CONTINUE:

---

7. The installation is now performed in the /opt/IBM\_DS directory. After a successful installation, a message similar to Example 4-6 is displayed.

*Example 4-6 Installation Complete*

---

Installation Complete

-----

Congratulations. IBM System Storage DS Storage Manger 10 has been successfully installed to:

/opt/IBM\_DS

PRESS <ENTER> TO EXIT THE INSTALLER:

#

---

## 4.2.4 Uninstalling DS Storage Manager on Linux

To uninstall DS Storage Manager on Linux, complete the following steps:

1. Run the following command at the command prompt:

```
# /opt/IBM_DS /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 the current directory by running the following command:

```
rpm -qil SMclient-10.83.G5.18-1.noarch
```

At the time of writing, that is the name of the package. The current name of the RPM console package should display after you run the following command:

```
rpm -qa |grep SM
```

2. Read the introduction and press Enter to continue, as shown in Example 4-7.

*Example 4-7 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 4-8.

*Example 4-8 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 4-9 is displayed.

*Example 4-9 Complete uninstallation*

```
=====
Uninstall Complete
-----
All items were successfully uninstalled.

linux:/opt/IBM_DS/Uninstall IBM System Storage DS Storage Manager 10 #
```

## 4.3 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 are not included in the installation package; it must be installed separately from the Storage Manager package. For more information about the AIX multipath drivers, see 19.2, “Configuring MPIO” on page 558.

As on Linux, the IBM Storage Manager can be installed by using a GUI or a text-based interface. Because we described the graphic and console mode in 4.2.2, “Installing Storage Manager by using the GUI” on page 141 and in 4.2.3, “Installing DS Storage Manager by using a text console” on page 146, we are going to include only the preparation for this installation.

After complete these steps, you will have installed the packages detailed below and be able to start IBM Storage Manager by running **SMclient** from a terminal window.

The following packages are shown by running the **ls1pp -a1 |grep SM** command:

SMclient.aix.rte	10.83.G5.18.0	COMMITTED	IBM DS Storage Manager 10
SMesm.aix.rte	10.83.G5.02.0	COMMITTED	IBM DS Storage Manager ESM
SMruntime.aix.rte	10.83.65.00.0	COMMITTED	IBM DS Storage Manager 10
SMutil.aix.rte	10.0.65.21.0	COMMITTED	IBM DS Storage Manager 10

### 4.3.1 Preparing for the installation

To prepare for the installation, complete the following steps:

1. Download the latest version of the IBM Systems Storage DS Storage Manager Client to the management station or host that is connected to the DCS3700 for inband or outband connections from the IBM Support website. For more information about the IBM Support website, see Appendix A, “IBM Support Portal website” on page 583.
2. Verify that 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 to a disk partition with enough free disk space instead of using `/tmp`.

**Note:** The setting of this variable changes *only* the temporary location where the IBM Storage Manager places the files for the installation.

To set the variable, enter one of the following commands at the UNIX CLI before running this installer again:

- For Bourne shell (sh), ksh, bash, and zsh, run the following commands:

```
$ IATEMPDIR=/your/free/space/directory
$ export IATEMPDIR
```

- For C shell (csh) and tcsh, run the following commands:

```
$ setenv IATEMPDIR /your/free/space/directory
```

3. Extract the file to a local directory, as shown in Example 4-10.

*Example 4-10 Extract package archive*

```
# gzip -d SM10.83_Aix_SMIA-10.83.x5.18.tgz
# tar -xvf SM10.83_Aix_SMIA-10.83.x5.18.tar
```

```
x Aix/Aix
x Aix/Aix/SMIA-AIX-10.83.65.18.bin, 707704559 bytes, 1382236 tape blocks
x Aix/SM10R2MIB
x Aix/SM10R2MIB/SM10_R2.MIB, 5298 bytes, 11 tape blocks
x Aix/SM10R2MIB/README_First.txt, 126 bytes, 1 tape blocks
```

---

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
- A subdirectory scripts that contains different scripts, such as Disable AVT

**Note:** This MIB file works only with the DS3500, DCS3700, DS4000, and DS5000 Storage Manager.

4. The SMIA-AIX-10.83.65.18.bin file is the installer that contains the installation files. To be able to run this installer, you need to type `sh` before the file or change the file permissions to be able to execute.
5. By default, the installer requires a graphical console to start, but it can also be run on a text console or left unattended.

Run this file with the option `-?` to display information about the available options to install Storage Manager.

Here are the options:

- **swing:** A GUI installation that allows you to select the language of the license agreement. 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.
6. Choose which method you want to use, and start the installation file, specifying your option with the `-i` parameter. For example:  

```
$sh SMIA-AIX-10.83.65.18.bin -i console
```

## 4.4 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, you must complete the following tasks:

- ▶ 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 perform the following tasks:

- ▶ 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 5, “Administration: Enterprise Management” on page 159.

## 4.4.1 Performing an automatic discovery of storage subsystems

To perform an initial automatic discovery of storage subsystems, complete the following steps:

1. From the Setup tab of the DS Storage Manager Enterprise Management window, click **Add Storage Subsystem**, as shown in Figure 4-22.

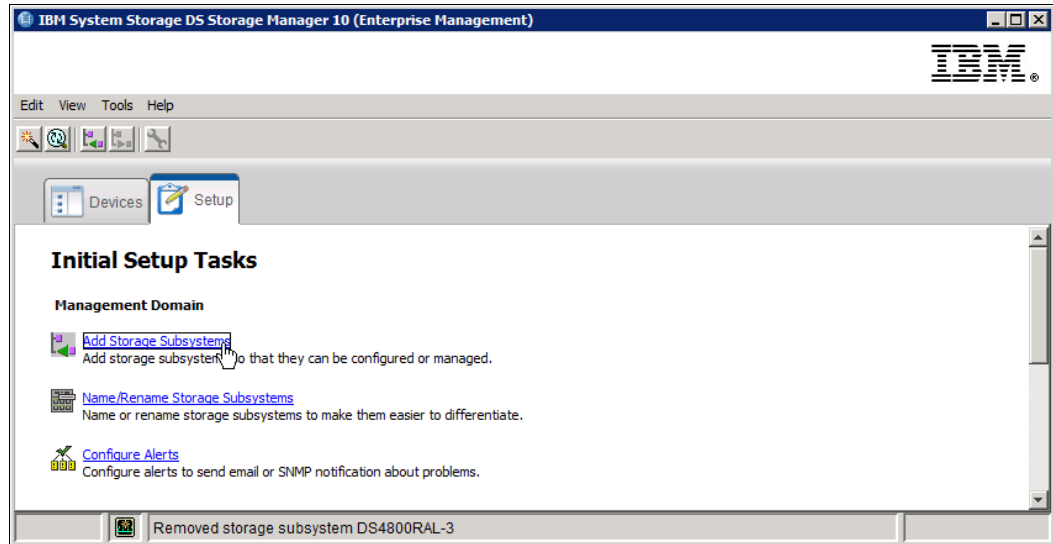


Figure 4-22 Add Storage Subsystem

2. Select **Automatic**, as shown in Figure 4-23, and click **OK**.



Figure 4-23 Select Addition Method - Automatic

3. Click **OK** to begin an initial automatic discovery of hosts and storage subsystems that are attached to the local subnetwork, as shown in Figure 4-24.

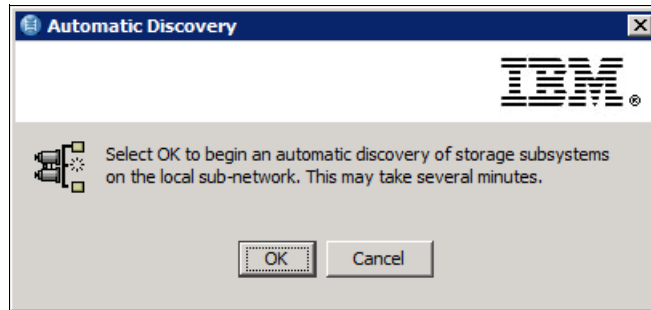


Figure 4-24 Automatic Discovery

After the initial automatic discovery is complete, the Devices tab of the Enterprise Management window displays all hosts and storage subsystems that are attached to the local subnetwork.

**Note:** The Enterprise Management window can take up to a minute to refresh after an initial automatic discovery.

#### 4.4.2 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, complete the following steps:

1. From the Setup tab DS Storage Manager Enterprise Management window, click the **Add Storage Subsystem** link, as shown in Figure 4-25.

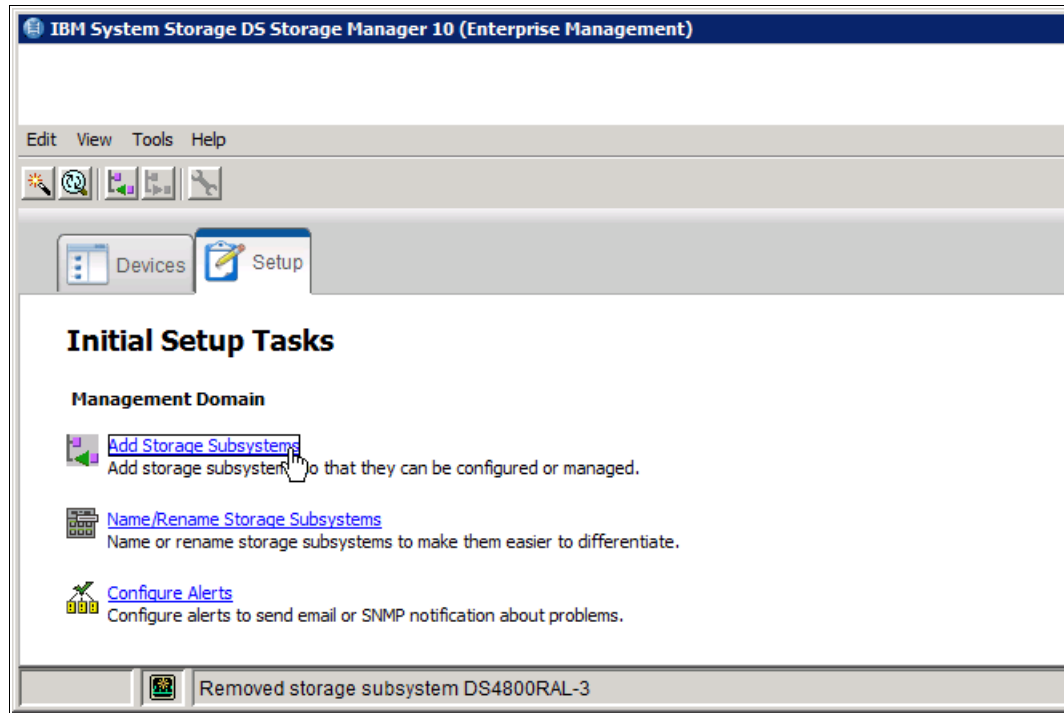


Figure 4-25 Add Storage Subsystem

2. Select **Manual**, as shown in Figure 4-26, and click **OK**.

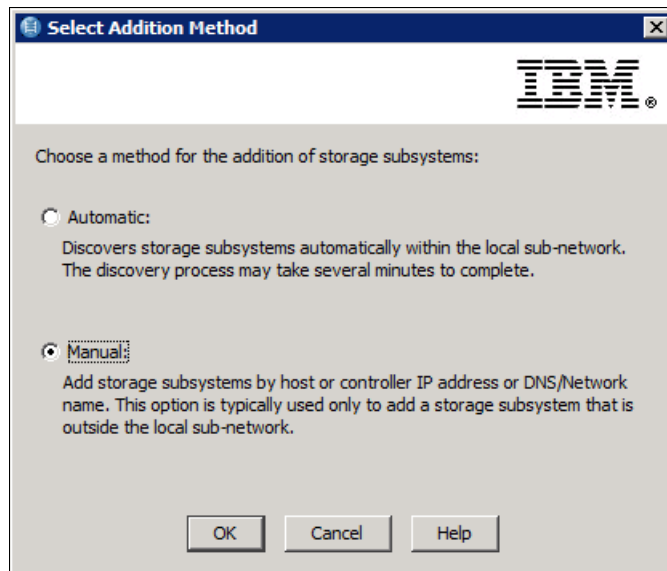


Figure 4-26 Select Addition Method - Manual



3. Select a management method and add the storage subsystem:
  - When using out-of-band management, select **Out-of-band management** and enter the IP addresses of both the controllers, as shown in Figure 4-27.

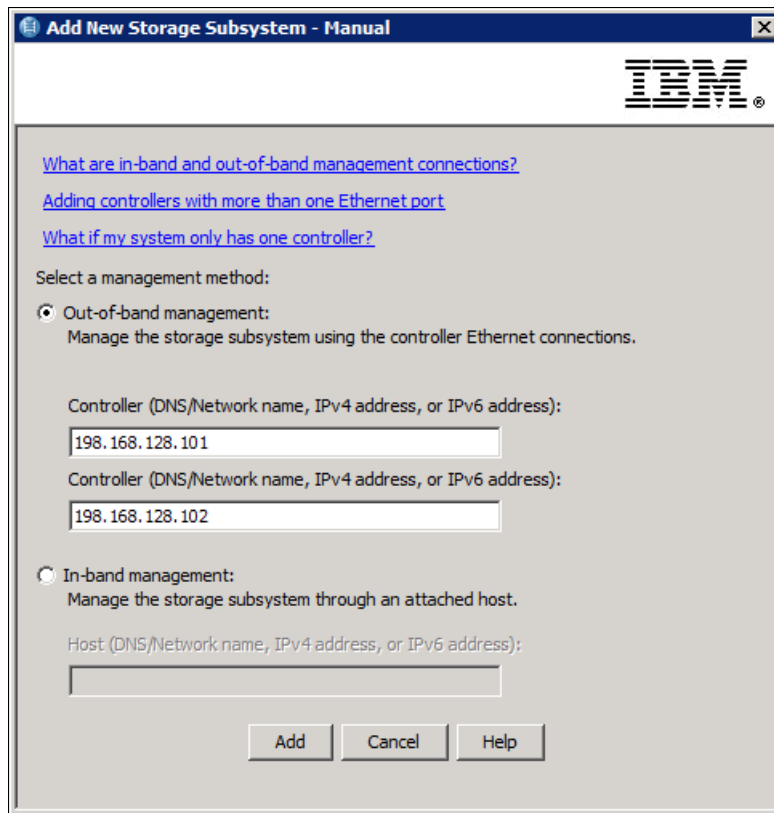


Figure 4-27 Out-of-band management

- For In-band management, select the **In-band management** option and enter the IP address of the attached host, as shown in Figure 4-28.




Figure 4-28 In-band management

4. Click **Add** to complete the process. After the storage subsystem is successfully added, you can add another storage subsystem, as shown in Figure 4-29, and you can repeat the process until all the storage subsystems that you want to add are added.

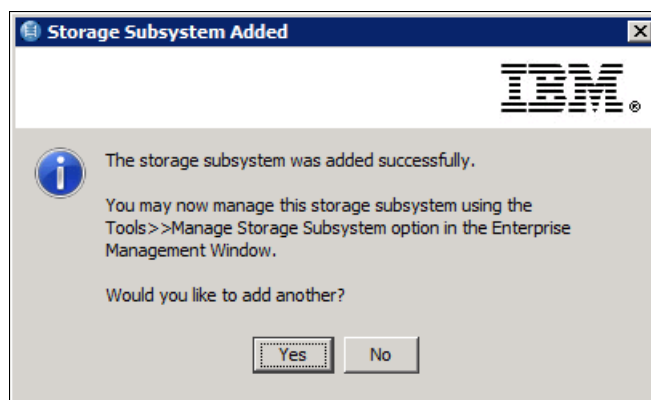


Figure 4-29 Storage Subsystem Added

### 4.4.3 Adding a 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, complete 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 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 system.

- ▶ If you are using the host-agent management method, complete the following steps:
  - a. Make sure that the SMagent is installed on the host.
  - b. Verify that you have a Fibre Channel connection from the storage subsystems to the host that has the SMagent installed.
  - c. Verify that all of the preparation steps are complete, and then complete the following steps:
    - i. Run the hot\_add utility.
    - ii. Restart the SMagent.
    - iii. 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.

You see an Enterprise Management window (Figure 4-30) when you add your systems successfully.

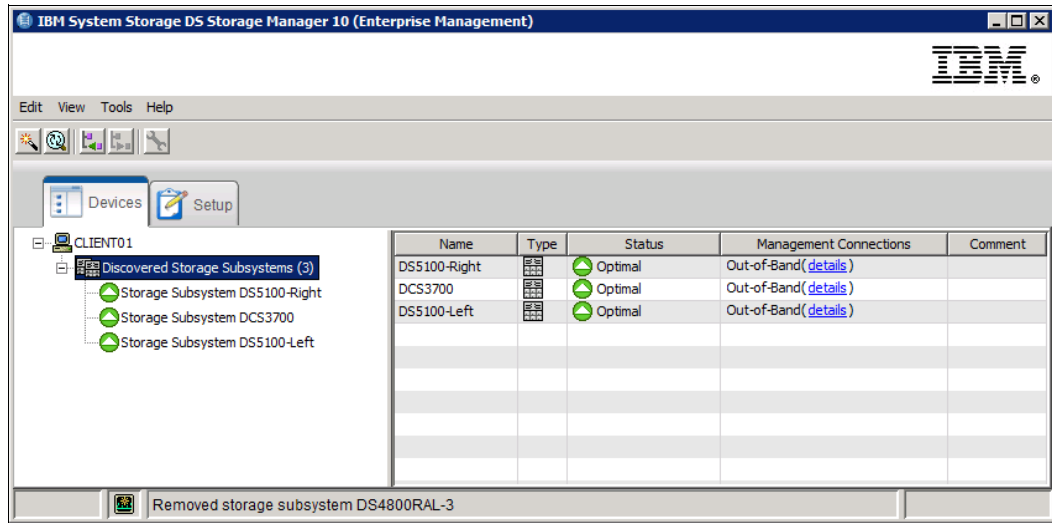


Figure 4-30 Enterprise Management window with discovered systems



# Administration: Enterprise Management

This chapter describes the IBM System Storage DCS3700 Storage Manager Enterprise Management window, including its functions and how to use it.

## 5.1 Enterprise Management window overview

When you start Storage Manager, you get the Enterprise Management window (shown in Figure 5-1), and the Setup tab (shown in Figure 5-2 on page 161) shows the Initial Setup Tasks window.

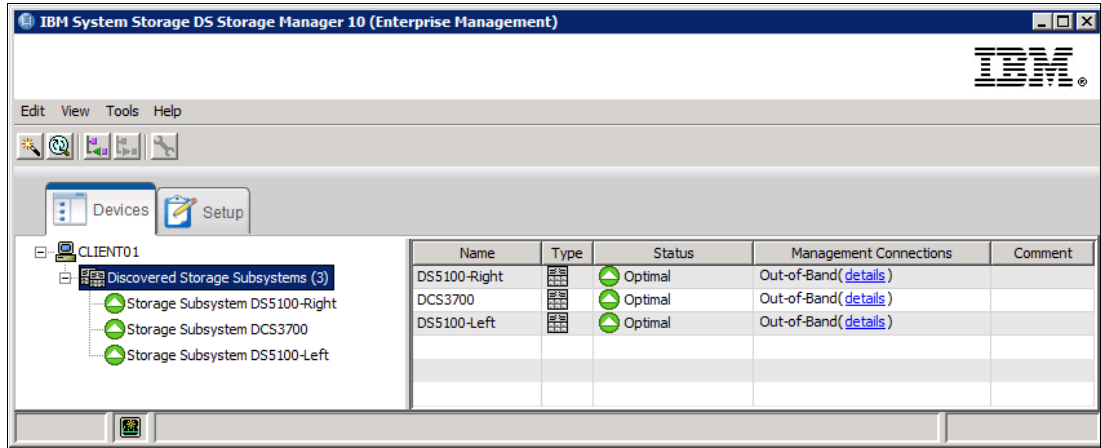


Figure 5-1 Enterprise Management window

### 5.1.1 Initial Setup Tasks

The Initial Setup Tasks window is shown in Figure 5-2 on page 161. The task list here gives you a quick way to perform common tasks.

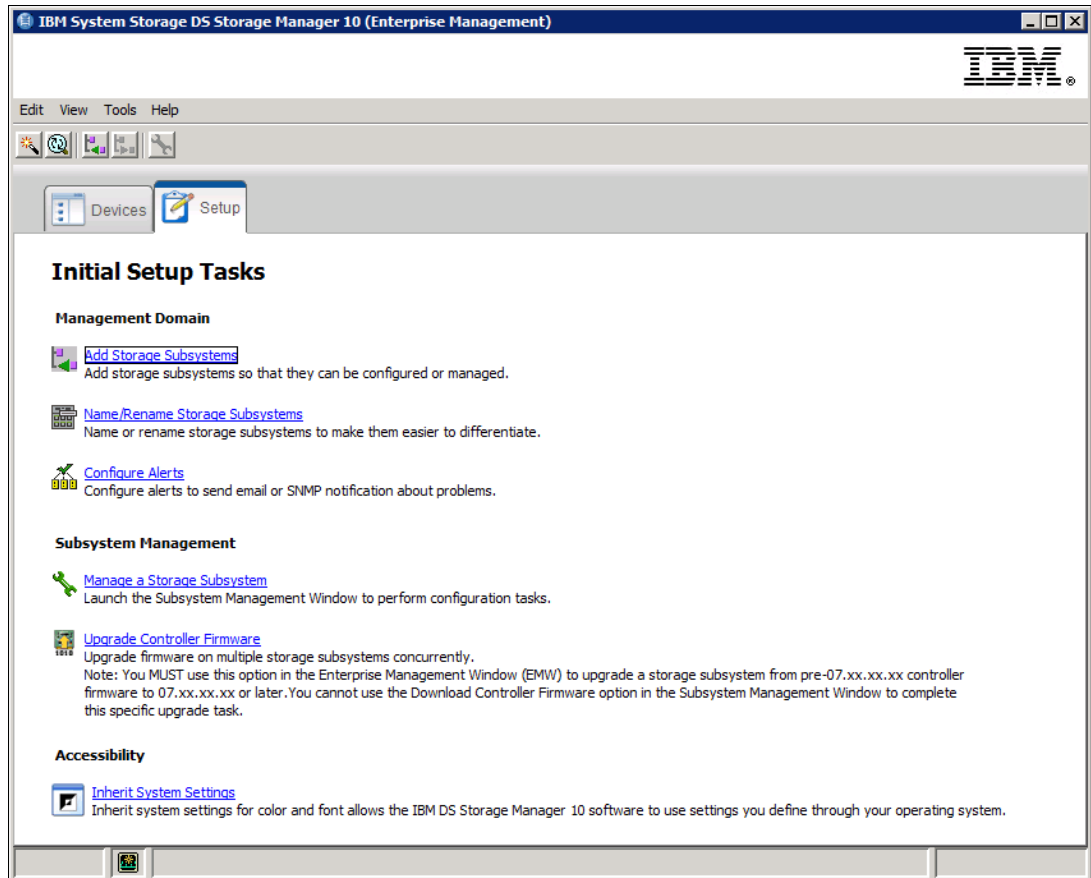


Figure 5-2 Initial Setup Tasks window

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

- ▶ Initial Setup Tasks:
  - Add Storage Subsystems: See 4.2, “Installing DS Storage Manager on Linux” on page 140 or “Add Storage Subsystem to view icon” on page 182.
  - Name/Rename Storage Subsystems: See 11.1.2, “Rename Storage Subsystem” on page 355.
  - Configure Alerts: See “Configure Alerts option” on page 168.
- ▶ Subsystem Management:
  - Manage a Storage Subsystem: Starts the main Storage Manager application, as described in this chapter.
  - Upgrade Controller Firmware: Upgrades firmware on multiple storage systems, as described in Chapter 13, “Administration: Monitor” on page 385.
- ▶ Accessibility:
  - Inherit System Settings: Inherits system settings for color and fonts, which allows the IBM System Storage DS software to use settings that are defined through the operating system.

## 5.1.2 Enterprise Management window

The Enterprise Management window (Figure 5-1 on page 160) is the entry point to manage each DCS3700 storage subsystem. We described how to add storage subsystems in 4.4, “Completing the DS Storage Manager installation” on page 151. After they are added, they appear every time that you start Storage Manager.

The Enterprise Management window displays a list of all DCS3700 and other DS3500, DS4000, and DS5000 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 more than once in the Enterprise Management window.

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

In the left 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 pane, you also see a short summary of this system in the right pane. If the subsystem appears red in this view, then it needs attention. What you do in this case is described in Chapter 13, “Administration: Monitor” on page 385.

Figure 5-3 shows the various status icons for a storage subsystem.








Status	Icon
Optimal	
Needs Attention	
Unresponsive	
Fixing Status	
Unsupported	
Software Unsupported	
Alarm Sounding	

Figure 5-3 Status icons



Figure 5-4 explains the meaning of each status icon.








Status	Icon	Indicates
Optimal		Each 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 your intervention to correct it.
Unresponsive		The storage management station cannot communicate with the storage subsystem or one controller or both controllers in the storage subsystem.
Fixing Status		A Needs Attention condition has been corrected and the managed storage subsystem is currently transitioning to an Optimal state.
Unsupported		The node is currently not supported by this version of the storage management software.
Software Unsupported		The storage subsystem is running a level of software that is no longer supported by the storage management software.
Alarm Sounding		When a controller enclosure alarm is enabled and the audible alarm for the controller enclosure is sounding, the <b>Alarm Sounding</b> icon appears for that storage subsystem. <b>Note:</b> The <b>Alarm Sounding</b> icon appears only in the EMW Tree view and the EMW Table view.

Figure 5-4 Status icons and their meaning

## 5.2 Functions in the Enterprise Management window

This section describes the various functions that are available from the Enterprise Management window.

## 5.2.1 Subsystem menu

Right-click one of the subsystems to see a menu with various options, as shown in Figure 5-5.

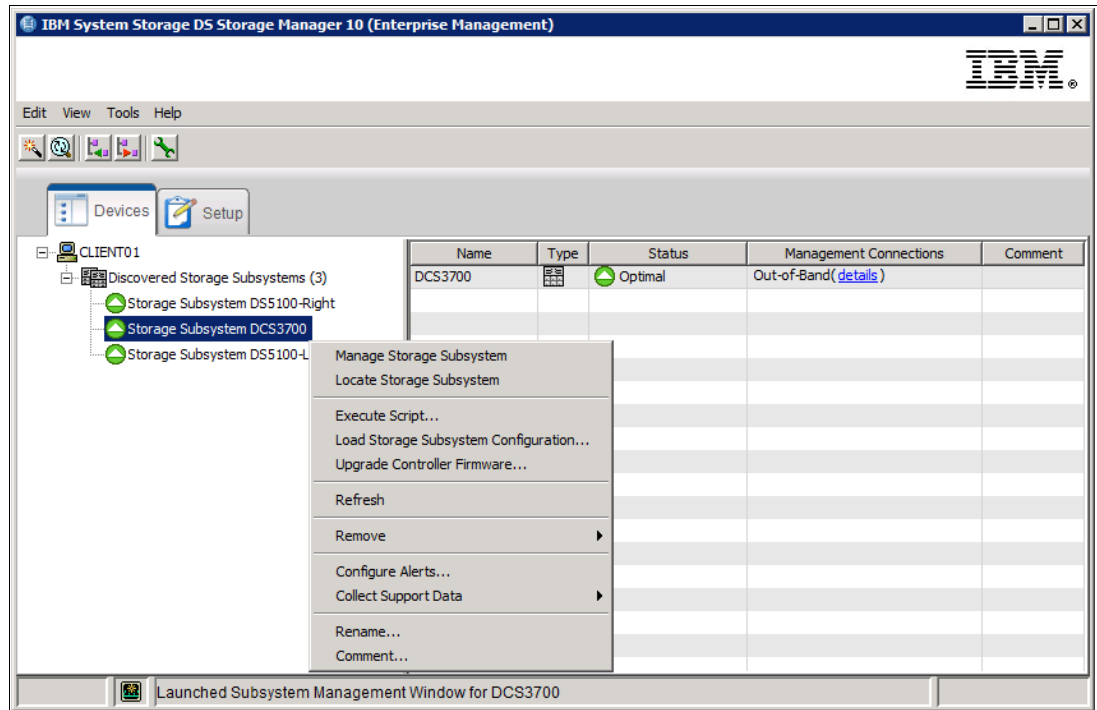


Figure 5-5 Subsystem menu

### Manage Storage Subsystem option

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

### Locate Storage Subsystem option

If you have several storage subsystems that are installed, it can be hard to find the right one. If you select this option, the indicator light on the selected subsystem starts flashing (Figure 5-6).

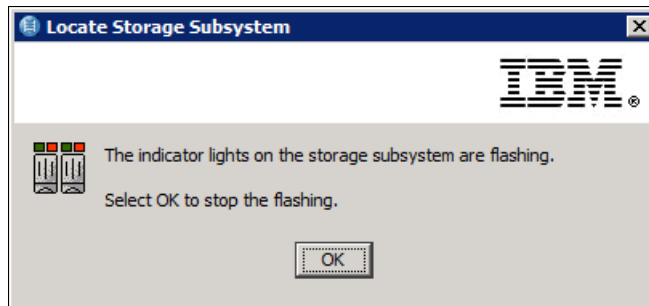


Figure 5-6 Indicator lights are flashing

### Execute Script option

Here you can open the DCS3700 Script Editor. This Editor can be an effective and useful tool to manage your subsystem. For more information about how to use it, see 15.3.1, “Script Editor” on page 465.

## Load Storage Subsystem Configuration option

Section 15.8.1, “Sample command: Save configuration script file” on page 478 describes how to save a configuration script file that is based on the current environment. This is done on the CLI by running 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

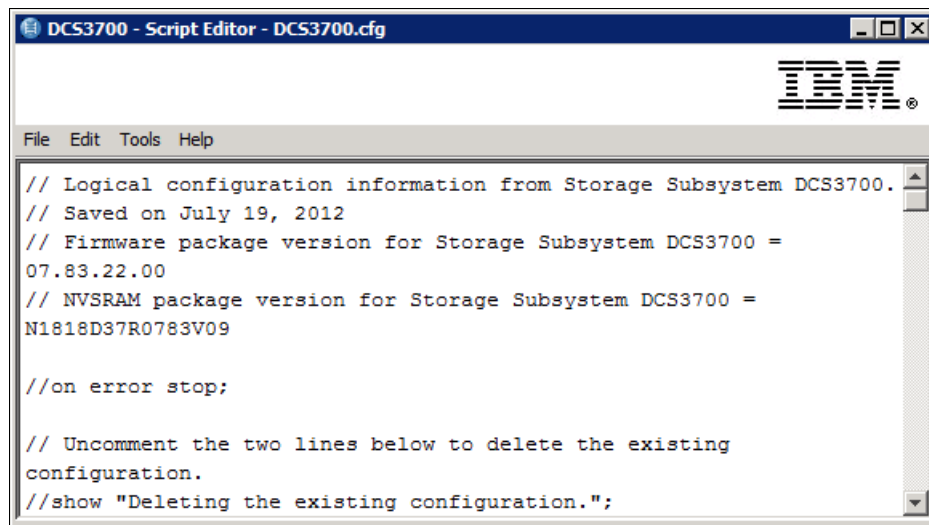
This option appends the uploaded configuration to your storage subsystem unless you do either of the following operations:

- ▶ 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, complete the following steps:

1. Select this option and the load configuration notes display. Read them carefully, and then click **Yes**.
2. The Script Editor opens. Select your configuration file and click **OK**.
3. The file loads in to the Script Editor and the Confirm Load Configuration window opens:
  - a. If you want to append a configuration to the existing one, click **Execute**, and the script appends 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 that are shown in Figure 5-7 and uncomment them. After uncommenting these two lines (shown below), click **Tools** → **Verify** → **Execute** from the menu bar to run the script and upload the new configuration.

```
show "Deleting the existing configuration.";
set storagesubsystem resetConfiguration = true;
```

The image shows a screenshot of a script editor window titled "DCS3700 - Script Editor - DCS3700.cfg". The window has a menu bar with "File", "Edit", "Tools", and "Help". The main text area contains the following script content:

```
// Logical configuration information from Storage Subsystem DCS3700.
// Saved on July 19, 2012
// Firmware package version for Storage Subsystem DCS3700 =
07.83.22.00
// NVSRAM package version for Storage Subsystem DCS3700 =
N1818D37R0783V09

//on error stop;

// Uncomment the two lines below to delete the existing
configuration.
//show "Deleting the existing configuration.";
```

Figure 5-7 Uncomment lines to delete the existing configuration before uploading a new one

**Existing configurations:** The script does not overwrite an existing configuration, that is, the script fails if it attempts to build an array by 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 by using this script because it initializes 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 DCS3700 storage subsystems by using a previously created file for an existing DCS3700. For example, if you have many 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 configure quickly the remaining storage subsystems. You could also use this option to re-create replacement DCS3700s if your original hardware was destroyed in a disaster.

## Upgrade Controller Firmware option

Select this option to upgrade the firmware of the controllers in the storage subsystem. While 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 be until it is activated.

The Upgrade Controller Firmware window and the Download Firmware dialog are shown in Figure 5-8.

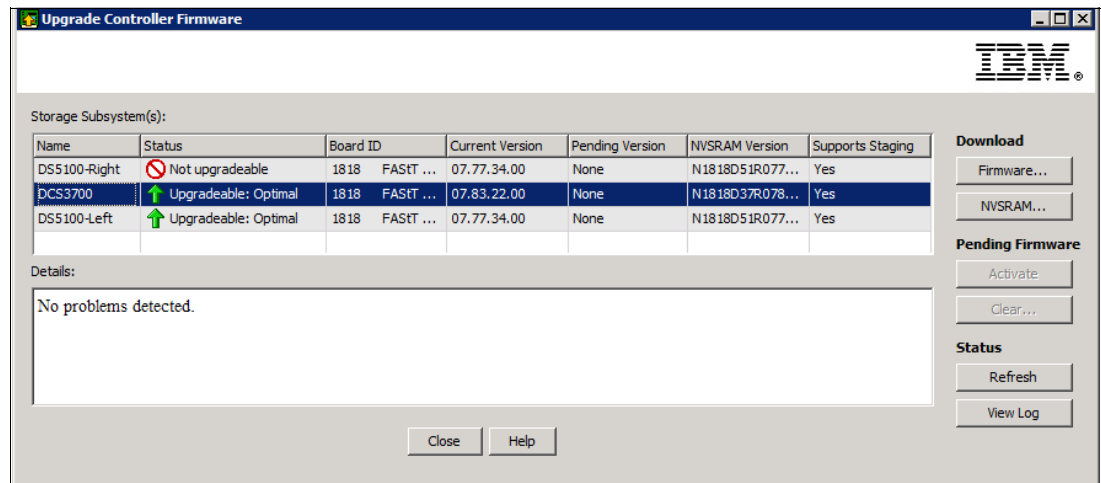


Figure 5-8 Upgrade Controller Firmware window

You also can use the command-line interface (CLI) to download and activate firmware to several storage subsystems. If the Status field shows that the system is not upgradeable, check the Details window for more information about the problem.

For more information about the upgrade procedures, see Chapter 14, “Administration: Upgrade” on page 421.

## Refresh option

Select this option to refresh the view of the Enterprise Management window (Figure 5-9). You might use this option if a system's status has changed, for example, after an error recovery.




Figure 5-9 Refresh View

## Remove option

Here you can remove the selected storage subsystem from your view. This does not affect the views of any other storage management stations. To remove a storage subsystem, select this option and click **OK** (Figure 5-10).

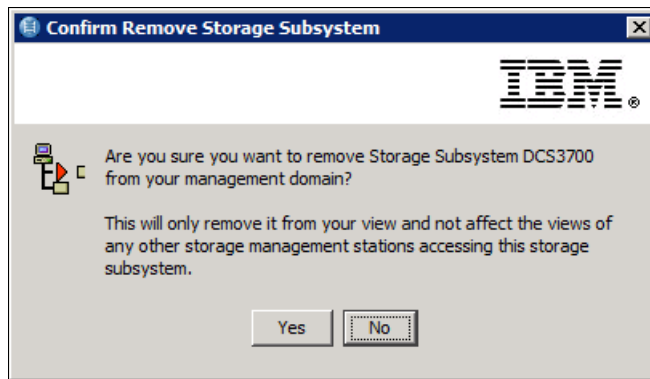


Figure 5-10 Remove Subsystem

## Remove Management Connections window

The Remove Management Connections window is shown in Figure 5-11.

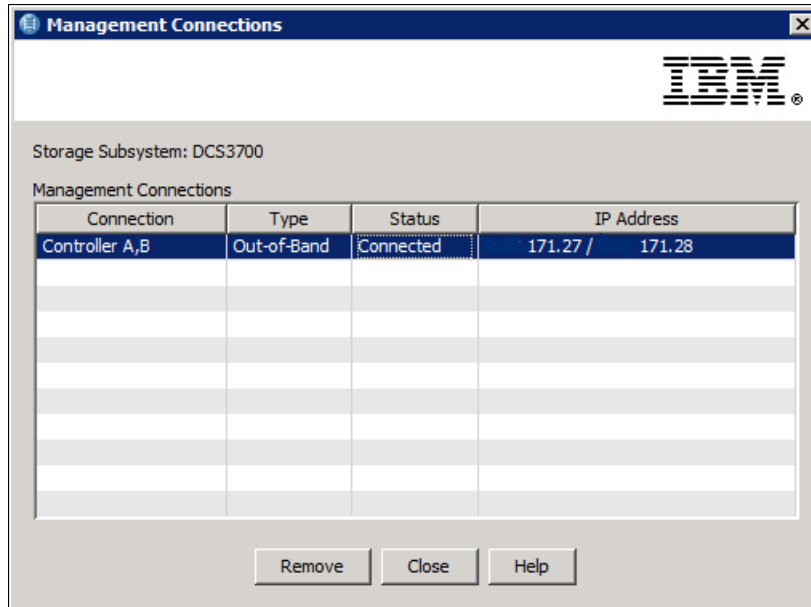


Figure 5-11 Remove Management Connections

## Configure Alerts option

The IBM System Storage DCS3700 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 DCS3700 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 correct alerting even if one server is down or a connection type has failed.

**Note:** The DCS3700 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 are sent.

## What systems are monitored

Depending on how you set up alerts, different storage subsystems are monitored by the Event Monitor. Here is how alerts can be configured:

- ▶ To configure alerts for all storage systems that are 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 5-12 on page 169.

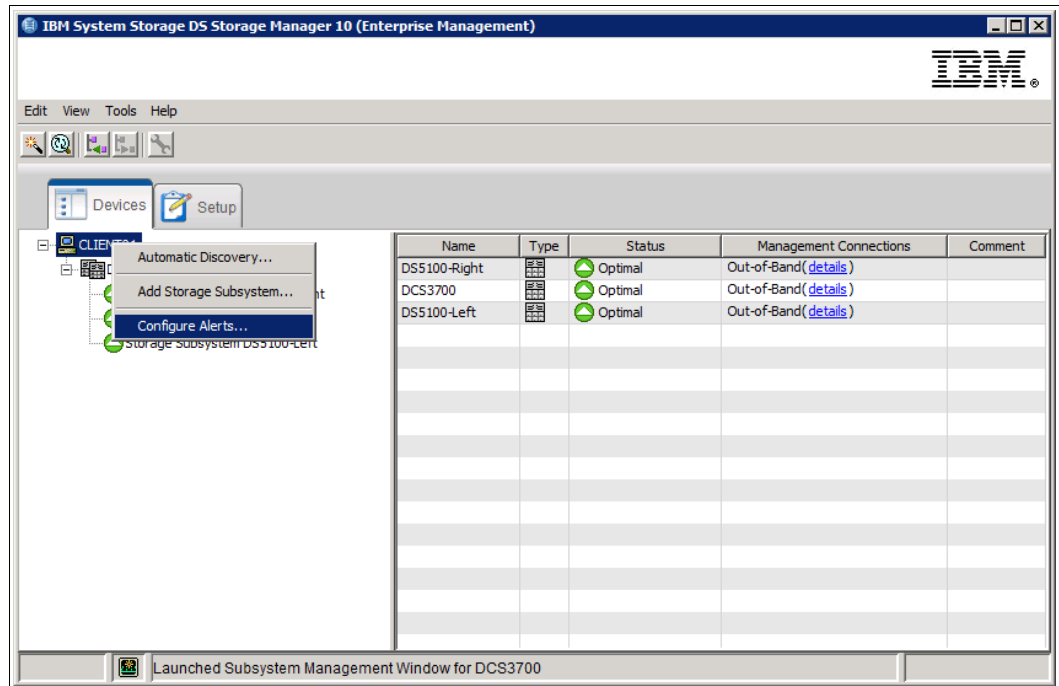


Figure 5-12 Configure Alerts for all systems

- ▶ If you can see the same storage subsystem through separate paths, directly attached, and through different hosts running the host agent, you receive multiple alerts. Right-click a specific storage system to define the alerting specifically for this storage server, as shown in Figure 5-13.

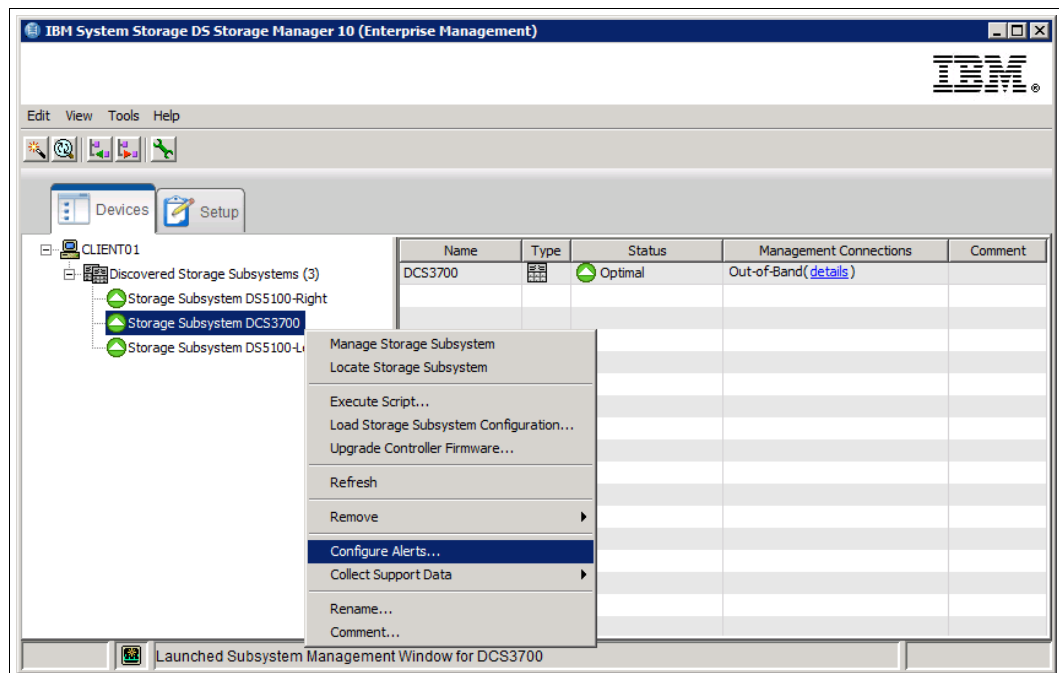


Figure 5-13 Configure Alerts for a particular subsystem

- ▶ Select **Configure Alerts** for one or more systems. The alert window opens (Figure 5-14).

### **Email alerts**

This section describes how to work with email alerts:

1. To send email alerts, you must first define an SMTP server. Enter the name or IP address of your mail server and the sender email address, as shown in Figure 5-14.

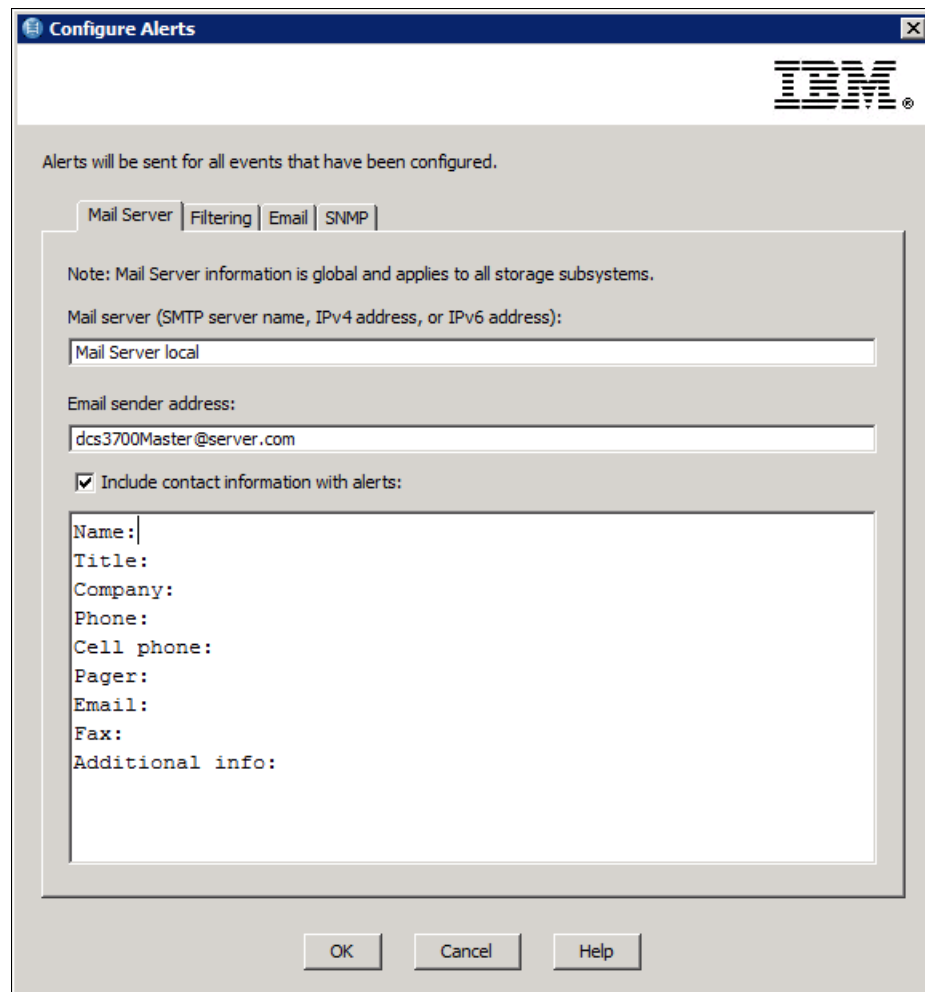


Figure 5-14 *Configure Alerts - define a mail server*



2. Select the **Filtering** tab to configure the events that will trigger an alert, as shown in Figure 5-15.

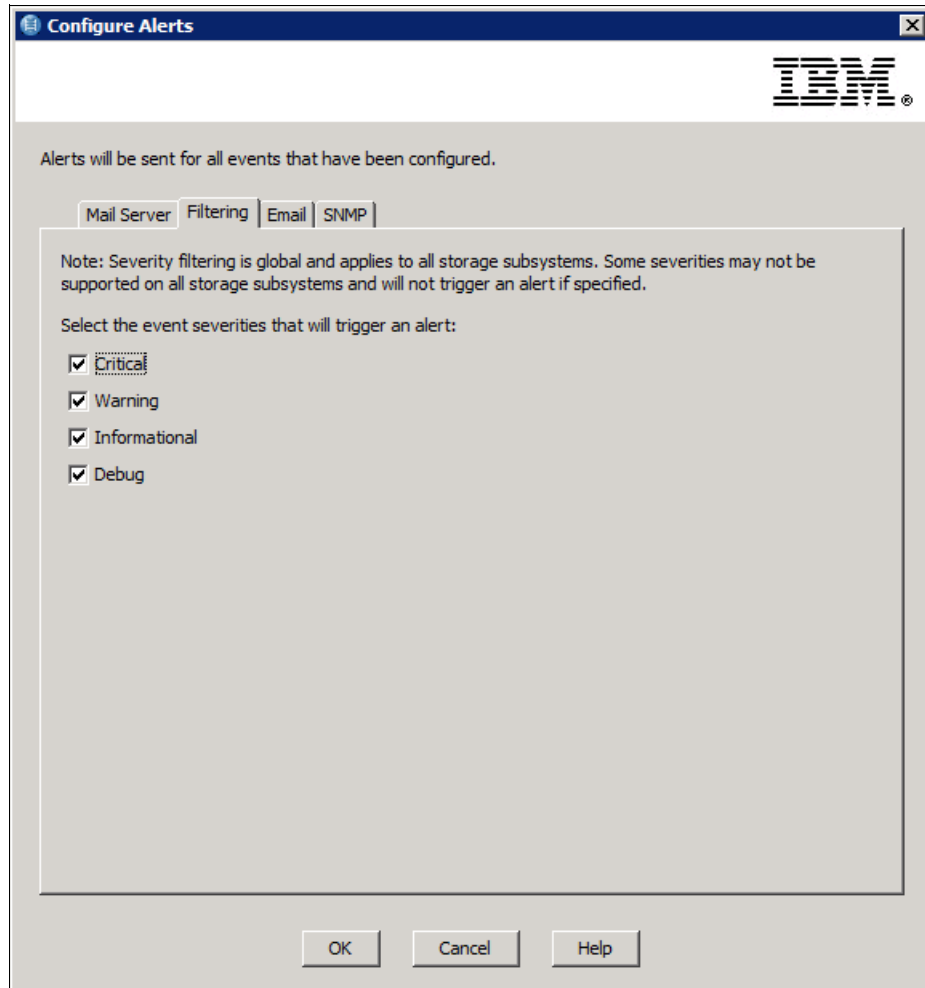


Figure 5-15 Filtering the event that will trigger the alert

3. 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 5-16.

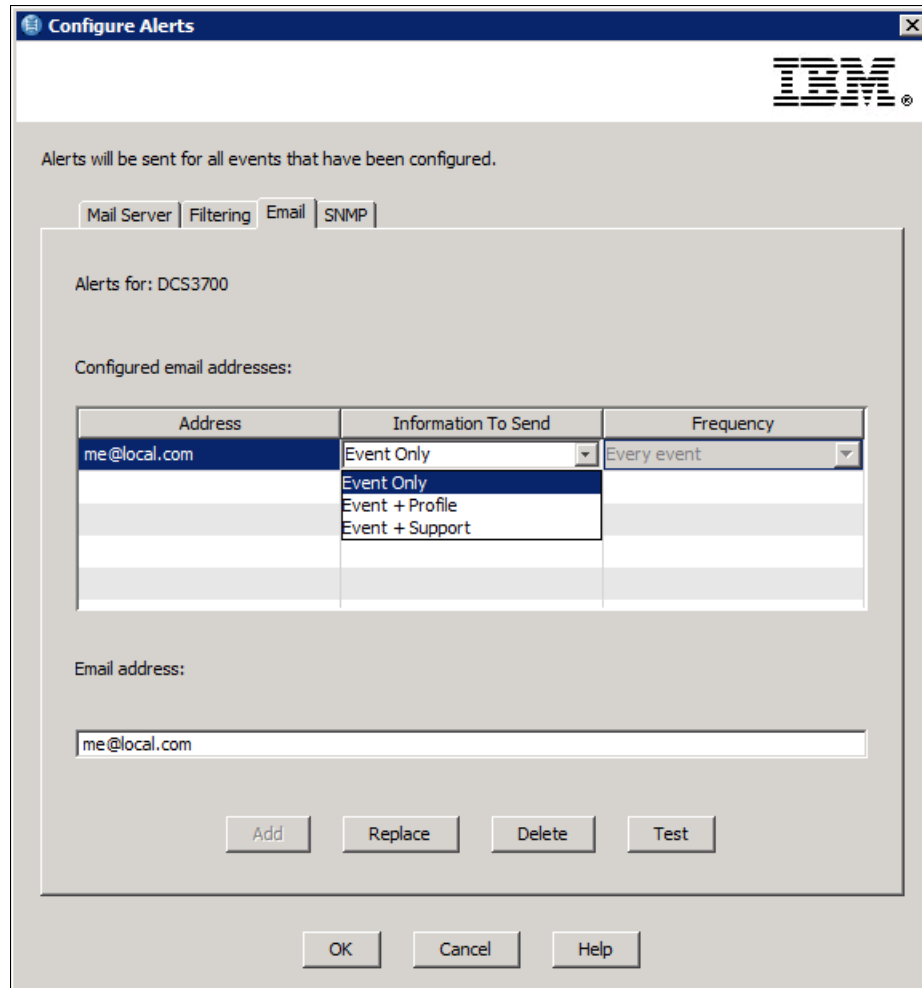


Figure 5-16 Configure Alerts - defining email address

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.

### SNMP alerts

Similarly, select the SNMP tab for receiving and handling the traps that are sent by the service. Enter your community name and the destination dress of your SNMP Server and click **Add**, as shown in Figure 5-17 on page 173.

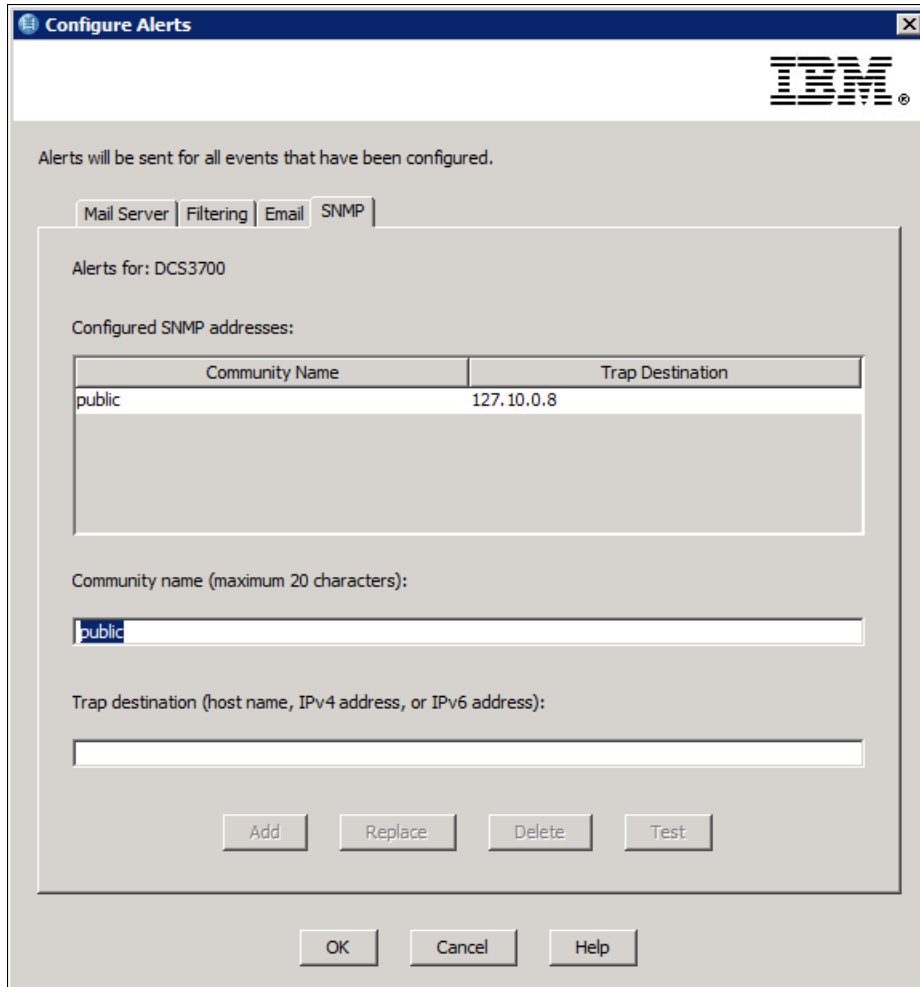


Figure 5-17 Configure Alerts - SNMP

**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.

## Collecting support data

The storage management software can automatically save a copy of the support data when the client monitor process detects an event. You can enable or disable the feature. This can be done either automatically if the user want to collect the data or scheduled. Figure 5-18 shows how to enable the automatic collection of support data.

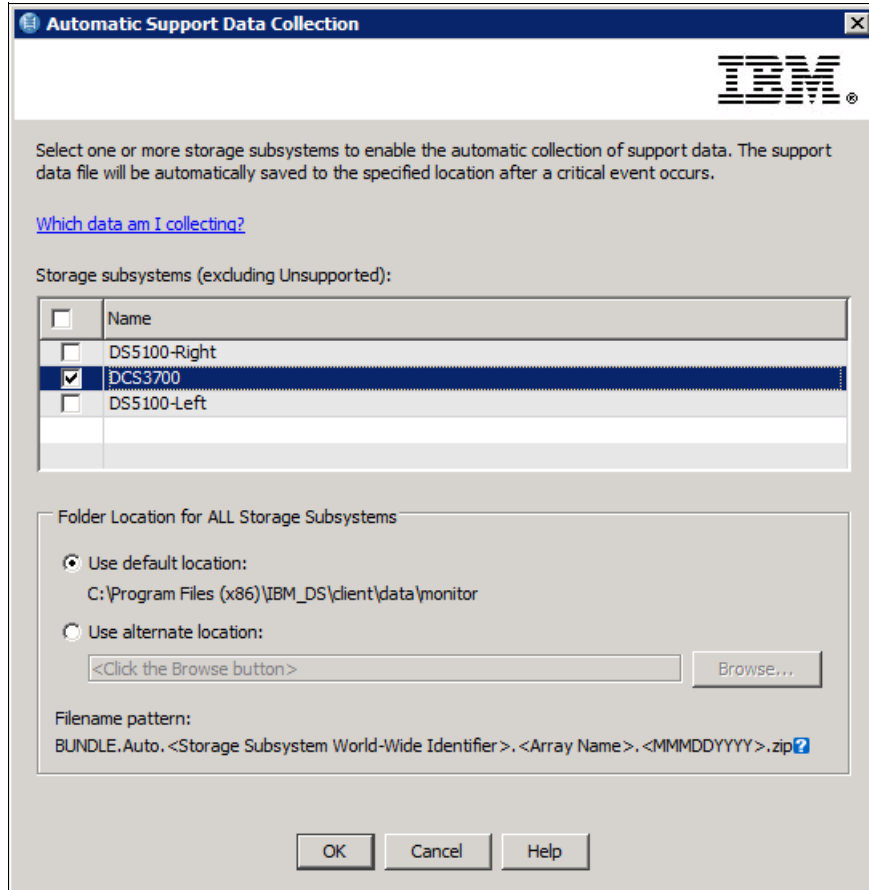


Figure 5-18 Automatic collection of support data

Figure 5-19 shows the window that you use to create/edit the schedule for collecting support data.

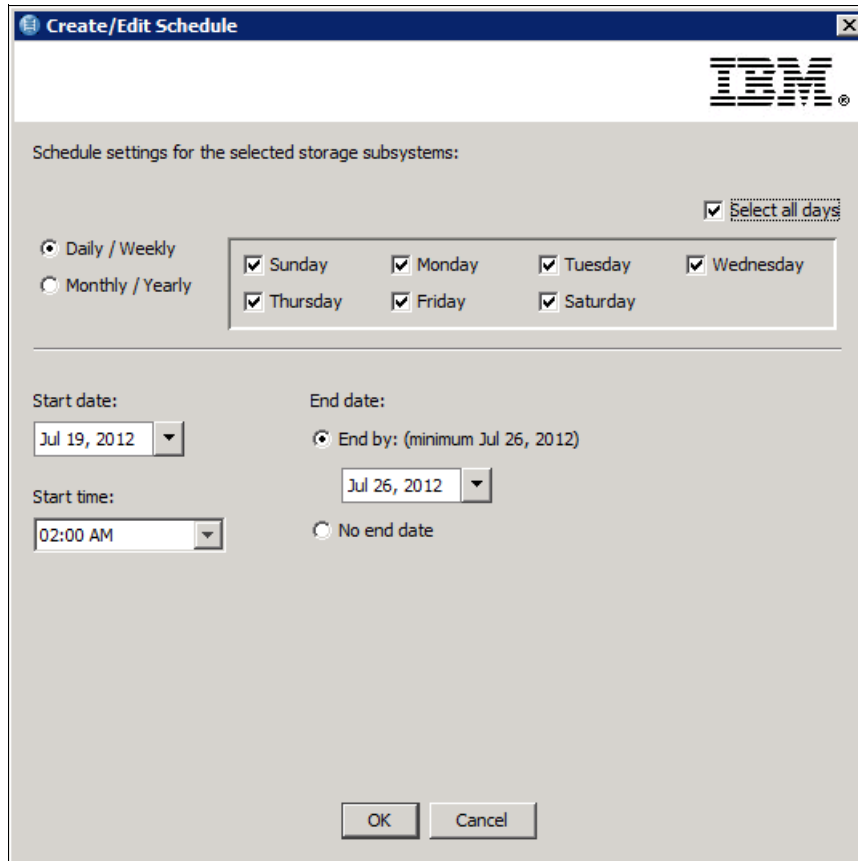


Figure 5-19 Create/Edit a schedule for collecting support data

After you create the schedule, click **OK** and check whether the schedules are updated and check whether you get a confirmation (Figure 5-20).

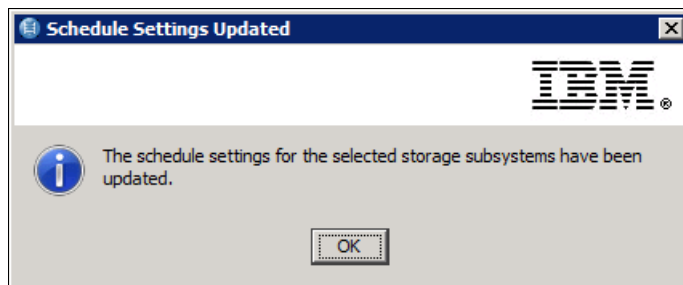


Figure 5-20 Schedule Settings Updated

In the Schedule Support Data Collection window (Figure 5-21), you can verify your settings, and edit, suspend, resume, or remove the schedule.

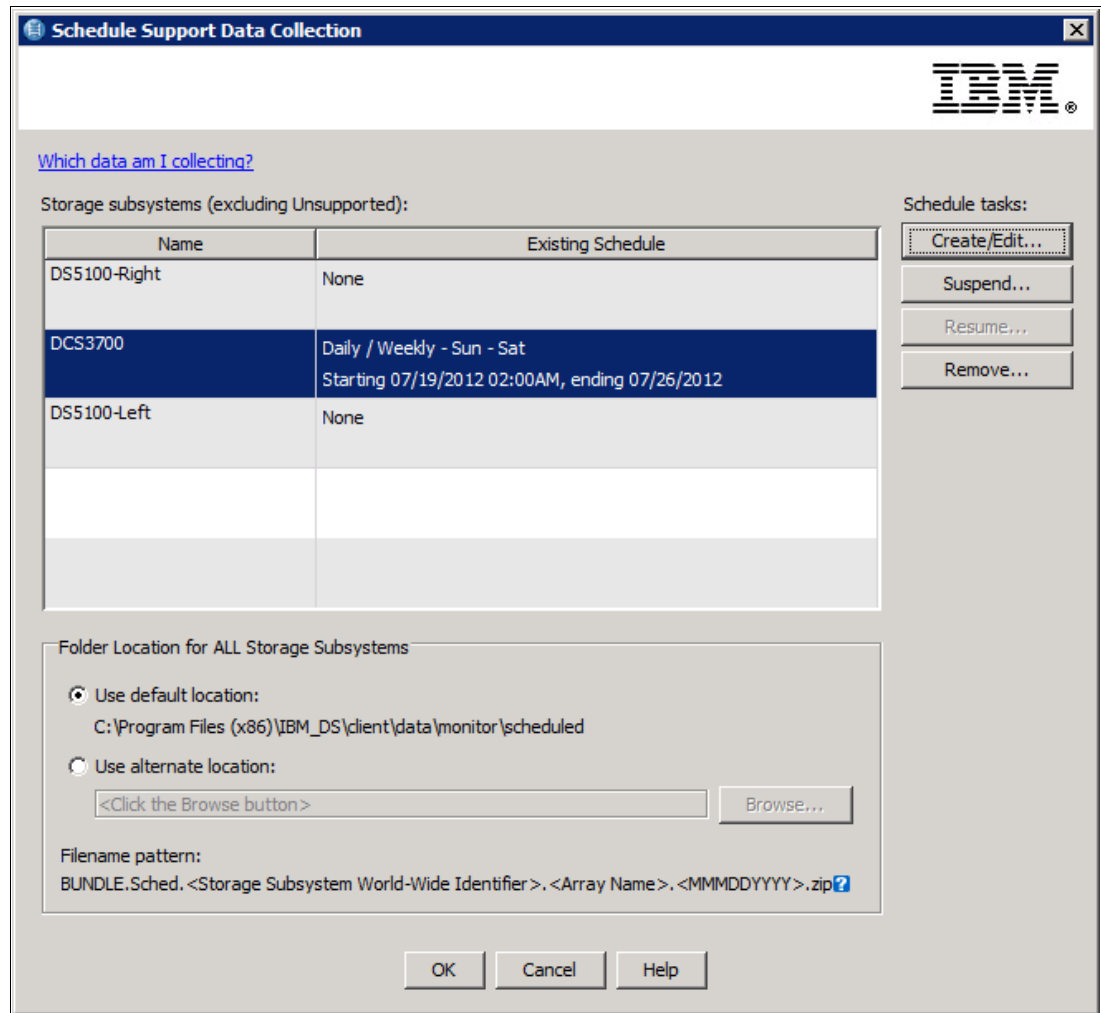


Figure 5-21 Schedule Support Data Collection

The default location for the saved data is C:\Program Files (x86)\IBM\_DS\client\data\monitor. The data is saved in a file starting with the prefix BUNDLE, for example, BUNDLE.Baseline.60080e500023f8f4000000004fffe472.DCS3700.Ju1192012. The file names of the saved files include the date. The operation keeps a maximum of five files. If five files already exist, the operation deletes the oldest file before saving a new one. If there is not enough disk space to save a new file, the operation deletes the oldest files until there is enough space. If your enclosure contains drawers, the diagnostic data for that enclosure is archived in a separate compressed file named EnclosureComponentsStateCapture\_<user-defined>.zip in the same location that is specified for the Customer Support Bundle. This data is requested in a support case.

The type of data that is collected is described in Table 5-1 on page 177.

Table 5-1 Collected support data

Type of data	File name and description
Drive cabling topology	Connections.txt: This file contains information about the connections between the drive expansion unit environmental services module (ESM) and the controller pair.
Drive command aging timeout values	driveCommandAgingTimeout.txt: This file contains the default and current values for the command aging timeout field for every drive.
Features of the storage subsystem	featureBundle.txt: A list of the number of logical drives, drives, and drive expansion units that are allowed in the storage subsystem, and a list of the Premium Features that are available and their limits.
Firmware inventory	firmwareInventory.txt: A list of all of the firmware versions for all of the components in the storage subsystem.
I/O path statistics	ioPathStatistics.zip: Contains raw performance data for each controller that can be used to analyze application performance issues.
iSCSI connections	iscsiSessionConnections.txt: A list of all of the current iSCSI sessions.
iSCSI statistics	iscsiStatistics.csv: Statistics for the Ethernet media access control (MAC), Ethernet Transmission Control Protocol (TCP)/Internet Protocol (IP), and iSCSI target.
Major event log	majorEventLog.txt: A detailed list of events that occur on the storage subsystem. The list is stored in reserved areas on the disks in the storage subsystem. The list records configuration events and component failures in the storage subsystem.
NVSRAM data	NVSRAMdata.txt: A controller file that specifies the default settings for the controllers.
Object bundle	objectBundle: A detailed description of the status of your storage subsystem and its components, which was valid at the time that the file was generated.
Performance statistics	performanceStatistics.csv: A file containing the number of read requests, write requests, cache read requests, and cache write requests for each logical drive in the storage subsystem.
Persistent reservations and registrations	persistentReservations.txt: A detailed list of logical drives on the storage subsystem with persistent reservations and registrations.
Read link status data	readLinkStatus.csv: A detailed list of errors that have been detected in the traffic flow between storage subsystems on the Fibre Channel loop. The archive also might include a file of historical read link status data.
Recovery Guru procedures	recoveryGuruProcedures.html: A detailed list of all of the Recovery Guru topics that are issued in response to problems detected on the storage subsystem.
Recovery profile	recoveryProfile.csv: A detailed description of the latest recovery profile record and historical data.
Switch-on-a-chip statistics	socStatistics.csv: A detailed list of all the statistics and error counters that are found in the traffic between controllers and switched FC devices.
State capture data	stateCaptureData.txt: A detailed description of the current state of your storage subsystem.
Storage subsystem configuration	storageSubsystemConfig.cfg: A detailed description of the logical configuration on your storage subsystem.
Storage subsystem profile	storageSubsystemProfile.txt: A description of all of the components and properties of a storage subsystem.

Type of data	File name and description
Trace buffer contents	traceBuffers.zip: This file contains the contents of the controllers' trace buffers that are used to record debug information.
Enclosure capture data	EnclosureComponentsStateCapture_<user-defined>.zip: If your enclosure contains drawers, the diagnostic data is archived in this compressed file. The compressed file contains a separate text file for each enclosure that contains drawers. The compressed file is saved as a separate file outside of the Customer Support Bundle.
Unreadable sectors	unreadableSectors.txt: A detailed list of all of the unreadable sectors that have been logged to the storage subsystem.

### Rename option

If you have multiple DCS3700 storage subsystems that are installed, give each one a user-specified name for easier identification. To rename a DCS3700, select this option, enter a new name for the subsystem, and click **OK**, as shown in Figure 5-22.

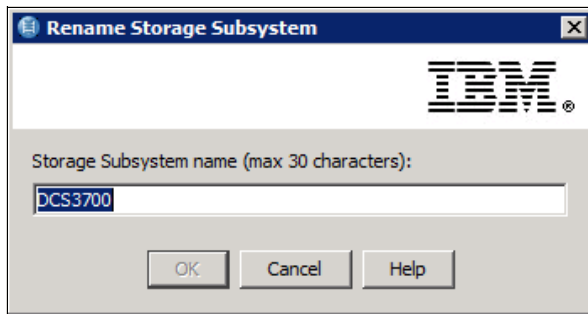


Figure 5-22 Rename Storage Subsystem

**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 option

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 5-23.

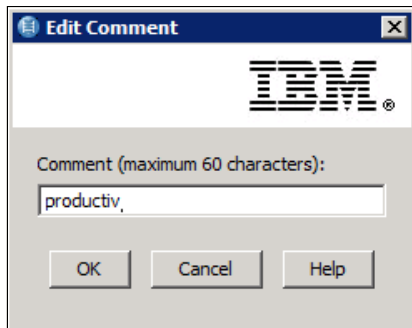


Figure 5-23 Edit comment





## Edit and Tools submenu

All the tasks that are described in 5.2.1, “Subsystem menu” on page 164 can also be accessed from the task menu bar. To do this, highlight a subsystem and open the appropriate menu. The Edit (Figure 5-25) and Tools (Figure 5-26) submenus contain all of the tasks. Performing these tasks is done exactly as described in the previous section.

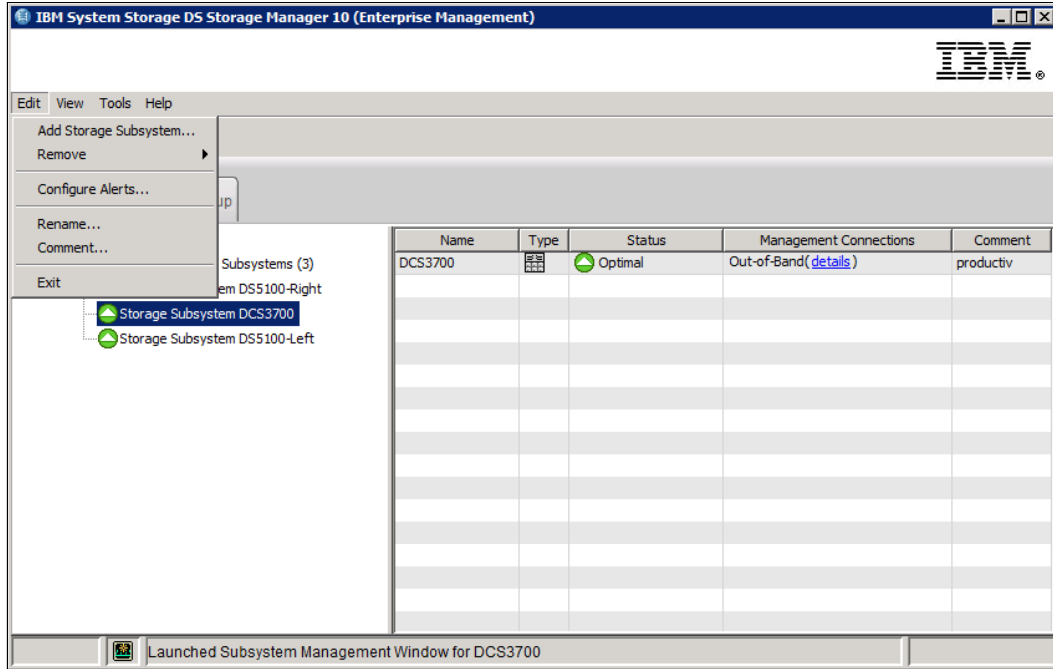


Figure 5-25 Edit menu

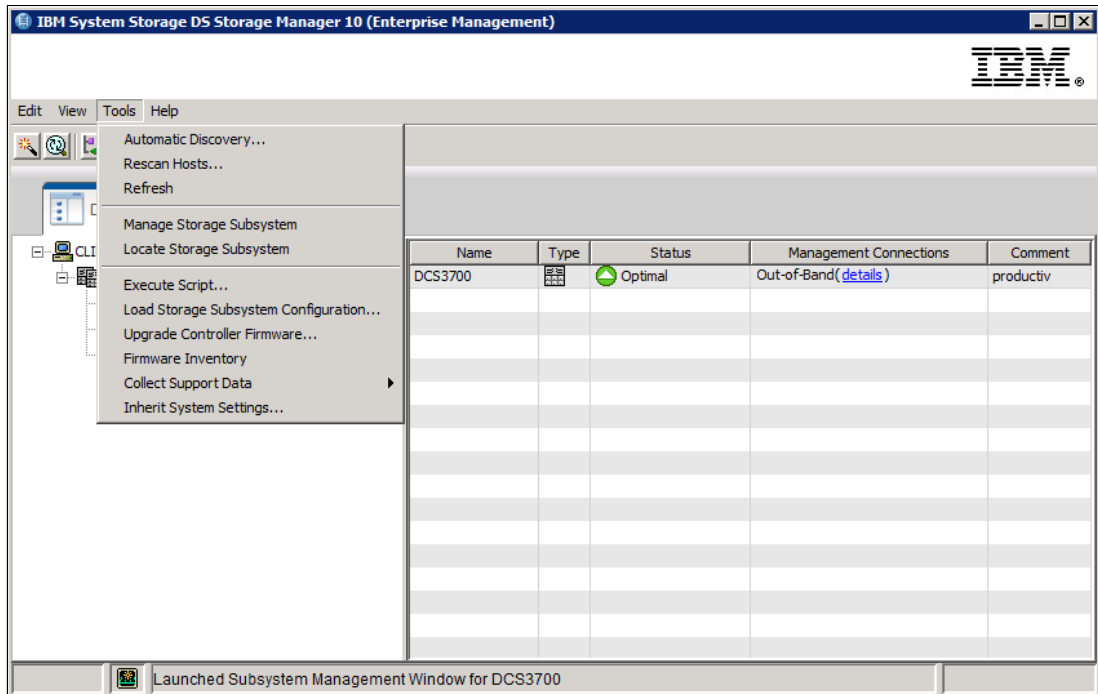


Figure 5-26 Tools menu

## The View submenu

Use this menu to start the Enterprise Management window Task Assistant (Figure 5-27) or to customize the view of your enterprise window by selecting one of the sort options. The Task Assistant is described in 5.1.1, “Initial Setup Tasks” on page 160.

### **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 Connection
- ▶ 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.

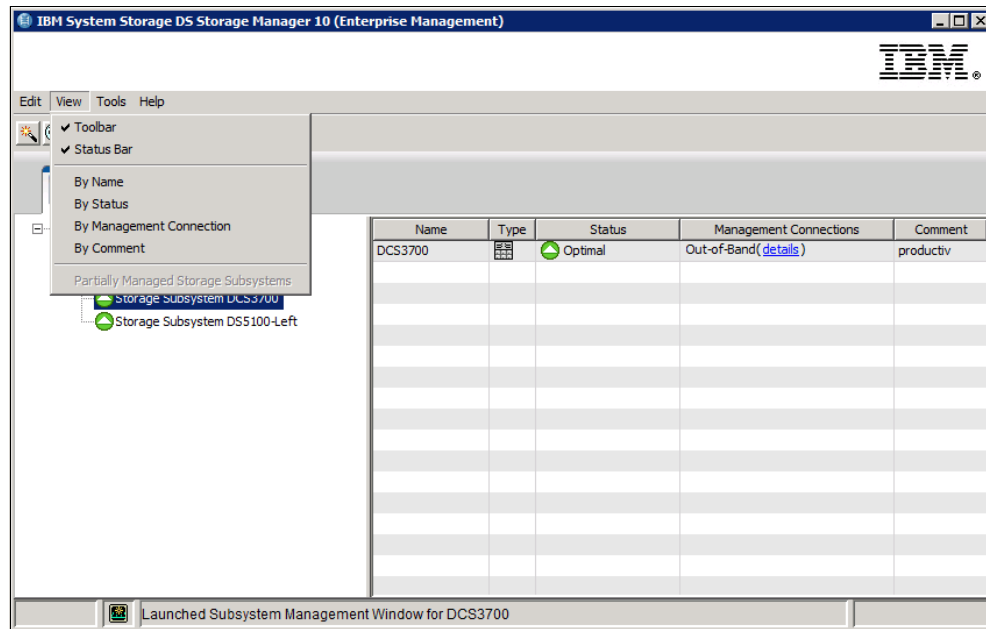


Figure 5-27 View submenu

## 5.2.3 The Quick Access icons

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

### **Automatically discover new storage subsystems icon**

The Automatic discovery icon for the new storage subsystems icon is shown in Figure 5-28.



Figure 5-28 Automatically discover new storage systems icon

This icon starts an automatic discovery of new storage subsystems (Figure 5-29) by sending in-band and out-of-band broadcasts. If it finds directly attached storage subsystems or hosts running the DCS3700 Storage Manager agent (with an attached storage subsystem), it adds these storage subsystems to the Enterprise Management window.

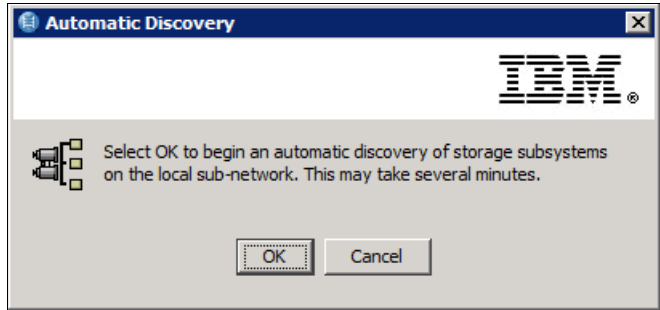


Figure 5-29 Automatic Discovery

**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 icon” on page 182.

### Rescan selected host for new storage subsystems icon

The Rescan selected host for new storage subsystems icon is shown in Figure 5-30.



Figure 5-30 Rescan Host icon

If you are using a host that is running an agent to manage your storage subsystems, you can use this icon to rescan only the selected host. To do this task, highlight the host, click this icon, and the Rescan window appears. Click **OK** and the host scans for newly attached storage devices.

### Add Storage Subsystem to view icon

The Add Storage Subsystem to view icon is shown in Figure 5-31.



Figure 5-31 Add Storage Subsystem to view icon

Click this icon to display the Add Storage Subsystem window. Here you can add systems manually to your Enterprise window view. For out-of-band management, you must provide the IP address or host names of the controllers (Figure 5-32). 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):  
192.168.128.101

Controller (DNS/Network name, IPv4 address, or IPv6 address):  
192.168.128.102

**In-band management:**  
Manage the storage subsystem through an attached host.

Host (DNS/Network name, IPv4 address, or IPv6 address):  
\_\_\_\_\_

Add Cancel Help

Figure 5-32 Manually add out-of-band system

For in-band management, specify the IP address or host name of the host that is attached to the storage server (Figure 5-33).

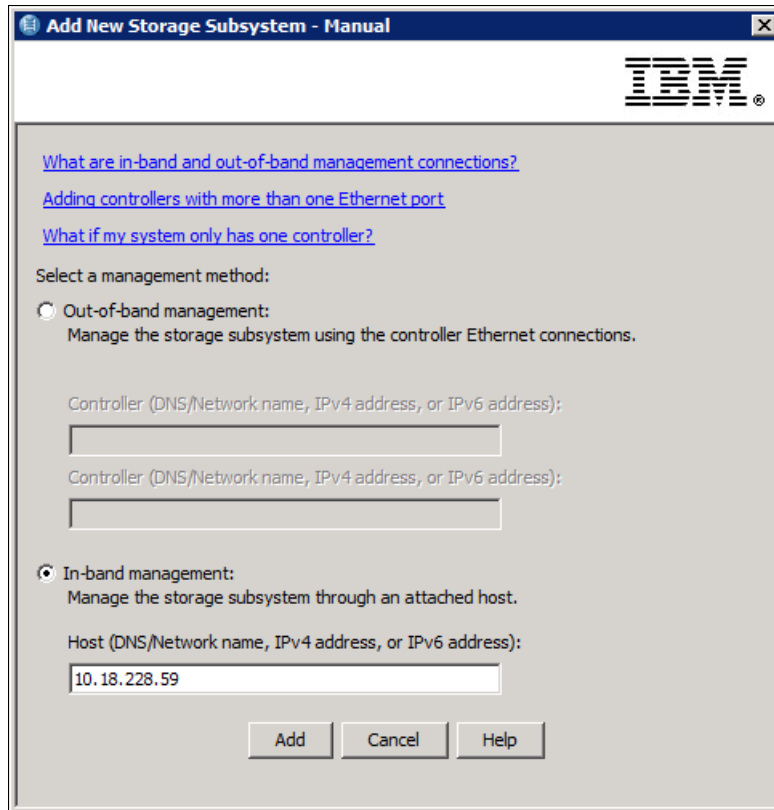


Figure 5-33 Manually add in-band system

## Remove Storage Subsystem from View icon

The Remove Storage Subsystem from View icon is shown in Figure 5-34.



Figure 5-34 Remove Storage Subsystem from View icon

To remove a storage subsystem from your view, highlight it, click this icon, and click **OK** to confirm (Figure 5-35). This subsystem is removed only from your view; the views of other management stations are not affected.

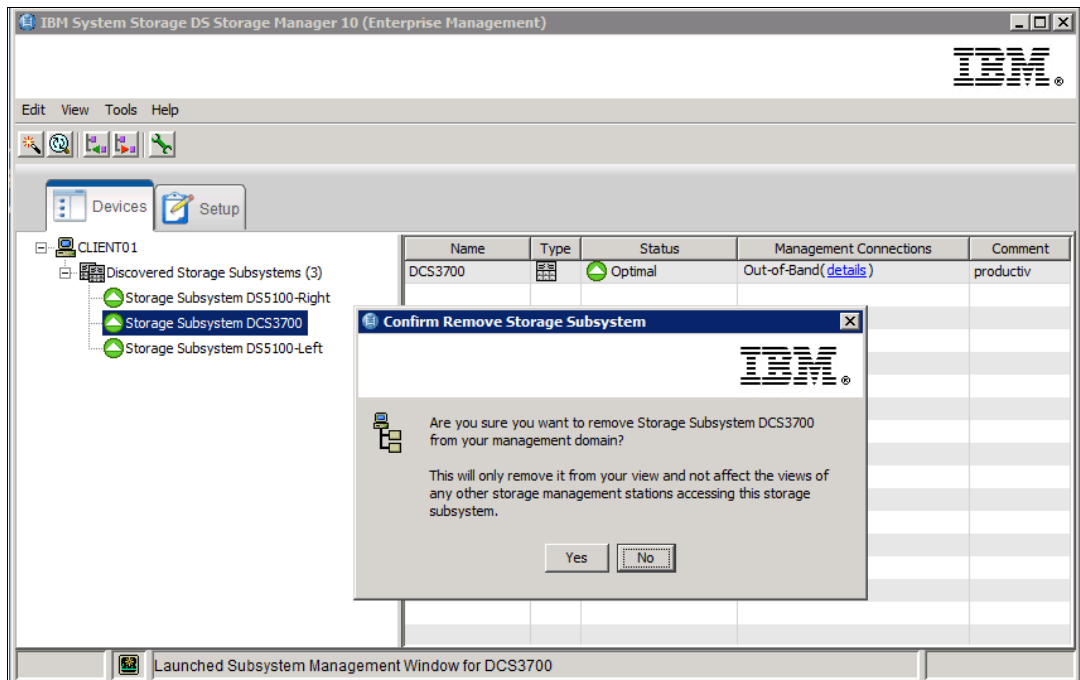


Figure 5-35 Remove Subsystem from view

### Launch Management window icon

The Launch Management window icon is shown in Figure 5-36.



Figure 5-36 Launch Management window icon

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







## Administration: Summary tab

This chapter gives you an overview of the Summary tab of the IBM System Storage DS (Subsystem Management) window.

The following topics are described:

- ▶ 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

## 6.1 DS Storage Manager: Summary tab

The first tab of the Subsystem Management window is the Summary tab. Figure 6-1 shows an IBM System Storage DCS3700 storage subsystem that is running in an optimal status.

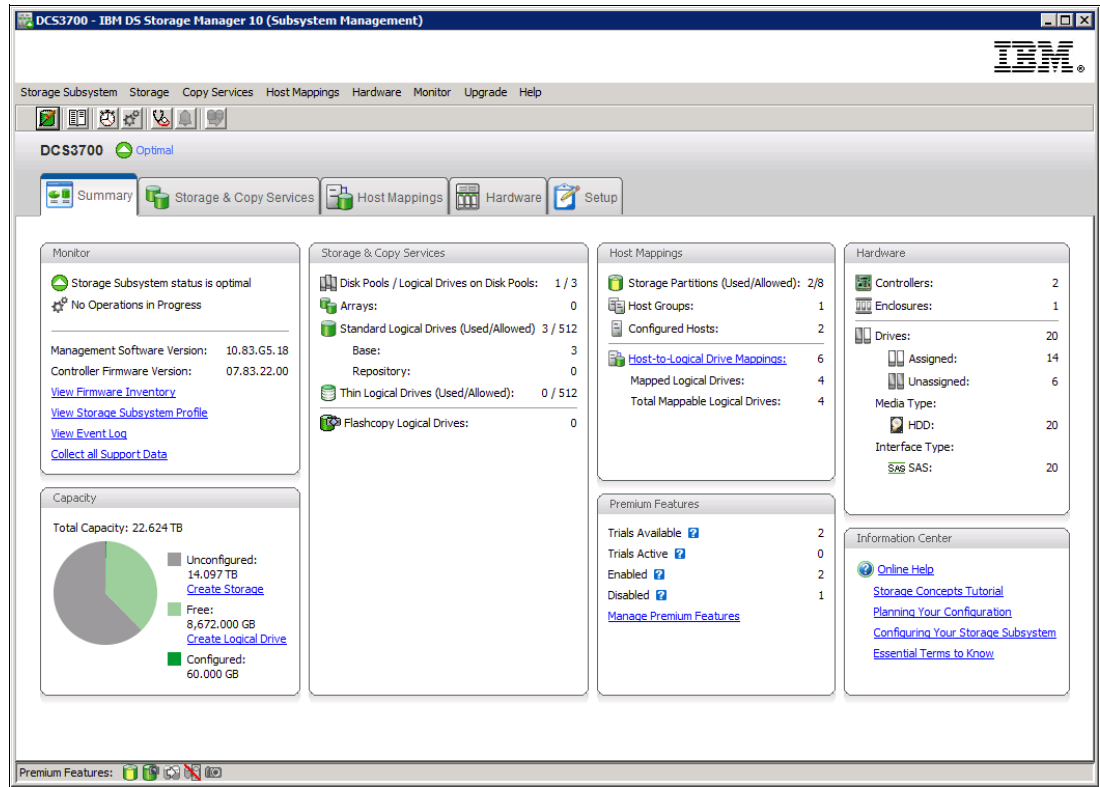


Figure 6-1 IBM System Storage DS Storage Manager Summary tab

The Summary tab is divided into seven panes, which are described in the following sections:

- ▶ Monitor
- ▶ Capacity
- ▶ Storage & Copy Services
- ▶ Host Mappings
- ▶ Premium Features
- ▶ Hardware
- ▶ Information Center

## 6.2 Monitor

The first pane is the Monitor section, which displays the basic status of your storage subsystem (Figure 6-2).

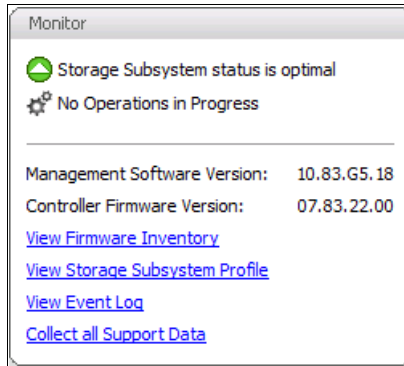


Figure 6-2 Monitor

### 6.2.1 Storage subsystem status

The first link in the Monitor pane shows the current subsystem status. If the status is green, as shown in Figure 6-2, the subsystem is running in an optimal status. If there is any problem with the subsystem, the status icon changes to red (Figure 6-3). Click **Storage Subsystem Needs Attention** to open the Recovery Guru. For more information about how to troubleshoot and recover from problems, see Chapter 13, “Administration: Monitor” on page 385. The same red icon is also visible in upper part of the Subsystem Manager window, close to the name of subsystem, and the Recovery Guru icon in the toolbar is flashing.




Figure 6-3 Monitor box - Storage Subsystem Needs Attention

The next two lines inform you about the installed Management and Controller Versions.

### 6.2.2 Storage Subsystem communication error

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 6-4). If only the management connection is lost, the attached hosts can still access data on the disks, but you cannot manage your subsystem or receive alerts. The Subsystem Management window then displays the last known state of the subsystem only, and you cannot perform any administrative operations.




Figure 6-4 Storage Subsystem communication error

## 6.2.3 Operations in Progress

The second link shows operations that are in progress (Figure 6-5). If there is at least one operation in progress, you can click the link to open a window with more detailed information (Figure 6-6). You can find more information about operations in progress in Chapter 8, “Administration: Storage & Copy Services tab” on page 225.




Figure 6-5 Operations in Progress

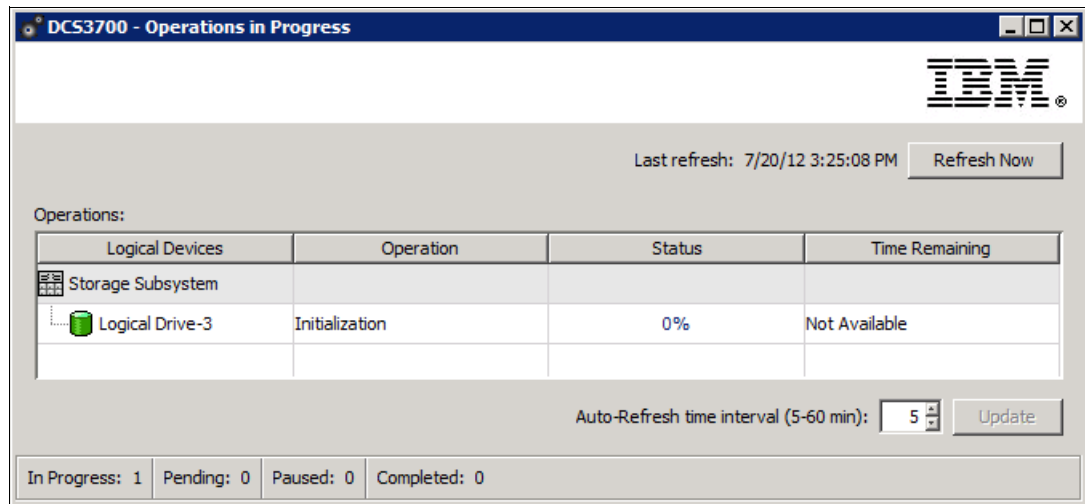


Figure 6-6 Detailed View Operations in Progress

## 6.2.4 View Firmware Inventory

This window informs you about all the firmware levels of the system. The window in Figure 6-7 on page 191 and Figure 6-8 on page 191 show examples of the View Firmware Inventory pane.

Storage Subsystem: DCS3700	
Firmware Inventory:	
<b>IBM DS Storage Manager 10</b>	
SMW Version:	10.83.G5.18
Report Date:	Fri Jul 20 13:51:13 EDT 2012
<b>Storage Subsystem</b>	
Storage Subsystem Name:	DCS3700
Current Package Version:	07.83.22.00
Current NVSRAM Version:	N1818D37R0783V09
Staged Package Version:	None
Staged NVSRAM Version:	None
<b>Controllers</b>	
Location:	Enclosure 99, Slot A
Current Package Version:	07.83.22.00
Current NVSRAM Version:	N1818D37R0783V09
Board ID:	2660
Sub-Model ID:	162
Location:	Enclosure 99, Slot B
Current Package Version:	07.83.22.00
Current NVSRAM Version:	N1818D37R0783V09
Board ID:	2660
Sub-Model ID:	162
<b>Power Supplies</b>	
Location:	Enclosure 99 Right
Firmware Version:	Not Available
Location:	Enclosure 99 UNKNOWN
Firmware Version:	Not Available

Figure 6-7 View Firmware Inventory - 1

Drive	Enclosure, Drawer, Slot:	Manufacturer:	Product ID:	Drive Type:	Capacity:	Drive Firmware Version:	FPGA Version: (SSD only)
	Enclosure 99, Drawer 1, Slot 1	SEAGATE	ST31000425SS	Serial Attached SCSI (SAS)	999668015104	Not Available	Not Available
	Enclosure 99, Drawer 1, Slot 4	SEAGATE	ST33000651SS	Serial Attached SCSI (SAS)	3000056111104	Not Available	Not Available
	Enclosure 99, Drawer 1, Slot 7	SEAGATE	ST31000425SS	Serial Attached SCSI (SAS)	999668015104	Not Available	Not Available
	Enclosure 99, Drawer 1, Slot 10	SEAGATE	ST31000425SS	Serial Attached SCSI (SAS)	999668015104	Not Available	Not Available
	Enclosure 99, Drawer 2, Slot 1	SEAGATE	ST33000651SS	Serial Attached SCSI (SAS)	3000056111104	Not Available	Not Available
	Enclosure 99, Drawer 2, Slot 4	SEAGATE	ST33000651SS	Serial Attached SCSI (SAS)	3000056111104	Not Available	Not Available
	Enclosure 99, Drawer 2, Slot 7	SEAGATE	ST31000425SS	Serial Attached SCSI (SAS)	999668015104	Not Available	Not Available
	Enclosure 99, Drawer 2, Slot 10	SEAGATE	ST31000425SS	Serial Attached SCSI (SAS)	999668015104	Not Available	Not Available
	Enclosure 99, Drawer 3, Slot 1	SEAGATE	ST31000425SS	Serial Attached SCSI (SAS)	999668015104	Not Available	Not Available
	Enclosure 99, Drawer 3, Slot 4	SEAGATE	ST31000425SS	Serial Attached SCSI (SAS)	999668015104	Not Available	Not Available
	Enclosure 99, Drawer 3, Slot 7	SEAGATE	ST3500415SS	Serial Attached SCSI (SAS)	499570991104	Not Available	Not Available
	Enclosure 99, Drawer 3, Slot 10	SEAGATE	ST31000425SS	Serial Attached SCSI (SAS)	999668015104	Not Available	Not Available
	Enclosure 99, Drawer 4, Slot 1	SEAGATE	ST31000425SS	Serial Attached SCSI (SAS)	999668015104	Not Available	Not Available
	Enclosure 99, Drawer 4, Slot 4	SEAGATE	ST31000425SS	Serial Attached SCSI (SAS)	999668015104	Not Available	Not Available
	Enclosure 99, Drawer 4, Slot 7	SEAGATE	ST31000425SS	Serial Attached SCSI (SAS)	999668015104	Not Available	Not Available
	Enclosure 99, Drawer 4, Slot 10	SEAGATE	ST31000425SS	Serial Attached SCSI (SAS)	999668015104	Not Available	Not Available
	Enclosure 99, Drawer 5, Slot 1	SEAGATE	ST33000651SS	Serial Attached SCSI (SAS)	3000056111104	Not Available	Not Available
	Enclosure 99, Drawer 5, Slot 4	SEAGATE	ST31000425SS	Serial Attached SCSI (SAS)	999668015104	Not Available	Not Available
	Enclosure 99, Drawer 5, Slot 7	SEAGATE	ST33000651SS	Serial Attached SCSI (SAS)	3000056111104	Not Available	Not Available
	Enclosure 99, Drawer 5, Slot 10	SEAGATE	ST31000425SS	Serial Attached SCSI (SAS)	999668015104	Not Available	Not Available

Figure 6-8 View Firmware Inventory - 2

## 6.2.5 View 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 that is known as a *subsystem profile*. This profile stores information (properties) about the controllers, attached drives, enclosures, microcode levels, arrays, logical drives, storage partitioning, and so on.

To generate and open the subsystem profile, click **View Storage Subsystem Profile** in the Monitor box (Figure 6-2 on page 189). The information is gathered from the various components and the profile window opens (Figure 6-9). The various tabs show the properties of various objects in storage subsystem. Click **Save As** to save the whole profile locally to disk to help document a change history of your storage subsystem. It is 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, see Chapter 13, “Administration: Monitor” on page 385.

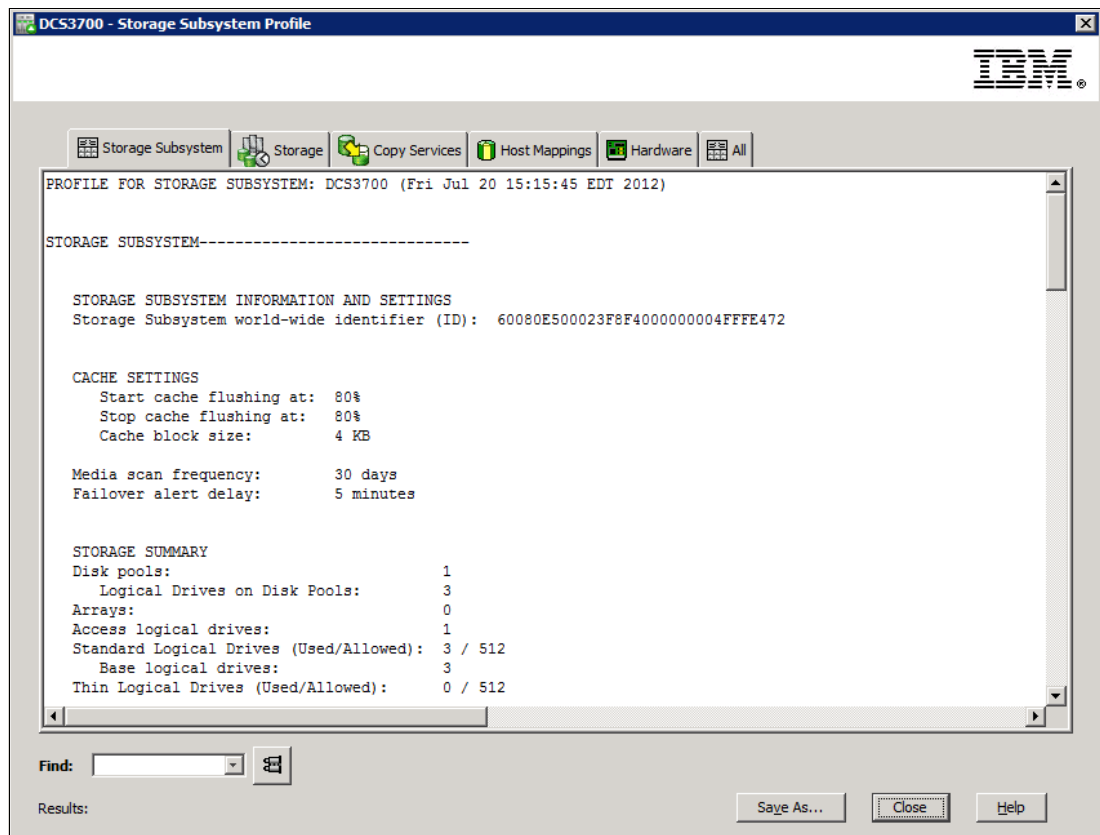
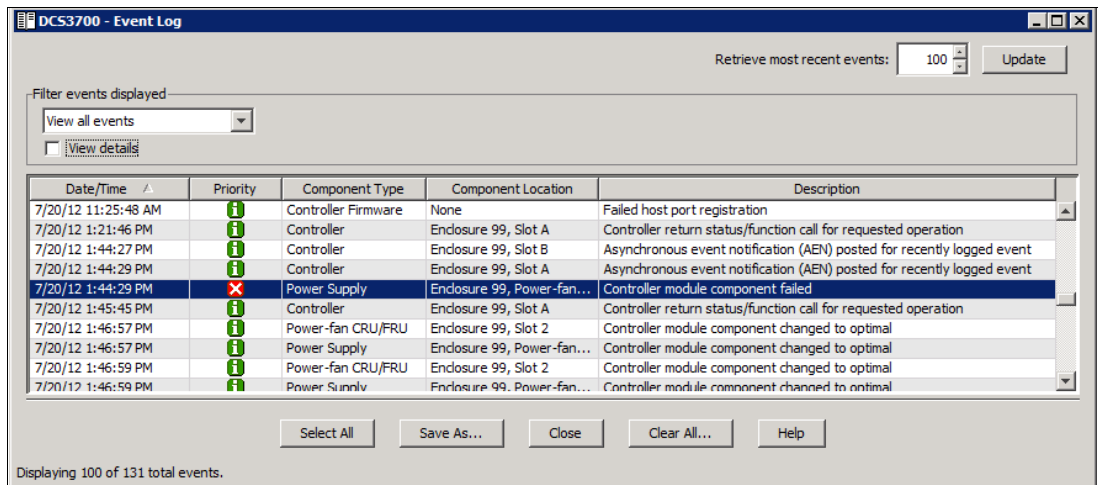


Figure 6-9 Storage Subsystem Profile window

**Note:** You should save a new version of the profile (or collect all support data (CASD) as described in Chapter 13, “Administration: Monitor” on page 385) and store it securely whenever the configuration changes. Even if there is a complete configuration loss, your data on disks might not be affected, and 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.

## 6.2.6 View Event Log

Use the Event Log Viewer to show a detailed list of events that occur in a storage subsystem. The Event Log is stored on reserved areas on the storage subsystem disks. It records configuration events and storage subsystem component failures. The Event Log stores approximately 8000 events before it replaces an event with a new event. If you want to keep the events, you may save them, and clear them from the Event Log. The Event Log window shows two types of event views. The first window displays only the event summary in table form, as shown in Figure 6-10.



The screenshot shows the 'DCS3700 - Event Log' window. At the top right, there is a 'Retrieve most recent events:' field set to '100' and an 'Update' button. Below this is a 'Filter events displayed' section with a dropdown menu set to 'View all events' and a 'View details' checkbox. The main area contains a table with the following columns: Date/Time, Priority, Component Type, Component Location, and Description. The table lists several events, with the most recent one (7/20/12 1:44:29 PM) highlighted in blue, indicating a critical failure of the Power Supply. Below the table are buttons for 'Select All', 'Save As...', 'Close', 'Clear All...', and 'Help'. At the bottom left, it says 'Displaying 100 of 131 total events.'

Date/Time	Priority	Component Type	Component Location	Description
7/20/12 11:25:48 AM	Information	Controller Firmware	None	Failed host port registration
7/20/12 1:21:46 PM	Information	Controller	Endosure 99, Slot A	Controller return status/function call for requested operation
7/20/12 1:44:27 PM	Information	Controller	Endosure 99, Slot B	Asynchronous event notification (AEN) posted for recently logged event
7/20/12 1:44:29 PM	Information	Controller	Endosure 99, Slot A	Asynchronous event notification (AEN) posted for recently logged event
7/20/12 1:44:29 PM	Critical	Power Supply	Endosure 99, Power-fan...	Controller module component failed
7/20/12 1:45:45 PM	Information	Controller	Endosure 99, Slot A	Controller return status/function call for requested operation
7/20/12 1:46:57 PM	Information	Power-fan CRU/FRU	Endosure 99, Slot 2	Controller module component changed to optimal
7/20/12 1:46:57 PM	Information	Power Supply	Endosure 99, Power-fan...	Controller module component changed to optimal
7/20/12 1:46:59 PM	Information	Power-fan CRU/FRU	Endosure 99, Slot 2	Controller module component changed to optimal
7/20/12 1:46:59 PM	Information	Power Supply	Endosure 99, Power-fan...	Controller module component changed to optimal

Figure 6-10 Event summary in table form

When you select the **View Details** check box (see Figure 6-11), the window splits into two sections views. The summary of a highlighted event is displayed, as shown in Figure 6-11. To resize the panes, select the splitter bar between the two views, and move it up or down.

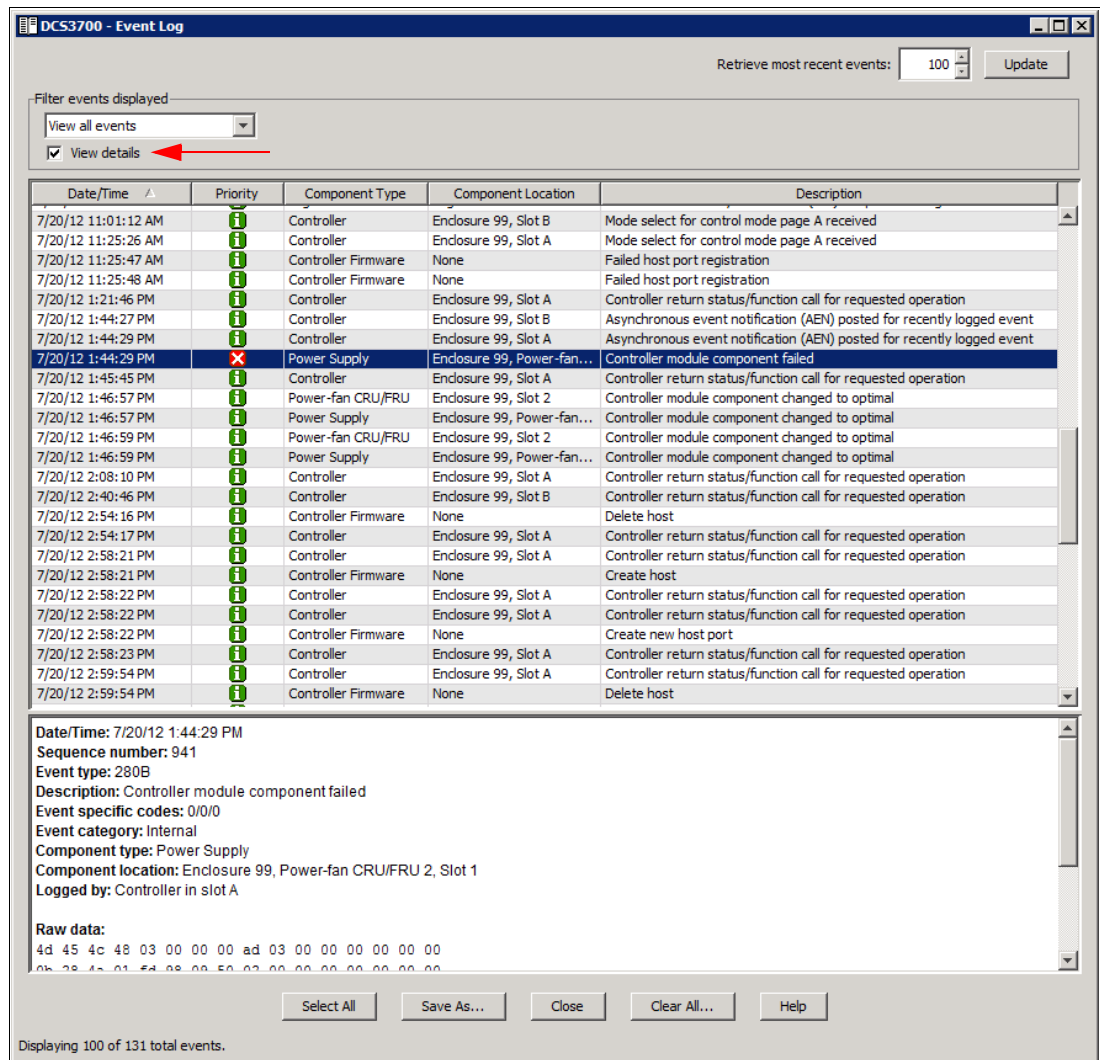


Figure 6-11 Detailed event view

## 6.2.7 Collect Support Data

Use the Collect Support Data link to gather various types of inventory, status, and performance data that can help troubleshoot any problems with your storage subsystem. All of the files that are gathered are compressed into a single archive in a compressed file format. Then, you can forward the archive file to your IBM Technical Support representative for troubleshooting and further analysis.

**Attention:** Use this option *only* under the guidance of your IBM Technical Support representative!

**Note:** If your enclosure contains drawers, the diagnostic data for that enclosure is archived in a separate compressed file named EnclosureComponentsStateCapture\_<user-defined>.zip in the same location that is specified for the support data file.



Figure 6-12 shows the collection of Support Data.

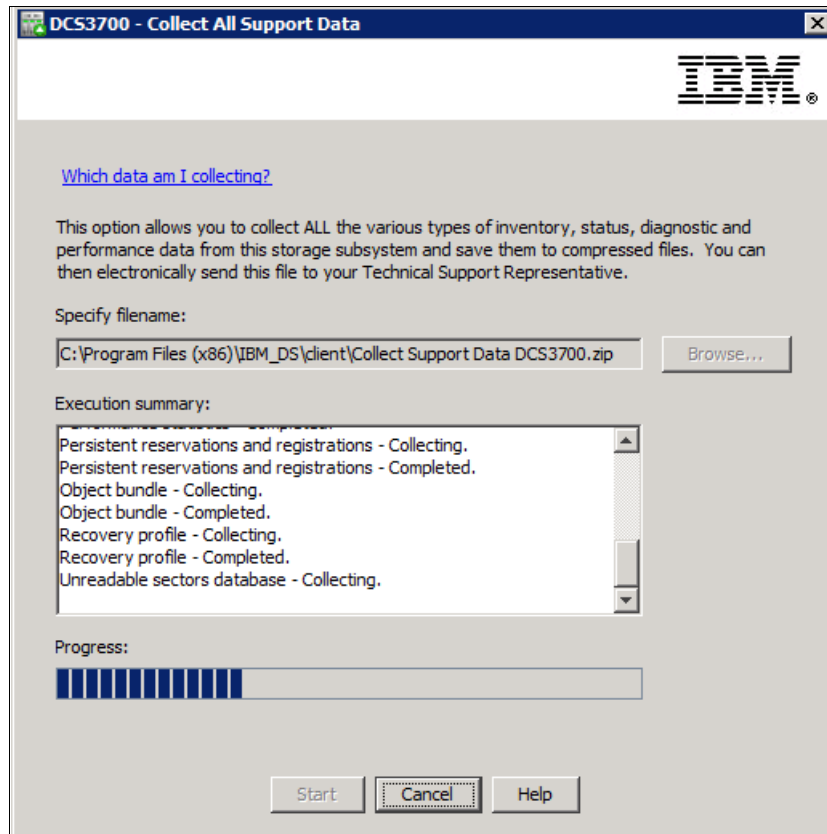


Figure 6-12 Collect Support Data

## 6.3 Capacity

The Capacity pane is at the lower left and is shown in Figure 6-13. It displays the overall capacity of your subsystem, that is, the total capacity of the installed disk drives. A darker green color is used to show configured capacity in logical volumes, and a light green color is used to show the capacity that can be used for the configuration of new logical drives. A gray color shows the space on unconfigured disks that are available on the system.

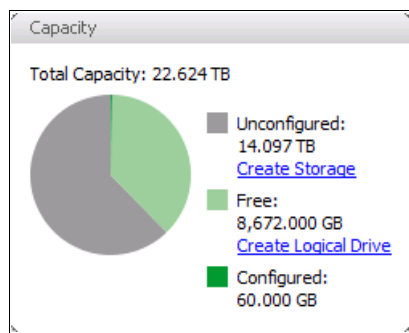


Figure 6-13 Capacity - configured, free, and unconfigured capacity

## 6.3.1 Create Storage

The Create Storage link allows for administration of the storage subsystem. Figure 6-14 shows the different possibilities. For more information about how to use the Automatic or Manual Configuration, see Chapter 8, “Administration: Storage & Copy Services tab” on page 225.

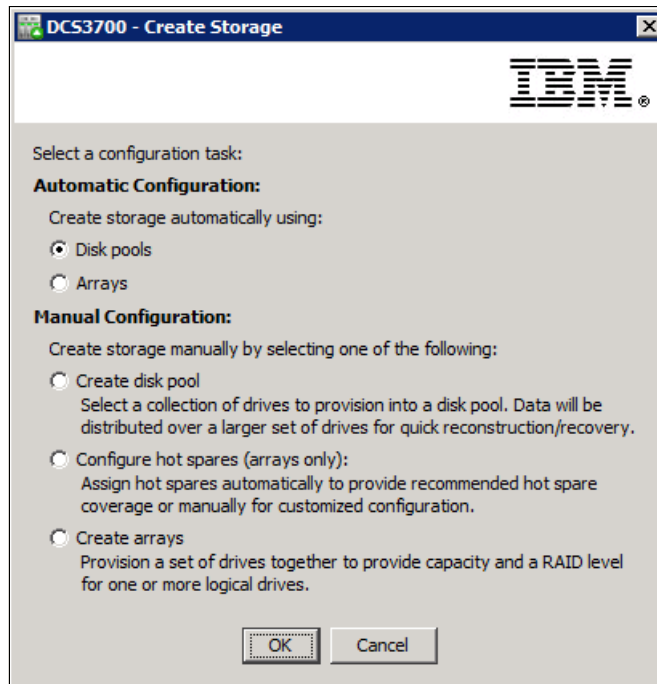


Figure 6-14 Create Storage

The following automatic configurations are possible:

- Create Disk Pools** A set of drives that is logically grouped. A disk pool provides the overall capacity that is needed to create one or more logical drives. A disk pool is similar to an array, with the following differences. The data in a disk pool is stored randomly on all of the drives in the disk pool, unlike data in an array, which is stored on the same set of drives. You do not specify a RAID level for a disk pool. A disk pool does not use hot spare drives. And, a disk pool allows a many drives to be grouped.
- Create Arrays** A set of drives that is logically grouped and assigned a RAID level. Each array that is created provides the overall capacity that is needed to create one or more logical drives.

The following manual configurations are possible:

- Create Disk Pools** A set of drives that is logically grouped. A disk pool provides the overall capacity that is needed to create one or more logical drives. A disk pool is similar to an array, with the following differences. The data in a disk pool is stored randomly on all of the drives in the disk pool, unlike data in an array, which is stored on the same set of drives. You do not specify a RAID level for a disk pool. A disk pool does not use hot spare drives. And, a disk pool allows many drives to be grouped.

### Configure hot spares (arrays only)

A spare drive that contains no data and that acts as a standby in case a drive fails in a Redundant Array of Independent Disks (RAID) 1, RAID 3, RAID 5, or RAID 6 logical drive. The hot spare drive can replace the failed drive in the logical drive. Hot spare drives are used only in arrays, not disk pools.

### Create arrays

A set of drives that is logically grouped and assigned a RAID level. Each array that is created provides the overall capacity that is needed to create one or more logical drives.

## 6.3.2 Create Logical Drive

In the Create Logical Drive window (see Figure 6-15), logical components are created for the host in order to access storage on the storage subsystem. A logical drive (LUN) is created from the capacity that is available on a disk pool or an array. Although a logical drive might consist of more than one drive, a logical drive appears as one logical component to the host.

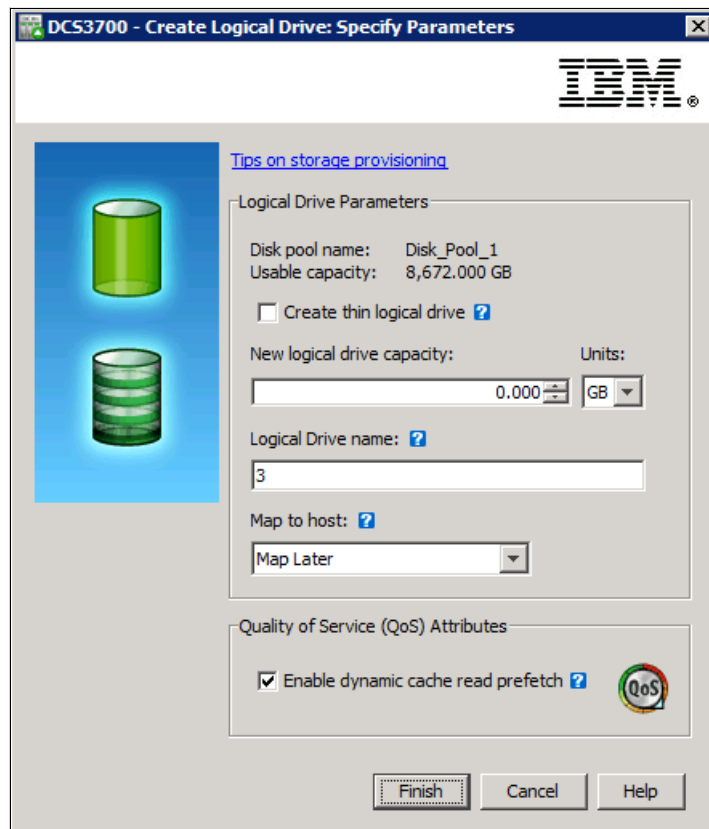
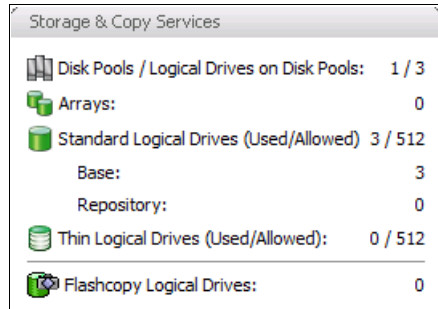


Figure 6-15 Create Logical Drive

## 6.4 Storage & Copy Services

The overview that is shown in Figure 6-16 provides information about the availability of Premium Features.



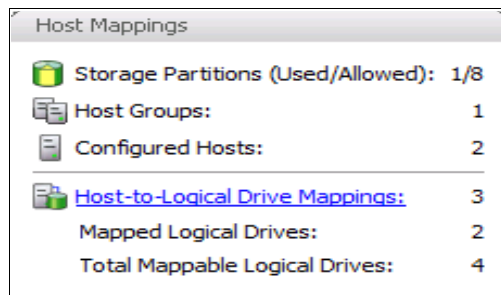
The screenshot shows a window titled "Storage & Copy Services" with the following data:

Disk Pools / Logical Drives on Disk Pools:	1 / 3
Arrays:	0
Standard Logical Drives (Used/Allowed)	3 / 512
Base:	3
Repository:	0
Thin Logical Drives (Used/Allowed):	0 / 512
Flashcopy Logical Drives:	0

Figure 6-16 Storage & Copy Services

## 6.5 Host Mappings

An overview of your Hosts & Mappings pane is shown in Figure 6-17.



The screenshot shows a window titled "Host Mappings" with the following data:

Storage Partitions (Used/Allowed):	1/8
Host Groups:	1
Configured Hosts:	2
<a href="#">Host-to-Logical Drive Mappings:</a>	3
Mapped Logical Drives:	2
Total Mappable Logical Drives:	4

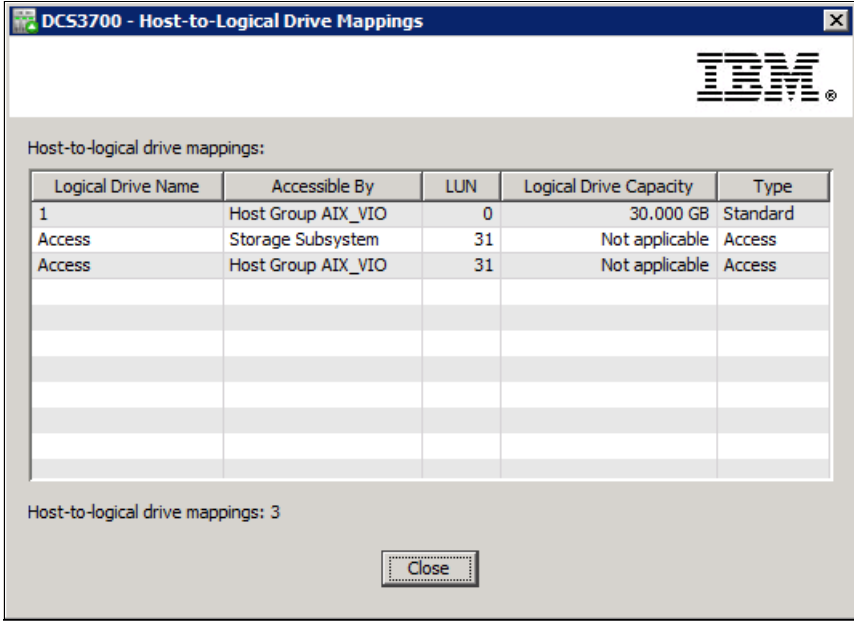
Figure 6-17 Host Mappings

Hosts are single server systems that can be mapped to a drive. This means that a host is a system that accesses the logical drive through the attached technology (SAS or Fibre Channel). If multiple hosts access one or a group of logical drives, they all 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 Chapter 10, "Administration: Mappings tab" on page 327.

## 6.5.1 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 mapped, which means the number of logical drives that are allowed to be seen by a specific host. Clicking the link displays a more detailed view (Figure 6-18).

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 10, “Administration: Mappings tab” on page 327.



Logical Drive Name	Accessible By	LUN	Logical Drive Capacity	Type
1	Host Group AIX_VIO	0	30.000 GB	Standard
Access	Storage Subsystem	31	Not applicable	Access
Access	Host Group AIX_VIO	31	Not applicable	Access

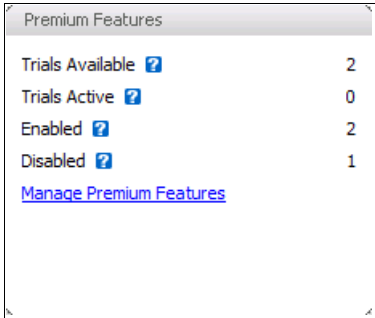
Host-to-logical drive mappings: 3

Close

Figure 6-18 Host-to-Logical Drive Mappings

## 6.6 Premium Features

This pane display the Number of Trial or activated Premium Features, as shown in Figure 6-19.



Trials Available	?	2
Trials Active	?	0
Enabled	?	2
Disabled	?	1

[Manage Premium Features](#)

Figure 6-19 Overview Premium Features

## 6.6.1 Manage Premium Features

To manage the Premium Features, click **Manage Premium Features**. The window that is shown in Figure 6-20 opens.

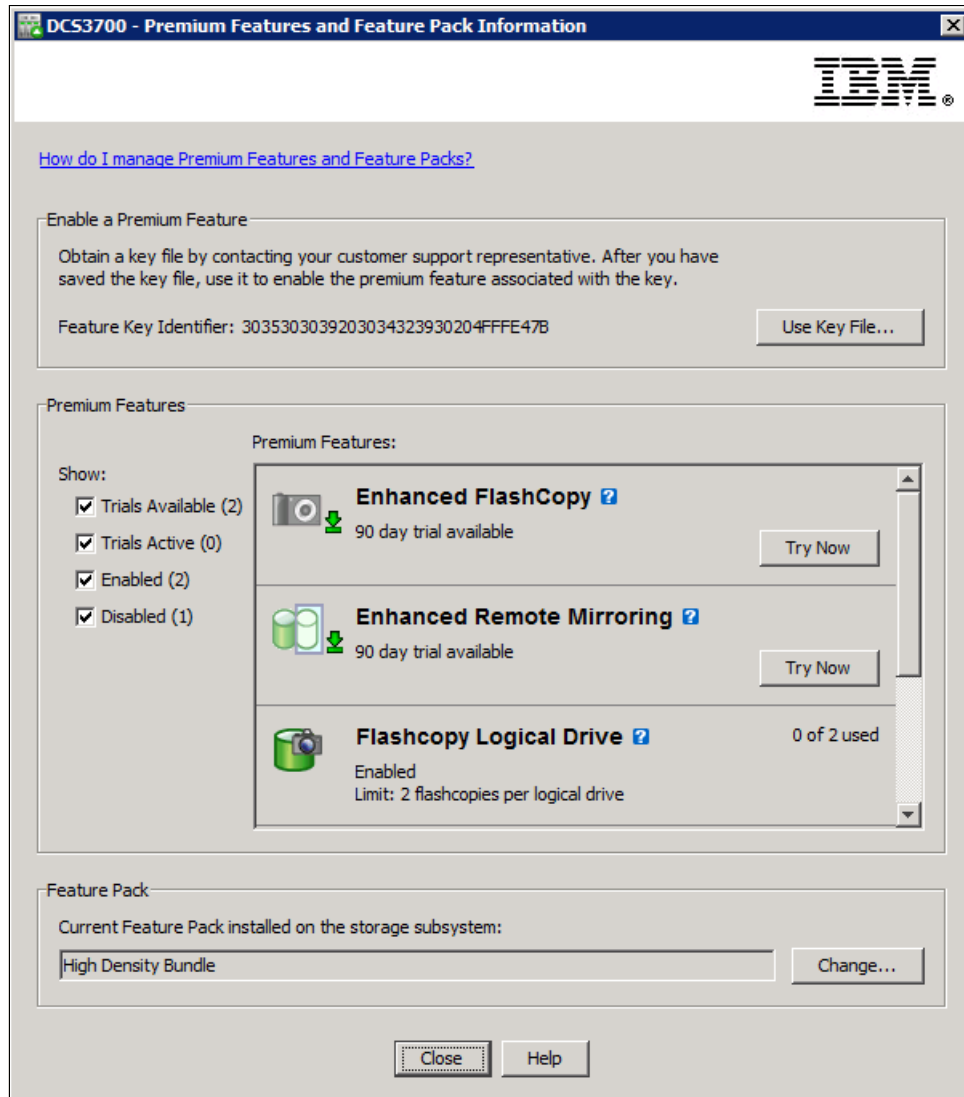


Figure 6-20 Manage Premium Features

More information about Premium Features is provided in 2.5, "Planning for Premium Features" on page 74 and "Premium Features" on page 205.

## 6.7 Hardware

The Hardware pane gives an overview of the Hardware Components (see Figure 6-21 on page 201) that are physically installed in this storage subsystem. The upper part of the pane shows the number of installed controllers and the number of connected enclosures (one in this case). The lower part of the pane shows information about installed hard disk drives (assigned and unassigned), media and number of drives, the drive type, and an overview of the hot spare drives.

Hardware	
Controllers:	2
Endosures:	1
<hr/>	
Drives:	20
Assigned:	14
Unassigned:	5
Media Type:	
HDD:	20
Interface Type:	
SAS:	20
Hot Spare Drives:	1
In-use:	0
Standby:	1

Figure 6-21 Hardware Overview

For more information about this pane, see Chapter 9, “Administration: Hardware tab” on page 305.

## 6.8 Information Center

At the right bottom edge of the Summary tab is the Information Center pane (Figure 6-22). It includes several links to help you understand and set up your storage subsystem, and answer any questions that you might have.

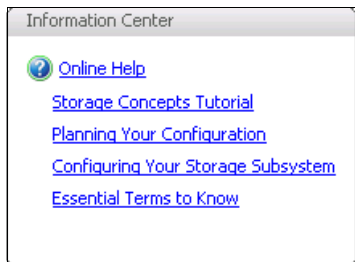


Figure 6-22 Information Center

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 6-23, topics about how to configure your DCS3700 storage subsystem are displayed if you click **Storage Concepts Tutorial**.

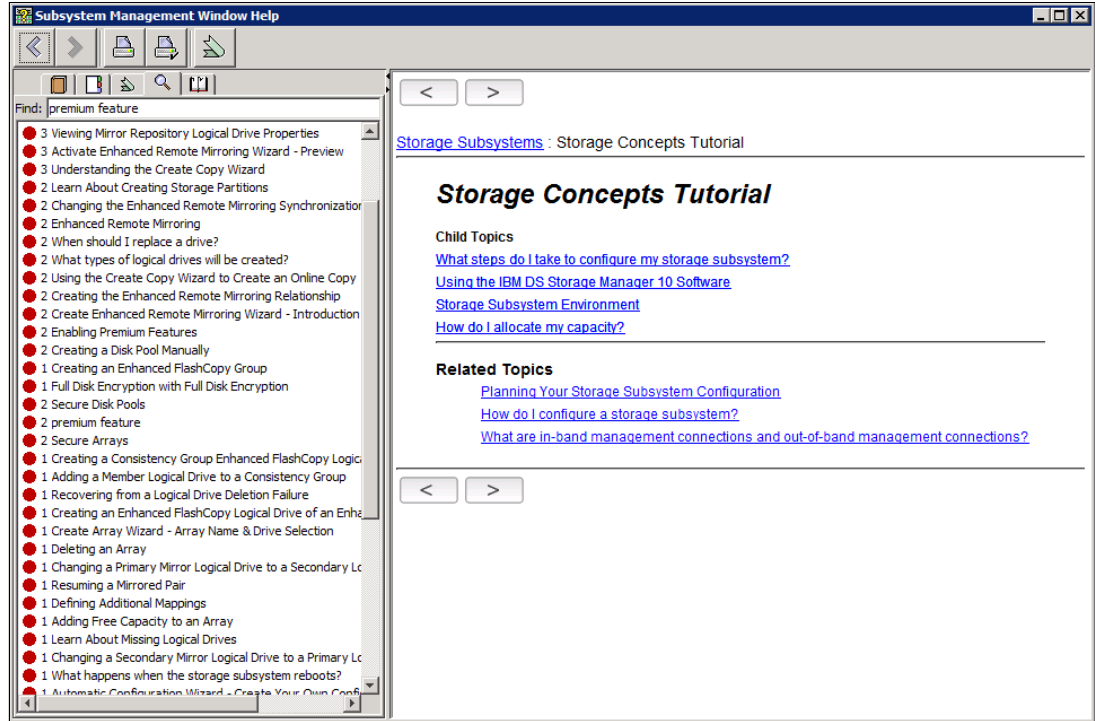


Figure 6-23 Information Center - Storage Concepts Tutorial





# Administration: Subsystem Management

Chapter 5, “Administration: Enterprise Management” on page 159 describes how to manage a storage subsystem from the IBM System Storage DS Storage Manager (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 window in the background and a small window that prompts for the password.

**Note:** Starting with DS Storage Manager V10.70, all managed subsystems must have a password set. Setting the password and a description of the new strong password rules are described in 11.1.4, “Set a Storage Subsystem Password” on page 356.

After entering the correct password, all five 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. For the Initial Setup Tasks for the first configuration steps on a new subsystem, see Chapter 11, “Administration: Setup tab” on page 353.

## 7.1 DS Storage Manager: Subsystem Management

The IBM System Storage DS Storage Manager Subsystem Management window has four main areas, as shown in Figure 7-1:

- ▶ Drop-down menu
- ▶ Toolbar
- ▶ Tabs
- ▶ Status bar

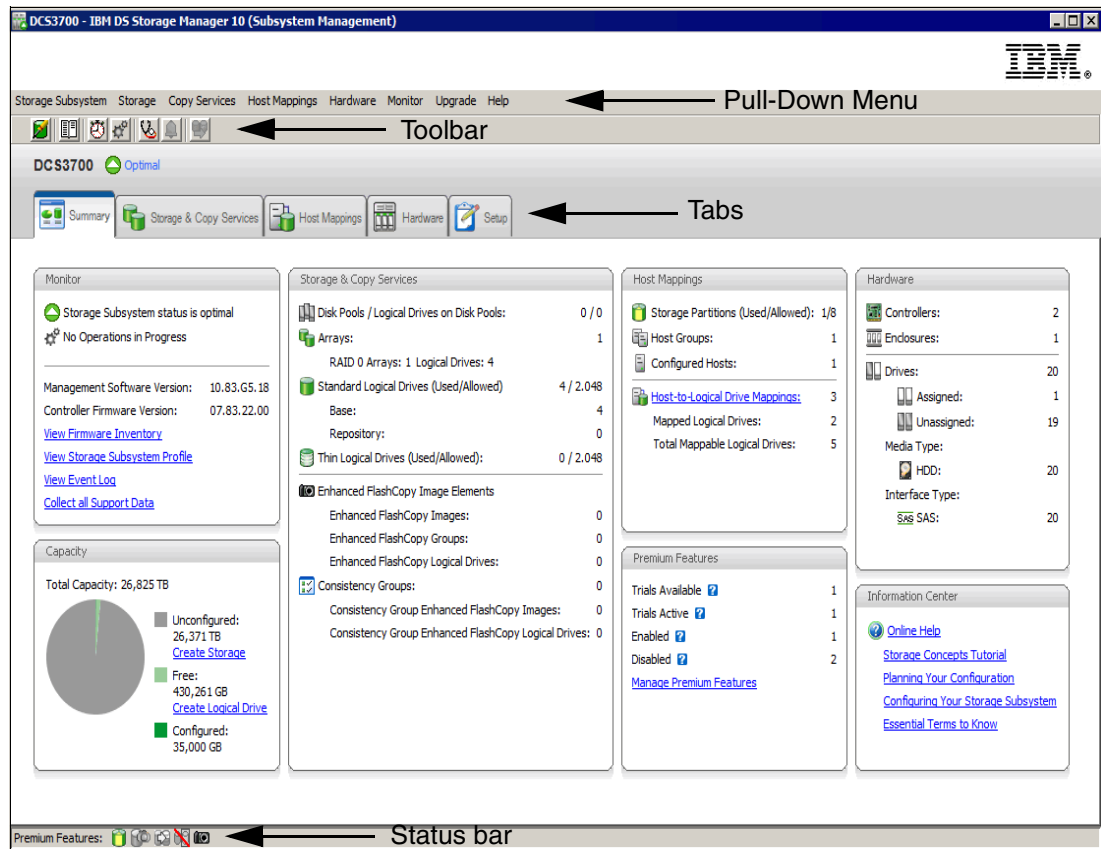


Figure 7-1 Subsystem Manager - Summary tab

## 7.2 Drop-down menu

The area at the top of Subsystem Management window is the drop-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 task. Another option is to configure the storage subsystem directly in the tab area. This section describes configuration tasks that are the easiest to do from the drop-down menu in this section. Other configuration tasks are described in the chapters that are dedicated to a specific tab.

This section covers the Storage Subsystem drop-down menu.

### 7.2.1 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 selected in the tab area.

Many items can be seen in this menu:

- ▶ Premium Features
- ▶ Security
- ▶ Change
- ▶ Configuration
- ▶ Rename
- ▶ Preferences
- ▶ Exit

Click **Storage Subsystem** to open the associated menu items, as shown in Figure 7-2. If you select items with the small right arrow on the right side of the menu item, it opens another submenu. The other items (with three dot characters at the end of menu text) directly open a window with the requested task.




Figure 7-2 Storage Subsystem menu

The following sections describe the most common items of this drop-down menu.

#### Premium Features

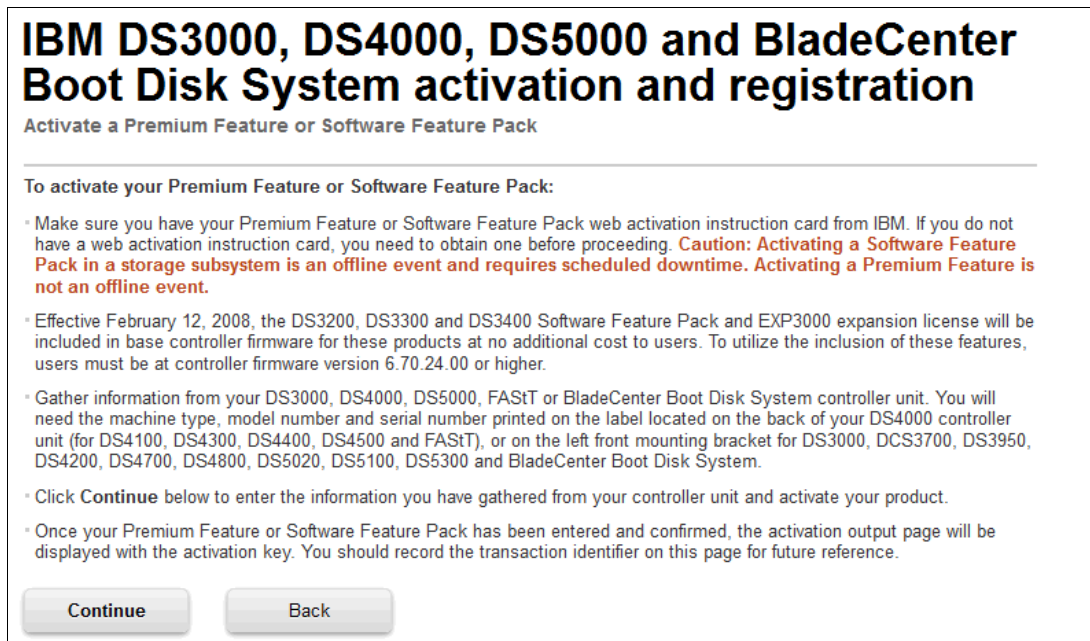
This item is used for viewing or activating software licenses for the Premium Features. You might have some of them already activated because certain basic licenses are delivered with each IBM System Storage DCS3700 storage subsystem as standard, and you can buy more as needed. Types of licenses are described in 2.5, “Planning for Premium Features” on page 74.

### Obtaining the license \*.key file

You need a special license \*.key file to activate more Premium Features in this step. This file must be downloaded from the IBM feature activation website (Figure 7-3), where each Premium Feature is registered based on the serial number of your DCS3700 enclosure and Feature Enable Identifier.

The IBM activation website is at the following URL:

<http://www.ibm.com/PremiumFeatures/jsp/keyPrereq.jsp>



**IBM DS3000, DS4000, DS5000 and BladeCenter Boot Disk System activation and registration**  
Activate a Premium Feature or Software Feature Pack

To activate your Premium Feature or Software Feature Pack:

- Make sure you have your Premium Feature or Software Feature Pack web activation instruction card from IBM. If you do not have a web activation instruction card, you need to obtain one before proceeding. **Caution: Activating a Software Feature Pack in a storage subsystem is an offline event and requires scheduled downtime. Activating a Premium Feature is not an offline event.**
- 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.
- Gather information from your DS3000, DS4000, DS5000, FASTT or BladeCenter Boot Disk System controller unit. You will need the machine type, model number and serial number printed on the label located on the back of your DS4000 controller unit (for DS4100, DS4300, DS4400, DS4500 and FASTT), or on the left front mounting bracket for DS3000, DCS3700, DS3950, DS4200, DS4700, DS4800, DS5020, DS5100, DS5300 and BladeCenter Boot Disk System.
- Click **Continue** below to enter the information you have gathered from your controller unit and activate your product.
- Once your Premium Feature or Software Feature Pack has been entered and confirmed, the activation output page will be displayed with the activation key. You should record the transaction identifier on this page for future reference.

Figure 7-3 Premium Feature activation website

The following Premium Features can be delivered with your DCS3700 storage subsystem or you can buy these Premium Features as optional (additional charge) Premium Features:

- ▶ Enhanced Remote Mirroring (ERM)
- ▶ Enhanced FlashCopy
- ▶ Storage Partitioning
- ▶ FlashCopy Logical Drive
- ▶ VolumeCopy

**Note:** In firmware V7.83 IBM offers a *Try and Buy* feature, where the Premium Features such as ERM or FlashCopy can be tested first, allowing you to test the functionality for 90 days before you purchase the Premium Feature.

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 DCS3700 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 website, it is difficult to get it again from IBM in a short time frame. You 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 is already registered and generated. You need only the serial number and the Feature Key Identifier of the DCS3700.

### View/Enable Premium Feature

Click **Storage Subsystem** → **Premium Feature** in drop-down menu to open a window where the list of possible features is displayed in the first pane, as shown in Figure 7-4. Some features are enabled at the plant, and some are disabled and can be enabled/activated by a \*.key file that must be downloaded from the IBM website (Figure 7-3 on page 206). The Feature Enable Identifier that is displayed is necessary (together with DCS3700 enclosure serial number) during the \*.key file generation at the IBM activation website.



Figure 7-4 Premium Feature window

Depending on the Premium Feature that is selected, you can enable a disabled feature, or disable an enabled feature. To enable a feature, click **Enable** and a window for selecting the \*.key file on your management station opens. When the key file is accepted, a window with more information opens (Figure 7-5).

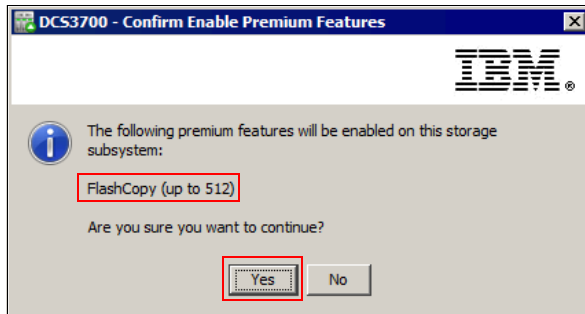


Figure 7-5 Enable Premium Feature - information

To disable a feature, select that feature and click **Disable**. A warning window opens (Figure 7-6).

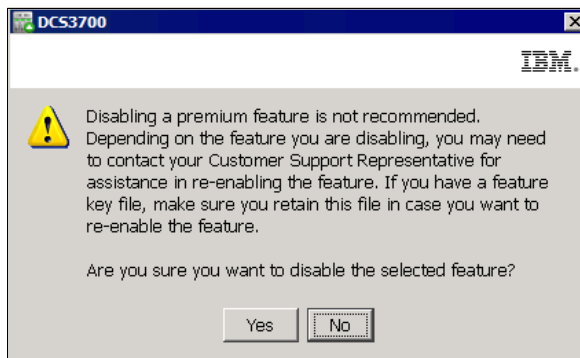


Figure 7-6 Disable Premium Feature

### Feature Pack

The second pane in the Premium Feature window (Figure 7-4 on page 207) is the Feature Pack window. It is not possible to change the Feature Pack on any DCS3700 storage subsystem.

### Security: Set Password

Starting with Storage Manager V10.70, new stronger password rules are in effect, so you must have passwords set for all DCS3700 storage subsystems. You must enter the password when Subsystem Manager is started for a storage subsystem, even if you just want to see the configuration. If you start Subsystem Manager for the first time, you must set a new password. In previous versions, the password was requested only in the case of configuration changes, such as deleting objects and creating a new array or logical volume.

A strong password must meet the following requirements:

- ▶ Must be 8 - 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 (!, @, #, \$, and so on).
- ▶ Must not contain any spaces, including trailing spaces.

Click **Storage Subsystem** → **Security** → **Set Password** to define a new password. You must set the current password and enter the new password twice.



Figure 7-7 Set Password

If you have forgotten the current password, contact your IBM Customer Support Representative. If a password was not set previously (for example, you have a new subsystem), the current password is not required to establish a new password.

## Change

The Change menu item in the Storage Subsystem menu is one of the most important menus for setting global parameters of the DCS3700 storage subsystem. Many changes that are done by this menu can also be done directly in tabs, but this section describes most of these changes.

Click **Storage Subsystem** → **Change**. The menu that opens (Figure 7-8) offers several options to configure:

- ▶ Cache Settings
- ▶ Failover Alert Delay

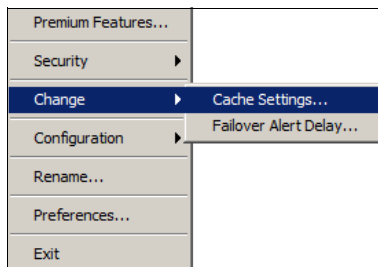


Figure 7-8 Change subsystem menu

## Cache Settings

This menu item sets the global cache settings of the DCS3700 storage subsystem. It is possible to set other cache settings that are dedicated for each logical drive. The recommended cache settings depend on the required host load and type of load, and can vary.

Click **Storage Subsystem** → **Change** → **Cache Settings**. The window that is shown in Figure 7-9 opens.

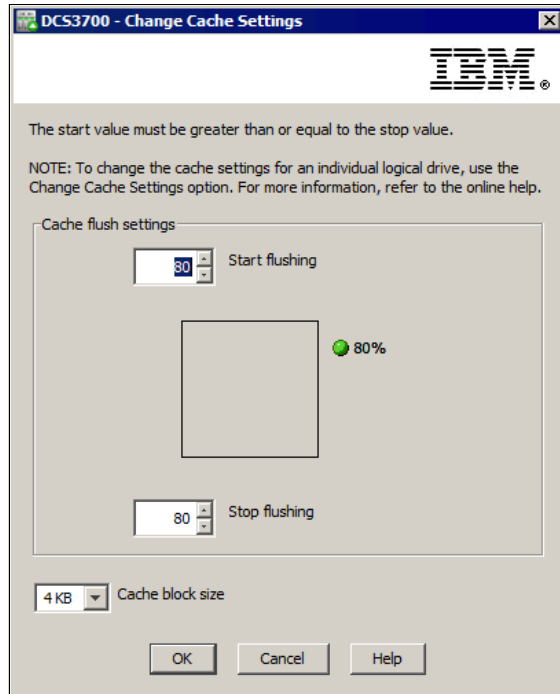


Figure 7-9 Cache Settings for a storage subsystem

The default value of the cache flushing is 80%/80%. You have a choice to set the Start and Stop level for cache flushing, and the Cache block size, as shown in Figure 7-9. For most 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 you click **OK**; you do not need to reset the DCS3700 storage subsystem, so it is possible to test your settings during regular production on the storage subsystem.

### ***Failover Alert Delay***

A failover alert delay can be specified on a DCS3700 storage subsystem. The setting can delay the logging of a critical event if the multipath failover driver transfers the logical drive to the non-preferred controller of the storage 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 that is shown in Figure 7-10 on page 211, you can set a range of 0 - 60 minutes. The default failover delay is 5 minutes.



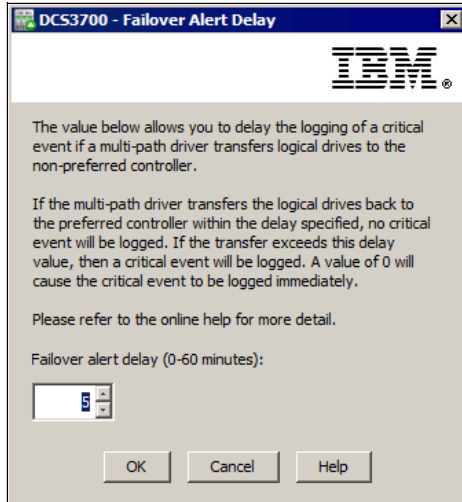


Figure 7-10 Failover Alert Delay

**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 it. For more information about alert sending, see 13.2, “View Health (Recovery Guru)” on page 387.

## iSCSI

This option (Figure 7-11) in the Storage Subsystem drop-down menu is available and visible only if iSCSI daughter cards are installed in the DCS3700 storage subsystem.



Figure 7-11 iSCSI drop-down menu

For a detailed description about how to set an attachment to iSCSI hosts, see Chapter 12, “Administration: iSCSI” on page 363.

## Automatic configuration of arrays

If the Automatic configuration option is selected, a new window with an introduction to automatic configuration opens that uses dedicated options to create tuned logical drives. Click **Next** and the window that is shown in Figure 7-12 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 open.

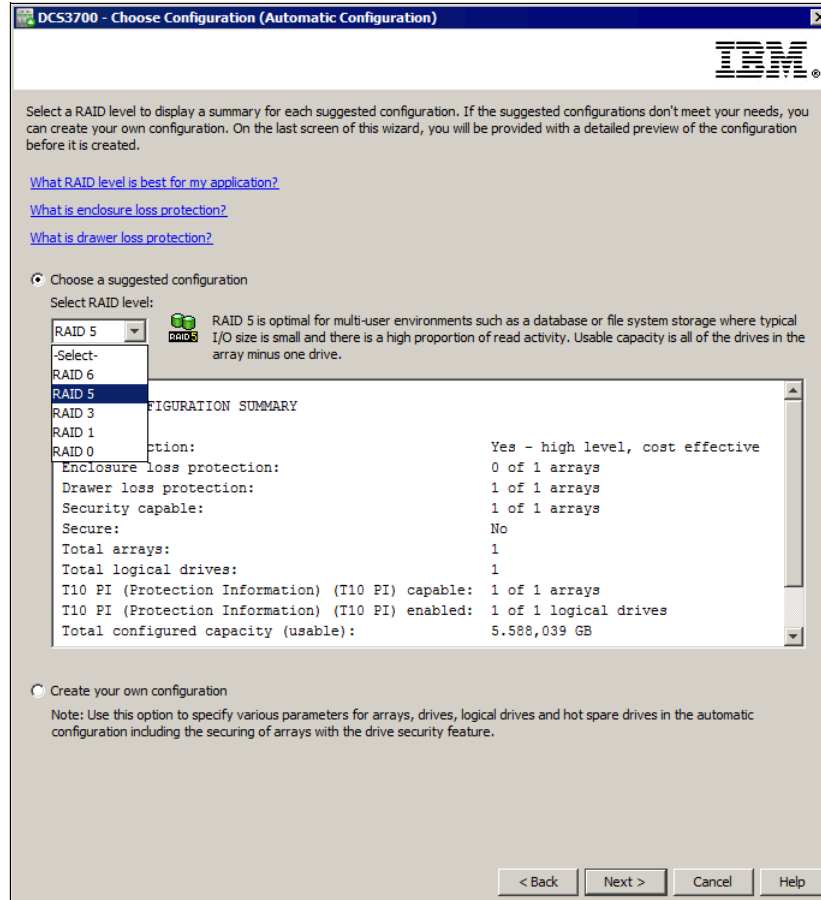


Figure 7-12 Fully automatic configuration of arrays

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 setting only the RAID level, you have no control of the number and size of logical disks and number of arrays that are created by 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 for automatic configuration.

You can preset the following items, as shown in Figure 7-13 on page 213:

- ▶ RAID level
- ▶ Number of disk drives per array
- ▶ Number of Hot spare drives
- ▶ Number of logical drives per array
- ▶ Change I/O Type

If you select Change I/O Type, you can preset the following items:

- ▶ 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 selecting disk drives from separate enclosures is impossible.

Set your required values in the wizard window that is shown in Figure 7-13 and click **Next**.

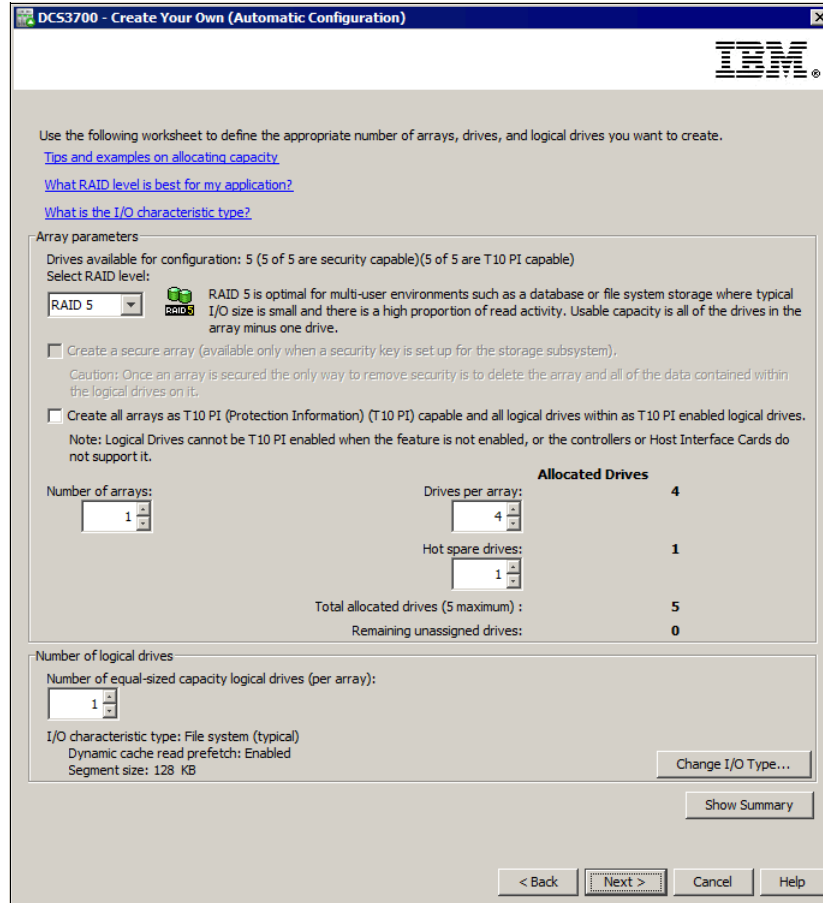


Figure 7-13 Automatic Configuration Parameters

A Summary window and a window with more information about the automatic configuration process (Figure 7-14) opens.

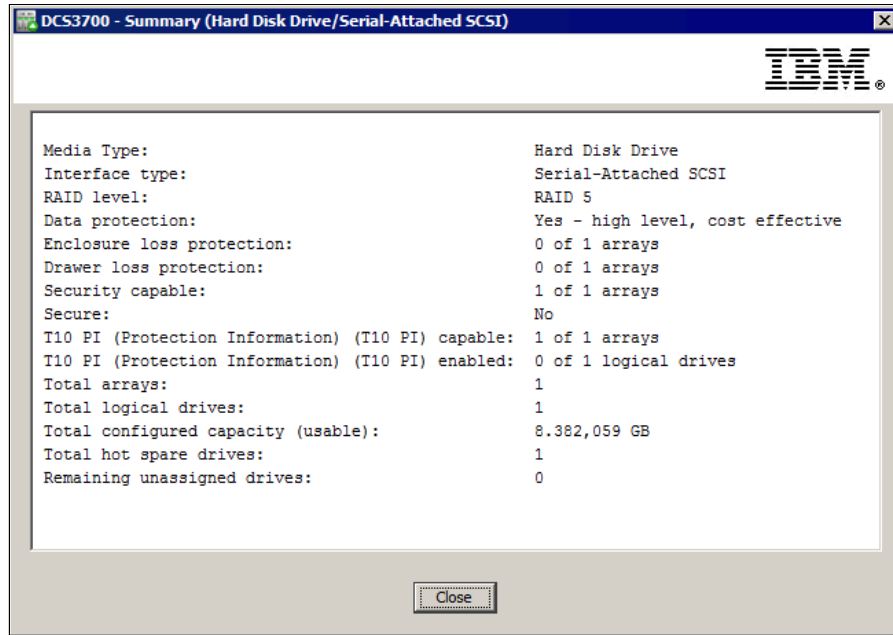


Figure 7-14 Automatic Configuration - Summary

## Dynamic Pool Automatic Configuration

The storage management software can detect the unconfigured capacity in a storage subsystem. When the unconfigured capacity is detected, the storage management software prompts you to create one or more disk pools, add the unconfigured capacity to an existing disk pool, or both. By default, the Automatic Configuration dialog box opens when one of these conditions are true (Figure 7-15 on page 215):

- ▶ The Subsystem Management window (SMW) is opened to manage a storage subsystem, disk pools do not exist in the storage subsystem, and there are enough similar drives to create a disk pool.
- ▶ New drives are added to a storage subsystem that has at least one disk pool. If there are enough eligible drives available, you can create a disk pool of different drive types than the existing disk pool.

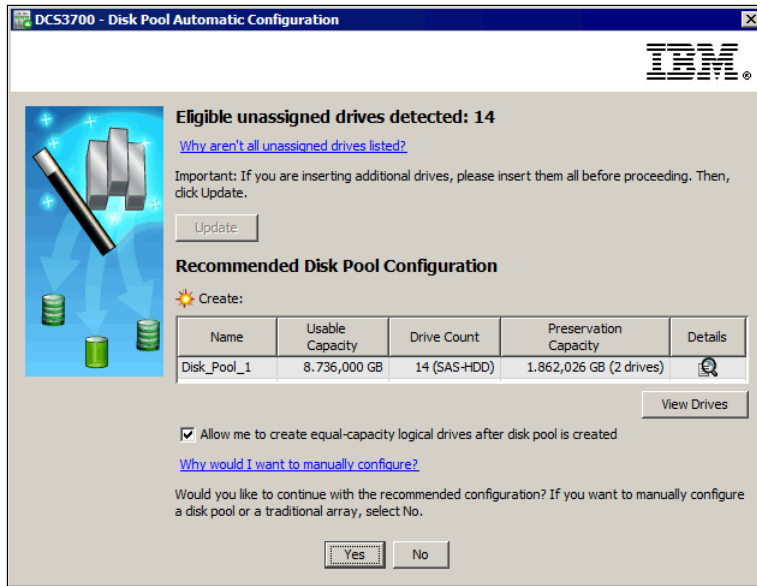


Figure 7-15 Automatic Configuration - Disk Pools

If you want to open the Automatic Configuration dialog box, click **Storage Subsystem** → **Configuration** → **Automatic** → **Disk Pools** (Figure 7-16).

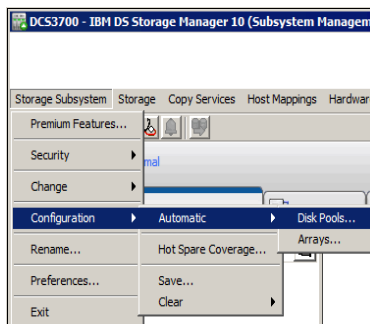


Figure 7-16 Invoke Automatic Configuration - Disk Pools

You can preset the following items (shown in Figure 7-15):

- ▶ Name: The default name that is given to the disk pool.
- ▶ Usable Capacity: The capacity of the disk pool that can be used to create logical drives.
- ▶ Drive Count: The number of drives that comprise the disk pool and their drive type and media type.
- ▶ Preservation Capacity: The disk pool capacity that is reserved for reconstruction of data when one or more drives fail.
- ▶ Details: An icon is displayed in this column. When you hover the mouse pointer over this icon, you can view more information about the disk pool recommendation, such as drive speed, security capability, and T10 PI (Protection Information) capability.

As shown in Figure 7-17, you can define the following items:

- ▶ Name
- ▶ Usable Capacity
- ▶ Free Capacity (% of Usable)
- ▶ Equal Capacity Logical Drives

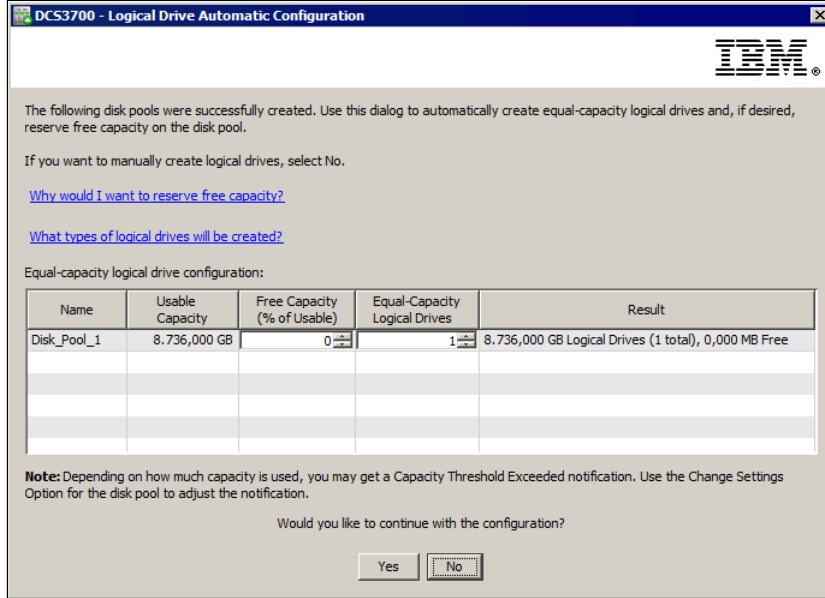


Figure 7-17 Logical Drive Automatic Configuration

### Hot spare coverage

If you select the **Configure hot spare drives** option in the Select Configuration Task window, a new window opens for hot spare configuration. In this window (Figure 7-18 on page 217), you have two choices: Automatic or manual configuration of hot spares.

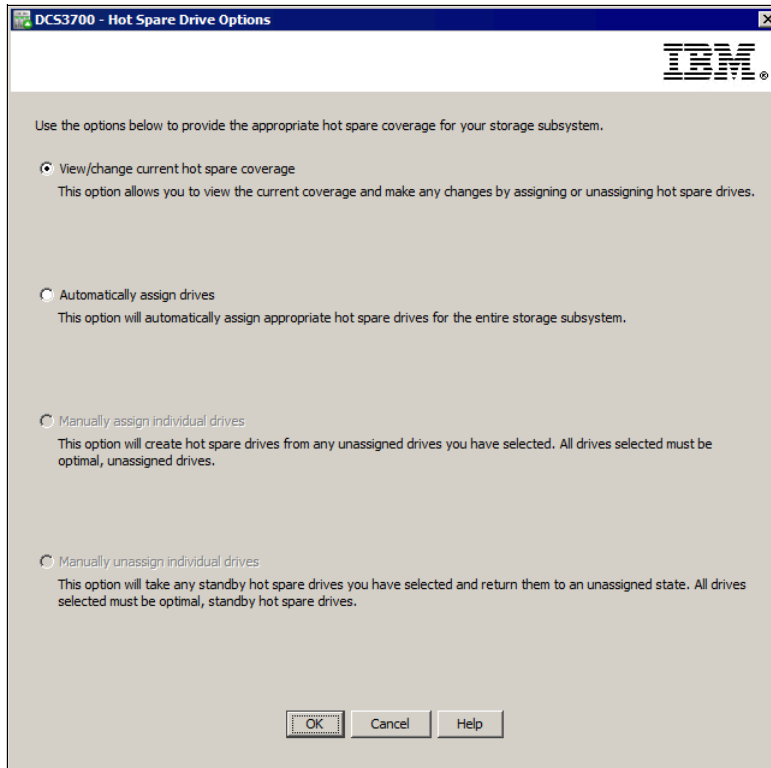


Figure 7-18 Hot Spare Coverage

The concept of a hot spare drive is explained in Chapter 9, “Administration: Hardware tab” on page 305. There you can also find a detailed description of the same Hot Spare Drive Option window, and how to set hot spare drives.

## Save Configuration

Click the **Save Configuration** link to open a window (Figure 7-19) 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, see “Load Storage Subsystem Configuration option” on page 165.

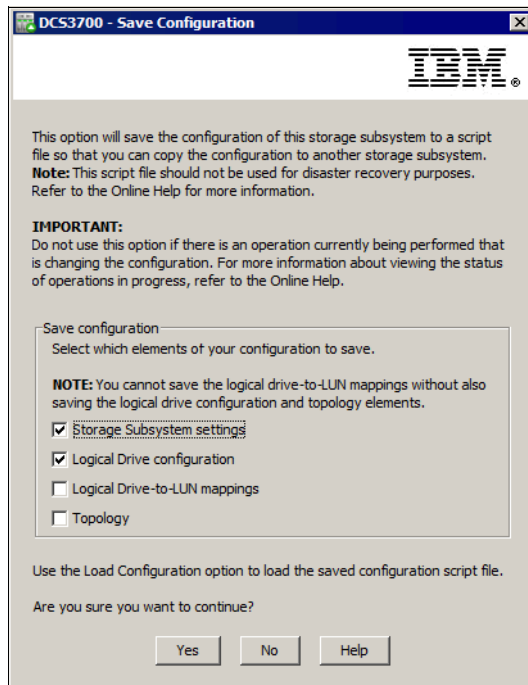


Figure 7-19 Save Configuration

When your DCS3700 Storage Subsystem is fully configured, the Save Configuration option lets you save storage subsystem logical settings, logical drive configuration, topology, and mappings to a text file in the form of script file that also can be used by the 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 logical device in your storage subsystem. For more information about SMcli, see Chapter 15, “Command-line interface (CLI)” on page 463.

**Save Configuration:** It is a preferred practice to save the Configuration after every change you make to the system. IBM Support can re-create the system after a disaster only if it has a valid configuration file.

Clicking **Configuration** → **Save** does not save the following settings:

- ▶ The age of the battery
- ▶ The controller time of day
- ▶ The nonvolatile static random access memory (NVSRAM) 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



## Clearing a configuration

Clicking **Configuration** → **Clear** in the Subsystem Management window provides the submenu options that are listed in the following sections.

### **Clear Storage System Configuration**

Click **Clear Storage System Configuration** to permanently delete the current logical drive configuration. Using this function cause these conditions to occur:

- ▶ All logical drives and arrays are deleted. This action destroys all data on these logical drives and arrays.
- ▶ The user-supplied DCS3700 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 that is 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.

To clear the storage subsystem configuration, complete the following steps:

1. In the DS Storage Manager (Subsystem Management) window, click **Configuration** → **Clear** → **Storage Subsystem**.

You are prompted to confirm that you want to proceed, as shown in Figure 7-20.

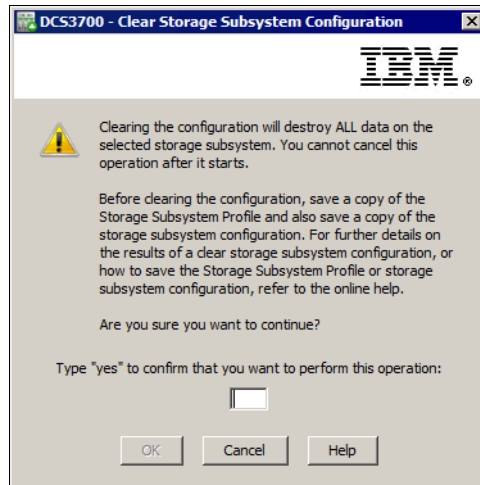


Figure 7-20 Clear Storage Subsystem Configuration

2. Type yes in the text box and click **OK** to clear the storage subsystem configuration.

### **Clear Array Configuration**

You can use the Logical Drive option to permanently delete the Logical Drive configuration on the DCS3700 storage subsystem.

**Attention:** All data that is 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.

To clear the storage subsystem array configuration, complete the following steps:

1. In the DS Storage Manager (Subsystem Management) window, click **Configuration** → **Clear** → **Logical Drive**.

You are prompted to confirm that you want to clear the Logical Drive configuration, as shown in Figure 7-21.

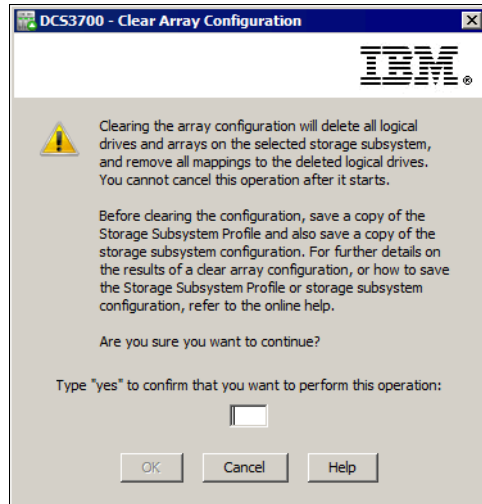


Figure 7-21 Clear Array Configuration

2. Type yes in the text box and click **OK** to clear the Logical Drive (array) configuration.

## Rename Storage Subsystem

When you add a storage subsystem to your Storage Manager, it is identified as “Unnamed”. If you have several storage subsystems that are installed in your environment, you can give each subsystem a meaningful name.

To change the name of subsystem, click **Storage Subsystem** → **Rename**, as shown in Figure 7-22.

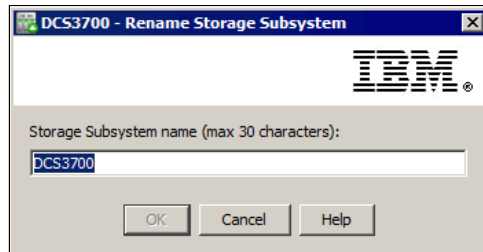


Figure 7-22 Rename Storage Subsystem

The subsystem name is immediately updated in Enterprise Manager.

**Note:** Renaming the storage subsystem can have an influence on services that are running outside the Storage Manager. For example, you can run some customized SMcli scripts where the storage subsystem is referenced by its name. Also, some types of multipath failover drivers on servers use the name of the attached storage subsystem. Rename a new storage subsystem as soon as possible before it is in production.

## Preferences

To change the User Interface Settings or to Reset Conditional Dialogs, click **Storage Subsystem** → **Preferences**, as shown in Figure 7-23.

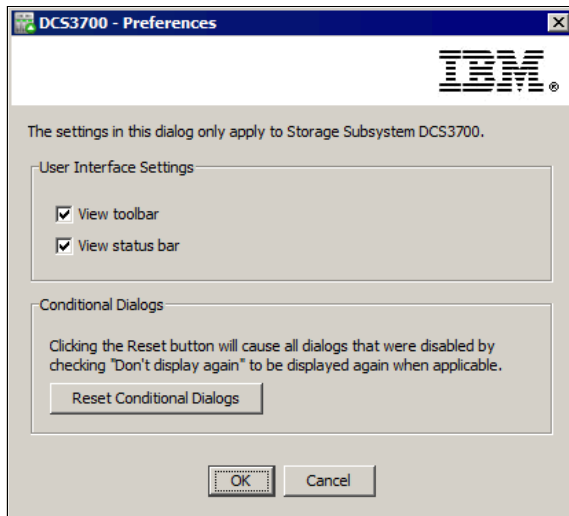


Figure 7-23 Preferences

If you choose **Reset Conditional Dialogs**, a confirmation window opens, as shown in Figure 7-24.

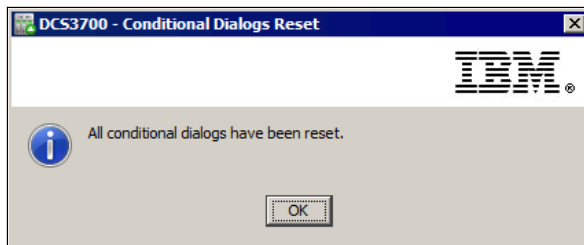


Figure 7-24 Conditional Dialogs Reset confirmation

## 7.3 Toolbar

Several icons are in the toolbar area and shown in Figure 7-25. These icons can help you easily navigate to certain management tasks.

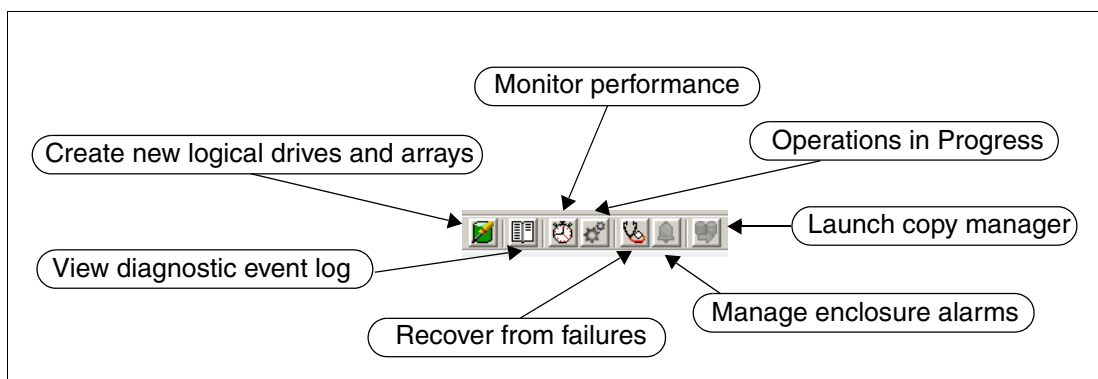


Figure 7-25 Toolbar icons

The seven icons in the toolbar are (from left to right):

- ▶ Create new logical drives and arrays
- ▶ View diagnostic event log
- ▶ Monitor Performance
- ▶ Operations in Progress
- ▶ Recover from failures
- ▶ Manage enclosure alarms
- ▶ Launch copy manager

## 7.4 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 7-26, this Premium Feature is not activated.

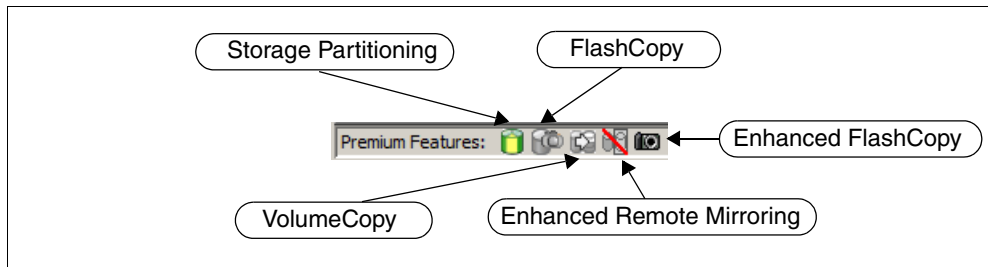


Figure 7-26 Status bar icons

If you hover your cursor over one of the icons, a small text box (Figure 7-27) is shown after about one second. It shows the description of the selected Premium Feature and its enablement status.



Figure 7-27 Premium Feature icon - description text

To activate new licenses (Premium Features), click **Storage Subsystem** → **Premium Features**, as described in “View/Enable Premium Feature” on page 207.

## 7.5 Tabs

This is the main management area of the IBM System Storage DS (Subsystem Management). In this area (see Figure 7-28 on page 223), you can see the highlighted name of the managed storage subsystem in the upper part, and close to the name is an icon that shows the 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.

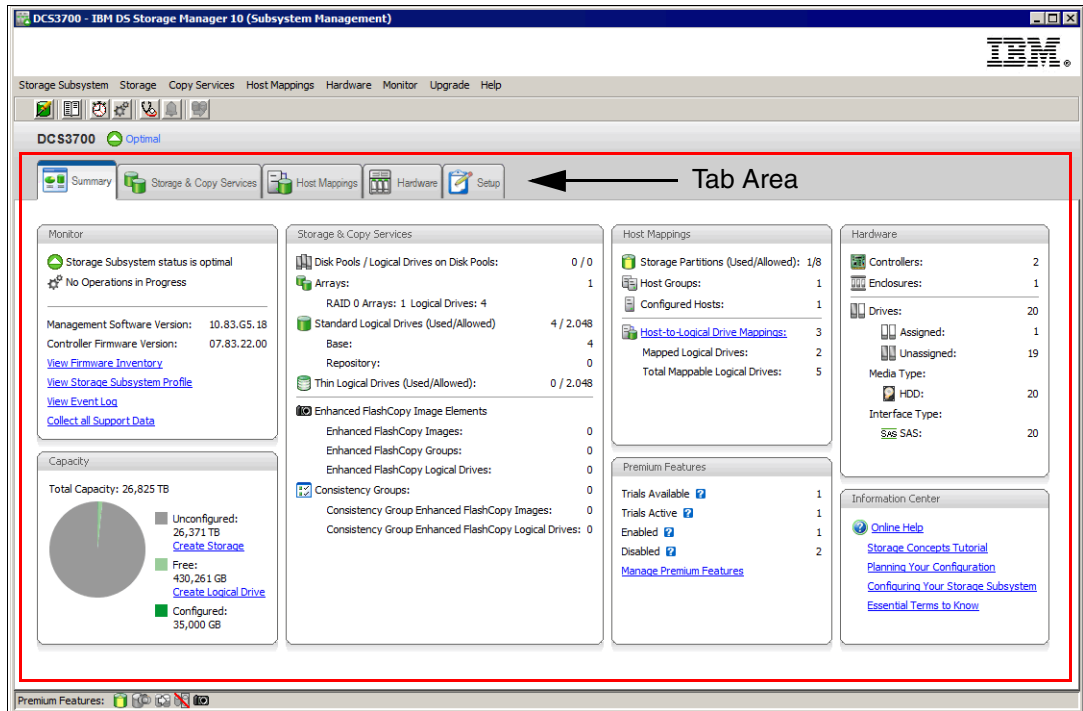


Figure 7-28 Tab area

Each of these five tabs is described in the following chapters in this book:

- ▶ Summary tab in Chapter 6, “Administration: Summary tab” on page 187
- ▶ Storage & Copy Services tab in Chapter 8, “Administration: Storage & Copy Services tab” on page 225
- ▶ Host Mappings tab in Chapter 10, “Administration: Mappings tab” on page 327
- ▶ Hardware tab in Chapter 9, “Administration: Hardware tab” on page 305
- ▶ Setup tab in Chapter 11, “Administration: Setup tab” on page 353





## Administration: Storage & Copy Services tab

This chapter describes the Storage & Copy Services tab of the DS Storage Manager. In this tab, you can logically configure an IBM System Storage DCS3700 storage subsystem. These configuration options include the following tasks:

- ▶ Creation of pools, arrays, and logical drives
- ▶ Disk pool technology
- ▶ Differences between a disk pool and an array
- ▶ Creation of FlashCopy, VolumeCopy, and Remote Mirror drives
- ▶ Display of the properties of these logical objects
- ▶ Usage of dynamic functions
- ▶ Change of controller ownership
- ▶ Change of modification priority, cache settings, and media scanning
- ▶ Consistency groups

## 8.1 Storage & Copy Services tab

The Storage & Copy Services tab (Figure 8-1) shows the organization of the storage subsystem by logical drives, pools, 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 13.10, “Storage Subsystem Profile” on page 407.

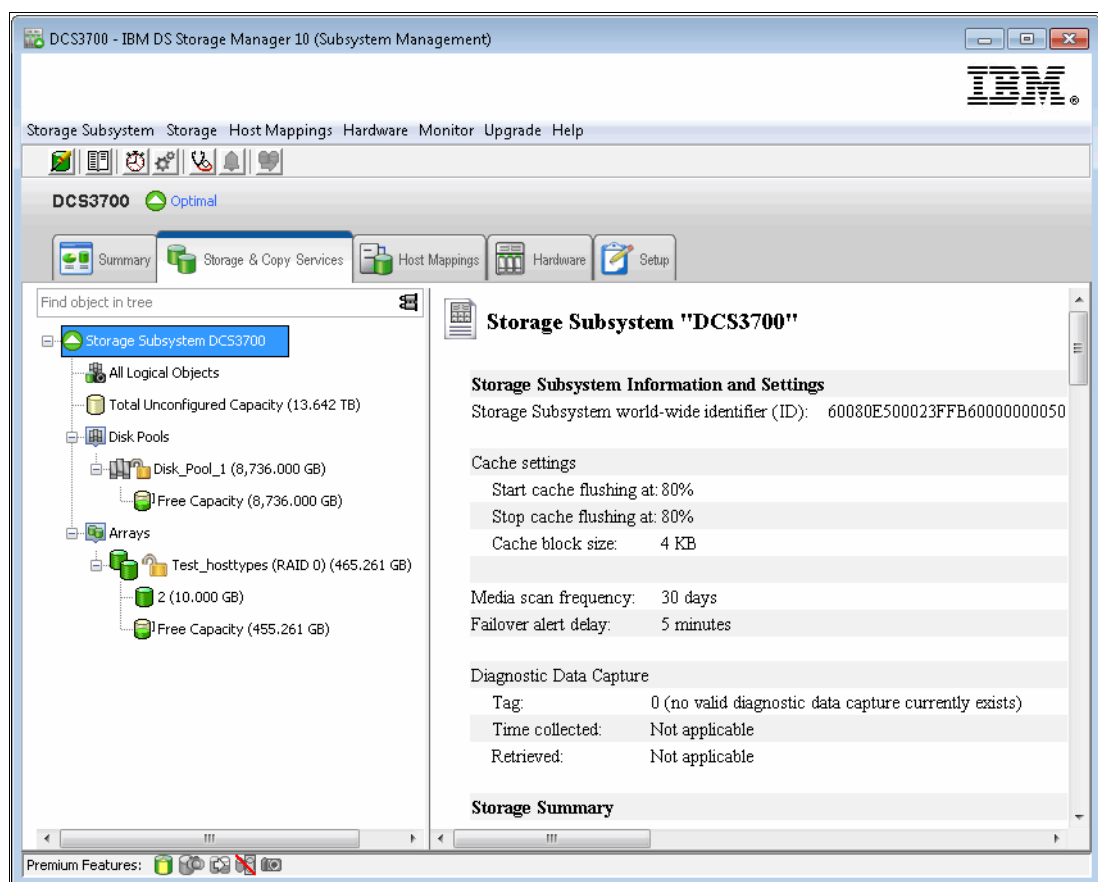


Figure 8-1 Service & Copy Services

The first line in the Properties pane is the View Associated Components link. When selected, it opens a window with the physical components that belong to the selected logical object in the Logical pane. The associated components are marked by blue bullets in this window.

If the first line of the Logical pane (Storage Subsystem) is selected, the link opens a window with the 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 that is selected in the Logical pane. All configuration options are shown as a menu when you right-click a particular object.

This chapter is split in to four basic sections to explain the following configuration options:

- ▶ Working with unconfigured capacity
- ▶ Working with arrays
- ▶ Creating a disk pool
- ▶ Working with a disk pool
- ▶ Working with free capacity in arrays
- ▶ Working with free capacity in a disk pool
- ▶ Working with logical drives
- ▶ Consistency groups

The first line in the Logical pane (the root of the logical directory tree) is Storage Subsystem (already selected in Figure 8-1 on page 226). The name of storage subsystem and basic status (green - Optimal, or red - Needs Attention) of this managed subsystem is shown there. In the Properties pane, you can see overall information about the whole DCS3700.

Right-click the storage subsystem in the logical pane and a full list of possible configurations opens, as shown in Figure 8-2. The first menu item is Locate. If you hover your cursor over it, new options appear: Storage Subsystem, which helps you locate your storage in a server room, and Stop All Indications opens, which you can use when the locate operation has not closed correctly and the locate LEDs continue to flash. All other menu items are the same as when you click the **Storage Subsystem** menu in the drop-down menu bar of DS Storage Manager. This menu is described in 7.2.1, “Storage Subsystem menu” on page 205.

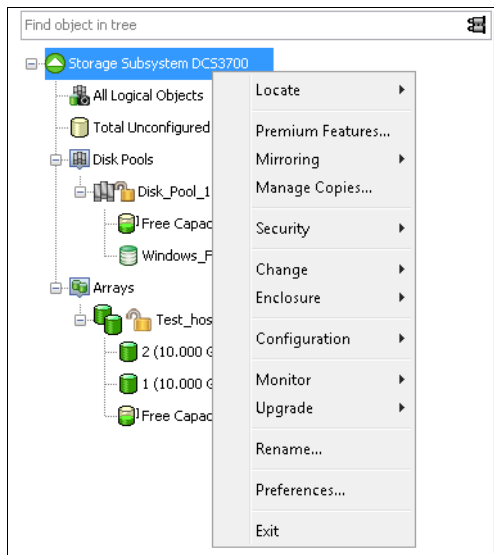


Figure 8-2 Storage Subsystem menu

There are several other objects in the Logical pane:

- ▶ All Logical Objects
- ▶ Total Unconfigured capacity
- ▶ Disk Pools
- ▶ Array
- ▶ Free capacity
- ▶ Standard Logical Drive
- ▶ Special Logical Drive (FlashCopy, FlashCopy Repository, Mirror Repository, and so on)

The following sections describe most of these objects. Others are described in chapters or sections that are dedicated to specific topics.

## 8.2 Working with unconfigured capacity

If you get a new IBM System Storage DCS3700, it probably has all unconfigured capacity. As you start to configure unassigned physical disk drives into arrays or disk pools, this unconfigured capacity decreases. When the whole storage subsystem is configured and no disks remains, this object disappears from the Logical pane.

The main task that can be completed on the unconfigured capacity is to create an array or disk pool. In Figure 8-3, you see an example of a partially configured DCS3700, where about 13.642 TB unconfigured capacity remains.

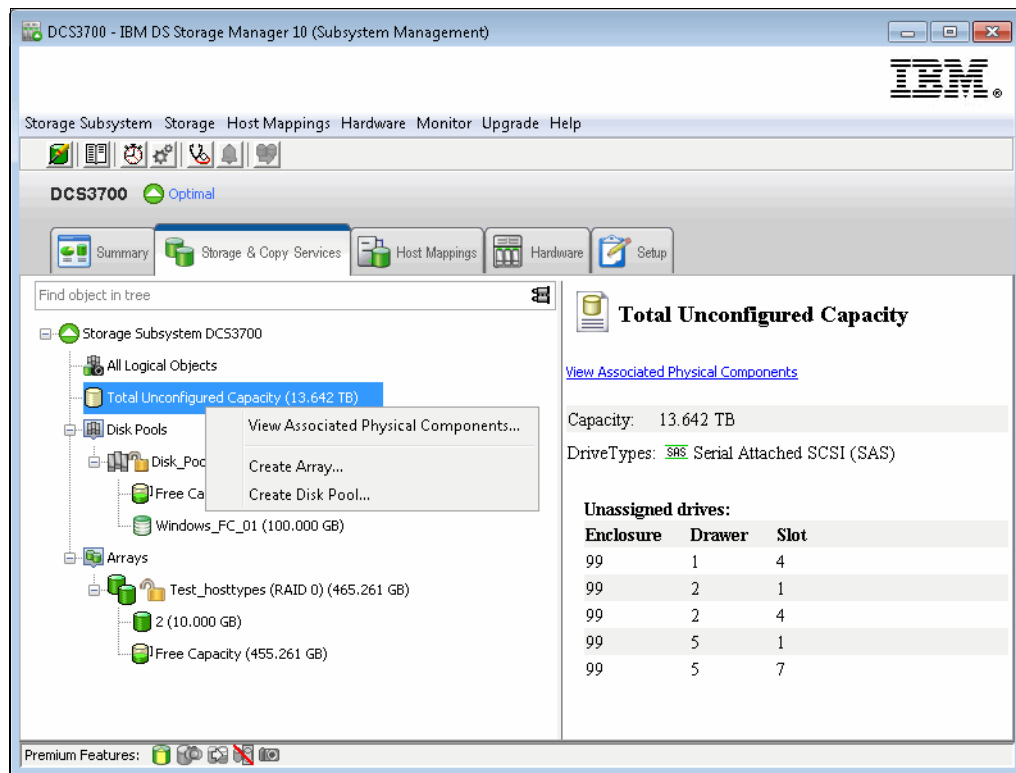


Figure 8-3 Unconfigured capacity example

### 8.2.1 View Associated Physical Components

Select the unconfigured capacity and in the Properties pane you can see a list of associated disk drives that are not configured yet. You can click the **View Associated Physical Components** link to open a window where the unconfigured disks are physically placed (see Figure 8-4 on page 229).

The same window opens if you right-click the unconfigured capacity and select **View Associated Physical Components**.

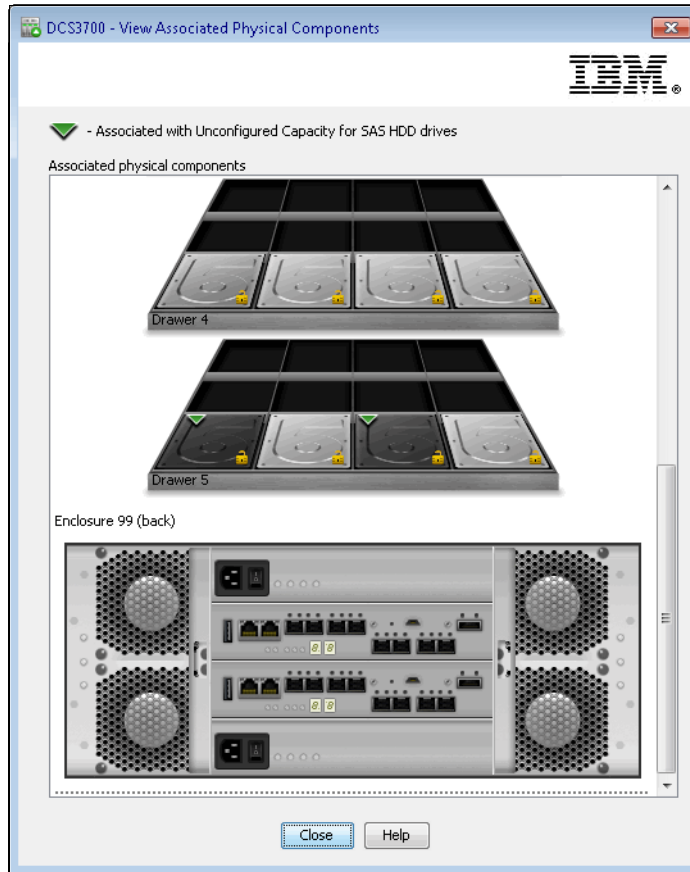


Figure 8-4 View Associated Physical Components

The disks that are marked with the green arrow are unconfigured. You can also recognize unconfigured capacity by looking for the icon that is used for unused disks (dark gray color).

Right-click the unconfigured capacity in the Logical pane, and start configuring an array or pool by clicking **Create Array** or **Create Disk Pool**.

If you can create a logical drive directly from unconfigured capacity (see the menu in Figure 8-3 on page 228), 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 directs you to create the array first.

## 8.2.2 Disk pool

The disk pool feature is a new way to deliver RAID protection and consistent performance. A disk pool is a set of drives that is logically grouped in the storage subsystem. The drives in each disk pool must be of the same drive type and drive media type, and they must be similar in size. As with an array, you can create one or more logical drives in the disk pool. However, the disk pool is different from the array by the way the data is distributed across the drives that comprise the disk pool.

A disk pool is similar to an array, with the following differences:

- ▶ The data in a disk pool is stored randomly on all of the drives in the disk pool, unlike data in an array, which is stored on the same set of drives.
- ▶ You do not specify a RAID level for a disk pool.
- ▶ A disk pool does not use hot spare drives.
- ▶ A disk pool allows many drives to be grouped.

In an array, the data is distributed across the drives based on a RAID level, as shown in Figure 8-5. You can specify the RAID level when you create the array. The data for each logical drive is written sequentially across the set of drives that comprise the array.




Figure 8-5 24-drive traditional RAID

In a disk pool, as shown in Figure 8-6, the storage management software distributes the data for each logical drive randomly across a set of drives that comprise the disk pool. Each disk pool must have a minimum of 11 drives. Although there is no limit on the maximum number of drives that can comprise a disk pool, the disk pool cannot contain more drives than the maximum limit for each storage subsystem. The storage management software automatically configures the RAID level when you create the disk pool. You cannot set or change the RAID level of disk pools or the logical drives in the disk pools.




Figure 8-6 24-drive pool

## Disk pool advantages

Here are the advantages of disk pools:

- ▶ Easy to create: It is easy to create a disk pool in the storage management software. To create a disk pool, select the drives from a list of eligible drive candidates. After a disk pool is created, you create logical drives. When you create disk pool logical drives, the only attribute you must specify is the logical drive capacity.

- ▶ **Better usage of drives:** When you add drives to a storage subsystem, the storage management software automatically detects the drives and prompts you to create a single disk pool or multiple disk pools based on the drive type and the current configuration. If disk pools were previously defined, the storage management software provides the option of adding the compatible drives to an existing disk pool. When new drives are added to an existing disk pool, the storage management software automatically redistributes the data across the new capacity, which now includes the new drives that you added. The data in the logical drives remain accessible when you add the drives to the disk pool. When you delete disk pool logical drives, the capacity of those logical drives is added to the total usable capacity of the disk pool and, therefore, can be reused.

**Note:** You can manually create a disk pool, if you prefer not to proceed with the automatic disk pool creation process.

- ▶ **Reduced hot spots:** A host might access some drives in the array for data more frequently than other drives because of the sequential manner in which the data is written to the drives. This frequency of access to drives creates hot spots in the array. In a disk pool, the hot spots are reduced because of the random manner in which the data is spread across many drives. The reduction of hot spots in the disk pool improves performance of the storage subsystem.
- ▶ **Faster reconstruction of data:** Disk pools do not use hot spare drives for data protection like an array does. Instead of hot spare drives, disk pools use spare capacity within each drive that comprises the disk pool, as shown in Figure 8-7.




Figure 8-7 Difference between drive reconstruction on an array and a disk pool

In hot spare drive coverage, the maximum drive IOPS limits the speed of reconstruction of data from the failed drive to the hot spare drive. In a disk pool, the reconstruction of data is much faster because the spare capacity in all of the drives that comprise the disk pool is used. Additionally, the data to reconstruct after a drive failure is reduced because the data is spread randomly across more drives in a disk pool. Faster reconstruction of data in a disk pool also reduces the risk of more drive failures during a reconstruction operation. Unlike arrays, the time period for which the disk pool is exposed to multiple drive failures during a reconstruction operation is reduced.

- ▶ **Reduced maintenance:** You can configure the storage management software to send alert notifications when the configured capacity of a disk pool is reaching a specified percentage of free capacity. Additionally, you do not need to manage any hot spare drives. You can replace a set of drives during a scheduled maintenance of the storage subsystem.

### Disk pool limitations

Here are the limitations of disk pools:

- ▶ Dynamic Segment Sizing (DSS) is not supported for disk pools.
- ▶ You cannot change the RAID level of a volume in a disk pool. They are always created as a RAID 6 volume.
- ▶ You cannot export a disk pool from a storage subsystem or import the disk pool to a different storage subsystem.
- ▶ All drive types (Fibre Channel, SATA, or SAS) in a disk pool must be the same.
- ▶ All drive media types in a disk pool must be the same. Solid-state disks (SSDs) are not supported.
- ▶ You can use the T10 PI (Protection Information) (T10 PI) capabilities of a drive set in a disk pool if all drives match in their T10 PI capabilities. You can use a drive set with mixed attributes, but the T10 PI capabilities of the drive cannot be used.
- ▶ If you downgrade the controller firmware version of a storage subsystem that is configured with a disk pool to a firmware version that does not support disk pools, the logical drives are lost and the drives are treated as unaffiliated with a disk pool.

## 8.2.3 Create Array

The example that follows assumes that there is unconfigured capacity on the DCS3700. This procedure illustrates the most common steps to follow in setting up an array from unconfigured drives. Start with the creation of an empty array, and then the logical drive creation is described in 8.6.1, “Creating a logical drive” on page 246.

To create an array, complete the following steps:

1. In the Logical pane (Figure 8-8 on page 233), right-click the **Total Unconfigured Capacity** and select **Create Array**.

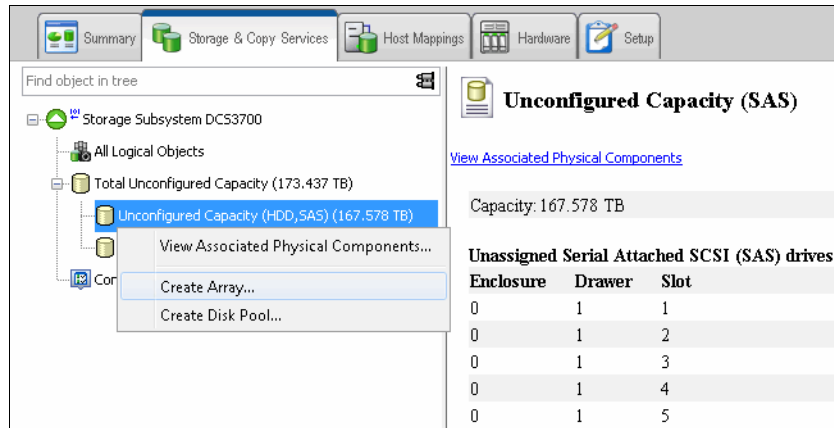


Figure 8-8 Using unconfigured capacity

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, enter 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 8-9. 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.

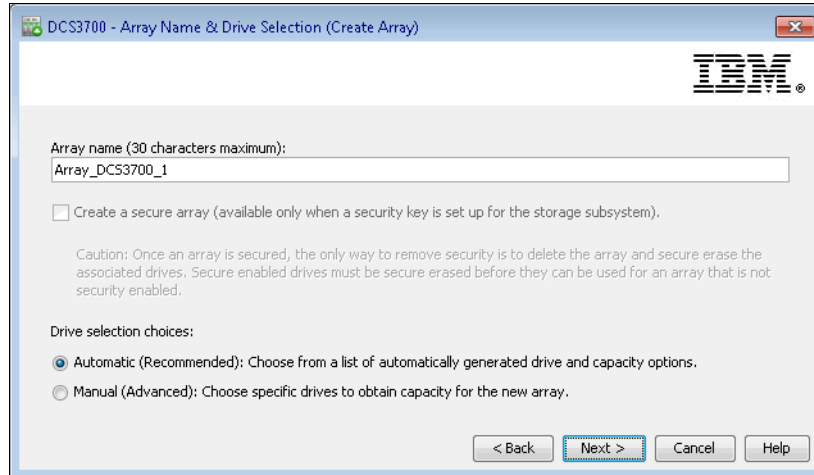


Figure 8-9 Create Array and assign a name and mode

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 must select all the drives individually. Make sure that you select them to maximize performance and availability.

You can also complete manual selection from the Physical tab. You can select multiple unconfigured disk drives in the Physical pane by pressing the Ctrl key before selecting, to create an array. If you right-click one of preselected disks and select **Create Array**, the same initiation window opens, as shown in Figure 8-9 on page 233. In this window, you can see that the Manual mode is preselected. If you use the Manual mode, follow the instructions in this chapter, and when you reach the window for disk drives selection, you see that your planned disks are already selected for a new array.

3. Click **Next** to select the RAID level.
  - Select the wanted RAID level. The window that is shown in Figure 8-10 now shows the capacity options, depending on the unconfigured drives that are available in your storage subsystem. If you have different disk sizes, you have more than one option for the same number of disks. To determine the best RAID level for your specific environment and application, see Chapter 2, “IBM System Storage DCS3700 storage subsystem planning tasks” on page 23.
  - Select the line with the required capacity from the list and click **Finish**. In our example (Figure 8-10), we create a RAID 1 array with four 1 5 K RPM SAS drives.

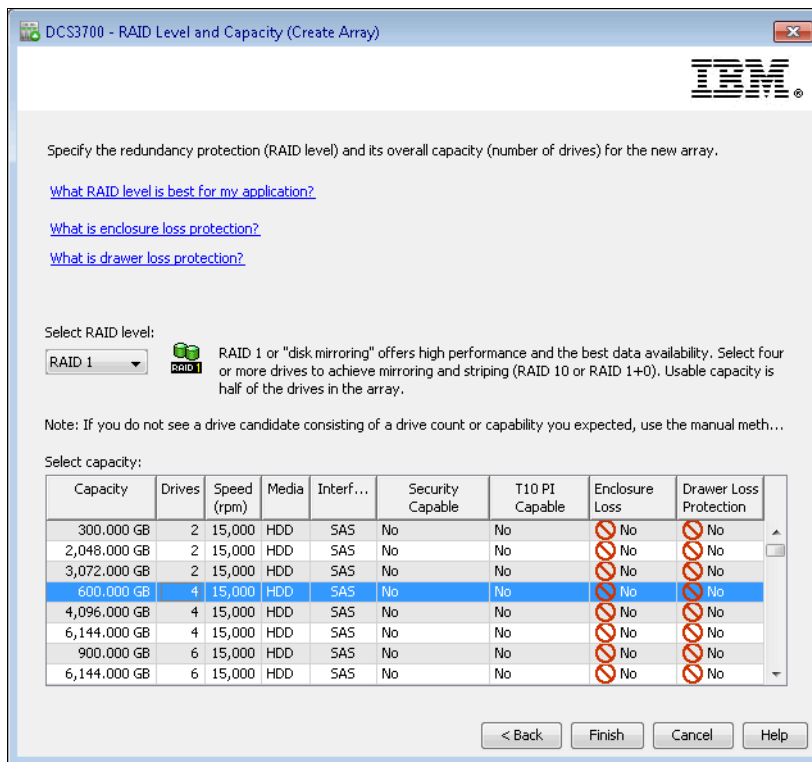


Figure 8-10 RAID level and capacity

4. The Array Success window (Figure 8-11 on page 235) opens and confirms that the array is now created. You can click **Yes** to continue with the creation of a logical drive (see 8.6.1, “Creating a logical drive” on page 246). Select **No** to finish for now.



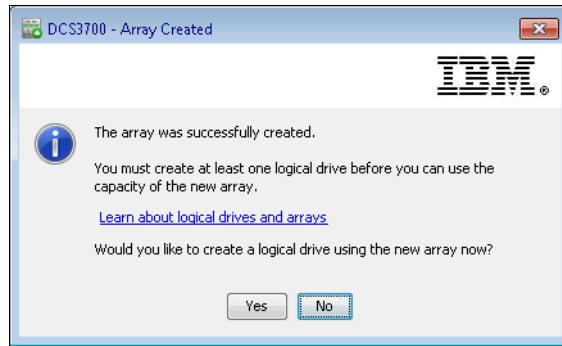


Figure 8-11 Array Created

5. Your empty array is created immediately because it does not need to format any logical drive capacity. You can repeat the same procedure for creating another array.

## 8.3 Working with arrays

An array (sometimes called a disk group) is a set of disks with a defined RAID level. You can configure one or more logical drives on it. The available capacity of the array is displayed in the Logical pane as *free capacity*. If no logical drives are configured, the free capacity that is displayed is also the total capacity of the array.

Actions that you can take on the array can be accessed by right-clicking the array and selecting from the menu (Figure 8-12).

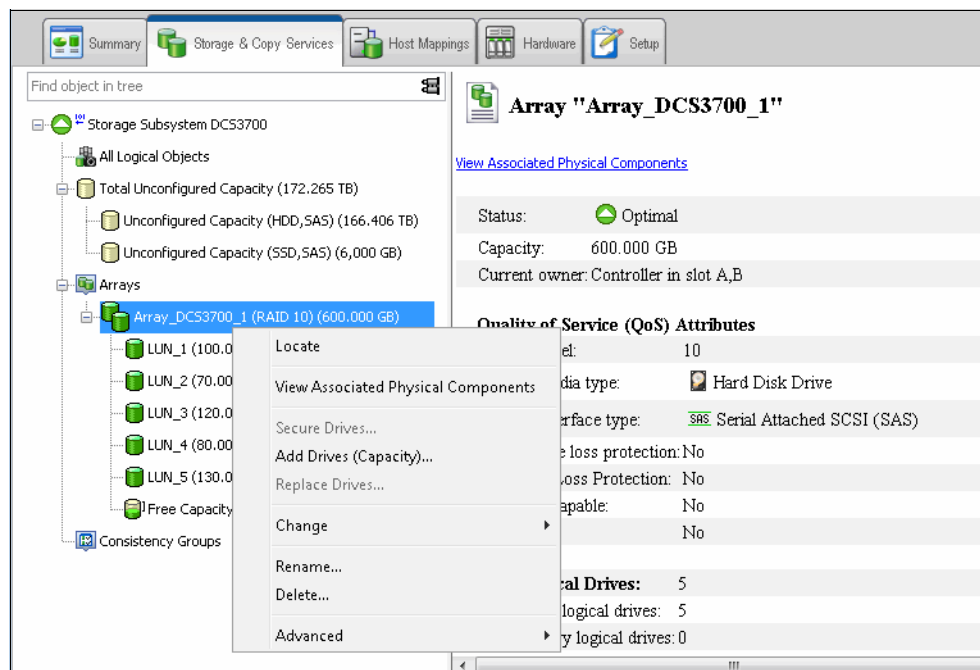


Figure 8-12 Array menu

These options are described in the following sections:

- ▶ Locate and View Associated Components
- ▶ Add Free Capacity (Drives)

- ▶ Replace Drive
- ▶ Change Ownership and RAID level
- ▶ Deleting and renaming arrays

**Note:** Because an array can be created from only from unconfigured capacity, the Create option in the Array menu is never active in this tab.

### 8.3.1 Locate and View Associated Components

Click **Locate** to find the disk drives in the rack from which the array is configured. The orange LED diodes on these disks flash.

Click **View Associated Physical Components** to view the components, as shown in Figure 8-13. The disks that are marked by a green arrow are members of the selected array. The green arrow can mark one or both controllers, if the logical drives are owned by both controllers.

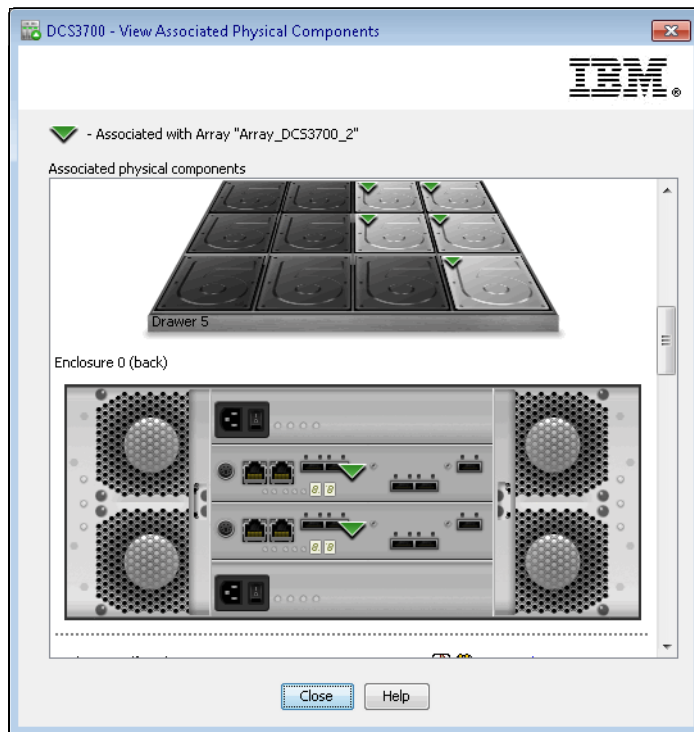


Figure 8-13 Associated components for array

### 8.3.2 Add Free Capacity (Drives)

In today's IT environment, the need for storage space grows constantly. Because this process must be nondisruptive 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 more 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)*, which is explained in 8.8.2, “Increase Capacity option” on page 258.

Use the Add Capacity option to expand the capacity of a selected array by adding unassigned drives. The expansion causes more free capacity to be included in the array. You can use this free capacity to create more logical drives. To add new drives into an array, right-click the array and select **Add Free Capacity (Drives)**. In the Add Drives window (Figure 8-14), 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.

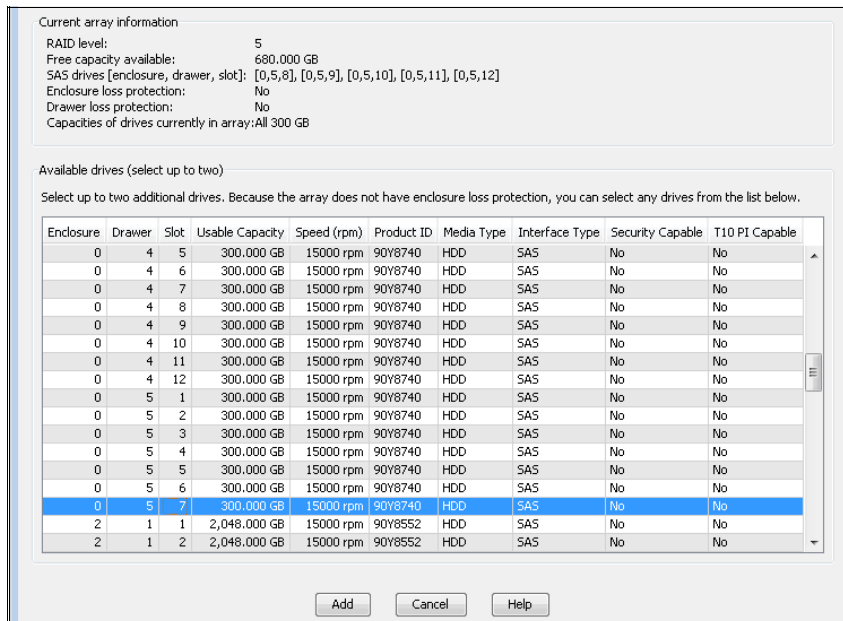


Figure 8-14 Add Free Capacity (Drives)

**Note:** It is not a general rule to add an SAS drive with slower rotational speed, as performance of the new disk can slow down existing faster disks in the array. In the other case, a new disk can be underutilized if its speed is higher than the speed of the current disks in the array

For RAID 3, 5, and 6, select one or two drives, and for RAID 1 and 10, select two drives. Pay attention to the following considerations:

- ▶ Only SAS drives are supported on DCS3700, so every unconfigured disk drive with the appropriate capacity can be listed as a candidate to add.
- ▶ Select drives that have a capacity equal to the current drive capacities in the array.

Drives larger than the other drives that are participating in the array can be added, but avoid this situation because their usable capacity is reduced so that they match the current drives capacities in the array.

**Note:** After the procedure for adding drive is started, it cannot be stopped because the subsystem must redistribute the data that is 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.

### 8.3.3 Replace Drive

This menu item is active if a disk drive in the array fails or is missing. In this case, you can replace the failed or missed drive with another one. This procedure is described in 9.4, “Failed disk drive replacement” on page 319.

### 8.3.4 Change Ownership and RAID level

This section describes changing the controller ownership and the RAID level.

#### Ownership

To change the logical drive controller ownership of all logical drives in the array to one selected controller, click **Change** → **Ownership/Preferred Path** → **Controller A/B**. The change of controller ownership must be confirmed, as shown in Figure 8-15.

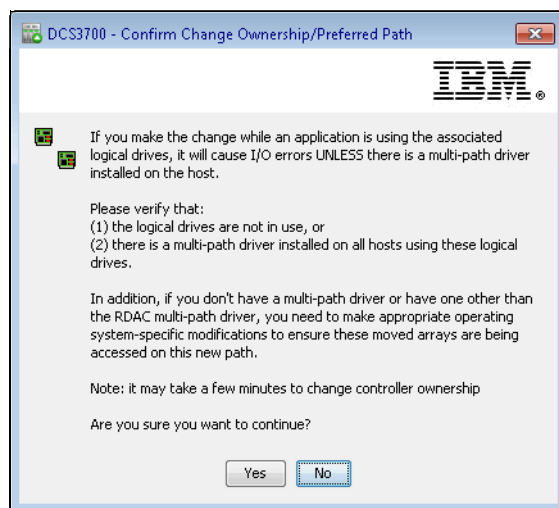


Figure 8-15 Change Ownership confirmation

#### RAID level

Changing the RAID level of an array is performed in a nondisruptive 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 limitations 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 is displayed that states 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 before you perform 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 much I/O happens, so there is 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. For more information, see 8.8.5, “Change Modification Priority” on page 262.

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 if there is a drive failure. This minimizes the critical time of a non-redundant operation that is 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, and name changes. Also, it is possible to define the mappings during RAID level migration.

**Note:** Even though RAID migration is a nondisruptive 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 shown in Figure 8-16. The current RAID level is marked by a black bullet, and you can select your new RAID level.

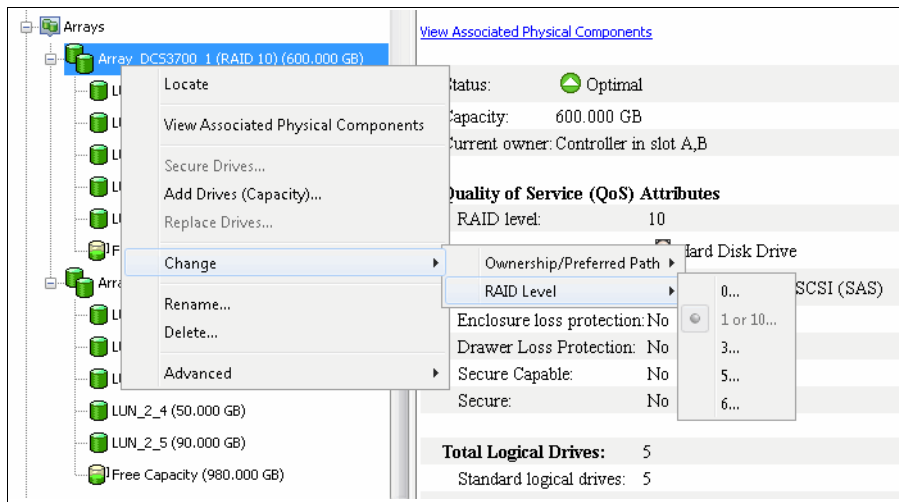


Figure 8-16 Change RAID level menu

The process of RAID level migration is nondisruptive and no redundancy is lost, but it can take a long time. A warning opens, as shown in Figure 8-17.

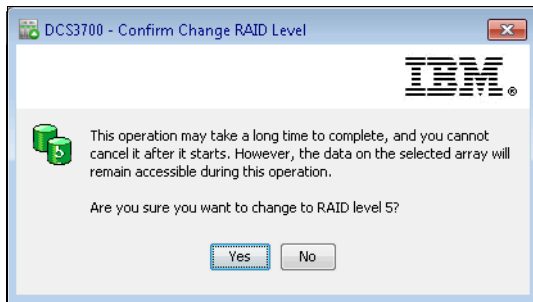


Figure 8-17 RAID level change confirmation

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 a clock icon. You cannot do certain other configuration changes on the array and logical drives until the process is completed.

### 8.3.5 Deleting and renaming arrays

Deleting the array automatically deletes all logical drives in this array (see Figure 8-18) if they are configured.

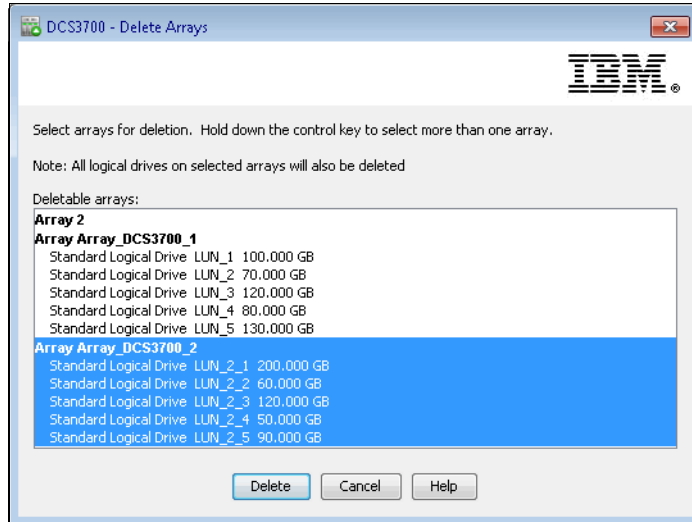


Figure 8-18 Array and logical drives deletion

If some of the logical drives cannot be deleted (because of any reason), the array cannot be deleted.

If you are renaming the selected array (Figure 8-19), you must use an 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.

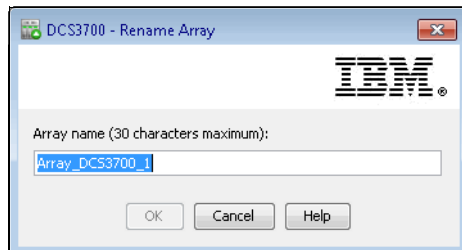


Figure 8-19 Rename array

## 8.4 Creating a disk pool

To create a disk pool, complete the following steps:

1. Right-click **Total Unconfigured Capacity** and select **Create Disk Pool** (see Figure 8-20). This action starts the wizard to create an array, which leads you to the disk pool creation process.

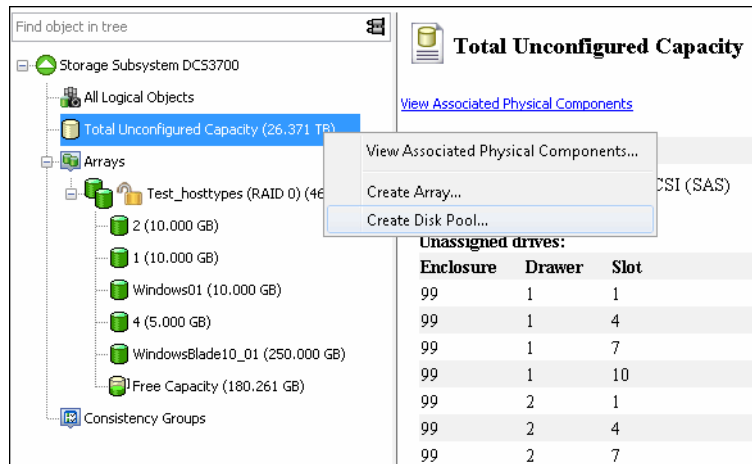


Figure 8-20 Create Disk Pool option



- In the Create Disk Pool window, enter a meaningful disk pool name that describes your set of disks in to the Disk Pool name field, as shown in Figure 8-21. From the Disk pool candidates (minimum 11 drives) table, choose a combination of disks.

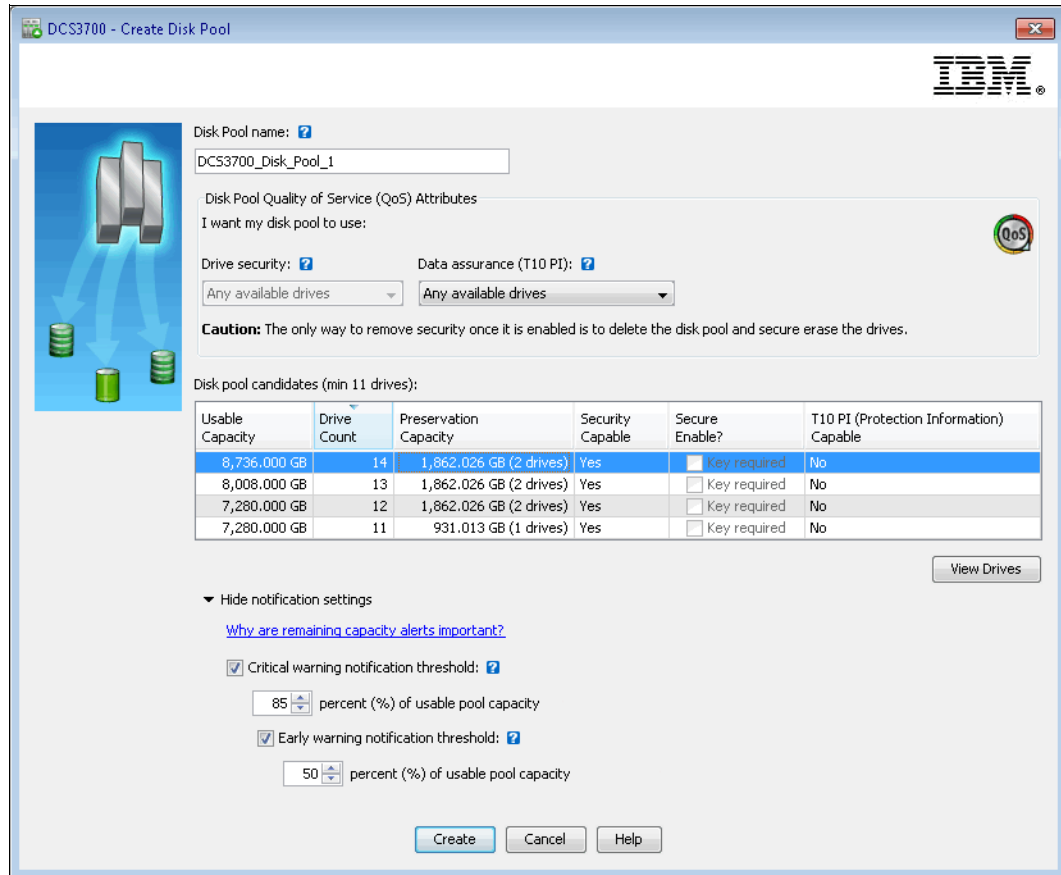


Figure 8-21 Create Disk Pool

- On the same window, in the Hide notification settings section, you can set a warning notification for the created disk pool.

**Note:** A disk pool needs sufficient remaining capacity to successfully perform storage subsystem operations. You can prevent interruptions to these important operations by knowing when to add drives (capacity) to the disk pool. Alert notifications inform you when your remaining capacity reaches a specified percentage, indicating that you might need to add drives (capacity).

4. Click **Create**, and a new disk pool displays in the Logical pane, as shown in Figure 8-22.

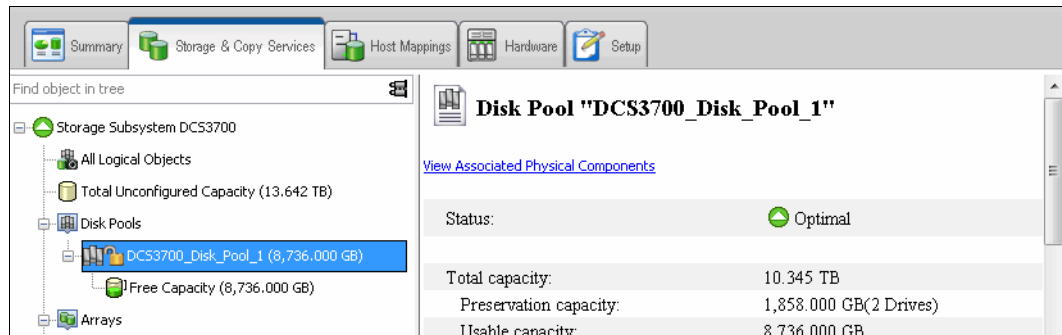


Figure 8-22 Created disk pool

## 8.5 Working with a disk pool

You can perform actions on the disk pool by right-clicking the disk pool and selecting an item from the menu that opens. The menu items are the same as the ones that are described in 8.3, “Working with arrays” on page 235. This section explains a few differences and limitations that are specific with disk pools.

### 8.5.1 Locate and View Associated Components

These options are the same as for an array, which is described in 8.3.1, “Locate and View Associated Components” on page 236

## 8.5.2 Add Free Capacity (Drive)

Use the Add Free Drives (Capacity) option to increase the free capacity of an existing disk pool by adding unassigned drives, as shown in Figure 8-23. After you add unassigned drives to a disk pool, the data in each logical drive of the disk pool is redistributed to include the additional drives.

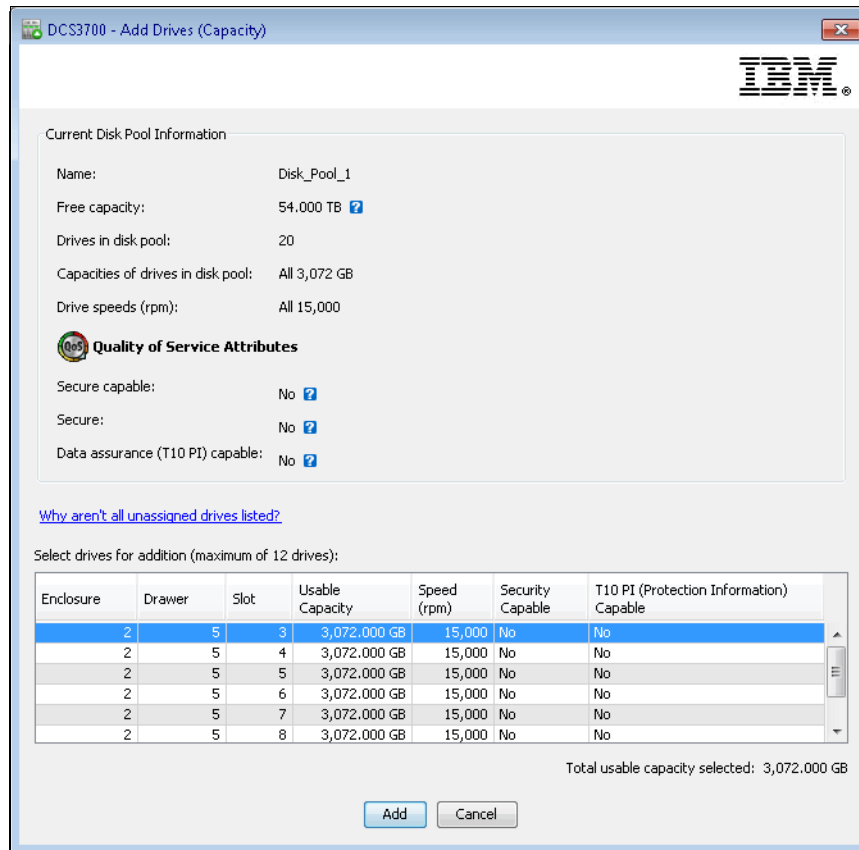


Figure 8-23 Add drives to a disk pool

During the process of adding disks to the disk pool, follow these guidelines:

- ▶ The status of the disk pool must be Optimal before you can add unassigned drives.
- ▶ You can add a maximum of 12 drives to an existing disk pool. However, the disk pool cannot contain more drives than the maximum limit for a storage subsystem.
- ▶ You can add only unassigned drives with an Optimal status to a disk pool.
- ▶ The data in the logical drives remains accessible during this operation.
- ▶ Adding drives with different capacities is possible.

## 8.6 Working with free capacity in arrays

The only way to use the free capacity of a specified array is to configure a new logical drive.

If you select **Free Capacity** in the Logical pane tree, as shown in Figure 8-24, you can see basic information about this capacity in the Properties pane. You can right-click **Free Capacity** and select **Create Logical Drive** to create a LUN by using that free capacity.

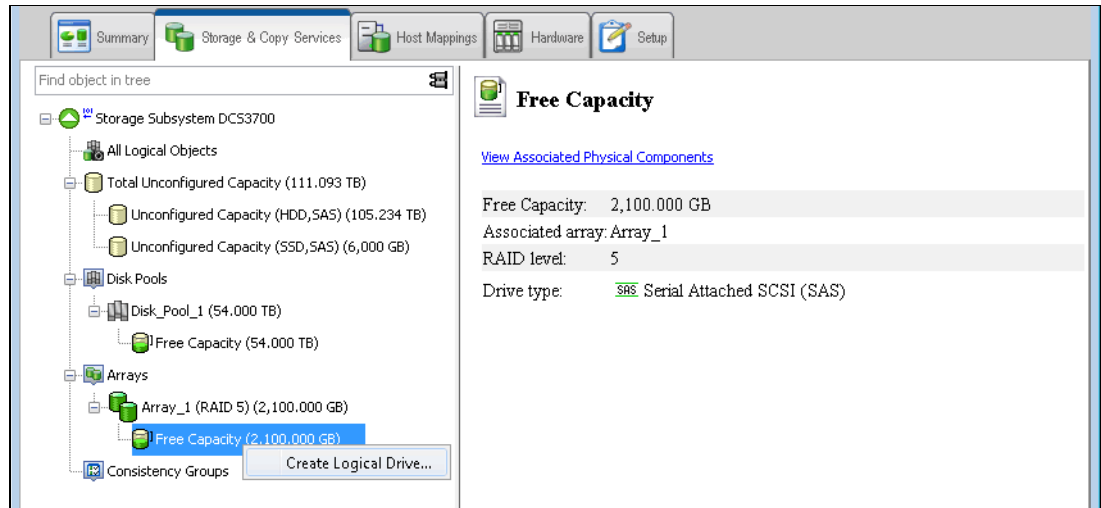


Figure 8-24 Free Capacity menu

### 8.6.1 Creating a logical drive

To create a logical drive (sometimes referred to as a volume), complete the following steps:

1. Right-click **Free Capacity** for the array where you want to create a new logical drive and select **Create Logical Drive**.

2. The Specify Capacity Parameters window opens, as shown in Figure 8-25.

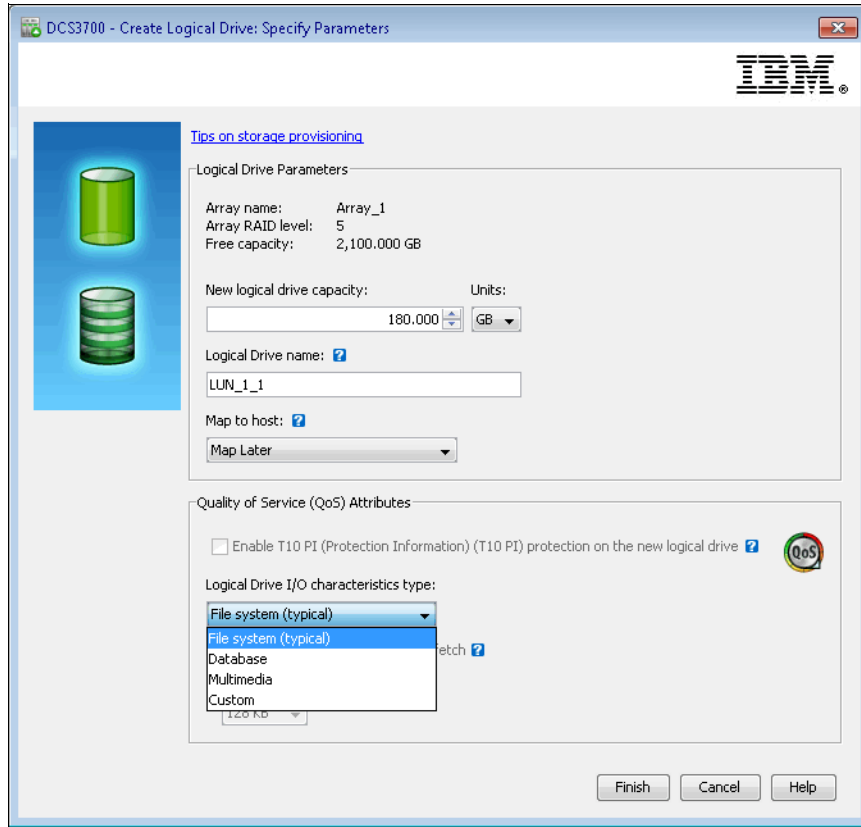


Figure 8-25 Specifying logical drive capacity, name, and I/O characteristic type

Complete the following steps:

- a. Enter the wanted size capacity, which must be less than the available free capacity of the array. Choose this number carefully if you want to define more than one logical drive in the array.

**Note:** Specify the size by entering the required number even if you can use all the capacity of an array. The offered number could be rounded, and you might have trouble 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. Specify the logical drive I/O characteristics for the logical drive based on your needs. The available options are shown in Table 8-1.

**Note:** *Dynamic cache read prefetch* allows the controller to copy more sequential data blocks into the cache while it is reading data blocks from a drive to the cache.

Table 8-1 Logical Drive I/O characteristics type

Option	Enable dynamic cache read prefetch	Segment size value
<b>File system (typical)</b>	Enabled.	128 KB
<b>Database</b>	Disabled.	128 KB
<b>Multimedia</b>	Enabled.	Greater than or equal to 256 KB
<b>Custom</b>	Select your own value.	Select your own: 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, or 512 KB

3. In the drop-down menu **Map to host**, leave the default as **Map later**. Mapping is explained in Chapter 10, “Administration: Mappings tab” on page 327.

**Important:** Manually mapping logical drives to hosts prevents unwanted mapping and is always the preferred choice.

4. Click **Finish** and the Creation Successful window opens, as shown in Figure 8-26. You are prompted with the option to create another logical drive. Click **No** if you are finished creating logical drives or want to continue later.

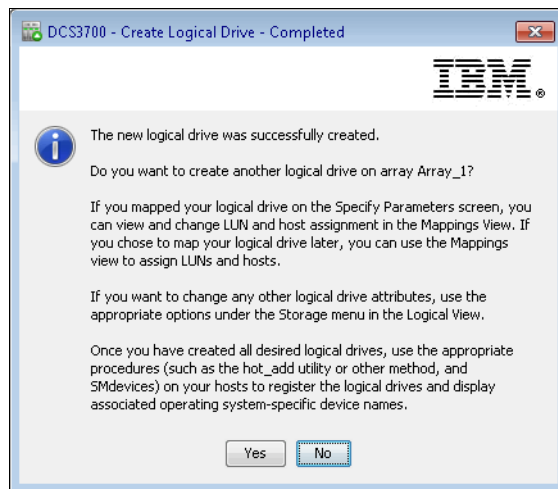


Figure 8-26 Create logical drive - Creation Successful

- The new logical drive is formatted. You can check the progress of this operation in the Properties pane of Storage Manager (Figure 8-27). The Logical Drive icon in the Logical pane is marked with a clock symbol during this process.

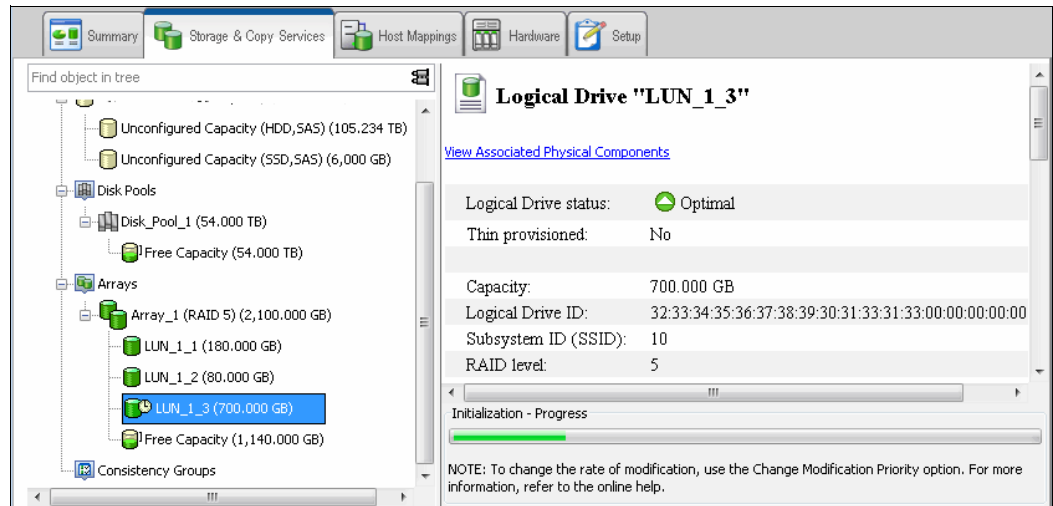


Figure 8-27 Logical drive that is formatted

- The new logical drive is immediately available to be mapped to the host. This function is called *immediately availability format (IAF)*. You can repeat the same process to create other arrays or logical drives.

## 8.7 Working with free capacity in a disk pool

As with arrays, the only way to use the free capacity of a specified disk pool is to configure a new logical drive.

If you select **Free Capacity** in the Logical pane tree, you can see basic information about this capacity in the Properties pane. You can right-click **Free Capacity** and select **Create Logical Drive** to start create a LUN by using that free capacity.

During the process of creating logical drives within a disk pool, you can create them as one of the following drives:

- ▶ A standard logical drive
- ▶ A thin logical drive

### 8.7.1 Creating a standard logical drive

To create a logical drive within a disk pool, complete the following steps:

- Right-click **Free Capacity** in the array where you want to create a logical drive and select **Create Logical Drive**.
- In the Specify Capacity Parameters window, as shown in Figure 8-28 on page 250, complete the following steps:
  - Enter the wanted size capacity, which must be less than the available free capacity of the disk pool. Choose this number carefully if you want to define more than one logical drive in the disk pool.

- b. Assign a unique name to the logical drive.
- c. Enable or disable dynamic cache read prefetch, as shown in Figure 8-28. Enabling this option improves storage performance.

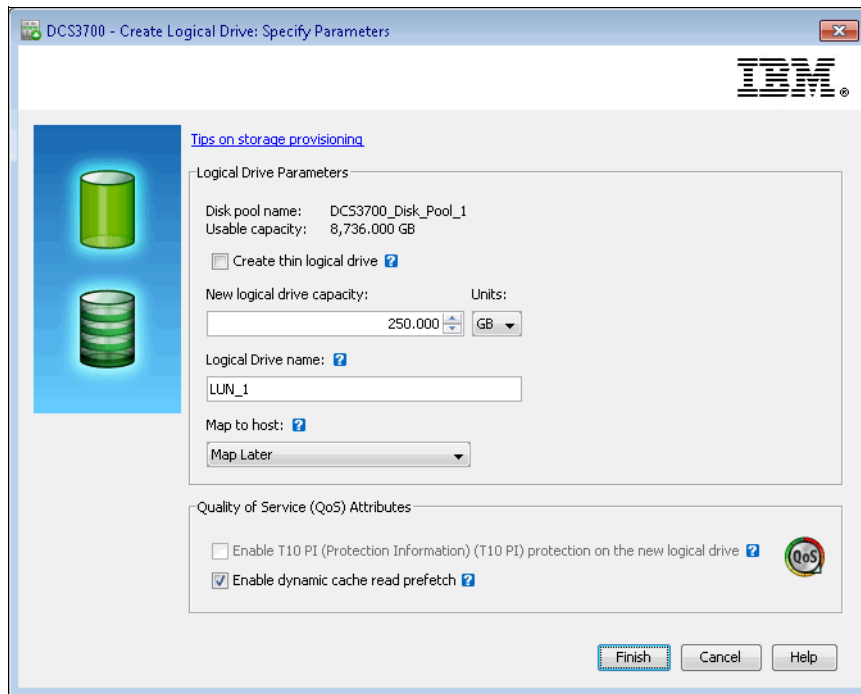


Figure 8-28 Specify parameters for logical drive in the disk pool

3. Click **Finish** and the Creation Successful window opens. You are prompted with the option to create another logical drive. Click **No** if you are finished creating logical drives or want to continue later.
4. The new logical drive is formatted and ready for use.

## 8.7.2 Thin logical drive

Thin logical drives report a virtual capacity to attached hosts that is not directly represented by actual storage, as shown in Figure 8-29.

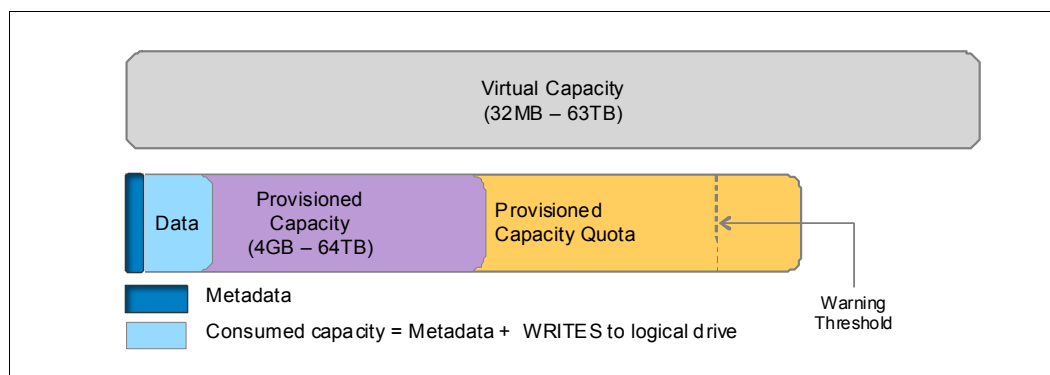


Figure 8-29 Thin logical drives



Thin volumes have three characteristic capacities:

- ▶ **Virtual capacity:** The capacity that is reported to the hosts. This capacity is presented to the host. You normally set the virtual capacity to be the maximum capacity to which you think the logical drive will grow.
- ▶ **Provisioned capacity:** The amount of physical space that is allocated to the volume. Physical capacity is allocated out of virtual capacity either manually or automatically. When the amount of data that was written gets close to the physical capacity, you can expand the physical capacity. The storage management software can automatically expand the physical capacity or you can do it manually. If expansion is done automatically, you can set a maximum expansion capacity. The maximum expansion capacity allows you to limit the logical drive's automatic growth below the virtual capacity.
- ▶ **Consumed capacity:** The amount of physical space that is written to the logical drive. This includes the user data and volume metadata.

Thin provisioning allows logical drives to be created with a large virtual capacity and relatively small physical capacity, which is beneficial for storage usage and efficiency. Thin logical drives can help simplify storage administration because the physical capacity can increase as the application needs change, without disrupting the application, allowing for better storage usage.

**Note:** A thin logical drive can be created only from a disk pool and not from the array.

The capacity limits for a thin logical drive are shown in Table 8-2.

Table 8-2 Thin logical drive capacity limits

Type of capacity	Size
Minimum virtual capacity	32 MB
Maximum virtual capacity	63 TB
Minimum physical capacity	4 GB
Maximum physical capacity	64 TB

**Note:** Any capacity that is not a multiple of 4 GB is allocated but not usable. To make sure that the entire capacity is usable, specify the capacity in 4 GB increments. If unusable capacity exists, the only way to regain it is to increase the capacity of the logical drive.

## Benefits of thin logical drives

Here are the benefits of thin logical drives:

- ▶ A thin logical drive requires only a minimum physical space to be allocated versus a standard logical drive, which requires a full allocation of the specified capacity.
- ▶ A thin logical drive can automatically expand the data repository as needed from available space in the pool.
- ▶ Expanding the virtual address space can be done with no physical allocation change versus a standard logical drive, which requires more physical space and usually is a long running operation.
- ▶ Improves storage usage and efficiency.
- ▶ Allows the deferral of storage acquisition.

- ▶ Lessens the impact of administrator uncertainty when sizing logical drives.
- ▶ Reduces the physical footprint and power and cooling requirements.

### Limitations for a thin logical drive

Thin logical drives support all of the operations that standard logical drives do with the following exceptions:

- ▶ You cannot change the segment size of a thin logical drive.
- ▶ You cannot enable the pre-read redundancy check for a thin logical drive.
- ▶ You cannot use a thin logical drive as the target logical drive in a VolumeCopy.
- ▶ You cannot use a thin logical drive in a FlashCopy operation.
- ▶ You cannot use a thin logical drive in an Enhanced Remote Mirroring operation.
- ▶ Ownership of all thin logical drive components must be by the same controller.

**Note:** If you want to change a thin logical drive to a standard logical drive, use the VolumeCopy operation to create a copy of the thin logical drive. The target of a VolumeCopy is always a standard logical drive.

## 8.7.3 Creating a thin logical drive

To create a thin logical drive within a disk pool, complete the following steps:

1. Right-click **Free Capacity** in the array where you want to create a new logical drive and select **Create Logical Drive**.
2. Select the **Create thin logical drive** check box, as shown in Figure 8-30.

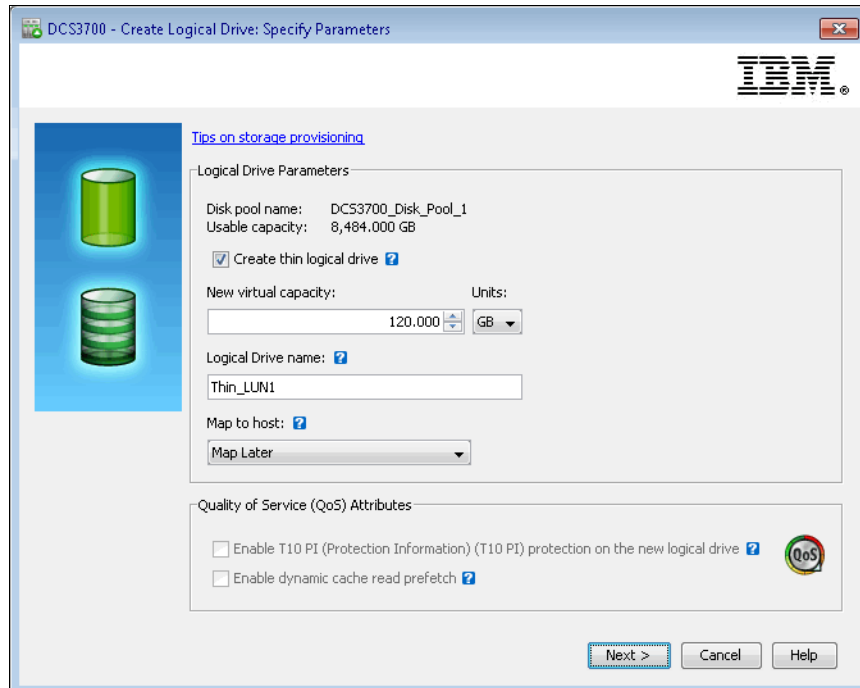


Figure 8-30 Create thin logical drive parameters

Complete the following steps:

- a. Enter the wanted size capacity, which must be less than the available free capacity of the disk pool. Choose this number carefully if you want to define more than one logical drive in the disk pool.
  - b. Assign a unique name to the logical drive.
  - c. In the Map to host list, specify how you want to map hosts to logical drives.
3. Click **Next**. Choose one of the following actions based on whether you want to customize the capacity settings:
- Use recommended settings: Click **Use recommended capacity settings**, as shown in Figure 8-31, and then click **Finish**.

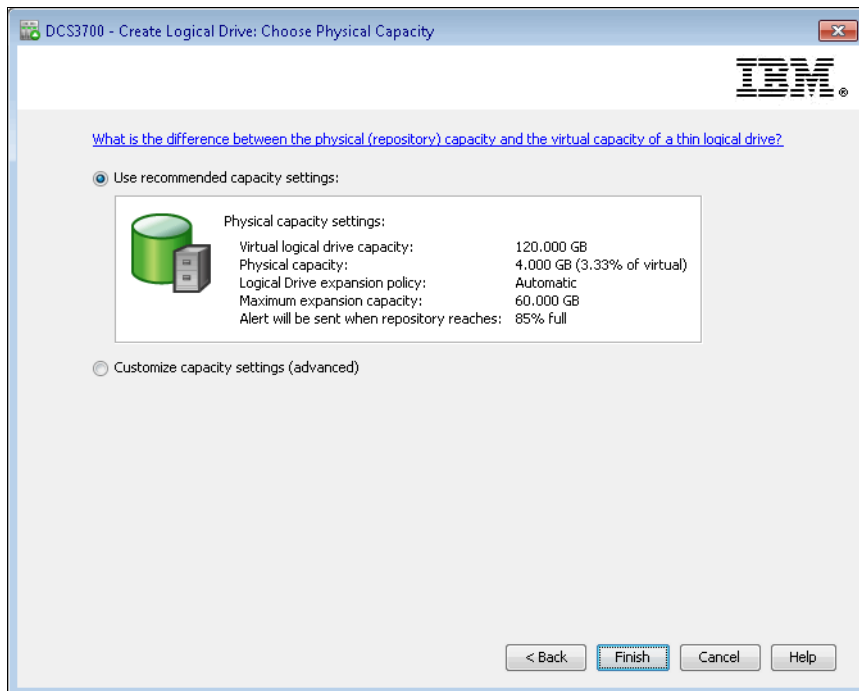


Figure 8-31 Thin logical drive recommended settings

- Choose your own settings: Click **Customize capacity settings (advanced)** and click **Next**.

- As shown in Figure 8-32, complete the **Preferred Capacity** field to indicate the initial physical capacity of the logical drive and the **Units** drop-down list to indicate the specific capacity units to use (MB, GB, or TB). The physical capacity is the amount of physical disk space that is reserved for write requests. The physical capacity must be at least 4 GB in size, and can be no larger than 256 GB in size.

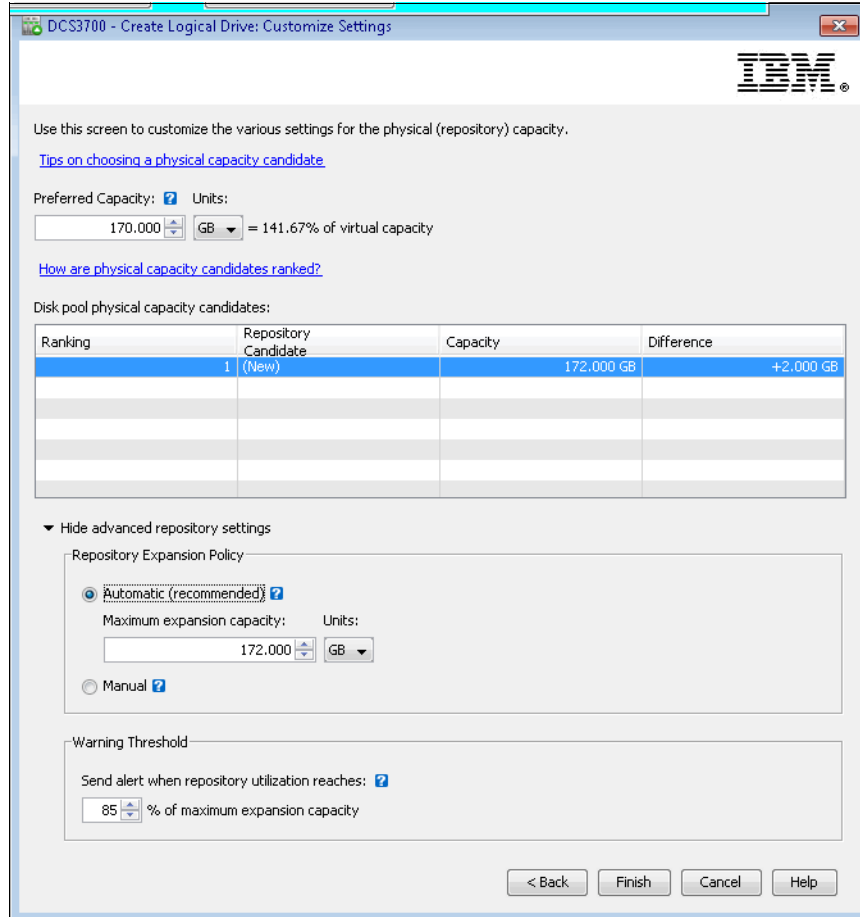


Figure 8-32 Customize capacity settings window

- Based on the values that you entered in step 4, the Disk pool physical capacity candidates table is populated with matching repository logical drives. The new repository candidates that are returned either match the capacity that you specify or are rounded up to the closest 4 GB increment to make sure all the repository capacity is usable. Select a repository from the table. Existing repositories are placed at the top of the list. The benefit to reusing an existing repository is that you can avoid the initialization process that occurs when you create a new one.
- (Optional) To change any of the following options, click **View advanced repository settings**:
  - Repository expansion policy: Select either **Automatic** or **Manual**. When the consumed capacity gets close to the physical capacity, you can expand the physical capacity. The storage management software can automatically expand the physical capacity or you can do it manually. If you select **Automatic**, you also can set a maximum expansion capacity. The maximum expansion capacity allows you to limit the logical drive's automatic growth below the virtual capacity. The value for the maximum expansion capacity must be a multiple of 4 GB.

- Warning threshold: In the **Send alert when repository utilization reaches** field, enter a percentage. The storage management software sends an alert notification when the physical capacity reaches the percentage full.
7. Click **Finish** and the Logical Drive Successfully Created dialog opens, as shown in Figure 8-33.

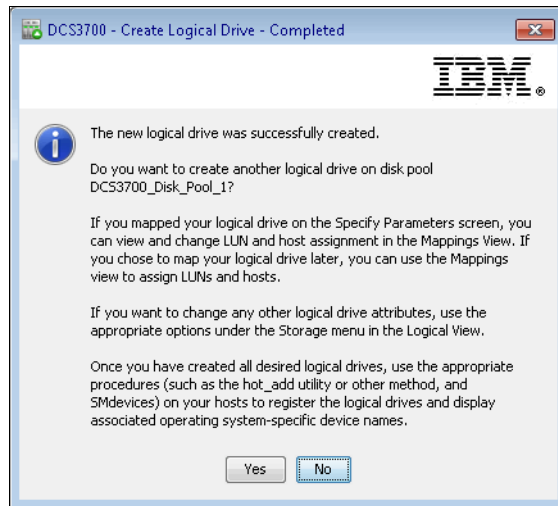


Figure 8-33 Create Logical Drive Completed

## 8.8 Working with logical drives

A logical drive is the logical component that hosts use for data storage. Hosts that are attached to the storage subsystem write data to the logical drives and read data from the logical drives.

You can create a logical drive from either an array or a disk pool. Before you create a logical drive, the array or a disk pool must exist and it must have enough free capacity to create the logical drive.

You can give hosts access to a logical drive either when you create the logical drive or later. All host access is managed through a logical unit number (LUN). Hosts detect LUNs that are, in turn, mapped to logical drives. If you are mapping a logical drive to multiple hosts, use clustering software to make sure that the logical drive is available to all of the hosts.

The logical drive can be tuned by using many parameters, which are described in the following sections.

The basic information about a logical drive is displayed in the Properties pane of the Logical tab (see Figure 8-34).

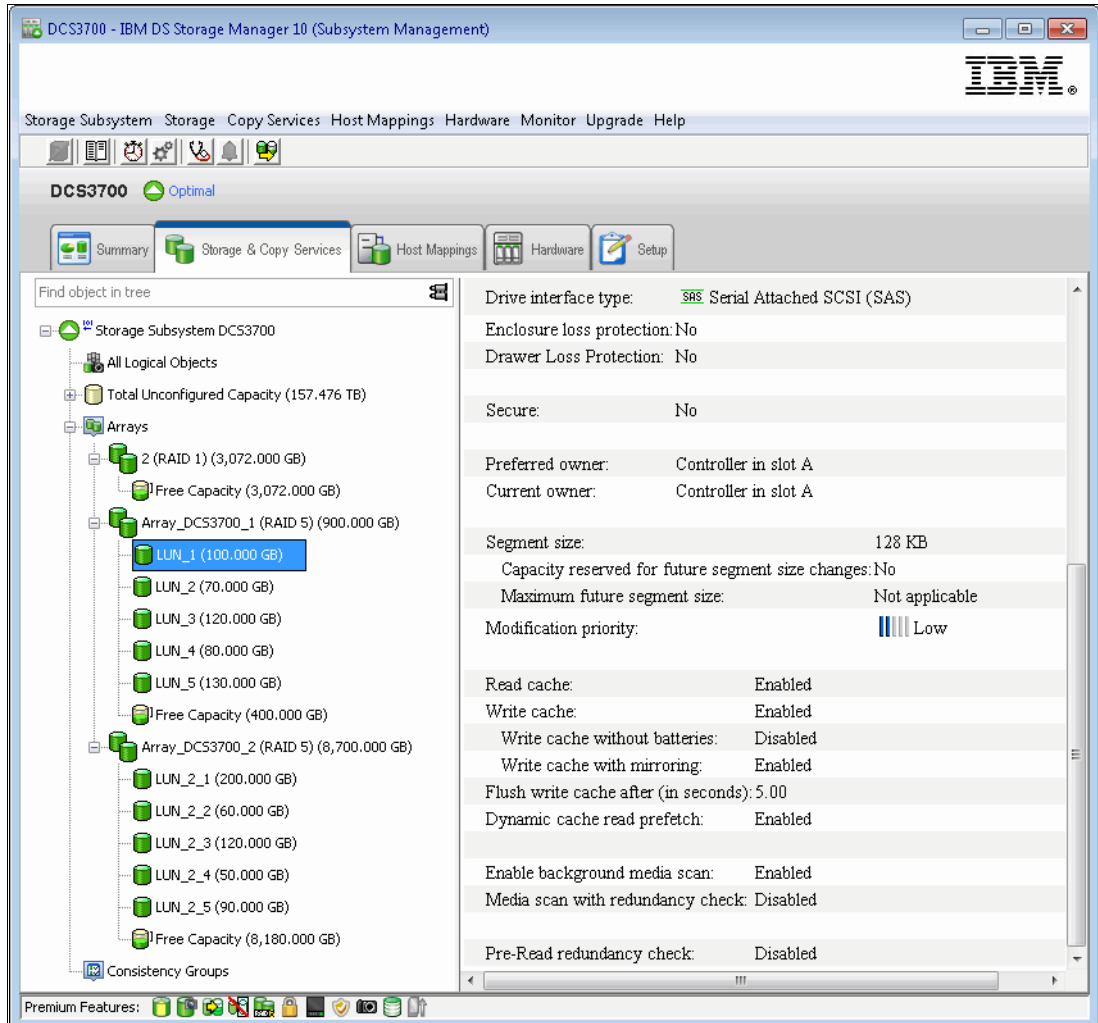


Figure 8-34 Logical Drive - Logical and Properties pane

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 8-34 on page 256). These services are described in *IBM System Storage DS Storage Manager Copy Services Guide*, SG24-7822.

When you right-click a logical drive that you want to modify, a LUN menu opens, as shown in Figure 8-35.

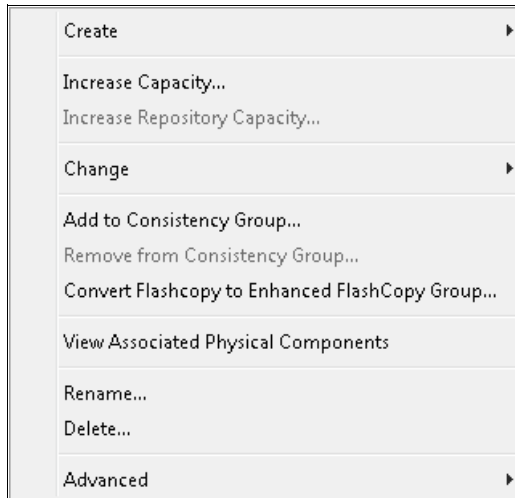


Figure 8-35 Logical drive menu

## 8.8.1 Create option

When you click the first option, **Create**, the following menu opens (Figure 8-36).



Figure 8-36 Create menu

All of these options are related to copy services and they are described in the following resources:

- ▶ *IBM System Storage DS Storage Manager Copy Services Guide*, SG24-7822
- ▶ *Copy Services User's Guide - IBM System Storage DS Storage Manager v10*, which can be found at:

<https://www-947.ibm.com/support/entry/myportal/docdisplay?brand=5000008&Indocid=MIGR-61173>

## 8.8.2 Increase Capacity option

The second option in the menu is Increase Capacity. Use this option to increase the storage capacity of an existing logical drive in array or disk pool. This action is called *Dynamic Volume Expansion* (DVE). You can increase storage capacity by completing one of these tasks:

- ▶ Use the free capacity that is available in the array or the disk pool.
- ▶ Add unconfigured capacity (in the form of unused drives) to the array of the logical drive. Use this option when no free capacity exists on the array. (This procedure applies only to logical drives in arrays.)

When you select **Increase Capacity**, the window that is shown in Figure 8-37 opens.

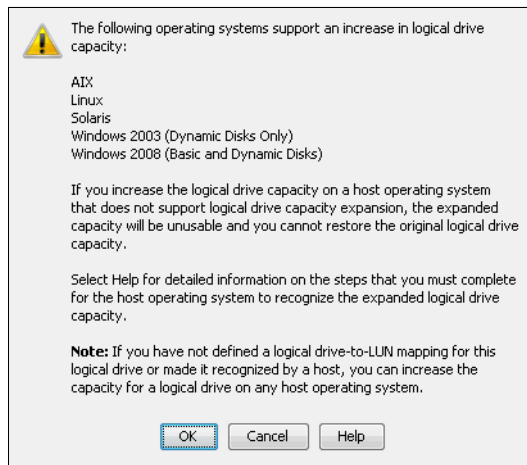


Figure 8-37 Increase Logical drive additional instruction

The window notifies you about operating systems that support logical drive increases.

Here are the operating systems that support a dynamic increase of capacity in a mapped logical drive:

- ▶ AIX
- ▶ Linux
- ▶ Solaris
- ▶ Windows 2003 (Dynamic Disks Only)
- ▶ Windows 2008 (Basic and Dynamic Disks)

**Note:** If you increase the logical drive capacity on a host operating system that is unsupported, the expanded capacity is unusable, and you cannot restore the original logical drive capacity.

**Attention:** Increase logical drive is an one-way operation. It is not possible to decrease space in a logical drive.



Click **OK** to confirm that your operating system supports logical drive increase, and in the window that is shown in Figure 8-38, choose the capacity that you want to add to your logical drive, or you can click **Add Drives** and choose the drives that are added to the array or disk pool where the logical drive belongs.

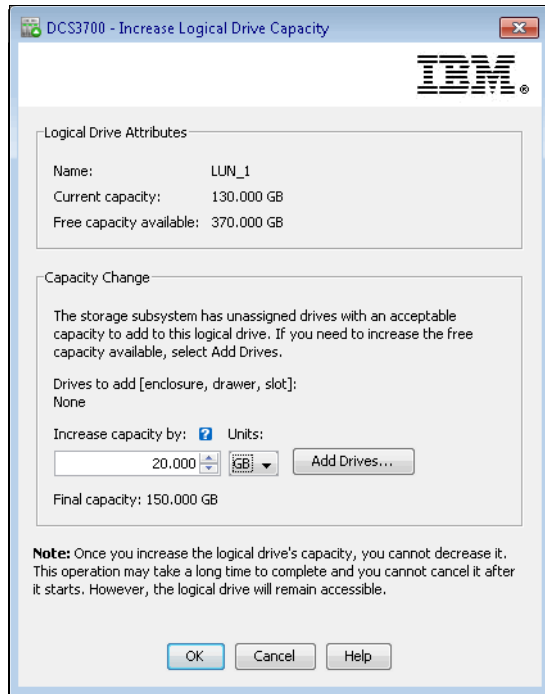


Figure 8-38 Increase Logical Drive Capacity

The **Free Capacity available** field shows you how much free space you have in an array or disk pool for your logical drive increase. In the **Increase capacity by:** field, enter the value for the logical drive increase, and the **Final capacity** field shows you how large your logical drive will be after the increase. Click **OK**, and the process starts. You can view the progress of the operation by clicking the logical drive, as shown in Figure 8-39 on page 260.

**Note:** Increasing the capacity of a standard logical drive is supported only on certain operating systems. If the logical drive capacity is increased on a host operating system that is not supported, the expanded capacity is unusable and the original logical drive capacity cannot be restored.

Increasing the capacity of a DDP volume is limited by the maximum volume size of 64 TB for the logical drive.

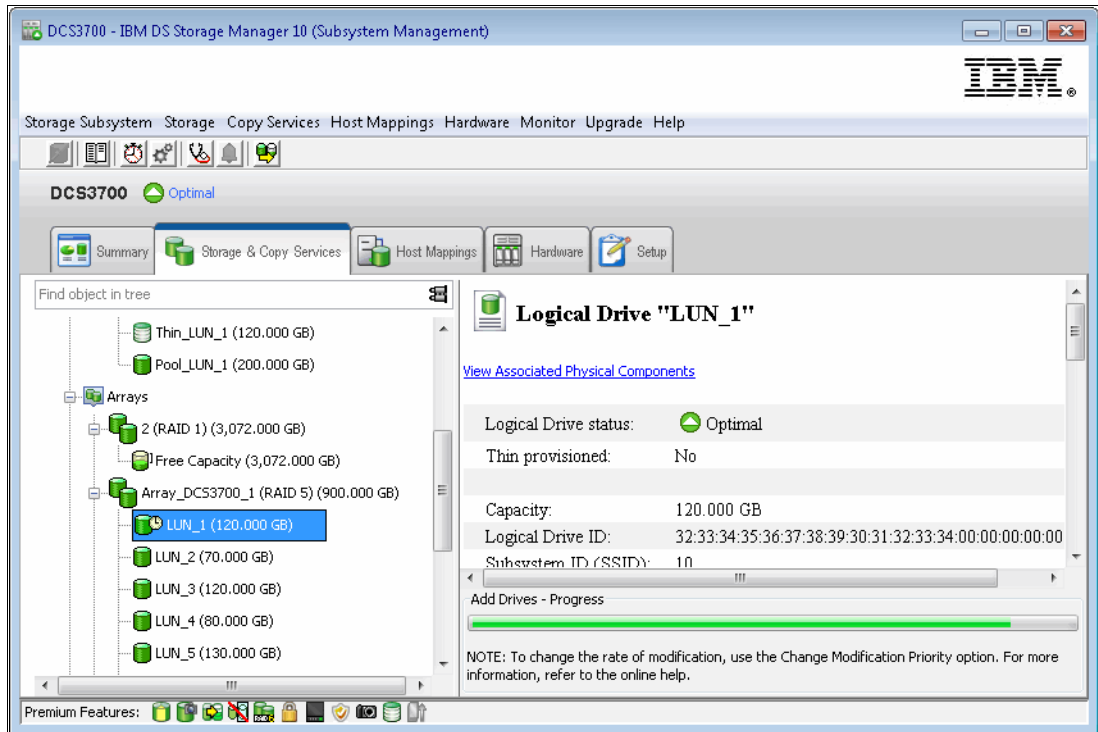


Figure 8-39 Logical drive increase progress

### 8.8.3 Increase Repository Capacity

Use this option to increase the physical capacity of a thin logical drive. Physical capacity is the amount of physical disk space that is allocated for writing data.

This option is available only for a thin logical drive within a disk pool, as shown in Figure 8-40.

**Note:** If free capacity is not available in the disk pool, you cannot increase the capacity of the logical drive. You must increase the size of the disk pool first or delete unused logical drives. If the maximum physical capacity for the thin provisioned volume of 64 TB was reached or the expansion policy is set to automatic, you cannot increase the capacity.

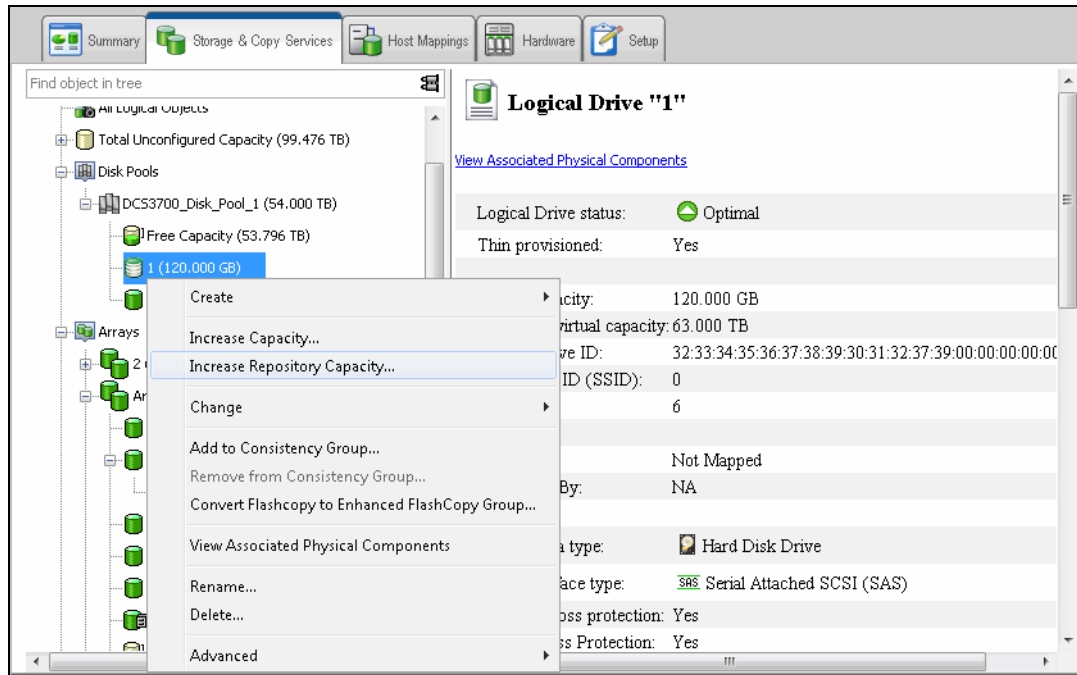


Figure 8-40 Increase Repository Capacity

After you select **Increase Repository Capacity**, the Increase Repository Capacity window opens, as shown in Figure 8-41.

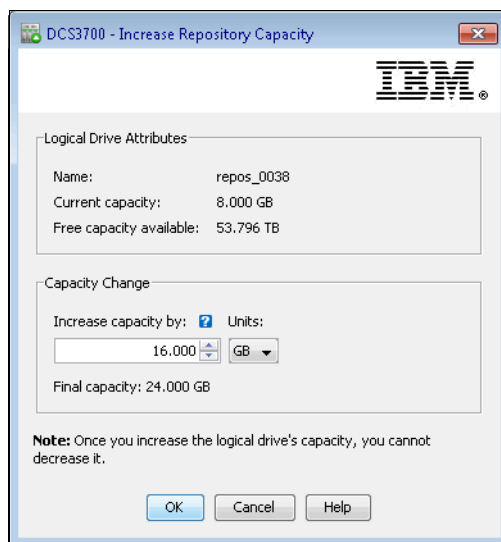


Figure 8-41 Increase Capacity Repository

**Note:** Regardless of the capacity that is specified, capacity in a disk pool is allocated in 4 GB increments. Any capacity that is not a multiple of 4 GB is allocated but not usable. To make sure that the entire capacity is usable, specify the capacity in 4 GB increments. If unusable capacity exists, the only way to regain it is to increase the capacity of the logical drive.

## 8.8.4 Change

Using this option, you can modify all the important parameters of the logical drive. Select this line and the next menu opens, as shown in Figure 8-42.

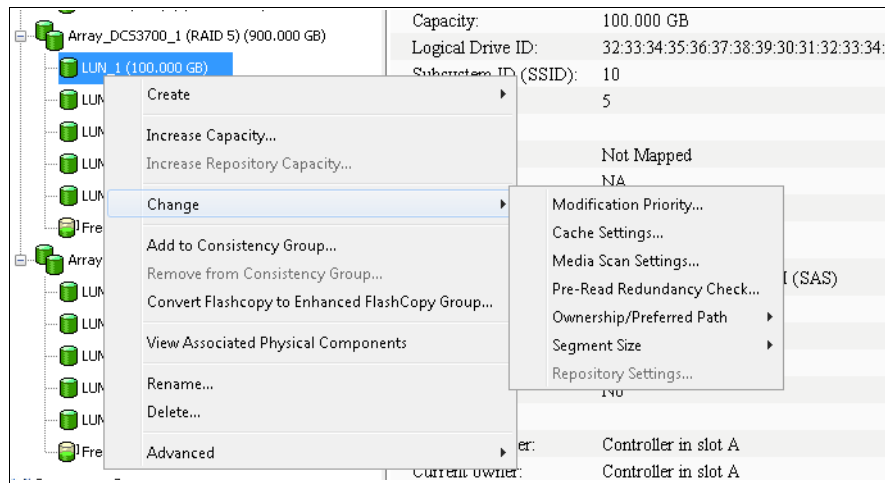


Figure 8-42 Logical Drive - Change menu

The possible changes of logical drive parameters are described in the following sections:

- ▶ Change Modification Priority
- ▶ Change Cache Settings
- ▶ Change Media Scan Settings
- ▶ Change Pre-Read Redundancy Check
- ▶ Change Ownership/Preferred Path
- ▶ Change Segment Size
- ▶ Change Repository settings

## 8.8.5 Change Modification Priority

The Change Modification Priority option defines how much processing time is allocated for operations by modifying the logical drive relative to the system performance. Here are operations that cause a logical drive modification:

- ▶ 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

These operations depend on the logical drive type (base or DDP), and not all the operations are valid for DDP.

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, complete 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 8-43.

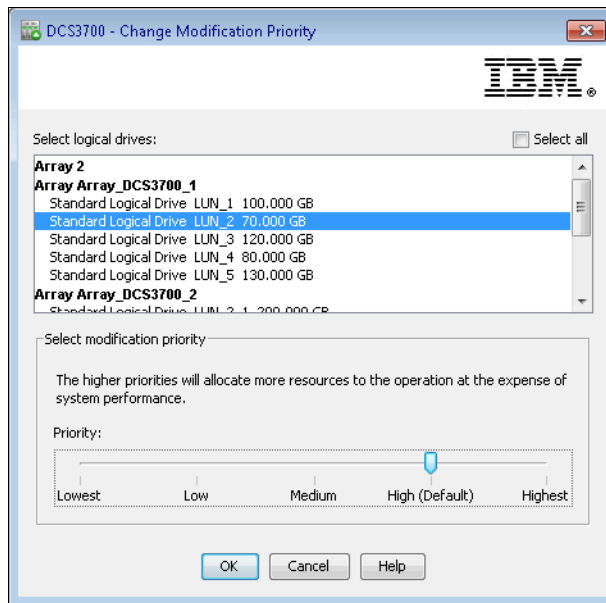


Figure 8-43 Modification priority for a logical drive

3. Click **OK** and then **Yes** in next window to confirm this change.

4. You then see the completion window, as shown in Figure 8-44.




Figure 8-44 Change Modification priority - Completed

You can select more drives in Figure 8-43 on page 263 by pressing the Ctrl or Shift key during selection or by clicking **Select All**, where *all* of the selected drives have the same modification priority after the change.

## 8.8.6 Change Cache Settings

Using this option (Figure 8-42 on page 262), you can change the cache settings on the logical drive:

- ▶ You can specify the DCS3700 global cache settings, as described in “Cache Settings” on page 209. These settings affect all arrays and logical drives that created on the storage 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 8-45 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.

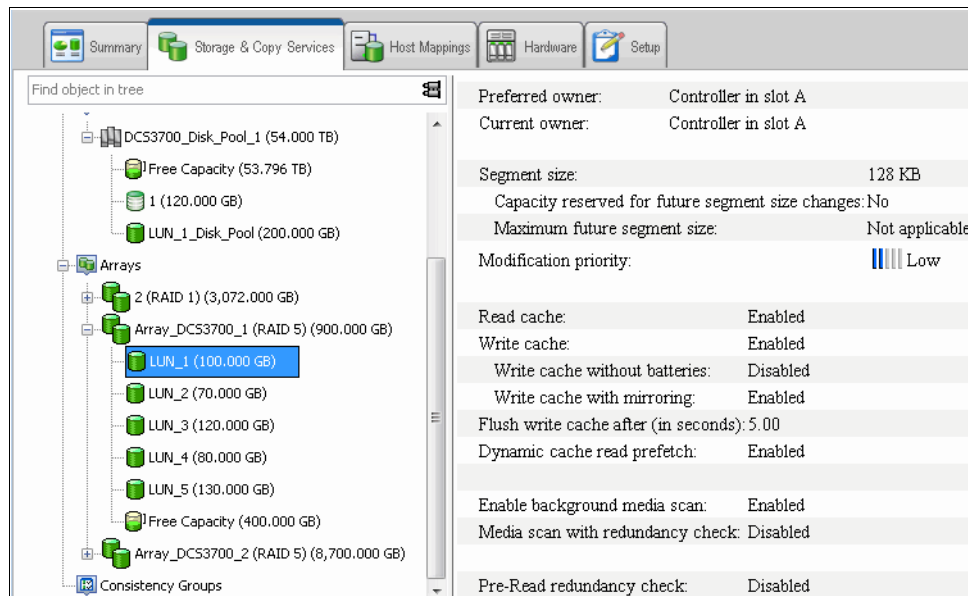


Figure 8-45 Default values for logical drives

**Note:** We recommend that you set the values manually during creation to meet 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 that is shown in Figure 8-46 on page 266 opens. The logical drive that you selected is highlighted.

You can select more logical drives in Figure 8-46 on page 266 by pressing the Ctrl or Shift key during selection or by clicking **Select All**, where all of the selected drives will have the same modification priority after the change.

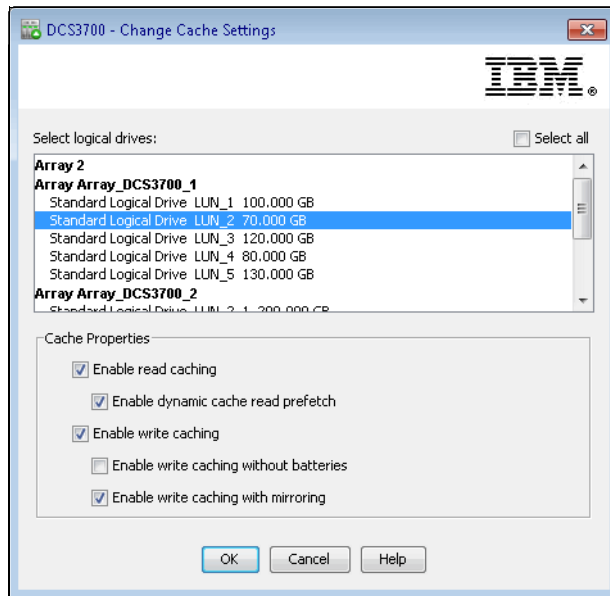


Figure 8-46 Logical Drive cache settings - defaults

These cache settings can have a large impact on the performance of the DCS3700 and on the availability of data. Performance and availability 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 alternative controller for all write data. The write cache is used only 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 the data was already sent to host. If a host initially requests data that is 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, 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 used only for read access, it might improve overall performance to disable the write cache for this logical drive.

### Write cache with mirroring

The DCS3700 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 that is disabled, but be sure that you can regenerate the data again easily because you can lose it if a controller fails.



By default, the write cache is always mirrored to the other controller, 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 DCS3700 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 written only 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. When the data is written to the cache, it must be flushed to the disks to make room for new data that is arriving in to the 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 more 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 action turns on monitoring of the I/O to the logical drive and enables the new algorithm to choose dynamically 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 tunes itself depending on the I/O characteristics. When sequential access is detected, the controller automatically starts using read ahead buffering. When random or non-sequential I/O is used, then it stops using the read ahead buffer. To disable this feature, clear the **Dynamic Cache Read Prefetch** check box for the relevant logical drive.

## 8.8.7 Change Media Scan Settings

Media scan is a background process that is enabled by default and 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 that are caused by bad sectors of the physical disks before they can disrupt normal data reads or writes. This process is sometimes known as *data scrubbing*.

Media scan is an option that is available for logical drive space that is 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 by 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 DCS3700 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 that is allocated. After the media scan completes, 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 improves 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 disk 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 that is calculated. In the case of a RAID 1 logical drive, the redundancy scan compares blocks between copies on mirrored drives.

**Note:** A media scan can resolve only media errors and data or parity mismatches. A media scan does not attempt to resolve any other error that is occurring during I/O operations.

We have seen no effect on I/O when we use a 30-day setting unless the processor is used 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 usage 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 8-3 shows the errors and describes several of the actions that the DCS3700 takes as a result of a media scan and redundancy check operations.

*Table 8-3 Media scan errors*

<b>Reported error</b>	<b>Description</b>	<b>Result</b>
Unrecovered media error	The data cannot be read on first attempt, or on any subsequent retries.	With redundancy check, data is reconstructed and scanned again. Without a redundancy check, there is no error correction.
Recovered media error	The drive cannot read the requested data on its first attempt, but succeeds on a subsequent attempt.	Data is written to the 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 ran.
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, so you can run a media scan only on logical drives that meet the following conditions:

- ▶ Optimal status.
- ▶ No modification operation is in progress.

To locate this setting, select the logical drive in the Logical pane, right-click it, and select **Change** → **Media Scan Settings**. The window that is shown in Figure 8-47 opens.

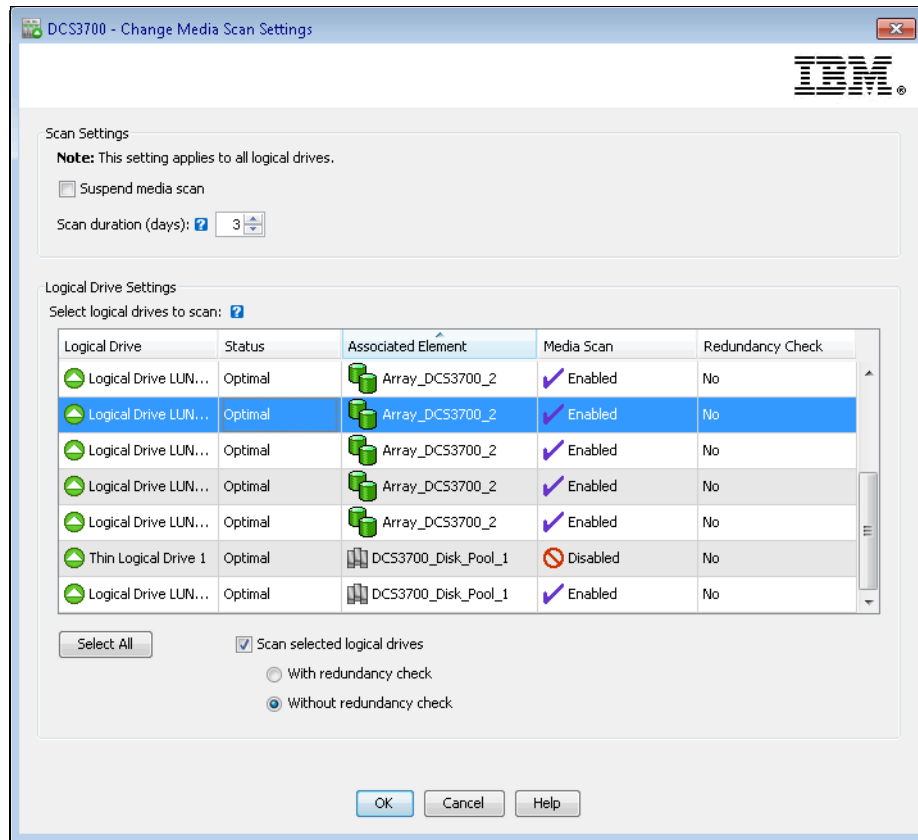


Figure 8-47 Media Scan Settings for a logical drive

The Change Media Scan Settings window has two sections. In the first, you can suspend a media scan or set the scan duration for the entire storage subsystem, and in the second, you can enable or disable scanning for selected logical drives. You can decide whether it is 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 are disabled for all logical drives after a possible scan is completed.

### Editing scan duration days

To edit the scan duration, make sure that **Suspend media scan** is clear 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. Select one or more logical drives that you want to set.
2. To enable the media scan, select the **Scan selected logical drives** check box, select **With redundancy check** or **Without redundancy check**, and click **OK**.

3. The changes are saved on the storage subsystem, and you are informed when it is complete.

### 8.8.8 Change Pre-Read Redundancy Check

You can use this option to define the capability of the DCS3700 firmware to pre-read logical drive redundancy information and determine whether the data of a 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, 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 the following items:

- ▶ 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 a 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 ignore pre-read redundancy check. If the logical drive is then migrated back to a RAID type that supports redundancy, the feature becomes 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 8.8.7, “Change Media Scan Settings” on page 268). 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 use this setting, complete the following steps:

1. Highlight the logical drive in the Logical pane tree, right-click it, and select **Change** → **Media Scan Settings**. The window that is shown in Figure 8-48 on page 272 opens.

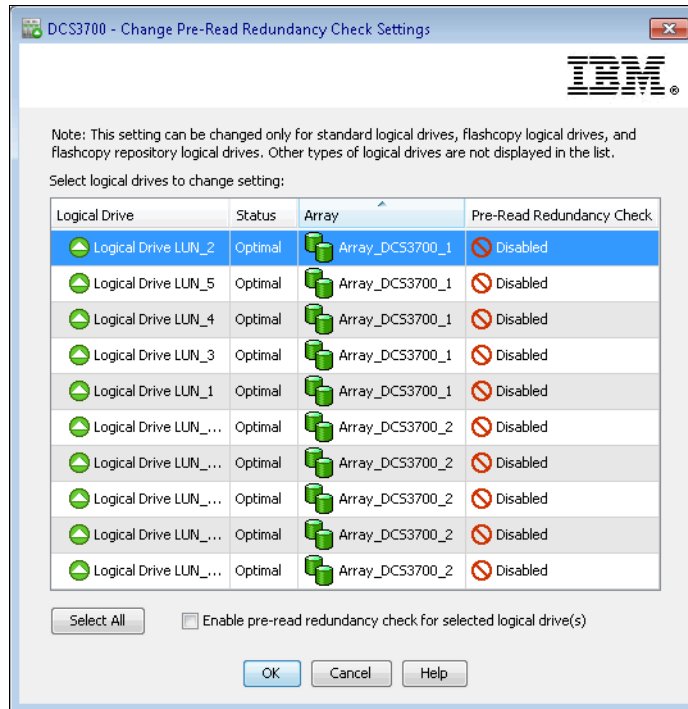


Figure 8-48 Change Pre-Read Redundancy Check settings

2. Select the logical drives that are required for a pre-read redundancy check. You can press 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 see a list of LUNs with changed settings.

## 8.8.9 Change Ownership/Preferred Path

If you are using a dual controller version of the DCS3700, each logical drive has a preferred controller of ownership. This controller normally handles all I/O requests for this particular logical drive. Each logical drive is owned by only one controller. The alternative controller takes over and handles the I/O requests only if there is a failure along the I/O path, for example, a defective host bus adapter or switch. When you define logical drives, the system normally alternates ownership between the two controllers as they are defined.

Situations can occur when all heavily stressed logical drives are 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 that is using the logical drive uses a multipath I/O driver or it will lose access to the logical drive if the ownership changes.

Balancing traffic is not always a trivial task. For example, if an application requires large disk space to be located and accessed in one chunk, it becomes harder to balance traffic by spreading the smaller volumes among controllers.

In addition, the load across controllers and logical drives 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 and 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 the guidelines for logical drive assignment and storage partitioning:

- ▶ Assign defined logical drives evenly across all controllers to balance controller usage.
- ▶ Use the manual method of creating logical drives. This allows for greater flexibility of configuration settings, such as enclosure loss protection and using both drive loops.
- ▶ If some logical drives are highly used, when possible, separate them by putting them on their own or another array. This reduces 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 owner of the logical drives. If the preferred controller must be replaced, the controller should be disabled first. This intentionally causes a failover of LUNs to the other controller and allows 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 is 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.

## 8.8.10 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 is 128 KB unless you are creating logical drives under a RAID 3, 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 items to correctly select a logical drive segment size that is 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 uses the logical drive.
  - Small I/O sizes allow for greater transaction (IOPS) rates.
  - Large I/O sizes allow 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 that are participating in the array.

### **Segment Size greater than the 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 greater than or equal to 2 x Block Size for high IOPS.

### **Segment Size equal to 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 less than 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 is 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 accessed only once, or the exact number of times if the block size is too large and you must use multiple stripes.

To calculate the best segment size in this environment, consider the RAID protection and the number of drives that are participating in the array by using the following formula:

Segment size = (Block size / X), where X is the number of drives that are 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 greater than 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 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 RAID 3 volumes, which are set to 256 KB.
- ▶ For database applications, block sizes 32 - 128 KB have been 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 16 - 64 KB.

The performance monitor can be used to evaluate how a segment size affects the workload.

**Note:** A performance testing schedule should be undertaken in the environment before you go 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 a segment size change, complete the following steps:

1. Select the logical drive in the Logical pane, right-click it, and select **Change** → **Segment Size** → **value**. The value that is highlighted in the menu is either twice or half of the current value, as shown in Figure 8-49 on page 276.

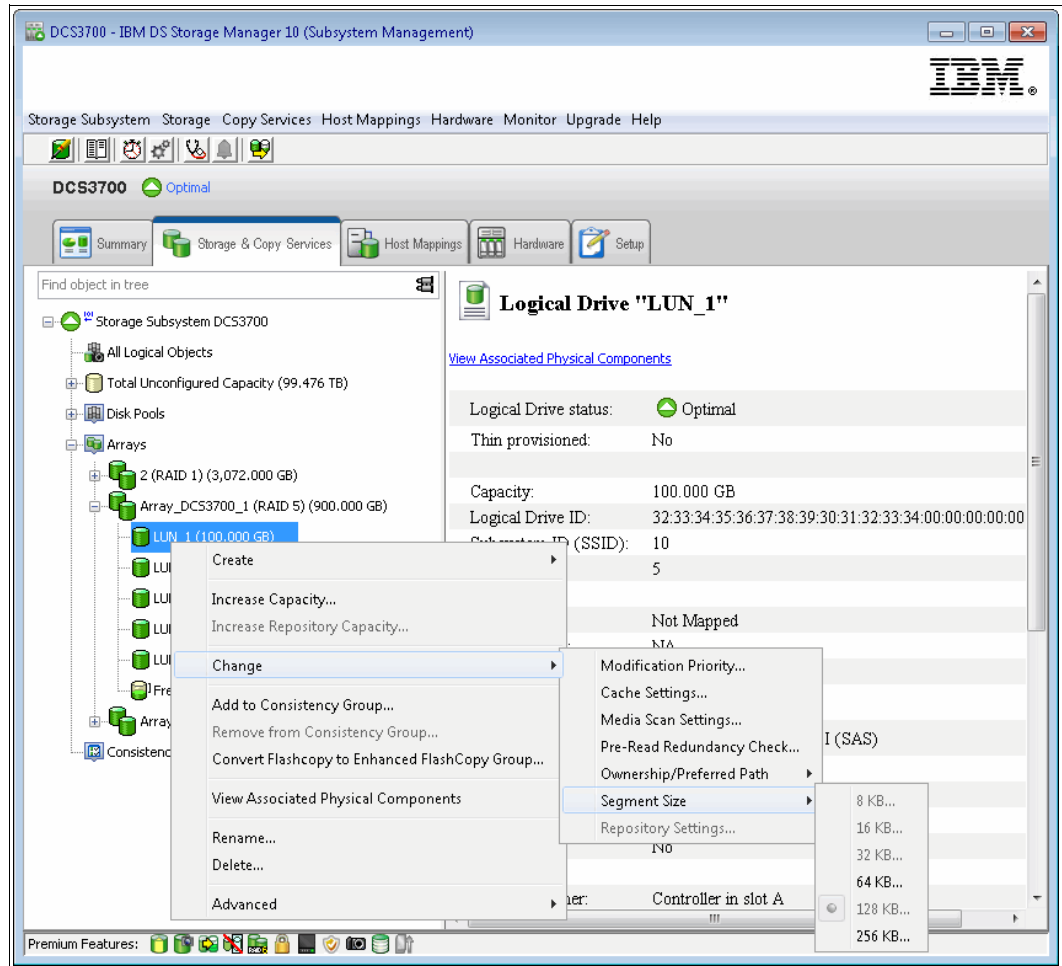


Figure 8-49 Change Segment Size menu

2. If you select a new segment size, the warning window that is shown in Figure 8-50 opens.

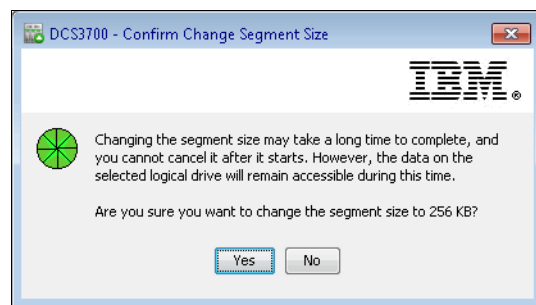


Figure 8-50 Segment Size warning

Segment size change is long-term process that can also influence the I/O operations of attached hosts. The process cannot be stopped. If you can change the 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 re-create it with your required segment size.

3. Select **Yes** and the logical drive begins to be modified. You can see the clock icon on it and the progress bar in the Properties pane, as shown in Figure 8-51.

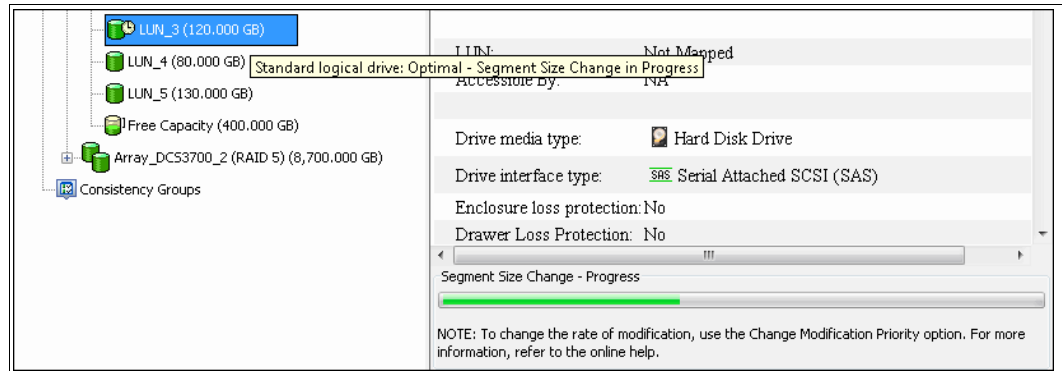


Figure 8-51 Progress bar during segment size modification

### 8.8.11 Change Repository settings

This option applies only to thin logical drives. The repository settings that you can change with this menu option are summarized in Table 8-4.

Table 8-4 Repository settings

Repository setting	Description	Possible settings
Repository expansion policy	When the consumed capacity gets close to the physical capacity, you can expand the physical capacity. You can automatically expand the physical capacity or you can do it manually.	Automatic or manual
Maximum expansion capacity	Allows you to limit the logical drive's automatic growth below the virtual capacity. This setting is available only if the repository expansion policy is set to automatic.	Capacity
Warning threshold	The level at which the storage management software sends an alert notification that the physical capacity is nearly full.	Percentage of maximum expansion capacity

To change the repository expansion policy, select either **Automatic** or **Manual**, as shown in Figure 8-52.

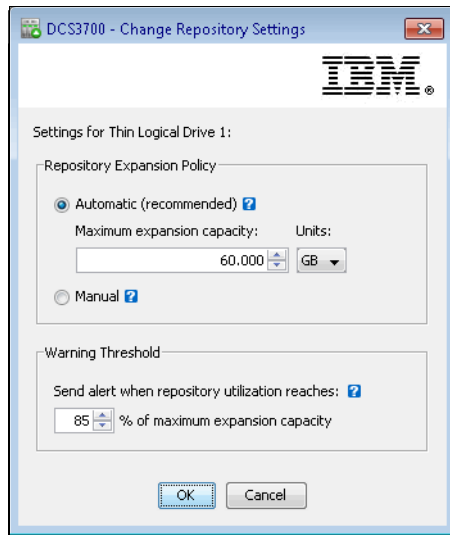


Figure 8-52 Change Repository Settings

If you select **Automatic**, you also can set a maximum expansion capacity. Use the **Maximum expansion capacity** box to indicate the capacity and the **Units** list to indicate the specific capacity units to use (MB, GB, or TB). The value for the maximum expansion capacity must be a multiple of 4 GB.

## 8.8.12 Add to Consistency Group and Remove from Consistency Group

These options are explained in 8.9, “Consistency groups” on page 285.

## 8.8.13 Convert FlashCopy to Enhanced FlashCopy Group

We are using this option to convert a FlashCopy logical drive and its associated repository to an Enhanced FlashCopy group. The system performs the following actions for each converted FlashCopy logical drive:

- ▶ Deletes the FlashCopy logical drive definition and creates a new Enhanced FlashCopy group (the new Enhanced FlashCopy group is created empty, with no Enhanced FlashCopy images).
- ▶ Converts the associated FlashCopy repository logical drive to an Enhanced FlashCopy group repository.
- ▶ Retains the same schedule (if a schedule is defined) for the new Enhanced FlashCopy group.
- ▶ Creates a read-only Enhanced FlashCopy logical drive with a Paused status. The new Enhanced FlashCopy logical drive inherits the worldwide name (WWN) and host mappings as the converted FlashCopy logical drive.

**Note:** If the number of FlashCopy copies that exist for a given base logical drive exceeds the maximum number of allowed Enhanced FlashCopy groups per base logical drive, then you should delete any excess FlashCopy copies before performing the conversion process.

The conversion process is performed on a given base logical drive, and applies to all Enhanced FlashCopy copies of a given base logical drive. If there is an online VolumeCopy related FlashCopy on the base logical drive, you must delete the online VolumeCopy job before initiating the conversion process. FlashCopy logical drives and Enhanced FlashCopy groups cannot exist on the same base logical drive; therefore, any FlashCopy logical drive that you do not select for conversion is deleted from the storage subsystem.

When you select this option, a window with important notes opens (Figure 8-53).

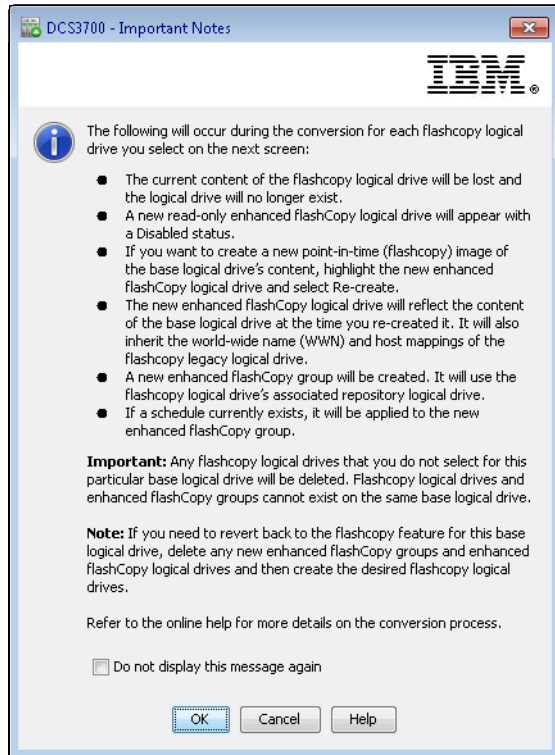


Figure 8-53 Conversion notice

After you click **OK**, a window opens, where you select the FlashCopy logical drives that you want to convert and then click **Add** to add them to the FlashCopy logical drives to convert table (Figure 8-54).

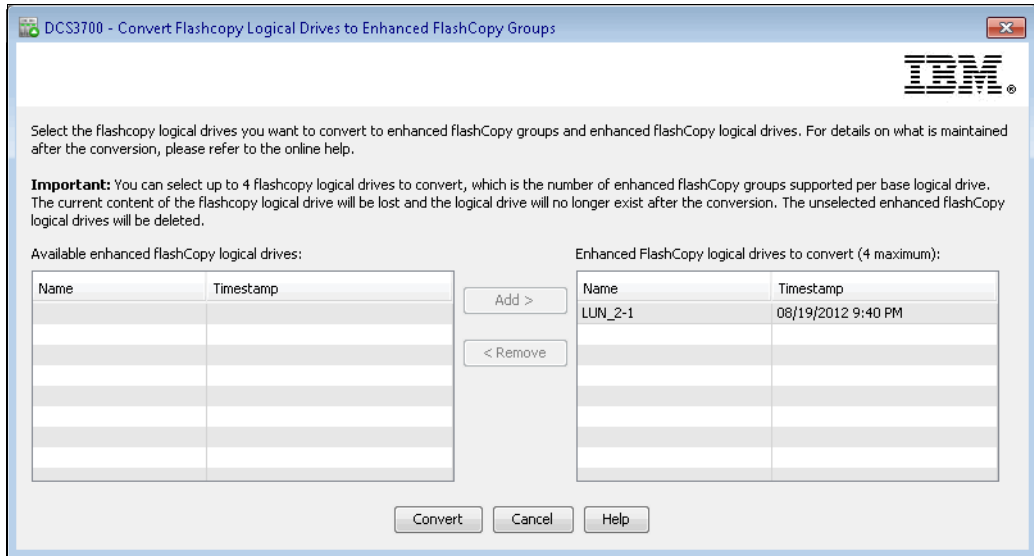


Figure 8-54 Convert table

Click **Convert** to convert the selected FlashCopy copies to Enhanced FlashCopy groups.

FlashCopy logical drives no longer appear in the logical pane of the Logical view, as shown in Figure 8-55.

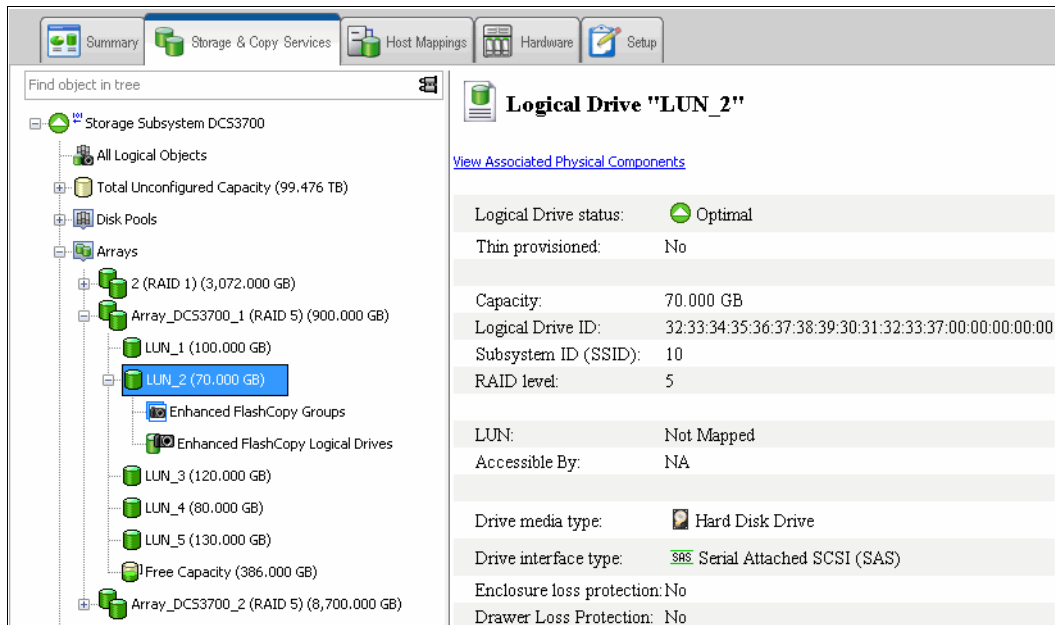


Figure 8-55 Enhanced FlashCopy groups

## 8.8.14 View Associated Physical Component

To check the disk drive configuration in the logical drive, right-click the logical drive that you want to check in the Logical pane tree and select **View Associated Physical Components**. You can also select the link on the second line of the Properties pane. A window opens and shows the components that are associated with the logical drive, such as 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 a green arrow.

## 8.8.15 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 DCS3700. Renaming has no impact on the I/O operations and the connection to the hosts. However, renaming impacts the SMcli scripts that might be used because logical drive-related commands use the name of the logical drive in their syntax.

When you delete a logical drive, the free space remains as a separate island, as shown in Figure 8-56. This free space can be used to create a logical drive of equal or smaller size than this free space. Defragmenting the array can eliminate these free capacity islands.

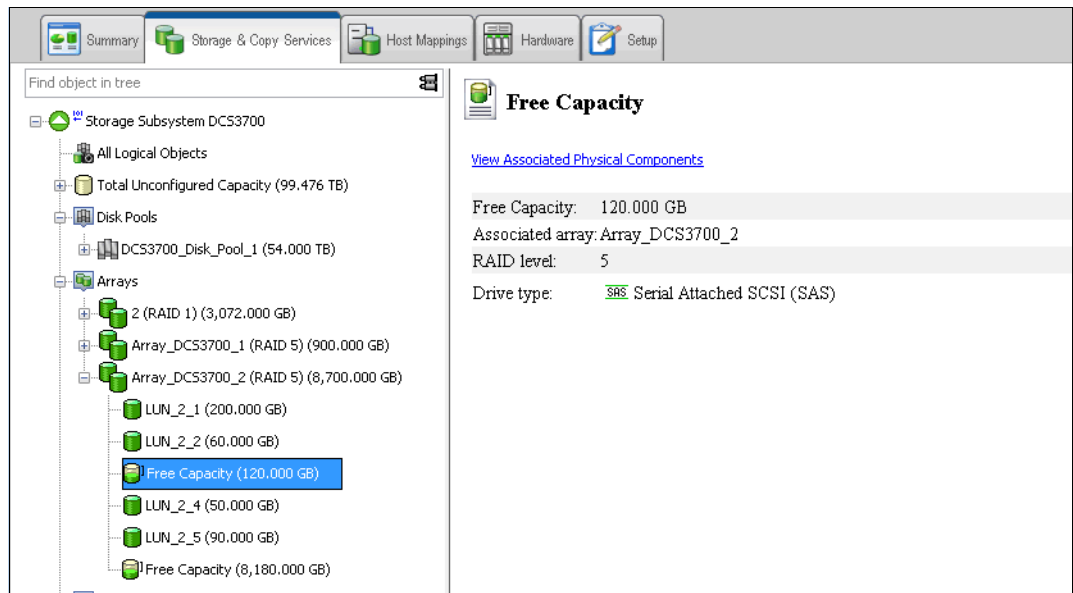


Figure 8-56 Free Capacity island in an array

## 8.8.16 Advanced

The Advanced menu has several options, as shown in Figure 8-57.




Figure 8-57 Advanced menu

## Disable T10 PI (Protection Information)

The T10 PI (Protection Information) (T10 PI) Premium Feature checks for and corrects errors that might occur as data is communicated between a host and a storage subsystem. T10 PI creates error-checking information, such as a cyclic redundancy check (CRC), and appends that information to each block of data. Any errors that might occur when a block of data is transmitted or stored is then detected and corrected by checking the data with its error-checking information.

When you select a T10 PI -enabled logical drive in the Subsystem Management window, the Disable T10 PI (Protection Information) (T10 PI) item is enabled in the Logical Drives menu. Use this option to permanently disable T10 PI on the selected logical drive.

## Initializing

A logical drive is automatically initialized when it is first created. However, the Recovery Guru might advise that you manually initialize a logical drive to recover from certain failure conditions, or if you want to reformat your logical drive from storage side.

Before you initialize a logical drive, consider the following items:

- ▶ 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.

**Attention:** Initializing a logical drive or thin logical drive erases all data from the logical drive.

### *Initializing a thin logical drive*

The thin logical drive keeps its WWN, host mappings, virtual capacity, and repository expansion policy. It also keeps the same T10 PI settings and security settings.



You can initialize a thin logical drive by using either of two options:

- ▶ **With the Same Physical Capacity:** If you want to keep the same physical capacity, the logical drive can keep its current repository logical drive, thus saving initialization time. This option is shown in Figure 8-58.

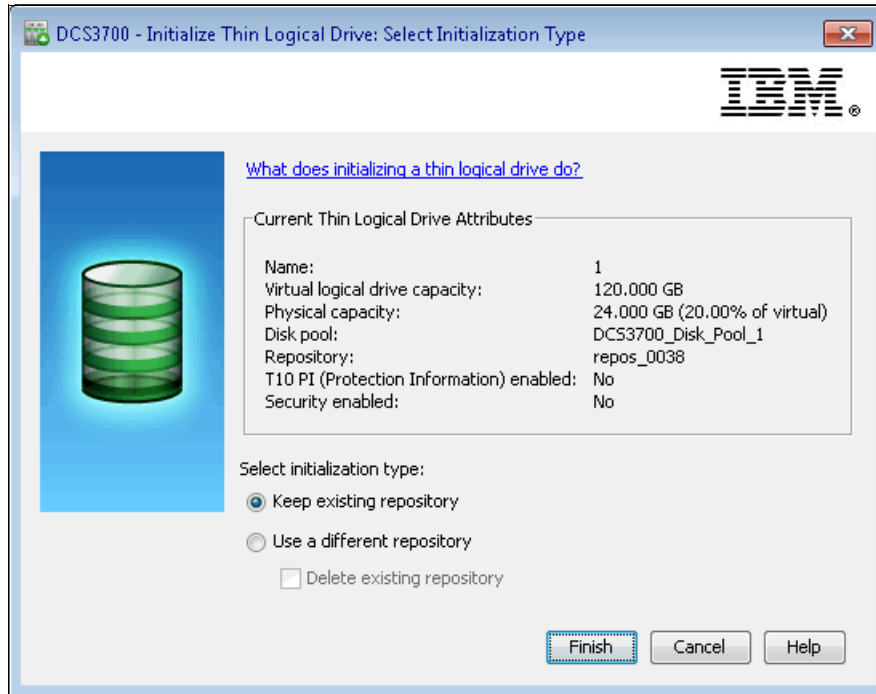


Figure 8-58 Initialize a thin logical drive with the same physical capacity

- ▶ **With Different Physical Capacity:** If you want to change the physical capacity of the thin logical drive, this option creates a new repository logical drive and you can optionally change the repository expansion policy and warning threshold. Based on whether you want to keep the current repository for future use (saving initialization time), select or clear the **Delete existing repository** check box, and click **Next**. An option with a different repository appears, as shown in Figure 8-59 on page 284.

Based on the value that you entered in the previous paragraph, the *Disk pool physical capacity candidates* table is populated with matching repositories. Select a repository from the table. Existing repositories are placed at the top of the list. The benefit of reusing an existing repository is that you can avoid the initialization process that occurs when you create one. To change the repository expansion policy or warning threshold, click **View advanced repository settings** (this setting is explained in 8.8.11, “Change Repository settings” on page 277). Click **Finish**. The Confirm Initialization of Thin Logical Drive dialog appears. To confirm that you want to initialize the thin logical drive, type yes, and click **OK**.

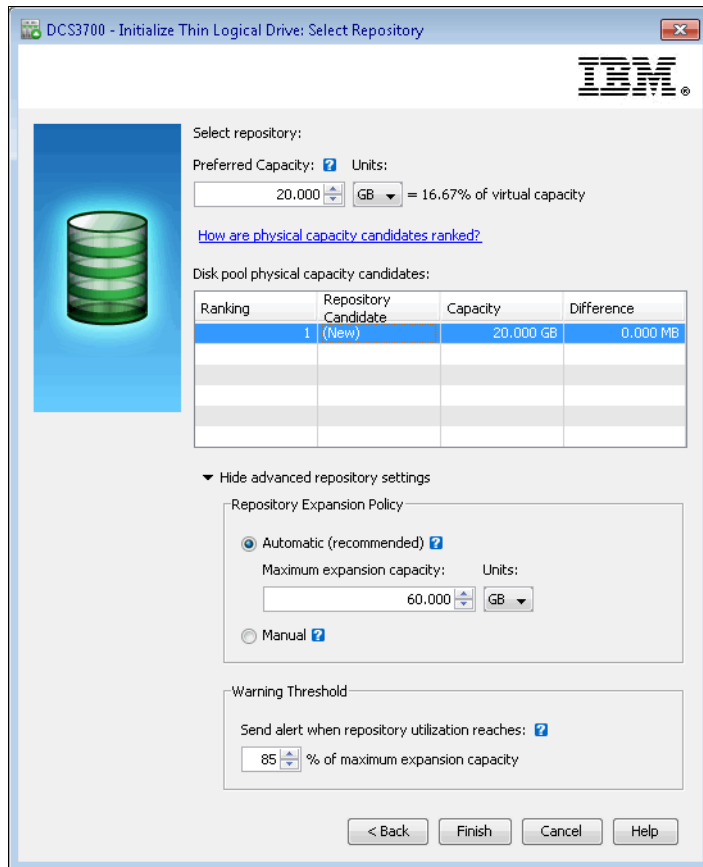


Figure 8-59 Initialize Thin Logical Drive with a different physical capacity

### Place Logical drive Online

Use the Place Logical drives Online option to return logical drives that are offline to an online status.

Cache data loss might occur after a storage subsystem experiences an interruption of power. If cache data cannot be recovered, the affected logical drives are automatically placed offline.

**Note:** You should place a logical drive online only if you are instructed to do so by the Recovery Guru or by your IBM Technical Support representative.

If a controller places a logical drive offline because of cache data loss, and that logical drive has a FlashCopy logical drive that is associated with it, the FlashCopy logical drive is marked as failed. You must remove the FlashCopy logical drive, place the original logical drive back online, and re-create the FlashCopy logical drive. If you place the original logical drive back online without completing the other steps, the FlashCopy logical drive remains in the failed status.

### Redistribute logical drive

Use this option to move logical drives back to their preferred owners. Usually, multipath drivers move logical drives from their preferred controller owners when a problem occurs along the data path between the host and the storage subsystem, but in some cases this action must be done manually.

**Risk of application errors:** If you use this option without a multipath driver on the hosts, stop I/O to the logical drives while this operation is in progress to prevent application errors.

Before you complete this procedure, consider the following items:

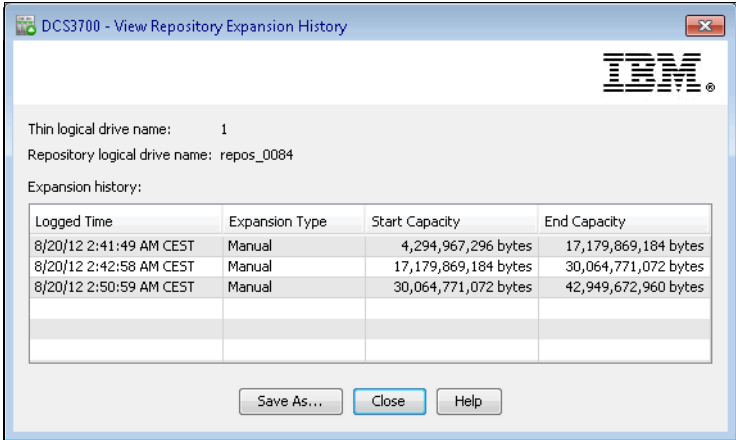
- ▶ This option is not available if all of the logical drives are owned by their preferred controllers, or if no logical drives exist on the storage subsystem.
- ▶ Under certain application host type environments, a reconfiguration of the multipath host driver might be an additional requirement. You might also need to make operating system modifications to recognize the new I/O path to the logical drives.

### View Repository Expansion History

This option applies only to thin logical drives. The controller maintains a history of the 16 most recent expansion operations for each thin logical drive on the storage subsystem.

As shown in Figure 8-60, the following information is tracked for each operation:

- ▶ Whether the expansion was an automatic or manual operation
- ▶ The date and time that the expansion occurred
- ▶ The before and after capacity of the repository



Logged Time	Expansion Type	Start Capacity	End Capacity
8/20/12 2:41:49 AM CEST	Manual	4,294,967,296 bytes	17,179,869,184 bytes
8/20/12 2:42:58 AM CEST	Manual	17,179,869,184 bytes	30,064,771,072 bytes
8/20/12 2:50:59 AM CEST	Manual	30,064,771,072 bytes	42,949,672,960 bytes

Figure 8-60 View Repository Expansion History

## 8.9 Consistency groups

With firmware Version 07.8x.xx.xx, IBM introduces a new feature in the storage manager: *consistency groups*. Consistency groups are available only on the IBM System Storage DS3500 and DCS3700.

A consistency group is a feature of the Enhanced FlashCopy Premium Feature. It is a collection of up to 32 base logical drives in your storage subsystem. These base logical drives, which are the sources of the Enhanced FlashCopy images, are referred to as member logical drives of a consistency group.

If you want to perform the same Enhanced FlashCopy image operations frequently on multiple logical drives, create a consistency group. A consistency group is a group of logical drives that you treat as one entity. An operation that you perform on the consistency group is performed simultaneously on all of the logical drives in the consistency group. Enhanced FlashCopy image operations that you can perform on a consistency group include creating, scheduling, and rolling back.

Each logical drive that belongs to a consistency group is referred to as a *member logical drive*. When you add a logical drive to a consistency group, the system automatically creates a new Enhanced FlashCopy group that corresponds to this member logical drive. A member logical drive is either a standard logical drive in a disk pool or an array, or a thin logical drive in a disk pool. A member logical drive must not have any FlashCopy logical drives defined. If any FlashCopy logical drives are defined in a member logical drive, then delete them before that logical drive becomes a part of a consistency group. This requirement does not apply to Enhanced FlashCopy logical drives. A consistency group member logical drive might have its own Enhanced FlashCopy images in addition to the Enhanced FlashCopy images as part of the consistency group.

If the Enhanced FlashCopy Premium Feature is enabled, you can have the following consistency group child nodes:

- ▶ Consistency Group: This node represents a grouping node, which includes all the child nodes that are created for this consistency group. Expand this node to see the child nodes.
- ▶ Enhanced FlashCopy Images: This node represents a collection of logical point-in-time images of the member logical drives of a consistency group.
- ▶ Member logical drives: This node is a collection of the logical drives that are members of this consistency group.
- ▶ Enhanced FlashCopy logical drives: This node represents the Enhanced FlashCopy images of the member logical drives that are visible to a host.

Before you create a consistency group, consider the following items:

- ▶ If the base logical drive is on a standard array, the repository members for any associated consistency group can be on either a standard array or a disk pool. If a base logical drive is on a disk pool, all repository members for any associated consistency group must be on the same disk pool as the base logical drive.
- ▶ You cannot create a consistency group on a failed logical drive.
- ▶ A consistency group contains one Enhanced FlashCopy group for each logical drive that is a member of the consistency group. You cannot individually manage an Enhanced FlashCopy group that is associated with a consistency group. Instead, you must perform the manage operations (create Enhanced FlashCopy image, delete Enhanced FlashCopy image or Enhanced FlashCopy group, and rollback Enhanced FlashCopy image) at the consistency group level.
- ▶ If you attempt to create a consistency group Enhanced FlashCopy image, the operation might remain in a Pending state because of the following conditions:
  - The base logical drive that contains this consistency group Enhanced FlashCopy image is a member of an enhanced Global Mirror group.
  - The base logical drive is in a synchronizing operation. The consistency group Enhanced FlashCopy image creation completes when the synchronization operation is complete.

- ▶ The consistency group repository logical drive must have the same T10 PI and quality of service (QoS) settings as the associated base logical drive for the consistency group member logical drive. For example, if a base logical drive that is a member logical drive in a consistency group is T10 PI enabled, then the associated consistency group repository logical drive must be T10 PI enabled.

**Note:** There is a maximum limit to the number of member logical drives that can belong to a consistency group (depending on the configuration). For a DCS3700, it is 32.

A consistency group Enhanced FlashCopy image, which is sometimes referred to as *consistency group PiT*, is a logical point-in-time image of the content of each member logical drive in the consistency group at the same point in time. This action creates synchronized Enhanced FlashCopy images of all the logical drives in a consistency group. Consistency group Enhanced FlashCopy images are stored consecutively based on creation time, with the oldest Enhanced FlashCopy image at the top of the list.

Before creating an Enhanced FlashCopy image for each member in a consistency group, the controllers suspend all pending I/O operations for each member logical drive of the consistency group. If the controller software cannot create an Enhanced FlashCopy image for any one member logical drive because I/Os cannot be suspended, no new consistency group Enhanced FlashCopy image is created.

The basic structure of consistency group is shown on Figure 8-61.

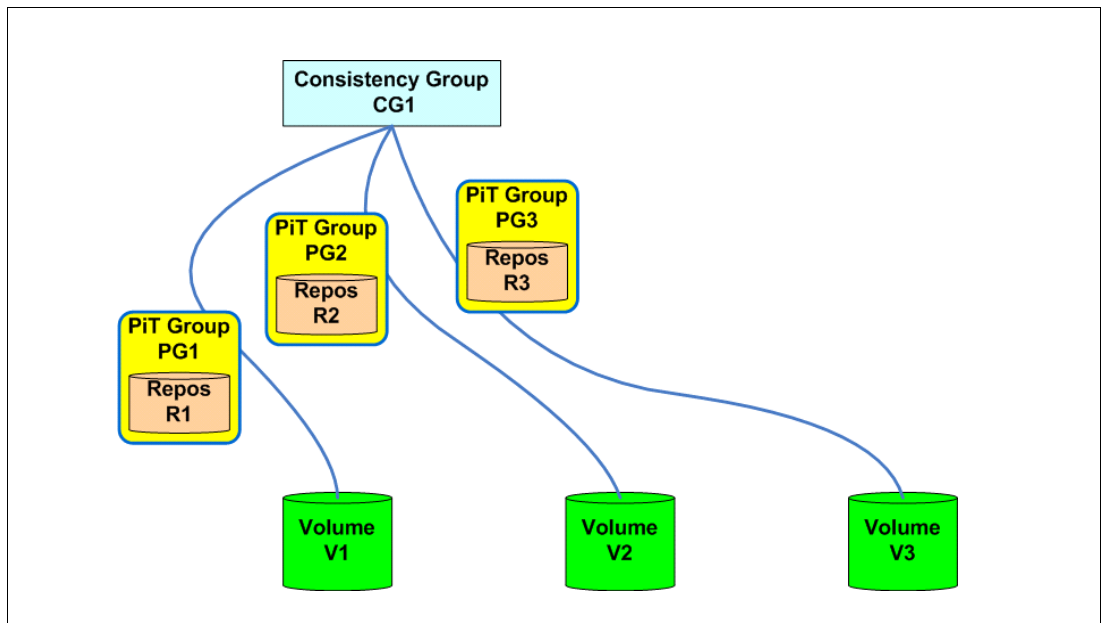


Figure 8-61 Basic structure of a consistency group

After Enhanced FlashCopy images are created for a consistency group, if you add a member logical drive, the existing member logical drives have a different number of stored Enhanced FlashCopy images compared to the newly added member logical drive. If you roll back an Enhanced FlashCopy image that was taken before new member logical drives were added to a consistency group, the rollback affects only the member logical drives that are in the consistency group only before the new logical drives are added to the consistency group.

The consistency group Enhanced FlashCopy image cannot be mapped to a host for I/O operations because the consistency group Enhanced FlashCopy image is used to save only the original data from the member logical drive. A consistency group Enhanced FlashCopy logical drive, which is sometimes referred to as consistency group Enhanced FlashCopy view or consistency group PIT view, must be created for each consistency group Enhanced FlashCopy image that is mapped to the host. The consistency group Enhanced FlashCopy logical drive repository is used to save any subsequent modifications that are made by the host application to the member logical drive, without affecting the referenced consistency group Enhanced FlashCopy image.

**Note:** As in the case with the Enhanced FlashCopy image logical drive, the consistency group Enhanced FlashCopy logical drive does not need to have the associated repository drive created if the host access is read-only. If the host access is read/write, a consistency group Enhanced FlashCopy logical drive repository logical drive must be created along with the consistency group Enhanced FlashCopy logical drive.

### 8.9.1 Creating a consistency group

As mentioned before, a consistency group is simultaneous Enhanced FlashCopy copies of multiple logical drives, thus ensuring consistent copies of a group of logical drives. Each logical drive that belongs to a consistency group is referred to as a member logical drive. When you add a logical drive to a consistency group, the system automatically creates an Enhanced FlashCopy group that corresponds to this member logical drive.

To create a consistency group, select **Consistency Groups** in the Logical pane and select **Create** as shown in as shown in Figure 8-62.

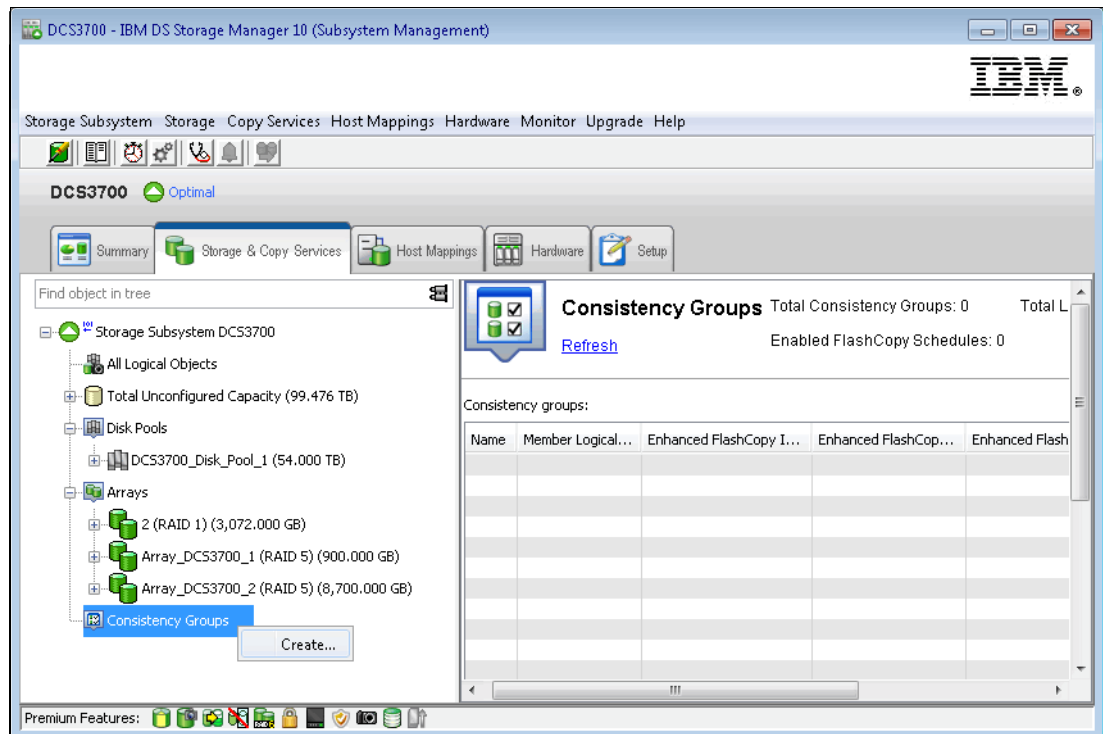


Figure 8-62 Create a consistency group

The Consistency Group Settings window opens (Figure 8-63). In the Consistency group name field, enter a unique name (30 characters maximum) that best describes the member logical drives that you want to add for this group.

**Note:** By default, the consistency group name is shown in the name text box as CG\_sequence-number.

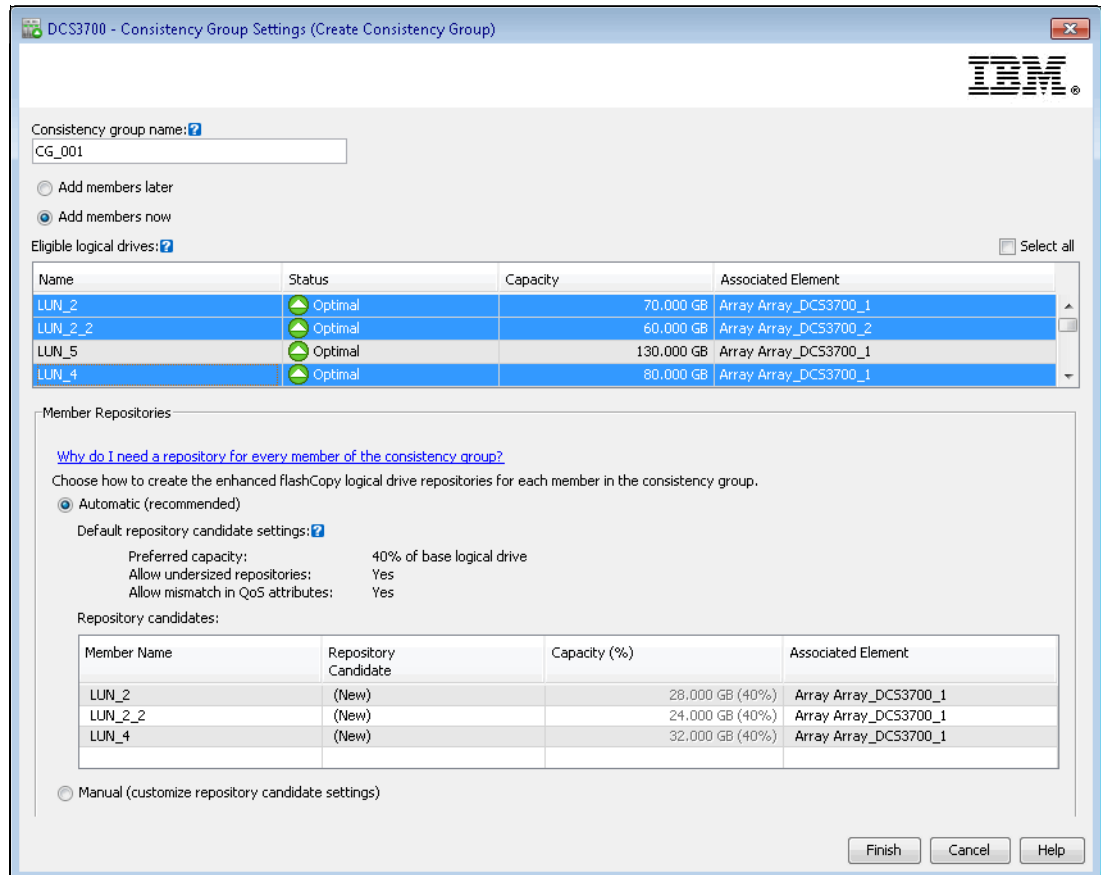


Figure 8-63 Consistency Group Settings

If you select **Add members later**, this is last step and you click **Finish** to complete the Consistency Group creation.

If you select **Add members now**, select, from the eligible member logical drives, the logical drives that you want to add as members to the consistency group, as shown on Figure 8-63.

**Note:** To be an eligible member of a consistency group, a logical drive must not be in the Failed status and must not have a maximum number of Enhanced FlashCopy groups that are associated with it.

It is a preferred practice to select **Automatic (recommended)** for each member of a consistency group and click **Finish** to complete consistency group creation. If you select **Manual (customize repository candidate settings)**, the window that is shown in Figure 8-64 opens.

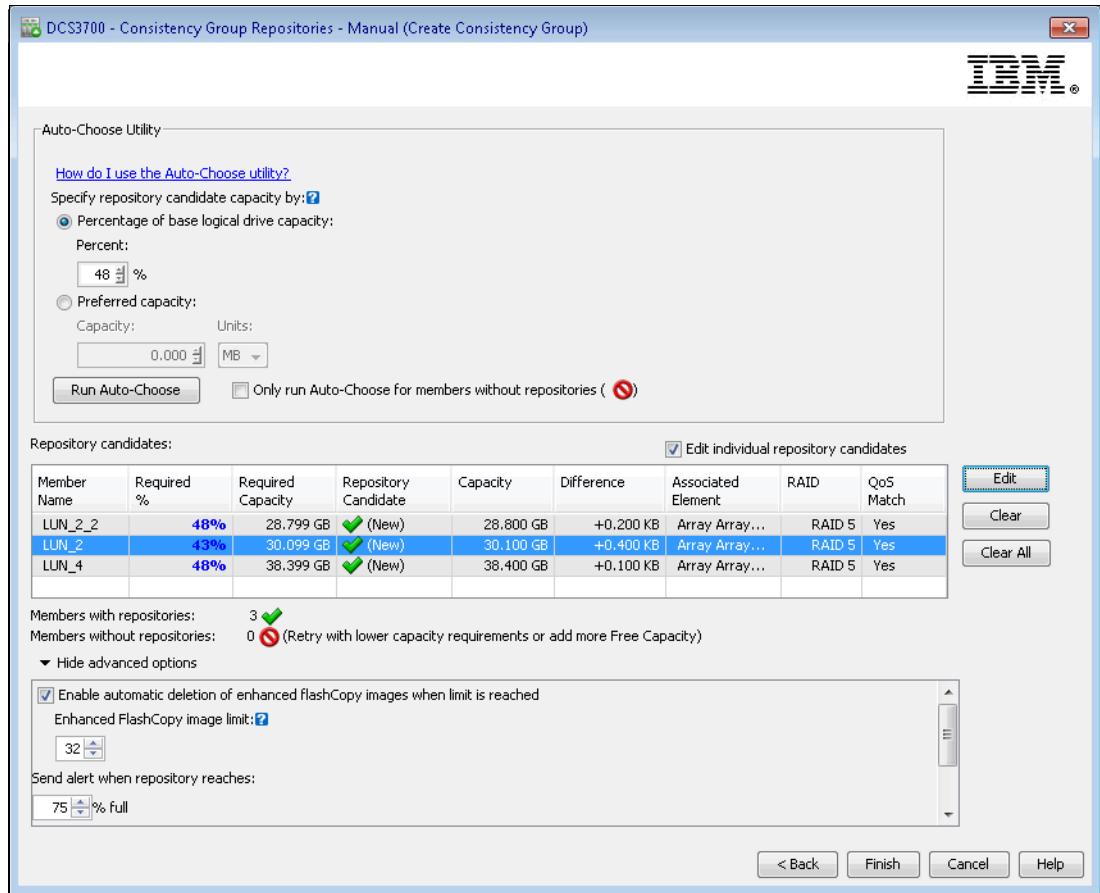


Figure 8-64 Manual (Create Consistency Group) settings

**Note:** The manual method is considered advanced and only those who understand drive redundancy and optimal drive configurations should use this method.



Choose how you want to filter the repository candidates for each member logical drive in the consistency group, based on either a percentage of the base logical drive capacity or by preferred capacity. Select one of the options that are described in Table 8-5.

Table 8-5 Specify the repository candidate's capacity

Filter options	Description
Percentage of base logical drive capacity	<p>Use the Percent spinner box to set the capacity percentage of the base logical drive.</p> <p>The default setting for the consistency group repository capacity is 20% of the capacity of the base, and usually this capacity is sufficient. The capacity that is needed varies, depending on the frequency and size of I/O writes to the base logical drive and how long you must keep the repository.</p> <p>In general, choose a larger capacity for the repository if one or both of these conditions exist:</p> <ul style="list-style-type: none"> <li>▶ You intend to keep the repository for a long period.</li> <li>▶ A large percentage of data blocks will change on the base logical drive during the life of the repository due to heavy I/O activity.</li> </ul> <p>Use historical performance monitor data or other operating system utilities to help you determine typical I/O activity to the base logical drive.</p>
Preferred capacity	<p>Use the Capacity spinner box to indicate the capacity of the repository and the Units spinner box to indicate the specific capacity units to use (MB, GB, or TB).</p> <p>The capacity value, which is shown by default, is the equivalent percentage of the base logical drive capacity. For example, if the default entry for percentage is 20%, the default entry for the preferred capacity is 20% * base logical drive capacity.</p>

Select the **Edit individual repository candidates** check box if you want to edit repository candidates for the member logical drives. Then, select the repository, from the repository candidates table, that you want to use for each member logical drive in the consistency group. Select a repository candidate that is closest to the capacity you specified.

- ▶ The repository candidates table shows both new and existing repositories that can be used for each member logical drive in the consistency group based on the value that you specified for percentage or the value you specified for preferred capacity.
- ▶ By default, the system displays the repositories for each member logical drive of the consistency group by using a value of 20% of the member logical drive's capacity. It filters out undersized repository candidates, and those with different QoS attributes. If appropriate candidates are not returned by using these settings, you can click **Run Auto-Choose** to provide automatic candidate recommendations.
- ▶ The Difference column shows the mathematical difference between your selected capacity and the actual capacity of the repository candidate. If the repository candidate is new, the system uses the exact capacity size that you specified and displays zero in the difference column.

To edit an individual repository candidate, select the candidate from the repository candidates table and click **Edit** to modify the capacity settings for the repository and then click **OK** to continue with the repository creation process.

Select **Advanced options** and then accept or change the default settings that are shown in Table 8-6 on page 292, as appropriate.

Table 8-6 Repository capacity advanced settings

Advanced settings	Description
Enable automatic deletion of Enhanced FlashCopy images when limit is reached	Enables automatic deletion of the oldest Enhanced FlashCopy image for the Enhanced FlashCopy group when the Enhanced FlashCopy image limit is reached.
Enhanced FlashCopy image limit	The Enhanced FlashCopy image limit is a configurable value that specifies the maximum number of Enhanced FlashCopy images that are allowed for a consistency group Enhanced FlashCopy group.
Send alert when repository reaches	A configurable value that determines when a warning is triggered when the capacity of a consistency group repository reaches the defined percentage.
Policy for full repositories	Select which policy to use when the capacity of the consistency group repository reaches its maximum defined percentage.

Click **Finish** to create the repository with the capacity settings and advanced settings that you specified.

The created consistency group is shown in the logical pane, and you can see members of that group in the right pane, as shown in Figure 8-65.

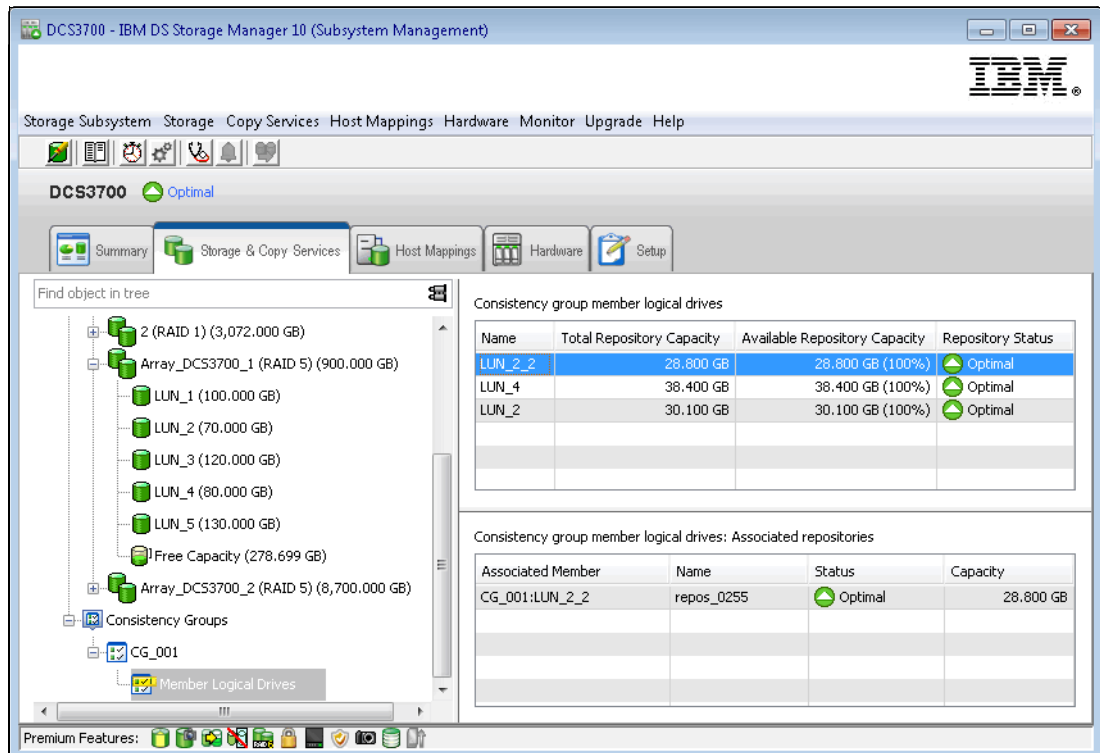


Figure 8-65 Consistency group member logical drives

## 8.9.2 Work with Consistency Group

After you create the consistency group, you have several options to manage that group. If you select a consistency group in logical pane and right-click, the menu that is shown in Figure 8-66 opens.



Figure 8-66 Consistency Group menu

### Consistency Group Enhanced FlashCopy Image

A consistency group Enhanced FlashCopy image is a logical point-in-time image of the content of each member logical drive in the consistency group.

The consistency group Enhanced FlashCopy image provides an Enhanced FlashCopy of all the member logical drives at the same point in time. This action creates a synchronized Enhanced FlashCopy of all the logical drives and is ideal for applications that span multiple logical drives (for example, a database application that has the logs on one logical drive and the database on another logical drive).

Here are the characteristics of consistency group Enhanced FlashCopy images:

- ▶ If a logical drive is a member of a consistency group, then the system creates an Enhanced FlashCopy group for that base logical drive, which counts towards the maximum allowable number of Enhanced FlashCopy groups per base logical drive.
- ▶ There is a maximum allowable number of Enhanced FlashCopy images that can be added to a consistency group (depending on your configuration).

Here are the items that you can select for this option:

- ▶ Create: The system takes a copy of the associated member logical drives and then displays the Consistency Group Enhanced FlashCopy Image Successfully Created dialog, as shown in Figure 8-67.

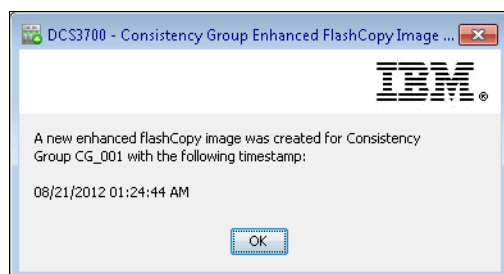


Figure 8-67 Consistency Group Enhanced FlashCopy Image Successfully Created dialog

- **Create/Schedule:** The system schedules a time of copy action on the associated member logical drives, as shown in Figure 8-68.

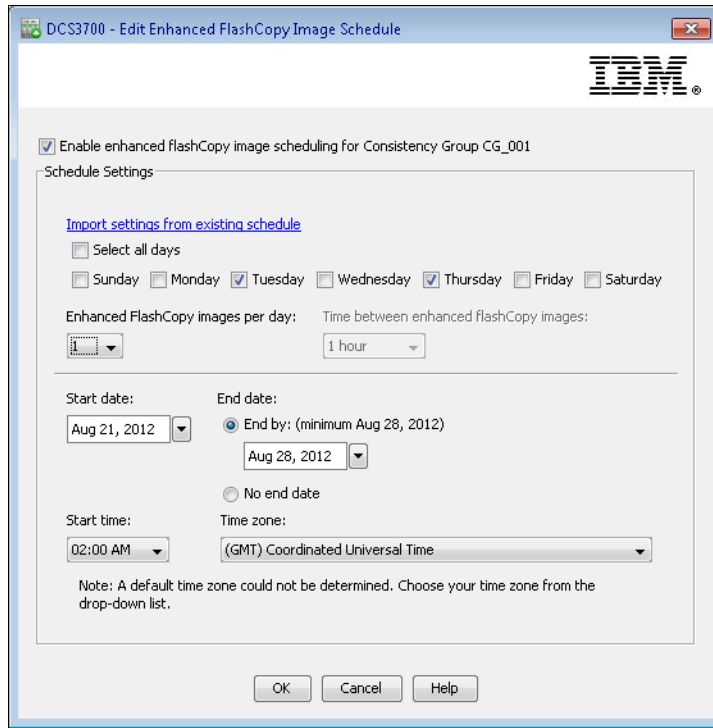


Figure 8-68 Edit Enhanced FlashCopy Image Schedule

### Create Consistency Group Enhanced FlashCopy Logical Drive

A consistency group Enhanced FlashCopy logical drive comprises multiple Enhanced FlashCopy logical drives to provide host access to an Enhanced FlashCopy image that was taken for each selected member logical drive at the same moment in time.

The consistency group Enhanced FlashCopy logical drive can be designated as either read-only or read/write. Read/write consistency group Enhanced FlashCopy logical drives require a repository for each member logical drive that you select in the wizard to save any subsequent modifications that are made by the host application to the base logical drive without affecting the referenced Enhanced FlashCopy image. Each member repository is created at the same time the consistency group Enhanced FlashCopy logical drive is created.

**Note:** A read/write Enhanced FlashCopy logical drive requires an overall repository. The overall repository is created initially with one individual repository logical drive. However, the overall repository can contain multiple repository logical drives in the future for expansion purposes, as shown in Figure 8-69.

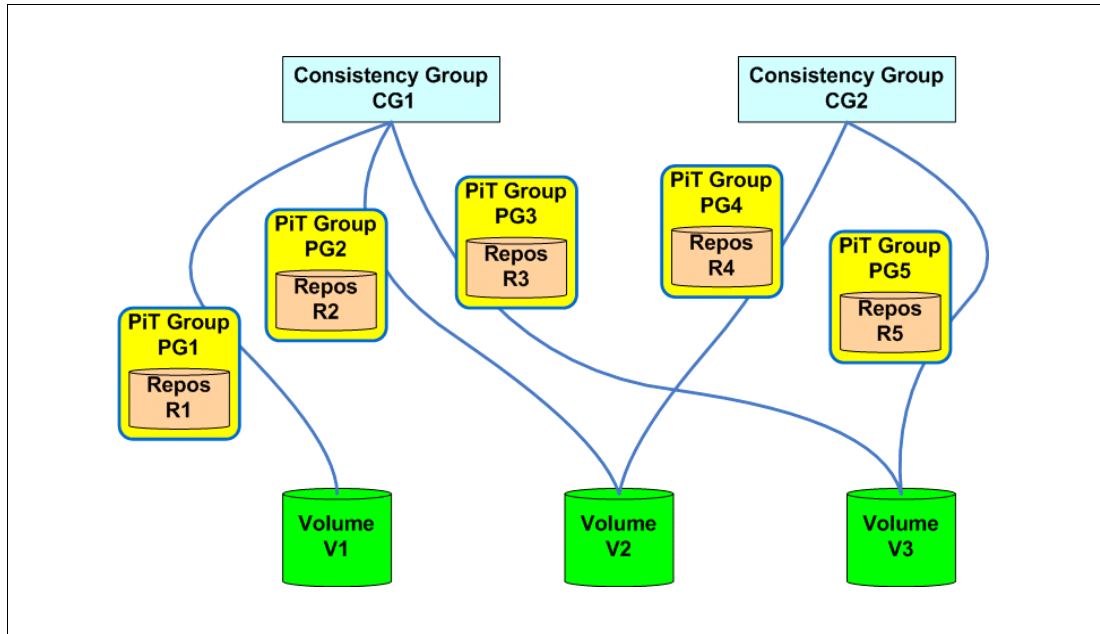


Figure 8-69 Consistency group - member overlap

Here are the prerequisites for an Enhanced FlashCopy logical drive:

- ▶ The Enhanced FlashCopy Premium Feature must be enabled on the storage subsystem.
- ▶ The consistency group must contain at least one member logical drive before you can create a consistency group Enhanced FlashCopy logical drive.

To create a consistency group Enhanced FlashCopy logical drive, complete the following steps:

1. Choose the consistency group Enhanced FlashCopy image for which you want to create an Enhanced FlashCopy logical drive. Do one of the following actions (Figure 8-70 on page 296):
  - Click **An existing Enhanced FlashCopy image** and then select an Enhanced FlashCopy image from the consistency group Enhanced FlashCopy images table and click **Next**.
  - Click **A new Enhanced FlashCopy image** and then select an Enhanced FlashCopy group from the existing Enhanced FlashCopy group table and click **Next**.

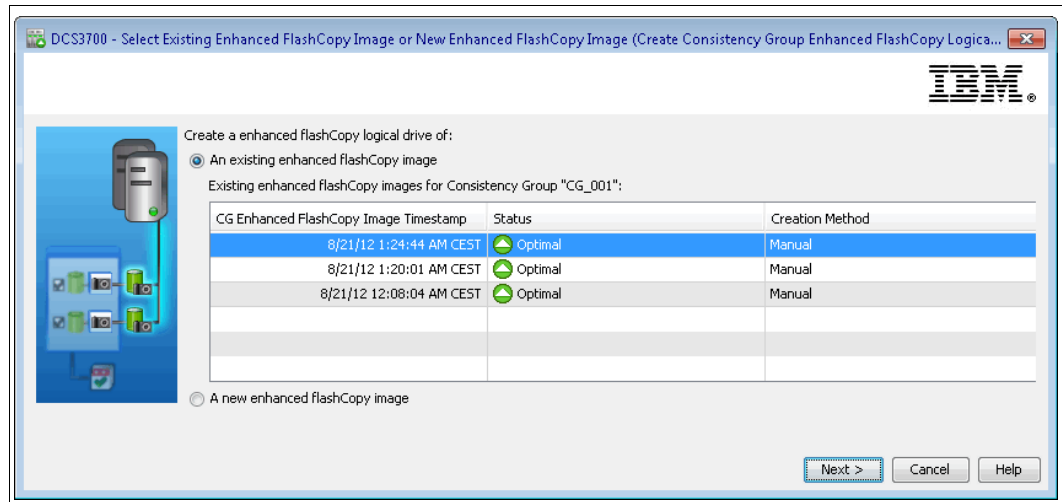


Figure 8-70 Create Consistency Group Enhanced FlashCopy Logical Drive

The system displays the Consistency Group Enhanced FlashCopy Logical Drive Settings dialog.

- As shown in Figure 8-71, in the Consistency group Enhanced FlashCopy logical drive name field, enter a unique name (30 character maximum) that best describes the consistency group that is selected for this Enhanced FlashCopy image, for example, Accounting Data.

By default, the consistency group Enhanced FlashCopy logical drive name is shown in the name text box as [consistency-group-name] - SV + *sequence-number*, where SV (Enhanced FlashCopy logical drive) is the appended suffix and sequence-number is the chronological number of the Enhanced FlashCopy logical drive relative to the consistency group.




Figure 8-71 Create Consistency Group Enhanced FlashCopy Logical Drive Settings

- In the **Map to host** drop-down menu, specify how you want to map the host for each Enhanced FlashCopy logical drive that is created for a selected member logical drive by using the following options. This map attribute is applied to every member logical drive you select in the consistency group.
  - Map Now to Default Group: The logical drive is automatically assigned a logical unit number (LUN) and is accessible by any hosts that are connected to the storage subsystem.
  - Map Later: The logical drive is not assigned a LUN and is not accessible by any hosts until you go to the Host Mappings tab and assign a specific host and LUN to this logical drive.
  - Select a specific host: You can select a specific host or host group from the list. This option is available only if Storage Partitioning is enabled.

4. Choose how to grant host access to each selected member logical drive's Enhanced FlashCopy logical drive. Complete one of the following tasks:
  - Select **Read/Write** to provide the host application with WRITE access to a copy of the data that is contained in the Enhanced FlashCopy image. A read/write Enhanced FlashCopy logical drive requires an associated repository.
  - Select **Read Only** to provide a host application with READ access to a copy of the data that is contained in the Enhanced FlashCopy image, but without the ability to modify the Enhanced FlashCopy image. A Read-Only Enhanced FlashCopy logical drive does not have an associated repository.
5. Select each member logical drive in the consistency group for which you want to create an Enhanced FlashCopy logical drive. You can select the **Select all** check box to create an Enhanced FlashCopy logical drive for each member logical drive that is displayed in the select members table. (If you selected Read-Only host access in step 4, you can skip this step because repositories are not required for Read-Only Enhanced FlashCopy logical drives.)
6. Choose how you want to create the Enhanced FlashCopy logical drive repositories for each member in the consistency group. Complete one of the following tasks:
  - Select **Automatic** and click **Finish** to create each Enhanced FlashCopy logical drive repository with the default capacity settings. This option is the recommended one.
  - Select **Manual** and click **Next** to define the properties for each Enhanced FlashCopy logical drive repository, and then click **Finish** to continue with the Enhanced FlashCopy logical drive creation process. You can click **Edit individual repository candidates** to manually edit a repository candidate for each member logical drive.



- In the Logical pane of Storage Manager, click **Enhanced FlashCopy Logical Drive**, and in the Properties pane, you see information about the Enhanced FlashCopy logical drive (Figure 8-72).

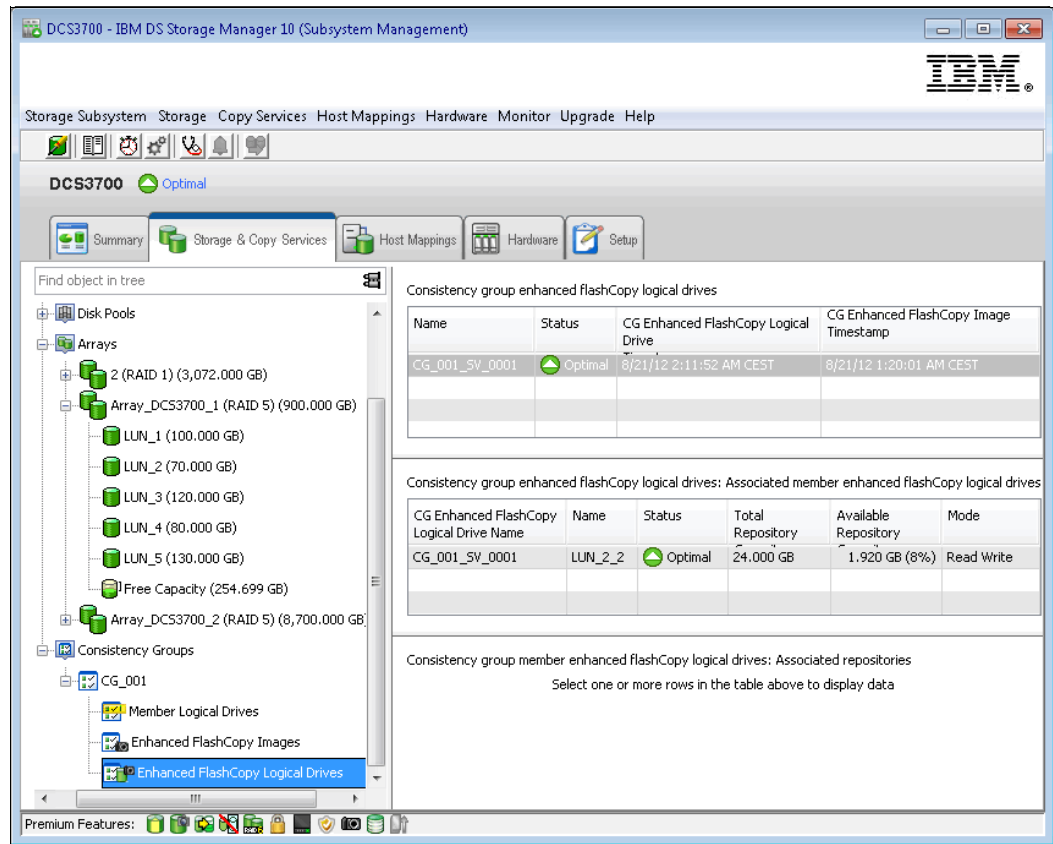


Figure 8-72 Enhanced FlashCopy Logical Drive

## Add Member Logical Drive

Use this option to add a member logical drive to an existing consistency group.

**Note:** When a member is added to a consistency group, you must also add a repository logical drive.

As described in 8.9.1, “Creating a consistency group” on page 288, only eligible logical drives display (Figure 8-73).

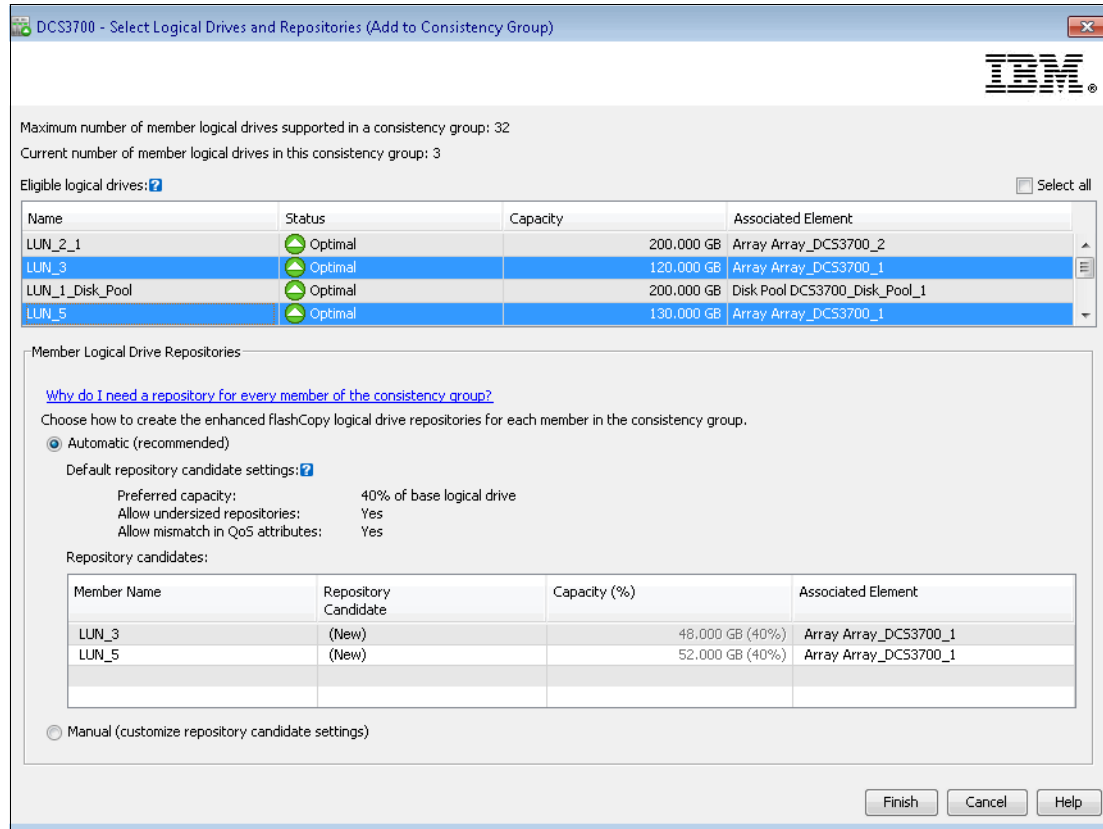


Figure 8-73 Add Member logical drive to a consistency group

## Remove Member Logical Drive

Use this option to remove a member logical drive from an existing consistency group, as shown in Figure 8-74. When you remove a member logical drive from a consistency group, the system automatically deletes the Enhanced FlashCopy group that is associated with that member logical drive.

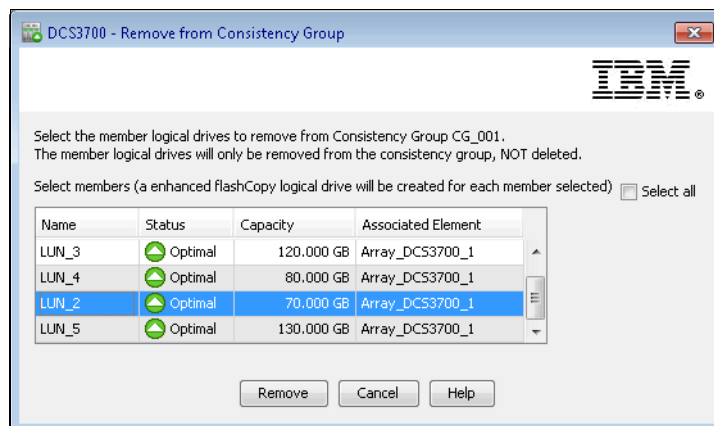


Figure 8-74 Remove from Consistency Group

In addition, you can choose whether you want to delete any repositories that are associated with the member logical drive (Figure 8-75).

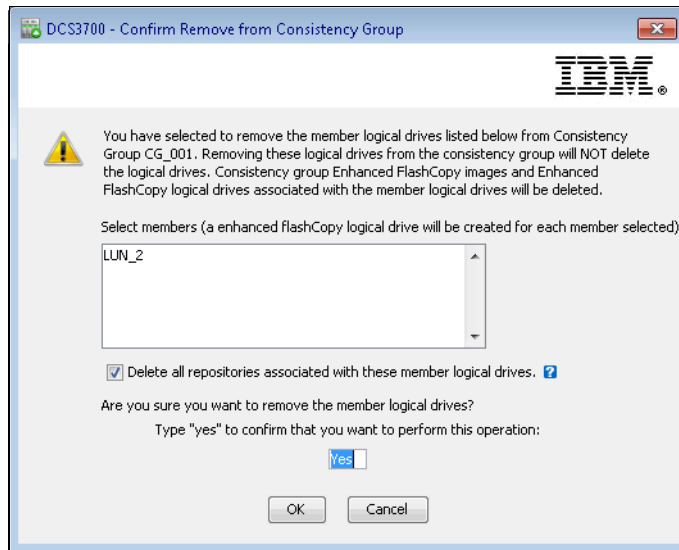


Figure 8-75 Confirm Remove from Consistency Group

## Change Settings

Use this option to modify the auto-delete settings and the consistency group member repository settings that were configured when you created the consistency group.

As shown in Figure 8-76 on page 302, you can change two options:

- ▶ **Auto-Delete Settings:** You can configure each consistency group to keep the total number of Enhanced FlashCopy images in the group at or below a user-defined maximum. When this option is enabled, the system automatically deletes the oldest Enhanced FlashCopy image in the group any time an Enhanced FlashCopy is created to comply with the maximum number of Enhanced FlashCopy images that are allowed for the group.
- ▶ **Consistency Group Repository Settings:** You can define a maximum percentage for the consistency group member repository that determines when a warning is triggered when the capacity of a consistency group member repository reaches the defined percentage. In addition, you can specify which policy to use when the capacity of the consistency group member repository reaches its maximum defined percentage:
  - **Auto-purge oldest Enhanced FlashCopy image:** The system automatically purges the oldest Enhanced FlashCopy image in the consistency group, which releases the member repository's reserve space for reuse within the group.
  - **Reject writes to base logical drive:** When the repository reaches its maximum defined percentage, the system rejects any I/O write request to the base logical drive that triggered the repository access.

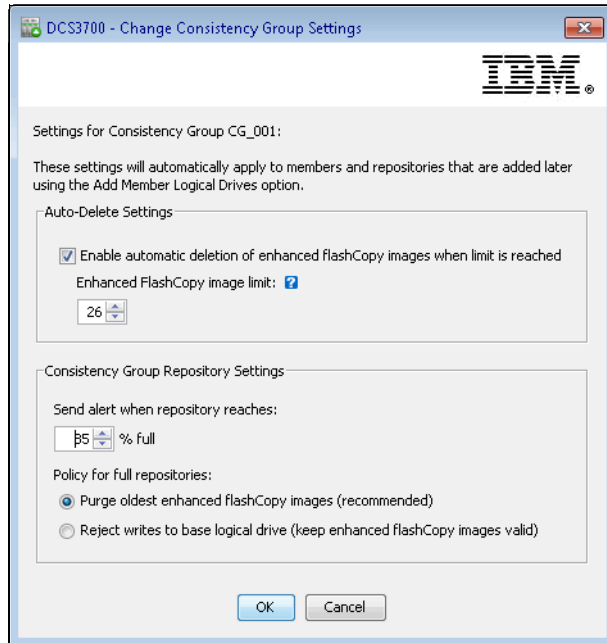


Figure 8-76 Change Consistency Group Settings

## Rename and Delete

Use the Rename option to rename a consistency group. You must use a unique name that has not been used for any other consistency group in the managed DCS3700.

As shown in Figure 8-77, use the Delete option to delete a consistency group.

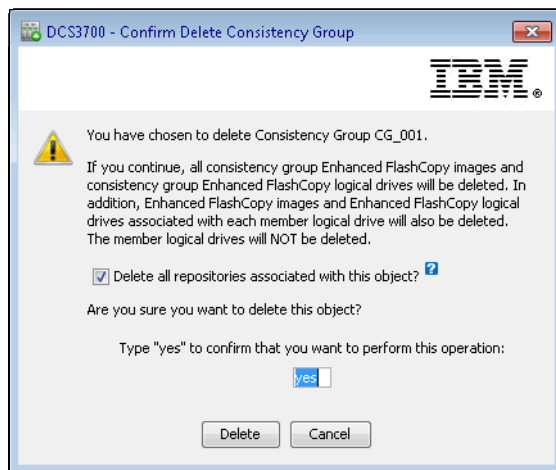


Figure 8-77 Delete Consistency Group

The system performs the following actions when a consistency group is deleted:

- ▶ Deletes all existing Enhanced FlashCopy images from the consistency group.
- ▶ Deletes all existing Enhanced FlashCopy logical drives from the consistency group.
- ▶ Deletes all the associated Enhanced FlashCopy images that exist for each member logical drive in the consistency group.

- ▶ Deletes all the associated Enhanced FlashCopy logical drives that exist for each member logical drive in the consistency group.
- ▶ Deletes all associated repositories that exist for each member logical drive in the consistency group (if selected).

### **Advanced**

The Advanced has one item, that is, Cancel Pending Enhanced FlashCopy Image.

Use this option to cancel an Enhanced FlashCopy image that was put in a Pending state when you attempted to create the Enhanced FlashCopy image for either an Enhanced FlashCopy group or a consistency group. The Enhanced FlashCopy image is in a Pending state due to the following concurrent conditions:

- ▶ The base logical drive for an Enhanced FlashCopy group or one or more member logical drives of a consistency group that contains this Enhanced FlashCopy image is a member of an Enhanced Global Mirror group.
- ▶ The logical drive or logical drives are currently in a synchronization operation.





## Administration: Hardware tab

We describe how to configure logical objects (arrays and logical drives) of the IBM System Storage DCS3700 storage subsystem in Chapter 8, “Administration: Storage & Copy Services tab” on page 225 and how to map it in Chapter 10, “Administration: Mappings tab” on page 327. This chapter describes the DS Storage Manager Hardware tab, where you can find options to configure selected hardware components and see their properties.

## 9.1 Hardware tab

The Hardware tab (Figure 9-1) 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.

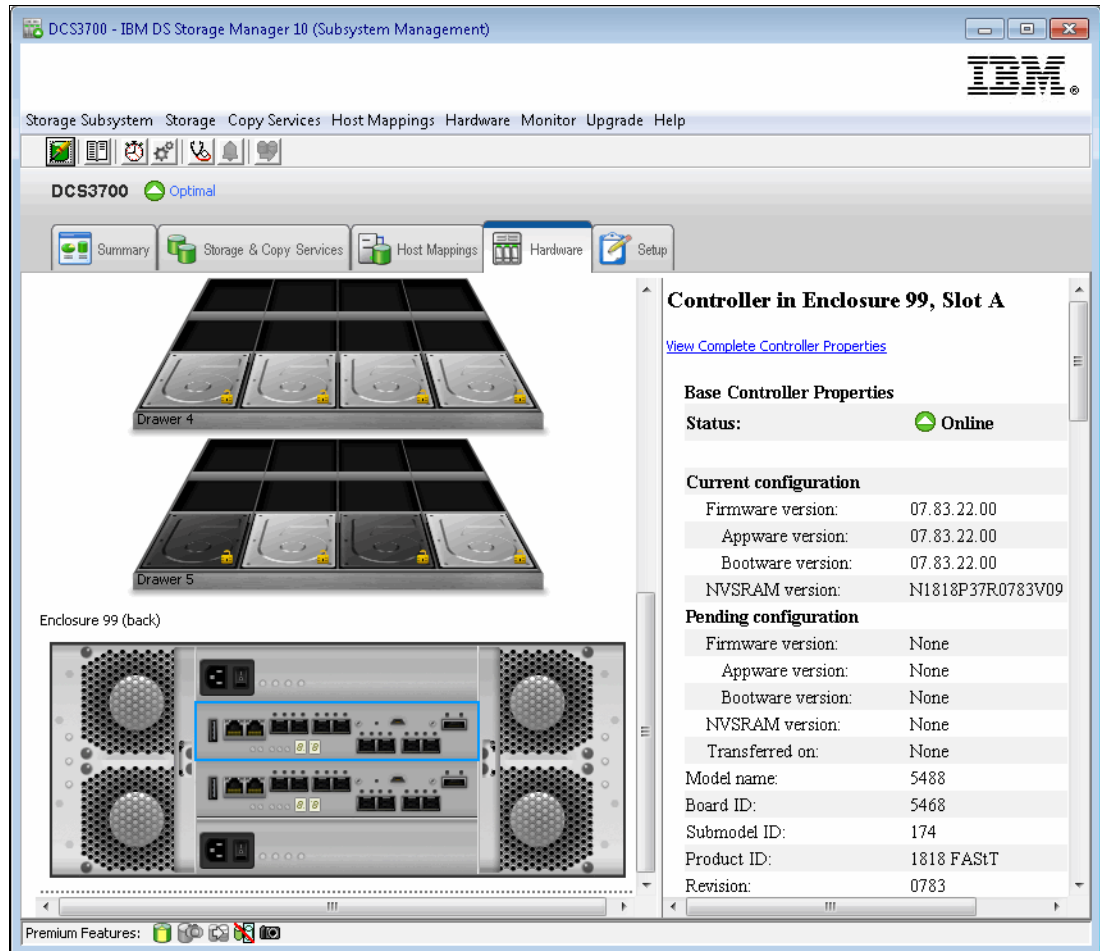


Figure 9-1 Hardware tab

The *Physical pane* shows a graphical representation of the physical components in the storage subsystem, such as the controller modules and the drive expansion units.

The *Properties pane* displays detailed information about the component that is selected in the Physical pane. As in the Storage & Copy Services 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, the following sections focus on the Physical pane:

- ▶ Discovering component properties and location
- ▶ Setting a hot spare drive
- ▶ Failed disk drive replacement
- ▶ Setting a preferred loop ID
- ▶ Setting remote access
- ▶ Setting the Ethernet management ports



The other configuration options in the Physical pane are described in previous chapters and we provide the references as needed.

## 9.2 Discovering component properties and location

In the left Physical pane, you can see a schematic layout of the physical enclosures of the DCS3700 storage subsystem. In Figure 9-1 on page 306, a subsystem with a DCS3700 controller enclosure unit with 60 disk drives is displayed. You can align this layout to the actual physical placement in your rack, as described in “Other enclosure options” on page 309.

### 9.2.1 Showing the disk type

In the upper part of the Physical pane, you can see a drop-down menu where you can select the type of the disk that you want to check. Select the type of the disk and click **Show**. All the disks that are marked with a green arrow belong to a selected disk type, as shown on Figure 9-2. Click **Hide** to remove the arrows from the disk drives.

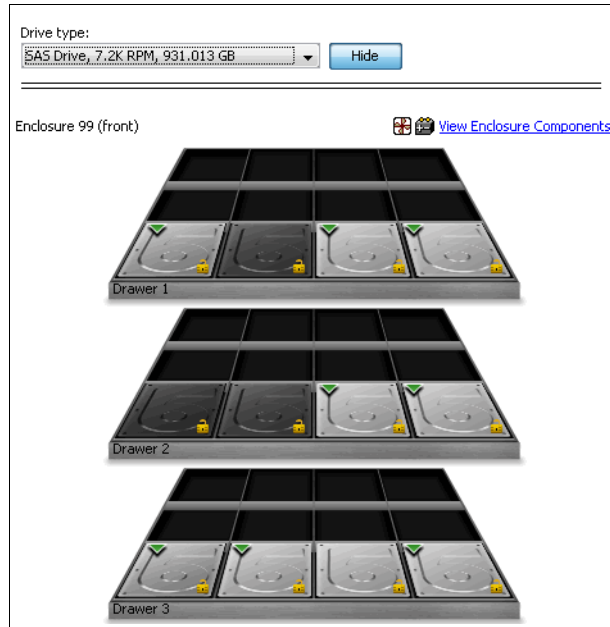


Figure 9-2 Drive type selection

## 9.2.2 View Enclosure Components

The second link in the Physical pane opens a window that has detailed information about the hardware components of a specific enclosure (Figure 9-3). If you click the buttons at the bottom of that window, you can also find this enclosure in the rack (the blue LED diode on the left side of enclosure starts to flash) or you can change the enclosure ID. For more information, see 3.2, “Enclosure ID settings” on page 101.

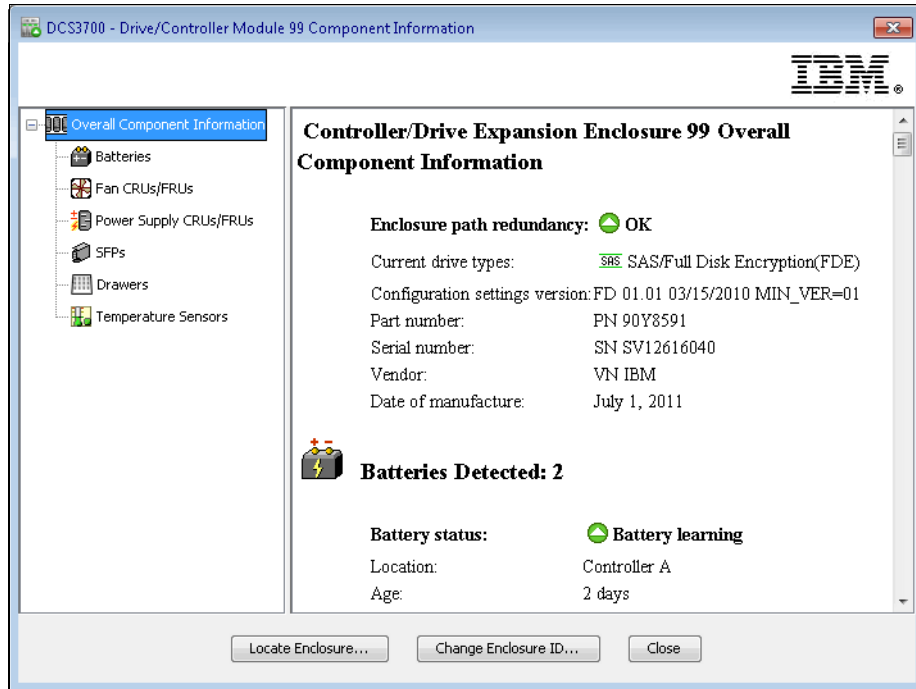


Figure 9-3 Component information

The remaining physical components in the Physical pane are disks and controllers. If you select a disk or a controller, you get relevant information about that disk or controller in the Properties pane (Figure 9-4 on page 309).

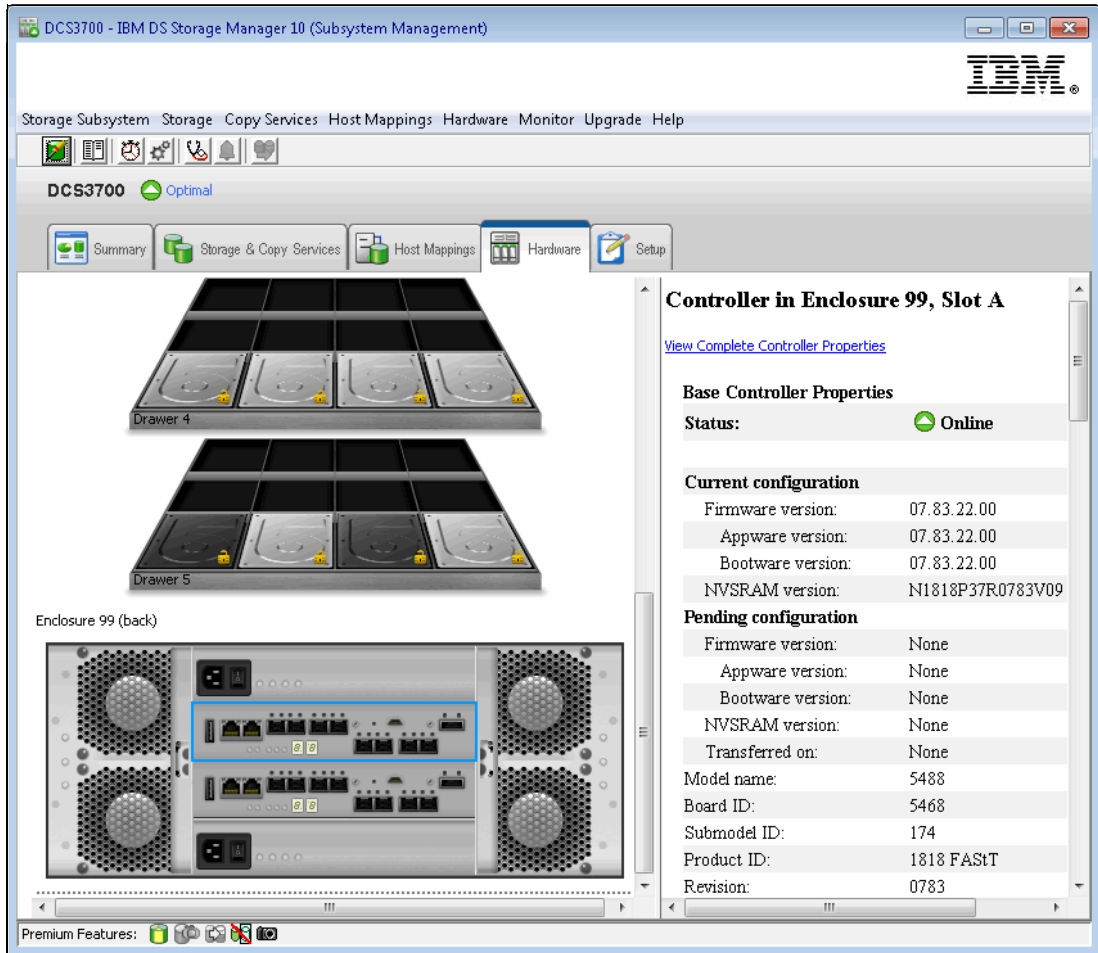


Figure 9-4 Controller information

## Other enclosure options

You can find more options that are related to enclosure in the Hardware menu at the top of the Subsystem Management window (Figure 9-5).

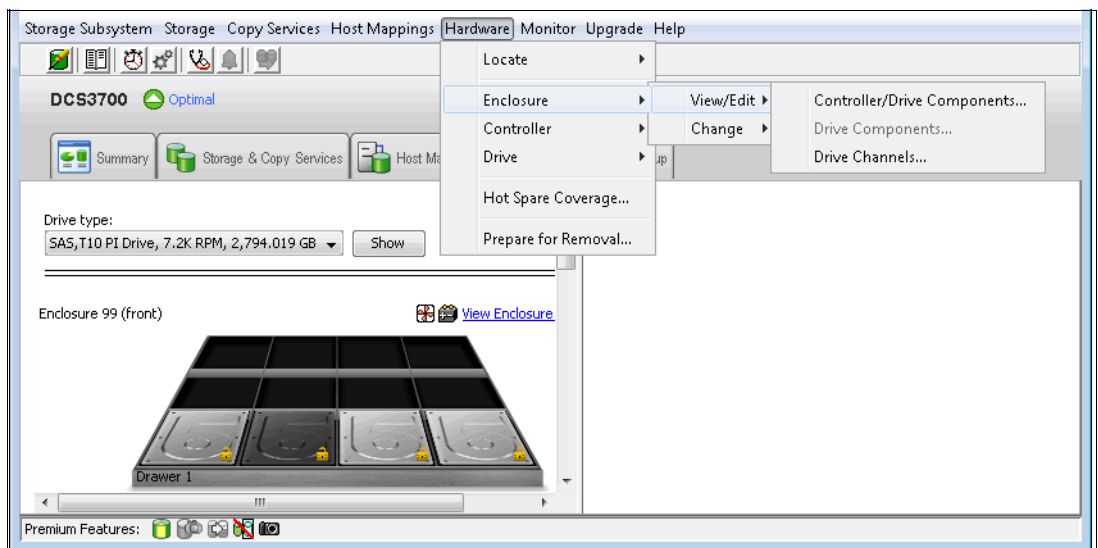


Figure 9-5 Hardware menu

Under the **View/Edit** menu item are the following options:

- ▶ **Controller Drive Components:** This option is the same as the View Enclosure Components option of the Physical pane (see 9.2.2, “View Enclosure Components” on page 308).
- ▶ **Drive Components:** This option provides information about expansion enclosure components.
- ▶ **Drive Channels:** This option provides information about physical links between controller enclosures and drive expansion units that transfer host I/O data (see Figure 9-6). By using this option, you can complete the following tasks:
  - You can find more information about a drive channel by clicking **Show Details**.
  - You can clear error counters on drive channels by clicking **Clear Counters Error**.
  - You can change the drive channel status by clicking **Set Degraded/Set Optimal**.

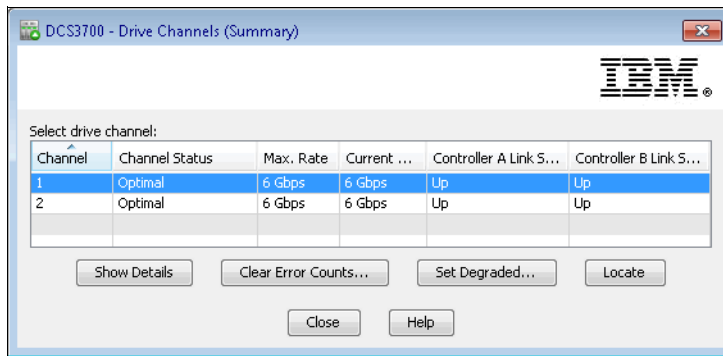


Figure 9-6 Drive Channels

Under the Change menu item, you can find following options:

- ▶ **ID:** This option allows you to change the ID number of the DCS3700 storage subsystem and the expansion units that are attached to it.
- ▶ **Hardware View Order:** Use this option to change the order of the controller enclosures and the drive expansion units in the Physical pane to match the hardware configuration in your storage subsystem (Figure 9-7).

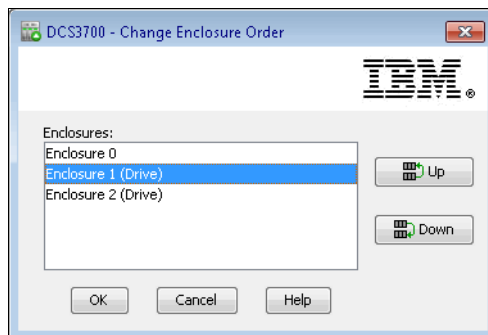


Figure 9-7 Change Enclosure Order

- ▶ **Alarm Settings:** Use this option to set an audible alarm that indicates that a controller enclosure or a drive expansion unit is experiencing a critical failure. You can configure the audible alarm to sound when a failure occurs, or you can disable the audible alarm.

- ▶ **Battery Settings:** With this option, you can schedule a battery learn cycle on both controllers (Figure 9-8). A learn cycle is an automatic cycle for calibrating the smart battery gas gauge. The learn cycle consists of these phases:
  - Controlled battery discharge
  - Reset period
  - Charge

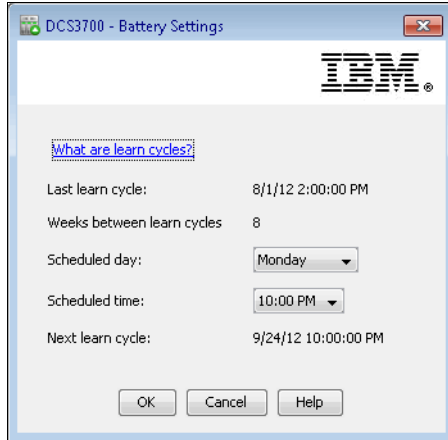


Figure 9-8 Battery Settings

### 9.2.3 Disk Drive menu

If you right-click a disk drive, the menu for other configuration options opens, as shown in Figure 9-9. Certain items in the menu might not be available because not all options are possible for all disk drive types. The options in the menu also depend on the status of the disk drive and how the disk is configured.




Figure 9-9 Disk drive options

**Note:** Use the Advanced option only if you are advised to do so in an action plan that is provided by IBM.

Here are the two menu options that are described later in this chapter:

- ▶ Hot Spare Coverage in 9.3, “Setting a hot spare drive” on page 314
- ▶ Replace Drive in 9.4, “Failed disk drive replacement” on page 319

The other menu items are covered in the following sections.

## Locate

This menu item is used to find where the selected disk is placed in your enclosure in the rack. You also can locate the enclosure where the disk is placed.

## Locate Drive Expansion Enclosure

Use the Locate Drive Expansion Enclosure option to physically locate and identify a drive expansion unit in the storage subsystem.

## 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 by pressing 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 opens, as shown in Figure 8-9 on page 233. In this window, manual mode is preselected. Follow the instructions that are described in 8.2.3, “Create Array” on page 232, and when you reach the window for disk drives selection, your planned disks are already selected for the new array.

The Create Logical Drive menu option must start with creating an array. When an array is created as described above, you can define a new logical drive by using the same procedure that is described in Chapter 8, “Administration: Storage & Copy Services tab” on page 225.

## 9.2.4 Controller menu

If you right-click one of the two controllers in the Physical pane, the menu for other controller configuration options opens (Figure 9-10).



Figure 9-10 Controller option

### Synchronize Clocks

If you select this menu item, you can synchronize the storage subsystem controller clocks with the storage management station. This option makes sure that the event time stamps that are written by the controllers to the Event Log match the event time stamps that are written to host log files. During this process, the controllers remain available during synchronization.

### Configure

Under this option in the menu, you can configure the management port on the selected controller, as shown in Figure 9-11 on page 313.

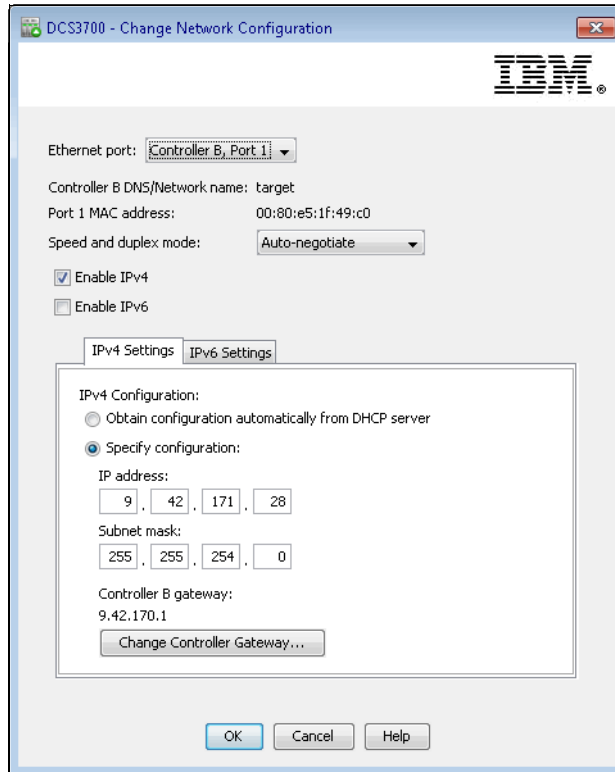


Figure 9-11 Management port configuration

## Change

When you click this menu item, the following configuration options are displayed:

- ▶ Preferred Loop ID, described in 9.5, “Setting a preferred loop ID” on page 321 (Fibre Channel version only)
- ▶ Remote Login, described in 9.6, “Setting remote access” on page 323

## Advanced

**Note:** All options under the Advanced menu should be used only if you are advised to do so by IBM Support.

This menu item has four options:

- ▶ Place: This option has the following suboptions:
  - 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.

**Note:** The controller might be in an Offline state if it failed a diagnostic check.

- Online: Placing a controller online sets it to the Optimal status and places it in an active mode. It is available for I/O operations.

**Note:** The detection of a restored preferred path by the multipath driver can take up to 10 minutes.

- In Service Mode: Use the Place Service Mode option to safely run controller diagnostic tests without affecting data availability.
- ▶ Run Diagnostics: This option has the following suboptions:
  - Controller: This option starts various internal tests to verify that a controller is functioning correctly.

**Note:** If you run diagnostic tests while a host is using the logical drives that are owned by the selected controller, the I/O directed to this controller path is rejected.

- Discrete Lines: Use this option if the Recovery Guru has a problem with Discrete Lines. The Discrete Lines are dedicated control and status lines between the controllers that are in the Interconnect-battery CRU/FRU.
- ▶ Reset: Use this option to reset a controller from within the storage management software.

**Note:** Resetting a controller makes the controller unavailable for I/O until the reset is complete. If a host is using logical drives that are owned by the controller that is being reset, the I/O directed to the controller is rejected. Before resetting the controller, either make sure that the logical drives that are owned by the controller are not in use, or make sure that a multipath driver is installed on all of the hosts that are using these logical drives.

- ▶ Enable Data Transfer: A controller can be in the disabled state because of diagnostic tests that are being run or firmware that is being downloaded. A controller in the disabled state no longer accepts I/O and is marked with a dashed red border on the Hardware tab. If a multipath driver is in use, I/O is 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 complete these steps to enable data transfer.

## 9.3 Setting a hot spare drive

A hot spare drive is like a replacement disk drive that is installed in advance. Hot spare disk drives provide more protection that can be essential if there is a disk drive failure in a fault-tolerant RAID array.

**Note:** Disk pools do not use hot spare drives for data protection like an array does. Instead of hot spare drives, disk pools use spare capacity within each drive that makes up the disk pool.

A configuration can have only SAS disk types, but with different capacities and speeds, you must plan how to provide the best hot spare protection for your storage subsystem.

To plan your hot spare coverage correctly, following these guidelines:

- ▶ 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 DCS3700 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.



**Important:** When assigning disks as hot spares, make sure that 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 spare drive 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 mode, 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 distribution avoids the risk of a single disk drive channel or expansion unit failure causing a loss of access to all hot spare drives in the storage subsystem.

### **Quantity and type of hot spare drives**

There is no fixed rule about the quantity of disk drives to assign as hot spares, and depends on the method that you use to create storage space.

If you used the automatic configuration function, the storage management software creates one hot spare drive for every 30 drives of a particular media type and interface type. If you configured the storage subsystem manually, you can create hot spare drives that are used among the arrays in the storage subsystem. The hot spare drives must be equal to or greater than the capacity of the largest drive in the drive set. The recommended number of hot spare drives is two for each drive set in the storage subsystem.

The following section covers the available methods of defining hot spare drives with Storage Manager for created arrays.

To start the configuration of a new hot spare from the Physical pane, select an unconfigured disk drive that you plan to use as hot spare disk drive, right-click it, and select **Hot Spare Coverage**. The window that is shown in Figure 9-12 opens.

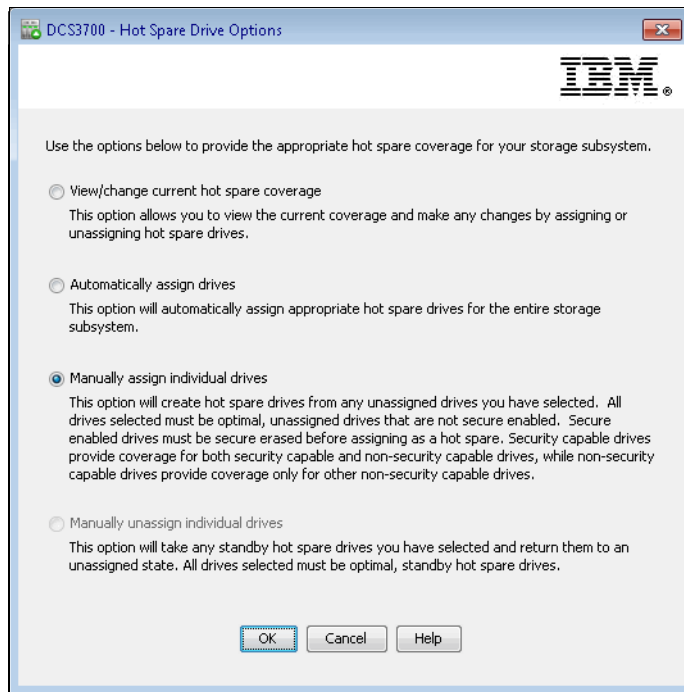


Figure 9-12 Hot Spare Drive Options

You have four options to define hot spare disk drives:

- ▶ View/Change coverage
- ▶ Automatic assignment
- ▶ Manual assignment (if an unassigned disk was selected in the Physical pane)
- ▶ Manual unassignment (if a hot spare disk drive was selected in the Physical pane)

### Automatic assignment

For automatic hot spare assignment, complete these steps:

1. To automatically create the hot spare coverage by using the drives that are available, select **Automatically assign drives**. The recommended quantity of spares drives that is needed for your configuration are created automatically.
2. Click the **Physical** tab to view the results of the automatic hot spare creation.

This automatic hot spare creation function on DCS3700 creates one hot spare drive for every 10 unassigned disk drives. Results can be unexpected in an environment with mixed drive types and if disk drives are already assigned to arrays. In such scenarios, we recommend using the manual assignment method instead of the automatic method.

### Manual assignment and unassignment

To perform manual hot spare assignment, complete the following steps:

1. Select an unconfigured disk drive that you plan to use as hot spare disk, right-click it, and select **Hot Spare Coverage**. The window that is shown in Figure 9-12 opens with the **Manually assign individual drives** option already preselected. Click **OK** and the unassigned drive is defined as a hot spare.

2. Select a hot spare disk drive that you do not want to use as a hot spare disk drive any more, right-click it, and select **Hot Spare Coverage**. The window that is shown in Figure 9-12 on page 316 opens with the **Manually unassign individual drives** option already preselected. Click **OK**, confirm your choice by typing Yes in the next window, and the former hot spare drive is an unassigned disk drive again.

**Tips 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 that contain many drives, it is necessary to define multiple hot spare drives, as the reconstruction of a failed drive to a hot spare can take a long time.
- ▶ Consider having hot spare drives for each speed and type of disk drives.

**View/change hot spare coverage option**

In the Hot Spare Drive Option window (Figure 9-12 on page 316), click the **View/change hot spare coverage** option. The Hot Spare Coverage window opens, as shown in Figure 9-13.

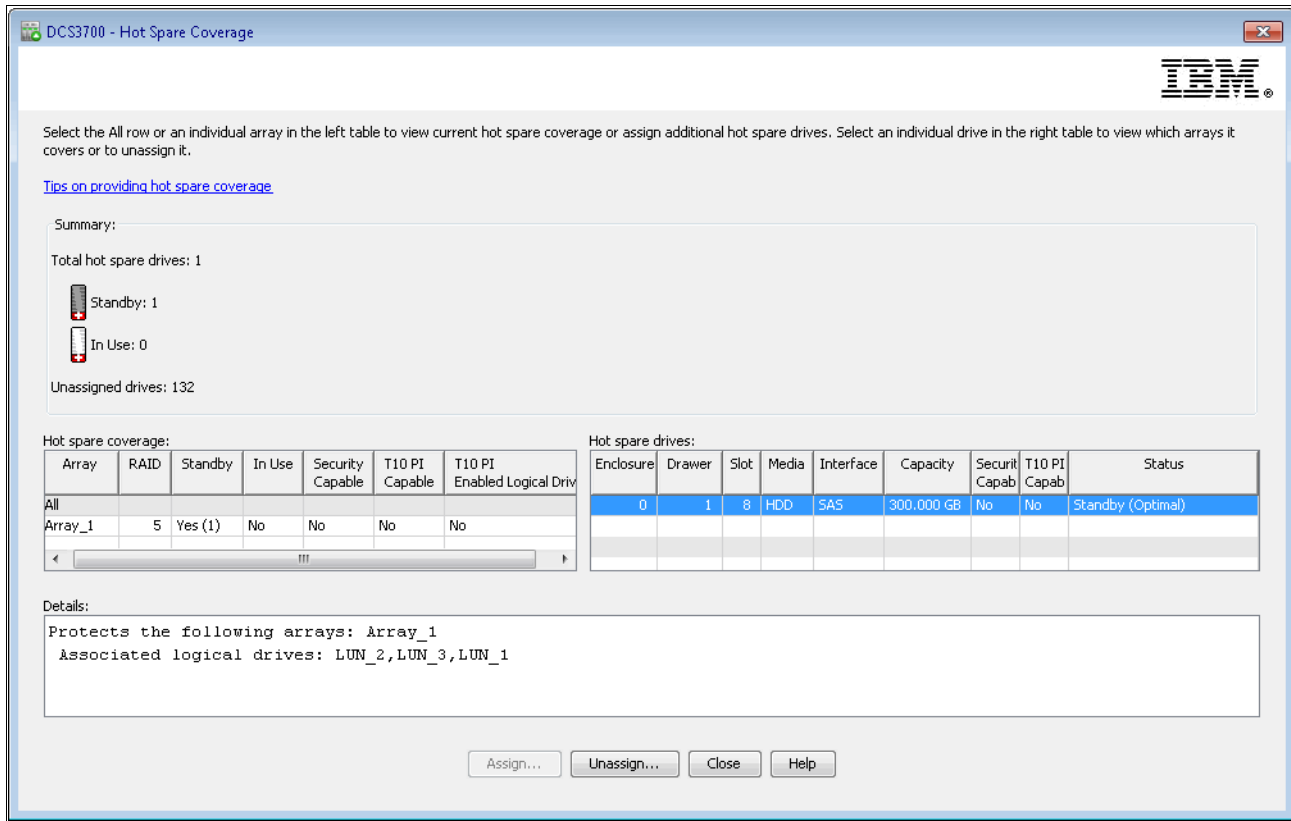


Figure 9-13 Hot Spare Coverage

This window has four panes:

- ▶ The Summary pane, which has the following items:
  - Numbers of hot spare disk drives
  - Number of hot spare disk drives in use
  - Number of unassigned disks that can be used as hot spare disk drives

- ▶ The Hot spare coverage pane shows all the arrays and whether they are protected by hot spare disk drives. If an array is selected in this left pane, in the right Hot spare drives pane, you see the hot spare disk drive as assigned. If you select an array that has no hot spare disk drive that is defined, you can click **Assign**, and a new window opens (Figure 9-14).

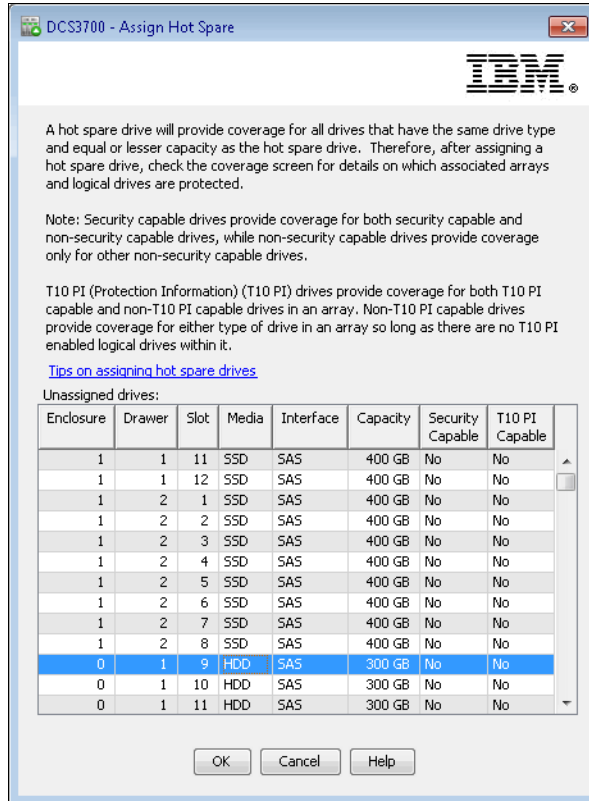


Figure 9-14 Assign Hot Spare

Here you can select a new hot spare disk drive for your unprotected array.

- ▶ The Hot spare drives pane shows hot spare disk drives that are used. If you have more hot spare disk drives than you need, you can select a hot spare disk drive here and click **Unassign** to unassign it. The **Unassign** button is unavailable if you can lose protection by unassigning this hot spare disk drive. If you still need to unassign this hot spare drive, do it manually, as described in “Manual assignment and unassignment” on page 316.
- ▶ The Details pane shows logical objects that belong to selected array or hot spare disk drives.

Make sure to use this window to check that all your arrays are protected. If they are not protected, implement complete hot spare disk drive protection for all arrays as soon as possible.

## 9.4 Failed disk drive replacement

You can replace a failed or missing disk drive that is used only by an array. Logical drives are in the “Degraded” status if parity for the failed disk is still calculated to the hot spare disk drive (see the example of the Recovery Guru window in Figure 9-15).

You can assign a working unassigned disk drive that is already inserted in the DCS3700 as replacement drive. In addition, the hot spare disk drive in use can be used as replacement drive but only after parity recalculation is completed.

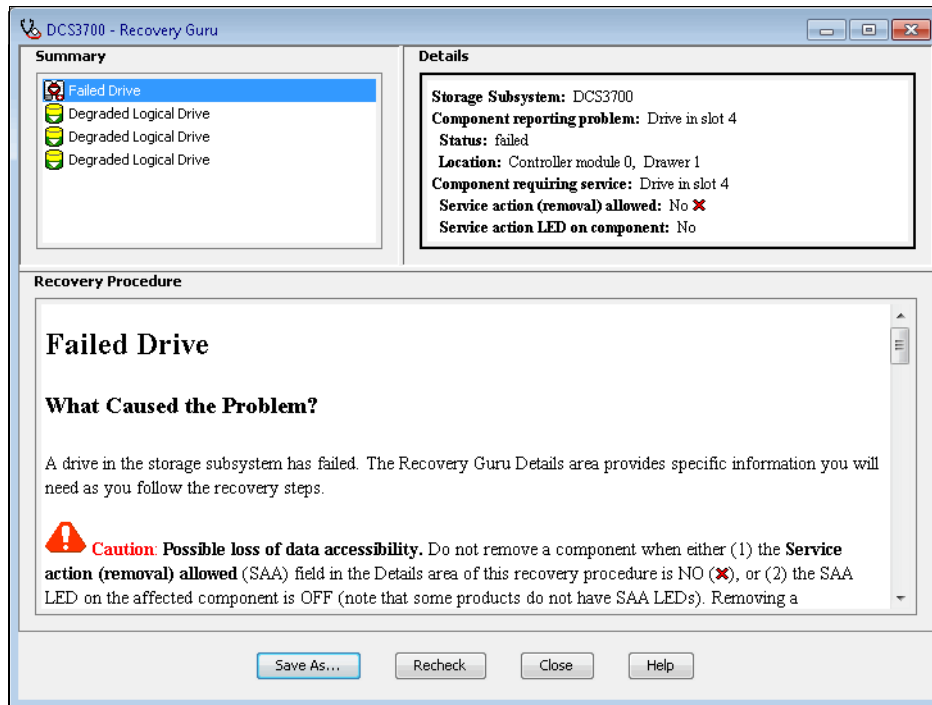


Figure 9-15 Recovery Guru for a failed drive

To replace a failed disk drive, complete 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 opens (Figure 9-16).

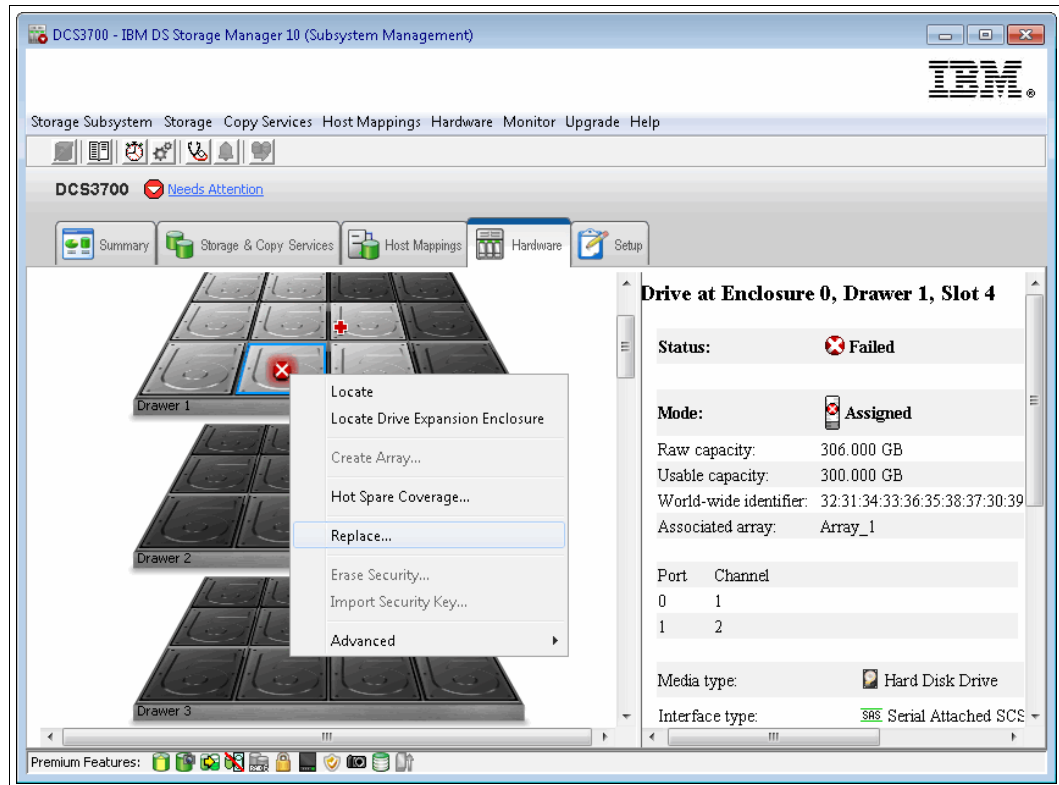


Figure 9-16 Replace Drives option

2. The list of eligible replacement disk drives displays in the window (Figure 9-17).

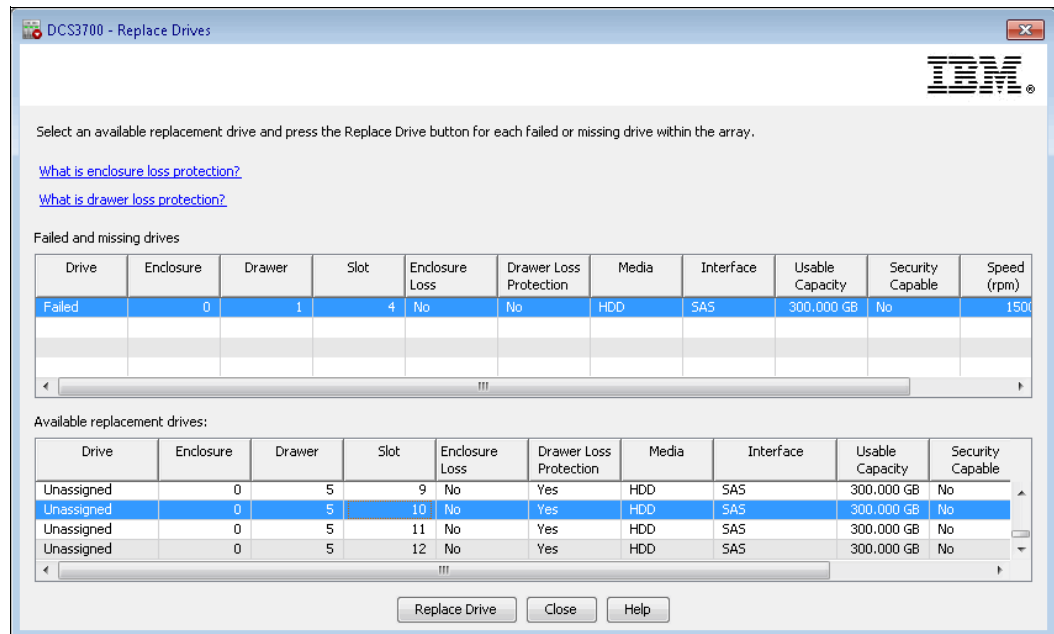


Figure 9-17 Replace Drives - Select disk

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 9-17 on page 320). If this hot spare disk drive is defined and is in use by the array with a missing or failed disk, you must wait until the DCS3700 controller reconstructs parity of all the logical drives in that array.

4. Click **Replace Drive** to start the operation. The confirmation window opens. Click **OK** and a new updated window similar to Figure 9-17 on page 320 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 disk drive. 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.

## 9.5 Setting a preferred loop ID

A loop ID is the unique ID of a port in a Fibre Channel loop topology. It is sometimes referred as a loop address. The Preferred Loop ID menu is not available for DCS3700 without an FC daughter card being installed.

**Attention:** If your Fibre Channel host adapter driver manages storage subsystems that are 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, complete the following 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 9-18).

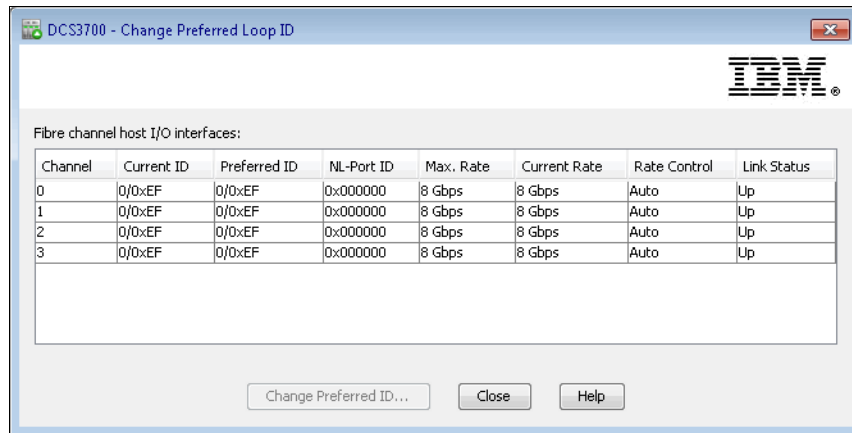


Figure 9-18 FC ports - Loop ID

2. Select the port that you want to change and click **Change Preferred ID**. The window that is shown in Figure 9-19 opens.

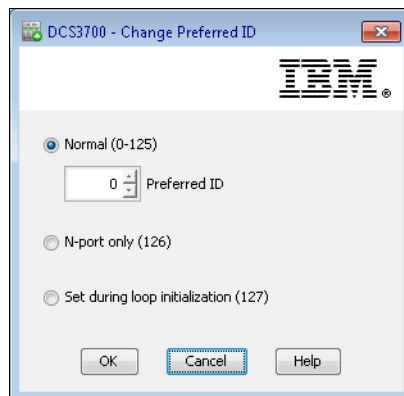


Figure 9-19 Loop ID setting

3. Select one of the following options:
  - Normal (0 - 125): This option assigns a normal preferred loop ID value to the controller port. If the controller port currently has a preferred loop ID 0 - 125, this radio button is selected by default. The current ID is shown in the field.
  - 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.



## 9.6 Setting remote access

This option lets IBM engineers access the subsystem controller through the telnet protocol. This action can be done from outside of the server room where the DCS3700 storage subsystem is installed, allowing the IBM engineer to service the storage subsystem without having to use the special serial management port of the DCS3700 storage subsystem, which is used for maintenance only.

To set remote access, complete the following steps:

1. Select the controller in the Physical pane, right-click it, and select **Remote Login** in the menu. The window that is shown in Figure 9-20 opens.

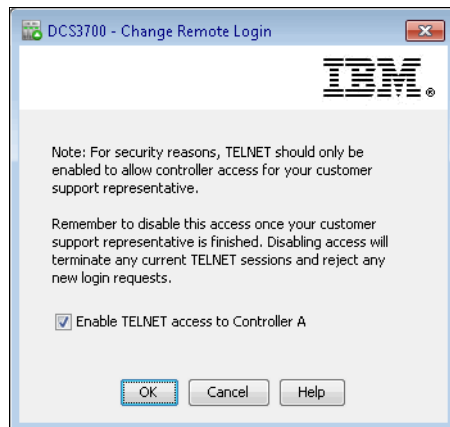


Figure 9-20 Remote login set

2. Read the warning in the window, select **Enable TELNET access to controller A** to set direct remote access to your controller, and click **OK**. You must do the same procedure for the second controller.

**Note:** A new storage subsystem comes with the TELNET function enabled by default. The login is protected by a user name and password. To increase security on the DCS3700 storage subsystem, you can disable the Remote Login option if you do not have a specific need or directions from IBM Support.

## 9.7 Setting the Ethernet management ports

Setting the IP addresses of Ethernet management ports can be done in the configuration window. By default, the DCS3700 storage subsystem tries to use 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 DCS3700 storage subsystem has two management Ethernet ports per controller. Use one Ethernet port for daily management of your DCS3700 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 so that service personnel can attach an isolated Ethernet switch if needed.

## 9.7.1 Initial settings for the default IP addresses in DS Storage Manager

To change the default IP addresses, you must connect to the DCS3700 storage subsystem. Complete 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 DCS3700 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 5, “Administration: Enterprise Management” on page 159.

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 started, the program presents a window to start discovery of the attached devices automatically.

- b. Select **OK** to initiate automatic discovery.

If the DCS3700 storage subsystem does not appear after this action, check the network connectivity to both default IP addresses of the storage subsystem by running **ping**. Each time a network cable is plugged in to an Ethernet port of a controller, it detects the link 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 subsystem during that time, you might receive an error, so wait at least 5 minutes after you plug 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 DCS3700 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 must 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 9-21.

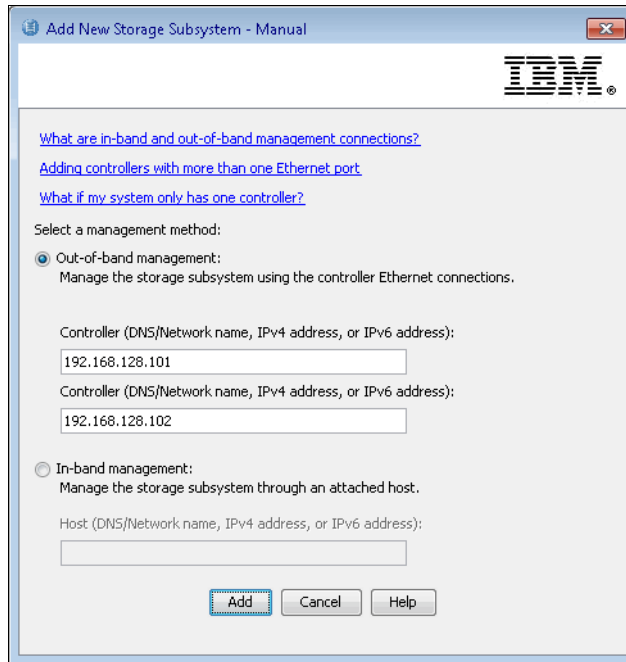


Figure 9-21 Manually adding a storage system

- c. A confirmation message opens and prompts you for more 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 the Devices tab to open the Subsystem Management window. If not already set, you are 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 you are prompted to do so.
- f. Now you are connected to the DCS3700 storage subsystem, but with the default IP addresses. Change the address to your specific ones. To configure the Ethernet ports on DCS3700 controllers, use the standard procedure that is described in 9.7.2, “Configuring the Ethernet management ports” on page 326.

## 9.7.2 Configuring the Ethernet management ports

To configure the Ethernet management ports of a DCS3700 storage subsystem, complete the following steps:

1. Select one controller in Physical pane, right-click it, and select **Configure** → **Ethernet Management Ports**. The window that is shown in Figure 9-22 opens. In this window, you can set IP addresses, IP masks, the default gateway, and the link speed of the Ethernet port, for all four Ethernet ports on both controllers.

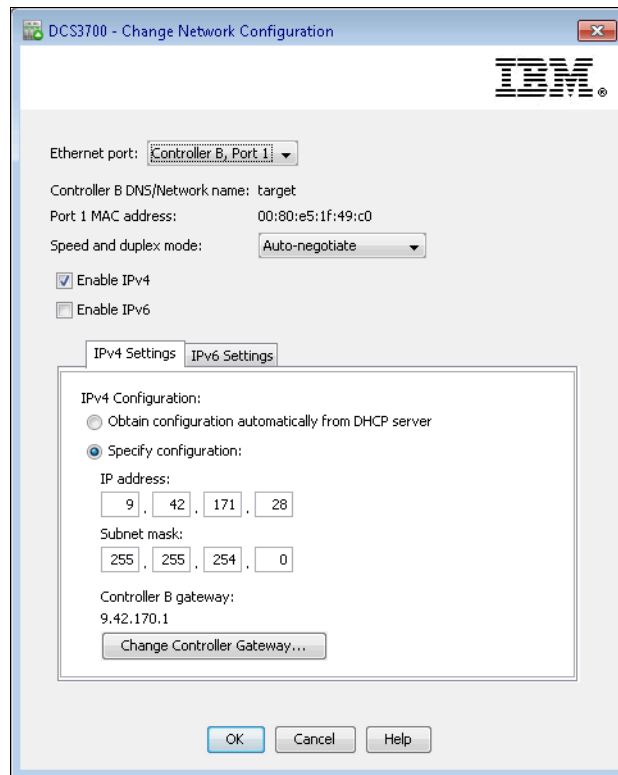


Figure 9-22 Change Ethernet settings

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 DCS3700 storage subsystem 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 and leave port 2 for service.

**Note:** After changing the IP addresses, DS Storage Manager might freeze because you have lost management connection to subsystem. Close the Storage Subsystem window, and remove the subsystem with the 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.



## Administration: Mappings tab

This chapter describes the Mappings tab of the 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 as a Logical Unit Number (LUN).

Section 2.4.9, “Storage partitioning” on page 67 explains the concept of storage partitioning. This chapter shows practical examples about how to configure storage partitioning for SAS and Fibre Channel (FC) connections.

Because heterogeneous hosts can be attached to the IBM System Storage DCS3700 storage subsystem, you must configure storage partitioning for two reasons:

- ▶ Each host operating system requires slightly different settings on the DCS3700 storage subsystem. You must 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 has access only to its assigned logical drives.

You can have up to 256 LUNs that are assigned to a single storage partition, and a maximum of 256 LUNs that are configured per storage subsystem.

## 10.1 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 that are 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:

- ▶ Defining a host
- ▶ Defining storage partitioning
- ▶ Defining a host group
- ▶ Manage Host Port Identifiers option
- ▶ Define Additional Mapping option
- ▶ Viewing unassociated ports
- ▶ Moving, removing, and renaming a host
- ▶ Changing a host type
- ▶ Changing and removing mapping

To start with mapping, select the **Mappings** tab in DS Storage Manager subsystem window. If you have not defined any storage partitions yet, the Mapping Start-Up Help window (shown in Figure 10-1) opens. The information in the window advises you to create only host groups if your plan includes sharing logical volumes across hosts (normally a cluster).

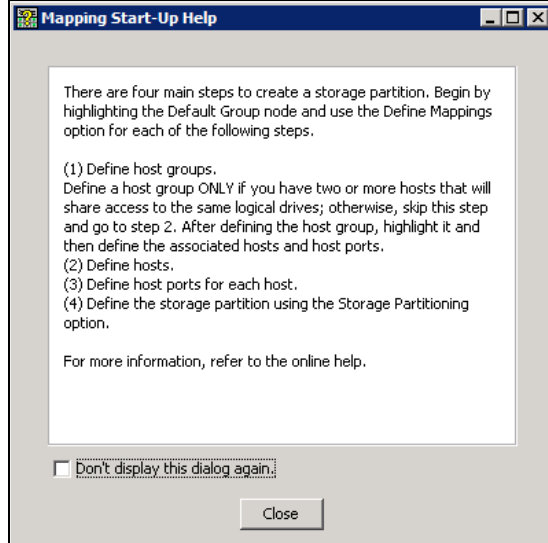


Figure 10-1 Mappings Start-Up Help

Figure 10-2 shows an example of the Mappings tab. The right pane lists all the mappings that are owned by the object that you selected in the left pane. 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.

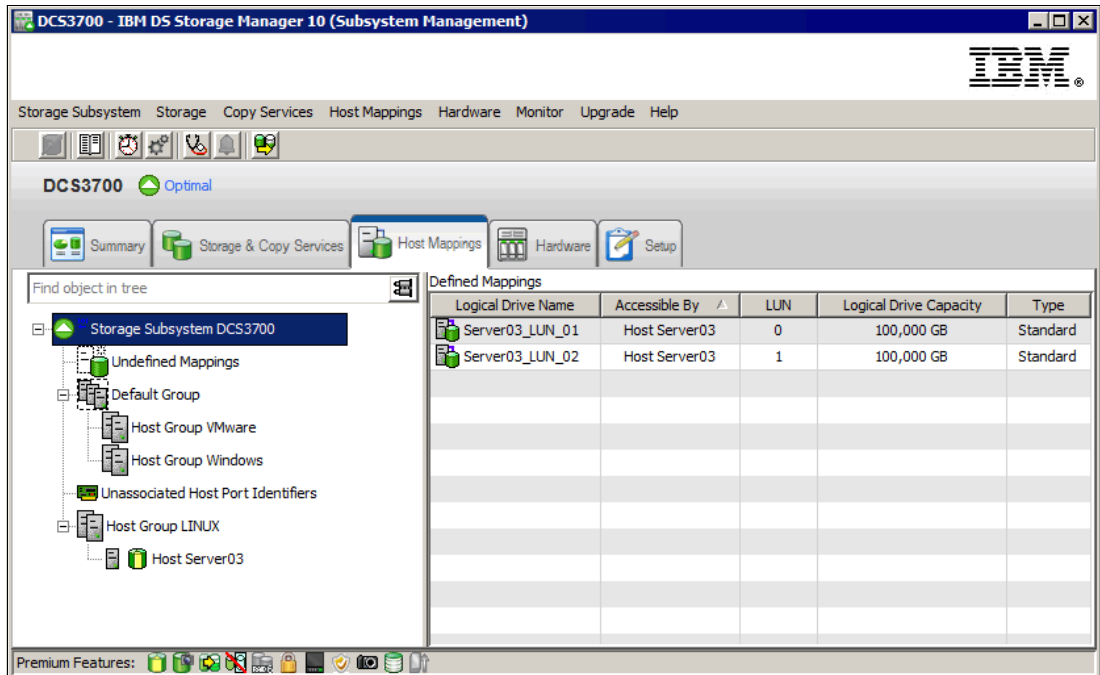


Figure 10-2 Mappings tab in the DS Subsystem Management window

The following sections show how to configure the storage partitioning.

**Notes:**

- ▶ Before configuring the storage partitions, if you have your storage connected to an FC switch, make sure to first define the SAN zones appropriately.
- ▶ Perform a hardware rescan on each host that must be mapped to the DCS3700 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 DCS3700 storage subsystem controller.

## 10.2 Defining a host

A host can be defined as stand-alone server or as a cluster member. To define a host, complete the following steps:

1. Either right-click your Storage Subsystem icon, right-click the Default Group icon, or right-click an existing host group. In the menu that opens, select **Define** → **Host**, as shown in Figure 10-3. Even with no host group created, it is easier to create the hosts first and then the host groups. In Figure 10-3, you can see that all general configuration options for the whole storage subsystem are accessible from the Mappings tab. These options are described in Chapter 7, “Administration: Subsystem Management” on page 203.

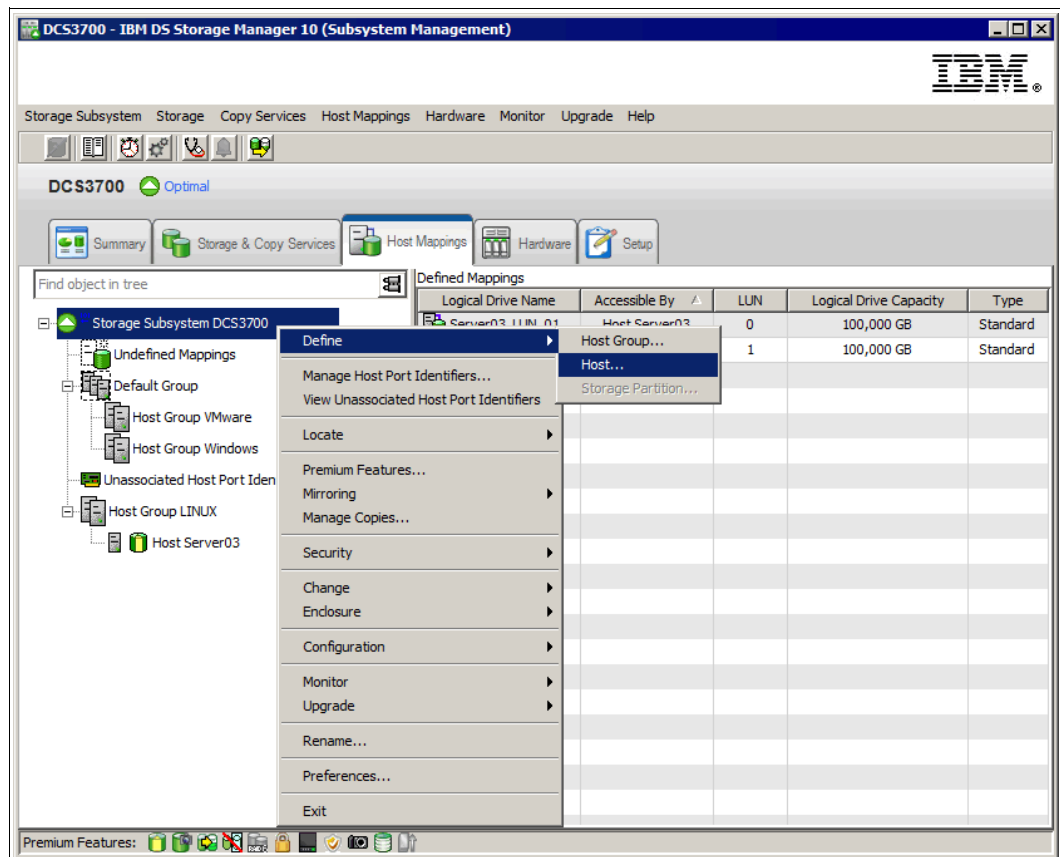


Figure 10-3 Define Host menu

When you select **Define** → **Host**, a wizard for defining a host opens. You must provide the following information during the process:

- Name of the new host
- Protocol attachment method:
  - SAS
  - Fibre Channel (FC)
  - iSCSI
- 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 10-4), you are asked to assign a host name to the host that is being defined, and whether you plan to use storage partitioning.



Figure 10-4 Define Host - Specify Host Name

3. If you plan to attach only a single host server, then you do not need storage partitioning. But if you plan to group host environments, or if clusters are sharing logical drives, click **Yes** to create different hosts groups later. Select **Next** to continue.

4. In the next window (Figure 10-5), you need the attachment protocol to specify the host interface type (SAS, iSCSI, or FC) and the HBA host port information (host port identifier and alias) or the Host Initiator name (for example iqn.1996-04.de.microsoft:01:27dc9458a936). Remember that the HBA host port identifier is the worldwide port name (WWPN) 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, the SANsurfer management tool for the QLogic HBA), use Fabric Name Server in FC fabric management tool, or you can obtain the WWPN from the BIOS during boot process of the server where the HBA is placed.

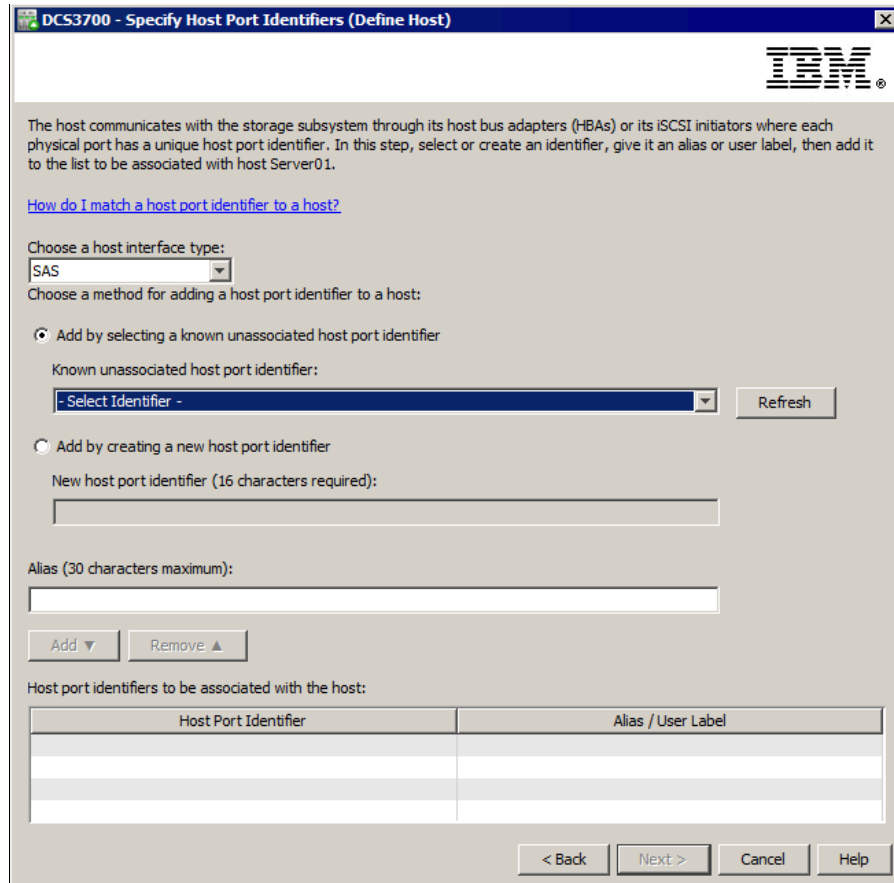


Figure 10-5 Define Host - interface type, known port identifier, and alias

Later in this book, we cover the basic usage of the QLogic SANSurfer HBA management tool in 17.1.3, “Installing QLogic SANSurfer” on page 506 and in 18.1.3, “Installing QLogic SANSurfer” on page 524.

Complete the following steps:

- a. Select your interface type (SAS in our example)
- b. Select an active port that is already detected by the storage subsystem in the Known unassociated host port identifier field, as shown in Figure 10-5, or enter the WWPN directly into the field (see Figure 10-6 on page 333).




Figure 10-6 Define Host - interface type, new port identifier, and alias

- c. Enter an alias for the specified host port identifier and click **Add**.
- d. Repeat the same process until you define all the HBAs. 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 the host.

5. The next window requires you to specify the host type (Figure 10-7). This is the operating system that is running on the host. It is vital that you select the appropriate host type because the multipath drivers 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.

**Note:** ALUA capabilities are introduced with firmware Version 7.83. The usage of this new feature relies on the correct host type configuration. For more information, see the IBM System Storage Interoperation Center (SSIC), found at:

<http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss>

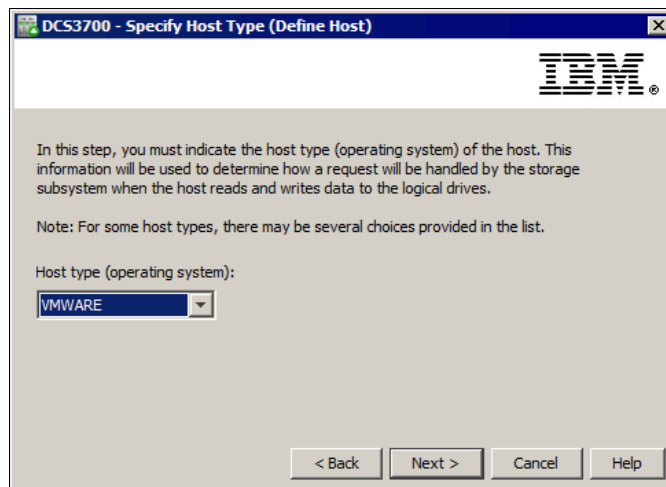


Figure 10-7 Define Host - Specify Host Type

6. In the next step, you are asked whether the host is a part of a cluster (Figure 10-8). Choose **Yes** if it will be clustered or **No** if it will be a stand-alone server. Click **Next** to continue.

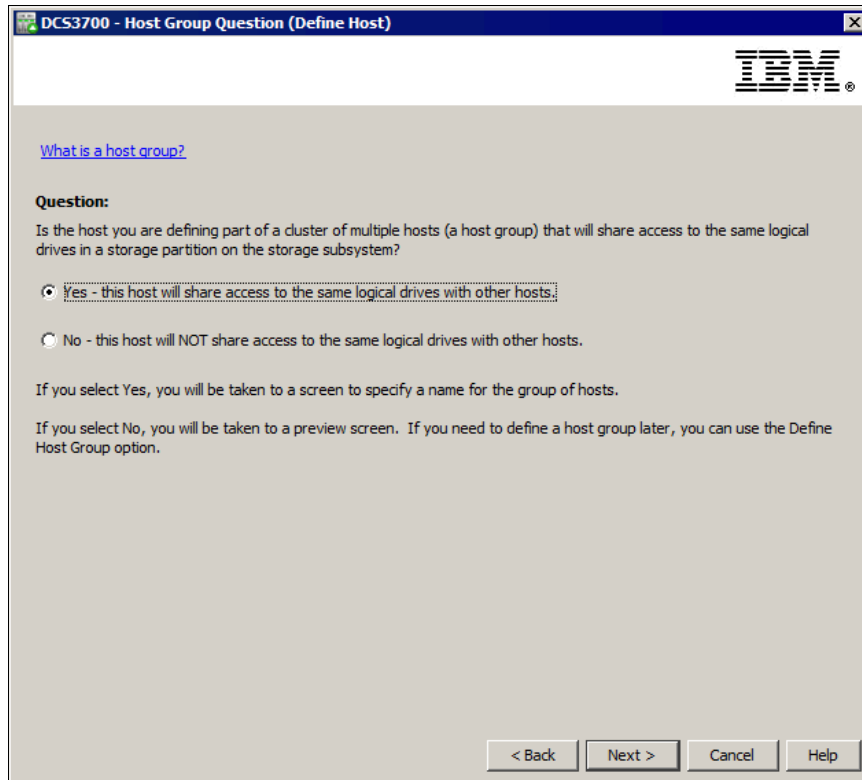


Figure 10-8 Define Host - specify sharing (clustering)

7. (Optional) If the answer is **Yes**, then you must 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 10-9, the new host group is created with this new host as a member. Select **Next** to continue.

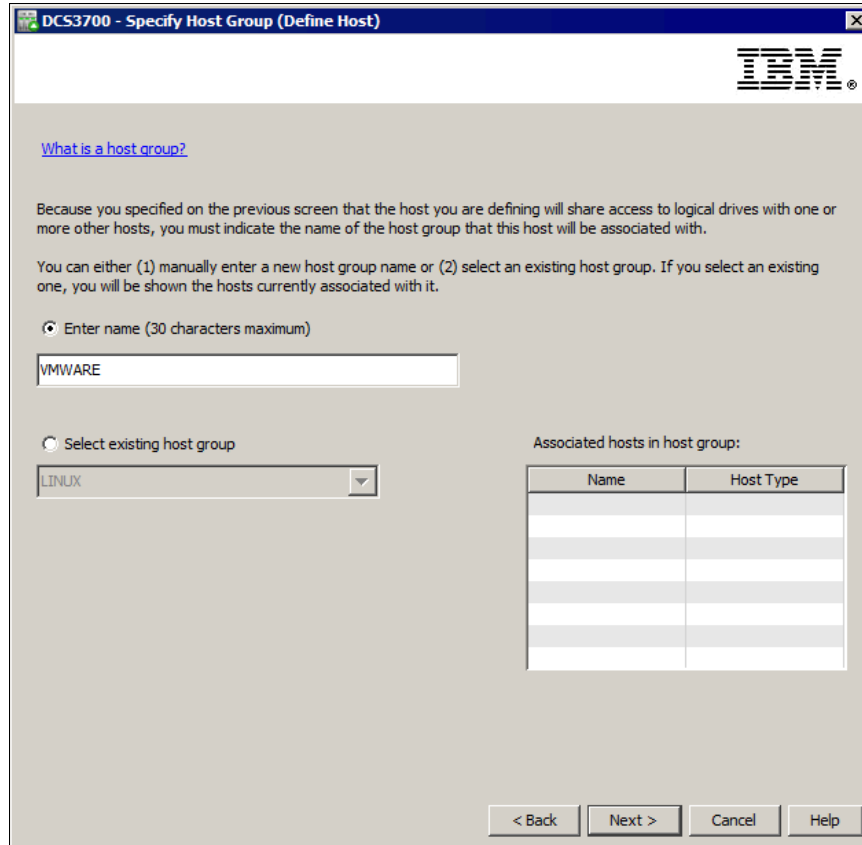


Figure 10-9 Define Host - specify a new host group

8. Finally, you have a chance to preview the new host definition (Figure 10-10). If all the selections are correct, click **Finish** to define the host. You can save the host definition as a script file for SMcli by clicking **Save As Script**. 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.

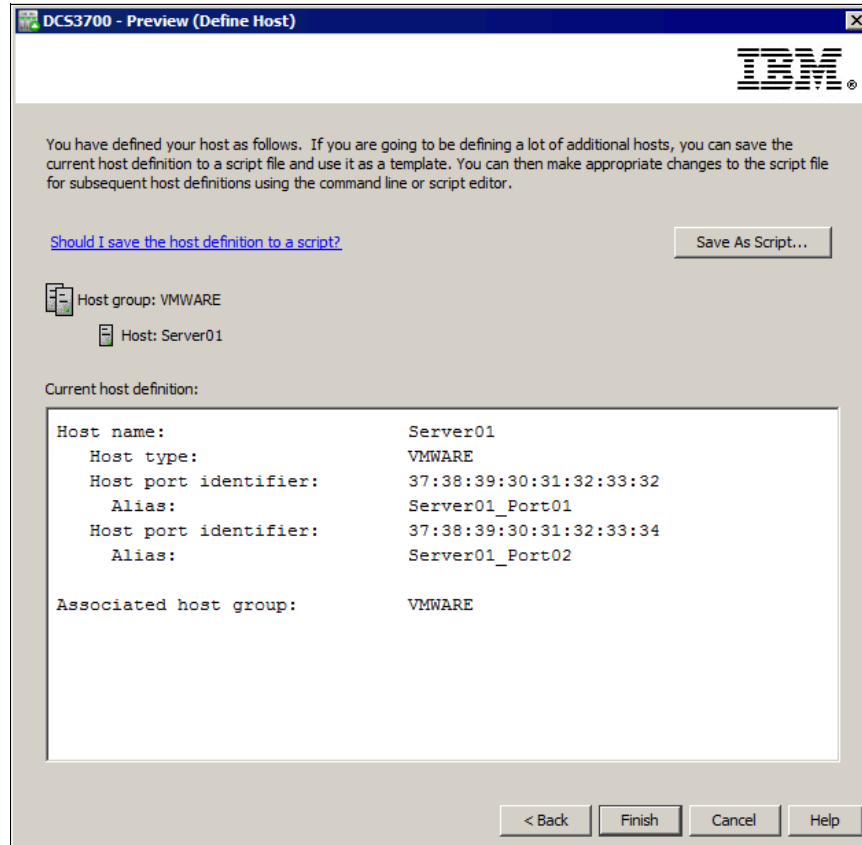


Figure 10-10 Define Host - new host definition preview

9. 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 10-11). It stays there until you create a storage partition by assigning the logical drives to that host (or host group).




Figure 10-11 Host defined

10. Repeat these steps to create more hosts and host groups.

### 10.2.1 Adding a host to an 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. When you put the new host in to the host group that the new host can access all logical drives that are mapped to this host group immediately.



To add a host to the VMWARE host group, complete the following steps:

1. Right-click the VMWARE host group and select **Define** → **Host**, as shown in Figure 10-12.

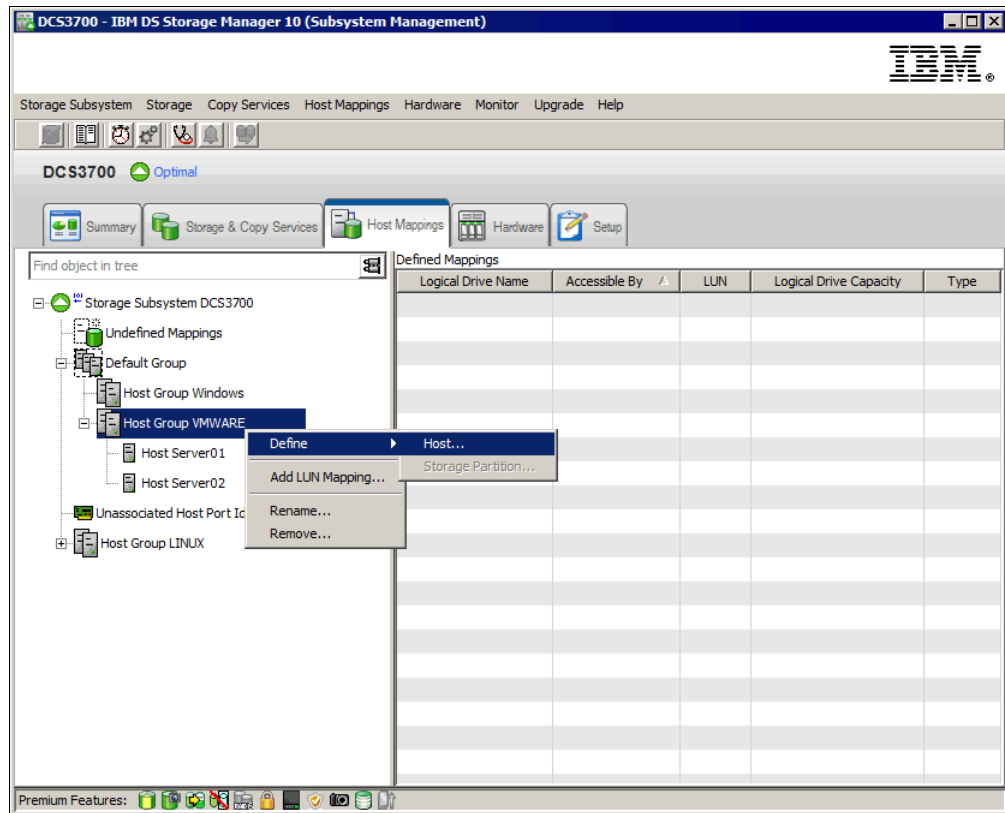


Figure 10-12 Defining hosts

2. Define the host by completing steps 1 on page 330 to 7 on page 336.

## 10.3 Defining storage partitioning

This section describes how to set a new storage partition definition for a host or host group that has no logical drives mapped yet. You must use the **Additional Mapping** menu to map a logical drive to a host or host group where logical drives are already mapped.

To create a storage partition by assigning the logical drives to new hosts or host groups, complete the following steps:

1. Right-click Default Group or any unmapped host or host group icon in the Topology tree below the Default Group and select **Define** → **Storage Partitioning**. The Storage Partitioning wizard leads you through the process. An unmapped logical drive must exist to start this wizard.

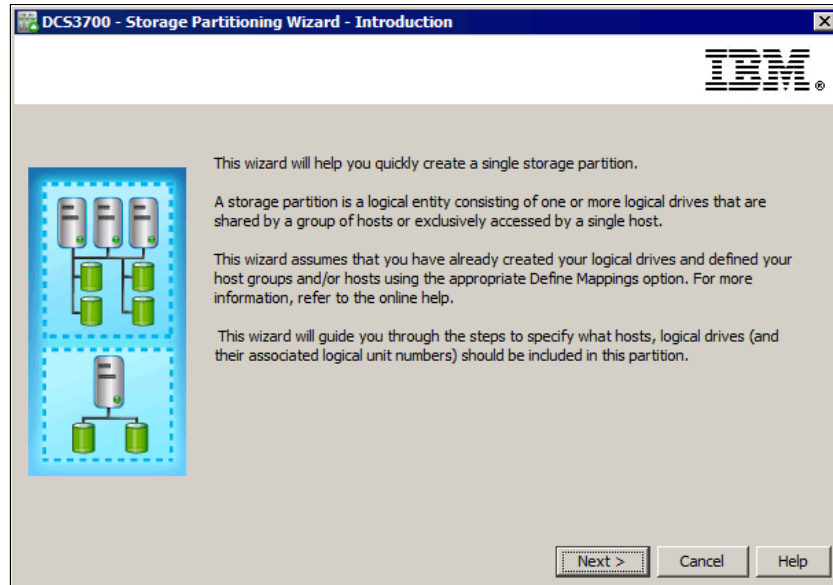


Figure 10-13 Storage Partitioning - Introduction

2. After the introductory window (Figure 10-13), the wizard asks you to select either a host or a host group (Figure 10-14). If you are creating a storage partition for clustered host servers, you must specify the appropriate host group. Otherwise, you can select an individual host. In this window, you can see only hosts and host groups with no logical drives mapped. If you select a host or host group before you start the Storage Partitioning wizard, it is selected in this window.

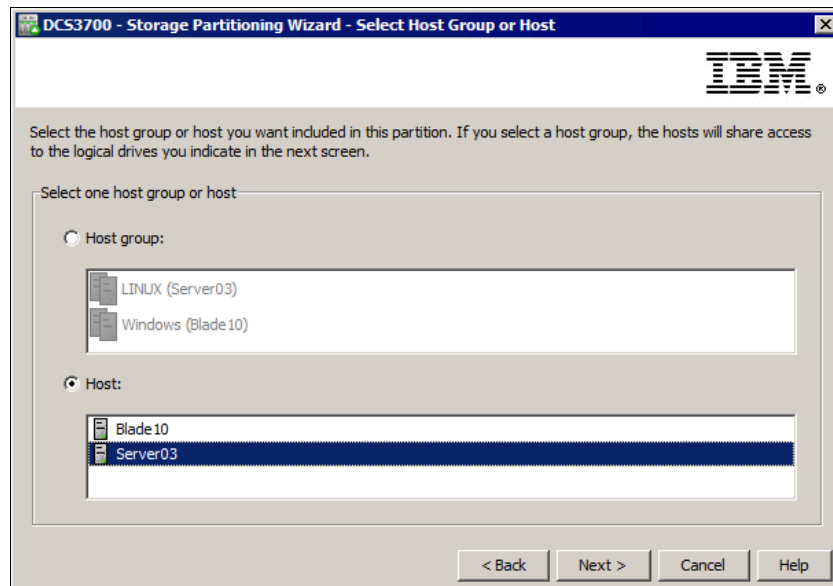


Figure 10-14 Storage Partitioning - Select Host or Host Group

3. The next window (Figure 10-15) 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 in the operating system of the mapped host.

The Access logical drive is used for in-band management (by Fibre Channel) only. For more information, see 1.6.1, “In-band management” on page 15. Do not map the access LUN if you do not plan on using in-band management.

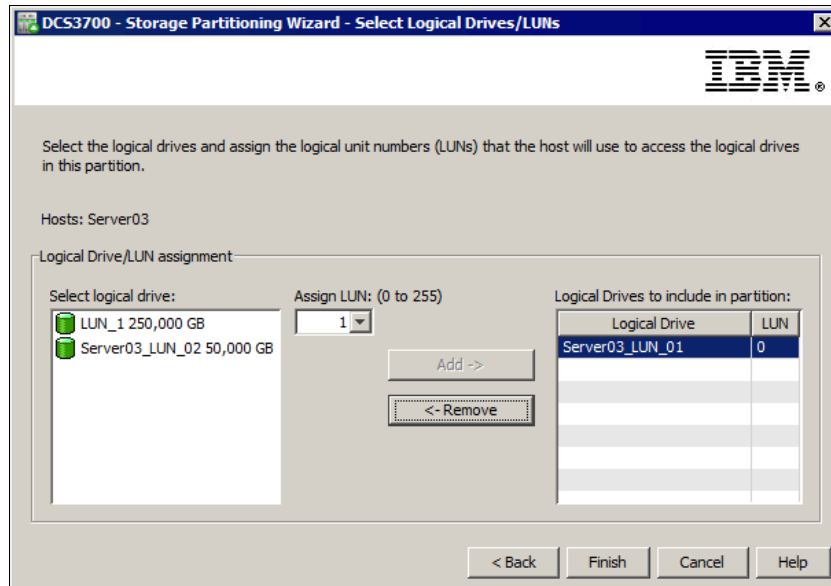


Figure 10-15 Storage Partitioning - Select Logical Drives/LUNs

4. Click **Finish** when you are done selecting the logical drives and assigning the LUNs. The process should be completed, as shown in Figure 10-16.

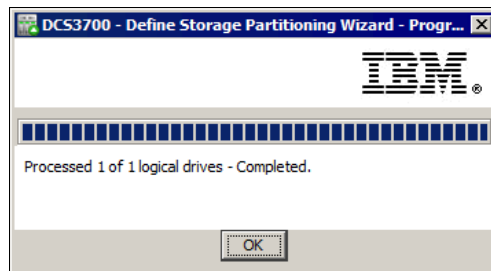


Figure 10-16 Storage Partitioning - Process Completed Status

- 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 that the mapping is there, as shown in Figure 10-17. In the Defined Mappings pane, you can see the logical drive with an assigned LUN number.

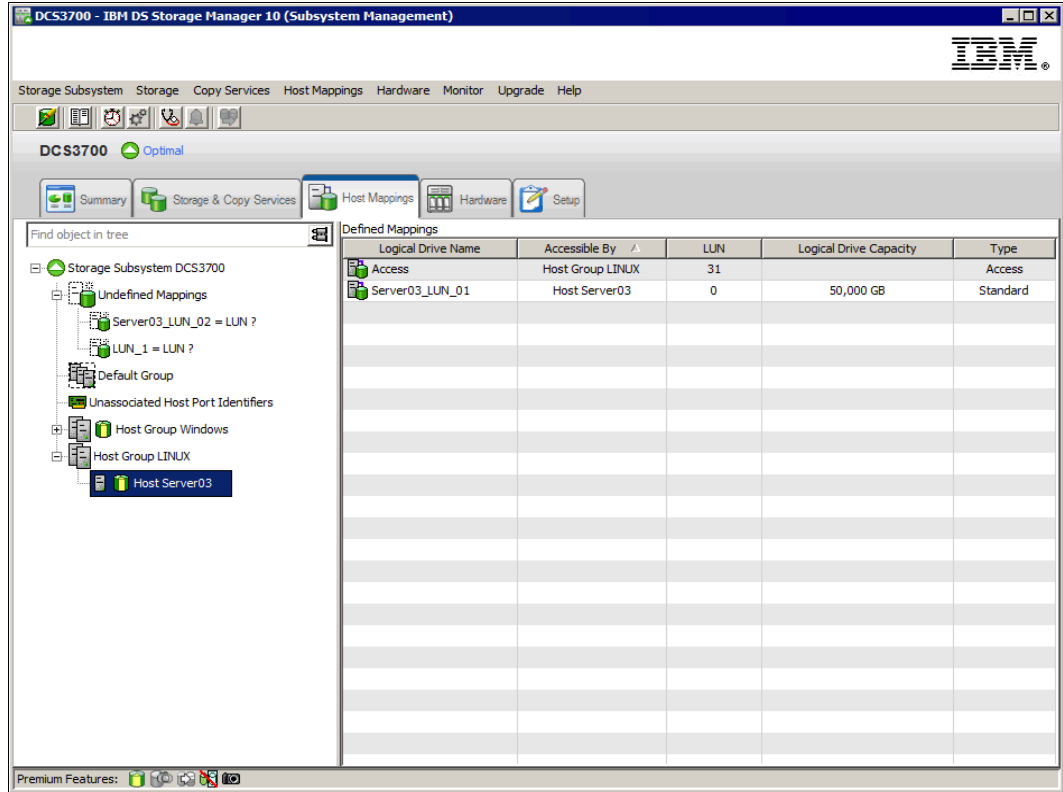


Figure 10-17 Displaying created mappings

Server03 is defined as a member of the host group Linux in our example, and the logical drive is mapped to the host group. The logical drive that is mapped to the host group is automatically mapped to all hosts that are 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 map logical drives to partitions only up to the maximum number that are 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 must 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 assign one particular logical drive so that it could be not shared by other hosts in this host group (for example SAN boot drive), you must use another storage partition for this logical drive and host.

## 10.4 Defining a host group

Because the host groups and hosts can be created together with the Define Host wizard, you might not need to define the host groups from outside the wizard. But, you still can do it and it is possible to start by creating an empty host group definition, and 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 in to it during host definition.

It is also possible to move defined hosts in to the defined host groups, and then move mapping from a host to the host group. We describe moving hosts and mapping changes in 10.8, “Moving, removing, and renaming a host” on page 348 and 10.10, “Changing and removing mapping” on page 351.

To define a host group, complete the following steps:

1. Right-click Default Group and select **Define** → **Host Group**, as shown in Figure 10-18.

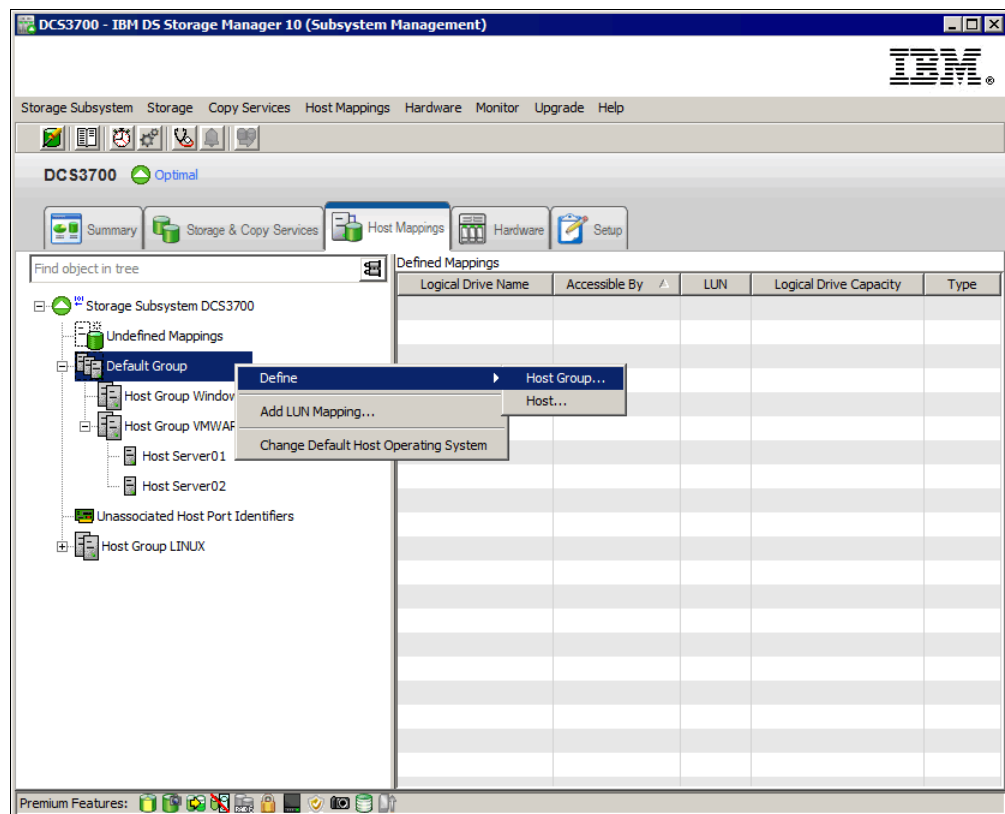


Figure 10-18 Define Host Group menu

- The Define Host Group window (Figure 10-19) opens. Enter the name of the host group that you want to define. Select every host that you want to add to the group and click **Add**. If you select no hosts to add, the host group stays empty and you can move existing hosts in to 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**.

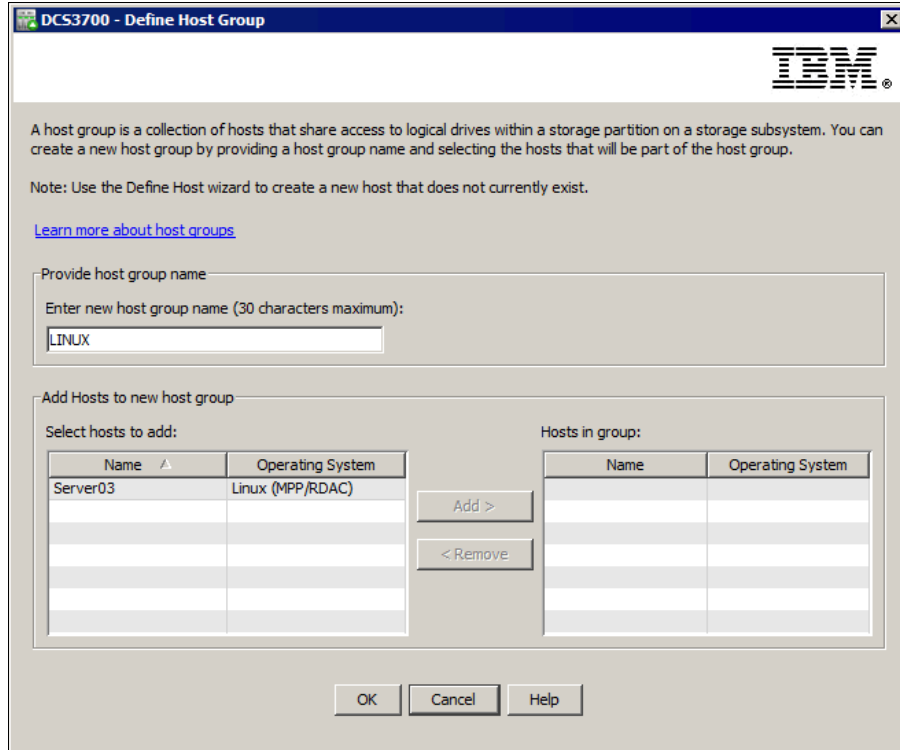


Figure 10-19 Define Host Group name

Because storage partitioning of the DCS3700 storage subsystem is based on the WWPNs of the host ports, the definitions for the host groups and the hosts represent only 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. For a final review before you confirm the operation, see Figure 10-20.

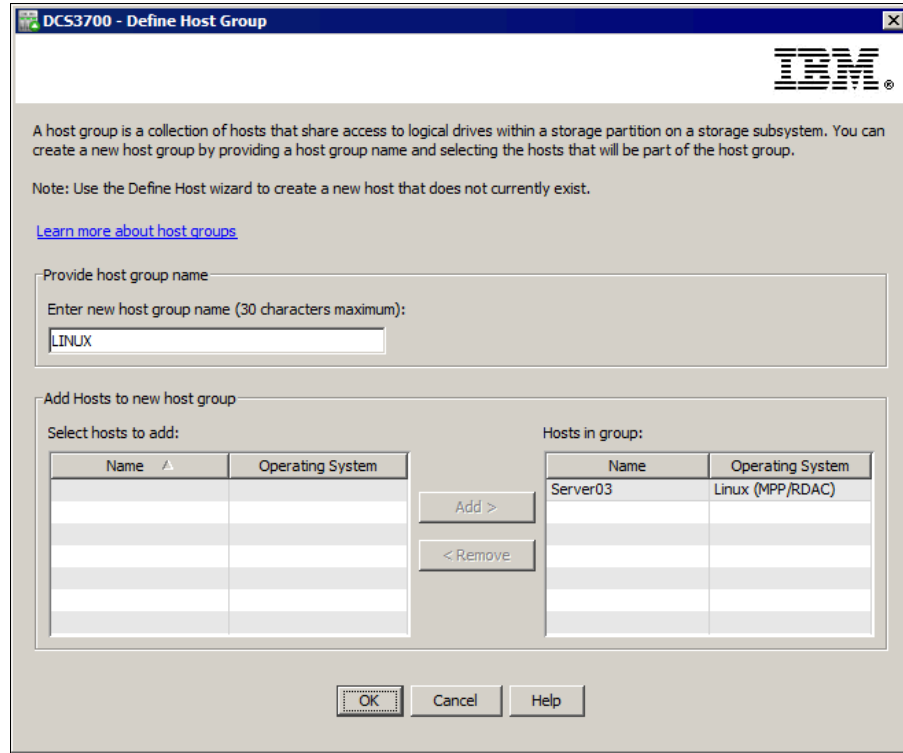


Figure 10-20 Define Host Group review

## 10.5 Manage Host Port Identifiers option

For environments with multiple hosts and attachment types, you can have a single source of information about the available host ports on your servers. To accomplish this task, 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 opens, as shown in Figure 10-21.

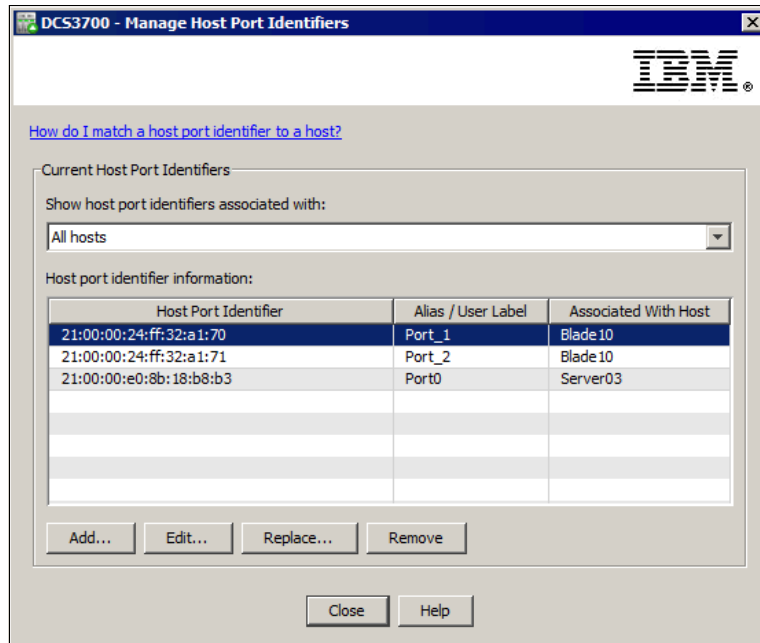


Figure 10-21 Manage Host Port Identifiers

Select this option if you need to review your port configuration assignment, add, remove, and change port settings, and whenever you must 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).

## 10.6 Define Additional Mapping option

A particular host or a host group can be part of a certain storage partition. This means that the logical drives are already mapped to this host or host group. If you must map more logical drives to this host or host group, use the Define Additional Mapping option by completing 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**. You can also select a Logical drive in the Undefined mappings section of the Topology tree, right-click it, and select **Define Additional Mapping**.
2. In the Define Additional Mapping window (Figure 10-22 on page 347), select the logical drive that you want to map to this host group or host, and assign the correct LUN number. Your host group or host is already preselected. If you select a logical drive in the Undefined Mapping section for additional mapping, you must select the host or the host group where it must be mapped and the correct LUN number. In our example, the logical drive is preselected.



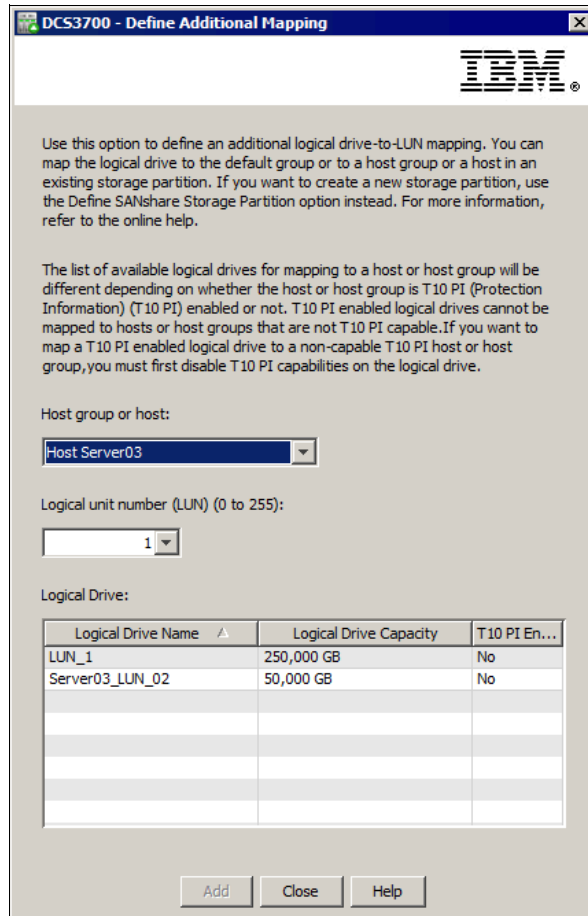


Figure 10-22 Define Additional Mapping - select logical drive

3. Click **Add**. The selected mapping is active immediately. You can repeat this action and add processes until no unmapped logical drives are left. All hosts and all host groups are available for additional mapping, and only unmapped logical drives can be selected in the lower part for mapping.
4. To make this added logical drive (or drives, if you select more) available to the host systems without rebooting, you must ensure that all host bus adapters are rescanned for new devices, and that the devices are assigned within the operating system. This can be done by using tools that are based in the operating system. This procedure is documented in the sample configuration chapters later in this book.

## 10.7 Viewing unassociated ports

Previous sections described how to assign host ports (HBA in servers) to the host. To check what other host ports are already recognized by your DCS3700 storage subsystem but are not assigned to a host yet, click **Host Mappings** → **View Unassociated Host Port Identifiers**, or click the **Unassociated Host Port Identifiers** icon in the Topology tree, right-click it, and select **View**. If there are any unassociated host ports, they are shown in the window that is shown in Figure 10-23.

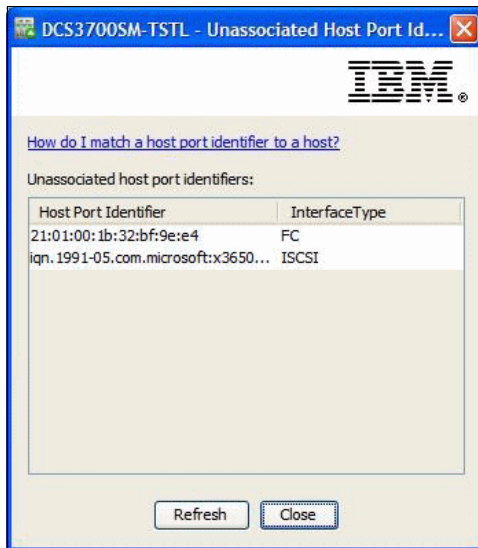


Figure 10-23 Unassociated port identifiers

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 is the case with Fibre Channel or SAS. For more information about how to configure iSCSI hosts, see Chapter 12, "Administration: iSCSI" on page 363.

## 10.8 Moving, removing, and renaming a host

If you right-click any host in the Topology tree, it opens a menu where you can move, remove, or rename the host.

### 10.8.1 Move option

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 10-24 on page 349).

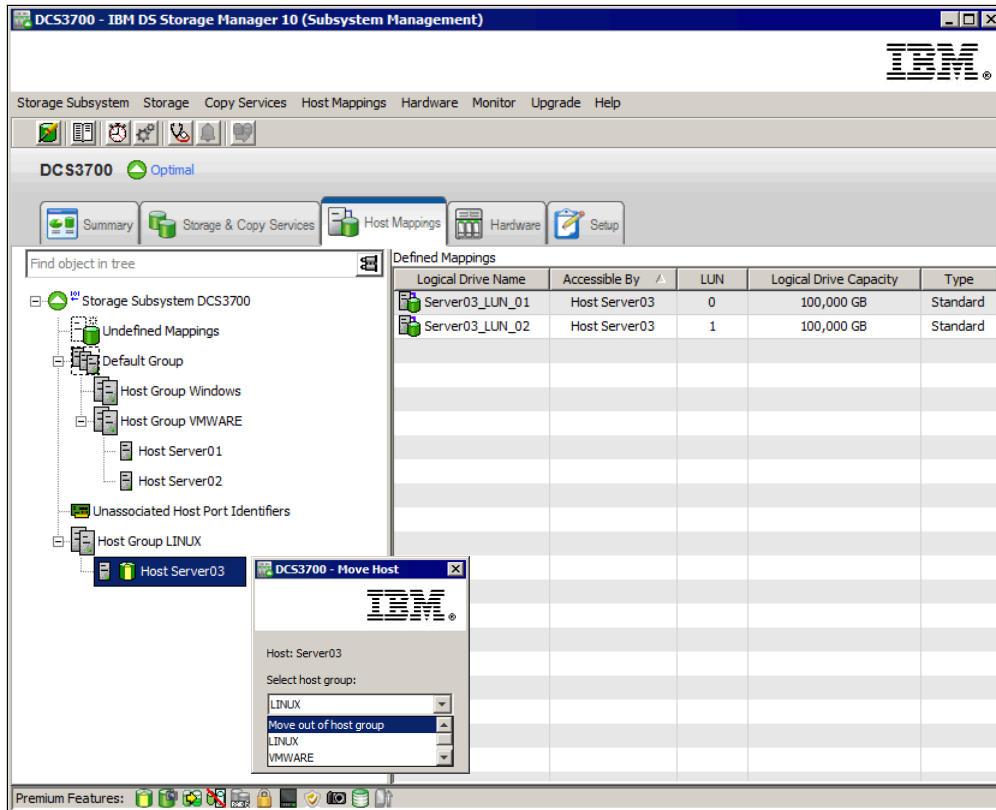


Figure 10-24 Move a host from a host group or to another group

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 that are 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 must change LUN number for the host and rediscover the logical drive in this host first.

## 10.8.2 Remove option

Use this option to delete host or host group definitions. Host port identifiers that are assigned to a deleted host are moved and they are still visible in the Unassociated host port identifiers window. Before you remove a host, a warning window opens, as shown in Figure 10-25.

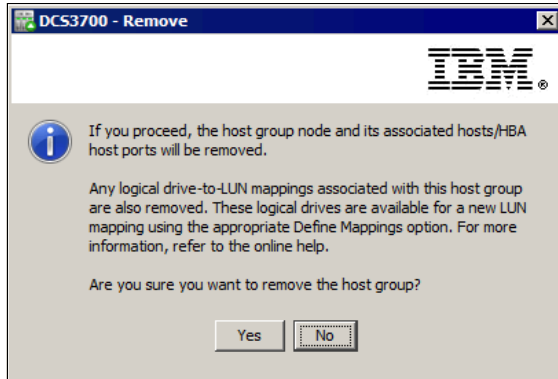


Figure 10-25 Remove a host

## 10.8.3 Rename option

This option is used to rename a host or host group, as shown in Figure 10-26.




Figure 10-26 Rename a host

**Note:** Be careful if SMcli scripts are applied on your DCS3700 storage subsystem. If you rename your host and its name is used in an SMcli script, the script will no longer work. Changes of the host mapping are used almost in every DR process.

## 10.9 Changing a host type

It is not difficult to reinstall the operating system on server hardware. If the OS is changed, you do not have to define this server as a new host. It is possible to change the host type (operating system), leaving the other definitions such as host ports and mapped LUNs unchanged.

To change the host operating system, complete the following steps:

1. Right-click the host on which the operating system must be changed in the Topology tree and select **Change Operating System**. The window that is shown in Figure 10-27 opens.

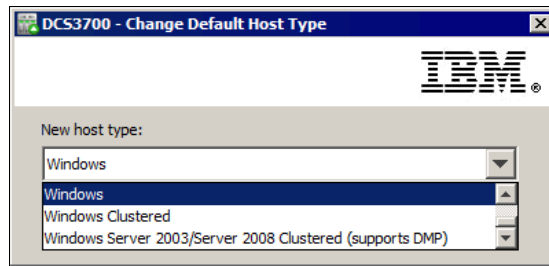


Figure 10-27 Change host type

2. Select the required operating system and click **OK**. The change is accepted immediately.

## 10.10 Changing and removing mapping

Until now, all described mapping configuration changes were provided in the Topology pane. In the Defined Mapping pane, you can see mapped logical drives. The logical drives displayed there are mapped to the host that is selected in the Topology pane. Select a logical drive, right-click it, and two options appear:

- ▶ Change Mapping
- ▶ Remove Mapping

### 10.10.1 Change Mapping option

Select this option to change a host or host group, where the logical drive is newly mapped, or you can change LUN number. The Change Mapping window is shown in Figure 10-28.

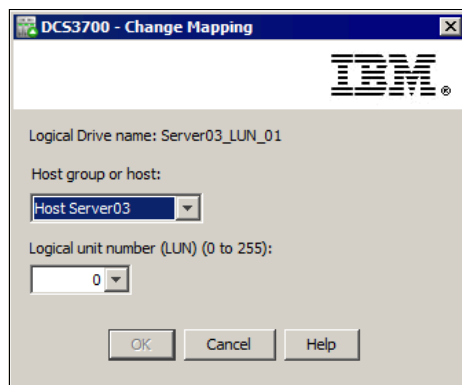


Figure 10-28 Change Mapping - host and LUN number

Click **OK** and a warning window opens (Figure 10-29).

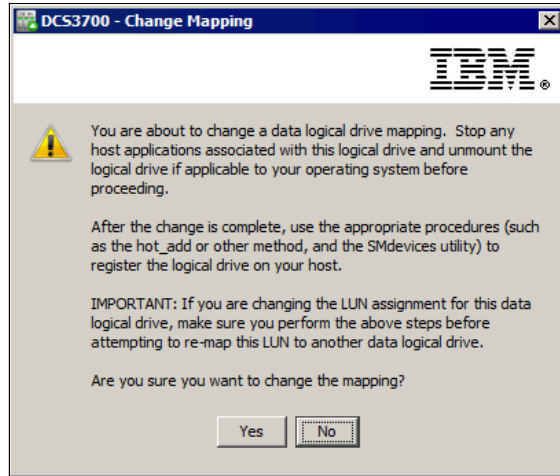


Figure 10-29 Change Mapping - warning

**Note:** Even if only a 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 the mapping on running hosts.

### 10.10.2 Remove Mapping option

If you remove mapping, the selected logical drive is no longer accessible by the host server. However, the data is still on the logical drive and it can be used later for other purposes. Before you remove any mapping, a warning window opens, as shown in Figure 10-30.

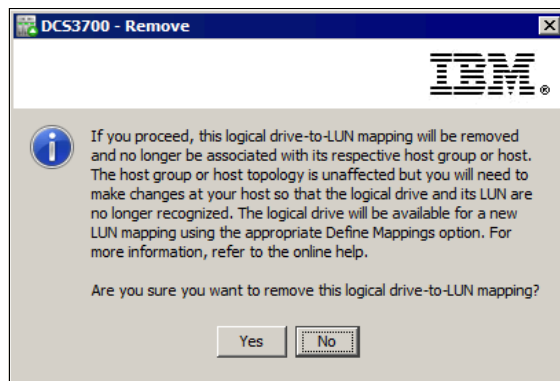


Figure 10-30 Remove Mappings - warning



## Administration: Setup tab

This chapter describes the Setup tab of the DS Storage Manager. This tab can be used in the initial configuration of an IBM System Storage DCS3700 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 capacities on the storage subsystem, and others. You can click a link to open the corresponding dialog or wizard.

This chapter covers the following tasks:

- ▶ Locate Storage Subsystem
- ▶ Rename Storage Subsystem
- ▶ Change Hardware View Order
- ▶ Set a Storage Subsystem Password
- ▶ Manage Premium Features
- ▶ Configure iSCSI Host Ports
- ▶ Create Storage
- ▶ Save Configuration
- ▶ Manually Define Hosts
- ▶ Map Logical Drives
- ▶ Configure Ethernet Management Ports
- ▶ Manage iSCSI Setting

If you successfully provide all the base configurations that are referenced in this chapter, you can be almost certain that your subsystem is configured correctly.

This chapter describes all these tasks and provides you with the appropriate links for more information about a particular task as necessary.

# 11.1 Setup tab

Figure 11-1 shows the Setup tab in the DS Storage Manager.

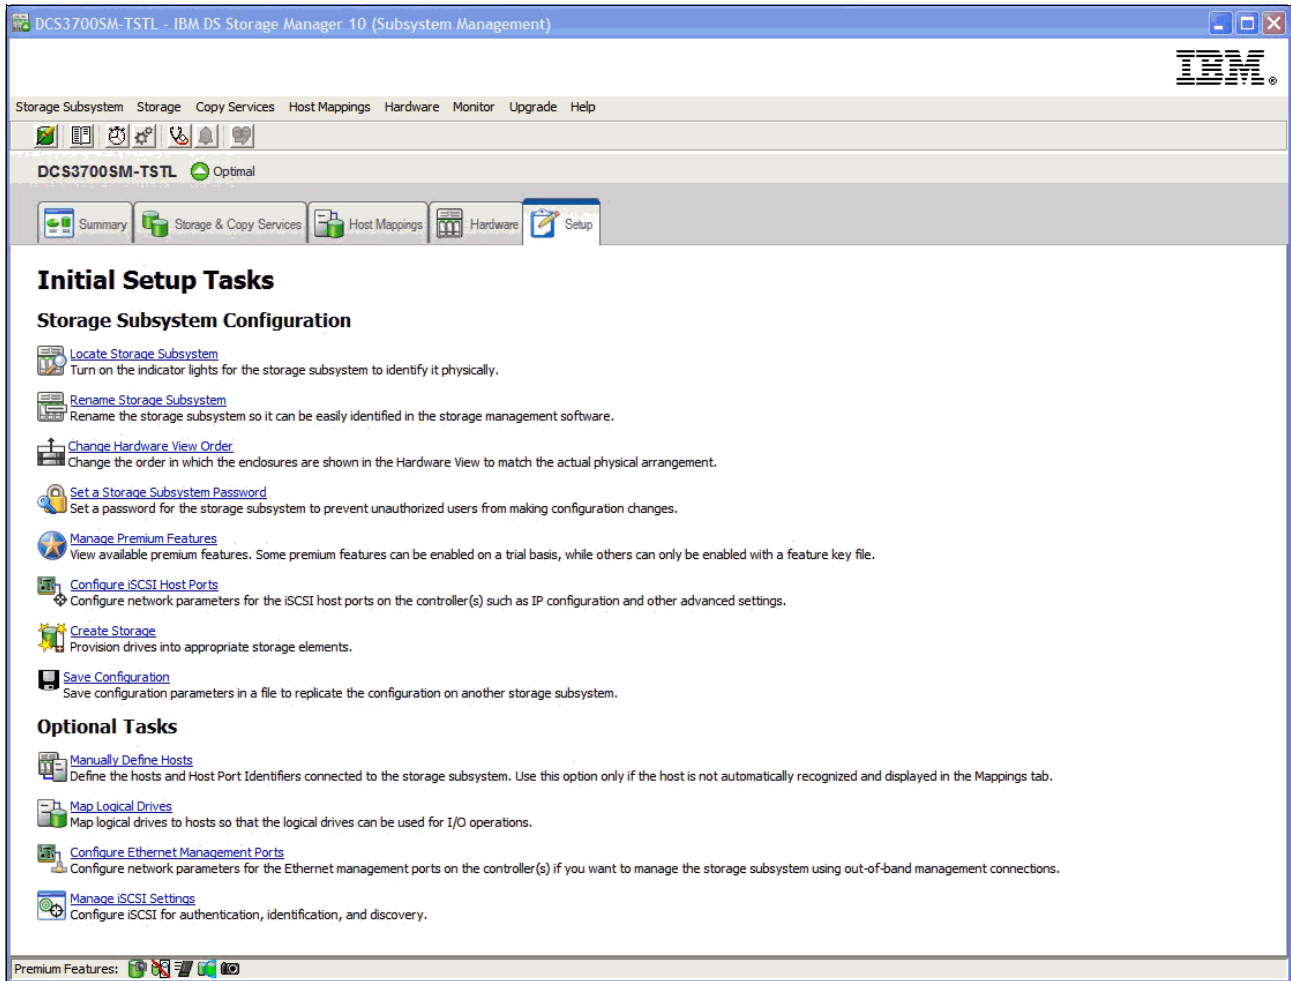


Figure 11-1 Setup tab of DS Storage Manager

The following sections describe the links from top to bottom.

## 11.1.1 Locate Storage Subsystem

The Locate Storage Subsystem link causes the blue LEDs of your DCS3700 storage subsystem and all its DCS3700 expansion units to flash so you can find where your subsystem is placed in the rack. To stop the flashing, click **OK**, as shown in Figure 11-2 on page 355.



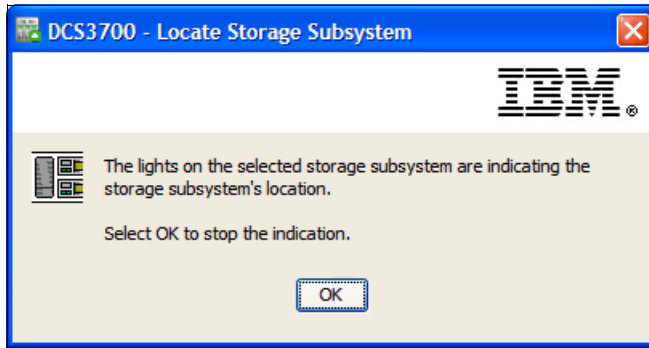


Figure 11-2 Locate Storage Subsystem

### 11.1.2 Rename Storage Subsystem

The Rename Storage Subsystem link opens a window where you can change the name of the DCS3700 storage subsystem. When you add a new storage subsystem to your Storage Manager, it is identified as Unnamed. If you have several storage subsystems that are installed in your environment, give each subsystem a meaningful name, as shown in Figure 11-3.

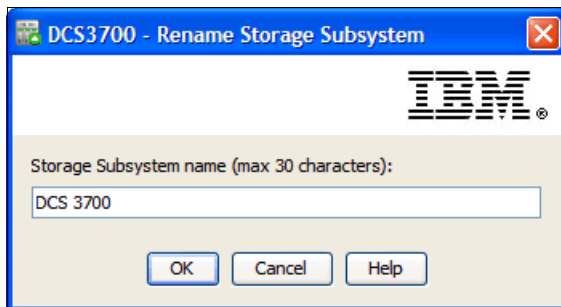


Figure 11-3 Rename Storage Subsystem

The subsystem name is immediately updated in the Enterprise Manager window.

**Note:** It is important to know that renaming the storage subsystem can influence certain services running outside the Storage Manager. For example, if you run customized SMcli scripts where the subsystem is referenced by its name, the script will no longer work correctly. Also, certain types of multipath failover drivers on servers use the name of the attached subsystem. Rename a new subsystem before it is in production.

### 11.1.3 Change Hardware View Order

Use this link to change the order of the controller enclosures and the drive expansion units in the Physical pane to match the hardware configuration in your storage subsystem. The Physical pane that initially appears is a default view that might not match your storage subsystem. The enclosure order change remains in effect until it is modified again (see Figure 11-4).

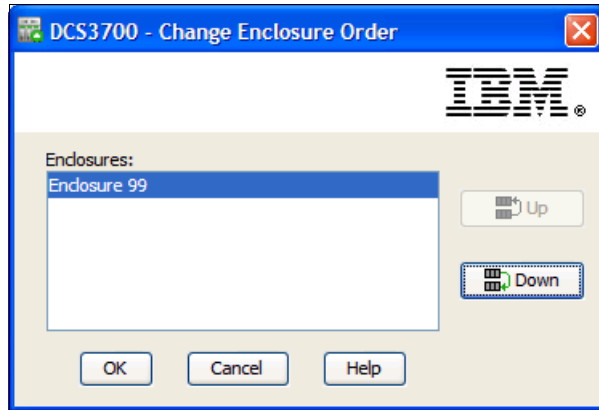


Figure 11-4 Change Enclosure Order

### 11.1.4 Set a Storage Subsystem Password

This link opens a window (Figure 11-5) where you can change the password for the storage subsystem. For a DCS3700 storage subsystem, strong password rules are set, so the password should be changed the first time that you open the DS Storage Manager for a new DCS3700 storage subsystem. You can find more details about DCS3700 password settings in “Security: Set Password” on page 208.



Figure 11-5 Set Password

If you are setting the password for the first time on a new DCS3700 storage subsystem, the Current password field remains empty.

## 11.1.5 Manage Premium Features

Select this link to activate your optional licenses for Premium Features, such as storage partitioning, copy services, and so on. This procedure is described in “Premium Features” on page 205. Also, you find your Feature Key Identifier here (see Figure 11-6).

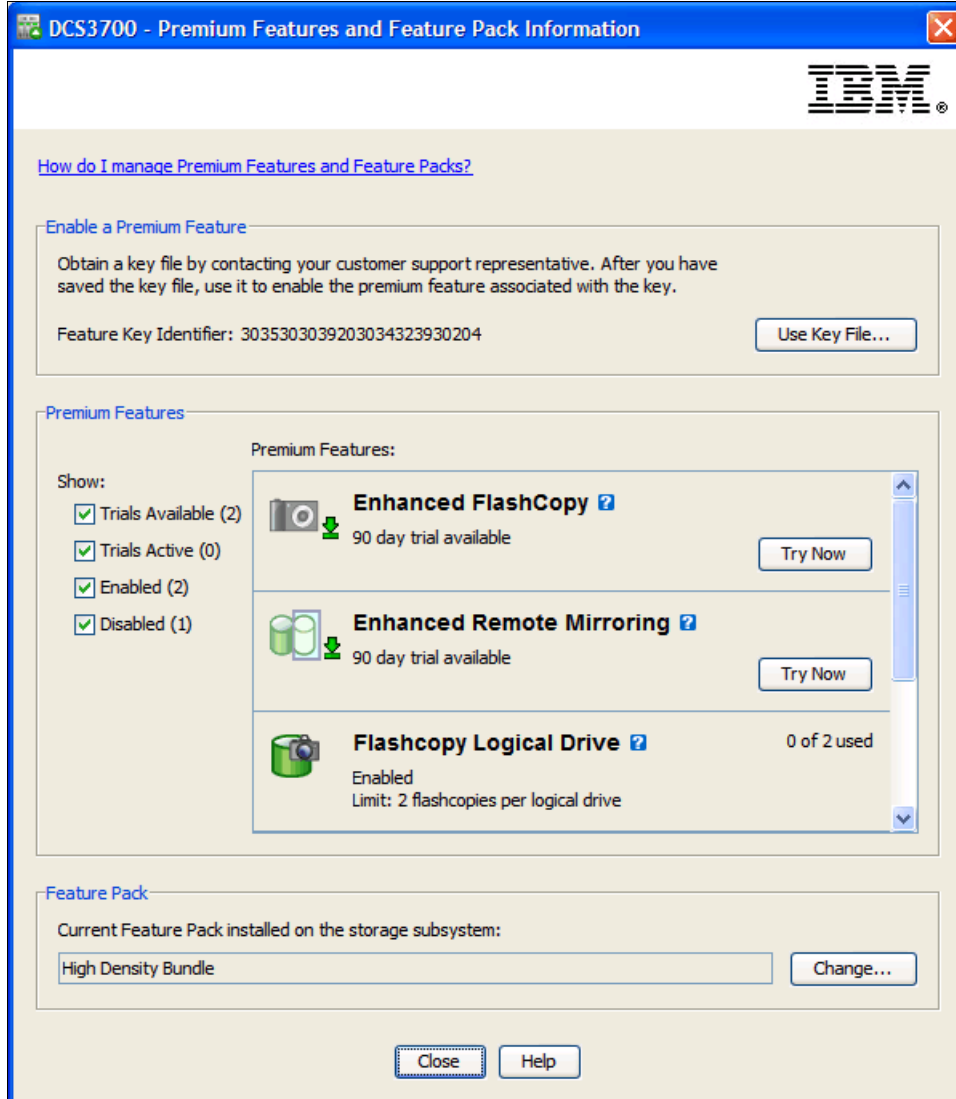


Figure 11-6 Premium Features, Feature Key Identifier, and Feature Pack Information

## 11.1.6 Configure iSCSI Host Ports

This link opens a window where you can configure iSCSI host ports (see Figure 11-7). This link is available only when you install the optional iSCSI Host Interface Card (HIC). For more information about configuring iSCSI Host Ports, see Chapter 12, “Administration: iSCSI” on page 363.

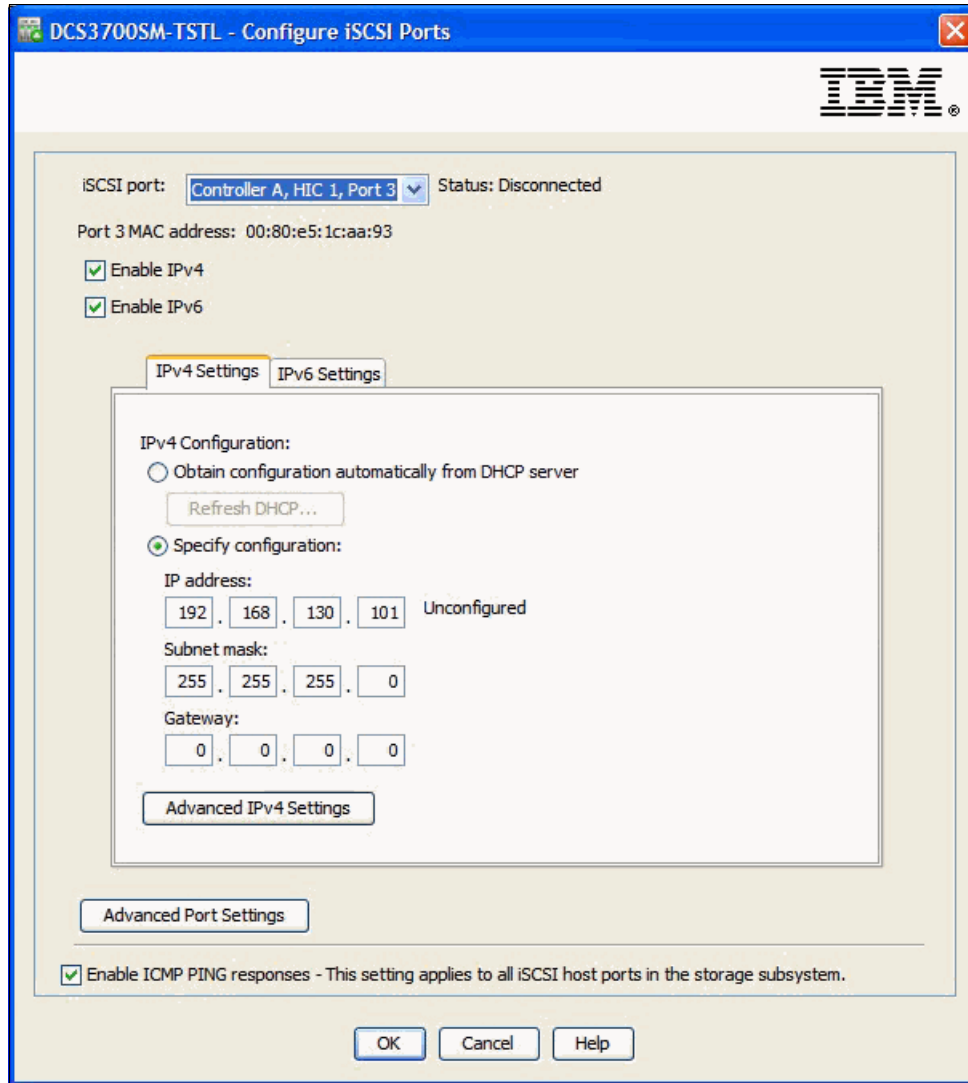


Figure 11-7 Configure iSCSI Host Ports

## 11.1.7 Create Storage

Click this link to create storage. Figure 11-8 on page 359 shows the different possibilities. For more information about how to use the Automatic or Manual Configuration, see Chapter 8, “Administration: Storage & Copy Services tab” on page 225.

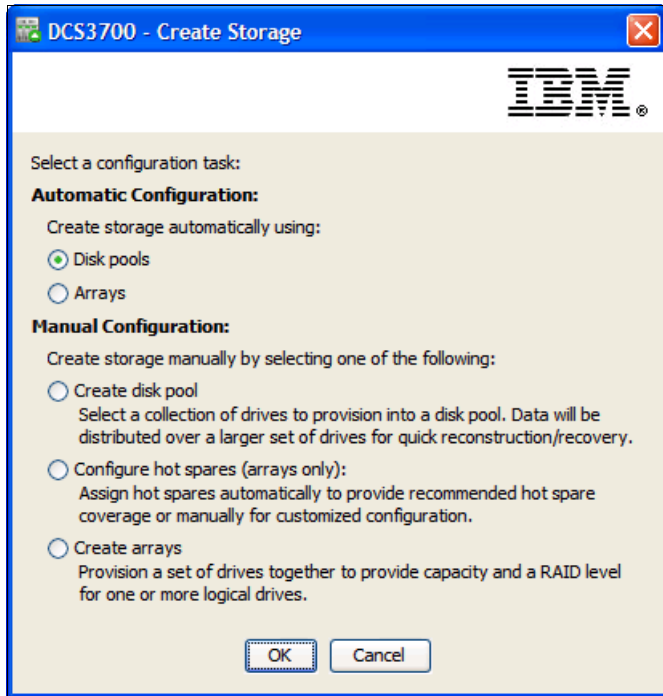


Figure 11-8 Create Storage

The following automatic configurations are possible:

- ▶ Create Disk Pools

A set of minimum 11 drives that is logically grouped. A disk pool provides the overall capacity that is needed to create one or more logical drives. A disk pool is similar to an array, with the following differences. The data in a disk pool is stored randomly on all of the drives in the disk pool, unlike data in an array which is stored on the same set of drives. You do not specify a RAID level for a disk pool. A disk pool does not use hot spare drives. And, a disk pool allows many drives to be grouped.

- ▶ Create Arrays

A set of drives that is logically grouped and assigned a RAID level. Each array that is created provides the overall capacity that is needed to create one or more logical drives.

The following manual configurations are possible:

- ▶ Create Disk Pools

A set of (minimum of 11) drives that is logically grouped. A disk pool provides the overall capacity that is needed to create one or more logical drives. A disk pool is similar to an array, with the following differences. The data in a disk pool is stored randomly on all of the drives in the disk pool, unlike data in an array, which is stored on the same set of drives. You do not specify a RAID level for a disk pool. A disk pool does not use hot spare drives. A disk pool allows many drives to be grouped.

- ▶ Configure hot spares (arrays only)

A spare drive that contains no data and that acts as a standby in case a drive fails in a RAID 1, RAID 3, RAID 5, or RAID 6 logical drive. The hot spare drive can replace the failed drive in the logical drive. Hot spare drives are used only in arrays, not disk pools.

► Create arrays

A set of drives that is logically grouped and assigned a RAID level. Each array that is created provides the overall capacity that is needed to create one or more logical drives.

### 11.1.8 Save Configuration

Use this link to open a window (Figure 11-9) 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, see Chapter 5, “Administration: Enterprise Management” on page 159.

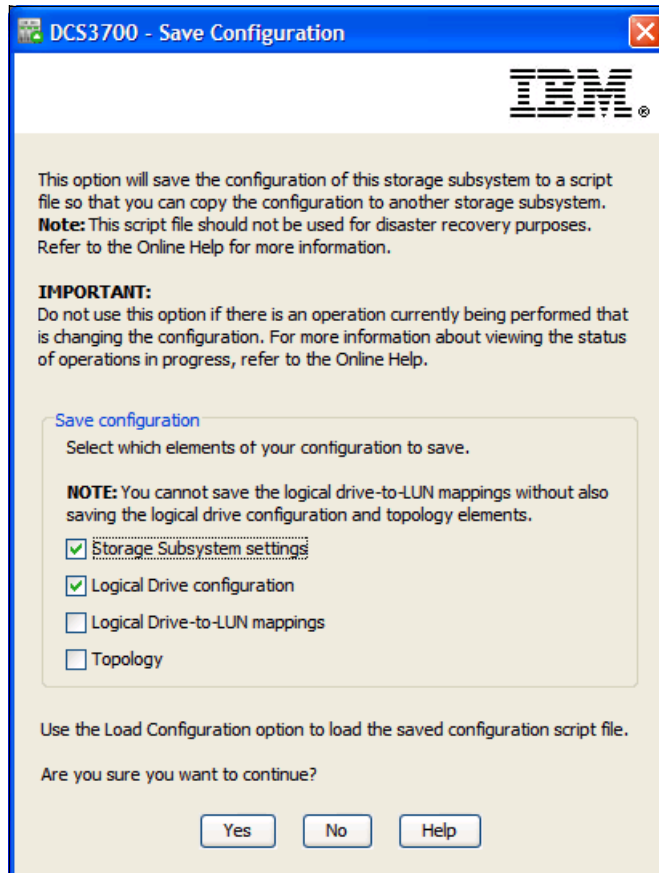


Figure 11-9 Saving a storage subsystem configuration

When your DCS3700 storage subsystem is fully configured, you can use the Save Configuration option to save storage subsystem logical settings, logical drive configuration, topology, and mappings to a text/script file that can be used by the SMcli. 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 logical devices in your storage subsystem. More information about SMcli is in Chapter 15, “Command-line interface (CLI)” on page 463.

**Note:** You should not use this generated CLI script for disaster recovery for storage subsystems running IBM DS Storage Manager 10 V10.83 or later. Instead, to do a system restore, use the configuration database backup file that you create manually or by having the Persistent Monitor running in the background.

Clicking **Configuration** → **Save** 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 (except Optimal) and states of the arrays

**Risk of application errors:** Do not use the Save Configuration option if the storage subsystem is undergoing an operation that changes any logical configuration setting. Examples of these operations include creating or deleting logical drives, changing storage partitioning settings, downloading controller firmware, assigning or modifying hot spare drives, or adding capacity (drives) to an array.

## 11.1.9 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 10.2, “Defining a host” on page 330. All other configuration steps that are needed for assigning logical drives to hosts are described in Chapter 10, “Administration: Mappings tab” on page 327.

## 11.1.10 Map Logical Drives

This link in the Setup tab of DS Storage Manager opens a window that is used to assign an existing logical drive to existing hosts. This window is described in 10.6, “Define Additional Mapping option” on page 346. All other configuration steps that are needed for assigning logical drives to hosts are described in Chapter 10, “Administration: Mappings tab” on page 327.

## 11.1.11 Configure Ethernet Management Ports

This link enables you to change the IP addresses of your storage subsystem. Follow the steps and comments that are described in 9.7, “Setting the Ethernet management ports” on page 323.

## 11.1.12 Manage iSCSI Setting

This link allows you to change your settings regarding Target Configuration, Remote Initiator Configuration, and Target Discovery (see Figure 11-10). This link is available when only you install the optional iSCSI HIC. For more information about managing iSCSI Settings, see Chapter 12, “Administration: iSCSI” on page 363.

DCS3700SM-TSTL - Manage iSCSI Settings

IBM

Target Configuration Remote Initiator Configuration Target Discovery

Target Identification

iSCSI name:  
iqn.1992-01.com.lsi:2365.60080e50001beed00000000506d6cb1

iSCSI alias (optional - max 30 characters):

Target Authentication

**Important:** If an initiator requires mutual authentication, you must also provide information in the Remote Initiator Configuration tab.

Select the authentication method that will be supported by the target:

None: An initiator attempting to access the target is not required to be authenticated by the target.  
**Warning:** This option does not provide data security because ANY initiator will be able to access this target.

CHAP: An initiator attempting to access the target using CHAP authentication must provide the target's CHAP secret.

CHAP secret defined: No CHAP Secret

OK Cancel Apply Help

Figure 11-10 Manage iSCSI Settings





## Administration: iSCSI

This chapter describes iSCSI management options that are available in the DS Storage Manager. These options are only for IBM System Storage DCS3700 storage subsystems with optional iSCSI daughter cards that are installed.

The iSCSI settings are configured from the DS Storage Manager Storage Subsystem drop-down menu, as shown in Figure 12-1 on page 366.

**Important:** This publication presents only the iSCSI configurations options for the DCS3700 storage subsystem. The host configuration procedures for iSCSI are similar to the DS3500 storage subsystem and can be found in *IBM System Storage DS3500 Introduction and Implementation Guide*, SG24-7914

## 12.1 Planning for iSCSI attachment

In addition to using SAS and Fibre Channel as the interface connection method, iSCSI host interfaces of DCS3700 storage subsystems allow servers with hardware iSCSI cards or with software running over regular Ethernet cards to run the iSCSI protocol as the connection method to attach the DCS3700 storage subsystems. As with SAS or FC, before beginning an iSCSI deployment, plan, understand, and document the network topology that will be used.

### 12.1.1 Hardware initiators

For more information about which iSCSI hardware initiators might be supported for attachment to a DCS3700 storage subsystem, periodically check the supported types and versions at the IBM interoperability website found at:

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

**Note:** At the time of the writing of this book, there is no iSCSI hardware initiator that is supported by the DCS3700 storage subsystem.

Before you install and configure these adapters, make sure that you install the latest management application and the latest firmware code that is supported by IBM.

### 12.1.2 Software initiators

You can check for the supported iSCSI software initiators for various operating systems on the IBM interoperability website found at:

<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 the writing of 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.

### 12.1.3 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, use 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.

To define hosts and storage partitions in iSCSI, you must plan for iSCSI addresses for the host ports, and use the iSCSI Qualified Name (IQN) of the host you want to map.

## 12.1.4 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, to protect the integrity of the data in your DCS3700 storage subsystem and its continuous access, try to isolate the iSCSI traffic in a dedicated network. The iSCSI multipathing architecture provides failover to the alternative controller if there is 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 by 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 for several optional configuration parameters, including enabling jumbo frames, configuring 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 that are supported are 1501 - 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 DCS3700 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 DCS3700 series of storage subsystems. You can set the Ethernet priority of the target iSCSI interfaces to increase the class of service that is received within the network itself.

## 12.1.5 Security

Unlike FC SANs or direct SAS connections, Ethernet networks can be more open, so to provide additional security, configure the following additional authentication protocols on the DCS3700 storage subsystems:

- ▶ The Internet Storage Name Service (iSNS) protocol allows for automated discovery, management, and configuration of iSCSI devices on a Internet Protocol 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 DCS3700 storage subsystem.

## 12.2 iSCSI configuration summary

This section covers the following settings of iSCSI configuration:

- ▶ Managing iSCSI Settings
  - Target Authentication
  - Mutual Authentication

- Target Identification
- Target Discovery
- ▶ Configuring iSCSI host ports on a DCS3700 storage subsystem
- ▶ Session and statistics
  - View iSCSI Statistics
  - View/End iSCSI Sessions

This section also covers differences in iSCSI host server attachment by describing the following topics:

- ▶ View unassociated iSCSI initiators.
- ▶ Define a iSCSI host.
- ▶ Manage iSCSI initiators.

The main menu for iSCSI configuration can be found in the Storage Subsystem drop-down menu, as shown in Figure 12-1.

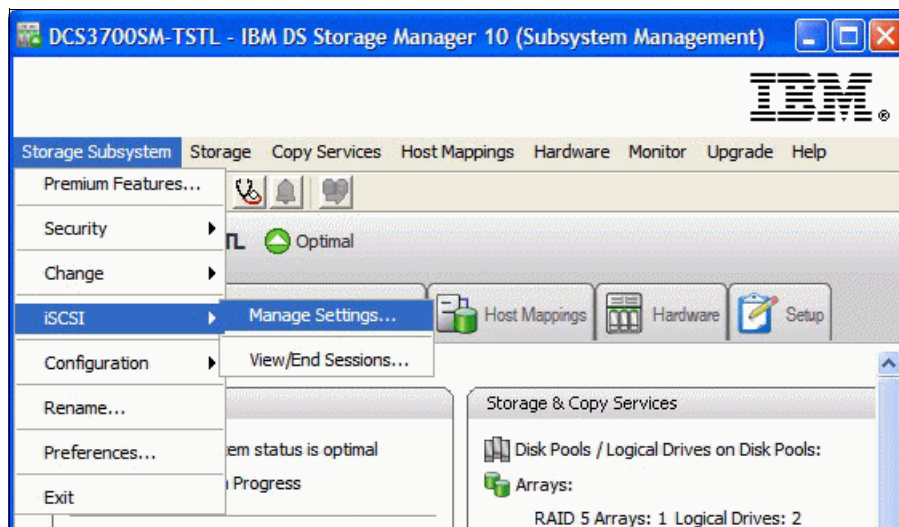


Figure 12-1 Basic iSCSI configuration setting menu

Links for the configuration of iSCSI ports on DCS3700 storage subsystems and for the configuration of iSCSI protocol settings are also accessible from the DS Storage Manager Setup tab that is shown in Figure 11-1 on page 354. These settings are the most important, and should be done on a new DCS3700 storage subsystem during the initial configuration, as shown in Figure 12-2.

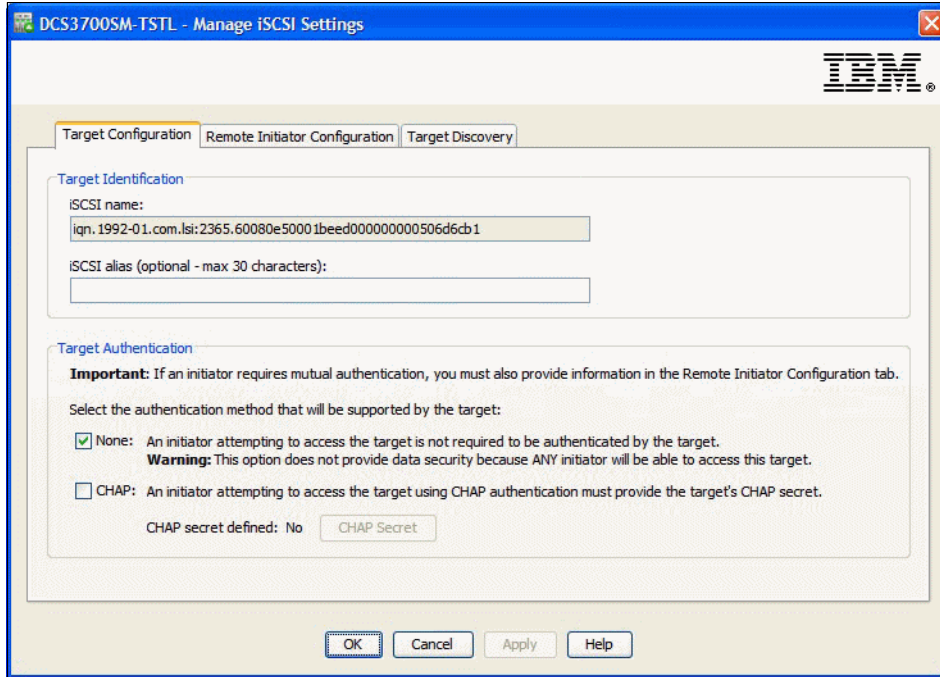


Figure 12-2 Manage iSCSI Settings in Initial Setup Tasks

## 12.3 Manage iSCSI protocol settings

Settings that affect the security and discovery of iSCSI connections can be done here. Section 12.1.5, “Security” on page 365 describes iSCSI security.

### 12.3.1 Target authentication

The Target Authentication (also called *unidirectional authentication*) setting allows only configured iSCSI initiators in servers to access the storage subsystem. When target authentication is enabled, the target (DCS3700 storage subsystem) must 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 called *mutual authentication* (also called *bidirectional authentication*) can be configured only on the iSCSI initiator when the initiator itself authenticates against the storage subsystem. If this is required, you must also configure an Initiator CHAP secret on the DCS3700 storage subsystem by completing the steps in 12.3.2, “Mutual authentication” on page 369.

To configure target authentication, complete the following steps:

1. Click **Storage Subsystem** → **iSCSI** → **Manage Settings** → **Target Configuration** and you see two options in the Target Authentication (Figure 12-3) pane.

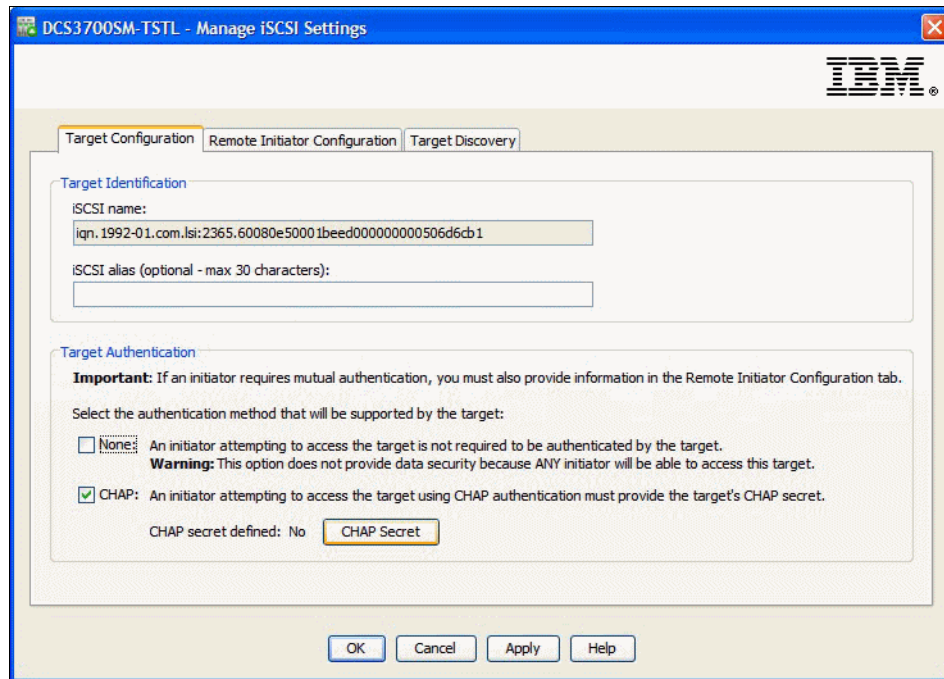


Figure 12-3 Manage iSCSI Settings - Target Authentication

Here are the two options:

- None
- CHAP

This setting affects the connection between an iSCSI initiator and DCS3700 iSCSI ports. The None option allows any iSCSI initiator to establish an iSCSI connection to this target. When you select **CHAP**, an initiator is required to provide a CHAP password to establish a session. CHAP must be enabled if mutual authentication is required by an iSCSI Initiator. Both options (None and CHAP) can be enabled concurrently; if so, 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 complexity, you should first set up and test all connections without CHAP, and implement CHAP later.

2. If CHAP is enabled, define the CHAP target secret by clicking **CHAP Secret**, as shown in Figure 12-3. The Enter Target CHAP Secret window that is shown in Figure 12-4 on page 369 opens.
3. Enter a CHAP secret that is at least 12 characters long in to the first field. Enter the secret in the second field to confirm it. Alternatively, click **Generate Random Secret** to generate a 16 character long secret.

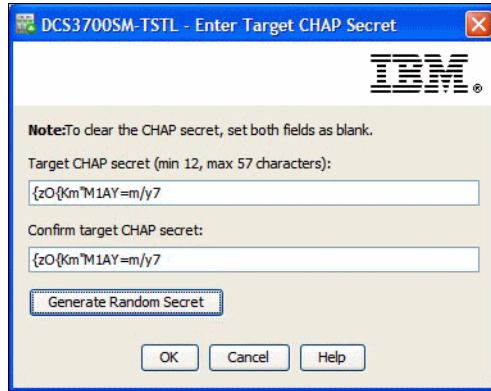


Figure 12-4 Target CHAP secret

4. Click **OK** to return to the Target Authentication window. Click **OK** to save the modifications.

### 12.3.2 Mutual authentication

With mutual authentication, you configure a CHAP secret that a specific initiator must use to gain access to the storage subsystem. Mutual authentication is also called initiator authentication or bidirectional authentication. Each initiator can have its own CHAP secret.

To activate mutual authentication, complete the following steps:

1. In the Manage iSCSI Settings window (Figure 12-3 on page 368), click the **Remote Initiator Configuration** tab. The Remote Initiator Configuration tab is shown in Figure 12-5.

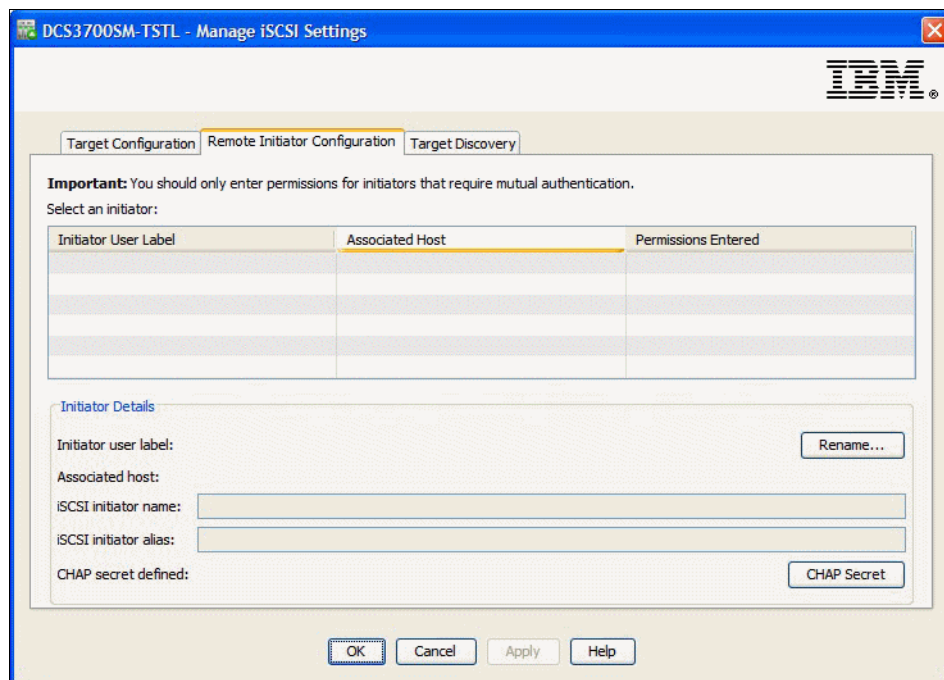


Figure 12-5 Remote initiator or mutual authentication configuration



- In the window that is shown in Figure 12-6, select a host port that is defined, as described in 12.7.2, “Defining a new iSCSI host” on page 382. The host port is referred to as the initiator label. In this example, there is no CHAP secret that is defined (No permissions entered) for this initiator, as shown in the Permission Set? column. Click **CHAP Secret**.

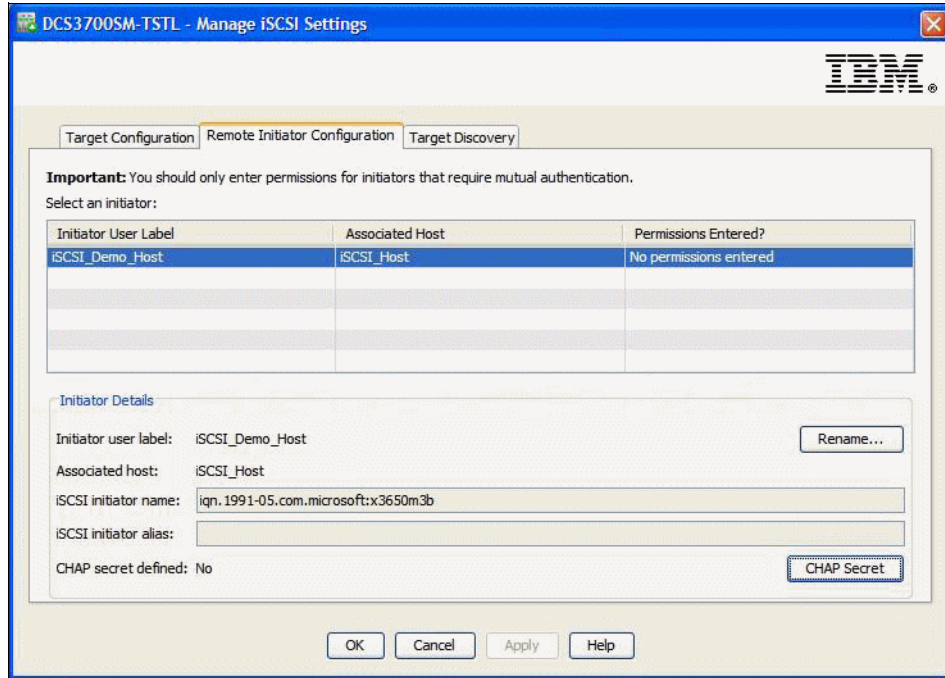


Figure 12-6 iSCSI mutual authentication

- In the window that is shown in Figure 12-7, enter an initiator CHAP secret with at least 12 characters twice. The same initiator CHAP secret must be set for an iSCSI initiator on the host server. Click **OK** to continue.

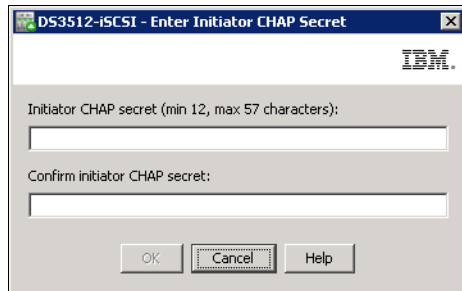


Figure 12-7 Storage Manager - initiator CHAP secret

- The Permission Set column in Figure 12-6 now indicates that the port has a CHAP secret defined. Define CHAP secrets for all the remaining host ports, if needed, by using the same procedure. Each host with iSCSI initiator can have a separate initiator CHAP secret. When this task is complete, click **OK** to exit this task.

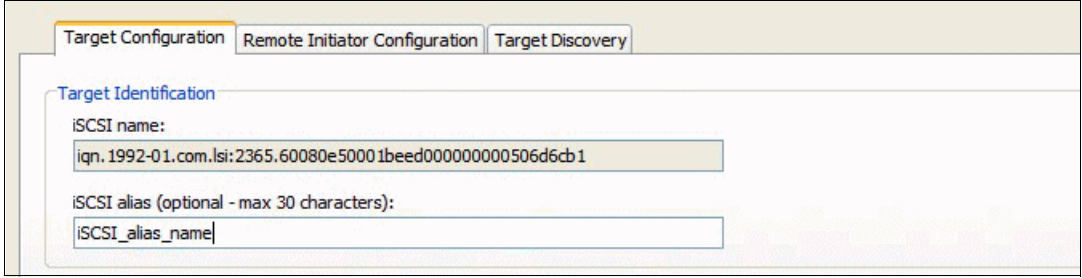


### 12.3.3 Target identification

You can use the Target Identification tab to set up an alias name for the iSCSI target device (your DCS3700 storage subsystem) that can be discovered by iSCSI initiators on hosts. In the iSCSI protocol, the iSCSI target is the storage, while on the host server the iSCSI initiator is the storage. You can check your DCS3700 storage subsystem IQN name in this tab as well.

Set up an iSCSI target alias by completing the following steps:

1. Click **Target Configuration**. You see the Target Identification pane.
2. Enter an alias name in the SCSI alias field, as shown in Figure 12-8, and click **OK**.



The screenshot shows a window titled "Target Configuration" with three tabs: "Target Configuration", "Remote Initiator Configuration", and "Target Discovery". The "Target Configuration" tab is active, showing a "Target Identification" section. It contains two text input fields: "iSCSI name:" with the value "iqn.1992-01.com.lsi:2365.60080e50001beed00000000506d6cb1" and "iSCSI alias (optional - max 30 characters):" with the value "iSCSI\_alias\_name".

Figure 12-8 iSCSI target identification

### 12.3.4 Target discovery

The techniques for target discovery are described in 1.3, “iSCSI disk attachment”, of *IBM System Storage DS3500 Introduction and Implementation Guide*, SG24-7914. You can use the Target Discovery tab in the Manage iSCSI Setting window (Figure 12-9 on page 372) to specify whether an Internet Storage Name Service (iSNS) is used to register this iSCSI target so that it is more easily found by iSCSI initiators.

To change the default target discovery, complete the following steps:

1. Click the **Target Discovery** tab.
2. Select **Use iSNS server**, as shown in Figure 12-9.

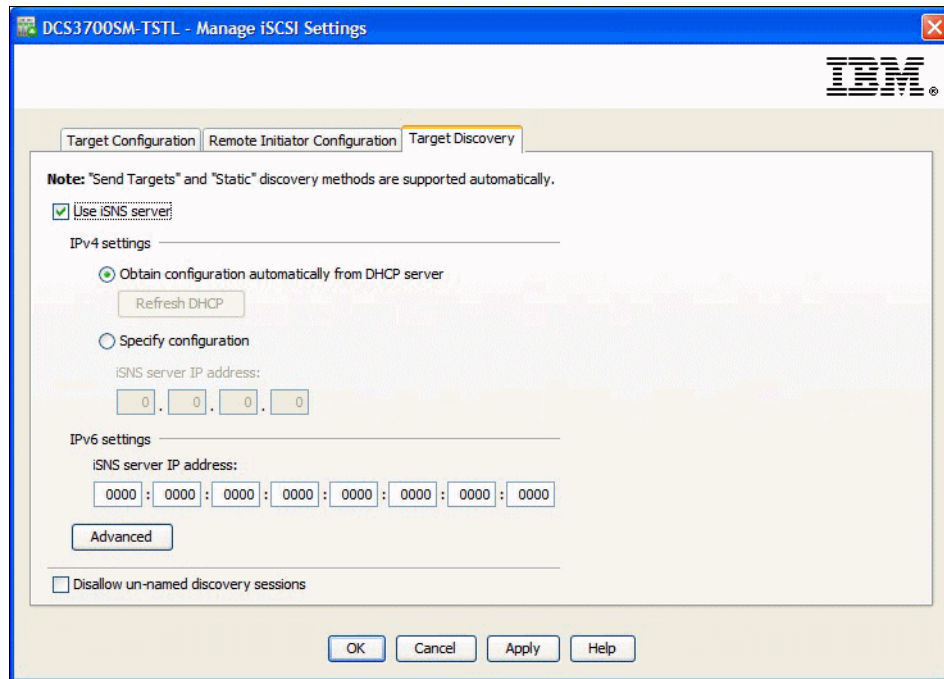


Figure 12-9 Target Discovery - Enable iSNS server

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 unnamed discovery sessions** if you want your storage subsystem to respond only to requests that are 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.

## 12.4 Configuring iSCSI host ports

The iSCSI host ports are the 10 Gb Ethernet ports on each DCS3700 controller with a iSCSI daughter card that are used to connect hosts to this storage subsystem by using the iSCSI protocol. Each controller has two RJ45 10 Gb Ethernet ports that are available on the iSCSI daughter card. The DCS3700 storage subsystem has two 10 Gb optical SFP connectors per controller.

The Configure iSCSI Host Ports task (see Figure 12-11 on page 374) is activated by clicking **Storage Subsystem** → **Setup** → **Configure iSCSI Host Ports** (see Figure 12-10 on page 373). 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.

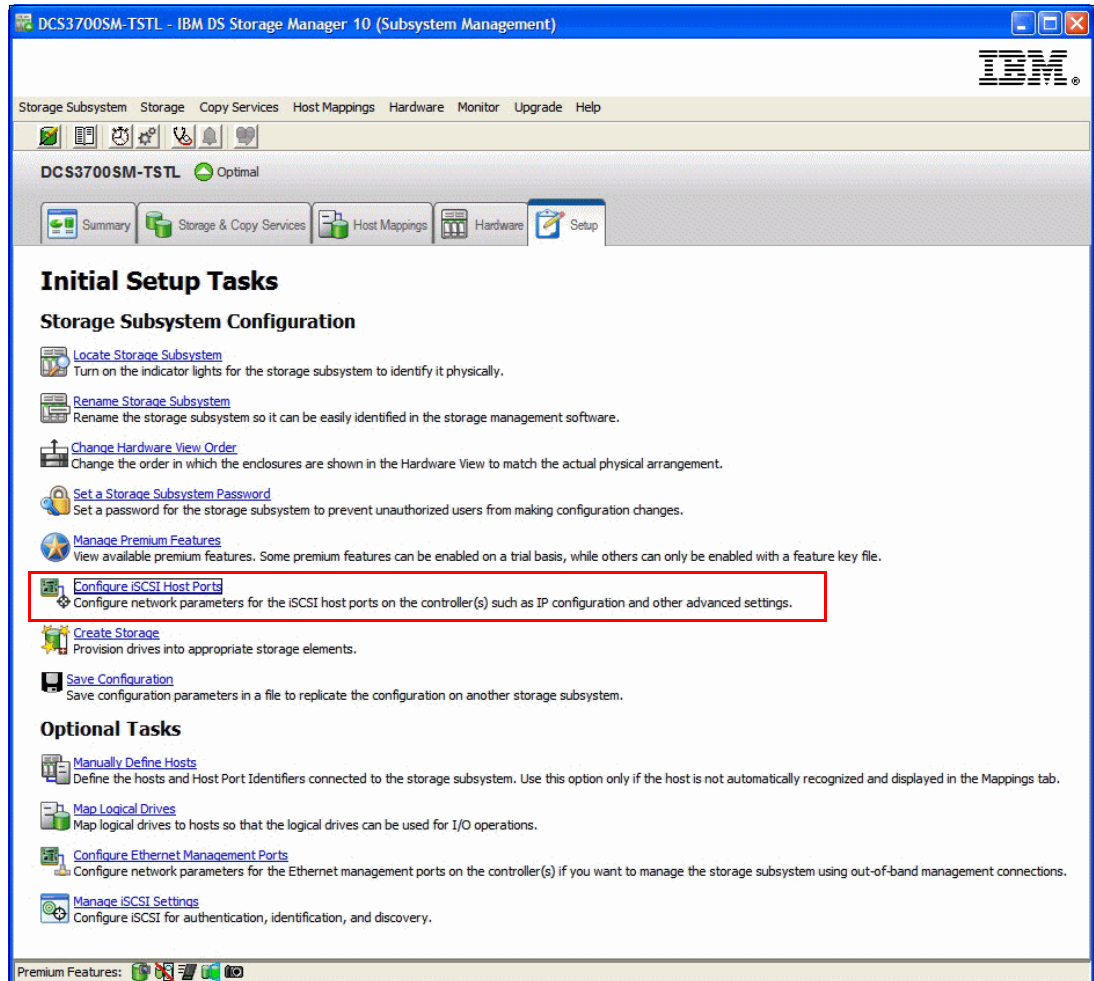


Figure 12-10 Configure iSCSI Host Ports

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 that are 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, complete the following steps:

1. Click **Storage Subsystem** → **Setup** → **Configure iSCSI Host Ports**. A window opens, as shown in Figure 12-11. Our example window shows an active IPv4 configuration with the default IP address for the iSCSI port.

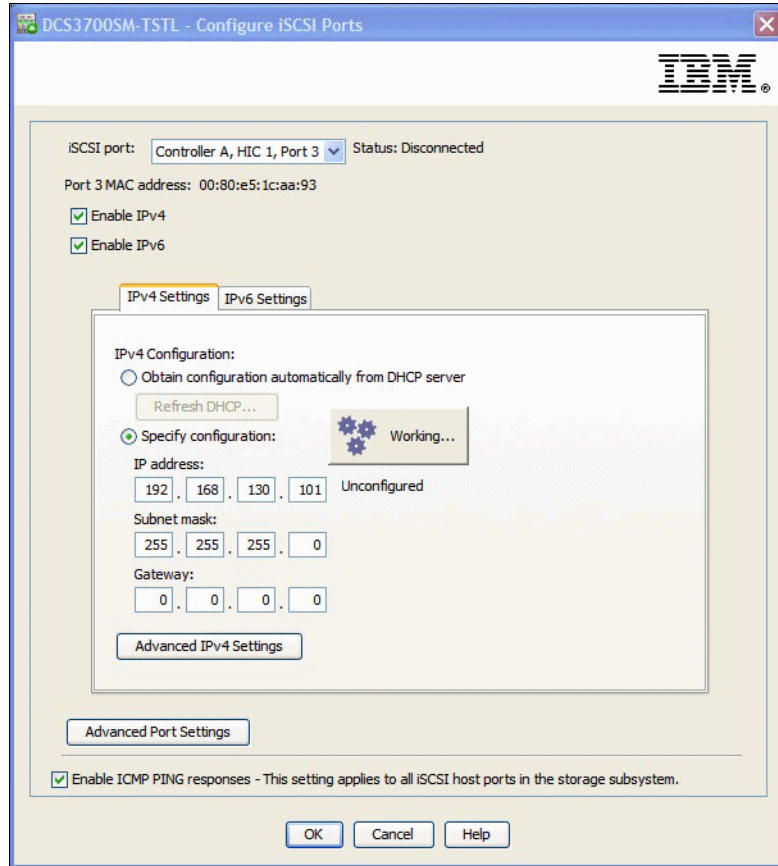


Figure 12-11 Configure iSCSI host ports

2. Select the iSCSI host port to be configured from the **iSCSI port** drop-down menu, as shown in Figure 12-12.

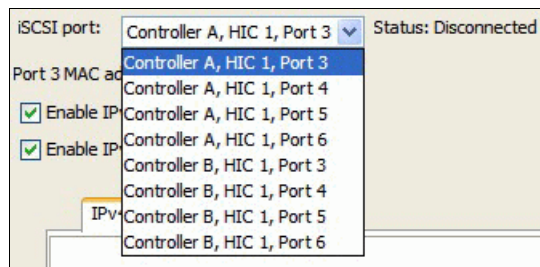


Figure 12-12 iSCSI host port selection

The default subnet mask is 255.255.255.0 and the default IP addresses for iSCSI ports are the following ones:

- 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, clear the **Enable IPv6** check box.
  4. On the IPv4 Settings tab, click **Obtain configuration automatically from DHCP server** if the host port address is 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 DCS3700 storage subsystem. If static IP reservations are not provided or you lose the DHCP configuration, the iSCSI initiators can lose connection to the DCS3700 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 the Advanced IPv4 Settings window (Figure 12-13), you can enable native VLAN support and specify a VLAN ID that is used by the selected port. In addition, the Ethernet priority can be set here. Similar configurations are available for the IPv6 settings.

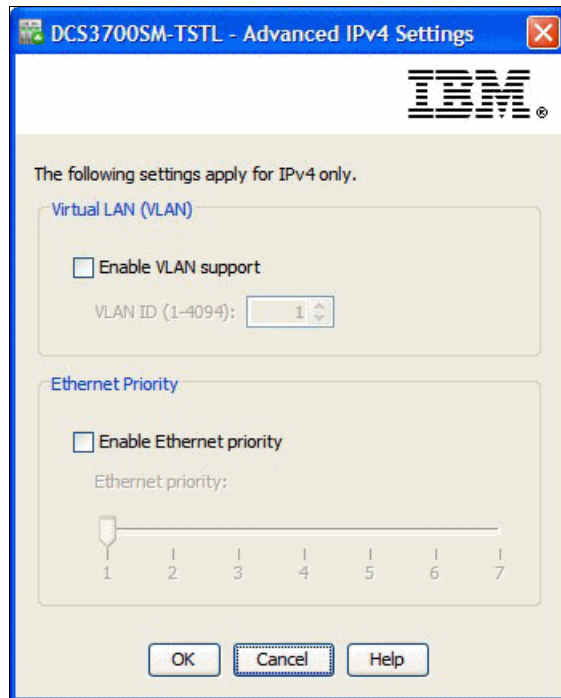


Figure 12-13 Advanced IPv4 Settings

- After configuring IPv4 and IPv6, you can select **Advanced Host Port Settings**, as shown in Figure 12-14. 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 49152 - 65535.

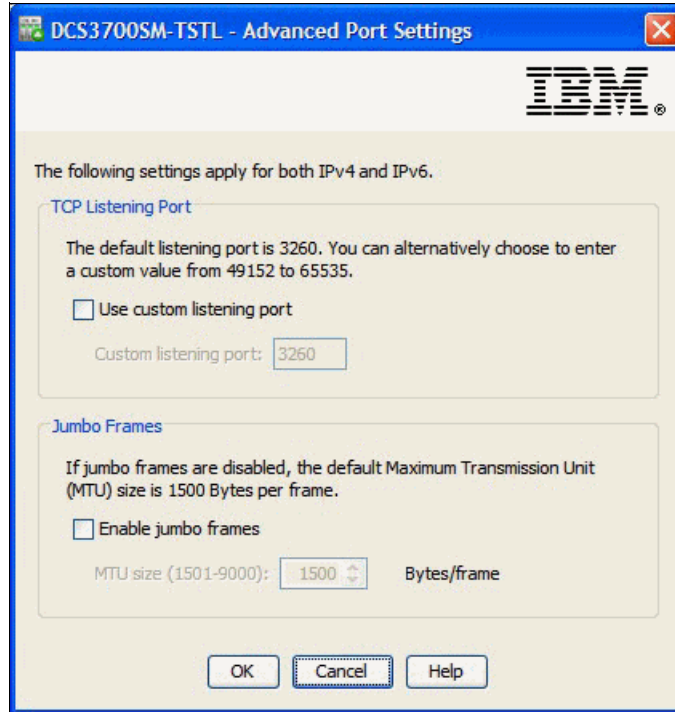


Figure 12-14 Advanced iSCSI host port settings

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 1501 - 9000 bytes per frame.

- Configure the remaining iSCSI host ports in the **iSCSI host port** drop-down menu (Figure 12-12 on page 374) in the same way.
- In bottom part of the Configure iSCSI Host Ports window (Figure 12-11 on page 374), you can disable ping (ICMP echo) requests for all iSCSI host ports. A reset of all ports is required to apply this setting.
- After setting all the options, click **OK** to save the modifications. In Figure 12-15 on page 377, you are advised which ports will be reset to apply the configuration. Active iSCSI sessions that are connected to these ports are closed. Click **Yes** to proceed.



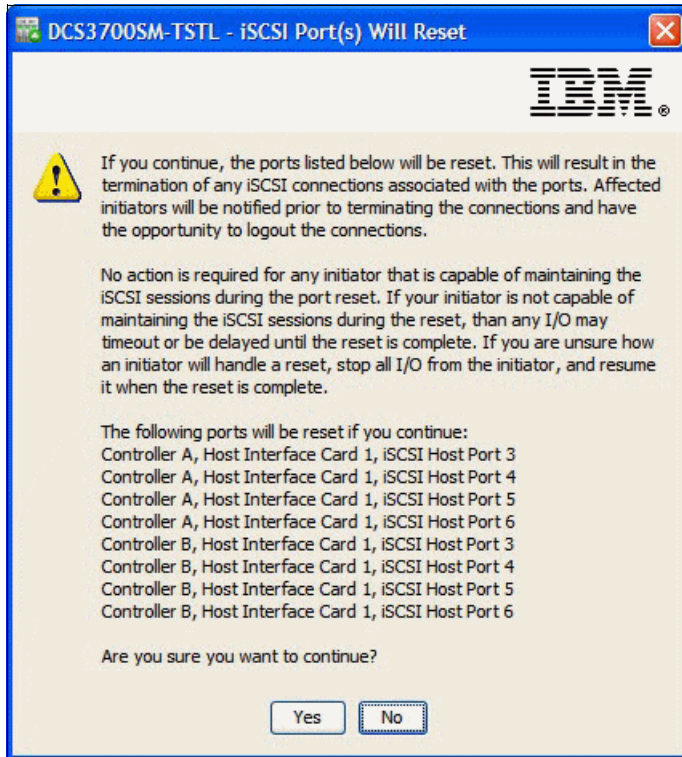


Figure 12-15 iSCSI port reset warning

## 12.5 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 by completing the following steps:

1. Click **Storage Subsystem** → **iSCSI** → **View/End iSCSI Sessions** (Figure 12-16).

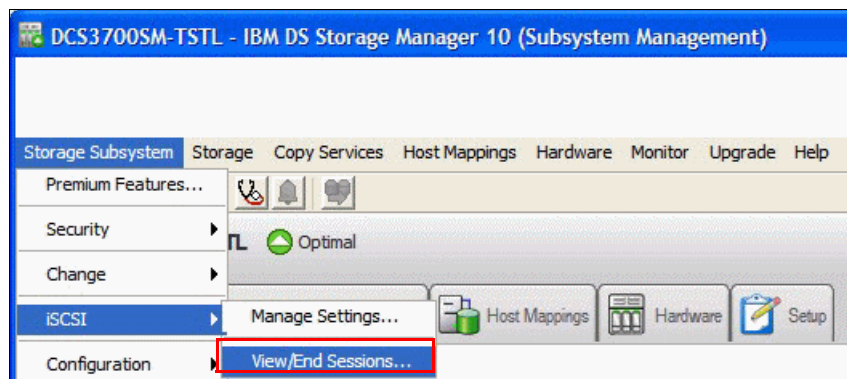


Figure 12-16 End an iSCSI session

- The window that is shown in Figure 12-17 opens. In the top pane, you can see all the sessions. Select a session to see its details in the lower pane.

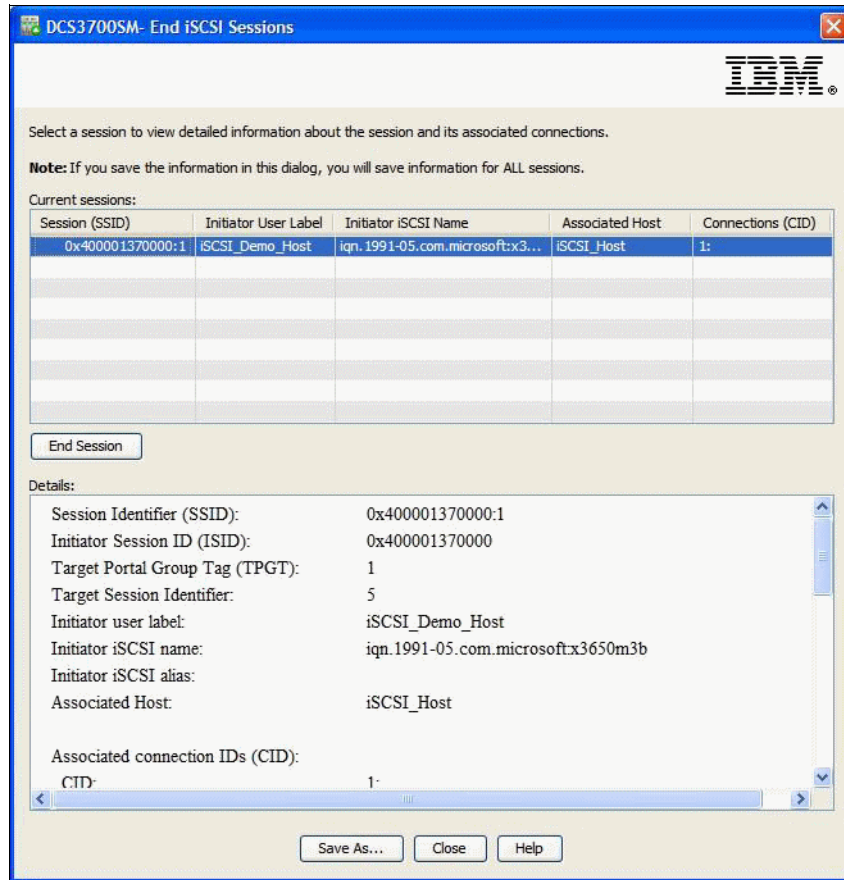


Figure 12-17 End iSCSI Sessions

- 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.
- In the bottom pane, click **Save As** to copy the session details into a text file.

Example 12-1 shows the complete session information for the selected initiator in Figure 12-17. This output can help verify that the host setup was done correctly.

*Example 12-1 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.1991-05.com.microsoftx3650m3b
  Initiator iSCSI label: port_iscsi_
  Initiator iSCSI alias:
  Host: iSCSI_Host

```

```

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: DCS3700-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.1991-05.com.microsoftx3650m3b

5. Click **Close**.

## 12.6 Viewing iSCSI statistics

To get information about transfers in OSI layers 2 - 5, click **Monitor** → **Health** → **iSCSI Statistics** (see Figure 12-18).

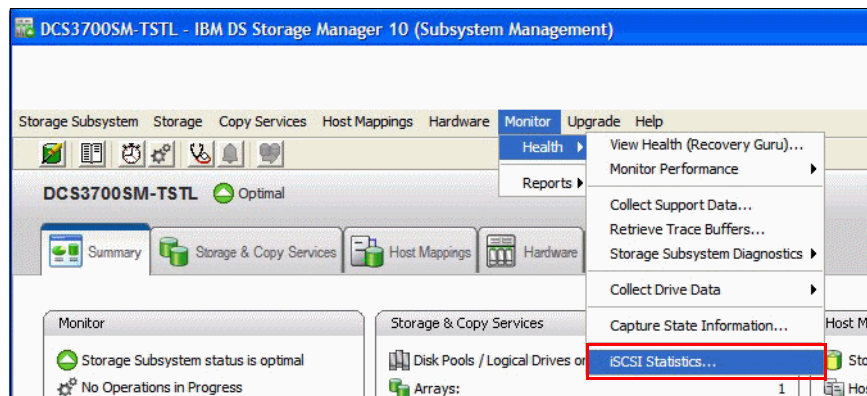


Figure 12-18 iSCSI statistics

Figure 12-19 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 uses a time that was defined by the user by using **Set Baseline**. This button resets all base line counters to zero and starts counting from the beginning.

Click **Save As** to save all the current statistics in a \*.csv file. Raw and baseline mode statistics for all levels are saved.

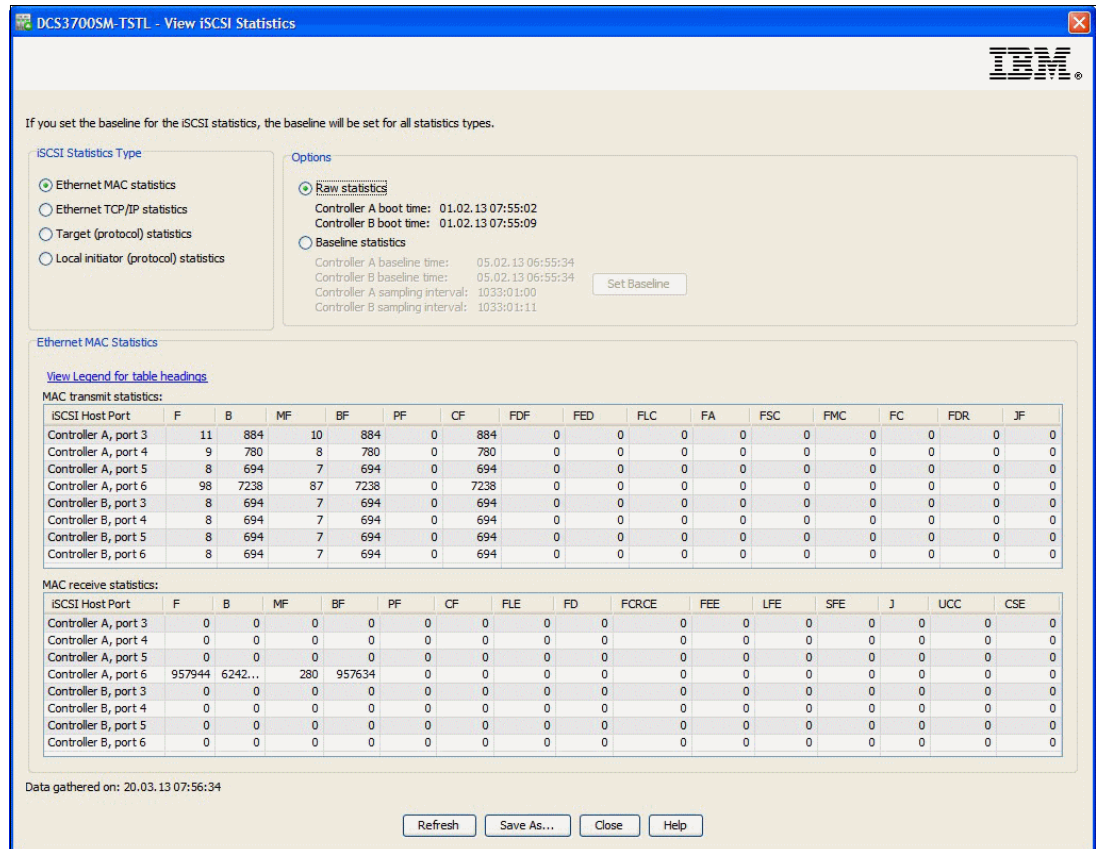


Figure 12-19 View iSCSI statistics for TCP/IP

Click **Help** to open the Online Help window, where a description is provided to translate column names into more meaningful names for each statistic type (see Figure 12-20 on page 381).

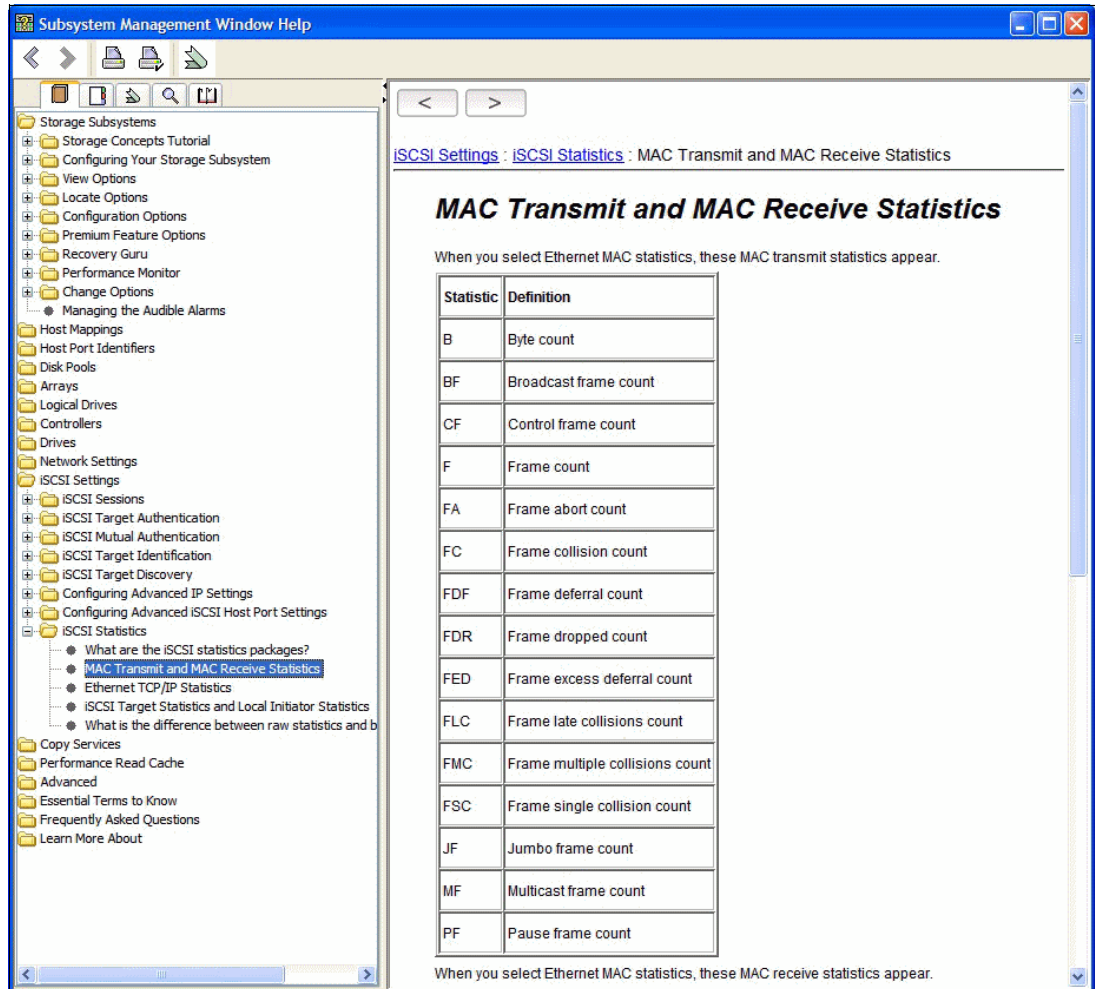


Figure 12-20 iSCSI help

## 12.7 Defining iSCSI hosts

Defining host using an iSCSI initiator is slightly different from defining a host with FC or SAS HBAs. In the case of iSCSI hardware initiators, you can see as many iSCSI host port identifiers as iSCSI HBAs that are installed in that host. In the case of an iSCSI software initiator, you see just one iSCSI host port identifier, even if you have more Ethernet cards that are used for iSCSI in your host server. Multiple redundant data paths are presented as multiple iSCSI sessions in this case, as described in 12.5, “View/End iSCSI Sessions” on page 377.

Defining a host with traditional SAS or FC HBAs is described in 10.2, “Defining a host” on page 330. The following sections show iSCSI host definition differences.

### 12.7.1 Viewing unassociated iSCSI initiators

A host definition is provided in the Mapping tab of DS Storage Manager. To see whether servers with an active iSCSI initiator are available to be defined as a host, view the unassociated host port identifiers.

There are several methods to show unassociated host ports. For example, you can click **Mappings** → **View Unassociated Host Port Identifiers** in the Subsystem Management drop-down menu. A window opens, as shown in Figure 12-21. You can see WWPNs for SAS or FC HBAs, and if an iSCSI initiator is configured and authenticated for access to your DCS3700 storage subsystem, you see the IQN name of this iSCSI Initiator. An iSCSI software initiator in Microsoft Windows is shown in Figure 12-21. 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.

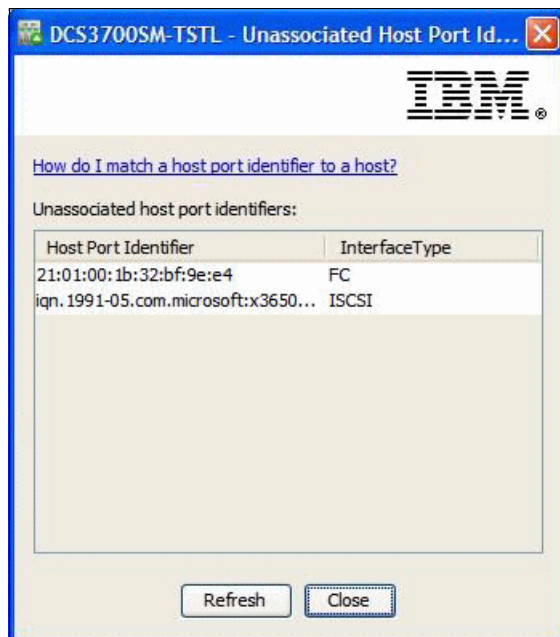


Figure 12-21 Unassociated iSCSI software initiator

**Note:** You can see the iSCSI initiator IQN name of your server in this window only if you already have discovered and authenticated your DCS3700 storage subsystem as an iSCSI target on your host server.

## 12.7.2 Defining a new iSCSI host

You can define a new iSCSI host by using the same procedure that is described in 10.2, “Defining a host” on page 330. The only differences are the setting in the Specify Host Port Identifier window, as shown in Figure 12-22 on page 383:

1. Select **ISCSI** as the host interface type, and if you select the **Known unassociated host port identifiers** drop-down menu, you can find your iSCSI Initiator IQN name.
2. Enter a User Label (alias), and click **Add** to add this iSCSI Initiator as a host port to your newly defined host. For an iSCSI software initiator, you add only one IQN name per host.
3. Click **Next** to continue and finalize the host definition.

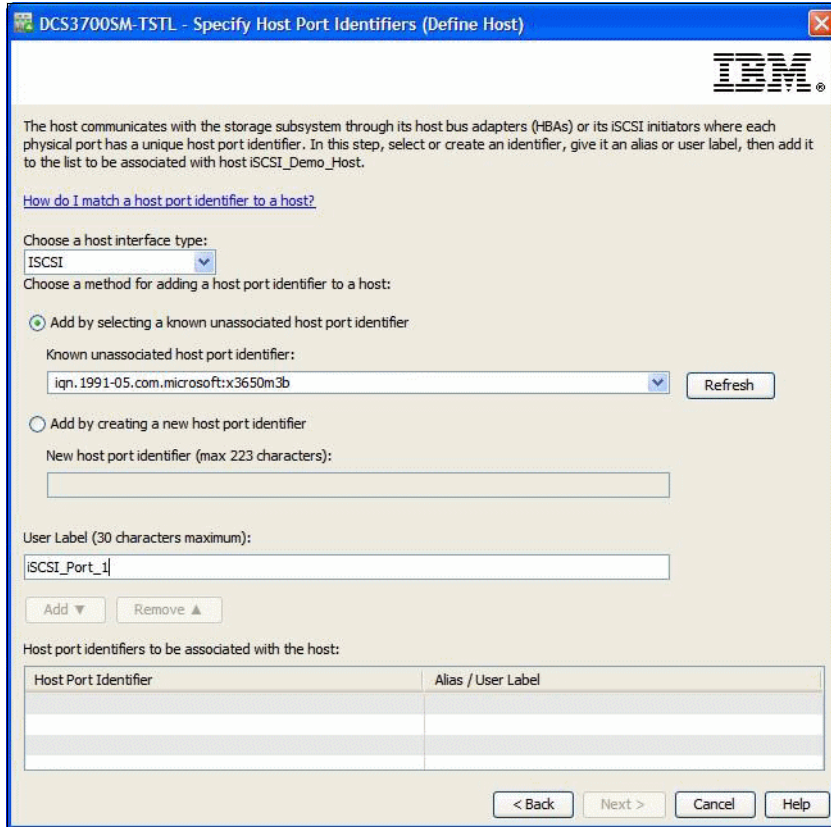


Figure 12-22 Define an iSCSI host - host port identifier specification

The mapping of logical drives to an iSCSI host is the same as for every other host. It does not depend on the type of host attachment interface. This procedure is described in Chapter 10, “Administration: Mappings tab” on page 327.

### 12.7.3 Managing iSCSI host ports

You can check or modify your current definition of the iSCSI hosts in the Manage Host Port Identifier menu of the DS Storage Manager. To do so, complete the following steps:

1. Click **Mappings** → **Manage Host Port Identifiers**. A window opens (Figure 12-23).
2. Select your host under the Show host port identifiers associated with: field. You can see details about your iSCSI host port identifier in the Host port identifier information field.

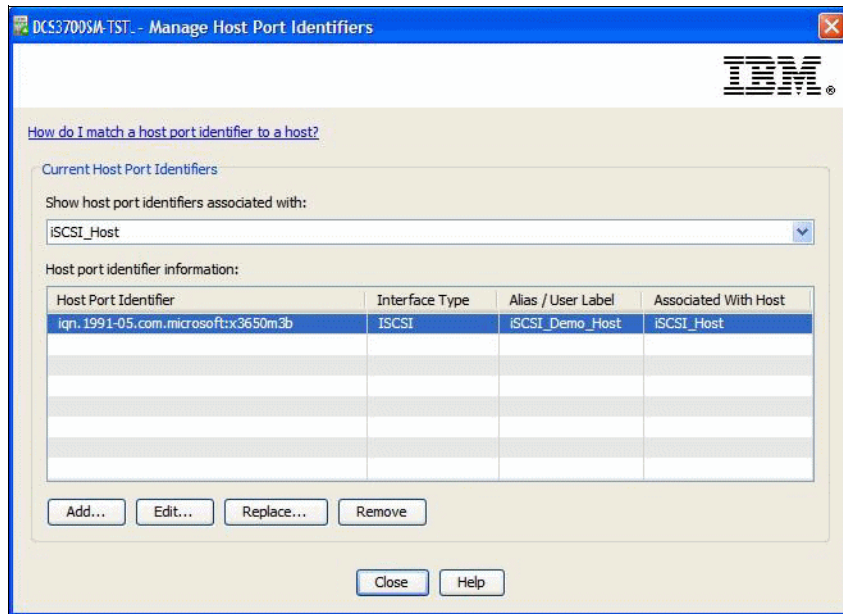


Figure 12-23 iSCSI host port information

3. Theoretically, you can add or replace the iSCSI port identifier, but practically, if you have an iSCSI software initiator that is tightly coupled with your server operating system, these options cannot be used.



## Administration: Monitor

This chapter describes the monitor options of the IBM System Storage DCS3700 storage subsystem. These options help with troubleshooting and reporting tasks. These options are accessed through the DS Storage Manager (Subsystem Management) Monitor menu.



## 13.1 The Subsystem Management Monitor menu

This section introduces the monitor functions that are accessed from the DS Storage Manager (Subsystem Management) Monitor menu that is shown in Figure 13-1.

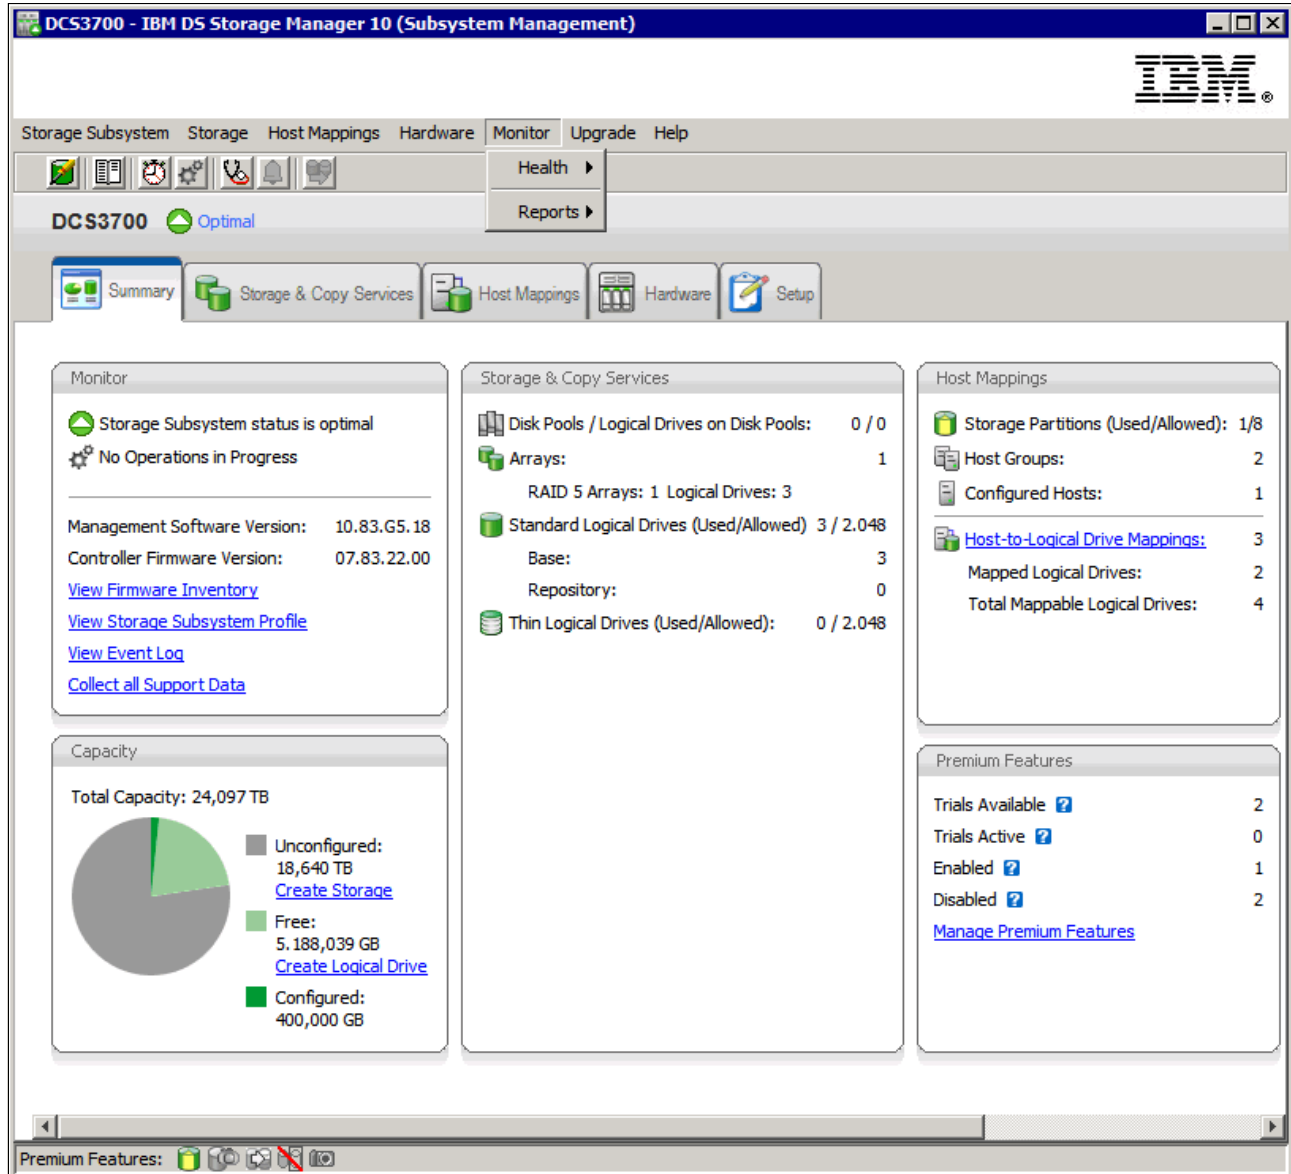


Figure 13-1 The DS Storage Manager (Subsystem Management) Monitor menu

The Monitor menu provides access to two sets of tools:

- ▶ Health
- ▶ Reports



## 13.2 View Health (Recovery Guru)

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 flashing 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 Needs Attention link appears in the DS Storage Manager (Subsystem Management), as shown in Figure 13-2.

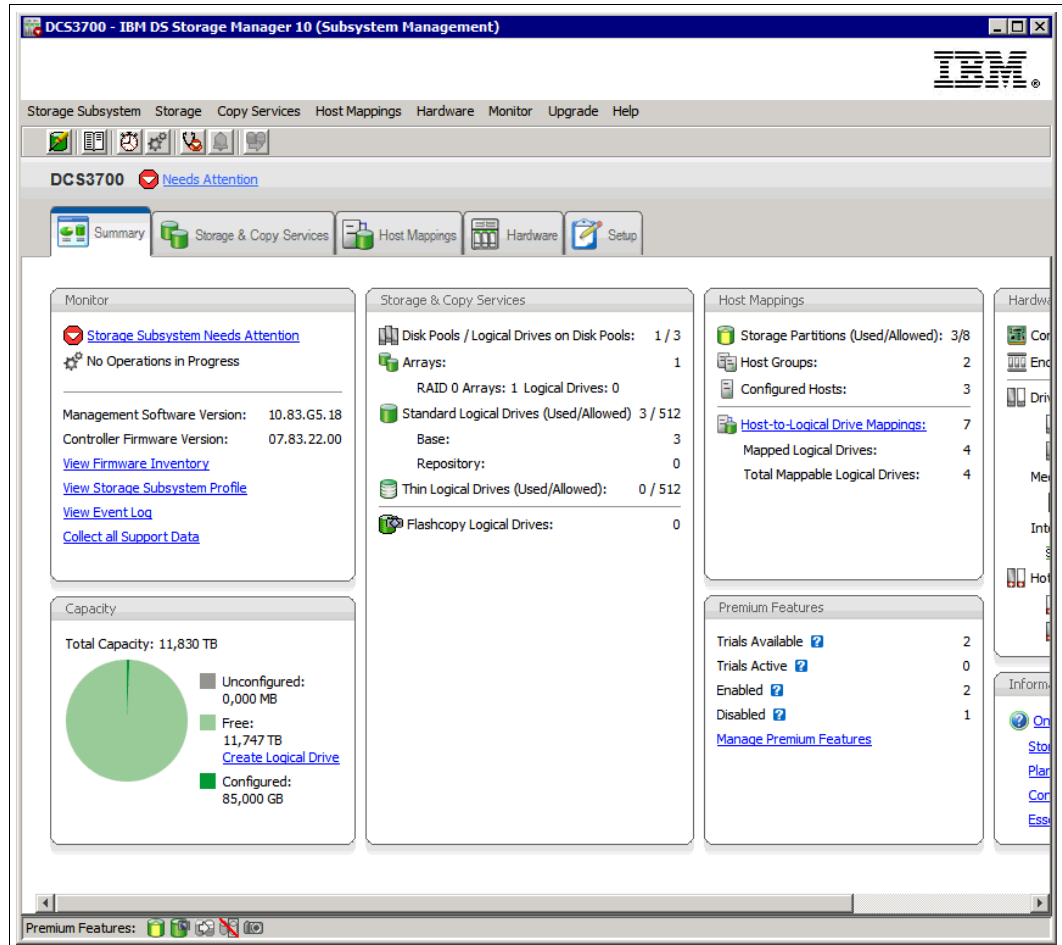


Figure 13-2 Summary tab - Storage subsystem needs attention

Click the **Needs Attention** link to start the Recovery Guru, as shown in Figure 13-3.

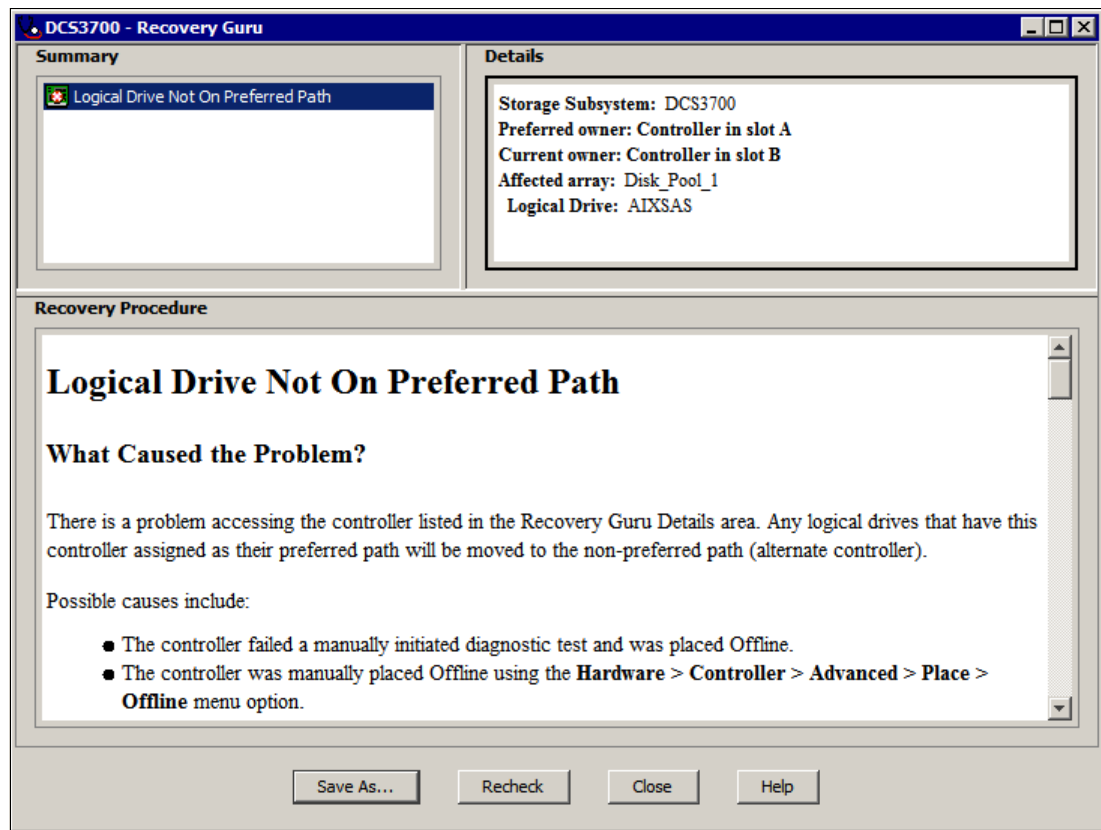


Figure 13-3 Recovery Guru

The Recovery Guru window has three main panes:

- ▶ Summary pane: All detected problems are listed here. There are two separate problems that are listed in the example that is shown in Figure 13-3.
- ▶ Details pane: Lists specific details for the problem that is selected in the Summary section. If there is more than one problem that is 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 complete the recovery steps, use this button to recheck the system.
- ▶ Close: Close the Recovery Guru.
- ▶ Help: Open the Recovery Guru Help.

## 13.3 Performance Monitor

This section presents the DS Storage Manager Performance Monitor software. It provides visibility into performance activity across monitored storage devices. Performance Monitor is used to perform these tasks:

- ▶ View in real time the values of the data that is collected for a monitored device. This capability helps you determine whether the device is experiencing any problems.
- ▶ See a historical view of a monitored device to identify when a problem started or what caused a problem.
- ▶ Specify various reporting attributes, such as time increments and filtering criteria, to examine performance trends and to pinpoint the cause of availability and performance issues.
- ▶ Display data in tabular format (actual values of the collected metrics) or graphical format (primarily as line-graphs), or export the data to a file.

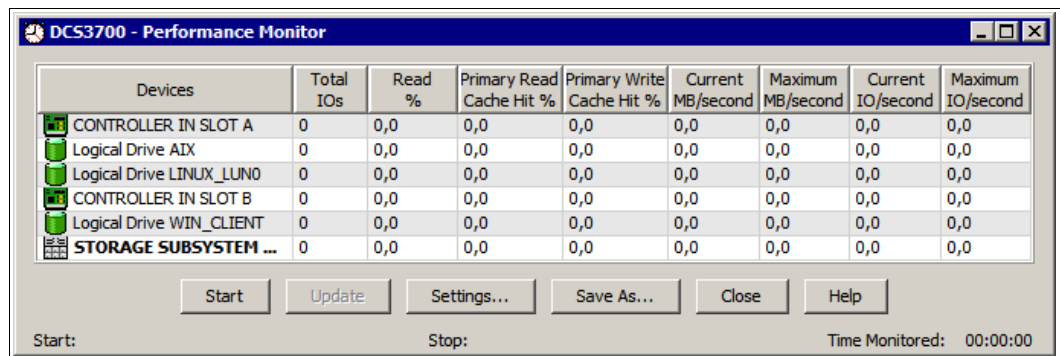
The Performance Monitor software allows two different approaches to display counters:

- ▶ Table View
- ▶ Graphical View

### 13.3.1 Performance Monitor: Table View

This topic describes how to view the performance data in tabular format:

1. Open the Monitor Performance Table View by clicking **Monitor** → **Health** → **Monitor Performance** → **Table View...** 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 13-4.



The screenshot shows a window titled "DCS3700 - Performance Monitor". It contains a table with the following data:

Devices	Total IOs	Read %	Primary Read Cache Hit %	Primary Write Cache Hit %	Current MB/second	Maximum MB/second	Current IO/second	Maximum IO/second
CONTROLLER IN SLOT A	0	0,0	0,0	0,0	0,0	0,0	0,0	0,0
Logical Drive AIX	0	0,0	0,0	0,0	0,0	0,0	0,0	0,0
Logical Drive LINUX_LUN0	0	0,0	0,0	0,0	0,0	0,0	0,0	0,0
CONTROLLER IN SLOT B	0	0,0	0,0	0,0	0,0	0,0	0,0	0,0
Logical Drive WIN_CLIENT	0	0,0	0,0	0,0	0,0	0,0	0,0	0,0
STORAGE SUBSYSTEM ...	0	0,0	0,0	0,0	0,0	0,0	0,0	0,0

Below the table are buttons for Start, Update, Settings..., Save As..., Close, and Help. At the bottom, there are fields for Start, Stop, and Time Monitored: 00:00:00.

Figure 13-4 Performance Monitor

2. Click **Settings** to configure the Performance Monitor (Figure 13-5 on page 391). In this window, you select the specific logical drives and controllers that you want to collect performance counters for. Click **Select All** 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 select the subsystem logical drives and controllers that you want to monitor.

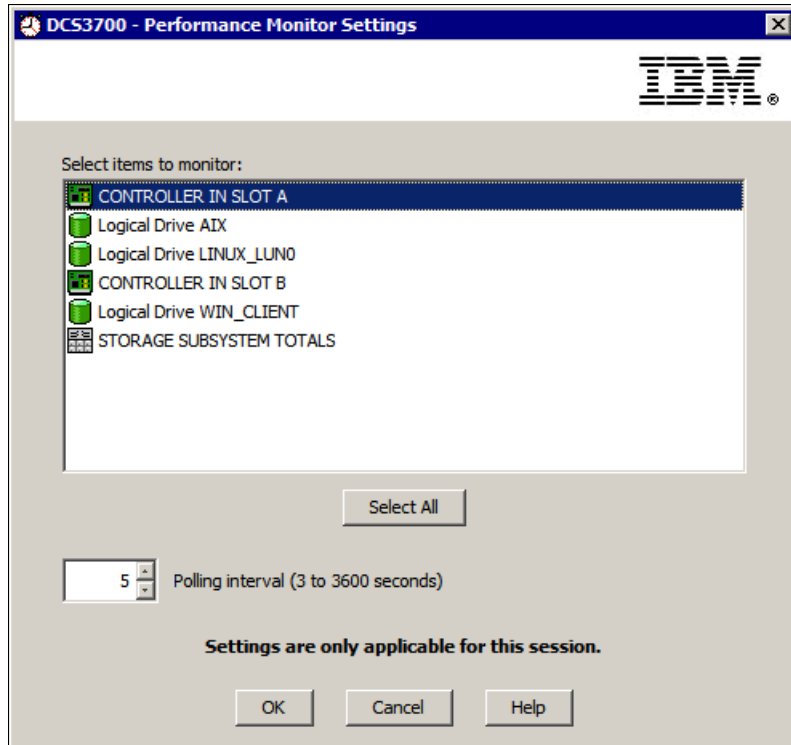


Figure 13-5 Performance Monitor Settings

3. Click **Start** to begin collecting the performance counters (Figure 13-6). 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 Figure 13-5. Click **Update** to force an immediate update of the performance counters.

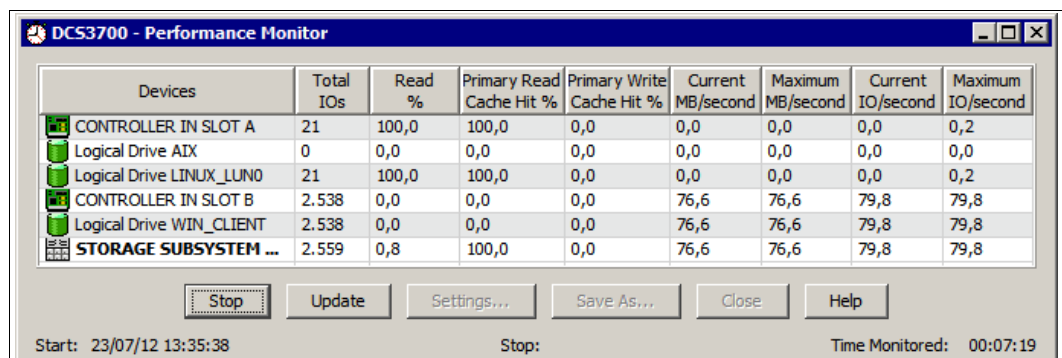


Figure 13-6 Performance Monitor collecting data

4. When you collect the performance data, click **Save As** 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.

## 13.3.2 Graphical Performance Monitor

This next topic describes how to view the performance data in graphical format. The data is presented as a graph displaying the value of a selected metric over a specified time period for a selected device. Open the Graphical Performance Monitor (as shown in Figure 13-7) by clicking **Monitor** → **Health** → **Monitor Performance** → **Graphical View...**

The View Graph window provides two basic graph contexts, as shown in Figure 13-7:

- ▶ **Real-time:** The real-time context plots the sampled data points in real time, so you see instantaneous changes in the tracked metrics as they occur. The data points on the graph refresh every 5 seconds. The system drops the oldest data point from the graph and adds the newest data point to the graph.
- ▶ **Historical:** The historical context plots sampled data points that are collected in the past and stored in the history buffer. The maximum length of a history data window is 7 days. For historical data graphs, the data points are static. After a historic monitoring session is started, the system polls the storage subsystem every 10 minutes.

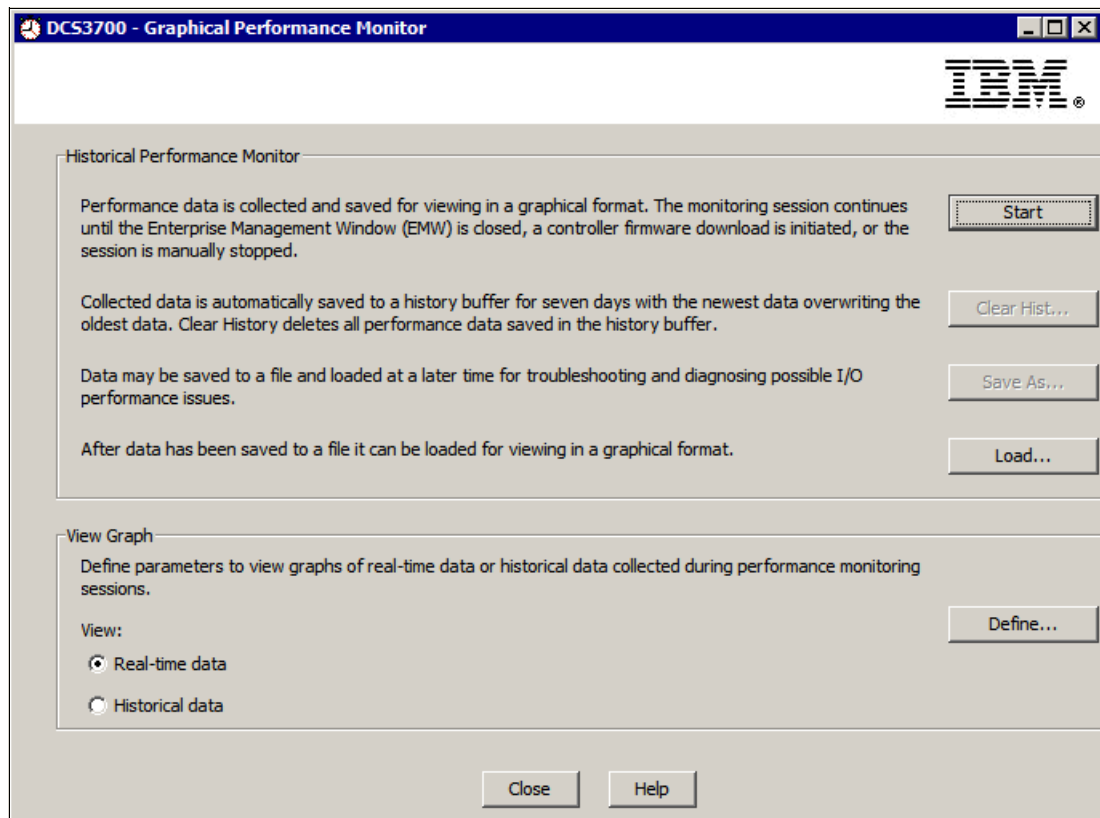


Figure 13-7 Graphical Performance Monitor options

**Note:** For historical data, you can select up to four devices and one metric to view per graph. For real-time data, you can select only one metric and one device to view per graph.

## 13.4 Collect Support Data

This wizard is an automated way to collect data for remote troubleshooting and analysis of problems with the storage management software and hardware. All gathered files are compressed into a single archive that is sometimes referred to as the CASD (Collect All Support Data) log. 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 occurs 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 **Collect all Support Data** link in the main Summary tab (Figure 13-1 on page 386) to open the Collect All Support Data window (Figure 13-8).

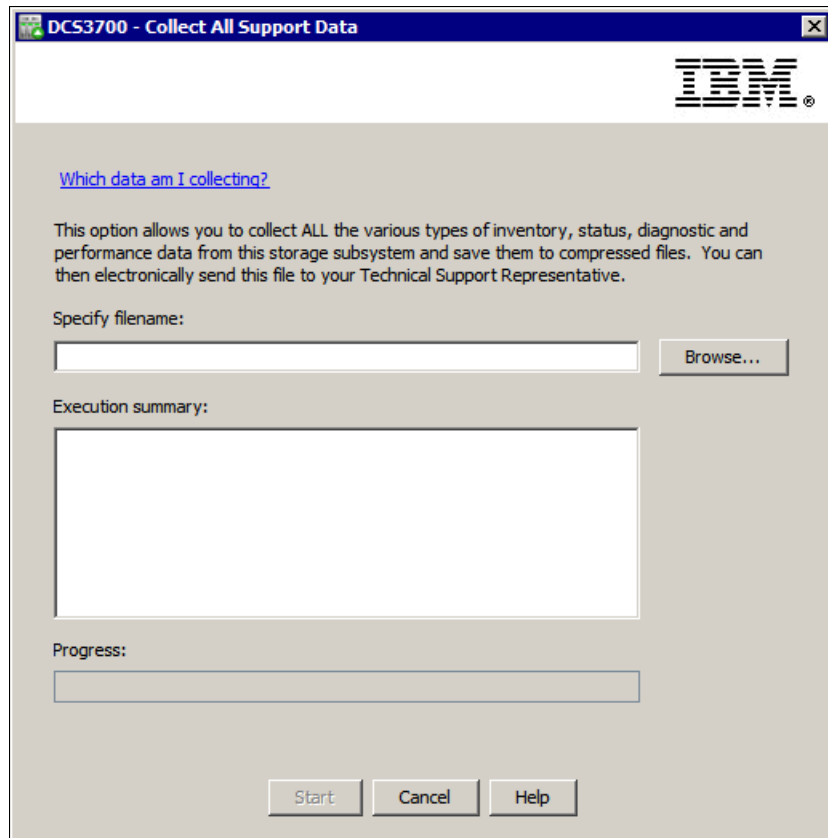


Figure 13-8 Collect All Support Data window

- ▶ Click **Monitor** → **Health** → **Collect Support Data**, as shown in Figure 13-9, to open the Collect All Support Data window.

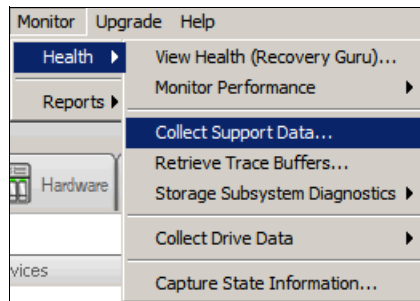


Figure 13-9 Open the Collect All Support Data window from the Monitor menu

The Collect All Support Data window allows you to specify the file name and the directory location where the DCS3700 support data is 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 that are collected by the wizard.

### 13.4.1 Saving the support data

To save the support data, complete the following steps:

1. Open the Collect All Support Data window and click **Browse** to specify the file name and folder location where the wizard saves the collected support data, as shown in Figure 13-10 on page 395.



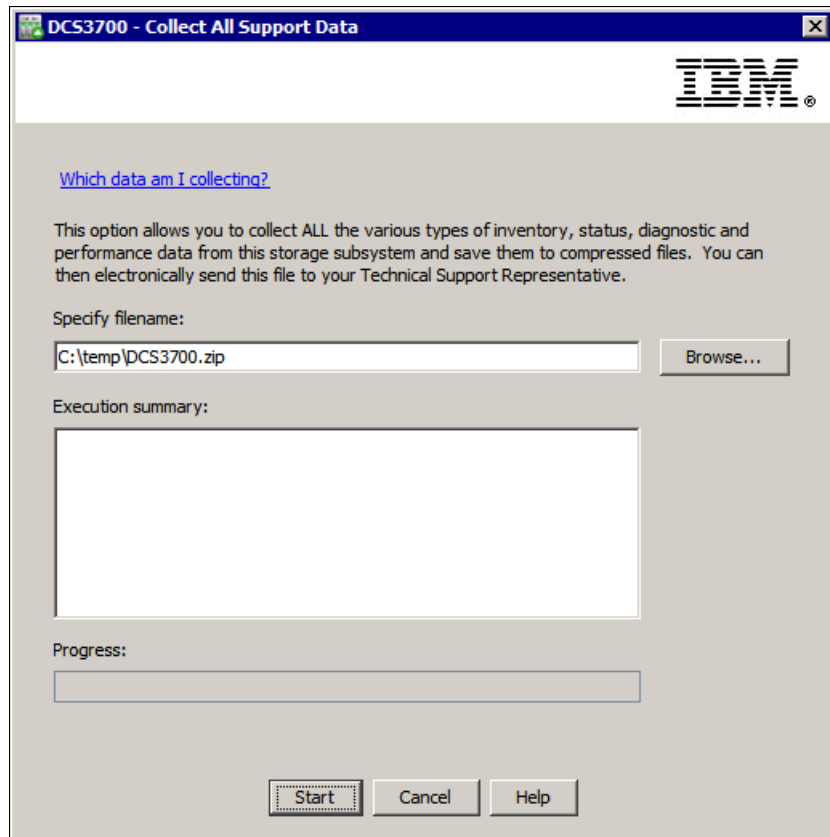


Figure 13-10 Collect All Support Data - specify the file name and folder location

2. Click **Start** to begin the data collection, as shown in Figure 13-10.

3. The Collect All Support Data window displays a progress bar while the data is being collected, as shown in Figure 13-11.

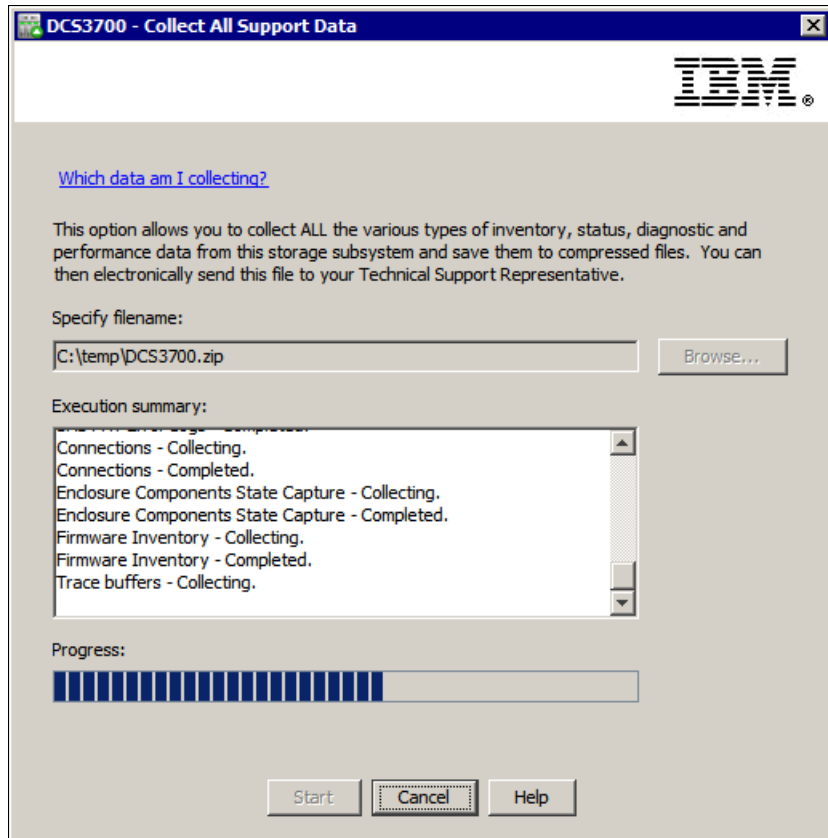


Figure 13-11 Collect All Support Data - data collection in progress

- The window displays a message indicating when the support data is collected and saved to the file that you specified in step 1 on page 394, as shown in Figure 13-12.

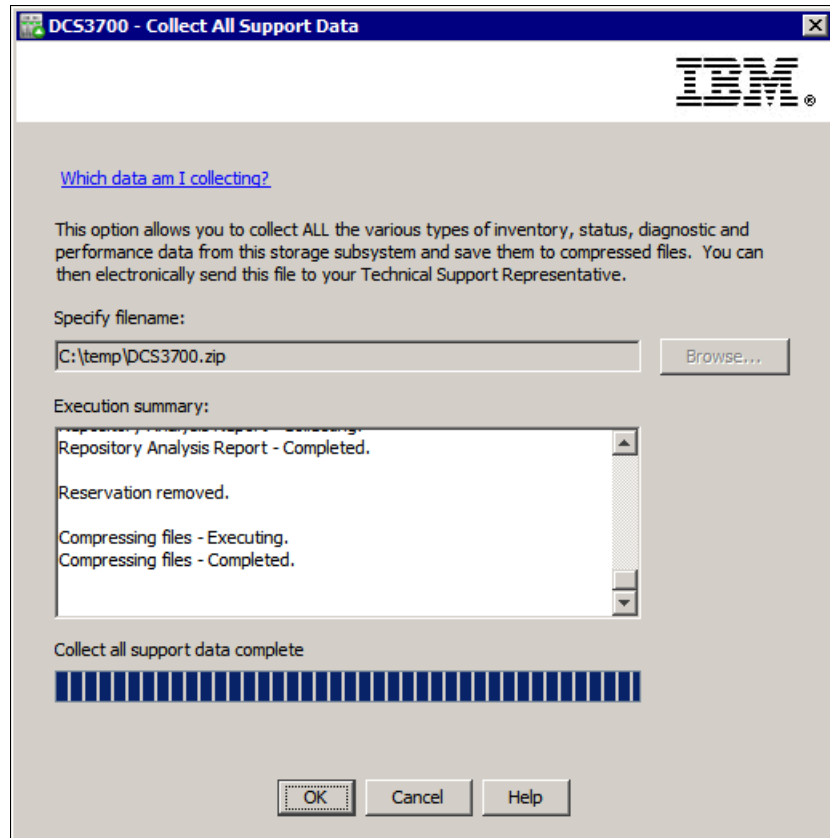


Figure 13-12 Collect All Support Data - save completed

- Send the compressed file containing the collected support data to your IBM Customer Support representative for troubleshooting and further analysis.

### 13.4.2 Automatic Support Data Collection settings

DS Storage Manager can enable automatic collection of the support data. When this option is enabled and a critical event occurs, the support data file is collected and automatically transferred to the directory that is specified and all the information that is relevant for troubleshooting by your support representative is preserved. If consecutive critical events occur, only the information that is related to the last event is stored in the auto save support bundle.

**Tip:** Enable the Automatic Support Data 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 do the following actions:

- Specify a directory outside your DCS3700 storage subsystem 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 available through the Storage Manager (Enterprise Management) window. Right-click the correct storage subsystem and select **Collect Support Data** → **Automatically...** (Figure 13-13) to open the Automatic Support Data Collection window.

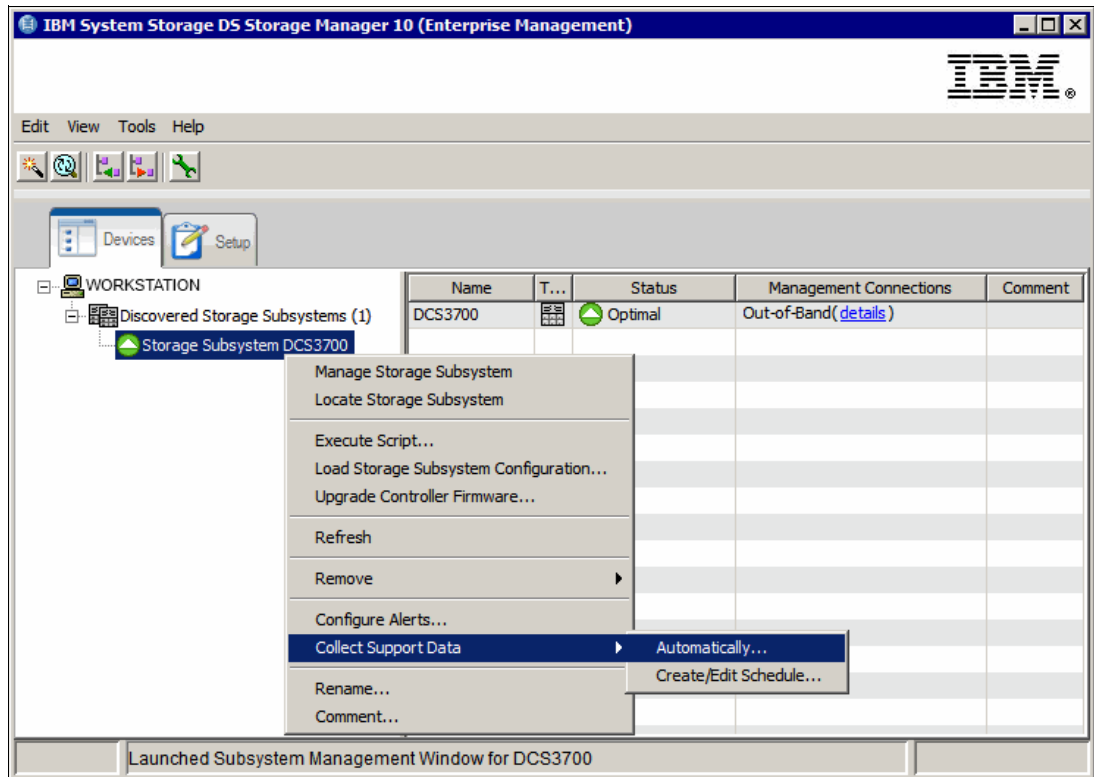


Figure 13-13 Automatic Support Data Collection function from the Collect Support Data menu

In the Folder Location for ALL Storage Subsystems pane, the **Use default location** option is selected (Figure 13-14 on page 399) Optionally, you can click **Use alternate location** to specify the destination directory for the collected data. The value of the Filename pattern: field cannot be changed.

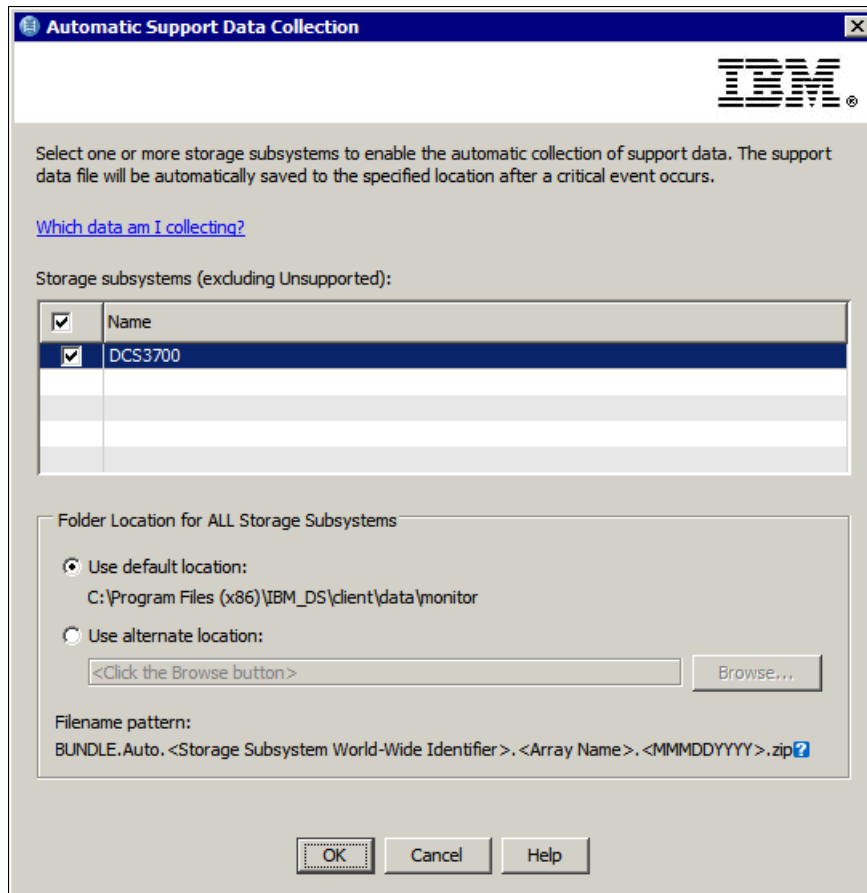


Figure 13-14 The Automatic Support Data Collection window

## 13.5 Retrieve Trace Buffers

Use the Retrieve Trace Buffers option to save trace information to a compressed file. The firmware uses the trace buffers to record processing, including exception conditions, that might be useful for debugging. Trace information is stored in the current buffer. You can move the trace information to the flushed buffer after you retrieve the information. (The option to move the trace information to the flushed buffer is not available if you select **Flushed buffer** from the Trace buffers list.) Because each controller has its own buffer, there might be more than one flushed buffer. You can retrieve trace buffers without interrupting the operation of the storage subsystem and with minimal effect on performance.

To open this window, click **Monitor** → **Health** → **Retrieve Trace Buffers...** as shown in Figure 13-15.

**Note:** Use this option only under the guidance of your IBM Technical Support representative.

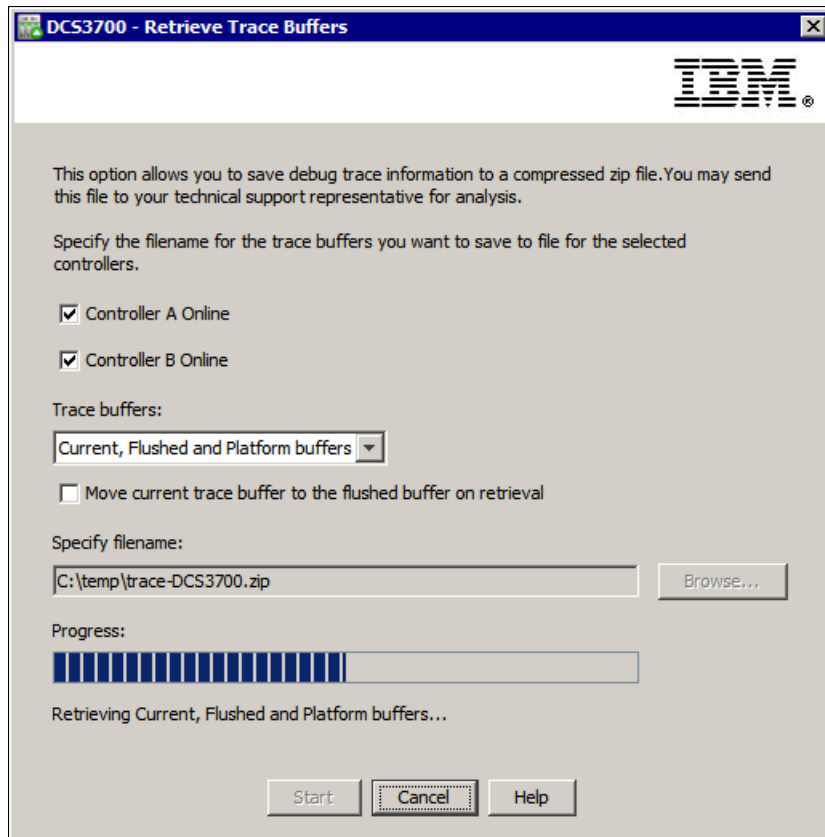


Figure 13-15 Retrieve Trace Buffers in progress

A compressed archive file is stored at the location that you specify. The archive contains trace files from one or both of the controllers in the storage subsystem along with a descriptor file named `trace_description.xml`. Each trace file includes a header that identifies the file format to the analysis software that is used by your IBM Technical Support representative.

## 13.6 Storage subsystem diagnostic tests

This section describes the functions that are used for specific health diagnostic tests of DCS3700 storage subsystem.

**Attention:** Perform storage subsystem diagnostic operations only under the guidance of your IBM Technical Support representative.

### 13.6.1 Read Link Status

The Read Link Status error refers to link errors that have been detected in the traffic flow of a Fibre Channel loop. The errors that are detected are represented as a count (32-bit field) of error occurrences that have accumulated over time. The errors provide a coarse measurement of the integrity of the components and storage subsystems on the loop. The Read Link Status dialog retrieves the error counts and shows the controllers, drives, environmental services modules (ESMs), and Fibre Channel ports in channel order.

**Note:** For further troubleshooting assistance, save the Read Link Status results and forward them to your IBM Technical Support representative for assistance.

### 13.6.2 Collecting I/O Path Statistics

Application performance issues can be caused by memory usage, processor usage, network latency, I/O latency, or other issues. You can find and solve these performance issues by analyzing the input/output (I/O) path statistics. The I/O path statistics are collected automatically during support data collection or you can collect them manually by following the procedure that documented in this section. The data from this procedure is saved in a binary format that can be analyzed only by an IBM Technical Support representative.

**Attention:** Use this option only under the guidance of your IBM Customer Support representative.

To start this procedure, click **Monitor** → **Health** → **Storage Subsystem Diagnostic** → **Collect I/O Path Statistics...**, which opens a window where a file name must be provided (Figure 13-16).

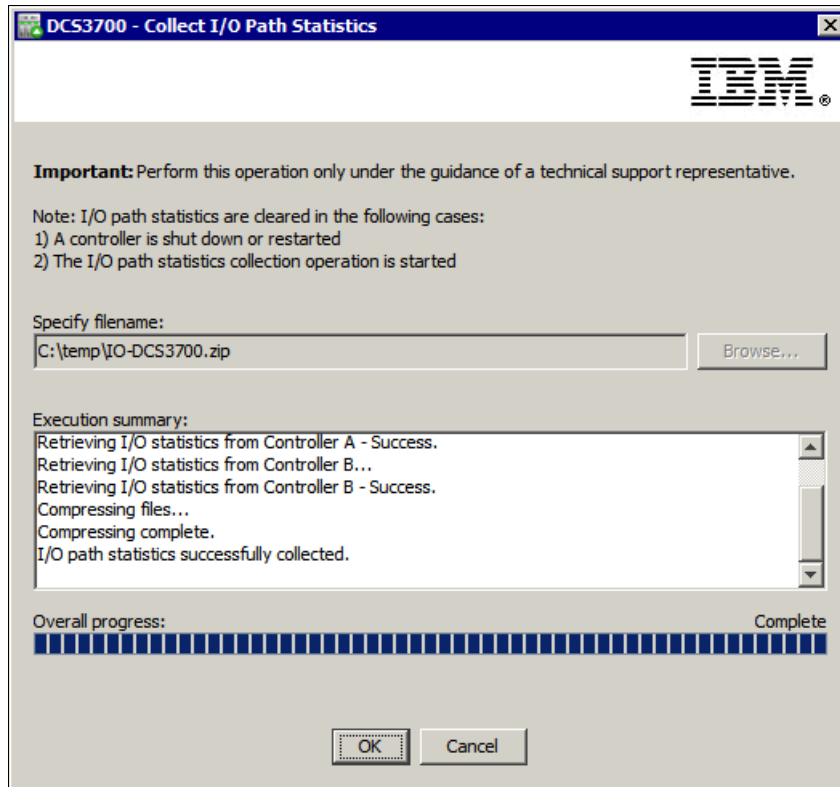


Figure 13-16 I/O Path Statistics collected

**Note:** After you click **OK** in Figure 13-16, the counters for the I/O path statistics are reset even if you later cancel the operation at the confirmation window that is shown in Figure 13-17.

The counters are also reset when the controller begins start of day processing.

After you enter the file name, click **OK**. A confirmation window opens (Figure 13-17 on page 403). If you want to proceed, enter yes in the field and click **OK**.



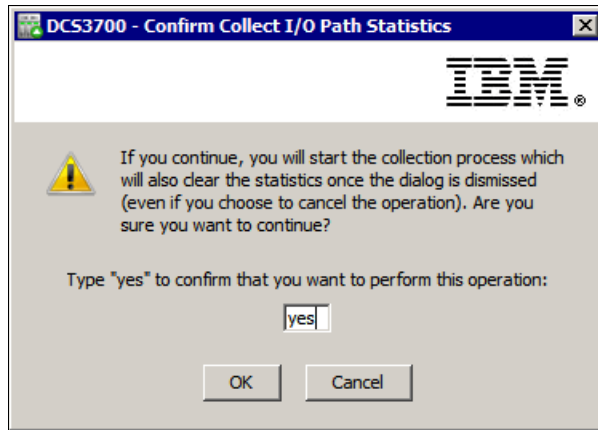


Figure 13-17 Confirmation window before you start Collecting I/O Path Statistics

### 13.6.3 Validating the configuration database

The configuration database stores information that is used by the storage subsystem to manage the controller firmware. Use the Run Configuration Database Diagnostics option to check that the configuration database is consistent.

**Attention:** Use this option only under the guidance of your IBM Technical Support representative.

## 13.7 Collecting drive data

You can save diagnostic data from all of the physical drives that are installed in the DCS3700 storage subsystem. This drive diagnostic data is only readable by your IBM Customer Support representative, and this option is normally used at their specific request.

**Note:** This log is saved as part of the Collect All Support Data log (for more information, see 6.2.7, “Collect Support Data” on page 194).

Each drive in the DCS3700 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 problems that might exist.

To collect the drive data, complete the following steps:

1. In the DS Storage Manager (Subsystem Management) window, click **Monitor** → **Health** → **Collect Drive Data** → **All Drives...** to open the Collect Drive Data window, as shown in Figure 13-18.

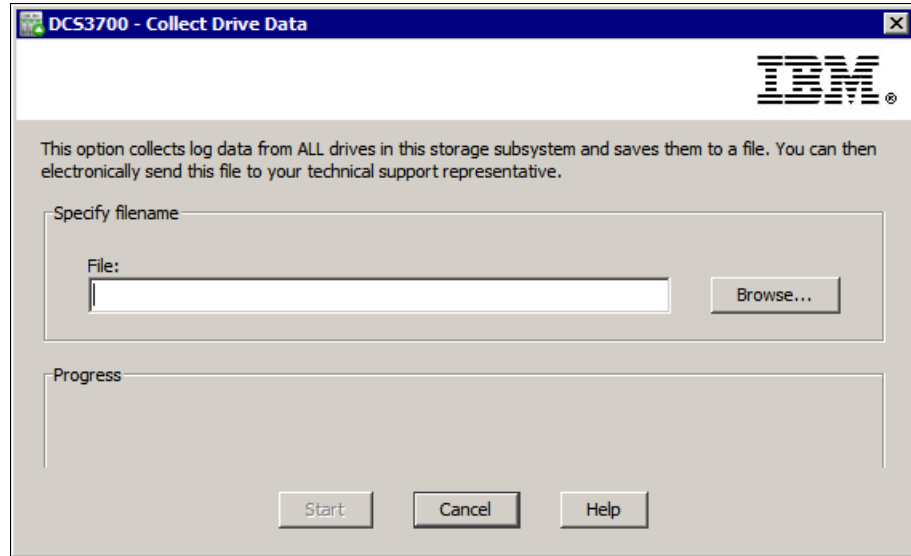


Figure 13-18 Collect Drive Data

2. Click **Browse** to specify the location and file name of the compressed file where you save the drive diagnostic data. Then, click **Start** to save the file. Figure 13-19 shows the window after the drive diagnostic data is saved.

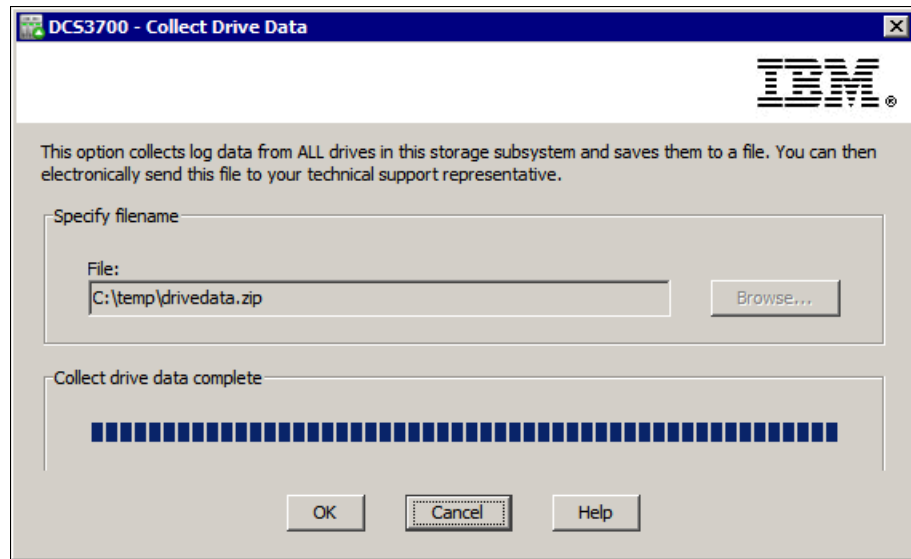


Figure 13-19 Collect drive data complete

## 13.8 Collecting state capture

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 (for more information, see 6.2.7, “Collect Support Data” on page 194).

**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 **Monitor** → **Health** → **Capture State Information...** to open the State Capture window, which shown in Figure 13-20.

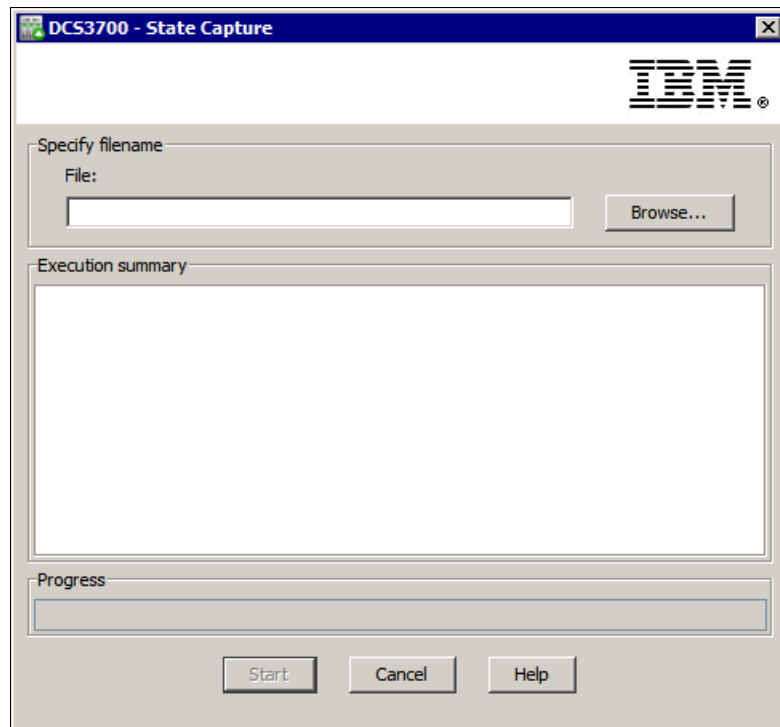


Figure 13-20 State Capture

2. Click **Browse** to specify the location and file name of the state capture information file (the default file extension is .dmp). Then, click **Start** to begin saving the file. You are prompted to confirm that you want to save the state capture data (Figure 13-21). Enter yes in the box to confirm and then click **OK** to continue.

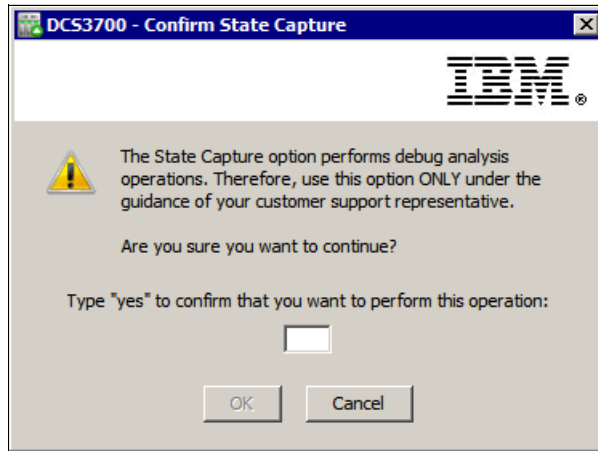


Figure 13-21 Confirm State Capture

3. Figure 13-22 shows that the State Capture window after the state capture information is saved to a file.

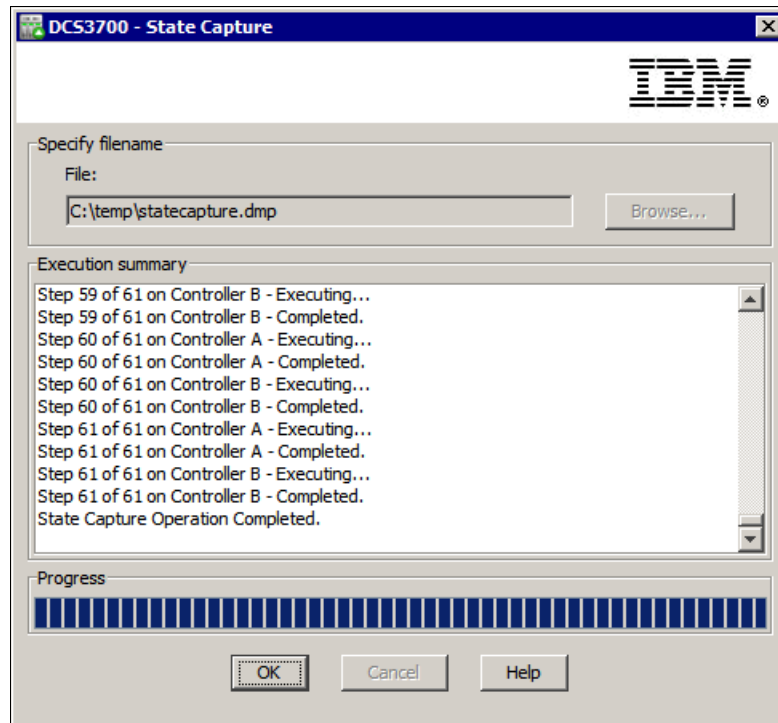


Figure 13-22 State capture completed

## 13.9 Operations in Progress

This tool can show all the internal operations that are being processed by the DCS3700 storage subsystem, such as a RAID level change or a disk being rebuilt. By clicking **Monitor** → **Reports** → **Operations in Progress**, you can see a table with information about the status of all operations (Figure 13-23).

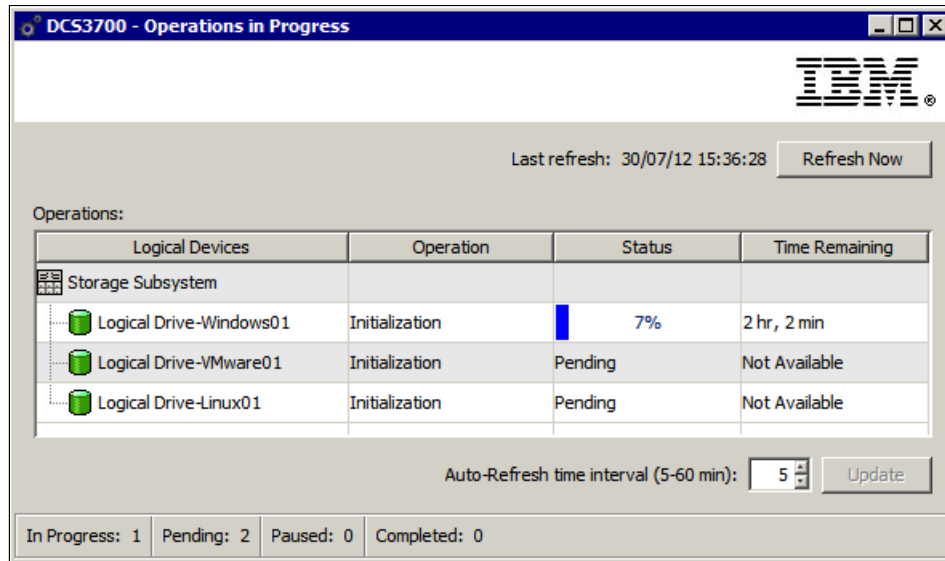


Figure 13-23 Operations in Progress window

## 13.10 Storage Subsystem Profile

Section 6.2.5, “View Storage Subsystem Profile” on page 192 introduced the Storage Subsystem Profile. 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 the following items:

- ▶ Storage subsystem
- ▶ Storage
  - Disk pools
  - Arrays
  - Logical drives
  - Missing logical drives
  - Logical hierarchy
- ▶ Copy services
  - FlashCopy
  - Repositories
- ▶ Host mappings
- ▶ Hardware
  - Controllers
  - Drives
  - Drive channels
  - Enclosures

Here you can see all of the information about the storage subsystem at once, or filter information to see a specified component. You can view the profile in the window or save it as a text file.

To view the DCS3700 storage subsystem profile, complete the following steps:

1. Click **View Storage Subsystem Profile** in the Summary tab under Monitor Group (Figure 13-1 on page 386). The Storage Subsystem Profile window opens, as shown in Figure 13-24.

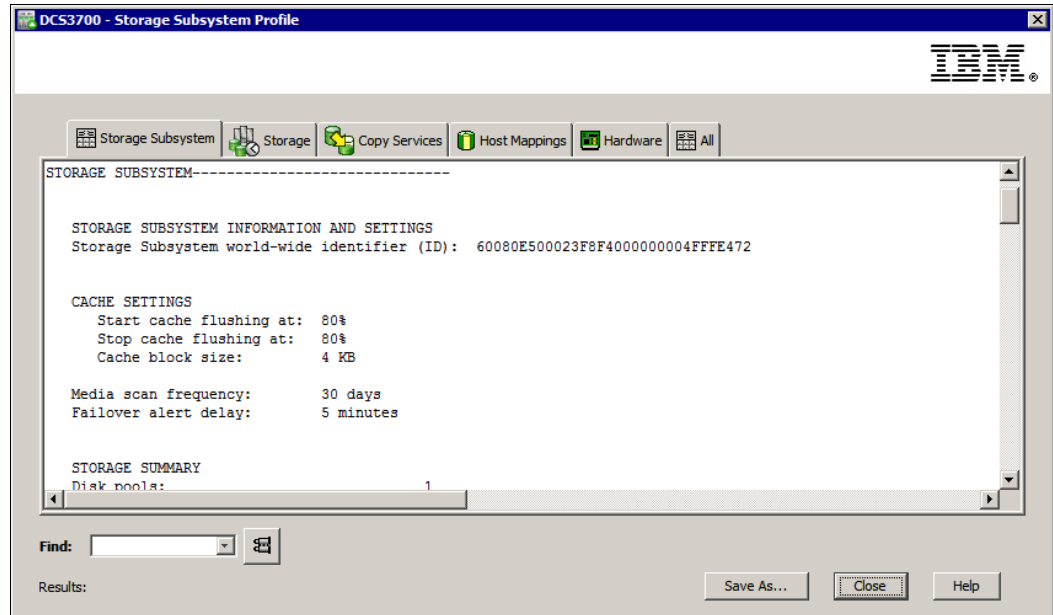


Figure 13-24 The default view when viewing the DCS3700 storage subsystem profile

2. Select the tab for the information you want to view. By default, you open the **Storage Subsystem** tab. Select a specific component or click the **All** tab to show everything, as shown in Figure 13-25.




Figure 13-25 The tab bar of the Profile window with the Storage Subsystem tab selected

Use the **Find** box (Figure 13-26) in the lower left corner of the profile window to search for the profile. Enter the required text and click the Binoculars icon to find the string in the current tab.




Figure 13-26 The Find box

3. To save the profile, click **Save As**. At the top of Figure 13-27 on page 409, choose which sections of the profile that you want to save. By default, the tab that you were viewing is 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 is stored as a text file.

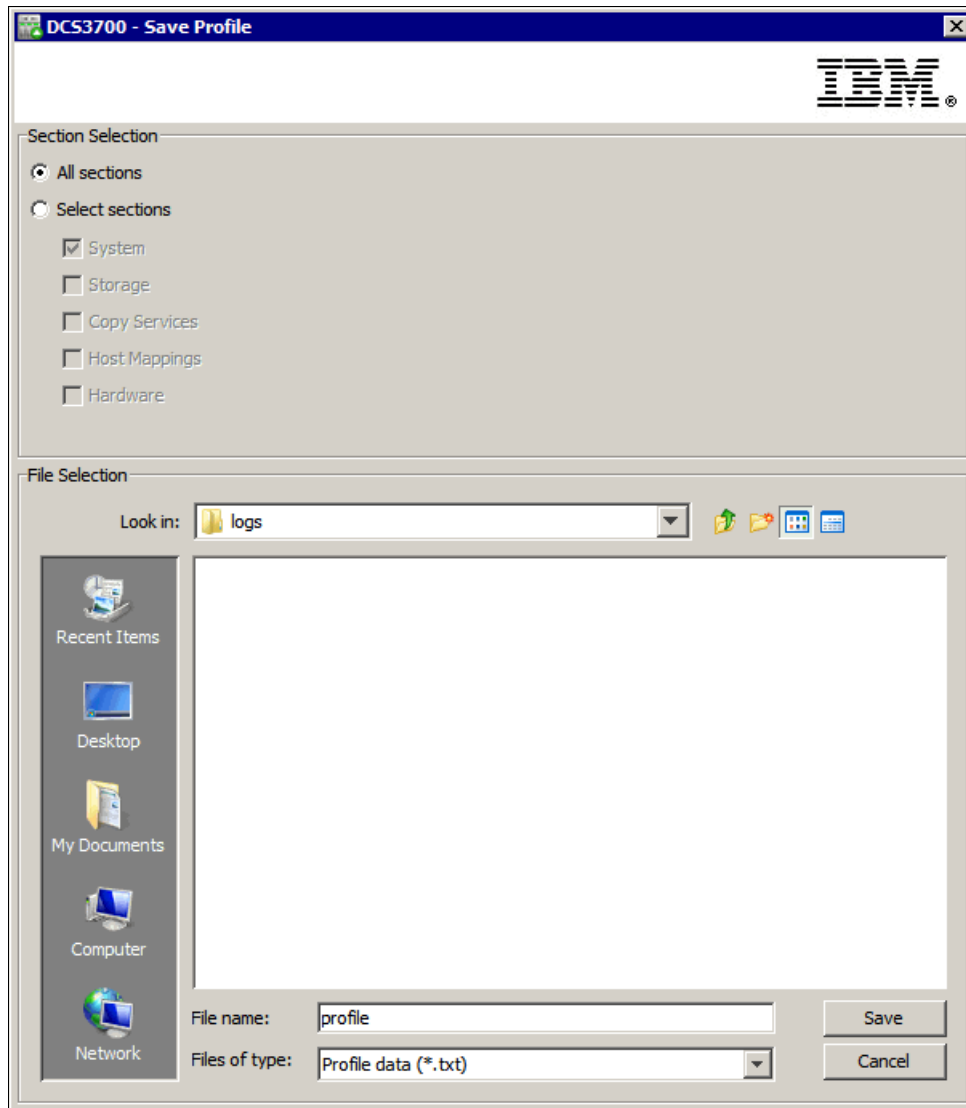


Figure 13-27 Save Profile

4. After the profile is saved, click **Close** (in Figure 13-24 on page 408) to return to the Support tab.

**Note:** More information is saved if the Collect All Support Data log is gathered (for more information, see 6.2.7, “Collect Support Data” on page 194).

## 13.11 Cable connection

This function is used only when an expansion unit is connected. If you receive an error message for a channel, you can use this function to identify the components on the channel that might have caused the error. By isolating these components, you prevent accidentally disconnecting components that are still in operation, which could cause an interruption in data flow.

The Connections dialog also provides information about drive channel trunking. If your configuration supports drive channel trunking, the Trunking Active column displays in the Connections dialog. The Trunking Active column indicates whether the trunk is active, if the trunk is not active, or if the trunk is not active because a miswire exists on a port. To open the Connections window, which is shown in Figure 13-28, click **Monitor** → **Reports** → **Cable Connections**.

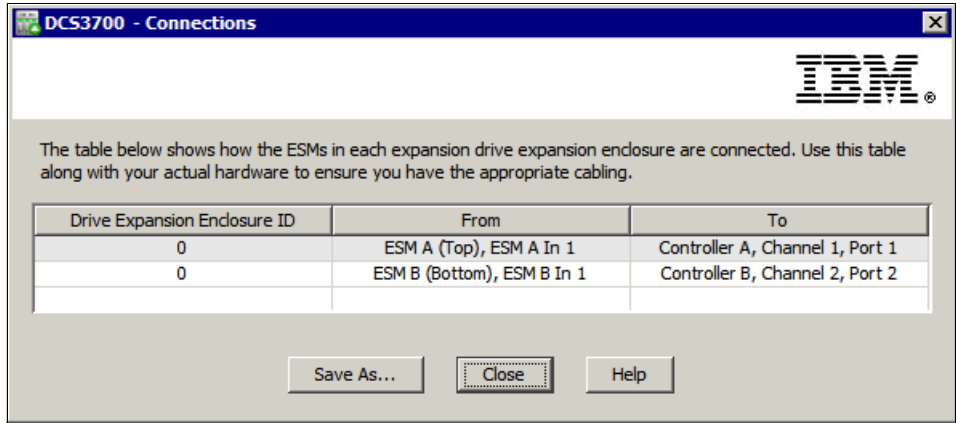


Figure 13-28 Cable Connections window

## 13.12 Event Log

This report option is used to view a list of events that have occurred on your DCS3700 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 an 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 configured to receive the alert notifications, as described in "Configure Alerts option" on page 168.

**Note:** The DCS3700 - Event Log can be accessed by clicking **Monitor** → **Reports** → **Event Log** or by clicking the **Summary** tab and clicking the **View Event Log** link in the Monitor pane, as shown in Figure 13-1 on page 386



The initial Event Log view is shown in Figure 13-29.

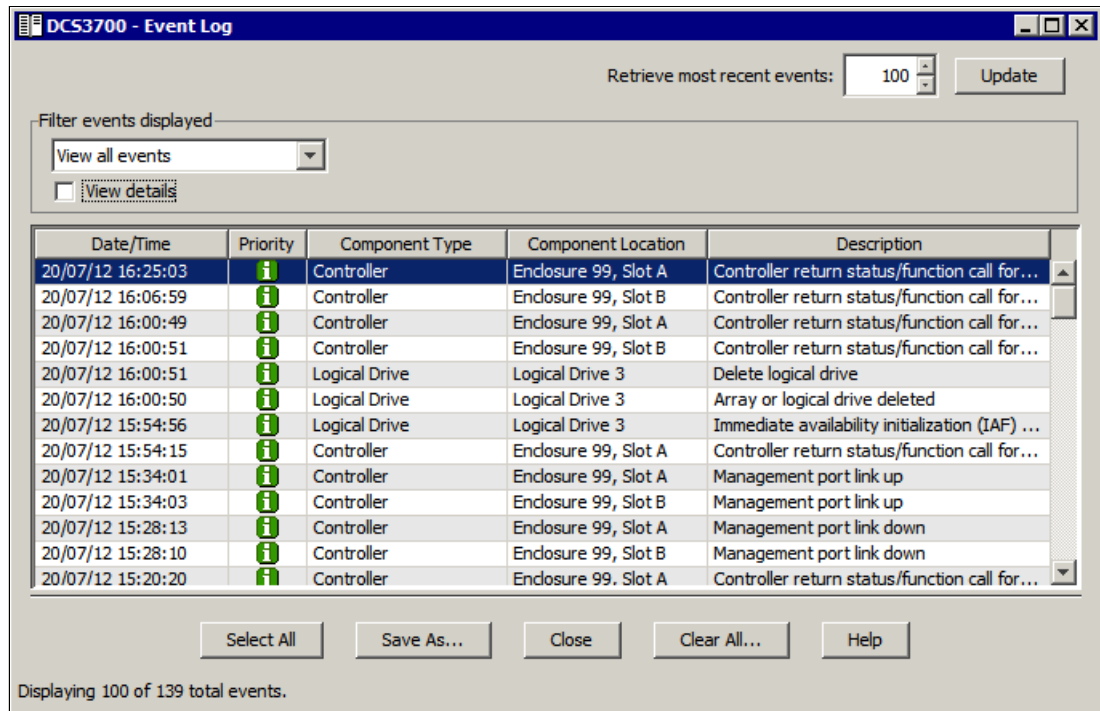


Figure 13-29 Event Log

The event log shows a summary for each event in the list:

- ▶ **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 13-30) that is in the upper left side of the Event Log window. The event details show more information about the event (Figure 13-31).




Figure 13-30 View details check box

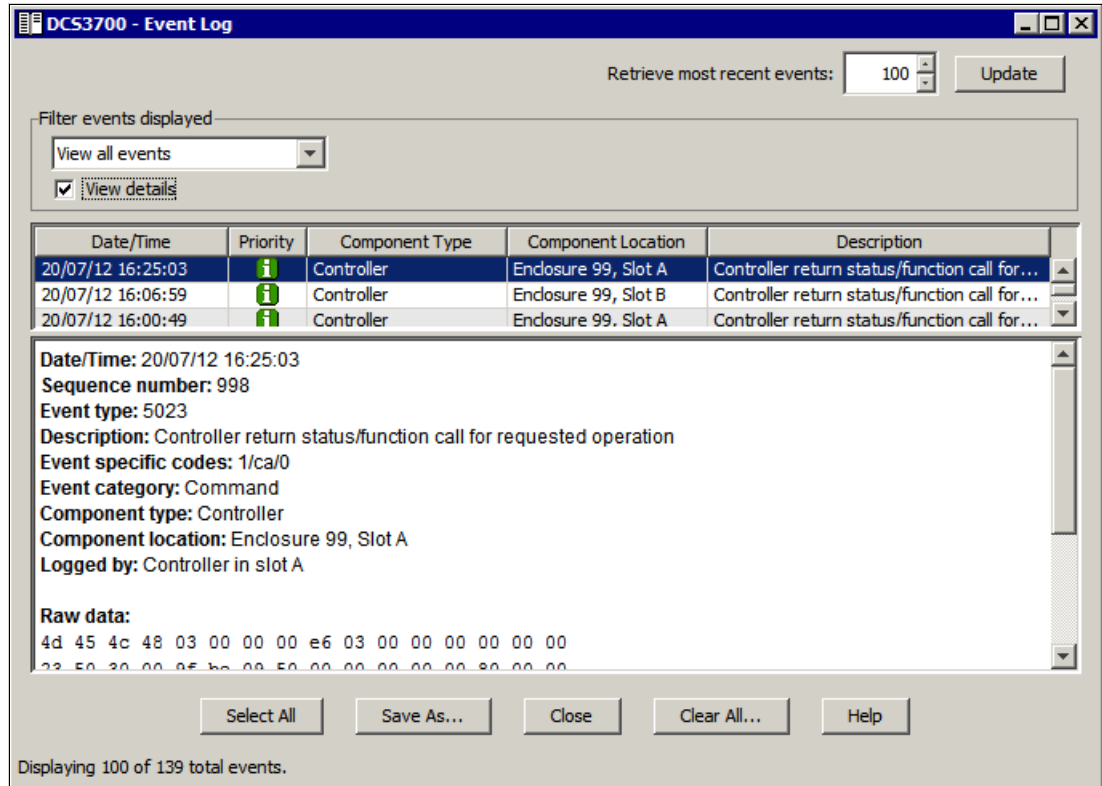



Figure 13-31 Viewing the Event Log details

The detailed 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 failed. For example, a disk drive or battery failed.
  - **State Change:** A logical device in the storage subsystem changed status. For example, an optimal logical drive is marked offline.

- Internal: These are internal controller operations that do not require your action.
  - Command: A command is issued to the storage subsystem. For example, a hot spare disk drive has been assigned.
  - Host: The host software 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 upper right corner, you can set how many events to display; the default initial value is 100. To change this value, enter another value and click **Update** (Figure 13-32). The lower left part of the window shows the total number of events that are stored on the storage subsystem. The maximum number of events is 8191. If more events occur, the older events are removed from the log on a first in, first out basis.



*Figure 13-32 Change the number of recent events to display*

In the upper left corner, there is a drop-down menu and a check box:

- ▶ The drop-down menu options are View only critical events, View only warning events, View only information events, and View all events, as shown in Figure 13-33.
- ▶ 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 13-31 on page 412.

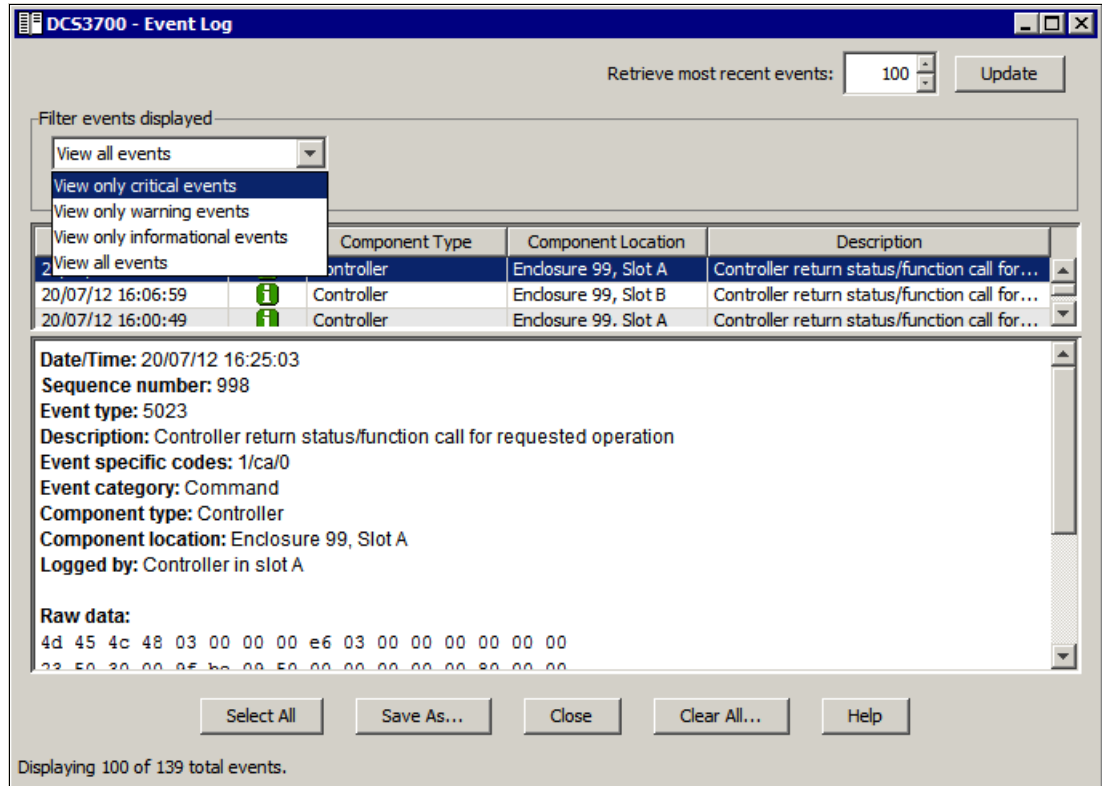


Figure 13-33 Drop-down menu to filter events

The buttons at the bottom of the window are:

- ▶ **Select All:** Click this button to highlight all events (Figure 13-34).

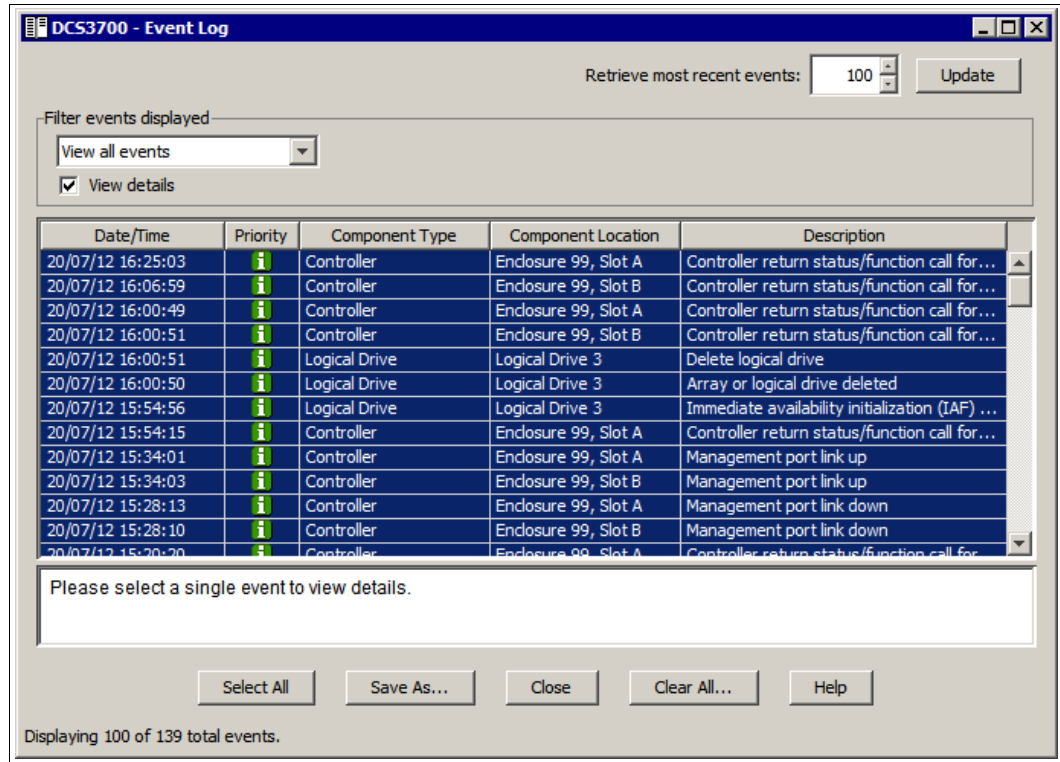


Figure 13-34 Select All events

- ▶ **Save As:** This button saves the highlighted events. To save several entries, press the Ctrl key to select them. To save all entries, click **Select All** and then click **Save As**. You are prompted for a location to save the events (the default extension is \*.log).

Figure 13-35 shows a sample log file that is viewed in Notepad.

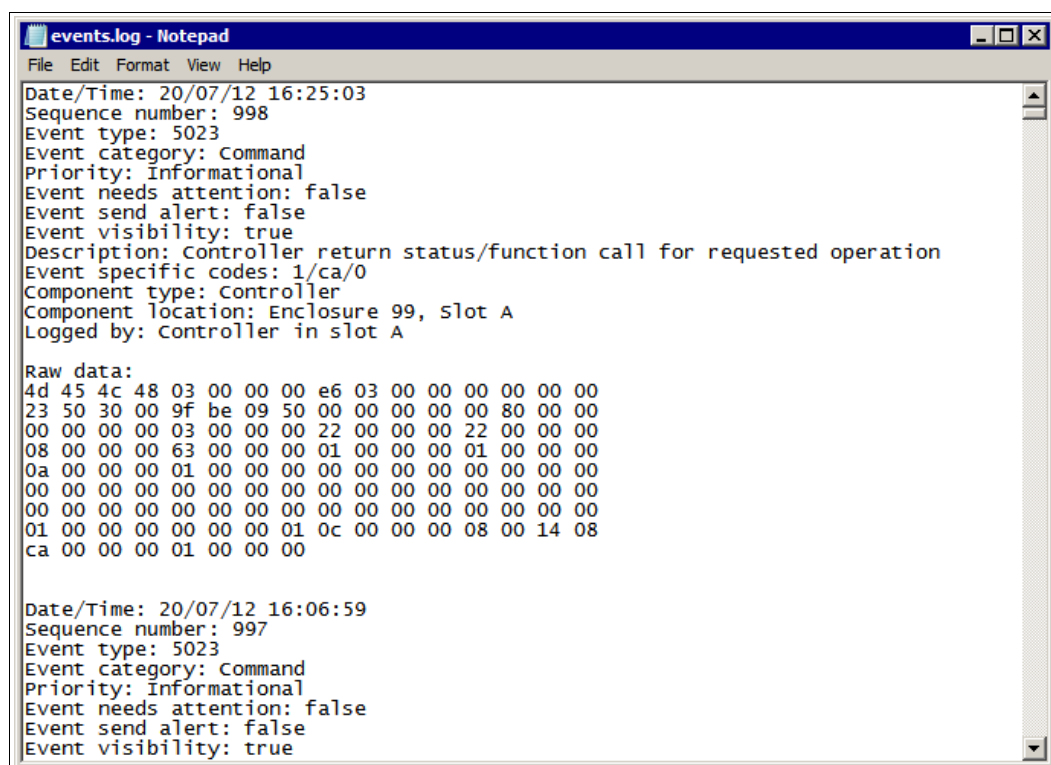


Figure 13-35 Event view - log

- ▶ Close: Closes the Event viewer.
- ▶ Clear All: Select this option to clear the event log. In the confirmation window (Figure 13-36), enter yes to delete all the events. Save all event logs before you clear the log.

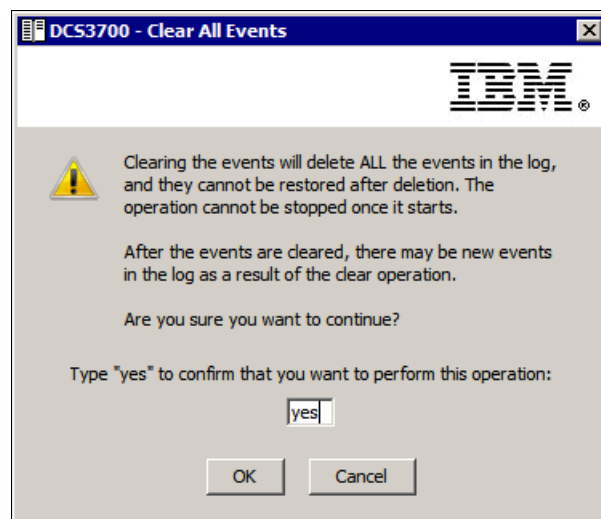


Figure 13-36 Event view - Confirm Clear All Events

- ▶ Help: Opens the online help.

## 13.13 Unreadable Sectors Log

Use the Unreadable Sectors Log dialog to show a detailed list of all of the known unreadable sectors that are detected on the storage subsystem. This dialog lets you clear and save unreadable sectors.

This option is accessed by clicking **Monitor** → **Reports** → **Unreadable Sectors Log**.

**Possible loss of data:** Any data in the Unreadable Sectors Log cannot be recovered and should be considered lost. Use this option only under the guidance of your IBM Customer Support representative.

## 13.14 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 subsystem power loss

**Note:** This function should normally be used only 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 (CLI).

To view the persistent reservations, complete the following steps:

1. Click **Monitor** → **Reports** → **Persistent Reservations** to open the Persistent Reservations window, as shown in Figure 13-37.

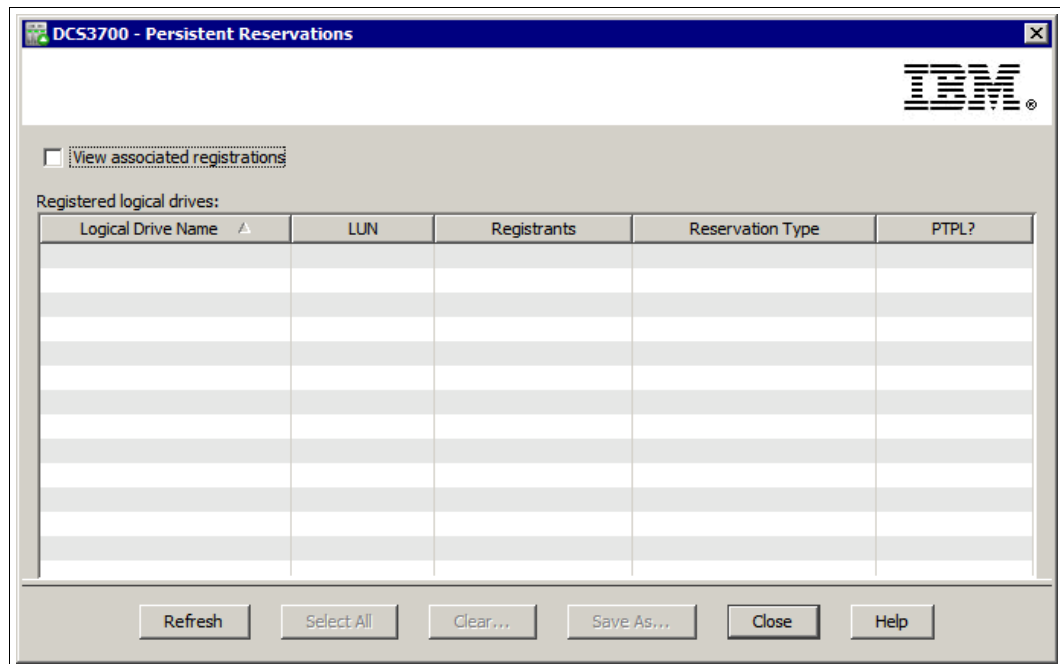


Figure 13-37 Persistent Reservations

2. To view the registrations that are associated with the logical drive reservation, either double-click the logical drive or highlight the logical drive and select the **View associated registrations** check box in the upper left of the window. You can view only one registration for a logical drive at a time.
3. Click **Refresh** to update the reservation and registration view.

You can also perform the following functions in this window:

- ▶ Click **Save As** 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 highlighted in the upper pane.
- ▶ Click **Clear** 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.

To clear persistent reservations, complete the following steps:

- a. In the upper left corner of the Persistent Reservations window, make sure that the **View associated registrations** check box is not selected.
- b. Click one or more logical drives. To select all of the logical drives, click **Select All**.
- c. Click **Clear**.
- d. In the text box in the Clear Registrations/Reservations window, enter yes and click **OK**.  
The reservation and registrations that were associated with the logical drives that you highlighted in the upper pane are cleared.



## 13.15 IBM Remote Support Manager for Storage

The IBM Remote Support Manager for Storage (RSM for Storage) software is a no additional charge software package that is installed on an IBM System x server that is running Novell SUSE Linux, Red Hat Enterprise Linux, and VMware. It provides problem reporting and remote access for IBM Support for the IBM System Storage DCS3700, DS3000, DS4000, and DS5000 families.

The RSM for Storage software can be downloaded from the IBM Support website, which can be found at:

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

After you open the website, select the appropriate IBM System Storage system and operating system. Also, check the Compatibility Guide (included in the RSM for Storage package) for information about dependencies and requirements regarding your hardware and operating system. 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 for Storage.

The problem reporting utility that is provided by RSM for Storage automatically creates an entry in the IBM call management system for each storage subsystem that reports a problem. This is the equivalent of placing a voice call to IBM Support for a problem. After they are in the system, problems are responded to with the priority that is specified by the maintenance agreement in place for the product.

**Note:** During off-shift hours, in addition to RSM for Storage reporting a problem, you can call your local IBM Customer Support.

For support telephone numbers in your country or region, go to the following URL:

<http://www.ibm.com/planetwide>

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.

Remote access to RSM for Storage by IBM Support is provided by either an external modem that is attached to the server or through an external SSH connection. This connection provides IBM Support with a CLI to the server. The newest version allows users to configure RSM for Storage with *One-way Call Home*, which does not require remote access by IBM Support. For most hardware failures, IBM Support can resolve problems and RSM for Storage can automatically clear alerts for storage problem that are repaired. All bulk data transfers for logs and other problem determination files are sent to IBM through email by using the server's Ethernet interface, as shown in Figure 13-38 on page 420. 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 cannot change any security features of the software.

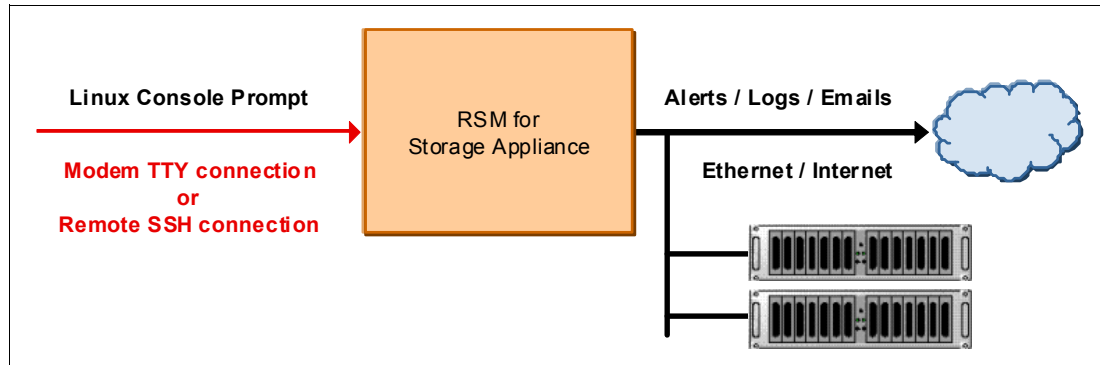


Figure 13-38 RSM layout

IBM System Storage DS Storage Manager V10, also known as Storage Manager, monitors the storage subsystems. This software is configured to send SNMP traps to the Remote Support Manager when critical events are detected.

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 also receives emails about status changes in these areas.

**Hardware withdrawal:** IBM System Storage DS Remote Support Manager Model RS3 and associated features have been withdrawn from marketing. For more information, go to:

<http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=an&subtype=ca&htmlfid=897/ENUS911-180&appname=iource&language=enus>

The RSM for Storage software is only available for download and installation on a client-provided server.

One RSM for Storage server can support up to 50 DCS3700, DS3500, DS4000, DS5000, FASiT 200, and FASiT 500 storage subsystems. Only IP connectivity to the Ethernet management ports of the subsystems is required: serial cables are not needed for access to the disk subsystems.

For more information about RSM for Storage, see the following resources:

- ▶ *IBM System Storage DS5000 Series Hardware Guide*, SG24-8023
- ▶ *IBM Remote Support Manager for Storage Version 2.8 Planning, Installation, and User's Guide*, found at:  
<http://www-947.ibm.com/support/entry/portal/docdisplay?lnocid=migr-5087766>
- ▶ IBM DCS3700 Fix Central page website:  
[http://www.ibm.com/support/fixcentral/swg/quickorder?parent=Mid-range+disk+systems&product=ibm/Storage\\_Disk/DCS3700+Storage+System&release=All&platform=All&function=all&source=fc](http://www.ibm.com/support/fixcentral/swg/quickorder?parent=Mid-range+disk+systems&product=ibm/Storage_Disk/DCS3700+Storage+System&release=All&platform=All&function=all&source=fc)



## Administration: Upgrade

This chapter describes the upgrade functions for the IBM System Storage DCS3700 storage subsystem. It covers the steps to upgrade the storage subsystem firmware and NVSRAM microcode. These functions are accessible through the DS Storage Manager (Enterprise Management) or DS Storage Manager (Subsystem Management).

## 14.1 The Subsystem Management Upgrade menu

This section lists the upgrade functions that are accessed from the DS Storage Manager (Subsystem Management) Monitor menu, as shown in Figure 14-1.

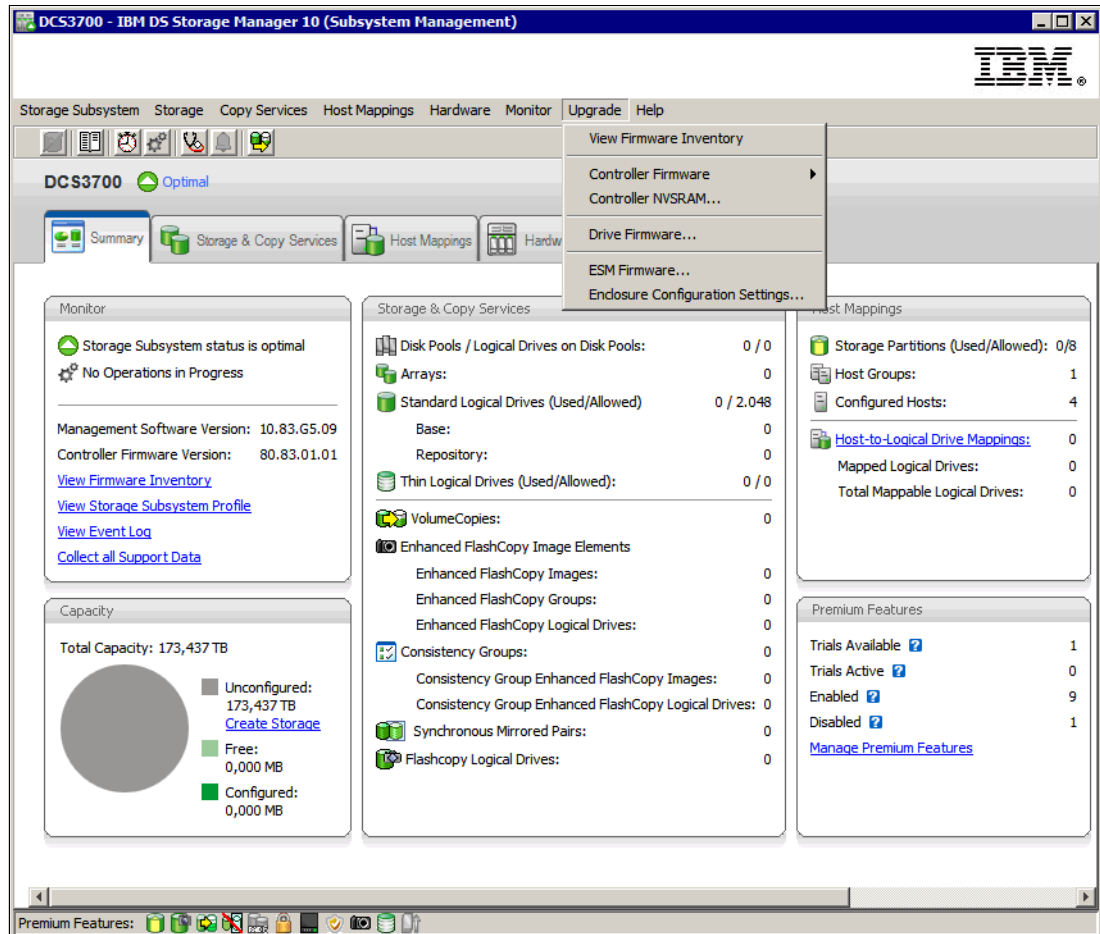


Figure 14-1 The DS Storage Manager (Subsystem Management) Upgrade menu

The Upgrade menu (Figure 14-1) consists of the following items:

- ▶ View firmware Inventory: Collects and displays all firmware information from the system, as described in 14.2.2, "Firmware inventory" on page 425.
- ▶ Controller Firmware: Used to start the firmware updates, activation, and cleaning, as described in 14.3, "Upgrading the DCS3700 controller firmware" on page 428.
- ▶ Controller NVSRAM: Used for NVSRAM update operations only, as described in 14.4, "Controller NVSRAM" on page 453.
- ▶ Drive Firmware: Updates the hard disk drive firmware version, as described in 14.5, "Drive firmware" on page 455.
- ▶ ESM Firmware: Updates the External Enclosure card firmware version, as described in 14.6, "ESM firmware" on page 460.
- ▶ Enclosure Configuration Settings: Allows the transfer of configuration settings to and from External Enclosures, as described in 14.7, "Enclosure configuration settings" on page 462.

## 14.2 Downloading firmware

This section describes the steps that are completed to obtain the DCS3700 storage subsystem firmware and NVSRAM microcode. You can download the latest firmware from the following website:

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

Appendix A, “IBM Support Portal website” on page 583 describes this website and its contents.

On the same website, you can select (as shown in Figure 14-2, for the DCS3700) the firmware (controller firmware, 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 DCS3700 storage subsystem are available at this site.

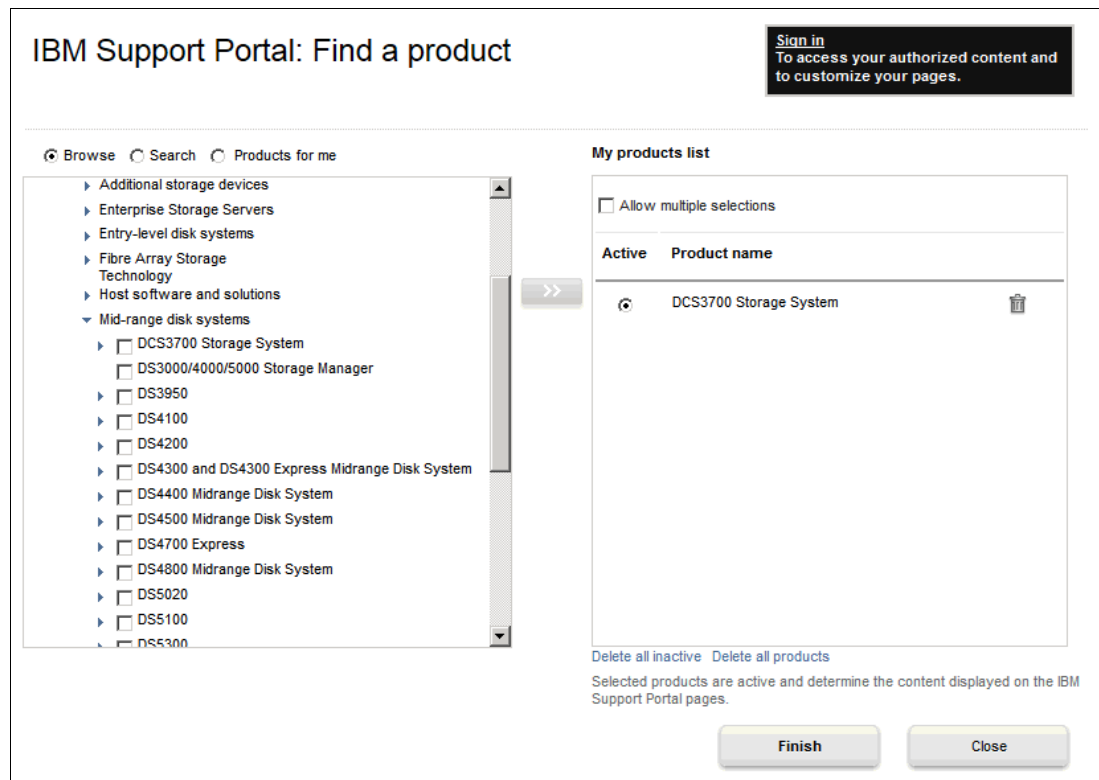


Figure 14-2 Code downloads for DCS3700 storage subsystem

The DCS3700 firmware version is closely connected to the Storage Manager version, and the HBA device driver and firmware are available for your hosts. You should always check the DCS3700 firmware readme and change-history files to see whether there are any host software dependencies. If there is a newer version of Storage Manager (including multipath device drivers) or HBA software is required for the hosts, download and install these items first before updating the DCS3700 firmware. The ESM firmware should also be upgraded to the required level (for more information, see 14.6, “ESM firmware” on page 460) before you start the controller firmware upgrade. In general, run at the latest levels unless advised otherwise.

**Note:** Updating firmware normally requires a maintenance window, so it is better to plan an appropriate outage time in advance.

Download all the packages that you need from the IBM Support website and extract and store them on a local disk.

**Important:** The procedures in this section are provided as a guide. However, these procedures can change. Always carefully check the readme files that are distributed with any firmware package and follow *those* procedures exactly.

## 14.2.1 Before you upgrade

There are precautions that must be taken 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 14-3 on page 425. If not, run the Recovery Guru, as described in 13.2, “View Health (Recovery Guru)” on page 387, 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 DCS3700 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 multipath 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 DCS3700 controller must be manageable the entire time.

**Note:** The overall DCS3700 Monitor status and current firmware version can be checked in the Summary tab of the Subsystem Management Window, as shown in Figure 14-3.

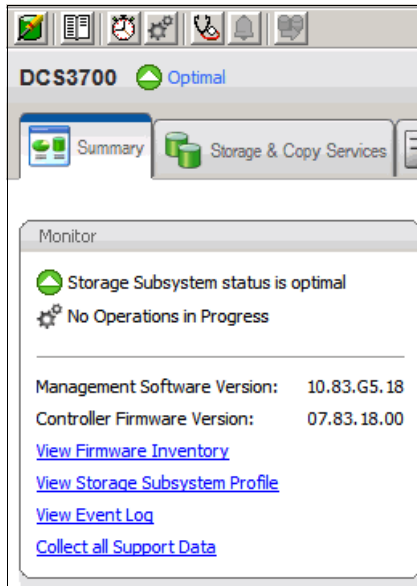


Figure 14-3 Viewing the DCS3700storage subsystem status in DS Storage Manager

**General rule:** Save the All Support Data collection described in 13.4, “Collect Support Data” on page 393 before starting any storage subsystem upgrades.

**Notes:**

- ▶ Update the hosts, and ensure that the correct version of the multipath driver is installed.
- ▶ Update the ESMs first to the latest level (as described in 14.6, “ESM firmware” on page 460) followed by the controller firmware. Outdated ESM firmware could make your expansion units inaccessible after the DCS3700storage subsystem firmware update. The required ESM firmware level for each controller firmware level is mentioned in the readme file for the controller firmware.
- ▶ Update the Controller firmware, followed by the NVSRAM. These two components must be updated together or immediately in this order.
- ▶ Any power or IP network, SAN, or SAS interruption during the update process might lead to configuration corruption. Therefore, do not power off the DCS3700storage subsystem or the management station during the update.

## 14.2.2 Firmware inventory

Before and after updating the firmware levels of the DCS3700storage subsystem, you can use the **Firmware Inventory** option to view the complete firmware levels for your storage subsystems.

To view the firmware inventory for one storage subsystem, complete the following steps:

1. Click **View Firmware Inventory** in the Monitor pane of the Summary tab of the Storage Subsystem window (Figure 14-4).

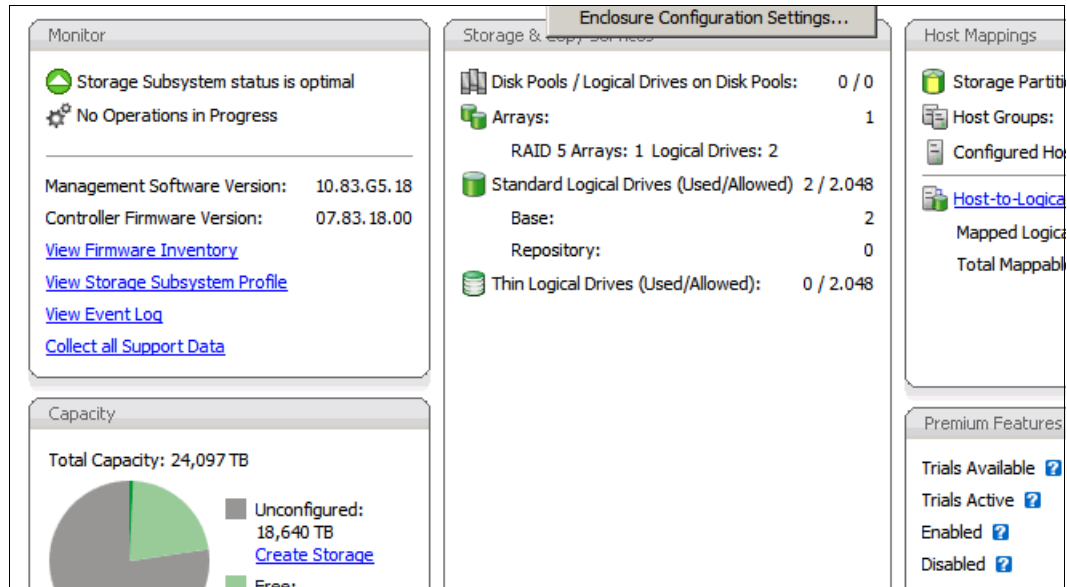


Figure 14-4 Storage Subsystem - Monitor

2. Figure 14-5 shows the Firmware Inventory window that opens. Save the firmware inventory to a text file by clicking **Save As**.

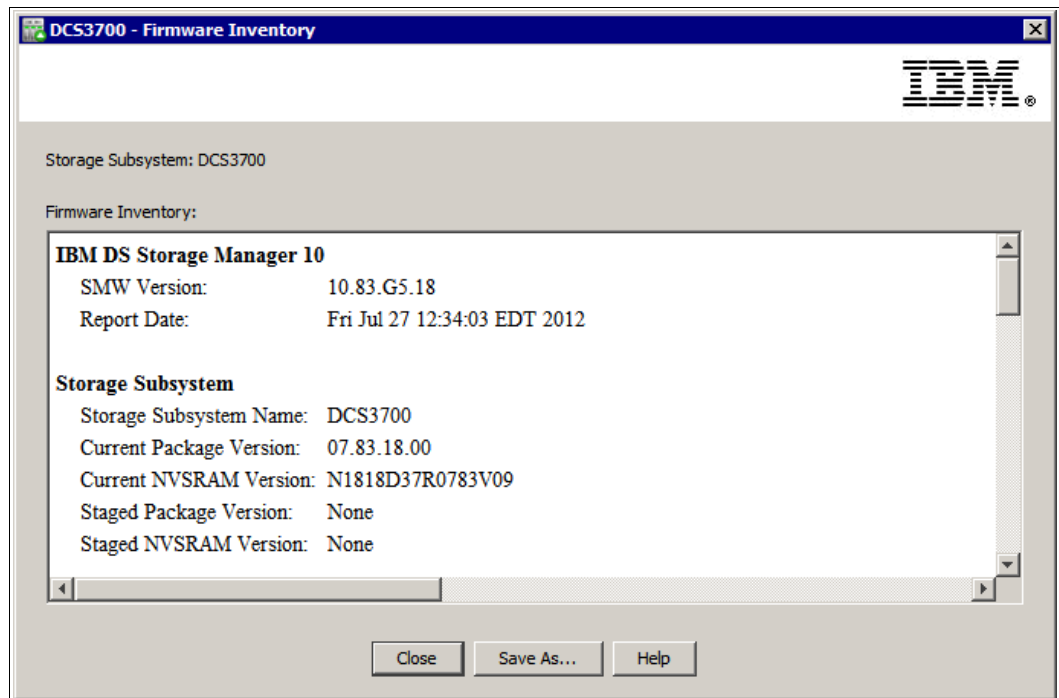


Figure 14-5 Storage Subsystem - Firmware Inventory



To view the firmware inventory for more than one storage subsystem, complete the following steps:

1. Click **Tools** → **Firmware Inventory** from the Enterprise Management window. Figure 14-6 shows the Firmware Inventory window that opens. From this window, it is possible to select different storage subsystems.

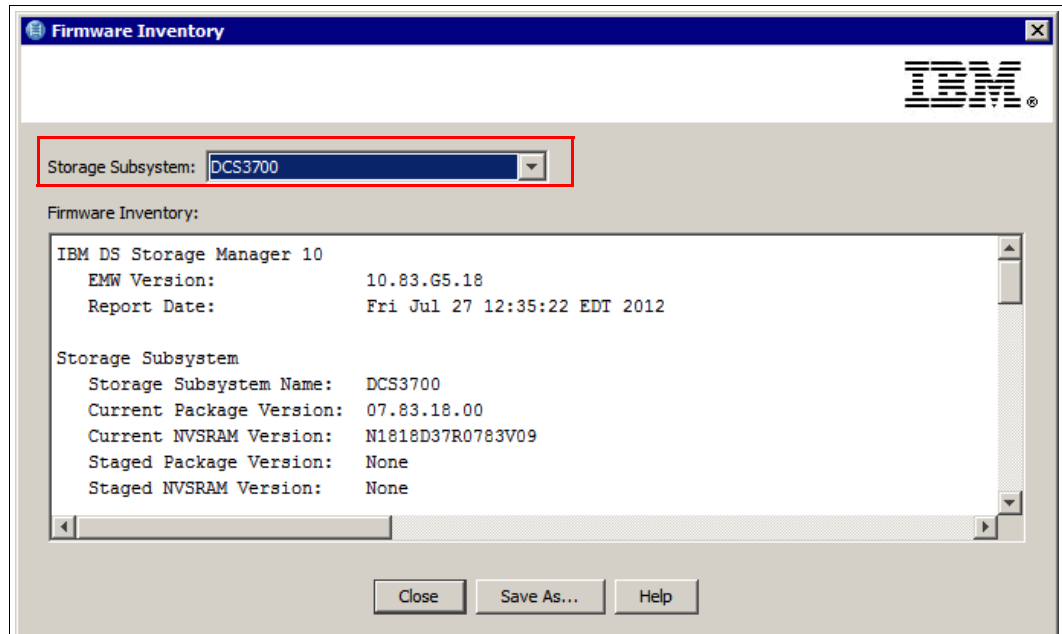


Figure 14-6 Enterprise Management - Firmware Inventory

2. Save the firmware inventory to a text file by clicking **Save As**.

**Note:** If you want to save the firmware inventory for more than one storage subsystem, you must select each of the storage subsystems in the drop-down box and then click **Save As**.

### 14.2.3 Updating the host

Before you update the firmware on your DCS3700 storage subsystem, 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 newer versions are available, download them and check the readme files for any special requirements before you update.

Update the HBA driver and Storage Manager if necessary. Updating Storage Manager is described in Chapter 4, “IBM System Storage DS Storage Manager installation” on page 131. Storage Manager for Windows also includes the associated multipath driver. The multipath RDAC driver is a separate package for Linux. Other host operating systems provide their own multipath drivers. The required multipath driver must 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 that it is at the latest supported level. You can also find this information at the IBM Support website.

## 14.3 Upgrading the DCS3700 controller firmware

This section describes how you can upgrade the firmware of the controllers in the storage subsystem by using the DS Storage Manager software.

**Important:** During the update of one controller, all Fibre Channel links for this controller are disconnected. When the upgrade of the first controller is finished, the second one starts immediately with the firmware upgrade and the second controller drops all FC links. Some FC topologies need some time to re-establish the link to the first controller. So, this task can result in a complete loss of access for a short period.

### 14.3.1 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 DCS3700 storage subsystem.
2. The DCS3700 storage subsystem is upgraded with the new firmware file. The process of upgrading 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. This is known as a *staged* upgrade.

### 14.3.2 Pre-upgrade checks

The storage management software checks for existing conditions in the storage subsystem that 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 DCS3700 storage subsystem as *non-upgradeable* (see Figure 14-8 on page 432). 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 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 corrupted.
- ▶ 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 that are 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 (Figure 14-9 on page 432).

### 14.3.3 Major and minor firmware upgrades

A version number exists for each firmware file. For example, at the time of the writing of this book, 07.83.22.00 is the current version number for the DCS3700 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, see 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 07.83.22.00, 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 07.83.22.00, 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 that are included in the DS Storage Manager software to update the controller firmware and NVSRAM:

- ▶ The firmware upgrade tool, which is 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 DCS3700 storage subsystem, 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 14.3.4, “Using the Enterprise Management upgrade tool” on page 430.
- ▶ The firmware upgrade tool, which is 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 cannot perform either staged firmware upgrades or major firmware upgrades. This tool is described in 14.3.5, “Using the DCS3700 Storage Manager (Subsystem Management)” on page 448.

**Note:** You also can use the CLI to download and activate firmware to the DCS3700 storage subsystems. For more information, see the About the Command Line Interface online help topic in the Enterprise Management window.

Table 14-1 lists the major differences in functionality between the DS Storage Manager firmware update tools.

*Table 14-1 Comparison of the firmware updaters functions*

<b>Update type</b>	<b>Enterprise Management</b>	<b>Subsystem Management</b>
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

#### **14.3.4 Using the Enterprise Management upgrade tool**

This section describes the procedure to update the DCS3700 controller firmware by using the firmware upgrade tool, which is accessible from the DS Storage Manager Enterprise Management window.

**Note:** You can use the upgrade tool to perform both minor and major controller firmware upgrades.

To use the Enterprise Management upgrade tool, complete the following steps:

1. From the DS Storage Manager Enterprise Management window, right-click a selected DCS3700 storage subsystem and select **Upgrade Controller Firmware**, as shown in Figure 14-7.

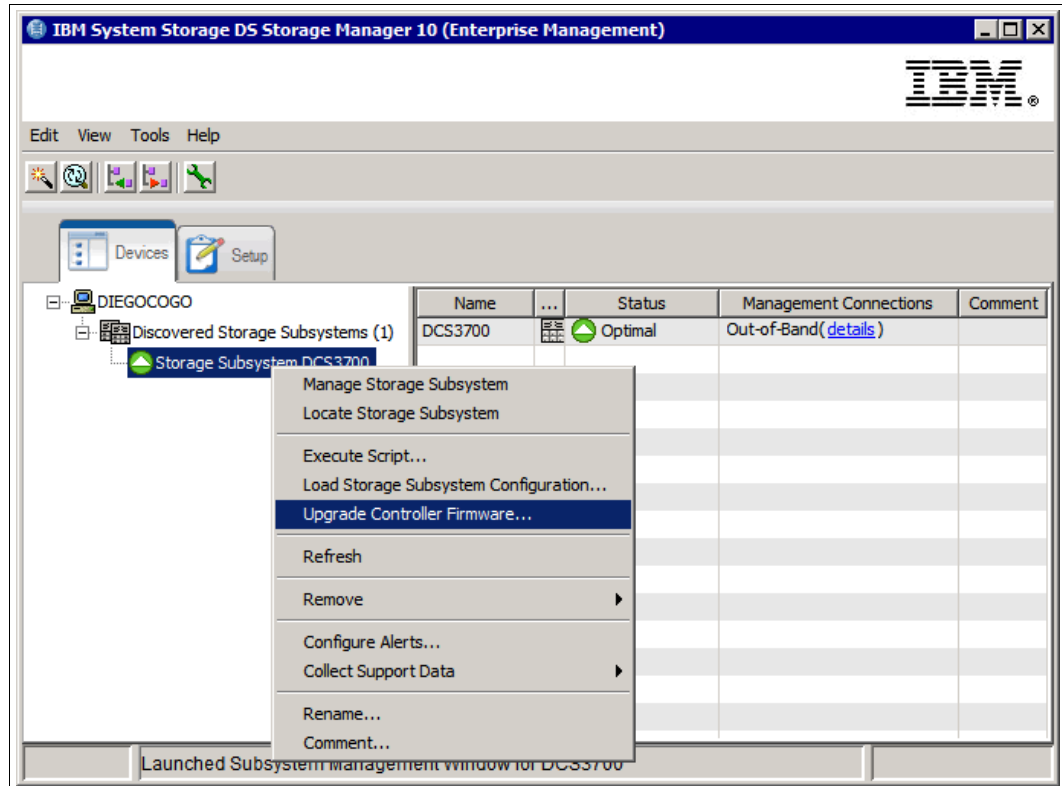


Figure 14-7 Upgrade Controller Firmware option in the Enterprise Management window

**Note:** This function is also accessible by clicking **Tools** → **Upgrade Controller Firmware...** in the DS Storage Manager (Enterprise Management) window.

- The Upgrade Controller Firmware window opens (Figure 14-8). In the example, 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.

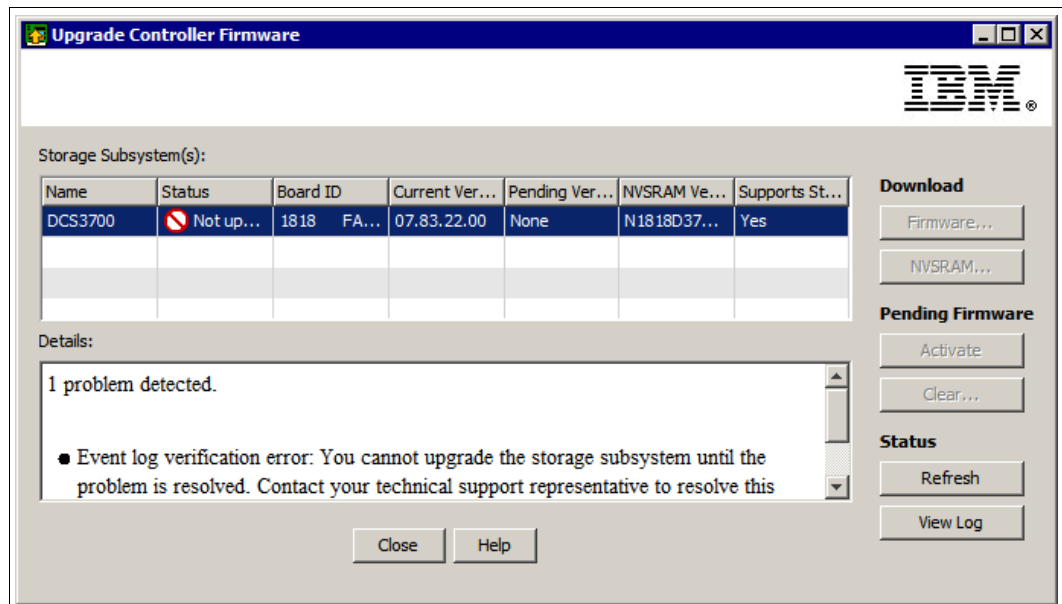


Figure 14-8 Upgrade Controller Firmware window

If the firmware upgrade tool reports a status of non-upgradeable, as shown in Figure 14-8, you can click **View Log** to review the cause of the non-upgradeable status. Figure 14-9 shows the upgrade log.

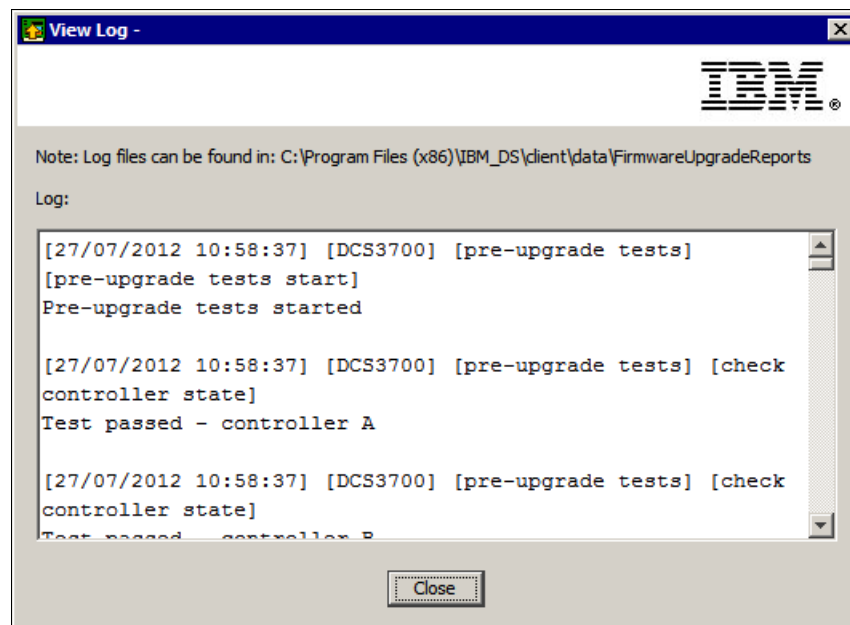


Figure 14-9 Viewing the firmware upgrade tool log

In this example, the problem that was preventing a firmware upgrade was fixed by clearing the DCS3700 storage subsystem event log. For instructions about how to clear the subsystem event log, see 13.12, “Event Log” on page 410. However, certain types of problems require that you contact your local IBM Customer 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 14-10.




Figure 14-10 Upgrade Controller Firmware

4. Click **Firmware** to open the Download Firmware window, which is shown in Figure 14-11.

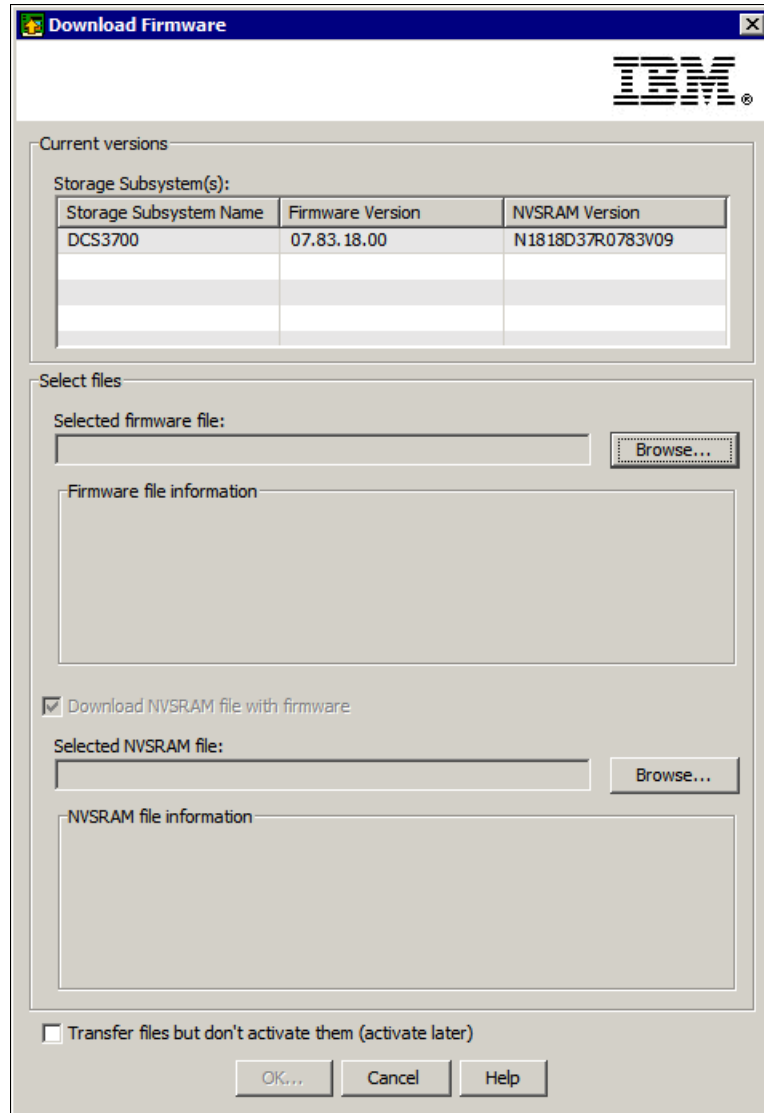


Figure 14-11 Download Firmware window

There are two main sections to this window:

- Current versions: This section reports the versions of controller firmware and NVSRAM that are installed.
- Select files: This section is where you select the firmware and NVSRAM files that you want to transfer to the DCS3700 storage subsystem. 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 **Browse** and select the firmware upgrade file, as shown in Figure 14-12.




Figure 14-12 Select firmware file

- b. In the Selected NVSRAM file pane, click **Browse** and select the NVSRAM file that you want to transfer (Figure 14-13).

**Attention:** Ensure that you choose the version of the NVSRAM file that is appropriate for your DCS3700storage subsystem. For example, if your DCS3700storage subsystem has two controllers, you must select the dual controller NVSRAM file. If your DCS3700storage subsystem has one controller, you must select the single controller NVSRAM file.

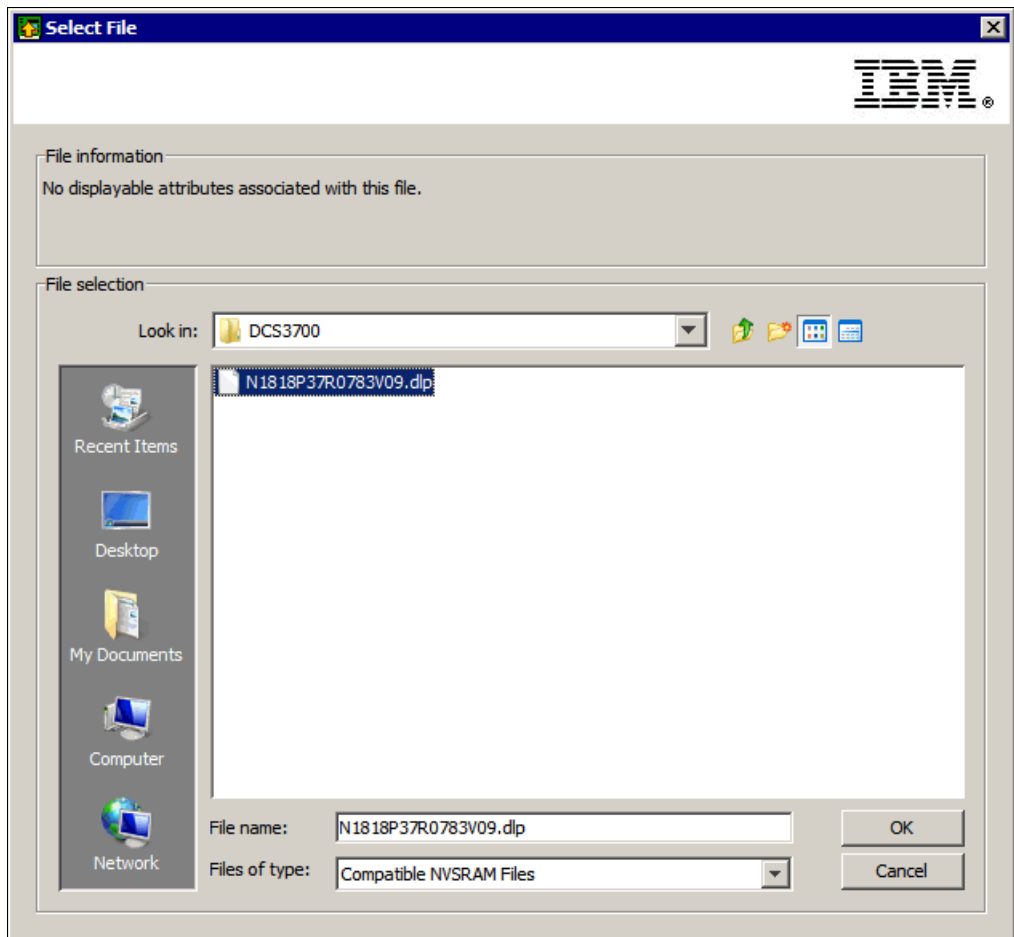


Figure 14-13 Select the NVSRAM file

5. If you want to transfer the firmware and NVSRAM and activate them immediately, then click **OK**. You are prompted to confirm that you want to transfer and activate the firmware and NVSRAM immediately (see Figure 14-14).

**Note:** The process to transfer the firmware and NVSRAM images to the DCS3700 storage subsystem, but activate them later, is described in step 6 on page 439.

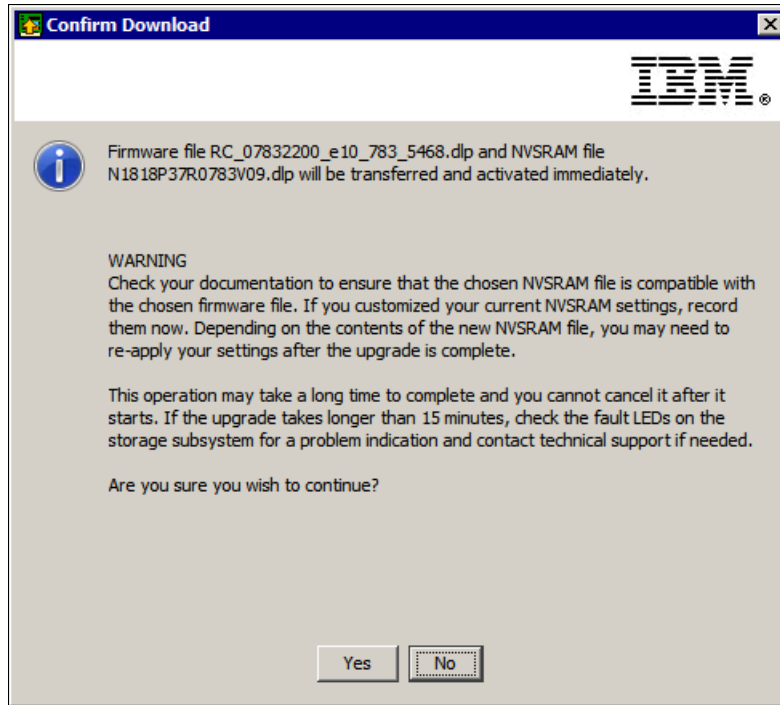


Figure 14-14 Transfer and activation confirmation window

Click **Yes** to start the transfer and activation. The firmware upgrade tool begins transferring the firmware images to the controllers. The status can be monitored in the tool (Figure 14-15).

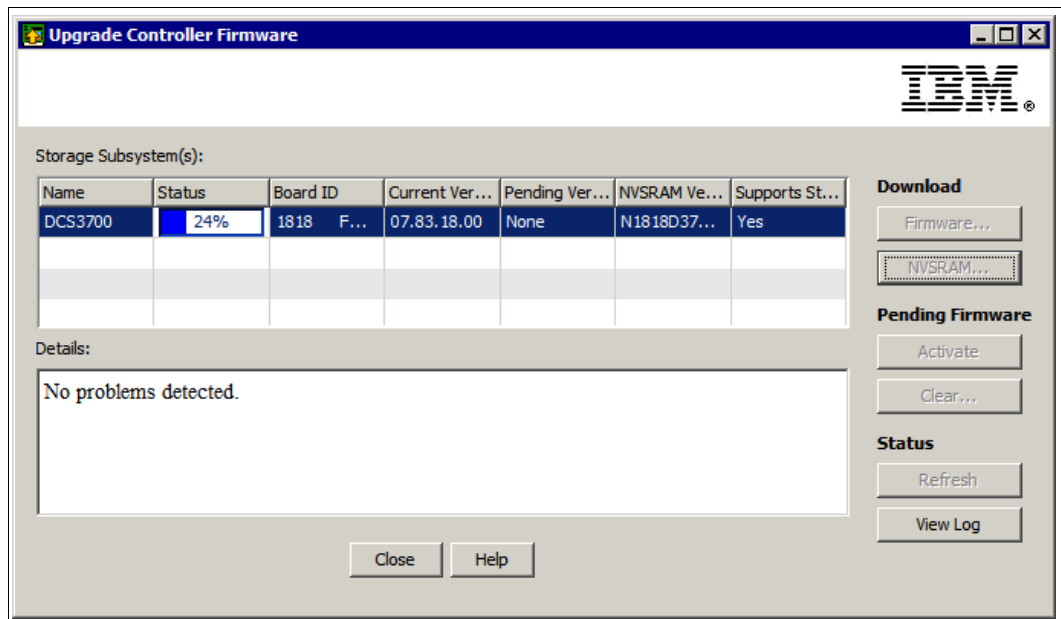


Figure 14-15 Firmware upgrade in progress

Figure 14-16 shows the status after the firmware is transferred and activated.

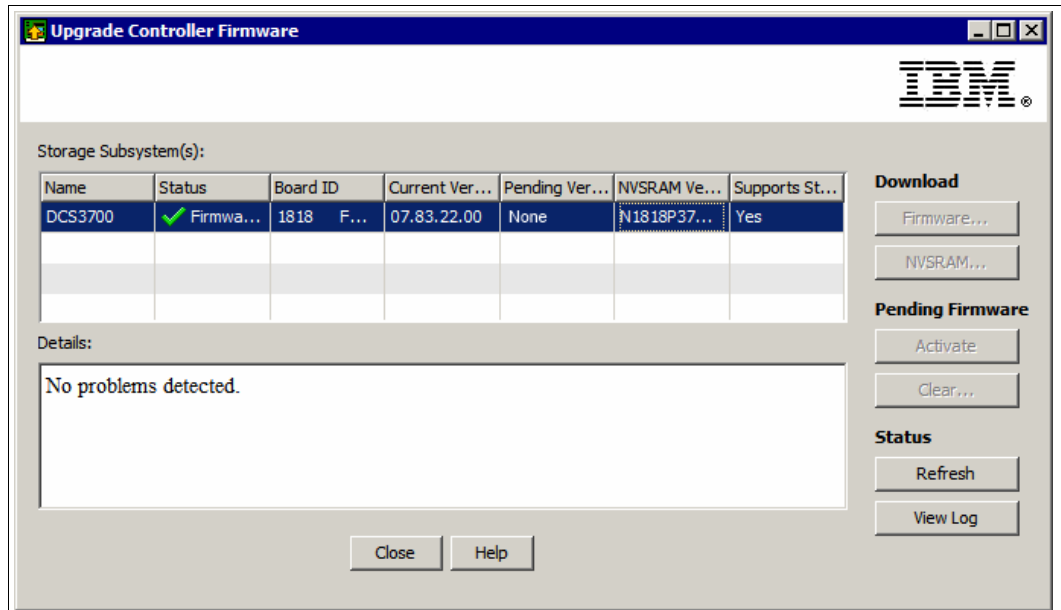


Figure 14-16 Successful firmware transfer and activation

6. The default procedure is to transfer the firmware and NVSRAM and activate the new versions immediately. However, you can decide to transfer only the firmware and NVSRAM now, and then activate the firmware and NVSRAM later by completing the following steps:
  - a. To transfer the firmware and NVSRAM, select the **Transfer files but don't activate them (activate later)** check box at the bottom of the Download Firmware window (see Figure 14-17).

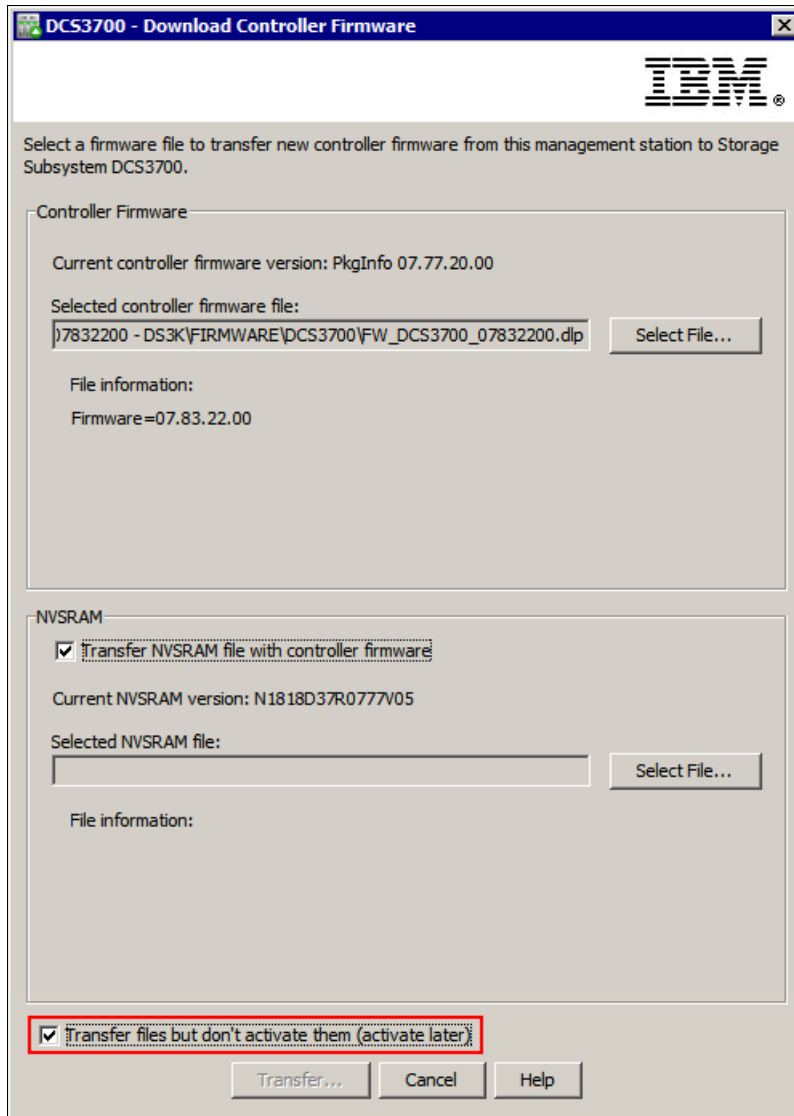


Figure 14-17 Transfer firmware and NVSRAM only (activate later)

- b. Click **OK**. You are prompted to confirm that you want to transfer the firmware and NVSRAM now but activate the new versions later (Figure 14-18).

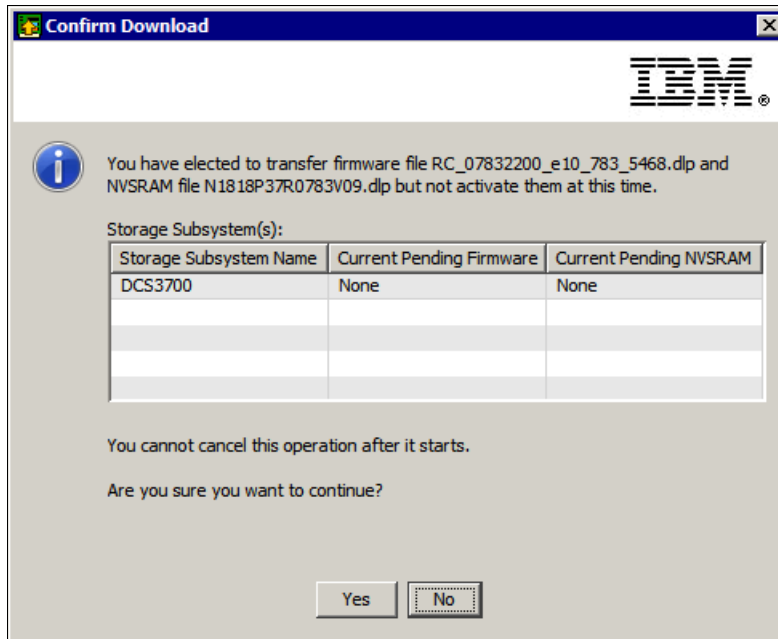


Figure 14-18 Confirm Download

- c. The status of the transfer of the firmware and NVSRAM can be monitored in the tool, as shown in Figure 14-19.




Figure 14-19 Transferring the firmware and NVSRAM

- d. Figure 14-20 shows the status after the firmware and NVSRAM are transferred. The status is reported as Firmware pending and the **Activate** button is now functional.




Figure 14-20 Firmware activation pending

**Note:** A pending firmware upgrade can also be activated in the DS Storage Manager (Subsystem Management) window by clicking **Upgrade** → **Controller Firmware** → **Activate...**

- e. Click **Activate** to activate the pending firmware. The window that is shown in Figure 14-21 opens.

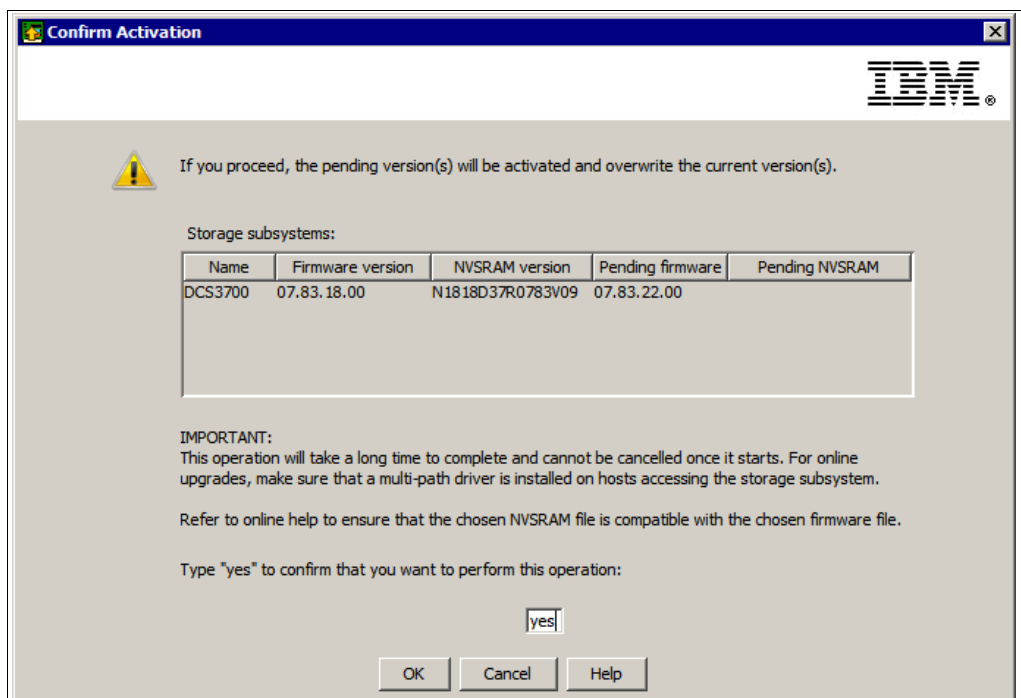


Figure 14-21 Confirm activation

- f. Enter yes in to the box to confirm that you want to activate the pending firmware and click **OK**.

The activation process can take some time to complete, but the status of the process can be monitored in the tool. Figure 14-22 shows the status after the firmware is successfully activated.

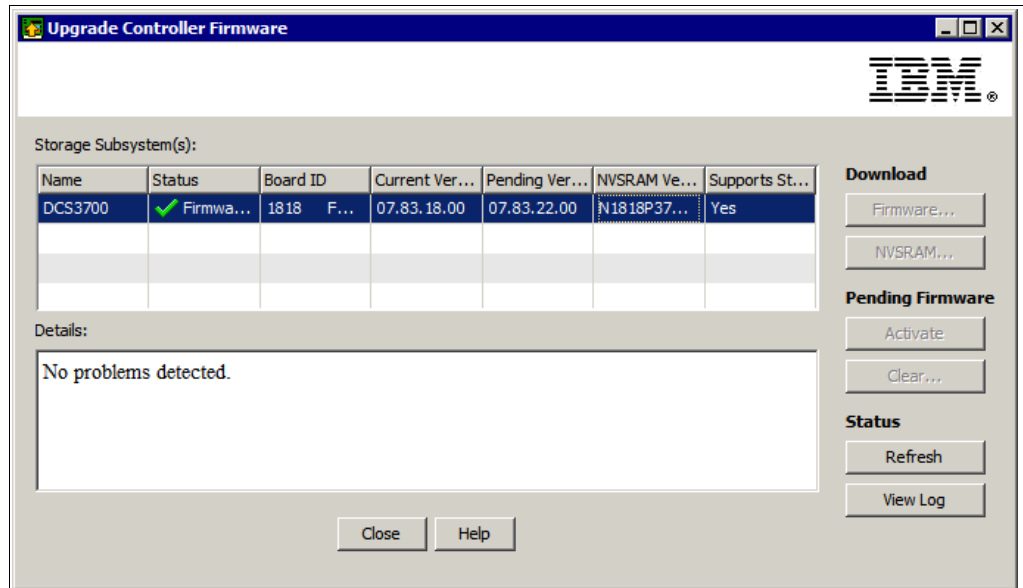


Figure 14-22 Firmware activated successfully



7. It is also possible to clear a pending firmware and NVSRAM upgrade image that has been transferred to the DCS3700 storage subsystem but not yet activated.

**Note:** You can also clear a pending firmware upgrade in the DS Storage Manager (Subsystem Management) window by clicking **Upgrade** → **Controller Firmware** → **Clear...**

**Note:** This step was taken for information purposes only; it is not a necessary step in the upgrade procedure.

- a. Click **Clear** in the Upgrade Controller Firmware window, as shown in Figure 14-23.

**Note:** You must click **Refresh** for the Clear button become enabled.

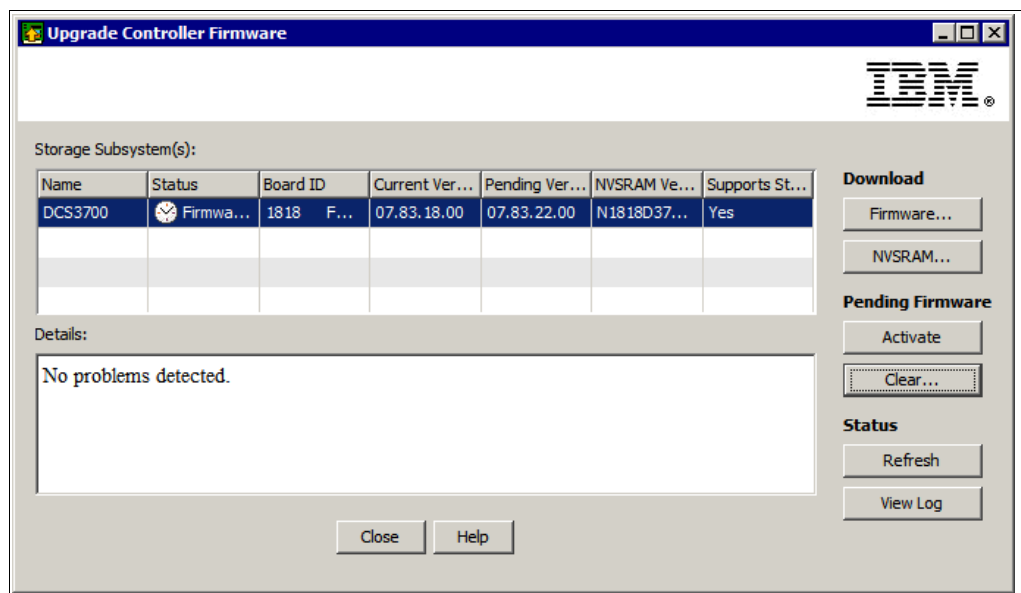


Figure 14-23 Pending firmware upgrade

- b. You are prompted to confirm that you want to clear the pending firmware upgrade, as shown in Figure 14-24.

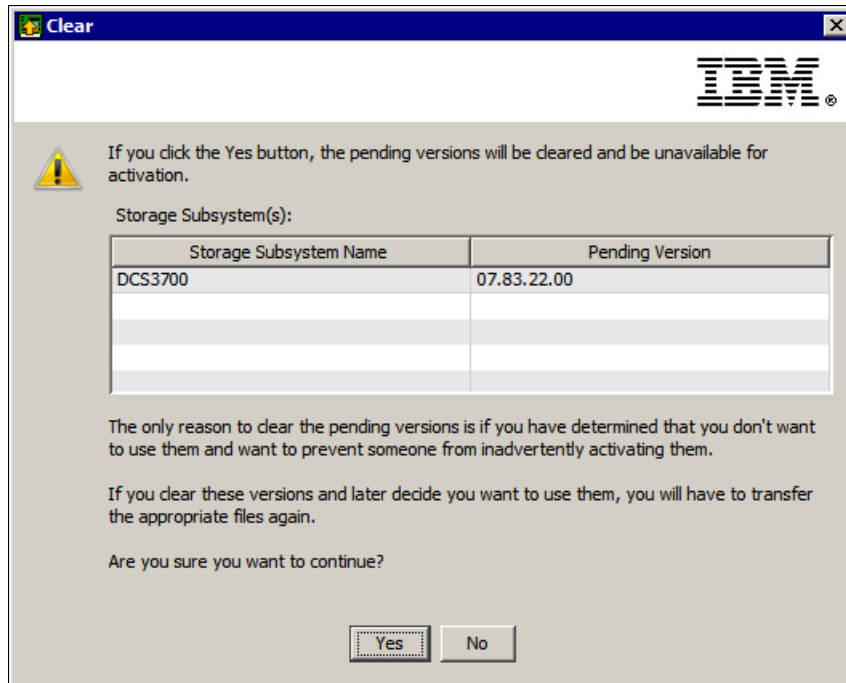


Figure 14-24 Clear pending firmware upgrade

- c. Click **Yes** to clear the pending firmware upgrade and the Clear Pending Firmware Progress window that is shown in Figure 14-25 opens.




Figure 14-25 Clear Pending Firmware Progress

- d. The Upgrade Controller Firmware window reports the status that is shown in Figure 14-26 after the firmware is cleared.

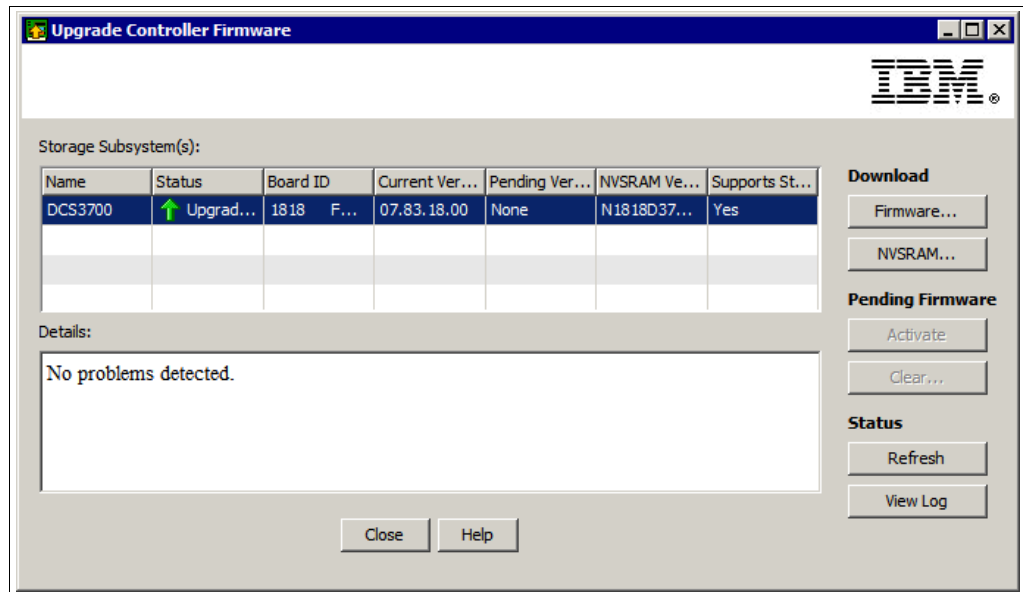


Figure 14-26 Status after the pending firmware is cleared

8. You can also use the upgrade tool to install only an updated NVSRAM file:
- Click **NVSRAM** in Figure 14-26. The Download NVSRAM window opens, as shown in Figure 14-27.

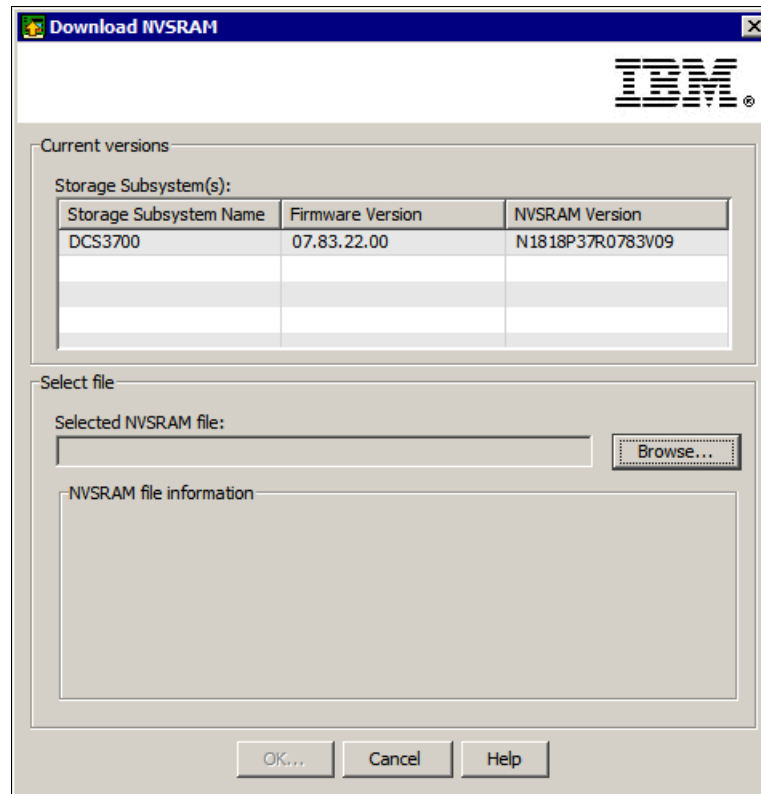


Figure 14-27 Download NVSRAM

- b. Click **Browse** to select the appropriate NVSRAM file (either single controller or dual controller NVSRAM) for your DCS3700 storage subsystem (Figure 14-28).

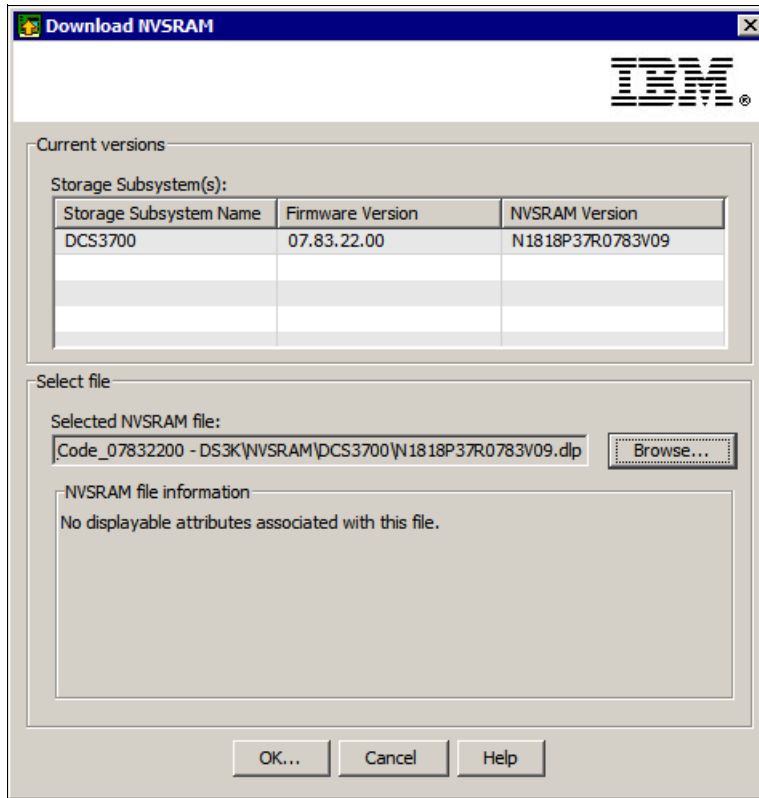


Figure 14-28 Download NVSRAM with the NVSRAM file selected

- c. Click **OK** to continue, and the **Confirm Download** window opens, as shown in Figure 14-29.



Figure 14-29 Confirm Download

- d. You can monitor the status of the NVSRAM upgrade in the Upgrade Controller Firmware window. Figure 14-30 shows the status after the NVSRAM is transferred and activated successfully.

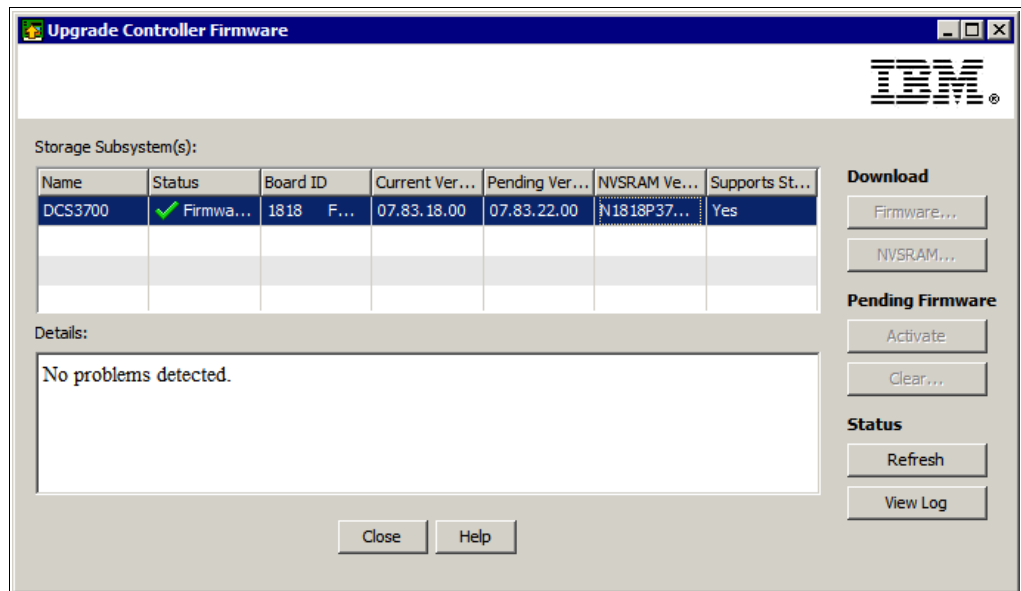


Figure 14-30 NVSRAM upgrade successful

### 14.3.5 Using the DCS3700 Storage Manager (Subsystem Management)

This section describes the upgrade procedure for the DCS3700firmware and NVSRAM by using the DS Storage Manager (Subsystem Management) application.

#### Updating the storage subsystem

After you have updated your hosts (as described in 14.2.3, “Updating the host” on page 427), you can update the code of your storage subsystem. To do this, click **Upgrade** → **Controller Firmware** → **Upgrade...**, as shown in Figure 14-31.

**Tip:** The current DCS3700firmware levels are listed in the storage subsystem profile. Clicking **Upgrade...** also displays the current firmware levels of the respective components.

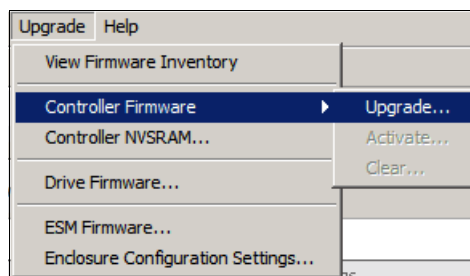


Figure 14-31 Upgrade menu

**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 website.

#### Download Controller Firmware

To install the DCS3700 storage subsystem firmware, complete the following steps.

**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 that is 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 that uses the Enterprise Management window to download and activate the controller firmware is described in 14.3.4, “Using the Enterprise Management upgrade tool” on page 430.

1. To update the controller firmware by using the DS Storage Manager, click **Upgrade...**, as shown in Figure 14-31.

- Every time that you click **Upgrade**, a warning message displays for an automatic pre-upgrade check, as shown in Figure 14-32.

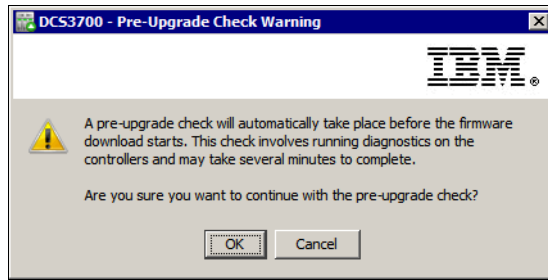


Figure 14-32 Pre-upgrade check warning message

- Sometimes after you click **OK**, a warning window might open because the DCS3700 storage subsystem has logged too many critical events.

You must view the DCS3700event log and resolve the errors (as described in 13.12, “Event Log” on page 410). In our testing lab environment, clearing the DCS3700storage subsystem event log fixed the condition that was causing the verification error. If you are unable to fix the condition that prevents the firmware upgrade, contact your local IBM Customer Support representative.

- If there are no event log verification errors, then the Download Controller Firmware window opens, as shown in Figure 14-33.

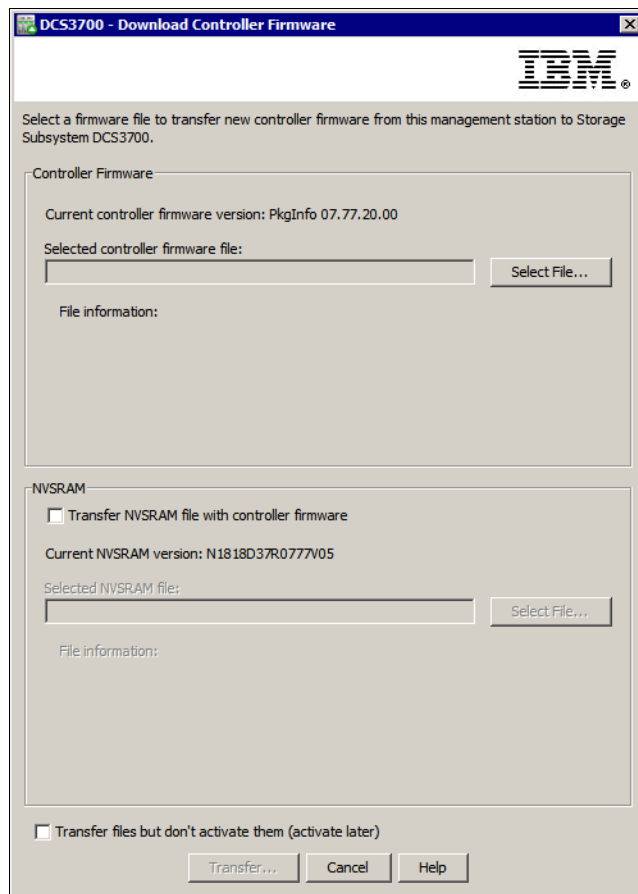


Figure 14-33 Download Controller Firmware

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 are done in one step.

**Note:** Install the controller firmware and controller NVSRAM updates in one step.

5. Click **Select File** in the Controller Firmware section and enter or browse to the file name.
6. 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 14-34).

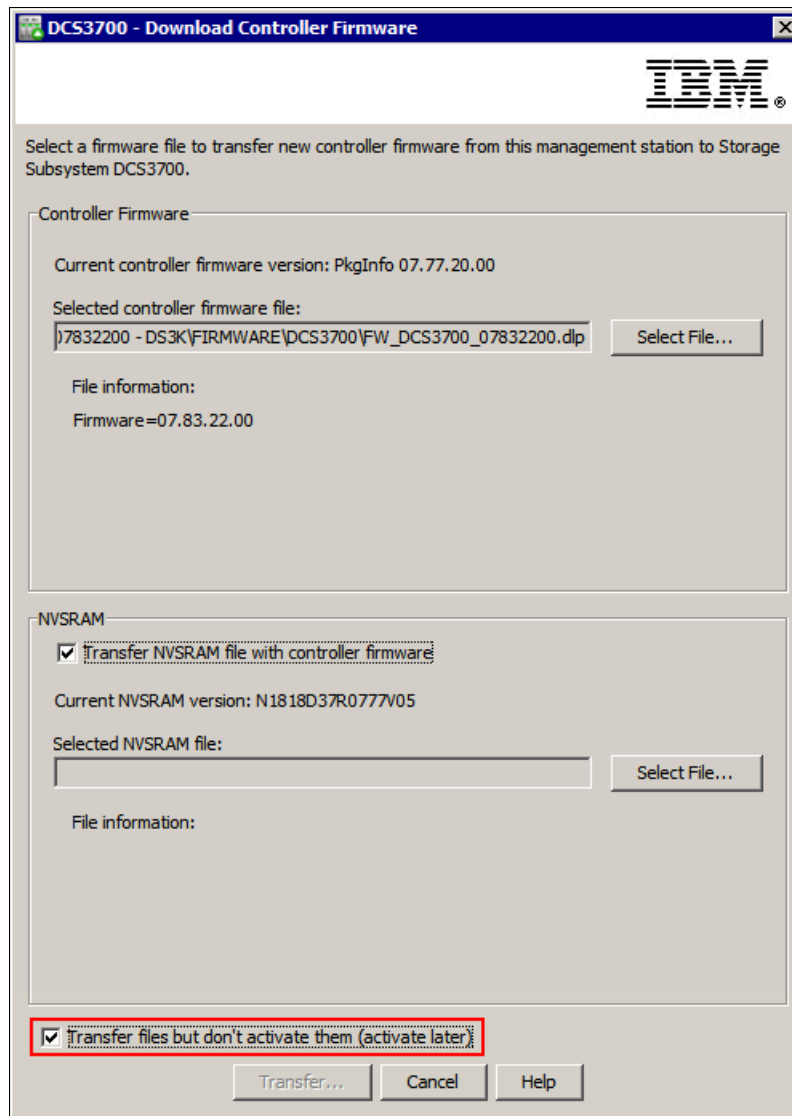


Figure 14-34 Download Controller Firmware with NVSRAM firmware

**Note:** It is a preferred practice to transfer files but not activate them in one step, as shown in Figure 14-34.



7. Click **Transfer** to submit the new microcode to the storage subsystem. The Confirm Download window opens (Figure 14-35). Read it and click **Yes** to continue.

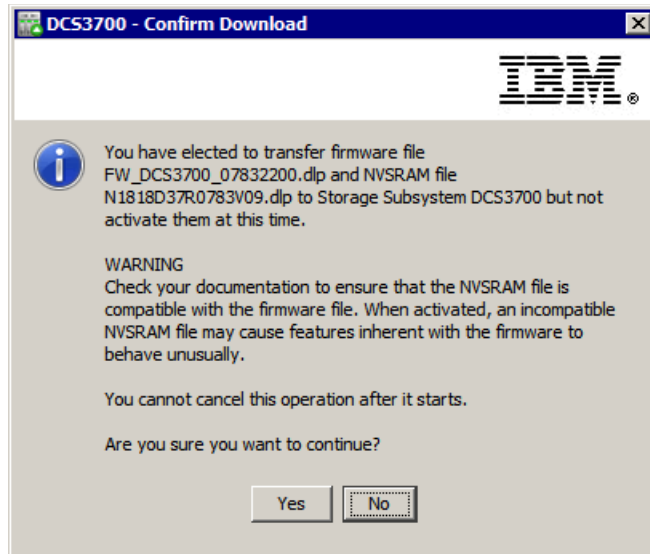


Figure 14-35 Confirm Download

8. The microcode is transferred to the storage subsystem and applied to both installed controllers. This procedure can take some time. You can monitor the status in the Downloading window that opens (Figure 14-36).



Figure 14-36 Downloading the firmware with activation postponed

After the controller firmware is transferred, it is necessary to activate both controllers. This step can also take some time.

To start the activation process, click **Upgrade** → **Controller Firmware** → **Activate...** A warning window opens, as shown in Figure 14-37.

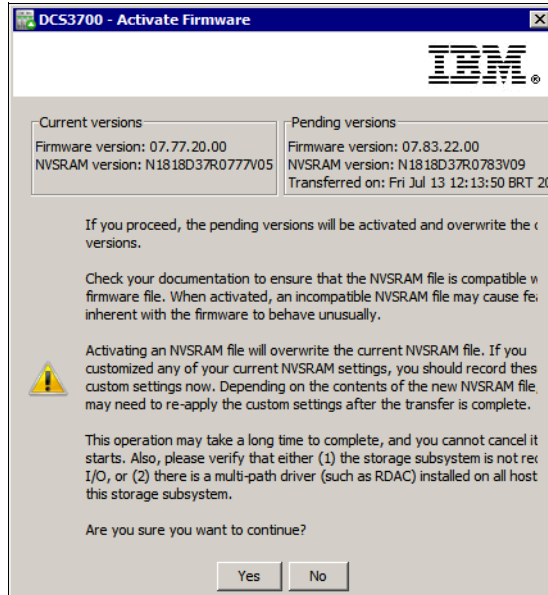


Figure 14-37 Activate Firmware warning window

Click **Yes** to proceed with the activation. The status is reported, as shown in Figure 14-38.



Figure 14-38 Activating the firmware

9. After the firmware upgrade completes successfully (Figure 14-39), click **Close**.



Figure 14-39 Download successful

**Note:** If the firmware version is going from Version 07.77.20.00 to 07.83.22.00, as it is in this example, you might see the warning show in Figure 14-39 prompting you to close and reopen Storage Manager, even if you are running the latest version. This step is needed to reload all the new features that are described in this book.

## 14.4 Controller NVSRAM

To install the controller NVSRAM, complete the following steps.

**Note:** To access the Download Controller NVSRAM function, click **Upgrade** → **Controller NVSRAM**.

1. To download a new NVSRAM file to your storage subsystem, open the associated window (Figure 14-40). Click **Select File** to select the NVSRAM file.

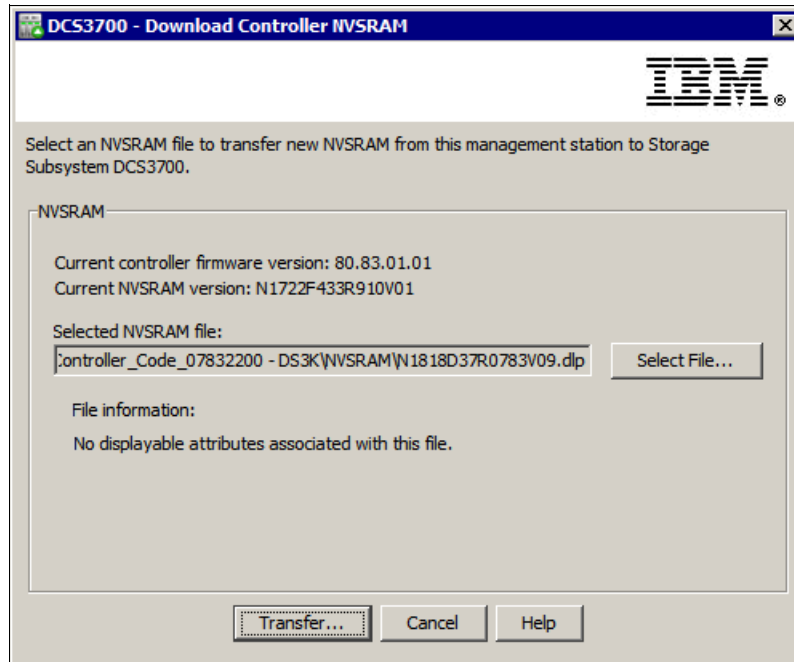


Figure 14-40 Download Controller NVSRAM

2. Click **Select File** to select the NVSRAM file, as shown in Figure 14-40.
3. After you select the NVSRAM file that you want to install, click **Transfer...** to begin installing the new NVSRAM file on the storage subsystem. The Confirm Download window opens (Figure 14-41). Read the information and click **Yes** to start the NVSRAM download.

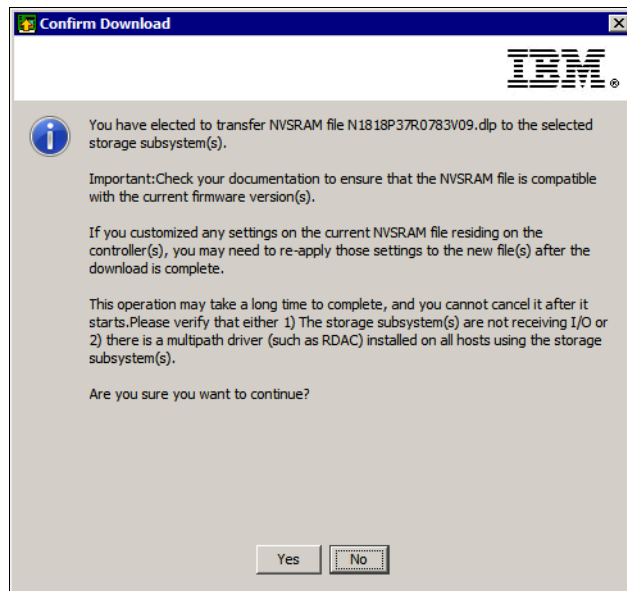


Figure 14-41 Confirm NVSRAM download

- The Downloading window opens and shows you the status of the transfer and the activation on the controllers, as shown in Figure 14-42.

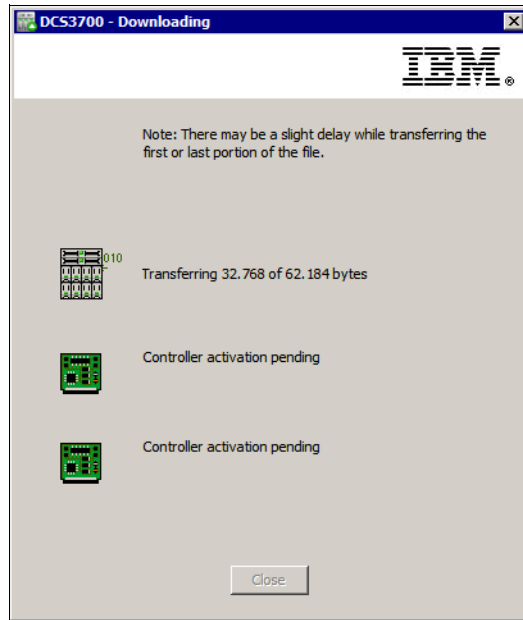


Figure 14-42 Download progress of NVSRAM

- After the NVSRAM file is transferred and activated successfully, click **Close**.

## 14.5 Drive firmware

The hard disk drive firmware upgrade process is disruptive and requires that all host I/Os be suspended during the update. Plan a production outage in advance and complete the following steps (when I/O operations are suspended) to perform the drive firmware upgrade.

**Note:** You can access the Download Drive Firmware function by clicking **Upgrade** → **Drive Firmware**.

1. With the Download Drive Firmware wizard open, read the information that is displayed on the introduction window (Figure 14-43) and then click **Next**.

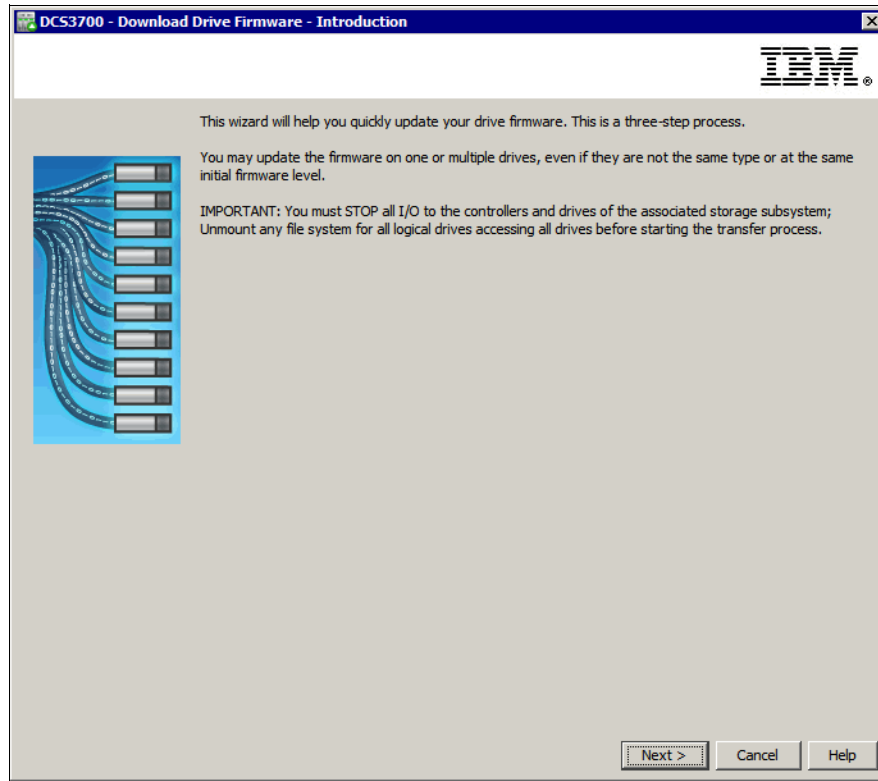


Figure 14-43 Download Drive Firmware - Introduction window

2. The Download Drive Firmware - Add Packages window opens (Figure 14-44). This window has two main sections:
  - The *Current package* section contains a list of all the drive models and the firmware levels that are 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.

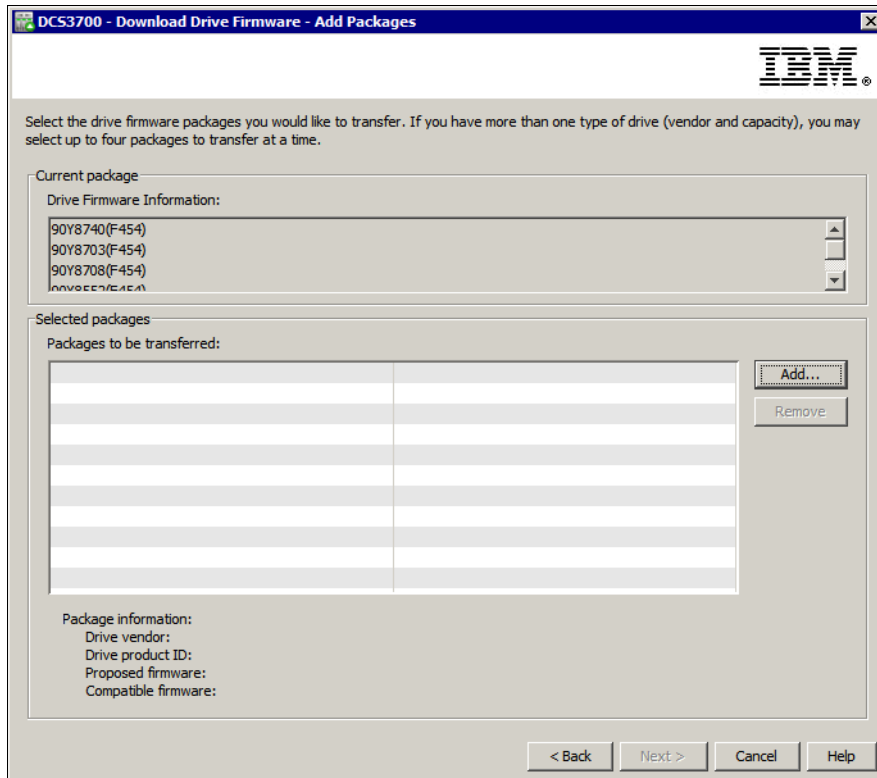


Figure 14-44 Select drive firmware package

3. Click **Add** to open the Select A Package window and browse to the folder that contains the drive package files.

As you highlight each package that is listed, the File Information section reports if the selected drive firmware package is Compatible (Figure 14-45) or Not Compatible (Figure 14-46) with the types of drives that are installed in your system.




Figure 14-45 Drive firmware is compatible




Figure 14-46 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 to continue.
5. Now you must select the drives in the subsystem you want to update. You see only the compatible drives in this window. Check **Select all** to update all compatible drives.
6. The drives that are not compatible with the selected updates are shown in the Incompatible Drives tab (Figure 14-47). These are other drive types, and you will have to use other firmware packages.

Vendor	Product ID	Enclosure	Slot	Device ...	Current ...	Propose...	Status	Array	Mode
IBM-ESXS	MBE214...	85	12	SAS	SC10		Optimal	SAS_HS2...	Assigned
IBM-ESXS	MBE214...	85	13	SAS	SC10		Optimal	SAS_HS2...	Assigned
IBM-ESXS	MBE214...	85	4	SAS	SC10		Optimal	SAS_HS2...	Assigned
IBM-ESXS	MBE214...	85	9	SAS	SC10		Optimal	SAS_HS2...	Assigned
IBM-ESXS	MBE214...	85	7	SAS	SC10		Optimal	SAS_HS2...	Assigned
IBM-ESXS	MBE214...	85	18	SAS	SC10		Optimal	SAS_HS2...	Assigned
IBM-ESXS	MBE214...	85	15	SAS	SC10		Optimal	SAS_HS2...	Assigned
IBM-ESXS	ST95004...	85	2	SAS	BD10		Optimal		Unassigned
IBM-ESXS	ST95004...	85	5	SAS	BD10		Optimal		Unassigned
IBM-ESXS	ST95004...	85	3	SAS	BD10		Optimal		Unassigned
IBM-ESXS	ST91467...	85	23	SAS	E621		Optimal		Unassigned
IBM-ESXS	ST91467...	85	19	SAS	E621		Optimal		Unassigned

Figure 14-47 Incompatible drives tab

7. Click **Finish** to update the drive firmware. A confirmation window opens. Enter yes and click **OK** to start the update.



- After the update completes, a summary window opens (Figure 14-48). If you want to update more drives, click **Transfer More** to select the compatible packages. Otherwise, click **Close** to exit the drive update window.

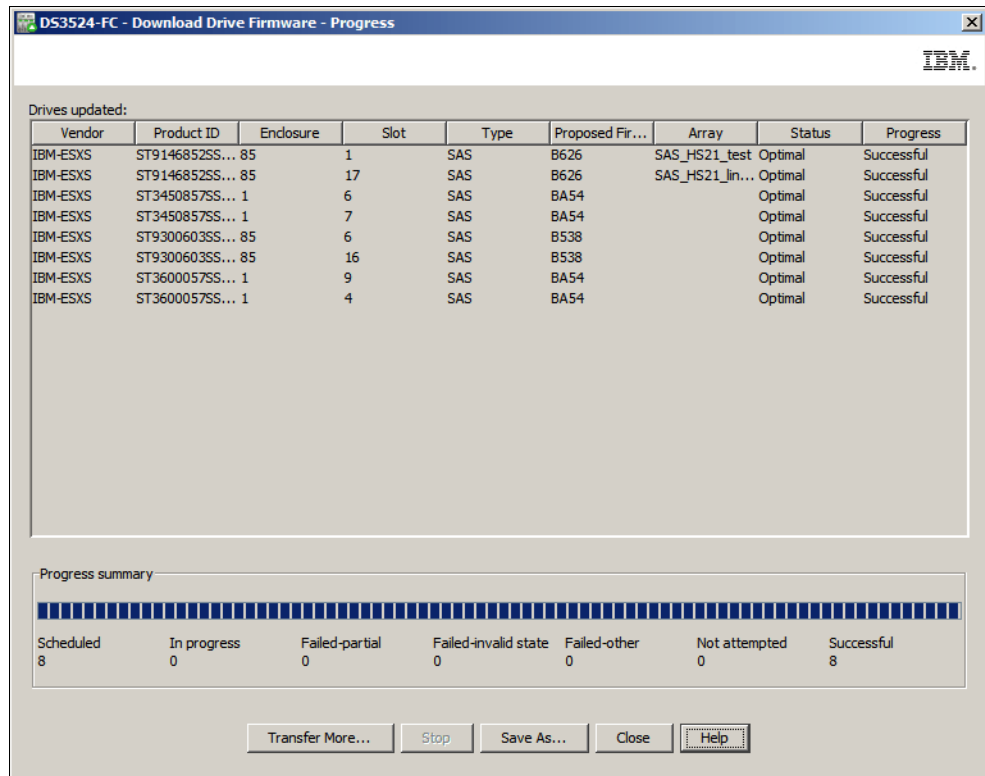


Figure 14-48 Drive update complete

## 14.6 ESM firmware

To install the ESM firmware for the EXP3500 enclosures, complete the following steps:

1. To update the ESM code, open the Download Environmental (ESM) Card Firmware window by clicking **Upgrade** → **ESM Firmware**. Figure 14-49 shows the window.

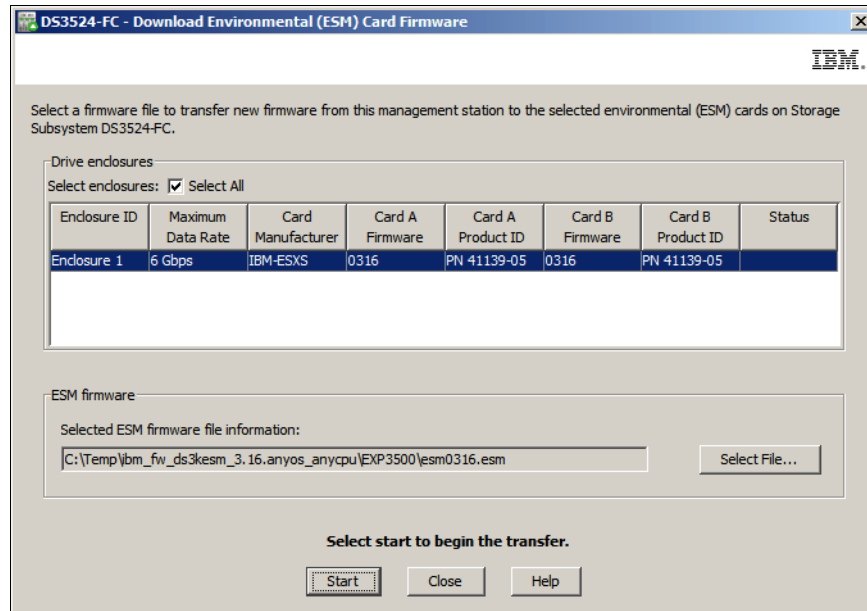


Figure 14-49 Download Environmental (ESM) Card Firmware

2. There are two sections in this view. The top section, *Drive enclosures*, lists all the installed enclosures. You can either select which enclosures to update or select the **Select All** check box to update all the connected enclosures.

In the second section, *ESM firmware*, select the ESM firmware file that you downloaded previously 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 opens. Read the information, confirm it by entering yes, and click **OK** to start the update (Figure 14-50).




Figure 14-50 Confirm Download

- 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 14-51.




Figure 14-51 Transferring the ESM firmware

- The status changes to Completed when the ESM firmware is downloaded, as shown in Figure 14-52.




Figure 14-52 ESM firmware transfer completed

## 14.7 Enclosure configuration settings

To install new ESM configuration settings from a file, complete the following steps.

**Note:** You can access the Download ESM configuration settings function by clicking **Upgrade** → **Enclosure Configuration Settings**.

1. To download ESM configuration settings, select the **Enclosure Configuration Settings** option, as shown in Figure 14-31 on page 448. The window that is shown in Figure 14-53 opens.

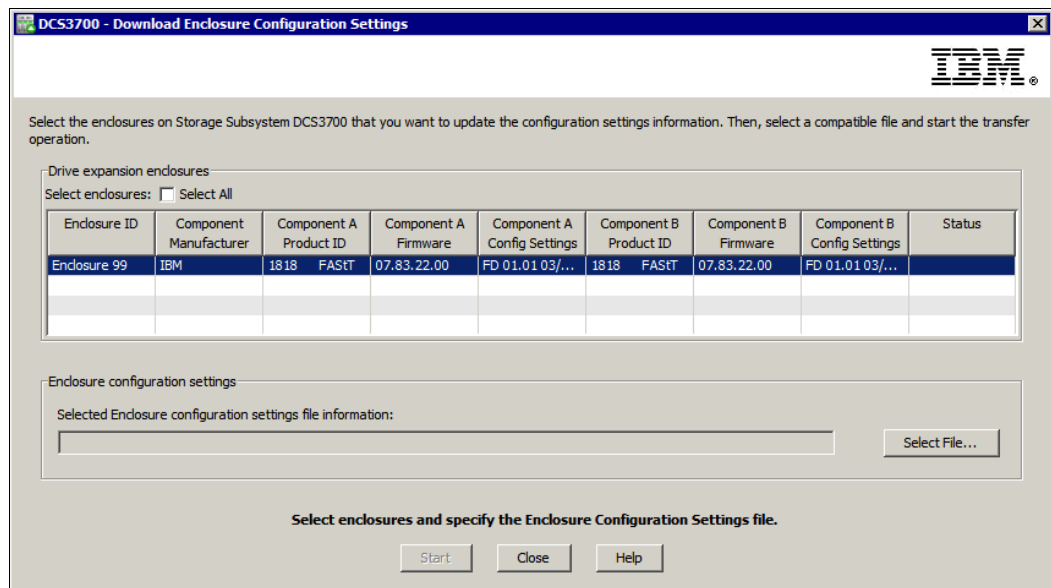


Figure 14-53 Download ESM card configuration

2. If there is more than one drive enclosure that is connected to the DCS3700 storage subsystem, 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, select only that enclosure in the list of enclosures.
3. Click **Select File** and specify the ESM configuration file that you want to transfer.
4. Click **Start**. You are prompted to confirm that you want to proceed. Enter yes in the text box and click **OK** to continue.



## Command-line interface (CLI)

This chapter explains the basics of the command-line interface (CLI) that can be used to manage an IBM System Storage DCS3700 storage subsystem. Along with the DCS3700 storage subsystem, the CLI can manage other models, such as DS3000, DS4000, and DS5000 storage subsystems.

This chapter is organized in to the following sections:

- ▶ Introduction to the command-line interface
- ▶ How to use the command-line interface
- ▶ Running the command-line interface
- ▶ General SMcli syntax
- ▶ Adding a storage subsystem to the Storage Manager configuration
- ▶ Showing the defined subsystems in the Storage Manager configuration
- ▶ Configuring alerts
- ▶ Issuing commands to the storage subsystem

The official guide for the CLI for the DCS3700 can be found at the following URL:

<https://www.ibm.com/support/entry/myportal/docdisplay?ln docid=MIGR-5076792&brandin d=500028>

## 15.1 Introduction to the command-line interface

The CLI is a software tool that storage subsystem installers and engineers use to configure and monitor storage subsystems by 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.

To run the script commands either through the script window, which is started from the IBM DS Storage Manager client Enterprise window, or through the CLI by using the SMcli program, you must install the IBM DS Storage Manager client. 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, the subdirectory is either C:\Program Files\IBM\_DS\Client (for 32-bit) or C:\Program Files (x86)\IBM\_DS\Client (for 64-bit). In Linux, the subdirectory is /opt/IBM\_DS/client.

The CLI program is called *SMcli* in Windows and Linux. Commands that can be used on the CLI to administer your DCS3700 storage subsystem are identical to the commands that are used in the Script Editor.

**Note:** Not all commands that are listed in the CLI help apply to all models of the storage subsystem. It depends on storage subsystem model, firmware, and DS Manager version.

The CLI gives you access to all the functions that are 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 must perform operations more than once. For example, you can run script files when you want to install the same configuration on several storage subsystems. You can use the CLI 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 by using the CLI. Similarly, the CLI is required to save the configuration file.

## 15.2 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 the operation environment parameters.

A CLI command consists of these elements:

- ▶ The SMcli
- ▶ The storage subsystem identifier
- ▶ Parameters
- ▶ Script commands

A CLI command takes this form:

```
SMcli storageSubsystem parameters script-commands;
```

Where:

- ▶ **SMcli** starts the CLI.
- ▶ *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.

### 15.2.1 Usage notes

If you enter **SMcli** and a storage subsystem name but do not specify CLI parameters, script commands, or a script file, the CLI runs in interactive mode. You can use Interactive mode to 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 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.

## 15.3 Running the command-line interface

There are two ways to enter commands into the DCS3700:

- ▶ Interactively by using the **SMcli**
- ▶ Creating saved scripts with the Scripting Editor.

Section 15.4, “General **SMcli** syntax” on page 468 describes **SMcli** .

In the following examples, we use a 64-bit version of Windows 7 for running IBM DS Manager and the client utilities.

### 15.3.1 Script Editor

You can use the Script Editor to create and save the files of commands from the CLI so you can batch run or schedule them.

To create a file by using the Script Editor, complete the following steps:

1. To start the Script Editor, from the Storage Manager window, click **Tools** → **Execute Script**, as shown in Figure 15-1.

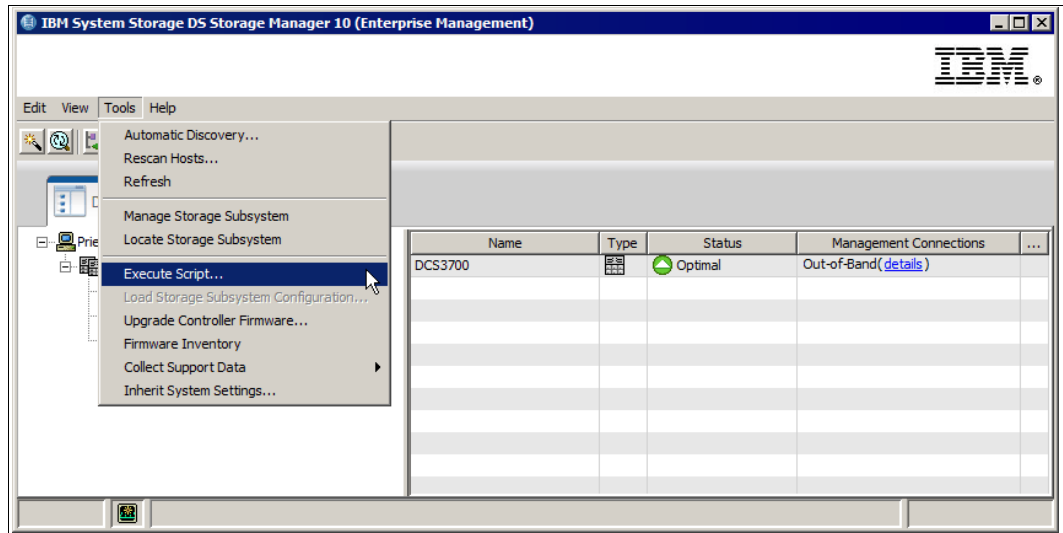


Figure 15-1 Script Editor

2. A blank Script Editor window opens, as shown in Figure 15-2. From here, you can enter commands and save them for later execution.

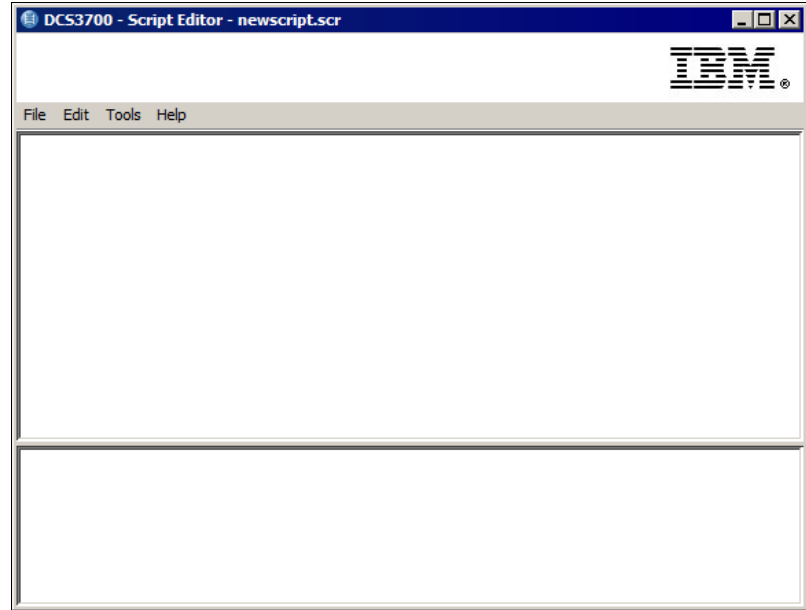


Figure 15-2 Script Editor window



- To save and retrieve script files that you create, click **File** → **Save Script** and **File** → **Load Script**. The Tools menu includes options to verify the syntax and run the scripts, as shown in Figure 15-3.

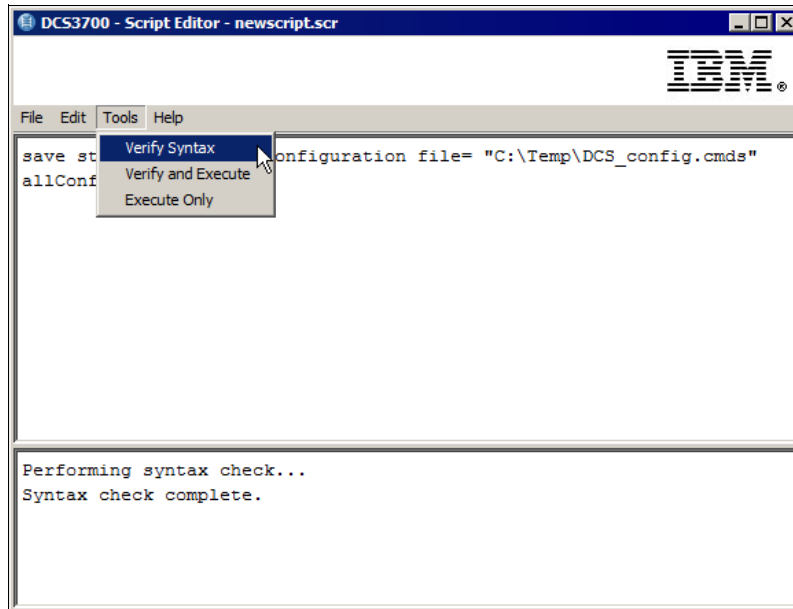


Figure 15-3 Tools menu - Verify Syntax

- The Script Editor contains an online help function that describes the usage of each of the available commands. To display the online help (Figure 15-4), click **Help** → **Command Reference**, or for help about using the Script Editor itself, click **Help** → **Overview**.

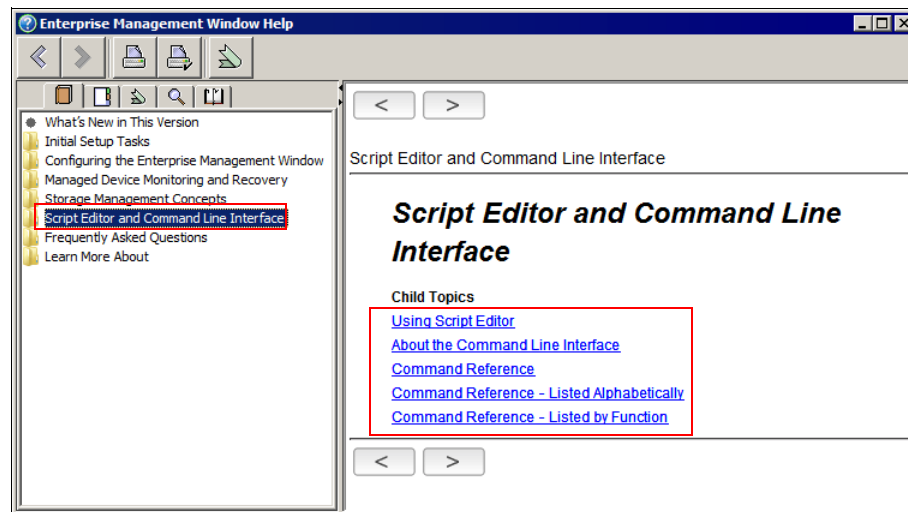


Figure 15-4 Command Reference - Help

## 15.4 General SMcli syntax

Table 15-1 lists the conventions that are used in the command statement syntax to manage the storage subsystem

Table 15-1 Command name syntax conventions

Convention	Definition
<b>(a   b   c)</b>	Alternative (“a” or “b” or “c”)
italicized-words	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 run SMcli from any window.

To list the full SMcli syntax, run **SMcli -?**, as shown in Example 15-1.

Example 15-1 SMcli syntax

```
C:\Program Files (x86)\IBM_DS\client>SMcli.exe -?
SMcli enable forceRestriction
SMcli disable forceRestriction
SMcli <DNS-network-name-or-IP-address>
    [<DNS-network-name-or-IP-address>]
    [-c "<command>; [<command2>;...]" ]
    [-n <storage-system-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-system-name> | -w <WWID>]
    [-o <outputfile>] [-p <password>] [-e] [-S] [-quick]
SMcli {-n <storage-system-name> | -w <WWID>}
    [-c "<command>; [<command2>;...]" ]
    [-o <outputfile>] [-p <password>] [-e] [-S] [-quick]
SMcli {-n <storage-system-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-system-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-system-name> | -w WWID | -h <hostName>]
    [-S]
SMcli -a email:<email address>
    [<hostname or IP address1> [<hostname or IP address2>]]
    [-n <storage-system-name> | -w <WWID> | -h <hostName>]]
    [-I <informationToInclude>] [-q <frequency>] [-S]
SMcli {-a | -x} trap:<community>,<hostname or IP address>
    [<hostname or IP address1> [<hostname or IP address2>]]
```

```

        [-n <storage-system-name> | -w <WWID> | -h <hostName>]
        [-S]
SMcli -alertSeverities [<severity>,<severity2>...]
SMcli -alertTest
SMcli -supportBundle auto (enable | disable) (all | storage-system-name) [data=pa
ath-name]
SMcli -supportBundle (auto | schedule) show
SMcli -supportBundle schedule (enable|disable) (all|storage-array-name)
        [data=path-name] [startTime=HH:MM] [startDate=MM:DD:YYYY] [endDate=MM:DD
:YYYY]
        (daysOfWeek=[Sunday Monday Tuesday Wednesday Thursday Friday Saturday] |
        months=[January February March April May June July August September Octo
ber November December] (onDays=[1-31] |
        weekNumber=(First|Second|Third|Fourth|Last) dayOfWeek=(Sunday|Mo
nday|Tuesday|Wednesday|Thursday|Friday|Saturday)))
SMcli -?

```

For more information, refer to your CLI documentation or use `man SMcli` or `help SMcli` depending on your operating system.

SMcli completed successfully.

---

The various syntax options that are shown by `SMcli -?` perform the following tasks:

- ▶ Run commands on a storage subsystem that are specified on the command line over an out-of-band management connection.
- ▶ Run commands that are specified in a command script over an out-of-band management connection.
- ▶ Run commands on a storage subsystem that are specified on the command line over an in-band management connection.
- ▶ Run 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 a defined email alert.
- ▶ Configure email alerts.
- ▶ Configure or remove SNMP alerts.

The CLI can perform all the functions that are 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 DCS3700 controllers. As a preferred practice, assign a new name after the controllers are added to the DS Manager.

Table 15-2 gives an explanation of the various parameters that can be specified with SMcli commands.

Table 15-2 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. IP address or fully qualified domain name of the host running the Storage Manager agent for in-band management. Use either IP addresses or host names, the subsystem name, or WWID.
-A -A <IP C1> <IP C2>	Use this parameter to add a storage array to the configuration files. If you do not follow the -A parameter with a host name or IP address, auto-discovery scans the local subnet for storage arrays.
-a	Use this parameter to add a Simple Network Management Protocol (SNMP) trap destination or an email address alert destination. <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 must 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: storage-system-name host-name1 host-name2
-e	Runs the commands without performing a syntax check first.
-F <email addr>	The email address that is used in the sent by field of an email alert. Responses to an alert are sent to this address. 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 is included in all email alert notifications.
-h	Specifies the host name that is running the SNMP agent to which the storage subsystem is connected.
-I eventOnly -I profile -I supportBundle	Define what information is included in an alert that gets sent through email. Use one of these options: <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 be used with only parameter -d. The file contents takes this format: storage-system-name IP-address1 IPaddress2
-m <mail server>	Specifies the host name or the IP address of the email server from which email alert notifications are sent.

Option	Description
<b>-n &lt;subsystem name&gt;</b>	The name of the storage subsystem that is shown by the <b>-d</b> options. Use either an IP or host names, subsystem name, or WWID.
<b>-o &lt;output file&gt;</b>	Specify the file that contains the output. Errors are sent to the standard error output, which is usually the console. Use redirection when required. The output is not appended to the output file.
<b>-p &lt;password&gt;</b>	The password of the storage subsystem to perform management tasks. This is an optional option that is only required when a password is defined for the specified storage subsystem. By default, no password is defined and this option is not required.
<b>-q</b>	Specifies the frequency that you want to include more 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 <b>-I</b> terminal to <b>eventOnly</b> , the only valid value for the <b>-q</b> terminal is every event. If you set the <b>-I</b> terminal to either the profile value or the <b>supportBundle</b> value, this information is included with the emails with the frequency specified by the <b>-q</b> terminal.
<b>-r &lt;direct_sa   host_sa&gt;</b>	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 <b>direct_sa</b> value. For storage subsystems that are managed in-band, use the <b>host_sa</b> value.
<b>-S (uppercase)</b>	Suppresses informational messages that describes the command progress that appears when you run script commands.
<b>-s (lowercase)</b>	Shows the alert settings in the configuration file when used with the <b>-d</b> terminal.
<b>-v</b>	Show the current global status of the known devices in a configuration file when used with the <b>-d</b> terminal.
<b>-w &lt;WWID&gt;</b>	Specifies the WWID of the storage subsystem. This terminal is an alternative to the <b>-n</b> terminal. Use the <b>-w</b> terminal with the <b>-d</b> terminal to show the WWIDs of the known storage subsystems. The file content takes this format: storage-system-name world-wide-ID IPaddress1 IP-address2
<b>-w</b>	This option should be used only with the <b>-d</b> parameter. It shows the WWID of the defined storage subsystems.
<b>-X (uppercase)</b>	Deletes a storage subsystem from a configuration.
<b>-x (lowercase)</b>	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.
<b>-quick</b>	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 during the command. Do not use this parameter for operations involving more than one single-line operation. Extensive use of this command can overrun the controller with more commands than the controller can process, which causes operational failure. Also, status and configuration updates that are normally collected from background processes are not available to the CLI. This parameter causes operations that depend on background information to fail.
<b>-?</b>	Shows usage information about the CLI commands.

We now show specific examples of using the CLI for tasks that were demonstrated by using the GUI.

## 15.5 Adding a storage subsystem to the Storage Manager configuration

Adding storage subsystems to the Storage Manager configuration can be done manually or by automated discovery as follows:

1. Upon installation, Storage Manager has no subsystems that are defined. Use option **-d** of the SMcli to list the defined subsystems. For a detailed description of the command that is shown in Example 15-2, see 15.6, “Showing the defined subsystems in the Storage Manager configuration” on page 473.

*Example 15-2 SMcli - list storage subsystems that are defined in the Storage Manager configuration*

---

```
C:\Program Files (x86)\IBM_DS\client>SMcli -d
There are currently no storage systems listed in the configuration file. Add
storage systems using the Add Storage System option in the storage management
software or by command line.

SMcli
failed.
```

---

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 15-3 shows both methods.

*Example 15-3 SMcli - Add storage subsystems to the Storage Manager configuration*

---

```
C:\Program Files (x86)\IBM_DS\client>SMcli -A 9.42.171.23 9.42.171.24 -n DCS3700
New storage system was discovered at address 9.42.171.27.

New storage system was discovered at address 9.42.171.28.

SMcli completed successfully.

C:\Program Files (x86)\IBM_DS\client>SMcli -A
Starting auto discovery.
.....
.....
Auto discovery operation successful.
SMcli completed
successfully.
```

---

## 15.6 Showing the defined subsystems in the Storage Manager configuration

Option `-d` of the `SMcli` gives a list of defined storage subsystems. There are more parameters that can be used together with option `-d` to display more details about the defined storage subsystems:

- `[-i]` Show IP addresses instead of host names.
- `[-s]` Show alert settings.
- `[-w]` Show the worldwide identifier.
- `[-v]` Show the status of the defined subsystems in the configuration file.
- `[-S]` Do not list progress information.

We use all these parameters in Example 15-4. The output shows that there is a global alert in place that sends email alerts to email address `dsc3700@rivers.local`. The mail server `nile.rivers.local` is used and a return address `dsc3700@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 DCS3700 storage subsystems.

SNMP alerts are shown for the DCS3700 storage subsystems. `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 WWID, host names, or IP addresses of the controllers and the status of the subsystem. This information also can be found in the Enterprise Management window of the Storage Manager GUI.

*Example 15-4 Displaying details about the defined storage subsystems*

---

```
C:\Program Files (x86)\IBM_DS\client>SMcli -d -i -s -w -v -S
[MAIL SERVER]
nile.rivers.local dcs3700@nile.rivers.local

[ALERT SETTINGS - DEVICE GROUPS]

All storage subsystems
    dcs3700@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:

DCS3700 60080e50001f49c00000000050114e2b          9.42.171.27      9.42.171.28
```

---

## 15.7 Configuring alerts

This section describes how to manage alert recipients.

### 15.7.1 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, run the following command:

```
SMcli -m <IP or host name of mail server> -F <email address>
```

For more information, see Example 15-5. To verify the setting, run `SMcli -d -s` to see the current settings.

*Example 15-5 SMcli - define email server and return email address*

---

```
C:\Program Files (x86)\IBM_DS\client>SMcli -m nile.rivers.local -F  
DCS3700@nile.rivers.local  
SMcli completed successfully.
```

---

To delete these settings, use empty strings, as shown in Example 15-6.

*Example 15-6 SMcli - delete mail server and return email address setting*

---

```
C:\Program Files (x86)\IBM_DS\client>SMcli -m "" -F ""  
SMcli completed successfully.
```

---

### 15.7.2 Defining email alert recipients

Alerts can be defined for a single DCS3700 storage subsystem, all out-of-band managed storage subsystems, all in-band managed storage subsystems, or all subsystems that are managed by a single management station. This section describes those settings.

**Note:** An email server should be defined as explained in 15.7.1, “Defining the mail server and email address to send out the email alerts” on page 474. Without a mail server, the local system running the Storage Manager is used to send email alerts, which might not always be possible.

Example 15-7 shows how to define an email alert recipient for a single DCS3700 storage subsystem.

*Example 15-7 SMcli - define email alert recipient for single system*

---

```
C:\Program Files (x86)\IBM_DS\client>SMcli -a email:dcs3700@nile.rivers.local -n "DCS3700"  
SMcli completed successfully.
```

---

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 15-8 shows an email recipient for all in-band managed storage subsystems.

*Example 15-8 SMcli - email recipient for a group of storage subsystems*

---

```
C:\Program Files (x86)\IBM_DS\client>SMcli -a email:dcs3700-inband@nile.rivers.local
SMcli completed successfully.
```

---

To send alerts that are generated by any single subsystem that is defined in the configuration file of Storage Manager, run the command that is shown in Example 15-9.

*Example 15-9 SMcli - email alert recipient for all subsystems*

---

```
C:\Program Files (x86)\IBM_DS\client>SMcli -a email:dcs3700@nile.rivers.local
SMcli completed successfully.
```

---

In older versions, global email alert recipients can be configured by omitting the target specification, such as `-r host_sa` or a host name. Example 15-10 shows a global email recipient configuration.

*Example 15-10 SMcli - configure global email recipient*

---

```
C:\Program Files (x86)\IBM_DS\client>SMcli -a email:dcs3700@nile.rivers.local
SMcli completed successfully.
```

---

### 15.7.3 Deleting email alert recipients

To delete an email alert recipient, use the `-x` option and the same syntax that was used to add the email alert recipient. Example 15-11 shows the commands to remove email recipients from the configuration. The first command shows an example of a configured management station system and all the following commands 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 then the global email recipient.

*Example 15-11 SMcli - delete email recipients*

---

```
C:\Program Files (x86)\IBM_DS\client>SMcli -d -s
[MAIL SERVER]
nile.rivers.local dcs3700@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:
DCS3700 9.42.171.27    9.42.171.28 None

SMcli completed successfully.
```

```
C:\Program Files (x86)\IBM_DS\client>SMcli -x email:dc3700@nile.rivers.local -n DCS3700
SMcli completed successfully.
```

```
C:\Program Files (x86)\IBM_DS\client>SMcli -x email:in-band@nile.rivers.local
SMcli completed successfully.
```

```
C:\Program Files (x86)\IBM_DS\client>SMcli -x email:out-band@nile.rivers.local
SMcli completed successfully.
```

```
C:\Program Files (x86)\IBM_DS\client>SMcli -x email:host@nile.rivers.local -h
colorado.rivers.local
SMcli completed successfully.
```

```
C:\Program Files (x86)\IBM_DS\client>SMcli -x email:global@nile.rivers.local
SMcli completed successfully.
```

---

## 15.7.4 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 15-12 shows how to configure SNMP trap recipients for a single DCS3700 storage subsystem, in-band and out-of-band managed DCS3700 storage subsystems, and a global trap recipient.

### *Example 15-12 SMcli - SNMP alert recipient*

---

```
C:\Program Files (x86)\IBM_DS\client>SMcli -a trap:dc3700,nile.rivers.local -n dcs3700
SMcli completed successfully.
```

```
C:\Program Files (x86)\IBM_DS\client>SMcli -a trap:in-band,nile.rivers.local
SMcli completed successfully.
```

```
C:\Program Files (x86)\IBM_DS\client>SMcli -a trap:out-band,nile.rivers.local
SMcli completed successfully.
```

```
C:\Program Files (x86)\IBM_DS\client>SMcli -a trap:global,nile.rivers.local
SMcli completed successfully.
```

---

## 15.8 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]
```

For information about the meaning of each option, see 15.4, “General SMcli syntax” on page 468.

The **-c** parameter requires that all commands be enclosed in single or double quotation marks. Each command must be terminated by a semicolon. Multiple commands can follow on one command line.

Example 15-13 shows a command to a storage subsystem to get the health status of that storage subsystem.

*Example 15-13 SMcli - command*

---

```
C:\Program Files (x86)\IBM_DS\client>SMcli -n DCS3700 -c "show storagesubsystem
healthstatus;" -S
Storage Subsystem health status = optimal
```

---

An alternative to specifying multiple commands in one line is to create a script file that has each command on a single line. You can do this task by using the Script Editor (which you can use to validate and run 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 15-14 shows a script file that creates logical drives and assigns them to a host.

*Example 15-14 Example of a command script file*

---

```
create array drives=(99,5,1 99,1,4 99,2,1) RAIDLevel=5 userLabel="Array01_DevTest"
securityType=capable T10PI=none;
create logicalDrive array="Array01_DevTest" userLabel="SERVER03_LUN01" owner=A
capacity=107374182400 Bytes segmentSize=128 dssPreAllocate=false T10PI=enabled
mapping=none;
create logicalDrive array="Array01_DevTest" userLabel="SERVER03_LUN02" owner=B
capacity=107374182400 Bytes segmentSize=128 dssPreAllocate=false T10PI=enabled
mapping=none;
```

---

To run this script file (outside of Script Editor), save it (in our example, as createDrive.cmd) and run it as shown in Example 15-15.

*Example 15-15 SMcli runs the script*

---

```
C:\Program Files (x86)\IBM_DS\client>SMcli -n DCS3700 createDrive.cmd -o hudson.out -S
```

---

Certain commands require the input of string values, such as the command that is shown in Example 15-16.

*Example 15-16 SMcli - character masking*

---

```
C:\Program Files (x86)\IBM_DS\client>SMcli -n DCS3700 -c "set logicalDrive
[\"SERVER03_LUN01\"] userLabel=\\\"SERVER03_LUN01_Renamed\";" -S -e
```

---

This command renames the logical drive SERVER03\_LUN01 to SERVER03\_LUN01\_Renamed. The parameter **logicalDrive** requires the name of a logical drive, which is 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 information about character masking. An alternative that is not available for Microsoft Windows is the usage 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 DCS3700 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. To get details about each command, see the Script Editor online help or the CLI online help.

### 15.8.1 Sample command: Save configuration script file

To create a script that contains the configuration statements for your environment, run **save StorageSubsystem configuration**. A sample of this command is shown in Example 15-17.

*Example 15-17 CLI command to save the storage subsystem configuration*

---

```
save StorageSubsystem configuration file="DCS3700_Config.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 15-18.

*Example 15-18 Sample storage subsystem configuration file*

---

```
// Logical configuration information from Storage Subsystem DCS3700.  
// Saved on 27 de julio de 2012  
// Firmware package version for Storage Subsystem DCS3700 = 07.83.18.00  
// NVSRAM package version for Storage Subsystem DCS3700 = N1818D37R0783V09  
  
//on error stop;  
  
// Uncomment the two lines below to delete the existing configuration.  
//show "Deleting the existing configuration."  
//clear storagesubsystem configuration;  
  
// Storage Subsystem global logical configuration script commands
```

```

show "Setting the Storage Subsystem user label to DCS3700.";
set storagesubsystem userLabel="DCS3700";

show "Setting the Storage Subsystem media scan rate to 30.";
set storagesubsystem mediaScanRate=30;

// Uncomment the three lines below to remove the default logical drive (if exists). NOTE:
Default logical drive name is always = "Unnamed".
//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.
show "Setting the Storage Subsystem cache block size to 8.";
set storagesubsystem cacheBlockSize=8;

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 Array Array01_DevTest, RAID 5.";
//This command creates array <Array01_DevTest>.
create array drives=(99,5,1 99,1,4 99,2,1) RAIDLevel=5 userLabel="Array01_DevTest"
securityType=capable T10PI=none;

show "Creating logical drive SERVER03_LUN01 on array Array01_DevTest.";
//This command creates logical drive <SERVER03_LUN01> on array <Array01_DevTest>.
create logicalDrive array="Array01_DevTest" userLabel="SERVER03_LUN01" owner=A
capacity=107374182400 Bytes segmentSize=128 dssPreAllocate=false T10PI=enabled
mapping=none;
show "Setting additional attributes for logical drive SERVER03_LUN01.";
// Configuration settings that can not be set during Logical Drive creation.
set logicalDrive["SERVER03_LUN01"] cacheFlushModifier=10;
set logicalDrive["SERVER03_LUN01"] cacheWithoutBatteryEnabled=false;
set logicalDrive["SERVER03_LUN01"] mirrorEnabled=true;
set logicalDrive["SERVER03_LUN01"] readCacheEnabled=true;
set logicalDrive["SERVER03_LUN01"] writeCacheEnabled=true;
set logicalDrive["SERVER03_LUN01"] mediaScanEnabled=true;
set logicalDrive["SERVER03_LUN01"] redundancyCheckEnabled=false;
set logicalDrive["SERVER03_LUN01"] cacheReadPrefetch=true;
set logicalDrive["SERVER03_LUN01"] modificationPriority=high;
set logicalDrive["SERVER03_LUN01"] preReadRedundancyCheck=false;

show "Creating logical drive SERVER03_LUN02 on array Array01_DevTest.";
//This command creates logical drive <SERVER03_LUN02> on array <Array01_DevTest>.
create logicalDrive array="Array01_DevTest" userLabel="SERVER03_LUN02" owner=B
capacity=107374182400 Bytes segmentSize=128 dssPreAllocate=false T10PI=enabled
mapping=none;
show "Setting additional attributes for logical drive SERVER03_LUN02.";
// Configuration settings that can not be set during Logical Drive creation.

```

```
set logicalDrive["SERVER03_LUN02"] cacheFlushModifier=10;
set logicalDrive["SERVER03_LUN02"] cacheWithoutBatteryEnabled=false;
set logicalDrive["SERVER03_LUN02"] mirrorEnabled=true;
set logicalDrive["SERVER03_LUN02"] readCacheEnabled=true;
set logicalDrive["SERVER03_LUN02"] writeCacheEnabled=true;
set logicalDrive["SERVER03_LUN02"] mediaScanEnabled=true;
set logicalDrive["SERVER03_LUN02"] redundancyCheckEnabled=false;
set logicalDrive["SERVER03_LUN02"] cacheReadPrefetch=true;
set logicalDrive["SERVER03_LUN02"] modificationPriority=high;
set logicalDrive["SERVER03_LUN02"] preReadRedundancyCheck=false;
// Creating Logical Drive-To-LUN Mappings
```

---

We show how to load this file to create the configuration on another system in “Load Storage Subsystem Configuration option” on page 165. Loading the configuration file overwrites any existing data on the array. You should do this only on a new array.

## 15.9 More information

For more information about all the CLI parameters, consult the Command Line reference included in the SMclient online help or *Command Line Interface and Script Commands Programming Guide - IBM System Storage DS3000, DS4000, and DS5000*, found at:

<https://www.ibm.com/support/entry/myportal/docdisplay?ln docid=MIGR-5076792&brandin d=5000028>



## Windows SAS configuration guide for IBM BladeCenter

This chapter describes and explains a sample configuration that shows how to connect logical drives that are configured on an IBM System Storage DCS3700 storage subsystem to a Windows Server 2008 operating system that is running on an IBM BladeCenter HS23 blade server that is connected with SAS Connectivity Modules.

For the Fibre Channel (FC) Windows host attachment guides, see the following documents:

- ▶ *IBM Midrange System Storage Implementation and Best Practices Guide*, SG24-6363
- ▶ *IBM System Storage DS Storage Manager Version 10 - Installation and Host Support Guide*, found at:

<http://www-947.ibm.com/support/entry/portal/docdisplay?brand=5000028&indocid=MI GR-5075652>

## 16.1 Required equipment

Use the following hardware and software components for this configuration:

- ▶ IBM BladeCenter H
- ▶ IBM BladeCenter HS23 with Windows Server 2008 installed on a local SAS disk
- ▶ IBM BladeCenter Advance Management Module (AMM) installed in a BladeCenter H
- ▶ Two IBM BladeCenter SAS Connectivity Modules
- ▶ IBM BladeCenter dual port SAS Expansion card (CIOv)
- ▶ A DCS3700 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 10.83)

Figure 16-1 shows our sample hardware setup for BladeCenter and DCS3700 connectivity.




Figure 16-1 Hardware setup - BladeCenter and DCS3700 connectivity



## 16.2 IBM BladeCenter setup

Check the IBM support website for the latest DCS3700 interoperability matrix to ensure that all the hardware that is used for this configuration is supported. Ensure that the BladeCenter, blade server, and modules are updated with the latest firmware. For the latest firmware, see the following IBM Support websites:

- ▶ For Storage Interoperability:  
<http://www.ibm.com/systems/support/storage/ssic/interoperability.wss>
- ▶ For BladeCenter Interoperability:  
<http://www.ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5073016>

### 16.2.1 Installing Windows Server 2008 R2

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 HS23 (7885) with Microsoft Windows Server 2008 R2 setup, use the instructions that are found at:

<http://www.ibm.com/support/entry/portal/docdisplay?ln docid=migr-5089571>

### 16.2.2 HS23 SAS Expansion Cards

You must install the IBM BladeCenter SAS Connectivity Card (CIOv) in the HS23 BladeCenter host before you 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 one of the following URLs:

- ▶ <http://www.ibm.com/systems/support/>
- ▶ [http://publib.boulder.ibm.com/infocenter/bladectr/documentation/topic/com.ibm.bladecenter.ec\\_ibm\\_sasconnectivity.doc/DW1HLMst.pdf](http://publib.boulder.ibm.com/infocenter/bladectr/documentation/topic/com.ibm.bladecenter.ec_ibm_sasconnectivity.doc/DW1HLMst.pdf)

**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.

### 16.2.3 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 DCS3700 storage subsystem:

1. Turn on or restart the HS23 host.

2. Press Ctrl+c to enter the LSI Logic Configuration Utility, as shown in Figure 16-2.

```

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 (2009.06.00)
Copyright 2000-2009 LSI Corporation.

Press Ctrl-C to start LSI Corp Configuration Utility...

```

Figure 16-2 SAS Expansion Card - LSI Logic Configuration Utility

3. The menu that is shown in Figure 16-3 opens after you press Ctrl+c.

```

LSI Corp Config Utility      v6.26.02.00 (2009.06.00)
Adapter List Global Properties
Adapter      PCI  PCI  PCI  PCI  FW Revision  Status  Boot
              Bus Dev Fnc Slot
SAS3020XD   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+1 = Help  
Alt+N = Global Properties    -/+ = Alter Boot Order    Ins/Del = Alter Boot List

Figure 16-3 SAS Expansion Card - LSI Logic Configuration Utility menu

- Press the Enter key to select the SAS adapter that is internally connected to the SAS Connectivity Module in I/O Module bay 3, as shown in Figure 16-4.

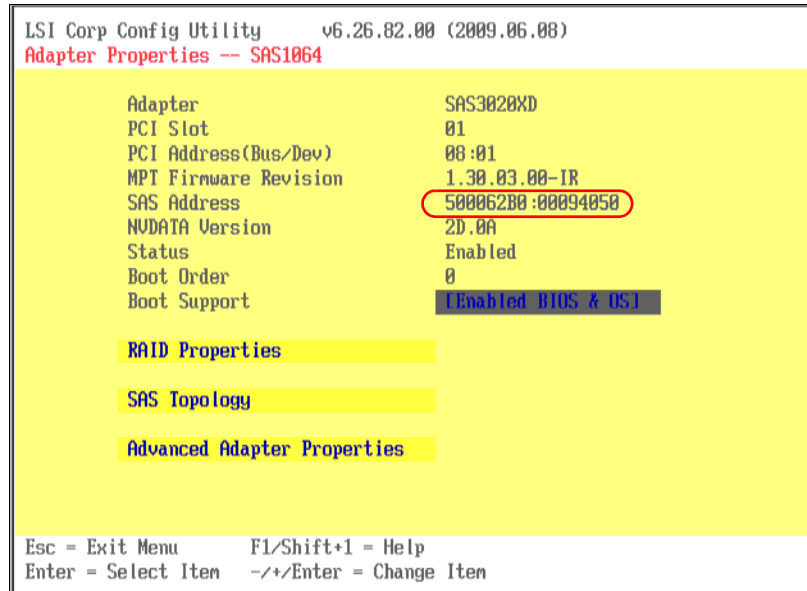


Figure 16-4 SAS Expansion Card - LSI Logic Config Utility adapter selected

- 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 DCS3700 storage subsystem. The WWPN can also be retrieved from the SAS Connectivity Module web interface.

The WWPN can also be recorded by clicking **Hardware VPD** → **Blade Bay** → **Expansion Card** ((in our example, the Blade Bay is Blade6). This action opens the BladeCenter Vital Product Data window. Click the **Ports** tab to record the WWPNs, as shown in Figure 16-5.

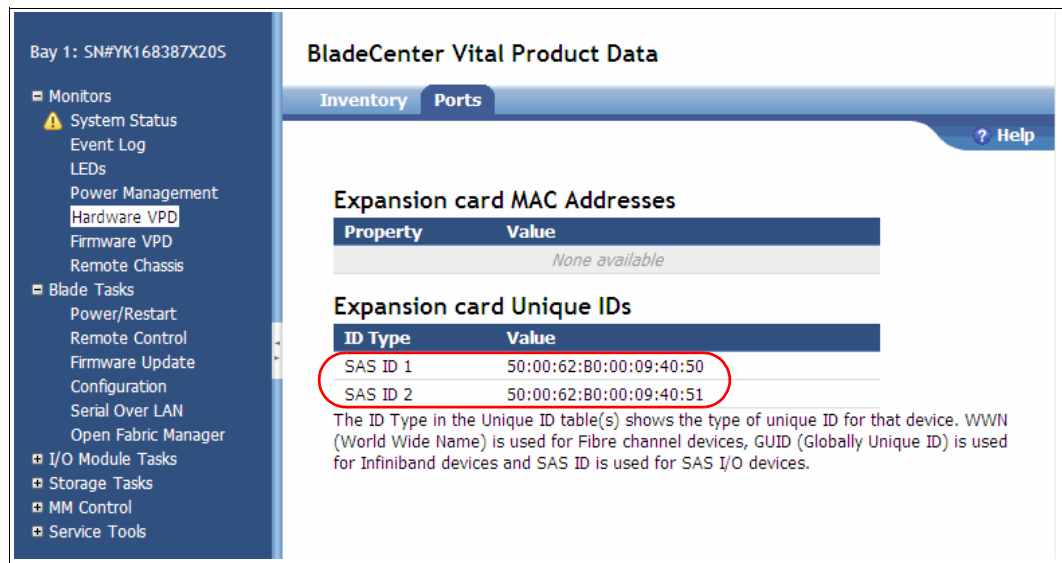


Figure 16-5 Expansion Card Unique IDs

- The name of the SAS adapter for the expansion card is SAS3020XD and is visible in the Adapter List window. To determine whether 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 16-6.

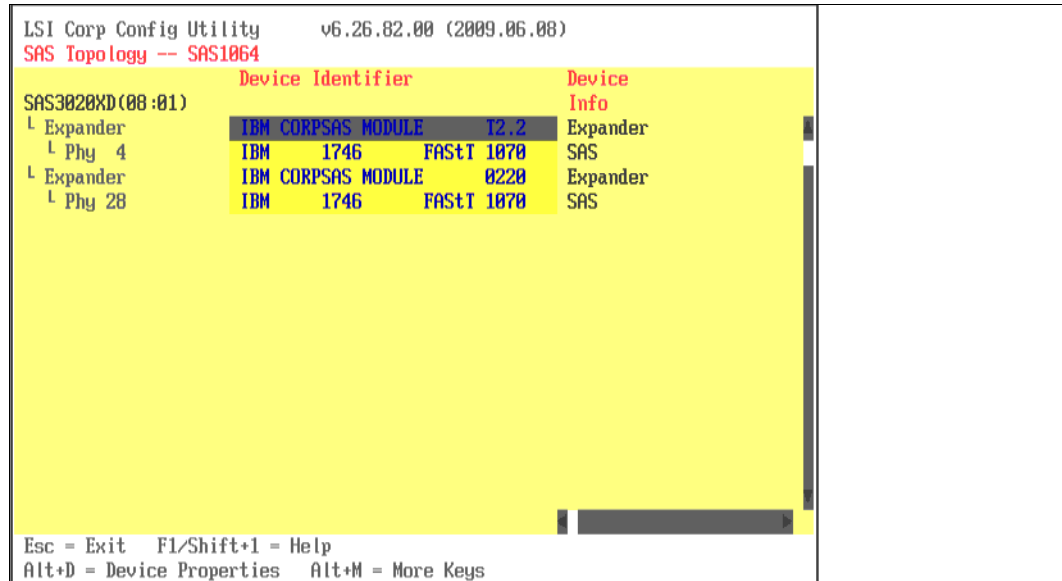


Figure 16-6 SAS Expansion Card - LSI Logic Config Utility adapter confirmation

## 16.2.4 HS23 SAS Expansion Card device driver

In this example, we use the SAS Expansion Card device driver that is part of the Windows Server 2008 R2 installation. The SAS Expansion Card is installed and Windows Server 2008 R2 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 R2 was available from following IBM website:

<http://www-933.ibm.com/support/fixcentral/systemx/selectFixes?parent=BladeCenter+HS23&product=ibm/systemx/7875&platform=Windows+2008+x64&function=all#Firmware%20Update>

## 16.2.5 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 E Type 8677
- ▶ BladeCenter HT Type 8740 / 8750
- ▶ BladeCenter H Type 7989/8852
- ▶ BladeCenter S Type 7779/8886

Installing a connectivity module in I/O module bay 3 or I/O module bay 4 provides connectivity to the SAS expansion cards installed in the blade servers in your BladeCenter unit. Installing two connectivity modules allows you to have two connections to the SAS expansion cards that are installed in the blade servers.

**Important:** The connectivity modules in I/O module bay 3 and I/O module bay 4 and all expansion cards in the BladeCenter unit must use the same interface type. Therefore, you must install SAS expansion cards before you install connectivity modules in the blade servers in your BladeCenter unit. For more information about the SAS expansion card, see the *Installation and User's Guide* for the SAS expansion card, found at:

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

Connect the cables from the DCS3700 storage subsystem controllers A and B to the external ports # 1 of the two SAS Connectivity Modules, as shown in Figure 16-1 on page 482.

## 16.2.6 SAS Connectivity Module firmware update

Ensure that your SAS Connectivity Module is updated with the latest firmware. For the latest firmware update, go to the following website:

<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 by using the web interface with the IP address that is defined for the connectivity module in the BladeCenter Advance Management Module (AMM), as shown in Figure 16-7. Use USERID as the user ID and PASSWORD as the password. (0 in PASSWORD means null). You can change the password under the Administer Users menu option after you are logged on.

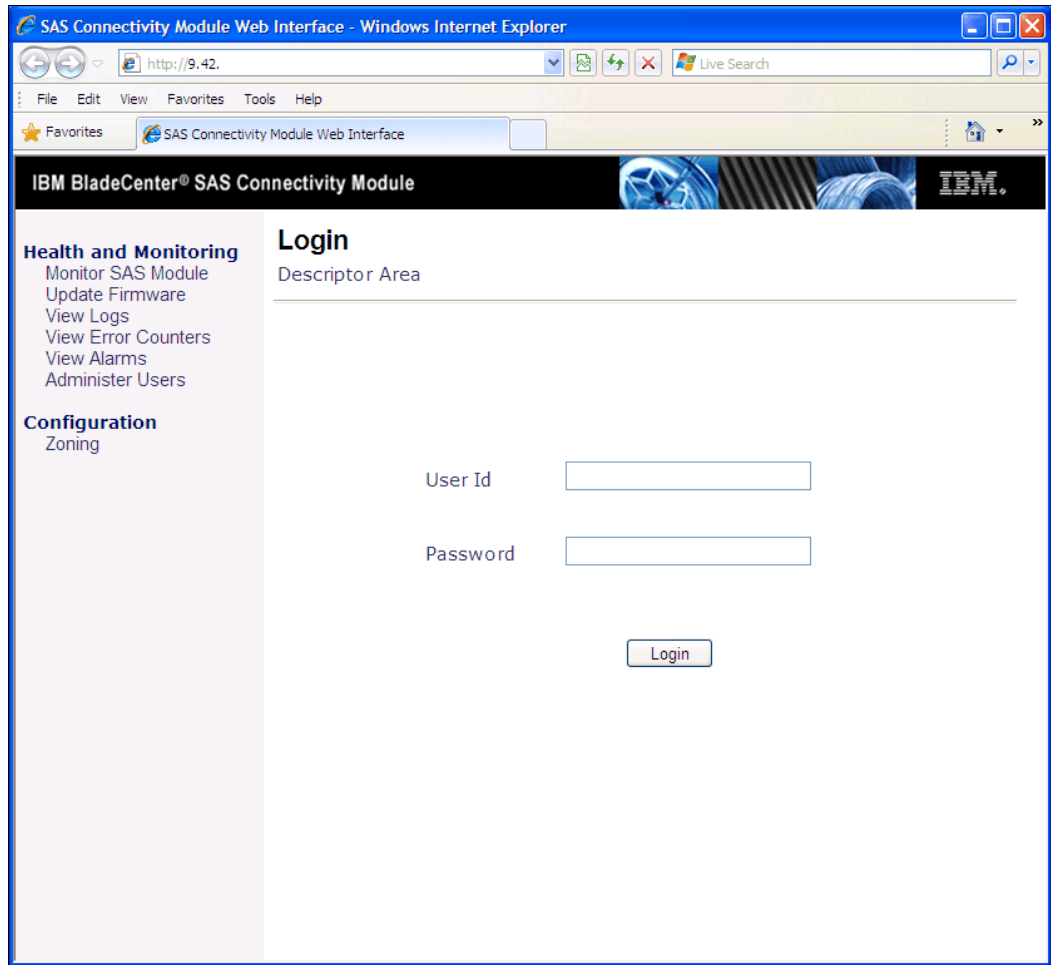


Figure 16-7 SAS Connectivity Module - login

2. In the Monitor Module window, click **Update Firmware**. The Update Firmware window opens, as shown in Figure 16-8. The current firmware level is also displayed.

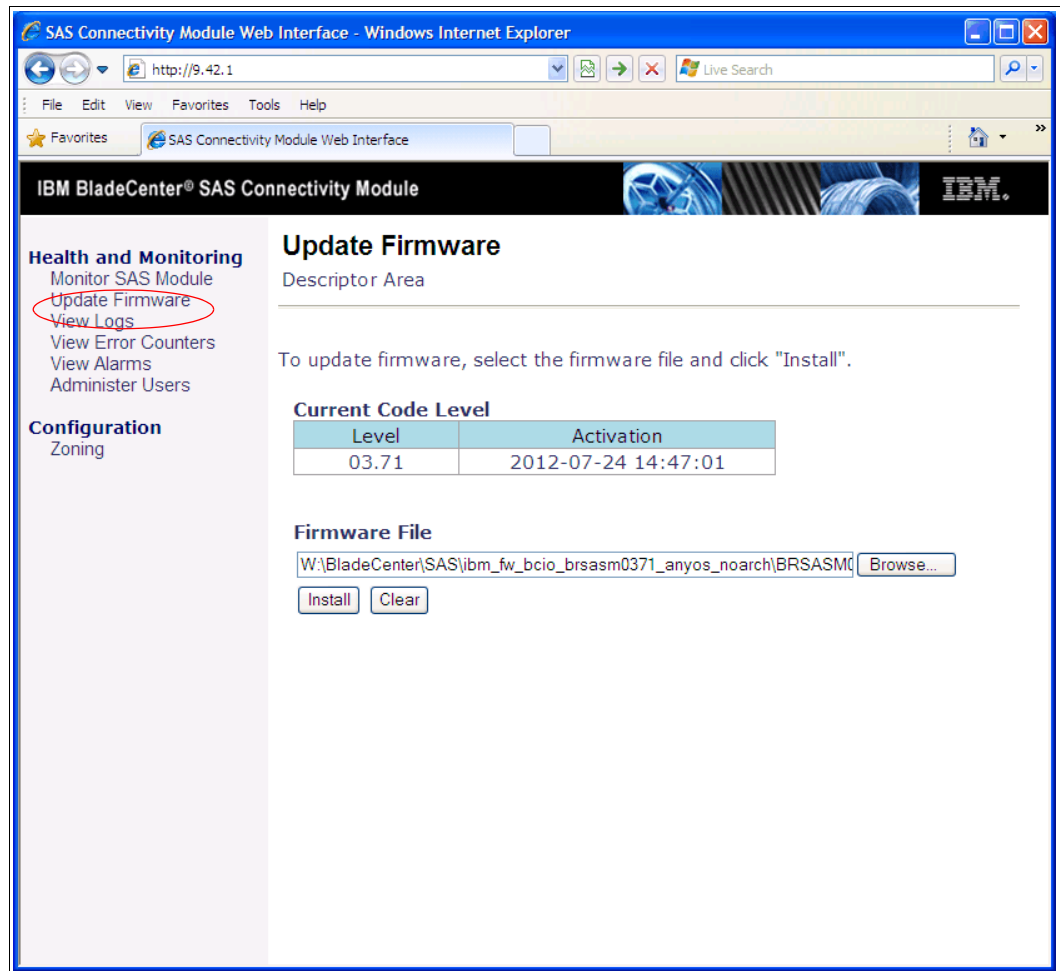


Figure 16-8 SAS Connectivity Module - Update Firmware

3. In the Firmware File field, enter the new firmware file name, or click **Browse** and locate the firmware file.

4. Click **Install** to install the new file. A firmware update progress window opens, as shown in Figure 16-9.

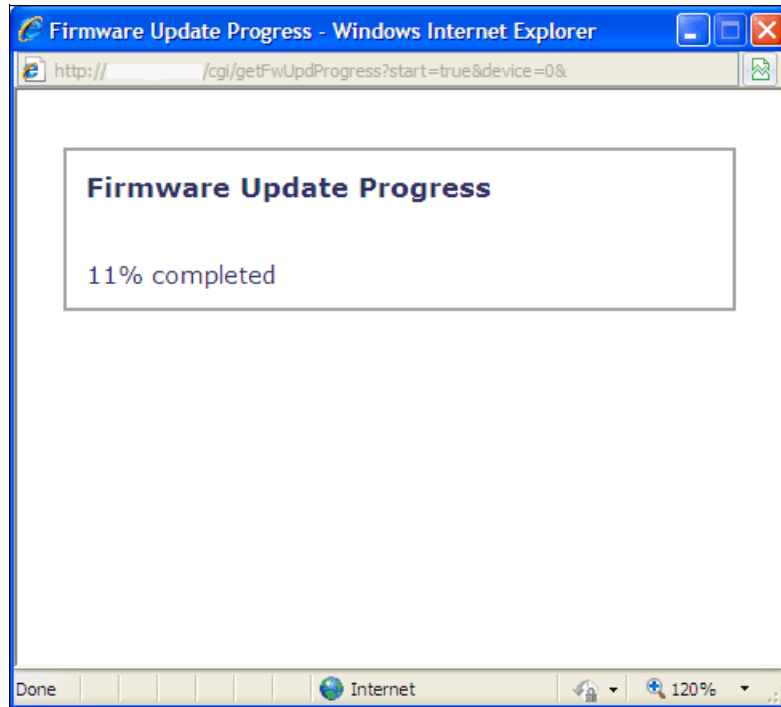


Figure 16-9 SAS Connectivity Module - installation confirmation



5. Avoid any power off or restart actions during the firmware upgrade, as noted in Figure 16-10.

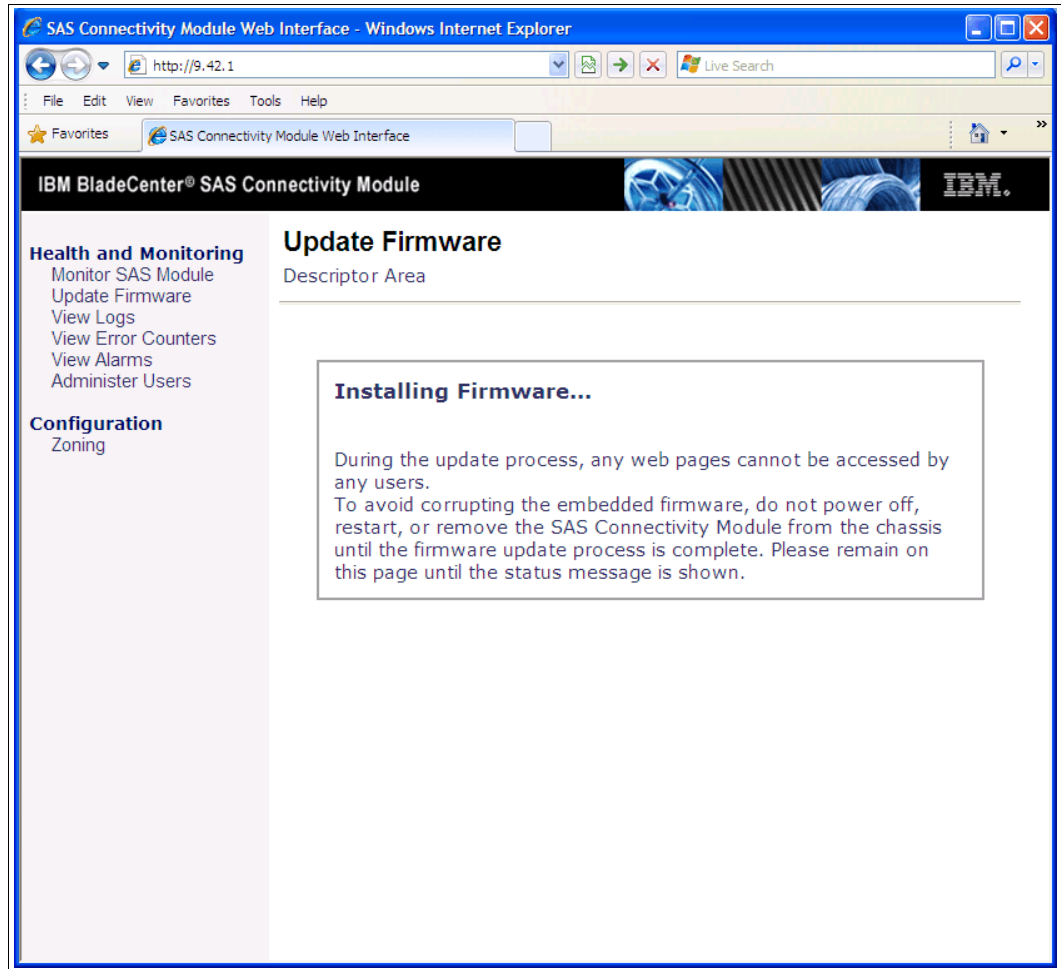


Figure 16-10 Installing firmware

6. Click **OK** or **Cancel**. If the installation of the new firmware file is successful, an installation confirmation window opens, as shown in Figure 16-11. If there are errors during the installation, an error message window opens.

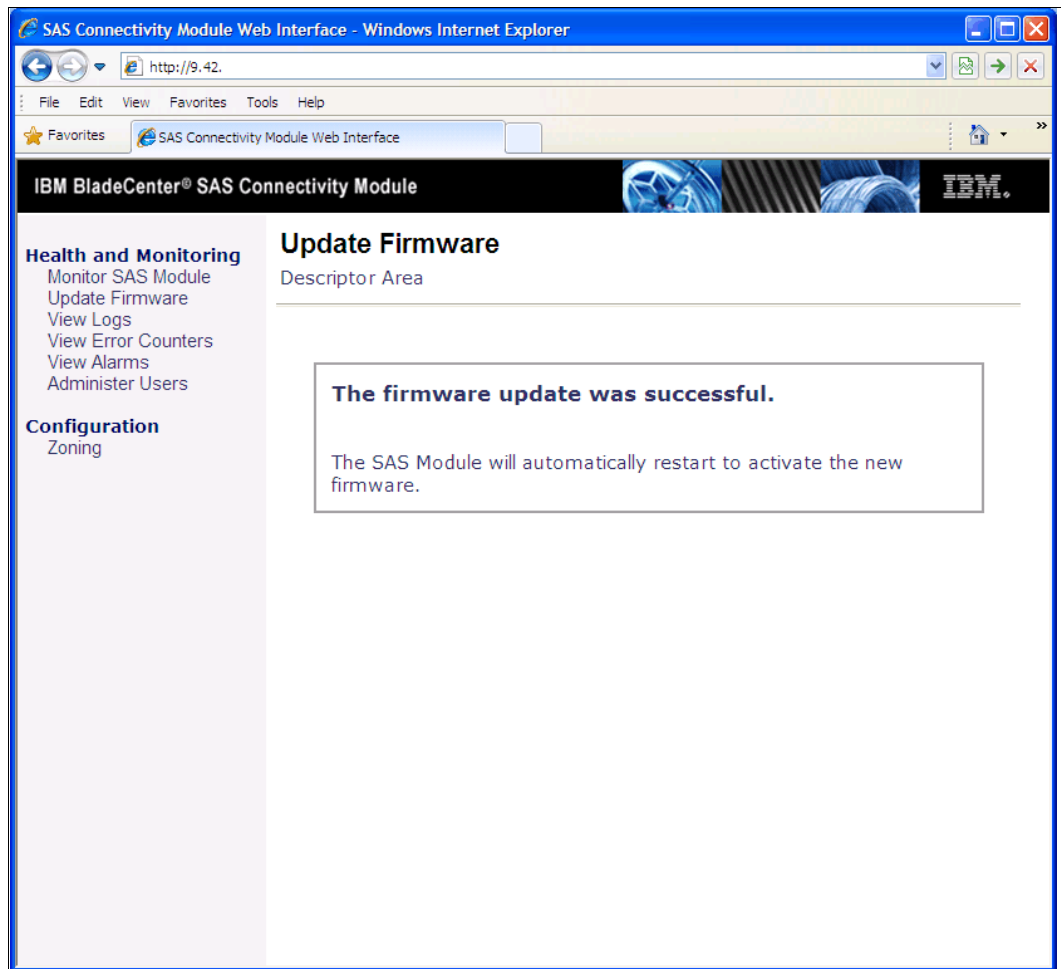


Figure 16-11 SAS Connectivity Module - update successful

## 16.2.7 Configuring the SAS Connectivity Module

Ensure that the external ports on the SAS Connectivity Modules are enabled to allow connectivity from the DCS3700 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 16-12.

The screenshot shows the IBM BladeCenter Advanced Management Module web interface. The main content area displays a table of modules in bays 1-10. Bays 3 and 4 contain SAS Conn Mod modules. Below the table, there are sections for 'Available actions' and 'I/O Module Advanced Setup'. In the 'I/O Module Advanced Setup' section, the 'External ports' dropdown menu is set to 'Enabled' and is highlighted with a red circle.

Bay	Type	Manufacturer	MAC Address	IP Address	Pwr	Unique ID	Type	ID	Stacking Mode	Protected Mode
1	Ethernet SM	DLNK (n/a)	00:0D:88:FA:F3:DD	<a href="#">View</a>	On	n/a		n/a	n/a	n/a
2	No Module									
3	SAS Conn Mod	IBM (n/a)	00:14:5E:C3:23:EA	<a href="#">View</a>	On	n/a		n/a	n/a	n/a
4	SAS Conn Mod	IBM (n/a)	00:14:5E:C3:0D:2A	<a href="#">View</a>	On	n/a		n/a	n/a	n/a
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: Disabled

External ports: **Enabled**

Figure 16-12 SAS Connectivity Module - configuration

## 16.2.8 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 that is connected to the SAN should be allowed access only 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 started by a simple selection from the Advanced Management Module (AMM). Zoning on the SAS Connectivity Module can be performed by 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 16-13.

**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	Date
<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	00/00/0000, 00:00:00
<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	00/00/0000, 00:00:00
<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	00/00/0000, 00:00:00
<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	00/00/0000, 00:00:00
<input type="radio"/>	<input checked="" type="checkbox"/>	Predefined Config 01	Pre-defined	Chassis: BCE, BCH, BCT and BCHT. SAS modules: 1 or 2. Zoned Blade Bays: 1-14. Each SAS module port belongs to its own zone. All Blades can access all external ports. Cannot be modified.	5	04/24/2007, 02:00:00

Figure 16-13 SAS Connectivity Module - predefined zones

You can select from five predefined zone configuration options. In this example, the option 5 predefined zone configuration is active (indicated by a check mark), as shown in Figure 16-13. With this option, each server bay is exclusively zoned with all the external ports, thus allowing access to one or more storage controller ports that are connected to the SAS Connectivity Module.

The example in Figure 16-14 on page 495 was captured from the SAS Connectivity Module web interface. It lists the Basic Zone Permission table for the HS23 blade server in slot 6 because the blade in slot 6 is used for this example. The blade in slot 6 is zoned with four external ports. The External port is set to True under the Connected column and Normal under the Status column. This is because the DCS3700 storage subsystem Controller A port is connected to the external port #6 for this example.

The screenshot displays the IBM BladeCenter SAS Connectivity Module configuration interface. It features a sidebar with navigation options like 'Health and Monitoring' and 'Configuration'. The main area is divided into 'Working Configuration' and 'Active Configuration' sections, both showing 'Predefined Config 01'. Below this, there are tabs for 'Zone Groups' and 'Basic Zone Permission Table'. The 'Basic Zone Permission Table' is currently active, showing a table of zone configurations. A red circle highlights the row for Zone Group ID 20, which is mapped to Blade Slot Connection 6 with a status of 'Normal'. Another red circle highlights the row for Zone Group ID 10, which is mapped to External Port 1 with a status of 'Normal'.

Port	Attached Port Add	Enabled	Connected	Status
External Port 1		true	true	Normal

Select	Zone Group ID	Port	Attached Port Add	Enabled	Connected	Status
<input type="checkbox"/>	10	External Port 1		true	true	Normal
<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	false	No Cable
<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	500507602316C401	true	true	Normal
<input type="checkbox"/>	19	Blade Slot Connection 5	0000000000000000	true	false	No Cable
<input type="checkbox"/>	20	Blade Slot Connection 6	500507604D000EDA	true	true	Normal

Figure 16-14 SAS Connectivity Module - zone configuration

**Note:** There is only a single path active from the DCS3700 storage subsystem controller A to the SAS Connectivity Module in BladeCenter I/O slot bay 3. The second path from DCS3700 storage subsystem controller B is to the SAS Connectivity Module in BladeCenter I/O slot bay 4.

## 16.3 Installing the DS Storage Manager host software

Here you install the host server components of DCS3700 Storage Manager. These components include the following items:

- ▶ DS Storage Manager 10 Utilities
- ▶ DS Storage Manager 10 Agent
- ▶ MPIO Device Specific Module (DSM)

You install DS Storage Manager 10 Utilities because then you can run the **hot\_add** utility; if you use this utility, you do not need to reboot the host server when you add logical drives.

Because the host server contains a 2-port SAS Expansion Card, you need multipath support. This support is installed as part of the Host selection when you install the DS Storage Manager. For the exact installation steps, see Chapter 4, "IBM System Storage DS Storage Manager installation" on page 131.

After the DCS3700 Storage Manager client is installed, add the DCS3700 storage subsystem in the Enterprise Management window and prepare the logical drives for the Windows server. For simplicity's sake, create just one logical drive of 20 GB.

Define your Windows Server 2008 host in the Storage Manager and map it to the logical drive.

We explained how to add storage subsystems, define hosts, prepare logical drives, and map them in Chapter 10, "Administration: Mappings tab" on page 327, so we do not repeat these steps here. Figure 16-15 shows the 200 GB logical drive named Data that is mapped to the host server Blade06.

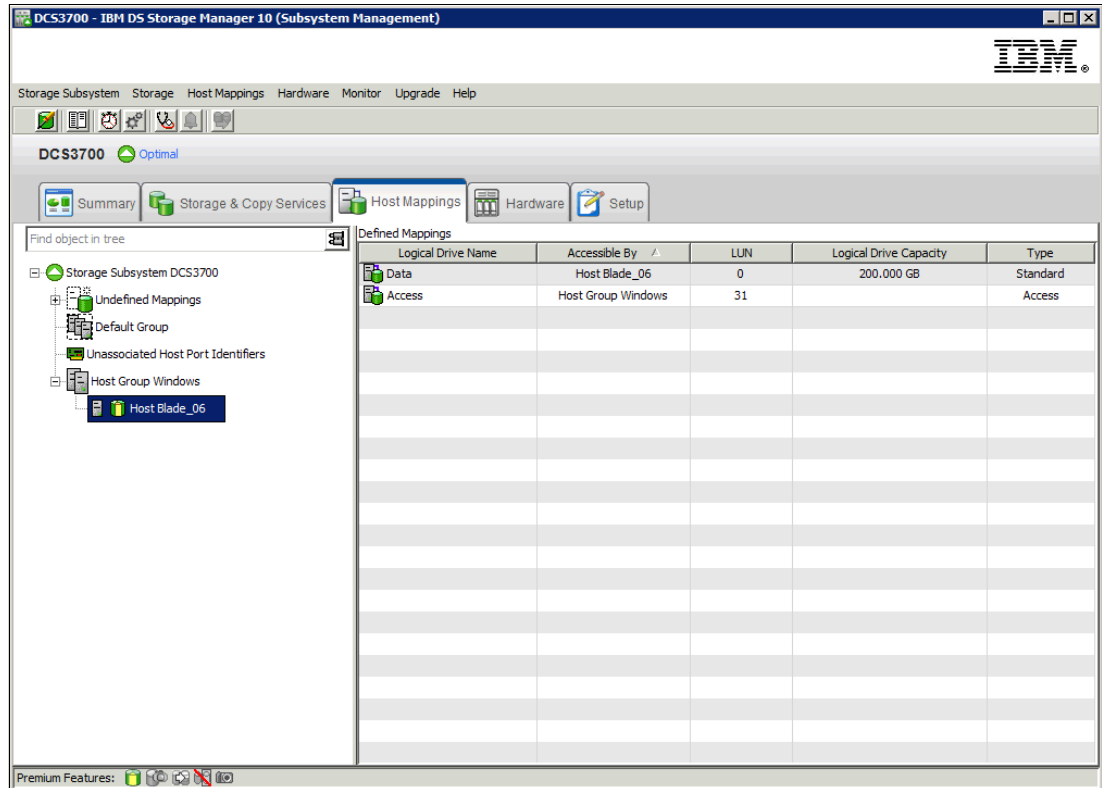


Figure 16-15 Drive mapping

## 16.4 Configuring the disk space in Windows Server 2008 R2

The logical drive must be recognized in Windows before you can start using the disk space. One possibility is to shut down and reboot the host server or to select **Rescan Disks** from the **Action** menu bar of the Disk Management window. A better alternative is to use the **hot\_add** utility (a part of SMutil).

By default, the **hot\_add** executable file is in C:\Program Files (x86)\IBM\_DS\util. In a command prompt window, change to this directory. If you installed the DCS3700 Storage Manager in a non-default directory, change to that directory instead, and then to subdirectory util. Then, run **hot\_add**.

When the utility completes, you can use the Disk Management applet in the Computer Management utility and configure the disk space for operating system use. Complete the following steps:

1. Initialize and optionally convert to dynamic disk.

This task is done by using the Initialize and Convert Disk wizard. When the wizard finishes, the Disk Management applet looks similar to Figure 16-16.




Figure 16-16 Disk Management applet - initialized disk

As you can see, the 200 GB logical drive is recognized as Disk 1, but it does not contain a partition yet. (It is deallocated.)

2. Initialize the disk by right-clicking Disk1 and selecting **Initialize Disk**, as shown in Figure 16-17.




Figure 16-17 Initialize Disk



3. Create and format a partition. Right-click in the deallocated disk space and select **New Simple Volume**, as shown in Figure 16-18.




Figure 16-18 Creating a partition

The New Simple Volume wizard starts.

- Follow procedure that is presented by the wizard to define the partition size, drive letter, file system, and volume label. The partition is created and formatted. The Disk Management applet now displays the new partition (allocated drive F) that is ready to be used (see Figure 16-19).

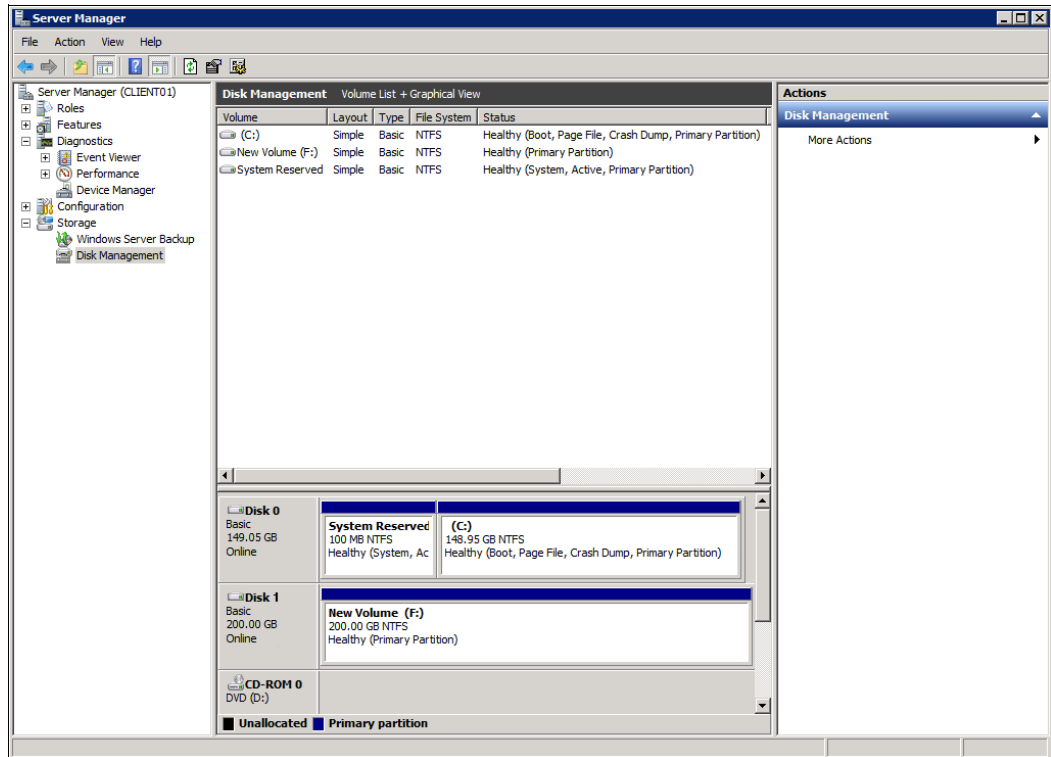


Figure 16-19 New partition on the logical drive

5. You can now access the logical drive on the DCS3700 storage subsystem as drive letter F, as shown in Figure 16-20.




Figure 16-20 Logical drive is ready for use





## Windows Fibre Channel configuration guide

This chapter explains how to configure an IBM System Storage DCS3700 storage subsystem by using Fibre Channel on Windows hosts. It also explains how to verify a correct installation and shows how to discover the storage logical drives. It describes how to use QLogic SANSurfer to identify the installed HBAs.

Before connecting your hosts to your DCS3700 storage subsystem, make sure to use the guidelines for levels of firmware and device driver levels of each component of the solution.

For the supported levels, see the IBM System Storage Interoperation Center (SSIC) website:

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

If you need to update any components, you can find details at the IBM System Storage support website:

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

## 17.1 Fibre Channel storage configuration

This section presents the steps that you must perform in your Windows host to configure and manage a DCS3700 storage subsystem that uses Fibre Channel connections.

We assume that you have already configured the DCS3700 storage subsystem by using the DS Storage Manager, and that the logical volumes are already mapped to the Windows host, as described in the Chapter 10, “Administration: Mappings tab” on page 327.

### 17.1.1 Planning for the installation

For our example setup, we used the following hardware and software components:

- ▶ One IBM BladeCenter HS23 blade server with a QLogic 8Gb Fibre Channel Expansion Card (CIOv)
- ▶ Two IBM BladeCenter Fibre Channel pass-through modules
- ▶ One DCS3700 storage subsystem
- ▶ Two Brocade 6510 FC Switches
- ▶ Windows Server 2008 R2 Datacenter with SP1
- ▶ The latest version of MPIO drivers for Windows

Regarding the *zoning configuration*, make sure that the SAN zoning is correctly set up according to your planned configuration. For more information about SAN FC zoning, see *Implementing an IBM b-type SAN with 8 Gbps Directors and Switches*, SG24-6116.

For more information about the HS23 blade server, components, and standard options, see *IBM BladeCenter HS23E*, TIPS0887.

Figure 17-1 shows the example setup.




Figure 17-1 BladeCenter and DCS3700 - Windows implementation FC interconnection layout

## 17.1.2 Installing the HBA drivers

To support the Fibre Channel adapters on your system, you must ensure that you have the correct drivers installed. To download the Windows drivers for both host bus adapters, go to the following website:

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

Select the appropriate IBM Storage System DS product and click the **Download** option to find all the available packages for the various HBAs supported by your storage subsystem. Select the appropriate package for your specific model and operating system version to download the drivers.

If you are not sure about the adapter model that is installed in your host, you can use your adapter management software to obtain the information, as shown in 17.1.3, “Installing QLogic SANSurfer” on page 506. You can also verify the list of installed storage adapters in your Windows host for the adapter type by opening the Server Manager and clicking **Diagnostics** → **Device Manager** → **Storage controllers**, as shown in Figure 17-2.

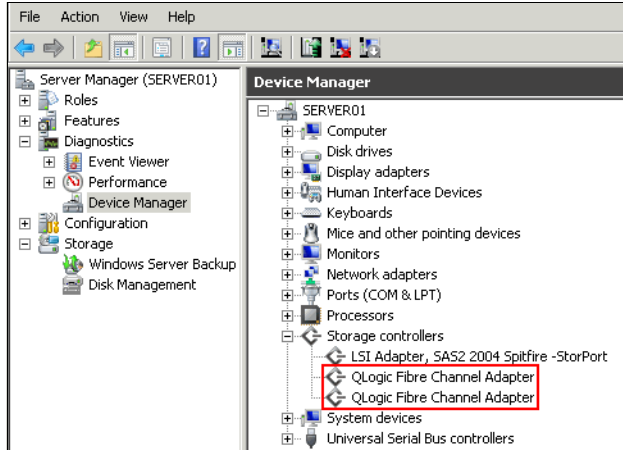


Figure 17-2 Server Manager window showing the storage controllers that are installed

Right-clicking any of the adapters and selecting **Properties** opens the Drivers tab, which displays the version that installed for both adapters.

In this example, for FC attachment we use QLogic FC cards with Windows 2008 R2 with SP1, which has the correct adapter driver of the QLogic Fibre Channel Host Bus Adapters, so there is no need to install it. If you are using any other supported HBAs, install the specific package for your adapter and Windows version from the IBM Support website.

**Note:** For redundancy, use multiple HBAs to access your storage subsystem. Only adapters of the same type can be configured as a failover pair.

Before you install a driver, make sure that you download the correct version for your operating system from the IBM Support website.

To change your adapter settings and update the firmware or BIOS of your adapter, use the management software that is provided by your HBA manufacturer. In our example, we use QLogic SANSurfer.

### 17.1.3 Installing QLogic SANSurfer

The QLogic SANSurfer component is the graphical user interface that you use to manage your QLogic adapters. You can use this application to perform the following actions:

- ▶ Determine your adapter type.
- ▶ Check its status.
- ▶ Determine the firmware, BIOS, and driver levels.
- ▶ Set specific parameters.
- ▶ Display the LUNs that are available

This section covers the installation on a Windows host.



QLogic SANsurfer contains the following components:

- ▶ QLogic SANsurfer GUI: A Java -based GUI application to manage HBAs on host servers.
- ▶ Windows QLogic SANsurfer Agent (or qlremote agent): This software is required on servers where HBAs are.

**Note:** Starting with Version 2.0.30b62, IBM no longer releases a specialized IBM version of the QLogic SANsurfer HBA Manager or SANsurfer PRO program. You can use the QLogic version 5.0.X buildX that supports Fibre Channel (FC) HBAs and iSCSI HBAs.

## Installation

The name of the installation file indicates the SANsurfer version. At the time of writing, the file name is `standalone_sansurfer5.0.1b75_windows_install.exe` for the FC version.

Complete the following steps:

1. Find the downloaded file (`standalone_sansurfer5.0.1b75_windows_install.exe`) and double-click it to start the installation wizard, as shown in Figure 17-3.

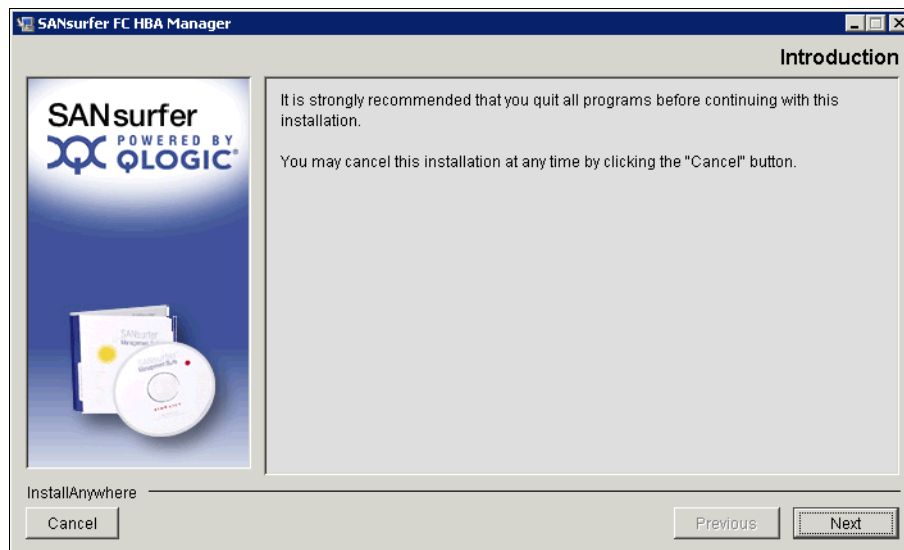


Figure 17-3 SANsurfer installation wizard

2. Click **Next** to open the Important Information window and **Next** again to proceed.

3. In the Choose Product Features window, click **FC HBA GUI and Agent**, as shown in Figure 17-4 and click **Next**.



Figure 17-4 Choose Product Features window

4. The following window asks for the path where the files will be installed; by default, the installation path is the one shown in Figure 17-5. Click **Next** to proceed.

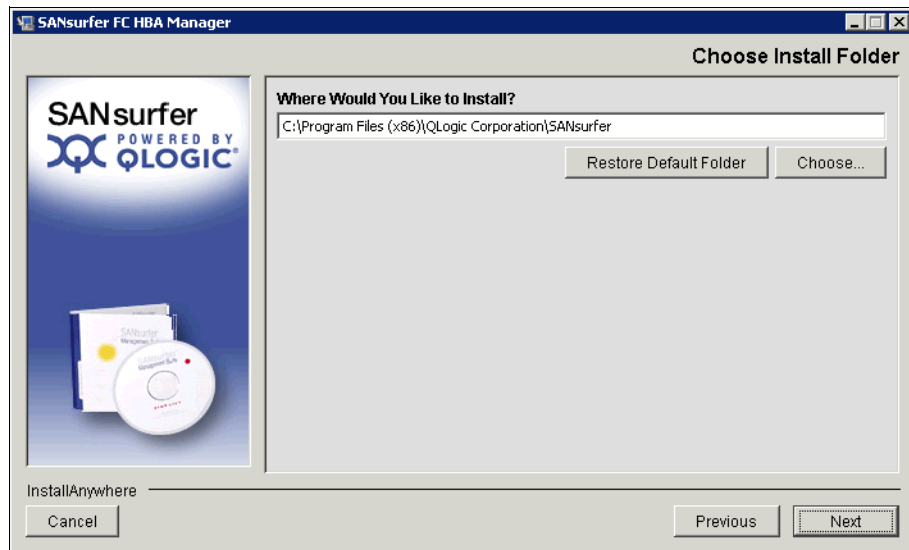


Figure 17-5 Choose Install Folder window

5. Choose whether to have a desktop icon. Click **Next** to proceed.

- The installation process begins and takes some time. You can see the progress in the window that is shown in Figure 17-6.

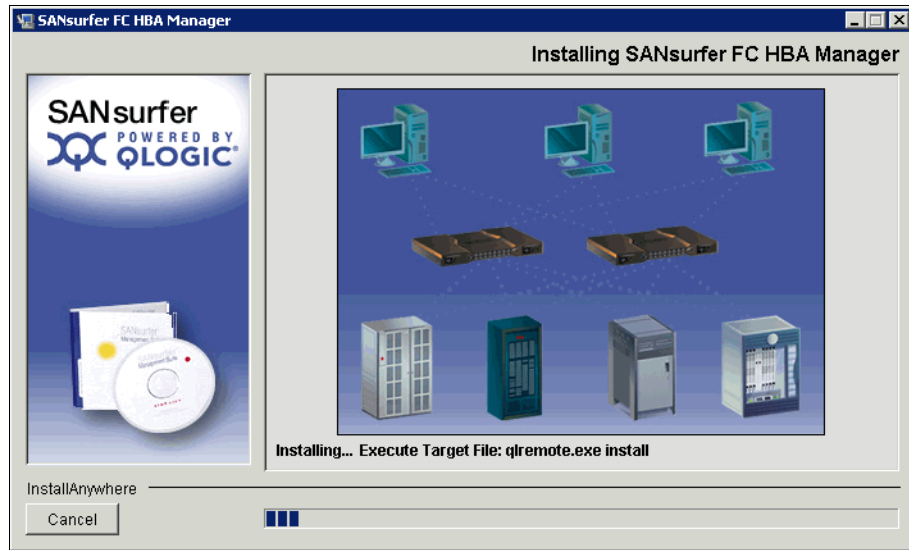


Figure 17-6 Installation progress window

- At the end of installation process, do *not* select **Enable QLogic Failover Configuration** because you are going to use MPIO as the default multipath driver (Figure 17-7).

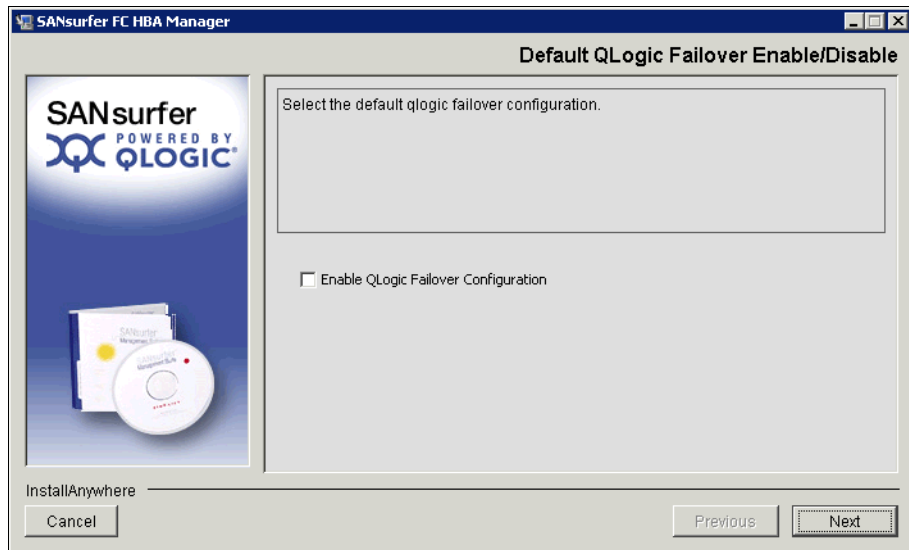


Figure 17-7 Leave the Enable QLogic Failover Configuration check box clear

## Running QLogic SANsurfer

QLogic SANsurfer can be started by double-clicking the desktop icon. The SANsurfer Connect to Host window opens, as shown in Figure 17-8.




Figure 17-8 Running QLogic SANsurfer

Click **Connect** to open SANsurfer HBA Manager, as shown in Figure 17-9.

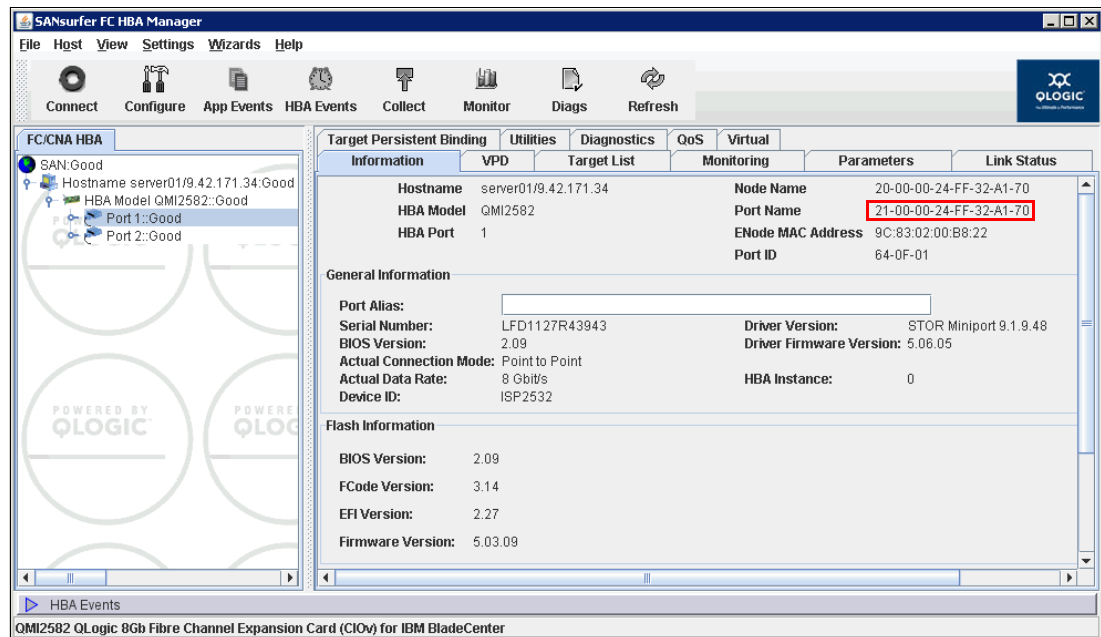


Figure 17-9 SANsurfer HBA Manager window with WWPN in red

The following tasks, activities, and information can be accomplished or acquired from the SANsurfer interface:

- ▶ Get FC HBA information:
  - HBA model
  - HBA Alias
  - Serial number
  - Driver version
  - Firmware version
  - WWN/WWPNs

- ▶ Maintenance tasks
  - Perform a firmware update.
  - Perform a driver update (only if the installed drivers are not the native OS driver; multipath drivers are managed separately).
- ▶ Advanced parameters configuration
  - Enable/disable HBA Port Loop ID
  - Enable/disable HBA Port BIOS
  - Enable/disable Loop Reset Delay
  - Enable/disable Target Reset
  - Configure LUN masking

For more information about SANsurfer capabilities, press the F1 key or click **Help**.

### 17.1.4 Installing the Windows multipath driver

The MPIO Device Specific Module (DSM) driver for Windows is the multipath supported driver for attaching your DCS3700 storage subsystem to your Windows host.

The MPIO driver for Windows is part of the DS Storage Manager software. You can download it from the following website:

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

A detailed description for installing the Storage Manager is in Chapter 4, “IBM System Storage DS Storage Manager installation” on page 131. When you use the procedure outlines there, select either the **Host** installation type or **Custom** installation.

It is also possible to download the Storage Manager package and extract its contents. To do so, look for the SMIA-WSX64-01.03.0305.0608.exe file (this was the version of the MPIO driver package at the time of writing) under the ..\WS03WS08\_10p83\_X64\Windows folder to install the driver only, as shown in Figure 17-10.



Figure 17-10 MPIIO Device Specific Module installation

After a successful installation, you must reboot the host, as shown in Figure 17-11.




Figure 17-11 Reboot needed after MPIIO DSM driver installation

After the reboot, click **Server Manager** → **Device Manager** → **Storage Controllers**, and you see a new device called “Microsoft Multi-Path Bus Driver”, as shown in Figure 17-12.

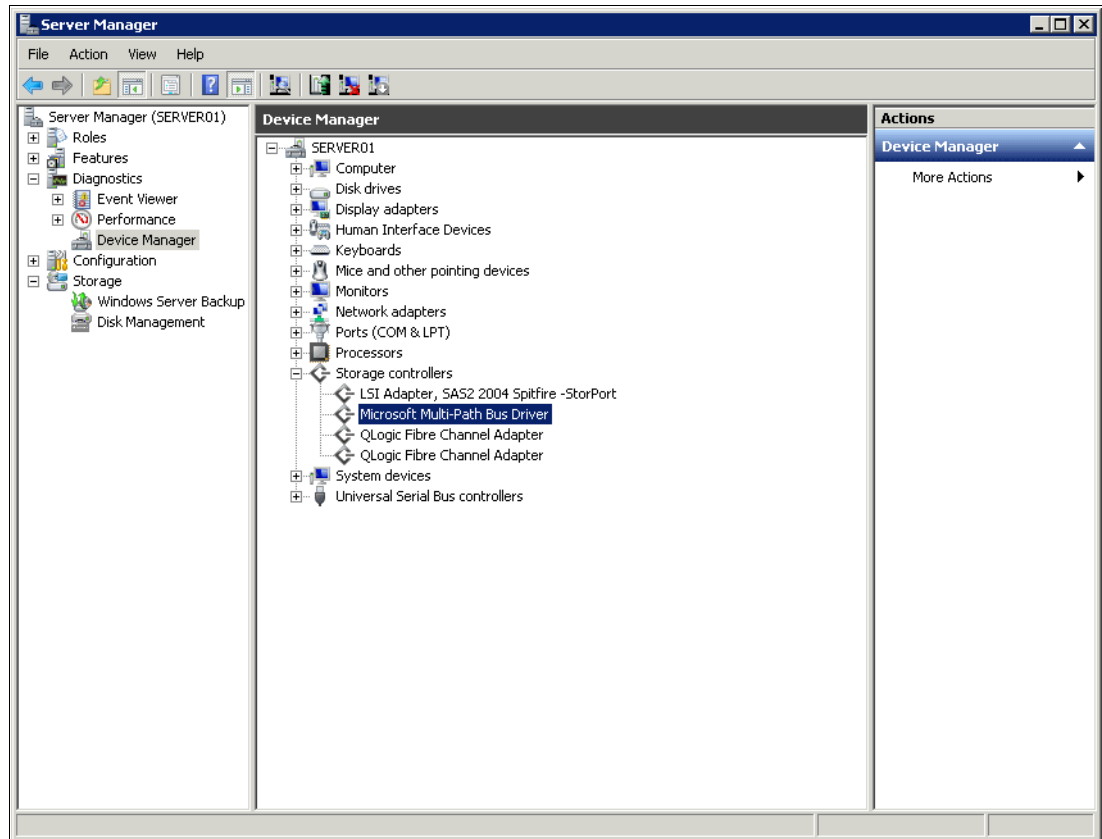


Figure 17-12 Microsoft multipath driver installed

To check the MPIO DSM driver installation and multipath status, right-click **Microsoft Multi-Path Bus Driver** and select **Properties**. Click the **MPIO** tab to see all the path options, as shown on Figure 17-13.

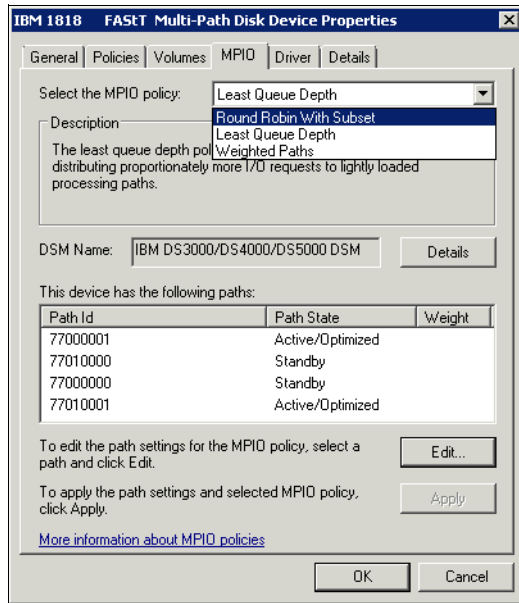


Figure 17-13 MPIO policy options

The load balance policy Least Queue Depth is selected by default, but it is possible to change it from this window, if necessary. You also can use the **MPCLAIM** command to manage the path policies on a single disk or all disks; for more information, go to the following website:

<http://technet.microsoft.com/en-us/library/dd851699.aspx>

## 17.1.5 Discovering devices (logical disk drives)

New and existing devices can be discovered by clicking **Server Manager** → **Storage**, right-clicking **Disk Management**, and selecting **Rescan Disk**, as shown in Figure 17-14.

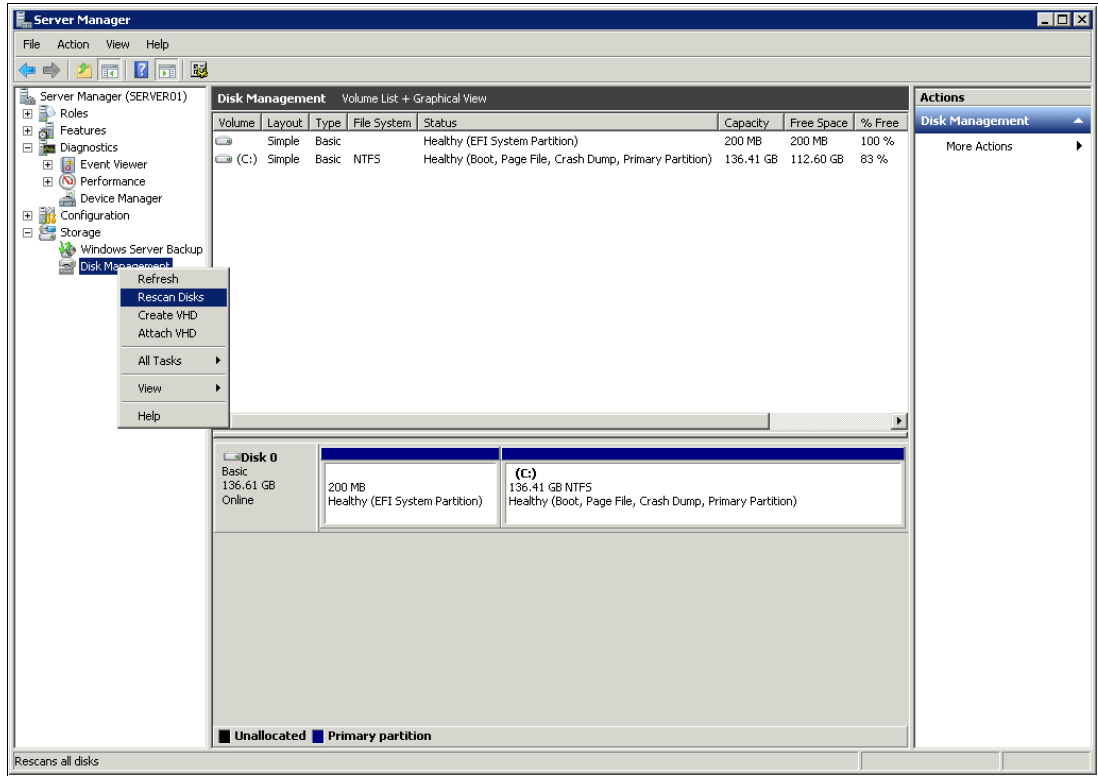


Figure 17-14 Rescan disks



If new disks are found, they appear as offline, such as Disk 1 in Figure 17-15.

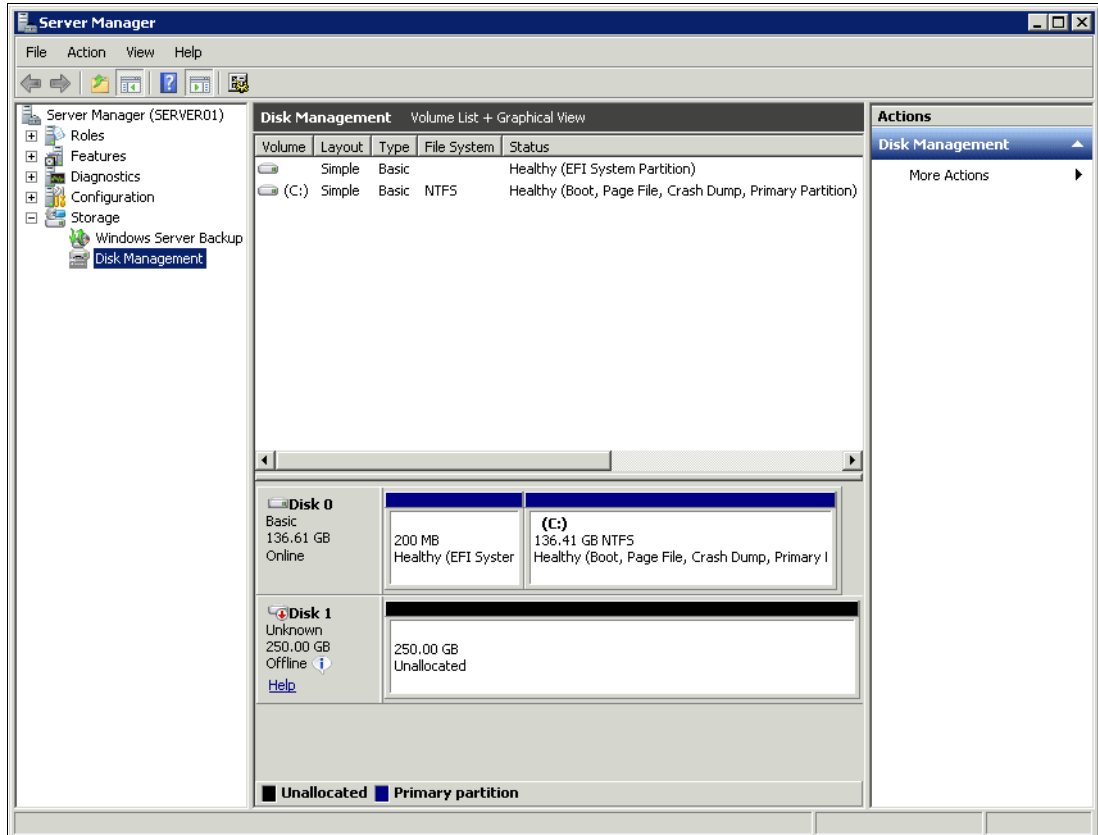


Figure 17-15 Adding a disk

To make the new disk available for use, right-click **Disk 1** and select **Online**, as shown in Figure 17-16.

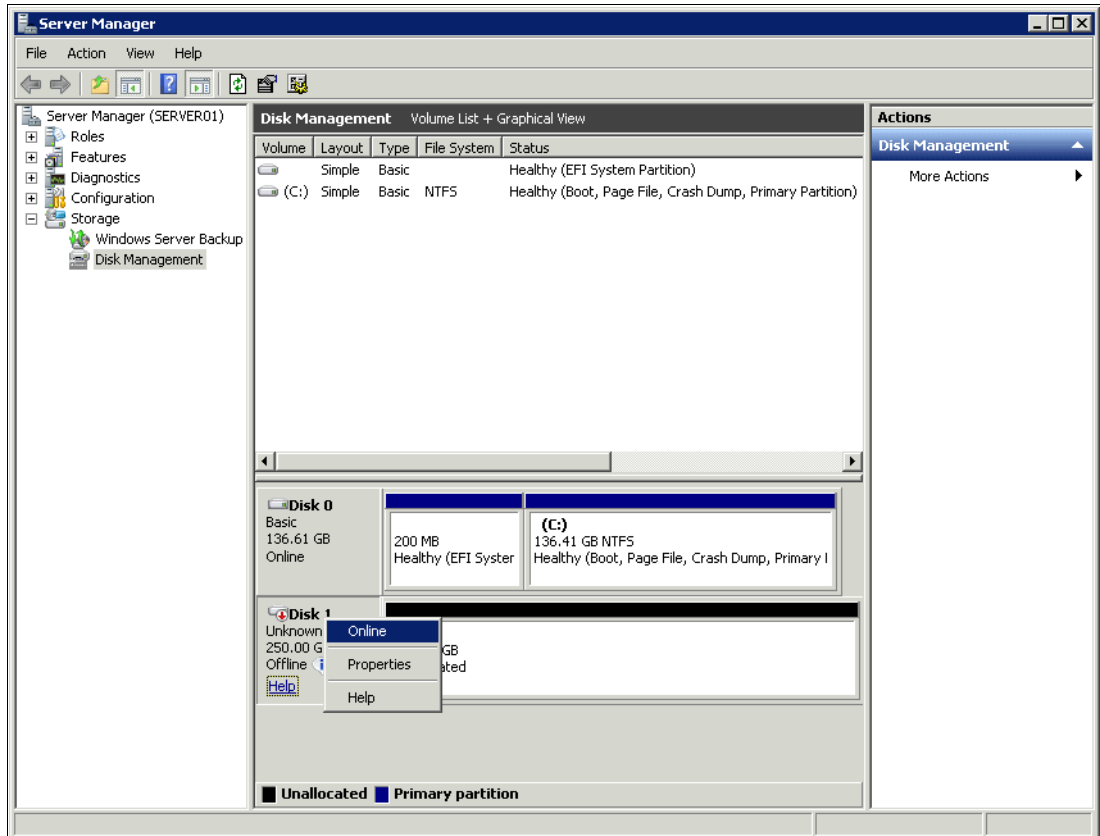


Figure 17-16 Bring Disk 1 online

Initialize the disk by right-clicking *Disk 1* and selecting **Initialize Disk**, as shown in Figure 17-17.

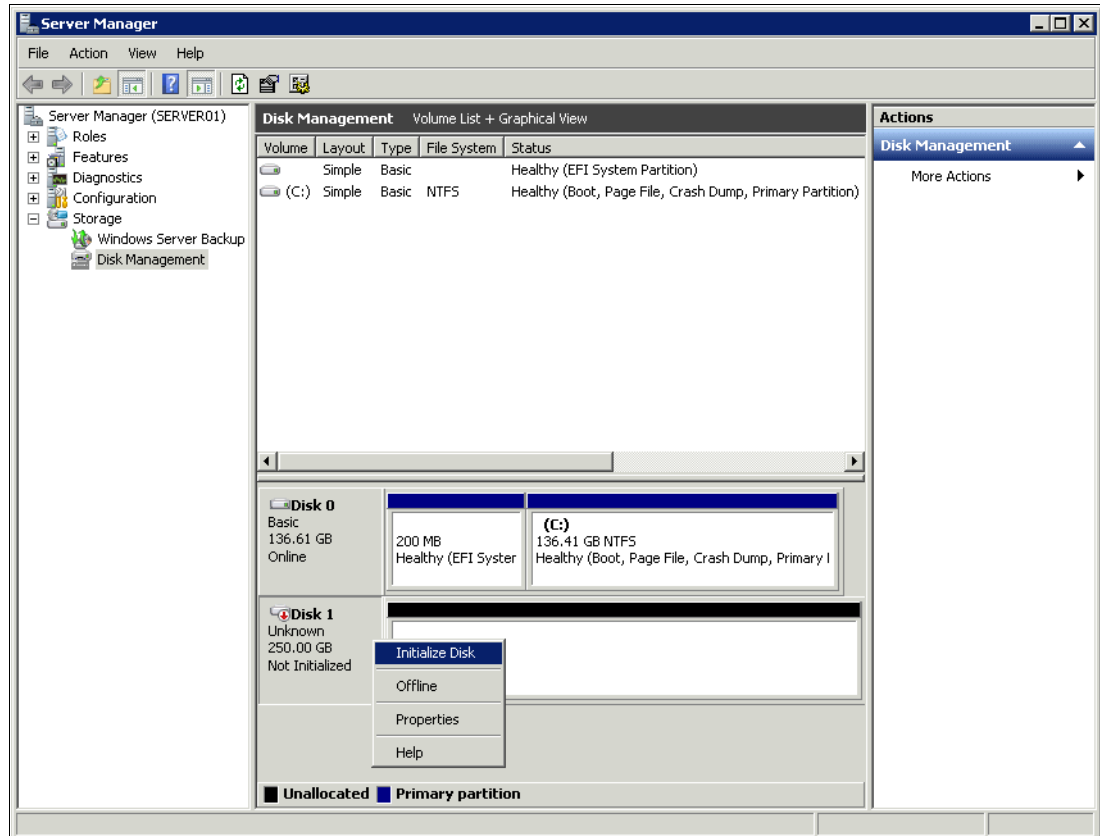


Figure 17-17 Initialize Disk

A window opens where you can initialize the disk and select the partition style (by default, Windows selects GPT, but it is only recommended for disks larger than 2 TB or Itanium -based hosts). MBR is the recommended style because of its compatibility, as shown in Figure 17-18.

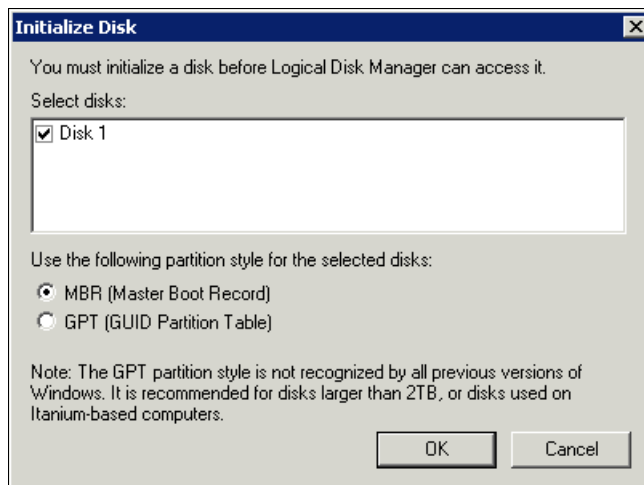


Figure 17-18 Initialize Disk wizard

To create a Simple Volume from the previously initialized disk, right-click the black disk area and select **New Simple Volume...**, as shown in Figure 17-19.

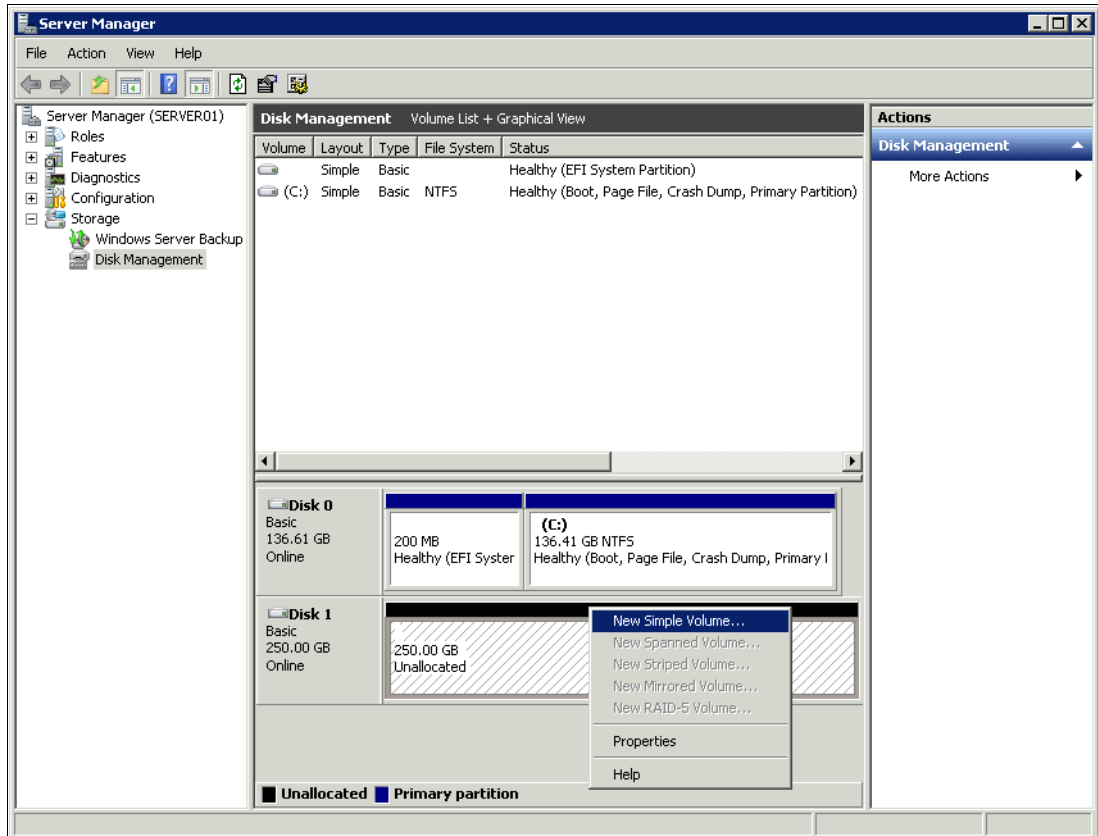


Figure 17-19 New Simple Volume

The New Simple Volume wizard starts with a window where you can click **Next** and follow the instructions to determine the Volume Size (Figure 17-20), Drive Letter (Figure 17-21 on page 519), and Format Partition (Figure 17-22 on page 519) options.

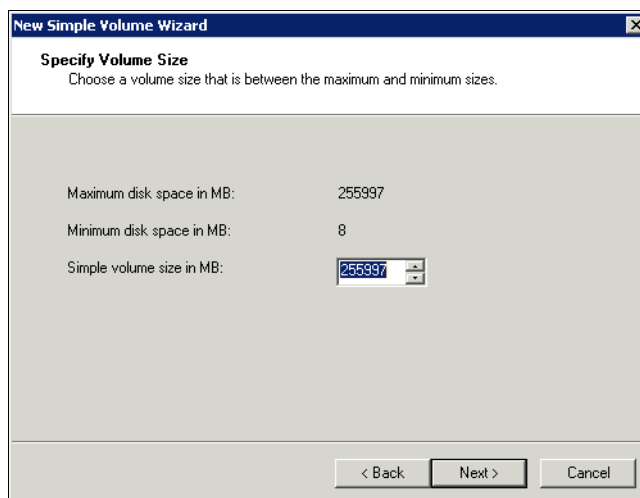


Figure 17-20 Volume Size

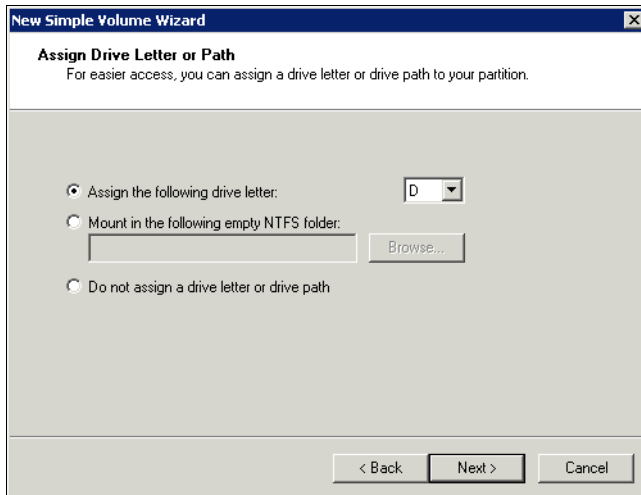


Figure 17-21 Drive Letter or Path

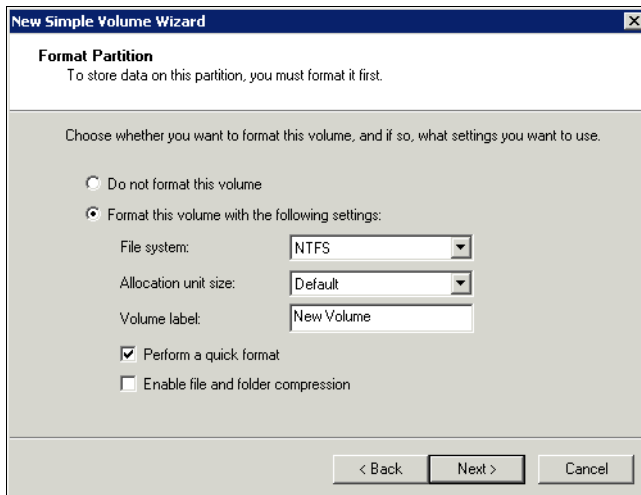


Figure 17-22 Format Partition

Formatting a partition can take some time depending on the disk size. You can observe the progress of the process by using the Disk Manager, as shown in Figure 17-23.

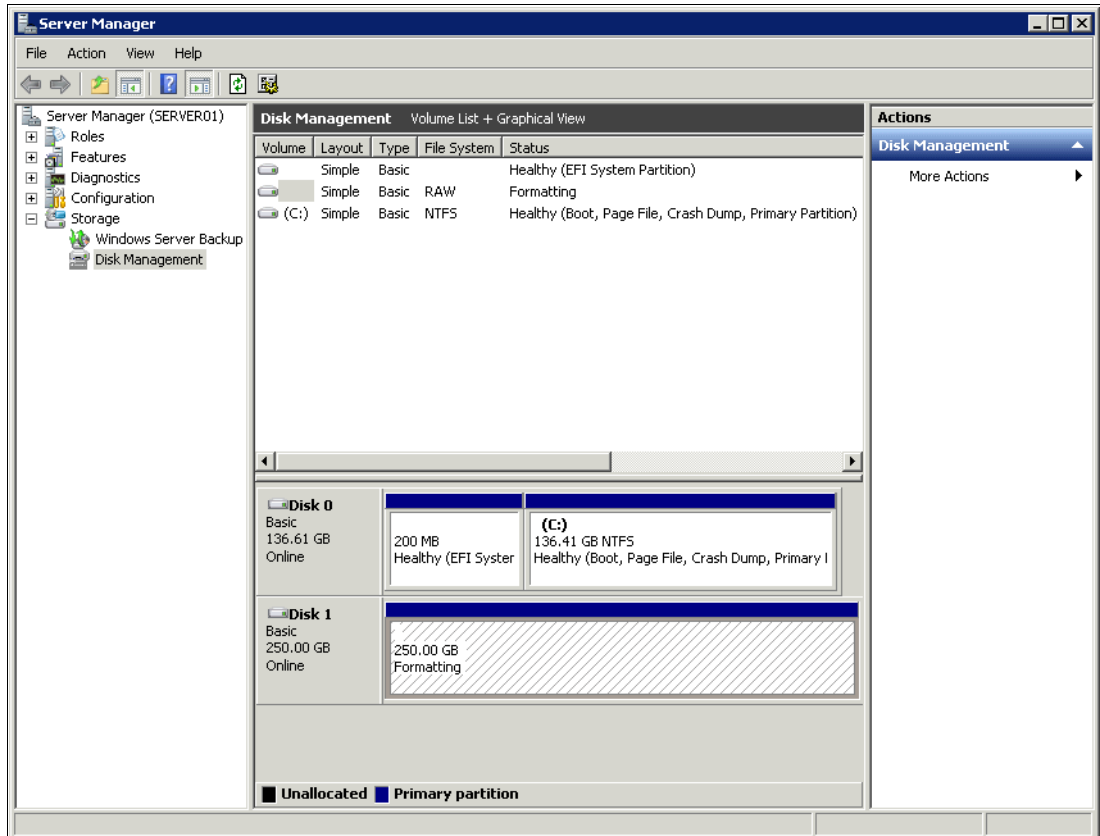


Figure 17-23 Formatting a new disk

After the formatting is complete, you can see the new disk by opening Windows Explorer, as shown in Figure 17-24.

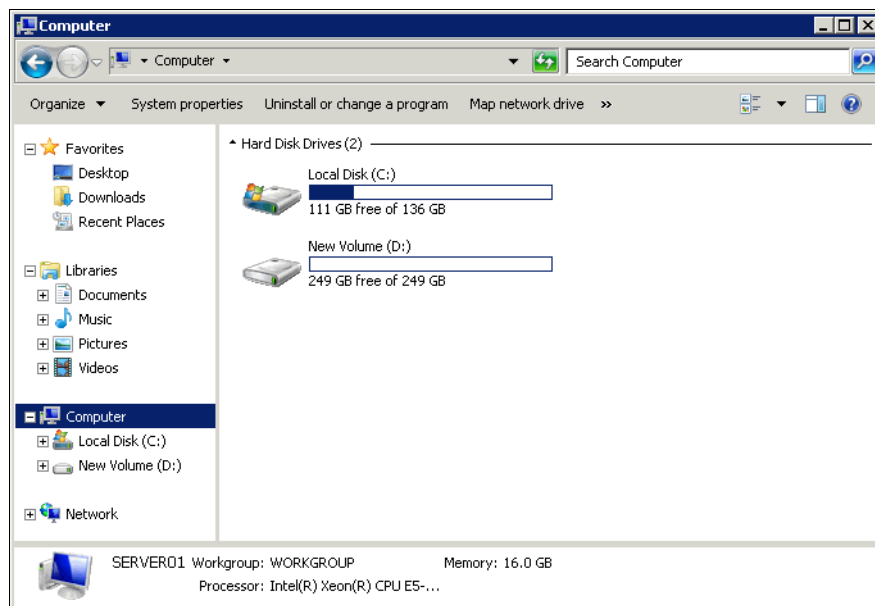


Figure 17-24 New disk available



## Linux configuration guide

This chapter explains how to configure the IBM System Storage DCS3700 storage subsystem on the Linux operating system. It then explains how to verify a correct installation and shows you how to discover the storage logical drives. This chapter also describes how to use Logical Volume Manager (LVM) to manage the configured storage. This chapter explains how to use QLogic SANSurfer to identify the installed HBAs.

Whatever supported Linux distribution you have to connect to your DCS3700 storage subsystem, make sure to follow the guidelines for the levels of firmware and device driver levels of each component of the solution.

For the supported levels, see the IBM System Storage Interoperation Center (SSIC) website:

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

If you need to update any components, you can find details at the IBM System Storage support website:

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

## 18.1 Fibre Channel storage configuration

This section presents the steps that you must perform in your Linux host to configure and manage your DCS3700 storage subsystem that is using Fibre Channel connections.

We assume that you have configured the DS3700 storage subsystem by using the DS Storage Manager and that the logical volumes are mapped to the Linux host, as described in Chapter 10, “Administration: Mappings tab” on page 327.

### 18.1.1 Planning for the installation

For our example setup, we used the following hardware and software components:

- ▶ One IBM BladeCenter HS23 blade server with a QLogic 8Gb Fibre Channel Expansion Card (CIOv)
- ▶ Two IBM BladeCenter Fibre Channel pass-through modules
- ▶ One DCS3700 storage subsystem
- ▶ Two Brocade 6510 FC Switches
- ▶ Red Hat Enterprise Linux Server 6 update 2
- ▶ The latest version of the IBM RDAC drivers for Linux (formerly provided by LSI/Netapp)

Regarding the *zoning configuration*, make sure that the SAN zoning is correctly set up according to your planned configuration. For more information about SAN FC zoning, see *Implementing an IBM b-type SAN with 8 Gbps Directors and Switches*, SG24-6116.

For more information about HS23 blade server and its components and standard options, see to *IBM BladeCenter HS23E*, TIPS0887.

The BladeCenter connection to the DCS3700 storage subsystem for this Linux FC interconnection layout is shown in Figure 18-1 on page 523.






Figure 18-1 BladeCenter and DCS3700 storage subsystem - Linux implementation FC interconnection layout

## 18.1.2 Installing HBA drivers

To support an FC or SAS adapter on your storage subsystem, you must ensure that you have the correct drivers installed. To download the Linux drivers for both host bus adapters, go to the following website:

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

Select the appropriate storage subsystem and click **Download** to find all the available packages for download for the various HBAs that are supported for connecting your storage subsystem. Select the appropriate package for your specific model and download the drivers.

If you are not sure of the adapter model that is installed in your host, you can use your adapter management software to obtain the information, as described in 18.1.3, “Installing QLogic SANSurfer” on page 524. You can also query the events in your Linux host for the adapter type, as shown in Example 18-1.

*Example 18-1 Determining the HBA type that is installed*

```
# dmesg | grep -i emulex
be2net 0000:16:00.0: Emulex OneConnect 10Gbps NIC(be3) port 0
be2net 0000:16:00.1: Emulex OneConnect 10Gbps NIC(be3) port 1
be2net 0000:16:00.4: Emulex OneConnect 10Gbps NIC(be3) port 2
be2net 0000:16:00.5: Emulex OneConnect 10Gbps NIC(be3) port 3
# dmesg | grep -i qllogic
```

```
qla2xxx [0000:00:00.0]-0005: QLogic Fibre Channel HBA Driver: 8.04.00.04.06.3-k.  
qla2xxx [0000:11:00.0]-00fb:1: QLogic QMI2582 - QLogic 8Gb Fibre Channel Expansion Card (CIOv) for IBM BladeCenter.  
qla2xxx [0000:11:00.1]-00fb:2: QLogic QMI2582 - QLogic 8Gb Fibre Channel Expansion Card (CIOv) for IBM BladeCenter.
```

In Example 18-1 on page 523, there are two QLogic QMI2582 adapters that are installed, and the device drive level is also displayed.

Depending on your adapter model, the driver might be packaged with your Linux distribution. The readme file mentions an “in-distro” driver because they are packaged together with the Linux distribution. If the driver is indicated as “standard” in the readme file for your adapter, then you must download and install the appropriate driver separately.

In this example, for FC attachment we used QLogic FC cards with Red Hat Enterprise Linux 6, which includes the adapter driver of the QLogic Fibre Channel Host Bus Adapters, so there is no need to install it. If you are using any other supported HBAs, install the specific package for your adapter and Linux distribution from the IBM Support website.

**Note:** For redundancy, use multiple HBAs to access your storage subsystem. Only adapters that are of the same type can be configured as a failover pair.

Before you install the driver, make sure that your system meets the requirements that are in the readme file of the driver package regarding kernel versions, host adapter settings, and other items. You might have to install the kernel headers and kernel source for the supported version of kernel before you install the drivers. You can download them from the following website:

<http://updates.redhat.com>

Set your adapter settings according to the readme file specifications. For our example, we set all the adapter settings, except for the following ones, which have the IBM defaults:

<b>Host Adapter settings</b>	Loop reset delay: 8
<b>Advanced Adapter Settings</b>	LUNs per target: 0
	Enable Target Reset: Yes
	Port down try count: 12

To change your adapter settings and update the firmware or BIOS of your adapter, you can use the management software that is provided by your HBA manufacturer. In our case, we used QLogic SANsurfer.

### 18.1.3 Installing QLogic SANSurfer

The QLogic SANsurfer component is the graphical user interface that is used to manage your QLogic adapters. You can use this application to perform the following tasks:

- ▶ Determine your adapter type.
- ▶ Check its status.
- ▶ Determine the firmware, BIOS, and driver levels.
- ▶ Set specific parameters.
- ▶ Display the LUNs available.

This section covers the installation on a Linux host.

**Important:** When you install SANsurfer, do not enable FC HBA adapter failover because you will be defining your own preferred multipath driver in this example, such as Linux RDAC or DM-Multipath.

The QLogic SANsurfer contains these components:

- ▶ QLogic SANsurfer GUI is a Java -based GUI application for the management of HBAs on host servers.
- ▶ Linux QLogic SANsurfer Agent (or qlremote agent) is software that is required on servers where HBAs are.

You can install the QLogic SANsurfer either from the GUI or from the CLI, as explained in the “GUI installation” and “Command-line interface installation” on page 526. If you do not have a GUI in your Linux system, then you can still install just the agent and run the GUI from a separate management station.

**Note:** Starting with Version 2.0.30b62, IBM no longer releases a specialized IBM version of the QLogic SANsurfer HBA Manager or SANsurfer PRO program. You can use the QLogic version 5.0.X buildX that supports Fibre Channel (FC) HBAs and iSCSI HBAs.

## GUI installation

The name of the installation file indicates the SANsurfer version. At the time of writing, the file package name is `standalone_sansurfer5.0.1b75_linux_install.tgz` for the FC version.

To complete the installation, complete the following steps:

1. Change to the directory where you downloaded the file and extract it by running the following command:  

```
tar -xvzf standalone_sansurfer5.0.1b75_linux_install.bin.tgz
```
2. Change the file attributes to be an executable file by running the following command:  

```
chmod a+x standalone_sansurfer5.0.1b75_linux_install.bin
```
3. Change to the directory where the installation file is and run the following command:

```
sh ./standalone_sansurfer5.0.1b75_linux_install.bin
```

This command initiates the installation process.

4. Confirm the first windows that open by clicking **Next**, and select the components to install. Continue the installation process. Do *not* select **Enable QLogic Failover Configuration** because we are planning to install RDAC as the multipath driver. For more information, see Figure 18-2.

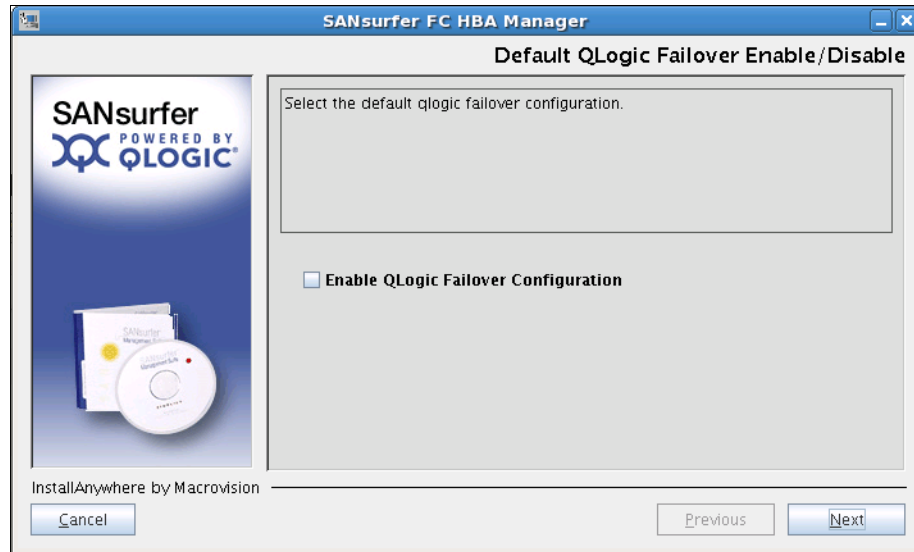


Figure 18-2 Linux SANsurfer installation

### Command-line interface installation

To install the QLogic combined version from a CLI, complete the following steps:

1. Open a shell and change to the directory that contains the sansurfer installation file.

2. If you want to install the Linux GUI Agent, run the following command:

```
sh ./standalone_sansurfer5.0.1b75_linux_install.bin -i silent  
-DSILENT_INSTALL_SET="QMSJ_G_LA"
```

3. If you want to install the Linux GUI, run the following command:

```
sh ./standalone_sansurfer5.0.1b75_linux_install.bin -i silent  
-DSILENT_INSTALL_SET="QMSJ_G"
```

4. If you want to install the Linux Agent, run the following command:

```
sh ./standalone_sansurfer5.0.1b75_linux_install.bin -i silent  
-DSILENT_INSTALL_SET="QMSJ_LA"
```

### Running QLogic SANsurfer

QLogic SANsurfer can run only on the GUI.

From the CLI, run **SANsurfer**, which opens the window that is shown in Figure 18-3 on page 527.

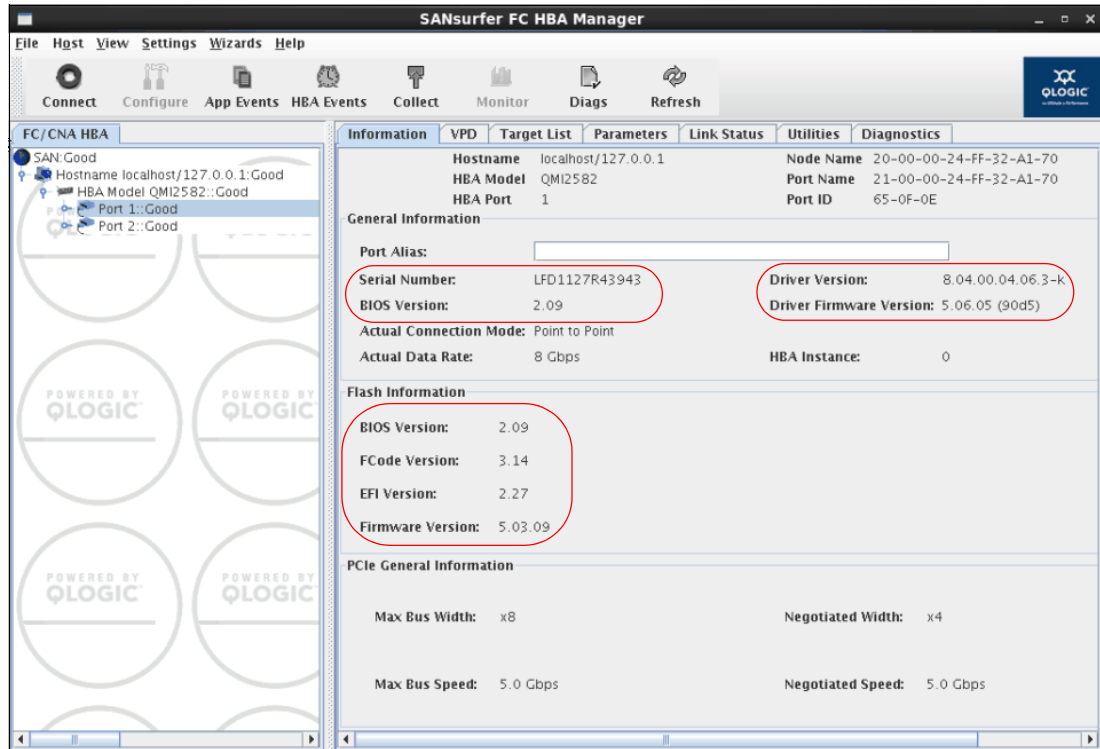


Figure 18-3 Running QLogic SANsurfer

The following tasks, activities, and information can be accomplished or acquired from the SANsurfer GUI interface:

- ▶ Get FC HBA Information:
  - HBA model.
  - HBA alias.
  - Serial number.
  - Driver version.
  - Firmware version.
  - WWN/WWPNs.
- ▶ Maintenance tasks
  - Perform a firmware update.
  - Perform a driver update (only if the installed drivers are not the native OS driver - multipath drivers, such as DMM and RDAC, which are managed separately).
- ▶ Advanced Parameters Configuration
  - Enable/Disable HBA port loop ID
  - Enable/Disable HBA port BIOS
  - Enable/Disable loop reset delay
  - Enable/Disable target reset
  - Configure LUN masking

For more information about SANsurfer capabilities, press the F1 key or run `man SANsurfer`.

## 18.1.4 Installing the Linux multipath driver

The MPP Linux RDAC and DM-Multipath are the current multipath supported drivers for attaching your DCS3700 storage subsystem to your Linux host.

The Linux RDAC provides redundant failover/failback support for the logical drives in the storage subsystems that are mapped to the Linux host server. The Linux host server must have host connections to the ports of both controllers A and B of the storage subsystem. It is provided as a redundant path failover/failback driver alternative to the Linux FC host bus adapter failover device driver.

In addition, RDAC provides functions and utilities to get much better information about the DCS3700 storage subsystem from the operating system. After the RDAC driver is installed, you have access to several commands, which are useful for problem determination and correction.

**Note:** In Linux, multipath device drivers such as RDAC or DM-Multipath are not installed together with the DS Storage Manager software. You must download and install these drivers separately.

The RDAC driver for Linux is not part of the DS Storage Manager software. You can download the RDAC driver from either of the following websites:

<http://www.ibm.com/servers/storage/support/disk>  
<http://support.netapp.com/NOW/public/apbu/oemcp/>

Device Mapper Multipathing (DM-Multipath) is the native Linux multipath I/O support that is added to the Linux 2.6 kernel tree with the release of Linux kernel 2.6.32-279. Section “Linux Device mapper multipathing driver (DMM / DM-Multipath)” on page 533 covers this driver.

### Linux RDAC Multipath Driver

Because in our example we cover Red Hat Enterprise Linux 6.2, we use RDAC driver version 09.03.0C05.0638, which is the latest supported version at the time of writing.

Make sure that your system meets the requirements that are stated in the readme file of the driver package regarding kernel versions, HBA settings, and so on. Also, pay attention to the installation sequence.

To see what the supported Linux OS and support multipath drivers are, visit the IBM System Storage Interoperation Center (SSIC) website:

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

**Note:** Before you install the RDAC driver, make sure that the HBA driver is installed and that the partitions and LUNs are configured and mapped to a Linux host type.

At the time of writing, the following Linux RDAC limitations apply. You can find more information about them in the readme file.

- ▶ The RDAC driver cannot coexist with an HBA-level failover driver, such as the Emulex, QLogic, or any other HBA failover driver.
- ▶ The Linux SCSI layer does not support *sparse/skipped* LUNs, which means that mapped LUNs must be sequentially numbered for the Linux host to see them.

- ▶ If you have multiple FC HBAs on the host server and each HBA port sees both controllers (through an unzoned switch), the Linux RDAC driver might return I/O errors during controller failover.

**Note:** The guideline here is for each HBA to have a path to both the controllers in the DCS3700 storage subsystem, either by using multiple SAN switches or by using one SAN switch with the appropriate zoning. The intent here is to not have multiple initiators in the same zone.

- ▶ The Linux RDAC reports I/O failures immediately if it detects a failure of all paths to the storage server. This behavior is unlike that of the HBA failover driver, which delayed for a certain period.
- ▶ The DCS3700 storage subsystem must be connected to the HBAs in order for the virtual HBA driver to be loaded.

**Tip:** Save a copy of the existing initial RAMdisk before you install the RDAC driver. This backup copy, together with a separate stanza in `/etc/lilo.conf` (SUSE OS) or `/etc/grub.conf` (Red Hat OS) is needed in case the system does not boot after you install the RDAC device driver.

Now you can install the RDAC driver. If you are planning to use your Linux system to configure your storage subsystem, then follow the instructions in 4.2, “Installing DS Storage Manager on Linux” on page 140 first, and then configure and map the logical drives to the Linux host to meet the previous guidelines.

**Note:** Auto Logical Drive Transfer (ADT/AVT) mode is now supported.

To install the Linux RDAC driver, complete the following steps (for the latest details, always check the readme file that comes with the driver package):

1. Install the kernel-source RPM files for the supported kernel by running the following commands. These files are available from your Linux vendor or your OS installation CD set. Also, be sure that the `ncurses` packages and the `gcc` packages are installed:

```
# rpm -ivh kernel-source*.rpm (SLES9/10/11)
# rpm -ivh kernel-syms*.rpm (SLES9/10/11)
# rpm -ivh kernel-devel*.rpm (RH)
# rpm -ivh kernel-utils*.rpm
# rpm -ivh kernel-header*.rpm
# rpm -ivh ncurses*.rpm
# rpm -iv ncurses-devel*.rpm
# rpm -iv gcc*.rpm
# rpm -iv libgcc*.rpm
```

Query the installed packages by running the following commands:

```
# rpm -qa | grep kernel
# rpm -qa | grep ncurses
# rpm -qa | grep gcc
```

2. The host server must have the non-failover Fibre Channel HBA device driver correctly built and installed before the Linux RDAC driver installation.
3. Change to the directory where the RDAC package was downloaded, and run the following command to extract the file:

```
# tar -xvzf rdac-LINUX-xx.xx.xx.xx-source.tar.gz
```

xx.xx.xx is the release version of the RDAC driver. In our case, we ran the following command:

```
# tar -xvzf rdac-LINUX-09.03.0C05.0638-source.tar.gz
```

4. Change to the directory where the files are extracted. In our case, we ran the following command:

```
# cd /linuxrdac-09.03.0C05.0638
```

5. Remove the old driver modules in case there were previous RDAC installations. Run the following command:

```
# make clean
```

6. To compile all the driver modules and utilities in a multiple-CPU server (SMP kernel), run the following command:

```
# make
```

7. To install the driver, run the following command:

```
# make install
```

This command does the following tasks:

- Copies the driver modules to the kernel module tree.
- Builds the new RAMdisk image (mpp-`uname -r`.img), which includes the RDAC driver modules and all driver modules that are needed at boot

8. A new line must be added to the /etc/grub.conf file with the new MPP drivers, as shown in Example 18-2. Add a line and then edit the title by adding “with RDAC MPP Support” at the end.

*Example 18-2 Edit /etc/grub.conf*

---

```
device (hd0) HD(1,800,64000,e35d67d6-9828-416c-b00b-fd8e8eb1ce45)
default=1
timeout=5
splashimage=(hd0,1)/grub/splash.xpm.gz
hiddenmenu
title Red Hat Enterprise Linux (2.6.32-279.el6.x86_64)
    root (hd0,1)
    kernel /vmlinuz-2.6.32-279.el6.x86_64 ro root=/dev/mapper/vg_server03-lv_root
rd_LVM_LV=vg_server03/lv_swap rd_NO_LUKS LANG=en_US.UTF-8 rd_LVM_LV=vg_server03/lv_root rd_NO_MD
SYSEFONT=latacyrheb-sun16 crashkernel=auto KEYBOARDTYPE=pc KEYTABLE=us rd_NO_DM rhgb quiet
    initrd /initramfs-2.6.32-279.el6.x86_64.img

title Red Hat Enterprise Linux - rdac MPP Support (2.6.32-279.el6.x86_64 - Linux rdac)
    root (hd0,1) <----- Add “rdac MP Support” on this line for better understanding
    kernel /vmlinuz-2.6.32-279.el6.x86_64 ro root=/dev/mapper/vg_server03-lv_root
rd_LVM_LV=vg_server03/lv_swap rd_NO_LUKS LANG=en_US.UTF-8 rd_LVM_LV=vg_server03/lv_root rd_NO_MD
SYSEFONT=latacyrheb-sun16 crashkernel=auto KEYBOARDTYPE=pc KEYTABLE=us rd_NO_DM rhgb quiet
    initrd/ mpp-2.6.32-279.el6.x86_64.img <----- Add the rdac mpp module at the end
```

---

9. Reboot the system.

**Note:** If you plan to use the Linux RDAC MPP Driver as the default during the boot, consider changing the *default* parameter to point to the correct driver support, as shown in Example 18-3 on page 531.



10. Verify that the following modules are loaded by running the `/sbin/lsmmod` command, as shown in Example 18-3:

- `sd_mod`
- `sg`
- `mpp_Upper`
- `qla2xxx`
- `mpp_Vhba`

*Example 18-3 Validating the loaded modules - lsmmod output*

---

```
qla2xxx          400728  4
scsi_transport_fc 55235  1 qla2xxx
scsi_tgt         12173  1 scsi_transport_fc
mpt2sas         182971  3
scsi_transport_sas 35940  1 mpt2sas
raid_class       4804   1 mpt2sas
mppVhba         139035  0
dm_mirror        14101  0
dm_region_hash  12170  1 dm_mirror
dm_log           10122  2 dm_mirror,dm_region_hash
dm_mod           81692  11 dm_mirror,dm_log
mppUpper        156935  1 mppVhba
sg               30124  0
sd_mod           39488  4
crc_t10dif       1541   1 sd_mod
```

---

11. Run the `mppUtil` command to display the RDAC driver version that you installed:

```
# mppUtil -V
Linux MPP Driver Version: 09.03.0C05.0638
```

12. Check the LUN mappings, as shown Example 18-4.

*Example 18-4 Checking the LUNs mappings*

---

```
# ls -lR /proc/mpp/
/proc/mpp/:
total 0
dr-xr-xr-x. 4 root root 0 Jul 19 00:16 DCS3700 <----- Storage subsystem

/proc/mpp/DCS3700:
total 0
dr-xr-xr-x. 4 root root 0 Jul 19 00:16 controllerA<----- DCS3700 Controller A
dr-xr-xr-x. 4 root root 0 Jul 19 00:16 controllerB<----- DCS3700 Controller B
-rw-r--r--. 1 root root 0 Jul 19 00:16 virtualLun0

/proc/mpp/DCS3700/controllerA: <----- DCS3700 Controller A - Discovered Ports
total 0
dr-xr-xr-x. 2 root root 0 Jul 19 00:16 qla2xxx_h1c0t1
dr-xr-xr-x. 2 root root 0 Jul 19 00:16 qla2xxx_h2c0t0

/proc/mpp/DCS3700/controllerA/qla2xxx_h1c0t1:
total 0
-rw-r--r--. 1 root root 0 Jul 19 00:16 LUN0 <----- Presented SAN Space - LUN0
-rw-r--r--. 1 root root 0 Jul 19 00:16 UTM_LUN31 <----- Access LUN

/proc/mpp/DCS3700/controllerA/qla2xxx_h2c0t0:
total 0
-rw-r--r--. 1 root root 0 Jul 19 00:16 LUN0
-rw-r--r--. 1 root root 0 Jul 19 00:16 UTM_LUN31

/proc/mpp/DCS3700/controllerB:<----- DCS3700 Controller B - Discovered Ports
```

```

total 0
dr-xr-xr-x. 2 root root 0 Jul 19 00:16 qla2xxx_h1c0t0
dr-xr-xr-x. 2 root root 0 Jul 19 00:16 qla2xxx_h2c0t1

/proc/mpp/DCS3700/controllerB/qla2xxx_h1c0t0:
total 0
-rw-r--r--. 1 root root 0 Jul 19 00:16 LUN0
-rw-r--r--. 1 root root 0 Jul 19 00:16 UTM_LUN31

/proc/mpp/DCS3700/controllerB/qla2xxx_h2c0t1:
total 0
-rw-r--r--. 1 root root 0 Jul 19 00:16 LUN0
-rw-r--r--. 1 root root 0 Jul 19 00:16 UTM_LUN31

```

---

13. You can run the **hot\_add** or **mppBusRescan** command, as shown in Example 18-5, to scan for the newly mapped LUNs. These commands can be run without rebooting the server.

*Example 18-5 Rescanning HBAs - mppBusRescan output*

```

# mppBusRescan
scan qla2 HBA host /sys/class/scsi_host/host1...
    found 1:0:0:1
    found 1:0:1:1
scan qla2 HBA host /sys/class/scsi_host/host2...
    found 2:0:0:1
    found 2:0:1:1
scan mpt2sas HBA host /sys/class/scsi_host/host0...
    no new device found
run /usr/sbin/mppUtil -s busscan...
scan mpp virtual host /sys/class/scsi_host/host4...
    found 4:0:0:1->/dev/sdc <----- Newly discovered SAN Space - LUN1

```

---

14. Verify that the RDAC driver discovered the LUNs by running the **ls -lR /proc/mpp** command and checking for virtual LUN devices, as shown in Example 18-6.

*Example 18-6 LUN verification*

```

# ls -lR /proc/mpp/
/proc/mpp/:
total 0
dr-xr-xr-x. 4 root root 0 Jul 19 00:16 DCS3700

/proc/mpp/DCS3700:
total 0
dr-xr-xr-x. 4 root root 0 Jul 19 00:16 controllerA
dr-xr-xr-x. 4 root root 0 Jul 19 00:16 controllerB
-rw-r--r--. 1 root root 0 Jul 19 00:16 virtualLun0
-rw-r--r--. 1 root root 0 Jul 19 00:16 virtualLun1

/proc/mpp/DCS3700/controllerA:
total 0
dr-xr-xr-x. 2 root root 0 Jul 19 00:16 qla2xxx_h1c0t1
dr-xr-xr-x. 2 root root 0 Jul 19 00:16 qla2xxx_h2c0t0

/proc/mpp/DCS3700/controllerA/qla2xxx_h1c0t1:
total 0
-rw-r--r--. 1 root root 0 Jul 19 00:16 LUN0
-rw-r--r--. 1 root root 0 Jul 19 00:16 LUN1 <-- New SAN Space - LUN1 - Controller A . Port0
-rw-r--r--. 1 root root 0 Jul 19 00:16 UTM_LUN31

/proc/mpp/DCS3700/controllerA/qla2xxx_h2c0t0:
total 0
-rw-r--r--. 1 root root 0 Jul 19 00:16 LUN0
-rw-r--r--. 1 root root 0 Jul 19 00:16 LUN1<-- New SAN Space - LUN1 - Controller A . Port1
-rw-r--r--. 1 root root 0 Jul 19 00:16 UTM_LUN31

```

```

/proc/mpp/DCS3700/controllerB:
total 0
dr-xr-xr-x. 2 root root 0 Jul 19 00:16 q1a2xxx_h1c0t0
dr-xr-xr-x. 2 root root 0 Jul 19 00:16 q1a2xxx_h2c0t1

/proc/mpp/DCS3700/controllerB/q1a2xxx_h1c0t0:
total 0
-rw-r--r--. 1 root root 0 Jul 19 00:16 LUN0
-rw-r--r--. 1 root root 0 Jul 19 00:16 LUN1<-- New SAN Space - LUN1 - Controller B. Port0
-rw-r--r--. 1 root root 0 Jul 19 00:16 UTM_LUN31

/proc/mpp/DCS3700/controllerB/q1a2xxx_h2c0t1:
total 0
-rw-r--r--. 1 root root 0 Jul 19 00:16 LUN0
-rw-r--r--. 1 root root 0 Jul 19 00:16 LUN1<-- New SAN Space - LUN1 - Controller B. Port1
-rw-r--r--. 1 root root 0 Jul 19 00:16 UTM_LUN31

```

---

In our example, we have a DCS3700 storage subsystem now with two LUNs, each mapped to this Linux host. The blade host has two HBAs, each mapped to both controllers for redundancy.

The Linux RDAC MPP driver configuration and advanced parameters are stored in the `/etc/mpp.conf` file. If any changes are made to the MPP configuration file `/etc/mpp.conf` or the persistent binding file `/var/mpp/devicemapping`, then the `mppUpdate` executable can be used to rebuild the RAMdisk.

The RDAC multipath driver consists of two parts. One part (`mppUpper`) is loaded before any HBA is loaded and prevents the HBA from advertising the LUNs to the system. The other part (`mppVhba`) is a virtual HBA on top of the real HBA, which is responsible for bundling the various paths to the same LUN to a single (multipath) device.

To dynamically reload the driver stack (`scsi_mod`, `sd_mod`, `sg`, `mpp_Upper`, `<physical HBA driver>`, and `mpp_Vhba`) without rebooting the system, complete the following steps:

1. Comment out all `scsi_hostadapter` entries in `/etc/module.conf`.
2. Run `modprobe -r mpp_Upper` to unload the driver stack.
3. Then, run `modprobe mpp_Upper` to reload the driver stack.

Reboot the system whenever you must unload the driver stack.

## Linux Device mapper multipathing driver (DMM / DM-Multipath)

You can use device mapper multipathing (DM-Multipath) to configure multiple I/O paths between server nodes and storage arrays in to a single device. These I/O paths are physical SAN connections that can include separate cables, switches, and controllers. Multipathing aggregates the I/O paths, creating a device that consists of the aggregated paths, providing redundant failover/failback support for the logical drives in the DCS3700 storage subsystems that are mapped to the Linux host server.

The DMM Driver can be used in the following ways:

- |                             |   |
|-----------------------------|---|
| <b>Redundancy</b>           | DM-Multipath can provide failover in an active/passive configuration. In an active/passive configuration, only half the paths are used at any time for I/O. If any element of an I/O path (the cable, switch, or controller) fails, DM-Multipath switches to an alternative path. |
| <b>Improved Performance</b> | DM-Multipath can be configured in active/active mode, where I/O is spread over the paths in a round-robin fashion. In some configurations, DM-Multipath can detect loading on the I/O paths and dynamically rebalance the load.   |

This section covers the DM-Multipath configuration for redundancy.

**Note:** DM-Multipath is only supported on IBM System Storage DS storage subsystems running firmware V7.77 and higher and using Red Hat Enterprise Linux 6 or Novell SUSE Linux Enterprise Server 11 SP1 as the operating system.

Red Hat Enterprise Linux 6 and Novell SUSE Linux Enterprise Server 11 SP1 include the DM-Multipath driver (DMM Driver) by default. In our example, we use Red Hat Enterprise Linux 6.

The DMM Driver and libraries are included in the installation media in the package directory as an RPM file. To check whether you have the DMM Driver and library dependencies installed, run the commands that are shown in Example 18-7.

*Example 18-7 Checking DM-Multipath Driver packages*

```
# rpm -qa | grep device-mapper-multipath
device-mapper-multipath-0.4.9-56.el6.x86_64
device-mapper-multipath-libs-0.4.9-56.el6.x86_64
```

If the DM-Multipath driver is not installed on your systems, run the .rpm installation commands, as shown in Example 18-8.

*Example 18-8 Basic DM-Multipath .rpm installation*

```
# mount /media/RHEL_6.2\ x86_64\ Disc\ 1/
# cd /media/RHEL_6.2\ x86_64\ Disc\ 1/
# cd Packages/
# pwd
/media/RHEL_6.2 x86_64 Disc 1/Packages
# rpm -ivh device-mapper-multipath-0.4.9-56.el6.x86_64
# rpm -ivh device-mapper-multipath-libs-0.4.9-56.el6.x86_64
```

**Note:** Before you start working with the DM-Multipath driver configuration, you must be sure that there is not any other multipath driver that is loaded on the system, such as LSI RDAC or any vendor HBA multipath driver.

### ***Basic DM-Multipath configuration (basic failover/redundancy)***

The DM-Multipath daemon requires a basic configuration file, which is stored in /dev/multipath.conf. You can set up multipath by using the mpathconf utility, which creates the multipath configuration file /etc/multipath.conf.

You can enable multipath support and create the basic configuration rules by running the commands that are shown in Example 18-9.

*Example 18-9 Enabling multipath on your system*

---

```
# mpathconf --enable --with_multipathd y
Starting multipathd daemon:                [ OK ]
# ll /etc/multipath.conf
-rw-----. 1 root root 2749 Jul 20 19:05 /etc/multipath.conf
```

---

After running these commands, you have the daemon loaded on your system and added to the init sequence.

The default settings for DM-Multipath are compiled in to the system and do not need to be explicitly set in this file. To check the DM-Multipath Status, run the command that is shown in Example 18-10.

*Example 18-10 Checking the DM-Multipath status*

---

```
# mpathconf
multipath is enabled
find_multipaths is enabled
user_friendly_names is enabled
dm_multipath module is loaded
multipathd is chkconfigured on
```

---

To discover the targeted devices, run the command that is shown in Example 18-11. The following example shows a 20 GB disk that is assigned, where the LUN ID is 360080e500023f8f4000009335004b137.

*Example 18-11 Running multipath queries by running the multipath command*

---

```
# multipath -ll
mpathb (360080e500023f8f4000009335004b137) dm-3 IBM,1818      FAStT
size=20G features='1 queue_if_no_path' hwhandler='1 rdac' wp=rw
|-+- policy='round-robin 0' prio=6 status=active
|  |- 2:0:0:0 sdc 8:32 active ready running
|  `-- 1:0:1:0 sde 8:64 active ready running
`-+- policy='round-robin 0' prio=1 status=enabled
   |- 1:0:0:0 sdb 8:16 active ghost running
   `-- 2:0:1:0 sdd 8:48 active ghost running
```

---

For troubleshooting purposes, DM-Multipath has an interactive interface console to the multipathd daemon. Running `multipathd -k` starts an interactive multipath console. After running this command, you can run `help` to get a list of available commands, you can enter an interactive command, or you can enter Ctrl-d to quit.

*Example 18-12 DM-Multipath interactive interface console*

---

```
show# multipathd -k
multipathd> help
multipath-tools v0.4.9 (04/04, 2009)
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 format $format
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
disablequeueing map|multipath $map
restorequeueing map|multipath $map
disablequeueing maps|multipaths
restorequeueing maps|multipaths
reinstate path $path
fail path $path
paths count
quit|exit
multipathd> show paths
hctl    dev dev_t pri dm_st  chk_st dev_st  next_check
0:0:0:0 sda 8:0   1  undef  ready  running orphan
1:0:0:0 sdb 8:16  1  active ghost  running XXXXXXXX.. 17/20
2:0:0:0 sdc 8:32  6  active ready  running XXXXXXXX.. 17/20
2:0:1:0 sdd 8:48  1  active ghost  running XXXXXXXX.. 17/20
1:0:1:0 sde 8:64  6  active ready  running XXXXXXXX.. 17/20
multipathd> show status
path checker states:
up           3
ghost       2

```

---

For more information about the DM-Multipath driver, its configuration, and troubleshooting it, go to the Red Hat documentation support website:

[http://docs.redhat.com/docs/en-US/Red\\_Hat\\_Enterprise\\_Linux/index.html](http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/index.html)

## 18.2 SAS storage configuration

This section describes how to add a logical drive to your Linux Server by using SAS attach and multipath drivers. This section also presents the basic actions that you must perform before and after you attach your DCS3700 storage subsystem to a physical server. For this configuration, we use an HS23 blade server that is running Red Hat Linux Enterprise Server 6 Update 2.

Before you map the logical drives to the server, consider the following items:

- ▶ Check at the IBM System Storage Interoperation Center (SSIC) website to be sure that it is a supported environment and also to be sure which multipath driver is needed. The SSIC provides the interoperability matrix, and is a selection driven tool where you provide the criteria for your configuration. At the time of writing, Linux rdac (LSI/Netapp) and native DMM are supported.

For more information, go to the following website:

<http://www-03.ibm.com/systems/support/storage/config/ssic>

- ▶ 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.

Verify the boot order in MPT SAS BIOS. If this setting is not correctly configured, the system tries to boot from the logical drive instead of the local hard disk.

In the following scenario, SAS zoning configuration is not necessary, unless you plan to configure your server to boot from SAN. The SAS I/O modules that we use in this example require a minor zone configuration. The configuration must be the same if you have more than one SAS I/O module that is connected.

## 18.2.1 Preparing for the installation

For our example setup, which is shown in Figure 18-4, we used the following hardware and software components:

- ▶ One HS23 blade server with a SAS Connectivity Card (CIOv)
- ▶ Two IBM BladeCenter SAS pass-through modules
- ▶ One DCS3700 storage subsystem
- ▶ Red Hat Enterprise Server 6 Update 2
- ▶ The latest version IBM RDAC drivers for Linux (formerly provided by LSI/Netapp).
- ▶ Linux RDAC multipath driver (MPP)




Figure 18-4 BladeCenter and DCS3700 - Linux Implementation SAS interconnection layout



## SAS HBA Drivers

Depending on your adapter model, the driver might be packaged with your Linux distribution. At the time of writing, the supported hardware configurations for the BladeCenter are shown in Figure 18-5.

<b>Product Family</b>	<b>Product Model</b>
IBM System Storage Midrange Disk	DCS3700
IBM System Storage Real-time Compression	DS3950
IBM System Storage SAN Volume Controller	DS4100
IBM System Storage Virtualization Engine for Tape	DS4200
<b>Product Version</b>	<b>Connection Protocol</b>
DCS3700 (07.77)	Fibre Channel
DCS3700 (07.83)	SAS
<a href="#">Export Selected Product Version (xls)</a>	
<b>Host Platform</b>	<b>Server Model</b>
IBM BladeCenter	IBM BladeCenter Servers (AMD)
	IBM BladeCenter Servers (Intel)
	IBM BladeCenter Servers (POWER)
<b>Operating System</b>	<b>Adapter (HBA, CNA, etc)</b>
Novell SUSE Linux Enterprise Server 11 SP1	IBM P/N 39Y9190
Red Hat Enterprise Linux 5.7	IBM P/N 43W4068
Red Hat Enterprise Linux 6.1	
Red Hat Enterprise Linux 6.2	
<b>Switch Module</b>	<b>SAN or Networking</b>
IBM BladeCenter SAS Connectivity Module (P/N 39Y9195)	
<b>Clustering</b>	<b>Multipathing</b>
<none>	<none>
Red Hat Cluster Suite	Linux Device Mapper Multipath (DMM)
SIOS LifeKeeper	Linux RDAC

Figure 18-5 SSIC supported configuration - firmware Version 7.83

For more information about the supported configuration, visit the IBM System Storage Interoperation Center (SSIC) at the following website:

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

## 18.2.2 Getting SAS Adapters address

Like any Host Bus Adapter, an SAS card has his own hardware address and there are different ways to identify those addresses.

To verify that the SAS expansion module is installed in the HS23 blade server, from the BladeCenter management console, click **Monitor** → **Hardware VPD (Vital Product Data)** → **Select your Blade Server bay** (in our case, we used bay 6), as shown in Figure 18-6.

Hardware Topology		Activity
<a href="#">Collapse all</a>   <a href="#">Expand all</a>		
Module Name	Module Description	Presence
Chassis and Chassis Managed Components		
<a href="#">Chassis</a>	BladeCenter-H	Installed
└─ [1] <a href="#">Media Module</a>	Media Tray	Installed
Blades		
└─ [1] Blade Bay	---	Not Installed
└─ [2] <a href="#">ADXblade1</a>	PS701/702 (Type 8406)	Installed
└─ [3] <a href="#">RHEL6.1 Server1</a>	HS22 (Type 7870)	Installed
└─ [4] Blade Bay	---	Not Installed
└─ [5] <a href="#">ADXblade2</a>	PS701/702 (Type 8406)	Installed
└─ [6] <a href="#">ITSO-Redhat</a>	HS23 (Type 7875)	Installed
└─ Processors		
└─ Memory		
└─ Storage		
└─ [1] Expansion Card	Expansion Card	Installed
└─ [1] Panel	Panel	Installed
└─ [1] Blade Expansion Module	---	Not Installed
└─ [2] Blade Expansion Module	---	Not Installed
└─ [3] Blade Expansion Module	---	Not Installed
└─ [4] Blade Expansion Module	---	Not Installed
└─ [1] <a href="#">Expansion Card</a>	SAS Expansion Option	Installed
└─ [1] High Speed Expansion Card	High Speed Expansion Card	Installed
└─ [1] Battery	---	Not Installed
└─ [1] Cable	---	Not Installed

Figure 18-6 BladeCenter Vital Product Data - hardware

After you identify the bay and the presence of an SAS expansion card, you can check the port address from the SAS Connectivity Module.

The SAS Connectivity Module can be reached by using the IP address that was previously configured in the BladeCenter. In this example, we have two SAS pass-through modules that are connected to I/O bay 3 and bay 4, as shown in Figure 18-7.

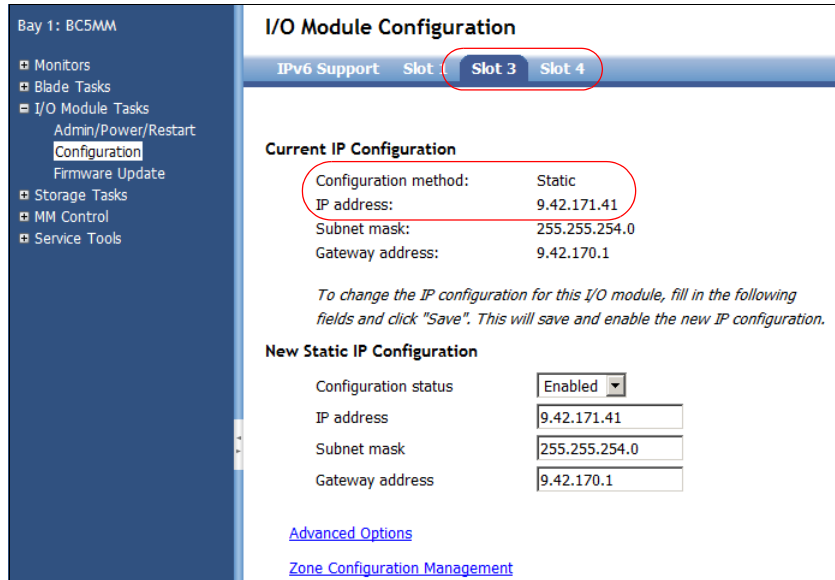


Figure 18-7 I/O Modules - network configuration

From the web browser, enter the IP address for your SAS Module and log in using your credentials. The default values are User: USERID / Password: PASSWORD (with a zero).

As shown in Figure 18-8, Blade Slot Connection 6 has an address assigned.

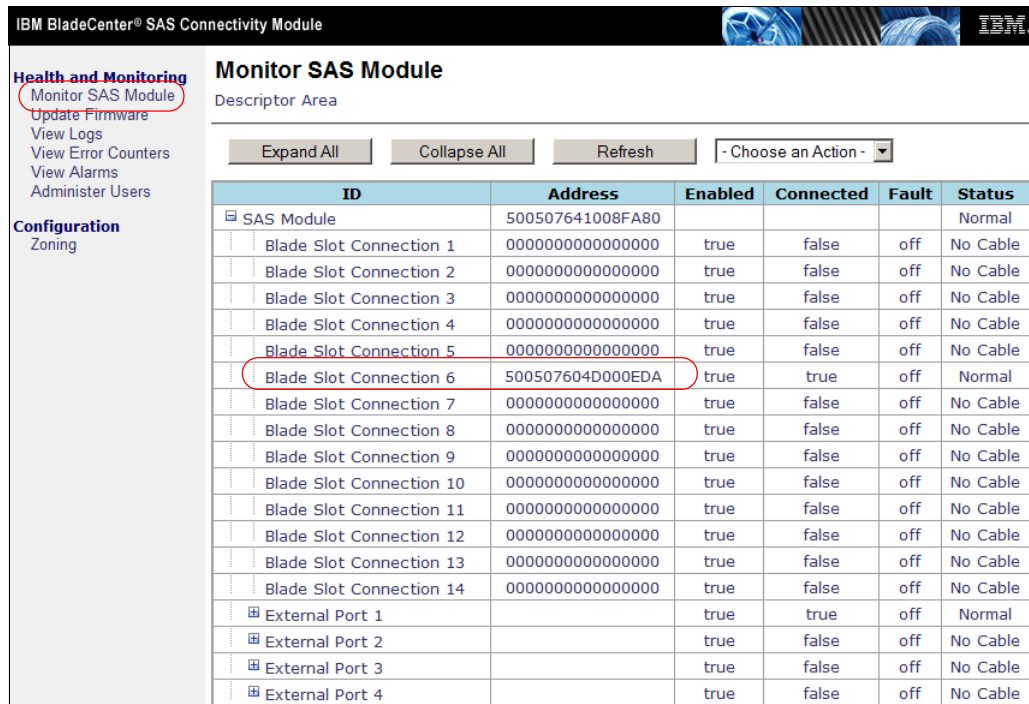


Figure 18-8 Discovering an SAS adapter address

Repeat the same procedure on the secondary SAS module to get the secondary SAS port address.

**Note:** This example assumes that you have port mapping that is configured on the storage side and the SAS module is connected to the DCS3700 Controller SAS Port.

If your DCS3700 storage subsystem is correctly connected and configured, and if you have some storage that is assigned, you should be able to discover the attached devices during the boot sequence, as shown in Figure 18-9.

```

LSI Corporation MPT SAS2 BIOS
MPT2BIOS-7.19.00.00 (2011.05.16)
Copyright 2000-2011 LSI Corporation.

```

PCI SLOT	ENCL SLOT	LUN NUM	VENDOR NAME	PRODUCT IDENTIFIER	PRODUCT REVISION	INT13 ENTRY	SIZE \ NUDATA
0	0	0	LSI	SAS2004-IR	10.00.08.00	0A:02:00:10	
0	0	0	IBM-ESXS	ST9146852SS	B62C	Boot	140 GB
0	28	31	IBM	Universal Xport	0777	81h	20 MB
0	28	31	IBM	Universal Xport	0777	82h	20 MB

```

LSI Corporation MPT2 boot ROM successfully installed!

```

Figure 18-9 Discovering disk space by using the SAS adapter during the boot stage

### 18.2.3 Installing the Linux multipath driver

The process for installing the RDAC multipath driver and native DMM support are similar to the Fibre Channel Implementation, which is described in the following sections:

- ▶ “Linux RDAC Multipath Driver” on page 528
- ▶ “Linux Device mapper multipathing driver (DMM / DM-Multipath)” on page 533

### 18.2.4 Discovering devices (disk drives)

New and existing devices can be discovered or listed by running the commands that are shown in Example 18-13.

Example 18-13 Discovering devices

```

# dmesg | grep -i LSISAS <----- Check if the SAS Driver module is loaded.
mpt2sas0: LSISAS2004: FWVersion(10.00.08.00), ChipRevision(0x03), BiosVersion(07.18.01.13)

# mppBusRescan <----- Rescan all Host Adapters.
scan mpt2sas HBA host /sys/class/scsi_host/host0...
no new device found
run /usr/sbin/mppUtil -s busscan...
scan mpp virtual host /sys/class/scsi_host/host1...
no new virtual device found

# multipath -ll <----- List logical volumes (by using the DMM driver).
mpathb (360080e500023f8f4000009335004b137) dm-3 IBM,VirtualDisk
size=30G features='0' hwhandler='0' wp=rw
~+- policy='round-robin 0' prio=1 status=active

  ~- 1:0:0:0 sdb 8:16 active ready running

# ls -lR /proc/mpp/ <----- List logical volumes (by using the RDAC driver).

```

```

/proc/mpp/:
total 0
dr-xr-xr-x. 4 root root 0 Jul 25 19:44 DCS3700 <----- DCS3700 storage subsystem
/proc/mpp/DCS3700:
total 0
dr-xr-xr-x. 3 root root 0 Jul 25 19:44 controllerA <----- DCS3700 Controller A
dr-xr-xr-x. 3 root root 0 Jul 25 19:44 controllerB <----- DCS3700 Controller B
-rw-r--r--. 1 root root 0 Jul 25 19:44 virtualLun0

/proc/mpp/DCS3700/controllerA:
total 0
dr-xr-xr-x. 2 root root 0 Jul 25 19:44 mpt2sas_h0c0t5

/proc/mpp/DCS3700/controllerA/mpt2sas_h0c0t5: <---- DCS3700 Controller A - Discovered Ports
total 0
-rw-r--r--. 1 root root 0 Jul 25 19:44 LUN0 <----- Presented SAN Space - Lun0 - Path A
-rw-r--r--. 1 root root 0 Jul 25 19:44 UTM_LUN31 <----- Access LUN

/proc/mpp/DCS3700/controllerB:
total 0
dr-xr-xr-x. 2 root root 0 Jul 25 19:44 mpt2sas_h0c0t4

/proc/mpp/DCS3700/controllerB/mpt2sas_h0c0t4: <---- DCS3700 Controller B - Discovered Ports
total 0
-rw-r--r--. 1 root root 0 Jul 25 19:44 LUN0 <----- Presented SAN Space - Lun0 - Path B
-rw-r--r--. 1 root root 0 Jul 25 19:44 UTM_LUN31 <----- Access LUN

```

---

## 18.3 Managing the disk space by using IBM DS Storage Manager and LVM

This section describes how to create a volume group and how to expand it by using native OS Logical Volume Manager (LVM) support.

First, check what space you have assigned on the storage subsystem by running `multipath -ll` on the server console. For more information, see Example 18-11 on page 535.

From the DS Storage Manager, click the **Storage & Copy Services** tab and click **Disk Pools**. Then, click LINUX\_LUN0, as shown in Figure 18-10.

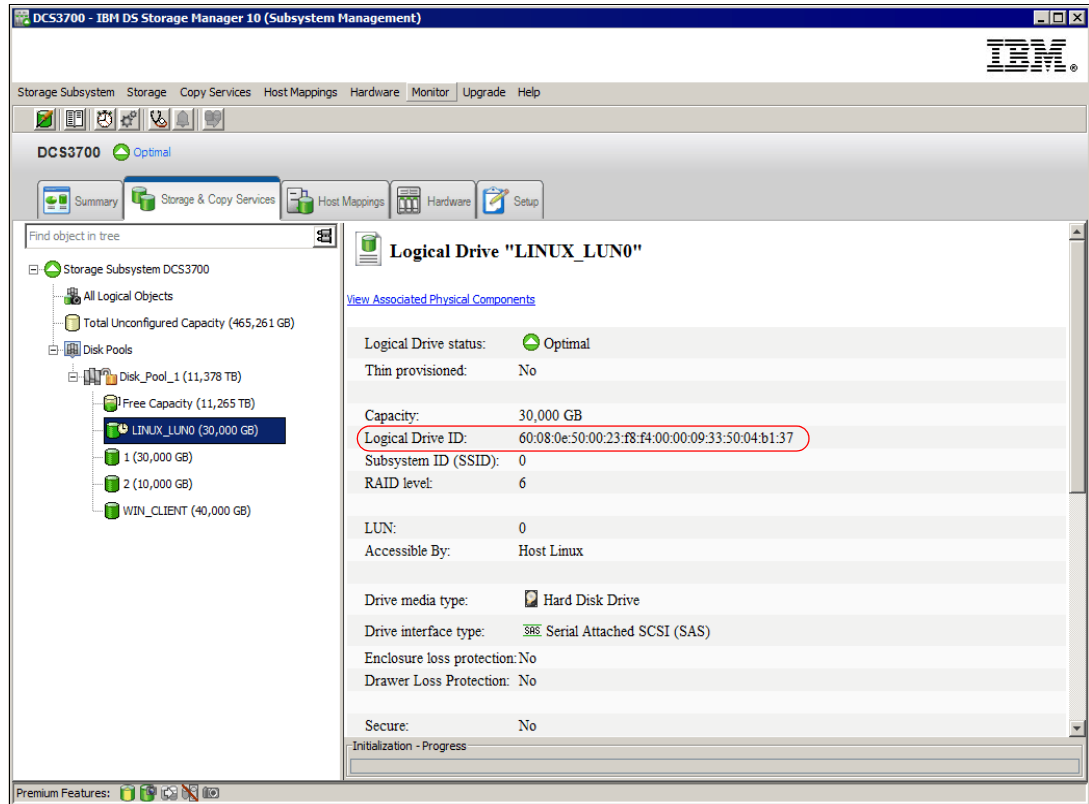


Figure 18-10 Matching disk space

For this test scenario, we assigned two disks with 20 GB that are configured with multipath support that is provided by the DM-Multipath driver (DMM). The `/dev/sda` device is the local disk that supports Red Hat Enterprise Linux V6. `/dev/sdb` and `/dev/dm-3` are the drives from the DSC3700 storage subsystem that are attached to the server.

To create an LVM Volume, run the commands that are shown in Example 18-14.

*Example 18-14 Creating an LVM Volume*

```
# cat /proc/partitions<----- Check active partitions.
major minor #blocks name

    8      0 143374000 sda
    8      1   204800 sda1
    8      2   512000 sda2
    8      3 142655488 sda3
  253     0 52428800 dm-0
  253     1   8224768 dm-1
  253     2   81997824 dm-2
    8     16 20971520 sdb
  253     3  20971520 dm-3

# pvcreate /dev/dm-3 <----- Create a physical volume, Initialize the disk.
Physical volume "/dev/dm-3" successfully created

# pvdisplay /dev/dm-3 <----- Check the physical volume.
"/dev/dm-3" is a new physical volume of "20.00 GiB"
--- NEW Physical volume ---
PV Name                /dev/dm-3
```

```

VG Name
PV Size          20.00 GiB
Allocatable      NO
PE Size          0
Total PE         0
Free PE          0
Allocated PE     0
PV UUID          qVtEww-foKI-moZb-wNSN-02tp-s8Wk-ganEj9

# vgcreate vg_DCS3700_Linux /dev/dm-3 <----- Create a volume group.
Volume group "vg_DCS3700_Linux" successfully created

# vgsdisplay vg_DCS3700_Linux <----- Check the volume group.
--- Volume group ---
VG Name          vg_DCS3700_Linux
System ID
Format           lvm2
Metadata Areas   1
Metadata Sequence No 1
VG Access        read/write
VG Status        resizable
MAX LV           0
Cur LV          0
Open LV          0
Max PV           0
Cur PV          1
Act PV           1
VG Size          20.00 GiB
PE Size          4.00 MiB
Total PE         5119
Alloc PE / Size  0 / 0
Free PE / Size   5119 / 20.00 GiB
VG UUID          cCm6bZ-DVCi-3NCO-fuax-jutX-ffqr-SqMzGj

# lvcreate -L 19.99GiB vg_DCS3700_Linux <----- Create a logical volume.
Rounding up size to full physical extent 19.99 GiB
Logical volume "lvo10" created

# lvdisplay vg_DCS3700_Linux <----- Check the logical volume.
--- Logical volume ---
LV Path          /dev/vg_DCS3700_Linux/lvo10
LV Name          lvo10
VG Name          vg_DCS3700_Linux
LV UUID          m3Nf5V-VWIU-JDKN-3xTw-vKiI-10lq-167iYZ
LV Write Access  read/write
LV Creation host, time server03, 2012-07-23 18:11:44 -0400
LV Status        available
# open           0
LV Size          19.99 GiB
Current LE       5118
Segments         1
Allocation       inherit
Read ahead sectors auto
- currently set to 256
Block device     253:4

```

After the SAN Space is being managed by LVM, you can format and mount the file system to the server by running the commands that are shown in Example 18-15.

---

*Example 18-15 Formatting and mounting disk partitions*

```

mkfs.ext4 /dev/vg_DCS3700_Linux/lvo10 <----- Format the logical volume.
mke2fs 1.41.12 (17-May-2010)
Filesystem label=
OS type: Linux

```

```

Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
1310720 inodes, 5240832 blocks
262041 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=4294967296
160 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done

```

This filesystem will be automatically checked every 27 mounts or 180 days, whichever comes first. Use tune2fs -c or -i to override.

```

# mkdir /opt/DCS3700_Space <----- Create a mount point.

# mount /dev/vg_DCS3700_Linux/lvo10 /opt/DCS3700_Space/ <----- Mount the file system.

# df -h <----- Check the mounted file system's capacity.
Filesystem      Size  Used Avail Use% Mounted on
/dev/mapper/vg_server03-lv_root
                50G  9.7G   38G  21% /
tmpfs            7.8G  88K   7.8G   1% /dev/shm
/dev/sda2        485M   70M  390M  16% /boot
/dev/sda1        200M  256K  200M   1% /boot/efi
/dev/mapper/vg_server03-lv_home
                77G  186M   73G   1% /home
/dev/mapper/vg_DCS3700_Linux-lvo10
                20G  172M   19G   1% /opt/DCS3700_Space

```

To extend an existing volume group by adding extra free space, run the commands that are shown in Example 18-16. The goal is to extend the LUN to 30 GB and resize the space dynamically on an existing file system that is running on an LVM.

*Example 18-16 Extending LVM and file system capacity*

```

# umount /opt/DCS3700_Space/ <----- Unmount the file system.

# pvresize /dev/dm-3 <----- Resize the physical volume (previously added to your VG.
Physical volume "/dev/mapper/mpathb" changed
1 physical volume(s) resized / 0 physical volume(s) not resized

# vgdisplay vg_DCS3700_Linux <----- Check the VG.
--- Volume group ---
VG Name          vg_DCS3700_Linux
System ID
Format           lvm2
Metadata Areas   1
Metadata Sequence No 6
VG Access        read/write
VG Status        resizable
MAX LV           0
Cur LV          1
Open LV          0
Max PV           0
Cur PV          1
Act PV           1
VG Size          30.00 GiB
PE Size          4.00 MiB

```



```

Total PE          7679
Alloc PE / Size  5118 / 19.99 GiB
Free PE / Size   2561 / 10.00 GiB
VG UUID          cCm6bZ-DVCi-3NCO-fuax-jutX-ffqr-SqMzGj

# lvextend -L +9.99GiB /dev/vg_DCS3700_Linux/lvo10 <----- Extend the logical volume.
Rounding size to boundary between physical extents: 9.99 GiB
Extending logical volume lvo10 to 29.98 GiB
Logical volume lvo10 successfully resized

# lvdisplay /dev/vg_DCS3700_Linux/lvo10
--- Logical volume ---
LV Path            /dev/vg_DCS3700_Linux/lvo10
LV Name            lvo10
VG Name            vg_DCS3700_Linux
LV UUID            m3Nf5V-VWIU-JDKN-3xTw-vKiI-101q-167iYZ
LV Write Access    read/write
LV Creation host, time server03, 2012-07-23 18:11:44 -0400
LV Status          available
# open             0
LV Size            29.98 GiB
Current LE         7676
Segments           1
Allocation         inherit
Read ahead sectors auto
- currently set to 256
Block device       253:4

# resize2fs -p /dev/vg_DCS3700_Linux/lvo10 <----- Resize the file system/
resize2fs 1.41.12 (17-May-2010)
Filesystem at /dev/vg_DCS3700_Linux/lvo10 is mounted on /opt/DCS3700_Space; online resizing required
old desc_blocks = 1, new_desc_blocks = 2
Performing an online resize of /dev/vg_DCS3700_Linux/lvo10 to 5238784 (4k) blocks.
The filesystem on /dev/vg_DCS3700_Linux/lvo10 is now 5238784 blocks long.

# df -h <----- Check the mounted file system's capacity.
Filesystem          Size  Used Avail Use% Mounted on
/dev/mapper/vg_server03-lv_root
                    50G  9.7G  38G  21% /
tmpfs                7.8G  88K  7.8G   1% /dev/shm
/dev/sda2            485M   70M  390M  16% /boot
/dev/sda1            200M  256K  200M   1% /boot/efi
/dev/mapper/vg_server03-lv_home
                    77G  186M  73G   1% /home
/dev/mapper/vg_DCS3700_Linux-lvo10
                    30G  172M  29G   1% /opt/DCS3700_Space

```

Because the goal of Example 18-16 on page 546 is to show the basic configuration steps for attaching SAN space and managing the space dynamically by using LVM, some OS recommendations were excluded:

- ▶ Label the logical volume by running the `e2label` command.
- ▶ Mount the file system permanently by adding the mount point to the `/etc/fstab` file.
- ▶ Run `blkid` to get more disk information, such as the UUID. You should have an UUID when you mount the partition in `/etc/fstab` (for stability).

For more information about using LVM and Linux commands, visit your vendor support documentation portal and `man` pages.

## 18.4 Troubleshooting guide

Here are a few common problems and ways to troubleshoot them.

### 18.4.1 `rescan_scsi_bus.sh` is not showing any new device

If the server did not recognize any newly mapped LUNs after running the `rescan_scsi_bus.sh` command a few times, reboot the server.

### 18.4.2 `multipath -ll` output does not show all the LUNs.

This error occurs because the Device Mapper probably is not seeing all the LUNs that are mapped to it. This can happen because of various reasons,

If mappings are edited on the host that is connected to the array, the host might fail to monitor for newly added LUNs and there will not be a SCSI block device that is created for these newly added LUNs. To fix this issue, rescan for the SCSI devices by running `rescan_scsi_bus`. This should create SCSI device for the newly added LUNs. Run `multipath` to add the SCSI devices to the device mapper table.

### 18.4.3 `multipath` does not show paths that are equal to the number of LUNs that are mapped

The first thing that you should check is the `lsscsi` output to see whether the SCSI midlayer is seeing (# of paths x # of LUNs) number of block devices. If not, run `rescan-scsi-bus.sh`. This should fix the issue. If not, reboot the host.

If `lsscsi` is seeing all the devices, then complete the following steps:

1. Run `multipath` so that device mapper can add the devices that are seen by `lsscsi`.
2. Flush all the mappings and readd them by running `multipath -F` and `multipath`.
3. Restart the multipath daemon by running `service multipathd restart`.

### 18.4.4 `multipath -ll` output is showing the paths as [failed][faulty]

This error can happen if the physical path to a controller fails, or if the controller itself fails and failover happens. To correct this error, complete the following steps:

1. Verify the disk UID and parameters in the `multipath.conf` file.
2. Verify the cables and zoning configuration.
3. If this error happens after an error in test and you want to make the configuration optimal, save the logs first, flush and readd multipath devices, and restart the multipath daemon. If the issue still persists, reboot the system.



# IBM AIX V7.1 configuration guide

This chapter describes sample configurations that you can use to configure AIX V7.1 to access a logical drive from an IBM System Storage DCS3700 storage subsystem. For this sample Fibre Channel (FC) configuration, we use an IBM Power 550 Express server with a dual-port FC adapter (in our case, Feature Code 5447) and two FC switches that are connected to a DCS3700 storage subsystem.

AIX V7, and AIX V6 and IBM AIX 5L™ V5 (which are all supported AIX versions), are award winning operating systems that deliver superior scalability, reliability, and manageability.

AIX runs across the entire range of IBM POWER®, from entry level servers to powerful supercomputers. With continuous development of Power Systems and their hardware capabilities, the latest version of AIX V7.1 further uses features and benefits of intensive computing, highly available, or fault-tolerant systems.

In addition, AIX has an excellent history of binary compatibility, which provides assurance that your critical applications continue to run as you upgrade to newer versions of AIX.

This chapter covers the following topics:

- ▶ Planning for the installation
- ▶ Configuring MPIO
- ▶ Setting up the DCS3700 logical drives and host mapping
- ▶ Scanning and managing the storage logical drive from AIX
- ▶ SAS attachment

**Note:** At the time of writing, only an FC configuration is supported on AIX with storage firmware Version 07.83.22.00.

An SAS configuration with AIX is supported by storage firmware Version 07.77.xx.xx.

Other supported configurations depend on the type of controller that you have. For more information, see the IBM SSIC website for the latest support matrix.

## 19.1 Planning for the installation

Proper planning is always a crucial step in every configuration, and here are some important considerations:

- ▶ AIX host attachment to the DCS3700 storage subsystem requires an additional purchase of the AIX operating system Host Kit Option or Feature Code. The AIX Host Kit Option contains the required IBM licensing that is needed to attach an IBM AIX host system to the DCS3700 storage subsystem. For purchasing information, contact your IBM service representative or IBM reseller.
- ▶ 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 DCS3700 storage subsystem. For detailed HBA and systems information regarding combinations, go to the System Storage Interoperation Center (SSIC) at the following website:

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

The SSIC provides the interoperability matrix, which is a selection driven tool where you provide the criteria for your configuration, as shown in Table 19-1.

*Table 19-1 SSIC options for AIX interoperability matrix - required information*

Option type	Select options
Product Family	IBM System Storage Midrange Disk.
Product Model	DCS3700.
Product Version	DCS3700 (select firmware version).
Connection Protocol	Fibre Channel, SAS.
Host Platform	IBM Power Systems.
Operation Systems	AIX (select version and TL level).
HBA Model	Select HBA model.
SAN Vendor	Any, Cisco, Brocade, McData (select vendor).
SAN or Networking Model	Select SAN Switch model.
Clustering	IBM PowerHA® (select version).
Multipathing	IBM MPIO, IBM SDD PCM (select version).

To complete the SSIC interoperability matrix, you must have all the required information. On new IBM POWER6® and IBM POWER7® systems, use the latest firmware versions. For existing systems, you must determine the prerequisites and upgrade paths to achieve a supported system.

For further interoperability information, go to the following website:

<http://www.ibm.com/systems/storage/product/interop.html>

Always make sure that the HBA that you plan to use is at the supported level for the model and firmware version that is installed on your DCS3700 storage subsystem and for your version of AIX. You can check the AIX version by running the command that is shown in Example 19-1.

*Example 19-1 Check AIX version*

---

```
# oslevel -s  
7100-00-03-1115
```

---

To verify the microcode level on the HBAs, run the following commands:

- ▶ **lscfg -v1 fcs0** (The Z9 or ZA fields contains the firmware level)
- ▶ **lsmcode -d -d fcs0**
- ▶ **fcstat fcs0**

To map the logical drive to the host and to configure the zoning, obtain the WWN of each HBA port (fcsX). Example 19-2 shows how to get the WWN on Virtual I/O Server (VIOS), and Example 19-3 shows how to get it on AIX:

*Example 19-2 WWN from VIOS*

---

```
$ lsdev -dev fcs0 -vpd  
fcs0          U78A5.001.WIH1752-P1-C10-T1  FC Adapter  
  
Network Address.....1000000C9721CCC  
ROS Level and ID.....02E82774  
Device Specific.(Z0).....2057706D  
Device Specific.(Z1).....00000000  
Device Specific.(Z2).....00000000  
Device Specific.(Z3).....03000909  
Device Specific.(Z4).....FFE01231  
Device Specific.(Z5).....02E82774  
Device Specific.(Z6).....06E12715  
Device Specific.(Z7).....07E12774  
Device Specific.(Z8).....2000000C9721CCC  
Device Specific.(Z9).....ZS2.71X4  
Hardware Location Code.....U78A0.001.DNWGPL0-P1-C3-T1
```

PLATFORM SPECIFIC

```
Name: fibre-channel  
Model: 10N7255  
Node: fibre-channel@0  
Device Type: fcp  
Physical Location: U78A0.001.DNWGPL0-P1-C3-T1
```

---

*Example 19-3 WWN from AIX*

---

```
# lscfg -v1 fcs0  
fcs0          U78A0.001.DNWGPL0-P1-C3-T1  4Gb FC PCI Express Adapter  
(df1000fe)  
  
Part Number.....10N7255  
Serial Number.....1F7510C950  
Manufacturer.....001F
```

```
EC Level.....A
Customer Card ID Number.....5774
FRU Number..... 10N7255
Device Specific.(ZM).....3
Network Address.....10000000C9721CCC
ROS Level and ID.....02E82774
Device Specific.(Z0).....2057706D
Device Specific.(Z1).....00000000
Device Specific.(Z2).....00000000
Device Specific.(Z3).....03000909
Device Specific.(Z4).....FFE01231
Device Specific.(Z5).....02E82774
Device Specific.(Z6).....06E12715
Device Specific.(Z7).....07E12774
Device Specific.(Z8).....20000000C9721CCC
Device Specific.(Z9).....ZS2.71X4
Device Specific.(ZA).....Z1F2.70A5
Device Specific.(ZB).....Z2F2.71X4
Device Specific.(ZC).....00000000
Hardware Location Code.....U78A0.001.DNWGPL0-P1-C3-T1
```

---

Unlike IBM System x (where you enable the HBA BIOS), Power Systems do not require any particular configuration for you to discover the logical drive.

### 19.1.1 Zoning considerations

A SAN zone is collection of WWNs that can communicate with each other. A zone is a group of fabric-connected devices that are arranged in to a specified group. Zones can vary in size depending on the number of fabric-connected devices, and devices can belong to more than one zone. Only members of one zone can communicate with each other. Only members of one zone can communicate with each other, and one port can be a member of multiple zones.

**Note:** Using zoning is a good idea for 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 HBA2 and controller B, and then create additional zones for access to other resources. For a server with single HBA, consider using a separate zone for controller A and a separate zone for controller B.

**Note:** If there are different traffic types (such as disk and tape) or different vendor environments, use separate HBAs.

There are two different types of zoning. You can use hard (switch domain and port number) zoning and soft (port WWN, which is also known as WWPN) zoning. These two types should not be mixed inside one fabric.

For more information about zoning, see *Implementing an IBM/Cisco SAN*, SG24-7545 and *Implementing an IBM b-type SAN with 8 Gbps Directors and Switches*, SG24-6116.

Here are the general zoning rules:

- ▶ Dual fabric: Wet up your zoning to avoid a single point of failure.
- ▶ Single initiator zoning: A zone with one host port and one or more storage ports.
- ▶ Naming conventions (for fabric members and zones): Always use a descriptive name for WWPNs.
- ▶ Configuration backup: Always keep latest updated SAN switch configuration because if there are any issues with the switch, recovery is much faster.
- ▶ Always check the interoperability matrix: Before you make any firmware updates in the SAN environment, always check the interoperability dependencies for all SAN components.
- ▶ Maintenance window: It is important to do the SAN firmware upgrade during the scheduled maintenance window when traffic is at a minimum to reduce any possible performance issue in your SAN environment.
- ▶ Firmware upgrade: Keep your fabric updated according to the interoperability matrix because new firmware brings new features and provides fixes.

## 19.1.2 Planning MPIO

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

**Note:** At the time of writing, even though both MPIO and RDAC were available, RDAC is being replaced by MPIO. Therefore, this chapter covers only 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 healthcheck capability that can be used to do the following tasks:

- ▶ Check the paths and determine which paths can be used to send 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 unused paths that might be used if a failover occurs (for example, when the algorithm attribute value is failover, the health check can test the alternative paths).

This device driver is packaged with the AIX set of installation media and is installed with AIX if the POWER system has HBAs physically on the machine or LPAR when AIX was installed. If the HBAs are physically installed afterward or dynamically associated with the LPAR, then there is a possibility that all licensed programs products (LPPs) and prerequisites for the MPIO (and the HBA device drivers) are not installed. If so, the AIX installation media might be required when you run `cfgmgr`. This command automatically configures the HBAs to AIX and, using 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 CD-ROM, and run `cfgmgr -i /dev/cd0`.

You might be prompted to insert another CD-ROM to complete the configuration, and then you can reapply the Technology Level (TL) package again, which contains updates to the base level file sets installed on the CD media. To do this task, follow your normal procedure of applying TL maintenance to AIX.

You can verify that the MPIO driver packages are correctly installed, as shown in Example 19-4.

*Example 19-4 Verify that the following packages are installed*

---

```
# lspp -al | grep -i mpio
  devices.common.IBM.mpio.rte
                                7.1.0.15  COMMITTED MPIO Disk Path Control Module
  devices.common.IBM.mpio.rte
                                7.1.0.15  COMMITTED MPIO Disk Path Control Module
```

---

### **Planning for single Virtual I/O Server**

Figure 19-1 on page 555 shows the recommended scenario for a single VIOS solution. This scenario was not supported previously for RDAC because each HBA controller did not have access to both storage controllers. The VIOS 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 that dedicated to each client partition.



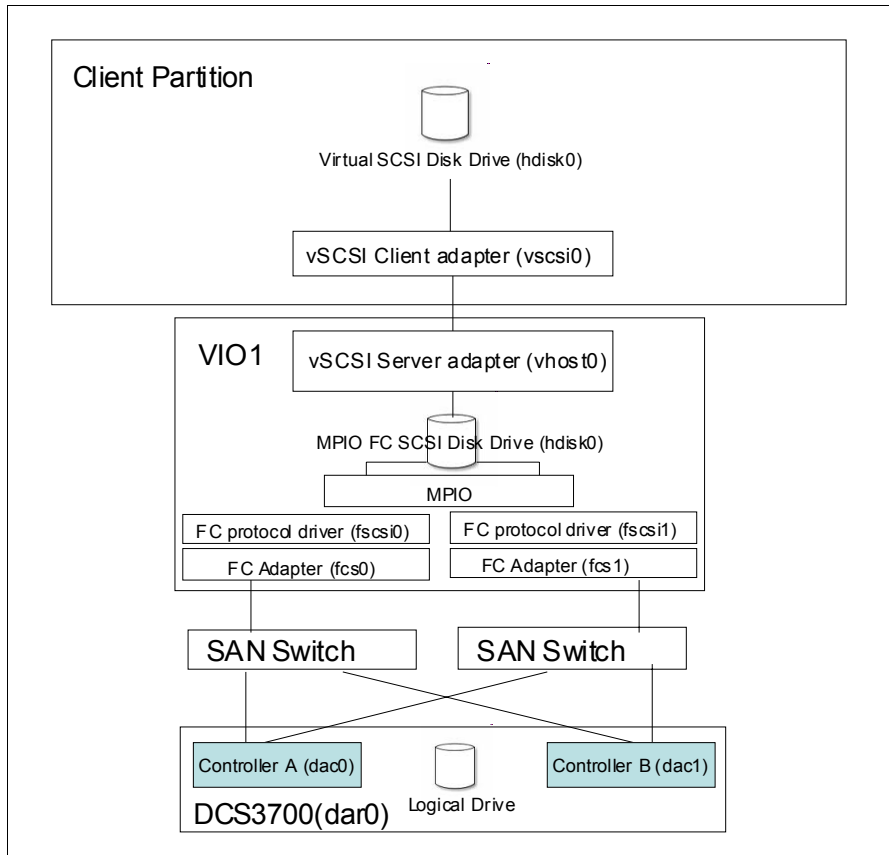


Figure 19-1 Single VIOS configuration with MPIO

The VIOS can 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 VIOS is not limited to using just one type of backing storage device.

If you use logical volumes for the backing storage device on the VIOS, a volume group (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 are mapped to the client partitions as virtual disks, with one logical volume for each virtual disk. When you create the logical volume, the size that is specified for the logical volume becomes the size of the virtual disk, which means you have 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 you use a whole physical disk for a backing device for your virtual disks, it is necessary to use a disk that the VIOS has not assigned to a volume group, which allows client partitions to share the physical storage controller and to use full performance of a physical disk. Failure of a single disk impacts only the single client partition that is that disk.

If you use SAN storage, the VIOS can use a LUN for the storage backing device. You can use the SAN environment to provision and allocate LUNs, and the VIOS 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 that is configured as a RAID array can be an effective solution. It is preferable to map SAN-attached LUNs as physical volume backing storage rather than configuring volume groups and logical volumes on top of them.

If you use the Integrated Virtualization Manager instead of the Hardware Management Console (HMC), 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 IBM BladeCenter environments.

### Planning for dual Virtual I/O Servers

Figure 19-2 shows the optimal scenario for a dual VIOS solution. This configuration provides multiple paths to the data at the client partition and at the VIOS. In addition, two SAN switches are included in the configuration to provide redundancy in the SAN.

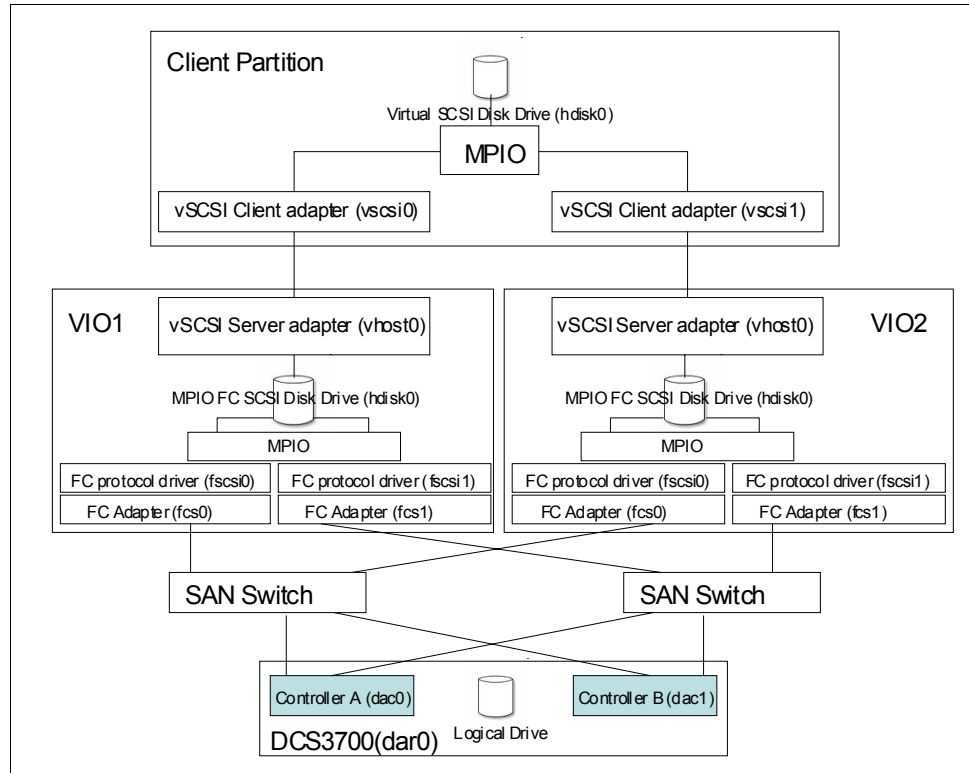


Figure 19-2 Dual VIOS configuration with MPIO

Although this configuration illustrates the virtualization of a single client disk, it can be extended to multiple disks if required. Each VIOS uses multiple paths to the SAN LUNs through two host bus adapters, each connected to a different SAN switch, which allows the VIOS 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 VIOS. The client partition uses MPIO to provide separate paths to the data disk.

This configuration is robust, but is more complicated than the other configurations. It is necessary to set up MPIO on the client. Two VIOs must be installed and configured. Two or more host bus adapters are needed in each VIOS partition. The VIOS 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 VIOS for a code update without interrupting the client partitions.

Even a SAN switch failure cannot prevent the client partitions from getting to their data. Finally, by using RAID in the SAN LUNs, there is redundancy if there is a disk failure. Attach an FC storage subsystem to an AIX V 7.1 Server.

In this scenario, because the disks are accessed through both VIOs concurrently, the `reserve_policy` for each disk must be set to `no_reserve` on both VIOs. Change the `reserve_policy` attribute to `no_reserve` by running `chdev`, as shown in Example 19-5.

*Example 19-5 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, you must add the `-perm` parameter at the end. This parameter updates ODM and the setting takes effect on the next reboot.

Run `lsdev` to make sure that the `reserve_policy` attribute is set to `no_reserve`, as shown in Example 19-6. If you used the `-perm` parameter, the `reserve_policy` is set to `no_reserve`, but a reboot is needed to apply the change.

*Example 19-6 `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
ww_name	0x20480080e51b0c90	FC World Wide Name	False

## 19.2 Configuring MPIO

The following parameters must be changed when you use MPIO with physical devices.

**Important:** Because these attributes cannot be changed while the device is in an 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 the `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. Run `chdev`, as shown in Example 19-7 and Example 19-8, to change the attributes.

*Example 19-7 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 19-8 Changing the attributes on AIX*

```
# chdev -l fscsi0 -a fc_err_recov=fast_fail -a dyntrk=yes -P
fscsi0 changed
```

It is important to change these two parameters because of the following reasons:

- `dyntrk`: The default for this setting is `no`. When it is set to `yes`, it enables 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 that is 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 controls 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 alternative paths when the 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, approximately 15 seconds, so that the fabric can stabilize. This situation 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 tries 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 DCS3700 controller host port.

- ▶ To enable load balancing across multiple Fibre Channel adapters within the VIOs when you use the base AIX MPIO support, set the algorithm to `round_robin` for each physical disk, as shown in Example 19-9 and Example 19-10.

**Note:** It is necessary to set the `reserve_policy` to `no_reserve` (as shown in Example 19-5 and Example 19-6) before setting the algorithm to `round_robin`.

*Example 19-9 Setting the `round_robin` parameter on VIOS*

---

```
$ chdev -dev hdisk2 -attr algorithm=round_robin -perm
hdisk2 changed
```

---

*Example 19-10 Setting the `round_robin` parameter on AIX*

---

```
# chdev -l hdisk2 -a algorithm=round_robin -P
hdisk2 changed
```

---

**Note:** MPIO for virtual SCSI devices only supports failover mode in AIX.

- ▶ Enable the healthcheck mode for the disk so that the status of the disks is updated automatically. Healthcheck mode is disabled by default (`hcheck_interval=0`). If it is disabled, the client partition does not update the path status if there is a failure of the active path. To activate the healthcheck function, run `chdev`, as shown in Example 19-11 and Example 19-12. In this example, we use a healthcheck interval of 60 seconds. To check for the attribute setting, run `lsattr`.

**Note:** The path switching also works if the `hcheck_interval` attribute is disabled, but it is recommended to have the status updated automatically.

*Example 19-11 Changing the health interval parameter on VIOS*

---

```
$ chdev -dev hdisk2 -attr hcheck_interval=60 -perm
hdisk2 changed
```

---

*Example 19-12 Changing the health interval parameter on AIX*

---

```
# chdev -l hdisk2 -a hcheck_interval=60 -P
hdisk2 changed
```

---

All these parameters can also be changed by using `smitty mpio` and clicking MPIO Device Management and then Change/Show MPIO Devices Characteristics.

## 19.3 Setting up the DCS3700 logical drives and host mapping

Map the logical drives to the AIX operating system by using the procedure that is described in Chapter 10, “Administration: Mappings tab” on page 327.

Use the correct Host type when mapping the logical drive to the host from the IBM Storage Manager, as shown in Figure 19-3.

When you define the host port, the host type of the attached host is defined as well. Using this information, the DS Storage System modifies the NVSRAM settings to behave according to the host type that is defined.

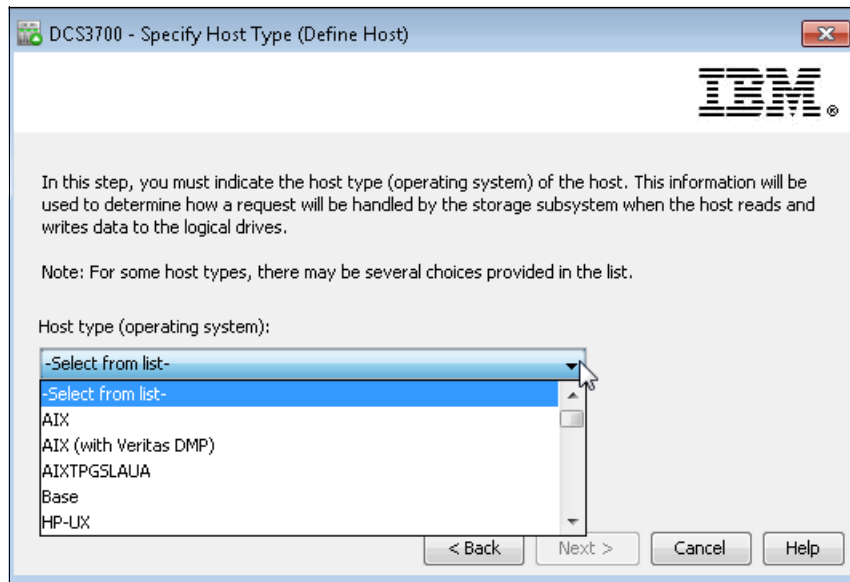


Figure 19-3 Defining the host type of the host

The most common of the specifications that differ according to the host type setting is the Auto Volume Transfer (AVT) function, which is enabled for certain host types, and disabled for other host types. It is important to choose carefully the correct host type from the list of available types because this is the part of the configuration that deals 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 a loss of the path failover function when attached to the storage server.

For the AVT function, use the following settings:

- ▶ AIX: For RDAC and MPIO, always use the AIX host type.
- ▶ AIX (with Veritas DMP): This host type is used when third-party multipath drivers are installed, such as using Dynamic Multi-pathing (DMP) as the default multipath driver. DMP is part of VERITAS Volume Manager.
- ▶ AIXTPGSLAUA: At the time of writing, this host type was still in the testing phase. If testing goes well, AIX will support the Asymmetric Logical Unit Access (ALUA) failover method in future releases of the controller firmware.

**Note:** Path failover is not supported during the AIX boot process. After an AIX host boots, failover operates normally.

## 19.4 Scanning and managing the storage logical drive from AIX

In this section, it is assumed that the logical drive is mapped to the system. For a detailed procedure, see 8.6.1, “Creating a logical drive” on page 246.

To discover new devices on AIX, run `cfgmgr`. This command configures devices and optionally installs those devices’ software in to the system.

To discover new devices on VIOS, run `cfgdev`. This command performs the same function as `cfgmgr` and is only applicable to VIOS.

Example 19-13 shows the disk that we configured for the DCS3700 storage subsystem. The third column shows the location code, which in our example is 02-00-02. 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 the device locations.

### Example 19-13 Disk status and location codes

---

```
# lsdev | grep -i mpio
hdisk1    Available 02-00-02    MPIO DS5080 Disk
```

---

Run `mpio_get_config` to see all the available paths to the logical drive, as shown in Example 19-14. 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 VIOS. To use the command, run `oem_setup_env` first to start a root AIX shell. Then, the `mpio_get_config -Av` command is available.

**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 19-14 `mpio_get_config` output

---

```
# mpio_get_config -Av
Frame id 0:
  Storage Subsystem worldwide name: 608e50023f8f400004fffe472
  Controller count: 2
  Partition count: 1
  Partition 0:
  Storage Subsystem Name = 'DCS3700'
    hdisk#      LUN #  Ownership      User Label
    hdisk1      0      B (preferred)  1
```

---

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 displays the information in columnar form. When no flags are specified that qualify the paths to display, the three columns are status, device, and parent. Example 19-15 shows the output for a dual controller DCS3700 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 19-15 lspath command output*

---

```
# lspath
Enabled hdisk0 vscsi0
Enabled hdisk1 fscsi0
Enabled dac0 fscsi0
Enabled hdisk1 fscsi0
Enabled dac1 fscsi0
Enabled dac1 fscsi1
Enabled dac0 fscsi1
```

---

Here are the values that can appear in the status column:

- ▶ **enabled**: Indicates that the path is configured and operational. It is considered when the paths are selected for I/O.
- ▶ **disabled**: Indicates that the path is configured, but not currently operational. It is manually disabled and is not considered when paths are selected for I/O.
- ▶ **failed**: Indicates that the path is configured, but it has IO failures that rendered it unusable. It is not considered when paths are selected for I/O.
- ▶ **defined**: Indicates that the path is not configured in 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 have this status only during boot, so this status should never appear as a result of the **lspath** command.

## 19.4.1 Ways to manage the paths

There are three commands that are available to manage the paths:

- ▶ **chpath**: The **chpath** command either changes the operational status of paths to the specified device or it changes one or more attributes that are associated with a specific path to the specified device. The required syntax is different depending upon the change that is made.
- ▶ **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.
- ▶ **rmpath**: The **rmpath** command unconfigures, and possibly undefines, one or more paths that are 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 results in all paths that are associated with the device being unconfigured or undefined, the command exits with an error and without unconfiguring or undefining any path. In this situation, the **rmdev** command must be run 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 definitions from the system.



All these commands, including **lspath**, are in the following menu:

```
# smitty mpio
```

In fact, with this menu, it is also possible to set the priority for each device's path by running **# smitty mpio** and then selecting 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 then type or select a valid MPIO path. The panel that is shown in Example 19-16 opens.

*Example 19-16 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	

---

## 19.5 SAS attachment

This section describes how to add a logical drive that is configured on a DCS3700 storage subsystem to an AIX V7.1 operating system that is running on an IBM BladeCenter PS701 blade server that is connected with a SAS Connectivity Module.

**Note:** At the time of writing, the SAS Configuration with AIX is supported by storage firmware Version 07.77.xx.xx.

### 19.5.1 Equipment required

For this example configuration, we use following hardware and software:

- ▶ IBM BladeCenter H
- ▶ IBM BladeCenter PS701 server with AIX V7.1 installed on a local disk
- ▶ Two IBM BladeCenter SAS Connectivity Modules
- ▶ IBM BladeCenter SAS connectivity card CIOv FC8246
- ▶ A DCS3700 storage subsystem within controllers that have the host SAS connectivity option

- ▶ Two SAS cables
- ▶ The latest version of DS Storage Manager running on an external management workstation (at the time of writing, this is Version 10.83.xx.18)

The connectivity scheme is shown in Figure 16-1 on page 482.

## 19.5.2 IBM BladeCenter setup

For latest DCS3700 storage subsystem interoperability matrix, which ensures that the hardware that is used for this implementation is supported, check the IBM System Storage Interoperation Center, found at the following website:

<http://www-947.ibm.com/support/entry/portal/docdisplay?brand=5000020&indocid=MIGR-5073016>

**Note:** Ensure that all the hardware and software that is used in your configuration is at the latest supported firmware level. For more information, consult IBM Support.

## 19.5.3 SAS Connectivity Modules

Installation of the IBM BladeCenter SAS Connectivity Modules is supported only in IBM 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 cards that are installed in the blade servers in your BladeCenter unit. Installing two connectivity modules allows you to have two connections to the SAS expansion cards that are installed in the blade servers.

You can find installation and user guide for SAS Connectivity Modules and the SAS Connectivity Modules firmware update procedure at the following website:

<http://www-947.ibm.com/support/entry/portal/docdisplay?brand=5000020&indocid=MIGR-5072374>

## 19.5.4 SAS card installation

In this example, you must install the IBM BladeCenter SAS Expansion Card in the PS701 BladeCenter host before you 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.

After you power on the PS701 BladeCenter host, rescan the hardware by running `cfgmgr`. After you rescan the new hardware, list the hardware by running the command that is shown in Example 19-17 on page 565.

Example 19-17 *lsdev* command - new hardware

---

```
## lsdev
L2cache0 Available          L2 Cache
cd0       Defined    2.1.2.3    USB DVD-COMBO Drive
cluster0 Available          Cluster Node
dlc8023   Available          IEEE Ethernet (802.3) Data Link Control
dlcether  Available          Standard Ethernet Data Link Control
dlcfddi   Available          FDDI Data Link Control
dlcqllc   Available          X.25 QLLC Data Link Control
dlcsdlc   Available          SDLC Data Link Control
dlctoken  Available          Token-Ring Data Link Control
en0       Available          Standard Ethernet Network Interface
en1       Available          Standard Ethernet Network Interface
ent0      Available          Logical Host Ethernet Port (lp-hea)
ent1      Available          Logical Host Ethernet Port (lp-hea)
et0       Defined           IEEE 802.3 Ethernet Network Interface
et1       Defined           IEEE 802.3 Ethernet Network Interface
fwdump    Defined           Logical volume
hd1       Defined           Logical volume
hd2       Defined           Logical volume
hd3       Defined           Logical volume
hd4       Defined           Logical volume
hd5       Defined           Logical volume
hd6       Defined           Logical volume
hd8       Defined           Logical volume
hd10opt   Defined           Logical volume
hd11admin Defined           Logical volume
hd9var    Defined           Logical volume
hdisk0    Available  00-08-00    SAS Disk Drive
inet0     Available          Internet Network Extension
iocp0     Defined           I/O Completion Ports
iscsi0    Available          iSCSI Protocol Device
lg_dumplv Defined           Logical volume
lhea0     Available          Logical Host Ethernet Adapter (l-hea)
livedump  Defined           Logical volume
lo0       Available          Loopback Network Interface
lvdd      Available          LVM Device Driver
mem0      Available          Memory
nsmb0     Available          N/A
pci0      Available          PCI Bus
pci1      Available          PCI Bus
pkcs11    Available          PKCS#11 Device
proc0     Available  00-00       Processor
proc4     Available  00-04       Processor
proc8     Available  00-08       Processor
proc12    Available  00-12       Processor
proc16    Available  00-16       Processor
proc20    Available  00-20       Processor
proc24    Available  00-24       Processor
proc28    Available  00-28       Processor
pty0      Available          Asynchronous Pseudo-Terminal
rcm0      Defined           Rendering Context Manager Subsystem
rootvg    Defined           Volume group
sas0      Available  00-08-00    Controller SAS Protocol
sata0     Available  00-08-00    Controller SATA Protocol
```

ses0	Available	00-08-00	SAS Enclosure Services Device
sfw0	Available		Storage Framework Module
sfwcomm0	Available	00-08-00-FF	SAS Storage Framework Comm
sissas0	Available	00-08	PCI-X266 Planar 3Gb SAS Adapter
sys0	Available		System Object
sysplanar0	Available		System Planar
usb0	Available		USB System Software
usbhc0	Available	01-08	USB Host Controller (33103500)
usbhc1	Available	01-09	USB Host Controller (33103500)
usbhc2	Available	01-0a	USB Enhanced Host Controller (3310e000)
vio0	Available		Virtual I/O Bus
vsa0	Available		LPAR Virtual Serial Adapter
vtty0	Available		Asynchronous Terminal

**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.

**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.

### 19.5.5 SAS Connectivity Module firmware update

Ensure that your SAS Connectivity Module is updated with the latest firmware. For the latest firmware update, go to the following website:

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

To update the SAS Connectivity Module firmware to the latest version, complete the following steps:

1. Log on to the SAS Connectivity Module by using the web interface with the IP address that is defined for the connectivity module in the BladeCenter Advance Management Module (AMM), as shown in Figure 19-4 on page 567.

**Note:** The default Used ID is USERID and the Password is PASSWORD. You can change the password under the Administer Users menu option after you are logged on.

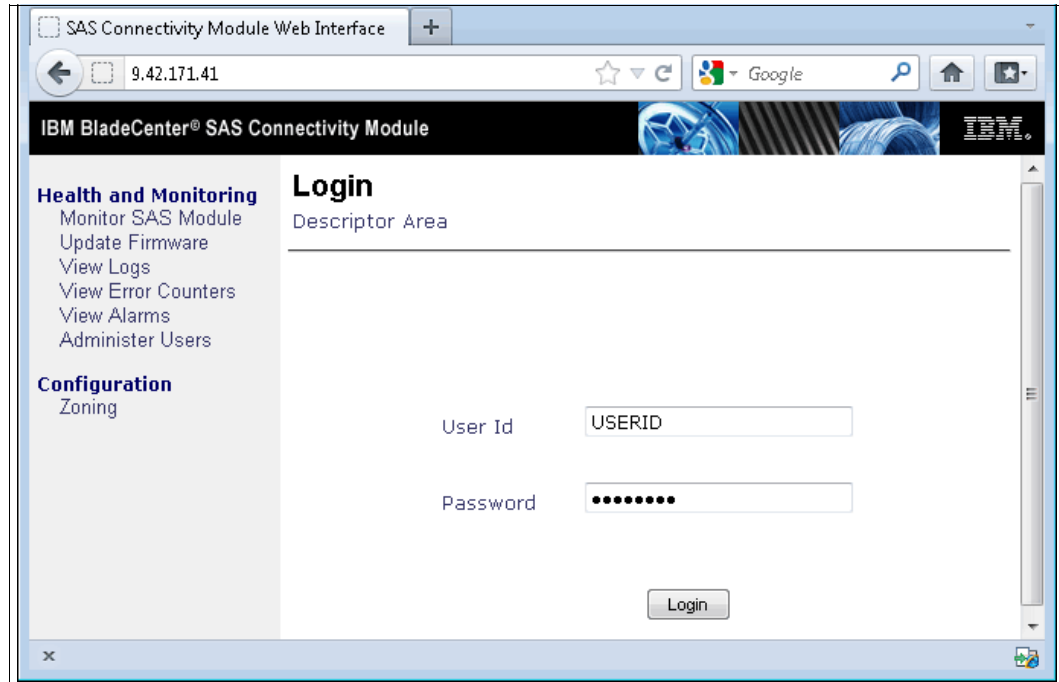


Figure 19-4 SAS Connectivity Module - Login

2. In the Monitor Module window, click **Update Firmware**. The Update Firmware window opens, as shown in Figure 19-5. The current firmware level is also displayed.

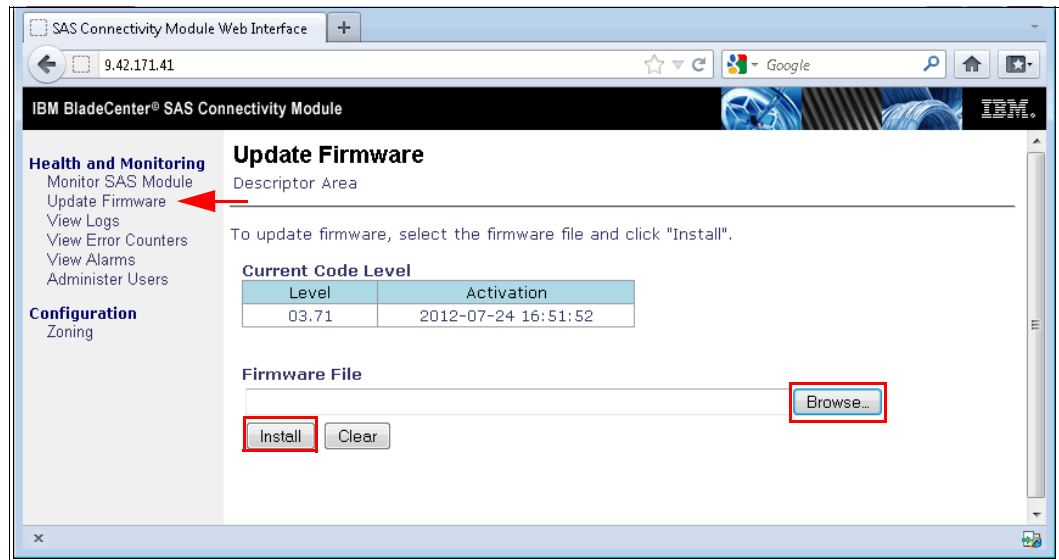


Figure 19-5 SAS Connectivity Module - Update Firmware

In the Firmware File field, enter the new firmware file name, or click **Browse** and locate the firmware file. Click **Install** to install the new file.

3. A firmware update progress window opens, as shown in Figure 19-6.




Figure 19-6 SAS Connectivity Module - installation confirmation

4. Click **OK** or **Cancel**. If the installation of the new firmware file is successful, an installation confirmation window opens and notifies you that SAS module will be rebooted automatically for new firmware activation.

## 19.5.6 SAS Connectivity Module zoning

Serial Attached SCSI (SAS) zoning is a way to limit which blades of an IBM BladeCenter can access which disks and external ports. SAS zoning values are maintained in SAS configuration stores on the SAS Connectivity Modules.

The following website has IBM BladeCenter information pertaining to SAS zoning:

[http://publib.boulder.ibm.com/infocenter/systems/topic/com.ibm.bladecenter.8886.doc/dw1fs\\_r\\_integrated\\_storage\\_planning.html](http://publib.boulder.ibm.com/infocenter/systems/topic/com.ibm.bladecenter.8886.doc/dw1fs_r_integrated_storage_planning.html)

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 by 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 that shows the predefined zone configuration options, as shown in Figure 19-7. You can select from five predefined zone configuration options. In this example, option 5, Predefined Configuration 01, is active (as indicated by the check mark). With this option, each server bay is exclusively zoned with all the external ports, thus allowing access to one or more storage controller ports that are 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"/>	<input checked="" type="checkbox"/>	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 19-7 SAS Connectivity Module - predefined zones

Figure 19-8 was captured from the SAS Connectivity Module web interface. It lists the Basic Zone Permission table for the PS701 blade in slot 8 because the blade in slot 8 is used for this example. The blade in slot 8 is zoned with four external ports. The External port is set to True under the Connected column and Normal under the Status column because the DCS3700 storage subsystem Controller A port is connected to external port #1 for this example.

ID	Address	Enabled	Connected	Fault	Status
SAS Module	500507641008FA80				Normal
Blade Slot Connection 1	0000000000000000	true	false	off	No Cable
Blade Slot Connection 2	0000000000000000	true	false	off	No Cable
Blade Slot Connection 3	0000000000000000	true	false	off	No Cable
Blade Slot Connection 4	0000000000000000	true	false	off	No Cable
Blade Slot Connection 5	0000000000000000	true	false	off	No Cable
Blade Slot Connection 6	500507604D000EDA	true	true	off	Normal
Blade Slot Connection 7	0000000000000000	true	false	off	No Cable
Blade Slot Connection 8	500507602316C401	true	true	off	Normal
Blade Slot Connection 9	0000000000000000	true	false	off	No Cable
Blade Slot Connection 10	0000000000000000	true	false	off	No Cable
Blade Slot Connection 11	0000000000000000	true	false	off	No Cable
Blade Slot Connection 12	0000000000000000	true	false	off	No Cable
Blade Slot Connection 13	0000000000000000	true	false	off	No Cable
Blade Slot Connection 14	0000000000000000	true	false	off	No Cable
External Port 1		true	true	off	Normal
External Port 2		true	false	off	No Cable

Figure 19-8 SAS Connectivity Module - zone configuration

**Note:** There is only a single path that is active from the DCS3700 storage subsystem controller A to the SAS Connectivity Module in BladeCenter I/O slot bay 3. The second path from DCS3700 storage subsystem controller B is to the SAS Connectivity Module in BladeCenter I/O slot bay 4.



## 19.5.7 Mapping a LUN from IBM DS Storage Manager

Now, you add mapping to your LUN by using the Host Mapping tab by completing the following steps:

1. Search for new host ports that are connected to the DCS3700 storage subsystem, right-click **Unassociated Host Port Identifiers**, and select **View**, as shown in Figure 19-9.

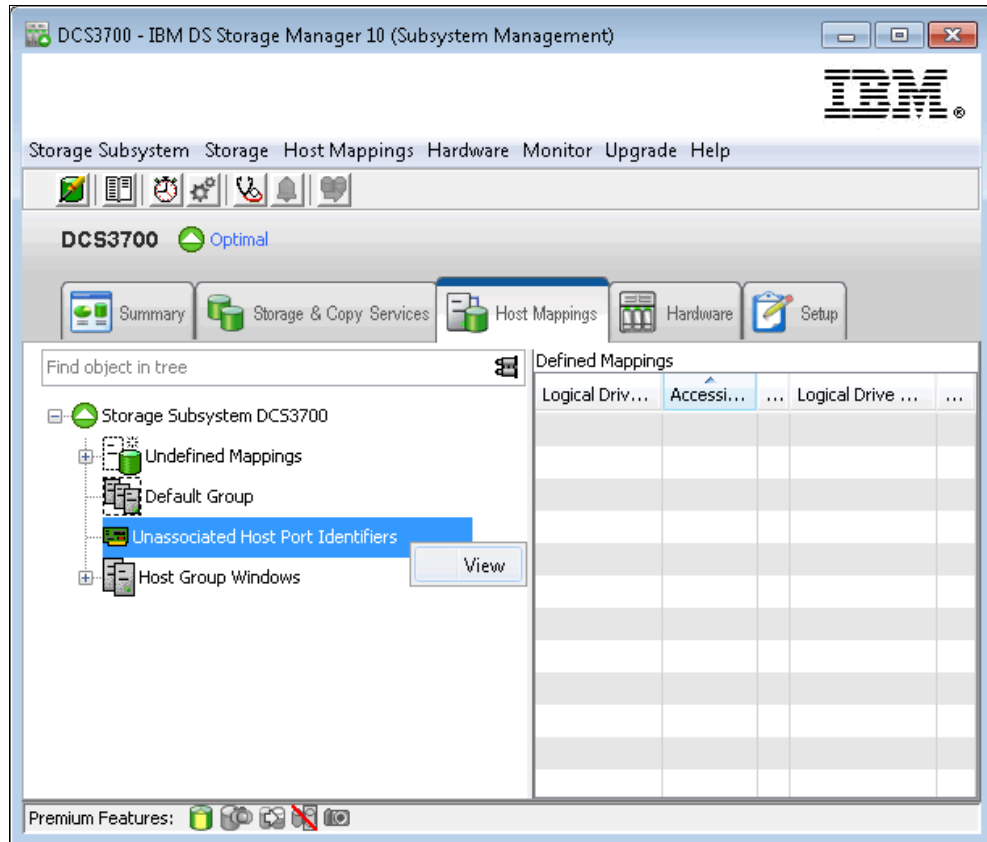


Figure 19-9 Unassociated host port

The windows shows information about the available host port and interface types, as shown in Figure 19-10.

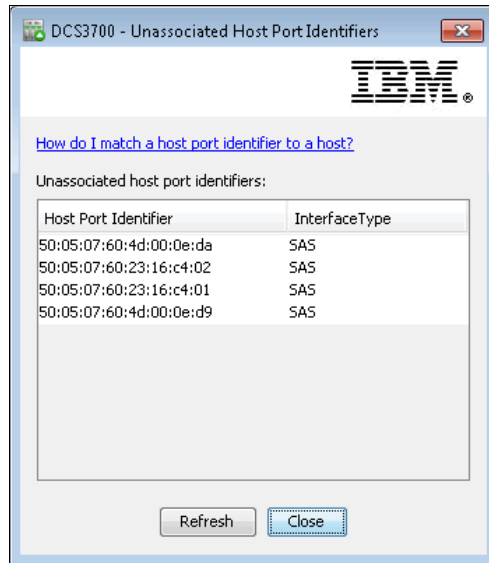


Figure 19-10 Available ports

2. Right-click **Default Group** and select **Define** → **Host Group...**. Enter the name of the host group in to the **Enter new host group name** field and click **OK**.
3. Right-click the AIX host group and select **Define** → **Host...**, as shown in Figure 19-11.

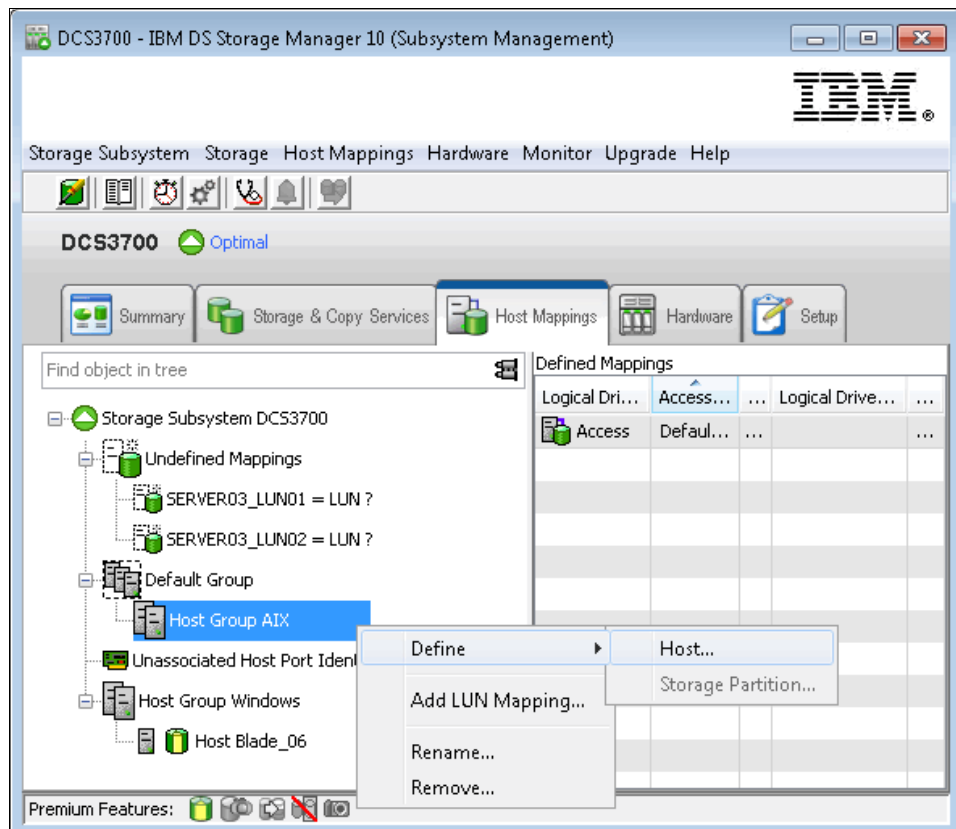


Figure 19-11 Defining a host

4. Enter the name of host, as shown in Figure 19-12.

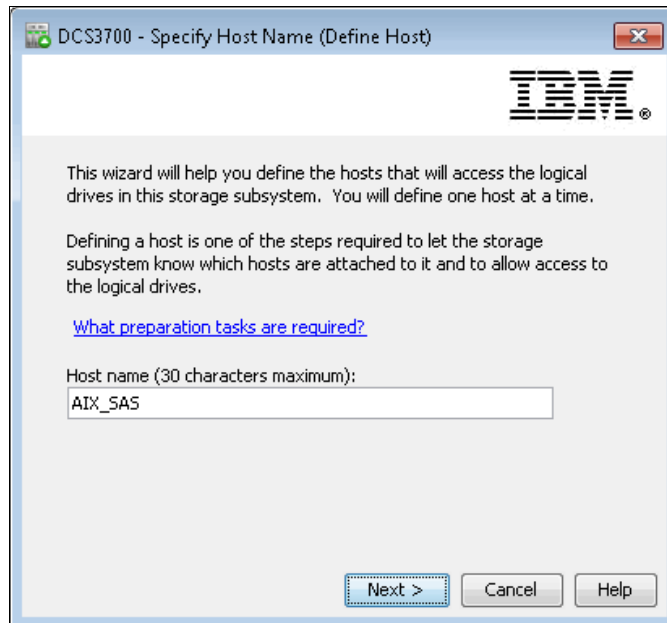


Figure 19-12 Defining a host name

5. In Figure 19-13, define the host interface type, choose a host port identifier, and create an alias for the host and add it.

DCS3700 - 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 AIX\_SAS.

[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:  
50:05:07:60:23:16:c4:02

Add by creating a new host port identifier

New host port identifier (16 characters required):

Alias (30 characters maximum):

Add Remove

Host port identifiers to be associated with the host:

Host Port Identifier	Alias / User Label
50:05:07:60:23:16:c4:01	Port_1
50:05:07:60:23:16:c4:02	Port_2

< Back Next > Cancel Help

Figure 19-13 Specify host port identifiers

6. Select the **AIX** host type from the Host type (operating system): menu, as shown in Figure 19-14.

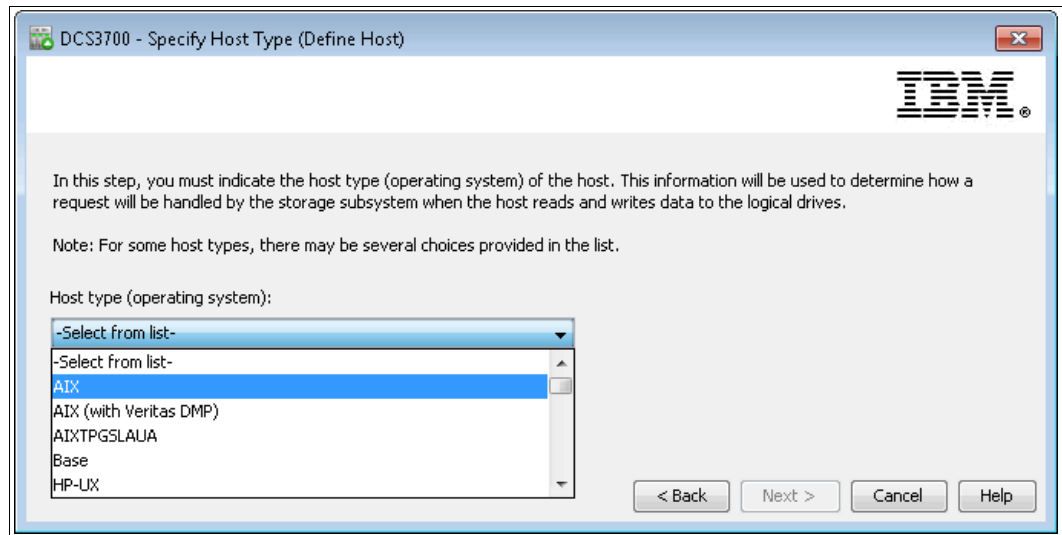


Figure 19-14 Specify host type

7. In Figure 19-15, you see the definition for the created host. Click **Finish** to complete the host mappings.

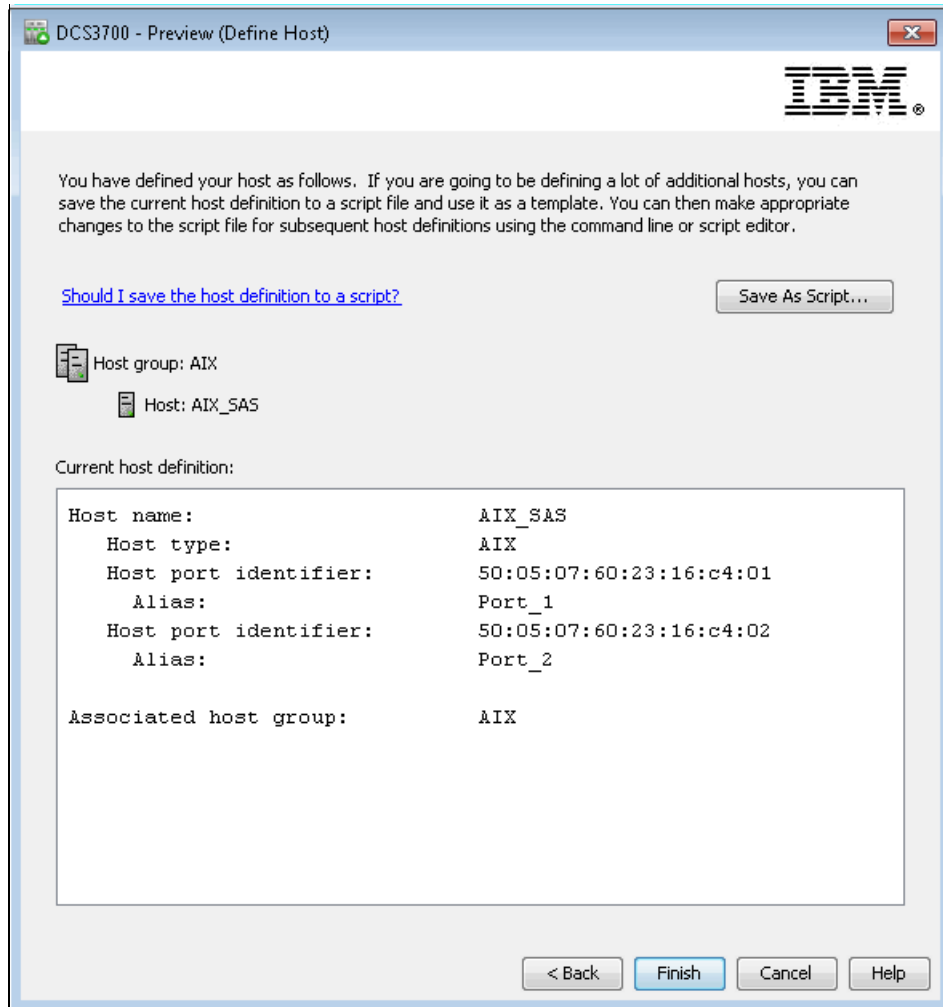


Figure 19-15 Preview (Define Host)

8. Create a logical drive LUN and map it to your AIX host, as shown in Figure 19-16 on page 577. LUN creation is described in Chapter 8, “Administration: Storage & Copy Services tab” on page 225.

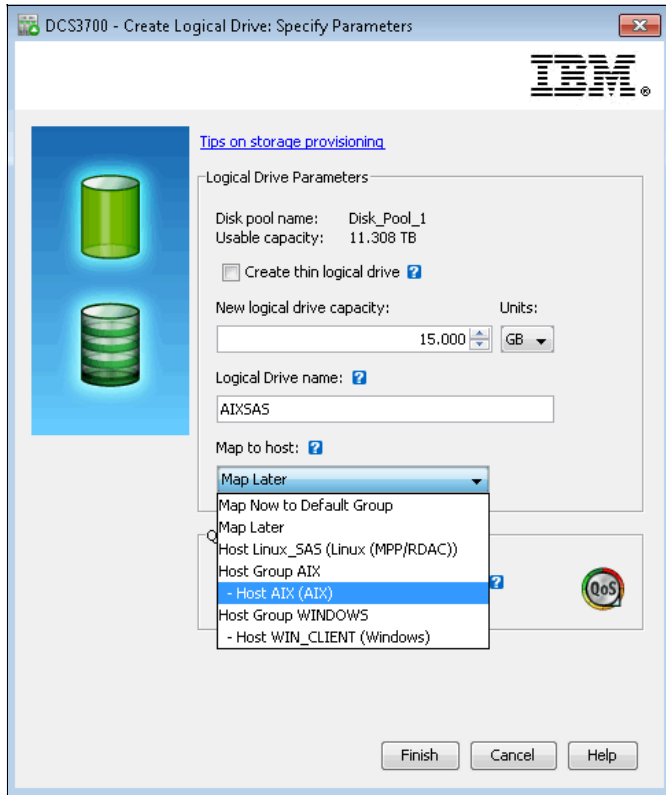


Figure 19-16 Creating and mapping LUN

## 19.5.8 Adding a LUN in AIX

After you finish with LUN creation and mapping by using the IBM DS Storage Manager, you must perform disk rescanning by using AIX. To discover devices on AIX, run `cfgmgr -v`. This command configures devices and (optionally) installs the devices' software in to the system.

Running `lsdev` lists all your resources in AIX, as shown in Example 19-18. You now have the paths to your new disk (`hdisk1`).

*Example 19-18 lsdev*

```
# # lsdev
L2cache0 Available L2 Cache
cd0 Defined 2.1.2.3 USB DVD-COMBO Drive
cluster0 Available Cluster Node
dac0 Available 00-08-00 DS3/4K PCM User Interface
dac1 Available 00-08-00 DS3/4K PCM User Interface
dlc8023 Available IEEE Ethernet (802.3) Data Link Control
dlcether Available Standard Ethernet Data Link Control
dlcfdi Available FDDI Data Link Control
dlcqllc Available X.25 QLLC Data Link Control
dlcsdlc Available SDLC Data Link Control
dlctoken Available Token-Ring Data Link Control
en0 Available Standard Ethernet Network Interface
en1 Available Standard Ethernet Network Interface
ent0 Available Logical Host Ethernet Port (lp-hea)
ent1 Available Logical Host Ethernet Port (lp-hea)
et0 Defined IEEE 802.3 Ethernet Network Interface
```

et1	Defined	IEEE 802.3 Ethernet Network Interface
fwdump	Defined	Logical volume
hd1	Defined	Logical volume
hd2	Defined	Logical volume
hd3	Defined	Logical volume
hd4	Defined	Logical volume
hd5	Defined	Logical volume
hd6	Defined	Logical volume
hd8	Defined	Logical volume
hd10opt	Defined	Logical volume
hd11admin	Defined	Logical volume
hd9var	Defined	Logical volume
hdisk0	Available 00-08-00	SAS Disk Drive
<b>hdisk1</b>	<b>Available 00-08-00</b>	<b>MPIO Other DS3K Array Disk</b>
inet0	Available	Internet Network Extension
iocp0	Defined	I/O Completion Ports
iscsi0	Available	iSCSI Protocol Device
lg_dump1v	Defined	Logical volume
lhea0	Available	Logical Host Ethernet Adapter (l-hea)
livedump	Defined	Logical volume
lo0	Available	Loopback Network Interface
lvdd	Available	LVM Device Driver
mem0	Available	Memory
nsmb0	Available	N/A
pci0	Available	PCI Bus
pci1	Available	PCI Bus
pkcs11	Available	PKCS#11 Device
proc0	Available 00-00	Processor
proc4	Available 00-04	Processor
proc8	Available 00-08	Processor
proc12	Available 00-12	Processor
proc16	Available 00-16	Processor
proc20	Available 00-20	Processor
proc24	Available 00-24	Processor
proc28	Available 00-28	Processor
pty0	Available	Asynchronous Pseudo-Terminal
rcm0	Defined	Rendering Context Manager Subsystem
rootvg	Defined	Volume group
sas0	Available 00-08-00	Controller SAS Protocol
sata0	Available 00-08-00	Controller SATA Protocol
ses0	Available 00-08-00	SAS Enclosure Services Device
sfw0	Available	Storage Framework Module
sfwcomm0	Available 00-08-00-FF	SAS Storage Framework Comm
sissas0	Available 00-08	PCI-X266 Planar 3Gb SAS Adapter
sys0	Available	System Object
sysplanar0	Available	System Planar
usb0	Available	USB System Software
usbhc0	Available 01-08	USB Host Controller (33103500)
usbhc1	Available 01-09	USB Host Controller (33103500)
usbhc2	Available 01-0a	USB Enhanced Host Controller (3310e000)
vio0	Available	Virtual I/O Bus
vsa0	Available	LPAR Virtual Serial Adapter
vtty0	Available	Asynchronous Terminal

---



Run `mpio_get_config -Av` to get information about the storage type that you attached to your PS701 blade server, as shown in Example 19-19.

*Example 19-19 mpio\_get\_config -Av*

---

```
# mpio_get_config -Av
Frame id 0:
  Storage Subsystem worldwide name: 608e50023f8f400004fffe472
  Controller count: 2
  Partition count: 1
  Partition 0:
    Storage Subsystem Name = 'DCS3700'
      hdisk#          LUN #   Ownership           User Label
      hdisk1          0     A (preferred)       AIXSAS
```

---

Run `lspath -l hdisk1` to list all the available path to hdisk1, as shown in Example 19-20.

*Example 19-20 lspath -l hdisk1*

---

```
# lspath -l hdisk1
Enabled hdisk1 sas0
Enabled hdisk1 sas0
```

---

The commands that are shown in Example 19-21 provide more information about hdisk1.

*Example 19-21 lsattr -El*

---

```
# lsattr -El dac0
PCM          PCM/friend/ds4k-ui          Path Control Module      False
unique_id    270ASV126050090FUniveyesl Xport03IBMsas Unique device identifier  False
ww_id        60080e500023ffb6000008b04fffe443    World Wide Identifier    False
# lsattr -El dac1
PCM          PCM/friend/ds4k-ui          Path Control Module      False
unique_id    270ASV127042900FUniversal Xport03IBMsas Unique device identifier  False
ww_id        60080e500023ffb6000008b04fffe443    World Wide Identifier    False
# lsattr -El hdisk1
PCM          PCM/friend/otherapdisk
Path Control Module      False
hdisk0 Available 00-08-00 SAS Disk Drive
Persistant Reserve Key Value  True
hdisk1 Available 00-08-00 MPI0 Other DS3K Array Disk
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_pcnt            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
max_transfer 0x100000
Maximum TRANSFER Size True
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
reserve_policy single_path
Reserve Policy True
rw_timeout 30
READ/WRITE time out value True
size_in_mb 0
Size in Megabytes False
unique_id 3E21360080E500023F8F400000B80500F46F10F1818 FAST03IBMsas
Unique device identifier False
ww_id 60080e500023f8f400000b80500f46f1
World Wide Identifier False

```

---

The commands that are shown in Example 19-22 provide more information about path availability and storage connectivity.

*Example 19-22 sissasraidmgr*

---

```

# sissasraidmgr -L -l sissas0 -j3
-----
Name      Resource  State      Description      Size
-----
sissas0   FFFFFFFF Available  PCI-X266 Planar 3Gb SAS Adapter

hdisk0    00020000 Available  SAS Disk Drive   600.0GB
hdisk1    0000C000 Available  MPIO Other DS3K Array N/A
hdisk1    0001C000 Available  MPIO Other DS3K Array N/A
ses0      00080000 Available  SAS Enclosure Service N/A
dac0      0000C01F Available  DS3/4K PCM User Inter N/A
dac1      0001C01F Available  DS3/4K PCM User Inter N/A
# sissasraidmgr -T -l hdisk1 -o 1

Adapter   Adapter Port  Path Active  Path State  Device
-----
sissas0   0              Yes          Operational  hdisk1

Node  SAS Address      Port Type  Phy  Status      Info
----  -
1     500507602316C401 Adapter    0     Operational 3.0 GBPS
2     500507641008FA80 Expander   14    Operational 3.0 GBPS
3     500507641008FA80 Expander   1C    Operational 3.0 GBPS
4     50080E523F8F4000 Device      3     Operational 3.0 GBPS

```


5 0000000000000000 LUN 1C Operational 00

Adapter	Adapter Port	Path Active	Path State	Device
-----	-----	-----	-----	-----
sissas0	1	Yes	Operational	hdisk1

Node	SAS Address	Port Type	Phy	Status	Info
----	-----	-----	----	-----	-----
1	500507602316C401	Adapter	1	Operational	3.0 GBPS
2	5005076410034C40	Expander	14	Operational	3.0 GBPS
3	5005076410034C40	Expander	1C	Operational	3.0 GBPS
4	50080E523FFB6000	Device	3	Operational	3.0 GBPS
5	0000000000000000	LUN	1C	Operational	00

---





## **IBM Support Portal website**

This appendix explains how to use the IBM Support Portal website to obtain downloads and support information for the IBM System Storage DCS3700 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 DCS3700 storage subsystem:

1. Enter the following URL in to the address field of your web browser:

<https://www.ibm.com/support/entry/myportal/overview>

A window similar to Figure A-1 opens.

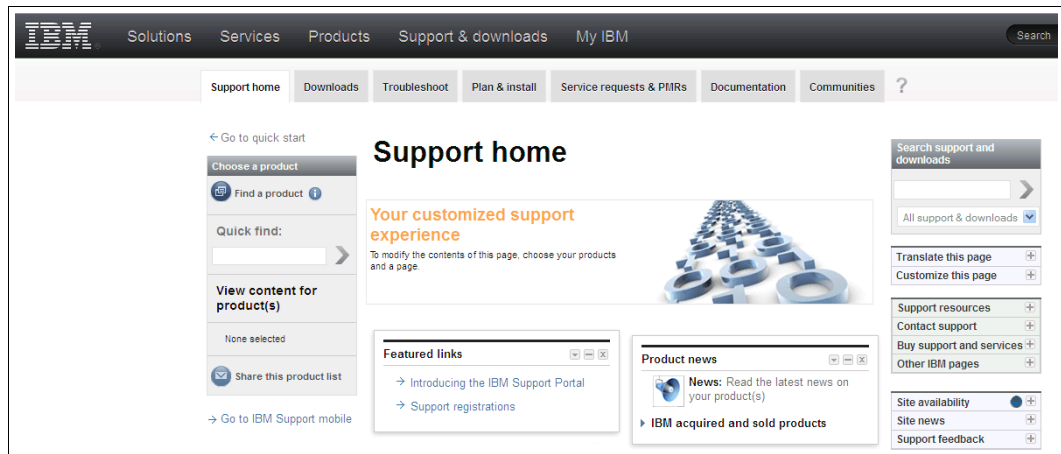


Figure A-1 IBM Support portal

2. Enter DCS3700 in to the Quick Find box and press Enter. A window similar to Figure A-2 opens.

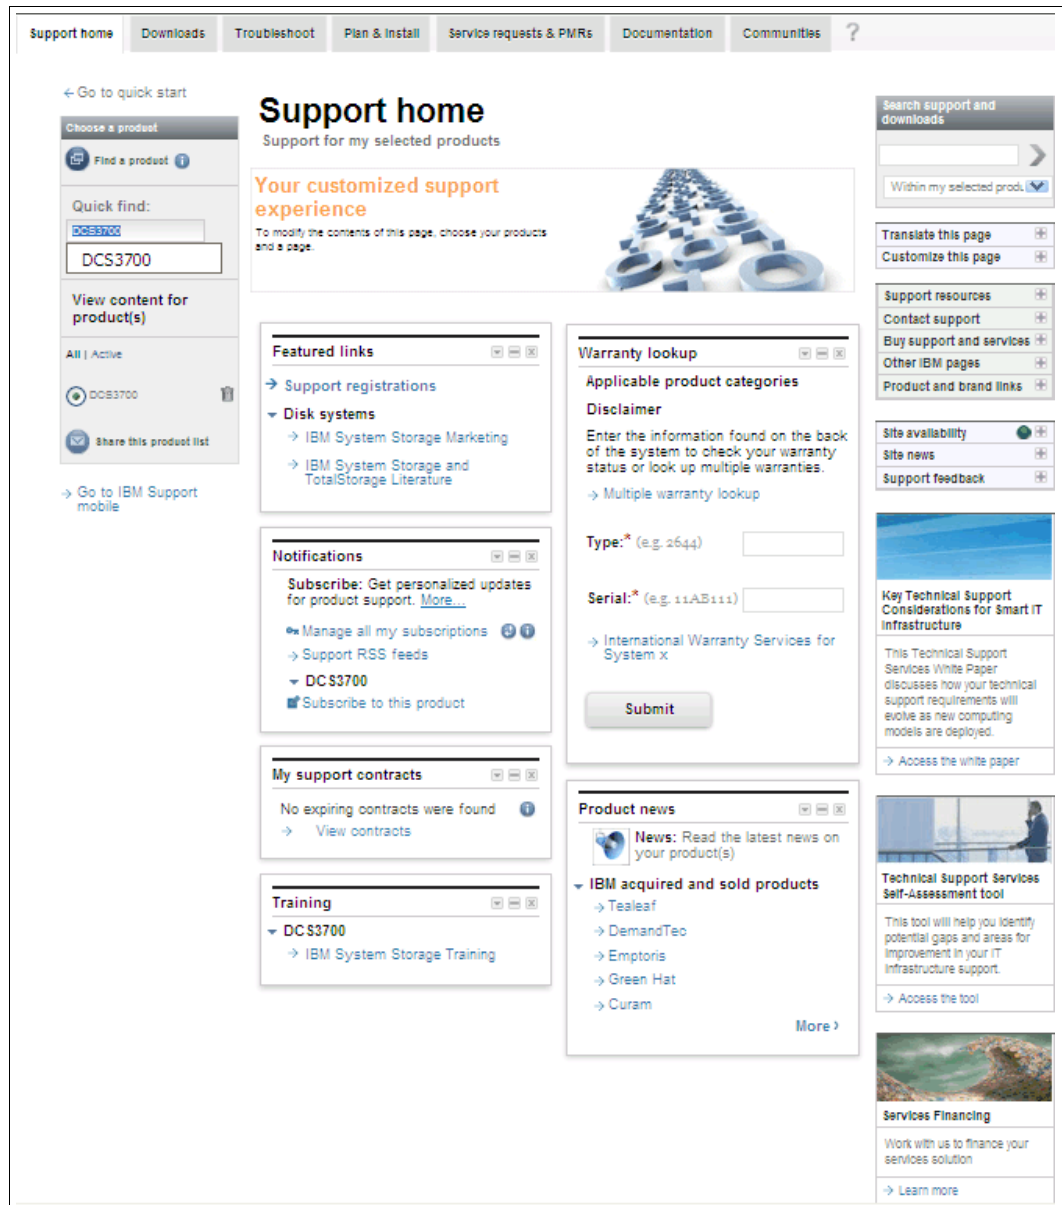


Figure A-2 Search for DCS3700

- You can have more than one product in your selection pane, so you can select multiple products, as shown in Figure A-3.




Figure A-3 Multiple product selection

- Select the **DCS3700** check box. The view refreshes automatically. A support website similar to the one shown in Figure A-4 opens.

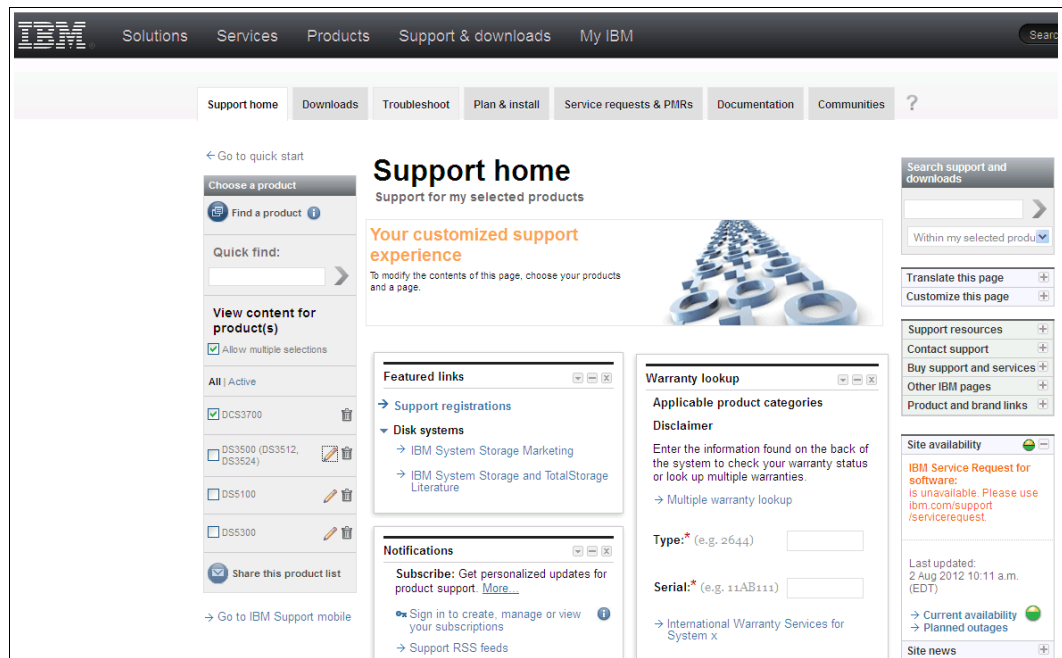


Figure A-4 IBM Support home portal - DCS3700



5. At the top of this page, you can select your function, as shown in Figure A-5. Use these links to select the support option that you want to access.




Figure A-5 Choose your task

## Downloading code updates

To download code updates, complete the following steps:

1. From the Choose your task list (Figure A-5), click **Downloads**. A web page similar to the one shown in Figure A-6 opens.

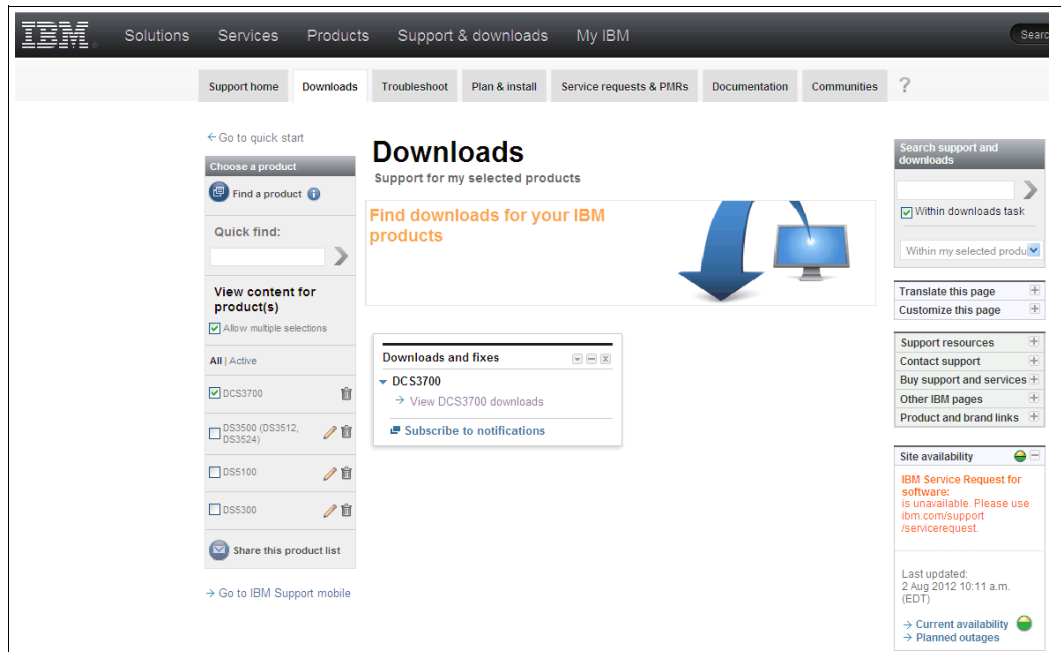


Figure A-6 Support portal - Downloads page

2. Click **View DCS3700 downloads**.

- If you have an IBM ID and are already signed in to the IBM Support Portal, you see the available downloads, as shown in Figure A-7. Here you can select which fixes and updates to download.

## Select fixes

Mid-range disk systems, DCS3700 (All releases, All platforms)

---

**Download options**

- Download method: Download Director [Change download options](#)
- Include requisites: Yes

---

**Select fixes category view**

The following results match your request. Select the fixes you want to download. [Share this download list](#)

- To try a different query, go to the [Identify fixes](#) page.

---

Continue Clear selections Show fix details | Hide fix details

---

Controller firmware   Linux RDAC Driver   Remote Support Manager for Storage   Storage Manager  
ESM & HDD firmware

---

**Controller firmware**

- 1. **fix pack:** [Disk-CntlFW-Jul-2012-version-7.83.22.00](#) ➔ Jul 9, 2012  
IBM DS Controller firmware version 7.83.22.00 code package  
[Change History](#)   [Readme](#)
- 2. **fix pack:** [Disk-CntlFW-Apr-2012-version-7.77.34.00](#) ➔ Apr 20, 2012  
IBM DS Controller firmware version 7.77.34.00 code package  
[Readme](#)   [Change History](#)

[Back to top](#)

---

**ESM & HDD firmware**

Figure A-7 Select fixes

- If you do not have an IBM ID, or you are not signed in to the IBM Support Portal, you see the web page that is shown in Figure A-8.

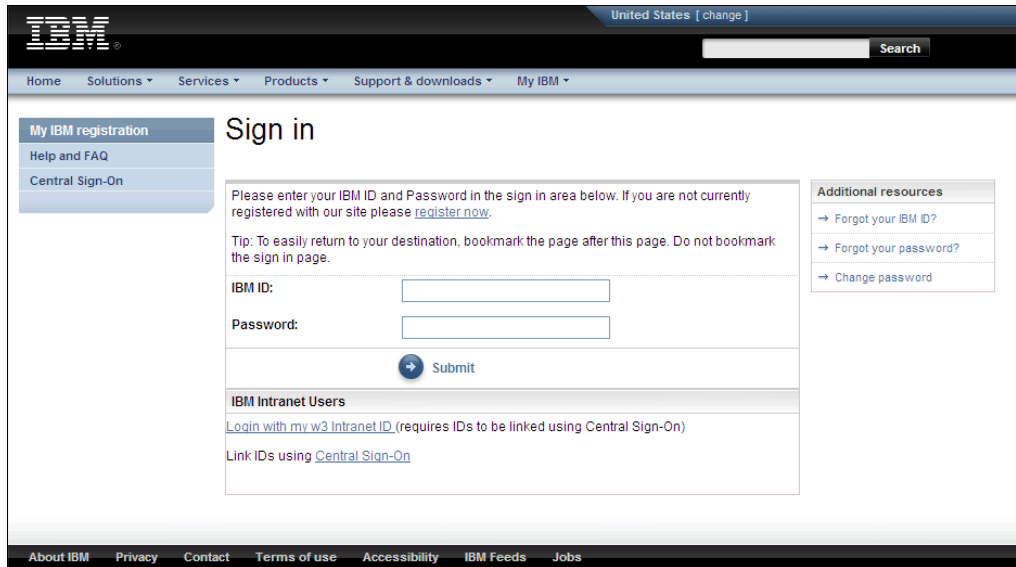


Figure A-8 IBM ID - Sign in

Here you can sign in to the IBM Support Portal by using an existing IBM ID, or create an IBM ID by clicking **Register now**. If you click **Register now**, then the IBM registration window opens, as shown in Figure A-9. Complete the registration form and click **Continue**.

The screenshot shows the IBM registration page. At the top, there is the IBM logo and a navigation bar with links for Home, Solutions, Services, Products, Support & downloads, and My IBM. A search bar is also present. The main heading is "My IBM registration" with a sub-heading "Step 1 of 2". A sidebar on the left contains links for "My IBM profile", "My IBM registration", "Help and FAQ", and "Help desk". The registration form itself includes a note about asterisks indicating required fields, a preferred language selection (English), a disclaimer about IBM's PC business sale to Lenovo, and a reminder that the IBM ID cannot be changed. It then lists the required fields: IBM ID (with a link "Why do I have to provide an email address as my IBM ID?"), Password (minimum 8 characters), Verify password, Security question, Answer to security question, and Email. At the bottom of the form, there is a dropdown menu for "Country/region of residence" and "Continue" and "Cancel" buttons. The footer contains links for About IBM, Privacy, Contact, Terms of use, Accessibility, IBM Feeds, and Jobs.

Figure A-9 IBM registration



5. Figure A-12 shows the different methods that you can use to download the fixes.

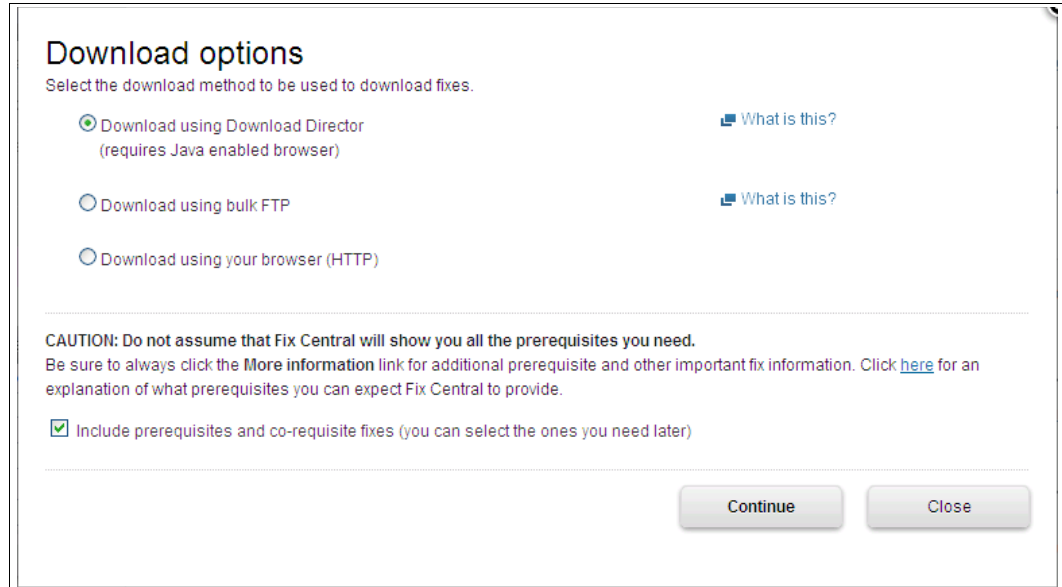


Figure A-12 Download options

You can download by using your web browser (HTTP) or download by 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 capabilities. It starts in a separate browser window and provides a GUI to specify the download location and view the download status.

Click **What is this?** to learn more information about the Download Director. Then, click **Continue**.

6. If you choose to use the Download Director, the Download Director window opens, as shown in Figure A-13.



Figure A-13 Download files by using Download Director

- a. Ensure the check box beside the code package is selected and then click **Download now**.  
The first time that you use the Download Director, a Java applet downloads and installs.
- b. You are prompted to specify a default location for all downloads, as shown in Figure A-14. Click **Browse** to change the folder location. Click **OK** when you choose the download location.




Figure A-14 Download Director - Initial Setup

The Download Director downloads the selected files, as shown in Figure A-15.




Figure A-15 Download Director - download progress

7. If you choose to download by using your browser (HTTP), the web page that is shown in Figure A-16 opens.

**Download files using HTTP**  
Mid-range disk systems, DCS3700 (All releases, All platforms)

---

**Download files using your web browser**  
Click the download link next to each file to download it.

Order number: 57380208  
Total size: 66.51 MB

---

**fix pack: Disk-CntlFW-Jul-2012-version-7.83.22.00** [Change History](#)  
[Readme](#)

IBM DS Controller firmware version 7.83.22.00 code package  
The following files implement this fix.

- [ibm\\_fw\\_ds3k\\_07832200\\_anyos\\_anycpu.txt \(61.1 KB\)](#)
- [ibm\\_fw\\_ds3k\\_07832200\\_anyos\\_anycpu.chg \(71.08 KB\)](#)
- [ibm\\_fw\\_ds3k\\_07832200\\_anyos\\_anycpu.zip \(66.38 MB\)](#)

---

[Back](#)

Figure A-16 Download files by using HTTP

When you use this option, you must click each link to save the file to your local disk.



# My notifications

The IBM My notifications website can be used to keep you informed about 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 DCS3700 storage subsystem:

1. From the main support portal window, which is shown in Figure A-17, click **Notifications**.

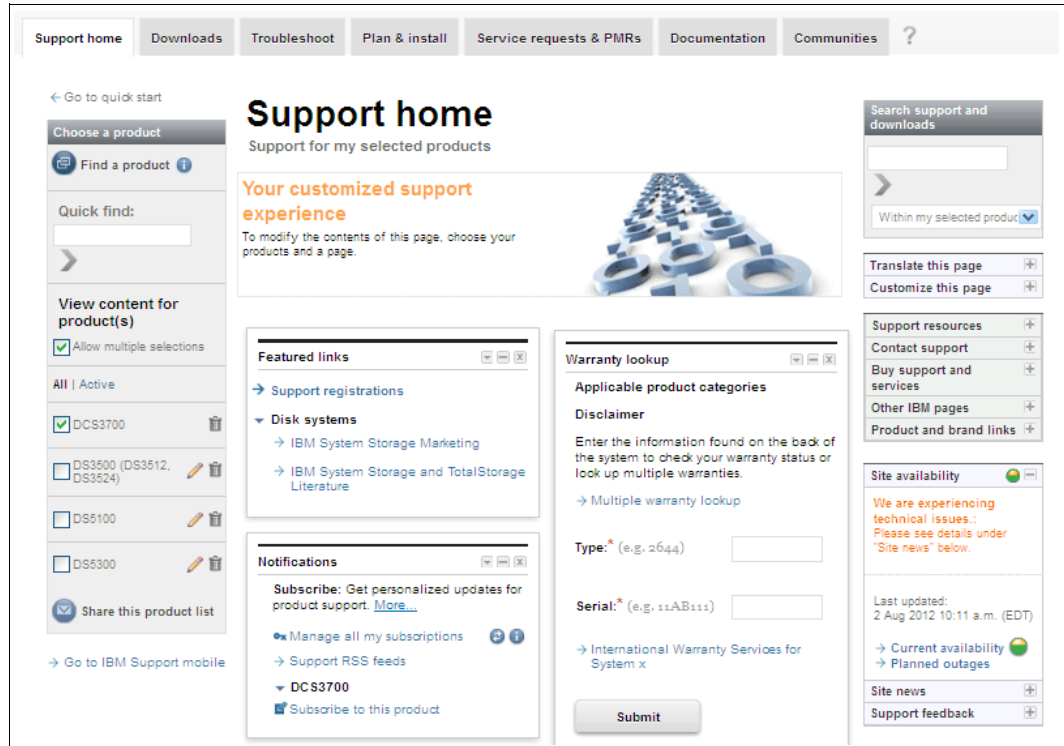


Figure A-17 IBM Support site - overview

2. The Sign in window opens, as shown in Figure A-8 on page 589. Enter your IBM ID and password and proceed. If you are not currently registered with the website, click **Register Now** and register.
3. The My notifications main page opens, as shown in Figure A-18.

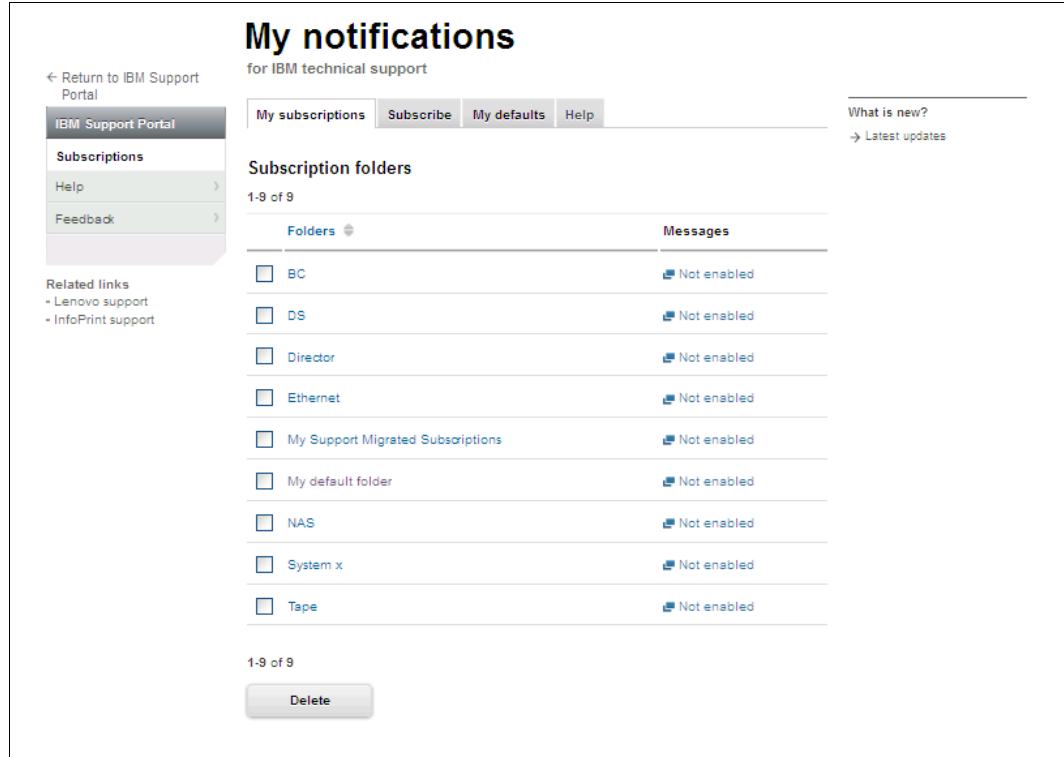


Figure A-18 My notifications

- Select the **Subscribe** tab, and you see the list of available product types, as shown in Figure A-19.

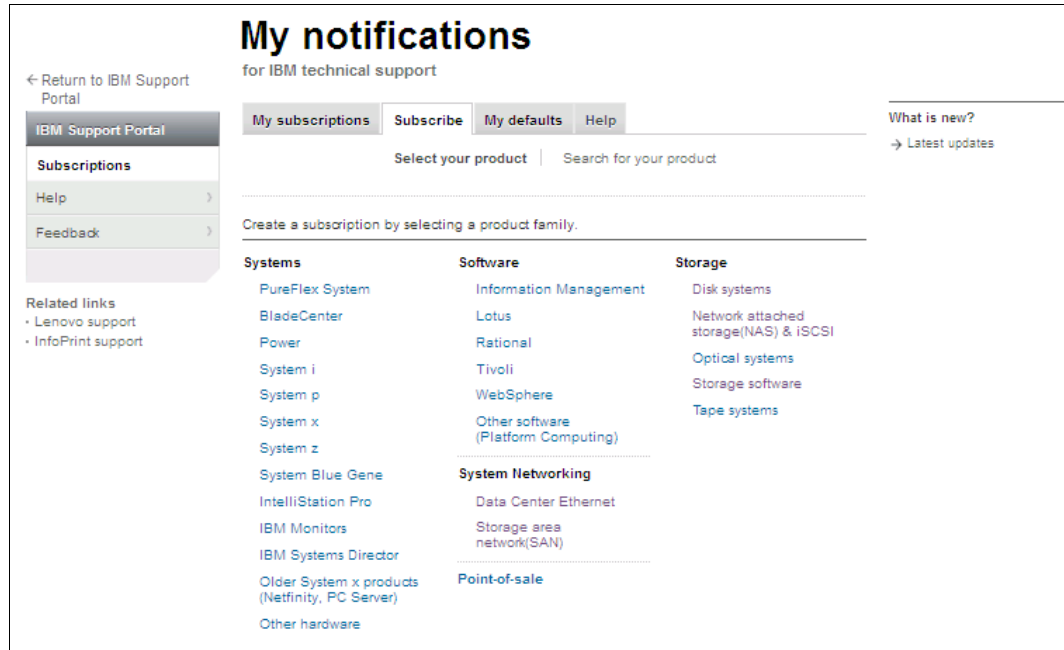


Figure A-19 My notifications - subscribe

- Click **Disk systems** under the Storage column and you see a list of the storage subsystems, as shown in Figure A-20.

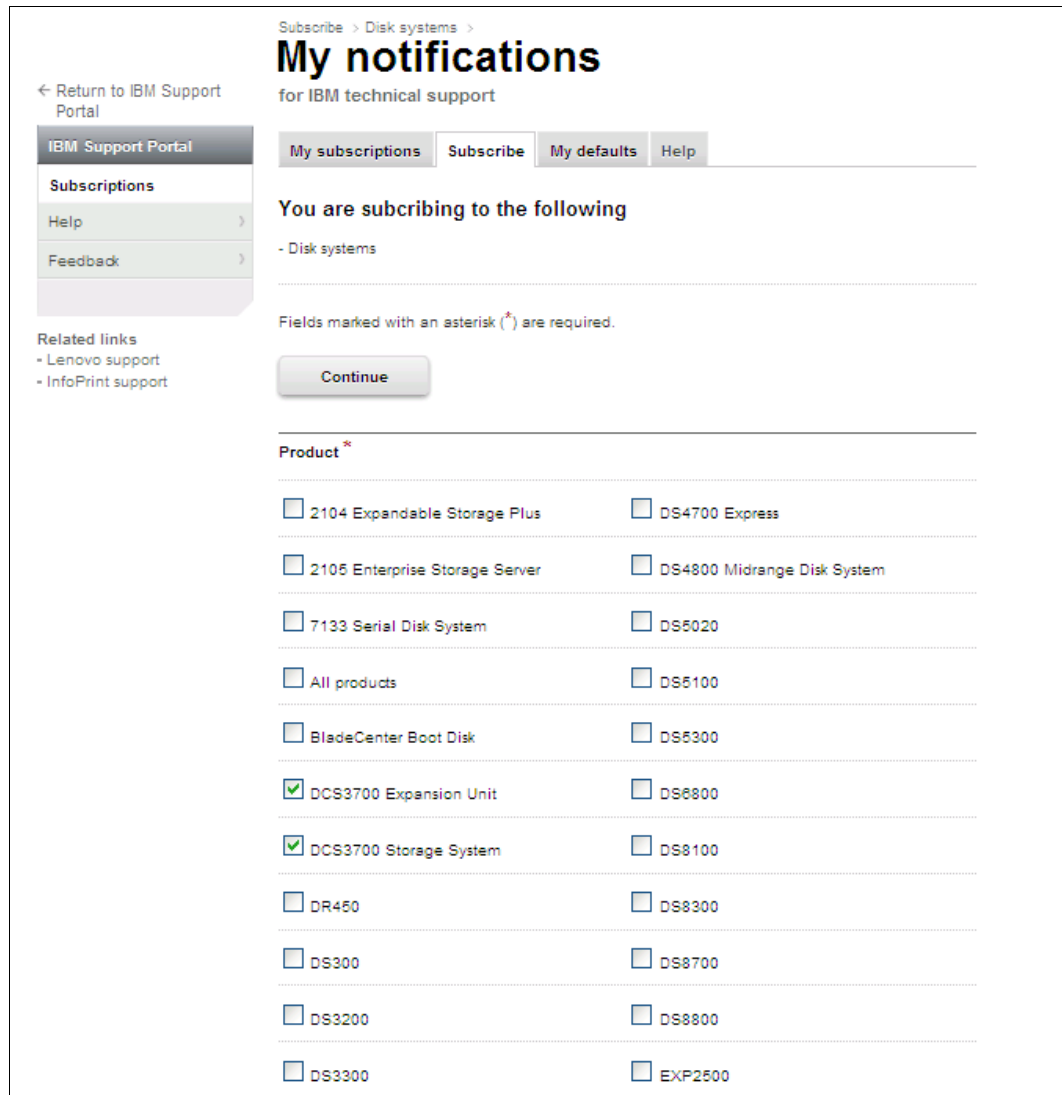


Figure A-20 My notifications - disk systems

- Select the check box beside the disk systems that want to receive notifications for and click **Continue**. You see the website that is shown in Figure A-21 on page 599. Here you can manage when and how you receive your product notifications.

In this example, we chose “DCS3700” as the name of the new folder where the notifications are saved.

After you configure your subscriptions, click **Submit**.

Subscribe > Disk systems >

# My notifications

for IBM technical support

← Return to IBM Support Portal

IBM Support Portal

Subscriptions

Help >

Feedback >

Related links

- Lenovo support
- InfoPrint support

My subscriptions   **Subscribe**   My defaults   Help

**You are subscribing to the following**

- Disk systems
- DCS3700 Storage System, DCS3700 Expansion Unit

---

Fields marked with an asterisk (\*) are required.

**Options**

Subscription name:\*

---

Save in existing or new folder:

Existing folder name:\*

New folder name:\*

---

**Notify me by**

Email

Daily email    Weekly email

Plain text email    Html email

---

Delivery to this folder

---

Delivery via syndication feed (RSS,Atom)

[What is this?](#)

---

Figure A-21 My notifications - setup

7. You see the window that is shown in Figure A-22 when you complete the My notifications setup.

My subscriptions >

# Subscription action was successful

The subscription has been successfully created

Please select an option:

- [Return to IBM Support Portal](#)
- [View all my subscription folders](#)

← Return to IBM Support Portal

IBM Support Portal

Subscriptions

Help >

Feedback >

Related links

- Lenovo support
- InfoPrint support

Figure A-22 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

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- ▶ *IBM Midrange System Storage Hardware Guide*, SG24-7676
- ▶ *IBM Midrange System Storage Implementation and Best Practices Guide*, SG24-6363
- ▶ *IBM System Storage b-type Multiprotocol Routing: An Introduction and Implementation*, SG24-7544
- ▶ *IBM System Storage DS3000: Introduction and Implementation Guide*, SG24-7065
- ▶ *IBM System Storage DS4000 and Storage Manager V10.30*, SG24-7010
- ▶ *IBM System Storage DS5000 Series*, SG24-8024
- ▶ *IBM System Storage DS5000 Series Hardware Guide*, SG24-8023
- ▶ *IBM System Storage DS Storage Manager Copy Services Guide*, SG24-7822
- ▶ *Implementing an IBM b-type SAN with 8 Gbps Directors and Switches*, SG24-6116
- ▶ *Implementing an IBM/Cisco SAN*, SG24-7545
- ▶ *SAN Boot Implementation and Best Practices Guide for IBM System Storage*, SG24-7958
- ▶ *VMware Implementation with IBM System Storage DS5000*, SG24-8052

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

[ibm.com/redbooks](http://ibm.com/redbooks)

## Other publications

These publications are also relevant as further information sources:

- ▶ *Copy Services User's Guide - IBM System Storage DS Storage Manager v10*, GC53-1136
- ▶ *IBM System Storage DS3000, DS4000, and DS5000 Command Line Interface and Script Commands Programming Guide*, GC52-1275
- ▶ *Installation and Host Support Guide - IBM Storage Manager v10*, GA32-0963
- ▶ *Installation, User's, and Maintenance Guide - IBM System Storage DCS3700*, GA32-0959
- ▶ *Quick Start Guide - IBM System Storage DCS3700*, GA32-0960

## Online resources

These websites are also relevant as further information sources:

- ▶ IBM System Storage 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>

- ▶ Support for IBM disk systems:

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

- ▶ System Storage Interoperation Center (SSIC):

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

## Help from IBM

IBM Support and downloads

[ibm.com/support](http://ibm.com/support)

IBM Global Services

[ibm.com/services](http://ibm.com/services)





**Redbooks**

# **IBM System Storage DCS3700 Introduction and Implementation Guide**

(1.0" spine)

0.875" x 1.498"

460 <-> 788 pages







# IBM System Storage DCS3700 Introduction and Implementation Guide



**Redbooks®**

**IBM System Storage  
DS Storage Manager  
V10.8x**

**Hardware planning,  
installation, and  
maintenance**

**Host configuration  
and troubleshooting**

The IBM System Storage DCS3700 consists of IBM System Storage DCS3700 storage subsystem and IBM System Storage DCS3700 expansion unit. The DCS3700 features the latest technologies, including 6 Gbps SAS and 8 Gbps Fibre Channel host interfaces, along with 6 Gbps SAS drives. The DCS3700 also features a 10 Gbps iSCSI host interface with an optional Performance Modules system.

The DCS3700 provides a simple, efficient, and flexible approach to storage that is based on seven generations of design knowledge and firmware development. The DCS3700 can act as a cost-effective and fully integrated complement to IBM System x servers, IBM BladeCenter, and IBM Power Systems for various intensive computing environments.

This IBM Redbooks publication specifically addresses the hardware features, configuration, and implementation of the DCS3700. It presents detailed descriptions of the hardware configurations and options that are offered with the DCS3700. It then presents the concepts and functions that are used in planning and managing the storage subsystems, such as multipathing and path failover. This book offers a step-by-step guide to using the IBM Storage Manager to create arrays, logical drives, and other basic (and advanced) management tasks.

This publication also contains practical information about diagnostic tests and troubleshooting, and includes practical examples about how to use scripts and the command-line interface (CLI).

**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)**

SG24-8037-00

ISBN 0738438243